



UNIVERSITY OF GENOA

MASTER'S PROGRAM IN BIOENGINEERING

Thesis submitted in partial fulfillment of the requirements for the title of
Master of Bioengineering

From biological to digital twin: a pipeline for electrical stimulation of human derived neuronal networks

Matilde Grossi

October 2026

Thesis advisor: Prof. Monica Frega

Thesis co-advisor: Dott. Leonardo Della Mea

Contents

Abstract	7
Chapter 1 Introduction	9
Chapter 2 State of the Art	13
2.1. Neurodevelopmental disorders and patient-derived in vitro models: the biological twin	13
2.1.1. High-Density CMOS Micro-Electrode Arrays.....	15
2.1.2. Fragmented network bursts as an in vitro disease phenotype	17
2.1.3. Electrical Stimulation	18
2.2. <i>In silico</i> models: Digital Twin.....	19
2.2.1. Morphologically Detailed and Point-Neuron Models	21
2.2.2. Extracellular Electrical Stimulation.....	22
Chapter 3 Materials and Methods.....	27
3.1. hiPSC generation and maintenance	27
3.1.1. High-density Micro-Electrode Arrays.....	28
3.1.2. Electrophysiological Data Analysis.....	29
3.2. Cluster Analysis.....	36
3.2.1. Mutual Information	37
3.2.2. Surrogate analysis.....	39
3.3. Neural Dynamics.....	41
3.3.1. Population state-space analysis (Principal Component Analysis).....	41
3.3.2. Effective dimensionality	44
3.3.3. Trajectory speed and state-space extent	45
3.3.4. Entropy and volume of the state repertoire	46
3.4. <i>In silico</i> model	48
3.4.1. Morphologies and three-dimensional-to-two-dimensional reduction	49
3.4.2. Biophysics: Rich human L5 model	52
3.4.3. Model variants	55
3.4.4. The axon and the axon initial segment	55
3.4.5. Numerical compartmentalisation	56
3.4.6. Extracellular field and electrode configuration	57
3.4.7. Stimulation protocol	59

3.4.8.	Resting potential, leak tuning and field-to-membrane coupling	59
3.4.9.	From deterministic single-neuron response to probability	61
3.4.10.	Spatial coordinates relative to the electrode array	63
3.4.11.	Activation region, simulated region and neuronal density	64
3.4.12.	Control simulations on the asymmetry between the two sides.....	64
3.4.13.	Software and computation	65
3.5.	Statistical analysis.....	66
Chapter 4	Results	67
4.1.	Whole-well Super Burst characterisation.....	67
4.1.1.	Super Burst occurrence	67
4.2.	Super Burst temporal organisation.....	69
4.2.1.	Cluster-level Super Burst analysis	70
4.3.	Mutual Information	72
4.3.1.	Surrogate analysis.....	72
4.4.	Neural Dynamics.....	75
4.4.1.	Principal Component Analysis.....	75
4.5.	<i>In silico</i> model	85
4.5.1.	Morphological reduction and numerical resolution	85
4.5.2.	Intrinsic properties of the neuron model	89
4.5.3.	Axon initial segment and stylised axon	91
4.5.4.	Single-pulse approximation.....	93
4.5.5.	Population response to stimulation	95
4.5.6.	Comparison with the preliminary campaign.....	103
4.5.7.	One sample of 2000 neurons.....	105
Chapter 5	Discussion	107
Chapter 6	Future perspectives	111
6.1.	Towards a digital twin of the biological twin.....	111
6.2.	Experimental validation on HD-MEA	112
6.3.	Coupling with an SBI-calibrated network digital twin	113
6.4.	<i>In silico</i> design of stimulation protocols	114
Chapter 7	Limitations	117
Chapter 8	Appendix.....	119
Bibliography		123

Abstract

Neurodevelopmental disorders with epilepsy are often resistant to drugs and the mechanisms by which genetic alterations disrupt the activity of neuronal networks remain largely unknown. Neuronal networks derived from patient human induced pluripotent stem cells (hiPSCs) and grown on multi-electrode arrays provide a biological twin of the patient whose activity can be recorded and perturbed. Several of these disorders share a common network phenotype *in vitro*, in which synchronous network bursts break up into sequences of fragments. This thesis takes Dup15q syndrome as a case study. Patient-derived networks were recorded on high-density multi-electrode arrays, where the phenotype appears as Super Bursts and were stimulated electrically with the electrodes of the array. After stimulation the Super Bursts became shorter and the intervals between them longer in all the wells in which they persisted. In parallel, the fragmented activity was characterised beyond its duration and recurrence, in the state space of the population and at the level of local groups of electrodes, to define what a network model of the culture will have to reproduce. An observed change, however, does not identify the mechanism behind it: this requires a digital twin that includes the stimulation itself, and no existing model describes what an extracellular stimulus on the array does to the neurons of a culture across their positions, orientations and morphologies. To fill this gap, an *in silico* pipeline was developed that couples reconstructed human pyramidal morphologies and human membrane biophysics to the extracellular field of the stimulating electrodes, simulates large samples of independent neurons in which position, orientation, morphology and dendritic complexity vary, and expresses the responses as probabilities of activation, subthreshold depolarisation and subthreshold hyperpolarisation as a function of position. Direct activation was confined to the neurons close to the electrodes, whereas a larger population around them was polarised without firing, depolarised on the side of the cathodes and hyperpolarised on the side of the anodes; the response was set mainly by the position of the neuron and only weakly by the details of the cell. The stimulus therefore acts directly on a small and spatially defined group of neurons, and the changes observed across the network must be carried by synaptic interactions. The probability maps specify how a stimulus acts on each neuron of a network according to its position and are the first component of a digital

twin in which electrical stimulation can be explored *in silico*, towards the identification of the mechanisms by which stimulation modifies pathological network activity.

Chapter 1

Introduction

Neurodevelopmental disorders affect brain development from its earliest stages and are frequently accompanied by epilepsy, which in a large fraction of patients does not respond to drugs [M. Mastrangelo, 2026]. Large sequencing studies have identified many of the genetic alterations involved [Deciphering Developmental Disorders Study, 2015; H.O. Heyne, 2018], but a genetic diagnosis rarely indicates a treatment: between the altered gene and the clinical picture lies the activity of neuronal networks, and how a given alteration changes that activity is in most cases unknown. Animal models have contributed fundamental knowledge, yet differences between species limit how far they reproduce human-specific mechanisms, and the patient's brain can be observed only indirectly.

Human induced pluripotent stem cells (hiPSCs) changed this picture. Somatic cells of a patient can be reprogrammed [K. Takahashi, 2006; K. Takahashi, 2007] and differentiated into neurons [S.M. Chambers, 2009; Y. Zhang, 2013] that carry the patient's genetic background and form functional networks in culture [M. Frega, 2017]. Grown on multi-electrode arrays (MEAs), these networks can be recorded and their activity is reproducible enough to distinguish genotypes [B. Mossink, 2021]. High-density arrays extend this to thousands of electrodes, from the single neuron to the whole network [L. Berdondini, 2009; M. Schröter, 2025; H. Amin, 2016]. A patient-derived network on an array can therefore be regarded as a biological twin of the patient: a controllable system in which the consequences of a genetic alteration on network activity can be measured directly.

In several neurodevelopmental disorders these measurements converge on a common phenotype. Healthy networks produce network bursts, brief episodes in which most neurons fire together; in patient-derived networks these events break up into sequences of fragments [N. Doorn, 2024; E.J.H. van Hugte, 2023; M. Frega, 2019; K.S. Pradeepan, 2024]. Dup15q syndrome, caused by the duplication of the 15q11.2–q13.1 region and associated with autism and drug-resistant epilepsy [M. Elamin, 2023; Y. Perez, 2025; L. Kalsner, 2015], is taken here as a case study. It has a recognised electrophysiological

signature in patients [J. Frohlich, 2016; J. Frohlich, 2019], and neurons derived from patients are hyperexcitable [J.J. Fink, 2021], but its network phenotype on high-density arrays is less well characterised. A phenotype, however, is a description and not an explanation: the same fragmented activity can arise from different combinations of cellular and synaptic alterations, and a recording alone cannot tell them apart.

This is the reason for pairing the biological twin with a digital twin, a computational model of the same network that can be interrogated in ways the culture cannot [L.S. Fekonja, 2024]. Network models fitted to MEA recordings have been used to propose which mechanisms produce fragmented bursts [N. Doorn, 2024; N. Doorn, 2023], and simulation-based inference now allows their parameters to be estimated directly from the recordings of a patient line [N. Doorn, 2025]. These models describe the spontaneous activity of the network. They do not describe what happens when the network is perturbed from outside.

Electrical stimulation is used clinically to modulate brain activity, including in drug-resistant epilepsy, although the mechanisms by which it acts are still debated. *In vitro*, the electrodes of the array can deliver current as well as record [D.R. Merrill, 2005; D.A. Wagenaar, 2004], and stimulation of hiPSC-derived networks evokes reproducible responses and can modify their activity [J. Wen, 2022; G. Zanini, 2023; G. Parodi, 2024]. If stimulation changes the pathological phenotype of a patient-derived network, the question becomes which mechanisms of the network it engages, and answering it requires a digital twin that includes the stimulus. Here the gap is specific. Detailed models of cortical neurons, including human ones, exist [E. Hay, 2011; S. Rich, 2021; H. Mohan, 2015; G. Eyal, 2016; G. Eyal, 2018], but they were built to study the neuron and not its response to an extracellular field. The response of single neurons to extracellular fields has been studied in theory and in simulations with realistic morphologies [F. Rattay, 1986; F. Rattay, 1999; A.S. Aberra, 2018; G.-S. Yi, 2017; M. Komarov, 2019], and at the level of single cells on high-density arrays [S. Ronchi, 2019; M. Radivojevic, 2016], but one neuron at a time, and not across the variability of positions, orientations and morphologies of a culture. And models of stimulation on MEAs describe the field [S. Joucla, 2014] or the network, without a neuron-level account of how the stimulus acts. The work has two parts. In the first, networks derived from a patient with Dup15q syndrome are recorded on high-density MEAs, their fragmented activity is characterised, and the effect of electrical stimulation on it is measured. In the second, the same

stimulation protocol is reproduced *in silico* on detailed models of human neurons, and the question is how the response of a neuron depends on its position relative to the electrodes, and how much this response changes with the morphology of the neuron and with the modelling choices. Two further analyses prepare the step that follows this thesis, a network model of the Dup15q culture in which the probability maps determine how each neuron responds to the stimulus according to its position. The first describes the fragmented events in the state space of the population activity, adding to the usual descriptors of network bursts. The second asks whether the activity of a well can be represented by a small group of electrodes, as a model of reduced size requires.

Chapter 2

State of the Art

This chapter is divided into two parts. The first presents the biological twin: neurodevelopmental disorders and patient-derived *in vitro* models, the high-density micro-electrode arrays used to record them, the fragmented network bursts taken here as the disease phenotype and electrical stimulation. The second presents the digital twin: *in silico* models of neurons at different levels of detail and the modelling of extracellular electrical stimulation.

