FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

# Improving closed-loop neurostimulation: developing novel strategies from the bottom up with microelectrode array technology

**Domingos Leite de Castro**




Programa Doutoral em Engenharia Electrotécnica e de Computadores

Supervisor: Paulo de Castro Aguiar

Co-supervisor: David B. Grayden

July 26, 2024

© Domingos Leite de Castro, 2024

# Improving closed-loop neurostimulation: developing novel strategies from the bottom up with microelectrode array technology

**Domingos Leite de Castro**

Programa Doutoral em Engenharia Electrotécnica e de Computadores

July 26, 2024

# Resumo

Estimulação elétrica neuronal permite a interação direta com o sistema nervoso central, tendo imenso potencial para tratar doenças neurológicos. No entanto, as terapias existentes são baseadas em estimulação fixa, otimizada empiricamente nas sessões clínicas. Esta abordagem de malha aberta não é eficaz para controlar sistemas tão complexos. Controlo de atividade neuronal em malha fechada é uma área ativa de investigação , mas a sua contribuição para aplicações clínicas é quase nulo. Neste trabalho, tentamos abordar este problema a partir das suas bases, utilizando arranjos de microeletrodos (MEAs) para testar diferentes estratégias de controlo com modelos neuronais *in vitro* e *in silico*. Vamos focarnos nas aplicações mais promissoras da neuroestimulação: a ablação da atividade neuronal patológica e a restauro de perceção sensorial perdida.

Quanto ao primeiro problema, implementamos, pela primeira vez em culturas neuronais, um método popular chamado *Delayed Feedback Control* (DFC) para supprimir a sincronização excessiva de neurónios observada em doenças como a doença de Parkinson, tremor essencial ou epilepsia. Os nossos resultados in vitro} mostraram que o DFC não suprime essas oscilações, mas amplfica-as. Apresentamos por isso uma versão melhorada do algoritmo, *Adaptive Delayed Feedback Control* (aDFC), que reduziu robustamente a intensidade da oscilação e a sincronia neuronal.

Quanto à substituição da perceção sensorial perdida, apresentamos uma solução para melhorar as próteses visuais. Os estimuladores de retina existentes não têm em conta a resposta da retina à estimulação, proporcionando resultados insatisfatórios. Aqui apresentamos uma estrutura de *deep learning*, treinada com dados da própria retina, para codificar imagens em padrões de estimulação que provocam a resposta neuronal desejada. A arquitetura codificador-decodificador tira provetio de padrões de estimulação multipolares, alcançando alta resolução espacial na ativação neural.

Finalmente, tendo apresentado novas estratégias de neuromodulação, procuramos explorar um paradigma de computação alternativo para melhorar a eficiência de processadores implantáveis. Testamos um tipo de dispositivo neuromórfico chamado *memristor*, um elemento de computação micrométrico que se comporta como uma sinapse. O memristor monitora os disparos neuronais de uma cultura *in vitro*, e ao detectar uma avalanche de disparos neuronais, ativa estimulação de uma cultura neuronal alvo, em tempo real. Esta prova de conceito demonstra o potencial dos memristores como possivel elemento base de processadores em estimuladores neuronais.

O trabalho apresentado nesta tese apresenta as bases para múltiplas estratégias que podem ser usadas em futuros neuroestimuladores implantáveis para restaurar e substituir a função neuronal.

# Abstract

Electrical neurostimulation allows for direct interaction with the central nervous system and its potential for treating neurological disorders is immense. However, existing therapies rely on fixed stimulation, empirically optimised during clinical sessions. This open-loop approach is not an effective way of controlling the behaviour of such a complex system. Closed-loop control of neuronal activity has been an active field of research for decades but its translation to clinical applications is almost non-existent. Here, we try to approach this problem from the bottom up, using microelectrode arrays (MEAs) to test different control strategies with *in vitro* and *in silico* neuronal models. We focused on the most promising applications of neurostimulation: the disruption of pathological neuronal activity and the replacement of lost sensory perception.

Regarding the first problem, we implemented, for the first time in neuronal cultures, a popular method called Delayed Feedback Control (DFC) meant to disrupt the excessive neuronal synchronization seen in diseases like Parkinson's disease, essential tremor or epilepsy. Our *in vitro* results showed that DFC does not suppress that oscillations but rather exacerbates them. We then present an improved algorithm, Adaptive Delayed Feedback control (aDFC), which robustly decreased the oscillation intensity and overall neuronal synchrony.

As for the replacement of lost sensory perception, we present a solution to improve visual prostheses. Existing retina stimulators do not take into account the retina response to stimulation providing unsatisfactory results. Here we present a deep-learning framework, trained with data from the target retina, to encode images into stimulation patterns that elicit the desired retinal response. Our encoder-decoder architecture exploited multipolar stimulation patterns achieving high spatial resolution of neural activation.

Finally, having presented these novel strategies for neuromodulation, we sought to explore an alternative computing paradigm to improve the efficiency of implantable chips. We tested a type of neuromorphic device called memristor, a micrometric computing element that behaves like a synapse. The memristor monitored the neuronal spikes of an *in vitro* culture and, upon detecting a neuronal burst, triggered stimulation of a target neuronal culture, in real-time. This proof of concept showcases the potential of memristors as the processing building block of neuronal prostheses.

The work presented in this thesis lays the foundations for multiple strategies that can be used in future implantable neurostimulators to restore and replace neuronal function.

# Acknowledgments

These last four years were challenging, arduous and, if you allow me the mouthful, existential-crisis-inducing at times. But as I look back, the only thing in my mind is: *well, this was actually so much fun!* And that was, in great part, due to the people I had to opportunity to share this journey with. I shall now express my heartfelt gratitude, for their lives intersected with mine in various meaningful ways.

First and foremost, I am deeply grateful to **Paulo**. Or, perhaps more accurately, I'm deeply grateful *for* Paulo. I have grown a lot during these last years, and you were one of the most influential people in this new stage of my life (adulthood, they say...). More than a supervisor, you have become someone I admire and strive to make proud – a role model of common sense, clear thinking, respect, integrity and humility, while always playful, jovial and down-to-earth. I thank you for the confidence and intellectual freedom you've always shown towards my work. I feel truly honoured knowing that you hold me to high standards "*It's not bad, but the Domingos I know wouldn't be satisfied with just this, would he?...*"

Your love for science is inspiring and your enthusiasm towards life is contagious. I don't know what the future holds, but if I end up having others working "for me", I will certainly have the motto "*What would Paulo do?*" running in a loop, somewhere in my brain. You have also made me way more nerdy, so we still need to talk about that.

To Professor **David**, my co-supervisor, who always went the extra mile to ensure I was welcomed at the University of Melbourne. And I truly felt welcomed, right from the get-go: from our first tour around the beautiful University campus to the last majestic burger with an amazing view over Melbourne. Your notorious knowledge and expertise were always accompanied by a rigorous and frontal approach, as well as extreme kindness and humility. Your availability was also much appreciated and respected, I have always felt attended to and supported (even when it required after-work intercontinental Zoom meetings!). And it truly was special to see you again in Baltimore, a week after leaving Australia, while I was still emotionally hungover from Melbourne. Touring around the aquarium, toasting with yet another epic burger at the Hard Rock Cafe – the best part of the conference!

Even though it is my name on the cover, this work would not exist without a team behind it. And oh boy, what a team. To **Miguel**, a true buddy. Always prompt to help me in the lab and always ready for adventures outside the lab, be it jumping from a plane, climbing walls, squat competitions or going for 10km runs (yeah about that...I'll join next year!). **Zé**, you and your annoying eternal youth genes. Thank you for your infinite patience every time I interrupted your work with my endless and naive question-streams "*Uuuh, what's that you're doing? What's that for? Why? Yeah... but why?*". And **Andreia**, the best PhD companion I could have asked for. You were way more than a colleague and a friend, you were a psychologist, personal medical assistant, a true

confidant. To **Sara**, the mommy from NCN, thank you for your kindness and the coffee supply, it helped me so much. To **Eduardo**, thank you for all the engaging discussions, from neuroscience to aliens, from Romans to apple-making-machines. To **Rui**, thank you for the genuine discussions and encouraging thoughts about my future professional ventures. To **Ana Marote**, what a great new NCNer, thank you for your love for music and your amazing connections within the industry! And to our bright master students **Bia**, **Jalves**, **Sofia** and **Gatta**, if only I was this smart and mature at your age. In particular, thank you **Bia** for being so cute and for helping me find Helena Coelho, a night I will always remember (poor Zé though...). To **Ana Gerós**, who was and will always be an NCNer, thank you for all the good vibes and for being such outstanding "social glue". I'd also like to thank all the **i3S staff** for making this institute such a pleasant place to work. In particular, I'd like to thank the sweet canteen ladies for the discounts on the afternoon sandwiches – most of my thesis would not exist without that afternoon caloric boost.

To **Márcia**, my former MIB colleague and one of the best friends I have ever had, and to **João**, another great ex-NCNer. What a lovely crossing of universes. What a refreshing summer that was after the COVID lockdowns, thank you for that. And thank you for hosting me in Paris during my troublesome flight to Australia, made it worth it!

Now back to the other side of the ocean, the great land of Australia. My stay in Australia was the best 7-month block of my life. So stimulating and challenging, so fun and exciting. I have grown a lot during those 7 months, both academically and personally. And there are a lot of people to thank for that.

First, I felt very welcomed by everyone at the Department of Biomedical Engineering of the University of Melbourne. It was such a stimulating and vibrant community, with regular meetings, discussions and presentations on diverse and fascinating neuroengineering topics such as epilepsy research, vision prosthesis and brain-computer interfaces. I'd first like to thank **Martin** and **Hamish** for the work we developed and for always considering my ideas with a critical and rigorous lens. To my office colleagues, **Lilith** and **Zharo**, so kind, cheerful and generous; **Gianni**, always up for a good laugh and willing to help in everything I needed (those anti-theft car screws worked like a charm!); and, in particular, to **Sam**, such a good lad, so much fun outside work with bouldering, playing dodge-ball, running (you almost killed me though!).

But, if the Melbourne experience was magical, it was thanks to the amazing group we gathered. To **Dani** and **Isi** the German duo, thank you for the great energy, the road trips, and cooking for us! To **Lloyd**, thank you for registering all the special moments with your beautiful drawings. To **Eliza**, such a beautiful person, thank you for always taking care of us. And, of course, to **Maria**, *you crazy child...* How lucky was I that our lives crossed once again, straight from Café Concerto to Melbourne. You and your magnetic energy made every moment count.

Back to Portugal. So many people to be thankful for. To those helping me keep my sanity, namely all my Jiu-Jitsu team – **Bruno**, **Sandra**, **Hugo**, **Pedro** and all the others – and the surf crew, **Gou**, **Sandra**, **João**, and **Xavi** (let´s keep Luis for later...). You have no idea how important these refuges were at times.

To my "primaços"! **Antonio**, my dearest cousin, while I was "worried" about my PhD, you were bravely facing death in the eyes. Thank you for your example of courage and honour. You also followed the last days of my Master's Thesis somewhere in Asia...now your brother Miguel accompanies me with the Doctoral Thesis back in Australia. **Miguel**, by now we share so many crazy stories and memories, so much to tell our grandkids. Thank you for being so full of love, I respect

you and owe you a lot (and thank you for cooking for me during those inexplicable productivity peaks before lunch and dinner!).

To **Ana Martins**, for sharing the anxiety of the approaching submission deadline. An unexpected new friendship that I'm very grateful for. Thank you for all the extra work of submitting my thesis while I'm away – the next PhD thesis is on me!

To my life friends and former band buddies with whom I share some of my favourite memories: **Luís**, **Praça** and **Gonçalo**. Luis, you have always been a brother to me, and I feel blessed to have your unconditional friendship. Thank you for being the friend I can always count on. Praça, thank you for all the honest conversations we've shared and late-night jam sessions. Gonçalo, thank you for the contagious grind and dedication you put into everything you do, and for reminding me of the pleasure of making music.

And finally to my amazing family. To my three beautiful sisters, Margarida, Teresinha and Rosarinho. **Margarida**, for your inspiring selflessness, generosity, courage and strong spirit. I admire you so much. **Teresinha**, for your nurturing soul and for starting the family expansion, you are and will be an exceptional mother. **Rosarinho**, my sweet little sister, thank you for your kind and pure heart. It was really special travelling with you in Australia, can't wait for our next road trip! To our most recent family "add-ons", **Pi** and my niece **Teresinha**, thank you for bringing even more love into the family. To my dear grandma **Maria Carlos**, for being one of the strongest people I know, so resilient, practical and straightforward, with such a positive outlook on life.

To my loving parents, my mother **Teresa** and father **Domingos**. I couldn't have asked for more supporting parents and better role models. Your unconditional love for each other and your family is moving. I know it doesn't always seem like it, but I'm doing my best to follow your steps. Thank you for putting up with my crazy ideas, always with a comprehensive and caring heart. To you, I dedicate this thesis.

With love,

Domingos

# Contents

**Resumo** i

**Abstract** iii

**1 Introduction** 1
- 1.1 Motivation . . . 2
- 1.2 Objectives and Contributions . . . 4
- 1.3 Thesis Outline . . . 6

**2 Background** 7
- 2.1 Neuronal activity: A tale of multiple scales . . . 7
- 2.2 Therapeutic electrical neurostimulation . . . 13
  - 2.2.1 The modulating effect of electrical stimulation . . . 13
  - 2.2.2 Available neurostimulation therapies . . . 16
  - 2.2.3 Closing the loop on neurostimulation . . . 22
- 2.3 *In Vitro* electrophysiology with MEAs . . . 33
  - 2.3.1 Intracellular VS extracellular recordings . . . 33
  - 2.3.2 Reading, writing and responding to neurons with MEAs . . . 35
  - 2.3.3 *In vitro* preparations . . . 43

**3 Controlling unwanted neuronal oscillations with closed-loop stimulation** 49
- 3.1 Introduction . . . 50
- 3.2 Methods . . . 51
  - 3.2.1 Cell Cultures . . . 51
  - 3.2.2 Electrophysiological Recordings . . . 51
  - 3.2.3 Closed-loop Computer Application . . . 52
  - 3.2.4 Delayed Feedback Control Algorithm . . . 52
  - 3.2.5 Experimental Protocol . . . 53
  - 3.2.6 Selection criteria . . . 54
  - 3.2.7 Signal Processing . . . 54
  - 3.2.8 Statistical Analysis . . . 55
  - 3.2.9 *In silico* model . . . 56
- 3.3 Results . . . 56
  - 3.3.1 Adaptive component is required due to the dynamic nature of neuronal populations . . . 57
  - 3.3.2 Network controllability is conditioned by intrinsic population dynamics . . . 59

3.3.3 aDFC disrupts oscillations and decreases overall synchrony in controllable networks . . . 60
3.3.4 Controllability is affected by synaptic strength and balance between excitation and inhibition . . . 62
3.3.5 aDFC can promote transitions to stable asynchronous states . . . 64
3.4 Discussion . . . 65
3.A Supporting Information . . . 68
3.A.1 Acknowledgements: . . . 68

**4 Deep learning framework to optimise sensory neurostimulation 77**
4.1 Introduction . . . 79
4.2 Methods . . . 81
4.2.1 ANN-based stimulation protocol . . . 81
4.2.2 Measurement Predictor Network . . . 82
4.2.3 Stimulus Generator Network . . . 85
4.2.4 Model of Retina Stimulation and Recording . . . 85
4.3 Results . . . 90
4.3.1 The Measurement Predictor Network accurately estimates retinal responses to simulated multipolar stimulation . . . 90
4.3.2 The Stimulus Generator Network produces effective stimulation patterns 90
4.3.3 The ANN method matches the accuracy of the AMI approach for a LNL model retina . . . 92
4.4 Discussion . . . 95
4.A Supporting Information . . . 98

**5 Envisioning the next generation of neuroprosthesis with Neuromorphic Engineering 105**
5.1 Introduction . . . 107
5.2 Methods . . . 109
5.2.1 Memristive Devices . . . 109
5.2.2 Cell Culture . . . 109
5.2.3 Electrophysiological Recordings . . . 109
5.2.4 Real-Time Spike Detection . . . 111
5.2.5 Mediating Electrical Circuit . . . 111
5.2.6 Electrical Stimulation . . . 112
5.2.7 Experimental Protocol . . . 112
5.2.8 Signal Processing . . . 112
5.3 Results and Discussion . . . 113
5.3.1 Memristive Interface Detects and Selectively Responds to Network Bursting Activity in Real Time . . . 114
5.3.2 Memristive Interface Promotes Robust Coupling and Modulation of Neuronal Populations . . . 114
5.4 Conclusions . . . 115
5.A Supporting Information . . . 118

**6 Conclusions 121**
6.1 Bird's-eye view . . . 121
6.2 Contributions to the field of neuroengineering . . . 122
6.3 Future perspectives . . . 126

**References** **129**

**A PowerMapp toolbox** **161**

**B NCN MEA Control Suite** **167**

# List of Figures

2.1 **Electrophysiology at multiple scales.** Electrophysiological recordings can be performed at multiple spatiotemporal scales: micro-scale, at the neuronal level; mesoscale, at the circuits level; and macro-scale, at the brain regions level. The type of electrical signal collected at each scale varies accordingly: the micro-scale captures intracellular activity, allowing the recording of action potentials, and the small excitatory and inhibitory inputs the neuron receives; the mesoscale captures extracellular spikes (the extracellular signal generated by an AP) and Local Field Potentials (LFPs), depending on the recording method and neuronal model; at the macroscale, it is only possible to record slow brain waves. The three scales can be studied with different neuronal models: *in vivo*, *in vitro* and *in silico*. Each scale and neuronal model requires different recording techniques, ranging from single-cell patch clampling, to all brain EEG. This elements used in this figure were adapted from Servier Medical Art. Servier Medical Art by Servier is licensed under a Creative Commons Attribution 3.0 Unported License (https://creativecommons.org/licenses/by/3.0/) . . . . . 8

2.2 **Available neurostimulation devices.** Stimulation of the nervous system can be used as means of restoring healthy brain function (blue) of replacing lost sensory perception (orange). Regarding the first group, the most common therapies are: Deep Brain Stimulation (DBS), with applications for Parkinson's disease, dystonia, tremor, obsessive compulsive disorder, depression, epilepsy, among others; the Responsive Neurostimultor from NeuroPace for epilepsy control, the only responsive neurostimulator device in the market; and Vagus Nerve Stimulation (VNS), for epilepsy as well. As for replacing lost sensory function, we have: cochlear implants, which take the spotlight as the most successful implantable neurostimulator strategy; and retina prosthesis, which despite having a presence on the market, the available products are still far from satisfactory. The layout of this figure was inspired by the first figure from (Denison and Morrell, 2022) . . . 17

2.3 **Open-loop and closed-loop neural stimulation.** In open-loop stimulation (top) the stimulus parameters, such as pulse frequency, amplitude, width, etc., are fixed and the stimulation runs continuously. The brain response to stimulation is not observable since these devices have no recording capabilities. Thus, the stimulation parameters are tuned by the physician during clinical appointments based on the observed symptoms. In closed-loop stimulation (bottom), the stimulation parameters are adjusted in real-time, based on the recorded brain activity. The clinician is in charge of tuning the control parameters (such as control gains or thresholds), based on the clinical symptoms. . . . . 21

2.4 **Types of closed-loop electrical stimulation systems in Neuroengineering**. Closed-loop stimulation can be applied in two completely different tasks: either controlling neuronal activity (left) or using neural activity to control external bodies (right). The first type of closed-loop system is typically seen in cases where the brain activity is pathological and needs to be controlled. The locus of control is a closed-loop algorithm that receives the recorded neuronal activity as input and adjusts the electrical stimulation as output. We refer to this application as "Therapeutic Neurostimulation". The systems and the algorithms can be developed and tested with different types of neuronal models, *in vivo*, *in vitro* or *in silico*. Closed-loop stimulation of neuronal activity can also be used to control experimental conditions *in vitro* to better unravel the input-output function of the network. We refer to this type of application as "Controlling Network Variables". The other completely different application of closed-loop neural stimulation is for controlling external bodies. Here, the locus of control is the neuronal system, and not an external algorithm. The classical example is closed-loop Brain Machine Interfaces (BMI), where the brain activity of a paralysed patient is decoded to control a robotic arm or mouse cursor on a computer. The loop is closed by incorporating sensors into the system and encoding the sensor information into electrical neural stimulation. Analogous setups have been developed *in vitro* to study the capacity of neuronal networks to learn and adapt to external environments. This type of *in vitro* application is typically referred to as "Neuronal Cultures Embodiment". Our contributions with this thesis are focused on Therapeutic Neurostimulation, with studies performed both *in vitro* – Chatpers 3 and 5 – and *in silico* – Chapters 3 and 4 – as shown in blue. . . . . 24

2.5 **Closed-loop electrical neuromodulation algorithms.** The algorithms are organised according to the nature of the monitored signal and the nature of the actuation. The type of monitored signal is presented along the X-axis and divided into "Magnitude" and "Phase". The first refers to algorithms that monitor and control the magnitude of a given metric of neuronal activity, such as the power of a given frequency band. In "phase", we group the algorithms that track the phase of oscillatory phenomena or the timing of specific neuronal events. As for the nature of the actuation, presented along the Y-axis, we group the algorithms into "Event Triggered" and "Continuous Control", distinguishing those that only actuate when specific events are detected from those that are continuously running a control algorithm. The algorithms were placed on the map according to how they were implemented in the specific studies for neuronal control. We identify the neuronal models on which the methods were tested with the keys: **cl**, **vv**, **vt** and **sl**, which stand for clinical, *in vivo*, *in vitro* and *in silico* models, respectively. Some studies merge different control methods, which we represent with grey lines linking the different boxes. Our contributions are presented with white boxes, representing both the method name of the method (aDFC, Neuromorphic, Enc-Dec), the Chapters (II, III, IV), and the neuronal model used (*in vitro* and *in silico*) . . . . . 25

2.6 **Intracellular and extracellular neuronal recordings**. A) Spatial relationships between an extracellular electrode (yellow, with blue circuitry), intracellular needle (red), and the probed neuron (pink); and the associated equivalent electrical circuit (black) established by the electrode, neuron and culture medium (scheme inspired by Figure 1 from Spira and Hai (2013)). The electrode impedance is represented by the electrode resistance, $R_e$ and capacitance, $C_e$; the medium resistance at the cleft is represented by $R_{seal}$; the neuronal membrane is divided into two portions: the section that faces the electrode is referred to as junctional membrane (purple) as modelled by $R_j$ and $C_j$, and the section that faces the culture medium, referred to as non-junctional membrane (pink), and modelled by the $R_{nj}$ and $C_{nj}$. The scheme is not to scale. B) Action potential propagation across the neuron (pink), and the spatial spatial relations with the extracellular electrode (yellow) and intracellular needle (red). The membrane polarity of the section undergoing an action potential is inverted compared to the rest of the membrane, being relatively more positive intracellularly. C) Schematic representation of the same action potential recorded intracellularly and intracellularly. The intracellular signal shows sub-threshold potentials (gray arrows) such as excitatory and inhibitory signals received from presynaptic cells, referred to as excitatory presynaptic potentials (EPSPs) and inhibitory presynaptic potentials (IPSPs), respectively. The action potential has a positive peak since it captures the positive shift in polarity at the membrane. The amplitude of the recorded action potential corresponds to the actual magnitude of the physiological signal, with a rise of around 100mV. The extracellular signal, on the other hand, captures mostly the supra-threshold activity, i.e., the action potentials (being dependent on the electrodes input resistance and the dipole arrangement of the cells). Also, the distance between the electrode and the neuron leads to an attenuation of the spike's amplitude by a factor of about 1/1000. The recorded spike has a negative peak since it captures the transient potential drop of the extracellular medium . . . . . 34

2.7 **MEAs for *in vitro* electrophysiology.** A) There are two main classes of MEAs, passive MEAs (top) which typically have 60 to 256 electrodes, and CMOS HD MEAs (bottom), which can have from 4000 to more than 26000 electrodes. The passive MEA image corresponds to a 120 electrode MEA from MCS, taken from the MCS User Manual. The CMOS HD MEA is the MaxOne MEA from MaxWell Biosystems and the image was taken from their website. B) MEA systems enable "reading" and "writing" into neuronal populations. Passive MEAs have all signal conditioning circuitry located in an external headstage (hence the MEA itself is passive). CMOS HD MEAs, on the other hand, feature on-chip circuitry for signal conditioning and channel multiplexing, enabling high-density electrode arrays. Most MEA systems that enable stimulation use the same electrodes for recording and stimulating. To avoid strong stimulus artefacts, the reading channel (blue wiring) is opened while the stimulation is active (red wiring). Besides "reading" and "writing", MEA systems can also "respond" to neuronal activity. For that, it must integrate a processor that has access to the neuronal stream in real-time and the ability to control the stimulators. The processor can be the CPU of the computer used for data storage or, for experiments requiring low closed-loop latencies, a microprocessor such as a DSP or programmable logic device such as an FPGA, with direct access to the datastreams and stimulators. In this thesis, werespond to neuronal activity using the CPU in Chatper 3 and the DSP and a neuromorphic element in Chapter 5, as shown in the blue square . . . . . 37

2.8 The most common biological preparatoins are acute slices, organotypic cultures and dissociated neuronal cultures. These differ in many relevant aspects: acute slices are typically thicker than organotypic slices whereas dissociated neuronal cultures consist of a single layer of cells; the signals acptured by brain slices are LFPs due to the increased cell density, while dissociated cultures capture extracellular spikes of individual cells; slices preserve the original brain circuitry whereas dissociated cultures do not; acute brain slices can only be used for a couple of hours, whereas the neuronal cultures, be it organotypic or dissociated, can be maintained for several weeks for long-term experiments; finally, dissociated neuronal cultures are the easiest to prepare, while organotypic cultures require significant experimental expertise. The slices images were taken from the MCS manuals for the preparation of acute slices and organotypic cultures, the dissociated culture was taken from Dranias et al. (2013) . . . . . 44

3.1 **Controlling pathological oscillations with Adaptive Delayed Feedback Control (aDFC).** (A) In vivo concept of implantable closed-loop neurostimulation and the analogous in vitro setup developed. Here, we have dissociated hippocampal neurons as a brain, microelectrode arrays as implantable electrodes, emergent periodic network bursting as pathological synchronous oscillations, a workstation as an implantable chip and an external stimulator as an implantable neurostimulator. (B) Scheme of the Delayed Feedback Control algorithms. The instantaneous population firing rate is calculated and filtered by a damped oscillator. The actuation signal (stimulation frequency) is obtained by subtracting the ongoing oscillation to the past half-cycle delayed oscillation. The natural frequency of the oscillator, $\omega$, and the duration of the half-cycle delay are updated online (adaptive component in blue) by detecting the synchronous events and their current periodicity, T. (C) Controller computations under different firing regimes – oscillatory (left), asynchronous (right) and mixed, i.e., oscillatory with induced sparse activity (middle). Spike detection is performed for the monitoring electrodes (top) and the spike train is convolved online with a square window to compute the instantaneous population firing rate (middle). Synchronous events (black circles) are detected using a threshold crossing method and used to update the oscillation periodicity. The actuation signal (red) is built by subtracting the oscillator and delayed oscillator signals (blue and grey, respectively) and translated to a stimulation frequency signal. . . . . 58

3.2 **Representative example of the effect of DFC and aDFC stimulation protocols.**(A,B) Wavelet transform of the instantaneous firing rate of a network under DFC (A) and aDFC (B) protocols. (C,E) Example of the instantaneous firing rate of the network before stimulation (black) and during DFC (C, green) and aDFC (E, blue) stimulation. The signals shown in C and E correspond to the narrow time windows marked in A and B, respectively. (D,F) Power Spectrum Density of the instantaneous firing rate of the network before stimulation (black) and during DFC (D, green) and aDFC (F, blue) stimulation. . . . . 59

3.3 **Comparison of controllable and uncontrollable networks.** (A) Results obtained from multiple repetitions of the three stimulation protocols for a single network. We calculated the synchrony, firing rate, and oscillation intensity for the three segments – pre-stimulation (OFF), stimulation (ON), post-stimulation (OFF) – of each experiment and normalized each metric to the value of the pre-stimulation period. For the oscillation intensity, this normalization corresponds to the difference and not ratio, since the SNR is given in dB (logarithmic scale). Each triplet of points corresponds to one experiment performed with either aDFC (blue), DFC (green) or Poisson (red) stimulation. (B) Three-dimensional (3D) feature space – Firing Rate vs. Synchrony vs. Oscillation Intensity – for the network in A. Each line corresponds to the metrics evolution from the initial OFF to ON (from the first to the second point in A). This network is considered controllable since each protocol consistently drove the network to a distinct subspace, meaning that it had a unique and reproducible effect. (C) Representative example of the 3D feature space for an uncontrollable network where the effect of each stimulation protocol is indistinguishable from the others. (D) Networks were separated into controllable (blue) and uncontrollable (grey) by computing the MANOVA of the ON metrics for the different stimulation protocols. Networks with p-values lower than 0.05 had separable effects in the 3D feature space of the metrics and were thus considered controllable (8/13). [*Continues on the following page*] . . . . . 61

3.4 **Effect of aDFC, DFC and Poisson stimulation in controllable networks.** Fold change in oscillation intensity (A), synchrony (B) and firing rate (C) of the 5 min stimulation period compared to the 5 min pre-stimulation period with aDFC (blue), DFC (green) and Poisson (red) stimulation. The control group (grey) corresponds to baseline recordings where there is no stimulation and thus portrays the natural variability of each metric within two subsequent 5 min blocks. Each point represents the average effect of a stimulation protocol across all trials for a given network. We used the one-way ANOVA with repeated measures as statistical test (n = 8; * $p < 0.05$; ** $p < 0.01$; details of statistical tests in Table S1). . . . . 63

3.5 **In silico networks also present a controllable subspace for intermediate baseline levels of firing rate and synchrony.** (A) Level of synchrony (left) and firing rate (right) for the period of spontaneous activity. Each pixel represents the average synchrony and firing rate of five different simulations under the same conditions of excitatory/inhibitory balance and synaptic weight. (B) Average change in synchrony (left) and firing rate (right) of aDFC ON period compared to the OFF period. (C) Representative example of the simulations under three distinct network parameterizations (1, 2 and 3 in A and B). Each panel has the instant spike count at the top, the raster plot at the bottom and the network bursts (green) and triggered stimuli (red) in the middle. . . . . 63

3.6 **Multi-stable network with sporadic transitions to asynchronous state.** (A) Spontaneous synchrony and firing rate over time (middle). Representative raster plots of the synchronous (top) and asynchronous states (bottom). The coloured boxes represent the sliding window over which the firing rate and synchrony are calculated. (B) Synchrony vs. Firing Rate phase plot evidencing the cyclic behaviour of the network, with sporadic transitions to the asynchronous state (blue cluster). The clusters were identified using unsupervised Gaussian mixture models. (C) Changes in the frequency domain correlate with the automatically identified clusters (top bar). . . . . 64

3.7 **aDFC can promote transition to asynchronous state.** (A) States of spontaneous activity for a given multi-stable network. Blue cluster represents the sporadically reached asynchronous state. (B) Frequency domain of the recording considered in A. Blue rectangles represent the regions associated with the automatically identified asynchronous state. (C, E, G) Percentage of time in asynchronous state during the OFF-ON-OFF protocols for aDFC (C), DFC (E) and Poisson (G) stimulation. Each triplet of points represents an experiment performed with the corresponding stimulation protocol. We compared the time spent in AS during the ON and the second OFF periods with the initial OFF period using one-way ANOVA with repeated measures (**** $p < 0.0001$; details of statistical tests in Table S2). (D, F, H) Representative example of an experiment performed with aDFC (D), DFC (F) and Poisson (H) stimulation portrayed in the frequency domain. The stimulation period is represented with the associated coloured bar on the top. The identified asynchronous states are represented by the blue rectangles, as in B. . . . . 66

3.8 **Representative example of aDFC adaptability in a network with fast changing dynamics.** A) Instantaneous network firing rate during spontaneous activity. The black dots represent the real-time detection of network bursts performed by the adaptive controller. B) Interburst interval predicted by the aDFC controller and the associated real interbursts. Every time the controller detects a new network burst (black dots in A) it estimates when the next burst would spontaneously occur by calculating the median of the previous five interburst intervals. This prediction is used to properly tune the controller's half-cycle delay and the natural frequency of the oscillators used to create the actuation signal. In this example the stimulation is turned OFF, so the controller is simply tracking the neuronal activity without interfering with it. C) Histogram of the prediction error for the representative example shown in A and B. The largest errors result from the prediction lag associated with the fast transitions between bursting regimes. The fact that the prediction error is typically small, even under these fast-changing network dynamics, means that the half-cycle delay and the oscillator's natural frequency of aDFC are properly self-tuned in real-time. . . . . 69

3.9 **Controllable and uncontrollable networks across cultures, DIVs and animal preparations.** A) Dendrogram of the networks according to the controllability measure (MANOVA p-values of the modulation results obtained with the different stimulation protocols). This representation shows that the controllable networks (p-values < 0.05) are more similar to each other than to the uncontrollable ones. B) Spontaneous firing rate and synchrony for the all the trials of the different networks, separated into controllable and uncontrollable networks. C) Detail of B, evidencing how some cultures have very well confined spontaneous dynamics. D) Division of the networks according to the associated neuronal culture. When probed in different days, the same neuronal culture can be considered a different dynamical system because the spontaneous dynamics are clearly different. Therefore, we consider these to be different networks. For example, culture 1 (blue) and culture 2 (orange) led to two different networks each and, in the case of culture 2 one network is controllable whereas the other is not. Cultures 3 to 11 (grey) only have one network each and are represented with the same colour for clarity. E) Evolution of the networks with the DIV. This representation suggests that there is a trend towards higher levels of firing rate and synchrony with the DIVs. Also, the DIV by itself does not predict which networks are controllable and uncontrollable as these groups span multiple overlapping DIVs. F) The networks used came from a total of five different animal preparations. The distinction of controllable and uncontrollable networks cannot be attributed to any particular preparation, as the same preparation can lead to both types of networks. . . . . 71

3.10 **Representative example of the effect of Poisson stimulation in the frequency domain.** A) Wavelet transform of the instantaneous network firing rate during the OFF and ON periods. B) Example of the instantaneous firing rate of the network before stimulation (black) and during Poisson stimulation. The signals shown in B correspond to the narrow time windows marked in A. (C) Power Spectrum Density of the instantaneous firing rate of the network before stimulation (black) and during Poisson stimulation. Since the random Poisson stimulation is an open-loop control algorithm, i.e., it does not adapt to the ongoing network dynamics, the frequency of the Poisson process (which triggers the stimulation) had tuned beforehand. For that, we set the average stimulation frequency to match the average stimulation frequency used in the antecedent aDFC experiment. It is important to note that, even if properly tuned, open-loop Poisson stimulation will not adjust as the network approaches the desired state or stop stimulating once it is reached. This is relevant for cases like the one displayed in Figure 7, where aDFC and Poisson stimulation led to very different network responses. . . . . 72

3.11 **DFC leads to an increase in Oscillation Intensity and Frequency.** A) Change in Oscillation Frequency (fold change) as a function of the change in Oscillation Intensity (SNR difference) from the OFF to ON periods. In this representation, the left and right quadrants are respectively associated with a decrease and increase in Oscillation Intensity when stimulation is turned ON. The lower and upper quadrants are respectively associated with a reduction and increase in the Oscillation Frequency during stimulation. The Oscillation Frequency corresponds to the frequency at the peak of the PSD. To account for the fact that oscillations with very low SNR may not have a meaningfull value of oscillation frequency, the opacity of the datapoints is proportional to the value of SNR during the ON period. Thus, the more transparent points have less meaningful values of Oscillation Frequency. Each datapoint corresponds to an individual experiment performed with a given controllable network and sitmulation protocol (aDFC in blue, DFC in green, Poisson in red, and control in gray). The experiments performed with DFC are mostly located on the upper right quadrant, meaning that it consistently leads faster and more consistent oscillations. aDFC and random stimulation lead to a decrease in Oscillations Intensity, and the Oscillation Frequency is typically neglectable (datapoints are mostly transparent). As expected, the control group, where no stimulation is applied, is located at the center of the map, since the activity is similar during the ON and OFF periods. B) Representative examples of the Power Spectrum Density, from where the Oscillation Intensity and Frequency are calculated, during the OFF and ON periods for the three stimulation protocols and control, marked with arrows in panel A. . . . . 73

3.12 **aDFC allows the network to maintain some autonomous firing whereas DFC and Poisson tend to confine most network activity to stimulus responses.** A) Stimuli (red) and respective response windows of 400 ms (blue) for representative examples of aDFC (top), DFC (middle) and Poisson (bottom) applied to the same network. B) Detail of the spikes (black) evoked after stimulation (red) inside the response window (blue). C) Fraction of evoked spikes during stimulation for all the trials performed with the three protocols for each controllable network (one plot per network). When the fraction of evoked spikes is close to 1, the network has almost no spontaneous activity and is simply replicating the actuation signal, which is not ideal. [*Continues on the following page*] . . . . . 74

3.13 **The controllable region of in silico networks is also associated with intermediate levels of spontaneous firing rate and synchrony.** The values of synchrony and firing rate during the initial OFF period were plotted for all the simulations performed, covering all the network parametrizations. The datapoints are coloured according to the change in synchrony level from OFF to ON, as a proxy for controllability. The in silico simulations cover a much larger metric space than the in vitro results, which are confined to the regions between a spontaneous synchrony of 10-1 and 100 and spontaneous firing rate of 100 and 101 (the in silico simulations with very low spontaneous synchrony generated firing dynamics that would not be elected for in vitro experiments). . . . . 75

4.1 **Restoring vision with retina electrical stimulation.** (A) Schematic representation of the stimulation workflow: a camera captures the desired field of view of the implantee; the raw image is subsampled to create a target for the stimulation algorithm; the generated stimulus pattern activates the retina so that the elicited visual perception resembles the original image. (B) Sequential unipolar stimulation (either cathodic-first or anodic-first stimulus in all electrodes) leads to broad retinal activations (left) while simultaneous multipolar stimulation (simultaneous cathodic-first and anodic-first stimulus in different electrodes) can attenuate unwanted current spreads by generating sharper activations (right). The top images show simulated retina activations elicited by the 3-by-3 multi-electrode array while the bottom images represent schematically how the stimulation amplitudes (red and blue rectangles) shape the current spread (traced lines for the contributions of positive and negative stimulus and filled line for the resulting potential). The three electrodes (black rectangles) represented in the bottom images correspond to the three electrodes crossed by a diagonal white dotted arrow in the top images. . . . . 81

4.2 **Proposed deep learning-based stimulation protocol for visual prostheses .** The protocol has three stages: (1) training the measurement predictor network (MPN) based on the stored implantee's retinal response to random multi-polar stimuli; (2) training the stimulus generator network (SGN) using the trained MPN and a large dataset of natural images; (3) implementing the trained SGN in the implant for real-time generation of stimulus patterns based on the real-world images captured with a camera. . . . . 83

4.3 **Training of the Measurement Predictor Network (MPN) .** The MPN is a decoder that predicts the measurement of the retinal response to a given stimulus. It is meant to be a surrogate of the system Retina-MEA (blue box). The MPN is trained with a dataset composed of multiple random stimulus patterns (input data, represented by a coloured grid of electrodes) and the associated retinal response measured with the MEA (ground truth data). The error between the estimated retina measurement (top right) and the actual measurement (bottom right) is used to train the MPN. Here, both the stimulation and the recording are performed with the implanted MEA (simulated in this study) but could also be done with separate stimulating and recording units. . . . . 84

4.4 **Training and testing of the Stimulus Generator Network (SGN) .** The SGN receives a subsampled version of the raw captured image as input—the measurable target; this is the desired retinal response that we want to measure with the MEAs. The SGN outputs a stimulation pattern that should elicit the desired retinal response. To train the SGN, we use the implantee-specific predictor of retinal response, the measurement predictor network (MPN), and a large dataset of natural images (black branch). The error between the predicted measurement and the measurable target ($\varepsilon_{train}$) is used for training. Once the SGN is trained, it can be tested in the real retinal implant (red branch, here simulated). The efficacy of the SGN is accessed at the test stage by comparing the measured retinal activity with the measurable target, obtaining the measurable error ($\varepsilon_{test}$). . . . . 86

4.5 **Retina model and MEA recording.**A) Representative depiction of the ERFs of two different retinal cells (blue and red stars). B) Representative depiction of two different transfer functions ($k_1$=5, $z_1$=0.6; $k_2$=10, $z_2$=0.4). C) Retinal responses to a stimulus of amplitude 1 on the red electrode, for the different spread regimes considered in this study. D) Retinal cells (red) recorded by each electrode (blue star). E) Retina activation that results from a stimulus of amplitude 1 applied in the red electrode (left) and the resulting MEA read-out (right). . . . . 89

4.6 **The Measurement Predictor Network (MPN) predicts the stimulus response in diverse spread conditions.** (A) Measurement prediction error in the test set for MPNs trained under different spread regimes. For each regime, the experiments were repeated five times with different initializations of the MPN and retina model, hence there are five violin plots for each spread value. (B) Measurement prediction errors of the test set for a retina with spread of 1.5 (arrow in A). (C) Example of two test stimuli and the associated retina activations, measured activities and measurement predictions, extracted from the tails of the histogram in B, associated with low and high prediction error (blue and red dots in B, respectively). This shows that the MPN predictions are qualitatively very similar to the actual measurements. . . . . 91

4.7 **The stimulus generator network (SGN) encodes efficient multipolar stimulation patterns.** A) Test results for all the conditions simulated, where each dot represents the average error of 10,000 test stimuli for a given spread, regularization lambda, and random initialization. B) Training loss curves for the five SGNs instantiations at a spread of 1.5 and regularization lambda of $10^{-3}$ (black arrow in A). C) Histogram of the error for the 10,000 test images for the best SGN instantiation at a spread of 1.5 and lambda of $10^{-3}$ (whose training is marked with a black arrow in B). D) Qualitative assessment of the results across the spectrum of errors, showing the original raw image, the measurable target that the SGN aims to reproduce, the stimulus pattern produced by the SGN, the retina activation (unattainable in a real scenario), and the measured retinal activity. The measurable error represented in B corresponds to the error $\varepsilon$ between the measured retina activity and the measurable target. . . . . 93

4.8 **Comparison between the conventional stimulation method and neural activity shaping (NAS) performed with the artificial neural networks (ANN) protocol and analytic model inversion (AMI).** (A) Power consumption of the three stimulation methods for the entire test set. AMI was parameterized to have the same power consumption as the ANN method. The power of a given stimulus corresponds to the sum of the squared electrodes amplitude of the stimulation pattern. Here, the three methods were tested for a retina with spread of 1.5 (as in Figures 4.7 C and D). (B) Measurable error based on the MEA recordings at the test stage for the three stimulation methods. To account for experimental imprecisions that would occur in a clinical setting, we added noise to the parameters of the retina model in the test phase. The noise added to the parameters of each retinal cell was extracted from normal distributions centred in 0, with standard deviations of either 0 (no noise), 0.5 $\sigma$ or 1.0 $\sigma$, where $\sigma$ corresponds to the standard deviation of the values of the considered parameter across all the cells of the retina used to train the measurement predictor network. The coloured points depict the error of the test image shown in C. (C) Qualitative representation of the retina activation (top) and measured retinal activation through the MEA using the three stimulation methods (conventional in red, ANN in blue, AMI in green) for a given image (top, left) and the associated measurable target (bottom, left). Statistical tests for A and B are not shown because the large sample size ($n$ = 10 000) causes even negligible differences to appear significant. . . . . 94

4.9 **The ANN-based approach produces efficient stimulation patterns in retinas with different response topologies.** Qualitative results obtained in three different retinas: Retina 1 – regular topology, identical to those presented in the main paper; Retina 2 – elongated phosphene topology, obtained by changing the ERF parametrizations (y sigma larger than x sigma, ERF angles confined to a smaller range, set as a function of cell coordinates, retina transfer function parameters set as a function of cell coordinates); Retina 3 – elongated phosphene topology, as Retina 2, but with larger spread. The top row shows the pulse response s of each retina to a single pulse with amplitude 1. The remaining rows show the results obtained for different test images. Each retina was trained with a different MPN/SGN. The SGN learns stimulation patterns that adjust to specific retina responses. . . . . 99

4.10 **Comparison of conventional and ANN-based stimulation protocols for different spread regimes.** A) Measurable error based on the MEA recordings. For the conventional method we tested different scaling factors for each spread and chose the most efficient one (see Figure 4.11). The ANN-based approach achieved significantly better results, both quantitatively and qualitatively, for all spreads. B) Retinal activation error, i.e., the error between the actual retinal activity and the original raw image. Since the original image had a 32-by-32 resolution, we approximated the activity of the 10,000 retinal cells to a 32-by-32 image and computed the RMSE between that and the original image. The qualitative results of retina activation using ANNs show a clear improvement in terms of spatial resolution when compared with the conventional method. . . . . 100

4.11 **Scaling factor selected for different spread conditions in the conventional stimulation method.** To ensure that we compared the ANN approach with the best version of the conventional method, we tested different scaling factors and selected the most efficient one for each spread. We tested each scaling in the 10,000 test images and calculated the RMSE between the retina activation and the measurable target. For each spread, we chose the scaling with lower mean RMSE (black arrow). The results of the conventional method shown in Figure 8 were based on the parametrization for a spread of 1.5. . . . . 101

4.12 **The ANN method approximates the solutions provided by Analytical Model Inversion.** The stimulation patterns obtained with our ANN method are very similar to those obtained with AMI (left columns within the blue rectangle). This means that the presented ANN protocol can efficiently estimate the forward and inverse retinal model without imposing any assumptions about the retinal response (unlike the AMI which assumes a LNL model). We also show the solutions provided by the conventional method, which are just a direct translation of the target image into a stimulation pattern (middle column). . . . . 102

4.13 **The SGNs that become stuck in local minima during training induce bipolar generator potentials $\vec{G}$ that may lead to artifacts in the retina activation.** SGN 1 was the result of a successful training, thus generating unipolar $\vec{G}$ that lead to the expected retina activations. SGN 2 became stuck in a local minimum during training (worst performance in Figure 4.7.B). The higher training/validation errors obtained with this network are due to the artifact caused by the transition between regions of positive and negative $\vec{G}$. That is, in one region, the retinal cells are being activated due to an overall negative current, whereas in the other region they are being activated due an overall positive current (since the retinal transfer function is a double sigmoid with regards to $\vec{G}$ both positive and negative $\vec{G}$ values lead to retinal activation). In the transition between regions, the current is null, and thus there is no retina activation that leads to a dark artifact (example of the left). However, in some situations, this transition region may not affect the results or even be beneficial (example on the right). A network that learns how to effectively exploit these "artifacts" could be used to create sharp regions of no retinal activation, as proposed by Spencer et al. (2019) . . . . . 103

5.1 ***In vivo* concept and *in vitro* setup.** The neuromorphic technology presented in this study is envisioned for closed-loop loop devices with strict requirements in terms of size, power consumption and computing capabilities. The main function of the device is monitoring a given source population, processing the recorded activity in real-time, and actuating upon a target population when required (left). This was achieved in this study using an *in vitro* setup with hippocampal neuronal network cultures on MEAs , being monitored by a neuromorphic element – the memristor . . . . . 106

5.2 **Real-time monitoring and adaptive control of neuronal populations using memristors.** (A) Biological and memristive synapses share an analogous STP, with dynamics interchanging between low (OFF, left) and high (ON, right) conductance states as a function of the activity history. (B) Envisioned neuroprosthesis in vivo. The source and target neuronal population can be distinct (left) or the same (right), providing, either way, an effective monitor-compute-actuate paradigm through the memristor-based device. (C) In vitro implementation used in this work, with a hardware interface bridging neuronal and memristor signals. (D) Memristor hysteretic loop showing "set" under positive and "reset" under negative voltage [$V(t) = 0.5sin(100t)$]. The multiple cycles show the intrinsic variability of these types of devices, similar to the variability found in biological synapses. (E) STP dynamics with potentiation for repetitive positive pulsed (300 mV, 100µs) stimuli (top) and recovery under constant negative (-150 mV) stimulus (bottom). (F) Schematic representation of the in vitro setting, integrating six physically independent neuronal cultures in a 6-well MEA (left) and a memristor to monitor the source population in real-time and dynamically modulate the activity of five target populations (right). [*Continues on the following page*] . . . . . 110

5.3 **Real-time coupling and modulation of neuronal networks using a memristor.** (A) Representative source and target neuronal networks cultured in different wells (scale bar = 500 µm). (B) Raster plot (top) and average firing rate per electrode on the distinct wells (bottom); the left and right panels correspond, respectively, to the system without and with the memristor-mediated connection. High-frequency spikes in the source electrode (black dots, top right) correlate with NBs in the source well (well B, in this representative example). (C) Synchronisation of high-frequency activity between the source and targets is evidenced by the CC curves between the average firing rate per electrode on the wells (B, bottom). There is a CC curve for each source-target pair per trial for the three memristors tested (Memristor #1: 6 trials, 3.7 $\pm$ 1.7 min each; Memristor #2: 8 trials, 6.8 $\pm$ 4.2 min each; and Memristor #3: 6 trials, 11.9 $\pm$ 4.5 min each). The coloured lines represent the CC curves for the trials with the memristor, gray curves are for trials without the memristor. The thicker lines represent the average. (D) Diagram describing the three parameters used for the network activity fingerprints: NB duration, interval, and maximum firing rate. (E) The memristive interface modulates the dynamics of the network activity in the target wells. Each dot represents an NB in a given well. Before inserting the memristor (left), the NBs of each well had consistent separable features (duration, interval, and maximum firing rate). With the memristive interface, the activity fingerprints of the source (well B) and target wells are indistinguishable (right), showing that the network dynamics that govern the target populations were modulated by the source. . . . . 116

5.4 **Representative electrical I-V hysteretic behaviour of 8 discrete memristive devices in the same chip.** The devices show low operation power and the typical device-to-device variability. A larger hysteresis translates in a higher separation between the resistance states. . . . . 118

5.5 **Representative short-term plasticity (STP) behaviour.** A) Effect of the voltage amplitude (0.1-1.0 V) on the resistance change at fixed pulse number (20), duration (100 µs) and frequency. B) Cumulative memristor voltage increase (resistance decrease) after each 320 ms pulse at 2 Hz. C) Effect of the negative voltage amplitude (-0.25 and -0.7 V) on the resistance change time . . . . . 119

5.6 **Electrical circuit used to connect the memristor with the interface board of the MEA2100 System from MultiChannel Systems (bottom) and a schematic representation of the TTL signal across the circuit (top).** The first block is responsible for the transformation of TTL signal to the lower operation voltage range of the memristor. A voltage divider and a low-noise instrumentation amplifier shift the pulse to have both a positive (for “set”) and a negative (for “reset”) component. The last block amplifies the signal at the memristor output, for it to be read by the interface board as a logical "0" when the memristor is OFF (TTL < 2 V) and a logical "1" when the memristor is ON (TTL > 2 V). The diodes are only present as a precaution, for the protection of the data acquisition system. . . 120

A.1 **Loading data and setting analysis parameters.** A) PowerMapp opening window, where the user can open the file chooser dialogue (1). B) File finder platform (Schwarz, 2024). Using this platform, the user can search for files directories (2), select and add the desired files from within the selected directories (3), and check the files selected (4). C) The user can then see the selected files in the PowerMapp window (5), choose a downsampling factor (6), windowing parameters for the band-power (7), set the desired band limits (8), choose the reference type (9). The files are only uploaded after pressing the "Start Analysis" button (10). . . . . 162

A.2 **Analysis panel for a given sEEG recording**: A) Initial panel with a list of all the channels (1), the sEEG raw data (2), and the band-power for the frequency ranges selected (3). *Continues in the following page...* . . . . . 164

A.3 **PowerMapp can be used to find relevant biomarkers such as High-Frequency Oscillations (HFOs):** A) HFOs can be identified as isolated yellow spots in the high-frequency band-power map (1). B) A zoomed-in view of the sEEg signal at the identified spot shows the HFO (2), which would not be identifiable by looking at the original sEEG time-series. . . . . 165

B.1 **Layered architecture of the NCN MEA Control Suite.** The C# solution is organised in three separate blocks/projects: the Mea_Devices layer (blue) which interfaces with the hardware via the McsUsbNet library (gray); the GUI layer (red) which incorporates the User-Interfaces, and thus communicates with all the other layers; and the General_Logic layer (green) with all the classes for signal analysis 168

B.2 **Main GUI.** A) opening window with the acquisition hardware configurations of on the left panel and the buttons to open the available functionalities (the other GUIs) on the right. B) Dropdown box to select the acquisition device and USB port from all the devices connected to the computer. C) Dropdown box to select the MEA layout associated with the select device. D) Dropdown box to set the acquisition sampling rate . . . . . 169

B.3 **ElecDataForm GUI for electrode data visualization**. A) Example of the layout for single well 60-electrode MEA device. B) Example of the layout for the 6-well 60-electrode device. The layout of the electrode grid is set automatically according to the device selected in the Main GUI. . . . . 170

B.4 **SpikesForm GUI for performing online spike detection**. A) Tab to define spike thresholds for each electrode. The threshold is set to N sigmas of background noise for each electrode after filtering the data with a high-pass filter. B) Tab to define the monitoring electrodes. Selected electrode are represented in green. C) Online spike detection tab. Each spike is represented by a blue dot associated with the corresponding electrode along the Y-axis. . . . . 172

B.5 **StimulatorForm GUI for manual control of the stimulators** A) Setting the stimulation electrodes for stimulator 1 (STG1 in green). B) Setting the stimulation electrodes for stimulator 2 (STG2 in red). C) Example of the stimulation electrodes layout for a MEA device with 256 electrodes and 9 wells. . . . . 173

# List of Tables

3.1 Details of the statistical tests used in Figure 3.4 to compare the modulation results of the different stimulation protocols in controllable networks. We used repeated measures one-way ANOVA with multiple comparisons . . . . . 70
3.2 Details of the statistical tests used in Figure 3.7 to compare percentage of time spent in asynchronous state for each stimulation protocol in a multistable network. We used repeated measures one-way ANOVA with multiple comparisons. . . . . 70
3.3 Details of the statistical tests used in Figure 3.10 to compare fraction of spontaneous spikes during stimulation for the three stimulation protocols. We used repeated measures one-way ANOVA with multiple comparisons. . . . . 70

# Abbreviations

| | |
|---|---|
| AI | Artificial Intelligence |
| ANN | Artificial Neuronal Networks |
| API | Application Programming Interface |
| DBS | Direct Brain Stimulation |
| DSP | Digital Signal Processor |
| DLL | Dynamics Link Library |
| ECoG | Electrocorticography |
| EEG | Electroencephalography |
| EPSP | Excitatory Postsynaptic Potential |
| FPGA | Field-Programmable Gate Array |
| GUI | Graphical User Interface |
| IPSP | Inhibitory Postsynaptic Potential |
| LFP | Local Field Potential |
| MEA | Micro-Electrode Array |
| NB | Network Burst |
| P | Proportional |
| PD | Parkinson's Disease |
| PI | Proportional-Integral |
| PID | Proportional-Integral-Derivative |
| PRC | Phase Response Curve |
| sEEG | Stereo-electroencephalography |
| SCS | Spinal Cord Stimulation |
| UPO | Unstable Periodic Orbit |
| VNS | Vagus Nerve Stimulation |

# Chapter 1

# Introduction

In the last three decades, the number of deaths associated with neurological disorders increased by 39% and disability-adjusted life-years (DALYs, which counts the years of healthy life lost) increased by 15%. This trend is expected to persist due to the global population growth and ageing (Feigin et al., 2020).

Despite the significant advances in pharmacological therapies for diseases such as epilepsy (Sankaraneni and Lachhwani, 2015; Kanner and Bicchi, 2022), Parkinson's Disease (PD) (Connolly and Lang, 2014; Cacabelos, 2017), or essential tremor (Zesiewicz et al., 2005), many patients are still drug-resistant or suffer the burden of severe side effects (Rush et al., 2006; Louis et al., 2010; Sankar and Lozano, 2011; Olanow et al., 2013). Since the discovery of new drugs for neurological diseases has been slow and with modest impact regarding drug resistance (Perucca et al., 2007; Perucca, 2019; Vijiaratnam et al., 2021), there is a strong motivation to find alternative therapies that may increase the quality of lifeand reduce the DALYs for these patients.

Electrical neurostimulation is a viable option for many conditions. The implantation of electrodes in direct contact with neuronal tissue enables the electrical modulation of neuronal activity. Some examples include the well-established Spinal Cord Stimulation (SCS) for chronic pain (Shealy et al., 1967; Knotkova et al., 2021) and cochlear implants to restore hearing (Simmons et al., 1965; Naples and Ruckenstein, 2020); the successful application of Deep Brain Stimulation (DBS) to treat drug-resistant neurological disorders such as PD (Burchiel et al., 1999; Schuurman et al., 2000; Bronstein et al., 2011), essential tremor (Hubble et al., 1996; Schuurman et al., 2000; Rehncrona et al., 2002), epilepsy (Kerrigan et al., 2004; Andrade et al., 2006; Vonck and Boon, 2012; Benedetti-Isaac et al., 2015), multiple sclerosis tremor (Schuurman et al., 2000), depression (Figee et al., 2022; Sheth et al., 2022), chronic pain Rasche et al. (2006); Boccard et al. (2013, 2015) and severe obsessive-compulsive disorder (Nuttin et al., 1999; Abelson et al., 2005; Wu et al., 2020); and the promising use of retinal or cortex stimulation to restore visual perception (Dobelle and Mladejovsky, 1974; Yue et al., 2016).

However, despite its remarkable effectiveness for some applications, such as cochlear implants, other use-cases are still unable to exploit the full potential of this neuromodulation strategy. Stimulation outcomes typically include poor efficacy, excessive stimulation, and the associated undesirable side effects (Schermer, 2011; Wojtecki et al., 2015; Zarzycki and Domitrz, 2019). Most of these limitations can be traced back to the fact that, currently, stimulation does not act in accordance with the associated neuronal response or the ongoing brain activity. Stimulation runs continuously based on parameters chosen empirically by the clinician. With this open-loop setup, the actuation is blind to the state of the system it tries to control. In theory, closed-loop control should be able to achieve a more efficacious and stable modulation of faulty neuronal activity. Also, moving from a continuous stimulation paradigm to an adaptive one should result in power savings, reducing the surgical interventions required to replace the battery, which is often associated with severe complications (Hariz, 2002).

Despite its potential, closed-loop neural stimulation is still in its infancy. Much is still unknown about how to efficiently implement it and the long-term efficacy and adverse effects of chronic implants remain to be determined (Krauss et al., 2020). So far, the clinical results of the few closed-loop stimulators available in the market show that the effectiveness is similar to that of simpler open-loop devices (Vassileva et al., 2018; Alcala-Zermeno et al., 2022). The algorithms implemented rely on very simple rules, such as an ON-OFF switch that triggers a predefined stimulus pattern when a given electrophysiological biomarker of neuronal activity is detected. Such control techniques do not leverage the extensive knowledge and principles from relevant engineering fields such as Control Theory, Data Science or even Computational Neuroscience. Artificial Intelligence (AI) techniques like Deep Learning have shown remarkable results for the analysis of neuronal data and now play a significant role in Neuroscience (Schirrmeister et al., 2017; Kiral-Kornek et al., 2018; Daoud and Bayoumi, 2019; Craik et al., 2019; Aggarwal and Chugh, 2022). Nonetheless, all this knowledge is not yet being effectively translated into implantable devices and there is still a significant gap between the technological tools available, from an engineering standpoint, and those actually implemented for neural stimulation.

## 1.1 Motivation

Many challenges hinder the effective implementation of advanced neuromodulation strategies. We know that brain activity correlates with mental states and behaviour, both in healthy and pathological conditions. But, being arguably the most complex system ever studied, it is not clear how different patterns of neuronal activity give rise to brain function (Doty, 1998; Hameroff, 2006; Koch et al., 2016; Zhang et al., 2011) and, more importantly, how we can optimally interfere with such patterns for therapeutic applications (Kuncel and Grill, 2004; Suppa et al., 2016).

This makes it incredibly challenging to conceive a robust, efficient, yet minimally interfering, control algorithm, especially given the inherent variability between individual brains. Moreover, even if the rules governing the relation between brain state and electrical stimulation were known, the

complex neuronal dynamics can never be fully observed by the controller, which only has access to a spatially restricted sampling of the brain through recording electrodes. Also, the actuation points are limited to the stimulation electrodes which only directly affect the surrounding neuronal tissue. On top of all these technical challenges, there is yet another significant constraint that does not apply to other systems in typical engineering problems: each brain is the core of an individual sentient organism. That raises severe ethical concerns regarding the *in vivo* exploration of innovative closed-loop brain stimulation algorithms.

We argue that the investigation of advanced control algorithms for neurostimulation should start with simple neuronal systems, namely *in silico* and *in vitro* models. The methods that are efficient in controlling the simpler systems are good candidates for *in vivo* testing. Also, by starting from first principles, with systems that are readily accessible and where the experimental conditions are easier to control, the underlying mechanisms through which electrical stimulation influences complex network dynamics may become clearer.

Computer simulations are typically the initial approach when testing control methods in engineering. That entails simulating the environment, referred to as *plant*, and exploring the algorithms and parametrizations required to control it. In the context of neuromodulation, the *plant* would be a computational model of the specific neuronal system under study, tuned to describe the observed pathological dynamics. The field of Computational Neuroscience provides us with a vast repertoire of computer models that can be used to simulate neural systems (Sejnowski et al., 1988; Migliore et al., 2003). The models can differ in complexity and spatiotemporal resolution, ranging from a single neuron to whole brain models. With this framework, the researcher is free to explore multiple approaches and compare them rigorously under highly controlled conditions. The downside is that computer simulations, despite being useful, will always be a simplistic representation of the actual neuronal system.

A step-up from *in silico* models is using actual biological neuronal populations. *In vitro* setups are popular in the various fields of Neuroscience research because they enable highly controllable experimental conditions with high throughput (Malik et al., 2014; Little et al., 2019). Neuronal populations cultured *in vitro* retain many neurophysiological properties of *in vivo* models: non-stationary spontaneous activity (Opitz et al., 2002; Wagenaar et al., 2006; Ham et al., 2007; Cohen et al., 2008; Orlandi et al., 2013), intrinsic excitability (Bertucci et al., 2019), inducible plasticity (Chiappalone et al., 2008; le Feber et al., 2010), and even fundamental computational properties (Feinerman et al., 2008). Most of these studies rely on microelectrode array (MEA) technology, which enables the concomitant recording and stimulation of neuronal populations with high spatiotemporal resolution. When cultured on MEAs, the networks can be maintained for several weeks or even months, enabling long-term studies. Thus, they provide a powerful test-bed to explore closed-loop stimulation protocols designed control to the dynamics of neuronal populations.

## 1.2 Objectives and Contributions

The main objective of this thesis was to develop intelligent electrical stimulation protocols specifically designed for therapeutic applications. My work was performed *in vitro* and *in silico*, leveraging on the highly controllable experimental conditions that these environments provide to explore uncovered ground in terms of stimulation strategies. The overarching goal was to harness the available tools from various engineering sub-fields and put them to the test in a neuromodulation context, with a vision toward their potential implementation in future neuroprosthesis. By working form the "bottom-up", using simple yet realistic neuronal models, we can precisely test, evaluate and compare more complex control strategies with less confounding variables.

Throughout this thesis I focused on two main aspects: the development and validation of the stimulation algorithms themselves and the exploration of efficient computational paradigms to be employed in future implantable devices. The key research questions can be laid out as follows:

1. **What type of control algorithm can be used to disrupt pathological oscillations?**

   There is a vast myriad of techniques that can be employed to control the behaviour of a given system. The adequate choice of controller plays a crucial role in the overall success. This choice should be guided, first of all, by the nature of the dynamics exhibited by the system. The dynamics of neuronal firing and their interactions with electrical stimulation are highly nonlinear, arguably chaotic. Techniques used for Chaos Control have been hypothesised as a way to disrupt pathological neuronal activity in computational studies (Pyragas, 1992; Popovych et al., 2005, 2006; Tukhlina et al., 2007; Batista et al., 2010; Popovych et al., 2017a; Popovych and Tass, 2019). However, the proposed methods were never tested with biological neuronal networks. So, I address this first research question by designing an *in vitro* framework for closed-loop control of neuronal populations and developing Chaos Control methods to disrupt unwanted neuronal dynamics. This work was published in the journal eLife (de Castro et al., 2024):

   **Publication I:** D. Castro et al., "Disrupting abnormal neuronal oscillations with adaptive delayed feedback control",*eLife* (2024), https://doi.org/10.7554/eLife.89151

2. **How can machine learning techniques guide electrical stimulation for optimal neuromodulation?** The complex rules governing neural response to stimulation may undermine the effective use of conventional control techniques. The rapid advancement of AI has opened up new possibilities for neuromodulation algorithms, enabling data-driven patient-specific stimulation protocols. But how should such a protocol be designed? What sort of data-driven methods would effectively find optimal electrical stimuli that elicit the desired neuronal activation? We answered this question in the context of retinal prostheses, where the goal is to apply stimulation patterns to the MEA implanted in the retina in order to elicit

the desired visual perception. Our AI-based approach showed promising results in a realistic computational model of retinal response to stimulation. The results were published in the Journal of Neural Engineering (Leite de Castro et al., 2024):

**Publication II:** D. Castro et al., "Neural activity shaping in visual prostheses with deep learning ", *Journal of Neural Engineering* (2024), https://doi.org/10.1088/1741-2552/ad6186

3. **How can intelligent stimulation protocols run efficiently in future neuroprosthesis?**

   As implantable devices, neurostimulators are inherently limited in size, power efficiency and, consequently, computational capabilities. Even though the latter is not a concern for devices that provide continuous stimulation, future implantable devices will require strong computational capabilities to process large quantities of neuronal data in real-time and make decisions based on that. So, my final research question focused on exploring possible solutions for a nano-chip that could meet the computational requirements of an intelligent neurostimulator. The solution we propose here is based on a Neuromorphic Engineering framework, a computational paradigm where event-driven processing and memory are co-localised within nano-scale neuron-like hardware elements. Neuromorphic computing is inherently parallelizable, clockless, low-power and for ideal data-driven solutions such as Deep Learning. We showed, for the first time, that a nano-scale neuromorphic element – the *memristor* – can selectively modulate the activity of neuronal cultures with electrical stimulation in real-time. Our results were published in the journal ACS Applied Electronic Materials (Dias et al., 2022):

   **Publication III:** C. Dias and D. Castro et al., "Memristor-based neuromodulation device for real-time monitoring and adaptive control of neuronal populations", *ACS Applied Electronic Materials* (2022), https://doi.org/10.1021/acsaelm.2c00198

My three first-author publications counted with the support and collaboration of multiple teams and organisations, namely my team Neuroengineering and Computational Neuroscience (NCN) at i3S, the Institute of Physics for Advanced Materials, Nanotechnology and Photonics (IFIMUP) from the Faculty of Sciences of the University of Porto, and Department of Biomedical Engineering of the University of Melbourne.

Additionally to the work published in peer-reviewed journals, we also present two software packages that are provided to the scientific community as open-source applications. The first software stems from a collaboration with the neurosurgery team from Centro Hospitalar do Porto (CHP). I made a MATLAB application with an intuitive Graphical User Inferface (GUI), the *PowerMapp* (see Apendix A), for clinicians to analyse stereoelectroencephalography (SEEG) data from epileptic patients recorded prior to resective surgery. PowerMapp is already being used by clinicians to support clinical decisions.

The second open-source software presented here is an Application Programming interface (API) and a set of GUI templates, the *NCN MEA Control Suite* (see Apendix B), for controlling the MEA's acquisition and stimulation hardware used in the Publication I. This C# API was designed to facilitate the development of future closed-loop applications for *in vitro* experimentation.

## 1.3 Thesis Outline

In Chapter 2, I will provide a brief introduction of fundamental concepts in Neuroscience and Neuroengineering. Section 2.1 opens with an overview of neuronal activity and how it is differntly manifested at multiple spatiotemporal scales. In section 2.2, I will focus on how electrical stimulation has been used as a therapeutic option and then detail the state of the art regarding closed-loop control of neuronal activity. Finally, in the last section of the introduction, section 2.3, I will narrow down into the specificities of *in vitro* electrophysiology using micro-electrode array technology. I intentionally wrote this first introductory chapter to be concise while broad in its content so that it encompasses the three following chapters. Each of the remaining chapters will then have its own specific introduction, where all the required contextualisation is provided.

Chapter 3 will be devoted to the Research Question 1. I will start by introducing the rationale behind using Chaos Control to modulate neuronal activity and then present my **Publication I**, where we implemented, for the first time with *in vitro* neuronal cultures, a popular Chaos Control method meant to disrupt excessive neuronal synchronization called Delayed Feedback Control (DFC).

In Chapter 4, I will focus on the Research Question 2. After a brief introduction motivating the need for AI-based algorithms in future neurostimulators, I will present my **Publication II**, where we used an encoder-decoder Deep Learning architecture to effectively modulate retinal activity in visual prosthesis *in silico*.

Chapter 5 will be dedicated to the Research Question 3. Here, I will start by explaining what Neuromorphic Engineering is, and what motivates its use for future intelligent neurostimulators. I will then present my **Publication III**, a proof-of-concept study where we used a neuromorphic element, the memristor, to modulate a target neuronal culture based on real-time feature detection of another neuronal culture *in vitro*.

I will conclude this dissertation in Chapter 6 with an overall discussion of the work developed and the future directions that can be taken from here onward. The two open-source software packages – PowerMapp and NCN MEA Control Suite – are presented in the Appendix.

# Chapter 2

# Background

## 2.1 Neuronal activity: A tale of multiple scales

Neurons are the fundamental functional units of the brain. Despite being outnumbered by non-neuronal glial cells, neurons are in charge of crucial brain functions such as sensory processing, motor control and cognition. But how do they do that? What is going on when things start to fail? And how can we engineer ways to compensate for such malfunctions? There are no trivial or defined answers to any of these questions, and any attempt to do so will most likely require us to understand brain function at multiple spatiotemporal scales.

