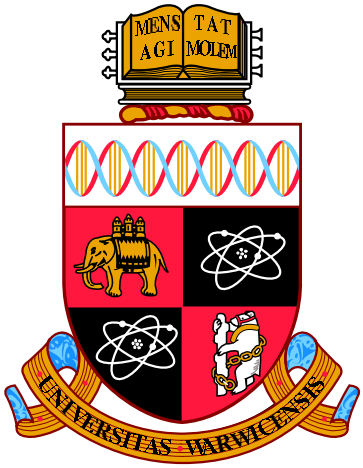

Digital capture of the histological microarchitecture in the myometrium and its implications for the propagation of electrophysiological excitation

by

E. Josiah Lutton

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

in

Systems Biology

University of Warwick, Department of Systems Biology
October 2016

THE UNIVERSITY OF
WARWICK

Contents

List of Figures

List of Tables

Acknowledgements

First and foremost I would show my gratitude to my supervisors Dr Andrew Blanks and Dr Hugo van den Berg for their support, guidance and patience. I would like extend my appreciation and thanks to Dr Wim Lammers for his collaboration and for providing electrode array data for my research. I would like to thank Sean James and the staff at Warwick Clinical Sciences Research Laboratories who enabled me to obtain data for my research. My thanks go to the Reproductive Health group for providing a fun and enjoyable environment. I am grateful to the Medical Research Council for funding my research.

I would like to thank my family for supporting me during my postgraduate studies. I would especially like to thank Jane Hollingsworth for her love and support throughout my PhD.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree. The work presented (including data generated and data analysis) was carried out by myself except where otherwise stated.

Abbreviations

$[Ca^{2+}]_i$ Intracellular calcium concentration

$[Ca^{2+}]_{SR}$ Calcium concentration in the sarcoplasmic reticulum

ATP Adenosine triphosphate

bgCl Background chloride current

bgK Background potassium current

BK Big unitary current potassium channel

Ca^{2+} Calcium ions

CaCC Calcium activated chloride channel

CaL L-type calcium channel

CaT T-type calcium channel

DAG Diacyl-glycerol

GPGPU General-purpose computing on graphics processing units

hERG Human *ether-a-go-go*-related gene potassium channel

H & E Haemotoxylin and eosin

IP_3 Inositol 1,4,5-trisphosphate

K^+ Potassium ions

Kir Inwardly rectifying potassium channel

Kv Potassium ion channel

MSMC Myometrial smooth muscle cell

Na^+ Sodium ions

NBF Neutral buffered formalin

NCX Sodium-calcium exchanger

PIP_2 Phosphatidylinositol 4,5-bisphosphate

PLC Phospholipase C

PMCA Plasma-membrane calcium exchanger

pNa/K Sodium-potassium exchanger

SK Small unitary current potassium channel

SR Sarcoplasmic reticulum

Abstract

Coordination of uterine contractions during labour is critical for successful delivery. The mechanisms underlying this coordination are not fully understood. Propagation of contraction signals has previously been observed to occur through transmission of electrical excitation waves. This thesis aims to examine the histological microarchitecture of the muscular layer of the uterus (myometrium) and determine how this structure affects the propagation of excitation by means of *in silico* three-dimensional reconstruction of the myometrium and numerical simulations of a spatially structured excitation-relaxation model.

A key aim of the *in silico* reconstruction of the smooth muscle architecture of the myometrium is to identify structural features that correspond to the control of excitation behaviour in the myometrium. This examination is aided by analysis of excitation patterns observed in multi-electrode array recordings. The reconstruction is subsequently used as a basis for simulating electrical activity in the myometrium.

Novel structural features are identified here that are located at the initiation points of electrical activity and are proposed to be the pacemaker sites in rat myometrium. Furthermore, boundary of low connectivity across the mesometrial border was observed in the rat, which corresponds to the termination of excitation waves observed in multi-electrode array recordings. In addition, bridges of smooth muscle cells connecting the inner and outer layers of the myometrium were observed in both rat and human myometrium. Taken together these three features suggest a novel mechanism for control of contraction in the rat myometrium; an analogous mechanism is proposed for the human myometrium.

The results presented in this thesis could provide an explanation for the patterns of excitation propagation observed in human and rat uteri. Further refinements of the methods used here are outlined and expected to generate a more detailed visualisation of the structures underpinning these mechanisms.

Chapter 1

Introduction

This thesis is concerned with the role the architecture of the uterine muscle (myometrium) plays in regulating excitation propagation. While the general structure of the myometrium has been well-characterised [111, 116], the connectivity of the fibrous structures is not fully understood. The spatial organisation of connectivity within the myometrial tissue determines the patterns of activation wave propagation through the tissue as well as synchronicity of excitation at the whole-organ level, which is a key determinant of successful labour [9]. This thesis presents a method for determining this micro-architecture, that is, resolving the patterns of connectivity at the cellular-level. Moreover, the impact of tissue micro-architecture on excitation propagation are investigated by means of numerical simulations implementing a spatially structured mathematical model for tissue excitation.

The main aim of this thesis is to understand the relationship between the structure and function of the myometrium. To achieve this goal the following three points shall be addressed. The first and main objective is the generation of a three-dimensional *in silico* reconstruction of the microarchitecture of the myometrium. The second objective is to use this reconstruction as a platform for detailed simulation of electrical activity in the myometrium. The third objective is the analysis of multi-electrode array recordings of electrical activity in the myometrium prior to reconstruction. The multi-electrode array analysis has two main goals: first, to obtain simulation parameter values to ensure that the simulated excitation reflects the real-world excitation, and second to provide insight into the locations of initiation and termination points, which can be examined in the reconstructions to determine if these points correspond to specific anatomical structures.

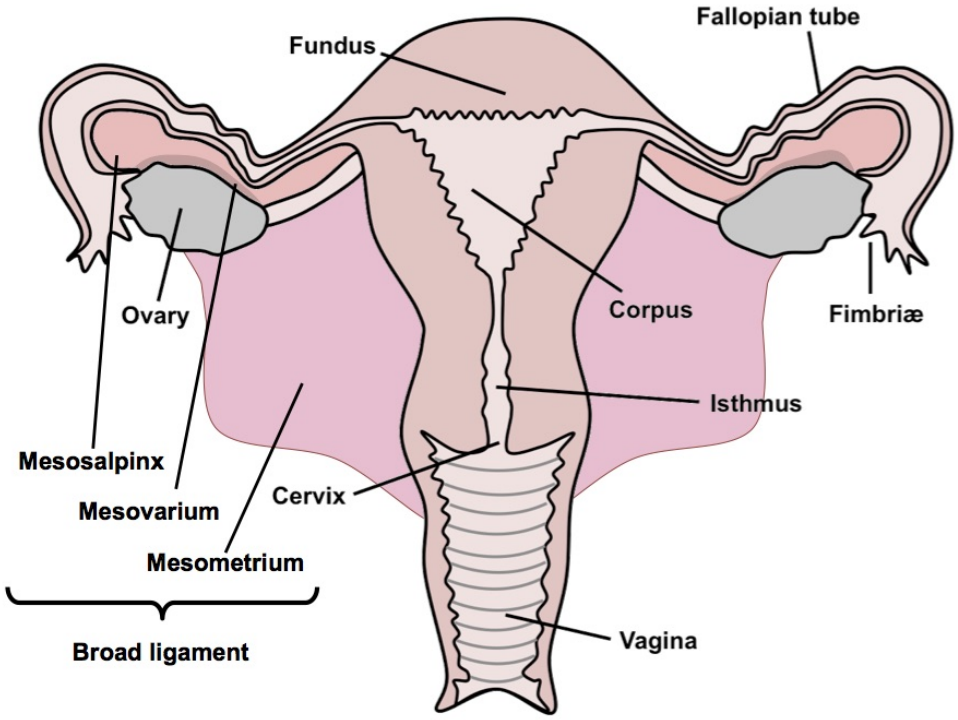

Figure 1.1: The female human reproductive system. During ovulation, the fimbriae draw out an ovum to pass into the Fallopian tube. For pregnancy to occur, sperm cells are passed through the uterus to the Fallopian tube, where they fuse with the ovum (fertilisation). The fertilised egg cell passes from the Fallopian tube into the cavity of the uterus. Adapted from Sheldon *et al.* [101].

1.1 Overview of the female reproductive system

The mammalian female reproductive system consists of the ovaries, Fallopian tubes, uterus, cervix, and vagina, as illustrated in Figure 1.1 [47]. Egg cells, or ova, mature in the ovaries in so-called ovarian follicles. During ovulation, a ripened follicle expels its single ovum, which subsequently passes into the Fallopian tubes, where it may be fertilised by a sperm cell. Following copulation, the male ejaculate enters the uterine cavity. Sperm cells in this ejaculate are transferred into the Fallopian tubes by the uterus, thus increasing the chances of successful fertilisation, which is the fusion of ovum and sperm cell. The product of this fusion is a zygote, which travels from the Fallopian tubes into the interior cavity of the uterus, and quickly develops into a blastocyst. The latter subsequently im-

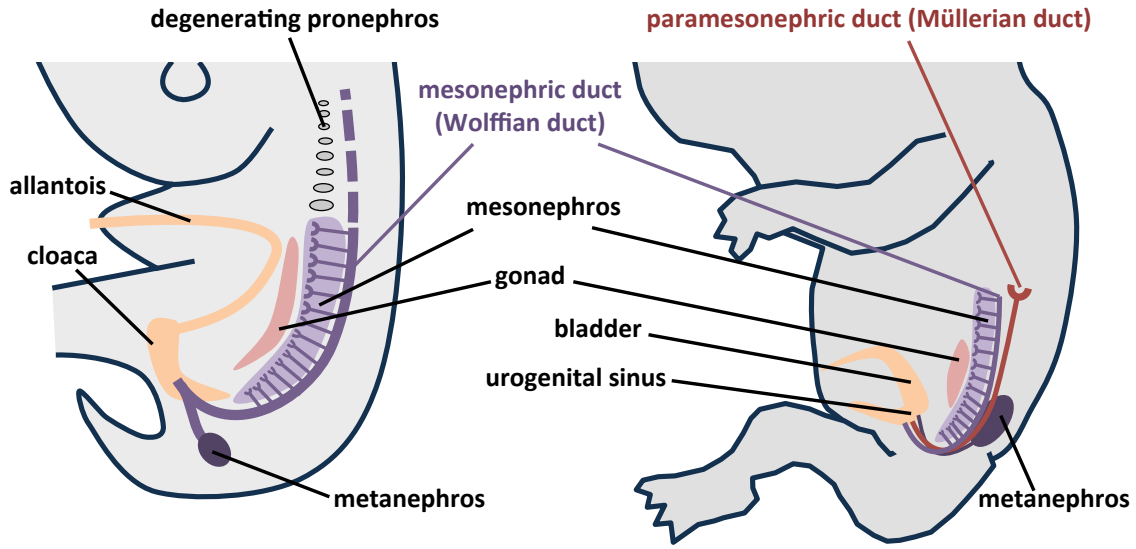

Figure 1.2: Longitudinal sections through a 5-week (left) and 7-week (right) human embryo. During these stages, the mesonephros is the functional excretory organ, while the metanephros is undergoing development into the permanent kidney. During week 5, primordial germ cells migrate from the yolk sac, via the allantois to the gonadal ridges, where they establish themselves in the gonadal primordium that runs alongside the mesonephros. The effluent conduit of the mesonephros is the mesonephric duct, also known as the Wolffian duct; it will ultimately become the vas deferens in the male and disappear in the female. An additional structure, called the paramesonephric duct or Müllerian duct, arises in the vicinity of the mesonephric duct; this structure will not persist in the male, but becomes the internal reproductive tract in the female. Adapted from Carson and Wells [27].

plants in the inner mucous lining of the uterus, an event known as implantation, which marks the start of pregnancy.

1.1.1 Gross anatomy of the human uterus

The human uterus is comprised of three layers called, from the inside moving out, endometrium, myometrium, and perimetrium [27, 95]. The endometrium is the inner mucous lining in which the blastocyst embeds itself. The perimetrium is a thin outer layer encasing the body of the uterus. The myometrium constitutes the bulk; it is the muscular portion of the uterus responsible for contraction [27].

The human uterus connects to the birth canal at its caudal end, where its cavity is continuous with the cervix, which leads on to the vagina [27]. The rostral ("top") end of

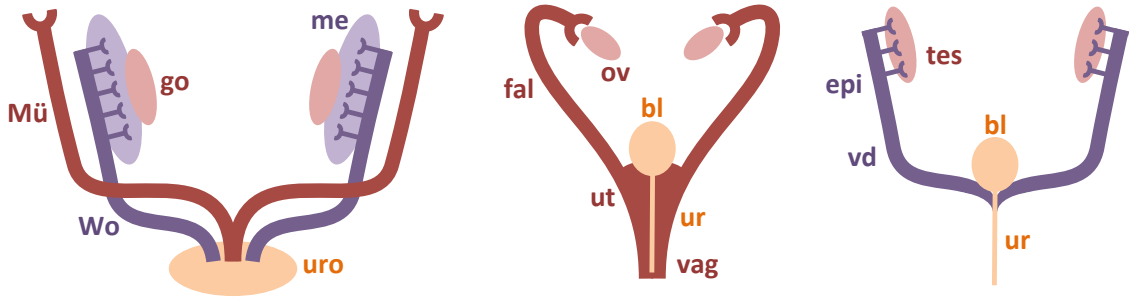

Figure 1.3: Left: indifferent gonadal and urogenital system, in which both mesonephric or Wolffian duct (**Wo**) and paramesonephric or Müllerian duct (**Mü**) are present; indicated are mesonephros (**me**), gonad (**go**), and urogenital sinus (**uro**), Middle: the default development into the female situation, where the Müllerian ducts give rise to vagina (**vag**), uterus (**ut**), and Fallopian tubes (**fal**); the gonad is now an ovary (**ov**). Right: the testosterone-dependent development into the male situation, where the Wolffian ducts give rise to epididymis (**epi**), and vas deferens (**vd**); the gonad is now a testis (**tes**). Also indicated are the bladder (**bl**) and the urethra (**ur**), which both arise out of the urogenital sinus. Adapted from Carson and Wells [27].

the uterus is known as the fundus; here the uterine cavity is continuous with the luminal spaces of the Fallopian tubes.

1.1.2 Development of the reproductive tract

The reproductive and excretory systems develop early in pregnancy from the intermediate mesoderm [27, 95]. In humans, three pairs of excretory organs develop: the pronephroi, the mesonephroi and the metanephroi. The pronephroi develop in week 4 and have degenerated by week 5, but the pronephric ducts only partially degenerate, with the remainder forming the mesonephric ducts. The mesonephroi begin development in week 4, and by week 5 function as the main excretory organs while the metanephroi develop. The metanephroi begin to develop in week 5 and eventually form the kidneys. The state of development of these organs at weeks 5 and 7 is shown in Figure 1.2.

The gonads are initially sexually indifferent, meaning that development proceeds identically in the female and in the male [97]. The reproductive cells arise as primordial germ cells in the yolk sac and travel to the site of the primitive gonad, where they are then called gonocytes, which become organised in string-like structures called sex cords. The default

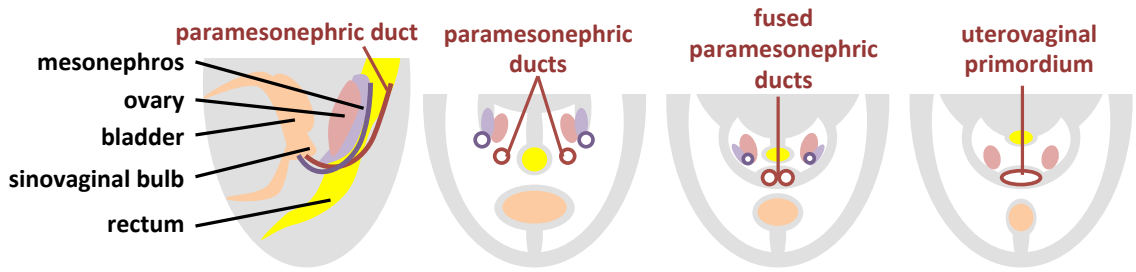

Figure 1.4: Far left: longitudinal section through an 8-week human female embryo. Right: three transverse sections, going from rostral to caudal, showing the fusion of the two paramesonephric ducts. In the female, the mesonephros and its associated duct are degenerating at this stage. The two caudal-most transverse sections also show that the urogenital ridges, which housed the mesonephros and the gonad, have now fused to form a wall separating the body cavity from a new cavity, the rectouterine pouch enveloping the ovaries. Moreover, at the level of the uterovaginal primordium, these walls form the beginning of the broad ligaments. Adapted from Carson and Wells [27].

development results in the female arrangement, in which the gonocytes become oocytes embedded in primordial follicles, and the gonad becomes an ovary [28]. The alternative development is dependent on the additional influence of testosterone, resulting in the male configuration of the reproductive system: the gonocytes become spermatogonia and the gonad develops as a testis.

During the indifferent stage, the gonad is closely associated with the mesonephros, which serves at this point as the functional kidney of the embryo [27]. Ultimately the excretory function will be taken over by the metanephros, which is developing into the later kidney, and then the mesonephros and its associated conduit (known as the mesonephric or Wolffian duct) will degenerate if the default female pattern prevails, while the paramesonephric or Müllerian duct develops into Fallopian tubes, uterus, and vagina [28]. The alternate male development sees the mesonephric duct develop into epididymis and vas deferens while the paramesonephric duct degenerates. These two alternate processes are illustrated in Figure 1.3.

In the human female embryo, the caudal sections of the paramesonephric duct fuse and become the uterovaginal primordium (Figure 1.4). Several congenital malformations can be understood as failures of this Müllerian fusion [27]. For instance, complete failure

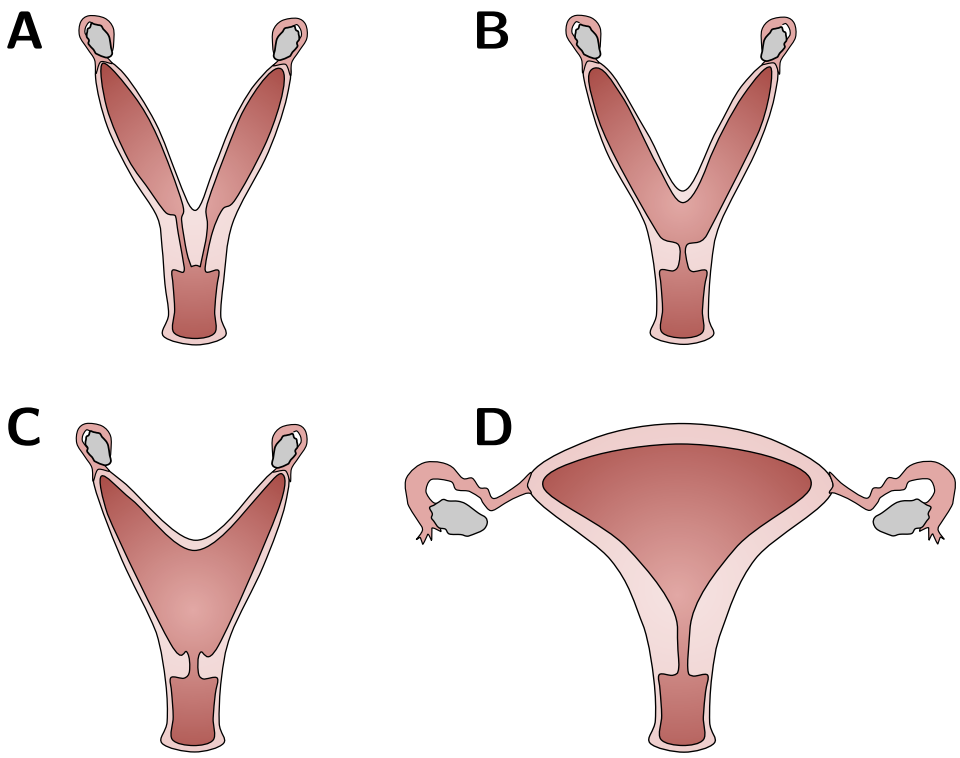

Figure 1.5: Various forms of uterus present in mammals. **A**: The duplex uterus, which is common to rodents. This form is characterised by complete separation of the uterus into two halves, each with its own cervix. **B**: The bicornuate uterus, characterised by two separate uteri joined above the cervix. **C**: The bipartite uterus, similar to the bicornuate uterus, but the uteri are partially fused. **D**: The simplex uterus, found in many primates, including humans. This form is characterised by complete fusion to form a single uterus. Image source: Sheldon [98].

results in *uterus didelphys* in which each duct gives rise to its own vagina, cervix, uterus, and Fallopian tube (only the vulva itself is not reduplicated since it arises from other primordia). If fusion only happens at the caudal-most level corresponding to the vagina, the result is a single vagina but two separate cervices leading to separate uteri; this condition is called a duplex uterus, which is the normal situation in the rat (Figure 1.5**A**) [47]. This form supports multiple pups in each horn of the uterus, as shown in Figure 1.6.

If there is a single cervix, but two uteri are present the situation is called bicornuate uterus (Figure 1.5**B**), which is the normal arrangement in many mammalian genera including farm animals such as *Sus*, *Bos*, *Equus*, and *Ovis* [39]. To distinguish the normal

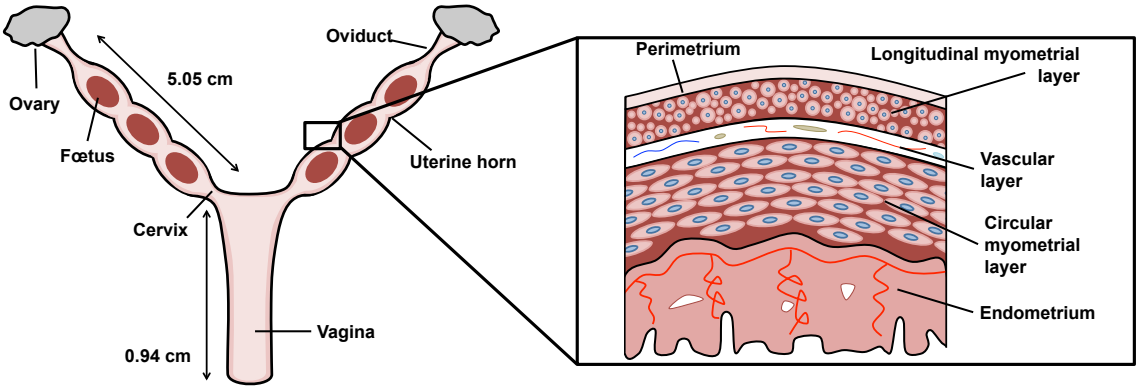

Figure 1.6: Female reproductive system in the rat. The separate uterine horns are each able to maintain multiple pups during pregnancy. The horns each have their own cervices, as is characteristic of the duplex uterus. The structure of the tissue of each horn is formed into layers. In particular, the myometrium, situated between the perimetrium and endometrium, is itself made up of orderly layers. The arrangement of these layers differs from that of the human, which is more complex. Image source: Sheldon [98].

human morphology in the context of comparative anatomy, it is referred to as a simplex uterus (Figure 1.5**D**).

Besides these malformations which correspond closely to the nominal situation in other species, there are several other conditions, such as a bicornuate uterus in which one of the horns remains rudimentary or fails to develop, or a *septate* uterus which is a simplex uterus in which merger of the uterus horns happens belatedly, resulting in a tissue wall down the midline of the uterus [27].

1.2 Excitable properties of myometrial smooth muscle cells

In order to understand the relationship between structure and function of the myometrium, it is important to gain an insight into the function and physiology of its constituent cells. Contraction in the myometrium occurs on a single-cell level, with activating signals being passed between cells via intercellular connections (gap junctions); this communication of activity between the myometrial cells is essential to contraction in the tissue as a whole [44].

Myometrial smooth muscle cells (MSMCs) are spindle-shaped, generally being ~ 5 μm

wide and ~ 50 μm long in the rodent [15], and of similar width in the human but with length ~ 500 μm [23]. Each MSMC contains a meshwork of myosin and actin filaments which slide alongside each other in an ATP-dependent manner [4], thus enabling the cell to contract. The contraction is strictly dependent on an increase in the calcium ion concentration in the cytoplasm [110]. This elevated calcium concentration is mediated by release from internal stores as well as fluxes of ions through conductive elements that span the cell membrane [5]. These conductive elements include ion channels and ion pumps. The flow of ions through ion channels is a passive process, with ions flowing down the electrochemical gradient. In general, each type of channel has an appreciable permeability for just one species of ion, a phenomenon known as selectivity. Of particular interest here are ion channels that are permeable to the ionic species calcium (Ca^{2+}), potassium (K^+), and sodium (Na^+) [37].

Electrical activity of the MSMC is characterised by a large positive deflection of the electrical potential difference across the cell membrane. This temporary excursion of the transmembrane potential difference is known as the action potential [37, 56].

Cells differ with respect to the ease with which such an action potential can be evoked by inciting stimuli. The relative level of susceptibility is referred to as excitability. It depends on several factors, such as the cell membrane's complement of conductive elements, as well as the value of the membrane potential at rest [8, 112].

1.2.1 Dynamic processes that underpin the action potential

In the absence of stimuli that may evoke an action potential, the membrane potential attains a value that reflects a balance between diffusive processes mediated by the ion channels and electric forces acting on ions passing through these channels, in addition to restorative fluxes of ions against their electrochemical gradient which ensure that differences in ionic concentrations between the cytoplasm and the extracellular medium can be maintained [104]. At rest, the sodium-potassium pump transfers potassium ions into the

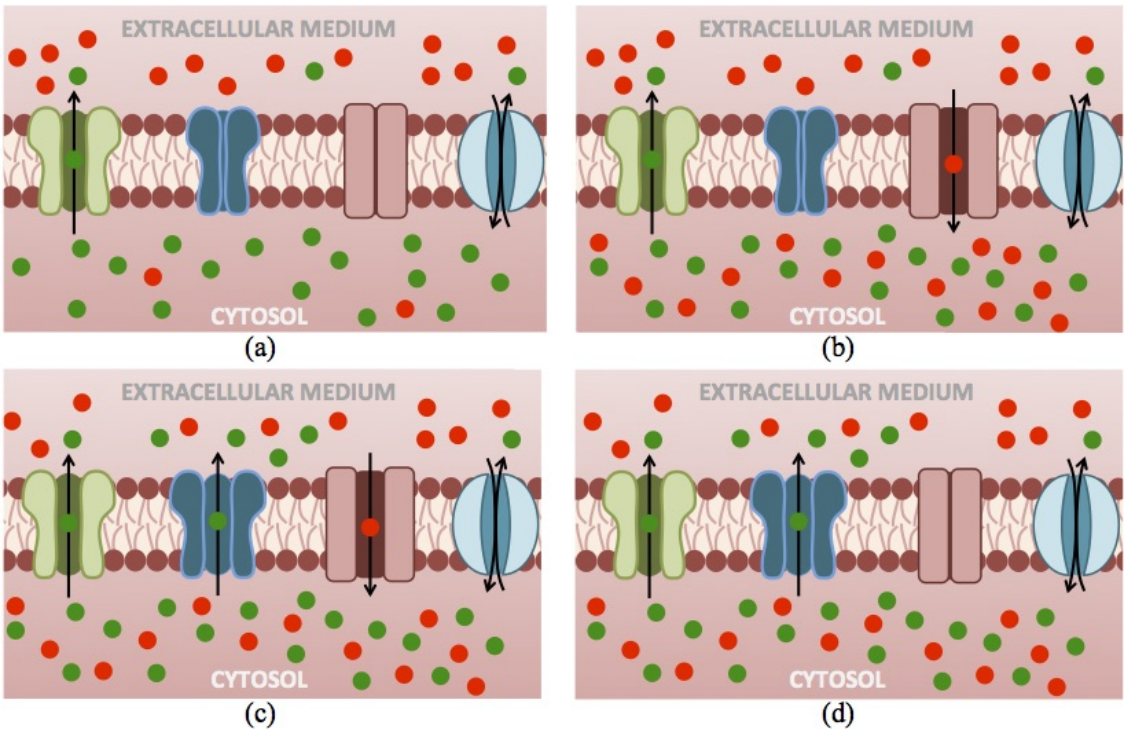

Figure 1.7: Gating events during an action potential. Potassium ions are indicated as green dots, calcium ions as red dots. A potassium "background leak" channel is shown in light green whereas a voltage-gated potassium channel is shown in dark blue. A voltage-gated calcium channel is indicated in mauve, and an active ion pump is indicated in light blue. (a) The resting state. (b) Voltage-gated channels open, leading to depolarisation which drives further channel opening, thus creating a positive feedback loop. (c) Voltage-gated potassium channels open and conduct a repolarising current. (d) A transient hyperpolarisation occurs as open potassium channels dominate the membrane potential. Adapted from Sheldon [98].

cell and sodium ions out of the cell [102]. This process moves two potassium ions for every three sodium ions. The exchange of sodium and potassium ions creates and maintains substantial concentration gradients in both ions.

The value of the membrane potential in this unexcited state is known as the resting membrane potential. The interplay between the gating dynamics of several different types of ion channels, conducting distinct ionic species, is what allows the membrane potential to depart temporarily from this resting value [55].

Ion channels are gated, that is, opened and closed, by specific stimuli, so that an individual channel may either be in a state where it allows the passage of ions through

its central pore, or alternatively blocks this passage. Gating stimuli include variations in membrane potential and binding to ligands, which can be external (hormones, neurotransmitters) or internal (cytoplasmic ions, second messengers) [54].

Voltage-gated ion channels respond to a change in the membrane potential in the vicinity of the channel. Sufficient variations in the potential induce conformational changes in the proteins that make up the channel and this can lead to a dramatic change in the ionic permeability (i.e., conductivity) of the channel. Ligand-gated ion channels change conformation in response to the binding of a specific binding partner [5].

If a cell is subjected to a sufficiently large increase in membrane potential, this triggers the opening of voltage-gated calcium and sodium ion channels. The opening of these channels lets calcium and sodium ions flow into the cell. This leads to a further depolarisation of the cell membrane, since the ion channels that allow these ionic species are themselves voltage-gated; thus closing a positive feedback loop that underpins an explosive excursion in the value of the membrane potential.

At the same time, however, the depolarisation also affects voltage-dependent potassium ion channels, which open. In addition, the favourable voltage-gating of the calcium channels ceases at elevated values of the potential. Both effects lead to a balance of permeability between the various ionic species that give the net current a repolarising character. Thus the membrane potential returns to its resting value and in fact becomes temporarily lower. During this hyperpolarised phase (or "undershoot") the resting balance between the various conductivities and pumping actions re-establishes itself.

1.2.2 Excitation-contraction coupling

Excitation is a membrane-localised effect, whereas contraction is mediated by the actin and myosin molecules located in the cell's interior. For one to lead to the other, there has to be a mechanistic connection, and in fact, this link is formed by the rapid influx of calcium ions which occurs during the action potential (Figure 1.8) [110].

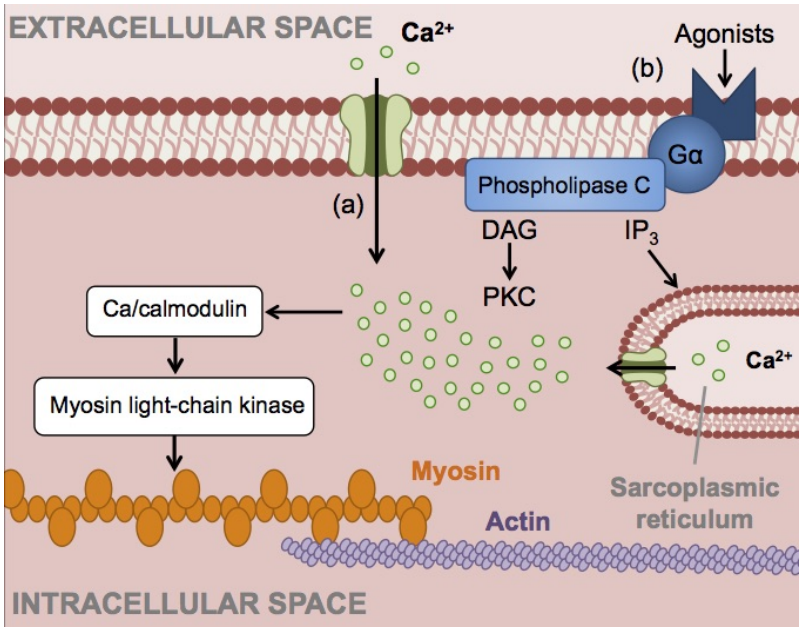

Figure 1.8: Intracellular calcium ions bind calmodulin, which subsequently activates the myosin light-chain kinase. The light-chain of myosin is phosphorylated, leading to contraction. Calcium ions enter the cell through two pathways: calcium entering via voltage-gated calcium channels; and endocrine stimulation of G-protein-coupled receptors. Image source: Sheldon [98].

Calcium ions in the cytosol bind calmodulin, a messenger protein, which in turn binds the myosin light-chain kinase. This activates the myosin light-chain kinase, which phosphorylates myosin light-chains. The phosphorylated myosin then interacts with actin filaments, allowing them to pull together. This pulling action sustains the contraction force that is exerted by the muscle as a whole [11, 77].

The cytosolic calcium contraction is further affected by a hormone-triggered signalling cascade [14]. If the first messenger oxytocin is present in the extracellular medium, it may bind to G-protein-coupled receptors which are embedded in the cell membrane. This binding activates the membrane-bound enzyme phospholipase C (PLC). PLC cleaves the membrane-bound lipid phosphatidylinositol 4,5-bisphosphate (PIP_2) into diacyl-glycerol (DAG) and the water-soluble molecule inositol 1,4,5-trisphosphate (IP_3). The DAG component can then activate protein-kinase C, an enzyme that inhibits myosin light-chain phosphatase, leading to an increased ability to contract. The IP_3 component diffuses into

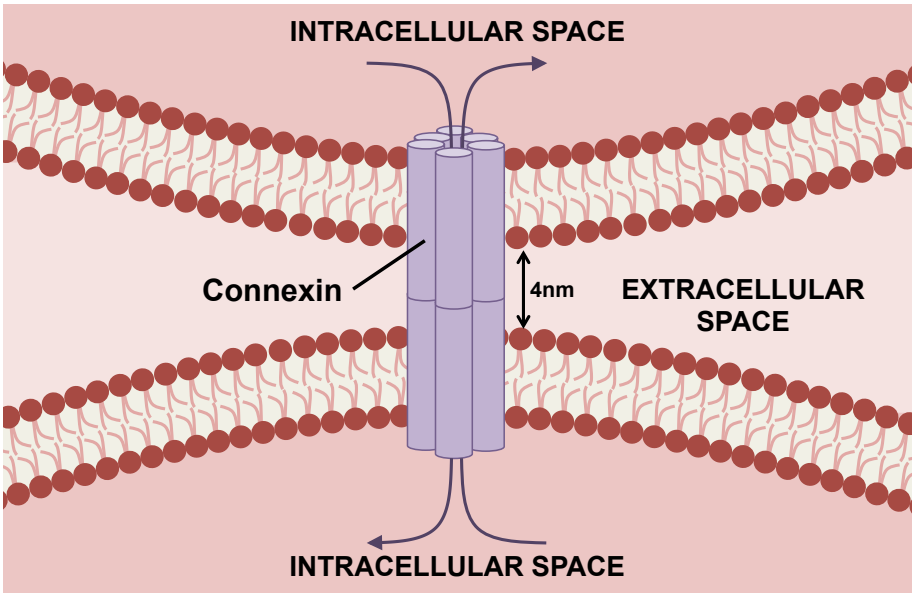

Figure 1.9: Structure of the gap-junctional channel, to which each cell contributes six connexin proteins. The pore formed by the connexin proteins allows the transfer of ions and small molecules between adjacent cells. Image source: Sheldon [98].

the cytosol, where it can interact with receptors located on the sarcoplasmic reticulum (SR), these receptors are also ion channels that allow calcium to flow out of the SR, diffusing down an extremely steep concentration gradient that is normally maintained between the cytosol and the lumen of the SR.

1.2.3 Gap junctions

Excitation signals are passed between cells through gap junctions [46]. A gap junction is an assembly of a large number of pores connecting the intracellular spaces of two cells. Each of these pores is formed by a gap-junctional channel, which is composed of two sets of six connexin proteins, one set from each cell, as illustrated in Figure 1.9. This structure forms a channel through which ions as well as small metabolites and second messenger molecules can diffuse from one cell to its neighbour. Thus, the gap junctions also form an electrical connection between the two adjacent cells. These connections are electrically passive and are also known as electrotonic couplings [6].

The conductance of a network of cells is determined by the density of gap junctions in the network and the shape of the cells [44, 78]. The resistance of the intracellular medium is negligible as compared to the resistance of the gap junctions [61]. This means that the magnitude of conductivity of a network of cells is determined at each point by the density of gap junctions. The directional components of the conductivity, however, are determined by the orientation of the individual cells at the given point. Conduction along the major axis of the cells requires passing current through relatively few gap junctions. Conduction perpendicular to the major axis, however, requires passing current through many more gap junctions, due to the fact that the cell membranes are closer together in that direction. This anisotropy gives rise to the conductivity tensor, which is a key concept in the present thesis and which will be discussed at length in what follows.

1.2.4 Regulation of excitability

Contraction is regulated by a host of different factors which are summarised in Figure 1.10. Endocrine, neurogenic, and myogenic factors can all modulate the excitability of the myometrium. These factors have been reviewed in depth by Garfield [43, 45] and Sheldon *et al.* [100].

Of particular interest in this thesis are the structural components of regulation, that is, the micro-architecture of the fibres and bundles that make up the functional units of the uterus as a contracting muscle. These structural components, together with the prevalence and conductivity of gap junctions, are what determine the conductive properties of the myometrium. The conductivity determines the ability of excitation propagation through electrochemical signalling.

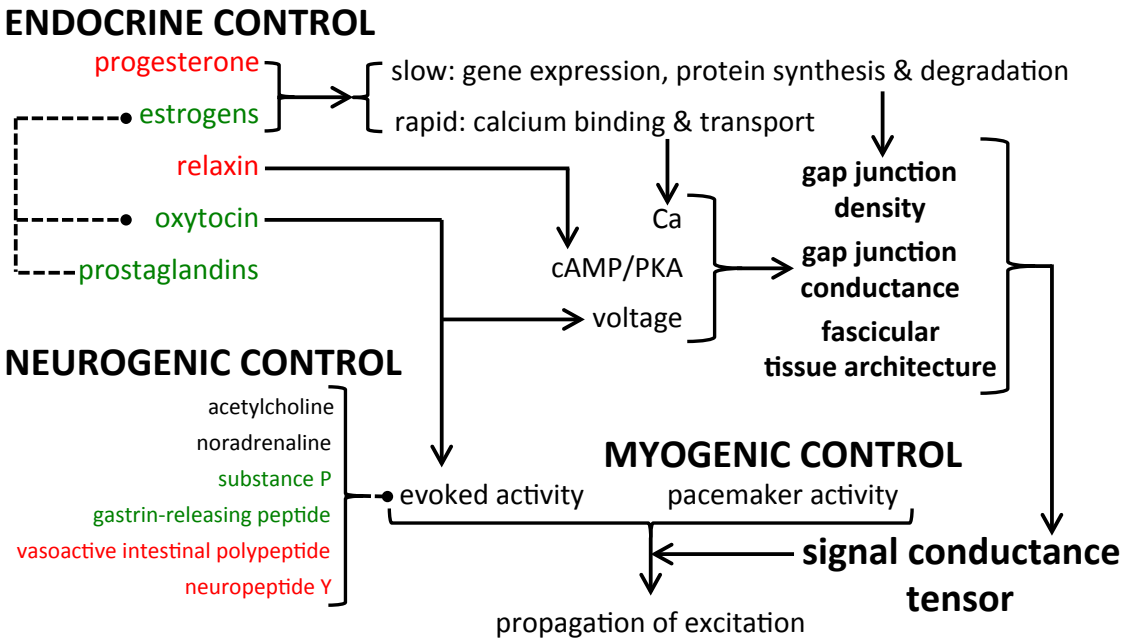

Figure 1.10: Factors which regulate contraction in the uterus. Green indicates factors which are up-regulated for labour, while red indicates factors which are down-regulated. Of particular interest in this thesis is the role of tissue architecture and how this affects the propagation of the action potential through the tissue.

1.3 The histological fine structure of the myometrium

The myometrium is a layer of smooth muscle tissue in the uterus which contracts during labour; without myometrial contraction, parturition could not take place [16]. The muscle fibres in this tissue are organised in distinct inner and outer layers, which are separated by a layer of vascular connective tissue [78]. In rat myometrium these layers form orderly structures, as illustrated in Figure 1.11. The inner layer, known as the circular layer, is comprised of rings of fibres surrounding the uterine cavity [18]. The fibres in the outer layer form roughly parallel tracts running the length of the uterine horn. This outer layer is known as the longitudinal layer. Sparse evidence points to the existence of fibres connecting these two layers, where possible communication between the layers has been observed [24, 86]; however, no such structures have been identified.

In marked contrast to the rat uterus, the human myometrium makes at first a much

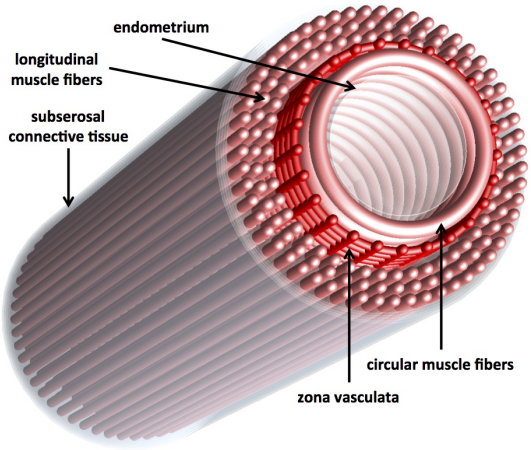

Figure 1.11: The basic structure of rat myometrium. The inner layer is formed of circular rings of fibres. A similar formation of the inner myometrium has been observed in the human uterus [111]. The longitudinal fibres in the outer myometrium typically form roughly parallel lines down the length of the uterus. The human myometrium differs greatly from that of the rat in the outer layer. This difference is characterised by fibres situated in a more disorderly formation as opposed to the uniform longitudinal pattern shown here. Between these two muscular layers is a layer of vasculature, which is present in both rat and human uteri. Image source: Sheldon [98].

more disorderly impression. It has been shown by Weiss *et al.* [111] that while the inner myometrium is comprised of circular fibres analogous to those found in the rat myometrium, the outer myometrium has little apparent structure.

The fibres within the myometrium are further divided into sub-fibrous structures, as has been demonstrated by Young and Hession [116]. The schematic shown in Figure 1.12 illustrates this subdivision. Each fibre is composed of fasciculi, which are smaller fibrous structures, and sheets. These fasciculi are joined to one another by connecting bridges, which are thought to enable excitatory signals to propagate between fasciculi. It therefore stands to reason that the density of these bridges play a key role in regulating the extent of connectivity in the fibre network.

The fasciculi are further subdivided into individual bundles, which are again smaller fibrous structures. The bundles within a fasciculus are also connected by bridging structures. These bridging structures determine connectivity within a fasciculus, which one would presume to control the level of synchronicity of excitation.

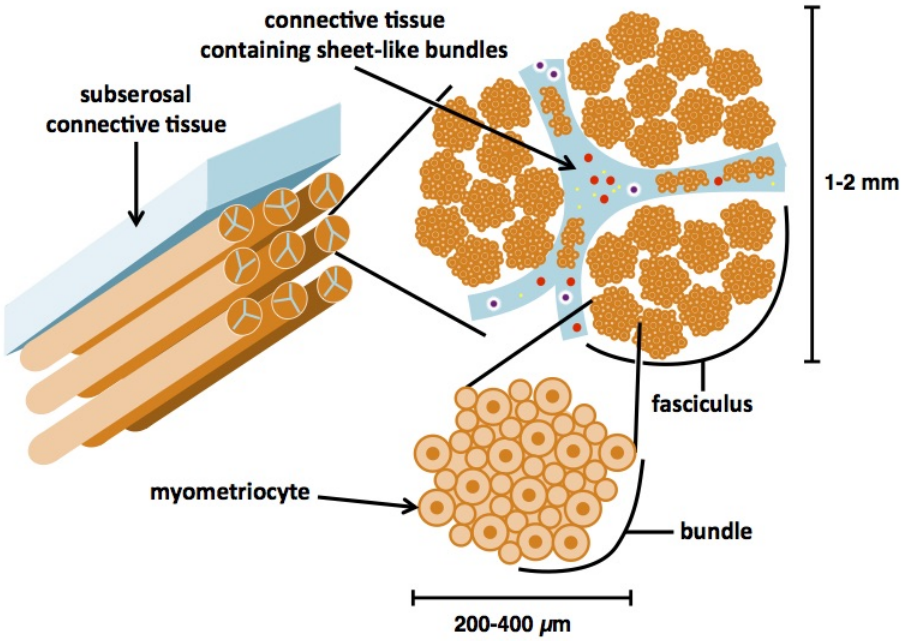

Figure 1.12: The general hierarchy of fibre structures in the myometrium according to Young and Hession [116]. Each bundle has width in the range 200–400 μm and is comprised of several smooth muscle cells, aligned with the direction of the bundle. These bundles cluster together to create larger fibrous structures known as fasciculi, which are themselves clustered, separated by sheets, to form fibres that have a typical width 1–2 mm. At each level there are bridges which allow excitation to propagate from one structure to another. Adapted from Young and Hession [116].

It is generally accepted that the level of connectivity on each of the above scales is an important factor in determining how the tissue as a whole functions. A poorly connected network of fibres would prevent excitation propagating throughout the network. Conversely, it has been shown by Sheldon *et al.* [99] that a high level of connectivity within an excitable network can prevent full excitation of the network from a stimulus. It is therefore important to determine the micro-architecture, that is, the detailed structure of the tissue at the cellular level, in order to determine the extent to which excitation can propagate through the network of fibres.

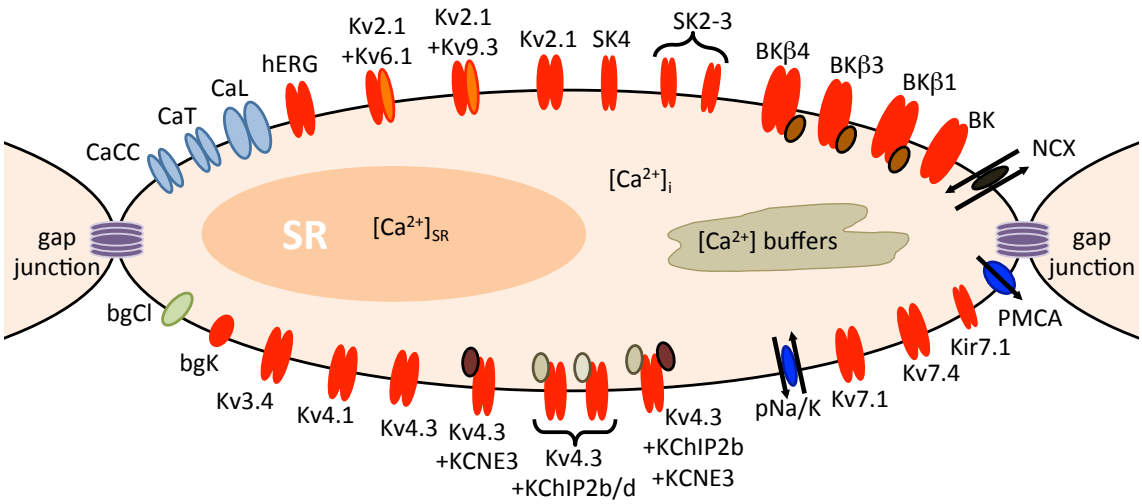

Figure 1.13: The full set of conductive elements present in the membrane of human MSMCs as presented by Chan *et al.* [22]. CaCC: calcium activated chloride channel; CaT: T-type calcium channel; CaL: L-type calcium channel; hERG: Human *ether-à-go-go*-related gene potassium channel; Kvx.x: potassium ion channels; SKx: small unitary current potassium channels; BKx: big unitary current potassium channels; NCX: sodium-calcium exchanger; PMCA: plasma-membrane calcium exchanger; Kirx.x: inwardly rectifying potassium channel; pNa/K: sodium-potassium exchanger; bgK: background potassium current; bgCl: background chloride current; [Ca]: calcium concentration. Image source: Chan *et al.* [22].

1.4 Mathematical modelling of excitation

The ion channels, ion exchangers, and other electrogenic proteins that are expressed in MSMCs have previously been identified by Chan *et al.* [22] and Tong *et al.* [107]. A chart summarising the "conductome" given by Chan *et al.* [22] is given in Figure 1.13. In principle, this cellular model could be employed as the fundamental unit of a detailed biophysical model of the dynamics of the membrane potential in a network of MSMCs. Such a model would incorporate the known gating kinetics of all these entities [10, 76, 107]. Since the gating characteristics of the gap-junctional channels (connexons) are also known [80, 100], one might assume that it is possible to model the propagation of electric activity in a biophysically detailed model of the entire MSMC network, by coupling numerous individual cell models together.

Unfortunately, two considerations stand in the way of such a direct approach. One of

these is theoretical, whereas the other is of a practical nature. The practical reason is that numerical simulation of a detailed model of a large number of cells greatly exceeds available computational resources. For example, simulation of the gating kinetics in a 1 mm^3 portion of tissue would require storing $\sim 10^9$ variables and parameters in active memory, taking up a total of ~ 4 gb. Modelling a excitation in a whole rat uterus at this level of detail would require more than 100 times this figure, while requirements for the human uterus would be even greater.

The theoretical difficulty is that while the conductome at the tissue level is known, this is not the case at the single-cell level; indeed, there is good reason to suppose that there exists considerable spatial heterogeneity in the expression levels of the channels represented in Figure 1.13. In particular, the overwhelming degree of functional redundancy inherent in this repertoire of ion channels, notably the K^+ channels [10], strongly suggests that individual MSMCs express at any given moment in time only a subset of the list identified by Chan *et al.* [22]. It is even conceivable that an individual MSMC cycles through distinct subsets as it turns over its ion channels and changes the relative levels of mRNA expression for these various types. Little remains known in detail at the present point in time.

In view of these impediments to detailed biophysical modelling, this thesis resorts to a highly simplified model of excitable tissue. In defence of this decision the arguments may be advanced that it is a well-established and popular generic model of an excitable system [9, 13, 26, 64], and that such a generic model may reasonably be expected to suffice for the present purposes, which are to investigate how tissue micro-architecture affects the patterns of activation wave propagation.

The model in question is the FitzHugh-Nagumo model [40]. It can be shown that this is a reasonable low-dimensional dynamical system that inherits the key qualitative properties of a high-dimensional biophysically detailed model [61]. The dynamics of this

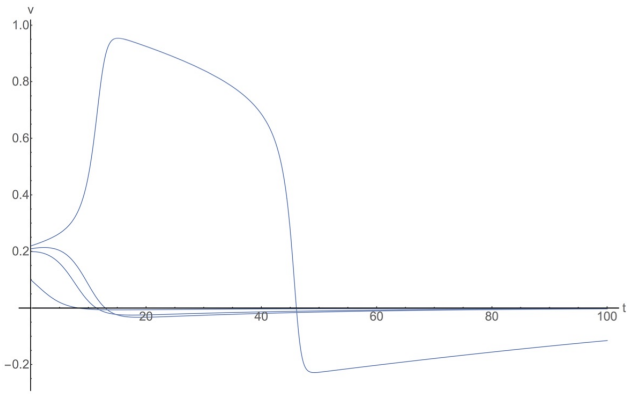

Figure 1.14: The behaviour of v in the FitzHugh-Nagumo equations. Each line represents the curve of v (y-axis) plotted against time (x-axis) for a given initial stimulus, given by the value at time 0. The value of α in this case was taken to be 0.2. If $v(0) \leq \alpha$, then v returns to 0 without any further increase in value. The top curve has $v(0) = 0.22$, and shows the full action potential.

model is expressed by the following equations:

$$\frac{\partial v}{\partial t} = f(v) - w$$

$$\frac{\partial w}{\partial t} = \epsilon(v - w - v_0)$$

where $\epsilon > 0$ (in the absence of spatial structuring these are really just ordinary differential equations, but rates of change have been represented as partial derivatives in anticipation of what follows). In these equations, $f(\cdot)$ is a function that remains to be specified. A standard choice is the following cubic function [61]:

$$f(v) = Av(v - \alpha)(1 - v)$$

and A, α, and v_0 are constants with $0 < \alpha < 1$. A detailed derivation of these equations will be given in Chapter 4.

The state variables v and w are both dimensionless. The correspondence between these two variables and those in more detailed models is that v corresponds to the membrane potential, whereas w corresponds to everything else. In a generalised sense, w represents

the "inertia" introduced by channel gating kinetics. Accordingly, v and w are referred to as the *fast* and *slow* variables.

The qualitative behaviour of these equations is well-captured by the term of art "excitation-relaxation cycle" (or even "excitation-relaxation oscillation" when the system's equilibrium point is unstable). Such a cycle can be equated, loosely speaking, to the depolarisation-repolarisation cycle that characterises the action potential. Suppose that the system is in equilibrium with $v = 0$ and $w = 0$. If v is perturbed any amount in the negative direction, then $f(v)$ becomes positive and the value of v is increased back to equilibrium. If v is increased a small amount such that $v \leq \alpha$, then $f(v)$ will acquire a negative value. This in turn will cause v to decrease back to 0. The situation where v is increased sufficiently beyond α, however, would result in $f(v)$ becoming positive. A large excursion to the other branch of the cubic then occurs, as shown in Figure 1.14; the qualitative kinship to an action potential is obvious.

A spatially structured version of this model is required to simulate an extended network of MSMCs. Current passes between cells in a network through the gap junctions joining the cells [44]. For a large network of cells, a continuum approach is reasonable [61], that is, the conductivity between cells is approximated by a continuous function of space. The following system of partial differential equations, known as the electrodiffusion equations, represent such a system:

$$\frac{\partial v}{\partial t} = \nabla D \nabla v + f(v) - w$$

$$\frac{\partial w}{\partial t} = \epsilon(v - w - v_0)$$

where D is the conductivity tensor. The conductivity tensor is determined by two factors: (i) the density of gap junctions in cell membranes, and (ii) the fibre direction of the tissue at a given point. Moreover, on the additional assumption that the gap junction density is homogeneous throughout the tissue, it is possible to determine the excitation behaviour of the tissue from the structure alone for a given conductivity constant. In this thesis,

computer-aided histological image analysis will be used to determine this structure and to derive the tensor D from serial sections of myometrial tissue.

1.5 Image analysis and histology

Various techniques in image analysis and their application to histological inference are central to the results presented in this thesis. This section reviews several basic ideas and discusses the caveats and pitfalls.

1.5.1 Images: greyscale, thresholding, convolution, and edge detection

A fundamental idea in image analysis is granularity. Even though an oil painting may appear a continuum of colours merging seamlessly together, zooming in sufficiently we would find that the picture is ultimately composed of pigment particles. One might wonder whether Nature itself is ultimately granular or smooth, but even if the latter, any method of creating an image (such as an optical system) will have inherent limitations on its resolution, which is the ability to distinguish nearby points as being distinct. These considerations lead to a standard abstract representation of an image as a regular two-dimensional array of picture elements ("pixels") each of which has the property that it assigns a single colour value to the region of space it represents. (Three-dimensional arrays of volume elements, "voxels," can also be considered as a type of image. Neither type of array has to be regular, but this is a convenient routine assumption.)

Accordingly, an 8-bit greyscale image I can be represented by a two-dimensional array, with entries $I(x, y)$ corresponding to grey values from 0 (black) to 255 (white) [83]. A colour image can be represented in the same fashion, but with each entry $I(x, y) = (r, g, b)$ encoding the red, green, and blue values. Each colour value lies between 0 and 255 in the same manner as the grey image. Taking each colour value separately yields a set of three

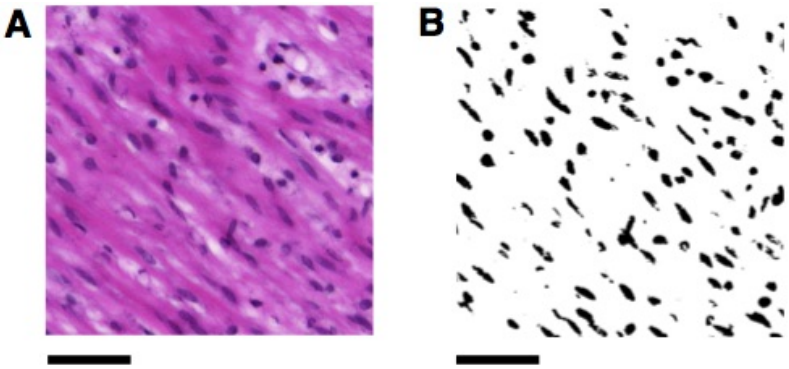

Figure 1.15: The effect of thresholding on a colour image. **A**: Histological image with nuclei stained dark purple and surrounding tissue stained pink. **B**: Binary image generated by applying a threshold to the individual colour values of **A**. The resulting image contains black pixels where the threshold condition is met for all colours. In this instance the colour thresholding has isolated the nuclei in **A**. Scale bars represent 50 μm. Figure 2.6 gives a more detailed representation of this thresholding procedure.

greyscale images, I_r, I_g, and I_b with the property that

$$I(x,y) = (I_r(x,y), I_g(x,y), I_b(x,y)).$$

This representation is used to handle each colour separately in the following chapters; this has the advantage of reducing all image-processing steps to the analysis of greyscale images. For this reason, the remainder of this section will only be concerned with the greyscale representation of images.

One specific type of image of importance is a binary image. This type of image has entries which only have values either 1 or 0. A binary image can be used to mark the position of particular features in an image; the forced choice between "all" or "none" implies that the features are set apart quite crisply. Thresholding provides a standard procedure to generate a binary image from a greyscale image [83]. Given a greyscale image I, the image I_t generated by applying the threshold t is defined by

$$I_t(x,y) := \begin{cases} 1 & \text{if } I(x,y) > t \\ 0 & \text{if } I(x,y) \leq t \end{cases}.$$

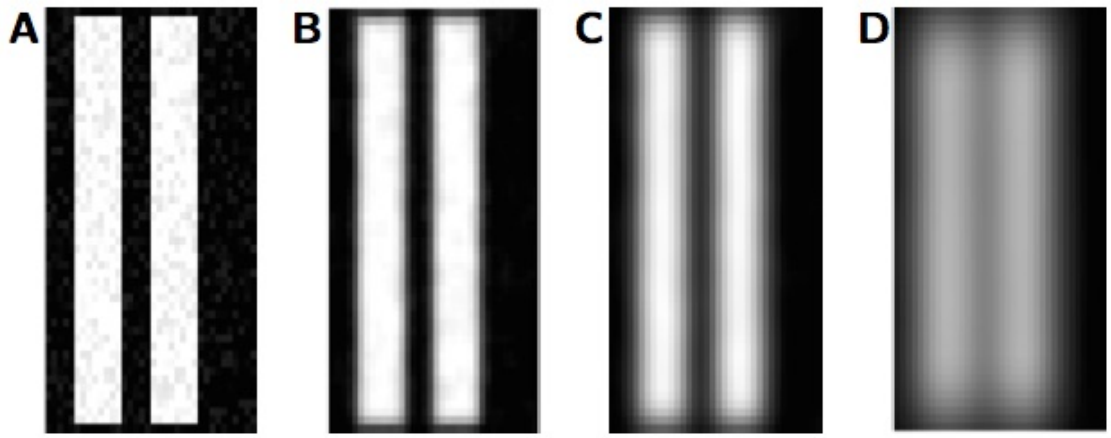

Figure 1.16: Gaussian blurring using differing values of σ. The original image **A** shows a pair of lines with some random noise added. These lines might represent myometrial smooth muscle fibres, for instance. In order to filter out noise, Gaussian smoothing is applied. The images produced in **B** ($\sigma = 2$) and **C** ($\sigma = 4$) represent low values of σ, and reduce noise without severely altering the image. The image in **D** used a much higher value of $\sigma = 8$, and has blurred the lines to become less distinguishable.

An example of the effect of applying a threshold to a histological image is given in Figure 1.15. This binary image is a useful representation of the original image because it clearly marks the pixels as being either part of a feature or part of the background.

A common technique used in image analysis is convolution [83]. Given a mask $M = (M(i,j))_{i,j=-r}^{r}$, the convolution of M with an image I is given by

$$(M * I)(x,y) = \sum_{i=-r}^{r} \sum_{j=-r}^{r} M(i,j) I(x+i, y+j).$$

Loosely speaking, the intuitive effect of applying a convolution is to suppress or emphasise certain features of an image, e.g. blurring, sharpening, detecting regions where pixel values change rapidly. Which of these effects is attained depends on the choice of M, and which of these effects is desired depends on the information that we wish to distil from the image. A technical difficulty is that this sum is not well-defined close to the edges of I; however, this problem can be handled through padding the images with a sufficiently broad outer edge of blank space so that the formula is valid everywhere.

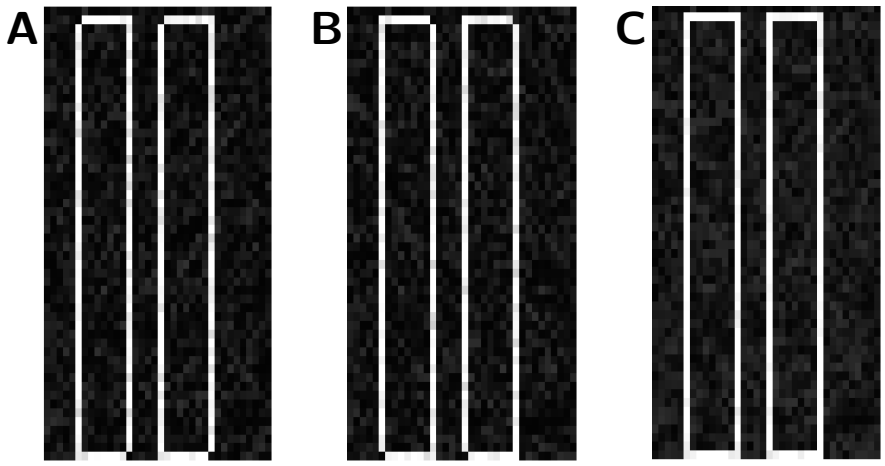

Figure 1.17: The Roberts cross operator applied to the original image in Figure 1.16. **A**–**B**: The images generated by applying M^+ and M^-. **C**: The image generated by taking the maximum value of the other two images at each pixel.

One important application of convolution is Gaussian smoothing [105]. An $n \times n$ Gaussian mask G with variance σ^2 is a square mask given by

$$G(i,j) := \frac{1}{2\sigma^2\pi} \exp\left\{\frac{-(i^2+j^2)}{2\sigma^2}\right\}.$$

This operator has the effect of blurring the image, as illustrated in Figure 1.16. This blurring can be useful for reducing the effect of noise prior to further processing of the images. However, smoothing also reduces the contrast of features within the image. The scale at which this reduction in contrast occurs is proportional to σ, as is evident in Figure 1.16. Therefore, when applying a Gaussian filter to an image the value of σ needs to be sufficiently small to maintain the contrast (and hence distinguishability) of the smallest features being measured, while remaining close to this limit in order to maximally reduce the noise in the image.

Another application of convolution is as a differencing operator. Consider the mask

$$M^- = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The convolution of M^- with an image I gives

$$(M^- * I)(x, y) = I(x + 1, y + 1) - I(x, y).$$

This is a discrete analogue of the derivative along the line $(1, 1)$ and the mask M^- is one of two templates used in the Roberts cross operator [93]. The other template, M^+, is the analogous matrix which computes the difference along the line $(1, -1)$. The image E given by

$$E(x, y) := \max\{(M^- * I)(x, y), (M^+ * I)(x, y)\}$$

gives the maximum of these two differences at each point. The effect of this operator is illustrated in Figure 1.17. This maximum is useful for determining edges of features; a point with a high value of E would correspond to a rapid change of colour in the image, which in many instances corresponds to the edge of a feature of interest.

More complex differencing operators can give rise to more precise edge detection algorithms. In particular, the Sobel operator [34] combines differencing with a local smoothing, as described in Appendix A. The Canny edge detector uses a combination of Gaussian smoothing and differencing, which further improves the edge detection process [21]. The algorithm for computing the Canny edge detection is discussed in full in Appendix A.

Whereas differencing operators are useful for generating images that can be used for registration (see below), they do not necessarily determine the entire contiguous border of a feature [83], as shown in Figure 1.18: there are breaks and gaps where the human eye can readily tell that the border line ought to have been continuous.

A key challenge in the *in silico* reconstruction of myometrial tissue micro-architecture is to define bundle-like structures. One obvious approach is to use the edges of the bundle as the defining feature. If the boundary of a bundle is incompletely marked, however, it becomes impossible to determine the precise set of points within the bundle. Furthermore, suppose that two bundles are situated adjacent to one another, without any conducting

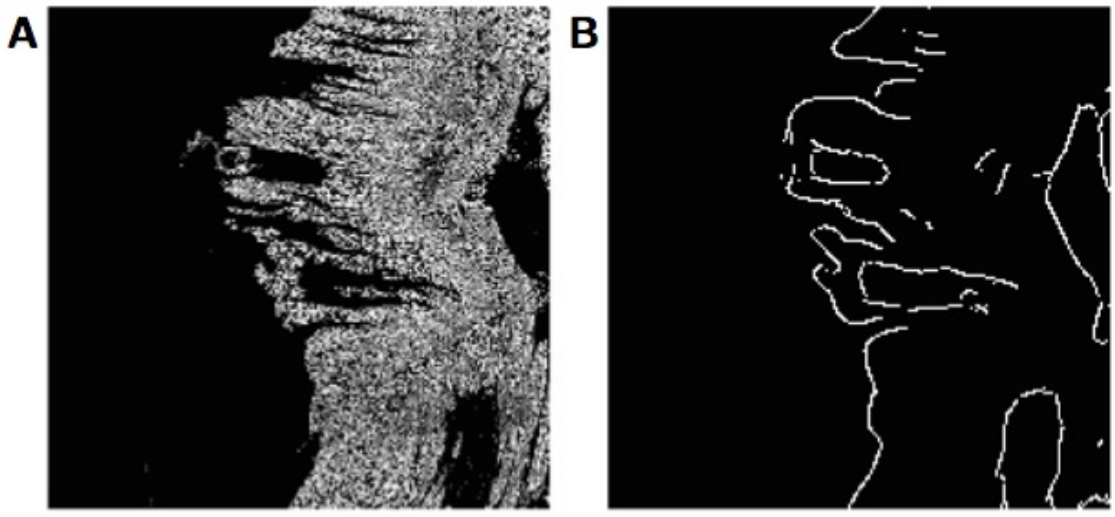

Figure 1.18: An example of a differencing edge detector failing to accurately determine edges. Here the Canny operator was used to determine the edges of **A**. The resulting edge image in **B** has lost some parts of the edge. Generated using the Canny edge detector algorithm given in Appendix A.

bridges forming electrotonic connections. From a physiological point of view, activity cannot propagate from one bundle to the other. However, if the inferred boundary between the bundles is incomplete, it would be possible for a simulated excitation wave to pass through the gap. "Gappy" edge detection would thus give rise to wave propagation artefacts. In view of this problem, an alternative method has been developed, as will be discussed in detail in Chapter 3.

1.5.2 Registering serial sections

A key concept for generating three-dimensional structure from a series of two-dimensional images is registration. Image registration is the process of aligning two images by matching either areas or features of the images [118]. These two approaches to registration, area-based and feature-based, both have strengths and weaknesses and the nature of the images being registered decides which of the two is the most appropriate.

In the context of serial histological sections, one might imagine that the tissue can be tagged with some kind of marker structure before section, such that it leaves a unique

(and uniquely orientable) feature on all slides. Then all the image analyst has to do is to match up these features on subsequent slides. Unfortunately, due to tearing and plastic deformation of the thin paraffin sections that can occur before they are safely affixed to their glass slides, the problem of histological registration is considerably more complex.

A common form of an area-based approach to registration is the comparison of the Fourier transform of each image [30]. The objective of such an algorithm is to transform one image in the spatial domain and to do this in such a way that the Fourier transform matches that of the reference image. An advantage is that no specific features need to be found in the image prior to registration. It might be expected that this makes these area-based methods more generally applicable than feature-based approaches, since they do not require information specific to the type of image in order to function. However, by the same token, images that would be judged to be very different from a feature-based point of view may well have quite similar Fourier transforms which could lead to good matches that are nonetheless false. In particular, a major drawback of area-based methods as applied to the registration of histological images is that they require some level of spatial heterogeneity in the image [118]. Images of smooth muscle tissue in particular generally lack this required heterogeneity to enable precise registration using area-based techniques.

Feature-based registration does not require spatial heterogeneity in the images themselves, provided that there is adequate distinction in the position and shape of the features. The basic concept in feature-based registration is to transform an image so that its transformed features constitute a close approximation to the features in the reference image [118]. This can be rephrased in terms of shape detection; if there is a unique feature in an image with a characteristic shape, then the registration algorithm aims to find a close approximation to a spatially transformed version of this shape in the reference image. A standard shape detector is the generalised Hough transform [33], which will be discussed in detail in Chapter 2.

Feature-based registration is applicable when features can be easily identified. A class

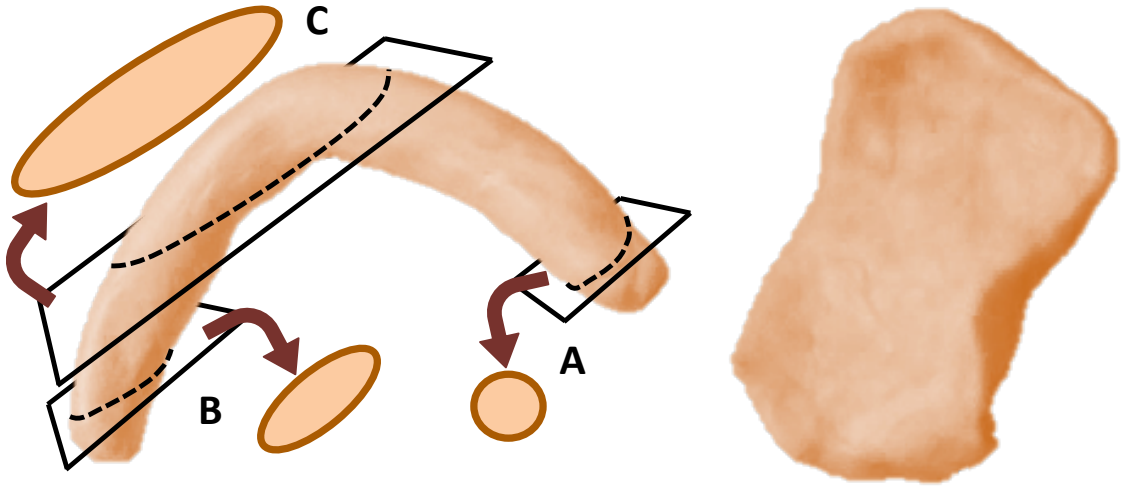

Figure 1.19: The angle of the plane of section relative to the orientation of the three-dimensional object (left) greatly affects how it appears on two-dimensional cross-section (examples **A**, **B**, **C**). Moreover, the cross-sections do not uniquely determine the shape of the object, since an entirely different object (right) may yield similar cross-sections.

of features that is commonly used for registration is that of edges determined by means of an edge-detecting algorithm [118]. This feature is useful for registering serial sections of tissue, inasmuch as the edge (outline) of the tissue can be considered roughly constant between two sections. An example of a feature specific to smooth muscle images is that of the nuclei of the smooth muscle cells. These nuclei are ideal for shape-matching because at sufficiently high resolutions they can be viewed as shapes in themselves. Moreover, they are often conveniently similar to an ellipse, which is a simple geometrical shape. Such additional shape data can allow for a higher accuracy in registration than can be achieved with position alone.

1.5.3 Challenges associated with three-dimensional interpretation and reconstruction

Histological sections are extremely thin and may effectively be regarded as two-dimensional slices through a three-dimensional structure within a biological tissue. Microscopists have long been familiar with the challenge posed by the interpretation of these two-dimensional

images in terms of the actual structure from which they have been derived [58]. Unfortunately, these challenges are inherited when an attempt is made to reconstruct these structures as data-objects in a computer. Consider for instance the curved cylindrical ('bent rod') object in Figure 1.19. The cross-sections will generally be elliptical, with circles as one limiting case and quite elongated shapes at the other extreme. It can be seen that greatest dimension of these two-dimensional shapes is a poor guide to the actual size of the object; ideally, the latter would be quantified by a parametric curve that goes through the centre line of the cylinder, associated with a function that specifies the diameter of the cylinder along every point of this mid-line curve. Only the least eccentric ellipses give a reliable impression of that width, and in other cases the minor axis is the best estimate, but the latter is in general only a lower bound since the plane of sectioning need not touch the centre line.

A different type of problem is that a qualitatively different kind of structure, such as a sheet ('pillow') as shown on the far right in Figure 1.19, can also present itself as circular or elliptical on cross-section, depending on the angle of sectioning. In some cases, when information is available in advance about the dominant directions of the structures of interest in a given piece of tissue, the microscopist may be able to orientate the paraffin or plastic block, in which the sample is embedded, relative to the microtome so as to obtain the most informative sections. This technique has proven to be effective for a broad range of tissues (e.g., brain, heart, liver, kidney) but less so in tissues such as the myometrium, in which fibres run in several directions.

A consideration of key importance in this thesis is the correct reconstruction of what might be termed the *topology* of the myometriocyte network. This network consists of discrete structures that split and merge [115]. Identifying regions where tubular or sheet-like structures come together is essential for a correct prediction of excitation wave propagation within the tissue. Figure 1.20 shows a sculpt suggesting such a region, and a series of cross-sections (Figure 1.20**A**–**C**) in which that three-dimensional structure is not at

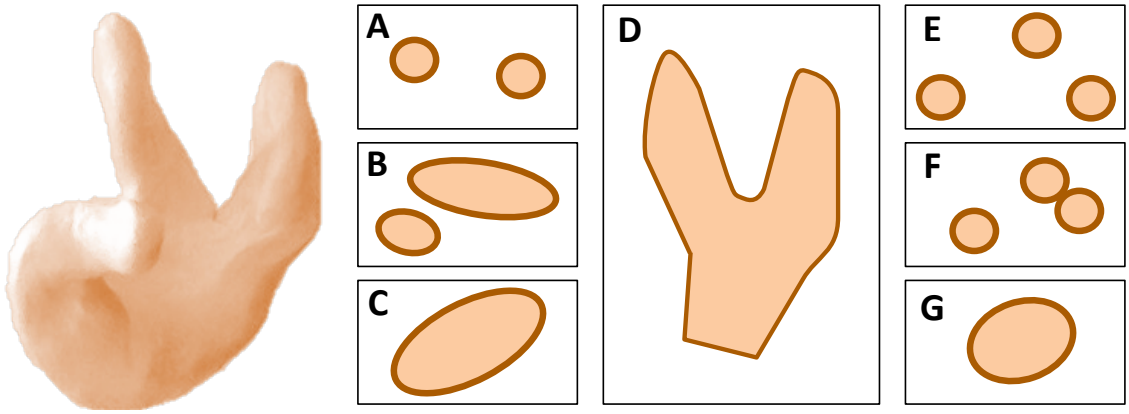

Figure 1.20: The spatial relationships between the components of a complex three-dimensional object (left) may not be apparent in a series of cross-sections (**A**, **B**, **C**). The bifurcation will only be apparent if it is in-plane with the sectioning microtome (**D**). Only a series of cross-sections within a narrow range of sectioning angles will make the trifurcation obvious (**E**, **F**, **G**).

all obvious. These difficulties are only worsened if the structure exists in a tissue neighbourhood that is closely packed with fibre-like structures belonging to distinct conducting tracts. A further layer of problems is associated with image registration (relative deformation of subsequent slides). The combined effect of all these difficulties is that some of these circular cross-sections are at risk of being incorrectly assigned.

In the example of Figure 1.20, there is no (planar) section that would capture the trifurcation, and only fortuitous planes of section (Figure 1.20**D**) are capable of representing one of the bifurcations directly in a single image. It is not altogether impossible that the trifurcation is immediately obvious from serial sections (Figure 1.20**E**–**G**) but only if the angle of section happens to be just right.

Even when fibre-like structures do not split or merge, their irregular shape may hamper the correct assignation of sections to three-dimensional shapes, particularly in cases where they come into close contact, as suggested by the sculpt on the far left of Figure 1.21. Feature-based registration used relies on classifying features as "the same" over subsequent images (this amounts to the assignation of cross-sections to a set of objects). The inherent (and true) variation in two-dimensional shape may lead to an incorrect assignment. The result might be that the wrong 'limbs' end up being stitched together by the computer,

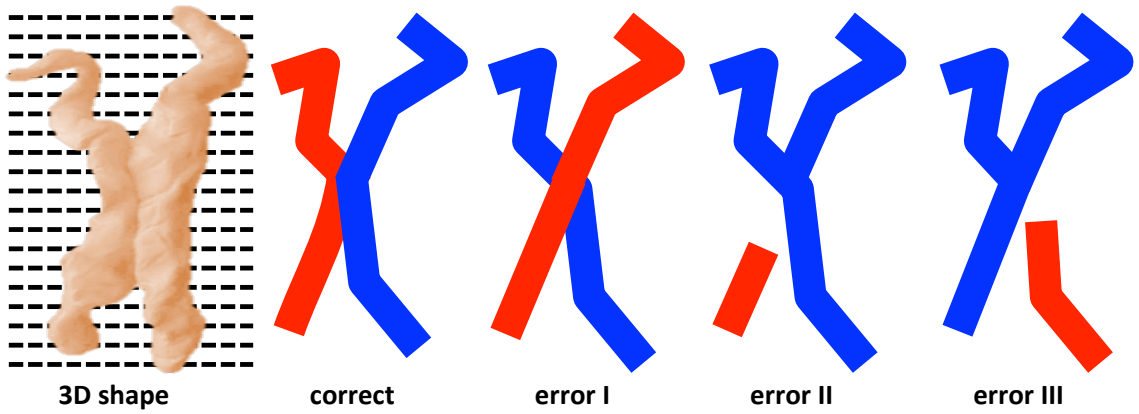

Figure 1.21: Close apposition of irregular three-dimensional objects may cause difficulties in reconstructing these shapes from serial sections. Incorrect assignment of cross-sectional areas between subsequent sections (along the dashed lines) leads to an erroneously inferred crossing (error I); incorrect segmentation at the point of close approximation can lead to erroneously inferred bifurcations (error II and error III).

leading to a reconstructed chiasma that is not real.

Alternatively, a segmentation error at the level of the close juxtaposition can lead an image-processing algorithm to impute a single cross-section at that level where in actual fact there were two adjacent but distinct shapes. This could have repercussions down the image-processing pipeline, such as an inferred bifurcation where none exists in reality.

Chapter 2

Reconstruction of the myometrium

This chapter describes the process of determining the three-dimensional structure of the myometrium. The structure was found through histological inference of serial sections of the tissue, using computational image analysis on the histological sections to determine fibre direction. This structure is a representation of the fibre directions in the tissue block at a resolution of ~ 50 μm per voxel length. The computational process was semi-automated, with the only manual input required being to check the integrity of slides after registration.

Section 2.1 presents the materials and methods used for the processing of tissue prior to computational analysis. Section 2.2 gives a summary of the number of slides used to generate the *in silico* reconstruction and the overall volume of tissue sampled in this sectioning. The analysis of the images generated is detailed in Section 2.3. This section describes how the two-dimensional fibre direction is obtained from the histological images. The registration of the slides based on this two-dimensional structure is described in Section 2.4. Section 2.5 describes techniques for combining these registered slides to form an evenly-spaced three-dimensional structure. The final step in determining the three-dimensional fibre structure is presented in Section 2.6. This uses the registered two-dimensional structures to generate a three-dimensional direction field. All code used in this chapter can be found at `https://figshare.com/s/651e592a91f309182edd`.

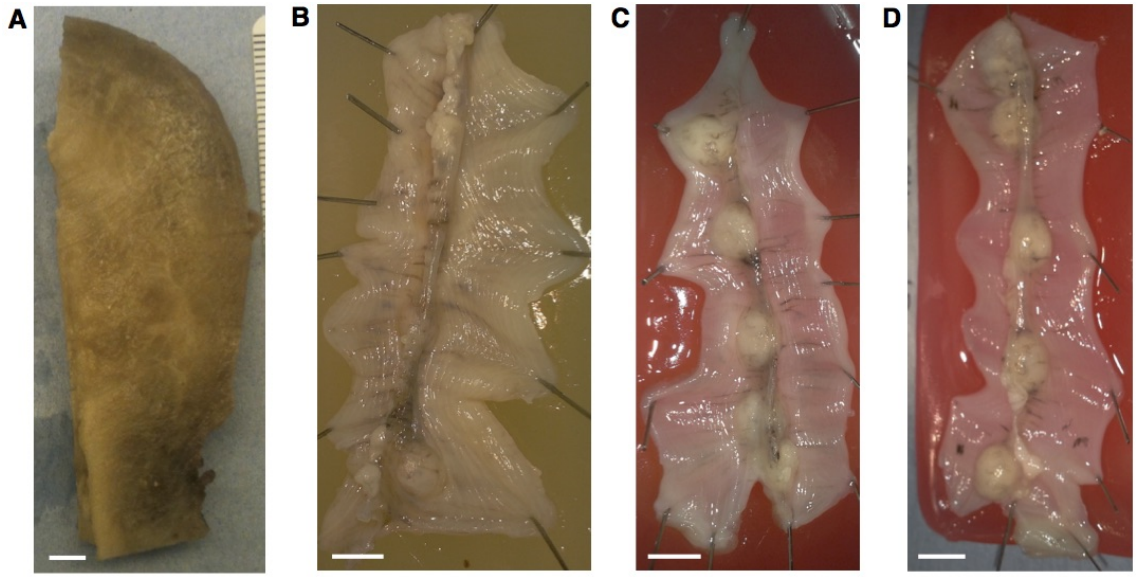

Figure 2.1: **A**: Human uterus tissue block, positioned with fundus at the top of the image and cervix at the bottom. Sectioning was performed in the orientation shown here, starting at the top (nearest) surface, up to a depth of 3.5 mm. **B**–**D**: Rat uteri processed, positioned with ovarian end at the top of the image and cervical end at the bottom. Each rat uterus was cut along the antimesometrial border, preserving the mesometrial border where implantation sites can be seen as yellow bumps near the middle of the tissue. All rat samples were sectioned in the orientation displayed, starting at the bottom (farthest) surface, and continuing through the entirety of the tissue. Scale bars represent 5 mm.

2.1 Tissue Preparation

All procedures on human tissue were conducted within the guidelines of The Declaration of Helsinki and were subject to local ethical approval (REC-05/Q2802/107). Prior to surgery, informed written consent for sample collection was obtained.

The experimental protocol for the use of rat tissue was approved by the Animal Research Ethics Committee, Faculty of Medicine & Health Sciences, United Arab Emirates University (institutional ethical approval: AE/03/30).

Three rat uteri and one human uterus were used. The human uterus was non-pregnant and a block of dimensions 7 cm × 3 cm × 0.35 cm was taken from the tissue for processing. This block was taken such that it included the length of the organ, as shown in Figure 2.1**A**. The full length, however, was slightly too large to fit on the slides used. For this reason

the cervical (caudal) end was clipped during sectioning. The rat uteri were obtained in the manner described below, and were subjected to electrode array recording prior to fixation. These steps up to fixation were performed by Wim Lammers as follows.

Eight virgin Wistar rats were time-mated, and pregnancy dated as day 0 of gestation if the sperm cells were observed in the vaginal lavage the next morning. On the day of the experiment, day 20 in the first two rats (Figures 2.1**B** and **C**) and day 19 in the third (Figure 2.1**D**), rats were euthanized by graded CO_2 inhalation, and the uterine horns were rapidly excised via a midline incision of the abdomen. The uterine horns were opened longitudinally along the mesometrial border and embryos removed, and the myometrium placed in a tissue bath containing modified Tyrode's solution. After equilibration the myometrium was pinned at resting length, serosal side facing upwards, to a final dimension of approximately 20×50 mm. The tissue was then perfused in modified Tyrode's solution at a rate of 100 ml/min. The modified Tyrode's solution had the following composition (mM): 130 NaCl, 4.5 KCl, 2.2 $CaCl_2$, 0.6 $MgCl_2$, 24.2 $NaHCO_3$, 1.2 NaH2PO_4, and 11 glucose, saturated with carbogen (95% O_2 - 5% CO_2) pH 7.35 ± 0.05, $37 \pm 0.5°$ C.

The electrophysiology experiments were performed as described in detail by Lammers *et al.* [67]. Electrical recordings were made using a custom rectangular 240-electrodes array (24×10; 2 mm inter-electrode distance), which covered the whole preparation. The electrodes consisted of Teflon-coated silver wires (0.3 mm diameter, Cooner Wire). Unipolar electrograms were recorded from each individual electrode with a silver plate located in the tissue bath serving as the common reference electrode [66]. All electrodes were connected through shielded wires to 240 AC amplifiers where the signals were amplified ($4000\times$), filtered (bandwidth 2–400 Hz), digitized (8 bits, 1 KHz sampling rate) and stored on a PC. Recordings were performed for 30 consecutive minutes. After the experiments, signals were digitally filtered (20-point moving average).

Both rat and human tissue samples were processed according to the chart in Figure 2.2. After recording, rat tissue samples were washed and fixed as pinned in 10% neutral buffered

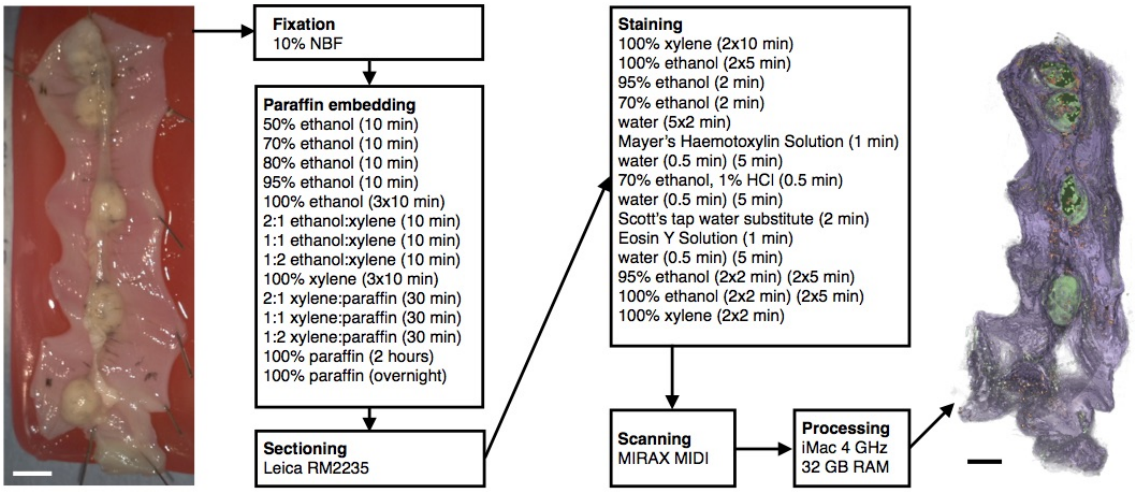

Figure 2.2: Steps performed to obtain *in silico* reconstructions from tissue samples. Left panel shows tissue prior to processing, right panel shows a representation of the final *in silico* reconstruction. Scale bars represent 5 mm. Equipment and solutions are shown in the relevant boxes. Series of graded solutions were required for paraffin embedding and staining, listed here in order of application. Steps using paraffin mixtures were performed within the range $54 - 58°$C to keep the paraffin in a liquid state without damaging the polymer structure.

formalin (NBF). The human tissue sample was likewise fixed in 10% NBF. All samples were dehydrated with graded ethanol solutions and washed with graded xylene solutions before the graded exchange of xylene with paraffin produced the paraffin embedded tissue blocks. The tissue blocks were sliced into serial sections 5 μm thick using a microtome (Leica RM2235). The slides were labelled with a unique identifier for the tissue block and numbered according to the sectioning order. Some sections produced were highly distorted through creasing or tearing and were therefore deemed unfit for processing and were discarded. Each such section was accounted for in the numbering of the slides. After discarding these slides, the number of remaining slides subjected to the following analysis was 284 in the human tissue block, and 221, 237, and 201 in the rat tissue blocks. All remaining slides were subsequently stained (automated using Leica Biosystems Autostainer) with haemotoxylin and eosin (H & E). The stained slides were scanned (scanner: MIRAX MIDI), producing virtual slides with pixel dimensions 0.465×0.465 μm^2 per pixel for one of the rat tissue blocks (Figure 2.1**B**), and 0.371×0.372 μm^2 per pixel

Step	Time per slide (rat)	Time per slide (human)
2D fibre orientation	19.9 ± 9.84	42.8 ± 8.76
Registration (rigid)	2.45 ± 1.69	1.98 ± 1.17
Registration (elastic)	7.72 ± 5.29	25.3 ± 7.76

Table 2.1: Mean time required to process a slide for the most time-consuming steps, in minutes. Standard deviation is given as the error value. Slides from the human tissue block contained more tissue than those from rat tissue blocks, which increased the processing time in processing histological images and elastic registration.

for all other blocks. The scanned slides were in MIRAX format, and were converted into TIFF files using either the MIRAX Pannoramic Viewer [1] or a combination of Vips [75], OpenSlide [50], and bigTIFFTools [32]. The conversion produced a set of tiles of size 2048×2048 pixels for each slide, where each tile represented a square of the overall image without overlap. These tiles were then used as the input for the ImageJ [2] plugin described below (Section 2.3).

The semi-automated processing pipeline for generating *in silico* reconstructions from these images was performed using an iMac with 4 GHz Intel Core i7 and 32 GB RAM. The processing steps that required the most computation time were the identification of two-dimensional fibre orientation (Section 2.3) and registration (Section 2.4), with all other steps taking a total of less than 2 hours for a tissue block. The time required for these slower steps is summarised in Table 2.1. Scanning times were not recorded, but approximately 30 minutes were required to scan each slide. Sectioning was performed at a rate of approximately 50 slides per hour.

2.2 Reconstructed volume bounds and slide representation

Not all the sections generated by the above procedure were viable for use in the reconstruction. The selection of slides was performed during the sectioning process, where slides were discarded if they were deemed distorted beyond repair, and during the registration process detailed in Section 2.4 below, where slides were discarded if they failed to accurately register. Additionally, the reconstructed volumes all have bounds that were created

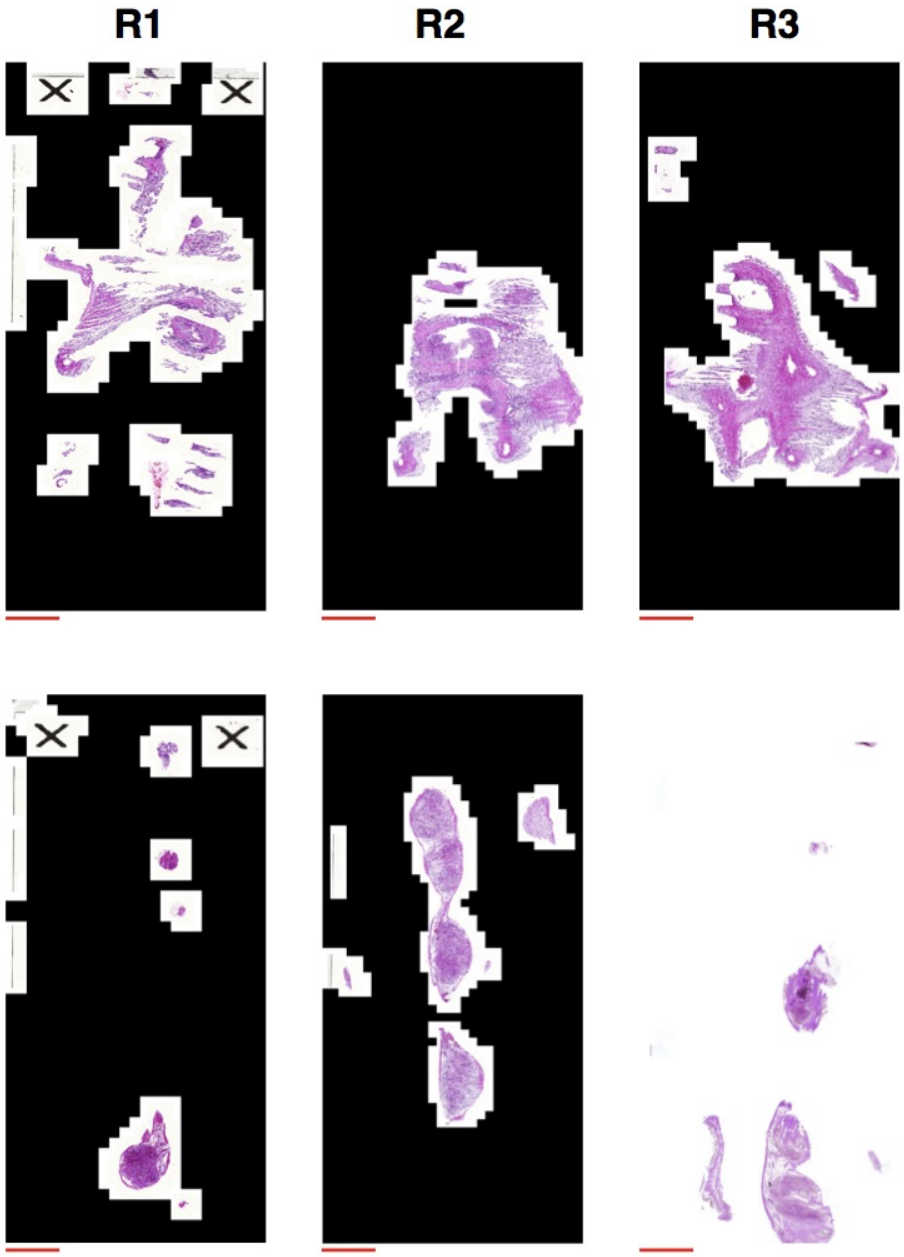

Figure 2.3: The first and last slides used in the reconstruction of the rat tissue blocks. The amount of tissue located prior to the first slide is unknown due to the lack of numbering prior to the first slide; however, the number of slides discarded prior to the first slide generated was less than 100, which puts an upper bound of 500 μm on the depth lost. The last slide for each tissue block represents the last slide which the registration algorithm could register. In each case, the remaining slides include tissue in and surrounding the placental beds. The amount of tissue remaining in each case is less than 500 μm, which was was determined the number of slides which could not be registered. Top row: first slides in each of the rat tissue blocks. Bottom row: last slides in each of the rat tissue blocks. Labelling indicates tissue block number. Scale bars represent 5 mm.

by missing sections at the top and bottom of the tissue block, as well as clipping the tissue to fit it onto the slides.

