

Model-driven stimulation for targeted functional restoration in chronic spinal cord injury



Conor Keogh
Trinity College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2023

Acknowledgements

I would like to express my gratitude to Professor James FitzGerald, my primary supervisor, for his constant support and encouragement while carrying out this work. The work environment that he created allowed me to develop new skills and explore new areas, leading to the work outlined here and much outside it. It has been a great pleasure to be a part of this group and I hope to continue to be involved in the next stages of my career.

I am also indebted to Professor Alex Green, my secondary supervisor, for his support. He was always available to discuss and provide insight into any aspect of my work and helped to make the experience of working in the department so enjoyable.

I would like to thank the broader Oxford Neural Interfacing group for their advice, suggestions and insight, as well as for challenging me and pushing me to advance my work any time I presented it. In particular, I would like to thank Professor Brian Andrews, Professor Jonathan Jarvis and Dr. Adrian Poulton, whose regular critical feedback was essential in shaping this work.

The Oxford Functional Neurosurgery group provided a further source of support. I would like to thank Mr. Martin Gillies for his regular input, and Professor Tipu Aziz for his advice and encouragement during the early stages of this work.

I would like to acknowledge the other students of the Oxford Neural Interfacing group. They offered a constant source of professional and personal support and many have become both collaborators and friends. It was a particular pleasure to see the group grow and the variety of work being done continue to expand. These include Linshan Chu, Siobhan Hall, Giovanni Rolandino and Max Stewart.

The students of the Oxford Functional Neurosurgery provided a further source of support, and I look forward to continuing to work with them in both academic and clinical settings. These include Ashley Raghu, Tariq Parker, John Eraifej, Amir Zand, Allie Deli, Sean Martin and Casey Rosso.

I would like to thank my parents, Martin and Patricia Keogh, and my sisters, Aoife and Orla Keogh. They provided a regular source of personal support.

I would particularly like to acknowledge the support of my fiancé, Laura Cox. Her constant support and encouragement was essential to the completion of this work.

I would like to thank a number of organisations for playing important roles in my life at various stages of this work. Trinity College Oxford provided support,

socialisation and accommodation during the early stages of this work. The Oxford University Boat Club provided challenge, structure and distraction during the middle stages. The Department of Neurosurgery at the John Radcliffe Hospital provided clinical insight, career direction and even greater distraction during the final stages.

Finally, I would like to thank the funders that made this work possible. These include the Engineering and Physical Sciences Research Council, the National Institute for Health and Care Research, the University of Oxford Clarendon Fund and Trinity College Oxford. Without their support this work could not have been carried out.

The support outlined above allowed this work to be carried out. However, I would also like to acknowledge that the presented work differs from that originally planned, and to acknowledge the support of my supervisory team, research group and funding bodies in adapting my research plan and navigating difficult circumstances. This work had originally been planned to include preliminary computational investigations followed by pilot validation studies in humans. The coronavirus pandemic and resulting periods of lockdown, closure of laboratory facilities and cessation of all experimental work meant that the focus of the work was changed to have a much stronger emphasis on the development of the computational work. While it was disappointing to not be able to carry this work through to validation studies during the period of the DPhil, the resulting emphasis on the computational work has been rewarding and valuable, and I am looking forward to carrying this on to validation studies as I progress my career.

Abstract

Spinal cord injury is a significant cause of disability worldwide. Interruption of white matter tracts results in a loss of function below the level of the injury with a devastating impact on quality of life. Recent work on epidural electrical stimulation combined with rehabilitation has, however, challenged the traditional narrative that the resulting disability is irreversible following the acute phase. Using spatiotemporally patterned stimulation to activate specific spinal nerve roots at corresponding phases of the gait cycle, weight bearing and volitional movement have been restored in patients with clinically complete, chronic spinal cord injuries.

Functional gains made during stimulation-aided rehabilitation were maintained even in the absence of ongoing stimulation, suggesting that stimulation facilitated repair of the damaged neural circuits. Interestingly, this raises the possibility of using non-invasive electrical stimulation during rehabilitation to provide functional gains without the need for a permanent implant. However, it is not clear whether non-invasive stimulation can reliably produce activation of spinal structures, whether the structures activated are the same as those activated by epidural stimulation, whether selective activation of specific regions for spatiotemporal targeting is possible, what the optimal electrode position is or how to account for inter-individual differences in anatomy. Answering these questions requires the ability to model the effects of variations in electrode positions and stimulation parameters on detailed models of individual anatomy and physiology. The techniques necessary to carry out this analysis currently do not exist in a form that allows these questions to be answered.

We therefore developed a set of techniques for modelling the effects of electrical neuromodulation on individuals. We present methods for deriving patient-specific computational models directly from routine clinical imaging with simulated placement of electrodes, for solving for the electric field produced by stimulation on these detailed models, for coupling these electric field models to detailed biophysical axon models to investigate the physiological effects of stimulation and for optimising electrode positioning and stimulation parameters to maximise activation of target structures. This provides a comprehensive set of computational methods for carrying out detailed investigations of the effects of electrical neuromodulation.

Using these techniques, we investigated the feasibility of targeted non-invasive spinal cord stimulation. We demonstrate that non-invasive stimulation of spinal

structures is feasible and is likely to mediate its effects via activation of nerve roots, and that targeted activation of segmental nerve roots at a single level unilaterally is possible with specific electrode positioning. We further outline the principles required to achieve targeted non-invasive spinal cord stimulation in individuals. These results demonstrate that targeted stimulation of the spinal cord with non-invasive methods is feasible and provides the principles necessary to do this, opening the possibility for individually targeted non-invasive spinal cord stimulation based on detailed computational models for stimulation-aided rehabilitation in spinal cord injury.

Contents

List of Figures	xvi
Outline	1
I Introduction	3
1 Neuromodulation for functional restoration	4
II Theory	15
2 A physical framework for therapeutic neuromodulation	16
III Implementation	42
3 Computation of activating fields for approximation of the orientation-specific neural response to electrical stimulation	43
4 Patient-specific computational models for simulating interventions	55
5 A finite element solver for neuromodulation problems	81
6 Individualised models for biophysically optimised neuromodulation	105
IV Application	134
7 Feasibility and mechanism of transcutaneous spinal cord stimulation: a computational study	135
V Conclusion	154
8 Conclusion	155

Outputs	168
Appendices	
A Derivation of variational forms for numerical solution of physical problems in neuromodulation	171
B NRRDosurgery: a generalisable tool for patient-specific neuromodulation	181
C A system for generating complex stimulation waveforms for human neuromodulation	196
References	206

Contents (detailed)

List of Figures	xvi
Outline	1
I Introduction	3
1 Neuromodulation for functional restoration	4
1.1 Spinal cord stimulation & neural repair	4
1.1.1 Spinal cord injury and loss of function	4
1.1.2 Neuromodulation for spinal cord injury	5
1.1.3 Augmenting rehabilitation	7
1.2 Noninvasive stimulation	8
1.2.1 Transcutaneous neuromodulation	8
1.2.2 Modelling of stimulation	11
1.3 Contributions	13
1.3.1 Targeted non-invasive spinal cord stimulation	13
1.3.2 Modelling for individualised therapy	13
II Theory	15
2 A physical framework for therapeutic neuromodulation	16
2.1 Electric fields	17
2.1.1 The electrodynamics of neuromodulation	17
2.1.2 Finite element approximations	19
2.2 Electrode-tissue interface	23
2.2.1 Coupling between electrical and neural systems	23
2.2.2 Charge transfer	23
2.2.3 Mathematical models	24
2.3 Biophysics	25
2.3.1 Multicompartment cable models	26
2.3.2 Hodgkin-Huxley dynamics	28

2.3.3	Myelinated axon models	31
2.3.4	Measures of activation	32
2.4	Approximations of neural activation	33
2.4.1	Simplified metrics	33
2.4.2	The activating function	34
2.5	Optimisation of stimulation parameters	35
2.5.1	Activation as a function of stimulation parameters	35
2.5.2	Numerical optimisation	37
2.6	Technical challenges	38
2.6.1	Embedding information in domains	39
2.6.2	Coupling electric fields and biophysics	40
2.6.3	Clinically relevant optimisation	41
III	Implementation	42
3	Computation of activating fields for approximation of the orientation-specific neural response to electrical stimulation	43
3.1	Introduction	43
3.2	Methods	46
3.2.1	Electric field computation	46
3.2.2	Biophysical model	48
3.2.3	Computation of activation	49
3.3	Results	49
3.3.1	Orientation sensitivity	50
3.3.2	Performance comparison	50
3.4	Discussion	53
4	Patient-specific computational models for simulating interventions	55
4.1	Introduction	55
4.2	Methods	58
4.2.1	Strategy	58
4.2.2	Implementation	59
4.2.3	Application	63
4.3	Results	66
4.3.1	Generation of patient-specific models from imaging	66
4.3.2	Simulation of interventions	67
4.3.3	Noninvasive spinal cord stimulation	69
4.3.4	Targeting interferential stimulation	72
4.4	Discussion	75

4.4.1	Targeted noninvasive spinal cord stimulation	75
4.4.2	A generalisable system for patient-specific modeling	77
4.4.3	Precision neuromodulation	78
4.4.4	Conclusion	79
5	A finite element solver for neuromodulation problems	81
5.1	Introduction	81
5.2	Methods	86
5.2.1	Physical problem	86
5.2.2	Implementation	87
5.2.3	Analytical validation	87
5.2.4	Numerical validation	89
5.2.5	Application to neuromodulation	90
5.2.6	Performance	91
5.3	Results	91
5.3.1	A custom finite element solver	91
5.3.2	Complex parameters	96
5.3.3	Performance	99
5.4	Discussion	101
5.4.1	Solving clinical neuromodulation problems	101
5.4.2	An open-source, high-performance solver	102
5.4.3	Conclusion	103
6	Individualised models for biophysically optimised neuromodulation	105
6.1	Introduction	106
6.2	Methods	112
6.2.1	Electric field computation	112
6.2.2	Biophysical axon models	113
6.2.3	Identifying trajectories	114
6.2.4	Mapping axons to trajectories	116
6.2.5	Computation of activation	118
6.2.6	Biophysical objective functions	118
6.3	Results	120
6.3.1	Derivation of axon trajectories	120
6.3.2	Coupling of models	123
6.3.3	Strength-duration curves	124
6.4	Discussion	126
6.4.1	Linking anatomy and physiology	126
6.4.2	Automated axon mapping	127

6.4.3	Modelling the effects of stimulation	129
6.4.4	Individually optimised neuromodulation	130
6.4.5	Conclusion	131

IV Application 134

7	Feasibility and mechanism of transcutaneous spinal cord stimulation: a computational study	135
7.1	Introduction	135
7.2	Methods	137
7.2.1	Electrode montages	137
7.2.2	Biophysical models	138
7.2.3	Measuring activation	140
7.2.4	Parameter sweeps	142
7.3	Results	144
7.3.1	Spinal cord activation	144
7.3.2	Activation of foraminal structures	147
7.3.3	Optimisation of electrode montage	149
7.4	Discussion	151
7.4.1	Efficacy and mechanism	151
7.4.2	Targeting foraminal structures	151
7.4.3	Conclusion	153

V Conclusion 154

8	Conclusion	155
8.1	Model-driven spinal cord stimulation	155
8.1.1	Transcutaneous nerve root stimulation	155
8.1.2	Targeting stimulation	156
8.1.3	Functional restoration	158
8.2	Computational modelling of neuromodulation	159
8.2.1	Physical basis of neuromodulation	159
8.2.2	Modelling of electrical stimulation	160
8.2.3	Patient-specific therapy	161
8.3	Outlook	162
8.3.1	Limitations	162
8.3.2	Further development	165
8.3.3	Conclusion	166

Outputs 168**Appendices****A Derivation of variational forms for numerical solution of physical problems in neuromodulation 171**

A.1	Introduction	171
A.2	Constant current	173
A.3	Time-varying current	176
A.4	Conclusion	179

B NRRDosurgery: a generalisable tool for patient-specific neuro-modulation 181

B.1	Introduction	181
B.2	Methods	184
B.2.1	Model creation	184
B.2.2	Interface	187
B.2.3	Implementation	188
B.2.4	Example	188
B.3	Results	188
B.3.1	Multiple interfaces	188
B.3.2	Abstracted system	190
B.3.3	Personalised models	190
B.4	Discussion	192
B.4.1	A tool for targeted neuromodulation	192
B.4.2	Limitations	194
B.4.3	Conclusion	194

C A system for generating complex stimulation waveforms for human neuromodulation 196

C.1	Introduction	196
C.2	Methods	198
C.2.1	Circuit design	198
C.2.2	PCB manufacture	199
C.2.3	Device assembly	200
C.2.4	Firmware	200
C.2.5	Testing	201
C.3	Results	201
C.3.1	Stimulation device	201
C.3.2	Custom stimulation patterns	203

C.4	Discussion	204
C.4.1	A system for complex neuromodulation	204
C.4.2	Applications	204
C.4.3	Conclusion	205
References		206

List of Figures

2.1	Model of electric potential in tissue	20
2.2	Multicompartment cable model of a neuron	26
2.3	Hodgkin-Huxley neuron model	28
2.4	Patient-specific model of activation	36
2.5	Optimisation of stimulation paramters	37
3.1	Geometry and sample electric field	48
3.2	Rotated Hessian approximations	50
3.3	Polar plot	51
3.4	Amplitude threshold and rotated Hessian	52
3.5	Amplitude threshold and rotated Hessian are highly correlation . .	52
4.1	Flowchart for patient-specific modeling	59
4.2	Electrode montages	64
4.3	Image segmentation	67
4.4	Electrode insertion	68
4.5	Electric fields produced by a series of electrode montages	70
4.6	Effect of transabdominal 30mA DC stimulation	71
4.7	Effect of stimulation with complex montage	73
4.8	Control of depth and steering with interfering fields	74
4.9	Control of depth and steering with interfering fields	75
5.1	Simple geometry	88
5.2	electrode geometry	90
5.3	Analytical vs. numerical, simple	92
5.4	Time dependent performance	93
5.5	Analytical vs. numerical, complex	94
5.6	Detailed comparison	95
5.7	Image-based model, axial slice	97
5.8	Activating function, axial slice	98
5.9	Solver performance	99
6.1	Targets for axon trajectories	121

6.2	Target anatomy as point clouds	122
6.3	Trajectories on imaging	123
6.4	Response to stimulation	124
6.5	Determination of strength-duration curves using binary search . . .	125
7.1	Targets for axon trajectories	139
7.2	Schematic of dorsal root axon model	141
7.3	Range of electrode positions	143
7.4	Strength-duration curves, midline electrodes	144
7.5	Variations in electrode position	146
7.6	Activation threshold as a function of electrode position	147
7.7	Strength-duration curves for nerve roots, midline electrodes	148
7.8	Regions of selectivity for root activation	149
7.9	Optimised electrode positions	150
B.1	Flowchart of overall system design	184
B.2	Flowchart of solver engine	185
B.3	Flowchart of finite element analysis workflow	186
B.4	Graphical user interface	189
B.5	Example documentation page	191
B.6	Personalised model example	192
C.1	Stimulation device schematic	198
C.2	Stimulation device PCB	199
C.3	Stimulation device hardware	202
C.4	Stimulation output example	203

Outline

This thesis presents the development of a set of methods for carrying out detailed computational modelling of the effects of electrical neuromodulation using standard clinical imaging and the application of these techniques to investigate non-invasive spinal cord stimulation as a potential therapeutic modality for spinal cord injury. We demonstrate that targeted non-invasive stimulation of neural structures is possible and describe the principles of achieving this as a means of developing a technique suitable for clinical application to spinal cord injury.

Section 1 provides an introduction to the application of electrical neuromodulation to functional restoration in spinal cord injury, and of the importance of computational modelling techniques in the development of these therapies.

Section 2 describes a mathematical framework for modelling electrical neuromodulation. This demonstrates how the clinical problem of targeting stimulation can be re-formulated as an optimisation problem using a biophysical objective function. Each step of the process from selecting stimulation parameters to the resulting physiological effects can be modelled mathematically to create an individualised function describing the effects of stimulation, which can then be optimised to determine the ideal stimulation parameters or electrode positions for a desired effect.

Section 3 outlines the implementation of the theoretical framework outlined in Section 2. In order to solve for the optimal electrode positions and stimulation parameters required for achieving a desired effect in an individual, we demonstrate techniques for deriving reliable simplified metrics for investigating the effects of stimulation (Chapter 3), for creating patient-specific computational models with detailed anatomy and for simulating arbitrary electrode placements (Chapter 4), for solving the physical problem posed by electrical neuromodulation (Chapter 5), and for automatically coupling finite element models of the electric field in

detailed anatomy to biophysical axon models for assessing the physiological effects of stimulation and for optimising stimulation parameters using biophysical objective functions to achieve physiologically and clinically meaningful measures.

Section 4 applies these developed techniques to non-invasive spinal cord stimulation. This demonstrates the value of these methods in answering scientifically and clinically meaningful questions and provides a method for developing a non-invasive stimulation technique potentially suitable for clinical application to spinal cord injury rehabilitation.

Section 5 provides an overall conclusion and outlook for potential future applications of these techniques.

A number of appendices are included. **Appendix A** outlines the derivation of the variational forms of the system of partial differential equations used to model the physical problem of electrical neuromodulation. **Appendix B** describes a generalisable tool for carrying out patient-specific modelling of neuromodulation using the techniques developed. **Appendix C** demonstrates the development of a hardware system for flexible control of stimulation parameters in order to allow for clinical testing and translation of the results of computational modelling.

Part I

Introduction

1

Neuromodulation for functional restoration

Contents

1.1 Spinal cord stimulation & neural repair	4
1.1.1 Spinal cord injury and loss of function	4
1.1.2 Neuromodulation for spinal cord injury	5
1.1.3 Augmenting rehabilitation	7
1.2 Noninvasive stimulation	8
1.2.1 Transcutaneous neuromodulation	8
1.2.2 Modelling of stimulation	11
1.3 Contributions	13
1.3.1 Targeted non-invasive spinal cord stimulation	13
1.3.2 Modelling for individualised therapy	13

1.1 Spinal cord stimulation & neural repair

1.1.1 Spinal cord injury and loss of function

Injury to the spinal cord results in interruption of motor and sensory pathways [1]. This produces loss of function below the level of the injury, with loss of the ability to walk, bladder and bowel control, autonomic regulation, arm movement and even breathing depending on the level of the injury. This produces significant morbidity [1, 2]. The resulting disability may improve in the immediate phase

post-injury, but following this acute phase the neurological deficit has traditionally been viewed as being permanent and irreversible, with most intervention at this stage focused on assistive technology and adaptation to disability rather than on regaining function [3–5].

Spinal cord injury is commonly traumatic, and as such tends to affect young people [1]. Spinal cord injured patients can then survive for many years following injury with significant neurological disability. There are an estimated 2500 new spinal cord injuries every year in the UK alone, with a total of 50,000 spinal cord injured patients in the country [6]. In addition to their established disability, spinal cord injured patients have increased healthcare needs and are at risk of complications such as autonomic dysregulation, deep vein thromboses and pneumonia [1]. This represents a significant burden of morbidity and of healthcare costs. The ability to restore function in this patient population would have a major impact on individuals' quality of life and on the healthcare system.

1.1.2 Neuromodulation for spinal cord injury

The emergence of epidural spinal cord stimulation as a treatment modality for improving function in chronic spinal cord injury has challenged the belief that neurological deficits are permanent following the acute phase of injury. The use of epidural electrode arrays in combination with physiotherapy has been demonstrated to restore weight bearing and volitional movements, with participants able to regain locomotor function even many years following their injury [7–12].

Further evidence has shown the ability of epidural stimulation to restore autonomic control and trunk stability [13], highlighting the potential for electrical neuromodulation as a promising means for restoring function in this patient group.

The ability to produce motor responses with epidural electrode arrays placed over the posterior aspect of the spinal cord is mediated by activation of motor reflexes [14, 15]. Activation of sensory dorsal nerve roots as they enter the spinal cord at high intensities can produce the H-reflex, or at lower intensities can produce posterior root-muscles reflexes through activation of proprioceptive afferents [14–

18]. This allows production of motor activity at specific segmental levels with stimulation of posterior spinal elements [11, 18, 19]. Stimulation has further been demonstrated to alter excitability of local spinal networks, resulting in a reduction in spasticity and activation of pattern generator circuits, producing involuntary locomotor activity [19–25]. While it is possible to produce patterned activity with targeted activation of specific spinal motor circuits through these reflex mechanisms, this is entirely involuntary. It is, however, possible using the same approach to restore volitional control of movement by combining this neuromodulation approach with rehabilitation [7–12].

The precise factors important to achieving functional restoration following spinal cord injury remain unclear. One group of studies have focused on the use of simple, tonic stimulation paradigms in combination with high volumes of rehabilitation [7, 8]. This approach assumes that the underlying pattern of stimulation is less important, and that the essential factor is the rehabilitation load.

All of these studies are uncontrolled case series demonstrating improvement in functional measures with intensive rehabilitation. There are no controlled studies. In particular, it remains unclear what the actual magnitude of the effect of electrical stimulation on rehabilitation outcomes is.

The volume of rehabilitation received by participants in these series [8] greatly exceeds the routine rehabilitation received by most spinal cord injured patients. It is currently not entirely clear whether functional gains would be seen over time with intensive rehabilitation alone in this group. Conventional wisdom has dictated that there is little functional gain following the acute phase of injury; it may be that these patients would benefit from intensive rehabilitation alone.

There is clearly a need for higher-quality studies in this area. An ongoing trial is currently aiming to address this question of whether tonic epidural electrical stimulation can improve functional outcomes [26]. Using stimulation and sham stimulation, this aims to determine whether stimulation offers a benefit on top of intensive rehabilitation in these patients.

1.1.3 Augmenting rehabilitation

These studies aim primarily to use tonic stimulation to augment the effect of intensive rehabilitation. Other groups aim to use specifically targeted stimulation to facilitate motor function [9–11].

These studies are again uncontrolled. However, they appear to show a functional benefit without the need for the level of intensive rehabilitation used in other trials. This is suggestive of a specific benefit of targeted stimulation, although the exact contribution of patterned stimulation is hard to define with the existing evidence.

Much of the work around optimising the potential effects of electrical stimulation has focused on trying to understand how it interacts with the damaged nervous system.

Although the precise mechanism of action remains unclear, it appears as though residual neural pathways across the level of injury remain, even in patients with clinically complete injuries [27, 28]. While these residual fibres are not sufficient to induce voluntary movement in isolation, the addition of electrical stimulation to targeted rehabilitation appears to be sufficient to allow signals to cross the level of the injury and produce movement distally, with reorganisation of local spinal circuits [12, 28–32]. Interestingly, it appears that this effect is maintained in the absence of stimulation [9, 10, 33]. Once functional gains are made during stimulation-assisted rehabilitation, these gains are maintained with the stimulation system deactivated. This suggests that electrical stimulation in combination with rehabilitation produces some kind of neuroplasticity or regenerative effect that restores communication across the level of injury and potentially provides a means for developing long-term, stimulator-independent gains in function.

The ability to facilitate rehabilitation and encourage the transmission of neural signals across the level of injury is based on the use of spatiotemporally patterned stimulation [9]. Using detailed computational modelling of spinal circuits, stimulation patterns were designed that targeted specific spinal levels at the point in the gait cycle that those neural subpopulations are usually engaged [11, 19]. This targeted stimulation approach allowed for the injured neural pathways to

become active and to restore function across the injury. Targeting of other spinal levels and the use of other stimulation patterns further allows for the restoration of other functions, such as autonomic regulation of blood pressure reflexes [13]. In this way, the use of specifically targeted stimulation based on neurophysiology and detailed computational modelling allows for the restoration of normal neural dynamics across the injured spinal cord [11, 14, 18].

These results suggest a good mechanistic basis for the use of spatiotemporally patterned stimulation. This is supported by the case series data showing a functional benefit. This approach differs from the use of tonic stimulation with intensive rehabilitation by more strongly emphasising the role of electrical stimulation with a mechanistic basis.

While more rigorous clinical studies are still needed to accurately characterise the specific benefit conferred by spatiotemporally patterned stimulation - in contrast to tonic stimulation, or no stimulation at all - there is sufficient mechanistic and case series data to suggest that this is a promising approach to improving functional outcomes following spinal cord injury.

These technologies offer much promise to patients with spinal cord injuries. However, these new treatment modalities must be evaluated with appropriate rigor and caution. Their exact contribution has yet to be determined. Properly randomised, blinded studies will be needed to unmask the true effects of stimulation on top of rehabilitation. In addition, the minimum rehabilitation load will need to be defined. It may not be practically possible to offer the kind of intensive rehabilitation used in some of these trials on a routine basis. The specific contributions of stimulation and of rehabilitation need to be characterised prior to this approach being impactful on a large scale.

1.2 Noninvasive stimulation

1.2.1 Transcutaneous neuromodulation

The ability to maintain functional gains achieved during stimulation-facilitated rehabilitation raises the possibility of applying stimulation only during rehabilitation

sessions. If sufficient targeting can be achieved using non-invasive electrical stimulation, this could be used in place of an implanted system during rehabilitation, providing the benefits of electrical stimulation without the need for a permanent implant. This would reduce the risk associated with neuromodulation for functional restoration in spinal cord injury and would make this emerging therapeutic approach available to a much larger group.

Non-invasive electrical stimulation of the spinal cord has been demonstrated to produce posterior root-muscle reflexes, analogous to those produced by epidural stimulation [21, 34]. This suggests that non-invasive stimulation may be able to produce activity in the spinal cord using similar mechanisms to that produced by epidural stimulation. Similarly, targeting different regions of the spinal cord along the rostrocaudal axis has been shown to produce activity in different muscle groups [34–36], and to produce step-like movements, suggesting activation of motor circuitry non-invasively [35, 37].

The application of non-invasive spinal cord stimulation to spinal cord injury has been investigated in a number of studies [35, 37–42]. A recent systematic review of these studies has indicated that while all showed some effect of stimulation, there were significant issues with the evidence generated [43]. The protocols and outcome measures used are inconsistent and poorly standardised, and in many cases unsuitable to determine whether there was any useful long-term effect, although they appear to indicate that activating motor circuitry is possible. Similarly, most studies did not have a control group or used only within-subject comparisons, although this difficulty with implementing a placebo control is common with many applications of neuromodulation.

While these results generally indicate a positive signal towards functional effects, small sample sizes and poor study design make confident statements about the effect of noninvasive stimulation challenging. Although many studies with widely varying protocols report positive effects, the level of evidence of any individual study is low in isolation, and the effects of publication bias need to be considered when assessing groups of low-quality studies in aggregate; it is likely that many of the studies

showing positive effects would not have been published had they not produced these effects. Equally, many negative studies not evident in the literature may exist. The more accessible nature of non-invasive technologies has meant greater access, but this has also resulted in a large volume of poorly standardised studies.

On reviewing these studies on non-invasive spinal cord stimulation for spinal cord injury, electrode positioning emerges as a key parameter for determining the response achieved [43]. Although not standardised, most studies use stimulating electrodes placed in the midline at the T11/T12 or L1/L2 levels, or at C3/C4 or C5/C6 for upper limb targets [35, 37–42]. Return electrode positions varied substantially, with most studies using either a dorsoventral arrangement with a return electrode on the anterior body wall or a rostrocaudal arrangement with a return electrode on the back either above or below the stimulating electrode.

This lack of standardisation of experimental protocols, stimulation methods and reporting of experimental methods and results is a significant issue with the literature surround noninvasive electrical stimulation. The determination of standardised protocols and electrode positions to achieve activation of specific targets remains a priority for advancing the field towards robust clinical studies.

Similarly, it is not clear what level of selectivity is possible. Alterations of electrode positioning suggest that some rostrocaudal selectivity is possible [34–36], and there is some evidence for lateralised effects [44]. However, it is not clear whether it is possible to achieve targeted activation of neural populations at specific segmental levels in order to achieve spatiotemporally patterned stimulation analogous to the methods used for rehabilitation with epidural stimulation.

Much of the work on noninvasive stimulation has not benefited from the same level of mechanistic motivation as has been seen in epidural stimulation. While advances in epidural stimulation have been driven forwards by detailed computational modelling and insights into the mechanisms underlying the effects of stimulation, much of the non-invasive literature relies on unprincipled placement of electrodes and inconsistent stimulation protocols.

Overall, the literature on noninvasive stimulation in spinal cord injury is highly varied. Systematic reviews suggest that this is a promising approach [43], however these results need to be interpreted cautiously. Effects have been reported over short follow-up times with a wide range of stimulation set ups. Importantly, some of these results suggest a level of spatial selectivity that is greater than might be anticipated for noninvasive methods.

Larger, more rigorous studies are required to better characterise the potential benefits of noninvasive approaches. In order to move in this direction, the field would benefit from a firmer mechanistic basis, as has been used to drive forwards progress in the area of spatiotemporally patterned epidural stimulation. Identifying potential mechanisms and building stimulation trials around these would be beneficial to provide greater direction to tests of noninvasive stimulation

1.2.2 Modelling of stimulation

Computational modelling of non-invasive stimulation has provided some insights into its mechanism of action [45–48]. Simple models suggest that non-invasive spinal cord stimulation can achieve activation of spinal nerve roots. However, these models are simplistic geometric models and do not represent true anatomy. While useful for these questions, they lack the detail required to examine the effects of electrode position and to determine the level of selectivity that it is possible to achieve. It has been repeatedly demonstrated that the inclusion of detailed anatomy in models of neuromodulation therapies produces more reliable results and allows individualisation of therapies [49, 50]. In order to determine whether spatiotemporally patterned stimulation using non-invasive methods is feasible and what electrode positions are needed to achieve this, we therefore require a method for explicitly representing detailed anatomy in our models. Further, in order to be able to account for inter-individual differences in anatomy, we require a method that can be easily applied to new individuals to optimise treatment in order to make this model-guided stimulation approach clinically viable.

The use of geometric models for investigation, as has been standard, is therefore unsuitable. Existing methods for individualised anatomy involve significant simplifications of anatomy and laborious workflows [49, 50], making them also unsuitable.

Existing methods used in noninvasive stimulation lack explicit representation of detailed anatomy; the use of techniques developed in other applications can potentially help to address this. However, existing methods will need to be adapted to maximise their utility in this area.

New methods are therefore required in order to thoroughly investigate this question and to develop an approach to model-guided non-invasive spinal cord stimulation that is potentially useful. An ideal method for modelling the response to stimulation for investigating the effects of neuromodulation, determining optimal parameters and individualising therapies should allow for computational models to be developed with explicit representation of detailed, patient specific anatomy. Further, it is necessary to be able to simulate interventions. Modelling solutions in other domains are dependent on post-procedure imaging [49–51]; this is unsuitable for individualisation of therapy and determination of optimal parameters for a given target as it requires electrodes to be held in a fixed position. Ideally, we need to be able to assess the effects of arbitrary electrode positions in detailed anatomy without the need for physically iterating over and imaging all electrode positions.

The use of biophysical axon models for modelling the effects of electrical stimulation has become standard in the field [11, 14, 18, 19, 50, 52]. We therefore require a method to link detailed anatomical models to these detailed biophysical models. Existing methods rely either on simple geometries [45–48], or on laboriously hand-drawn trajectories [14, 18]. Neither is suitable for the task of carrying out individualisation of therapy for model-driven treatment. Novel methods are therefore also required for coupling anatomical and physical models for this purpose. Lastly, in order to determine the optimal electrode positions required to achieve targeted spinal cord stimulation and to carry out individualisation of therapy, it is necessary to be able to use these models in order to determine the optimal parameters for a specific goal. The selection of electrode positions and stimulation parameters from a

set of clinically acceptable options is, in an informal sense, an optimisation problem; it is necessary to formalise this clinical problem as a mathematical optimisation problem and to solve this in order to answer the question of whether spatiotemporal patterned stimulation is possible with non-invasive methods and whether automated individualised of therapy is feasible.

1.3 Contributions

1.3.1 Targeted non-invasive spinal cord stimulation

We demonstrate that non-invasive spinal cord stimulation can reliably achieve activation of spinal structures. We show that it mediates its effects via activation of nerve roots. Further, we demonstrate that targeted activation of segmental nerve roots is possible, including control of the rostrocaudal level and laterality of stimulation. We outline the principles of electrode positioning for achieving this targeted stimulation, which differ from those used in the field up to now. This approach opens the possibility for spatiotemporally patterned stimulation using non-invasive methods. This represents a novel approach to non-invasive spinal cord stimulation which has significant potential applications to stimulation-facilitated rehabilitation for functional restoration following spinal cord injury. This approach will require validation in experimental studies. Off-target effects not modelled here, such as the activation of cutaneous afferent fibres and resultant discomfort, may limit the application of noninvasive stimulation in this area. However, this approach allows the design of principled, mechanism-focused studies for the assessment of noninvasive electrical stimulation, bringing an increased level of rigor to studies in this area.

1.3.2 Modelling for individualised therapy

We present a set of novel methods for carrying out individualised computational modelling of the effects of electrical stimulation using routine clinical imaging. This allows for automated derivation of detailed models of individual anatomy, simulation of arbitrary electrode placements and solving for the electric field produced by these electrodes in that individual. We further demonstrate methods for coupling these

detailed anatomical models to detailed biophysical axon models to investigate the effects of stimulation on specific neural structures. We show how these biophysical measures can be used to derive clinically and physiologically meaningful objective functions for numerical optimisation, allowing scientific questions to be investigated and for therapies to be optimised for individuals. This provides a comprehensive and generalisable set of computational methods for carrying out detailed investigations of the effects of electrical neuromodulation. Care must be taken to ensure that the generated models are adapted to fit the target area before interpreting results. However, this provides a useful means of generating patient-specific models of anatomical regions with coupled biophysical models while retaining local anatomy in a way which adds to established methods.

Part II

Theory

2

A physical framework for therapeutic neuromodulation

Contents

2.1	Electric fields	17
2.1.1	The electrodynamics of neuromodulation	17
2.1.2	Finite element approximations	19
2.2	Electrode-tissue interface	23
2.2.1	Coupling between electrical and neural systems	23
2.2.2	Charge transfer	23
2.2.3	Mathematical models	24
2.3	Biophysics	25
2.3.1	Multicompartment cable models	26
2.3.2	Hodgkin-Huxley dynamics	28
2.3.3	Myelinated axon models	31
2.3.4	Measures of activation	32
2.4	Approximations of neural activation	33
2.4.1	Simplified metrics	33
2.4.2	The activating function	34
2.5	Optimisation of stimulation parameters	35
2.5.1	Activation as a function of stimulation parameters	35
2.5.2	Numerical optimisation	37
2.6	Technical challenges	38
2.6.1	Embedding information in domains	39
2.6.2	Coupling electric fields and biophysics	40
2.6.3	Clinically relevant optimisation	41

2.1 Electric fields

Electrical neuromodulation is based on the application of electricity to influence the activity of the nervous system. By understanding how our selection of stimulation parameters interacts with the body to produce an electric field within the tissues and how that electric field interacts with excitable tissues to produce a physiological effect, we can start to build an integrated physical framework that describes how our selection of stimulation parameters is linked with our desired effects. This allows us to mathematically model how activation in specific structures is related to our clinical decisions. By re-expressing this as an optimisation problem, we can then develop a means of rigorously identifying the mathematically optimal electrodes and stimulation parameters to activate specific target structures.

2.1.1 The electrodynamics of neuromodulation

The behaviour of electrical charges is well characterised experimentally and can be accurately described by mathematical theory. The laws of electromagnetism describe how electrical charges behave and interact with one another [53, 54]. Applied charges produce an electric field. This is broadly defined as the region within which the charges exert a force on other charges. This effect drives changes in biological tissues by exerting an effect on specific cellular components, producing a physiological effect.

Applying a voltage, or difference in electric potential, results in the flow of current within the underlying tissues. This current, effectively a "flow" of charge, results in changes in the electric potential throughout the applied electric field. These potential differences can then drive a biological effect in excitable tissues with voltage-sensitive components.

The amplitude and the direction of the current produced by stimulation is determined by the voltage applied, by the positions of the electrodes and by the geometry and electric properties of the underlying tissues. This current is closely related to the effects produced by stimulation. The "dose" of charge delivered is

determined by the amplitude of the applied current and the time for which it is applied. As the impedance between the electrodes is difficult to approximate and tends to change over time, resulting in changes in the applied current as the two are related by Ohm's law, it is common in modern systems to apply a regulated current rather than a regulated voltage. This results in alteration of the voltage applied at the stimulating electrodes as the impedance changes in order to maintain a fixed current output. This is current-controlled rather than voltage-controlled stimulation.

The behaviour of electric fields is well understood and can be described using systems of partial differential equations based on the laws of electromagnetism. Appendix A describes in detail how the system of equations describing the behaviour of electrical fields produced by stimulation are derived, beginning from physical conservation laws. Solving these equations allows for the behaviour and influence of electric charges to be accurately predicted. In some simple scenarios, such as individual point charges in homogeneous media, the system of equations can be solved directly to produce a general answer to how charges behave. In the context of clinical neuromodulation, the scenario is significantly more complex.

The electric field produced by stimulation can broadly be considered to be determined by stimulation factors and patient factors. Stimulation factors include the position, shape and material of the electrodes and the stimulation parameters used, including the amplitude of the voltage or current applied and the pulse width used for square-wave stimulation. Patient factors include the patient's anatomy and the electrical properties of the tissues. These factors interact to produce the electric field within the tissue which mediates the effects of stimulation. In practice, it is not possible to fully describe the complex interactions between applied currents and human anatomy mathematically in a way that permits a straightforward general solution. Rather, it is necessary to take advantage of well-developed methods for numerical approximation of solutions to complex systems of partial differential equations in order to determine results and predict effects for specific situations. This allows us to take advantage of our understanding of

the physics of electric fields to predict their behaviour even in situations which are too complex to solve directly mathematically.

2.1.2 Finite element approximations

The finite element method offers a means of approximating the solution to systems of partial differential equations [55]. It provides a well-developed means of computing solutions to arbitrarily complex systems of equations, allowing solutions to be determined even where the system is too complex to be solved analytically. This method is of enormous importance in science and engineering.

The finite element method is based on solving the complex system of equations in a specific domain under specific conditions rather than generally. In the case of neuromodulation, this means we can solve for the electric field produced by specific stimulation settings in specific anatomy, but we cannot produce a general expression for all stimulation settings in all anatomy. This is done computationally by, in essence, breaking down the complex domain into many small, simple problems, solving each of these and then combining them back together into an overall solution.

To do this, we must first define the set of equations describing the behaviour of electric fields produced by neuromodulation. We can then specify the conditions used, i.e. the stimulation applied, and the domain of interest, i.e. the anatomical region being investigated, and solve computationally. In order to use this method, the system of equations must be defined in a specific mathematical form to allow methods from variational calculus to be used to determine a numerical solution. The derivation of the equations describing clinical neuromodulation and their variational forms are outlined in appendix A.

Figure 2.1 shows the basic system of equations describing steady-state stimulation. This is relevant for non-time-dependent stimulation, for example when examining the effect of the "on" part of square-wave stimulation or for constant-current stimulation. This same approach can also be used for stimulation which changes over time if we assume that the changes are sufficiently slow that the system is in "steady state" at all times; this is quasistatic approximation. This may




Figure 2.1: System of partial differential equations determining the electric potential produced in tissue by electrical stimulation.

introduce errors around timepoints where the stimulation changes, but provides a computationally much simpler approach to time-varying stimulation.

This method can be straightforwardly extended to include time-dependent stimulation responses, such as when examining complex waveforms and responses with short timescales. This is done by altering this system of equations to incorporate time-dependent changes as outlined in appendix A.

The behaviour of steady-state currents in materials governed by Ohm's Law is given in the first equation shown in figure 2.1. σ is the electrical conductivity of the tissues, while V is the electric potential produced by stimulation. Ω is the domain within which we want to compute a solution. In our case, this is the region of anatomy of interest. The second equation describes the conditions on the external boundary of the domain, $\partial\Omega$. Here, we set an insulating condition, stating that no current passes through the external skin boundary. This defines the behaviour at the edges of the domain.

The last two lines of figure 2.1 describe the boundary conditions defining the stimulation we wish to apply. The surfaces of active, Γ_A , and return, Γ_R , electrodes are defined. We then apply stimulation by defining the voltage between the stimulating electrodes. In the case of voltage-controlled stimulation, the externally applied voltage, V_e , is defined directly.

In the case of current-controlled stimulation, the applied voltage V_e can be adjusted such that the total current across the active electrode is equal to the

desired input current: $\int_{\Gamma_A} \frac{\partial V}{\partial n} = J_{applied}$, where $J_{applied}$ is the applied current. In this way, voltage-controlled and current-controlled stimulation can be applied by determining how V_e is calculated.

Note that it is also possible to define a Neumann-type boundary condition on the active electrode to define current density directly for current-controlled stimulation. However, this generates a uniform current density across the electrode surface, which may not be appropriate.

This system of equations can then be solved computationally in order to calculate the electric potential throughout the domain of interest given the specified boundary conditions. This provides us with a general expression for the physical system of clinical neuromodulation. However, in order to provide a useful means of investigating the effects of stimulation with specific electrodes and parameters in specific individuals, we need to consider how to adapt this framework to the specific situations we are investigating.

We therefore define the set θ of all controllable parameters, i.e. those that are clinically determined.

$$\theta = \langle \text{geometry, position, amplitude, pulse width, } \dots \rangle$$

This includes the geometry and position of the electrodes and the stimulation parameters used. These are not fixed components of the physical model of neuromodulation and will change in any new scenario that is modelled. These set parameters will then result in changes to other parts of the physical framework:

$$\Omega = f_{\theta \rightarrow \Omega}(\theta) \tag{2.1}$$

$$J_{applied} = \theta_{amplitude} \tag{2.2}$$

Equation 2.1 shows that the domain we solve over is dependent on our selected parameters θ . While we cannot control individual anatomy, though this will change from person to person and is variable in that sense, we can control electrode position. This changes the overall domain by altering the position of electrode boundaries and therefore the field produced. These complex interactions between electrodes

and anatomy can be modelled mathematically by incorporating electrodes into the computational domain. The construct of the physical region that we consider therefore changes whenever the electrodes are moved. The overall computational domain, which is used to define a finite element mesh as part of the overall method, is linked to the clinically controllable parameters by some function $f_{\theta \rightarrow \Omega}$ which generates a new domain Ω for a given set of clinical parameters θ .

Equation 2.2 shows that the amplitude applied is determined directly by the clinical parameters selected. In the case of voltage-controlled stimulation, voltage is determined in an analogous way by setting this parameter based on the corresponding θ parameter. In time-dependent stimulation, specific waveforms can also be specified by setting the voltage or current at each timepoint.

We can then re-express the system of equations describing neuromodulation in terms of our clinically controllable parameters.

$$\begin{aligned} \nabla \cdot \sigma \nabla V &= 0 && \text{in } f_{\theta \rightarrow \Omega}(\theta), \\ \frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega, \\ V &= \theta_{\text{amplitude}} && \text{on } \Gamma_A, \\ V &= 0 && \text{on } \Gamma_R \end{aligned} \tag{2.3}$$

This gives us a set of partial differential equations describing the potential field produced by stimulation in a given individual for some set of clinical parameters defined by θ . This system can then be solved in a standard fashion to compute the electric potential throughout the domain of interest. Given that the clinical parameters θ are the only parameters not fixed by the physical system and all equations describing the relationships between parameters are constant, we can then define this as a fixed mapping to compute the electric potential produced for a given set of clinical parameters.

$$V = f_{\theta \rightarrow V}(\theta) \tag{2.4}$$

By making the appropriate selections of the parameters θ , we can then use this function to solve for the electric potential produced for any set of electrode geometries, positions or stimulation parameters. This gives us a general method for

determining the physical effects of stimulation in any individual for any given clinical parameters.

2.2 Electrode-tissue interface

The physical effects of applied stimulation are also influenced by the physics of the electrode-tissue interface. When considering the effects of electrical stimulation on biological systems, we must also consider how these two systems interact. There are multiple reactions at the interface with an important effect on the physical effects of stimulation which must be considered when modelling clinical neuromodulation.