#### Microscale: single neuron

At the microscale, we have all the phenomena that occur at the neuronal level (and below), from which the most relevant is the capacity neurons have to generate electrical signals. They do that by integrating the electrical inputs received from neighbouring neurons through their synapses, at the soma. If the received signals lead to a decrease in intracellular potential, resting at -70 mV, they are referred to as inhibitory postsynaptic potentials (IPSP). On the other hand, if the synaptic inputs make the intracellular potential less negative, they are called excitatory postsynaptic potentials (EPSP). If this excitatory input is strong enough and the intracellular potential rises above -55 mV, something drastic happens: a cascade of ionic exchanges is suddenly triggered at the cellular membrane, which locally shifts its polarity generating an electrical current, referred to as *action potential* or *spike*. The spike is propagated through the axon via ionic channels in a wave of membrane depolarisation until it reaches its output terminals (and is also backpropagated to the soma and dendrites). It is then sent chemically to the neurons it branches to, either exciting or inhibiting them depending on the type of neurotransmitter released. The entire process occurs in milliseconds, involving changes in the transmembrane potential of tens of millivolts.

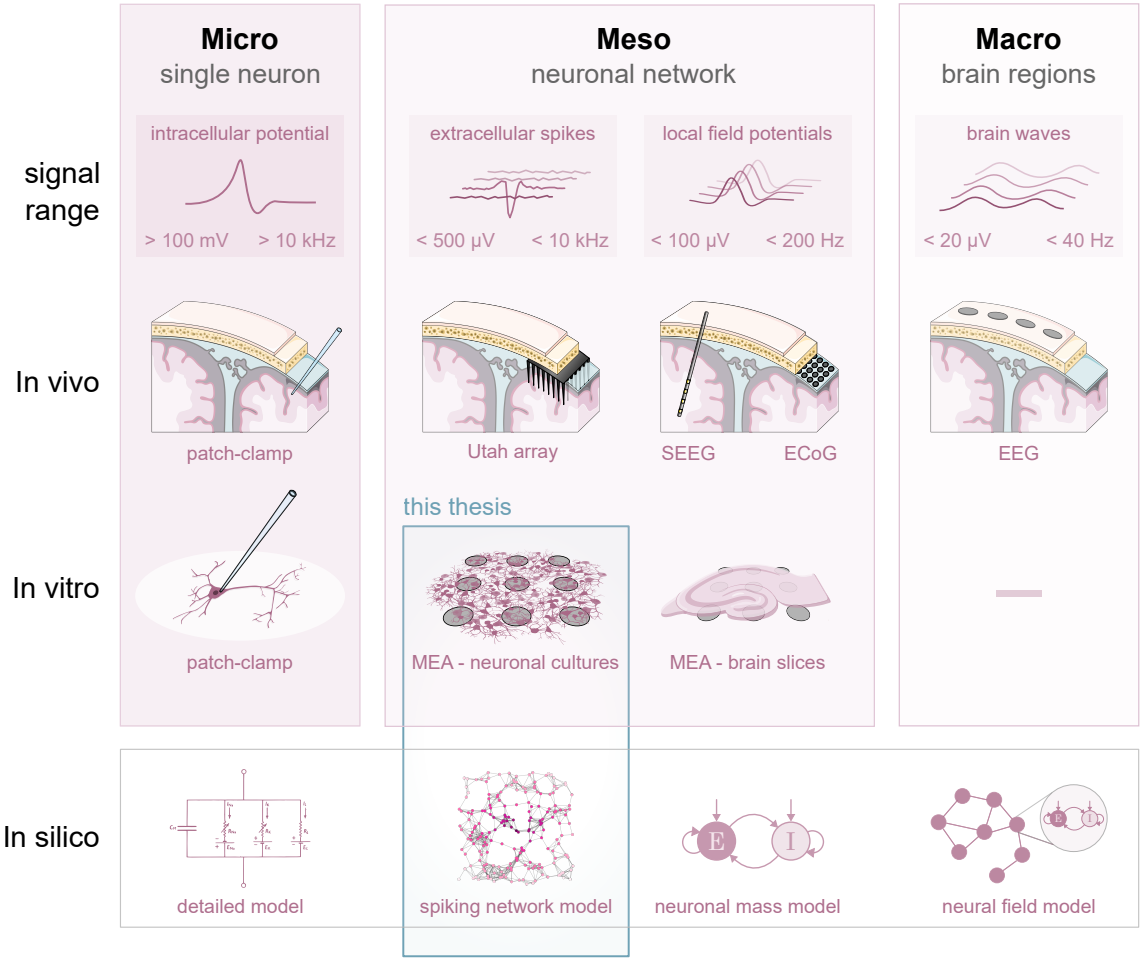



Figure 2.1: **Electrophysiology at multiple scales.** Electrophysiological recordings can be performed at multiple spatiotemporal scales: micro-scale, at the neuronal level; mesoscale, at the circuits level; and macro-scale, at the brain regions level. The type of electrical signal collected at each scale varies accordingly: the micro-scale captures intracellular activity, allowing the recording of action potentials, and the small excitatory and inhibitory inputs the neuron receives; the mesoscale captures extracellular spikes (the extracellular signal generated by an AP) and Local Field Potentials (LFPs), depending on the recording method and neuronal model; at the macroscale, it is only possible to record slow brain waves. The three scales can be studied with different neuronal models: *in vivo*, *in vitro* and *in silico*. Each scale and neuronal model requires different recording techniques, ranging from single-cell patch clampling, to all brain EEG. This elements used in this figure were adapted from Servier Medical Art. Servier Medical Art by Servier is licensed under a Creative Commons Attribution 3.0 Unported License (https://creativecommons.org/licenses/by/3.0/)

This seemingly simple all-or-none, spike or no-spike behaviour is enough to lay the foundations for a general computing machine through the assembly of small circuits of excitatory and inhibitory neurons (McCulloch and Pitts, 1943). Furthermore, neurons are not stereotypical structures. Each neuron has unique morphology, ion channel expression, neurotransmitter release and reception capabilities, connectivity with neighbouring neurons, etc., which promotes a rich repertoire of firing capabilities and expands the computing capabilities of neuronal networks.

Currently, there is a vast understanding of the microscopic physiology and morphology of neurons, with contributions spanning multiple scientific fields. Electrophysiology played a major role here, providing techniques that enable the recording of intracellular currents of individual neurons or even a single ion channel, both *in vitro* and *in vivo* (Figure 2.1, left). Patch clamping was at the forefront of electrophysiology and is still the go-to technique for those interested in studying the microscale activity of neurons. It was developed in the 70's by E. Neher and B. Sakmann who were later awarded the Nobel Prize for Physiology or Medicine, in 1991. Patch clamping techniques measure the current or voltage drop generated by the flow of ions crossing the cellular membrane. The measurements are performed with a glass needle whose small tip (diameter around 2-3 $\mu$m) seals a small area of the cellular membrane.

Even before the invention of patch clamping, Hodgkin and Huxley were already able to record the intracellular neuronal activity using glass electrodes that fit inside the squid's giant axon. This precursor of patch clamp allowed them to study the ionic mechanisms behind the generation of action potentials enabling them to develop their Nobel Prize-worthy neuronal model (Hodgkin and Huxley, 1952), the most famous *in silico* model of microscale neuronal activity (Figure 2.1 bottom left). The Hodgkin-Huxley (HH) model describes the dynamics of different ion channels through an equivalent electrical circuit which captures the generation and propagation of action potentials. The parameters of the model can be fit to experimental data in order to build different types of neurons (Pospischil et al., 2008). The microscopic analysis can dive much deeper by looking at the molecular and genetic factors that determine specific neuronal properties, but that is out of the scope of this thesis.

**Macroscale: brain regions**

Despite the undeniable importance of understanding such microscale details of neurophysiology, brain disorders eventually manifest at the macroscopic level. At this end of the spectrum, the focus typically lies in associating the observed symptoms with the underlying neuronal activity recorded from whole brain regions. For example, an epileptic seizure can be described as "a transient occurrence of signs and/or symptoms due to abnormal excessive or synchronous neuronal activity in the brain" (Fisher et al., 2005).

The macroscopic neuronal activity can be recorded with electroencephalography (EEG), a popular and non-invasive technique invented almost 100 years ago (Berger, 1929). The electrodes are

placed on the scalp and capture the collective electrical activation of entire cortical regions (centimetre range). The frequency bandwidth covered by EEG ranges from 0.5 to 40 Hz, with signal amplitudes between 10 – 20 $\mu$V (Schalk and Leuthardt, 2011) (Figure 2.1, right).

The signals collected by EEG appear as slow oscillating waves with characteristic frequency components that correlate with different brain states, in both healthy and pathological conditions. Brain waves are typically divided into different frequency classes. The waves captured within the bandwidth of EEG are delta (below 4 Hz), theta (4–8 Hz), alpha (8–12 Hz) and beta waves (12–30 Hz). Higher frequency waves such gamma (above 30 Hz) are hardly measured with EEG.

Some brain disorders affect the way neuronal circuits communicate leading to measurable phenotypes in the recorded brain activity, commonly named *biomarkers*. Examples of biomarkers include the increase in beta band observed in PD patients or the transient high amplitude and high-frequency waves that dominate the EEG during epileptic seizures. The macroscale lens gives us a broad understanding of brain dynamics and states (Stam, 2005), brain activity decoding (Schirrmeister et al., 2017), communication pathways, organisational and functional aspects of different brain regions and how they may correlate with function (Sporns, 2011; van den Heuvel and Yeo, 2017). However, there is still an explanatory gap between the observed microscopic cellular activity and the observed macroscopic dynamics of brain regions (Devor et al., 2013; Zhang et al., 2011). This explanatory gap is due to the mismatch of orders of magnitude, both in time and space, between the micro and macro scales. For example, epilepsy research has a strong focus on mapping brain-wide dynamics but that is typically done at a scale too far removed from the circuit-level dynamics that generate and sustain seizures (Farrell et al., 2019).

Our current computational models are also circumscribed to a tight spatiotemporal scale (Figure 2.1 bottom) so they cannot explain how changes at the micro level may lead to their macro manifestation. For example, a detailed model describing the intracellular dynamics of a pyramidal neuron will not be able to explain the emergence and propagation of spontaneous hippocampal seizures. For that, you may need a computational model that operates at the macroscale, and these can take several forms. One example is Neural Field Models (Figure 2.1 bottom), which try to model the spatial and temporal neuronal dynamics within a brain region. These are typically built by coupling multiple excitatory and inhibitory populations, each modelled by its average firing rate, a type of model known as Neural Mass Model (NMM), introduced by Wilson and Cowan in their seminal publication (Wilson and Cowan, 1972). Neural field models have been used to provide a high-level mechanistic explanation for the EEG and MEG (magnetoencephalography) signals (Jirsa et al., 2002; Jirsa, 2009; Bojak et al., 2010) and study the effects of brain stimulation (Giannakakis et al., 2020). NMMs, which are concerned only with temporal dynamics, have already been used to model the different types of brain waves by tuning the kinetics of different excitatory and inhibitory populations, their connectivity and connection delay (David and Friston, 2003).

Another popular approach is simulating the observed dynamics directly using abstract mathematical models that capture the temporal transitions and bifurcations, with no concerns regarding the underlying physiology. The most famous example is the Epileptor (Jirsa et al., 2014), a model that can simulate the transitions to seizure using only five state variables linked with integral-differential equations. But again, these computational models may help us understand the macroscale electrophysiology data but do not allow us to establish any causal link to its microscale underpinnings.

**Mesoscale: neuronal networks**

Looking at an intermediate level – the mesoscale – may fill some of these gaps (Figure 2.1, middle). The mesoscale can be accessed using multiple electrodes in direct contact with the neuronal tissue. This provides a high-resolution readout of neuronal activity, both in time and space. Unlike the microscopic techniques, at the mesoscale, the recordings are performed extracellularly and cover multiple points of the neuronal tissue instead of a single neuron; unlike the macroscopic techniques, mesoscopic recordings capture the local interactions within a specific neuronal network. The recorded activity is constricted to the surroundings of the electrodes, whose reach depends on the specific recording modality.

In clinical and *in vivo* settings, electrodes are typically implanted either in the neocortex (electrocorticography, ECoG, also referred to as intracranial EEG) or in deeper brain structures (e.g. SEEG) (Figure 2.1, *in vivo* mesoscale on the right). The signals collected are known as Local Field Potentials (LFP) and are thought to correspond to the summation of synchronized synaptic currents of neurons in the vicinity of the electrodes (millimetre range) (Obien et al., 2015). Besides the higher spatial resolution, ECoG has a much broader frequency bandwidth than EEG, ranging from 0.5 to 200 Hz, with amplitudes reaching up to 100 $\mu$V (Lago and Cester, 2017).

The frequency bandwidth of intracortical recordings can be increased even further using Utah arrays. These devices were developed in the 1980s at the University of Utah (Campbell et al., 1991) and consist of a grid of 100 needle electrodes that penetrate the cortex (Figure 2.1, *in vivo* mesoscale on the left). Instead of LFPs, these devices capture spiking data from individual neurons providing rich mesoscale information. For this reason, this has been the technology of choice in remarkable clinical studies where the brain activity of tetraplegic patients is decoded in real-time to control robotic arms (Hochberg et al., 2006, 2012), cursor on a screen (Simeral et al., 2011) and even fast handwriting prediction (Willett et al., 2021).

When working *in vitro*, the spatial resolution can get much higher. Brain slices or dissociated neuronal populations can be recorded using microelectrode arrays (MEAs) which may have dozens, hundreds or even thousands of microscopic electrodes covering a large portion of the neuronal ensemble (see section 2.3.2 for a brief overview of MEA technology or Obien et al. (2015) and Cerina et al. (2023) for comprehensive reviews). MEAs can capture both LFPs or individual

spikes, depending on the neuronal population cultured above: with brain slices, the high cell density generates LFPs; when using dissociated neuronal cultures, the sparsity of the network enables the recording of individual spikes, i.e., the extracellular current generated by the action potentials of neurons placed over a given electrode.

The mesoscale offers a window into the emergent firing dynamics of neuronal populations, allowing for the analysis of activity patterns and group phenomena like synchronization and oscillations, at the level of specific neuronal circuits. These collective dynamics are ubiquitous in healthy brain activity but, in certain conditions, are also the hallmark of specific neurological disorders. Typical examples include changes in synchronization of certain brain areas observed in disorders like epilepsy and PD (Uhlhaas and Singer, 2006). *In vitro* preparations are interesting models to study some of these mesoscopic activity patterns. Hippocampal brain slices, for example, may show spontaneous or induced epileptiform activity, which propagates according to the anatomical structures preserved in the slice. Hippocampal brain slices have played a crucial role in unravelling some of the mechanisms behind seizure generation, propagation and termination (Avoli et al., 2005; Pal et al., 2001; Librizzi et al., 2017; Weissinger et al., 2005; Losi et al., 2016).

Dissociated neuronal cultures are simpler *in vitro* models but have the advantage of allowing long-term experiments, as explained in more detail in section 2.3.3. These networks show emergent collective behaviour manifested as rich patterns of network-wide activation. These are similar to the bursts of activity observed in neonatal brains (Ben-Ari, 2001) or the inter-ictal epileptiform spikes (le Feber, 2019), i.e., synchronized epileptiform activity that occurs outside seizure periods.

*In silico* mesoscale models have also played an important role in posing possible mechanisms behind the phenomena recorded *in vivo* or *in vitro*. At the mesoscale, computational models are typically expressed as networks of individual spiking neurons. These networks can be built using a multitude of neuronal models. The HH model is a possible candidate and has been used to simulate *in vivo* mesoscale dynamics, such as the slow oscillations observed in the cortex during sleep or anesthesia (Compte et al., 2003). However, this model is computationally expensive, so it is not ideal for modelling large neuronal networks. Most mesoscale *in silico* studies focus on variations of Leaky Integrate and Fire (LIK) models, which simplify the neuron to a leaky integrator of synaptic inputs with a resetting threshold for the action potential.

Common non-linear variations of the LIK model are the Izhikevich neurons (Izhikevich, 2003) or Adaptive and Exponential LIK neurons (Brette and Gerstner, 2005). LIK models have been used to describe computational and physiological properties of neuronal systems recorded *in vivo* such as the Up and Down states (Holcman and Tsodyks, 2006; Destexhe, 2009) (shifts between two membrane potential levels that modulate neuronal excitability Steriade et al. (1993); Wilson and Kawaguchi (1996)); emergence of network synchrony (Vreeswijk and Hansel, 2001) and its employment for efficient information transfer (Buehlmann and Deco, 2010); computational mechanisms for working memory (Wang, 1999; Compte, 2000) and decision making (Wang, 2002); and

also pathological features such as the propagation of epileptic seizure waves (Ursino and La Cara, 2006).

The mesoscale analysis using *in vivo*, *in vitro* and *in silico* models has significantly expanded our understanding of the higher brain function. More importantly for this thesis, it had a significant contribution for our understanding of the pathological network dynamics behind neurological disorders like PD (Singh, 2018) or epilepsy (Farrell et al., 2019).

Finally, mesoscale electrophysiology technology offers an interesting opportunity: since the electrodes are in direct contact with the neural tissue, besides recording, they can also be used to electrically stimulate the network. At this scale, the causal links between external interventions and changes in population dynamics become clearer. More importantly, having the ability to directly stimulate specific neuronal populations opens the possibility of modulating malfunctioning circuits, possibly leading the brain to a healthier state, as we will discuss in the next section.

In this thesis...

The work presented in this thesis is focused on neuronal network dynamics and how we can modulate them using tailored stimulation protocols. As such, it will be circumscribed to the mesoscale. Throughout the three following chapters, we will present experiments performed *in vitro* using dissociated neuronal networks cultured on MEAs (Chapters 3 and 5) and *in silico* models of neuronal networks (Chapters 3 and 4), as represented in the blue box of Figure 2.1

## 2.2 Therapeutic electrical neurostimulation

This section will be focused on the use of direct electrical neurostimulation for therapeutic applications. It is divided into three parts: first, we explain how electrical stimulation is thought to affect the activity of the neuronal networks, then we elaborate on the currently available neurostimulation therapies and, finally, we present the literature on closed-loop stimulation algorithms, which we expect to be the direction of future neuromodulation.

### 2.2.1 The modulating effect of electrical stimulation

Neuronal communication is based on electrochemical signals. As such, it is possible to interfere with it using either chemical or electrical stimuli. Experiments using electrical stimulation have a long, intriguing and somewhat dark history that dates back centuries (Kellaway, 1946), long before modern medicine and pharmacology. The generation of bioelectricity in animal bodies was famously disputed between the anatomist and physician Luigi Galvani (1737-1798) and physicist Alessandro Volta (1745-1827) back in the 1780s. Galvani noticed that when he connected two metals to the exposed muscles of frog legs, these would move, suggesting the generation of electrical current within the biological tissue. His theory was not accepted then, as Volta, a renowned

scientist at the time and the father of the electric battery, argued that the two different types of metals were the source of the current. Aldini Giovani (1762-1834), Galvani's nephew, set to prove his uncle's theory of body electricity in a series of theatrical demonstrations throughout Europe where he used (ironically) the Voltaic battery to electrically stimulate the bodies of decapitated animals and even executed prisoners. The results shocked audiences as they thought they were witnessing the dead coming back to life.

The first experiment reporting the use of electrical brain stimulation in a conscious human is attributed to Robert Bartholow. Bartholow induced limb movements in a 30-year old patient, Mary Rafferty, using a needle electrode through her 50mm skull hole, caused by a cancerous ulcer (Bartholow, 1874). He reported how brain stimulation affected motor function, but his protocols and research goals raised serious ethical concerns. Rafferty died from a severe seizure, after waking up from a three-day comma.

Direct/invasive electrical stimulation of specific brain regions is now used to alleviate motor symptoms such as tremors and bradykinesia, but the underlying processes are not well understood yet. These may involve several physiological mechanisms operating at different temporal scales depending on the disease and stimulation site (Herrington et al., 2016). The long-term effects of brain stimulation, whether good or bad, can be, in part, attributed to the adaptation of the neuronal networks to the stimulus – a property known as *plasticity*. This process unfolds over periods of days, weeks – even months – and involves physiological and morphological changes in the neuronal substrate. Plasticity at the level of synapses may occur at a shorter time-scale, within minutes or hours (Herrington et al., 2016), as shown by (Shen et al., 2003), where they induced synaptic plasticity phenomena, such as short-term potentiation (STP), long-term potentiation (LTP) and long-term depression (LTD), by stimulation rat brain slices at 100 Hz.

Some long-term effects are reported after months of stimulation, such as improved mood, anxiety, sleep and somatic symptoms, in patients with severe depression (Mayberg et al., 2005; Lozano et al., 2008), which are most likely due to morphological changes. However, if properly tuned, electrical stimulation has an acute effect on the observed clinical symptoms which cannot be due to morphological changes in the network. Given the short time scales, the rapid effects of stimulation can only be explained by direct changes in the firing dynamics of the affected networks, whose spontaneous activity was likely nefarious. Indeed, certain neurological disorders are associated with malfunctioning brain circuits from where pathological neuronal activity may emerge. While microscale neurodegeneration may explain some features of the disease, many symptoms are thought to stem directly from the abnormal firing patterns within the affected networks (see McGregor and Nelson (2019) and Farrell et al. (2019) for a review on the circuit mechanisms of PD and epilepsy, respectively). Excessive neuronal synchronization is a typical example of abnormal firing dynamics and a hallmark of neurological disorders such as Parkinson's Disease (Volkmann et al., 1996; Hammond et al., 2007) and epilepsy – although recent evidence shows that the dynamics of synchronisation during seizures may be more complex than previously thought,

where desynchronization too plays a role in the unrolling of seizures (Jiruska et al., 2013).

Electrical stimulation interferes with these collective firing dynamics in ways that are not yet fully understood. Different – and apparently contradictory — hypotheses have been suggested to explain these short-term effects of stimulation. Some argue that stimulation has an inhibitory effect analogous to that of a brain lesion (Benazzouz et al., 1995) since it reduces/suppresses firing in the vicinity of the stimulation electrode (Boraud et al., 1996; Meissner et al., 2005; Hamani et al., 2010; Kiss et al., 2002; Anderson et al., 2006). The hypotheses put forward are based on 1) neuronal deactivation by an abrupt depolarization of their membranes, known as depolarization block (Hamani et al., 2010), 2) activation of inhibitory populations, or 3) masking the activity of the so-called "tremor cells" (those shown to fire in tune with the observed tremor) with high-frequency stimulation (Kiss et al., 2002).

On the other hand, there is plenty of evidence showing that high-frequency stimulation has an excitatory effect (Windels et al., 2000, 2003; Hashimoto et al., 2003) likely due to the axonal activation, including antidromic activation (propagation of action potential in the "opposite" direction, i.e., from the terminal portion to the soma), as reviewed by Deniau et al. (2010). The fact that electrical stimulation excites neuronal populations seems straightforward – what is less intuitive is why this activation leads to a decrease in clinical symptoms, which were associated with excessive neuronal activation in the first place.

An increasingly accepted theory suggests that, rather than inhibiting or exciting their target populations, the fundamental therapeutic effect of high-frequency stimulation is due to the disruption of information flow in the malfunctioning circuits, thereby suppressing the pathological activity (Chiken and Nambu, 2015; Florence et al., 2015). This disruption of collective neuronal dynamics could explain why high-frequency stimulation leads to network desynchronization in epilepsy, reducing pathological epileptic discharges like interictal spikes (bursts of activity in the periods between seizures) and high-frequency oscillations (Yu et al., 2018). It is also hypothesised that reducing synchronization may reduce seizure generation and propagation (Yu et al., 2018; de Castro Medeiros and Moraes, 2014).

Some have studied the transitions to pathological states, such as epileptic seizures, from a Dynamical Systems stand-point and argue that these transitions can result from falling into an attractor outside the boundaries of healthy brain function (da Silva et al., 2003; Jirsa et al., 2014; El Houssaini et al., 2020). Within this more abstract framework, electrical stimulation could be seen as the external input that drives the system away from such an abnormal state.

Having said all this, we must recognise that, despite being an active field of research, the implementation of stimulation therapies is not really built from first principles regarding stimulus-brain interactions. It is instead based on empirical trial-and-error optimisation of stimulation parameters, as we will see in the next section where we will briefly go through the current clinical solutions available.

Finally, as a side note on neuromodulation, emerging strategies are being developed to enable modulation with different types of stimuli, besides electrical. With optogenetics (Boyden et al., 2005), neurons are genetically modified to respond to light, enabling precise control over neural activity with cell-type specificity. Ultrasound can also be used to mechanically stimulate neuronal activity, being an effective non-invasive neuromodulation method (Tufail et al., 2011). Both methods show promising results, which could lead to innovative approaches to treating neurological disorders. Still, further advancements are necessary to translate these techniques into practical and safe clinical applications for human therapies. It is worth noting, however, that this thesis focuses solely on electrical modulation and does not cover pharmacological treatments, optogenetics or other potential therapeutic interventions.

### 2.2.2 Available neurostimulation therapies

The use of electrical brain stimulation for therapeutic applications dates back almost 100 years when the neurologist Ugo Cerletti first applied electroconvulsive therapy on a schizophrenic patient, triggering seizures that relieved his hallucinations (Cerletti, 1950). Despite evidence suggesting the efficacy of electroconvulsive therapy in controlling certain neurological and mental disorders, its clinical application was controversial. Electrical brain stimulation had a resurgence in 1987 when the French surgeon Alim Benabid discovered that the symptoms of PD patients could be improved by electrically stimulating the basal ganglia (Benabid et al., 1987). This marked the birth of DBS therapy, arguably the most iconic modality of brain stimulation.

Direct neural stimulation is used for multiple clinical applications and comes in multiple formats. Currently, there are two main areas where this intervention is seen as a viable clinical option: 1) **restoring** healthy neuronal function, such as in DBS for neurological disorders; and 2) **replacing** dysfunctional sensory input, such as cochlear implants for severe hearing loss. This thesis makes contributions to both categories – *restore* in Chapter 3 and 5 and *replace* in Chapter 4 – so we will describe them in more detail, along with the clinical options that are currently available. Other applications of neural stimulation, such as therapies for chronic pain, are out of the scope of this thesis.

#### Restoring healthy brain function

In this category, we place the situations in which some domain of brain activity is hijacked by pathological dynamics. And, even though we refer to this category as *restoring healthy brain function*, the neuronal activity is not the actual readout used in clinical appointments, but instead the associated observable symptoms. Still, these are the downstream manifestations of an underlying neuronal malfunction which typically correlates with identifiable biomarkers (as mentioned before for the cases of PD, tremmors or epilepsy). At this point, only a small subset of all potentially targetable diseases is electable for electrical stimulation therapy. We will briefly review those.

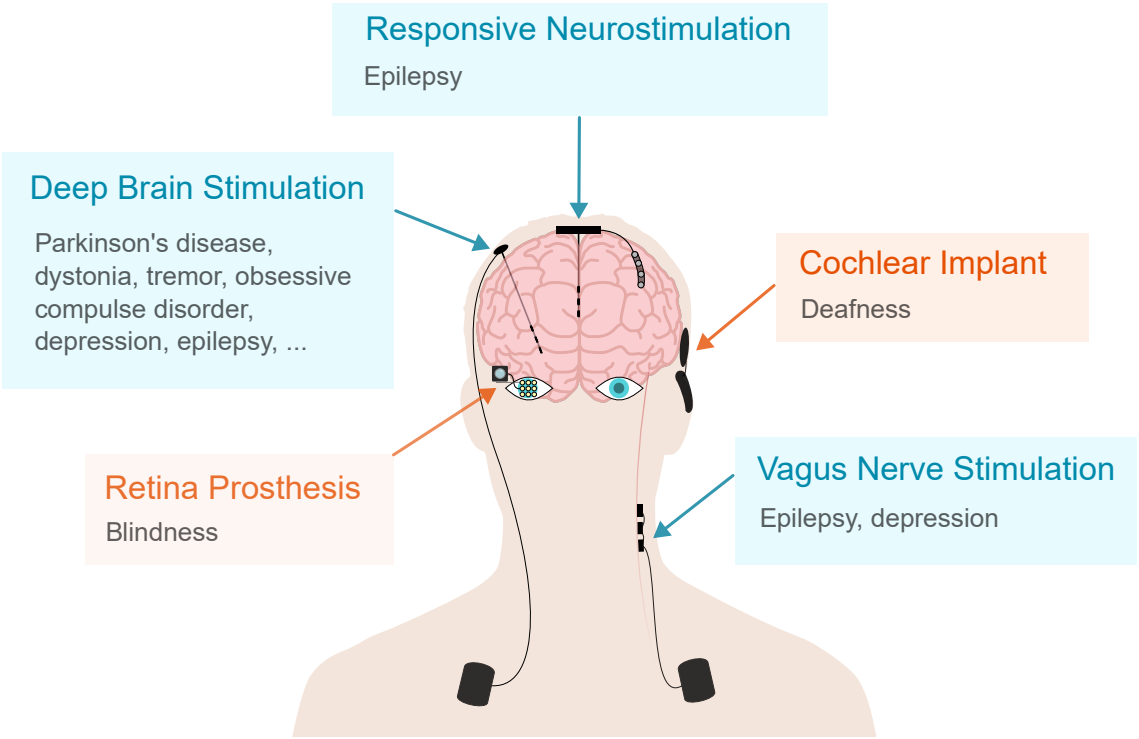



Figure 2.2: **Available neurostimulation devices.** Stimulation of the nervous system can be used as means of restoring healthy brain function (blue) of replacing lost sensory perception (orange). Regarding the first group, the most common therapies are: Deep Brain Stimulation (DBS), with applications for Parkinson's disease, dystonia, tremor, obsessive compulsive disorder, depression, epilepsy, among others; the Responsive Neurostimultor from NeuroPace for epilepsy control, the only responsive neurostimulator device in the market; and Vagus Nerve Stimulation (VNS), for epilepsy as well. As for replacing lost sensory function, we have: cochlear implants, which take the spotlight as the most successful implantable neurostimulator strategy; and retina prosthesis, which despite having a presence on the market, the available products are still far from satisfactory. The layout of this figure was inspired by the first figure from (Denison and Morrell, 2022)

There are different stimulation modalities, depending on the regions targeted – deep brain, cortex, vagus nerve or even spinal cord. DBS is the most popular invasive brain stimulation technique and is now a standard approach to alleviate motor symptoms in many PD patients (Mao et al., 2019) and essential tremor (Flora et al., 2010) as well as tremors in Multiple Sclerosis (Giordano et al., 2020) or tics in Tourette syndrome (Coulombe et al., 2019; Baldermann et al., 2016). It has also shown promising results in mitigating higher-order impairments, leading to memory improvement and reduction of cognitive decline in Alzheimer's disease (Luo et al., 2021), and improvement in life quality for treatment-resistant depression (Figee et al., 2022; Sheth et al., 2022). But, so far, DBS therapies have only received FDA approval for essential tremor, PD, dystonia and obsessive-compulsive disorder. Vagus Nerve Stimulation (VNS) has also received FDA approval as an alternative for refractory epilepsy, intractable depression, cluster headaches, and migraines (Wang et al., 2021). In VNS the stimulus is not applied directly to the brain, but instead to the vagus nerve, in the neck (Figure 2.2).

These recent developments underscore the growing significance of electrical neural stimulation in medical care. However, despite having the potential to treat several neurological conditions, the practical implementation of direct electrical stimulation is technically challenging. Currently, the rationale behind these therapies relies on identifying the brain region and stimulation parameters that alleviate symptoms. However, there is a large degree of freedom of parameters and anatomical targets which must be empirically optimised (Kuncel and Grill, 2004).

After implanting the electrodes, the clinician initiates a trial-and-error protocol to determine the appropriate stimulation parameters (pulse width, amplitude, stimulation frequency, etc.) that effectively alleviate symptoms. This iterative process entails a qualitative evaluation of the patient's clinical responses, and once effective parameters are identified, stimulation remains unchanged until the next appointment (Figure 2.3 top). One of the main problems of this continuous open-loop stimulation approach is that it fails to adapt to the dynamic nature of neurological diseases and the progression of the patient's condition. Moreover, the continuous stimulation can lead to unwanted adverse effects that are hard to mitigate, such as depression, cognitive decline and verbal fluency deficits (Wojtecki et al., 2015; Zarzycki and Domitrz, 2019).

In theory, these limitations could be overcome using closed-loop electrical stimulation that effectively adapts to the ongoing brain activity in real-time, assuring minimal stimulation (Figure 2.3 bottom). With such a closed-loop setup, the stimulation parameters (stimulation frequency, amplitude, etc.) are determined by the control algorithm. The clinician would instead be in charge of the parameters of the controller (thresholds, control gains, etc.).

Currently, there is only one closed-loop stimulator available in the market: the Responsive NeuroStimulator (RNS) from Neuropace®, for seizure control (Figure 2.2). This device performs on-chip data processing to detected seizure onset and trigger a predefined stimulus burst, meant to disrupt the seizure. The parameters for seizure detection must be tuned by the clinician, as well as the stimulus pattern. The RNS device is with no doubt innovative in this regard, but is still based

on the simplest form of closed-loop control. In section 2.2.3 we will discuss in detail the different closed-loop strategies proposed over the last three decades for neurostimulation. This will also set forth the motivation for the work presented in Chapters 3 and 5.

#### Replacing lost sensory perception

Another domain where neurostimulation has great potential is for restoring lost sensory function. Our bodies have the ability to transduce physical stimuli, like air vibration and electromagnetic radiation, into discernible electrical signals that the brain decodes into sensory perceptions, such as sound and imagery, respectively. That requires specialised sensory cells like the hair cells in the inner ear and photoreceptors in the retina. Certain diseases affect these cell types leading to deafness or blindness and direct electrical stimulation can be used to encode the sensory information directly into the nervous system.

This problem is conceptually different from the previous one in relevant ways, and, accordingly, the challenges in neuromodulation challenges are also different. In the previous category, the challenge was finding the ideal brain region and choosing, "blindly", effective stimulation parameters. Here, the region to stimulate is well defined and so is the intended neuronal activation: one wants to activate a particular sensory pathway in a way that mimics the response to a natural stimulus. The challenge lies in knowing how the natural stimulus is encoded into neural activity and how one can trigger the same activation using external electrical stimulation.

Cochlear implants are a prime example of such technology and, with hundreds of thousands of registered devices implanted (Nassiri et al., 2022), is by far the most successful neural stimulator (Figure 2.2). The implants exploit the fact that the different frequency components of sound waves are encoded separately along the cochlea. These devices have an external microphone to capture ambient sound, a microprocessor that divides the sound into its frequency components, and a set of electrodes implanted along the cochlea which target different portions of the cochlea according to its frequency-encoding disposition (Zeng et al., 2008).

Stimulation in cochlear implants is performed in open-loop since it does not adjust according to the neural response elicited by the stimulus. Just like DBS and other brain stimulation techniques, cochlear implants require a parametrization stage performed by the technician, based on patient feedback (Laughlin et al., 2012). Optimising the stimulation parameters for a given patient can take several hours and often requires recurrent adjustments during the first months of use (Dorman et al., 1992; Hughes et al., 2001; Gordon et al., 2004).

Another promising application of neural stimulation is restoring visual perception. Diseases such as macular degeneration or retinitis pigmentosa lead to the death of photoreceptors, responsible for transducing light into "neuronal language". Precise spatiotemporal electrical stimulation of the remaining retinal layers using implantable microelectrodes could be used to restore visual perception. These devices have an external camera whose image is processed and transformed into a stimulus pattern, sent to the retina through an implanted MEA (Figure 2.2). However, so far,

the products in the market are not capable of providing a satisfactory experience for the patients (Erickson-Davis and Korzybska, 2021). Users report that their "artificial vision" was actually an acquired skill-set that required interpreting the flashes induced by the stimulator. Using the device was reported as being cognitively fatiguing and with no practical benefits for their daily life (Erickson-Davis and Korzybska, 2021).

The two major limitations are the very low spatial resolution due to the reduced number of electrodes and the fact that stimulation, being open-loop, does not take into account the complex ways through which the retina responds to stimulation. These will be some of the points addressed in the publication presented in Chapter 4

**Conclusions on the state of the art of conventional stimulation therapies**

Given the absence of first-principle knowledge guiding the parametrization of electrical stimulation, the clinical procedure is based on laborious trial-and-error optimisation protocols. Research is now moving towards the development of closed-loop approaches, which are expected to improve efficacy while decreasing the amount of stimulation provided.

As we mentioned before, the RNS device stands out in market for being the only one featuring closed-loop capabilities. However, recent developments indicate that this approach may be adopted by other companies. The latest DBS device from Medtronic, the Activa™ PC + S, has recording capabilities and its hardware supports the implementation of closed-loop algorithms, but only for research purposes at the moment.

The trajectory of emergent and promising start-ups point in the same direction. Two recent players in the market are Synchron and Neuralink. Sychron's device, the Stentrode, developed at the University of Melbourne, is capable of recording (Oxley et al., 2016) and stimulating (Opie et al., 2018) neural activity using a stent-based multielectrode platform that accesses the brain activity through its vessels, requiring no surgical intervention. At this point is envisioned mostly for Brain Machine Interface (BMI) applications, where the goal is to decode brain activity to control external devices, however, its stimulating capabilities open the opportunity for neuromodulation as well (John et al., 2019). Neuralink, Elon Musk's BMI startup, is focused on expanding the bandwidth of the brain-device interaction by significantly increasing the number of recording/stimulating electrodes up to more than 3000 channels Musk (2019). Even though this device is still in early developmental stages, it has on-chip processing capabilities meaning that, in principle, it can be used for closed-loop applications.

These advancements show the interest of the community in moving towards more advanced stimulation devices and protocols. However, that will require the development of closed-loop algorithms, and their subsequent validations in increasingly more realistic neuronal models. This is going to be a long process involving contributions from multiple fields. As we will see in the next section, this collaborative effort started a few decades ago and is still far from being a settled issue.

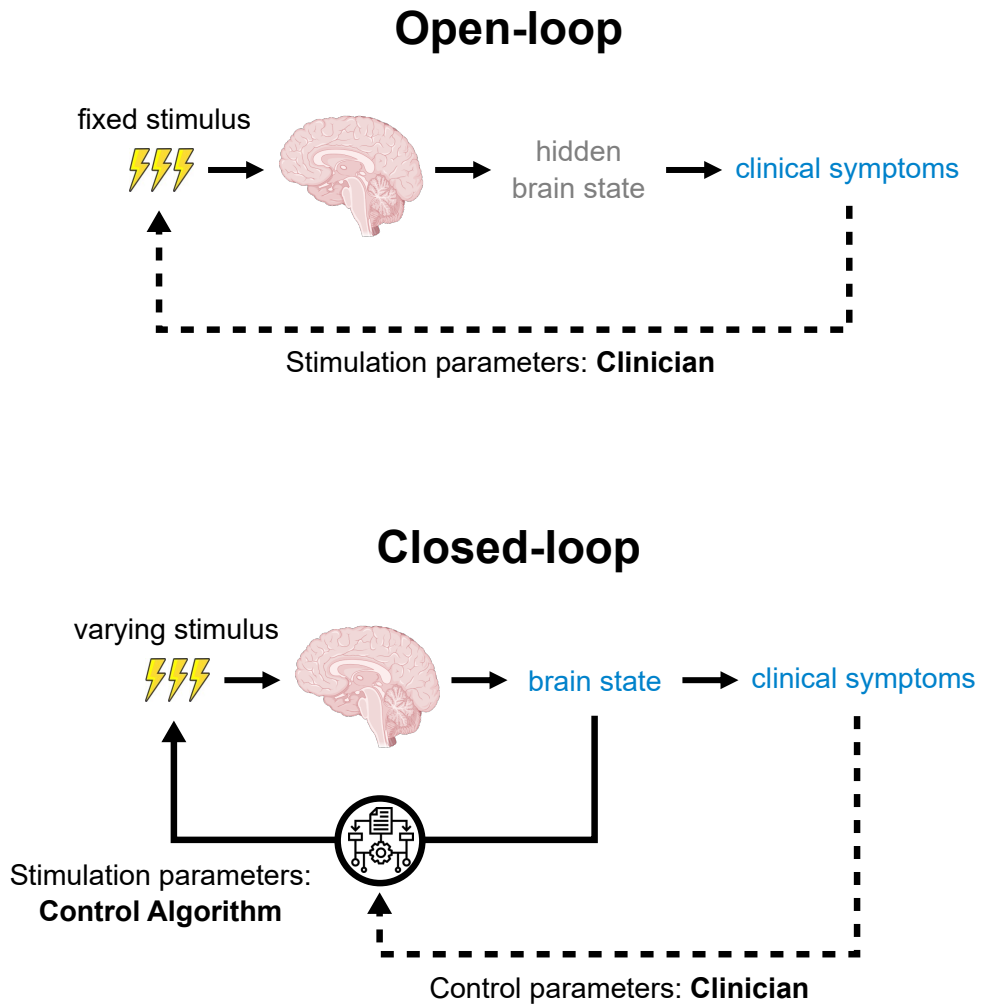



Figure 2.3: **Open-loop and closed-loop neural stimulation.** In open-loop stimulation (top) the stimulus parameters, such as pulse frequency, amplitude, width, etc., are fixed and the stimulation runs continuously. The brain response to stimulation is not observable since these devices have no recording capabilities. Thus, the stimulation parameters are tuned by the physician during clinical appointments based on the observed symptoms. In closed-loop stimulation (bottom), the stimulation parameters are adjusted in real-time, based on the recorded brain activity. The clinician is in charge of tuning the control parameters (such as control gains or thresholds), based on the clinical symptoms.

### 2.2.3 Closing the loop on neurostimulation

A shared limitation between the two neurostimulation categories described in the previous section is that the fact that the stimulus does not adapt to the neural response or ongoing activity. With such open-loop designs, it is virtually impossible to achieve optimal stimulation over long time spans due to the complexity and plasticity of neuronal systems.

Throughout this thesis, we will see neurostimulation as a control task: we have a system that must be controlled and the means to modulate its behaviour. Within this framework, stimulation could be improved with a closed-loop setup, as is usually done in typical control problems in engineering. In a closed-loop paradigm, the input provided to the system is a function of the desired state (reference) and the current state of the system. That requires essentially two additional components: a sensor to measure the state of the system and a rule that dictates how to adjust the input provided to the system. The first component is easily addressable since the electrodes used to stimulate can also be used to measure neural activity. The second component – the rule used to adjust the input – is the big unknown. The challenge stems from the fact that 1) it is not clear how one should characterise the *state* of a neural system based on the recorded neural activity, and 2) the interactions between stimulus and neural activity are highly nonlinear and difficult to predict, especially across long periods of time.

In this section, we will expand on the closed-loop algorithms – the *rules* – that have been explored over the last two decades for neurostimulation. We want to lay out the landscape of control strategies already implemented so that the novel work presented in this thesis can be better contextualised. But, before moving on to the algorithms themselves, it is important to categorise the different types of closed-loop setups found in Neuroengineering studies and to distinguish those relevant to this thesis. We can separate two main groups (Figure 2.4): 1) algorithms to control neuronal activity; and 2) systems that use neuronal activity to control an external body.

The first category is focused on developing closed-loop methods to control the activity of neuronal systems using electrical stimulation (Figure 2.4 left). The most obvious context is therapeutic closed-loop brain stimulation, which is also the realm of this thesis. This is where we would include closed-loop DBS and other analogous systems envisioned to treat neurological disorders, which can be developed and tested using different neuronal models, be it *in vivo*, *in vitro* or *in silico*. Other uses of closed-loop neuronal stimulation can be simply to better control experimental conditions by assuring stable neuronal activity. We named this use-case "Control Network Variables", and is typically used in an *in vitro* context, to study fundamental neuronal physiology.

The second category is focused on using neuronal activity to control external objects such as prostheses or robots (Figure 2.4 right). The prime example is Brain Machine Interfaces (BMIs), also referred to as Brain Computer Interfaces (BCIs), where a brain decoder is used to interpret the patient's intentions and control a robotic arm or the cursor on a screen. This is typically done in

open-loop (only from brain to computer), but the loop can also be closed by incorporating sensors and an algorithm that encodes the sensors' data into direct electrical stimulation. Closed-loop BMI is out of the scope of this thesis so we direct interested readers to reviews on the subject such as Pan et al. (2022); Belkacem et al. (2023). This category of closed-loop system can also be seen *in vitro*, in the realm we refer to as "Neuronal Cultures Embodiment". Here, an *in vitro* neuronal network is used to control the actions of an external entity, be it a physical robot (**?**), or an agent in a simulated environment (DeMarse and Dockendorf, 2005; Kagan et al., 2022). The goal of these studies is generally to unravel the fundamental mechanisms behind emergent neuronal processing. This too is out of the scope of this thesis and we direct readers interested in Culture Embodiment to the book chapter written by Bakkum et al. (2004).

Now that we are better contextualised on the class of closed-loop algorithms relevant to this thesis – those that focus on controlling neuronal activity – we present a brief review of the methods developed over the last decades. We devised a conceptual map, Figure 2.5, which categorises the different control algorithms along two dimensions: 1) Magnitude/Phase Responsive; and 2) Event Triggered/Continuous Control.

The first dimension is associated with the nature of the signal used to characterise the system. Magnitude-responsive methods monitor and control the magnitude of a given metric (or metrics) across time. This could be the strength of the neuronal response to stimulation or the power at a given frequency band, for example. Phase-responsive methods operate on the assumption that the monitored signal is periodic (or quasi-periodic) and act according to the phase of the ongoing oscillation. The second dimension is related to the nature of the actuation. Here we differentiate between the methods where a predefined stimulus is triggered by specific events – Event Triggered – from the methods where the stimulus is an explicit function of the monitored variable – Continuous Control.

It is important to note that, despite establishing a useful conceptual map, the boundaries between these four categories are not always clear-cut. Also, the methods were displaced on the map according to the specific context in which they were implemented for neuronal control, and not simply based on the nature of the algorithm itself. For example, Intelligent Control algorithms are placed in the Magnitude-Responsive Continuous Control quadrant because that is how AI has been used in the neuromodulation literature so far. Nonetheless, AI methods are flexible enough to be applicable in Phase Control as well. Yet, this flexibility does not apply to methods such as ON-OFF, which are inherently Event Triggered, or Phase-Locking which is intrinsically phasic, for example. Finally, we reinforce that this map focuses only on the control methods employed and not on the methods used to derive the metrics/variables they may follow. For instance, a hypothetical implantable device that triggered a predefined stimulus based on real-time seizure prediction performed by a sophisticated AI model would still be considered an ON-OFF controller.

The different methodologies were implemented in models that differ in complexity, ranging from clinical studies, to *in silico* simulations. As such, for each class of algorithms, we will address

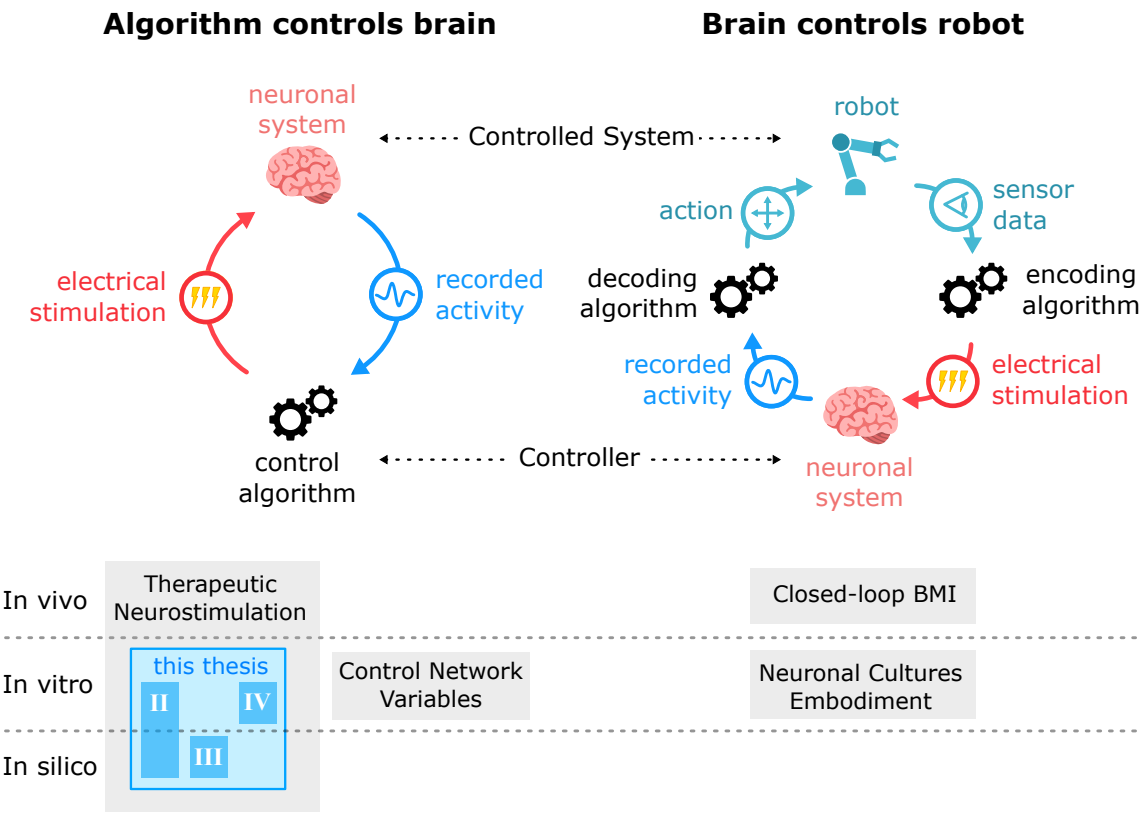



Figure 2.4: **Types of closed-loop electrical stimulation systems in Neuroengineering**. Closed-loop stimulation can be applied in two completely different tasks: either controlling neuronal activity (left) or using neural activity to control external bodies (right). The first type of closed-loop system is typically seen in cases where the brain activity is pathological and needs to be controlled. The locus of control is a closed-loop algorithm that receives the recorded neuronal activity as input and adjusts the electrical stimulation as output. We refer to this application as "Therapeutic Neurostimulation". The systems and the algorithms can be developed and tested with different types of neuronal models, *in vivo*, *in vitro* or *in silico*. Closed-loop stimulation of neuronal activity can also be used to control experimental conditions *in vitro* to better unravel the input-output function of the network. We refer to this type of application as "Controlling Network Variables". The other completely different application of closed-loop neural stimulation is for controlling external bodies. Here, the locus of control is the neuronal system, and not an external algorithm. The classical example is closed-loop Brain Machine Interfaces (BMI), where the brain activity of a paralysed patient is decoded to control a robotic arm or mouse cursor on a computer. The loop is closed by incorporating sensors into the system and encoding the sensor information into electrical neural stimulation. Analogous setups have been developed *in vitro* to study the capacity of neuronal networks to learn and adapt to external environments. This type of *in vitro* application is typically referred to as "Neuronal Cultures Embodiment". Our contributions with this thesis are focused on Therapeutic Neurostimulation, with studies performed both *in vitro* – Chatpers 3 and 5 – and *in silico* – Chapters 3 and 4 – as shown in blue.

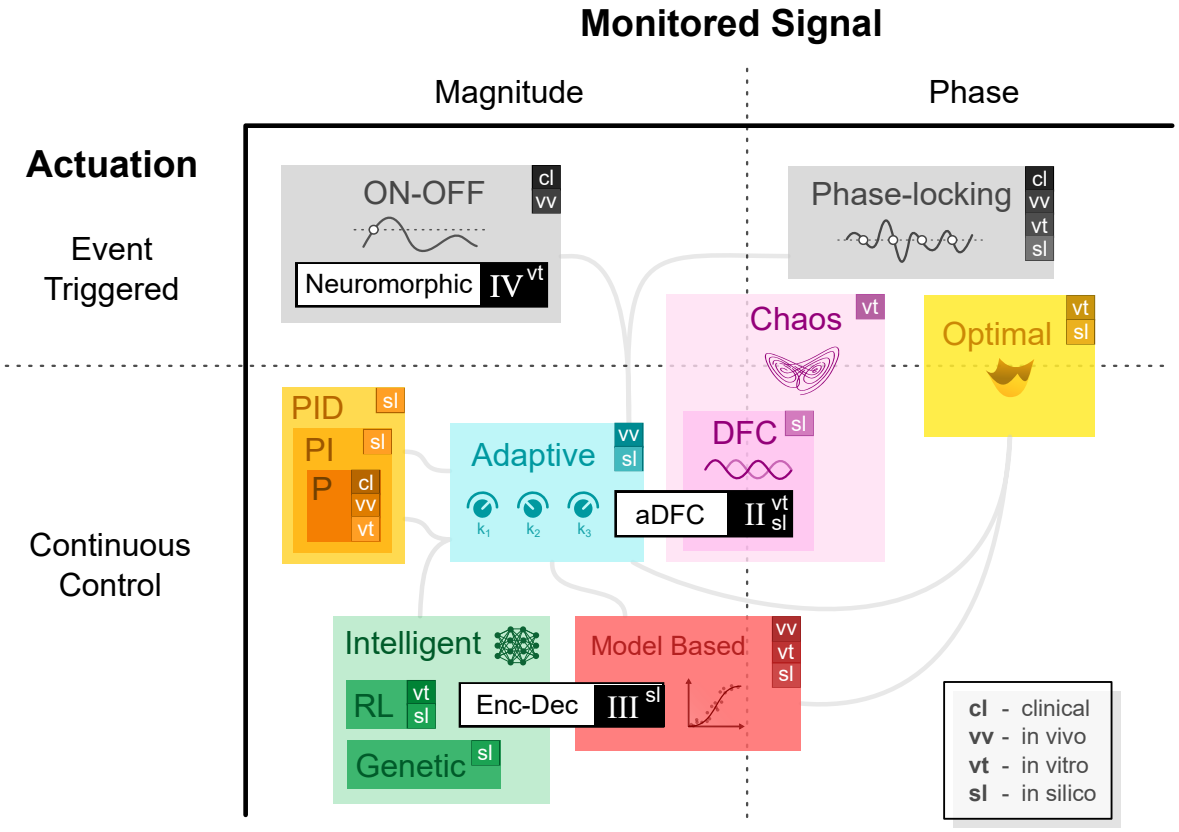



Figure 2.5: **Closed-loop electrical neuromodulation algorithms.** The algorithms are organised according to the nature of the monitored signal and the nature of the actuation. The type of monitored signal is presented along the X-axis and divided into "Magnitude" and "Phase". The first refers to algorithms that monitor and control the magnitude of a given metric of neuronal activity, such as the power of a given frequency band. In "phase", we group the algorithms that track the phase of oscillatory phenomena or the timing of specific neuronal events. As for the nature of the actuation, presented along the Y-axis, we group the algorithms into "Event Triggered" and "Continuous Control", distinguishing those that only actuate when specific events are detected from those that are continuously running a control algorithm. The algorithms were placed on the map according to how they were implemented in the specific studies for neuronal control. We identify the neuronal models on which the methods were tested with the keys: **cl**, **vv**, **vt** and **sl**, which stand for clinical, *in vivo*, *in vitro* and *in silico* models, respectively. Some studies merge different control methods, which we represent with grey lines linking the different boxes. Our contributions are presented with white boxes, representing both the method name of the method (aDFC, Neuromorphic, Enc-Dec), the Chapters (II, III, IV), and the neuronal model used (*in vitro* and *in silico*)

the most influential studies according to the following hierarchy: devices currently available in the market; clinical studies; *in vivo* studies; *in vitro* studies; and finally *in silico* studies. Naturally, as the sophistication of the algorithms increases, the models they are tested on become simpler. We will describe this landscape according to the complexity of the methods employed, starting with the simplest ones, already implemented in clinical settings: ON-OFF control, Phase-Locking and Proportional Control.

### ON-OFF Control

The simplest and most common approach to closed-loop electrical stimulation is ON-OFF Control. This controller follows a simple rule: if the magnitude of the monitored metric reaches a certain threshold, a predefined stimulus pattern is activated. The only control metric is then the level of the threshold (the stimulus parameters must also be optimised by the clinician, but these are not parameters of the controller). The only FDA-approved closed-loop brain stimulator – the Response NeuroStimulator, RNS from Neuropace®– operates with such an ON-OFF controller. Devised for the treatment of drug-resistant epilepsy, this device monitors one or two seizure foci and when the recorded activity meets certain criteria (optimised by the clinician for each patient) brief pulses of electrical current are delivered to reestablish normal activity. Despite its simplicity, RNS has shown remarkable results, with a reduction of 41.5% mean seizure reduction 5 months post-implant, and a 75% median reduction after 9 years (Jarosiewicz and Morrell, 2020). Several *in vivo* studies reinforced the efficacy of ON-OFF stimulators for the control of epileptic seizures (Salam et al., 2015b,a; van Heukelum et al., 2016; Kozák and Berényi, 2017). Despite differences in monitored variables, stimulation parameters and sites, they all rely on a predefined stimulus triggered by seizure detection.

ON-OFF controllers have also been explored for the treatment of PD. Clinical studies have used the power of the beta band (Little et al., 2013, 2014, 2015; Velisar et al., 2019), gamma band (Swann et al., 2018) or the rectified beta band-passed LFP signal (Tinkhauser et al., 2017) as control variables. Results show a reduction in motor symptoms with considerably less stimulation time. Some have proposed variations to the vanilla ON-OFF controller, such as a dual threshold on the beta band (Velisar et al., 2019) or a controller whose stimulation shifts between high and low voltage, instead of ON and OFF (Swann et al., 2018). The latter two studies were performed using the Activa™ PC + S neurostimulator from Medtronic, which has closed-loop capabilities incorporated into the hardware, only available for research purposes at the moment (Parastarfeizabadi and Kouzani, 2017).

The beta band is typically used as a control signal because it correlates with motor symptoms (Kühn et al., 2006; Kuhn et al., 2008). Nonetheless, experiments performed on a non-human primate PD model noticed that, in reaching tasks, only continuous DBS improved bradykinesia, probably because this movement led to a natural decrease of the beta band, affecting the closed-loop DBS (Johnson et al., 2016). Thus they questioned whether following a physiological signal, such as the beta band power, would be enough to control motor symptoms. As a side note, in the

PD literature, closed-loop stimulation is commonly referred to as *adaptive* stimulation or *adaptive* DBS (aDBS), because the stimulation parameters (stimulation timing, frequency, amplitude, etc.) adapt to the recorded activity in real-time. However, we will not use the term "adaptive" in this context because, in Control Theory, Adaptive Control is not a synonym for closed-loop control. It instead refers to the class of controllers that self-tune their internal parameters (such as the control gain), as we will see later.

### Phase-Locking

After the ON-OFF controller, the simplest control method for neuronal control is Phase-Locking. We are now on the left panel of Figure 2.5, with phase-responsive stimulation methods, which assume that the signal is periodic and the effects of stimulation are phase-dependent. In phase-locking, the stimulus is triggered by a specific phase of the neuronal oscillation cycle, the sole parameter to optimise. This method has been explored for diseases characterised by the presence of pathological oscillations and was proposed to be an alternative to continuous high-frequency brain stimulation by using precisely timed low-frequency stimuli. Cagnan et al. (2013) tested phase-locked stimulation to treat the symptoms of essential tremor and reported decreases in tremor amplitude of up to 87% (Cagnan et al., 2016). However, the monitored signal was not brain activity but, instead, the wrist movement captured with an accelerometer, and the phase-locked stimulation was only tested for 30 seconds, on average. Rosin et al. (2011) tested different open-loop and closed-loop stimulation protocols on a primate model of PD and showed that closed-loop stimulation was optimal when the stimulus was delivered 80 ms after an action potential was detected on the recorded brain region. The delay of 80 ms after the spiking activity was thought to interfere with the pathological synchronization cycle. Computational studies support this hypothesis using Phase Response Curves (PRCs) (Azodi-Avval and Gharabaghi, 2015; Holt et al., 2016). PRC analysis reveals how external perturbations shift the phase of an oscillating neuronal system, which may lead to an increase or decrease in network synchronization (Smeal et al., 2010). All the advanced phase-responsive stimulation methods presented later are based on the analysis of PRC.

### Proportional Control

The last control method ever tested in a clinical setting is Proportional Control (marked as P in Figure 2.5, inside the PID box), a Magnitude-responsive method and the simplest form of Continuous Control. In proportional control, the intensity of the actuation signal (stimulation voltage or frequency, for example) is proportional to the monitored metric or the error regarding a target value. It is the simplest form of continuous control where the only tuning parameter is the gain that converts the monitored signal into an actuation input. Proportional control was tested on PD patients, and the results showed an improvement in motor symptoms compared to continuous stimulation (Rosa et al., 2015; Arlotti et al., 2018). In these studies, the DBS frequency was constant and voltage was proportional to the power of the beta band. Proportional stimulation was

also tested in the context of epilepsy *in vivo* using a model of penicillin-induced episodic seizures in rats, where the stimulation signal was proportional to the recorded EEG signal (Colpan et al., 2007). As for *in vitro* experiments, Gluckman et al. (2001) used a proportional controller to either suppress or enhance epileptiform activity in brain slices, and Wagenaar et al. (2005b) showed that the average firing rate of neuronal networks could be reliably controlled using a proportional controller to determine the stimulation voltage of multiple microelectrodes, that activated sequentially at a fixed stimulation frequency.

#### Proportional-Integral-Derivative Control

An upgrade to proportional control is Proportional-Integral-Derivative Control (PID). This is probably the most popular type of controller used in engineering problems due to its simplicity and efficacy in tracking a given target value. In PID, the actuation signal is not only proportional to the error between the current value and the target but is also proportional to the integral and derivative of the error. Each component has an associated gain that must be properly tuned. When the derivative gain is set to zero, the controller is referred to as a PI controller. This is a common practice in neuronal control applications because the noisy nature of neuronal activity often makes the derivative component ineffective. It is worth noting that although Proportional-Derivative is also possible (usually referred to as PD), it has never been implemented for neuronal control. Therefore, throughout this thesis, the term 'PD' is only used to refer to Parkinson's Disease. PI was already used *in vitro* with patch clamp stimulation when experiments required a given neuron to fire with a fixed inter-spike interval (Miranda-Domínguez et al., 2010; Couto, 2015). For therapeutic applications, PID (and PI) was only tested *in silico*, both in the context of Parkinson's Disease (Pyragas et al., 2007; Gorzelic et al., 2013; Fleming et al., 2020a; Liu et al., 2016, 2021) and epilepsy (Wang et al., 2016a).

#### Adaptive Control

All the control parameters presented so far require some sort of parametrization that must be optimised. A simple ON-OFF controller requires setting the threshold that triggers stimulation; phase-locking requires setting the stimulus phase delay; PID requires setting the P, I and D gains; etc. These parameterizations are generally performed empirically or based on some stability criteria that assure predictable and stable behaviour of the system. Still, in real-world problems, the systems may change over time, or present complex non-linear dynamics. In these situations, the fixed parameterization may lead to poor performance and unstable behaviour. Adaptive Control aims to solve this problem with methods that self-tune the control parameters online.

The adaptive approaches presented here are based on control methods already presented, to which an adaptive layer was added. These were all tested in computer simulations, probably due to the increased complexity of the approaches implemented. Also, we placed this control type in the Continuous Magnitude Responsive Control quadrant of Figure 2.5 because that is how it has been commonly implemented. For example, Fleming et al. (2020b) developed a P controller whose

gain was defined by a differential equation that was a function of the pathological activity. Yang et al. (2022) merged P control with AI to suppress epileptic seizures, using a Radial Basis Function ANN to compensate for the limitations of the P controller in real-time. Rouhollahi et al. (2019) implemented an adaptive version of the Sliding Model control technique, a non-linear control method. An interesting exception to the Continuous Magnitude-Responsive adaptive algorithms presented before is the Event-Driven, Magnitude and Phase-Responsive adaptive control algorithm presented by Grado et al. (2018). Here, they merged ON-OFF control (stimulation is triggered when the beta power is above a certain threshold) with phase-locking (stimulus is only sent at a specific phase of the oscillation cycle) and used Bayesian optimisation to tune the oscillation power threshold, the oscillation phase trigger and the stimulus amplitude.

#### Model-Based Control

Another important type of control paradigm is Model-Based Control. Here, the neuronal system is first described by a mathematical model that captures the stimulus-response relationships. The model is designed by probing the system with a wide range of stimulation settings and storing the associated neuronal responses. In its simplest form, once the input-output model is obtained, the stimulus that achieves a desired neuronal activation is determined by inverting the model. Ward et al. (2015) used this model-based approach to adjust the stimulation parameters of Vagus Stimulation in a rat model, allowing them to control the level of activation of different types of nerve fibres. This type of empirical curve-fitting model is often called a *descriptive model* in the Computational Neuroscience literature. Besides this *in vivo* experiment, the other model-based studies we found were performed *in silico*. Model estimation and inversion were also implemented in the context of retina stimulation for visual prosthesis Spencer et al. (2019) using a descriptive model of retina response to multipolar stimulation previously obtained with *in vitro* experiments (Maturana et al., 2016).

The fitted model can also be used to tune another controller analytically. For example, Zhang et al. (2022) developed an auto-regressive model to estimate the dynamic response of the neuronal system's firing rate to electrical stimulation. They then used this model to parametrize PI for seizure control using the Routh-Hurwitz stability criterion. Similarly, Su et al. (2019) and Santaniello et al. (2010) built an auto-regressive model to estimate the stimulus-response relations and used it for online Adaptive Control of pathological activity by adjusting the parameters of a PI controller, as in the case of Su et al. (2019), or adjusting the estimated model itself in real-time, as in the case of Santaniello et al. (2010). Another common approach is Kalman-filtering, which has been used in multiple studies to estimate the unobservable states of the neuronal system and design/tune different control strategies with applications for PD and epilepsy (Schiff and Sauer, 2007; Schiff, 2010; Liu and Gao, 2013).

Another way of characterising the stimulus-response relations in neuronal systems is through a Phase-Response Curve (PRC). As described above, the PRC describes the phase-shift caused by stimulating the system at a particular phase of the oscillation cycle. This curve can be estimated

experimentally by probing the system at multiple phases, and then used to control neuronal activity. Faramarzi and Netoff (2021) used the PRC-based algorithm proposed by Matchen and Moehlis (2018) to divide the synchronous oscillations of a brain slice into four distinct clusters of synchronous oscillations. PRC-based algorithms have been mostly implemented in the context of Optimal Control, as we will see in the next section.

### Optimal Control

In many real-world applications, the controller must be able to reach a desired target value while being submitted to a defined set of constraints, such as power consumption. Optimal Control focuses on determining analytical control laws that minimise a given cost function or maximise some performance criteria, based on the model of the system under control. In the context of Neuroengineering, Optimal Control has been mostly used to modulate the phase or firing time of oscillating neuronal activity, based on PRCs as the input-output model of the system. Nabi et al. (2013) used this strategy to control the firing time of individual *in vitro* neurons after the detection of an action potential using patch clamp, similar to previous theoretical studies (Moehlis et al., 2006; Nabi and Moehlis, 2009; Dasanayake and Li, 2011). Similar Event-Based Optimal Control strategies have been presented in multiple *in silico* studies where the goal is to find the stimulus waveforms that drive neurons to a phaseless state after the detection of an action potential, enabling network desynchronisation (Nabi et al., 2012; Danzl et al., 2009).

An exception to these Phase-Responsive methods is the approach proposed by Liu et al. (2016). Similarly to the model-based adaptive algorithm from Santaniello et al. (2010), they used an auto-regressive model to estimate the input-output relations of their PD computational model, which was tuned at every step, making it adaptive. Then they derived the optimal control law which balanced the efficacy in suppressing pathological activity with the energy expenditure of the stimulation.

### Intelligent Control

In the previous section, we saw methods that derived the optimal control law according to a given cost function, based on an estimated mathematical model of the system's response to stimulation. But the control law can also be extracted from the data directly, without an explicit estimation of the system under control. This is the rationale behind AI-based control, from which Reinforcement Learning (RL) is the prime example. In RL, a machine learning model (an ANN, for example) is trained by interacting with the environment, rewarding desired control actions and punishing undesired ones, given the current state of the system. In the context of neuronal control, RL was first implemented in an *in vitro* model of limbic seizures (Pineau and Avoli, 2009; Guez et al., 2008) with experiments performed on acute hippocampal slices. The algorithm learned the quality of each action for a given state, the *Q function*, using random trees. However, the training was conducted offline after collecting a vast stimulus-response dataset, and not in an online trial-and-error fashion, as is typical of RL. Later, the same group tested a model trained offline on brain

slices never seen by the controller, and the results showed no improvement in seizure reduction compared with open-loop periodic stimulation, although using less stimulation (Panuccio et al., 2013). Still, the training was not made by trial-and-error exploration of the controller. Wülfing et al. (2019) used RL in an online setting to control the response strength of neuronal cultures to stimulation using MEAs. They showed that deep Q learning, where an ANN is used to estimate the Q function, could learn the rules governing stimulus-response in neuronal cultures through real-time interaction with the system. Until now, RL was tested mostly *in silico*, where the constraints of limited data or training time do not apply (unless they are purposefully built into the study setup to simulate clinical conditions). The computational studies focused on applications for seizure control using TD-learning (Nagaraj et al., 2017), and treatment of PD with Actor-Critic RL (Lu et al., 2019; Gao et al., 2020, 2022).

Other types of AI used for control of pathological neuronal activity include: simply ANN trained with back-propagation (Liu et al., 2020a), without the RL framework; Fuzzy Logic (Liu et al., 2015), which provides a flexible approach for decision-making in situations involving imprecise/uncertain information; Genetic Algorithms (Feng et al., 2007), an iterative searching method inspired by natural selection that tries to optimise a given cost function.

#### Chaos Control

Finally, some studies focused on using principles from Chaos Theory to control neuronal activity. Chaos is found in systems that, despite being deterministic, are very difficult to predict due to their sensitivity to small perturbations. Chaos Control exploits the observed behaviour of a chaotic system to stabilise very complex dynamics using minimal perturbations. The rationale behind employing Chaos Control for neuromodulation is based on the fact that neuronal systems display some features that can be categorised as chaotic, in particular, unstable fixed-point attractors that lead to Unstable Periodic Orbits (UPOs). This behaviour can be observed by examining the relations between the timings of sequential events. Using patch clamp recordings, Schiff et al. (1994) argued that the activity of hippocampal slices presents chaotic behaviour, evidenced by plotting each interspike-interval (ISI) versus the previous ISI (first-return plot). The evolution of the ISIs exhibited a convergence of ISIs along a stable trajectory towards an unstable fixed point, from where the trajectory diverges along an unstable direction. This behaviour is characteristic of a saddle-node bifurcation. They then exploited this local geometry to stabilise the ISI by timing stimulation pulses so that the triggered ISI stays on the stable trajectory, therefore increasing the periodicity of the system. They also performed "anti-control" where they pushed the system away from the stable trajectory, making it aperiodic.