A special case where slides from the rat tissue block failed to register and were subsequently discarded was at the top of the stack, where the only tissue present was placenta surrounded by a small amount of smooth muscle, as illustrated in the bottom row of Figure 2.3. In each tissue block the amount of smooth muscle tissue present in the slides will at some point become insufficient for accurate registration. For this reason, the *in silico* reconstructions of the rat tissue blocks do not contain the topmost slides. The reduction in stack size is noted in Table 2.2, and is equivalent to a depth of roughly 0.5 mm for each tissue block. Additionally, the bottom of the stack, where the sectioning began, is also missing tissue, because a number of slides had to be discarded prior to the first numbered slide. The number of slides lost in this manner cannot be quantified precisely because these slides were not numbered, but is fewer than 100, which is the approximate number lost at the top of the stack. These two bounds on the volume represent a truncation of the tissue in the final volume representation, and the effect of this truncation shall be noted in the analysis in of the reconstructions in Chapter 7. Similarly, the reconstruction of the human tissue block also represents a volume with artificial bounds at the start and end of the stack, and the effect of this truncation will likewise be noted.

The human tissue block, and second and third rat tissue blocks, also had bounds on the cross-sectional area imposed because the tissue was too large to fit entirely on the slides used. In order to allow the tissue to fit on the slides, the tissue was clipped, meaning that any tissue extending over the edge of a slide was trimmed off. The human tissue block was clipped at the cervical end to fit the tissue to the slide, and additionally the tissue block was made to fit the width of the slides by cutting the tissue in two prior to embedding in paraffin. This had the result of creating an artificial bound at the centre of the uterus and a bound at the cervical end, as illustrated in Figure 2.4. Similarly, the second rat tissue block was too long for to fit the length of the slides, and sections were

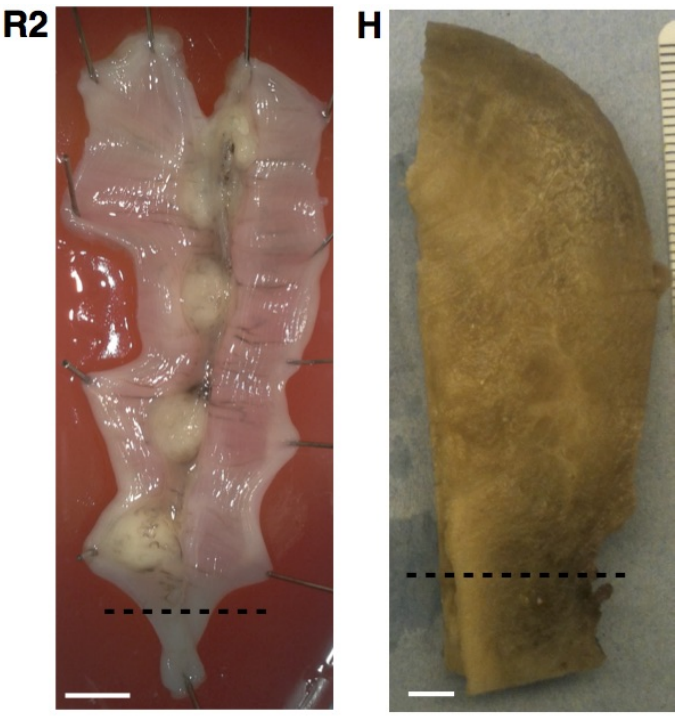

Figure 2.4: The approximate location of the clipping boundaries applied to the second rat block (**R2**) and human block (**H**). The clipping was applied during the sectioning process, and varied between slides, but the overall effect was a truncation of the tissue along the dashed lines shown. Scale bars represent 5 mm.

clipped along the line indicated in Figure 2.4. The third rat tissue block was also too long for the sections to fit on the slides; however, it was possible to preserve much of the smooth muscle architecture by exploiting the three-dimensional positioning of the tissue in the paraffin block. The cervical end was slightly raised in the block, which caused the initial slides to contain tissue mostly from the cervical end. Accordingly, the ovarian end of the tissue was clipped to obtain the smooth muscle at the cervical end for theses slides. As the sectioning progressed, the smooth muscle became less prevalent at the cervical end and more prevalent at the ovarian end. For this reason, the clipping was switched to the cervical end beyond a given slide deemed to be a transition between these two. While this process optimised preservation of the smooth muscle tissue for the reconstruction, the transfer of clipping did introduce slightly more complex artificial bounds, as illustrated in Figure 2.5. The effect of these bounds will be noted in subsequent analysis of the *in silico*

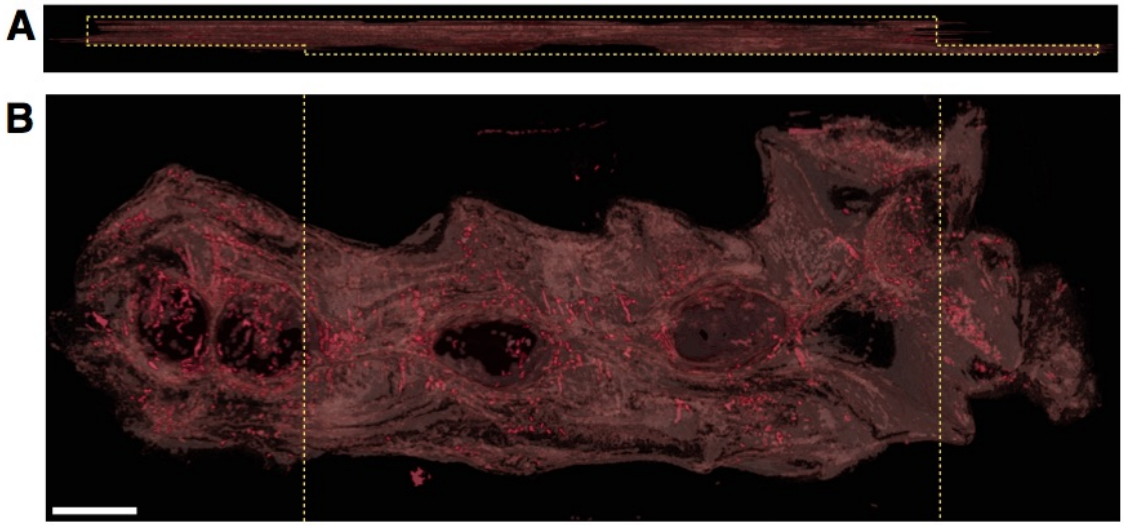

Figure 2.5: The bounds induced in the third rat tissue block by clipping the tissue while sectioning, as demonstrated on a representation of the three-dimensional reconstruction of the tissue. Here the tissue was clipped at the ovarian end for the first 120 slides, before switching to clipping at the cervical end. **A**: The approximate bounds created by this switch in clipping. **B**: The approximate locations where the tissue was clipped at each end. Scale bar represents 5 mm.

reconstruction.

During the sectioning process, it was noted that the final reconstruction did not require the full set of slides in the human tissue block. Accordingly, the number of sections discarded from the human tissue block was increased after the 300th slide to two out of every three in order to reduce the time required to process the slides. This reduced slide count is evident in the number of slides present prior to registration, given in Table 2.2; however, the effect of this purposeful discarding on the final reconstructed volume is minimal, due to the slides being coarse-grained, as described in Section 2.5.

Table 2.2 summarises the number of slides present at each step, and the total number of slides which would be present, had none of the slides been discarded. About half of the sections in each of the rat tissue blocks were deemed unfit for registration. The proportion of sections discarded in the human tissue block was considerably greater than half; however, this was due to the reduction in slides mentioned in the previous paragraph. As detailed in Section 2.5 below, the *in silico* reconstruction was coarse-grained in the z-

Tissue block	Total number of slides	Slides present before registration	Slides present after registration	Proportion of slides used
Rat block 1	389 (485)	221 (271)	141	36%
Rat block 2	411 (488)	237 (279)	174	42%
Rat block 3	418 (442)	201 (210)	118	28%
Human block	711	284	185	26%

Table 2.2: Number of slides at each step of processing. Values in brackets indicate the number of slides including the top slides discarded due to an insufficient amount of tissue (rat blocks only, see text for details). Proportion shown is the percentage of the original slides used to construct the final volume.

direction, which means that each slide in the final three-dimensional volume represents approximately 10 slides from the original tissue. This coarse-graining means that each slide of the reconstructed tissue is represented by an average of 3–4 slides from the original tissue blocks. This strongly suggests that the *in silico* reconstruction is representative of the tissue through the stack, and each slide in the final volume represents a reliable sample of the original tissue.

2.3 Two-dimensional extraction of cell orientation

Fibre direction in smooth muscle tissue is well-approximated by the orientation of the cell nuclei within the tissue [70,71,84,103]. Accordingly, the image analysis process began with the acquisition of shape data for the nuclei. This process was fully automated, and was performed by a bespoke ImageJ plugin. In addition, a second process of coarse graining these data was applied, as detailed below. The coarse graining algorithm is likewise fully automated, but performed by a Java class independent of ImageJ.

To isolate the nuclei from the rest of the image, a thresholding procedure was used, as summarised in Figure 2.6. The H & E staining colours nuclei purple and the remaining tissue pink [47]. Accordingly, the thresholding algorithm targeted red and blue colour values in the images.

Variation in staining of slides necessitated the use of different thresholds for each

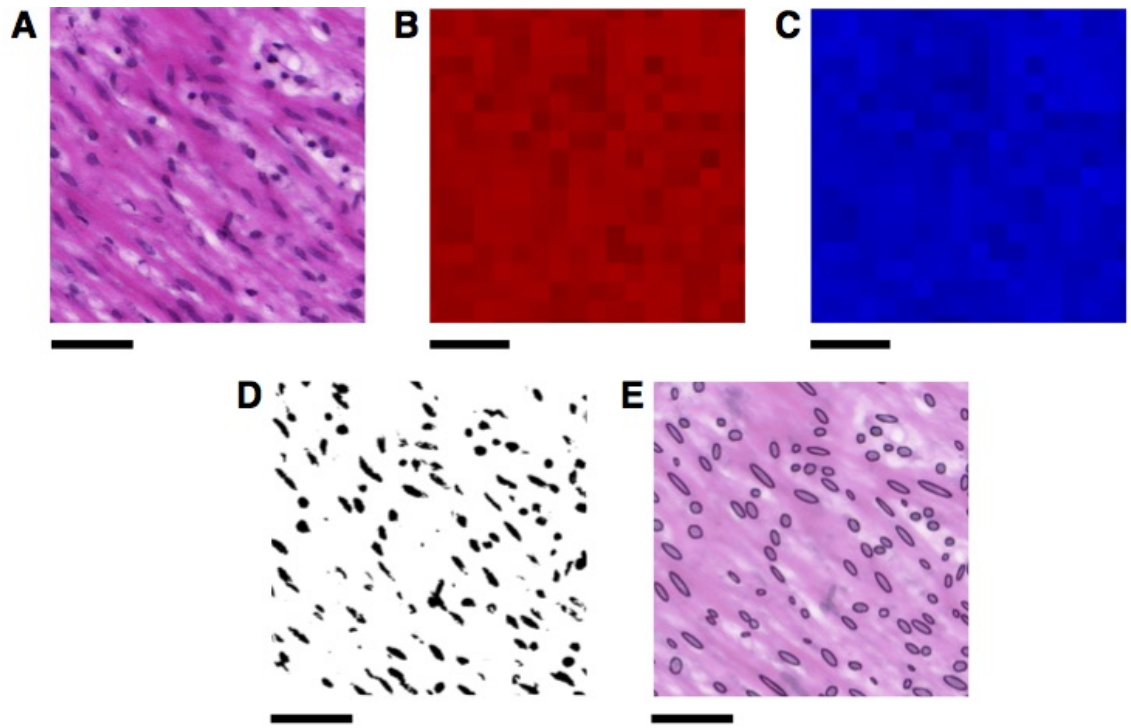

Figure 2.6: **A**: A 500×500 pixel region of tissue. **B** & **C**: Local red and blue thresholds respectively, found by considering the average values within 32×32 pixel regions. **D**: Binary image obtained by applying these thresholds to the original image **A**. **E**: Ellipses found by applying the ImageJ "Analyze Particles" function [2] to **D**, superimposed onto a lower intensity copy of **A**, showing the nuclei isolated from the rest of the tissue. Scale bars represent 50 μm.

slide. An additional difficulty was variation of the contrast between the nuclei and the surrounding tissue within each slide, as can be seen in Figure 2.6**A**, so that threshold values needed to be defined locally within each slide. In view of these considerations, each slide was divided into regions of 32×32 pixels (12×12 μm^2 high-resolution, 15×15 μm^2 low-resolution). This size was chosen as it is slightly larger than the width of a nucleus and each region containing a nucleus will therefore contain a portion of the surrounding tissue.

The threshold for each colour was determined by taking the average value of each within the region: for red values, the upper threshold was set at 5/6 times the average value (Figure 2.6**B**). The nuclei are highlighted by setting the threshold slightly below the average red value; this also removes the surrounding tissue from the binary image.

The blue values have an upper threshold equal to the average value (Figure 2.6**C**), which serves to reduce the noise in the resulting binary image, while preserving the nuclei. The resulting binary image is shown in Figure 2.6**D**.

The binary images obtained contain the approximate shapes of the nuclei as seen within the plane. Next, the images were analysed using the ImageJ "Analyze Particles" function [2]. This function fits an ellipse to each nucleus, as shown in Figure 2.6**E**. The function also allows the ellipses to be filtered by size, which enables further noise reduction by eliminating ellipses which are outside the feasible size range of nuclei. Ellipses with area smaller than 40 pixels (5.5 μm^2 high-resolution, 8.6 μm^2 low-resolution) in size were discarded as noise, while ellipses larger than 600 pixels (82.6 μm^2 high-resolution, 129.7 μm^2 low-resolution) in size were discarded on the assumption that these corresponded to clusters of nuclei, which would not provide accurate shape data for the purposes of determining fibre direction.

The length and orientation of the major axis of the ellipse was defined by the largest diameter of the nucleus, whereas the minor axis was fixed by setting the area of the ellipse equal to that of the nucleus [2]. For each of these ellipses, the position, aspect ratio, area, and angle of the major axis from the x-axis were recorded as illustrated in Figure 2.7.

The nuclear shape data indicates the orientation of individual cells in the tissue. This orientation gives an indication of the direction of the fibre in which the cell was embedded; however, the orientation does not directly correspond to the fibre direction. Thus, to determine the fibre structure, a coarse-graining step was required. For this reason, regional averages were calculated to reduce the data set. Each slide was divided into regions of size 128×128 pixels (47.5×47.5 μm^2 high-resolution, 59.5×59.5 μm^2 low-resolution). All ellipses with centres within a given region were considered as being contained within that region.

The ellipses used to generate this regional representation were filtered by size: ellipses within the size range 75–450 pixels (low-resolution 59–356 pixels) were considered to be

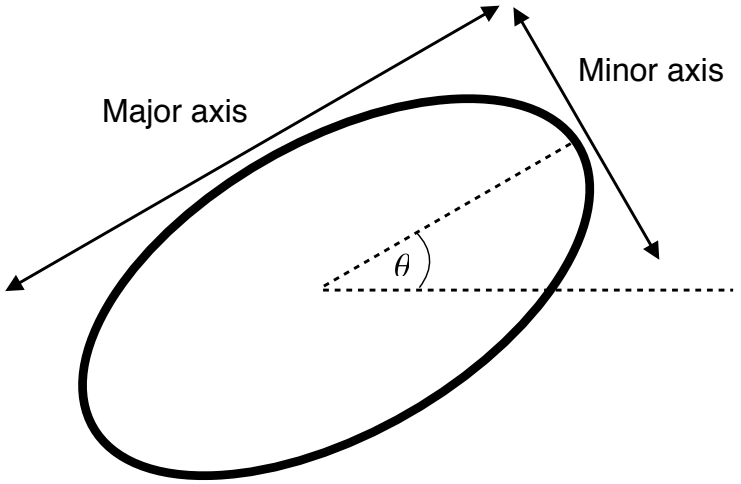

Figure 2.7: Shape characteristics of an ellipse. Aspect ratio was calculated as major axis divided by minor axis.

smooth muscle nuclei. This refinement of the earlier size filter removed smaller cells, such as red blood cells, and larger clusters from the calculations. The smaller size range was used in detecting vasculature, while the full range of sizes was used for detecting placental tissue, as described in Chapter 3.

The ellipses within a region were categorised as being *planar*, i.e. nucleus lying in the plane, *vertical*, i.e. nucleus lying vertical to the plane, or indeterminate. A long thin ellipse is more likely to represent a nucleus in the plane than a short round one. Accordingly the aspect ratio was used as a criterion to distinguish between these categories: nuclei with aspect ratio greater than 2.0 were classified as planar, nuclei with aspect ratio less than 1.6 were classified as vertical, and nuclei with aspect ratio between 1.6 and 2.0 were considered indeterminate, and were included in both *planar* and *vertical* categories. This middle set served to increase the number of nuclei represented in the coarse-grained approximation described in the next paragraph.

In this manner, regions are categorised as being *planar*, *vertical*, or *empty*. If a region contained fewer than two ellipses, it was categorised as being *empty*. If 5/9 of the ellipses were classified as *vertical*, the region was classified as *vertical* and no planar direction was

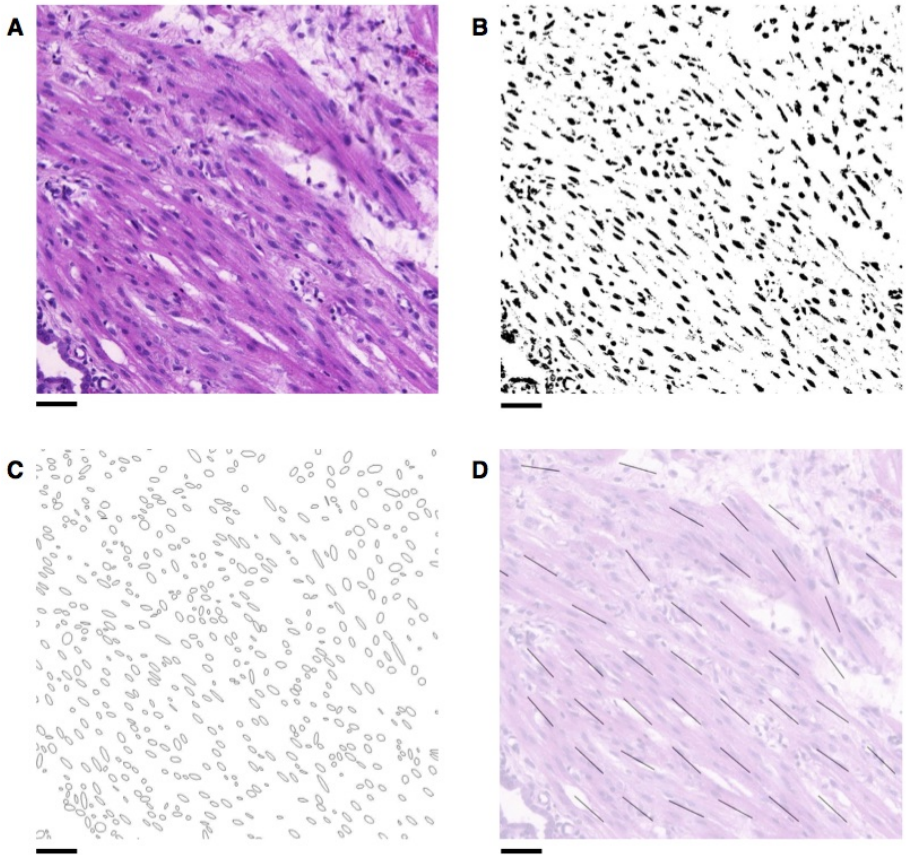

Figure 2.8: The process of obtaining directional data from histology as applied to a 0.5×0.5 mm^2 area. **A**: Example area of an H & E stained slide taken from rat myometrium. **B**: Image found after applying the thresholding procedure to **A**. This binary image shows the nuclei in black. **C**: Ellipses approximating the nuclei found using the ImageJ "Analyze Particles" function on **B** [2]. **D**: Directional data obtained by averaging directions of ellipses in **C**, shown in comparison with the original image **A**. Scale bars represent 50 μm.

assigned to the region. Otherwise, the region was classified as *planar* and a direction for the region was also assigned, as follows. The direction of a region was defined to be the median angle of all ellipses in the region. The median was used rather than the mean as the number of ellipses within a region was generally small, which renders the mean highly sensitive to outliers, whereas the median is more robust. The interquartile range of the angles within a given region was also taken; if this was greater than 45° then the direction was deemed unreliable, and consequently the region was marked as *vertical* and therefore was not assigned a direction. The resulting directional data can be seen in Figure 2.8**D**.

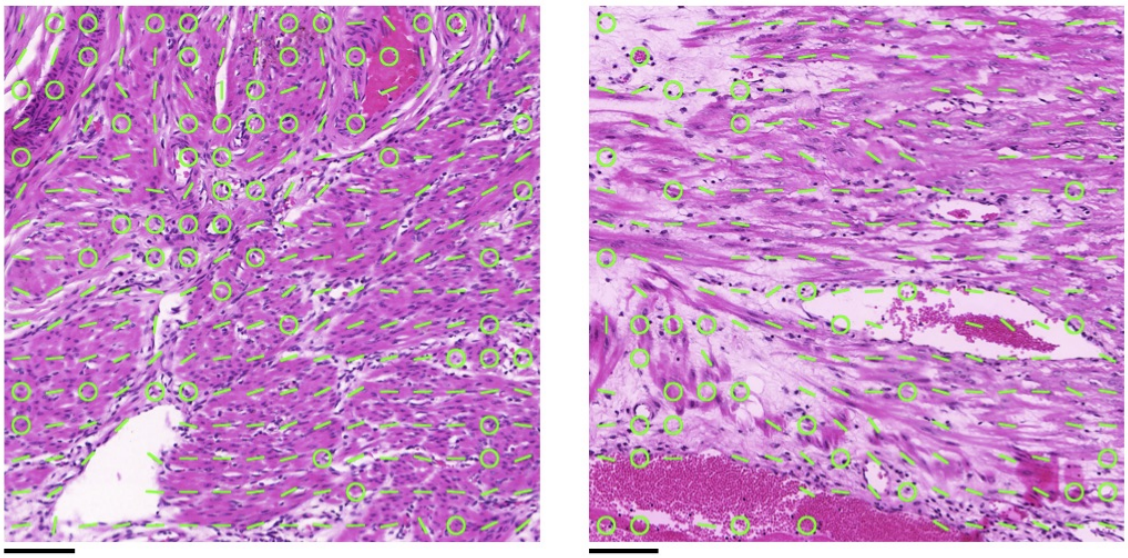

Figure 2.9: Examples of directions assigned to the tissue based on the nuclei. Green lines indicate the planar orientation detected at that point, while green circles indicate that no planar direction could be assigned to the given point, but the point is still marked as containing nuclei. Left: human tissue block. Right: rat tissue block 3. Scale bars represent 100 μm.

Visual verification of the effectiveness of this procedure was performed on randomly selected 2048 $\times$ 2048 pixel tiles, as shown in Figure 2.9, with direction indicated by green lines. The close agreement indicates that the planar directions found by the image analysis represent the tissue direction. Regions classified as *vertical*, where no planar direction could be assigned but contain nuclei, are shown as green circles. These arise when either the cross sections of the detected nuclei closely approximate a circle, so the direction is unreliable, or when the detected nuclei in the region have widely varying direction. While smooth muscle fibres are well-represented by the direction vectors in the images, it is also evident that a portion of these vectors have in fact been assigned to the vasculature which represents a confounding factor to the *in silico* reconstruction of the myometrial direction field. To eliminate these nuisance vectors, vascular regions are removed later in the processing pipeline by a separate algorithm, as described in Chapter 3. The remaining tiles which were compared in this manner can be found at `https://figshare.com/s/9a54e1ae6699cf713670`.

2.4 Image Registration

In order to obtain an *in silico* representation of the three-dimensional structure of the tissue, the histological slides first need to be registered. Two types of registration were applied sequentially to achieve this objective. The first is a rigid registration, where each slide is translated and rotated based on a single transformation that is applied to the whole slide, so that the slide maintains its original shape. The second is an elastic registration, where a grid of points is laid over the slide and each point can be subjected to a different transformation. This allows the slide to change shape.

In both registration processes a generalised form of the Hough transform was used [12], which will be discussed in detail in Section 2.4.1. The registration algorithms are outlined in Sections 2.4.2 and 2.4.3. These registration algorithms enable the registration of one slide to a reference slide. The order in which the elastic registration is applied to obtain the best representation of the tissue is detailed in Section 2.4.4. Finally, Section 2.4.5 deals with the manual identification of slides which fail to register accurately.

2.4.1 Generalised Hough Transform

The Hough transform provides a means to find instances of a given feature in an image; see Ballard [12] for a detailed discussion and various applications. The simplest application of the Hough transform is finding a straight line in an image [57]. A straight line has the general form

$$y = mx + c$$

where m and c are the line parameters. This may be rewritten as follows:

$$c = -xm + y.$$

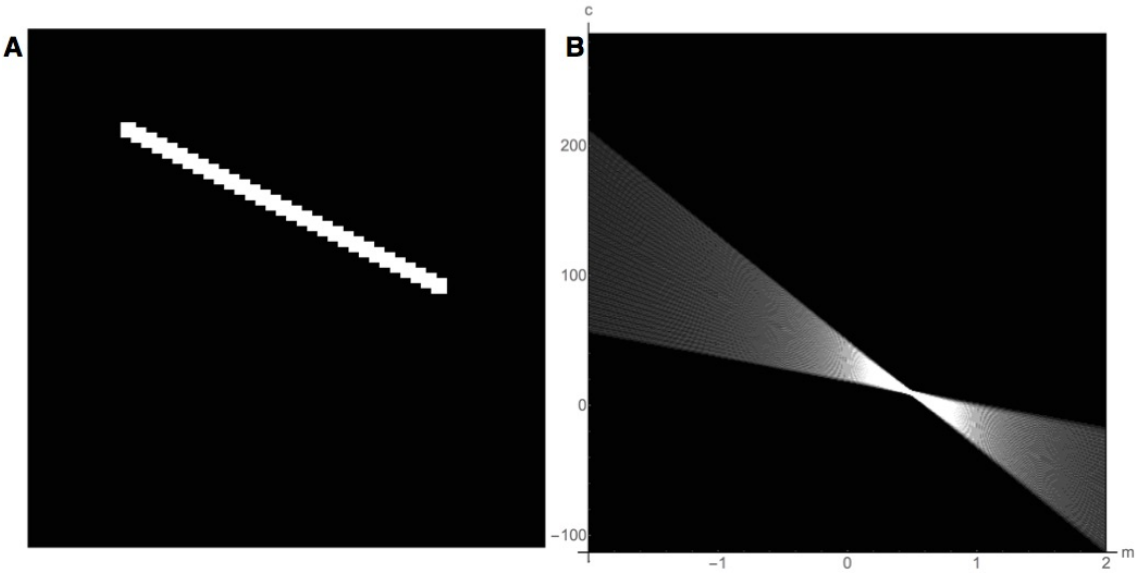

Figure 2.10: The Hough transform for line detection. **A**: A 100×100 pixel binary image containing a line. **B**: the voting array in the Hough space corresponding to **A**, visualised as a greyscale image. The point of highest intensity corresponds to the parameters of the line given in **A**.

This shows that each point (x, y) in an image generates a corresponding relationship between the parameters c and m that defines every straight line passing through (x, y). This linear relationship between c and m can be plotted in what is referred to as the Hough space. Thus, given an image containing a straight line, each point on that line can be mapped to its image in the Hough space, which also happens to be a straight line. A two-dimensional array of votes is generated by discretising the Hough space and counting the number of these lines passing through each discrete area, as illustrated in Figure 2.10. The highest entry in this array corresponds to the parameters defining the line in the original image. This method can also be applied to images containing multiple lines by applying a threshold to the voting array [33]. Alternative parametrisations of a line may be more appropriate for line detection [33]; however, line detection is not relevant to the following methods and thus this parametrisation will suffice for the purposes of illustration.

To understand how this can be used in image registration, consider two arbitrary shapes, S_1 and S_2. Suppose there exists a rotation angle θ and a translation vector t such

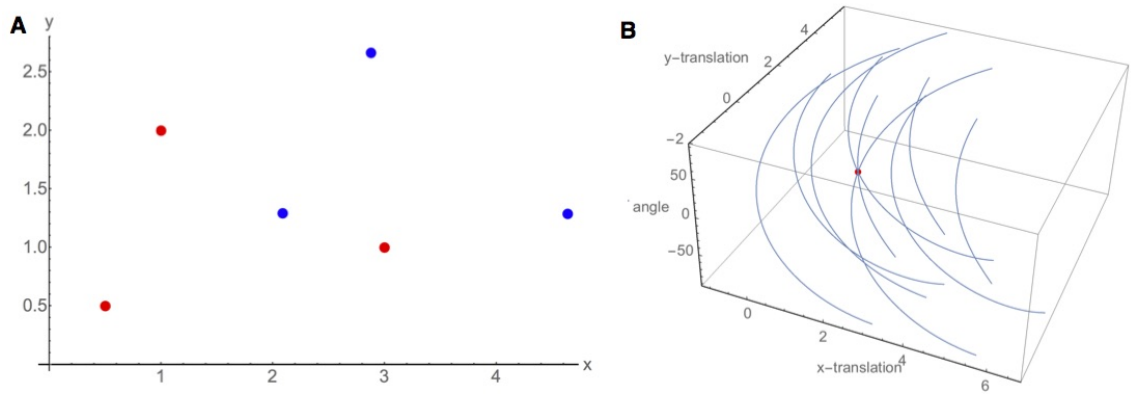

Figure 2.11: An example of applying the generalised Hough transform to registration of two three point images. **A**: The red and blue points in this plot can be thought of as images, which can be registered using the generalised Hough transform. **B**: Each line in the Hough space shown here represents the set of all possible transformations, given by a rotation θ and a translation t, which transform a given red point to a given blue point. The only intersection of three of these curves is highlighted. This intersection represents the only transformation which translates each red point to a blue point.

that rotating S_1 by θ and then translating by t gives S_2. More formally,

$$M_\theta S_1 + t = S_2$$

where M_θ is the rotation matrix of θ, defined by

$$M_\theta := \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}.$$

Consider a point p in S_1 and the point q in S_2 to which p maps in the above transformation. In formula,

$$M_\theta p + t = q,$$

which may be rewritten in the form

$$t = -M_\theta p + q.$$

This gives a relationship between t and θ for each pair p and q. Thus, each pair of points p and q, in S_1 and S_2 respectively, generates a relationship between t and θ, which is represented by a curve in the Hough space. The optimal transformation is the point at which the density of these curves achieves a global maximum, since this indicates a maximum agreement of the two shapes. The process is illustrated in Figure 2.11.

In principle, this technique enables registration of any two shapes. However, mapping each pair of points to the Hough space in this fashion is impractical for shapes with a large number of points. For this reason it is necessary to use additional techniques to reduce the computational effort required. These techniques will be detailed below in Sections 2.4.2 and 2.4.3.

2.4.2 Rigid slide registration

The rigid registration was performed by using the edge of the tissue in each slide as the shape to be used as the input for the generalised Hough transform. To obtain the edge image of a given slide, a greyscale image was generated by considering the regional data obtained in Section 2.3. Each region was assigned a pixel value equal to the number of nuclei within the region. These pixel values were then scaled up by a constant factor given in Table 2.3 to improve contrast for visualisation. An example of this representation is given in Figure 2.12**A**.

A Canny edge detection algorithm [21] was subsequently applied to these images to produce images of the edges of the tissue (Figure 2.12**B**), which were then subjected to the rigid registration procedure. For more information on the Canny edge detection algorithm see Appendix A. Parameter values for this edge detection are given in Table 2.3.

Computing a full Hough space to find the optimal transformation is infeasible from a computational point of view. For this reason the transformation was split into a translation and a rotation, and the Hough density optimisation was performed for each separately. The registration algorithm alternated, in an iterative fashion, between rotating and translating

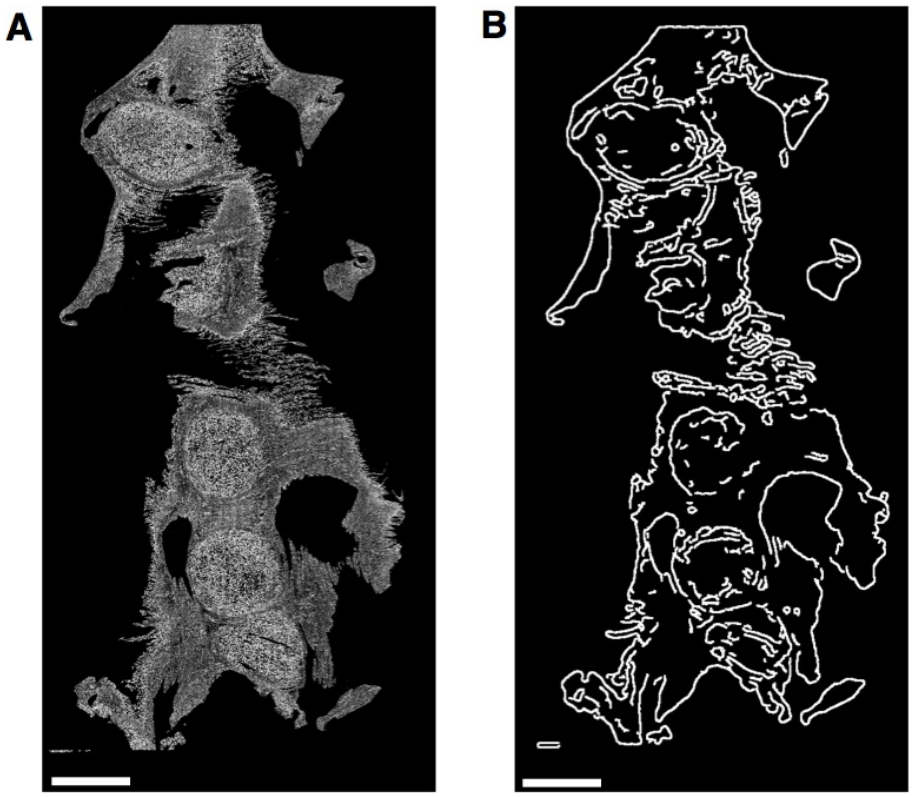

Figure 2.12: The process for obtaining edge images. **A**: Grey image representing the nuclear count assigned to each pixel. **B**: The edge image produced using a Canny edge detection algorithm on **A**. Scale bars represent 5 mm.

the image to find the optimal transformation. This sequence of iterations was subdivided into 7 *transformation steps*, where a maximum of 10 such iterations were performed in each step. Furthermore, with each step the range of possible translations and rotations was decreased, as described below.

For the optimisation of translation, a point p in the edge image being registered maps to all points t in the translation Hough space with

$$t = q - p,$$

where q is a point in the reference edge image. This Hough space was bounded for ease of computation, and the bounds for the Hough space were narrowed with each transformation

Parameter	value
Count scale factor	20
Gaussian standard deviation	3 pixels
Sobel window radius	1 pixels
Upper hysteresis threshold	40
Lower hysteresis threshold	20

Table 2.3: The Canny edge detector uses Gaussian smoothing, a Sobel differencing operator, and hysteresis thresholding with the parameter values specified here.

step, to further improve efficiency. The sequence of bounds on translation in each direction is given by

$$\frac{W}{2^i},$$

where $W = 800$ is the image width and i is the transformation step from 1 to 7. Votes were counted in this Hough space as described in Section 2.4.1.

For the optimisation of rotation, a point p in the edge image being registered maps to all angles θ in the rotation Hough space with

$$M_\theta p = q,$$

where q is a point in the reference image. This formulation of rotation specifically sets the centre of rotation at the origin of the image space. To improve performance a range of rotation centres, C, was used, as illustrated in Figure 2.13. This grid of points was dependent on transformation step: the central coordinate of the image c_0 was included in each such grid, and all other grid points had the form

$$c_0 + 50i(x, y),$$

where i is the transformation step and x and y are integers. This grid was confined to the area within 200 pixels from c_0 for the first transformation step, and to within the bounds of the image in all later transformation steps. This increase in spacing of grid points reduced

Figure 2.13: The centres of rotation used for rigid slide registration. Each point shown in red is used as a centre of rotation for the algorithm. The size and spacing of this grid varied depending on the transformation step. Scale bar represents 5 mm.

the number of calculations required for each transformation step, and hence reduced the computation time.

For each $c \in C$, the mapping of p to the rotation Hough space is thus defined by

$$M_\theta(p - c) + c = q.$$

Optimising for each centre of rotation in turn enabled a faster overall optimisation. To calculate the values for the rotational Hough space, the images were converted to polar coordinates, as shown in Figure 2.14. For each centre of rotation, c, the polar space was defined by taking c as the origin.

A translation along the angle axis in the polar space is equivalent to a rotation about c in Cartesian space. Thus a point p in the edge image with polar coordinates (r_p, ϕ_p) is mapped to all angles θ in the rotation space with

$$(r_p, \phi_p) + (0, \theta) = (r_q, \phi_q)$$

Figure 2.14: Polar transform of the edge image given in Figure 2.12**C**.

where (r_q, ϕ_q) is a point in the polar-transformed reference edge image. As was the case with the translation Hough space, the angle Hough space was bounded, and this bound was reduced with each transformation step. The sequence of bounds on rotation was given by

$$360^\circ, 10^\circ, 5^\circ, 2.5^\circ, 1.25^\circ, 1^\circ.$$

This sequence starts at 360° to allow for the possibility that a slide was rotated by $\sim 180^\circ$ during the sectioning process. Votes were counted in this Hough space as described in Section 2.4.1.

The iteration procedure for each transformation step is as follows:

(i) Find optimal translation and translate image

(ii) For each centre of rotation, find optimal rotation and rotate image

(iii) If no rotation or translations change the image, end, else go to (i).

This iterative sequence was repeated up to 7 times, with each transformation step reducing the range of translations and rotations as described above. This technique was applied to all pairs of slides sequentially in the stack, yielding a set of slides that were globally aligned with each other. The elastic registration procedure was applied to this stack to correct

localised distortions induced by the sectioning procedure, as described in the following section.

2.4.3 Elastic slide registration

The physical process of sectioning distorts the tissue on a local level. In order to correct for this distortion, a second registration algorithm was used that allows a small amount of elasticity within the slides. The elastic registration process enabled the registration of one slide to a reference slide using the generalised Hough transform. The regional representation of the slides, as detailed in Section 2.3, was used to provide the shape data for the generalised Hough transform. This representation categorises regions as either one of *planar*, *vertical*, or *empty*. The registration algorithm specifically registers *planar* regions, while *vertical* and *empty* regions are considered to be indistinguishable for the purposes of registration, both being regarded as *empty*.

The distortion in a slide cannot be completely eliminated, since attempting to do so would inevitably lead to each slide being transformed into a near copy of the reference slide to which it was registered. For this reason a lower limit was set on the degree of distortion to be corrected. This limit was set locally, and was determined on the basis of the spatial heterogeneity of the neighbourhood, as follows. Consider the two areas illustrated in Figure 2.15. Figure 2.15**A** shows a relatively heterogeneous region of tissue, which translates into a limited range of potential transformations onto the reference slide. On the other hand, Figure 2.15**B** shows a more homogeneous region; this homogeneity might allow this region to be transformed to many regions on the reference slide and still appear to be properly aligned.

The heterogeneity of a tile is expressed by the variance of the regional angles within the tile. Only *planar* regions are assigned angles; thus, to obtain a complete measure of heterogeneity, each non-*planar* region was temporarily assigned the angle perpendicular to the mean angle of all the planar regions in the tile. These temporary angle values were

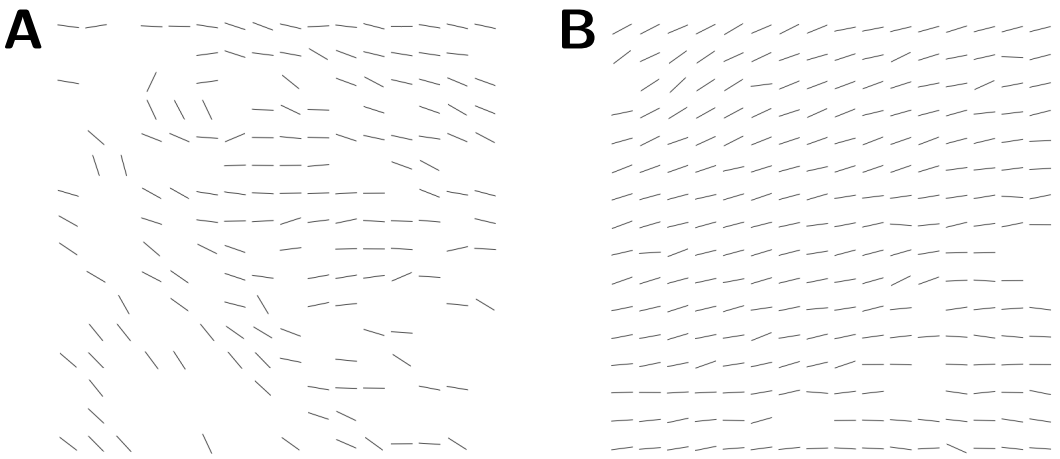

Figure 2.15: The varying degrees of heterogeneity within a slide. The lines represent directions of *planar* regions. Each image shows the planar directions in a 16×16 pixel (approx. 1×1 mm^2) tile. **A**: An example of a tile of high heterogeneity. The large amount of empty space and variation in fibre direction allows for a precise registration of the tile. **B**: An example of a tile of low heterogeneity. Here the tile comprises vectors with little variation in direction. This homogeneity can readily lead to an unreliable registration, and accordingly, the area of interest needs to be expanded to obtain a higher level of heterogeneity.

discarded once the variance was calculated. This assignment is illustrated in Figure 2.16.

The level of heterogeneity was used to determine how large an area around a tile was required for registration. If the angle variance was below $(45°)^2$, the area being registered was expanded to include neighbouring tiles. This expansion continued until the required level of heterogeneity was achieved, to minimise mapping to spurious areas on adjacent slides.

The following steps require some quantitative expression that indicates how well a given area matches the reference slide for a given transformation. This comparison was quantified in terms of an energy function. Consider an area A containing N regions and a transformation T. For a region at point p, the transformed point $T(p)$ will be surrounded by a set of points $\{q_i\}$ corresponding to regions in the reference slide. These surrounding points are given by

$$q_i = \lfloor T(p) \rfloor + (x, y)$$

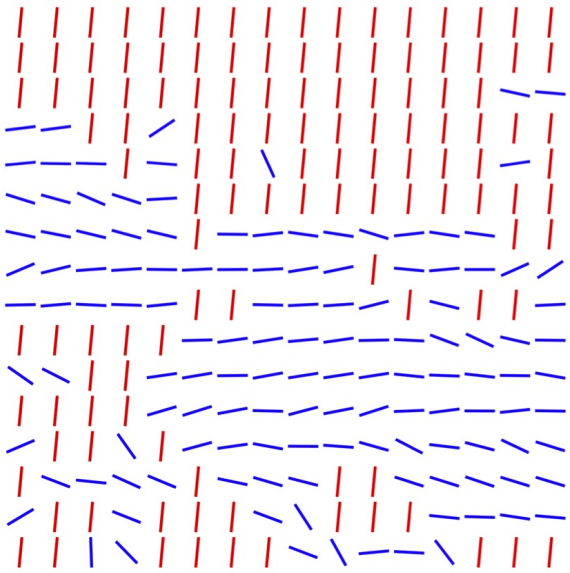

Figure 2.16: An example of the temporary assignment of angles to non-*planar* regions to determine heterogeneity. Here the blue lines indicate the directions of the *planar* regions, while the red lines indicate the assignment of direction to non-*planar* regions. This direction is perpendicular to the mean direction of all the *planar* regions, thus representing a maximal difference in direction. These temporary directions allow the measure of heterogeneity to factor in non-*planar* regions. Once the variance has been calculated the red lines are discarded to return to the original set of planar directions.

where $\lfloor \cdot \rfloor$ is a rounding down of the coordinates to the next region point, and $x, y \in \{0, l\}$ for region length l. The point p is defined to be *matched* if, for any q_i, either both p and q_i are non-*planar*, or both are planar and

$$|a(q_i) - (a(p) + \theta_T)| < a_{\max},$$

where $a(\cdot)$ is the planar angle of a region, θ_T is the angle of translation, and $a_{\max}$ is the maximum angle difference allowed, as given in Table 2.4. The energy e of the transformation T in A is defined to be

$$e(A, T) := 1 - \frac{N_m}{N} \ ,$$

where N_m is the number of regions which are matched according to the criteria described above.

Parameter	Value
$a_{\max}$	20°
$e_{\min}$	0.2
b_0 (human)	32
b_0 (rat)	16

Table 2.4: Parameter values for elastic registration.

The first step in optimising the transformation of an area A was to determine if it requires a transformation. The area was deemed to require transformation if

$$e(A, T_0) > e_{\min},$$

where T_0 is the transformation defined by no rotation and no translation and $e_{\min}$ is the minimum energy required to perform a transformation, as given in Table 2.4. This initial test reduced the computational workload by reducing the number of areas needing to be registered.

For an area requiring transformation, the directional data within this area were used for the Hough transform. To create the Hough space, the regional data in the area of interest and the corresponding data on the reference slide were taken to be the shape data. Denote the points corresponding to *planar* regions in the area of interest and in the reference slide by S_1 and S_2, respectively. The curves in the Hough space are defined by the expression

$$t = -M_\theta p + q.$$

for any $p \in S_1$ and $q \in S_2$, as in Section 2.4.1. Additionally, these curves are bounded by comparison of the directional data: if the region at point p has angle a_1, and the region at point q (on the reference slide) has angle a_2, then a further requirement is

$$|a_2 - (a_1 + \theta)| < a_{\max},$$

where $a_{\max}$ is the maximum difference in direction allowed, given in Table 2.4. This ensures

that the direction of p, after rotation, gives an angle which approximates that of q. This second constraint bounds the length of the curve in the Hough space corresponding to the transformation of p to q. In addition to this bound, the distance an individual region can be transformed and the angle of rotation were also bounded. The angle of rotation was bounded by 30°, while the bound on transformation distance was determined as follows. For an area A the distance bound is defined to be

$$b(A) := b_0 e(A, T_0)$$

where T_0 is the zero transformation and b_0 is a constant. The energy factor $e(A, T_0)$ in this bound means that the search for the transform is more localised when the area closely matches the reference slide, and a wider search is used when the area does not resemble the corresponding area in the reference slide. The bound $b(A)$ is the elasticity of the area A, with b_0 fixing the maximum elasticity of any given area. The value of b_0 was selected for a tissue block based on the cross-sectional area of the tissue; a tissue with a larger cross-section has more area to distort, and hence requires a larger amount of elasticity. This elasticity limit serves not only to reduce computational workload, but also to prevent the registration algorithm distorting the tissue through a widely inaccurate transformation. Subject to these bounds, the generalised Hough transform was implemented as outlined in Section 2.4.1.

This process assigned a transformation to each tile within a slide. In order to ensure accuracy of each of these transformations, they were tested against neighbouring tiles. For a given tile, the transformations assigned to tiles situated within a local neighbourhood of the initial tile were compared to the transformation assigned to the initial tile. If one of these was categorised as superior, it replaced the original transformation. To compare transformations, the energy function was used. For a given neighbourhood, the tiles within that neighbourhood were assigned areas of interest $\{A_i\}$ and transformations $\{T_j\}$ by the process detailed above. The total energy of a transformation T_j over the neighbourhood

is given by

$$\sum_i e(A_i, T_j).$$

The transformation that minimised this sum was deemed the optimal transformation. The initial tile was assigned this optimal transformation.

The application of this test has two main advantages. First, it reduces the chance of a tile being assigned an incorrect transformation. Secondly, it allows for a wider range of transformations for an individual tile: the range of transformations available to each tile is bounded by the local elasticity. This local elasticity differs with each tile, therefore a tile with a high level of elasticity could attain a transformation outside the range of tiles with lower levels of elasticity. By allowing each tile to potentially be assigned transformations from neighbouring tiles, the range of transformations can be increased.

Finally, each tile was transformed based on these allocated transformations. The transformation was applied to the original nuclear data (see Section 2.3). Each data point was transformed using a weighted average of the transformations from the nearest tiles. For a nucleus at point p, let $\{c_i\}$ denote the set of the four closest tile centres to p. The transformed point is defined as

$$\sum_{i=0}^{3} w_i T_i(p),$$

where T_i is the transformation associated with the tile centred at c_i, and w_i is the weighting function, given by

$$w_i = \frac{1/|p - c_i|^2}{\sum_{j=0}^{3} 1/|p - c_j|^2} \ .$$

This final transformation completes the process used to elastically register two slides. The next section describes how to apply this algorithm and the rigid registration algorithm to obtain the best representation of the original tissue.

2.4.4 Registration order

The techniques outlined in the foregoing sections provide a means for registering slides and reducing distortion. What remains is to implement them to obtain the best representation of the original tissue. Initially slides were registered using the rigid registration, with each slide being registered to the previous slide in the z-stack. This produced a rough alignment which allowed the comparison of slides for distortion levels. The levels of distortion were then used to select the optimal reference slides for the elastic registration.

The elastic registration process registered each slide to a reference slide. The sectioning process created some distortion in each slide. For this reason, the least distorted slides were designated most suitable to provide the reference for other slides. The process of selecting reference slides started with selecting a *global* reference slide. This global reference slide ultimately served as a reference slide for all other slides. *Local* reference slides were subsequently selected, working outwards from the global reference slide based on how well they matched the previous reference slide. Each reference slide was registered to the previous reference slide and as a result, all reference slides were effectively registered to the global reference slide. Finally, all non-reference slides were registered to their nearest reference slide to complete the registration process.

To determine a measure of how well two slides matched, a match score was calculated as follows:

$$1 - \frac{\#\text{matched regions}}{\text{average } \#\text{non-}\textit{empty} \text{ regions}},$$

where regions are defined as matched if they are non-empty in both slides. This match score is a measure of the relative distortion of the two slides, with 0 being perfectly matched and 1 being completely mismatched.

The *global* reference slide was determined by identifying a slide with a large quantity of tissue present. The amount of tissue present was automatically quantified by counting the number of non-*empty* regions in a slide. This provides a reliable global reference as

it gives a general cross-sectional shape of the tissue. To ensure that this global reference slide was a robust representation of the tissue block, a substack of slides containing the largest average number of non-*empty* regions was determined. All slides in this substack were then compared with each other to obtain an approximation of the level of distortion in each slide. The *global* reference slide was taken to be the slide in the substack with the lowest approximate distortion.

2.4.5 Slide integrity

An accurate representation of slides with low levels of distortion could be achieved using the registration algorithms detailed above. If the level of distortion of a slide was too high, however, it would not register to a reference slide. The cut-off point where the distortion becomes too high is not always evident from visual inspection of the slides prior to registration. Therefore, it is only after the registration is complete that a slide can be checked to determine if it has a sufficiently low level of distortion. This check was performed manually after each of the rigid and elastic registration steps had completed. If a slide tested was deemed to have failed to register accurately, it was discarded. If this registration failure occurred during the rigid slide registration, the next slide in the stack was registered to the previous slide to account for the inaccurate registration. If this registration failure occurred during the elastic registration and the slide was not a reference slide, no further registration was required. If, however, the slide was a (local) reference slide, then the elastic registration algorithm was run again from the last reference slide that had been accurately registered. This was necessary because the elastic registration is dependent on the registered version of the reference slides.

The manual testing of slide integrity after the rigid slide registration was performed by comparing the edge images defined in Section 2.4.2. The edge image of a slide being tested was added to the edge image of its reference slide. This composite image was then examined by eye, and was deemed unacceptable if the edges were almost completely

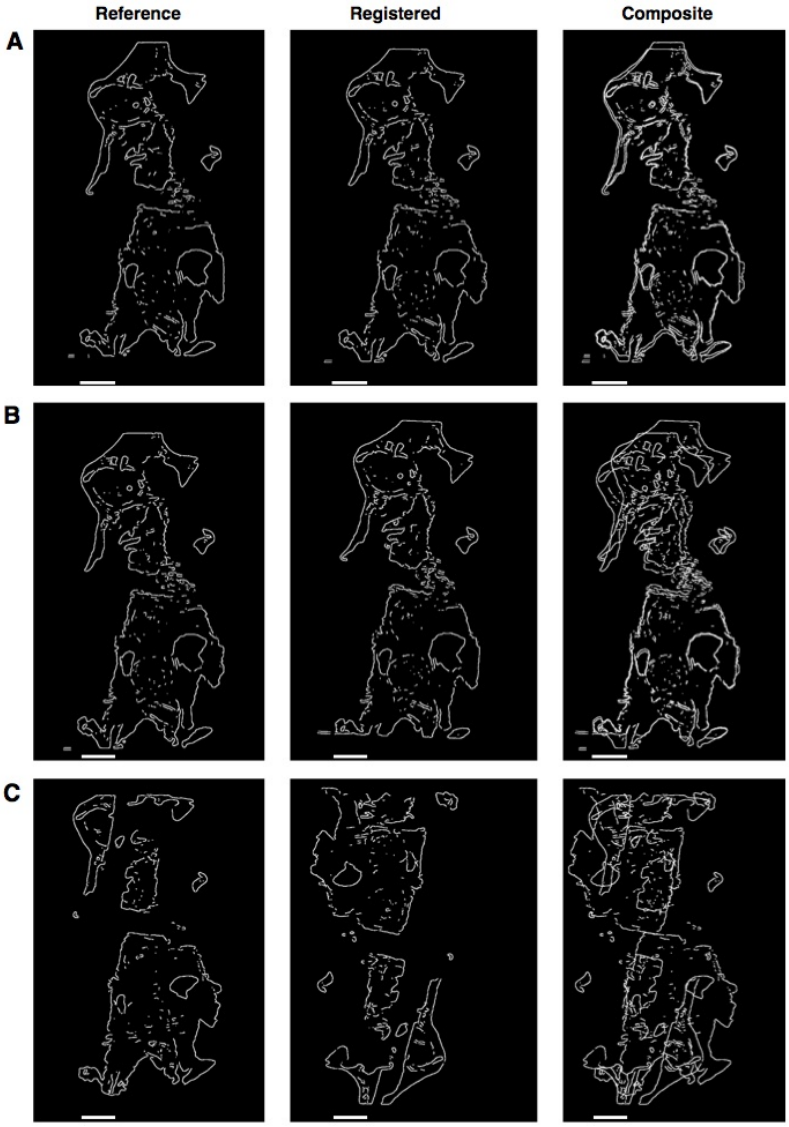

Figure 2.17: The possible outcomes of the rigid registration algorithm. Each row shows three panels: a reference slide, the result of registering a slide to the reference slide, and the sum of these two images. Scale bars represent 5 mm. **A**: An example of successful rigid registration. The two images are well aligned, which means they pass the visual integrity test. This situation is the typical outcome of registration (88%). **B**: An example of partially successful rigid registration. In this example the slide being registered is distorted compared to the reference slide. This level of distortion exceeds the maximum elasticity allowed by the elastic registration algorithm, and therefore cannot be corrected in the elastic registration step, which means that this distorted slide is discarded. **C**: An example of failed rigid registration. In this case the slide being registered was rotated $\sim 180°$ during sectioning. While the registration algorithm does allow 180° rotations, the shape of the slide in this configuration is similar to that of the reference slide, and therefore remained in this configuration. In this case the slide being registered is discarded.

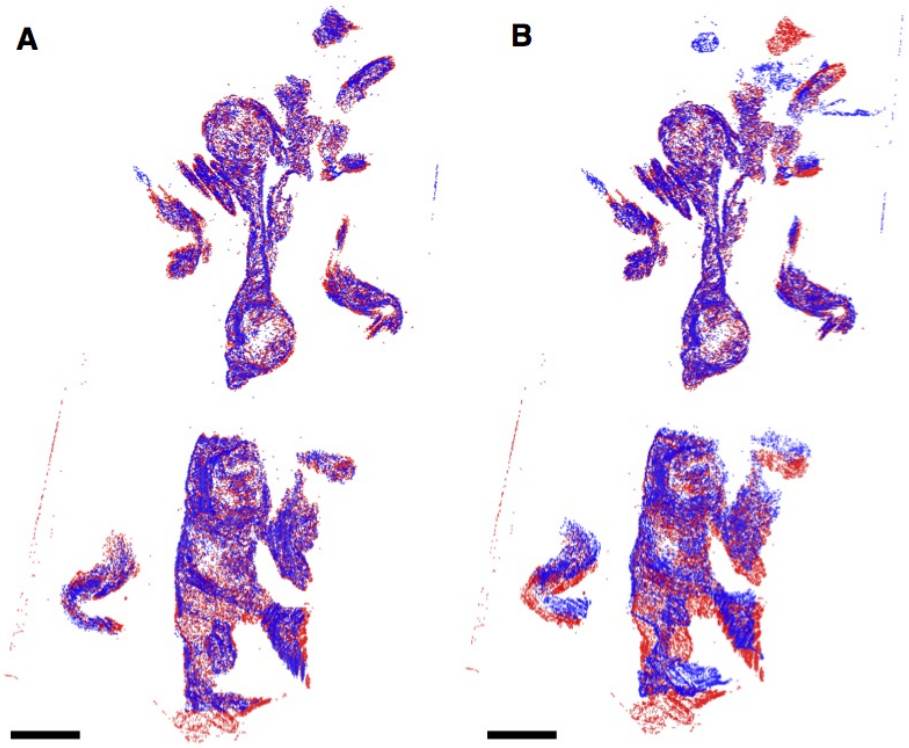

Figure 2.18: The possible outcomes of the elastic registration algorithm. The images here compare the *planar* regions of a reference slide in red with the *planar* regions of a slide being registered in blue. **A**: An example of successful elastic registration. Here the slides show perfect alignment over the full length of the slide. This is the typical outcome of the elastic registration algorithm (75%). **B**: An example of failed elastic registration. Here the slides are aligned for part of the length of the slide, but have failed to correctly register on the lower half. This occurs when the slide being registered is too distorted for the elastic registration algorithm to fix. Accordingly the slide being registered here is discarded. Scale bars represent 5 mm.

mismatched, or if a large level of distortion was observed. Examples of both of these cases are shown in Figure 2.17.

The testing of slide integrity after the elastic slide registration was performed by comparing planar regions of slides, as shown in Figure 2.18. This visualisation was used to compare elastically registered slides instead of edge images because a slide might fail to register accurately on the interior of the tissue, while remaining accurately registered at the edges of the tissue. Each slide was compared to the reference slide to which it was registered. A slide was deemed inaccurately registered if its relative distortion compared to the reference slide was clearly visible from the overlaid planar regions. An example of a

slide deemed inaccurately registered and an example of a similar slide deemed accurately registered are shown in Figure 2.18.

2.5 Averaging in the z-direction

The next step in the semi-automated processing pipeline is to create a uniformly spaced stack of slides from the registered slides. Slides were discarded during tissue preparation and registration, which means that the remaining slides form an irregularly-spaced substack of the full set of slides taken from the tissue, as shown in Fig 2.19**A**–**C**. In addition, one issue to contend with is that the minimum distance between slides (5 μm) is an order of magnitude smaller than the length of the slide regions ($\sim$ 50 μm) used to generate the coarse-grained two-dimensional representation of fibre direction in Section 2.3, making the z-resolution up to 10 times higher than the effective slide resolution. Another issue is the reliability of the fibre directions in individual slides. The elastic registration (Section 2.4) corrected large distortions of the tissue, but distortion in the order of 50 μm would be too small to be corrected in this manner. Small distortions of this kind would change the fibre direction of individual regions, which reduces the reliability of the inferred direction in each region. This distortion varies between slides, and therefore combining directional data from multiple slides can be used to improve the reliability of the inferred fibre direction. For these reasons, the slides were condensed into a new stack of slides, with uniform separation equal to the length of the regions generated in Section 2.3, as follows. The stack of slides was first divided into a set of substacks of uniform thickness, as illustrated in Fig 2.19. Each of these substacks was then mapped onto a representative slide in a new stack. The substack size was chosen so that the total thickness of a substack was equal to the length of a region. Thus the range of a substack, with representative slide at position z, was given by

$$zl_r/l_s \leq s < (z+1)l_r/l_s, s \in S$$

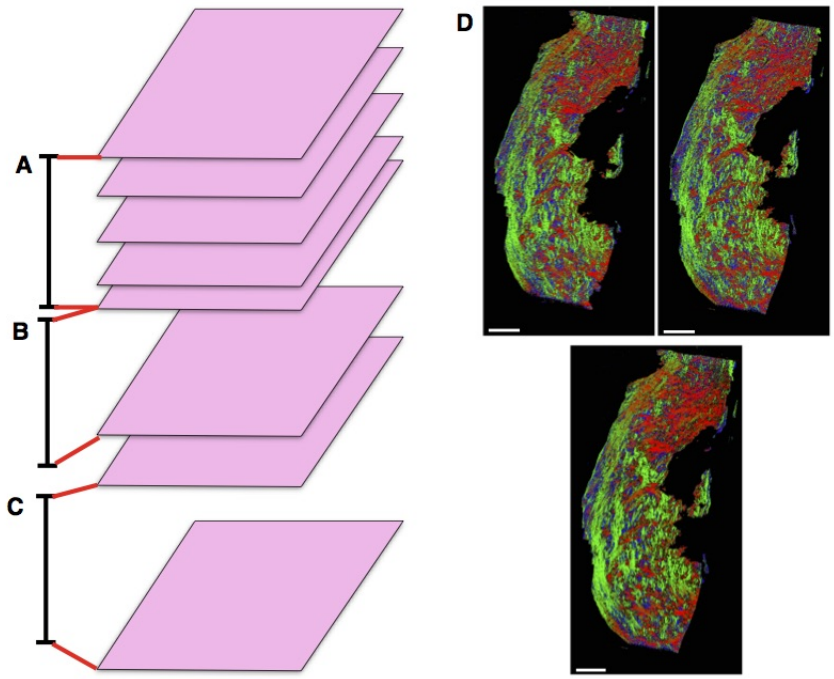

Figure 2.19: Generating a stack of representative slides. **A**–**C**: Selecting slides to combine to generate representative slides. The stack of irregularly-spaced slides is divided into ranges (black bars) containing substacks. Red bars indicate the range of slides used to generate the representative slides. When there is more than one slide present in a range (**A**), the substack contained in the range is used to generate the representative slide. If only one slide is present in the range (**B**), the nearest slide outside of the range is added to the substack. If no slides are present in the range (**C**) then the nearest slides either side of the range are taken as the substack. **D**: Representative slide (bottom) generated from two registered slides (top), where red indicates *planar* regions with left-right direction in the image plane, green indicates *planar* regions with up-down direction, and blue indicates *vertical* regions. Scale bars represent 5 mm.

where S is the set of all slide numbers in the original stack, l_r is the length of a region in μm, and l_s is the thickness of the sections in μm.

Each representative slide was generated by combining nuclear data from each slide in the substack, yielding a stack of evenly-spaced slides that form a three-dimensional volume. To reflect the three-dimensional nature of the stack, the discrete units of the representative slides will be referred to as voxels, which form the same two-dimensional layout as the regions generated in Section 2.3. These voxels were assigned categories and direction in the same manner as the regions in Section 2.3, based on the nuclei present in the volume represented by the given voxel. The average cross-sectional nuclear count was

also assigned to each voxel as opposed to the total nuclear count assigned in Section 2.3, given by the total number of nuclei present divided by the number of slides in the substack.

The procedure as outlined above could result in a substack with fewer than two slides. In the case of a substack containing one slide, the nearest slide outside the given range is added to this substack (Fig 2.19**B**). In the case of an empty substack, the nearest slides present after registration on either side of the given range are used to generate the representative slide (Fig 2.19**C**). Extending the range in each of these cases allows the generation of reliable representative slides for any sparsely populated areas of the stack.

The resulting representative slide generated from two individual slides is shown in Fig 2.19**D**. The representative slide shows higher contrast in colours, which represent directions, while maintaining the overall structure. Higher contrast represents a more uniform local direction within fibrous structures, which indicates that these directions are more reliable representations of the original fibrous structure than the input slides.

2.6 Determining the three-dimensional direction field

The final step in this process is determining the three-dimensional fibre directions in the tissue. The nuclear data obtained in Section 2.3 provided orientations of nuclei within the plane. The shape of the nuclei allowed a broad categorisation of the z-direction, but did not provide a reliable direction, due to the variability in shape and size of the individual nuclei. To obtain the z-direction for a voxel, the local bundle structure around the voxel was utilised.

At this point in the processing, the bundle structure has not been explicitly determined. For this reason, this structure was inferred from the two-dimensional directional data. To determine the local bundle structure, consider two adjacent voxels. These voxels are defined as belonging to the same bundle if all of the following criteria hold:

(i) Both voxels are non-*empty*.

(ii) The difference in planar direction is below 30°.

(iii) The difference in z-direction is below 30°.

In order to apply these criteria to voxels which have yet to be assigned three-dimensional directions, some exceptions need to be made. If one of the voxels being compared is *planar*, then it has no defined z-direction, so the third criterion is excluded. If one of the voxels being compared is *vertical*, then it has no defined planar direction, so the second criterion is excluded. A *vertical* voxel also has no defined z-direction. However, for the purposes of this algorithm *vertical* voxels are considered to have a z-direction 60° from the (x, y)-plane. In the algorithm detailed below, the comparison of two voxels always has at least one voxel which had previously been assigned a three-dimensional direction, so both planar and z-directions are already defined. Thus at most one of the above criteria is excluded through an undefined direction of a voxel.

Each step in the algorithm described below requires a means to compare potential direction vectors for a voxel in order to obtain the optimal direction vector. The measure used to compare vectors is the distance along a direction vector from the voxel to the nearest bundle edge, as illustrated in Figure 2.20. This distance is bounded above by a constant value, in order to ensure that it is measuring local bundle structure. Consider a voxel at point p and a direction vector v. The line l passing through p with direction v is defined by

$$l(\lambda) := p + \lambda v$$

for $\lambda \in \mathbb{R}$. The line l can be discretised to a sequence of voxels at points $\{p_i\}_{i\in\mathbb{Z}}$, defined by

$$p_i := [p + \lambda_i v],$$

where $[\cdot]$ rounds a point to the coordinates of the nearest voxel, and the sequence $\{\lambda_i\}_{i\in\mathbb{Z}}$ is defined such that

$$\lambda_i - \lambda_{i-1} = \min \frac{1}{v_j}$$

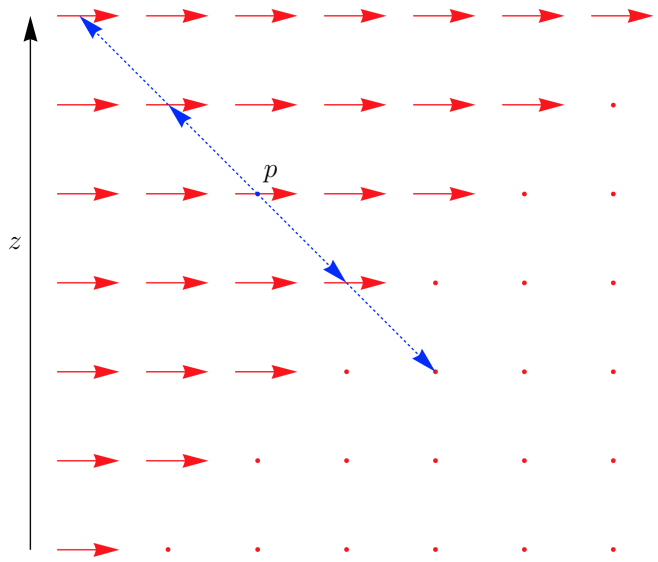

Figure 2.20: Measuring the distance to a bundle edge along a direction vector. Red arrows represent the planar directions at each point, red dots represent a planar direction through the plane being viewed. The direction vector v in this example is taken to be $-45°$ to the slide plane. The distance $L(p, v)$ is measured as the number of discrete steps, shown here with blue arrowheads, along v between p and a stopping point where either a dramatic change in direction occurs, as shown here, or an empty voxel is reached. This stopping point is considered to be the edge of the bundle containing p.

and

$$\lambda_0 = 0.$$

The distance along l to the nearest bundle edge is determined by sequentially comparing the voxels at $\{p_i\}$ outward from $p_0 = p$, using the above criteria. Suppose the voxel at p_i has been determined to be in the bundle containing p. For the sake of convenience assume that $i \geq 0$. This voxel is compared with p_{i+1} to determine if p_{i+1} is also in the bundle containing p. The planar and z-directions of the voxel at p_i, for the purposes of this comparison, are defined as follows. If the voxel already has a defined planar and z-direction, these directions are used. If the voxel is *vertical*, the direction v is used. If the voxel is *planar*, the planar direction of the voxel is used, and the angle v makes with the (x, y)-plane defines the z-direction. By the criteria above, the voxel at p_i cannot be *empty*. If the planar and z-directions of the voxel at p_{i+1} are already defined, these are used for

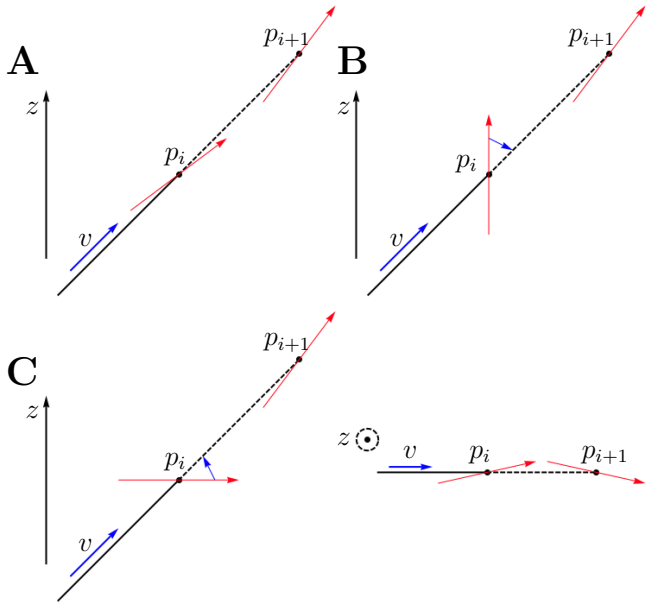

Figure 2.21: The ways in which the direction at p_i is temporarily modified to compare with p_{i+1}. **A**: The direction at p_i has already been determined in three dimensions, therefore no modifications are needed. **B**: The voxel at p_i has category *vertical*, meaning it has no predefined direction, so the direction vector v is used. **C**: The voxel at p_i has category *planar*, meaning it already has a fixed planar direction, so the vertical direction of v is used with the planar direction of p_i.

comparison. Otherwise, the comparison is performed according to the criteria detailed above, with the appropriate criterion excluded. These various options are illustrated in Figure 2.21. This sequential comparison is performed in both directions simultaneously, and stops when one of $p_{\pm i}$ is no longer contained in the same bundle. Denoting this stopping point by i_0, the length $L(p, v)$ of the line from p to the nearest bundle edge in the direction v, is defined to be

$$L(p, v) := i_0.$$

This measure is used below to compare direction vectors: if, for a voxel at p, the direction vectors v_0 and v_1 are such that

$$L(p, v_0) > L(p, v_1),$$

then v_0 is considered more likely than v_1 to represent the true direction at p.

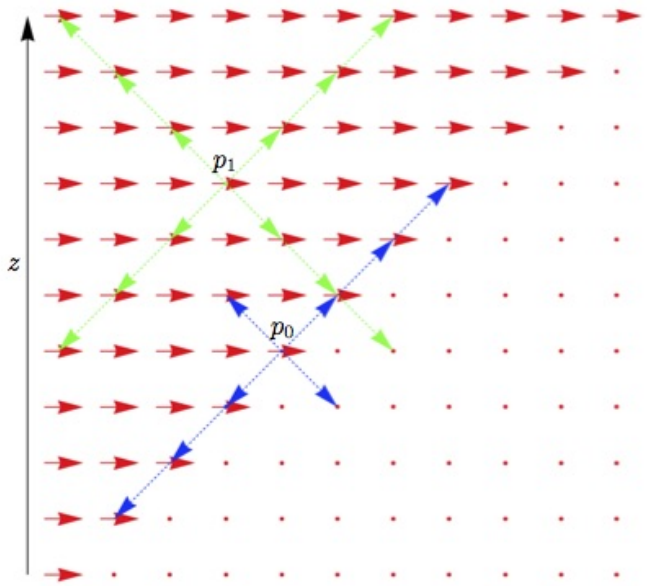

Figure 2.22: Comparison score of extreme z-angles. The point p_0 is situated on the boundary between areas of dramatically differing directions and therefore the local structure will impose a tight restriction on the z-direction at p_0. The point p_1 is farther from this boundary, and therefore the z-direction is not clearly determined by the local planar directions. The value of s for each of these points reflects this effect by comparing the distance to such boundaries at the extreme angles shown.

The process for determining the three-dimensional direction of a voxel differed depending on whether the voxel was categorised as *planar* or *vertical*. Since the *planar* voxels already have a planar direction assigned, all *planar* voxels were assigned z-directions prior to assigning directions to the *vertical* voxels. In both of these sets of voxels the directions were assigned in a specific order, as described below. This ordering enabled the voxels with a stronger requirement for a specific direction to be assigned that direction first, so that the voxels with a weaker directional requirement had their directions reinforced by the presence of the directions defined in the other voxels.

The following measure was used to determine the ordering of the *planar* voxels. For a *planar* voxel at p, with planar direction defined by θ, the function $s(p)$ is given by

$$s(p) = |L(p, v_+) - L(p, v_-)|$$

Where v_+ and v_- are defined by

$$v_+ = (\cos\theta\cos 60, \sin\theta\cos 60, \sin 60)$$

$$v_- = (\cos\theta\cos 60, \sin\theta\cos 60, -\sin 60).$$

This measure compares the lengths of the lines at 60° to the plane, with planar direction defined by the voxel. These two angles are used for this measure because, as detailed below, these values define the boundaries of the range of potential angles assigned to a *planar* voxel. Thus the score s compares the line lengths corresponding to the z-directions at either extreme. A high value of s corresponds to a higher chance of the z-angle being toward one of these extremes, which means that the voxel is more likely to have a well-defined z-direction based on the original regional data, as illustrated in Figure 2.22. For this reason, the planar voxels are sorted into descending order of s-values.

The three dimensional vector assigned to a *planar* voxel was bounded by the angles ±60°. This bound was used because the voxel had already been categorised as having direction closer to the plane than the vertical when it was first assigned this category.

For a *planar* voxel at point p, with planar direction defined by θ, the optimal z-angle $\hat{\phi}$ is defined such that

$$L(p, (\cos\theta\cos\hat{\phi}, \sin\theta\cos\hat{\phi}, \sin\hat{\phi})) = \max_{\phi\in[-60,60]} L(p, (\cos\theta\cos\phi, \sin\theta\cos\phi, \sin\phi)).$$

To approximate $\hat{\phi}$, an evenly spaced set of angles $\{\phi_i\}_{i=0}^{n-1}$, ranging from $\phi_0 = -60°$ to $\phi_{n-1} = 60°$ with separation 15° was used. For each angle ϕ_i the length

$$L_i := L(p, (\cos\theta\cos\phi_i, \sin\theta\cos\phi_i, \sin\phi_i)$$

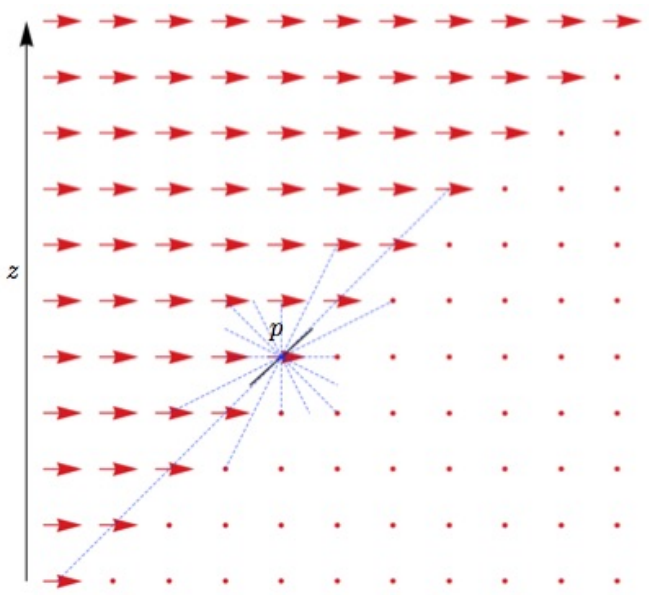

Figure 2.23: An example of the z-angle averaging to determine the approximate optimal angle. The lines corresponding to a z-angle below 0 all have length below the length threshold, and consequentially are not used in the calculation of the average angle. The weighted average of the directions with sufficiently large lengths is shown here as a black line. This direction is a good approximation to the actual direction along the edge of the bundle.

was calculated. The weighted average

$$\bar{\phi} := \sum_{i \in S} \frac{L_i \phi_i}{\sum_{i \in S} L_i}$$

was then taken to be the approximation of $\hat{\phi}$, where S is the set of all indices i for which L_i is larger than 2. This averaging is illustrated in Figure 2.23. This approximation, taken with the value θ, defines the three-dimensional direction at p, given by

$$v(p) := (\cos\theta \cos\bar{\phi}, \sin\theta \cos\bar{\phi}, \sin\bar{\phi}).$$

Unlike *planar* voxels, *vertical* voxels have no pre-defined planar direction. For this reason, both planar and z-directions must be found for the *vertical* voxels. To determine the three-dimensional direction of the *vertical* voxels, first the voxels are ordered in a

manner analogous to that of the *planar* voxels described above. For a voxel at p the function $s(p)$ is given by

$$s(p) := \max_{u,v \in V} |L(p,u) - L(p,v)|,$$

where the set V is defined to be the set of vectors

$$\{(\pm\cos 45, 0, \sin 45), (0, \pm\cos 45, \sin 45)\}.$$

This set is composed of the direction vectors which are at $45°$ to the (x,y)-plane and perpendicular to one of the x- or y-axes. This angle was selected because it allows the matching criterion to be met between p and *vertical* voxels, which have default angle 60, and with *planar* voxels which have been assigned angles away from the plane. As is the case for *planar* voxels, a high s instils more confidence that p has direction toward one of these extremes, and accordingly, the *vertical* voxels are sorted into descending order of s-values.

For a given *vertical* voxel at point p, the optimal planar angle $\hat{\theta}$ and z-angle $\hat{\phi}$ are defined such that

$$L(p, (\cos\hat{\theta}\cos\hat{\phi}, \sin\hat{\theta}\cos\hat{\phi}, \sin\hat{\phi})) = \max_{\theta\in[0,360)} \max_{\phi\in[30,90]} L(p, (\cos\theta\cos\phi, \sin\theta\cos\phi, \sin\phi)).$$

To approximate $\hat{\theta}$ and $\hat{\phi}$, an evenly spaced set of planar angles $\{\theta_i\}_{i=0}^{n-1}$, ranging from $\theta_0 = 0°$ to $\theta_n = 360°$ with separation $11.25°$ was used. For each planar angle θ_i, the approximate optimal z-angle $\bar{\phi}(\theta_i)$ was calculated using the same technique as detailed above for *planar* voxels, with the exception that the range of z-angles was defined as $[30, 90]$ as opposed to $[-60, 60]$. The length

$$L_i := L(p, (\cos\theta_i\cos\bar{\phi}(\theta_i), \sin\theta_i\cos\bar{\phi}(\theta_i), \sin\bar{\phi}(\theta_i))$$

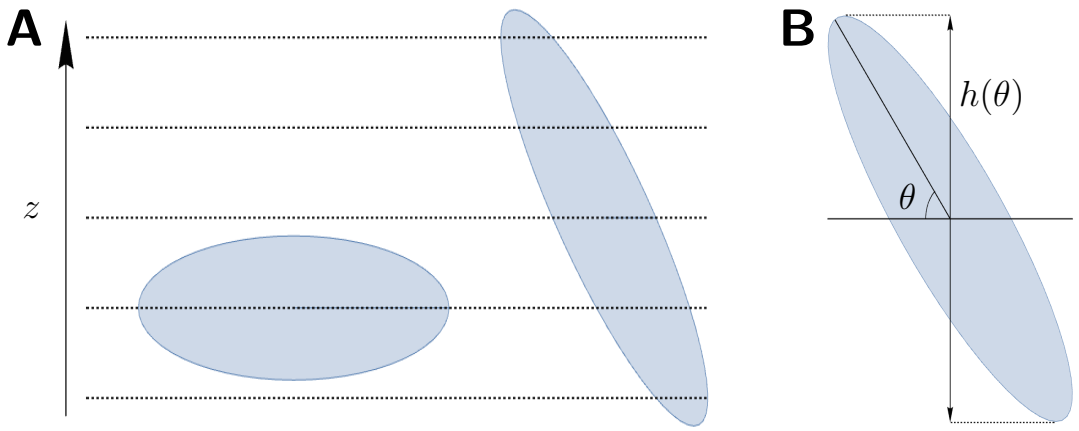

Figure 2.24: The effect of the angle a nucleus makes with the plane on the probability of being present in a slide. **A**: Comparison of the sectioning of two nuclei at different angles to the plane. The nucleus lying within the plane is only sampled once, whereas the nucleus which is more vertical is sampled multiple times in sectioning. **B**: The height function h of angle θ. The value of $h(\theta)$ is the difference in height between the highest and lowest points on a nucleus at angle θ.

was calculated. The approximation $\bar{\theta}$ of $\hat{\theta}$ was defined to be the value θ_i with maximal associated length L_i. This approximation, taken with the associated value of $\bar{\phi}$ defined the three-dimensional direction at p, given by

$$v(p) := (\cos\bar{\theta}\cos\bar{\phi}, \sin\bar{\theta}\cos\bar{\phi}, \sin\bar{\phi}).$$

Using this technique the three-dimensional fibre directions were determined from the regional data.

2.7 Weighting the direction vector

The final step in quantifying the three-dimensional fibre structure is the weighting of the voxels. The conductance of a given voxel in the tissue depends on the amount of tissue contained within a voxel and on the density of gap junctions between adjacent cells. On the assumption that the latter was uniform, each voxel was weighted using the nucleus count present in the regional data as a measure of tissue density. This count is the average

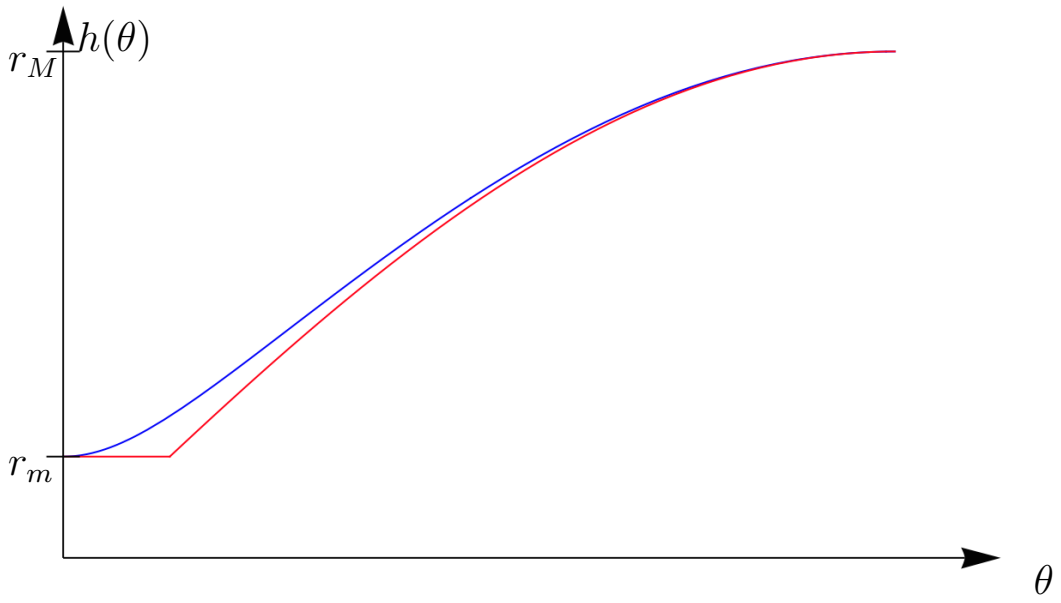

Figure 2.25: A comparison of the value of $h(\theta)$ and a scaled Sine curve which has been cut off at $\sin(\theta) = r_m$. Here r_M is the length of a nucleus, and r_m is the width. This approximation is acceptable in this case because of the high eccentricity of the nuclei.

number of nuclei observed in a two-dimensional slice through the voxel, as described in Section 2.5. Because this value represents the average number of nuclei observed within a plane, the direction of the nuclei in relation to this plane affect the number of nuclei observed, as illustrated in Figure 2.24**A**. Within a given voxel, a nucleus contained within the voxel with angle from the plane ϕ is observed in any given plane with probability

$$P(\text{observed}) = \frac{h(\phi)}{l}$$

where l is the height of a voxel, and $h(\phi)$ is the length of the vertical projection of the nucleus as illustrated in Figure 2.24**B**. For ease of calculation, the function h was approximated by

$$h(\phi) \approx \max\{r_M \sin\phi, r_m\}$$

Parameter	Value
r_M/l	0.3
r_m/l	0.1
N_0	48

Table 2.5: Parameter values for the weighting function.

where r_M is the length of the major axis and r_m is the length of the minor axis of a nucleus. This approximation is illustrated in Figure 2.25. Using this probability, the average number of nuclei observed in a plane N_p can be expressed as a function of the total number of nuclei in the voxel N:

$$N_p = E(\text{number of observed}) = \sum_i P(\text{nucleus } i \text{ observed}) = \sum_i \frac{h(\phi_i)}{l} = \frac{h(\phi)}{l} N,$$

where $\{\phi_i\}$ are the vertical angles of the nuclei and ϕ is the average vertical angle of the nuclei. The value of ϕ is given by the three-dimensional direction vector d assigned to the voxel:

$$\sin\phi = d_z$$

Thus the total number of nuclei within a voxel can be estimated given the average number of nuclei in the plane and the z-component of the direction. The weight of a voxel with inferred nuclear count N is accordingly defined as

$$w(N) := \min\left\{\frac{N}{N_0}, 1\right\}$$

where N_0 is selected such that if $N \geq N_0$ then the voxel is filled with conducting tissue. The parameter values for the weighting are given in Table 2.5. This weighting is applied in Chapter 3 to obtain a more refined three-dimensional bundle image and in Chapter 4 to reduce conductance in voxels which contain less than the maximum conductive tissue.

This concludes the methods used to determine three-dimensional structure of the uterus from serial histological sections. Chapter 3 provides refinements to this structure

with the aim of providing a more physiologically representative platform for simulated excitation and enabling measurement and visualisation of the fibre microarchitecture.

Chapter 3

Segmentation and analysis of three-dimensional structure

The aim of this chapter is to describe methods for segmenting and analysing the fibrous structure of the myometrium from the directional data found in Chapter 2. This structure will be used in this and later chapters; first to determine widths of the fibres in this structure, which will be used to for comparison with measurements made by Young and Hession [116]. A second use for this structure is to determine how the network of fibres connects various parts of the tissue. A comparison of this connectivity with multi-electrode array recordings described in Chapter 6 will enable the identification of structural characteristics that are relevant to excitation propagation. A third use for this structure is in excitation simulation. In particular, the segmentation of the fibrous architecture provides clear boundaries between distinct fibres; these boundaries can be assumed to have zero conductance since they are separated by electrically inactive connective tissue. This assumption ensures that if two structures are deemed to be distinct, then no current will pass between them.

An additional technique described in this chapter will be used to determine widths of the fibres in the two dimensional slide planes. These widths will be used to verify the registration of the slides: a registered slide that is representative of the original slides will have similar distributions of fibre width, and therefore fibre width is used as a measure of deformation induced by the registration process.

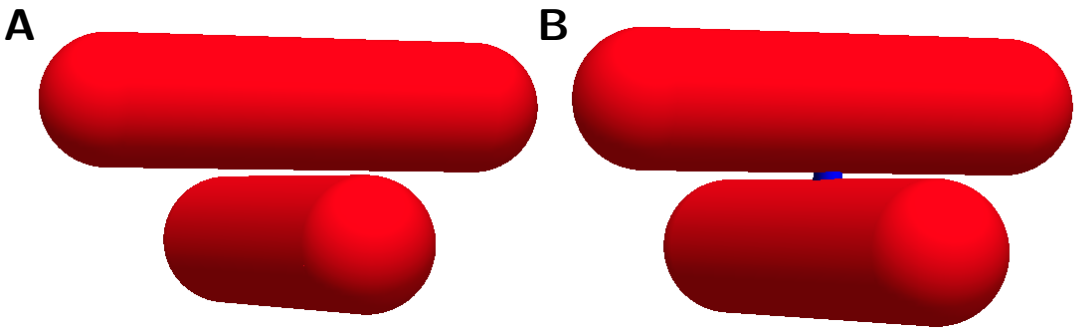

Figure 3.1: **A**: The two tubes shown here represent fibrous structures. They are close together, but not communicating. At resolutions comparable to the size of the gap, these structures may appear to touch. This is a problem when the aim is to resolve individual structures and to determine conductivity in the tissue. **B**: A bridge structure (blue) has been added. In this situation the fibres are connected by a filamentous structure and therefore there should be a connection in the digital representation. Even though the two situations are similar, they have vastly different diffusive properties, and therefore need to be differentiated.

Two forms of segmentation are described in this chapter: segmentation of tissue types and segmentation of fibres. The segmentation of tissue types, described in Section 3.4, is the process of separating the smooth muscle tissue from connective tissue, vasculature and, in the case of the pregnant rat tissue, placenta. The fibrous structure of the *in silico* tissue is determined by segmenting the fibres through the use of an edge detection algorithm. These edges must have the property that they completely encase a fibrous structure, enabling the distinction between two adjacent fibres which are not connected by a bridge, as illustrated in Figure 3.1. It turns out that edge detectors which rely on a local differencing operator, such as the Canny edge detector, are not suitable for this kind of distinction [59]. Even though they are able to determine the general shape of the edges, preliminary analysis suggests that they are susceptible to local noise and can therefore fail to detect all the edge points. To overcome this issue, a variant of a watershed algorithm is used, as described in the following sections. Watershed algorithms have the general property that any boundary is continuous [105], and will therefore enable the distinction between fibres to be absolute.

The techniques used to determine these boundaries in two dimensions differ from those

used in three dimensions, since the nature of the direction vectors is fundamentally different between the two cases; in two dimensions the directional data is in the form of two-dimensional planar angles and vertical points, whereas in three dimensions all directions are represented by a three-dimensional vector. The general formulation of the algorithm in both cases is the same, however. Initially, a grey image or volume is generated from the directional data using a form of Gaussian smoothing. This grey image is used to perform the modified watershed algorithm, which generates the fibre edges.

Section 3.1 describes the modified watershed algorithm when applied to a grey two-dimensional image. Section 3.2 details the edge detection in two dimensions, while Section 3.3 covers edge detection in three dimensions. The process of isolating the smooth muscle tissue from other tissue types is described in Section 3.4. This segmented structure is subsequently analysed using the tools described in Sections 3.5 and 3.6; Section 3.5 details the procedure for determining connectivity, while fibre widths in both two and three dimensions are measured using the methods described in Section 3.6. All code for the methods described in this chapter can be found at `https://figshare.com/s/651e592a91f309182edd`.

3.1 The modified watershed

A watershed algorithm uses the grey values of an image to determine boundaries between local maxima [105]. Consider the image in Figure 3.2**A**. This image has two local maxima, with a dark region between these two points. The watershed algorithm finds a line along that dark region which separates the two maxima, as illustrated in Figure 3.2**B**. These two local maximal can be considered to be two separate features, and in this way the watershed algorithm determines the solid line of points which separates these two features.

The watershed algorithm determines these separating lines as follows. Each point p in the image is sorted in order of increasing gradient values, to give the list $\{p_i\}_{i=0}^{n-1}$. This list is subsequently separated into sets of points, which shall be referred to as pools.

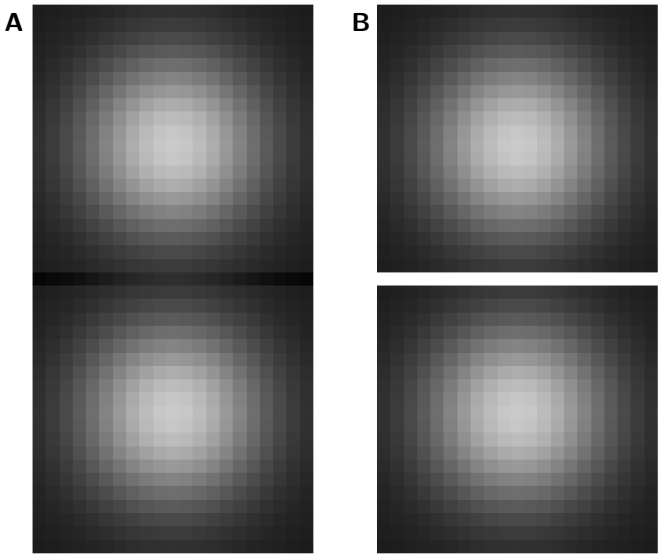

Figure 3.2: An example of the watershed process. The image in **A** contains two areas of high intensity separated by an area of low intensity. The watershed algorithm determines the line which separates these two high intensity points, as shown by a white line in **B**.