2.2.1 Coupling between electrical and neural systems

The interface between electrodes and biological tissue is complex, with many factors affecting these interactions [56, 57]. These interactions are characterised by electrochemical reactions to facilitate charge transfer between electronic and biological systems.

Charge is carried in electronic systems by free electrons. When we apply electrical stimulation, we apply a current by changing the flow of electrons out of the surface of the electrodes. However, in biological tissues charge is usually carried as charged ions. It is therefore not possible to directly transfer charge between a stimulating electrode and the underlying tissue. This requires either capacitive charge transfer or a series of reactions at the electrode-tissue interface to facilitate the interaction between electronic and biological systems.

2.2.2 Charge transfer

Charge transfer primarily occurs via two mechanisms: capacitive charging and Faradaic reactions [56–58]. The relative importance of each type of charge transfer differs depending on the electrode material used, and each has differing behaviour in terms of ability to transfer charge and potential for tissue damage. Electrode material therefore has significant impacts on stimulation, including on charge

injection capacity, side effects and total stimulation limits beyond simply the gross electrical properties of the material selected.

Capacitive charging occurs where build-up of electrons at the electrode-tissue interface due to stimulation causes local ion species to form an electrode-electrolyte double layer [59]. This double layer can then charge and discharge, passing electrons between the metallic conductor of the stimulating electrode and ions in the electrolyte surrounding the electrode. This produces a stable interface over time, even with repetitive stimulation. However, this mechanism has a relatively low charge injection capacity, limiting the amount of charge that can be injected for a given electrode surface area.

Faradaic charge transfer is characterised by chemical reactions at the electrode-tissue interface. Oxidation and reduction of ionic species at the electrode-electrolyte interface results in rapid charge transfer between electrode and tissue. This allows for a large charge injection capacity, facilitating high stimulation volumes even with very small electrodes. However, this comes at the cost of potentially irreversible chemical reactions at the interface causing gradual degradation of the electrodes over time, and the potential for tissue damage. Biphasic stimulation waveforms, where the polarity of stimulation is rapidly reversed after stimulation, are often used to try to take advantage of rapid Faradaic charge transfer mechanisms while rapidly reversing electrochemical reactions to avoid long-term irreversible degradation [56–58].

2.2.3 Mathematical models

The complex interactions at the electrode-tissue interface can be incorporated into models of the effects of stimulation. Particularly in the case of microelectrodes, these interactions can have important effects on the behaviour of stimulation [60].

The most common model of the electrode-tissue interface is a simple capacitor in line with a resistor [58]. The capacitive component represents the formation of the capacitive double layer for passive charge transfer. The resistive component represents the resistance to transfer of metal ions into the electrolyte due to electrochemical reactions. Altering the relative proportions of these two components

determines the behaviour of the electrode-tissue interface. More complex models have also been developed which attempt to take into account other factors such as the diffusion of ions around the electrode-tissue interface [61–63].

By altering the boundary conditions used in our system of partial differential equations, it is possible to explicitly incorporate the impedance at the electrode-tissue interface [64]. This can be done by applying a Robin boundary condition at the interface.

$$\sigma \frac{\partial V}{\partial n} = g(V_{\text{electrode}} - V) \quad (2.5)$$

This describes an interface where the current through the electrode is proportional to the voltage difference across the electrode, $V_{\text{electrode}} - V$, and the surface conductance g .

Note that where the current density over the surface of the electrode is close to uniform, this interface can be described using a Neumann boundary condition as described earlier. For this reason, it is more common to model the electrode-tissue interface as either a Neumann-type boundary condition for describing current-controlled stimulation or a Dirichlet-type boundary condition for describing voltage-controlled stimulation [65–69].

2.3 Biophysics

The laws of electromagnetism and an understanding of the electrode-tissue interface can tell us how stimulation will interact with the body and what the resulting electric field will look like. This does not tell us what the physiological effect of stimulation will be. We need to then understand the influence that the electric field produced by stimulation will have on excitable tissues. Biophysical models allow us to describe mathematically the electrical behaviour of excitable tissues and their interaction with electric fields. This allows us to predict the physiological effects produced by electrical neuromodulation.

2.3.1 Multicompartment cable models

In order to understand how electric fields relate to neural activity, we need to understand how external electric fields influence the electrical activity of excitable tissues. Neurons, or other excitable tissues, can be modelled using multicompartment cable models [70]. These models are the gold standard for describing the physiological activity of neurons. The "cable theory" of excitable tissues is based on the recognition that the cable equation, a partial differential equation derived for applications to transatlantic telegraph cables by Lord Kelvin, could be applied to describe the electrical behaviour of parts of neurons [71–73].

To describe the activity of complex cells, excitable tissues are considered to be made up of a series of "compartments". Each compartment is electrotonically compact, i.e. has uniform electrical properties, and is described by a set of equations. These compartments are then coupled together to describe how adjacent segments influence each other. This allows for even very complex cells with variable morphology and electrical properties to be described by considering individual segments separately.

Based on the law of conservation of currents, we assume that the currents into and out of each segment are equal, i.e. that no current is spontaneously created or destroyed inside a cell. The currents influencing any given segment of tissue can then be described as a sum of a local capacitive current, active ionic currents, a passive leak current and an axial current along the neuron, as shown in figure 2.2.

$$C_m \frac{d(V_{i,n} - V_{e,n})}{dt} + I_{ionic,n} + G_a (V_{i,n} - V_{i,n-1}) + G_a (V_{i,n} - V_{i,n+1}) = 0$$

Figure 2.2: Equation describing the behaviour of a segment of a multicompartment cable model neuron with capacitive, ionic and axial currents.

The capacitive current describes a passive current across the cell membrane, as the cell's phospholipid bilayer causes it to act as a capacitor, i.e. an insulating material with a difference in charge across it. This current is determined by the membrane capacitance, a fixed physical quantity for a given segment, and the change in the potential across the membrane over time. $V_{i,n}$ is the intracellular potential, i.e. the potential inside the cell relative to a distant ground, at the n th segment, while $V_{e,n}$ is the extracellular potential, i.e. the potential just outside the cell membrane relative to a distant ground, at the same segment. As the local intracellular and extracellular potentials change, this passive current flows in or out of the cell.

The ionic current describes the current into or out of the cell driven by active ion channel dynamics. This produces much of the rich electrophysiological behaviour of neurons. These dynamics result in the generation of action potentials and other behaviours, often in response to passive changes in membrane potential. The dynamics of this ionic current are described in more detail below.

The axial current is determined by the axial conductance, G_a , i.e. the reciprocal of the resistance to current flow along the inside of the cell, and by the difference in intracellular potential between the present compartment, $V_{i,n}$ and adjacent compartments, $V_{i,n-1}$ and $V_{i,n+1}$. Therefore, if there is a change in intracellular potential in a given compartment, driven by a change in ionic currents or capacitive currents, this will drive a current into or out of adjacent compartments. This allows, for example, action potentials to propagate by driving changes in intracellular potential at nerve segments adjacent to those that have experienced a sudden ionic current.

Note that in any given segment, the local membrane voltage is given by the difference between the local intracellular and extracellular potentials, each measured relative to a distant ground.

$$V_m = V_i - V_e \quad (2.6)$$

This formulation of membrane voltage V_m is more commonly used in neurophysiology. Changes in this membrane voltage is what triggers action potential generation.

The goal of electrical neuromodulation is therefore frequently to manipulate the membrane voltage of specific target cell populations.

The effects of electrical neuromodulation are mediated by generating an electric field in the tissue, governed by the laws of electromagnetism as described previously. This electric field then alters the extracellular electric potential, V_e . Based on the cable model of neurons, this can then drive a passive capacitive current into the cell by altering the membrane voltage. These changes in membrane voltage can then produce complex electrophysiological responses by inducing activation of local ion channels and action potential generation.

2.3.2 Hodgkin-Huxley dynamics

The behaviour of voltage-sensitive ion channels and their contribution to active membrane potential dynamics such as action potential generation can be described mathematically. This can be done by describing the behaviour of active ion channels as a set of ionic currents driven by alterations in the membrane conductance of specific charged ion species in response to changes in membrane potential. As described by Hodgkin and Huxley, this behaviour can be captured by a set of ordinary differential equations [74]. Solving the cable formulation of a neuron model with Hodgkin-Huxley dynamics allows us to predict the physiological effects of stimulation.

The diagram illustrates the Hodgkin-Huxley model equation for ionic current, I_{ionic} , with various components and their physical interpretations highlighted by colored arrows and labels:

- Leakage current:** Indicated by a blue bracket over the first term $\bar{g}_L (V_m - E_L)$.
- Leakage conductance:** Points to \bar{g}_L (blue arrow).
- Membrane potential:** Points to V_m (blue arrow).
- Leakage equilibrium:** Points to E_L (blue arrow).
- Potassium conductance:** Points to \bar{g}_K (pink arrow).
- Potassium gating variable:** Points to n^4 (purple arrow).
- Potassium equilibrium:** Points to E_K (red arrow).
- Ionic current:** Indicated by a red bracket over the second and third terms.
- Sodium equilibrium:** Points to E_{Na} (green arrow).
- Sodium gating variables:** Points to m^3h (olive arrow).
- Sodium equilibrium:** Points to E_{Na} (light green arrow).

$$I_{ionic} = \bar{g}_L (V_m - E_L) + \bar{g}_K n^4 (V_m - E_K) + \bar{g}_{Na} m^3 h (V_m - E_{Na})$$

Figure 2.3: Equation describing the Hodgkin-Huxley model of ion channel dynamics.

Figure 2.3 shows the Hodgkin-Huxley equation describing the overall ionic current, I_{ionic} . This is made up of a sum of a constant leakage current and a number of variable ionic currents. Each current is defined by a conductance and a potential, where the potential driving each current is simply the difference between the current membrane potential V_m and the equilibrium potential for that current. This equilibrium potential is exactly equivalent to the reversal potential or Nernst potential, commonly used in electrophysiology.

The leakage potential is a current with a constant conductance which aims to drive the membrane potential to the leakage equilibrium. This contributes to the overall resting potential of the cell. This component of the ionic current is driven by channels open at rest and describes an overall, non-specific ionic current.

The potassium and sodium currents are explicitly represented in the equation due to their complex dynamics and importance in determining the electrical behaviour of excitable cells. The potential driving each current is determined by the current membrane potential, V_m , and the equilibrium potentials for potassium, E_K , and sodium, E_{Na} . The actual ionic current is determined by the driving potential and the conductance to that ionic species. Physically, this is determined by the number of ion channels that permit passage of that ion which are open at any given time. This will alter the conductance and therefore the current that flows for any given driving potential.

These dynamic ionic conductances are described by the maximum conductance and a gating term. The maximum potassium and sodium conductances, \bar{g}_K and \bar{g}_{Na} respectively, are fixed quantities for a given segment of cell membrane and are determined by the density of ion channels in that segment. These terms give the conductance to the relevant ion if all ion channels in the relevant segment were open simultaneously.

The actual ion conductance at any given time is determined by the gating variables, n for potassium and both m and h for sodium. These gating terms vary between 0 and 1, where 0 represents all channels closed, i.e. no conductance, and 1 represents all channels open, i.e. maximum possible conductance. These terms

vary over time and are dependent on the membrane potential. It is these terms that govern the complex dynamics of ionic currents in cell membranes.

Potassium channels contain four n gates which operate independently; these open and close slowly in response to changes in membrane voltage. They gradually open on membrane depolarisation, causing repolarisation; when the membrane has been repolarised, it closes slowly. This contributes to the after-hyperpolarisation and to the relative refractory period of the neuron.

The sodium channels dynamics are determined by three n gates, or activation gates, and one h , or inactivation, gate. On depolarisation, the n gates open rapidly, causing activation of the channel and passage of a sodium current. The h gate, which closes in response to depolarisation, then closes to inactivate the channel, preventing further depolarisation. The m gates will then close when the membrane is repolarised. However, even if the membrane is depolarised again, the channel cannot reopen while the h gate remains closed; this is the absolute refractory period. The h gate will then re-open, a process called de-inactivation, returning the channel to its initial state, able to respond again to depolarisation.

The dynamics of these gating terms are described in the remainder of the Hodgkin and Huxley equations:

$$\begin{aligned}\frac{dn}{dt} &= \alpha_n(V_m)(1 - n) - \beta_n(V_m)n \\ \frac{dm}{dt} &= \alpha_m(V_m)(1 - m) - \beta_m(V_m)m \\ \frac{dh}{dt} &= \alpha_h(V_m)(1 - h) - \beta_h(V_m)h\end{aligned}\tag{2.7}$$

where α and β are rate constants for each ion channel type. These describe how channel conductances vary with time and membrane potential. They were determined empirically and provide a highly accurate model of how excitable tissues behave.

The combination of a multicompartment cable model with Hodgkin-Huxley active ion channel dynamics provides a highly accurate means of describing the behaviour of neurons and other excitable tissues. Cable theory helps us to understand how electrical impulses are transmitted along axons with varying resistive and

capacitive features. By dividing complex axonal morphologies into multiple discrete segments and describing each as a "compartment" with uniform physical properties, coupled together along their lengths, we can model the behaviours of complex axons with varying physical parameters. By introducing ion channels with Hodgkin-Huxley dynamics to these compartments, we can accurately model the physiological behaviours of complex neurons.

In combination with an accurate model of electric field generation, this allows the physiological effects of electrical neuromodulation to be described mathematically with great precision. However, these models require some adaptation to represent the specific dynamics present in myelinated axons, often the target of neuromodulation interventions.

2.3.3 Myelinated axon models

Myelinated axons present more complex dynamics than models with uniform biophysical properties. The restriction of active ion channels to the nodes of Ranvier and the specific dynamics of segments insulated by myelin changes the behaviour of these cells. These factors must be taken into account to produce a realistic description of the behaviour of these cells, including dynamics such as saltatory conduction of action potentials. Many models of myelinated axons have been developed [75]. The McIntyre-Richardson-Grill axon model remains one of the most widely used [76].

This defines multiple repeating segments with differing dynamics, representing the nodes of Ranvier with active ion channel dynamics, the myelin attachment zone, the fluted sections between the attachment zone and the primary myelinated segment, and the fully insulated myelinated region, all with specific biophysical properties. This model further describes a passive current running between the myelin and the axonal membrane by adding an additional "external cable" in the periaxonal space defined by a myelin resistance and capacitance and the periaxonal axial conductance, as well as a means for this passive current to discharge through the myelin attachment zones.

This model has been shown to accurately describe the behaviour of myelinated axons [76] and has been shown to capture the response of these axons to electrical stimulation [77–79].

Multiple other models exist which capture different aspects of the behaviour of axons. For example, the Bostock model captures differences in the expression of ion channels and gating parameters between motor and sensory fibres to characterise the differences in electrophysiological behaviour between these fibre types [80, 81]. The selection of a specific model and the parameters will influence the results produced and needs to be well-matched to the question being investigated. However, the McIntyre-Richardson-Grill model is well suited to extending over long physical distances to simulate long axons and is well validated in modelling the response of axons to external electrical stimulation [76, 82–87], making it suitable for examining the link between externally applied electrical stimulation and activity in target structures. Adaptations of this model to incorporate features in other models, such as differential responses of motor and sensory fibres, may be useful in some circumstances.

2.3.4 Measures of activation

These descriptions of the electrical behaviour of neurons allow us to determine how the electric field produced by stimulation influences the membrane potential of specific target neurons. This allows us to define a mapping between the applied field and the membrane response.

$$V_m = f_{V \rightarrow V_m}(V) \quad (2.8)$$

Using these biophysical models, we can define a function $f_{V \rightarrow V_m}$ which describes the response of the membrane potential, V_m at our target to any given electric field, V , produced by stimulation. The link between the externally applied electric field and the membrane response is determined by the potential field produced in the tissue and the response of the neural tissue to this field, i.e. the relationship described in equation 2.8 is biophysically fixed, so we can use the mapping from

clinical parameters, θ , to external electric field described in equation 2.4 to define a mapping from clinical parameters to membrane response.

$$V_m = f_{\theta \rightarrow V_m}(\theta) \quad (2.9)$$

This provides us with a function $f_{\theta \rightarrow V_m}$ which links, via the electromagnetic behaviour of stimulation and the biophysics of cell membranes, our selection of stimulation parameters to the membrane response at our target structures. This gives us a rigorous mathematical description of the physiological effects of our clinical decisions and allows us to begin to robustly examine the effects of these parameters.

2.4 Approximations of neural activation

While a model such as equation 2.9, which links our selection of parameters to membrane activity, provides a highly realistic and rigorous means of evaluating the effects of stimulation, this comes at the cost of high computational complexity, and to date reliable implementation has been an issue. This can limit application in some settings.

2.4.1 Simplified metrics

For this reason, there has been significant interest in finding other metrics that predict neural activation. If we consider some metric of activation of our target structure, A , this therefore aims to find some function $f_{V \rightarrow A}$ which links the potential field produced by some model of electrical stimulation to a prediction of the effect of stimulation. This is often much simpler than modelling the intervening biophysics, but comes at the cost of reduced accuracy.

There has been much work on the idea of "volume of activated tissue" methods for approximating the effects of stimulation [78]. These aim to predict the neural response by using a range of simplified models to determine the area in which activation of neural structures is likely. A wide range of models have been used, from extremely simple models such as assuming a sphere of activation around a fixed point to sophisticated biophysical measures.

Metrics derived directly from the potential field have been used to approximate neural activation [51, 77, 88]. This is based on the assumption that the activity of neurons is directly related to the amplitude of the applied electric field. While simple to compute and providing a rapid response, this has little relationship to biophysical reality.

Other measures attempt to use simplified biophysical models, such as axons placed at fixed distances from the electrode, while some use an understanding of the biophysics of stimulation to derive measures that are related to neural activation [51, 77, 88].

2.4.2 The activating function

The activating function is a metric which describes the "driving force" of the potential field which produces a neural response [89, 90]. By manipulating the equation shown in figure 2.2, it is possible to show that the current flow into a given segment of a neuron is proportional to this activating function. This describes the influence of the external electric field on the membrane potential:

$$f(x) = \frac{d^2V_e}{dx^2} \quad (2.10)$$

Neural activation is therefore related to the second spatial derivative of the external potential along the axon. While this does not allow direct determination of the effects of stimulation, this provides a useful approximation of the effects of neuromodulation. However, this method was developed for analysis of the behaviour of one-dimensional neurons. It does not translate in a straightforward way to the setting of complex three-dimensional geometries.

Simplified metrics for determining measures of activation, A , based on a potential field, V , are valuable to the analysis of the effects of neuromodulation. These are, however, limited and current methods have significant issues with limit their use.

2.5 Optimisation of stimulation parameters

Creating a mathematical model which accurately describes the potential field produced by stimulation and the physiological effect produced by this potential field for a given set of clinical parameters allows us to rigorously investigate the effects of altering these parameters. Further, by deriving mathematical expressions which capture our clinical objectives, e.g. activation of specific target structure, we can apply optimisation methods to determine the optimal clinical parameters to achieve our end goal in any specified individual.

2.5.1 Activation as a function of stimulation parameters

After defining the effects of electrical neuromodulation in this mathematical framework, we can express our clinical goals in a mathematical expression. For example, we can aim to produce significant membrane potential responses, representing action potentials, in a proportion of the axons in a target structure. We can also express more complex goals, such as achieving activation at a given target while avoiding activation at a second location to minimise side effects. These goals can be expressed in terms of the desired effect on the membrane voltage, V_m , of target cells. We can then define our target metric, A , and a function that maps the membrane voltage of our target structures to this quantity which captures our therapeutic goal:

$$A = f_{V_m \rightarrow A}(V_m) \quad (2.11)$$

This function $f_{V_m \rightarrow A}$ gives us a measure of how well we have achieved our therapeutic goal following stimulation based on the membrane potential responses.

As shown in equation 2.9, we can link our selection of clinical parameters, θ , to the physiological effect of stimulation on membrane voltage, V_m . By combining this with equation 2.11, we can create a mapping from parameter selection to activation of our target structures:

$$A = f_{\theta \rightarrow A}(\theta) \quad (2.12)$$

This provides us with a rigorous mathematical model which explicitly links our selection of clinical parameters to activation at our target structures.

This allows us to use this framework to create detailed, patient-specific models which express our clinical goals, i.e. activation at specific targets, as a function of our selection of clinical parameters.




Figure 2.4: Derivation of patient-specific models of activation as a function of stimulation parameters by composition of numerical methods for simulating interventions, coupling the computed electric field to biophysical models and generating measures of activation from biophysical parameters.

Figure 2.4 shows how composition of the models we have described into an integrated physical system allows us to derive patient-specific models which determine the clinical effects of stimulation using specified parameters. This provides a mapping, $f_{\theta \rightarrow A}$, from our selection of clinical parameters, θ , to our expression describing our clinical goal, A .

The first term, introduced in equation 2.4, simulates a therapeutic intervention by determining the electric field produced in the patient's anatomy given simulated electrode placements with specified geometry and stimulation parameters. This can be solved using the finite element method to determine effect of electrical stimulation in the region of interest, giving us the potential field produced by stimulation with our desired parameters in specific patient anatomy.

The second term, introduced in equation 2.8, couples the potential field produced by stimulation to the physiological effects of stimulation. A detailed biophysical model is used which determines the effect of electrical stimulation on the membrane voltage at the target structures.

The third term, introduced in equation 2.11, expresses our measure of activation. This describes how we determine success based on membrane response, such as achieving action potential generation in target structures.

Combining these models gives us a means of expressing how our selection of clinical parameters determines the physiological effect of stimulation. This then allows us to systematically vary these input parameters in order to rigorously evaluate their effects on our target metrics, providing a firm physical basis for individualised parameter selection for individualised electrical neuromodulation.

2.5.2 Numerical optimisation

The derivation of patient-specific models which map clinical parameters to activation also allows us to go further than examining the effects of varying clinical parameters on activation. By explicitly representing the links between parameters and activation, we can apply optimisation methods to determine individually optimised parameters.

The system of equations making up this framework is complex and it is not possible to solve this system analytically. However, methods from numerical optimisation are well developed for this purpose [91]. This allows us to use our numerical function, $f_{\theta \rightarrow A}$, to determine the optimal set of clinical parameters to achieve our therapeutic goal in a specific individual by systematically varying the inputs and applying optimisation algorithms to identify the best possible set of inputs.

$$\theta_{optimal} = \arg \max_{\theta \in \theta'} f_{\theta \rightarrow A}(\theta) \quad (2.13)$$

Figure 2.5: Determination of optimum stimulation parameters within some constraints by numerical optimisation of a patient-specific function mapping selected parameters to biophysical measures of activation.

Figure 2.5 shows the application of this approach of framing our clinical goal as an optimisation problem in a rigorous physical framework. The patient-specific model, $f_{\theta \rightarrow A}$, describes how our selected stimulation parameters produce activation

at our target structures as described above. Our set of clinically controllable parameters, θ , is used as an input, producing a measure of the level of activation achieved given our current set of clinical parameters.

These input parameters, θ , are then allowed to vary over some set, θ' . This constrained optimisation approach allows clinical information to be incorporated into our optimisation approach. For example, we can define safe limits for stimulation of a possible set of permissible electrode locations. This ensures that the resulting set of parameters is both mathematically optimal and also clinically acceptable.

This optimisation approach will vary our parameters over the defined acceptable range in order to identify the set of parameters which produce the maximum level of activation. This allows us to determine $\theta_{optimal}$, the set of stimulation parameters which best achieves our therapeutic goal in that individual.

This provides a general physical framework for describing the effects of electrical neuromodulation and for determining individually optimised stimulation parameters by expressing clinical goals in this framework. This opens possibilities for individualised neuromodulation, targeting of specific structures, and the investigation and development of novel targeted neuromodulation techniques.

2.6 Technical challenges

This chapter has outlined an integrated physical framework for describing the physiological effects of electrical neuromodulation. It has further shown how clinical goals and decisions can be expressed in this framework, and how this allows optimisation methods to be applied to identify individualised parameters. However, while well-developed methods for numerically solving the systems of partial differential equations representing the behaviour of electrical fields [55], the ordinary differential equations describing the biophysical behaviour of excitable tissues [92], or the optimal inputs to maximise the output for a function [91], realisation of the potential impact of this integrated system for electrical neuromodulation requires several challenges to be overcome. Table 2.1 details the factors which need to be

implemented to allow this physical framework to be applied in practice, along with the chapters in which each part of this system are implemented.

Table 2.1: Numerical implementations required to solve problem

Function	Description	Implementation
$f_{V \rightarrow A}$	Simplified metric for approximating activation	Chapter 3
$f_{\theta \rightarrow \Omega}$	Generation of patient-specific meshes	Chapter 4
$f_{\theta \rightarrow V}$	Solver for finite element problems	Chapter 5
$f_{V \rightarrow V_m}$	Coupling of FEM and biophysical models	Chapter 6
$f_{V_m \rightarrow A}$	Definition of objective function for optimisation	Chapter 6
$f_{\theta \rightarrow A}$	Composition of solver, coupling and objective function	Appendix B

2.6.1 Embedding information in domains

The use of finite element models to simulate electrical stimulation is well-established [9, 11, 14, 18, 19, 49, 50]. However, this has generally been restricted to simplified geometric models [45, 47]. Attempts to include more detailed patient-specific anatomy have consistently shown that model performance is improved by representing anatomy in detail [49, 50]. This is because of the interactions between the tissue geometry and the electric field, which are lost in simplifications of anatomy. However, existing methods are restricted to either highly simplified geometries or laborious reductions in anatomic complexity in order to "parameterise" specific aspects of anatomy so that they can be expressed mathematically and modelled. This is highly limiting and prevents these methods from being used in a robust way for individualised modelling of neuromodulation.

Further, existing methods have been restricted to either simple geometries or inclusion of fixed electrodes based on post-operative positioning [45, 50, 51]. This is highly limiting, as it prevents examination of the potential effects of altering electrode position to inform clinical decision making. However, simulation of electrode positioning to model the effect of stimulation is limited by the same restrictions around need for highly simple geometries or parameterisations. This prevents the use of this system for true optimisation of electrodes.

Chapter 4 outlines the development of a method which derives detailed, patient-specific models suitable for finite element analysis directly from clinical imaging. Further, this shows how electrode positioning using arbitrarily complex electrode arrangements can be simulated and incorporated into the computational domain. By embedding information on highly detailed patient anatomy and electrode positioning directly in the computational domain, this avoids the need for simplifications and parameterisations. Chapter 5 shows how these detailed, individualised models can then be solved using standard methods, with results at least as good as gold-standard finite element solvers.

These developments allow us to automatically generate a computational domain, Ω , containing detailed individualised information on anatomy and electrodes, given a set of clinical parameters, θ . This function, $f_{\theta \rightarrow \Omega}$, can be combined with the custom solver to define a function $f_{\theta \rightarrow V}$ which generates the potential field produced by stimulation in highly detailed individual anatomy given a set of clinical parameters.

2.6.2 Coupling electric fields and biophysics

Having solved for the potential field produced by stimulation, it is necessary to be able to examine the physiological effects of stimulation using biophysical models. The use of these models to examine the effects of stimulation is well-established [9, 11, 14, 18, 19, 49, 50], however their use has again been limited to highly idealised geometries, limiting realism and application to individualised therapies. Where more detailed models have been used, these are limited by the need for hand-drawn trajectories, limiting practical application.

In order to implement an integrated framework for individualised neuromodulation, it is necessary to be able to map complex neuron geometries into individual anatomy and to couple these models to the simulated electric field to examine its effect. Chapter 6 describes a method for achieving this. This shows that complex biophysical models can be automatically coupled with complex anatomic models and potential fields, allowing the physiological effects of stimulation in detailed individualised models to be investigated. This provides a function $f_{V \rightarrow V_m}$ which

allows us to link the potential field produced by stimulation to the membrane voltage response.

2.6.3 Clinically relevant optimisation

In order to be useful for individualised neuromodulation, we need to be able to apply this framework to clinically meaningful problems. Chapter 6 shows how clinical goals can be expressed in a useful way in this framework, and how these clinically meaningful quantities can be used as goals for optimisation. This provides us with a method $f_{V_m \rightarrow A}$ which links the physiological response to stimulation to our clinical goals.

Chapter 3 shows how measures of activation can also be derived directly from the potential field using simplified metrics, i.e. a function $f_{V \rightarrow A}$. This provides an alternative, more computationally simple, means of carrying out rapid assessments and optimisation of stimulation in some settings.

Appendix B shows how these implementations for individualised modelling, solving the physical system produced by neuromodulation, coupling with biophysical models and clinically relevant optimisation can be combined into an integrated framework for individualised neuromodulation. This provides an overall function $f_{\theta \rightarrow A}$ which allows us to investigate the effects of clinical parameters on meaningful measures of activation in individualised models. This opens possibilities for individualised neuromodulation and the development of novel targeted electrical neuromodulation interventions.

Part III

Implementation

3

Computation of activating fields for approximation of the orientation-specific neural response to electrical stimulation

Contents

3.1	Introduction	43
3.2	Methods	46
3.2.1	Electric field computation	46
3.2.2	Biophysical model	48
3.2.3	Computation of activation	49
3.3	Results	49
3.3.1	Orientation sensitivity	50
3.3.2	Performance comparison	50
3.4	Discussion	53

3.1 Introduction

Computational methods for determining the neural response to extracellular stimulation are of increasing importance, both in the development of novel devices [9, 18, 52, 93] and the optimization of stimulation parameters for neuromodulation therapies [49, 50].

The biophysical response to extracellular stimulation can be modelled very

accurately using multi-compartment neuron models with Hodgkin-Huxley dynamics [74]. However, this is computationally intensive and does not easily provide visualizations of the regions activated by stimulation. In scenarios where more rapid iterations of designs and parameters are required, more computationally efficient metrics are desirable.

The activating function provides a classic simplified metric for approximating the response to stimulation [89, 90]. A detailed neuronal model shows that current flow in the n th segment of the model is determined by a capacitive current, an ionic current and a current along the inside of the axon:

$$C_m \frac{d(V_{i,n} - V_{e,n})}{dt} + I_{ionic,n} + G_a(V_{i,n} - V_{i,n-1}) + G_a(V_{i,n} - V_{i,n+1}) = 0$$

By introducing the reduced voltages

$$V_n = V_{i,n} - V_{e,n} - V_{rest}$$

We get:

$$\frac{dV_n}{dt} = \frac{G_a(V_{n-1} - 2V_n + V_{n+1} + V_{e,n-1} - 2V_{e,n} + V_{e,n+1}) - I_{ionic,n}}{C_m}$$

We then insert $G_a = \pi d^2 / 4\rho_i \Delta x$ and $C_m = \pi d L c_m$ (where L is the active membrane length, i.e. the length of the nodes of Ranvier, or $L = \Delta x$ in the case of unmyelinated axons) and introduce the ionic current density $i_{ionic,n}$:

$$\frac{dV_n}{dt} = \left(\frac{d\Delta x}{4\rho_i L} \left(\frac{V_{n-1} - 2V_n + V_{n+1}}{\Delta x^2} + \frac{V_{e,n-1} - 2V_{e,n} + V_{e,n+1}}{\Delta x^2} \right) - i_{ionic,n} \right) / c_m$$

Examining the source term, we can see that the contribution of the extracellular field produced by stimulation is given by

$$f_n(t) = \frac{V_{e,n-1} - 2V_{e,n} + V_{e,n+1}}{\Delta x^2}$$

The current flow in the n th segment is therefore proportional to this activating function, defined as the second difference of the extracellular potential along the axon. As $\Delta x \rightarrow 0$, this becomes the standard expression for the activating function:

$$f(x, t) = \frac{\partial^2 V_e(x, t)}{\partial x^2}$$

This simple metric provides a useful means of approximating the neural response to stimulation. However, this metric is defined as a second difference along a specific 1D trajectory (i.e. the length of the axon considered). It therefore only allows us to consider axons in specific locations with specific orientations, limiting the application to situations where we wish to compute the region likely to be activated by stimulation and to scenarios with homogeneous axon properties.

Metrics such as the volume of tissue activated [78] have attempted to overcome this issue for applications such as deep brain stimulation. Previous attempts have been made to derive measures analogous to the activating function from the simulated potential field [82], and to account for the relative orientations of electrode and axons in the calculation of the metric for approximating activation [77]. However, these metrics either disregard axon orientation or assume a tangential orientation of the axon with respect to the electrode. These assumptions make these methods unsuitable for accurate evaluation of the effect of stimulating fields on axons with specific orientations, such as those encountered in peripheral nerve stimulation, spinal cord stimulation or deep brain stimulation when specific fibre tracts are targeted.

We therefore require the ability to calculate an activating function in a manner that can be generalized to a field, while maintaining explicit representation of the target axons' orientation. We consider the Hessian matrix (H) of the electric potential produced by stimulation, containing second derivative information for all directions at each point in the field; we then apply a rotation $R_n : x \rightarrow n$:

$$H' = R_n^{-1} H R_n$$

Where R_n is a rotation matrix that rotates the x-axis of the co-ordinate frame onto the orientation of the axon (n), defined by the angles between the x-axis and the axon of interest. We can then extract the first element of the rotated Hessian, giving us the second partial derivative along the axon's orientation, i.e. the activating function, defined for a field:

$$(H')_{11} = \frac{\partial^2 V}{\partial n^2}$$

We demonstrate that this metric can be computed and that it provides a computationally efficient method of approximating activation for fields with axons of specified orientations. It accurately predicts the level of neural activation as a function of axon orientation, with high measures of this rotated Hessian metric indicating a low activation threshold (i.e. only a small amount of current required to induce an action potential), and low levels of this metric predicting a high threshold, where a large current needs to be used to produce an action potential.

This provides a computationally efficient means of linking the electric field produced by stimulation, V , to the activation of specific fibre tracts, i.e. $f_{V \rightarrow A}$. This allows us to investigate the physiological effects of stimulation in detailed anatomy in a meaningful way while bypassing the computationally expensive steps involved in biophysical modelling of axons by taking advantage of the geometry of the electric field. While detailed biophysical models remain a gold standard in modelling of the physiological effects of stimulation, this provides a simple to compute means of rapidly assessing the effects of stimulation in even very complex models.

3.2 Methods

We calculated the electric field produced by stimulation using a simple geometric model of surface stimulation of a peripheral nerve. Activation of axons of varying orientations was assessed using detailed biophysical models and the rotated Hessian activating field. These metrics were compared to evaluate the ability of the rotated Hessian to approximate neural activation. The correlation coefficient between these measures was used to assess the ability of the rotated Hessian to capture the orientation-dependence of biophysical activation thresholds.

3.2.1 Electric field computation

A three-layer tissue model (skin, subcutaneous fat, muscle) was created. A bipolar electrode arrangement was simulated on the skin surface, with 5cm x 5cm electrodes separated by 5cm (figure 3.1). The physical properties, derived from a standard database [94], applied to each tissue type are shown in table 3.1.

Table 3.1: Physical properties of tissue types

Tissue type	Conductivity (S/m)
Skin	$2E^{-4}$
Fat	$3.77E^{-2}$
Soft tissue	$2.02E^{-1}$
Nerve	$1.71E-2$

The domain is made up of a thin skin surface overlying a 5mm layer of fat. Beneath this is a large domain of homogeneous muscle tissue (figure 3.1A). Within this tissue, a nerve volume is inserted 5cm below the surface with its midpoint halfway between the active and return electrodes (figure 3.1B). The nerve volume is rotated through 360 degrees with its midpoint remaining fixed beneath the electrodes (figure 3.1C) in order to evaluate the effects of changing the orientation of the nerve relative to the electrodes on the response to the generated electric field (figure 3.1D).

The electric field produced by stimulation was computed by solving the Laplace equation with variable coefficients with homogeneous Neumann boundary conditions on the external skin boundaries using the finite element method with COMSOL Multiphysics [95]:

$$\begin{aligned}\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega\end{aligned}$$

where V is the electric potential generated by stimulation, σ is the tissue conductance, Ω is the domain modelled and $\partial\Omega$ is the external (skin) boundary of the domain.

Additional Dirichlet boundary conditions were added to simulate the effect of stimulation through the simulated electrodes:

$$\begin{aligned}V &= V_e && \text{on } \Gamma_A, \\ V &= 0 && \text{on } \Gamma_R\end{aligned}$$

Where Γ_A and Γ_R are the surfaces of the active and return electrodes respectively. and V_e is the voltage applied through the active electrode. The applied voltage V_e is computed such that $\int_{\Gamma_A} \frac{\partial V}{\partial n} = J$, where J is the target current to deliver.

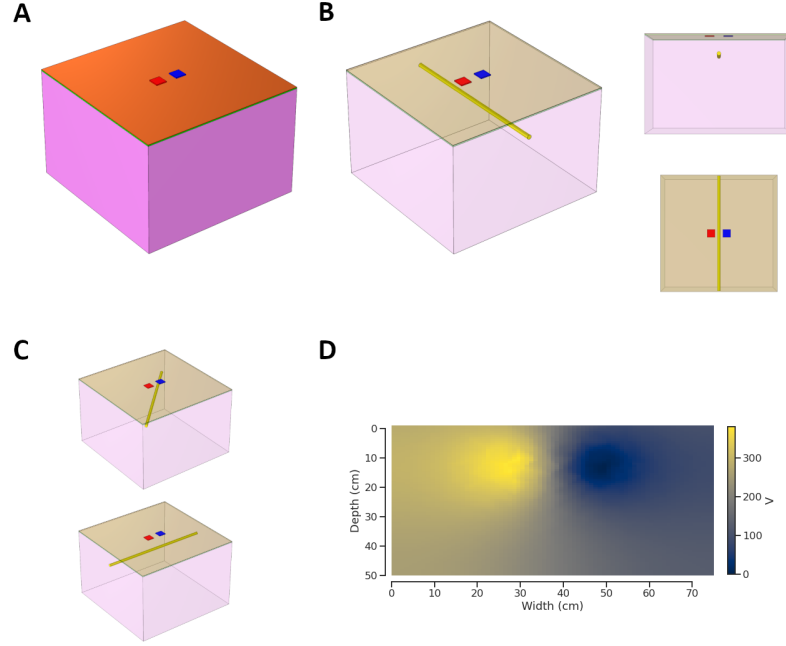


Figure 3.1: Geometry and sample electric field. **A** Simple three-layer model of tissue used for assessing activation metrics. Active (red) and return (blue) electrodes overlie the skin (orange) surface. Beneath this is a 5mm uniform layer of fat (green). Below this is a large uniform layer of muscle (pink). **B** Insertion of simulated nerves. A nerve volume, into which biophysical axon models are mapped, is inserted below the skin surface. This is centred below the electrodes. **C** Altering axon orientation. The target nerve volume is rotated, keeping the midpoint fixed under the electrodes. This allows evaluation of the effect of orientation on the response to stimulation. **D** A 2D slice through the computed potential field showing the potential produced by stimulation.

3.2.2 Biophysical model

The effects of the computed electric field on an axon at 5cm depth under the electrodes was assessed by coupling the results of the finite element simulation to a biophysical model of a myelinated axon with Hodgkin-Huxley dynamics [74]. The orientation of the axon was varied in order to assess the effect of axon orientation. Target axons were mapped into the nerve volume of the tissue model, with a total length of 10cm and an orientation following that of the nerve volume.

The McIntyre-Richardson-Grill model of a myelinated axon was used to simulate stimulation of an $11.5\mu\text{m}$ axon of a peripheral nerve using the default parameters defined by McIntyre et al. [76] All biophysical models were implemented in NEURON [92]. Extracellular fields were coupled using the cable equation with the extracellular

field defined at each point of the axon. Simulations were run and whether the model generated an action potential was measured.

Thresholds for the amplitude required for activation were computed using a binary search algorithm. Stimulation amplitude was altered and simulations re-run until activation thresholds were identified to within 1mA.

3.2.3 Computation of activation

The Hessian matrix of the electric potential produced by stimulation was approximated numerically by computing the second differences of the field at each point over a regularly spaced grid sampled every 0.05mm, producing a 3x3 matrix of second partial derivative information at each point in the field.

A rotation matrix was computed to rotate the x-axis of the original co-ordinate frame onto the trajectory of the axon of interest. This rotation matrix was then applied to the Hessian field.

The activating function for an axon with the given orientation was extracted from the Hessian matrix at each point in the field by taking out its first element, corresponding to the second partial derivative along the orientation of the axon, i.e. the activating function. This was performed for every point in the field, providing a metric over the whole field.

This approximation was then compared to the thresholds required for stimulation for axons with a range of orientations in order to evaluate its ability to approximate neural activation using a simple metric while maintaining orientation information.

3.3 Results

The geometric tissue model used is shown in figure 3.1. The potential field in the tissue model due to stimulation can be computed and visualized. These results were then used to compute the activation metrics assessed, demonstrating that the rotated Hessian approximation provides a computationally efficient, accurate method of assessing the effect of stimulation on neural tissue.

3.3.1 Orientation sensitivity

The rotated Hessian approximation can be computed for axons of arbitrary orientations, providing a metric defined everywhere on the domain modelled of the activation of an axon with that orientation at that position.

Activating fields for a selection of axons of differing orientations are shown in figure 3.2. The rotated Hessian approximation is able to capture the orientation selectivity of axons and to provide straightforward visualizations of the activating fields produced by stimulation.




Figure 3.2: Rotated Hessian approximations for axons of varying orientations; slice taken through centre of electrodes. Yellow indicates high activation, blue indicates low activation. 90° to the x- axis, 45° to the x-axis and 0° to the x-axis are shown. The pattern of activation differs for axons of differing orientations. This highlights the importance of using metrics for approximating activation the retain orientation information.

3.3.2 Performance comparison

The rotated Hessian approximation shows a similar orientation selectivity to the biophysical models (figure 3.3). The amplitude required to induce an action potential

in an axon model changes as a function of the orientation of the axon with respect to the electrodes (figure 3.4). The rotated Hessian approximation captures this rotation and provides a computationally efficient means of evaluating the response to stimulation for axons with a specific orientation.




Figure 3.3: Polar plot of amplitude threshold and rotated Hessian approximation for a range of orientations. Angles represent the orientation of the axon of interest. The radius indicates the normalized magnitude of the activation metric for an axon of that orientation, i.e. the amplitude required to activate it and the rotated Hessian metric. The rotated Hessian successfully captures the orientation selectivity seen in the biophysical model. High rotated Hessian metrics occur where the amplitude threshold is lower.

Where the electric field is oriented along the primary axis of the axon, little current is required to induce an action potential. Its threshold is therefore low, and its activation metrics are high, indicating a high likelihood of firing in response to stimulation. When the electric field is not aligned with the trajectory of the axon, it requires relatively more current to activate it. Its threshold is higher. The rotated Hessian metric captures this orientation-dependence.

The amplitude required for producing an action potential in an axon model and the rotated Hessian approximation at the location of the axon are highly




Figure 3.4: Amplitude threshold and rotated Hessian approximation shown over a range of axon orientations; the rotated Hessian captures the behaviour of the biophysical model, with high rotated Hessian metrics predicting low amplitude thresholds.

correlated ($r = -0.92$; $p < 0.0001$), as shown in figure 3.5. The rotated Hessian provides an accurate means of approximating the orientation-specific neural response to electrical stimulation.




Figure 3.5: The amplitude threshold and the rotated Hessian metric are highly correlated over a range of axon orientations ($r = -0.92$; $p < 0.0001$).

3.4 Discussion

These results show that it is possible to rapidly compute a metric that predicts the neural response to electrical stimulation that can be generalized to a field and that accounts for the orientation selectivity of axons.

This method is, however, only applicable to fibres with a constant trajectory, i.e. regions where the fibres are locally straight. This is useful for cases such as the spinal cord or peripheral nerves or central tracts outside of local changes of orientation. Accounting for activation with locally complex morphology is more challenging – requiring, for example, projection of the field onto the trajectory of interest. These methods, while useful, lend themselves less well to generalization to a field.

Biophysical measures are useful and highly accurate, but are slow to compute and are not easily applied to the field produced by an extracellular electrode in order to evaluate the regions where activation is likely to occur. The metric presented here provides a simple means to generalize a classic approximation of the effect of stimulation on neural tissue to a field, while retaining information on orientation.

