Translational pipelines for closed-loop neuromodulation

Robert Toth

Brasenose College

A thesis submitted for the degree of

Doctor of Philosophy

March 2024



MRC Brain Network Dynamics Unit

Nuffield Department of Clinical Neurosciences

University of Oxford

Oxford, United Kingdom

Acknowledgements

First and foremost, I would like to express my gratitude to my supervisors, Andrew Sharott and Tim Denison, for their patient support and the creative freedom they afforded me throughout the years.

Each of my projects relied on a number of collaborators, none more so than work on the DyNeuMo. I am extremely grateful to Mayela Zamora for her work on filter design and for being a good friend. I am also grateful to Moaad Benjaber for his detailed evaluation work, Brook Perry and Yangfan Peng for their advice on 3D printing, Jonathan Baker and Keith Purpura for carrying out pilot experiments, and Allie Deli for advice on data analysis. I'd also like to recognise the team at Bioinduction for their partnership on the DyNeuMo, in particular Tom Gillbe for his support in creating our experimental firmware releases.

I am grateful to Colin McNamara for entrusting me with unpublished work on phase estimation, and to Jess Myatt for her immense experimental efforts. I also thank Karen Wendt for carrying out a literature review on phase estimation algorithms.

I'd also like to acknowledge Bettina Schwab and Antonio Valentín for contributing clinical data to help specify the various presented designs.

Without particular order, I wish to thank Bence Halpern, Dmitrij Szamozvancev, Aime Wren, Giulio Spagnol and Duy An Tran for their friendship. Finally, I would like to extend my deepest appreciation to my family – I could not have done this without you.

Abstract

Closed-loop neuromodulation systems have shown significant potential for addressing unmet needs in the treatment of disorders of the central nervous system, yet progress towards clinical adoption has been slow. Advanced technological developments often stall in the preclinical stage by failing to account for the constraints of implantable medical devices, and due to the lack of research platforms with a translational focus. This thesis presents the development of three clinically relevant research systems focusing on refinements of deep brain stimulation therapies.

First, we introduce a system for synchronising implanted and external stimulation devices, allowing for research into multi-site stimulation paradigms, cross-region neural plasticity, and questions of phase coupling. The proposed design aims to sidestep the limited communication capabilities of existing commercial implant systems in providing a stimulation state readout without reliance on telemetry, creating a cross-platform research tool.

Next, we present work on the Picostim-DyNeuMo adaptive neuromodulation platform, focusing on expanding device capabilities from activity and circadian adaptation to bioelectric marker–based responsive stimulation. Here, we introduce a computationally optimised implementation of a popular band power–estimation algorithm suitable for deployment in the DyNeuMo system. The new algorithmic capability was externally validated to establish neural state classification performance in two widely-researched use cases: Parkinsonian beta bursts and seizures. For in vivo validation, a pilot experiment is presented demonstrating responsive neurostimulation to cortical alpha-band activity in a non-human primate model for the modulation of attention state.

Finally, we turn our focus to the validation of a recently developed method to provide computationally efficient real-time phase estimation. Following theoretical analysis, the method is integrated into the commonly used Intan electrophysiological recording platform, creating a novel closed-loop optogenetics research platform. The performance of the research system is characterised through a pilot experiment, targeting the modulation of cortical theta-band activity in a transgenic mouse model.

Arising publications

R. Toth *et al.*, "Frequency and phase synchronization in distributed (implantable-transcutaneous) neural interfaces," in *Proceedings of the 2019 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3831–3834, Jul. 2019.

R. Toth *et al.*, "DyNeuMo Mk-2: An investigational circadian-locked neuromodulator with responsive stimulation for applied chronobiology," in *Proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 3433–3440, Oct. 2020.

M. Zamora *et al.*, "Case report: Embedding "digital chronotherapy" into medical devices – a canine validation for controlling status epilepticus through multi-scale rhythmic brain stimulation," *Frontiers in Neuroscience*, vol. 15, p. 734265, Sep. 2021.

M. Zamora *et al.*, "DyNeuMo Mk-1: Design and pilot validation of an investigational motion-adaptive neurostimulator with integrated chronotherapy," *Experimental Neurology*, vol. 351, p. 113977, May 2022.

A. Kavoosi *et al.*, "Computationally efficient neural network classifiers for next generation closed loop neuromodulation therapy – a case study in epilepsy," in *Proceedings of the 2022 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 288–291, Jul. 2022.

J. L. Baker *et al.*, "Regulation of Arousal and Performance of a Healthy Non-Human Primate Using Closed-Loop Central Thalamic Deep Brain Stimulation," in *Proceedings of the 2023 International IEEE/EMBS Conference on Neural Engineering (NER)*, Apr. 2023.

Table of contents

List of figures **10**

List of tables **12**

List of abbreviations **13**

1 Introduction **15**

1.1 The origin of neural signals . . . 17
1.1.1 Single unit activity . . . 17
1.1.2 Synaptic signalling . . . 18
1.1.3 Multi-unit activity . . . 19
1.1.4 Background-unit activity . . . 20
1.1.5 Local field potential . . . 20
1.2 Neural plasticity . . . 21
1.2.1 Hebbian plasticity . . . 21
1.2.2 Induced plasticity . . . 22
1.3 The recording of neural signals . . . 25
1.3.1 Single-unit activity . . . 26
1.3.2 Multi-unit activity . . . 26
1.3.3 Local field potential . . . 27
1.3.4 Stereotactic electroencephalography . . . 28
1.3.5 Electrocorticography . . . 28
1.3.6 Electroencephalography . . . 29
1.3.7 Non-electrical methods . . . 29
1.4 Neurostimulation . . . 31
1.4.1 Non-invasive stimulation . . . 31
1.4.2 Intracranial stimulation . . . 33
1.5 Closed-loop systems . . . 38
1.5.1 Overview . . . 38
1.5.2 Safety consideration . . . 39
1.5.3 Stimulation artefacts . . . 41
1.5.4 Technological constraints . . . 44
1.6 Overview . . . 46
1.7 Aims . . . 48

2 Multi-site stimulation for therapy research **49**
2.1 Introduction . . . 49
2.2 System proposal . . . 52
2.2.1 Signal range . . . 53
2.2.2 External stimulation . . . 53
2.2.3 Artefact tolerance . . . 55
2.2.4 Frequency range . . . 56
2.2.5 Resolution . . . 56
2.2.6 Accuracy . . . 57
2.2.7 Latency . . . 57
2.2.8 Medical safety . . . 57
2.3 Technical method . . . 59
2.3.1 Phase-locked loop . . . 59
2.3.2 Realisation . . . 60
2.4 External validation . . . 62
2.5 Discussion . . . 64
2.5.1 Overview . . . 64
2.5.2 Limitations . . . 64
2.5.3 Future development . . . 65
2.6 Contribution . . . 66

3 Real-time power estimation for closed-loop stimulation **67**
3.1 Introduction . . . 67
3.1.1 The Picostim-DyNeuMo platform . . . 68
3.1.2 Research need . . . 73
3.1.3 Aims . . . 74
3.2 Technical methods . . . 75
3.2.1 Signal Chain . . . 75
3.2.2 Algorithm outline . . . 75
3.3 Development . . . 77
3.3.1 General purpose filter block . . . 77
3.3.2 Offset removal high-pass filter . . . 78
3.3.3 Low-pass filter . . . 80
3.3.4 Processing chain . . . 84
3.4 External validation . . . 85
3.4.1 Validation setup . . . 85
3.4.2 Tests . . . 85
3.5 Experimental validation . . . 90
3.5.1 Experimental methods . . . 90
3.5.2 Results . . . 94
3.6 Discussion . . . 97
3.7 Contribution . . . 99

4 Real-time phase estimation for closed-loop stimulation **101**
4.1 Introduction . . . 101
4.1.1 The phase-magnitude duality . . . 102
4.1.2 Real-time phase estimation . . . 104
4.1.3 Aims . . . 106
4.2 Technical methods . . . 107
4.2.1 The state estimator method . . . 107
4.2.2 Efficient implementation . . . 112
4.2.3 Evaluation . . . 116
4.3 Development . . . 119
4.3.1 User interface . . . 119
4.3.2 FPGA interaction . . . 124
4.3.3 Parameter selection . . . 125
4.3.4 Stimulation control . . . 126
4.3.5 Magnitude estimate . . . 127
4.4 Validation . . . 128
4.4.1 Experiment overview . . . 128
4.4.2 Experimental methods . . . 129
4.4.3 Results . . . 133
4.5 Discussion . . . 136
4.6 Future development . . . 137
4.6.1 Control policy refinements . . . 137
4.6.2 Dynamic gain control . . . 137
4.6.3 Multi-site studies . . . 138
4.6.4 Complex target states . . . 139
4.7 Contribution . . . 139

5 Discussion **141**

Appendix

Signal power estimation **145**
A1 The uncertainty principle . . . 145
A2 Signal transforms . . . 147
A3 Bandpass filters . . . 148

References **148**

List of figures

1.1 Extracellular potential field of a spiking neuron 19
1.2 Spike-timing-dependent plasticity measured in cultured hippocampal neurons . . 21
1.3 Classification of neurostimulation topologies 23
1.4 The origin and scale of neural activity 25
1.5 Strength-duration curve of neuronal excitability 35
1.6 Outline of a general-purpose neural control system 39
1.7 The IEC 60601-1-10 physiologic control framework 40
1.8 Comparison of electrical and optical stimulation artefacts 42
1.9 Spatial and temporal resolution of the techniques available to record and manipulate brain activity 47

2.1 Subclavicular deep brain stimulation system 51
2.2 Paired-pulse stimulation setup 52
2.3 DBS artefact visible in an electrocardiogram 54
2.4 Measurement artefact from transcranial alternating current stimulation 55
2.5 Paired-pulse stimulation system overview 58
2.6 Phase-locked loop operation theory 59
2.7 Phase-locked loop implementation for DBS frequencies 61
2.8 Phase partitioning example 62
2.9 Locking dynamics of the PLL system on a 140 Hz pulse train 63

3.1 The Picostim-DyNeuMo device 68
3.2 Capabilities of the Picostim-DyNeuMo platform 69
3.3 The Picostim-DyNeuMo clinician programming interface 70
3.4 Chronic validation of LFP sensing in the DyNeuMo 72
3.5 The aDBS algorithm architecture 77
3.6 Digital filter architectures 79
3.7 Steady state frequency response of the offset filter 80
3.8 Steady state frequency response of the exponential moving average filter 82
3.9 Performance characteristics of the exponential moving average filter 83
3.10 Error types in the DyNeuMo Parkinsonian beta-oscillation detector 86
3.11 Receiver operating characteristic of the DyNeuMo aDBS algorithm for seizure detection 88
3.12 Response of the DyNeuMo in aDBS algorithm to a seizure event 89
3.13 Design and realisation of the externalised IPG adapter parts 91
3.14 Externalised closed-loop research setup 93
3.15 Alpha-band training data 94
3.16 Results of non-human primate experiment 95
3.17 Correlation of alpha-detector activation and eye-tracking 96

4.1 Analytical continuation of oscillating signals . . . 103
4.2 Conventions of phase angle measurement . . . 105
4.3 Visualisation of the phase estimation algorithm . . . 109
4.4 Quadrature demodulation scheme . . . 110
4.5 Example phase tracking data . . . 111
4.6 Operation of a CORDIC algorithm, visualised . . . 114
4.7 Steady state frequency response of the state estimator model . . . 117
4.8 Steady state frequency response of the digital high-pass filter . . . 118
4.9 Default view of the Intan recording interface . . . 119
4.10 Phase tracking visualisation in the customised Intan recording interface . . . 120
4.11 Stimulation output visualisation in the customised Intan recording interface . . . 121
4.12 Closed-loop configuration tab in the custom Intan interface . . . 122
4.13 Class structure of the customised Intan graphical user interface . . . 123
4.14 Structure of the customised Intan FPGA firmware . . . 124
4.15 Conditional stimulation capabilities of the custom Intan system . . . 127
4.16 Post-hoc phase histograms of stimulation events during closed-loop optogenetic manipulation . . . 134
4.17 Stimulation triggered averages of extracellular potentials in response to closed-loop optogenetic manipulation . . . 135

A1 The time-frequency uncertainty principle . . . 146
A2 Trade-offs of time-frequency analysis methods . . . 147
A3 Frequency and impulse response of two narrow band-pass filters . . . 149

List of tables

2.1 Performance characteristics of the prototyped PLL designs 62

3.1 Example filter chain configurations for three neural biomarkers 84
3.2 Computational cost of various filter chains 84

List of abbreviations

AC Alternating current
API Application programming interface

BCI Brain-computer interface
BP Band-pass
BUA Background-unit activity

CORDIC Coordinate rotation digital computer
CPI Clinician programmer interface
CT Computed tomography

DBS Deep brain stimulation
DC Direct current
DF Direct-form
DTI Diffusion tensor imaging
DyNeuMo Dynamic neuromodulator

ECoG Electrocorticography
EEG Electroencephalography
ESE Empirical state-estimator

FIR Finite impulse response
fMRI Functional magnetic resonance imaging
FPGA Field-programmable gate array

GUI Graphical user interface

HP High-pass

IIR Infinite impulse response
IPG Implantable pulse generator
IR Infrared
ISI Inter-stimulus intervals

JSON JavaScript object notation

LCD Liquid crystal display
LED Light emitting diode
LFP Local field potential
LP Low-pass
LTS Long term support
LUT Look-up table

MCU Microcontroller unit
MEG Magnetocorticography
MICS Medical implant communication system
MRI Magnetic resonance imaging
MUA Multi-unit activity

NA Numerical aperture
NHP Non-human primate
NMR Nuclear magnetic resonance

PBS Phosphate-buffered saline
PEEK Polyether ether ketone
PLL Phase-locked-loop
PV Parvalbumin

RBP4 Retinol binding protein 4
RMS Root mean square
ROC Receiver operating characteristic

SD Standard deviation
SNR Signal-to-noise ratio
SOS Second-order sections
STDP Spike-timing-dependent plasticity
STN Subthalamic nucleus
SUA Single-unit activity

tACS Transcranial alternating current stimulation
tDCS Transcranial direct current stimulation
tES Transcranial electrical stimulation
tFUS Transcranial focused ultrasound stimulation
TMS Transcranial magnetic stimulation
tRNS Transcranial random noise stimulation

UI User interface
USB Universal serial bus
UV Ultraviolet

VCO Voltage-controlled oscillator

Chapter 1

Introduction

Overcoming disorders of the brain is one of the greatest societal challenges of our times [1]. With recent advances in integrated manufacturing and machine learning, technological solutions are seemingly within our reach, yet clinical adoption remains elusive: impressive technologies regularly appear in literature only to never be heard of again [2]. Research into modern therapeutic technologies is fragmented, and held back by the lack of comprehensive platform technologies available to the community, severely limiting the adoption of advanced brain-computer interfaces (BCI). Over a decade passed since the landmark first-in-human implantation of the Utah array for motor decoding in 2012 [3], all the while dozens of independent attempts have been made at designing broadly identical and incompatible neural signal processing systems [2]. There also appears to be a gap between hardware and software, with machine learning models chasing 'state-of-the-art' performance at the expense of ever-increasing complexity [4]. Designers should be conscious of the computational constraints of real-world medical devices in order to provide real-term clinical benefits [5]. Research into medical algorithms should also focus on providing transparent quality assurances to ensure patient safety [6]. Overall, the neurotherapy research space is held back significantly by the lack of pipelines from research to clinical implementation of technological methods.

In this work, we present the development and validation of three clinically-minded research tools for the refinement of deep brain stimulation therapy, at three levels of complexity. Our first design aims to expand the capabilities of existing commercial implantables to explore new research questions and inform the next generation of systems. Subsequent work details the deployment and validation of a long-studied neuromodulation scheme in a contemporary deep brain stimulation device, with the final chapter presenting a research platform to aid in the preclinical validation of an advanced, yet computationally optimised closed-loop control algorithm.

This introductory chapter is intended to bring readers of various backgrounds up to speed on the key pieces of knowledge for discussing neuromodulation technologies. As an overview, the chapter first introduces the origins of neural signals and neural plasticity, discusses common recording methods, then presents a review of contemporaneous stimulation technology as a means of interacting with the brain. The chapter concludes by introducing a unified model of physiological control, tying together measurement and response in preparation for exploring closed-loop technologies. Experienced readers might consider skipping ahead to a brief overview in Section 1.6.

The core chapters of this thesis present three experimental toolsets aimed at advancing neuromodulation therapy research. Each chapter shares a common structure, providing first a high-level introduction to the technology discussed, followed by in-depth technological explanations and development details, finally concluding with the presentation of experimental validation work. The structure is meant to aid the readers in identifying aspects of the work relevant to their interests – experimentalists should find it easy to get an overview of the capabilities of the introduced methods, then focus on experimental use, while those wishing to build on the engineering work are provided with the details to get started.

1.1 The origin of neural signals

1.1.1 Single unit activity

What makes neural tissue unique is the ability of neurons to respond to, elicit, and propagate electrical activity. The electrical activity of single cells is referred to as the action potential and is often referred to as 'spiking' or 'firing' as it is a short (millisecond-scale) impulse-like signal. Artificial neural networks commonly reported in contemporary machine learning research do not capture the spiking nature of real-world neuronal activity but do highlight the power of connected networks [7]. What gives neurons the ability to create action potentials though?

The key to neuronal function is a structured manipulation of ionic charges. As a crude functional division, neurons comprise three major parts: the dendrites where input signals are received, the soma where inputs are aggregated, and the axon where an output signal is propagated. At rest, the cytosol of neurons is abundant in potassium (positive), and low in sodium ions (positive) compared to the extracellular fluid. This imbalance of ionic concentrations creates a potential difference across the cell membrane, called the resting potential (nominally -60 to $-70\,\mathrm{mV}$). The concentration gradient is maintained against leakage by specialised 'ion pump' transmembrane protein structures, exchanging intracellular sodium ions for extracellular potassium in an energy-intensive active process, estimated to account for up to 40 % of the power budget of neural tissue [8]. Input signals arriving at the dendrite have the potential to open up small numbers of ion channels across the cell membrane, allowing an influx of sodium ions, and shifting the transmembrane potential. Given insufficient input, active pumping will soon restore the resting potential. If significant enough input is received within a short period of time and the membrane potential exceeds a certain threshold, voltage-gated ion channels are activated, allowing a greater influx of sodium ions, increasing the membrane potential further. The conformation (i.e. opening and closing) of voltage-gated channels is conditioned on the transmembrane potential. As the potential increases further, sodium channels close again (and become inactivated until the resting potential is restored). Simultaneously, potassium channels open, providing an opposing ionic current and decreasing the membrane potential once again, closing only once further below the resting potential (hyperpolarisation). Finally, active pumping restores the resting potential. While hyperpolarised, the cell is difficult to excite and will not produce further action potentials, creating a post-firing refractory period. Changes in the

membrane potential spread out across the cell membrane through the activation of neighbouring ion channels, forming a propagating potential wave starting at the soma and moving towards the axon terminals, where the action potential will serve as an input to other connected cells.

An often overlooked element in models of neural networks is the presence of glial cells. Various neuroglial cell types are known in the central nervous system. One particular type, oligodendrocytes, plays an important role in facilitating neuronal signalling by forming an insulating layer around axons, the so-called myelin sheath, whose presence speeds up the propagation of action potentials. This simple mechanistic view of oligodendrocyte function has led to them usually being ignored in network modelling, commonly 'lumped into' parameters of axonal conductivity. Their significance in healthy neuronal function should not be forgotten. Myelination is essential to healthy neural function with its disruptions leading to serious disease conditions [9]. Another glial cell type, astrocytes, have recently been proposed to take part in neural signalling themselves, besides their important structural roles [10–12].

The activity of individual neurons is commonly thought of as the fundamental unit or 'building block' of the functioning nervous system. Beyond the oft-ignored glial cells though, one should consider how solitary neurons would be relatively useless: the power of neurons lies in their ability to form connections and organise into complex functional networks. This brings us to synaptic signalling.

1.1.2 Synaptic signalling

Synapses are electrochemical junctions that facilitate the transmission of signals between neurons [13], belonging to two major classes. Gap junctures, also called 'electrical synapses' are essentially channels formed between neighbouring cells, that allow a direct exchange of ionic materials and electrical activity [14]. The other class are chemical synapses, formed between the axons of pre-synaptic ('sending') and the dendrites of post-synaptic ('receiving') neurons. In a chemical synapse, a pre-synaptic axon terminal (bouton), upon receiving an electrical impulse, releases vesicles (essentially small packets) of so-called neurotransmitter substances. A corresponding spine on the post-synaptic neuron (specialised synaptic structures on the dendrite), expressing various chemical receptors, can sense the neurotransmitter release and initiate a small electrical discharge via opening connected ion channels. The electrical activity can have either polarising (inhibitory) or depolarising (excitatory) effect on the membrane, depending on whether the opened ion channel allows passage for positive (sodium) or negative (chloride) ions. Following a synaptic transmission, the pre-synaptic neuron reabsorbs the released neurotrans-

mitters. The post-synaptic neuron integrates the received impulses over time – should sufficient activity reach the dendrites in a small time window, an action potential is initiated, leading to neurotransmitter release in the downstream axon terminals, and the cycle repeats [13].

While a crucial part of neural function, known measurement techniques are limited to either too coarse or too fine a resolution to be useful for interfacing with the nervous system at this level – a handful of single synapses [15] using implanted carbon fibre microelectrode amperometry, or centimetre-scale averages using nuclear magnetic resonance spectroscopy [16, 17].

1.1.3 Multi-unit activity

Multi-unit activity (MUA) refers to the electrical signal recorded simultaneously from a small population of neurons. As action potentials comprise ionic currents propagating from the cell body along the axon, small electric fields are generated associated with these ionic currents, that can be sensed through electrodes inserted into the extracellular space. As shown in Fig. 1.1, the extracellular action potential of a neuron can vary significantly in shape and magnitude based on a combination of cell morphology and relative location from the recording electrode. Due to the high density of neurons in tissue, extracellular potentials will tend to show a mixture of single unit activity.

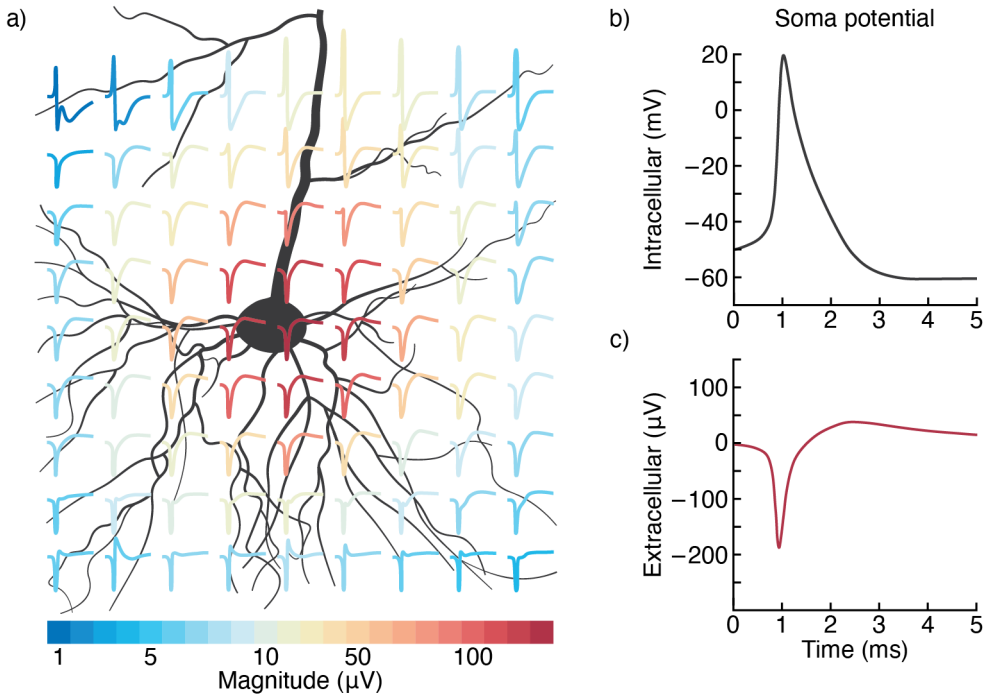


Figure 1.1. Extracellular potential field of a spiking neuron. **a)** The local shape, and peak magnitude of extracellular action potential measurements are shown overlaid on a tracing of a rat cortical neuron. **b)** The transmembrane action potential measured at the soma. **c)** Extracellular measurement of the action potential over the soma. Note how extracellular potential, due to a lack of ionic flow at rest, is typically near 0 V, and the potential excursion appears inverted due to the direction of ionic flows. Adapted from a simulation result by Linden et al. [18].

In the context of BCI decoding, multi-unit activity benefits from greater stability over time than SUA [19], as well as from redundant population-codes used throughout various brain regions [20]. MUA-based BCI algorithms have previously been demonstrated, including but not limited to movement [21] and auditory decoding [22].

1.1.4 Background-unit activity

The background-unit activity (BUA) is signal of medium-scale local activity. BUA can be obtained via measurements largely identical to those of MUA, with some post-processing. First, the signal is filtered to remove low-frequency activity with high spectral power (wide local activity), while also removing MUA (narrow local activity) via thresholding. BUA remains relatively underutilised compared to other signal types discussed in this chapter, commonly considered to be a part of background neural noise [23], BUA has in fact been shown to correlate with more 'mainstream' decoding targets [24–26]. As the measurement of this signal involves deliberately disregarding unit activity, it could provide a robust measure of medium-scale population activity benefiting of long term stability.

1.1.5 Local field potential

While the synaptic and extracellular potentials of individual cells are typically small, synchronous activity over a broader local population adds up to a macro-scale signal [2, 27, 28], the Local Field Potential (LFP) [29]. Trying to infer the exact origin of LFPs from measurements, however, can be both non-trivial and counter-intuitive: despite their name, LFPs are not necessarily local in origin. A detailed review of the caveats of interpreting LFPs is beyond the scope of this introductory chapter, interested readers are referred to a dedicated review [30].

1.2 Neural plasticity

A key feature of neural circuits is plasticity, or their ability to change. Several high-level methods for network learning have been proposed over the years, particularly in the context of artificial neural networks, with the most prominent of these being the method of backpropagation [31,32]. While it remains an open research question whether proposed learning methods take place in neural tissue [31, 32], we can certainly say a few things about the nature of synaptic plasticity at a lower level.

1.2.1 Hebbian plasticity

The most well-known principle guiding synaptic plasticity is Hebb's rule. The learning rule was originally phrased by Donald Hebb in his seminal work on neuropsychology in 1949 [33] as follows. "When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" [33]. The key insight of Hebbian plasticity is temporal causality: synaptic connections are strengthened when the activation of a pre-synaptic neuron precedes that of a post-synaptic neuron, while non-coherent and anti-causal patterns will tend to decrease synaptic strength.

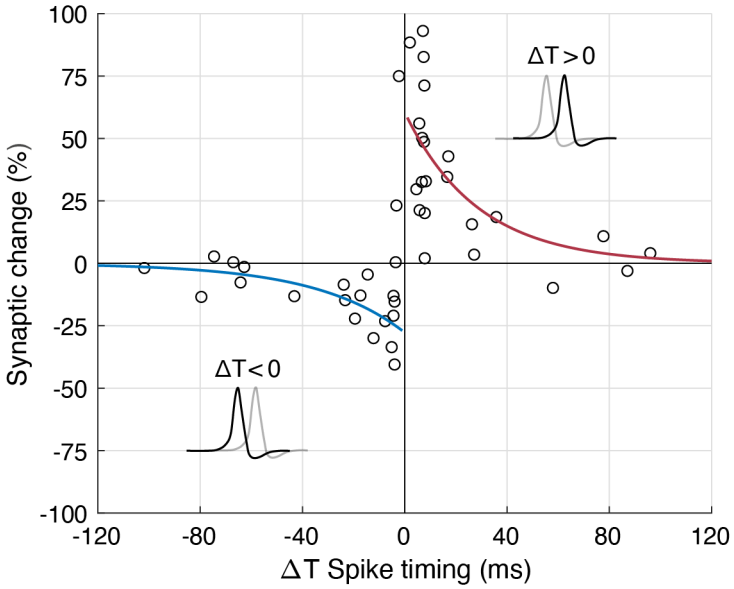


Figure 1.2. Spike-timing-dependent plasticity measured in cultured hippocampal neurons. Synaptic change reported as the percentage change in amplitude of the excitatory post-synaptic potential (subthreshold), recorded 30 minues after a one minute long training period where correlated firing was induced at a rate of 1 Hz. First order exponential model fits are shown for the potentiating and depressing phases. Data from [34].

Hebbian learning is commonly phrased in the mnemonic form "cells that fire together wire together" after Carla Shatz (misattributed to Hebb) [35, 36]. In the context of neural network theory, this learning rule is commonly referred to as spike-timing-dependent plasticity (STDP). Fig. 1.2 presents data from a famous experiment demonstrating STDP-like behaviour in hippocampal neurons *in vitro* [34].

Hebbian plasticity would, however, imply that single neuron representations cannot be reliable sources of information for BCI decoding and that throughout the use of observed neural circuits, plasticity would cause slow changes in the neural representation of function [19]. Notably, SUA-based decoders have indeed been shown to suffer from representational drift, even over the course of mere hours [37]. The question naturally arises: how does the brain maintain stability of function then? Recent work suggests mechanisms for synaptic homeostasis exist, at least within small populations of neurons, limiting changes to firing rate and other dynamic properties [38, 39], implying functional representations are preserved on the population rather than single cell level [19].

1.2.2 Induced plasticity

Externally driving targeted changes in network activity by exploiting innate plasticity offers the potential to create lasting therapeutic effects. While Hebb's rule is, as discussed, incomplete in its explanation of neural dynamics, it still serves as a useful first-order framework in exploring induced plasticity.

As the core mechanism works on synchronised activity between pre- and post-synaptic neurons, targeted modification of synaptic strength requires methods to drive this synchrony. One obvious approach is based on neurostimulation. Literature categorises stimulation setups into three main topologies, as illustrated in Fig. 1.3. The most straightforward approach is repeated stimulation of a pre-synaptic target (Fig. 1.3a), which relies on driving the network at a high-enough intensity that post-synaptic spiking is elicited. Over time, successful paired spiking will lead to long term potentiation of the connection. The technique has been primarily studied in human motor rehabilitation after nervous system injuries [40]. This approach, while often viable, is not particularly efficient as the induced pre-synaptic activity will frequently be anti-correlated with post-synaptic spiking, working against the desired effect while connectivity is still weak. As a further draw-back, repetitive pre-synaptic stimulation does not lend itself well to depression of synaptic strength.

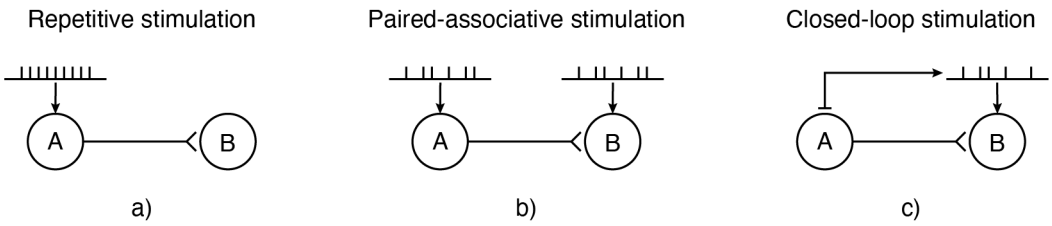


Figure 1.3. Classification of neurostimulation topologies. **a)** Repetitive stimulation to affect local, and downstream neural activity can lead to long term potentiation at sufficient driving strengths. **b)** Multi-site stimulation can drive the activity of both pre- and post-synaptic nodes, creating a forced state of synchrony (or asynchrony). **c)** Closed-loop stimulation uses neural activity as a feedback signal to deliver stimulation at the optimal time for driving the network for a desired effect. Based on a review by Jackson and Zimmermann [40].

The next major technique, multi-site stimulation (Fig. 1.3b) solves this problem by forcing the activity of each target to follow a desired pattern, allowing for both potentiation and depression of the connection. In human evidence for the feasibility of this technique has been reported in Parkinson's disease patients, where paired co-stimulation of motor cortex and a connected deep brain structure was shown to increase the excitability of the targeted pathway, with a lasting effect on the order of an hour [41].

Finally, closed-loop stimulation (Fig. 1.3c), a more technologically challenging approach, solves the issue by monitoring neural activity to optimally align the timing of stimulation in the target [40]. A closed-loop stimulation pilot study in Parkinsonian patients, based on modulating the activity of a deep brain structure based on its activity, demonstrated the potential of this method to achieve long-term changes in neural dynamics in deep brain areas. It is reported that the frequency of detector activations gradually decreased throughout the experimental session, while therapeutic effect remained unchanged [42]. While further work is needed to establish whether and on what time scales these changes these changes may exert beneficial effects, this study represents a key early demonstration of the potential of closed-loop therapies.

Neurostimulation is not always necessary to achieve plastic changes in the brain though. BCIs have been successfully deployed in non-human primates that enable control of a cursor on a computer screen by decoding activity from the motor cortex [43]. Notably, these experiments have demonstrated adaptation of neural activity while learning to use the interface – aspects of activity analysed by the decoder became more prominent over the course of several days of training, correlating with improved accuracy. Similar effects on plasticity were observed in both LFP [44] and spiking activity decoding systems [43]. In these experiments, the presence of a clear goal (moving the cursor in a desired way) creates a mental representation of success, and

an associated measure of performance via visual feedback. This allows the brain to harness existing pathways for refining motor activity, and facilitates learning. Similar results have been obtained in a more limited number of human studies [45].

The success of sensory feedback approaches in driving plasticity is likely reserved to specific neural targets based on the presence of inherent adaptation pathways, and is not likely generalise to all brain regions. Feedback pathways may also be of limited help when aiming to reliably drive plasticity in disease conditions. Overall, neurostimulation will likely remain the most viable, and most generally available technique when exploring the therapeutic capabilities of neural plasticity in the near future.

1.3 The recording of neural signals

Choosing the right scale of neural signalling to interact with the brain at is a key challenge of contemporary BCI research due to limitations of recording hardware in both power consumption and data throughput [46]. Designers must attempt to maximise information content, while aiming to minimise redundancy to meet these limitations. Fig. 1.4 presents a preview of commonly used organisational levels and their recording techniques, that we explore in more detail in the following section.

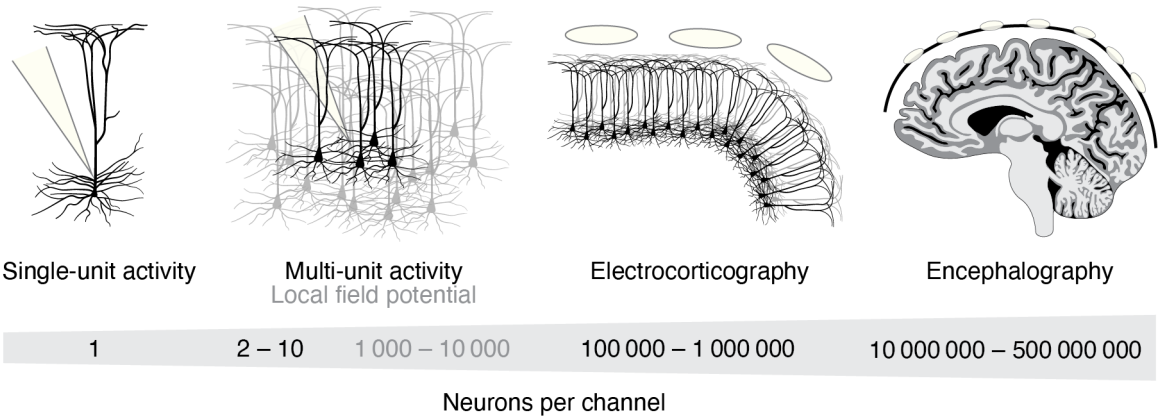


Figure 1.4. The origin and scale of neural activity. The figure shows common forms of electrical measurements of neural activity, ordered by the neural resolution. The potential from single unit activity (SUA) of individual neurons is spatially summed into increasingly coarse representations of population activity, with decreasing localisation. Local field potentials (LFP) are considered to be relatively small population activity signals. Electrocorticograms (ECoG) are no longer able to distinguish the activity of cortical layers and are generally taken with sparser spacial resolution, while magneto- and electrocorticography (MEG and EEG) are performed non-invasively, providing typically the coarsest spacial sampling, though the whole-cortex can receive coverage. EEG signals undergo significant 'degradation' passing through the cranium. Note how generally coarser resolution techniques tend be less invasive. Based on [47,48]. Neuron traced after Ramón y Cajal [49].

1.3.1 Single-unit activity

The highest fidelity method for measuring the electrical activity of single neurons is the patch-clamp technique, where an electrolyte-filled glass micropipette is brought into mechanical contact with a neuron. The pipette clamps onto the cell membrane using negative pressure, and forms a tight, high-impedance seal, providing excellent signal-to-noise ratio [50,51]. Successful clamping requires visual feedback and careful micromanipulation. A very similar method, juxtacellular recording, is sometimes used where only a loose seal is formed with the membrane, slightly decreasing the complexity of the preparation [50, 52]. While patch-clamp and juxtacellular recordings have been extremely powerful tools for studying a wide range of neural dynamics, including receptor and ion channel function [53], as well as microcircuit interactions [54], state-of-the-art equipment *in vitro* still only allows simultaneously recording up to a dozen cells [51], making these technique ill-suited for brain-computer interface (BCI) applications.

1.3.2 Multi-unit activity

Multi-unit recordings are largely performed via microelectrode arrays, of which several major classes exist. These include twisted microwire bundles (e.g. tetrodes, hextrodes), cortical needle arrays, and micropatterned silicon probes. Silicon probes offer the greatest flexibility in electrode site placement and could help improve SUA acquisition performance [55], however, these electrodes can suffer from delamination and other long-term issues with stability [56]. Microwires offer the greatest depth, but the number of electrodes implantable can be more limited using this method. Cortical needle arrays are while mechanically stable, prone to scar tissue buildup, limiting long-term recording quality [57]. An overview of the technologies can be found in an external review [56].

In practice, SUA is most commonly obtained by post-processing of MUA due to the practical limitations of single-cell recording techniques [51, 54]. This post-processing step is referred to as spike-sorting, that is, the categorisation of spiking activity in the MUA into distinct groups. In practice, this process consists of 1. filtering the broad-band neural signal, 2. thresholding to isolate events from background noise, then 3. identifying events belonging to a common unit (clustering) [58]. This last step can be a highly technical process with great computational cost, though generally speaking it is based on the shape of the extracellular action potentials (Fig. 1.1). A nigh limitless number of clustering algorithms exist with various trade-offs, theoretical research suggests a 'perfect' algorithm might be impossible to construct [59].

Recent work has demonstrated that common MUA recording arrays are arranged in suboptimal geometries, and fail to maximise the recorded information content for SUA clustering [55]. Interestingly, it has been suggested that using modern recording systems, the size (and therefore impedance) of common microelectrode contacts does not materially affect the quality of unit acquisition [60].

For successful clustering, sampling rates of 15–30 kHz are necessary to provide sufficient detail of the shape of individual potential waveforms for identification. Where unsorted MUA can be used directly, sampling rates on the order of 5 kHz are sufficient, significantly improving the information density of data channels in population coding interfaces.

Over the past decade, we have seen the first demonstrations of real-time spike sorting in BCI systems, though for now these have invariably been unidirectional technologies with no feedback path [46, 61, 62]. In these real-time systems, the complexity of the spike sorting problem is side-stepped by performing clustering offline and basing subsequent detections on template-matching. Template matching approaches run into difficulties in long-term use: the observed spike shapes slowly change over time due to tissue relaxation and inflammation around the electrode. More computationally complex implementations have proposed methods for slowly adapting the templates over time to compensate for drift [62].

Interestingly, we seem to be able to observe significantly fewer units in the local population than expected based on tissue density [48]. Methods such as current source density calculation allow estimating the distance of recorded units from an electrode array, suggesting that recorded cells can be routinely found up to 100 µm away from an electrode. The volume enclosed by a 100 µm radius sphere should contain on the order of 500–1000 neurons in cortical areas [48, 63], while we typically observe only 10–15 single units over such scales in neural recordings [48]. This observation has been attributed to a combination of most units being inactive under the specific experimental conditions, limitations of spike sorting methods, and the high spatial orientation-dependence of extracellular potential signals [48].

1.3.3 Local field potential

Local field potential activity is commonly measured by implanted macroelectrodes directly as an electrical signal, both in research and clinical applications [42]. In practice, researchers tend to use differential measurements between two electrodes enclosing the neuronal population of interest to mitigate volume conductivity effects and exclude far-field activity from the LFP

[64]. While LFP already appears on microelectrodes, long-term recordings can benefit from the reduction in electrode impedance, as glial scar build-up progresses, especially in systems where spiking activity is not of interest.

LFPs are commonly recorded at 500–1000 Hz sampling rates to capture broad-band activity, though rates significantly lower can be sufficient if targeting certain low-frequency LFP sub-bands.

1.3.4 Stereotactic electroencephalography

Stereotactic EEG (sEEG), not to be confused with regular EEG (discussed later), is a 'minimally invasive' recording option for LFP measurements, where a small number of sparse electrodes are transcranially implanted, aimed at covering a number of distinct brain regions, rather than providing detailed sampling of a few. Stereotactic surgery allows for highly precise implantation based on anatomical maps, and can be greatly refined based on pre-surgical anatomical scanning of individual patients [65]. Both cortical and deep brain structures can be targeted in sEEG, with signal properties essentially identical to other electrical measures of LFPs.

sEEG is in active clinical use for localising focal sites in epilepsy, i.e. brain regions where seizures emanate from [66]. While the typical sparsity of recording sites limits BCI decoding using sEEG alone, the technique has been proposed to have great utility when used in conjunction with high-density cortical recordings, by offering a refinement of decoding based on deep-brain signals [67].

1.3.5 Electrocorticography

Electrocorticography (ECoG) is an intracranial technique for measuring neural activity. It is distinct from LFP recordings primarily based on the recording setup: typically grid- or strip-shaped flexible electrode arrays laid out over the cortical surface [29]. This recording arrangement precludes differentiating the activity of cortical layers, and provides coarser spatial sampling with typically mm-scale resolutions, though sub-mm arrays have been successfully demonstrated *in vivo* [68]. ECoG has seen clinical use both in pre-surgical planning and intraoperatively in epilepsy and tumour surgeries [69].

Due to its ability to provide large-scale yet detailed coverage of specific cortical areas, ECoG signals have formed the basis of some of the highest quality functional brain decoding systems, including recent attempts at decoding articulatory information in speech neuroprosthetic prototypes [70, 71].

1.3.6 Electroencephalography

Electroencephalography (EEG) is among the most well-known neural recording modalities, widely used in both clinical and research applications. EEG primarily reveals high-amplitude, large population synchrony (Fig. 1.4). Due to its non-invasive nature, EEG has possibly seen a greater number of attempts at BCI applications than any other neural signal type [72]. EEG also benefits from the availability of relatively low-cost recording setups, while allowing recorded subjects (relative) freedom of movement. That said, EEG suffers heavily from low spatial resolution, as well as from significant reductions in signal power due to the bone and tissue layers separating the electrodes from the cortical surface. A further issue can be the presence of significant motion artefacts [72].

EEG is normally measured via a set of electrodes in mechanical contact with the scalp, typically attached to a head cap. In so-called wet electrode systems, typically a conductive gel can be applied, though skin abrasion can be necessary to reduce impedance to suitable levels for reliable measurement [73, 74]. Dry systems, while offering significantly reduced preparation complexity before measurements (as well as discomfort), have classically suffered from poor reliability and SNR, though significant progress has been made on this front [75]. Another method for improving signal quality in dry EEG systems is available through the use of on-electrode active signal processing circuits. The active components constrain the smallest size at which active dry electrodes can be realised, limiting the resolution to about 20 recording channels at present [76, 77], compared to passive systems where 32 channels have long been a minimum standard, with systems up to 265 channels available in research use [75].

1.3.7 Non-electrical methods

As changing electric fields are associated with magnetic fields, magnetoencephalography (MEG) has been developed as a higher resolution alternative of EEG [78, 79]. Recent technological advances in optically pumped magnetometers [80] led to the development of high-channel-count portable MEG headsets [81]. While of significant research and clinical interest, MEG requires a magnetically shielded room to eliminate noise [81], making it unsuitable for practical BCI work at present.

Neural firing is an energy-intensive process, thus elevated levels of local activity require significant oxygen and energy (sugar) uptake from the cerebral vasculature [82]. Two recording techniques commonly seen in neuroscience literature operate on the basis of measuring changes in blood-oxygenation in the brain, functional magnetic resonance imaging (fMRI) and functional

near-infrared spectroscopy (fNIRS). While fMRI has utility in non-invasive recording of deep brain areas, it is limited in both spacial and temporal resolution (millimetre – second scale) [83, 84], not to mention lacking a pathway for out-of-clinic use. Studies have also highlighted that fMRI results can be particularly susceptible to noise unless statistical analysis is carefully planned [85]. In the case of fNIRS, recording is limited to the brain surface, measurements provide an estimate of average cortical activity on the centimetre scale, with temporal resolutions up to 10 Hz [86]. As the signal is not sensitive to electrical stimulation artefacts, fNIRS has some potential utility in closed-loop systems for the simultaneous monitoring of local cortical activity in stimulation studies [87].

1.4 Neurostimulation

In previous sections we covered some of the most established means of measuring the nervous system activity at various scales. On our journey to designing closed-loop neurotherapies, we now turn our focus to exploring some of the most common techniques available to interact with the brain. Neurostimulation encompasses a wide range of methods aimed at modulating neural activity. Through targeted application of electromagnetic fields, electrical currents, or light in conjunction with genetic manipulation, neurostimulation exerts control over neuronal behaviour. Depending on the method and specific parameters employed (frequency, waveform), intense stimulation tends to induce action potentials, while low intensity stimulation exerts an effect by altering the excitability threshold of neurons. Although most stimulation methods excite neurons on the level of cells, specific combinations of stimulation patterns and network connectivity (e.g. proportions of excitatory and inhibitory synapses) can result in either suppression or activation on the population scale. Optogenetic techniques are unique in their ability to directly silence particular types of neurons by introducing new ion channels to control the membrane potential.

As already apparent, the design of stimulation therapies is a complex problem. To steer network dynamics in a desired way, stimulation methodology and parameters must be carefully selected given the electrophysiology and neuroanatomy of any given target area. As one of the primary considerations in therapy design, in the following sections we categorised methods based on their level of invasiveness.

1.4.1 Non-invasive stimulation

Transcranial electrical stimulation

Transcranial electrical stimulation (tES), or the application of small electrical currents to the scalp to modulate the activity of cortical tissue beneath, is in widespread experimental use in human research in contemporary literature [87]. The method's popularity can be primarily attributed to its safe, non-invasive nature, along with low instrumentational complexity. Several stimulation types are in use, including direct current (tDCS), alternating current (tACS) and random noise (tRNS), though complex multi-frequency modulation patterns have also been reported [88].

Despite the wide popularity, the research community is divided on the merits of tES techniques, due to questions about whether sufficient current can reach the cortex for neural activation. Scepticism is primarily driven by a famous human cadaver study, where implanted

electrodes were used to directly measure the stimulation intensity achieved within the cortical tissue, showing that most of the externally applied current was shunted by the scalp and other interposing tissues [89, 90].

Follow-up studies in non-human primates, however, suggest that low intracortical current levels do not preclude successful neuromodulation. While low-intensity tACS at 0.5 mA failed to elicit correlated neuronal firing, as stimulation amplitude increased proportionally more and more neurons were successfully entrained, reaching 30 % of sorted units at 4.5 mA [91]. This experiment led to the hypothesis that tES produces selective neuronal activation due to the high impact of specific cell morphology and orientation relative to electric field [90, 91].