Interestingly, in the following year, this breakthrough study was challenged by Christini and Collins (1995) who repeated the experiments using simple stochastic non-chaotic neuronal models and obtained very similar results to those presented by Schiff et al. (1994). They argued that Schiff et al. (1994) may not have controlled actual chaos in the brain slices, while reinforcing the idea that chaos control techniques may be applicable even for non-chaotic systems. Khosravani

et al. (2003) used similar concepts to abort seizures in hippocampal brain slices. They analysed the dynamical structure of the inter-ictal to ictal transitions with ISI first-return plots and used it to trigger stimulation once the transition was detected. It is important to note that, even though their return plots showed signs of chaos, they do not claim to be dealing with a chaotic system due to the difficulty in assuming that from a time series.

One of the most popular methods to control chaotic systems is Delayed Feedback Control (DFC), first proposed by Pyragas (1992), which is a good example of a Continuous Control method that is both magnitude and phase-responsive. DFC was designed to stabilise one of the multiple UPOs of a chaotic attractor to a periodicity T by feeding back the delayed state of the system from a period $T$ into the past. Neuronal systems may also present UPOs, which in the case of many neurological disorders would be manifested as pathological synchronous oscillations. Multiple *in silico* studies showed that DFC can be used to disrupt these orbits (instead of fixing them) by feeding back the difference between the ongoing oscillation and the previous oscillation, typically delayed by $T/2$, where $T$ is natural oscillation periodicity. Multiple variations of DFC have been proposed for treatment of PD, tremors and epilepsy (Rosenblum and Pikovsky, 2004; Hauptmann et al., 2005; Popovych et al., 2005, 2006; Rosenblum, 2006; Tukhlina et al., 2007; Batista et al., 2010; Popovych et al., 2017a,b; Daneshzand et al., 2018b,a; Popovych and Tass, 2019; Zhou et al., 2020; Liu et al., 2023). Unlike the chaos control methods presented above for the *in vitro* experiments, DFC is not trigger-based, since it operates under continuous control, based both on the Magnitude and Phase of the system's oscillation.

**In this thesis...**

Now that we have laid out the conceptual map of control algorithms used for neuronal control, it is easier to understand where our contributions expand the current state of the art. As portrayed in Figure 2.5, the publication from Chapter 3 merges DFC, a Chaos Control Method, with Adaptive Control. We tested our novel algorithm termed Adaptive Delayed Feedback control *in vitro* and *in silico*. In Chapter 4, we present an algorithm for retinal stimulation in visual prosthesis that combines Model-based Control with AI. Following an encoder-decoder architecture, we used an ANN that first learns the forward model of the retina response (decoder), which is then used to train another ANN that learns the ideal stimulation patterns for visual reconstruction (encoder). We validate our algorithm *in silico* using a realistic computational model of retinal response to multipolar stimulation. Finally, in Chapter 5, we implement an ON-OFF controller that, instead of running a conventional pre-programmed algorithm, performs the control computation in a physical substrate, in a neuromorphic fashion, ideal for future low-power small-scale neuromodulation devices. We showed that such a neuromorphic controller can selectively modulate neuronal activity *in vitro*.

## 2.3 *In Vitro* electrophysiology with MEAs

As we saw in the previous section, the closed-loop algorithms tested *in vivo* are very simple, mostly limited to ON-OFF, Phase-locking and Proportional Control. As for the remaining control methods, most were developed *in silico*, with only a few exceptions being tested *in vitro*. When applicable, the step-up from *in silico* to *in vitro* experiments is an important milestone that should not be overlooked if one intends to move to more sophisticated control algorithms. *In silico* tests are based on simplistic representations of neuronal systems and many of the possible shortcomings of complex algorithms may only be evidenced after testing with real neuronal environments. *In vitro* platforms provide this "real neuronal environment" under well-controlled experimental conditions.

In this section, we will discuss how MEA technology and *in vitro* neuronal models can be used to expand the current state of the art on closed-loop algorithms for neuronal control. We will begin by discussing how the extracellular signals acquired with extracellular electrodes, such as those present in MEAs, differ from the intracellular signals, such as those captured with intracellular recording techniques such as patch-clamp. We will then detail how MEAs can be used to "read", "write" and "respond" to neuronal activity. Finally, we will conclude this section by briefly reviewing biological preparations that can be used for closed-loop experiments *in vitro*.

### 2.3.1 Intracellular VS extracellular recordings

Intracellular recordings played a pivotal role in our understanding of neuronal physiology. The seminal work of Hodgkin and Huxley – arguably the most influential collaboration in electrophysiology – was centred on intracellular recordings, a method they developed to measure the potential difference across the membrane of the squid giant axon (Hodgkin and Huxley, 1939). Nowadays, intracellular recordings are typically performed by either perforating the cell with a sharp electrode or by sealing the cell membrane with a micropipette. The latter is called patch-clamping and enables the recording of intracellular current/voltage flow through a small set or even individual ion channels (Figure 2.6.A, red).

In contrast with intracellular recordings, which establish direct contact with the neuron, recordings performed with extracellular electrodes are made without ever touching the cell. There is a cleft between the neuron and the electrode filled with ionic solution which dissipates the current generated by the neuron, as shown in the equivalent circuit in Figure 2.6.A, in black. For planar electrodes, this cleft has a thickness of around 40-100 nm, which leads to a resistance, $R_{seal}$, in the $M\Omega$ order of magnitude, and extracellular potentials of tens to few hundred microvolts (Braun and Fromherz, 1997; Gleixner and Fromherz, 2006; Wrobel et al., 2007). Some have focused on developing 3D electrodes with protrusions such as nanopillars (Brüggemann et al., 2011; Robinson et al., 2012) or mushroom-shaped electrodes (Hai et al., 2010; Mateus et al., 2019) that can be engulfed by the neuronal membrane, reducing the thickness of the cleft and, therefore, increasing the signal-to-noise ratio.

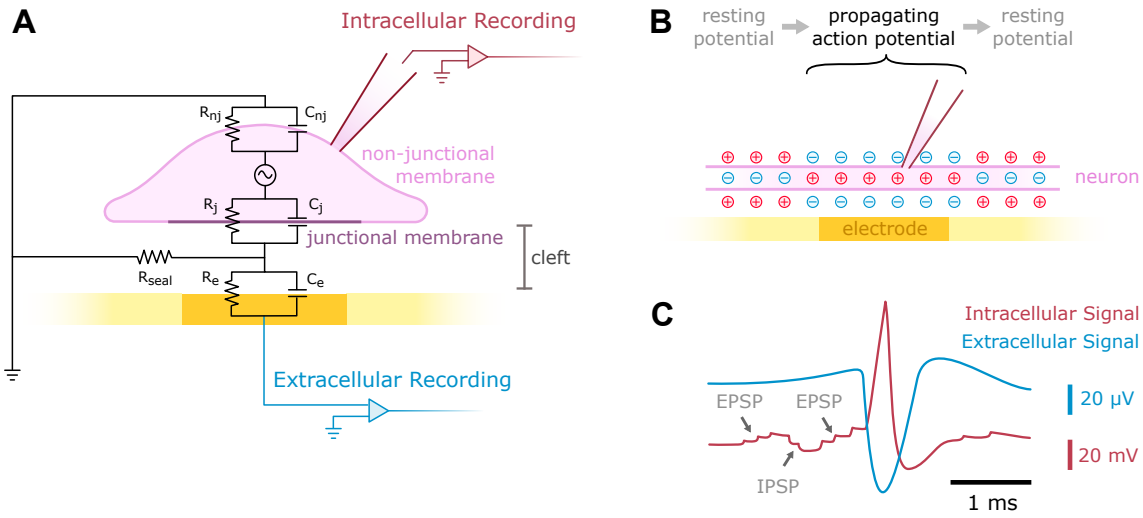



Figure 2.6: **Intracellular and extracellular neuronal recordings**. A) Spatial relationships between an extracellular electrode (yellow, with blue circuitry), intracellular needle (red), and the probed neuron (pink); and the associated equivalent electrical circuit (black) established by the electrode, neuron and culture medium (scheme inspired by Figure 1 from Spira and Hai (2013)). The electrode impedance is represented by the electrode resistance, $R_e$ and capacitance, $C_e$; the medium resistance at the cleft is represented by $R_{seal}$; the neuronal membrane is divided into two portions: the section that faces the electrode is referred to as junctional membrane (purple) as modelled by $R_j$ and $C_j$, and the section that faces the culture medium, referred to as non-junctional membrane (pink), and modelled by the $R_{nj}$ and $C_{nj}$. The scheme is not to scale. B) Action potential propagation across the neuron (pink), and the spatial spatial relations with the extracellular electrode (yellow) and intracellular needle (red). The membrane polarity of the section undergoing an action potential is inverted compared to the rest of the membrane, being relatively more positive intracellularly. C) Schematic representation of the same action potential recorded intracellularly and intracellularly. The intracellular signal shows sub-threshold potentials (gray arrows) such as excitatory and inhibitory signals received from presynaptic cells, referred to as excitatory presynaptic potentials (EPSPs) and inhibitory presynaptic potentials (IPSPs), respectively. The action potential has a positive peak since it captures the positive shift in polarity at the membrane. The amplitude of the recorded action potential corresponds to the actual magnitude of the physiological signal, with a rise of around 100mV. The extracellular signal, on the other hand, captures mostly the supra-threshold activity, i.e., the action potentials (being dependent on the electrodes input resistance and the dipole arrangement of the cells). Also, the distance between the electrode and the neuron leads to an attenuation of the spike's amplitude by a factor of about 1/1000. The recorded spike has a negative peak since it captures the transient potential drop of the extracellular medium

The signals recorded intra and extracellularly are remarkably different. To better understand this difference, we must look at what is actually being captured with each recording modality. Just like other animal cells, the interior of neurons is negatively charged compared with the extracellular environment. However, this charge balance is shifted during the generation and propagation of an action potential (Figure 2.6.B) due to the ionic flow through specific ion channels. As such, in the portion of the axon undergoing an action potential, the intracellular medium is temporarily positive compared with the extracellular medium. Intracellular recordings capture this rise of intracellular potential or the charging of the membrane (depending on the recording modality), as shown in Figure 2.6.C, in red.

Extracellular electrodes, on the other hand, capture the extracellular signal. As shown in Figure 2.6.B, during an action potential, the charge balance at the outer surface of the axon is transiently negative. The electrode records this drop in extracellular potential, leading to a negative spike in the recorded signal (Figure 2.6, in blue).

Another important distinction between intracellular and extracellular signals is the richness of the information captured by each. Intracellular recordings capture the "full picture" of internal electrical currents at the perfurated/patched region, with virtually no attenuation between the generated and the recorded signal. Thus, it is possible to measure the small excitatory and inhibitory signals received by neighbouring cells, referred to as excitatory presynaptic potentials (EPSPs) and inhibitory presynaptic potentials (IPSPs), which can be recorded at the soma (Figure 2.6.C, red). The action potential is triggered when the excitatory input leads to an increase in membrane potential above a certain threshold (around -50 mV). Extracellular recording can only capture the extracellular manifestation of action potentials (Figure 2.6, in blue), as all the sub-threshold events are invisible intracellularly (an exception would be the recording of *field* EPSPs, which can also be captured with extracellularly). Also, the magnitude of the signal is orders of magnitude smaller due to the strong attenuation caused by the cell-electrode cleft.

### 2.3.2 Reading, writing and responding to neurons with MEAs

As we have just seen, intracellular recording techniques grant researchers the ability to observe and analyse the details of neuronal activity with great precision, enabling a detailed examination of their physiological processes. However, the invasiveness of patch clampling allows for only acute single-cell recordings. It is also a time-consuming technique that requires highly trained electrophysiologists and an extensive experimental setup.

In the context of neuronal control, patch clamping enables rigorous control of a single neuron's activity, which explains its use in many closed-loop control experiments (Schiff et al., 1994; Miranda-Domínguez et al., 2010; Nabi et al., 2013; Panuccio et al., 2013; Couto, 2015). However, intracellular recording and stimulation do not mimic the modulation capabilities achievable with implantable devices, which operate extracellularly. Also, control algorithms that use patch clamp to interact with the neuronal system are only concerned about the temporal control, and

have no consideration for spacial control, since they operate at a single recording/stimulation site. However, problems such as the choice of effective electrodes from a pool of candidate stimulation electrodes may become increasingly relevant as future implantable devices are expected to offer many stimulation electrodes.

Microelectrode arrays enable non-invasive recording and stimulation of neuronal cultures extracellularly at dozens, hundreds or thousands of sites, with a simple experimental setup. As we will see throughout this section, MEAs open a window into the mesoscale dynamics of neuronal systems by enabling non-invasive long-term recordings (up to months) of entire cell populations with micrometre and millisecond resolution. If coupled with a stimulator, MEAs can be used to "write" information into the network, enabling neuromodulation at the mesoscale as well. If a processor is included into the setup for real-time analyses of the recorded signals, stimulation can be driven as a response to the ongoing neural activity, enabling closed-loop experiments. Throughout this section, we will expand on MEA technology and the associated equipment used for reading, writing and responding to neuronal activity.

**Reading with MEAs**

The electrical activity of the neuronal population can be read extracellularly using MEAs. A "vanilla" MEA consists of a grid of electrodes with diameters ranging from 10 to 30 $\mu$m with integrated wiring for each electrode (Figure 2.7.A top). The opposite end of the wire is connected to amplifiers and analog-to-digital converters (Figure 2.7.B). The electrodes are typically embedded at the bottom of a glass substrate, surrounded by a well, forming a chamber for the neuronal culture to grow on. The substrate can also be made of malleable materials such as polymers, which are ideal for *in vivo* applications. The spatial resolution can vary depending on the number of electrodes (typically 60 to 256) and the spacing between them (usually from 100 to 500 $\mu$m). MEAs can have a single well or multiple independent culture wells (ranging from 6 to 96, depending on the device), each with an associated electrode grid. Multi-well MEAs enable high-throughput experiments, although at the cost of having fewer electrodes for each well.

Having dozens or even a few hundred electrodes gives a broad view of the overall network activity. Still, it is important to be aware that the discrepancy between the size of neurons (around 10 $\mu$m) and the spacing between electrodes leads to an under-sampling of the network by several orders of magnitude. Solving this would require a much higher electrode density. However, the density of electrodes in passive MEAs is spatially constrained by the embedded electrode wiring. A significant increase in electrode count requires a different approach to MEA technology and new fabrication techniques.

High-density MEAs (HD MEAs) were introduced in the early 2000's. These can have thousands of electrodes separated by less than 20 $\mu$m (Figure 2.7.A bottom) achieving single or sub-cell resolution. For that, they require specialised hardware such as signal multiplexing due to spacial constraints of electrode wiring and on-chip signal amplification to minimise signal-conduction

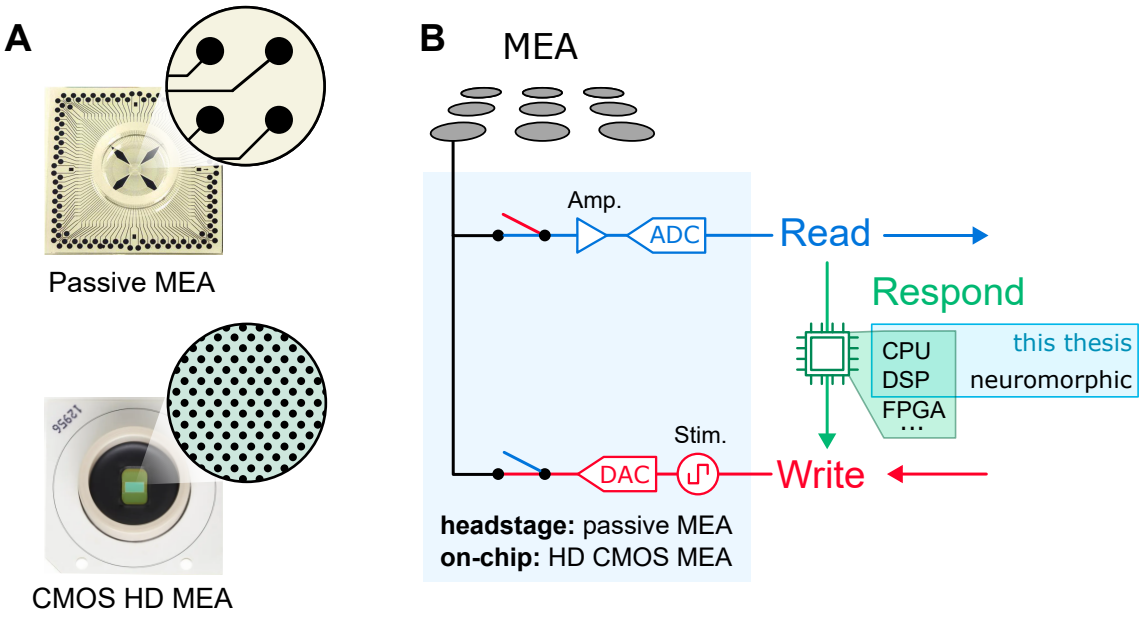



Figure 2.7: **MEAs for *in vitro* electrophysiology.** A) There are two main classes of MEAs, passive MEAs (top) which typically have 60 to 256 electrodes, and CMOS HD MEAs (bottom), which can have from 4000 to more than 26000 electrodes. The passive MEA image corresponds to a 120 electrode MEA from MCS, taken from the MCS User Manual. The CMOS HD MEA is the MaxOne MEA from MaxWell Biosystems and the image was taken from their website. B) MEA systems enable "reading" and "writing" into neuronal populations. Passive MEAs have all signal conditioning circuitry located in an external headstage (hence the MEA itself is passive). CMOS HD MEAs, on the other hand, feature on-chip circuitry for signal conditioning and channel multiplexing, enabling high-density electrode arrays. Most MEA systems that enable stimulation use the same electrodes for recording and stimulating. To avoid strong stimulus artefacts, the reading channel (blue wiring) is opened while the stimulation is active (red wiring). Besides "reading" and "writing", MEA systems can also "respond" to neuronal activity. For that, it must integrate a processor that has access to the neuronal stream in real-time and the ability to control the stimulators. The processor can be the CPU of the computer used for data storage or, for experiments requiring low closed-loop latencies, a microprocessor such as a DSP or programmable logic device such as an FPGA, with direct access to the datastreams and stimulators. In this thesis, werespond to neuronal activity using the CPU in Chatper 3 and the DSP and a neuromorphic element in Chapter 5, as shown in the blue square

noise and external interferences (Figure 2.7.B). This is made possible with complementary metal oxide semiconductor (CMOS) technology. The spatial resolution of CMOS HD MEAs is such that they enable tracking an action potential along a single axon extracellularly (Bakkum et al., 2013). However, recordings cannot be performed with all electrodes at the same time as the number of recording channels is typically smaller than the number of electrodes (the exact number varies according to the manufacturer of the device). As a drawback, CMOS HD MEAs are rigid devices, which limits their application to *in vitro* experiments, and opaque, not allowing the use of optical microscopy techniques. Also, unsurprisingly, their fabrication is far more expensive.

Currently, the companies manufacturing the MEA systems mentioned above are:

- **Multi-Channel Systems (MCS, Reutlingen, Germany)** with systems for both passive MEAs, single and multi-well, and single-well CMOS HD MEAs. The passive MEAs can have 60, 120 or 256 electrodes, all available for recording and stimulation. This is arguably the most popular system used for *in vitro* research. The CMOS HD MEAs have 4225 electrodes for simultaneous recording and 1024 stimulating electrodes.
- **Alpha MED scientific (Osaka, Japan)**, with a system for their passive MEAs with 64 electrodes, all available for recording and stimulation.
- **Axion Biosystems (Atlanta, USA)**, focused on multi-well MEA systems, with up to 768 passive electrodes for recording and stimulation, distributed across 6, 12, 24, 48, or 96 wells. The MEAs are inserted in an incubator assuring a well-controlled environment for the cell cultures in terms of temperature, humidity and $CO_2$ concentration.
- **3Brain (Pfaffikon, Switzerland)**, with single-well CMOS systems of 4096 simultaneous recording electrodes and up to 4096 on-chip stimulation sites, and 6-well systems of up to 2304 electrodes per well. The data acquisition system also features an incubator for the MEAs.
- **Maxwell Biosystems (Zurich, Switzerland)**, with single and multi-well CMOS HD MEA systems with 26400 electrodes, although with only 1020 recording channels. The data acquisition system is embedded within an incubator as well.

Another relevant consideration regarding MEA design is the material of the electrodes, as it can have a strong influence on the quality of the recordings. The most common electrode materials are conductive metals such as gold (Au), platinum (Pt), titanium nitride (TiN), aluminium (Al), and alloys like iridium oxide (IrOx). These are typically planar electrodes, whose background noise is compromised by the small electrode area and the fact that planar metal surfaces do not promote good cell adhesion. The effective electrode area can be increased by modifying the surface with porous conductive materials such as Pt (Park et al., 2010; Boehler et al., 2015, 2020), Pt-black (Rui et al., 2011; Lee et al., 2016), carbon nanotubes (Wang et al., 2006; Shein et al., 2008; Keefer et al., 2008; David-Pur et al., 2013), conductive polymers like PEDOT and PEDOT:PSS (Asplund et al., 2008; Bolin et al., 2009; Ludwig et al., 2011; Venkatraman et al., 2011; Cellot et al., 2016), or gold nanostructures (Hai et al., 2010; Brüggemann et al., 2011; Mateus et al., 2019). Besides the beneficial area increase, these rougher surfaces also promote improved cell-electrode coupling by providing 3D micro-structures for the cells to adhere to and may confer improved biocompatibility. The reading of *in vitro* neuronal networks may be further refined by combining MEA technology with microfluidic platforms (Lopes et al., 2018; Mateus et al., 2021). Introduced 20 years ago, microfluidic structures can be used to guide axonal growth in culture and induce the compartimentalization of the somatodendritic and axonal portions (Taylor et al., 2002).

### Writing with MEAs

MEAs can also be used to "write" into the neuronal populations using electrical stimulation. This has been used to study the elicited responses in models such as brain slices (Borkholder et al., 1997; van Bergen et al., 2003; Panuccio et al., 2018), dissociated neuronal cultures (Jimbo et al., 2000; Wu et al., 2020; Callegari et al., 2022), retina explants (Grumet et al., 2000; Ahuja et al., 2008; Ghaffari et al., 2021), among others.

The monitoring electrodes can be turned into stimulating electrodes by connecting them to external stimulators. The elicited response is then recorded by the remaining monitoring electrodes. In most experiments, electrophysiologists are interested in recording from the stimulating electrode(s) as well, to capture the local response to stimulation. That requires the stimulator to be linked to the stimulating electrode during the stimulation period exclusively (Figure 2.7, red wiring), and connected to the amplifiers for the remainder of the time (Figure 2.7, blue wiring).

The stimulus waveform – pulse amplitude, width, polarity, etc. – must be carefully selected to ensure an effective neuronal response that avoids long-term damage to both the neurons and the electrodes. That must take into account the stimulation type (voltage or current), the electrode material, and the reliability of the elicited neuronal response. Wagenaar et al. (2004) showed that voltage was more effective than current stimulation, providing reliable responses for amplitudes above 600 mV using biphasic pulses with 400 $\mu$s per phase (Wagenaar et al., 2004). However, we must note that the biphasic pulses used for voltage stimulation are not advised given the type of electrodes used (Titanium nitride, TiN). According to the manufacturer (MCS), when using TiN electrodes, one should ensure that no positive charges accumulate in the electrode's surface as these lead to electrolysis and consequent electrode damage. This means that voltage stimulation should rely on monophasic negative pulses only, whereas current stimulation should use biphasic pulses with the negative phase first. The resistance to electrolysis is highly dependent on the material of the electrode, and thus the specificities of the waveforms must be confirmed according to the suggestions of the manufacturer regarding pulse polarity, amplitude and duration.

There is a technical handicap all electrophysiologists face when performing experiments with electrical stimulation – the presence of stimulation artefacts. Indeed, whenever a stimulus is triggered, the monitoring electrodes (which may include the stimulation electrode, once the stimulus is over) capture an artefact orders of magnitude larger than the neuronal spikes. This artefact is generated by the involved internal circuitry and the propagation of the pulse through the conductive culture medium. Multiple strategies can be used to mitigate these artefacts. The most common intervention is right at the recording stage as most devices disconnect the amplifiers from the monitoring electrodes for a few milliseconds once stimulation is triggered (Figure 2.7.B, red wiring). Also, the high-pass filtering used to capture neuronal activity removes most of the artefact, which has a strong low-frequency component. Other approaches typically involve software and rely on some sort of template matching of the artefact followed by its subtraction from the recorded time series,

which can be performed offline or online (Wagenaar and Potter, 2002; Mena et al., 2017; O'Shea and Shenoy, 2018).

### Responding with MEAs

MEA data acquisition systems are designed for the recording of neuronal activity and most have incorporated stimulators enabling both reading and writing into the neuronal cultures. In conventional experimental setups, the recording and stimulation components are uncoupled – the stimulus is set *a priori* and is thus independent of the recorded activity. It is undeniable that these systems have played an important role in unravelling neuronal function. However, such an open-loop interaction is not ideal for studying systems as complex as neuronal networks. An *in vitro* platform that enables bi-directional interaction with the neuronal cultures through responsive stimulation is a desirable tool for three types of applications in Neuroscience research (Newman et al., 2013), as we pointed out in the *in vitro* section of Figure 2.4:

- **Clinical neurostimulation**: closed-loop stimulation therapies are expected to provide better outcomes than the currently used open-loop strategies. *In vitro* setups provide a simplified and highly controlled environment to develop and validate novel control algorithms, as reviewed in detail in section 2.2.3. This is the research area addressed throughout this thesis.

- **Control of network variables**: studying the rules that govern neuronal response to stimulation is a challenging task because the input-output relations will vary as a function of the high-dimensional history-dependent network state. Such rules can be better disentangled if one adjusts the stimulus in accordance to the ongoing observable network state, however it is defined. *In vitro* closed-loop stimulation is thus a powerful methodology to study the mechanisms behind the input-output function of neuronal networks (Wallach et al., 2011; Weihberger et al., 2013).

- **Embodiment of neuronal cultures**: the main function of a brain is processing information from the external world and coordinating motor actions accordingly. *In vitro* neuronal models are, metaphorically speaking, unembodied brains. Still, these models can be very helpful to study how neuronal ensembles self-organise for information processing by connecting them to an external agent, be it a simulated body or a robot. That requires a closed-loop interaction where information about the network activity is converted into the agent's actions and sensory information about the agent's environment is encoded into the network with electrical stimulation (Reger et al., 2000; DeMarse et al., 2001; Bakkum et al., 2004; DeMarse and Dockendorf, 2005; Kagan et al., 2022).

To answer these research questions, besides recording and stimulating capabilities, the involved hardware needs to have another crucial component: a processor for analysing the recorded activity and controlling stimulators in real-time (Figure 2.7, green).

The *in vitro* studies performed within the three aforementioned closed-loop domains used the CPU of a PC as the processing unit. However, a more efficient alternative to a PC is using dedicated real-time hardware such as a Digital Signal Processor (DSP). These microprocessors are highly efficient and designed specifically for real-time digital signal processing such as filtering, modulation, and analysis. DSPs perform high-speed computations for very specific tasks and can be brought close to the acquisition and stimulation equipment in the hardware pipeline, enabling minimal closed-loop latency. Currently, out of all the commercial MEA systems listed above, only one incorporates such built-in capabilities: the MEA2100-System from MCS, which includes the TMS320C6454 DSP from Texas Instruments, running at 819.2 MHz.

The DPS from the MEA2100-System has full access to the data stream from the electrodes and can be used to analyse neuronal data and activate the stimulators in real-time. The software provided by MCS, the Multi Channel Experimenter, allows the user to easily define simple ON-OFF closed-loop protocols, triggered by spike detection or the level of instantaneous firing rate ($N$ spikes detected in a defined time window), in a pre-selected set of electrodes. When the threshold is crossed, the predefined stimulus is sent to the pre-selected electrodes. For more sophisticated closed-loop control algorithms, the DSP is "opened" to be programmed by the user.

Despite the versatility provided by this MEA system, it has not yet been used to answer relevant research questions – its use has been limited to simple proof-of-concept studies. George et al. (2015) used the DSP functionality in the context of neuronal cultures embodiment. The robot's movement was a function of the network activity, and the neuronal stimulation was a function of the robot's sensors. The DSP of the MEA2100-System was thus programmed for real-time processing of the neuronal signals and controlling the stimulation. As a side note, despite also dealing with the recording and stimulation of neuronal cultures, the studies on the embodiment of neuronal cultures are out of the scope of this thesis because these do not aim to control the neuronal system, but instead exploit the neuronal activity to control an external agent – the locus of control is not an external algorithm, but the network itself , as already explaining in section 2.2.3 and in Figure 2.4.

Tanskanen et al. (2019) recently used the DSP from MEA2100-System in a proof of concept study to showcase the potential of this system for closed-loop stimulation. In their experiments, they stimulated a pre-selected electrode when 3 spikes were detected within a 500 ms window in another pre-selected electrode. This study did not aim to answer any research question, they just wanted to prove the feasibility of closed-loop experiments with the MEA2100-System.

Despite its potential, the community has not yet adopted this device as a standard tool for *in vitro* closed-loop experiments. That could be due to one of the three following facts: first and foremost, this acquisition system with an integrated DSP is a recent player in the market and closed-loop experiments started more than three decades ago so, naturally, most literature relied on different solutions; secondly, the default closed-loop functionalities available in the associated acquisition software are limited to variations of ON-OFF control; finally, the development of personalised

control algorithms requires programming the DSP directly, which is not a trivial task for an electrophysiologist without a strong engineering background (and the documentation provided by the manufacturer, besides being quite limited, is not presented in a tutorial-like fashion).

Given the interest in closed-loop solutions for MEA experiments, some research groups have focused on developing their MEA setups with integrated responsive capabilities. A popular approach is using FPGAs to handle real-time processing of the large quantities of data acquired with MEAs. In contrast to microprocessors such as CPUs and DSPs, where the functionality is determined by fixed hardware architectures and controlled by software instructions, in FPGAs the computation is integrated into the hardware itself. An FGPA is basically a matrix of logic gates that can be reconfigured by the user to construct custom digital circuits that perform the intended task – its hardware is programmable. These devices offer extensive application flexibility, are inherently parallel devices (which is highly convenient for processing MEA data), and enable significantly faster and more energy-efficient computations than CPUs.

Hafizovic et al. (2007) integrated an FPGA into a CMOS MEA system with 128 electrodes for acquisition and stimulation. The FPGA was in charge of data acquisition and spike detection in real-time, allowing the PC to trigger stimulation in less than 2 ms. The system developed by Müller et al. (2013) used the FPGA not only for spike detection but also for stimuli generation in a CMOS MEA with 11,011 electrodes. It was able to simultaneously process data from 126 channels and generate stimulus on 42 output channels with a latency below 1 ms. Park et al. (2017) used an FPGA for both spike detection and spike sorting in 128 acquisition channels simultaneously. As a side note, spike sorting consists of classifying spikes detected in the same electrode as belonging to different neurons, by comparing their shape. The FPGA was also in charge of triggering closed-loop stimulation in 8 stimulation channels but the closed-loop latency was not disclosed. (Seu et al., 2018) developed an FPGA-based system for a CMOS MEA with 4096 recording channels and 16 stimulation electrodes with a closed-loop latency below 2 ms.

These advances in closed-loop setups for MEA experiments reflect the interests of the community and showcase the possibilities available for electrophysiologists. However, as the proposed systems and associated hardware are custom-made by the specific laboratories, it is very optimistic to assume that these solutions will be replicated and adopted by other neurobiology groups that do not share the same electrical engineering expertise.

At the end of the day, all the *in vitro* studies with MEAs mentioned in section 2.2.3 used a PC as a processing unit to close the control loop. It is the easiest solution to implement and does not require any specialised hardware besides the acquisition and stimulation equipment already incorporated in most MEA systems. This strategy comes with a cost though: the closed-loop latency is larger (several milliseconds) and less consistent due to jitter. Still, latencies of a few milliseconds are not critical if the goal is to control population dynamics, as these do not change on the millisecond scale.

Some research groups made their software freely available for others to use. Steve Potter's lab was at the forefront of using MEAs to analyse and interact with the neuronal networks *in vitro* and shared two of the platforms they developed: MeaBench (Wagenaar et al., 2005a) and NeuroRigther (Rolston, 2009). MeaBench is an open-source Linux application written in C++ that mostly focuses on online MEA signal analysis and visualisation. It includes functions such as filtering, suppression of stimulus artefacts, spike detection, etc. With MeaBech it is also possible to activate external stimulators enabling closed-loop experiments. This software is compatible with the MEA equipment from MCS.

NeuroRighter advances on the previous contributions from the group, such as MeaBench and other software applications for *in vitro* experiments (Wagenaar and Potter, 2002; Wagenaar et al., 2004), by extending its applicability to freely moving animals and focusing on user-friendly interfaces. In line with the primary objective of facilitating open-source use and adoption across various laboratories, NeuroRighter was developed using C#, a Microsoft-supported, type-safe, object-oriented language, which offers a more straightforward approach to desktop application development compared to C/C++. The software was designed to allow the user to implement custom closed-loop algorithms. The closed-loop latency was tested for different hardware setups, ranging from around 7 ms to almost 50 ms, depending on the specific configuration.

The Closed Loop Experiment Manager (CLEM) is yet another example of an open-source Windows desktop application developed for closed-loop experiments. Similarly to NeuroRighter, CLEM has a user-friendly user interface and allows the user to develop custom closed-loop algorithms. It was developed in C and C++, having a closed-loop latency below 2 ms.

These three platforms were developed for the 60-electrodes MEAs from MCS and the associated acquisition hardware (even though NeuroRighter also supports an *in vivo* setup). We note, however, that these applications have not been adopted by the community at large. Most citations refer to these studies as examples of available applications to use in experiments, and not as reference to its use in the study in question. What could explain that? Perhaps it may be easier for a lab to implement its own software, specifically adapted to the available hardware, than to acquire/develop all the equipment required to run someone else's program which, without disregarding the immense value in providing free open-source tools, can be a risky investment.

### 2.3.3 *In vitro* preparations

In the last section, we discussed the relevant equipment to perform closed-loop experiments *in vitro* with MEAs. The other crucial piece is, of course, the *in vitro* biological preparation. The three most common types of *in vitro* neuronal models are acute slices, organotypic cultures and dissociated neuronal cultures. These vary across multiple dimensions, such as in terms of resemblance to the *in vivo* tissue, ease of preparation, longevity, and type of extracellular signals recorded. We will briefly describe these neuronal models and how they can be used for electrophysiological experiments.

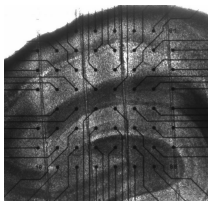
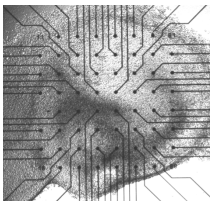
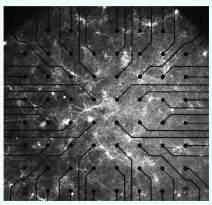
| | | | In this thesis |
|---|---|---|---|
| | **Acute Slice** | **Organotypic Culture** | **Dissociated Culture** |
| |  |  |  |
| **Thickness** | Thick (~300-500 µm) | Thin (~100-300 µm) | 2D |
| **Extracellular signals** | LFPs | LFPs | Extracellular spikes |
| **Original Circuitry** | Preserved | Preserved | Lost |
| **Viability** | Short term (hours) | Long term (weeks) | Long term (weeks) |
| **Preparation** | Moderate | Hard | Easy |

Figure 2.8: The most common biological preparatoins are acute slices, organotypic cultures and dissociated neuronal cultures. These differ in many relevant aspects: acute slices are typically thicker than organotypic slices whereas dissociated neuronal cultures consist of a single layer of cells; the signals acptured by brain slices are LFPs due to the increased cell density, while dissociated cultures capture extracellular spikes of individual cells; slices preserve the original brain circuitry whereas dissociated cultures do not; acute brain slices can only be used for a couple of hours, whereas the neuronal cultures, be it organotypic or dissociated, can be maintained for several weeks for long-term experiments; finally, dissociated neuronal cultures are the easiest to prepare, while organotypic cultures require significant experimental expertise. The slices images were taken from the MCS manuals for the preparation of acute slices and organotypic cultures, the dissociated culture was taken from Dranias et al. (2013)

### Acute Slices

The most commonly used type of brain slices used in electrophysiology research are acute brain slices (Figure 2.8 left). These slices have a thickness of around 300 to 500 $\mu$m, typically from the hippocampus or cortex, and are placed on a saline medium with carbogen (5% O2 and 95% CO2), shortly after euthanizing the animal. These are short-lived preparations, not cultures, and their electrophysiological activity can be recorded both intracellularly with path-clamp or extracellularly with MEAs.

The main advantage of acute brain slices is that they preserve the local cytoarchitecture of the structures under study. The circuits of interest can be easily accessed and manipulated enabling well-controlled experimental studies while removing all the confounding variables that would naturally occur in a living animal (motion artifacts for example).

Brain slices still preserve many of the activity biomarkers of *in vivo* brains. They have been widely used to study a myriad of mechanisms behind epileptic seizures such as seizure onset (Pal et al., 2001; Librizzi et al., 2017), propagation (Weissinger et al., 2005; Losi et al., 2016) and termination, as well as strategies for seizure control, both using drugs (Hongo et al., 2015), electrical stimulation (Schiff et al., 1994; Gluckman et al., 2001; Khosravani et al., 2003; Guez et al., 2008; Panuccio et al., 2013; Matchen and Moehlis, 2018), light or even temperature (Motamedi et al., 2006). Experimental protocols have been developed to induce seizure-like activity in brain slices (Schwartzkroin and Prince, 1978; Rutecki et al., 1987; Harrison et al., 2004).

MEA recordings of brain slices capture LFPs similar to those recorded *in vivo*. However, effective adhesion of the slice to MEAs may be technically difficult to achieve, which compromises MEA experiments. Protocols for MEA experiments are available detailing the experimental details required to achieve prolonged high-quality experiments(Panuccio et al., 2018). Probing these slices intracellularly is more robust then intracellularly due to a better tissue-electrode adhesion. As such, it is not surprising that patch clamping has been the technique of choice for closed-loop stimulation experiments performed with acute brain slices (Schiff et al., 1994; Gluckman et al., 2001; Khosravani et al., 2003; Guez et al., 2008; Panuccio et al., 2013; Matchen and Moehlis, 2018). These used hippocampal slices as models of epilepsy and used different types of algorithms to control seizure activity.

The main disadvantage of brain slices is that they are only viable for a few hours. Their thickness and lack of circulatory system lead to the rapid death of cells within the slice. Also, the lesion provoked to extract the slice may activate cellular stress responses that could compromise the experiment, depending of the research question at hand.

### Organotypic cultures

Another type of brain slice preparation is organotypic cultures (Figure 2.8 center). Unlike the short-lived acute preparations, organotypic slices are cultures that can be maintained for several

weeks. These also retain the cytoarchitecture of the *in vivo* tissue and can be used to study the maturation and regeneration of neuronal circuitry in long-term experiments (Li et al., 1996; Hofmann et al., 2004), .

Organotypic cultures are typically thinner than acute brain slices, which facilitates the perfusion of gases and nutrients to the inner cells or the slice. As such, organotypic slices can be cultured on MEAs for several weeks enabling long-term studies, which are not feasible with acute slices. Their extracellular activity also consists of LFPs.

The main limitation of organotypic slices is their technically challenging experimental preparation. The slices are very thin making them extremely fragile and difficult to manipulate which compromises the adhesion to the MEA. Also, these are easier to prepare and maintain when obtained from a young animal, since the tissue is physiologically more resistant and adaptable. Slices from older brains may prove more challenging to maintain in culture, Limiting the applications of this type of *in vitro* model. To this day, we have not yet seen this type of neuronal model being used for closed-loop electrical stimulation experiments. However, we do think it is a model worth exploring for algorithms that require long-term experiments, such as intelligent data-driven methods like Reinforcement Learning.

#### Dissociated Neuronal Cultures

Compared with acute brain slices, dissociated neuronal cultures are much simpler to prepare and maintain (Figure 2.8 rigth). The cells are extracted from embryonic or newborn animals, typically rats or mice. Using young cells is very important for this type of model because the cells retain the ability to go through the process of network formation.

The most commonly used cell types are also the cortical and hippocampal cells. The reason why these cultures are referred to as dissociated is because the target brain region is extracted and the tissue is enzymatically and mechanically dissociated into a cell suspension immersed in culture medium. The dissociated cells are then plated on the MEA. Over the course of several days, neurons start expanding their axons establishing synapses with neighbouring neurons. The initially dissociated neurons form a functional 2D network with a brand new connectivity.

Just like in developing brains, the activity of dissociated neuronal networks evolves considerably at its early stages. During the first days of development, the culture goes through a maturation process that involves a drastic shift in the synaptic function. Neurons can be either excitatory or inhibitory depending on the type of neurotransmitter released at their synapses. However, in immature brains and cultures, the most common inhibitory neurons – GABAergic neurons – have an excitatory effect (Ben-Ari, 2002). Developing neurons have high intracellular $Cl^-$; as the network matures and intracellular $Cl^-$ decreases, the effect of GABA transitions to inhibitory, marking the end of maturation. On the electrophysiological side, the network activity evolves from sporadic spiking of individual neurons to collective network events. A matured network is ready for the experiments and can be maintained for several weeks, even months.

The cytoarchitecture of a dissociated culture is different from that of the original brain tissue. However, the functionality of neurons is retained as well as their ability to interact with neighbouring cells (other neurons, glial cells, etc.), enabling the study of fundamental physiology such as neuronal communication, neuro-glia interaction, drug screening, network dynamics and their response to electrical stimulation.

This type of neuronal preparation can also be used to model specific features of disease. For example, hippocampal and cortical cultures exhibit spontaneous synchronous bursts that are remarkably similar to interictal spikes (le Feber, 2019). Interictal spikes are biomarkers of pathologic activity that occur in the epileptic brain during seizure-free periods (interictal period). Seizure-like activity can be also induced with additional manipulations that affect the excitation-inhibition (Furshpan and Potter, 1989). Dissociated neuronal networks are inherently hyperexcitable systems and can be used to study the effect of anticonvulsant drugs. For example, the specific excitatory receptors that must be blocked to effectively prevent network synchronization in hippocampal cultures are the same used in anticonvulsant therapy for epilepsy (Rogawski, 2011).

In conclusion, there are many possibilities of neuronal preparations – even besides acute brain slices, organotypic cultures and dissociated cultures – and these can be extracted from different animal sources – typically rat, but now human induced pluripotent stem cells are a possibility that allows studying human neuronal cells extracted from non-neural cells. Also, there are multiple electrophysiological recording modalities available, such as patch clamping, passive MEAs, high-density MEAs, etc. The choice of neuronal model and recording setup must be carefully selected according to the research question at hand.

In conclusion, there are various types of neuronal preparations available for research purposes, even besides acute brain slices, organotypic cultures, and dissociated cultures (such as organoids and steroids, which are out of the scope of this thesis, because, so far, they are not suitable for closed-loop experiments). These can be obtained from different animal sources (typically rats), however, with the advent of human induced pluripotent stem cells, it is now possible to study cultures of human neuronal cells extracted from non-neural cells. Also, there are several electrophysiological recording modalities to choose from, such as patch clamping, passive MEAs, high-density MEAs. It is crucial to carefully select the appropriate neuronal model and recording setup based on the research question at hand.

In this thesis...

As we have seen, there are many MEA systems available for *in vitro* electrophysiology. Throughout this thesis, we worked with multi-well passive MEAs with 256 TiN electrodes and the 2100MEA-System from MCS which, as mentioned before, includes not only all the acquisition and stimulation hardware but also a DSP, enabling closed-loop experiments. In Chapter 3 we developed our own C# application to close the control loop as the available open-source tools were not compatible with our hardware and the control time-scale did not justify using the DSP. The software had a closed-loop latency of around 20 ms and always below 40 ms, perfectly acceptable given the control task. In Chapter 5 we used the DSP included in the interface board of the 2100MEA-Sytem. The DSP was in charge of online filtering and spike detection enabling real-time interaction with the neuromorphic element, the memristor, which was the focus of that study. Regarding the type of neuronal preparations, in Chapters 3 and 5 we used dissociated hippocampal neuronal cultures extracted from rat embryos.

# Chapter 3

# Controlling unwanted neuronal oscillations with closed-loop stimulation

As we have seen in the previous chapter, most algorithms designed to control pathological neuronal activity focus on disrupting excessive synchrony and the associated firing oscillations, the reason being that these can be traced back to the severity of motor symptoms (Schnitzler et al., 2009; Levy et al., 2000). Having identified the control metric (the intensity of a pathological oscillation) and the control objective (disrupting the oscillation), one must ensure that the control algorithm conforms with the intrinsic dynamics of the system. But *what type of control algorithm can be used to disrupt pathological neuronal activity?* That was our first research question. Each algorithm poses a set of assumptions about the system it aims to control and, typically, the simpler the method, the stronger the assumptions. For example, when using a PID controller one assumes response linearity to external input – a strong assumption for complex systems like neuronal networks. Sophisticated methods, like the adaptive model-based methods, are more flexible and thus capable of adjusting themselves to the non-linear time-changing dynamics of neuronal systems. But the practical implementation of these algorithms in real experimental conditions is hindered by their complexity, so they were only tested in computer simulations.

We believe that simple algorithms capable of controlling complex behaviour are the best candidates for the next generation of closed-loop devices. Chaos Control techniques pose an interesting alternative. The assumption here is that the system presents chaotic dynamics which, once identified, can be stabilised/destabilised with small external perturbations. Even though the identification of chaos in neuronal activity is somewhat controversial, Christini and Collins (1995) showed that chaos control techniques may also work for systems that only *appear* chaotic.

In this chapter, we implemented a popular control technique which was first designed to control

chaotic systems – Delayed Feedback Control (DFC). We used an *in vitro* model that showed *apparently* chaotic UPOs in the form of network oscillations (these may also result from quasi-periodic dynamics, an intermediate state between periodicity and chaos). This work resulted in a first-author publication titled *"Disrupting abnormal neuronal oscillations with adaptive delayed feedback control"*, currently under review. This work counted with support of the following co-authors: Miguel Aroso, from the NCN Lab at i3S; A. Pedro Aguiar from FEUP; David B. Grayden, from the Department of Biomedical Engineering, University of Melbourne, Australia; and Paulo Aguiar, PI of NCN at i3S. This work resulted in a first-author publication in the journal eLife (de Castro et al., 2024): `https://doi.org/10.7554/eLife.89151`.

## 3.1 Introduction

The excessive oscillatory synchronization that emerges in neurological disorders, such as Parkinson's disease and epilepsy, can be disrupted using direct electrical stimulation. Despite its potential, conventional brain stimulation therapies still rely on continuous open-loop stimulation, leading to excessive stimulation and burdensome side-effects (Kuo et al., 2018). In theory, closed-loop stimulation guided by an efficient control algorithm should provide better outcomes, disrupting the pathological neuronal dynamics with minimal intervention. But, given the complexity of neuronal systems, it is very challenging to foresee what type of control algorithms would be effective. Multiple *in silico* studies focused on employing closed-loop control methods to desynchronize pathological oscillations in computational models of neurological disorders (Beverlin II and Netoff, 2013; Holt et al., 2016; Weerasinghe et al., 2021; Witt et al., 2013), but these have not been translated into studies with real neuronal systems to access its viability.

A particularly prevalent method in the field is Delayed Feedback Control (DFC), a technique from Chaos Theory developed to stabilize/destabilize chaotic systems using minimal perturbations (Pyragas, 1992; Rosenblum and Pikovsky, 2004). DFC controls the system by applying a feedback signal proportional to the difference between the current oscillation amplitude and the oscillation amplitude from a fixed delay into the past. DFC has been extensively applied in the context of *in silico* neuroscience, with studies focusing on specific disorders such as Parkinson's Disease (Rosenblum and Pikovsky, 2004; Vlachos et al., 2016; Popovych et al., 2017a; Daneshzand et al., 2018b; Popovych and Tass, 2019; Yu et al., 2021) and epilepsy (Zhou et al., 2020). However, there is still controversy in the literature regarding its efficacy and whether it may instead promote synchronization, depending on the baseline levels of synchrony of the modelled network (Dovzhenok et al., 2013). Thus, it is critical to test the effectiveness of DFC with biological neuronal networks that display synchronized oscillatory firing. Neuronal cultures on microelectrode arrays (MEAs) are a useful testbed for this task: these cultures self-organize into complex interconnected networks with quasiperiodic synchronized bursting, and MEAs enable closed-loop modulation of the neuronal activity, providing concurrent recording and stimulation at multiple points of the population with high spatial and temporal resolution (Obien et al., 2015).

In this work, we implemented, for the first time, the DFC algorithm to disrupt the synchronized oscillatory bursting of biological neuronal populations. The cells were cultured on MEAs and the network activity was monitored and controlled in real time with electrical stimulation. We show that traditional DFC lacks the adaptability to cope with the malleable dynamics of biological networks and may actually impose a new firing periodicity. As an alternative, we developed and implemented an adaptive DFC (aDFC) algorithm that self-tunes to better predict the current oscillation phase and periodicity. We demonstrate that aDFC reduces the synchronization of the network and efficiently disrupts the baseline oscillation. While doing so, we identified and characterized a new feature of biological neuronal networks: the intrinsic firing dynamics of a network unveils its propensity for neuromodulation. To better understand this relationship, we went back to the *in silico* neuronal networks and revealed that their controllability can be conditioned by specific parameters of the network architecture, such as the excitatory/inhibitory balance and overall synaptic weights of the network connections. Taking all together, we show that theoretical predictions regarding the efficacy of DFC might be misleading and we present an improved algorithm for closed-loop electrical stimulation. The knowledge gathered while testing these algorithms *in vitro* was relevant to recognise short-comings of *in silico* simulations and may guide the development of better computational neuronal models. Our results revealed fundamental properties and considerations regarding the controllability of neuronal systems.

## 3.2 Methods

### 3.2.1 Cell Cultures

The experiments followed both the European legislation regarding the use of animals for scientific purposes and the protocols approved by the ethical committee of i3S. The Animal Facility of i3S follows the FELASA guidelines and recommendations concerning laboratory animal welfare, complies with the European Guidelines (Directive 2010/63/EU) transposed to Portuguese legislation by Decreto-Lei no 113/2013 and is licensed by the Portuguese official veterinary department (DGAV, Ref 004461). Embryonic (E18) rat hippocampal neurons were cultured on 6-well or 9-well chamber MEA (256-6well MEA200/30iR-ITO-rcr and 256-9wellMEA300/30iR-ITO, respectively) (Multichannel System MCS, Germany) with a density of $5 \times 10^4$ cells/well and 3 x $10^4$ cells/well, respectively. Each well of the 6-well MEA has 42 TiN electrodes (array of $7 \times 6$), and the 9-well contain arrays of 26 TiN electrodes (6 x 5 without corners). The experiments were performed between 13 and 55 days in vitro.

### 3.2.2 Electrophysiological Recordings

The recordings were performed with the MEA2100-256 System (Multichannel System, MCS, Germany). The cell cultures were kept at 37 °C and 5% CO2 using a stage-top incubator (ibidi GmbH, Germany) adapted to the headstage of the MEA2100-256 system. The electrophysiology

signals were sampled at 10 kHz and filtered with a HP at 200 Hz (second order high pass butterworth filter). The neuronal spikes were detected and recorded using the Experimenter software from MCS. Peaks that were 6 STD (sometimes up to 8 STD, depending on the background noise) below or above baseline were considered neuronal spikes. There was a deadtime of 3 ms after each detection to prevent biphasic spikes from being identified twice. This spike data was used for posterior offline analysis.

### 3.2.3 Closed-loop Computer Application

The developed C# software was built using the McsUsbNet dynamic link library from MCS (https://github.com/multichannelsystems/McsUsbNet). The software ran in parallel with the Experimenter software in a quad-core Intel i7-6700 with 32.0 GB RAM workstation. The raw neuronal stream of data was received in batches at 10 kHz. Each batch had 1 sample point per electrode and was transferred directly from MEA2100-256 system via USB 3.0 high speed cable. The signal processing was analogous to that running on Experimenter: the raw data was high-pass filtered at 200 Hz and the spike detection was performed using the same algorithm (section 5.2). The C# application had access to the MEA2100-256 System stimulators via USB 3.0, allowing us to trigger stimulation by software. The app also commanded the beginning of the recordings in Experimenter. The overall closed-loop latency associated with the activation of the stimulus was below 40 ms. The application can be is freely available for download in GitHub (`https://github.com/NCN-Lab/aDFC`).

### 3.2.4 Delayed Feedback Control Algorithm

The C# application ran the DFC algorithm online. The application only monitored electrodes that had an average firing rate above of 0.1 Hz, measured before starting the experiments. The instantaneous firing rate ($R$) of an active electrode was calculated in a square moving window of 0.1 s (sometimes tuned to 0.2 or 0.5 before starting the experiments depending on the dynamics of the network). The $R$ signal corresponds to the number of spikes per active electrode per unit of time. The $R$ of the monitored network was filtered using a damped oscillator (Popovych et al., 2017a; Daneshzand et al., 2018b),

$$\ddot{x} + \omega\dot{x} + \omega^2 x = kR(t), \tag{3.1}$$

where $\omega$ denotes the natural frequency of the oscillator and $k$ is a scaling coefficient. Here, we considered $k = \omega$, which prevents the oscillator amplitude from scaling with the different ranges of periodicities exhibited by different cultures. We used $\dot{x}$ as the output of Equation 3.1. The stimulation frequency ($S$) is obtained by subtracting the current oscillation to the half-period shifted oscillation,

$$S(t) = K\left(\dot{x}\left(t - \frac{T}{2}\right) - \dot{x}(t)\right), \tag{3.2}$$

where $K$ represents the stimulation gain and $T$ is the oscillation period. A stimulus is triggered when

$$t > t_{laststim} + \frac{1}{S(t)}, \tag{3.3}$$

where t is the current time and $t_{laststim}$ is the time stamp of the last stimulus. Equation 3.3 was bounded so that the stimulus was only sent when $1 < S(t) < 20$ . The upper bound is required due to the latency in the stimulus activation (below 40 ms) and the lower bound guarantees that low background values of $S$ do not trigger stimulation (and that there is no negative stimulation frequency). Here, $T$ and $\omega$ are adapted based on the real-time detection of synchronous events. The synchronous events (network bursts, NB) were detected when the $R$ signal was above a certain threshold, typically 10 Hz, considering a minimum interval between NB of 0.1 s. The threshold and minimum NB interval were sometimes adjusted before starting the experiments to better suit the observed data.

### 3.2.5 Experimental Protocol

Before starting the experiments, we monitored the networks for up to 10 min to identify the active electrodes. The application calculates the average firing rate for all the electrodes and selects those firing above 0.1 Hz. During this preliminary period, we also tuned, the parameters for the detection of synchronous events, if necessary.

Each stimulus unit was a negative electrical pulse with 200 µs duration. The choice of stimulation electrode and pulse amplitude was based on the premise that the stimulus should promote a response in the neighbouring electrodes using minimal stimulation. We probed different electrodes manually using Experimenter software, starting with the most active ones. For each electrode, we tested different stimulus amplitudes from -400 mV to -1000 mV in steps of -100 mV. When a given electrode and stimulation amplitude evoked a consistent response in nearby neurons, we fixed the stimulus amplitude to 120% of that value to account for eventual adaptations of the network, which could render the stimulation ineffective.

The experiments started with a recording of 15 min of spontaneous activity. The choice of stimulation protocols – DFC, aDFC and Poisson – for the following trials was randomized (for the first four networks tested, the order of stimulation protocols was also arbitrary, but not explicitly random). Each trial had three periods: 5 mins without stimulation (OFF), 5 mins with stimulation (ON), and 5 mins without stimulation again (OFF). For the DFC and aDFC protocols, the stimulation followed the $S(t)$ signal determined by each algorithm. For the Poisson protocol, the

stimulation followed a Poisson distribution of stimuli with an average frequency equivalent to that used in the previous aDFC trial.

### 3.2.6 Selection criteria

We only performed experiments on networks that displayed consistent oscillatory bursting activity with periodicities between 0.5 and 5 seconds, approximately. Some cultures presented the required dynamics but were rejected because they were (or became) irresponsive to stimulation. In order to have statistical robustness in the evaluation of the different stimulation protocols, we only analysed networks that were submitted to at least three trials with each protocol. A total of 14 networks fulfilled the criteria to proceed with the experiments and analysis. Here, we define a network as being a neuronal culture in a given day in vitro. We did this because the neuronal dynamics of the same culture of neurons may change drastically from day to day, so we considered it a different dynamical system. The 14 networks came from 11 independent cultures, from a total of five different animal preparations (Supplementary Figure 3.9).

### 3.2.7 Signal Processing

*Features of neuronal activity*

The offline analysis of neuronal data was based on the spikes recorded with Experimenter. We used three features to characterize the activity of the different networks: average firing rate, synchrony and the oscillation intensity. The average firing rate corresponds to the average number of spikes per electrode per unit of time. The synchrony measure is based on the $\chi$ metric, which compares the variance of the network activity with the variance of each neuron's activity (Golomb and Rinzel 1994),

$$\chi^2 = \frac{Var\left(\frac{1}{N}\sum_{i=1}^{N} V_i(t)\right)}{\frac{1}{N}\sum_{i=1}^{N} Var(V_i(t))}, \tag{3.4}$$

where $N$ is typically the number of neurons, but here corresponds to the number of electrodes, and $V_i(t)$ is typically the voltage of neuron $i$, but here corresponds to the spike train of electrode $i$ convolved with a square wave of 50 ms and binarized. This signal defines the periods on which a given electrode is considered active. We are thus measuring the degree of coactivation of the multiple electrodes, regardless of the number spikes exhibited by different electrodes at the moment of synchronization. We did not want this factor to leak to the synchrony measure because different electrodes may be synchronized but with different number of spikes per activation (which could be due to different number of neurons surrounding each electrode, for example). The oscillation intensity is measured by calculating SNR of fundamental frequency of the Power Spectrum Density. This was calculated using the Matlab functions *pwelch* for the PSD and *snr* for the SNR, which computes the ratio of the summed squared magnitude of the fundamental frequency to that of the noise. The snr function automatically detects the fundamental frequency and assures that

the harmonics and DC component are not mistaken as noise. The fundamental frequency corresponds to the frequency at the peak of the PSD. For the firing rate and synchrony, the effect of each stimulation protocol was evaluated by calculating the metric's fold change form OFF to ON. For the oscillation intensity we calculated the difference between the OFF and ON periods, since it is measured in dB, which is a logarithmic scale.

*Unsupervised clustering of neuronal dynamics*

The different states of neuronal activity of a given recording (section 3.3.5) were identified by calculating the firing rate and synchrony in a moving window. The width of the moving window corresponded to the average duration of five periodic events in the analysed recording. The moving step was half the window width. The clusters of neuronal dynamics were automatically identified with Gaussian Mixture Models (GMMs) applied to the bidimensional matrix of firing rate and synchrony. We tested GMMs with 1 to 4 groups and each condition was repeated 10 times with different initial values. The ideal clustering was chosen based on the Bayesian Informative Criterion. We calculated the centroid of each cluster and considered the asynchronous states as being those with mean synchrony below 0.5 and mean firing rate below the average for that network.

### 3.2.8 Statistical Analysis

*Controllable and uncontrollable networks*

The classification of networks as controllable or uncontrollable (Figure 3.3) was based on the assumption that controllable networks could be driven to specific subspaces of neuronal activity with a given stimulation protocol whereas the uncontrollable networks could not. We started by calculating the fold change of the four metrics (synchrony, firing rate, oscillation frequency and intensity) for all trials performed with each network. To evaluate whether the spaces reached during stimulation were statistically different, we calculated the Multivariate ANOVA (MANOVA) considering the fold change of the three metrics, simultaneously. Here, the null hypothesis was that the modulation results were identical for all stimulation protocols.

*Neuromodulation results for controllable networks*

For a given network, we calculated the average fold change in synchrony and firing rate, and the difference in SNR of the main oscillation across all the trials performed with each stimulation protocol. Then, the comparisons between groups were performed using one way ANOVA with repeated measures.

*Time spent in asynchronous state*

To compare the efficacy of the different stimulation protocols in promoting the transition to a stable asynchronous state in a multi-stable network, we calculated the percentage of time spent

in the asynchronous state during the OFF and ON periods for all the trials. We then compared if there was a significant change from OFF pre-stimulation to ON and from OFF pre-stimulation to OFF post-stimulation using one way ANOVA with repeated measures

### 3.2.9 *In silico* model

Our in silico networks were composed of 1000 randomly connected Izhikevich neurons with varying percentages of regular spiking excitatory neurons (a = 0.02; b = 0.2; c = -65; d = 8) and fast spiking inhibitory neurons (a = 0.1; b = 0.2; c = -65; d = 2) (Izhikevich, 2003). The neurons were connected in an all-to-all architecture of varying synaptic weights according to the connectivity matrix S. The overall synaptic weights of the different simulations were modulated by scaling S. Each condition was simulated five times. The model was based on the script provided in https://www.izhikevich.org/publications/net.m. The simulations had an initial period of 500 ms for the network activity to stabilize. After the stabilization period, we had 2000 ms of spontaneous activity and 2000 ms under the aDFC stimulation protocol. The aDFC algorithm was identical to that applied in vitro, considering only the positive component of the actuation signal. A pool of 100 neurons was randomly selected for stimulation. Each electrical pulse corresponded to an instantaneous input of 20 mA to the stimulated neurons. The firing rate and synchrony measures were computed using the same algorithms as those used in vitro (section 5.7) and averaged across the five repetitions of each simulation condition.

## 3.3 Results

To assess the potential of DFC as a closed-loop brain stimulation protocol (Figure 3.1.A, top) for suppressing pathological neuronal oscillations, we developed an *in vitro* setup composed of hippocampal networks cultured on microelectrode arrays (MEAs), a workstation running the control algorithms and specialized hardware (2100MEA-System, Multichannel Systems) interfacing both environments at very low latencies (Figure 3.1.A, bottom). Our *in vitro* setup provides a powerful testbed to analyse DFC algorithms as it incorporates neuronal populations that display synchronous oscillations (here in the form of quasiperiodic network bursts with periodicities in the seconds range), an acquisition system that measures the population's extracellular activity at high spatiotemporal resolution enabling the inference of synchrony among neurons, on-demand electrical pulse generators to stimulate the network at high spatiotemporal resolution and a feedback loop with low actuation latency in the milliseconds range (orders of magnitude below that of the neuronal dynamics under control).

The core of the DFC algorithm used (Figure 3.1.B, black) is analogous to implementations reported in the literature (see Methods section for details). After spike detection (Figure 3.1.C, top row), the instantaneous population firing rate is calculated (Figure 3.1.C, middle row) and filtered by a damped oscillator (Figure 3.1.C, bottom row in blue). The resulting signal is a good representation of the oscillatory behaviour of the network, such as that exhibited while the stimulation

is turned OFF (Figure 3.1.C, left). The feedback signal (Figure 3.1.C, left, bottom in red) is then created by subtracting the ongoing oscillation to the half-cycle delayed oscillation (Figure 3.1.C, bottom row in grey). This way, when the activity is oscillatory, the actuation signal is maximal at the antiphase of the neuronal oscillation (Figure 3.1.C, left and middle). The obtained feedback signal is linearly translated to stimulation frequency by applying a fixed gain. Since the actuation signal is explicitly dependent upon the neuronal oscillation, stimulation is highly activated when synchrony is high (Figure 3.1.C, middle) and decreases as this firing pattern vanishes (Figure 3.1.C, right). Expanding on the DFC algorithm, and to respond to the dynamic nature of biological neuronal systems, we improved canonical DFC to adaptive DFC (aDFC) by adding an extra component that detects the synchronous firing events (Figure 3.1.C, black circles in middle row). This way, aDFC monitors the current periodicity of the neuronal oscillation (T) and updates the delay duration (T/2) and natural frequency of the oscillator ($\omega$) accordingly (Figure 3.1.B, blue and Supplementary Figure 3.8).

The neuromodulation performance of the control algorithms was analysed in a series of experiments with multiple neuronal populations on MEAs. Each experiment had three stages: 5 minutes of spontaneous activity pre-stimulation (OFF), 5 minutes under stimulation (ON) and 5 minutes of spontaneous activity post-stimulation (OFF). Every network, here defined as a neuronal culture on a given day *in vitro*, was submitted to three different stimulation algorithms: standard DFC, aDFC and random stimulation. During random stimulation, the electrical pulses were triggered by a Poisson process with an average frequency identical to that obtained in the preceding aDFC test. In all algorithms, stimulation was provided as monophasic negative voltage pulses at controlled timings. To ensure that the individual pulses had a comparable effect on the different neuronal networks, we chose a single stimulation electrode that triggered a neuronal response measured by neighbouring electrodes and the minimum voltage amplitude that achieved that response (see Methods section).

### 3.3.1 Adaptive component is required due to the dynamic nature of neuronal populations

Changes in the neuronal dynamics caused by the stimulation protocols can be evaluated using the wavelet transform of the instantaneous network firing rate (Figure 3.2.A and 3.2.B). This representation evidences how the bursting periodicity of the network changes over time. Consistent inter-burst intervals (i.e., consistent neuronal oscillations) will appear as an horizontal line in the wavelet domain, corresponding to the fundamental frequency of the oscillation. When stimulation was activated, both DFC and aDFC showed clear effects on the neuronal dynamics. DFC stimulation imposed a change to a faster and very consistent oscillation (DFC periods in Figure 3.2.A and 3.2.C). This happens because DFC initially forces the network to fire in its antiphase, creating a change in the bursting periodicity. Since this change is not monitored to update the controller's delay, the delayed oscillator (Figure 3.1.C grey) no longer matches the anti-phase of the ongoing oscillation. Therefore, stimulation becomes prejudicial as it is sent in phase with the

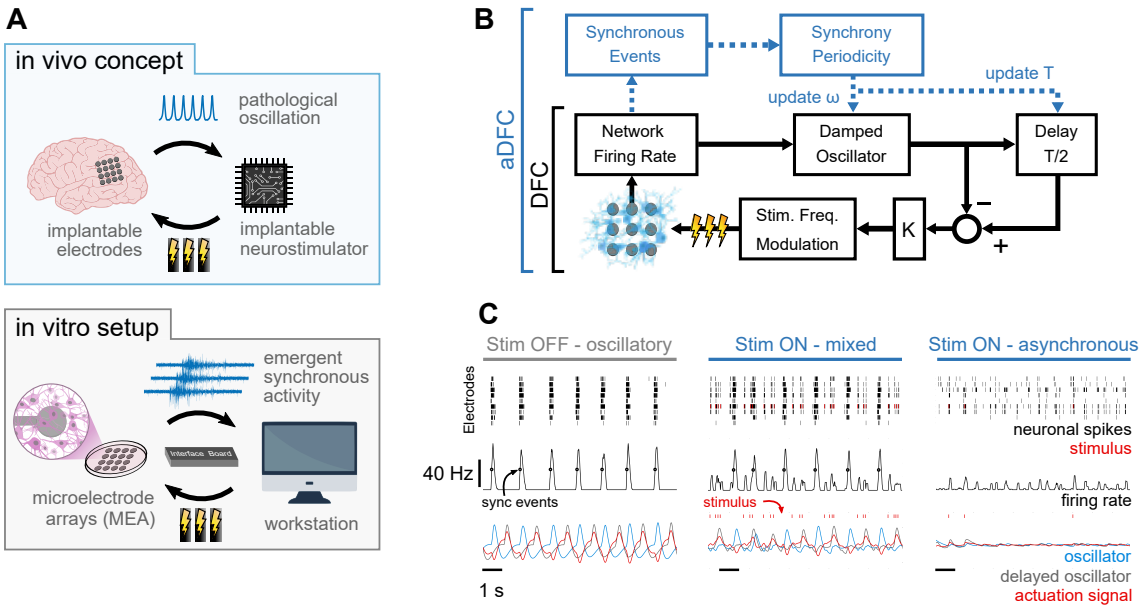



Figure 3.1: **Controlling pathological oscillations with Adaptive Delayed Feedback Control (aDFC).** (A) In vivo concept of implantable closed-loop neurostimulation and the analogous in vitro setup developed. Here, we have dissociated hippocampal neurons as a brain, microelectrode arrays as implantable electrodes, emergent periodic network bursting as pathological synchronous oscillations, a workstation as an implantable chip and an external stimulator as an implantable neurostimulator. (B) Scheme of the Delayed Feedback Control algorithms. The instantaneous population firing rate is calculated and filtered by a damped oscillator. The actuation signal (stimulation frequency) is obtained by subtracting the ongoing oscillation to the past half-cycle delayed oscillation. The natural frequency of the oscillator, $\omega$, and the duration of the half-cycle delay are updated online (adaptive component in blue) by detecting the synchronous events and their current periodicity, T. (C) Controller computations under different firing regimes – oscillatory (left), asynchronous (right) and mixed, i.e., oscillatory with induced sparse activity (middle). Spike detection is performed for the monitoring electrodes (top) and the spike train is convolved online with a square window to compute the instantaneous population firing rate (middle). Synchronous events (black circles) are detected using a threshold crossing method and used to update the oscillation periodicity. The actuation signal (red) is built by subtracting the oscillator and delayed oscillator signals (blue and grey, respectively) and translated to a stimulation frequency signal.