2.1. Neurodevelopmental disorders and patient-derived *in vitro* models: the biological twin

Neurodevelopmental disorders (NDDs) are a heterogeneous group of disorders characterised by impairments in cognitive, behavioural, motor, or language function. Autism spectrum disorder (ASD), intellectual disability, and epilepsy frequently coexist, and large-scale genetic studies have shown that severe NDDs can result from alterations in numerous genes involved in neuronal differentiation, maturation, synaptic transmission, and neuronal function [Deciphering Developmental Disorders Study, 2015]. Among the most severe forms are developmental and epileptic encephalopathies (DEEs), in which neurological impairment is accompanied by early-onset epilepsy. Their genetic basis is highly heterogeneous and extends beyond ion-channel genes to include many genes regulating synaptic function and neuronal signalling [M. Mastrangelo, 2026; H.O. Heyne, 2018].

Although these alterations originate at the molecular and cellular levels, many of the features most relevant to epilepsy emerge from the synchronised activity of neuronal populations. Linking a genetic alteration to its consequences at the level of the network therefore requires a human, patient-specific experimental model in which population activity can be measured.

Human induced pluripotent stem cells (hiPSCs) are generated by reprogramming adult somatic cells into a pluripotent state, enabling the production of patient-specific human neurons and glial cells for experimental studies [K. Takahashi, 2006; K. Takahashi, 2007]. Their main advantage for disease modelling is that they preserve the donor's complete genetic background, allowing the functional consequences of pathogenic variants and copy-number alterations to be investigated directly in human neurons. Combined with genome-editing technologies, hiPSC models also enable the generation of genetically corrected isogenic controls, substantially reducing the confounding effects of inter-individual genetic variability [M. Elamin, 2023; P.L. Germain, 2017].

Several differentiation strategies have been developed to generate neuronal cultures from hiPSCs. Developmentally guided protocols based on dual inhibition of SMAD signalling generate neural progenitor cells that subsequently differentiate into mixed neuronal and glial populations, providing a physiologically relevant model of early corticogenesis but requiring prolonged differentiation and exhibiting greater variability across cell lines and experimental batches [S.M. Chambers, 2009]. A faster and more reproducible alternative is direct neuronal programming through forced expression of the transcription factor Ngn2, which rapidly converts hiPSCs into relatively homogeneous populations of glutamatergic excitatory neurons [Y. Zhang, 2013]. The protocol was subsequently adapted for multi-electrode array (MEA) recordings, enabling the generation of functionally mature neuronal networks within a few weeks while maintaining tight control over cell density and experimental conditions [M. Frega, 2017]. Although Ngn2 induction does not reproduce the full cellular diversity, its high reproducibility, rapid maturation and reduced experimental variability make it particularly suitable for comparative electrophysiological studies.

Recorded on MEAs, these networks provide a robust and reproducible readout of the relationship between genotype and network phenotype across several NDDs [B. Mossink, 2021]. Within this framework, a patient-derived culture constitutes a biological twin: it preserves the patient's genetic background while providing an experimentally accessible human model in which network activity can be quantitatively recorded and experimentally perturbed, for instance by a defined electrical stimulus delivered to the same culture. A biological twin, however, measures the phenotype without directly identifying the biophysical mechanisms that produce it: it shows what changes, not how.

The case study that motivates the present work is duplication 15q11.2–q13.1 (Dup15q) syndrome, which is investigated in our laboratory with patient-derived Ngn2 networks. Dup15q syndrome is a rare NDD caused by an increased copy number of the imprinted 15q11.2–q13.1 region and is characterised by hypotonia, developmental delay, intellectual disability, ASD, and a high prevalence of epilepsy, which is one of the major determinants of disease severity [L. Kalsner, 2015]. Studies on patient-derived neurons have described a hyperexcitable cellular phenotype [J.J. Fink, 2021] that is largely rescued by correction of the duplication [M. Elamin, 2023].

2.1.1. High-Density CMOS Micro-Electrode Arrays

Microelectrode arrays (MEAs) are devices composed of microfabricated electrodes on which neuronal cultures are grown, allowing extracellular recordings of the electrical activity generated by nearby neurons [L. Berdondini, 2009]. Since the recordings are extracellular and non-invasive, the same neuronal culture can be monitored during maturation, allowing changes in spontaneous network activity to be analysed over time. MEAs complement techniques such as whole-cell patch clamp and calcium imaging by providing direct extracellular recordings from many sites simultaneously with millisecond temporal resolution. Rather than resolving the properties of individual neurons, they capture the collective dynamics of neuronal populations, enabling the quantification of network activity, synchronisation, and bursting [L. Berdondini, 2009; B. Mossink, 2021].

As neuronal cultures mature, spontaneous activity progressively evolves from isolated spikes into bursts and synchronised network bursts. Consequently, neuronal cultures can be characterised through multiple electrophysiological features, including firing rate, burst activity, network-burst dynamics, synchronisation and functional connectivity. Rather than relying on a single metric, the combination of these features provides a more comprehensive description of network organisation.

The robustness of this approach was evaluated by Mossink et al., who analysed Ngn2-derived excitatory neuronal networks from ten healthy donors across multiple experimental batches [B. Mossink, 2021]. Despite biological and technical variability, control cultures showed reproducible developmental trajectories, while disease-specific

alterations remained distinguishable in several neurodevelopmental disorders, including Kleefstra syndrome and MELAS. Importantly, the study also demonstrated that mean firing rate alone is insufficient to describe network behaviour, supporting the use of multiparametric electrophysiological characterisation.

Traditional passive MEAs contain relatively few recording electrodes, limiting their spatial resolution. High-density complementary metal-oxide-semiconductor (HD-CMOS) technology overcomes this limitation by integrating amplifiers and readout electronics directly beneath the electrode array, enabling thousands of densely packed recording sites within a compact area [M. Schröter, 2025; H. Amin, 2016]. This increases spatial sampling while preserving simultaneous recordings across the entire neuronal culture. Under appropriate recording and analysis conditions, HD-MEAs can resolve extracellular electrical footprints associated with individual neurons while maintaining access to population-wide network activity [M. Schröter, 2025; H. Amin, 2016].

These capabilities make HD-MEAs particularly well suited to the present work. Compared with conventional MEAs, they provide a more detailed description of both local and global network organisation while maintaining the non-invasive recordings required for studying neuronal maturation and disease-associated network phenotypes.

A further advantage of both MEA and HD-MEA platforms is their ability to deliver electrical stimulation through selected electrodes while simultaneously recording network activity [H. Amin, 2016]. The same neuronal culture can therefore be characterised under spontaneous conditions, perturbed by controlled electrical inputs and subsequently re-recorded to quantify stimulation-induced changes in network dynamics. This bidirectional interface forms the experimental basis of the stimulation protocol investigated in this work.

However, MEAs also have limitations. Extracellular recordings provide access only to the observable output of the network and do not directly measure membrane potentials, synaptic currents, or other intracellular mechanisms. Moreover, electrical stimulation introduces recording artefacts and its biological effects depend on electrode geometry, stimulus parameters, and the spatial relationship between electrodes and neurons [H. Amin, 2016; D.R. Merrill, 2005]. Consequently, electrophysiological recordings describe the functional phenotype of the network, whereas understanding the underlying biophysical mechanisms requires complementary computational modelling, introduced later in this chapter.

2.1.2. Fragmented network bursts as an in vitro disease phenotype

A network burst is a transient episode of synchronised population activity involving a substantial fraction of the recorded network. While network bursts can differ considerably in frequency, duration and temporal organisation across experimental conditions and disease models, one recurrent phenotype is the presence of multiple closely spaced firing-rate peaks within a single network burst [N. Doorn, 2024].

Doorn et al. defined this pattern as a fragmented network burst [N. Doorn, 2024]. Each local maximum of the population firing-rate profile represents a fragment, whereas the fragmented network burst corresponds to the complete population event containing all fragments. Fragmentation therefore reflects the internal temporal organisation of a single network burst rather than the occurrence of multiple independent bursts.

Fragmented network bursting has been reported in several patient-derived neuronal models, including SCN1A deficiency, Kleefstra syndrome and Rett syndrome, although different studies adopted partially different terminology and detection criteria [N. Doorn, 2024; E.J.H. van Hugte, 2023; M. Frega, 2019; K.S. Pradeepan, 2024]. Together, these observations suggest that fragmentation represents a recurrent alteration of network dynamics rather than a phenotype restricted to a single genetic disorder.

To investigate the mechanisms underlying this phenomenon, Doorn et al. combined experimental recordings with a biophysical network model [N. Doorn, 2024]. Their simulations showed that fragmentation emerged when enhanced short-term synaptic depression repeatedly interrupted recurrent excitation within the same network event. Pharmacological enhancement of synaptic depression in control cultures induced a comparable fragmented phenotype, supporting the proposed mechanism. The model further indicated that asynchronous neurotransmitter release, NMDA receptor-mediated currents and short-term facilitation modulate the number and temporal spacing of fragments. These findings provide a mechanistic framework for fragmented bursting, although they do not imply that identical mechanisms underlie fragmentation in every disease model.

In Dup15q networks recorded in our laboratory, fragments organise into prolonged Super Bursts. In six Dup15q wells, after low-frequency electrical stimulation Super Bursts were shorter in all five wells in which they persisted, and the intervals between them tended to

be longer; control networks were not stimulated. Stimulation therefore modifies a disease-associated network phenotype but the recordings do not show how the stimulus acts on the neurons. This is the question addressed by the computational pipeline developed in this thesis.

2.1.3. Electrical Stimulation

Spontaneous MEA recordings characterise the intrinsic activity of neuronal networks in the absence of external inputs. Electrical stimulation complements this characterisation by introducing a controlled perturbation through selected electrodes while recording the resulting network response on the same platform. In this work, stimulation is used to evaluate the responsiveness of patient-derived neuronal networks and to determine whether the spontaneous Dup15q network phenotype is modified following perturbation. Electrical stimulation must balance biological efficacy with electrochemical safety. At the electrode–electrolyte interface, charge is transferred through both capacitive and Faradaic processes. While capacitive charge transfer is reversible, irreversible Faradaic reactions may alter the electrode surface and damage the surrounding biological preparation. For this reason, stimulation protocols generally employ charge-balanced biphasic pulses, in which the second phase compensates for the charge delivered during the first, limiting residual electrode polarisation and reducing cumulative electrochemical stress [D.R. Merrill, 2005; D.A. Wagenaar, 2004].

In the present work, stimulation is delivered under current control. Under this configuration, the injected charge is determined directly by the pulse amplitude and duration, whereas the electrode voltage varies according to the impedance of the electrode–electrolyte interface. Current-controlled stimulation therefore provides reproducible charge delivery across electrodes and cultures, provided that electrode polarisation remains within the safe operating range of the electrode material [D.R. Merrill, 2005]. The stimulation waveform adopted here consists of positive-then-negative rectangular biphasic pulses, a configuration previously shown to provide effective neuronal recruitment in dissociated neuronal cultures while maintaining electrochemical safety [D.A. Wagenaar, 2004].