This separation is performed sequentially from p_0 using the following steps. For each p_i, consider all points $\{p_j\}$ adjacent to p_i. If $i < j$ for all j, then none of the points p_j have been placed into a pool, and so p_i is placed into a new pool. If all points p_j with $j < i$ are contained in the same pool, then p_i is added to that pool. Otherwise, p_i is on the boundary between two pools, and therefore is not placed into any pools. Such unassigned points form the lines separating the features, and the set of all these points is referred to as the watershed.

The following sections outline a method of generating greyscale images that represent local homogeneity in two and three dimensions. Areas of high homogeneity are more likely to correspond to areas within fibrous structures, while areas of low homogeneity are more likely to correspond to points on the boundary between fibrous structures. For this reason, the watershed segmentation algorithm sorts points in order of decreasing local homogeneity, rather than increasing gradient. As will become apparent in the following sections, two local maxima in the greyscale homogeneity image do not necessarily correspond to two separate features. If the above watershed algorithm were performed on such an image, this

would lead to erroneous boundaries forming, a process known as over-segmentation [25]. For this reason, an additional step is employed. For two local maxima p and q, a comparison function $f(p,q)$ is used which represents how well the local directions about p and q are matched. The specific formulation of f is dependent on the dimensionality of the problem, and shall therefore be described separately in the following sections. This function can be used to modify the watershed algorithm in the following manner. Consider a point p_i which has been found to lie on the boundary of the two pools, and let p and q be the local maxima in each of these pools. The point p_i is considered a watershed point if

$$f(p,q) < f_{\min},$$

where $f_{\min}$ is a constant threshold for the function f. This means that if the direction fields about p and q are sufficiently well-matched, then no barrier is drawn between them. The modification can be extended to the case where p_i lies on the boundary of more than two pools. In this case, all such pairs of local maxima p and q are compared, and p_i is deemed a watershed point if any of the pairs satisfy the above inequality.

This modified watershed algorithm is used to determine fibre edges. The following sections detail how the grey image or volume is generated in order to perform the watershed algorithm and how the comparison function f is formulated.

3.2 Determining boundaries in two-dimensional structure

The aim of this section is to describe how to generate a grey image and comparison function based on two-dimensional regional slide data for use in the modified watershed algorithm. The regional data for each slide consists of *planar*, *vertical*, and *empty* regions. Prior to registration, these categorisations represent an approximation of the fibre direction perpendicular to the plane, with *planar* representing a fibre running parallel to the plane, and *vertical* representing a fibre running perpendicular to the plane. Accordingly, the fibre

Parameter	Value
Gaussian variance	1.0
Vertical grey level threshold	10.0
$f_{\min}$	0.15
Planar grey level threshold	0.01

Table 3.1: Parameter values for two-dimensional fibre segmentation. The same Gaussian operator was applied in both *planar* and *vertical* segmentation processes.

edges were determined separately for each of the two categories.

3.2.1 Vertical grey image

The *vertical* regions of a slide have no planar direction. The absence of a vector-valued attribute means that these regions of each slide can be considered to represent a grey image. The pixel values of this image are taken to be the nuclear count of the regions. In order to reduce the presence of noise, a Gaussian smoothing operator was applied to the image. To determine the edges of these vertical structures, a lower threshold on the grey values of the image is set. A point is deemed to be an edge point if its grey value is below this threshold and adjacent to a point that has a grey value above the threshold. The result of this procedure is shown in Figure 3.3**B**. Parameters for this procedure are listed in Table 3.1.

3.2.2 Planar grey image

The *planar* regions of an image have an associated direction vector, which is represented by an angle. This means that the image made up of such regions is vector-valued, and must be converted into a greyscale image before it can be subjected to the watershed algorithm. The technique used for this conversion uses Gaussian smoothing to determine the local homogeneity of the planar directions. The level of homogeneity is a measure of how close a point is to being on the interior of a fibrous structure: regions at the edges of fibrous structures have low levels of homogeneity due to the region containing nuclei from outside of this structure which are likely to have different orientation to that of the fibre.

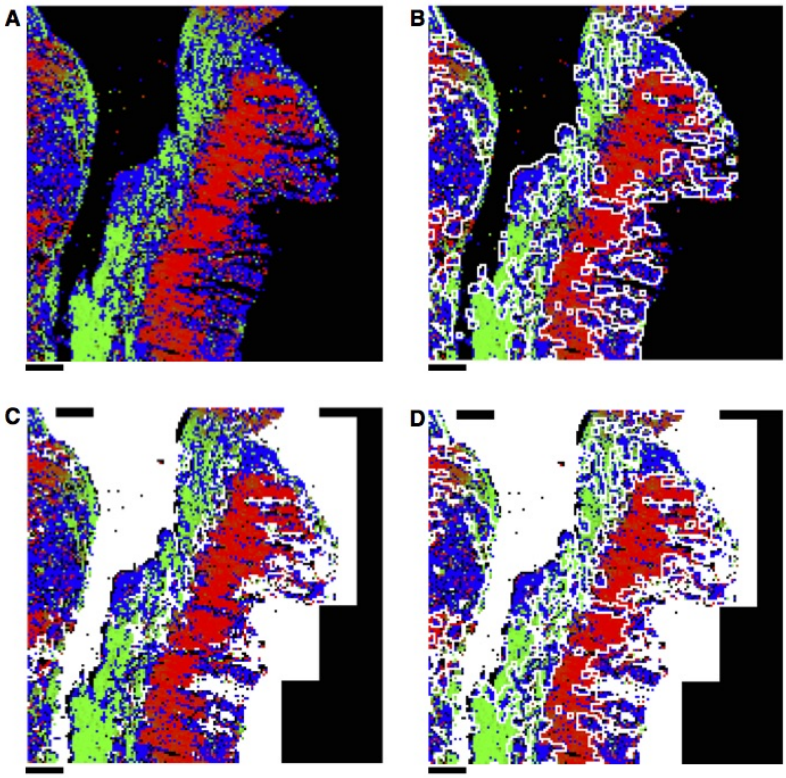

Figure 3.3: An example of the edges generated from the two separate edge detection algorithms. **A**: A colour representation of the regions, where blue indicates vertical regions, while red and green indicate planar regions. The red and green levels indicate direction; the green level is the y-component of the direction and the red level is the x-component. **B**: The edges obtained from the vertical regions. Here the edges, shown in white, have been drawn around areas of blue with sufficiently high nuclear counts. **C**: The edges obtained from the planar regions. Here a thresholding algorithm has turned all points in the grey image below the threshold white, and edges are formed between areas of mismatched directions. **D**: The combined edges from **B** and **C**. Scale bars represent 1 mm.

This is precisely the measure that can be used in the watershed algorithm to determine the edges of the structures. The parameters for the following procedure are given in Table 3.1.

A *planar* region has planar direction v given by an angle θ, such that

$$v = (\cos\theta, \sin\theta).$$

The angle $\theta + 180$ represents the same direction as θ, because rotating the direction vector 180° produces the same fibre direction. Thus an equally valid vector representation of the

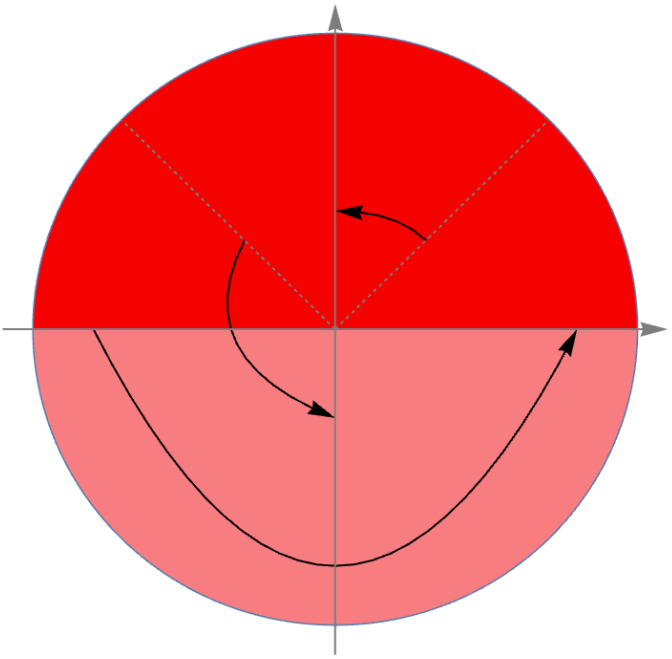

Figure 3.4: The process of doubling angles of vectors. Each vector is transformed by a rotation about the origin, doubling the polar angle. The dark red indicates the original position of the vectors, while the lighter red indicates the extra range included with transformation.

planar direction would be

$$\tilde{v} = (\cos(\theta + 180), \sin(\theta + 180)) = -v.$$

This means that the additive inverse of a direction vector represents the same fibre direction as the original vector. Applying a Gaussian mask to this image would therefore generate grey values which do not necessarily represent the local homogeneity: convolution with a Gaussian mask involves adding vector components in a local neighbourhood, and in this system vectors representing the same direction could sum to zero. In fact, even if the angles were reduced to a 180° interval there would always be a direction along which the homogeneity would not be accurately represented.

In order to formulate an image to use for Gaussian smoothing, the angle 2θ is used instead, as illustrated in Figure 3.4. The direction vector v' generated by this formulation

is given by

$$v' = (\cos(2\theta), \sin(2\theta)).$$

The vector $\tilde{v}'$ corresponding to the angle $\theta + 180$ is thus given by

$$\tilde{v}' = (\cos(2(\theta{+}180)), \sin(2(\theta{+}180)) = (\cos(2\theta{+}360), \sin(2\theta{+}360)) = (\cos(2\theta), \sin(2\theta)) = v'.$$

This means that vectors corresponding to equivalent angles are represented by the same direction vectors in this system.

Consider the vector image V' given by the above double-angle formulation. For a given point p, the vector obtained by the convolution of a Gaussian kernel G with V' is given by

$$(G * V')(p) = \sum_k g_k v'_k,$$

where g_k is the value of the Gaussian kernel corresponding to point p_k, v'_k is the vector value of p_k in V', and the set $\{p_k\}$ is the set of points in the window about p. Taking the norm squared of this vector yields

$$\begin{aligned} |(G * V')(p)|^2 &= \left| \sum_k g_k v'_k \right|^2 \\ &= \left(\sum_j g_j v'_j \right) \cdot \left(\sum_k g_k v'_k \right) \cdot \\ &= \sum_{jk} g_j g_k (v'_j \cdot v'_k) \end{aligned}$$

The scalar product $v'_j \cdot v'_k$ that occurs in the summand of this last expression is given by

$$v'_j \cdot v'_k = \cos(2\theta_j - 2\theta_k),$$

where θ_j and θ_k are the angles corresponding to the original vectors at p_j and p_k respec-

tively. This can be expressed in terms of these original vectors v_j and v_k:

$$v'_j \cdot v'_k = 2\cos^2(\theta_j - \theta_k) - 1 = 2(v_j \cdot v_k)^2 - 1.$$

Thus, the image I defined by

$$I(p) := |(G * V')(p)|^2$$

has entries

$$I(p) = \sum_{jk} g_j g_k (2(v_j \cdot v_k)^2 - 1).$$

This sum is a smoothed measure of the local homogeneity of the slide about p; the closer together the directions $\{v_k\}$ are, the larger this sum will be, and the terms are weighted by the Gaussian coefficients.

One final amendment made to this image is of weighting. In particular, the vectors are weighted by the nuclear count, giving the weighted form I_w of the image I to have entries

$$I_w(p) = \left| \sum_k g_k w_k v'_k \right|^2,$$

where w_k is the nuclear count at p_k, and all other terms are as before. Writing the entries in terms of scalar products:

$$I_w(p) = \sum_{jk} g_j g_k w_j w_k (2(v_j \cdot v_k)^2 - 1).$$

This weighting ensures that regions with lower nuclear counts, which are more likely to be at the edge of a fibrous structure, yield smaller terms than regions with higher nuclear counts. Reducing the values of terms from the edges of structures further sharpens these edges. An illustration of this weighted image is shown in Figure 3.5**C**. The image is a scalar-valued image that has been subjected to Gaussian smoothing, which is of the form required for the modified watershed algorithm. Prior to edge detection, a low level

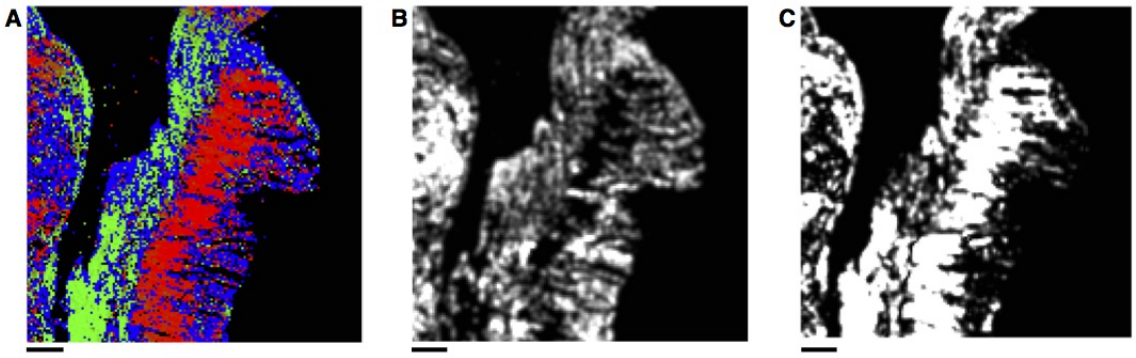

Figure 3.5: An example of the grey images generated from a vector valued image. **A**: A representation of the directions vectors, where red and green indicate the direction of planar vectors, and blue indicates vertical regions. **B**: The grey image generated from smoothing the nuclear counts present in the vertical regions. **C**: The grey image generated from the planar vectors represented in **A**. Scale bars represent 1 mm.

threshold is applied to the grey image to remove points with high levels of heterogeneity. This threshold ensures that any points which are manifestly not an interior part of any fibre are removed.

To apply the modified watershed algorithm, the comparison function f must also be defined. Consider the points p and q in the slide. The aim is to find a function which represents the following weighted average:

$$A(p,q) = \sum_{jk} \frac{g_j g_k w(p_j) w(q_k)}{\sum_{nm} g_n g_m w(p_n) w(q_m)} (v_j \cdot u_k)^2$$

where p_j are the points in the Gaussian window about p, with weights $w(p_j)$, vectors $V(p_j) = v_j$, and Gaussian coefficients g_j, and likewise for q_k. This sum is the weighted average of the squared scalar product of the vectors about p with the vectors about q, which represents the similarity of the directions in the windows about p and q. The following function f is directly representative of A:

$$f(p,q) := \frac{|(G * V'_w)(p) + (G * V'_w)(q)|^2 - I_w(p) - I_w(q)}{(G * W)(p)(G * W)(q)}$$

where V'_w is the weighted vector image with doubled angles, and W is the weight im-

age, which is equal to the nuclear count at each point. Expanding the first term in the numerator gives the following:

$$\begin{aligned}
|(G * V'_w)(p) + (G * V'_w)(q)|^2 &= \left| \sum_k g_k (w(p_k)v'_k + w(q_k)u'_k) \right|^2 \\
&= \sum_{jk} g_k g_j (w(p_k)v'_k + w(q_k)u'_k) \cdot (w(p_k)v'_k + w(q_k)u'_k) \\
&= \sum_{jk} g_k g_j w(p_k) w(p_j) v'_k \cdot v'_j + \sum_{jk} g_k g_j w(q_k) w(q_j) u'_k \cdot u'_j \\
&+ 2 \sum_{jk} g_k g_j w(p_k) w(q_j) v'_k \cdot u'_j \\
&= I_w(p) + I_w(q) + 2 \sum_{jk} g_k g_j w(p_k) w(q_j) v'_k \cdot u'_j \\
&= I_w(p) + I_w(q) + 2 \sum_{jk} g_k g_j w(p_k) w(q_j) (2(v_k \cdot u_j)^2 - 1).
\end{aligned}$$

Thus the comparison function f becomes

$$f(p,q) = \frac{2 \sum_{jk} g_k g_j w(p_k) w(q_j) (2(v_k \cdot u_j)^2 - 1)}{\sum_{nm} g_n g_m w(p_n) w(p_m)}$$

where the convolution in the denominator has been expanded. This is equivalent to

$$f(p,q) = 4A(p,q) - 2,$$

which means that f directly represents the similarity in the directions about p and q. Therefore, applying a threshold $f_{\min}$ to f determines whether the regions about p and q are similar enough to be considered part of the same fibrous structure.

The watershed determined by using this grey image and comparison function represents the edges of fibrous structures. An example of the full effect of this algorithm is illustrated in Figure 3.3**C**. This watershed is combined with the edges found for the *vertical* regions, to give the final edge image as illustrated in Figure 3.3**D**. These edges can be used to determine structures in the slides prior to registration, which can be used to verify the

structure of the three-dimensional structure.

3.3 Determining boundaries of the three-dimensional structure

The aim of this section is to determine how to generate the grey volume and comparison function to be used in the modified watershed algorithm in three dimensions. This process is analogous to determining these entities for the *planar* regions in the two-dimensional case. The grey image is determined by performing Gaussian smoothing on the vector-valued volume to determine the local heterogeneity. Gaussian smoothing in three dimensions is similar to that in two dimensions, but with the mask $G = G(x, y, z)$ now being given by a cube, with entries

$$G(x - w_{1/2}, y - w_{1/2}, z - w_{1/2}) = \frac{1}{\sigma^3 (2\pi)^{3/2}} \exp\left\{\frac{-(x^2 + y^2 + z^2)}{2\sigma^2}\right\}$$

where $w_{1/2}$ is the half-length of the cube.

The three-dimensional volume V obtained using the techniques described in Chapter 2 consists of a set of voxels each with a vector and weight assigned. These direction vectors are either zero, where there is no tissue present, or they are unit vectors in the direction of the tissue in the given voxel. The assigned weight represents an approximation of the number of nuclei present in the tissue at the given voxel. A three-dimensional direction vector v can be written in terms of spherical polar coordinates:

$$v = (\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi),$$

where $\phi \in [0, 180)$ and $\theta \in [0, 360)$. In the same manner as the two-dimensional case, the direction vector $-v$ represents the same tissue direction as v. Consider the vector v', given

by

$$v' = (\cos\theta \sin(2\phi), \sin\theta \sin(2\phi), \cos(2\phi)).$$

The vector $-v$ can be written

$$-v = (\cos\theta \sin(\phi + 180), \sin\theta \sin(\phi + 180), \cos(\phi + 180)).$$

Applying the double-angle transformation to $-v$ gives

$$\begin{aligned}(-v)' &= (\cos\theta \sin(2(\phi + 180)), \sin\theta \sin(2(\phi + 180)), \cos(2(\phi + 180))) \\ &= (\cos\theta \sin(2\phi), \sin\theta \sin(2\phi), \cos(2\phi)) \\ &= v',\end{aligned}$$

which is the analogous result to the two-dimensional case. This double-angle transformation is illustrated in Figure 3.6**A**. While this transformation converts vectors in opposite directions to vectors in the same direction, it does not necessarily convert perpendicular vectors to opposing vectors. Consider the vectors $u_1 = (1, 0, 0)$ and $u_2 = (0, 1, 0)$. The above transformation gives $u_1' = u_2' = (0, 0, -1)$. For this reason it is necessary to consider the analogous transformations shown in Figures 3.6**B** and 3.6**C**. These transformations use the same double-angle formulation, but the angle being doubled is the angle from the x- or y-axis. Applying each of these transformations to the volume V generates three volumes $V^{(i)}$, where $i \in \{0, 1, 2\}$ indicates the axis along which the angle is doubled.

Consider the convolution of $V^{(i)}$ with the three-dimensional Gaussian kernel G:

$$(G * V^{(i)})(p) = \sum_k g_k v_k^{(i)},$$

where g_k is the entry of the Gaussian kernel corresponding to the point p_k, $v_k^{(i)}$ is the direction vector of p_k in $V^{(i)}$, and the set $\{p_k\}$ is the set of points in the window about p.

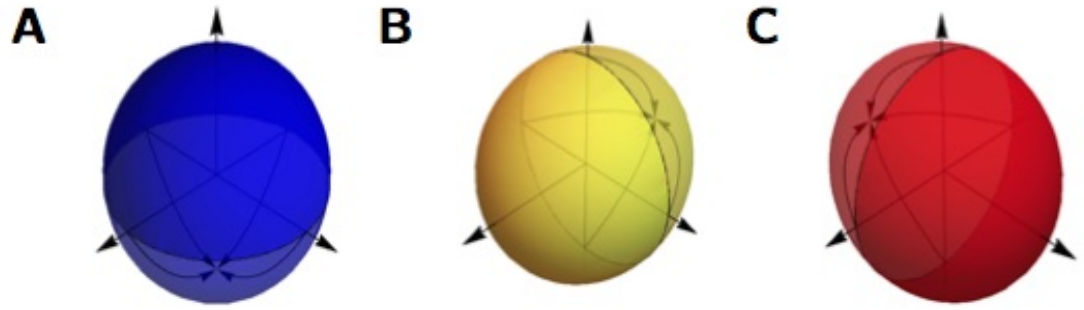

Figure 3.6: The three double angle transformations, characterised by doubling the angle from each of the axes separately. The darker colours represent untransformed vectors, taken to be positive in the relevant index. These hemispheres are stretched by the doubling of the angles, which pushes the domain to encompass the whole sphere.

Taking the norm squared of this vector gives:

$$\begin{aligned} |(G * V^{(i)})(p)|^2 &= \left| \sum_k g_k v_k^{(i)} \right|^2 \\ &= \left(\sum_j g_j v_j^{(i)} \right) \cdot \left(\sum_k g_k v_k^{(i)} \right) \\ &= \sum_{jk} g_j g_k v_j^{(i)} \cdot v_k^{(i)}. \end{aligned}$$

Let the angles (θ_j, ϕ_j) and (θ_k, ϕ_k) denote the spherical polar coordinates for v_j and v_k respectively, with ϕ corresponding to the angle from axis i. The scalar product in the above sum is given by

$$\begin{aligned} v_j^{(i)} \cdot v_k^{(i)} &= \cos\theta_j \sin(2\phi_j) \cos\theta_k \sin(2\phi_k) + \sin\theta_j \sin(2\phi_j) \sin\theta_k \sin(2\phi_k) + \cos(2\phi_j)\cos(2\phi_k) \\ &= \cos(2\phi_j)\cos(2\phi_k) + \sin(2\phi_j)\sin(2\phi_k)\cos(\theta_j - \theta_k) \\ &= (2\cos^2\phi_j - 1)(2\cos^2\phi_k - 1) + 4\cos\phi_j \cos\phi_k \sin\phi_j \sin\phi_k \cos(\theta_j - \theta_k) \\ &= (2v_{ji}^2 - 1)(2v_{ki}^2 - 1) + 4v_{ji}v_{ki}\sin\phi_j \sin\phi_k \cos(\theta_j - \theta_k) \end{aligned} \tag{3.1}$$

where v_{*i} is the ith component of v_*. The scalar product $v_j \cdot v_k$ is given by

$$\begin{aligned} v_j \cdot v_k &= \cos\theta_j \sin\phi_j \cos\theta_k \sin\phi_k + \sin\theta_j \sin\phi_j \sin\theta_k \sin\phi_k + \cos\phi_j \cos\phi_k \\ &= v_{ji} v_{ki} + \sin\phi_j \sin\phi_k \cos(\theta_j - \theta_k). \end{aligned}$$

Substituting this into equation (3.1) yields:

$$\begin{aligned} v_j^{(i)} \cdot v_k^{(i)} &= (2v_{ji}^2 - 1)(2v_{ki}^2 - 1) + 4v_{ji}v_{ki}(v_j \cdot v_k - v_{ji}v_{ki}) \\ &= 4v_{ji}^2 v_{ki}^2 - 2v_{ji}^2 - 2v_{ki}^2 + 1 + 4v_{ji}v_{ki}v_j \cdot v_k - 4v_{ji}^2 v_{ki}^2 \\ &= 4v_{ji}v_{ki}v_j \cdot v_k - 2v_{ji}^2 - 2v_{ki}^2 + 1. \end{aligned}$$

Define the volume image I by

$$I(p) := \sum_i |(G * V^{(i)})(p)|^2.$$

Using the above formulation of the scalar product for each i yields

$$\begin{aligned} I(p) &= \sum_i \sum_{jk} g_j g_k (4v_{ji}v_{ki}v_j \cdot v_k - 2v_{ji}^2 - 2v_{ki}^2 + 1) \\ &= \sum_{jk} g_j g_k \left(4v_j \cdot v_k \sum_i v_{ji}v_{ki} - 2\sum_i v_{ji}^2 - 2\sum_i v_{ki}^2 + 3 \right) \\ &= \sum_{jk} g_j g_k (4(v_j \cdot v_k)^2 - 1) \end{aligned}$$

using the fact that $|v_j| = |v_k| = 1$. This image differs from the two-dimensional case in that the squared scalar product term has a coefficient of 4 instead of 2, meaning that each term of the sum is in the range $[-1, 3]$, as opposed to $[-1, 1]$. This is an artefact of the fact that three double-angle transformed vectors are being summed: if two vectors u and v are contained in the (x, y)-plane, then

$$u^{(0)} \cdot v^{(0)} = u^{(1)} \cdot v^{(1)} = u' \cdot v',$$

where u' and v' have been transformed by the double-angle formulation used in Section 3.2. Additionally $u^{(2)} = v^{(2)} = (0, 0, -1)$, so

$$\sum_i (u^{(i)} \cdot v^{(i)}) = 2u' \cdot v' + 1.$$

Thus the image $\tilde{I}$, given by

$$\tilde{I}(p) = (I(p) - 1)/2$$

is equivalent to the image generated in Section 3.2 for two-dimensional direction vectors. As with the two-dimensional case, this volume image represents the smoothed homogeneity of the vector-valued volume. For ease of calculation, the following will use the original formulation of I; however, the linear relationship between I and $\tilde{I}$ means that the following derivations for I are compatible with $\tilde{I}$.

In Chapter 2, the volume was also assigned weights, determined by approximating the nuclear count within each voxel. This weight volume W can be used to modify the vector volumes $V^{(i)}$, generating the weighted vector volumes $V_w^{(i)}$, given by

$$V_w^{(i)}(p) = W(p)V^{(i)}(p).$$

The weighted volume I_w is given by

$$I_w(p) := \sum_i |(G * V_w^{(i)})(p)|^2$$

which can be rewritten in terms of scalar products of vectors:

$$I_w(p) = \sum_{jk} g_j g_k w_k w_j (4(v_j \cdot v_k)^2 - 1)$$

where w_k is the weight at p_k, and all other terms are as above. This weighted volume is the volume used for the modified watershed algorithm.

The comparison function f associated with this volume is given by

$$f(p,q) := \frac{\sum_i |(G * V_w^{(i)})(p) + (G * V_w^{(i)})(q)|^2 - I_w(p) - I_w(q)}{(G * W)(p)(G * W)(q)}.$$

In the same manner as the two-dimensional case, this comparison function represents the similarity of the vectors in the smoothed windows about p and q. Here, the threshold for f was set to 2.0.

3.3.1 Adaptive smoothing

One final modification to the above process is the incorporation of adaptive smoothing. The Gaussian windows in the above formulation are of arbitrary size; furthermore, the above formulation of the watershed algorithm can be used with a smoothing kernel that varies in size with position. This means that the Gaussian mask can be set locally to optimise the smoothing for each point. If a point is contained in a wide fibre, then a large smoothing kernel would allow greater noise reduction. Conversely, if a point were contained in a narrow fibre, then a small smoothing kernel would reduce the blurring of the boundaries of the fibre.

The particular adaptive smoothing method used here was previously described by Gomez [49]. The aim is to select Gaussian filters to minimise the following expression at each point p:

$$d = \frac{\lambda}{\sigma(p)^2} + \epsilon(p)^2,$$

where λ is a constant, $\sigma(p)^2$ is the variance of the Gaussian at p, and $\epsilon(p)^2$ is given by

$$\sum_{q \in N_p} \sum_i |G(q) * V_w^{(i)}(q) - V_w^{(i)}|^2,$$

where N_p is the first order neighbourhood of voxels, as shown in Figure 3.7. The minimisation is performed by computing the value of d for a range of values σ at each point, and selecting the value of σ corresponding to this minimum.

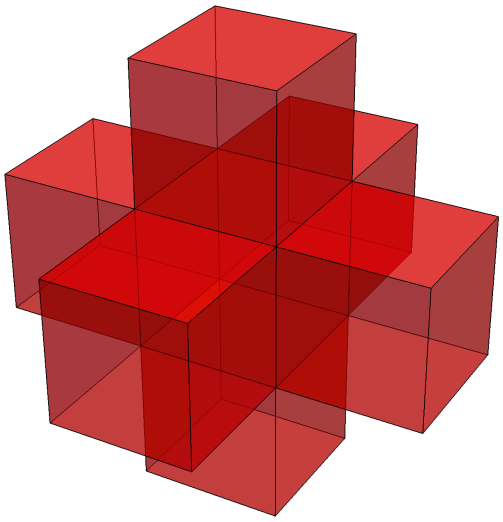

Figure 3.7: The first order neighbourhood of the central voxel in this image is given by the voxels coloured in red. This is analogous to the first order neighbourhood of a pixel in an image, which includes four neighbouring pixels.

Parameter	Value
λ	1.22
σ_{min}	0.4
σ_{max}	4.0
$\Delta\sigma$	0.1

Table 3.2: Parameter values used for adaptive smoothing.

The value of λ determines the behaviour of d with respect to σ. A large value of lambda would cause the value of d to be determined by σ with ϵ^2 having little influence. This would fix σ at the largest possible value for all points. Conversely, if λ were too small, d would be determined by ϵ^2, which would be minimal when σ is smallest. The value of λ used here was determined by locating the range of values between these two extremes through inspection of the behaviour of d and selecting a value within this range. The parameter values used for this adaptive smoothing are listed in Table 3.2.

The adaptive Gaussian mask generated in this manner refines the smoothing used in the edge detection by optimising the smoothing used in the generation of the grey volume used in the watershed algorithm. This in turn leads to a more precise segmentation. The effect of this refined edge detection algorithm is shown in Figure 3.8.

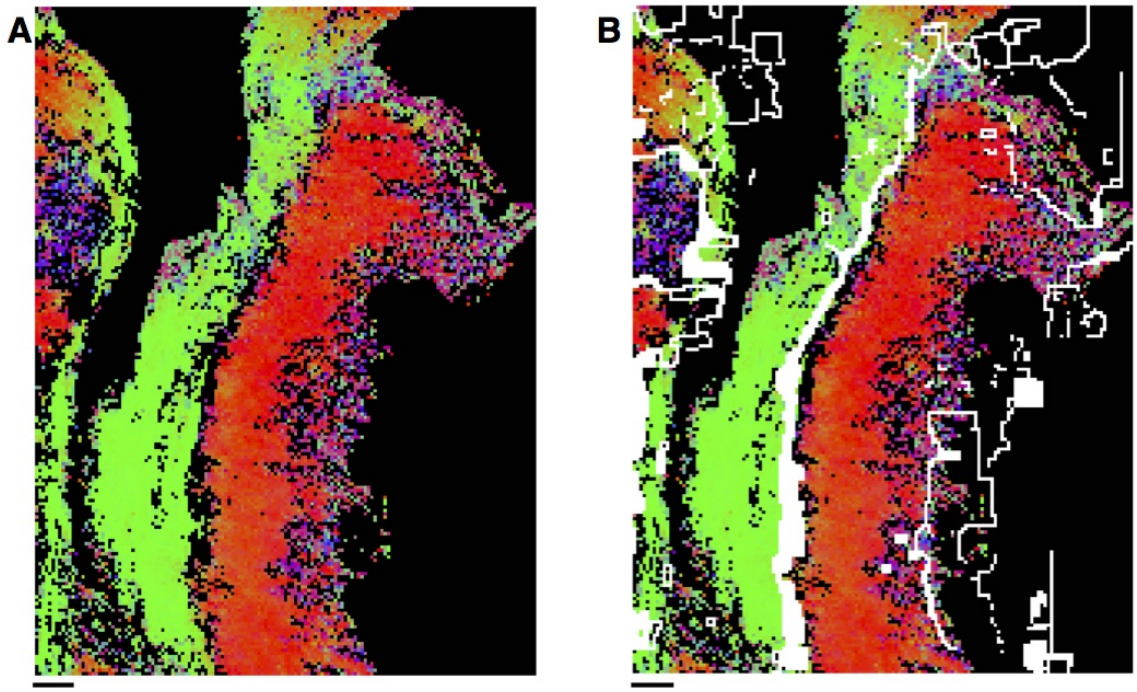

Figure 3.8: An example of the effect of the three-dimensional watershed-based edge detector. Here the image has been pseudocoloured by direction, with red representing the x-component, green representing the y-component, and blue representing the z-component. Scale bars represent 1 mm.

This completes the process of fibre segmentation in three dimensions. The edges generated by this process have been constructed to ensure that they completely encase each fibrous structure. These edges can be used to determine structural shape characteristics, as will be described in the following section. They are also used in excitation simulation to ensure that no communication occurs between bundles which are touching but not connected, as described in Chapter 4.

3.4 Isolating smooth muscle tissue

Throughout the processing of the tissue so far there has been no explicit differentiation of tissue types. It is essential that the smooth muscle tissue is isolated from the other tissue types to ensure that analysis of the structure and simulated excitation propagation specifically relate to the myometrium. The aim of this section is to separate the myometrium

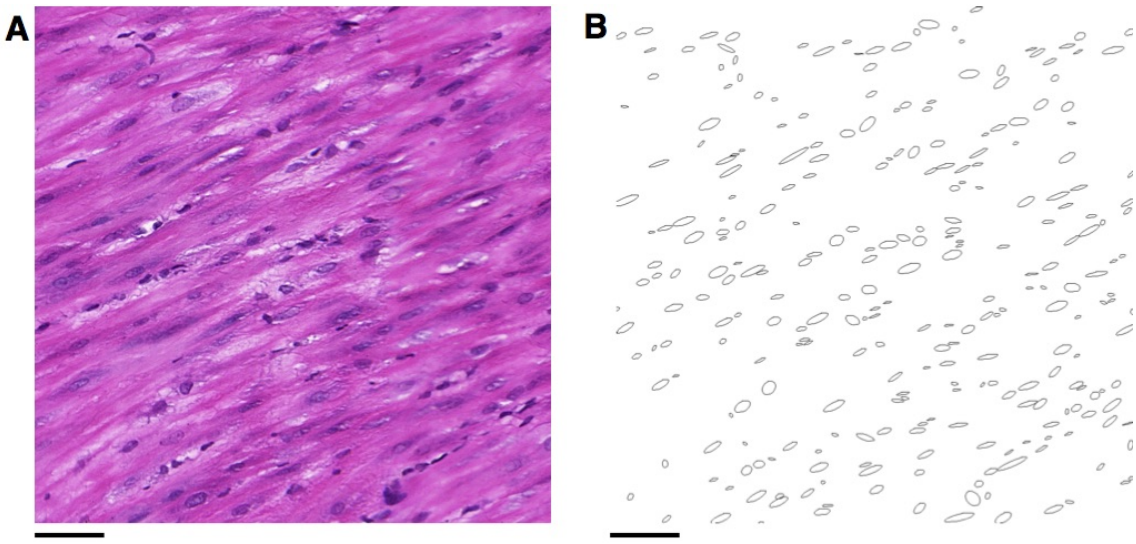

Figure 3.9: Within a sufficiently small volume the nuclei in smooth muscle tissue have similar orientation. **A**: A histological image of smooth muscle tissue. **B**: The orientation of the nuclei within the slide plane from **A**. Scale bars represent 50 μm.

from connective tissue, vasculature, and the placenta. This goal is achieved by analysing characteristics of the tissue that differ between myometrial smooth muscle and other tissue types.

The main representative characteristic of smooth muscle tissue used in this process is homogeneity of nuclei within each voxel. The direction of nuclei within smooth muscle tissue are relatively homogeneous in small volumes, as illustrated in Figure 3.9. Within a voxel, which has width $\sim$ 50 μm, it is reasonable to assume that the orientation of nuclei within smooth muscle tissue is approximately constant. Therefore, a low level of homogeneity within a voxel would indicate that the voxel does not contain smooth muscle. The method used to measure homogeneity is similar to the method for smoothing directions in two-dimensional segmentation (Section 3.2). For a given voxel, let S be the set of angles of orientation of the nuclei contained in the voxel and ϕ be the z-angle assigned to the voxel. Define the vector-valued function $v(\theta)$ by

$$v(\theta) := (\sin\phi\cos(2\theta), \sin\phi\sin(2\theta), \cos\phi).$$

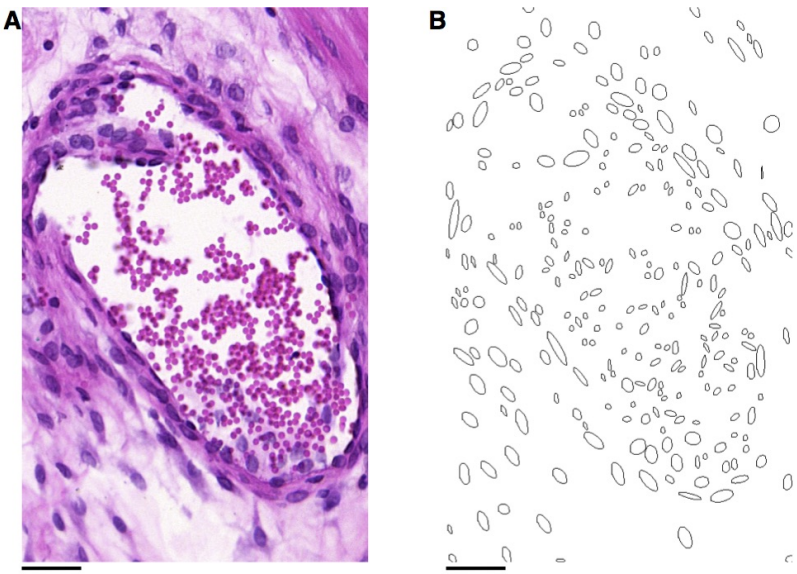

Figure 3.10: An example of red blood cells detected by the image processing algorithm. Normally these nuclei are not included in structural computation because they lie within a size range which is below the threshold used for smooth muscle cell nuclei. Scale bars represent 20 μm.

The measure of homogeneity for a voxel is defined to be

$$\frac{\left|\sum_{\theta \in S} v(\theta)\right|^2}{|S|}.$$

While three-dimensional vectors are used in this measure, it is not necessary to double the angles in the z-direction, because the value of z is constant within the voxel. The product of this homogeneity function with the original weighting function described in Chapter 2 yields a refined weighting function. This weighting function is used to generate the weight images used in Section 3.3. Additionally this function is used below to generate the final myometrial smooth muscle volume.

The above measure can be used to highlight smooth muscle fibres in the tissue. Unfortunately, this measure does not differentiate myometrial smooth muscle from vascular smooth muscle. In order to locate the vasculature, an alternative approach is required. The image analysis methods for detecting nuclei also detects red blood cells, as illustrated in Figure 3.10. This is because red blood cells are located in space which contains no other

Parameter	Value
Maximum red blood cell size	14 μm^2
Minimum vessel size	9 voxels
Gaussian variance	1.0 voxels

Table 3.3: Parameter values for detecting vasculature. Cell size is measured in the original histological images, vessel size is measured in terms of the voxels used in the reconstruction

cells for the algorithm to use as a background reference. The red blood cells are filtered out of the main tissue detection algorithm when the nuclei are thresholded by size, but are still recorded (see Chapter 2). For this reason they can be located using a threshold which targets the lower size range in which they are present. Using this threshold to locate red blood cells, the proportion of cells within a voxel which are deemed red blood cells is recorded. A Gaussian mask is applied to the red blood cell image. This smoothed image is thresholded to obtain potential locations for blood vessels. These points form clusters in the volume, and clusters with sufficiently large size are deemed to be points within blood vessels. These interior points are dilated by one voxel to include the tissue forming the wall of the vessels. While this technique provides some information on the location of blood vessels, it is not sufficient to remove all vascular smooth muscle. This is because not all blood vessels are filled with red blood cells to provide an accurate set of interior points. The intention of this process, however, is not to completely eradicate all blood vessels, but rather to reduce the impact of vascular smooth muscle on the overall structure so that it provides fewer erroneous connections between myometrial smooth muscle fibres. The parameter values used in this process are listed in Table 3.3.

The final technique used to reduce the reconstructions to myometrial smooth muscle tissue is specific to the pregnant rat uteri being processed, namely detecting placental tissue. Generally, local heterogeneity of the tissue is not sufficient to identify placenta, as shown in Figure 3.11. It is therefore necessary to determine the location of the sites using an alternative characteristic. The density of nuclei within the placenta is greater than in smooth muscle tissue, as can be observed in Figure 3.11. Accordingly, the nuclear density

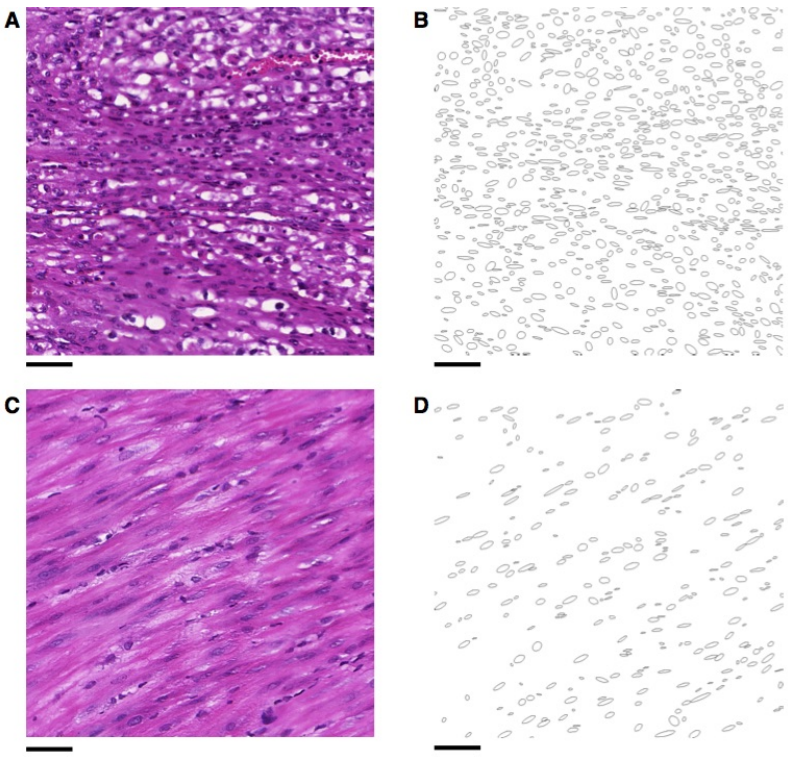

Figure 3.11: An example of the nuclei detected in an implantation site (**A**-**B**), compared to smooth muscle (**C**-**D**). The nuclei orientation in **B** are locally homogeneous. They are, however, much more dense than in smooth muscle (**D**), which enables differentiation of the implantation site voxels. Scale bars represent 50 μm.

was used to identify placental tissue.

Each voxel in a slide represents nuclei within a reduced size range, as described in Chapter 2. This reduction in size range was used because it represented the range of nuclear sizes in smooth muscle tissue. For the purpose of detecting placental tissue, the range of sizes was increased to include all nuclei detected in the histological images. The nuclear density of a voxel was defined to be the total number of nuclei from this extended range contained within the voxel. The threshold given in Table 3.4 was subsequently applied to these voxels to obtain a binary volume.

The aim of the remaining steps of this process is to identify large clusters of voxels that correspond to the placental tissue in the implantation sites. The binary volume was smoothed by using a large Gaussian kernel with $\sigma = 4$ voxels. This volume was

Parameter	Value
Minimum density	4500 nuclei per mm^2
Lower threshold	0.3
Upper threshold	0.6
Minimum cluster size	10000 voxels

Table 3.4: Parameter values for identifying implantation sites. Threshold values here are based on voxel values within the range $[0, 1]$.

subjected to a hysteresis thresholding procedure, as is used in Canny edge detection (see Appendix A). Hysteresis thresholding generates connected sets of points which all have a value above a lower threshold, and each connected set contains at least one point above an upper threshold. The upper and lower thresholds used here are given in Table 3.4. All such clusters above the size threshold given in Table 3.4 were deemed sufficiently large to represent implantation sites.

Using the above volume representations of various features in the tissue, it is now possible to obtain the final volume comprising of only myometrial smooth muscle. The process of obtaining this representation begins by taking the vector-valued volume obtained in Chapter 2 and removing points corresponding to edges (from Section 3.3), blood vessels, and placenta (if present). The vectors in this volume are then multiplied by the weighting functions described above and in Chapter 2. This vector volume is smoothed by convolving with a Gaussian kernel of radius 1 voxel. All points identified as edge points are excluded from this smoothing to maintain boundaries between distinct fibres. The directions of the resulting vectors are taken to be the fibre direction, with weights equal to the lengths of the vectors. The effect of this process is shown in Figure 3.12.

For the purpose of detecting fibre widths in registered and unregistered slides, it was also necessary to perform a similar operation on these two-dimensional slides. Since the registered slides directly correspond to the representative slides in the volume, the locations of vasculature and implantation sites correspond to the same features in the registered slides. These features were projected back onto the unregistered slides by identifying the points in the unregistered slides that are transformed to points corresponding to these

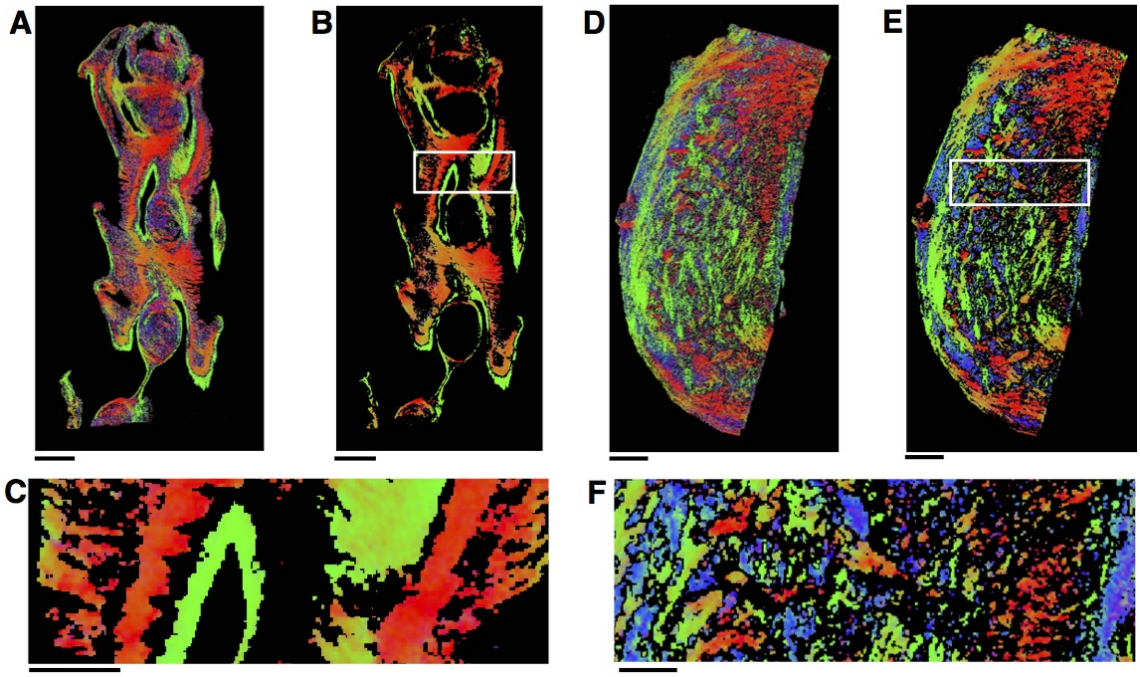

Figure 3.12: The combined effect of fibre segmentation, tissue segmentation, and smoothing on the structure on a slide from rat myometrium (**A**–**C**) and a slide from human myometrium (**D**–**F**). Images are pseudocoloured by direction: red represents fibres oriented horizontally in the image plane, green represents fibres oriented vertically in the image plane, and blue represents fibres oriented perpendicular to the image plane. **A**: Rat fibre directions prior to segmentation and smoothing. **B**: Rat fibre directions after segmentation and smoothing. **C**: Enlarged view of the white panel shown in **B**. **D**: Human fibre directions prior to segmentation and smoothing. **E**: Human fibre directions after segmentation and smoothing. **F**: Enlarged view of the white panel shown in **E**. Scale bars represent 5 mm in the top row and 2 mm in the bottom row.

features during registration. These features and the edges found by the segmentation algorithm described in Section 3.2 were removed from the tissue for the purposes of measuring fibre widths. The remaining planar directions were subsequently smoothed using a Gaussian kernel of radius 1 region length. These smoothed *planar* regions were used to measure fibre width, as described in Section 3.6.

The steps outlined in this section ensure that the volume analysed using the following steps is composed of myometrial smooth muscle. Furthermore, this new volume is used as the basis for simulations described in Chapter 4.

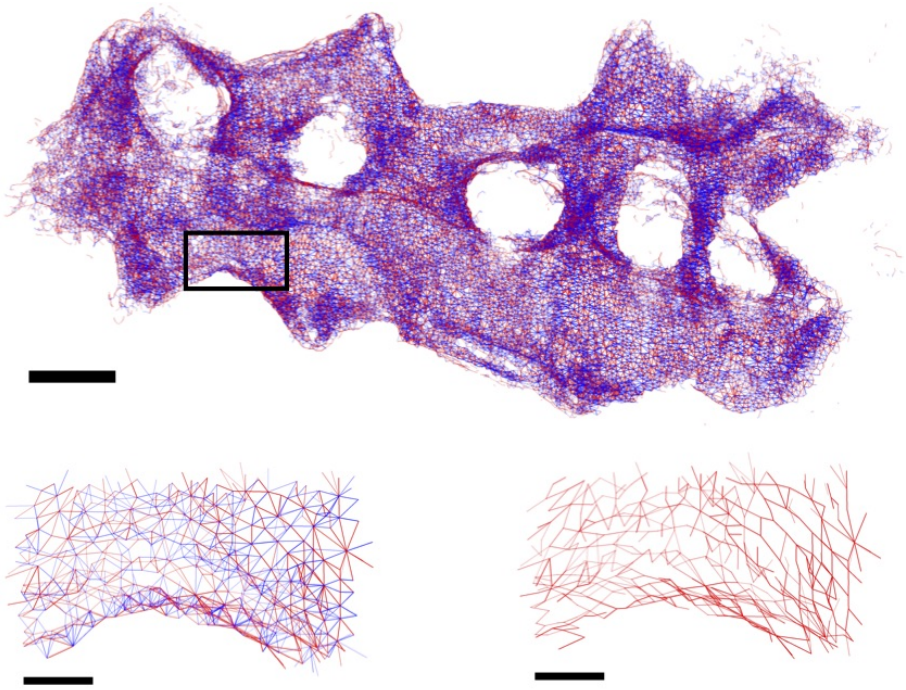

Figure 3.13: Top: an example of the network graph of the whole tissue. Bottom left: the network graph of a small portion of the tissue. Bottom right: the same portion of tissue with only connections along the fibre direction. Red arcs are connections along the fibre direction, blue arcs are connections away from the fibre direction. Top scale bar represents 1 mm, lower scale bars represent 200 μm.

3.5 Generating network representations of the tissue

This section describes the process of generating a representation of the reconstruction in terms of connectivity. This network model of the tissue is generated by using a watershed algorithm in a similar manner to that used for the segmentation of the tissue presented in Section 3.3. In this case, however, the structures isolated by the watershed algorithm are used to represent nodes of the network, with watershed boundaries representing the connections between these nodes. This network representation provides a visualisation of the tissue that enables the visual identification of areas of low connectivity between fibres. Furthermore, it enables the identification of pathways connecting separate points in the tissue.

The three-dimensional scalar image used as the basis for the watershed algorithm is the weighting image generated in Section 3.4. A standard watershed algorithm is applied to this image, as described in Section 3.1. This segmentation yields a set of isolated areas in the reconstruction. Each of these areas corresponds to a node in the network representation. The arcs of the network representation correspond to connections between these areas: if two segmented areas share a boundary, then an arc is drawn between the corresponding nodes. This process enables the generation of a basic network representation of the tissue.

Each node in this network corresponds to a cluster of voxels in the reconstructed tissue. For each node the size, centre of mass, and mean direction of the corresponding cluster are recorded. Each connecting arc is assigned a weight equal to the largest weight value in the set of border voxels between the two clusters connected by the arc. These measurements enabled the visualisation of orientation of the connections relative to fibre direction: the mean direction at two nodes connected by an arc represents the fibre direction along that arc, which can be compared to the direction of the arc to determine how close this arc is to the fibre direction. If a fibre is closer than 45° to the fibre direction, then it is classified as being along the fibre direction. Otherwise it is classified as being perpendicular to the fibre direction. This classification presents the network as a set of lines which either represent fibres or connections between them. Additionally, the weighting of each arc represents the strength of the connection. An example of a network representation of rat myometrium is shown in Figure 3.13.

3.6 Measuring the widths of fibres

This section describes the process for measuring the fibre widths in two and three dimensions. Comparison of the fibre widths in the three dimensional structure with those taken from the unregistered two-dimensional slides will provide verification of the structural model. The process for determining these width measurements follows the same general

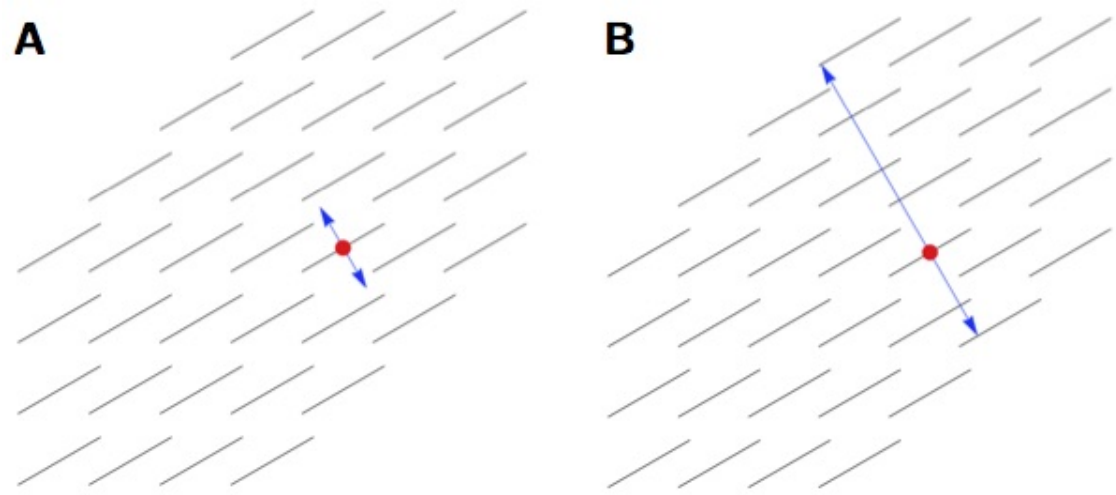

Figure 3.14: Measuring fibre widths using directional data. **A**: The direction of the width measurement taken at the grid point shown is perpendicular to the regional direction at the given point. **B**: The width measurement is taken by extending the line out from the initial grid point in both directions until it reaches the edge of the bundle.

form in both cases; random points are selected in the structure and the width is measured by drawing a line perpendicular to the tissue direction.

The process for measuring fibre width in two dimensions is as follows. The smoothed regional representation of the slides is used, as described in Section 3.4. A uniform grid of predetermined separation is placed at random on the tissue image. If the region at a given grid point is classified as *planar* in the smoothed regional representation, a measurement is recorded at that point as follows. The direction perpendicular to the planar direction of the region is used for width measurement. This direction selection is illustrated in Figure 3.14**A**. The length of the line segment along this width direction through the centre of the given region within the tissue is taken to be the width measurement at this point, as illustrated in Figure 3.14**B**. These width measurements are taken for each of the registered and unregistered slides, giving a set of width measurements in two dimensions for the whole tissue block in both cases. Comparison of the distributions of these widths will determine if the registration process has substantially deformed the tissue block.

The process for measuring fibre width in the three-dimensional volume is performed

in much the same manner as the two-dimensional case. The volume being measured is the smoothed volume described in Section 3.4. Voxels with weight below a threshold of 0.2 are excluded for the purpose of these measurements. This step is necessary in three dimensions because, in contrast to the two dimensional case, no threshold was applied to the weights during the segmentation process. The volume is divided by a uniform grid of predetermined width, and a point is randomly selected from each cube of the grid. For a given point, the width direction is selected at random from the set of directions perpendicular to the tissue direction at that point. The line within the tissue in the given direction is measured in the same manner as the two-dimensional case. These width measurements give a sample of the fibre widths in the three-dimensional volume.

The techniques presented in this chapter enable the analysis and verification of the three-dimensional structure determined in Chapter 2. Furthermore, the smoothed and segmented fibre representation provides a more accurate structure on which to base simulated excitation described in Chapter 4.

Chapter 4

Modelling excitation propagation

This chapter presents the mathematical and computational techniques used to solve the electrodiffusion equations. As described in Chapter 1, the FitzHugh-Nagumo model is used to approximate the behaviour of the cell membrane. This is combined with a diffusion element to give the full equations. The system uses the structure found in Chapter 3 as the foundation for the anisotropic diffusion. The numerical solution of these electrodiffusion equations is performed using the finite element method. This technique produces a system of linear equations to be solved. These equations are solved using the Gauss-Seidel iterative method. These iterations are parallelised and processed using general-purpose computing on graphics processing units (GPGPU). All code used in these calculations can be found at `https://figshare.com/s/651e592a91f309182edd`.

4.1 Electrodiffusion equations

The aim of this section is to derive the equations used in the model of excitation in the myometrium. This derivation is based on the formulation developed for cardiac modelling [26], coupled with the FitzHugh-Nagumo model of cell membrane conductance [82]. Excitation is defined as a substantial temporary excursion of the membrane potential of a cell, away from its resting value. Consider an anisotropically conductive medium with conductivity tensor $\mathbf{G}$, which represents the conductance of the medium in all directions

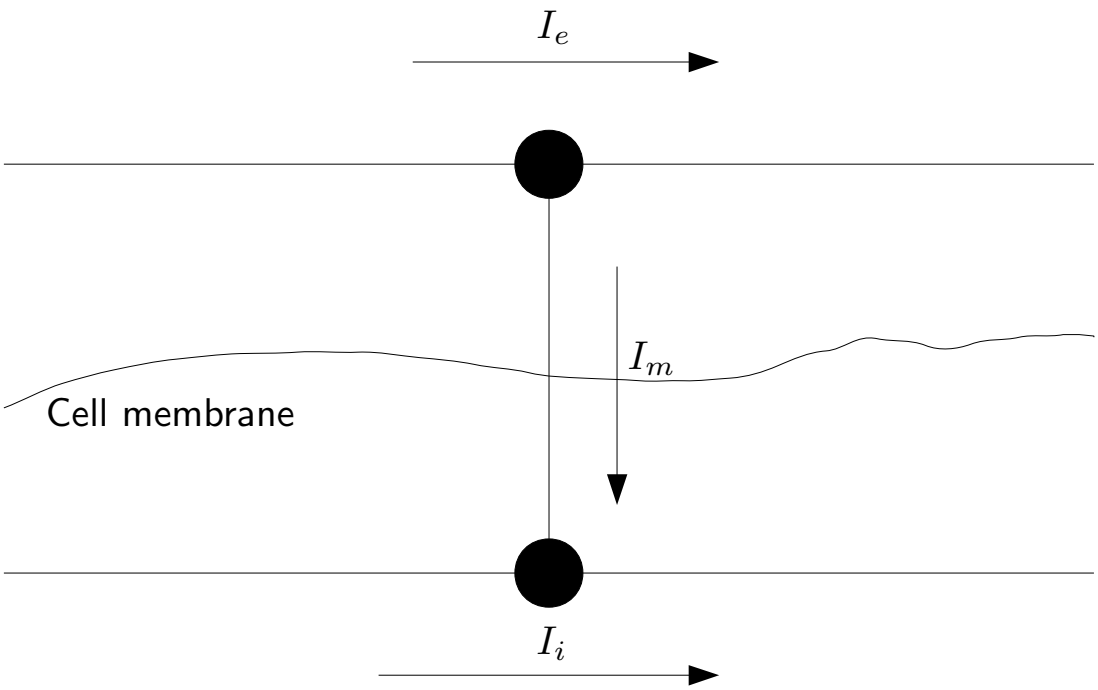

Figure 4.1: The relationship between intra-cellular current I_i, extra-cellular current I_e, and transmembrane current I_m. The current I_m is governed by the conductive elements present in the cell membrane.

at any given point. The current density, $\mathbf{J}$, is given by Ohm's law to be

$$\mathbf{J} = -\mathbf{G}\nabla\phi, \tag{4.1}$$

where ϕ is the electric potential.

The present model accounts for these quantities in both intra- and extra-cellular media, which are denoted with subscript i and e respectively. Applying this notation to the equation (4.1), the intra- and extra-cellular current densities are given by

$$\mathbf{J}_i = -\mathbf{G}_i\nabla\phi_i$$

$$\mathbf{J}_e = -\mathbf{G}_e\nabla\phi_e.$$

The two media interact by way of a transmembrane current, I_m, as shown in Figure 4.1.

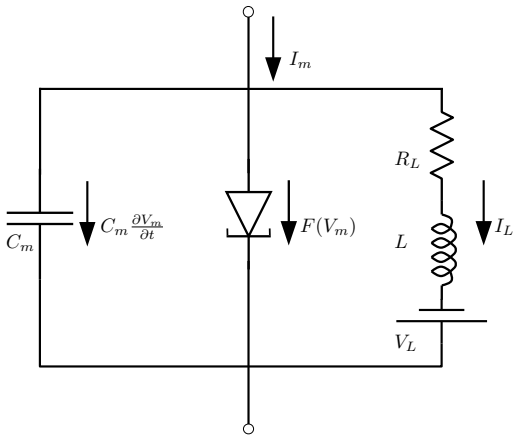

Figure 4.2: The circuit used to model membrane potential in the FitzHugh-Nagumo model [82]. The membrane current I_m is split into three parallel currents: the capacitive current, given by $C_m \partial V_m / \partial t$; the current through the tunnel diode, given by $F(V_m)$; and the current through the resistor-inductor circuit, denoted I_L.

This current is the only local source/sink of current in both media. Accordingly, Kirchhoff's first law gives

$$\nabla \cdot \mathbf{J_i} = -I_m \tag{4.2}$$

$$\nabla \cdot \mathbf{J_e} = I_m, \tag{4.3}$$

where I_m is the transmembrane current per unit volume. Summing these two equations gives the following relationship:

$$\nabla \cdot (\mathbf{J_i} + \mathbf{J_e}) = 0. \tag{4.4}$$

The FitzHugh-Nagumo model is used to simulate the behaviour of the transmembrane current [82]. This model splits the transmembrane current into various contributions: a capacitive component, simulating the capacitive nature of the cell membrane; a voltage-dependent resistor, such as a tunnel diode; and a resistor-inductor circuit, as illustrated in Figure 4.2. The latter two of these components model the involvement of ion channels in conducting the transmembrane current. In this model, the transmembrane current is

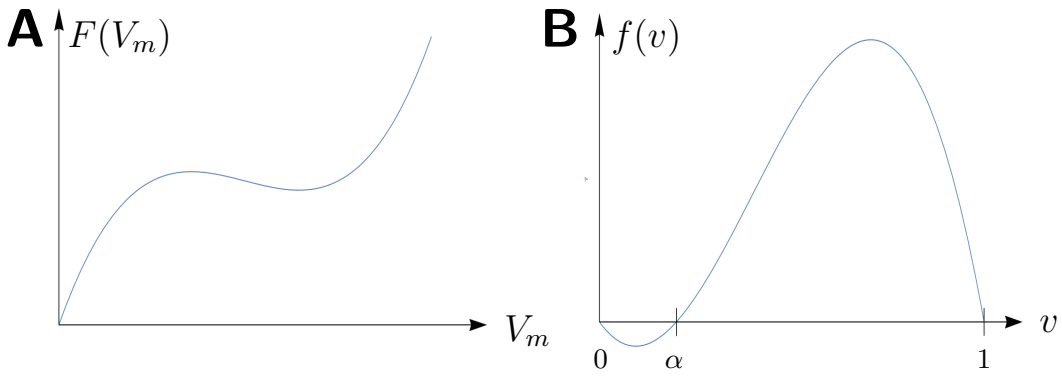

Figure 4.3: The current through the tunnel diode as a function of potential difference **A**: An example curve $F(V_m)$ defining the current through the tunnel diode as a function of membrane voltage. **B**: The nondimensional function $f(v)$, which is a rescaled and translated version of $-F(V_m)$ such that $f(0) = 0$ and $f(1) = 0$.

the sum of the currents in each of these parallel components, which is given by

$$I_m = \beta_m \left(C_m \frac{\partial V_m}{\partial t} + F(V_m) + I_L \right), \tag{4.5}$$

where β_m is the surface area to volume ratio of a cell, C_m is the membrane capacitance per unit area, $F(V_m)$ is the current through the voltage-dependent resistor per unit area, I_L is the current through the resistor-inductor circuit per unit area, and V_m is the membrane voltage, given by

$$V_m = \phi_i - \phi_e. \tag{4.6}$$

The precise mathematical specification of the function F is not essential for the correct behaviour of this model, provided that the general shape is as illustrated in Figure 4.3**A**. In the simulations a following cubic is used, given in Equation 4.13 below.

Kirchhoff's second law on the resistor-inductor circuit yields

$$L \frac{\partial I_L}{\partial t} = V_m - V_L - R_L I_L, \tag{4.7}$$

where L is the inductance of the inductor, V_L is the potential induced by the battery,

and R_L is the resistance of the resistor in the inductor circuit.

Combining the definitions of the current densities and Equations (4.2), (4.4), (4.5), and (4.6) gives

$$\nabla \cdot \mathbf{G}_i(\nabla V_m + \nabla \phi_e) = \beta_m \left(C_m \frac{\partial V_m}{\partial t} + F(V_m) + I_L \right) \tag{4.8}$$

$$\nabla \cdot ((\mathbf{G}_i + \mathbf{G}_e)\nabla \phi_e) = -\nabla \cdot (\mathbf{G}_i \nabla V_m). \tag{4.9}$$

Equations (4.7), (4.8), and (4.9) are the dimensional electrodiffusion equations. These equations are nondimensionalised by defining

$$v = \frac{V_m}{V_1}, \quad u = \frac{\phi_e}{V_1}, \quad w = \frac{AR_L I_L}{V_1}, \quad \tau = \frac{t}{AC_m R_L}, \quad f(v) = \frac{-AR_L F(V_1 v)}{V_1}.$$

Where V_1 is the peak voltage of the system and A is the surface area of a cell. In terms of these scaled quantities, the equations take on the following form:

$$\frac{\partial v}{\partial t} = \nabla \cdot (\mathbf{D}_i \nabla v) + \nabla \cdot (\mathbf{D}_i \nabla u) + f(v) - w \tag{4.10}$$

$$\nabla \cdot (\mathbf{D}_{i+e} \nabla u) = -\nabla \cdot (\mathbf{D}_i \nabla v) \tag{4.11}$$

$$\frac{\partial w}{\partial t} = \epsilon(v - \gamma w + v_0), \tag{4.12}$$

where $\mathbf{D}_i = AR_L \mathbf{G}_i / \beta_m$, $\mathbf{D}_{i+e} = AR_L(\mathbf{G}_i + \mathbf{G}_e)/\beta_m$, $\epsilon = AC_m R_L^2 / L$, $\gamma = 1/A$, $v_0 = V_L / V_1$, and t has been substituted for τ for ease of notation. The cubic function f is defined by

$$f(v) = Av(1 - v)(v - \alpha). \tag{4.13}$$

where A and α are positive constants with $0 < \alpha < 1$. This function is shown graphically in Figure 4.3**B**. Equations (4.10)–(4.12) are known as the *bidomain* equations [26]. These equations can be simplified by approximating $\mathbf{G}_e$ by a constant multiple of $\mathbf{G}_i$:

$$\mathbf{G}_e \approx \lambda \mathbf{G}_i. \tag{4.14}$$

This yields the equations

$$\frac{\partial v}{\partial t} = \nabla \cdot (\mathbf{D}\nabla v) + f(v) - w \tag{4.15}$$

$$\frac{\partial w}{\partial t} = \epsilon(v - \gamma w + V_0), \tag{4.16}$$

where $\mathbf{D} = \lambda A R_L \mathbf{G}_i / ((1+\lambda)\beta_m)$ and all other constants are as before.

Equations (4.15) and (4.16) are known as the *monodomain* equations [26]. The variable v is called the *fast* variable and the variable w is called the *slow* variable in this system.

The final step in the set up of the electrodiffusion equations is the boundary condition. The boundary of the system is the edge of the tissue, and therefore no current will flow through the boundary, giving the boundary condition

$$\frac{\partial v}{\partial x} \cdot n = 0,$$

where n is the direction normal to the boundary.

These equations form the mathematical basis for excitation simulation for the remainder of this chapter. The fast variable v represents membrane potential, and thus the following sections will present a methodology to approximate the value of v numerically.

4.2 Conductivity Tensor

The conductivity tensor is determined by the anisotropy of the medium [78]. Chapters 2 and 3 provide a means of determining fibre direction in the myometrium. The diffusion tensor is constructed from this fibre direction as follows.

For a current density J and electric potential ϕ, the conductivity tensor D satisfies Ohm's Law:

$$J = -D\nabla\phi.$$

Figure 4.4 shows the ellipsoidal representation of the conductivity tensor. In this representation, the conductance along each unit vector u_i is κ_i. This is represented in formula

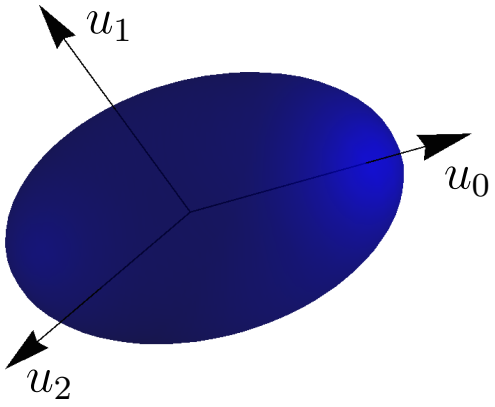

Figure 4.4: The ellipsoid form of the conductivity tensor D. The eccentricity of the ellipsoid along each vector u_i is the conductance along u_i, given by κ_i. In the case of conductance in the myometrium, the conductance is uniform in the transverse direction, so $\kappa_1 = \kappa_2$.

by

$$Du_i = \kappa_i u_i.$$

This means that D is a matrix with eigenvectors $\{u_i\}$ and corresponding eigenvalues $\{\kappa_i\}$. Therefore D can be written

$$D = UKU^T,$$

where U is the matrix with ith column u_i and K is the diagonal matrix with ith value κ_i. This form can be simplified further by using the fact that $\kappa_1 = \kappa_2$, as illustrated in Figure 4.4. Subtracting $\kappa_1 I$:

$$D - \kappa_1 I = UKU^T - \kappa_1 I = UKU^T - \kappa_1 UIU^T = U(K - \kappa_1 I)U^T = (\kappa_0 - \kappa_1)u_0^T u_0.$$

Thus for a given point x, the matrix $\mathbf{D}(x)$ is given by

$$\mathbf{D}(x) = \kappa_t I + (\kappa_l - \kappa_t)u^T(x)u(x),$$

where $\kappa_l = \kappa_0$ is the longitudinal conductance, $\kappa_t = \kappa_1$ is the transverse conductance,

and $u(x)$ is the fibre direction at x.

The vectors $u(x)$ are directly represented in the structure found using the process described in Chapter 2. The transverse and longitudinal conductances, however, are not inherently determined from the structure alone. This is because they are dependent on the density of gap junctions in the cell membranes, as explained in Chapter 1. The local structure can be used to assign weights to the conductance values, as described in Chapters 2 and 3. Given a weighting function $w(x)$, the weighted form of the diffusion tensor is given by

$$\mathbf{D}(x) = w(x)(\kappa_i I + (\kappa_l - \kappa_t)u^T(x)u(x).$$

This weighted form is used in the simulation of the electrodiffusion equations. The values for κ_l and κ_t are approximated using techniques given in Chapter 5.

4.3 The Finite Element Method

The finite element method (FEM) is a technique for numerically solving partial differential equations [69]. This section presents the application of FEM to the monodomain equations derived in Section 4.1. The monodomain equations are given by

$$\frac{\partial v}{\partial t} = \nabla(\mathbf{D}\nabla v) + f(v) - w \tag{4.17}$$

$$\frac{\partial w}{\partial t} = \epsilon(v - w) \tag{4.18}$$

The objective is to develop a method for obtaining a sequence of arrays $\{V_i(t_m)\}$, which approximate the values of v on a discretisation $\Omega_h = \{x_i\}$ of the spatial domain Ω for the sequence of discrete times $\{t_m\}$. Sections 4.3.1 and 4.3.2 describe the process of discretising the spatial domain, while Section 4.3.3 discusses discretisation in time.

The methods presented here give an overview of the techniques required; however, much of the theoretical background has been omitted for the sake of brevity. A full

description of the mathematical concepts required for a rigorous formulation can be found in [69].

4.3.1 Spatial discretisation

This section will present the methods used to approximate v for a given spatial discretisation of the domain Ω. Throughout this section, the notation $v(x)$ is used in place of $v(x,t)$ for the sake of convenience. Since equation (4.18) has no spatial derivative, it will not feature in this section, which will be focussed on the numerical approximation of the following version of equation (4.17):

$$\frac{\partial v}{\partial t} = \nabla(\mathbf{D}\nabla v) + f(v,w), \tag{4.19}$$

where

$$f(v,w) := f(v) - w.$$

The numerical methods presented here use the *weak* form of equation (4.19). This requires the following definition.

Definition 1 *The weak derivative, u, of the function v, is defined on a domain Ω by*

$$\int_\Omega \psi u = -\int_\Omega v\nabla\psi$$

for all differentiable functions ψ with $\psi|_{\partial\Omega} = 0$.

In the above definition, if v is sufficiently smooth, then $u = \nabla v$. The *weak* solution of equation (4.19) is defined to be v such that, for any differentiable function ψ with $\psi|_{\partial\Omega} = 0$,

$$\int_\Omega \frac{\partial v}{\partial t}\psi = -\int_\Omega \mathbf{D}\nabla v\cdot\nabla\psi + \int_\Omega f(v,w)\psi, \tag{4.20}$$

For ease of notation the following definition is used.

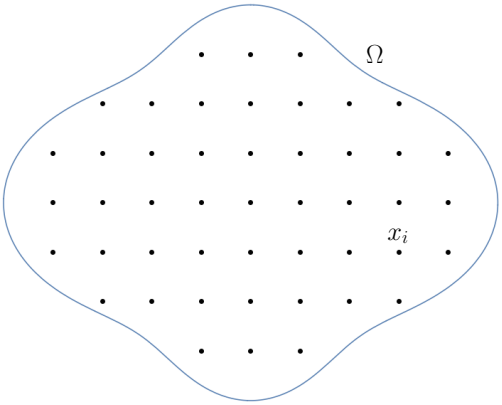

Figure 4.5: Discretisation of the two-dimensional domain Ω. The set of points is a discrete set of equally spaced points $\Omega_h = \{x_i\}$ such that $\Omega_h \subset \Omega$. This set is used to formulate the discrete approximation of the reaction-diffusion equations in space.

Definition 2 *The inner product* $(\cdot, \cdot)$ *is defined by*

$$(u, v) := \int_\Omega u(x)v(x)\mathrm{d}x$$

In addition, the functional $a(\cdot, \cdot)$ is defined to be

$$a(u, v) := -(\mathbf{D}\nabla u, \nabla v).$$

Thus equation (4.20) becomes

$$\left(\frac{\partial v}{\partial t}, \psi\right) = a(v, \psi) + (f(v, w), \psi). \tag{4.21}$$

This formulation will be used below to define the approximate solution to equation (4.19).

The next step is to reformulate the problem in terms of piecewise linear functions. The discretisation of the domain Ω is given by the lattice $\Omega_h = \{x_i\}$, with spacing h, as illustrated for two dimensions in Figure 4.5. For ease of notation the subscript h shall be used in the rest of this section to denote piecewise linear functions relating to Ω_h. Define

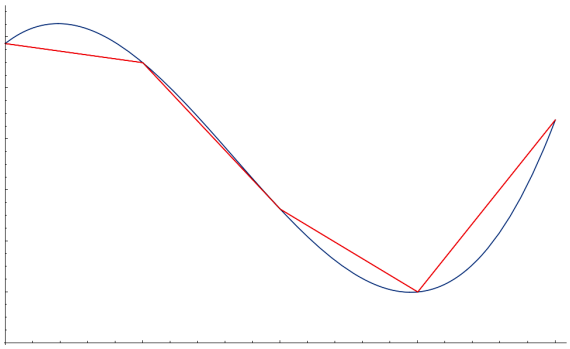

Figure 4.6: Spatial discretisation of a function in one dimension. Here the function v, shown in blue, is approximated by the continuous piecewise linear function v_h, shown in red.

the continuous piecewise linear function v_h such that

$$v_h(x_i) = v(x_i).$$

An example of this approximation for one dimension is illustrated in Figure 4.6. Similar formulations for w_h and ψ_h yield the piecewise linear approximation to equation (4.21):

$$\left(\frac{\partial v_h}{\partial t}, \psi_h\right) = a(v_h, \psi_h) + (f(v_h, w_h), \psi_h) \tag{4.22}$$

for any continuous piecewise linear function ψ_h with $\psi_h|_{\partial\Omega} = 0$. This formulation can now be converted into a system of linear equations, as follows. Suppose that the set of functions $\{\phi_i\}$ is a basis for the space of continuous piecewise linear functions on Ω. In order for it to be a basis it must conform to the following two conditions. The first condition is that $\{\phi_i\}$ *spans* the space of continuous piecewise linear functions on Ω. This means that for any continuous piecewise linear function u_h, there exists a vector U such that

$$u_h(x) = \sum_i U_i \phi_i$$

for all $x \in \Omega$. The second condition is that $\{\phi_i\}$ are *linearly independent*. This means that

if a set of scalar values $\{\lambda_i\}$ has the property

$$\sum_i \lambda_i \phi_i(x) = 0$$

for all $x \in \Omega$, then

$$\lambda_i = 0$$

for all i. Using the spanning condition, any ψ_h can be written

$$\psi_h = \sum_i \lambda_i \phi_i$$

for some vector (λ_i). Thus, if all the elements ϕ_i in the basis satisfy equation (4.22):

$$\left(\frac{\partial v_h}{\partial t}, \phi_i\right) = a(v_h, \phi_i) + (f(v_h, w_h), \phi_i)$$

then, by linearity, ψ_h satisfies it as well:

$$\begin{aligned}\left(\frac{\partial v_h}{\partial t}, \psi_h\right) &= \sum_i \lambda_i \left(\frac{\partial v_h}{\partial t}, \phi_i\right) \\ &= \sum_i \lambda_i (a(v_h, \phi_i) + (f(v_h, w_h), \phi_i)) \\ &= a(v_h, \psi_h) + (f(v_h, w_h), \psi_h).\end{aligned}$$

Conversely, if equation (4.22) holds for all ψ_h then it holds for $\psi_h = \phi_i$. Therefore, the new formulation of equation (4.22) is

$$\left(\frac{\partial v_h}{\partial t}, \phi_i\right) = a(v_h, \phi_i) + (f(v_h, w_h), \phi_i) \tag{4.23}$$

for all ϕ_i.

The spanning property of $\{\phi_i\}$ means that there exists a vector V such that

$$v_h = \sum_i V_i \phi_i.$$

Similarly there exists a vector F defined by

$$f_h = \sum_i F_i \phi_i.$$

With this notation, equation (4.23) becomes

$$\left(\frac{\partial v_h}{\partial t}, \phi_j\right) = \sum_i a(\phi_i, \phi_j) V_i + \sum_i (\phi_i, \phi_j) F_i = (KV + MF)_j \tag{4.24}$$

for all ϕ_j, where we have defined the *stiffness* matrix K by

$$K_{ij} := a(\phi_i, \phi_j)$$

and the *mass* matrix M by

$$M_{ij} := (\phi_i, \phi_j).$$

Equation (4.24) is a set of linear equations in V_i and F_i. These equations can be solved numerically, given a time discretisation and a set of basis functions $\{\phi_i\}$, as shall be described in Section 4.4.1.

4.3.2 Constructing the basis of piecewise linear functions

The next step is to select a suitable set of piecewise linear basis functions $\{\phi_i\}$. The aim is to define $\{\phi_i\}$ such that K and M are sparse matrices. This aim is motivated by the fact that, if K and M are sparse, then the equations to be solved will have a relatively small number of elements, and so will require a minimal amount computational effort. In order to fix the set $\{\phi_i\}$, the following definition is used.

Definition 3 *For a function ϕ, the* support *of ϕ is defined to be the set of points x with $\phi(x) \neq 0$.*

Suppose that, for a given pair of basis functions ϕ_i and ϕ_j, the supports of these functions do not intersect. Then for any $x \in \Omega$, either $\phi_i(x) = 0$ or $\phi_j(x) = 0$, so

$$(\phi_i, \phi_j) = \int_\Omega \phi_i(x)\phi_j(x) = 0$$

and

$$a(\phi_i, \phi_j) = \int_\Omega \mathbf{D}\nabla\phi_i(x) \cdot \nabla\phi_j \mathrm{d}x = 0.$$

For this reason, $\{\phi_i\}$ is constructed such that the supports of these functions intersect at a minimal amount of points. Note that these intersections cannot be completely eliminated because this would lead to the set $\{phi_i\}$ not satisfying the spanning condition. This motivates the definition a set of piecewise linear functions $\phi_i(x)$ such that

$$\phi_i(x_j) = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{if } j \neq i \end{cases}. \tag{4.25}$$

For the purposes of illustration, the remainder of this section will use the two-dimensional analogue of these three-dimensional functions. While this reduction in dimensionality does not affect the underlying theory, it does reduce the complexity of the system. The support of each ϕ_i in this two-dimensional analogue is illustrated in Figure 4.7**A**. This area has been triangulated into six separate regions, such that ϕ_i is a linear function ϕ_i^k in each region k. Each ϕ_i^k has the general form

$$\phi_i^k(y) = 1 + \sum_j C(k,j)\frac{y_j - x_{i,j}}{h},$$

Where $C(k,j) \in \{-1, 0, 1\}$ is such that each ϕ_i^k satisfies equation (4.25). This is illustrated in Figure 4.7**B**. The overlapping supports of $\{\phi_i\}$ are illustrated in Figure 4.7**C**. The

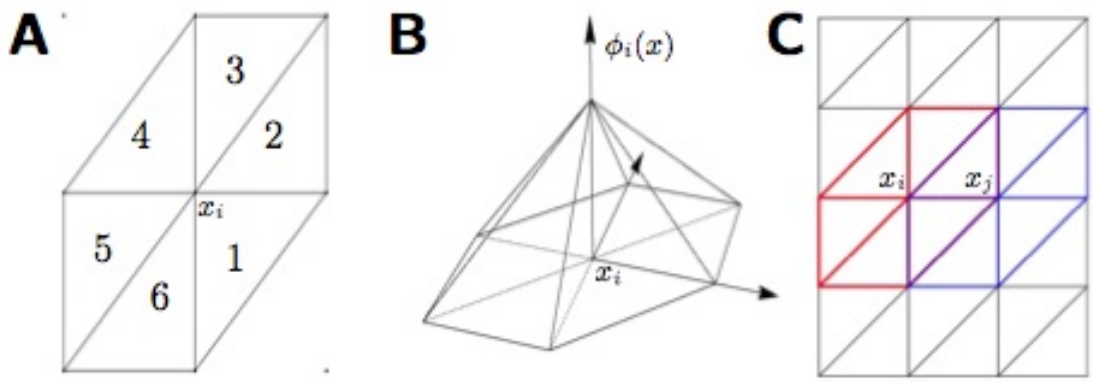

Figure 4.7: The two-dimensional function ϕ_i and its support. **A**: The support of the function $\phi_i(x)$. The support is triangulated, with each numbered triangle corresponding to a linear element of ϕ_i, as shown in **B**. **B**: The function $\phi_i(x)$; the linear components are defined such that $\phi_i(x_j) = 0$ for $j \neq i$. **C**: The supports of two functions ϕ_i, with triangulation coloured red, and ϕ_j, with triangulation coloured blue. The triangulation of these supports is defined such that the triangles align. This means that within any given triangle k, each function ϕ_i is given by a linear function ϕ_i^k.

support of each ϕ_i only intersects the support of ϕ_j if x_i and x_j are adjacent. In these overlaps, the triangulations of the supports line up, so that each entry M_{ij} in the mass matrix corresponding to these functions can be written

$$M_{ij} := (\phi_i, \phi_j) = \int_\Omega \phi_i(y)\phi_j(y)\mathrm{d}y = \sum_k \int_{\Omega_k} \phi_i^k(y)\phi_j^k(y)\mathrm{d}y,$$

where Ω_k is defined to be the area of triangle k, ϕ_i^k and ϕ_j^k are the linear functions of ϕ_i and ϕ_j defined on triangle k, and the sum is over all triangles k contained in the supports of both ϕ_i and ϕ_j. This defines the mass matrix for this set of functions, and a similar formulation defines the stiffness matrix.

For the set of functions $\{\phi_i\}$ to be used in this FEM formulation, it must be a basis for the space of piecewise linear functions on Ω. The set $\{\phi_i\}$ spans the set of all piecewise linear functions in Ω, since any piecewise linear function g_h can be written

$$g_h(x) = \sum_i g_h(x_i)\phi_i(x).$$

The set is also linearly independent by definition, since each ϕ_i is the only function in the set with non-zero value at x_i. Thus the set $\{\phi_i\}$ forms a basis for the piecewise linear function space of Ω.

This basis is used to generate the stiffness and mass matrices as described above. These matrices can be used in equation (4.24), which approximates the values of v on a spatial discretisation of Ω.

The spatial discretisation used for excitation simulation in the *in silico* reconstruction was taken to be the discrete regions which form the volume, as described in Chapter 2. This completes the process for spatial discretisation, what remains is to formulate the time discretisation.

4.3.3 Time discretisation

This section details the time discretisation of equation (4.24), and the application of the time and space discretisation to equation (4.18), to obtain the final set of linear equations used to approximate the original system. In this section the function $v_h(t)$ will be used to denote $v_h(x,t)$ for ease of notation.

Let the sequence $\{t_m\}$ denote the time discretisation for the system of equations (4.17) and (4.18), with

$$t_{m+1} - t_m = \Delta t$$

$$t_0 = 0.$$

The time derivative $\partial u/\partial t$ can be approximated by

$$\frac{\partial u(t)}{\partial t} \approx \frac{u(t+\Delta t) - u(t)}{\Delta t}.$$

Applying this approximation to the left-hand side of equation (4.24):

$$\begin{aligned}\left(\frac{\partial v_h(t_m)}{\partial t}, \phi_j\right) &\approx \left(\frac{v_h(t_{m+1}) - v_h(t_m)}{\Delta t}, \phi_j\right) \\ &= \frac{1}{\Delta t}\left((v_h(t_{m+1}), \phi_j) - (v_h(t_m), \phi_j)\right) \\ &= \frac{1}{\Delta t}\left(\sum_i (\phi_i, \phi_j) V_i(t_{m+1}) - \sum_i (\phi_i, \phi_j) V_i(t_m)\right) \\ &= \frac{1}{\Delta t}\left([MV(t_{m+1})]_j - [MV(t_m)]_j\right)\end{aligned} \tag{4.26}$$

For ease of notation, the vector $V(t_m)$ will be denoted by V^m. Substituting this approximation into equation (4.24) and rearranging gives

$$MV^{m+1} = MV^m + \Delta t(KV^m + MF^m). \tag{4.27}$$

This system of linear equations can be used to approximate the solution to equation (4.17).

The final step is to apply the time and space discretisation to equation (4.18). For a given time step t_m, define the vector W^m such that

$$W_i^m = w(x_i, t_m).$$

The discretised version of equation (4.18) is defined to be

$$\frac{W^{m+1} - W^m}{\Delta t} = \epsilon(V^m - W^m - v_0).$$

Rearranging this gives

$$W^{m+1} = W^m + \epsilon\Delta t(V^m - W^m - v_0).$$

In conclusion, the solutions of equations (4.17) and (4.18) are approximated by the solu-

tions of the following system of linear equations:

$$MV^{m+1} = MV^m + \Delta t(KV^m + MF^m) \tag{4.28}$$

$$W^{m+1} = W^m + \epsilon\Delta t(V^m - W^m - v_0). \tag{4.29}$$

This completes the derivation of the linear system of equations approximating the reaction-diffusion equations. This system is characterised by sparse mass and stiffness matrices, which will be utilised in the following section to create an efficient algorithm for solving this system.

4.4 Parallelisation and computation

The computation of the solution of equations (4.28) and (4.29), for given initial conditions, will now be outlined. Equation (4.29) gives W^{m+1} in terms of elements from the previous time step t_m. This means that the computation of W^{m+1} is straightforward given the values of V^m and W^m. On the other hand, finding the value of V^{m+1} in equation (4.28), given the values of V^m and F^m, is more involved. The approximation of this solution is found via a parallelised implementation of the Gauss-Seidel method. This is performed by means of general-purpose computing on graphics processing units (GPGPU). In Section 4.4.1, the Gauss-Seidel method and its application to this problem is described. Section 4.4.2 details the methods used to enable this computation. Section 4.4.3 describes the GPU architecture and how this supports the rapid parallel computation of the solution.

4.4.1 The Gauss-Seidel method

The Gauss-Seidel method is a technique for iteratively computing the solution to a set of linear equations [20]. This method uses an iterative sequence to approximate the value an unknown vector x of length N with

$$Ax = b$$

for some $N \times N$ matrix A and vector b of length N. Define L to me the lower triangular matrix of A and U to be the strictly upper triangular matrix such that $L + U = A$. On these definitions the system of linear equations can be rewritten

$$Lx = b - Ux.$$

For a given x^0, the iterative sequence $\{x^k\}$ is defined by

$$Lx^{k+1} = b - Ux^k.$$

If this sequence converges, then it will converge to x, since if $y = \lim_{k\to\infty} x^k$, then

$$Ay = Ly + Uy = L \lim_{k\to\infty} x^{k+1} + U \lim_{k\to\infty} x^k = \lim_{k\to\infty} (Lx^{k+1} + Ux^k) = b = Ax.$$

This sequence converges if A is positive-definite [20]. The sequence can converge if A is not positive-definite; however, this condition is sufficient for the purpose of solving the present problem.

The next step is to show how this sequence can be computed. For a given index i,

$$\sum_{j=0}^{N-1} L_{ij} x_j^{k+1} = b_i - \sum_{j=0}^{N-1} U_{ij} x_j^k. \tag{4.30}$$

By the definition of L and U,

$$L_{ij} = \begin{cases} A_{ij} & \text{if } i \geq j \\ 0 & \text{if } i < j \end{cases}$$

$$U_{ij} = \begin{cases} A_{ij} & \text{if } i < j \\ 0 & \text{if } i \geq j \end{cases}$$

Therefore equation (4.30) becomes

$$\sum_{j=0}^{i} A_{ij} x_j^{k+1} = b_i - \sum_{j=i+1}^{N-1} A_{ij} x_j^k.$$

Rearranging gives:

$$x_i^{k+1} = \frac{1}{A_{ii}} \left(b_i - \sum_{j=0}^{i-1} A_{ij} x_j^{k+1} - \sum_{j=i+1}^{N-1} A_{ij} x_j^k \right). \tag{4.31}$$

Thus each value x_i^{k+1} can be computed using the values of x^k and all values in $\{x_j^{k+1}\}_{j=0}^{i-1}$ with $A_{ij} \neq 0$. It follows that x^{k+1} can be computed sequentially from x_0^{k+1}.

To apply this technique to equation (4.28), for a given time step t_m, let

$$b = MV^m + \Delta t(KV^m + MF^m)\ A = M\ x^0 = V^m\ x = V^{m+1}.$$

The mass matrix M defined as $M_{ij} = \int \phi_i \phi_j$, as outlined in Section 4.3, is positive definite. Thus the sequence $\{V^{m+1(k)}\}$ defined by the equation

$$V_i^{m+1(k+1)} = \frac{1}{M_{ii}} \left([MV^m + \Delta t(KV^m + MF^m)]_i - \sum_{j=0}^{i-1} M_{ij} V_j^{m+1(k+1)} - \sum_{j=i+1}^{N-1} M_{ij} V_j^{m+1(k)} \right) \tag{4.32}$$

converges to V^{m+1}. This equation can be iterated over k at each time step to obtain an approximation for V^{m+1}.