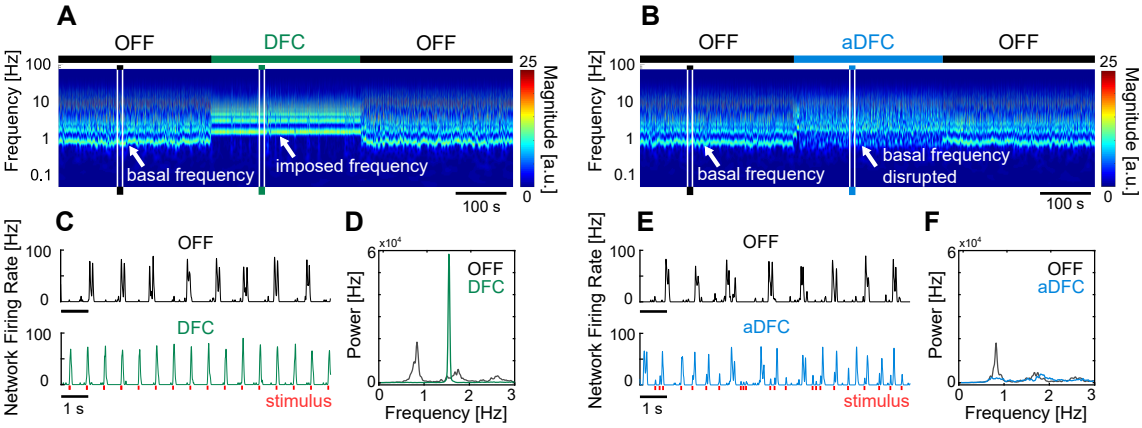



Figure 3.2: **Representative example of the effect of DFC and aDFC stimulation protocols.**(A,B) Wavelet transform of the instantaneous firing rate of a network under DFC (A) and aDFC (B) protocols. (C,E) Example of the instantaneous firing rate of the network before stimulation (black) and during DFC (C, green) and aDFC (E, blue) stimulation. The signals shown in C and E correspond to the narrow time windows marked in A and B, respectively. (D,F) Power Spectrum Density of the instantaneous firing rate of the network before stimulation (black) and during DFC (D, green) and aDFC (F, blue) stimulation.

new oscillation, reinforcing it (DFC period in Figure 3.2.C). On the other hand, aDFC disrupts the spontaneous oscillation frequency exhibited by the network, creating a firing regime where there is no predominant oscillation (aDFC periods in Figure 3.2.B and 3.2.E). In this case, the controller is constantly forcing the network to fire in antiphase and aDFC adapts as the network oscillation changes with the stimulation. To quantify this effect, we calculated the magnitude of the main oscillation by calculating the signal-to-noise ratio (SNR) of the Power Spectrum Density (PSD) of the instantaneous firing rate for the ON and OFF segments (Figure 3.2.D and 3.2.F).

### 3.3.2 Network controllability is conditioned by intrinsic population dynamics

The effect of the different stimulation protocols was evaluated using three metrics calculated for the OFF and ON periods: synchrony, firing rate, and oscillation intensity. Synchrony measures the degree of co-activation of multiple electrodes (see Methods); firing rate corresponds to the average number of spikes per electrode per unit of time (not to be confused with the instantaneous network firing rate signal presented in Figures 3.2.C and D); oscillation intensity corresponds to the SNR of the PSD of the instantaneous firing rate signal (Figures 3.2.D and F). The three stimulation protocols were repeated multiple times for each network following the OFF-ON-OFF routine. Their effects on a given network are evidenced by normalizing the values of each feature to the initial OFF period (before stimulation) across the multiple trials (Figure 3.3.A). This also reveals how the dynamics are recovered once the stimulation is turned OFF. Note that, for the Oscillation Intensity, this normalization corresponds to the difference between the ON and OFF periods, since the SNR scale is logarithmic [dB].

Some networks were consistently driven to distinct regions of the features space with each stimulation protocol (Figure 3.3.B). We refer to these as *controllable* networks – hypothetically, an ideal controller could lead such networks to any desired reachable state. On the other hand, some networks did not exhibit consistent modulation when stimulated using the different protocols (Figure 3.3.C). We termed these *uncontrollable* networks. Based on this definition, each network was classified as either controllable or uncontrollable by calculating the multivariate ANOVA for the ON values of the three metrics (see Methods for details). Networks with low p-values ($p < 0.05$) (i.e., networks with significantly distinct neuromodulation results for the different stimulation protocols) where considered controllable (Figure 3.3.D). Having recognized that some networks are controllable, we questioned if there are any common traits that make them more prone to neuromodulation. Even though there may be an indefinite number of hidden variables that influence the controllability of a neuronal population, some of these properties could be translated into specific firing dynamics, which we can access with the MEAs. If that is the case, then we expect the controllable and uncontrollable networks to have distinct fingerprints in their spontaneous firing dynamics. To evaluate this hypothesis, as well as ensure that the uncontrollable networks did not simply result from an ineffective stimulation electrode, we compared the spontaneous firing rate and synchrony during the initial OFF period (before stimulation) for all the experiments performed with all networks. This revealed that, as hypothesised, controllable and uncontrollable networks lie in different regions of the feature space (Figure 3.3.E and Supplementary Figure 3.9). Furthermore, controllable networks are in the centre, which is associated with intermediate levels of spontaneous firing rate and synchrony. This can be further evidenced by calculating the principal components (using Principal Components Analysis, PCA) of the standardized map and projecting the values of each experiment along the principal component. A histogram shows that the experiments performed with controllable networks are in the centre, surrounded by a bimodal distribution of the experiments performed with uncontrollable networks (Figure 3.3.F).

### 3.3.3 aDFC disrupts oscillations and decreases overall synchrony in controllable networks

The performances of DFC, aDFC and Poisson stimulation were evaluated by comparing their average effects on the three metrics across all the controllable network (Figure 3.4), where each datapoint corresponds to the metric's mean for all the experiments performed with a given network and stimulation protocol. Uncontrollable networks are not considered here as their results are, by definition, not reproducible. We included trials without stimulation for each network as a control group to capture the natural variability of each metric; in these, we calculated the metrics for two consecutive 5 min blocks of spontaneous activity, representing the initial OFF and ON periods. The most significant distinction between control algorithms was evidenced at the level of the bursting oscillations (Figure 3.4.A). Both aDFC and random stimulation led to a significant decrease in Oscillation Intensity (see also Supplementary Figure 3.10 for a representative wavelet transform under Poisson stimulation). The standard DFC algorithm, on the other hand, led to a

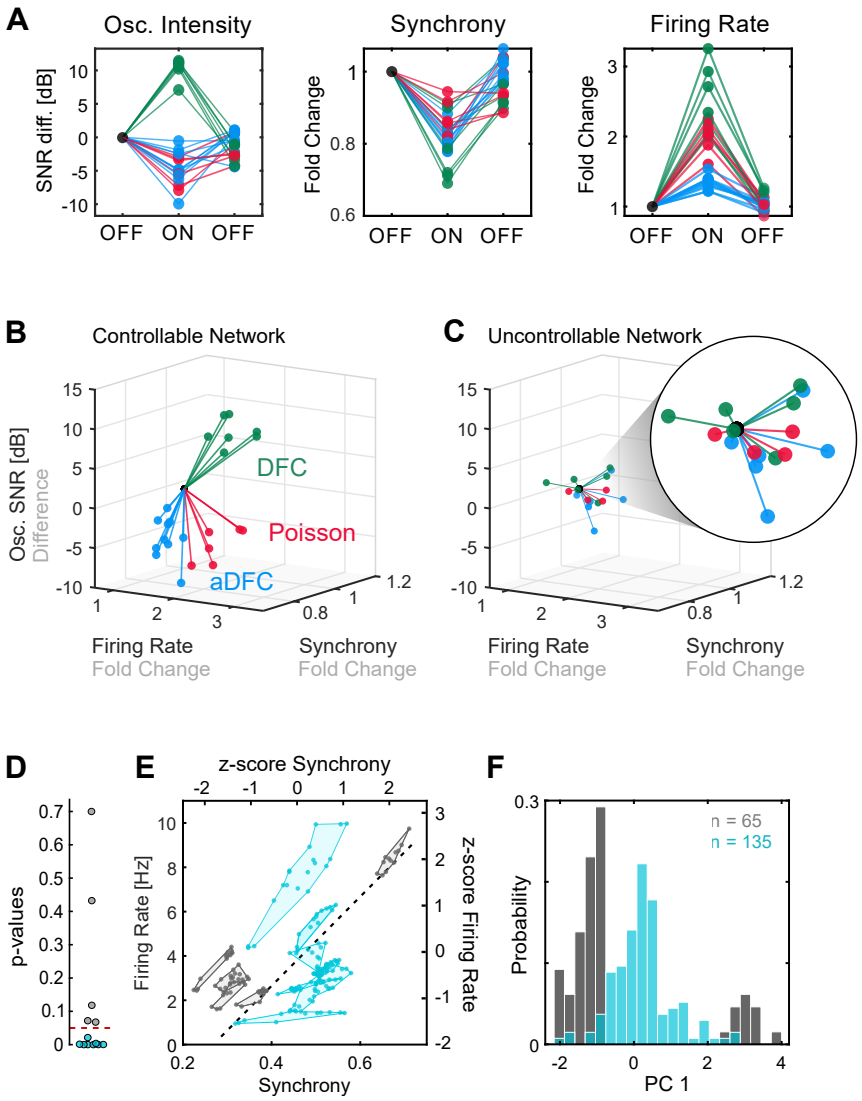



Figure 3.3: **Comparison of controllable and uncontrollable networks.** (A) Results obtained from multiple repetitions of the three stimulation protocols for a single network. We calculated the synchrony, firing rate, and oscillation intensity for the three segments – pre-stimulation (OFF), stimulation (ON), post-stimulation (OFF) – of each experiment and normalized each metric to the value of the pre-stimulation period. For the oscillation intensity, this normalization corresponds to the difference and not ratio, since the SNR is given in dB (logarithmic scale). Each triplet of points corresponds to one experiment performed with either aDFC (blue), DFC (green) or Poisson (red) stimulation. (B) Three-dimensional (3D) feature space – Firing Rate vs. Synchrony vs. Oscillation Intensity – for the network in A. Each line corresponds to the metrics evolution from the initial OFF to ON (from the first to the second point in A). This network is considered controllable since each protocol consistently drove the network to a distinct subspace, meaning that it had a unique and reproducible effect. (C) Representative example of the 3D feature space for an uncontrollable network where the effect of each stimulation protocol is indistinguishable from the others. (D) Networks were separated into controllable (blue) and uncontrollable (grey) by computing the MANOVA of the ON metrics for the different stimulation protocols. Networks with p-values lower than 0.05 had separable effects in the 3D feature space of the metrics and were thus considered controllable (8/13). [*Continues on the following page*]

Figure 3.3: (E) Pre-stimulation firing rate and synchrony for all experiments performed with each network, divided into controllable and uncontrollable networks. Each dot corresponds to one experiment performed with either DFC, aDFC or Poisson in a given network. The experiments performed with the same network are confined in a polygon. The black traced line corresponds to the first principal component (PC1) of the standardized (z-score, i.e., mean equals zero and standard deviation equals one) synchrony and firing rate defining a relevant descriptor of the neuronal dynamics. (F) Distribution of the baseline neuronal dynamics along PC1, defined in E, showing that the controllable networks have an intermediate level of spontaneous firing rate and synchrony.

significant increase in the oscillations. This increase in magnitude was followed by an increase in the oscillation frequency (Figure 3.2A-C and Supplementary Figure 3.11). The only stimulation protocol that led to a significant decrease in synchrony was aDFC (Figure 3.4.B) and its effect was significantly different from that of random stimulation. On the other hand, it was also the only method that significantly increased the average firing rate (Figure 3.4.D), even though the average stimulation frequency was similar for the three methods (Supplementary Figure 3.12).

### 3.3.4 Controllability is affected by synaptic strength and balance between excitation and inhibition

We then sought to understand whether the controllability subspace found for the neuronal cultures portrayed any fundamental property of these type of systems and what parameters could govern such behaviour. We modelled randomly connected networks of 1000 Izhikevich neurons (Izhikevich 2003) with different proportions of regular spiking excitatory neurons and fast spiking inhibitory neurons. We also varied the overall synaptic weight (see Methods for details). Each condition was simulated five times. The simulations were composed of two different periods: spontaneous activity (OFF) and controlled activity (aDFC ON). For each simulation, we computed the overall synchrony and firing rate for the OFF (Figure 3.5.A) and aDFC ON periods (Figure 3.5.B). Low percentages of excitatory neurons and weak synaptic weights within the network led to sparse and chaotic activity (regions 1 in Figures 3.5.A-C). In this scenario, the controller is sporadically activated due to fluctuations of neuronal activity, creating brief moments of synchronization (Figure 3.5.C, top). High percentages of excitatory neurons and strong synaptic weights led to highly ordered network dynamics with strong bursting activity (regions 3 in Figures 3.5.A-C). Here, the controller could not disrupt the dominating spontaneous activity (Figure 3.5.C, bottom). However, in the transition between these two network states, there was a parameter subspace that generated intermediate baseline levels of firing rate and synchrony (regions 2 in Figure 3.5.A-C). In parallel to what was found *in vitro*, the networks were controllable in this intermediate region (see also Supplementary Figure 3.13). Here, aDFC could desynchronize the network shortly after activating the stimulation (Figure 3.5.C, middle). Thus, among the multiple possible parameters governing network controllability, the excitatory/inhibitory balance and overall network connectivity strength could explain the separation into controllable and uncontrollable networks seen *in vitro*.

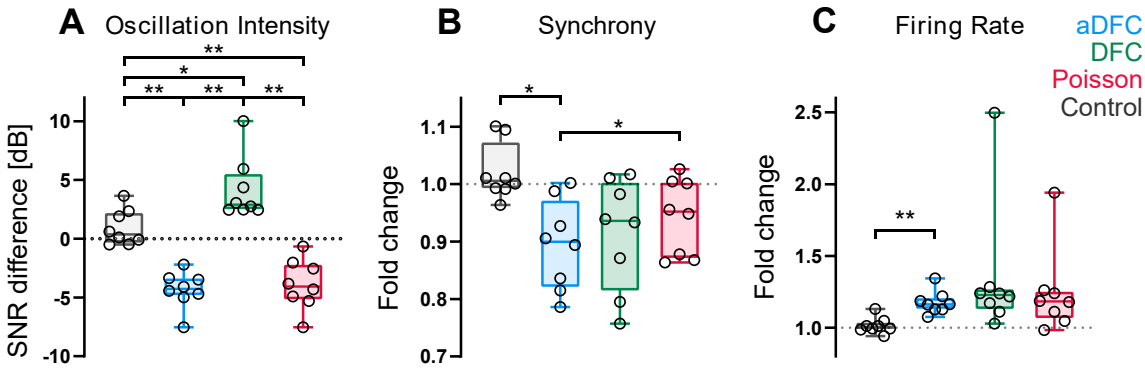



Figure 3.4: **Effect of aDFC, DFC and Poisson stimulation in controllable networks.** Fold change in oscillation intensity (A), synchrony (B) and firing rate (C) of the 5 min stimulation period compared to the 5 min pre-stimulation period with aDFC (blue), DFC (green) and Poisson (red) stimulation. The control group (grey) corresponds to baseline recordings where there is no stimulation and thus portrays the natural variability of each metric within two subsequent 5 min blocks. Each point represents the average effect of a stimulation protocol across all trials for a given network. We used the one-way ANOVA with repeated measures as statistical test (n = 8; * $p < 0.05$; ** $p < 0.01$; details of statistical tests in Table S1).

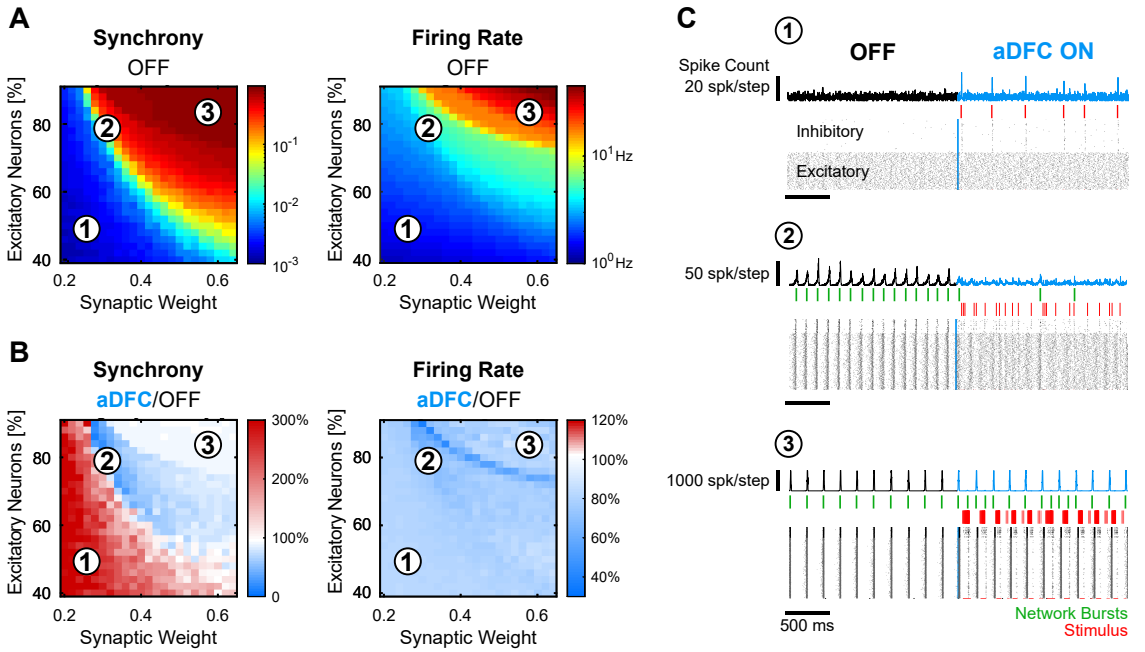



Figure 3.5: **In silico networks also present a controllable subspace for intermediate baseline levels of firing rate and synchrony.** (A) Level of synchrony (left) and firing rate (right) for the period of spontaneous activity. Each pixel represents the average synchrony and firing rate of five different simulations under the same conditions of excitatory/inhibitory balance and synaptic weight. (B) Average change in synchrony (left) and firing rate (right) of aDFC ON period compared to the OFF period. (C) Representative example of the simulations under three distinct network parameterizations (1, 2 and 3 in A and B). Each panel has the instant spike count at the top, the raster plot at the bottom and the network bursts (green) and triggered stimuli (red) in the middle.

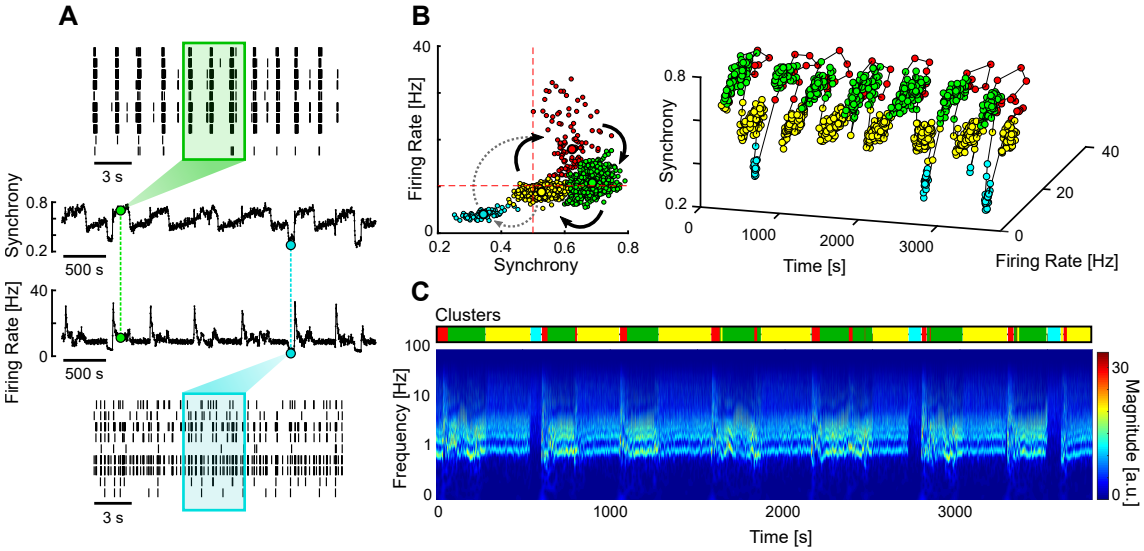



Figure 3.6: **Multi-stable network with sporadic transitions to asynchronous state.** (A) Spontaneous synchrony and firing rate over time (middle). Representative raster plots of the synchronous (top) and asynchronous states (bottom). The coloured boxes represent the sliding window over which the firing rate and synchrony are calculated. (B) Synchrony vs. Firing Rate phase plot evidencing the cyclic behaviour of the network, with sporadic transitions to the asynchronous state (blue cluster). The clusters were identified using unsupervised Gaussian mixture models. (C) Changes in the frequency domain correlate with the automatically identified clusters (top bar).

### 3.3.5 aDFC can promote transitions to stable asynchronous states

The majority of the *in vitro* cultures had monostable dynamics, presenting a very consistent firing pattern over time (e.g., OFF periods of Figure 3.2). In these networks, even though aDFC reduced neuronal synchronization (Figure 3.4.C) and disrupted baseline oscillatory activity (Figure 3.2.D-F and Figure 3.4.A), it generally did so by creating sparser bursting with no particular periodicity. In monostable networks, aDFC did not promote a phase transition to a stable asynchronous state as opposed to what is typically seen in DFC computational studies (M. Rosenblum and Pikovsky 2004; Vlachos et al. 2016; Popovych et al. 2017; Daneshzand, Faezipour, and Barkana 2018; Popovych and Tass 2019; Yu et al. 2021; Zhou et al. 2020). We also found a less common type of networks that had multi-stable dynamics, spontaneously transitioning between different firing regimes. Multi-stable networks were particularly interesting when they presented sporadic transitions to an asynchronous state (AS) as it showed that these networks could autonomously sustain asynchronous activity, even if only for short periods of time (Figure 3.6). The different stable states were identified using unsupervised clustering on the values of firing rate and synchrony calculated over a sliding window (Figure 3.6.A,B). A given state was considered asynchronous if its average synchrony was below 0.5 and the mean firing rate was below the average of that network. The obtained clusters correlate with the different regimes seen in the frequency domain (Figure 3.6.C).

To evaluate the effectiveness of each stimulation protocol in promoting transitions to AS, we compared the percentage of time spent in AS during the OFF and ON periods in multi-stable networks (Figure 3.6). Since our experimental protocol was composed of 5 min blocks (OFF-ON-OFF), the desynchronizing effect of the closed-loop stimulation was only clear for networks that spontaneously exhibited rare and brief (considerably less than 5 mins) AS. This way, a significant increase in the time spent in AS during the ON periods was most likely caused by the closed-loop stimulation. Out of all the networks tested, one presented the required criteria to evaluate this effect (Figure 3.7.A and B). aDFC consistently promoted desynchronization in all the trials performed with this network, leading to a considerable increase of time spent in AS (Figure 3.7.C,D). DFC stimulation, on the other hand, tended to impose a new bursting periodicity (as already reported), sometimes interrupted by transitions to AS (Figure 3.7.E,F). However, the percentage of time spent in AS during DFC stimulation is within the range of spontaneous transitions and thus cannot be directly attributed to DFC. Random Poisson stimulation with identical stimulation frequency promoted sparser bursting with some transitions to AS, but these are also within the range of spontaneous transitions (Figure 3.7.G,H). This means that the precisely timed stimuli determined by aDFC are causing the transitions to the asynchronous state. Also, the random Poisson stimulation, being an open-loop protocol, does not stop once the desired asynchronous state is achieved, which could be a significant drawback in a clinical setting.

## 3.4 Discussion

Controlling brain activity is a long-standing goal in neuroengineering and holds great potential for multiple clinical applications. This is a challenging problem because neuronal systems have complex dynamics, virtually impossible to fully observe, and the space of actions available to an external controller is highly restricted. Most of these constraints are circumvented in computational studies, allowing researchers to explore solutions in a highly controlled environment. However, under such simplified conditions, these solutions may converge to algorithms that fail when applied to biological neuronal populations. It is, therefore, important to assess the performance of such algorithms with simple biological neuronal systems before moving to *in vivo* trials. In this work, we tested, for the first time, the efficacy of conventional DFC in an *in vitro* setting and reported that it promoted a new bursting frequency, instead of ablating it. This is in line with the claim that linear DFC could lead to synchronization if the delay deviated substantially from the half-period (Pyragas, 2006). This was shown *in silico* by testing different delays with the same network. Here, however, the deviation from the half-period delay was caused by the fact that DFC stimulation induced a faster periodicity, while the delay was kept constant. The neuronal dynamics aligned with the non-adaptive controller to the point that almost all activity was confined to stimulus responses (Supplementary Figure 3.12). This was not predicted by computational studies, possibly because the simple *in silico* networks used were not as dynamic as biological ones and, therefore, could not sustain different synchronization periodicities.

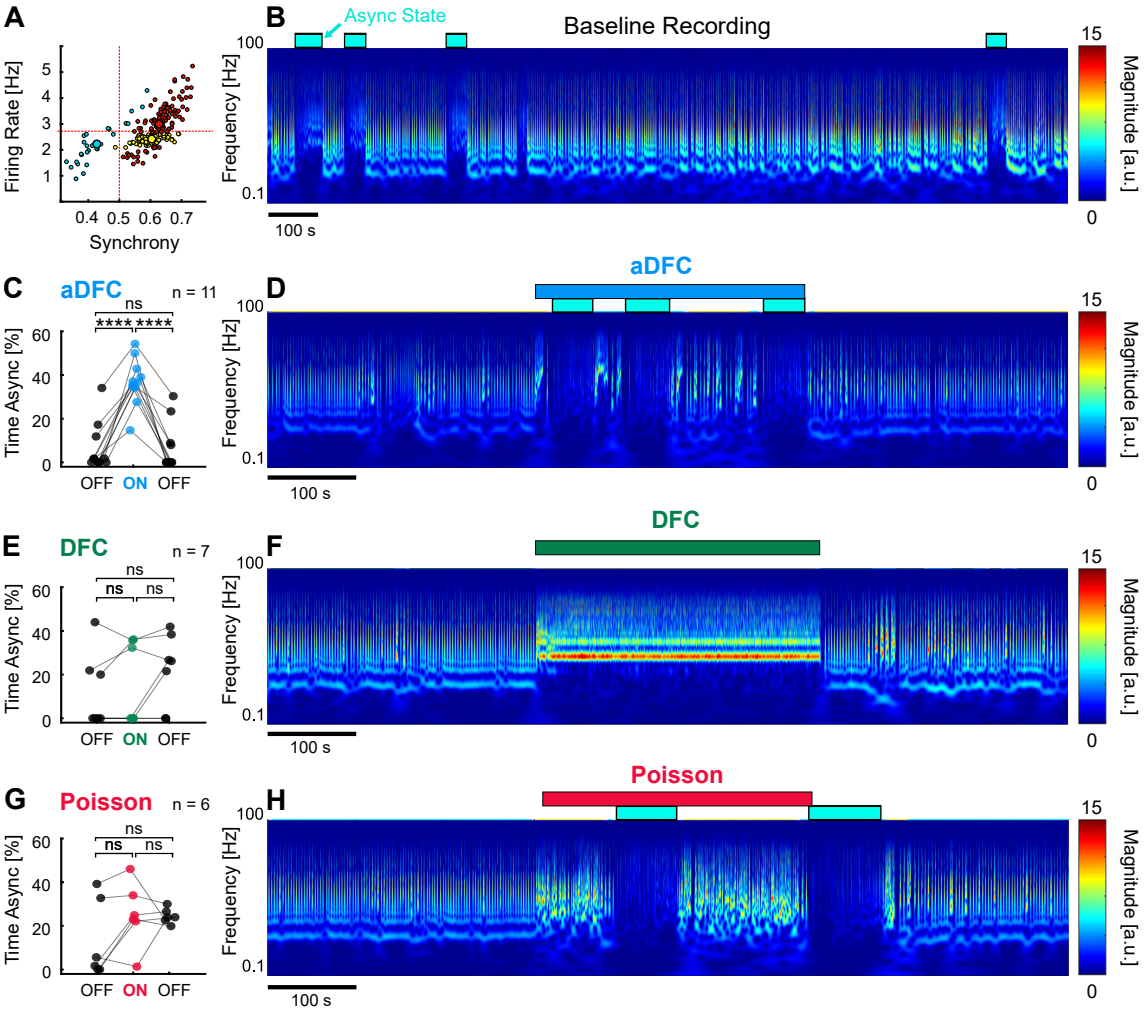



Figure 3.7: **aDFC can promote transition to asynchronous state.** (A) States of spontaneous activity for a given multi-stable network. Blue cluster represents the sporadically reached asynchronous state. (B) Frequency domain of the recording considered in A. Blue rectangles represent the regions associated with the automatically identified asynchronous state. (C, E, G) Percentage of time in asynchronous state during the OFF-ON-OFF protocols for aDFC (C), DFC (E) and Poisson (G) stimulation. Each triplet of points represents an experiment performed with the corresponding stimulation protocol. We compared the time spent in AS during the ON and the second OFF periods with the initial OFF period using one-way ANOVA with repeated measures (**** $p < 0.0001$; details of statistical tests in Table S2). (D, F, H) Representative example of an experiment performed with aDFC (D), DFC (F) and Poisson (H) stimulation portrayed in the frequency domain. The stimulation period is represented with the associated coloured bar on the top. The identified asynchronous states are represented by the blue rectangles, as in B.

The adaptive version, aDFC, effectively disrupted the oscillations and decreased network synchrony. Interestingly, random stimulation, using similar stimulation frequency, was also able to achieve this, although without significantly decreasing synchronization. Electrical stimulation may be acting as a disrupter of the emergent network dynamics, as has been proposed for high-frequency deep brain stimulation (Chiken et al. 2016). We showed that only a subset of networks was differently modulated by each stimulation protocol. However, it is important to note that controllability is contingent on the experimental protocol. In particular, the fact that we only used one stimulating electrode may restrict the pool of controllable networks. We did so to follow the principle that DFC should control the system using minimal intervention (M. G. Rosenblum and Pikovsky 2004). Also, despite the fact that aDFC decreased the synchrony of controllable networks, the neuronal activity was still mostly dominated by synchronous events. These results were far from those obtained in *in* silico studies, where the synchronous events vanish almost completely. We point to two possible causes. First, with our setup, we could only apply excitatory electrical stimulation, much like existing implantable stimulators. For this reason, when the actuation signal (here modulated as a stimulation frequency) was negative, we sent no input to the network. Nonetheless, in theory, using only the positive phase of the actuation signal should still prevent the formation of synchronous events, as it did in our *in silico* tests. So, it is possible that the limiting factor is the neuronal model itself. The unconstrained growth of our hippocampal cultures may have led to highly interconnected networks whose neurons simply cannot fire asynchronously. This is further supported by the fact that aDFC consistently promoted the asynchronous state in a network that could already sustain it autonomously. Again, these considerations are conditioned to the method used, whose main goal was to achieve desynchronization with minimal intervention. Other methods could possibly achieve network asynchrony by following more invasive approaches. Actually, it was already shown that using continuous high frequency stimulation (10-50 Hz, whereas our method uses around 1 Hz) distributed across multiple electrodes managed to prevent network synchronization, although at the cost of highly increasing the overall network firing rate (Wagenaar et al., 2005b), thus possibly establishing another highly non-physiological state (bursts disappeared for firing rates above 20 times the level of spontaneous firing).

Despite being able to disrupt the baseline oscillation and decreasing synchrony levels, there is still margin to improve the aDFC algorithm by making it self-adaptive in terms of the actuation gain, pulse amplitude and selection of stimulation electrode(s), all of which were tuned manually before starting the experiments. It remains to be explored whether such an adaptive controller would consistently achieve full desynchronization in this type of networks.

To conclude, the technological advances made in implantable neurostimulators open the door for the use of algorithms capable of controlling pathological neuronal dynamics. We show that aDFC has the potential to be a promising candidate and believe that the considerations extracted from this *in vitro* study will support the translation to future *in vivo* experiments.

## 3.A Supporting Information

### 3.A.1 Acknowledgements:

This work was partially financed by FEDER—Fundo Europeu de Desenvolvimento Regional funds through the COMPETE 2020—Operacional Programme for Competitiveness and Internationalisation (POCI), Portugal 2020, and by Portuguese funds through FCT—Fundação para a Ciência e a Tecnologia/Ministério da Ciência, Tecnologia e Ensino Superior, in the framework of project PTDC/EMD-EMD/31540/2017 (POCI01-0145-FEDER-031540).
Domingos Leite de Castro was funded by FCT – Fundação para a Ciência e a Tecnologia, grant contract SFRH/BD/143956/2019.

### Data Availability

The data that support the findings of this study are openly available at the following URL/DOI: https://doi.org/10.5281/zenodo.10138446.

## Supplementary Figures

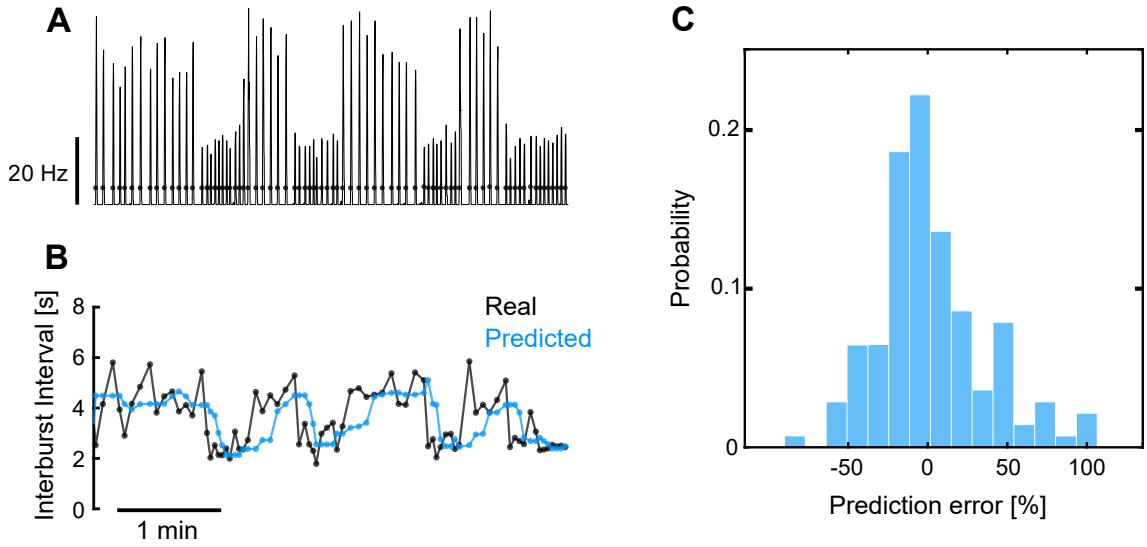



Figure 3.8: **Representative example of aDFC adaptability in a network with fast changing dynamics.** A) Instantaneous network firing rate during spontaneous activity. The black dots represent the real-time detection of network bursts performed by the adaptive controller. B) Interburst interval predicted by the aDFC controller and the associated real interbursts. Every time the controller detects a new network burst (black dots in A) it estimates when the next burst would spontaneously occur by calculating the median of the previous five interburst intervals. This prediction is used to properly tune the controller's half-cycle delay and the natural frequency of the oscillators used to create the actuation signal. In this example the stimulation is turned OFF, so the controller is simply tracking the neuronal activity without interfering with it. C) Histogram of the prediction error for the representative example shown in A and B. The largest errors result from the prediction lag associated with the fast transitions between bursting regimes. The fact that the prediction error is typically small, even under these fast-changing network dynamics, means that the half-cycle delay and the oscillator's natural frequency of aDFC are properly self-tuned in real-time.

| | | **Oscillation Intensity** | | **Synchrony** | | **Firing Rate** | |
|---|---|---|---|---|---|---|---|
| Comparison | n | Mean Diff. | p values | Mean Diff. | p values | Mean Diff. | p values |
| Control vs aDFC | 8 | 5.290 | 0.0048 | 0.1266 | 0.0454 | -0.1600 | 0.0082 |
| Control vs DFC | 8 | -3.242 | 0.0164 | 0.1075 | 0.1313 | -0.3332 | 0.2651 |
| Control vs Poisson | 8 | 4.835 | 0.0039 | 0.07756 | 0.1817 | -0.2281 | 0.2071 |
| aDFC vs DFC | 8 | -8.532 | 0.0012 | -0.01904 | 0.5039 | -0.1732 | 0.6420 |
| aDFC vs Poisson | 8 | -0.4550 | 0.9098 | -0.04899 | 0.0142 | -0.06806 | 0.8459 |
| DFC vs Poisson | 8 | 8.077 | 0.0029 | -0.02995 | 0.4456 | 0.1051 | 0.4577 |

Table 3.1: Details of the statistical tests used in Figure 3.4 to compare the modulation results of the different stimulation protocols in controllable networks. We used repeated measures one-way ANOVA with multiple comparisons

| | **aDFC** | | | **DFC** | | | **Poisson** | | |
|---|---|---|---|---|---|---|---|---|---|
| **Comparison** | **n** | **Mean Diff.** | **p values** | **n** | **Mean Diff.** | **p values** | **n** | **Mean Diff.** | **p values** |
| OFFpre - ON | 11 | -0.3072 | <0.0001 | 7 | -0.02274 | 0.771 | 6 | -0.1208 | 0.3023 |
| ON - OFFpos | 11 | 0.3052 | <0.0001 | 7 | -0.07382 | 0.3143 | 6 | 0.07317 | 0.9952 |
| OFFpre - OFFpos | 11 | -0.002 | 0.9989 | 7 | -0.09656 | 0.1706 | 6 | -0.1135 | 0.3444 |

Table 3.2: Details of the statistical tests used in Figure 3.7 to compare percentage of time spent in asynchronous state for each stimulation protocol in a multistable network. We used repeated measures one-way ANOVA with multiple comparisons.

| **Comparison** | **n** | **Mean Diff.** | **p values** |
|---|---|---|---|
| aDFC vs DFC | 8 | -0.2372 | 0.0089 |
| aDFC vs Poisson | 8 | -0.1909 | 0.0181 |
| DFC vs Poisson | 8 | 0.04628 | 0.4459 |

Table 3.3: Details of the statistical tests used in Figure 3.10 to compare fraction of spontaneous spikes during stimulation for the three stimulation protocols. We used repeated measures one-way ANOVA with multiple comparisons.

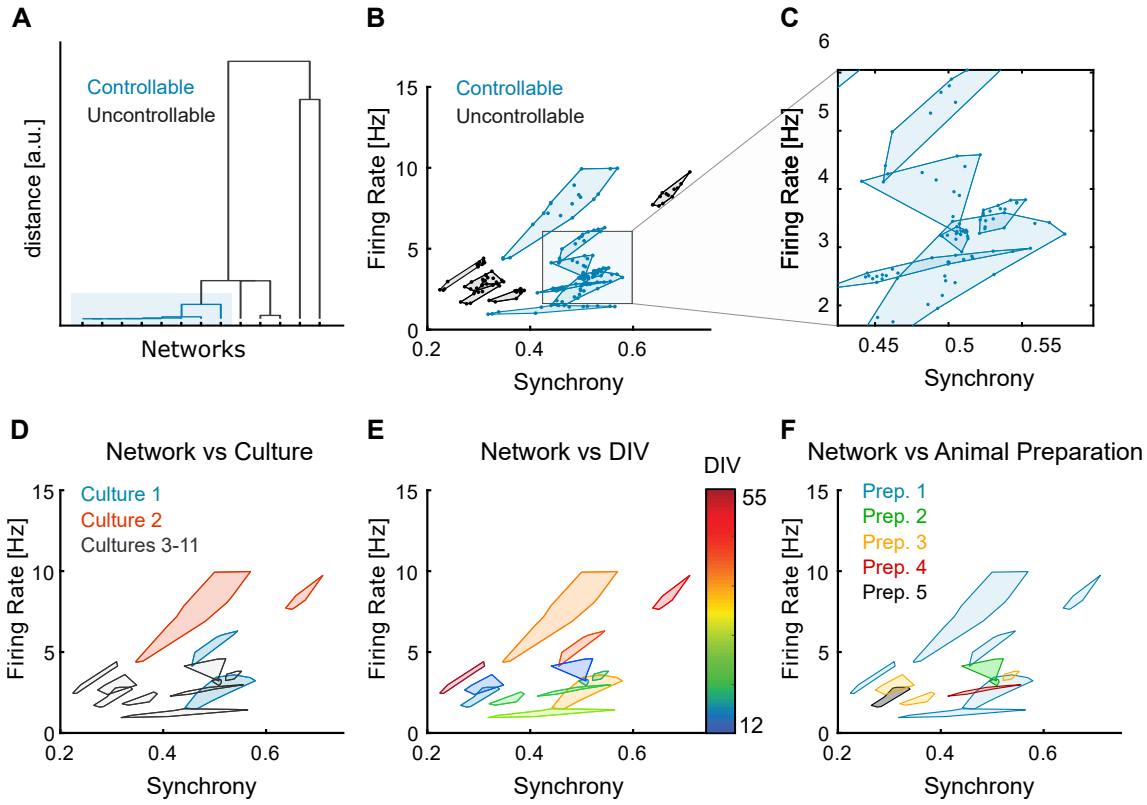



Figure 3.9: **Controllable and uncontrollable networks across cultures, DIVs and animal preparations.** A) Dendrogram of the networks according to the controllability measure (MANOVA p-values of the modulation results obtained with the different stimulation protocols). This representation shows that the controllable networks (p-values $< 0.05$) are more similar to each other than to the uncontrollable ones. B) Spontaneous firing rate and synchrony for the all the trials of the different networks, separated into controllable and uncontrollable networks. C) Detail of B, evidencing how some cultures have very well confined spontaneous dynamics. D) Division of the networks according to the associated neuronal culture. When probed in different days, the same neuronal culture can be considered a different dynamical system because the spontaneous dynamics are clearly different. Therefore, we consider these to be different networks. For example, culture 1 (blue) and culture 2 (orange) led to two different networks each and, in the case of culture 2 one network is controllable whereas the other is not. Cultures 3 to 11 (grey) only have one network each and are represented with the same colour for clarity. E) Evolution of the networks with the DIV. This representation suggests that there is a trend towards higher levels of firing rate and synchrony with the DIVs. Also, the DIV by itself does not predict which networks are controllable and uncontrollable as these groups span multiple overlapping DIVs. F) The networks used came from a total of five different animal preparations. The distinction of controllable and uncontrollable networks cannot be attributed to any particular preparation, as the same preparation can lead to both types of networks.

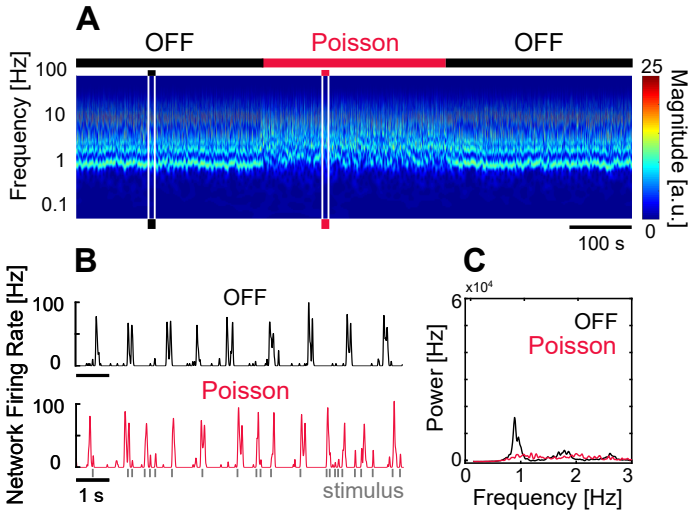



Figure 3.10: **Representative example of the effect of Poisson stimulation in the frequency domain.** A) Wavelet transform of the instantaneous network firing rate during the OFF and ON periods. B) Example of the instantaneous firing rate of the network before stimulation (black) and during Poisson stimulation. The signals shown in B correspond to the narrow time windows marked in A. (C) Power Spectrum Density of the instantaneous firing rate of the network before stimulation (black) and during Poisson stimulation. Since the random Poisson stimulation is an open-loop control algorithm, i.e., it does not adapt to the ongoing network dynamics, the frequency of the Poisson process (which triggers the stimulation) had tuned beforehand. For that, we set the average stimulation frequency to match the average stimulation frequency used in the antecedent aDFC experiment. It is important to note that, even if properly tuned, open-loop Poisson stimulation will not adjust as the network approaches the desired state or stop stimulating once it is reached. This is relevant for cases like the one displayed in Figure 7, where aDFC and Poisson stimulation led to very different network responses.

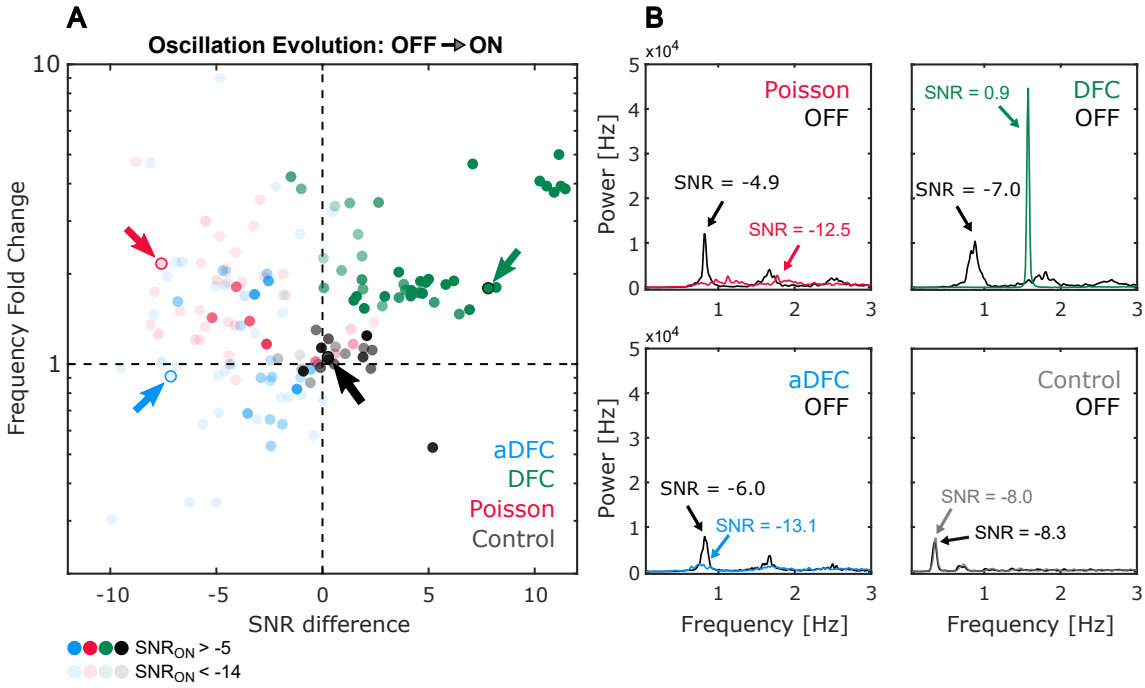



Figure 3.11: **DFC leads to an increase in Oscillation Intensity and Frequency.** A) Change in Oscillation Frequency (fold change) as a function of the change in Oscillation Intensity (SNR difference) from the OFF to ON periods. In this representation, the left and right quadrants are respectively associated with a decrease and increase in Oscillation Intensity when stimulation is turned ON. The lower and upper quadrants are respectively associated with a reduction and increase in the Oscillation Frequency during stimulation. The Oscillation Frequency corresponds to the frequency at the peak of the PSD. To account for the fact that oscillations with very low SNR may not have a meaningfull value of oscillation frequency, the opacity of the datapoints is proportional to the value of SNR during the ON period. Thus, the more transparent points have less meaningful values of Oscillation Frequency. Each datapoint corresponds to an individual experiment performed with a given controllable network and sitmulation protocol (aDFC in blue, DFC in green, Poisson in red, and control in gray). The experiments performed with DFC are mostly located on the upper right quadrant, meaning that it consistently leads faster and more consistent oscillations. aDFC and random stimulation lead to a decrease in Oscillations Intensity, and the Oscillation Frequency is typically neglectable (datapoints are mostly transparent). As expected, the control group, where no stimulation is applied, is located at the center of the map, since the activity is similar during the ON and OFF periods. B) Representative examples of the Power Spectrum Density, from where the Oscillation Intensity and Frequency are calculated, during the OFF and ON periods for the three stimulation protocols and control, marked with arrows in panel A.

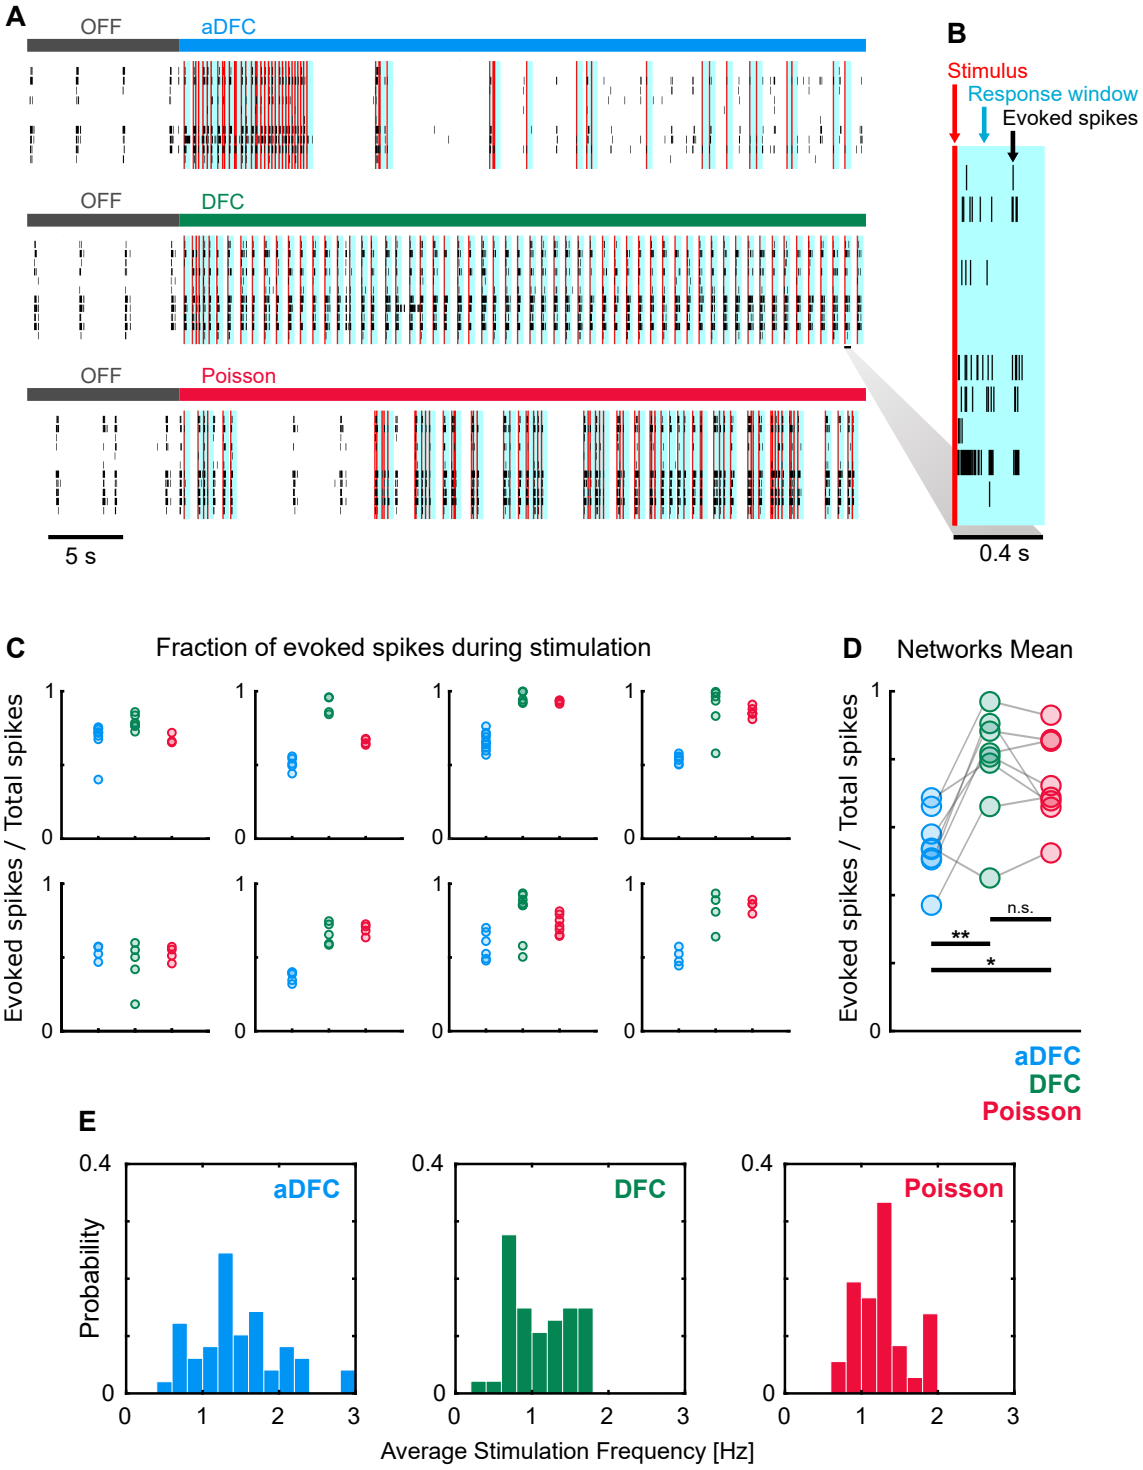



Figure 3.12: **aDFC allows the network to maintain some autonomous firing whereas DFC and Poisson tend to confine most network activity to stimulus responses.** A) Stimuli (red) and respective response windows of 400 ms (blue) for representative examples of aDFC (top), DFC (middle) and Poisson (bottom) applied to the same network. B) Detail of the spikes (black) evoked after stimulation (red) inside the response window (blue). C) Fraction of evoked spikes during stimulation for all the trials performed with the three protocols for each controllable network (one plot per network). When the fraction of evoked spikes is close to 1, the network has almost no spontaneous activity and is simply replicating the actuation signal, which is not ideal. [*Continues on the following page*]

Figure 3.12: D) Comparison of the mean fraction of evoked spikes for the controllable networks. Each triplet of point corresponds to one network. aDFC has a significantly lower percentage of evoked activity, meaning that it allows the network to maintain some intrinsic activity. We repeated measures one-way ANOVA as statistical test (n = 8; **, $p < 0.01$; details of statistical tests in Table S3). E) Average stimulation frequency used by each method for all trials performed with all networks. The overall stimulation frequency is within the same range for all methods. This shows that the results found in D are not explained by differences in the amount of stimulation but are rather due to the specific timing of the stimuli determined by each algorithm.

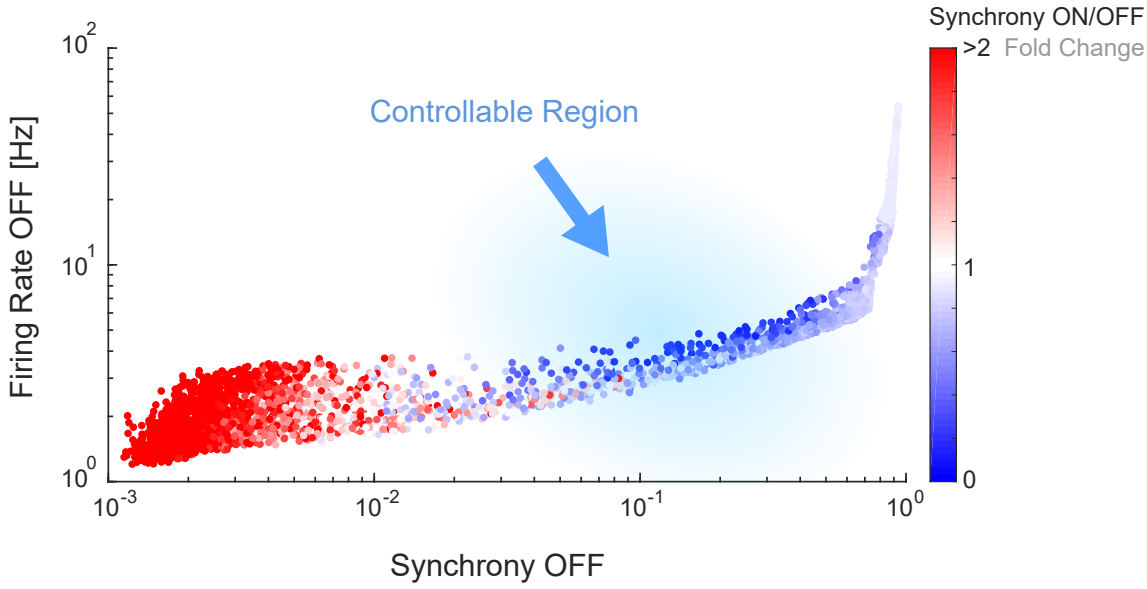



Figure 3.13: **The controllable region of in silico networks is also associated with intermediate levels of spontaneous firing rate and synchrony.** The values of synchrony and firing rate during the initial OFF period were plotted for all the simulations performed, covering all the network parametrizations. The datapoints are coloured according to the change in synchrony level from OFF to ON, as a proxy for controllability. The in silico simulations cover a much larger metric space than the in vitro results, which are confined to the regions between a spontaneous synchrony of 10-1 and 100 and spontaneous firing rate of 100 and 101 (the in silico simulations with very low spontaneous synchrony generated firing dynamics that would not be elected for in vitro experiments).

# Chapter 4

# Deep learning framework to optimise sensory neurostimulation

With the advance of implantable MEA technology, future multi-electrode closed-loop stimulators will operate with high-dimensional inputs and define stimulation patterns that include the spatial component as well. That poses a significant challenge for traditional control methods. Intelligent data-driven algorithms may be the most suitable approach for such problems. With that in mind, our second research question was: *How can machine learning techniques guide electrical stimulation for optimal neuromodulation?*

The paper presented in the previous chapter used MEAs to monitor the collective neuronal dynamics. However, since the algorithm was solely focused on the temporal dynamics, it operated only on a single, carefully selected, electrode. However, in many applications, there is a need to control the spatial activation of specific neuronal cells. The most evident cases relate to the recovery of lost sensory perception, such as cochlear implants, where stimulation must activate different segments of the auditory nerve that encode specific sound frequencies. Another application that will greatly benefit from refined spatial control of neuronal activity is the restoration of visual perception, a problem that, unlike hearing loss, is far from being solved. An efficient retinal prosthesis must achieve precise spatial activation of retinal cells using a grid of implantable electrodes. This is not a trivial problem because the response of retinal cells to electrical stimulation is highly non-linear and there may be considerable interactions between the currents of neighbouring electrodes.

In this chapter, we propose a solution based on a Deep Learning framework which can be used for different type of sensory neuroprosthesis. Here, we will apply it in the context of retinal prosthesis. Our method first learns the spatial rules governing retinal response to multipolar stimulation and then exploits them to define multi-electrode stimulation patterns. This work is only focused on the spatial dimension, so it differs from the previous manuscript in relevant ways. Since we are not

concerned with changes in neuronal activity across time, we are not operating under the paradigm of Dynamical Systems. That translates into the type of neuronal model used which is considerably different from the *in silico* model used in Chapter 3. Instead of a *mechanistic model*, where each neuron is described by a differential equation that determines how its state varies across time, here, the neurons are modelled using a *descriptive model* obtained empirically from *in vitro* experiments (Maturana et al., 2016).

This work was performed at the Department of Biomedical Engineering of the University of Melbourne and resulted in the first-author manuscript titled *Neural Activity Shaping in Visual Prostheses with Deep Learning*, currently in the submission process. It counted with the support of the co-authors David B. Grayden, Hamish Meffin and Martin Spencer, from the Department of Biomedical Engineering, University of Melbourne, Australia. The work resulted in a first-author publication in the Journal of Neural Engineering (Leite de Castro et al., 2024): https://doi.org/10.1021/acsaelm.2c00198.

## 4.1 Introduction

Vision loss affects approximately 230 million people worldwide (Wang et al., 2023). Electrical stimulation of the retina or cortex has emerged as a promising approach to restore some visual perception for people with severely impaired vision. In diseases such as age-related macular degeneration or retinitis pigmentosa, injecting current into the retina allows for the direct stimulation of the surviving inner retinal cells, bypassing the deteriorated outer retina, where the photoreceptors lie (Humayun, 1996; Weiland et al., 2005). Visual prostheses typically comprise a camera that captures the image meant to be seen by the patient, an algorithm that transforms this image into a stimulation pattern and a multi-electrode array (MEA) that delivers the stimuli at the specified locations (Figure 4.1.A). Since visual perception is correlated with the spatial activation of the neuronal cells, the goal of existing retinal implants is to activate neurons that encode for bright regions of the image.

Currently, stimulation patterns are formed by directly translating the brightness values of the target image into stimulation amplitudes at the MEA. The idea is to recreate the target image by combining individual phosphenes (the visual percept elicited by a stimulus) like pixels in a bitmap. However, this method is limited by two main factors: (1) it does not take into account the heterogeneities of the retinal responses to stimulation, which typically lead to distorted phosphenes (Beyeler et al., 2019), and (2) as electrode density increases, the retinal activation induced by neighbouring electrodes overlap, imposing a limit on the sharpness of the perceived images (Tong et al., 2020). This problem is exacerbated when the stimulation is performed with multiple electrodes simultaneously due to spread of current to the inter-electrode regions (Wilke et al., 2011; Schmid et al., 2013).

The blurring problem due to excessive current spreads can be minimized using stimuli of opposite polarities applied simultaneously, leading to an improved visual experience (Figure 4.1.B). This multipolar strategy has been tested in cochlear implants allowing for focused stimulation (van den Honert and Kelsall, 2007; Smith et al., 2013; Shepherd et al., 2017). Multipolar stimulation has already been tested in retinal implants *in vivo* and was shown to produce spatially restricted phosphenes (Spencer et al., 2016). To properly shape the electric field, the electrode amplitudes were tuned according to the measurable stimulation current spreads. In a computational study, (Spencer et al., 2019) proposed that multipolar stimulation could instead be tuned according to the actual retinal responses. This multipolar strategy, named neural activity shaping (NAS), can be used to define optimal multipolar stimulus, as long as one can determine the stimulus pattern that will elicit a desired retinal activation. This non-trivial problem can be solved analytically by inverting a fitted linear-nonlinear model of the retina response to multipolar stimulation (Spencer et al., 2019, 2021). Throughout this manuscript we will refer to this method as analytical model inversion (AMI). However, such approaches rely on strong assumptions regarding retinal response to stimulation and are computationally expensive.

In this work, we develop a non-analytical stimulation algorithm based on artificial neuronal networks (ANNs) that uses NAS to generate sharp retinal activations without any assumptions regarding the retinal response. The ANN, which we refer to as stimulus generator network (SGN), must learn the inverse model of retinal response to multipolar stimulus so that it can then output the stimulus pattern that elicits the desired retina activation. However, there is no ground truth data for the 'correct' stimulus patterns. As such, our method needs to first train a decoder network—measurement predictor network (MPN)—that learns the forward model of retinal response based on pre-recorded responses to multipolar stimulus. Then, the SGN can learn the inverse model of the trained MPN, as an encoder. If the MPN is an accurate representation of the measured retinal response, learning its inverse is equivalent to learning the inverse of the retinal model.

Similar encoder-decoder frameworks have been recently introduced in the visual prosthesis context (Granley et al., 2022; Relic et al., 2022; Wu et al., 2023). However, in these studies, the goal of the encoders was to define stimulation patterns that accounted for phosphene distortions, characteristic of epiretinal implants, using multi-electrode unipolar stimulation. Our encoder, on the other hand, uses NAS (multipolar stimulation tuned according to the measured retina response) to solve the blurring problem that occurs when the phosphenes overlap. Irregular phosphene shapes are thus learned as well, but implicitly. Also, the decoders presented in the previously mentioned studies (Granley et al., 2022; Relic et al., 2022; Wu et al., 2023) estimate an intermediate mathematical model that approximates the phosphenes elicited by stimulation, as described by the implantee. Imprecisions in this intermediate model, either due to the implantee's description or the limited flexibility of the model, leak into the estimated decoder, which then propagate to the encoder that generates the stimulus patterns. Also, since the mathematical model used described phosphenes generated with single electrode stimulation (Beyeler et al., 2019), it may not be suitable to describe response to multi-electrode stimulation because it does not account for the crosstalk between electrodes (Wilke et al., 2011; Schmid et al., 2013; Tong et al., 2020; Moleirinho et al., 2021) . Our decoder network, on the other hand, is trained directly on the measured neuronal responses to multipolar stimulation. Measuring the retinal response also enables its comparison with the target activation which is useful to evaluate performance in a real clinical setting and allows for online retraining of the encoder network to account for long-term changes in retinal response.

We test our algorithm in a simulated environment using a realistic linear non-linear (LNL) mathematical model of retinal response to multipolar stimulation developed with *in vitro* experiments (Maturana et al., 2016) and tested *in vivo* (Halupka et al., 2016). We show that our method learns to produce effective stimulation patterns that elicit the desired retinal response by combining positive and negative currents. The stimulation results are significantly better than those obtained with the conventional approach used in existing devices, being equivalent to the theoretically optimal NAS results obtained with AMI for a LNL retina model, but at a much lower computational cost and without being restricted to an LNL model. This is because artificial networks are considered general function approximators and make no assumptions regarding the underlying model. There-

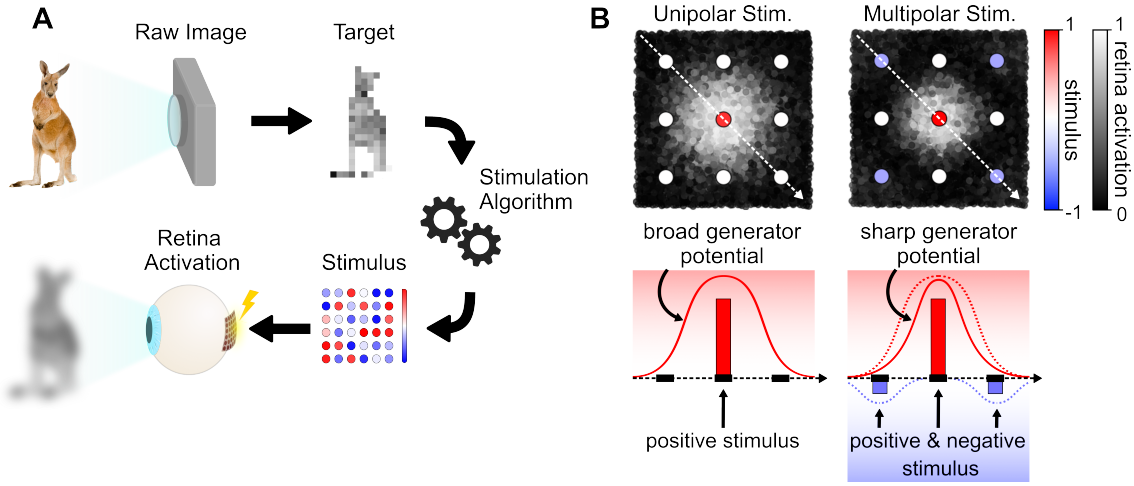



Figure 4.1: **Restoring vision with retina electrical stimulation.** (A) Schematic representation of the stimulation workflow: a camera captures the desired field of view of the implantee; the raw image is subsampled to create a target for the stimulation algorithm; the generated stimulus pattern activates the retina so that the elicited visual perception resembles the original image. (B) Sequential unipolar stimulation (either cathodic-first or anodic-first stimulus in all electrodes) leads to broad retinal activations (left) while simultaneous multipolar stimulation (simultaneous cathodic-first and anodic-first stimulus in different electrodes) can attenuate unwanted current spreads by generating sharper activations (right). The top images show simulated retina activations elicited by the 3-by-3 multi-electrode array while the bottom images represent schematically how the stimulation amplitudes (red and blue rectangles) shape the current spread (traced lines for the contributions of positive and negative stimulus and filled line for the resulting potential). The three electrodes (black rectangles) represented in the bottom images correspond to the three electrodes crossed by a diagonal white dotted arrow in the top images.

fore, our protocol should lead to an effective and personalized stimulation algorithm regardless of the type of retinal response, MEA architecture or electrode density. We believe that our strategy will greatly benefit future retinal implants by leveraging the increased electrode densities—and subsequent stimulation overlap—of high electrode density devices.

## 4.2 Methods

### 4.2.1 ANN-based stimulation protocol

The proposed data-driven stimulation protocol requires training two ANNs sequentially (Figure 4.2). The first network to be trained is the decoder: the MPN; Figure 4.2, top. The goal of the MPN is to predict how a specific implantee's retina responds to a given multipolar stimulus provided by the implanted MEA. In a clinical context, training the MPN would require a dataset composed of multiple stimulus patterns and the associated retinal responses as training data (see section 4.2.2 for details). This dataset would be created by applying a set of random multipolar stimuli via the implant and recording the retinal responses that are measured directly by the MEA. This stage would take less than 20 min, using around 10 000 stimuli at a frame rate of 10 stimuli

per second. Here, we simulate this step using a computational model of the retina (see section 4.2.4 for details on the model). The MPN can then be trained offline.

Once the MPN is trained (i.e. once it can accurately predict the measured retinal response to a given multipolar stimulus pattern), it can be used to train the encoder: the SGN; Figure 4.2, middle. This network outputs the multipolar stimulus pattern that elicits the desired retinal response for a given target image. It can be trained purely offline using a large dataset of natural images and the MPN, which is used to evaluate how similar the evoked retinal response is to the target image (see section 4.2.3 for details).

Once the SGN is trained, its parameters can be integrated into the implant to run in real time (Figure 4.2 bottom). The images captured with a camera are pre-processed and sent to the SGN, which would produce effective multipolar stimulation patterns. This stage is also simulated in our work using the original retina model from step 1 (Figure 4.2 top), but with varying levels of noise to account for experimental errors.

### 4.2.2 Measurement Predictor Network

The MPN is the data-driven, implantee-specific predictor of the measured retinal responses to stimulus (Figure 4.3). It is a fully connected ANN having the stimulus pattern as input, two hidden layers, and the predicted measured retina response as output layer. In this study, we consider a square MEA with an 8-by-8 matrix of electrodes which is within the range of existing retinal implants (Luo and da Cruz, 2016; Rosenfeld et al., 2020; Fernández et al., 2021; Petoe et al., 2021). Thus, the input and output layers are vectors with 64 neurons. The two hidden layers have 128 neurons each. We used the Rectified linear unit (ReLU) activation function in the hidden and output layers. This architecture was optimized empirically with the aim of obtaining a small but effective network. We used the open-source Keras API (version 2.10.0) to build and train the ANNs, running on a i7-1165G7 processor at 2.80 GHz.