Another criticism of tES studies has been a frequent lack of controls to discern whether effects of tES arise due to direct electrical activation of the cortex, or indirectly through activating nerve endings in the scalp that couple into the sensory cortex. Once again, primate studies have been conducted to allay these concerns about the validity of tES study protocols, demonstrating that tACS carried out under local anaesthesia still produce consistent tissue activation [92].

Further supporting the validity of tES techniques, large-scale meta analyses particularly in the form of tACS, have shown significant therapeutic effects in a range of neurological conditions [93].

In terms of implementation, tES requires an electrode pair to inject current into the tissue; several choices are available. The earliest type in use is a simple sterile sponge wetted with a saline solution. Solid conductive electrodes are used in conjunction with conductive gels to achieve good contact with the scalp. A common alternative choice of electrodes is conductive silicone, typically shaped in large (10 cm) rectangular or ring shaped patches. So-called high-density silver/silver-chloride electrodes are sometimes preferred for their more versatile placement, such as multiple cathodes surrounding a common anode in a star shape, called the 4x1 arrangement [94]. The small size also lends itself well to interleaving EEG and tACS electrodes on a head-cap, though simultaneous recording and stimulation is not typically feasible due to the magnitude of the tACS artefact.

The choice of electrode materials, their size, and their placement all significantly affect the electric field density inside the cranium. The size of the electrodes guides the current injection limit, larger and lower impedance electrodes can be used to safely inject more current at an equal density per surface area, achieving greater field of potentiation. The placement of electrodes, referred to as a 'montage', plays possibly the greatest role in shaping the current path. Montages can be divided into two fundamentally different groups, cephalic montages where both the

source and sink electrodes are placed on the head, and extracephalic montages where one of the electrodes is located away from the head, typically near the shoulder. Cephalic montages can be further divided into several topologies with varying field effects. For a comprehensive review with illustrations readers are referred to [95]. To illustrate the significance of the impact of montage selection, an extracephalic return path contralateral to the crainal electrode was shown to significantly entrain tremor activity, while ispilateral return, and various cephalic montages failed to demonstrate an effect [96].

The choice of stimulation frequency in tACS also has an impact on, as tissue impedance is non-linear with frequency. Overall, field modelling during the study planning stage is recommended.

Transcranial magnetic stimulation

As a less controversial method, transcranial magnetic stimulation (TMS) is based on the generation of local electric currents in the brain using strong, externally applied magnetic fields. While early systems have suffered from poor field spread, both novel electromagnetic coil designs and magnetic pulse shaping have helped improve focusing stimulation [97, 98]. TMS has shown significant potential in clinical use, though the systems are lacking in portability as a constraint of the power-electronics components required to generate large strong fields [98].

Unlike tES, where blinded stimulation can be difficult without applying topical anaesthesia, TMS systems benefit from the availability of a convincing control condition based on auditory feedback [98].

1.4.2 Intracranial stimulation

As we have seen, non-invasive stimulation techniques tend to be restricted in their spatial resolution, and depth. Deep brain structures, however, offer some of the best therapeutic value as targets in a variety of disease conditions, including movement disorders, epilepsy, and disorders of consciousness. At present, a trade-off exists between the portability and the achievable depth in non-invasive stimulation techniques – transcranial focused ultrasound (tFUS) and TMS being unsuited for even short-term wearable use, let alone chronic therapy, while tES methods are restricted to cortical stimulation. Intracranial stimulation modalities do not suffer from these limitations, and despite the inherent surgical risks, form an essential part of contemporary neuromodulation therapy and research.

Electric stimulation

Electrical stimulation is the exclusive approved form of intracranial neurostimulation at present, though other techniques have long existed in research use [99]. As neural tissue is naturally electrically excitable, electrical stimulation has a straightforward mechanism of action, though the safe chronic interfacing of stimulation electrodes and the brain presents important safety considerations worth discussing.

Most constraints stem from the fact that the tissue-electrode interface should not be thought of as a simple electrical connection. Electron current in the stimulator system has to couple to an ionic flow in the interstitial fluid to complete the circuit. The coupling occurs primarily through electrostatic interaction, making the electrode surface area a dominant factor in the connection's impedance. This presents us with a design trade-off: larger electrodes decrease impedance while reducing our ability to precisely target tissue at high spatial resolution [100]. Why is lowering electrode impedance important though?

To understand the importance of low tissue-electrode impedance, we need to consider that stimulation voltages beyond certain limits can cause irreversible chemical changes in the electrode microenvironment. This is thought to be a contributing factor in creating a chronic state of inflammation around the electrodes that leads to the formation of scar tissue, degrading the interface [101]. The most important electrochemical reaction in this context is the oxidation ($0.8\,\mathrm{V}$) and reduction ($-0.6\,\mathrm{V}$) of water molecules, referred to as the 'water window' [102]. In practice, the impedance of the tissue-electrode interface, on the order of $\sim 1\,\mathrm{k}\Omega$ [103], does not permit delivering the current levels required for tissue activation within the water window. The use of advanced material technologies, such as conducting polymers and hydrogels has the potential to reduce tissue-electrode impedance by greatly expanding the effective surface for the coupling of electronic and ionic currents compared to metal electrodes [101].

The same concern of electrochemical stability leads to further safety considerations: stimulation should not cause an imbalance of local charge [104]. As we need to deliver stimulation somehow, practically this constraint means that stimulation should not cause *lasting* imbalance of charge, thus stimulation pulses should be kept relatively short, rapidly followed up by an equivalent pulse of opposing current. The injecting and extracting (also referred to as cathodic and anodic) phases only need to equal in injected charge, thus phases of different current magnitudes are acceptable so long as pulse widths are appropriately adjusted [102,104]. Stimulation that only repeatedly delivers unidirectional current is referred to as monopolar, while charged

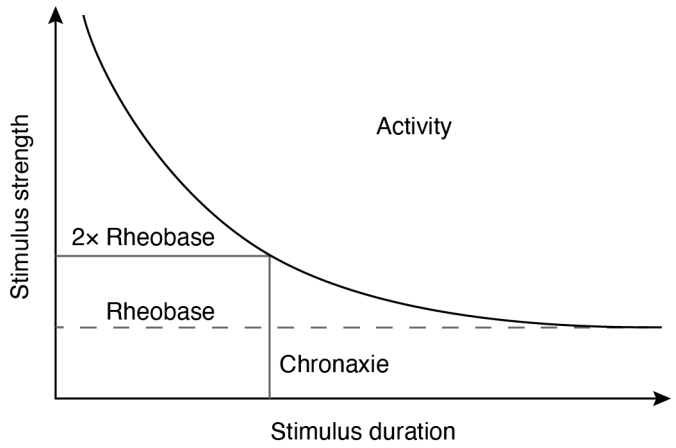


Figure 1.5. Strength-duration curve of neuronal excitability. The curve shows the threshold of the lowest stimulation intensity at a given stimulation length to elicit neural firing. The rheobase is the minimum intensity for arbitrarily long stimulus, establishing an asymptotic baseline, while the chronaxie, the time where the activation threshold is twice the rheobase, is commonly used to establish the time constant of the curve.

balance waveforms are bipolar. The balancing phase of bipolar stimulation can happen through active (driven) or passive (capacitive discharging) means. For further details of common stimulation waveforms, readers are recommended to consult external reviews [102, 104, 105].

All of that said, we should consider how much current is necessary to cause reliable tissue activation. As neuronal activity is primarily event-based with refractory periods present, the relationship between stimulation intensity and activation is not immediately obvious – the key parameter missing to make the connection is time. The shorter a stimulus is applied, the more intense it needs to be to elicit a response from a neuron. Accordingly, neural excitability is characterised by the so-called strength-duration curve, providing the lowest stimulation intensity to evoke an action potential when applied for any particular duration. A generalised characteristic curve is shown in Fig. 1.5. The lowest stimulation intensity required to elicit spiking from a neuron if present in an infinitely long stimulation pulse is referred to as the rheobase. To capture the time constant of the curve, a second point is defined, called the chronaxie, the length at which the intensity is twice the rheobase [106]. Typical deep brain stimulation (DBS) stimulation uses currents on the order of milliamperes, at pulse widths of 50–450 µs [107].

Optical stimulation

While most neuronal cell types do not naturally respond to light, research into genetic modifications has produced a wide range of techniques for inducing photosensitivity, collectively referred to as optogenetics. The basis of optogenetics involves introducing new genetic information into cells that encode the production of light-sensitive ion channels, allowing light stimulus to act as

a novel pathway for polarising and depolarising the cell membrane. The genetic information is most commonly delivered via viral vectors (carriers) in a process called transduction. Through more extensive engineering, the expression of opsins can be made conditional on the presence of certain proteins, enzymes, and metabolites, which allows researchers to target specific neuronal cell types. This selectivity has made optogenetics an invaluable tool for studying neural circuits, by enabling researchers to disentangle the role of different cell types within an area through targeted activation and suppression. For a comprehensive overview of the technology, interested readers are referred to a review in [108].

While it is precisely what makes optical stimulation a powerful technique, the necessity of genetic modification, and the use of viral vectors pose a significant barrier to adoption in human therapeutic use. At the time of writing, no approved human optogenetics technique exists – demonstrations in human tissue have been carried out in tissue samples ex vivo, and stem cell-derived cultures *in vitro*. The first clinical trial into human optogenetics has recently been approved, aiming at the partial restoration of visual function in degenerative retinal conditions [109]. It has been suggested that safety evaluations of optogenetics for use in the human brain could be carried out in epileptic patients, by deploying the viral vectors into neural tissue already set to be surgically resected as part of their therapy [110], though no such research has been reported so far.

Optogenetics benefits from the widespread availability of laser and LED technology. Stimulation is typically delivered to the tissue via implanted optical fibres, sometimes integrated into conventional neural microelectrode arrays referred to as optrodes [111]. Recent technological developments have made it possible to integrate micrometre-scale LEDs onto the surface of silicon electrodes, significantly improving the potential for chronic implantation [48, 112]. Such integrated recording-stimulation probes have paved the way for closed-loop optogenetics in recent years [61, 113, 114], the last complication to overcome remaining computational complexity [61, 114].

A particular safety concern of optogenetic stimulation is excessive tissue heating. Irreversible tissue damage is thought to occur around a temperature difference of 6–8 °C above baseline [115]. A recent methodological study in mice has shown that common finite element models tend to underestimate the heating effect of optical stimulation, with intensities of $320\,\mathrm{mW/mm^2}$ at wavelengths below 550 nm (blue, green) rapidly heating cortical tissue beyond safe temperatures [115], though short pulsed patterns seem to perform significantly more favourably [116].

As a further constraint of the technology, opsins can be limited by their channel dynamics to lower activation rates than what direct electrical stimulation would typically elicit in neural tissue [117]. Overall, while optogenetics is an extremely powerful tool for neuromodulation research, it should not generally be thought of as a model of electrical stimulation.

Thermal stimulation

After the previous discussion of the safety concerns of tissue heating, it might be surprising that thermal stimulation would be used intentionally. Thermal stimulation can essentially be thought of as an opsin-free form of optical stimulation. Instead of light activation of artificially introduced opsins, action potentials are elicited through the thermal excitation of ion channels. Effective mm-scale local tissue activation has recently been demonstrated in a mouse model [118], while resulting in peak temperature differentials under 5 °C [118], at illumination intensities generally considered safe in optogenetic setups [116]. Unlike optogenetic techniques, thermal stimulation could provide an interesting combination of properties for future research. Opsin-free indiscriminate tissue activation could be more closely representative of electrical stimulation, while still benefitting from some of the desirable proper optical stimulation including low artefact levels in electrical recordings.

1.5 Closed-loop systems

1.5.1 Overview

In the previous sections, we explored the major technologies for monitoring and modulating neural activity, their capabilities and limitations. Combining these techniques into complete bidirectional recording-stimulation systems, we need one more component: a control system. Without getting into the mathematical details, let's have a look at the main strategies for process control.

Generally speaking, the purpose of a control system is to steer the state of a process towards a desired target state and to maintain it there by counteracting disturbances. Fig. 1.6 presents this idea in a high level schematic view. In the context of neuromodulation, the system under control is a neural circuit, though the nature of the system state is slightly more complex. While we aim to control disease state overall, it can be quite difficult to discern whether disease-correlated biomarkers of neural activity are causative or merely epiphenomena, byproducts of a more fundamental disease process. For example, pathological neural activity is widely accepted as the causative mechanism of seizures. In this sense, the state to be monitored (epileptiform neural activity) and the process to be controlled (the generation and spread of epileptiform activity) are directly related [119, 120]. In the case of Parkinson's disease, where LFP oscillations have been shown to serve as a biomarker of symptom severity, even with successful patient outcomes of responsive neurostimulation, debate persists whether stimulation simply happens at the right time to reduce symptoms due to correlation, or whether it is the suppression of oscillatory activity itself that results in symptomatic improvement [121]. Recent studies with stimulation locked to the oscillatory activity might provide some clues [25, 122–124].

The key to the successful operation of feedback-type control is timing. Unless the sensor can establish a low-latency readout of the state, and the controller can rapidly respond to deviations from the target, compensatory action might no longer be appropriate by the time it is delivered, causing instability in the system. A separate timing consideration for successful feedback control is the rate at which the feedback process should attempt adjustments to the state such that it is not over-driven, yet actuation (therapy) can be delivered on time. The rate of adjustment ties into selecting appropriate maximum limits on actuation action.

As shown in the overview of Fig. 1.6, there exists a second control pathway: feedforward control. In this pathway, we use external information, and potentially predictive models, to compensate for disturbances that act on the system. To illustrate how external knowledge

can help counteract disturbances, several disorders show recurring periodicities, circadian and infradian (slower) rhythms of symptom severity. One such condition is epilepsy [125], where information about seizure patterns can be used to adapt neurostimulation in advance as a preventative measure [126, 127]. In practice, feedback and feedforward control can complement each other well, and can be used in conjunction.

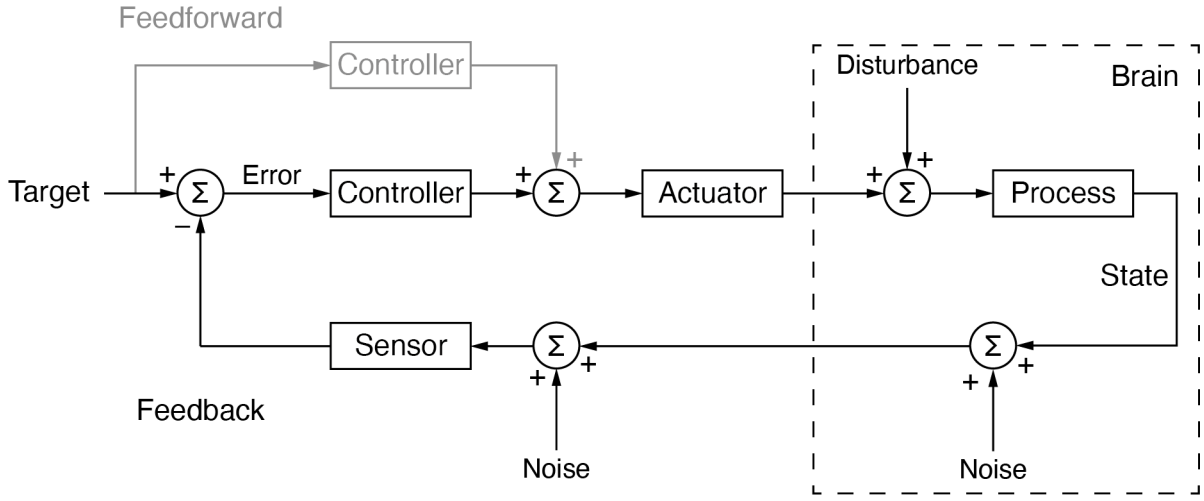


Figure 1.6. Outline of a general-purpose neural control system. Control systems aim to steer a target process to reach a desired system state. In a feedback-control mode, the target state is compared to a measurement of the system state to inform the controller. In a feedforward mode, the controller is informed by information on the disturbance to provide compensatory actuation. Sensing of the true process state is complicated by 'neural noise', activity unrelated to the process, as well as sensing noise.

1.5.2 Safety consideration

The deployment of closed-loop, automated systems for patient use carries with it critical safety considerations. The research community has recently proposed a unified framework for the safe design and operation of general purpose physiological control systems, termed the IEC 60601-1-10 [128]. A compliant system model is presented in Fig. 1.7 with some of the specifics of neural systems included, while kept largely applicable for general purpose physiological control. As expected, the framework is built around a feedback loop–style mental model of control (dashed block), with specific considerations for implantable devices.

The presented model divides system functionality into three distinct sets based on responsibilities. The implantable device fulfils the role of the control system including the sensor and control elements. Interactions with the device are separated into two groups, the patient and the clinician programmer, which can in practice have separate physical realisations as well [129] to prevent inadvertent or unsafe configuration changes. While both have capabilities for state

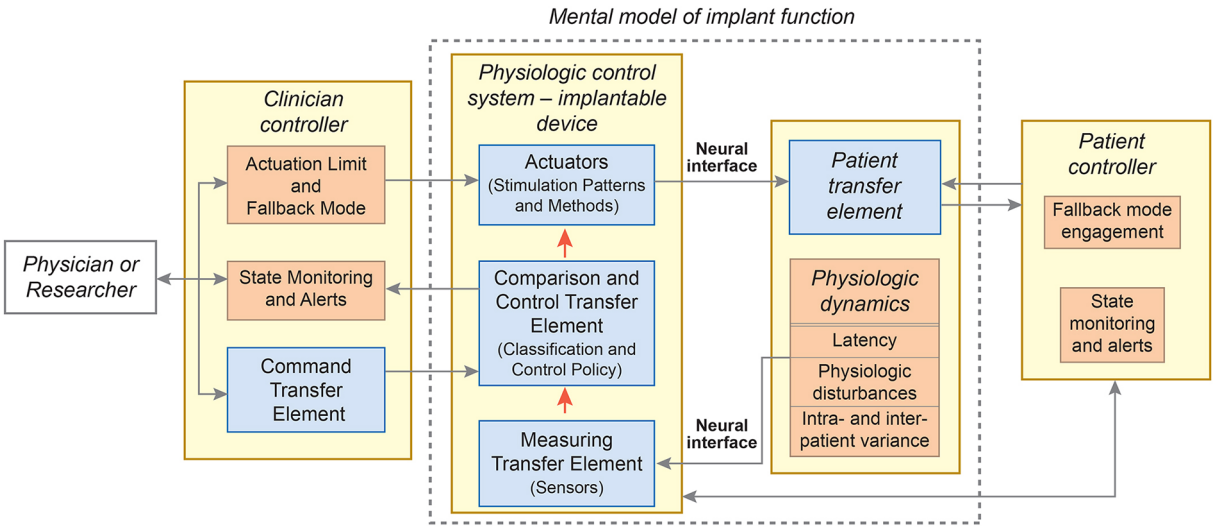


Figure 1.7. The IEC 60601-1-10 physiologic control framework [128]. The framework allows to balance user needs (blue) and risk mitigation strategies (orange) in a systematic way for a broad range of physiological control systems. Reproduced from [129] under CC BY 4.0.

readout, the clinician programmer would have access to a deeper level of diagnostic information and has the exclusive ability to define new operating modes. The patient controller is limited to engaging and disengaging, or selecting between clinician-approved algorithmic modes.

Automated therapy control, especially where a deep modelled understanding of disease processes might not be available, can be a disconcerting idea. To alleviate patient (and regulatory) concerns and provide guarantees on device safety, actuation limits should be put in place, restricting the parameter ranges within which automated adjustments are possible to known-safe levels. In the case of DBS therapy, the parameter space can be intractably broad with stimulation frequency, waveform and intensity all tuneable components of therapy. From a practical standpoint, we can differentiate between therapy setup and therapy tuning, assigning the former to clinicians in establishing a basic operating space, while allowing the control system to carry out the latter. In the case of DBS therapy, there appears to be broad consensus in the research space in limiting automated control to stimulation magnitude, while keeping other parameters constant [42, 130, 131]. An alternative to within-range tuning is for clinicians to pre-define a discrete set of known-safe programs for their patients, should they be linked to identifiable use cases [129]. Discrete programmes benefit from including more complex parameter changes, such as both magnitude and frequency [127], without risking automated control drifting into non-validated state space. Note that pre-set programmes should not be limited to the set of operating modes the control loop is allowed to select from, fall-back and emergency modes should also be available to the patient and carers to initiate from the patient controller device [127–129]. Other

constraints on automatic parameter tuning can include rate-limiting to prevent rapid changes in operating state [127]. This approach was shown to reduce stimulation-induced sensory side-effects (paraesthesia) in Parkinson's patients [132, 133].

As a final general safety measure, the stimulation state should always be observable via the patient controller – keeping the patient informed of the delivered therapy is of critical importance [128, 129]. Stimulation states might be associated with side effects [127, 132, 133] which the patient should be able to account for in their activities, override if needed, or report to their clinician if problematic. Consider the case of a responsive seizure suppression system: even when an episode is successfully averted, a patient might want to take certain precautions for their safety – they must have the freedom to do so [134].

1.5.3 Stimulation artefacts

Stimulation artefacts are an ever-present challenge in neuromodulation therapies. In the control framework of Fig. 1.6, artefacts can be represented as an additional additive noise term before the sensor. Due to the high energy imparted to the tissue, common algorithms can misinterpret artefacts as high-energy states of a monitored neural marker. In a closed-loop controller aimed at suppressing a high-energy process (e.g. seizure), an activation of the stimulation might lead to a series of false-positive detections, driving the control loop into a self-sustaining state of continuous stimulation. This pathological state of a feedback controller is referred to as a 're-entrant loop'.

As a further complication, 'actuation noise' in a correctly-tuned closed-loop system will tend to be highly correlated with the state of the process under control, making its removal difficult without losing information on the process itself. It should be clear that the success of delivering safe real-time neuromodulation depends on our ability to deal with stimulation artefacts, thus we should spend some time to go over the signal characteristics of common artefact types.

Electrical artefacts

As an important draw-back of electrical stimulation, it can impact our ability to measure the brain's electrical activity – stimulation artefacts typically exceed neural signals by several orders of magnitude (single unit and LFP signals 10 µV–1 mV) [2, 28]. The tissue-electrode interface through common human-use platinum DBS electrodes has a typical impedance on the order of kiloohms ($\sim$1 kΩ) [103]. Assuming a purely resistive impedance, from Ohm's Law ($V = R\,I$) we can estimate stimulation artefacts to be on the order of volts (1 V) for stimulation currents

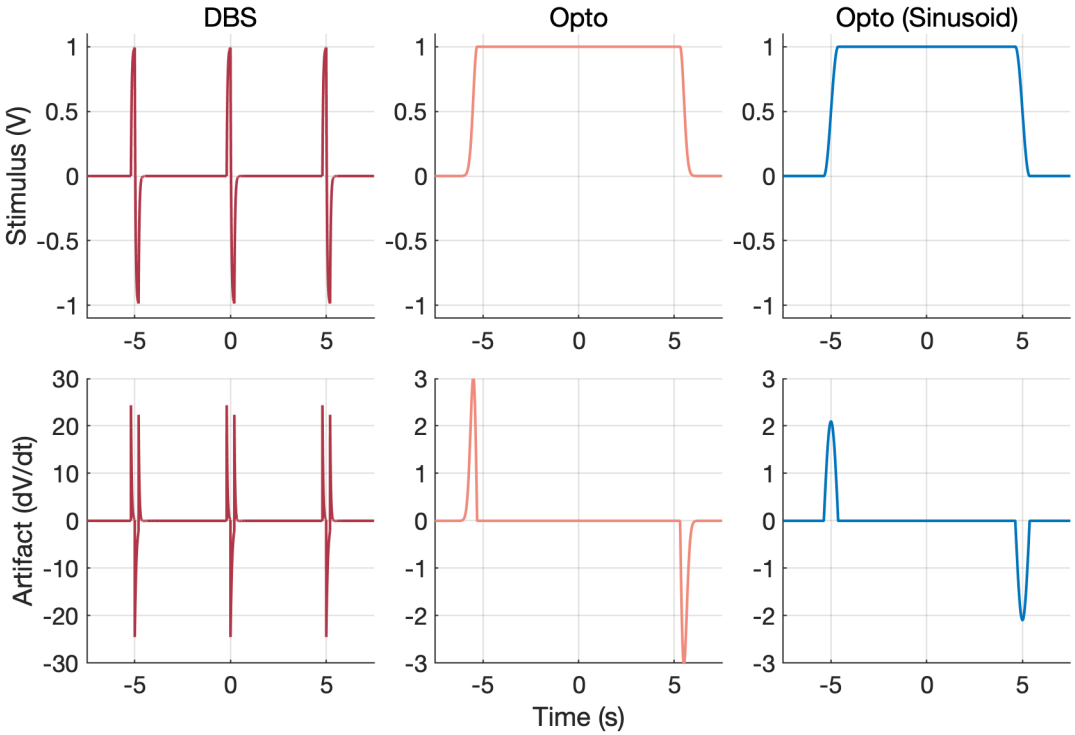


Figure 1.8. Comparison of electrical and optical stimulation artefacts. Both DBS and optical artefacts present as a voltage on the recording electrodes – DBS artefacts from current injected across the tissue-electrode interface, optical stimulation artefacts because of the photoelectric effect. The capacitive coupling of neural signal amplifiers acts as a differentiator on signal components below its time constant, creating the characteristic artefact shapes observed in recordings. Sinusoidal modulation of transitions in optical stimulation can shape the artefact for removal by narrow-band filtering.

on the order of milliamperes (1 mA), typical in human DBS therapy [28]. Amplifiers designed for measuring small signals might run into issues with saturation – that is, the amplified output would have to exceed the supported output range and can no longer track the input signal. No meaningful signal can be recorded while the amplifier is in saturation, thus the original signal cannot be recovered through post-processing. To overcome this limitation and allow for recovering the input, instrumentation with great dynamic range is necessary, so the superposition of the highest artefact magnitude and the neural signal never reaches or exceeds the operating range [135]. Even with a high dynamic range though, artefact removal can be challenging. Stimulation pulses are relatively short (commonly 50–150 µs), and the induced artefact are sharp, impulse-like events. Fig. 1.8 demonstrates the characteristic shape of electrical artefacts. Impulses carry broad harmonic frequency content, which can make appropriate filtering challenging to achieve without distorting the underlying neural signals.

Optical artefacts

Optical stimulation can also contribute artefacts to electrical measurements, due to the photoelectric effect: an electric current arising in materials when exposed to light. Photons absorbed by metallic and semiconducting materials above a certain energy threshold (proportional to frequency) can excite and 'dislodge' electrons to form an electric current [136]. The photoelectric threshold is a material-dependent property, and typically falls in the ultraviolet range for most common metals, however, most opsins respond to light in the visible range, which raises the question of why the photoelectric effect is an issue in electrophysiology at all. When electrons absorb a low-energy photon, the excess energy is soon re-released. The lifetime of the short-lived excited state is long enough though, that with sufficiently intense illumination there is a chance of subsequent photons being absorbed and making up the energy gap overall. This significantly lowers the photoelectric threshold in practice, with photoelectric artefacts prominent towards the blue (higher frequency) end of the visible electromagnetic spectrum, and decreasing in magnitude towards the red (lower frequency) end [136]. For further information on photoelectric artefacts, readers are referred to a detailed review article [136]. Interestingly, neural tissue tends to be more transparent to lower frequency photons, meaning the use of red-receptive opsins would simultaneously produce smaller artefacts, and waste less energy through absorption into the tissue. Reduced tissue absorption not only allows for greater energy efficiency, it also improves the safety of stimulation by reducing heating.

As light stimulation is switched, the photoelectric effect creates a potential excursion reflecting the illumination intensity over the conductive electrode surfaces, as shown in the top panels of Fig. 1.8. This is not the characteristic artefact shape we most commonly see in neural recordings (shown in the lower panels), however. The difference stems from the design of neural amplifiers, which contain a capacitive coupling element between the electrodes and the amplifiers [137–139]. The capacitors will charge to match the voltage steps, creating a short decaying impulse-like signal governed by the time constant of charging. As the impulses are generally not of pure frequencies, filtering is usually not the method of choice for their elimination [140]. Optical stimulation artefacts tend to be significantly smaller than their electrical counterparts, thus amplifier saturation is typically avoided. Most commonly artefact removal is attempted by creating an average 'artefact template' for each channel recorded in the dataset that can be aligned to and subtracted from the recording.

Alternatively, the use of sinusoidal transitions instead of (approximately) square edges results in artefacts of an equal sinusoidal frequency, shown in the right column of Fig. 1.8. Given that the artefact can be constrained to a singular frequency, a relatively narrow band-stop filter can be designed to recover the neural activity with minimal contamination. Further, the frequency of the sinusoidal segment, i.e. the transition rate can be purposefully selected such that the artefact and the associated elimination filter avoid any frequencies of analytical interest in a given study.

1.5.4 Technological constraints

In designing implantable therapeutic systems, we encounter challenging trade-offs between algorithmic complexity, performance and battery life. Open-loop implantables operate with static therapy configurations and receive only infrequent parameter updates, so complex digital parts can stay in a low-power sleep mode most of the time. The bulk of power is then used to deliver the stimulation current itself, which under typical settings is both small (order of mA) and very low duty cycle (e.g. 1–2 % in Parkinson's disease therapy), allowing for substantial battery life. In fact, open-loop implantables commonly use primary cell (non-rechargeable) batteries, while still being able to offer operating times on the order of years [139, 141]. In contrast, rechargeable open-loop systems require approximately weekly or biweekly charging [139]. The differences are due to a combination of increasing device complexity, differences in energy density in single-use and rechargeable batteries, and shrinking device sizes.

As battery capacity is so severely constrained, existing implantables are equipped with processing units quite limited in computational power by today's standards – saving on baseline power consumption by reducing processing capability. This hardware trade-off leaves severely limited headroom to implement neural state decoder algorithms.

Off-loading processing from the implantable to more powerful external devices, e.g. smartphone, tablet or desktop computer might seem to mitigate the power and complexity constraints. In practice however, the radio link connecting an implantable to the outside world suffers from high power consumption in itself [142, 143], making continuous data stream-out an unattractive option at present. In the distributed systems convention, such an external yet local computing device could be referred to as an edge-processor. Disregarding the power-cost of data streaming, edge devices can offer several orders of magnitude greater processing complexity compared to implantable processors and could become powerful tools where near-real-time feedback is sufficient for the therapeutic need [142].

Edge-processing can be supplanted by non-local high-performance computing resources, referred to as cloud computing. As a centralised resource, a cloud computing node would typically be shared between a number of users. Depending on momentary demand, processing requests might have to wait in a queue causing delays that are unpredictable on the user's side. Overall, cloud computing is best reserved for time-insensitive, long term refinements [144], e.g. detecting long term periodicity or trends in disease progression, used to compensate therapy settings over the course of weeks, months or years. While the theoretical latency introduced by edge-to-cloud data uplink can be as low as tens of milliseconds, the lack of standards on network reliability severely limit the potential for providing continuous feedback in a medical setting, and should only be used for safety non-critical monitoring for the foreseeable future. As a further consideration, the transmission of neural data opens up concerns of both security and patient privacy.

1.6 Overview

The preceding sections are meant to provide a snapshot of the state of the neural monitoring and stimulation technologies available to research, allowing us to systematically identify techniques with closed-loop potential. A summary of the explored methods is presented in Fig. 1.9. On both the measurement and stimulation side, non-invasive techniques are the most prevalent in the human research literature. This popularity can be largely attributed to the non-invasive nature itself, rather than to intrinsic benefits these techniques offer. In terms of everyday utility, magnetic measurements remain completely impractical due to instrumentation complexity (fMRI) and a requirement of shielding (MEG) despite great success at miniaturisation [81]. With optical sensing raising concerns about tissue heating under prolonged use (fNIRS), overall, active-dry EEG comes the closest to a convenient long-term method, at the cost of especially sparse spatial resolution [77]. As for intracranial measurement, electrical recording of neural activity is possible across various scales ranging from synaptic activity [15] to large population activity [29]. Techniques for finer resolutions suffer from long-term stability [56, 57] and prohibitively high data-processing burden [145], leaving coarser population activities the most attractive targets with near-term BCI potential [29, 70]. Overall, LFPs recorded from sEEG, DBS probes, and ECoG are of primary interest to our work.

As for stimulation, the considerations are broadly similar. Non-invasive methods tend to suffer from a trade-off between portability and spatial resolution. While tACS might offer lower stimulation intensities and coarser stimulation focusing than TMS, it has still produced strong effects in human participants [88]. As tACS systems can be made relatively portable, it could provide a complementary method to facilitate neural plasticity during other therapies, e.g. physiotherapy [146, 147] or conventional DBS [148], making it our preferred external stimulation method in this work. Electrical stimulation remains the only approved method among intracranial techniques and has already seen widespread adoption in DBS therapy. While optical stimulation systems continue to be developed in the form of optogenetics [117] and label-free modalities [118, 149], they are yet to address significant questions of safety in chronic use. Such systems will necessarily be restricted to non-human research, thus failing to offer short-term clinical relevance. All of that said, a research pipeline can benefit significantly from supporting optical stimulation due to its potential to help in disentangling therapeutic mechanisms of action on the neural circuit level in animal models of disease.

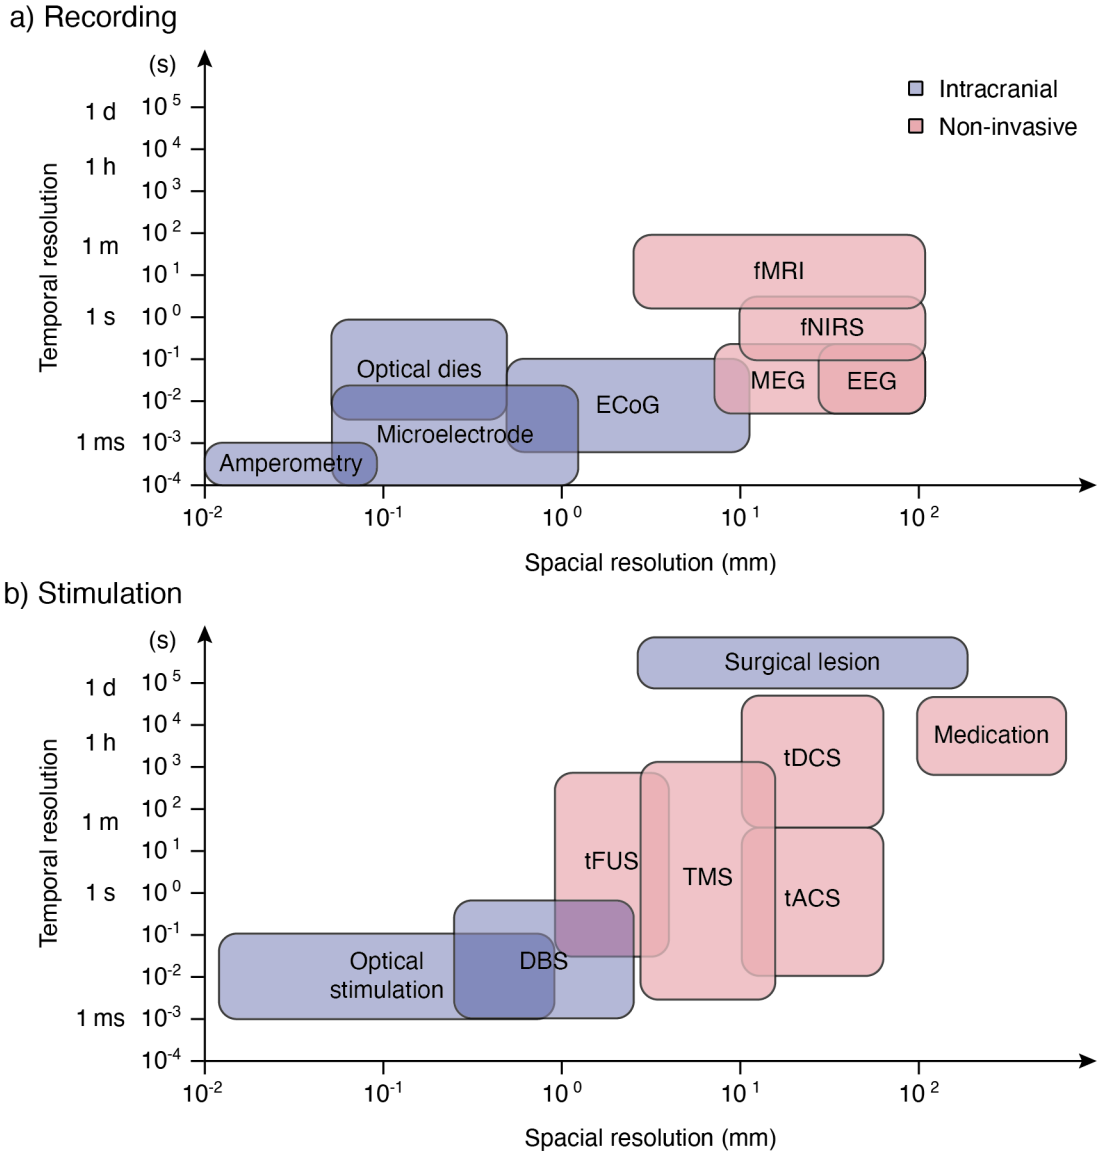


Figure 1.9. Spatial and temporal resolution of the techniques available to **a)** record and **b)** manipulate brain activity. Intracranial techniques are marked in blue, non-invasive ones in red.

1.7 Aims

Our work aims to provide technological refinements to DBS therapy, with clear pathways to clinical adoption. Our systems should ideally be suited for chronic therapy, but at the very least for at-home use to offer maximal patient benefits. These goals provide us with a set of useful constraints. Our developments should target incremental refinements of existing clinical systems to leverage a proven track record of patient safety, rather than attempting to develop complete neuromodulation platforms from the ground up out of academic curiosity. To offer compatibility with existing therapy systems, algorithmic refinements should be conscious of power consumption and computational cost. Based on our constraints and the technological review conducted, we identified three avenues of refinement, thematically aimed at targeting the past, present and near-future of DBS systems.

Chapter 2: DBS systems are typically restricted to interacting with two brain areas, often the same target bilaterally. Research suggests therapy could be improved by involving additional regions [70, 150], with multi-site stimulation potentially allowing for lasting therapeutic effects through the induction of neural plasticity [41]. We aim to design a portable system to synchronise external non-invasive neurostimulators (tACS) to an implanted DBS device through a manufacturer-agnostic interface, targeting patients with existing implantables.

Chapter 3: Stimulation therapy remains largely constrained to open-loop neurostimulation, despite a decade of evidence on the effectiveness of closed-loop responsive control [42, 151]. We aim to demonstrate that through careful algorithm design closed-loop control is already attainable to modulate the power of spectral biomarkers in a clinically suited implantable system.

Chapter 4: A further refinement to modulating oscillatory neural activity could be phase-targeted stimulation, promising significant power savings, and the ability to bidirectionally modulate activity unlike conventional high-frequency DBS [25,123]. A method is identified to provide computationally efficient real-time phase estimation, suitable for next-generation implantables. We aim at integrating the method into an optogenetics research platform to enable studying its efficacy, and mechanism of action prior to clinical research.

Chapter 2

Multi-site stimulation for therapy research

2.1 Introduction

Precise timing of activity throughout the brain is necessary for carrying out complex behaviours and processes. Neural oscillations are prevalent across all spatiotemporal scales and provide a natural mechanism for temporally coordinating activity within and across distributed brain regions [152–156], thereby affecting the transfer of information [157–163]. The dynamics of this rhythmic activity can contribute to synaptic plasticity [164], which relies on the specific timing of activity between interconnected brain regions.

Disruptions of the network can upset finely-tuned feedback loops, creating pathological oscillatory activity that is thought to be behind several common disorders of the nervous system [165]. When the timing of neuronal activity has been pathologically disrupted, chronic neurostimulation has been successfully used to help restore the dynamics. Over the past decades, the primary approach in clinical care has been to break-up pathological network synchronisation using continuous high-frequency stimulation [166, 167].

One widely studied condition is epilepsy, characterised by large-scale, amplifying waves of synchronised neural activity (i.e. seizures). Therapy can focus on suppressing the activity of the originating point of the activity wave, when identifiable [66]. Alternatively, stimulation can target highly connected nodes in the central nervous system, such as the centromedian thalamus, which may act as a key 'distribution point' in the evolution of a seizure, regardless of its source [168].

Another prominent case where stimulation therapy has been successful is Parkinson's disease. Parkinson's is also thought to be a network-condition, where the decay of dopaminergic neurons in the substantia nigra biases one of the pathways of the so-called cortico-thalamic loop. This

This chapter is based in part on prior published work [148], permissible under thesis regulations, endorsed by the co-authors. The content is adapted under CC BY 4.0.

network imbalance is characterised by exaggerated oscillatory activity in the beta-band [121]. High-frequency stimulation of targets related to the cortico-thalamic loop (globus pallidus or subthalamic nucleus) have been shown to suppress the oscillatory activity [169], significantly improving motor symptoms [25, 167].

Most implantable therapy systems currently use a static open-loop stimulation pattern preset by the clinician. As a refinement to therapy protocols, mainstream clinical research is now primarily focused on deploying closed-loop stimulation capability [150, 170]. Future stimulation algorithms may use more complex patterns designed to interact with the timing of neural activity in order to more selectively target pathological activity or achieve longer-lasting therapeutic effects [171].

In this chapter we focus on an alternative approach to therapy research, aiming to harness induced neural plasticity to achieve lasting effects of network weights to compensate for pathological dynamics in the cortico-thalamic loop in Parkinson's disease. In our prior discussions, we identified three distinct stimulation topologies capable of manipulating plasticity (Section 1.2.2), two of which we can quickly discount. High-intensity repetitive stimulation suffers from low efficiency, and is limited in the directionality of plastic changes. As an alternative, closed-loop neuromodulation for plasticity would require leaving well-established electrode implantation procedures for novel experimental targeting – an unacceptable risk to patients. These considerations leave us with multi-site (paired) stimulation as an efficient and clinically feasible approach.

Paired-pulse stimulation has been shown to successfully alter the excitability of the cortico-thalamic pathway in Parkinson's patients [41], by delivering paired DBS and TMS. However, the study did not demonstrate meaningful clinical outcomes due to several limitations worth addressing. The study relied on a small number of patients, available for only a single short session (minutes) with the paired stimulation protocol, and observation of post-stimulation effects for about an hour.

More extensive research is currently held back by the capabilities of clinically approved DBS systems. Implantables do not provide suitable telemetry protocols for interfacing to cooperate with external systems. DBS devices are designed predominantly with power consumption in mind; as a trade-off wireless interfaces suffer from high latencies for stimulation program updates and offer little functionality not needed in everyday clinical use: lacking an option for external pulse-to-pulse triggering. These factors limit opportunities for experimental work to newly implanted patients in restrictive time window of possibly a few days. Depending on local clinical

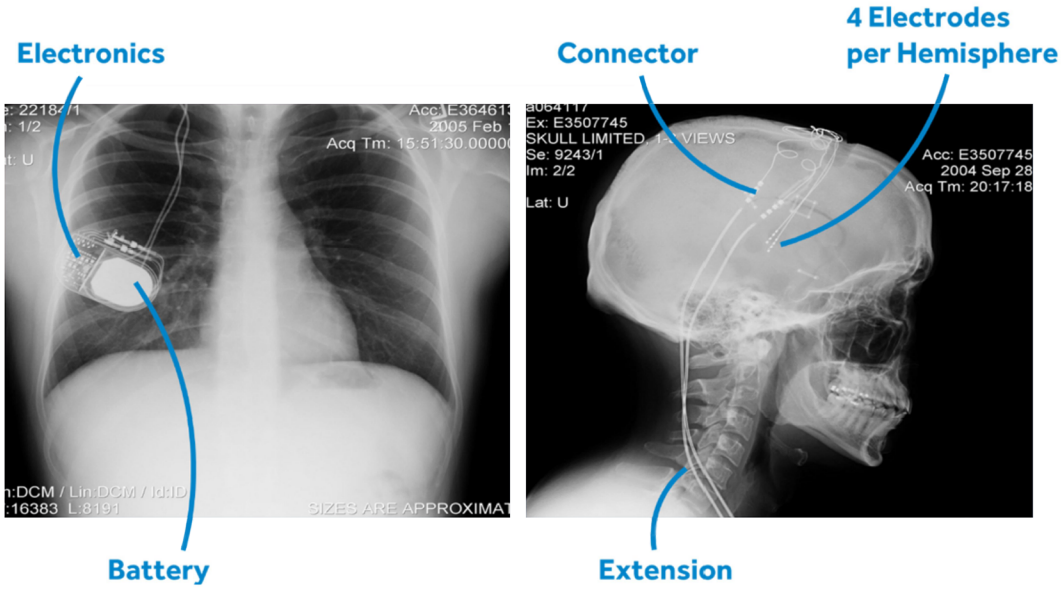


Figure 2.1. Subclavicular deep brain stimulation system. Adapted from [173] with permission from the authors.

best-practices, implantation may take place in two separate stages: inserting the DBS electrodes in a primary neurosurgery, and placing the subclavicular implantable pulse generator (IPG) in shorter procedure on a separate occasion (Fig. 2.1). The DBS leads may be accessed in between the two operations. Alternatively, externalised leads are briefly available in patients undergoing IPG replacement [42, 166, 172].

The following sections detail the design considerations, and feasibility analysis of a method for synchronising external cortical neurostimulation to fully implanted DBS systems. The proposed research setup sidesteps the conventional interfacing limitations of IPGs, making the recruitment of study participants significantly easier, and allowing for evaluating the effect of repeated stimulation session on symptom scores. While our planned application in Parkinson's disease, a toolset capable of manipulating deep-brain to cortical connectivity would be equally relevant in epilepsy and stroke rehabilitation research [146].

2.2 System proposal

To fulfil our aims of creating an IPG-agnostic research tool for multi-site stimulation, we propose a synchronisation system based on non-invasive sensing of the pulse train of an IPG, extracting timing information from this signal via a phase-locked loop, and using this timing to trigger an external stimulation unit with known and adjustable latency to the IPG. A high-level overview of the proposed system is represented in Fig. 2.2.

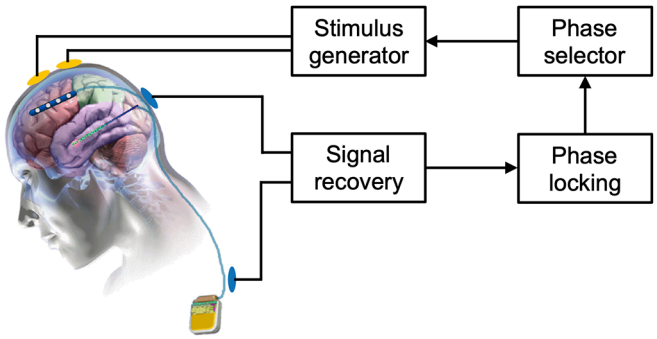


Figure 2.2. Paired-pulse stimulation setup. To form an IPG-agnostic research system, the stimulation pulse train of the implanted device is extracted through non-invasive electrical sensing, processed via a phase-locked loop. The obtained signal serves as a frequency and phase reference, stabilised by a phase-locked loop, and is used to adjust the timing of an external stimulator.

Our primary study of reference [41], demonstrating plastic modulation of the thalamo-cortical loop via paired stimulation, concluded that their operating frequency (3 Hz DBS) was likely one of the limiting factors towards a clinically valuable result. Thus, rather than aiming for a replication of the study in chronic conditions, we should explore whether existing literature might provide guidance on selecting a stimulation pattern for our intended study.

Structures within the thalamo-cortical loop are known to be rich in beta-band oscillations [121]. The suppression of this oscillatory activity [169] is directly linked to an improvement in motor symptoms [25, 167]. This suggests beta frequencies are a natural resonant mode of the network, thus attempting to modulate coupling at this rate could benefit from greater entrainment, leading to more significant results that the previous attempt using low-frequency stimulation. In this setup, the DBS pulse train will act to entrain the endogenous beta oscillations into a known phase. When stimulating at beta frequencies, DBS intensity should be adjusted not to significantly potentiate the pathological oscillations in order to minimise discomfort in patients. Notably, beta-band stimulation has also been shown to effectively entrain cortical activity in a non-invasive stimulation study [174].