The response of neuronal networks to electrical stimulation depends not only on pulse parameters but also on network organisation and cellular composition. Previous studies

on hiPSC-derived neuronal cultures have shown that stimulation efficacy is influenced by pulse amplitude, stimulation frequency, stimulation site, and the balance between excitatory and inhibitory neurons. Wen et al. demonstrated that repeated high-frequency stimulation progressively reduces network recruitment through short-term synaptic depression, whereas low-frequency stimulation elicits more reliable responses [J. Wen, 2022]. Zanini et al. further showed that response probability increases with stimulation amplitude, while Parodi et al. reported that homogeneous excitatory networks respond more reliably than mixed excitatory/inhibitory cultures under current-controlled stimulation on high-density MEAs [G. Zanini, 2023; G. Parodi, 2024]. Together, these studies provide the experimental rationale for the stimulation protocol adopted in the present work.

Electrical stimulation provides a controlled perturbation of patient-derived neuronal networks through which changes in spontaneous network organisation can be quantified. The primary outcome is therefore not only the immediate evoked response but also the comparison of spontaneous activity before and after stimulation, allowing changes in firing, bursting, network-burst fragmentation and network synchronisation to be assessed. Although extracellular recordings alone cannot identify the cellular mechanisms responsible for these changes, they provide the experimental observations that motivate the computational analyses developed in the following chapters.

2.2. In silico models: Digital Twin

Experimental recordings provide direct access to neuronal and network activity, but an observed phenotype does not necessarily identify the cellular and biophysical mechanisms responsible for it. Similar patterns of firing, bursting, or synchronisation may arise from different combinations of intrinsic membrane properties, synaptic interactions, and network organisation. Mechanistic *in silico* models provide a complementary framework in which these mechanisms can be represented explicitly, allowing individual parameters to be manipulated and alternative hypotheses to be tested in a controlled environment [N. Doorn, 2023; N. Doorn, 2025]. In this context, computational models are not intended to replace experimental approaches, but rather to complement them by providing access to mechanisms and perturbations that may be difficult to isolate experimentally.

This approach has been applied to patient-derived neuronal networks. Doorn et al. developed a biophysically detailed *in silico* model of patient-derived human neuronal networks and used it to investigate the mechanisms underlying disease-associated network phenotypes [N. Doorn, 2023]. In a model of Dravet syndrome, the simulations showed that the known $Na_v 1.1$ dysfunction alone was insufficient to reproduce the experimentally observed network phenotype, indicating that additional cellular and synaptic mechanisms contributed to the altered network activity. The model therefore provided a means of linking a population-level phenotype measured by MEA to candidate mechanisms at the cellular and synaptic levels.

More recently, the same modelling framework was extended using simulation-based inference (SBI), allowing distributions of biophysical parameters to be inferred from patient-derived neuronal network activity [N. Doorn, 2025]. Rather than identifying a single parameter set, SBI provides posterior distributions of parameters compatible with the observed activity, explicitly accounting for uncertainty and parameter degeneracy. This extends the role of *in silico* modelling from forward simulation towards probabilistic mechanistic inference and provides a framework for testing how changes in specific cellular mechanisms may alter network behaviour [N. Doorn, 2025].

The broader concept of the digital twin extends mechanistic modelling towards a computational representation that is linked to a corresponding physical system and can be used to simulate, analyse, and predict its behaviour under different conditions [L.S. Fekonja, 2024]. In this context, the distinction between a generic computational model and a digital twin is important: the latter is intended to maintain a meaningful correspondence with a specific biological system through the integration of system-specific experimental information.

For patient-derived neuronal models, this creates a complementary relationship between the biological and computational systems. Experimental measurements provide the observed phenotype and constrain the computational representation, while the *in silico* model provides a controllable environment in which mechanisms and interventions can be explored. Such a framework is particularly relevant when the ultimate objective is not only to characterise disease-associated network dysfunction, but also to predict how the phenotype may change following an external perturbation.

2.2.1. Morphologically Detailed and Point-Neuron Models

The level of detail required in an *in silico* model depends on the biological question being addressed. Multicompartmental neuronal models explicitly represent the spatial structure of the neuron and allow membrane properties and external inputs to vary across soma, dendrites and axon. This spatial resolution is essential when neuronal morphology contributes to the response being modelled, but it also increases computational cost and limits the direct scalability of such models to large neuronal populations.

The importance of morphology in conductance-based modelling was demonstrated by Hay et al., who developed detailed multicompartmental models of layer 5b pyramidal neurons constrained by experimentally measured somatic and dendritic properties [E. Hay, 2011]. Importantly, they showed that the same biophysical parameterisation does not necessarily produce the same electrophysiological behaviour when transferred to a different morphology. Differences in dendritic load and transfer impedance altered the resulting responses, demonstrating that neuronal morphology and membrane biophysics cannot be considered independently [E. Hay, 2011].

This consideration is particularly relevant for human cortical neurons. Rich et al. developed a full-spiking, multicompartmental model of a human layer 5 pyramidal neuron using morphological and electrophysiological measurements obtained from a primary human neuron [S. Rich, 2021]. The model incorporates ten active conductance types and calcium dynamics and was developed with particular emphasis on the properties of the human hyperpolarisation-activated cation current I_h , whose kinetics differ substantially from those described in rodent neurons [S. Rich, 2021].

The original Rich model is associated with the morphology of the individual neuron used for its development. The combination of morphological and electrophysiological measurements from the same cell was essential for constraining the model but does not by itself represent the structural variability present within a population of human pyramidal neurons [S. Rich, 2021]. Human cortical pyramidal neurons exhibit substantial variability in dendritic and axonal architecture, which can be quantified from independent morphological reconstructions [H. Mohan, 2015].

This distinction is particularly relevant when conductance distributions vary with anatomical distance. Applying the same distance-dependent biophysical rules to neurons

with different dendritic extents can place a given conductance profile in different anatomical regions and consequently modify dendritic load and excitability. Rich et al. highlighted the limited availability of detailed human information for constraining the spatial distribution of I_h , making such transfers an important modelling consideration [S. Rich, 2021]. The use of multiple morphologies therefore makes this source of variability explicit rather than embedding it within a single morphological reconstruction.

At the network level, however, a full multicompartmental representation for every neuron rapidly becomes computationally expensive. Reduced point-neuron models provide a complementary solution by describing neuronal dynamics using a small number of state variables while omitting the explicit spatial structure of the cell. Integrate-and-fire and adaptive exponential integrate-and-fire (AdEx) models are widely used for this purpose because they retain key features of neuronal excitability while allowing simulations of large neuronal populations at substantially lower computational cost.

The two representations therefore serve different roles. A detailed multicompartmental model resolves the spatial and biophysical mechanisms required to determine the response of an individual neuron [A.S. Abera, 2018], whereas a reduced point-neuron model provides a tractable representation for large-scale network simulations [N. Doorn, 2024]. For extracellular stimulation, this distinction becomes particularly important because the spatial coupling between neuronal morphology and the external field must first be resolved before the resulting response can be transferred to a reduced network representation.

2.2.2. Extracellular Electrical Stimulation

Electrical stimulation introduces a controlled external perturbation by generating an extracellular electric field around the neuronal membrane. Unlike intracellular current injection, the stimulus is not directly applied to the cell. Instead, the extracellular potential varies in space and interacts with the spatially extended neuronal membrane, making the neuronal response dependent on the position, orientation, and morphology of the cell relative to the stimulating electrodes.

The theoretical basis for modelling extracellular excitation was established through cable-based descriptions of neurons coupled to external electric fields [F. Rattay, 1986;

F. Rattay, 1999]. Rattay showed that the spatial variation of the extracellular potential along a neuronal process contributes to the transmembrane drive, providing the basis for the activating-function framework [F. Rattay, 1986; F. Rattay, 1999]. Consequently, extracellular stimulation cannot in general be described only by the absolute value of the extracellular potential at one location: the spatial relationship between the field and the neuronal structure determines where excitation is preferentially generated.

This principle has been incorporated into detailed computational models of electrical stimulation. Aberra et al. developed biophysically realistic multicompartmental models of cortical neurons coupled to extracellular electric fields and demonstrated that neuronal morphology and active membrane properties influence the resulting stimulation thresholds and membrane polarisation [A.S. Aberra, 2018]. Similarly, Yi et al. showed that the response of multicompartmental hippocampal neurons to a uniform electric field depends strongly on morphological properties, including dendritic geometry and branching [G.-S. Yi, 2017]. Together, these studies demonstrate that a given extracellular field does not necessarily produce the same neuronal response in cells with different morphologies.

The computational modelling of extracellular stimulation generally combines two components: calculation of the electric field generated by the stimulating electrodes and calculation of the neuronal response to this field using a spatially resolved membrane model. Joucla et al. reviewed these approaches in the context of extracellular neural stimulation and microelectrode arrays, highlighting the coupling between electrode geometry, extracellular potential, compartmental cable models, and neuronal activation [S. Joucla, 2014].

The electrode configuration is therefore an integral part of the stimulation model. Electrode number, polarity, spatial arrangement, and distance determine the geometry of the extracellular field and consequently influence the spatial distribution of neuronal recruitment. This becomes particularly relevant for multi-electrode stimulation, where the resulting field may be spatially heterogeneous and anisotropic.

Electrical stimulation has also been investigated directly in human pluripotent stem cell (hPSC)-derived neuronal networks. Wen et al. combined hPSC-derived neuronal networks with MEA recordings, electrical stimulation, and *in silico* network modelling, demonstrating that stimulation responses depend on both temporal and spatial stimulation parameters [J. Wen, 2022]. In particular, local low-frequency stimulation produced

reliable network responses, while multi-site stimulation provided an additional degree of control over network recruitment [J. Wen, 2022]. Zanini et al. further investigated the reliability of evoked responses in hiPSC-derived neuronal networks and showed that response probability depends on stimulation amplitude and frequency [G. Zanini, 2023]. More recently, Parodi et al. examined electrical and chemical modulation of homogeneous and heterogeneous hiPSC-derived networks on high-density arrays, demonstrating that the organisation and cellular composition of the network influence the response to electrical stimulation [G. Parodi, 2024].

These studies establish the feasibility of using electrical stimulation as a controlled perturbation of human neuronal networks. However, the stimulation response in such network-level frameworks is generally represented at the scale of the population rather than derived from the detailed interaction between an extracellular field and the morphology and membrane biophysics of each individual human neuron.

A further reduction of the stimulation problem is to represent neuronal recruitment probabilistically. Komarov et al. developed a framework in which the probability of neuronal activation was described as a function of neuronal properties, cortical layer, and spatial position relative to the stimulation source and the resulting recruitment was incorporated into network simulations [M. Komarov, 2019]. This provides an important conceptual bridge between single-neuron stimulation and large-scale network modelling: rather than explicitly simulating every detailed neuron in the network, the cellular response can be represented by a compact recruitment function.

A closely related experimental development is provided by high-density CMOS MEAs, which enable spatially selective extracellular stimulation of individual neurons and can resolve the relationship between stimulation site and somatic spiking [S. Ronchi, 2019]. These approaches further support the importance of spatially resolved stimulation.