The MPN is trained using random multipolar stimulation patterns. The amplitudes of the stimuli are randomly sampled from a Gaussian distribution with mean 0 and standard deviation 0.3 to limit the stimulation amplitudes between 1 and +1. For the sake of simplicity, in this simulation study, the units of stimulation amplitude are normalized, where 0 corresponds to no current, and +1 or -1 corresponds to currents that elicit strong retinal activation. This simplification has no implication on the applicability of the MPN since, in a real clinical setting, the stimulation units (for example, mA) would still need to be scaled accordingly, as the input/output of deep learning models are typically normalized. The ground truth data is the measured retinal responses to stimulation, expressed in terms of, for example, the evoked compound action potentials amplitude (ECAPs; see section 4.2.4 for details). In a clinical setting, these are obtained by recording the retinal activation after each stimulus, requiring the use of an implant with neural recording capabilities. Here, we simulate the retinal response to stimulus as well as the recording performed with the MEA (see section 4.2.4 for details). To emulate a clinical setting, the ground truth data in our simulated

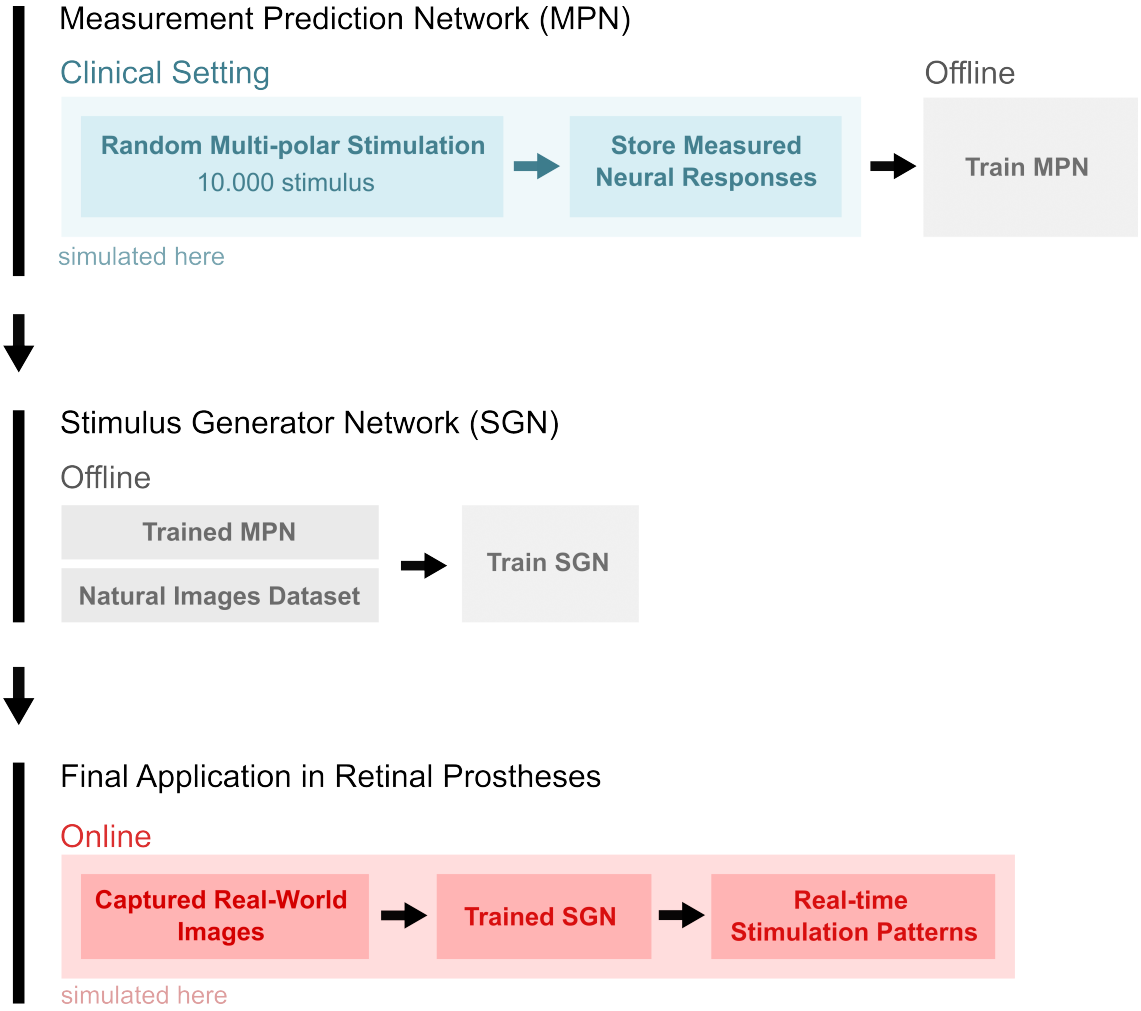



Figure 4.2: **Proposed deep learning-based stimulation protocol for visual prostheses .** The protocol has three stages: (1) training the measurement predictor network (MPN) based on the stored implantee's retinal response to random multi-polar stimuli; (2) training the stimulus generator network (SGN) using the trained MPN and a large dataset of natural images; (3) implementing the trained SGN in the implant for real-time generation of stimulus patterns based on the real-world images captured with a camera.

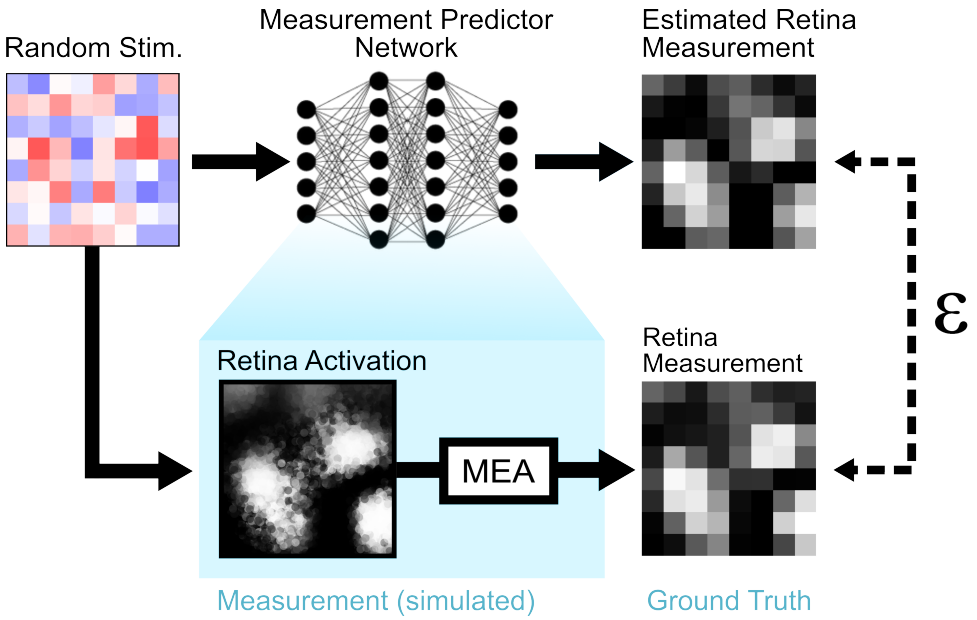



Figure 4.3: **Training of the Measurement Predictor Network (MPN) .** The MPN is a decoder that predicts the measurement of the retinal response to a given stimulus. It is meant to be a surrogate of the system Retina-MEA (blue box). The MPN is trained with a dataset composed of multiple random stimulus patterns (input data, represented by a coloured grid of electrodes) and the associated retinal response measured with the MEA (ground truth data). The error between the estimated retina measurement (top right) and the actual measurement (bottom right) is used to train the MPN. Here, both the stimulation and the recording are performed with the implanted MEA (simulated in this study) but could also be done with separate stimulating and recording units.

environment is not the activity of the individual retinal cells, but rather the responses recorded with the MEA as well (see section 4.2.4). Just like with the stimulation units, we bound the retinal activity of each recording channel between 0 (no activity) and 1 (maximal evoked amplitude) to simplify the workflow with the ANNs.

The split between training, validation, and testing sets are 10 000, 2000 and 1000 stimulus patterns, respectively. We use the ADAM optimizer with an initial learning rate of 0.0005 (the remaining parameters were set to their default value). As loss function, we use the mean square error between the retinal response predicted by the MPN and the actual retinal response measured with the MEA (Figure 4.3).

To evaluate the robustness of the MPN, we train it using different initial conditions by changing both the seeds of the random number generator and retinal activation spread. The latter is a parameter that controls the spatial extension of the retinal activation for a given stimulus (see details in section 4.2.4).

### 4.2.3 Stimulus Generator Network

The SGN receives a target image and outputs an implantee-specific multipolar stimulation pattern. The target image provided to the SGN is down-sampled to have the resolution of the MEA because it needs to be compared with the measured retinal response or its prediction during training (Figure 4.4, test and train branches, respectively). We refer to this down-sampled target as the *measurable target*, as it refers to the target that can actually be measured with the MEA, thus distinguishing it from the original full-resolution target image captured with the camera. Here, the measurable targets are created by filtering the original images with a low-pass filter to avoid aliasing and down-sampling it to the MEA resolution.

The SGN has three hidden layers, each with 128 neurons and ReLU activation functions. The values of the output layer correspond to the stimulation amplitudes of each electrode of the MEA, thus having 64 neurons. The output layer has a hyperbolic tangent activation function to bound the stimulation amplitudes between 1 and +1. The architecture of the network is optimized empirically to create a small but effective network (optimization results not shown). Finally, we add L2 regularization to the output layer to penalize high amplitude stimulus. We tested the following values for the *lambda* parameter of the L2 regularization: 0, $10^5$, $10^4$, $10^3$, $10^2$ and $10^1$.

To train the SGN, we merge it with the trained MPN. The weights and biases of the MPN are frozen so that the only trainable parameters are those from the SGN. Stacking these two networks allows the training of the SGN to be performed offline using a large dataset of natural images since the MPN predicts how the implantee's retina responds to the stimulus outputted by the SGN (black branch in Figure 4.4). For this study, we use the CIFAR 100 dataset, which contains 60 000 32-by-32 images (Krizhevsky et al., 2009). The split between training, validation, and testing sets is 37 500, 12 500, and 10 000 images, respectively. The training is performed for 300 epochs using the ADAM optimizer with an initial learning rate of 0.0005. As the loss function, we use the Mean-Squared Error between the measurable target and the predicted measurement outputted by the MPN. The training is repeated five times for each condition to evaluate the training convergence. After the training is complete, the best SGN is selected and tested on the actual retina (here simulated). The measured retinal responses are then compared with the measurable targets of the test set (red branch in Figure 4.4).

### 4.2.4 Model of Retina Stimulation and Recording

The proposed stimulation protocol is tested in a simulated environment composed of $Nr$ retinal cells evenly distributed in a 2D surface, overlaid with a matrix of Ne stimulating/recording electrodes. In our simulations, we consider $Nr = 10,000$ and $Ne = 64$. The cells are placed on a square grid and randomly jittered to create a uniform, yet random, retina. For mathematical convenience, the grid has dimensions 32-by-32 (unitless) to match the 32-by-32 image dataset. The response rate of the retina to electrical stimuli, $\vec{r}$, is simulated using a previously presented mathematical

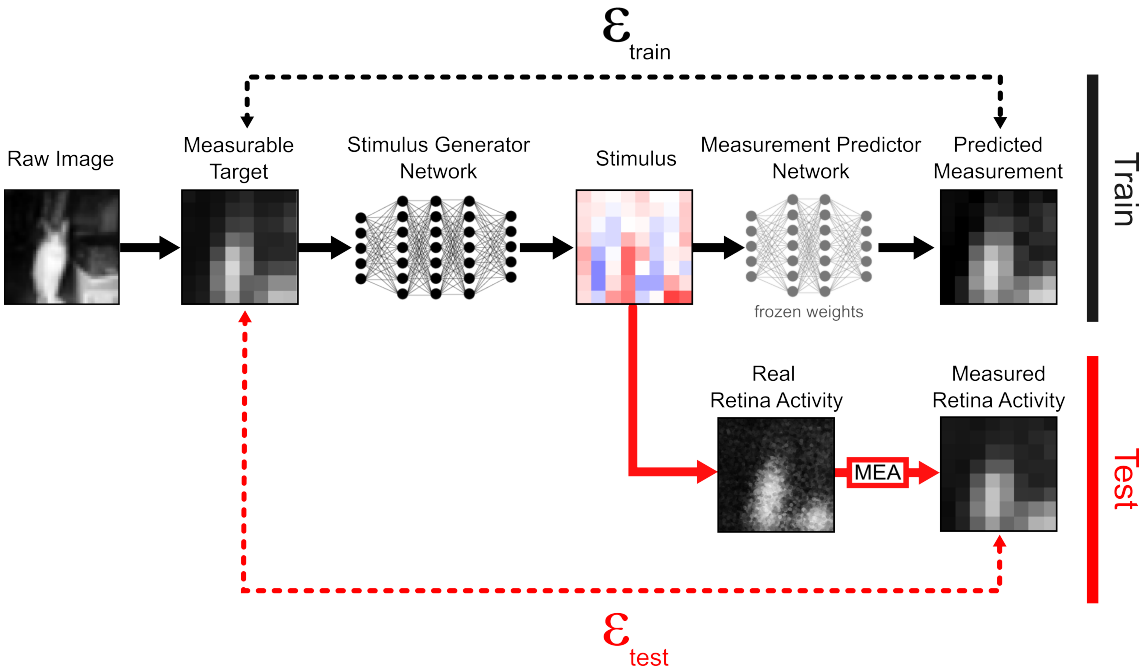



Figure 4.4: **Training and testing of the Stimulus Generator Network (SGN) .** The SGN receives a subsampled version of the raw captured image as input—the measurable target; this is the desired retinal response that we want to measure with the MEAs. The SGN outputs a stimulation pattern that should elicit the desired retinal response. To train the SGN, we use the implantee-specific predictor of retinal response, the measurement predictor network (MPN), and a large dataset of natural images (black branch). The error between the predicted measurement and the measurable target ($\varepsilon_{train}$) is used for training. Once the SGN is trained, it can be tested in the real retinal implant (red branch, here simulated). The efficacy of the SGN is accessed at the test stage by comparing the measured retinal activity with the measurable target, obtaining the measurable error ($\varepsilon_{test}$).

model that describes the short-latency response (<5 ms) of individual retinal cells to simultaneous multipolar electrical stimulation based on *in vitro* experiments (Maturana et al., 2016). The value of $\vec{r}$ can be interpreted as the amplitude of ECAPs, which we are normalizing between 0 (no response) and 1 (maximal response amplitude) for the sake of simplicity and ease of integration with the ANNs (which would also require data normalization in a real scenario). The model is composed of a linear and a non-linear part (Equation 4.1). The linear part is referred to as the Generator Potential, $\vec{G}$, where each element of this vector may be interpreted biophysically as relating to the current crossing the membrane of a given retinal cell as a result of the linear sum of the multi-electrode stimulation (Esler et al., 2018). The contribution of each electrode to the activation of a neuron is given by its electrical receptive field (ERF). The ERFs of all retinal cells are stored as rows of the matrix W with dimensions $NrNe$. The Generator Potential, $\vec{G}$, of the entire retina is then given by the dot product of $W$ with the vector of the applied electrode current magnitudes, $\vec{s}$.

$$\vec{r} = F(\vec{G}) = F(W\vec{s}), \tag{4.1}$$

The values of $\vec{s}$ range between –1 and +1, representing maximal anodic and cathodic electrode stimulation, respectively. The ERFs are modelled as 2D ellipsoidal Gaussians. The ERF of each cell $i$ has a different angle, $\theta[i]$, and different standard deviations, $\sigma_x[i]$ and $\sigma_y[i]$, for the two axes of the ellipsoids (Figure 4.5.A). The ellipsoidal ERFs are compiled in $W$ according to

$$\begin{aligned} W[i,j] = (1+\eta[i])\exp\big(&-a[i](x_r[i]-x_e[j])^2 \\ &-2b[i](x_r[i]-x_e[j])(y_r[i]-y_e[j]) \\ &-c[i](y_r[i]-y_e[j])^2, \end{aligned} \tag{4.2}$$

where $W[i,j]$ is a weighting factor describing the influence of current from the electrode $j$ on the retinal cell $i$, $\eta[i]$ is a factor relating the amplitude of response of neuron $i$, here a random number uniformly distributed between –1 and 1, to account for heterogeneities of the retina. $x_r[i]$ and $y_r[i]$ are the coordinates of the cell $i$ and $x_e[j]$ and $y_e[j]$ are the coordinates of electrode $j$. The values of $a$, $b$, and $c$ are

$$a[i] = \frac{cos^2\theta[i]}{2\sigma_x[i]^2} + \frac{sin^2\theta[i]}{2\sigma_y[i]^2} \tag{4.3}$$

$$b[i] = \frac{sin2\theta[i]}{4\sigma_x[i]^2} - \frac{sin2\theta[i]}{4\sigma_y[i]^2} \tag{4.4}$$

$$c[i] = \frac{sin^2\theta[i]}{2\sigma_x[i]^2} + \frac{cos^2\theta[i]}{2\sigma_y[i]^2} \tag{4.5}$$

The values of $\sigma_x[i]$ and $\sigma_y[i]$ determine the spatial extension of the ERFs. Each $\sigma$ is obtained using

$$\sigma_x[i] = l\lambda_x[i] \tag{4.6}$$

$$\sigma_y[i] = l\lambda_y[i] \tag{4.7}$$

where $l$ is the distance between electrodes and $\lambda$ is a dimensionless parameter that determines the extension of the ERF with regards to the electrodes. Large values of $\lambda$ create large ERFs, meaning that the retinal cells are affected by stimuli that are several electrodes away. The values of $\lambda_x[i]$ and $\lambda_y[i]$ for each cell $i$ are uniformly sampled from the interval $\Lambda \pm 25\%$ where $\Lambda$, also referred to as spread, modulates the current overlap of neighbouring electrodes. It can also be seen as a parameter that scales the dimensions of the MEA with regards to the retina (a high-density MEA would be in a high spread regime, for example). We perform simulations under different spread regimes, with $\lambda$ set to 1.0, 1.5, 2.0, 2.5, and 3.0 electrode pitches.

The non-linear part of the model converts the Generator Potential into retinal activation. This non-linear Transfer Function can be simplified to a double sigmoid function, as previously described in our in vitro study (Maturana et al., 2016). Here, we assume that the double sigmoid is symmetric,

$$F(G) = \frac{1}{(1 + \exp(-k(|-z)))}. \tag{4.8}$$

Each cell has a different Transfer Function (Figure 4.5.B), obtained by uniformly sampling $k$ and $z$ from the intervals [5, 10] and [0.4, 0.6], respectively. Implementing the realistic linear/non-linear model under different spread regimes allows us to evaluate how our stimulation protocol performs in diverse conditions (Figure 4.5.C). Elongated phosphenes can also be created by forcing $\sigma_x$ and $\sigma_y$ to be different and imposing a limited range on the angle $\theta$ (Figure 4.9). We also model the recording of the retinal activity with the MEAs. For that, we assume that the electrodes have a diameter equal to half the electrode pitch (centre to centre) and that each electrode's read-out is proportional to the average activity of all the cells within its radius (Figure 4.5.D and 4.5.E).

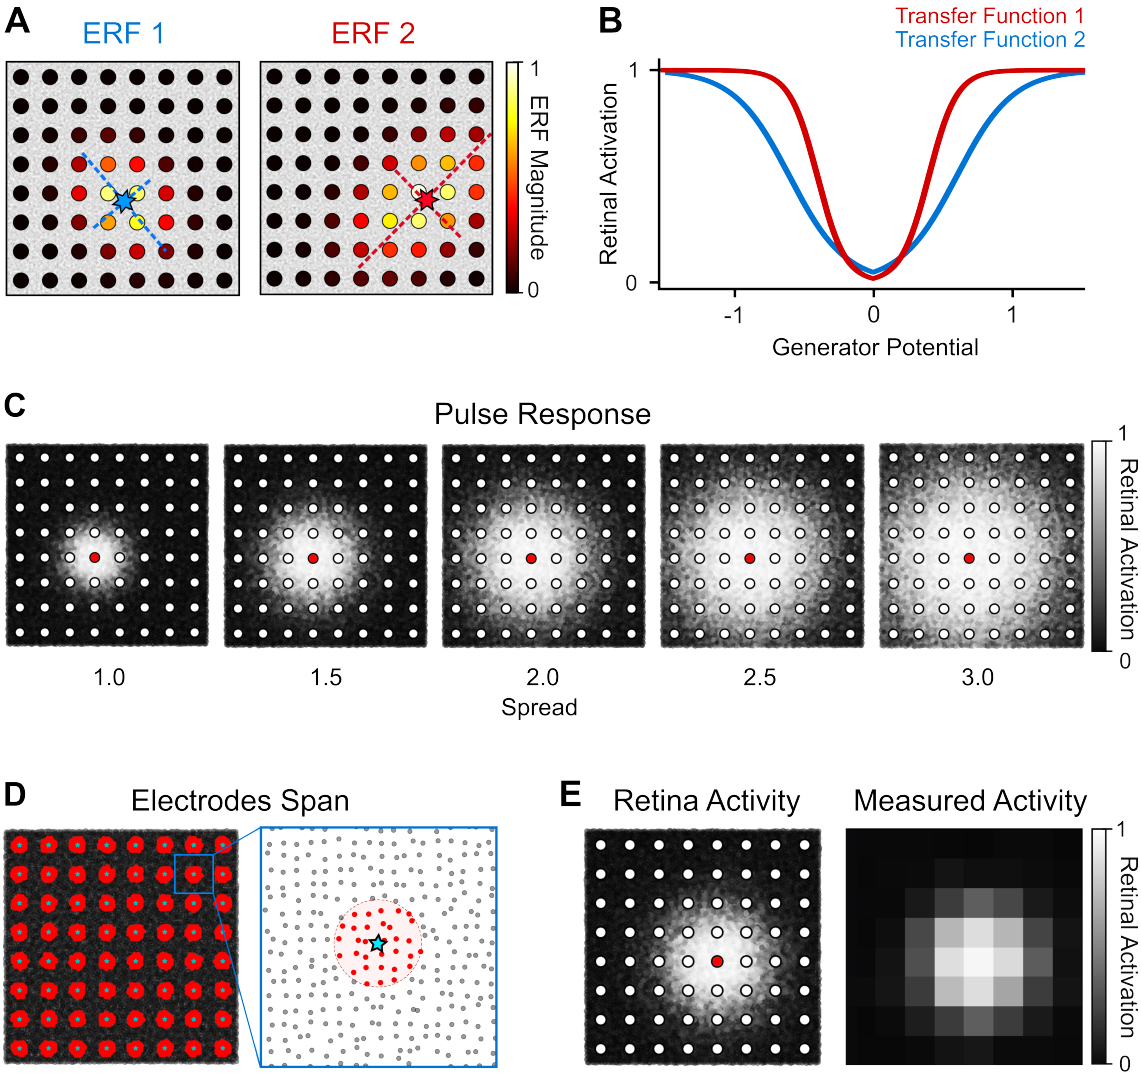



Figure 4.5: **Retina model and MEA recording.**A) Representative depiction of the ERFs of two different retinal cells (blue and red stars). B) Representative depiction of two different transfer functions ($k_1$=5, $z_1$=0.6; $k_2$=10, $z_2$=0.4). C) Retinal responses to a stimulus of amplitude 1 on the red electrode, for the different spread regimes considered in this study. D) Retinal cells (red) recorded by each electrode (blue star). E) Retina activation that results from a stimulus of amplitude 1 applied in the red electrode (left) and the resulting MEA read-out (right).

## 4.3 Results

### 4.3.1 The Measurement Predictor Network accurately estimates retinal responses to simulated multipolar stimulation

The first step of our stimulation protocol consists of obtaining a implantee-specific predictor of the retinal response to multipolar stimulation (Figure 4.2 top). The dataset used to train the MPN was obtained by gathering the responses to 10 000 random multipolar stimuli, here using the simulated retina. We trained MPNs under different spread regimes, to account for different experimental conditions in a clinical setting. To check the consistency of the results for each spread, we replicated each condition five times, each with a different parametrization of the MPN and retinal model (the random number generator was initialized with a different seed for each repetition). The training converged smoothly for all conditions.

Once trained, we submitted the MPNs to 1000 test stimuli. We used the root mean square error (RMSE) to evaluate the performance of the network in predicting the measured retinal response. The prediction error was similar for all spreads and consistent across the multiple initializations of the same spread (Figure 4.6). The RMSE was typically below 0.06 [a.u.], whereas a random predictor would have an error of 0.45 ± 0.04 [a.u.] ($n$ = 1,000 test stimuli). The MPN results were qualitatively very similar to the measurement of the simulated MEA (Figures 4.6 B and C). This demonstrates that the ANNs accurately predicted the MEA measurement by looking only at the stimulus pattern. This is a necessary requirement if the MPN is to be used for training the SGN. The MPN was efficient for all levels of spread (Figure 4.6). This is an important observation since this variable cannot be controlled in a real scenario as it is intrinsic to the retina-implant system.

### 4.3.2 The Stimulus Generator Network produces effective stimulation patterns

The SGN was trained offline using the previously trained MPN and a large dataset of natural images, as explained in section 4.2.3. Each MPN-SGN pair was trained within the same simulation, meaning that the spread regimes and initializations used for the MPN were also applied for the SGN. After training, the SGN was tested with 10 000 test images by applying the inferred stimulus patterns to the original retina model (the one used to train the MPN). This corresponds to the last step of the protocol in Figure 4.2 (online testing) and the red branch in Figure 4.4. The retinal response to the SGN stimuli were measured with the MEA and compared with the measurable targets (red section in Figure 4.4), obtaining the *Measurable Error*. The measurable error is based on the response measured with the MEAs so it can also be calculated in a clinical setting (as opposed to the error between the original full resolution image and the activation of individual ganglion cells, which can only be compared in a simulation). For each condition (spread regime and random initialization), we tested six levels of L2 regularization (Figure 4.7). Applying this regularization on the SGN output has two effects: it explicitly forces the network to choose lower stimulation amplitudes, which is important in a clinical application, and it implicitly smooths a

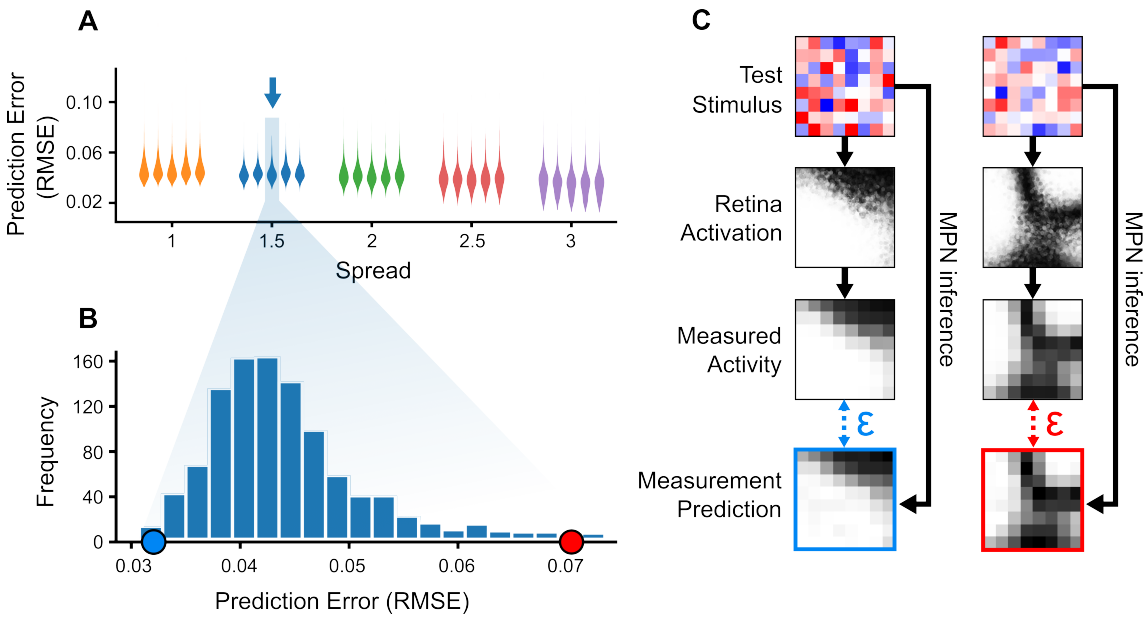



Figure 4.6: **The Measurement Predictor Network (MPN) predicts the stimulus response in diverse spread conditions.** (A) Measurement prediction error in the test set for MPNs trained under different spread regimes. For each regime, the experiments were repeated five times with different initializations of the MPN and retina model, hence there are five violin plots for each spread value. (B) Measurement prediction errors of the test set for a retina with spread of 1.5 (arrow in A). (C) Example of two test stimuli and the associated retina activations, measured activities and measurement predictions, extracted from the tails of the histogram in B, associated with low and high prediction error (blue and red dots in B, respectively). This shows that the MPN predictions are qualitatively very similar to the actual measurements.

loss function, making the training less prone to land in sharp local minimum, which improves the generalization of the network.

We can see that, as the relative magnitude of the L2 regularization increases, the results became more consistent across the multiple initializations for all the spreads. This is due to the smoothing of the loss function. On the other hand, large regularizations induced lower stimulus amplitudes, which led to less effective stimulation patterns (high errors for lambda equal to $10^2$ and $10^1$). We found a good compromise between these two effects for a regularization lambda of $10^3$. Yet, for smaller spreads (1 and 1.5), the SGN may still become stuck in a local minimum even after 300 training epochs (Figures 4.7 A and B). This means that, in a real application where the actual spread in not known - since it is an intrinsic variable of the retina-implant system - the SGN would need to be trained several times with different initializations and the instance with the best performance would be implemented in the retinal prosthesis. Analogously, here we considered the results obtained with the best SGN instance for a spread of 1.5, as an example. The training loss curves show that only one out of five instantiations of the SGN became stuck in a high-error local minimum (Figure 4.7.B). The best SGN instance for this spread produced results that were qualitatively very similar to the original image for the entire spectrum of test errors (Figures 4.7 C and D). Finally, larger activation spreads increased the training convergence (smaller dispersion of points in Figure 4.7.A). This suggests that the loss function becomes smoother as the overlap between electrodes increases.

### 4.3.3 The ANN method matches the accuracy of the AMI approach for a LNL model retina

We then compared our ANN-based protocol with two other stimulation strategies: the conventional stimulation strategy and the AMI. In the conventional method, the stimulation magnitude of a given electrode is proportional to the brightness of the image in the location underlying the electrode. The stimulus is typically applied one electrode at a time (sequential stimulation) but fast enough so that the implantee perceives a single image. Here, since we are not considering the temporal dimension, we simplified this approach by applying the stimulus in all electrodes simultaneously and assuming that the visual experience would be identical. In these tests, we considered a retina with a spread of 1.5 (see Figure 4.10 with the results obtained with different spread values). For the conventional method, we tested 10 different scaling factors, ranging from 0.02 to 0.2 in steps of 0.02 (Figure 4.11) and selected the most efficient one. The AMI strategy estimates the inverse retinal model analytically assuming a LNL response and finds optimal NAS solutions, as presented by Spencer et al. (2019). The forward model estimation for AMI is based on the method presented by Maturana et al. (2016). Thus, since we are actually using a LNL model for each retinal cell in this study, the results of the AMI strategy serve as a ground truth for our ANN method. To ensure a fair comparison between the ANN and AMI methods, we regularized AMI using quadratic programming to constrain stimulation amplitudes as in Spencer et al. (2019),

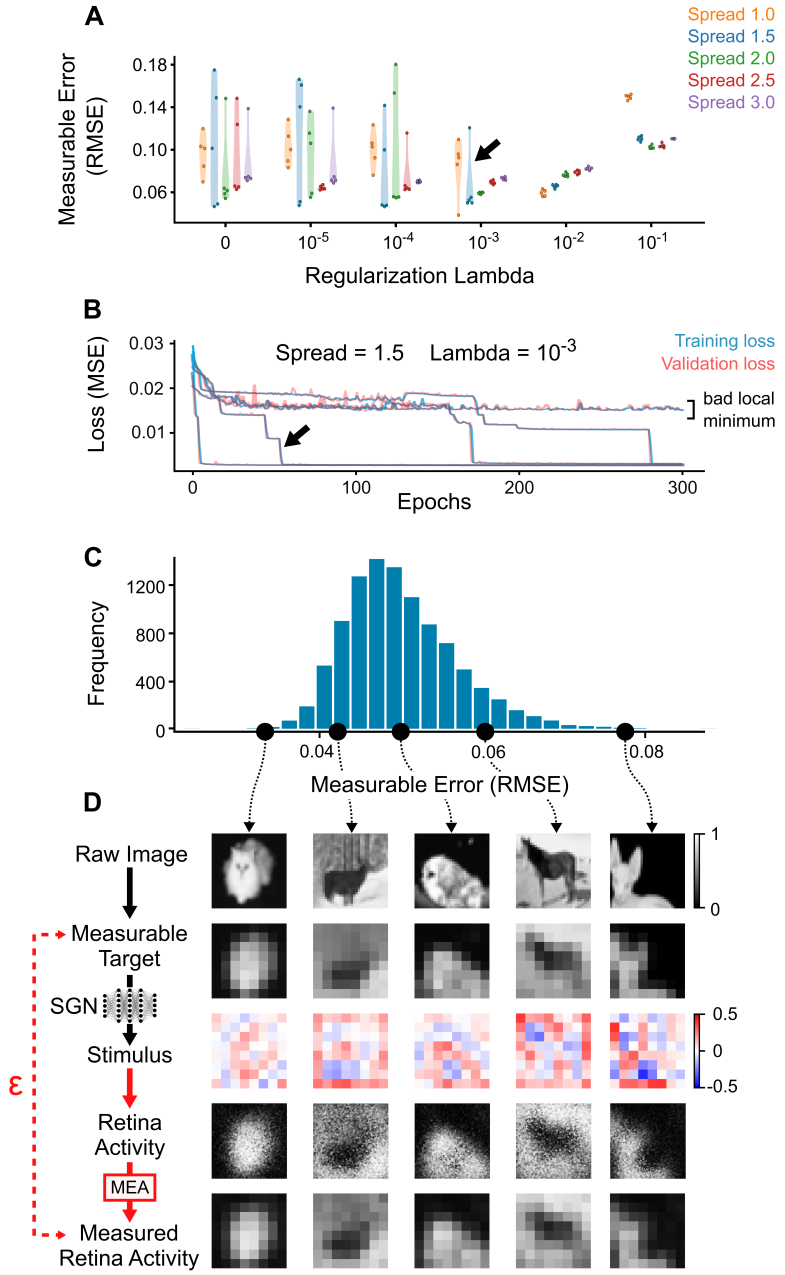



Figure 4.7: **The stimulus generator network (SGN) encodes efficient multipolar stimulation patterns.** A) Test results for all the conditions simulated, where each dot represents the average error of 10,000 test stimuli for a given spread, regularization lambda, and random initialization. B) Training loss curves for the five SGNs instantiations at a spread of 1.5 and regularization lambda of $10^{-3}$ (black arrow in A). C) Histogram of the error for the 10,000 test images for the best SGN instantiation at a spread of 1.5 and lambda of $10^{-3}$ (whose training is marked with a black arrow in B). D) Qualitative assessment of the results across the spectrum of errors, showing the original raw image, the measurable target that the SGN aims to reproduce, the stimulus pattern produced by the SGN, the retina activation (unattainable in a real scenario), and the measured retinal activity. The measurable error represented in B corresponds to the error $\varepsilon$ between the measured retina activity and the measurable target.

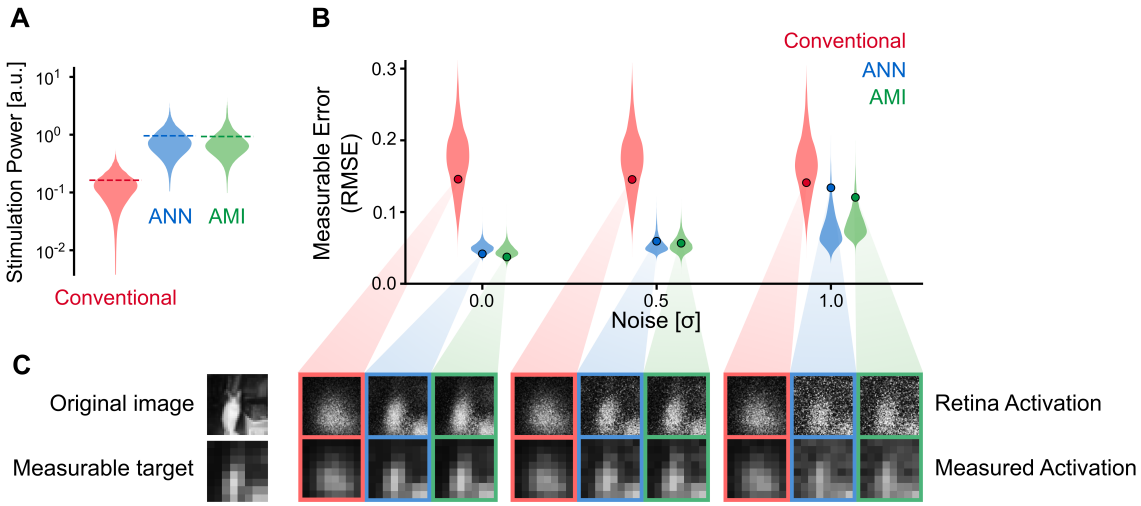



Figure 4.8: **Comparison between the conventional stimulation method and neural activity shaping (NAS) performed with the artificial neural networks (ANN) protocol and analytic model inversion (AMI).** (A) Power consumption of the three stimulation methods for the entire test set. AMI was parameterized to have the same power consumption as the ANN method. The power of a given stimulus corresponds to the sum of the squared electrodes amplitude of the stimulation pattern. Here, the three methods were tested for a retina with spread of 1.5 (as in Figures 4.7 C and D). (B) Measurable error based on the MEA recordings at the test stage for the three stimulation methods. To account for experimental imprecisions that would occur in a clinical setting, we added noise to the parameters of the retina model in the test phase. The noise added to the parameters of each retinal cell was extracted from normal distributions centred in 0, with standard deviations of either 0 (no noise), 0.5 $\sigma$ or 1.0 $\sigma$, where $\sigma$ corresponds to the standard deviation of the values of the considered parameter across all the cells of the retina used to train the measurement predictor network. The coloured points depict the error of the test image shown in C. (C) Qualitative representation of the retina activation (top) and measured retinal activation through the MEA using the three stimulation methods (conventional in red, ANN in blue, AMI in green) for a given image (top, left) and the associated measurable target (bottom, left). Statistical tests for A and B are not shown because the large sample size ($n$ = 10 000) causes even negligible differences to appear significant.

so that the overall power consumption of AMI was comparable with the ANN method (Figure 4.8.A).

We tested the three methods under varying noise conditions by making the test retina slightly different from the retina used for training (Figure 4.8.B). The results show that, for the three noise conditions, the conventional method clearly gives the lowest accuracy of the three methods. The ANN and AMI methods obtain very similar results, thus showing that the ANN-based protocol is capable of non-explicitly approximating the optimal inverted model of the retina response to stimulation. Although the differences in errors between the ANN and AMI are statistically significant due to the very high sample size (10,000 samples per group), the actual magnitude of these differences is negligible (mean RMSE differences of 0.005, 0.002 and 0.006 [a.u.], for the noise levels of 0, 0.5 and 1 , respectively). This can be further confirmed by examining how the retina

responds to a representative test image for the three stimulation methods (Figure 4.8.C), as well as examining the underlying stimulation pattens, which are very similar for the ANN and AMI strategies (Figure 4.12). It is also clear that the ANN and AMI results get worse as the noise in the test retina increases, especially for a high contrast image as the one shown in Figure 4.8. The same does not happen for the conventional method because the broad retinal activations make it more robust to noise.

Despite providing comparable stimulation patterns, the ANN and AMI methods differ significantly in terms of computational efficiency. The inference time of a trained SGN for 10 000 images running on a i7-1165G7 processor at 2.80 GHz was 23.4 ± 3.7 ms ($n$ = 10 test runs, each with 10 000 test images, using a batch size of 10 000). The AMI method, on the other hand, took on average 69.6 ± 2.0 s ($n$ = 10 test runs, each with 10 000 test images), thus being three orders of magnitude slower. That may be a significant limitation for an implantable device, which should be as small and efficient as possible.

Finally, it is important to note that the ANN is not limited to any particular model of retinal response, unlike the AMI which assumes a LNL response. This means that, in a real application, where the actual retinal response may be more complex, the presented ANN protocol should still be capable of finding an adequate solution, even if it requires changing the network architecture.

## 4.4 Discussion

Our ANN-based approach solves the problem of encoding target images into personalized multipolar stimulation patterns following a model-free NAS approach. We showed that this strategy is a great improvement over the conventional unipolar stimulation, matching the accuracy of the theoretically optimal AMI for a LNL retina model, with the advantage of not being bounded to any model type and being 3 orders of magnitude faster. The significance of this novel approach will increase when used with higher density electrode implants since, unlike the conventional method, it is not limited by the stimulus overlap, but rather takes advantage of it. As such, the method has great potential to improve the visual experience of retinal prostheses. Nonetheless, our work assumes that the neural activation pattern in the retina is closely aligned to the pattern of light perceived by the implantee. Extending the method to target perceptual patterns would require a psychophysical approach in which back-propagation can be performed based on user feedback. Early work achieved this with a "dialogue" approach (Eckmiller et al., 2005, 2007) or more recently a "Human in the Loop" approach (Granley et al., 2023).

Image pre-processing algorithms for retinal prostheses can enhance specific parts of the image (edge detection, enhancing contrast, background removal, etc) to create simpler and more informative target images (Wang et al., 2023). As suggested by Leong et al. (2024), our method could be easily coupled downstream with such pre-processing strategies, as these are independent tasks.

Also, the overall strategy could be applied in any type of device with stimulating and recording capabilities. Still, the specific architecture of the networks used may need to be adjusted; a device with a high-density MEA would generate a more complex dataset and could thus require a bigger network, for example.

We tested our approach using simulated ground truth data, using a realistic model of retinal response to multipolar stimulation based on in vitro experiments (Maturana et al., 2016). Despite the heterogeneity of ERFs and transfer functions used, the responses elicited were approximately circular blobs, such as those obtained with suprachoroidal implants (Sinclair et al., 2016). These are also the types of implants that would benefit the most from our strategy because the larger distance between electrodes and ganglion cells leads to considerable current overlap. Irregular phosphenes shapes, such as the elongated ones often experienced by implantees, will impose no problem on the ANNs training, so long as they are fully captured by the recording electrodes (see Figure 4.9 for examples of results with elongated phosphenes). Likewise, real retinas may present higher non-linearities that are not captured by our LNL model but, as general function approximators, ANNs should be capable of handling these irregularities, even if it requires larger networks and more intensive training.

A limitation of the model considered here is that it is purely spatial and ignores the temporal interactions of the retinal response. We neglect neuronal adaptations such as Long/Short Term Potentiation/Depression, which could change the response to stimulation as a function of time. Another simplification is the assumption that the signal recorded with each electrode corresponds to the average activity of the cells over it. Also, our method does not take into account the fact that the electrical stimulus simultaneously excites cells that may encode different light information (such as ON and OFF ganglion cells), as that is not possible with existing devices. Such specificity would only be possible by recording and stimulating with single cell resolution as well as having the adequate algorithms to identify and target the different cell types (Shah et al., 2019; Madugula et al., 2022; Lotlikar et al., 2023). Finally, the mathematical model used to simulate the retinal cells was based on an in vitro study (Maturana et al., 2016) that evaluated the short-term retinal response to stimulation (<5 ms), which is attributed to the direct activation of ganglion cells (Eickenscheidt et al., 2012; Maturana et al., 2018). The long-latency ganglion responses due to the direct activation of photoreceptors (Eickenscheidt et al., 2012; Maturana et al., 2018) are not considered here since these cells are assumed to be degenerated.

Nonetheless, it is important to note that, even though we are testing this protocol in a simulated environment (where we, as developers, can access all the variables of the simulation), all the methods presented in this work were designed to be dependent only on the information accessible through the implant; i.e., the delivered stimulus patterns and the measured retinal responses. This ensures that the algorithms can be directly applied in a clinical setting. One aspect worth exploring in the future is choosing a more suitable error metric for the loss function of the ANNs. Here, we used a standard MSE, but we speculate that metrics sensitive to the information content of the image

(such as Structural Image Index (Wang et al., 2003)) may lead to better outcomes. Also, despite using multipolar stimulus, the final values of the generator potentials created with our stimulation method are always unipolar; i.e., the values of $\vec{G}$ are only positive (or negative) for all the retinal cells and, thus, land only on one side of their Transfer Function (Figure 4.5.B). Actually, the SGN instances that become stuck in an inefficient local minimum during training (high error curve in Figure 4.7.B) were the ones leading to bipolar generator potentials and, therefore, using both sides of the non-linear Transfer Function of retinal cells. When that happens, there are artifacts in the transitions between positive and negative $\vec{G}$ regions (Figure 4.13). To deal with this, we trained the SGN multiple times for each condition and chose the most efficacious SGN. We suggest that a similar approach should be taken in a clinical setting, since the prevalence of these artifacts cannot be determined a priori (it is a function of, at least, the activation spread and regularization lambda, as shown in Figure 4.7.A). Nonetheless, a previous study showed that it can be highly beneficial to exploit both sides of the non-linearity of the Transfer Function (Spencer et al., 2021) (Figure 4.13). We suggest that future work could explore how to force the network to use both sides of the non-linearity in order to create higher contrast retinal activations.

We did not show how the number of electrodes affects the stimulation performance (each different MEA configuration would require its own MPN-SGN architectures and parametrizations, which would need to be individually optimized). However, naturally, the higher the number of electrodes, the higher the spatial resolution attainable with the stimulation. Furthermore, as mentioned before, higher electrode density leads to a larger overlap between electrodes, which not only makes our approach more relevant but also more stable, since larger overlaps make the training converge better (even though, for the same number of electrodes, smaller the spreads lead the better image resolution, as shown in Figure 4.10). Also, our method requires a device with independent current sources for each electrode (Tran et al., 2014) or an algorithm that maps appropriately selected sub-group of electrodes to the available stimulus generators in real-time.

In a clinical setting, the networks can be retrained periodically to adapt to the long-term changes of the retinal response. The only requirement is that a subset of the real target images, stimulus patterns and retinal responses are stored in the device. These could be periodically downloaded for retraining the MPN and SGN as before, but now using the actual captured images (instead of random stimuli for the MPN and an arbitrary dataset for the SGN). Another solution would be to do this retraining online, where the weights of the SGN would be changed in real time in a full closed-loop fashion. This way, there would be no need for the MPN during retraining, since the SGN could use the responses measured in real time as ground truth data to compute the loss and update the SGN parameters. Finally, this framework is not limited to retinal implants; it may also be adapted for cortical visual protheses, so long as the transformation from target image to cortex representation is known in order to create the target activity patterns.

# 4.A Supporting Information

## Acknowledgements

This project was supported by the Australian National Health and Medical Research Council (NHMRC)—Development Grant 2014380. D C acknowledges funding from Fundação para a Ciência e a Tecnologia—scholarship SFRH/BD/143956/2019.

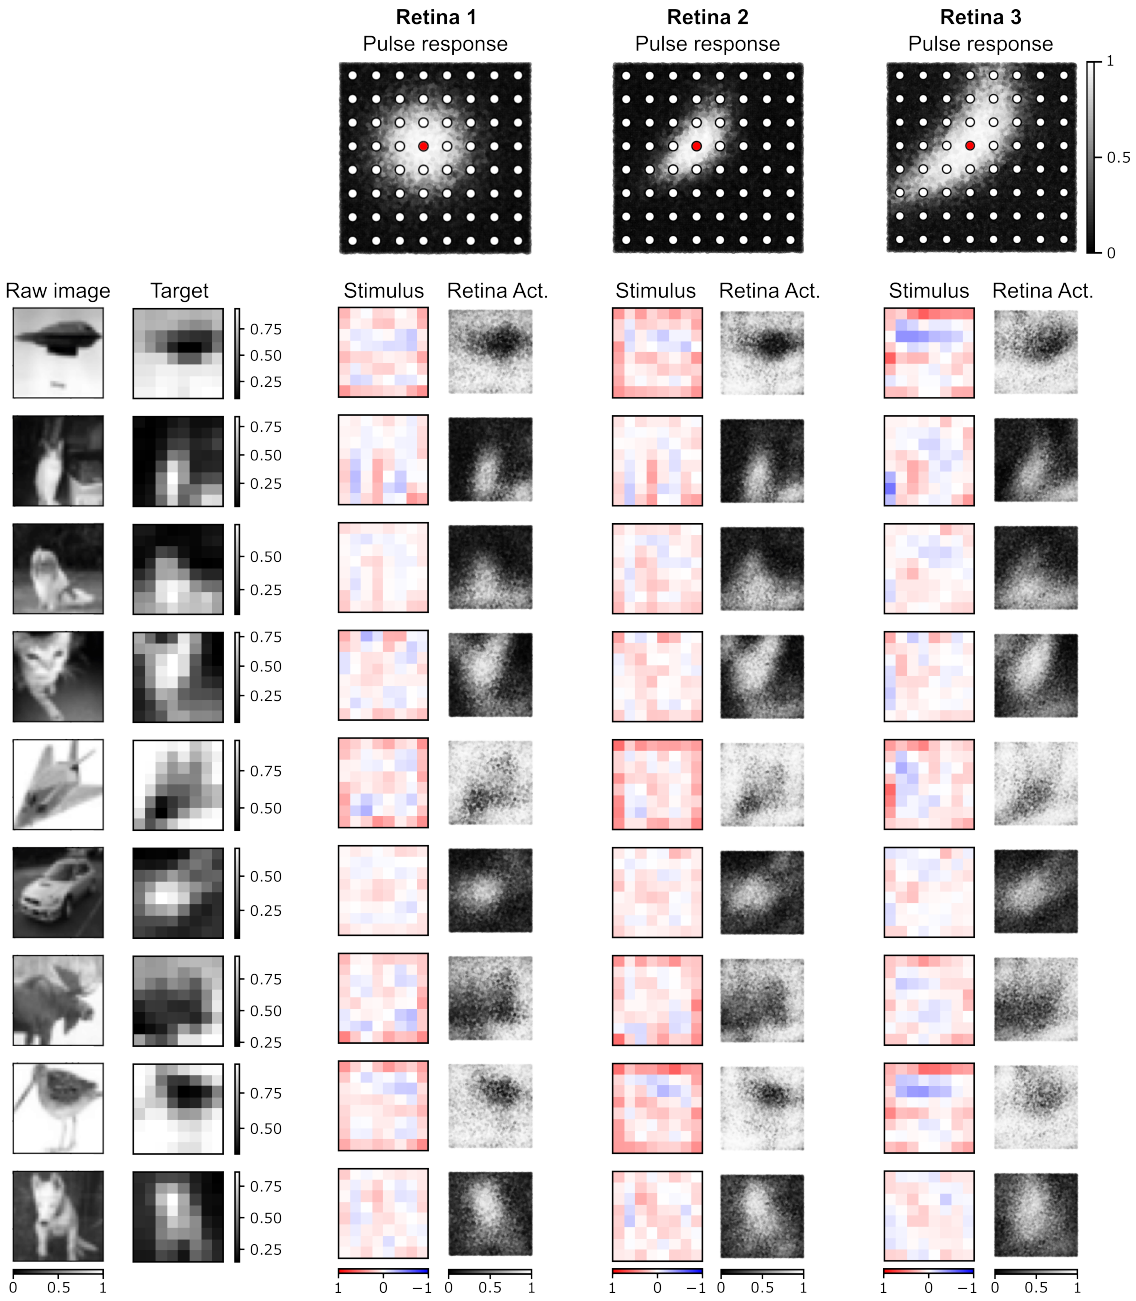



Figure 4.9: **The ANN-based approach produces efficient stimulation patterns in retinas with different response topologies.** Qualitative results obtained in three different retinas: Retina 1 – regular topology, identical to those presented in the main paper; Retina 2 – elongated phosphene topology, obtained by changing the ERF parametrizations (y sigma larger than x sigma, ERF angles confined to a smaller range, set as a function of cell coordinates, retina transfer function parameters set as a function of cell coordinates); Retina 3 – elongated phosphene topology, as Retina 2, but with larger spread. The top row shows the pulse response s of each retina to a single pulse with amplitude 1. The remaining rows show the results obtained for different test images. Each retina was trained with a different MPN/SGN. The SGN learns stimulation patterns that adjust to specific retina responses.

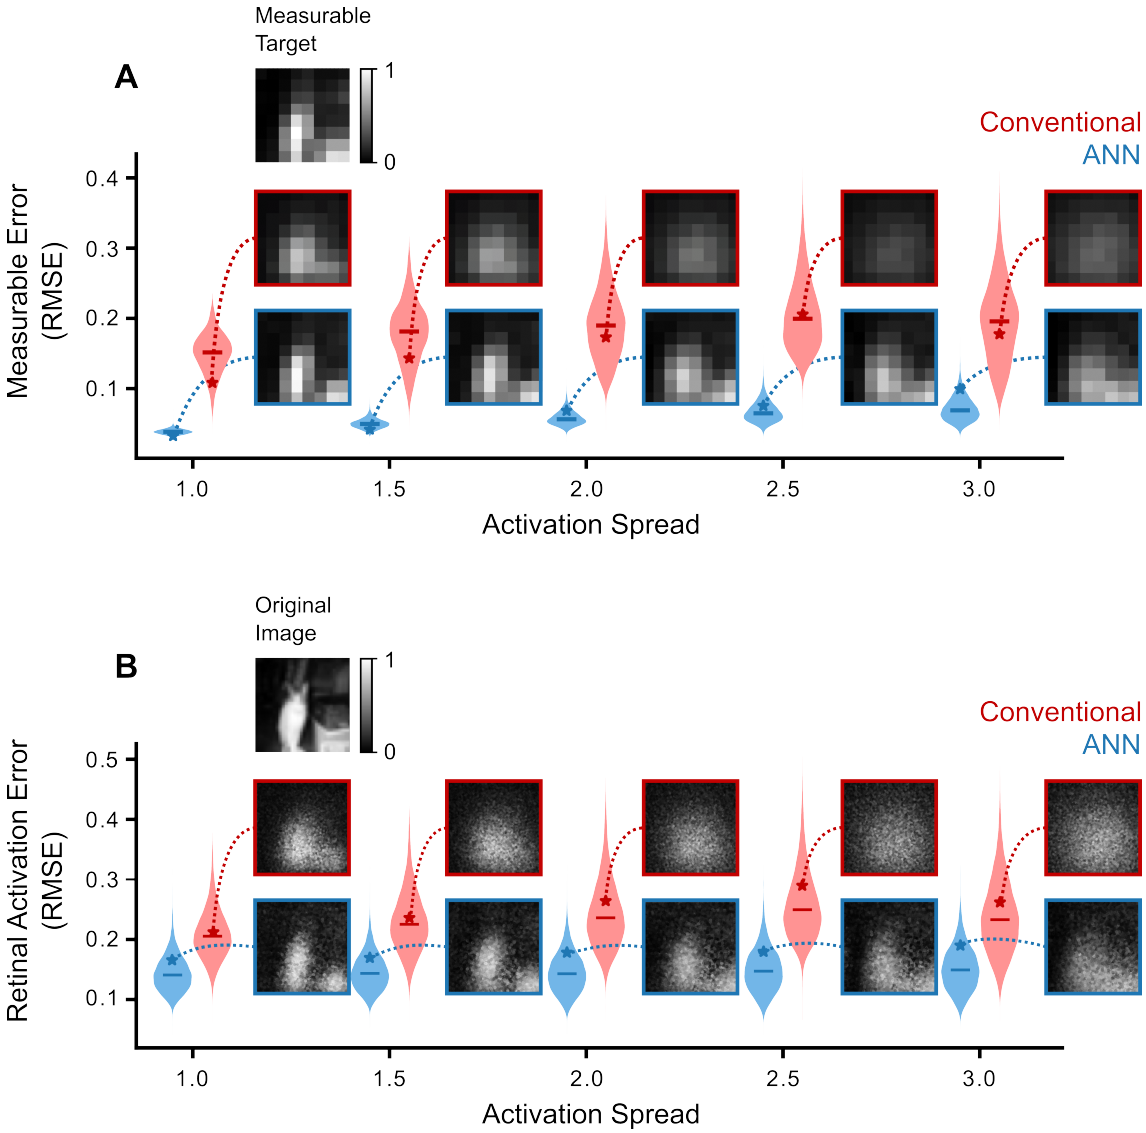



Figure 4.10: **Comparison of conventional and ANN-based stimulation protocols for different spread regimes.** A) Measurable error based on the MEA recordings. For the conventional method we tested different scaling factors for each spread and chose the most efficient one (see Figure 4.11). The ANN-based approach achieved significantly better results, both quantitatively and qualitatively, for all spreads. B) Retinal activation error, i.e., the error between the actual retinal activity and the original raw image. Since the original image had a 32-by-32 resolution, we approximated the activity of the 10,000 retinal cells to a 32-by-32 image and computed the RMSE between that and the original image. The qualitative results of retina activation using ANNs show a clear improvement in terms of spatial resolution when compared with the conventional method.

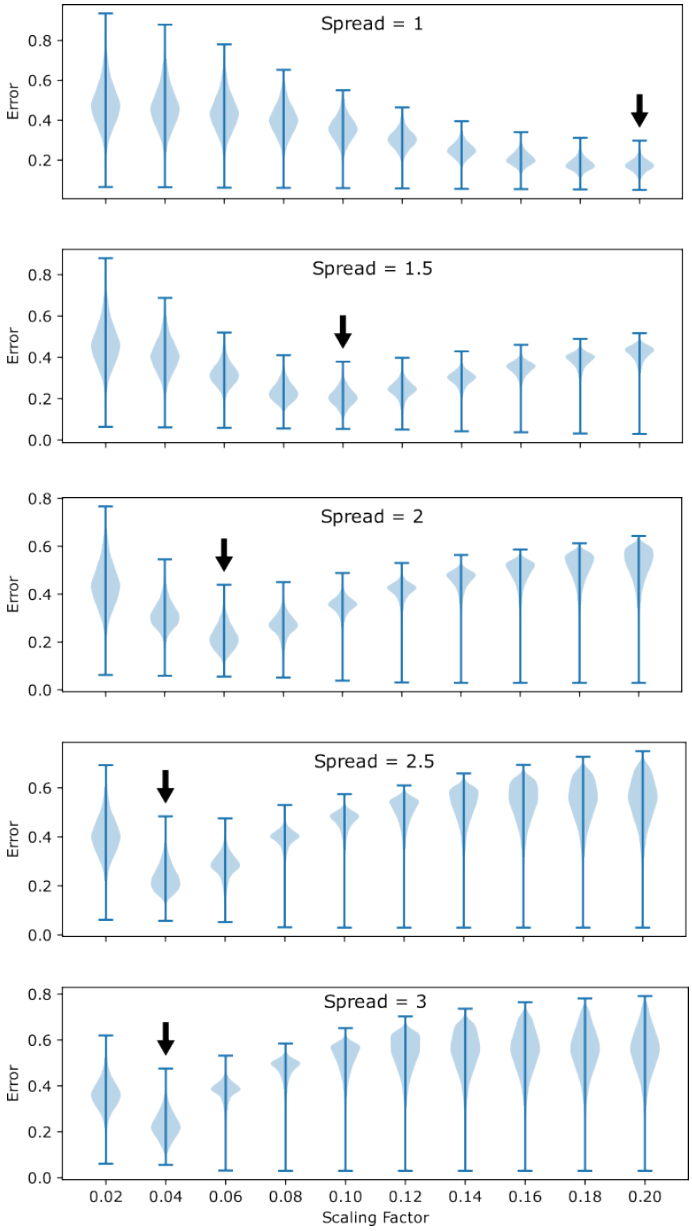



Figure 4.11: **Scaling factor selected for different spread conditions in the conventional stimulation method.** To ensure that we compared the ANN approach with the best version of the conventional method, we tested different scaling factors and selected the most efficient one for each spread. We tested each scaling in the 10,000 test images and calculated the RMSE between the retina activation and the measurable target. For each spread, we chose the scaling with lower mean RMSE (black arrow). The results of the conventional method shown in Figure 8 were based on the parametrization for a spread of 1.5.

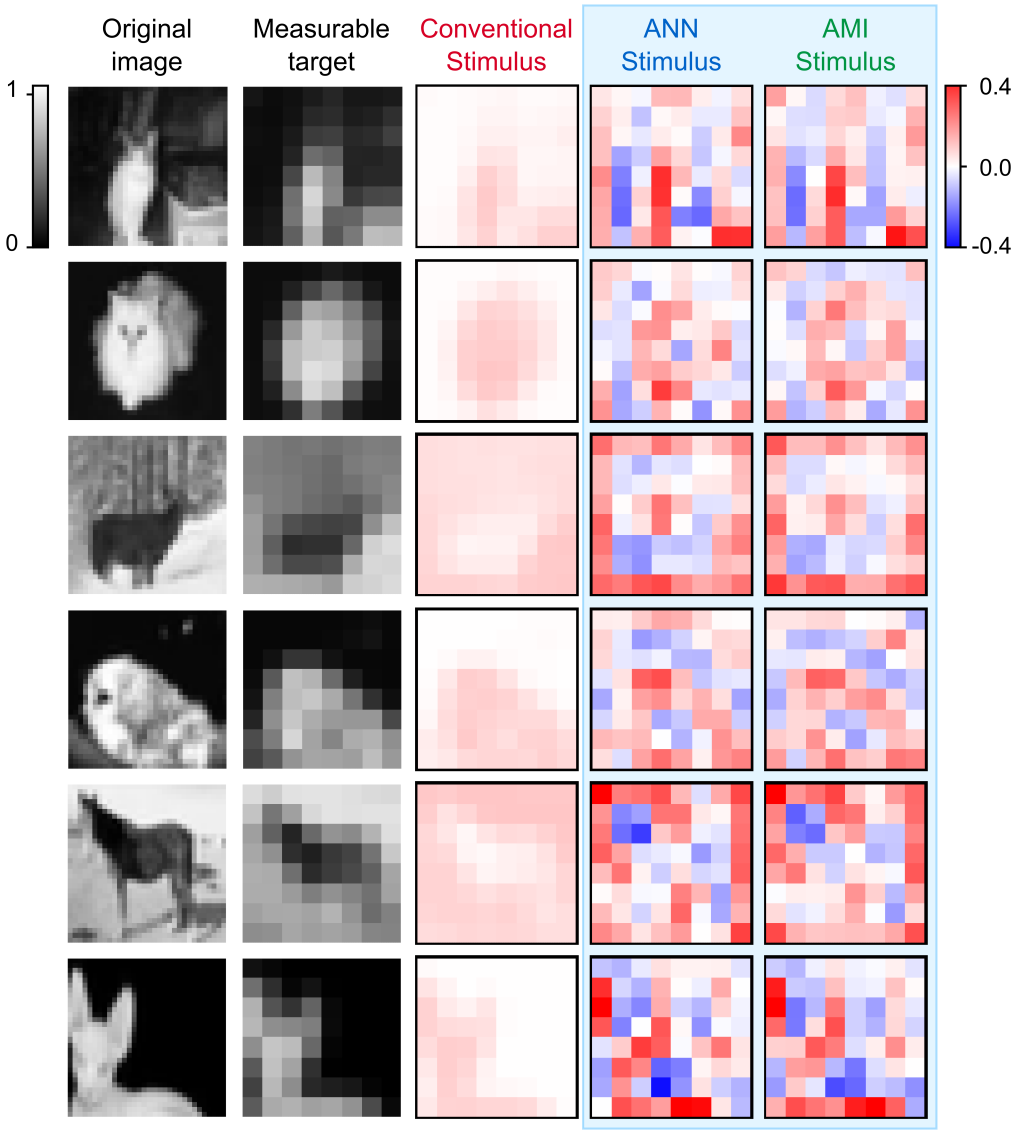



Figure 4.12: **The ANN method approximates the solutions provided by Analytical Model Inversion.** The stimulation patterns obtained with our ANN method are very similar to those obtained with AMI (left columns within the blue rectangle). This means that the presented ANN protocol can efficiently estimate the forward and inverse retinal model without imposing any assumptions about the retinal response (unlike the AMI which assumes a LNL model). We also show the solutions provided by the conventional method, which are just a direct translation of the target image into a stimulation pattern (middle column).

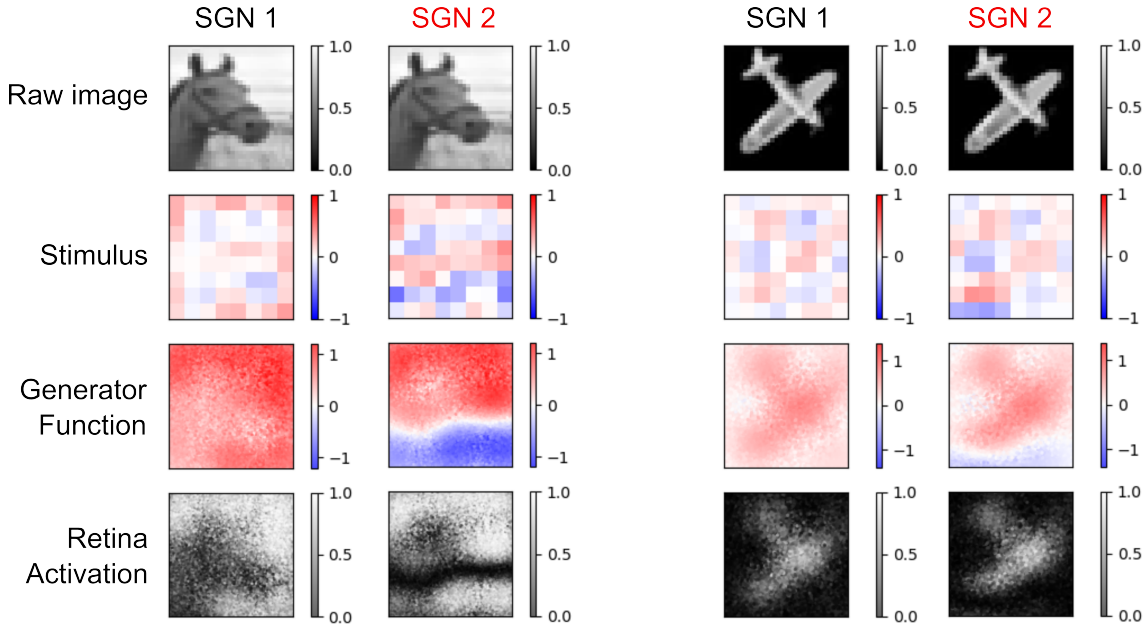



Figure 4.13: **The SGNs that become stuck in local minima during training induce bipolar generator potentials $\vec{G}$ that may lead to artifacts in the retina activation.** SGN 1 was the result of a successful training, thus generating unipolar $\vec{G}$ that lead to the expected retina activations. SGN 2 became stuck in a local minimum during training (worst performance in Figure 4.7.B). The higher training/validation errors obtained with this network are due to the artifact caused by the transition between regions of positive and negative $\vec{G}$. That is, in one region, the retinal cells are being activated due to an overall negative current, whereas in the other region they are being activated due an overall positive current (since the retinal transfer function is a double sigmoid with regards to $\vec{G}$ both positive and negative $\vec{G}$ values lead to retinal activation). In the transition between regions, the current is null, and thus there is no retina activation that leads to a dark artifact (example of the left). However, in some situations, this transition region may not affect the results or even be beneficial (example on the right). A network that learns how to effectively exploit these "artifacts" could be used to create sharp regions of no retinal activation, as proposed by Spencer et al. (2019)

# Chapter 5

# Envisioning the next generation of neuroprosthesis with Neuromorphic Engineering

The neuromodulation methods presented in the previous chapters were based on algorithms that explicitly dictate how to stimulate the monitored neuronal system. These algorithms were implemented in standard programming languages and ran on conventional CPUs (in a real implant, these would be running on the implanted microprocessor). However, traditional chips are approaching the limits of size and computational efficiency, and these two requirements are fundamental for implantable devices meant to process neuronal data in real-time and run powerful closed-loop algorithms. In the current chips, the processing occurs in a different physical location from where the memory and functions are stored – von Neumann architecture. This forces the data to be moved from memory to CPU for computation which significantly compromises speed and imposes an efficiency barrier commonly known as the *von Neumann bottleneck* (Zou et al., 2021). Also, the size of transistors is approaching its theoretical limit meaning that more computational power requires bigger processors. So, future generations of implantable neuromodulators may require a paradigm shift in the way we see computing devices. So, as our final research question, we asked: *How can intelligent stimulation protocols run efficiently in future neuroprosthesis?*

Beginning in the 1980s, the field of Neuromorphic Engineering proposed a biologically inspired analog computing paradigm that mimics the structure and function of neuronal systems (Ham et al., 2021). In biological neuronal systems, memory and processing are intertwined – *memory* can be seen as the synaptic strengths established within the network through learning mechanisms and *processing* would correspond to the event-driven propagation of firing patterns across the different neurons. In Neuromorphic Engineering, this principle is followed by devising trainable computing substrates where event-driven (clockless) computation is co-localized with sparse representations of the training data. An ideal neuromorphic chip can be thought of as a physical

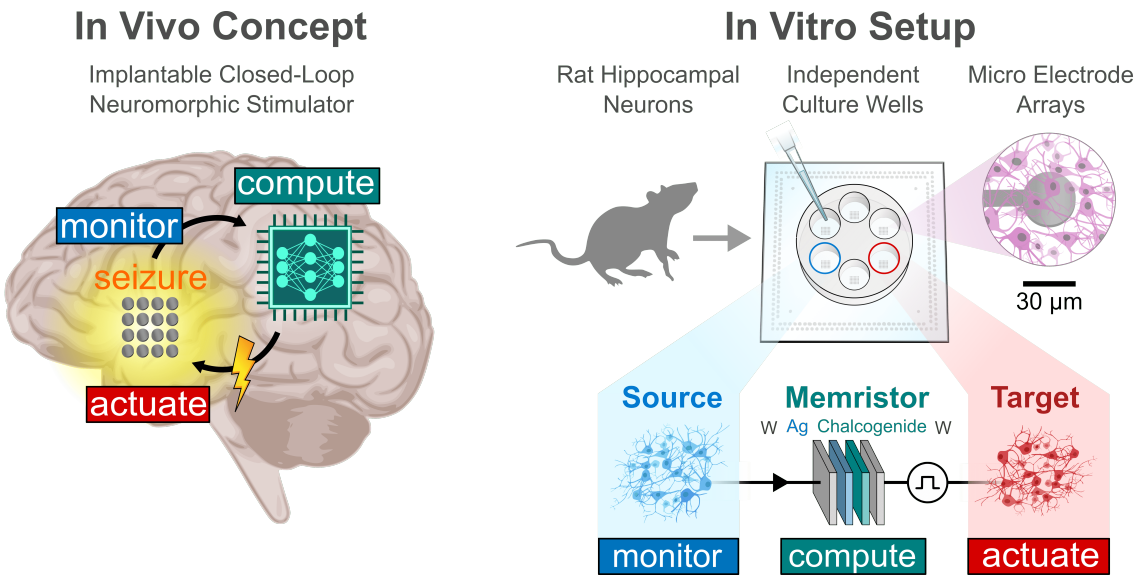



Figure 5.1: ***In vivo* concept and *in vitro* setup.** The neuromorphic technology presented in this study is envisioned for closed-loop loop devices with strict requirements in terms of size, power consumption and computing capabilities. The main function of the device is monitoring a given source population, processing the recorded activity in real-time, and actuating upon a target population when required (left). This was achieved in this study using an *in vitro* setup with hippocampal neuronal network cultures on MEAs , being monitored by a neuromorphic element – the memristor

instantiation of a Spiking Neuronal Network where each neuron is a physical, nanoscale, low-power, neuromorphic element. One of the challenges in the field is the envisioning, design, and fabrication of such a computing unit – the "neuron" in our analogy. One of the most promising candidates is the Memristor, a two-terminal resistance switch whose resistance value is a function of past voltage/current applied (Yang et al., 2012a; Upadhyay et al., 2019).