In summary, our initial research protocol will deliver beta-frequency DBS, coupled with beta-frequency external stimulation at different phase-offsets. The following sections aim to derive and elaborate tangible target specifications for this example protocol.

As an overarching requirement though, we aim to design a research tool that is compatible with approved clinical systems and practices, yet is flexible enough to merit adoption in the clinical research community.

2.2.1 Signal range

An immediate question arising from the design topology, is whether DBS pulses can be reliably sensed transcutaneously. To evaluate feasibility of signal pickup in advance of in-patient testing, we can turn to electrocardiograms recorded from patients receiving DBS therapy under routine clinical care [175, 176]. Fig. 2.3 shows such a measurement, taken in a Parkinsonian patient implanted with an IPG according to the classical subclavicular arrangement introduced previously in Fig. 2.1. The most contaminated channels show an artefact magnitude of approximately 1 mV in monopolar mode, while artefact is not visible (estimated $< 20\,\mu V$) in bipolar mode [175, 176]. As the recording was carried out according to the standard practices of electrocardiography, none of the measurement dipoles align well with the direction of current flow in the implanted system. Intentionally placing the measurement points along the lead, e.g. one above the IPG and another near the cranial burr hole, could provide DBS artefact magnitudes on the order of 10 mV, significantly exceeding the cardiac activity. Still, we expect the system to be limited to sensing monopolar stimulation programmes (Fig. 2.3). Monopolar stimulation is a popular therapy option in Parkinson's disease [177], patients normally receiving bipolar stimulation can still be safely transitioned to monopolar stimulation [177] for the duration of relatively brief research and therapy sessions with multi-site stimulation. While bipolar stimulation has been shown to significantly improve the battery longevity of non-rechargeable IPGs [178], the brevity of sessions would also preclude any material adverse effect in this regard.

2.2.2 External stimulation

The choice of external stimulation technique has a significant impact on further design requirements. While our reference study was based on TMS, the author's have singled out the inflexibility of this method as one of their most significant limitations [41]. Most commercial TMS systems are held back by the complex power electronics required to drive the magnetic field, limiting stimulation delivery to short pulsed patterns at very low frequencies. The main

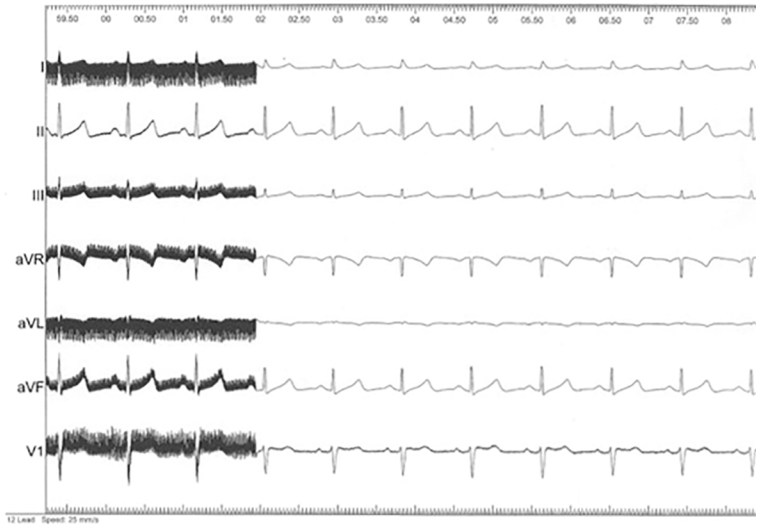


Figure 2.3. DBS artefact visible in an electrocardiogram. The measurement was reportedly taken via surface electrocardiography in a patient with early-onset Parkinson's disease receiving DBS therapy. The stimulation transitions from monopolar to bipolar mode at about a third of the recording length, resulting in a marked decrease in artefact magnitude. Note that no classical electrocardiography measurement dipoles are collinear with the implanted DBS leads. Intentional placement of skin measurement electrodes for sensing the DBS artefact would result in greater magnitude. Reproduced from [176] under CC BY-NC-ND 4.0.

benefit of TMS remains its ability to directly elicit neuronal firing. While recent designs have recently overcome these limitations [98], the requirement of as yet experimental stimulation equipment for high-quality TMS studies is not at present compatible with our aim to design a widely adaptable research tool.

In comparison, tACS is thought to act more along the lines of biasing the membrane potential of neurons, and affecting their excitability. This biasing can contribute to the strengthening of a particular endogenous oscillation via larger scale entrainment. Successful entrainment also depends on stimulating each particular region at the right frequency [179], and the availability of a stimulation montage that allows for sufficient depth of activation, while minimising impact on reliable detection of DBS pulses.

The gentler modulatory effect of tACS is not necessarily a limiting factor in the success of our proposed study though, since one of our major improvements compared to the original experiment [41] is interoperability with chronically implanted IPGs. This enables providing each participant with multiple stimulation sessions over the course of a few days, compensating for the potentially lower efficiency of tACS-facilitated neuronal activation. Overall, the stimulation flexibility that tACS allows for along with its widespread availability makes it our preferred method of choice.

As a technical requirement for the selection of a specific tACS unit, a hardware input channel is required to time stimulation. Based on a brief survey, most devices lack this capability and emphasise software control. One suitable tACS system identified on the market with approval for clinical research is the Starstim family of combined EEG–tACS devices (Neuroelectrics SLU, Barcelona, Spain), another the neuroConn DC Stimulator Plus (neuroConn GmbH, Ilmenau, Germany).

2.2.3 Artefact tolerance

A further consideration is whether the external stimulation source would interfere with our measurements of the DBS signal. Preliminary data shown in Fig. 2.4 suggests that tACS operated in a cephalic montage does not pose a significant challenge in the sensing of monopolar DBS stimulation, with stimulation artefact not exceeding the order of 5 mV in an clavicular-referenced signal pickup configuration.

While our primary study of reference relied on TMS [41], the strong magnetic fields are expected to introduce a large pulse-like disturbances to measurement, potentially difficult to distinguish and separate from DBS pulses. On the other hand, tACS delivered at pure sinusoid frequencies can be easily filtered in the signal pickup front-end, suggesting that tACS is a more appropriate external stimulation device for our system.

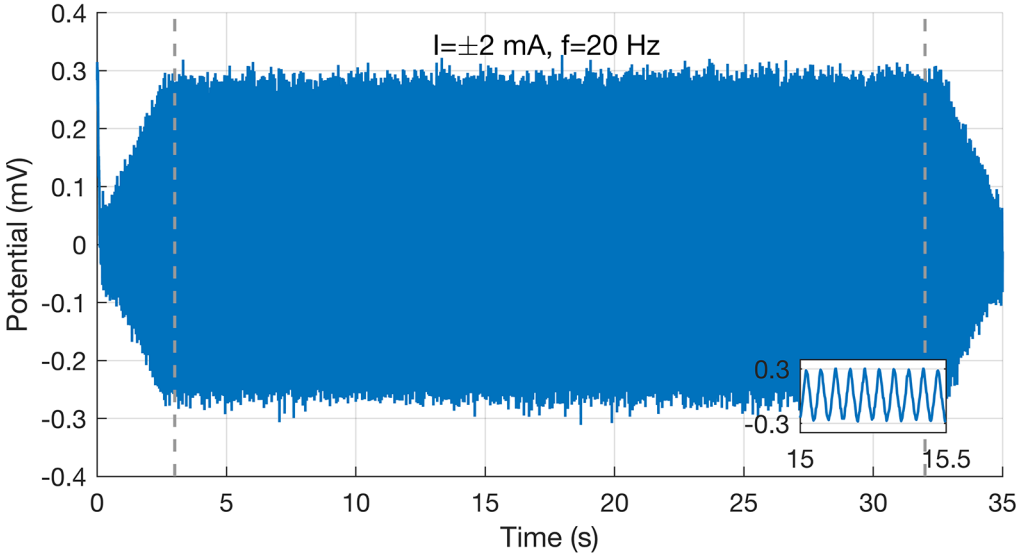


Figure 2.4. Measurement artefact from transcranial alternating current stimulation. The recording was taken over a clavicula-cranium measurement dipole in a healthy human subject without an IPG. Stimulation and recording were carried out using a Starstim 8 system in a 4x1 cephalic montage. Data courtesy of Dr Bettina Schwab (University of Twente, Enschede, Netherlands).

2.2.4 Frequency range

We have previously specified our target application as plastic modulations targeting the coupling of thalamic beta-band activity to cortical regions in Parkinson's disease by paired pulse stimulation at beta-band frequencies, this is not a typical therapeutic frequency. DBS is most commonly used at frequencies on the order of 120–140 Hz. At the same time tASC studies most frequently report settings within the classical EEG frequency band of 0.5–40 Hz (slow-wave to low-gamma range) [95]. Overall, it is not immediately clear what might be an optimal target range for our synchronisation system.

Giving relevance to accommodating higher frequencies, cortical gamma-band activity has been shown to serve as a meaningful feedback signal to subthalamic DBS delivery in a closed-loop pilot study, reducing stimulation-induced side effects in Parkinson's patients [150]. A recent tACS study further indicates that it might be the coupling of gamma-band activity in the cortex to low-frequency (theta-band) processes that significantly couples motor and deep brain areas [88]. Potential follow-ups of these studies in Parkinson's could be feasible with our synchronisation system, delivering packets of gamma-frequency tACS with a known phase relation to theta or beta frequency DBS.

Overall, we note a frequency range of immediate interest for tracking DBS, with a lower limit in the theta band (nominally 5 Hz), and a high frequency limit in the gamma band around 75 Hz — that is to allow for evaluating 1:1 coupling to high-frequency tACS as potential control conditions. As there is broader research interest in tACS at up to frequencies of 140 Hz (ripple band) [180], that conveniently includes high-frequency therapeutic DBS, we can define an extended frequency range of interest over the range 5–140 Hz.

2.2.5 Resolution

A key question in choosing an appropriate implementation for our synchronisation system is our targeted timing requirements. As a physiological point of reference, the axonal conduction delay between primary motor cortex and striatum is reported on the order of 6–7 ms in mice [174], suggesting a reasonable target resolution might lie in the millisecond range. Considering the strongly decaying efficiency of plastic manipulations outside the timing optimum (±20 ms, Fig. 1.2), we settle on a requirement to accommodate latency granularity between our two stimulation sites down to 1 ms. As our specification is set in time and not phase units, it is noteworthy that our requirement for phase will decrease with increasing DBS frequency.

2.2.6 Accuracy

The accuracy and uncertainty of our phase-tracking signal are important requirements for reproducible research. Without existing guidance, we set our phase uncertainty as $\pm 1\,\%$ as a strict guarantee, which translates to 70 µs at our high-frequency limit of 140 Hz, and is more relaxed at lower frequencies.

2.2.7 Latency

Latency incurred by communicating with external equipment, namely the tACS unit of choice, could easily render highly accurate measurements of DBS pulse timing meaningless. Given that paired-pulse stimulation is a repetitive protocol with regular stimulation frequency, our requirements for success are reduced largely to having a constant and finite latency: any deterministic latency can be accounted for by adjusting our target phase.

As an example, Starstim units were measured to operate with a computer-dependent, but constant latency on the order of 2 ms for receiving data, and 20 ms for issuing stimulation changes. With a further requirement that DBS pulses and tACS periods be cycle-to-cycle corresponding, Starstim units would be limited to operate at a maximal 50 Hz paired frequency — the cost of an unavoidable computerised interface.

2.2.8 Medical safety

To meet our goal of creating a a research tool ready for experimental use, we must consider integration into the clinical environment, particularly implications to patient safety. Given that the study relies on electrical stimulation, Galvanic isolation can be used to break up any stray current paths that may exist between connecting systems or the patient. Assuming the external neurostimulator of choice is a Starstim device, a minimal fully isolated setup is proposed in Fig. 2.5. Recent developments in low-noise isolation amplifiers allows for the reliable sensing of millivolt-scale signals without preamplification. Before use, the complete system shall be evaluated for electrical safety according to the relevant ISO 60601 standard.

For simplifying the approval process, an analogue output commercial ExG amplifier might be added to the input stage. In this case, both the sensing and stimulation unit have existing medical ratings, and the custom instrumentation has no patient-facing connections at all. Unfortunately the commercial space of analogue ExG is devices is on the decline in favour of purely digital interfacing units. A particular limitation in both analogue and digital systems is the restriction of recording bandwidth: for a device to function in our application front-end

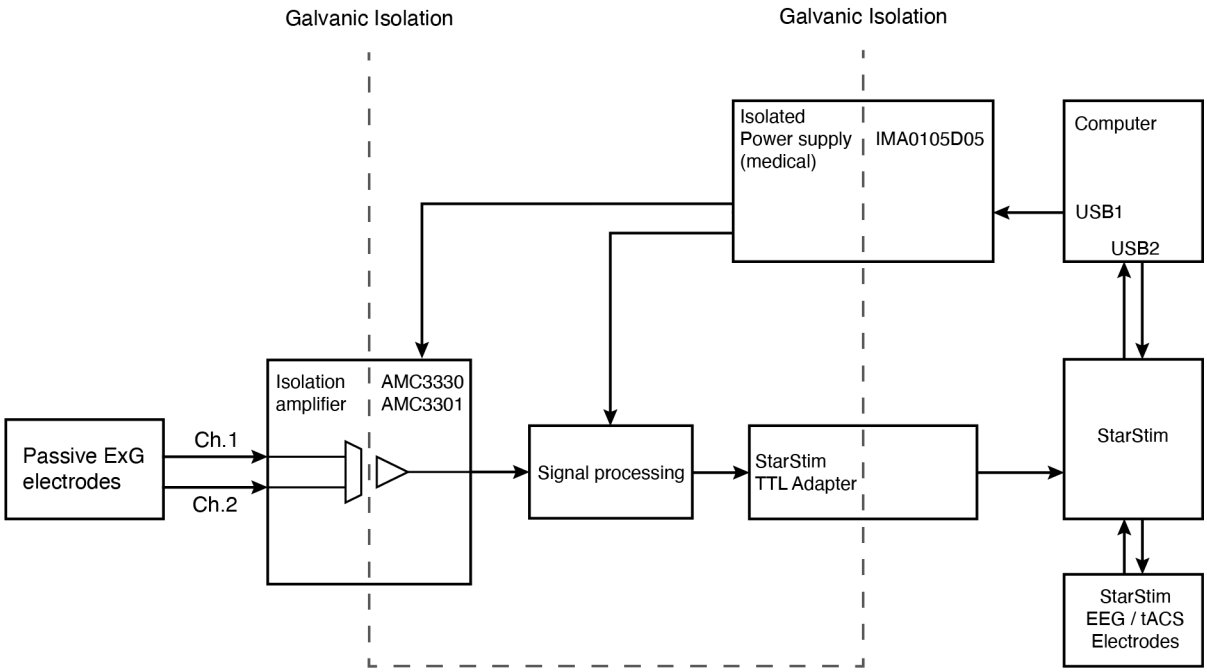


Figure 2.5. Paired-pulse stimulation system overview. To ensure compliance with patient safety protocols, galvanic isolation is provided on all inputs and outputs of the synchronisation device. The double isolation prevents ground loops as well as undesired (and potentially unsafe) leakage current between any of the patient, the custom instrumentation and the external neurostimulator.

filter options must be available to allow the propagation of short DBS pulses, i.e. beyond 17 kHz low-pass corner at 60 µs pulses. One suitable device identified with medical research approval is the Digitimer D360 (Digitimer Ltd., Welwyn Garden City, United Kingdom).

2.3 Technical method

2.3.1 Phase-locked loop

Phase-locked loops (PLLs) are commonly used to provide synchronisation between two digital clocks (periodic signals). In order to provide synchronisation, PLLs rely on a feedback path to tune an internal clock signal (S') such that it matches a provided reference signal (S). The phase of S' is continually compared to the target signal (S) via the feedback path, with the resulting phase error signal used to adjust the oscillator frequency. Phase error between two digital signals can be measured using an exclusive or operation when both are near equal in their duty cycle, or via an edge-triggered detector in the more general case where only consistent recurrence is evaluated. In either case, the error signal requires filtering to limit the effect of rapid changes in phase error on S', ensuring the stability of the feedback loop. Once the system settles, S and S' are matched in both phase and frequency. The analogue realisation of a PLL is shown in Fig. 2.6a, where the internal clock is typically created by a voltage-controlled oscillator (VCO).

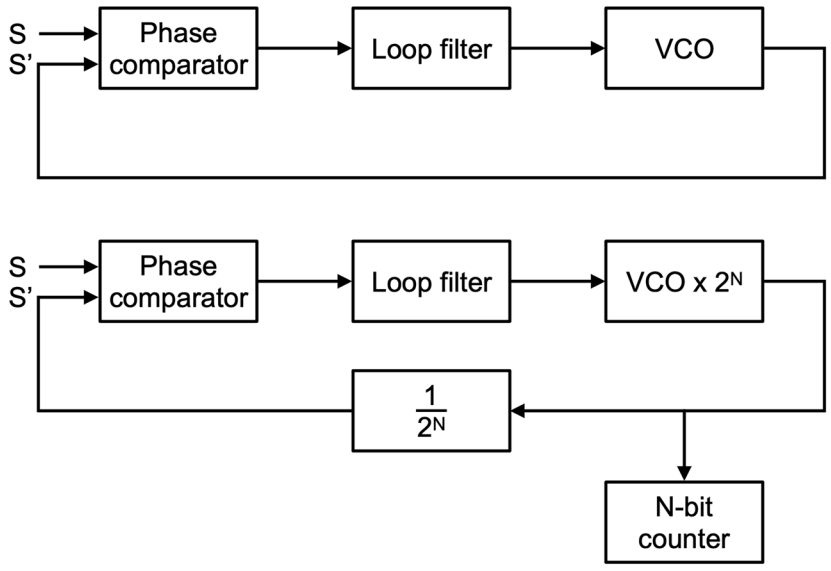


Figure 2.6. Phase-locked loop operation theory. **a)** Conventional analogue phase-locked loop based on tuning a voltage-controlled oscillator (VCO) element to provide an internal clock signal S' matching the reference signal S in both phase and frequency. **b)** By first adjusting the VCO to operate at a multiple of the expected reference frequency range, then dividing the clock by an equal amount, the loop remains in a stable lock while creating an intermediate high-frequency signal that subdivides each cycle of S at regular intervals. This intermediate signal can drive a digital counter to provide a real-time quantised estimate of phase progression.

The system discussed so far, however, is not yet sufficient to measure the phase of the input signal, and one might wonder what even is the point of creating an internal copy of an already existing clock signal. The answer lies in our ability to manipulate the feedback path to create arbitrary internal clock frequencies at multiples of the reference frequency, allowing us to subdivide each cycle of the input clock into regular intervals. By first adjusting the VCO to operate in a frequency range k times greater than the target signal S, then inserting an k-fold frequency divider in the S' feedback path, the loop remains stable, with S and S' still matching in both phase and frequency. However, the intermediate signal now performs k cycles for each period of S. A simple counter can be used to establish our progression in the cycle, providing a real-time estimate of phase. This modification of the loop is shown in Fig. 2.6b, with $k = 2^N$ as a convenient choice for digital counting.

2.3.2 Realisation

The full design of the phase-locked loop is presented in Fig. 2.7. Given the low duty cycle of DBS pulses, the system uses a sequential, positive edge-triggered phase detector. The lead-lag compensation filter was implemented to help balance the low-frequency disturbances while maintaining the stability of the loop. To generate the phase of the input signal in real-time, we used a frequency synthesiser topology. In this configuration, the voltage-controlled oscillator's (VCO) centre frequency is set as 2^N times the average input frequency (N being the number of bits in binary phase counting). Feeding this high-speed clock into a binary counter, each period of the input signal can be subdivided into 2^N segments, thereby quantising the phase. By choosing an edge-triggered counter (matching the phase detector), bits 0 through $N - 1$ represent the phase code, where bit 0 is the high-speed clock itself. Bit N can be used to reset the counter. To enable successful locking, the high-speed VCO clock has to be divided by 2^N to be of similar magnitude to the input – this signal is available on bit $N - 1$ of the counter.

Using an $N = 8$-bit phase accumulator, we were able to keep track of the phase at a resolution of 1.4° at all frequencies within the lock-in range. At a first glance this resolution might appear be excessive in the context of high-frequency signals. However, 8-bit is the lowest resolution available to satisfy our latency granularity specification of at most 1 ms at the low-frequency limit of 5 Hz, required for fine-tuned plasticity studies.

The phase selector commands the external stimulation generator, which can be gated by the lock-in detector signal, to prevent stimulation while the phase estimate is not in steady-state.

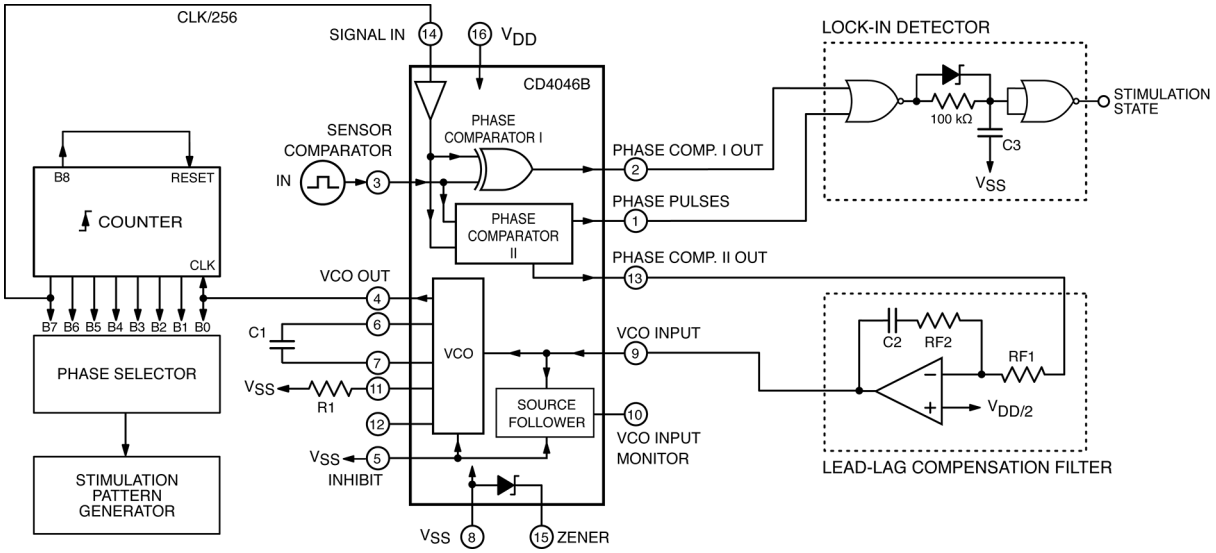


Figure 2.7. Phase-locked loop implementation for DBS frequencies. Phase comparators and voltage-controlled oscillators were provided by a CD4046B integrated PLL device [181]. The counter used was a CD4040B 12-bit counter converted to positive edge triggering with an inverter-configured CD4007UB. The 8-bit phase code is available as parallel output to interface with stimulation controllers. Lock-in status detection is implemented with a CD4001B NOR gate chip.

2.4 External validation

The design of Fig. 2.7 was evaluated using validation signals representing challenging end-cases of monopolar brain stimulation (Fig. 2.8), where the potential measurement is taken between the burr hole on the surface of the head and the implantable pulse generator. Key characteristics identified in the design aims are presented in Table 2.1.

Table 2.1. Performance characteristics of the prototyped PLL designs.

Specification	Value (Passive)	Value (Active)	Units
Capture range	$5-145$	$55-235$	Hz
Lock range	$5-145$	$55-235$	Hz
Lock-in time	380	128	ms
Minimum input pulse	14	14	µs
Jitter (peak)	± 4	± 19	µs
Phase resolution	1.4	1.4	°

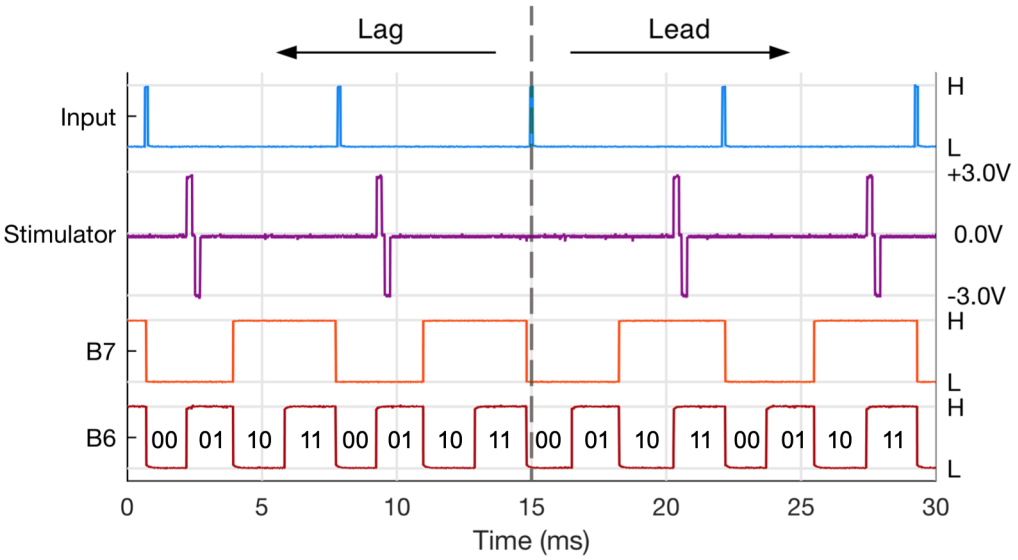


Figure 2.8. Phase partitioning example. Four-quadrant phase partitioning is demonstrated by observing bit B6 of the counter along with the synchronised clock. Relative timing for the external stimulation pattern generator is selected through binary word recognition in the phase selector. The system was successfully interfaced with an in-house biphasic pulse generator, demonstrating the ability to apply both phase lead and lag ($\pm 90°$ in this example) compared to the reference signal.

The system shows quick settling dynamics due to the choice of an underdamped loop-filter response (Fig. 2.9). Overshoot in the VCO input signal (by extension in the output frequency) while often undesirable, is mitigated by the lock indicator signal that can be used to prevent applying stimulus during the settling phase. As the experienced jitter translates to a phase uncertainty of about ±0.8° at our high-frequency limit of 140 Hz, 8-bit phase resolution reaches the limitations of this circuit. The use of an edge-triggered phase comparator ensures that the capture range and lock range of the circuit fully overlap, meaning that the system is able to lock on any signal within its range without a priori knowledge of the stimulator frequency, without significantly affecting the lock-in time.

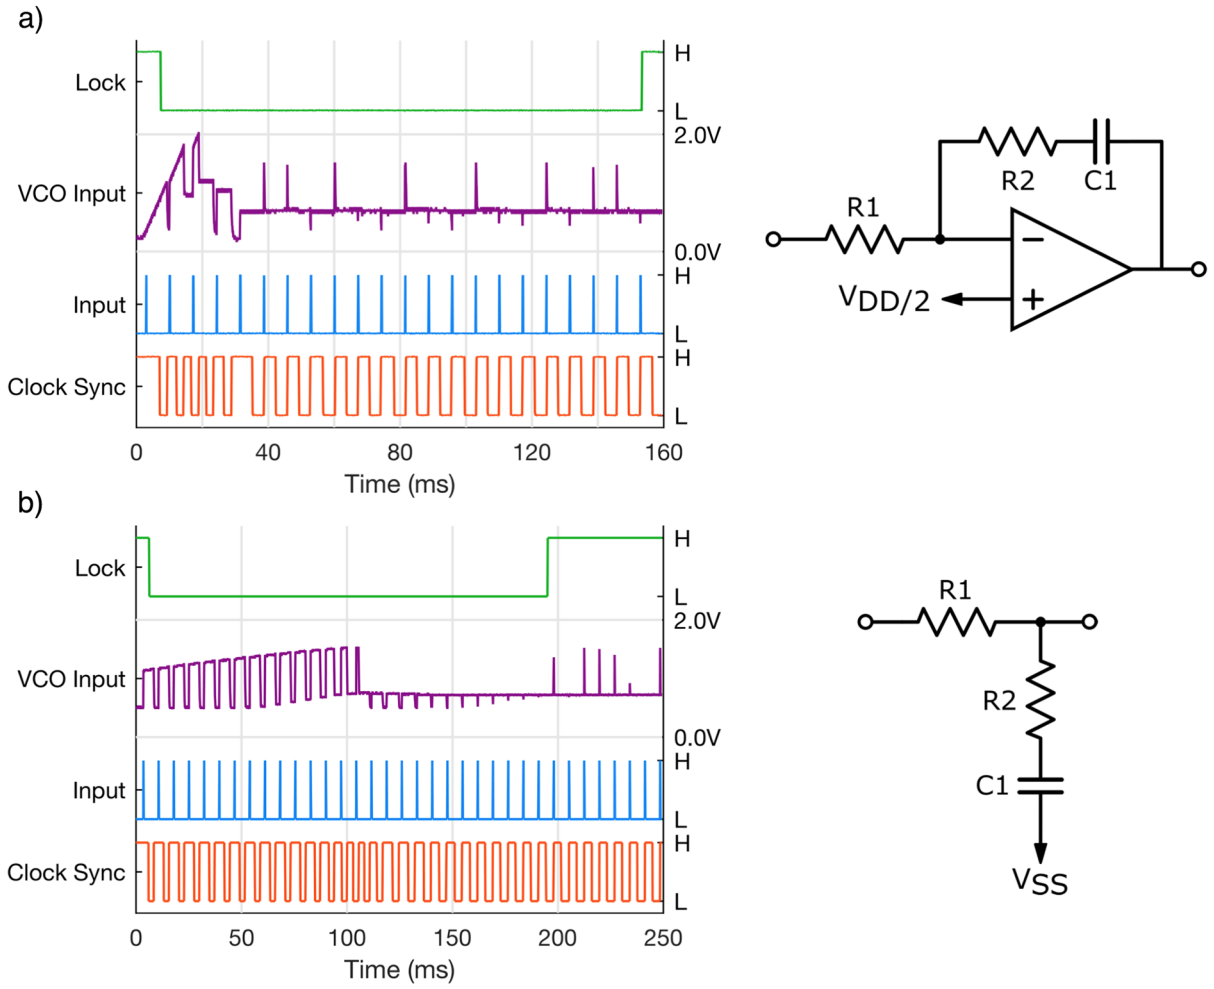


Figure 2.9. Locking dynamics of the PLL system on a 140 Hz pulse train, with **a)** active **b)** passive loop filter. Once the device is powered on and an input signal (blue) is available, the lock-status indicator signal (green) goes low within one period of the input signal. The phase mismatch between the divided oscillator output (orange) and the input signal causes the VCO input (purple) to change, thus adjusting the VCO frequency. A successful lock is marked by the status indicator signal going high. The overshoot dynamic observed is dictated by the damping ratio of the filter.

2.5 Discussion

2.5.1 Overview

We presented the successful development and characterisation of a method for phase synchronisation that coordinates implanted and external neural stimulators. Phase synchronisation between stimulation pulses and sensed physiological oscillatory signals have been effective at specifically targeting rhythmic activity in Parkinson's disease and essential tremor [25, 122], and we wish to explore further applications with distributed neural systems, targeting plasticity in the thalamo-cortical loop. The active implantable can provide access to subcortical structures not readily reached with existing minimally invasive techniques, while the external stimulator can provide broader cortical coverage that can be difficult to achieve with an active implantable system [41].

We chose to implement a phase-locked loop (PLL) to allow for flexible programming of relative stimulation clocks immune to drifts and other variations. The time to lock can be traded off against the stability to clock perturbations, and our design using a lead-lag filter attempts to optimise this choice while maintaining loop stability. The design is implementable with a small circuit that can operate in portable systems with minimum computational overhead. The system is applicable in a wide range of distributed stimulation setups, as it is possible to compensate for any deterministic phase error by programming the phase selector appropriately, allowing to take into account factors such as propagation delays between peripheral and central nervous system stimulators.

2.5.2 Limitations

Our specific design choices create limitations for the general applicability of the method.

Signal pick-up

In order to form a reliable measurement dipole, a long measurement axis parallel to the stimulation lead is required, ideally between the clavicle and the cranium. To prevent interference, this practically limits tACS to cephalic montages, significantly decreasing the field of activation. Complete implementations of the system might consider adding a front-end filter to the signal chain, as tACS artefacts at pure sinusoid frequencies can be easily removed by low-pass filtering without affecting DBS pulses, whose effective frequency is in the range of 10–20 kHz (based on the exact pulse with).

Reliability

The use of a sequential phase detector makes the design susceptible to a missing clock or excess noise. For the well-defined pulses from monopolar brain stimulation, this is an acceptable trade-off. However, for tracking physiological signals and perhaps bipolar stimulation we would use alternative phase comparators such as an analogue mixer. As measured, the passive loop filter takes 150 ms to recover from a single missing pulse, while recovery after two missed pulses is already identical to startup conditions, representing complete loss of tracking. On the other hand, the PLL design tolerates up to 2 % drift per cycle of the reference clock, well over the design specification of modern clinical DBS systems.

The locking of lower frequency signals pushes the boundary of acceptable analogue components, and an alternative variant using digital signal processing might provide more degrees of freedom, at the expense of computational energy [171].

External stimulation method

The application of stimulation from an external source might perturb our clock measurement, particularly in the case of TMS, where artefacts magnitudes can be significant, and appear similar (sharp pulses) to the expected DBS pulse train. Care must be taken to avoid a 're-entrant' where the PLL might clock off of stimulation artefacts, then proceeds to instruct stimulation further, creating a self-clocking pathological state.

To prevent this mode of operation, appropriate isolation methods must be taken, such as blanking the input signal during stimulation [170]. This is not easily achieved in analogue phase-locked loops, however easily implementable in digital designs. Finally, and most importantly, it is imperative that each use-case be evaluated for patient safety, and any undesirable interaction between the stimulators identified and mitigated.

2.5.3 Future development

While it is unlikely that 8-bit phase-offsets in the higher frequency operating region will be of experimental interest, our ability to offer rapid lock-in stable readout of stimulation state still carries significant research potential. The achieved lock-in time on the order of 150 ms for 130 Hz stimulation offers an avenue to study multi-site stimulation while the IPG is running adaptive DBS protocols, such as those piloted in Parkinson's disease [42], where stimulation can activate and deactivate on timescales of ∼500 ms. In case faster lock-in times were required, a field-programmable gate array or microprocessor-based digital phase-locked loop design should

be considered. Analogue VCOs are predominantly free-running, with PLLs employing them relying purely on frequency adjustments to reach phase alignment. In contrast, digital controlled oscillators can allow direct adjustment to both phase and frequency, offering even faster lock-in times.

2.6 Contribution

The analogue PLL circuit and its validation are my own work, with contributions from the co-authors of the original publication on this work: Abbey Holt conducted a literature review on multi-modal stimulation experiments, Moaad Benjaber provided design review, Andrew Sharott helped define the physiologically relevant parameter space, Tim Denison conceptualised the project and taught me PLL theory. Building on the publication, the example use case presented here was defined by the group of Bettina Schwab at the University of Twente (Enschede, Netherlands) who also provided artefact pilot data, with Nil Ramón i García characterising the Starstim latencies.

Chapter 3

Real-time power estimation for closed-loop stimulation

3.1 Introduction

The advent of deep brain stimulation therapy lead to great success in the treatment of a broad range of neurological and psychiatric conditions. First approved for the symptomatic treatment of essential tremor in 1997, approved usage rapidly expanded to other movement disorders, such as advanced Parkinson's disease (2002) and dystonia (2003) [165, 182]. The next major condition to see significant attention was refractory epilepsy, with clinical approval following in 2014 [168]. Since then, research into targeting a broader class of neural targets and disorders has been on the way, including disorders of consciousness, and psychiatric conditions such as major depression [165].

Despite the diversity of medical use cases, DBS therapy invariably tends to rely on fixed patterns of stimulation (most commonly 130 Hz) [165]. With only a handful of exceptions, the stimulation programme is tuned only by clinicians at patient follow-up visits, with no element of dynamic day-to-day adaptation, irrespective of symptomatic state. For detailed specifications of contemporary DBS device capabilities, readers are referred to an external review [183].

Several modalities of stimulation adaptation are yet to be exploited for refining DBS therapy, including adaptations to motion state long a staple of cardiac pacemakers [184], or circadian scheduling to take advantage of periodicities of disease states [126, 127].

In this chapter, we present our efforts to deploy a simple, but clinically relevant framework of neural state adaptive, closed-loop deep brain stimulation in the Picostim-DyNeuMo family of IPG systems.

3.1.1 The Picostim-DyNeuMo platform

The DyNeuMo (Dynamic NeuroModulator) is an experimental use firmware extension of the Picostim cranial-mount implantable pulse generator system developed by Bioinduction Ltd. (Bristol, UK) for adaptive DBS therapy research, shown in Fig. 3.1. The Picostim predicate device and its capabilities were designed to match conventional DBS systems while providing refinements to the implantation procedure [127].

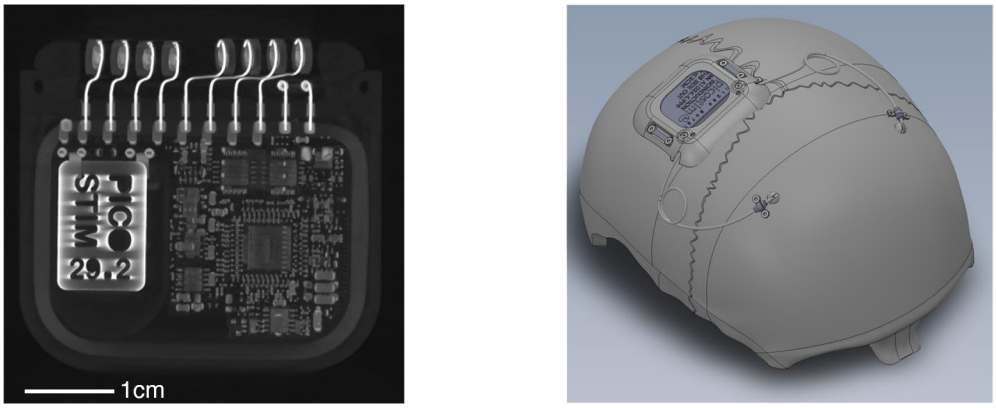


Figure 3.1. The Picostim-DyNeuMo device, and its intended cranial-mount placement. Images courtesy of Bioinduction Ltd., Bristol, UK.

Fig. 3.2 outlines the key technological features of the platform. As shown, the device ecosystem consists of three main hardware elements: the IPG, a handheld controller for patient use, and a clinician programmer interface (CPI). The handheld controller provides a MICS-band (Medical Implant Communication System) radio link [185] to the implantable, giving patients the ability to select between a set of clinician-approved programmes or to disable stimulation altogether. The handheld unit also provides a readout of the currently active stimulation programme, an important safety consideration. When used in conjunction with the CPI, the wireless link allows clinicians to securely programme therapy settings, download diagnostic logs, and perform firmware updates.

The Picostim IPGs were designed to be rechargeable via an inductive wireless charging protocol, offering a typical charge cycle of approx. 1 week for classical Parkinsonian DBS therapy, or between 1–2 days in advanced adaptive modes. For detailed technical specifications, readers are referred to consult the original publications on the DyNeuMo systems [126, 129].

Design and validation of the adaptive platform extension were carried out with ISO 13485 design controls with ISO 14971 compliant risk management standards and followed the IEC 60601-1-10 physiological control framework introduced in Section 1.3.2 (see Fig. 1.7). Functionality is partitioned between the handheld patient controller and the CPI (Fig. 3.2). The DyNeuMo operates

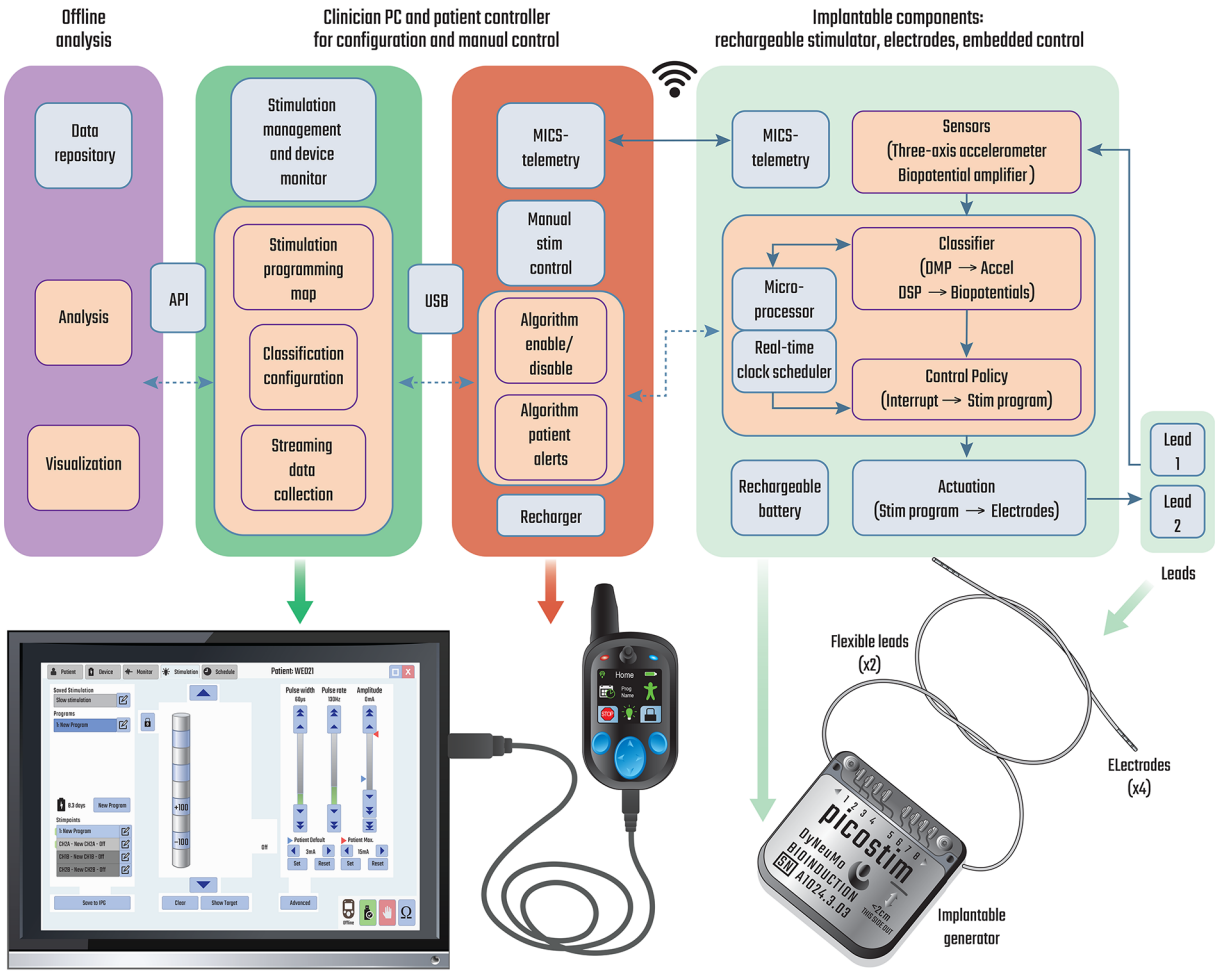


Figure 3.2. Capabilities of the Picostim-DyNeuMo platform. Note how the system architecture (top) conforms to the IEC 60601-1-10 physiologic control framework [128], shown in Fig. 1.7. The complete Picostim-DyNeuMo research ecosystem comprises a clinician programmer device for configuration and diagnostics, a handheld patient controller for at-home recharge and manual adjustments, and the IPG hardware. Reproduced from [126] under CC BY 4.0.

strictly on clinician-programmed stimulation patterns, including in all its adaptive modes. The CPI allows for configuring stimulation limits and fallback modes (see Fig. 3.3), while the handheld unit displays the current programme state, allowing patients to enable or disable, as well as change between approved stimulation patterns.

To support the case for first-in-human deployment, the device provides stimulation capabilities compatible with those of existing implantable neurostimulation systems, ensuring patients still have access at minimum to conventional therapy protocols even if they are found not to benefit from experimental modes [129].

To facilitate adoption as a research platform, the development team aims to provide a unified set of reference materials, including documentation of functionality and safety testing, to aid in supporting the preparation of clinical trial applications for various disorders of the central nervous system. Similarly, to support the work of clinicians, validated reference configurations are provided for common use cases of each supported algorithmic adaptive mode [129].

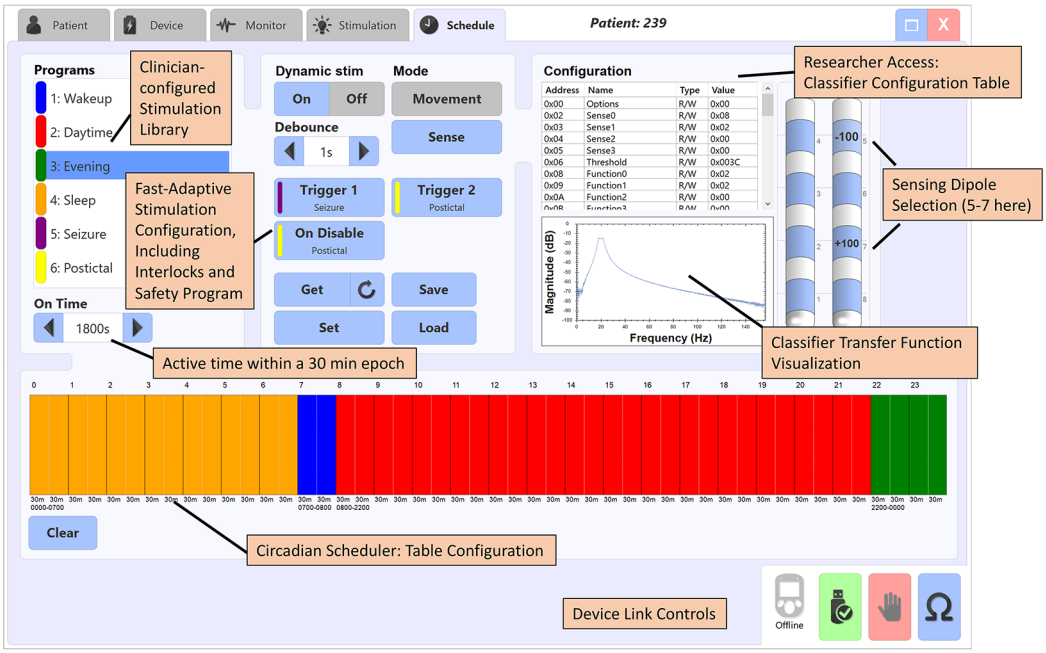


Figure 3.3. The Picostim-DyNeuMo clinician programming interface. The programming interface allows for defining up to eight stimulation programmes, as well as configuring their activation based on motion state, neural state sensing, and time of day in 30-minute blocks. Reproduced from [126] under CC BY 4.0.

Circadian adaptation

Taking advantage of circadian patterns with simple feed-forward models, even as simple as a circadian scheduler, could offer an effective first-order approximation to more complex neural marker–based adaptive treatment in Parkinson's disease [186] and epilepsy [125].

The DyNeuMo offers a simple circadian (24-hour) scheduler to adapt stimulation settings with a 30-minute granularity [126], accessible in the clinician programmer interface in Fig. 3.3. Chronic use of the scheduler system reached pilot validation in a case study involving an epileptic canine patient with highly periodic seizure patterns [127]. Other uses of more granular daytime adaptations might include the tapering of stimulation intensity in patients taking timed medication in addition to their DBS therapy.

Activity-based adaptation

While broader clinical applications are emerging, at the time of writing DBS is most established as a therapy option for movement disorders, such as Parkinson's disease, dystonia and essential tremors. As a short-acting pathway to therapy refinement, motion adaptation offers a simple way to make patient experience with DBS therapy feel more seamless in everyday life.