Taken together, the literature has therefore established several complementary components of the problem. Human-specific biophysical models provide mechanistic descriptions of neuronal excitability [S. Rich, 2021], multicompartmental stimulation models establish the spatial coupling between morphology and extracellular fields [F. Rattay, 1986; F. Rattay, 1999; A.S. Aberra, 2018; G.-S. Yi, 2017; S. Joucla, 2014], patient-derived hiPSC models demonstrate how *in silico* approaches can explain and infer mechanisms underlying network phenotypes [N. Doorn, 2023; N. Doorn, 2025], and hiPSC neuronal networks can be experimentally stimulated and modelled at the

network level [J. Wen, 2022; G. Zanini, 2023; G. Parodi, 2024]. Probabilistic recruitment models additionally demonstrate that direct stimulation can be reduced to a computationally tractable population-level representation [M. Komarov, 2019].

What remains less developed is the integration of these components into a single framework in which morphologically detailed human neurons with human-specific biophysics are explicitly coupled to an extracellular stimulation field, and the resulting deterministic responses are subsequently translated to a population-level stimulation rule suitable for large-scale patient-derived network simulations. This is particularly relevant for a patient-specific modelling framework, in which the objective is not simply to reproduce spontaneous network activity but to predict how the network may respond to an external intervention.

Chapter 3

Materials and Methods

This chapter follows the two parts of the work. The first part describes the experimental methods: the generation and culture of the hiPSC-derived networks, the recordings and the electrical stimulation on HD-MEA, and the analysis of the recordings (Super Bursts, clusters and mutual information, population state space). The second part describes the *in silico* model: the morphologies and their reduction, the membrane biophysics, the extracellular field, the stimulation protocol and the construction of the probability maps. The statistical analysis is described at the end of the chapter.

3.1. hiPSC generation and maintenance

A patient-derived human induced pluripotent stem cell (hiPSC) line carrying a 15q11.2–q13.1 duplication was generated by Neurolentech GmbH (Austria). It was derived from an individual with a severe, drug-resistant epileptic phenotype, corresponding to the most clinically challenging end of the Dup15q severity spectrum. This line, characterised at the network level together with a non-isogenic control line, is the line on which the electrophysiological analyses and the subsequent computational modelling presented in this thesis are based.

The hiPSCs were cultured on human laminin-coated plates (5 µg/mL Biolaminin 521 LN, LN521, diluted in phosphate-buffered saline) in StemFlex medium (Thermo Fisher Scientific, A3349401) supplemented with G418 (50 µg/mL, Sigma-Aldrich) and puromycin (0.5 µg/mL, Sigma-Aldrich). Cells were maintained at 37 °C and 5% CO₂. To differentiate hiPSCs into early-stage induced neurons (iNs), colonies were detached using ReLeSR (StemCell Technologies) and re-plated in supplemented medium enriched with RevitaCell (1%, Thermo Fisher Scientific) to enhance cell survival. Doxycycline (4 µg/mL, Sigma-Aldrich) was added to induce Ngn2 expression; induction was considered complete after approximately 3 days of exposure, and the neurons then matured for

several weeks. All cultures were maintained under standard incubator conditions (37 °C, 5% CO₂, and 95% humidity), with half-medium changes performed every 2 days.

Prior to cell seeding (DIV 0), the devices were cleaned with 5% Extran and rinsed with tap water and Milli-Q water. They were subsequently sterilised with 70% ethanol for 1 hour and exposed to UV light for 45 minutes. This process was followed by an overnight coating in the incubator with a mixture of human laminin (20 µg/mL) and poly-L-ornithine (50 µg/mL) diluted in DPBS to promote neuronal adhesion and maturation.

3.1.1. High-density Micro-Electrode Arrays

To investigate the electrophysiological activity of both the Dup15q and control lines, high-density Micro-Electrode Array devices (HD-MEAs) were employed. For this study, the platform used is a 3Brain HyperCAM Alpha (3Brain AG, Switzerland), a CMOS-based HD-MEA built on the CorePlate multiwell format. Each well integrates 4,096 bidirectional microelectrodes at an inter-electrode pitch of 60 µm, of which up to 2,304 can be recorded simultaneously.

3.1.1.1. Electrophysiological Recordings and Electrical Stimulation

The electrical stimulation was performed at DIV 63. We tested biphasic (positive-then-negative) pulses characterised by ± 50 µA in amplitude, 500 µs in duration, and delivered at 0.2 Hz for 3 minutes. Recordings were acquired at a sampling rate of 10110.09 Hz using the HyperCAM Alpha system. The experiment was performed under stable environmental conditions (37 °C with 5% CO₂). Initially, spontaneous neuronal activity was recorded for 20 minutes. During this phase, six electrodes per well exhibiting sustained spiking and bursting activity were selected, as they were considered functionally connected to the network and representative of the overall collective dynamics. These electrodes were subsequently used as stimulation sites. The neuronal networks were then stimulated using the above-mentioned stimulation protocol for 3 minutes (Figure 3.1), and then the activity was recorded again for 10 minutes.




Figure 3.1. Electrical stimulation protocol. (A) Timeline of one experiment: 20 min of spontaneous activity (PRE), during which the six stimulating electrodes were selected, 3 min of stimulation and 10 min of spontaneous activity after stimulation (POST). (B) Stimulation train: 36 pulses delivered at 0.2 Hz over 3 min. (C) Single pulse: biphasic, positive phase first, $\pm 50 \mu\text{A}$, 250 μs per phase (500 μs in total).

3.1.2. Electrophysiological Data Analysis

Off-line data analysis was performed using a combination of BrainWave software (3Brain GmbH) and custom MATLAB scripts (The MathWorks, Natick, MA, USA).

3.1.2.1. Data Preprocessing

Raw electrophysiological signals were band-pass filtered within the 200 Hz – 3 kHz range. Action potentials were identified via the PTSD spike-detection algorithm [A. Maccione, 2009], employing a detection threshold set to 10 standard deviations above the baseline noise level. Both the peak lifetime (spike duration) and refractory period (minimum inter-spike interval) were defined at 2 ms.

Spike detection returned, for each recording electrode i , a binary spike train.

Electrodes rejected by the 3Brain preprocessing pipeline due to excessive noise or unreliable signal quality were excluded. For every remaining electrode, the mean firing rate (MFR) was computed as

$$\text{MFR}_i = \frac{1}{T} \sum_{m=1}^M s_i[m]$$

with M the total number of acquired samples and T the recording duration. Following the activity criterion of Mossink et al. [B. Mossink, 2021], an electrode was classified as active when $\text{MFR}_i \geq 0.1$ Hz. Only electrodes satisfying both the signal-quality and the activity criteria contributed to the subsequent analyses.

3.1.2.2. *Super Burst Analysis*

In this work, the term Super Burst is used as an operational definition for the temporal envelope containing a sequence of closely spaced firing-rate peaks. Individual peaks within this envelope are referred to as fragments, whereas peaks occurring in isolation, outside any Super Burst, are classified as isolated network bursts. The aim of this analysis was to identify and characterise synchronised population activity at these two temporal levels: isolated network bursts and the higher-order fragmented events, the Super Bursts.

Because fragmented network activity is an emergent property of the population, it cannot be inferred reliably from individual electrodes but only from a signal that aggregates the synchronised recruitment of many neurons. Let C denote the set of active electrodes used to reconstruct such a population signal. The same procedure was applied at two spatial scales: C comprised either all active electrodes within a well (Global analysis) or a local 3×3 subset (Cluster analysis).

The population spike train was obtained by summing the binary trains

$$p[m] = \sum_{i \in C} s_i[m]$$

so that $p[m]$ represents the number of electrodes on which a spike was detected simultaneously. Since spikes occupy only a small fraction of the samples at the acquisition sampling rate, the population spike train is highly sparse. To obtain a more robust representation of population activity, it was discretised into consecutive non-

overlapping bins of duration $\Delta = 10ms$, corresponding to $L = \text{round}(\Delta f_s)$ samples per bin.

The 10 ms bin width reduces the sparsity of the original spike train while preserving sufficient temporal resolution to distinguish successive episodes of collective activation. This temporal resolution was inherited from previous analyses performed on the same hiPSC-derived neuronal preparation.

To estimate the temporal envelope of population activity, the binned firing-rate signal was smoothed using locally weighted linear regression (LOWESS), implemented through MATLAB's built-in `lowess` smoothing function with a span of seven bins (approximately $70ms$). The resulting instantaneous firing rate (IFR) provides a continuous representation of collective network activity while suppressing fluctuations. Because fragmented activity is defined at the population level rather than at the level of individual spikes, all subsequent event detection was performed on the IFR.

Episodes of synchronised population activity appear as transient increases in the IFR. Accordingly, local maxima of the IFR were identified as elementary population activation events. Importantly, at this stage each detected event simply identifies one episode of synchronised population activation. It is not yet possible to determine whether that event represents an isolated Network Burst or one fragment belonging to a longer fragmented event. This classification is established subsequently from the temporal relationship between neighbouring events.

First, an adaptive amplitude threshold

$$\text{IFR}[n] \geq \mu_{\text{IFR}} + \sigma_{\text{IFR}}$$

with μ_{IFR} and σ_{IFR} the mean and standard deviation of the IFR. Second, a minimum separation of 0.2 s between consecutive events. Third, a minimum prominence equal to 10% of the maximum IFR in the recording. The adaptive amplitude threshold accounts for differences in the overall activity level between recordings, whereas the prominence and minimum-distance criteria suppress minor local fluctuations and prevent multiple detections of the same activation episode.

Each detected activation event is associated with its occurrence time. Ordering the detected events chronologically yields the sequence:

$$\tau_1 < \tau_2 < \dots < \tau_P$$

where P denotes the total number of detected population activation events and τ_j the occurrence time of the j -th event.

To characterise this temporal organisation, the interval separating two consecutive events was defined as the Network Inter-Burst Interval (NIBI):

$$\text{NIBI}_j = \tau_{j+1} - \tau_j, j = 1, \dots, P - 1$$

If every activation event corresponds to an isolated Network Burst, all intervals belong to the same temporal regime and the empirical NIBI distribution is expected to be approximately unimodal as shown in Figure 3.2B. Conversely, if several activation events belong to the same fragmented episode, two characteristic temporal scales coexist. Short intervals separate fragments belonging to the same event (intra-event intervals), whereas longer intervals separate independent events (inter-event intervals). Consequently, fragmented activity produces a bimodal NIBI distribution (Figure 3.2A).

Rather than imposing a fixed temporal threshold, the boundary between these two temporal regimes was estimated independently for each recording from the empirical NIBI distribution, following the adaptive rationale proposed by Pasquale et al. [V. Pasquale, 2010]. NIBIs were represented on a logarithmic axis using 80 bins spanning the interval between 0.01 and 25 s. The resulting histogram was normalised to obtain an empirical probability distribution and smoothed with a Gaussian kernel.

When two modes were identified, their separation was quantified using the void parameter:

$$\nu = 1 - \frac{h_v}{\min(h_1, h_2)}$$

where h_1 and h_2 denote the heights of the two modes and h_v the height of the intervening minimum. Larger values of ν indicate a clearer separation between the two temporal regimes.

When the separation was sufficient, the intervening minimum $\theta = x_v$, was adopted as the adaptive temporal threshold. Intervals satisfying $\text{NIBI} < \theta$ were classified as intra-event intervals, whereas intervals satisfying $\text{NIBI} \geq \theta$ were classified as inter-event intervals.