Approaches to calculating the volume of tissue activated in response to deep brain stimulation have provided a number of methods of computing the regions likely to be activated by stimulation on a field [77]. These include very simple metrics like the electric field norm, methods based on Rattay’s activating function and complex methods based on detailed biophysical models. However, these methods are largely invariant to axon orientation. This provides a poor approximation for scenarios where local, directional axons are targeted, such as in the case of peripheral nerve stimulation. The orientation of axons with respect to the electric field is an important consideration, and any simplified metrics should maintain this relationship as much as possible.

The present work demonstrates a method derived from generalizing Rattay’s activating function to a field. By using rotations of the Hessian matrix of the electric potential produced by stimulation, it is possible to compute the activation for axons of arbitrary orientation and to generalize this to the entire domain modelled. This

allows rapid computation and simple visualizations of the regions of activation for axons of specific orientations.

Higher values of the activating function metric are related to a higher probability of generating an action potential. Positive values are related to a net "flow" of current into a given point, while negative values are related to a net "flow" of current away. However, it is difficult to determine an absolute threshold for activation of an axon using this metric alone - rather, it provides a means for predicting where action potential initiation is likely to occur by comparing relative values across conditions and locations.

This provides a useful, computationally efficient method for approximating activation for the development of novel neural interfacing techniques and the optimization of parameters for neuromodulation therapies.

In the context of developing an integrated system for targeted neuromodulation, this provides us with a simplified metric $f_{V \rightarrow A}$, which links the extracellular field produced by stimulation to the physiological effects of stimulation. This allows for the activation produced by stimulation to be approximated in a simple way, without the need for the intermediate steps of coupling and solving detailed biophysical models. This comes at a cost in terms of the accuracy of models using this simplified metric, but provides a valuable method for conducting work examining the effects of varying patterns of stimulation, with potential applications to the optimisation of stimulation parameters and the development of novel neuromodulation therapies. The development of truly reliable models for designing and optimising therapies will require methods for deriving individualised computational models, computing the extracellular field produced by stimulation and either coupling this to biophysical models or using simplified metrics such as the one outlined here.

4

Patient-specific computational models for simulating interventions

Contents

4.1	Introduction	55
4.2	Methods	58
4.2.1	Strategy	58
4.2.2	Implementation	59
4.2.3	Application	63
4.3	Results	66
4.3.1	Generation of patient-specific models from imaging . . .	66
4.3.2	Simulation of interventions	67
4.3.3	Noninvasive spinal cord stimulation	69
4.3.4	Targeting interferential stimulation	72
4.4	Discussion	75
4.4.1	Targeted noninvasive spinal cord stimulation	75
4.4.2	A generalisable system for patient-specific modeling . .	77
4.4.3	Precision neuromodulation	78
4.4.4	Conclusion	79

4.1 Introduction

Stimulation of the spinal cord using implanted electrodes has been well established in the management of pain [96]. Following this paradigm, implanted stimulation systems have shown promise in functional restoration after spinal cord injury [7–9].

These implanted systems carry several important limitations. These include the small risks of bleeding or infection associated with the procedure, and longer-term risks associated with device migration; as many as one-third of patients with an implanted system experience an adverse event, which in a small number of cases can cause worsening of neurological status [97], with up to 38% requiring revision in some series [98].

The ability to achieve spinal cord stimulation noninvasively would potentially allow for reduction of these risks. Even used in conjunction with implanted systems, the ability to reliably stimulate noninvasively would open the possibility of carrying out trial periods noninvasively, reducing risks and improving the efficiency of existing spinal cord stimulation care pathways for pain, while widening access to these treatments. Noninvasive methods would also make emerging indications, such as rehabilitation, more viable and easily scalable, making these technologies available to more patients.

Despite its potential clinical value, it is not clear that it is possible to achieve meaningful stimulation of the spinal cord noninvasively. There are reports of series using noninvasive stimulation in the context of rehabilitation after spinal cord injury [40], including improvements in the function of the lower limbs [41], upper limbs [99] and urinary tract [42], as well as temporary improvements in spasticity [100]. Recent systematic reviews of these series further suggest that there may be a functional impact of noninvasive stimulation in rehabilitation [43].

However, the combination of stimulation with other rehabilitation interventions without any control groups, and the complex outcome measures made up of a statistical construct of a large number of functional outcome measures, render interpretation difficult. Further, there are large variations in the electrode montages and stimulation parameters used, making it unclear whether an electric field within the canal is consistently being produced, and what the essential components of a stimulation configuration are for achieving improved outcomes. It remains unclear whether it is possible to reliably generate an electric potential within the spinal canal noninvasively. Emerging techniques such as interferential stimulation [101],

using multiple interfering electric fields applied simultaneously, offer a potential means to achieve greater penetration and targeting, but it remains equally unclear whether these can achieve sufficient depth in humans, and whether meaningful targeting is possible.

Computational modeling offers a promising means of answering these questions around the electric field generated by stimulation. The physics underlying the behaviour of electric fields are well characterised, and sophisticated methods for solving the associated equations numerically have been developed [95], allowing even very complex systems to be assessed.

Some modeling work on noninvasive stimulation has suggested it may be possible to stimulate the spinal cord [45]. However, these approaches used highly idealised geometries, which were not representative of true anatomy. Recent work on implanted stimulators has highlighted the importance of individual variations in anatomy on the effects produced [49], while work in deep brain stimulation has shown that even small variations in individual anatomy can produce important differences in the stimulation parameters required [50]. These effects are likely to be even more pronounced with noninvasive stimulation. The need for precise, accurate assessments of the electric fields produced by stimulation is further underlined by the importance of spatiotemporally precise stimulation patterns seen in applications such as rehabilitation [9].

To definitively evaluate the feasibility of noninvasive spinal cord stimulation, computational models are required that take account of explicit, detailed anatomy. Methods are required to create such realistic models with complex electrode montages and stimulation parameters in order to simulate their effects on a realistic domain. An approach using such detailed modeling would allow us to more rigorously address the question of whether it is possible to generate an electric field within the spinal canal in real anatomy, and whether meaningful targeting of stimulation is possible noninvasively. This would allow for a thorough, accurate assessment of the feasibility of targeted noninvasive spinal cord stimulation.

We aimed to develop a system for determining the potential field produced by simulated placement of electrodes on patient-specific models with highly detailed anatomy. We demonstrate a method for deriving models directly from routine clinical imaging, allowing for simulation on realistic models of individual anatomy. This allows for a realistic investigation into the feasibility of noninvasive spinal cord stimulation. We further investigate whether the developed methods allow for targeting of stimulation, potentially allowing for development of systems for individually optimised, targeted noninvasive stimulation.

In the context of the development of a system for model-based neuromodulation, these methods provide us with the ability to generate a patient-specific domain Ω over which to solve for the potential field produced by stimulation given a set of desired electrodes, locations and stimulation parameters θ , i.e. $f_{\theta \rightarrow \Omega}$. This can then be combined with methods for solving for the potential field over these domains to produce patient-specific models of the effects of neuromodulation, allowing for the effects of specific electrode geometries and positions to be assessed in detailed individual anatomy.

4.2 Methods

4.2.1 Strategy

To investigate the effect of stimulation using patient-specific models, computational models with simulated electrode placements were created directly from clinical imaging and solved to determine the potential fields produced.

Electrode geometries were mapped into the image volume to simulate intervention. A tetrahedral mesh, suitable for physics simulation using the finite element method, was then formed from the processed imaging. Physical properties and stimulation parameters were then applied to the model, and the resulting potential field solved, providing a realistic simulation of the potential field produced by a given set of stimulation parameters in patient-specific anatomy.

The developed method was then applied to the case of non-invasive spinal cord stimulation. A series of models were tested to assess the ability to simulate complex electrode arrangements and stimulation parameters in patient-specific models.

4.2.2 Implementation

These methods were implemented in a purpose-built software package (Figure 4.1).

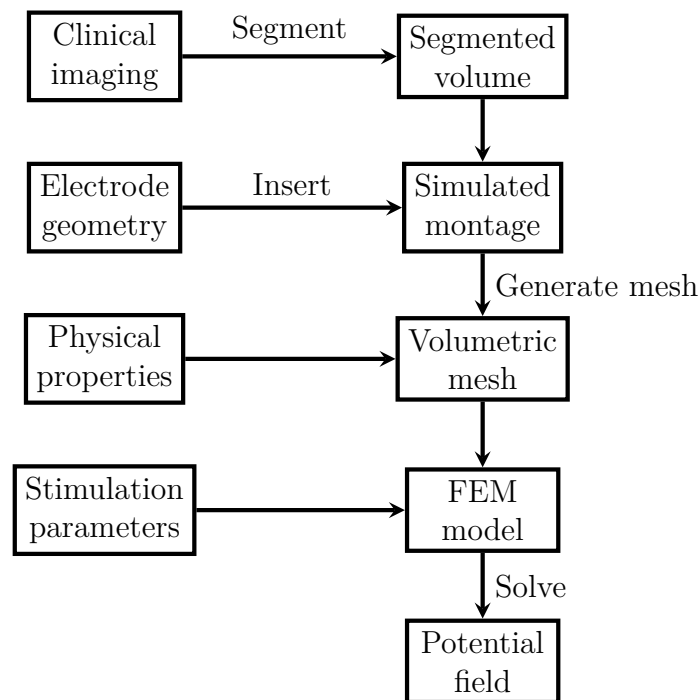


Figure 4.1: Flowchart of implemented method. Arbitrary electrode placements are simulated using clinical imaging, and the resulting potential fields solved. This allows for the testing of the physical effects of bioelectronic interventions in patient-specific models.

Image segmentation

Clinical imaging was segmented into discrete tissue types using 3D Slicer segmentation software [102].

Semi-automated tools were used to convert the three-dimensional image volume into a volume of tissue IDs. Segmentations were then inspected and manually corrected where necessary. Custom software tools were then used to process the segmented volume.

The original image volume was processed into a segmented volume, M , with the same dimensions as the original volume, made up of n segments, K_i , such that these segments do not overlap, i.e. each voxel is assigned a single tissue ID, and every voxel is assigned an ID, i.e. there are no unassigned voxels:

$$\text{int}(K_i) \cap \text{int}(K_j) = \emptyset,$$

$$\cup_i K_i = M$$

In this application, images were segmented into bone, soft tissue, lung, skin, cerebrospinal fluid and spinal cord as outlined below. Overlapping segments were avoided by ensuring every voxel is assigned a single tissue ID; in the case of overlaps between segments within a voxel, the order of priority of tissues was spinal cord, cerebrospinal fluid, skin, bone, lung, soft tissue. This ensured all voxels were assigned and no segments overlapped while ensuring that the morphology of the most functionally important tissues were maintained.

The end result of the segmentation process was a three-dimensional volume of the same dimensions as the original image volume, where each voxel is assigned a single scalar tissue ID. The anatomy is unchanged, while the image is decomposed into n segments based on tissue type, where n is the number of tissue types used in the model.

This allows for explicit representation of the specific anatomy of differing tissue types within the image volume.

Simulating intervention

Electrode placement was simulated by mapping electrode geometries into the 3D image space.

Geometries were defined using a simple image editor to outline the 2D footprint of each electrode. These geometries were then mapped into the image volume. The geometry footprint was rescaled such that the aspect ratio and size of the image pixels defining the geometry matched the aspect ratio and size of the image voxels in the target plane. The centrepoint of the electrode geometry was then mapped onto the 2D coordinates of its desired position on the tissue type of interest.

The electrode geometry was then projected onto the surface of the tissue of interest such that it is contact with it at every point. During projection, the geodesic distance over the anatomic surface of projection is used to adjust the electrode dimensions, allowing for projection onto non-planar anatomic structures without warping the dimensions of the electrode. This is achieved by rescaling the electrode geometry such that the ℓ_1 norm from the geometry centrepoint along the projection surface in each direction is equal to the desired length of the geometry in that direction:

$$||d||_1 = |x_n - x_0| + |y_n - y_0|$$

where d is the desired electrode size from the centrepoint, (x_0, y_0) are the coordinates of the centrepoint on the projection surface and (x_n, y_n) are the coordinates of the n^{th} step along the surface.

The electrode geometry was then extruded to a defined thickness to give it a 3D volume within the image, in contact with the defined surface with its centrepoint at the defined position. This allowed for arbitrary electrode geometries to be accurately placed within the image volume in order to simulate the effects of any desired geometry or montage. This allows this approach to be generalised to applications involving implanted electrodes, although only surface electrodes were used in this application.

Mesh generation

A three-dimensional tetrahedral volume mesh, suitable for finite element simulation, was derived from the processed image volume.

The processed volume was discretised using a Delaunay triangulation algorithm, as implemented by the Computational Geometry Algorithm Library [103], via a wrapper built on the iso2mesh library [104].

During the mesh generation process, a mapping from the tetrahedral mesh elements to the voxelwise tissue IDs in the original volume was created, i.e.

$$f(E_i) = K_j$$

where E_i is a given mesh element and K_j is the corresponding tissue ID.

The raw volumetric mesh data was then further processed to associate this tissue ID with each discrete mesh element based on its location within the volume. This method allowed the generation of a volumetric mesh from the processed image, while maintaining the ability to identify the tissue type represented at any location within the volume. In this way, the underlying patient-specific anatomy remains explicitly represented in the volumetric mesh.

Finite element analysis

To simulate the effect of stimulation through the electrode mapped into the image, physics simulations were carried out on the generated mesh using finite element analysis.

Conductivity and permittivity values at the frequency of interest can be defined for each tissue type based on standard databases [94]. This allowed the definition of a mapping:

$$f(K_i) = \langle \sigma_j, \varepsilon_j \rangle$$

where K_i is a given tissue type and $\langle \sigma_j, \varepsilon_j \rangle$ is a tuple of the associated conductivity and permittivity values.

By using our mapping from mesh element to tissue type and from tissue type to physical properties, values could then be assigned to each individual mesh element based on its anatomic location:

$$f(E_i) = \langle \sigma_j, \varepsilon_j \rangle$$

where E_i is a mesh element and $\langle \sigma_j, \varepsilon_j \rangle$ are the associated physical properties. This allows the patient-specific anatomy and its corresponding physical properties to be explicitly represented within the physics simulation.

The potential field within the tissue in response to stimulation was then calculated by solving the Laplace equation for linear materials with Neumann boundary

conditions numerically using a quasistatic solver in COMSOL Multiphysics [95]:

$$\begin{aligned}\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega\end{aligned}$$

where V is the electric potential generated by stimulation, σ is conductance, Ω is the domain modeled and $\partial\Omega$ is the external surface of the region modeled (i.e. skin at limits of modeled region).

Note that as we are assuming electrostatic conditions, the tissue permittivity is not used; this becomes necessary if this approach is used to evaluate time-varying fields without this assumption.

Additional expressions are then added to specify the boundary conditions applied on the electrode surfaces:

$$\begin{aligned}V &= V_e && \text{on } \Gamma_A, \\ V &= 0 && \text{on } \Gamma_R\end{aligned}$$

where Γ_A is the boundary of the active electrode, V_e is the applied electric potential, computed such that the target stimulation current applied is the integral of the current density over the active electrode surface $\int_{\Gamma_A} \frac{\partial V}{\partial n} = J_{applied}$, and Γ_R is the boundary of the return electrode.

Simulations on imaging were implemented using custom software. This allowed for automated solving for the electric potential for any arbitrary electrode geometry, placement or stimulation parameters on a detailed patient-specific anatomic domain.

The physiological effects of stimulation were then approximated using the second spatial derivative of the electric potential along the orientation of the target fibres as described in chapter 3. This activating field approximation allowed the effects of stimulation on the neural elements to be straightforwardly visualised.

4.2.3 Application

The developed methods were applied to a series of increasingly complex models of noninvasive spinal cord stimulation in order to assess whether this approach can simulate the effects of complex stimulation modalities and answer the question

of whether these methods can produce an potential field within the spinal canal with properties that produce action potential generation.

A simple transabdominal electrode arrangement, a more complex dorsal electrode arrangement with bilateral return electrodes and two independent pairs of electrodes were evaluated, as shown in figure 4.2A. The interferential electrode arrangements were evaluated a range of electrode spacings, as shown in figure 4.2B.




Figure 4.2: Electrode montages. **A** Simple transabdominal electrode montage, more complex dorsal montage and two electrode pairs for interferentials stimulation. **B** Interferential electrode pairs were evaluated at a range of electrode spacings.

Transabdominal stimulation

A CT myelogram was segmented into bone, soft tissue, lung, skin, cerebrospinal fluid and spinal cord tissue segments and the associated physical properties applied as shown in table 4.1.

A 3.2cm diameter cathode was placed in the midline over the T11 vertebra. A 3.2cm diameter anode was placed at the same level in the midline on the anterior surface of the model.

A constant current of 30mA was applied via the cathode and the resulting electric field calculated to assess whether this simple arrangement could generate an electric field within the spinal canal.

Table 4.1: Physical properties of tissue types for constant current stimulation

Tissue type	Conductivity (S/m)
Bone	$2E^{-2}$
Soft tissue	$2.02E^{-1}$
Lung	$3.89E^{-2}$
Skin	$2E^{-4}$
CSF	2
Spinal cord (transverse)	$8.30E^{-2}$
Spinal cord (longitudinal)	0.6

Complex montage

A model with a more complex multi-electrode arrangement was set up using the same image and physical properties, with parameters based on Sayenko et al.'s work [40] to assess whether these parameters could feasibly produce an electric field within the spinal canal.

A 3.2cm diameter cathode was placed in the midline over the T11 vertebra. 7.5cm x 13cm anodes were placed over the iliac crests bilaterally. 150mA of current was applied via the cathode.

Interferential stimulation

A model was created using the physical properties shown in table 4.2.

Table 4.2: Physical properties of tissue types for 2kHz stimulation

Tissue type	Conductivity (S/m)
Bone	$2E^{-2}$
Soft tissue	$3.3E^{-1}$
Lung	$8.35E^{-2}$
Skin	$2E^{-4}$
CSF	2
Spinal cord (transverse)	$8.30E^{-2}$
Spinal cord (longitudinal)	0.6

3.2cm diameter cathodes were placed lateral to the midline bilaterally at the

upper border of the T11 vertebra. Corresponding 3.2cm diameter anodes were placed lateral to the midline bilaterally at the lower border of the vertebra.

30mA of current was applied to each anode. The first electrode pair was driven at 2kHz, and the second at 2.02kHz, producing an envelope amplitude modulation of 20Hz in the region of interference between the two fields. The electric fields produced by each electrode pair were solved separately.

The resulting overall electric field was then determined, as well as the spatial distribution of the envelope modulation amplitude, given by:

$$E_{AM}(n, r) = ||(E_1(r) + E_2(r)) \cdot n| - |(E_1(r) - E_2(r)) \cdot n||$$

where $E_1(r)$ and $E_2(r)$ are the electric fields generated by the first and second electrode pairs at position r , and n is a unit vector along the direction of interest [101].

The amplitude modulation of the activating function was also calculated in order to evaluate the regions in which axonal activation was likely to occur [105]:

$$AF_{AM} = ||AF_1(r) + AF_2(r)| - |AF_1(r) - AF_2(r)||$$

where $AF_1(r)$ and $AF_2(r)$ are the activating functions generated by the first and second electrode pairs at position r .

This allows for extension of the described methods to complex electrode montages and physics, including the targeting of interferential stimulation.

4.3 Results

4.3.1 Generation of patient-specific models from imaging

The developed methods allow for the generation of patient-specific computational models directly from clinical imaging. Segmentation allows clinical imaging to be decomposed into sections corresponding to specific tissue types; these segmentations can then be combined over the image volume to directly create 3D models of individual anatomy, as shown in Figure 4.3.

These models can be automatically generated from clinical imaging, allowing for straightforward creation of models containing detailed explicit representations

of individual anatomy. The 3D structures created can then be processed to automatically generate a volumetric mesh suitable for physics simulation, as shown in Figure 4.4. This allows for detailed models for physics simulation to be derived directly from clinical imaging.




Figure 4.3: Image segmentation. **A** Example of an axial slice of a CT myelogram. **B** Segmentation of example slice into discrete tissue types; each colour represents one tissue type. **C** Surface mesh of segmented volume demonstrating 3D structure of segmented tissue types

4.3.2 Simulation of interventions

We can successfully create detailed 3D models that include simulated electrode placements. This allows for the effects of arbitrary electrode geometries and montages to be assessed and compared.

As shown in figure 4.4, even arbitrarily complex electrode geometries can be converted into a 3D volume. These can then be mapped into the image space at specified positions. Rescaling the electrode geometry according to image resolution

and aspect ratio and accounting for warping during projection onto anatomic surfaces allows electrode placement to be realistically simulated, with electrodes in contact with anatomic surfaces throughout. This allows for any electrode geometry or position to be simulated.




Figure 4.4: Electrode insertion. **A** Example of an arbitrarily complex electrode geometry, defined as a 2D image (left) and converted converted into a 3D structure for simulation (right). **B** Axial slice of segmented volume demonstrating insertion of electrode geometry into image space; electrodes (white) visible in axial section anteriorly and posteriorly on the midline. **C** 3D volume mesh derived from image with electrode inserted; muscle and soft tissue removed for visualisation. Demonstrates ability to insert arbitrary electrode geometries at specified locations in image volume and simulate their effects. Note figure shows a simple round electrode rather than a complex geometry.

As shown in figure 4.4C, a volumetric mesh for physics simulation can then be generated from the combined volume. This allows for assessment of the effects of stimulation through the applied electrodes. In this way, the effects of arbitrary interventions can be assessed using a patient-specific model built with pre-intervention imaging, as shown in figure 4.5.

This shows that we are able to carry out physics simulations on patient-specific models with simulated interventions. Figure 4.5 further shows that this method can be readily extended to complex systems with multiple electrodes and to systems with complex, interacting, time-varying electric fields.

Using these techniques, we were then able to assess the effects of variations in electrode montage and stimulation parameters on noninvasive stimulation of the spinal cord.

The arrangement of electrodes had a clear impact on the electric field generated (Figure 4.5). A transabdominal arrangement, with 3.2cm diameter round electrodes in the midline at T11 on the anterior and posterior skin surfaces, generates a potential across the body, with an electric field oriented orthogonal to the spinal cord. A more complex multi-electrode arrangement with a single anode and bilateral cathodes on the posterior wall produces a more localised distribution of potential, as well as a field with a different orientation. The combination of two interfering fields, as 2kHz and 2.02kHz, applied along the long axis of the spinal cord lateral to the midline bilaterally also produces a more focused potential distribution, and an overall electric field oriented along the point of overlap of the two independent fields.

Each of these methods demonstrates the ability generate a potential within the spinal canal. While the orientation of the electric fields produced varies, the natural bony canals formed by the intervertebral foramina and the spinal canal itself tend to direct the electric field along the long axis of the spinal cord. This ability to generate a potential within the canal and the orientation of the resulting field along the axis of the spinal cord suggests the possibility of generating a gradient along the cord, which could produce activation noninvasively.

4.3.3 Noninvasive spinal cord stimulation

A more detailed analysis of the effects of transabdominal 30mA DC stimulation is shown in figure 4.6. Figure 4.6A shows the same view as figure 4.5, with 3.2cm diameter electrodes across the abdomen and streamlines representing the field orientation. Here the vertebrae have been removed for visualisation and the spinal cord is shown in its anatomic position, coloured by the electric potential. The spinal cord, coloured by the electric potential, is also shown in isolation for visualisation.

Simple transabdominal stimulation with midline electrodes produces a potential gradient along the spinal cord. The magnitude of the potential gradient produced

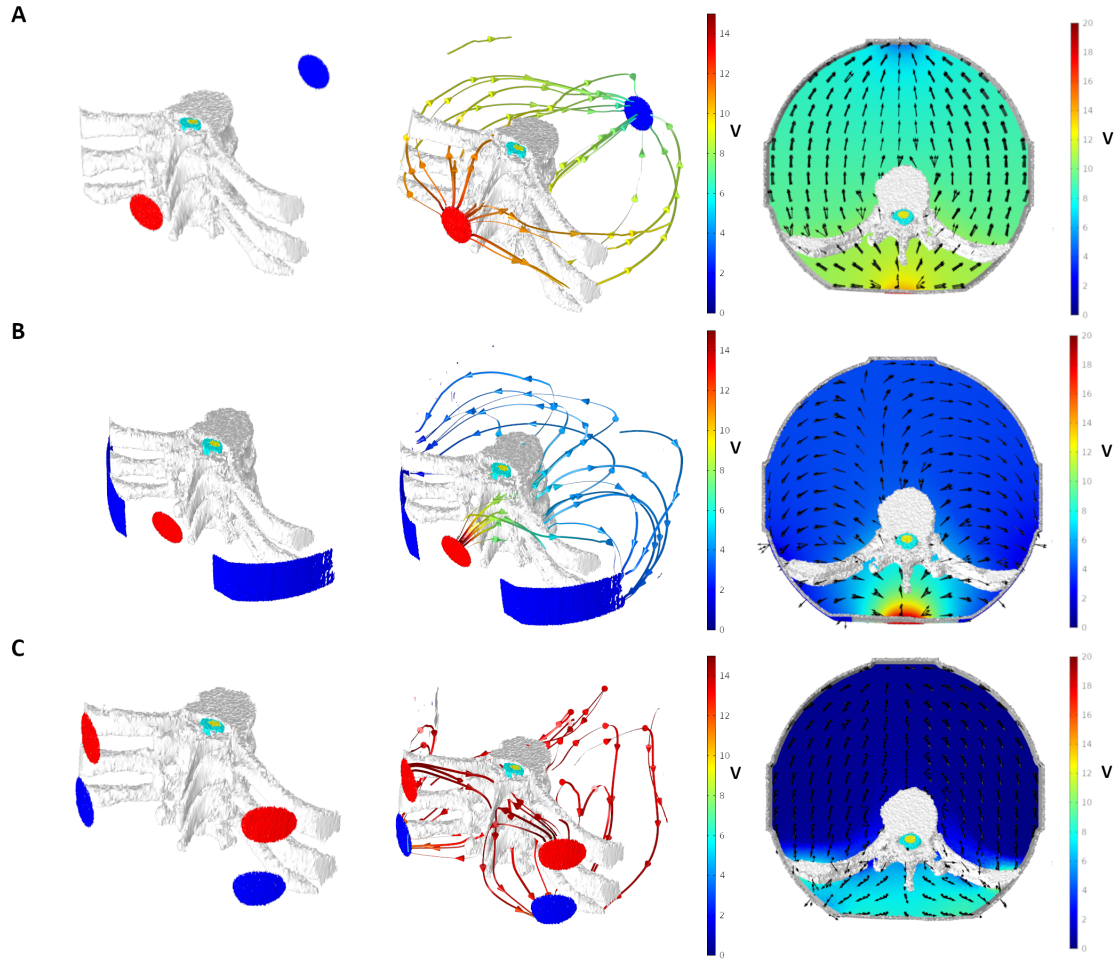


Figure 4.5: Electric fields produced by a series of electrode montages. Demonstrates ability to solve for the effects of complex electrode configurations and stimulation parameters. Streamlines in upper figures represent electric field orientation, coloured by potential (max 15V). Arrows in lower figures show electric field orientation; colourmap represents potential (max 15V). **A** 3.2cm diameter cathode over midline at T11, 3.2cm diameter anode anteriorly in midline at same level; 30mA of DC current applied. **B** 3.2cm diameter anode over midline at T11, 7.5cm x 13cm anodes bilaterally above iliac crests; 150mA of DC current applied. **C** Electrode pairs placed bilaterally for interferential stimulation; 30mA applied to each pair, at 2kHz and 2.02kHz.

along the surface of the spinal cord is on the order of 0.2V/m. The electric field appears to be directed through the natural bony foramina at the posterior aspect of the vertebral column, and then down along the spinal canal; this produces a potential gradient along the long axis of the spinal cord.

Figure 4.6B shows the same view, with the spinal cord coloured by a generalisation of Rattay's activating function [89] to a field, i.e. $\frac{\partial^2 V}{\partial n^2}$, where n is a unit vector along the principal axis of the axon of interest, as described in chapter 3.

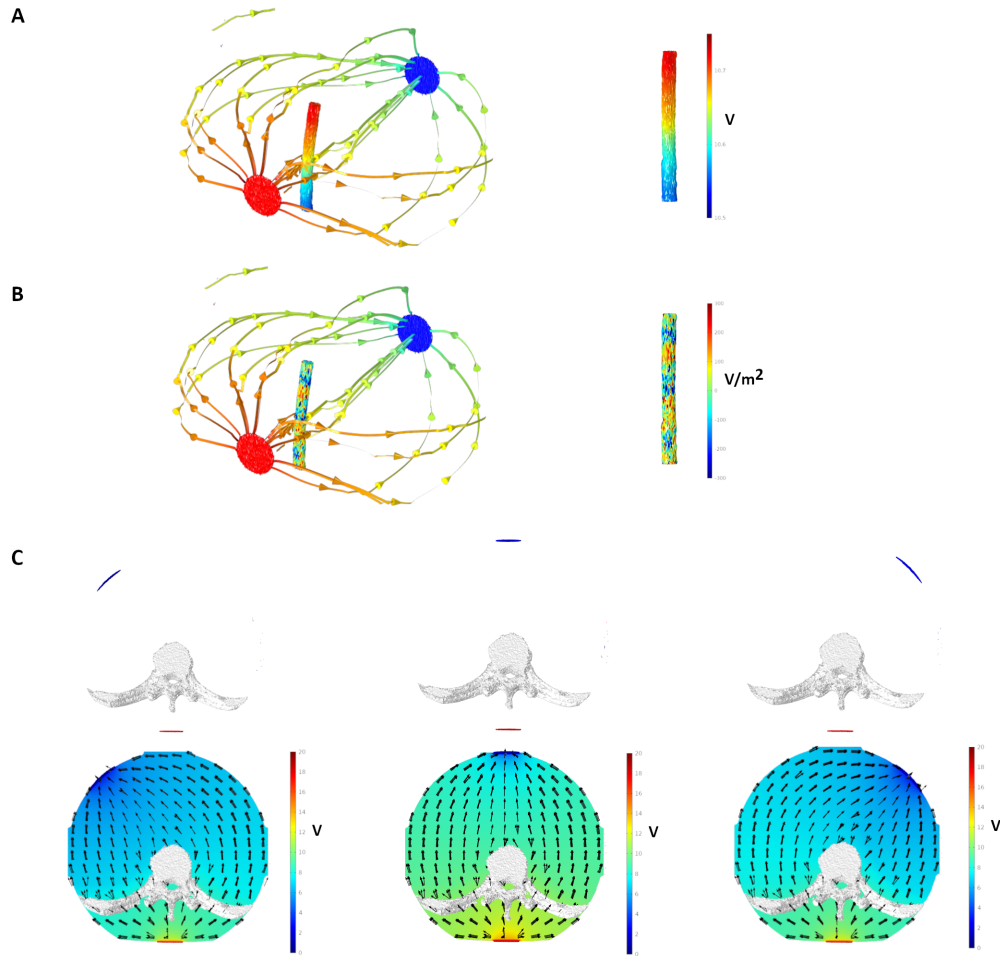


Figure 4.6: Effect of transabdominal 30mA DC stimulation. **A** Projection of the electric field generated; cathode (posterior midline) shown in red, anode (anterior midline) in blue. Streamlines represent the orientation of the electric field, coloured by potential (max 15V). Spinal cord is shown within the field streamlines in its anatomic position, coloured by potential. On the right the surface of the spinal cord is shown in isolation; the geometry shown generates a 0.2V/m gradient along the long axis of the spinal cord. **B** Projection of the electric field with spinal cord in situ, coloured according to activating function ($\frac{\partial^2 V}{\partial n^2}$, where n is the long axis of the spinal cord). On the right, the surface of the spinal cord is shown in isolation; the geometry shown generates segmental activation where the electric field is directed along the spinal canal. **C** Demonstration of solving for arbitrary electrode montages; solutions shown for anodes placed on the left, midline and right.

This provides a straightforward approximation for areas likely to be activated by stimulation. We see here areas of segmental activation (high activating function) corresponding to entry and exit points of the field along the canal. This shows that the gradient produced by simple transabdominal stimulation can produce a

pattern of segmental activation of the spinal cord based on the level at which stimulation is applied.

Figure 4.6C shows a series of alternative transabdominal electrode arrangements. By altering the locations of the electrodes, we can "shape" the resulting electric field to better direct stimulation. This demonstrates the ability to easily simulate arbitrary montages and assess their effects, and also shows that careful assessment of electrode positions may help to shape the resulting field to maximise stimulation.

Figure 4.7 shows the effect of a complex multi-electrode arrangement using a montage previously described for transcutaneous spinal cord stimulation [40] (figure 4.7A), with a 3.2cm diameter midline anode and bilateral cathodes on the posterior skin surface.

Figure 4.7B shows the same view as figure 4.5B, with vertebrae removed for visualisation. The spinal cord is shown, coloured by potential. The spinal cord is also shown in isolation, demonstrating the generating of a potential gradient along its surface. The gradient generated by this multi-electrode arrangement is on the order of 1V/m. This again appears to correspond to redirection of the electric field along the spinal canal.

Figure 4.7C shows the same view, with the spinal cord coloured by activating function. This again shows a pattern of segmental activation, suggesting that this electrode arrangement may be able to produce activation of the underlying spinal cord.

4.3.4 Targeting interferential stimulation

The same techniques can be extended to the case of multiple interfering electric fields. Figure 4.5C shows the creation of a model with pairs of electrodes aligned superior-inferior along the long axis of the spinal cord, parallel and lateral to the midline bilaterally. This shows the orientation of the resulting electric field, with orientation of the overall electric field along the point of interference, producing modulation of the envelope amplitude along the axis of the spinal canal.

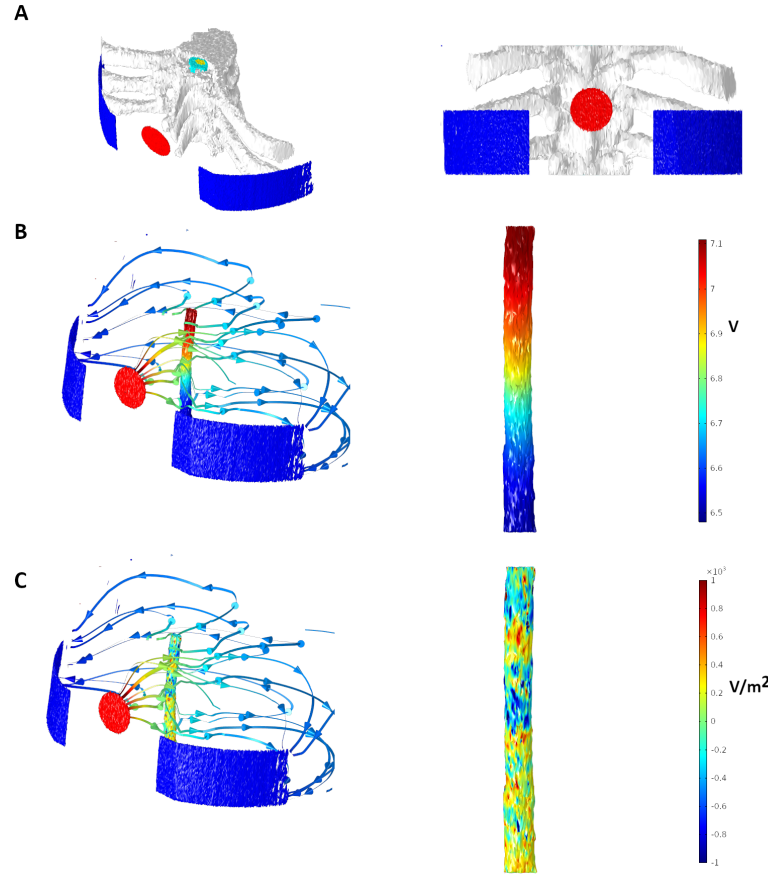


Figure 4.7: Effect of stimulation with parameters used by Sayenko et al. [40], i.e. 150mA through a midline anode with cathodes interolaterally on each side. **A** Projection of the electric field generated; cathode (posterior midline) shown in red, anodes (inferolateral) in blue. **B** Streamlines represent the orientation of the electric field, coloured by potential (max 15V). Spinal cord is shown within the field streamlines in its anatomic position, coloured by potential. On the right the surface of the spinal cord is shown in isolation; the geometry shown generates a 1V/m gradient along the long axis of the spinal cord. **C** Projection of the electric field with spinal cord in situ, coloured according to activating function ($\frac{\partial^2 V}{\partial n^2}$, where n is the long axis of the spinal cord). On the right, the surface of the spinal cord is shown in isolation; the geometry shown generates segmental activation where the electric field is directed along the spinal canal.

Figures 4.8 and 4.9 shows that we can modulate the depth and location of the focus of interference. These plots visualise E_{AM} , the modulation of the envelope amplitude, i.e. the area of interference, where stimulation is likely to occur.

Figure 4.8 shows the effect of altering the distance between the electrode pairs. As the pairs are moved further from the midline (figure 4.8A), the focus of interference becomes deeper within the tissue (figure 4.8B). The focus of activation also becomes deeper, with less cutaneous activation while still maintaining a high activating

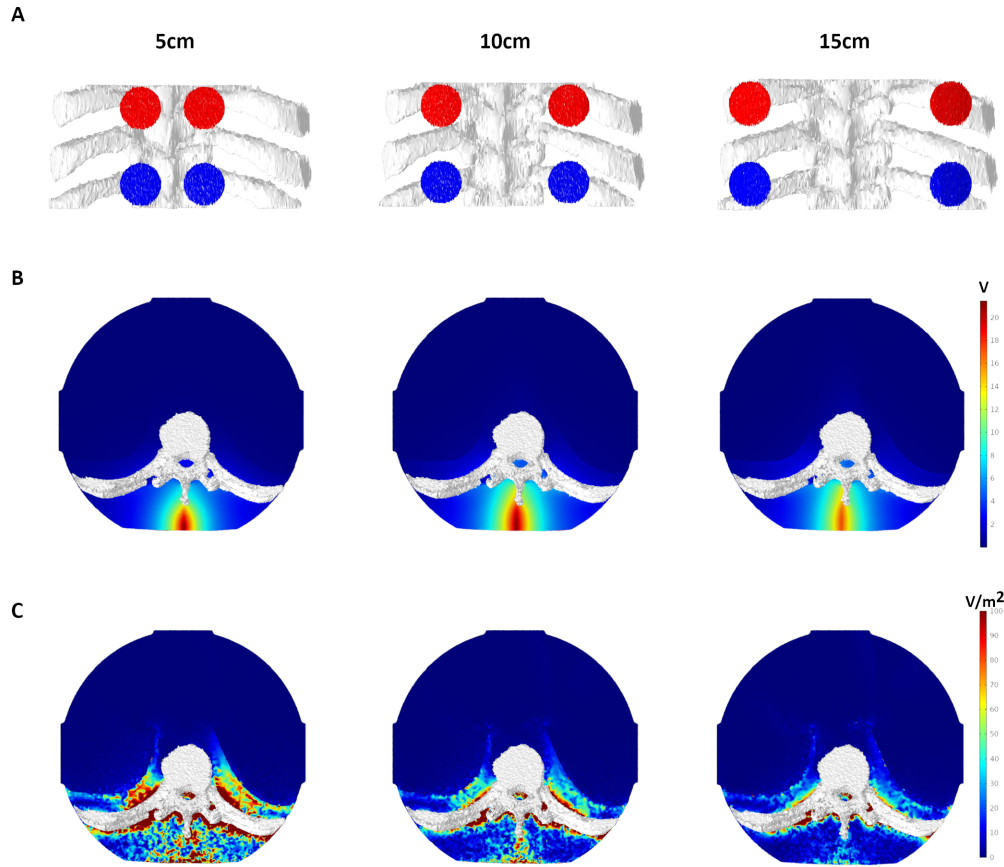


Figure 4.8: Control of depth and steering with interfering fields. Effects of varying positions and parameters on an index of field amplitude modulation shown. **A** Moving electrode pairs further from the midline moves modulates depth of maximum intensity. **B** Altering the ratio of stimulation amplitude between electrode pairs allows steering of the focus of maximum intensity with fixed electrode positions.

function at depth (figure 4.8C).

Figure 4.9 shows the effect of varying the ratio of the stimulation amplitudes between the electrode pairs with the electrode positions fixed. By changing the stimulation parameters, we are able to steer the focus of interference from left to right. Figure 4.9A shows that as the ratio of stimulation between two electrode pairs at a fixed distance is varied, the focus of interference of the electric field is steered from left to right. In line with this, figure 4.9B shows that the focus of activation can also be steered using the same approach. In this way, activation of




Figure 4.9: Control of depth and steering with interfering fields. Effects of varying positions and parameters on an index of field amplitude modulation shown. **A** Moving electrode pairs further from the midline moves modulates depth of maximum intensity. **B** Altering the ratio of stimulation amplitude between electrode pairs allows steering of the focus of maximum intensity with fixed electrode positions.

different spatial targets can be achieved with electrodes in a fixed position.

4.4 Discussion

4.4.1 Targeted noninvasive spinal cord stimulation

The results of the simulations indicate that noninvasive spinal cord stimulation has the potential to produce an effect by generating an electric field within the spinal canal. Further, it is possible to target stimulation by altering the electrode geometries and stimulation parameters used to manipulate the electric field produced.

Noninvasive stimulation creates a potential within the spinal canal in both simple transabdominal stimulation and with more complex multi-electrode montages at stimulation amplitudes which are in line with those used in experimental studies, and which are likely to be within a tolerable range for practical use, though further testing to determine limitation due to activation of cutaneous skin afferents and associated discomfort is needed. This indicates that it is possible to produce an electrical effect on the spinal cord with purely noninvasive stimulation using acceptable parameters.

In all of the simulations, the bony anatomy of the vertebral column resulted in redirection of the electric field produced. The field was directed into the spinal canal through the intervertebral foraminae, and then channeled along the long axis of the spinal cord by the bony channel of the spinal canal. This resulted in the creation of a potential gradient along the surface of the spinal cord, which supports the possibility of generating a physiological response.

Neural activation is classically proportional to the second spatial derivative of the extracellular potential along the long axis of the axon of interest [89, 90]; the extracellular field acts as virtual sources along the extent of the axon where this activating function is high, inducing transmembrane current flow. The activating function shows segmental areas of high value where the field has been directed to produce a potential gradient along the surface of the spinal cord, suggesting that the channeling of the electric field through natural bony openings and along the canal produces a field that could plausibly produce activation of the spinal cord or the nerve roots.

Although it appears to be possible to generate a field with characteristics that would produce activation, the magnitude of activating function is relatively low with 30mA transabdominal stimulation. This may not be sufficient to consistently produce a physiological response. Higher amplitude stimulation with 150mA produced a much stronger gradient; however, stimulation at these kinds of amplitudes is unlikely to be well tolerated in patients with intact sensation.

While additional work, including biophysical modeling, will be required to more accurately assess the threshold required for producing activation, our results indicate that it is possible to generate an electric field with useful characteristics within the spinal canal. This provides significant support to the results demonstrated in rehabilitation series [40, 41, 99], and raises the possibility of noninvasive spinal cord stimulation in the clinical setting, with potential application to pain, as well as emerging indications such as spinal cord injury rehabilitation.

These results also suggest a potential strategy for stimulation: by designing electrode geometries and montages to better direct the resulting field into the

natural foraminae of the vertebral column and allowing the bone to redirect the field along the axis of the neural elements, it may be possible to optimise the electrode montage to allow effective stimulation at lower amplitudes.

Simulation allows for electrode montages to be designed to direct an electric field into foraminae and to produce activation at specific levels. Designing montages based on individual anatomy and assessment of the resulting field potentially allows for accurate targeted stimulation, overcoming issues with lack of proper targeting with undirected surface stimulation.