The DyNeuMo makes use of an integrated accelerometer sensor and digital motion processor in the Picostim IPG that allow changing between eight predefined stimulation programmes tailored to various types of activity. The DyNeuMo supports adaptations to posture (e.g. sitting or standing) and motion state (e.g. still or walking), as well as options for emergency stimulation modes in cases of sudden shocks and falls [129]. A somewhat non-conventional use of the accelerometer is detecting a tap at the implantation site, which could offer patients a discrete way to interact with their devices without the handheld remote, and allow carers a quick and easy way to activate emergency stimulation programmes [127].

The motion-adaptive mode reached the stage of first-in-human validation in a dystonia patient under a humanitarian exemption granted by the Medicines and Healthcare products Regulatory Agency. The case study showed promising outcomes and is reported in [129].

Neural recording

The Picostim devices are equipped to record LFP-type neural activity from the implanted DBS electrodes. This capability was first validated *in vivo* in the case study of [127]. The sensing capabilities proved suitable for capturing seizure episodes and distinguishing common behavioural states (see Fig. 3.4). Streaming out neural activity from chronically implanted patients is exciting in its own right as a tool for longitudinal studies of disease [186].

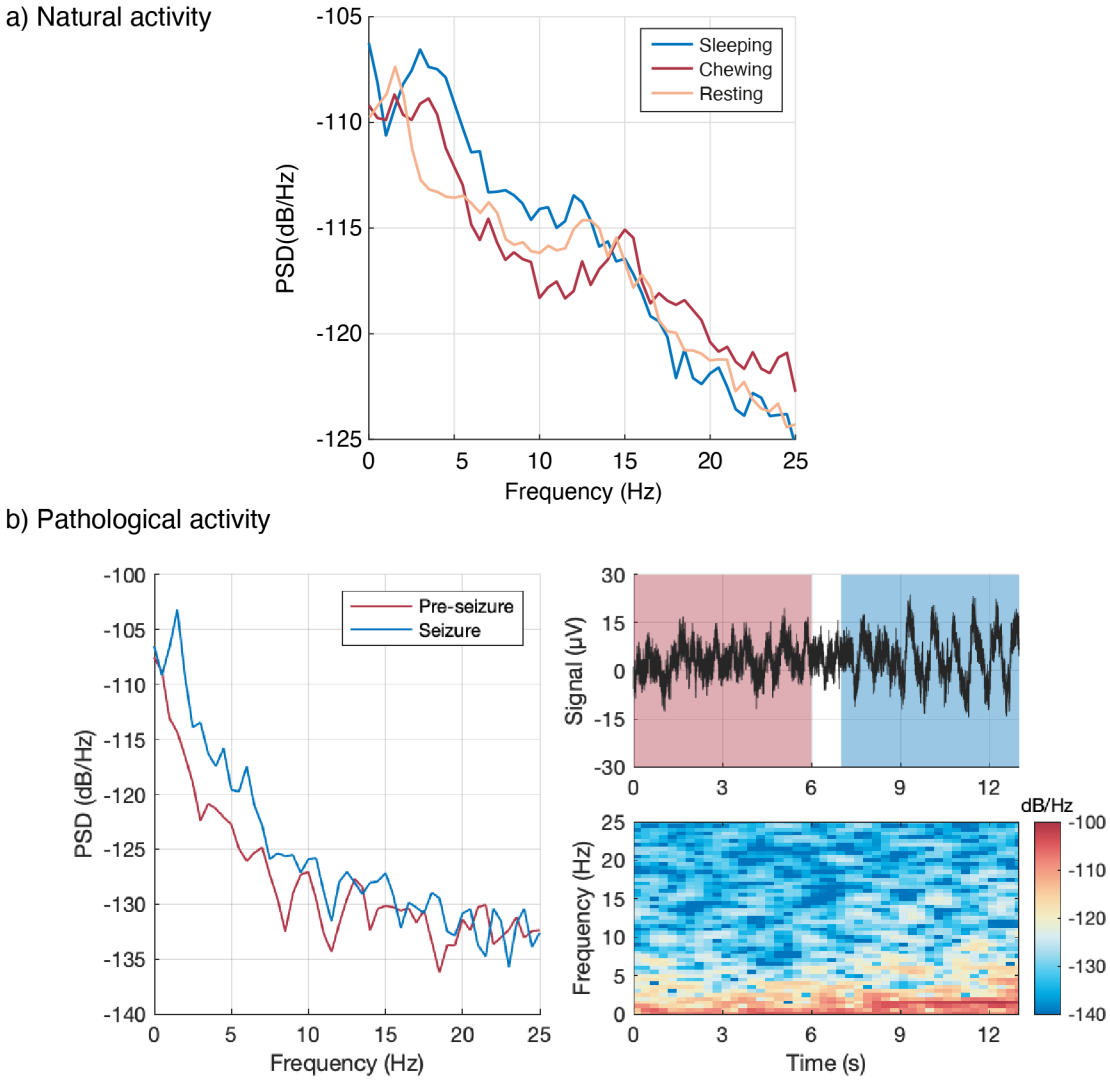


Figure 3.4. Chronic validation of LFP sensing in the DyNeuMo. Signals were recorded from the centromedian thalamic nucleus of a canine with naturally occurring epilepsy, approximately three months after implantation, in the presence of a 0.5 mA baseline stimulation programme. Parts of this data were previously reported in [127].

3.1.2 Research need

Research into Parkinson's disease has provided significant evidence for the effectiveness of adaptive DBS therapy (aDBS), shown to provide at least equivalent, if not better therapeutic outcomes as conventional DBS [42,187]. Research into adaptive aDBS-style as a responsive therapy in epilepsy is ongoing, with promising results [120].

While an obvious candidate in the treatment of nervous system disorders, neural-state adaptation has so far faced considerable challenges preventing clinical adoption. For one, neural signals tend to be small, making them challenging targets for low-power implantable instrumentation, constrained significantly in signal-to-noise ratio compared to external research equipment. The presence of stimulation further complicates reliable measurements [135, 188]. Despite the challenges, electrical biomarker-based adaptive DBS therapy was shown to be feasible [42,138,170] using experimental electronics. This leads us to the second major limitation.

Given the severe constraints of energy consumption in IPGs, current devices utilise low-power processing units to conserve energy (see Section 1.5.4). This trade-off severely limits the computational capabilities of implantables, holding up clinical research. To illustrate this point, the first successful demonstration of adaptive neurostimulation in Parkinson's disease has been published a decade ago as of the time of writing, yet no clinically approved implantable offers real-time DBS control for Parkinson's, and no study has been published on the chronic validity of this approach.

A common strategy in low-power wearable devices, off-loading signal processing tasks to an external device is unfortunately not viable in IPG systems at present. Recording-capable implantables are for now severely limited in operating time (order of hours) by the high power consumption of their radio links [142,143].

As demonstrated by preliminary results in the previous section (Fig. 3.4), the Picostim-DyNeuMo is capable of maintaining the fidelity of neural signals recorded in the presence of stimulation. Consulting with the manufacturer, the IPGs had a small amount of on-board processing power to spare. Overall, this made the platform an attractive target for attempting the implementation of an aDBS controller algorithm.

3.1.3 Aims

This chapter focuses on the following deliverables for advancing the Picostim-DyNeuMo platform towards clinically available responsive neurostimulation:

- Expand device capabilities for general purpose power-adaptive closed-loop control
- Validate algorithm performance in a tissue-electrode interface model environment
- Design a setup supporting externalised use of the IPG in a non-human primate model
- Pilot the algorithm *in vivo*

3.2 Technical methods

3.2.1 Signal Chain

The DyNeuMo utilises a capacitively coupled differential amplifier configuration for sensing neural signals across DBS electrode pairs. Previous research into neural amplifiers demonstrated that capacitive coupled differential circuits, commonly used in both experimental [137] and clinical systems [138, 139], benefit from greater signal-to-noise ratio (SNR) compared to other topologies in similar application [188].

A conventional analogue neural recording front-end [188], along with the electrode channel switch matrix can be operated as a sample-and-hold circuit to suppress stimulation artefacts. Synchronising the timing of measurements and stimulation pulses at a constant latency allows for disconnecting the sensing circuit from the electrodes during high potential excursions. The practical limit on the sampling rate at a particular level of measurement accuracy is going to be determined on the first order by the sample-and-hold capacitance. The DyNeuMo offers a sampling rate of 625 Hz while retaining a low measurement noise floor in the presence of 125 Hz DBS stimulation [126].

3.2.2 Algorithm outline

A seminal paper by Little et al. in 2013 [42] set the direction of research into adaptive brain stimulation for Parkinson's disease for the decade to follow [131, 138]. The neural state marker forming the basis of adaptation was beta-band power in the local field potential in the region targeted by stimulation, the subthalamic nucleus (STN).

Although a multitude of techniques exists for analysing signal power in a certain frequency band, all are subject to a common limitation: the time-frequency uncertainty principle. Readers unfamiliar with the concept may wish to consult the Appendix. The study by Little et al. used a simple filter-based algorithm to track and respond to changes in the oscillatory power. Due to the uncertainty principle, the more commonly seen signal transform–based approaches would offer little benefit in terms of responsiveness or band selectivity at the expense of significantly greater computational cost. An outline of this filter chain algorithm, as well as a demonstration of its step-by-step operation, are shown in Fig. 3.5. Processing chains of a similar structure have been explored for the responsive treatment of epilepsy, aimed at detecting seizure episodes based on their difference in spectral characteristics from normal neural activity [119]. To construct an automated oscillation detector, broadly four processing steps are involved.

The algorithm first uses a narrow band-pass filter to isolate the beta-band oscillation from the broad-band neural activity, allowing an estimate of oscillatory power to be formed by observing the signal magnitude. This approach is well suited for tracking oscillatory neural oscillations whose dominant frequencies are stable over time. In the case of Parkinson's disease, significant patient-to-patient variation in the target frequency was observed, however within each patient this activity was found to be stable over the long term. As human observers we can visualise the signal envelope of an isolated oscillating signal (a trace connecting the peaks on each side) as an estimate of the instantaneous magnitude of activity, an automated process to acquire an envelope is slightly more involved.

The signal envelope can be considered as a low-frequency analogue modulation of the oscillatory activity, implying we should be able to obtain it through low-pass filtering. Importantly though, since an isolated oscillatory signal shows symmetry about the zero amplitude level, its longer timescale averages will tend to zero. To make estimating the magnitude possible, a non-linear operation is necessary to break the symmetry, options include truncating the negative part of the signal, taking the absolute value, or the square of the signal. This operation is referred to as rectification in the algorithm. Following rectification, the signal can be low-pass filtered to obtain an estimate of the envelope. In practice, the absolute value operation is the most commonly used, being less computationally complex than squaring the signal, while truncation leaves the signal 'sparse', resulting in a time average significantly smaller than the true envelope. Note that even with absolute value rectification, the resulting envelope estimate will be strictly smaller than, but proportional to the true envelope.

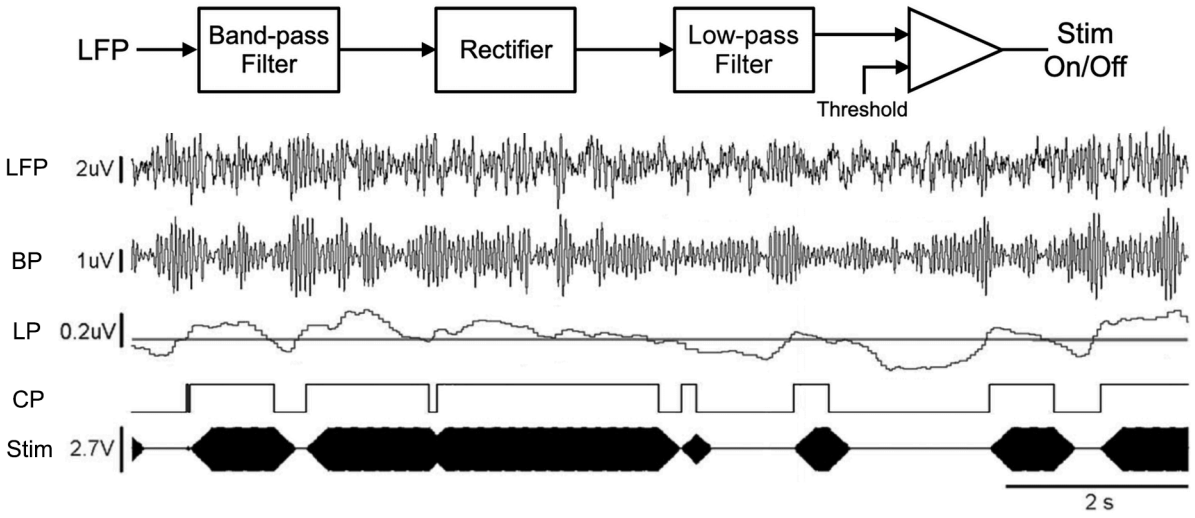


Figure 3.5. The aDBS algorithm architecture (top). The lowers panel shows the signal transformation performed by each stage. Lower panels adapted from [42] under CC BY 3.0.

3.3 Development

The Picostim-DyNeuMo IPGs are equipped with a 16-bit microcontroller unit (MCU) of the MSP430 family (Texas Instruments Inc., Dallas, TX, USA), a standard choice for ultra-low-power medical devices [165]. The microprocessor has a 32-bit hardware multiplier module available for integer arithmetic, and no floating-point unit. The microprocessor's limited capabilities by today's computational standards set up the constraints for implementing precision real-time control algorithms. The following sections describe the optimisations and design trade-offs taken to implement each of the signal processing modules of Fig. 3.5, to form an efficient processing chain for real-time band-power estimation in our IPGs.

3.3.1 General purpose filter block

Digital filters are primarily categorised as having either finite or infinite impulse responses (FIR, IIR). Digital filters operate on the basis of a delay line, keeping a number of past signal samples in memory, each having a weighted contribution (coefficient) to the filter output. The weighted samples of the delay line are referred to as 'taps'.

In the case of FIR filters, the filter output at any time only depends on the values of the input signal (Fig. 3.6a). This structure means the effect of a single sample impulse can affect the output at most as long as the number of taps, hence the name of the filter family. As a common algorithm for 'data smoothing', the moving average operation can be thought of as an FIR low-pass filter, with samples in the moving window being uniformly weighted taps.

In contrast, in the case of IIR filters, the output depends not only on the input signal and its history but also on past values of the output itself via a feedback path (Fig. 3.6b). IIR filters tend to require fewer delay units (and coefficients) to achieve similar performance to FIR filters, however, the presence of the feedback path introduces the potential for instability. To conserve memory in the microprocessor, we settled on using an IIR architecture to implement a customisable filter module.

Based on the grouping of commutative operations, it is possible to implement the same IIR filter in essentially four different topologies, referred to as direct- and transposed direct-forms of either type I or II, each with its own benefits and drawbacks [189]. The topology used here is the so-called direct-form I (DF-I), shown in Fig. 3.6b, for its unique property among the forms of being relatively resistant to internal overflows when implemented with two's complement fixed-point arithmetic. As a trade-off, the DF-I topology uses more internal memory than indirect forms and is more sensitive to coefficient precision. The latter can be largely mitigated by limiting the order of IIR blocks to second order. Higher order filters can then be formed from second-order sections (SOS) [189]. For the purposes of a sensitive medical device, reducing the likelihood of overflow was deemed a preferable trade-off.

3.3.2 Offset removal high-pass filter

Quantised band-pass filters are not guaranteed to have ideal DC (0 Hz) rejection. This can be problematic in our use case, as both fixed and slow-varying (sub-0.1 Hz) offsets are present in neural recordings, the former due to modest impedance mismatches in the measurement chain (manufacturing variability), the latter due to fluctuations in the electrode micro-environment. To guarantee the elimination of these effects, a specialised processing block was implemented dedicated to offset removal, as an additional filter stage preceding the canonical band-power estimation chain.

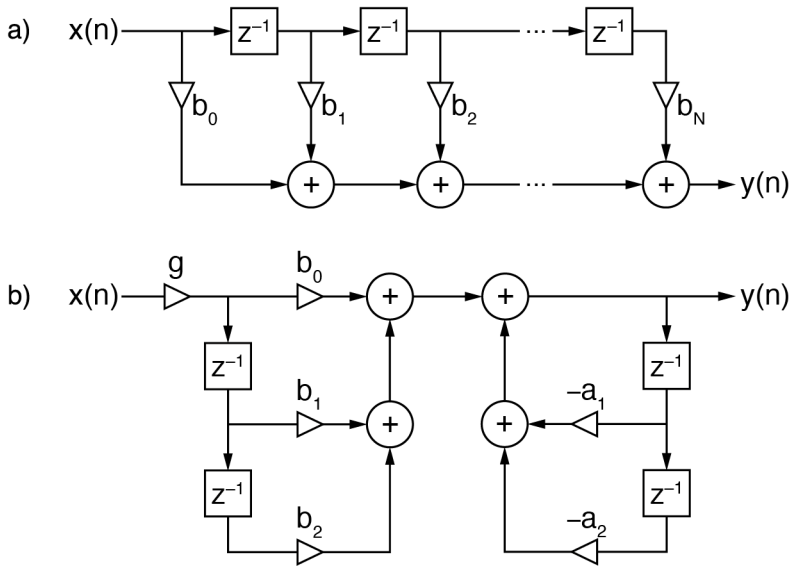


Figure 3.6. Digital filter architectures. **a)** Finite impulse response filter. **a)** Infinite impulse response filter in direct-form I arrangement. The z^{-1} blocks represent time delays of one sample each, triangles denote multiplicative coefficients, (+) blocks are summing nodes between two signal paths.

A family of first-order digital high-pass filters can be described as follows, for input at timestep n, x_n, and output y_n, where a is a positive coefficient in the range $0 \leq a < 1$.

$$y_n \leftarrow a \times (y_{n-1} + x_n - x_{n-1}) \tag{3.3.1}$$

While the structure could be expressed using the general purpose second order IIR function, a significantly more efficient implementation is available. Constraining the coefficient a as

$$a = \frac{2^G - 1}{2^G}$$

where G is a natural number, provides a series of values asymptotically approaching unity. The nominator and denominator of the fraction can be distributed between the computation steps as seen below, in the form of binary shifts. This arrangement entirely precludes stability errors, given a sufficient memory size for the internal states (e.g. 32 bits for y', z, for 16-bit x, y).

$$\begin{aligned} z_n &\leftarrow y'_{n-1} + (x_n - x_{n-1}) \ll G \\ y'_n &\leftarrow z_n - (z_n + 2^{G-1} \gg G) \\ y_n &\leftarrow (y'_n + 2^{G-1}) \gg G \end{aligned} \tag{3.3.2}$$

The binary scale was found to be a convenient trade-off between flexibility and simplicity, as it eliminated the need for fixed-point coefficients while still allowing a good range of cut-off frequency options. This allows users to use the offset filter for more than its strictly intended purpose of DC-offset removal, and boost the low-frequency rejection of the band-pass stage more broadly. This can be an especially desirable feature for certain subtle spectral targets, that do not rise significantly above the broader $1/f$ spectral density background expected of neural activity, or in the case of significant low-frequency artefact contamination of the signal [135]. The filter characteristics, as well as the available cut-off frequencies, are shown in Fig. 3.7.

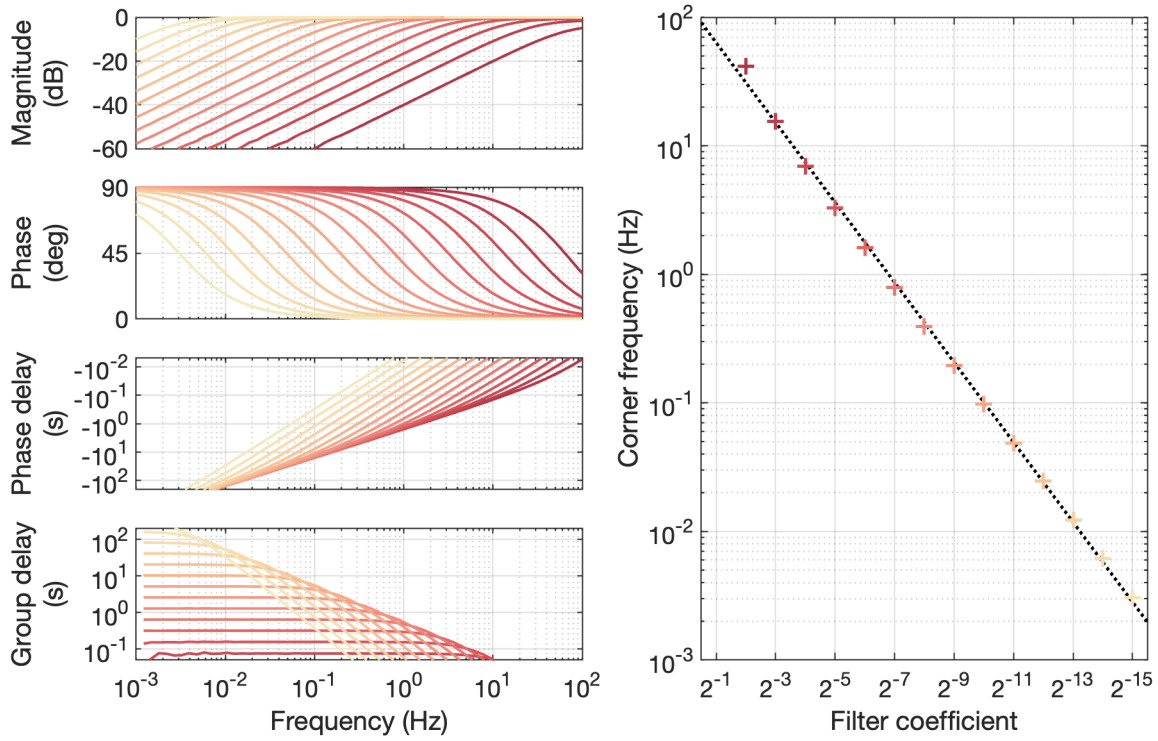


Figure 3.7. Steady state frequency response of the offset filter. The characteristics are presented for the range of filter coefficients $a = 2^{-1} - 2^{-15}$. Corner frequencies were measured at -3 dB. Input signals were synthesised at a sampling rate of 625 Hz.

3.3.3 Low-pass filter

Although the original implementation of the band-power estimation algorithm used a windowed moving average function as the low-pass stage [42], we opted to depart from this choice in our resource-constrained setting.

The classical moving average filter, as mentioned previously, is essentially a special case of FIR filters, formed of uniformly weighted taps over the length of the window. Though the FIR-nature imparts stability, a critical consideration for digital filters, the cut-off frequency can only

be adjusted by changing the window size. This is not ideal for an envelope demodulation block where sub-Hz cut-offs are required, as this translates to $N > f_s$ taps. To reduce the amount of memory required, we can once again look for an IIR-based solution.

The exponential moving average provides one such simple, low-resource alternative. The filter equation is provided as:

$$y_n \leftarrow y_{n-1} + \frac{x_n - y_{n-1}}{w} \tag{3.3.3}$$

or alternatively:

$$y_n \leftarrow \frac{1}{w} x_n + (1 - \frac{1}{w}) y_{n-1} \tag{3.3.4}$$

where w is a parameter analogous to the window size in a classical moving average. In fact, an equivalent window size may be defined based on the average age of data in the filter memory [190], as:

$$w_{eq} = \frac{N+1}{2} \tag{3.3.5}$$

where N is the window size in samples of a classical moving average filter.

One might also observe that the second-order IIR block that we already decided to implement in Section 3.3.1 can be configured to calculate an exponential moving average through a specific choice of coefficients.

$$\begin{gathered} y_n \leftarrow g\, b_0\, x_n + g\, b_1\, x_{n-1} + g\, b_2\, x_{n-2} + a_1 y_{n-1} + a_2 y_{n-2} \\ \begin{aligned} g &= 1 & b_0 &= 1/w \\ a_1 &= 1 - 1/w & b_1 &= 0 \\ a_2 &= 0 & b_2 &= 0 \end{aligned} \end{gathered} \tag{3.3.6}$$

However, as w tends to be large, $a_1 \sim 1$ and $b_0 \sim 0$, which creates issues when implementing the algorithm using fixed point arithmetic.

A custom fixed-point implementation based on Eq. 3.3.3 can provide an efficient form also resistant to numeric errors, so long as we are satisfied with restricting w to integer powers of 2, i.e. $w = 2^{w_{shift}}$, conversely $w_{shift} = ceil(log_2(w))$. In this constrained case, we are able to exploit the equivalence of multiplication and binary shifts to dynamically scale our representation such that our internal products are always represented with sufficient bit-resolution: the scaling factor

is w_{shift} itself. This leads to the following realisation:

$$\begin{aligned} x'_n &\leftarrow x_n \ll w_{shift} \\ z_n &\leftarrow z_{n-1} + ((x'_n - z_{n-1}) \gg w_{shift}) \\ y_n &\leftarrow z_n \gg w_{shift} \end{aligned} \tag{3.3.7}$$

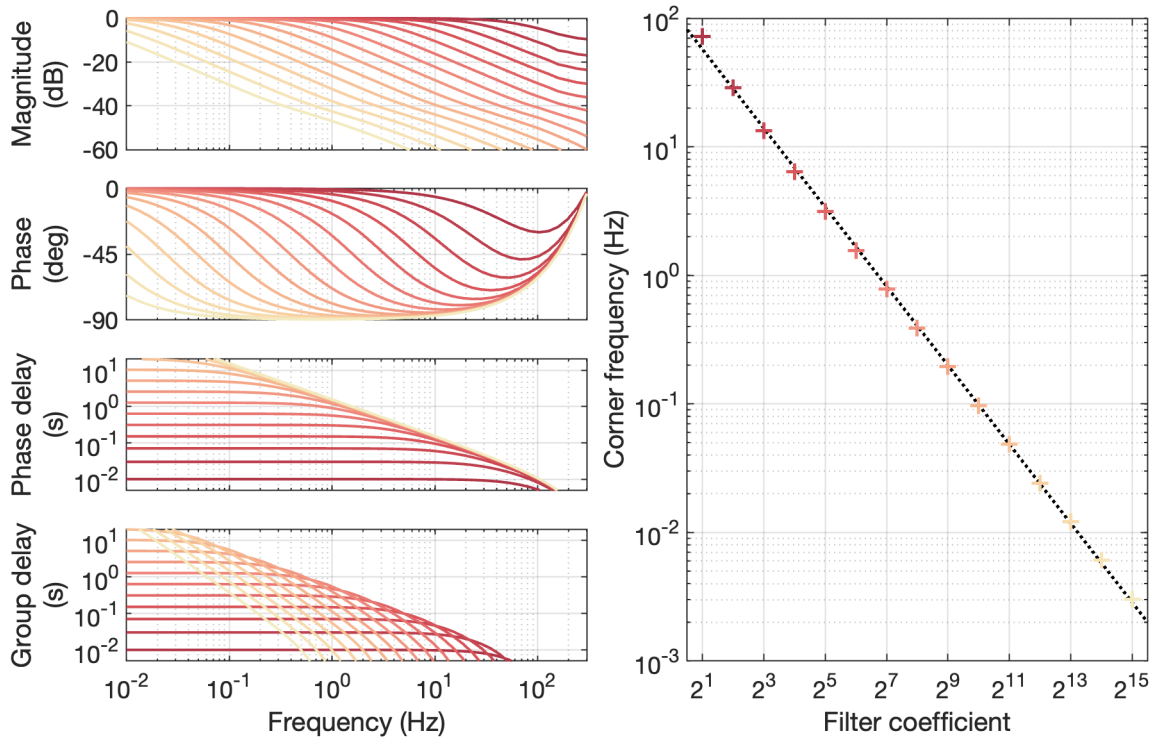


Figure 3.8. Steady state frequency response of the exponential moving average filter. The characteristics are presented for the range of filter coefficients $\lambda = w^1 - 2^{15}$. Corner frequencies were measured at $-3\,\mathrm{dB}$. Input signals were synthesised at a sampling rate of 625 Hz.

The characteristics of the moving average filter implemented according to this scheme are shown in Fig. 3.8. Since the exponential moving average assigns more weight to more recent samples, it will responds more quickly to changes in the signal compared to the classical moving average, something that becomes more pronounced at larger equivalent window sizes (Fig. 3.9). It should be noted that the exponential moving average is a first order filter, and will thus provide less stop-band attenuation than a SOS-based low-pass filter.

Reduced suppression of signals beyond the cut-off frequency allows some of the oscillation signal to pass through the filter, which adds ripple to the envelope estimate (see coloured traces in Fig. 3.9b). This effect is more significant the lower the value of w, establishing a trade-off between the time delay of the filter and the magnitude of ripple in the extracted envelope. In the future, the single-level threshold used in the detector could be replaced by a double-threshold

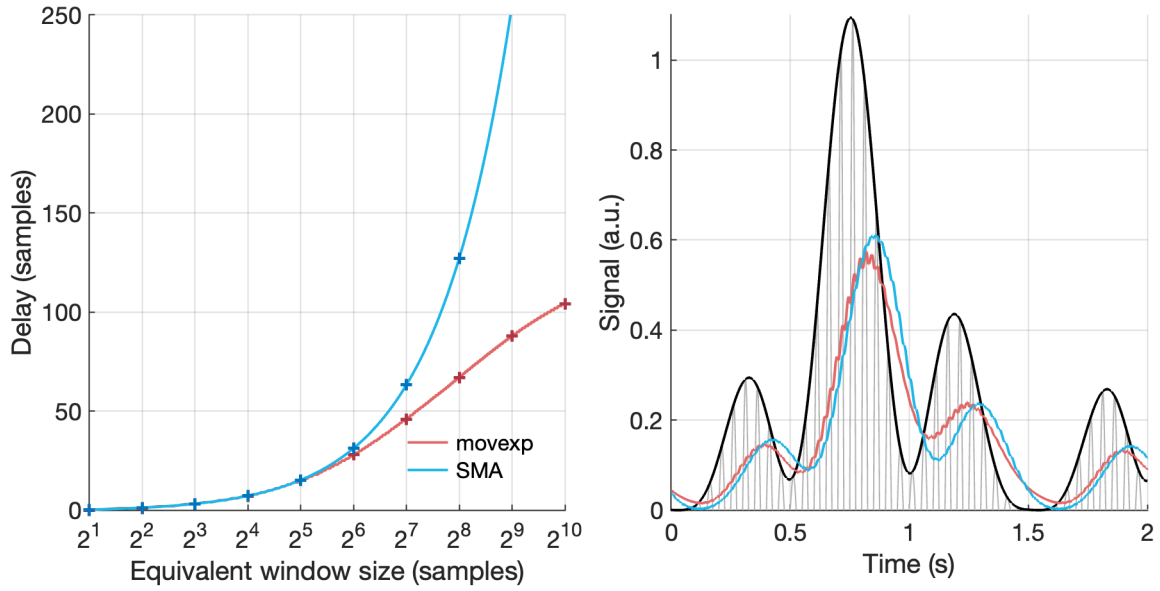


Figure 3.9. Performance characteristics of the exponential moving average filter. **a)** Time delay comparison of the classical and exponential moving average filters at a given equivalent window size. The filter delays were measured as the lag at peak cross-correlation relative to an offline-calculated Hilbert envelope. The test signal consisted of a 20 Hz base tone with arbitrary analogue modulation, sampled at 625 Hz. **b)** Time-domain characteristics of the classical and exponential moving averages at equal delays ($w_{shift} = 6$, $N = 127$ samples), on a segment of the same test signal.

comparator with hysteresis, where the two threshold levels could be set proportional to the peak-to-peak ripple magnitude, allowing to bias the detector toward faster response times while preventing unnecessary rapid state changes.

3.3.4 Processing chain

The implemented processing functions are enumerated at compile time and can be dynamically called. An 8-byte array stores the module call order, i.e. classifiers of up to eight stages can be defined arbitrarily as desired. The core of the algorithm is a looping function that sequentially calls the processing functions defined in the array, passing on the result of the preceding stage. For a standardised low-overhead pipeline, all function return values were restricted to 16 bits. 80 bytes of memory were allocated to be available for the algorithm that can be partitioned dynamically for the various available modules as needed. This amount of memory is sufficient for four IIR stages; or three IIR blocks along with five offset filters, moving average or absolute value stages. Note that the absolute value function does not require configuration memory, only an entry in the call sequence. We expect most applications will use an initial offset filter, a two or three-stage (4th or 6th order) IIR band-pass filter, an absolute value function, and one or two moving average blocks depending on the low-pass order needed.

Table 3.2 shows the total computational cost of various realisations of the filter chain. The results support replacing certain IIR blocks with custom filter types, as the modest increase in code complexity is offset by very significant improvements in execution time.

Table 3.1. Example filter chain configurations for three neural biomarkers.

Target	High-pass		Band-pass		Rectifier	Low-pass	
	Order	Cut-off	Order	Pass-band		Order	Cut-off
Cortical alpha	1	1.6 Hz	4	5–15 Hz	+	2	0.6 Hz
Parkinsonian beta	1	6.9 Hz	4	18–22 Hz	+	2	1.0 Hz
Thalamic seizures	1	1.6 Hz	4	5–24 Hz	+	2	0.6 Hz

Table 3.2. Computational cost of various filter chains. Compiled for MSP430FR5739 MCU in IAR Workbench 7.20, with small code and data models, using low compiler optimisation.

Filter chain configuration	Code (bytes)	Execution time (cycles)	Data (bytes)
Original: IIR BP[2], abs(), IIR LP[1]	518	2325	121
IIR HP[1], IIR BP[3], abs(), IIR LP[1]	0	+1440	+90
IIR HP[1], IIR BP[3], abs(), movexp()	+100	+745	+53
IIR HP[1], IIR BP[2], abs(), IIR LP[1]	0	+715	+48
IIR HP[1], IIR BP[2], abs(), movexp()	+50	+40	+11
Offset HP, IIR BP[3], abs(), IIR LP[1]	+98	+895	+57
Offset HP, IIR BP[3], abs(), movexp()	+148	+220	+20
Offset HP, IIR BP[2], abs(), IIR LP[1]	+98	+175	+15
Offset HP, IIR BP[2], abs(), movexp()	+148	–505	–22

3.4 External validation

3.4.1 Validation setup

A systematic and reproducible external validation method was critical for accurately characterising the operation of the DyNeuMo, to support the case for clinical trials.

To provide a realistic *ex vivo* environment that captures many of the complexities of measurements across the tissue-electrode interface, we opted to carry out our experiments inside a standard electrophoresis tank, filled with physiological saline solution (0.9 % w/v) [191]. The tank's built-in electrodes were used to inject test signals into the solution via a purpose-built current-mode signal generator device. The DyNeuMo's standard DBS leads were submerged inside the saline solution, along with a reference wire connecting to the IPG case. Similar setups are widely used in externalised research into implantables [188].

A commercial TMSi multi-channel ExG recording system (Porti7-24e4b4at; Twente Medical Systems International B.V., Ej Oldenzaal, Netherlands) was used as a reference, simultaneously recording the injected test signals, and the artefacts generated inside the saline tank by the DyNeuMo's stimulation response.

3.4.2 Tests

Critically, validation began by measuring the noise floor of the DyNeuMo, to confirm that neural signal sensing will not suffer materially from the presence of stimulation artefacts. Testing established a measurement noise floor of 190 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ over the 10–40 Hz range in the presence of biphasic stimulation up to 4.0 mA, when operating at a sampling rate of 625 Hz using a sample-and-hold scheme, with synchronised stimulation at a rate of 125 Hz. This stimulation frequency is close to typical high-frequency rates used in epileptic and Parkinsonian DBS therapy [165]. The combination of 625 Hz recording was selected for providing a Nyquist rate of > 300 Hz, covering the entire LFP band.

Long-term validation tests on algorithmic performance were carried out for two proposed use cases. The protocols and results are presented in the following sections.

Special thanks to Mayela Zamora and Moaad Benjaber (Department of Engineering Science, University of Oxford) for carrying out the long-term validation tests, and their work on the formal validation protocol.

Beta-oscillation detection

While the nominal operation of the algorithmic classifier can be tested against qualitative expectations on pre-recorded data containing Parkinsonian beta activity, ground-truth information on the bursts is hard, if at all possible to establish – clinician labelling is not typically performed on this signal type; reference filter–based approaches will invariably suffer from an arbitrary choice regarding specific filter characteristics.

Since neither clinicians nor automatic processing establishes ground truth to perform validation against, we settled on characterising detection characteristics against an artificially generated test signal. The synthesised data consisted of a 20 Hz test signal modulated with a 200 mHz square wave at 80 % depth and 50 % duty cycle, i.e. switching the 20 Hz oscillation between 20 % and 100 % magnitude every 2.5 s. The DyNeuMo's stimulation response was recorded along with the test signal in the reference TMSi system. Calculating an envelope around the stimulation artefact and the synthesised beta bursts (Fig. 3.10) allowed for automatic validation of detector performance.

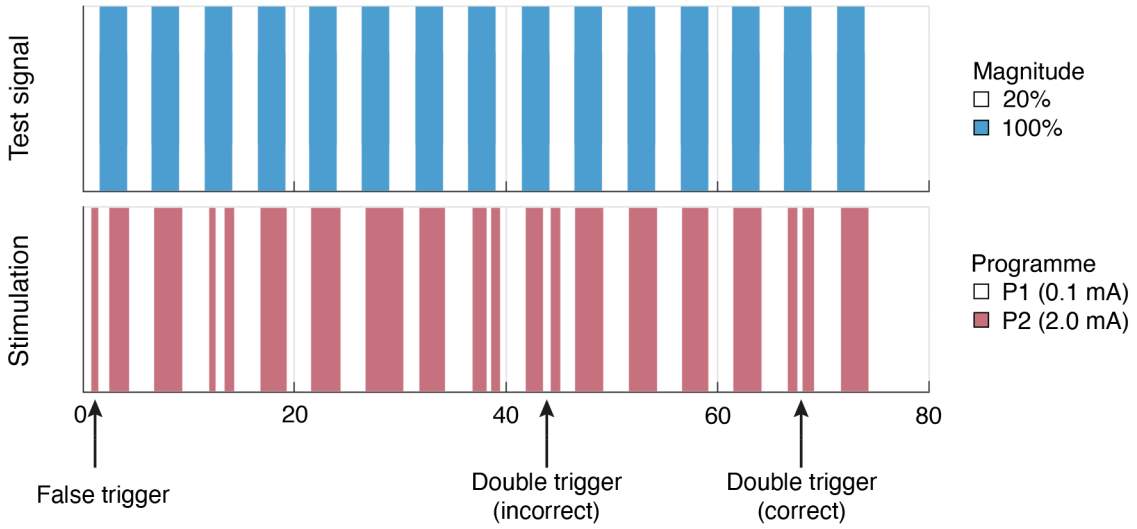


Figure 3.10. Error types in the DyNeuMo Parkinsonian beta-oscillation detector. The figure shows the stimulation state of the DyNeuMo (red) and the timing of test beta bursts (blue). Tests were carried out inside a saline tank, with test waveforms provided by a purpose-built signal generator. The test signal consisted of a 1.5 μV_{RMS}, 20 Hz test signal with added 'brown' noise, modulated by a 200 mHz square wave between 20 % and full magnitude. Rapid re-triggering was accepted as a single correct detection so long as they both covered a true beta burst. The DyNeuMo was configured with a 4th order Butterworth band-pass filter with a 22–24 Hz pass-band, and a 2nd order Butterworth low-pass filter with a corner frequency of 1 Hz, at a recording and processing rate of 625 Hz.

The DyNeuMo was found to be able to reliably detect a lowest beta burst magnitude of $1.5\,\mu V_{RMS}$. To establish performance in the small signal limit, instead of ROC-style (receiver operating characteristic) characterisation, the threshold was set to the lowest operational level. On a long-term 24-hour test, a best true positive rate of 98.9 % at a false positive rate of 12.0 % was established.

Seizure detection

Unlike beta-bursts, clinician-labelling is routinely carried out on electrophysiological recordings of epileptic patients, making ground-truth data easily available. A 10-hour data sample, recorded from the human centromedian thalamus containing 30 labelled seizure episodes was used for validating device performance. The measured detector response to a single representative event is shown in Fig. 3.12.

The evaluated classifier was tuned to detect elevated power over a broad theta-alpha-beta activity band (5–24 Hz). The full filter chain is specified in Table 3.1.

Based on the collected ROC performance of the detector (Fig. 3.11), the optimal operating point was selected using Youden's J-statistic [192]. At the Youden-threshold, the filter chain achieved a true positive rate of 82.7 % at a false positive rate of 9.7 % on the trial dataset.

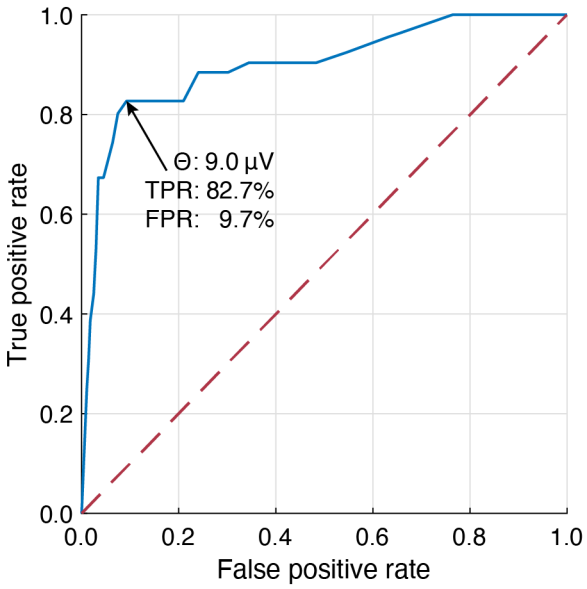


Figure 3.11. Receiver operating characteristic (ROC) of the DyNeuMo aDBS algorithm for seizure detection. The algorithm was evaluated by replaying a clinician-evaluated ground truth recording in a saline tank using a purpose-built signal generator. The test signal was recorded from the centromedian thalamic nucleus of a human patient over 10 hours, containing 30 labelled seizure episodes lasting on average 7.2 s. Youden's J-statistic [192] was used to identify a potential ROC optimum at a threshold level of $\Theta = 9.0\,\mu V$. Clinical data courtesy of Dr Antonio Valentín, King's College London.

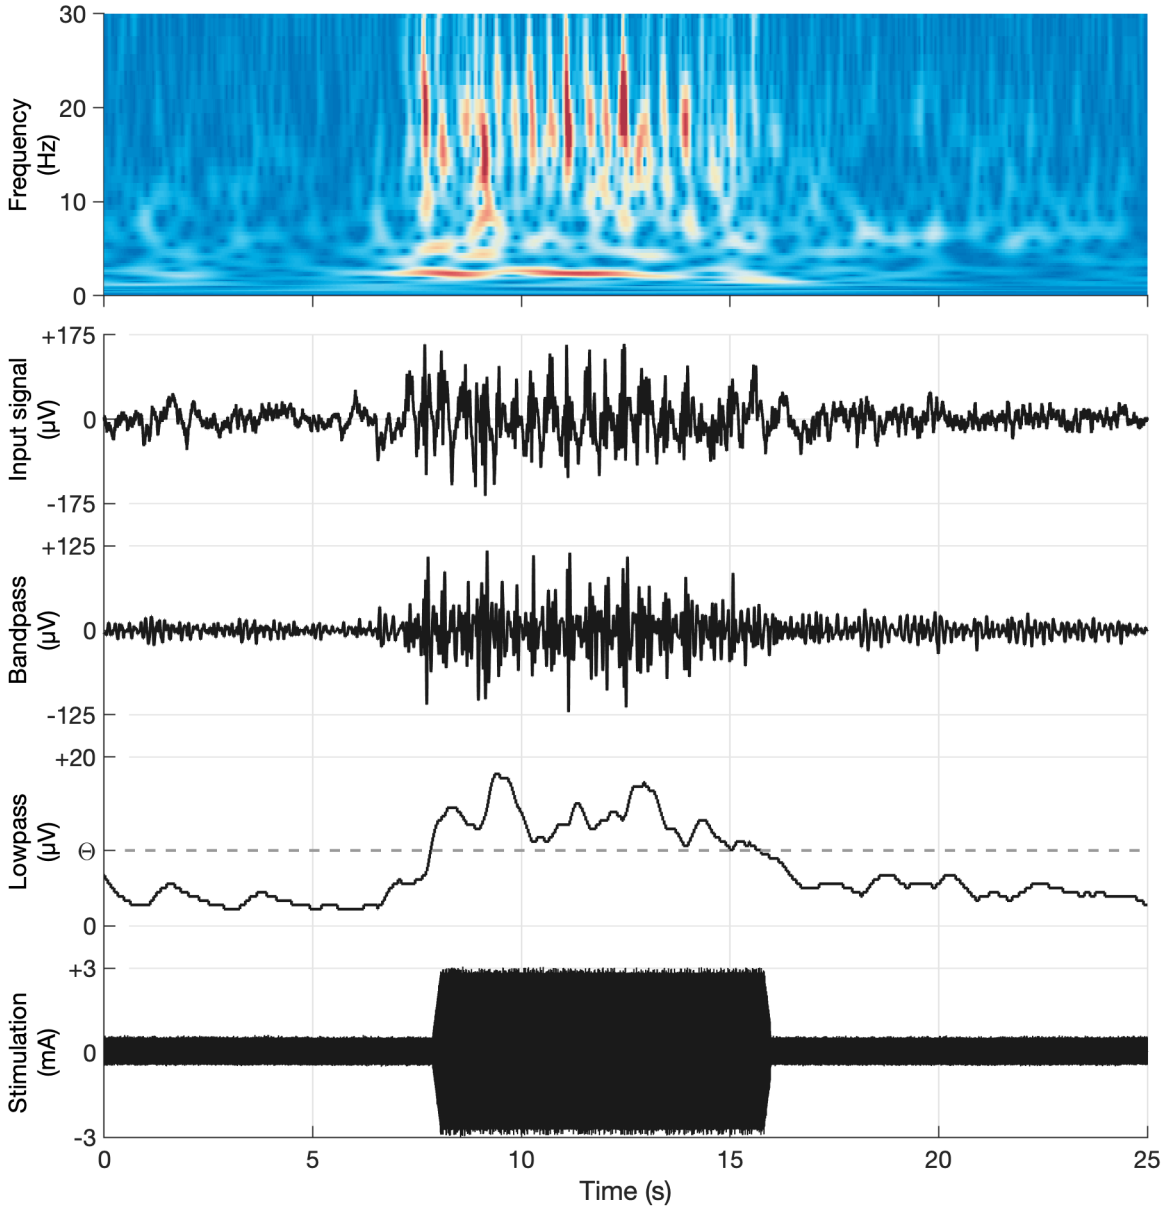


Figure 3.12. Seizure-response of the DyNeuMo in aDBS mode. Top: continuous wavelet transform spectrum of a single clinician-labelled seizure event, recorded from the left centromedian nucleus of the thalamus in a human patient. Lower panels: seizure event in the time domain, and key stages of the algorithmic response. The pre-recorded signal was replayed for evaluating the DyNeuMo aDBS algorithm in a saline tank by a purpose-built signal generator. The algorithm was configured with a 4th order Butterworth band-pass filter with a 5–24 Hz pass-band, and a 2nd order Butterworth low-pass stage at a corner frequency of 0.6 Hz (lowest achievable), at a recording and processing rate of 625 Hz. Stimulation was delivered in a bipolar-biphasic mode with symmetrical 90 µs anodic and cathodic pulses, with no inter-phase delay at a rate of 125 Hz. The detector threshold was set at 9 µV, triggering a stimulation state transition between a baseline of 0.5 mA and 3.0 mA intensities, with a ramp-rate of 1.25 mA/s Clinical data courtesy of Dr Antonio Valentín, King's College London.