3.1.2.3. *Super Burst definition and characterisation*

The adaptive NIBI threshold partitions consecutive population activation events into intervals occurring either within the same fragmented event or between distinct events. Consequently, consecutive activation events connected by successive intra-event intervals were assigned to the same higher-order activity episode.

A Super Burst was operationally defined as the complete temporal envelope encompassing a sequence of activation events linked by intra-event intervals. A Super Burst represents a single fragmented episode of synchronised population activity composed of multiple closely spaced activation events.

Adaptive thresholding identifies which activation events belong to the same fragmented episode, but additional criteria are required to distinguish sustained fragmented activity from short or sparse sequences of peaks. Candidate Super Bursts were therefore retained only when they satisfied two operational criteria inherited from the analysis pipeline previously established for the same hiPSC-derived neuronal preparation.

First, the envelope was required to contain at least five activation events $P_k \geq 5$, where P_k denotes the number of activation events contained within the k -th envelope.

Second, the temporal density of activation events within the envelope was required to satisfy:

$$\frac{P_k}{e_k - a_k} \geq 1.5 \text{ Hz}$$

where a_k and e_k denote the beginning and the end of the corresponding envelope.

The first criterion excludes short activation sequences that are unlikely to represent sustained fragmented activity, whereas the second rejects loosely organised event groupings that are more consistent with independent activation events than with a single fragmented episode. Maintaining these criteria ensured methodological consistency with previous analyses performed on the same neuronal preparation.

The quantity:

$$\frac{P_k}{e_k - a_k}$$

is referred to as the intra-envelope peak rate, as it describes the temporal density of activation events occurring within a single Super Burst rather than the occurrence of independent Network Bursts.

The same Super Burst identification procedure was applied independently to local neuronal populations and to the complete recording well. Consequently, local and global analyses differed only in the electrode population used to reconstruct the IFR, whereas identical temporal criteria were used to define fragmented events.

For recordings containing K Super Bursts, the average duration and recurrence interval were calculated as:

$$\overline{\text{SBD}} = \frac{1}{K} \sum_{k=1}^K \text{SBD}_k$$

and

$$\overline{\text{ISBEI}} = \frac{1}{K-1} \sum_{k=1}^{K-1} \text{ISBEI}_k$$

Together, SBD and ISBEI provide descriptions of fragmented network dynamics, as illustrated in Figure 3.2C. Whereas SBD characterises the duration of fragmented episodes, ISBEI quantifies how frequently such episodes recur. These descriptors therefore summarise the temporal organisation of fragmented population activity.




Figure 3.2. Detection of Super Bursts. (A) Bimodal distribution of the Network Inter-Burst Intervals (NIBI), showing two temporal regimes. The adaptive threshold (θ), defined as the minimum between the two modes, separates short intra-event intervals ($NIBI < \theta$) from longer inter-event intervals ($NIBI \geq \theta$). (B) Unimodal distribution, showing a single temporal regime, in which no distinct separation between intra- and inter-event intervals is present and therefore no adaptive threshold is defined. (C) Illustrative definition of Super Bursts. The instantaneous firing rate (IFR) is shown as a function of time. Isolated peaks occurring outside the highlighted envelopes represent isolated Network Bursts. Four illustrative candidate envelopes, labelled A to D in the panel, are shown. Envelopes A and B satisfy the operational criteria for Super Burst identification, each containing five activation peaks with an intra-envelope peak rate greater than 1.5 Hz (A: $5/2.5 = 2.0$ Hz; B: $5/2.4 = 2.08$ Hz), and are therefore classified as valid Super Bursts. Envelope C is rejected because it contains fewer than five

activation peaks ($P = 3$). Envelope D is rejected because, although it contains five activation peaks, its intra-envelope peak rate ($5/4.4 = 1.14$ Hz) is below the acceptance threshold (1.5 Hz). The Super Burst Duration (SBD) is the temporal interval between the first and last activation peaks belonging to the same Super Burst, whereas the Inter-Super-Burst Event Interval (ISBEI) is the time separating two consecutive Super Bursts.

3.2. Cluster Analysis

The analyses in this section were carried out in view of a computational model of the network, which is not developed in this thesis. A model cannot represent the activity of the 2,304 electrodes of a well for every parameter combination explored during fitting; the first step is to establish whether the activity of the well can be reduced to a small group of electrodes without losing the phenotype. The question is therefore whether a compact local population preserves the temporal features of the whole-well activity and remains functionally coupled to the rest of the network. If it does, it can serve as the biological reference for a network model of reduced size.

Candidate local populations were identified on the 48×48 electrode array according to the spatial distribution of electrode activity. Only electrodes satisfying both the signal-quality criterion and the activity criterion $MFR \geq 0.1$ Hz were considered eligible.

A local population was defined as a compact 3×3 neighbourhood of active electrodes. This geometry represented the largest contiguous square region that could be identified consistently across all recordings while simultaneously satisfying the activity criterion, the signal-quality requirements and the spatial-separation constraints between neighbouring populations.

For a candidate centre located at coordinates (r, c) , the corresponding neighbourhood was defined as

$$W(r, c) = \{(r', c') : |r' - r| \leq 1, |c' - c| \leq 1\}$$

Only neighbourhoods containing nine active and valid electrodes were retained.

Candidate neighbourhoods were ranked according to the activity of their central electrode. Since only neighbourhoods containing nine active electrodes were retained, a highly active central electrode identified local populations located within highly active regions of the network. The selected clusters therefore represent the most active regions of the well and not a random sample of it. To prevent spatial overlap, neighbouring

populations were selected through an algorithm based on the Chebyshev distance D . Candidate centres were accepted only when $D \geq 4$, thereby ensuring that neighbouring 3×3 populations neither overlapped nor directly shared adjacent electrodes. The selection procedure continued until either ten spatially independent clusters had been identified or no additional eligible candidates remained (Figure 3.3).




Figure 3.3. Cluster selection. (A, B) Example of cluster selection on the HD-MEA 48×48 recording grid in a representative well (well A1, MultiWell 2186). Green squares indicate active electrodes ($\text{MFR} \geq 0.1$ Hz) eligible for cluster selection, red squares indicate valid electrodes that did not satisfy the activity criterion, and grey squares represent electrodes automatically excluded by the BrainWave (3Brain GmbH) preprocessing pipeline because of poor signal quality. Ten non-overlapping clusters are shown for the same well before stimulation (A) and after stimulation (B).

3.2.1. Mutual Information

Similarity in Super Burst duration and recurrence does not necessarily imply that two neuronal populations participate in the same network dynamics. Two clusters may exhibit comparable temporal descriptors while remaining temporally independent. The objective of this analysis was therefore to determine whether local neuronal populations were statistically dependent on one another and on the surrounding network.

To address this question, Mutual Information (MI) was used to quantify the statistical dependence between the instantaneous firing-rate (IFR) signals. Unlike linear correlation, which measures only linear relationships, mutual information quantifies both linear and

non-linear statistical dependence and is therefore better suited to the non-Gaussian, event-like dynamics of neuronal population activity. It should be noted that MI is symmetric and non-directional; consequently, a positive MI indicates statistical dependence but does not imply causal interaction or information flow.

Two complementary analyses were performed. First, cluster-to-cluster MI quantified the statistical dependence between pairs of local neuronal populations within the same well. Second, cluster-to-network MI quantified the statistical dependence between each local neuronal population and the activity of the surrounding network. To avoid trivially inflating the estimated dependence, the contribution of the selected cluster was removed from the whole-network signal before MI estimation. Throughout the remainder of this section, cluster-to-network MI therefore refers to the mutual information between a cluster and the activity of the whole network excluding that cluster.

Mutual information was estimated using the first k -nearest-neighbour estimator proposed by Kraskov et al. [A. Kraskov, 2004]. This estimator was selected because it is a non-parametric method that estimates mutual information directly from continuous data without requiring explicit probability-density estimation or arbitrary histogram discretisation, making it particularly suitable for continuous IFR signals. The analysis was performed in MATLAB.

Before MI estimation, all IFR signals belonging to the same recording were truncated to a common length and standardised. Although mutual information is theoretically invariant under linear rescaling, standardisation places all signals on comparable numerical scales, improving the numerical stability of the nearest-neighbour search.

Because prolonged silent periods generate repeated identical IFR values, a negligible uniformly distributed random perturbation was added to avoid degenerate nearest-neighbour distances. The perturbation was several orders of magnitude smaller than the signal itself and is therefore not expected to affect the estimated MI.

The primary analysis used $k = 10$. This value provides a compromise between estimator variance and local resolution. Smaller values of k produce more local estimates but are generally associated with greater variability, whereas larger values produce smoother estimates at the expense of local sensitivity.

For every pair of local clusters i and j , mutual information was calculated between their IFR signals. Repeating this calculation for all cluster pairs yielded a symmetric mutual-information matrix describing the statistical dependence among local neuronal

populations within the same recording. For visualisation only, the off-diagonal entries were normalised by the maximum pairwise MI within the same well.

While cluster-to-cluster MI establishes whether neighbouring populations share information, it does not directly determine whether an individual cluster remains functionally integrated with the dynamics of the entire network. To address this question, cluster-to-network MI was computed for every local neuronal population. The whole-well population spike train was first reconstructed. The contribution of the selected cluster was then removed, yielding a signal representing the activity of the surrounding network. The corresponding IFR was obtained using the same binning, smoothing and standardisation procedure described previously.

A local population whose Super Burst features were comparable with those of the whole well, and whose cluster-to-network MI exceeded the surrogate distribution described below, was considered a candidate reference population for a future network model.

3.2.2. Surrogate analysis

The absolute value of the estimated mutual information is not interpretable on its own: it depends on the estimator and on its parameter k , on the number of samples, and on the properties of each signal and it is positive in finite data even for signals that are not related. The estimate must therefore be compared with a null hypothesis, here that the activity of the cluster and the activity of the surrounding network are statistically independent. Testing it requires a null model, from which surrogate data are drawn that keep the properties of each signal taken alone and remove any relation between the two; the mutual information computed on the surrogates gives the distribution expected under the null hypothesis, against which the observed value is compared.

The null model used here is a circular time shift. For each cluster, the IFR of the surrounding network was circularly shifted by a random temporal delay, generating $S = 1000$ surrogate signals. The temporal displacement was sampled uniformly between $0.2 \leq |\delta_s| \leq 2$ s, with positive and negative shifts selected with equal probability. A circular shift leaves each signal unchanged, including its amplitude distribution and its autocorrelation, and changes only the temporal relation between the two. The mutual information of the surrogates therefore estimates the distribution expected under the null hypothesis.

The observed cluster-to-network MI was compared with the surrogate distribution using a Monte Carlo permutation test. The observed value was considered significant when $p < 0.05$. Shifts of up to 2 s are shorter than a Super Burst, so in these surrogates two long events can still partly overlap and part of the mutual information is preserved. For this reason the test was repeated with a second, stricter null model, in which the shift was drawn uniformly over the whole recording (excluding the first and last 10%), which removes this overlap. Each cluster was tested separately, without correction for multiple comparisons: the tests describe the individual clusters and are not used to draw conclusions about the well. Since the Kraskov estimator depends on the number of neighbours k , the whole analysis was repeated for $k \in \{3, 5, 7, 10\}$.