4.4.2 Parallelisation

The methods for enabling parallel computation of the Gauss-Seidel iteration process will presently be described in detail. The entries of $V^{m+1(k+1)}$ as defined in equation (4.32) cannot be computed in parallel, because the sum defining $V_i^{m+1(k+1)}$ includes other entries of $V^{m+1(k+1)}$, and therefore would incur errors due to concurrent modification and use of elements. However, the matrix M is sparse, and so most of the coefficients of the

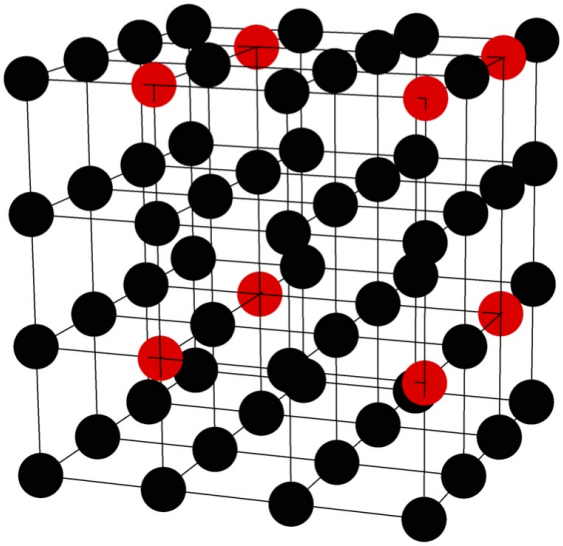

Figure 4.8: A three-dimensional partition P_n. Here the points in red denote points with index contained in P_n, given by a double-spaced sub-lattice of Ω_h. This sub-lattice is constructed in such a manner that the supports of all ϕ_i with $i \in P_n$ do not overlap.

elements $V_j^{m+1(k+1)}$ in equation 4.32 are zero. Utilising this fact enables the parallelisation of large sections of the computation.

Recall that the volume Ω is discretised by the lattice $\Omega_h = \{x_i\}$ with internodal distance h. For any two points $x_i = (x_{i,0}, x_{i,1}, x_{i,2})$ and $x_j = (x_{j,0}, x_{j,1}, x_{j,2})$ in the lattice with

$$|x_{i,k} - x_{j,k}| > h$$

for $k \in \{0, 1, 2\}$, the supports of ϕ_i and ϕ_j do not overlap, so

$$M_{ij} = \int \phi_i \phi_j = 0.$$

This result can be used in parallelisation as follows. Define the partitioning $P = \{P_n\}_{n=0}^{7}$ of the set of indices $\{i\}$, such that each P_n defines a sub-lattice $\{x_i\}_{i \in P_n}$ with internodal distance $2h$, as illustrated in Figure 4.8. This construction ensures that

$$|x_{i,k} - x_{j,k}| > h$$

for all $i, j \in P_n$. Suppose that the lattice $\Omega_h = \{x_i\}$ is re-indexed so that for each n

$$P_n = \{i \mid ln \leq i < l(n+1)\},$$

where l is defined to be the length of a partition. Strictly speaking, the partitions need not have equal length, but this is assumed here for ease of notation. Under this new indexing, equation (4.32) can be rewritten in the following form:

$$\begin{aligned} V_i^{m+1(k+1)} &= \tfrac{1}{M_{ii}} \left([MV^m + \Delta t(KV^m + MF^m)]_i \right. \\ &\quad \left. - \textstyle\sum_{j=0}^{ln-1} M_{ij} V_j^{m+1(k+1)} - \sum_{j=l(n+1)}^{N-1} M_{ij} V_j^{m+1(k)} \right) \end{aligned} \tag{4.33}$$

where $i \in P_n$. This means that the value of $V_i^{m+1(k+1)}$ can be calculated independently of all other values $V_j^{m+1(k+1)}$ with $j \in P_n$. Thus equation (4.33) can be computed in parallel for all values of $V^{m+1(k+1)}$ with indices within a given partition. The implementation of this parallel algorithm is performed using GPGPU, as outlined in the following section.

4.4.3 Implementation using GPGPU

The parallelised equations given above can be processed on a graphics processing unit (GPU). To appreciate how this can be achieved, consider the architecture of a GPU, which is illustrated in Figure 4.9. This arrangement of multiple small processing elements enables the GPU to perform many small calculations in parallel. This property can be exploited to perform the iterative calculations in equation (4.33) in parallel for each partition. The software developed to solve these equations is based on the Aparai package, which enables the GPU architecture to be accessed through Java code.

The memory required for storing the full stiffness and mass matrices far exceeds the memory available on a GPU, because for a spatial discretisation with n nodes each matrix has n^2 entries. Given current GPU technology, this bounds the number of nodes to $\sim 10^4$, which is ~ 100 times smaller than the number of nodes present in the spatial discretisation.

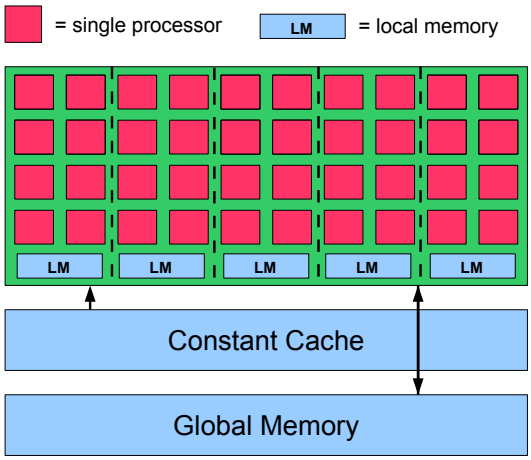

Figure 4.9: A simplified version of the general structure of the GPU. The GPU has several small processing units, each capable of performing small calculations. The overall effect is that several small calculations can be performed in parallel by the GPU. The processing units are split into blocks, with each block having a shared local memory which is used by the processors while performing calculations. The global memory is a large cache used to store the arrays V^*, W, and C. The constant cache is a large memory cache which is physically closer to the processing units than the global memory cache, allowing the processors more rapid access to the constant memory. However, data stored in the constant cache cannot be modified by communication with the GPU processors, and therefore is not used to store any of the variable arrays. This cache is used to store the mass, stiffness, and neighbour arrays because these remain constant throughout the simulation.

These matrices could feasibly be stored outside the GPU, and only loaded when required; however, the constant transfer of data on and off the device would considerably reduce the speed of the calculations. Therefore it is necessary to reduce the space required for these matrices.

Reduction of the memory required for the mass and stiffness matrices, is possible as both these matrices are sparse. Furthermore, for both matrices each row has an upper bound on the number of non-zero entries. This bound is dependent on the finite element approximation and the dimension of the system. For a three-dimensional system using the methods outlined in Section 4.3 this bound is 15. For row i of M, consider the set of nodes N_i such that

$$M_{ij} \neq 0$$

for all $j \in N_i$. By definition, if $M_{ij} = 0$ then $K_{ij} = 0$, so N_i is the set of all non-zero indices of the ith row in both matrices. Therefore the arrays $(M_{ij})_{j\in N_i}$ and $(K_{ij})_{j\in N_i}$, combined with the array N_i accurately represent the complete rows of both matrices. This new representation can be used for each row of M and K to produce three matrices:

$$N = (N_i)_{i=0}^{n-1}, \quad \tilde{M} = (M_{ij})_{j\in N_i, i=0}^{n-1}, \quad \tilde{K} = (K_{ij})_{j\in N_i, i=0}^{n-1}.$$

These matrices each have size $15 \times n$, therefore the memory required to store these matrices has been reduced from $O(n^2)$ to $O(n)$.

For $i \in P_n$, equation (4.33) can be written as

$$V_i^{m+1(k+1)} = [C(m) - M^* V^{m+1(k+1)*}]_i, \tag{4.34}$$

where

$$C(m) = \frac{1}{M_{ii}}(MV^m + \Delta t(KV^m + MF^m)); \; M^* = \frac{1}{M_{ii}}M$$

and the array $V^{m+1(k+1)*}$ has entries given by

$$V_j^{m+1(k+1)*} = \begin{cases} V_j^{m+1(k+1)} & \text{if } j \in P_l, l \leq n \\ V_j^{m+1(k)} & \text{if } j \in P_l, l > n \end{cases}.$$

A minor technicality is that $V_j^{m+1(k+1)*}$ is not well-defined for $j \in P_n$, since the value is being computed at the same time as $V_i^{m+1(k+1)}$. This does not present a problem, however, because P_n is constructed such that $M_{ij} = 0$ for $i, j \in P_n$. In this new formulation, the value of $C(m)$ can be computed prior to the iteration. In addition, the values of $V^{m+1(k+1)*}$ are precisely the last computed iterations of V^{m+1}. This means that all iterations can be stored in a single array V^{m+1*}.

With this new formulation, each increment in time from t_m to t_{m+1} is computed in the following manner:

κ_l	time to simulate 100 ms (minutes)
1	1
9	7
20	19
50	46
100	91
200	229
400	537
500	676

Table 4.1: Approximate computation times for simulation of 100 ms in rat block 3 for a range of longitudinal conductivity constant applied. Increasing the of κ_l increased the propagation speed of the excitation wave, which required smaller time steps to accurately compute this more rapidly changing system.

(i) Compute the value of $C(m)$

(ii) Set $V^{m+1*} := V^m$

(iii) For each partition P_n, compute the values of the next iteration of V_i^{m+1*} for $i \in P_n$

(iv) Repeat (iii) for the required number of iterations

(v) Compute the value of W^{m+1}.

4.5 Simulating excitation in the reconstructed tissue

The techniques described in this chapter provide a means to solve the electrodiffusion equations given the anisotropy of the *in silico* reconstructions described in Chapter 2. Simulations reported in this thesis were restricted to the third rat tissue block. This reconstruction is composed of voxels of side 47.5 μm, and this spatial discretisation was used as the basis for the simulations, as described in Section 4.3.1. The time step described in Section 4.3.3 was varied depending on the conductivity constants; higher conductivity leads to a faster propagation of action potentials through the tissue, which means that a finer time discretisation is required to represent this more rapidly changing system. The

time step Δt in the dimensionless time units used above was given by

$$\Delta t = 0.04/\kappa_l,$$

where κ_l is the longitudinal conductivity constant. One final value to be fixed is the maximum number of iterations to apply in the Gauss-Seidel approximation presented in the previous section. This value was fixed at 16 for all simulations based on observations of the behaviour of the system when varying this quantity.

Computation times for these simulations are given in Table 4.1. These times were increased for higher values of the longitudinal conductivity constant κ_l because the time step was increased, and therefore more calculations were performed for the same time increment. The results of these simulations are described in detail in Chapter 9.

Chapter 5

Analysis of multi-electrode array recordings

This chapter presents the methods used to process multi-electrode array recordings in order to obtain parameters for comparison with the simulated electrodiffusion discussed in Chapter 4. This processing consists of two steps. The first step is to measure bursting behaviour at each electrode by analysing the electrogram recorded by the given electrode. The second step is to analyse the bursting behaviour across the electrode array in order to determine initiation points and propagation speeds. Together, these two steps enable the estimation of speed and determination of initiation and termination point parameters for a given multi-electrode array recording. All code for the methods described in the chapter can be found at `https://figshare.com/s/651e592a91f309182edd`.

5.1 Processing multi-electrode array electrograms

This section details the methods used to determine characteristics of the electrical activity in the tissue from the recordings at each electrode. The methods used to obtain the electrode array recordings are described in Chapter 2. For each electrode, the activity was recorded as a time sequence of voltage values, which is referred to in the following as an electrogram. The time difference between subsequent values was 1 ms. Accordingly, the following shall assume that time is integer-valued, with each time-step equal to 1 ms.

An action potential is characterised in the electrogram by a rapid deflection in the

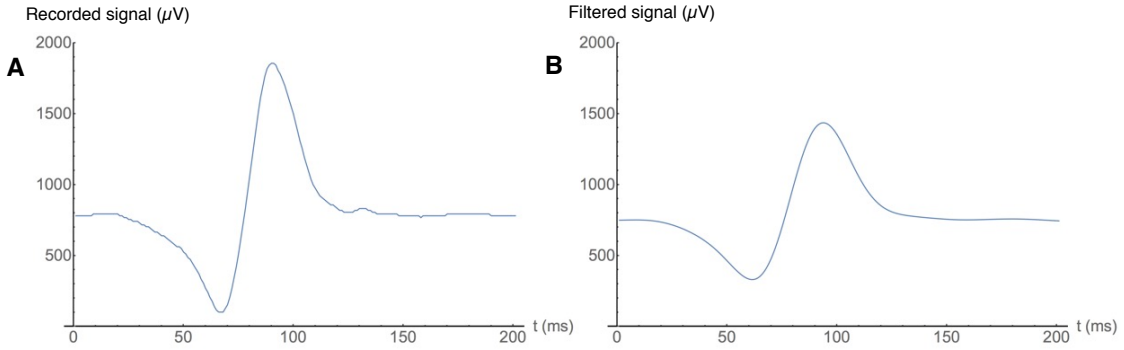

Figure 5.1: An example of an action potential in an electrogram, both before (**A**) and after (**B**) smoothing.

negative direction immediately followed by a rapid deflection in the positive direction, creating a trough-peak pair, as illustrated in Figure 5.1**A** [61]. For this reason, it was deemed appropriate to use the difference between a trough and the subsequent peak to determine at which point in time an action potential occurs. These differences were then subjected to two thresholding procedures to obtain the timing of each action potential. These action potentials were grouped together to determine the location in time of activity bursts for each electrogram.

5.1.1 Obtaining trough-peak differences

In order to reduce noise, a Gaussian filter was applied to the raw data [20]. A Gaussian filter M with variance σ^2 and radius $w_{1/2}$ has entries

$$M(i + w_{1/2}) := \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-i^2}{2\sigma^2}\right\}.$$

The filter used to smooth the data had standard deviation $\sigma = 10$ ms and radius $w_{1/2} = 20$ ms. These values were selected so as to smooth the electrogram locally while preserving the general shape of the curves. An example of this smoothing is shown in Figure 5.1**B**. The smoothing can displace the troughs and peaks slightly, but this movement is invariably within the bounds of the filter radius, which is sufficiently small to be contained within

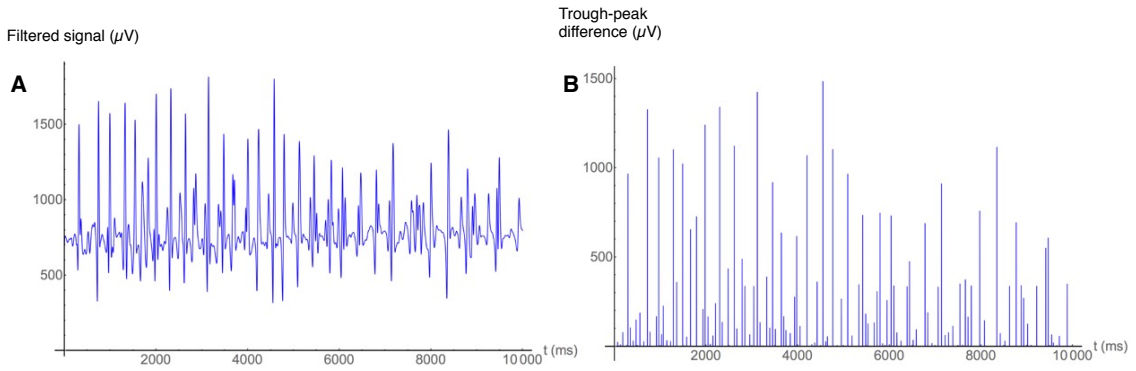

Figure 5.2: The trough-peak pairing used to determine action potentials. **A**: A 10-second interval of an electrogram displaying bursting behaviour. **B**: The differences between troughs and peaks, with each value placed at the time of the trough.

the width of a peak or trough representing an action potential. Applying this filter to the raw data yields the smoothed values $g(t)$.

The peaks and troughs of a smoothed electrogram were located by numerically approximating the first derivative. In the present instance, a seven-point derivative formula was used [20]. The application of this filter to $g(t)$ produces the estimated derivative $d(t)$. A time t with derivative $d(t)$ was defined to be a turning point if

$$d(t)d(t-1) \leq 0.$$

Each turning point was categorised as being either a peak or a trough. To determine whether the turning point was a peak or trough, the second derivative was used. The second derivative was approximated by applying the derivative filter a second time. While it is possible to use a second derivative filter to obtain a more precise estimation of this value [20], a high level of precision is not necessary here and therefore the second application of the first derivative filter was used for convenience. Each turning point was categorised as being a trough if the second derivative was less than zero, otherwise it was categorised as being a peak. This categorisation includes points of inflection in the peaks. This distinction is of no consequence, because points of inflection are not relevant

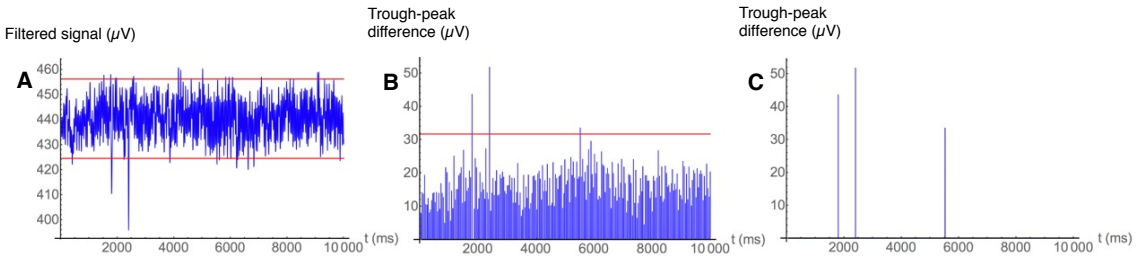

Figure 5.3: The effect of applying the baseline threshold to an electrogram with no bursting behaviour. **A**: A 10 second interval of an electrogram with no bursting behaviour, with the first and fourth quartiles marked in red. **B**: The trough-peak pairing, with the threshold equal to twice the interquartile range marked in red. **C**: The result of applying the threshold to **B**. While a few peaks remain, they are not sufficient in number to constitute bursting behaviour, and will therefore be discarded when the bursting behaviour is analysed.

to action potentials and therefore are generally only present at points which are not part of an action potential.

Each action potential is characterised by a trough followed by a peak. Accordingly, each trough was paired with the next peak to be detected in the time sequence to obtain all possible times at which an action potential may have occurred. A sequence $\{p_i = p(t_i)\}$ of amplitudes was thus generated, where each t_i is the time at which a trough occurred, and

$$p(t_i) := g(s_i) - g(t_i)$$

where s_i is the time of the peak paired to the trough at t_i. The effect of this pairing is illustrated in Figure 5.2. This sequence was thresholded to obtain the action potential times, as described in the following section.

5.1.2 Thresholding

Two forms of thresholding were used to determine the action potential times. The first was a broad threshold applied to the entire sequence $\{p_i\}$, enabling low-level noise to be filtered out of the data. The second procedure was a localised form of Otsu thresholding [88], which gives a more refined set of threshold values to identify the action potentials.

The threshold applied to the whole sequence was set at three times the interquartile

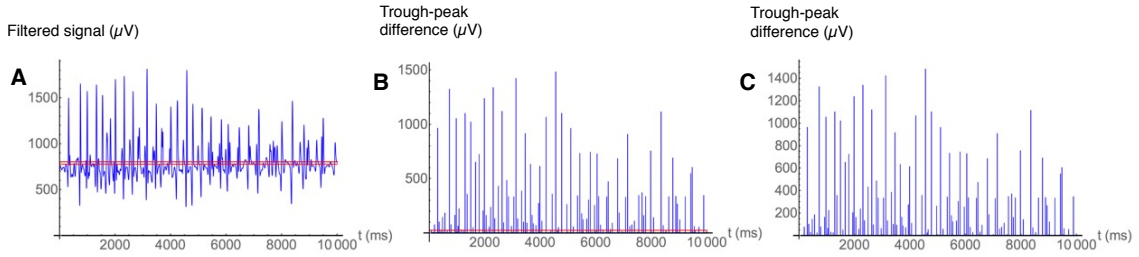

Figure 5.4: The effect of applying the baseline threshold to an electrogram displaying bursting behaviour. **A**: A 10 second interval of and electrogram with no bursting behaviour, with the first an fourth quartiles marked in red. **B**: The trough-peak pairing, with the threshold equal to three times the interquartile range marked in red. **C**: The result of applying the threshold to **B**. This threshold has little effect in this case; therefore a second thresholding procedure is required.

range of the smoothed data $g(t)$. The interquartile range was used because it represents the spread of the data and it may generally be expected that it is not greatly affected by the extreme values. This means that a threshold using the interquartile range will even detect values which are outside the normal spread of values. The factor of 3 was used because the sequence $\{p_i\}$ is a set of differences between extreme values of $g(t)$, so the threshold used must represent the full range of non-extreme values. The effect of this thresholding procedure on electrograms with no bursting is illustrated in Figure 5.3. The number of peak-trough pairs is greatly reduced, which, when combined with burst detection below, means that these electrograms will be recorded as having no bursting behaviour. The effect of the thresholding procedure on bursting electrograms is illustrated in Figure 5.4. While some low-level noise has been eliminated, a large proportion of the peak-trough pairs remaining do not represent action potentials. In order to remove these pairs, a localised form of Otsu thresholding was used.

Otsu thresholding is a general method of determining a threshold which separates foreground from background in image processing [88]. This technique is equally applicable to signal processing, as it does not rely on spatial data to determine the threshold. The aim of Otsu thresholding is to separate the values of p_i into two sets such that the weighted average of the variance of the sets is minimised. For a threshold value m, the sets $S_0(m)$

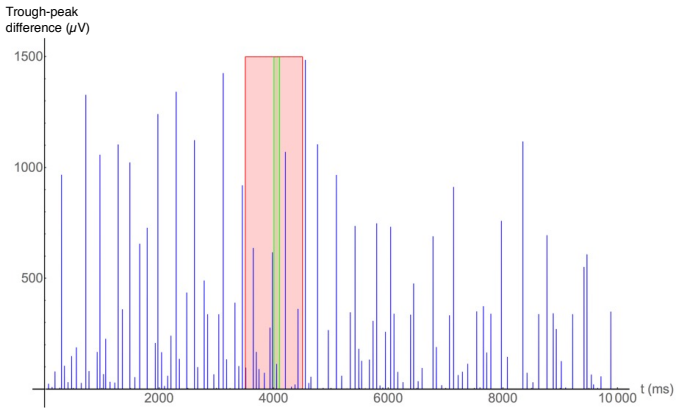

Figure 5.5: An example of a time band used for Otsu thresholding. The time band marked in green is subject to the Otsu threshold based on values in the extended radius marked in red.

and $S_1(m)$ are given by

$$S_0 = \{p_i \mid p_i < m\}$$

$$S_1 = \{p_i \mid p_i > m\}.$$

The intra-class variance $\sigma_w^2(m)$ is defined to be

$$\sigma_w^2(m) := w_0(m)\sigma_0^2(m) + w_1(m)\sigma_1^2(m),$$

where $\sigma_i^2(m)$ is the variance of $S_i(m)$ and $w_i(m)$ is the proportion of points belonging to $S_i(m)$, given by

$$w_i(m) := \frac{|S_i(m)|}{|S_0(m)| + |S_1(m)|} .$$

The value of m which minimises $\sigma_w^2(m)$ is the point at which the weighted average of the distribution of points in the sets $S_i(m)$ is smallest. This minimising value of m therefore represents the most likely boundary between two separate distributions of points. Minimising the intra-class variance is equivalent to maximising the inter-class variance $\sigma_b^2(m)$ [88], given by

$$\sigma_b^2 = w_0 w_1(\mu_1(m) - \mu_0(m))^2$$

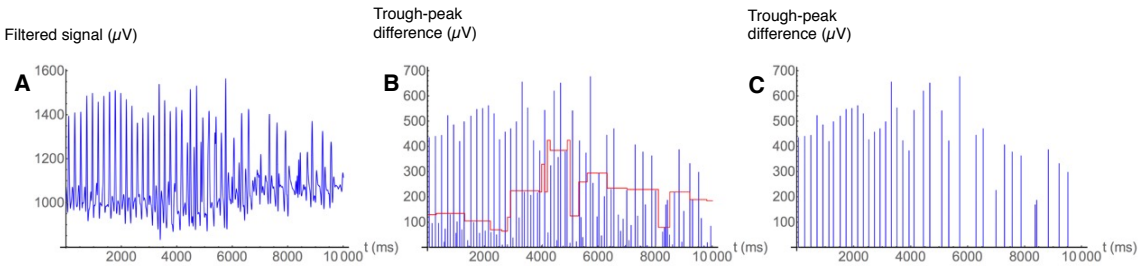

Figure 5.6: The effect of applying the Otsu threshold to an electrogram displaying bursting behaviour. **A**: A 10 second interval of an electrogram with no bursting behaviour. **B**: The trough-peak pairing, with the Otsu threshold values marked in red. **C**: The result of applying the threshold to **B**. This threshold adapts to the local magnitudes, which enables the removal of lower magnitudes while retaining the magnitudes representing action potentials.

where $\mu_i(m)$ is the mean of $S_i(m)$. This equivalence is due to the following relation:

$$\sigma_b^2(m) + \sigma_w^2(m) = \sigma^2$$

where σ^2 is the variance of the set $\{p_i\}$. The optimal threshold is the value m which minimises σ_b, and this threshold is selected to be the threshold applied to the set S.

While this procedure is capable of finding the optimal threshold for the general separation of foreground and background, the problem at hand requires a modified version. The values of $\{p_i\}$ corresponding to action potentials vary during the burst, as can be seen in Figure 5.2. In particular, action potentials in one region of the electrogram are not necessarily greater in magnitude than background levels in other regions. For this reason the threshold was set locally, which was achieved in the following way. The time variable was divided into bands of 100 ms, with the threshold calculated separately for each band. For each time band, the set of points used to determine the Otsu threshold was given by points p_i such that t_i was within 400 ms from the time band, as illustrated in Figure 5.5. This localisation enabled the threshold to be precisely tuned for a short time period, while the extended radius sufficiently increased the sample size to give an accurate threshold value. Using this extended set of values, the threshold in each time band was determined using the Otsu thresholding procedure detailed above.

Parameter	Value
$F_{\min}$	2 Hz
$D_{\max}$	1 ms
$T_{\min}$	10 s

Table 5.1: Burst detection parameter values.

The modified Otsu thresholding procedure produced a set of localised thresholds as illustrated in Figure 5.6. This localised thresholding procedure values yielded an accurate set of action potential times for each electrogram. These times were used to determine the bursting behaviour of the system as described in the following section.

5.1.3 Determining bursting behaviour

The action potentials detected by the electrode array formed bursts, as illustrated in Figure 5.7**A** [74]. These bursts were characterised by a set of action potentials occurring in rapid succession. Accordingly, to determine the bursts from the action potential times a minimum frequency, a maximum time difference, and a minimum burst length were selected. For each electrode with an associated sequence of action potential times $\{t_i\}_{i=0}^{n}$, a burst was defined to be any subsequence S of $\{t_j\}_{j=i_0}^{i_1}$ such that

$$(t_{i_1} - t_{i_0}) \geq T_{\min}, \frac{|S|}{(t_{i_1} - t_{i_0})} \geq F_{\min}, \min_{i_0 \leq i < i_1} (t_{i+1} - t_i) \leq D_{\max},$$

where $T_{\min}$ is the minimum burst length, $F_{\min}$ is the minimum burst frequency, and $D_{\max}$ is the maximum time difference allowed between two action potentials within a burst. These values are listed in Table 5.1. Applying these criteria to each electrogram yielded the burst times and lengths, as shown in Figure 5.7. This automated analysis enabled the processing of the multi-electrode array data obtained from rat uteri. In the next section the behaviour of the individual electrodes is analysed as part of the whole array in order to determine how the excitation wave propagates through the tissue.

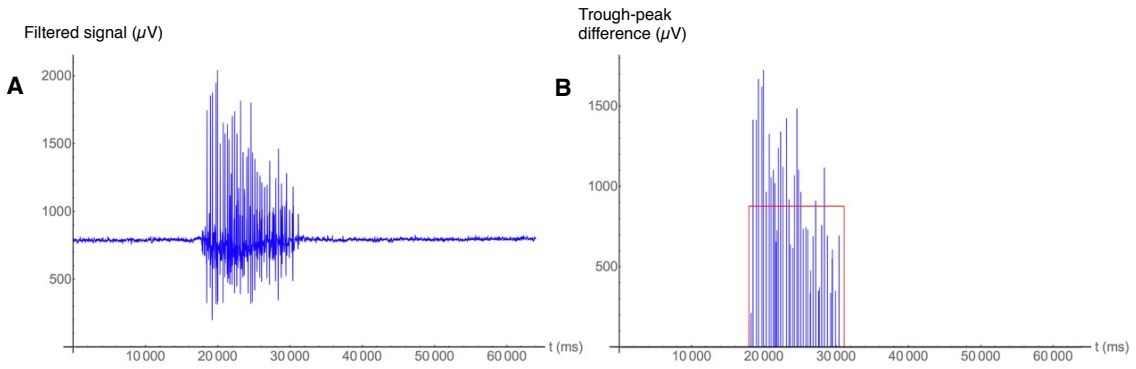

Figure 5.7: An example of burst detection. The electrogram **A**, after being subjected to the above processing, yields the set of action potentials in **B** (blue), which form a burst (red).

5.2 Determining propagation behaviour

The bursting activity determined using the above process yielded a time sequence of action potentials for each electrode on the multi-electrode array. The next step is to use these individual bursting data in order to determine how the signal propagates through the tissue. The main challenge with analysing propagation is determining which electrodes have recorded bursts from the same propagating wave. To solve this problem, all pairs of electrodes were compared to determine if the bursting behaviour of the given pair is similar enough to be part of the same wave. This pairing of electrodes created a graph of matched electrodes, with each connected component of the graph representing a group of electrodes with a common origin. Once electrodes were grouped in this manner, it was then possible to determine the point of origin and the propagation speeds for each burst.

The measure of similarity between two electrodes was determined with the assumption that if two bursts are part of the same propagating wave then the bursting behaviour in one electrode will be exhibited at a later time point in the other electrode. This similarity may be due to a direct propagation from one electrode to another or from a common origin where the wave has reached one point before the other. The similarity measure specifically compared the sequence of action potential times in the respective electrograms. While these times may vary as the wave propagates, due to variations in conductivity, the aim is

to group electrodes such that each group is definitely part of the same wave, and therefore speeds determined from this grouping will be accurate.

The electrogram of an electrode was compared to that of a reference electrode by comparing time intervals of 2000 ms. This time interval was selected to include a sufficient number of action potentials for comparison, while remaining short to minimise the effect of time-dependent changes in conductance on the comparison. The idea is to slide the action potentials detected in a given time interval backwards in time to compare with the action potentials detected by the reference electrode. The range of times which the interval can be projected back was restricted to the range $[50d, 1000d]$ ms, where d is the distance between the electrodes in cm. This range of times represents a propagation speed within the range $[1, 20]$ cm s^{-1}, which has been selected based on previous observations [91]. While technically these projections are not necessarily restricted to propagation between electrodes, restricting the projections to ones which could potentially be propagations improves the robustness of the matching. In addition, the projection speeds will be determined later based on the matchings produced in this step.

The time intervals were projected back in time by steps of 10 ms. At each time step, the projected action potentials were marked as matched if the reference electrogram contained an action potential within 10 ms of the given action potential. The action potentials in the reference electrogram were marked as matched in the same way. If the proportion of matched action potentials within the projected time interval in both electrograms was above a fixed threshold, the interval was marked as matched. Additionally, the interval was assigned the speed corresponding to the given time step. If more than one speed presented a match, then the speed with the largest proportion of the matched action potentials was taken to be the given speed. The speeds found in this manner may not necessarily represent actual propagation, but the following step presents a method for reducing this set to genuine propagation speeds.

The above process of matching electrodes produced a graph of the electrodes, where

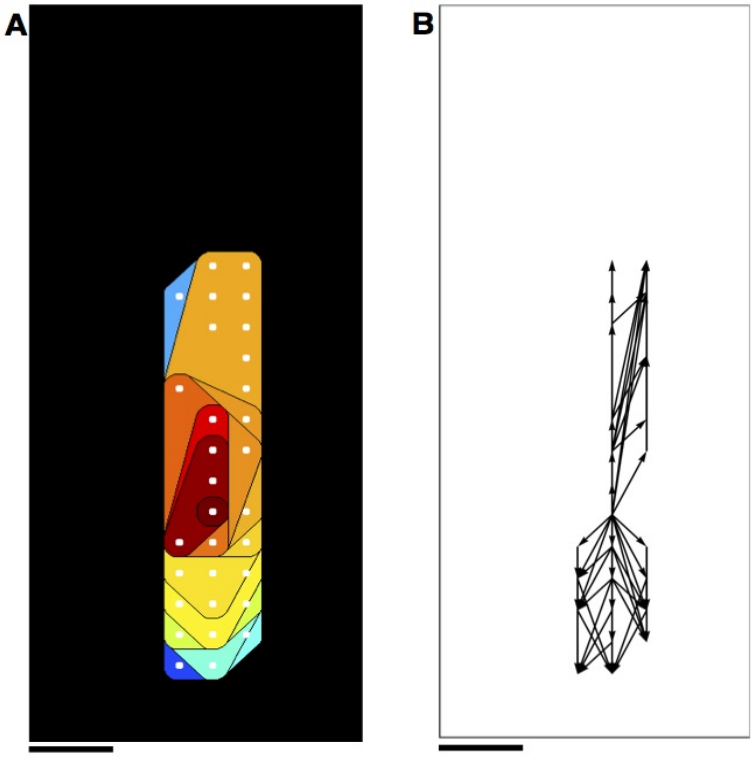

Figure 5.8: **A**: An example of an isochrone map generated. The isochrones are coloured according to time time at which the burst is first detected by the electrode, with red corresponding to earliest and blue corresponding to latest. **B**: The set of all propagation directions allowed based on the isochrones in **A**. Propagation is allowed down the colour gradient, and only in a direction within 45° of the direction from the centre of the initiation points.

each arc of the graph represents a matched pair of electrodes. The connected sub-graphs of this grouping with more than two nodes were taken to be points which exhibit bursting behaviour as part of a common excitation wave. This segmentation was used to generate isochrone maps, as illustrated in Figure 5.8**A**. The actual propagation speeds were now selected utilising these isochrone maps. Previously, all possible propagation speeds were recorded as part of the electrode matching algorithm. For each connected set of electrodes, these propagation speeds were reduced to the set of propagations which approximate the direction outward from the source electrode given in the isochrone image, as illustrated in Figure 5.8.**B** This reduction ensures that all recorded speeds are genuinely propagating from the relevant source electrode.

In order to put these calculated speeds in the context of fibre direction, the structural representation determined in Chapter 2 was utilised. The average fibre direction of the tissue at the surface around each electrode was calculated. This enables the categorisation of each propagation vector by the approximate angle that the vector makes with the fibre direction. This categorisation of the speeds allows comparison of speeds following a similar path in terms of fibre direction.

The methods detailed in this chapter provide the tools for automated analysis of multi-electrode array recordings. The processing of each electrogram provides a representation of the times at which action potentials occur, and whether the electrogram exhibits bursting behaviour. The analysis of the array as a whole enables the generation of isochrone maps, which show initiation points of propagating bursts, along with locations where propagation is blocked. Additionally a sample of the propagation speeds of each excitation wave can be found. These speeds are generated as a function of time, meaning that changes in propagation speeds can also be determined.

Chapter 6

Excitation behaviour observed in multi-electrode array recordings

Excitation in the myometrium is characterised by bursts of high frequency action potentials [74]. These bursts have been shown to propagate through the myometrium [79], forming an excitation wave. The aim of this chapter is to analyse the initiation points, bounds, and propagation speeds of these excitation waves.

Excitation is initiated by spontaneous activity in myometrial smooth muscle cells [74]. This behaviour is exhibited and well-documented in cardiac and other smooth muscle tissues. In cardiac muscle, pacemaker cells centrally located in the sino-atrial node initiate electrical activity [90]. Conversely, in the bladder and parts of the gastro-intestinal tract, pacemaker activity is more widely distributed [53,66]. Recent work by Lammers *et al.* [67] has shown that, in rat uteri, excitation waves are initiated at multiple locations. These locations tend to be adjacent to the mesometrial border, suggesting that pacemaker cells are more predominantly situated in that location. In this chapter, multi-electrode array data is compared to the *in silico* tissue reconstruction described in Chapter 2 and to the histological images used to generate this reconstruction to determine if there are any physiological structures present at the initiation points.

It has been shown that the excitation waves in rat myometrium are frequently bounded in space [65] i.e. the propagation appears to be blocked which prevents full excitation of the tissue. This bounding will also be examined by comparing multi-electrode array

recordings with histological slides and the *in silico* reconstruction to determine if there are any physiological structures common to these bounds.

The speed with which bursts propagate through the myometrium in rat tissue has previously been measured in longitudinal fibres [51,60,68,79], and these values are summarised in [91]. The propagation speeds found in these studies are in the range of 4.9–13.5 cm s^{-1} in the direction of the fibre and 2.3–4.0 cm s^{-1} perpendicular to the fibre direction. The speed has been observed to vary depending on proximity to the placenta [60], with samples taken from longitudinal fibres in the placental region having a lower propagation speed than samples from longitudinal fibres from a non-placental region. The analysis in this chapter determines the propagation velocity vectors through the longitudinal myometrium as a function of space. This means that variability of speed throughout the tissue is quantified to determine if any areas of high speed or low speed are present in each recording.

The methods used to determine electrical activity from multi-electrode array recordings are detailed in Chapter 5. Recordings were taken from three rat uteri, the first two of which were 20 days pregnant at the time of recording, while the third was 19 days pregnant. For each tissue block, three recordings were analysed, separated by approximately 10 minutes, and are numbered in the order in which they were recorded. This analysis was fully automated. For each recording, the individual electrograms were analysed to determine the temporal location of action potentials and these times were grouped to detect bursts comprised of those individual action potentials.

For each recording, the electrodes were segmented according to similarity of bursts; the bursts recorded by each electrode were compared and deemed part of the same excitation event if they exhibited similar action potential frequencies for sufficiently large portions of the recordings. In the following, each segmentation performed in this manner is assessed through visual inspection to determine if the segmentation is accurate. If there are indeed separate events, this would indicate that excitation is initiated at multiple points simultaneously. The segmentation enables the generation of isochrone maps, which are a routine

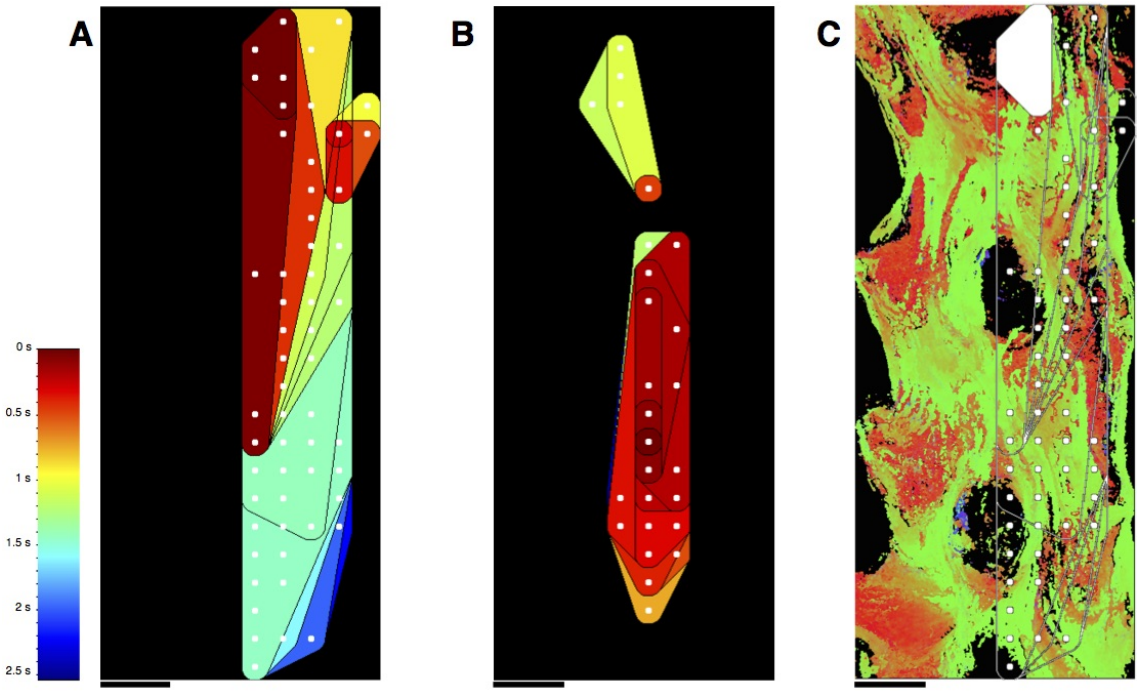

Figure 6.1: Images used to visualise Excitation propagation through the tissue. **A**: Isochrone map. This is a representation of the spread of the excitation wave. The colour represents the time at which the excitation wave reaches the given area, with dark red indicating the earliest and dark blue representing the latest times, as indicated by the colour chart. The white dots indicate that the burst was detected at the electrodes located at those points. **B**: Isochrone map for two distinct excitation waves. Here the two separate events have been separated to clearly show excitation propagation in each. **C**: Isochrone map in comparison to the top surface of the reconstructed tissue. Here the colour-coding of the isochrones has been removed, with the exception of the earliest region, which has been coloured white. The tissue colour-coding represents fibre direction, with green indicating fibres running vertical in the image plane, red indicating fibres running horizontal in the image plane, and blue indicating fibres running perpendicular to the image plane. Scale bars represent 5 mm.

visualisation method in the study of excitable tissues [53, 65, 67, 85], and which display the propagation of the excitation wave through the tissue, in particular mapping points of origin and termination of the excitation waves. An example of an isochrone map is shown in Figure 6.1**A**. These isochrone maps were generated as follows. The first action potential detected at each electrode was taken to be the time at which the excitation wave reached the given electrode. These times were separated into bands of width 200 ms, with the first band starting at the earliest detected start time of all bursts in the recording. For each of these time bands a convex hull was drawn around electrodes represented in all time

bands up to and including the given time band. The area inside that convex hull which had not been previously coloured was coloured in the appropriate colour, as represented by the colour colour chart in Figure 6.1**A**. If multiple excitation waves were recorded then each of these was processed in the above manner separately, as shown in Figure 6.1**B**. Additionally, these isochrone maps were compared to the *in silico* tissue structure, as illustrated in Figure 6.1**C**. This representation of the tissue structure is coloured according to the fibre direction of the topmost voxel of the volume at each planar coordinate, with green representing fibres running vertically in the image plane, red representing fibres running horizontally in the image plane, and blue representing fibres running perpendicular to the image plane. This comparison shows the initiation and termination points of the excitation wave in the context of the reconstructed tissue, which is used to select areas of interest in the *in silico* reconstruction and the histological slides. This enables the visual inspection of the tissue at the locations relevant to initiation and termination.

Propagation speeds were calculated according to the procedure detailed in Chapter 5. Specifically, these speeds were represented as a time sequence of speeds between pairs of electrodes. The fibre direction found in Chapter 2 was used to categorise these speeds by direction relative to the tissue; if the straight line drawn between a given pair of electrodes made an average angle of less than 45° to the fibre direction, then the propagation was deemed to be along the fibre direction, otherwise it was deemed to be perpendicular to the fibre direction. Propagation speed along the fibre direction is the quantity of interest in this analysis, as it is more directly related to conductivity, whereas propagation perpendicular to the fibre direction is also dependent on connectivity between fibres, so its relationship to conductivity is less direct. It turns out that in general the automated analysis more readily detected propagation along the fibre direction. This is because propagation was much more prevalent along the fibre direction, as is evident in the isochrone maps shown in the following sections. This had the effect that propagation speeds detected perpendicular to fibre direction were much more sensitive to miscategorisation, i.e. prop-

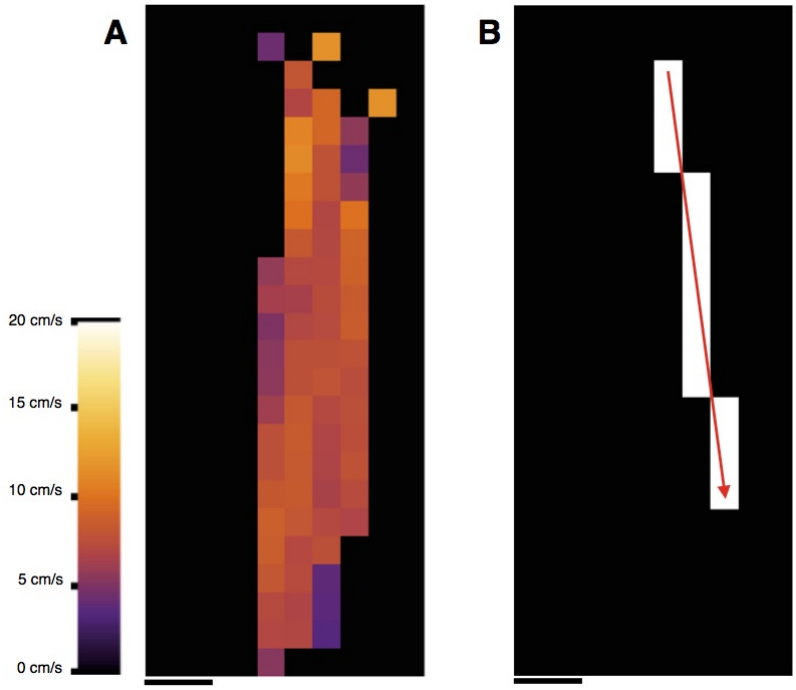

Figure 6.2: Representation of propagation speed in electrode array recordings. **A**: Spatial distribution of speeds. Here the colour represents the average propagation speed passing through the given point, with brighter colours representing higher speeds, as shown by the colour chart. **B**: Illustration how propagation speeds are assigned to pixels in **A**. The arrow indicates propagation speeds were recorded between its two endpoints, while white squares indicate points selected as the path between the two points. For each square on the path, the propagation speeds recorded between the two endpoints are combined with speeds from all other propagation paths passing through the given square to yield the average speeds represented in **A**. Scale bars represent 5 mm.

agation along the fibre direction which was categorised as being perpendicular to the fibre direction heavily influenced the distribution of speeds perpendicular to the fibre direction due to the small sample size. For this reason, all analysis of propagation speeds in the following sections will focus solely on propagation along the fibre direction. The propagation speeds along the fibre direction are represented below in the form of empirical cumulative distribution functions, generated in Mathematica. These enable the comparison of the distribution of speeds between recordings. Additionally, the speeds along the fibre direction for each recording are spatially represented by images showing the average propagation speed through regions of the tissue corresponding to electrode locations, as illustrated in Figure 6.2**A**. These images were generated in the following manner. Each pixel represents

the cross-sectional square area in the tissue centred at an electrode. These pixels were coloured according to the average propagation speed of all excitation waves deemed to be passing through the given area. For each pair of electrodes where propagation speeds were recorded, a straight line was drawn between the electrodes. The excitation wave passing between these two electrodes was deemed to be passing through each pixel which the line passed through, as illustrated in Figure 6.2**B**. While excitation propagation between electrodes is unlikely to form a straight line, this representation yields an approximation to the path and hence the visualisation of the spatial distribution represents an approximation of the true spatial distribution of the speeds. The pixels were subsequently coloured according to the colour chart shown in Figure 6.2**A**.

The results are presented here by tissue block. For each block, any segmentation applied is visually assessed, followed by analysis of the location of initiation points, areas where propagation terminates, and propagation speeds. The electrode array data used in this analysis can be found at `https://figshare.com/s/9a54e1ae6699cf713670`.

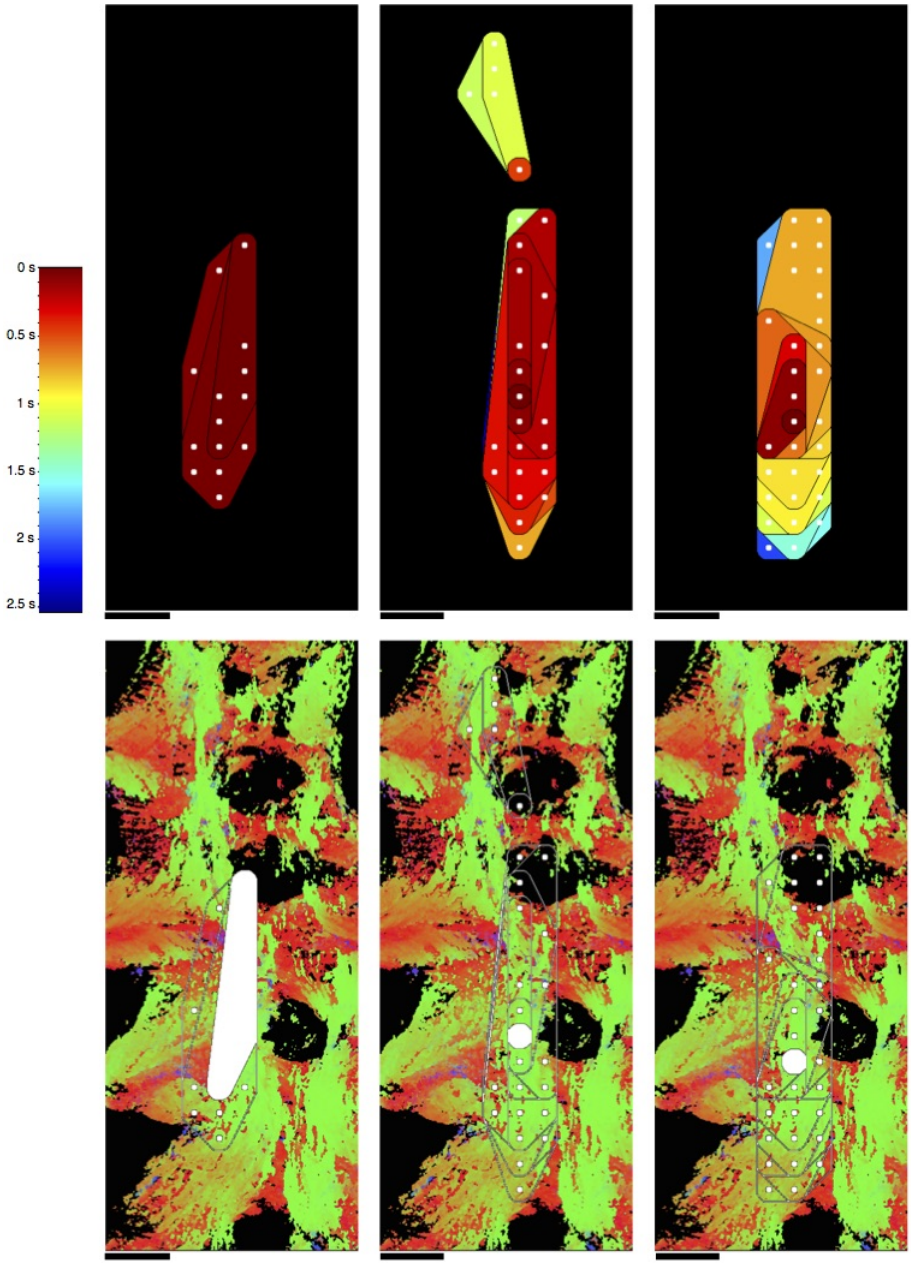

Figure 6.3: Isochrones for tissue block 1, in order of time of recording. Top row: isochrones taken at 200 ms intervals, with colour indicating time at which the excitation wave reached the electrode. White points indicate electrodes where bursts were recorded. Bottom row: comparison of isochrone initiation points and bounds with the top surface of virtual tissue. The tissue is pseudocoloured based on fibre direction, with green indicating the vertical direction in the image plane, red indicating the horizontal direction in the image plane, and blue indicating the direction perpendicular to the image plane. The ovarian end of the tissue is located at the top of each image. Scale bars represent 5 mm.

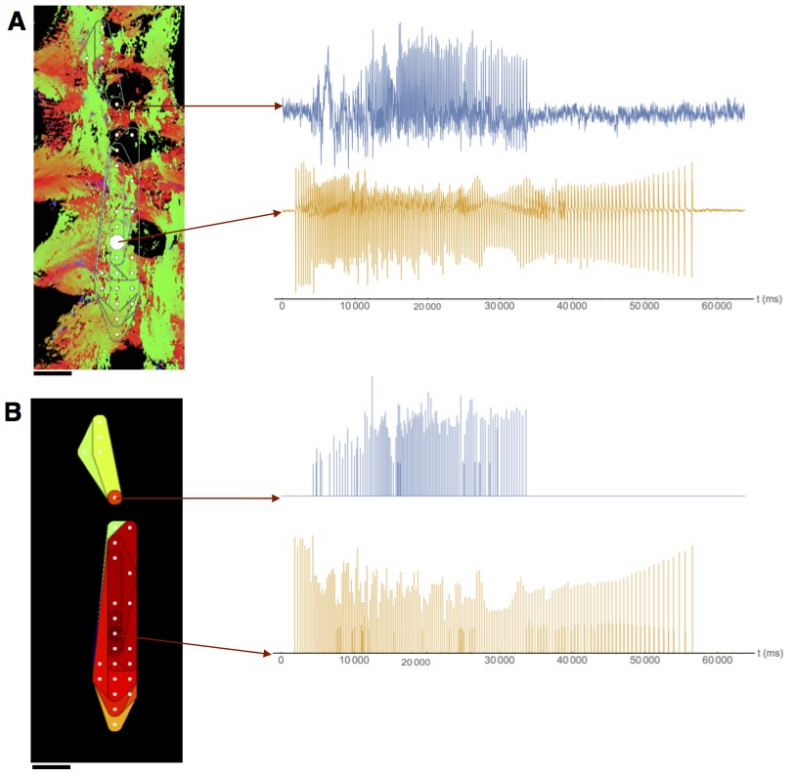

Figure 6.4: Comparison of electrograms taken from events detected as being separate in recording 2 from tissue block 1. **A**: The original recordings and the location of the electrodes in the tissue, given here as points on the *in silico* reconstruction. **B**; The individual action potentials detected by the procedure described in Chapter 5 and the position of the electrodes in the context of the isochrone map of the excitation in the tissue. Here the electrograms show markedly different bursting behaviour, suggesting that these are indeed two separate events. For convenience, electrogram comparisons for the other tissue blocks will only contain the isochrone map and original recordings to visualise this comparison. Scale bar represents 5 mm.

6.1 Results

6.1.1 Tissue block 1

The isochrone maps generated for the first tissue block are shown in the top row of Figure 6.3. The second recording shows a burst which has been deemed separate from the main excitation wave, which may indicate a second initiation point. A comparison of electrograms from the main excitation wave and this secondary wave is shown in Figure 6.4. The difference in burst length and start time between these two electrograms suggests

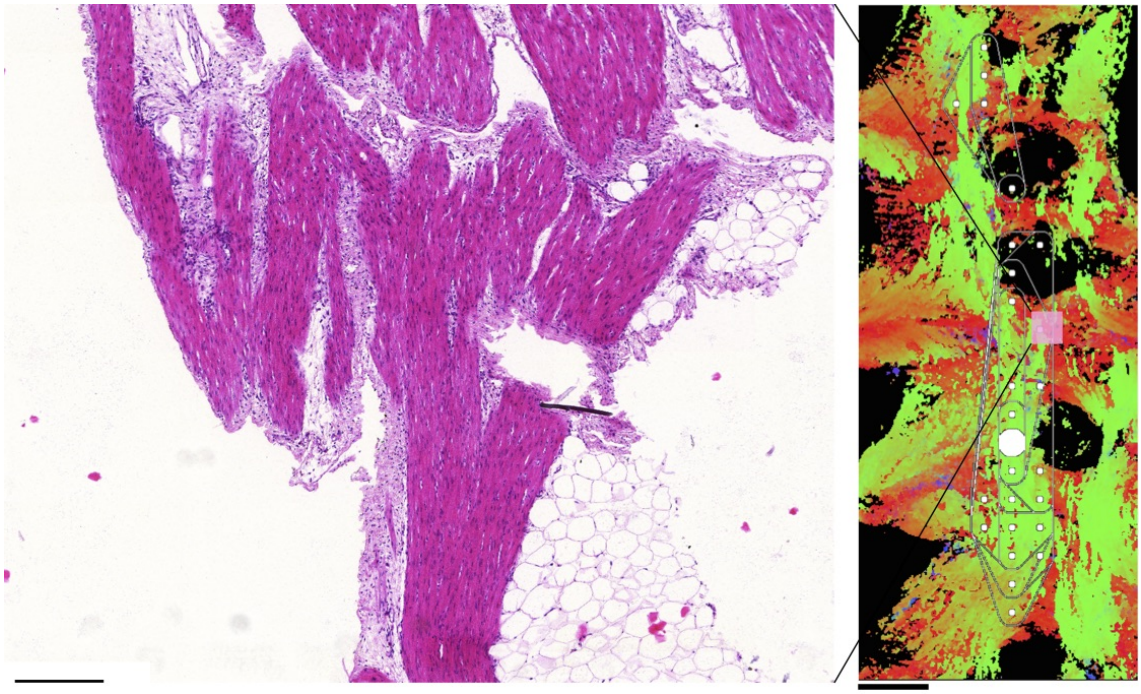

Figure 6.5: An example of fibres situated across the mesometrial border in tissue block 1. Here the fibres appear to connect; however the precise level of connectivity cannot be determined from two-dimensional images alone. The three-dimensional structure is required to determine the full connectivity, which is discussed in Chapter 7. Left panel scale bar represents 200 μm, right panel scale bar represents 5 mm.

that these are indeed two separate excitation waves. The presence of two excitation waves indicates that activity is being initiated from multiple sources simultaneously.

The second row of Figure 6.3 shows the isochrones in comparison to the top surface of the *in silico* tissue block. The direction of the fibres on the surface in the area covered by the excitation waves in each recording is generally vertical in the image plane, indicated by green. The main excitation wave in all three isochrone maps is restricted to the same approximate area. This might indicate a mechanism preventing current flowing past the given bounds for the duration of the recordings. The bounds perpendicular to the fibre direction are narrower than along the fibre direction, which suggests of a low level of communication between adjacent fibres. In particular, there is no communication across the mesometrial border. Figure 6.5 shows fibres connected across the mesometrial border;

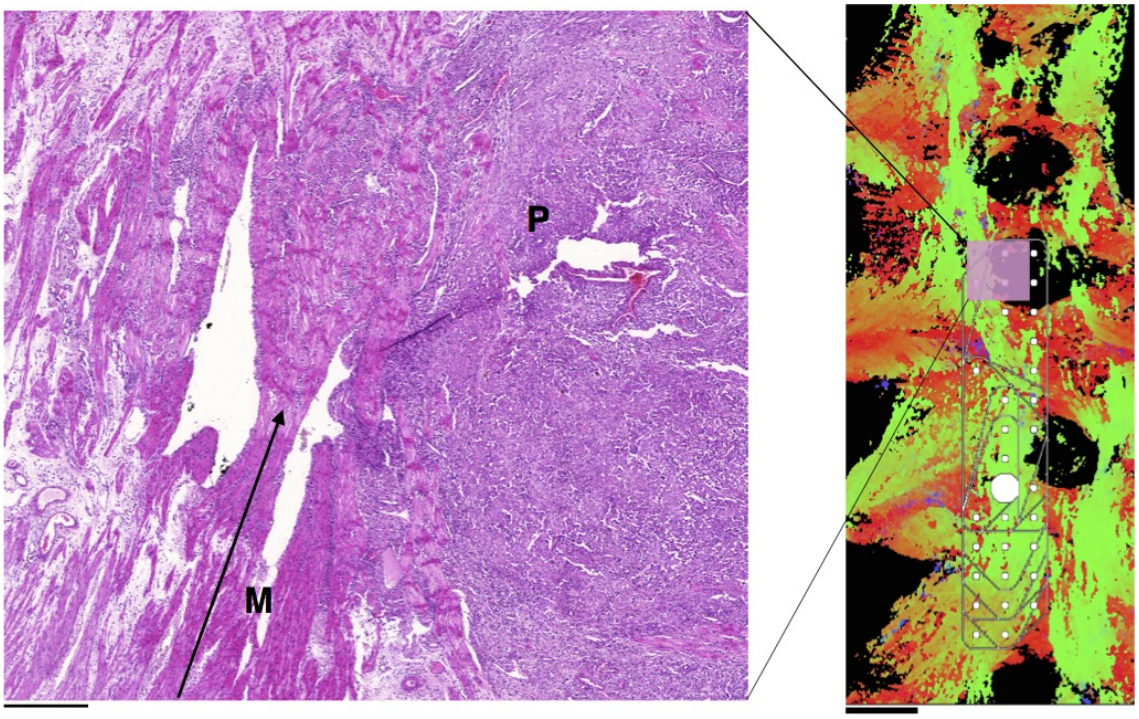

Figure 6.6: Fibres crossing into the placental bed in tissue block 1 in close proximity to a termination point of excitation along the fibre direction. The arrow indicates the direction of propagation in the myometrium. The conduction is blocked at some point within the volume containing this cross-sectional area. This slide shows the fibres are embedded in the placenta, suggesting that the placenta could play a role in preventing further propagation. **M** indicates myometrium, **P** indicates placenta. Left panel scale bar represents 500 μm, right panel scale bar represents 5 mm.

however the level of connectivity for the whole of the mesometrial border cannot be determined from two-dimensional histological slides alone. In Chapter 7 the three-dimensional *in silico* reconstruction will be used to determine the connectivity across the mesometrial border to assess whether this could be a factor in preventing propagation. The bounds along the fibre direction suggests a mechanism restricting propagation along the fibres. Figure 6.6 shows the tissue at the approximate location of a conduction block along the fibre direction. This shows that a possible path of conduction passes through fibres which penetrate the placental bed. The precise path that the excitation wave takes to reach this point is unknown due to the two-dimensional nature of the recording, but if the wave was restricted to these fibres then modification of the conductivity through these fibres by the

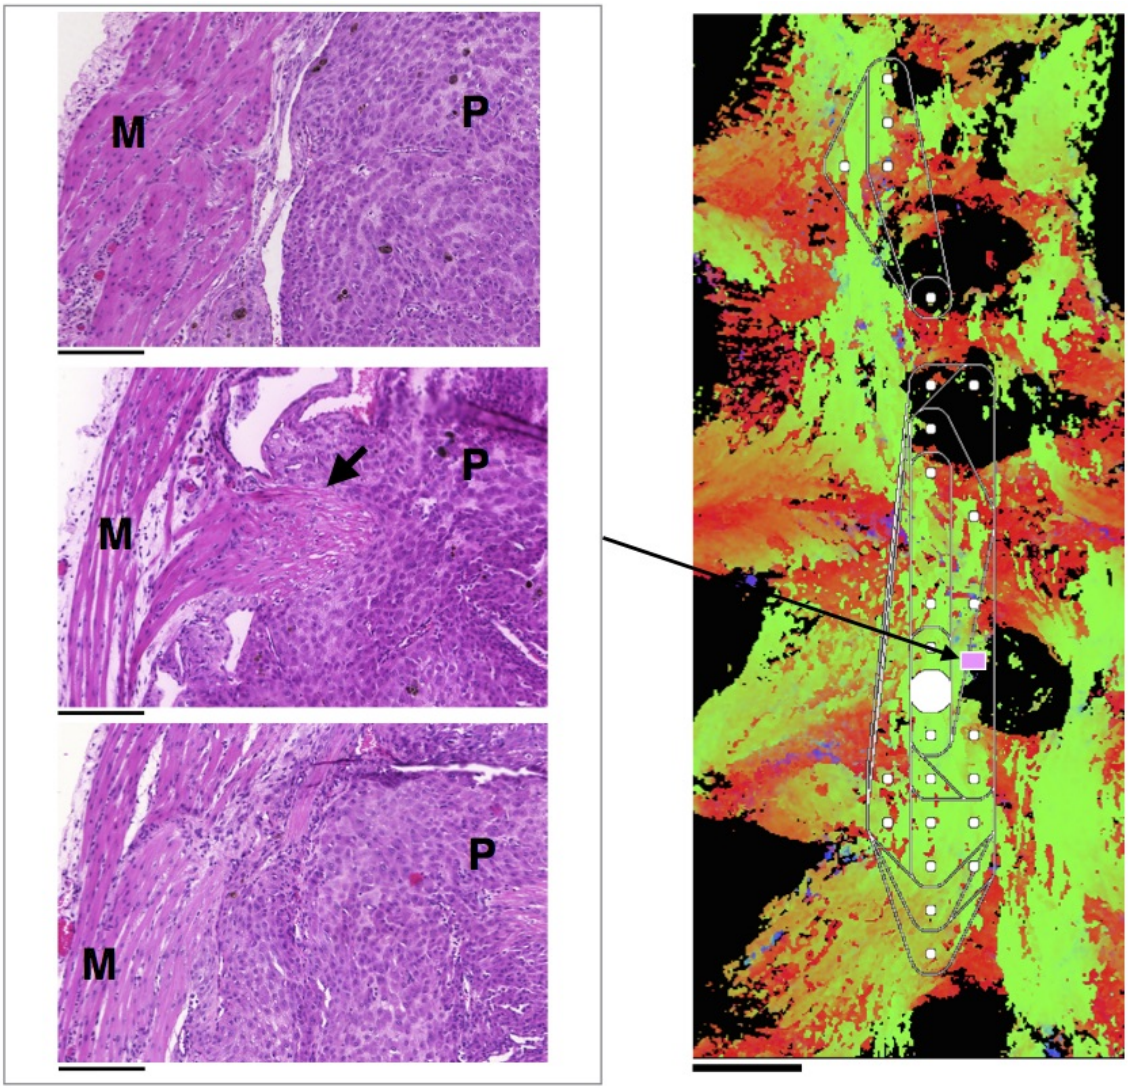

Figure 6.7: Fibres extending from the placental bed into the myometrium in tissue block 1. The middle histological image on the left shows this fibre, as indicated by the arrow. The images above and below this image are from sections above and below the middle section in the stack. These three images approximately occupy the space in the reconstructed tissue indicated by the pink square, which is in close proximity to the initiation point of the second and third recordings in tissue block 2. **M** indicates myometrium, **P** indicates placental bed. Left panel scale bars represents 200 μm, right panel scale bar represents 5 mm.

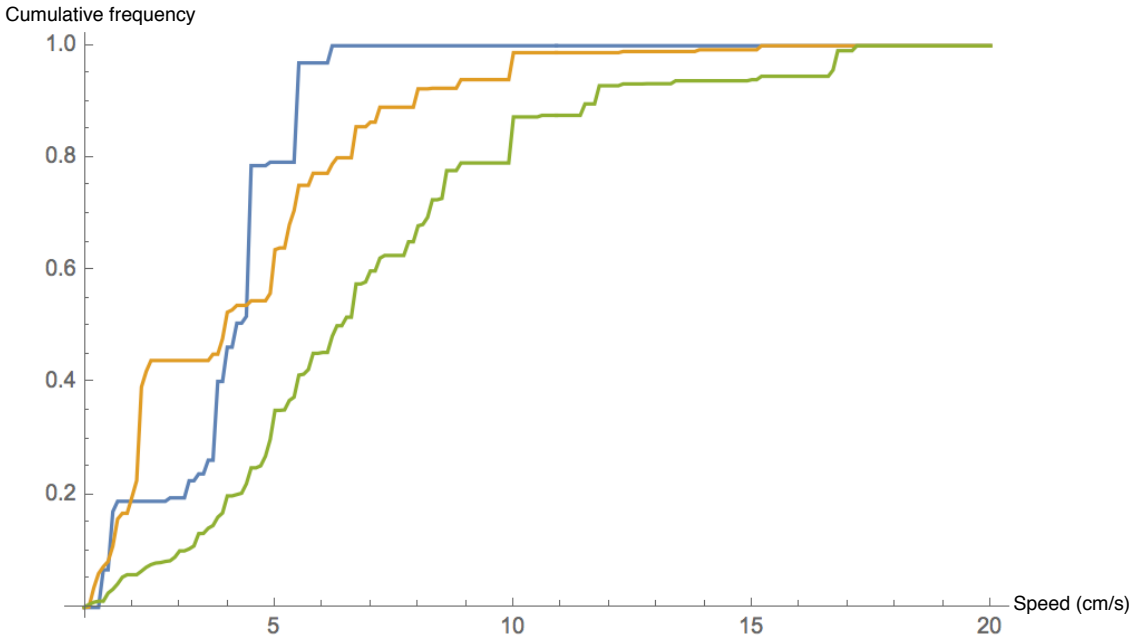

Figure 6.8: Empirical CDF of propagation speeds for each recording in tissue block 1. Blue represents recording 1, yellow represents recording 2, and green represents recording 3. The general shift to the right with time indicates that the speed is increasing between recordings; however, these recordings were made $\sim$ 10 minutes apart, which suggests that this increase is coincidental.

placental tissue could cause the wave to terminate at this point.

The initiation points for all recordings appear to lie in the vicinity of the mesometrial border, with the second and third recording showing initiation close to a specific implantation site. Moreover, the range of potential initiation points for the first recording includes locations adjacent to implantation sites. The secondary excitation in the second recording is also located at an implantation site. This suggests that the implantation sites determine the location at which excitation is initiated in these recordings. Figure 6.7 shows a fibre which extends into the placental bed from the myometrium. This fibre is located in close proximity to the initiation points of the second and third recordings, which suggests that it could be the location at which the excitation waves are initiated.

The mean propagation speeds for the recordings are summarised in Table 6.1. The distribution of the propagation speeds along the fibre direction for each recording is represented in Figure 6.8, showing variable distributions of speeds between recordings, which

Recording number	Mean speed ± CI (cm s^{-1})	n
1	3.90 ± 1.32	164
2	4.27 ± 1.01	978
3	6.96 ± 0.94	2502

Table 6.1: Mean propagation speeds along the fibre direction for each recording made for tissue block 1. Error denotes a 95% confidence interval.

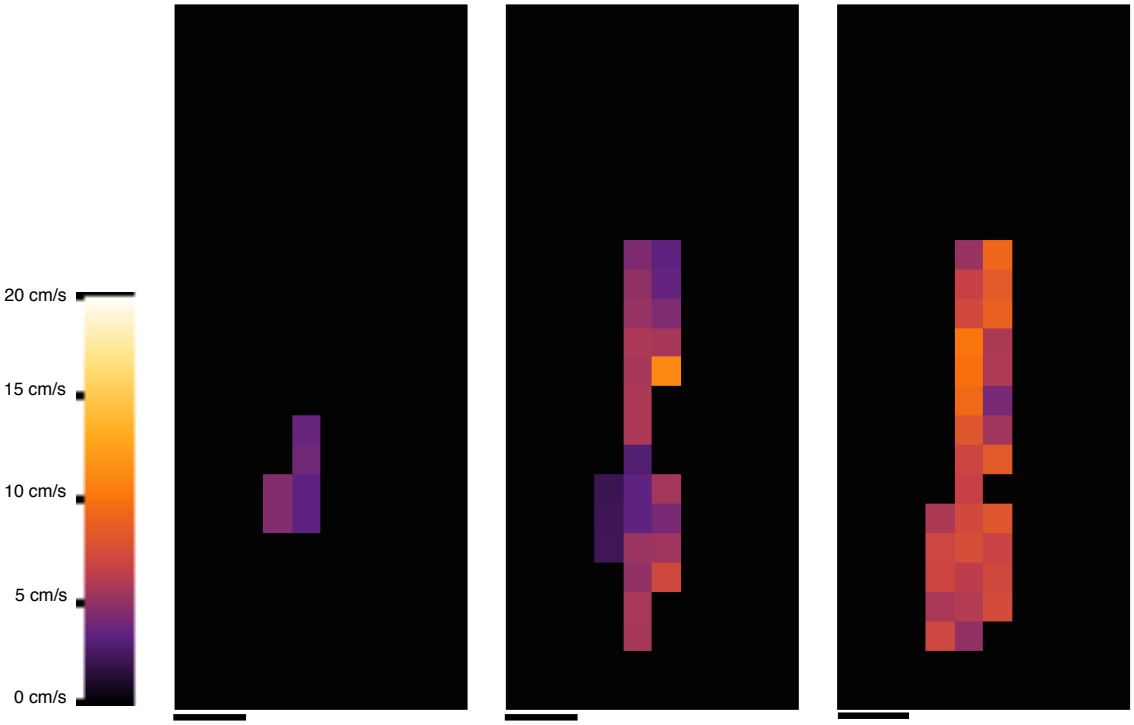

Figure 6.9: Representation of the propagation speeds as a function of space in tissue block 1. Scale bars represents 5 mm.

suggests that the propagation speeds vary with time. The images in Figure 6.9 represent the average propagation speed along the fibre direction through each point where a speed was recorded. The general increase in brightness between recordings suggests that this increase in speed is generally representative of the active tissue as a whole. There does not appear to be a general pattern for the spatial variation in speed within the recordings. This would suggest that the spatial variation does not strongly correlate with discernible anatomical features.

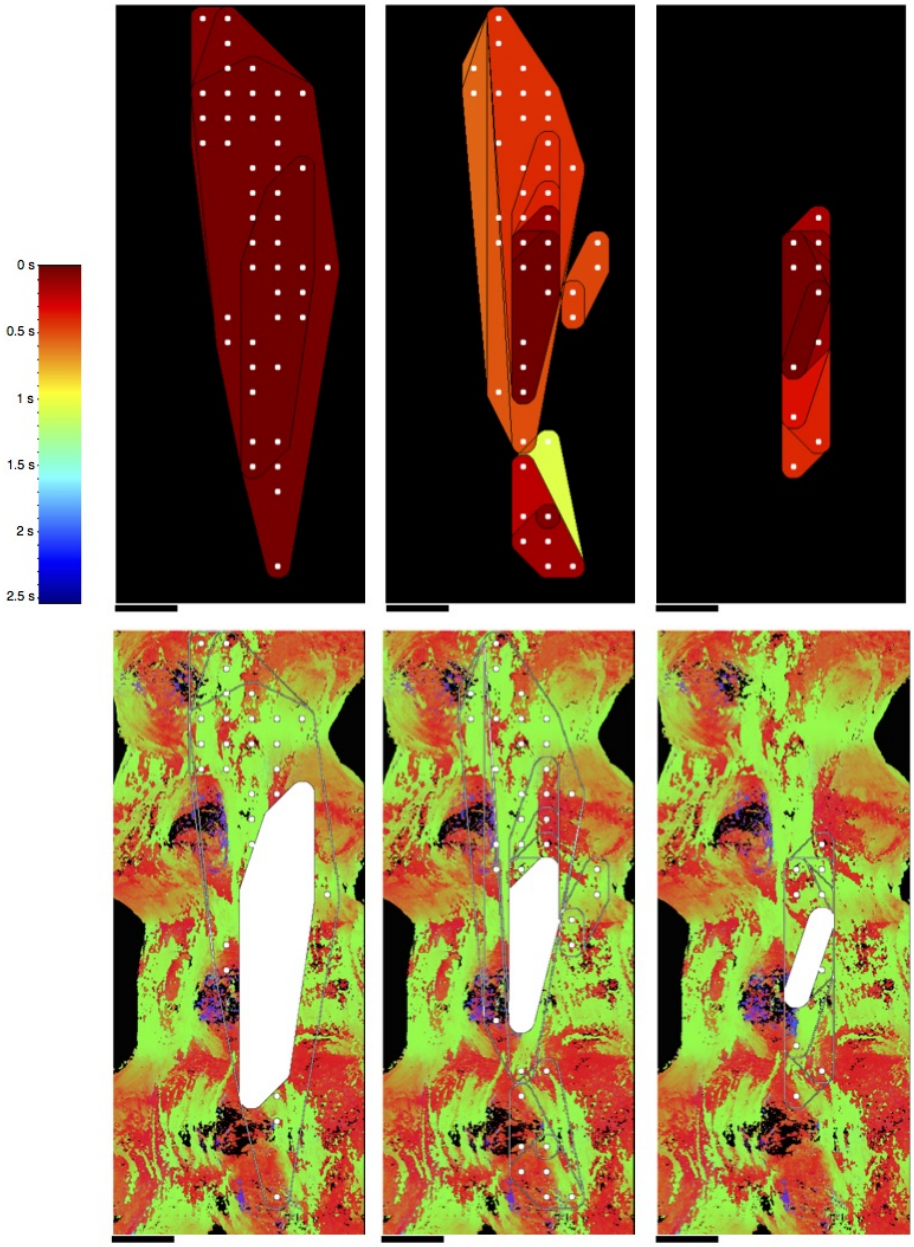

Figure 6.10: Isochrones for tissue block 2, in order of time of recording. Top row: isochrones taken at 200 ms intervals, with colour indicating time at which the excitation wave reached the electrode. White points indicate electrodes where bursts were recorded. Bottom row: comparison of isochrone initiation points and bounds with the top surface of virtual tissue. The tissue is pseudocoloured based on fibre direction, with green indicating vertical direction in the image plane, red indicating horizontal direction in the image plane, and blue indicating the direction perpendicular to the image plane. The ovarian end of the tissue is located at the top of each image. Scale bars represent 5 mm.

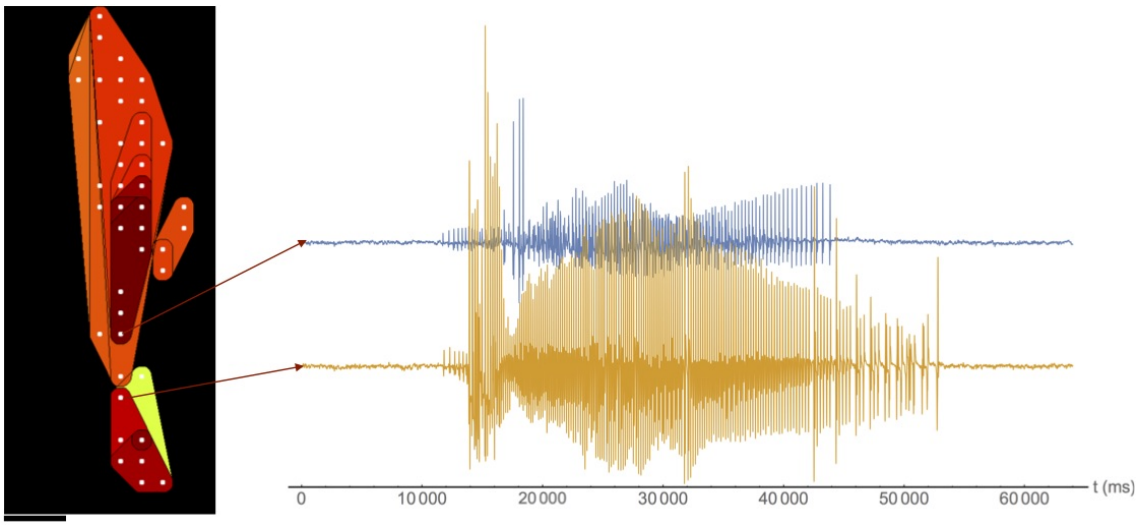

Figure 6.11: Comparison of electrograms taken from events detected as being separate in recording 2 from tissue block 2. The bursts appear similar toward the start of the recording, but the burst in the lower electrogram extends beyond the end of the burst in the upper electrogram. This would suggest that the two recordings are indeed from separate events. Scale bar represents 5 mm.

6.1.2 Tissue block 2

The isochrone maps generated for tissue block 2 are shown in the top row of Figure 6.10. The second recording shows two events separate from the main excitation wave, which could indicate up to three separate initiation points. The electrograms for these events are compared with the electrograms in the main excitation wave in Figures 6.11 and 6.12. The electrograms in Figure 6.11 appear similar at the start of the burst, which would indicate that the excitation recorded at the lower electrode could initially be part of the same wave as the burst recorded by the upper electrode. The lower burst extends for a period of time after the upper burst has ended, which suggests that the lower electrode is recording excitation from a different wave to the upper electrode toward the end of the burst. Accordingly, there appear to be two separate excitation events associated with this recording. The electrograms in Figure 6.11 appear different in length and frequency; however, the relatively low frequency of the action potentials in the lower recording suggests that it does not represent a full burst. This reduction in action potential frequency could

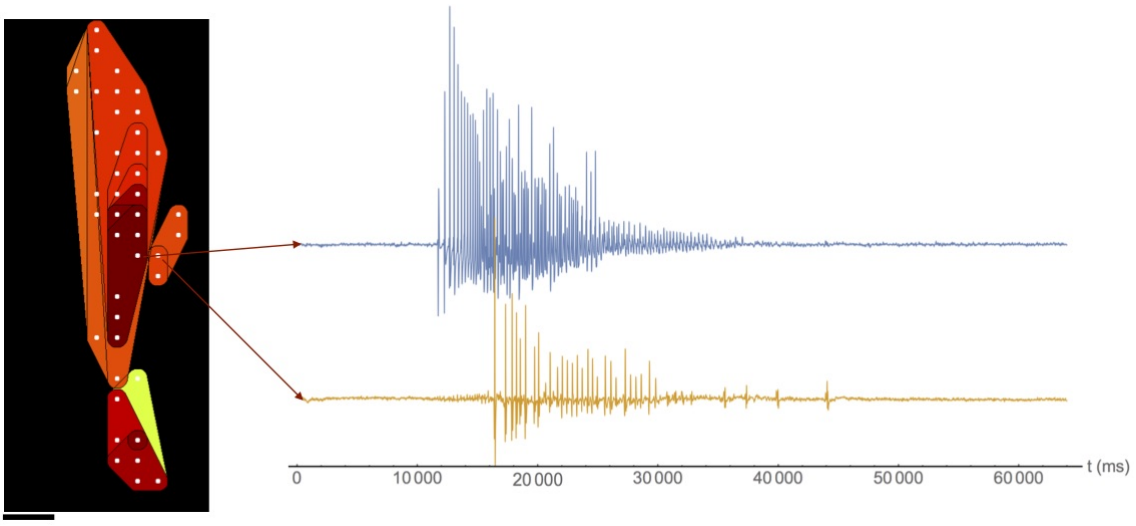

Figure 6.12: Comparison of electrograms taken from events detected as being separate in recording 2 from tissue block 2. The lower electrogram here has a markedly lower frequency of action potentials, which suggests that it has not recorded a complete burst. Scale bar represents 5 mm.

be due to a failure of propagation in the tissue between the electrodes for some of the action potentials, or alternatively because some physical factor such as tissue movement prevented the recording of the full set of action potentials.

The second row of Figure 6.10 shows the isochrones in comparison to the top surface of the *in silico* tissue reconstruction. The direction of the fibres covered by the excitation waves are generally vertical in the image plane, shown in green, with some horizontal fibre direction, shown in red, covered by the excitation wave in the first recording. Again, there appears to be a bound on propagation across the mesometrial border. This is shown clearly in the first and second recordings, with the bound covering the full length of the tissue. The histological image in Figure 6.13 shows that the fibres are connected across the mesometrial border; however it is difficult to determine how well-connected these fibres are from histological images alone. In Chapter 7 the three-dimensional connectivity across the mesometrial border shall be examined in detail in order to determine the level of connectivity between these fibres. The excitation wave in the first recording appears unbounded along the direction of the fibres, while the waves in the second and third

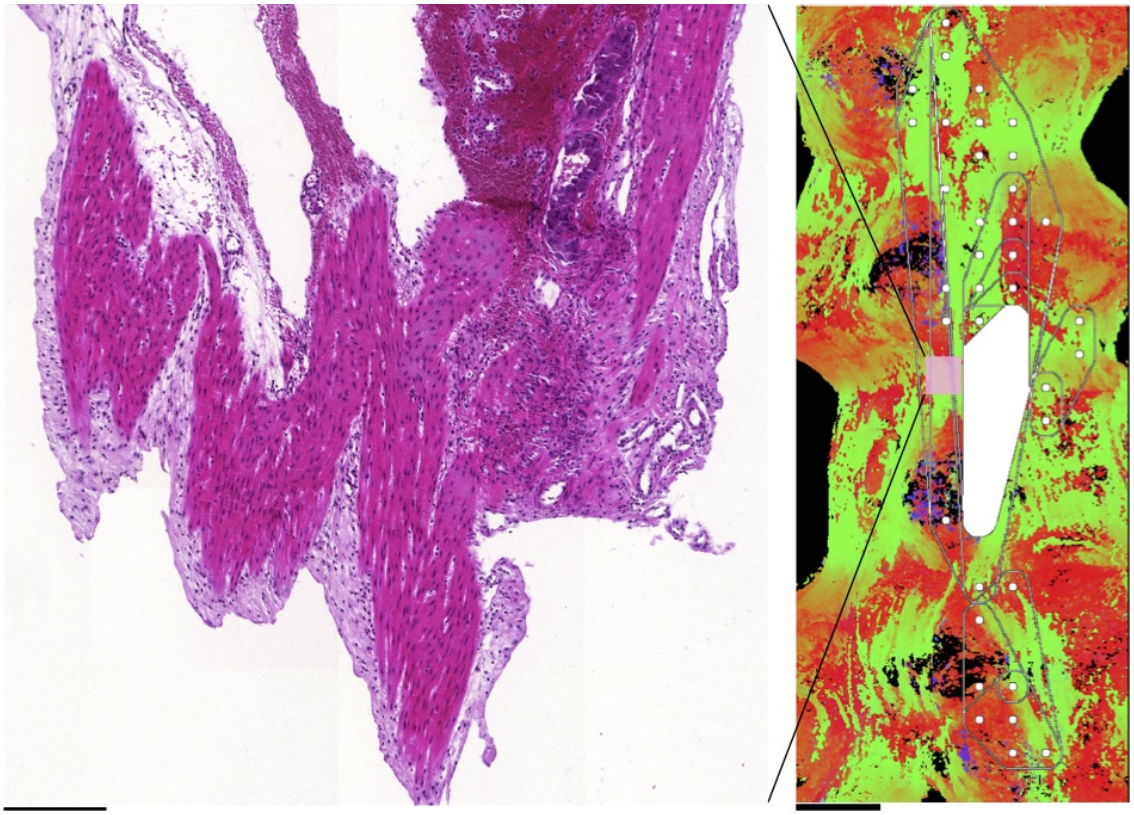

Figure 6.13: An example of fibres situated across the mesometrial border in tissue block 2. Here the fibres appear to connect; however the precise level of connectivity cannot be determined from two-dimensional images alone. The three-dimensional structure is required to determine the full connectivity, which is discussed in Chapter 7. Left panel scale bar represents 200 μm, right panel scale bar represents 5 mm.

recordings are bounded along the fibre direction. The bound at the lower end of the main excitation wave in the second recording is in the same approximate location as the lower bound on the wave in the third recording. This bound is located at the top of an implantation site, although in contrast to tissue block 1, there does not appear to be any clear evidence in the histological slides of fibres passing into the placenta at this point. Figure 6.14 shows fibres in this location passing adjacent to the placental bed. Again, it is not possible to determine the precise path taken by the excitation wave, but if the wave was restricted to these fibres then modification of the conductivity through these fibres by the placenta could cause the wave to terminate at this point.

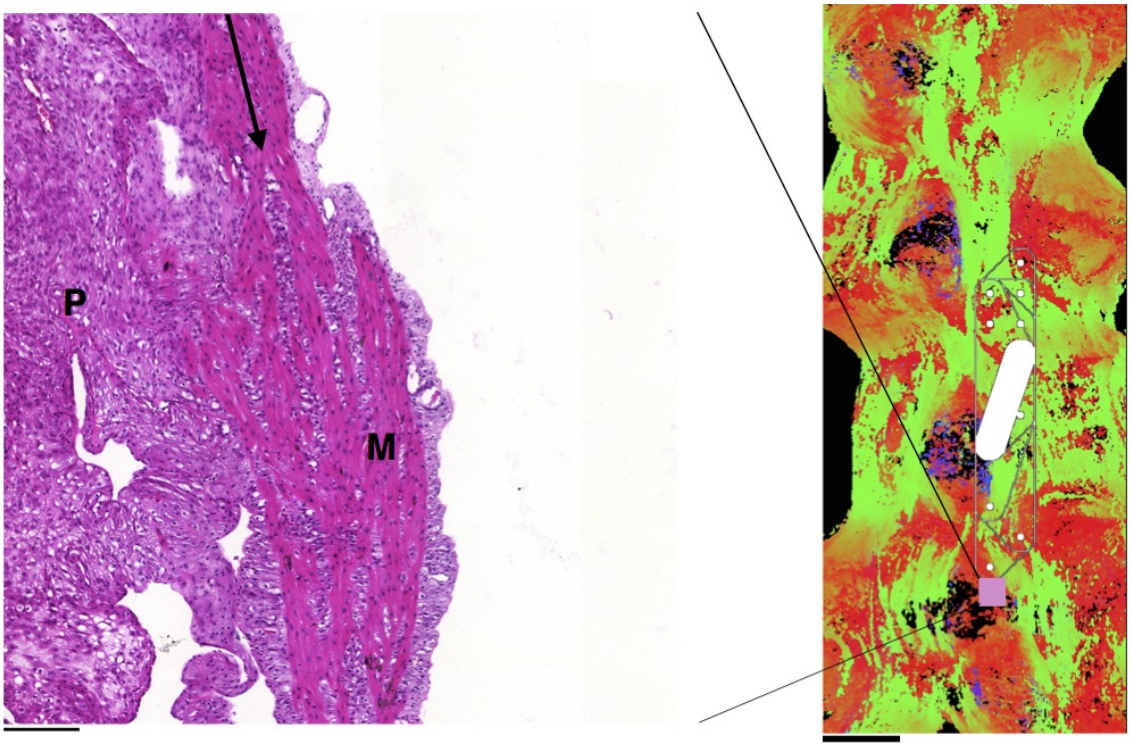

Figure 6.14: Fibres located in the region where the excitation wave terminates in the second and third recordings of tissue block 1. These fibres are situated adjacent to the placental bed, which could cause a reduced conductance in these fibres, which in turn could prevent propagation. **M** indicates myometrium, **P** indicates placental bed. Left panel scale bar represents 200 μm, right panel scale bar represents 5 mm.

The possible locations for initiation points of the main excitation wave in each recording appear to be in a similar location; however, it is difficult to identify the initiation point with any degree of confidence for the recording based on the electrograms, as illustrated in Figures 6.15 and 6.16. The set of possible initiation points in all three recordings is close to the mesometrial border. One potential location of initiation common to all three recordings is adjacent to an implantation site, which could indicate that the waves originate at some point close to the implantation site. Figure 6.17 shows a fibre extending from the placenta into the myometrium, located close to a possible initiation point. The proximity of this fibre to potential initiation points in each recording suggests that the excitation wave could be initiated in this fibre.

The mean propagation speeds for each of the recordings are summarised in Table 6.2.

Recording number	Mean speed ± CI (cm s^{-1})	n
1	5.37 ± 1.58	320
2	6.42 ± 1.62	540
3	5.78 ± 2.86	180

Table 6.2: Mean propagation speeds for each recording made for tissue block 2. Error denotes a 95% confidence interval.

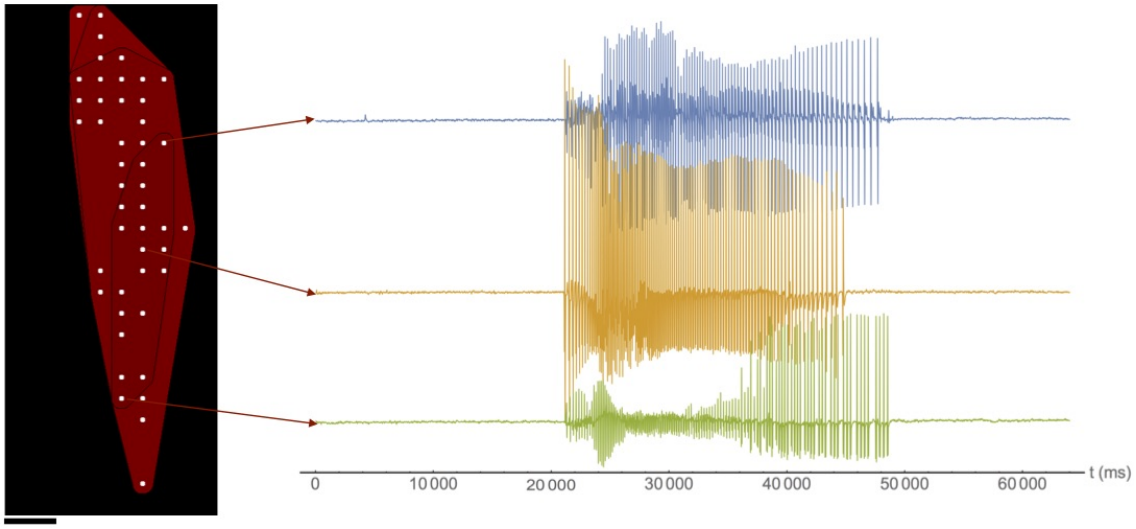

Figure 6.15: Comparison of electrograms from the area covered by the first isochrone of recording 1 from tissue block 2. All three bursts shown here have similar start times, and therefore cannot be accurately be placed in chronological order. Scale bar represents 5 mm.

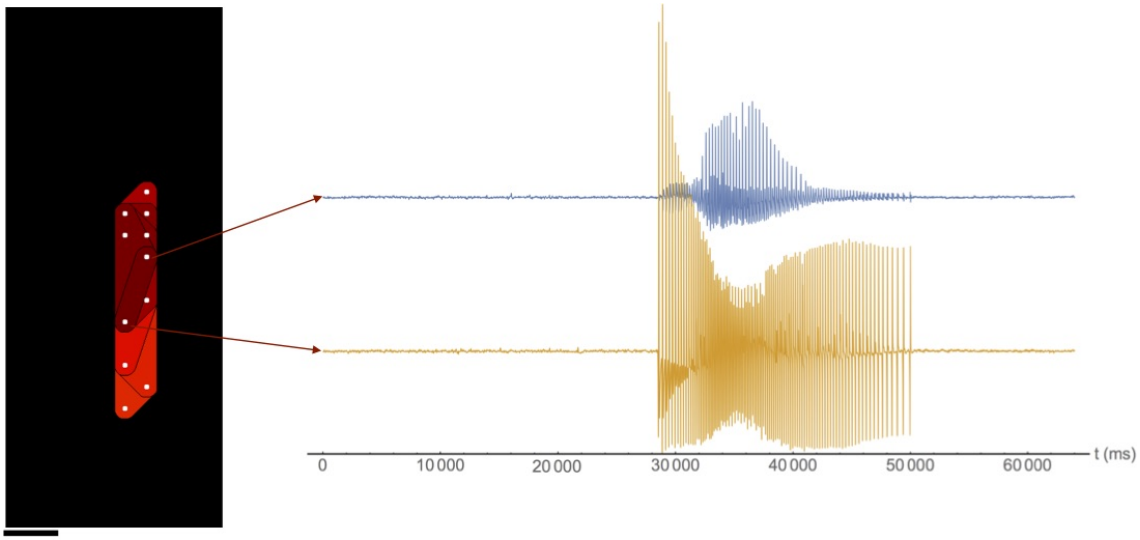

Figure 6.16: Comparison of electrograms from the area covered by the first isochrone of recording 3 from tissue block 2. Both recordings shown have similar start times, and therefore the order in which they started cannot be accurately determined. Scale bar represents 5 mm.