3.5 Experimental validation

3.5.1 Experimental methods

Validation data was collected from a healthy rhesus macaque (Macaca mulatta), previously implanted and trained in a study of visual attention. All work involving the animal was carried out in strict accordance with the National Institutes of Health Guidelines for Use of Animals in Research, under a protocol reviewed by the Weill Cornell Medical College Institutional Animal Care and Use Committee. Details of the surgical planning, the complete implantation procedure and the original study are detailed in an arising publication [193].

Relevant for our validation effort, the animal was equipped with a set of bilateral DBS electrodes (Heraeus Medical Components, Hanau, Germany) with four independent rings of electrodes, where each ring consisted of three semi-annular platinum/iridium contacts. Post-operative imaging was carried out to confirm the positioning of the leads in the medial central thalamus, following prior work [194]. Two titanium grounding plates were placed to serve as electrode reference. An epidurally implanted linear 8-channel ECoG array (Fetz Spinal Cord 8; CorTec GmbH, Freiburg, Germany) over the supplementary motor and premotor regions of the left hemisphere was also available.

Adapting the DyNeuMo

To enable experiments with in the existing NHP setup, the IPG had to be adapted to the constraints of animal use, operating overnight in a freely behaving context. Further, it had to conform to the standard implantation procedure in place for the animal subjects. To allow for long-term data streaming, the IPG case was opened up and connected to a high-capacity 600 mA h single-cell rechargeable lithium-polymer type battery, reaching an operating time of 10 hours. Custom parts were designed and manufactured to house and protect the IPG, the external battery, as well as a USB Type-C charging circuit (Fig. 3.13b).

The implanted electrodes were not immediately compatible with the Picostim-DyNeuMo device in terms of form factor, and had a greater channel count (4 rings of 3 split segments, i.e. 12 channels/lead) than the IPG was designed for (4 channels/lead). The implanted leads were terminated in 12-contact Omnetics Micro 360 connectors (Omnetics Connector Corp., MN, USA), and affixed to the implanted headstage. A set of custom adapters were designed, to bridge regular Picostim 4-channel DBS leads with the implanted ones, shown in Fig. 3.13a. As platinum is neither easily soldered nor welded to, the adapters use spring-loaded contacts to

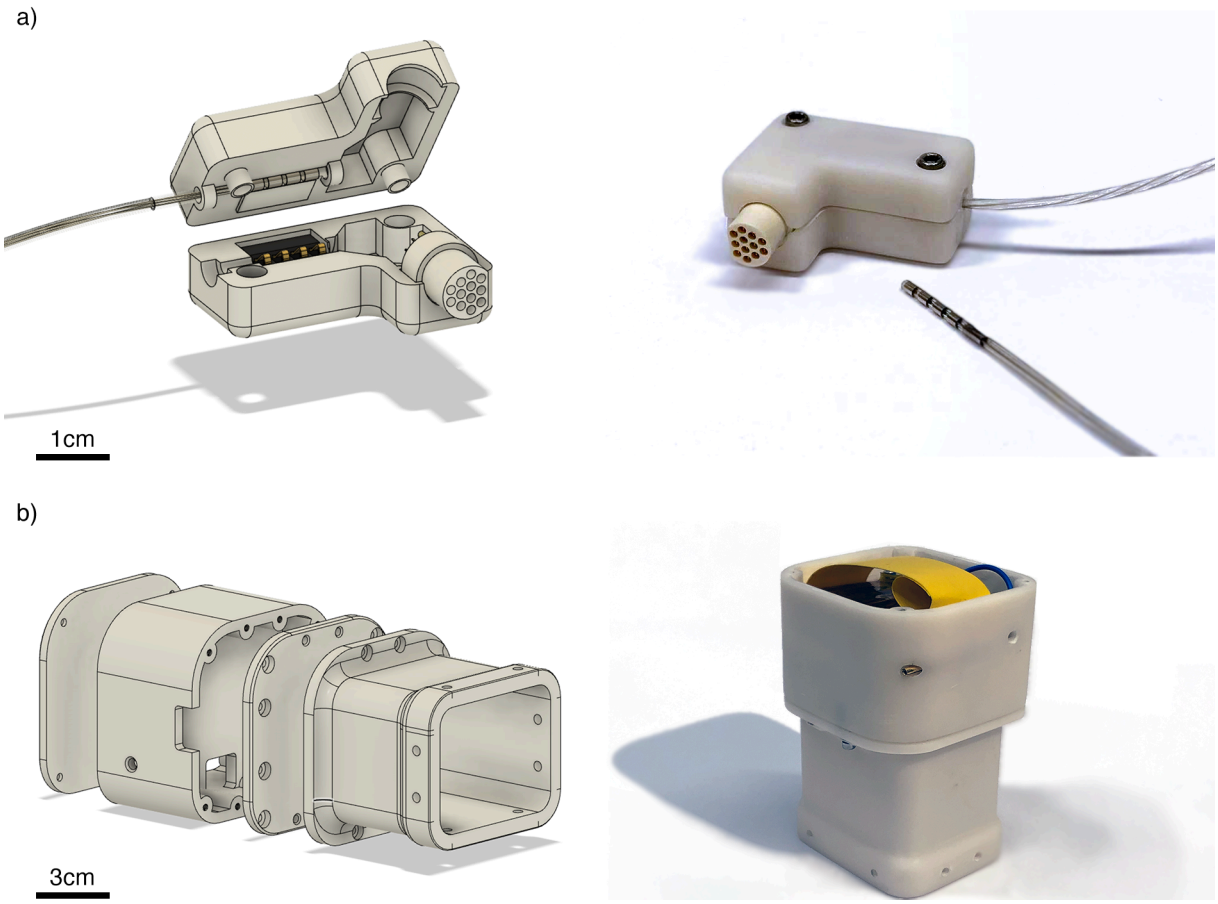


Figure 3.13. Design and realisation of the externalised IPG adapter parts. **a)** Adapter used for interfacing 4-channel externalised DBS leads with 12-channel implanted leads. A pair of guide rings helps align the external lead with a spring connector inside the adapter. **b)** IPG case assembly. A two-compartment case, holding the IPG and a rechargeable battery, was designed to fit atop a microdrive cap. Parts were manufactured via a stereolithographic process out of Formlabs 'Rigid 10K' resin and assembled using a standard ISO M1.6 stainless steel screw-set.

form a reliable mechanical contact with the electrode rings. Each IPG electrode is connected to the three segments of one of the electrode bands of the implanted lead via the adapter to provide matching channel counts.

Parts were modelled in Autodesk Fusion 360 (Autodesk Inc., San Rafael, CA, USA), then manufactured in-house using a standardised additive manufacturing toolset (Formlabs Inc., Somerville, MA, USA). The models were sliced in PreForm using default settings at a resolution of 50 µm, then manufactured in a Form 2 liquid resin stereolithographic printer using 'Rigid 10K v1' resin. The printed parts were first washed for 20 minutes in 98 % isopropyl alcohol in a Form Wash tank, allowed to air dry, and then cured under ultraviolet light at 70 °C for 60 minutes using a Form Cure unit. The curing process left parts with a powdery finish which was sanded away to avoid potential irritation on skin contact, as a precautionary measure.

To prevent damage to the IPG setup, a protective cap was formed and fitted from a sheet of Klarity-RG thermoplastic mesh (Klarity Medical Products Inc., Heath, OH, USA), established as safe and compatible with non-human primate (NHP) experimental use [195].

Electrophysiological protocol

Closed-loop recordings were carried out using the setup shown in Fig. 3.14. Using the custom parts described in the previous section (Fig. 3.13), a Picostim-DyNeuMo IPG was connected to DBS lead 2 for delivering stimulation. DBS lead 1 was recorded with an RZ2 data acquisition system (Tucker Davis Technologies, Alachua, FL, USA). As the IPG did not support data-streaming while running in closed-loop mode, the ECoG channels were split between the two devices in an interdigitated pattern to allow an approximation of the LFP seen by the IPG during offline analysis.

Eye tracking signals were recorded using Oculomatic Pro v1.9.7 (Neuro-Software Developers Inc., Jacksonville, FL, USA) [196]). The output of a vibration sensor was also used to track large movements within the behavioural setup, synchronised to the RZ2.

The externalised IPG was set to deliver current-mode stimulation to DBS lead 2, with two predefined stimulation programmes (P1, P2) that differed only in current magnitude. Stimulation was delivered in a bipolar biphasic pattern with equal 90 µs anodic and cathodic phases, and no interphase delay. The overall stimulation frequency was set to 125 Hz. The current magnitudes were 0.1 mA for pattern P1, and 2.0 mA for pattern P2. The closed-loop algorithm would be able to initiate transitions from the low-intensity baseline programme P1 into the high-intensity P2 mode. Transitions between the programmes were ramped at a rate of 0.01 mA/pulse, i.e. 1.25 mA/s to reduce paraesthesia. The stimulation state was locked for a duration of 3 s to prevent rapid switching. The contact impedances were typically in the range of 2–4 kΩ for each three-segment electrode ring.

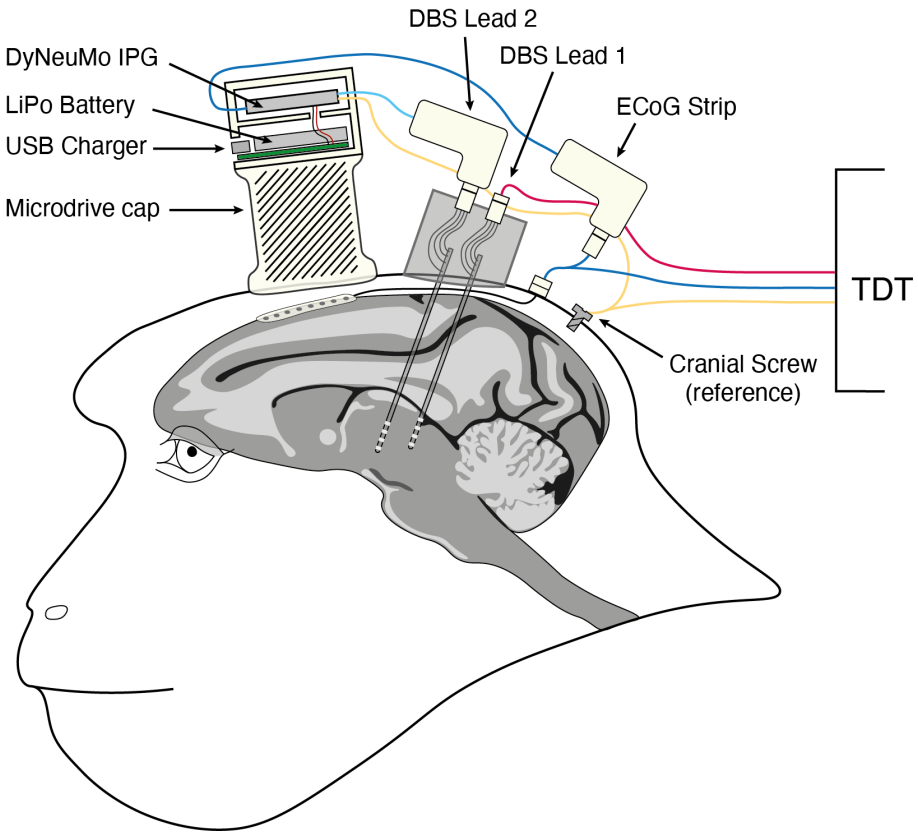


Figure 3.14. Externalised closed-loop research setup. Two 12-channel DBS leads were implanted in the medial central thalamus, along with an 8-channel epidural ECoG array over the left prefrontal cortex of a healthy NHP subject. To keep the setup compact and suited for freely moving observation, the DyNeuMo IPG was fitted securely atop an existing microelectrode drive, not in use during this study. One DBS lead and half of the ECoG channels were connected to the IPG, while the rest were connected to a Tucker-Davis Technologies (TDT) RZ2 recording system for reference measurements.

3.5.2 Results

Data collection

The Picostim-DyNeuMo IPG was first used to record a sample of ECoG activity in the absence of stimulation, which in conjunction with behavioural observation allowed for closed-loop parameter selection. A spectral sample of the 'training data' is presented in Fig. 3.15, showing periods of markedly increased alpha-band activity. Such activity was observed to occur during periods of low motivation, preceding prolonged eye closures. The algorithm was configured with the parameters described in Table 3.1 for alpha-band activity, and a threshold of 6.2 µV was experimentally obtained.

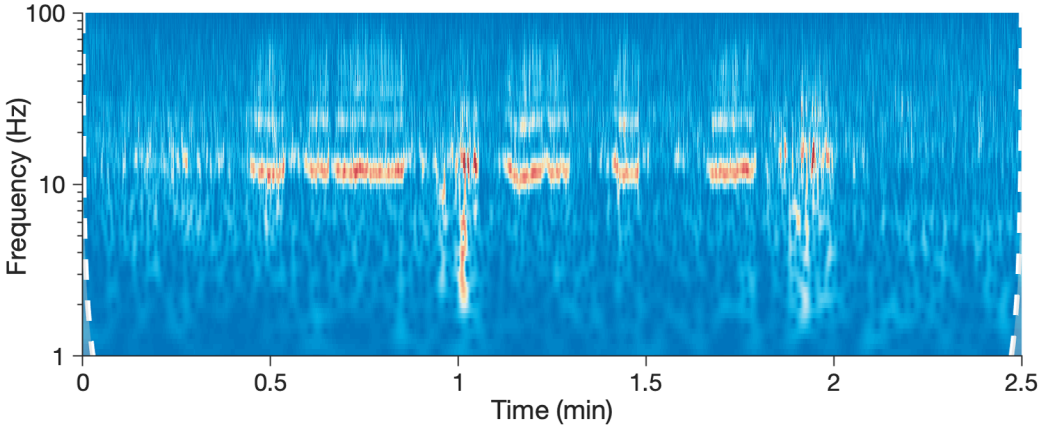


Figure 3.15. Alpha-band training data. Sample of unstimulated baseline cortical activity, showing prominent alpha-band features during inattentiveness, used for close-loop parameter selection.

Closed-loop operation

Prolonged periods of eye-closures were used as a quantifiable indicator of inattentiveness, as a proxy measure of low vigilance mental state. The selection of a cut-off threshold between natural blinks, and 'significant' lengths of eye closure was not immediately clear, therefore analysis was carried out in a parametric way. Stimulation transition times were stored in the IPG memory and could be downloaded following each session. The stimulation state could be validated by observing electrical artefacts present in the signal recorded from the non-stimulated DBS lead via the reference system. The outcome of a pilot trial is shown in Fig. 3.16.

As events were relatively short (order of seconds) and sparse (10–15 % of total recording time), we would not expect very high values of correlation given lag or jitter-type uncertainty in the system. To account for persistent lag, the peak of the cross-correlation of the signals was taken. To allow tolerance for jitter, the 'spiky' nature of the signals can be reduced by convolving

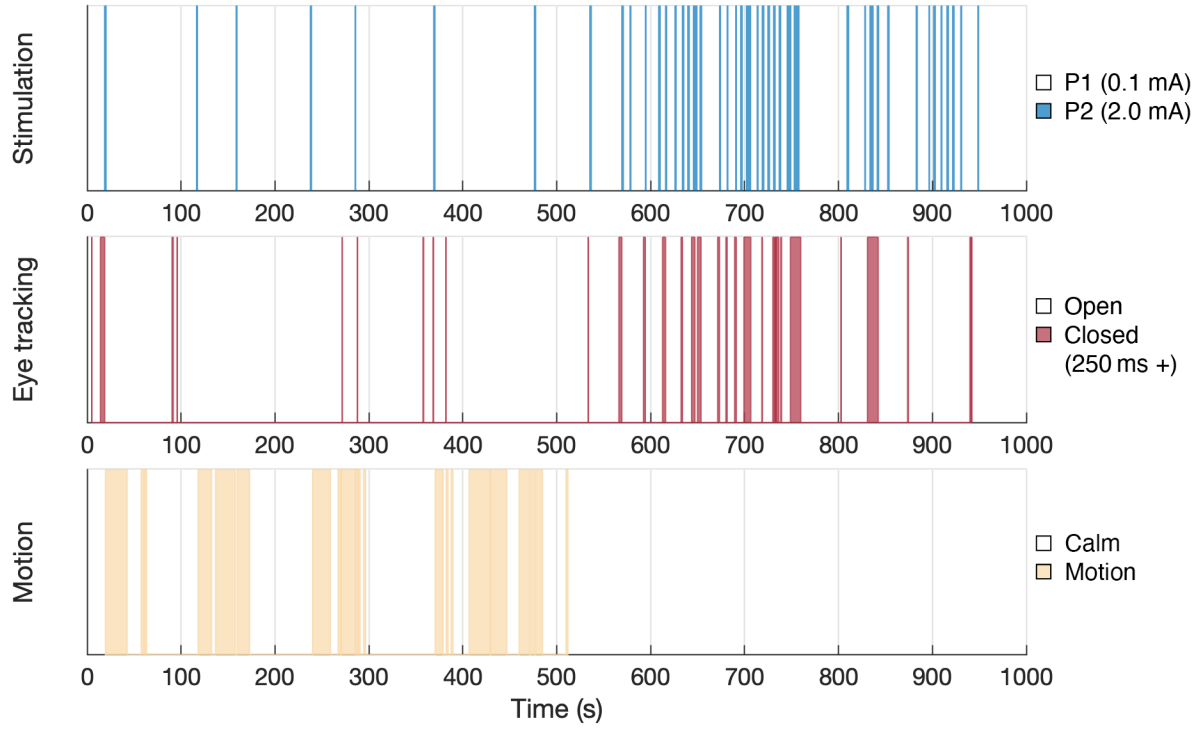


Figure 3.16. Results of non-human primate experiment. Stimulation state, eye-tracking and motion data are shown from a 16-minute closed-loop session. Loss of attention (eyes closed and lack of movement) shows a qualitative correlation with activations of the alpha-detector.

them with a Gaussian kernel. Fig. 3.17 shows the peak cross-correlation evaluated at varying degrees of smoothing and a broad range of eye-closure thresholds. For moderate smoothing ($\sigma = 1\,\text{s}$), and a threshold of 250 ms, the peak cross-correlation was found to be 0.63, indicating strong association. Notably, a step-change in correlation can be observed at a threshold of $\sim$3–4 s, indicating an approximate timescale for alpha-activity associated eye-closures.

Experimental limitations

A limitation of the Picostim-DyNeuMo is that data stream-out is not available during closed-loop operation, meaning for now only an indirect, qualitative confirmation of the expected operation can be claimed.

The question naturally arises whether the closed-loop system affects the network state in a meaningful way, activating from loss of attention, then causing the behavioural observation of the eyes reopening, or it is in fact the eye-state that modulates closed-loop activation. Restating in plain terms, whether the closed-loop device is simply acting as an expensive and invasive 'blink-detector'. In future sessions, the direction of causality could be supported through reference ECoG measurements (as proposed in Fig. 3.14), by establishing the exact time-ordering between alpha band activity, detector activation and prolonged eye closures. Analysing task performance in the alert (eyes open) intervals following a high-intensity stimulation period could also provide

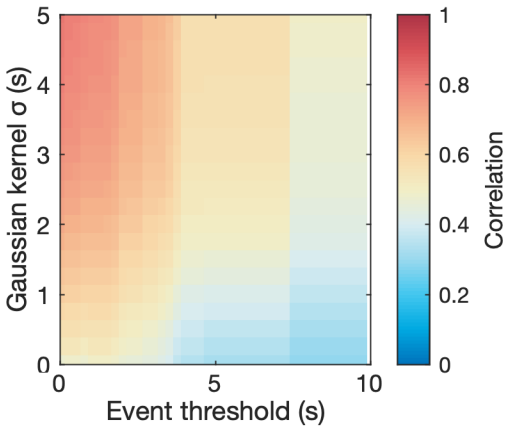


Figure 3.17. Correlation of alpha-detector activation and eye-tracking. Eye-closures are used to establish a reference for loss of attention, as a validation-proxy for the closed-loop system. To exclude natural blinking, an event threshold is necessary to distinguish inattention-related eye closures. A Gaussian kernel is applied prior to calculating peak cross-correlation to account for latency between binary signals, that might otherwise mask the presence of correlation. The correlation values are shown at various degrees of smoothing and event thresholds. The step change in correlation at ∼3–4 s likely indicates the physiological timescale for alpha-activity associated eye closures.

support for a mechanistic hypothesis of closed-loop neuromodulation, helping distinguish states of true vigilance or arousal, from states where eyes are reopened for other reasons, e.g. general boredom, feelings of stimulation-induced paraesthesia.

3.6 Discussion

In this chapter, we demonstrated the feasibility of deploying real-time neural state adaptation for DBS therapy in the Picostim-DyNeuMo family of implantables, based on a highly optimised implementation of a classical band-power estimation algorithm.

Devices successfully targeting large-signal neural state classification *in vivo* have been demonstrated before, starting with the since discontinued sensing-only NeuroVista seizure-forecasting system [134], followed by stimulation-capable devices, such as the Medtronic Activa PC+S [170] and Summit RC+S [197] (Medtronic plc, Minneapolis, MN, USA), and the NeuroPace RNS (NeuroPace Inc., Mountain View, CA, USA).

However, most experiments targeting small-signal spectral markers, as in the case of PD, would rely on chains of custom desk bench equipment, contributing to variation in reproducibility [151]. As most closed-loop applications have been implemented externally, investigators have to be able to use computationally-complex algorithms, rather than having to deal with the constraints imposed by IPGs. This reliance on external systems has so far prevented long-term validation of the fast-adaptive closed-loop DBS paradigm. In the commercial space, the Newronika AlphaDBS system (Newronika S.p.A., Milan, Italy) is currently undergoing clinical trials for aDBS in Parkinson's disease [131, 198].

The introduced algorithmic capabilities of the DyNeuMo were validated *ex vivo* with the device placed in a tank of isotonic saline solution (0.9 % w/v), serving as a realistic model of the tissue-electrode interface [199]. The testing established a measurement noise floor of 190.0 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ over the 10–40 Hz range in the presence of biphasic stimulation up to 4.0 mA. In the large-signal case for seizure detection, over a long-term test carried out on a clinician-labelled human thalamic signal, the filter chain achieved a true positive rate of 82.7 % and a false positive rate of 9.7 %. On a challenging 1.5 $\mu\mathrm{V_{RMS}}$ artificial beta-burst signal in long-term tests, the algorithm reached a true positive rate of 98.9 % at a false positive rate of 12.0 %, indicating that the system would also be suited for performing aDBS for Parkinson's disease in a majority of patients.

As the DyNeuMo algorithmic capabilities can be made available to Picostim users via an in-clinic firmware update, clinical trials into the long-term effects of aDBS therapy can be substantially de-risked – patients found not benefiting can easily be reverted to an approved version of firmware providing conventional therapy. The performance validation presented takes the first steps in supporting the case for clinical evaluation.

The capacity for low-noise recording during stimulation along with wireless data streaming makes the DyNeuMo an ideal candidate for a research platform for translational neurotherapies more broadly. In this chapter, we described the design steps taken to adapt the IPG for externalised research use, and demonstrated its successful implementation *in vivo* through experimental recordings in a freely moving NHP model. The presented pilot data collected in this setting demonstrates successful deployment of the aDBS algorithm targeting alpha-band activity, tuned on the go to a subject's neural activity patterns. The pilot study represents an important experimental milestone toward bringing a novel therapeutic protocol closer to patients suffering from impairments of consciousness.

Despite our extensive efforts at reducing the computational footprint of the adaptive algorithm, IPGs running the experimental DyNeuMo firmware were limited to operating times of around 30 h in the aDBS mode. While this duration makes the system feasible to use with daily charging and includes a 6 h safety margin, the overall power consumption is somewhat unimpressive given the typical weekly recharge rate of the baseline Picostim. While previous studies suggested reduction in stimulation delivery by about 50 % via the aDBS method are possible without affecting clinical outcomes in Parkinson's disease [42], savings on stimulation energy are insignificant when compared to presently available computational costs.

In future generations of IPGs, computational energy can be reduced by designing low-power co-processor cores, off-loading the repetitive algorithmic computation from general purpose microcontrollers. The availability of aDBS algorithms in a chronic implantable device today should not be discounted, even if inconvenienced by daily charging though: this development takes a major step towards large-scale clinical evaluation of responsive neurostimulation approaches.

3.7 Contribution

Work on the Picostim-DyNeuMo project has been a significant team effort. For hardware development credit is owed to Bioinduction Ltd. Initial device capabilities, such as the circadian scheduler, motion-adaptation and simultaneous sensing and stimulation were piloted by the Denison group at the University of Oxford. My contribution to this initial validation was data analysis for feasibility testing for seizure detection (Fig. 3.4). My main contribution to the present state of the DyNeuMo is the selection and implementation of low-resource filter algorithms (Section 3.3), in equal contribution with Mayela Zamora of the Denison group, and consultation with Tom Gillbe of Bioinduction Ltd. for assessing device capabilities and firmware deployment. Moaad Benjaber of the Denison group conducted the saline tank validations, with Mayela Zamora and I contributing the algorithm configurations.

The *in vivo* experiment was conducted by Jonathan L. Baker and Keith Purpura at the Weill Medical College of Cornell University, with Mayela Zamora leading algorithm configuration. My contribution to the experiment was the design and creation of the IPG adapter system detailed in Fig. 3.13, and the data analysis presented in this chapter (Fig. 3.15 – 3.17).

Chapter 4

Real-time phase estimation for closed-loop stimulation

4.1 Introduction

We previously introduced the phenomenon of oscillatory activity in the central nervous system, emerging from the coherent, synchronised firing of large populations of neurons. Based on their role in the network, however, the activity of different cell types will align to different phases of the LFP activity [200].

In open-loop stimulation paradigms amplifying and suppressing effects can depend on complex patterns of frequency alignment [127, 201], while classical high-frequency DBS is only capable of eliciting local suppression of activity, not amplification thereof.

In contrast, stimulation targeting specific phases of local oscillatory activity offers the potential to interact with neural activity in a more efficient way, by strengthening or disrupting local neurons at the exact times critical for establishing population-wide synchrony [25, 122, 164, 202]. In conjunction with modern optogenetics techniques, experiments could be designed to investigate the specific network contribution of individual cell types to LFP oscillations.

Prior work achieving successful phase-locked neurostimulation is sparse, with most attempts over the past decade relying on combinations of custom equipment and high computational complexity, which has so far limited adoption in the scientific community. The majority of work in this space focused on non-invasive recording and stimulation systems [203–207], or non-invasive feedback paths such as visual neurofeedback [208].

Closed-loop optogenetic stimulation has previously been reported in non-phase-locked contexts, such as based on single-neuron activity [61], or more complex pattern recognition algorithms for the suppression of epileptiform activity in rodent models of epilepsy [209, 210].

A successful combination of phase-estimation and optogenetic stimulation was first reported in 2014, aiming to manipulate hippocampal theta oscillations in mice [211]. The method used for phase tracking relied on a computer-in-the-loop processing approach, yielding a relatively poor round trip latency of ∼20 ms (15–25 % error). The algorithm used did not offer uniform performance throughout the oscillatory cycle either, performing with significantly less accuracy when aiming at the trough of theta oscillations as compared to the peak [211].

Over the years number of computer-in-the-loop studies have reported very similar latencies despite using a wide range of algorithmic techniques [61], suggesting an inherent overhead of ∼10–20 ms exists in the communication between typical recording systems and the processing computer [114]. To overcome this limitation, the closed-loop detector algorithms should aim to be implemented on the recording interface level, necessitating advanced, resource-conscious programming approaches.

In this chapter, we focus on the design, validation, and first deployment of a general-purpose research tool for closed-loop phase-locked optogenetic stimulation studies. The proposed method aims to improve upon previously reported phase-tracking capabilities by eliminating the need for computer-based processing, thereby significantly reducing latency.

4.1.1 The phase-magnitude duality

While oscillatory (periodic, recurring) processes can be described in terms of two natural quantities: phase and magnitude, typical measurements do not readily provide these quantities of a real-time behaving system. Though one can imagine the past of a periodic process can be used to estimate the present phase within a cycle, this breaks down when processes are non-stationary, as most biological systems of analytical interest are. In offline non-causal processing, i.e. when both the past and future of points within a signal are already known (save for the edges), the Hilbert transform offers an analytical method for revealing the instantaneous magnitude and instantaneous phase of each point within the signal. To gain intuition on why this is possible, the reader is encouraged to consult the visualisation of Fig. 4.1. Consider when the amplitude of an oscillating signal momentarily reaches zero: we intuitively assume the oscillation is to continue on its prescribed trajectory from this point. This is because of our understanding that even at zero amplitude the underlying process maintains the same amount of energy, i.e. it retains the 'momentum' to continue – this energy is simply present in a mode not observable to us. In a mathematical sense, oscillatory processes can be described as complex-valued time series, with our measurements of them capturing the real-valued part. Methods such as the Hilbert

transform calculate the analytical continuation of our real-valued signal, in essence reconstructing its complex nature. The complex signal then allows us to measure instantaneous phase and magnitude based on various projections, as shown in the demonstration of Fig. 4.1.

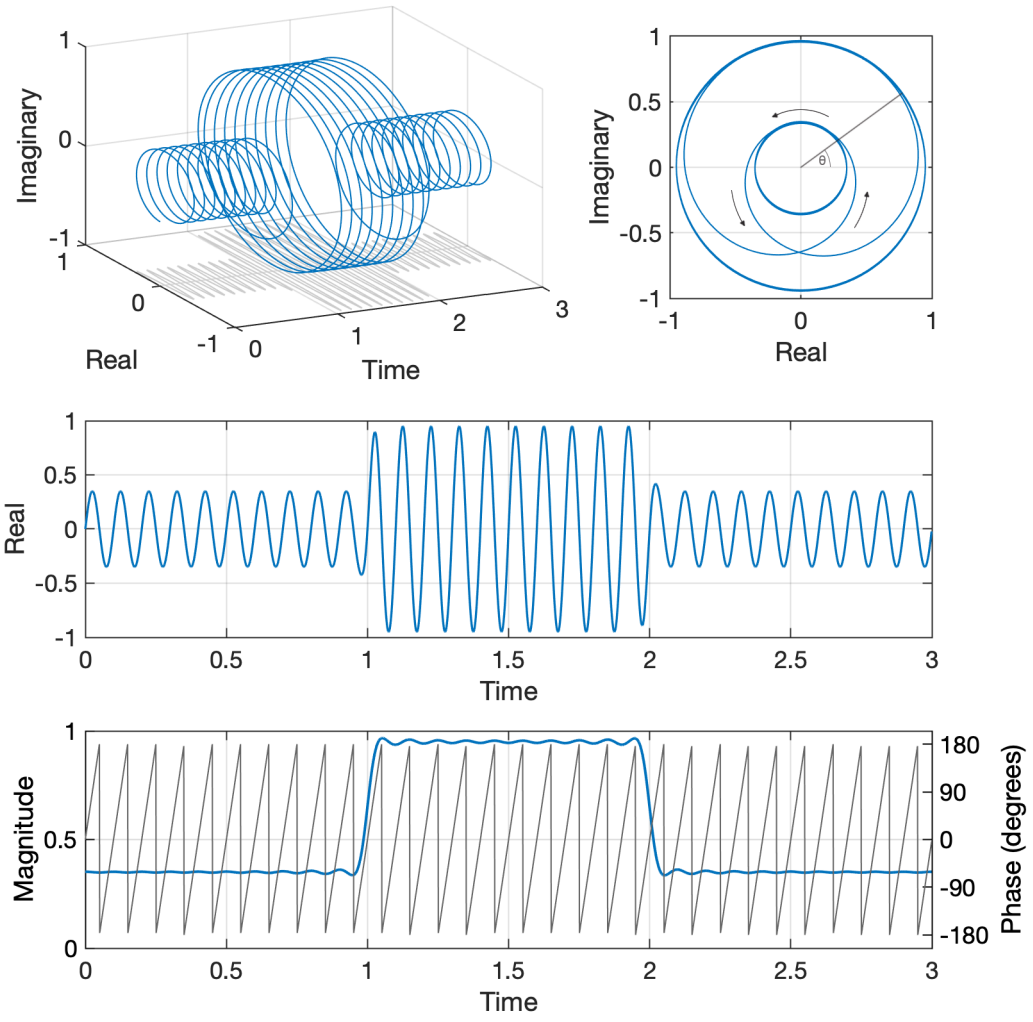


Figure 4.1. Analytical continuation of oscillating signals. The middle panel shows a 10 Hz oscillatory process, with a baseline amplitude of ±0.4 units, transitioning to and from a higher intensity burst with an amplitude of ±0.9 units, lasting 1 s. The top left panel shows the complex signal obtained by performing a Hilbert transform on the test signal. Note how real measurements are akin to observing the complex process from a limited perspective. The top right panel shows the same complex signal from a view orthogonal to the time axis, showing how stationary sinusoidal oscillations collapse onto circular trajectories, lending themselves to a natural definition of instantaneous phase. The bottom panel shows the real-time phase and magnitude estimates obtained via the Hilbert transform.

4.1.2 Real-time phase estimation

We identified broadly three classes of methods for real-time phase-estimation, known to have been previously trialled in neuroscience research.

The oldest of these methods is the so-called endpoint-corrected Hilbert transform. The Hilbert transform is commonly computed via a manipulation of the discrete Fourier transform of signals. The endpoint-correction method performs an additional filtering operation in the frequency domain prior to calculating the inverse transform to decrease the distortion at the two edges of the transform window, improving the phase-estimate accuracy at the last signal sample [207, 212].

Another method used autoregressive forward prediction of the input signal to create an extended window on which to calculate the Hilbert transform [203–206, 208, 211].

Both methods require relatively large running window lengths for accurate prediction, precipitating into large memory consumption for embedded device implementations in addition to relatively high processing costs. While methods typically lock on to an input signal to within $\pm 10\,\%$ error in $\sim$2–3 oscillatory cycles, they are highly sensitive to parameter selection and are easily knocked out of the steady state by stimulation artefacts.

A newer approach uses a Kálmán filter to estimate the state of a modelled oscillatory process [213]. A Kálmán filter can be implemented with lower computational cost than previous Hilbert transform-based methods and offers several additional benefits. Importantly, a method for parameter selection can be derived analytically based on SNR and oscillatory variability. As another interesting property, Kálmán filters intrinsically provide confidence bounds on the state estimate, which could be used to establish a mechanism for disengaging stimulation when estimates are of low-quality [213].

In general though, it can be an arduous task trying to compare the performance of various phase tracking models. For one, authors of different methods tend to report their performance characteristics using different measures. A more fundamental issue, however, is the problem of establishing a ground truth for phase [214]. While the Hilbert transform is widely regarded as the 'gold standard' method for calculating instantaneous phase, it only really provides meaningful results for 'pure' oscillatory processes. In the context of neuroscience, the quality factor of neural signals is low, and oscillations are embedded in a broad mixture of off-target activity. In practice, oscillatory processes of interest are isolated from the broad-band signal via band-pass filtering. It is easy to see though, how the exact selection of band-pass filters is going to affect

Special thanks to Karen Wendt (Department of Engineering Science, University of Oxford) for carrying out an internal literature review of real-time phase tracking methods, this section draws on her work.

the phase estimate we obtain. As no theoretically 'most correct' choice of band-pass filter exists, testing phase tracking accuracy on any spectrally complex real-world data can only serve as a qualitative guideline for correctness in general. Phase tracking methods ultimately have to be evaluated on test signals closely representative of the intended use case, and the exact trade-offs (e.g. lock-in time vs steady-state error) have to be balanced by the experimenter.

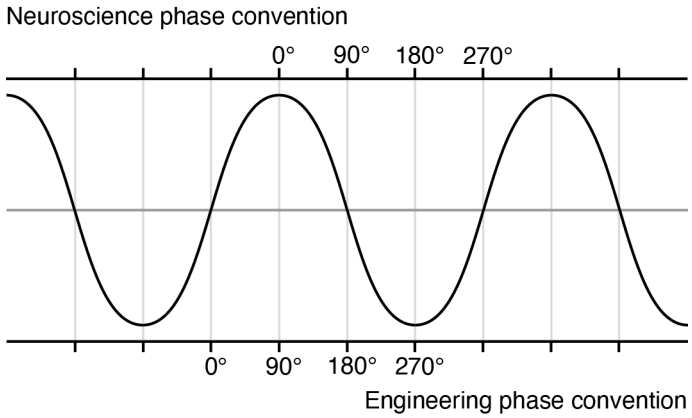


Figure 4.2. Conventions of phase angle measurement

4.1.3 Aims

McNamara et al. [123, 124] recently presented a novel method for real-time phase estimation, and successfully demonstrated its use in a neuroscience context through closed-loop modulation of beta-band LFP activity in a Parkinsonian rat model. The original study did not make use of a magnitude estimate alongside the phase information.

Notably, the method was arrived at from an algorithm development flow rather than from analytical principles, and as such does not posit to be an optimal estimator of oscillatory phase, making its general purpose performance characteristics largely unknown. Based on the original work, parameter selection is left largely empirical. For the method to see widespread adoption, it is necessary to establish a theoretical framework behind the algorithm. This understanding will make it possible to generalise parameters for high-quality operation under new experimental settings.

As the original study used electrical stimulation exclusively as the feedback element, it remains to be seen how the method performs in an optogenetics setup, of great practical interest to the research community.

With these observations in mind, we define the following aims:

- Analyse the theoretical operation of the phase-estimation algorithm
- Provide methods to optimise meta-parameters for arbitrary targets
- Characterise phase-tracking performance over the parameter space
- Provide a user interface for configuring settings and visualising the state estimates
- Pilot the method and its newly developed interface with optogenetic stimulation

4.2 Technical methods

4.2.1 The state estimator method

The method selected to form the basis of this work was first used in [123], and was originally described under patent no. WO/2020/183152. In the context of the work presented here, the method can be regarded as an empirical state estimator model (ESE model for short) for the purposes of analysis. The method can be described using the steps outlined below.

For a measured input signal s_n at timestep n, the algorithm estimates a complex state vector $\vec{r} = x + iy = [x, y]^{\mathrm{T}}$,

$$\vec{r}_n = \begin{bmatrix} \sin\theta_n & \cos\theta_n \\ -\cos\theta_n & \sin\theta_n \end{bmatrix} \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} \tag{4.2.1}$$

with θ_n progressing at a constant step size of

$$\theta_{n+1} - \theta_n = \frac{2\pi f_c}{f_s} \tag{4.2.2}$$

where f_c represents the target frequency, and f_s the sampling frequency.

The parameters of our state estimation model, α_n and β_n, are updated to minimise the error (Δ_n) between the input signal (s_n) and the real part of the reconstructed vector on the current sample.

$$\begin{gathered} \Delta_n = s_n - \mathrm{Re}(\vec{r}_n) \\ \begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} + \gamma\Delta_n \begin{bmatrix} \sin\theta_n \\ \cos\theta_n \end{bmatrix} \end{gathered} \tag{4.2.3}$$

where γ is referred to as the error gain.

The quantities of interest, phase and magnitude can be obtained from the estimated complex vector as follows.

$$\phi_n = \mathrm{angle}(\vec{r}_n) = \arctan(\mathrm{Im}(\vec{r}_n)/\mathrm{Re}(\vec{r}_n)) \tag{4.2.4}$$

$$M_n = \mathrm{abs}(\vec{r}_n) = \sqrt{\mathrm{Re}(\vec{r}_n)^2 + \mathrm{Im}(\vec{r}_n)^2} \tag{4.2.5}$$

with angle defined as anticlockwise from the positive real axis to the vector.

Phasor interpretation

Let us first consider the 2×2 matrix of Eq. 4.2.1. The columns of the matrix can be thought of as two vectors, $\vec{v}_1 = [\sin\theta_n, -\cos\theta_n]^{\mathrm{T}}$ and $\vec{v}_2 = [\cos\theta_n, \sin\theta_n]^{\mathrm{T}}$, with θ_n the phase at time sample n. The two vectors are permanently locked in an orthogonal arrangement by their definitions, and both are of unit length. As θ_n progresses, $\vec{v}_1$ and $\vec{v}_2$ essentially from a rotating coordinate frame, as shown in Fig. 4.3 (left panel). The rate of rotation is equal to the target frequency, due to Eq. 4.2.2. The parameters α_n and β_n can then be considered to be the coordinate measures of the estimated state vector $\vec{r}$, in the $\vec{v}_1$, $\vec{v}_2$ rotating reference frame, i.e. $\vec{r}_n = [\alpha_n, \beta_n]_v^{\mathrm{T}}$.

The 2×2 matrix formed by column vectors $\vec{v}_1$, $\vec{v}_2$ can further be recognised as equivalent to the canonical two-dimensional rotation matrix $\mathbf{R}_\theta$ to within a 90° shift in angle θ:

$$\mathbf{R}_\theta = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} = \begin{bmatrix} \sin(\theta - 90°) & \cos(\theta - 90°) \\ -\cos(\theta - 90°) & \sin(\theta - 90°) \end{bmatrix} = \begin{bmatrix} \sin\theta' & \cos\theta' \\ -\cos\theta' & \sin\theta' \end{bmatrix} \tag{4.2.6}$$

where $\theta' = \theta - 90°$. Note that this is a consequence of the 'neuroscience notation' of phase angle, explained previously in Fig. 4.2.

Let us now consider the meaning of the equations used for refining the state estimate (Eq. 4.2.3). As we can only observe the real part of a signal through direct measurement, hence the very need for an estimation process for the complex state, we can understand that the Δ_n error term can only compare the measured samples s_n to the real part of the state estimate $\mathrm{Re}(\vec{r}_n)$. The error term can be thought of as a vector quantity collinear with the positive real axis since the measurement happens along it. The vector term included in the equation, $[\sin\theta_n, \cos\theta_n]^{\mathrm{T}}$, simply serves to project the error term from the complex coordinate axes to into the rotating reference frame. Finally, a weight term, γ is included to adjust the relative magnitude of cycle-to-cycle error correction. Overall, the weighted sample error $\gamma\Delta_n$ is used to either advance, or retreat the state estimate $\vec{r}$ in its rotation with respect to the coordinate frame. The process of the update cycle is shown graphically in the right-hand panel of Fig. 4.3.

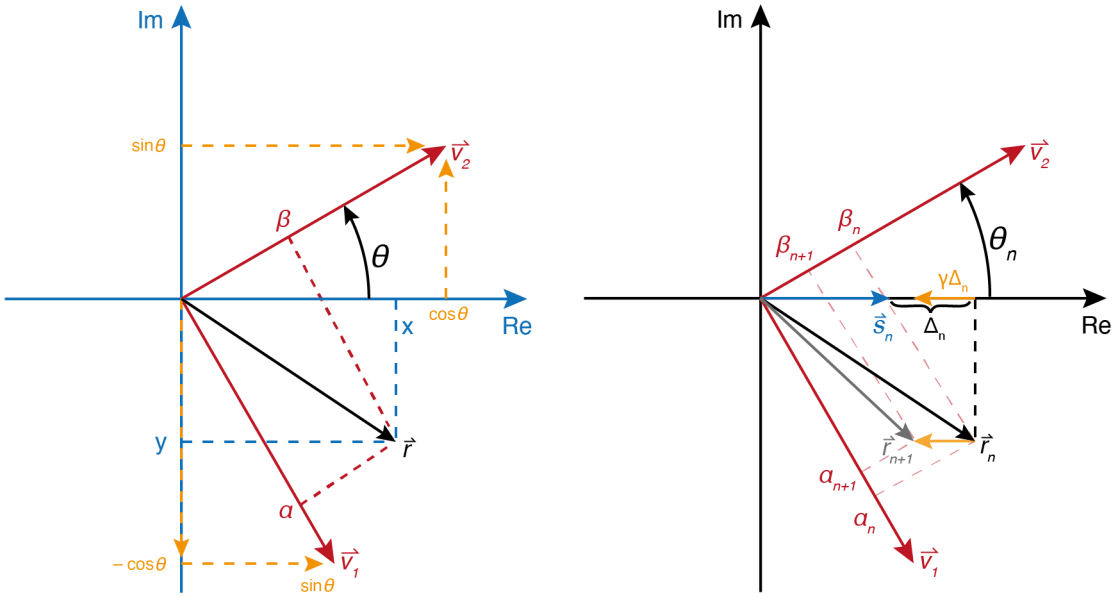


Figure 4.3. Visualisation of the phase estimation algorithm. **Left:** Projections of the state estimate between coordinate systems. The estimate, rotating in the complex plane to model an oscillating process, can be regarded as stationary in the $\vec{v}_1$, $\vec{v}_2$ rotating reference frame. The relationships between θ and the rotating frame are shown. **Right:** State update cycle of the estimation algorithm. The weighted sample error $\gamma\Delta_n$ is used to either advance or retreat $\vec{r}$ in its rotation with respect to the coordinate frame.

Quadrature interpretation

Some readers might find a different visualisation of the ESE scheme more helpful. Instead of considering the rotating frame of reference presented in the phasor interpretation, we can think of the rotating coordinate frame as a pair of signals 90° out of phase (one quarter-cycle, i.e. quadrature). Fig. 4.5 demonstrates the operation of the ESE method on a real-world signal example, with the lower panels visualising the internal sinusoidal reference signals ($\sin\theta_n$, $\cos\theta_n$) present in the algorithm's governing equations previously introduced, with the lower panel showing the two components of state estimate as signals themselves (α_n, β_n). Simplistically, we can describe the state estimation process as asking the question of how much $\sin\theta_n$ and $\cos\theta_n$ the input signal contains at any particular point in time. In more formal terms, the estimation process attempts to describe the input signal as a linear combination of the two orthogonal reference signals.

This is remarkably similar to the method of quadrature demodulation used in communication signalling, shown in Fig. 4.4, where an internal reference oscillator and a phase-shifted copy thereof are used to decompose an input signal into two orthogonal parts [215]. The state estimate $[\alpha_n, \beta_n]_v^{\mathrm{T}}$ then corresponds to the $[I, Q]^{\mathrm{T}}$ signal pair in the classical notation of quadrature demodulation.

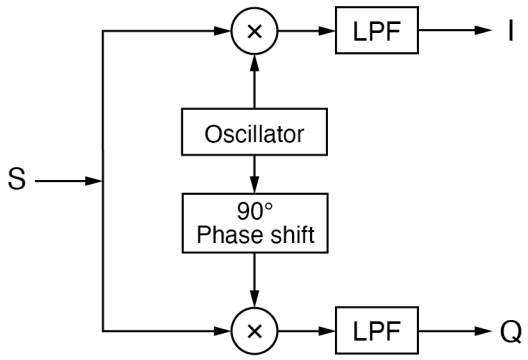


Figure 4.4. Quadrature demodulation scheme. The input signal S is multiplied by a pair of internal oscillator signals 90° out of phase (in quadrature), to provide an orthogonal decomposition in the form of the I/Q signal pair [215].