3.3. Neural Dynamics

This section introduces a complementary framework for characterising neuronal network activity by considering the system as a dynamical process evolving in a state space. Rather than focusing exclusively on conventional firing and bursting metrics, the analysis describes how the collective activity of the neuronal population evolves over time and explores the structure of the states visited by the network. The activity is therefore examined from both a temporal and a dynamical-systems perspective, considering its low-dimensional representation, the evolution and variability of the resulting trajectories and the organisation of the network dynamics in state space. To this end, Principal Component Analysis (PCA), trajectory-based measures and entropy are introduced as complementary tools to characterise different aspects of the underlying network dynamics. These analyses are exploratory and are not required by the stimulation pipeline developed in this thesis: they provide additional quantities that a network model of the Dup15q culture should reproduce, and against which it can be compared.

3.3.1. Population state-space analysis (Principal Component Analysis)

High-density microelectrode arrays sample neuronal population activity across many sites simultaneously, so that the instantaneous firing rate of each electrode contributes one dimension to the description of the network state. For a recording with N active electrodes, each time point is a single point in an N -dimensional state space and the succession of these points over time traces a population trajectory describing how the collective activity evolves [J.P. Cunningham, 2014]. This shifts the description from individual electrodes to the population: a network event is then characterised not only by scalar quantities such as firing rate or duration, but also by the sequence of population states it generates and by the geometry and temporal evolution of the resulting trajectory [J.P. Cunningham, 2014].

During coordinated activity many electrodes vary together, so that their measurements are partly redundant and a large share of the population activity is captured by a few dominant modes. Dimensionality reduction exploits this redundancy to recover a small set of collective variables that describe these dominant modes [J.P. Cunningham, 2014].

Here, dimensionality reduction was performed by PCA, which identifies the orthogonal directions of maximal variance in the population activity. Given the standardised activity matrix $\tilde{R} \in \mathbb{R}^{T \times N}$, with T temporal samples and N electrodes, the electrode covariance matrix is

$$C = \frac{1}{T-1} \tilde{R}^\top \tilde{R} \in \mathbb{R}^{N \times N},$$

and its eigen decomposition $C w_j = \lambda_j w_j$ yields the loading vectors w_j and their variances λ_j . Projecting the activity onto the first three loadings,

$$Z = \tilde{R} W, \quad W = [w_1 \ w_2 \ w_3], \quad Z \in \mathbb{R}^{T \times 3},$$

gives for each time point the population state in the reduced space. A network burst is represented as a movement of the population through this low-dimensional space rather than as a simultaneous rise in single-electrode firing, and the trajectory exposes collective properties such as the extent of state space explored and the temporal organisation of the transitions, which are not accessible from scalar burst descriptors [J.P. Cunningham, 2014; J.A. Gallego, 2017]. This reduced representation relates to the notion of a neural manifold, a lower-dimensional structure embedded in the space of possible activity patterns [J.A. Gallego, 2017]. PCA does not describe the manifold itself: it gives a linear representation of the subspace within which the population dynamics unfold, while the manifold traced in that subspace is in general curved. Principal components are chosen to capture as much variance as possible, so a component does not necessarily correspond to a specific phenotype. In a strongly synchronised culture, in particular, the first component largely reflects the whole network activating together during bursts, a global pattern rather than a specific phenotype; consistently, nearly all electrodes contributed to it with the same sign, so that its score tracked the overall amplitude of population activation and peaked during bursts. Higher components are orthogonal directions that capture progressively less variance. The second and third components may be more closely related to the spatial part of the spatiotemporal structure captured by the decomposition, that is, to differences between electrodes in when and how strongly they take part in an event, but they are not assigned to a specific phenotype here. The representation is therefore used as a compact summary of the dominant coordinated activity, not as evidence about the true dimensionality or the underlying mechanisms of the network.

The instantaneous firing rate of each electrode was obtained using the same 10-ms binning and LOWESS smoothing (span of 7 bins) described for the Super Burst analysis; here, each electrode's smoothed rate was kept separate rather than summed into a single population signal, to preserve the per-electrode structure required by PCA. Comparing PRE and POST trajectories requires both to be expressed in the same coordinate system, and therefore to preserve electrode identity: because each electrode has a fixed weight in every loading vector, the two conditions must use the same electrodes in the same column order [J.A. Gallego, 2018]. The analysis was accordingly restricted to the intersection of the electrodes classified as active ($\text{MFR} \geq 0.1 \text{ Hz}$) in both conditions, kept in the same order throughout. Each electrode's binned rate was then standardised in time, placing electrodes with different absolute rates on a comparable scale so that the PCA reflects coordinated modulation rather than the few most active electrodes. The standardisation was applied to each condition separately, each with its own mean and standard deviation. This removes from the comparison any change in the overall amplitude of the activity of an electrode between conditions and it makes the scale of the POST trajectory depend on the variability of the activity in the POST recording. Because the PRE and POST recordings were of unequal length, both were truncated to the common number of bins before projection, so that the trajectory analysis uses matched-duration segments. As a control, the measures that depend on this scale (peak speed, generalised variance and differential entropy) were also computed after standardising the POST activity with the mean and standard deviation of the PRE recording (common standardisation), so that the two conditions are expressed in the same units.

PCA was computed on the PRE condition alone, giving the loading matrix W_{pre} . The PRE trajectory was obtained by projecting the standardised PRE activity onto these loadings, and the POST activity was projected onto the same PRE-derived basis,

$$Z_{pre} = \tilde{R}_{pre} W_{pre}, \quad Z_{post} = \tilde{R}_{post} W_{pre}.$$

This anchoring is essential, as recomputing PCA separately for each condition would produce two differently rotated coordinate systems, so that apparent trajectory differences could reflect a change of basis rather than a change of dynamics. A fixed PRE basis places both conditions in the same coordinates, so that geometric differences can be interpreted as changes in how activity is distributed within a common state space [J.A. Gallego, 2018].

The number of components retained for the quantitative analysis was determined from the data (see Results): the activity was dominated by a single mode, with a second, smaller but reliable component, so the trajectory measures were computed in the PC1–PC2 plane. The third component contributed negligible variance and was poorly preserved between conditions; it is shown only in the three-dimensional trajectory illustrations.

3.3.2. Effective dimensionality

The number of components that carry coordinated activity was estimated with Parallel Analysis [J.L. Horn, 1965], which compares the eigenvalues of the recording with those expected without coordination. The reference without coordination was built from the recording itself. The train of each electrode was shifted in time by a random amount, different for each electrode, with the end of the recording wrapped around to its start (circular shift). Each electrode keeps its own firing rate and temporal pattern, but the timing between electrodes is lost, so any coordination disappears. The shifted data went through the same steps as the real data (smoothing, standardisation, PCA). This was repeated 100 times for each well and condition, with PRE and POST over the same duration.

Two quantities were obtained from the comparison. The number of significant components counts how many directions are above noise: a component of the recording was counted when its eigenvalue exceeded the 95th percentile of the surrogate eigenvalues of the same order (first with first, second with second, and so on), counting from the first component until the first one that did not. The effective dimensionality describes how the coordinated activity is shared among these directions. The mean surrogate eigenvalue was subtracted from each eigenvalue of the recording, keeping only the excess (negative values set to zero), and the excesses were expressed as fractions of their total, $p_j = \lambda_j^{signal} / \sum_i \lambda_i^{signal}$ and combined into the participation ratio [P. Gao, 2017]

$$D_{eff} = \frac{1}{\sum_j p_j^2},$$

D_{eff} equals 1 when all the coordinated activity lies along a single component and n when it is shared equally among n components.

3.3.3. Trajectory speed and state-space extent

The geometry of a population trajectory describes the path followed by the network state through the reduced state space, whereas its temporal dynamics describe how rapidly the population state moves along that path. All speed measures were computed in the PC1–PC2 plane.

In dynamical-systems terms the trajectory is a time-dependent state vector $Z(t)$ and the speed at which the state moves at each instant is the magnitude of its first temporal derivative:

$$v(t) = \|\dot{Z}(t)\|, \quad \dot{Z}(t) = \frac{dZ(t)}{dt}.$$

Because the trajectory is sampled at discrete time points, this derivative was approximated by central finite differences on samples separated by $\Delta t = 10$ ms:

$$v(t) = \frac{\|\Delta Z(t)\|}{\Delta t} = \frac{1}{\Delta t} \sqrt{\Delta PC_1(t)^2 + \Delta PC_2(t)^2}.$$

The instantaneous speed $v(t)$ is a moment-by-moment time series: it is high when the population state is changing rapidly and low at baseline [J.P. Cunningham, 2014]. The trajectory was smoothed with a Gaussian window of 15 samples (150 ms) before differentiation, and the speed profile with a Gaussian window of 10 samples (100 ms). Each event was summarised by a single value: the maximum of $v(t)$ within each event was retained as the event-level peak trajectory speed, v_{peak} .

Trajectory speed depends, however, on the spatial scale of the trajectory. If the whole trajectory is rescaled by a factor a , its speed is rescaled by the same factor,

$$Z'(t) = a Z(t) \Rightarrow v'(t) = a v(t),$$

so an increase in v_{peak} may reflect either a faster reconfiguration of the population or a larger excursion traversed at the same rhythm. To separate these, the characteristic amplitude of the excursion was defined from the covariance matrix Σ_Z of the population states visited during that event:

$$A = \sqrt{\text{tr}(\Sigma_Z)},$$

where the trace is the sum of the variances along the axes of the state distribution, so that A measures the overall size of the excursion. The peak speed was then normalised by this amplitude, giving the scale-normalised speed

$$\tilde{v} = \frac{v_{peak}}{\sqrt{\text{tr}(\Sigma_Z)}}.$$

If the state distribution is expanded by a factor a , both v_{peak} and $\sqrt{\text{tr}(\Sigma_Z)}$ increase by the same factor, leaving $\tilde{v}$ unchanged. The scale-normalised speed therefore measures the rate of population-state change relative to the size of the excursion; it is a relative quantity and not a physical velocity. The two measures are complementary: the peak speed depends on both the amplitude and the time course of the event, whereas the scale-normalised speed depends only on its time course. Both were averaged over the events of each well.

3.3.4. Entropy and volume of the state repertoire

Beyond the size and speed of the trajectory, the analysis asked how large and how organised the repertoire of visited states is. Differential entropy quantifies the spread of this repertoire, negentropy its structure, and the generalised variance its geometric extent. The entropy and the volume were computed on the cloud of points formed, in the PC1–PC2 plane, by the time points belonging to the events of one type (network bursts or Super Bursts). The differential entropy of this distribution was estimated non-parametrically with the Kozachenko–Leonenko k -nearest-neighbour estimator [L.F. Kozachenko, 1987],

$$\hat{H} = -\psi(k) + \psi(N) + \log c_d + \frac{d}{N} \sum_{i=1}^N \log \varepsilon_k(i),$$

where ψ is the digamma function, N the number of state points, $d = 2$ the dimension of the state space, $\varepsilon_k(i)$ the Euclidean distance from point i to its k -th nearest neighbour ($k = 5$), and c_d the volume of the unit d -ball (π in two dimensions).