The use of interferential stimulation also relies on the ability to accurately simulate the electric fields produced. In this case, the interactions between two overlapping, interfering fields allows for control of the depth and position of the focus of the resulting field. We show that interferential stimulation, using electrode pairs oriented along the long axis of the spinal cord, can produce an electric field within the spinal canal, and the focus of activation of the field can be steered by altering the ratio of amplitudes of the electrode pairs.

Simulation on patient-specific anatomy allows for the assessment of the effect of any electrode arrangement, providing a means for determining the positions and amplitudes required to activate our target structures. This opens the possibility of true targeted noninvasive spinal cord stimulation, with individualised electrode montages and stimulation parameters based on imaging and simulation to target specific structures. The ability to steer the focus of the field potentially offers a level of spatiotemporal precision required for useful neuromodulation in precision-critical contexts such as rehabilitation [9]. The ability to achieve this noninvasively would greatly increase the access to these technologies by providing an easily accessible and low-risk means of augmenting spinal cord injury rehabilitation.

4.4.2 A generalisable system for patient-specific modeling

Beyond the present application to noninvasive spinal cord stimulation, the techniques developed have the potential to be applied to other regions and to patient-specific as-

assessments. This offers a valuable tool for the development of novel neuromodulation systems, as well as a potentially useful system for patient-specific therapies.

The ability to derive computational models directly from imaging allows these techniques to be easily applied to other anatomic regions, allowing for investigation of neuromodulation therapies throughout the body. The same process can also be applied to assess variations in stimulation modality, or even to assess other physics, such as fluid dynamics through segmentation of blood vessels. This potentially offers a generalisable system for modelling physical phenomena in detailed anatomic models.

Beyond the investigation of the feasibility or effect of physical interventions, the ability to easily conduct detailed physics simulations on anatomic models is potentially useful for developing novel stimulation systems. It is straightforward to assess the effects of changes to electrode geometry or simulation parameters in realistic models, allowing for rapid testing of feasibility, and for rapid iteration of designs based on simulations on real anatomy. This has the potential to offer a useful tool for accelerating translational work in neuromodulation.

4.4.3 Precision neuromodulation

The ability to carry out detailed simulations on individual anatomy using standard clinical imaging offers the future possibility of truly individualised precision stimulation. Using imaging to create a model and defining a target region, the ideal electric field characteristics can be determined to cause activation; it is then possible to iteratively determine the electrode geometry, position and stimulation parameters that will cause maximal activation in that patient. This sort of optimisation at the level of the individual patient allows for individual differences in anatomy to be taken into account, and potentially opens the possibility of genuine precision neuromodulation.

The current simulation methods extend to solving for the electric field produced by a given set of parameters. This allows derivation of metrics such as the activating function [89] or its extension to a three-dimensional field as described in chapter

3, an approximation for neural activation based on a cable approximation of the axon which links depolarisation and activation to transmembrane currents induced by virtual sources along the one-dimensional extent of the idealised axon. However, it does not allow us to directly assess the neurophysiological effects of the fields generated, and interpretation is rendered challenging by the lack of a clear threshold for activation.

By linking the present simulations to realistic biophysical axon models based on the Hodgkin-Huxley formalism [74], such as the McIntyre-Richardson-Grill model of myelinated nerves [76], it is possible to extend the present system to include automated biophysical measures. This would allow direct assessment of the effects of a given set of stimulation parameters on the underlying electrophysiology, providing more detailed, and clinically useful, information than is provided by the electric field alone.

The methods developed here allow the determination of electric field characteristics from electrode location and stimulation parameters. The ease of iterating simulations with arbitrary parameters offers the possibility of optimisation of stimulation. Using biophysically relevant metrics, such as the focus of interferential stimulation or biophysical measures of nervous system activation, optimisation algorithms could be applied iteratively to compute the optimal configuration for a given patient for a given target. This ability to provide individually optimised stimulation could then offer the potential for reliable stimulation by overcoming the effects of variations in anatomy.

4.4.4 Conclusion

The present results demonstrate the feasibility of noninvasive stimulation of the spinal cord, including via interferential stimulation, in simulation. Further, the methods developed offer a potentially rigorous means of assessing the effects of a bioelectronic effects on an individual patient using routine pre-intervention imaging. These techniques, if properly validated in physical models and in human testing,

could open the possibility of optimised neuromodulation using noninvasive stimulation for a range of existing and emerging indications.

In the context of the development of a broader system for model-based neuromodulation, this work provides us with a method for deriving a patient-specific model, Ω , with simulated placement of electrodes or other interventions specified by a set of parameters θ , i.e. $f_{\theta \rightarrow \Omega}$. By combining this with a method for numerically solving for the electric field produced by stimulation, this allows us to determine the physical effects of stimulation in detailed individual anatomy for any given set of stimulation parameters or electrode positions.

This method can then be extended by the use of either simplified metrics to link the electric field to neural activation, i.e. $f_{V \rightarrow A}$, or by coupling these individualised results to detailed biophysical axon models to model the physiological effects of stimulation. By considering this physiological effect of stimulation in individualised models, we open the potential for individualised, targeted neuromodulation therapies based on biophysical targets, as well as a potential platform for the evaluation of novel neuromodulation devices.

5

A finite element solver for neuromodulation problems

Contents

5.1	Introduction	81
5.2	Methods	86
5.2.1	Physical problem	86
5.2.2	Implementation	87
5.2.3	Analytical validation	87
5.2.4	Numerical validation	89
5.2.5	Application to neuromodulation	90
5.2.6	Performance	91
5.3	Results	91
5.3.1	A custom finite element solver	91
5.3.2	Complex parameters	96
5.3.3	Performance	99
5.4	Discussion	101
5.4.1	Solving clinical neuromodulation problems	101
5.4.2	An open-source, high-performance solver	102
5.4.3	Conclusion	103

5.1 Introduction

Clinical neuromodulation involves the controlled application of electricity to influence the activity of the nervous system. The generation of an electric field in the tissue

produces a transmembrane potential at cells within the field. The excitable cells of the nervous system are sensitive to changes in transmembrane potential, with depolarisation of a sufficient magnitude producing opening of voltage-sensitive ion channels and generation of action potentials.

The ability to predict the electric field produced by neuromodulation therapies is potentially valuable for planning targeted modulation of specific structures and the design of novel stimulation systems [11, 18, 49, 50, 106]. The use of computational models of stimulation has shown predictive ability in the use of spinal cord stimulation [49], deep brain stimulation [50], and peripheral nerve stimulation [107, 108], as well as the design and testing of new targeted approaches [11, 18, 19, 106, 109].

The ability to predict the electric field produced by stimulation is made possible as the behaviour of electricity is well understood. The characteristics of electric fields are accurately described by the physical laws of electromagnetics [53]. This allows us to mathematically describe the physical problem encountered in human therapeutic neuromodulation. In the simplest case, the electric potential within the body produced by stimulation will satisfy:

$$\begin{aligned}\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega\end{aligned}$$

where V is the electric potential produced by stimulation, σ is the electric conductivity of the tissues, Ω is the region of interest, $\frac{\partial V}{\partial n}$ is the current density normal to the boundary and $\partial\Omega$ is the external (skin) boundary of the area of interest.

This describes a physical system where all materials obey Ohm's law with insulating conditions on the outer boundary.

This provides a steady-state solution where any current applied is constant. This provides a good approximation of the response to stimulation in human tissue as the tissue is largely resistive and the length of stimulation applied is typically sufficiently larger than the time constant of neural membranes such that transient responses during time-varying fields can be largely neglected [110]. In situations where the time-varying component varies sufficiently slowly relative to the neural

membrane time constant and the effects of transient responses are negligible, we can assume that the system is in steady state at each time-step; these quasistatic conditions allow time-varying fields to be evaluated in a simpler manner

However, in the case of complex time-varying fields where transient effects are important, it becomes necessary to consider the reactive component of the response to stimulation. In this case, the potential produced by a time-varying field can be described by:

$$\begin{aligned} \nabla \cdot \sigma \nabla V - \frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} &= 0 & \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 & \text{on } \partial\Omega \end{aligned}$$

where the first term describes the linear response to stimulation governed by Ohm's law and the second term describes the time-varying response. ε_0 represents the permittivity of free space and ε_r represents the relative permittivity of the tissue.

This allows us to accurately describe the distribution of electric potential in tissue using a straightforward partial differential equation with a homogeneous Neumann boundary condition on the external border to define an insulating condition on the outer boundary. We can then consider the effects of stimulation by adding additional boundary conditions.

Clinically, electrical stimulation can be applied to tissue using either voltage-controlled or current-controlled methods. In voltage-controlled stimulation, a fixed potential is applied across two electrodes and this will produce some current in the tissue, where the amount of current produced varies according to the resistance of the tissue. This can be described using a simple Dirichlet boundary condition:

$$\begin{aligned} V &= V_{\text{applied}} & \text{on } \Gamma_A, \\ V &= 0 & \text{on } \Gamma_R \end{aligned}$$

where V_{applied} is the potential applied, Γ_A is the surface of the active (stimulating) electrode and Γ_R is the surface of the return (ground) electrode. Here, there is simply a fixed potential defined on the electrode, and the current is then determined by whatever impedance is encountered within the body.

These variations in the current delivered over time due to changes in tissue resistance can lead to unpredictable clinical effects. As a result, it has become more common for clinical stimulation systems to deliver current-controlled stimulation. Under this paradigm, a fixed current is applied between two electrodes; the potential at the electrode surface is allowed to vary (up to some maximum value, termed the compliance voltage) in order to maintain a fixed total current despite changes in tissue impedance. This produces a more predictable and consistent effect. In this case, stimulation with a fixed current can be described using Neumann boundary conditions on the electrodes:

$$\begin{aligned}\frac{\partial V}{\partial n} &= J_{\text{applied}} && \text{on } \Gamma_A, \\ \frac{\partial V}{\partial n} &= -J_{\text{applied}} && \text{on } \Gamma_R\end{aligned}$$

where J_{applied} is the current density normal to the electrode surface and $\frac{\partial V}{\partial n}$ is the derivative of V normal to the electrode-tissue interface. This describes a fixed current source through the electrode surface, with a matching return current through the return electrode.

By combining the partial differential equation describing the physical problem with appropriate boundary conditions describing the nature of stimulation applied, we can accurately mathematically represent the physical problem of therapeutic neuromodulation.

Although we can accurately describe the potential produced by stimulation using a system of equations, this system is not analytically solvable. We cannot straightforwardly solve this system to tell us what the electric field in the tissue looks like. However, we can approximate the solution to this system using numerical methods. This allows us to compute a solution for some specific set of conditions, even though we do not have access to a more general global solution. In the context of neuromodulation, this means that although we cannot derive a mathematical expression that will describe the electric field produced by stimulation in all situations, we can compute a solution for any given individual.

The finite element method is a commonly used numerical method for solving systems of partial differential equations [55]. By specifying the domain we are interested in, Ω , and the boundary conditions, we can calculate a solution for that domain given those boundary conditions. By embedding patient-specific anatomy and the details of electrode position and geometry within our model of the domain, we can then calculate a patient-specific solution for the electric field produced by stimulation using this method. This allows us to accurately predict the effects of electrical stimulation applied to that individual.

Commercial systems for numerically solving systems of partial differential equations using the finite element method exist [95]. While extremely flexible and powerful, these systems are limited by the need for expensive licenses and the complexity of setting up problems to solve, producing significant barriers to their widespread application to neuromodulation problems. However, the physical system that needs to be solved in the context of therapeutic neuromodulation is relatively constant; the factors that change, such as patient-specific anatomy, electrode geometry and position and stimulation parameters can be contained in the domain and in boundary conditions. This opens the possibility of a simpler, single-purpose solver for problems in human neuromodulation that overcomes the issues with applying more general solutions to this area.

We therefore aimed to develop a numerical solver specifically for computing solutions to the physical system of therapeutic neuromodulation. We demonstrate a purpose-built solver for solving systems with voltage- and current-controlled stimulation for both steady-state and time-varying currents on patient-specific anatomy. This allows for detailed simulations of the electric fields produced by stimulation using a scalable, open source approach with the potential to be applied on a large scale to problems in human neuromodulation.

In the context of a broader system for model-based neuromodulation, this solver provides us with a method of solving the electric field V for some patient specific domain Ω , i.e. $f_{\Omega \rightarrow V}$. In combination with a method for deriving patient-specific models with simulated electrodes, this allows us to generate a function

$f_{\theta \rightarrow V}$, which predicts the electric potential produced in a specific patient using some chosen stimulation parameters θ ; this then forms the basis of a system for model-based individualised therapy.

5.2 Methods

5.2.1 Physical problem

The electric potential produced by stimulation was computed using a steady current formulation:

$$\begin{aligned}\nabla \cdot \sigma \nabla V &= 0 & \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 & \text{on } \partial\Omega\end{aligned}$$

For the case of time-varying fields, the following system was solved:

$$\begin{aligned}\nabla \cdot \sigma \nabla V - \frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} &= 0 & \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 & \text{on } \partial\Omega\end{aligned}$$

To simulate the effects of voltage-controlled stimulation, Dirichlet boundary conditions were added to electrode surfaces:

$$\begin{aligned}V &= V_{\text{applied}} & \text{on } \Gamma_A, \\ V &= 0 & \text{on } \Gamma_R\end{aligned}$$

For the case of current-controlled stimulation, Neumann boundary conditions were added to the electrode surfaces:

$$\begin{aligned}\frac{\partial V}{\partial n} &= J_{\text{applied}} & \text{on } \Gamma_A, \\ \frac{\partial V}{\partial n} &= -J_{\text{applied}} & \text{on } \Gamma_R\end{aligned}$$

This corresponds to the Laplace equation with variable coefficients with homogeneous Neumann boundary conditions on the external skin boundary and variable boundary conditions depending on the type of stimulation applied on the electrode surfaces.

Variational forms of the resulting systems of partial differential equations to describe this physical system were derived as outlined in appendix A.

5.2.2 Implementation

A solver for the variational problem was developed in the Python programming language [111] using the FEniCS platform [112] for developing numerical solvers for systems of partial differential equations.

The solver was implemented to solve the physical problems encountered in therapeutic neuromodulation using the above formalisms. The solver requires the definition of a domain, Ω , i.e. the region of interest. This takes the form of a finite element mesh. The physical properties of the subdomains must also be defined, i.e. the conductivity and relative permittivity of the tissues included in the domain.

The nature of the problem must then be defined. Having defined the domain and its physical properties, solutions can be computed for voltage- or current-controlled stimulation under steady-state or time-varying assumptions by varying the system of equations that is solved. An equivalent variational problem is defined. This is then solved numerically on the defined finite element mesh using a function space of first-degree Lagrange polynomials. The resulting linear system is assembled and solved using a sparse solver.

5.2.3 Analytical validation

Initial validation of the accuracy of the numerical solutions produced by the solver was carried out by assessing its ability to reproduce known solutions.

Simple models with known analytical solutions were created and solved. The numerical solutions produced with our custom solver were compared to the known exact solutions.

In the case of steady currents, the Laplace equation with variable coefficients was solved on the unit square with Dirichlet boundary conditions on the upper and lower borders:

$$\begin{aligned} -\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\ V &= 0 && \text{on } \Gamma_L, \\ V &= 1 && \text{on } \Gamma_U \end{aligned}$$

where Γ_U and Γ_L are the upper and lower borders of the unit square respectively. The square was divided into two equal subdomains with differing physical properties at its midpoint along the y -axis as shown in figure 5.1, i.e.:

$$\sigma = \begin{cases} \sigma_1, & \text{if } y < 0.5 \\ \sigma_2, & \text{if } y > 0.5 \end{cases}$$

where $\sigma_1 = 1.5$ and $\sigma_2 = 50$ for our simple model.




Figure 5.1: Simple geometry. Unit square divided into two equal subdomains at its midpoint along the y -axis. Different physical properties can be applied to each subdomain and Dirichlet boundary conditions applied to the upper and lower borders; this simple system can then be solved analytically and the results compared to numerical solutions to validate their accuracy.

The solution can then be approximated numerically using our solver, and the result compared to the known analytical solution:

$$y = \begin{cases} \frac{2\sigma_1 y^2}{\sigma_1 + \sigma_2} & \text{if } y < 0.5 \\ \frac{2\sigma_2 y^2}{\sigma_1 + \sigma_2} & \text{if } y > 0.5 \end{cases}$$

Using a similar approach for the time-dependent case, a simple problem with a known analytical solution was solved and compared to the numerical solution using the solver. In this case, the time-dependent Poisson equation was solved on the unit circle with time-varying Dirichlet boundary conditions:

$$\begin{aligned} \frac{\partial V}{\partial t} &= \nabla^2 V + f && \text{in } \Omega, \\ V &= V_D && \text{on } \partial\Omega, \\ V &= V_0 && \text{at } t = 0 \end{aligned}$$

where V varies with time and space, f is a source function, V_D is the value applied to the exterior boundary as a Dirichlet boundary condition and V_0 , the initial condition, is the value of V at the first timepoint.

For this simple model, we define an external boundary condition that varies linearly with time and quadratically with space:

$$V_D = 1 + x^2 + \alpha y^2 + \beta t$$

where $\alpha = 3$ and $\beta = 1.2$. This gives us the initial condition $V_0 = 1 + x^2 + \alpha y^2$, and it can be calculated that $f = \beta - 2 - 2\alpha$.

The solution to this problem can be approximated by the solver and compared to the exact solution, i.e. $V = 1 + x^2 + \alpha y^2 + \beta t$.

5.2.4 Numerical validation

To assess the ability of the solver to approximate solutions on more complex domains, where an analytical solution cannot be computed, a test problem was designed that captures the essential components of problems seen in neuromodulation, reduced to an idealised two-dimensional geometry as shown in figure 5.2. Problems on this idealised model were then solved using the custom solver and COMSOL Multiphysics [95] and solutions were compared in order to evaluate how well the custom solver approximated solutions to complex problems.

The two-dimensional geometry was made up of a tissue layer with two overlying electrodes in contact with the tissue surface. The electrodes were each surrounded by an insulating substrate on all surfaces except the surface that was in contact with the tissue. The electrode substrate and the remaining tissue surface were surrounded with air.

The physical properties used for each material type are shown in table 5.1.

Problems using Dirichlet and Neumann boundary conditions on the electrode-tissue interfaces were defined and solved. The results were compared with COMSOL Multiphysics to provide a measure of the accuracy of the numerical approximation achieved on neuromodulation problems relative to a known baseline.




Figure 5.2: Electrode geometry. A two-dimensional reduction of the features introduced by neuromodulation problems, including multiple domains, boundary conditions on internal domain boundaries and electrode-tissue interfaces. The lower part of the domain represents tissue. The large rectangle in the upper part of the domain represents electrode substrate; within this are two smaller electrode domains in contact with the skin surface. The remaining dark area is surrounding air; this does not form part of the modelled domain. This captures the key elements of problems on large, complex geometries. Solvers can then be assessed on this more easily analysable geometry in order to assess their accuracy at dealing with these features before scaling to large neuromodulation problems.

Table 5.1: Physical properties of tissue types for idealised geometries

Tissue type	Conductivity (S/m)
Tissue	$2.02E^{-2}$
Electrode	$5.99E^7$
Substrate	$1E^{-14}$

5.2.5 Application to neuromodulation

The ability of the custom solver to deal with complex neuromodulation problems on highly detailed models was evaluated. Image-based models were generated and solved using the solver as with the idealised geometries. Results were then assessed in order to evaluate the performance of the custom solver on detailed three-dimensional models.

Using the methods outlined in chapter 4, a detailed model of the lower thoracic spine was generated with a transabdominal electrode montage using round 3.2cm

diameter electrodes as in figure 4.2A. Physical properties were assigned to tissue segments according to table 4.1.

The resulting model was then solved using COMSOL as in chapter 4 and using the custom solver in order to assess whether comparable results were produced using this approach.

5.2.6 Performance

The performance of the solver was assessed by measuring the time taken to compute a solution for models of a range of sizes. This provides an estimate of the practical limits on the application of this approach to real problems. The ability to use a graphics processing unit to speed up the solution of the assembled linear system by parallelising the solution process over many compute nodes was also evaluated.

5.3 Results

5.3.1 A custom finite element solver

A custom solver is able to implement and solve the physical system describing neuromodulation problems. Assessment on problems with known analytical solutions and complex neuromodulation problems demonstrate the ability to produce reliable estimates of the electric field produced by stimulation, to a level at least equivalent to industry standard finite element solvers.

This solver can implement voltage- and current-controlled stimulation by varying boundary conditions, and can handle both static and time-varying stimulation waveforms. This allows the full range of modern neuromodulation techniques to be accurately simulated using an open-source numerical solver.

Figure 5.3A visualises the analytical solution to the Laplace equation with variable coefficients on the unit square with Dirichlet boundary conditions, representing a highly simplified case with a physical structure equivalent to voltage-dependent stimulation. This allows us to assess the performance of our solver on this kind of problem in comparison to a known exact solution.

The numerical results produced by our solver are shown in figure 5.3B. Visually, the results are close to identical to the true analytical solution. This shows that our solver is able to solve the physical problem posed by neuromodulation and can produce a good approximation of the true, known result.




Figure 5.3: Analytical vs. numerical, simple. **A** Analytical solution to the Laplace equation with variable coefficients on the unit square with Dirichlet boundary conditions. Provides a known benchmark with a physical structure similar to neuromodulation problems to compare to the results of numerical solvers. **B** Numerical results produced by a custom solver. The results are visually equivalent to the analytical solution, suggesting that the custom solver can accurately solve neuromodulation problems. **C** Point-by-point comparison of the solver’s performance to the exact solution. The solver’s numerical result matches the known solution at every point. The numerical solver is able to accurately approximate the true result.

A point-by-point comparison of the solver’s performance compared to the exact solution is shown in figure 5.3C. At all points, the result predicted by our numerical solver matches the known exact solution. The maximum error between the numerical result and the known solution was 3.3×10^{-16} . This is within the acceptable bounds of error for numerical methods, and indicates that our solver is able to accurately solve this physical system.

Our custom solver demonstrates similarly good performance on time-dependent problems. By solving a simplified time-dependent problem with a known analytical solution, we can compare our solver’s performance over time to the known solution. This allows us to assess the accuracy of our numerical methods over time. The difference between the numerical and analytical solutions over time is shown in figure 5.4. This demonstrates a gradual accumulation of error as time progresses. This is a known behaviour of the implicit Euler integration scheme used to integrate over time. However, the maximum error reached even after a large number of time steps is 2.5×10^{-14} . This is well within acceptable numerical ranges. This demonstrates

that our solver can solve systems with arbitrarily complex, time-varying waveforms with performance close to the theoretical minimum for the numerical techniques used, and well within acceptable margins.




Figure 5.4: Time dependent performance. Error between the numerical approximation and the known exact solution accumulates over time steps. This is a known behaviour of implicit Euler integration schemes. The maximum error is well within acceptable margins. The custom solver accurately solves time-dependent problems with error behaviour close to the theoretical minimum.

The custom solver performs at a level equivalent to industry standards on simplified versions of the target problem. Figure 5.5A shows the potential and electric field simulated in a 2D simplified geometry with multiple domains representing cathodic and anodic electrodes, surround electrode substrate, tissue and air, with current-controlled stimulation simulated via Neumann boundary conditions on the electrode surfaces, simulated in COMSOL. This provides a two-dimensional reduction of the target problem of neuromodulation, while maintaining the important features of multiple domains, boundary conditions on internal electrode-tissue interfaces and current-controlled stimulation.

Figure 5.5B shows the results of the custom solver on the same problem. We are able to solve the neuromodulation problem on a complex domain with the characteristics required to apply to large anatomical models. The results are visually similar to those produced by COMSOL. Figure 5.5C shows the results

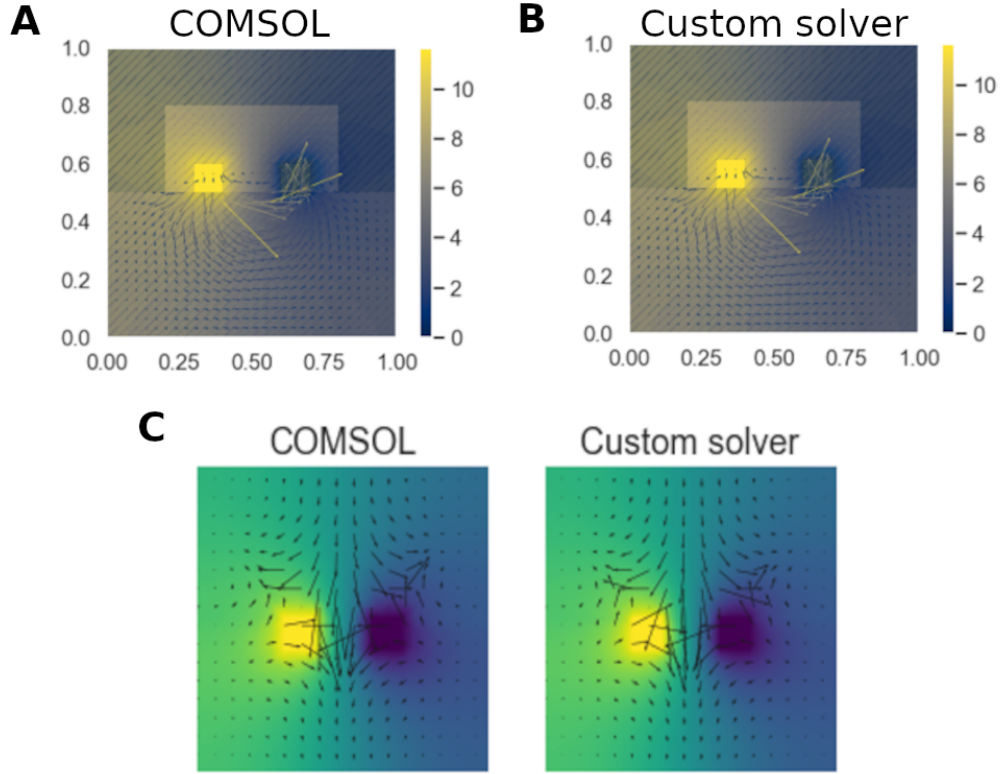


Figure 5.5: Analytical vs. numerical, complex. **A** Potential and electric field simulated in a simplified neuromodulation model using COMSOL. This provides a numerical benchmark to assess solvers against an industry standard commercial solver. **B** Potential and electric field simulated in a simplified neuromodulation model using our custom solver. The results are visually equivalent to those produced by COMSOL, suggesting that the custom solver is able to accurately solve problems with the key physical features introduced by neuromodulation. **C** The potential and electric field produced by stimulation for COMSOL and the custom solver visualised side-by-side. The results are equivalent, showing that the custom solver performs at a level similar to industry standard.

produced by COMSOL and our custom solver side-by-side. These are visually equivalent, suggesting that our solver produces reliable results on this kind of problem. This demonstrates that we can solve problems with the complex features introduced by neuromodulation with results similar to those produced by industry-standard commercial solvers.

A more detailed analysis of the differences in the results produced by our solver compared to an industry standard commercial solver is shown in figure 5.6. Figure 5.6A shows the absolute difference between the numerical results produced by each solver. This shows that the error is highest around the corners of the electrode

volumes, while remaining very low throughout the tissue volume. The greatest magnitude of difference between the solvers was 0.6V, which represents a small proportional error, within the acceptable range for numerical methods.

This concentration of error around the electrode corners may be due to the selection of boundary conditions, with uniform current density. It is also likely that the meshing around these corners contributed to the differences seen. While the geometries were identical, each solver used a different mesher; differences in meshing around these corners may have contributed to the error seen.

Figure 5.6B shows the potential measured along the electrode-tissue interface for each solver. As expected, there is a high potential at the stimulating electrode surface and the potential is equal to zero at the return electrode surface. Both solvers produce close to identical results along this cut line. This demonstrates that the custom solver produces reliable results along critical tissue boundaries, a necessary feature for application to neuromodulation problems.




Figure 5.6: Detailed comparison. **A** Absolute difference between the numerical results produced by COMSOL and the custom solver. The error is low throughout the volume, with highest values around the corners of the electrodes. This may be due to differences in meshing in addition to the boundary conditions used. The maximum error is within acceptable numerical margins. The two solvers produce equivalent results within reasonable error margins. **B** Potential measured along a cut line through the electrode-tissue interface. There is a high potential at the stimulating electrode surface and a low potential at the return electrode surface. Both COMSOL and the custom solver capture the expected behaviour along the electrode-tissue interface. **C** Potential measured along a vertical cut line midway between the electrodes. The COMSOL results demonstrates considerable numerical noise, while the custom solver shows much smoother variation between adjacent points. There is relatively little change in potential expected along this line. It is likely that differences in meshing, rather than the solver itself, contributed to the small difference seen here.

The potential measured along a cut line midway between the electrodes is shown in figure 5.6C. The potential along this line is roughly consistent along

its length, with a slight variation as the tissue boundary is crossed. Our custom solver captures this expected behaviour well, with smooth, consistent results within domains. However, the COMSOL results demonstrate significantly more numerical inconsistency, with significant variations in potential between adjacent points. This is physically unlikely, and represents noise introduced by the numerical methods used. It is likely that this represents small differences introduced by differences in meshing the geometry rather than differences in solver accuracy.

5.3.2 Complex parameters

Our custom solver can be applied to complex three-dimensional geometries. Using a detailed anatomical model derived from clinical imaging with simulated electrode placement as described in chapter 4, we can use our solver to numerically approximate the electric field produced by stimulation within the detailed three-dimensional tissue volume of the thoracic spine with transbdominal stimulation using 3.2cm diameter round electrodes as in figure 4.2A.

Figure 5.7 shows an axial slice of a CT scan with a contour map of the electric potential and a vector field showing the electric field simulated by our solver overlaid. This demonstrates that it is possible to solve for the electric field produced in individuals with arbitrary electrode positions and stimulation parameters. This allows individualised examination of the interactions between detailed anatomy and the electric field produced by stimulation, and allows for useful visualisation of the resulting electric field in the context of clinical anatomy.

Pointwise comparison of the results of COMSOL and the custom solver demonstrate that the maximum error was $<2\%$ of the result. This suggests that the results of the custom solver are analogous to those produced by COMSOL.

This approach can then be extended to other derived metrics. This allows the solver to be used to investigate the effects of stimulation. For example, metrics derived from the simulated electric field can be used to approximate the response of neural tissue to stimulation. By taking the Hessian matrix of the electric field and applying appropriate rotations, we can derive the activating function for axons




Figure 5.7: Image-based model, axial slice. Contour plot of electric potential (V) and vector field of the electric field produced by stimulation superimposed on clinical imaging. Neuromodulation problems can be solved on complex anatomical domains with simulated electrode placements and the results visualised in the context of clinical anatomy. This allows the effects of stimulation in individuals to be evaluated in great detail, opening the possibility for individually optimised neuromodulation therapies.

of specific orientations as described in chapter 3, allowing us to examine areas of likely action potential generation.

Figure 5.8 demonstrates the application of this rotated Hessian metric to examine areas of activation for axons in the spinal cord and nerve roots within the intervertebral foraminae using the results derived from the custom numerical solver. This shows an axial slice with the corresponding activating functions overlaid. This is a relative metric, where high values are predictive of initiation of action potentials, but there are no absolute values predictive of activation.




Figure 5.8: Activating function, axial slice. Rotated Hessian activating metrics, derived from the simulated electric field, for the spinal cord (left) and the nerve roots in the intervertebral foraminae (right) superimposed on clinical imaging (V/m^2). It is possible to derive useful metrics for investigating the effects of stimulation directly from the solver results. These images show that the activating function metric which predicts activation in axons aligned with the spinal cord is low within the spinal canal, but that the activation function predicting activation in axons aligned with the dorsal roots is high in the intervertebral foraminae. This can be seen by clustering of high values in the heatmap around these areas. This suggests that stimulation with the simulated parameters is unlikely to produce a direct effect in the spinal cord, but may produce activation of nerve roots within the intervertebral foraminae, where the activating function for axons oriented along this trajectory are high.

There are low values of the cord-focused activating function throughout, while there are regions of high activation for the root-focused metric, with particularly high levels seen within the intervertebral foraminae. This suggests that the simulated electric field is relatively unlikely to produce action potentials within the spinal cord, but may produce activation within nerve roots traversing the intervertebral foraminae.

This demonstrates that this custom solver can be used to solve for the electric field produced by stimulation in complex individual anatomy, and that these results can be used to derive useful metrics for examining the effects of stimulation on individuals.

This represents a numerical solver customised for solving problems in neuromodulation, which can simulate voltage- and current-controlled stimulation with fixed or time-varying stimulation waveforms on even highly complex domains at a level equivalent to industry standard, and which can produce results that are

potentially valuable to the investigation of the effects of neuromodulation and to the development of individually targeted treatments.

5.3.3 Performance

While this numerical approach is able to produce accurate results on highly complex models to simulate the effects of neuromodulation, the time taken to produce a result is limited by computational resources. Figure 5.9 shows the time taken to produce a result using our custom solver for models of varying sizes using a standard desktop workstation and a graphics processing unit, which offers the potential for solving the resulting linear system efficiently by parallelisation.




Figure 5.9: Solver performance. The time taken to solve a model is heavily dependent on model size. Up to several million elements, models can be solved in seconds to minutes on standard hardware. Beyond this, when large models start to exceed the RAM available, the time taken to solve increases exponentially. This can be overcome by increasing the memory and computational resources available. The use of dedicated graphics processing units for solving the resulting linear system in parallel is limited by the high efficiency achieved on standard hardware, the need to transfer large volumes of data to the dedicated hardware and the limited memory available on these systems. Solving models of even one million elements was not possible on these systems due to memory constraints. CPU: central processing unit (standard hardware); GPU: graphics processing unit.

The time taken to produce a result is dependent on the size of the model. However, even models with millions of elements, enough to accurately capture individual anatomical features, can be solved in seconds to minutes. There is an

exponential increase in the time taken to produce a result as the model size starts to exceed what can be held in random access memory. Here, for a system with 16GB of RAM, we see a model with one million elements reaching a solution in under a minute, while extending to ten million elements results in a solution time of thirty minutes, with sizes beyond fifty million failing to converge. This highlights a known limitation of these numerical approaches; convergence of highly complex models is heavily dependent on resources available. However, highly detailed anatomical models can be accurately represented with element counts in the low millions, and can be solved using our dedicated solver in seconds to minutes on systems with even moderate resources. In the event of requiring much larger models, the addition of greater computational resources can overcome this barrier.

The use of dedicated graphics processing units demonstrated little advantage over standard hardware. Although the resulting parallelisation theoretically allows the linear system assembled by the finite element method to be solved more efficiently, in reality this produced no advantage in total solution time. This is because our solver produced results sufficiently quickly on standard hardware that the advantage of dedicated graphics processing units was lost due to the time taken to transfer data to the hardware. This approach would theoretically become advantageous as the linear systems became larger and took longer to solve on standard hardware, however this would require large amounts of RAM on the graphics units. In the present assessment, even models with one million elements could not be solved on our graphics units due to memory limitations.

These results demonstrate that our solver is able to produce simulations of the effects of stimulation on highly complex models using easily available hardware in short times. This represents a highly performant, accurate numerical method for solving neuromodulation problems, potentially suitable for applications to individually optimised treatment strategies.

5.4 Discussion

5.4.1 Solving clinical neuromodulation problems

These results demonstrate that it is possible to accurately represent the physical problem encountered in neuromodulation using a system of partial differential equations and to solve this system numerically to simulate the effect of stimulation.

We demonstrate a purpose-built numerical solver designed to solve this physical problem using the finite element method. This custom solver is able to reproduce known solutions and performs equivalently to industry standard finite element packages, providing a reliable means of simulating the effects of stimulation.

The ability to automatically solve for the electric field produced by stimulation in individual patients' anatomy with arbitrary simulated electrode placements is particularly interesting. This allows for the effects of stimulation on individual patients to be evaluated pre-intervention. Using this method, we can evaluate the effects of variations in electrode geometry, electrode position and stimulation parameters on a given individual. This provides potentially valuable information for the planning of neuromodulation interventions and the selection of stimulation parameters. Further, this offers a potentially valuable means of testing proposed novel electrodes and stimulation methods. By simulating the effects of changes in stimulation methods, we can compare different stimulation modalities computationally to predict their effects, providing a valuable resource for the design of new neuromodulation therapies.

This approach offers a good degree of flexibility. We demonstrate the ability to solve static and time-dependent problems for both voltage- and current-controlled stimulation. This approach can be scaled to arbitrary numbers and arrangements of electrodes and can be applied to any patient-specific anatomy. By taking advantage of the mathematical structure of the physical problem encountered in neuromodulation and embedding the factors that vary between cases in the formation of the finite element mesh and in the boundary conditions, we can produce a general solver for all neuromodulation problems with sufficient flexibility

to deal with any type of electrical stimulation to any region while automatically constructing and solving the mathematical problem. This significantly reduces the technical barriers to the application of this approach to clinical neuromodulation and opens the possibility of more widespread use of computational modelling for planning individualised treatments.

Although the results of the custom solver varied only little from those produced by COMSOL, some deviations did exist. Particularly in the case of the simplified 2D geometry, it is likely that some of these differences were due to differences in meshing rather than due to the solver. Additional steps to optimise meshing, or to generate a uniform mesh for comparison between solvers to control for this, would be likely to reduce some of this deviation between models.

Even when using identical meshes as in the case of image-based models, small deviations continued to exist. This is likely to be the result of numerical differences in the solvers. Additional tests to characterise these differences would be valuable to inform potential sources of error; additional optimisation of the underlying models may help to reduce these deviations.

5.4.2 An open-source, high-performance solver

We demonstrate an open-source, freely available numerical solver for neuromodulation problems. This provides an easily available tool that can rapidly solve problems for individual patients, allowing prediction of the electric field produced by stimulation. This tool removes the technical barriers to using general-purpose commercial finite element tools, while retaining sufficient flexibility to be able to accurately represent all currently used electrical stimulation modalities with straightforward extension to novel methods. This allows rapid assessment of the patient-specific effects of stimulation using a simple tool.

This solver scales well to very large problems by using sparse methods for solving the resulting linear system, allowing even very detailed models to be solved on systems with limited memory. Solutions can also be carried out using graphics processing units, although the resulting performance gains are largely offset by the

need to transfer data and the typically low memory available on these systems. However, this custom solver is well positioned to be scaled up on large computing clusters. The solver can be directly set up and run on high-performance computing systems without the need for alterations to the model. This allows straightforward scaling to very large problems, or accelerating solutions in order to solve large numbers of models in a short period, for example for testing a wide range of stimulation parameters for individual optimisation.

This provides a tool suitable for widespread use for the investigation of the effects of neuromodulation. Without the requirements for licenses or technical barriers to setting up problems and with the ability to run on both low-resource and high-performance environments with minimal changes, this custom solver is well positioned to be applied on a widespread basis to problems in neuromodulation, and potentially to the individualised optimisation of therapies in a clinical context.

5.4.3 Conclusion

We demonstrate the ability to solve for the electric field produced by stimulation on arbitrary models using the finite element method. The solver developed for this purpose represents a simple to use tool that can solve the physical problems encountered in neuromodulation. This allows for the complex interactions between electrodes, stimulation parameters and detailed individual anatomy to be taken into account and the effects of stimulation evaluated for any arbitrary set of stimulation parameters or electrode positions.

This represents a valuable step towards the development of model-based neuromodulation therapies. By evaluating the electric field produced by stimulation, we can assess how changes in the electrodes or stimulation parameters affect the electric field produced. This allows us to investigate the effects of neuromodulation and to individually optimise neuromodulation therapies to individual patients by predicting the effects of stimulation computationally.

In the context of a broader system for model-driven stimulation, this provides a function $f_{\Omega \rightarrow V}$ which computes the electric potential produced by solving the

system of partial differential equations describing electrical neuromodulation of a patient-specific domain Ω , which also includes information on the position and geometry of electrodes. By combination with a function for deriving patient-specific models from a set of stimulation parameters, $f_{\theta \rightarrow \Omega}$, this gives us the ability to compute the electric field produced in an individual for a given set of stimulation parameters, i.e. this gives us the function $f_{\theta \rightarrow V}$.

In order to move towards truly targeted neuromodulation based on computational models, we need to be able to evaluate the effects of the produced electric field on the nervous system in order to predict its physiological effects. To achieve this, we need to develop a method that links our electric potential produced by stimulation, V , to the membrane potential of our target neural structures, V_m , i.e. a function $f_{V \rightarrow V_m}$. By then defining these biophysical measures as an objective function, we can link our produced electric fields, calculated using the presented numerical solver, to a target physiological effect and move towards patient-specific optimisation of neuromodulation.

6

Individualised models for biophysically optimised neuromodulation

Contents

6.1	Introduction	106
6.2	Methods	112
6.2.1	Electric field computation	112
6.2.2	Biophysical axon models	113
6.2.3	Identifying trajectories	114
6.2.4	Mapping axons to trajectories	116
6.2.5	Computation of activation	118
6.2.6	Biophysical objective functions	118
6.3	Results	120
6.3.1	Derivation of axon trajectories	120
6.3.2	Coupling of models	123
6.3.3	Strength-duration curves	124
6.4	Discussion	126
6.4.1	Linking anatomy and physiology	126
6.4.2	Automated axon mapping	127
6.4.3	Modelling the effects of stimulation	129
6.4.4	Individually optimised neuromodulation	130
6.4.5	Conclusion	131

6.1 Introduction

The application of electrical stimulation produces an electric field within the tissue. This electric field then has an effect on the neural structures within the tissue. It is this change in activity of the neural tissue, such as the production of an action potential or blocking or normal activity, that mediates the desired clinical effect. Equally, it is changes in the activity of off-target structures that mediates side effects, such as pain due to activation of cutaneous afferents in surface stimulation or motor effects due to activation of internal capsule fibres in deep brain stimulation.

Knowledge of the electric field produced by stimulation is therefore useful, as this influences neural activity, but it is not sufficient to predict the physiological effects of neuromodulation. We need to be able to predict not only the electric field that is produced by stimulation, but also how this electric field will influence the underlying neural elements. This then allows us to potentially predict the physiological effects of electrical stimulation.

The approximation of the electric field produced by neuromodulation has received significant attention. Computationally simple measures are used in some areas, such as assuming a quadratic rate of decay of potential around an implanted electrode [77, 88], but the use of finite element models has produced the most reliable results [18, 49, 50]. This allows us to compute the potential in the tissue with complex arrangements of electrodes and anatomy. This approach has been applied to predict the electric fields produced by a range of stimulation modalities, including transcutaneous direct current stimulation [113, 114], deep brain stimulation [50], and spinal stimulation [49]. However, while idealised geometries are common, this method also extends to models that capture detailed individual anatomy. It has been shown that capturing this detail allows more accurate simulations of the effects of stimulation [49, 50]. We have demonstrated in chapter 4 that it is possible to automatically derive models that capture detailed patient-specific anatomy and in chapter 5 that these models can be used to accurately approximate the electric field produced by arbitrary neuromodulation interventions on detailed individual anatomy.

However, the ability to predict the physiological effects of stimulation has received comparatively little attention. In order to evaluate this, it is necessary to derive some measure that captures neural activity from the computed electric field. Multiple simplified metrics are used to try to link the predicted electric field to a physiological effect. These include the $L2$ norm of the electric field [51, 77, 88], which is simple to compute but has a limited physiological link to the activity of neural tissue, and variations of the volume of tissue activated [78], which provides a useful but limited metric depending on how it is computed. The activating function [89, 90] represents the classic simplified metric linking extracellular potentials to neural activity. We have demonstrated in chapter 3 that this can be generalised to a field, providing a method of deriving a simple metric that predicts the neural response to electrical stimulation from a field computed using the finite element method. This overcomes the issue of orientation-dependent responses to stimulation that is not taken into account by usual simplified metrics, increasing the accuracy of these approaches.

However, the most reliable method of determining the neural response to stimulation requires the use of a detailed biophysical model of neural activity [11, 14, 18, 19, 49, 50]. This involves coupling the system of partial differential equations describing the electric potential produced by stimulation with a system of ordinary differential equations describing the response of the neural transmembrane potential to the application of an extracellular potential.