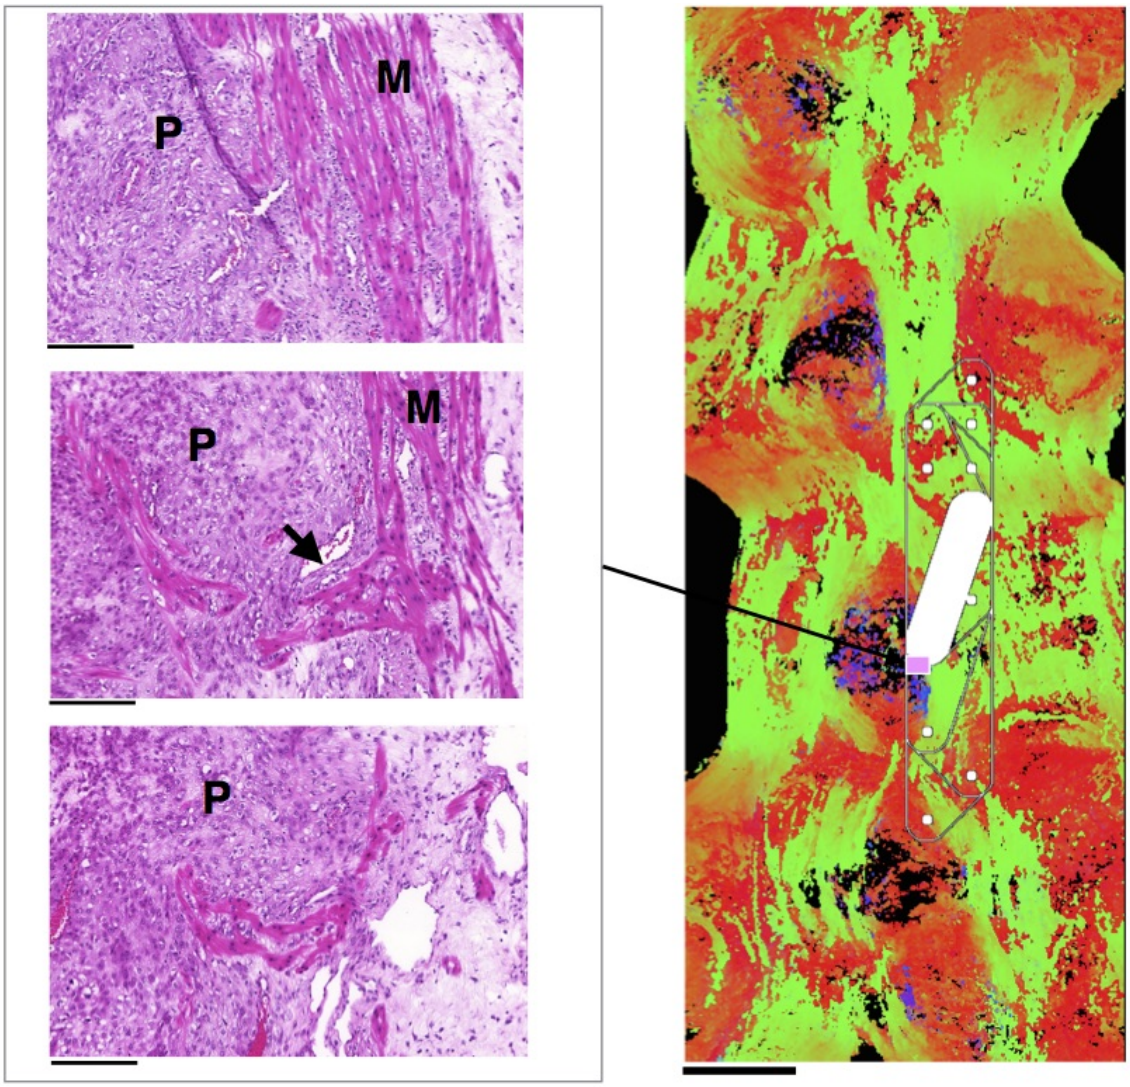

Figure 6.17: Fibres extending from the placental bed into the myometrium in tissue block 2. The middle histological image on the left shows this fibre, as indicated by the arrow. The images above and below this image are from sections above and below the middle section in the stack. These three images approximately occupy the space in the reconstructed tissue indicated by the pink square, which is in close proximity to a potential initiation point for all recordings in tissue block 2. **M** indicates myometrium, **P** indicates placental bed. Left panel scale bars represents 200 μm, right panel scale bar represents 5 mm.

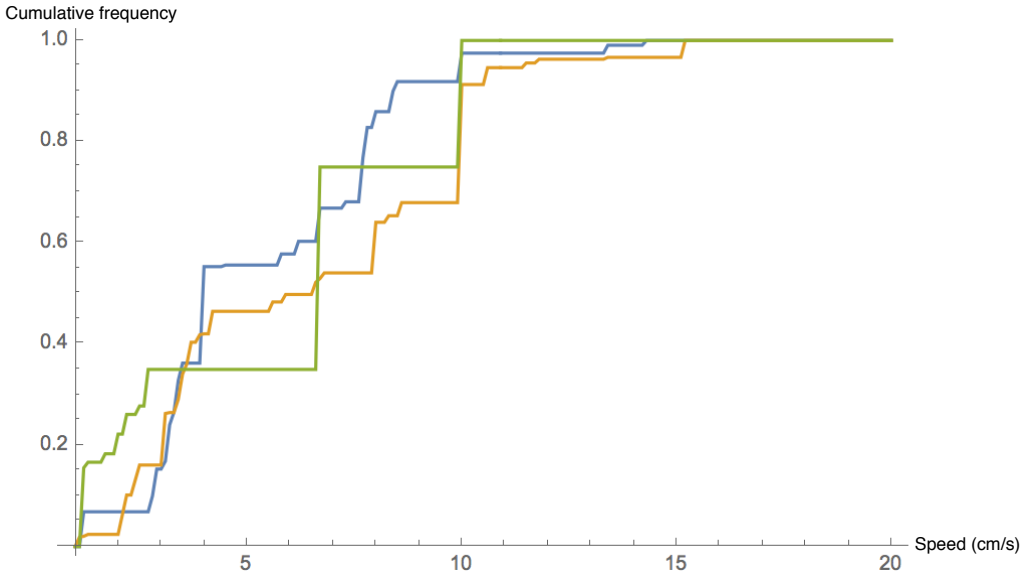

Figure 6.18: Empirical CDF of propagation speeds for each recording in tissue block 2. Blue represents recording 1, yellow represents recording 2, and green represents recording 3. The distribution of speeds varies between recordings, suggesting that the conductive properties of the tissue vary with time.

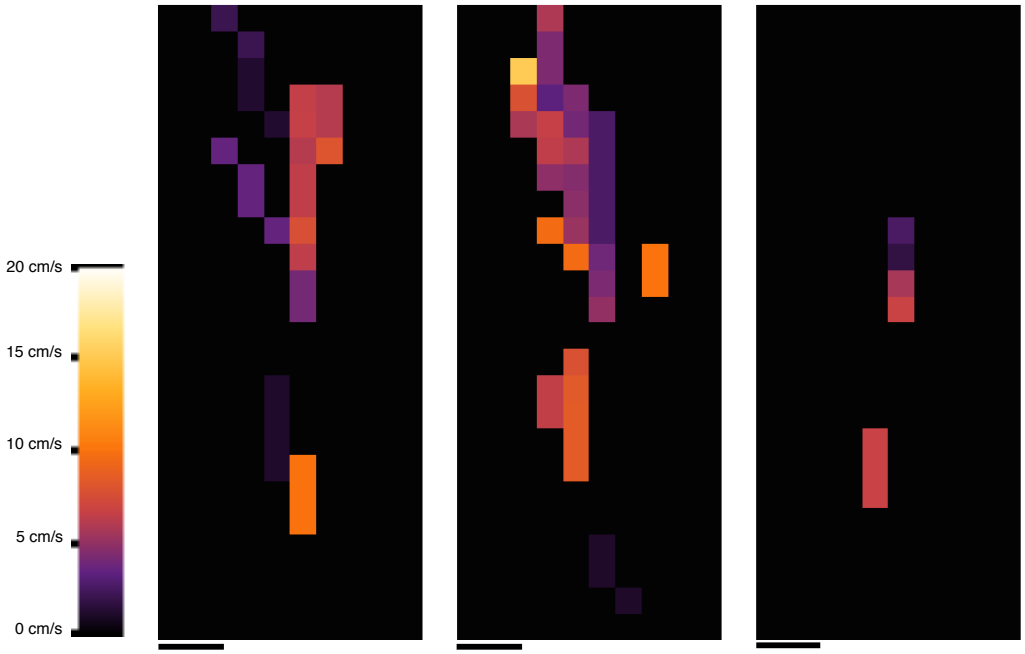

Figure 6.19: Representation of the propagation speeds as a function of space in tissue block 2. Scale bars represents 5 mm.

The propagation speeds for each recording are represented in Figure 6.18 as the empirical cumulative distribution functions. This shows a shift to the right of the distribution of speeds between recordings 1 and 2, suggesting that the conduction speed has increased between these two. This trend is not continued with recording 3, which has a distribution situated between the first two, suggesting that the propagation speed does not necessarily increase with time. This images in Figure 6.19 represent the average propagation speed through each point where a speed was recorded. There does not appear to be a general pattern for the spatial variation in speed within the recordings. This would suggest that the spatial variation is not representative of any specific anatomical feature.

6.1.3 Tissue block 3

The isochrone maps generated for tissue block 3 are shown in the top row of Figure 6.20. The first and third recordings show events outside the main excitation wave, which could indicate multiple initiation events. A comparison of electrograms in the first recording from each event is shown in Figure 6.21. These electrograms appear to be dissimilar in length and frequency, which would indicate that these are indeed two separate events. The electrograms for the secondary events in the third recording are compared with the electrograms in the main excitation wave in Figures 6.22 and 6.23. The electrograms in Figure 6.22 appear dissimilar in length and frequency of the burst; however, the relatively short length of the burst and low frequency of action potentials in the upper electrogram suggests that the recording in the upper electrogram is not representing a complete burst. The reduction in action potentials recorded may be due to a genuine reduction in the amount of propagation through the tissue, or because some physical aspect such as tissue movement prevented the electrode from recording some of the action potentials at that point. The electrograms in Figure 6.23 appear dissimilar in frequency and duration, which suggests that they are indeed separate events.

The second row of Figure 6.20 shows the isochrones in comparison to the top surface

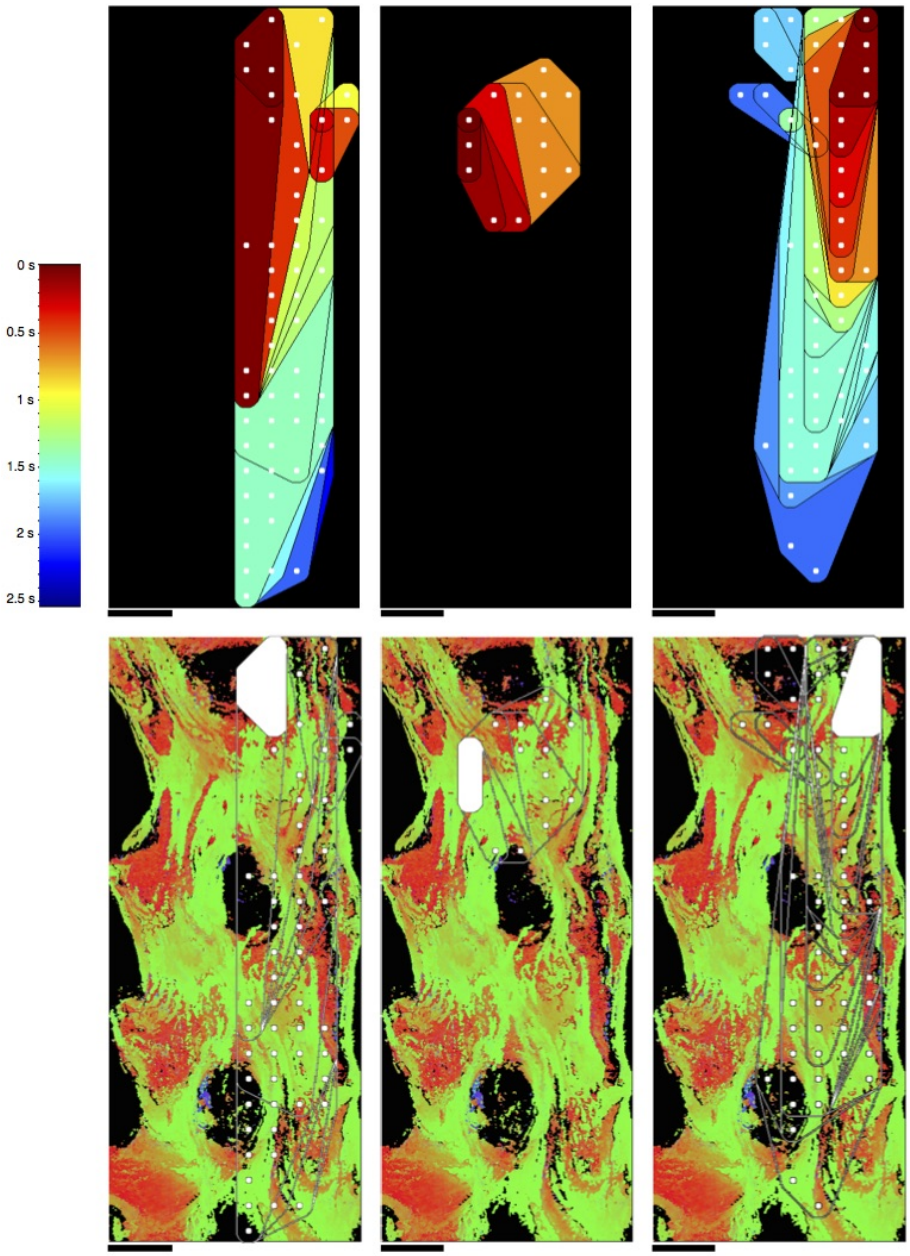

Figure 6.20: Isochrones for tissue block 3, in order of time recorded. Top row: isochrones taken at 200 ms intervals, with colour indicating time at which the excitation wave reached the electrode. White points indicate electrodes where bursts were recorded. Bottom row: comparison of isochrone initiation points and bounds with the top surface of virtual tissue. The tissue is pseudocoloured based on fibre direction, with green indicating the vertical direction in the image plane, red indicating the horizontal direction in the image plane, and blue indicating the direction perpendicular to the plane. The ovarian end of the tissue is located at the top of each image. Scale bars represent 5 mm.

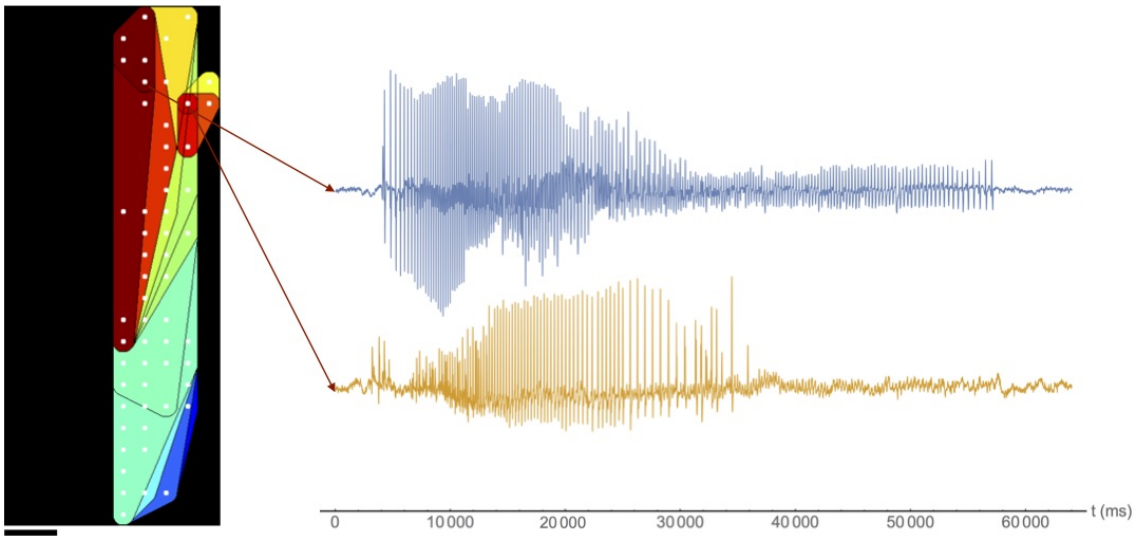

Figure 6.21: Comparison of electrograms taken from events detected as being separate in recording 1 from tissue block 3. These electrograms appear dissimilar in duration of burst and action potential frequency, and are therefore likely to be form separate excitation waves. Scale bar represents 5 mm.

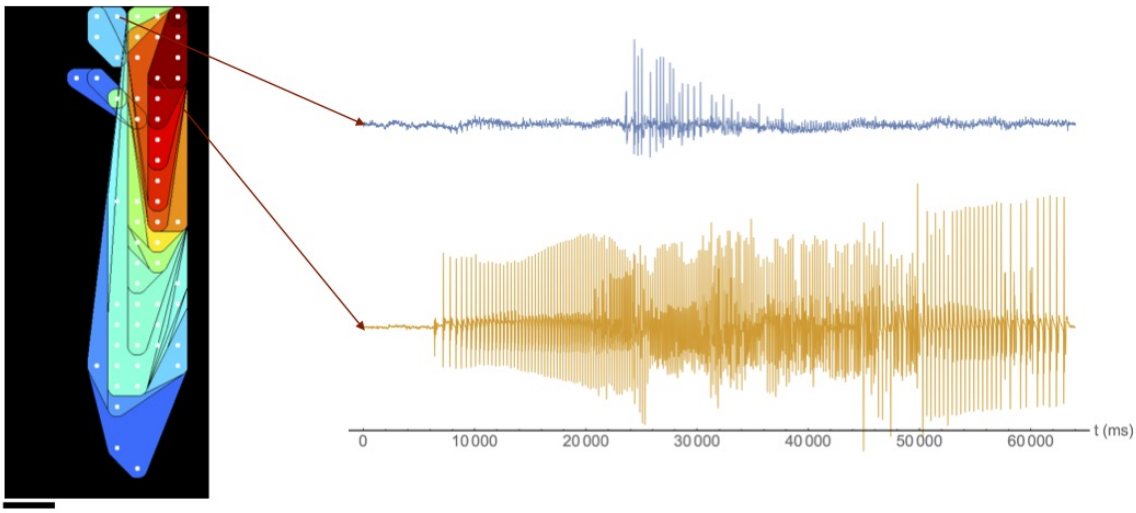

Figure 6.22: Comparison of electrograms taken from events detected as being separate in recording 3 from tissue block 3. The upper electrogram shows a relatively low frequency of action potentials, suggesting that it is an incomplete representation of a burst. Scale bar represents 5 mm.

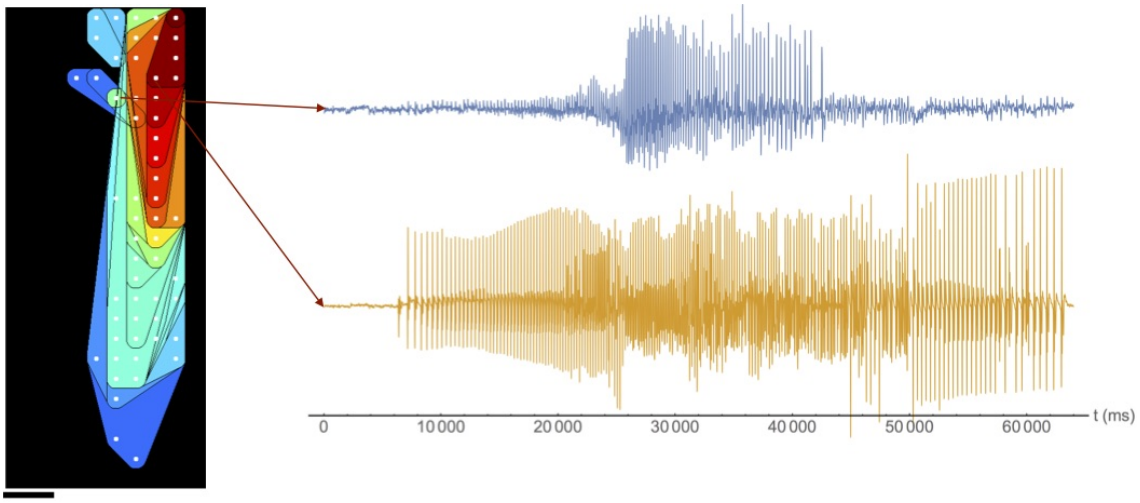

Figure 6.23: Comparison of electrograms taken from events detected as being separate in recording 3 from tissue block 3. The bursts in these electrograms appear dissimilar in length and frequency, suggesting that they are indeed separate events. Scale bar represents 5 mm.

of the *in silico* tissue reconstruction. Generally the fibres within the area covered by the excitation wave run vertically in the image plane, indicated by green. The excitation in the first and third recordings appears to be unbounded along the fibre direction, while the excitation in the second recording is restricted to an area between two implantation sites. The excitation in the first and third recordings are bounded by the mesometrial border in the direction perpendicular to the fibre direction. The second recording, however, propagates across the mesometrial border. This suggests that conduction across the mesometrial border is not completely eliminated, but may be prevented under certain conditions. Figure 6.24 shows fibres across the mesometrial border. As was the case with the previous blocks, these fibres are connected, and a detailed analysis of the connectivity across the mesometrial border will be presented in Chapter 7.

The initiation point of the first recording is located adjacent to an implantation site. The initiation point of the second recording is located in close proximity to the same implantation site. The third recording appears to have an initiation point toward the anti-mesometrial border; however, this initiation point is also at the edge of the electrode array, and so the excitation could originate at a point beyond the array. Figure 6.25 shows

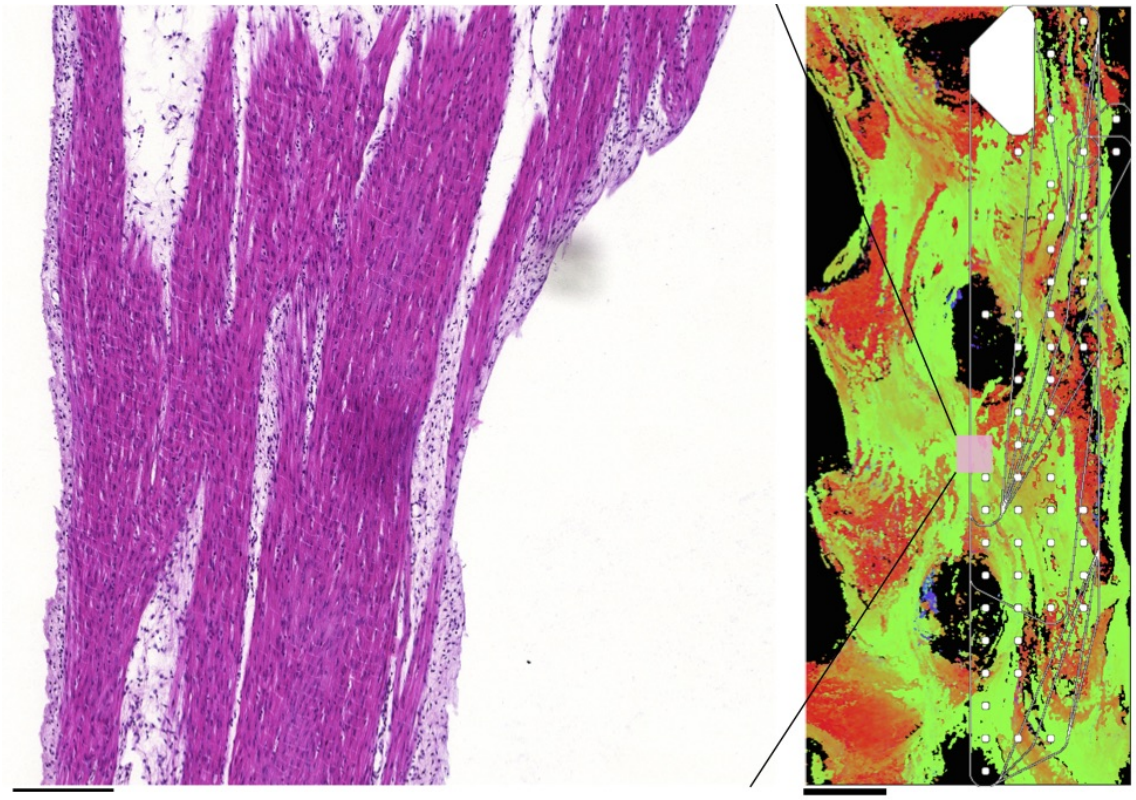

Figure 6.24: An example of fibres situated across the mesometrial border in tissue block 3. Here the fibres appear to connect; however the precise level of connectivity cannot be determined from two-dimensional images alone. The three-dimensional structure is required to determine the full connectivity, which is discussed in Chapter 7. Left panel scale bar represents 200 μm, right panel scale bar represents 5 mm.

the tissue at the location of the initiation point in the first recording. The figure shows fibres connecting the interior of the placenta with the adjacent fibres. This fibre could be the location where the excitation originates.

The mean propagation speeds are summarised in Table 6.3. The propagation speeds for each recording are represented in Figure 6.26 as the empirical cumulative distribution functions. This shows that the distribution of propagation speeds is similar between recordings 1 and 3, while recording 2 generally has slower speeds. The reduction in speed for recording 2 suggests that the difference in propagation behaviour described above could be linked to the propagation speed. The images in Figure 6.27 represent the average propagation speed through each point where a speed was recorded. Comparing the images

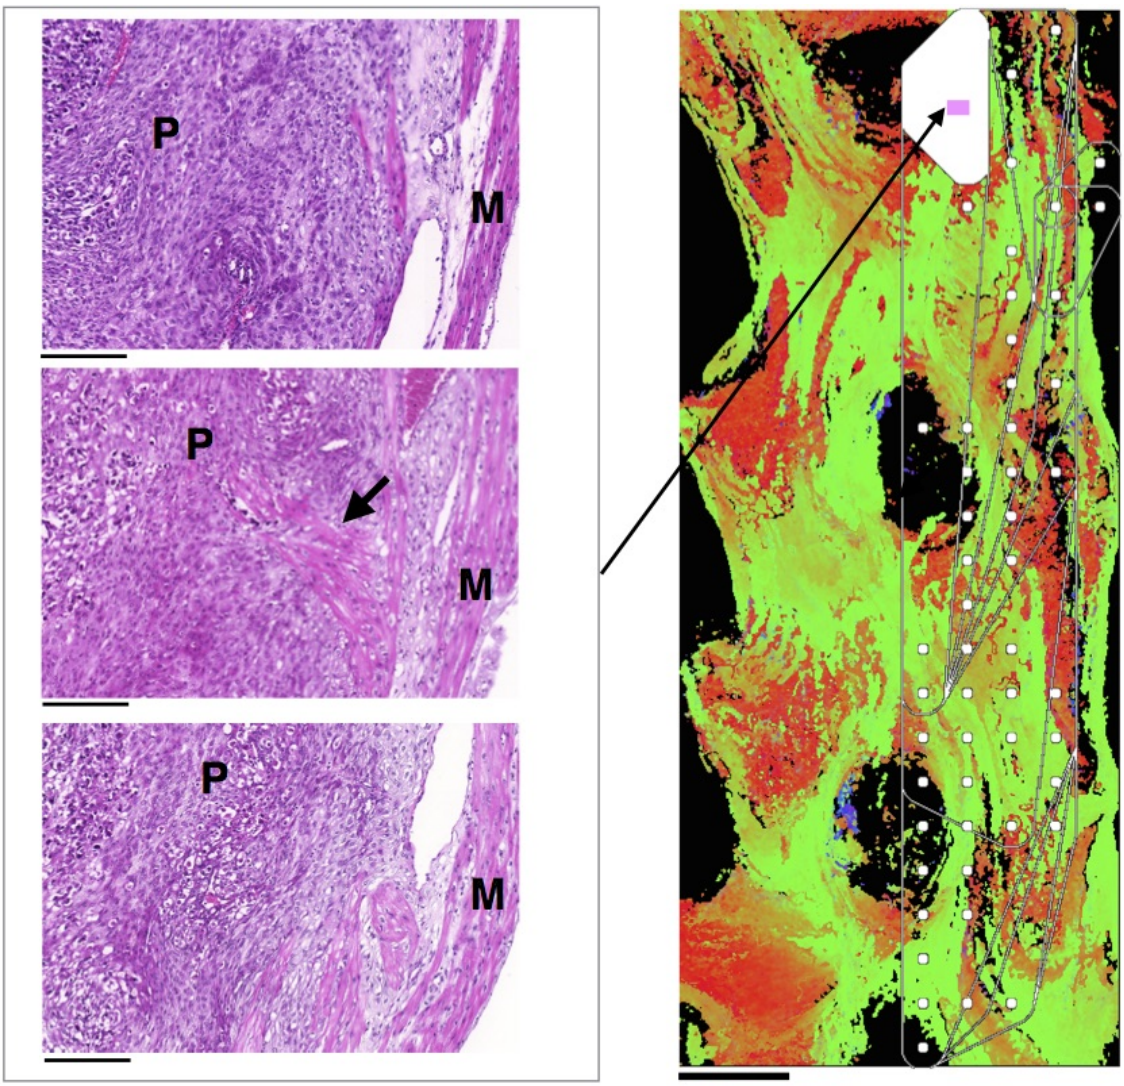

Figure 6.25: Fibres extending from the placental bed into the myometrium in tissue block 3. The middle histological image on the left shows this fibre, as indicated by the arrow. The images above and below this image are from sections above and below the middle section in the stack. These three images approximately occupy the space in the reconstructed tissue indicated by the pink square, which is in close proximity to the initiation point of the first recording in tissue block 2. **M** indicates myometrium, **P** indicates placental bed. Left panel scale bars represents 200 μm, right panel scale bar represents 5 mm.

Recording number	Mean speed ± CI (cm s^{-1})	n
1	7.45 ± 0.43	15727
2	4.74 ± 2.62	226
3	6.95 ± 0.54	9713

Table 6.3: Mean propagation speeds for each recording made for tissue block 3. Error denotes a 95% confidence interval.

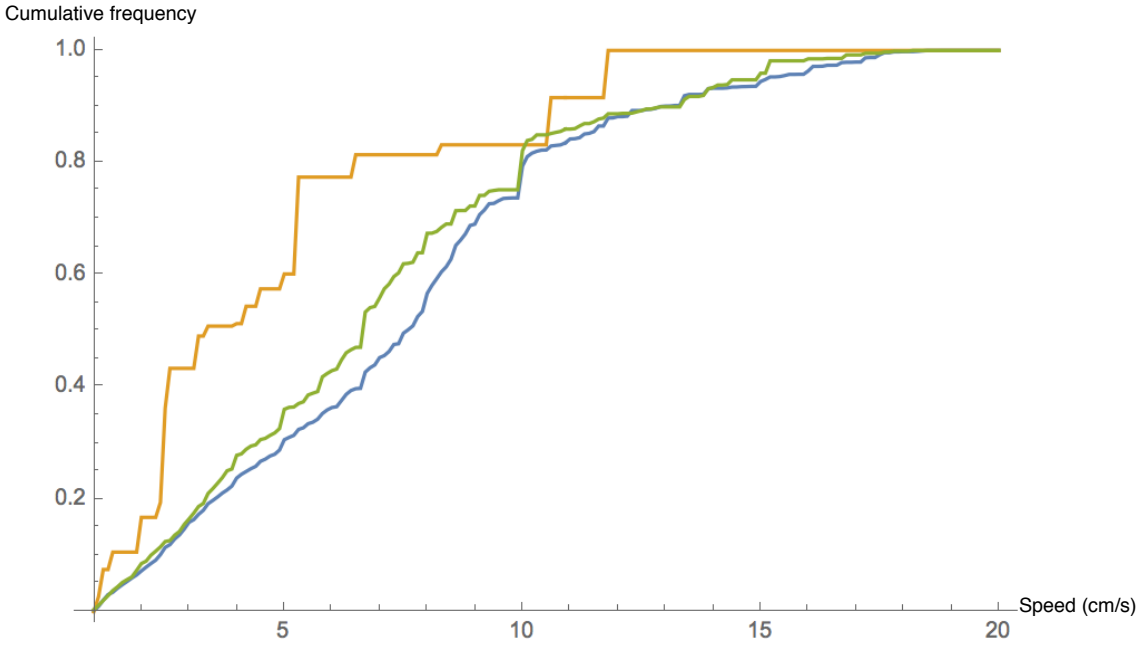

Figure 6.26: Empirical CDF of propagation speeds for each recording in tissue block 3. Blue represents recording 1, yellow represents recording 2, and green represents recording 3. Recordings 1 and 3 show similar distributions; however, the spatial distribution of speeds shown in Figure 6.27 shows that the speed does vary locally between recordings.

for recordings 1 and 3, the areas where the propagation speed is highest, indicated by pale red, are different between the recordings. This suggests that, while the propagation speed varies with time, the spatial variation is not dependent on location in the tissue.

6.1.4 Summary

In all three tissue samples, the recorded excitation waves did not extend over the full area of the tissue. In eight of the recordings, the excitation wave terminated at the mesometrial border. While there appears to be a connection across the mesometrial border in all tissue

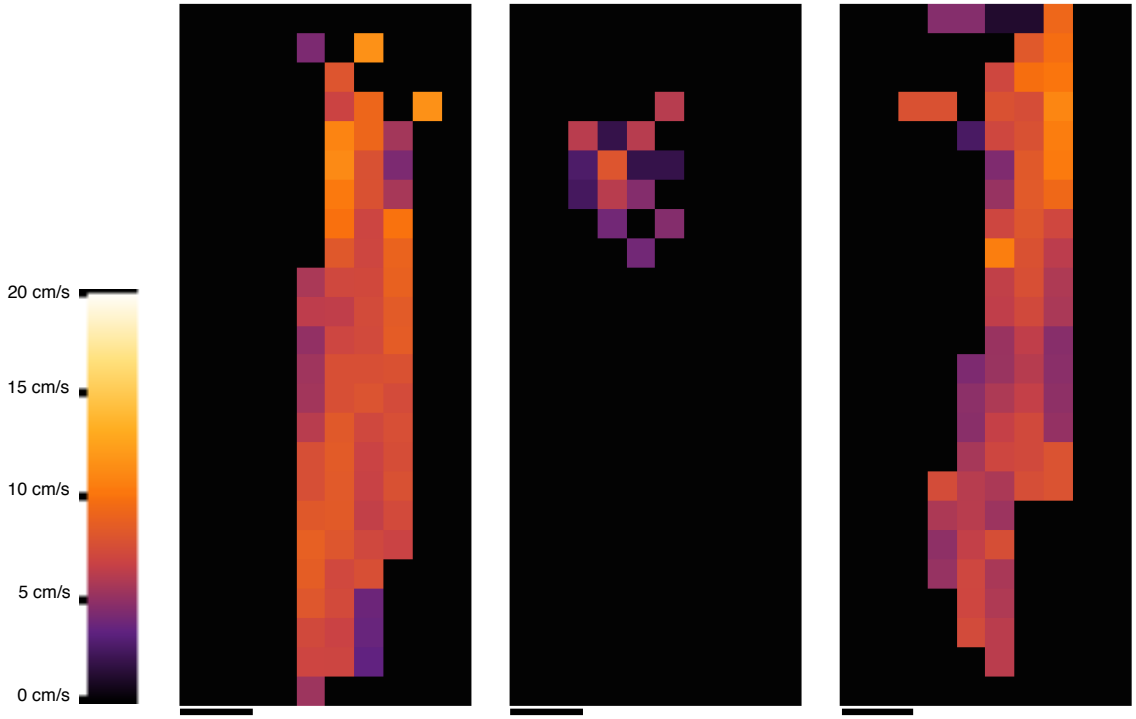

Figure 6.27: Representation of the propagation speeds as a function of space in tissue block 3. Scale bars represents 5 mm.

blocks, the three-dimensional analysis of connectivity presented in Chapter 7 shall be used to determine how well-connected these fibres are. All three recordings from tissue block 1 were bounded along the fibre direction at the same locations. Examination of the histology of tissue block 1 at that point showed that some fibres extended into the placental bed. A reduction of conductivity in these fibres could cause termination of the excitation wave at that point. Two of the recordings from tissue block 2 were bounded along the fibre direction. These bounded recordings shared a bound. Examination of the histology of tissue block 2 at that point showed that some fibres extended into the placental bed. Again, reduction of conductivity in these fibres could cause termination of the excitation wave. The initiation points of activity in eight of the recordings were located in close proximity to implantation sites. Examination of the histology around these points showed that there exist fibres extending out of the placental bed into the adjacent tissue, which could be a location of stimulation of the excitation. The propagation speeds varied

between recordings, indicating that the conductive properties of the tissue vary with time. The propagation speeds varied in space in each recording; however, this variation does not follow any distinguishable pattern for all recordings in any of the tissue blocks.

Chapter 7

Analysis of reconstructed tissue architecture

In this chapter the *in silico* reconstructions of rat and human myometrium generated using the methods described in Chapters 2 and 3 will be analysed. This reconstruction was generated from histological sections, which were registered and analysed by a semi-automated procedure to obtain the final volume. The slides were coarse-grained into regions, which were assigned planar directions based on the orientation of nuclei in the given region. This regional representation was used to register the slides to obtain a three-dimensional representation. The registered slides were subsequently coarse-grained in the z-direction in order to create a uniformly spaced volume, where each voxel is assigned a three-dimensional direction vector based on the direction of the nuclei within the given voxel and the local fibre architecture. The volume generated in this manner was subsequently segmented to generate the *in silico* reconstruction of the myometrium. The *in silico* reconstructions described here can be found at `https://figshare.com/s/9a54e1ae6699cf713670`.

All three-dimensional reconstructions were visualised using 3D Slicer [38] (www.slicer.org), with the exception of Figures 7.20 and 7.21, which were generated in Osirix [94] (www.osirix-viewer.com). These reconstructions were facilitated by generating DICOM files using the dcm4che package (www.dcm4che.org). The code used to generate these DICOM files can be found at `https://figshare.com/s/651e592a91f309182edd`.

7.1 Verification of reconstructed tissue

The *in silico* reconstruction of the myometrial tissue generated according to the methods described in Chapter 2 is based on histological sections. This reconstruction specifically and explicitly represents fibre direction. In order to verify that the final *in silico* version represents the original tissue, the following verification steps were performed. First, to verify the planar direction vectors assigned to the tissue, tiles were sampled at random and were visually compared to the assigned direction vectors. Examples of this comparison are given in Chapter 2, and further examples can be found at `https://figshare.com/s/9a54e1ae6699cf713670`. Second, in order to verify that the registration did not deform the original slides beyond a given threshold, the fibre widths in registered and unregistered tiles were compared to show equivalence. The third verification step was to compare images of the original tissue blocks prior to sectioning with images of the *in silico* reconstruction. This comparison was performed manually, in order to verify that the shape and features of the reconstructed tissue matched the original tissue.

The second step in the verification process was to compare the registered slides with the unregistered slides, to ensure that the registration procedure has not deformed the reconstructed tissue. The quantity used to compare these slides was *fibre width*. Fibre width was automatically measured at points on a randomly placed grid, as described in Chapter 3, for each registered and unregistered slide. The grids had a separation corresponding to 240 μm in the first rat tissue block, and 190 μm in all other tissue blocks. These measurements were taken from each slide present before and after registration. The distributions of these datasets are positively skewed, so the log transform was applied to enable comparison of means. The distributions of the resulting log-transformed data sets for registered and unregistered slides in each tissue block are shown as box plots in Figure 7.1. The log-transformed widths from unregistered and registered slides were compared to test for equivalence of the mean in the following manner. The maximum de-

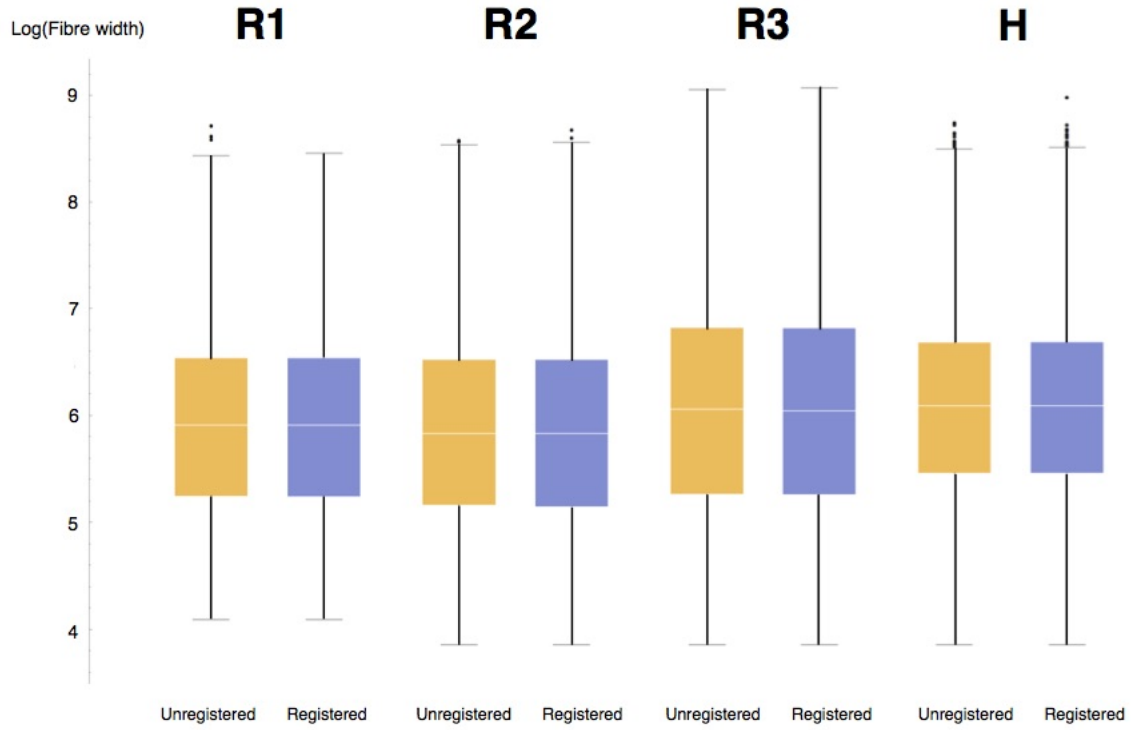

Figure 7.1: Distribution of log-transformed fibre widths in registered and unregistered slides. **R1**, **R2**, and **R3** represent each of the rat tissue blocks, **H** represents the human tissue block. The general similarity in shape between registered and unregistered slides in each block suggests that any deformation incurred by the registration process has little effect on the overall structure of the tissue block.

formation allowed was set to 5%, meaning that if the means of fibre widths in unregistered and registered slides differ by more than 5%, the registration process is deemed to have substantially deformed the tissue and the tissue is no longer representative of the original slides. The log transform of this effect size is given by the interval

$$(\log(1) - \log(1.05), \log(1) - \log(0.95)) = (-0.0488, 0.0513).$$

If the difference between registered and unregistered tissue blocks of the mean of the log-transform of the fibre widths is within the above interval with 95% confidence, then the deformation incurred by the registration process is deemed acceptable. The difference in log-transformed means and 95% confidence intervals are shown in Table 7.1. The

Tissue Block	Mean Log difference	Confidence interval
Rat Block 1	0.00283	$(-0.00402, 0.00968)$
Rat Block 2	0.00483	$(-0.000649, 0.0103)$
Rat Block 3	0.00505	$(-0.00137, 0.0115)$
Human Block	-0.000848	$(-0.00313, 0.00144)$

Table 7.1: Difference in means of log-transformed width measurements between unregistered and registered slides. Confidence interval represents a 95% probability of being in the given interval. The interval representing deformation of up to 5% is $(-0.0488, 0.0513)$ for all tissue blocks. This interval contains all the above confidence intervals, suggesting that the tissue has been deformed by no more than 5% by the registration process.

confidence intervals are all contained within the above effect size interval, and therefore the deformation incurred is substantially below 5%. This suggests that any deformation incurred by the registration process has not significantly changed the mean fibre width.

To ensure that the distribution of widths was unchanged by the registration process, a bootstrap method was used. For each tissue block, random samples of 1000 values were taken from the set of widths in each of the registered and unregistered slides. For each pair of samples, the Kolmogorov-Smirnov test statistic was calculated. The distribution of these values was used to determine the 90% confidence interval for the actual value of the Kolmogorov-Smirnov test statistic. If the confidence interval was found to lie below the critical value for the Kolmogorov-Smirnov test at a 95% confidence level, the distributions were deemed to be equivalent. The Kolmogorov-Smirnov test statistic is defined to be

$$D_{n,m} := \sup_x |F_{u,n}(x) - F_{r,m}(x)|,$$

where $F_{u,n}(x)$ is the empirical CDF of a sample of n values from the unregistered slide widths and $F_{r,m}(x)$ is the empirical CDF of a sample of m values form the registered slide widths. The probability of interest is

$$P\left(\sqrt{\frac{nm}{n+m}} D_{n,m} < c(\alpha)\right),$$

where $c(\alpha)$ is the critical value for the Kolmogorov-Smirnov test at significance level α.

Tissue Block	n	m	$K_{n,m}$	$\mathrm{P}(K_{1000,1000} > 1.36)$
Rat block 1	242856	239319	0.96	0.0496
Rat block 2	453123	449007	1.56	0.0514
Rat block 3	404011	398464	1.32	0.0456
Human block	2321528	2153015	1.00	0.0468

Table 7.2: Value of the normalised Kolmogorov test statistic, $K_{n,m} = \sqrt{nm/n+m}D_{n,m}$, for comparing registered and unregistered slides. The probabilities were generated by bootstrap sampling to obtain a distribution of $K_{1000,1000}$. The critical value $c(0.05) = 1.36$ represents a 95% confidence interval for the normalised test statistic. The probabilities show that the 90% confidence interval for the value of $K_{m,n}$ in each pair of datasets is contained within the 95% confidence interval, suggesting that the pairs of values represent similar distributions.

The critical value used in the equivalence test was $c(0.05) = 1.36$ [48]. The values of these tests are summarised in Table 7.2. The table shows that the above probability is greater than 90% for all tissue blocks, strongly suggesting that the distribution of widths is similar between registered and unregistered tissue blocks.

While the foregoing confirms that the tissue blocks have not been substantially deformed in general, it is also necessary to determine if any systematic deformation has occurred. Systematic deformation would be the case where a slide has been inaccurately deformed by the registration process, causing all slides registered to the given slide to be deformed in the same manner. This systematic deformation would be observable as a substantial increase or decrease in fibre width toward either end of the stack. The graphs in Figure 7.2 show the difference in the mean of the log-transformed widths between registered and unregistered slides as a function of slide position. The quadratic line of best fit for each set is shown in red. The quadratic fit was used here as opposed to linear fit to allow for the possibility that the systematic distortion is in the same direction at each end, meaning that the values can increase or decrease at either end independently. The dashed lines indicate the values for a deformation of $\pm 5\%$. In all cases the quadratic lines lie within these bounds, suggesting that any systematic deformation is sufficiently small to be within a 5% effect size, so the systematic distortion was deemed insignificant. There does appear to be some systematic deformation in all tissue blocks within these bounds, which suggests that larger stacks could incur a significant amount of systematic

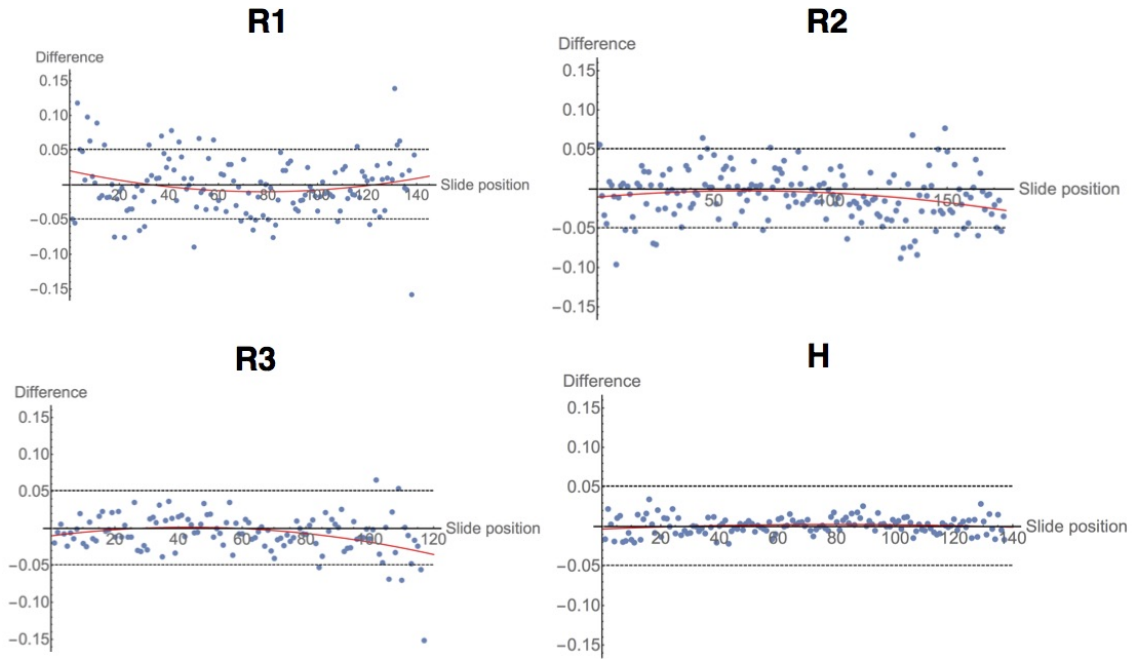

Figure 7.2: Difference in mean log-transformed fibre widths between registered and unregistered tissue blocks by slide. Dashed lines indicate a 5% change in width, red lines indicate the quadratic line of best fit. All lines of best fit remain within the 5% deformation bounds, suggesting that any deformation trend toward the ends of the tissue blocks is insignificant. Additionally, the vast majority of points lie within the 5% bounds, suggesting that even at an individual slide level the deformation is not significant.

deformation.

Comparison between the original tissue blocks and the *in silico* reconstructions are shown in Figures 7.3–7.6. This shows that the general shape of the tissue blocks has been largely maintained in the reconstruction, with slight differences observable at the edges of the tissue. Furthermore, the blood vessels in the rat tissue blocks, visible as darker areas in the original tissue and red in the reconstructed tissue, have also been maintained in the reconstruction. These two similarities between tissue and *in silico* reconstruction suggest that the general shape of the reconstruction represents that of the original tissue blocks.

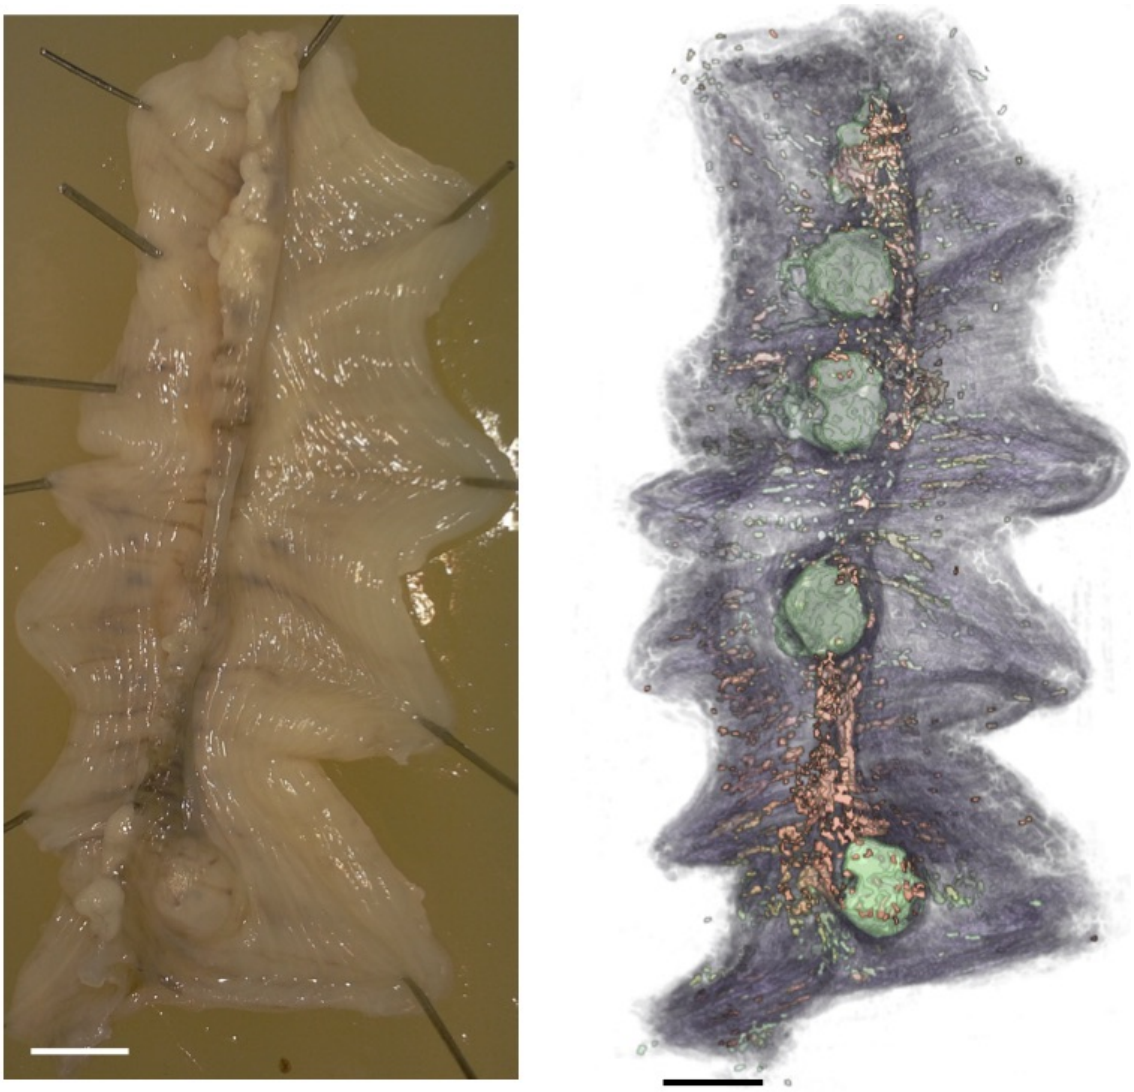

Figure 7.3: Comparison of the general shape of the reconstructed tissue with the original tissue in rat block 1. Left: Image of the original tissue block taken between fixation and embedding in paraffin. Right: Representation of the *in silico* reconstruction of the tissue block, where red indicates blood vessels, green indicates implantation sites, and purple indicates smooth muscle tissue. Here the general shape of the tissue and positioning of blood vessels appears to align with the original tissue block, suggesting that the reconstruction is a good representation of the overall shape of the tissue. Scale bars represent 5 mm.

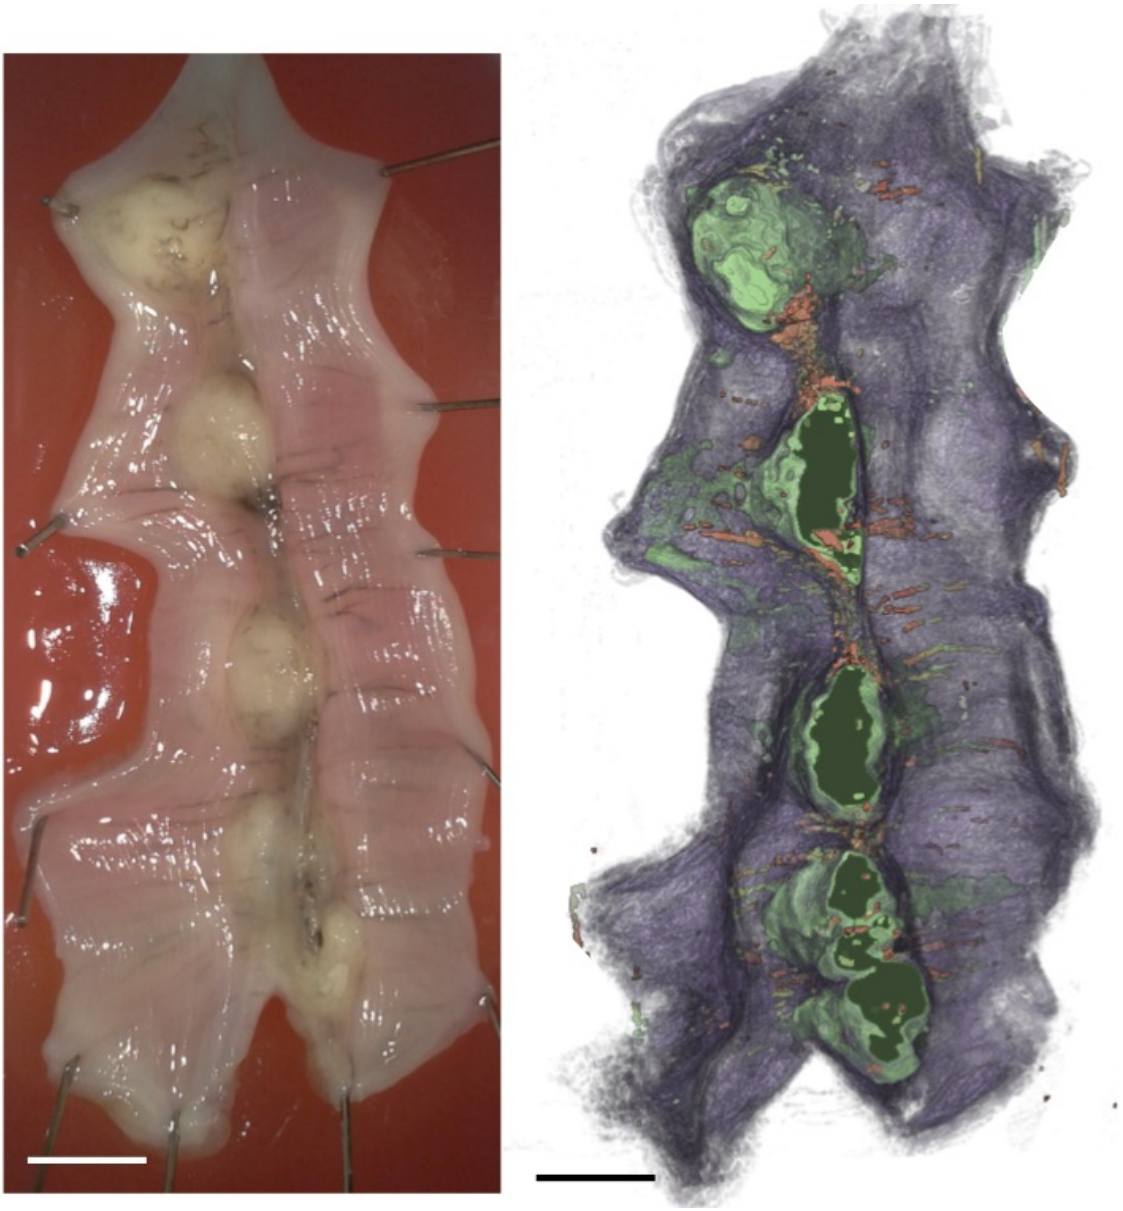

Figure 7.4: Comparison of the general shape of the reconstructed tissue with the original tissue in rat block 2. Left: Image of the original tissue block taken between fixation and embedding in paraffin. Right: Representation of the *in silico* reconstruction of the tissue block, where red indicates blood vessels, green indicates implantation sites, and purple indicates smooth muscle tissue. Here the general shape of the tissue and positioning of blood vessels appears to align with the original tissue block, suggesting that the reconstruction is a good representation of the overall shape of the tissue. Scale bars represent 5 mm.

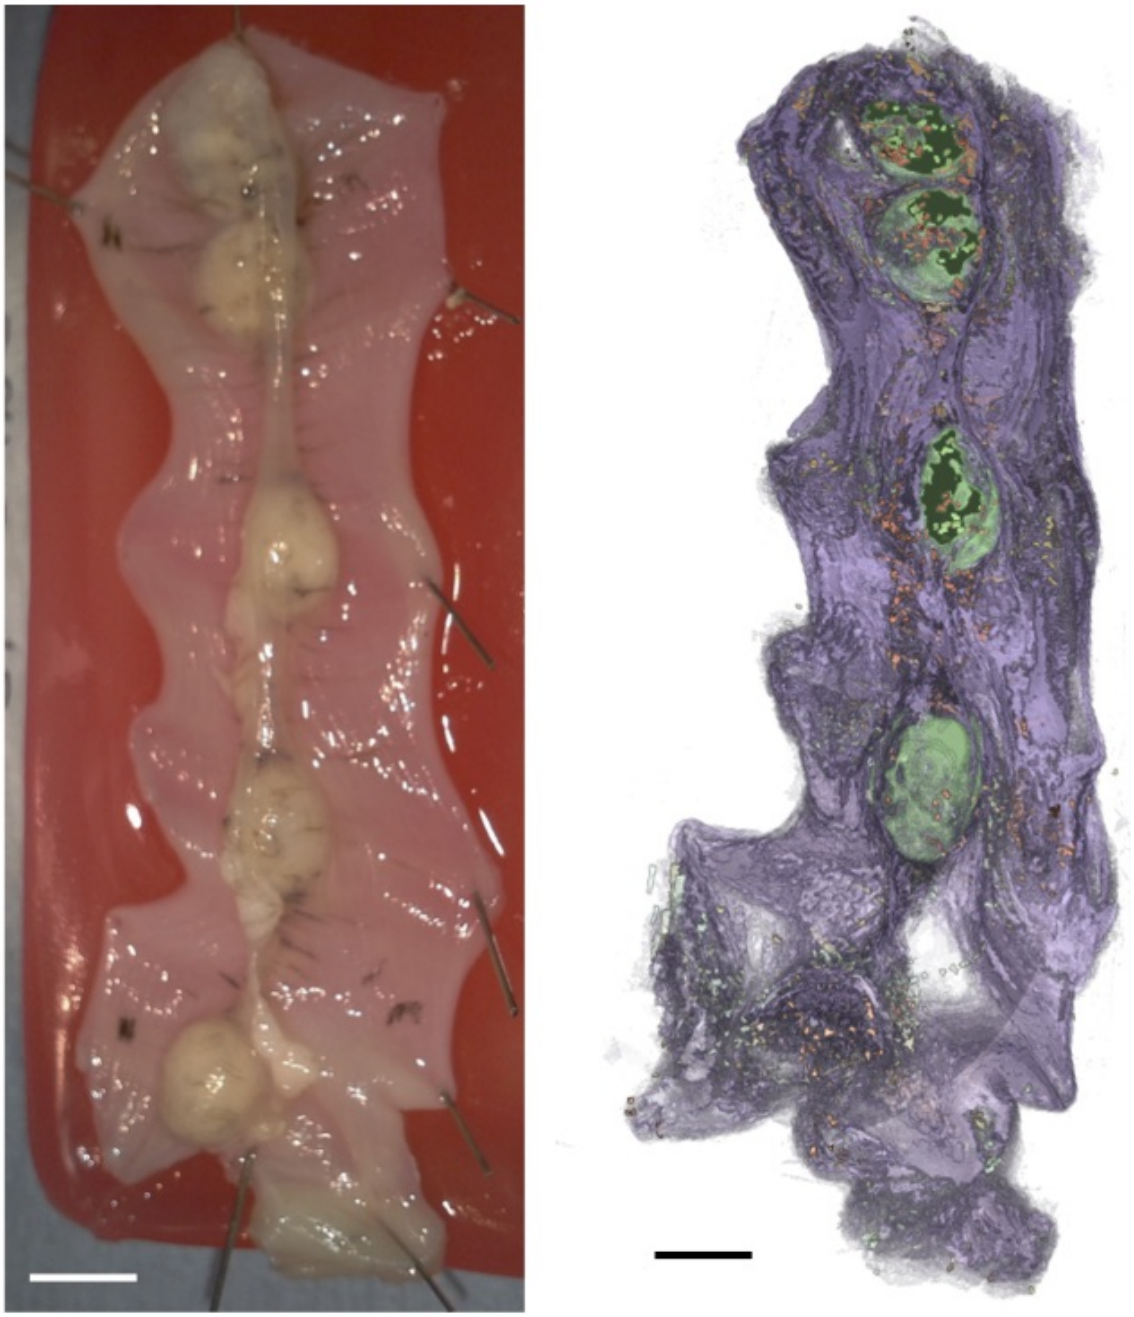

Figure 7.5: Comparison of the general shape of the reconstructed tissue with the original tissue in rat block 3. Left: Image of the original tissue block taken between fixation and embedding in paraffin. Right: Representation of the *in silico* reconstruction of the tissue block, where red indicates blood vessels, green indicates implantation sites, and purple indicates smooth muscle tissue. Here the general shape of the tissue and positioning of blood vessels appears to align with the original tissue block, suggesting that the reconstruction is a good representation of the overall shape of the tissue. Scale bars represent 5 mm.

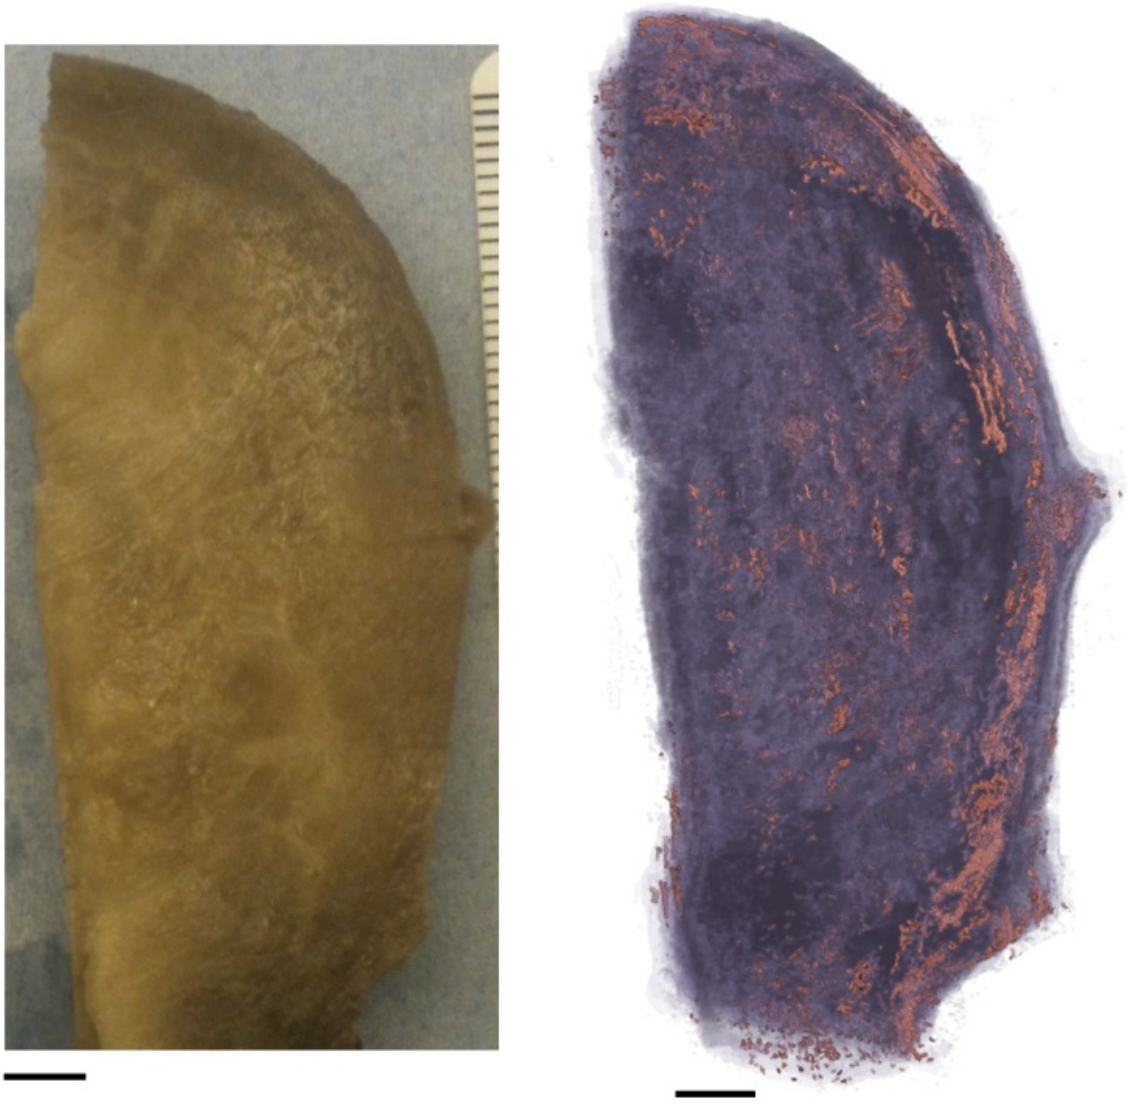

Figure 7.6: Comparison of the general shape of the reconstructed tissue with the original tissue in the human tissue block. Left: Image of the original tissue block taken between fixation and embedding in paraffin. Right: Representation of the *in silico* reconstruction of the tissue block, where red indicates blood vessels, and purple indicates smooth muscle tissue. Here the general shape of the tissue appears to align with the original tissue block, suggesting that the reconstruction is a good representation of the overall shape of the tissue. Scale bars represent 5 mm.

7.2 Analysis of fibre widths

The methods for obtaining fibre widths from the three-dimensional volume have been described in Chapter 3. In all blocks, the fibre width was measured at nodes of a randomly positioned lattice in a randomly selected direction perpendicular to the direction of the given voxel. The separation of nodes in the lattice was set to 4 voxels in the slide plane, which represents 240 μm in the first rat tissue block and 190 μm in all other tissue blocks, and 1 voxel in the z-direction, which represents 60 μm in the first rat tissue block and 47.5 μm in all other tissue blocks. The higher sampling rate in the z-direction was required because the depth of the tissue blocks were typically $\sim 10\times$ smaller than the other dimensions. The distributions of measured fibre widths in the rat and human tissue blocks are compared in Figure 7.7. The median, 95% confidence interval, and sample size of each of these data sets for each tissue block are listed in Table 7.3. These measurements were compared in a similar manner to the fibre width comparison performed in the previous section; the data was log-transformed, and pairwise comparisons were performed on the log-transformed data sets to test for substantial differences. The difference in means of the log transformed data between widths in each pair of tissue blocks is shown in Table 7.4, with 95% confidence intervals. This shows that the difference in means between all tissue blocks is non-zero with 95% confidence, which suggests that the fibre widths vary between all tissue blocks.

The magnitude of the difference in widths between two tissue blocks A and B as a proportion of the overall width of A is given by

$$1 - \exp(\mu_A - \mu_B)$$

where μ_A and μ_B are the means of the log-transformed widths for the tissue blocks A and B. The difference in means between rat blocks 1 and 2 corresponds to an decrease in width of $5\% \pm 1\%$ from block 1 to block 2, which suggests that there is no substantial

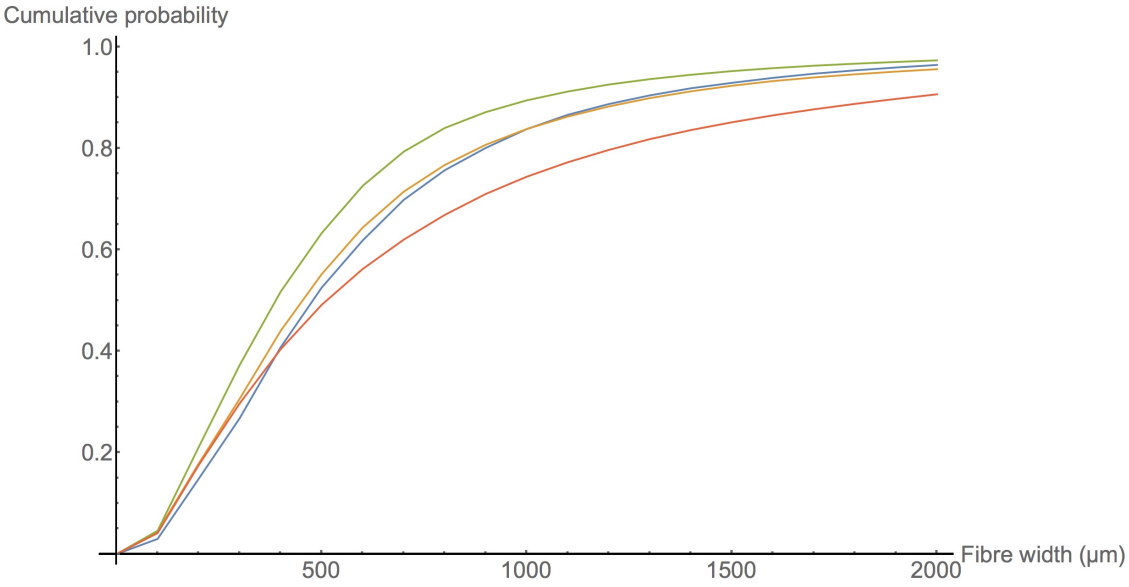

Figure 7.7: The empirical CDF of the distribution of fibre widths in each of the three rat uteri and in the human tissue block. This shows that the widths in the human tissue block are markedly higher than those in the rat tissue blocks, as indicated by the shift to the right of the distribution. Additionally, the third rat tissue block appears to have smaller widths than the other two rat tissue blocks, as indicated by a shift to the left of the distribution.

difference in fibre width between these two blocks. The difference in means between rat block 3 and rat blocks 1 and 2 correspond to a decrease in width of $19\% \pm 1\%$ from block 1 to block 3, and a decrease in width of $14\% \pm 1\%$ from block 2 to block 3. This suggests that rat block 2 is substantially different from the other two rat blocks. One confounding factor which might contribute to this difference, however, is the more complex artificial bounds on rat block 3, as described in Section 2.2. This effect is apparent in the analysis of fibre widths as a function of depth in the stack, which is covered in more depth below. The difference in means between the human tissue block and the rat tissue blocks correspond to a decrease in width of $10\% \pm 1\%$, $15\% \pm 1\%$, and $35\% \pm 1\%$ from the human block to rat blocks 1, 2, and 3 respectively. This suggests that there is a substantial difference between the human tissue block and the rat tissue blocks. However, the difference between the widths in the human block and rat block 1 is less than the difference between rat block 3 and rat block 1, which means that the difference observed between human and rat fibre

Tissue Block	Median ± CI (μm)	n
Rat Block 1	482 ± 2	91438
Rat Block 2	452 ± 2	189958
Rat Block 3	388 ± 2	166666
Human Block	513 ± 2	922534

Table 7.3: Median fibre width for each tissue block, error shown is the 95% confidence interval, rounded up to the nearest μm.

Tissue Block	Rat Block 1	Rat block 2	Rat block 3	Human block
Rat Block 1	-	0.051 ± 0.006	0.21 ± 0.006	-0.092 ± 0.006
Rat Block 2	-0.051 ± 0.006	-	0.15 ± 0.005	-0.14 ± 0.004
Rat Block 3	-0.21 ± 0.006	-0.15 ± 0.005	-	-0.29 ± 0.004
Human Block	0.092 ± 0.006	0.14 ± 0.004	0.29 ± 0.004	-

Table 7.4: The difference of means of the log-transformed fibre widths between each tissue block. Error shown is 95% confidence interval.

widths may be within the bounds of variation between individuals.

The spatial distribution of fibre widths in the slide plane for each tissue block is shown in Figure 7.8. These images were generated by calculating the median width along the z-axis for each point in the grid given by the slide plane cross-section of the lattice and representing this median value according to the LUT shown. In the rat tissue blocks, these indicate that the fibre width is variable throughout the tissue, with some areas of the tissue having larger widths, but no apparent correlation with any specific anatomical feature. In the human tissue block, there does appear to be an area of wider fibres toward the perimetrium.

A comparison of widths by depth in the z-direction for the rat tissue blocks is shown in Figure 7.9. There appears to be a difference in width between the higher and lower fibres in the block, with widths toward the circular end being higher. This suggests that fibres are wider in the circular layer than in the longitudinal layer. The orientation of the human tissue block motivates the comparison of widths along the arcs shown in Figure 7.10**A**. These arcs represent steps away from the endometrium, and hence can be used to compare inner and outer layers. The comparison of widths measured along these arcs is shown in Figure 7.10**B**. This shows that the median width increases width

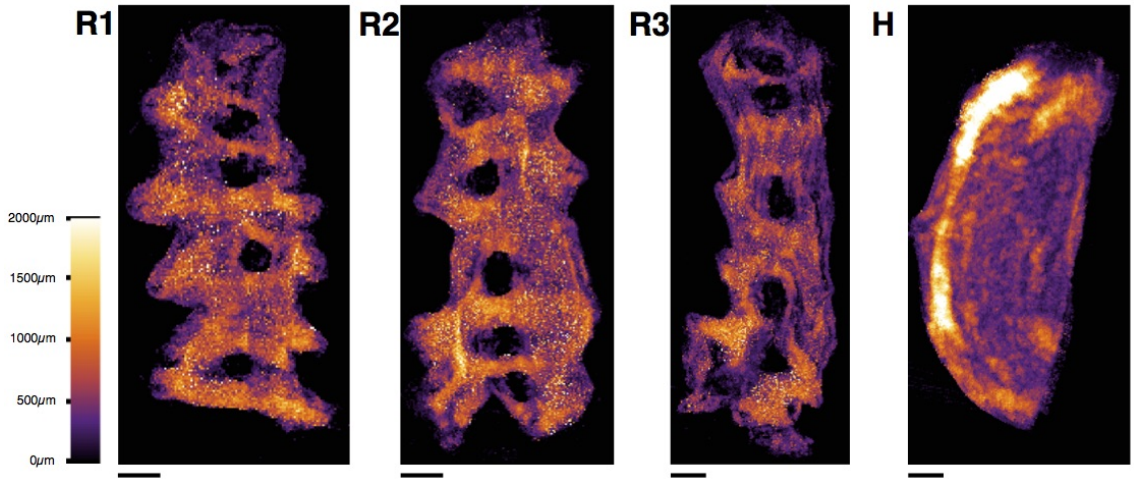

Figure 7.8: Mean fibre widths along the z-axis for each tissue block. The images were generated by taking the median value of recorded widths in the lattice along the z-axis, and coloured according to the LUT displayed on the left. **R1**, **R2**, **R3**: Rat tissue blocks 1, 2, and 3 respectively. **H**: Human tissue block. Scale bars represent 5 mm.

distance from the endometrium, with a large spike toward the perimetrium. It appears that the fibres are wider in the longitudinal layer than in the circular layer, with larger fibres occupying the outermost regions of the longitudinal layer. This stands in contrast to the observations made in the rat myometrium, which shows larger fibres in the circular layer and thinner fibres in the longitudinal layer.

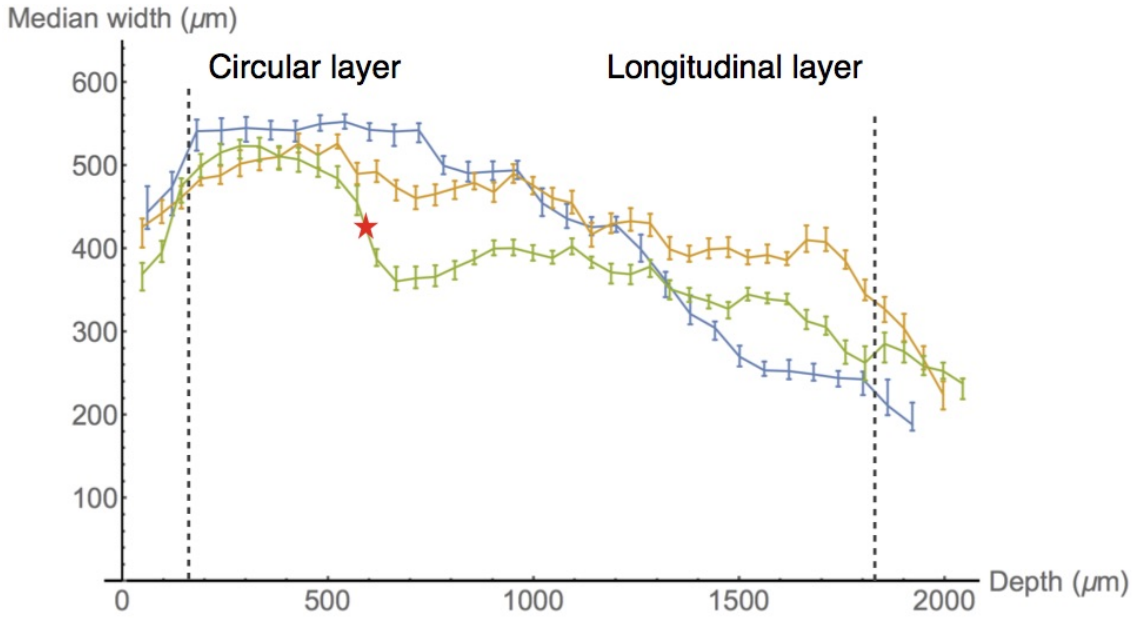

Figure 7.9: The median widths in the rat tissue blocks as a function of z. Each point represents a slide in the *in silico* reconstruction, sorted in increasing slide number such that the lower values represent slides containing mostly circular smooth muscle tissue, while the represent slides containing mostly longitudinal smooth muscle tissue. Blue represents tissue block 1, orange represents tissue block 2, and green represents tissue block 3. The general trend is a reduction in width from circular to smooth muscle, suggesting that fibres are thicker in the circular layer.

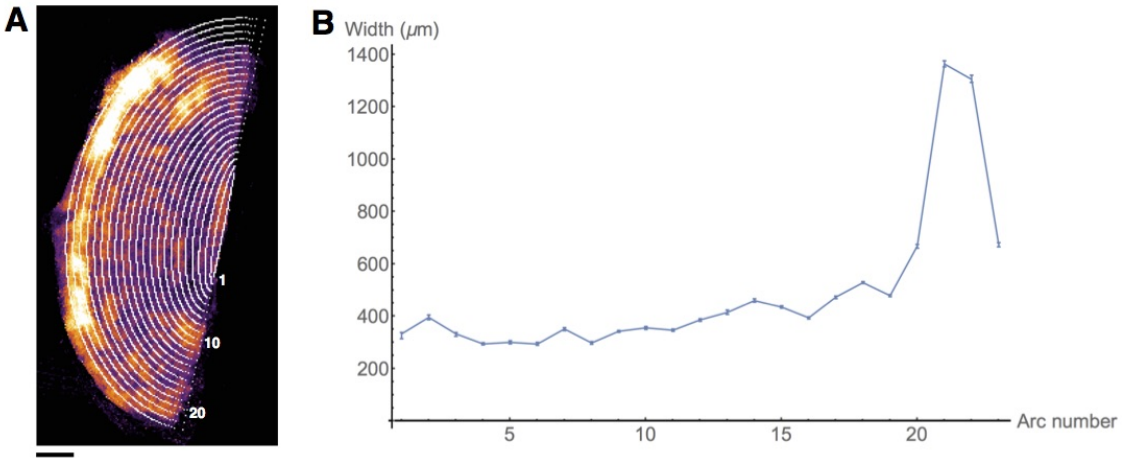

Figure 7.10: Comparison of fibre widths along arcs located between the perimetrium and endometrium. **A**: The arcs used to compare widths at varying distance from the endometrium. **B**: Median width on each arc. The large peak toward the perimetrium is the layer of subserosal fibres observed by Weiss *et al.* [111], and the values shown here indicate that the fibres are approximately 3 times wider than in the rest of the tissue. Scale bar represents 5 mm.

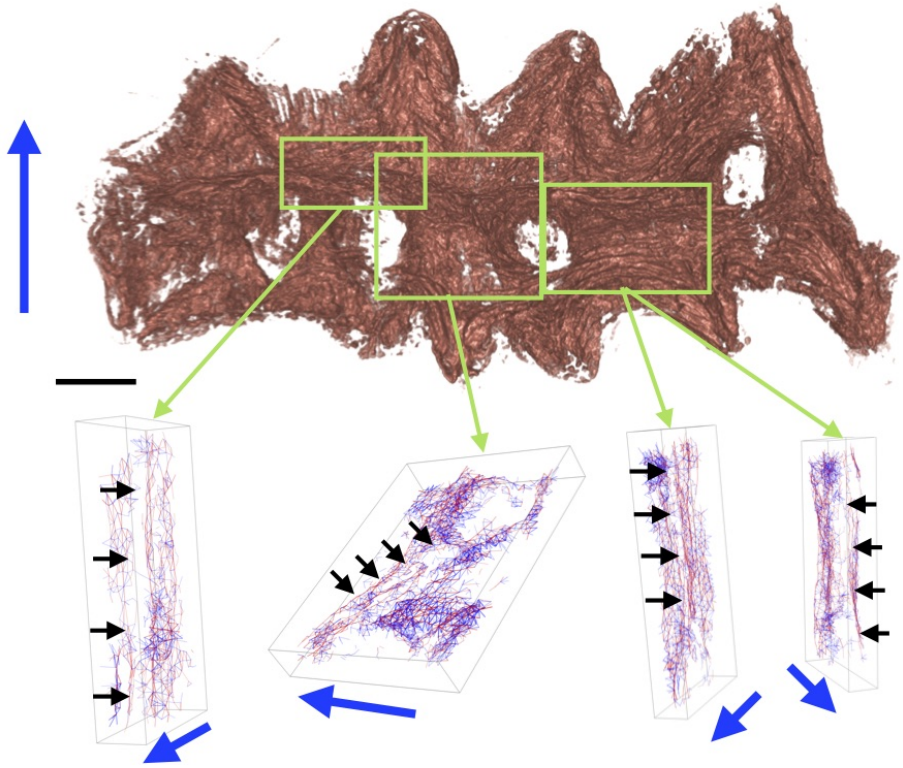

Figure 7.11: Gaps in the network between the longitudinal mesometrial border fibres and the longitudinal layer of the myometrium in tissue block 1. Black arrows indicate a boundary between these two features where little or no connectivity was detected. This suggests that the fibres in the longitudinal layer along the mesometrial border are poorly connected to the rest of the longitudinal layer at the points indicated. Blue arrow indicates orientation of the networks relative to the reconstruction. Scale bar represents 5 mm.

7.3 Analysis of connectivity

Connectivity was measured by generating networks of connected points, as described in Chapter 3. The networks were generated by performing a watershed-type algorithm, where the pools are represented by nodes in the network, and the points where these pools make contact are represented by arcs in the network. The average direction vector and size of each pool was recorded, and each arc was assigned an angle based on the difference between the direction of the given arc and the average direction of the two pools represented by the nodes which the arc connected. An angle below 45° assigned in this manner represents a connection along the fibre direction, whereas an angle above 45° represents a connection

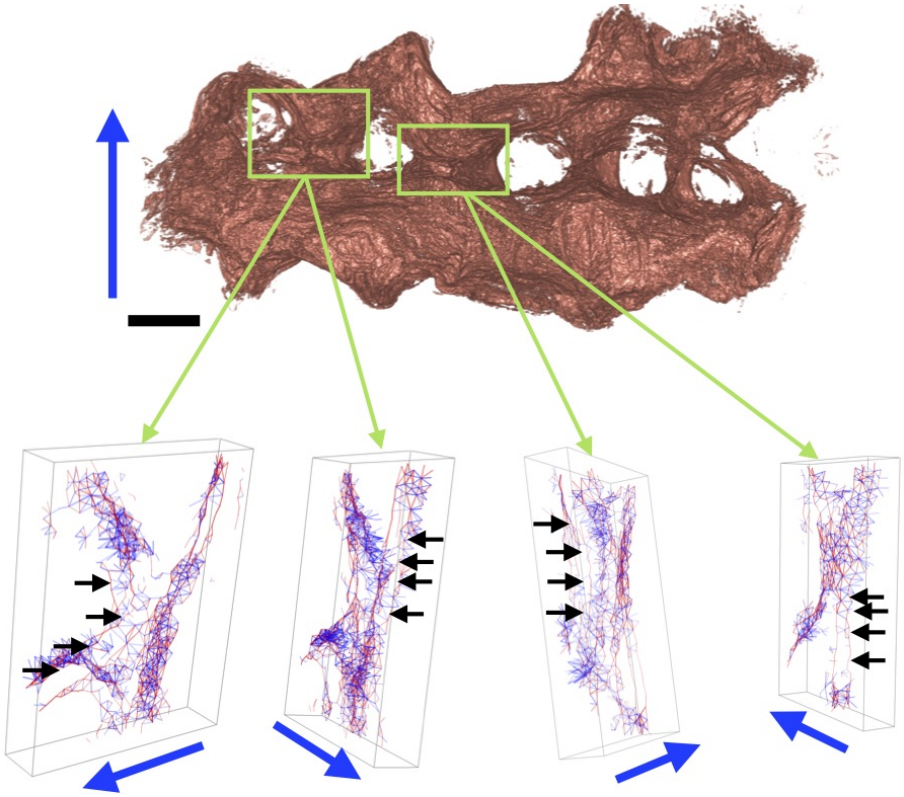

Figure 7.12: Gaps in the network between the longitudinal mesometrial border fibres and the longitudinal layer of the myometrium in tissue block 2. Black arrows indicate a boundary between these two features where little or no connectivity was detected. This suggests that the fibres in the longitudinal layer along the mesometrial border are poorly connected to the rest of the longitudinal layer at the points indicated. Blue arrow indicates orientation of the networks relative to the reconstruction. Scale bar represents 5 mm.

away from the fibre direction. This angle will be used below to differentiate between arcs representing connections between fibres and arcs representing the actual fibres. The arcs were also assigned weights given by the maximum weight of all points joining the two pools represented by the end nodes of a given arc. This weighting quantifies the strength of the connection represented by the arc.

The aim of the present section is to determine the connectivity properties of the tissue from the network graph. Specifically, the following properties will be considered. The connectivity across the mesometrial borders will be assessed visually to determine if there is a change in connectivity in this location, which could provide a possible explanation for the

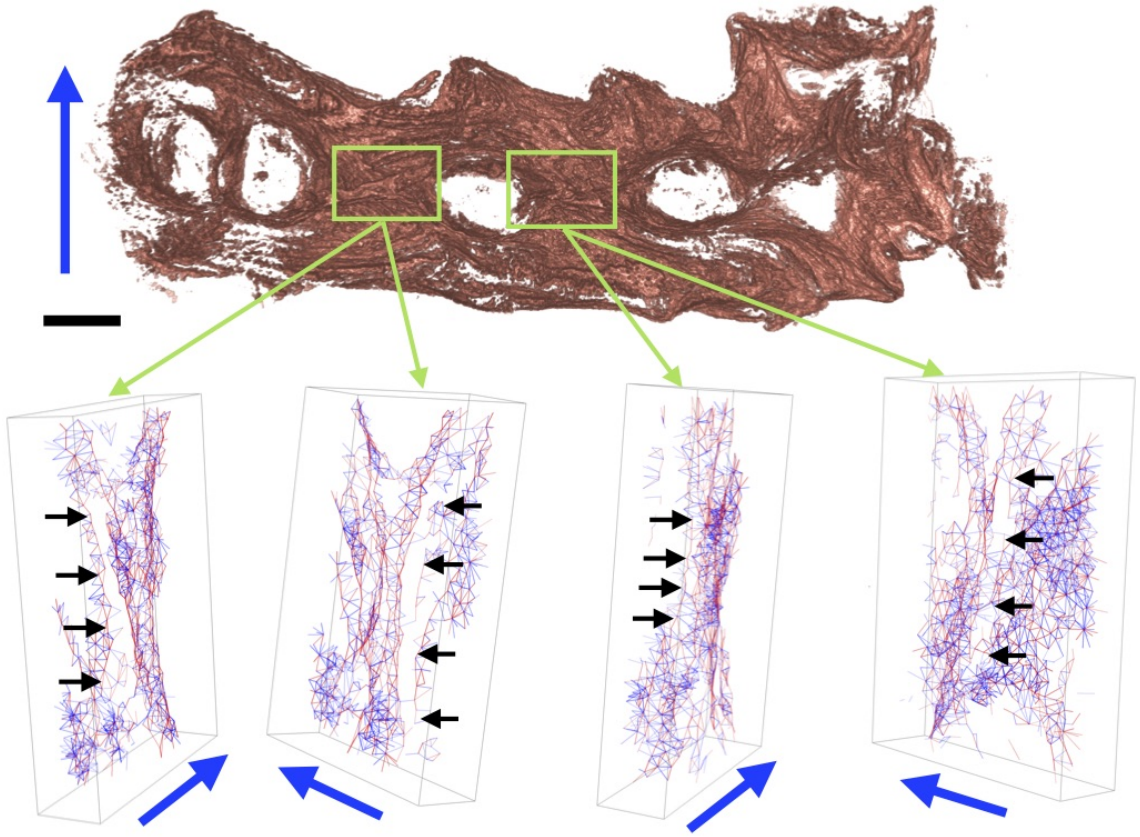

Figure 7.13: Gaps in the network between the longitudinal mesometrial border fibres and the longitudinal layer of the myometrium in tissue block 3. Black arrows indicate a boundary between these two features where little or no connectivity was detected. This suggests that the fibres in the longitudinal layer along the mesometrial border are poorly connected to the rest of the longitudinal layer at the points indicated. Blue arrow indicates orientation of the networks relative to the reconstruction. Scale bar represents 5 mm.

conduction blocks noted in Chapter 6. The connectivity between circular and longitudinal layers will be assessed, which could provide a possible interface for communication between the layers, as suggested in [86].