The estimate is in nats. The entropy grows when the state points are more spread out, that is, when the network visits a wider variety of states. The estimator assumes independent samples, whereas the state points are consecutive samples of a smoothed trajectory, so that neighbours in state space are often neighbours in time; the subsampling described below reduces this dependence but does not remove it, and the absolute values should be read with this limit in mind.

Differential entropy depends on the scale of the coordinates: if the axes are rescaled, it changes by an additive constant. Because the trajectory has been standardised and rotated by PCA, its units are arbitrary, so the absolute value of $\hat{H}$ carries no biological meaning on its own; only comparisons within one fixed coordinate system do. $\hat{H}$ was therefore

used only to compare PRE and POST within the same well. This comparison is independent of the units only if the two conditions are expressed in the same units, which is why it was repeated with the common standardisation described above. The second measure is the negentropy, defined as the entropy of a Gaussian with the same covariance minus the entropy of the data [A. Hyvärinen, 2000]:

$$\Delta H = H_{gauss} - \hat{H}, \quad H_{gauss} = \frac{1}{2} \log[(2\pi e)^d \det \Sigma_Z]$$

Because the Gaussian has the largest entropy among all distributions with the same covariance [T.M. Cover, 2006], the negentropy is never negative and is zero only for a Gaussian cloud. It does not change when the axes are rescaled or rotated, so it does not depend on the units.

The two quantities describe different things: $\hat{H}$ measures the size of the state repertoire, how spread out the visited states are, whereas ΔH measures its structure, how far the cloud of states departs from a Gaussian cloud of the same size. A repertoire can grow while becoming less structured, which is why both are reported.

The geometric extent of the state cloud was quantified by the generalised variance, $V = \sqrt{\det \Sigma_Z}$, proportional to the area of the covariance ellipse, which has no free parameter. POST/PRE ratios were used as the interpretable quantities, since absolute values depend on scale.

Both the entropy and the volume estimators are biased at finite sample size and depend systematically on N . The PRE and POST clouds contain different numbers of state points, because the events occupy a different number of time bins in the two conditions, so a direct comparison would confound the effect of stimulation with the difference in sample size. To remove this artefact, every metric was estimated inside a common size-matched resampling. On each of $B = 200$ resamples, an equal number of points, $N_{target} = \min(N_{PRE}, N_{POST})$, was drawn at random without replacement from PRE and from POST, and the entropy, the negentropy and the generalised variance, together with their PRE–POST differences, were recomputed. The reported estimates are the means over resamples. The fraction of resamples in which the sign of the PRE–POST difference is preserved is reported as a measure of the stability of the estimate within a well; it is not a test across biological replicates.

The quantities defined in this section are computed here on the recordings only. Once a network model of the culture is available, the same procedure can be applied to the

simulated activity, with the same binning, the same PRE-derived basis and the same measures, and the two compared directly.

3.4. *In silico* model

The aim of the *in silico* pipeline is to translate a biophysically detailed model of human pyramidal neurons into compact quantities that can be used at network scale: the probability that a neuron is directly activated by an extracellular electrical stimulation pulse and the probabilities that it is depolarised or hyperpolarised below threshold, as a function of its position relative to the stimulating array. The central methodological problem is that extracellular stimulation is spatially distributed. The effect of the field depends on how the extracellular potential varies along the extended neuronal membrane and therefore on neuronal morphology, position and orientation. A point neuron can later represent the resulting response, but it cannot explicitly reconstruct this spatial morphology–field coupling. The pipeline therefore starts from reconstructed three-dimensional human layer 2/3 (L2/3) pyramidal morphologies, applies a human layer 5 (L5) biophysical model, simulates the extracellular stimulation at compartment level, and ultimately reduces the detailed response to a population-level probabilistic stimulation rule.

The distinction between the detailed model and the final network representation is intentional. Morphology, channel distributions, axonal structure, compartmentalisation and field geometry are used to determine the response of each simulated neuron; they are not carried into the network model. This reduction is what makes the result computationally usable: a downstream network model, which is a future step and not part of this thesis, would contain thousands of simplified neurons, for which explicitly simulating a multicompartment human neuron for every cell would be computationally prohibitive. Each network neuron therefore needs only a compact rule describing its probability of direct activation and of subthreshold polarisation by the stimulus.

The pipeline consists of six main stages: (1) obtaining and reducing reconstructed three-dimensional morphologies to the quasi-planar geometry of the culture; (2) assigning human L5 biophysics and a controlled stylised axon; (3) defining a spatially resolved compartmentalisation; (4) calculating the extracellular potential generated by the

stimulating electrodes; (5) simulating the membrane response and classifying the outcome at the soma as activation, depolarisation, hyperpolarisation or no effect; and (6) aggregating the responses of an ensemble of simulated neurons into spatial probabilities.

3.4.1. Morphologies and three-dimensional-to-two-dimensional reduction

The reconstructed morphologies are human L2/3 pyramidal neurons from the temporal cortex, obtained from tissue resected during neurosurgery (ex vivo), filled with biocytin in acute slices and reconstructed in three dimensions [H. Mohan, 2015; G. Eyal, 2016; G. Eyal, 2018]. The reconstructions, in NeuroLucida format, are imported into NEURON with Import3d, which converts the geometry into sections and compartments. Six specimens are used: 60308, 130303, 60303, 60311, 130305 and 130306. The pipeline is written so that additional reconstructions can be included without changing the downstream framework. Their role is to provide realistic variation in neuronal size, branch length, branching structure and apical extent.

This differs from the original Rich et al. [S. Rich, 2021] human L5 model. Rich et al. developed their model around a primary human L5 pyramidal neuron for which both a morphological reconstruction and extensive electrophysiological measurements were available. The morphology therefore belonged to the experimental cell used to constrain the model.

In the present pipeline, the morphology source and the channel source are instead decoupled: the geometry comes from the Eyal reconstructions, whereas the human biophysics comes from Rich et al. This allows the same human-specific biophysical framework to be evaluated across multiple human morphologies. Because a complete human L2/3 channel set is not available, L2/3 geometries are combined with L5 biophysics.

The Rich model was designed to reproduce and constrain the electrophysiological behaviour of a specific human L5 neuron. The present approach is instead more appropriate for a population-level stimulation problem because it explicitly samples morphological variability. At the same time, transferring distance-dependent channel distributions from the Rich model to morphologies with different anatomical extents introduces an additional modelling assumption and is therefore a declared limitation.

A reconstructed neuron is an extended three-dimensional object. Its dendrites and axon occupy different positions in space, and therefore different parts of the membrane are exposed to different values of extracellular potential. Under extracellular stimulation, the relevant coupling depends not simply on whether a neuron is close to an electrode, but on how the extracellular potential varies along the neuronal arbour.

The morphology is therefore not a visual representation. Branch length, diameter, orientation and spatial arrangement determine the electrotonic structure of the neuron and influence how the external field is converted into transmembrane polarisation. The detailed neuronal tree is consequently required at the stage at which the morphology–field interaction is computed. A point or single-compartment model cannot reproduce this spatial coupling by itself because it contains no distributed membrane on which the extracellular potential can vary.

3.4.1.1. Reduction of the three-dimensional morphology

The experimental culture is approximately planar: cells grow on a substrate and the electrodes lie in that plane. A fully three-dimensional reconstructed neuron would therefore introduce out-of-plane structure that is not directly representative of the geometry of the culture. The morphology is consequently reduced.

First, a finite-thickness slab is defined around the soma. If z_c is the soma's vertical coordinate and T is the chosen thickness, the slab is

$$z_c - \frac{T}{2} \leq z \leq z_c + \frac{T}{2}.$$

The slab is therefore centred on the soma, rather than on an arbitrary electrode plane.

For each morphological section, the fraction of its arc length lying inside this interval is calculated, and the decision is made at the level of the whole section:

$$f_{\text{in}} = \frac{L_{\text{in}}}{L_{\text{section}}}.$$

A section is retained when $f_{\text{in}} \geq 0.80$, if it is non-axonal and remains connected to the soma through an already retained parent section.

Importantly, the slab does not impose a maximum branch length. A section may extend beyond the slab boundaries and still be retained if at least 80% of its total arc length lies inside the slab. Conversely, when the fraction is below 80%, the section is removed as a whole. Because the tree is traversed from the soma outwards, once a section is excluded,

its downstream subtree is no longer retained, preserving connectivity. The 80% threshold is therefore a criterion on length fraction, not on the fraction of sampled points. Two consequences follow: a retained section can protrude from the slab by up to 20% of its length, and a dendrite lying inside the slab is lost when its parent section leaves it (Figure 3.4).




Figure 3.4. Example of slab reduction ($T = 80 \mu\text{m}$). (A) Side view: the slab $z_c \pm T/2$ and, for each section, the fraction of its arc length inside the slab (f_{in}); blue: retained, red: removed, yellow: axon (always removed and replaced by the stylised axon). (B) The same sections as a tree: a section is retained if $f_{\text{in}} \geq 0.80$ and its parent is retained; B1b-i lies mostly inside the slab but is removed with its parent B1b. (C) Top view after flattening to $z = z_c$; grey: removed. The toy tree is drawn for illustration; the decision uses the same functions as the pipeline.

The reduction acts only along z : the retained arbour is not truncated laterally and can extend far from the soma. For specimen 60308, for example, the retained dendrites reach up to $267\text{ }\mu\text{m}$ from the soma at $T = 40\text{ }\mu\text{m}$ and up to $774\text{ }\mu\text{m}$ at $T = 80$ and $120\text{ }\mu\text{m}$.

In a second operation, all surviving points are flattened onto the soma plane, $z = z_c$, producing the planar morphology used in the subsequent simulations, since in vitro the neurites grow on the substrate. Flattening shortens the branches that have a vertical component, slightly reducing the membrane area.

3.4.1.2. *Layer thickness*

The actual effective thickness of the cultured neuronal layer is not known precisely. Three values, $T = 40, 80$ and $120\text{ }\mu\text{m}$, are therefore simulated for every morphology. This is primarily a sensitivity analysis: if the conclusions are stable across the three values, they do not depend strongly on the precise choice of layer thickness. Whether the outcomes depend on T is reported in the Results.

The three retained morphologies are nested, $K_{40} \subseteq K_{80} \subseteq K_{120}$, so increasing the slab thickness adds previously excluded branches rather than reshuffling the morphology.

The layer thickness also changes the electrotonic load of the reduced cell, because a thicker slab retains more of the original out-of-plane dendritic arbour. It therefore acts not only as a geometric sensitivity parameter but also as a controlled source of variation in cellular excitability. This dependence became particularly relevant in the analysis of depolarisation block.

A limitation of the reduction is that excluding sections can create artificial dendritic terminations. These terminals are treated as sealed ends, so they can develop exaggerated local polarisation under an extracellular field. This is a genuine cable effect at a real terminal but becomes artificial when the apparent terminal exists only because the original dendrite was cut by the reduction. The present analysis assumes that this effect is secondary for the outcome at the soma, because spike initiation is dominated by the soma and the proximal region rather than by the artificial distal tips.