The classic method for modelling neural dynamics is given by the Hodgkin-Huxley model [74], where the ionic current across the cell membrane, I_{ionic} , is given by the sum of the currents through each of the ion channel species in the cell membrane and an ongoing "leak" current:

$$I_{ionic} = \bar{g}_L(V_m - E_L) + \bar{g}_k n^4(V_m - E_k) + \bar{g}_{Na} m^3 h(V - E_{Na})$$

where \bar{g}_L , \bar{g}_k and \bar{g}_{Na} are the conductances for the leak current, potassium channels and sodium channels respectively; E_L , E_k and E_{Na} are the equilibrium potentials for each current and n , m and h are time-dependent gating variables on

the active ion channels, which give the membrane its active dynamics by allowing the conductances of ions to vary over time.

Each current is therefore determined by the difference between the current membrane potential and that current's equilibrium potential and the time-varying conductance for that current. The total transmembrane current is then simply given by the sum of these currents.

The complex time-dependent dynamics of the membrane in response to changes in membrane potential, allowing phenomena such as action potential generation and refractory periods, are determined by the dynamics of these gating variables which alter the conductance of the membrane to ionic species. These gating variables are themselves a function of the transmembrane potential, leading to the voltage-gated conductances that allow for action potential generation:

$$\begin{aligned}\frac{dn}{dt} &= \alpha_n(V_m)(1 - n) - \beta_n(V_m)n \\ \frac{dm}{dt} &= \alpha_m(V_m)(1 - m) - \beta_m(V_m)m \\ \frac{dh}{dt} &= \alpha_h(V_m)(1 - h) - \beta_h(V_m)h\end{aligned}$$

where α and β are rate constants for each ion channel type, and are determined empirically.

This provides a detailed model of the active dynamics of the neural membrane. In order to model the activity of a cell in an external electric field, the neuron can be modelled as a multicompartment cable [71–73], where the current in each compartment of the cell is determined by a passive capacitive current across the cell membrane, a voltage-dependent active ionic current governed by Hodgkin-Huxley dynamics and a passive axial current between neighbouring compartments within the cell:

$$C_m \frac{d(V_{i,n} - V_{e,n})}{dt} + I_{ionic,n} + G_a(V_{i,n} - V_{i,n-1}) + G_a(V_{i,n} - V_{i,n+1}) = 0$$

where C_m is the membrane capacitance, $V_{i,n}$ is the intracellular potential in compartment n , $V_{e,n}$ is the extracellular potential in compartment n (noting that

the transmembrane potential $V_m = V_i - V_e$ for any given compartment) and G_a is the axial conductance, i.e. the intracellular conductance along the length of the cell.

This models the electrical behaviour of a neural cell as being composed of an active time-varying ionic current, which allows the generation of action potentials through voltage-gated dynamics, an axial current which allows propagation of activity along the length of a cell as a change in membrane potential in one compartment will drive a change in neighbouring compartments, leading to transmission of action potentials, and a passive current across the cell membrane.

In order to model the effect of extracellular stimulation on neural tissue, we can couple our calculated electric field to a model of the neural tissue by setting the V_e parameter of this capacitive current, i.e. the potential just external to the cell membrane, using our calculated potential. We can then solve this system of ordinary differential equations in order to compute the biophysical response of the cell to this extracellular potential.

Many variations of neuron models exist that can be used to model the effects of stimulation on neural tissue [75, 77–79]. One of the most commonly used in the McIntyre-Richardson-Grill model of myelinated axons [76]. This models the activity of myelinated axons by limiting the presence of an active ionic current to the nodes of Ranvier, as occurs physiologically. Further, this models the myelinated segments using a double-cable model by introducing an additional "external cable" outside the standard cable model made up of a passive current determined by the myelin resistance and capacitance and an axial current along the periaxonal space, between the myelin and the cell membrane. An additional current is then added at the myelin attachment zone at the borders of the nodes of Ranvier allowing this current to escape passively following action potential generation. This model has been shown to accurately capture the behaviour of myelinated axons [76], and has been validated in the context of extracellular stimulation with a distant electrode [77–79].

The use of coupled finite element and biophysical models is relatively well established for investigating the effects of neuromodulation [11, 18, 49, 50]. However,

these approaches typically rely on idealised geometries [45, 50] and laboriously hand-drawn trajectories [14, 18]. The response of neurons to stimulation is heavily dependent on the physical position of the cell within the electric field and on its orientation [115, 116]; the neural response to stimulation is therefore determined by a complex interaction between individual anatomy and the electric field produced by stimulation. It has been demonstrated that accounting for this detailed anatomy improves the performance of modelling approaches [49, 50]. The ability to accurately map axons into realistic anatomy is therefore important for the generation of individualised models of the effect of stimulation.

Multiple methods have been developed to represent the complexity of anatomical trajectories in finite element models. Idealised models for mechanistic investigation allow direct specification of the anatomy of interest by defining this manually [14, 18, 45]. This approach ensures the features of interest are represented. This effectively allows for maximum accuracy in controlling and representing anatomic features. However, this approach is labour intensive and is not scalable to patient-specific, image-based models.

Automated positioning of axon trajectories in idealised models has shown promise [117, 118]. This aims to use known principles of the anatomy of the structures of interest to define anatomically consistent trajectories. Axon models can then be mapped onto these trajectories. In this way, anatomically realistic models can be generated without the need for detailed manual specification. However, this approach is defined in relation to specific structures within specific idealised models. While adaptable, it is not straightforwardly applicable to other anatomical regions or to new models.

Simplification of individual anatomy to fit simplified trajectories has been used in other domains with some success [50]. This is a promising approach, however the loss of anatomical realism results in potentially less accurate results. In the context of spinal cord stimulation, mapping axons into patient-specific domains has been demonstrated [49]. This allows the use of known principles of the underlying

anatomy to be used to create trajectories within the domains of interest. As with idealised geometries, this is not easily applicable to other areas.

We aimed to build on these techniques to attempt to develop a method for mapping axons trajectories into complex geometries which is potentially useful in multiple regions. While initially applicable to the spinal cord and dorsal roots, we aimed for this technique to be generalisable to other applications, including peripheral nerves, by limiting the principle-based assumptions used. It is unlikely to be possible to generate a truly generalisable method as model adaptations to local anatomical features will always be required, however we wished to define a method that would provide a useful initial model in patient-specific anatomy for many potential anatomical regions.

We therefore developed a technique for the automated determination of axon trajectories that follow anatomic structures in realistic, complex models. We show that axon models can be coupled onto these trajectories and the effects of stimulation on these target structures can be modelled. This provides a means for carrying out simulations of the effects of neuromodulation while accounting for detailed anatomy and physiology. This approach is largely automated, potentially making it suitable for more widespread use in a clinical context by removing the time required to define axonal trajectories.

Having determined the biophysical effects of stimulation, we still need a target for evaluating the intended effect of stimulation. We can capture the clinical goal of our neuromodulation intervention using a measure based on these biophysical parameters. For example, we may aim to find the lowest amplitude current that we need to apply in order to activate some target structure, or to find the electrode position that maximises activation in our target structure. By deriving an objective function using biophysical measures, we can capture the intended physiological effect of stimulation. We can then apply techniques from numerical optimisation in order to determine the parameters we need to use in order to maximise our measure of activation.

This approach allows the physiological effects of stimulation to be predicted and for stimulation parameters to be optimised to achieve our target response.

This opens the possibility for truly individually optimised neuromodulation using detailed computational models.

In the context of developing an integrated system for targeted neuromodulation, this allows us take an extracellular potential V , computed using the finite element method using a patient-specific model, and calculate the membrane potential V_m in target neural structures - that is, it gives us a function $f_{V \rightarrow V_m}$ that relates the potential produced by stimulation to the membrane potential in our target structure. We can then create a function $f_{V_m \rightarrow A}$ that relates our measured membrane potential to some objective function A that captures our clinical intent, such as maximising activation in our target structure.

Using our previously constructed functions to compute the extracellular potential produced in an individual with a given set of stimulation parameters, $f_{\theta \rightarrow V}$, we can construct a function $f_{\theta \rightarrow A}$ which relates our choice of stimulation parameters to our chosen objective function. We can then apply optimisation techniques in order to calculate the optimal stimulation parameters for achieving our clinical aim [91].

6.2 Methods

6.2.1 Electric field computation

Finite element models of detailed anatomy with simulated electrode positions were derived using the methods outlined in chapter 4. A model of the thoracolumbar spine was derived. A dorsoventral electrode montage was simulated with a 5cm x 5cm cathode positioned in the midline over the T10/T11 intervertebral space and a 5cm x 5cm anode positioned at the same level in the midline anteriorly.

The steady-state electric field produced by stimulation was computed by solving the Laplace equation with variable coefficients using the methods detailed in chapter 5. Dirichlet boundary conditions were used on the electrode-tissue interfaces to simulate the application of stimulation:

$$\begin{aligned}
\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\
\frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega, \\
V &= V_e && \text{on } \Gamma_A, \\
V &= 0 && \text{on } \Gamma_R
\end{aligned}$$

where V is the electric potential produced by stimulation, σ is the tissue conductivity, Γ_A and Γ_R are the surfaces of the active and return electrode respectively, Ω is the domain modeled, and $\partial\Omega$ is the external skin boundary. V_e is the electric potential applied at the active electrode; this is computed such that $\int_{\Gamma_A} \frac{\partial V}{\partial n} = J_{applied}$, so that current $J_{applied}$ is delivered via the active electrode.

For the present model, a current of 100mA was simulated. The field produced by other amplitudes was then derived by simple scaling. Since the Laplace equation is a linear partial differential equation, it follows that a scaled input current αJ produces a scaled potential αV . We can therefore determine the field produced by other amplitudes by scaling the results such that $f(J_{new}) = V_{new} = V \cdot \frac{J_{new}}{J_{applied}}$, where J_{new} is the new current to be investigated and V_{new} is the potential produced by stimulation at that amplitude, all other factors held constant.

6.2.2 Biophysical axon models

The effects of the computed electric field was evaluated by coupling the results of the finite element simulation to a biophysical axon model. A detailed multicompartment cable model was implemented in NEURON [92]. The extracellular potential was applied to each segment of the multicompartment model by setting the value of the extracellular potential term in the cable equation using the results of the finite element model. This biophysical model was then solved numerically in order to evaluate the effect of stimulation on the membrane potential.

The McIntyre-Richardson-Grill model of a myelinated axon was used [76]. This model is made up of four repeating segment types, representing the nodes of Ranvier (NODE), the paranodal myelin attachment zones (MYSA), the paranodal fluted segment (FLUT) and the internodal segments (STIN). Each segment has distinct biophysical behaviours, with active ion channel dynamics limited to the

nodes of Ranvier and explicit representation of the myelin sheath using a double-cable structure in other segments. $11.5\mu\text{m}$ axons were modelled with biophysical parameters shown in table 6.1.

Table 6.1: Biophysical parameters for myelinated axon model

Model parameter	Symbol	Value
Nodal capacitance	C_n	$2\ \mu\text{F}/\text{cm}^2$
Internodal capacitance	C_i	$2\ \mu\text{F}/\text{cm}^2$
Myelin capacitance	C_m	$0.1\ \mu\text{F}/\text{cm}^2$
Axoplasmic resistivity	ρ_a	$70\ \Omega\text{cm}$
Periaxonal resistivity	ρ_p	$70\ \Omega\text{cm}$
Myelin conductance	g_m	$0.001\ \text{S}/\text{cm}^2$
MYSA conductance	g_a	$0.001\ \text{S}/\text{cm}^2$
FLUT conductance	g_f	$0.0001\ \text{S}/\text{cm}^2$
STIN conductance	g_i	$0.0001\ \text{S}/\text{cm}^2$
Fast Na^+ conductance	g_{Naf}	$3.0\ \text{S}/\text{cm}^2$
Slow K^+ conductance	g_{Ks}	$0.08\ \text{S}/\text{cm}^2$
Persistent Na^+ conductance	g_{Nap}	$0.01\ \text{S}/\text{cm}^2$
Nodal leakage conductance	g_L	$0.007\ \text{S}/\text{cm}^2$
Na^+ equilibrium potential	E_{Na}	$50.0\ \text{mV}$
K^+ equilibrium potential	E_K	$-90.0\ \text{mV}$
Leakage reversal potential	E_L	$-90.0\ \text{mV}$
Rest potential	V_{Rest}	$-80.0\ \text{mV}$

This model was implemented in Python [111]. Axon models were implemented to facilitate mapping onto arbitrary trajectories; the number and arrangement of segments was automatically adjusted in order to fit onto complex trajectories while maintaining the correct model geometry. This provides a detailed physiological model of the effects of stimulation while representing complex anatomy.

6.2.3 Identifying trajectories

Axon trajectories were automatically determined using image-based models. For our present model, trajectories for axons were derived in the spinal canal representing the spinal cord and in the intervertebral foraminae representing the spinal nerve

roots bilaterally. This allows the physiological effects of stimulation on spinal cord axons and on nerve roots within the foraminae to be investigated in detailed anatomical models.

The target structures, i.e. the spinal cord and the intervertebral foraminae, were marked on the original clinical imaging during the segmentation process. The segmented volumes representing the targets were then transformed into point-cloud representations, i.e. an $N \times 3$ matrix made up of the x , y and z coordinates of N points within the target structure.

We then apply an eigendecomposition to the point cloud representation of the target structure. The target points are then projected onto each of the ranked eigenvectors, providing a change of basis to a new orthogonal coordinate system where the first basis vector lines up with the primary eigenvector of the target point cloud. This corresponds to the direction of maximal variation of the anatomic structure of interest, providing us with a new coordinate system oriented along the primary axis of the target structure.

In the new coordinate system, we take a seed point as a starting position for our trajectory. We position this at the lowest value of the primary eigenvector, i.e. at the lowest point along the structure's primary axis. In the case of the spinal cord, this corresponds to the most inferior aspect of the cord.

In this plane, we determine the position within the target structure by finding the minimum and maximum values along the orthogonal basis vectors; this gives us, in the case of the spinal cord, the coordinates of the most anterior, posterior and lateral points of the cord at the current plane. We randomly select a point between these extremes; this gives us a randomly selected point within the spinal cord on a plane taken through its most inferior aspect.

The trajectory is then continued by moving progressively along the primary eigenvector (i.e. along the spinal cord from inferior to superior, or along the intervertebral foramen from outside to inside). At each position along this primary axis, we find a point along the orthogonal bases that is the same proportional distance between the anteroposterior and lateral extremes as the original seed point.

In this way, we derive a set of points defining a trajectory that passes along the primary axis of the target structure and follows its natural curvatures, remaining in a consistent internal position throughout.

This set of points is finally transformed back into the original coordinate system by reversing the change of basis, producing a set of points defining a trajectory along our target structure in the original image coordinates. This can then be used to map biophysical axon models into the model space. This provides an automated means of deriving anatomically realistic axon trajectories from complex anatomical models, allowing detailed biophysical models to be automatically coupled to detailed finite element models. This can then be repeated multiple times to provide a distribution of axons of differing sizes and biophysical properties within the target structures.

6.2.4 Mapping axons to trajectories

Having defined a trajectory, we then need to map axon models onto the targets. This requires determination of the length of the required model, calculating the required number of model compartments and establishing the physical coordinates of each segment along the target trajectory.

The length of the trajectory is determined by simply calculating the cumulative $L2$ norm of the full trajectory. This provides a measure of the geodesic distance along the trajectory:

$$d(p) = \sum_{n=2}^N \sqrt{p_n^2 + p_{n-1}^2}$$

where p is a set of points defining the trajectory and d is cumulative distance along the trajectory, or total trajectory length. This defines the total length, in the model coordinate system, of the required axon model. Each model segment has a fixed length (e.g. the nodes of Ranvier are a fixed size and a fixed distance apart, regardless of axon length). We can therefore calculate the number of segments required to create a model of the required a length. A model with the required number of compartments with the appropriate arrangements can then be constructed.

This gives us an axon model with a total length equal to our target trajectory and a set of points along that trajectory. We can then calculate the actual physical coordinates of each segment of the model. This "maps" the axon model into the finite element model, providing every segment with a physical position that corresponds to its actual anatomic location.

Using our set of points p and measures of distance along the trajectory d , we can define a linear interpolation function $f_{d \rightarrow p}$ that maps any given distance to the corresponding point along the trajectory. Given that we know the physical length of each axon model segment, we can calculate the distance of each segment along the trajectory, d_{axon} . We can then apply our interpolation function to these distance measures, i.e. $f_{d \rightarrow p}(d_{axon}) = p_{axon}$. This produces a set of x , y and z coordinates for each segment of the axon model such that every segment has a physical position along the target trajectory, mapping our detailed physiological model into a complex anatomical geometry.

Lastly, we can use our function $V(x, y, z)$ to determine the extracellular potential produced by stimulation at each compartment of the axon model. This allows us to couple the finite element and biophysical models by placing the axon model within the model anatomy and computing the potential experienced by the axon at every point along its length. We can then use this value to drive the extracellular potential term in our system of biophysical equations and calculate the physiological response to stimulation.

We used this approach to simulate populations of axons in the dorsal columns of the spinal cord and in the dorsal roots. 100 seed points were selected at random within the dorsal column segmented volume. Trajectories were then determined within the dorsal columns from these seed points. Axon models were mapped onto each trajectory, producing 100 simulated dorsal column axons distributed throughout the dorsal columns volume, each following anatomically realistic courses. The superior extent of the dorsal column axons was at the superior limit of the domain; the inferior extent was at the inferior limit of the domain. Mean length

of simulated dorsal column axons was $324\text{mm} \pm 2\text{mm}$. Axon diameters were randomly distributed between $5.7\mu\text{m}$ and $11.5\mu\text{m}$ [49, 119, 120].

Dorsal root axons were created using a similar approach. 100 trajectories were generated within the target volume. Each trajectory started just lateral to the spinal cord in the dorsal root entry zone. They extended laterally through the spinal canal in realistic trajectories, passing through the intervertebral foramen and beyond. The mean length of dorsal root trajectories was $58\text{mm} \pm 6\text{mm}$. Axon models were mapped onto each of the 100 trajectories. Axon diameters were again randomly distributed between $5.7\mu\text{m}$ and $11.5\mu\text{m}$ [49, 119, 120].

6.2.5 Computation of activation

Having mapped the axon model into the finite element model, we can measure its response. The extracellular potential at each compartment of the axon model is used to set the amplitude of a rectangular stimulation waveform. The extracellular potential term of the cable equation at each segment is then driven using this rectangular waveform, simulating the effects of applying stimulation of some duration to an axon at the target position using the simulated electrode arrangement. The model is then solved using a backwards Euler implicit integration scheme with an integration time step of 0.001 ms .

In this way, the membrane potential at each point in the axon is recorded as stimulation is applied. The generation of action potentials in response to stimulation was then determined by evaluating the membrane potential over time following stimulation.

6.2.6 Biophysical objective functions

In order to link our intended effect to the biophysical measure derived from simulation, we defined a set of objective functions that represented our desired outcomes. This defines a set of functions $f_{V_m \rightarrow A}$, where A is our objective function.

We define a simple function to determine whether our target structure has been activated by stimulation or not:

$$A_{activated} = \begin{cases} 1, & \text{if } \max(V_m) > 0 \text{ in } > 70\% \text{ of axons} \\ 0, & \text{if } \max(V_m) > 0 \text{ in } < 70\% \text{ of axons} \end{cases}$$

i.e. we consider our target structure to have been "activated" by stimulation if $> 70\%$ of the population of 100 axons produce an action potential due to stimulation, defined by the membrane potential becoming positive at any point following the application of stimulation. 70% was selected as the threshold for considering the target activated as this represents significant activation of the target structure to a level that is not easily explained by noise, without requiring total activation of all axons, which could be limited by issues with individual axon trajectories.

By combining our function for getting the potential produced by stimulation, $f_{\theta \rightarrow V}$, our function for coupling this to biophysics, $f_{V \rightarrow V_m}$ and this biophysical objective function, we can derive a function $f_{\theta \rightarrow A}$ that tells us, for any given set of stimulation parameters, whether our target structure was or was not activated in that individual.

We can then use this function to perform numerical optimisation. In order to achieve our clinical goal of determining the lowest amplitude required to produce activation of our target, $J_{threshold}$, we apply a binary search algorithm to determine the lowest value of J that produces a positive value of $A_{activated}$.

We first set J_{max} and J_{min} , the maximum and minimum allowable amplitudes, to 0mA and 150mA respectively, and we set a desired precision, p , of 1mA, i.e. we want to determine the threshold between 0mA and 150mA to an accuracy of within 1mA.

$$J_{max} = 150$$

$$J_{min} = 0$$

$$p = 1$$

At each iteration, we set the applied current, θ_J , to the midpoint between the minimum and maximum amplitudes.

$$\theta_J = \frac{J_{max} + J_{min}}{2}$$

We can then compute the response $f_{\theta \rightarrow A}(\theta_J) = A$. If $A = 1$, we set $J_{max} = \theta_J$, or if $A = 0$ we set $J_{min} = \theta_J$:

$$\begin{cases} \text{if } A_{activated} = 1 : J_{max} = \theta_J \\ \text{if } A_{activated} = 0 : J_{min} = \theta_J \end{cases}$$

The maximum and minimum acceptable amplitudes are therefore iteratively altered according to the individual response to stimulation based on whether a response was produced. This process was repeated until the minimum and maximum agreed to within the specified precision, at which point the activation threshold was detected.

$$\text{if } (J_{max} - J_{min}) < p : J_{threshold} = \theta_J$$

This gives us the threshold current, the minimum current required to produce activation of our target neural structure, at a given pulse width. We can then repeat this process for a range of pulse widths in order to produce a strength-duration curve. This allows us to characterise the stimulation parameters required to activate our neural target using a standard physiological measure and to visualise the variation in threshold current as the stimulation pulse width is varied.

This provides us with a method to calculate the threshold current for activating our target where θ_J , the applied current, is the only free parameter. This can then be extended to optimise other parameters.

6.3 Results

6.3.1 Derivation of axon trajectories

Complex axon trajectories can be derived directly from clinical imaging. Target anatomy can be specified directly on the original imaging, and from this realistic trajectories can be determined within the anatomic structures. This allows for targeting of specific structures even in detailed anatomy.

Figure 6.1A shows an axial slice of a CT scan with stimulation targets highlighted. In this case, we are targeting the spinal cord and the intervertebral foraminae, where the spinal nerve roots exit, bilaterally. A sagittal slice at the level of the intervertebral foraminae is shown in figure 6.1B, demonstrating the placement of targets in the intervertebral foraminae at three spinal levels. Complex combinations

of targets are possible in order to assess the relative activation of each structure, for example to determine the selectivity of stimulation in order to avoid off-target effects.

Once the target anatomy has been defined, these segments can be represented as a set of points in 3D space. Figure 6.1C shows a 3D model of the spinal cord and the bilateral intervertebral foraminae at one spinal level along with a dorsoventral electrode montage in order to visualise the anatomic relationship between our target structures and the stimulating electrodes.




Figure 6.1: Targets for axon trajectories. **A** Axial slice of a CT scan with the spinal cord and bilateral intervertebral foraminae, containing traversing spinal nerve roots, highlighted as stimulation targets. **B** Sagittal slice of a CT scan showing the intervertebral foraminae at three spinal levels highlighted. Complex combinations of targets can be used to investigate selectivity and off-target effects. **C** 3D model of spinal cord and intervertebral foraminae targets shown in relation to the dorsoventral stimulating electrodes. Trajectories can be generated which pass through these structures and beyond. Complex target anatomy can be visualised in 3D space to clarify relationships between structures.

Target structures can be explicitly represented as a set of points defining the target volume. Figure 6.2A shows the spinal cord and bilateral intervertebral foraminae at three spinal levels plotted as point clouds.

Using this point cloud representation of target structures, we can identify anatomically realistic axon trajectories by performing an eigendecomposition of the point cloud, then sequentially stepping along the primary eigenvector and selecting points which maintain a fixed proportional distance from the extrema along the orthogonal axes from a seed point. An example trajectory along the spinal cord, superimposed in red, is shown in figure 6.2B. This automatically derived trajectory accurately follows the anatomic course of the target structure.

This method performs equally well for complex structures with orientations that would make manual drawing of trajectories challenging, such as intervertebral foraminae. Figure 6.2C shows an intervertebral foramen point cloud with an automatically derived trajectory superimposed in red. This accurately represents the course of spinal nerve roots in the foramen, derived automatically from the segmented clinical imaging. Many trajectories can then be automatically generated which follow defined structures and beyond; here, passing from adjacent to the spinal cord to beyond the intervertebral foramen laterally.




Figure 6.2: Target anatomy as point clouds. **A** Point cloud representation of the spinal cord and bilateral intervertebral foraminae, containing traversing nerve roots, at three spinal levels. Target structures can be represented as point clouds. **B** Point cloud representation of the spinal cord with an example axon trajectory superimposed in red. Point cloud representations of targets allow automatic derivation of anatomically realistic axon trajectories by performing eigendecompositions of the point clouds. **C** Point cloud representations of the spinal cord and a single intervertebral foramen with an example axon trajectory superimposed in red. The eigendecomposition method allows anatomically realistic trajectories to be determined even in complex anatomic regions.

Axon trajectories can then be transformed back into the original imaging space. This allows us to visualise the trajectories in the context of detailed anatomy and to represent them in a coordinate system that allows them to be easily coupled to the results of the finite element model. Figure 6.3 shows multiple example trajectories within the spinal cord, visualised in the axial, sagittal and coronal planes. Using this approach, arbitrary numbers of anatomically realistic axon trajectories can be automatically derived from clinical imaging.

This allows targeting of complex anatomical structures while removing the need for simplifications, idealised geometries or time-consuming processing for individually, manually adding trajectories. This opens the possibility for automatic




Figure 6.3: Trajectories on imaging. Axial, sagittal and coronal slices showing multiple axon trajectories within the spinal cord on imaging, determined by eigendecomposition of the point cloud representation of the segmented spinal cord. This allows straightforward derivation of realistic trajectories for coupling detailed physiological models of axon function to models of detailed anatomy.

coupling of complex anatomical and physiological models for investigating the effects of neuromodulation.

6.3.2 Coupling of models

By mapping each point of a biophysically realistic axon model onto a point along one of these trajectories, we can map the axon models into the complex anatomical model. This allows us to couple a detailed model of axonal physiology to a detailed anatomical model of our target structures. We can then link the results of the finite element model to the biophysical model, and therefore can compute the physiological response to electrical stimulation. We can therefore combine detailed models of anatomy and physiology to assess the response to electrical stimulation.

Figure 6.4A shows the membrane potential recorded at 20 consecutive nodes of Ranvier of a 59cm length, $11.5\mu\text{m}$ axon positioned along a trajectory within the spinal cord, extending from the superior to inferior extents of the image volume. Electrical stimulation is applied at $t = 0$. An action potential is initiated at node 10, and can be seen to propagate both orthodromically and antidromically along adjacent nodes of Ranvier.

These models can be used to analyse the response of target structures to stimulation. Figure 6.4B shows a raster plot of the responses of 100 axons placed within the spinal cord in response to non-invasive stimulation. The lower plot

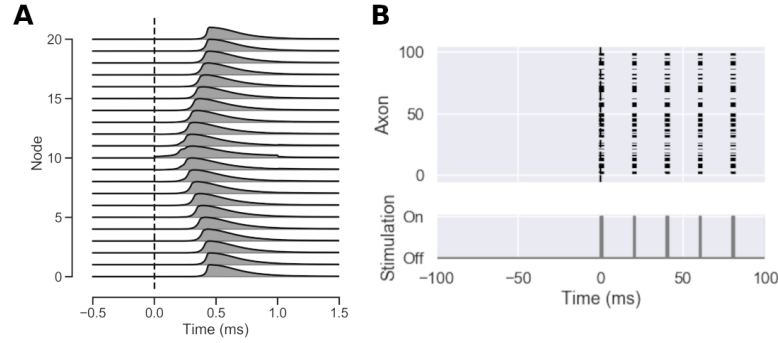


Figure 6.4: Response to stimulation. **A** Membrane potential in 20 consecutive nodes of Ranvier in a 59cm, $11.5\mu\text{m}$ spinal cord axon. Electrical stimulation is applied at $t = 0$. An action potential is initiated at node 10 and propagates bidirectionally along the adjacent nodes. **B** Raster plot of the response to stimulation in 100 spinal cord axons. The upper plot shows a black dot when an action potential is initiated in the corresponding axon. The lower plot shows the timing of stimulation. There are bursts of activity within the spinal cord following the application of non-invasive stimulation.

shows the timing of stimulation pulses, while the upper plot shows a raster plot of the response to stimulation, where a black dot indicates an action potential at that time generated in the corresponding axon on the y axis. There are bursts of activity throughout the spinal cord following each burst of stimulation. By coupling finite element and biophysical models, we can get detailed insights into the response to electrical stimulation in individuals.

Using this approach, we can derive highly detailed, biophysically realistic measures of the effect of electrical stimulation in detailed models by coupling these axon models to finite element models of detailed anatomy. However, care must be taken to ensure these are interpreted appropriately. Automatically generated trajectories and models may fail to take into account biophysically important features of the local microanatomy and axon branching structure. These isolated axons may need adaptation to the specific biophysical problem at hand in order to ensure reliable results are produced.

6.3.3 Strength-duration curves

By applying numerical optimisation methods to the coupled finite element and biophysical models, we can compute the minimum current required to produce

activation in our target structure for a given electrode position using a binary search approach. This allows us to calculate the stimulation threshold for our neural target.

Figure 6.5A shows the results of this binary search algorithm as it converges on the threshold current. After eight iterations, the numerical method has determined that the minimum current required to produce activation is 45.1mA. It is possible to quickly produce individualised thresholds for activating specific target structures using coupled computational models.

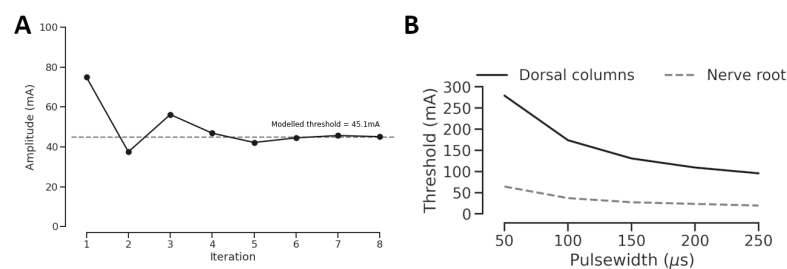


Figure 6.5: Determination of strength-duration curves using binary search. **A** Convergence of the binary search algorithm on a threshold value. The model is run repeatedly with different stimulation amplitudes and the response measured - in this case, activation of $>70\%$ of 100 axons distributed throughout the dorsal root volume with diameters of $5.7\mu\text{m}$ - $11.5\mu\text{m}$. At each iteration, the algorithm converges towards the true threshold until this is found within a given tolerance. **B** Strength-duration curves for dorsal column and nerve root axons measured using a binary search method. The process of identifying the threshold for activation is repeated for each pulse width in order to characterise the strength-duration relationship. Modelling and optimisation techniques can be used to derive physiologically meaningful quantities in individualised models.

This approach can be extended to characterise the responsiveness of target structures to stimulation. By computing the minimum current required for activation for a range of pulse widths, we can generate a strength-duration curve for the target structure. This represents a physiological characteristic of the target structure and allows us to investigate the stimulation parameters required to produce our desired effect.

Figure 6.5B shows computed strength-duration curves for producing activation in populations of axons in the spinal cord and in the nerve roots within the intervertebral foraminae. Notably, the amplitude required to produce activation within the nerve roots is consistently much lower than that to produce activation in the spinal cord. This suggests that the effects produced by non-invasive stimulation

are mediated by activation of spinal nerve roots rather than by direct activation of cord structures. However, care must be taken when interpreting these results as the model has not been optimised to account for local microanatomy and axon branching structures. This may result in artefactual activation at low amplitudes due to edge effects in unconnected axon terminals within the stimulated area.

Automatic coupling of finite element and biophysical models allows us to interrogate the effects of electrical stimulation in a detailed model of individual anatomy and physiology. This allows us to investigate individual-specific stimulation thresholds and to examine the nature of the neural elements mediating the effects of stimulation.

6.4 Discussion

6.4.1 Linking anatomy and physiology

These results demonstrate that it is possible to use automatically coupled finite element and biophysical models to examine in detail the effects of varying stimulation parameters on the resulting physiological effect. This further demonstrates the ability to rapidly compute individually optimised stimulation parameters by simulating the effects of stimulation with a variety of parameters, allowing us to produce individually optimised stimulation for targeting specific neural structures using standard clinical imaging.

The electric field produced by stimulation is influenced by individual differences in anatomy. The differences in the resulting electric field due to these interindividual anatomical variations have a significant impact on the effects of stimulation, making the stimulation parameters needed to produce a response in any individual difficult to predict. Accounting for these individual differences in anatomy in models of electrical stimulation improves the performance of these models [49, 50]. The use of detailed individual anatomy in computational models is therefore important to produce individually optimised treatment predictions.

However, even with highly detailed models of the electric field produced by stimulation in an individual, the physiological response to stimulation is unpredictable.

Simplified metrics that link electric field parameters to neural responses exist, but are limited [77, 88]. The most reliable models are validated multi-compartment cable models of axons. These models accurately represent neural physiology and have been validated in the context of predicting the response to external electrical stimulation [11, 18, 19, 49, 50].

Coupling of finite element models of the electric field and biophysical models of the neural response to stimulation are the gold standard for combining detailed models of anatomy and physiology for predicting the response to stimulation. The application of this approach has so far been limited by the need for either simplified, idealised anatomy [45, 47, 48], limiting widespread applications or use for individualised predictions, or complex, laborious workflows involving manual positioning of axon trajectories [18, 50], limiting practical applications.

Our results demonstrate a method for automatically combining highly detailed models of individual anatomy and physiology. The techniques we have developed provide a means for creating realistic computational models of individuals directly from clinical imaging in a way that is automated, accurate and scalable. This makes this approach uniquely suitable for application in a clinical context to inform decision-making in neuromodulation and for the investigation of novel neuromodulation techniques. However, care must be taken that relevant microstructural features and branching patterns are taken into account when applying these models in order to ensure results are reliable.

6.4.2 Automated axon mapping

We demonstrate that biophysical axon models can be mapped automatically onto realistic trajectories within target anatomical structures on clinical imaging. This allows detailed physiological models to be "placed" in 3D space within complex anatomy, facilitating the modelling of complex responses to stimulation. This automated coupling step is a key advance in the use of individualised computational models for predicting the effects of intervention.

Existing techniques for achieving this mapping of axonal trajectories into complex anatomy are limited in their scope for widespread use or application in a clinical context. Manual definition of axon positions has the advantage of being highly user-controlled, providing a great degree of flexibility and accuracy regarding placement of known axon trajectories. This makes this approach valuable for the investigation of neuromodulation mechanisms, where a single highly accurate model can provide useful insights. However, the difficulty and time investment required to produce accurate axon trajectories makes this approach unsuitable for simulating large populations within a region or for individually optimising stimulation.

Other automated methods for determining trajectories are potentially valuable to bridge this gap and make automated model coupling possible. Tractographic methods [121] can be used to define axon locations. This can then be combined with the coupling techniques we describe to automatically link models. However, this approach is limited by the need for specialised imaging modalities rarely used outside the brain, as well as the documented limitations of current tractography methods.

The method we describe here is based on the eigendecomposition of segmented clinical imaging and derivation of trajectories in the new eigenbasis, followed by transformation back to the original image space. Axon models are then mapped into the anatomical space by projecting each point of the axon model onto this trajectory.

This method of determining axon trajectories has the advantage of not requiring any specialist imaging, allowing it to be applied to any region and imaging modality. This makes application to new areas and integration into existing clinical pathways straightforward. Further, it is relatively computationally simple and fully automated, removing potential barriers to its application.

This approach is, however, potentially limited in its application to complex axon trajectories. As it relies on determining a trajectory along the target anatomical region's primary eigenvector, a target which includes significant turns, i.e. greater than 90 degrees, can result in unrealistic trajectories. This is not an issue for axons within the spinal cord or most peripheral nerves, however in some cases it may limit application of this method. This issue can be overcome by dividing complex

trajectories into multiple shorter trajectories around each significant deviation and then coupling these, though this then adds to the complexity of applying this approach. Such complex trajectories, such as long association fibres within the brain, may be better approached by tractography, with subsequent mapping of axons onto the resulting trajectory using the techniques we describe here.

Additionally, the limitation of this approach to modelling gross anatomical features needs to be considered. Modelling isolated axons using anatomical regions is useful for many applications, and is potentially generalisable to domains outside of the spinal cord such as peripheral nerve stimulation. However, local anatomical features can have an important impact on the effects of stimulation.

By generating trajectories based on defined regions, potentially important nuances can be lost. For example, the angle of entry of dorsal roots into the spinal cord [122], which influences the response to stimulation [49] is not directly represented in this way. Similarly, the intradural course of these fibres is not represented as well as with automated methods based on specific principles for this fibre type [117].

These issues need to be taken into account when interpreting the results of these models. Further, the use of isolated trajectories needs to be considered. In the absence of branching within the dorsal columns, the results of stimulation may be inaccurate [49, 118]. Apparent activation may be due to artefactual activation at axon terminals. While useful, adaptation of the specifics of the model to the biophysical problem being assessed needs to be considered in order to ensure results are accurate.

6.4.3 Modelling the effects of stimulation

By modelling the interaction between electric fields produced by stimulation and detailed anatomy and the effects of this on biophysically realistic axons, we can accurately predict the response to stimulation. This approach has been applied to the development of novel neuromodulation systems and the investigation of the mechanisms of action of existing systems [9, 11, 14, 18, 19, 52].

We have shown that we can automatically create detailed individualised models for performing this kind of evaluation. This significantly reduces the technical barriers to carrying out this kind of evaluation by allowing simulation on highly detailed anatomy directly from clinical imaging with automatic coupling to detailed physiological models. This potentially allows the evaluation of novel neuromodulation approaches on detailed models of arbitrary anatomical regions. The applications of spinal and peripheral neuromodulation are constantly increasing, and this approach offers an attractive means of rigorously evaluating the effects of these stimulation approaches and of developing new approaches.

6.4.4 Individually optimised neuromodulation

The ability to assess the effects of stimulation in detailed models can be further extended to determine the effects of electrical neuromodulation in specific individuals. We show that it is possible to determine the physiological effects of stimulation to inform individually optimised stimulation parameters in order to achieve activation of specific target structures in a given individual.

The clinical goal of stimulation can be expressed as a biophysical objective function. This represents a quantity that we aim to maximise, such as the activation of axons within a specific region. More complex quantities can also be constructed, for example by including cost functions in order to maximise activation of a target while minimising activation of other, off-target structures.

We can then apply simple optimisation methods in order to allow derivation of physiologically and clinically meaningful measures from these individualised computational models. We show that it is straightforward to measure a strength-duration curve for a target neural structure in an individual. This provides a useful physiological characterisation of the responsiveness to stimulation of a target structure in that individual, allowing us to calibrate our stimulation parameters.

This approach also allows us to investigate the level of selectivity and the "therapeutic window" of stimulation parameters for an individual. We show that it is possible to selectively activate the nerve roots traversing the intervertebral

foramen; this approach further allows us to determine the electrode positions and stimulation parameters required to activate this target while avoiding activation of other nerve roots or the spinal cord.

This represents an approach to individualised, "precision" neuromodulation that is based on a known physical system with explicit representations of anatomy and physiology, where clinical intent is captured in a physiologically meaningful objective function. This provides a truly personalised model for individualised optimisation. This is in contrast to other approaches which have gained popularity in "precision" medicine, such as deep learning techniques based on interpolation within a high-dimensional latent space learned from large population data. The method we present here avoids the need for function approximations linking stimulation parameters and clinical outcomes and the potential biases this introduces. We do this by explicitly modelling the link between parameter selection and clinical outcome as a mathematical system based on physics and patient anatomy, which we then solve using well-established techniques from numerical optimisation.

This represents a clinically feasible approach to individualised optimisation of neuromodulation therapies. We demonstrate a method for determining optimal stimulation parameters which requires only routine clinical imaging and which is explicitly individualised and interpretable, with no need to defer to "black box" functions derived from population data. This approach is automated, removing the need for time-consuming clinical workflows and opening the possibility for large-scale applications in a clinical context. The application of these methods requires some computational resources, however the requirements are well within the limits of what is currently used in clinical practice for clinical tractography and deep brain stimulation planning. This therefore represents a potentially viable approach for clinically relevant optimisation of neuromodulation treatments.

6.4.5 Conclusion

We demonstrate the ability to automatically couple finite element models of the electric field produced by stimulation to biophysically realistic axon models. The

methods we have developed do not require any specialised imaging, are automated and are generalisable to any anatomical region. This approach allows for the simulation of the effects of stimulation in individuals with arbitrary electrode locations and stimulation parameters, while accounting for detailed individual anatomy and physiology. Care must be taken to ensure any necessary anatomical or biophysical features are accurately represented before interpreting resulting in order to avoid inaccuracies due to edge effects or loss of microstructural features.

This represents an important advance in the development of individually optimised neuromodulation. By combining models of the electric field produced by stimulation with models of the physiological effect of that stimulation, we can assess how changes in our selected parameters affect neural physiology. By then expressing our clinical goals in terms of a desired physiological effect, we can optimise our chosen stimulation parameters in order to find the set of parameters that allows us to achieve this goal. This allows us to investigate the physiological effects of neuromodulation in great detail, and opens the possibility for individually optimised neuromodulation therapies by numerically optimising for a desired physiological effect using detailed computational models derived directly from standard clinical imaging.

In the context of a broader system for model-driven stimulation, this allows us to derive a function $f_{V \rightarrow V_m}$, which links the electric field produced by stimulation to the membrane potential of a target neural structure by solving a multi-compartment cable model with Hodgkin-Huxley dynamics. This allows us to assess the physiological effect of stimulation.

We can then extend this by defining a function $f_{V_m \rightarrow A}$, which links the membrane potential of our neural target to some measure of activation A , capturing our clinical intent in terms of biophysics. By using a function $f_{\theta \rightarrow V}$, which computes the electric field V produced in an individual for a given set of stimulation parameters θ , we can derive a composite function $f_{\theta \rightarrow A}$, which links our selection of stimulation parameters to our measure of activation. By utilising techniques from numerical optimisation, we can then determine the optimal values of θ in order to maximise A , i.e. the stimulation parameters which best achieve our clinical goal in the given individual.

This represents a major step towards model-driven neuromodulation. By combining the steps from deriving an individualised finite element model, to solving for the electric field produced, to coupling with biophysics and optimisation with physiologically meaningful measures, we can derive a fully personalised model of the effects of neuromodulation on a given individual. By combining the steps from deriving an individualised finite element model, to solving for the electric field produced, to coupling with biophysics and optimisation with physiologically meaningful measures, we can derive a fully personalised model of the effects of neuromodulation on a given individual. The integration of these techniques into a single system for individualised modelling of neuromodulation using highly detailed models is described in appendix B.

This approach allows us to automatically optimise stimulation parameters to achieve our clinical goals in an individual using a highly detailed individualised model. We show that this allows us to rigorously evaluate the effects of electrical stimulation on neural tissue and to optimise therapies for individual patients.

Part IV

Application

7

Feasibility and mechanism of transcutaneous spinal cord stimulation: a computational study

Contents

7.1	Introduction	135
7.2	Methods	137
7.2.1	Electrode montages	137
7.2.2	Biophysical models	138
7.2.3	Measuring activation	140
7.2.4	Parameter sweeps	142
7.3	Results	144
7.3.1	Spinal cord activation	144
7.3.2	Activation of foraminal structures	147
7.3.3	Optimisation of electrode montage	149
7.4	Discussion	151
7.4.1	Efficacy and mechanism	151
7.4.2	Targeting foraminal structures	151
7.4.3	Conclusion	153

7.1 Introduction

Stimulation of the spinal cord using implanted epidural electrodes has recently shown significant promise in augmenting rehabilitation to restore naturalistic locomotion

after spinal cord injury [7–10]. This technique has the potential to offer new opportunities for functional restoration following significant neurological injuries. However, the invasive nature of this approach, requiring surgical implantation of epidural electrodes, has limited more widespread application of this technique. While implantation of epidural stimulation electrodes is well established in the treatment of chronic pain [96] and is typically well tolerated, this remains a surgical procedure which carries risk [97, 98]. This risk may be further raised in patients following spinal cord injuries, who may have distorted anatomy.