In this chapter, I will focus on the use of a memristor as the computing element responsible for monitoring and modulating neuronal activity (Figure 5.1). This proof of concept work resulted in a first co-author publication in the journal ACS Applied Electronic Materials titled *"Memristor-Based Neuromodulation Device for Real-Time Monitoring and Adaptive Control of Neuronal Populations"* (Dias et al., 2022).

This work was first-authored with Catarina Dias from IFIMUP at the Faculty of Sciences of the University of Porto (FCUP), and counted with the support of the following co-authors: Miguel Aroso from NCN lab at i3S;, João Ventura from IFIMUP at FCUPM and Paulo Aguiar, PI of NCN lab at i3S. This paper is available at `https://doi.org/10.1021/acsaelm.2c00198`

## 5.1 Introduction

Neurological disorders are a major cause of death and disability worldwide, and their burden in society continues to increase with population aging and growth (Feigin et al., 2019). Today's therapeutic strategies still rely heavily on pharmacological approaches, with important problems regarding nonspecificity and side effects. Furthermore, progress has been notably slow in discovering new drugs for diseases such as Parkinson's, Alzheimer's, or epilepsy. Recognizing that neuronal function is intimately related to electrophysiology, attention is steadily growing toward a different therapeutic strategy: direct modulation of neuronal electrical activity. Deep brain stimulation for Parkinson's disease, spinal cord stimulation for chronic pain, and cochlear implants for hearing loss are examples of success stories demonstrating the potential of this approach.

Neurotechnologies for therapies based on electrical modulation have, however, important challenges that still need to be addressed. Constraints in terms of size, power signature, signal transduction, and computational capabilities are currently limiting the progress in implantable medical devices. For example, the cost, risk (e.g., infection), and idiosyncrasy of surgeries solely for battery replacement cannot be underestimated. Development of novel devices should take into account that neurons and conventional electronics do not use the same electrical signals to encode information, and could take advantage of event-based operation modes to reduce energy consumption. Furthermore, any effective stimulation device requires adaptive computation capabilities to cope with the dynamic nature of neuronal activity. A device that is unable to dynamically respond to changes in neuronal activity patterns, falls short in its therapeutic potential and effectiveness (George et al., 2020).

It is only natural to address these neuroprosthesis challenges by mimicking key properties of the nervous system's solution to communication between neurons: the synapse. In this respect, memristive devices (Chua, 1971; Strukov et al., 2008; Jo et al., 2010; Wang et al., 2016b, 2018; Williamson et al., 2013; Pickett et al., 2012; Li et al., 2020) are a promising candidate to play the role of artificial synapses and integrate, as a core component, neuroprosthesis systems Jang et al. (2018); Fu et al. (2020); Broccard et al. (2017). Memristors are electrical components whose present conduction state depends on the electrical stimulus that has been previously applied to them. Although their properties rely on the particular combination of metal-insulator-metal materials used, the physical mechanisms behind conductance switching are typically related to the creation and rupture of nanoscale conductive filaments under the applied bias (Celano et al., 2015; Yuan et al., 2017; Buckwell et al., 2015; Barbera et al., 2015; Yang et al., 2012b). Besides the low operation power involved in these switching mechanisms, the conductance state is also nonvolatile, staying unchanged when the power supply is removed, with additional benefits regarding power consumption. Among the most studied materials such as metal oxides (e.g., TiO2 and HfOx) and semiconductors (e.g., Si), the chalcogenide family (e.g., GeSe and Ag2S) has shown promising neuromorphic properties (Li et al., 2013, 2014; Tuma et al., 2016; Xu et al., 2020; Brückerhoff-Plückelmann et al., 2021). Furthermore, as nanoscale two-terminal devices, their small feature

size (<100 nm) enables high-density integration architectures and hardware implementation of powerful signal processors such as artificial neuronal networks.

The dynamics of a memristor's response to stimulation is analogous to that of synapses, presenting learned transitions from low to high conductivity states, and vice versa (Figure 5.2..A). Importantly, the timescale of these transitions can be made analogous to the plasticity timescale of learning mechanisms in the brain, such as synaptic short-term plasticity (STP), which acts on the range from seconds to minutes. Furthermore, memristors can be integrated into the back-end-of-line of complementary metal-oxide semiconductor technologies in parallel with neuronal probe manufacturing (Gupta et al., 2016; Mikhaylov et al., 2020). This means that memristive technology can leverage the recent developments in high-density microelectrode arrays (MEAs) and neuronal probes carrying hundreds to thousands of recording/stimulating electrodes at an unprecedented spatiotemporal resolution (Steinmetz et al., 2020).

The recognition of the exceptional combination of properties in memristors has fostered important proof-of-concept studies demonstrating the feasibility of direct neuron-memristor connection (Gater et al., 2013; Lunelli et al., 2019; Chiolerio et al., 2017; Shchanikov et al., 2021). Memristors can not only be made to respond to neuronal activity but they can also act as an effective interface between biological and artificial neurons, implemented either using software or using very large-scale integration hardware (Serb et al., 2020). A unidirectional, activity-dependent, direct coupling between two neurons in brain slices has also been recently achieved via organic memristive devices (Juzekaeva et al., 2018). In that study, however, the memristive devices only progressed from a low to high conductance state (linking the neurons), limiting the functional usefulness of the coupling system. Neuronal activity in these systems is recorded/stimulated using either MEAs (Gater et al., 2013; Lunelli et al., 2019), patch-clamp pipettes (Juzekaeva et al., 2018), or a combination of the two (Serb et al., 2020). Memristors not only emulate fundamental synaptic properties but they can also be combined to achieve nontrivial computations. Important examples, with particular interest for memristor-based neuroprosthesis, include the use of memristors to perform real-time processing of neuronal spikes Gupta et al. (2016) or the use of memristor arrays to implement the filtering and identification of epilepsy-related neural signals (Liu et al., 2020b). Although these proof-of-concept studies have covered important separate aspects of the potential role of memristors in implantable neuroprosthetic devices, to the best of our knowledge, no work has yet demonstrated the use of memristors in a fully functional configuration and in a clinically relevant setting.

It has already been well-established that electrical stimulation can ameliorate the symptoms of epilepsy (Fisher and Velasco, 2014). As such, in this work, we take as our motivation the development of memristor-based closed-loop neurostimulators for intractable epilepsy. We believe such devices should perform three core tasks in a monitor-compute-actuate paradigm (Figure 5.2.B): (1) online monitoring of a neuronal population prone to seizures, (2) real-time seizure detection, and (3) stimulation of specific (inhibitory or interfering) neuronal populations to suppress seizure pro-

gression. To produce our control setting, we developed a memristor-based interface (fully implemented in hardware) capable of performing real-time monitoring and adaptive coupling between two neuronal populations (Figure 5.2.C). Our system establishes direct communication between neuronal populations (and not just individual neurons) and in vitro experiments are carried out with neurons from the hippocampus, a brain region frequently involved in epilepsy. The reversible short-term dynamics of the memristors are used both for the detection of network bursts (NBs) (a hallmark of epileptic seizures) and for the gating of electrical stimulation to the target neuronal population. NBs, defined as periods of strongly synchronized high-level activity in the neuronal population, share fundamental characteristics in both *in vitro* and *in vivo* conditions (Suresh et al., 2016).

## 5.2 Methods

### 5.2.1 Memristive Devices

The memristive devices used were acquired from KNOWM Inc., and are composed of stacks of W/C + $Ge_2Se_2$/SnSe/$Ge_2Se_3$ Mix/Ag/$Ge_2Se_2$ Adhesion/W (Knowm, 2019; Campbell, 2017, 2019). They are analog devices with very low switching energy and fast switching response. The inherent threshold voltage of these devices varies in the range of approximately 0.25-0.45 V (Figure 5.2.C). As the plasticity dynamics can be tuned by the applied frequency, amplitude and duration of pulse stimulation, to properly choose the values of the components in the electrical circuit (Figure S3), we performed a preceding characterization of the devices' response to pulse stimulation, as can be seen in Figure 5.2.D.

### 5.2.2 Cell Culture

All the experiments were performed in accordance with the European legislation for the use of animals for scientific purposes and protocols approved by the ethical committee of i3S. The Animal Facility of i3S is licensed by the Portuguese official veterinary department (DGAV, Ref 004461), complies with the European Guidelines (Directive 2010/63/EU) transposed to Portuguese legislation by Decreto-Lei no 113/2013, and follows the FELASA guidelines and recommendations concerning laboratory animal welfare. Embryonic (E18) rat hippocampal neurons were seeded and cultured at a density of $5 \times 10^4$ cells/well on a 6-well round chamber MEA with a macrolon ring 10 mm high (256-6well MEA200/30iR-ITO-rcr) (Multichannel System MCS, Germany). The 6-well MEA electrode has an array of $7 \times 6$ TiN electrodes in each well, with a total of 252 recording electrodes. Half-medium changes of Neurobasal TM Plus (Thermo Fisher Scientific) were performed every 2-3 days.

### 5.2.3 Electrophysiological Recordings

The experiments were performed at 11 and 13 DIV using the MEA2100-256 system (Multichannel System MCS, Germany) at a sampling rate of 10 kHz. The electrophysiological signals were high-

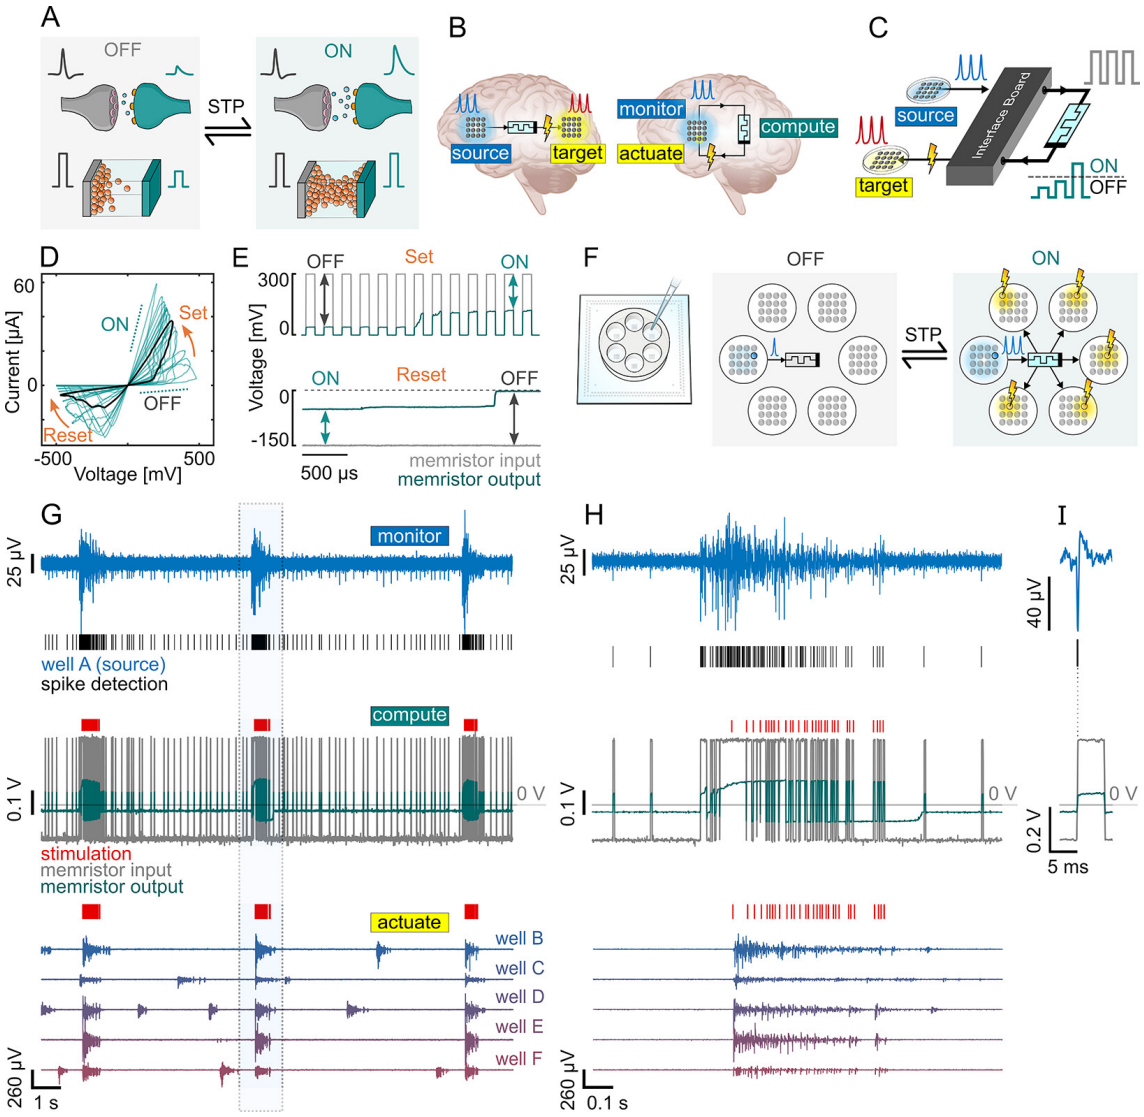



Figure 5.2: **Real-time monitoring and adaptive control of neuronal populations using memristors.** (A) Biological and memristive synapses share an analogous STP, with dynamics interchanging between low (OFF, left) and high (ON, right) conductance states as a function of the activity history. (B) Envisioned neuroprosthesis in vivo. The source and target neuronal population can be distinct (left) or the same (right), providing, either way, an effective monitor-compute-actuate paradigm through the memristor-based device. (C) In vitro implementation used in this work, with a hardware interface bridging neuronal and memristor signals. (D) Memristor hysteretic loop showing "set" under positive and "reset" under negative voltage [$V(t) = 0.5sin(100t)$]. The multiple cycles show the intrinsic variability of these types of devices, similar to the variability found in biological synapses. (E) STP dynamics with potentiation for repetitive positive pulsed (300 mV, 100μs) stimuli (top) and recovery under constant negative (-150 mV) stimulus (bottom). (F) Schematic representation of the in vitro setting, integrating six physically independent neuronal cultures in a 6-well MEA (left) and a memristor to monitor the source population in real-time and dynamically modulate the activity of five target populations (right). [*Continues on the following page*]

Figure 5.2: (G) Representative example of the real-time neuronal monitoring and modulation performed by the in vitro memristive interface. Actuation upon the target population is gated by sustained/consistent high-level activity at the source electrode (blue, top). The detected spikes (black traces, top) are converted to pulses and fed to the memristor (gray, middle). The memristor output (green, middle) increases when the source bursts, changing to the ON state. Consequently, this triggers electrical stimulation (red traces, middle and bottom) in the target population, which is therefore activated by the source. The recordings shown for the targets (bottom) are from electrodes neighbouring the stimulation electrode. When the source bursts end, the memristor transitions back to the OFF state and stops propagating the source spikes to the target. (H) Zoom of panel G, evidencing the OFF-ON transition at the beginning of the source burst and the ON-OFF transition at the end. (I) Detail of a source neuronal spike and the associated detection performed by the interface board in real-time hardware (top). The spike creates a pulse that is applied to the memristor (bottom).

pass filtered at 200 Hz. The recordings were acquired using Experimenter software from MCS. Cell culture conditions (37 °C and 5 % CO2) were maintained with a stage-top incubator (ibidi GmbH, Germany) adapted to the headstage of the MEA2100-256 system.

### 5.2.4 Real-Time Spike Detection

The detection of spikes in the source well was performed in real time (latency of less than 20 µs) with the digital signal processor (DSP) included in the interface board of the MEA2100-256 system, using a threshold crossing method on the already filtered signal. The negative threshold had five standard deviations (sometimes manually tuned at the beginning of the trials). Such a threshold is significantly higher than the background noise, ensuring that only neuronal spikes are sent to the memristor. The DSP was configured using Experimenter software before starting the trials, making it independent of the computer from there on. The interface board sends a Transistor-Transistor Logic (TTL) pulse to the memristor whenever the potential at the recording electrodes exceeds a pre-established threshold. The duration of the TTL was adjusted at the beginning of the trial to either 1, 5, 10, or 20 ms. Once configured, the entire system was fully independent of the computer.

### 5.2.5 Mediating Electrical Circuit

The 3.3 V TTL pulses arriving from the interface board (digital output) were attenuated to fit the memristor operation range (millivolt range) using a dedicated circuit (see Figure S3 in 5.A). The circuit was tuned to ensure that only a high frequency of incoming pulses would set the memristor from OFF to ON in the required temporal scales. Also, the baseline level was set to a small negative value that ensured the transition from ON to OFF after a significant period without neuronal spikes. The signal processed by the memristors was converted back to the operating range of the digital input of the interface board.

### 5.2.6 Electrical Stimulation

A stimulation electrode was selected for each target well. The electrical stimulus was a negative monophasic voltage pulse of 500 mV and 200 µs. The electrical stimulation was triggered for all the stimulation electrodes every time the interface board received a high amplitude TTL pulse (above 2 V) coming from the memristor. This way, when the memristor was ON, each spike detected in the source well triggered an electrical stimulation in each target well. The interface board operates at 50 kHz, meaning that it takes 20 µs to activate the stimulation once the memristor response is received. The closed-loop latency associated with the spike detection performed by the DSP and the activation of the stimulator is lower than the sampling period (0.1 ms). Stimulation artifacts are removed by blanking the recording electrodes for a few instants (2 ms) after each stimulus (blanking operation available in the MEA2100-256 system).

### 5.2.7 Experimental Protocol

The experiments began with a baseline recording of 10-30 min without memristor intervention. An electrode that had both bursting activity and isolated spikes was chosen as the source electrode and its well, therefore, as the source well. To select effective stimulation electrodes, the four most active electrodes of each well were individually stimulated with 20 electrical pulses at an interval of 10 s. For each well, we chose the electrode that triggered a response in the largest number of neighboring electrodes. The memristor was then inserted into the conditioning circuit connected to the interface board. Before starting the memristor experiment, the duration of the TTL pulses was adjusted to either 1, 5, 10, or 20 ms to assure proper STP on the desired temporal scales. Each trial had a maximum duration of 15 min, but could be stopped earlier if the memristor got stuck in an ON or OFF state for several minutes. The entirety of the recordings was considered in the analysis, including the periods where the memristor was apparently stuck in a state. A computer running MCS Experimenter software (Multichannel System MCS, Germany) was used solely to (1) record neuronal data and the stimulation time stamps for later offline analysis and (2) upload to the MEA2100 system the parametrization for the stimulator and the DSP (after this upload, the closed-loop system itself is independent of the computer).

### 5.2.8 Signal Processing

The neuronal signals were filtered with a 200 Hz high-pass and spike detection was performed for all the electrodes (not in real-time) using positive and negative six standard deviation thresholds with 3 ms dead time. The mean firing rate of each well was obtained by convolution of the spike trains of the electrodes with a Gaussian kernel of 0.05 s sigma and averaging across electrodes. The cross-CC curves were calculated between the mean firing rate profiles of each source-target well pair. The NBs consisted of periods when several electrodes burst simultaneously. Events that included less than five bursting electrodes were not considered. To identify the bursting periods of each individual electrode associated with the NB, we considered groups of at least five spikes with an interspike interval of less than 100 ms. The interspike interval considered for the burst detection

in individual electrodes is larger than what is typically considered for bursting neurons (usually around 5 ms) because we wanted each electrode burst to encompass the full contribution of that electrode to the NB (instead of multiple bursts in the same electrode for the same NB).

## 5.3 Results and Discussion

In this work, we used commercial memristors composed of stacks of W/C + Ge2Se2/SnSe/Ge2Se3 Mix/Ag/Ge2Se2 Adhesion/W (Knowm, 2019; Campbell, 2017, 2019). These memristors rely on the movement of silver ions (Ag+) into channels within the active layer, which has been doped with carbon to enhance and optimise their properties and are characterised by low-power binary switching. During the initial step of electroforming, under an applied positive voltage on the top electrode, Sn ions from the SnSe layer are generated and diffused into the active Ge2Se3 layer, where a metal-catalysed reaction distorts the glass network to provide conductive channels for the movement of Ag+. Because the amount of Ag within the channel determines the resistance of the device, the resistance is then tunable by the movement of silver into or away from these channels by applying positive ("set", ON state) or negative ("reset", OFF state) potential, respectively (Figure 5.2.D) (Campbell, 2017). The typical device-to-device variability is reflected in the electrical I-V behaviour, where a larger hysteresis translates into a higher separation between the resistance states (Figure 5.4). Notably, these types of memristors are capable of short-term memory dynamics. Of particular importance for this work is the reversibility of the ON-OFF conductance transitions. With conditioning positive pulses and constant negative voltage, we were able to produce interchangeable "set" and "reset" transitions, respectively, with adequate timescales for a neurobiology setting (Figure 5.2.E). These transitions are shown in the cumulative voltage increase from low (OFF state) to high (ON state) at each positive pulse ("set") and the change from a constant negative voltage (ON state) to almost zero (OFF state) under constant applied bias ("reset"). This cumulative voltage increase after each positive pulse produces functionally relevant STP behaviour (Figure 5.5). Aiming at population-level control, our designed system relies on MEAs, which allow long-lasting recordings and modulation of neuronal activity (as opposed to patch-clamp electrodes). To discriminate the memristive device's specific contribution to neuronal activity modulation, we used neuronal cultures on MEAs with a 6-well configuration, forcing the existence of six physically independent neuronal populations (Figure 5.2.F). Communication between the source population/well and the remaining populations/wells was mediated by the memristor, whose selectivity commanded electrical stimulation of the target populations according to the patterns of activity detected in the source population.

Currently, there are still no memristors available working in a voltage/current range compatible with a direct connection to neurons (μV and nA). As such, and for now, instead of a direct/passive circuit neurons-microelectrode-memristor-microelectrode-neurons, our system includes hardware (e.g., amplifiers and electrical stimulators) to translate between microelectrodes and memristors' voltage amplitudes (Figure 5.6).

### 5.3.1 Memristive Interface Detects and Selectively Responds to Network Bursting Activity in Real Time

Initially, the memristor is in a low conductance state (OFF), analogous to a weak synaptic connection. Bursting activity at the source electrode induced a gradual increase in the memristor conductivity, changing its state to ON. When the memristor was ON, source spikes triggered electrical stimulation in the target population, inducing them to fire with the source population (Figure 5.2.G,H). Importantly, the established neuronal coupling/modulation is reversible and the connection is dissolved when the spiking rate of the source decreases back to baseline activity. Note that the evolution of the memristor's state is automatic and unsupervised, which means that no additional system is used to actively change its conductance - the memristive interface does the NB detection computation autonomously. The time needed for the transition (<100 ms) is still low enough for the burst triggered in the targets to be in synchrony with the source burst, as the latter can last between several hundreds of milliseconds to seconds.

The memristor receives a pulse for every spike in the source electrode (Figure 5.2.I). A conditioning circuit (Experimental Section and Figure S3) was fine-tuned prior to the experiments to guarantee that the different patterns of activity exhibited by the neurons (patterns of received pulses) induce the memristor to interchange between its conductive states, ensuring the desired selective response to bursts. We lowered the amplitude of the arriving signal to assure that each pulse had a positive amplitude, lower than the memristor's "set" threshold. A pulse with an amplitude above the "set" threshold would immediately transition the memristor state to ON with a single source spike, eliminating any computation capabilities. When no spikes were being sent, the voltage across the memristor was negative with a small amplitude, slowly changing its resistance back to OFF. These parameters - amplitude of the pulse and negative baseline level - were optimized a priori and kept unchanged for all the experiments. This was performed taking into consideration the natural stochasticity of the memristive behaviour. In terms of material dynamics, the positive pulses are gradually and cumulatively creating a metallic filament across the device, whereas the negative constant voltage is slowly dissolving the filament. Note that the latter would not be needed in the case of volatile memristors.

### 5.3.2 Memristive Interface Promotes Robust Coupling and Modulation of Neuronal Populations

The memristor communicates with the neuronal cultures through a source electrode in the source well and a stimulation electrode in each target well (Figure 5.3.A). The bursting periods detected by the memristor are associated with network-wide events that dominate the culture activity (Figure 5.3.B), referred to as NBs. Although the memristor responds to the spikes from a single source electrode because this extracellular electrode records from multiple neighbouring cells Yger et al. (2018) (see zoom in Figure 5.3.A), the memristor is in fact tuned to detect the patterns of strong and synchronous network bursting activity. When there is no memristor-modulated communication between the source and target wells, the network activities of the six cultures are independent

of each other (Figure 5.3.B, left). When the memristor is inserted, the NBs of the source are detected due to the sudden increase in activity in the source electrode. This activity pattern sets the memristor state to ON, closing the modulatory bridge between the source and target wells via electrical stimulation. In turn, the electrical stimuli applied to the stimulation electrodes modulate neuronal activity in the target wells, thus establishing a communication pathway between previously isolated neuronal populations (Figure 5.3.B, right). The short-term memory of the memristor operates at the temporal scales of the NB dynamics and guarantees that after the NB finishes in the source well, the memristor transitions back to OFF, avoiding isolated source spikes to interfere with the activity of the target populations.

Three different memristors were tested in separate trials, each monitoring a different source well. Qualitatively, the three memristors maintained a strong coupling between the high-frequency activity patterns in the source and target wells when compared to the baseline (without memristor). To quantify this dynamic coupling, we computed the cross-correlation (CC) between the mean firing rate of each source-target well pair for each trial (Figure 5.3.C). Despite the inherent variability of the memristor, the CC curves showed a significant correlation between the firing rates of source and target wells, with a delay in the millisecond range. The memristive interface was thus capable of maintaining a robust low-latency coupling of specific activity patterns between independent neuronal populations.

Besides quantifying the temporal coupling of the different networks, we also evaluated the impact of the memristive control on the modulation of the activity fingerprints of the target networks. Specifically, we focused on the changes in properties such as duration and intervals between NBs, and maximum firing rates inside each NB (Figure 5.3.D). In the system without the memristor, each well had its own, independent, network activity fingerprint. In the memristor-based system, the network activity profile in all the target wells followed, within physiological bounds, the command activity of the source well (Figure 5.3.E). This shows that the dynamic modulation provided by the memristor aligned the activations of the target populations with the NBs of the source population. In the context of the detection-suppression system for epilepsy, the target neuronal population (or brain region) would be chosen so as to have an inhibitory (or interfering) effect on the source population.

## 5.4 Conclusions

Neuromorphic devices provide promising properties to address the size, power, and computational requirements critical for innovative implantable neuroprosthesis. Memristors, in particular, exhibit neuronal-like dynamics and can be tuned to operate on equivalent temporal scales. Here, we used memristors as an effective alternative to microprocessors to monitor biological neuronal populations and selectively modulate their activity in a long-term autonomous setting. Our in vitro memristor-based system performs the three core tasks required for an efficient feedback

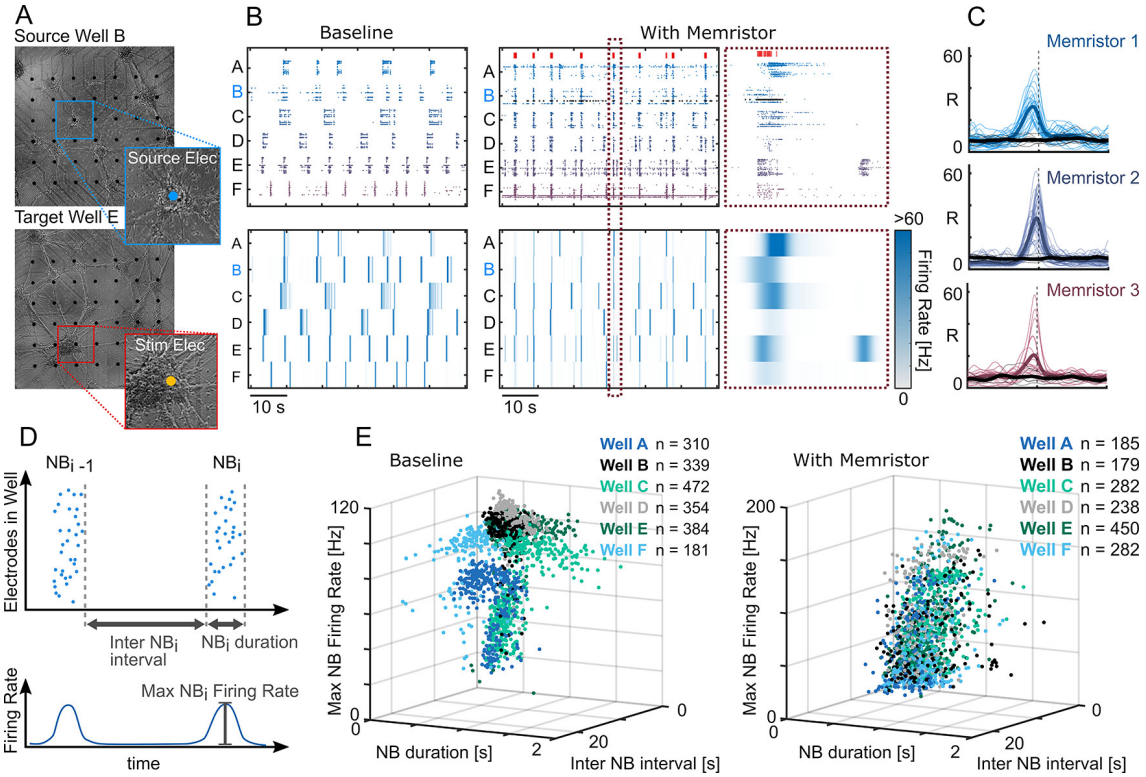



Figure 5.3: **Real-time coupling and modulation of neuronal networks using a memristor.** (A) Representative source and target neuronal networks cultured in different wells (scale bar = 500 µm). (B) Raster plot (top) and average firing rate per electrode on the distinct wells (bottom); the left and right panels correspond, respectively, to the system without and with the memristor-mediated connection. High-frequency spikes in the source electrode (black dots, top right) correlate with NBs in the source well (well B, in this representative example). (C) Synchronisation of high-frequency activity between the source and targets is evidenced by the CC curves between the average firing rate per electrode on the wells (B, bottom). There is a CC curve for each source-target pair per trial for the three memristors tested (Memristor #1: 6 trials, 3.7 $\pm$ 1.7 min each; Memristor #2: 8 trials, 6.8 $\pm$ 4.2 min each; and Memristor #3: 6 trials, 11.9 $\pm$ 4.5 min each). The coloured lines represent the CC curves for the trials with the memristor, gray curves are for trials without the memristor. The thicker lines represent the average. (D) Diagram describing the three parameters used for the network activity fingerprints: NB duration, interval, and maximum firing rate. (E) The memristive interface modulates the dynamics of the network activity in the target wells. Each dot represents an NB in a given well. Before inserting the memristor (left), the NBs of each well had consistent separable features (duration, interval, and maximum firing rate). With the memristive interface, the activity fingerprints of the source (well B) and target wells are indistinguishable (right), showing that the network dynamics that govern the target populations were modulated by the source.

neurostimulator-monitor, compute, and actuate-and it does so with a nanosized, low-power, neuromorphic computing element. The robust low-latency detection and modulation of network activity patterns presented in this proof-of-concept is fundamentally important for an implantable neuronal stimulator in many clinically relevant contexts. Here, the computations were performed by a single memristor, allowing the detection of simple neuronal patterns. Future studies should focus on the integration of memristive arrays with biological neurons to perform real-time detection of more refined patterns of activity. Also, here the communication between neurons and memristors was mediated by an interface board and a conditioning circuit. With memristors' diversity increasing every day, future improvements will focus on establishing a passive direct connection between neurons and memristors.

## 5.A Supporting Information

### Acknowledgements

This work was financially supported by Project PTDC/EMD-EMD/31540/2017. The authors thank P. Cruz for all the help with the electrical circuit and its components.

### Supplementary Figures

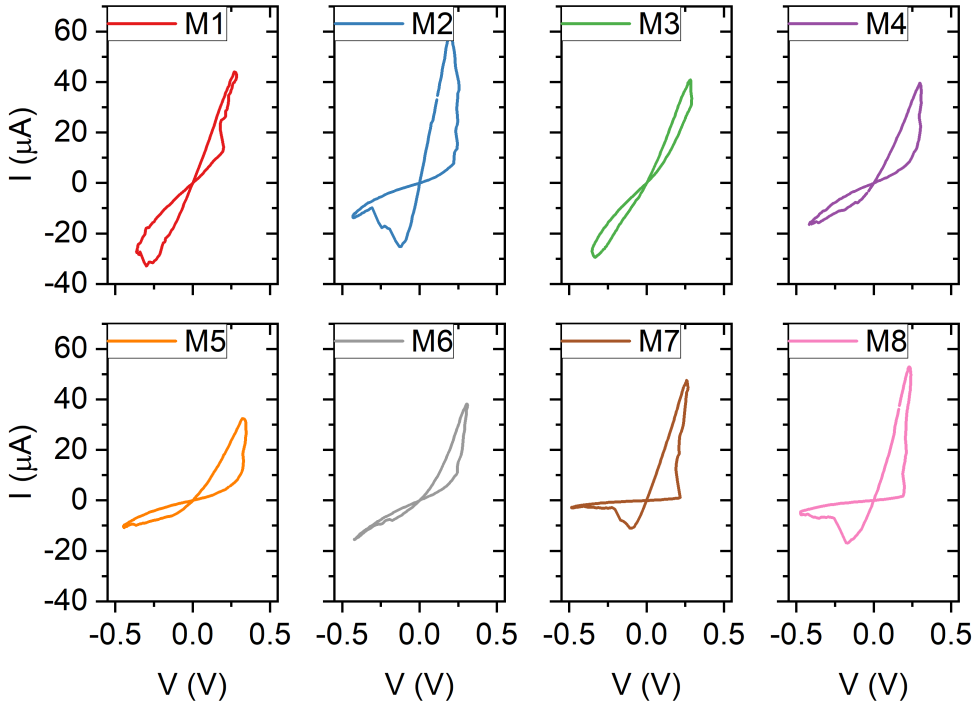



Figure 5.4: **Representative electrical I-V hysteretic behaviour of 8 discrete memristive devices in the same chip.** The devices show low operation power and the typical device-to-device variability. A larger hysteresis translates in a higher separation between the resistance states.

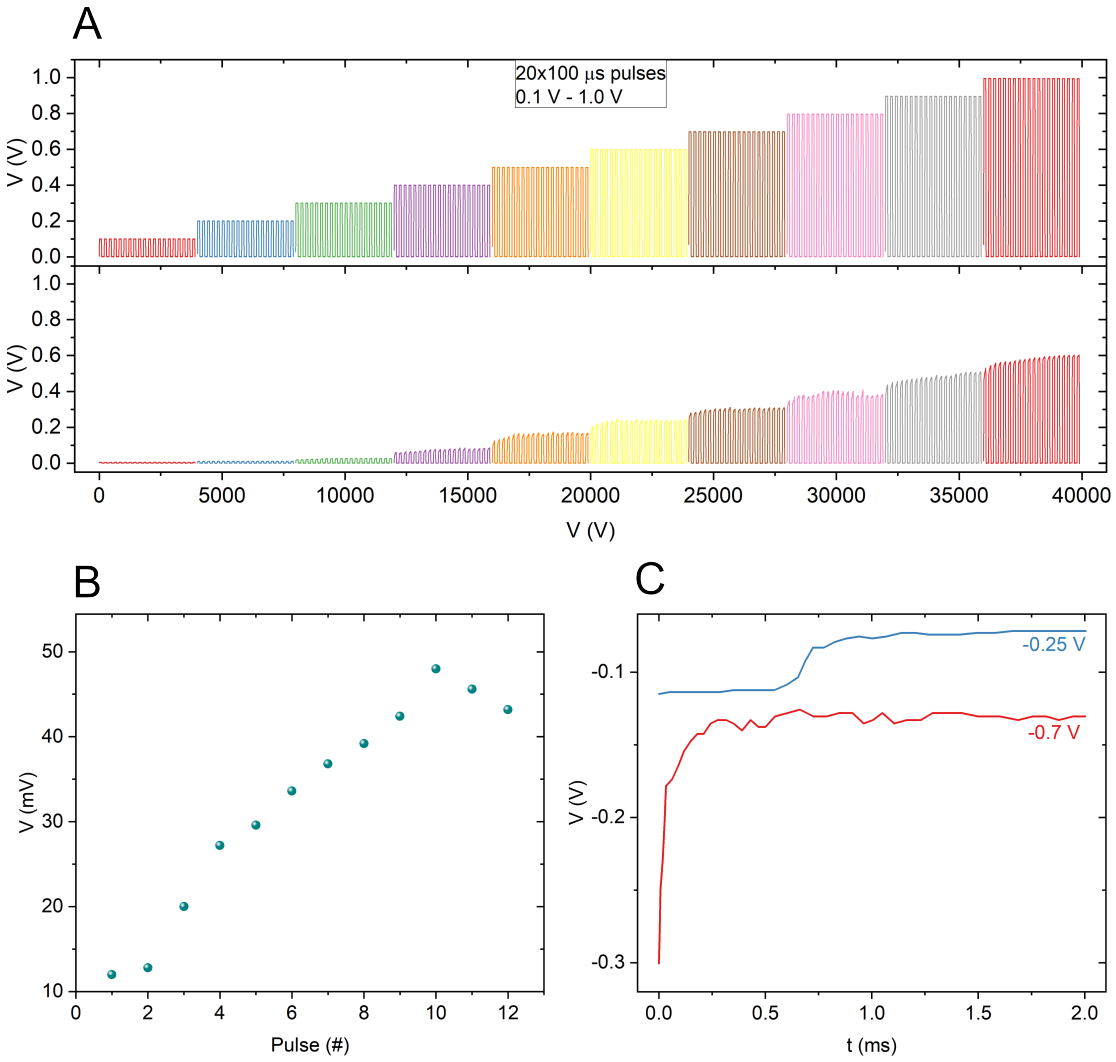



Figure 5.5: **Representative short-term plasticity (STP) behaviour.** A) Effect of the voltage amplitude (0.1-1.0 V) on the resistance change at fixed pulse number (20), duration (100 µs) and frequency. B) Cumulative memristor voltage increase (resistance decrease) after each 320 ms pulse at 2 Hz. C) Effect of the negative voltage amplitude (-0.25 and -0.7 V) on the resistance change time

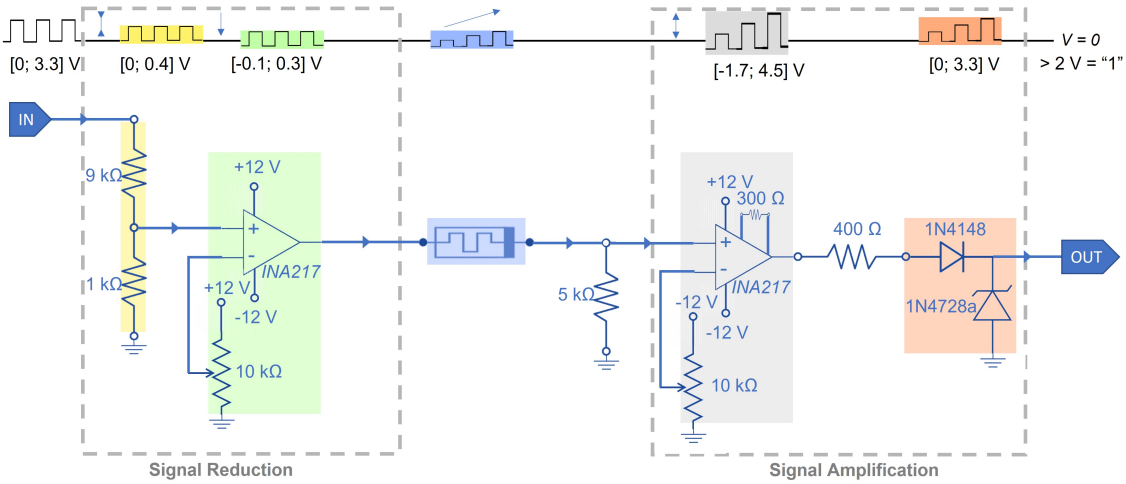



Figure 5.6: **Electrical circuit used to connect the memristor with the interface board of the MEA2100 System from MultiChannel Systems (bottom) and a schematic representation of the TTL signal across the circuit (top).** The first block is responsible for the transformation of TTL signal to the lower operation voltage range of the memristor. A voltage divider and a low-noise instrumentation amplifier shift the pulse to have both a positive (for "set") and a negative (for "reset") component. The last block amplifies the signal at the memristor output, for it to be read by the interface board as a logical "0" when the memristor is OFF (TTL < 2 V) and a logical "1" when the memristor is ON (TTL > 2 V). The diodes are only present as a precaution, for the protection of the data acquisition system.

# Chapter 6

# Conclusions

## 6.1 Bird's-eye view

The prevalence of neurological disorders is increasing with the rise of life expectancy. These disorders shorten the patient's longevity while progressively reducing the quality of life. So much so that the number of *healthy* years lost due to neurological disorders increased by 15% in the last three decades (Feigin et al., 2020). This last point is of particular importance for the work presented in this thesis. Electrical stimulation therapies will not cure neurological disorders and, even if they prevent premature deaths, it would be only by mitigating symptom-related accidents – epileptic patients, for example, may suffer severe injuries during seizures (Wirrell, 2006). What these therapies promise instead, is to alliviente or suppress the harsh and limiting clinical symptoms: the PD patient would no longer have tremors; the epileptic patient would have no more seizures. Other applications where stimulation would not directly alter *life expectancy* but rather improve *life quality* involve replacing lost sensory function with implantable neuroprostheses, such as cochlear or retinal implants – the deaf would hear and the blind would see once again.

Significant technical challenges still stand in the way of this almost utopian future. Current neurostimulation therapies have limited efficacy and work only for some patients, to the extent that its implementation is only seen as a last resource, when everything else fails, and must always be carefully weighed against the burden of possible side-effects or eventual post-surgical problems. There are multiple fronts on which technical improvements must unfold. I will briefly outline where I think most of these fronts lie, and which ones were expanded with the work presented in this thesis.

Firstly, there are several aspects of neural stimulation that have great margins for improvement but were not addressed here. For example, the fundamental knowledge about the specific molecular and cellular mechanisms of disease and its interaction with stimulus will certainly play a huge role, allowing the optimisation of therapies from first principles, instead of on a trial-and-error basis. Also, improving the mechanical and surficial biocompatibility of the implanted leads will greatly

improve the longevity of the device and its long-term efficacy (Lawrence et al., 2002). However, the work presented here did not focus on these factors, so there is not much potential for this thesis to expand the state of the art on this front.

On the more technological side, the trend in neuromodulation is to move towards a closed-loop stimulation paradigm and that will require innovations in the devices' hardware and software. Regarding the hardware, the goal is somewhat clear: the device must be endowed with many electrodes to enable brain recordings from multiple sites with high spatial precision; must have a stimulator, ideally with access to multiple electrodes to expand the actuation landscape; and some sort of highly efficient implantable processor to analyse the incoming data in real-time and respond with adequate stimulation. Some companies are already taking the lead on this front, with devices capable of closing the control loop, even if only for research purposes, as we discussed in section 2.2.2. As for the software component, the crucial aspect is the control algorithm – and here the possibilities are endless. As discussed in section 2.2.3, the algorithms ever tested in humans are based on the simplest forms of control and, so far, the results are not surprising. Still, in theory, closed-loop control must perform better than open-loop – the key is having the right control algorithm, aligned with the adequate actuation means. This is the main front that this thesis aims to push forward.

## 6.2 Contributions to the field of neuroengineering

I believe that closed-loop stimulation will play an important role in two domains of neurostimulation applications: 1) suppressing pathological brain activity and 2) replacing lost sensory function. It is still an open question, however, what sort of control algorithms should be used in each application and how to take advantage of Control Theory, Data Science and other relevant engineering areas for optimal neuromodulation. As such, these were the problems targeted with our first two research questions, in Chatpers 3 and 4, respectively. With the third research question, in Chapter 5, we sought to address alternative ways of computing that could be implemented in closed-loop neurostimulators to improve efficiency and computational capabilities while reducing size.

In this section, I will expand on the specific contributions this work brought to the field of neuroengineering by revisiting these three questions and trying to evaluate the proposed solutions with a critical lens.

#### Chapter 3 - Research question #1

The symptoms of many neurological disorders can be traced back to some pathological feature of neuronal activity – some biomarker that can be recorded and detected through the implantable electrodes. These are typically manifested at the level of brain oscillations, which may contain pathological components due to excessive neuronal synchronization at particular frequencies. The goal of brain stimulation is to disrupt these pathological dynamics and restore healthy neuronal

dynamics. Closed-loop stimulation is a potential solution to suppress these abnormal dynamics, but the appropriate approach is unknown. As such, our first research question was:

*What type of control algorithm can be used to disrupt pathological oscillations?*

There is a strong literature about control methods for complex systems with highly nonlinear dynamics, and one of the most popular methods is Delayed Feedback Control (DFC) (Pyragas, 1992). Based on computer simulations, many have argued that this is an efficient method to control pathological brain oscillations (Rosenblum and Pikovsky, 2004; Hauptmann et al., 2005; Popovych et al., 2005; Batista et al., 2010; Popovych et al., 2017a; Daneshzand et al., 2018a; Zhou et al., 2020; Liu et al., 2023), so we sought to test this hypothesis for the first time with actual biological neuronal networks, which also displayed synchronous oscillations. What we noticed right from the first experiments was that the results obtained did not match the theoretical prediction – DFC did not suppress the target oscillation but instead created a faster and more consistent oscillation with approximately twice the frequency. We believe that one of the reasons for the success of DFC in computer simulations and its failure *in vitro* is the fact that *in silico* networks cannot sustain multiple oscillation frequencies, whereas *in vitro* networks certainly can.

Since DFC was never tested *in vivo*, we do not know yet if its vanilla implementation (without the adaptive component) would also cause the amplification of faster oscillations in real brains. It is a possibility though, since the neuronal activity *in vivo* is even more dynamic than *in vitro*. I think the most important "meta-contribution" of this work is serving as a cautionary tale for the direct translation of neuromodulation algorithms from *in silico* to *in vivo*. *In silico* models do not fully capture the complex nonlinear temporal dynamics of real neuronal networks and their interaction with stimulus. Using these models to evaluate the effects of stimulation algorithms may result in unforeseen outcomes with serious consequences for the test subjects.

Regarding more concrete contributions, we presented Adaptive Delayed Feedback Control (aDFC), a DFC algorithm that adjusts the feedback delay to match the ongoing oscillation periodicity and disrupt it. The delay is a critical parameter of DFC, to the point that it determines whether the algorithm suppresses the oscillation or establishes it (Pyragas, 2006). Our *in vitro* tests showed that aDFC can effectively disrupt the unwanted neuronal oscillation without imposing a new one. So, *what type of control algorithm can be used to disrupt pathological oscillations?* According to our *in vitro* studies, aDFC seems like a good candidate.

This study also introduced two relevant concepts for the study of *in vitro* neuronal cultures: 1) network controllability and 2) the wavelet transform of the instantaneous population firing rate. Network Controllability is an active area of research in physics and engineering and has already been considered for the study of neuronal networks (Su et al., 2017; Karrer et al., 2020), although not without some controversy (Gu et al., 2015; Tu et al., 2018; Suweis et al., 2019). Still, in all closed-loop control studies we analysed (see section 2.2.3), we never saw network controllability being mentioned as a possible confounding factor. Our experiments, however, showed that

this may be a crucial element to consider and analyse when developing closed-loop protocols for neuromodulation. We also present a statistical method to identify controllable and uncontrollable networks. We expect similar approaches and considerations to be relevant is future studies.

Regarding the second point, I believe we introduced a novel and very useful way of analysing multi-electrode spiking data. We were interested in evaluating the collective neuronal dynamics in long-term recordings, and looking at rasterplots from dozens of electrodes is only useful for short time-scales (seconds, to a few minutes). Beyond that, there is not much information to be extracted from this type of representation. Our novel representation is based on wavelet transforms, which are typically applied directly to LFP (Roux et al., 2007; Aguera et al., 2011; Zhang et al., 2018) or EEG data (Akin et al., 2001; Adeli et al., 2003; Aguera et al., 2011; Hamad et al., 2016). However, here we took a different approach – we applied this transform to the average instantaneous firing rate of the entire population. The resulting transform evidences the bursting dynamics of the population, its periodicity, regularity, multi-stable states, and how these features vary across long time-spans. We believe that this type of representation will be very useful for other researchers working with multi-electrode spiking data.

Finally, we also provide an open-source application for closed-loop stimulation compatible with the MEA2100-256-System. Our app has an intuitive user interface and general-purpose functionalities – such as online filtering, spike detection, burst detection, instantaneous firing rate – and the control algorithm block can be easily modified to test different control methods. The code is freely available at `https://github.com/NCN-Lab/aDFC`.

**Chapter 4 - Research question #2**

Abraham Maslow said that when you only have a hammer, you see every problem as a nail. But Deep Learning is a really good hammer. Artificial Neuronal Networks (ANNs) are very efficient general function-approximation algorithms, which explains their explosion in all sorts of applications in the last decade. So, we believe that there is great potential in using such machine learning techniques to learn the input-output function of neuronal systems, and exploit that to conceive optimal stimulation protocols. That was the goal of our second research question:

*How can machine learning tools be used to guide optimal electrical stimulation?*

In many neuroscience problems, predicting the neuronal response to a given stimulus – the forward model – is much easier than predicting the stimulus that produces a given neuronal activation – the inverse model. The inverse model, however, is what you need if you want a neurostimulator the restore lost sensory perception.

In Chapter 4 we presented our solution for retinal prosthesis which is based on an encoder-decoder deep learning architecture to estimate the inverse retinal model. This architecture is not new, and recent publications on retinal stimulation had a similar approach (Granley et al., 2022; Relic et al., 2022; Wu et al., 2023). Here, we innovated by using this architecture to exploit a stimulation

paradigm known as Neural Activity Shaping, which uses multi-polar stimulus patterns to increase the spatial resolution of neural activation (Spencer et al., 2019). This method showed promising results using a realistic computational model of retinal response to multi-polar stimulation (Maturana et al., 2016).

In theory, the encoder-decoder framework presented here can be applied to different types of sensory prostheses where precise spatial neuronal activation is crucial, such as for the retina or cochlea. However, currently, our method does not consider the temporal dimension. This means that it can only be applied to "memoryless" neuronal systems, i.e., systems where the state at time $t$ depends only on the input provided at time $t$, being independent of past stimulus. Sensory systems behave like this, otherwise you would always see motion blur and hear echoing sounds. This also means that this method cannot be used for neuronal systems presenting complex recurrent connections, such as, arguably, most of the brain. For example, this stimulation paradigm would not be applicable for the sort of problems presented in the research question #1, since the neuronal dynamics involved are inherently dependent on past stimulus.

**Chapter 5 - Research question #3**

In the previous two chapters, we focused on two different neuromodulation approaches, for two different neuromodulation problems. These were focused on pushing the software frontiers a bit further. Our final research question was focused on the hardware frontier. With more complex control algorithms comes a need for more efficient computing devices which should be as small as possible and have minimal power requirements. So, we asked:

*How can intelligent stimulation protocols run efficiently in future neuroprosthesis?*

Neuromorphic engineering proposes a promising alternative with *in situ* neuro-inspired processing. These strategies could be used for hardware implementations of ANNs such as those used in Chapter 4. Within the neuromorphic universe, memristors are particularly interesting as these nanometric devices resemble the memory mechanisms of synapses, and can be tuned to work on the same current and time scales. In chapter 5 we explored the possibility of using memristors to process neuronal events in real-time and trigger modulating neurostimulation. Our *in vitro* setup showed that memristors could indeed be used to detect features of neuronal activity (neuronal bursts) and be used to condition the activity of target neuronal populations. This was the first time memristors were used in such a dynamic setting, interacting dynamically and in real-time with biological neurons. However, tuning the memristors to the temporal scales of the neuronal population required a laborious optimization of an interfacing electrical circuit, which would not be feasible in implantable devices. Still, this study served as a proof-of-concept for the possibility of using memristors as dynamic and highly efficient processors of neuronal activity which, hypothetically, could be directly connected to neurons, powered by the neuronal activity itself without any interfacing system, as they operate with similar power signatures.

So, *how can intelligent stimulation protocols run efficiently in future neuroprosthesis?* At this point, it is not possible to say whether memristors are the way to go. The commercial memristors used are still difficult to work with and not particularly resilient or long-lasting. But, in theory, the potential is there. It is worth mentioning that this study used a single memristor for feature detection. Similarly to an isolated percetron (the artificial neuron of ANNs) that can only classify two linearly separable classes, a single memristor can only do binary detection of simple neuronal patterns. To discern more complex neuronal patterns, an array of memristors would be necessary, just like an ANN (a network of perceptrons) is needed for complex non-linear classification tasks. Still, even if the memristors can be made reliable and robust, it is not totally clear how such a memristor array would be trained. Implementing backpropagation in memristor networks is still an active field of research and, to the best of our knowledge, only tested in computer simulations (Hasan and Taha, 2014; Zhang et al., 2017; Krestinskaya et al., 2018b,a).

## 6.3 Future perspectives

As we have built upon the work of others, I anticipate and encourage others to further develop the ideas presented here. Each of the chapters has a lot to build from so I will present what I see as future perspectives organised, once again, by chapter.

#### Chapter 3 - Future direction of aDFC and open questions on network dynamics

In my opinion, this study opened some interesting research opportunities. We introduced adaptive DFC as a possible way of disrupting pathological oscillations without creating new ones. The adaptive component was in charge of adjusting the time constant of the control algorithm, but there are other relevant parameters that could be made adaptive, such as control gain, or choice of stimulation electrode choice. However, we must point out that opening up the degrees of freedom of the adaptive controller will make the algorithm significantly more complex and difficult to auto-tune. Also, it could be made multi-output, if it controls not only the stimulation frequency but also the stimulation amplitude. It would be interesting to test the control capabilities of such an algorithm.

Still, it is important to recognise that, no matter how adaptive DFC is made, if it is based on electrical stimulation, it will always suffer from one important handicap: electrical stimulation is always excitatory, meaning that it cannot provide a negative control signal. Indeed, in our algorithm, when the control signal was negative, we did not stimulate (there is no such thing as negative stimulation frequency). However, with optogenetics, the actuation signal can be made positive (excitatory stimulation) or negative (inhibitory stimulation), so we expect that testing DFC with optogenetics will provide interesting results.

On the computational front, we pointed out the discrepancy between the efficiency of the previous *in silico* studies and our *in vitro* results using DFC. We posited a possible explanation on the fact

that, so far, computational models of spiking networks cannot support multiple oscillation periodicities, whereas *in vitro* (and, of course *in vivo*) networks can. Developing computational models of neuronal networks capable of sustaining more dynamic behaviour would have remarkable importance for the development of future control algorithms and validation of the currently proposed ones.

We also presented the notion of controllable and uncontrollable networks in neuronal cultures. We showed *in silico* that this distinction could (at least) be due to excitatory/inhibitory balance. It would be a compelling study to test this hypothesis *in vitro*. The excitatory/inhibitory balance could be modulated using drugs or by building networks with different percentages of excitatory and inhibitory neurons. The latter could be made using induced pluripotent stem cells (iPSC) (Wang et al., 2022, 2023). Then these different cultures would be submitted to different stimulation protocols to check whether the controllable networks do indeed match those that have an intermediate balance of excitation and inhibition.

We also detected one network that presented cyclical transitions between well-defined dynamical states. This was very intriguing on its own. What are the network requirements to achieve such dynamics? One possibility is the connectivity within the network, which we cannot control in our experiments. I wonder if having a neuronal network built with multiple weakly connected subpopulations would increase the probability of seeing this type of complex behaviour. This could be achieved by combining MEAs with microfluidic technology (Lopes et al., 2018).

Arguably, the most important question regarding the future perspectives of this study is whether or not aDFC should be tested *in vivo*. I would say yes. The activity of living brains is inherently multi-stable and is certainly capable of sustaining asynchronous activity. It is possible that *in vivo* results resemble the remarkable effects seen in the multi-stable network, which was the only one capable of sustaining asynchronous activity.

**Chapter 4 - Moving to *in vitro* tests and improving the computer model**

The deep learning algorithm presented in Chapter4 is ready to be tested *in vitro*, using retinal explants over MEAs. These tests will prove the true potential of the encoder-decoder framework to learn the forward and inverse models of retinal response to multipolar stimulation. Just like we warned for the previous research question, *in silico* studies may get unrealistic results due to the oversimplification of the neuronal model. Therefore, testing the algorithm with a real retina will be a crucial milestone. We are confident that the algorithm may perform well with real neurons, even if the forward and inverse models are more complex *in vitro* (which they most likely are), because the method is based on ANNs, which are general function approximation algorithms. As long as the underlying functions governing retinal activation remain constant over time, the ANN should be able to learn them. And, since the *in vitro* models use intact retina explants, success in the *in vitro* experiments sets a promising path for the *in vivo* translation.

If the *in vitro* results are not successful, the algorithm will need to be improved – and there is, indeed, still margin for improvement. One possibility for failure is that the non-linear response *in vitro* is considerably more complex and cannot be learned by the current ANNs. If that is the case, my suggestion is to go back to the *in silico* environment and create an even model complex retinal model, in particular regarding the non-linear response to stimulation. Here, we assumed that the activation function was modelled by symmetric sigmoid functions. However, *in vitro* tests show that the non-linearity may be more complex, due to second-order effects like indirect ganglion cell activation by neighbouring cells (Maturana et al., 2018). If the *in silico* results fail, that means that the architecture of the ANNs needs to be adjusted. Once this step is overcome in computer simulations, then the natural trajectory of moving to back *in vitro* and then proceed to *in vivo* studies.

**Chapter 5 - Improving memristor technology and testing memristive arrays *in vitro***

The last chapter also opened a lot of research avenues worth pursuing. We first point out that, based on our experience using commercial memristive devices, the robustness of the devices needs to be improved for these to be reliable building blocks of any useful processor. Still, if these future processors are based on large memristive arrays, then redundancy will play an important role (just like it does in brain networks), meaning that each individual element does not need to behave like an ideal memristor.

Regarding the specific *in vitro* setup of this study, a promising future step will be repeating the experiments without any interfacing elements – no interface board, no DPS, no conditioning electrical circuit. Direct connection between neuron, electrode and memristor. This will be no trivial task, but if it proves successful, that means that memristors can indeed be used as microscopic *in-situ* processors of neuronal activity, and ideally powered by it. Future implants would require only additional hardware and power supply for the stimulus generator.

Ultimately, the full potential of memristors for neurostimulation applications will only be unlocked when memristor arrays can be realised into effective and trainable processors. Having memristor networks will be crucial for detecting relevant neuronal activity patterns (Zhang et al., 2021; Li et al., 2022, 2023). Once such networks are developed, they can be tested *in vitro* using a setup analogous to the one presented in our study. If this second iteration of *in vitro* experiments is successful, then *in vivo* experiments are justified. As a first *in vivo* use case we envision an ON-OFF controller for seizure detection and responsive stimulation, where the processing is performed in real-time with a memristive network.

# References


Abelson, J. L., Curtis, G. C., Sagher, O., Albucher, R. C., Harrigan, M., Taylor, S. F., Martis, B., and Giordani, B. (2005). Deep brain stimulation for refractory obsessive-compulsive disorder. *Biological Psychiatry*, 57(5):510–516.

Adeli, H., Zhou, Z., and Dadmehr, N. (2003). Analysis of eeg records in an epileptic patient using wavelet transform. *Journal of Neuroscience Methods*, 123(1):69–87.

Aggarwal, S. and Chugh, N. (2022). Review of machine learning techniques for EEG based brain computer interface. *Archives of Computational Methods in Engineering*, 29(5):3001–3020.

Aguera, P.-E., Jerbi, K., Caclin, A., and Bertrand, O. (2011). Elan: A software package for analysis and visualization of meg, eeg, and lfp signals. *Computational Intelligence and Neuroscience*, 2011:1–11.

Ahuja, A. K., Behrend, M. R., Kuroda, M., Humayun, M. S., and Weiland, J. D. (2008). An in vitro model of a retinal prosthesis. *IEEE Transactions on Biomedical Engineering*, 55(6):1744–1753.

Akin, M., Arserim, M., Kiymik, M., and Turkoglu, I. (2001). A new approach for diagnosing epilepsy by using wavelet transform and neural networks. In *2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 2, pages 1596–1599. IEEE.

Alcala-Zermeno, J. L., Gregg, N. M., Starnes, K., Mandrekar, J. N., Gompel, J. J. V., Miller, K., Worrell, G., and Lundstrom, B. N. (2022). Invasive neuromodulation for epilepsy: Comparison of multiple approaches from a single center. *Epilepsy & Behavior*, 137:108951.

Anderson, T. R., Hu, B., Iremonger, K., and Kiss, Z. H. T. (2006). Selective attenuation of afferent synaptic transmission as a mechanism of thalamic deep brain stimulation-induced tremor arrest. *The Journal of Neuroscience*, 26(3):841–850.

Andrade, D. M., Zumsteg, D., Hamani, C., Hodaie, M., Sarkissian, S., Lozano, A. M., and Wennberg, R. A. (2006). Long-term follow-up of patients with thalamic deep brain stimulation for epilepsy. *Neurology*, 66(10):1571–1573.

Arlotti, M., Marceglia, S., Foffani, G., Volkmann, J., Lozano, A. M., Moro, E., Cogiamanian, F., Prenassi, M., Bocci, T., Cortese, F., Rampini, P., Barbieri, S., and Priori, A. (2018). Eight-hours adaptive deep brain stimulation in patients with parkinson disease. *Neurology*, 90(11):e971–e976.

Asplund, M., von Holst, H., and Inganäs, O. (2008). Composite biomolecule/pedot materials for neural electrodes. *Biointerphases*, 3(3):83–93.

Avoli, M., Louvel, J., Pumain, R., and Köhling, R. (2005). Cellular and molecular mechanisms of epilepsy in the human brain. *Progress in Neurobiology*, 77(3):166–200.

Azodi-Avval, R. and Gharabaghi, A. (2015). Phase-dependent modulation as a novel approach for therapeutic brain stimulation. *Frontiers in Computational Neuroscience*, 9.

Bakkum, D. J., Frey, U., Radivojevic, M., Russell, T. L., Müller, J., Fiscella, M., Takahashi, H., and Hierlemann, A. (2013). Tracking axonal action potential propagation on a high-density microelectrode array across hundreds of sites. *Nature Communications*, 4(1).

Bakkum, D. J., Shkolnik, A. C., Ben-Ary, G., Gamblen, P., DeMarse, T. B., and Potter, S. M. (2004). *Removing Some 'A' from AI: Embodied Cultured Networks*, page 130–145. Springer Berlin Heidelberg.

Baldermann, J. C., Schüller, T., Huys, D., Becker, I., Timmermann, L., Jessen, F., Visser-Vandewalle, V., and Kuhn, J. (2016). Deep brain stimulation for tourette-syndrome: A systematic review and meta-analysis. *Brain Stimulation*, 9(2):296–304.

Barbera, S. L., Vuillaume, D., and Alibart, F. (2015). Filamentary switching: Synaptic plasticity through device volatility. *ACS Nano*, 9(1):941–949.

Bartholow, R. (1874). Art. i.–experimental investigations into the functions of the human brain. *The American Journal of the Medical Sciences (1827-1924)*, (134):305.

Batista, C., Lopes, S., Viana, R., and Batista, A. (2010). Delayed feedback control of bursting synchronization in a scale-free neuronal network. *Neural Networks*, 23(1):114–124.

Belkacem, A. N., Jamil, N., Khalid, S., and Alnajjar, F. (2023). On closed-loop brain stimulation systems for improving the quality of life of patients with neurological disorders. *Frontiers in Human Neuroscience*, 17.

Ben-Ari, Y. (2001). Developing networks play a similar melody. *Trends in Neurosciences*, 24(6):353–360.

Ben-Ari, Y. (2002). Excitatory actions of gaba during development: the nature of the nurture. *Nature Reviews Neuroscience*, 3(9):728–739.

Benabid, A., Pollak, P., Louveau, A., Henry, S., and de Rougemont, J. (1987). Combined (thalamotomy and stimulation) stereotactic surgery of the vim thalamic nucleus for bilateral parkinson disease. *Stereotactic and Functional Neurosurgery*, 50(1–6):344–346.

Benazzouz, A., Piallat, B., Pollak, P., and Benabid, A.-L. (1995). Responses of substantia nigra pars reticulata and globus pallidus complex to high frequency stimulation of the subthalamic nucleus in rats: electrophysiological data. *Neuroscience Letters*, 189(2):77–80.

Benedetti-Isaac, J. C., Torres-Zambrano, M., Vargas-Toscano, A., Perea-Castro, E., Alcalá-Cerra, G., Furlanetti, L. L., Reithmeier, T., Tierney, T. S., Anastasopoulos, C., Fonoff, E. T., and Lopez, W. O. C. (2015). Seizure frequency reduction after posteromedial hypothalamus deep brain stimulation in drug-resistant epilepsy associated with intractable aggressive behavior. *Epilepsia*, 56(7):1152–1161.

Berger, H. (1929). Über das elektroenkephalogramm des menschen. *Archiv für psychiatrie und nervenkrankheiten*, 87(1):527–570.

Bertucci, C., Koppes, R., Dumont, C., and Koppes, A. (2019). Neural responses to electrical stimulation in 2d and 3d in vitro environments. *Brain Research Bulletin*, 152:265–284.

Beverlin II, B. and Netoff, T. I. (2013). Dynamic control of modeled tonic-clonic seizure states with closed-loop stimulation. *Frontiers in neural circuits*, 6:126.

Beyeler, M., Nanduri, D., Weiland, J. D., Rokem, A., Boynton, G. M., and Fine, I. (2019). A model of ganglion axon pathways accounts for percepts elicited by retinal implants. *Scientific Reports*, 9(1).

Boccard, S. G., Pereira, E. A., and Aziz, T. Z. (2015). Deep brain stimulation for chronic pain. *Journal of Clinical Neuroscience*, 22(10):1537–1543.

Boccard, S. G., Pereira, E. A., Moir, L., Aziz, T. Z., and Green, A. L. (2013). Long-term outcomes of deep brain stimulation for neuropathic pain. *Neurosurgery*, 72(2):221–231.

Boehler, C., Stieglitz, T., and Asplund, M. (2015). Nanostructured platinum grass enables superior impedance reduction for neural microelectrodes. *Biomaterials*, 67:346–353.

Boehler, C., Vieira, D. M., Egert, U., and Asplund, M. (2020). Nanopt—a nanostructured electrode coating for neural recording and microstimulation. *ACS Applied Materials & Interfaces*, 12(13):14855–14865.

Bojak, I., Oostendorp, T. F., Reid, A. T., and Kötter, R. (2010). Connecting mean field models of neural activity to eeg and fmri data. *Brain topography*, 23:139–149.

Bolin, M. H., Svennersten, K., Wang, X., Chronakis, I. S., Richter-Dahlfors, A., Jager, E. W., and Berggren, M. (2009). Nano-fiber scaffold electrodes based on pedot for cell stimulation. *Sensors and Actuators B: Chemical*, 142(2):451–456.

Boraud, T., Bezard, E., Bioulac, B., and Gross, C. (1996). High frequency stimulation of the internal globus pallidus (gpi) simultaneously improves parkinsonian symptoms and reduces the firing frequency of gpi neurons in the mptp-treated monkey. *Neuroscience Letters*, 215(1):17–20.

Borkholder, D., Bao, J., Maluf, N., Perl, E., and Kovacs, G. (1997). Microelectrode arrays for stimulation of neural slice preparations. *Journal of Neuroscience Methods*, 77(1):61–66.

Boyden, E. S., Zhang, F., Bamberg, E., and Nagel (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9):1263–1268.

Braun, D. and Fromherz, P. (1997). Fluorescence interference-contrast microscopy of cell adhesion on oxidized silicon. *Applied Physics A: Materials Science &; Processing*, 65(4–5):341–348.

Brette, R. and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, 94(5):3637–3642.

Broccard, F. D., Joshi, S., Wang, J., and Cauwenberghs, G. (2017). Neuromorphic neural interfaces: from neurophysiological inspiration to biohybrid coupling with nervous systems. *Journal of Neural Engineering*, 14(4):041002.

Bronstein, J. M., Tagliati, M., Alterman, R. L., Lozano, A. M., Volkmann, J., Stefani, A., Horak, F. B., Okun, M. S., Foote, K. D., Krack, P., Pahwa, R., Henderson, J. M., Hariz, M. I., Bakay, R. A., Rezai, A., Marks, W. J., Moro, E., Vitek, J. L., Weaver, F. M., Gross, R. E., and DeLong, M. R. (2011). Deep brain stimulation for parkinson disease. *Archives of Neurology*, 68(2).

Brückerhoff-Plückelmann, F., Feldmann, J., Wright, C. D., Bhaskaran, H., and Pernice, W. H. P. (2021). Chalcogenide phase-change devices for neuromorphic photonic computing. *Journal of Applied Physics*, 129(15):151103.