In Chapter 6 it was observed that action potentials failed to propagate across the mesometrial border in most cases. The network model was used to determine if this is caused by a change in connectivity across the mesometrial border. In the following, the sub-networks associated with selected sub-volumes of the reconstructed tissue will be assessed visually for levels of connectivity. These selected sub-volumes represent locations

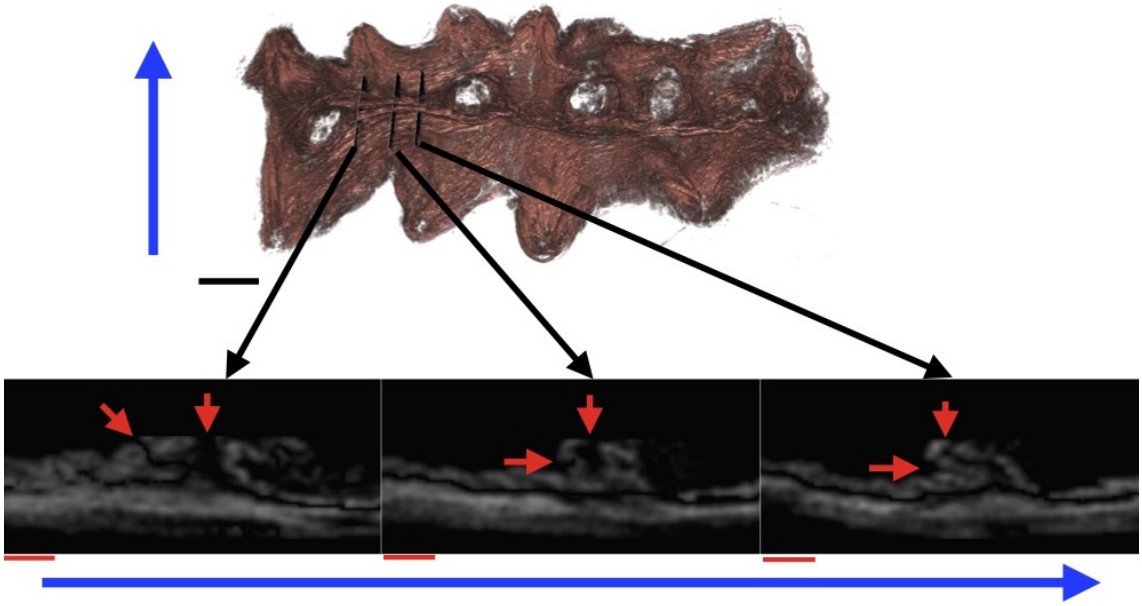

Figure 7.14: Resectioning of the reconstructed tissue showing boundaries of low connectivity between the longitudinal fibres and the mesometrial border in tissue block 1. The lower panels show transverse resections of the reconstruction in the upper panel. Red arrows indicate the boundaries of low connectivity between mesometrial border fibres and the rest of the longitudinal layer. Blue arrows indicate orientation of the resections relative to the reconstruction. Black scale bar represents 5 mm, red scale bars represent 1 mm.

where excitation failed to propagate, as described in Chapter 6. Figures 7.11–7.13 show the level of connectivity across the mesometrial border at selected locations. The sub-networks are colour-coded, with red indicating connections which represent fibres, and blue indicating connections between adjacent fibres. The level of opacity indicates how strong the connection is, as determined by the weighting described above. In each of the sub-networks there is a boundary of low connectivity between the longitudinal fibres along the mesometrial border and the rest of the longitudinal layer, indicated by the arrows. These boundaries are visible as areas of space which contain little or no connecting arcs, indicating that there is a low level of connectivity across the boundaries. It thus appears that the longitudinal fibres on the mesometrial border are poorly connected to the rest of the longitudinal layer at these selected locations.

Figures 7.14–7.16 show examples of these boundaries as seen in the reconstructed tis-

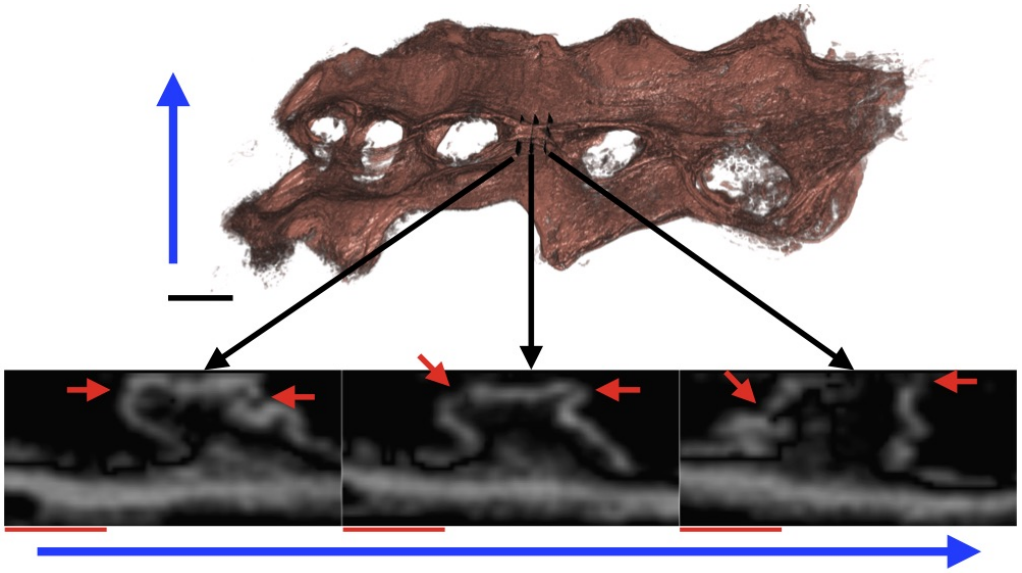

Figure 7.15: Resectioning of the reconstructed tissue showing boundaries of low connectivity between the longitudinal fibres and the mesometrial border in tissue block 2. The lower panels show transverse resections of the reconstruction in the upper panel. Red arrows indicate the boundaries of low connectivity between mesometrial border fibres and the rest of the longitudinal layer. Blue arrows indicate orientation of the resections relative to the reconstruction. Black scale bar represents 5 mm, red scale bars represent 1 mm.

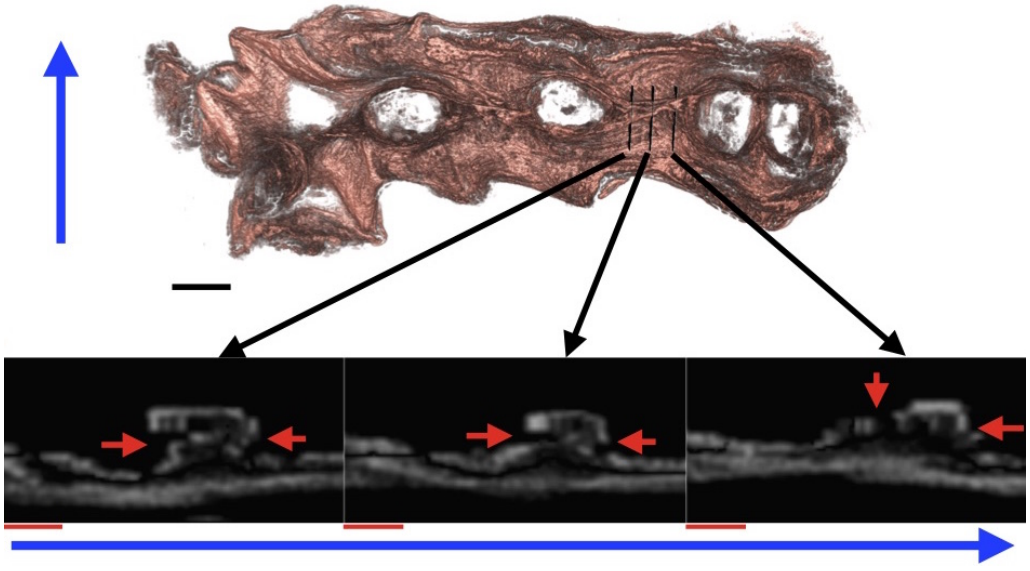

Figure 7.16: Resectioning of the reconstructed tissue showing boundaries of low connectivity between the longitudinal fibres and the mesometrial border in tissue block 3. The lower panels show transverse resections of the reconstruction in the upper panel. Red arrows indicate the boundaries of the low connectivity mesometrial border fibres and the rest of the longitudinal layer. Blue arrows indicate orientation of the resections relative to the reconstruction. Black scale bar represents 5 mm, red scale bars represent 1 mm.

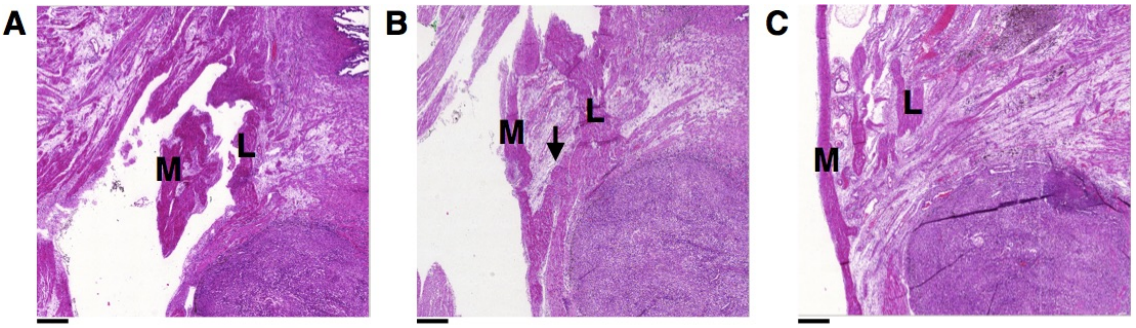

Figure 7.17: Histology of the tissue in Figure 7.14, separated by ~ 150 μm. **A** shows the start of the mesometrial border fibres as distinct from the longitudinal layer. **B** shows the mesometrial border fibres and longitudinal fibres remaining largely distinct, with a connection being made as indicated by the arrow. **C** shows the mesometrial border fibres now running the full length of the image, with longitudinal fibres receding. There is little contact between longitudinal fibres and fibres along the mesometrial border in this sequence of slides, which represent the transition between longitudinal fibres and mesometrial border fibres in the given location. **M**: Mesometrial border fibres, **L**: longitudinal fibres. Scale bars represent 500 μm.

sue. The boundaries are visible as gaps between the longitudinal layer and the mesometrial border, which supports the hypothesis that there is a functional boundary of low connectivity between the mesometrial border and the longitudinal layer. The histological images in Figure 7.17–7.19 show slides from the reconstructions shown in Figures 7.14–7.16. Specifically, they show a clear separation of fibres in the transition from longitudinal fibres to mesometrial border fibres, again indicative of poor connection between the mesometrial border and the general population of fibres in the longitudinal layer. The paucity of connectivity between longitudinal fibres along the mesometrial border and those away from the mesometrial border could be the cause of the propagation block observed in Chapter 6.

The network model was used to determine if any connections exist between circular and longitudinal layers. This was performed by first classifying the nodes of the network as circular and longitudinal based on the average direction; pools with direction less than 45° from the (x, z)-plane were classified as being circular, all others were classified as being longitudinal. This classification is not necessarily an accurate representation of the tissue, because the tissue was pinned, which stretched some of the longitudinal fibres into an orientation which would be considered circular by this classification. Given this classi-

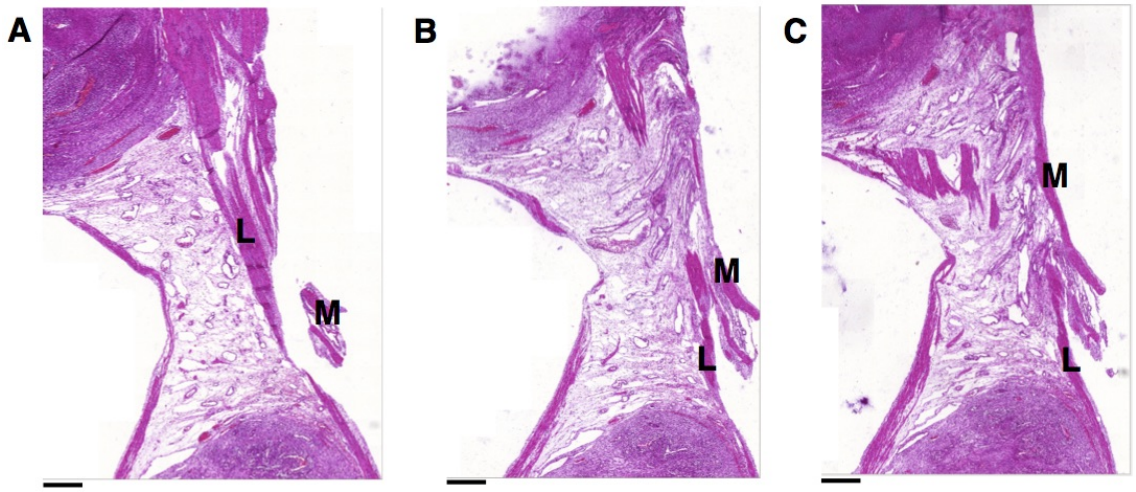

Figure 7.18: Histology of the tissue in Figure 7.15, separated by ~ 100 μm. **A** shows the start of the mesometrial border fibres as distinct from the longitudinal layer. **B** shows the mesometrial border fibres and longitudinal fibres remaining distinct. **C** shows the mesometrial border fibres now running the almost completely between the two implantation sites, with longitudinal fibres receding. There is little contact between longitudinal fibres and fibres along the mesometrial border in this sequence of slides, which represent the transition between longitudinal fibres and mesometrial border fibres in the given location. **M**: Mesometrial border fibres, **L**: longitudinal fibres. Scale bars represent 500 μm.

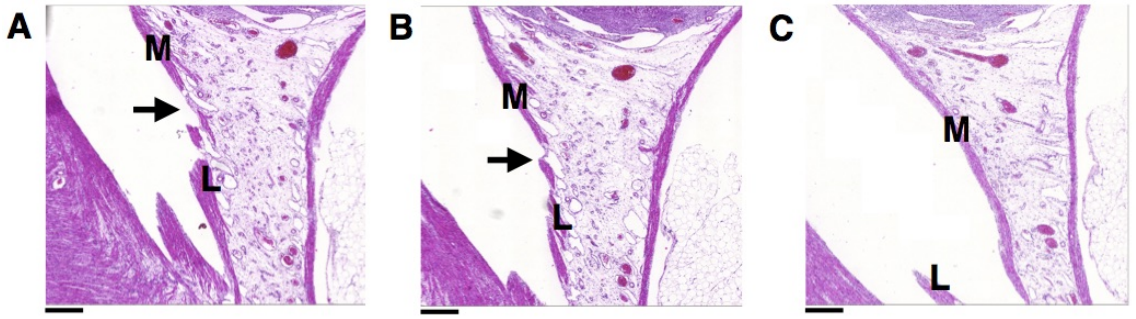

Figure 7.19: Histology of the tissue in Figure 7.14, separated by ~ 100 μm. **A** shows the start of the mesometrial border fibres as distinct from the longitudinal layer. **B** shows the mesometrial border fibres and longitudinal fibres remaining distinct. **C** shows the mesometrial border fibres now running the full length of the image, with longitudinal fibres no longer present at that proximity to the mesometrial border. There is little contact between longitudinal fibres and fibres along the mesometrial border in this sequence of slides, which represent the transition between longitudinal fibres and mesometrial border fibres in the given location. **M**: Mesometrial border fibres, **L**: longitudinal fibres. Scale bars represent 500 μm.

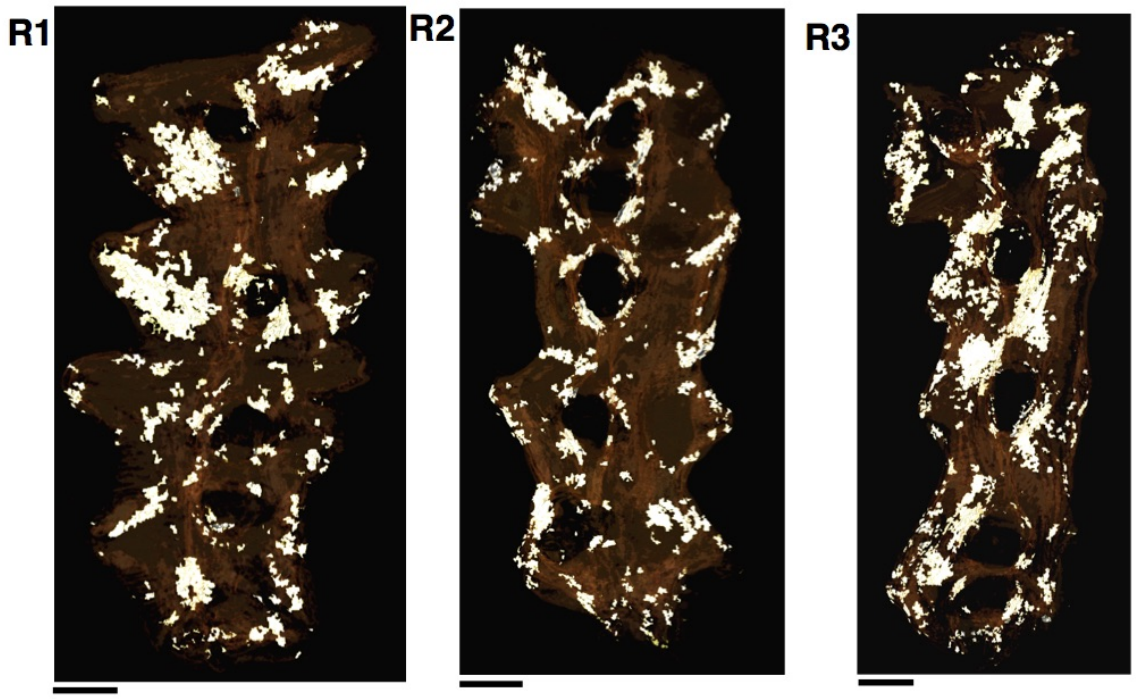

Figure 7.20: The spatial distribution of detected bridges between the circular and longitudinal layers of the myometrium in each of the rat tissue blocks. The areas highlighted correspond to the nodes of the network which have been classified as either circular or longitudinal and are connected to a node in the other category. This method of bridge detection highlights all possible points which could be bridges between the layers, but may also include points which do not correspond to a bridge between the two layers. Scale bars represent 5 mm.

fication, each arc connecting two nodes of differing categories was recorded: the spatial distribution of these arcs is shown in Figure 7.20. This detection algorithm was used to determine all possible locations of connections between circular and longitudinal layers; however, in order to ensure these connections were genuine, it was necessary to verify each connection manually by examining the histological slides. The clustering which can be observed in Figure 7.20 enabled the connections to be reviewed by cluster, and the result of this manual observation are illustrated in Figure 7.21, where confirmed connections have been circled. Some examples of the histology of clusters which were rejected are shown in Figure 7.22.

The green circles in Figure 7.21 indicate clusters which do appear to represent connections between the two layers, and some examples of the histological slides for these clusters

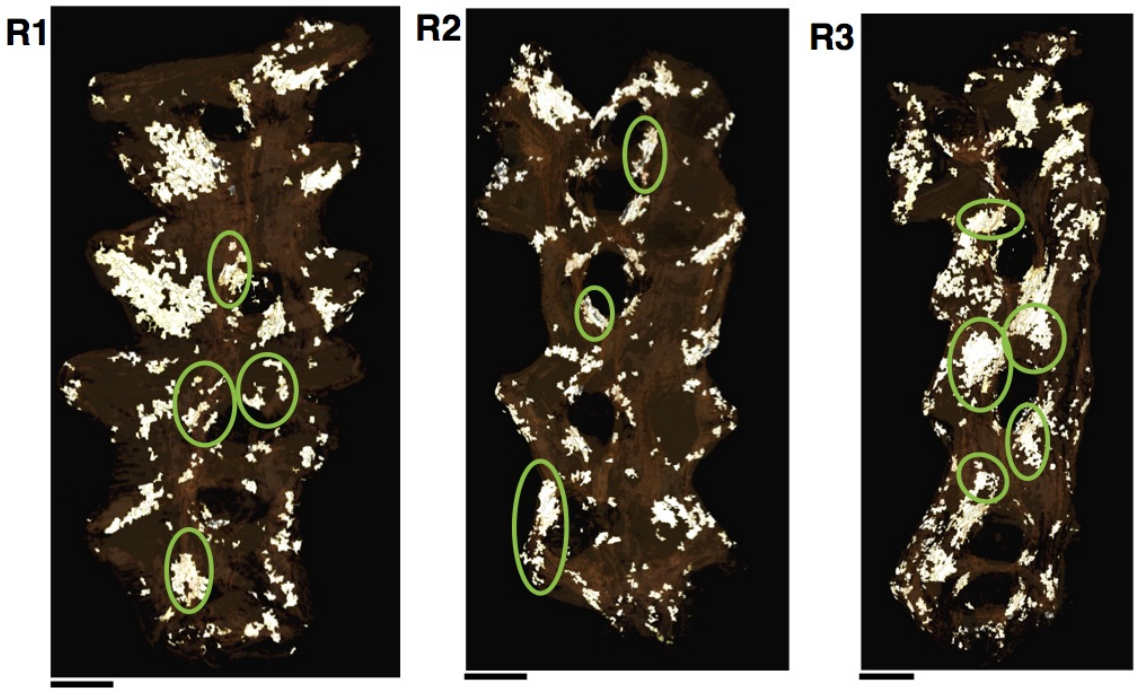

Figure 7.21: Clusters which have been verified as bridges between the circular and longitudinal layers in each rat tissue block. The circled clusters have been verified by examining the histological slides to determine if there is clear evidence of a bridge between the circular and longitudinal layers within the volume indicated by the clusters. Scale bars represent 5 mm.

are shown in Figure 7.23. The connections are generally formed from bundles 20–100 μm thick, with several such bundles located at some clusters, as illustrated in Figure 7.23. While the clusters highlighted in Figure 7.21 exhibit a spatial distribution of connections which lie close to the mesometrial border, there may be other connections not identified in this manner; some of the clusters marked as false positives may in fact be genuine, but without unambiguous verification of the connection have been mislabelled. Additionally, the clustering may be an artefact of the reconstruction process; an isolated connection is more likely to be lost through segmentation, for example, due to the relative low density of clearly directed smooth muscle tissue at that point. The spatial distribution of connections which have been detected, however, suggests that there are multiple locations at which the two layers connect.

In the human myometrium the inner layer is formed of circular fibres, while the outer

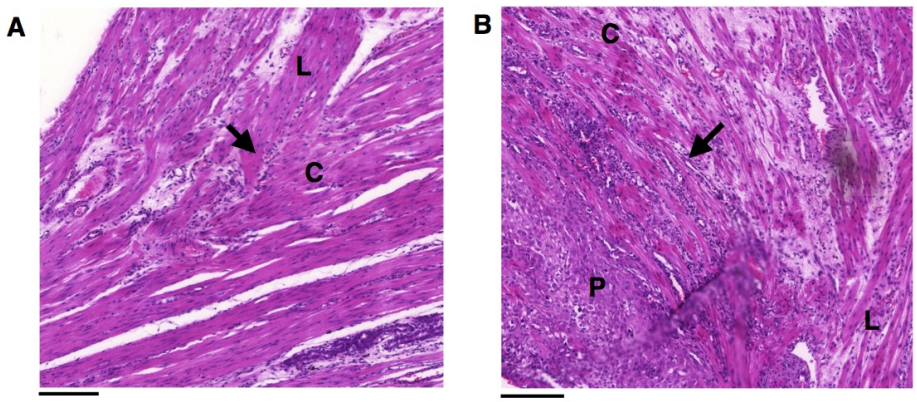

Figure 7.22: Areas where there is insufficient evidence to show the existence of bridges between the circular and longitudinal layer. **A**: The area indicated by the arrow corresponds to a bridge detected; however, close inspection of the area indicated reveals that there is a distinct boundary between the circular and longitudinal fibres. This area is taken from tissue in close proximity to the pins, which means that the layers are sufficiently distorted by the pinning process to cause the layers to have similar fibre direction at this point. This distortion reduces the effectiveness of the segmentation algorithm, allowing a connection between the fibres in the reconstructed tissue. **B**: The area indicated by the arrow corresponds to decidualized myometrium. While this may be the location of a bridge between the circular and longitudinal layers, these connections are also rejected because the connections occur within the placenta. **L**: longitudinal fibres, **C**: circular fibres, **P**: implantation site. Scale bars represent 200 μm.

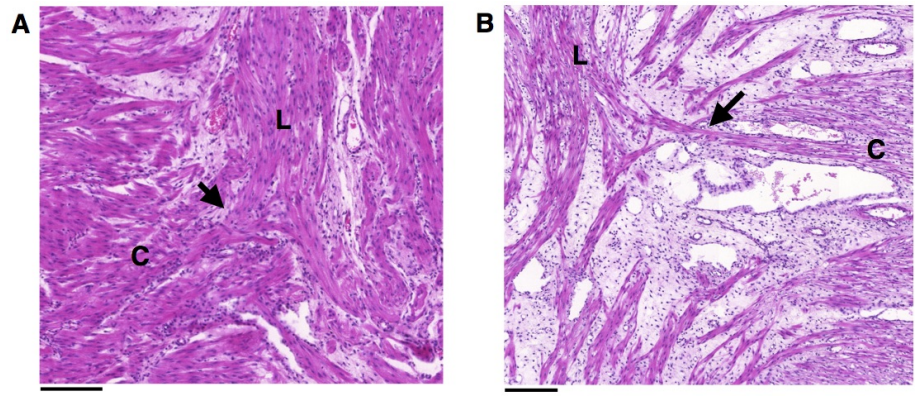

Figure 7.23: Examples of areas where bridges between the circular and longitudinal layer are present. **A**: The bridge indicated by the arrow shows a clear transition from circular to longitudinal layers. In this example there does not appear to be evidence of other bridges nearby. **B**: The bridge indicated by the arrow again shows a clear transition from circular to longitudinal fibres, but the connection is longer than in **A** ($\sim$ 1 mm). Additionally, similar structures can be seen near the indicated bridge, suggesting that there could be multiple connections in close proximity to this point. **L**: longitudinal fibres, **C**: Circular fibres. Scale bars represent 200 μm.

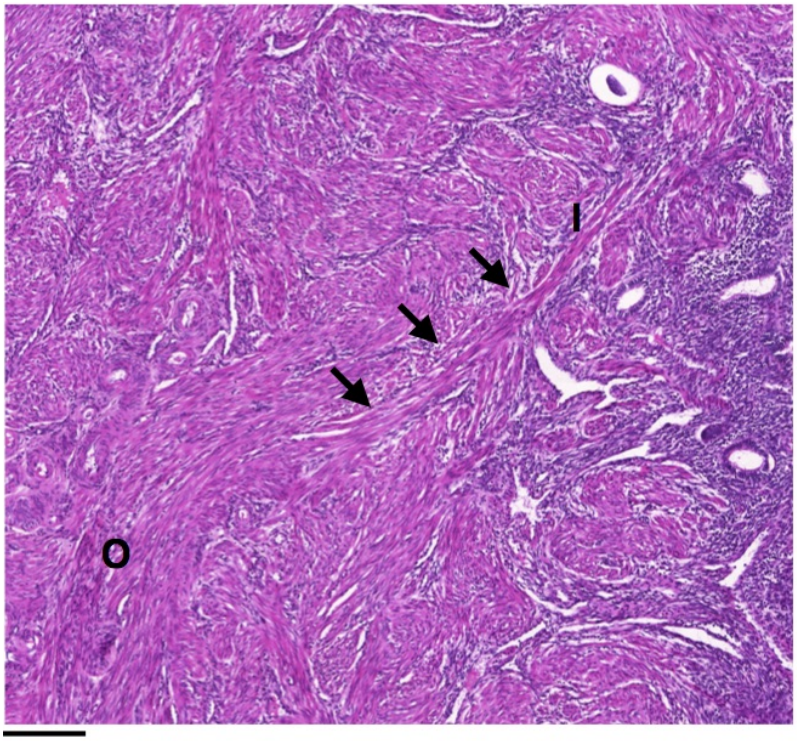

Figure 7.24: A bridge connecting the inner layer of the myometrium to the outer layer in the human uterus. Other structures such as this one could be identified at the histological level; however, accurate identification of the full extent of these bridges was not possible (see text for details). These bridges could underpin coordinated excitation between the inner and outer myometrium. Arrows indicate the connecting bridge. **I**: Inner myometrium, **O**: Outer myometrium. Scale bar represents 200 μm.

portion is more disordered, with the exception of the outermost layer, which contains longitudinal fibres [111]. An analogous structure to the connecting bridges found in rat uteri would be bridges joining the inner layer to the outer layer. Such structures are present in the human myometrium examined in this thesis, as demonstrated in Figure 7.24. The precise identification of the locations of these connective bridges in the human was complicated by the inherent structure of the myometrium: the disorderly arrangement of fibres in the outer myometrium leads to numerous instances of circular fibres in the outer myometrium. This prevented the accurate identification of the boundary between inner and outer myometrium. The separation of inner and outer myometrium is observable in MRI scans [19], which suggests that it would be possible to identify the boundary

between the layers accurately, and hence the fibres connecting them, by comparison of a reconstruction of the form presented in this thesis and MRI scans.

Chapter 8

Anatomy of the myometrium at pacemaker sites

In Chapter 6 fibres were observed extending into the placenta from the longitudinal layer in close proximity to excitation initiation points, as determined by analysis of multi-electrode array recordings. These longitudinal interpenetrating fibres could be the location of pacemaker activity in the myometrium. This chapter will focus on investigating the structural characteristics of these interpenetrating fibres. The *in silico* reconstruction described in Chapter 2 does not include these interpenetrating fibres because the resolution is too low to accurately represent them. In order to visualise the interpenetrating fibres in three dimensions, it was necessary to create high resolution reconstructions of the myometrium at each implantation site. In this chapter the methods for generating these reconstructions is described and the resulting structures are analysed. All code described in this chapter can be found at `https://figshare.com/s/651e592a91f309182edd`.

Myometrial smooth muscle cells have previously been observed in the decidual layer of the placenta in rats [31]. These decidualized myometrial cells are derived from the circular myometrium, and are incorporated into the decidua during pregnancy. During placentation, the circular layer at the implantation site is dispersed to allow access to the maternal vasculature [113]. Small parts of the longitudinal myometrium have also been observed in the mesometrial triangle in late pregnancy [108]. A schematic of the fully developed placenta is shown in Figure 8.1. This shows the dispersed circular layer situated in the

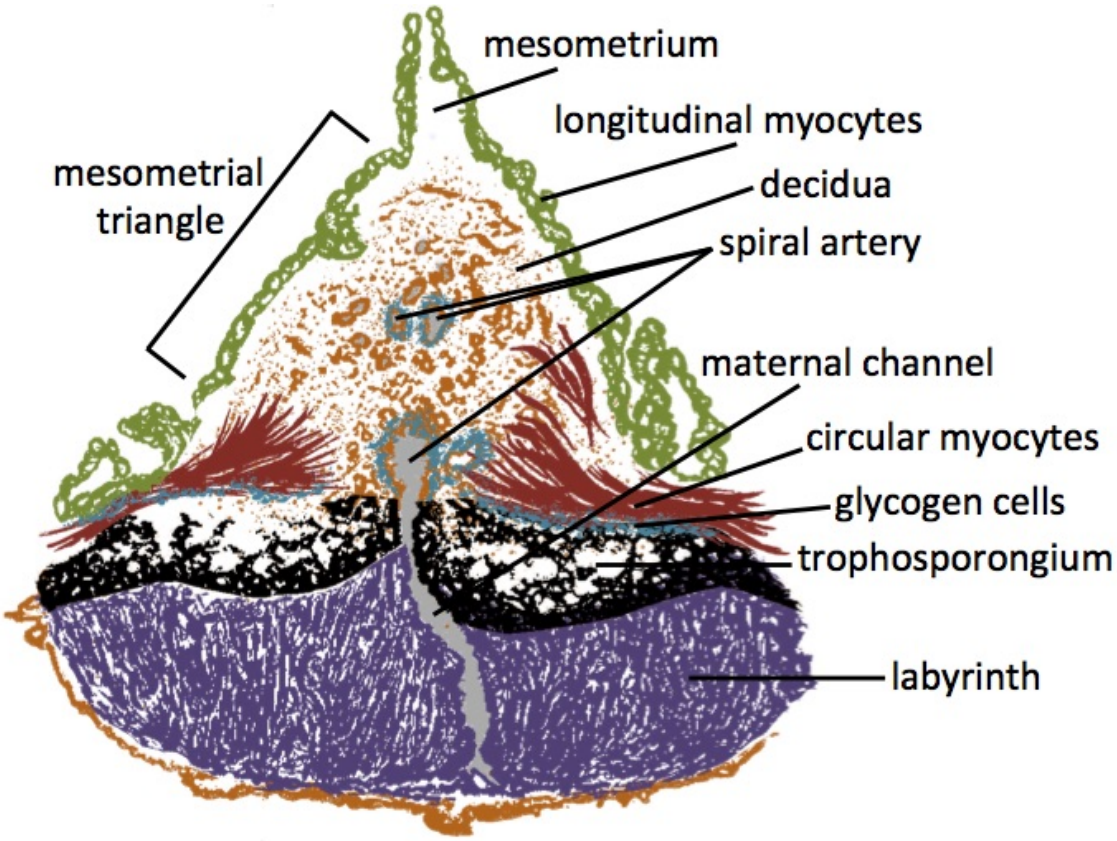

Figure 8.1: The relationship between the placenta and the myometrium. The longitudinal layer is displaced to line the mesometrial triangle. Circular smooth muscle fibres are dispersed in the decidua basalis. The spongiotrophoblast layer is internally adjacent to the decidua basalis. Interior to the spongiotrophoblast is the labyrinth. Drawing based on images from Vercruysse *et al.* [108].

decidua basalis directly exterior to the spongiotrophoblast, with the labyrinth positioned interior to the spongiotrophoblast. One important characteristic of the longitudinal interpenetrating fibres that will be investigated in the following is the depth to which these fibres penetrate and which layers of the placenta are affected.

8.1 Generating high-resolution reconstructions

The high resolution reconstruction of myometrial smooth muscle tissue at an implantation site was generated by differentiating between smooth muscle cell nuclei and placental cell nuclei, and using the smooth muscle cell nuclei as markers for the presence of smooth

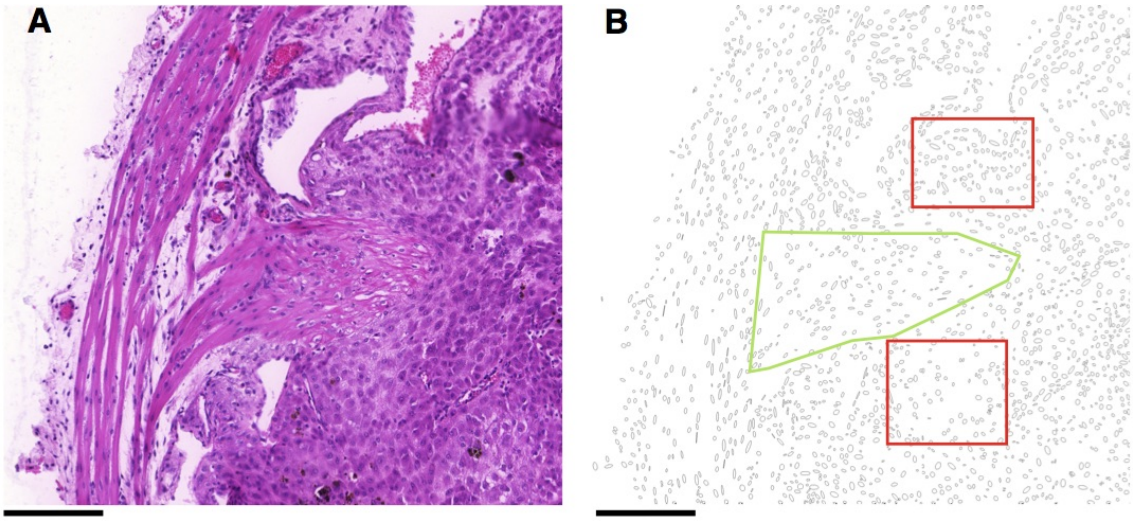

Figure 8.2: An example of how nuclear density and homogeneity was used to isolate fibres, applied to the histological image given in Figure 6.7. **A**: Histological image containing a longitudinal interpenetrating fibre. **B**: Ellipses approximating the nuclei in **A**. The area within the green polygon approximately corresponds to the fibre shown in **A**. The nuclei in this area are generally aligned along the same direction and the area is sparsely populated. Both red boxes are outside of the fibre shown in **A**, and fail on one of the two criteria. The upper red box contains nuclei which all have similar alignment, but it is densely populated. Conversely, the lower red box is sparsely populated, but the alignment of the nuclei is inconsistent. Scale bars represent 200 μm.

muscle tissue. This was achieved by considering the local density and homogeneity of nuclear direction in the area surrounding each nucleus; an area with relatively few nuclei in a similar orientation is likely to contain smooth muscle tissue, whereas an area which is densely populated with nuclei, or where the nuclei do not have a uniform direction is unlikely to contain smooth muscle, as shown in Figure 8.2.

The implantation sites in each slide were mapped onto images representing nuclear direction, where the length of each pixel corresponded to $\sim$ 10 μm in the histological slide. This resolution was used because it set one pixel to the approximate area of a single nucleus. The direction image for an implantation site in a slide has pixel values equal to the direction angle of the nucleus corresponding to the given pixel's position, if present. If multiple nuclei were present within the area corresponding to a pixel, one of the nuclei was selected at random for this representation. Local density and homogeneity within each of these images was calculated as follows. For each pixel p corresponding to a cell

Parameter	Value
Homogeneity score threshold	0.80
Density threshold	0.33
Smooth score threshold	0.05

Table 8.1: Parameters used for reconstruction of implantation sites.

nucleus, a rectangle R_p of length 30 pixels and width 3 pixels was drawn along the pixel direction, centred at the given pixel. A homogeneity score of this rectangle was calculated, as follows:

$$H(R_p) := \frac{1}{|r_p|} \sum_{q \in r_p} |\cos(a(p) - a(q))|,$$

where r_p is the set of all pixels in R_p which correspond to cell nuclei and $a(\cdot)$ is the pixel value. The density of R_p is given by

$$D(R_p) := |r_p|/|R_p|.$$

Threshold values were selected for $H(R_p)$ and $D(R_p)$ by comparison with visual inspection of the tissue, and are listed in Table 8.1. If the value of $H(R_p)$ was above the given score threshold and the value of $D(R_p)$ was below the given density threshold, then R_p was deemed to represent smooth muscle tissue. Figure 8.3**A** demonstrates this process in the context of the nuclear ellipses shown in Figure 8.2**B**. The set S of all such rectangles was used to create a homogeneity score image as follows. Each pixel q contained in at least one of the above rectangles was assigned the score

$$\max_{R_p \in S} |\cos(a(p) - a(q))|.$$

This score represents how well the given nucleus is aligned with a rectangle identified as being aligned with the direction of a smooth muscle fibre, and therefore is a measure of how well the nucleus fits into the category of smooth muscle fibre. Figure 8.3**B** demonstrates this scoring in the context of nuclear ellipses.

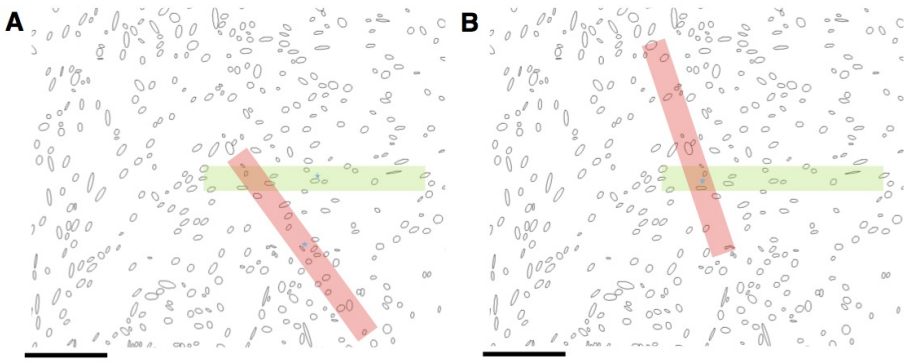

Figure 8.3: Rectangles used to generate score images. **A**: The initial selection of valid areas. Rectangles are placed over each nucleus along the direction given by the nucleus. The blue stars indicate the two nuclei which correspond to the rectangles shown. The green rectangle contains nuclei which closely approximate the direction of the central nucleus, so this rectangle is categorised as being within a fibre. Conversely the red rectangle contains nuclei which show little similarity to the central nucleus, so this rectangle is not categorised as being within a fibre. **B**: Scoring valid nuclei. For a given nucleus, the rectangles which contained this nucleus and were deemed to represent fibrous tissue by the process illustrated in **A** were compared to the nucleus. The rectangle with the closest orientation to the nucleus was selected to generate the score of the nuclei (see text). In this example, the nucleus (blue star) is closer in orientation to the green rectangle than the red, and thus the green rectangle is used to generate the score. Scale bars represent 100 μm.

The score images were generated from registered slides; accordingly, the stack of these score images forms a three-dimensional volume, with the z-coordinate for each slide equal to the position of the slide in the stack. This volume has variable voxel height, because the slides are not evenly spaced; however, this issue will be addressed by rescaling once all other processing is completed. The aim of the following steps is to obtain a set of clusters from this score volume which represent possible smooth muscle fibres. These clusters can subsequently be selected manually to reduce the set to clusters which generally represent smooth muscle fibres. Each of these steps is shown in Figure 8.4.

The score volume is sparsely populated with positive values due to the density of nuclei being restricted to 33% (Table 8.1. In order to convert this volume into a better connected representation of the fibrous structures, the volume was smoothed. The structures to be detected are long and thin, which motivates the use of steerable smoothing [41]. This smoothing was performed by applying a Gaussian kernel with a longer radius in the approximate direction of the given voxel, as illustrated in Figure 8.5. This method was used

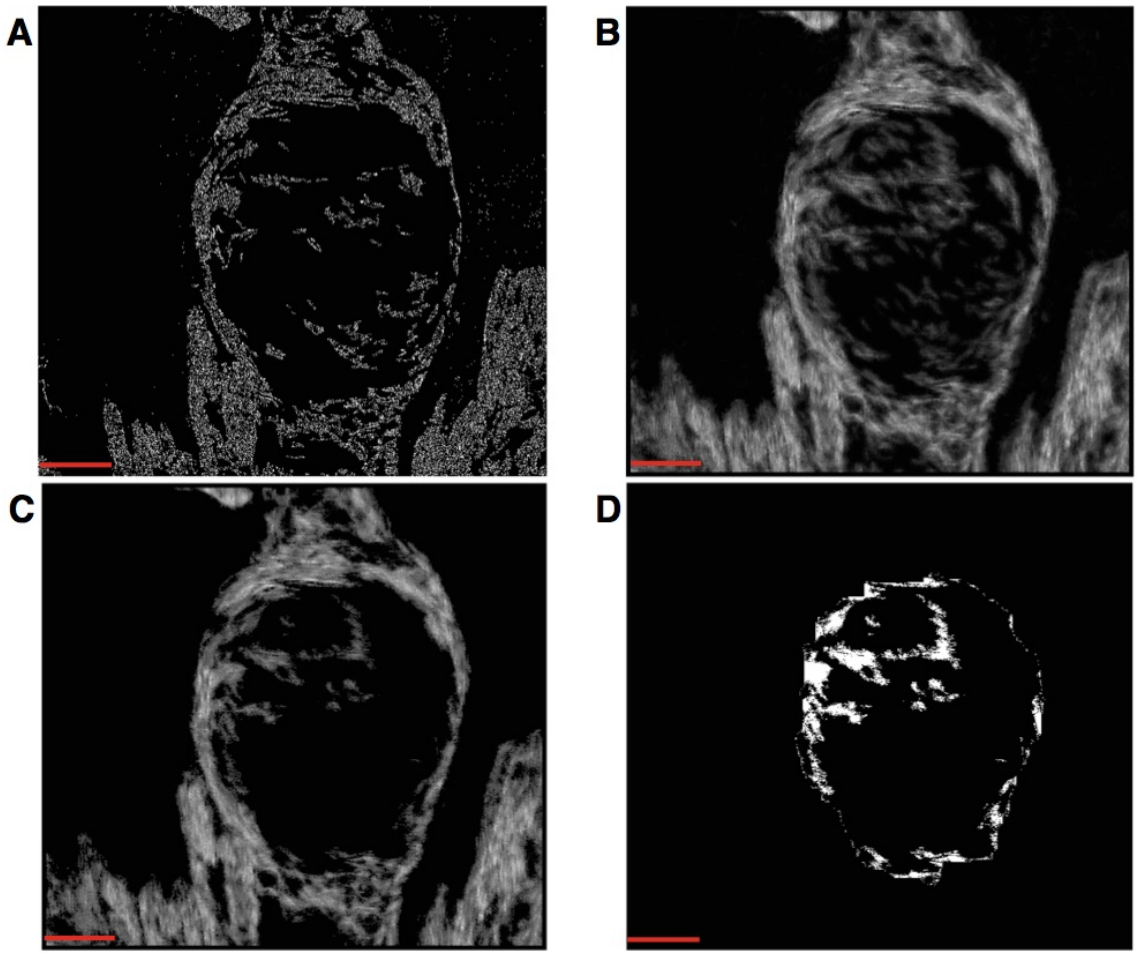

Figure 8.4: Steps taken to obtain potential fibres on the interior of the placenta. **A**: Score image, generated based on local homogeneity and density. **B**: Smoothed score image, generated by anisotropic smoothing of the score volume. Smoothing is performed in three dimensions, which can lead to the introduction of features not present in the original score image for this slide. **C**: Clusters connected to the exterior structure by some path through the volume. **D**: Clusters on the interior of the placenta. Scale bars represent 1 mm.

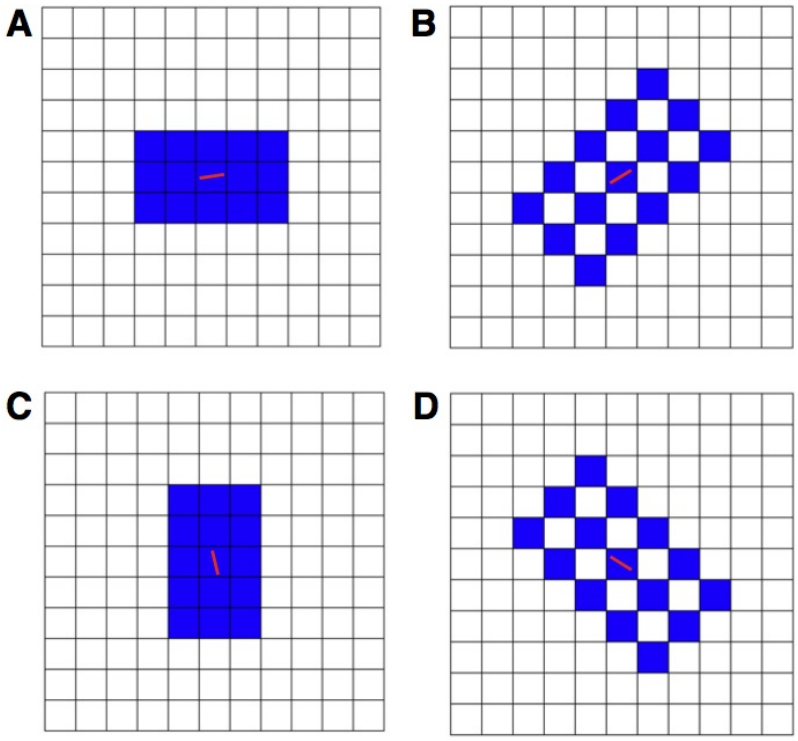

Figure 8.5: The shape of steerable smoothing kernels used to smooth the score volume. The kernel used at each point depends on the direction at the given point. The angle the direction made with the x-axis is rounded to the nearest 45°, and the corresponding kernel is selected. **A**: $0°, 180°$; **B**: 45°; **C**: 90°; **D**: 135°. While kernels **B** and **D** contain empty space between the points, these kernels are symmetric shapes of the same size as **A** and **C**, which ensures that they have similar effects.

to smooth at voxels with positive scores, but voxels with a score of 0 do not have a direction. The direction used to select a suitable kernel for voxels with a score of 0 was taken to be the direction of the nearest positive scoring voxel. An example of a slide after this smoothing is shown in Figure 8.4**B**.

The next step was to reduce the set of putative fibres to ones that were connected to the fibres on the outside of the placenta. This was achieved by separating voxels into sets of voxels above the smooth score threshold value given in Table 8.1 which are connected by adjacency, and setting all voxels not in the largest of these sets to 0. As long as the volume bounds around the placenta are sufficiently wide, the largest of these sets will include the fibres outside the placenta and any potential fibres which are attached to it. This step

ensures that manually selected fibres are guaranteed to be connected to the fibres outside the placenta, as these are the fibres of interest. An example of a slide after this step is shown in Figure 8.4**B**.

The final step required to enable manual selection of fibres is to detect putative fibres on the interior of the placenta and segment this set into clusters. Fibres are represented by clusters of connected points, therefore grouping these points into clusters ensures that manually selecting the portion of a fibres observed in one slide will correspond to the selection of the entirety of the fibre. This segmentation was achieved by identifying the interior of the overall structure identified above and segmenting individual structures embedded in this interior volume. The interior of the structure was found by identifying the interior within each two-dimensional slide plane. For each of these two-dimensional image planes, the interior was found by searching the image for connected sets of pixels below the smooth score threshold value given in Table 8.1. The central pixel in the image was taken as the starting point. If this pixel had value above the threshold value, then the starting point was moved to a random point within 5 pixels from the previous point. This process was repeated until a pixel with value below the threshold value was found. The interior set was found by expanding outward from this starting point, adding all points adjacent to the interior set which had value below the given threshold to the interior set. When this expansion was completed, if the set found had size less than 1/8th the total image area, then the starting point was randomly displaced as above and the process was repeated until sufficient points were found or a time out was reached. While this procedure is able to identify the interior of the structure in most image planes, there are two exceptional cases for which it would fail: not all image planes had an interior for the structure, as can be observed in Figure 8.6**A**, and, moreover, a number of image planes contained barriers which were too large to be crossed by the above random displacement, as can be observed in Figure 8.6**B**. Image planes with the former property were identified by checking the image boundary for interior points, while image planes with the latter

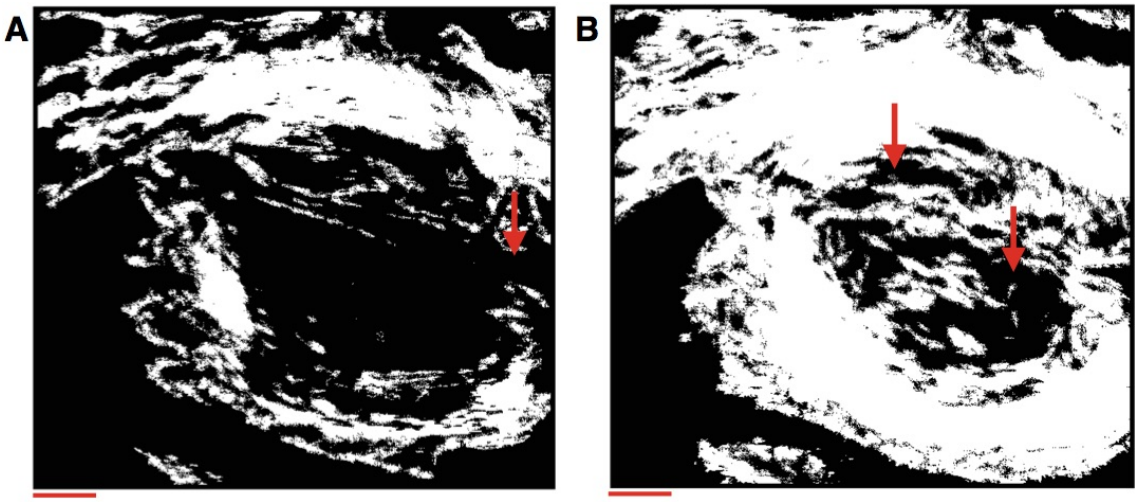

Figure 8.6: Examples of situations where the interior cannot be identified. Each image shows a thresholded image of a plane from the smoothed volume, where white represents the myometrium outside the placenta and potential interior myometrium. **A**: The outer structure does not completely enclose the interior, with a gap in the smooth muscle indicated by a red arrow. This means that the interior is not well-defined in this example. **B**: Two interior areas are present, as indicated by red arrows, with a barrier separating the two areas. This barrier could be overcome by the random displacement described in the text; however, if this does not occur in the finite sequence of random displacements, the interior area is identified as too small and the interior is classified as ill-defined. Scale bars represent 1 mm.

property were identified by the algorithm timing out before sufficient area was covered. In each of these cases, the interior set for the image was replaced by the interior set of the nearest image below the given image in the stack which had a well-defined interior. If all images below the given image in the stack did not have a well defined interior, then the interior set was taken from an image above. The interior sets of the image planes found in this manner were reassembled into a volume to give an interior volume.

Structures embedded in this interior volume were found by dilating and eroding the interior volume [105]. Dilation was performed by adding all voxels adjacent to an interior voxel to the set of interior voxels. Erosion, conversely, was performed by removing all interior voxels adjacent to a non-interior voxel from the set of interior voxels. The interior set was dilated 20 times and subsequently eroded 20 times. This corresponds to expanding and contracting the interior volume by ~ 200 μm. Any structures extending into the interior volume of width less than ~ 400 μm would be incorporated into the interior

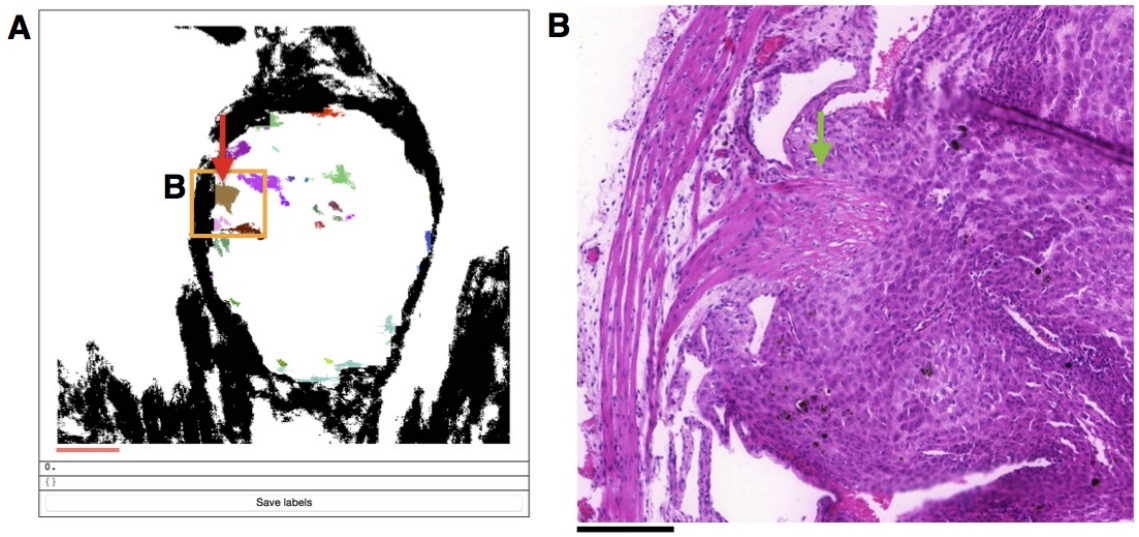

Figure 8.7: Manual selection of fibres, demonstrated here on the fibres previously displayed in Figure 6.7. **A**: The user interface used to select fibres. The upper panel displays segmented interior fibres in random colours for the given slide; the lower panels are used to write fibres selected by the user to generate the final volume. **B**: A small portion of the implantation site in the given slide. The fibre that is indicated in **B** by the green arrow corresponds to the cluster indicated by the red arrow in **A**, and is selected for the final representation. All other clusters in the area corresponding to **B** do not correspond to fibres, as can be seen in the histology, and thus are not selected. The clusters are three-dimensional: the highlighted cluster is part of a larger structure extending through multiple slides. Therefore, manually identifying this structure in one slide leads to the automatic selection of the structure in other slides. It is therefore possible to verify the fibres on a sparse selection of slides and use the clustering to interpolate the overall structure. Red scale bar represents 1 mm, black scale bar represents 200 μm.

volume upon dilation, but remain in the interior volume after erosion. An example of the resulting interior clusters is shown in Figure 8.4**D**. The structures that remain in the interior after dilation are the potential fibres which need to be verified manually.

The structures were labelled to enable visual comparison with individual slides. This was achieved by grouping points with value above the smooth score threshold given in Table 8.1 by adjacency in the three-dimensional volume, yielding clusters which represent structures, as shown in Figure 8.7**A**. These clusters were compared to the histological slides, and this structure was included in the final volume if smooth muscle tissue was found at the location of a labelled structure. The final volume was thus comprised of all exterior tissue identified by the above process and any interior structures identified

visually.

8.2 Results

The interpenetrating fibres examined in this section are distinct from the decidualized circular myometrium and are differentiated in the following manner. The newly observed interpenetrating fibres appear to be a structure branching off from the smooth muscle fibres on the exterior of the placenta, with the branch appearing to extend in a perpendicular direction to the external fibres. Decidualized smooth muscle, on the other hand, tends to appear as an expansion of the circular layer into the decidua. Additionally, the fibres examined here branch off from the longitudinal layer. The differences between these two types are shown in Figure 8.8.

Smooth muscle fibres entering the placenta from the longitudinal layer were identified in three implantation sites in tissue block 1, three implantation sites in tissue block 2, and one implantation site in tissue block 3. Additionally, there were structures that could represent these fibres, but it was not possible to confirm the status, because either it was unclear if these structures were attached to the longitudinal layer or slides containing tissue which could identify these fibres were discarded during sectioning for quality control. Examples of each of these situations are shown in Figure 8.9. Figure 8.10 charts the location of all interpenetrating fibres in each block in the context of the whole tissue, including locations where absolute confirmation was not possible due to the aforementioned factors. Positively identified interpenetrating fibres had lengths within the range 0.5–3 mm and widths at the point of entry within the range 100–200 μm, as demonstrated in the figures below. Two of the implantation sites in tissue block 1 and one of the implantation sites in tissue block 2 were reconstructed using the above technique. The observed fibres are presented below by tissue block, with three-dimensional reconstructions where applicable. In the following, sites are numbered according to the convention established in Figure 8.11.

All three-dimensional reconstructions were visualised using 3D Slicer [38] (www.slicer.org).

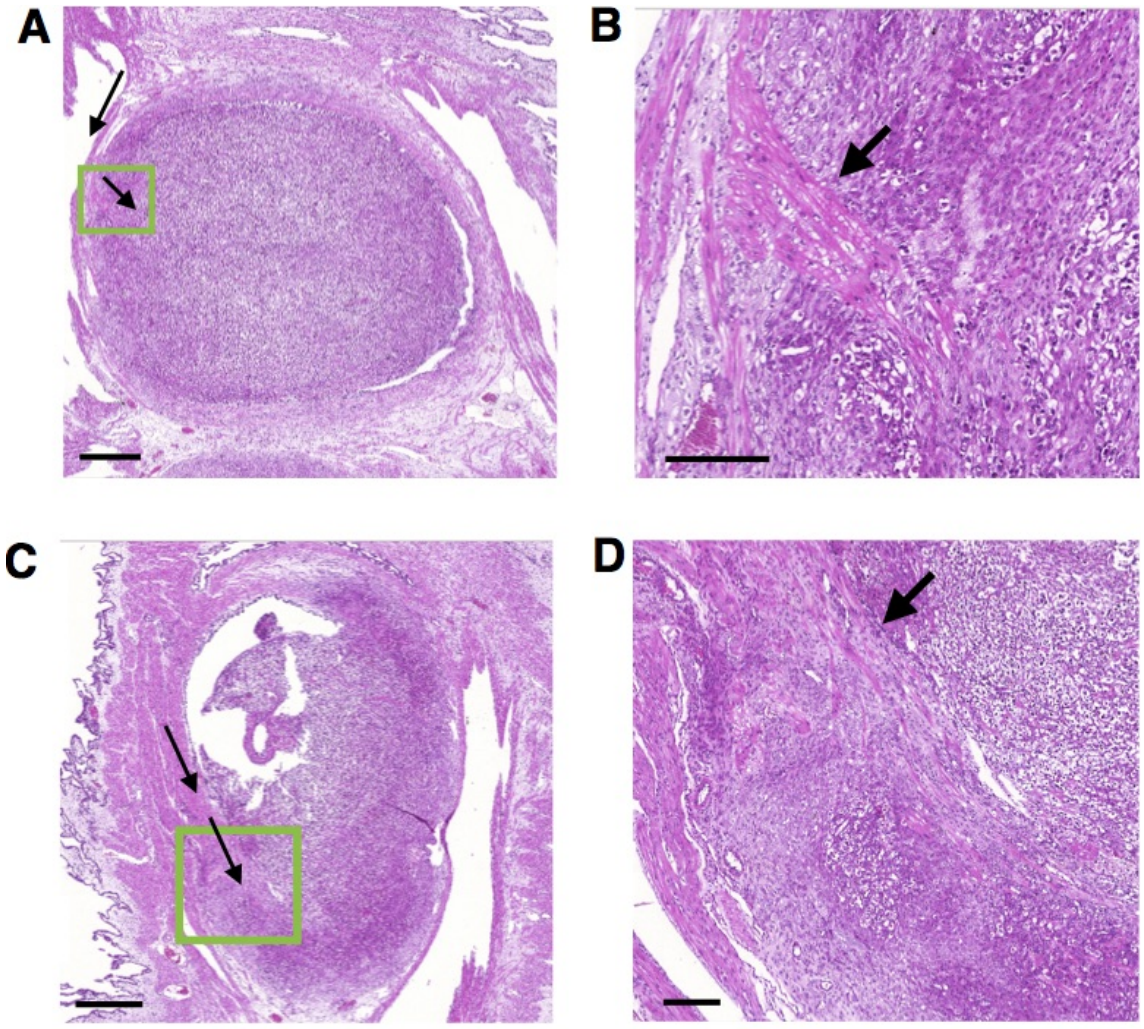

Figure 8.8: Examples of longitudinal interpenetrating fibres discussed here (top) and decidualized circular myometrium (bottom). **A**: Implantation site containing an interpenetrating fibre. Arrows indicate direction of the fibre and of the myometrium to which it is attached. Green panel indicates the location of **B**. **B**: Higher-resolution image of the panel in **A**. Arrow indicates the fibre of interest. **C**: Implantation site containing decidualized circular myometrium. Arrows indicate the fibre direction along the circular layer and the decidualized cells. Green panel indicates the location of **D**. **D**: Higher-resolution image of the panel in **C**. The interpenetrating fibre in **A** and **B** is shown in more detail in Figure 8.25. Arrow indicates decidualized smooth muscle tissue. Scale bars in **A** and **C** represent 1 mm, scale bars in **B** and **D** represent 200 μm.

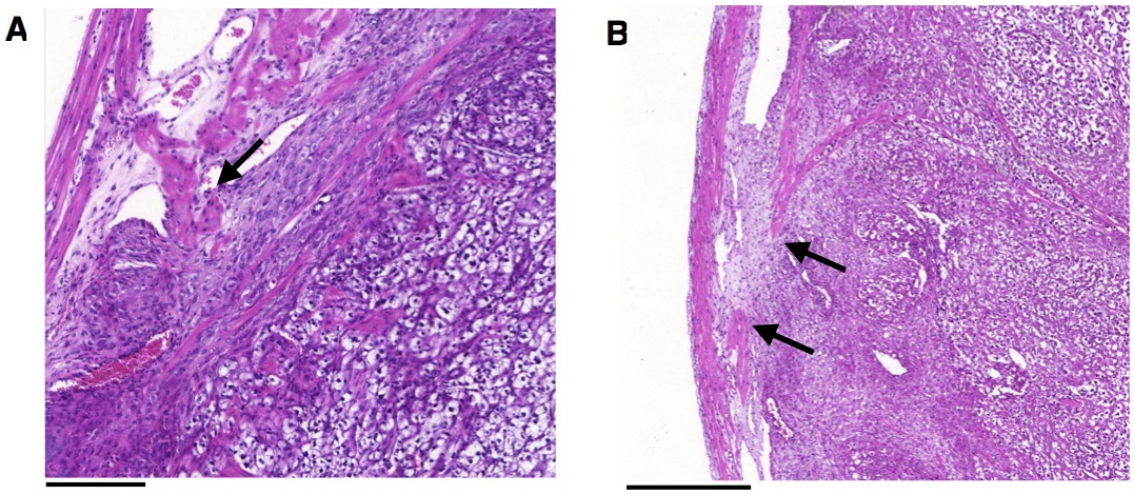

Figure 8.9: Structures which could belong to the longitudinal interpenetrating fibres but could not be confirmed. **A**: A fibre enters the placenta from the longitudinal layer (arrow), but the presence of circular fibres adjacent to this fibre mean that it is not possible to determine if fibres in this area are circular or longitudinal. **B**: The smooth muscle indicated by the arrows appears to be part of the same fibre; however the slides immediately above this slide in the stack were discarded and hence the connection could not be confirmed. Scale bar in **A** represents 200 μm, scale bar in **B** represents 500 μm.

These reconstructions were facilitated by generating DICOM files using the dcm4che package (www.dcm4che.org). The code used to generate these DICOM files can be found at `https://figshare.com/s/651e592a91f309182edd`.

8.2.1 Tissue block 1

Generally, the proportion of the placenta occupied by smooth muscle tissue was greater in tissue block 1 than in the other two blocks. This was noticeable in both the interpenetrating fibres of interest here and the decidualized circular muscle, leading to larger areas covered by both of these types. This increased size resulted in clearer three-dimensional reconstructions of the interpenetrating fibres; however, the increased presence of decidualized circular smooth muscle made separation of the two types more difficult, both manually and automatically. For this reason, the reconstruction shown in Figure 8.15 also highlights the decidualized circular smooth muscle along with the interpenetrating fibres, which are of interest here. Additionally, the implantation sites not featured here may contain the

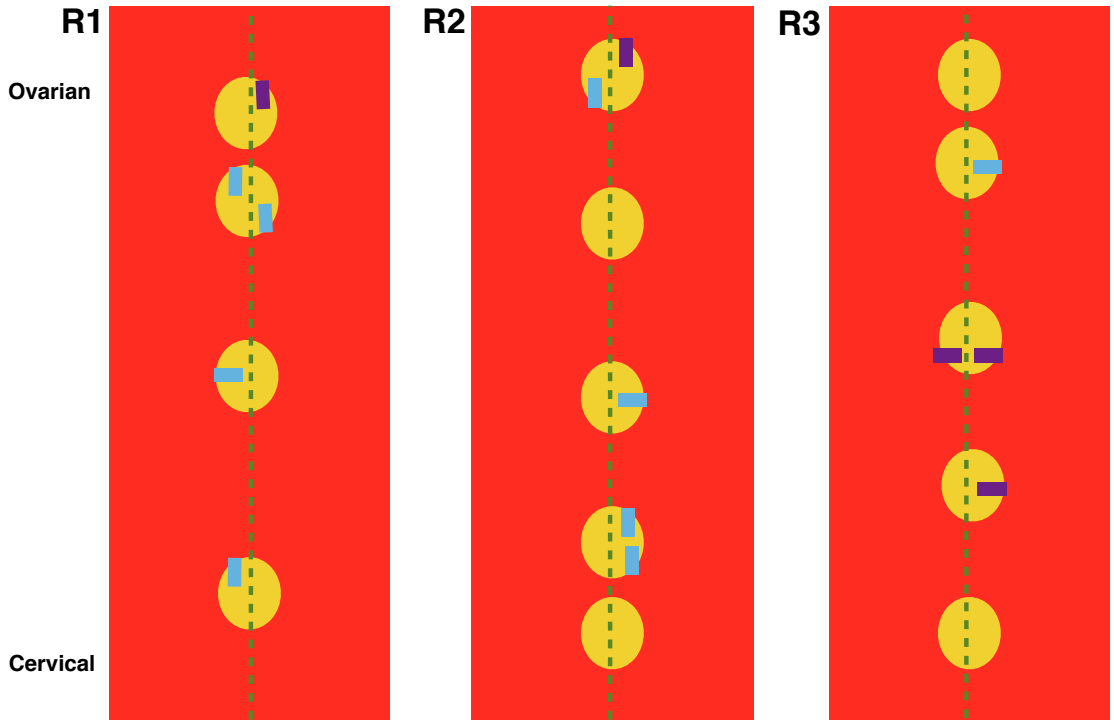

Figure 8.10: The locations of all putative longitudinal interpenetrating fibres in the rat myometrium, located through examination of the histology. **R1**, **R2**, and **R3** denote tissue blocks 1, 2, and 3 respectively. Yellow spots indicate the position of the implantation sites, green dashed lines indicate the approximate location of the mesometrial border, and blue rectangles indicate the approximate location and orientation of confirmed interpenetrating fibres, while purple rectangles indicate suspected interpenetrating fibres (see text for details). These fibres are present on both sides of the mesometrial border. The orientation appears to have no discernible pattern.

fibres of interest but could not be differentiated from decidualized circular smooth muscle.

The reconstruction of implantation site 2 is shown in Figure 8.12. This site contains two fibres extending from the longitudinal layer. Both fibres are attached to the longitudinal layer, as indicated in Figure 8.12. The upper of these two fibres, marked **F1**, is shown in histological context in Figure 8.13. This shows the fibre extending into the spongiotrophoblast of the placenta, but not reaching the labyrinth. The lower of the fibres, marked **F2**, is shown in histological context in Figure 8.14. Again, the fibre extends into the spongiotrophoblast of the placenta, but does not reach the labyrinth.

The reconstruction of implantation site 3 is shown in Figure 8.15. The fibre is attached to the longitudinal layer, as indicated. This fibre is shown in histological context in Fig-

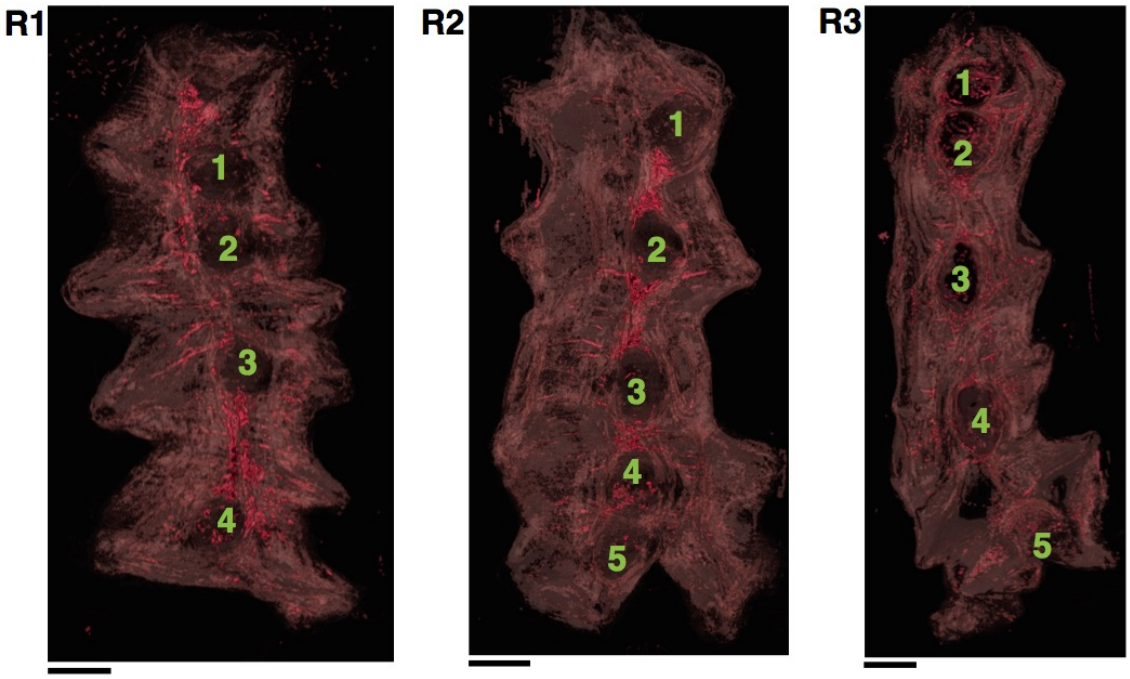

Figure 8.11: Numbering convention for implantation sites in the tissue blocks studied. Scale bars represent 5 mm.

ures 8.16 and 8.17. Figure 8.16 shows the fibre at the point of entry to the placenta, where it can be seen attaching to the longitudinal fibres adjacent to the placenta. Figure 8.17 shows the fibre at the innermost point of the placenta in the reconstruction, where it can be seen extending into the spongiotrophoblast, but not penetrating the labyrinth.

8.2.2 Tissue block 2

The reconstruction of implantation site 4 is shown in Figure 8.18. This reconstruction shows two distinct fibres extending into the placenta, as was the case for site 2 in tissue block 1. There is one point where these fibres apparently join, indicated in the lower image in the figure with an arrow; however, Figure 8.19 shows the histology at that point, with no presence of smooth muscle tissue. This joint must be considered to be an artefact of the cluster detection, and therefore has to be be ignored. The fibre situated lower in the stack, labelled **F1**, is connected to the longitudinal layer as shown in Figure 8.20.

This connection forms a perpendicular branch from the longitudinal fibre. The fibre extends into the spongiotrophoblast of the placenta, as shown in Figure 8.21. The fibre situated higher in the stack, labelled **F2**, is connected to the longitudinal layer as shown in Figure 8.22. This branch is at a smaller angle to the longitudinal layer than the branch in **F1**; however, a change of direction in the branch is observable. The fibre extends into the spongiotrophoblast of the placenta, as shown in Figure 8.23.

8.2.3 Tissue block 3

One implantation site was observed to contain a fibre extending from the longitudinal layer into the placenta. The histology containing this fibre is shown in Figure 8.25. The fibre extends into the spongiotrophoblast.

An interpenetrating fibre was observed in implantation site 3, but the reconstruction of the three-dimensional structure of this fibre was not possible. Figure 8.24 shows this fibre, which extends into the spongiotrophoblast of the placenta

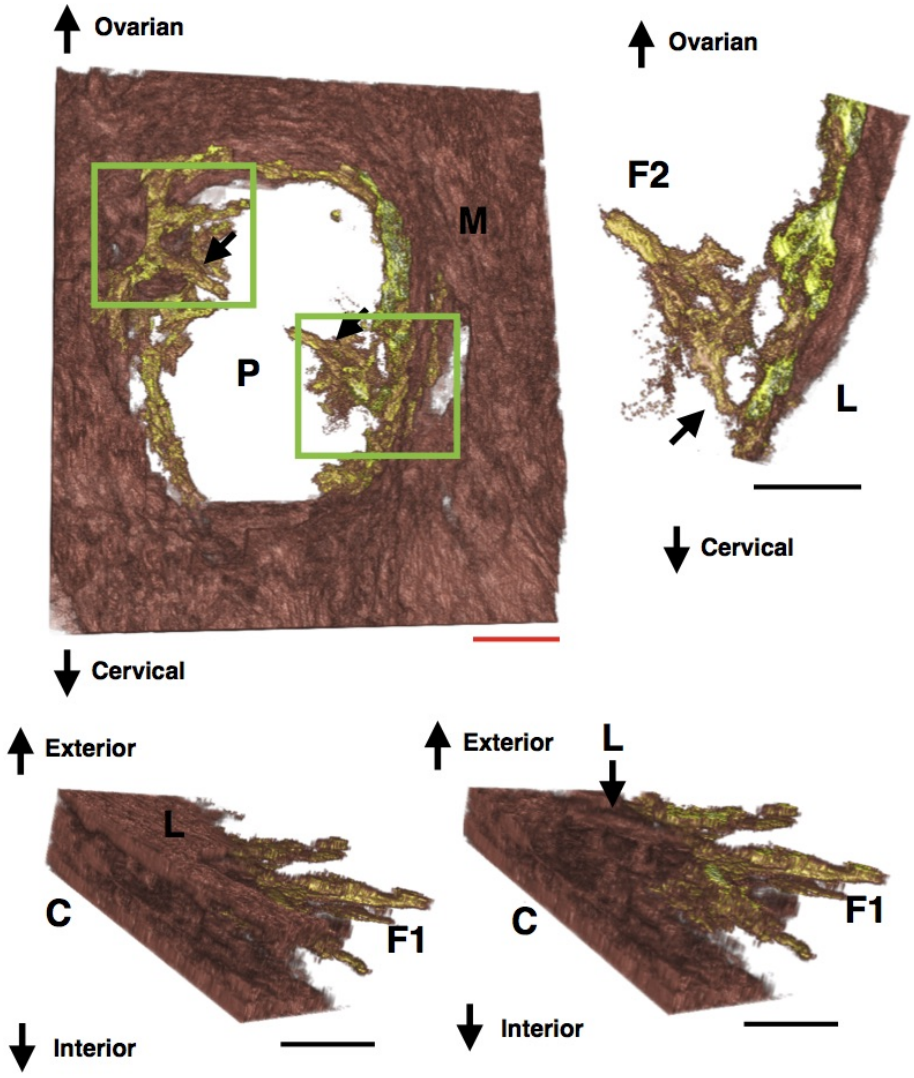

Figure 8.12: Three-dimensional *in silico* reconstruction of implantation site 2 in tissue block 1. Top left: reconstruction of the entire implantation site, with the smooth muscle fibres extending into the placenta highlighted in yellow. In this reconstruction it was not possible to separate the fibres of interest completely from decidualized circular smooth muscle, which has lead to excess highlighting of the interior. The fibres of interest are indicated by the labels **F1** and **F2**. Top right: isolated view of fibre **F2**, with arrows indicating points at which the fibre connects with the longitudinal smooth muscle fibres on the exterior of the placenta. Bottom right: **F1** viewed from the side of the tissue block. Bottom left: the same view of **F1** with most of the longitudinal fibres removed. This view shows the fibre attaching to a fibre in the longitudinal layer. **M**: myometrium, **P**: placenta, **L**: longitudinal myometrium, **C**: circular myometrium. Red scale bar represents 1 mm, black scale bars represent 500 μm.

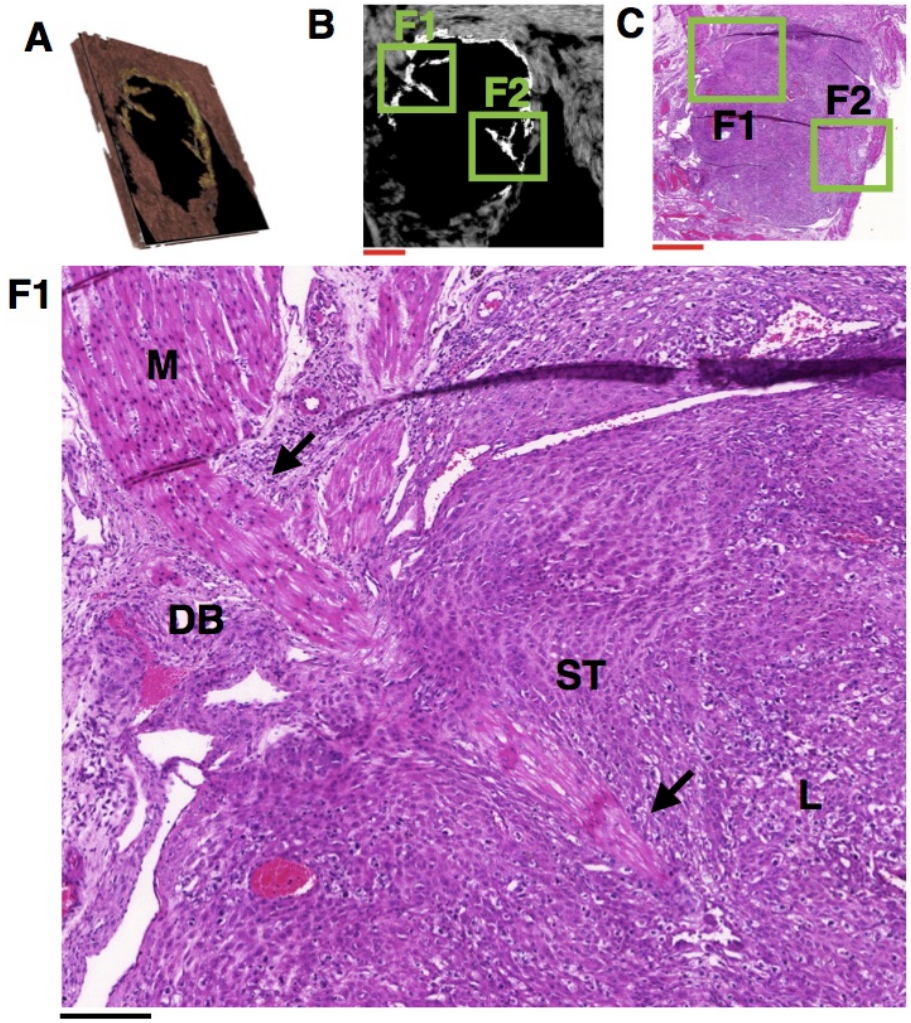

Figure 8.13: Histology of the placenta surrounding the entry point of the fibre **F1** shown in Figure 8.12. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frames indicating the approximate position of **F1** and **F2**. **C**: of implantation site 2 is entirely visible in this slide plane, with frames indicating the position of **F1** and **F2**. **F1**: high-resolution image of the implantation site around the fibre indicated by the arrows. The upper arrow indicates the point of entry of the fibre, attaching to the longitudinal layer. The lower arrow indicates the point of deepest observed penetration of the fibre, which appears to be contained within the spongiotrophoblast. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

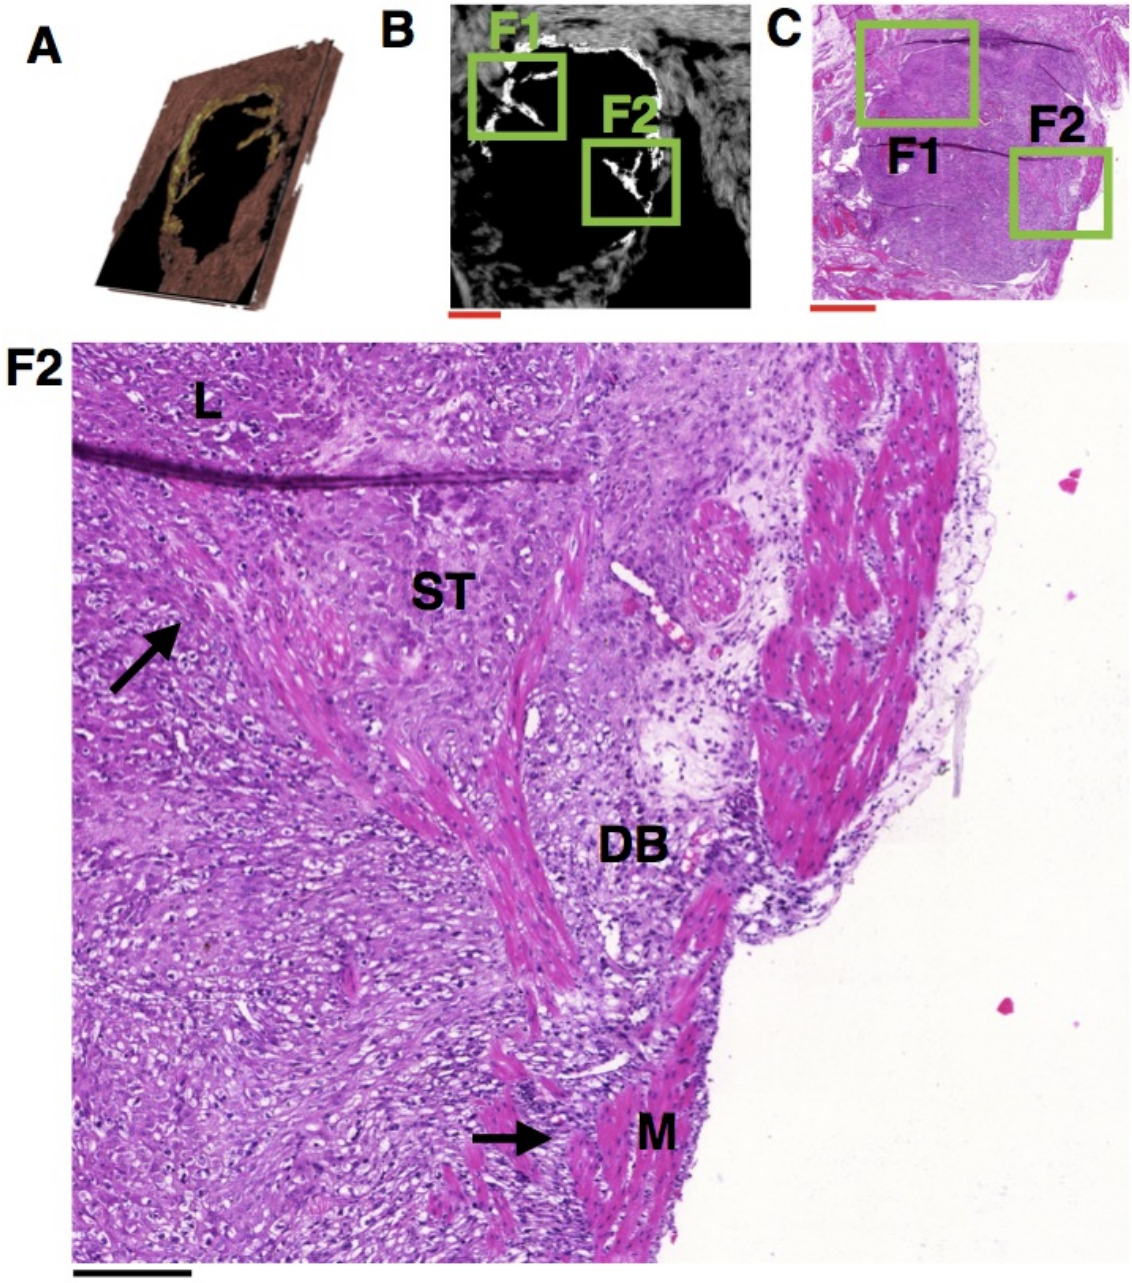

Figure 8.14: Histology of the placenta surrounding the entry point of the fibre **F2** shown in Figure 8.12. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frames indicating the approximate position of **F1** and **F2**. **C**: implantation site 2 is entirely visible in this slide plane, with frames indicating the position of **F1** and **F2**. **F2**: high-resolution image of the implantation site around the fibre indicated by the arrows. The lower arrow indicates the point of entry of the fibre, attaching to the longitudinal layer. The upper arrow indicates the point of deepest observed penetration of the fibre, which appears to be contained within the spongiotrophoblast. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

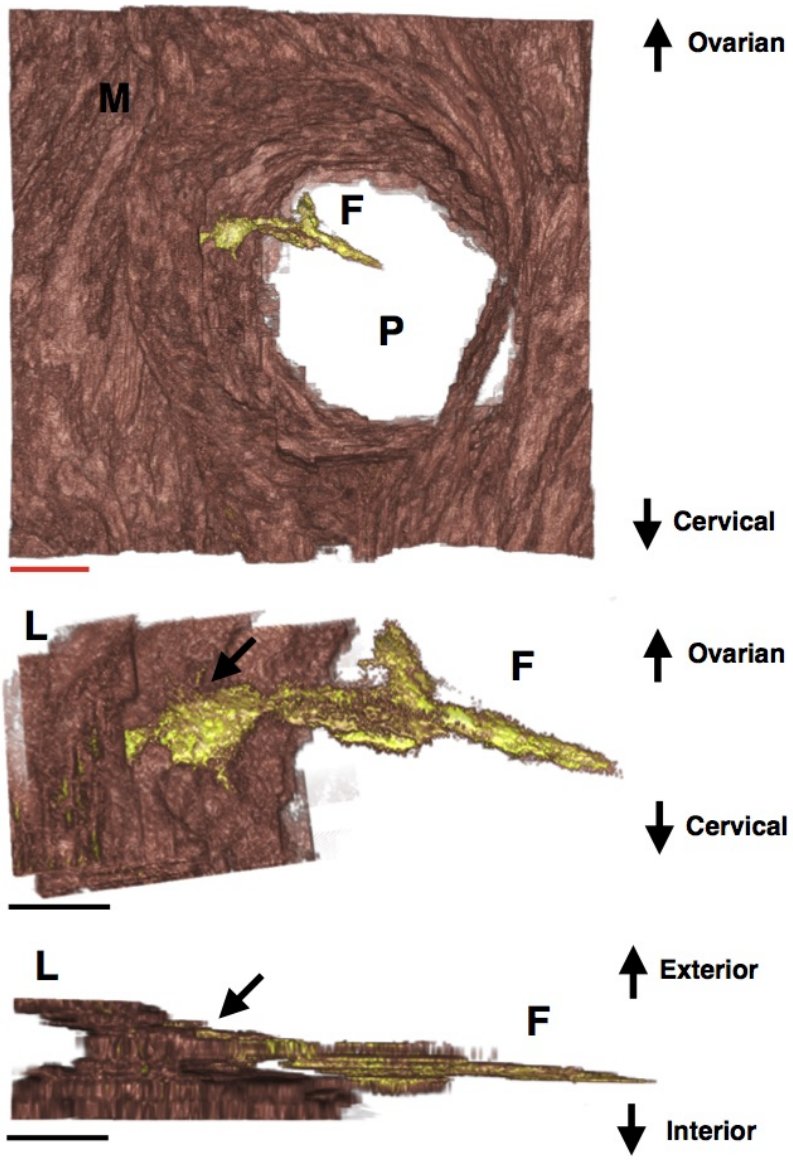

Figure 8.15: Three-dimensional *in silico* reconstruction of implantation site 3 in tissue block 1, previously identified in close proximity to excitation activity in Chapter 6. Top: reconstruction of the whole site, with the smooth muscle fibre extending into the placenta highlighted in yellow. Middle: view of the isolated fibre from the exterior surface, with an arrow indicating the point at which the fibre connects with the adjacent longitudinal fibres. Bottom: the isolated fibre from the caudal viewpoint, with an arrow indicating the point at which the fibre connects with the adjacent longitudinal fibres. **M**: myometrium, **F**: fibre of interest, **P**: placenta, **L**: longitudinal myometrium. Red scale bar represents 1 mm, black scale bars represent 500 μm.

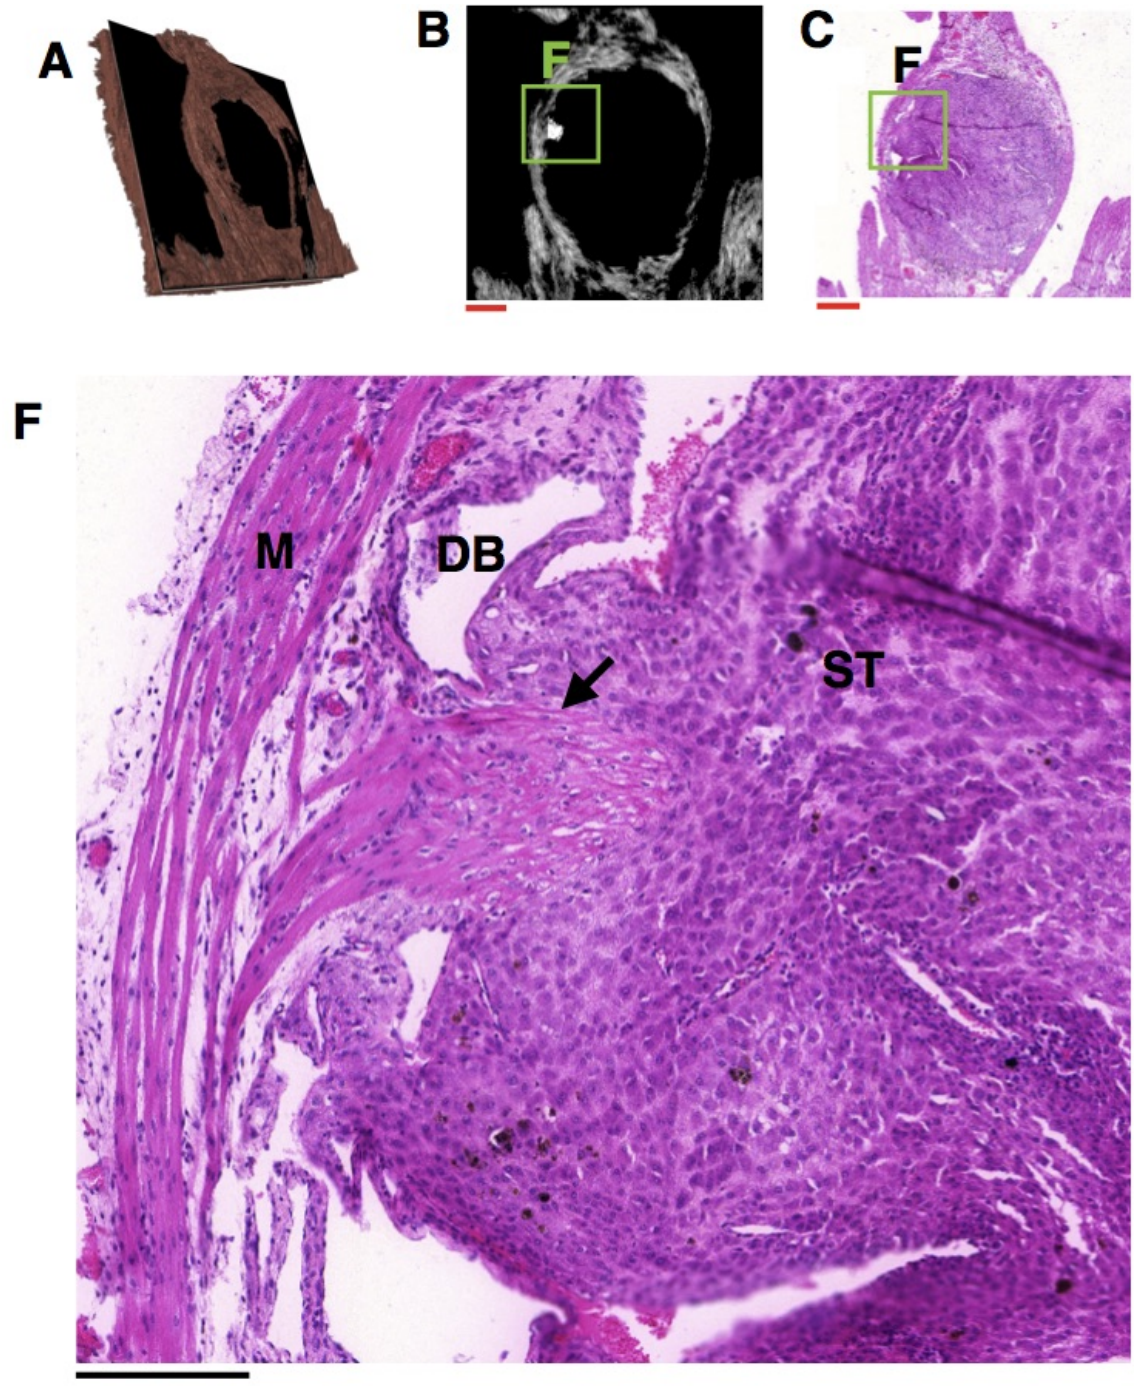

Figure 8.16: Histology of the placenta at the entry point of the fibre shown in Figure 8.15. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 3 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high resolution image of the implantation site at the entry point of the fibre indicated by the arrow (previously displayed with less detail in Figure 6.7). **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast. Red scale bars represent 1 mm, black scale bar represents 200 μm.