Further, it appears that benefits gained during stimulation-assisted rehabilitation are maintained even when stimulation is turned off [9, 10]. This means that it is not necessary to have persistently present electrodes with chronic stimulation at all times; rather, stimulation only needs to be applied during the period of rehabilitation to augment recovery. This negates many of the benefits of a chronic implanted system, which can be left in situ for chronic stimulation following the initial insertion. Alternatives have therefore been investigated to augment rehabilitation while reducing the risks of invasive electrode insertion.

Transcutaneous spinal cord stimulation has been investigated for augmenting rehabilitation to aid functional recovery after spinal cord injury [35, 38–41, 43, 99]. However, questions remain regarding the validity of this approach. While many studies show a positive response, these are uniformly small, poorly controlled studies with variable outcome parameters and little rationale for the electrode montages and stimulation parameters used. There has been little mechanistic investigation into the feasibility of transcutaneous spinal cord stimulation or principled exploration of the best means of delivering this therapy.

Despite these limitations, clinical trials are currently ongoing examining the application of transcutaneous spinal cord stimulation to functional recovery after spinal cord injury [123]. These aim to investigate, in a more rigorous way, the effects of stimulation on rehabilitation outcomes. This is an important step towards better evaluation of these potential therapies. However, the selection of electrode montages and stimulation parameters is still not based on any specific understanding

of the underlying mechanisms or therapeutic goals. A more rigorous approach to the evaluation of transcutaneous spinal cord stimulation has the potential to inform pilot clinical studies in order to maximise the potential benefit from this promising neuromodulation modality.

Using the techniques developed in the preceding chapters, we can model the effects of transcutaneous spinal cord stimulation using detailed models of anatomy and physiology. We can apply these methods to determine if existing methods of transcutaneous stimulation are likely to produce an effect on the spinal cord, whether any other neural structures are affected and what electrode positions and stimulation parameters are best suited to optimising target activation. This provides a means of using our knowledge of electrical neuromodulation and its effects to investigate transcutaneous spinal cord stimulation and to use these results to inform the application of this technique to functional restoration.

Using this approach, we show that transcutaneous stimulation is unlikely to produce direct activation of the spinal cord. Rather, we show that any effect of transcutaneous stimulation is likely to be mediated by activation of segmental nerve roots. Further, we show that the optimal parameters for targeting nerve roots are different to those applied in trials to date. We show that the application of optimised electrode positions for targeting segmental nerve roots provides a therapeutic window of selectivity, allowing activation of neural structures at specific spinal levels and lateralities. This allows spatially specific patterns of stimulation, with potential benefits for applications to functional restoration.

7.2 Methods

7.2.1 Electrode montages

A detailed model of the thoracolumbar spine was created using the methods outlined in chapter 4. A 5cm x 5cm cathode was simulated in the midline over the T10/T11 intervertebral space with a corresponding 5cm x 5cm anode at the same level in the midline anteriorly (figure 7.1).

The electric field produced by stimulation was computed by solving the system of equations describing the behaviour of electric fields in tissue using the methods outlined in chapter 5. Boundary conditions were applied to the electrode-tissue interfaces to simulate electrical stimulation through the applied electrodes:

$$\begin{aligned}\nabla \cdot \sigma \nabla V &= 0 && \text{in } \Omega, \\ \frac{\partial V}{\partial n} &= 0 && \text{on } \partial\Omega, \\ V &= V_e && \text{on } \Gamma_A, \\ V &= 0 && \text{on } \Gamma_R\end{aligned}$$

where V is the electric potential produced by stimulation, σ is the tissue conductivity, Γ_A and Γ_R are the surfaces of the active and return electrode respectively, Ω is the domain modeled, and $\partial\Omega$ is the external skin boundary. V_e is the electric potential applied at the active electrode; this is computed such that $\int_{\Gamma_A} \frac{\partial V}{\partial n} = J_{applied}$, so that current $J_{applied}$ is delivered via the active electrode.

The electrical properties of the tissue within the detailed anatomical model are shown in table 7.1 [94].

Table 7.1: Physical properties of tissue types for constant current stimulation

Tissue type	Conductivity (S/m)
Bone	$2E^{-2}$
Soft tissue	$2.02E^{-1}$
Lung	$3.89E^{-2}$
Skin	$2E^{-4}$
CSF	2
Spinal cord (transverse)	$8.30E^{-2}$
Spinal cord (longitudinal)	0.6

This allowed us to model the effects of electrical stimulation in a highly detailed anatomical model.

7.2.2 Biophysical models

This model of electrical stimulation was coupled to a detailed biophysical model using the methods outlined in chapter 6.

The McIntyre-Richardson-Grill model of myelinated axons was used [76]. Detailed multicompartment models were solved in NEURON [92] using the biophysical parameters outlined in table 6.1 [76].

To evaluate the effects of transcutaneous spinal cord stimulation, target structures were placed in the spinal cord and in the nerve roots at three levels bilaterally as shown in figure 7.1. This allowed for the physiological effects of stimulation to be assessed in detail, and for the selectivity of stimulation on spatially separated neural elements to be investigated. Figure 7.1A shows the spinal cord and foraminal targets in axial section, while figure 7.1B shows these targets in a paramedian sagittal section demonstrating targeting of nerve roots at three consecutive levels. Figure 7.1C shows the arrangement of electrodes used for stimulating the target structures.




Figure 7.1: Targets for axon trajectories. **A** Axial slice of a CT scan with the spinal cord and bilateral intervertebral foraminae, containing traversing spinal nerve roots, highlighted as stimulation targets. **B** Sagittal slice of a CT scan showing the intervertebral foraminae at three spinal levels highlighted. Complex combinations of targets can be used to investigate selectivity and off-target effects. **C** 3D model of midline stimulating electrodes at the T10/T11 level demonstrating the arrangement used for initial testing.

As outlined in chapter 6, anatomically appropriate axon trajectories were automatically determined within each target structure and biophysical axon models were adapted and automatically mapped onto each trajectory.

100 trajectories were generated within each target volume. Axon models were then mapped onto each trajectory. This produced 100 axons within the dorsal columns and within each of the modelled nerve roots. Axon diameters were randomly distributed between $5.7\mu\text{m}$ and $11.5\mu\text{m}$.

The dorsal column axons extended from the superior limit of the modelled domain to the inferior limit of the domain. The average axon length was 323mm +/- 2mm.

Dorsal root axons extended medially from the dorsal root entry zone, adjacent to the spinal cord, to beyond the lateral extent of the intervertebral foramen. The average axon length for dorsal roots was 57mm +/- 5mm.

In order to account for important structural features for determining root activation which are not automatically incorporated into volume-based trajectories, the root models were adjusted. From its termination at the dorsal root entry zone, the dorsal root trajectory was extended medially and superiorly to intersect with the spinal cord volume. A new dorsal column axon, with trajectories and characteristics identical to those above, was generated for each dorsal column axon. The root axon was then extended from its entry into the spinal cord to join its corresponding dorsal column axon. The biophysical axon model were then connected at this point. A schematic of the dorsal root axon model is shown in figure7.2.

This allowed explicit representation of the anatomy of the root entry zone and of the branching patterns of the dorsal roots within the spinal cord. These branch fibres in the dorsal columns were not evaluated as part of the assessment of dorsal column activation; they were generated to represent root branching structures and avoid artefactual activation and terminal nodes in axons ending in the root entry zone.

These were then coupled to the electric field model and the membrane potential response assessed by solving the corresponding system of equations.

7.2.3 Measuring activation

We define our target structures as activated by stimulation if more than 70% of axons within the target produce an action potential following stimulation. This can be expressed mathematically:

$$A_{activated} = \begin{cases} 1, & \text{if } \max(V_m) > 0 \text{ in } > 70\% \text{ of axons} \\ 0, & \text{if } \max(V_m) > 0 \text{ in } < 70\% \text{ of axons} \end{cases}$$

Using the methods outlined in chapter 6, we can then apply a binary search algorithm to identify the lowest applied current that produces activation within our




Figure 7.2: Schematic of dorsal root axon model. A trajectory is automatically determined based on the anatomical volume which passes from the root entry zone through the spinal canal and intervertebral foramen following an anatomical course. It passes out of the intervertebral foramen and terminates in the thorax. The termination in the root entry zone is extended superiorly and medially to intersect the spinal cord. Within the spinal cord, the dorsal root axon joins a rostrocaudal dorsal column axon and branches here.

target structures. This allows us to identify the threshold current for activation in these anatomically and physiologically detailed models. Threshold were identified between 0mA and 150mA to an accuracy of within 1mA.

The process of threshold identification was repeated over a range of pulse widths in order to characterise a strength-duration curve for each target structure. This allowed the physiological effects of stimulation parameters on our target structures to be investigated in a clinically meaningful way.

7.2.4 Parameter sweeps

Stimulation parameters were systematically altered in order to examine the effect of alterations in electrode position and stimulation settings on the response to stimulation.

Strength-duration curves were characterised for all targets with electrodes in the midline position. This allowed the relative effects of stimulation using a standard electrode montage on our target structures to be investigated. We could then characterise the structures activated by stimulation, the relative spatial selectivity of transcutaneous stimulation, and the therapeutic windows of stimulation parameters for achieving targeted activation of specific structures.

We then varied the position of electrodes in the model in order to assess the effects of alterations in electrode position on the response to stimulation and the level of selectivity that can be achieved.

From the starting position over the T10 spinous process shown in figure 7.1C, the active electrode position was varied in the rostrocaudal and mediolateral directions. The return electrode remained fixed in position in the midline anteriorly opposite the starting position.

The stimulating electrode was positioned from 6cm above to 6cm below this starting point at the T10 spinous process in 1cm steps. The position was varied laterally from 8cm to the left to 8cm to the right of the starting position in 1cm steps. This produced a total of 221 electrode positions within the evaluated area. This is shown in figure 7.3.

In order to determine the optimal electrode position for a given individual to activate our target structure, we want to find the electrode position that causes activation with the lowest applied current. To do this, we set our measure of activation $A = -J_{threshold}$, i.e. our measure of "activation" is high where the threshold is low. Using the methods described previously, we can then define a function $f_{\theta \rightarrow A}$ which returns this activation measure for any given electrode position θ .

We then optimise for the electrode position with the highest activation:




Figure 7.3: Range of electrode positions. The electrode was centred over the T10 spinous process. Its position laterally was varied up to 8cm in either direction in 1cm steps. Its rostrocaudal position was varied up to 6cm in either direction in 1cm steps. The results of stimulation at every position in this grid was evaluated.

$$\theta_{optimal} = \arg \max_{\theta \in \theta'} f_{\theta \rightarrow A}(\theta)$$

where θ is the set of electrode and stimulation parameters, where only the electrode position is allowed to vary, θ' is the set of all clinically allowable electrode positions and A is our measure of activation as defined above. By solving this numerically, this allows us to determine the electrode position for this individual that produces activation of our target structure at the lowest required current. In this way, we can carry out individually optimised targeting of specific structures automatically, allowing us to solve clinically meaningful problems.

We used this approach, in line with the methods described in chapter 6, to

identify whether selective stimulation could be achieved through activation of one target without activating others. The size of these regions of selective activation were characterised.

This variation in electrode position was used to identify the optimal electrode positions for activation of specific target structures, i.e. the electrode positions where the current required to produce activation for a given target was lowest. Strength-duration curves were then characterised for each target at these optimal positions. This provided a detailed assessment of the level of selectivity of stimulation that it is possible to achieve with transcutaneous stimulation and the extent of the therapeutic window for stimulation.

7.3 Results

7.3.1 Spinal cord activation

Stimulation with a midline electrode montage at T10/T11 can produce activation of axons within the spinal cord. However, less current is required to activate the exiting nerve roots at this level. The strength-duration curves for the spinal cord and for the local nerve roots are shown in figure 7.4.




Figure 7.4: Strength-duration curves for axons in the spinal cord and the left T10 nerve root with a midline T10 cathode, calculated using a binary search method. The nerve roots are activated at lower stimulation amplitudes, suggesting that they are the first structures to be activated by transcutaneous stimulation.

For all pulse widths, the amplitude required to produce activation of the spinal cord remains higher than the amplitude required to produce activation of the nerve roots. This suggests that the first structures activated by transcutaneous stimulation with a midline montage are the nerve roots, rather than the spinal cord itself.

As the position of the cathode is varied, the current required to produce activation changes. By using the threshold current as an optimisation target, we can allow the electrode position to be varied in order to compute the electrode position that requires the lowest stimulation amplitude to produce activation in the spinal cord and in the nerve roots. This allows us to determine the individually optimised electrode positions to activate target structures in a given individual, taking into account their detailed anatomy.

The cathode position was varied while the anode position was held constant. Figure 7.5A shows 3D models of the vertebral column with variations in the cathode position, while the ventral anode is held fixed in the midline at the T10/T11 intervertebral space. Figure 7.5B shows a contour map of the potential and a vector field of the electric field produced by stimulation superimposed on an axial slice through the lower border of T10 for each of the electrode arrangements shown in figure 7.5A. As the cathode position is varied over the dorsal surface, the response to stimulation varies. The electric field produced by stimulation in that position is computed and used to calculate the threshold required to activate the spinal cord and the nerve roots with a cathode in that position.

By calculating the electric field produced by stimulation in each electrode position and coupling this to a biophysical model to compute the threshold current for activating the spinal cord and the nerve roots, we can express the activation threshold for each target structure as a function of electrode position. Figure 7.6 shows a contour map of the minimum current required to produce activation of the spinal cord and the left-sided T10 nerve roots for cathode positions around the T10 vertebral level.

The left side of figure 7.6 shows the variation in the threshold for producing activation in the spinal cord as the cathode position is altered. The current required

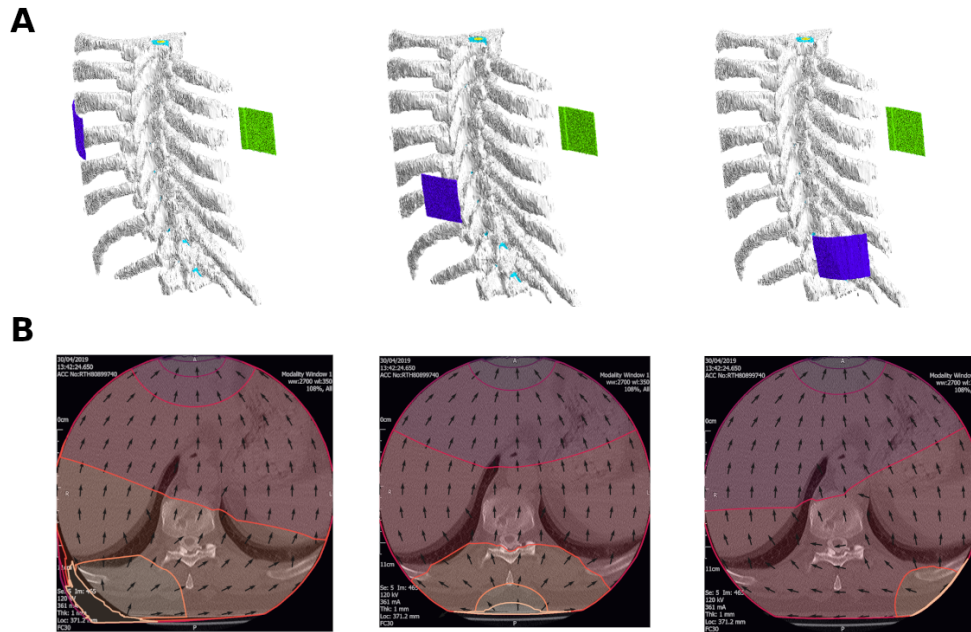


Figure 7.5: Variations in electrode position. **A** 3D models of the vertebral column with variations in electrode position. The ventral anode is held fixed at the level of the T10/T11 intervertebral space. The dorsal cathode's position is allowed to vary freely over the skin surface. **B** The potential (contours; V) and electric field (vectors) produced by stimulation for each of the electrode arrangements above. The electric field varies significantly as a function of electrode position. Electrode position must therefore be optimised in order to maximally activate target structures.

to activate the dorsal columns is roughly symmetrical across the midline, as there is no strong lateralisation of the effect of stimulation. The slight asymmetries are likely a result of asymmetries in the anatomy used to generate the model and due to rotation of the patient and gantry angle during acquisition, resulting in the model not being perfectly symmetrical.

There is a gradient in the current required to produce activation in the rostrocaudal axis, likely accounted for by the natural thoracic kyphosis resulting in variation in the amount of subcutaneous tissue and the thickness of the CSF layer between the cathode and the target structures at different rostrocaudal levels.

Notably, in all cathode positions the nerve roots are activated at a lower amplitude than the spinal cord. There is no electrode arrangement that causes primary activation of the spinal cord before activation of the nerve roots.




Figure 7.6: Activation threshold as a function of electrode position. The current required to produce activation of the spinal cord is roughly symmetrical across the midline and has a rostrocaudal gradient likely accounted for by the natural thoracic kyphosis. The activation threshold for the left-sided T10 nerve roots shows a region of low threshold 4cm left of the midline at the T10 level, with a rapid increase in threshold in the mediolateral and rostrocaudal directions, suggesting the possibility of achieving unilateral activation of nerve roots at a single segmental level.

7.3.2 Activation of foraminal structures

The nerve roots within the intervertebral foraminae are activated by transcutaneous stimulation. Even with midline stimulation, figure 7.4 shows that activation occurs at lower currents than spinal cord axons.

The right side of figure 7.6 shows the variation in the threshold for activation of the left T10 nerve root as an example of the variation in current required for activation as the cathode position is changed. In contrast to the spinal cord, activation of the nerve roots shows strong selectivity in both mediolateral and rostrocaudal axes. There is a clear focus of low threshold four centimetres off the midline at the target vertebral level, with a rapid increase in threshold as we move away from this position. Similarly, there is a rapid increase in threshold as the cathode is moved away from the target level in the rostrocaudal axis. This suggests that, with properly selected stimulation parameters, it may be possible to produce selective activation of nerve roots unilaterally at a single segmental level, potentially overcoming some of the issues with the selectivity of non-invasive

stimulation methods.

Even with midline stimulation, the current required to activate nerve roots at other spinal levels is much higher than for roots at the level of the electrode. Strength-duration curves for the target nerve root and the nerve root at the levels above and below are shown in figure 7.7.




Figure 7.7: Strength-duration curves for axons in the left T8, T9 and T10 nerve roots with a midline T10 cathode, calculated using a binary search method. The T10 nerve root is activated at a lower amplitude than the nerve roots at the levels above and below, suggesting the possibility for selective activation of segmental nerve roots.

This demonstrates a clear window of selectivity in activation of a single segmental nerve root in the rostrocaudal axis. This suggests that some selectivity in the stimulation of nerve roots may be possible. However, with midline stimulation, the nerve roots are activated bilaterally with similar amplitudes of stimulation, limiting selectivity in the mediolateral axis.

To further characterise the level of selectivity possible and to identify the regions that allow selective stimulation, we examined the spatial distribution of the difference between the threshold for target activation and the threshold for activation of off-target structures. Figure 7.8 shows the region where selective activation is possible and the difference between the current needed to activate an example target nerve root, the left T10 root in this case, and the lowest off-target threshold at each electrode position.

This shows that an area of selective nerve root activation exists. All areas with a positive value indicate that it is possible to produce selective activation of the target




Figure 7.8: Selectivity of nerve root activation. The colour map on the left highlights the area where the left T10 nerve root can be selectively activated without activating the contralateral nerve root or the roots at other levels. This shows a clear area where selective stimulation is possible. The colour map on the right indicates the difference between the threshold for activation of the left T10 nerve root, the target, and the threshold for activation of any off-target structures at each point. Positive values indicate a therapeutic window within which selective activation of a unilateral segmental nerve root can be achieved at that location using transcutaneous stimulation.

root with a cathode in this position, with no off-target activation. The value shows the size of the therapeutic window within which there is no off-target activation. Conversely, negative values show regions where off-target structures are activated before the target. This demonstrates that with careful electrode positioning with cathodes 4cm off the midline at the level of the target intervertebral foramen, it is possible to achieve selective activation of specific segmental nerve roots.

7.3.3 Optimisation of electrode montage

To improve our ability to target stimulation, we used the variation in the position of the cathode to identify the optimal electrode arrangement to achieve activation of each target structure. The optimal cathode positions to activate the spinal cord and the left T10 nerve root are shown in figure 7.9A.

The optimal electrode arrangement for activating the spinal cord is a midline arrangement. There is little advantage to more lateral arrangements for targeting the spinal cord. There is some rostrocaudal difference in threshold, but this needs to be traded off against the targeting of specific spinal levels. However, as shown previously, the spinal cord cannot be selectively activated regardless of electrode arrangement.

In contrast, the nerve roots are best targeted with an off-midline cathode. A region roughly 4cm from the midline at the level of the intervertebral foramen,

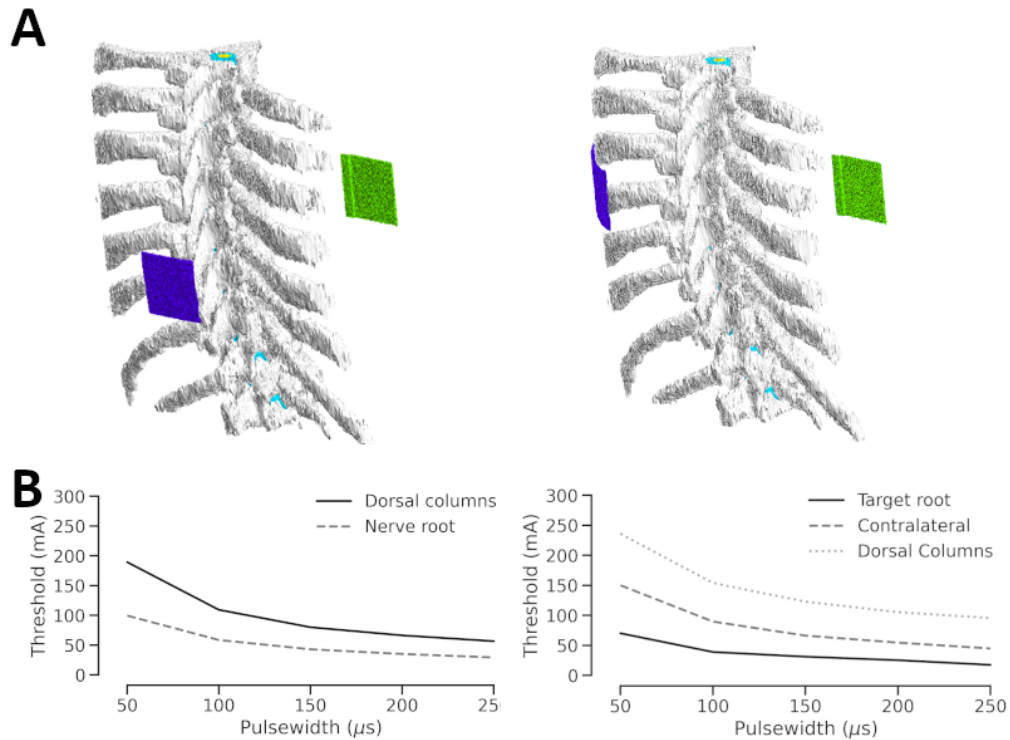


Figure 7.9: Optimised electrode positions for specific target structures. **A** 3D models of the vertebral column with optimised electrode position. The electrode arrangement on the left shows the midline montage optimised for activation of the spinal cord. The arrangement on the right shows the off-midline cathode position that is optimal for activation of a specific segmental nerve root. **B** Strength-duration curves for activation of target structures for optimised electrode positions. Even with optimised positioning, nerve roots are activated at lower thresholds than the spinal cord. In a position optimised for nerve root activation, there is a window of selectivity in both the rostrocaudal and mediolateral axes, allowing activation of segmental nerve roots unilaterally.

with roughly 2cm around this centrepont, allows for selective activation of the target nerve root with minimal off-target effects. This allows for selective targeting in the lateral and rostrocaudal planes.

The strength-duration curves at the optimal electrode locations were characterised as shown in figure 7.9B. The behaviour of the spinal cord montage remains unchanged from the initial midline arrangement shown in figure 7.4.

Figure 7.9B also shows the strength-duration curves for the target nerve root and off-target structures with the cathode in the optimal position. This shows that there is a clear therapeutic window of stimulation parameters for producing activation of the target without off-target effects. This allows us to identify an

electrode montage for specific activation of targeted nerve roots to facilitate spatially specific transcutaneous stimulation of neural elements.

7.4 Discussion

7.4.1 Efficacy and mechanism

Transcutaneous stimulation has drawn wide interest for its potential applications to rehabilitation following the success of epidural stimulation [7–10]. The effects of transcutaneous stimulation were assumed to be due to activation of spinal cord structures [43]. The technique has therefore been referred to as "transcutaneous spinal cord stimulation" or "noninvasive spinal cord stimulation", and electrode arrangements have designed largely in the midline with the goal of targeting the spinal cord in mind.

Our results suggest that, while activation of spinal cord structures is possible, nerve root activation occurs at much lower amplitudes. It is therefore likely that the effects of transcutaneous stimulation are mediated by activation of the spinal nerve roots rather than by activation of spinal cord axons. When stimulation is applied at very high amplitudes, direct activation of the spinal cord is possible, however in these cases the nerve roots will also be activated earlier in any case.

These results provide mechanistic evidence that transcutaneous stimulation of neural elements produces physiological effects. However, these effects appear to be mediated by a different mechanism than was previously thought. This has important implications for the design of electrode montages for current and future trials of transcutaneous stimulation, and may have implications for the level of selectivity that it is possible to achieve with transcutaneous stimulation.

7.4.2 Targeting foraminal structures

Our results suggest that activation of segmental spinal nerve roots is the primary mechanism of action of transcutaneous stimulation. These structures are consistently activated at lower thresholds than the spinal cord regardless of electrode arrangement. Any effects seen are therefore likely to be caused by this activation. Recent evidence

from epidural stimulation for rehabilitation suggests that this is also mediated by activation of nerve roots, with secondary activation of spinal reflexes [11, 18, 19]. This raises the possibility that transcutaneous stimulation may mediate its effects by the same mechanism as epidural stimulation for functional restoration.

Targeting spinal nerve roots rather than the spinal cord may allow greater selectivity in the regions activated. Activation of specific regions of the spinal cord is extremely challenging using transcutaneous stimulation due to the small size of the target structures relative to the electrodes, their close proximity to off-target regions and the large volume of intervening tissue. Spinal nerve roots are, however, spatially separated; this potentially lends itself to more targeted activation.

Our results demonstrate that with careful selection of electrode position, activation of specific target segmental nerve roots is possible without off-target effects. This is important as spatially patterned stimulation, with activation of specific regions at specific phases of the gait cycle, is essential for applications to rehabilitation [9, 11, 19, 117]. Achieving this kind of spatial selectivity with spinal cord targets is challenging. However, our results suggest that targeting specific nerve roots may allow for spatially patterned stimulation to be applied transcutaneously. This makes the possibility of transcutaneous stimulation for augmenting rehabilitation more viable.

Our results have important implications for the design of electrode montages. We have shown that the standard midline arrangements that are commonly used in investigations of transcutaneous spinal cord stimulation are not optimal. These require a higher current to produce activation of nerve roots, which likely mediate the effects of stimulation. Given this understanding of the likely mechanism, we show that the use of off-midline cathodes may allow for stimulation to be delivered at lower intensities while minimising off-target effects, increasing the efficacy, tolerability and efficiency of stimulation.

These results need to be validated in experimental studies to verify the mechanism of stimulation. However, they suggest the likely need for novel electrode configurations to achieve optimised transcutaneous stimulation, as well as the

possibility of spatially patterned stimulation by taking advantage of the root-mediated mechanism of action. Further, this opens the possibility of targeting other structures within the intervertebral foramen, such as the dorsal root ganglion, with many potential applications.

These results are based on detailed computational models of the effects of electrical stimulation. Further refinement of these models may allow better characterisation of the mechanisms underlying the effect of electrical stimulation and may help to facilitate design of stimulation protocols for experimental studies.

However, the results of these studies are only as reliable as the underlying models. The models used here were highly detailed in terms of anatomy and physiology. However, explicit inclusion of more tissue types and addition of specific microstructural features which were neglected here, such as the angle of entry of the roots into the spinal cord [49, 122] may help to further improve the reliability of the results generated from these models.

7.4.3 Conclusion

Our results suggest that transcutaneous stimulation of neural elements is feasible, but that the mechanism of action of stimulation is likely due to activation of segmental nerve roots rather than direct activation of the spinal cord. Adapting our approach to transcutaneous stimulation based on this understanding and the differences in electrode positioning and stimulation parameters it implies is important to ensure that transcutaneous stimulation is applied effectively. Taking advantage of the spatial distribution of nerve roots may further allow for spatially patterned stimulation. All of these results support the viability of model-driven, spatially targeted transcutaneous activation of nerve roots, with potential applications to stimulation-facilitated rehabilitation for functional restoration in spinal cord injury.

Part V

Conclusion

8

Conclusion

Contents

8.1	Model-driven spinal cord stimulation	155
8.1.1	Transcutaneous nerve root stimulation	155
8.1.2	Targeting stimulation	156
8.1.3	Functional restoration	158
8.2	Computational modelling of neuromodulation	159
8.2.1	Physical basis of neuromodulation	159
8.2.2	Modelling of electrical stimulation	160
8.2.3	Patient-specific therapy	161
8.3	Outlook	162
8.3.1	Limitations	162
8.3.2	Further development	165
8.3.3	Conclusion	166

8.1 Model-driven spinal cord stimulation

8.1.1 Transcutaneous nerve root stimulation

We demonstrate that it is possible to produce a nervous system response using non-invasive stimulation methods. We show this using detailed computational models of the electric fields produced by stimulation and the effect of these electric fields on the nervous system. This provides rigorous computational and mechanistic evidence to support the application of non-invasive stimulation of spinal structures in ongoing

trials and pilot studies, such as for aiding functional restoration following spinal cord injury [123]. Our results show that electric fields generated via non-invasive stimulation can produce nervous system responses, providing further evidence of the feasibility of non-invasive spinal stimulation techniques.

However, our results show that the response to stimulation is mediated by activation of dorsal nerve roots. Direct activation of spinal cord structures requires significantly higher stimulation amplitudes with all electrode arrangement and stimulation parameters assessed. This suggests that any effects produced by non-invasive stimulation are likely to be produced by stimulation of spinal nerve roots, with indirect activation of spinal structures. This is in line with previous computational studies on the effects of epidural spinal cord stimulation, suggesting that non-invasive stimulation produces effects via the same mechanism as epidural stimulation [14, 15, 18]. This is an interesting result, as it raises the possibility that non-invasive stimulation may be able to achieve the same effects as epidural stimulation by acting via the same structures. While implanted stimulation systems retain significant advantages in many domains, the ability to achieve the effects of epidural stimulation without the associated risks and the need for surgical intervention is promising.

These results provide strong evidence for the possibility of achieving non-invasive stimulation of spinal structure via activation of dorsal nerve roots. This provides mechanistic support to the application of non-invasive stimulation in this context and raises the possibility of further extending the applications of non-invasive stimulation of dorsal nerve roots to other applications of spinal cord stimulation.

8.1.2 Targeting stimulation

We show that the ability to stimulate nerve roots transcutaneously offers the ability to achieve highly selective stimulation. Targeting specific neural subpopulations within the spinal cord using transcutaneous stimulation is challenging, and is unlikely to be feasible using standard stimulation methods. Due to the higher currents required to achieve activation from the skin surface relative to epidural electrodes,

larger electrodes are required in order to reduce the current density at the electrode surface and avoid damage to the skin. As a consequence, the electrodes used for transcutaneous stimulation are significantly larger than the structures they are targeting, making fine selectivity challenging. Further, there is a significant volume of intervening tissue between the skin surface and the spinal cord, and the spinal cord is mobile during movements. This produces significant restrictions on the level of selectivity that it is possible to achieve in targeting specific segments of the spinal cord. While this can be somewhat overcome using techniques such as interferential stimulation [101], significant limitations remain.

We show that targeting the dorsal nerve roots using transcutaneous stimulation is possible with high levels of selectivity in both rostrocaudal and mediolateral axes. This then produces activation of spinal structures in a manner analogous to the effects of epidural spinal cord stimulation [14, 15, 18]. The anatomical separation of segmental nerve roots makes these a far more suitable target than direct targeting of neural subpopulations within the spinal cord. By utilising this root-mediated mechanism for targeting specific segmental levels and lateralities, it is possible to achieve far greater selectivity using surface stimulation than previously thought. These results are in line with previous work that suggested differential effects of transcutaneous stimulation at different rostrocaudal levels [34–36], as we here demonstrate that this differential effect is likely produced by activation of different segmental levels via nerve root stimulation.

We further demonstrate the electrode montages and stimulation parameters required to achieve targeted stimulation at specific segmental levels. Notably, the montages required to achieve selectivity are a significant departure from the electrode arrangements that have been used as standard in the field of non-invasive spinal stimulation. Most existing and ongoing trials of non-invasive spinal cord stimulation rely on the use of midline electrodes [43, 123]. We demonstrate that achieving targeted stimulation required the use of off-midline electrode positions. The design of new montages based on these principles is likely to provide better selectivity of stimulation as well as lowering the amplitude required to produce a

response, reducing cutaneous discomfort and improving tolerability of stimulation. A stimulation approach based on these principles is therefore likely to offer significant practical advantages over existing methods.

8.1.3 Functional restoration

The ability to achieve spatially targeted activation via stimulation of segmental nerve roots is attractive from an application perspective. Recent developments in the use of epidural stimulation for functional restoration have been based on the development of spatiotemporally targeted stimulation methods [9, 11, 19]. By using the principles we have developed to design stimulation systems that target segmental nerve roots, analogous spatiotemporal selectivity may be achievable using transcutaneous stimulation alone.

The application of spatiotemporally patterned non-invasive spinal cord stimulation requires the use of model-derived electrode montages and stimulation parameters for accurate targeting. By designing systems to use this approach, we open the possibility of applying model-derived spatiotemporally patterned stimulation for stimulation-facilitated rehabilitation. This would allow the application of this emerging functional restoration technique without the need for surgery, reducing the associated risk and greatly increasing accessibility. This would potentially allow for model-derived non-invasive stimulation methods to be incorporated into rehabilitation to make stimulation-facilitated rehabilitation available on a much larger scale.

While this approach would potentially facilitate rehabilitation, the precise role of spatiotemporally patterned stimulation remains to be determined. Restoration of locomotor function has been based on the use of patterned stimulation to facilitate activity during walking [9]. Other work has shown that simpler stimulation during rehabilitation aids functional recovery, and some of those functional gains are maintained without stimulation [8]. Whether spatiotemporally patterned stimulation augments rehabilitation to a greater degree and the extent of dependence on ongoing stimulation need to be determined.

Implanted therapies have a clear benefit in the event of stimulator-dependent activity, where functional restoration is dependent on the presence of patterned stimulation. However, if stimulation continues to be shown to augment the effects of rehabilitation with repair and maintenance of functional gains even without stimulation, non-invasive approaches that can achieve this potentially bring great value to the area of functional restoration after spinal cord injury.

8.2 Computational modelling of neuromodulation

8.2.1 Physical basis of neuromodulation

In order to allow the detailed investigation of neuromodulation techniques, we have outlined a mathematical framework for describing the effects of electrical stimulation on the nervous system and have outlined how clinical questions can be framed as optimisation problems within this framework. This couples a system of partial differential equations describing the behaviour of electric fields in tissue with a system of ordinary differential equations describing the influence of external electric fields on neural membrane potentials. This approach allows for rigorous mathematical reasoning about difficult problems in neuromodulation. More interestingly, the system of equations we describe can be solved numerically with significant practical implications. We demonstrate that by solving this system describing the effects of neuromodulation, we can investigate the effects of stimulation on detailed models and can optimise stimulation parameters to achieve targeted effects in specific individuals. This provides significant mechanistic insight into the physical systems underlying the effects of electrical stimulation and provides the practical ability to model the effects of varying stimulation parameters in individuals to investigate the effects.

This approach to modelling the effects of neuromodulation differs to that used in other modelling approaches in its rigorous physical realism. The electric fields produced by stimulation are computed using a system of partial differential equations derived directly from the physics of electromagnetics without introducing simplifications or assumptions about the nature of stimulation. These equations are solved using the finite element method on highly detailed models of anatomy.

The effects of these fields on nervous tissue are then calculated using explicit biophysical models of axons, again avoiding any simplifications or assumptions about the nature of the response of neurons to electric fields. This provides significant advantages in the accuracy and reliability of results obtained by avoiding any inaccuracies introduced by simplifications in the mathematical system linking stimulation to observed response.

8.2.2 Modelling of electrical stimulation

We present a set of methods for carrying out detailed patient-specific modelling of the effects of electrical stimulation. By numerically solving the mathematical system describing the effects of neuromodulation in an individual, we can accurately determine the effects of stimulation with arbitrary electrode positions and stimulation parameters in that individual. The value of detailed modelling has been repeatedly demonstrated in the design of novel stimulation approaches [11, 14, 18, 19] and guiding patient-specific treatment using existing therapies [49, 50].

The approach we have developed allows detailed models of individual anatomy to be derived from standard clinical imaging. This allows the effects of stimulation on individuals to be evaluated, and provides significant additional accuracy over simplified, generic geometric models. By deriving detailed models directly from imaging, we remove the need for simplifications of anatomy or highly complex, time-consuming analysis pipelines to create patient-specific models. This represents an approach that maintains the accuracy required for individualised modelling while providing a method that is potentially scalable for application in clinical environments.

Having computed the electric field produced by stimulation, we demonstrate that this can automatically be coupled to biophysical axon models mapped into anatomically realistic geometries. This allows for highly accurate assessment of the effects of electrical stimulation on nervous tissue. The use of biophysical models overcomes many of the issues associated with using simplified models linking electric fields and stimulation effects. The automated coupling methods we have developed allow for these highly detailed models to be used without the need for highly

simplified geometries or laborious manual positioning of axons. This opens the possibility for highly detailed, individualised models of anatomy and physiology in a manner that is easy to apply on a large scale without compromising accuracy.

These methods provide an integrated set of techniques for modelling the effects of electrical stimulation on specific neural targets in detailed models of individualised anatomy and physiology. This allows for the effects of arbitrary electrode positions and stimulation parameters to be assessed in a given individual using only standard clinical imaging, without the need for highly laborious analysis pipelines. This approach has demonstrated significant value in our investigation of the effects of non-invasive spinal cord stimulation, and also has potential to provide valuable insights into a range of problems in clinical neuromodulation.

8.2.3 Patient-specific therapy

By modelling the response to stimulation in detailed individualised models, these techniques have potential applications to patient-specific therapies. These approaches can be used for detailed targeting of noninvasive stimulation in individuals, allowing for more rapid optimisation of therapies and selection of stimulation parameters. The same approaches can also be applied to the development of novel stimulation methods and for the investigation of the feasibility or mechanism of potential therapies, as described in the case of noninvasive stimulation for rehabilitation in spinal cord injury.

The models created using this approach are highly detailed and specific to individuals. They are generated automatically, allowing clinically useful applications without extensive time investment in preparing individualised models. This overcomes much of the existing issues with these modelling techniques.

The application of optimisation methods then allows for automatic determination of individualised stimulation parameters. This potentially provides a useful means of identifying a starting point for stimulation programming and electrode placement, potentially reducing programming time and ensuring more accurate electrode placement. Approaches such as this become increasingly valuable as the complexity

of stimulation systems increase, with increasing numbers of controllable parameters and options for stimulation.

This offers a rigorous computational approach to neuromodulation based on highly realistic individualised models. This has significant potential applications to electrical neuromodulation and potentially offers a useful toolset for studying neuromodulation and developing targeted therapies.

8.3 Outlook

8.3.1 Limitations

These developments offer a valuable set of techniques for modelling the effects of electrical neuromodulation. However, the conclusions drawn are only as reliable as the underlying model. Applications to complex targets with poorly understood biophysical properties can therefore be unreliable. Model accuracy is therefore dependent on the accuracy of the characterisation of the properties of each of the tissue types modelled. This includes physical properties such as the impedances of tissue types at various frequencies and the properties of different axon fibre types, such as distribution of ion channel types. These parameters are not always well characterised in healthy tissue; it is also not entirely clear what parameters change as a consequence of disease states. For example, chronic loss of input to spinal circuits following spinal cord injury may result in changes in the expression of ion channels and a resultant change in the response to electrical stimulation.

Overcoming these issues would require basic science research to characterise the specific biophysical changes that occur in pathologies of the nervous system. In the absence of this, we are forced to proceed with modelling using the best information available, acknowledging that this may deviate from the true state of the nervous system, particularly in pathological states.

Further, the quality of the model is limited by the quality of the imaging and resulting segmentation to isolate structures of interest. Imaging artefacts need to be identified and dealt with to avoid them influencing the results of simulations. However, even with high-quality imaging, whether the images represent the anatomy

of interest needs to be considered. For example, the position of the spinal cord - and therefore the thickness of the CSF layer posterior to the cord - varies significantly with positioning. Even a very high-quality image and segmentation runs the risk of producing inaccurate results if we attempt to apply these results to other positions without considering whether the position during scan acquisition has influenced the anatomy of interest. It is important to take care in the application of results, even with apparently high-quality anatomical models.

Related to this, application of this approach requires visualisation of target structures. If it is visible on imaging, it can be segmented and included in an anatomical model. However, this limits application to structures not easily visible on imaging or structures with specific microanatomy of interest. Although it is possible to work around this issue using heuristics such as the likely location of structures based on known landmarks or normative models of microstructure embedded within the detailed anatomy, this limits the accuracy of the results as the model is no longer truly patient-specific. Better characterisation of anatomy with dedicated imaging sequences and image fusion where necessary may help to overcome some of these issues.

The segmentation of imaging to isolate structures of interest is a key step in the development of patient-specific models. This process can be partially automated using simple tools, while fully automated segmentation of clinical imaging continues to develop. However, it is necessary to carefully examine and correct the segmentations produced. This can be time consuming. This issue is likely to reduce as tools for automated segmentation continue to improve, but currently quality control of individualised models requires careful attention to ensure that erroneous results are not produced due to errors in the physical model.

Even with accurate biophysical data, imaging and segmentation, the application of modelling results can also be influenced by the structures modelled. Off-target effects are a major limiting factor in electrical neuromodulation. If the appropriate structures that mediate these effects - such as cutaneous afferent fibres - are not included in models, effects can be suggested by models which are ultimately not

applicable in an experimental setting as stimulation is limited by effects mediated by structures not included in the original models. It is therefore important to consider what structures are included and if relevant off-target structures are excluded to acknowledge that these may limit actual application.

The accuracy of the representation of the biophysically important features of the target structures is an important consideration. While automatically generated trajectories in anatomical volumes is a valuable approach to rapidly model complex targets while preserving anatomical accuracy, loss of microstructural features is an important consideration. Model structures may need to be adapted based on existing knowledge of the target structures to ensure results are reliable and to avoid inaccuracies being introduced by automatically generated trajectories.

Similarly, the branching structures of axon models can affect the response to stimulation. Isolated axons with unconnected terminals can behave unpredictably when exposed to stimulation and can produce artefactual activation. Where the anatomical connections of a target structure are known and where this branching occurs within the region of interest, it is necessary to adapt the model to account for this to avoid inaccurate results.

In this way, our explicit knowledge of the target region, drawn from basic science, helps us to build reliable computational models. By leveraging this knowledge to produce accurate models, we maximise our ability to generate useful insights and to use these approaches to inform experimental studies to drive forward the field of electrical neuromodulation.