Brüggemann, D., Wolfrum, B., Maybeck, V., Mourzina, Y., Jansen, M., and Offenhäusser, A. (2011). Nanostructured gold microelectrodes for extracellular recording from electrogenic cells. *Nanotechnology*, 22(26):265104.

Buckwell, M., Montesi, L., Hudziak, S., Mehonic, A., and Kenyon, A. J. (2015). Conductance tomography of conductive filaments in intrinsic silicon-rich silica RRAM. *Nanoscale*, 7(43):18030–18035.

Buehlmann, A. and Deco, G. (2010). Optimal information transfer in the cortex through synchronization. *PLoS Computational Biology*, 6(9):e1000934.

Burchiel, K. J., Anderson, V. C., Favre, J., and Hammerstad, J. P. (1999). Comparison of pallidal and subthalamic nucleus deep brain stimulation for advanced parkinson's disease: results of a randomized, blinded pilot study. *Neurosurgery*, 45(6):1375.

Cacabelos, R. (2017). Parkinson's disease: From pathogenesis to pharmacogenomics. *International Journal of Molecular Sciences*, 18(3):551.

Cagnan, H., Brittain, J.-S., Little, S., Foltynie, T., Limousin, P., Zrinzo, L., Hariz, M., Joint, C., Fitzgerald, J., Green, A. L., Aziz, T., and Brown, P. (2013). Phase dependent modulation of tremor amplitude in essential tremor through thalamic stimulation. *Brain*, 136(10):3062–3075.

Cagnan, H., Pedrosa, D., Little, S., Pogosyan, A., Cheeran, B., Aziz, T., Green, A., Fitzgerald, J., Foltynie, T., Limousin, P., Zrinzo, L., Hariz, M., Friston, K. J., Denison, T., and Brown, P. (2016). Stimulating at the right time: phase-specific deep brain stimulation. *Brain*, 140(1):132–145.

Callegari, F., Brofiga, M., Poggio, F., and Massobrio, P. (2022). Stimulus-evoked activity modulation of in vitro engineered cortical and hippocampal networks. *Micromachines*, 13(8):1212.

Campbell, K. A. (2017). Self-directed channel memristor for high temperature operation. *Microelectronics Journal*, 59:10–14.

Campbell, K. A. (2019). The self-directed channel memristor: Operational dependence on the metal-chalcogenide layer. In *Handbook of Memristor Networks*, pages 815–842. Springer International Publishing.

Campbell, P. K., Jones, K. E., Huber, R. J., Horch, K. W., and Normann, R. A. (1991). A silicon-based, three-dimensional neural interface: manufacturing processes for an intracortical electrode array. *IEEE Transactions on Biomedical Engineering*, 38(8):758–768.

Celano, U., Goux, L., Degraeve, R., Fantini, A., Richard, O., Bender, H., Jurczak, M., and Vandervorst, W. (2015). Imaging the three-dimensional conductive channel in filamentary-based oxide resistive switching memory. *Nano Letters*, 15(12):7970–7975.

Cellot, G., Lagonegro, P., Tarabella, G., Scaini, D., Fabbri, F., Iannotta, S., Prato, M., Salviati, G., and Ballerini, L. (2016). Pedot:pss interfaces support the development of neuronal synaptic networks with reduced neuroglia response in vitro. *Frontiers in Neuroscience*, 9.

Cerina, M., Piastra, M. C., and Frega, M. (2023). The potential of in vitro neuronal networks cultured on micro electrode arrays for biomedical research. *Progress in Biomedical Engineering*, 5(3):032002.

Cerletti, U. (1950). Old and new information about electroshock. *American Journal of Psychiatry*, 107(2):87–94.

Chiappalone, M., Massobrio, P., and Martinoia, S. (2008). Network plasticity in cortical assemblies. *European Journal of Neuroscience*, 28(1):221–237.

Chiken, S. and Nambu, A. (2015). Mechanism of deep brain stimulation. *The Neuroscientist*, 22(3):313–322.

Chiolerio, A., Chiappalone, M., Ariano, P., and Bocchini, S. (2017). Coupling resistive switching devices with neurons: State of the art and perspectives. *Frontiers in Neuroscience*, 11.

Christini, D. J. and Collins, J. J. (1995). Controlling neuronal noise using chaos control.

Chua, L. (1971). Memristor-the missing circuit element. *IEEE Transactions on Circuit Theory*, 18(5):507–519.

Cohen, E., Ivenshitz, M., Amor-Baroukh, V., Greenberger, V., and Segal, M. (2008). Determinants of spontaneous activity in networks of cultured hippocampus. *Brain Research*, 1235:21–30.

Colpan, M. E., Li, Y., Dwyer, J., and Mogul, D. J. (2007). Proportional feedback stimulation for seizure control in rats. *Epilepsia*, 48(8):1594–1603.

Compte, A. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, 10(9):910–923.

Compte, A., Sanchez-Vives, M. V., McCormick, D. A., and Wang, X.-J. (2003). Cellular and network mechanisms of slow oscillatory activity (<1 hz) and wave propagations in a cortical network model. *Journal of Neurophysiology*, 89(5):2707–2725.

Connolly, B. S. and Lang, A. E. (2014). Pharmacological treatment of parkinson disease. *JAMA*, 311(16):1670.

Coulombe, M.-A., Elkaim, L. M., Alotaibi, N. M., Gorman, D. A., Weil, A. G., Fallah, A., Kalia, S. K., Lipsman, N., Lozano, A. M., and Ibrahim, G. M. (2019). Deep brain stimulation for gilles de la tourette syndrome in children and youth: a meta-analysis with individual participant data. *Journal of Neurosurgery: Pediatrics*, 23(2):236–246.

Couto, J. (2015). On the firing rate dependency of the phase response curve of rat purkinje neurons in vitro. *PLOS Computational Biology*, 11(3).

Craik, A., He, Y., and Contreras-Vidal, J. L. (2019). Deep learning for electroencephalogram (EEG) classification tasks: a review. *Journal of Neural Engineering*, 16(3):031001.

da Silva, F. L., Blanes, W., Kalitzin, S. N., Parra, J., Suffczynski, P., and Velis, D. N. (2003). Epilepsies as dynamical diseases of brain systems: Basic models of the transition between normal and epileptic activity. *Epilepsia*, 44(s12):72–83.

Daneshzand, M., Faezipour, M., and Barkana, B. (2018a). Towards frequency adaptation for delayed feedback deep brain stimulations. *Neural Regeneration Research*, 13(3):408.

Daneshzand, M., Faezipour, M., and Barkana, B. D. (2018b). Robust desynchronization of parkinson's disease pathological oscillations by frequency modulation of delayed feedback deep brain stimulation. *PLOS ONE*, 13(11):e0207761.

Danzl, P., Hespanha, J., and Moehlis, J. (2009). Event-based minimum-time control of oscillatory neuron models. *Biological Cybernetics*, 101(5-6):387–399.

Daoud, H. and Bayoumi, M. A. (2019). Efficient epileptic seizure prediction based on deep learning. *IEEE Transactions on Biomedical Circuits and Systems*, 13(5):804–813.

Dasanayake, I. and Li, J.-S. (2011). Optimal design of minimum-power stimuli for phase models of neuron oscillators. *Physical Review E*, 83(6).

David, O. and Friston, K. J. (2003). A neural mass model for meg/eeg:. *NeuroImage*, 20(3):1743–1755.

David-Pur, M., Bareket-Keren, L., Beit-Yaakov, G., Raz-Prag, D., and Hanein, Y. (2013). All-carbon-nanotube flexible multi-electrode array for neuronal recording and stimulation. *Biomedical Microdevices*, 16(1):43–53.

de Castro, D. L., Aroso, M., Aguiar, A. P., Grayden, D. B., and Aguiar, P. (2024). Disrupting abnormal neuronal oscillations with adaptive delayed feedback control. *Elife*, 13:e89151.

de Castro Medeiros, D. and Moraes, M. F. D. (2014). Focus on desynchronization rather than excitability: A new strategy for intraencephalic electrical stimulation. *Epilepsy and Behavior*, 38:32–36.

DeMarse, T. B. and Dockendorf, K. P. (2005). Adaptive flight control with living neuronal networks on microelectrode arrays. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 3, pages 1548–1551. IEEE.

DeMarse, T. B., Wagenaar, D. A., Blau, A. W., and Potter, S. M. (2001). *Autonomous Robots*, 11(3):305–310.

Deniau, J., Degos, B., Bosch, C., and Maurice, N. (2010). Deep brain stimulation mechanisms: beyond the concept of local functional inhibition. *European Journal of Neuroscience*, 32(7):1080–1091.

Denison, T. and Morrell, M. J. (2022). Neuromodulation in 2035: The neurology future forecasting series. *Neurology*, 98(2):65–72.

Destexhe, A. (2009). Self-sustained asynchronous irregular states and up–down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *Journal of Computational Neuroscience*, 27(3):493–506.

Devor, A., Bandettini, P. A., Boas, D. A., Bower, J. M., Buxton, R. B., Cohen, L. B., Dale, A. M., Einevoll, G. T., Fox, P. T., Franceschini, M. A., et al. (2013). The challenge of connecting the dots in the brain. *Neuron*, 80(2):270–274.

Dias, C., Castro, D., Aroso, M., Ventura, J., and Aguiar, P. (2022). Memristor-based neuromodulation device for real-time monitoring and adaptive control of neuronal populations. *ACS Applied Electronic Materials*, 4(5):2380–2387.

Dobelle, W. H. and Mladejovsky, M. G. (1974). Phosphenes produced by electrical stimulation of human occipital cortex, and their application to the development of a prosthesis for the blind. *The Journal of Physiology*, 243(2):553–576.

Dorman, M. F., Smith, L. M., Dankowski, K., McCandless, G., and Parkin, J. L. (1992). Long-term measures of electrode impedance and auditory thresholds for the ineraid cochlear implant. *Journal of Speech, Language, and Hearing Research*, 35(5):1126–1130.

Doty, R. W. (1998). The five mysteries of the mind, and their consequences. *Neuropsychologia*, 36(10):1069–1076.

Dovzhenok, A., Park, C., Worth, R. M., and Rubchinsky, L. L. (2013). Failure of delayed feedback deep brain stimulation for intermittent pathological synchronization in parkinson's disease. *PLoS One*, 8(3):e58264.

Dranias, M. R., Ju, H., Rajaram, E., and VanDongen, A. M. J. (2013). Short-term memory in networks of dissociated cortical neurons. *The Journal of Neuroscience*, 33(5):1940–1953.

Eckmiller, R., Baruth, O., and Borbe, S. (2007). *Towards Learning Retina Implants: How to Induce Visual Percepts with Electrical Stimulation Patterns*, page 787–791. Springer Netherlands.

Eckmiller, R., Neumann, D., and Baruth, O. (2005). Tunable retina encoders for retina implants: why and how. *Journal of Neural Engineering*, 2(1):S91–S104.

Eickenscheidt, M., Jenkner, M., Thewes, R., Fromherz, P., and Zeck, G. (2012). Electrical stimulation of retinal neurons in epiretinal and subretinal configuration using a multicapacitor array. *Journal of Neurophysiology*, 107(10):2742–2755.

El Houssaini, K., Bernard, C., and Jirsa, V. K. (2020). The epileptor model: A systematic mathematical analysis linked to the dynamics of seizures, refractory status epilepticus, and depolarization block. *eneuro*, 7(2):ENEURO.0485–18.2019.

Erickson-Davis, C. and Korzybska, H. (2021). What do blind people "see" with retinal prostheses? observations and qualitative reports of epiretinal implant users. *PLoS One*, 16(2):e0229189.

Esler, T. B., Maturana, M. I., Kerr, R. R., Grayden, D. B., Burkitt, A. N., and Meffin, H. (2018). Biophysical basis of the linear electrical receptive fields of retinal ganglion cells. *Journal of Neural Engineering*, 15(5):055001.

Faramarzi, S. and Netoff, T. I. (2021). Closed-loop neuromodulation for clustering neuronal populations. *Journal of Neurophysiology*, 125(1):248–255.

Farrell, J. S., Nguyen, Q.-A., and Soltesz, I. (2019). Resolving the micro-macro disconnect to address core features of seizure networks. *Neuron*, 101(6):1016–1028.

Feigin, V. L., Nichols, E., Alam, T., Bannick, M. S., Beghi, E., Blake, N., Culpepper, W. J., Dorsey, E. R., Elbaz, A., Ellenbogen, R. G., Fisher, J. L., Fitzmaurice, C., Giussani, G., Glennie, L., James, S. L., Johnson, C. O., Kassebaum, N. J., Logroscino, G., Marin, B., Mountjoy-Venning, W. C., Nguyen, M., Ofori-Asenso, R., Patel, A. P., Piccininni, M., Roth, G. A., Steiner, T. J., Stovner, L. J., Szoeke, C. E. I., Theadom, A., Vollset, S. E., Wallin, M. T., Wright, C., Zunt, J. R., Abbasi, N., Abd-Allah, F., Abdelalim, A., Abdollahpour, I., Aboyans, V., Abraha, H. N., Acharya, D., Adamu, A. A., Adebayo, O. M., Adeoye, A. M., Adsuar, J. C., Afarideh, M., Agrawal, S., Ahmadi, A., Ahmed, M. B., Aichour, A. N., Aichour, I., Aichour, M. T. E.,

Akinyemi, R. O., Akseer, N., Al-Eyadhy, A., Salman, R. A.-S., Alahdab, F., Alene, K. A., Aljunid, S. M., Altirkawi, K., Alvis-Guzman, N., Anber, N. H., Antonio, C. A. T., Arabloo, J., Aremu, O., Ärnlöv, J., Asayesh, H., Asghar, R. J., Atalay, H. T., Awasthi, A., Quintanilla, B. P. A., Ayuk, T. B., Badawi, A., Banach, M., Banoub, J. A. M., Barboza, M. A., Barker-Collo, S. L., Bärnighausen, T. W., Baune, B. T., Bedi, N., Behzadifar, M., Behzadifar, M., Béjot, Y., Bekele, B. B., Belachew, A. B., Bennett, D. A., Bensenor, I. M., Berhane, A., Beuran, M., Bhattacharyya, K., Bhutta, Z. A., Biadgo, B., Bijani, A., Bililign, N., Sayeed, M. S. B., Blazes, C. K., Brayne, C., Butt, Z. A., Campos-Nonato, I. R., Cantu-Brito, C., Car, M., Cárdenas, R., Carrero, J. J., Carvalho, F., Castañeda-Orjuela, C. A., Castro, F., Catalá-López, F., Cerin, E., Chaiah, Y., Chang, J.-C., Chatziralli, I., Chiang, P. P.-C., Christensen, H., Christopher, D. J., Cooper, C., Cortesi, P. A., Costa, V. M., Criqui, M. H., Crowe, C. S., Damasceno, A. A. M., Daryani, A., la Cruz-Góngora, V. D., la Hoz, F. P. D., Leo, D. D., Demoz, G. T., Deribe, K., Dharmaratne, S. D., Diaz, D., Dinberu, M. T., Djalalinia, S., Doku, D. T., Dubey, M., Dubljanin, E., Duken, E. E., Edvardsson, D., El-Khatib, Z., Endres, M., Endries, A. Y., Eskandarieh, S., Esteghamati, A., Esteghamati, S., Farhadi, F., Faro, A., Farzadfar, F., Farzaei, M. H., Fatima, B., Fereshtehnejad, S.-M., Fernandes, E., Feyissa, G. T., Filip, I., Fischer, F., Fukumoto, T., Ganji, M., Gankpe, F. G., Garcia-Gordillo, M. A., Gebre, A. K., Gebremichael, T. G., Gelaw, B. K., Geleijnse, J. M., Geremew, D., Gezae, K. E., Ghasemi-Kasman, M., Gidey, M. Y., Gill, P. S., Gill, T. K., Girma, E. T., Gnedovskaya, E. V., Goulart, A. C., Grada, A., Grosso, G., Guo, Y., Gupta, R., Gupta, R., Haagsma, J. A., Hagos, T. B., Haj-Mirzaian, A., Haj-Mirzaian, A., Hamadeh, R. R., Hamidi, S., Hankey, G. J., Hao, Y., Haro, J. M., Hassankhani, H., Hassen, H. Y., Havmoeller, R., Hay, S. I., Hegazy, M. I., Heidari, B., Henok, A., Heydarpour, F., Hoang, C. L., Hole, M. K., Rad, E. H., Hosseini, S. M., Hu, G., Igumbor, E. U., Ilesanmi, O. S., Irvani, S. S. N., Islam, S. M. S., Jakovljevic, M., Javanbakht, M., Jha, R. P., Jobanputra, Y. B., Jonas, J. B., Jozwiak, J. J., Jürisson, M., Kahsay, A., Kalani, R., Kalkonde, Y., Kamil, T. A., Kanchan, T., Karami, M., Karch, A., Karimi, N., Kasaeian, A., Kassa, T. D., Kassa, Z. Y., Kaul, A., Kefale, A. T., Keiyoro, P. N., Khader, Y. S., Khafaie, M. A., Khalil, I. A., Khan, E. A., Khang, Y.-H., Khazaie, H., Kiadaliri, A. A., Kiirithio, D. N., Kim, A. S., Kim, D., Kim, Y.-E., Kim, Y. J., Kisa, A., Kokubo, Y., Koyanagi, A., Krishnamurthi, R. V., Defo, B. K., Bicer, B. K., Kumar, M., Lacey, B., Lafranconi, A., Lansingh, V. C., Latifi, A., Leshargie, C. T., Li, S., Liao, Y., Linn, S., Lo, W. D., Lopez, J. C. F., Lorkowski, S., Lotufo, P. A., Lucas, R. M., Lunevicius, R., Mackay, M. T., Mahotra, N. B., Majdan, M., Majdzadeh, R., Majeed, A., Malekzadeh, R., Malta, D. C., Manafi, N., Mansournia, M. A., Mantovani, L. G., März, W., Mashamba-Thompson, T. P., Massenburg, B. B., Mate, K. K. V., McAlinden, C., McGrath, J. J., Mehta, V., Meier, T., Meles, H. G., Melese, A., Memiah, P. T. N., Memish, Z. A., Mendoza, W., Mengistu, D. T., Mengistu, G., Meretoja, A., Meretoja, T. J., Mestrovic, T., Miazgowski, B., Miazgowski, T., Miller, T. R., Mini, G., Mirrakhimov, E. M., Moazen, B., Mohajer, B., Mezerji, N. M. G., Mohammadi, M., Mohammadi-Khanaposhtani, M., Mohammadibakhsh, R., Mohammadnia-Afrouzi, M., Mohammed, S., Mohebi, F., Mokdad, A. H., Monasta, L., Mondello, S., Moodley, Y., Moosazadeh, M., Moradi, G., Moradi-Lakeh, M., Moradinazar, M., Moraga, P., Velásquez, I. M., Morrison, S. D., Mousavi, S. M., Muhammed, O. S., Muruet, W., Musa, K. I., Mustafa, G., Naderi, M., Nagel, G., Naheed, A., Naik, G., Najafi, F., Nangia, V., Negoi, I., Negoi, R. I., Newton, C. R. J., Ngunjiri, J. W., Nguyen, C. T., Nguyen, L. H., Ningrum, D. N. A., Nirayo, Y. L., Nixon, M. R., Norrving, B., Noubiap, J. J., Shiadeh, M. N., Nyasulu, P. S., Ogah, O. S., Oh, I.-H., Olagunju, A. T., Olagunju, T. O., Olivares, P. R., Onwujekwe, O. E., Oren, E., Owolabi, M. O., PA, M., Pakpour, A. H., Pan, W.-H., Panda-Jonas, S., Pandian, J. D., Patel, S. K., Pereira, D. M., Petzold, M., Pillay, J. D., Piradov, M. A., Polanczyk, G. V., Polinder, S., Postma, M. J., Poulton, R., Poustchi, H., Prakash, S., Prakash, V., Qorbani, M., Radfar,

A., Rafay, A., Rafiei, A., Rahim, F., Rahimi-Movaghar, V., Rahman, M., Rahman, M. H. U., Rahman, M. A., Rajati, F., Ram, U., Ranta, A., Rawaf, D. L., Rawaf, S., Reinig, N., Reis, C., Renzaho, A. M. N., Resnikoff, S., Rezaeian, S., Rezai, M. S., González, C. M. R., Roberts, N. L. S., Roever, L., Ronfani, L., Roro, E. M., Roshandel, G., Rostami, A., Sabbagh, P., Sacco, R. L., Sachdev, P. S., Saddik, B., Safari, H., Safari-Faramani, R., Safi, S., Safiri, S., Sagar, R., Sahathevan, R., Sahebkar, A., Sahraian, M. A., Salamati, P., Zahabi, S. S., Salimi, Y., Samy, A. M., Sanabria, J., Santos, I. S., Milicevic, M. M. S., Sarrafzadegan, N., Sartorius, B., Sarvi, S., Sathian, B., Satpathy, M., Sawant, A. R., Sawhney, M., Schneider, I. J. C., Schöttker, B., Schwebel, D. C., Seedat, S., Sepanlou, S. G., Shabaninejad, H., Shafieesabet, A., Shaikh, M. A., Shakir, R. A., Shams-Beyranvand, M., Shamsizadeh, M., Sharif, M., Sharif-Alhoseini, M., She, J., Sheikh, A., Sheth, K. N., Shigematsu, M., Shiri, R., Shirkoohi, R., Shiue, I., Siabani, S., Siddiqi, T. J., Sigfusdottir, I. D., Sigurvinsdottir, R., Silberberg, D. H., Silva, J. P., Silveira, D. G. A., Singh, J. A., Sinha, D. N., Skiadaresi, E., Smith, M., Sobaih, B. H., Sobhani, S., Soofi, M., Soyiri, I. N., Sposato, L. A., Stein, D. J., Stein, M. B., Stokes, M. A., Sufiyan, M. B., Sykes, B. L., Sylaja, P., Tabarés-Seisdedos, R., Ao, B. J. T., Tehrani-Banihashemi, A., Temsah, M.-H., Temsah, O., Thakur, J. S., Thrift, A. G., Topor-Madry, R., Tortajada-Girbés, M., Tovani-Palone, M. R., Tran, B. X., Tran, K. B., Truelsen, T. C., Tsadik, A. G., Car, L. T., Ukwaja, K. N., Ullah, I., Usman, M. S., Uthman, O. A., Valdez, P. R., Vasankari, T. J., Vasanthan, R., Veisani, Y., Venketasubramanian, N., Violante, F. S., Vlassov, V., Vosoughi, K., Vu, G. T., Vujcic, I. S., Wagnew, F. S., Waheed, Y., Wang, Y.-P., Weiderpass, E., Weiss, J., Whiteford, H. A., Wijeratne, T., Winkler, A. S., Wiysonge, C. S., Wolfe, C. D. A., Xu, G., Yadollahpour, A., Yamada, T., Yano, Y., Yaseri, M., Yatsuya, H., Yimer, E. M., Yip, P., Yisma, E., Yonemoto, N., Yousefifard, M., Yu, C., Zaidi, Z., Zaman, S. B., Zamani, M., Zandian, H., Zare, Z., Zhang, Y., Zodpey, S., Naghavi, M., Murray, C. J. L., and Vos, T. (2019). Global, regional, and national burden of neurological disorders, 1990–2016: a systematic analysis for the global burden of disease study 2016. *The Lancet Neurology*, 18(5):459–480.

Feigin, V. L., Vos, T., Nichols, E., Owolabi, M. O., Carroll, W. M., Dichgans, M., Deuschl, G., Parmar, P., Brainin, M., and Murray, C. (2020). The global burden of neurological disorders: translating evidence into policy. *The Lancet Neurology*, 19(3):255–265.

Feinerman, O., Rotem, A., and Moses, E. (2008). Reliable neuronal logic devices from patterned hippocampal cultures. *Nature Physics*, 4(12):967–973.

Feng, X.-J., Shea-Brown, E., Greenwald, B., Kosut, R., and Rabitz, H. (2007). Optimal deep brain stimulation of the subthalamic nucleus - a computational study. *Journal of Computational Neuroscience*, 23(3):265–282.

Fernández, E., Alfaro, A., Soto-Sánchez, C., Gonzalez-Lopez, P., Lozano, A. M., Peña, S., Grima, M. D., Rodil, A., Gómez, B., Chen, X., Roelfsema, P. R., Rolston, J. D., Davis, T. S., and Normann, R. A. (2021). Visual percepts evoked with an intracortical 96-channel microelectrode array inserted in human occipital cortex. *Journal of Clinical Investigation*, 131(23).

Figee, M., Riva-Posse, P., Choi, K. S., Bederson, L., Mayberg, H. S., and Kopell, B. H. (2022). Deep brain stimulation for depression. *Neurotherapeutics*, 19(4):1229–1245.

Fisher, R. S., van Emde Boas, W., Blume, W., Elger, C., Genton, P., Lee, P., and Engel, J. (2005). Epileptic seizures and epilepsy: Definitions proposed by the international league against epilepsy (ILAE) and the international bureau for epilepsy (IBE). *Epilepsia*, 46(4):470–472.

Fisher, R. S. and Velasco, A. L. (2014). Electrical brain stimulation for epilepsy. *Nature Reviews Neurology*, 10(5):261–270.

Fleming, J. E., Dunn, E., and Lowery, M. M. (2020a). Simulation of closed-loop deep brain stimulation control schemes for suppression of pathological beta oscillations in parkinson's disease. *Frontiers in Neuroscience*, 14.

Fleming, J. E., Orłowski, J., Lowery, M. M., and Chaillet, A. (2020b). Self-tuning deep brain stimulation controller for suppression of beta oscillations: Analytical derivation and numerical validation. *Frontiers in Neuroscience*, 14.

Flora, E. D., Perera, C. L., Cameron, A. L., and Maddern, G. J. (2010). Deep brain stimulation for essential tremor: A systematic review. *Movement Disorders*, 25(11):1550–1559.

Florence, G., Sameshima, K., Fonoff, E. T., and Hamani, C. (2015). Deep brain stimulation. *The Neuroscientist*, 22(4):332–345.

Fu, T., Liu, X., Gao, H., Ward, J. E., Liu, X., Yin, B., Wang, Z., Zhuo, Y., Walker, D. J. F., Yang, J. J., Chen, J., Lovley, D. R., and Yao, J. (2020). Bioinspired bio-voltage memristors. *Nature Communications*, 11(1).

Furshpan, E. and Potter, D. (1989). Seizure-like activity and cellular damage in rat hippocampal neurons in cell culture. *Neuron*, 3(2):199–207.

Gao, Q., Naumann, M., Jovanov, I., Lesi, Grill, W. M., and Pajic, M. (2020). Model-based design of closed loop deep brain stimulation controller using reinforcement learning. In *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS)*, pages 108–118. IEEE.

Gao, Q., Schmidt, S. L., Kamaravelu, K., Turner, D. A., Grill, W. M., and Pajic, M. (2022). Offline policy evaluation for learning-based deep brain stimulation controllers. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS)*, pages 80–91. IEEE.

Gater, D., Iqbal, A., Davey, J., and Gale, E. (2013). Connecting spiking neurons to a spiking memristor network changes the memristor dynamics. In *2013 IEEE 20th International Conference on Electronics, Circuits, and Systems (ICECS)*. IEEE.

George, J. B., Abraham, G. M., Amrutur, B., and Sikdar, S. K. (2015). Robot navigation using neuro-electronic hybrid systems. In *2015 28th International Conference on VLSI Design*, pages 93–98. IEEE.

George, R., Chiappalone, M., Giugliano, M., Levi, T., Vassanelli, S., Partzsch, J., and Mayr, C. (2020). Plasticity and adaptation in neuromorphic biohybrid systems. *iScience*, 23(10):101589.

Ghaffari, D. H., Akwaboah, A. D., Mirzakhalili, E., and Weiland, J. D. (2021). Real-time optimization of retinal ganglion cell spatial activity in response to epiretinal stimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:2733–2741.

Giannakakis, E., Hutchings, F., Papasavvas, C. A., Han, C. E., Weber, B., Zhang, C., and Kaiser, M. (2020). Computational modelling of the long-term effects of brain stimulation on the local and global structural connectivity of epileptic patients. *PLOS ONE*, 15(2):e0221380.

Giordano, M., Caccavella, V. M., Zaed, I., Manzillo, L. F., Montano, N., Olivi, A., and Polli, F. M. (2020). Comparison between deep brain stimulation and magnetic resonance-guided focused ultrasound in the treatment of essential tremor: a systematic review and pooled analysis of functional outcomes. *Journal of Neurology, Neurosurgery & Psychiatry*, 91(12):1270–1278.

Gleixner, R. and Fromherz, P. (2006). The extracellular electrical resistivity in cell adhesion. *Biophysical Journal*, 90(7):2600–2611.

Gluckman, B. J., Nguyen, H., Weinstein, S. L., and Schiff, S. J. (2001). Adaptive electric field control of epileptic seizures. *The Journal of Neuroscience*, 21(2):590–600.

Gordon, K. A., Papsin, B. C., and Harrison, R. V. (2004). Toward a battery of behavioral and objective measures to achieve optimal cochlear implant stimulation levels in children. *Ear and Hearing*, 25(5):447–463.

Gorzelic, P., Schiff, S. J., and Sinha, A. (2013). Model-based rational feedback controller design for closed-loop deep brain stimulation of parkinson's disease. *Journal of Neural Engineering*, 10(2):026016.

Grado, L. L., Johnson, M. D., and Netoff, T. I. (2018). Bayesian adaptive dual control of deep brain stimulation in a computational model of parkinson's disease. *PLOS Computational Biology*, 14(12):e1006606.

Granley, J., Fauvel, T., Chalk, M., and Beyeler, M. (2023). Human-in-the-loop optimization for deep stimulus encoding in visual prostheses. *Advances in Neural Information Processing Systems*, 36:79376–79398.

Granley, J., Relic, L., and Beyeler, M. (2022). Hybrid neural autoencoders for stimulus encoding in visual and other sensory neuroprostheses. *Advances in Neural Information Processing Systems*, 35:22671–22685.

Grinenko, O., Li, J., Mosher, J. C., Wang, I. Z., Bulacio, J. C., Gonzalez-Martinez, J., Nair, D., Najm, I., Leahy, R. M., and Chauvel, P. (2017). A fingerprint of the epileptogenic zone in human epilepsies. *Brain*, 141(1):117–131.

Grumet, A. E., Wyatt, J. L., and Rizzo, J. F. (2000). Multi-electrode stimulation and recording in the isolated retina. *Journal of Neuroscience Methods*, 101(1):31–42.

Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., Medaglia, J. D., Vettel, J. M., Miller, M. B., Grafton, S. T., and Bassett, D. S. (2015). Controllability of structural brain networks. *Nature Communications*, 6(1).

Guez, A., Vincent, R. D., Avoli, M., and Pineau, J. (2008). Adaptive treatment of epilepsy via batch-mode reinforcement learning. In *AAAI*, volume 8, pages 1671–1678.

Gupta, I., Serb, A., Khiat, A., Zeitler, R., Vassanelli, S., and Prodromakis, T. (2016). Real-time encoding and compression of neuronal spikes by metal-oxide memristors. *Nature Communications*, 7(1).

Hafizovic, S., Heer, F., Ugniwenko, T., Frey, U., Blau, A., Ziegler, C., and Hierlemann, A. (2007). A cmos-based microelectrode array for interaction with neuronal cultures. *Journal of Neuroscience Methods*, 164(1):93–106.

Hai, A., Shappir, J., and Spira, M. E. (2010). Long-term, multisite, parallel, in-cell recording and stimulation by an array of extracellular microelectrodes. *Journal of Neurophysiology*, 104(1):559–568.

Halupka, K. J., Shivdasani, M. N., Cloherty, S. L., Grayden, D. B., Wong, Y. T., Burkitt, A. N., and Meffin, H. (2016). Prediction of cortical responses to simultaneous electrical stimulation of the retina. *Journal of Neural Engineering*, 14(1):016006.

Ham, D., Park, H., Hwang, S., and Kim, K. (2021). Neuromorphic electronics based on copying and pasting the brain. *Nature Electronics*, 4(9):635–644.

Ham, M. I., Bettencourt, L. M., McDaniel, F. D., and Gross, G. W. (2007). Spontaneous coordinated activity in cultured networks: Analysis of multiple ignition sites, primary circuits, and burst phase delay distributions. *Journal of Computational Neuroscience*, 24(3):346–357.

Hamad, A., Houssein, E. H., Hassanien, A. E., and Fahmy, A. A. (2016). Feature extraction of epilepsy eeg using discrete wavelet transform. In *2016 12th international computer engineering conference (ICENCO)*, pages 190–195. IEEE.

Hamani, C., Dubiela, F. P., Soares, J. C., Shin, D., Bittencourt, S., Covolan, L., Carlen, P. L., Laxton, A. W., Hodaie, M., Stone, S. S., Ha, Y., Hutchison, W. D., Lozano, A. M., Mello, L. E., and Oliveira, M. G. M. (2010). Anterior thalamus deep brain stimulation at high current impairs memory in rats. *Experimental Neurology*, 225(1):154–162.

Hameroff, S. R. (2006). The entwined mysteries of anesthesia and consciousness. *Anesthesiology*, 105(2):400–412.

Hammond, C., Bergman, H., and Brown, P. (2007). Pathological synchronization in parkinson's disease: networks, models and treatments. *Trends in Neurosciences*, 30(7):357–364.

Hariz, M. I. (2002). Complications of deep brain stimulation surgery. *Movement Disorders*, 17(S3):S162–S166.

Harrison, P. K., Sheridan, R. D., Green, A. C., Scott, I. R., and Tattersall, J. E. H. (2004). A guinea pig hippocampal slice model of organophosphate-induced seizure activity. *Journal of Pharmacology and Experimental Therapeutics*, 310(2):678–686.

Hasan, R. and Taha, T. M. (2014). Enabling back propagation training of memristor crossbar neuromorphic processors. In *2014 International Joint Conference on Neural Networks (IJCNN)*, pages 21–28. IEEE.

Hashimoto, T., Elder, C. M., Okun, M. S., Patrick, S. K., and Vitek, J. L. (2003). Stimulation of the subthalamic nucleus changes the firing pattern of pallidal neurons. *The Journal of Neuroscience*, 23(5):1916–1923.

Hauptmann, C., Popovych, O., and Tass, P. A. (2005). Effectively desynchronizing deep brain stimulation based on a coordinated delayed feedback stimulation via several sites: a computational study. *Biological Cybernetics*, 93(6):463–470.

Herrington, T. M., Cheng, J. J., and Eskandar, E. N. (2016). Mechanisms of deep brain stimulation. *Journal of Neurophysiology*, 115(1):19–38.

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., van der Smagt, P., and Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375.

Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., and Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442(7099):164–171.

Hodgkin, A. L. and Huxley, A. F. (1939). Action potentials recorded from inside a nerve fibre. *Nature*, 144(3651):710–711.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544.

Hofmann, F., Guenther, E., Hämmerle, H., Leibrock, C., Berezin, V., Bock, E., and Volkmer, H. (2004). Functional re-establishment of the perforant pathway in organotypic co-cultures on microelectrode arrays. *Brain Research*, 1017(1–2):184–196.

Holcman, D. and Tsodyks, M. (2006). The emergence of up and down states in cortical networks. *PLoS Computational Biology*, 2(3):e23.

Holt, A. B., Wilson, D., Shinn, M., Moehlis, J., and Netoff, T. I. (2016). Phasic burst stimulation: A closed-loop approach to tuning deep brain stimulation parameters for parkinson's disease. *PLOS Computational Biology*, 12(7):e1005011.

Hongo, Y., Takasu, K., Ikegaya, Y., Hasegawa, M., Sakaguchi, G., and Ogawa, K. (2015). Heterogeneous effects of antiepileptic drugs in an in vitro epilepsy model – a functional multineuron calcium imaging study. *European Journal of Neuroscience*, 42(2):1818–1829.

Hubble, J. P., Busenbark, K., Wilkinson, S., Penn, R. D., Lyons, K., and Koller, W. (1996). Deep brain stimulation for essential tremor. *Neurology*, 46(4):1150–1153.

Hughes, M. L., Werff, K. R. V., Brown, C. J., Abbas, P. J., Kelsay, D. M. R., Teagle, H. F. B., and Lowder, M. W. (2001). A longitudinal study of electrode impedance, the electrically evoked compound action potential, and behavioral measures in nucleus 24 cochlear implant users. *Ear and Hearing*, 22(6):471–486.

Humayun, M. S. (1996). Visual perception elicited by electrical stimulation of retina in blind humans. *Archives of Ophthalmology*, 114(1):40.

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572.

Jang, S., Jang, S., Lee, E.-H., Kang, M., Wang, G., and Kim, T.-W. (2018). Ultrathin conformable organic artificial synapse for wearable intelligent device applications. *ACS Applied Materials & Interfaces*, 11(1):1071–1080.

Jarosiewicz, B. and Morrell, M. (2020). The rns system: brain-responsive neurostimulation for the treatment of epilepsy. *Expert Review of Medical Devices*, 18(2):129–138.

Jimbo, Y., Kawana, A., Parodi, P., and Torre, V. (2000). The dynamics of a neuronal culture of dissociated cortical neurons of neonatal rats. *Biological Cybernetics*, 83(1):1–20.

Jirsa, V. K. (2009). Neural field dynamics with local and global connectivity and time delay. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1891):1131–1143.

Jirsa, V. K., Jantzen, K. J., Fuchs, A., and Kelso, J. S. (2002). Spatiotemporal forward solution of the eeg and meg using network modeling. *IEEE transactions on medical imaging*, 21(5):493–504.

Jirsa, V. K., Stacey, W. C., Quilichini, P. P., Ivanov, A. I., and Bernard, C. (2014). On the nature of seizure dynamics. *Brain*, 137(8):2210–2230.

Jiruska, P., de Curtis, M., Jefferys, J. G. R., Schevon, C. A., Schiff, S. J., and Schindler, K. (2013). Synchronization and desynchronization in epilepsy: controversies and hypotheses. *The Journal of Physiology*, 591(4):787–797.

Jo, S. H., Chang, T., Ebong, I., Bhadviya, B. B., Mazumder, P., and Lu, W. (2010). Nanoscale memristor device as synapse in neuromorphic systems. *Nano Letters*, 10(4):1297–1301.

John, S. E., Grayden, D. B., and Yanagisawa, T. (2019). The future potential of the stentrode. *Expert Review of Medical Devices*, 16(10):841–843.

Johnson, L. A., Nebeck, S. D., Muralidharan, A., Johnson, M. D., Baker, K. B., and Vitek, J. L. (2016). Closed-loop deep brain stimulation effects on parkinsonian motor symptoms in a non-human primate - is beta enough? *Brain Stimulation*, 9(6):892–896.

Juzekaeva, E., Nasretdinov, A., Battistoni, S., Berzina, T., Iannotta, S., Khazipov, R., Erokhin, V., and Mukhtarov, M. (2018). Coupling cortical neurons through electronic memristive synapse. *Advanced Materials Technologies*, 4(1):1800350.

Kagan, B. J., Kitchen, A. C., Tran, N. T., Habibollahi, F., Khajehnejad, M., Parker, B. J., Bhat, A., Rollo, B., Razi, A., and Friston, K. J. (2022). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969.e8.

Kanner, A. M. and Bicchi, M. M. (2022). Antiseizure medications for adults with epilepsy. *JAMA*, 327(13):1269.

Karrer, T. M., Kim, J. Z., Stiso, J., Kahn, A. E., Pasqualetti, F., Habel, U., and Bassett, D. S. (2020). A practical guide to methodological considerations in the controllability of structural brain networks. *Journal of Neural Engineering*, 17(2):026031.

Keefer, E. W., Botterman, B. R., Romero, M. I., Rossi, A. F., and Gross, G. W. (2008). Carbon nanotube coating improves neuronal recordings. *Nature Nanotechnology*, 3(7):434–439.

Kellaway, P. (1946). The part played by electric fish in the early history of bioelectricity and electrotherapy. *Bulletin of the History of Medicine*, 20(2):112–137.

Kerrigan, J. F., Litt, B., Fisher, R. S., Cranstoun, S., French, J. A., Blum, D. E., Dichter, M., Shetter, A., Baltuch, G., Jaggi, J., Krone, S., Brodie, M., Rise, M., and Graves, N. (2004). Electrical stimulation of the anterior nucleus of the thalamus for the treatment of intractable epilepsy. *Epilepsia*, 45(4):346–354.

Khosravani, H., Carlen, P. L., and Velazquez, J. L. P. (2003). The control of seizure-like activity in the rat hippocampal slice. *Biophysical Journal*, 84(1):687–695.

Kiral-Kornek, I., Roy, S., Nurse, E., Mashford, B., Karoly, P., Carroll, T., Payne, D., Saha, S., Baldassano, S., O'Brien, T., Grayden, D., Cook, M., Freestone, D., and Harrer, S. (2018). Epileptic seizure prediction using big data and deep learning: Toward a mobile system. *EBioMedicine*, 27:103–111.

Kiss, Z., Mooney, D., Renaud, L., and Hu, B. (2002). Neuronal response to local electrical stimulation in rat thalamus: physiological implications for mechanisms of deep brain stimulation. *Neuroscience*, 113(1):137–143.

Knotkova, H., Hamani, C., Sivanesan, E., Beuffe, M. F. E. L., Moon, J. Y., Cohen, S. P., and Huntoon, M. A. (2021). Neuromodulation for chronic pain. *The Lancet*, 397(10289):2111–2124.

Knowm (2019). Self directed channel memristors data sheet.

Koch, C., Massimini, M., Boly, M., and Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, 17(5):307–321.

Kozák, G. and Berényi, A. (2017). Sustained efficacy of closed loop electrical stimulation for long-term treatment of absence epilepsy in rats. *Scientific Reports*, 7(1).

Krauss, J. K., Lipsman, N., Aziz, T., Boutet, A., Brown, P., Chang, J. W., Davidson, B., Grill, W. M., Hariz, M. I., Horn, A., Schulder, M., Mammis, A., Tass, P. A., Volkmann, J., and Lozano, A. M. (2020). Technology of deep brain stimulation: current status and future directions. *Nature Reviews Neurology*, 17(2):75–87.

Krestinskaya, O., Salama, K. N., and James, A. P. (2018a). Analog backpropagation learning circuits for memristive crossbar neural networks. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE.

Krestinskaya, O., Salama, K. N., and James, A. P. (2018b). Learning in memristive neural network architectures using analog backpropagation circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(2):719–732.

Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.

Kuhn, A. A., Kempf, F., Brucke, C., Doyle, L. G., Martinez-Torres, I., Pogosyan, A., Trottenberg, T., Kupsch, A., Schneider, G.-H., Hariz, M. I., Vandenberghe, W., Nuttin, B., and Brown, P. (2008). 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*, 28(24):6165–6173.

Kuncel, A. M. and Grill, W. M. (2004). Selection of stimulus parameters for deep brain stimulation. *Clinical Neurophysiology*, 115(11):2431–2441.

Kuo, C.-H., White-Dzuro, G. A., and Ko, A. L. (2018). Approaches to closed-loop deep brain stimulation for movement disorders. *Neurosurgical focus*, 45(2):E2.

Kühn, A. A., Kupsch, A., Schneider, G.-H., and Brown, P. (2006). Reduction in subthalamic 8-35 hz oscillatory activity correlates with clinical improvement in parkinson's disease. *European Journal of Neuroscience*, 23(7):1956–1960.

Lago, N. and Cester, A. (2017). Flexible and organic neural interfaces: A review. *Applied Sciences*, 7(12):1292.

Laughlin, M. M., Lu, T., Dimitrijevic, A., and Zeng, F.-G. (2012). Towards a closed-loop cochlear implant system: Application of embedded monitoring of peripheral and central neural activity. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):443–454.

Lawrence, S. M., Larsen, J. O., Horch, K. W., Riso, R., and Sinkjær, T. (2002). Long-term biocompatibility of implanted polymer-based intrafascicular electrodes. *Journal of Biomedical Materials Research*, 63(5):501–506.

le Feber, J. (2019). *In Vitro Models of Brain Disorders*, page 19–49. Springer International Publishing.

le Feber, J., Stegenga, J., and Rutten, W. L. C. (2010). The effect of slow electrical stimuli to achieve learning in cultured networks of rat cortical neurons. *PLoS ONE*, 5(1):e8871.

Lee, Y. J., Kim, H.-J., Do, S. H., Kang, J. Y., and Lee, S. H. (2016). Characterization of nerve-cuff electrode interface for biocompatible and chronic stimulating application. *Sensors and Actuators B: Chemical*, 237:924–934.

Leite de Castro, D. D. C., Grayden, D. B., Meffin, H., and Spencer, M. (2024). Neural activity shaping in visual prostheses with deep learning. *Journal of Neural Engineering*.

Leong, F., Rahmani, B., Psaltis, D., Moser, C., and Ghezzi, D. (2024). An actor-model framework for visual sensory encoding. *Nature Communications*, 15(1).

Levy, R., Hutchison, W. D., Lozano, A. M., and Dostrovsky, J. O. (2000). High-frequency synchronization of neuronal activity in the subthalamic nucleus of parkinsonian patients with limb tremor. *The Journal of Neuroscience*, 20(20):7766–7775.

Li, C., Lammie, C., Amirsoleimani, A., Rahimi Azghadi, M., and Genov, R. (2023). Simulation of memristive crossbar arrays for seizure detection and prediction using parallel convolutional neural networks. *Software Impacts*, 15:100473.

Li, C., Lammie, C., Dong, X., Amirsoleimani, A., Azghadi, M. R., and Genov, R. (2022). Seizure detection and prediction by parallel memristive convolutional neural networks. *IEEE Transactions on Biomedical Circuits and Systems*, 16(4):609–625.

Li, D., Field, P., and Raisman, G. (1996). Connectional specification of regenerating entorhinal projection neuron classes cannot be overridden by altered target availability in postnatal organotypic slice co-culture. *Experimental Neurology*, 142(1):151–160.

Li, X., Tang, J., Zhang, Q., Gao, B., Yang, J. J., Song, S., Wu, W., Zhang, W., Yao, P., Deng, N., Deng, L., Xie, Y., Qian, H., and Wu, H. (2020). Power-efficient neural network with artificial dendrites. *Nature Nanotechnology*, 15(9):776–782.

Li, Y., Zhong, Y., Xu, L., Zhang, J., Xu, X., Sun, H., and Miao, X. (2013). Ultrafast synaptic events in a chalcogenide memristor. *Scientific Reports*, 3(1).

Li, Y., Zhong, Y., Zhang, J., Xu, L., Wang, Q., Sun, H., Tong, H., Cheng, X., and Miao, X. (2014). Activity-dependent synaptic plasticity of a chalcogenide electronic synapse for neuromorphic systems. *Scientific Reports*, 4(1).

Librizzi, L., Losi, G., Marcon, I., Sessolo, M., Scalmani, P., Carmignoto, G., and de Curtis, M. (2017). Interneuronal network activity at the onset of seizure-like events in entorhinal cortex slices. *The Journal of Neuroscience*, 37(43):10398–10407.

Little, D., Ketteler, R., Gissen, P., and Devine, M. J. (2019). Using stem cell–derived neurons in drug screening for neurological diseases. *Neurobiology of Aging*, 78:130–141.

Little, S., Beudel, M., Zrinzo, L., Foltynie, T., Limousin, P., Hariz, M., Neal, S., Cheeran, B., Cagnan, H., Gratwicke, J., Aziz, T. Z., Pogosyan, A., and Brown, P. (2015). Bilateral adaptive deep brain stimulation is effective in parkinson's disease. *Journal of Neurology, Neurosurgery & Psychiatry*, 87(7):717–721.

Little, S., Pogosyan, A., Neal, S., Zavala, B., Zrinzo, L., Hariz, M., Foltynie, T., Limousin, P., Ashkan, K., FitzGerald, J., Green, A. L., Aziz, T. Z., and Brown, P. (2013). Adaptive deep brain stimulation in advanced parkinson disease. *Annals of Neurology*, 74(3):449–457.

Little, S., Pogosyan, A., Neal, S., Zrinzo, L., Hariz, M., Foltynie, T., Limousin, P., and Brown, P. (2014). Controlling parkinson's disease with adaptive deep brain stimulation. *Journal of Visualized Experiments*, (89).

Liu, C., Wang, J., Deng, B., Wei, X., Yu, H., Li, H., Fietkiewicz, C., and Loparo, K. A. (2016). Closed-loop control of tremor-predominant parkinsonian state based on parameter estimation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 24(10):1109–1121.

Liu, C., Wang, J., Deng, B., Wei, X.-L., Yu, H.-T., and Li, H.-Y. (2015). Variable universe fuzzy closed-loop control of tremor predominant parkinsonian state based on parameter estimation. *Neurocomputing*, 151:1507–1518.

Liu, C., Zhao, G., Meng, Z., Zhou, C., Zhu, X., Zhang, W., Wang, J., Li, H., Wu, H., Fietkiewicz, C., and Loparo, K. A. (2021). Closing the loop of dbs using the beta oscillations in cortex. *Cognitive Neurodynamics*, 15(6):1157–1167.

Liu, C., Zhao, G., Wang, J., Wu, H., Li, H., Fietkiewicz, C., and Loparo, K. A. (2020a). Neural network-based closed-loop deep brain stimulation for modulation of pathological oscillation in parkinson's disease. *IEEE Access*, 8:161067–161079.

Liu, X. and Gao, Q. (2013). Parameter estimation and control for a neural mass model based on the unscented kalman filter. *Physical Review E*, 88(4).

Liu, Y., Zhu, R., Zhou, Y., Lü, J., and Chai, Y. (2023). Improved control effect of pathological oscillations by using delayed feedback stimulation in neural mass model with pedunculopontine nucleus. *Brain and Behavior*, 13(10).

Liu, Z., Tang, J., Gao, B., Yao, P., Li, X., Liu, D., Zhou, Y., Qian, H., Hong, B., and Wu, H. (2020b). Neural signal analysis with memristor arrays towards high-efficiency brain–machine interfaces. *Nature Communications*, 11(1).

Lopes, C. D., Mateus, J. C., and Aguiar, P. (2018). Interfacing microfluidics with microelectrode arrays for studying neuronal communication and axonal signal propagation. *Journal of Visualized Experiments*, (142).

Losi, G., Marcon, I., Mariotti, L., Sessolo, M., Chiavegato, A., and Carmignoto, G. (2016). A brain slice experimental model to study the generation and the propagation of focally-induced epileptiform activity. *Journal of Neuroscience Methods*, 260:125–131.

Lotlikar, A., Shah, N. P., Gogliettino, A. R., Vilkhu, R., Madugula, S., Grosberg, L., Hottowy, P., Sher, A., Litke, A., Chichilnisky, E., et al. (2023). Partitioned temporal dithering for efficient epiretinal electrical stimulation. In *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 1–5. IEEE.

Louis, E. D., Rios, E., and Henchcliffe, C. (2010). How are we doing with the treatment of essential tremor (ET)? *European Journal of Neurology*, 17(6):882–884.

Lozano, A. M., Mayberg, H. S., Giacobbe, P., Hamani, C., Craddock, R. C., and Kennedy, S. H. (2008). Subcallosal cingulate gyrus deep brain stimulation for treatment-resistant depression. *Biological psychiatry*, 64(6):461–467.

Lu, M., Wei, X., Che, Y., Wang, J., and Loparo, K. A. (2019). Application of reinforcement learning to deep brain stimulation in a computational model of parkinson's disease. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(1):339–349.

Ludwig, K. A., Langhals, N. B., Joseph, M. D., Richardson-Burns, S. M., Hendricks, J. L., and Kipke, D. R. (2011). Poly(3,4-ethylenedioxythiophene) (pedot) polymer coatings facilitate smaller neural recording electrodes. *Journal of Neural Engineering*, 8(1):014001.

Lunelli, L., Collini, C., Jimenez-Garduño, A., Roncador, A., Giusti, G., Verucchi, R., Pasquardini, L., Iannotta, S., Macchi, P., Lorenzelli, L., Pederzolli, C., Musio, C., and Potrich, C. (2019). Prototyping a memristive-based device to analyze neuronal excitability. *Biophysical Chemistry*, 253:106212.

Luo, Y., Sun, Y., Tian, X., Zheng, X., Wang, X., Li, W., Wu, X., Shu, B., and Hou, W. (2021). Deep brain stimulation for alzheimer's disease: Stimulation parameters and potential mechanisms of action. *Frontiers in Aging Neuroscience*, 13.

Luo, Y. H.-L. and da Cruz, L. (2016). The argus® ii retinal prosthesis system. *Progress in Retinal and Eye Research*, 50:89–107.

Madugula, S. S., Gogliettino, A. R., Zaidi, M., Aggarwal, G., Kling, A., Shah, N. P., Brown, J. B., Vilkhu, R., Hays, M. R., Nguyen, H., Fan, V., Wu, E. G., Hottowy, P., Sher, A., Litke, A. M., Silva, R. A., and Chichilnisky, E. J. (2022). Focal electrical stimulation of human retinal ganglion cells for vision restoration. *Journal of Neural Engineering*, 19(6):066040.

Malik, N., Efthymiou, A. G., Mather, K., Chester, N., Wang, X., Nath, A., Rao, M. S., and Steiner, J. P. (2014). Compounds with species and cell type specific toxicity identified in a 2000 compound drug screen of neural stem cells and rat mixed cortical neurons. *NeuroToxicology*, 45:192–200.

Mao, Z., Ling, Z., Pan, L., Xu, X., Cui, Z., Liang, S., and Yu, X. (2019). Comparison of efficacy of deep brain stimulation of different targets in parkinson's disease: A network meta-analysis. *Frontiers in Aging Neuroscience*, 11.

Matchen, T. D. and Moehlis, J. (2018). Phase model-based neuron stabilization into arbitrary clusters. *Journal of Computational Neuroscience*, 44(3):363–378.

Mateus, J., Lopes, C., Aroso, M., Costa, A., Gerós, A., Meneses, J., Faria, P., Neto, E., Lamghari, M., Sousa, M., and Aguiar, P. (2021). Bidirectional flow of action potentials in axons drives activity dynamics in neuronal cultures. *Journal of Neural Engineering*, 18(6):066045.

Mateus, J. C., Lopes, C. D. F., Cerquido, M., Leitão, L., Leitão, D., Cardoso, S., Ventura, J., and Aguiar, P. (2019). Improved in vitro electrophysiology using 3d-structured microelectrode arrays with a micro-mushrooms islets architecture capable of promoting topotaxis. *Journal of Neural Engineering*, 16(3):036012.

Maturana, M. I., Apollo, N. V., Garrett, D. J., Kameneva, T., Cloherty, S. L., Grayden, D. B., Burkitt, A. N., Ibbotson, M. R., and Meffin, H. (2018). Electrical receptive fields of retinal ganglion cells: Influence of presynaptic neurons. *PLOS Computational Biology*, 14(2):e1005997.

Maturana, M. I., Apollo, N. V., Hadjinicolaou, A. E., Garrett, D. J., Cloherty, S. L., Kameneva, T., Grayden, D. B., Ibbotson, M. R., and Meffin, H. (2016). A simple and accurate model to predict responses to multi-electrode stimulation in the retina. *PLOS Computational Biology*, 12(4):e1004849.

Mayberg, H. S., Lozano, A. M., Voon, V., McNeely, H. E., Seminowicz, D., Hamani, C., Schwalb, J. M., and Kennedy, S. H. (2005). Deep brain stimulation for treatment-resistant depression. *Neuron*, 45(5):651–660.

McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133.

McGregor, M. M. and Nelson, A. B. (2019). Circuit mechanisms of parkinson's disease. *Neuron*, 101(6):1042–1056.

Meissner, W., Leblois, A., Hansel, D., Bioulac, B., Gross, C. E., Benazzouz, A., and Boraud, T. (2005). Subthalamic high frequency stimulation resets subthalamic firing and reduces abnormal oscillations. *Brain*, 128(10):2372–2382.

Mena, G. E., Grosberg, L. E., Madugula, S., Hottowy, P., Litke, A., Cunningham, J., Chichilnisky, E. J., and Paninski, L. (2017). Electrical stimulus artifact cancellation and neural spike detection on large multi-electrode arrays. *PLOS Computational Biology*, 13(11):e1005842.

Migliore, M., Morse, T. M., Davison, A. P., Marenco, L., Shepherd, G. M., and Hines, M. L. (2003). ModelDB: Making models publicly accessible to support computational neuroscience. *Neuroinformatics*, 1(1):135–140.

Mikhaylov, A., Pimashkin, A., Pigareva, Y., Gerasimova, S., Gryaznov, E., Shchanikov, S., Zuev, A., Talanov, M., Lavrov, I., Demin, V., Erokhin, V., Lobov, S., Mukhina, I., Kazantsev, V., Wu, H., and Spagnolo, B. (2020). Neurohybrid memristive CMOS-integrated systems for biosensors and neuroprosthetics. *Frontiers in Neuroscience*, 14.

Miranda-Domínguez, O., Gonia, J., and Netoff, T. I. (2010). Firing rate control of a neuron using a linear proportional-integral controller. *Journal of Neural Engineering*, 7(6):066004.

Moehlis, J., Shea-Brown, E., and Rabitz, H. (2006). Optimal inputs for phase models of spiking neurons. *Journal of Computational and Nonlinear Dynamics*, 1(4):358–367.

Moleirinho, S., Whalen, A. J., Fried, S. I., and Pezaris, J. S. (2021). The impact of synchronous versus asynchronous electrical stimulation in artificial vision. *Journal of Neural Engineering*, 18(5):051001.

Motamedi, G. K., Salazar, P., Smith, E. L., Lesser, R. P., Webber, W. R., Ortinski, P. I., Vicini, S., and Rogawski, M. A. (2006). Termination of epileptiform activity by cooling in rat hippocampal slice epilepsy models. *Epilepsy Research*, 70(2–3):200–210.

Musk, E. (2019). An integrated brain-machine interface platform with thousands of channels. *Journal of Medical Internet Research*, 21(10):e16194.

Müller, J., Bakkum, D. J., and Hierlemann, A. (2013). Sub-millisecond closed-loop feedback stimulation between arbitrary sets of individual neurons. *Frontiers in Neural Circuits*, 6.

Nabi, A., Mirzadeh, M., Gibou, F., and Moehlis, J. (2012). Minimum energy desynchronizing control for coupled neurons. *Journal of Computational Neuroscience*, 34(2):259–271.

Nabi, A. and Moehlis, J. (2009). Charge-balanced optimal inputs for phase models of spiking neurons. In *ASME 2009 Dynamic Systems and Control Conference, Volume 1*, DSCC2009. ASMEDC.

Nabi, A., Stigen, T., Moehlis, J., and Netoff, T. (2013). Minimum energy control for in vitro neurons. *Journal of Neural Engineering*, 10(3):036005.

Nagaraj, V., Lamperski, A., and Netoff, T. I. (2017). Seizure control in a computational model using a reinforcement learning stimulation paradigm. *International Journal of Neural Systems*, 27(07):1750012.

Naples, J. G. and Ruckenstein, M. J. (2020). Cochlear implant. *Otolaryngologic Clinics of North America*, 53(1):87–102.

Nassiri, A. M., Sorkin, D. L., and Carlson, M. L. (2022). Current estimates of cochlear implant utilization in the united states. *Otology & Neurotology*, 43(5):e558–e562.

Newman, J. P., Zeller-Townson, R., Fong, M.-F., Arcot Desai, S., Gross, R. E., and Potter, S. M. (2013). Closed-loop, multichannel experimentation using the open-source neurorighter electrophysiology platform. *Frontiers in Neural Circuits*, 6.

Nuttin, B., Cosyns, P., Demeulemeester, H., Gybels, J., and Meyerson, B. (1999). Electrical stimulation in anterior limbs of internal capsules in patients with obsessive-compulsive disorder. *The Lancet*, 354(9189):1526.

Obien, M. E. J., Deligkaris, K., Bullmann, T., Bakkum, D. J., and Frey, U. (2015). Revealing neuronal function through microelectrode array recordings. *Frontiers in Neuroscience*, 8.

Olanow, C. W., Kieburtz, K., Rascol, O., Poewe, W., Schapira, A. H., Emre, M., Nissinen, H., Leinonen, M., and and, F. S. (2013). Factors predictive of the development of levodopa-induced dyskinesia and wearing-off in parkinson's disease. *Movement Disorders*, 28(8):1064–1071.

Opie, N. L., John, S. E., Rind, G. S., Ronayne, S. M., Wong, Y. T., Gerboni, G., Yoo, P. E., Lovell, T. J., Scordas, T. C., Wilson, S. L., et al. (2018). Focal stimulation of the sheep motor cortex with a chronically implanted minimally invasive electrode array mounted on an endovascular stent. *Nature biomedical engineering*, 2(12):907–914.

Opitz, T., Lima, A. D. D., and Voigt, T. (2002). Spontaneous development of synchronous oscillatory activity during maturation of cortical networks in vitro. *Journal of Neurophysiology*, 88(5):2196–2206.

Orlandi, J. G., Soriano, J., Alvarez-Lacalle, E., Teller, S., and Casademunt, J. (2013). Noise focusing and the emergence of coherent activity in neuronal cultures. *Nature Physics*, 9(9):582–590.

Oxley, T. J., Opie, N. L., John, S. E., Rind, G. S., Ronayne, S. M., Wheeler, T. L., Judy, J. W., McDonald, A. J., Dornom, A., Lovell, T. J., et al. (2016). Minimally invasive endovascular stent-electrode array for high-fidelity, chronic recordings of cortical neural activity. *Nature biotechnology*, 34(3):320–327.

O'Shea, D. J. and Shenoy, K. V. (2018). Eraasr: an algorithm for removing electrical stimulation artifacts from multielectrode array recordings. *Journal of Neural Engineering*, 15(2):026020.

Pal, S., Sun, D., Limbrick, D., Rafiq, A., and DeLorenzo, R. (2001). Epileptogenesis induces long-term alterations in intracellular calcium release and sequestration mechanisms in the hippocampal neuronal culture model of epilepsy. *Cell Calcium*, 30(4):285–296.

Pan, H., Song, H., Zhang, Q., and Mi, W. (2022). Review of closed-loop brain–machine interface systems from a control perspective. *IEEE Transactions on Human-Machine Systems*, 52(5):877–893.

Panuccio, G., Colombi, I., and Chiappalone, M. (2018). Recording and modulation of epileptiform activity in rodent brain slices coupled to microelectrode arrays. *Journal of Visualized Experiments*, (135).

Panuccio, G., Guez, A., Vincent, R., Avoli, M., and Pineau, J. (2013). Adaptive control of epileptiform excitability in an in vitro model of limbic seizures. *Experimental Neurology*, 241:179–183.

Parastarfeizabadi, M. and Kouzani, A. Z. (2017). Advances in closed-loop deep brain stimulation devices. *Journal of NeuroEngineering and Rehabilitation*, 14(1).

Park, J., Kim, G., and Jung, S.-D. (2017). A 128-channel fpga-based real-time spike-sorting bidirectional closed-loop neural interface system. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(12):2227–2238.

Park, S., Song, Y. J., Boo, H., and Chung, T. D. (2010). Nanoporous pt microelectrode for neural stimulation and recording: In vitro characterization. *The Journal of Physical Chemistry C*, 114(19):8721–8726.

Perucca, E. (2019). Antiepileptic drugs: evolution of our knowledge and changes in drug trials. *Epileptic disorders*, 21(4):319–329.

Perucca, E., French, J., and Bialer, M. (2007). Development of new antiepileptic drugs: challenges, incentives, and recent advances. *The Lancet Neurology*, 6(9):793–804.

Petoe, M. A., Titchener, S. A., Kolic, M., Kentler, W. G., Abbott, C. J., Nayagam, D. A. X., Baglin, E. K., Kvansakul, J., Barnes, N., Walker, J. G., Epp, S. B., Young, K. A., Ayton, L. N., Luu, C. D., and Allen, P. J. (2021). A second-generation (44-channel) suprachoroidal retinal prosthesis: Interim clinical trial results. *Translational Vision Science amp; Technology*, 10(10):12.

Pickett, M. D., Medeiros-Ribeiro, G., and Williams, R. S. (2012). A scalable neuristor built with mott memristors. *Nature Materials*, 12(2):114–117.

Pineau, J. and Avoli, M. (2009). Treating epilepsy via adaptive neurostimulation: A reinforcement learning approach. *International Journal of Neural Systems*, 19(04):227–240.

Popovych, O. V., Hauptmann, C., and Tass, P. A. (2005). Effective desynchronization by nonlinear delayed feedback. *Physical Review Letters*, 94(16).

Popovych, O. V., Hauptmann, C., and Tass, P. A. (2006). Control of neuronal synchrony by nonlinear delayed feedback. *Biological Cybernetics*, 95(1):69–85.

Popovych, O. V., Lysyansky, B., Rosenblum, M., Pikovsky, A., and Tass, P. A. (2017a). Pulsatile desynchronizing delayed feedback for closed-loop deep brain stimulation. *PLOS ONE*, 12(3):e0173363.

Popovych, O. V., Lysyansky, B., and Tass, P. A. (2017b). Closed-loop deep brain stimulation by pulsatile delayed feedback with increased gap between pulse phases. *Scientific Reports*, 7(1).

Popovych, O. V. and Tass, P. A. (2019). Adaptive delivery of continuous and delayed feedback deep brain stimulation - a computational study. *Scientific Reports*, 9(1).

Pospischil, M., Toledo-Rodriguez, M., Monier, C., Piwkowska, Z., Bal, T., Frégnac, Y., Markram, H., and Destexhe, A. (2008). Minimal hodgkin–huxley type models for different classes of cortical and thalamic neurons. *Biological Cybernetics*, 99(4–5):427–441.

Pyragas, K. (1992). Continuous control of chaos by self-controlling feedback. *Physics Letters A*, 170(6):421–428.

Pyragas, K. (2006). Delayed feedback control of chaos. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1846):2309–2334.

Pyragas, K., Popovych, O. V., and Tass, P. A. (2007). Controlling synchrony in oscillatory networks with a separate stimulation-registration setup. *Europhysics Letters (EPL)*, 80(4):40002.

Rasche, D., Rinaldi, P. C., Young, R. F., and Tronnier, V. M. (2006). Deep brain stimulation for the treatment of various chronic pain syndromes. *Neurosurgical Focus*, 21(6):1–8.

Reger, B. D., Fleming, K. M., Sanguineti, V., Alford, S., and Mussa-Ivaldi, F. A. (2000). Connecting brains to robots: An artificial body for studying the computational properties of neural tissues. *Artificial Life*, 6(4):307–324.

Rehncrona, S., Johnels, B., Widner, H., Törnqvist, A.-L., Hariz, M., and Sydow, O. (2002). Long-term efficacy of thalamic deep brain stimulation for tremor: Double-blind assessments. *Movement Disorders*, 18(2):163–170.

Relic, L., Zhang, B., Tuan, Y.-L., and Beyeler, M. (2022). Deep learning–based perceptual stimulus encoder for bionic vision. In *Augmented Humans 2022*, AHs 2022. ACM.

Robinson, J. T., Jorgolli, M., Shalek, A. K., Yoon, M.-H., Gertner, R. S., and Park, H. (2012). Vertical nanowire electrode arrays as a scalable platform for intracellular interfacing to neuronal circuits. *Nature Nanotechnology*, 7(3):180–184.

Rogawski, M. A. (2011). Revisiting ampa receptors as an antiepileptic drug target: Revisiting ampa receptors as an antiepileptic drug target. *Epilepsy Currents*, 11(2):56–63.

Rolston, J. D. (2009). A low-cost multielectrode system for data acquisition enabling real-time closed-loop processing with rapid recovery from stimulation artifacts. *Frontiers in Neuroengineering*, 2.

Rosa, M., Arlotti, M., Ardolino, G., Cogiamanian, F., Marceglia, S., Fonzo, A. D., Cortese, F., Rampini, P. M., and Priori, A. (2015). Adaptive deep brain stimulation in a freely moving parkinsonian patient. *Movement Disorders*, 30(7):1003–1005.

Rosenblum, M. (2006). Delayed feedback suppression of collective rhythmic activity in a neuronal ensemble. *International Journal of Bifurcation and Chaos*, 16(07):1989–1999.

Rosenblum, M. G. and Pikovsky, A. S. (2004). Controlling synchronization in an ensemble of globally coupled oscillators. *Physical Review Letters*, 92(11).

Rosenfeld, J. V., Wong, Y. T., Yan, E., Szlawski, J., Mohan, A., Clark, J., Rosa, M., and Lowery, A. (2020). Tissue response to a chronically implantable wireless, intracortical visual prosthesis (gennaris array). *Journal of Neural Engineering*.

Rosin, B., Slovik, M., Mitelman, R., Rivlin-Etzion, M., Haber, S. N., Israel, Z., Vaadia, E., and Bergman, H. (2011). Closed-loop deep brain stimulation is superior in ameliorating parkinsonism. *Neuron*, 72(2):370–384.

Rouhollahi, K., Emadi Andani, M., Askari Marnanii, J., and Karbassi, S. M. (2019). Rehabilitation of the parkinson's tremor by using robust adaptive sliding mode controller: a simulation study. *IET Systems Biology*, 13(2):92–99.

Roux, S. G., Cenier, T., Garcia, S., Litaudon, P., and Buonviso, N. (2007). A wavelet-based method for local phase extraction from a multi-frequency oscillatory signal. *Journal of Neuroscience Methods*, 160(1):135–143.

Rui, Y., Liu, J., Wang, Y., and Yang, C. (2011). Parylene-based implantable pt-black coated flexible 3-d hemispherical microelectrode arrays for improved neural interfaces. *Microsystem Technologies*, 17(3):437–442.

Rush, A. J., Trivedi, M. H., Wisniewski, Stewart, J. W., Warden, D., Niederehe, G., Thase, M. E., Lavori, P. W., Lebowitz, B. D., et al. (2006). Acute and longer-term outcomes in depressed outpatients requiring one or several treatment steps: a star* d report. *American Journal of Psychiatry*, 163(11):1905–1917.

Rutecki, P. A., Lebeda, F. J., and Johnston, D. (1987). 4-aminopyridine produces epileptiform activity in hippocampus and enhances synaptic excitation and inhibition. *Journal of Neurophysiology*, 57(6):1911–1924.

Salam, M. T., Kassiri, H., Genov, R., and Perez Velazquez, J. L. (2015a). Rapid brief feedback intracerebral stimulation based on real-time desynchronization detection preceding seizures stops the generation of convulsive paroxysms. *Epilepsia*, 56(8):1227–1238.