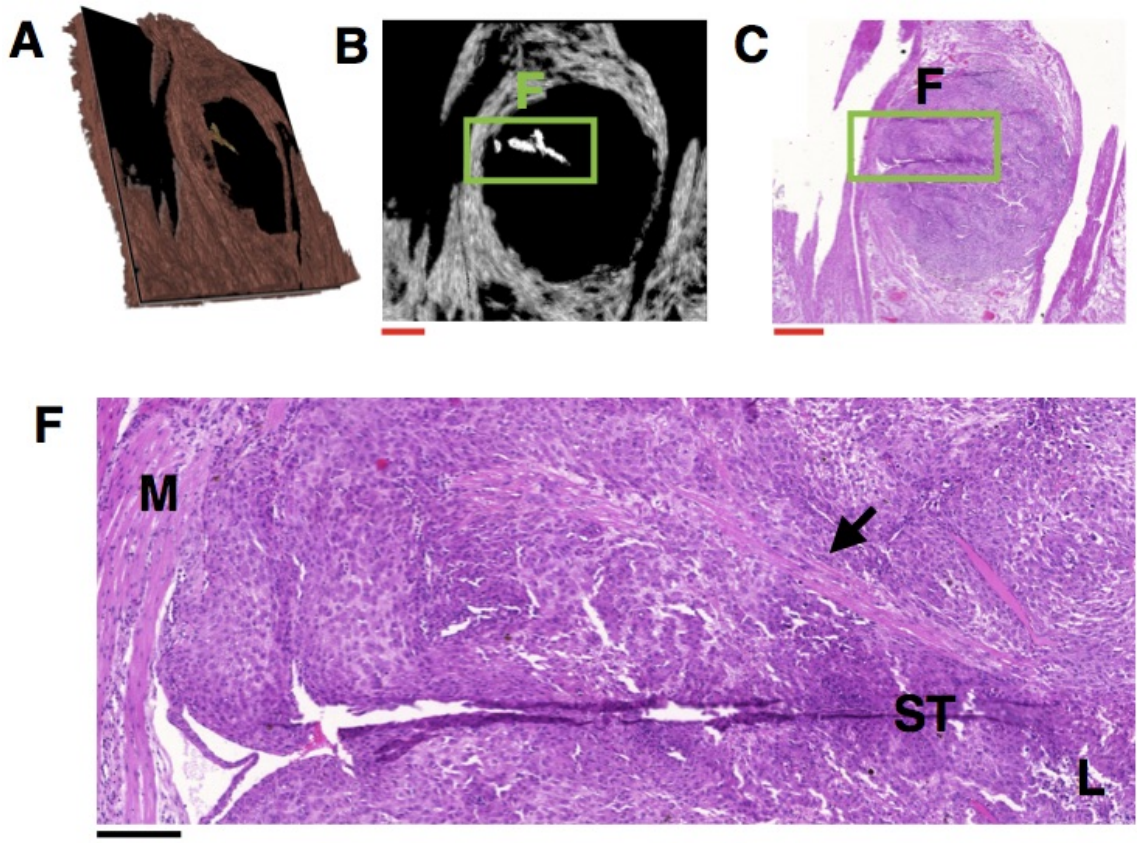

Figure 8.17: Histology of the placenta at the point of deepest penetration of the fibre shown in Figure 8.15. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 3 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high resolution image of the implantation site in this slide, with the fibre indicated by the arrow. This appears to penetrate the spongiotrophoblast, but no point of contact with the labyrinth is apparent. **M**: myometrium, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

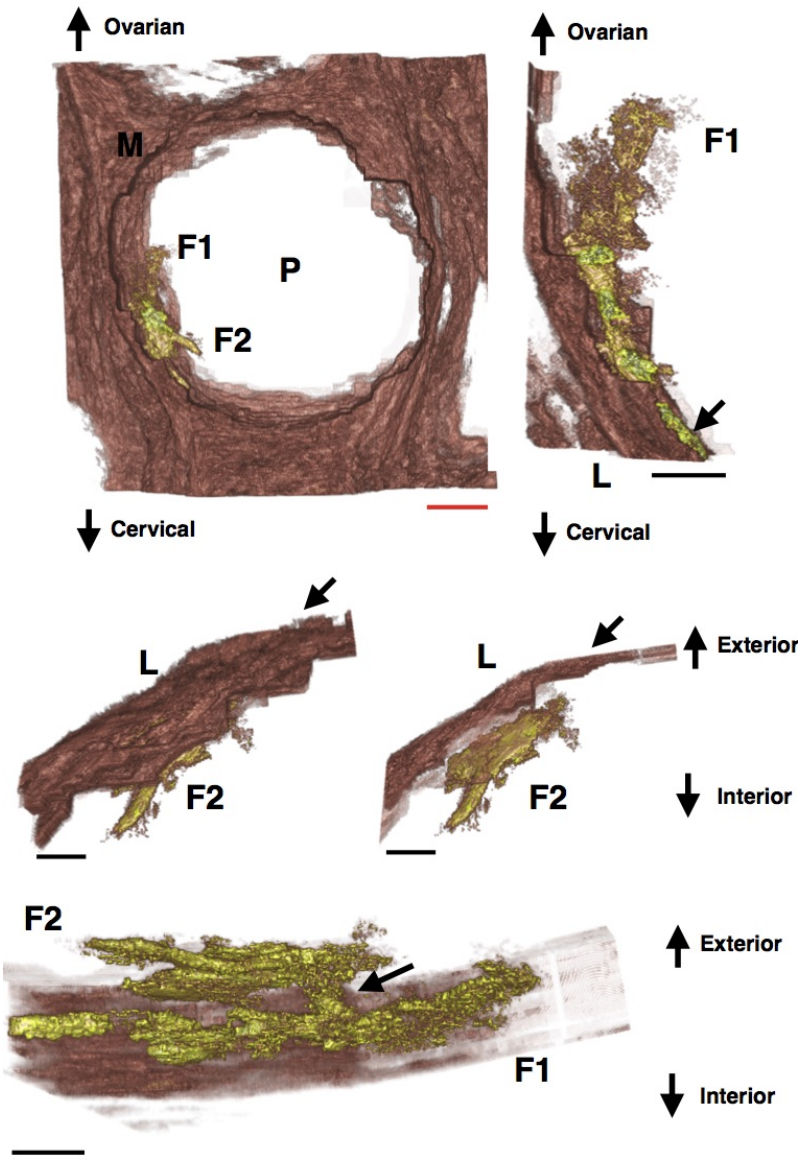

Figure 8.18: Three-dimensional *in silico* reconstruction of implantation site 4 in tissue block 2. Top left: reconstruction of the entire implantation site, with the smooth muscle fibres extending into the placenta highlighted in yellow. The fibres of interest are indicated by the labels **F1** and **F2**. Top right: isolated view of **F1**, with an arrow indicating where this fibre attaches to the adjacent longitudinal myometrium. Middle: isolated view of fibre **F2**, with arrows indicating the point at which the fibre connects with the longitudinal smooth muscle fibres on the exterior of the placenta. The left image shows the longitudinal layer present directly exterior to the fibre, while the right image shows the fibre without this cover. The reconstruction algorithm failed to detect the fibres joining **F2** to the point at which it attaches to the longitudinal myometrium. Bottom: both fibres viewed from the interior of the placenta. The arrow indicates the point which was detected as joining these fibres; however, this is shown to not contain smooth muscle tissue in Figure 8.19. **M**: myometrium, **P**: placenta, **L**: longitudinal myometrium. Red scale bar represents 1 mm, black scale bars represent 500 μm.

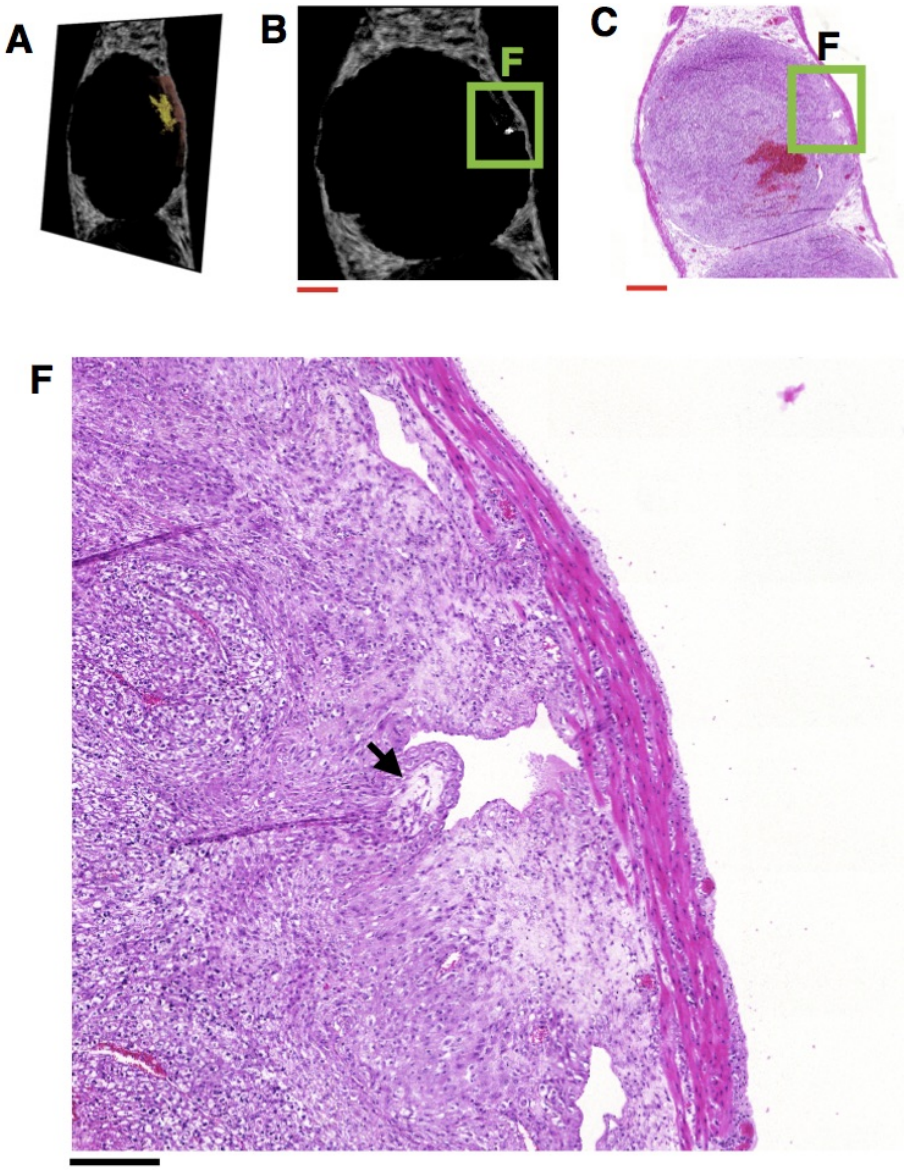

Figure 8.19: Histology of the placenta at the point indicated by the arrow in the lower panel of Figure 8.18. **A**: slide plane in the context of the three-dimensional reconstruction shown in the lower panel of Figure 8.18. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 4 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high resolution image of the implantation site at the detected bridge between **F1** and **F2** in Figure 8.18. The point detected as connecting these fibres is indicated by the arrow, which contains no smooth muscle, suggesting that this connection is a false positive Red scale bars represent 1 mm, black scale bar represents 200 μm.

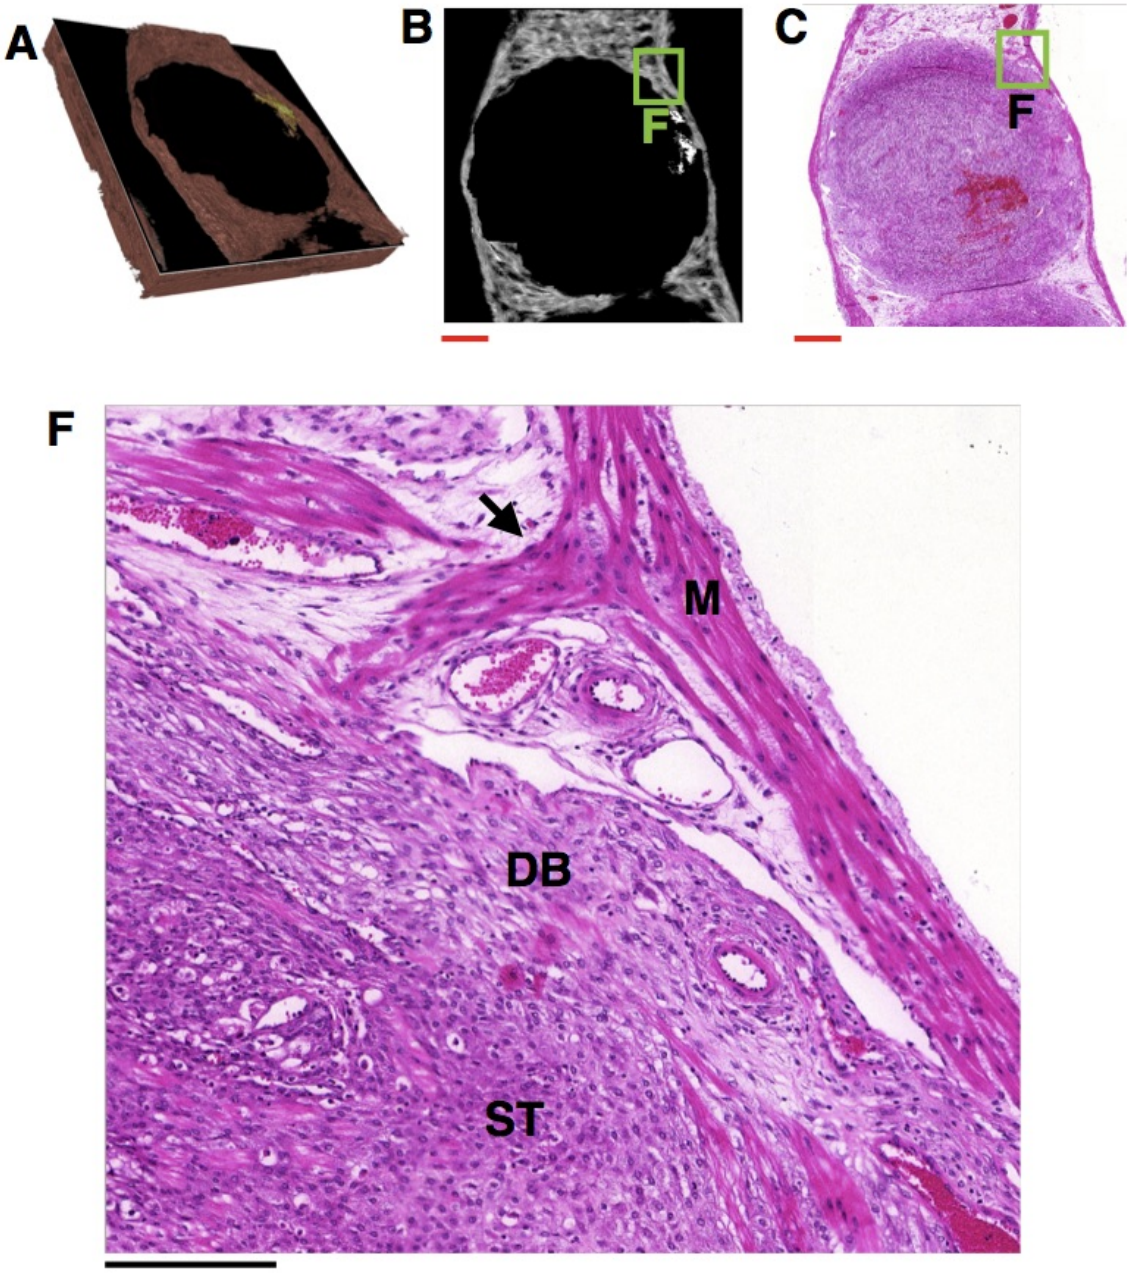

Figure 8.20: Histology of the placenta surrounding the entry point of the fibre **F1** shown in Figure 8.15. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 4 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high-resolution image of the implantation site at the entry point of the fibre **F1**, indicated by the arrow. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast. Red scale bars represent 1 mm, black scale bar represents 200 μm.

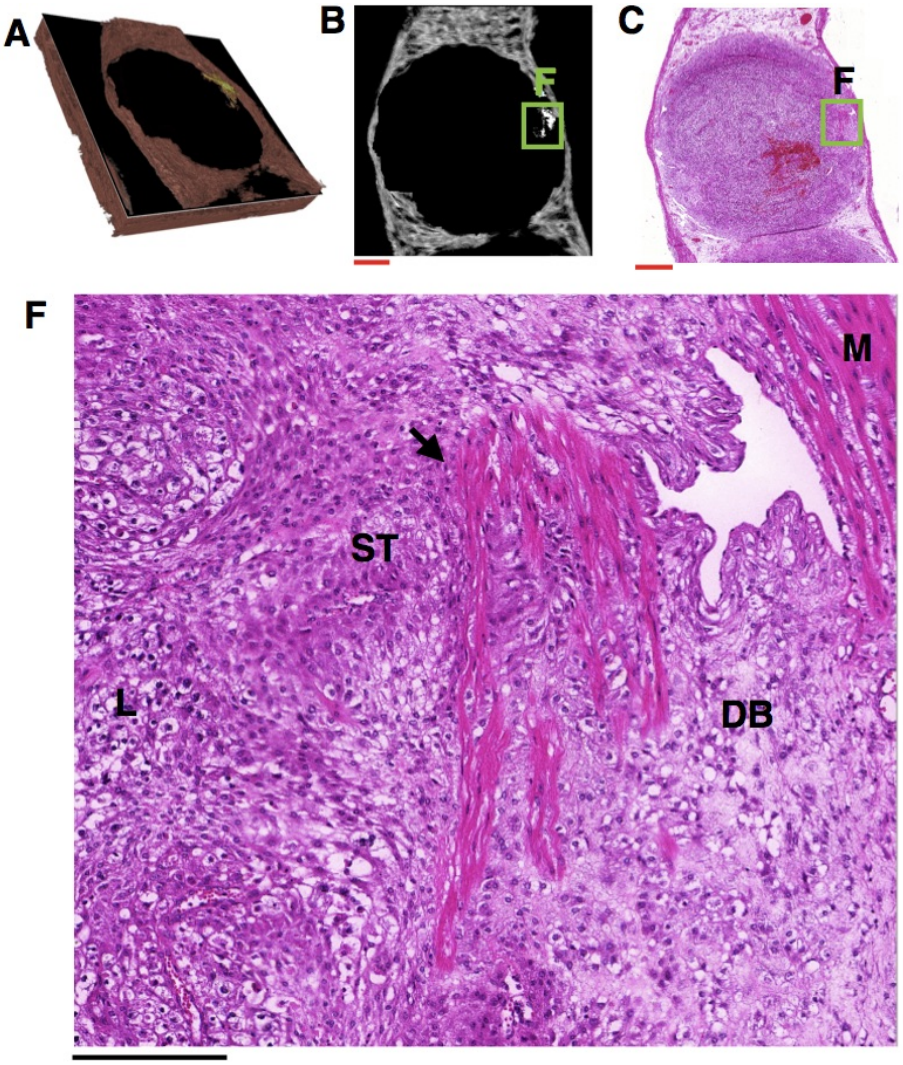

Figure 8.21: Histology of the placenta surrounding the point of deepest penetration of fibre **F1** shown in Figure 8.18. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 4 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high-resolution image of the implantation site surrounding the fibre **F1**, indicated by the arrow. This shows that the fibre penetrates the spongiotrophoblast, but does not reach the labyrinth. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

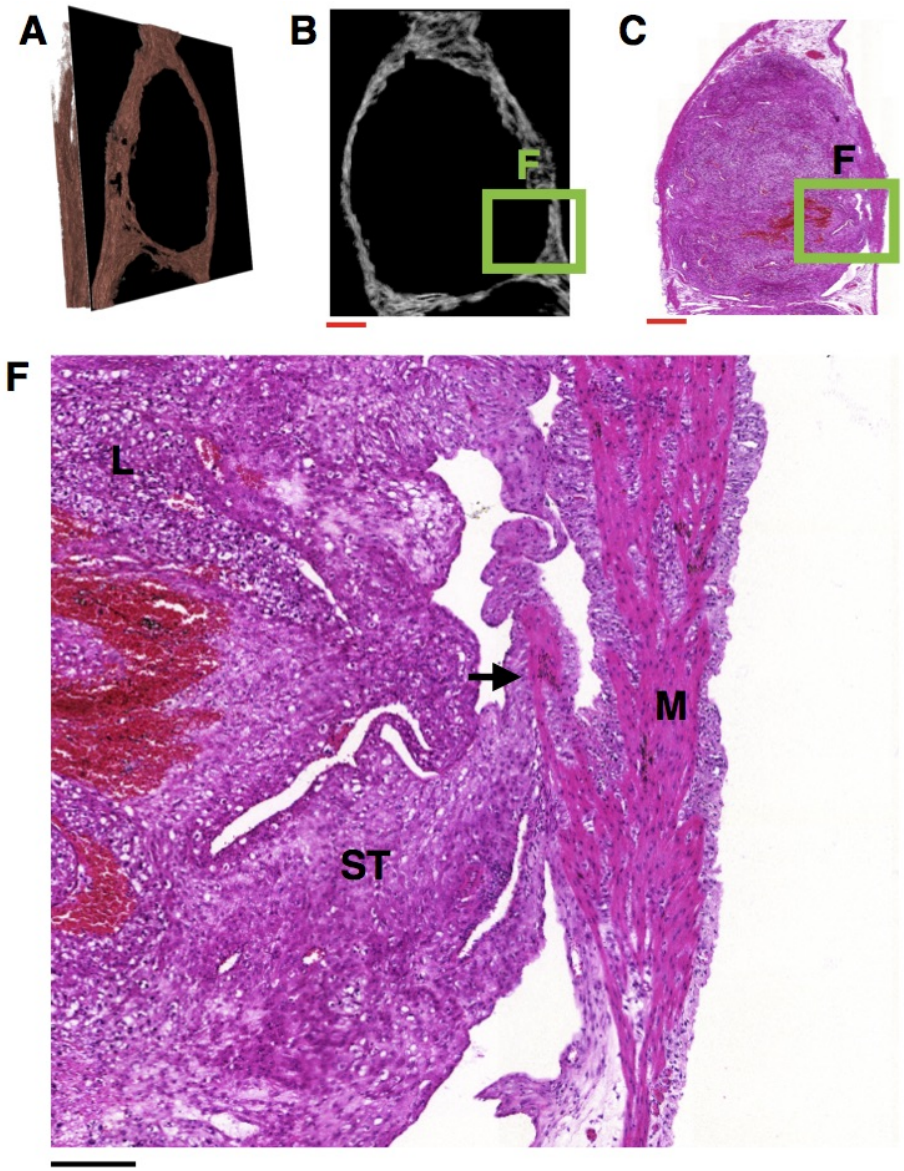

Figure 8.22: Histology of the placenta surrounding the entry point of the fibre **F2** shown in Figure 8.18. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 4 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high-resolution image of the implantation site at the entry point of the fibre indicated by the arrow. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

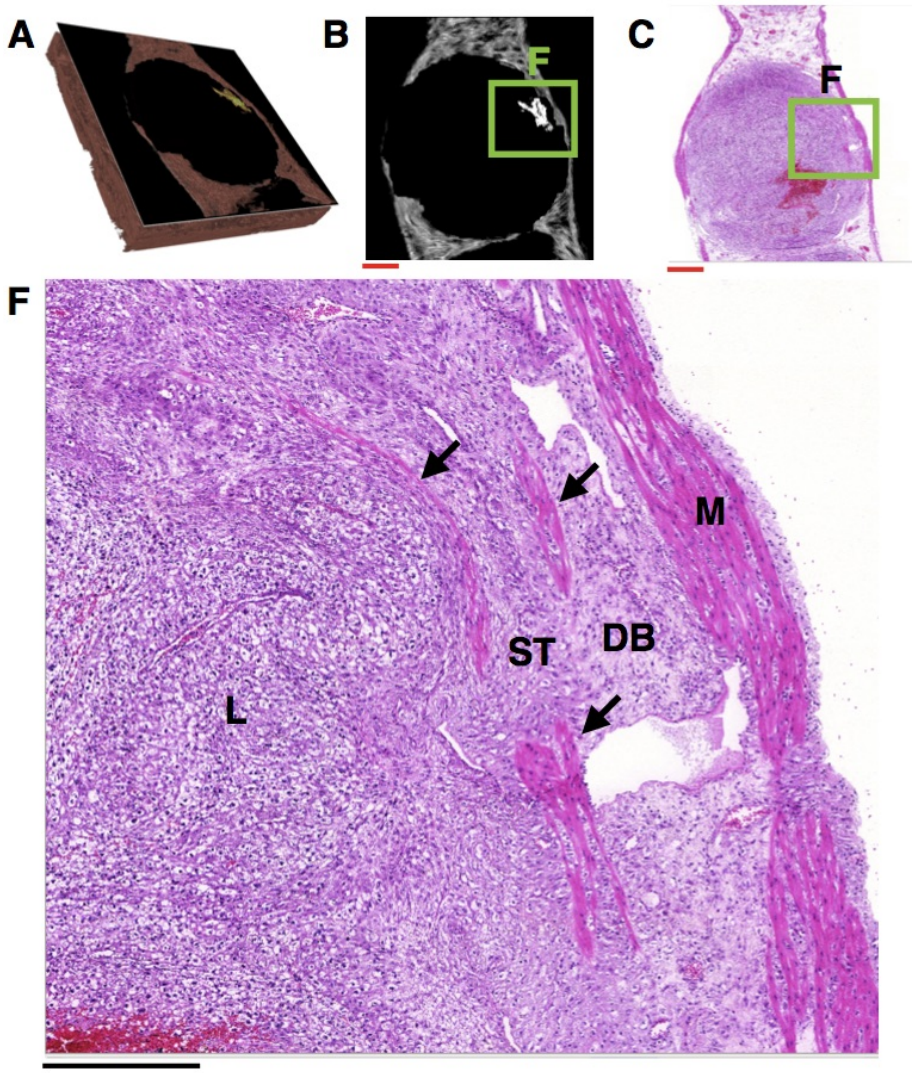

Figure 8.23: Histology of the placenta surrounding the point of deepest penetration of fibre **F2** shown in Figure 8.15. **A**: slide plane in the context of the three-dimensional reconstruction. **B**: two-dimensional slice of the reconstruction corresponding to **C**, with frame indicating the approximate position of **F**. **C**: implantation site 4 is entirely visible in this slide plane, with frame indicating the position of **F**. **F**: high-resolution image of the implantation site surrounding the fibre **F2**, indicated by the arrows. The portion of the fibre on the left appears to attain the deepest inward penetration of the fibre, which is located in the spongiotrophoblast. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bars represent 1 mm, black scale bar represents 200 μm.

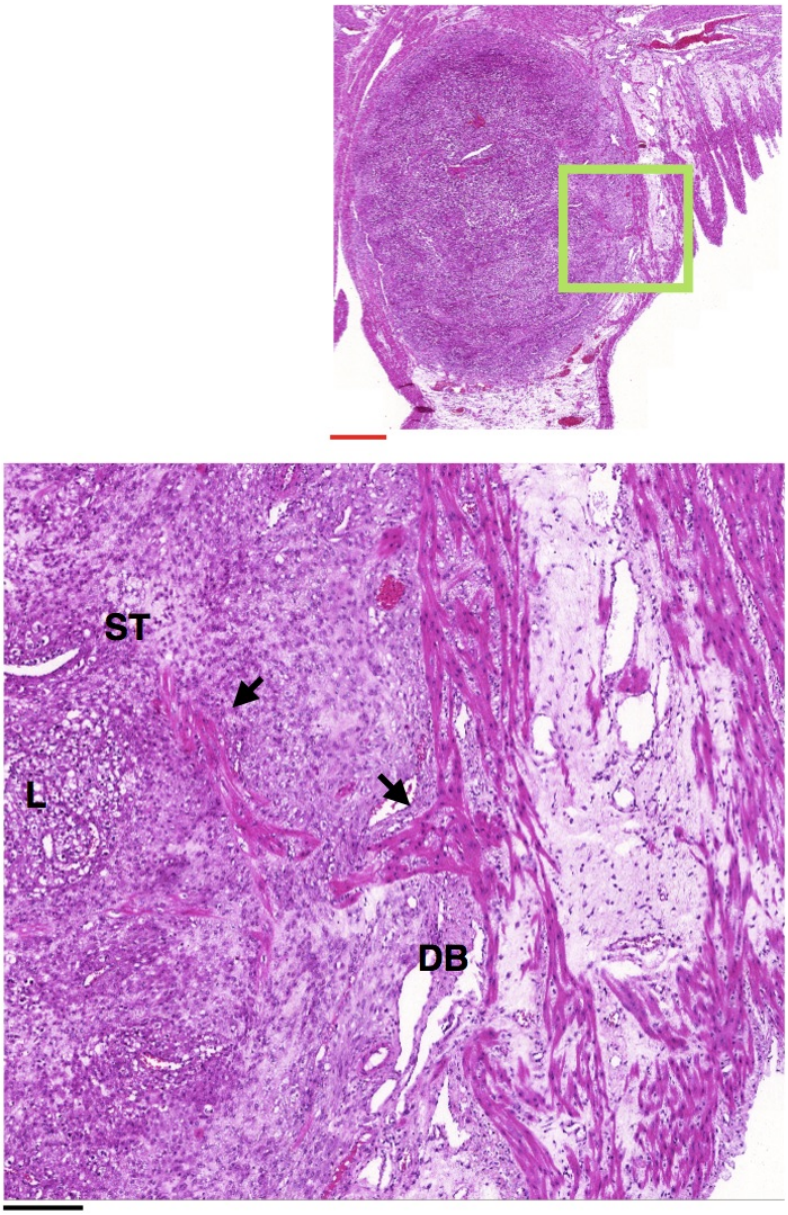

Figure 8.24: Histology of the placenta surrounding a fibre in site 3, which could not be accurately reconstructed in three dimensions. This fibre is located in close proximity to the initiation point of excitation waves, as described in Chapter 6. Top: implantation site 3 is entirely visible in this slide plane, with frame indicating the position of the bottom panel. Bottom: high resolution image of the implantation site at the entry point of the fibre identified in site 3, indicated by the arrows (previously displayed with less detail in Figure 6.17). The arrow on the left indicates the deepest penetration identified, which is located in the spongiotrophoblast. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bar represents 1 mm, black scale bar represents 200 μm.

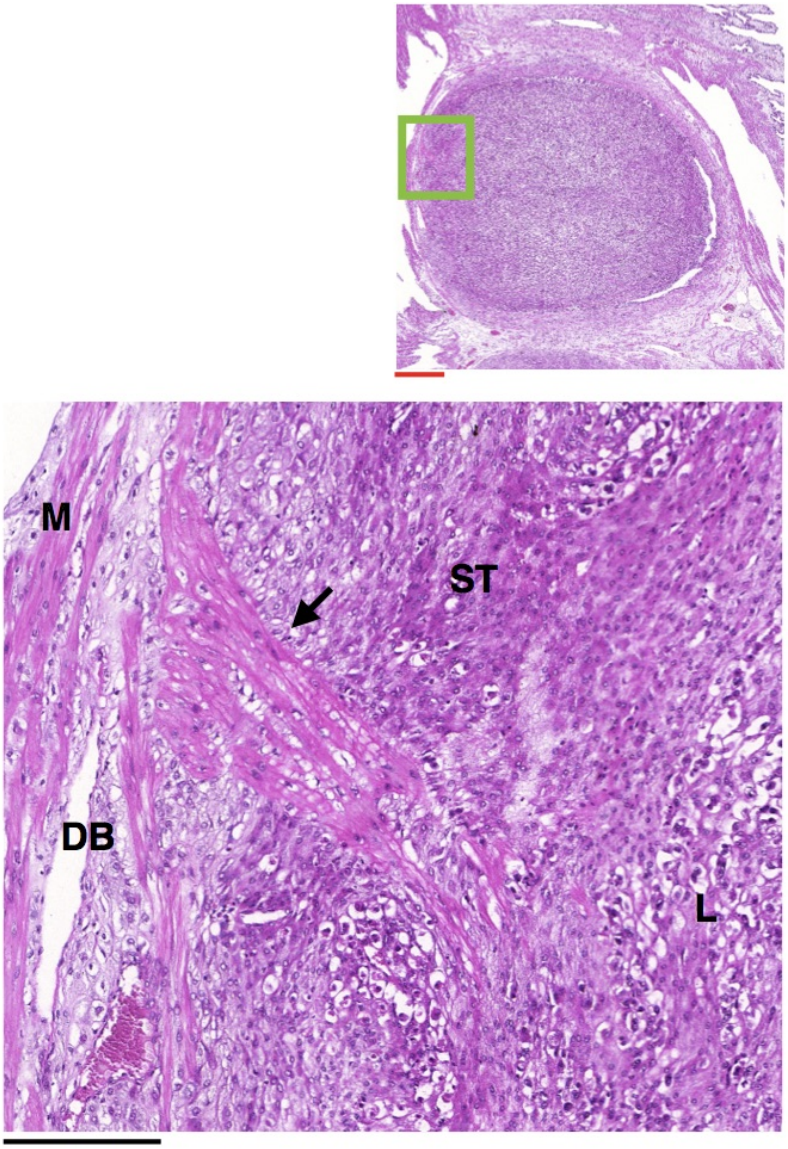

Figure 8.25: Histology of the placenta surrounding a fibre in site 2, which could not be accurately reconstructed in three dimensions. This fibre is located in close proximity to the initiation point of excitation waves, as described in Chapter 6. Top: implantation site 2 is entirely visible in this slide plane, with frame indicating the position of the bottom panel. Bottom: high-resolution image of the implantation site at the entry point of the fibre identified in site 2, indicated by the arrow. This fibre appears to penetrate into the spongiotrophoblast, but not the labyrinth. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast, **L**: labyrinth. Red scale bar represents 1 mm, black scale bar represents 200 μm.

Chapter 9

Numerical simulation of the propagation of excitation in *in silico* reconstructions of the myometrial tissue

The simulation of excitation in cardiac tissue using the electrodiffusion equations given below is well-documented [13, 26, 109]. This has also been used to perform simulations in the myometrium [9, 64, 92]. Sheldon *et al.* [99] showed that excitation propagation can fail when the tissue is too conductive. This phenomenon is investigated here, using values for the conductivity that are in accordance with propagation speeds found in Chapter 6.

The simulations were performed on the reconstruction of the third rat tissue block from Chapter 7. In particular, the first and third recordings presented in Chapter 6 were used for calibration below. These recordings were selected as they represented excitation through a larger proportion of the tissue than other recordings, which suggests that the propagation speeds found from these recordings are the most reliable.

The electrodiffusion equations are as follows:

$$\frac{\partial v}{\partial t} = \nabla \cdot (\mathbf{D} \nabla v) + f(v) - w$$

Parameter	value
A	2
α	0.1
γ	1
ϵ	0.01
v_0 (non-pacemaker)	0.0
v_0 (pacemaker)	0.2

Table 9.1: Table of parameter values used for simulations.

$$\frac{\partial w}{\partial t} = \epsilon(v - \gamma w + v_0),$$

where $f(v)$ is the following cubic function:

$$f(v) = Av(1 - v)(v - \alpha),$$

where A, α, γ, and ϵ are fixed parameters, v_0 is non-zero in pacemaker locations, and $\mathbf{D}$ is the conductivity tensor, which represents the diffusive anisotropy of the system. The fixed parameters for the simulations are listed in Table 9.1. The locations of pacemaker cells and the conduction constant κ were varied across a range of values. Pacemaker cells are simulated by increasing the value of v_0 at the given location to the value stated in Table 9.1. In Chapter 6 the initiation point of excitation was determined for each electrode array recording. Simulations performed here were initiated in the same approximate locations as observed in the electrode array recordings.

The first step was to determine the relationship between the conductivity constant and the propagation speed. This calibration enables the simulation to be compared to the electrode array recordings presented in Chapter 6, by selecting a conductivity constant which corresponds to the mean propagation speed of any given recording. The next step was to determine the minimum volume of pacemaker cells required for a range of conductivity constants. The remainder of this chapter will consider the cases where κ is constant throughout the tissue. All simulations described here can be found at `https://figshare.com/s/9a54e1ae6699cf713670`.

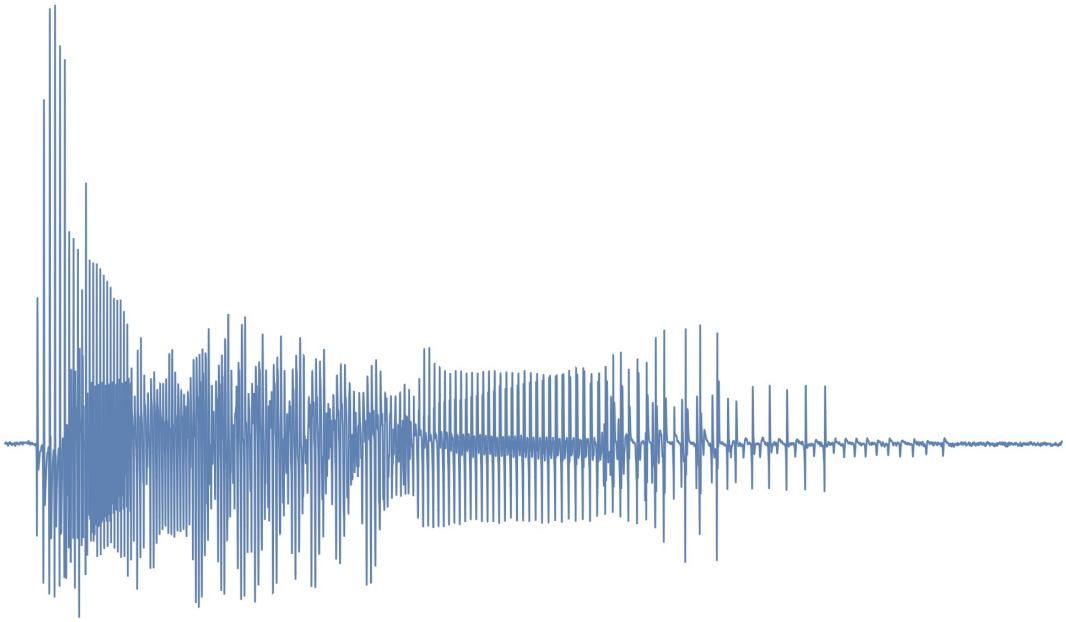

Figure 9.1: An example of a burst recorded by an electrode (see Chapter 6). The frequency of action potentials varies with time. Calibration of the simulation time step to real-time used the high-frequency portion of the burst for comparison with the simulated burst shown in Figure 9.3.

All three-dimensional simulations were visualised using 3D Slicer [38] (www.slicer.org). These reconstructions were facilitated by generating DICOM files using the dcm4che package (www.dcm4che.org). The code used to generate these DICOM files can be found at `https://figshare.com/s/651e592a91f309182edd`.

9.1 Calibration

In order to compare results from simulations to results determined by electrode array analysis, units must be converted in an appropriate manner. Since the simulations were performed on a *in silico* reconstruction of the original tissue, which is directly related to the histological images, the spatial conversion factor is 47.5 μm/voxel length (see Chapter 2). The temporal factor was calculated by comparing simulated and recorded bursts. An example of a recorded burst is shown in Figure 9.1. The time period between action potentials varies with time, as is typical of the recorded bursts presented in 6. The

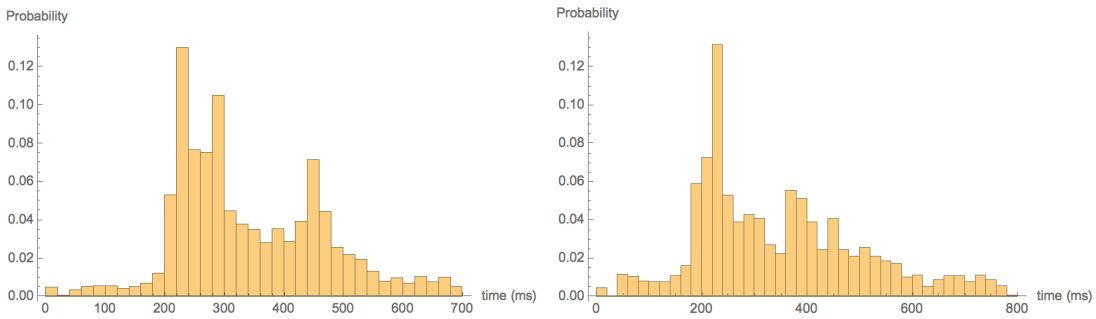

Figure 9.2: The distributions of the time difference between action potentials in a burst. These represent the first and third recordings from the third tissue block.

simulated bursts are chosen to be of a fixed period, and therefore a fixed time period must be determined to represent the recorded bursts in the simulations. The distribution of the time difference between action potentials is shown for the first and third recordings in Figure 9.2. These time periods show multiple modes, which represent separate action potential frequencies. In both cases, the first modal set is at 220 ms, which represents the high frequency portion of the burst. This value was taken to be the representative value for the period of an action potential. The corresponding value for the simulated burst is 77 time units, as demonstrated in Figure 9.3. It follows that one time unit is equal to 2.9 ms.

The final calibration requirement is the relationship between the dimensionless conductivity constant κ and the conduction speed measured in the electrode array analysis. This relationship was determined by measuring the time taken for a simulated excitation wave to travel between the two points indicated in Figure 9.4. The excitation waves in these simulations travelled along the mesometrial border fibres. While this form of excitation propagation was not observed in the electrode array recordings (see Chapter 6), the connective isolation of these fibres reduces the likelihood of propagation failure through dispersal of current into the surrounding tissue, as previously observed [99]. The resulting

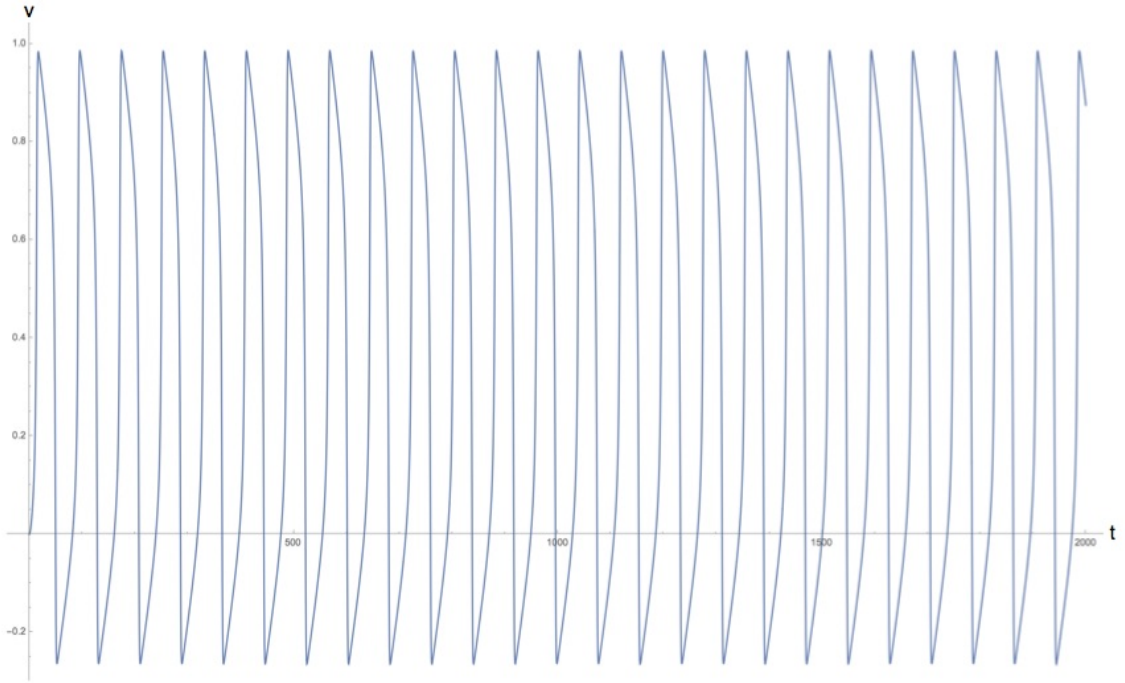

Figure 9.3: Simulated action potentials without diffusion at a point (i.e. $\kappa = 0$). Repeated excitation is simulated by setting $v_0 = 0.2$. A total of 26 action potentials cycles are completed between 0 and 2000, giving a time difference of 77 dimensionless time units between each action potential.

calibration curve is shown in Figure 9.5. The logarithmic scale shows the relationship

$$\log(t) = 5.1 - 0.37 \log(\kappa).$$

This relationship, combined with the distance between the two points and the time scale factor, yields the following relationship:

$$v = 0.74\kappa^{0.37}$$

where v is the propagation speed in cm s^{-1}.

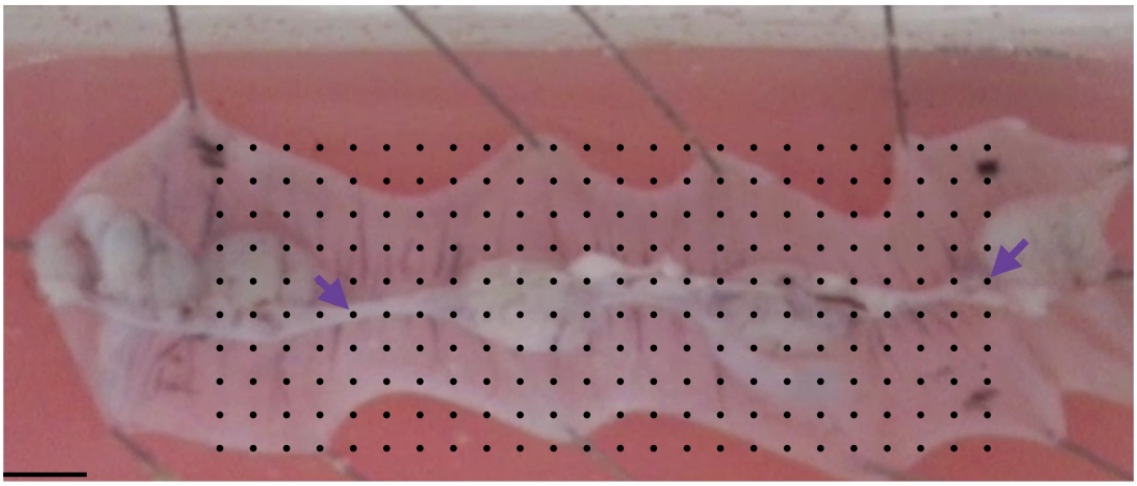

Figure 9.4: The points used for calibration, indicated by the arrows. The simulated excitation was initiated at the left calibration point, and the time taken to reach the right calibration point was recorded. The distance between these points is 20 mm, and the time difference was converted from simulation time units to seconds using the time scaling factor given in the text. Scale bar represents 5 mm.

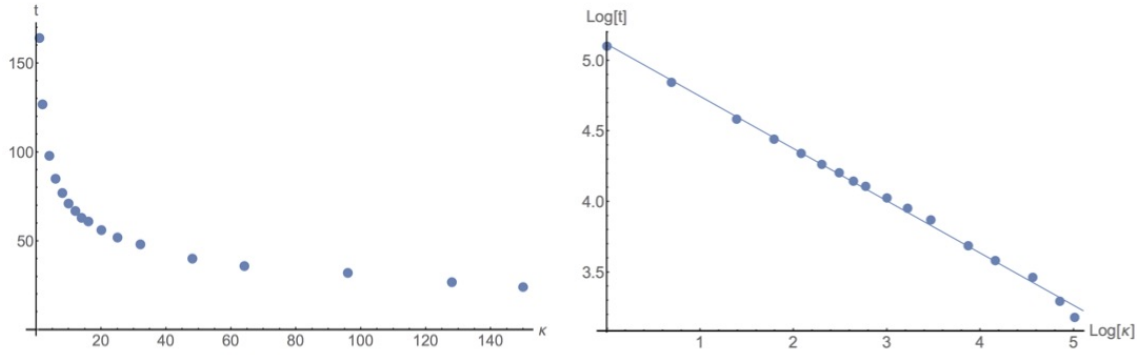

Figure 9.5: Calibration curves for determining the relationship between the conduction constant κ and the time taken to travel a fixed distance in dimensionless time units.

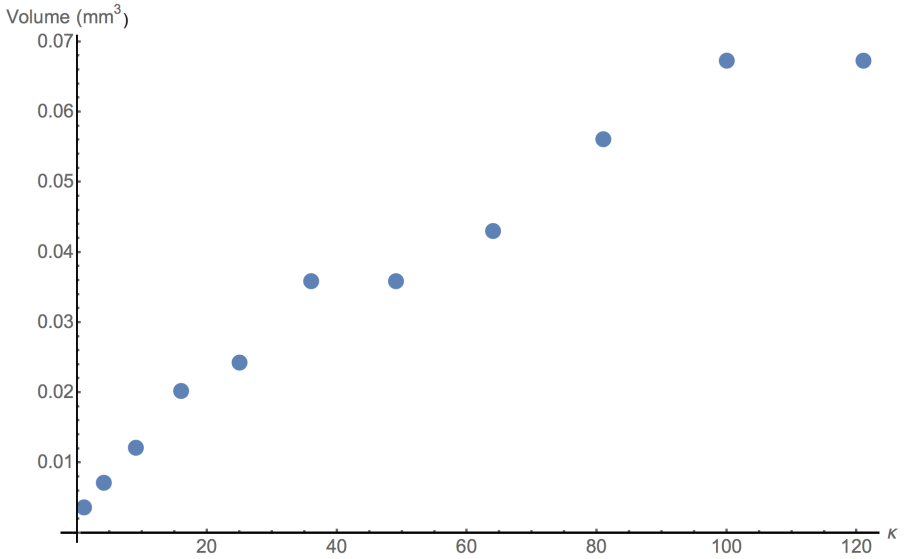

Figure 9.6: Minimum volume required for a group of pacemakers to initiate an excitation wave for a given κ.

9.2 Minimum volume

As described by Sheldon *et al.* [99] the ability for excitation to propagate from a source is dependent on the magnitude of the initial excitation and the value of the coupling constant, which is represented in these simulations by κ. A low value of κ will prevent excitation propagation by blocking conduction, while a high value of κ will cause surrounding tissue to act as a current sink, preventing the formation of an action potential.

In order to evaluate this effect on values of κ, brief simulations were performed for fixed values of κ with pacemaker region situated in an ellipsoid with a given volume. The excitation wave was deemed to have sufficiently spread if a wavefront was present outside the initial ellipsoid. The minimum volume is plotted against the value of κ in Figure 9.6. This shows that, whereas low values of κ corresponding to speeds of up to 0.74 cm s^{-1} allow the excitation of the whole tissue from a single voxel, the values of κ corresponding to speeds in the range found in Chapter 6 require a much larger initiation volume. This would suggest that either pacemaker cells are densely clustered over a large volume to initiate

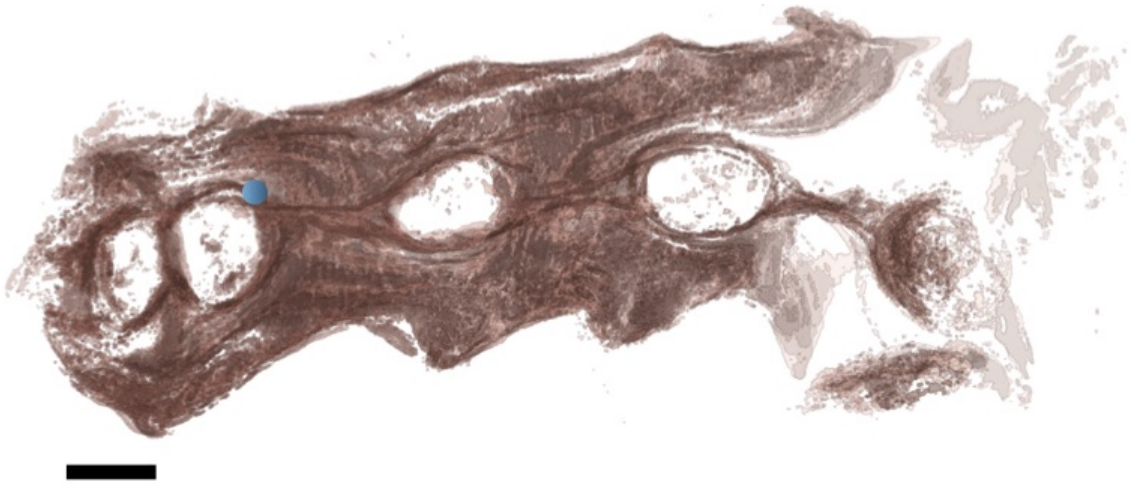

Figure 9.7: The pacemaker area for simulations is indicated by the blue circle. Scale bar represents 5 mm.

excitation, or that the conductivity around the pacemaker cells is reduced to facilitate a excitation from a smaller cluster. The longitudinal interpenetrating fibres described in Chapter 8 are embedded in the placenta, which means that, with the exception of the point of entry to the placenta, these fibres are isolated from the longitudinal fibres. This isolation could prevent the current sink effect observed here, which would enable an excitation wave to form in the interpenetrating fibres.

9.3 Constant conductivity

The foregoing sections presented approximate conditions required to simulate the excitation wave for a given speed. These parameter values were used to simulate conductivity in the tissue for a range of speeds. The set of pacemaker voxels present in these simulations was a sphere of radius 16 voxels (~ 1 mm) in the approximate location of observed excitation (see Chapter 6), as indicated in Figure 9.7. This volume was used because it is sufficiently large to prevent any current sink artefacts observed in Section 9.2. The aspect ratio of the conductivity tensor was set to 0.3.

Snapshots from the simulation for $\kappa = 500$ are shown in Figure 9.8, using the lookup

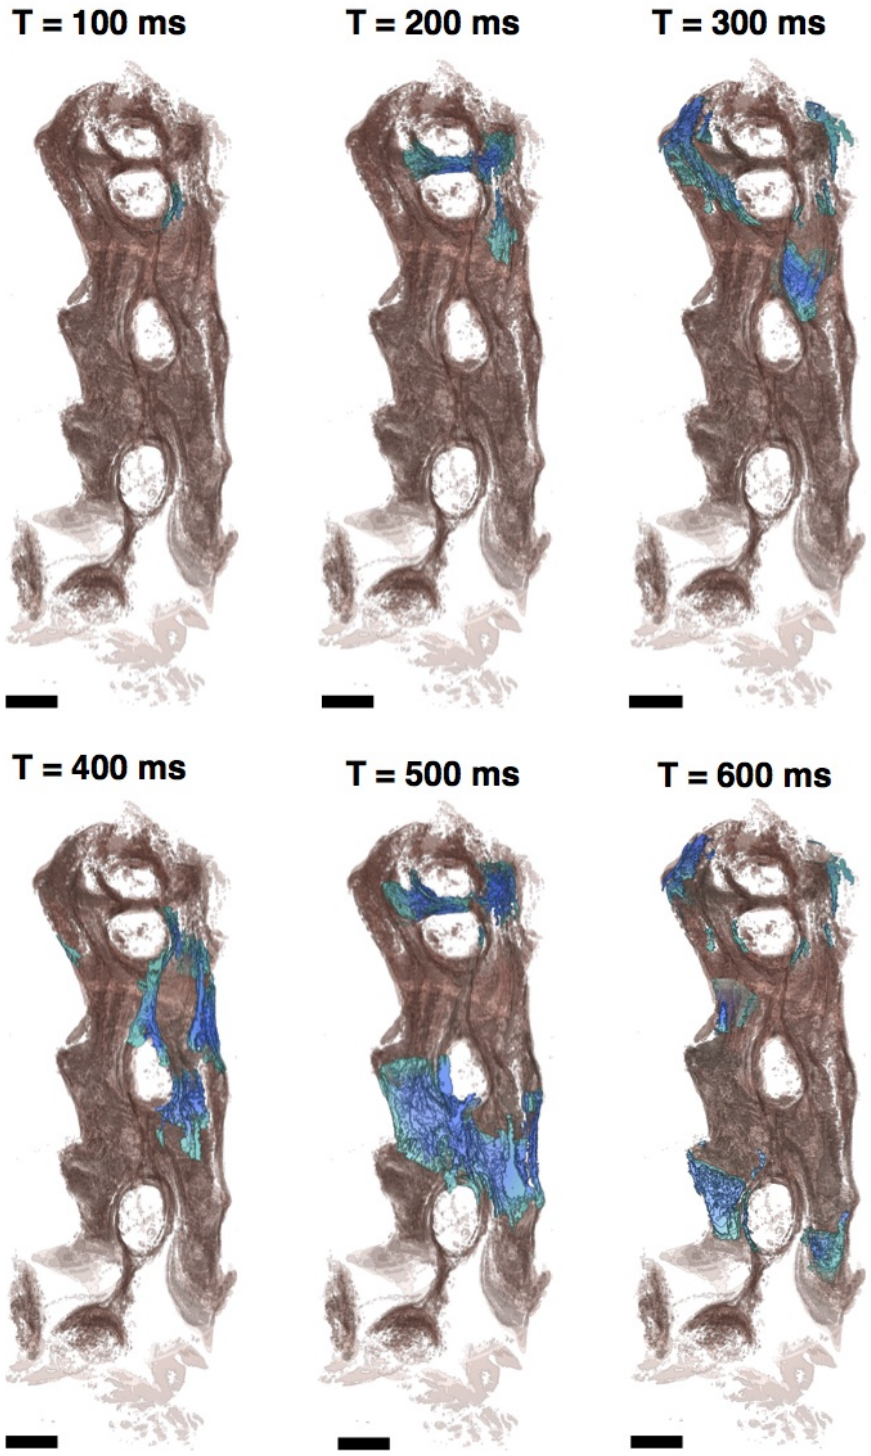

Figure 9.8: Simulation frames for $\kappa = 500$. The ovarian end of the tissue is at the top of all images. Green shows the position of the excitation wavefront, blue shows the trailing recovery of the wave, inactive tissue is shown in red. Scale bar represents 5 mm.

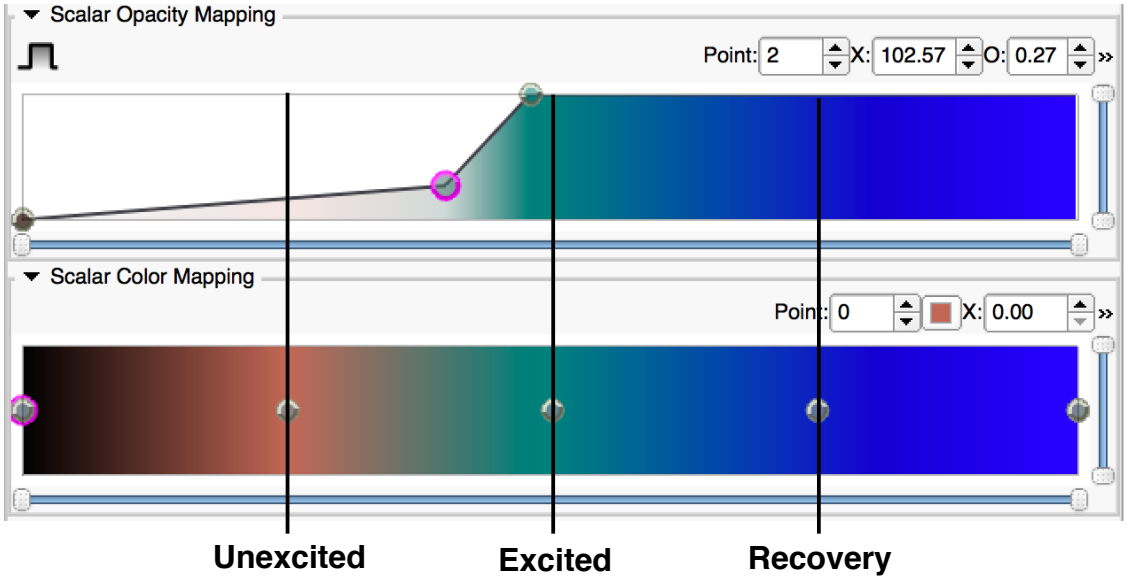

Figure 9.9: Simulation opacity (top) and colour (bottom) lookup tables.

tables shown in Figure 9.9. The propagation speed for this value of κ is given by the calibration function in Section 9.1 to be 7.4 cm s $^{-1}$. Figure 9.10 shows the excitation behaviour of the tissue in the pacemaker volume and at points along the path of the excitation wave, confirming that the action potential shape is maintained. The propagation behaviour of this simulation is shown in Figure 9.11.

The excitation wave propagates from the initiation site at the ovarian end of the tissue to the cervical end of the tissue via the longitudinal layer as indicated by the long arrow in Figure 9.11. This was also observed in the electrode array recordings presented in Chapter 6. The time taken for the excitation wave to reach the cervical end of the tissue in the simulated excitation (600 ms) does not match that of the electrode array recordings ($\sim$ 15 s); however, comparison of the recordings shown in Figure 9.12 reveals that the longer time delay observed in the electrode array recordings could correspond to a failure of propagation during the first $\sim$ 15 seconds. This delayed propagation could be caused by variations in conductivity of the tissue during the time of recording.

The simulated excitation wave propagates across the mesometrial border through the

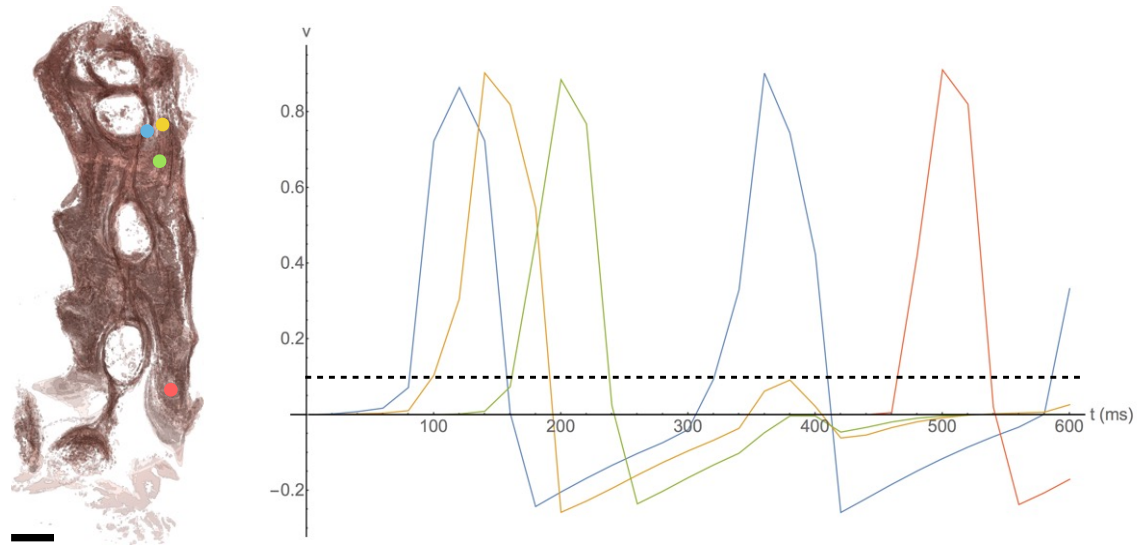

Figure 9.10: Excitation behaviour at points along the path of the simulated excitation wave. The left-hand panel shows the locations sampled to obtain the action potentials shown in the right-hand panel, where each point is colour-coded to match the colour of the corresponding sampled recording. The dashed line represents the threshold value of the FitzHugh-Nagumo system. The blue curve was sampled from the pacemaker region in the simulation, exhibiting bursting behaviour with period 220 ms, as specified. The next point sampled in the path (yellow) shows a failure to achieve the threshold on the second action potential, causing all points further along the path to fail to attain the second action potential. The shape of the action potentials from reaching the threshold value through to the start of recovery is maintained along the path.

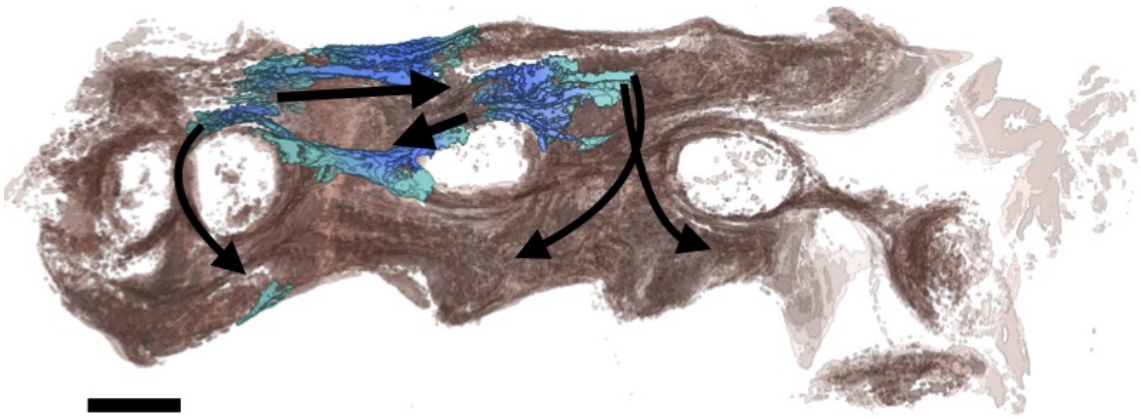

Figure 9.11: The paths of propagation in the simulation with $\kappa = 500$. Long arrow: excitation propagation through the longitudinal layer toward the cervical end of the tissue. Short arrow: the excitation wave passing into the mesometrial border through fibres adjacent to an implantation site. Curved arrows: excitation propagating across the mesometrial border via the circular layer. The ovarian end of the tissue is on the left in this image. Scale bars represent 5 mm.

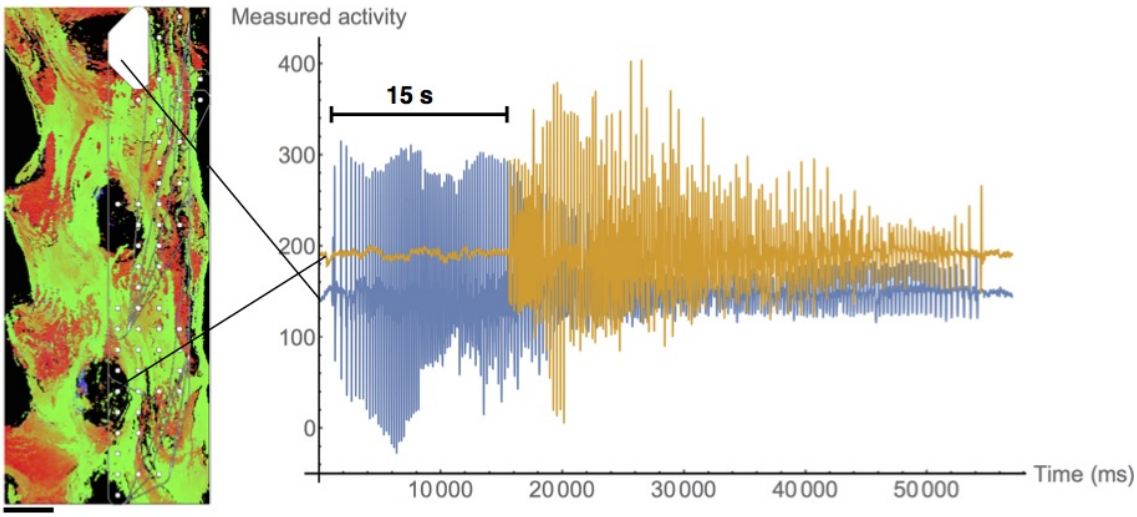

Figure 9.12: Electrograms from the burst origin at the ovarian end of the tissue (blue) and an electrode toward the cervical end of the tissue (orange). The later burst starts approximately 15 s after the original burst. However, the action potential frequency in the later burst closely resembles that of the original burst. This suggests that the excitation wave fails to reach the later electrode during the first 15 s, after which point propagation occurs at a speed similar to that generated by simulated excitation. Scale bars represent 5 mm.

circular layer ($T = 200$ ms and $T = 600$ ms), as indicated by the curved arrows in Figure 9.11. This is facilitated by the connecting bridges observed in Chapter 7. This excitation wave does not re-enter the longitudinal layer. This suggests that this activity could feasibly have occurred during the electrode array recordings, since the recordings were only made in the longitudinal layer. Alternatively, this behaviour could be an artefact of spatial homogeneity of κ, since a reduction in conductivity through these bridges could prevent this event.

The simulated excitation wave enters the mesometrial border fibres, as indicated by the short arrow in Figure 9.11. This does not agree with the electrode array recordings, which show no such behaviour, suggesting that this behaviour is inaccurate. This inaccuracy could be caused by reconstruction errors close to the implantation site, as discussed in Chapter 7. Alternatively, there could be a lower level of conductivity in the mesometrial border fibres, which would prevent propagation into the mesometrial border fibres.

The results presented in this chapter provide parameters for simulating excitation be-

haviour in rat myometrium based on electrode array recordings. The simulated excitation based on these parameters exhibits behaviour similar to that seen in multi-electrode array recordings. Future implementations of this technique should include spatial and temporal variation of conductivity to improve this correspondence.

Chapter 10

Discussion

In this chapter the results presented in previous chapters are summarised, implications of the research done are surveyed, and future work is proposed. The structural features identified in the *in silico* reconstruction of rat uteri which could affect propagation of excitation are discussed in Sections 10.1–10.3. Section 10.4 draws on the discussion in the preceding sections to identify the overall effect of these features. Section 10.5 discusses the general characteristics of the myometrium in both rat and human uteri as uncovered in this thesis, and puts these findings in perspective given what had previously been reported in the literature. The computational methods developed and applied in this thesis and possible future extensions are discussed in Section 10.7.

10.1 Mesometrial border connectivity

Low conductivity in myometrial fibres located in close proximity to the placenta has previously been observed in rat [60] and cat [29] uteri. The excitation propagation along the mesometrial border in guinea pigs is sparse and fragmented [66]. Previous work on the conductivity in rat myometrium has shown the existence of transient conduction blocks [65], where propagation through some areas of the tissue was prevented intermittently. These observations are in keeping with the transient conduction blocks that were observed in the electrode array analysis presented in Chapter 6. Additionally, a fixed boundary which appeared to prevent propagation was demonstrated in all uteri along the mesometrial border. Inspection of the *in silico* reconstruction of the tissue blocks revealed a reduced level of

connectivity between longitudinal fibres along the mesometrial border and the remainder of the longitudinal layer. This low level of connectivity was also observed in the histological slides, and could explain why excitation fails to propagate across the mesometrial border.

The permanent connective isolation of the longitudinal fibres along the mesometrial border could serve to prevent synchronised contraction across the mesometrial border prior to labour, in a similar manner to that suggested Kanda and Kuriyama [60]. Small, localised bursts of excitation in uteri of rats in mid-gestation have previously been noted [65]. If these localised excitation events were to propagate across the mesometrial border, this would cause synchronised contraction across the mesometrial border. Synchronised contraction across the mesometrial border could lead to a localised increase in pressure across an implantation site, which in turn could lead to the placenta becoming dislodged. The connective boundary observed here could serve to prevent this from happening prior to parturition.

The boundary along the mesometrial border could not be accurately defined, since in close proximity to the implantation sites, the presence of nuclei in the placental tissue reduced the accuracy of the reconstruction at these locations. Additionally, the level of connectivity away from the implantation sites could not be precisely determined, due to the reconstruction being of a lower resolution than cell width, as discussed below in Section 10.7.1. A higher resolution reconstruction is proposed in Section 10.7.1, which could be used to resolve this imprecision.

If an analogous structure were present in human myometrium, it could serve a similar purpose to the proposed mechanism outlined above. Normal implantation occurs along the midline of the upper part of the uterus in humans [113]. If a boundary to excitation were to perform a similar function in the human uterus to the rat, then this boundary should form along the midline of the uterus, and possibly only in the upper part of the uterus. No such boundary could be identified in the human uterus examined in this thesis;

however, the volume of the tissue used was bounded at a location in close proximity to this midline, and therefore a connectively isolated region could be present outside of the sampled volume.

Previous work has shown that excitation does propagate across the midline of the human uterus [52]. Thus, if there is a boundary of low connectivity along the midline of the uterus, it may not be present along the entirety of the midline in the outer myometrium. Moreover, lateral propagation has been detected in more instances in the lower portion of the uterus than in the upper portion [52]. This could indicate that a boundary of low connectivity is present only in the upper portion of the uterus.

Malformations of the human uterus caused by incomplete fusion of the Müllerian ducts during development, such as the bicornuate uterus, could prevent the formation of this conductive barrier: the proposed location of the barrier is along the line of fusion of the Müllerian ducts, and therefore would not be present if this fusion failed [27]. These types of malformations are associated with preterm birth [3]. Absence of a connective boundary across the implantation site could lead to highly coordinated excitation across the placenta in pregnancy during preterm excitation, which could produce contractile forces leading to preterm delivery. This mechanism could also be implicated in preterm birth associated with lateral implantation and placenta previa, both of which involve placentation away from the upper midline of the uterus [42].

10.2 Connectivity between longitudinal and smooth layers

Communication between circular and longitudinal layers in rat myometrium has previously been observed [24, 87]. In addition, previous research has shown the presence of connective bridges between the circular and longitudinal layers in pregnant ewes [89]. Bridges between the circular and longitudinal layers were identified here in rat uteri, as described in Chapter 7. These connecting bridges were all verified by visual inspection of the tissue. All potential bridges which could not be verified in this manner were discounted.

The verified bridges were distributed throughout the tissue, suggesting that the layers are connected at numerous points. The bridges tended to form in clusters, as shown in Figure 10.1. The clustering of the bridges suggests a high level of connectivity at these bridge points. The bridges were 20–100 μm in width, and observed to have lengths up to 1 mm.

Cell-to-cell coupling in the myometrium is mediated by gap junctions [44]. The density of these gap junctions in rat myometrium increases dramatically on day 21 of pregnancy, and decreases to preterm levels post-partum [46]. An increase in gap junctions in the connecting bridges at term could increase cell-to-cell coupling in these bridges, which would facilitate the spread of excitation from the longitudinal layer to the circular layer, as demonstrated in the simulations mentioned above.

Previous work by Chen and Chiang [24] has shown that contraction in the circular and longitudinal layers appears to be coordinated from day 21 of pregnancy, with less coordinated contractions present earlier in gestation. Rat uteri examined in this thesis were taken prior to day 21 of pregnancy, which means that these tissue samples do not exhibit this coordinated contraction. However, bridges connecting the circular and longitudinal layers were present in all uteri, which suggests that these bridges are non-conducting prior to day 21 of pregnancy. The coordination could be a result of increased cell-to-cell coupling in the connecting bridges.

Gap junction permeability has previously been observed to be regulated by various molecular factors [44]. These factors could regulate the permeability of gap junctions locally in the connecting bridges, which in turn would regulate conductance of the connecting bridges. These localised effects would regulate the ability for excitation to pass between layers in the myometrium, which would have the effect of regulating the global excitation behaviour of the myometrium, as discussed in Section 10.4.1.

Gap junctional channels in the myometrium can be formed from a range of different types of connexin proteins [63]. Differing composition of connexin proteins can lead to

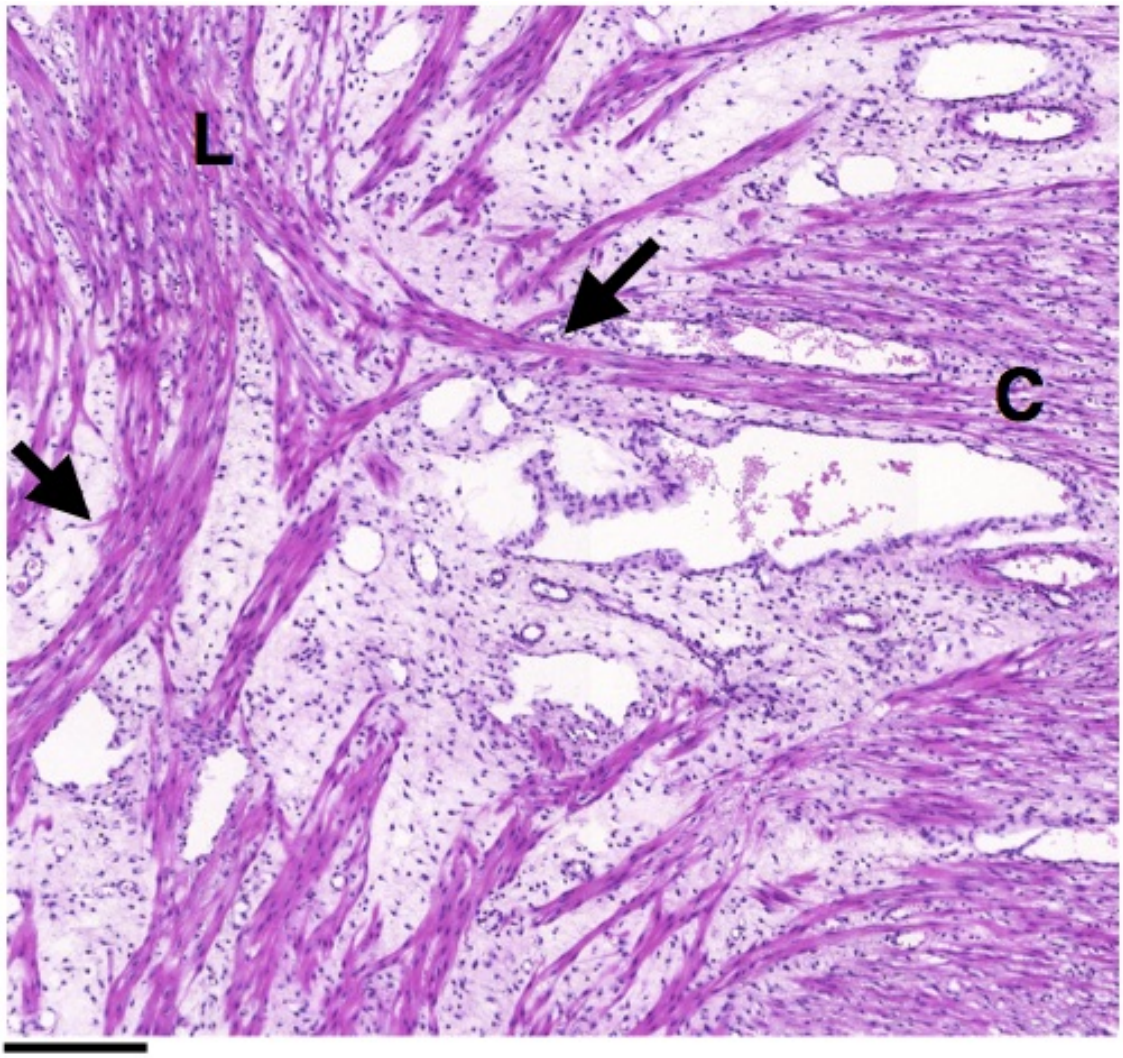

Figure 10.1: An example of a cluster of bridges connecting the circular and longitudinal layers in rat myometrium (previously shown in Figure 7.23). This is the typically observed formation of connecting bridges, with several bridges connecting the longitudinal and circular layers. The bridges have widths in the range 20–100 μm, and have been observed to have lengths of up to 1 mm. The arrow on the left of the image points to a bridge of width 100 μm, while the arrow in the middle of the image indicates a bridge with an approximate length of 1 mm. **L**: Longitudinal layer, **C**: circular layer. Scale bar represents 200 μm.

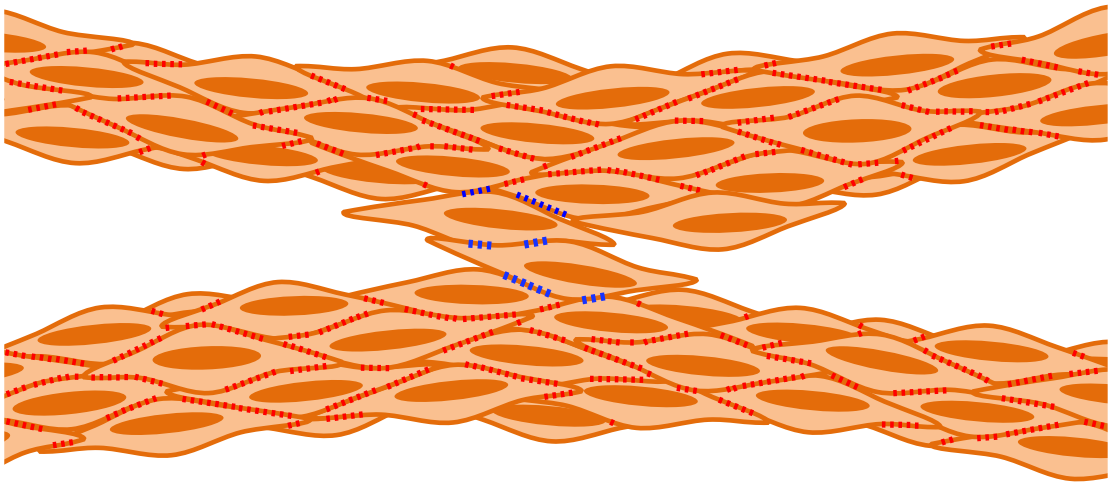

Figure 10.2: Bridges connecting circular and longitudinal layers may contain gap junctions of different permeability to those in the rest of the tissue. Reduced permeability in the gap junctions in the connecting bridge would block the passage of excitation waves through the bridge. Control of the permeability of gap junctions within these connecting bridges could enable control of communication between the layers. Blue dots represent gap junctions in a connecting bridge, red dots represent gap junctions in the circular and longitudinal layers.

different excitation behaviour [81]. Furthermore, the composition of connexin proteins in the gap junctional channel determines the permeability response to specific molecular controllers [7]. The gap junctions in the connecting bridges observed here could be constituted by an altered mixture of connexin proteins as compared to the circular and longitudinal layers, as illustrated in Figure 10.2. This could enable the precise modulation of excitation behaviour of these connecting bridges by exploiting the unique properties of the gap junctional channels present in the connecting bridges. The connexin make-up of the gap junctions in the connecting bridges could be observed using techniques described by Kilarski *et al.* [63], which could lead to a better understanding of the conductive behaviour of these bridges.

The accurate verification of bridges between the circular and longitudinal layers was not always possible; in some of the locations inspected the bridges could not be confirmed to connect in the slide plane, but may connect between slide planes. This could mean that not all locations of bridges were recorded. The identification of all bridges would require

the ability to observe them in three dimensions. The *in silico* reconstruction presented in this thesis has voxel length ~ 50 μm, which is larger than the smallest width of these connecting bridges (20 μm). A method for generating a higher resolution reconstruction is proposed in Section 10.7.1. This higher resolution reconstruction could provide a means for observing the detailed architecture present in these connecting bridges.

In the human myometrium the inner layer is formed of circular fibres, while the outer portion is more disordered, with the exception of the outermost layer, which contains longitudinal fibres [111]. An analogous structure to the connecting bridges found in rat uteri would be bridges joining the inner layer to the outer layer. Such structures are present in the human myometrium examined in this thesis. The precise identification of the locations of these connective bridges in the human was complicated by the inherent structure of the myometrium: the disorderly arrangement of fibres in the outer myometrium leads to numerous instances of circular fibres in the outer myometrium. This prevented the accurate identification of the boundary between inner and outer myometrium. The separation of inner and outer myometrium is observable in MRI scans [19], which suggests that it would be possible to identify the boundary between the layers accurately, and hence the fibres connecting them, by comparison of a reconstruction of the form presented in this thesis and MRI scans.

10.3 Longitudinal interpenetrating fibres

Pacemaker locations in rat myometrium have been identified previously by Lammers *et al.* [67] at the ovarian end of the uterus and adjacent to the implantation sites. The initiation points found in Chapter 6 were located in close proximity to implantation sites, with one exception, which appeared to be located at the ovarian end of the tissue. These findings are in accordance with those of Lammers *et al.* [67]. Histological examination revealed fibres entering the placenta from the longitudinal layer in close proximity to the initiation points located near the implantation sites.

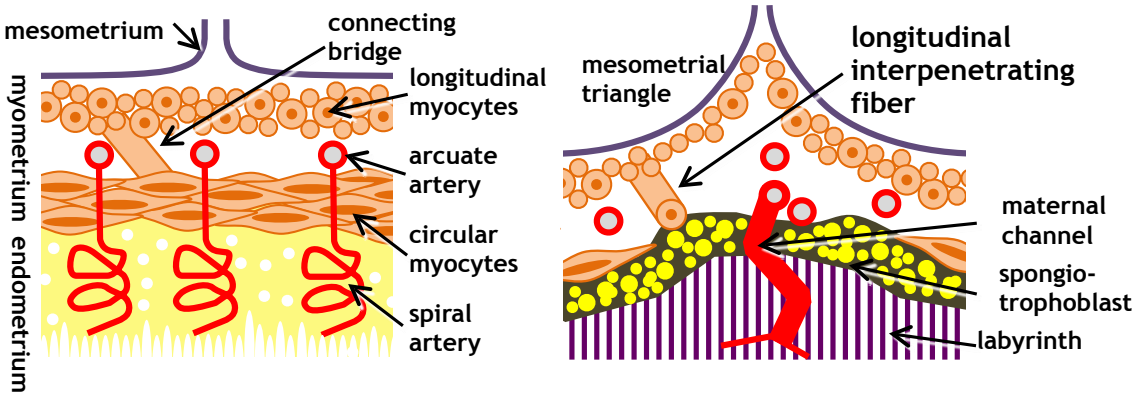

Figure 10.3: Comparison of the uterine wall before (left panel) and after (right panel) the development of the maternal-embryonic interface. The circular layer has been largely dispersed during the course of pregnancy, although it is still present in the decidua basalis, as can be seen in Figure 10.5. The longitudinal layer has been displaced to line the mesometrial triangle. Longitudinal interpenetrating fibres (right panel) extend into the spongiotrophoblast, and are generally located in close proximity to the decidualized circular cells. Bridges connecting circular and longitudinal layers were observed in the pregnant myometrium, and are conceivably present prior to implantation (left panel). Structural similarities between these connecting bridges and the interpenetrating fibres suggest the interpenetrating fibre could be derived from a connecting bridge.

The progressive incorporation of myometrium into the decidua basalis from day 10 of pregnancy onwards had been noted in rat uteri [31]. This process involves the formation of a distinct layer of decidua containing myometrial smooth muscle, separated from the spongiotrophoblast by a layer of decidua which does not contain myometrium. The invasion of small parts of the longitudinal layer by the trophoblast has previously been observed at days 20–21 of pregnancy [108].

The interpenetrating fibres observed in this thesis extend from the longitudinal layer into the spongiotrophoblast, as shown in the right-hand panel of Figure 10.3. These fibres have widths within the range of 100–200 μm at the point of entry, and have lengths within the range of 0.5–3 mm. The interpenetrating fibres was observed to extend into the placenta from both sides of the mesometrial border, and from the ovarian, cervical, and lateral directions. This suggests that the point of entry is not subject to tight regulation during placental development.

The longitudinal interpenetrating fibres could not be identified in all implantation sites

for any of the tissue blocks. However, in most of the implantation sites where the fibres could not be identified, there were structures which could correspond to interpenetrating fibres but which could not be accurately confirmed as being such structures. This ambiguity was due to either uncertainty with regards to the layer of the myometrium to which they were attached, or the slides in which the structures were observed were adjacent to a large sub-stack of slides (~ 10) which had been discarded during sectioning for the purposes of quality control, which means that the slides containing the full extent of any possible fibre was lost.

The cells in the interpenetrating fibres appear to be stellate, having thin extensions of the cytoplasm, as shown in Figure 10.4. Cells with pacemaker-like properties present in gastrointestinal tissue, known as interstitial cells of Cajal, also have this stellate shape [96]. This could indicate a functional relationship between the interpenetrating fibres observed here and interstitial cells of Cajal, specifically, the cells in the interpenetrating fibres could serve as pacemaker cells in the rat myometrium. However, the precise morphology of the cells in the interpenetrating fibres cannot be determined from the slides presented here, because the staining does not accurately show cell boundaries. More detailed observations could be made using electron microscopy and alternative stains, as was performed by Duquette *et al.* [35].

The precise origin of the longitudinal interpenetrating fibres cannot be determined from the samples used in this thesis, since all uteri were 19–20 days pregnant, and therefore the developmental process leading to the formation of these fibres in earlier pregnancy could not be observed. However, the interpenetrating fibres do show structural similarities with the bridges connecting the circular and longitudinal layers discussed in Section 10.2, which could indicate that these fibres are formed through remodelling of the connecting bridges. All but two of the fibres were observed to enter the placenta in close proximity to decidualized circular smooth muscle, as shown in Figure 10.5. The proximity of the interpenetrating fibres to the circular layer in the decidua basalis is consistent with the

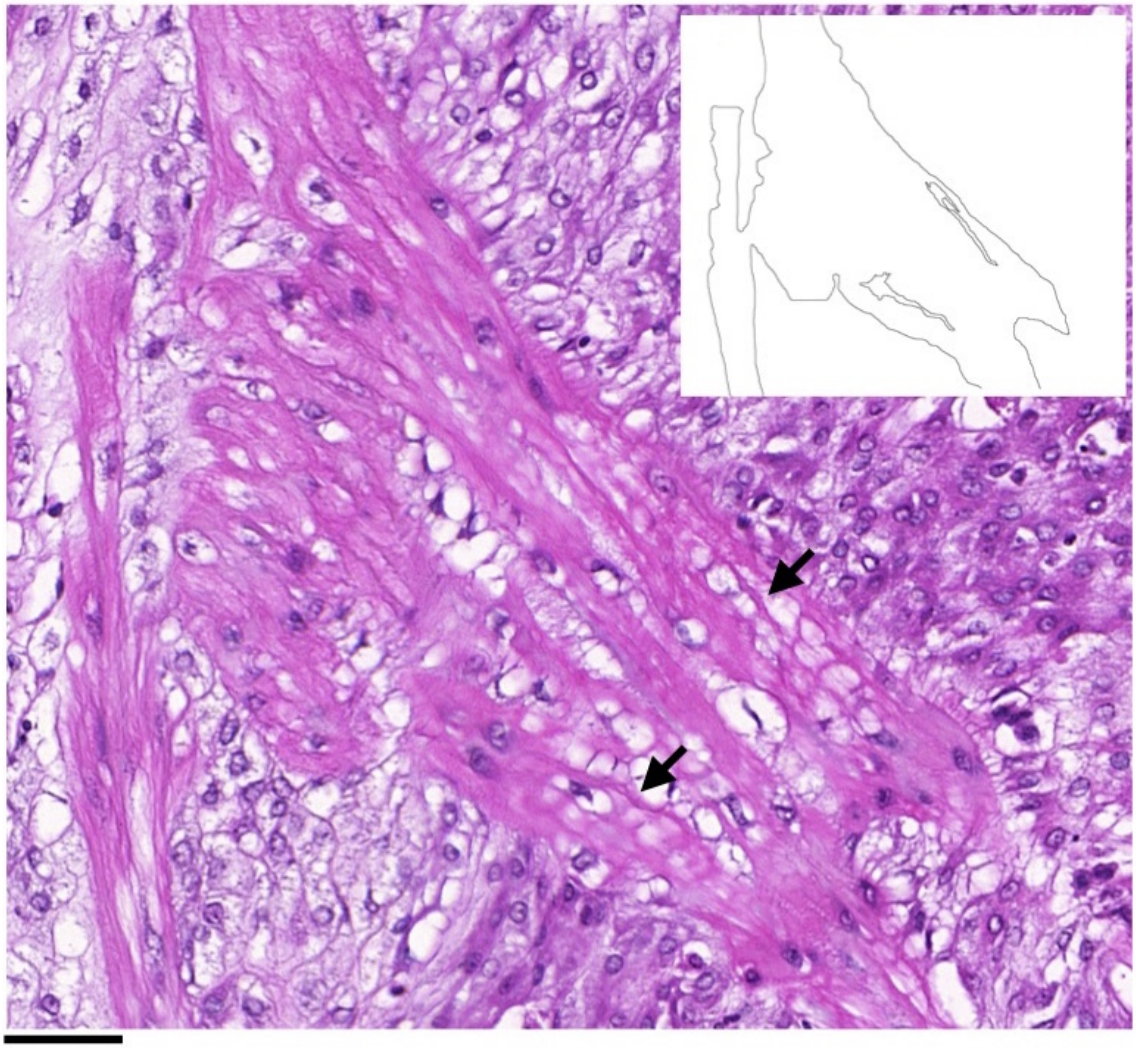

Figure 10.4: An enlarged view of a longitudinal interpenetrating fibre, previously shown in Figure 8.25. The cells appear have a stellate structure, meaning that thin extensions of the cytoplasm (arrows) extend out from the main cell body. Inset shows putative structure of the two indicated cells based on examination of the image shown. Scale bar represents 50 μm.

notion that the fibres were attached to the circular layer prior to implantation. An apparent change in fibre direction associated with the interpenetrating fibres can be observed in these bridges, as shown in Figure 10.6. Some interpenetrating fibres branched from the myometrium outside the placental bed, as demonstrated in Figure 10.7. The absence of invading trophoblast at the branch point in these instances could indicate that the structure was not formed as a consequence of trophoblast invasion. If this branching was not a result of placental invasion, it could be caused by a pre-existing connecting bridge.