The link between electrophysiological activity and physiological also needs to be considered in the use of modelling approaches. While it may be possible to generate action potentials in structures of interest, the link between this activity and clinical benefit is not always clear. Further clarity is needed on the exact link between neural activity and the desired physiological response, in health and in pathology, is needed to be able to reliably predict treatment outcomes with electrical neuromodulation. The assumption that generating action potentials are linked to a desirable physiological effect is reasonable, but it is less clear how much activity is

required to produce a clinically meaningful benefit, and its factors such as volume of action potentials or patterns of activation are important.

These results provide computational evidence to support the application of noninvasive stimulation to augment rehabilitation after spinal cord injury. However, it is necessary to validate these results in healthy volunteers and pilot clinical studies. These results provide a valuable advance in our understanding of noninvasive stimulation for rehabilitation and can guide the development of further studies, but are not sufficient to recommend application of noninvasive neuromodulation in a rehab context in isolation. Detailed validation studies are required in order to be able to reliably apply any of the outcomes of these computational studies.

8.3.2 Further development

These results provide a valuable basis for further developments in targeted neuromodulation.

We are working on initial trials of model-driven targeted noninvasive spinal cord stimulation in healthy volunteers on the basis of these results. We have started to recruit volunteers to validate the ability of noninvasive stimulation to produce targeted activation, measured by paraesthesia generation in target dermatomes and motor activation to indicate activation of specific segmental levels. In parallel, we are also testing the potential for varying stimulation waveforms to differentially activate axon fibre types. This aims to widen the window of selective activation of motor or sensory fibres, limiting off-target effects. It also aims particularly to limit the activation of cutaneous sensory afferents, which produce discomfort which can be a significant limiting factor in healthy volunteers or patients with intact sensation at the target spinal level. Preliminary results of this work have been promising, highlighting the potential translational utility of detailed computational approaches to helping to inform experimental work in electrical neuromodulation.

These techniques can then be refined and developed towards pilot clinical studies, with a view to translating these techniques into a viable intervention to augment rehabilitation in spinal cord injured patients. These tools can also be applied to other

areas of neuromodulation, such as spinal cord and dorsal root ganglion stimulation for pain, or to peripheral nerve stimulation for a wide range of applications.

The techniques demonstrated can be further developed to offer useful tools for investigating clinical neuromodulation. Appendix B outlines the integration of the techniques developed here into a standalone framework for individualised modelling of neuromodulation, with potential applications to targeted therapies and the development of novel interventions. These tools can continue to be improved to provide greater insight into electrical neuromodulation techniques. Further, appendix C shows the development of a system for generating complex stimulation waveforms for clinical testing. This allows simulated stimulation parameters to be straightforwardly applied and tested. In combination, these systems allow for arbitrary electrode arrangements and stimulation parameters to be simulated and investigated mechanistically and to be applied experimentally to validate results.

Further aspects of the techniques shown can be developed to widen the scope of application. The application of optimisation techniques can be improved to offer better individualisation of therapies. Integration with tractography based on diffusion tensor imaging can allow biophysical axon models to be mapped into anatomy based on specific fibre tracts, allowing targeting of specific fibre bundles. These and other developments potentially allow the computational framework of electrical neuromodulation to be expanded to impact a wide range of existing and novel therapies.

8.3.3 Conclusion

We demonstrate that noninvasive spinal cord stimulation can achieve targeted activation of neural structures at specific spinal levels and lateralities. We show that this is achieved by targeting nerve roots in line with the suggested mechanism of action of epidural stimulation for rehabilitation [9, 11, 14, 18, 19]. We characterise the electrode arrangements required to achieve this clinically. This provides computational and mechanistic support to the application of transcutaneous stimulation for rehabilitation in pilot studies and for its application in other domains.

We also show that a detailed physical model of neuromodulation allows clinically meaningful questions to be addressed. We outline a series of new techniques and tools to carry out individualised simulations of electrical neuromodulation. This represents a significant advance in the computational science of neuromodulation. We present an integrated system for carrying out individualised simulation. This opens the possibility for individualised, model-driven neuromodulation therapies and the principled development of novel stimulation systems. This is something that we intend to continue to develop going forwards, including incorporation of our knowledge of the microanatomy of the nervous system and axonal branching patterns, in order to develop a valuable resource for work in translational neuromodulation.

Outputs

Publications

Published

- C. Keogh, F. Saavedra, B. Andrews and J. J. FitzGerald, "Computation of Activating Fields for Approximation of the Orientation-Specific Neural Response to Electrical Stimulation," 2022 44th Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC), Glasgow, Scotland, United Kingdom, 2022, pp. 5152-5155, doi: 10.1109/EMBC48229.2022.9871706.
- Based on chapter 3.

In preparation

- C. Keogh and J.J. FitzGerald, "Neuromodulation for functional restoration".
- Review article based on chapter 1.
- C. Keogh and J.J. FitzGerald, "Patient-specific computational modelling for electrical neuromodulation". - Methods paper based on chapters 2, 4, 5 and 6.
- C. Keogh and J.J. FitzGerald, "Spatiotemporally patterned activation using transcutaneous spinal cord stimulation". - Based on chapter 7.

Contributions to other work

The techniques developed here have also been used in designing the stimulation protocols used in the following works:

- C. Keogh, F. Saavedra, S. Dubo, P. Aqueveque, P. Ortega, B. Gomez, Non-invasive phrenic nerve stimulation to avoid ventilator-induced diaphragm dysfunction in critical care. *Artif. Organs.* 2022; 46: 1988–1997. <https://doi.org/10.1111/aor.142>

- C. Keogh, F. Saavedra, S. Dubo, P. Aqueveque, P. Ortega, B. Gomez, et al. Closed-loop parameter optimization for patient-specific phrenic nerve stimulation. *Artif. Organs*. 2024; 48: 274–284. <https://doi.org/10.1111/aor.14593>
- "Pre-operative electrical stimulation in carpal tunnel syndrome (PROSPECT)", an ongoing randomised controlled trial of electrical neuromodulation for improving outcomes following carpal tunnel decompression surgery; ISRCTN registration number: ISRCTN94245451. <https://doi.org/10.1186/ISRCTN94245451>

Presentations

- "Spatiotemporally patterned activation using transcutaneous spinal cord stimulation", International Neuromodulation Society, Vancouver, 2024
- "Patient-specific modelling for precision neuromodulation", International Neuromodulation Society, Barcelona, 2022
- "Computation of activating fields for approximation of the orientation-specific neural response to electrical stimulation", IEEE Engineering in Medicine and Biology Society, Glasgow, 2022

Appendices

A

Derivation of variational forms for numerical solution of physical problems in neuromodulation

Contents

A.1	Introduction	171
A.2	Constant current	173
A.3	Time-varying current	176
A.4	Conclusion	179

A.1 Introduction

The electric field produced in the human body by neuromodulation therapies can be described using a system of partial differential equations. As described in chapter 2, we can use the physical laws of electromagnetics to describe the effects of these therapies.

These mathematical systems are, however, not analytically solvable. We therefore make use of a set of techniques to compute numerical solutions for specific cases. In this way, we can use this physical description of the behaviour of electric fields to predict the effect that it will have in any specific individual by computing the

numerical solution for that individual, even though we cannot derive an analytical expression that describes this effect exactly in all participants.

The finite element method is one of most successful methods for solving systems of partial differential equations [55]. We describe our physical problem independent of the domain to which it is applied. We can then compute a solution for any arbitrary domain by splitting even highly complex domains up into a large number of more elementary structures with a simpler shape - the "finite elements" - and solving our system iteratively on these smaller elements. By subsequently recombining all of the elements that constitute the domain of interest, we can get a solution on the whole domain.

More concretely, the laws of electromagnetics are well-described and can be written succinctly as systems of partial differential equations. However, it is difficult to numerically calculate the effect of stimulation on an individual patient's anatomy as the area being considered is a complex structure with many tissue types and intricate anatomy. However, by breaking this up into very small parts and considering these one at a time, the system of equations can be solved on these simpler parts then recombined.

In order to use this numerical approach, we must re-express our problem in a specific form, referred to as the variational form, the weak formulation or the bilinear form of the system of partial differential equations. This expression takes the general form

$$a(u, v) = L(v)$$

Where u is the unknown (test) function and v is a (trial) function used as part of the numerical method. By re-expressing our problem in this form, we can apply the finite element method to compute a numerical solution.

This appendix outlines the derivations of the systems of partial differential equations describing problems in neuromodulation for both static and time-varying currents and shows the derivation of the variational forms used to solve these systems numerically. These derivations form the basis of the solver outlined in chapter 5 and are central to computing the electric fields produced by neuromodulation.

A.2 Constant current

To describe the electric fields produced by neuromodulation, we begin with the law of conservation of current:

$$\nabla \cdot \vec{J} = -\frac{\partial \rho}{\partial t} \quad (\text{A.1})$$

where \vec{J} is a vector field representing the current density and ρ is the charge density. This states that the current density can only change as a function of the change in the amount of charge in the volume.

In the case of steady current problems, i.e. where direct currents are applied and we are interested only in the "steady state", ignoring the transient periods around the onset and offset of stimulation, then we can simplify this by stating that there is no change in the charge density at steady state:

$$\frac{\partial \rho}{\partial t} = 0 \quad (\text{A.2})$$

By substituting equation A.2 into equation A.1, we get:

$$\nabla \cdot \vec{J} = 0 \quad (\text{A.3})$$

i.e. in the steady-state approximation, the current density field has no sources or sinks.

We can then relate the current density to the electric field using Ohm's law:

$$\vec{J} = \sigma \vec{E} \quad (\text{A.4})$$

where σ is the electrical conductivity and \vec{E} is the electric field.

In combination with the definition of the electric field:

$$\vec{E} = -\nabla V \quad (\text{A.5})$$

where V is the scalar electric potential, equation A.4 gives us:

$$\vec{J} = -\sigma \nabla V \quad (\text{A.6})$$

which describes the relationship between the electric potential, conductivity and the current density.

By substituting equation A.6 into equation A.3, we get an equation describing the behaviour of steady-state currents in materials described by Ohm's law:

$$-\nabla \cdot (\sigma \nabla V) = 0 \quad (\text{A.7})$$

We can then solve for the electric field produced by stimulation by solving this equation with suitable boundary conditions. In order to solve this problem using the finite element method, we must first convert it to the standard variational form. In order to do this, we will use the standard nomenclature, i.e. our function to be approximated, the electric potential V , will be referred to as the trial function u , while v will designate the test functions used as part of the finite element method. This step, performed purely for the purposes of consistency with the standard notation used in the literature on the finite element method, changes equation A.7 to:

$$-\nabla \cdot (\sigma \nabla u) = 0$$

We can further generalise this by adding a source term, f , to the right hand side, allowing for the possibility of current sources:

$$-\nabla \cdot (\sigma \nabla u) = f \quad (\text{A.8})$$

We then multiply equation A.8 by the test function v and integrate over the domain Ω :

$$-\int_{\Omega} \nabla \cdot (\sigma \nabla u) v \, dx = \int_{\Omega} f v \, dx \quad (\text{A.9})$$

We then apply Gauss-Green's lemma to reduce the order of the derivative of u , thereby weakening the requirement of the polynomial space used for the trial functions to be twice-differentiable:

$$\int_{\Omega} -\nabla^2 u v \, dx = \int_{\Omega} \nabla u \cdot \nabla v \, dx - \int_{\partial\Omega} \frac{\partial u}{\partial n} v \, dS \quad (\text{A.10})$$

By applying this lemma to equation A.9, we get:

$$-\int_{\Omega} \sigma \nabla u \cdot \nabla v \, dx = \int_{\Omega} f v \, dx - \int_{\partial\Omega} \frac{\partial u}{\partial n} v \, dS$$

This is simplified by using the fact that we, by definition, require that the test function v is compactly supported on the domain Ω , i.e. v vanishes on the boundary, eliminating the surface integral on the boundary. This allows us to arrive at the weak formulation of our problem:

$$-\int_{\Omega} \sigma \nabla u \cdot \nabla v \, dx = \int_{\Omega} f v \, dx \quad (\text{A.11})$$

Our variational problem is then to find $u \in V$ such that:

$$a(u, v) = L(v) \quad \forall v \in \hat{V} \quad (\text{A.12})$$

where the function space V is our trial space, i.e. the space from which our solution can be drawn, and \hat{V} is the test space, from which the test functions are drawn. By imposing restrictions on the function spaces V and \hat{V} , we control the kinds of solutions that are possible. In our case, we define:

$$\begin{aligned} V &= \{v \in H^1(\Omega) : v = u_D \quad \text{on } \partial\Omega\}, \\ \hat{V} &= \{v \in H^1(\Omega) : v = 0 \quad \text{on } \partial\Omega\} \end{aligned}$$

where $H^1(\Omega)$ is a Sobolev space, imposing the requirement that u and v are continuous over the domain Ω . However, this requirement allows u to have discontinuous derivatives, imposing a weaker continuity requirement than the underlying partial differential equation, shown in equation A.7, which requires solutions that are twice-differentiable.

In order to express our problem in the standard format shown in equation A.12, we define:

$$\begin{aligned} a(u, v) &= -\int_{\Omega} \sigma \nabla u \cdot \nabla v \, dx, \\ L(v) &= \int_{\Omega} f v \, dx \end{aligned}$$

By collecting all terms containing our unknown function u in $a(u, v)$, the bilinear form, and all terms containing only the (known) test function v in the linear form

$L(v)$, we can then apply differing test functions v from the function space \hat{V} until we have assembled a linear system with as many equations - each with a different, known v - as there are unknown values of u . Boundary conditions can then be applied to this linear system and it can then be solved using standard methods.

This approach allows us to calculate the potential produced by stimulation using a standard method even on very complex geometries.

A.3 Time-varying current

The same approach can be applied in order to study the case of time-varying electric fields produced by a stimulating current that changes over time.

Here we again start with the law of conservation of current, shown in equation A.1, here rearranged:

$$\nabla \cdot \vec{J} + \frac{\partial \rho}{\partial t} = 0 \quad (\text{A.13})$$

In the case where there is a time-varying current, we can no longer assume that $\frac{\partial \rho}{\partial t} = 0$ in order to simplify this expression as in equation A.2.

Instead, we take Gauss's law, which relates the electric displacement field \vec{D} to the spatial distribution of charge:

$$\nabla \cdot \vec{D} = \rho$$

By substitution into equation A.13, we get:

$$\nabla \cdot \vec{J} + \frac{\partial \nabla \cdot \vec{D}}{\partial t} = 0 \quad (\text{A.14})$$

We can then relate the electric displacement field to the electric field using the constitutive law for linear materials:

$$\vec{D} = \varepsilon_0 \varepsilon_r \vec{E} \quad (\text{A.15})$$

where ε_0 is the permittivity of free space and ε_r is the relative permittivity of the material in the domain.

By substituting equation A.15 and Ohm's Law, given in equation A.4, into equation A.14, we get:

$$\nabla \cdot \sigma \vec{E} + \frac{\partial \nabla \cdot (\varepsilon_0 \varepsilon_r \vec{E})}{\partial t} = 0$$

By again using the definition of the electric field given in equation A.5, we get an equation describing the behaviour of time-varying currents in materials described by Ohm's law:

$$-\nabla \cdot \sigma \nabla V - \frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} = 0$$

As in equation A.8, we can generalise this by allowing a non-zero source term f :

$$-\nabla \cdot \sigma \nabla V - \frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} = f \quad (\text{A.16})$$

This then describes the electric field produced by a time-varying current. This partial differential equation contains both spatial and temporal derivatives. While the spatial problem can be discretised and solved numerically using the finite element method, we also need to discretise the temporal aspect of this problem in order to solve it.

Here we will use a simple finite difference discretisation of time. We define a notation where superscript n indicates that value at time t_n , where n is a discrete step in time, e.g. u^n indicates the value of the function u at time step n . The temporal derivative can then be approximated by a finite difference, i.e. as the difference in the value of the expression between the current and the next timepoint (i.e. a backward difference):

$$\left(\frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} \right)^{n+1} \approx \frac{\nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^{n+1} - \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^n}{\Delta t} \quad (\text{A.17})$$

We can rearrange equation A.16 to isolate the term with the temporal derivative:

$$\frac{\partial \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V}{\partial t} = f - \nabla \cdot \sigma \nabla V$$

By substituting in the finite difference approximation from equation A.17, we get:

$$\frac{\nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^{n+1} - \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^n}{\Delta t} = f - \nabla \cdot \sigma \nabla V^{n+1}$$

Multiplying across by Δt , we get:

$$\nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^{n+1} - \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^n = f \Delta t - \nabla \cdot \sigma \nabla V^{n+1} \Delta t$$

We can then rearrange so that all terms containing the unknown V^{n+1} are collected on the left-hand side, while known and previously computed terms are collected on the right-hand side:

$$\nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^{n+1} + \nabla \cdot \sigma \nabla V^{n+1} \Delta t = f \Delta t + \nabla \cdot \varepsilon_0 \varepsilon_r \nabla V^n \quad (\text{A.18})$$

This gives us an expression which is discretised in time and can be iteratively solved given some initial conditions, i.e. V^0 . In order to solve for the electric field produced by a time-varying current using this approach, we again must derive a variational form for our expression. Using the same convention as in equation A.8, we will designate our quantity to be computed, the electric potential V , as the trial function u and will allow v to designate the test functions. By replacing the terms in equation A.18, we get:

$$\nabla \cdot \varepsilon_0 \varepsilon_r \nabla u^{n+1} + \nabla \cdot \sigma \nabla u^{n+1} \Delta t = f \Delta t + \nabla \cdot \varepsilon_0 \varepsilon_r \nabla u^n$$

We then multiply by the test function v and integrate over the domain Ω :

$$\int_{\Omega} \nabla \cdot (\varepsilon_0 \varepsilon_r \nabla u^{n+1}) v \, dx + \int_{\Omega} \nabla \cdot (\sigma \nabla u^{n+1} \Delta t) v \, dx = \int_{\Omega} f \Delta t v \, dx + \int_{\Omega} \nabla \cdot (\varepsilon_0 \varepsilon_r \nabla u^n) v \, dx$$

By applying Gauss-Green's lemma, given in equation A.10, as in the case of steady currents, we get:

$$\int_{\Omega} \varepsilon_0 \varepsilon_r \nabla u^{n+1} \cdot \nabla v \, dx + \int_{\Omega} \sigma \nabla u^{n+1} \cdot \nabla v \Delta t \, dx = \int_{\Omega} f v \Delta t \, dx + \int_{\Omega} \varepsilon_0 \varepsilon_r \nabla u^n \cdot \nabla v \, dx$$

We can then express this in the standard form:

$$a(u^{n+1}, v) = L^{n+1}(v)$$

where

$$\begin{aligned} a(u^{n+1}, v) &= \int_{\Omega} (\varepsilon_0 \varepsilon_r \nabla u^{n+1} \cdot \nabla v + \sigma \nabla u^{n+1} \cdot \nabla v \Delta t) \, dx, \\ L^{n+1}(v) &= \int_{\Omega} (f v \Delta t + \varepsilon_0 \varepsilon_r \nabla u^n \cdot \nabla v) \, dx \end{aligned}$$

with initial conditions

$$\begin{aligned} a_0(u, v) &= \int_{\Omega} \varepsilon_0 \varepsilon_r \nabla u \cdot \nabla v \, dx, \\ L^0(v) &= \int_{\Omega} \varepsilon_0 \varepsilon_r \nabla u^0 \cdot \nabla v \, dx \end{aligned}$$

We can then solve for the electric field produced by a time-varying current by finding $u^0 \in V$ such that $a_0(u^0, v) = L^0(v)$ is true for all $v \in \hat{V}$. We can then find $u^{n+1} \in V$ such that $a(u^{n+1}, v) = L^{n+1}(v)$ for all $v \in \hat{V}$ for $n = 0, 1, 2, \dots$

In this way, we can discretise the temporal component of the problem using a finite difference approach, which produces a sequence of static problems that can be solved sequentially using the finite element method. This then allows us to calculate the electric field produced at each time point even with the application of time-varying currents.

A.4 Conclusion

The physical laws describing the electric fields produced by the application of electric stimulation can be expressed in the form of partial differential equations. In doing so, we can see that these equations have a standard form that is independent of the domain they are applied to. If we can solve these equations, we can therefore compute the electric field produced in a domain of any complexity for any set of boundary conditions - for example, a highly detailed model of individual anatomy for a given set of stimulation parameters.

In order to achieve this, these equations must be rewritten in variational, or bilinear, form. The primary partial differential equations for constant current and time-varying currents are derived here, starting from physical laws. We have then shown the derivation of variational forms of these systems of partial differential equations that are suitable for the numerical solution of problems in neuromodulation.

These systems can then be solved using the finite element method through assembly and subsequent solution of a linear system. In the case of time-varying currents, the problem can be discretised into time steps using a finite difference,

or implicit Euler, discretisation. The problem can then be iteratively solved using the same approach at each time step.

This approach allows us to solve, in a general way, this mathematical system describing the electric fields produced by neuromodulation. This has significant practical applications. By allowing the electric field produced by stimulation to be computed on any domain and for any set of stimulation parameters, this approach can be applied to produce patient-specific simulations of neuromodulation therapies for any set of electrode geometries or stimulation parameters.

The systems of equations outlined here are extremely general, and can be widely applied using a standard method to solve complex, patient-specific problems in neuromodulation without any alteration to the underlying mathematical system. This provides a valuable tool to the study of neuromodulation and its clinical application.

Chapter 4 describes the derivation of domains that incorporate detailed patient specific information; these equations can then be solved on such domains to yield patient-specific simulations. Chapter 5 describes a method for numerically solving the system of equations described here. This allows the practical application of the methods discussed here and their use to investigate real-world problems in neuromodulation.

B

NRRDosurgery: a generalisable tool for patient-specific neuromodulation

Contents

B.1	Introduction	181
B.2	Methods	184
B.2.1	Model creation	184
B.2.2	Interface	187
B.2.3	Implementation	188
B.2.4	Example	188
B.3	Results	188
B.3.1	Multiple interfaces	188
B.3.2	Abstracted system	190
B.3.3	Personalised models	190
B.4	Discussion	192
B.4.1	A tool for targeted neuromodulation	192
B.4.2	Limitations	194
B.4.3	Conclusion	194

B.1 Introduction

Electrical neuromodulation has become widely applied to a range of clinical conditions, such as pain [96] and movement disorders [124, 125], with emerging applications in a number of areas such as rehabilitation [7, 9, 10]. These therapies are

based on the application of electrical stimulation to produce controlled activation of the excitable tissue of the nervous system. Computational modelling of the electric fields produced by stimulation has demonstrated value in the development of novel stimulation strategies [9, 11, 18, 19] and in the prediction of the optimal stimulation parameters for individuals [49, 50].

Existing methods for modelling the effects of electrical stimulation are limited by the simplistic geometries used and the need for production of new models for all new applications. The use of highly detailed models of anatomy has been demonstrated to improve the accuracy of simulations, allowing more reliable prediction of individually optimised stimulation parameters [49, 50]. However, these more detailed methods are limited by the complex, time consuming workflows required to apply them, making them unsuitable for more widespread use or for clinical application.

Even following the simulation of the electric field produced by stimulation in complex anatomy, we cannot draw conclusions regarding the effects of stimulation. We need to link our measures of the electric field to some measure of activation. Many simplified metrics exist, from the $L2$ norm of the electric field [51, 77, 88] to classical metrics such as the activating function [89, 90]. In order to most accurately predict the response to stimulation, we need to link our electric field results to biophysical axon models in order to predict the physiological response of the nervous system to external stimulation [74].

By carrying out this simulation of the effects of stimulation in an individual for a given set of stimulation parameters, we can then apply optimisation techniques in order to identify the set of stimulation parameters that maximise activation of our target structure for that individual. This converts our clinical question into an individualised optimisation problem based on our computational model. However, the application of this kind of approach to individualising therapies is limited by the need for detailed individualised models and automated model generation and solving in order to allow iterations over stimulation parameters to be applied, which is not possible with current tools.

The clinical application of this computational approach to neuromodulation, in optimisation of existing therapies for individuals and in the development of new treatments, is limited primarily by technical barriers regarding the existing of automated tools with the correct capabilities. A tool that is able to use routine clinical imaging and provides a simple interface for defining stimulation parameters and can then automatically simulate the physiological effects of stimulation and perform parameter optimisation therefore has the potential to have great impact in the field of neuromodulation.

We have demonstrated in chapter 4 that it is possible to derive detailed computational models directly from clinical imaging, overcoming these issues and providing a method for automatically deriving detailed models with explicit representation of complex anatomy. We showed in chapter 5 that the electric field produced by stimulation can be solved on these models, simulating the effects of intervention, and in chapter 6 we showed that these detailed anatomic models can be coupled with biophysical axon models to predict the physiological effects of stimulation and optimise the stimulation parameters used for that individual.

We then aimed to develop a generalisable tool for detailed individualised biophysical modelling of the effects of electrical neuromodulation using these methods. This would provide a platform for the development of novel neuromodulation techniques and the individual optimisation of therapies.

We demonstrate a tool that allows for the creation of individualised models directly from clinical imaging, including the use of biophysical models and optimisation for truly optimised stimulation in individuals. This tool has an easy to use interface, with a graphical interface as well as text-based and scripting interfaces for complex analyses. This tool is optimised for potential use in clinical workflows, requiring only clinical imaging and a simple interface for defining parameters. This represents a potentially valuable tool for the development of new neuromodulation therapies and translational neuromodulation research more generally.

B.2 Methods

B.2.1 Model creation

In order to simulate the effects of stimulation on individual anatomy, a set of stimulation parameters must be specified. An individualised computational model with simulated electrode placement according to these parameters can then be generated. The electric field produced by stimulation is then solved using this model. The results of this simulation are then coupled to biophysical axon models mapped into the image volume. The physiological effects of stimulation can then be measured, and numerical optimisation methods applied to identify the optimal set of stimulation parameters to maximise activation of the target.

In order to achieve this, results must be passed from one part of the overall simulation system to another. In its simplest form, all user-specified parameters, including stimulation parameters and the image volume to be used, can be specified as a single data structure. This data structure can then be passed to a solver engine which uses these parameters to sequentially complete each part of the simulation process and returns the results. This is shown in figure B.1.



Figure B.1: Flowchart of overall system design. Models are specified using a graphical interface, a plain text file or by directly manipulating model data structures. The resulting data structure is passed to the solver engine, which carries out all required analyses and returns the desired results. This approach abstracts the technical details of analysis and provides a simple interface, greatly reducing the technical barriers to applying these methods to clinical questions.

This forces all configurable parameters to be determined in advance. This allows for detailed checking that all parameters are valid before a model can be run, avoiding issues with errors in the late stages of long-running simulations. Further, it allows the complex details of the model's solver to be abstracted away. All results are automatically computed and passed to the next stage, so there is no requirement to have a detailed understanding of the workings of the solver engine beyond the input parameters required to run a simulation.

The solver engine itself is made up of a number of stages, as shown in figure B.2. Each of these stages implements some transformation of the data and passes the result automatically to the next stage using the configurable parameters specified at the beginning. This allows highly complex simulation workflows to be undertaken without the need for detailed knowledge of each of the individual steps.



Figure B.2: Flowchart of solver engine design. Images are prepared and electrode placement is simulated to provide an image with simulated electrodes in situ. A tetrahedral mesh is then generated using this image volume. Boundary conditions are applied and physical characteristics applied to each tissue type. The electric field is then solved using finite element analysis (FEA). The results of this analysis are coupled to detailed biophysical axon models in order to evaluate the physiological response to stimulation with the simulated electrodes.

The image processing step involves preparing the clinical imaging and mapping the desired electrode geometry into the image volume. This allows for simulation of neuromodulation procedures using routine pre-procedure imaging, and allows

for iteration over electrode positions and measurement of the response in order to identify the optimal position for that individual. The images are then meshed to form a 3D tetrahedral mesh suitable for finite element analysis. This process is described in detail in chapter 4.

The electric field produced by stimulation using the simulated electrodes in detailed patient-specific anatomy can then be solved. This is carried out using finite element analysis. As shown in figure B.3, this involves specification of the mathematical problem in its variational form, assembly of a linear system on the mesh specified in the previous step with application of boundary conditions according to the model specification and solution of the linear system using numerical methods. The derivation of the variational form of the physical problem posed by electrical neuromodulation is described in appendix A, and the process for solving for the electric field on detailed patient-specific models is described in chapter 5.




Figure B.3: Flowchart of finite element analysis workflow. The physical problem to be solved is stated mathematically and expressed in variational form. Using the finite element method, a linear system is then assembled using the values at each node of the tetrahedral mesh. This linear system is then solved using numerical methods.

The electric field produced by stimulation can then be used to derive measures related to activation of target structures. This can be carried out using simplified metrics, which link features of the electric field to neural activation. This is computationally simple and suitable for some scenarios where it is necessary to

produce results rapidly or to test high numbers of possible parameters. This process is described in detail in chapter 3.

More accurate measures of activation can be derived by mapping detailed biophysical axon models into the image volume. The results of the electric field simulation can then be coupled to these models and the effect of stimulation on physiologically realistic axons assessed. Measures of activation can then be derived and used to evaluate the effects of stimulation with specific parameters. Methods from numerical optimisation can then be applied in order to determine the optimal input parameters to produce maximal activation within a specified target. This technique is described in detail in chapter 6.

Following application of each of these techniques, the results can then be saved or visualised as required.

B.2.2 Interface

Model parameters can be specified using a number of interfaces, suitable for different applications. The essential criteria is that they describe a data structure which contains valid entries for all required parameters such that it can be taken in by the model parser and sent to the solver engine. Models can be specified using a graphical interface, a file format or via a scripting language.

A simple graphical user interface allows all required parameters to be specified. By working through each section, all parameters are filled out with explanatory notes for each. In this way, it is straightforward for users with limited technical experience to specify and run complex simulations of neuromodulation interventions. The values entered into this field are then used to create the model specification data structure.

Models can be specified as a plain text file using key-value pairs. This allows for models to be quickly edited within a text editor. This allows for rapid model updates, as well as saving and transferring models, or creating large queues of models to be run. Similarly, it allows automatic generation of model specification files to run analyses such as parametric sweeps. These files can then be loaded by the solver engine and parsed to create the required data structure.

For more advanced users, the data structures can be manipulated directly. This allows for rapid and automated generation of models, including for optimisation and parametric sweeps, without the need for intervening temporary files. This also allows the solver methods to be incorporated into other analysis pipelines, allowing highly complex analyses to be performed.

In all cases, the required data structure must specify all required parameters to create a patient-specific mesh, map in the target electrode geometries to the required locations, specify the physical characteristics of each tissue type and the stimulation applied to solve for the electric field, specify the target segments and the results required. The solver can then run the model with the specified parameters and return the required results.

B.2.3 Implementation

All software was implemented in Python [111], using the numpy library [126] for handling of multidimensional arrays, matplotlib [127] for generation of plots, FENiCS [112] for specification of systems of partial differential equations, NEURON [92] for solving biophysical models and scikit-learn [128] for optimisation algorithms.

B.2.4 Example

As a simple usage example, we demonstrate the use of this approach to simulate the effects of stimulation with 100mA through a 3.2cm diameter cathode in the midline over the T11 vertebra with a corresponding 3.2cm diameter anode at the same level in the midline anteriorly. We position target axons within the dorsal columns of the spinal cord in order to evaluate whether non-invasive stimulation at this position with this intensity can produce activation of spinal cord structures in this individual.

B.3 Results

B.3.1 Multiple interfaces

We demonstrate a system for patient-specific modelling of neuromodulation using multiple interfaces. A common solver engine performs all image processing, meshing,

finite element analysis, model coupling, biophysics and optimisation steps, while multiple routes can be used to specify model parameters depending on the intended use.

Figure B.4 shows an example graphical user interface for model specification. All required parameters can be filled out by working through each tab and the model can then be run directly from the graphical interface. This provides a simple means of carrying out complex neuromodulation analyses, significantly reducing technical barriers to applying these techniques. However, this approach only allows single models to be run and is limited in its flexibility when compared to other methods of model specification.

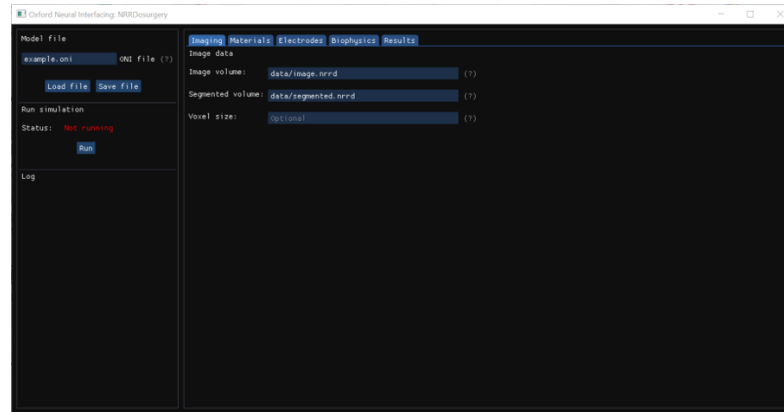


Figure B.4: Graphical user interface for model specification. The data structure required for running a model can be populated using a simple graphical interface. By working through each tab, all of the required parameters are assigned. Explanatory notes and links to further resources are provided to help inform parameter selection. This provides an easy to use interface for carrying out complex neuromodulation analyses.

Models can be specified using plain text files or directly using a scripting language. This allows more advanced workflows and integration with other tools. Additionally, this allows automated adjustment of stimulation parameters by automatically creating new data structures with updated parameters. This allows for sweeps through large ranges of parameters or the application of optimisation techniques to identify individually optimised parameters. Using this approach, models can be run and the results assessed, then the input parameters adjusted to increase the target activation. This can then be performed automatically for large numbers of iterations until the ideal parameters are identified.

This approach allows easy use of these complex techniques, lowering the barrier for applying them to new areas, while also allowing advanced workflows and extensability as well as integration with other software tools through the scripting interface.

B.3.2 Abstracted system

This approach to the modelling of neuromodulation allows for the complex details of image processing, finite element analysis, biophysics and optimisation to be abstracted away and replaced with a simple model specification interface. This allows users to carry out complex models without needing to deal with the details of how to carry out any of the individual analyses.

Detailed documentation is provided, a screenshot of which is shown in figure B.5. This outlines the steps required to specify and run models and to interpret the results, as well as the core theory required to understand the techniques used.

By abstracting away the technical details, we have significantly lowered the barrier to entry for carrying out detailed computational work in neuromodulation. This allows greater standardisation across models by using a common analysis pipeline and opens these techniques up to a greater audience. Further, this kind of abstraction is necessary for software for potential clinical application, where the technical details will need to be hidden with only what is required for clinical use exposed for the user.

B.3.3 Personalised models

We demonstrate that these software tools allow for complex, meaningful analyses to be carried out with relative ease. In the present example, we demonstrate that these tools can be used to demonstrate the activation of axons in the dorsal columns of the spinal cord using high amplitude stimulation at the T11 level with a dorsoventral electrode montage.



Figure B.5: Example of system documentation. All steps in model specification are documented in detail, including the parameters required and acceptable values. The implementation of the tools and the methods used are also outlined in order to allow more complex workflows to be developed using the scripting interface. This provides useful guidance for the specification and development of neuromodulation models.

Figure B.6A shows an axial slice of the standard clinical imaging used for analysis. Figure B.6B shows the segmentation of this image volume into discrete tissue types. These are then combined to form a 3D volumetric model as shown in figure B.6C.

Simulated electrodes were then mapped into the volume at the T11 level and the electric field produced by stimulation was solved. A 3D model of the vertebrae and electrodes with the electric field shown as streamlines is shown in figure B.6D, while figure B.6E shows the electric potential and the electric field superimposed on an axial slice of the original clinical imaging. This demonstrates that a potential is generated within the spinal canal with stimulation at this level.

Figure B.6F shows a sagittal slice of the imaging with example axon trajectories

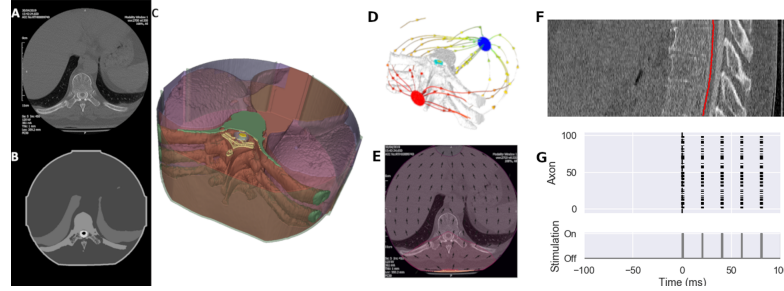


Figure B.6: Example of a personalised model. **A** Axial slice of original clinical imaging. **B** Axial slice of segmented imaging showing separation into discrete tissue types. **C** 3D model of tissue types reconstructed from segmentation. **D** 3D model of vertebrae and electrodes showing electric field produced by stimulation as streamlines. **E** Results of finite element analyses superimposed on clinical imaging, showing the electric potential and electric field produced by stimulation. **F** Sagittal slice of imaging showing example axon trajectories within the spinal cord. Biophysical models are projected onto these trajectories. **G** Raster plot showing the response of 100 axons within the spinal cord to stimulation. When stimulation is off there is no activity. When stimulation is turned on there are bursts of activity within the spinal cord. Non-invasive stimulation produces activity within the spinal cord in this individual using the electrode arrangement shown.

within the spinal cord shown in red. Biophysical axon models were mapped onto these trajectories and the effects of stimulation on these axons simulated. Figure B.6G shows a raster plot of the response of 100 axons within the spinal cord to multiple pulses of stimulation. There is no activity while stimulation is off, with bursts of activity throughout the spinal cord when stimulation is delivered. This shows that axons within the spinal cord are activated directly by non-invasive stimulation at T11 in this individual.

This demonstrates that it is possible to carry out complex analyses and to answer clinically meaningful questions using these tools, using a straightforward model specification interface.

B.4 Discussion

B.4.1 A tool for targeted neuromodulation

We present a simple to use tool that allows for complex models to be derived from routine clinical imaging. This allows for detailed modelling of neuromodulation to be carried out without the need for specific imaging sequences, extensive time spent

developing individualised models or expensive external software. This provides an integrated system for carrying out individualised modelling of neuromodulation for a range of applications.

Our proposed system has a well-abstracted model solver, allowing for computational work on neuromodulation to be carried out without requiring detailed knowledge of implementing imaging processing, finite element analysis, biophysics or optimisation analyses. This opens these approaches up to a far greater audience by reducing technical barrier to their use and brings these methods closer to clinical application by removing the need for extensive technical work to implement models. We provide a simple interface that allows models to be specified and run, while the solver engine ensures that the internal models and coupling between sections is carried out appropriately.

The system is also capable of carrying out highly complex analysis pipelines and being integrated into the workflows of other software. The scripting interface allows for any parameters to be altered and externalised for other algorithms. This allows, for example, automated optimisation of stimulation parameters based on achieving activation of a specified target structure.

This approach allows for models of a range of complexity to be developed, with potential applications to clinical neuromodulation and patient-specific therapies, as well as the development of novel neuromodulation techniques and devices and scientific applications to the investigation of the mechanisms underlying the efficacy of neuromodulation therapies.

In effect, this approach produces a function that maps user-selected stimulation parameters to activation of target structures. The activation produced by any set of stimulation parameters can then be interrogated. This provides a detailed individualised model of the effects of neuromodulation. This individualised mapping function can then be optimised in order to individualise therapy, or can be used to develop new therapeutic approaches or answer scientific questions. This provides a valuable new approach in the field of translational neuromodulation.

B.4.2 Limitations

While the present approach is almost entirely automated, some manual steps are required. Specifically, while the need to develop model geometries is eliminated, clinical imaging must be segmented in order to run the required analyses. While this can be carried out using semi-automated tools, manual checking and correction remain important to ensure that results are reliable. While significantly faster than existing methods, this still requires some time commitment for large models where accuracy of results is critical. The development of a method for more accurate automated segmentation of clinical imaging for neuromodulation simulation is an area of ongoing research.

The existing system abstracts away the details of running complex analyses and makes this straightforward. However, carrying out complex custom analyses may require the use of the scripting interface. While well-documented, this requires learning some details of the system and may not be immediately accessible to users with little technical training. This is provided to allow for additional flexibility and for the system’s powerful analysis methods to be applied to complex custom analyses, but there is a tradeoff between usability and flexibility; carrying out complex analysis workflows outside of standard methods may require the user to learn some of the details of the scripting interface and how the tools work.

B.4.3 Conclusion

We present a generalisable system for carrying out detailed patient-specific modelling of neuromodulation. This implements and brings together techniques for deriving patient-specific models from clinical imaging, simulating electrode positioning to allow intervention simulation using pre-intervention imaging, solving for the electric field produced by stimulation using finite element analysis, biophysical modelling of axons and optimisation for patient-specific therapies. These methods are coupled together into a solver engine, and the details abstracted away to provide a simple user interface for specifying models.

This represents a significant advance in techniques in translational neuromodulation research. We demonstrate a method for producing highly detailed models from clinical imaging, which is easily extensible and can be incorporated into complex workflows. This provides a valuable platform for developing novel neuromodulation therapies, for investigating the mechanisms underlying neuromodulation and for developing patient-specific neuromodulation approaches.

C

A system for generating complex stimulation waveforms for human neuromodulation

Contents

C.1 Introduction	196
C.2 Methods	198
C.2.1 Circuit design	198
C.2.2 PCB manufacture	199
C.2.3 Device assembly	200
C.2.4 Firmware	200
C.2.5 Testing	201
C.3 Results	201
C.3.1 Stimulation device	201
C.3.2 Custom stimulation patterns	203
C.4 Discussion	204
C.4.1 A system for complex neuromodulation	204
C.4.2 Applications	204
C.4.3 Conclusion	205

C.1 Introduction

The application of electrical stimulation to activate nervous system structures for therapeutic purposes is well established, with applications to a range of pathologies

[10, 96, 125]. Chapter 1 describes the emerging application of these techniques to the restoration of function following spinal cord injury. Chapter 2 describes the physical basis of these therapies, while chapters 3, 4, 5 and 6 show how neuromodulation can be modelled to develop new therapies and optimise existing ones. Chapter 7 then shows how these modelling techniques can be applied to answer meaningful clinical questions and develop new insights for developing targeted neuromodulation strategies for achieving functional restoration.

These modelling approaches allow the effects of neuromodulation on individuals to be evaluated in detail. However, in order to be able to translate the insights gained from these models into clinical benefit, it is necessary to be able to produce stimulation with complex waveforms with tight control over the stimulation parameters in order to replicate the stimulation conditions modelled.

A simple example of this is the use of burst stimulation for non-invasive spinal cord stimulation. Burst stimulation was popularised in the context of implanted spinal cord stimulation as a step towards paraesthesia-free stimulation [129]. However, it has since been adopted by practitioners of non-invasive spinal cord stimulation and has gained widespread use in the area of functional restoration [39, 130]. This is intended to reduce discomfort by minimising activation of cutaneous afferents, allowing for stimulation at higher intensities. In order to be able to apply this as well as other complex stimulation waveforms, it is necessary to be able to generate them using a standard stimulator with appropriate safety ratings for experimental use while maintaining control of the stimulation parameters.

The Digitimer DS8R is a widely used commercially available electrical stimulator, appropriate for human use. However, similar to most other lab-grade stimulator systems, it produces only single stimulation pulses following a trigger input. These systems allow strict control over stimulation parameters, but an external system must be used to define the stimulation pattern and stimulation parameters. We therefore require a reliable custom system which provides detailed control of stimulation patterns in order to be able to translate the insights gained from simulation.

We present a system for generating complex stimulation patterns for application to human neuromodulation. This provides a method for easily adjusting the parameters of complex stimulation modalities and represents a significant aid to experimental neuromodulation work. Further, this makes testing the hypotheses generated in computational models possible by providing tight control over stimulation with complex waveforms without requiring the development of custom hardware for every possible set of parameters.

C.2 Methods

C.2.1 Circuit design

The overall circuit design is shown in figure C.1. The system is controlled by a 32-bit ATmega4809 microcontroller. This is powered over a 5V USB connection. It is connected to the stimulator system's trigger input via a standard BNC connection, through which it generates complex stimulation patterns by controlling stimulation timing using TTL logic levels.