Salam, M. T., Velazquez, J. L. P., and Genov, R. (2015b). Seizure suppression efficacy of closed-loop versus open-loop deep brain stimulation in a rodent model of epilepsy. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 24(6):710–719.

Sankar, T. and Lozano, A. M. (2011). Surgical approach to l-dopa-induced dyskinesias. In *International Review of Neurobiology*, pages 151–171. Elsevier.

Sankaraneni, R. and Lachhwani, D. (2015). Antiepileptic drugs—a review. *Pediatric Annals*, 44(2).

Santaniello, S., Fiengo, G., Glielmo, L., and Grill, W. M. (2010). Closed-loop control of deep brain stimulation: a simulation study. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(1):15–24.

Schalk, G. and Leuthardt, E. C. (2011). Brain-computer interfaces using electrocorticographic signals. *IEEE Reviews in Biomedical Engineering*, 4:140–154.

Schermer, M. (2011). Ethical issues in deep brain stimulation. *Frontiers in integrative neuroscience*, 5:17.

Schiff, S. J. (2010). Towards model-based control of parkinson's disease. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1918):2269–2308.

Schiff, S. J., Jerger, K., Duong, D. H., Chang, T., Spano, M. L., and Ditto, W. L. (1994). Controlling chaos in the brain. *Nature*, 370(6491):615–620.

Schiff, S. J. and Sauer, T. (2007). Kalman filter control of a model of spatiotemporal cortical dynamics. *Journal of Neural Engineering*, 5(1):1–8.

Schirrmeister, R. T., Springenberg, J. T., Fiederer, L. D. J., Glasstetter, M., Eggensperger, K., Tangermann, M., Hutter, F., Burgard, W., and Ball, T. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping*, 38(11):5391–5420.

Schmid, E. W., Fink, W., and Wilke, R. (2013). Simultaneous vs. sequential and unipolar vs. multipolar stimulation in retinal prostheses. In *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 190–193. IEEE.

Schnitzler, A., Münks, C., Butz, M., Timmermann, L., and Gross, J. (2009). Synchronized brain network associated with essential tremor as revealed by magnetoencephalography. *Movement Disorders*, 24(11):1629–1635.

Schuurman, P. R., Bosch, D. A., Bossuyt, P. M., Bonsel, G. J., van Someren, E. J., de Bie, R. M., Merkus, M. P., and Speelman, J. D. (2000). A comparison of continuous thalamic stimulation and thalamotomy for suppression of severe tremor. *New England Journal of Medicine*, 342(7):461–468.

Schwartzkroin, P. A. and Prince, D. A. (1978). Cellular and field potential properties of epileptogenic hippocampal slices. *Brain Research*, 147(1):117–130.

Schwarz, D. (2024). uipickfiles: uigetfile on steroids. https://www.mathworks.com/matlabcentral/fileexchange/10867-uipickfiles-uigetfile-on-steroids.

Sejnowski, T. J., Koch, C., and Churchland, P. S. (1988). Computational neuroscience. *Science*, 241(4871):1299–1306.

Serb, A., Corna, A., George, R., Khiat, A., Rocchi, F., Reato, M., Maschietto, M., Mayr, C., Indiveri, G., Vassanelli, S., and Prodromakis, T. (2020). Memristive synapses connect brain and silicon spiking neurons. *Scientific Reports*, 10(1).

Seu, G. P., Angotzi, G. N., Boi, F., Raffo, L., Berdondini, L., and Meloni, P. (2018). Exploiting all programmable socs in neural signal analysis: A closed-loop control for large-scale cmos multielectrode arrays. *IEEE Transactions on Biomedical Circuits and Systems*, 12(4):839–850.

Shah, N. P., Madugula, S., Grosberg, L., Mena, G., Tandon, P., Hottowy, P., Sher, A., Litke, A., Mitra, S., and Chichilnisky, E. (2019). Optimization of electrical stimulation for a high-fidelity artificial retina. In *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 714–718. IEEE.

Shchanikov, S., Zuev, A., Bordanov, I., Danilin, S., Lukoyanov, V., Korolev, D., Belov, A., Pigareva, Y., Gladkov, A., Pimashkin, A., Mikhaylov, A., Kazantsev, V., and Serb, A. (2021).

Designing a bidirectional, adaptive neural interface incorporating machine learning capabilities and memristor-enhanced hardware. *Chaos, Solitons & Fractals*, 142:110504.

Shealy, C. N., Mortimer, J. T., and Reswick, J. B. (1967). Electrical inhibition of pain by stimulation of the dorsal columns: preliminary clinical report. *Anesthesia & Analgesia*, 46(4):489–491.

Shein, M., Greenbaum, A., Gabay, T., Sorkin, R., David-Pur, M., Ben-Jacob, E., and Hanein, Y. (2008). Engineered neuronal circuits shaped and interfaced with carbon nanotube microelectrode arrays. *Biomedical Microdevices*, 11(2):495–501.

Shen, K., Zhu, Z., Munhall, A., and Johnson, S. W. (2003). Synaptic plasticity in rat subthalamic nucleus induced by high-frequency stimulation. *Synapse*, 50(4):314–319.

Shepherd, R. K., Wise, A. K., Enke, Y. L., Carter, P. M., and Fallon, J. B. (2017). Evaluation of focused multipolar stimulation for cochlear implants: a preclinical safety study. *Journal of neural engineering*, 14(4):046020.

Sheth, S. A., Bijanki, K. R., Metzger, B., Allawala, A., Pirtle, V., Adkinson, J. A., Myers, J., Mathura, R. K., Oswalt, D., Tsolaki, E., Xiao, J., Noecker, A., Strutt, A. M., Cohn, J. F., McIntyre, C. C., Mathew, S. J., Borton, D., Goodman, W., and Pouratian, N. (2022). Deep brain stimulation for depression informed by intracranial recordings. *Biological Psychiatry*, 92(3):246–251.

Simeral, J. D., Kim, S.-P., Black, M. J., Donoghue, J. P., and Hochberg, L. R. (2011). Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *Journal of Neural Engineering*, 8(2):025027.

Simmons, F. B., Epley, J. M., Lummis, R. C., Guttman, N., Frishkopf, L. S., Harmon, L. D., and Zwicker, E. (1965). Auditory nerve: Electrical stimulation in man. *Science*, 148(3666):104–106.

Sinclair, N. C., Shivdasani, M. N., Perera, T., Gillespie, L. N., McDermott, H. J., Ayton, L. N., and Blamey, P. J. (2016). The appearance of phosphenes elicited using a suprachoroidal retinal prosthesis. *Investigative Opthalmology amp; Visual Science*, 57(11):4948.

Singh, A. (2018). Oscillatory activity in the cortico-basal ganglia-thalamic neural circuits in parkinson's disease. *European Journal of Neuroscience*, 48(8):2869–2878.

Smeal, R. M., Ermentrout, G. B., and White, J. A. (2010). Phase-response curves and synchronized neural networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1551):2407–2422.

Smith, Z. M., Parkinson, W. S., and Long, C. J. (2013). Multipolar current focusing increases spectral resolution in cochlear implants. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2796–2799. IEEE.

Spencer, M. J., Kameneva, T., Grayden, D. B., Burkitt, A. N., and Meffin, H. (2021). Neural activity shaping utilizing a partitioned target pattern. *Journal of Neural Engineering*, 18(4):046025.

Spencer, M. J., Kameneva, T., Grayden, D. B., Meffin, H., and Burkitt, A. N. (2019). Global activity shaping strategies for a retinal implant. *Journal of Neural Engineering*, 16(2):026008.

Spencer, T. C., Fallon, J. B., Thien, P. C., and Shivdasani, M. N. (2016). Spatial restriction of neural activation using focused multipolar stimulation with a retinal prosthesis. *Investigative Ophthalmology & Visual Science*, 57(7):3181–3191.

Spira, M. E. and Hai, A. (2013). Multi-electrode array technologies for neuroscience and cardiology. *Nature Nanotechnology*, 8(2):83–94.

Sporns, O. (2011). The human connectome: a complex network. *Annals of the New York Academy of Sciences*, 1224(1):109–125.

Stam, C. (2005). Nonlinear dynamical analysis of EEG and MEG: Review of an emerging field. *Clinical Neurophysiology*, 116(10):2266–2301.

Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., Chen, S., Colonell, J., Gardner, R. J., Karsh, B., Kostadinov, D., Mora-Lopez, C., Park, J., Putzeys, J., Sauerbrei, B., van Daal, R. J. J., Vollan, A. Z., Welkenhuysen, M., Ye, Z., Dudman, J., Dutta, B., Hantman, A. W., Harris, K. D., Lee, A. K., Moser, E. I., O'Keefe, J., Renart, A., Svoboda, K., Häusser, M., Haesler, S., Carandini, M., and Harris, T. D. (2020). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings.

Steriade, M., Nunez, A., and Amzica, F. (1993). Intracellular analysis of relations between the slow (lt; 1 hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *The Journal of Neuroscience*, 13(8):3266–3283.

Strukov, D. B., Snider, G. S., Stewart, D. R., and Williams, R. S. (2008). The missing memristor found. *Nature*, 453(7191):80–83.

Su, F., Kumaravelu, K., Wang, J., and Grill, W. M. (2019). Model-based evaluation of closed-loop deep brain stimulation controller to adapt to dynamic changes in reference signal. *Frontiers in Neuroscience*, 13.

Su, F., Wang, J., Li, H., Deng, B., Yu, H., and Liu, C. (2017). Analysis and application of neuronal network controllability and observability. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(2).

Suppa, A., Huang, Y.-Z., Funke, K., Ridding, M., Cheeran, B., Lazzaro, V. D., Ziemann, U., and Rothwell, J. (2016). Ten years of theta burst stimulation in humans: Established knowledge, unknowns and prospects. *Brain Stimulation*, 9(3):323–335.

Suresh, J., Radojicic, M., Pesce, L. L., Bhansali, A., Wang, J., Tryba, A. K., Marks, J. D., and van Drongelen, W. (2016). Network burst activity in hippocampal neuronal cultures: the role of synaptic and intrinsic currents. *Journal of Neurophysiology*, 115(6):3073–3089.

Suweis, S., Tu, C., Rocha, R. P., Zampieri, S., Zorzi, M., and Corbetta, M. (2019). Brain controllability: Not a slam dunk yet. *NeuroImage*, 200:552–555.

Swann, N. C., de Hemptinne, C., Thompson, M. C., Miocinovic, S., Miller, A. M., Gilron, R., Ostrem, J. L., Chizeck, H. J., and Starr, P. A. (2018). Adaptive deep brain stimulation for parkinson's disease using motor cortex sensing. *Journal of Neural Engineering*, 15(4):046006.

Tanskanen, J. M., Ahtiainen, A., and Hyttinen, J. A. (2019). Extracellular electrical stimulation-based in vitro neuroscience: A minireview of methods and a paradigm shift proposal. In *2019*

*26th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, pages 883–886. IEEE.

Taylor, A. M., Rhee, S. W., Tu, C. H., Cribbs, D. H., Cotman, C. W., and Jeon, N. L. (2002). Microfluidic multicompartment device for neuroscience research. *Langmuir*, 19(5):1551–1556.

Tinkhauser, G., Pogosyan, A., Little, S., Beudel, M., Herz, D. M., Tan, H., and Brown, P. (2017). The modulatory effect of adaptive deep brain stimulation on beta bursts in parkinson's disease. *Brain*, 140(4):1053–1067.

Tong, W., Meffin, H., Garrett, D. J., and Ibbotson, M. R. (2020). Stimulation strategies for improving the resolution of retinal prostheses. *Frontiers in neuroscience*, 14:262.

Tran, N., Bai, S., Yang, J., Chun, H., Kavehei, O., Yang, Y., Muktamath, V., Ng, D., Meffin, H., Halpern, M., et al. (2014). A complete 256-electrode retinal prosthesis chip. *IEEE Journal of Solid-State Circuits*, 49(3):751–765.

Tu, C., Rocha, R. P., Corbetta, M., Zampieri, S., Zorzi, M., and Suweis, S. (2018). Warnings and caveats in brain controllability. *NeuroImage*, 176:83–91.

Tufail, Y., Yoshihiro, A., Pati, S., Li, M. M., and Tyler, W. J. (2011). Ultrasonic neuromodulation by brain stimulation with transcranial ultrasound. *Nature Protocols*, 6(9):1453–1470.

Tukhlina, N., Rosenblum, M., Pikovsky, A., and Kurths, J. (2007). Feedback suppression of neural synchrony by vanishing stimulation. *Physical Review E*, 75(1).

Tuma, T., Pantazi, A., Gallo, M. L., Sebastian, A., and Eleftheriou, E. (2016). Stochastic phase-change neurons. *Nature Nanotechnology*, 11(8):693–699.

Uhlhaas, P. J. and Singer, W. (2006). Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology. *neuron*, 52(1):155–168.

Upadhyay, N. K., Jiang, H., Wang, Z., Asapu, S., Xia, Q., and Yang, J. J. (2019). Emerging memory devices for neuromorphic computing. *Advanced Materials Technologies*, 4(4):1800589.

Ursino, M. and La Cara, G.-E. (2006). Travelling waves and eeg patterns during epileptic seizure: Analysis with an integrate-and-fire neural network. *Journal of Theoretical Biology*, 242(1):171–187.

van Bergen, A., Papanikolaou, T., Schuker, A., Möller, A., and Schlosshauer, B. (2003). Long-term stimulation of mouse hippocampal slice culture on microelectrode array. *Brain Research Protocols*, 11(2):123–133.

van den Heuvel, M. P. and Yeo, B. T. (2017). A spotlight on bridging microscale and macroscale human brain architecture. *Neuron*, 93(6):1248–1251.

van den Honert, C. and Kelsall, D. C. (2007). Focused intracochlear electric stimulation with phased array channels. *The Journal of the Acoustical Society of America*, 121(6):3703–3716.

van Heukelum, S., Kelderhuis, J., Janssen, P., van Luijtelaar, G., and Lüttjohann, A. (2016). Timing of high-frequency cortical stimulation in a genetic absence model. *Neuroscience*, 324:191–201.

Vassileva, A., van Blooijs, D., Leijten, F., and Huiskamp, G. (2018). Neocortical electrical stimulation for epilepsy: Closed-loop versus open-loop. *Epilepsy Research*, 141:95–101.

Velisar, A., Syrkin-Nikolau, J., Blumenfeld, Z., Trager, M., Afzal, M., Prabhakar, V., and Bronte-Stewart, H. (2019). Dual threshold neural closed loop deep brain stimulation in parkinson disease patients. *Brain Stimulation*, 12(4):868–876.

Venkatraman, S., Hendricks, J., King, Z. A., Sereno, A. J., Richardson-Burns, S., Martin, D., and Carmena, J. M. (2011). In vitro and in vivo evaluation of pedot microelectrodes for neural stimulation and recording. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(3):307–316.

Vijiaratnam, N., Simuni, T., Bandmann, O., Morris, H. R., and Foltynie, T. (2021). Progress towards therapies for disease modification in parkinson's disease. *The Lancet Neurology*, 20(7):559–572.

Vlachos, I., Deniz, T., Aertsen, A., and Kumar, A. (2016). Recovery of dynamics and function in spiking neural networks with closed-loop control. *PLoS computational biology*, 12(2):e1004720.

Volkmann, J., Joliot, M., Mogilner, A., Ioannides, A. A., Lado, F., Fazzini, E., Ribary, U., and Llinas, R. (1996). Central motor loop oscillations in parkinsonian resting tremor revealed magnetoencephalography. *Neurology*, 46(5):1359–1359.

Vonck, K. and Boon, P. (2012). A decade of experience with deep brain stimulation for patients with refractory medial temporal lobe epilepsy. *International Journal of Neural Systems*, 23(01):1250034.

Vreeswijk, C. v. and Hansel, D. (2001). Patterns of synchrony in neural networks with spike adaptation. *Neural Computation*, 13(5):959–992.

Wagenaar, D., DeMarse, T. B., and Potter, S. M. (2005a). Meabench: A toolset for multi-electrode data acquisition and on-line analysis. In *Conference Proceedings. 2nd International IEEE EMBS Conference on Neural Engineering, 2005.*, pages 518–521. IEEE.

Wagenaar, D. A., Madhavan, R., Pine, J., and Potter, S. M. (2005b). Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *The Journal of Neuroscience*, 25(3):680–688.

Wagenaar, D. A., Pine, J., and Potter, S. M. (2004). Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *Journal of Neuroscience Methods*, 138(1–2):27–37.

Wagenaar, D. A., Pine, J., and Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience*, 7(1).

Wagenaar, D. A. and Potter, S. M. (2002). Real-time multi-channel stimulus artifact suppression by local curve fitting. *Journal of Neuroscience Methods*, 120(2):113–120.

Wallach, A., Eytan, D., Gal, A., Zrenner, C., and Marom, S. (2011). Neuronal response clamp. *Frontiers in Neuroengineering*, 4.

Wang, C., Fang, C., Zou, Y., Yang, J., and Sawan, M. (2023). Artificial intelligence techniques for retinal prostheses: a comprehensive review and future direction. *Journal of Neural Engineering*.

Wang, J., Niebur, E., Hu, J., and Li, X. (2016a). Suppressing epileptic activity in a neural mass model using a closed-loop proportional-integral controller. *Scientific Reports*, 6(1).

Wang, K., Fishman, H. A., Dai, H., and Harris, J. S. (2006). Neural stimulation with a carbon nanotube microelectrode array. *Nano Letters*, 6(9):2043–2048.

Wang, S., Rhijn, J.-R. v., Akkouh, I., Kogo, N., Maas, N., Bleeck, A., Ortiz, I. S., Lewerissa, E., Wu, K. M., Schoenmaker, C., Djurovic, S., van Bokhoven, H., Kleefstra, T., Nadif Kasri, N., and Schubert, D. (2022). Loss-of-function variants in the schizophrenia risk gene setd1a alter neuronal network activity in human neurons through the camp/pka pathway. *Cell Reports*, 39(5):110790.

Wang, X.-J. (1999). Synaptic basis of cortical persistent activity: the importance of nmda receptors to working memory. *The Journal of Neuroscience*, 19(21):9587–9603.

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968.

Wang, Y., Zhan, G., Cai, Z., Jiao, B., Zhao, Y., Li, S., and Luo, A. (2021). Vagus nerve stimulation in brain diseases: Therapeutic applications and biological mechanisms. *Neuroscience & Biobehavioral Reviews*, 127:37–53.

Wang, Z., Joshi, S., Savel'ev, S., Song, W., Midya, R., Li, Y., Rao, M., Yan, P., Asapu, S., Zhuo, Y., Jiang, H., Lin, P., Li, C., Yoon, J. H., Upadhyay, N. K., Zhang, J., Hu, M., Strachan, J. P., Barnell, M., Wu, Q., Wu, H., Williams, R. S., Xia, Q., and Yang, J. J. (2018). Fully memristive neural networks for pattern classification with unsupervised learning. *Nature Electronics*, 1(2):137–145.

Wang, Z., Joshi, S., Savel'ev, S. E., Jiang, H., Midya, R., Lin, P., Hu, M., Ge, N., Strachan, J. P., Li, Z., Wu, Q., Barnell, M., Li, G.-L., Xin, H. L., Williams, R. S., Xia, Q., and Yang, J. J. (2016b). Memristors with diffusive dynamics as synaptic emulators for neuromorphic computing. *Nature Materials*, 16(1):101–108.

Wang, Z., Simoncelli, E. P., and Bovik, A. C. (2003). Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee.

Ward, M. P., Qing, K. Y., Otto, K. J., Worth, R. M., John, S. W. M., and Irazoqui, P. P. (2015). A flexible platform for biofeedback-driven control and personalization of electrical nerve stimulation therapy. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 23(3):475–484.

Weerasinghe, G., Duchet, B., Bick, C., and Bogacz, R. (2021). Optimal closed-loop deep brain stimulation using multiple independently controlled contacts. *PLOS Computational Biology*, 17(8):e1009281.

Weihberger, O., Okujeni, S., Mikkonen, J. E., and Egert, U. (2013). Quantitative examination of stimulus-response relations in cortical networks in vitro. *Journal of Neurophysiology*, 109(7):1764–1774.

Weiland, J. D., Liu, W., and Humayun, M. S. (2005). Retinal prosthesis. *Annu. Rev. Biomed. Eng.*, 7:361–401.

Weissinger, F., Buchheim, K., Siegmund, H., and Meierkord, H. (2005). Seizure spread through the life cycle: Optical imaging in combined brain slices from immature, adult, and senile rats in vitro. *Neurobiology of Disease*, 19(1–2):84–95.

Wilke, R., Moghadam, G. K., Lovell, N., Suaning, G., and Dokos, S. (2011). Electric crosstalk impairs spatial resolution of multi-electrode arrays in retinal implants. *Journal of neural engineering*, 8(4):046016.

Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., and Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature*, 593(7858):249–254.

Williamson, A., Schumann, L., Hiller, L., Klefenz, F., Hoerselmann, I., Husar, P., and Schober, A. (2013). Synaptic behavior and STDP of asymmetric nanoscale memristors in biohybrid systems. *Nanoscale*, 5(16):7297.

Wilson, C. and Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *The Journal of Neuroscience*, 16(7):2397–2410.

Wilson, H. R. and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24.

Windels, F., Bruet, N., Poupard, A., Feuerstein, C., Bertrand, A., and Savasta, M. (2003). Influence of the frequency parameter on extracellular glutamate and -aminobutyric acid in substantia nigra and globus pallidus during electrical stimulation of subthalamic nucleus in rats. *Journal of Neuroscience Research*, 72(2):259–267.

Windels, F., Bruet, N., Poupard, A., Urbain, N., Chouvet, G., Feuerstein, C., and Savasta, M. (2000). Effects of high frequency stimulation of subthalamic nucleus on extracellular glutamate and gaba in substantia nigra and globus pallidus in the normal rat. *European Journal of Neuroscience*, 12(11):4141–4146.

Wirrell, E. C. (2006). Epilepsy-related injuries. *Epilepsia*, 47(s1):79–86.

Witt, A., Palmigiano, A., Neef, A., El Hady, A., Wolf, F., and Battaglia, D. (2013). Controlling the oscillation phase through precisely timed closed-loop optogenetic stimulation: a computational study. *Frontiers in neural circuits*, 7:49.

Wojtecki, L., Moldovan, A.-S., Groiss, S., Elben, S., Südmeyer, M., and Schnitzler, A. (2015). The treatment of parkinson/s disease with deep brain stimulation: current issues. *Neural Regeneration Research*, 10(7):1018.

Wrobel, G., Höller, M., Ingebrandt, S., Dieluweit, S., Sommerhage, F., Bochem, H. P., and Offenhäusser, A. (2007). Transmission electron microscopy study of the cell–sensor interface. *Journal of The Royal Society Interface*, 5(19):213–222.

Wu, H., Hariz, M., Visser-Vandewalle, V., Zrinzo, L., Coenen, V. A., Sheth, S. A., Bervoets, C., Naesström, M., Blomstedt, P., Coyne, T., Hamani, C., Slavin, K., Krauss, J. K., Kahl, K. G., Taira, T., Zhang, C., Sun, B., Toda, H., Schlaepfer, T., Chang, J. W., Régis, J., Schuurman, R., Schulder, M., Doshi, P., Mosley, P., Poologaindran, A., Lázaro-Muñoz, G., Pepper, J., Schechtmann, G., Fytagoridis, A., Huys, D., Gonçalves-Ferreira, A., D'Haese, P.-F., Neimat, J., Broggi, G., Vilela-Filho, O., Voges, J., Alkhani, A., Nakajima, T., Richieri, R., Djurfeldt, D., Fontaine, P., Martinez-Alvarez, R., Okamura, Y., Chandler, J., Watanabe, K., Barcia, J. A., Reneses, B., Lozano, A., Gabriëls, L., Salles, A. D., Halpern, C. H., Matthews, K., Fins, J. J., and Nuttin, B. (2020). Deep brain stimulation for refractory obsessive-compulsive disorder (OCD): emerging or established therapy? *Molecular Psychiatry*, 26(1):60–65.

Wu, Y., Karetic, I., Stegmaier, J., Walter, P., and Merhof, D. (2023). A deep learning-based in silico framework for optimization on retinal prosthetic stimulation. *arXiv preprint arXiv:2302.03570*.

Wülfing, J. M., Kumar, S. S., Boedecker, J., Riedmiller, M., and Egert, U. (2019). Adaptive long-term control of biological neural networks with deep reinforcement learning. *Neurocomputing*, 342:66–74.

Xu, M., Mai, X., Lin, J., Zhang, W., Li, Y., He, Y., Tong, H., Hou, X., Zhou, P., and Miao, X. (2020). Recent advances on neuromorphic devices based on chalcogenide phase-change materials. *Advanced Functional Materials*, 30(50):2003419.

Yang, J. J., Strukov, D. B., and Stewart, D. R. (2012a). Memristive devices for computing. *Nature Nanotechnology*, 8(1):13–24.

Yang, M., Wang, J., Liu, C., and Yue, W. (2022). Adaptive supervisory control of epilepsy in a neural mass model. In *2022 41st Chinese Control Conference (CCC)*, pages 5699–5704. IEEE.

Yang, Y., Gao, P., Gaba, S., Chang, T., Pan, X., and Lu, W. (2012b). Observation of conducting filament growth in nanoscale resistive memories. *Nature Communications*, 3(1).

Yger, P., Spampinato, G. L., Esposito, E., Lefebvre, B., Deny, S., Gardella, C., Stimberg, M., Jetter, F., Zeck, G., Picaud, S., Duebel, J., and Marre, O. (2018). A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife*, 7.

Yu, H., Meng, Z., Li, H., Liu, C., and Wang, J. (2021). Intensity-varied closed-loop noise stimulation for oscillation suppression in the parkinsonian state. *IEEE Transactions on Cybernetics*, 52(9):9861–9870.

Yu, T., Wang, X., Li, Y., Zhang, G., Worrell, G., Chauvel, P., Ni, D., Qiao, L., Liu, C., Li, L., Ren, L., and Wang, Y. (2018). High-frequency stimulation of anterior nucleus of thalamus desynchronizes epileptic network in humans. *Brain*.

Yuan, F., Zhang, Z., Liu, C., Zhou, F., Yau, H. M., Lu, W., Qiu, X., Wong, H.-S. P., Dai, J., and Chai, Y. (2017). Real-time observation of the electrode-size-dependent evolution dynamics of the conducting filaments in a SiO layer. *ACS Nano*, 11(4):4097–4104.

Yue, L., Weiland, J. D., Roska, B., and Humayun, M. S. (2016). Retinal stimulation strategies to restore vision: Fundamentals and systems. *Progress in Retinal and Eye Research*, 53:21–47.

Zarzycki, M. Z. and Domitrz, I. (2019). Stimulation-induced side effects after deep brain stimulation – a systematic review. *Acta Neuropsychiatrica*, 32(2):57–64.

Zeng, F.-G., Rebscher, S., Harrison, W., Sun, X., and Feng, H. (2008). Cochlear implants: system design, integration, and evaluation. *IEEE reviews in biomedical engineering*, 1:115–142.

Zesiewicz, T. A., Elble, R., Louis, E. D., Hauser, R. A., Sullivan, K. L., Dewey, R. B., Ondo, W. G., Gronseth, G. S., and Weiner, W. J. (2005). Practice parameter: Therapies for essential tremor [RETIRED]. *Neurology*, 64(12):2008–2020.

Zhang, H., Chen, Y., Xie, Y., and Chai, Y. (2022). Closed-loop controller based on reference signal tracking for absence seizures. *Scientific Reports*, 12(1).

Zhang, J., Xu, H., Long, H., Hou, K., Li, Z., Yang, Z., Wang, Y., and Li, Q. (2021). Eeg signal epilepsy detection system based on convolutional neural network and memristor array. In *2021 IEEE 4th International Conference on Electronics Technology (ICET)*, pages 851–855. IEEE.

Zhang, M., Schwemmer, M. A., Ting, J. E., Majstorovic, C. E., Friedenberg, D. A., Bockbrader, M. A., Jerry Mysiw, W., Rezai, A. R., Annetta, N. V., Bouton, C. E., Bresler, H. S., and Sharma, G. (2018). Extracting wavelet based neural features from human intracortical recordings for neuroprosthetics applications. *Bioelectronic Medicine*, 4(1).

Zhang, Y., Wang, X., and Friedman, E. G. (2017). Memristor-based circuit design for multilayer neural networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 65(2):677–686.

Zhang, Z., Valiante, T., and Carlen, P. (2011). Transition to seizure: From "macro"- to "micro"-mysteries. *Epilepsy Research*, 97(3):290–299.

Zhou, B., Che, Y., Qin, Q., Qin, Y., and Han, C. (2020). Seizure suppression in a thalamocortical computational model of absence epilepsy by linear delayed feedback control. In *2020 13th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 872–876. IEEE.

Zou, X., Xu, S., Chen, X., Yan, L., and Han, Y. (2021). Breaking the von neumann bottleneck: architecture-level processing-in-memory technology. *Science China Information Sciences*, 64(6).

# Appendix A

# PowerMapp toolbox

The main scientific outputs of this thesis are presented in Chapters 3, 4, and 5, which addressed our primary research questions. Nonetheless, we have also contributed to the scientific community by creating software toolboxes that are freely available for other researchers or clinicians. Here, we present the software PowerMapp, a Matlab application designed to analyse stereoelectroencephalography (sEEG) data. This work resulted from a collaboration with the neurosurgery team from Centro Hospitalar do Porto (CHP).

Some patients with drug-resistant epilepsy must be submitted to resective surgery to remove seizure-prone portions of the brain. To do that, clinicians carefully analyse brain activity to understand what critical regions are involved in generating seizures (the epileptogenic zone). This analysis is typically based on the visual inspection of SEEG recordings. However, these recordings comprise dozens of channels (up to 120 here), making the analysis challenging and labour-intensive.

PowerMapp facilitates this operation by providing a user-friendly platform for multi-channel analysis of SEEG on the frequency domain, across time, focused essentially on the calculation of band-power and wavelet transforms. The software is freely available at `https://github.com/DomingosLC/PowerMapp`. In this appendix, we briefly describe how to use the PowerMapp application. We also welcome the reader to access this link for a detailed video demonstration of the software.

PowerMapp was designed to allow batch analysis of multiple SEEG files. These can be selected in the opening interface (Figure A.1.A). After selecting the files (Figure A.1.B), the user can choose to downsample the recording to improve performance, parameterise the band-power algorithms, set the limits of the intended bands and select the reference type (Figure A.1.C).

Each file will have an associated analysis panel displaying all the channel labels, the sEEG recordings and the band-power plots for the selected frequency ranges (Figure **??**.A). The user can choose a subset of channels to analyse in detail in a new panel (Figure A.2.B) or calculate their wavelet

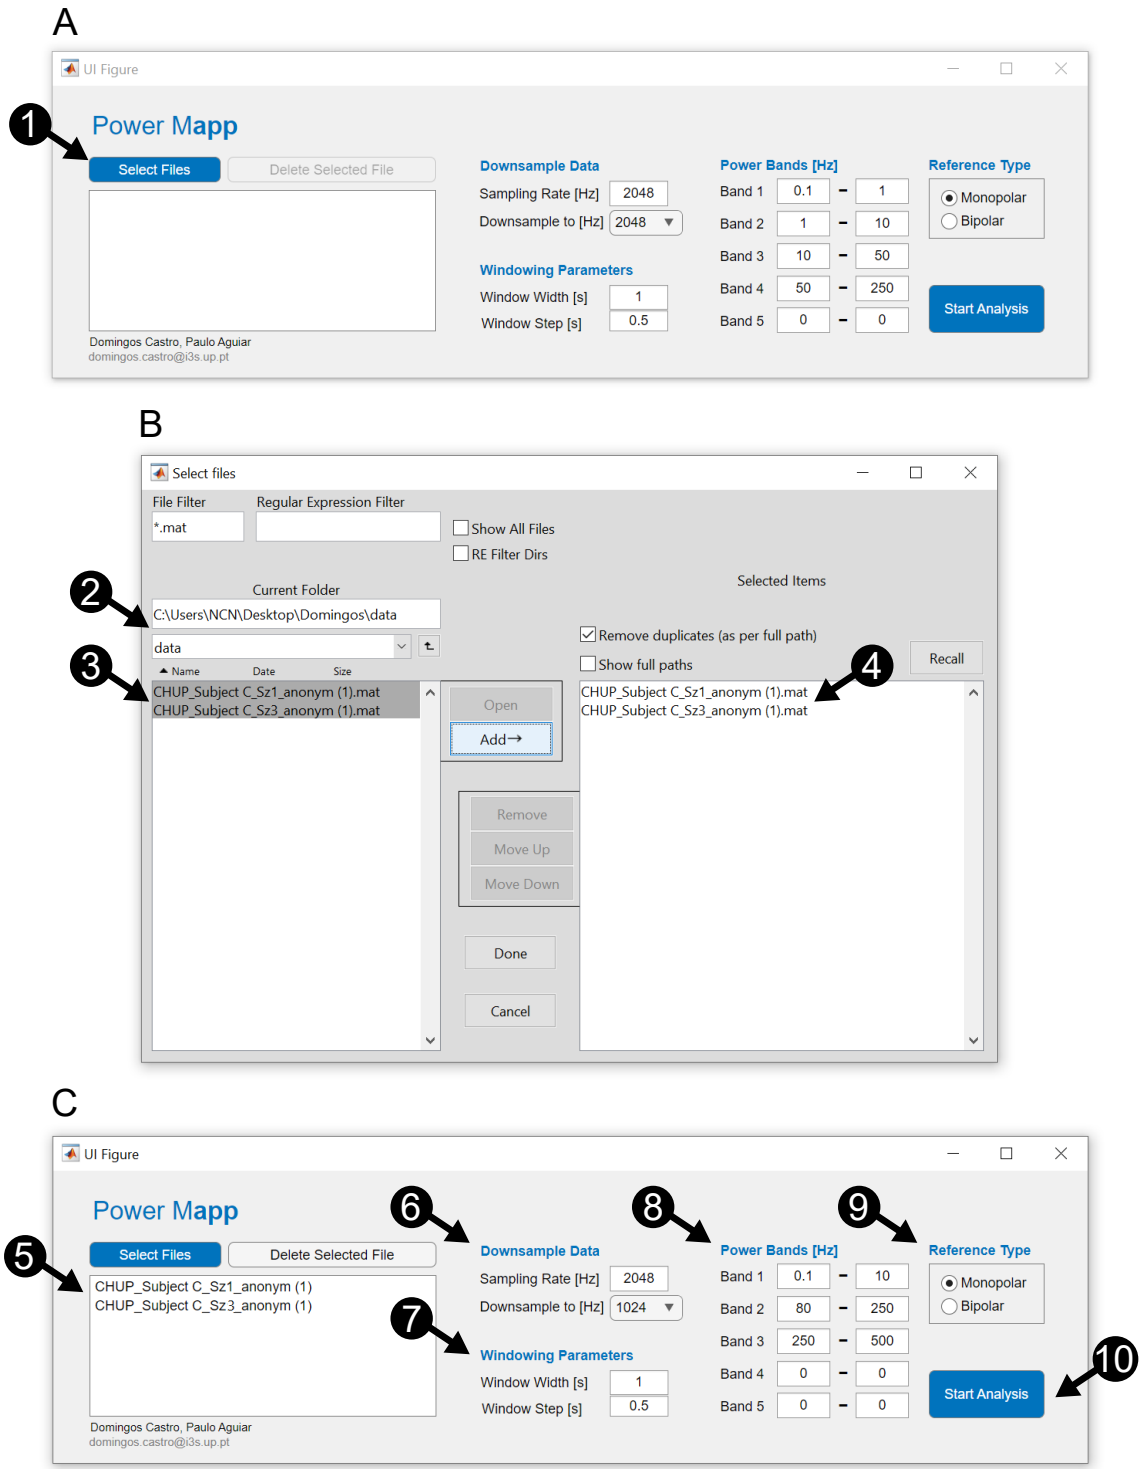



Figure A.1: **Loading data and setting analysis parameters.** A) PowerMapp opening window, where the user can open the file chooser dialogue (1). B) File finder platform (Schwarz, 2024). Using this platform, the user can search for files directories (2), select and add the desired files from within the selected directories (3), and check the files selected (4). C) The user can then see the selected files in the PowerMapp window (5), choose a downsampling factor (6), windowing parameters for the band-power (7), set the desired band limits (8), choose the reference type (9). The files are only uploaded after pressing the "Start Analysis" button (10).

transform (Figure A.2.C). The wavelet transform provides a very informative representation of the signal over the temporal and frequency domains. This type of representation is very useful as it portrays identifiable biomarkers of the epileptogenic zone (Grinenko et al., 2017).

Another relevant biomarker for clinicians is the presence of High-Frequency Oscillations (HFOs). These can be easily identified as isolated yellow spots in the high-frequency band-power maps (Figure A.3).

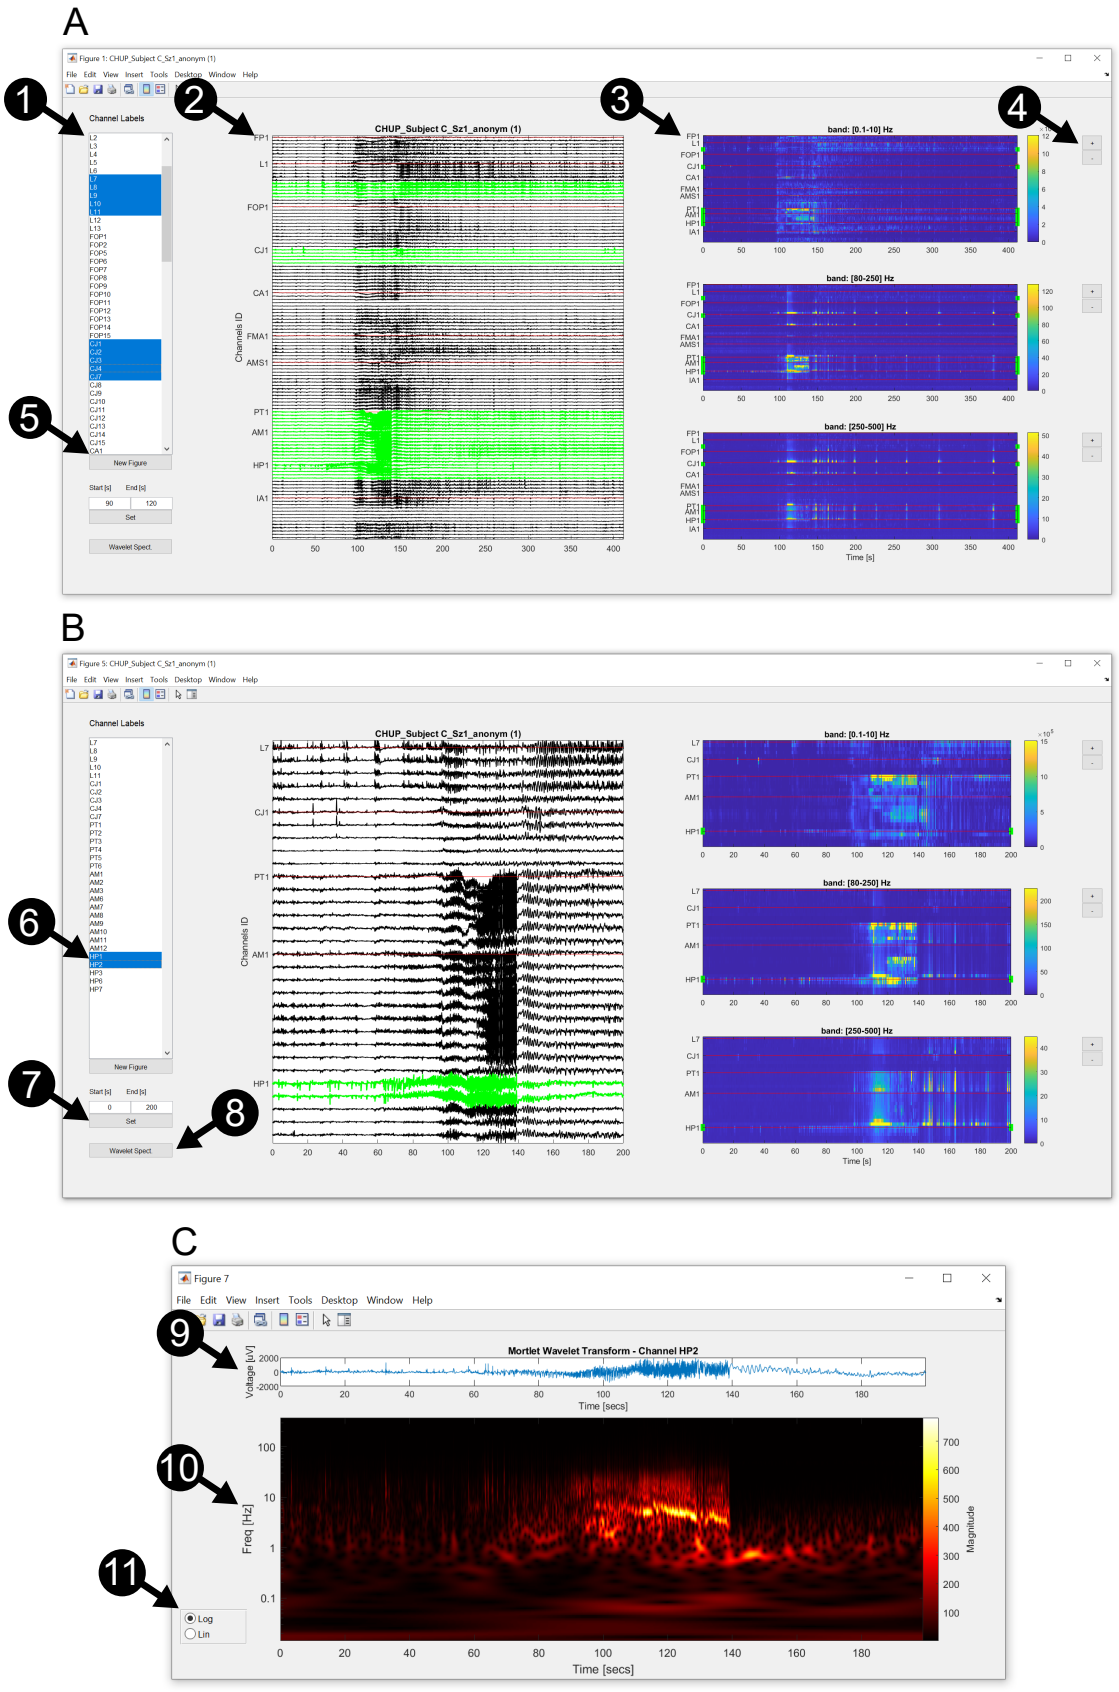
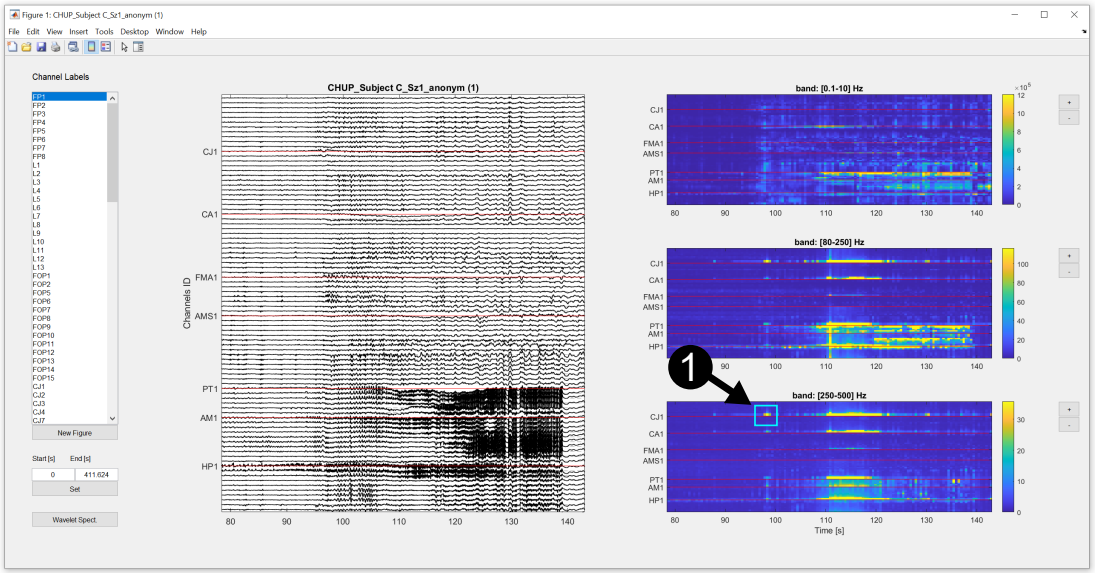



Figure A.2: **Analysis panel for a given sEEG recording**: A) Initial panel with a list of all the channels (1), the sEEG raw data (2), and the band-power for the frequency ranges selected (3). *Continues in the following page...*

Figure A.2: The colour scale of the band-power maps can be adjusted manually (4). The user can select a subset of channels and press "New Figure" (5), to create a new window with the selected channels. B) To calculate the wavelet transform, the user must select the desired channels (6), set the intended time frame limits for the analysis (7) and press "Wavelet Transf" (8). C) A wavelet transform panel will be generated for each of the selected files. In this panel, the user can see the raw voltage data (9), the associated wavelet transform (10), and set the scale of the y-axis to logarithmic of linear (11).*Continues in following page*

A




B

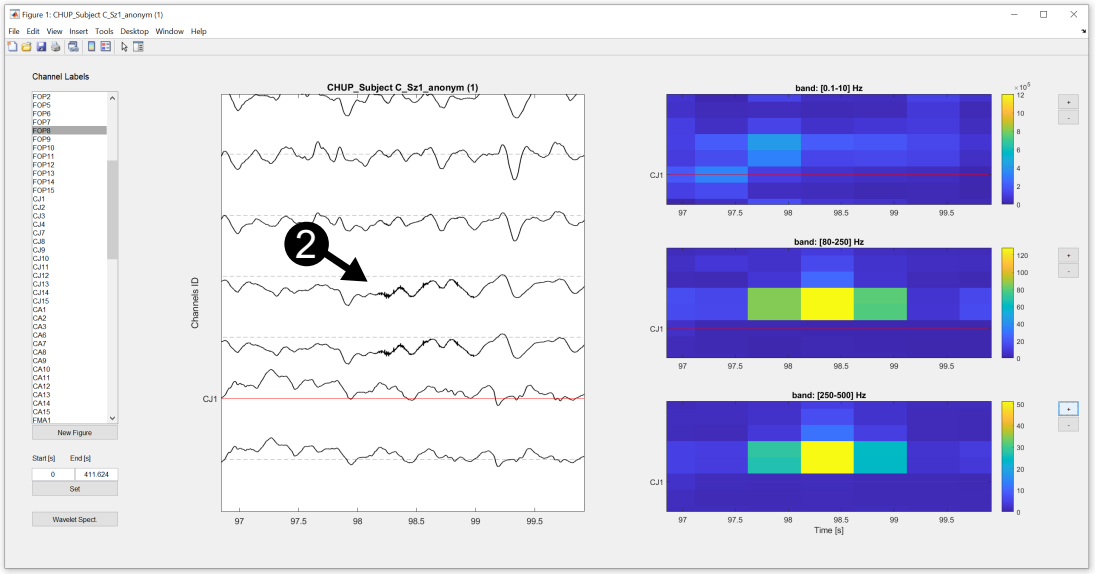



Figure A.3: **PowerMapp can be used to find relevant biomarkers such as High-Frequency Oscillations (HFOs):** A) HFOs can be identified as isolated yellow spots in the high-frequency band-power map (1). B) A zoomed-in view of the sEEg signal at the identified spot shows the HFO (2), which would not be identifiable by looking at the original sEEG time-series.

# Appendix B

# NCN MEA Control Suite

The experiments presented on Chapter 3 were performed using a purpose-built automation software which controlled the acquisition/stimulation hardware and ran the closed-loop algorithms. This program was developed in C# and based on the McsUsbNet dynamic link library (dll) from MCS. despite being undoubtedly useful, the application developed for that work was highly specific, meaning that it could hardly be translated to other closed-loop applications. However, I believe that having full control over the acquisition/stimulation of the MEA devices through software opens many possibilities for closed-loop studies and applications. As such, part of my work was also focused on building a high-level C# API to facilitate the development of applications for the MEA devices from MCS. I named this API of NCN MEA Control Suite, as it was developed in the context of the NCN Lab. The code is available in this link.

The overall goal was to develop a set of libraries for interfacing with the MCS hardware, a set of logic functions to process the acquired data and simple GUI templates for the most basic functionalities, such as recording data, performing spike detection and controlling the stimulators. These three basic primitive operations enable the development of complex closed-loop control setups. The code was to be general for, at least, MEA2100-Mini-System and the MEA2100-256-System, compatible with the multiple MEA layouts of each device, such as 1-well, 6-well and 9-well MEAs.

The code is structured in three separate blocks (three projects within the overall C# Solution in Visual Studio) which have different levels of abstraction. The block/project in charge of interfacing with the hardware is the Mea_Devices (blue block in Figure B.1). This is the lowest-level layer of the NCN MEA Control Suite, interacting directly with the hardware through the McsUsbNet dll (grey block in Figure B.1). This block has three main classes: 1) MeaDaq.cs, which defines the object and functions to interface with the acquisition hardware ; 2) Stimulator.cs, which defines the object and functions to interface with the stimulator; 3) Filter.cs, which defines the object

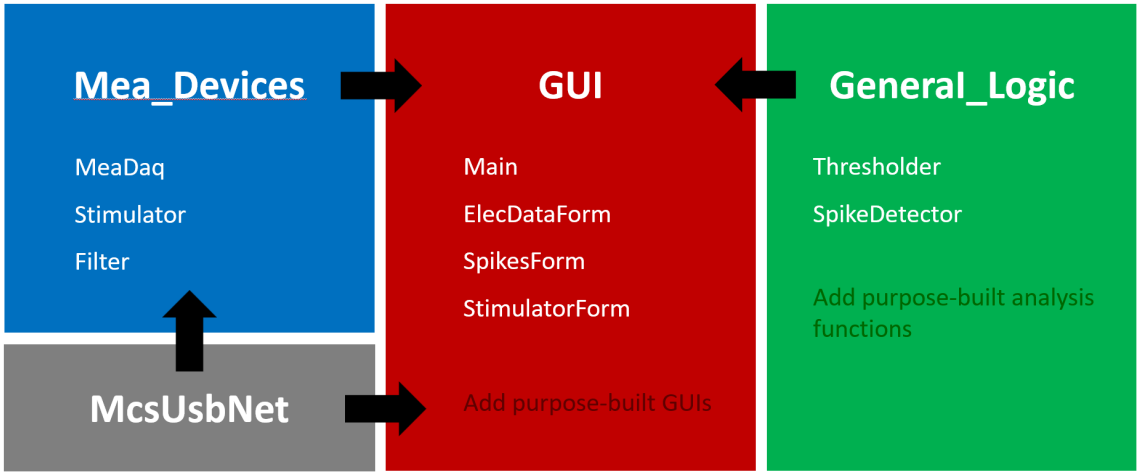



Figure B.1: **Layered architecture of the NCN MEA Control Suite.** The C# solution is organised in three separate blocks/projects: the Mea_Devices layer (blue) which interfaces with the hardware via the McsUsbNet library (gray); the GUI layer (red) which incorporates the User-Interfaces, and thus communicates with all the other layers; and the General_Logic layer (green) with all the classes for signal analysis

and functions to interface with the digital filter implemented in the DSP of the MEA2100 Systems. This low-level layer provides a very intuitive use of the hardware, hiding all the low-level configurations from the user.

The General_Logic project gathers the classes for signal analysis such as the SpikeDetector for spike detection and the Thresholder for finding the spikes threshold (green block in Figure B.1). Other purpose-built functionalities (for example for burst detection, firing rate, spectral analysis, etc.) can also be added here. These are high-level classes, being fully independent of the hardware (they have no dependencies with the Mea_devices project or the McsUsbNet dll). This means that they are general for any device architecture.

Finally, the GUI block/project is where the User-Interfaces are defined (red block in Figure B.1). Here, we implemented the most important UIs that serve as a template for future purpose-built applications. These include the Main.cs for setting up the hardware configurations (Figure B.2); the ElecDataForm.cs for real-time visualization of electrode data (Figure B.3); SpikesForm.cs for performing and visualizing real-time spike detection (Figure B.4); and the StimulatorForm.cs for controlling the stimulators (Figure B.5).

The GUI Main is the opening window (Figure B.2) which allows the user to see the devices connected to the PC and select the desired device and USB Port, as well as selecting the layout of the MEA chip used (1-well, 6-well or 9-well), and acquisition setting the sampling rate. To use the software Multichannel Experimenter from MCS in parallel with the NCN MEA Control Suite, one must select the USB Port B (regardless of the device selected), because the USB Port A is reserved for Experimenter. The user can open the other GUIs using the buttons in the right panel. The button "Electrode Data" forwards the user to the ElecDataForm (Figure B.3). This

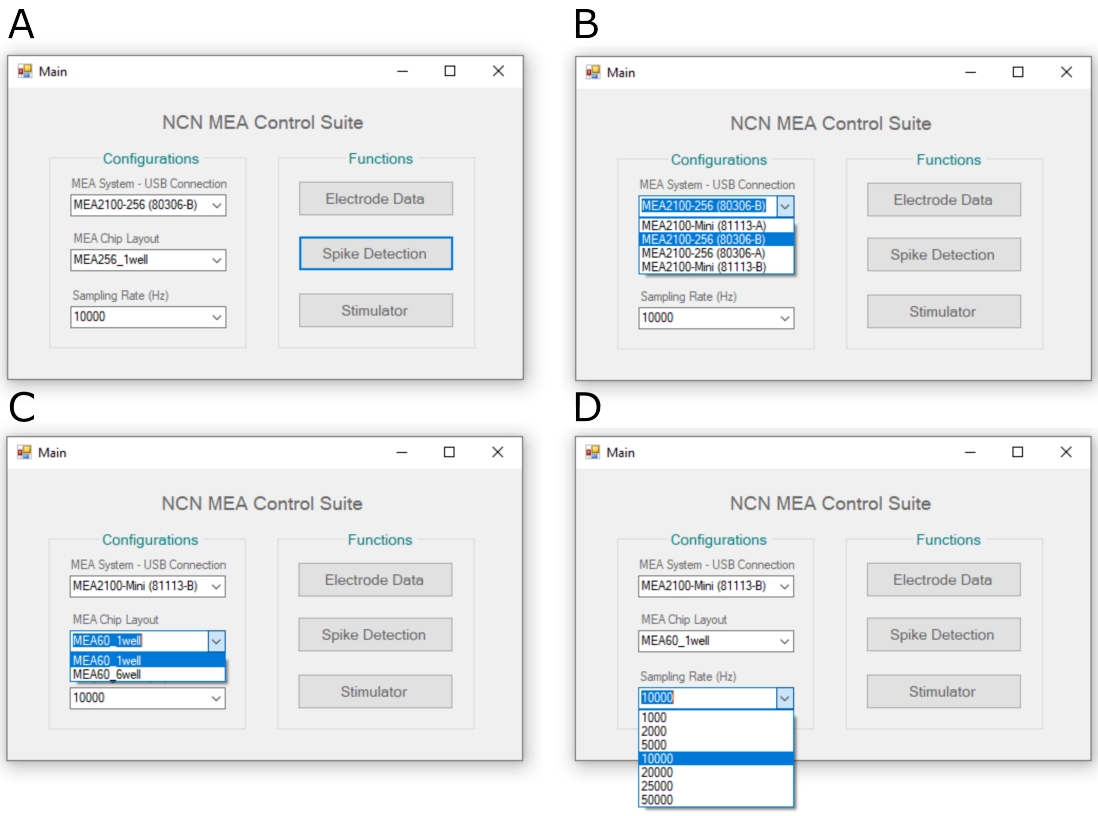



Figure B.2: **Main GUI.** A) opening window with the acquisition hardware configurations of on the left panel and the buttons to open the available functionalities (the other GUIs) on the right. B) Dropdown box to select the acquisition device and USB port from all the devices connected to the computer. C) Dropdown box to select the MEA layout associated with the select device. D) Dropdown box to set the acquisition sampling rate

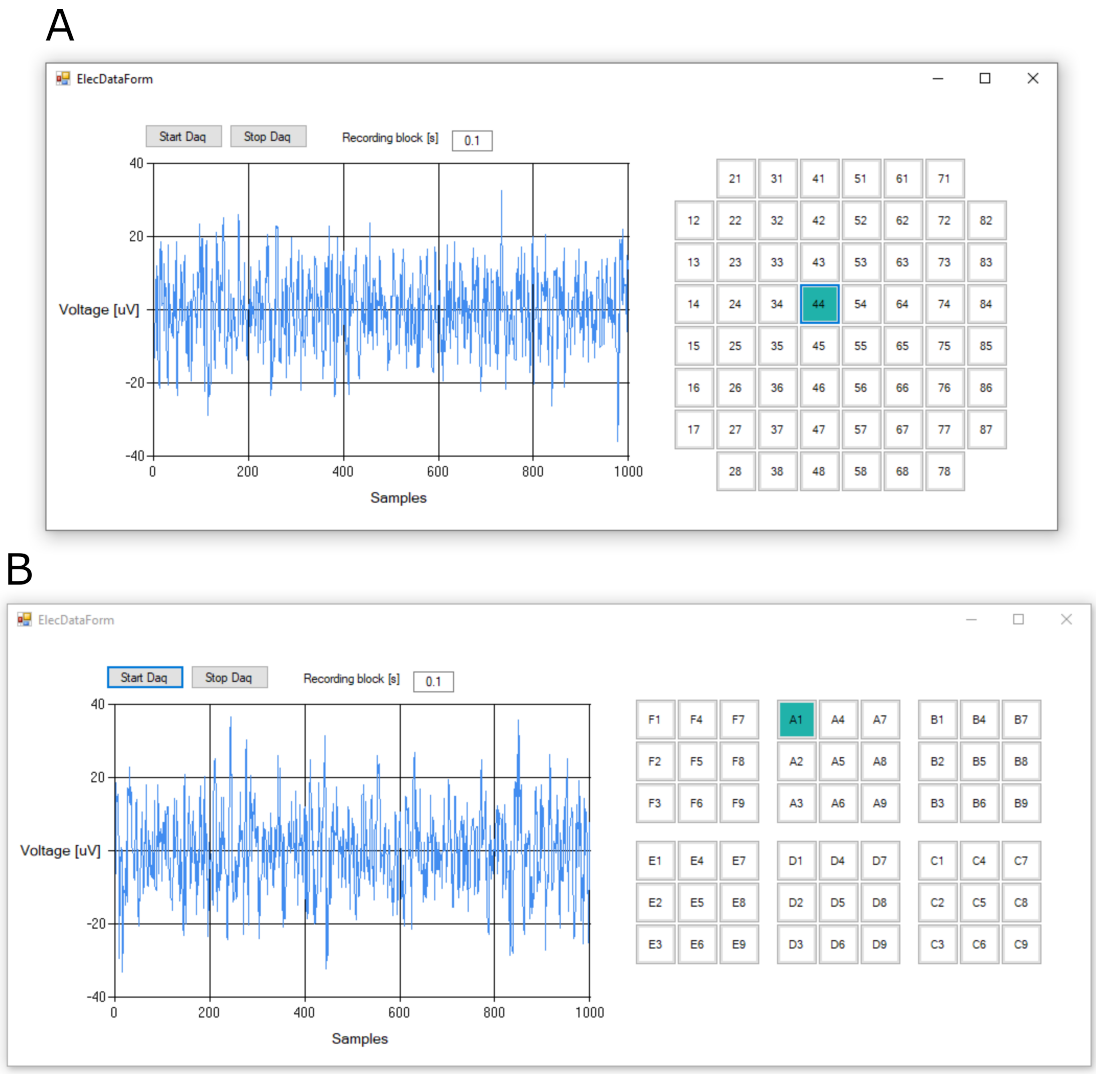



Figure B.3: **ElecDataForm GUI for electrode data visualization**. A) Example of the layout for single well 60-electrode MEA device. B) Example of the layout for the 6-well 60-electrode device. The layout of the electrode grid is set automatically according to the device selected in the Main GUI.

Form allows the user to select an electrode from the MEA an visualise the collected activity in real-time. The matrix of electrodes is set automatically according to the chip layout selected in the Main UI (Figure B.2.C), as shown in Figure B.3 A and B, meaning that the code for setting the electrode buttons positions is not part of the ElecDataForm, but is rather embedded in a lower-level auxiliary class within the GUI project (auxiliary functions are not shown in Figure B.1).

The SpikesForm GUI allows the user to perform and visualise online spike detection (Figure B.4). It is divided in three tabs: Spike Thresholds; Monitoring Electrodes; Spike Detector. The Spike Thresholds tab is used to define the voltage thresholds for the spike detection (Figure B.4.A). The layout of the grid of electrodes showing the recorded signal is automatically set by a lower-level

auxiliary GUI class based on the selected MEA layout. The tab Monitoring Electrodes is used to select the monitoring electrodes where the spike detection will be performed (Figure B.4.B). The tab Spike Detector is where the user will run the spike detection and the results online in a raster plot (Figure B.4.C), where a dot represents a spike detected in a given electrode (Y-axis) at a given time (X-axis).

The StimulatorForm GUI allows the user to control the stimulator manually. The user can download a stimulus waveform to stimulator 1 (green, B.5.A) or 2 (red, B.5.B) and choose the associated stimulus electrodes on the electrode grid. The button "Stimulate" triggers the stimulus on the selected stimulator/s, sending the downloaded waveform/s to the corresponding stimulation electrodes. Just like with all the other UI, the electrode grid of the stimulator is also automatically adjusted to the MEA layout using an auxiliary background function. Figure B.5.C shows the stimulator layout for a 256-electrode MEA with 9 wells with both stimulators activated (green and red buttons on top right corner, STG1 and STG2, are activated).

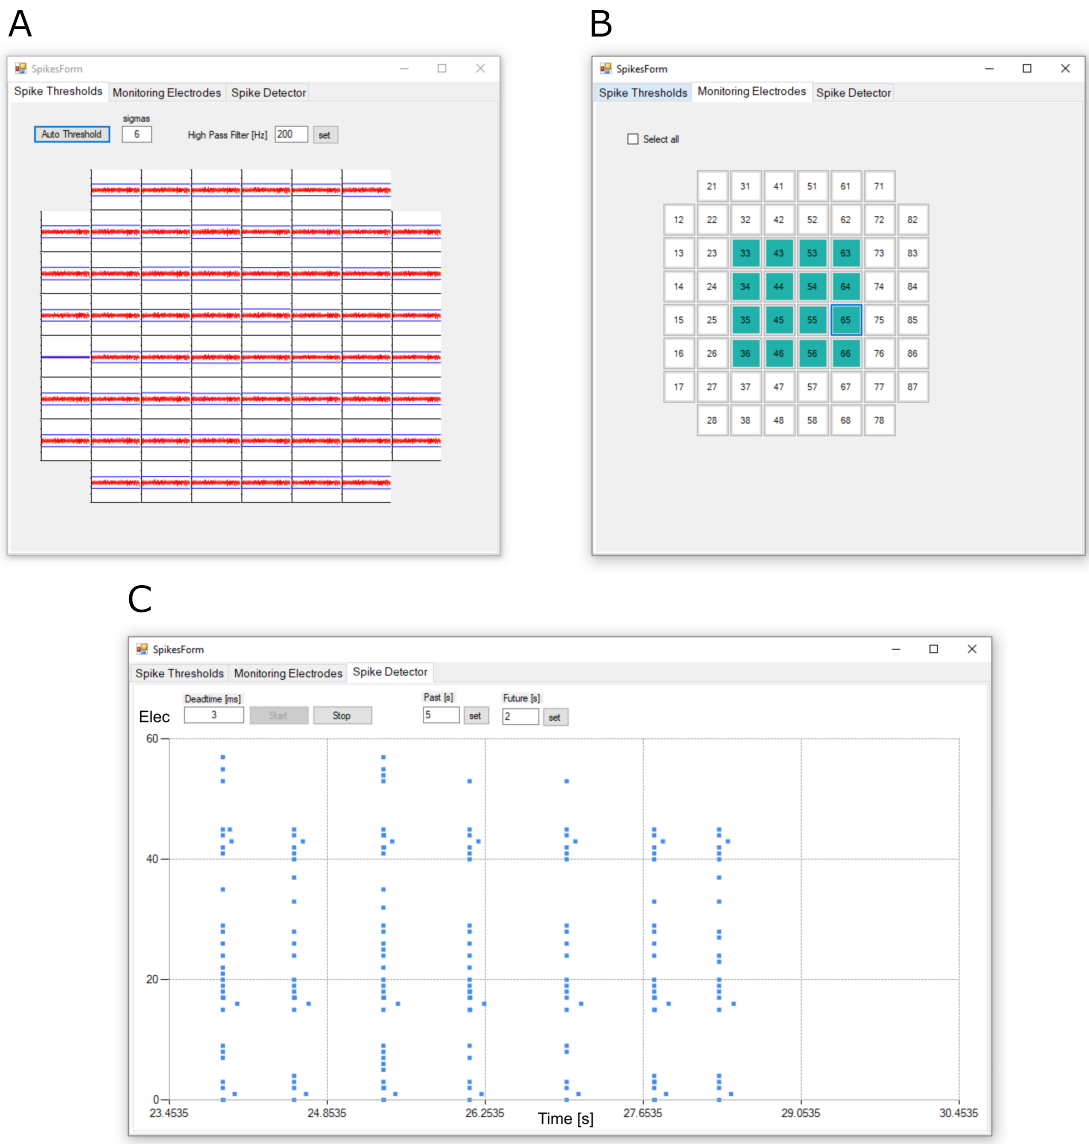



Figure B.4: **SpikesForm GUI for performing online spike detection**. A) Tab to define spike thresholds for each electrode. The threshold is set to N sigmas of background noise for each electrode after filtering the data with a high-pass filter. B) Tab to define the monitoring electrodes. Selected electrode are represented in green. C) Online spike detection tab. Each spike is represented by a blue dot associated with the corresponding electrode along the Y-axis.

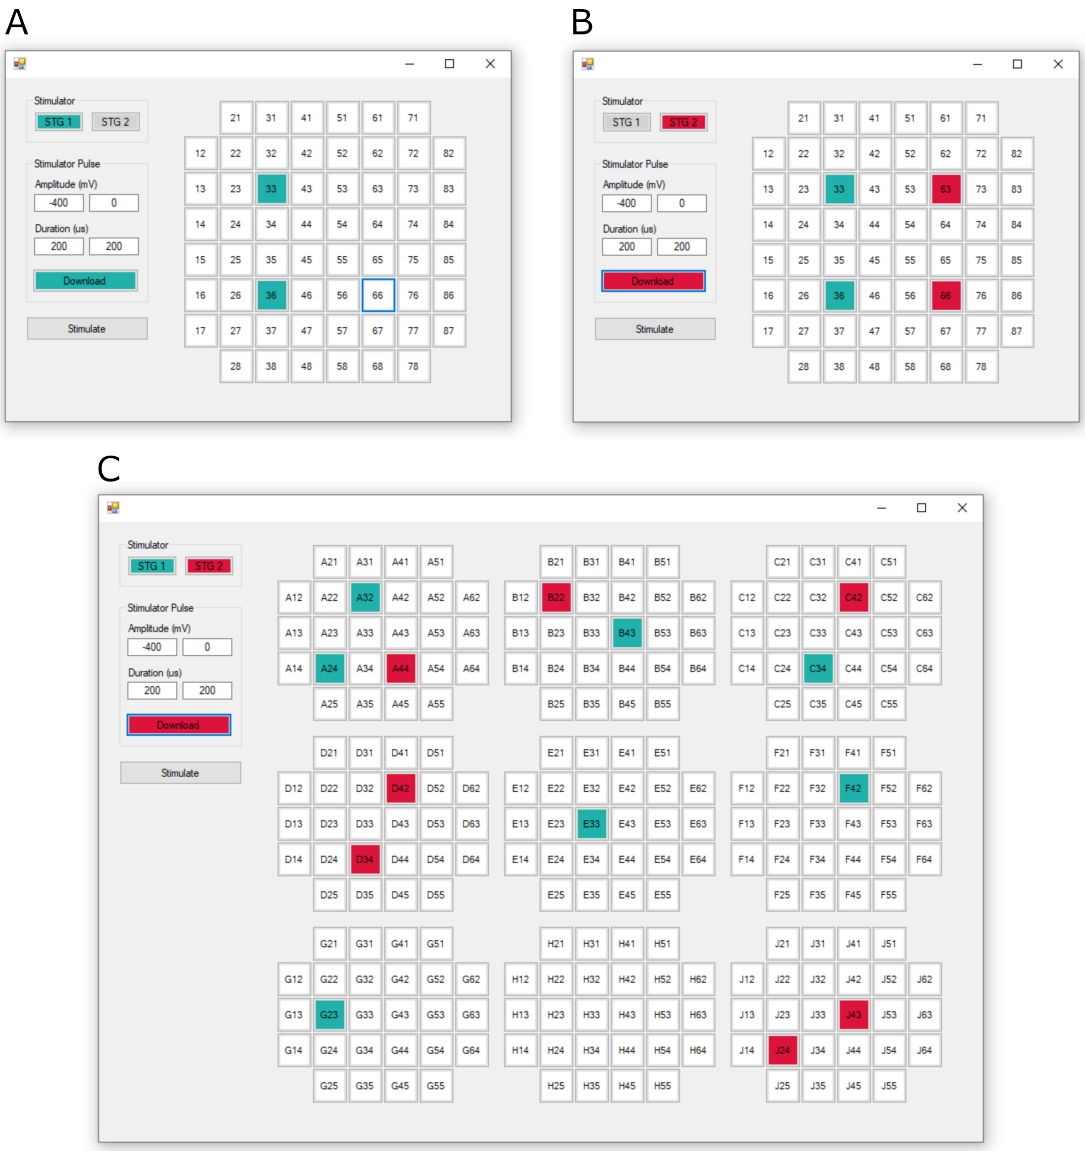



Figure B.5: **StimulatorForm GUI for manual control of the stimulators** A) Setting the stimulation electrodes for stimulator 1 (STG1 in green). B) Setting the stimulation electrodes for stimulator 2 (STG2 in red). C) Example of the stimulation electrodes layout for a MEA device with 256 electrodes and 9 wells.