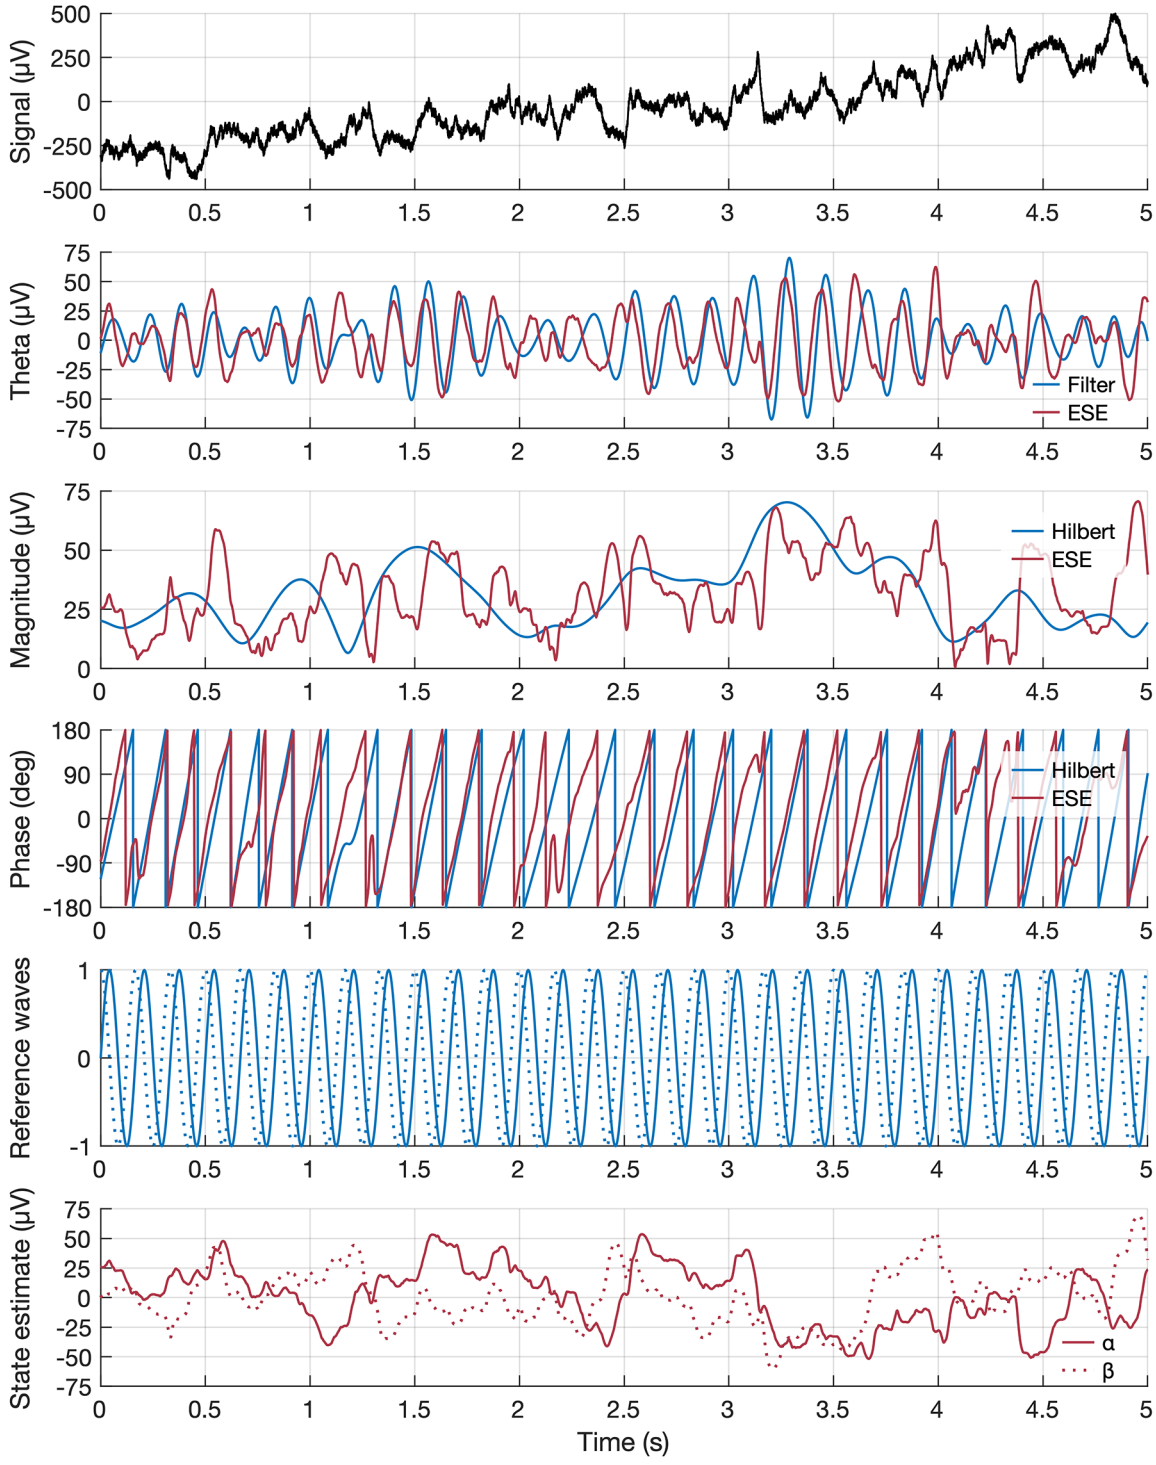


Figure 4.5. Example phase tracking data. A 5 s LFP recording segment from mouse primary motor cortex is shown, comparing the real-time ESE to an offline Hilbert transform. The Hilbert transform was performed after applying a 4th order zero-phase Butterworth filter at 6 ± 2 Hz, the ESE was configured with a feedback gain of 2^{-5}. The raw recording was sampled at 20 kHz. The real-time estimate of the theta oscillation shows close correspondence to the outcome of offline zero-phase band-pass filtering. The final panels show the instantaneous weights assigned to each of the quadrature reference signals, with solid and dotted lines corresponding between panels. Note the presence of interfering off-centre oscillatory frequencies, suggesting the feedback gain might be further decreased to increase tracking specificity (see Section 4.2.3).

4.2.2 Efficient implementation

The algorithm particularly favours constrained low-power processing systems, due to the availability of several shortcuts to normally complex computation – the lack of floating point acceleration, and trigonometric function support are both easily mitigated without significant impact on run-times. The most important strategies are outlined in the following sections.

Downsampling

For the most immediate saving on computational complexity, we should look at Eq. 4.2.2, and consider how the ESE models an oscillatory process. The rate of phase progression in the quadrature reference is set in proportion to the target frequency and the sampling rate, therefore we get an uneven model resolution based on the operating point. In fact at a typical sampling rate of 20 kHz, for the entire field potential frequency range of interest (0.1–200 Hz) the effective phase resolution ends up significantly greater ($< 4.5°$) than the limit of what can be considered practicable for targeting ($\sim 360°/16 = 22.5°$), given the relatively low-quality factor of neural oscillations. Therefore, it is proposed to downsample the input signal proportionally to the input frequency, so the number of points modelled in the phase progression is at least within the same order of magnitude. Downsampling also has the advantageous side-effect of making performance characteristics more or less normalised over a range of operating points, facilitating initial parameter estimation.

A relatively crude form of n-fold downsampling by averaging was used, implemented in the following form using integer arithmetic and bit shifts, to prevent a need for floating point operations and division:

$$q_j = \sum^{d} p_i \times 2^{-ceil(log_2(d))} \tag{4.2.7}$$

where d is the rate of downsampling, and $ceil()$ represents the ceiling function, that is a function mapping any $x \in \mathbb{R}$ to the least integer greater than or equal to x. The parameter d should be chosen to for any target frequency such that the reduced sampling rate should provide 40–80 points per cycle.

This downsampled process produces a biased estimate of a true rate-reduced signal – specifically, the output is equal to the average signal when d is an integer power of 2, and is strictly underestimated otherwise. The bias only materially affects the magnitude estimate though, and

is expected not to have a significant impact on the overall algorithm operation given at least moderate SNR. Importantly, the magnitude bias is deterministic and can thus be accounted for at the data interpretation stage.

Wavetables

The downsampling process as described before, provides a rough regularisation of the number of sample points per cycle of the target frequency (chosen to fall in the range 40–80 samples), making it possible to pre-compute a number of reference waves within this range and to store them in relatively little memory. Due to symmetry, a stationary sinusoid signal can be fully reconstructed from a quarter of a cycle, allowing for significant compression of the look-up table.

As elaborated in Section 4.3.3, the rate of downsampling and the optimal look-up table resolution for the quadrature references are dependent parameters and can be optimally selected for any desired operating point of the algorithm.

CORDICs

The CORDIC (**CO**ordinate **R**otation **DI**gital **C**omputer) family of algorithms, created for early computer navigation systems by Volder in 1959, offer a convenient way to reduce complex trigonometric functions to a set of classical digital operations: addition/subtraction, bit shifts and table look-ups [216]. The algorithms use a binary quantised notation of phase angle, i.e. units of $360°/2^N$. The scheme of operation, as outlined in the following text, is shown graphically in Fig. 4.6. First, the input vector is pre-rotated to have non-negative real and imaginary parts, placing it into the upper right quadrant of the coordinate space. The algorithm then operates by performing a sequence of smaller and smaller rotations on the input vector, with the amount of rotation decreasing with orders of magnitude on the binary scale $(45°, 22.5°, 11.25°...)$. At each step the sign of the imaginary part of the vector is evaluated, determining the direction of the rotation to follow – the next rotation is anticlockwise if the vector has been rotated 'too much', pushing it below the positive real axis, clockwise otherwise. The total amount of positive and negative (clockwise and anticlockwise) rotations performed are accumulated, providing a sequentially refined estimate of the phase angle as the input vector is rotated closer and closer to the positive real axis. This method of operation simultaneously provides an estimate of the vector magnitude without any additional computation, eliminating the need for evaluating a square root, as per Eq. 4.2.5. As the vector is rotated arbitrarily close to the positive real axis, the real part will approximately equal the vector magnitude. The number of approximation steps can

be selected on a case-by-case basis depending on the angular resolution and accuracy required in any particular application, providing a further level of resource optimisation. For information on the accuracy gained over subsequent steps, as well as for sample code for the algorithm, the reader is referred to a widely used application note on the matter from Turkowski [217]. Our implementations of the ESE method used 11 stages.

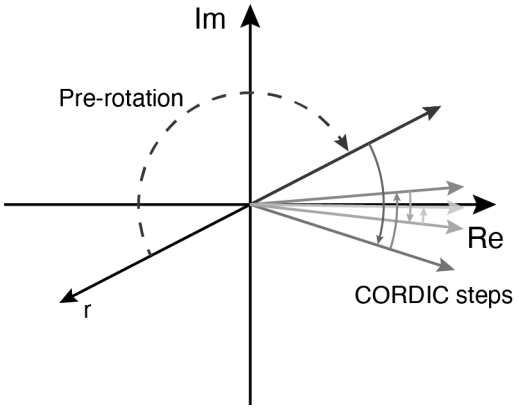


Figure 4.6. Operation of a CORDIC algorithm, visualised. Graduated colours represent the sequence of operations, lighter shades correspond to later steps. The input vector r is pre-rotated into the upper right quadrant, then progressively rotated closer and closer to the positive real axis, where the angle of rotation decreases two-fold with each step. The direction of rotation is reversed when r' has a negative imaginary part.

Restricted gain

A critical part of the algorithm is selecting appropriate values for the error gain. In practice, it appeared to be sufficient to use gains restricted to negative integer powers of two, i.e. $\gamma = 2^{-g}$, $g \in \mathbb{N}$. This decreasing scale of binary orders offers a tractable space for parameter search in practice, while still capturing a wide range of different process noise to measurement noise ratios. Importantly, it greatly simplifies the implementation in both software and hardware, as multiplication by a small number (requiring high precision) can be replaced by simple bit shifts.

High-pass filter

Raw signal values taken from the analogue measurement front-end can contain significant, and unpredictable DC offsets on each channel. To aid the operation of the ESE model, a very simple first-order digital high-pass filter was implemented. To minimise the computational burden of this stage, the filtering should take place after the downsampling step (but before the ESE stage). As a trade-off in configuration complexity, this means the corner frequencies available are going to change with both different sampling rates and target frequencies.

For an input signal q_n, the filter computes the output s_n as

$$s_n = q_n - x_n \tag{4.2.8}$$

$$x_{n+1} = x_n + \lambda s_n \tag{4.2.9}$$

where x_n is an internal memory state of the filter, and the parameter $\lambda \in [0, 1]$ is used to select the corner frequency, depending on the sampling rate f_s, and the ratio of downsampling d is

$$f_{HP} = \frac{\lambda f_s}{2\pi d}. \tag{4.2.10}$$

For computational efficiency, it is desirable to constrain the parameter λ as negative integer powers of two, i.e. $\lambda = 2^{-l}, \quad l \in \mathbb{N}$, or in practice to $l \in \{0, 15\}$. This allows implementation of the algorithm with integer arithmetic and bit shifts, without having to introduce either fixed point notation or floating point variables.

As a general guideline, the high-pass filter need not be reserved for offset removal though. The corner frequency can be selected as the highest available in the binary scale that is below the lower end of the oscillation band of interest. This provides additional suppression of the more intense lower interfering frequencies in the $1/f$ spectral context of neural noise.

4.2.3 Evaluation

Closed-loop specificity

While the operation of the ESE model has hopefully been cleared up in the previous sections, it is not immediately apparent how it would function on real-world signals containing a mixture of frequencies. To find the specificity of the state estimate to the target frequency, let's explore the method from a new perspective. We have previously focused on extracting phase and magnitude estimates from the ESE as our primary aim. Recall, however, that phase and magnitude information, as a pair, are sufficient to reconstruct a signal. If the algorithm has any specificity at all, the signal reconstructed from the instantaneous phase and magnitude estimates should be distinct from the measured input signal, and ideally contain frequencies close to the target frequency – we would expect the reconstructed ESE signal to resemble the output of a band-pass filter.

We can intuitively expect the γ gain term to influence the narrowness of the band, by considering the process of the update cycle (Eq. 4.2.3), visually in Fig. 4.3. At a gain of unity, the real part of the state estimate is adjusted to equal the measured signal at each time step, making it similar to an all-pass filter, constrained only by the sampling rate. Conversely, at gains close to zero the model receives barely any adjustment at each cycle, meaning any effect has to accumulate over long periods of time, favouring only coherent processes (a match in frequency).

In a more formal sense, the gain γ is similar to the Kálmán gain in classical control theory and should be selected to balance process noise with measurement noise. In this context, process noise would mean the variability of the biological oscillation – including variations in amplitude (waxing and waning), as well as any jitter present in the cycles. Measurement noise is then the combination of out-of-band neural activity and noise from the measurement instrumentation. In the case of large process noise, the gain should be set higher to allow adjusting the state estimate closer to the measurement. Conversely, with large measurement noise, the gain should be set lower to favour the modelling process. While lower gains can provide greater immunity to measurement noise, it imposes a limit on the rate of convergence between the state estimate, and the external process, impacting the temporal dynamics. Importantly, the gain should not reach (or exceed) unity, as this would mean ignoring the model completely, simply passing through the measured values instead.

The results of Fig. 4.7. present the frequency response of the ESE model at various gains and target frequencies, and seem to confirm our expectations.

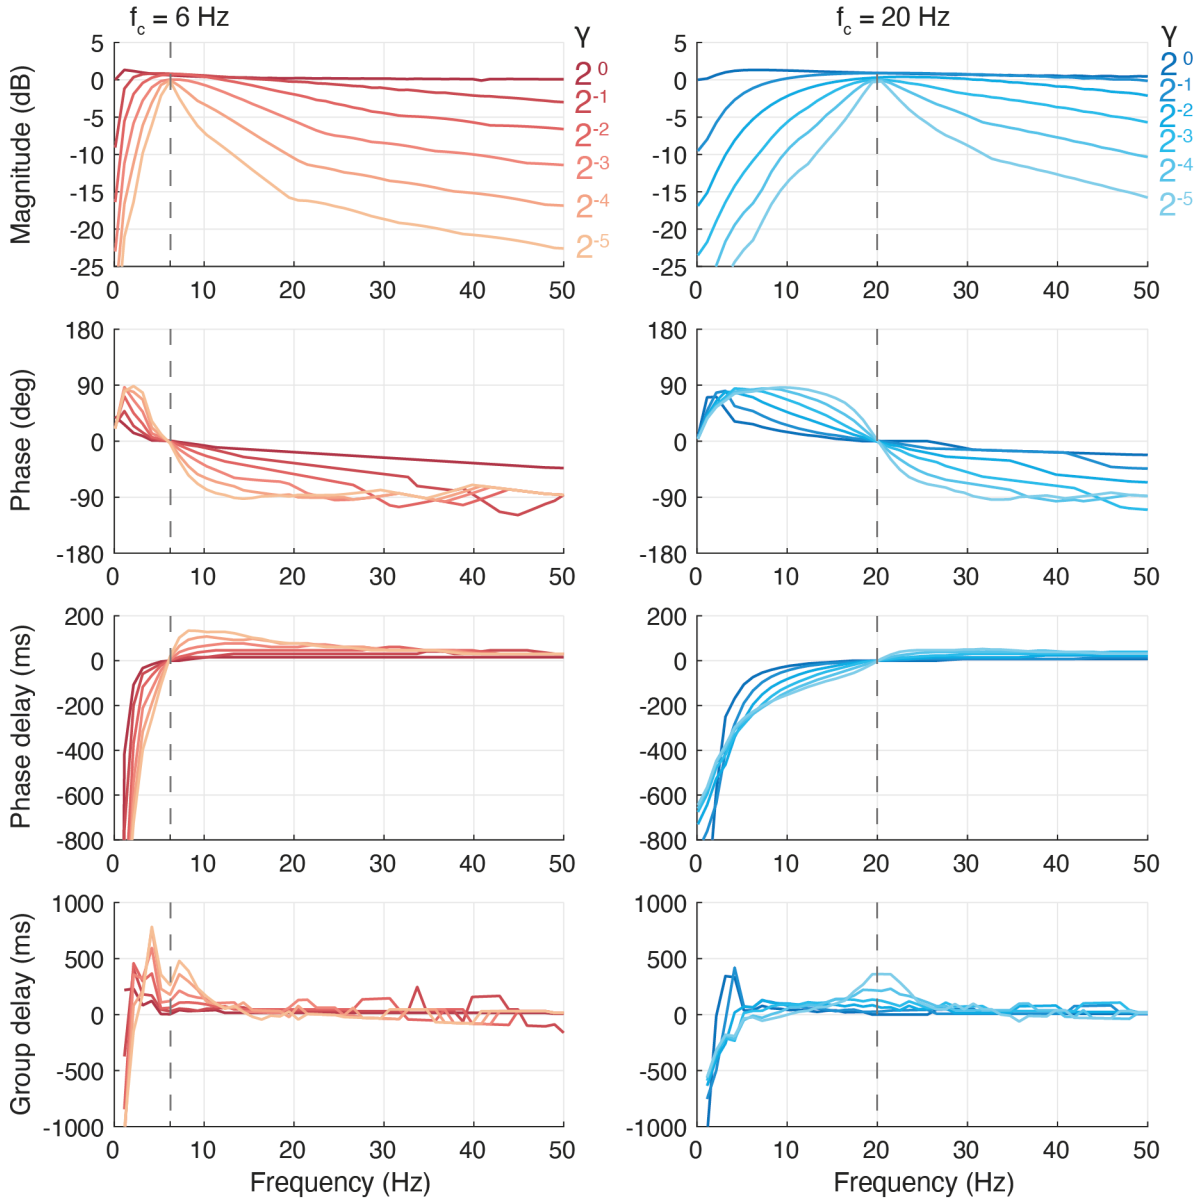


Figure 4.7. Steady state frequency response of the state estimator model. Four common measures used to characterise filters are presented for the ESE model, for two different centre frequencies (5 Hz, 20 Hz), and six values of the error gain parameter ($2^0 - 2^{-5}$). Note that while magnitude and phase responses are consistent with common real-time band-pass filters of various quality factors, the ESE model produces zero phase delay at the centre frequency, and very low (< 100 ms) group delays for moderate selectivity ($\gamma = 2^{-3}$).

High-pass filter performance

The frequency characteristics of the simple digital high-pass filter, described in Section 4.2.2, are shown in Fig. 4.8. As the results demonstrate, the simple filter structure remains stable, and does not suffer in performance from the integer arithmetic implementation – either in the expected cut-off frequencies or in the roll-off rate (20 dB/decade).

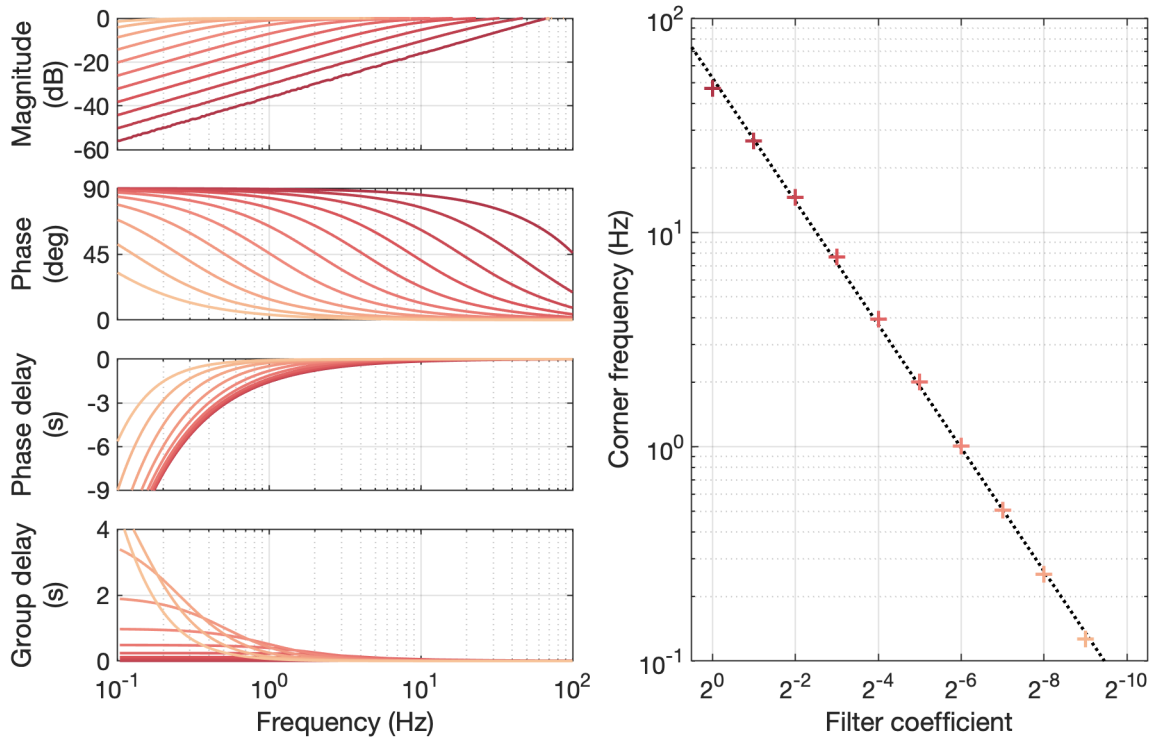


Figure 4.8. Steady state frequency response of the digital high-pass filter. The characteristics are presented for the range of filter coefficients $\lambda = 2^0 - 2^{-10}$. Corner frequencies were measured at -3 dB. Input signals were synthesised at 20 kHz, then downsampled 49× to $\sim$ 408 Hz, found to be the optimal processing rate for tracking 6 Hz theta oscillations in Section 4.3.3.

4.3 Development

4.3.1 User interface

We aimed to integrate the phase tracking method into the commercial Intan RHD2000 USB evaluation board (Intan Technologies Ltd., Los Angeles, CA, USA). The system was selected for its widespread use, low cost for the number of recording channels available, and its open source software. On the computer side, Intan provides a graphical user interface (GUI) written in C++ with the Qt cross-platform graphical framework, to configure their devices and manage recording data. Fig. 4.9 shows a typical view of the Intan user interface, displaying basic recording settings (sampling rate, filtering, etc.) with a live display of recorded neurological activity.

To maintain consistency with the data format of existing Intan recordings, the obtained phase and magnitude estimates were piped into existing analogue recording channels (Fig. 4.10). Similarly, the delivered closed-loop stimuli were made available on auxiliary digital channels recorded by the Intan system (Fig. 4.11). As a minor incongruity in the interface, the stimulation signal was connected to both a digital output and a digital input channel, as the Intan interface

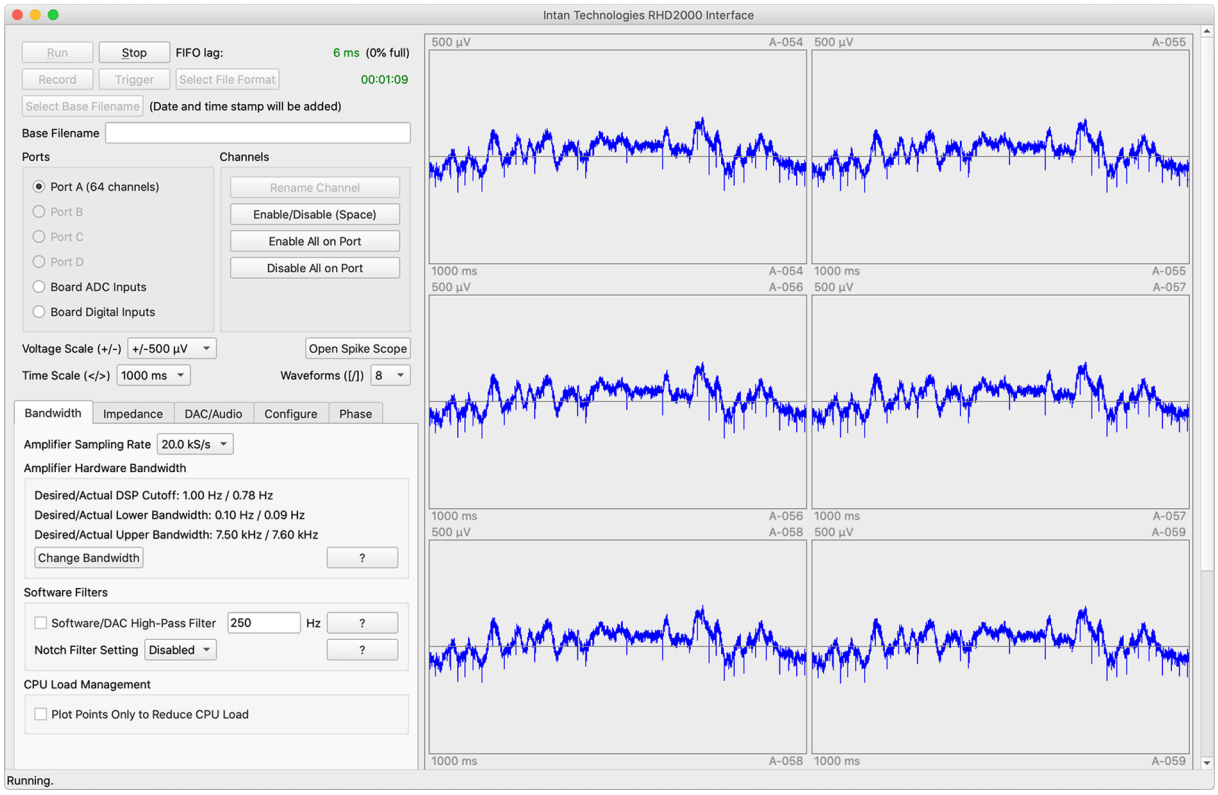


Figure 4.9. Default view of the Intan recording interface. Recording settings and visualisation options are shown in the left column, while six channels recording neural activity are visible on the right.

does not allow real-time visualisation of output channels, only input ones. The ability for users to validate stimulation delivery in real time made sacrificing a single extra channel a worthy trade-off in our view.

Configuring the phase tracker module as well as the closed-loop stimulation waveform is possible through a newly created setting tab in the interface (Fig. 4.12).

While it might be tempting to use the GUI's capabilities for rapid parameter prototyping, this approach is discouraged. Real-time estimates of the signal magnitude are displayed in arbitrary units, and thus serve only as visual confirmation of qualitative trends during experimental use. The phase estimate may provide a visual indication of noise. Generally speaking though, deriving the numerical configuration of the phase tracker is best performed offline.

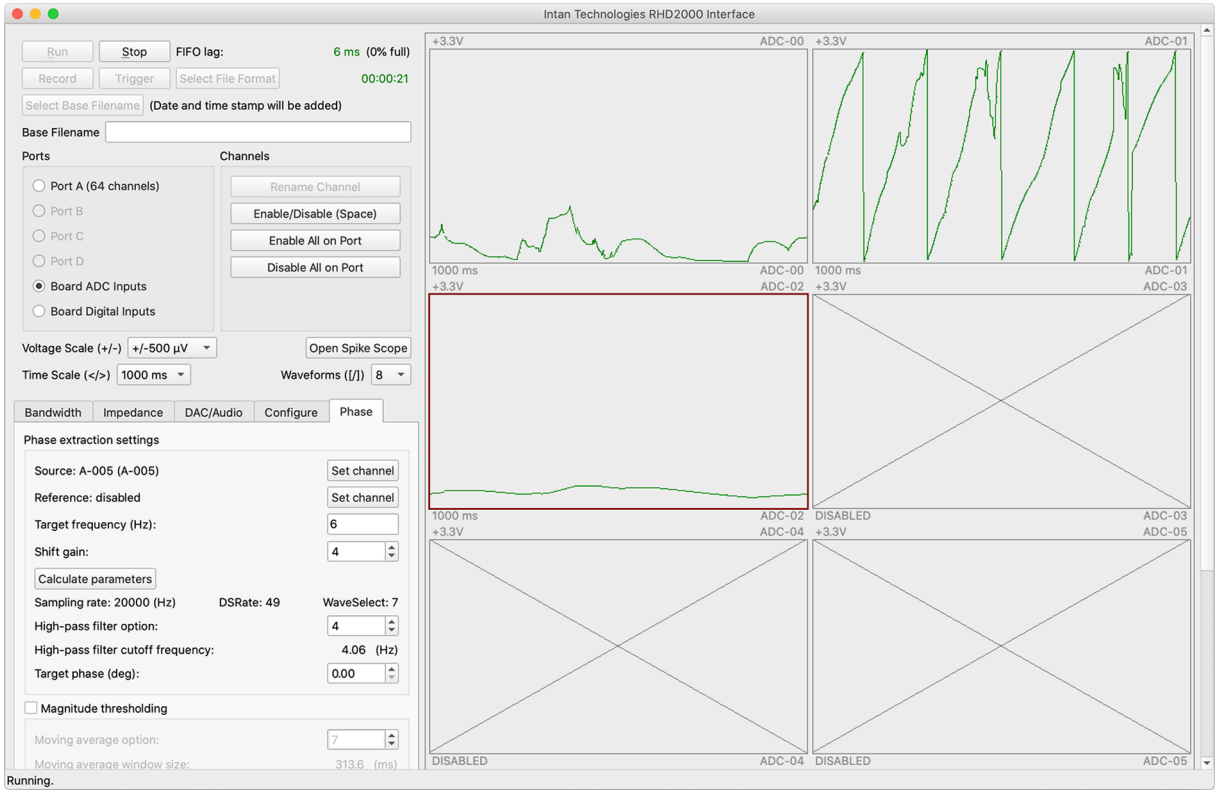


Figure 4.10. Phase tracking visualisation in the customised Intan recording interface. Three of the eight available analogue input channels were repurposed to display the output of the phase tracking module in real time. The raw magnitude estimate is shown in the top left panel before, and in the middle left panel after denoising. The phase of the signal is shown in the top right panel.

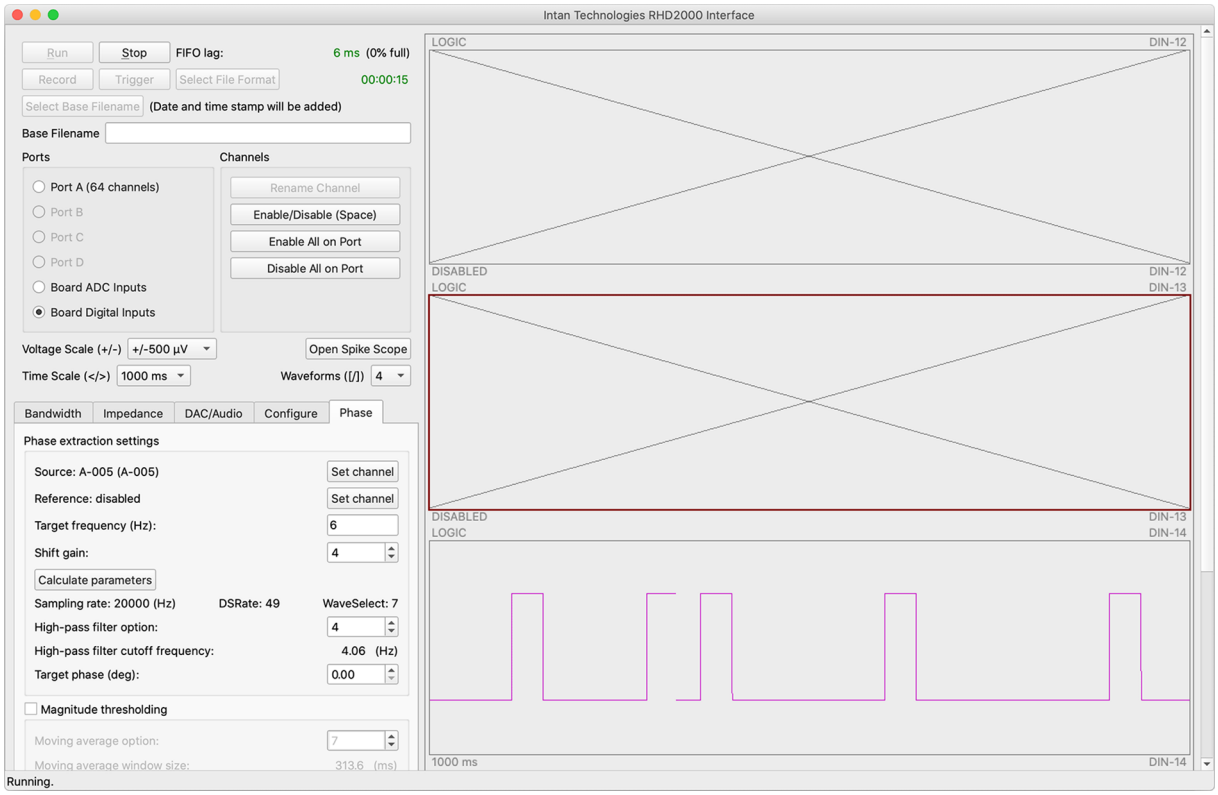


Figure 4.11. The stimulation state is shown on one of the 16 available digital input channels, to allow the user to verify the correct operation of the closed-loop system throughout an experiment.

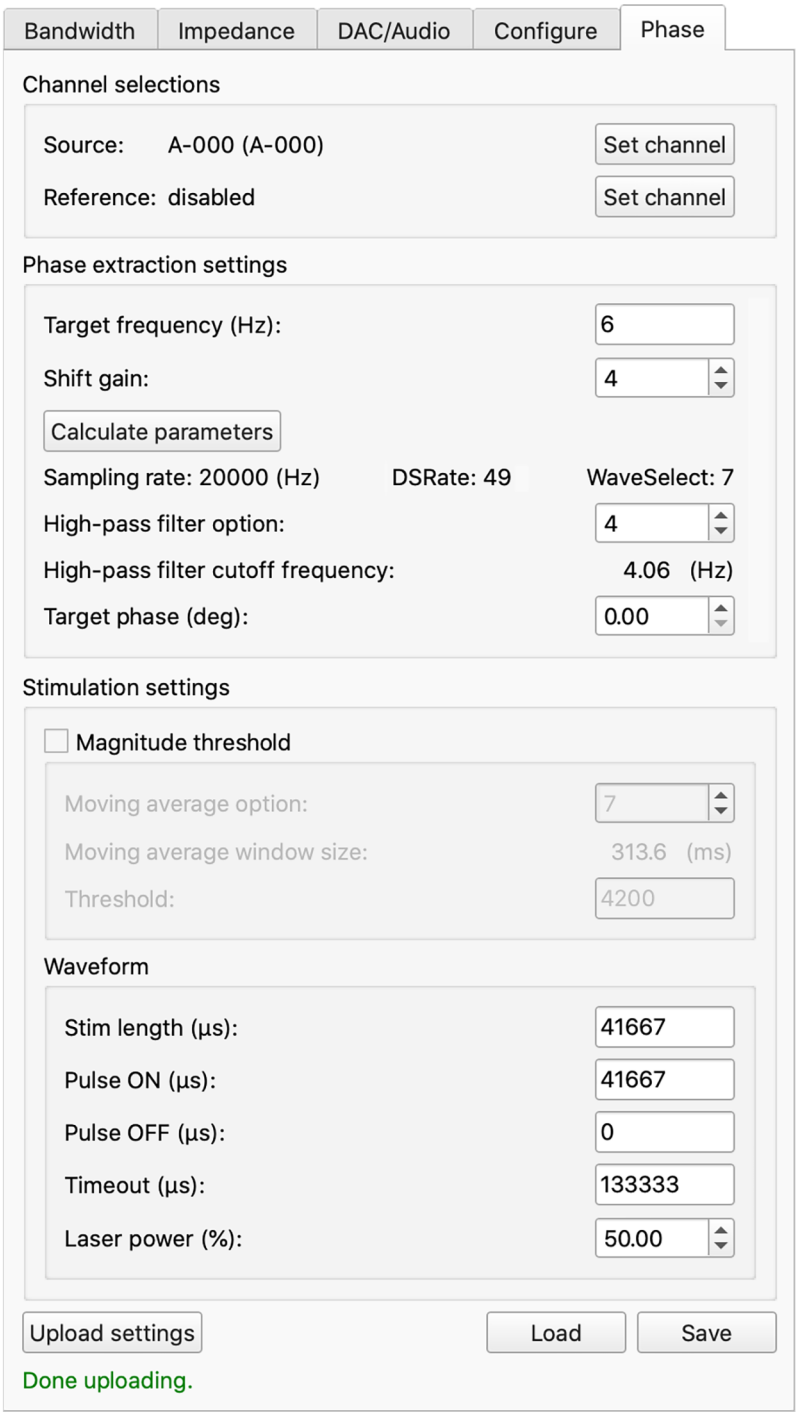


Figure 4.12. Closed-loop configuration tab in the custom Intan interface. The source signal for phase tracking (and an optional reference channel for differential measurement) are user selected by clicking on channels in the recording port view (Fig. 4.9). Phase tracking options include the choice of high-pass filter settings, target frequency, target phase, and loop gain. Stimulation waveforms, including sub-pulse timing, are user-programmable. Filtering and thresholding the magnitude estimate as an extra condition to deliver the phasic stimulation are optionally available. Settings profiles can be imported and exported.

Phase tracking and stimulation parameters can be exported and imported in the form of JavaScript Object Notation (JSON) files, a lightweight and human-readable format. Imported settings are validated before they can be sent to the device. If necessary, imported values are rounded to the nearest supported parameter values to ensure correct operation, in which case a pop-up window notifies the user of the executed changes.

Fig. 4.13 shows the structure of the extended interface software. As developers of the original code offer are free to make extensive changes on their end, the custom modules were designed to interact with the original codebase in as few places as possible. Keeping customisations heavily compartmentalised allows for reduced maintenance burden, and minimises the impact of future updates.

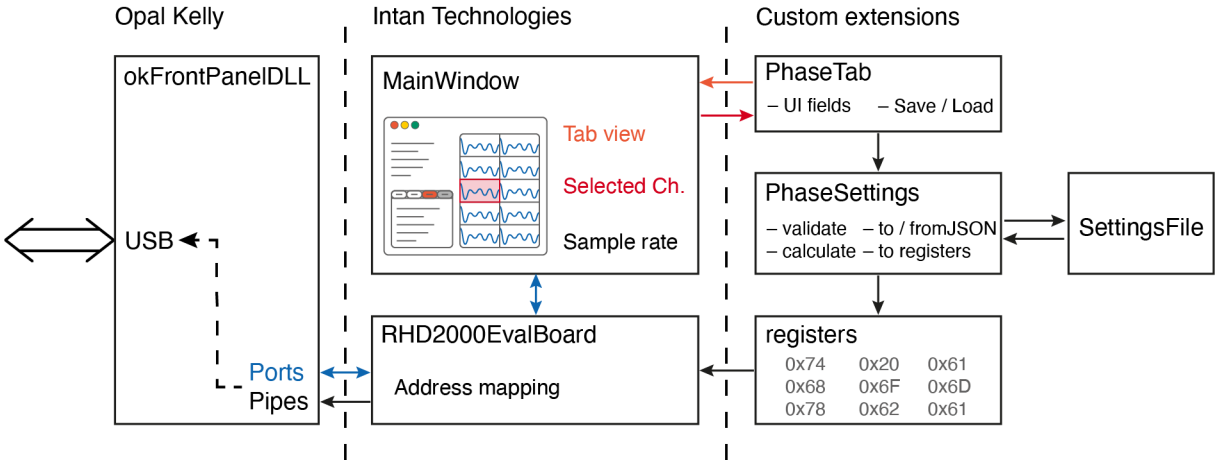


Figure 4.13. Class structure of the customised Intan graphical user interface. Key modules (classes) are highlighted and grouped by their developers. Standard Opal Kelly libraries were used to provide high-speed USB communication with the recording hardware, while Intan provided the basic user interface for saving and visualising electrophysiological data. Custom extensions comprised a high-level interface module, deployed as a settings tab into the Intan GUI. The phase tracking module extracts sampling rate and channel selection information from the main application window. Closed-loop settings were kept in high-level, user-friendly variables for UI configuration and storage, validated and converted into low-level register fields when being loaded into the recording device.

Illustrating the power of cross-platform frameworks, the interface was compiled under the Qt 5.15 Long Term Support (LTS) version and was successfully deployed on Windows 10 (update 21H2), Mac OS 10.15, Ubuntu 18.04.6 LTS, 20.04.4 LTS, 22.04 LTS, as well as Fedora 29 Workstation Edition, all in active use in various research groups.

4.3.2 FPGA interaction

The Intan USB recording interface board is built around an Opal Kelly XEM6010 USB/FPGA interface module (Opal Kelly Inc., Portland, OR, USA). The module houses a user-configurable field-programmable gate array (FPGA), with both supporting hardware and software modules for high-speed USB communication. The particular interface module uses a Xilinx Spartan-VI FPGA (XC6SLC45-2; Advanced Micro Devices Inc., Santa Clara, CA, USA) with *sim*45 thousand 6-input look-up tables (LUTs). The core Opal Kelly communication modules, and the Intan recording engine use approximately half of the LUTs, leaving significant room for custom extensions. Firmware development for Spartan-VI devices was possible via the now legacy Xilinx ISE 14.7 development environment, available under long-term support as of the time of writing.

Fig. 4.14 presents an overview of the functional blocks designed to achieve closed-loop phasic stimulation in the existing Intan platform.

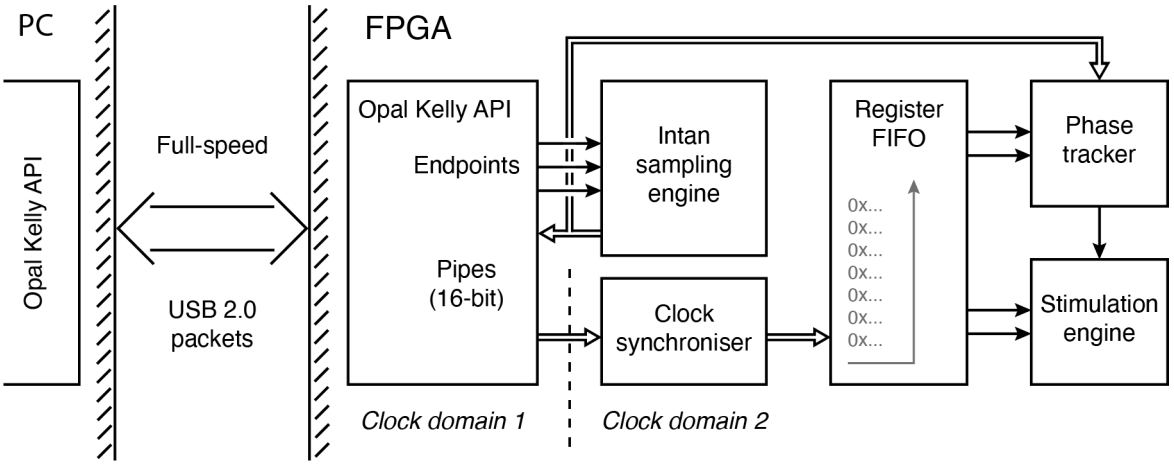


Figure 4.14. Structure of the customised Intan FPGA firmware. The embedded firmware communicates with the data collection application running on a computer via a full-speed USB 2.0 link. Data transfer is handled by proprietary application programming interface (API) modules provided by Opal Kelly, developers of the FPGA development module used. The API maps endpoints and pipes between the GUI and the FPGA fabric. A classical dual flip-flop clock-domain synchroniser was implemented to allow configuring the phase tracker and stimulation modules, that operated at a lower clock rate (matching the Intan sample rate), than the USB transaction rate.

4.3.3 Parameter selection

While it is the most natural from a user perspective to configure the phase tracking module by selecting the target phase, target frequency, and specificity, the algorithm operates on lower level parameters of wave table sizes and sampling rate ratios, as described in Section 4.2.2. For both facilitating adoption, and simplifying day-to-day use, it is imperative to abstract these parameters away from the user.

In most experimental context, the sampling frequency simply derives from our choice of recording electrode size and the scale of neural activity of interest. Intan systems operate with a default sampling rate of 20 kHz, commonly used for recording spiking activity in microelectrodes. Note that the base sampling rate does not affect the performance of the phase-tracker in any significant way above 1 kHz, the lowest value available in the Intan user interface.

The selection of the centre frequency is going to be a key user decision, and should be based on experimental measurements of pilot data before deploying the phase-tracker, to ensure an oscillatory process at the target frequency in fact exists. Popular methods for automating the detection of oscillatory processes in neural data are provided in references [218, 219]. Selecting a centre frequency becomes problematic for widely defined oscillations though, as standard definitions of phase are no longer applicable, e.g. a loosely defined 'ripple-band' of 100–250 Hz in hippocampal function. Such special cases are beyond the scope of this work. Users are advised to use well-defined centre-frequencies for narrow-band oscillations.

For a given sampling frequency f_s, desired target frequency f_c, and a progressive quarter-wave table with at least s_{min} and at most s_{max} points, the following algorithm can be employed to find the combination of an integer downsampling ratio d and quarter-wave resolution s, that approximates f_c with the smallest absolute error. Ties in the error are broken in favour of the parameter pair with greater wave resolution.

Given that most users will interact with the phase-tracking system via the GUI, $s_{min} = 10$ and $s_{max} = 19$ can be assumed as fixed parameters, leaving only the centre frequency and sampling rate to decide.

```
get_params (f_s, f_c, s_min, s_max)
    i_min ← 0;
    for i = 0 to (s_max − s_min) do
        d_list(i) ← f_s/f_c/(s_min + i)/4;
        f_list(i) ← f_s/d_list(i)/(s_min + i)/4;
        err(i) ← |f_list(i) − f_c|;
        if err(i) ≤ err(i_min) then
            i_min ← i;
        end
    end
    d ← d_list(i_min);
    s ← s_min + i_min;
    return d, s;
end
```

4.3.4 Stimulation control

The simplest option to turn the phase tracking module into a closed-loop system would have been to provide a short digital pulse on one of the Intan digital output channels each time the target phase is reached or approached. In our experience optogenetics light sources from various vendors differ greatly in their support for waveform customisation, however, they invariably support direct digital control from an external device. It was therefore deemed preferable to equip the custom interface with a simple waveform synthesis module of its own, to ensure protocols are transferable between research groups. The simple module allows for delivering stimulation centred around the target phase without any significant added latency. Waveforms can be defined based on pulse length and sub-pulse timing for burst-like patterns (parameters shown in Fig. 4.12). To moderate stimulation rates, an optional 'timeout' parameter was included to prevent further stimulation pulses in a desired post-stimulus interval. This can be useful as a general sanity constraint, i.e. the stimulation rate should not normally exceed the nominal frequency of the target oscillation, or it can also prove useful when trying to study evoked potentials lasting multiple oscillatory cycles. The aforementioned features of the stimulation module are demonstrated in Fig. 4.15.