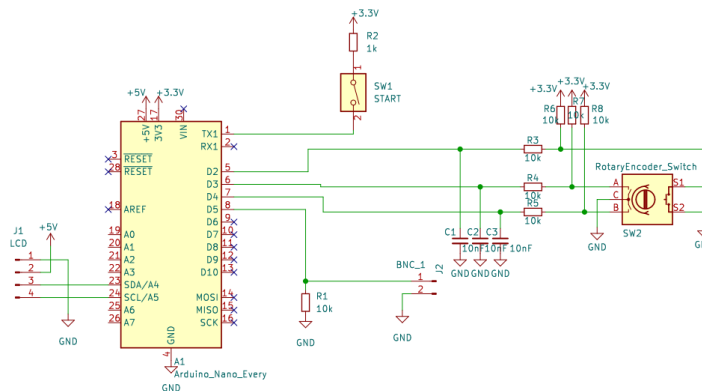
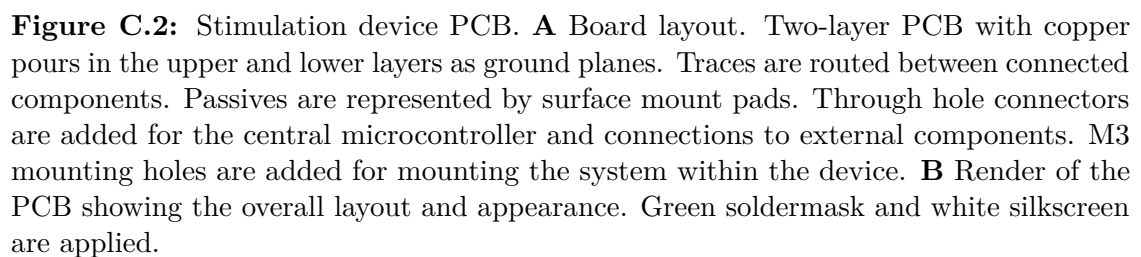


Figure C.1: Stimulation device schematic. A central microcontroller generates the stimulation patterns according to user input using a rotary encoder. The current settings are displayed on an LCD screen. The device is attached to the trigger input of a stimulator device using a BNC port to provide an output stage. An enable switch is included, which can be used to terminate output as a safety mechanism if it is switched off.

Stimulation settings are controlled using a rotary encoder with push-button switch. This allows for stimulation parameters to be altered and multiple settings

An enable switch is included to allow rapid cessation of stimulation if necessary. Triggers will only be sent to the stimulator device if this physical switch is set. This provides an additional layer of control and safety.

The central circuit board was manufactured as a two-layer PCB. Figure C.2A shows the PCB layout, while figure C.2B shows a render of the resulting circuit board.



Copper pours were used as ground planes on both layers. Traces were then routed between connected components. All passive components were included as surface mount pads. Through-hole connectors were added for mounting the central microcontroller. Through hole 2.54mm pitch connectors were added for connecting to external components such as the BNC port, enable switch, rotary encoder and LCD screen.

M3 mounting holes were placed in the corners to allow the PCB to be mounted within the device. Green soldermask and white silkscreen were applied to the finished board.

C.2.3 Device assembly

All surface-mount components were positioned using a pick-and-place machine and soldered using a reflow oven. Surface mount components were hand-soldered. The central microcontroller was mounted on the circuit board and Molex snap-on connectors were soldered to provide connections to external components. Wires and corresponding Molex snap-on connectors were soldered to the BNC port, enable switch, rotary encoder and LCD screen.

The device was constructed within a custom enclosure. A 3D printed insert was produced with mounting holes for the PCB. This was then inserted within the enclosure to ensure it remained fixed in position. Custom 3D printed front and back panels were produced with mounting holes for external components. The BNC port, enable switch, rotary encoder and LCD screen were then mounted on these panels. A female-to-female USB expander was attached to the microcontroller and mounted to the back panel in order to allow it to be powered from a panel connection.

All components were then connected internally and the enclosure was screwed shut.

C.2.4 Firmware

Custom firmware was written in C++ to control the device. Drivers were written to read the rotary encoder, write to the LCD screen and control the trigger output.

Turning and pressing the rotary encoder were programmed to trigger interrupts. Turning the rotary encoder was set to change the target parameter, while pressing the push-button changed the target parameter. The target parameter and current value were displayed on the LCD screen.

The device was programmed to allow stimulation with a number of different stimulation modalities, including standard stimulation, high frequency stimulation

and burst stimulation. The current mode was changed by altering the mode parameter using the rotary encoder.

Standard stimulation allowed for stimulation in the sub-1 kHz range, corresponding to most standard stimulation protocols. In this mode, the frequency parameter could be altered. In high frequency mode, stimulation at multiple kilohertz could be applied. In this mode, the frequency parameter could be altered over a much wider range, although with less resolution. This allowed for the use of high-frequency methods. In burst stimulation mode, bursts of pulses are applied. In this mode, the burst frequency, intra-burst frequency and burst duration could be controlled. This allowed for the application of complex burst patterns of stimulation.

In all modes, stimulation was triggered by altering the logic level at the BNC trigger output when the enable switch was set. If the enable switch was turned off, all output was disabled.

C.2.5 Testing

The output of the device was verified by attaching the trigger output to an oscilloscope and comparing the intended trigger patterns to the measured patterns. The device was then attached to a Digitimer DS8R stimulator and the ability to generate the target stimulation patterns was evaluated.

C.3 Results

C.3.1 Stimulation device

The final stimulation device allows precise control of complex stimulation patterns using a straightforward user interface, allowing for the testing of wide ranges of stimulation patterns in humans.

Figure C.3A shows the device within the enclosure with the top removed to visualise the internals of the device. The circuit board is mounted centrally on a 3D printed insert. The central microcontroller and connectors for external components are soldered to this circuit board. Corresponding connectors and wires connect this board to external components mounted on custom front and back panels to allow easy

access for users. The microcontroller is powered via its USB port, which is connected via a female-to-female expander to a panel-mount USB connector on the back panel.




Figure C.3: Stimulation device hardware. **A** Device within enclosure with top removed. The circuit board is mounted on a 3D printed insert. The microcontroller and connectors for external components are soldered to the circuit board. Connectors and wires connect the board to external components, which are mounted on the front and back panels. A USB cable connects the microcontroller's power input to a panel-mount USB connector. **B** Front panel of the device. The rotary encoder allows the user to alter the target stimulation parameter, while the current settings are displayed on the LCD screen. **C** Back panel of the device. The enable switch determines whether any output is delivered. The device is connected to the trigger input of a stimulator system using the BNC port. It is powered using the USB connector.

Figure C.3B shows the fully assembled device with front panel visible. The user interacts with the device to set stimulation parameters using the rotary encoder, with the present stimulation mode and current settings displayed on the LCD screen. Figure C.3C shows the back panel. The enable switch can be used to activate or deactivate stimulator output, providing an additional layer of safety if stimulation needs to be terminated. The device is connected to the output stimulator's trigger input using the BNC port. The device is powered using 5V via the USB port.

This provides an easy-to-use device that can produce complex, useful stimulation patterns. The stimulation parameters can be easily controlled using a simple user interface. This allows for wide ranges of parameters using complex stimulation modalities to be easily tested in humans without the need for extensive custom hardware.

C.3.2 Custom stimulation patterns

The custom device produces complex stimulation patterns determined by user input. Figure C.4A shows an example of standard stimulation. Stimulation pulses are delivered at regular intervals according a frequency set by user input. Figure C.4B shows a corresponding example of high-frequency stimulation. Regular stimulation is applied at high frequencies according to user selected parameters, allowing the effects of high frequency stimulation to be tested.




Figure C.4: Example output of stimulation device. **A** Standard stimulation. Example output with regular, low frequency stimulation. The frequency of stimulation is controller by the user. **B** High frequency stimulation. Example output with regular, high frequency stimulation. The frequency of stimulation is controlled by the user. **C** Burst stimulation. Example output with burst stimulation. The burst frequency, intra-burst frequency and burst duration are controlled by the user.

Figure C.4C shows an example of burst stimulation. The burst frequency, the intra-burst frequency and the burst duration are set by the user. Stimulation according to these parameters is then delivered.

This provides a means of delivering complex stimulation patterns with tight control over parameters using a simple user interface. This allows for wide ranges of parameters to be easily tested in humans using clinically relevant stimulation modalities.

C.4 Discussion

C.4.1 A system for complex neuromodulation

We demonstrate a device that allows for the delivery of complex neuromodulation protocols using a standard laboratory-grade stimulator system. This allows for straightforward adjustment of stimulation parameters and reliable delivery of clinically relevant stimulation patterns according to user-defined settings.

A simple microcontroller system with a straightforward user interface allows for the pattern of trigger outputs to be altered over a wide range of potential values without the need for reprogramming or altering hardware to make changes. This allows a single standalone device to be deployed in conjunction with a stimulator as an output stage in order to deliver a wide variety of stimulation patterns over many variations in parameters.

This system overcomes the major issue with standard single-shot stimulator systems that are suitable for human use by removing the need for complex pipelines or custom hardware for making adjustments to the stimulation protocol or adjusting stimulation parameters. We show that this simple device can reliably produce complex, clinically relevant stimulation patterns over a wide range of parameters with strict user control and can be used to drive stimulator systems to provide complex neuromodulation therapies in humans. This opens the possibility of providing highly detailed individual calibration of stimulation and for evaluation of novel stimulation modalities by providing easy control over a wide range of parameters.

C.4.2 Applications

Non-invasive stimulation is increasingly being applied to the spinal cord [38–40]. However, only a small range of parameters have been thoroughly tested. Epidural

spinal cord stimulation has seen great advances due to an expanding number of clinically useful stimulation modalities, such as burst stimulation [129] and high frequency stimulation [131]. The ability to explore a wider range of stimulation patterns and stimulation modalities has the potential to increase the application of non-invasive spinal cord stimulation, and also allows for greater individualised calibration of stimulation parameters by offering finer control over stimulation.

These developments are equally relevant to areas outside of the spinal cord. Stimulation of peripheral nerves is increasingly relevant in a number of contexts. However, application of non-invasive stimulation to other areas is limited by many of the same factors limiting its application to the spinal cord, including the generation of cutaneous discomfort. The ability to use other stimulation modalities may allow for greater comfort and increase the therapeutic window of tolerated stimulation parameters.

The ability to finely control a wide range of parameters is useful to the development of novel stimulation modalities and parameter combinations throughout the body, potentially opening up new therapeutic targets and possibilities. It is also valuable for the individualisation of existing therapies, with finer control allowing tight individual calibration of stimulation, preventing unnecessary discomfort.

C.4.3 Conclusion

We demonstrate a system for generating clinically relevant stimulation patterns using a device with a simple user interface. This allows for strict control of a wide variety of parameters, opening the possibility for exploring new stimulation modalities in many therapeutic targets.

Further, this system's tight control allows for the detailed results of individualised computational models of neuromodulation to be replicated, allowing individualised stimulation parameters to be translated into an experimental setting for the assessment of new individualised approaches to neuromodulation.

References

- [1] Christopher S. Ahuja et al. *Traumatic spinal cord injury*. Apr. 2017.
- [2] Michael G. Fehlings et al. *A Clinical Practice Guideline for the Management of Acute Spinal Cord Injury: Introduction, Rationale, and Scope*. Sept. 2017.
- [3] CD Witiw and MG Fehlings. “Acute spinal cord injury”. In: *Journal of Spinal Disorders* (2015). URL: <https://www.ingentaconnect.com/content/wk/jspit/2015/00000028/00000006/art00002>.
- [4] J. W. Fawcett et al. “Guidelines for the conduct of clinical trials for spinal cord injury as developed by the ICCP panel: spontaneous recovery after spinal cord injury and statistical power needed for therapeutic clinical trials”. In: *Spinal cord* 45 (3 Mar. 2007), pp. 190–205. URL: <https://pubmed.ncbi.nlm.nih.gov/17179973/>.
- [5] Kim D. Anderson. “Targeting recovery: priorities of the spinal cord-injured population”. In: *Journal of neurotrauma* 21 (10 Oct. 2004), pp. 1371–1383. URL: <https://pubmed.ncbi.nlm.nih.gov/15672628/>.
- [6] David McDaid et al. “Understanding and modelling the economic impact of spinal cord injuries in the United Kingdom”. In: *Spinal Cord* 2019 57:9 57 (9 May 2019), pp. 778–788. URL: <https://www.nature.com/articles/s41393-019-0285-1>.
- [7] Susan Harkema et al. “Effect of epidural stimulation of the lumbosacral spinal cord on voluntary movement, standing, and assisted stepping after motor complete paraplegia: a case study”. In: *The Lancet* 377 (9781 2011), pp. 1938–1947. URL: [http://dx.doi.org/10.1016/S0140-6736\(11\)60547-3](http://dx.doi.org/10.1016/S0140-6736(11)60547-3).
- [8] Justin Vogt et al. “Recovery of Over-Ground Walking after Chronic Motor Complete Spinal Cord Injury”. In: (2018), pp. 1–7.
- [9] Fabien B. Wagner et al. “Targeted neurotechnology restores walking in humans with spinal cord injury”. In: *Nature* 2018 563:7729 563 (7729 Oct. 2018), pp. 65–71. URL: <https://www.nature.com/articles/s41586-018-0649-2>.
- [10] Henri Lorach et al. “Walking naturally after spinal cord injury using a brain–spine interface”. In: *Nature* 618 (7963 June 2023), pp. 126–133.
- [11] Andreas Rowald et al. “Activity-dependent spinal cord neuromodulation rapidly restores trunk and leg motor functions after complete paralysis”. In: *Nature Medicine* 28 (2 Feb. 2022), pp. 260–271.
- [12] Claudia Kathe et al. “The neurons that restore walking after paralysis”. In: *Nature* 611 (7936 Nov. 2022), pp. 540–547.
- [13] Jordan W. Squair et al. “Implanted System for Orthostatic Hypotension in Multiple-System Atrophy”. In: *New England Journal of Medicine* 386 (14 Apr. 2022), pp. 1339–1344.

- [14] Marco Capogrosso et al. “A Computational Model for Epidural Electrical Stimulation of Spinal Sensorimotor Circuits”. In: 33 (49 2013), pp. 19326–19340.
- [15] Eduardo Martin Moraud et al. “Mechanisms Underlying the Neuromodulation of Spinal Circuits for Correcting Gait and Balance Deficits after Spinal Cord Injury”. In: *Neuron* 89 (4 Feb. 2016), pp. 814–828.
- [16] Walter Troni et al. “Improved methodology for lumbosacral nerve root stimulation”. In: *Muscle and Nerve* 19 (5 May 1996), pp. 595–604.
- [17] M. R. Dimitrijevic et al. “Neurophysiological evaluation of chronic spinal cord stimulation in patients with upper motor neuron disorders”. In: *International rehabilitation medicine* 2 (2 1980), pp. 82–85. URL: <https://pubmed.ncbi.nlm.nih.gov/6969710/>.
- [18] Marco Capogrosso et al. “A brain–spine interface alleviating gait deficits after spinal cord injury in primates”. In: *Nature* 539 (7628 Nov. 2016), pp. 284–288. URL: <http://www.ncbi.nlm.nih.gov/pubmed/27830790><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5108412><http://www.nature.com/doifinder/10.1038/nature20118>.
- [19] Nikolaus Wenger et al. “Spatiotemporal neuromodulation therapies engaging muscle synergies improve motor control after spinal cord injury”. In: *Nature Medicine* 22 (2 Feb. 2016), pp. 138–145.
- [20] Claudia A. Angeli et al. “Altering spinal cord excitability enables voluntary movements after chronic complete paralysis in humans”. In: *Brain* 137 (5 2014), pp. 1394–1409.
- [21] K. Minassian et al. “Human lumbar cord circuitries can be activated by extrinsic tonic input to generate locomotor-like activity”. In: *Human Movement Science* 26 (2 Apr. 2007), pp. 275–295.
- [22] K. Minassian et al. “Stepping-like movements in humans with complete spinal cord injury induced by epidural stimulation of the lumbar cord: Electromyographic study of compound muscle action potentials”. In: *Spinal Cord* 42 (7 July 2004), pp. 401–416.
- [23] Rubia Van Den Brand et al. *Restoring Voluntary Control of Locomotion after Paralyzing Spinal Cord Injury*. 2012. URL: <https://www.science.org>.
- [24] Megan L. Gill et al. “Neuromodulation of lumbosacral spinal networks enables independent stepping after complete paraplegia”. In: *Nature Medicine* 24 (11 Nov. 2018), pp. 1677–1682.
- [25] Peter J. Grahn et al. “Enabling Task-Specific Volitional Motor Functions via Spinal Cord Neuromodulation in a Human With Paraplegia”. In: *Mayo Clinic Proceedings* 92 (4 Apr. 2017), pp. 544–554.
- [26] D Darrow et al. “Effect of epidural spinal cord stimulation after chronic spinal cord injury on volitional movement and cardiovascular function: study protocol for the phase II open label controlled E-STAND trial”. In: *BMJ Open* (2022).
- [27] Grégoire Courtine and Michael V. Sofroniew. *Spinal cord repair: advances in biology and technology*. 2019.

- [28] Gregoire Courtine et al. “Recovery of supraspinal control of stepping via indirect propriospinal relay connections after spinal cord injury”. In: *Nature Medicine* 14 (1 Jan. 2008), pp. 69–74.
- [29] Grégoire Courtine et al. “Transformation of nonfunctional spinal circuits into functional states after the loss of brain input”. In: *Nature Neuroscience* 12 (10 Oct. 2009), pp. 1333–1342.
- [30] Aya Takeoka et al. “Muscle spindle feedback directs locomotor recovery and circuit reorganization after spinal cord injury”. In: *Cell* 159 (7 Dec. 2014), pp. 1626–1639.
- [31] Lucia Friedli et al. “Pronounced species divergence in corticospinal tract reorganization and functional recovery after lateralized spinal cord injury favors primates”. In: *Science translational medicine* 7 (302 Aug. 2015). URL: <https://pubmed.ncbi.nlm.nih.gov/26311729/>.
- [32] Katherine C. Murray et al. “Recovery of motoneuron and locomotor function after spinal cord injury depends on constitutive activity in 5-HT_{2C} receptors”. In: *Nature Medicine* 16 (6 June 2010), pp. 694–700.
- [33] Isabela Peña Pino et al. “Long-Term Spinal Cord Stimulation After Chronic Complete Spinal Cord Injury Enables Volitional Movement in the Absence of Stimulation”. In: *Frontiers in Systems Neuroscience* 14 (June 2020).
- [34] M. Krenn et al. “Selectivity of transcutaneous stimulation of lumbar posterior roots at different spinal levels in humans”. In: *Biomedizinische Technik* 58 (SUPPL. 1 TRACK-A Aug. 2013).
- [35] Dmitry G Sayenko et al. “Spinal segment-specific transcutaneous stimulation differentially shapes activation pattern among motor pools in humans”. In: *J Appl Physiol* 118 (2015), pp. 1364–1374. URL: <http://www.jappphysiol.org>.
- [36] A. Maertens De Noordhout et al. “Percutaneous electrical stimulation of lumbosacral roots in man.” In: *Journal of Neurology, Neurosurgery Psychiatry* 51 (2 Feb. 1988), pp. 174–181. URL: <https://jnnp.bmj.com/content/51/2/174><https://jnnp.bmj.com/content/51/2/174.abstract>.
- [37] U. S. Hofstoetter et al. “Effects of transcutaneous spinal cord stimulation on voluntary locomotor activity in an incomplete spinal cord injured individual”. In: *Biomedizinische Technik* 58 (SUPPL. 1 TRACK-A Aug. 2013).
- [38] Yazı Al’joboori et al. “The Effects of Adding Transcutaneous Spinal Cord Stimulation (tSCS) to Sit-To-Stand Training in People with Spinal Cord Injury: A Pilot Study”. In: *Journal of clinical medicine* 9 (9 Sept. 2020), pp. 1–22. URL: <https://pubmed.ncbi.nlm.nih.gov/32858977/>.
- [39] Yury P. Gerasimenko et al. “Noninvasive Reactivation of Motor Descending Control after Paralysis”. In: *Journal of neurotrauma* 32 (24 Dec. 2015), pp. 1968–1980. URL: <https://pubmed.ncbi.nlm.nih.gov/26077679/>.
- [40] Dmitry G. Sayenko et al. “Self-assisted standing enabled by non-invasive spinal stimulation after spinal cord injury”. In: *Journal of Neurotrauma* 36 (9 May 2019), pp. 1435–1450.

- [41] Parag Gad et al. “Weight bearing over-ground stepping in an exoskeleton with non-invasive spinal cord neuromodulation after motor complete paraplegia”. In: *Frontiers in Neuroscience* 11 (JUN June 2017). URL: <https://pubmed.ncbi.nlm.nih.gov/28642680/>.
- [42] Parag N. Gad et al. “Non-invasive Neuromodulation of Spinal Cord Restores Lower Urinary Tract Function After Paralysis”. In: *Frontiers in Neuroscience* 12 (June 2018). URL: <https://pubmed.ncbi.nlm.nih.gov/30008661/>.
- [43] Alvaro Megía García et al. “Transcutaneous Spinal Cord Stimulation and Motor Rehabilitation in Spinal Cord Injury: A Systematic Review”. In: *Neurorehabilitation and Neural Repair* 34 (1 Jan. 2020), pp. 3–12. URL: <https://journals.sagepub.com/doi/full/10.1177/1545968319893298>.
- [44] Jonathan S Calvert et al. “Preferential activation of spinal sensorimotor networks via lateralized transcutaneous spinal stimulation in neurologically intact humans”. In: *J Neuro-physiol* 122 (2019), pp. 2111–2118. URL: www.jn.org.
- [45] Simon M Danner et al. “Can the human lumbar posterior columns be stimulated by transcutaneous spinal cord stimulation? A modeling study”. In: (2011).
- [46] Andreas Kuhn et al. “A model for transcutaneous current stimulation: simulations and experiments”. In: *Medical Biological Engineering Computing* 47:3 47 (3 Nov. 2008), pp. 279–289. URL: <https://link.springer.com/article/10.1007/s11517-008-0422-z>.
- [47] Josef Ladenbauer et al. “Stimulation of the human lumbar spinal cord with implanted and surface electrodes: a computer simulation study”. In: *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society* 18 (6 Dec. 2010), pp. 637–645. URL: <https://pubmed.ncbi.nlm.nih.gov/21138794/>.
- [48] F. Rattay, K. Minassian, and M. R. Dimitrijevic. “Epidural electrical stimulation of posterior structures of the human lumbosacral cord: 2. quantitative analysis by computer modeling”. In: *Spinal Cord* 38:8 38 (8 Aug. 2000), pp. 473–489. URL: <https://www.nature.com/articles/3101039>.
- [49] Scott F. Lempka et al. “Patient-Specific Analysis of Neural Activation During Spinal Cord Stimulation for Pain”. In: *Neuromodulation: Technology at the Neural Interface* 23 (5 July 2020), pp. 572–581.
- [50] Kabilar Gunalan et al. “Creating and parameterizing patient-specific deep brain stimulation pathway-activation models using the hyperdirect pathway as an example”. In: *PLoS ONE* 12 (4 2017), pp. 1–19.
- [51] Andreas Horn et al. “Lead-DBS v2: Towards a comprehensive pipeline for deep brain stimulation imaging”. In: *NeuroImage* 184 (Jan. 2019), pp. 293–316. URL: <https://pubmed.ncbi.nlm.nih.gov/30179717/>.
- [52] Jan Holsheimer. “Computer modelling of spinal cord stimulation and its contribution to therapeutic efficacy”. In: *Spinal cord* 36 (8 1998), pp. 531–540. URL: <https://pubmed.ncbi.nlm.nih.gov/9713921/>.
- [53] James Clerk Maxwell. “A dynamical theory of the electromagnetic field”. In: *A dynamical theory of the electromagnetic field* (1865).

- [54] David J. Griffiths. “Introduction to Electrodynamics”. In: *Introduction to Electrodynamics* (June 2017).
- [55] Hans Petter Langtangen and Kent-Andre Mardal. “Introduction to Numerical Methods for Variational Problems”. In: 21 (2019). URL: <http://link.springer.com/10.1007/978-3-030-23788-2>.
- [56] Conor Keogh. “Optimizing the neuron-electrode interface for chronic bioelectronic interfacing”. In: *Neurosurgical Focus* 49 (1 July 2020).
- [57] Stuart F Cogan. “Neural stimulation and recording electrodes.” In: *Annual review of biomedical engineering* 10 (2008), pp. 275–309.
- [58] DR Merrill, M Bikson, and JGR Jefferys. “Electrical stimulation of excitable tissue: design of efficacious and safe protocols”. In: *Journal of Neuroscience Methods* (2005). URL: https://www.sciencedirect.com/science/article/pii/S0165027004003826?casa_token=jrFvKN01MW0AAAAA:WgalSp4L80XYXk8w3uXjTR1_XC9JfjNRdodHrSK9RgIg6kGI0RlWPmi30EeJULrr06sFf-cbAg.
- [59] H. Helmholtz. “Studien über electrische Grenzschichten”. In: *Annalen der Physik* 243 (7 1879), pp. 337–382.
- [60] MO Heuschkel. *Fabrication of multi-electrode array devices for electrophysiological monitoring of in-vitro cell/tissue cultures*. 2001. URL: https://www.researchgate.net/profile/Marc_Heuschkel/publication/37413110_Fabrication_of_multi-electrode_array_devices_for_electrophysiological_monitoring_of_in-vitro_celltissue_cultures/links/02e7e5257dac024ccc000000.pdf.
- [61] David Leonard Chapman. “LI. A contribution to the theory of electrocapillarity”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 25 (148 Apr. 1913), pp. 475–481.
- [62] Hugo Fricke. “XXXIII. The theory of electrolytic polarization”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 14 (90 Aug. 1932), pp. 310–318.
- [63] E. Warburg. “Ueber das Verhalten sogenannter unpolarisirbarer Elektroden gegen Wechselstrom”. In: *Annalen der Physik* 303 (3 1899), pp. 493–499.
- [64] Sébastien Joucla and Blaise Yvert. “Modeling extracellular electrical neural stimulation: From basic understanding to MEA-based applications”. In: (2011).
- [65] Christopher R. Butson and Cameron C. McIntyre. “Role of electrode design on the volume of tissue activated during deep brain stimulation”. In: *Journal of Neural Engineering* 3 (1 Dec. 2005), p. 1. URL: <https://iopscience.iop.org/article/10.1088/1741-2560/3/1/001%20https://iopscience.iop.org/article/10.1088/1741-2560/3/1/001/meta>.
- [66] C. C. McIntyre and W. M. Grill. “Finite element analysis of the current-density and electric field generated by metal microelectrodes”. In: *Annals of Biomedical Engineering* 29 (3 2001), pp. 227–235.
- [67] Cameron C. McIntyre et al. “Cellular Effects of Deep Brain Stimulation: Model-Based Analysis of Activation and Inhibition”. In: *Journal of Neurophysiology* 91 (4 Apr. 2004), pp. 1457–1469.

- [68] F Rattay, RN Leao, and H Felix. “A model of the electrically excited human cochlear neuron. II. Influence of the three-dimensional cochlear structure on neural excitability”. In: *Hearing Research* (2001). URL: https://www.sciencedirect.com/science/article/pii/S0378595500002574?casa_token=QyCHk9FUfL0AAAAA:FnyU0JL9z_NckcL5njhnFkPBf1mC9J3u43qw5T5nT2EdBqZWJ4KV0LpvIywTXWpIVVwCBwuYKQ.
- [69] Svjetlana Miocinovic et al. “Computational analysis of subthalamic nucleus and lenticular fasciculus activation during therapeutic deep brain stimulation”. In: *Journal of Neurophysiology* 96 (3 2006), pp. 1569–1580.
- [70] Christof Koch and Idan Segev. *Methods in Neuronal Modeling*. 2003. URL: <https://mitpress.mit.edu/9780262517133/methods-in-neuronal-modeling/>.
- [71] A. L. Hodgkin and W. A. Rushton. “The electrical constants of a crustacean nerve fibre”. In: *Proceedings of the Royal Society of Medicine* 134 (873 Dec. 1946), pp. 444–479. URL: <https://pubmed.ncbi.nlm.nih.gov/20281590/>.
- [72] W. R. Holmes and W. Rall. “Estimating the electrotonic structure of neurons with compartmental models”. In: <https://doi.org/10.1152/jn.1992.68.4.1438> 68 (4 1992), pp. 1438–1452. URL: <https://journals.physiology.org/doi/10.1152/jn.1992.68.4.1438>.
- [73] Wilfrid Rall. “Core Conductor Theory and Cable Properties of Neurons”. In: *Comprehensive Physiology* (Dec. 1977), pp. 39–97. URL: <https://onlinelibrary.wiley.com/doi/full/10.1002/cphy.cp010103%20https://onlinelibrary.wiley.com/doi/abs/10.1002/cphy.cp010103%20https://onlinelibrary.wiley.com/doi/10.1002/cphy.cp010103>.
- [74] A. L. Hodgkin and A. F. Huxley. “A quantitative description of membrane current and its application to conduction and excitation in nerve”. In: *The Journal of Physiology* 117 (4 Aug. 1952), p. 500. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1392413/>.
- [75] Donald R. McNeal. “Analysis of a model for excitation of myelinated nerve”. In: *IEEE transactions on bio-medical engineering* 23 (4 1976), pp. 329–337. URL: <https://pubmed.ncbi.nlm.nih.gov/1278925/>.
- [76] Cameron C. McIntyre, Andrew G. Richardson, and Warren M. Grill. “Modeling the excitability of mammalian nerve fibers: influence of afterpotentials on the recovery cycle”. In: *Journal of neurophysiology* 87 (2 2002), pp. 995–1006. URL: <https://pubmed.ncbi.nlm.nih.gov/11826063/>.
- [77] Gordon Duffley et al. “Evaluation of methodologies for computing the deep brain stimulation volume of tissue activated”. In: *Journal of Neural Engineering* 16 (6 Oct. 2019), p. 066024. URL: <https://iopscience.iop.org/article/10.1088/1741-2552/ab3c95%20https://iopscience.iop.org/article/10.1088/1741-2552/ab3c95/meta>.
- [78] Christopher R. Butson et al. “Patient-Specific Analysis of the Volume of Tissue Activated During Deep Brain Stimulation”. In: *NeuroImage* 34 (2 Jan. 2007), p. 661. URL: <https://pubmed.ncbi.nlm.nih.gov/1794656/>.

- [79] Christopher R. Butson and Cameron C. McIntyre. “Tissue and electrode capacitance reduce neural activation volumes during deep brain stimulation”. In: *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology* 116 (10 Oct. 2005), pp. 2490–2500. URL: [/pmc/articles/PMC3653320/%20/pmc/articles/PMC3653320/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3653320/](https://pmc/articles/PMC3653320/%20/pmc/articles/PMC3653320/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3653320/).
- [80] Hugh Bostock and J. C. Rothwell. “Latent addition in motor and sensory fibres of human peripheral nerve.” In: *The Journal of Physiology* 498 (1997). URL: <https://api.semanticscholar.org/CorpusID:8056437>.
- [81] H Bostock J Howells L Trevillion and D Burke. “The voltage dependence I(h) in human myelinated axons.” In: *The Journal of Physiology* 590 (2012).
- [82] E J Peterson, O Izad, and D J Tyler. “Predicting myelinated axon activation using spatial characteristics of the extracellular field”. In: *Journal of Neural Engineering* 8.4 (July 2011), p. 046030. URL: <https://dx.doi.org/10.1088/1741-2560/8/4/046030>.
- [83] Matthew A. Schiefer, Ronald J. Triolo, and Dustin J. Tyler. “A Model of Selective Activation of the Femoral Nerve With a Flat Interface Nerve Electrode for a Lower Extremity Neuroprosthesis”. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 16 (2008), pp. 195–204. URL: <https://api.semanticscholar.org/CorpusID:10315198>.
- [84] Amorn Wongsarnpigoon, John P. Woock, and Warren M. Grill. “Efficiency Analysis of Waveform Shape for Electrical Excitation of Nerve Fibers”. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 18 (2010), pp. 319–328. URL: <https://api.semanticscholar.org/CorpusID:9158355>.
- [85] K.L. Kilgore and N. Bhadra. “Nerve conduction block utilising high-frequency alternating current.” In: *Medical and Biological Engineering and Computing* 42 (2004), pp. 394–406.
- [86] L.E. Medina and W.M. Grill. “Nerve excitation using an amplitude-modulated signal with kilohertz-frequency carrier and non-zero offset.” In: *Journal of NeuroEngineering and Rehabilitation* (2016).
- [87] J D Miles et al. “Effects of ramped amplitude waveforms on the onset response of high-frequency mammalian nerve block”. In: *Journal of Neural Engineering* 4.4 (Nov. 2007), p. 390. URL: <https://dx.doi.org/10.1088/1741-2560/4/4/005>.
- [88] Mattias Åström et al. “Relationship between neural activation and electric field distribution during deep brain stimulation”. In: *IEEE Transactions on Biomedical Engineering* 62 (2 Feb. 2015), pp. 664–672.
- [89] Frank Rattay. “Analysis of Models for External Stimulation of Axons”. In: *IEEE Transactions on Biomedical Engineering* BME-33 (10 1986), pp. 974–977.
- [90] Frank Rattay. “Analysis of Models for Extracellular Fiber Stimulation”. In: *IEEE Transactions on Biomedical Engineering* 36 (7 1989), pp. 676–682.
- [91] Jorge Nocedal and Stephen Wright. *Numerical Optimization*. Springer New York, 2006. URL: <http://link.springer.com/10.1007/978-0-387-40065-5>.
- [92] N.T. Carnevale and M.L. Hines. *The NEURON Book*. Cambridge University Press, 2006. URL: <https://www.neuron.yale.edu/phpBB/viewtopic.php?t=73>.

- [93] J. Noailly, A. Malandrino, and F. Galbusera. “Computational modelling of spinal implants”. In: *Computational Modelling of Biomechanics and Biotribology in the Musculoskeletal System: Biomaterials and Tissues* (Jan. 2014), pp. 447–484.
- [94] PA Hasgall et al. *IT’IS Database for thermal and electromagnetic parameters of biological tissues*. 2018. URL: itis.swiss/database.
- [95] *COMSOL - Software for Multiphysics Simulation*. URL: <https://www.comsol.com/>.
- [96] Krishna Kumar et al. “Spinal cord stimulation versus conventional medical management for neuropathic pain : A multicentre randomised controlled trial in patients with failed back surgery syndrome”. In: 132 (2007), pp. 179–188.
- [97] Judith A Turner et al. “Spinal cord stimulation for patients with failed back surgery syndrome or complex regional pain syndrome : a systematic review of effectiveness and complications”. In: 108 (2004), pp. 137–147.
- [98] Nagy A Mekhail, Manu Mathews, and Fady Nageeb. “Retrospective Review of 707 Cases of Spinal Cord Stimulation : Indications and Complications”. In: 11 (2 2011), pp. 148–153.
- [99] Parag Gad et al. “Non-invasive activation of cervical spinal networks after severe paralysis”. In: *Journal of Neurotrauma* 35 (18 Sept. 2018), pp. 2145–2158. URL: <https://pubmed.ncbi.nlm.nih.gov/29649928/>.
- [100] Ursula S. Hofstoetter et al. “Transcutaneous Spinal Cord Stimulation Induces Temporary Attenuation of Spasticity in Individuals with Spinal Cord Injury”. In: *Journal of Neurotrauma* 37 (3 Feb. 2020), pp. 481–493. URL: <https://pubmed.ncbi.nlm.nih.gov/31333064/>.
- [101] Nir Grossman et al. “Noninvasive Deep Brain Stimulation via Temporally Interfering Electric Fields”. In: *Cell* 169 (6 June 2017), 1029–1041.e16. URL: <http://dx.doi.org/10.1016/j.cell.2017.05.024>.
- [102] Ron Kikinis, Steve D. Pieper, and Kirby G. Vosburgh. “3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support”. In: *Intraoperative Imaging and Image-Guided Therapy* (2014), pp. 277–289. URL: https://link.springer.com/chapter/10.1007/978-1-4614-7657-3_19.
- [103] The CGAL Project. *{CGAL} User and Reference Manual*. 5.1.1. CGAL Editorial Board, 2020. URL: <https://doc.cgal.org/5.1.1/Manual/packages.html>.
- [104] Qianqian Fang and David A. Boas. “Tetrahedral mesh generation from volumetric binary and grayscale images”. In: 2009, pp. 1142–1145.
- [105] Ehsan Mirzakhali et al. “Biophysics of Temporal Interference Stimulation”. In: *Cell Systems* 11.6 (2020), 557–572.e5. URL: <https://www.sciencedirect.com/science/article/pii/S2405471220303720>.
- [106] Staniša Raspopović et al. “Finite element and biophysics modelling of intraneural transversal electrodes: Influence of active site shape”. In: *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC’10* (2010), pp. 1678–1681.

- [107] Stanisa Raspopovic et al. “Experimental validation of a hybrid computational model for selective stimulation using transverse intrafascicular multichannel electrodes”. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20 (3 2012), pp. 395–404.
- [108] Stanisa Raspopovic, Marco Capogrosso, and Silvestro Micera. “A computational model for the stimulation of rat sciatic nerve using a transverse intrafascicular multichannel electrode”. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 19 (4 2011), pp. 333–344.
- [109] James J. Fitzgerald. “Suppression of scarring in peripheral nerve implants by drug elution”. In: *Journal of Neural Engineering* 13 (2 Feb. 2016).
- [110] Laurence Abbott and Peter Dayan. *Theoretical Neuroscience*. 2001. URL: <https://mitpress.mit.edu/9780262041997/theoretical-neuroscience/>.
- [111] Guido Van Rossum and Fred Drake. *The Python Language Reference — Python 3.12.0 documentation*. 2009. URL: <https://docs.python.org/3/reference/index.html>.
- [112] Matthew W. Scroggs et al. “Construction of Arbitrary Order Finite Element Degree-of-Freedom Maps on Polygonal and Polyhedral Cell Meshes”. In: *ACM Transactions on Mathematical Software (TOMS)* 48 (2 May 2022). URL: <https://dl.acm.org/doi/10.1145/3524456>.
- [113] Guilherme B. Saturnino et al. “SimNIBS 2.1: A Comprehensive Pipeline for Individualized Electric Field Modelling for Transcranial Brain Stimulation”. In: *Brain and Human Body Modeling: Computational Human Modeling at EMBC 2018* (Jan. 2019), pp. 3–25. URL: https://link.springer.com/chapter/10.1007/978-3-030-21293-3_1.
- [114] Yu Huang et al. “ROAST: An Open-Source, Fully-Automated, Realistic Volumetric-Approach-Based Simulator For TES”. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference 2018* (Oct. 2018), pp. 3072–3075. URL: <https://pubmed.ncbi.nlm.nih.gov/30441043/>.
- [115] Julia P. Slopsema et al. “Clinical deep brain stimulation strategies for orientation-selective pathway activation”. In: *Journal of neural engineering* 15 (5 Sept. 2018). URL: <https://pubmed.ncbi.nlm.nih.gov/30095084/>.
- [116] Daria Nesterovich Anderson et al. “Anodic stimulation misunderstood: preferential activation of fiber orientations with anodic waveforms in deep brain stimulation”. In: *Journal of neural engineering* 16 (1 Feb. 2019). URL: <https://pubmed.ncbi.nlm.nih.gov/30275348/>.
- [117] Marco Capogrosso et al. “Advantages of soft subdural implants for the delivery of electrochemical neuromodulation therapies to the spinal cord”. In: *Journal of Neural Engineering* 15 (2 Apr. 2018), p. 026024. URL: <http://www.ncbi.nlm.nih.gov/pubmed/29339580%20http://stacks.iop.org/1741-2552/15/i=2/a=026024?key=crossref.771e50657e4c4ca1b709bb5b90a267b5>.
- [118] N Greiner et al. “Recruitment of upper-limb motoneurons with epidural electrical stimulation of the cervical spinal cord”. In: *Nature Communications* 435 (2021).

- [119] H. K. P. Feirabend et al. “Morphometry of human superficial dorsal and dorsolateral column fibres: significance to spinal cord stimulation”. In: *Brain* 125.5 (May 2002), pp. 1137–1149. eprint: <https://academic.oup.com/brain/article-pdf/125/5/1137/17865033/1251137.pdf>. URL: <https://doi.org/10.1093/brain/awf111>.
- [120] Carlos J. Anaya et al. “Evoked Potentials Recorded From the Spinal Cord During Neurostimulation for Pain: A Computational Modeling Study”. In: *Neuromodulation: Technology at the Neural Interface* 23.1 (2020), pp. 64–73. URL: <https://www.sciencedirect.com/science/article/pii/S1094715921020705>.
- [121] Marco Catani and Michel Thiebaut de Schotten. “A diffusion tensor imaging tractography atlas for virtual in vivo dissections”. In: *Cortex; a journal devoted to the study of the nervous system and behavior* 44 (8 2008), pp. 1105–1132. URL: <https://pubmed.ncbi.nlm.nih.gov/18619589/>.
- [122] Johannes J. Struijk, Jan Holsheimer, and Herman B.K. Boom. “Excitation of dorsal root fibers in spinal cord stimulation: a theoretical study”. In: *IEEE transactions on bio-medical engineering* 40 (7 1993), pp. 632–639. URL: <https://pubmed.ncbi.nlm.nih.gov/8244424/>.
- [123] Elizabeth A. Bye et al. “Transcutaneous spinal cord stimulation combined with locomotor training to improve walking ability in people with chronic spinal cord injury: study protocol for an international multi-centred double-blinded randomised sham-controlled trial (eWALK)”. In: *Spinal Cord* 2021 60:6 60 (6 Jan. 2022), pp. 491–497. URL: <https://www.nature.com/articles/s41393-021-00734-1>.
- [124] Michael S. Okun. “Deep-Brain Stimulation for Parkinson’s Disease”. In: <https://doi.org/10.1056/NEJMct1208070> 367 (16 Oct. 2012), pp. 1529–1538. URL: <https://www.nejm.org/doi/full/10.1056/NEJMct1208070>.
- [125] Günther Deuschl et al. “A Randomized Trial of Deep-Brain Stimulation for Parkinson’s Disease”. In: <https://doi.org/10.1056/NEJMoa060281> 355 (9 Aug. 2006), pp. 896–908. URL: <https://www.nejm.org/doi/full/10.1056/nejmoa060281>.
- [126] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 2020 585:7825 585 (7825 Sept. 2020), pp. 357–362. URL: <https://www.nature.com/articles/s41586-020-2649-2>.
- [127] John D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science and Engineering* 9 (3 2007), pp. 90–95.
- [128] Fabian Pedregosa FABIANPEDREGOSA et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (85 2011), pp. 2825–2830. URL: <http://jmlr.org/papers/v12/pedregosa11a.html>.
- [129] Dirk De Ridder et al. “Burst spinal cord stimulation: toward paresthesia-free pain suppression”. In: *Neurosurgery* 66 (5 May 2010), pp. 986–990. URL: <https://pubmed.ncbi.nlm.nih.gov/20404705/>.

- [130] Parag Gad et al. “Transcutaneous Spinal Neuromodulation Reorganizes Neural Networks in Patients with Cerebral Palsy”. In: *Neurotherapeutics* 18 (3 July 2021), p. 1953. URL: [/pmc/articles/PMC8608961/%20/pmc/articles/PMC8608961/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8608961/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8608961/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8608961/).
- [131] Leonardo Kapural et al. “Novel 10-kHz High-frequency Therapy (HF10 Therapy) Is Superior to Traditional Low-frequency Spinal Cord Stimulation for the Treatment of Chronic Back and Leg Pain: The SENZA-RCT Randomized Controlled Trial”. In: *Anesthesiology* 123 (4 Oct. 2015), pp. 851–860. URL: <https://pubmed.ncbi.nlm.nih.gov/26218762/>.