There are dissimilarities between the interpenetrating fibres and the bridges connecting the circular and longitudinal layers. The cells in the bridges are spindle-shaped, as opposed to the stellate structure observed in the interpenetrating fibres, as can be seen in Figure 10.6. However, that cells in the connecting bridges might undergo morphological alteration upon trophoblast invasion, causing the cells to transform into the stellate cells observed in the interpenetrating fibres. The apparent clustering of the bridges is not present in the interpenetrating fibres, which appear as isolated fibres in the placenta. Additionally, the width range of the connecting bridges (20–100 μm) is below the width range of the interpenetrating fibres (100–200 μm). One possible explanation for these differences is that most of the clustered connecting fibres may be dispersed in the same manner as the circular layer, while the remaining bridges are merged to produce a thicker fibre of similar dimensions to that of an interpenetrating fibre.

The possible correspondence between the connecting bridges and interpenetrating fibres cannot be determined with certainty on the basis of the data that are currently available. To obtain the required data, future research could apply the methods described in this thesis to generate *in silico* reconstructions of uteri at earlier times in pregnancy and examine the myometrial structure at the implantation site to determine a pathway for the development of these interpenetrating fibres. If the interpenetrating fibres are derived from the connecting bridges, the bridges would be present in the placenta during the process of circular muscle dispersal, which occurs between days 10 and 14 [31].

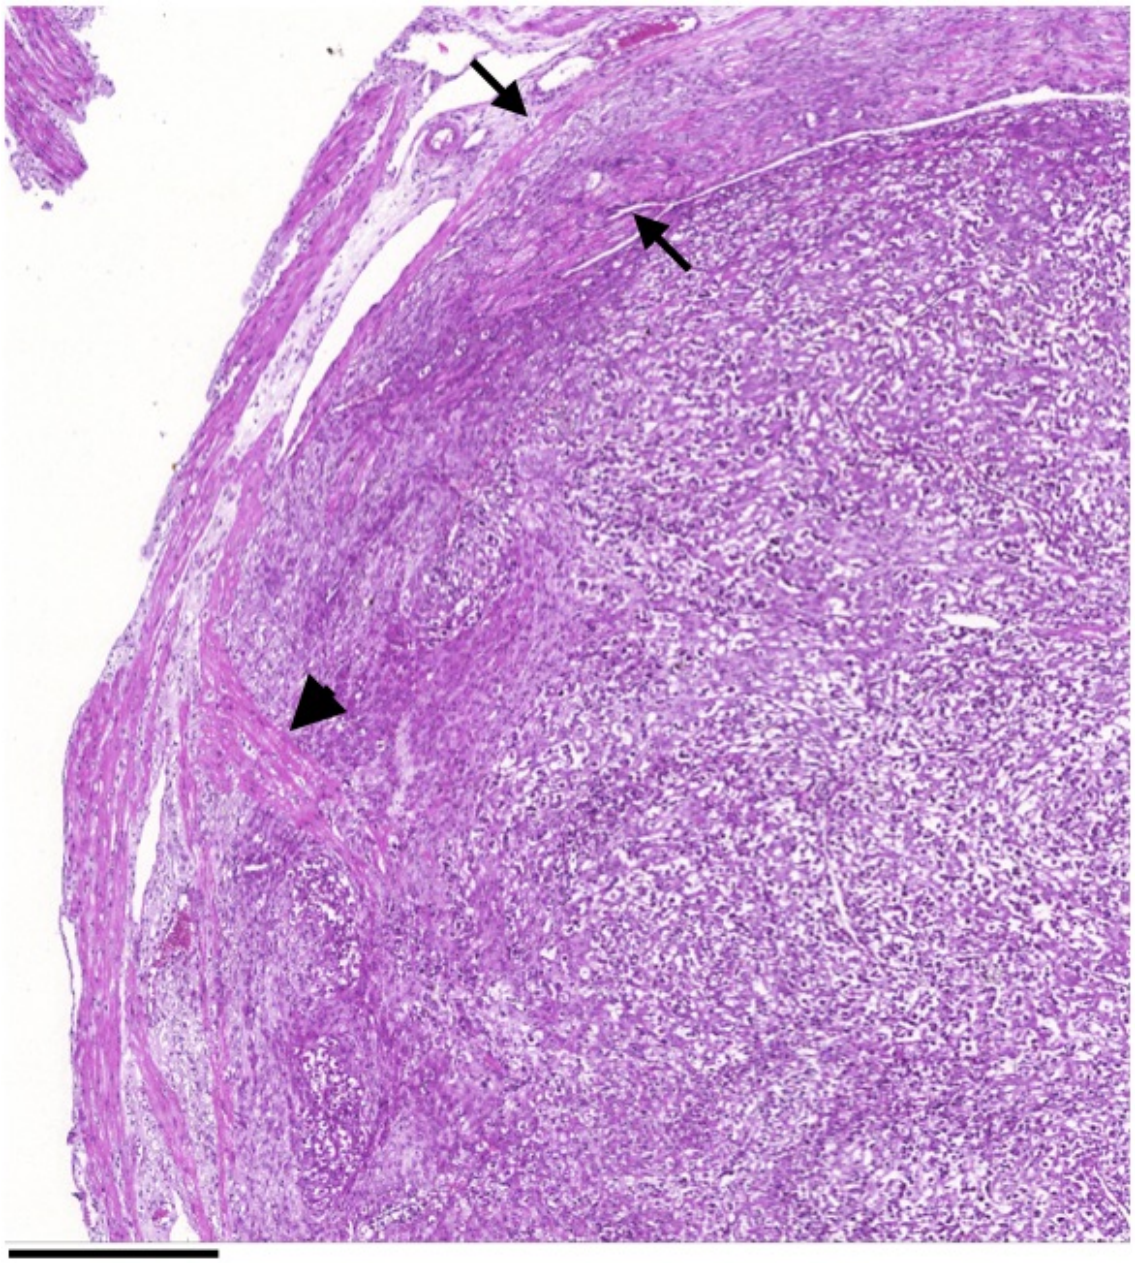

Figure 10.5: An example of a longitudinal interpenetrating fibre (arrowhead) being located in close proximity to decidualized circular fibres (arrows), previously shown in Figure 8.25. This close proximity is consistent with the hypothesis that the interpenetrating fibre was formed by remodelling of a bridge connecting the circular and longitudinal layers. Scale bar represents 500 μm.

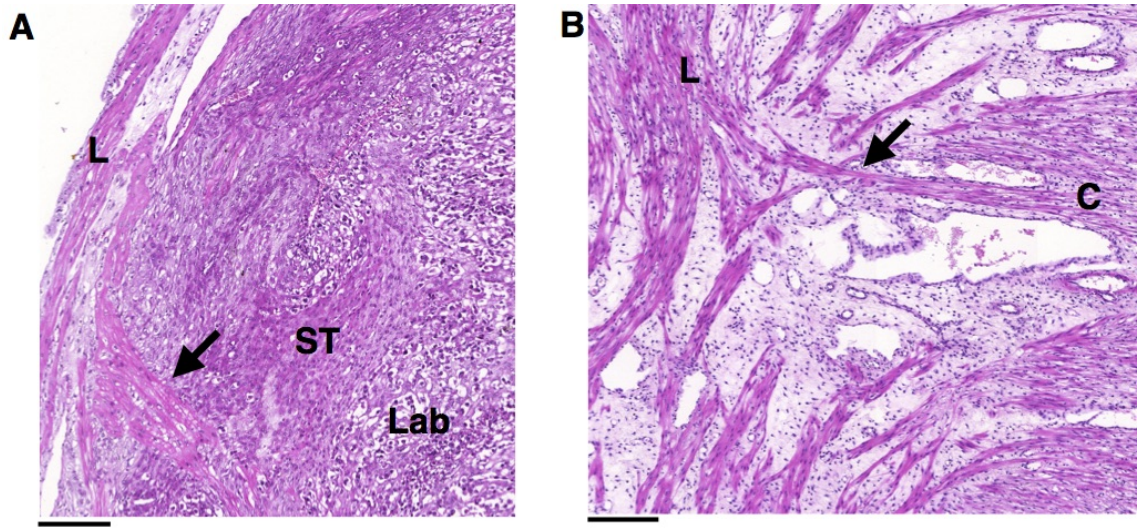

Figure 10.6: Comparison of the structure of an interpenetrating fibre shown in **A** (previously shown in Figure 8.25) with the structure of a cluster of bridges connecting the circular and longitudinal myometrium, shown in **B** (previously shown in Figure 7.23). Both interpenetrating fibres and bridges appear to branch away from the fibre direction, as exhibited by the interpenetrating fibre and bridge indicated by arrows. The interpenetrating fibres tend to have a larger width than the connecting bridges, and the clustering present at bridge points is not observed in the interpenetrating fibres. **L**: longitudinal myometrium. **ST**: spongiotrophoblast, **Lab**: labyrinth, **C**: circular myometrium. Scale bars represent 200 μm.

Spacing of implantation sites is controlled at implantation by contractile forces in the circular myometrium [113]. These contractions cause the implantation sites to form at evenly spaced locations along the length of the uterine horn, as is evident in the rat uteri examined in this thesis. Consequently, the interpenetrating fibres are evenly spaced along the length of the horn. The even distribution of pacemaking zones along the length of the myometrium in the horn could aid localised excitation at specific points along the horn. Such localised contractile behaviour could provide fine control of the contractile forces during parturition.

The analogous zone to the spongiotrophoblast layer in human placenta is the basal plate. The presence of myometrial fibres in the basal plate of the placenta in humans is associated with an increased risk of placenta accreta [62]. This suggests that analogous structures to these interpenetrating fibres, if present in humans, would not penetrate into the basal plate in normal pregnancy. If so, these fibres would only be observable in the

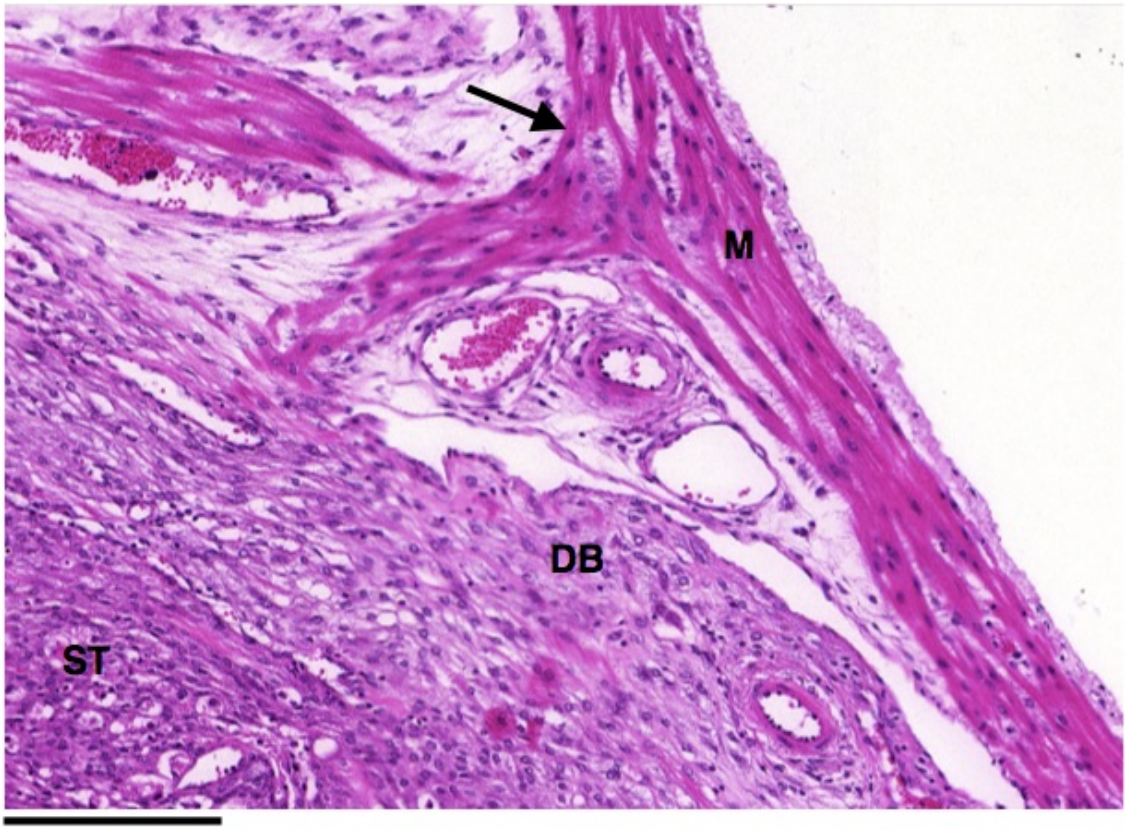

Figure 10.7: An interpenetrating fibre branching from the longitudinal layer outside of the placenta, as indicated by the arrow (previously shown in Figure 8.20). This branching is present outside the leading edge of the invading trophoblast, suggesting that it was a pre-existing structure. **M**: myometrium, **DB**: decidua basalis, **ST**: spongiotrophoblast. Scale bar represents 200 μm.

placental bed, which consists of the decidua basalis and a layer of inner myometrium adjacent to the decidua basalis [19]. Numerous studies have been performed on placental bed biopsies (summarised by Lyall [72]), but none of these appear to have reported any myometrial cells with the stellate structure shown in Figure 10.4. This suggests that structures analogous to the interpenetrating fibres observed here in rat uteri is not present in human uteri.

10.4 Proposed mechanisms controlling contraction in pregnancy

The previous three sections discussed three separate features identified in this thesis which could greatly shape the generation of excitation behaviour in the pregnant rat uterus. A synthesis of these findings suggests a cohesive picture of the overall effect of the presence of these structures.

The structures discussed in the preceding sections are summarised in Figure 10.8. The combination of these features could explain the contraction patterns, as observed by Lammers *et al.* [67] in guinea pig uteri, as described below in Section 10.4.1. Two of these features could also control excitation behaviour in the human myometrium, as discussed below in Section 10.4.2.

10.4.1 Relationships with previous observations in guinea pig uteri

Recent work by Lammers *et al.* [67] showed the following pattern of contractile behaviour in the guinea pig myometrium. Most contractions occur as local contractions close to the mesometrial border. These are often followed by contraction in the circumferential direction, propagating outward in both circumferential directions from the site of the local contraction. In a minority of cases this leads to contraction in other areas of the myometrium, causing full contraction of the organ. The structures shown in Figure 10.8, which have been observed here in rat uteri, could explain this pattern of contraction.

The process of local contraction based on excitation propagation is illustrated in Figure 10.9. Interpenetrating fibres present in the implantation sites were identified in Section 10.3 as putative pacemaker cells. Pacemaking activity in one of these fibres would generate excitation waves propagating out from the implantation site and into the adjacent longitudinal fibres (Figure 10.9**B**). The connective isolation of the mesometrial border fibres, observed in Section 10.1, prevents this excitation wave from propagating across the

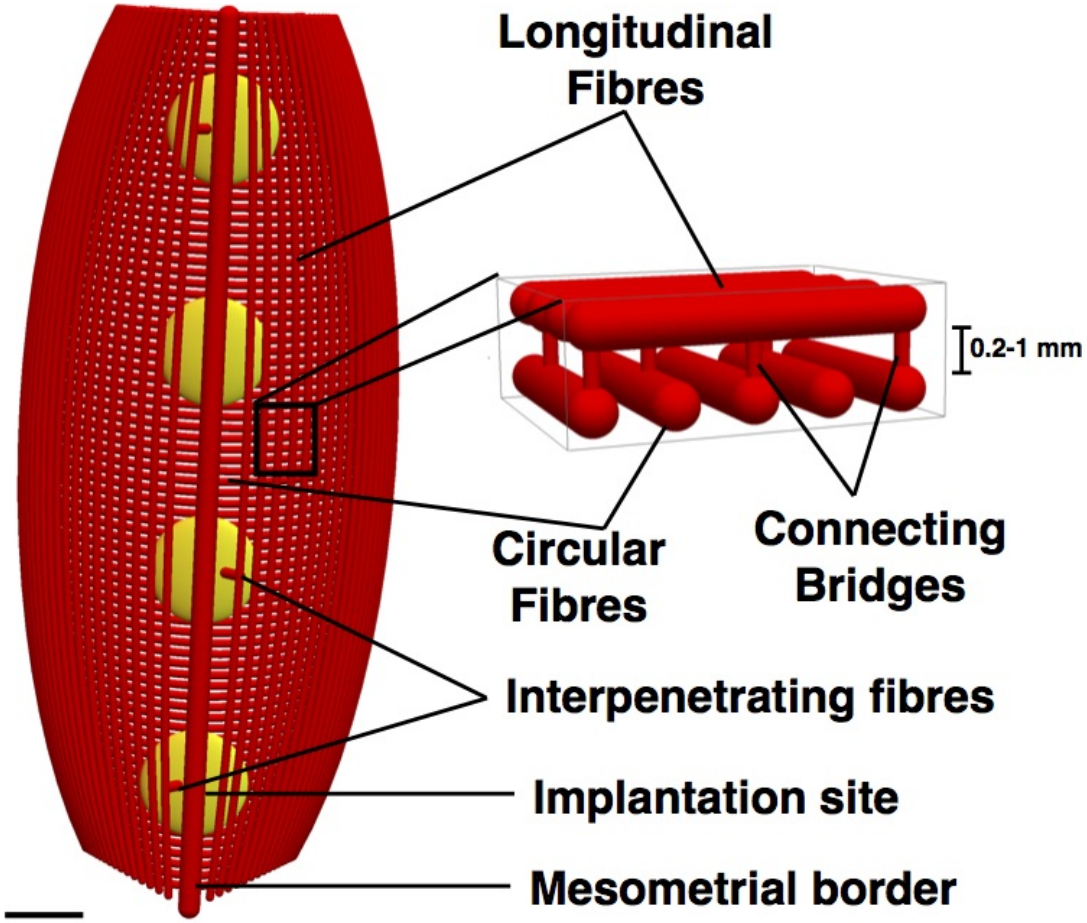

Figure 10.8: Summary of the structural features discussed in Sections 10.1–10.3. The longitudinal fibres in the mesometrial border are connectively isolated from the rest of the longitudinal layer, which could act as a boundary to propagation across the mesometrial border (Section 10.1). Bridges connecting the circular and longitudinal layers could provide a means to coordinate contraction between circular and longitudinal layers (Section 10.2). Interpenetrating fibres entering the placenta from the longitudinal layer are proposed here as a putative site of excitation initiation (Section 10.3). Scale bar represents ~ 5 μm.

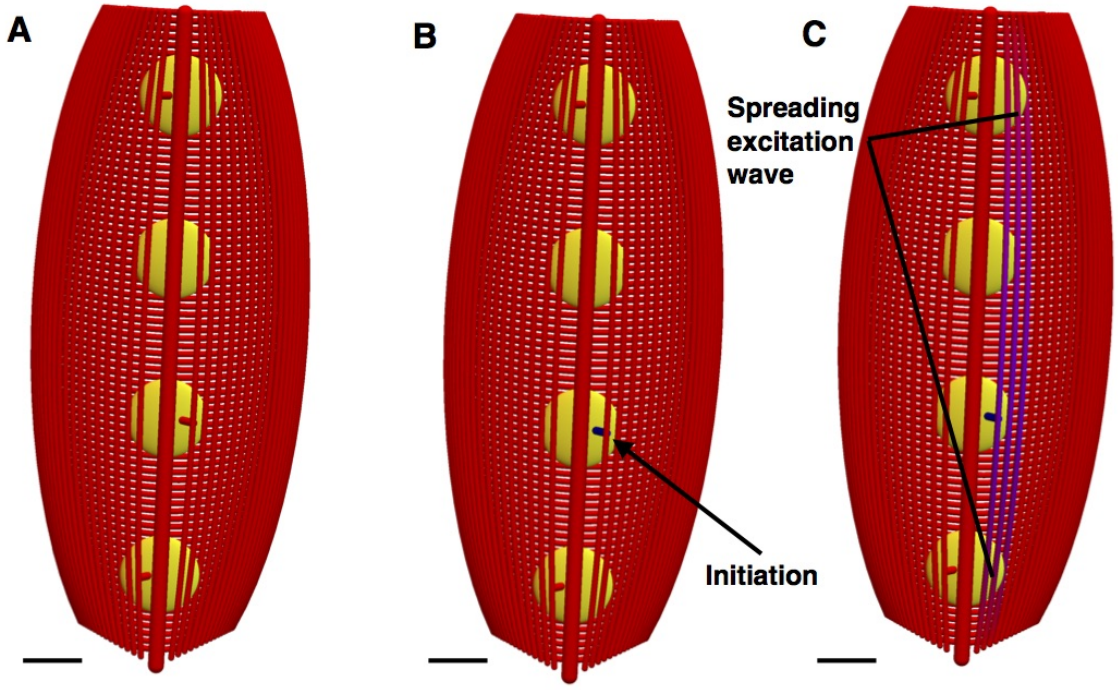

Figure 10.9: The localised excitation waves observed here (Chapter 6) and in [67]. Blue indicates electrically excited fibres, red indicates quiescent fibres. **A**: tissue prior to excitation. **B**: initiation of excitation, originating from an interpenetrating fibre. **C**: excitation spreading into the longitudinal myometrium. The excitation wave fails to propagate across the mesometrial border due to the connective isolation of the fibres along the mesometrial border. The area excited to the right of the mesometrial border is varied between recordings, and is possibly modulated by changes in gap junction permeability in the longitudinal muscle fibres. Scale bars represent ~ 5 μm.

mesometrial border. Thus, all excitation propagating through the longitudinal layer is confined to one side of the mesometrial border. Excitation propagating along the longitudinal fibres to the right of the mesometrial border is shown in Figure 10.9**C**. This form of excitation was observed in the electrode array recordings presented in Chapter 6, and could correspond to the local contractions observed in [67]. The area of tissue excited during these local contractions varied between recordings, which suggests that this area is not predetermined by the myometrial architecture. The cell-to-cell coupling in the myometrium occurs through gap junctions [44]. Temporal variations in the permeability of gap junctions in the tissue could explain this variable area of excitation, in the same manner as discussed in Section 10.2.

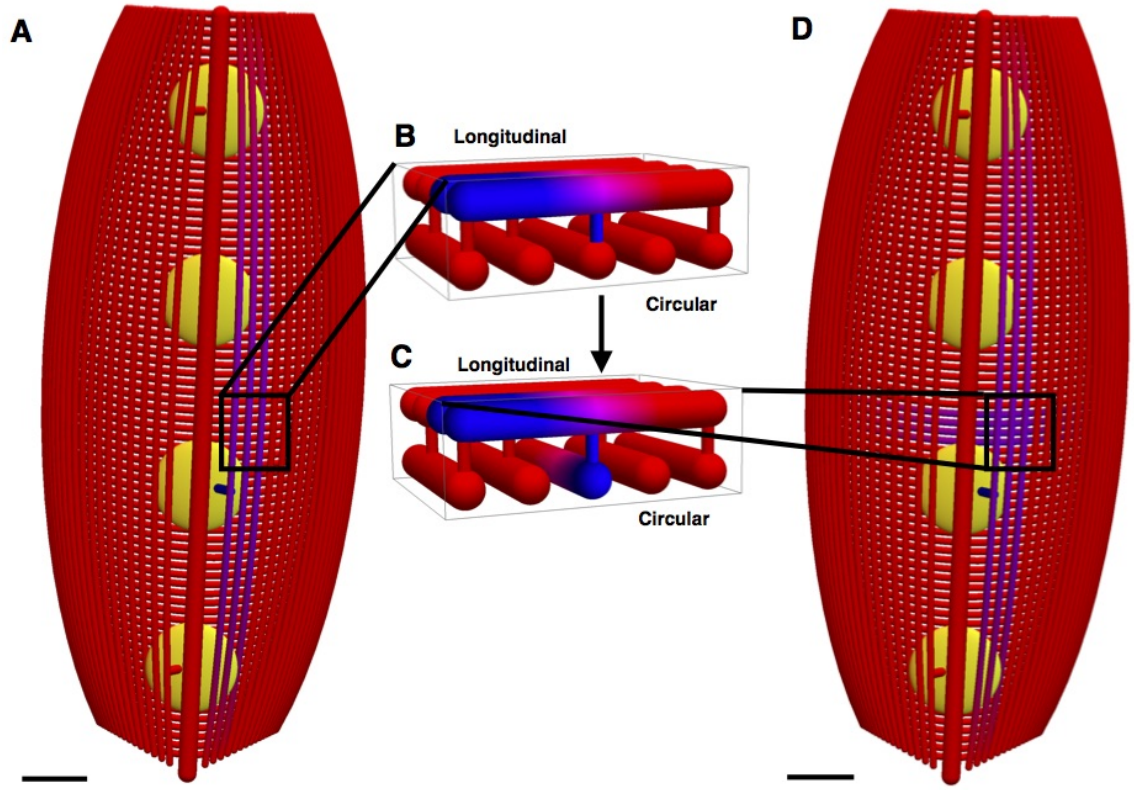

Figure 10.10: Spread of an excitation wave into the circular layer via connecting bridges. Blue represents electrically excited fibres, red indicates quiescent fibres. The local excitation wave shown in **A** spreads from the longitudinal layer into a connecting bridge (**B**). This allows the excitation wave to propagate into the circular layer (**C**), where it spreads out further (**D**). Scale bars represent ~ 5 μm.

The bridges connecting the circular and longitudinal fibres, which are discussed in Section 10.2, could serve as a conducting interface to enable coordination of contractions between the longitudinal and circular layers of the myometrium. Specifically, these bridges could enable excitation from a localised contraction in the longitudinal layer to propagate into the circular myometrium. The excitation wave would subsequently propagate along the circular fibres, leading to stimulation of circumferential contraction, as observed in guinea pig uteri [67]. The process of excitation passing from the longitudinal layer to the circular layer is shown in Figure 10.10. This behaviour could not be observed in the electrode array recordings presented in Chapter 6 because excitation was only recorded in the longitudinal layer. However, propagation of excitation between the longitudinal and

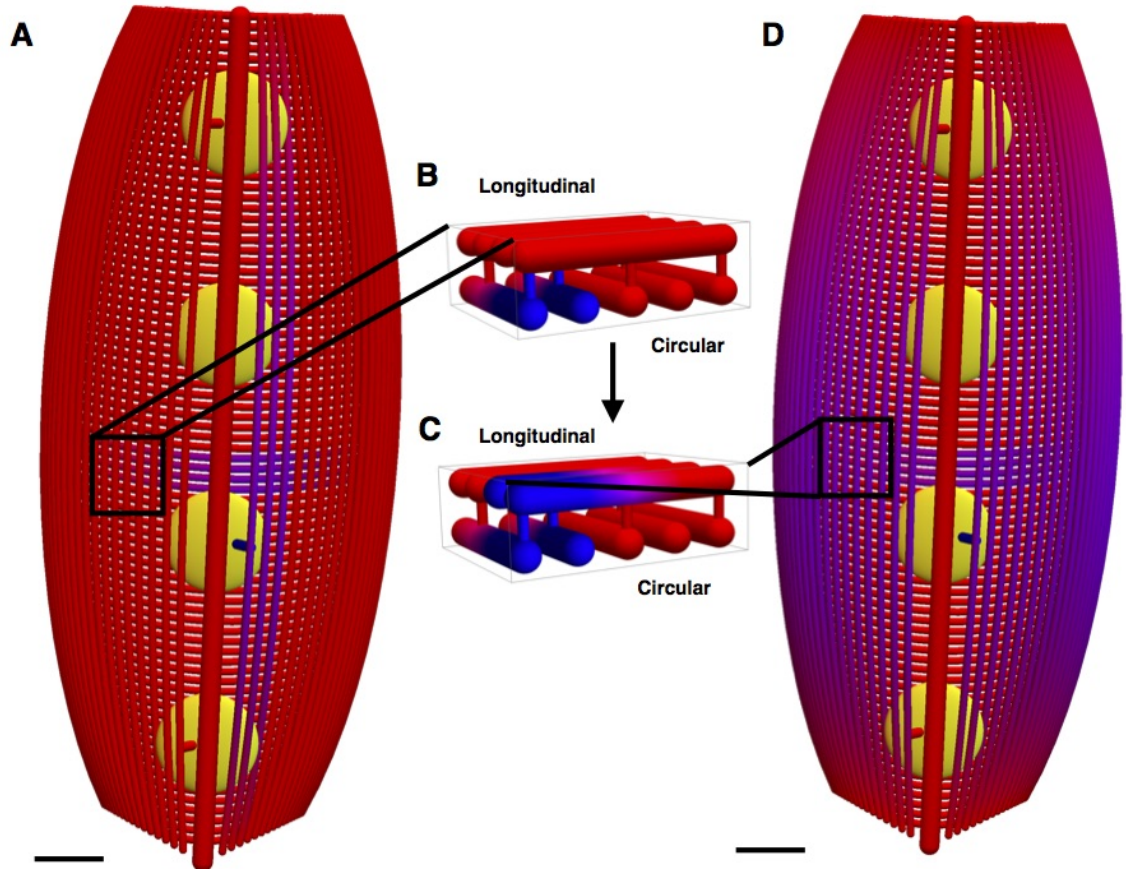

Figure 10.11: Full excitation of the myometrium. Blue represents electrically excited fibres, red indicates quiescent fibres. The excitation wave in the circular layer shown in **A** propagates through the connecting bridges (**B**) back into the longitudinal layer (**C**). This leads to excitation of the remaining parts of the longitudinal myometrium, as shown in **D**. Scale bars represent ~ 5 μm.

circular layers in rat myometrium has previously been observed [24, 87], which suggests that this method of interaction is feasible.

The excitation wave propagating through the circular layer may spread into areas of the longitudinal myometrium that were not recruited by the initial local excitation shown in Figure 10.9. This excitation propagation is facilitated by the connecting bridges in the same manner as propagation into the circular layer, as shown in Figure 10.11. This final step would produce full excitation of the tissue, as observed by Lammers *et al.* [67]. This behaviour was not observed in the electrode array recordings discussed in Chapter 6; however, this event was only observed in a minority of cases ($\sim 30\%$) in guinea pigs [67],

which indicates that this kind of event may be observed given a larger sample of recordings than was analysed here (9).

This sequence of events could explain the contractions described by Lammers *et al.* in guinea pig uteri [67]. However, the structural features shown in Figure 10.8 were only observed in rat uteri. In order to confirm this link, future research might include using the same methods used here to produce an *in silico* reconstruction of the guinea pig uterus. This *in silico* reconstruction could subsequently be examined to determine if the features mentioned here are present in the guinea pig uterus.

10.4.2 Implications for contractile behaviour in human pregnancy

The human myometrium is composed of an inner and an outer layer of smooth muscle, with the inner layer composed of circular fibres and the outer layer composed of a disordered collection of fibres and vasculature [111]. This contrasts with the orderly structure of circular, vascular, and longitudinal layers present in the rat [18]. The structures presented in this thesis relate to the longitudinal layer in the rat myometrium, which is structurally dissimilar to the corresponding outer layer in the human myometrium. However, possible analogues of two of these structures are suggested in Sections 10.1 and 10.2.

The placenta develops along the mesometrial border in normal rat implantation [113], while placentation typically occurs along the upper portion of the midline of the uterus in normal human pregnancy [17]. A boundary of low connectivity in rat uteri was observed along the mesometrial border. An analogous structure in the human uterus would therefore be located along the midline of the uterus. This would appear to contradict previous findings which show excitation propagation across this midline [36]. There are two possible explanations for this discrepancy. First, a boundary of low connectivity along the midline of the uterus may not cover the entire length of the uterus; an incomplete coverage of the midline by a boundary of low connectivity would enable the passage of excitation waves across some parts of the midline. If this were the case, the boundary may only

be present in the upper portion of the outer myometrium, where placental development normally occurs in pregnancy [17]. This structural feature could explain the relatively few instances of propagation observed across the midline in the upper portion of the uterus by Escalona-vargas *et al.* [36]. A second possibility is that excitation across the midline of the myometrium could be facilitated by connecting bridges between the inner and outer myometrium.

Bridges connecting the inner and outer myometrium were observed here in the non-pregnant human uterus, which possibly correspond to the connecting bridges found in rat uteri discussed in Section 10.2. These bridges are here proposed to act as a conductive interface between inner and outer myometrium, which in turn would allow coordinated electrical activity between the layers. In the model proposed in Section 10.4.1, conduction of excitation through the circular layer enabled the excitation of multiple portions of the longitudinal layer from a small initiation event. A similar mechanism could be employed in the human myometrium to coordinate the electrical activity in the outer myometrium. Furthermore, the localised control of the cell-to-cell coupling in these bridges discussed in Section 10.2 could be employed to control global excitation patterns in the same manner as proposed in Section 10.4.1.

The structures proposed here, as observed in the rat uterus (Sections 10.1 and 10.2), could present a means for the control of coordinated contraction in the human myometrium through the propagation of electrical excitation. However, there are multiple factors involved in this process that require further investigation. The boundary of low connectivity across the upper portion of the midline of the myometrium could not be observed in the sample of human myometrium that was inspected in this thesis. Reconstruction of the fundus of the human myometrium, using the techniques developed here, may provide evidence to confirm the existence of such a boundary. The bridges connecting the inner and outer myometrium were observed here in the human myometrium. The proposed mechanisms for fine control of the conductivity of these bridges depend on the gap junctions present

in these bridges. Immunofluorescent staining could present a means for determining gap junction density in these bridges [80]. Furthermore, staining for specific connexin proteins, as has previously been performed [63], would allow identification of any specific connexin proteins with a higher prevalence in these bridges. Differences in connexin composition of gap junctions in these bridges compared to other myometrial tissue could enable the targeted control of conductivity in the bridges, which would facilitate localised control of global excitation behaviour.

10.5 General architecture of the myometrium

Previous work by Weiss *et al.* [111] has shown that the human myometrium can be divided into a well-defined circular layer and outer layer of thick fibres, with an apparently disorderly collection of fibres lying between these two layers. The rat myometrium is composed of an inner circular layer and outer longitudinal layer separated by a vascular layer [78]. These structural layouts were observed for human and rat myometrium in the *in silico* reconstructions presented in this thesis.

Three-dimensional fibre structure of the human myometrium has previously been measured by Young and Hession [116]; fasciculi are large structures with a typical width of 1–2 mm, composed of bundles of width 200–400 μm, which are separated by connective tissue of thickness $\sim$ 20 μm. These fasciculi are separated by layers of connective tissue approximately 75 μm thick, which contains sheet-like bundles. Bridges of myometriocytes connect bundles within a fasciculus, while larger bridges connect adjacent fasciculi. A schematic of this structure is shown in Figure 10.12, with the width range changed to reflect the fasciculus widths measured in the *in silico* reconstruction, as presented in Chapter 7. In the human myometrium, fasciculi of width greater than 1 mm were abundant in the outermost layer, which agrees with observations made by Young and Hession [116]. Fasciculi in the intermediate and inner layers, however, were thinner, having median width between 0.4 and 0.5 mm. This suggests that fasciculi in the intermediate and inner layers

of the myometrium are much smaller than those observed by Young and Hession [116].

In the rat myometrium, fibre width was measured to be 500 ± 50 μm in the circular layer and 300 ± 100 in the longitudinal layer, as demonstrated in Figure 10.13**B**. This contrasts with the human myometrium, where the outermost layer of myometrium contains fasciculi of much greater width than the inner layer, shown in Figure 10.13**A**. Without the inclusion of the fibre widths in the outermost layer, the human myometrium still shows increased width with distance from the inner layer.

The variable widths of the fibres observed here could influence the excitability of the fibres [106]. The ratio of cell-to-cell coupling to volume is lower in thinner fibres than in thicker fibres: a larger proportion of cells in a thinner fibre are on the edge of the fibre due to the increased surface are to volume ratio, and these cells have uncoupled areas on their surface. This means that the total area available for cell-to-cell coupling per unit volume is lower in thinner fibres than thicker fibres. It has previously been shown that excitation propagation through cell networks with high levels of cell-to-cell coupling requires a large input of current [99]. This suggests that a lower input current would be required to excite thinner fibres than thicker fibres. Pacemaking activity in the rat uterus was observed in the longitudinal layer here (Chapter 6) and by Chen and Chiang [24]. Thinner fibres in the longitudinal layer in the rat myometrium could facilitate the propagation of excitation waves into the longitudinal layer from small initiation areas. Pacemaking regions in the human myometrium have not previously been identified [44]. The pattern of decreasing fibre width toward the inner myometrium could indicate that excitation is more likely to be initiated in the inner myometrium.

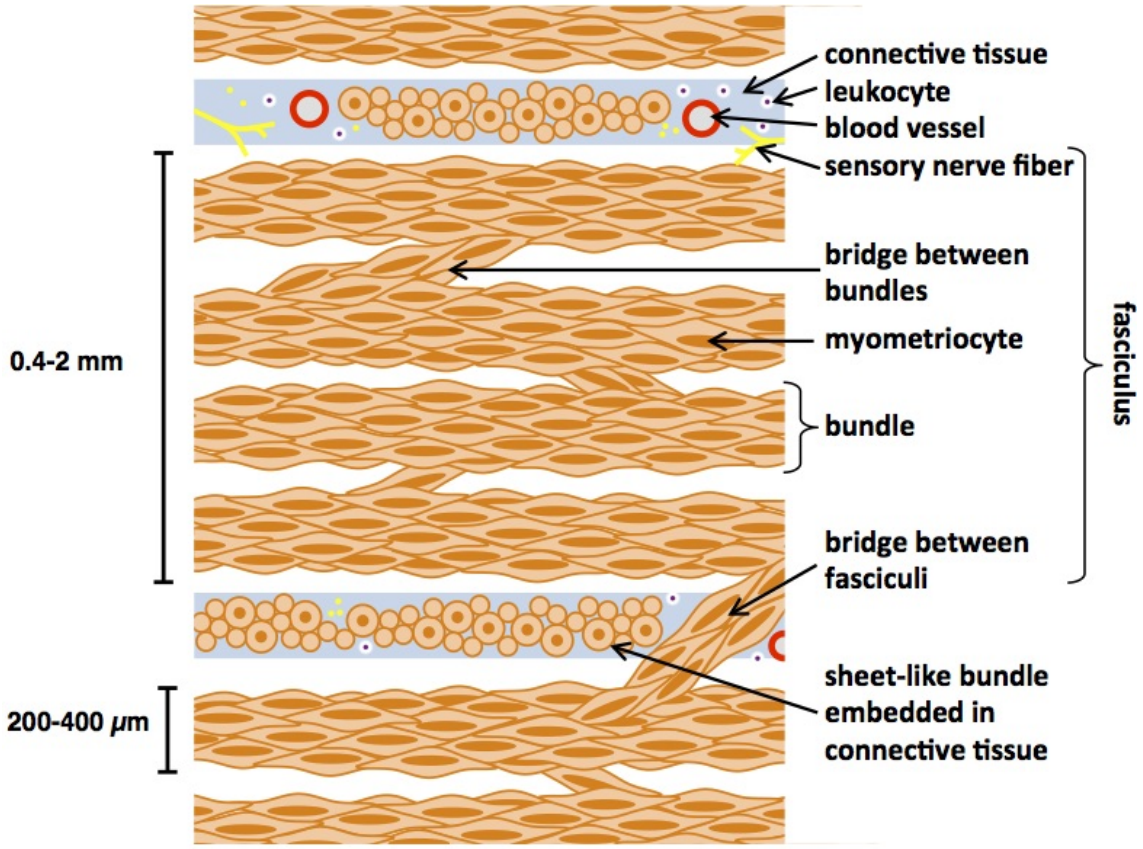

Figure 10.12: Schematic of the structure of a fasciculus in the myometrium as proposed by Young and Hession [116], with the width range modified based on observations made here. In close proximity to the perimetrium, the fasciculus is a 1–2 mm thick collection of bundles, each of width 200–400 μm. In deeper portions of the human uterus reconstructed in this thesis, fasciculi have median width approximately 0.4 μm, in contrast to the minimum width of 1 mm proposed by Young and Hession [116]. The bundles within the fasciculus are connected by bridges of myometriocytes. Fasciculi are separated by connective tissue, which contains sheet-like bundles, Bridges of myometriocytes between fasciculi are present, and tend to be thicker than those between individual bundles.

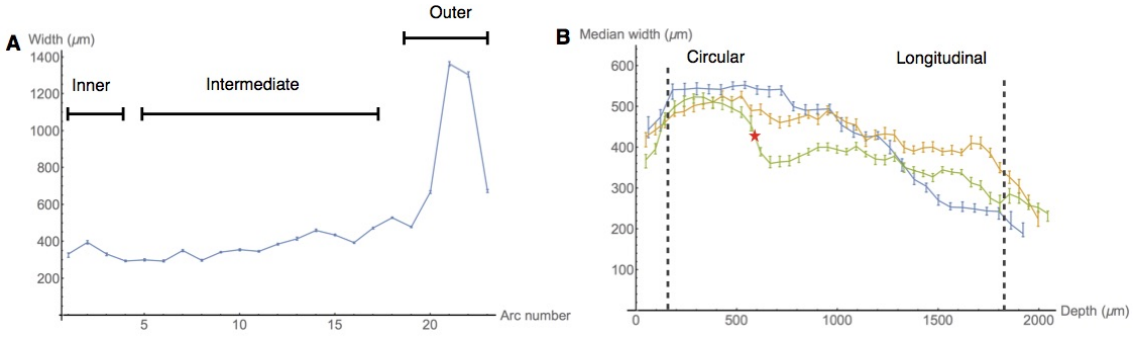

Figure 10.13: Variation of fibre widths between inner and outer myometrium in the human uterus (**A**), and between circular and longitudinal myometrium in rat uteri (**B**), as previously presented in Chapter 7. Numbered arcs were used for measurement of the myometrium in human tissue (see Chapter 7), with arc number representing the distance from the interior of the uterus. The fibre width increases with distance from the inner myometrium, with a large increase in fibre width corresponding to the outer layer of sub-serosal longitudinal fibres. The reverse trend is present in the rat uteri: wider fibres are present in the (inner) circular myometrium, while fibres in the (outer) longitudinal myometrium are thinner. Dashed lines indicate the points where fibre width is reduced by close proximity to the start and end of sectioning, while the red star indicates the location in the third rat block (green) where fibre widths are reduced by complications in the sectioning process (see Chapter 7) for details.

10.6 Limitations

The separation of bundles within the fasciculi is smaller than the voxel size of the *in silico* reconstruction, which means that this separation cannot be directly observed in the reconstruction. The segmentation process described in Chapter 3 created boundaries based on changes in fibre direction. The bundles within the fasciculi, however, are generally oriented in the same direction, and therefore the segmentation algorithm could not separate these bundles. For this reason, finer structures within the fasciculi could not be accurately measured in the *in silico* reconstructions. A means for generating higher resolution reconstructions is discussed in Section 10.7.1, which could be used to obtain a reconstruction which accurately represents these finer structures.

The rat tissue samples were cut along the antimesometrial border and pinned prior to reconstruction. These steps distorted the tissue, which in turn means that the *in silico* reconstructions do not represent the original structure of the uteri. The stretching could

have impacted the width measurements presented in Chapter 7: longitudinal fibres were stretched transverse to their direction by the pinning, while circular fibres were stretched along the length, both of which could affect the width of the structures measured. Additionally, areas where the longitudinal fibres were distorted to align with the circular layer may contain erroneous connections between the layers in the reconstruction that would not be identified by the segmentation algorithm, which only identifies boundaries where there is a substantial change in fibre direction. The failure to accurately segment these parts of the tissue could affect simulated excitation in the tissue, since it would enable propagation between parts of the tissue that are not connected. Similarly, connectivity in these areas may not be accurately represented in the reconstruction. However, the reconstructions were in direct correspondence with the histological slides, which enabled the verification of the features identified here. While the histological slides were distorted by the pinning, connecting bridges and longitudinal interpenetrating fibres would not form as a result of this stretching. The argument could be made that the stretching could have severed the connections along the mesometrial border, leading to the observed boundary. This scenario is unlikely, however, since such force on a uniformly connected network of fibres would create tears in other parts of tissue, which were not observed. One possible future application of the reconstruction methods presented here is the reconstruction of rat uteri that more closely represent the original configuration. This would produce more accurate representations of the connectivity in these tissues, and could provide insight into the distribution of connections between the circular and longitudinal layers.

10.7 Software

This thesis has required the development of dedicated software to analyse the biological data. Specifically, tools have been developed for *in silico* reconstruction, simulation, and electrode array analysis.

10.7.1 Reconstruction

The reconstruction of myometrium biopsies has previous been studied by Young *et al.* [116]. This reconstruction used the connective tissue present in the myometrium to segment fibrous structures. Imaging of the uterus using MRI has previously been performed [111]. This method enabled the visualisation of fibre structures on an organ level.

The methods presented in this thesis provide a means to reconstruct the whole myometrium, as was performed on the rat uteri, at a higher resolution (~ 50 μm) than the fibrous structures identified by Young and Hession [116] (0.1–2 mm). This high resolution enabled the identification of low connectivity fibres across the mesometrial border and of connective bridges between circular and longitudinal layers in rat myometrium. Furthermore, the reconstructed high-resolution tissue microarchitecture provided a structure for simulated excitation.

The resolution of the *in silico* reconstruction yielded a voxel length similar to the typical length of a myometrial smooth muscle cell (50 μm), but significantly higher than the typical width (10 μm) [15]. This led to uncertainty in the precise location of the boundaries of fibres: if two fibres of the same orientation were situated less than 50 μm apart, the reconstruction would fail to differentiate between the two. In Chapter 8, a higher resolution reconstruction (~ 10 μm per voxel) was used to visualise myometrial fibres present in the placenta. Preliminary results of this technique from the use on implantation sites in Chapter 8 show that this technique can be used to represent fibre structure, but will require some adaptation. The reconstruction in Chapter 8 required manual input to differentiate smooth muscle tissue from placental tissue. Therefore, one adaptation required of this technique would be to develop means for automated detection of different tissue types. This could include techniques developed for detecting separate tissue types here (see Chapter 3). Alternatively, additional information could be acquired from histological images, such as local colour variations, which could aid in the identification of separate tissue types in a similar manner to visual verification. Another requirement would

be determining the three-dimensional direction of the fibres. Again, methods developed for the current reconstruction could be adapted to fit this higher-resolution reconstruction.

10.7.2 Electrode array analysis

The methods for multi-electrode array analysis presented in this thesis provide an automated means of tracking the propagation of excitation bursts recorded by a multi-electrode array. This enabled the detection of initiation and termination points of excitation waves, and handled multiple initiation events where present.

While the techniques used in this thesis enable accurate tracking of the burst propagation in the tissue, more subtle effects which may be observable through tracking individual action potentials could not be detected in this manner, since the algorithm is based on exploiting the frequency of action potentials in an individual burst to track its propagation.

10.7.3 Simulated excitation

General-purpose computing on graphics processing units (GPGPU) for simulating cardiac electrophysiology is a well-established technique [109, 117]. This technique utilises the highly parallelised nature of the computing architecture in the GPU to perform large quantities of calculations in parallel. The methods described in this thesis utilise the directional data present in the *in silico* reconstruction to simulate excitation propagation with a high level of accuracy. The high resolution and large volume of the reconstruction yielded a large number of equations to be solved ($\sim 10^6$), which motivated the use of GPGPU. The resulting simulations were able to represent excitation propagation through the tissue at a resolution of ~ 50 μm.

Simulation of excitation in the whole myometrium has previously been performed using a variety of models and techniques [9, 64, 73, 92, 114]. In particular, the mathematical model and numerical methods used to simulate excitation in [64] are similar to those used here. The simulation uses a direction field that can be varied spatially to represent the

actual fibre structure. This model uses a coarser spatial discretisation of the myometrium than used in this thesis; however, as discussed below, the large volume of the human myometrium necessitates the use of a coarser model.

The simulations performed in this thesis used the FitzHugh-Nagumo model of an action potential, which is represented by two partial differential equations [40]. In reality, the dynamics of excitation in the myometrium are far more complex, requiring ~ 100 partial differential equations [10]. The increase in equations and variables by a factor of ~ 50 that would be introduced by using this more complete model would increase the memory requirements of the GPU to above what is practically available at present. Additionally, any increase in resolution, such as described in Section 10.7.1, or increase in size of the tissue, which would be needed for modelling the whole human uterus, would further increase the memory requirements. A conservative estimate for the memory requirements for this more detailed model of excitation on a whole human uterus at the resolution used here would be roughly $100\times$ the current capacity of the most advanced GPU architecture. One possible method of coarse-graining the reconstruction would be to use the network representation described in Chapter 3. This representation contains approximately $1000\times$ fewer nodes than the number of voxels used to represent the tissue, and would therefore be a viable structure on which to base these simulations. The network representation has direct correspondence with the original reconstruction, which means that the network could be used to project properties of localised simulated electrophysiology onto the whole organ model. This would act as a bridge between cell-level excitation and organ-level excitation, enabling precise models of excitation to be mapped onto larger structures.

10.8 Conclusion

The *in silico* reconstruction presented in this thesis enabled the location of possible pacemaker cells, regions of low connectivity along the mesometrial border, and connective bridges between longitudinal and circular layers in rat myometrium. Additionally, the

reconstructions enabled general comparison between rat and human myometrium. The reconstruction was used as a basis for simulated excitation.

Bibliography

[1] 3DHISTECH. Pannoramic Viewer. [Online; accessed 3-February-2016].

[2] M. D. Abràmoff, P. J. Magalhães, and S. J. Ram. Image processing with ImageJ. *Biophotonics International*, 11(7):36–42, 2004.

[3] P. Acién. Reproductive performance of women with uterine malformations. *Human Reproduction*, 8(1):122–126, 1993.

[4] H. N. Aguilar and B. F. Mitchell. Physiological pathways and molecular mechanisms regulating uterine contractility. *Human Reproduction Update*, 16(6):725–744, 2010.

[5] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell*, chapter Chapter 11. Membrane Transport of Small Molecules and the Electrical Properties of Membranes, pages 651–694. Garland Science, 5$^{\text{th}}$ edition, 2008.

[6] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell*, chapter Chapter 15. Mechanisms of Cell Communication, pages 879–964. Garland Science, 5$^{\text{th}}$ edition, 2008.

[7] J. S. Alstrøm, D. B. Hansen, M. S. Nielsen, and N. MacAulay. Isoform-specific phosphorylation-dependent regulation of connexin hemichannels. *Journal of Neurophysiology*, 114(5):3014–3022, 2015.

[8] G. F. Anderson, T. Kawarabayashi, and J. M. Marshall. Effect of indomethacin and aspirin on uterine activity in pregnant rats: Comparison of circular and longitudinal muscle. *Biology of Reproduction*, 24:359–372, 1981.

[9] O. Aslanidi, J. Atia, A. P. Benson, H. A. van den Berg, A. M. Blanks, C. Choi, S. H. Gilbert, I. Goryanin, B. R. Hayes-Gill, A. V. Holden, P. Li, J. E. Norman, A. Shmygol, N. A. B. Simpson, M. J. Taggart, W. C. Tong, and H. Zhang. Towards a computational reconstruction of the electrodynamics of premature and full term human labour. *Progress in Biophysics and Molecular Biology*, 107(1):183–192, 2011.

[10] J. Atia. *Development of a method to characterise the expression profile of electrogenic transmembrane proteins in excitable cells.* PhD thesis, University of Warwick, 2014.

[11] Y. S. Babu, J. S. Sack, T. J. Greenhough, C. E. Bugg, A. R. Means, and W. J. Cook. Three-dimensional structure of calmodulin. *Nature*, 315(6014):37–40, 1985.

[12] D. H. Ballard. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2):111–122, 1981.

[13] E. Bartocci, E. M. Cherry, J. Glimm, R. Grosu, S. A. Smolka, and F. H. Fenton. Toward real-time simulation of cardiac dynamics. In *Proceedings of the 9th International Conference on Computational Methods in Systems Biology*, pages 103–112. ACM, 2011.

[14] M. J. Berridge. Inositol trisphosphate and calcium signalling. *Nature*, 361(6410):315–325, 1993.

[15] S. T. Blackburn. *Maternal, Fetal, & Neonatal Physiology: A clinical perspective.* Saunders, Elsevier, Missouri, USA, 3rd edition, 2007.

[16] A. M. Blanks, A. Shmygol, and S. Thornton. Myometrial function in prematurity. *Best Practice & Research Clinical Obstetrics & Gynaecology*, 21(5):807–819, 2007.

[17] R. Booth, C. Wood, R. Beard, J. Gibson, and J. H. Pinkerton. Significance of site of placental attachment in uterus. *British Medical Journal*, 1(5294):1732–1734, 1962.

[18] J. R. Brody and G. R. Cunha. Histologic, morphometric, and immunocytochemical analysis of myometrial development in rats and mice: I. normal development. *American Journal of Anatomy*, 186(1):1–20, 1989.

[19] J. J. Brosens, R. Pijnenborg, and I. A. Brosens. The myometrial junctional zone spiral arteries in normal and abnormal pregnancies: a review of the literature. *American Journal of Obstetrics and Gynecology*, 187(5):1416–1423, 2002.

[20] R. L. Burden and J. D. Faires. *Numerical Analysis*. Brooks/Cole, Cengage Learning, Belmont, CA, USA, 2005.

[21] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Inteligence*, 8(6):679–698, 1986.

[22] Y.-W. Chan, H. A. van den Berg, J. D. Moore, S. Quenby, and A. M. Blanks. Assessment of myometrial transcriptome changes associated with spontaneous human labour by high throughput RNA-seq. *Experimental Physiology*, 99(3):510–524, 2014.

[23] T. Chard and J. G. Grudzinskas. *The Uterus*. Cambridge University Press, Cambridge, UK, 1995.

[24] C. J. Chen and S. T. Chiang. The coordination of contractions of circular and longitudinal muscle layers of rat uterus in late pregnancy and during delivery. *The Chinese Journal of Physiology*, 30(1):35–47, 1986.

[25] L. P. Clarke, R. P. Velthuizen, M. A. Camacho, J. J. Heine, M. Vaidyanathan, L. O. Hall, R. W. Thatcher, and M. L. Silbiger. MRI segmentation: methods and applications. *Magnetic Resonance Imaging*, 13(3):343–368, 1995.

[26] R. H. Clayton, O. Bernus, E. M. Cherry, H. Dierckx, F. H. Fenton, L. Mirabella, A. V. Panfilov, F. B. Sachse, G. Seemann, and H. Zhang. Models of cardiac tissue electrophysiology: Progress, challenges and open questions. *Progress in Biophysics and Molecular Biology*, 104(1):22–48, 2011.

[27] K. Coward and D. Wells, editors. *Textbook of Clinical Embryology.* Cambridge University Press, Cambridge, UK, 2013.

[28] G. R. Cunha. The dual origin of vaginal epithelium. *American Journal of Anatomy*, 143(3):387–392, 1975.

[29] E. E. Daniel. Effect of the placenta on the electrical activity of the cat uterus in vivo and in vitro. *American Journal of Obstetrics and Gynecology*, 80(2):229–244, 1960.

[30] E. De Castro and C. Morandi. Registration of translated and rotated images using finite fourier transforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):700–703, 1987.

[31] E. P. C. T. De Rijk, E. Van Esch, and G. Flik. Pregnancy dating in the rat: placental morphology and maternal blood parameters. *Toxicologic Pathology*, 30(2):271–282, 2002.

[32] C. Deroulers, D. Ameisen, M. Badoual, C. Gerin, A. Granier, and M. Lartaud. Analyzing huge pathology images with open source software. *Diagnostic Pathology*, 8(1):92, 2013.

[33] R. O. Duda and P. E. Hart. Use of the Hough transformation to detect lines and curves in pictures. *Graphics and Image Processing*, 15(1):11–15, 1972.

[34] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis.* Wiley, New York, USA, 1973.

[35] R. A. Duquette, A. Shmygol, C. Vaillant, A. Mobasheri, M. Pope, T. Burdyga, and S. Wray. Vimentin-positive, c-kit-negative interstitial cells in human and rat uterus: a role in pacemaking? *Biology of Reproduction*, 72(2):276–283, 2005.

[36] D. Escalona-Vargas, R. B. Govindan, A. Furdea, P. Murphy, C. L. Lowery, and H. Eswaran. Characterizing the propagation of uterine electrophysiological signals recorded with a multi-sensor abdominal array in term pregnancies. *PloS One*, 10(10):e0140894, 2015.

[37] P. Fatt and B. L. Ginsborg. The ionic requirements for the production of action potentials in crustacean muscle fibres. *Journal of Physiology*, 142:516–543, 1958.

[38] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, et al. 3D slicer as an image computing platform for the quantitative imaging network. *Magnetic Resonance Imaging*, 30(9):1323–1341, 2012.

[39] G. A. Feldhamer, L. C. Drickamer, S. H. Vessey, J. F. Merritt, and C. Krajewski. *Mammalogy: Adaptation, Diversity, and Ecology.* The Johns Hopkins University Press, Maryland, USA, 2007.

[40] R. FitzHugh. Thresholds and plateaus in the Hodgkin-Huxley nerve equations. *The Journal of General Physiology*, 43:867–896, 1960.

[41] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 9:891–906, 1991.

[42] T. Y. Fung, D. S. Sahota, T. K. Lau, T. Y. Leung, L. W. Chan, and T. K. H. Chung. Placental site in the second trimester of pregnancy and its association with subsequent obstetric outcome. *Prenatal Diagnosis*, 31(6):548–554, 2011.

[43] R. E. Garfield. Gap junction formation in myometrium: control by estrogens, progesterone, and prostaglandins. *American Journal of Physiology*, 238(3):C81–C89, 1980.

[44] R. E. Garfield, M. G. Blennerhassett, and S. M. Miller. Control of myometrial contractility: role and regulation of gap junctions. *Oxford Reviews of Reproductive Biology*, 10:436–490, 1987.

[45] R. E. Garfield, G. Saade, C. Buhimschi, I. Buhimschi, L. Shi, S.-Q. Shi, and K. Chwalisz. Control and assessment of the uterus and cervix during pregnancy and labour. *Human Reproduction Update*, 4(5):673–695, 1998.

[46] R. E. Garfield, S. Sims, and E. E. Daniel. Gap junctions: their presence and necessity in myometrium during parturition. *Science*, 198(4320):958–960, 1977.

[47] L. P. Gartner and J. L. Hiatt. *Color Atlas of Histology.* Lippincott Williams & Wilkins, Maryland, USA, 5$^{\text{th}}$ edition, 2009.

[48] J. D. Gibbons and S. Chakraborti. *Nonparametric Statistical Inference.* Springer, 2011.

[49] G. Gomez. Local smoothness in terms of variance: The adaptive gaussian filter. In *BMVC*, pages 1–10, 2000.

[50] A. Goode, B. Gilbert, J. Harkes, D. Jukic, and M. Satyanarayanan. OpenSlide: A vendor-neutral software foundation for digital pathology. *Journal of Pathology Informatics*, 4(1):27, 2013.

[51] M. Goto, H. Kuriyamah, and Y. Abe. Refractory period and conduction of excitation in the uterine muscle cells of the mouse. *Japanese Journal of Physiology*, 11:369–377, 1961.

[52] R. B. Govindan, E. Siegel, S. Mckelvey, P. Murphy, C. L. Lowery, and H. Eswaran. Tracking the changes in synchrony of the electrophysiological activity as the uterus approaches labor using magnetomyographic technique. *Reproductive Sciences*, 22(5):595–601, 2014.

[53] F. T. Hammad, B. Stephen, L. Lubbad, J. F. B. Morrison, and W. J. E. P. Lammers. Macroscopic electrical propagation in the guinea pig urinary bladder. *American Journal of Physiology – Renal Physiology*, 307(2):F172–F182, 2014.

[54] B. Hille. *Ionic Channels of Excitable Membranes*. Sinauer Associates, Inc., 2nd edition, 1992.

[55] A. L. Hodgkin and A. F. Huxley. Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *Journal of Physiology*, 116(4):449–472, 1952a.

[56] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 177(4):500–544, 1952d.

[57] P. V. C. Hough. Method and means for recognizing complex patterns, December 18 1962. US Patent 3,069,654.

[58] C. V. Howard and M. G. Reed. *Unbiased Stereology: three-dimensional measurement in microscopy*. Springer-Verlag, New York, USA, 1998.

[59] H. G. Kaganami and Z. Beiji. Region-based segmentation versus edge detection. In *Intelligent Information Hiding and Multimedia Signal Processing, 2009. IIH-MSP '09. Fifth International Conference on*, pages 1217–1221. IEEE, 2009.

[60] S. Kanda and H. Kuriyama. Specific features of smooth muscle cells recorded from the placental region of the myometrium of pregnant rats. *Journal of Physiology*, 299:127–144, 1980.

[61] J. Keener and J. Sneyd. *Mathematical Physiology*. Springer-Verlag, New York, USA, 2004.

[62] T. Y. Khong and A. C. Werger. Myometrial fibers in the placental basal plate can confirm but do not necessarily indicate clinical placenta accreta. *American Journal of Clinical Pathology*, 116(5):703–708, 2001.

[63] W. M. Kilarski, S. Rothery, G. M. Roomans, U. Ulmsten, M. Rezapour, S. Stevenson, S. R. Coppen, E. Dupont, and N. J. Severs. Multiple connexins localized to individual gap-junctional plaques in human myometrial smooth muscle. *Microscopy Research and Technique*, 54(2):114–122, 2001.

[64] P. S. La Rosa, H. Eswaran, H. Preissl, and A. Nehoraj. Multiscale forward electromagnetic model of uterine contractions during pregnancy. *BMC Medical Physics*, 12(4):1–16, 2012.

[65] W. J. E. P. Lammers, K. Arafat, A. el-Kays, and T. Y. el-Sharkawy. Spatial and temporal variations in local spike propagation in the myometrium of the 17-day pregnant rat. *Pflgers Archiv - European Journal of Physiology*, 267(5):C1210–C1223, 1994.

[66] W. J. E. P. Lammers, H. Mirghani, B. Stephen, S. Dhanasekaran, A. Wahab, M. A. H. Al-Sultan, and F. Abazer. Patterns of electrical propagation in the intact pregnant guinea pig uterus. *American Journal of Physiology – Regulatory, Integrative and Comparative Physiology*, 294:R919–R928, 2008.

[67] W. J. E. P. Lammers, B. Stephen, M. A. H. Al-Sultan, S. B. Subramanya, and A. M. Blanks. The location of pacemakers in the uteri of pregnant guinea pigs and rats. *American Journal of Physiology – Regulatory, Integrative and Comparative Physiology*, 309(11):R1439–R1446, 2015.

[68] W. J. E. P. Lammers, B. Stephen, R. Hamid, and D. W. G. Harron. The effects of oxytocin on the pattern of electrical propagation in the isolated pregnant uterus of the rat. *Pflgers Archiv - European Journal of Physiology*, 437(3):363–370, 1999.

[69] M. G. Larson and F. Bengzon. *The finite element method: theory, implementation, and applications*, volume 10. Springer Science & Business Media, Berlin, Germany, 2013.

[70] M. Lei, H. Ghezzo, M. F. Chen, and D. H. Eidelman. Airway smooth muscle orientation in intraparenchymal airways. *Journal of Applied Physiology*, 82(1):70–77, 1997.

[71] T. Luo, H. Chen, and G. S. Kassab. 3D reconstruction of coronary artery vascular smooth muscle cells. *PloS One*, 11(2):e0147272, 2016.

[72] F. Lyall. The human placental bed revisited. *Placenta*, 23(8):555–562, 2002.

[73] C. Marque, J. Laforêt, C. Rabotti, A. Alexandersson, G. Germain, J. Gondry, B. Karlsson, B. Leskosek, M. Mischi, C. Muszinski, G. Oei, J. Peuscher, and D. Rudel. A multiscale model of the electrohysterogram the BioModUE_PTL project. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 7448–7451. IEEE, 2013.

[74] J. M. Marshall. Regulation of activity in uterine smooth muscle. *Physiological Reviews Supplement*, 5:213–227, 1962.

[75] K. Martinez and J. Cupitt. Vips — a highly tuned image processing software architecture. In *IEEE International Conference on Image Processing*, volume 2, pages 574–577. IEEE, 2005.

[76] C. McCloskey, C. Rada, E. Bailey, S. McCavera, H. A. van den Berg, J. Atia, D. A. Rand, A. Shmygol, Y.-W. Chan, S. Quenby, J. J. Brosens, M. Vatish, J. Zhang, J. S. Denton, M. J. Taggart, C. Kettleborough, D. Tickle, J. Jerman, P. Wright, T. Dale, S. Kanumilli, D. J. Trezise, S. Thornton, P. Brown, R. Catalano, N. Lin, S. K. England, and A. M. Blanks. The inwardly rectifying K^+ channels Kir7.1 controls

uterine excitability throughout pregnancy. *EMBO Molecular Medicine*, 6(9):1161–1174, 2014.

[77] A. R. Means, I. C. Bagchi, M. F. VanBerkum, and B. E. Kemp. Regulation of smooth muscle myosin light chain kinase by calmodulin. *Advances in Experimental Medicine and Biology*, 304:11–24, 1991.

[78] R. N. Miftahof and H. G. Nam. *Biomechanics of the Gravid Human Uterus*. Springer Science & Business Media, Berlin, Germany, 2011.

[79] S. M. Miller, R. E. Garfield, and E. E. Daniel. Improved propagation in myometrium associated with gap junctions during parturition. *American Journal of Physiology*, 256(1):C130–C131, 1989.

[80] H. Miyoshi, M. B. Boyle, L. B. MacKay, and R. E. Garfield. Voltage-clamp studies of gap junctions between uterine muscle cells during term and preterm labor. *Biophysical Journal*, 71(3):1324–1334, 1996.

[81] A. P. Moreno, J. G. Laing, E. C. Beyer, and D. C. Spray. Properties of gap junction channels formed of connexin 45 endogenously expressed in human hepatoma (SKHep1) cells. *American Journal of Physiology – Cell Physiology*, 268(2):C356–C365, 1995.

[82] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IEEE*, 50(10):2061–2070, 1962.

[83] M. Nixon and A. Aguado. *Feature Extraction & Image Processing*. Newnes, Oxford, UK, 2002.

[84] M. K. O'Connell, S. Murthy, S. Phan, C. Xu, J. Buchanan, R. Spilker, R. L. Dalman, C. K. Zarins, W. Denk, and C. A. Taylor. The three-dimensional micro-and nanostructure of the aortic medial lamellar unit measured using 3D confocal and electron microscopy imaging. *Matrix Biology*, 27(3):171–181, 2008.

[85] G. O'Grady, P. Du, L. K. Cheng, J. U. Egbuji, W. J. E. P. Lammers, J. A. Windsor, and A. J. Pullan. Origin and propagation of human gastric slow-wave activity defined by high-resolution mapping. *American Journal of Physiology – Gastrointestinal and Liver Physiology*, 299(3):G585–G592, 2010.

[86] T. Osa and T. Katase. Physiological comparison of the longitudinal and circular muscles of the pregnant rat uterus. *Japanese Journal of Physiology*, 25(2):153–164, 1975.

[87] T. Osa. An interaction between the electrical activities of longitudinal and circular smooth muscles of pregnant mouse uterus. *The Japanese Journal of Physiology*, 24(2):189–203, 1974.

[88] N. Otsu. A threshold selection method from gray-level histograms. *Automatica*, 11(285–296):23–27, 1979.

[89] H. C. Parkington, R. Harding, and J. N. Sigger. Co-ordination of electrical activity in the myometrium of pregnant ewes. *Journal of Reproduction and Fertility*, 82(2):697–705, 1988.

[90] R. Putz and R. Pabst. *Sobotta - Atlas of Human Anatomy.* Elsevier, Munich, 14$^{\text{th}}$ edition, 1989.

[91] C. Rabotti and M. Mischi. Propagation of electrical activity in uterine muscle during pregnancy: A review. *Acta Physiologica*, 213(2):406–416, 2014.

[92] C. Rabotti, M. Mischi, L. Beulen, G. Oei, and J. W. M. Bergmans. Modeling and identification of the electrohysterographic volume conductor by high-density electrodes. *Biomedical Engineering, IEEE Transactions on*, 57(3):519–527, 2010.

[93] L. G. Roberts. *Machine perception of three-dimensional solids.* PhD thesis, Massachusetts Institute of Technology, 1963.

[94] A. Rosset, L. Spadola, and O. Ratib. Osirix: an open-source software for navigating in multidimensional DICOM images. *Journal of Digital Imaging*, 17(3):205–216, 2004.

[95] T. W. Sadler. *Langman's Medical Embryology.* Lippincott Williams & Wilkins, Maryland, USA, 12$^{\text{th}}$ edition, 2012.

[96] K. M. Sanders. A case for interstitial cells of cajal as pacemakers and mediators of neurotransmission in the gastrointestinal tract. *Gastroenterology*, 111(2):492–515, 1996.

[97] G. C. Schoenwolf, S. B. Bleyl, P. R. Brauer, and P. H. Francis-West. *Larsen's Human Embryology.* Elsevier, Amsterdam, Netherlands, 4$^{\text{th}}$ edition, 2008.

[98] R. E. Sheldon. *The role of gap junctions in the excitability of the myometrial smooth muscle network.* PhD thesis, University of Warwick, 2014.

[99] R. E. Sheldon, M. Baghdadi, C. McCloskey, A. M. Blanks, A. Shmygol, and H. A. van den Berg. Spatial heterogeneity enhances and modulates excitability in a mathematical model of the myometrium. *Journal of the Royal Society Interface*, 10(86):20130458, 2013.

[100] R. E. Sheldon, C. Mashayamombe, S.-Q. Shi, R. E. Garfield, A. Shmygol, A. M. Blanks, and H. A. van den Berg. Alterations in gap junction connexin43/connexin45 ratio mediate a transition from quiescence to excitation in a mathematical model of the myometrium. *Journal of the Royal Society Interface*, 11(101):20140726, 2014.

[101] R. E. Sheldon, A. Shmygol, H. A. van den Berg, and A. M. Blanks. Functional and morphological development of the womb throughout life. *Science progress*, 98(2):103–127, 2015.

[102] J. C. Skou. The influence of some cations on an adenosine triphosphate from peripheral nerves. *Biochimica et Biophysica Acta*, 23(2):394–401, 1957.

[103] B. Spronck, J. J. Merken, K. D. Reesink, W. Kroon, and T. Delhaas. Ureter smooth muscle cell orientation in rat is predominantly longitudinal. *PloS One*, 9(1):e86207, 2014.

[104] W. D. Stein and T. Litman. *Channels, Carriers, and Pumps*, chapter 6 - Primary Active Transport Systems, pages 247–328. Elsevier, London, UK, 2015.

[105] K. D. Toennies. *Guide to Medical Image Analysis*. Springer Berlin Heidelberg, Berlin, Germany, 2012.

[106] T. Tomita. Electrophysiology of mammalian smooth muscle. *Progress in Biophysics and Molecular Biology*, 30:185–203, 1976.

[107] W.-C. Tong, C. Y. Choi, S. Karche, A. V. Holden, H. Zhang, and M. J. Taggart. A computational model of the ionic currents, Ca^{2+} dynamics and action potentials underlying contraction of isolated uterine smooth muscle. *PloS One*, 6(4):e18685, 2011.

[108] L. Vercruysse, S. Caluwaerts, C. Luyten, and R. Pijnenborg. Interstitial trophoblast invasion in the decidua and mesometrial triangle during the last third of pregnancy in the rat. *Placenta*, 27(1):22–33, 2006.

[109] W. Wang, H. H. Huang, M. Kay, and J. Cavazos. GPGPU accelerated cardiac arrhythmia simulations. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 724–727. IEEE, 2011.

[110] R. C. Webb. Smooth muscle contraction and relaxation. *Advances in Physiology Education*, 27(1–4):201–206, 2003.

[111] S. Weiss, T. Jaermann, P. Schmid, P. Staempfli, P. Boesiger, P. Niederer, R. Caduff, and M. Ba. Three-dimensional fiber architecture of the nonpregnant human uterus determined ex vivo using magnetic resonance diffusion tensor imaging. *The Anatom-*

ical Record Part A: Discoveries in Molecular, Cellular, and Evolutionary Biology, 288(1):84–90, 2006.

[112] G. Wolfs and H. Rottinghuis. Electrical and mechanical activity of the human uterus during labour. *Archiv für Gynäkologie*, 208:373–385, 1970.

[113] P. Wooding and G. Burton. *Comparative placentation: structures, functions and evolution*. Springer Science & Business Media, Berlin, Germany, 2008.

[114] J. Xu, S. N. Menon, R. Singh, N. B. Garnier, S. Sinha, and A. Pumir. The role of cellular coupling in the spontaneous generation of electrical activity in uterine tissue. *PloS One*, 10(3):e0118443, 2015.

[115] R. C. Young. Myocytes, myometrium, and uterine contractions. *Annals of the New York Academy of Sciences*, 1101:72–84, 2007.

[116] R. C. Young and R. O. Hession. Three-dimensional structure of the smooth muscle in the term-pregnant human uterus. *Obstetrics & Gynecology*, 93(1):94–99, 1999.

[117] D. Yu, D. Du, H. Yang, and Y. Tu. Parallel computing simulation of electrical excitation and conduction in the 3D human heart. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 4315–4319. IEEE, 2014.

[118] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21:977–1000, 2003.

Appendix A

Canny edge detection

The Canny edge detector is a well-established technique for obtaining feature edges in an image [83]. Given a greyscale image, as shown in Figure A.1**A**, the Canny edge detector is able to determine the edges, as shown in Figure A.1**B**. To obtain this edge image, the edge detector has four main steps [21]:

(i) Gaussian smoothing

(ii) Sobel edge detection

(iii) Non-maximal suppression

(iv) Hysteresis thresholding

These steps are described in the following sections. The first two steps use convolution of an image $I = [I(x, y)]_{(x,y)\in\Omega_I}$ with a square mask $M = [M(i, j)]_{i,j=-r}^{r}$. This operation is given by

$$(M * I)(x, y) := \sum_{i,j=-r}^{r} M(i, j)I(x + i, y + j),$$

where r is the radius of M. In the Sections A.1 and A.2, the process for finding the relevant mask is described. Sections A.3 and A.4 detail how to convert the output of the application of these masks to the image into the final edge image.

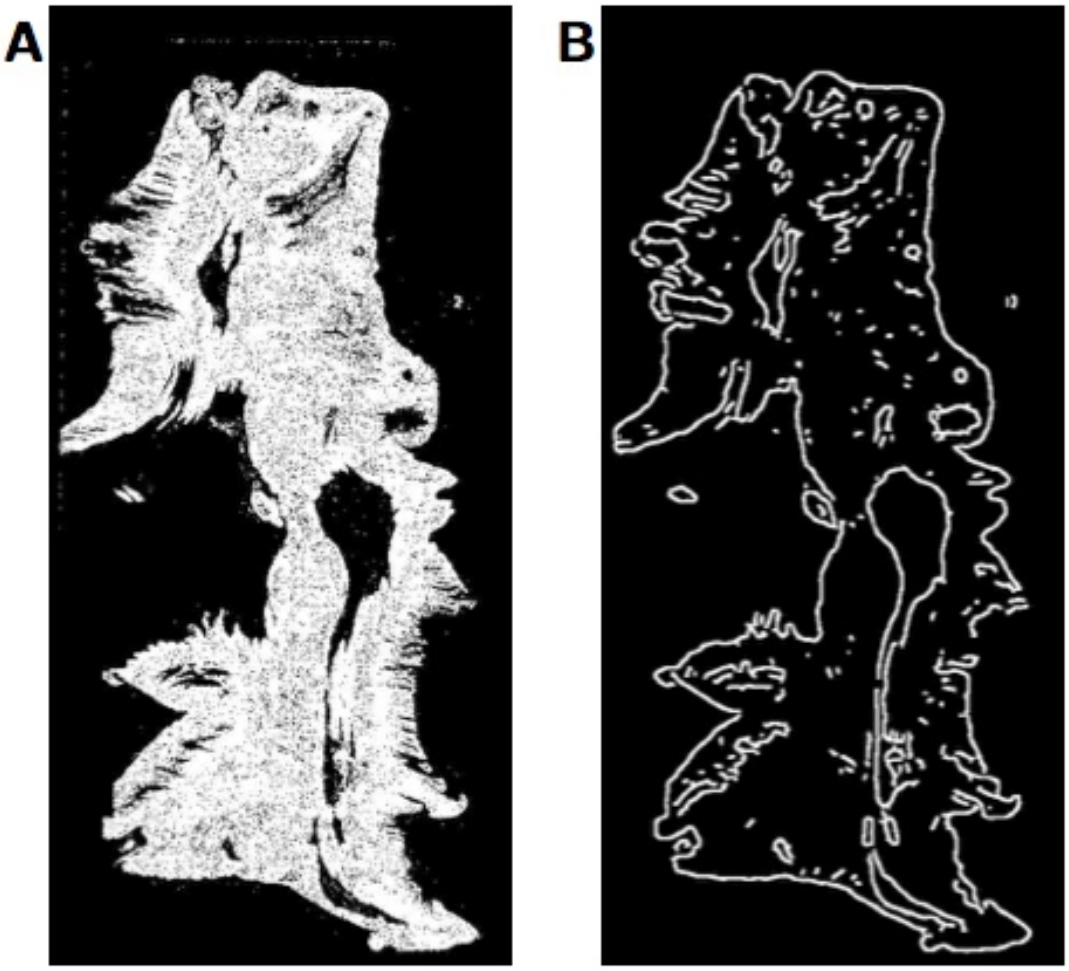

Figure A.1: **A**: An example of a greyscale slide image, taken from Chapter 2. **B**: The edge image found by applying the Canny edge detector to **A**.

A.1 Gaussian smoothing

The process of Gaussian smoothing reduces the noise present in the image [83]. The Gaussian filter with variance σ^2 and radius r has a mask M defined by

$$M(i,j) := \frac{1}{2\sigma^2\pi} \exp\left\{\frac{-(i^2+j^2)}{2\sigma^2}\right\}.$$

A 5×5 Gaussian mask is given in Figure A.2**A**.

A

0.003	0.011	0.024	0.011	0.003
0.011	0.045	0.101	0.045	0.011
0.024	0.101	0.225	0.101	0.024
0.011	0.045	0.101	0.045	0.011
0.003	0.011	0.024	0.011	0.003

B

0.05	0.21	0.48	0.21	0.05

Figure A.2: Examples of the Gaussian mask in one and two dimensions. **A**: The approximate values of a two-dimensional Gaussian mask, with $\sigma = 1.0$. **B**: The approximate values of a one-dimensional Gaussian mask, with $\sigma = 1.0$.

The performance of Gaussian smoothing can be improved by using one-dimensional masks in each direction. This improvement utilises the fact that higher dimensional Gaussian kernels are separable. Consider the convolution of an image I with M at a point (x, y):

$$\begin{aligned}(M * I)(x,y) &= \sum_{i,j=-r}^{r} \frac{1}{2\sigma^2\pi} \exp\left\{\frac{-(i^2+j^2)}{2\sigma^2}\right\} I(x+i, y+j) \\ &= \sum_{i,j=-r}^{r} \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-j^2}{2\sigma^2}\right\} \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-i^2}{2\sigma^2}\right\} I(x+i, y+j) \\ &= \sum_{i=-r}^{r} \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-i^2}{2\sigma^2}\right\} \sum_{j=-r}^{\lfloor w/2 \rfloor} \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-j^2}{2\sigma^2}\right\} I(x+i, y+j)\end{aligned}$$

Thus, by defining the one-dimensional row mask M_x and one-dimensional column mask M_y by

$$M_x(i) := \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-i^2}{2\sigma^2}\right\}$$

$$M_y(j) := \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{\frac{-j^2}{2\sigma^2}\right\},$$

the convolution can be rewritten as

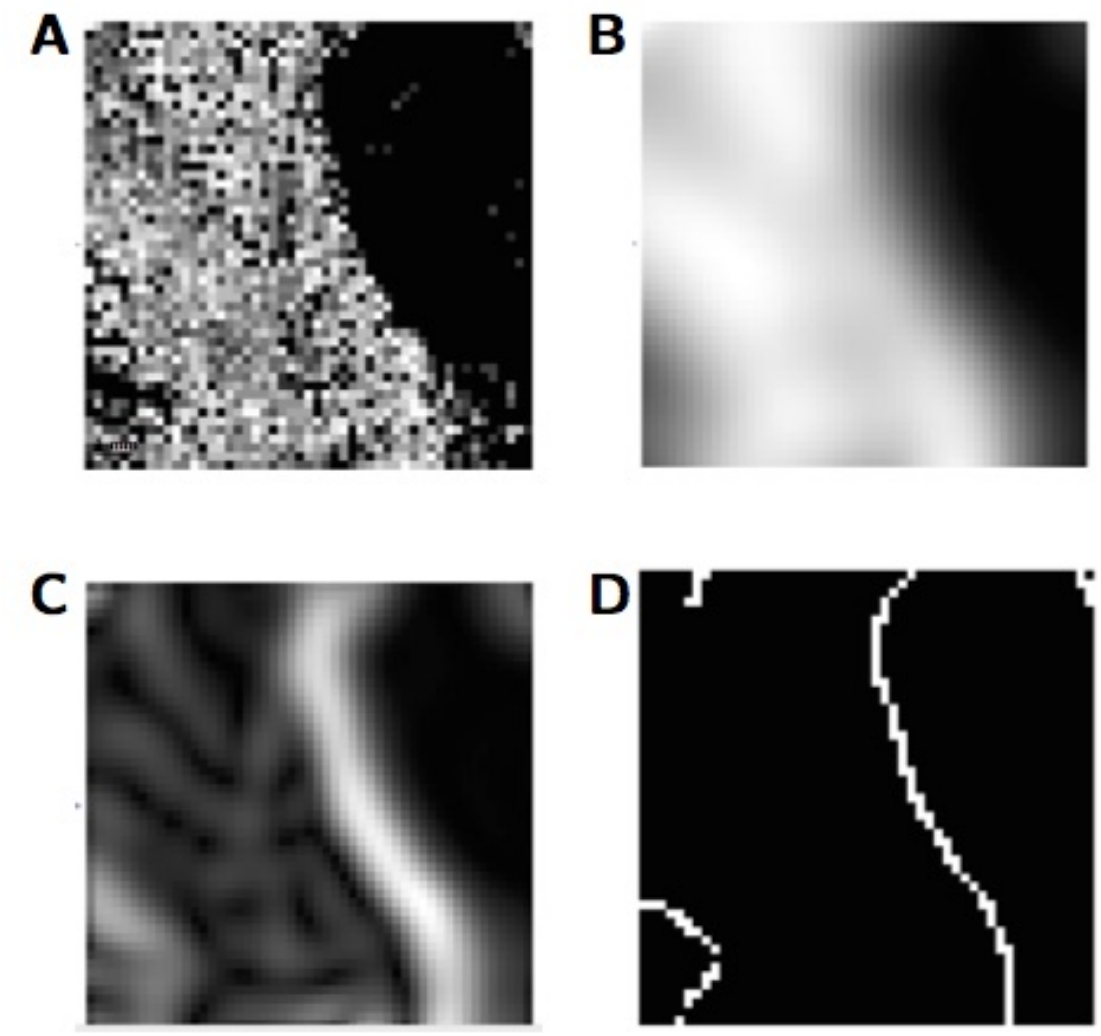

Figure A.3: The effect of the steps in the Canny edge detection algorithm. **A**: A close-up region of the greyscale image shown in Figure A.1**A**. **B**: The effect of convolving **A** with a Gaussian kernel. This convolution blurs the image and hence reduces the influence exerted by noise. **C**: The magnitude image given by combining the images generated by convolving **B** with the Sobel kernels. This image shows the magnitude of the gradient of the grey values. **D**: The edge image generated by performing non-maximum suppression and hysteresis thresholding. The application of non-maximum suppression reduces the edges to well-defined lines, and the hysteresis thresholding reduces the level of noise in the image.

$$(M * I)(x, y) = (M_x * (M_y * I))(x, y).$$

An example of a one-dimensional row mask M_x is given in Figure A.2**B**. The effect of Gaussian filtering on an image is illustrated in Figure A.3**B**.

A.2 Sobel edge detection

The Sobel edge detector determines the gradient of the greyscale values in the image [34]. The gradients in the x and y directions are calculated separately. Thus this edge detector actually uses two masks, one for the x-gradient and one for the y-gradient. The two masks, however, are defined in the same way, as swapping the axes of one mask gives the other.

We shall define the mask M_x, with window length w, which is used to determine the gradient in the x-direction. The entries of the mask are defined by a combination of smoothing and differencing. The smoothing uses the binomial coefficients as an approximation to the Gaussian function. The binomial coefficients are defined as follows:

$$\binom{n}{k} := \begin{cases} \frac{n!}{(n-k)!k!} & \text{if } 0 \leq k \leq n \\ 0 & \text{otherwise} \end{cases}.$$

Thus smoothing function $s(y)$ is given by

$$s(y) := \binom{w-1}{y},$$

for $0 \leq y < w$. The upper bound of y is $w-1$, so setting $n = w-1$ in the above identity is required to make $s(y)$ approximate a Gaussian function centred at $y = \lceil w/2 \rceil$. The difference function $d(x)$ takes the difference between the smoothing coefficients, given by

$$d(x) := \binom{w-2}{x} - \binom{w-2}{x-1},$$

for $0 \leq x < w$. Since the value of d is defined as the difference between two binomial coefficients, setting $n = w-2$ is required to ensure that $d(x)$ approximates the difference between two points of a Gaussian function centred at $x = \lceil w/2 \rceil$. The entries of M_x are

1	0	-1
2	0	-2
1	0	-1

Figure A.4: An example of a Sobel mask for calculating the gradient of an image in the x-direction.

thus defined by a smoothing in y and difference of smoothed values in x, given by

$$M_x(x-r, y-r) := s(y)d(x),$$

where $r = (w-1)/2$ is the radius of M_x. An example of this mask is shown for $w = 3$ in Figure A.4. The result of convolving an image with this mask, and the analogous mask in the y-direction, is shown in Figure A.3**C**. These two images can be combined to determine a magnitude of the gradient. Let I denote an image, and I_x and I_y be the images generated after convolving I with M_x and M_y, respectively. The magnitude image, I_m, of the gradient images is then given by

$$I_m(i,j) = \sqrt{I_x(i,j)^2 + I_y(i,j)^2}.$$

This magnitude image is shown in Figure A.3**C**.

A.3 Non-maximal suppression

Once the gradients have been found, the next step is to ensure uniqueness of the edge points. Consider the magnitude image shown in Figure A.3**C**. This image shows an edge. However, the exact position of the points on this edge are not clearly defined. Non-maximal

suppression is used to determine the precise line defining the edge.

For a given point (x, y) in an image I, if (x, y) is an edge point, then there is a direction, n, which is normal to the edge at (x, y). For (x, y) to be a true edge point, the magnitude of the gradient at (x, y), given by $I_m(x, y)$, will be maximal along n. To determine n, the gradient images I_x and I_y are used. The gradient vector, $(I_x(x, y), I_y(x, y))$, is normal to the edge at (x, y), and so gives the direction of n, as illustrated in Figure A.5. Since I_m is a discrete set of values, the nearest values of I_m along n need to be interpolated. The precise coordinates used for this interpolation depend on the direction of n. The possible situations arising from this dependency are analogous to each other, however, and so the case in Figure A.5 can be assumed. The aim is to interpolate a value for I_m at point p_0, as shown in Figure A.5. The distance from $(x - 1, y)$ to p_0 is given by

$$|p_0 - (x - 1, y)| = \frac{I_y(x, y)}{I_x(x, y)}$$

and the distance from $(x - 1, y - 1)$ to p_0 is given by

$$|p_0 - (x - 1, y - 1)| = \frac{I_x(x, y) - I_y(x, y)}{I_x(x, y)}.$$

Thus the interpolated value $\hat{I}_m(p_0)$ is given by

$$\hat{I}_m(p_0) = \frac{I_y(x, y)}{I_x(x, y)} I_m(x - 1, y) + \frac{I_x(x, y) - I_y(x, y)}{I_x(x, y)} I_m(x - 1, y - 1).$$

The interpolation of $\hat{I}_m(p_1)$ is determined in the same way. The criterion for (x, y) being an edge point is defined as

$$I_m(x, y) > \max(\hat{I}_m(p_0), \hat{I}_m(p_1)). \tag{A.1}$$

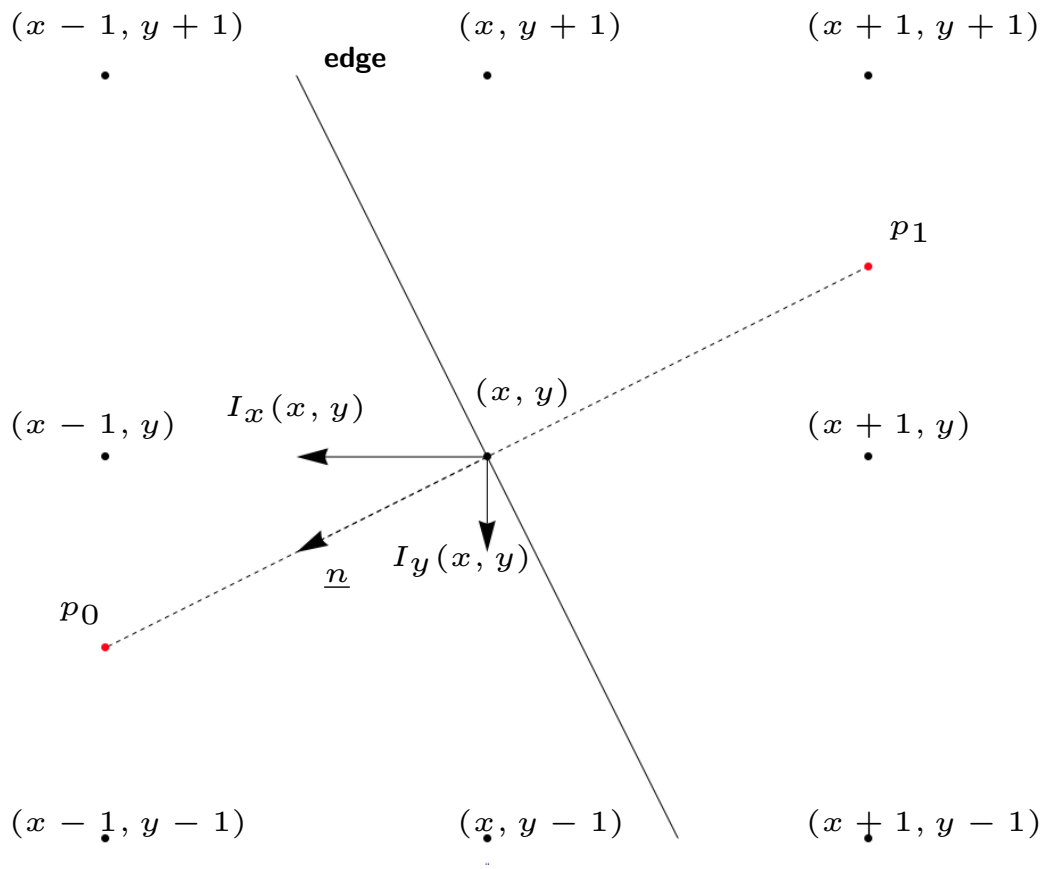

Figure A.5: The geometry used for calculating the interpolated values of I_m. The normal vector n is defined to be the unit normal in the direction of the gradient vector $(I_x(x,y), I_y(x,y))$. The points p_0 and p_1 are the points at which the line through (x,y) along the normal vector intersect the square formed by the points adjacent to (x,y). The interpolated values of I_m at p_0 and p_1 are used to determine if I_m is locally maximal at (x,y) along n. Adapted from Nixon and Aguado [83].

This test is applied to all points in I_m, and produces a new image, I_e, defined by

$$I_e(x,y) = \begin{cases} I_m(x,y) & \text{if criterion A.1 holds for } (x,y) \\ 0 & \text{otherwise} \end{cases} .$$

This image I_e is used in the next step to produce the final edge image.

A.4 Hysteresis thresholding

The final step to find the edges in the image is hysteresis thresholding. This thresholding technique uses two thresholds, an upper threshold t_u and a lower threshold t_l. All pixels in I_e with grey value greater than t_u are marked as edge pixels. The remaining pixels are marked by the following recursive procedure:

(i) Given a pixel p marked as an edge pixel, consider the set of pixels S_p adjacent to p.

(ii) For each $q \in S_p$, if $I_e(q) > t_l$, then mark q as an edge pixel.

(iii) Repeat this process for q.

The use of this thresholding procedure accomplishes reduction of noise but is less likely to discard valid edge points than thresholding alone. The result of this thresholding procedure is illustrated in Figure A.3**D**. This step completes the process, and the result of applying the full edge detection algorithm to the image in Figure A.1**A** can be seen in Figure A.1**B**.