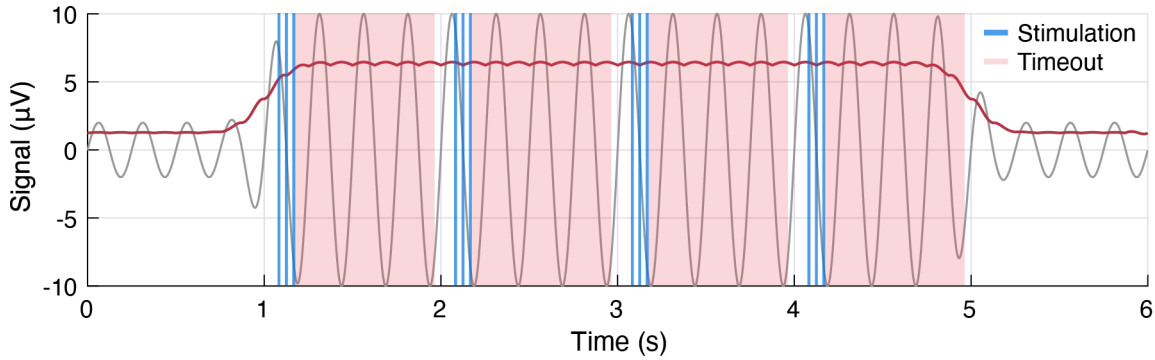


Figure 4.15. Conditional stimulation capabilities of the custom Intan system. The test signal consists of a pure 4 Hz oscillation, that transitions between a 5 μV_{pp} and a 20 μV_{pp} magnitude state, forming a burst. Stimulation (blue) can be targeted at arbitrary phases of the oscillation (90°), with the option of defining burst-like stimulation patterns. The delivered maximum rate of stimulation can be constrained using a 'timeout' parameter (blocked periods in pale red). The magnitude estimate (post-filtering), shown in red, can be used to gate stimulation with a threshold (5 µV).

4.3.5 Magnitude estimate

Observing the magnitude estimate provided by the algorithm (Fig. 4.5), we can conclude it is subject to high-frequency noise, as compared to reference calculations, which can be corrected via low-pass filtering at the expense of added delay. To keep computational costs low an exponential moving average filter was implemented. The properties of this filter have been previously explored in detail in Chapter 2. In the GUI users are provided with the option to record both the filtered and unfiltered magnitude estimates, as well as the ability to impose a threshold on the filtered magnitude estimate to either enable or disable the delivery of phasic stimulation (Fig. 4.15).

4.4 Validation

In this section, we describe a novel, closed-loop optogenetics-based research protocol. We present pilot data to showcase the real-world performance of the ESE model, and to serve as a demonstration of the types of research questions the system makes possible to investigate. Based on our aims, our questions concern whether the phase-tracking method is viable under optogenetic stimulation, and how accurately it performs at low-SNR oscillatory targets.

4.4.1 Experiment overview

To serve as *in vivo* test case for our work, we took advantage of an ongoing neurostimulation study conducted by the Sharott group at the University of Oxford. The experiment in question used transgenic mouse models to optogenetically activate two distinct cell types of the primary motor cortex (M1), suspected to have opposing effects on macro-scale network activity: an inhibitory interneuron subtype (RBP4), and a line of parvalbumin-expressing pyramidal neurons (PV). Viral targeting selectively introduced Channelrhodopsin-2 (ChR2) to these cell types, an opsin causing membrane depolarisation, when exposed to light thus contributing to neural firing. Preliminary data showed a modest, but consistently present 6 Hz oscillatory target in M1 while the mice were at rest – an ideal candidate for our evaluation aims. The multiple optogenetic targets conveniently provide an opportunity to reflect on the repeatability of our methods within one experiment.

As the study is ongoing at the time of writing and as such unpublished. The experimental methods, provided by Jessica Myatt, are reported in the next section for completeness. Readers interested in focusing on the validation results may consider skipping ahead to Section 4.4.3.

4.4.2 Experimental methods

Experiments involving animals were conducted in accordance with institutional guidelines and the UK Animals (Scientific Procedures) Act 1986 under personal and project licenses issued by the Home Office following ethical review.

Subjects

Adult male PV-Cre and PV-RBP4 mice (weighing 28–35 g) were housed (initially in groups) in a dedicated environmentally controlled room under a 12-hour light/dark cycle, with food and water provided *ad libitum*.

Anaesthetic protocol

The animals underwent three recovery stereotaxic surgical procedures under deep anaesthesia using isoflurane in oxygen (4 % induction, 0.5–2 % maintenance). Once anaesthesia was successfully induced, local anaesthetic (Marcaine, 2 mg/kg) as well as an opioid analgesic (Vetergesic, 0.08 mg/kg) were administered subcutaneously. All surgical procedures were carried out aseptically.

Viral Injection Surgery

The first surgical procedure involved the preparation of a small cranial window, to allow a unilateral viral injection into the right motor cortex (AP +1.5, ML +1.3, DV +0.8) using a glass micropipette (tip diameter 15–20 µm). A volume of 100 nl of viral suspension was dispensed at a rate of 60 nl/min, at a concentration of 6.5×10^{12} vg/ml (vg = 'viral genomes'). Post injection, the micropipette was left in place for 5–10 minutes before being removed to allow for the viral particles to diffuse. In animals to be transduced with opsin, the virus used was AAV5-Ef1a-DIO-ChR2(E123T/T259C)-eYFP expressing channelrhodopsin-2 (ChR2), whereas in opsin-negative controls the virus AAV5-Ef1a-DIO-eYFP was used (Designed by Deisseroth, Vector Core, University of North Carolina). Post surgery, the animals were provided with a subcutaneous injection of 0.4 ml warm glucose solution and moved to a heated recovery home cage with supplementary food and monitored until regaining ambulatory capability, then further every 24 hours until fully recovered and within 5 % of their original weight.

Headplate surgery

At least 10 days after the viral injection, the animals were implanted with a titanium headplate under deep anaesthesia (as above). The headplate was aligned to the bregma and lambda points of the cranium and affixed with clear dental cement (Super Bond C&B, Parkell). A reference screw was placed in one hemisphere over the cerebellum, usually ipsilaterally to the injection site. Following recovery, animals were singly housed and regularly monitored.

Habituation

Mice were initially handled to become comfortable with the experimenter, then habituated to the recording set up over a period of 5–10 days, by applying head-fixation for increasing amounts of time until the mouse could sit comfortably for an hour. A periodic reward of 20 % sucrose solution was provided throughout. Following each session, the mouse would be held for 10–15 minutes to ensure being comfortable with the experimenter.

Craniotomy

Upon successful habituation, a third and final anaesthetised recovery surgery was carried out to create a small cranial window (2 mm^2) above the recording site. A layer of soft silicone was placed over the cranial window to act as a protective barrier over the exposed tissue.

Experiments

Data acquisition and behavioural testing took place in a dedicated recording room in a total of 3-5 sessions of 3-4 hours each, over the course of 5 days. Upon completion, the mice were anaesthetised using 4 % isoflurane as before, given an intraparietal pentobarbital injection (3 ml, Pentoject, 200 mg/ml), then transcardially perfused with phosphate-buffered saline (PBS) followed by fixative (4 %wt/vol paraformaldehyde in PBS).

Closed-loop recording

Closed-loop experiments were carried out in sessions lasting 2–4 hours, comprising parameter sweeps in blocks of 3–5 minutes per parameter, presented in a pseudo-random order.

Neural signals were acquired using 32-channel silicon microelectrode arrays with 4 linear shanks of 8 electrodes each, 100 µm electrode spacing and 200 µm shank spacing (A4×8-100-200-177-OA32; NeuroNexus, Michigan, USA). The probes were inserted acutely in a head-fixed setup using a stereotaxic frame, at a depth of 900–1300 µm, until unit activity became visible, allowing 10-15 minutes for the tissue to settle.

Signals were amplified and digitised with a 64-channel Intan RHD2164 unit (Intan Technologies, California, USA), and recorded using an RHD2000 USB interface board, configured with modified firmware as described in the 'Development' section previously. The microelectrode array and the amplifier required an adapter for interfacing (Adpt-A32-OM32; NeuroNexus).

For closed-loop operation the digital high-pass filter (for removing channel offset) was configured with a cut-off frequency of 1 Hz. Raw signals were recorded at a rate of 20 kHz, with analogue filters in the instrumentation configured to a wide pass-band of 0.09 Hz–7.6 kHz. The closed-loop target frequency was set to 6 Hz, and the loop-gain was set to 2^{-4}. The optimal downsampling ratio was programmatically selected as 49×, with a reference wave resolution of 68 points per cycle, providing a true target frequency of 6.002 Hz. Recordings were carried out inside a Faraday cage, connected to the ground pin of the headstage amplifier.

The optical stimulus was sourced from a 473 nm solid-state diode-pumped laser (Cobolt 06-MLD; Cobolt Laser, Sweden), transmitted via a multi-mode fibre optic patch cable (FG105UCA, 105 µm core diameter, 0.22 NA; Thorlabs, New Jersey, USA), and delivered from a flat profile optical fibre ending on the second shank of the silicon probe 200 µm above the uppermost electrode contact. Laser power levels were calibrated prior to each recording session using an integrating sphere photodiode power sensor (S140C; Thorlabs). The optical stimulus was modulated digitally using a shutter built into the laser module with a digital signal from the Intan interface board, between zero and a user-configured intensity value.

Stimulation protocol

Stimulation consisting of an uninterrupted pulse delivered around the target phase for a duration of approximately 1/4th of the 6 Hz theta cycle was delivered at four target phases of the theta cycle, at $0, 90, 180, 270^{\circ}$ (refer to Fig. 4.2).

The implementation of the closed-loop system allowed for setting up a lock-out period after stimulation, under which no further stimulation might take place. This parameter was set to 133.333 ms. The value was selected as 80% of the period of a 6 Hz theta cycle as recommended by McNamara et al. [124]. This value allows for ample tolerance to successfully stimulate a subsequent oscillatory cycle even if the current estimate was imprecise while helping to reduce spurious triggering.

Stimulation was initially tested at three laser intensities, $1, 2, 5$ mW, with the 5 mW level abandoned after the first two animals for significantly affecting the broad-band LFP spectra.

4.4.3 Results

To establish whether the ESE method can produce a useful, accurate phase estimate during optogenetic stimulation, we first looked at a set of histograms, placing each stimulation event into a phase-bin calculated via offline zero-phase bandpass filtering ($6.0 \pm 0.5\,\text{Hz}$, 4th order Butterworth) and a Hilbert transform. Results were compiled from $n = 7$ PV-Cre and $n = 1$ RBP4-Cre opsin positive and $n = 4$ opsin-negative PV-Cre animals. At a first glance, the four target phases are visually well-separated and centred around the expected target phases.

Since this study explored targeting four equally spaced phases over an oscillatory cycle, it made sense to analyse achieved accuracy based on the proportion of stimulation delivered within a quarter-cycle ($360°/4$, or $\pm 360°/8 = \pm 45°$) of the intended phase, i.e. being closer to the target than to any neighbouring quarter-phase point. The results of Fig. 4.16, we show that in opsin-expressing animals under unstimulated conditions, and in control animals under both stimulated and unstimulated conditions the ESE method delivered stimulation within a quarter phase of the target $38.3 \pm 2.7\,\%$ of the time (mean $\pm$ SD). The unstimulated and control distributions did not show significant differences under the Watson-Williams test at the $\alpha = 0.05$ level.

The case of stimulated conditions of opsin-expressing animals was more complicated, though. While the apparent accuracy changed compared to the unstimulated baseline, we posit this could be a natural consequence of successful manipulation of oscillatory power – increased power increases regularity, thus making it more likely for the reference method and the ESE phase to align, while under suppressing conditions oscillations are disrupted, and phase becomes less tractable (and meaningful).

Stimulation triggered averages of the narrow-band field potential, presented in Fig. 4.17 also visualise the achieved post hoc phase, further supporting the accuracy of the method. The envelope of the average oscillatory activity appears to decay in either direction about the time of stimulation, due to cycle-to-cycle jitter and slight mismatch from the nominal band-pass centre frequency adding up over time.

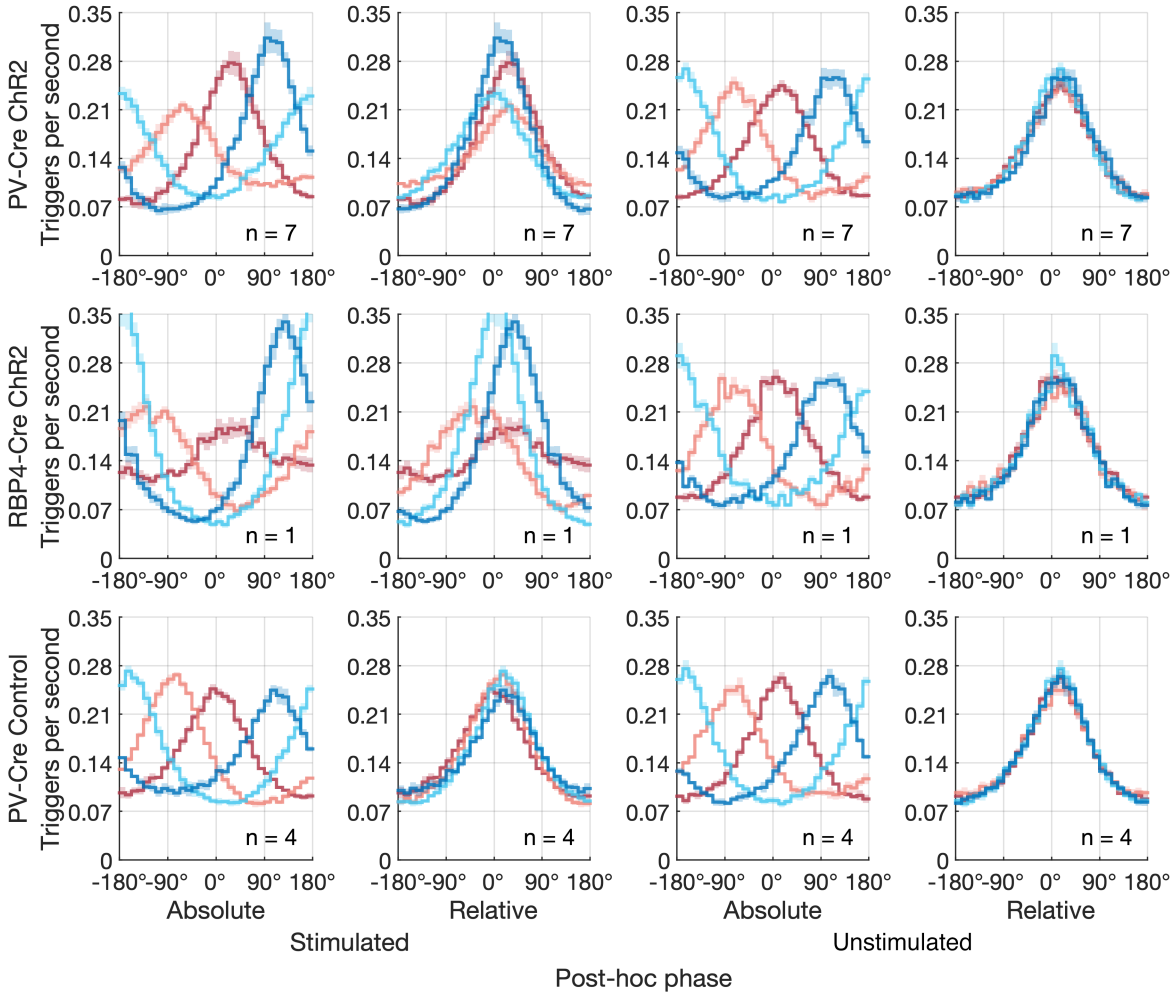


Figure 4.16. Post-hoc phase histograms of stimulation events during closed-loop optogenetic manipulation. Post hoc phase was established by filtering the extracellular potentials in the $6.0 \pm 0.5\,\mathrm{Hz}$ band using an offline, zero-phase 4th order Butterworth filter, then calculating the Hilbert transform. Columns showing the absolute histograms demonstrate clear separation between the four targeted phases of $0, 90, 180, 270°$, while alternate columns showing the distributions normalised to their intended target phases are shown to evaluate overall accuracy. The histograms represent time-normalised frequencies, grouped into 32 bins. Results are shown from three animal types: PV-Cre and RBP4-Cre mice transduced with ChR2, and control PV-Cre mice without ChR2. The results show the average of subject means. Shaded areas represent the between-subject standard error of the mean in PV-Cre animals, and within-subject standard deviation in the single RBP4-Cre animal. Results from the 1 and 2 mW laser intensities were pooled for this figure, as no significant differences were observed when separated. Comparing the results between stimulated and unstimulated conditions, we can conclude that accuracy did not suffer during stimulation. Changes in the histogram peaks are observed in the presence of stimulation due to amplification and suppression of the targeted oscillatory process.

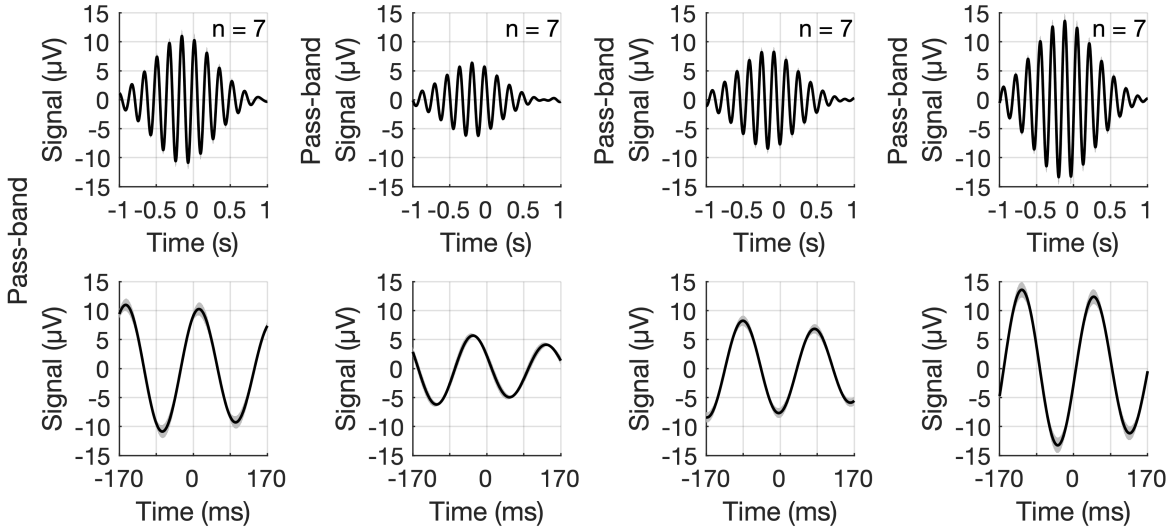


Figure 4.17. Stimulation triggered averages of extracellular potentials in response to closed-loop optogenetic manipulation of ChR2 transduced PV-Cre mice. The four columns separate the four target phases of $0, 90, 180, 270°$ respectively. Zero on the time axis represents the centre-point of the delivered stimulus, and falls on the targeted phase of the theta cycle. Average activity in the 6.0 ± 0.5 Hz band was filtered offline with a zero-phase 4th order Butterworth filter. The two rows differ only in the scale of the time axis. The black lines represent the means of subject-means for each condition. Error bars in the form of light-grey shading represent the standard error of the mean, with the number of animals indicated in the panels. Results from the 1 and 2 mW laser intensities were pooled for this figure, as no significant differences were observed when separated.

4.5 Discussion

In this chapter, we investigated a recently introduced closed-loop algorithm for real-time phase estimation. The established theory behind the operation of McNamara's empirical state estimator method [124] was used to systematically explore the parameter space of the algorithm, making it possible to automate configuring phase tracking for arbitrary oscillatory processes. The development work finally resulted in the creation of a graphical user interface extension for the Intan neural signal recording platform, with the potential to make the method available to the research community in a standardised format.

The ESE method was shown to offer a high degree of accuracy for phase-locked optogenetic stimulation at a new oscillatory target, despite its low SNR (approx. 1 dB above $1/f$). Post hoc analysis revealed that stimulation times on average successfully followed the target phase (Fig. 4.16). Phase at stimulation followed von Mises distributions, with $\sim$40 % of stimuli falling within a quarter-cycle of the target, suggesting three phase divisions might be more appropriate if higher probability of targeting accuracy is required.

While apparent accuracy decreased under stimulation at certain phases (Fig. 4.16), this can be attributed in large part to the successful modulation of oscillatory power. At low power and waning amplitude phase becomes ill-defined, as such the ESE and post hoc calculations cannot be expected to be in agreement.

The data heretofore collected has been confirmed to contain identifiable unit activity in the recordings, which could be used to establish a 'ground truth' for the network phase. For stimulated trials, a template-based approach can be utilised to remove the spike-like stimulation artefacts to facilitate the automatic processing of unit data.

As demonstrated, the ESE method is highly successful in tracking weak oscillatory processes in the presence of neural noise. This advantage of the method can easily turn into a limitation should the algorithm be set up improperly. A phase estimate will always be returned, which may contribute to unreliable findings in case the oscillation targeted is not present. As a preparation for future studies, applying a threshold to the so far unused magnitude estimate is available in the user interface to test for the presence of a stable oscillation.

4.6 Future development

While the pilot experiments show the implementation of the ESE method to be a capable research tool, we can identify several points for continued development, ranging from minor improvements to major extensions of the supported research space.

4.6.1 Control policy refinements

In the PID (proportional-integral-derivative) control notation, the ESE method uses purely proportional control to maintain an accurate state estimate (see Eq. 4.2.3). The main consequence of this is that the estimator is not able to cope with certain kinds of constant residual errors, such as the error arising when the centre frequency of the oscillatory process and the target frequency provided to the model are mismatched. Integral control would aid in correcting such steady-state offsets, adding weight to error persistent over time, effectively allowing for unsupervised adjustments to the target frequency. This is not necessarily a desirable trait for a phase estimator in the context of neural signal processing, however. As seen in the results of McNamara et al. (Fig. 1g) [124], targeted suppression of activity at a specific frequency can inadvertently lead to amplification of nearby frequencies. On successful suppression of the target oscillation, integral feedback would likely end up adjusting the effective target frequency to newly dominant off-target activity, potentially leading to further and further drift of the target frequency, and other unpredictable behaviour.

In the case of other physiological signals, where the target frequency is less stable or predictable, especially at high signal-to-noise ratios, the inclusion of integral control could be desirable. A prominent physiological example would be tremors, where phase-locked muscle stimulation is an active area of study [220].

4.6.2 Dynamic gain control

In the case of optical stimulation artefacts, the ESE method has proven to be highly resilient, at least in part due to the brevity of the individual artefacts (see Fig. 4.17) compared to the dynamically reduced sampling rates. Electrical artefacts, however brief themselves, could not be ignored as easily, as they can easily saturate neural recording systems. The post-saturation recovery adds considerably to the length of signal contamination.– gain dependent on stimulation state instead of blanking. In the work of McNamara et al., electrical artefacts were mitigated by adding a sample-and-hold type blanking module to the phase tracker [124].

Another major class of artefacts to consider are motion artefacts arising from the electrical activity of muscle activation, which might prove problematic in behaviourally active animals. Motion is commonly tracked via accelerometry to exclude recording segments from analysis. Some Intan headstages are available in accelerometry-capable versions, thus motion information is readily available within the FPGA fabric. Activity levels could be thresholded to establish when the presence of motion artefacts is likely, to engage or disengage phase tracking.

A unified means of interacting with all of the artefact types above would be to introduce dynamic adaptation in the ESE loop-gain – based on the severity of the impact of any given artefact type on phase tracking accuracy, the weight given to the recorded input signal compared to the internal oscillator model could be proportionally adapted over time.

4.6.3 Multi-site studies

The presented pilot experiment focused on manipulating local oscillatory activity, that is, stimulating and recording from a singular brain area. Another class of experiments would separate the recording and stimulation sites to study macro-scale functional connectivity and extended network effects. Experiments aimed at manipulating coupling strength between multiple synaptically connected regions could attempt to affect plasticity based on stimulating one area based on the oscillatory phase of another or based on paired-site stimulation for a potentially more prominent effect.

To account for synaptic conduction delays between regions, it would be beneficial to implement a customisable delay element between the phase tracking and stimulation modules, though in some experimental setups this can be bypassed without complications by biasing the target phase. Implementation of a delay element in an embedded system will be primarily memory constrained, mitigation techniques involve reduced temporal granularity for stimulation waveforms, or transitioning to an event-based asynchronous model. An event table can be conveniently restricted in size by natural limits on maximum stimulation frequency and useful range of time delays, e.g. supporting a 1 s delay, while allowing up to 200 Hz stimulation would use 200 units of memory, several orders of magnitude less than a classical sampled signal delay line.

4.6.4 Complex target states

Due to the highly optimised computational complexity of the ESE model, several instances could be deployed in a single Intan system in parallel. To complement the deployment of multiple phase trackers, a 'logical mixer' module should be implemented, to allow for flexible definitions of stimulation conditions via the user interface, including combinations of amplitude and phase states from multiple tracker modules.

Optogenetic manipulation allows both for the simultaneous tagging of multiple brain regions, and multiple cell types with different opsins. Contemporary technology already allows for simultaneous stimulation at multiple light wavelengths *in vivo* [111]. Complex conditional stimulation combined with multi-colour optogenetic setups could open up new avenues for studying phase-amplitude coupling in animal models.

4.7 Contribution

The state estimator method was originally developed, and first implemented for FPGAs by Colin McNamara. The interpretation of the method and subsequent parameter-scape exploration are original works from this thesis. The user interface and its connection to the FPGA module are also my own work. The experiment presented in this chapter was carried out by Jessica Myatt, who kindly provided the description of methods. The preliminary data analysis regarding the accuracy of the phase tracker is my own.

Chapter 5

Discussion

Much of current research focuses on the development of highly advanced classifier algorithms, pursuing 'state-of-the-art' performance at the cost of immense computational complexity [221]. Systems as computationally demanding as conventional neural network algorithms are unlikely to be available in clinical use devices in the near future, and are therefore unable to provide real-term patient benefits [5]. Accordingly, our work focused on designs within the constraints of existing device capabilities, aiming to provide methods with realistic implementation potential in implantable systems in clinical use today. We identified and explored three levels of refinements in increasing technological complexity for safe and reliable in-patient research into advanced neuromodulation therapies. We hope these methods will provide meaningful refinements to therapy with positive impacts on the quality of life of patients across a wide diversity of therapeutic needs.

Though strictly speaking an open-loop system, our first design of a non-invasive phase-locked loop system for the synchronisation of implanted and external neurostimulators aims to enable research into some aspects of neuromodulation classically considered in closed-loop protocols, including the induction of cross-brain neural plasticity [41, 42], and assessing the contribution of multiple network nodes in disease [41, 150]. So far, research into multimodal and multi-site stimulation has required access to the implanted DBS leads, severely restricting experimental time to intraoperative cases, and largely preventing chronic studies. The availability of a device-agnostic, non-invasive system has the potential to greatly facilitate in-human research by enabling the observability of binary stimulation states and stimulation patterns in chronically implanted patients.

The identification and study of secondary modulation targets could offers finer control over the network, potentially improving the effectiveness of DBS or reducing stimulation-related side-effects [132]. By modulating the relative timing of multiple network nodes contributing to the

pathological neural state, multi-site stimulation has the potential to harness neural plasticity in creating lasting therapeutic effects [41]. A next generation of implantable pulse generators capable of multi-site protocols might then be able to drastically improve device longevity, by delivering infrequent plastic stimulation protocols to correct for network state, instead of applying continuous stimulation.

In principle symptoms in disorders of the nervous system should invariably be detectable (and sometimes predictable) through analysis of network states. In practice, however, the complexity of the state space, and of the decoding required [71], will prove prohibitive in most applications in the near future. To provide the most immediate benefit to patients, we should focus on proxy measures of network state, until high channel-count systems become available for human experimental medicine, and decoding technologies mature so real-time processing becomes viable. Our previous work on the DyNeuMo platform [129] offered several avenues to target simple proxy measures of disease state, such as time of day for circadian periodicities of symptom severity, or motional activity states. As part of this thesis, we presented the design and validation work undertaken to expand the platform's capabilities to neural state classification following the classical fast, power-adaptive DBS paradigm [42]. Following validation in a reference model environment, the IPG was adapted to and successfully deployed in a non-human primate experimental setting for closed-loop therapy research.

More generally, the Picostim-DyNeuMo as a platform technology presents the potential to establish a unified translational pipeline for experimental neuroscience from animal research to clinical use, in four stages. Conceptually, research into novel therapy options can begin by 1) exploring novel neural targets or algorithms in animal models, then progress by 2) conducting preliminary human validation under externalised device conditions in clinic, both as part of research studies, and under emergency medical intervention exemptions (such as in a reported case study of treatment-resistant dystonia [129]), providing data to 3) inform formalised clinical trials, and finally 4) bring newly approved therapy modes to existing implant users via software updates throughout the lifetime of the devices. All of these steps are facilitated by adapting the same IPG system to all stages of the research pipeline. In this setting, protocols that progress past a proof-of-concept stage are ensured not to face technological roadblocks on the path to clinical validation; thus data collected at each stage can meaningfully contribute to ethical and clinical decision-making further down the line.

To support research into more complex, emerging methods in neurostimulation methods, we investigated the generalisation of a recently reported algorithm for tracking the phase of neural oscillations in real-time [123, 124]. Oscillatory phase is emerging as a target for more refined modulation of activity [25, 122]. The method was deployed in standard laboratory recording systems via a reversible firmware update, with a graphical user interface provided to facilitate configuration. Parameter selection is aided by our newly developed understanding of the algorithm's operation, and automated parameter optimisation. The performance of the algorithm and the user ecosystem was validated on a new frequency target (theta oscillation) in an optogenetic mouse model. The experiment demonstrated that the state-estimation method (originally demonstrated on beta-band activity) carries over to new oscillatory targets well, and is capable of bi-directional modulation of neural activity. The computational complexity of the algorithm used also makes it particularly well suited for low-power microprocessor deployment.

While aDBS has been suggested to offer significant reductions in the amount of stimulation delivered, providing power savings up to 45 % in patients [150], phasic stimulation could greatly outperform this. Consider the suppression of a 20 Hz target in Parkinson's disease. Targeting an identified 'most suppressing' phase would require 20 Hz stimulation instead of 130 Hz, offering a 6.5× reduction in stimulation [124], without even considering that phasic stimulation could be coupled with power-adaptation as well. As the state-estimator method used for phase tracking also provides an estimate of real-time oscillatory magnitude, the algorithm can already implicitly perform complete aDBS functionality with minimal added signal conditioning. A potential avenue toward clinical adaptation would be to use the state estimator algorithm as a drop-in replacement for the conventional aDBS algorithm in a firmware update to the DyNeuMo platform. First-in-human validation of phase-locked DBS could be enabled via an option to switch between magnitude-based and phasic stimulation control in a research mode, under clinician-supervised experimental use.

While the focus of this thesis has been primarily closed-loop therapeutic technologies, it is easy to see how our developments in low-power methods could have significant impact in refining a broader class of BCI systems. To illustrate with an example, speech neuroprosthetics currently rely on very high complexity decoders, running on powerful computers in an edge-processing setup – while sub-second scale latency is not required, near-real time second-scale decoder output is desirable. For a fully functional system this requires continuous monitoring of cortical activity, in practice necessitating a tethered setup [71]. Recording-capable implantable devices are currently limited in operating time by the high power consumption of the radio

link [142, 143]. These devices could be made compatible with a chronic speech decoding setup by introducing a simple neural activity classifier to detect the presence or absence of speech-intent based on localised cortical activity, initiating the power-intensive wireless stream-out only as needed.

Beyond simple decision making functionality, on board extraction of relevant signal characteristics can significantly reduce the amount of data to be streamed to external processing units, providing for either a greater effective number of recording channel (finite streaming bandwidth), or extend recording time (finite battery life). These improvements will benefit open- and closed-loop systems alike, by facilitating the collection of training data and contributing to refined algorithm performance in the future.

Appendix

Signal power estimation

The following sections provide an overview of the theoretical constraints of trying to measure oscillatory strength (band-power) in real-time.

A1 The uncertainty principle

Readers are likely to already be familiar with the famous uncertainty principle of quantum mechanics, introduced by Heisenberg in 1927, about the impossibility of simultaneously resolving the position and momentum of particles to arbitrary precision [222]. However, uncertainty principles apply to a broad range of other conjugate pairs of variables.

First described by Dennis Gábor in 1946, an uncertainty principle limits our ability to resolve the time and frequency characteristics of signals more generally, with profound implications on signal processing [223]. The relationship, since known as the Gábor limit is stated below:

$$\sigma_t \; \sigma_f > \frac{1}{4\pi} \approx 0.08 \, \text{cycles}$$

where σ_t, σ_f denote the standard deviation of time- and frequency-domain estimates of a signal. Analogous results have been proven for discrete (sampled) signal processing as well [224].

Simply put, the time-frequency uncertainty principle states that signals can either be time- or frequency-localised by an observer, not both (see Fig. A1a). To illustrate, we can think of several real-world phenomena. In the visual domain, we might observe temporary periods of apparent synchronisation between turn signals while waiting at a red light in traffic – extended observation reveals subtle differences in frequency. In the auditory domain, consider how difficult it is to define the pitch of highly time-localised sounds such as claps compared to held-out sounds such as music [223, 225, 226].

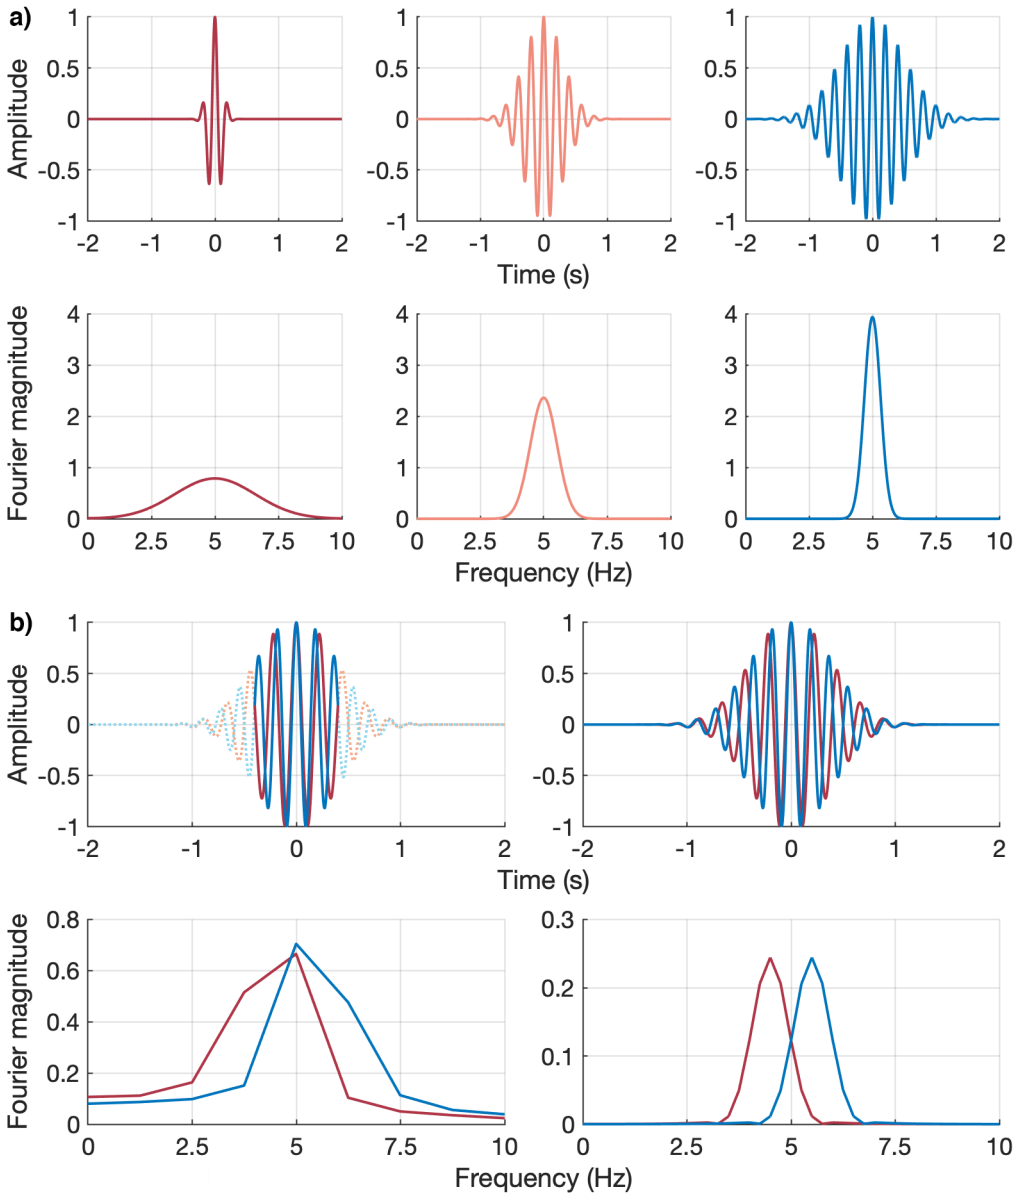


Figure A1. The time-frequency uncertainty principle. **a)** Time and frequency domain representations of a series of 5 Hz analytical test signals modulated by Gaussian envelopes of increasing spreads, demonstrating the inverse relationship between time- and frequency-domain localisation. **b)** Time course and discrete Fourier transform of the same test signals (4.5 Hz and 5.5 Hz) over two window lengths (1 s and 4 s), demonstrating the importance of observation window length for increasing frequency resolution.

A2 Signal transforms

The most well-known mathematical operation for determining the frequency content of a time-domain signal is the Fourier transform. Fig. A1a shows the Fourier transform of three analytical test signals of decreasing time-localisation, demonstrating the constraints imposed by the uncertainty principle. Note how the Fourier transform only provides a single view of frequency content for time series of any length, i.e. it does not help identify changes over time.

When both time and frequency localisation of a signal is important, we can divide our time series into smaller segments, and perform a Fourier transform on each one – this technique is referred to as the short-time Fourier transform (see Fig. A2). While the Fourier transform does generalise to finite length discrete-time (sampled) signals, two notable limitations appear. First, the highest frequency component possible to resolve is going to be dictated by the sampling rate, called the Nyquist rate, at a value of $f_{Nyquist} = f_s/2$. Second, stemming from the uncertainty principle, the smallest frequency difference resolvable will depend not on the sampling rate, but on the length of the signal under transform, as $1/T_{length}$. This phenomenon is illustrated in Fig. A1b).

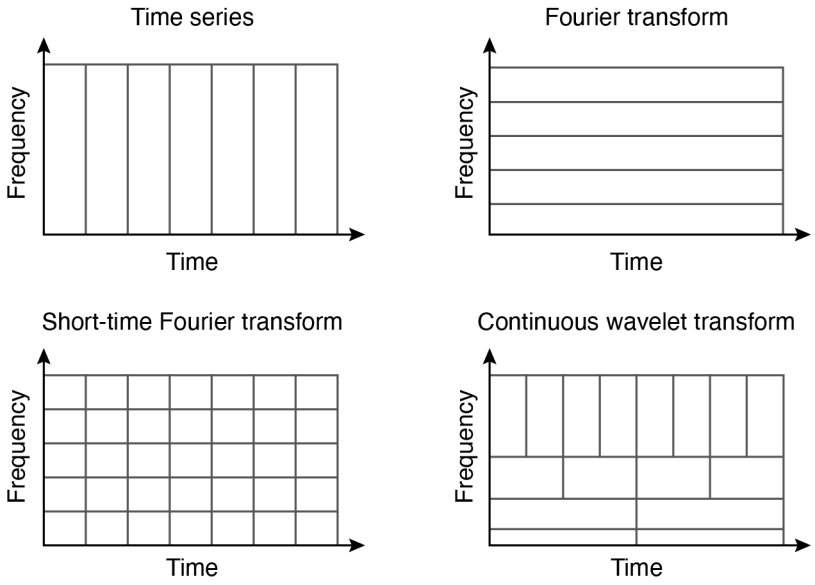


Figure A2. Trade-offs of time-frequency analysis methods. Representations of a signal 'tile' the time-frequency space in various ways. A raw time signal does not distinguish between frequency components present at any point in time. Conversely, the Fourier transform of a signal separates frequency content but does not provide time-localisation. The short-time Fourier transform can be used to perform both time and frequency localisation within the constraints of the uncertainty principle, providing a uniform tiling of the space, favouring either time or frequency resolution. Wavelet transforms provide an alternative tiling of the space, providing finer time-localisation at higher frequencies. Figure based on [227].

Wavelet transforms can provide more information in certain applications by employing an alternative, adaptive tiling of the time-frequency space (Fig. A2) [228]. Other techniques for time-frequency analysis worth mentioning include empirical mode decompositions, such as the Huang-Hilbert transform [229]. Unfortunately, these methods are non-causal (i.e. to evaluate at a given point, samples measured later in time are required), and are thus restricted to *post hoc* processing.

A3 Bandpass filters

Bandpass filters can be used to selectively suppress components of a signal in real time, while preserving frequencies of interest. In the narrow-band limit, passing a signal through a parallel array of band-pass filters (filter bank) and measuring the output magnitude would provide a real-time approximation of the frequency content of the signal. In practice, when we are only interested in specific parts of the frequency spectrum, a small number of filters can prove to be more computationally efficient than performing short-time Fourier transforms. Filtering operations are also subject to the uncertainty principle, however. The settling time of a filter is going to be inversely proportional to the width of its pass-band, i.e. the more selective the filter, the longer it takes to reach steady state. This trade-off can be visualised by observing the step and impulse responses (Fig. A3). Importantly in the context of neurostimulation, impulse-like artefacts can lead to significant enough filter ringing to overshadow the signal of interest when using particularly selective filters [135].

Neural networks, both classical and convolutional, can work in a similar fashion by learning filter-like behaviour, but are limited to high–SNR use at small network sizes, posing computational barrier to implantable device applications [5].

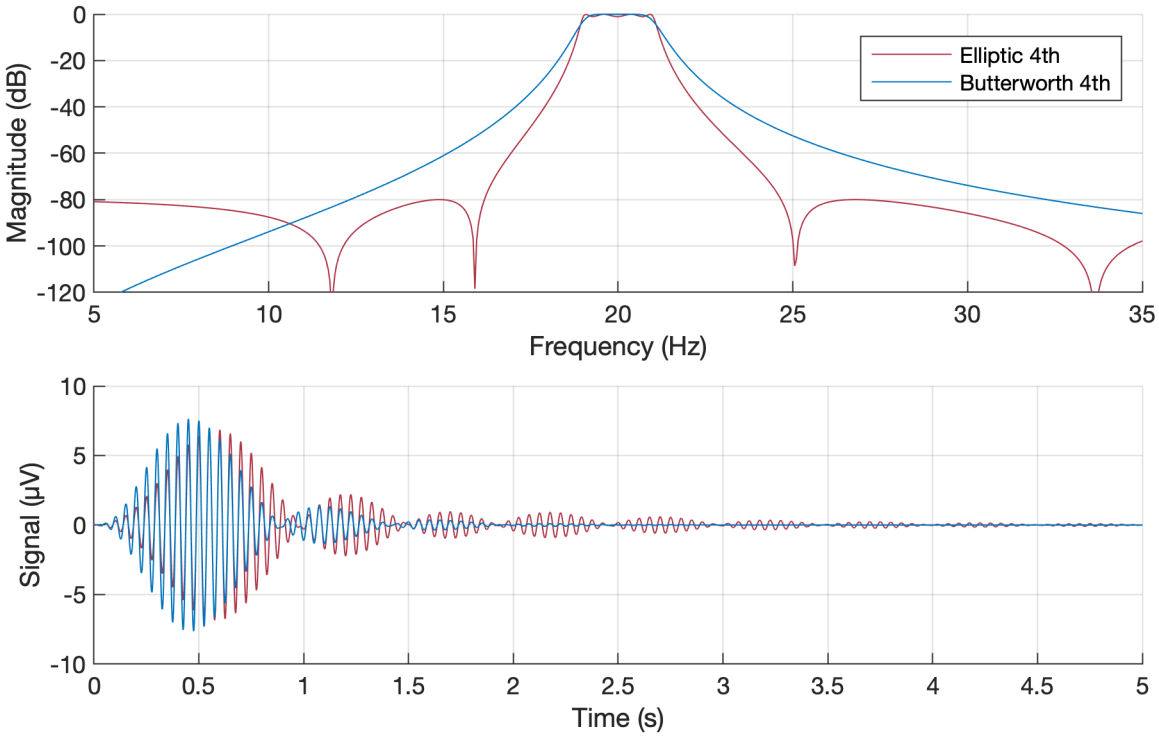


Figure A3. Frequency and impulse response of two narrow band-pass filters. Elliptic filters can provide steeper roll-off than Butterworth filters of equal order (top), thus greater frequency-domain selectivity, at the cost of more extensive ringing apparent in the impulse response (bottom) [135]. Both filters were 4th order, configured with a 20 ± 1 Hz pass-band. The elliptic figure was designed for 80 dB out-of-band attenuation and 1 dB pass-band ripple. The test impulse magnitude was 1 mV, lasting a single sample at $t = 0$ s at a sampling rate of 625 Hz.

Bibliography

[1] M. DiLuca and J. Olesen, "The cost of brain diseases: A burden or a challenge?" *Neuron*, vol. 82, no. 6, pp. 1205–1208, Jun. 2014.

[2] K. A. Ng, E. Greenwald, Y. P. Xu, and N. V. Thakor, "Implantable neurotechnologies: A review of integrated circuit neural amplifiers," *Medical & Biological Engineering & Computing*, vol. 54, no. 1, pp. 45–62, Jan. 2016.

[3] L. R. Hochberg *et al.*, "Reach and grasp by people with tetraplegia using a neurally controlled robotic arm," *Nature*, vol. 485, no. 7398, pp. 372–375, May 2012.

[4] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in *2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. Association for Computing Machinery, Mar. 2021, pp. 610–623.

[5] A. Kavoosi *et al.*, "Computationally efficient neural network classifiers for next generation closed loop neuromodulation therapy – a case study in epilepsy," in *44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Glasgow, Scotland, UK: IEEE, Jul. 2022, pp. 288–291.

[6] B. Haibe-Kains *et al.*, "Transparency and reproducibility in artificial intelligence," *Nature*, vol. 586, no. 7829, pp. E14–E16, Oct. 2020.

[7] M. Pfeiffer and T. Pfeil, "Deep learning with spiking neurons: Opportunities and challenges," *Frontiers in Neuroscience*, vol. 12, p. 774, Oct. 2018.

[8] N. Berndt and H.-G. Holzhütter, "The high energy demand of neuronal cells caused by passive leak currents is not a waste of energy," *Cell Biochemistry and Biophysics*, vol. 67, no. 2, pp. 527–535, Nov. 2013.

[9] M. Dubey *et al.*, "Myelination synchronizes cortical oscillations by consolidating parvalbumin-mediated phasic inhibition," *eLife*, vol. 11, p. e73827, Jan. 2022.

[10] T. A. Fiacco, C. Agulhon, and K. D. McCarthy, "Sorting out astrocyte physiology from pharmacology," *Annual Review of Pharmacology and Toxicology*, vol. 49, no. 1, pp. 151–174, Feb. 2009.

[11] M. Santello, N. Toni, and A. Volterra, "Astrocyte function from information processing to cognition and cognitive impairment," *Nature Neuroscience*, vol. 22, no. 2, pp. 154–166, Feb. 2019.

[12] H.-G. Lee, M. A. Wheeler, and F. J. Quintana, "Function and therapeutic value of astrocytes in neurological diseases," *Nature Reviews Drug Discovery*, vol. 21, no. 5, pp. 339–358, May 2022.

[13] D. Choquet and A. Triller, "The dynamic synapse," *Neuron*, vol. 80, no. 3, pp. 691–703, Oct. 2013.

[14] G. Söhl, S. Maxeiner, and K. Willecke, "Expression and functions of neuronal gap junctions," *Nature Reviews Neuroscience*, vol. 6, no. 3, pp. 191–200, Mar. 2005.

[15] B. A. Patel, M. Arundell, K. H. Parker, M. S. Yeoman, and D. O'Hare, "Microelectrode investigation of neuroneal ageing from a single identified neurone," *Physical Chemistry Chemical Physics*, vol. 12, no. 34, p. 10065, Sep. 2010.

[16] C. J. Stagg, V. Bachtiar, and H. Johansen-Berg, "What are we measuring with GABA magnetic resonance spectroscopy?" *Communicative & Integrative Biology*, vol. 4, no. 5, pp. 573–575, Sep. 2011.

[17] A. Dennis *et al.*, "An ultra-high field magnetic resonance spectroscopy study of post exercise lactate, glutamate and glutamine change in the human brain," *Frontiers in Physiology*, vol. 6, Dec. 2015.

[18] H. Lindén, E. Hagen, S. Leski, E. S. Norheim, K. H. Pettersen, and G. T. Einevoll, "LFPy: a tool for biophysical simulation of extracellular potentials generated by detailed model neurons," *Frontiers in Neuroinformatics*, vol. 7, Jan. 2014.

[19] M. E. Rule and T. O'Leary, "Self-healing codes: How stable neural populations can track continually reconfiguring neural representations," *Proceedings of the National Academy of Sciences*, vol. 119, no. 7, p. e2106692119, Feb. 2022.

[20] Y. Sakurai, "Population coding by cell assemblies — what it really is in the brain," *Neuroscience Research*, vol. 26, no. 1, pp. 1–16, Sep. 1996.

[21] E. Stark and M. Abeles, "Predicting movement from multiunit activity," *Journal of Neuroscience*, vol. 27, no. 31, pp. 8387–8394, Aug. 2007.

[22] C. Heelan *et al.*, "Decoding speech from spike-based neural population recordings in secondary auditory cortex of non-human primates," *Communications Biology*, vol. 2, no. 1, p. 466, Dec. 2019.

[23] J. Martinez, C. Pedreira, M. J. Ison, and R. Quian Quiroga, "Realistic simulation of extracellular recordings," *Journal of Neuroscience Methods*, vol. 184, no. 2, pp. 285–293, Nov. 2009.

[24] S. K. Maran, Y. Cao, D. Heck, and D. Jaeger, "Background firing rate affects the signal transfer of behavior locked input patterns from Purkinje cells to the cerebellar nuclei," *BMC Neuroscience*, vol. 13, no. S1, pp. P112, 1471–2202–13–S1–P112, Jul. 2012.

[25] A. B. Holt *et al.*, "Phase-dependent suppression of beta oscillations in Parkinson's disease patients," *Journal of Neuroscience*, vol. 39, no. 6, pp. 1119–1134, Feb. 2019.

[26] H. Cagnan *et al.*, "Temporal evolution of beta bursts in the parkinsonian cortical and basal ganglia network," *Proceedings of the National Academy of Sciences*, vol. 116, no. 32, pp. 16 095–16 104, Aug. 2019.

[27] E. W. Schomburg, C. A. Anastassiou, G. Buzsáki, and C. Koch, "The spiking component of oscillatory extracellular potentials in the rat hippocampus," *Journal of Neuroscience*, vol. 32, no. 34, pp. 11 798–11 811, Aug. 2012.

[28] M. Zhang, Z. Tang, X. Liu, and J. Van der Spiegel, "Electronic neural interfaces," *Nature Electronics*, vol. 3, no. 4, pp. 191–200, Apr. 2020.

[29] G. Buzsáki, C. A. Anastassiou, and C. Koch, "The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes," *Nature Reviews Neuroscience*, vol. 13, no. 6, pp. 407–420, Jun. 2012.

[30] O. Herreras, "Local field potentials: Myths and misunderstandings," *Frontiers in Neural Circuits*, vol. 10, Dec. 2016.

[31] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, "Backpropagation and the brain," *Nature Reviews Neuroscience*, vol. 21, no. 6, pp. 335–346, Jun. 2020.

[32] Y. Song, T. Lukasiewicz, Z. Xu, and R. Bogacz, "Can the brain do backpropagation? — Exact implementation of backpropagation in predictive coding networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, 2020, pp. 22 566–22 579.

[33] D. O. Hebb, *The organization of behavior; a neuropsychological theory.* New York: Wiley, 1949.

[34] G. Bi and M. Poo, "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type," *Journal of Neuroscience*, vol. 18, no. 24, pp. 10 464–10 472, Dec. 1998.

[35] C. J. Shatz, "The developing brain," *Scientific American*, vol. 267, no. 3, pp. 60–67, Sep. 1992.

[36] C. Keysers and V. Gazzola, "Hebbian learning and predictive mirror neurons for actions, sensations and emotions," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 369, no. 1644, p. 20130175, Jun. 2014.

[37] A. D. Degenhart *et al.*, "Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity," *Nature Biomedical Engineering*, vol. 4, no. 7, pp. 672–685, Jul. 2020.

[38] E. Marder and J.-M. Goaillard, "Variability, compensation and homeostasis in neuron and network function," *Nature Reviews Neuroscience*, vol. 7, no. 7, pp. 563–574, Jul. 2006.

[39] T. O'Leary and D. J. A. Wyllie, "Neuronal homeostasis: Time for a change?" *The Journal of Physiology*, vol. 589, no. 20, pp. 4811–4826, Oct. 2011.

[40] A. Jackson and J. B. Zimmermann, "Neural interfaces for the brain and spinal cord—restoring motor function," *Nature Reviews Neurology*, vol. 8, no. 12, pp. 690–699, Dec. 2012.

[41] K. Udupa *et al.*, "Cortical plasticity induction by pairing subthalamic nucleus deep-brain stimulation and primary motor cortical transcranial magnetic stimulation in Parkinson's disease," *Journal of Neuroscience*, vol. 36, no. 2, pp. 396–404, Jan. 2016.

[42] S. Little *et al.*, "Adaptive deep brain stimulation in advanced Parkinson disease," *Annals of Neurology*, vol. 74, no. 3, pp. 449–457, Sep. 2013.

[43] D. Moran, "Evolution of brain–computer interface: Action potentials, local field potentials and electrocorticograms," *Current Opinion in Neurobiology*, vol. 20, no. 6, pp. 741–745, Dec. 2010.

[44] A. G. Rouse and D. W. Moran, "Neural adaptation of epidural electrocorticographic (EECoG) signals during closed-loop brain computer interface (BCI) tasks," in *31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Minneapolis, MN: IEEE, Sep. 2009, pp. 5514–5517.

[45] D. B. Silversmith *et al.*, "Plug-and-play control of a brain–computer interface through neural map stabilization," *Nature Biotechnology*, vol. 39, no. 3, pp. 326–335, Mar. 2021.

[46] S. Luan *et al.*, "Compact standalone platform for neural recording with real-time spike sorting and data logging," *Journal of Neural Engineering*, vol. 15, no. 4, p. 046014, Aug. 2018.

[47] S. Waldert, T. Pistohl, C. Braun, T. Ball, A. Aertsen, and C. Mehring, "A review on directional information in neural signals for brain-machine interfaces," *Journal of Physiology – Paris*, vol. 103, no. 3–5, pp. 244–254, May 2009.

[48] G. Buzsáki *et al.*, "Tools for probing local circuits: High-density silicon probes combined with optogenetics," *Neuron*, vol. 86, no. 1, pp. 92–105, Apr. 2015.

[49] S. Ramón y Cajal, *Textura del sistema nervioso del hombre y de los vertebrados*, 1st ed. Madrid: Nicolás Moya, 1899–1904, vol. 2, p. 837.

[50] E. Cid and L. M. de la Prida, "Methods for single-cell recording and labeling in vivo," *Journal of Neuroscience Methods*, vol. 325, p. 108354, Sep. 2019.

[51] Y. Peng, F. X. Mittermaier, H. Planert, U. C. Schneider, H. Alle, and J. R. P. Geiger, "High-throughput microcircuit analysis of individual human brains through next-generation multineuron patch-clamp," *eLife*, vol. 8, p. e48178, Nov. 2019.

[52] Q. Tang, M. Brecht, and A. Burgalossi, "Juxtacellular recording and morphological identification of single neurons in freely moving rats," *Nature Protocols*, vol. 9, no. 10, pp. 2369–2381, Oct. 2014.

[53] E. Neher and B. Sakmann, "Single-channel currents recorded from membrane of denervated frog muscle fibres," *Nature*, vol. 260, no. 5554, pp. 799–802, Apr. 1976.

[54] Y. Peng, N. Schöneberg, M. S. Esposito, J. R. Geiger, A. Sharott, and P. Tovote, "Current approaches to characterize micro- and macroscale circuit mechanisms of Parkinson's disease in rodent models," *Experimental Neurology*, vol. 351, p. 114008, May 2022.

[55] R. Toth, A. M. Barth, A. Domonkos, V. Varga, and Z. Somogyvári, "Do not waste your electrodes — principles of optimal electrode geometry for spike sorting," *Journal of Neural Engineering*, vol. 18, no. 4, p. 0460a8, Aug. 2021.

[56] M. Hejazi, W. Tong, M. R. Ibbotson, S. Prawer, and D. J. Garrett, "Advances in carbon-based microfiber electrodes for neural interfacing," *Frontiers in Neuroscience*, vol. 15, p. 658703, Apr. 2021.

[57] K. Woeppel *et al.*, "Explant analysis of Utah electrode arrays implanted in human cortex for brain-computer-interfaces," *Frontiers in Bioengineering and Biotechnology*, vol. 9, p. 759711, Dec. 2021.

[58] C. Rossant *et al.*, "Spike sorting for large, dense electrode arrays," *Nature Neuroscience*, vol. 19, no. 4, pp. 634–641, Apr. 2016.

[59] J. Kleinberg, "An impossibility theorem for clustering," in *15th International Conference on Neural Information Processing Systems (NeurIPS)*. Vancouver, Canada: MIT Press, Jan. 2020, pp. 463–470.

[60] J. P. Neto *et al.*, "Does impedance matter when recording spikes with polytrodes?" *Frontiers in Neuroscience*, vol. 12, p. 715, Oct. 2018.

[61] T. K. T. Nguyen *et al.*, "Closed-loop optical neural stimulation based on a 32-channel low-noise recording system with online spike sorting," *Journal of Neural Engineering*, vol. 11, no. 4, p. 046005, Aug. 2014.

[62] J. J. Jun, C. Mitelut, C. Lai, S. L. Gratiy, C. A. Anastassiou, and T. D. Harris, "Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction," *bioRxiv*, Jan. 2017, doi: 10.1101/101030.

[63] J. G. Kelly and M. J. Hawken, "Quantification of neuronal density across cortical depth using automated 3D analysis of confocal image stacks," *Brain Structure and Function*, vol. 222, no. 7, pp. 3333–3353, Sep. 2017.

[64] O. Marmor *et al.*, "Local vs. volume conductance activity of field potentials in the human subthalamic nucleus," *Journal of Neurophysiology*, vol. 117, no. 6, pp. 2140–2151, Jun. 2017.

[65] S. Hemm and K. Wårdell, "Stereotactic implantation of deep brain stimulation electrodes: a review of technical systems, methods and emerging tools," *Medical & Biological Engineering & Computing*, vol. 48, no. 7, pp. 611–624, Jul. 2010.

[66] D. Satzer, Y. T. Esengul, P. C. Warnke, N. P. Issa, and D. R. Nordli, "SEEG in 3D: Interictal source localization from intracerebral recordings," *Frontiers in Neurology*, vol. 13, p. 782880, Feb. 2022.

[67] C. Herff, D. J. Krusienski, and P. Kubben, "The potential of stereotactic-EEG for brain-computer interfaces: Current progress and future directions," *Frontiers in Neuroscience*, vol. 14, p. 123, Feb. 2020.

[68] T. Kaiju *et al.*, "High spatiotemporal resolution ecog recording of somatosensory evoked potentials with flexible micro-electrode arrays," *Frontiers in Neural Circuits*, vol. 11, Apr. 2017.

[69] D. J. Englot, E. F. Chang, and C. J. Vecht, "Epilepsy and brain tumors," in *Handbook of Clinical Neurology*. Elsevier, 2016, vol. 134, ch. 16, pp. 267–285.

[70] C. Herff *et al.*, "Brain-to-text: Decoding spoken phrases from phone representations in the brain," *Frontiers in Neuroscience*, vol. 9, Jun. 2015.

[71] D. A. Moses *et al.*, "Neuroprosthesis for decoding speech in a paralyzed person with anarthria," *New England Journal of Medicine*, vol. 385, no. 3, pp. 217–227, Jul. 2021.

[72] R. Abiri, S. Borhani, E. W. Sellers, Y. Jiang, and X. Zhao, "A comprehensive review of EEG-based brain–computer interface paradigms," *Journal of Neural Engineering*, vol. 16, no. 1, p. 011001, Feb. 2019.

[73] T. C. Ferree, P. Luu, G. S. Russell, and D. M. Tucker, "Scalp electrode impedance, infection risk, and EEG data quality," *Clinical Neurophysiology*, vol. 112, no. 3, pp. 536–544, Mar. 2001.

[74] T. Choy, E. Baker, and K. Stavropoulos, "Systemic racism in EEG research: Considerations and potential solutions," *Affective Science*, vol. 3, no. 1, pp. 14–20, Mar. 2022.

[75] P. Fiedler, C. Fonseca, E. Supriyanto, F. Zanow, and J. Haueisen, "A high-density 256-channel cap for dry electroencephalography," *Human Brain Mapping*, vol. 43, no. 4, pp. 1295–1308, Mar. 2022.

[76] J. D. Slater, G. P. Kalamangalam, and O. Hope, "Quality assessment of electroencephalography obtained from a "dry electrode" system," *Journal of Neuroscience Methods*, vol. 208, no. 2, pp. 134–137, Jul. 2012.

[77] J. J. Halford *et al.*, "Comparison of a novel dry electrode headset to standard routine EEG in veterans," *Journal of Clinical Neurophysiology*, vol. 33, no. 6, pp. 530–537, Dec. 2016.

[78] M. Hämäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa, "Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain," *Reviews of Modern Physics*, vol. 65, no. 2, pp. 413–497, Apr. 1993.

[79] S. Baillet, "Magnetoencephalography for brain electrophysiology and imaging," *Nature Neuroscience*, vol. 20, no. 3, pp. 327–339, Mar. 2017.

[80] T. M. Tierney *et al.*, "Optically pumped magnetometers: From quantum origins to multi-channel magnetoencephalography," *NeuroImage*, vol. 199, pp. 598–608, Oct. 2019.

[81] R. M. Hill *et al.*, "Multi-channel whole-head OPM-MEG: Helmet design and a comparison with a conventional system," *NeuroImage*, vol. 219, p. 116995, Oct. 2020.

[82] R. B. Buxton, "The thermodynamics of thinking: Connections between neural activity, energy metabolism and blood flow," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 376, no. 1815, p. 20190624, Jan. 2021.

[83] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann, "Neurophysiological investigation of the basis of the fMRI signal," *Nature*, vol. 412, pp. 150–157, Jul. 2001.

[84] Z. Shi *et al.*, "On the relationship between mri and local field potential measurements of spatial and temporal variations in functional connectivity," *Scientific Reports*, vol. 9, no. 1, p. 8871, Dec. 2019.

[85] L. Lyon, "Dead salmon and voodoo correlations: Should we be sceptical about functional MRI?" *Brain*, vol. 140, no. 8, pp. e53–e53, Aug. 2017.

[86] V. Quaresima and M. Ferrari, "Functional near-infrared spectroscopy (fNIRS) for assessing cerebral cortex function during human behavior in natural/social situations: A concise review," *Organizational Research Methods*, vol. 22, no. 1, pp. 46–68, Jan. 2019.

[87] D. Yang, Y.-I. Shin, and K.-S. Hong, "Systemic review on transcranial electrical stimulation parameters and EEG/fNIRS features for brain diseases," *Frontiers in Neuroscience*, vol. 15, p. 629323, Mar. 2021.

[88] H. Akkad *et al.*, "Increasing human motor skill acquisition by driving theta-gamma coupling," *eLife*, vol. 10, p. e67355, Nov. 2021.

[89] M. Vöröslakos *et al.*, "Direct effects of transcranial electric stimulation on brain circuits in rats and humans," *Nature Communications*, vol. 9, no. 1, p. 483, Dec. 2018.

[90] A. Liu *et al.*, "Immediate neurophysiological effects of transcranial electrical stimulation," *Nature Communications*, vol. 9, no. 1, p. 5092, Dec. 2018.

[91] V. Beliaeva and R. Polania, "Can low-intensity tACS genuinely entrain neural activity in vivo?" *Brain Stimulation*, vol. 13, no. 6, pp. 1796–1799, Nov. 2020.

[92] P. G. Vieira, M. R. Krause, and C. C. Pack, "tACS entrains neural activity while somatosensory input is blocked," *PLOS Biology*, vol. 18, no. 10, p. e3000834, Oct. 2020.

[93] S. Grover, R. Fayzullina, B. M. Bullard, V. Levina, and R. M. Reinhart, "A meta-analysis suggests that tACS improves cognition in healthy, aging, and psychiatric populations," *Science Translational Medicine*, vol. 15, no. 697, p. eabo2044, May 2023.

[94] B. M. Hampstead, M. Ehmann, and A. Rahman-Filipiak, "Reliable use of silver chloride HD-tDCS electrodes," *Brain Stimulation*, vol. 13, no. 4, pp. 1005–1007, Jul. 2020.

[95] L. Wu, T. Liu, and J. Wang, "Improving the effect of transcranial alternating current stimulation (tACS): A systematic review," *Frontiers in Human Neuroscience*, vol. 15, p. 652393, Jun. 2021.

[96] A. R. Mehta, A. Pogosyan, P. Brown, and J.-S. Brittain, "Montage Matters: The Influence of Transcranial Alternating Current Stimulation on Human Physiological Tremor," *Brain Stimulation*, vol. 8, no. 2, pp. 260–268, Mar. 2015.

[97] M. M. Sorkhabi, J. Frounchi, P. Shahabi, and H. Veladi, "Deep-brain transcranial stimulation: A novel approach for high 3-D resolution," *IEEE Access*, vol. 5, pp. 3157–3171, Mar. 2017.

[98] M. M. Sorkhabi and T. Denison, "A neurostimulator system for real, sham, and multi-target transcranial magnetic stimulation," *Journal of Neural Engineering*, vol. 19, no. 2, p. 026035, Apr. 2022.

[99] J. K. Krauss *et al.*, "Technology of deep brain stimulation: current status and future directions," *Nature Reviews Neurology*, vol. 17, no. 2, pp. 75–87, Feb. 2021.

[100] F. Alonso, M. Latorre, N. Göransson, P. Zsigmond, and K. Wårdell, "Investigation into deep brain stimulation lead designs: A patient-specific simulation study," *Brain Sciences*, vol. 6, no. 3, p. 39, Sep. 2016.

[101] R. A. Green, S. Baek, L. A. Poole-Warren, and P. J. Martens, "Conducting polymer-hydrogels for medical electrode applications," *Science and Technology of Advanced Materials*, vol. 11, no. 1, p. 014107, Feb. 2010.

[102] S. F. Cogan, "Neural stimulation and recording electrodes," *Annual Review of Biomedical Engineering*, vol. 10, no. 1, pp. 275–309, Aug. 2008.

[103] D. Satzer, D. Lanctin, L. E. Eberly, and A. Abosch, "Variation in deep brain stimulation electrode impedance over years following electrode implantation," *Stereotactic and Functional Neurosurgery*, vol. 92, no. 2, pp. 94–102, May 2014.

[104] D. R. Merrill, M. Bikson, and J. G. Jefferys, "Electrical stimulation of excitable tissue: design of efficacious and safe protocols," *Journal of Neuroscience Methods*, vol. 141, no. 2, pp. 171–198, Feb. 2005.

[105] D. N. Anderson, G. Duffley, J. Vorwerk, A. D. Dorval, and C. R. Butson, "Anodic stimulation misunderstood: preferential activation of fiber orientations with anodic waveforms in deep brain stimulation," *Journal of Neural Engineering*, vol. 16, no. 1, p. 016026, Feb. 2019.

[106] E. J. Tehovnik, A. S. Tolias, F. Sultan, W. M. Slocum, and N. K. Logothetis, "Direct and indirect activation of cortical neurons by electrical microstimulation," *Journal of Neurophysiology*, vol. 96, no. 2, pp. 512–521, Aug. 2006.

[107] R. Ramasubbu, S. Lang, and Z. H. T. Kiss, "Dosing of electrical parameters in deep brain stimulation (DBS) for intractable depression: A review of clinical studies," *Frontiers in Psychiatry*, vol. 9, p. 302, Jul. 2018.

[108] O. Yizhar, L. Fenno, T. Davidson, M. Mogri, and K. Deisseroth, "Optogenetics in neural systems," *Neuron*, vol. 71, no. 1, pp. 9–34, Jul. 2011.

[109] J.-A. Sahel *et al.*, "Partial recovery of visual function in a blind patient after optogenetic therapy," *Nature Medicine*, vol. 27, no. 7, pp. 1223–1229, Jul. 2021.

[110] A. Adamantidis *et al.*, "Optogenetics: 10 years after ChR2 in neurons—views from the community," *Nature Neuroscience*, vol. 18, no. 9, pp. 1202–1212, Sep. 2015.

[111] K. Kampasi *et al.*, "Dual color optogenetic control of neural populations using low-noise, multishank optoelectrodes," *Microsyst. Nanoeng*, vol. 4, no. 10, Dec. 2018.

[112] H. Yasunaga *et al.*, "Development of a neural probe integrated with high-efficiency MicroLEDs for in vivo application," *Japanese Journal of Applied Physics*, vol. 60, no. 1, p. 016503, Jan. 2021.

[113] A. E. Mendrela *et al.*, "A miniature headstage for high resolution closed-loop optogenetics," in *2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. Torino, Italy: IEEE, Oct. 2017, pp. 1–4.

[114] A. E. Mendrela *et al.*, "A high-resolution opto-electrophysiology system with a miniature integrated headstage," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 5, pp. 1065–1075, Oct. 2018.

[115] G. Arias-Gil, F. W. Ohl, K. Takagaki, and M. T. Lippert, "Measurement, modeling, and prediction of temperature rise due to optogenetic brain stimulation," *Neurophotonics*, vol. 3, no. 4, p. 045007, Nov. 2016.

[116] S. Senova *et al.*, "Experimental assessment of the safety and potential efficacy of high irradiance photostimulation of brain tissues," *Scientific Reports*, vol. 7, no. 1, p. 43997, Apr. 2017.

[117] K. Paralikar *et al.*, "An implantable optical stimulation delivery system for actuating an excitable biosubstrate," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 321–332, Jan. 2011.

[118] Á. C. Horváth *et al.*, "Infrared neural stimulation and inhibition using an implantable silicon photonic microdevice," *Microsystems & Nanoengineering*, vol. 6, no. 1, p. 44, Dec. 2020.

[119] I. Osorio, M. G. Frei, and S. B. Wilkinson, "Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset," *Epilepsia*, vol. 39, no. 6, pp. 615–627, Jun. 1998.

[120] A. Schulze-Bonhage, "Long-term outcome in neurostimulation of epilepsy," *Epilepsy & Behavior*, vol. 91, pp. 25–29, Feb. 2019.

[121] S. Little and P. Brown, "The functional role of beta oscillations in Parkinson's disease," *Parkinsonism & Related Disorders*, vol. 20, pp. S44–S48, Jan. 2014.

[122] H. Cagnan *et al.*, "Stimulating at the right time: Phase-specific deep brain stimulation," *Brain*, vol. 140, no. 1, pp. 132–145, Jan. 2017.

[123] C. G. McNamara, M. Rothwell, and A. Sharott, "Phase-dependent closed-loop modulation of neural oscillations in vivo," *bioRxiv*, May 2020, doi: 10.1101/2020.05.21.102335.

[124] C. G. McNamara, M. Rothwell, and A. Sharott, "Stable, interactive modulation of neuronal oscillations produced through brain-machine equilibrium," *Cell Reports*, vol. 41, no. 6, p. 111616, Nov. 2022.

[125] M. O. Baud *et al.*, "Multi-day rhythms modulate seizure risk in epilepsy," *Nature Communications*, vol. 9, no. 1, p. 88, Dec. 2018.

[126] R. Toth *et al.*, "DyNeuMo Mk-2: An investigational circadian-locked neuromodulator with responsive stimulation for applied chronobiology," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. Toronto, ON, Canada: IEEE, Oct. 2020, pp. 3433–3440.

[127] M. Zamora *et al.*, "Case report: Embedding "digital chronotherapy" into medical devices – a canine validation for controlling status epilepticus through multi-scale rhythmic brain stimulation," *Frontiers in Neuroscience*, vol. 15, p. 734265, Sep. 2021.

[128] A. Gunduz *et al.*, "Adding wisdom to 'smart' bioelectronic systems: A design framework for physiologic control including practical examples," *Bioelectronic Medicine*, vol. 2, no. 1, pp. 29–41, Mar. 2019.

[129] M. Zamora *et al.*, "DyNeuMo Mk-1: Design and pilot validation of an investigational motion-adaptive neurostimulator with integrated chronotherapy," *Experimental Neurology*, vol. 351, p. 113977, May 2022.

[130] J. E. Fleming, E. Dunn, and M. M. Lowery, "Simulation of closed-loop deep brain stimulation control schemes for suppression of pathological beta oscillations in Parkinson's disease," *Frontiers in Neuroscience*, vol. 14, p. 166, Mar. 2020.

[131] M. Arlotti *et al.*, "A new implantable closed-loop clinical neural interface: First application in Parkinson's disease," *Frontiers in Neuroscience*, vol. 15, p. 763235, Dec. 2021.

[132] B. D. Swan, L. B. Gasperson, M. O. Krucoff, W. M. Grill, and D. A. Turner, "Sensory percepts induced by microwire array and DBS microstimulation in human sensory thalamus," *Brain Stimulation*, vol. 11, no. 2, pp. 416–422, Mar. 2018.

[133] J. A. Herron, M. C. Thompson, T. Brown, H. J. Chizeck, J. G. Ojemann, and A. L. Ko, "Cortical brain-computer interface for closed-loop deep brain stimulation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 2180–2187, Nov. 2017.

[134] M. J. Cook *et al.*, "Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: A first-in-man study," *The Lancet Neurology*, vol. 12, no. 6, pp. 563–571, Jun. 2013.

[135] J. Ansó *et al.*, "Concurrent stimulation and sensing in bi-directional brain interfaces: A multi-site translational experience," *Journal of Neural Engineering*, vol. 19, no. 2, p. 026025, Mar. 2022.

[136] T. D. Y. Kozai and A. L. Vazquez, "Photoelectric artefact from optogenetics and imaging on microelectrodes and bioelectronics: New challenges and opportunities," *Journal of Materials Chemistry B*, vol. 3, no. 25, pp. 4965–4978, Apr. 2015.

[137] R. R. Harrison, "A versatile integrated circuit for the acquisition of biopotentials," in *2007 IEEE Custom Integrated Circuits Conference (CICC)*. San Jose, CA: IEEE, Sep. 2007, pp. 115–122.

[138] M. Arlotti, L. Rossi, M. Rosa, S. Marceglia, and A. Priori, "An external portable device for adaptive deep brain stimulation (aDBS) clinical research in advanced Parkinson's Disease," *Medical Engineering & Physics*, vol. 38, no. 5, pp. 498–505, May 2016.

[139] A. Goyal *et al.*, "The development of an implantable deep brain stimulation device with simultaneous chronic electrophysiological recording and stimulation in humans," *Biosensors and Bioelectronics*, vol. 176, p. 112888, Mar. 2021.

[140] T. Hashimoto, C. M. Elder, and J. L. Vitek, "A template subtraction method for stimulus artifact removal in high-frequency deep brain stimulation," *Journal of Neuroscience Methods*, vol. 113, no. 2, pp. 181–186, Jan. 2002.

[141] A. Sette *et al.*, "Battery longevity of neurostimulators in Parkinson disease: A historic cohort study," *Brain Stimulation*, vol. 12, no. 4, pp. 851–857, Jul. 2019.

[142] M. T. Salam, H. Kassiri, N. Soltani, H. He, J. L. Perez Velazquez, and R. Genov, "Tradeoffs between wireless communication and computation in closed-loop implantable devices," in *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*. Montreal, QC: IEEE, May 2016, pp. 1838–1841.

[143] A. Zhou *et al.*, "A wireless and artefact-free 128-channel neuromodulation device for closed-loop stimulation and recording in non-human primates," *Nature Biomedical Engineering*, vol. 3, no. 1, pp. 15–26, Dec. 2018.

[144] S. Baldassano *et al.*, "Cloud computing for seizure detection in implanted neural devices," *Journal of Neural Engineering*, vol. 16, no. 2, p. 026016, Apr. 2019.

[145] S. Koganemaru *et al.*, "Gait-synchronized rhythmic brain stimulation improves poststroke gait disturbance: A pilot study," *Stroke*, vol. 50, no. 11, pp. 3205–3212, Nov. 2019.

[146] G. Di Pino *et al.*, "Modulation of brain plasticity in stroke: A novel model for neurorehabilitation," *Nature Reviews Neurology*, vol. 10, no. 10, pp. 597–608, Oct. 2014.

[147] Y. Inukai *et al.*, "Comparison of three non-invasive transcranial electrical stimulation methods for increasing cortical excitability," *Frontiers in Human Neuroscience*, vol. 10, Dec. 2016.

[148] R. Toth, A. B. Holt, M. Benjaber, A. Sharott, and T. Denison, "Frequency and phase synchronization in distributed (implantable-transcutaneous) neural interfaces," in *41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Berlin, Germany: IEEE, Jul. 2019, pp. 3831–3834.

[149] M. Kiss, P. Földesy, and Z. Fekete, "Optimization of a Michigan-type silicon microprobe for infrared neural stimulation," *Sensors and Actuators B: Chemical*, vol. 224, pp. 676–682, Mar. 2016.

[150] N. C. Swann *et al.*, "Adaptive deep brain stimulation for Parkinson's disease using motor cortex sensing," *Journal of Neural Engineering*, vol. 15, no. 4, p. 046006, Aug. 2018.

[151] S. Little and P. Brown, "Debugging adaptive deep brain stimulation for Parkinson's disease," *Movement Disorders*, vol. 35, no. 4, pp. 555–561, Apr. 2020.

[152] L. M. Ward, "Synchronous neural oscillations and cognitive processes," *Trends in Cognitive Sciences*, vol. 7, no. 12, pp. 553–559, Dec. 2003.

[153] P. Fries, J. H. Reynolds, A. E. Rorie, and R. Desimone, "Modulation of oscillatory neuronal synchronization by selective visual attention," *Science*, vol. 291, no. 5508, pp. 1560–1563, Feb. 2001.

[154] P. Lakatos, G. Karmos, A. D. Mehta, I. Ulbert, and C. E. Schroeder, "Entrainment of neuronal oscillations as a mechanism of attentional selection," *Science*, vol. 320, no. 5872, pp. 110–113, Apr. 2008.

[155] G. Buzsáki and J. J. Chrobak, "Temporal structure in spatially organized neuronal ensembles: A role for interneuronal networks," *Current Opinion in Neurobiology*, vol. 5, no. 4, pp. 504–510, Aug. 1995.

[156] G. Buzsáki and A. Draguhn, "Neuronal oscillations in cortical networks," *Science*, vol. 304, no. 5679, pp. 1926–1929, Jun. 2004.

[157] M. J. Kahana, R. Sekuler, J. B. Caplan, M. Kirschen, and J. R. Madsen, "Human theta oscillations exhibit task dependence during virtual maze navigation," *Nature*, vol. 399, no. 6738, pp. 781–784, Jun. 1999.

[158] A. K. Engel and P. Fries, "Beta-band oscillations — signalling the status quo?" *Current Opinion in Neurobiology*, vol. 20, no. 2, pp. 156–165, Apr. 2010.

[159] S. D. Muthukumaraswamy, "Functional properties of human primary motor cortex gamma oscillations," *Journal of Neurophysiology*, vol. 104, no. 5, pp. 2873–2885, Nov. 2010.

[160] P. Fries, "Rhythms for cognition: Communication through coherence," *Neuron*, vol. 88, no. 1, pp. 220–235, Oct. 2015.

[161] I. Grothe, S. D. Neitzel, S. Mandon, and A. K. Kreiter, "Switching neuronal inputs by differential modulations of gamma-band phase-coherence," *Journal of Neuroscience*, vol. 32, no. 46, pp. 16 172–16 180, Nov. 2012.

[162] C. Bosman *et al.*, "Attentional stimulus selection through selective synchronization between monkey visual areas," *Neuron*, vol. 75, no. 5, pp. 875–888, Sep. 2012.

[163] J.-M. Schoffelen, R. Oostenveld, and P. Fries, "Neuronal coherence as a mechanism of effective corticospinal interaction," *Science*, vol. 308, no. 5718, pp. 111–113, Apr. 2005.

[164] S. Zanos, I. Rembado, D. Chen, and E. E. Fetz, "Phase-locked stimulation during cortical beta oscillations produces bidirectional synaptic plasticity in awake monkeys," *Current Biology*, vol. 28, no. 16, pp. 2515–2526.e4, Aug. 2018.

[165] Y.-C. Wu, Y.-S. Liao, W.-H. Yeh, S.-F. Liang, and F.-Z. Shaw, "Directions of deep brain stimulation for epilepsy and Parkinson's disease," *Frontiers in Neuroscience*, vol. 15, p. 680938, Jun. 2021.

[166] B. Wingeier, T. Tcheng, M. M. Koop, B. C. Hill, G. Heit, and H. M. Bronte-Stewart, "Intra-operative STN DBS attenuates the prominent beta rhythm in the STN in Parkinson's disease," *Experimental Neurology*, vol. 197, no. 1, pp. 244–251, Jan. 2006.

[167] A. A. Kuhn *et al.*, "High-frequency stimulation of the subthalamic nucleus suppresses oscillatory activity in patients with Parkinson's disease in parallel with improvement in motor performance," *Journal of Neuroscience*, vol. 28, no. 24, pp. 6165–6173, Jun. 2008.

[168] N. G. Laxpati, W. S. Kasoff, and R. E. Gross, "Deep brain stimulation for the treatment of epilepsy: Circuits, targets, and trials," *Neurotherapeutics*, vol. 11, no. 3, pp. 508–526, Jul. 2014.

[169] A. A. Kühn *et al.*, "Pathological synchronisation in the subthalamic nucleus of patients with Parkinson's disease relates to both bradykinesia and rigidity," *Experimental Neurology*, vol. 215, no. 2, pp. 380–387, Feb. 2009.

[170] S. Stanslaski *et al.*, "Design and validation of a fully implantable, chronic, closed-loop neuromodulation device with concurrent sensing and stimulation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 20, no. 4, pp. 410–421, Jul. 2012.

[171] J. C. Jackson, R. Corey, G. Loxtercamp, S. Stanslaski, H. Orser, and T. Denison, "Computationally efficient, configurable, causal, real-time phase detection applied to local field potential oscillations," in *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*. Montpellier, France: IEEE, Apr. 2015, pp. 942–947.

[172] S. A. Shah, G. Tinkhauser, C. C. Chen, S. Little, and P. Brown, "Parkinsonian tremor detection from subthalamic nucleus local field potentials for closed-loop deep brain stimulation," in *40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Honolulu, HI: IEEE, Jul. 2018, pp. 2320–2324.

[173] D. Bourget *et al.*, "An implantable, rechargeable neuromodulation research tool using a distributed interface and algorithm architecture," in *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*. Montpellier, France: IEEE, Apr. 2015, pp. 61–65.

[174] A. Koralek, R. Costa, and J. Carmena, "Temporally precise cell-specific coherence develops in corticostriatal networks during learning," *Neuron*, vol. 79, no. 5, pp. 865–872, Sep. 2013.

[175] C. Constantoyannis and B. Heilbron, "Electrocardiogram artifacts caused by deep brain stimulation," *Canadian Journal of Neurological Sciences*, vol. 31, no. 3, pp. 343–346, Aug. 2004.

[176] Y. Bader and J. Weinstock, "Successful implantation of a subcutaneous cardiac defibrillator in a patient with a preexisting deep brain stimulator," *Heart Rhythm Case Reports*, vol. 1, no. 4, pp. 241–244, Jul. 2015.

[177] A. Boutet *et al.*, "Predicting optimal deep brain stimulation parameters for Parkinson's disease using functional MRI and machine learning," *Nature Communications*, vol. 12, no. 1, p. 3043, Dec. 2021.

[178] L. Almeida *et al.*, "Deep brain stimulation battery longevity: Comparison of monopolar versus bipolar stimulation modes," *Movement Disorders Clinical Practice*, vol. 3, no. 4, pp. 359–366, Jul. 2016.

[179] F. Frohlich and J. Riddle, "Conducting double-blind placebo-controlled clinical trials of transcranial alternating current stimulation (tACS)," *Translational Psychiatry*, vol. 11, no. 1, p. 284, May 2021.

[180] A. Antal and W. Paulus, "Transcranial alternating current stimulation (tACS)," *Frontiers in Human Neuroscience*, vol. 7, no. 317, Jun. 2013.

[181] D. K. Morgan and G. W. Steudel, "COS/MOS phase-locked-loop – a versatile building block for micro-power digital and analog applications," *RCA Engineer*, vol. 18, pp. 69–76, Sep. 1972.

[182] S. Miocinovic, S. Somayajula, S. Chitnis, and J. L. Vitek, "History, applications, and mechanisms of deep brain stimulation," *JAMA Neurology*, vol. 70, no. 2, pp. 163–171, Feb. 2013.

[183] C. Sarica *et al.*, "Implantable pulse generators for deep brain stimulation: Challenges, complications, and strategies for practicality and longevity," *Frontiers in Human Neuroscience*, vol. 15, p. 708481, Aug. 2021.

[184] K. den Dulk, L. Bouwels, F. Lindemans, I. Rankin, P. Brugada, and H. J. J. Wellens, "The Activitrax rate responsive pacemaker system," *Americal Journal of Cardiology*, vol. 61, no. 1, pp. 107–112, Jan. 1988.

[185] M. N. Islam and M. R. Yuce, "Review of Medical Implant Communication System (MICS) band and network," *ICT Express*, vol. 2, no. 4, pp. 188–194, Dec. 2016.

[186] J. J. van Rheede *et al.*, "Diurnal modulation of subthalamic beta oscillatory power in Parkinson's disease patients during deep brain stimulation," *npj Parkinson's Disease*, vol. 8, no. 1, p. 88, Dec. 2022.

[187] T. Bocci *et al.*, "Eight-hours conventional versus adaptive deep brain stimulation of the subthalamic nucleus in Parkinson's disease," *npj Parkinson's Disease*, vol. 7, no. 1, p. 88, Dec. 2021.

[188] E. J. Peterson, D. A. Dinsmoor, D. J. Tyler, and T. J. Denison, "Stimulation artifact rejection in closed-loop, distributed neural interfaces," in *2016 European Solid-State Circuits Conference (ESSCIRC)*. Lausanne, Switzerland: IEEE, Sep. 2016, pp. 233–236.

[189] J. Proakis and D. Manolakis, *Digital Signal Processing – Principles, Algorithms, and Applications*, 4th ed. New York: Pearson, 2006, pp. 653–640.

[190] R. G. Brown, *Smoothing, forecasting and prediction of discrete time series*, 1st ed. Englewood Cliffs, N.J.: Prentice-Hall, 1962, pp. 97–122.

[191] A. G. Rouse *et al.*, "A chronic generalized bi-directional brain-machine interface," *Journal of Neural Engineering*, vol. 8, no. 3, p. 036018, Jun. 2011.

[192] W. J. Youden, "Index for rating diagnostic tests," *Cancer*, vol. 3, no. 1, pp. 32–35, Jan. 1950.

[193] J. L. Baker *et al.*, "Regulation of arousal and performance of a healthy non-human primate using closed-loop central thalamic deep brain stimulation," in *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)*. Baltimore, MD, USA: IEEE, Apr. 2023.

[194] J. L. Baker, J.-W. Ryou, X. F. Wei, C. R. Butson, N. D. Schiff, and K. P. Purpura, "Robust modulation of arousal regulation, performance, and frontostriatal activity through central thalamic deep brain stimulation in healthy nonhuman primates," *Journal of Neurophysiology*, vol. 116, no. 5, pp. 2383–2404, Nov. 2016.

[195] B. A. Perry *et al.*, "Protective cranial implant caps for macaques," *Journal of Neuroscience Methods*, vol. 348, p. 108992, Jan. 2021.

[196] J. Zimmermann, Y. Vazquez, P. W. Glimcher, B. Pesaran, and K. Louie, "Oculomatic: High speed, reliable, and accurate open-source eye tracking for humans and non-human primates," *Journal of Neuroscience Methods*, vol. 270, pp. 138–146, Sep. 2016.

[197] S. Stanslaski *et al.*, "A chronically implantable neural coprocessor for investigating the treatment of neurological disorders," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 6, pp. 1230–1245, Dec. 2018.

[198] S. Marceglia *et al.*, "Double-blind cross-over pilot trial protocol to evaluate the safety and preliminary efficacy of long-term adaptive deep brain stimulation in patients with Parkinson's disease," *BMJ Open*, vol. 12, no. 1, p. e049955, Jan. 2022.

[199] S. Stanslaski, H. Farooqi, D. E. Sanabria, and T. I. Netoff, "Fully closed loop test environment for adaptive implantable neural stimulators using computational models," *Journal of Medical Devices*, vol. 16, no. 3, p. 034501, Sep. 2022.

[200] T. Klausberger and P. Somogyi, "Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations," *Science*, vol. 321, no. 5885, pp. 53–57, Jul. 2008.

[201] J. J. Sermon *et al.*, "Sub-harmonic entrainment of cortical gamma oscillations to deep brain stimulation in Parkinson's disease: Predictions and validation of a patient-specific nonlinear model," *bioRxiv*, Mar. 2022, doi: 10.1101/2022.03.01.482549.

[202] D. Escobar Sanabria *et al.*, "Real-time suppression and amplification of frequency-specific neural activity using stimulation evoked oscillations," *Brain Stimulation*, vol. 13, no. 6, pp. 1732–1742, Nov. 2020.

[203] L. L. Chen, R. Madhavan, B. I. Rapoport, and W. S. Anderson, "Real-time brain oscillation detection and phase-locked stimulation using autoregressive spectral estimation and time-series forward prediction," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 3, pp. 753–762, Mar. 2013.

[204] E. Blackwood, M. Lo, and A. S. Widge, "Continuous Phase Estimation for Phase-Locked Neural Stimulation Using an Autoregressive Model for Signal Prediction," in *40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Honolulu, HI, USA: IEEE, Jul. 2018, pp. 4736–4739.

[205] C. Zrenner, D. Desideri, P. Belardinelli, and U. Ziemann, "Real-time EEG-defined excitability states determine efficacy of TMS-induced plasticity in human motor cortex," *Brain Stimulation*, vol. 11, no. 2, pp. 374–389, Mar. 2018.

[206] N. Schaworonkow, J. Triesch, U. Ziemann, and C. Zrenner, "EEG-triggered TMS reveals stronger brain state-dependent modulation of motor evoked potentials at weaker stimulation intensities," *Brain Stimulation*, vol. 12, no. 1, pp. 110–118, Jan. 2019.

[207] S. R. Schreglmann *et al.*, "Non-invasive suppression of essential tremor via phase-locked disruption of its temporal coherence," *Nature Communications*, vol. 12, no. 1, p. 363, Dec. 2021.

[208] U. Rutishauser, A. Kotowicz, and G. Laurent, "A method for closed-loop presentation of sensory stimuli conditional on the internal brain-state of awake animals," *Journal of Neuroscience Methods*, vol. 215, no. 1, pp. 139–155, Apr. 2013.

[209] Y. Takeuchi, M. Harangozó, L. Pedraza, T. Földi, G. Kozák, and A. Berényi, "Closed-loop stimulation of the medial septum terminates epileptic seizures," *Brain*, vol. 144, no. 3, pp. 885–908, Mar. 2021.

[210] K. Hristova *et al.*, "Medial septal GABAergic neurons reduce seizure duration upon optogenetic closed-loop stimulation," *Brain*, vol. 144, no. 5, pp. 1576–1589, Jun. 2021.

[211] J. H. Siegle and M. A. Wilson, "Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus," *eLife*, vol. 3, p. e03061, Jul. 2014.

[212] L. Marple, "Computing the discrete-time "analytic" signal via FFT," *IEEE Transactions on Signal Processing*, vol. 47, no. 9, pp. 2600–2603, Sep. 1999.

[213] A. Wodeyar, M. Schatza, A. S. Widge, U. T. Eden, and M. A. Kramer, "A state space modeling approach to real-time phase estimation," *eLife*, vol. 10, p. e68803, Sep. 2021.

[214] C. Zrenner, D. Galevska, J. O. Nieminen, D. Baur, M.-I. Stefanou, and U. Ziemann, "The shaky ground truth of real-time phase estimation," *NeuroImage*, vol. 214, p. 116761, Jul. 2020.

[215] J. Proakis and D. Manolakis, *Digital Signal Processing – Principles, Algorithms, and Applications*, 4th ed. New York: Pearson, 2006, pp. 422–425.

[216] J. E. Volder, "The CORDIC trigonometric computing technique," *IRE Transactions on Electronic Computers*, vol. EC-8, no. 3, pp. 330–334, Sep. 1959.

[217] K. Turkowski, "Fixed-point trigonometry with CORDIC iterations," in *Graphics Gems*, 1st ed., A. S. Glassner, Ed., 1990, ch. 10.2, pp. 494–497.

[218] H. Wen and Z. Liu, "Separating fractal and oscillatory components in the power spectrum of neurophysiological signal," *Brain Topography*, vol. 29, no. 1, pp. 13–26, Jan. 2016.

[219] M. Gerster *et al.*, "Separating neural oscillations from aperiodic 1/f activity: Challenges and recommendations," *Neuroinformatics*, vol. 20, no. 4, pp. 991–1012, Oct. 2022.

[220] B. S. Arruda, C. Reis, J. J. Sermon, A. Pogosyan, P. Brown, and H. Cagnan, "Identifying and modulating distinct tremor states through peripheral nerve stimulation in parkinsonian rest tremor," *Journal of NeuroEngineering and Rehabilitation*, vol. 18, no. 1, Dec. 2021.

[221] W. Barry, S. Arcot Desai, T. K. Tcheng, and M. J. Morrell, "A high accuracy electrographic seizure classifier trained using semi-supervised labeling applied to a large spectrogram dataset," *Frontiers in Neuroscience*, vol. 15, p. 667373, Jun. 2021.

[222] W. Heisenberg, "Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik," *Zeitschrift für Physik*, vol. 43, no. 3–4, pp. 172–198, Mar. 1927.

[223] D. Gábor, "Theory of communication. Parts 1–2," *Journal of the Institution of Electrical Engineers – Part III: Radio and Communication Engineering*, vol. 93, no. 26, pp. 429–445, Nov. 1946.

[224] T. Tao, "An uncertainty principle for cyclic groups of prime order," *Mathematical Research Letters*, vol. 12, no. 1, pp. 121–127, Jan. 2005.

[225] R. L. Freyman and D. A. Nelson, "Frequency discrimination as a function of tonal duration and excitation-pattern slopes in normal and hearing-impaired listeners," *Journal of the Acoustical Society of America*, vol. 79, no. 4, pp. 1034–1044, Apr. 1986.

[226] I.-H. Hsieh and K. Saberi, "Imperfect pitch: Gabor's uncertainty principle and the pitch of extremely brief sounds," *Psychonomic Bulletin & Review*, vol. 23, no. 1, pp. 163–171, Feb. 2016.

[227] S. Agrawal and M. H. Dhend, "Wavelet transform based voltage quality improvement in smart grid," in *2016 International Conference on Automatic Control and Dynamic Optimization Techniques (ICACDOT)*. Pune, India: IEEE, Oct. 2016, pp. 289–294.

[228] A. Grossmann and J. Morlet, "Decomposition of hardy functions into square integrable wavelets of constant shape," *SIAM Journal on Mathematical Analysis*, vol. 15, no. 4, pp. 723–736, Jul. 1984.

[229] N. E. Huang *et al.*, "The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time-series analysis," *Proceedings of the Royal Society of London A: Mathematical Physical and Engineering Sciences*, vol. 454, no. 1971, pp. 903–995, Mar. 1998.