3.4.2. Biophysics: Rich human L5 model

The biophysical component is based on the human L5 pyramidal neuron model of Rich et al. The model contains ten active conductance types derived from the channel set used in

the Hay L5 modelling lineage [E. Hay, 2011] and adjusted using human electrophysiological data. Rich's model was specifically developed to investigate human h-current kinetics and their contribution to intrinsic neuronal behaviour. The ten conductances used in the present implementation are:

1. fast transient sodium, NaTa_t ;
2. persistent sodium, Nap_Et2 ;
3. fast delayed-rectifier potassium, SKv3_1 ;
4. transient potassium, K_Tst ;
5. slow (persistent) potassium, K_Pst ;
6. low-voltage-activated calcium, Ca_LVAst ;
7. high-voltage-activated calcium, Ca_HVA ;
8. calcium-activated potassium, SK_E2 ;
9. M-current, I_m ;
10. hyperpolarisation-activated h-current, I_h .

The calcium-related mechanisms are accompanied by intracellular calcium dynamics.

The roles of the conductances are distributed across different voltage ranges and cellular regions. The sodium currents provide the regenerative upstroke of the action potential; potassium currents contribute to repolarisation; calcium currents provide dendritic nonlinearities; SK_E2 contributes to the calcium-dependent afterhyperpolarization; and I_m and I_h contribute to slower subthreshold dynamics. The present implementation retains the spatially heterogeneous channel distributions of the Rich model: the soma carries all the conductances except I_m ; the apical dendrite carries a low density of NaTa_t , SKv3_1 , SK_E2 , Ca_LVAst , Ca_HVA , I_m and I_h ; the basal dendrites carry I_h only. The placement and the maximal conductances are summarised in Table 3.1.

Table 3.1. Where the channels are (Full Tuned model; maximal conductances in S/cm²).

Conductance	Soma	Apical dendrite	Basal dendrites	Stylised axon
NaTa_t (fast sodium)	2.1	0.001	—	—
Nap_Et2 (persistent sodium)	1×10^{-6}	—	—	—
SKv3_1 (fast potassium)	0.04	0.04	—	—
K_Pst (slow potassium)	0.065	—	—	—
K_Tst (transient potassium)	2×10^{-5}	—	—	—
SK_E2 (calcium-activated potassium)	2.45×10^{-9}	2.45×10^{-9}	—	—
Ca_LVAst (low-threshold calcium)	9.98×10^{-4}	$9.98 \times 10^{-4} (\times 100 \text{ in the hot spot})$	—	—
Ca_HVA (high-threshold calcium)	5.69×10^{-9}	$5.69 \times 10^{-9} (\times 10 \text{ in the hot spot})$	—	—
I_m (M-current)	—	2×10^{-4}	—	—
I_h (h-current)	5.14×10^{-5}	$5.14 \times 10^{-5} \times f(x)$, increasing with distance	5.14×10^{-5}	—
Leak	1.75×10^{-5}	1.75×10^{-5}	1.75×10^{-5}	1.75×10^{-5}

3.4.2.1. How the channels are placed on each morphology

The Rich model was built on one neuron. To use it on the six Eyal morphologies, each channel is assigned by rule rather than compartment by compartment: every section is recognised by its type (soma, apical dendrite, basal dendrite, axon) and receives the conductances of that type (Table 3.1). Two conductances additionally depend on the path distance x from the soma, measured along the dendrite, so that their density is computed for every compartment of every morphology from its own distance.

The first is the exponential gradient of the h-current along the apical dendrite,

$$\bar{g}_h(x) = \bar{g}_h(-0.8696 + 2.0870 e^{3.6161 x/x_{max}}), \quad x_{max} = 1000 \mu\text{m},$$

so that the density is 1.2 times the somatic value at the soma, 2.1 times at 100 μm , 5.3 times at 300 μm and 17 times at 600 μm . The second is the apical calcium hot spot: between 360 and 600 μm from the soma, Ca_LVAst is multiplied by 100 and Ca_HVA by 10. Both rules are written in micrometres, with the same x_{max} and the same hot-spot interval for all morphologies.

Three further adjustments apply to all morphologies: the activation of NaTa_t and SKv3_1 is shifted by 5 and 10 mV, as in the source implementation; the stylised axon is

attached to the soma; and, in Full Tuned, the leak reversal potential is recomputed for each compartment of each morphology.

The passive properties are those of the Rich model: specific leak conductance $g_{pas} = 1.75 \times 10^{-5}$ S/cm², leak reversal potential $E_{pas} = -84.4$ mV (before the tuning described below), axial resistance $R_a = 495.7$ Ω ·cm and specific membrane capacitance $c_m = 1.6$ μ F/cm² (1.0 μ F/cm² at the soma). The reversal potentials are $E_{Na} = +50$ mV and $E_K = -85$ mV.

A crucial consequence of transferring these rules from the original Rich morphology to different Eyal morphologies is that the same path-distance function can occupy different anatomical portions of each cell. The Rich paper explicitly notes that the non-uniform h-channel distribution was derived from rodent observations because sufficiently detailed human information was unavailable, and that differences between human and rodent morphology may affect such a transfer.

In the present study this has two consequences. First, multiple human morphologies allow the dependence of excitability on morphological extent to be explored rather than hidden inside a single reconstruction. Second, the transfer of the same channel-distance rules onto differently extended morphologies is an explicit modelling limitation and a potential source of morphology-dependent excitability.

3.4.3. Model variants

Three variants of the membrane model were implemented.

- **Soma Only:** the full Rich active set at the soma only; dendrites and stylised axon passive.
- **Full Active:** the full Rich active set in soma and dendrites, plus sodium and potassium channels on the axon initial segment (AIS) of the stylised axon (used for the tests of the axon initial segment).
- **Full Tuned** (used for the results presented here): the full Rich active set in soma and dendrites, with a passive stylised axon, as in the Rich model, and a tuned leak.

3.4.4. The axon and the axon initial segment

The reconstructed axons are not retained. In the reconstructions the axon is represented only by a short unbranched segment far from the extent of a cortical axon. The reduction

therefore discards it, and a controlled stylised axon is attached to the soma, as in the published models based on the same cells [G. Eyal, 2016; G. Eyal, 2018]: a straight cable of approximately 6 mm by default, with parametrisable length and a defined AIS. Its length makes it electrically long (about ten length constants), so that its end does not affect the soma. The purpose is not to claim that this geometry reproduces the true human axon but to create a reproducible structure across all morphologies.

The stylised axon is part of the cell in all three variants, including those in which it carries no voltage-gated channels, because it belongs to the template from which the cell is built. A passive axon is in principle not an absent axon, for three reasons. First, it is an electrical load on the soma: it draws axial current and therefore contributes to the input resistance and to the resting state of the cell. Second, under extracellular stimulation the axon is a cable exposed to the field like any other neurite: the field polarises it, and the resulting axial current reaches the soma. Third, keeping the same structure in all variants means that they differ only in where the voltage-gated channels are placed, not in geometry.

In the Full Tuned model used for the campaign, the stylised axon is passive, as in the original Rich implementation; the soma is then the only region with a strong sodium conductance, and spikes initiate there.

In the Full Active variant, the AIS is instead made electrically active with sodium and potassium channels, providing an axonal initiation site. Rich could leave the axon passive in the context of intracellular somatic stimulation, for which the current is injected directly into the soma, whereas under extracellular stimulation the field polarises the initial segment differently from the soma, depending on the position and orientation of the neuron. The synthetic axon remains a declared approximation.

3.4.5. Numerical compartmentalisation

The membrane potential is governed by the cable equation, in which the potential varies continuously along each neurite. A computer cannot solve it at infinitely many points: NEURON divides each morphological section into a finite number of compartments and computes one value of the membrane potential for each. Every compartment is therefore treated as isopotential, i.e. the variation of the potential inside it is neglected. This is an approximation, and it is acceptable only if the compartments are short compared with the

distance over which the potential changes. The d_λ rule [M.L. Hines, 2001] sets this compromise: the compartment length is chosen as a fixed fraction of the AC length constant, the distance over which a signal of frequency f decays along the neurite,

$$\lambda_f = \frac{1}{2} \sqrt{\frac{d}{\pi f R_a c_m}},$$

evaluated at a reference frequency of 100 Hz. The length constant at 100 Hz is used instead of the steady-state one because fast signals decay over shorter distances: for a dendrite of 1 μm diameter with the parameters of the model, the steady-state length constant is about 540 μm , whereas λ_{100} is about 100 μm . The numerical criterion used in the model is $L_{\text{seg}}/\lambda_{100} < 0.1$, so each compartment is about 10% of the length constant or shorter. The segmentation is determined from the passive cable properties so that the geometry is spatially resolved according to this predefined numerical criterion. The rule is applied to each section after the morphological reduction. It is applied to every simulated morphology and layer. The soma is represented by a single compartment, about 22 μm long.

3.4.6. Extracellular field and electrode configuration

The stimulation array is represented by six electrodes in a bipolar configuration: three anodes and three cathodes, each arranged in a column, at the experimental pitch of 60 μm and with an electrode size of 25 μm . The anodes lie at $x = -30 \mu\text{m}$ and the cathodes at $x = +30 \mu\text{m}$, in rows at $y = +30, -30$ and $-90 \mu\text{m}$. All six electrodes are active simultaneously: in the first phase the anodes inject the current and the cathodes return it, and in the second phase the signs are reversed; “anode” and “cathode” refer to the first phase. The field is generated by the superposition of the contributions of the six current sources.

For a single point source on the insulating surface of the MEA, in a homogeneous conductive half-space, the extracellular potential is

$$V_e(r, t) = \frac{I(t)}{2\pi\sigma r}$$

where the factor 2π corresponds to the half-space geometry and $\sigma = 1.5 \text{ S/m}$ is the conductivity of the medium. For multiple electrodes, the contributions add linearly:

$$V_e(\mathbf{x}, t) = \sum_e \frac{s_e I(t)}{2\pi\sigma (\|\mathbf{x} - \mathbf{x}_e\|)}$$

where $s_e = +1$ for an anode and $s_e = -1$ for a cathode, and the distance is bounded below by $r_{min} = 12.5 \mu\text{m}$ (half the electrode size) to avoid the singularity of the point source.

This description relies on five assumptions: the medium is ohmic and linear, so that the six contributions add; the field is quasi-static, which is standard for pulses of hundreds of microseconds; the medium is homogeneous, ignoring the cells, the cell layer on the substrate and the walls of the well; the electrodes are point sources, an approximation that fails within about $12.5 \mu\text{m}$ of an electrode, where the potential is clamped; and the neurons do not modify the field. Under current control, the nonlinearity of the electrode–electrolyte interface does not change the delivered current.

By symmetry, $V_e = 0$ along the midline between the anode and cathode columns, in both phases; the electric field is not zero there: it points from the anode to the cathode column, but it is weaker than near the electrodes.

The field has two zones of opposite sign, which exchange between the two phases (Figure 4.10A). In the first phase the anode column injects current and the cathode column returns it: the extracellular potential is positive on the anode side and negative on the cathode side, with the same magnitude at mirror positions.

The extracellular potential itself is not the membrane potential. The membrane potential satisfies $V_m = V_i - V_e$, and the external stimulus enters the cable equation through the spatial variation of V_e along the neuron. In the compartmental form solved by NEURON, for compartment k ,

$$C_m \frac{dV_{m,k}}{dt} = \sum_j g_{kj} (V_{i,j} - V_{i,k}) - I_{ion,k}, \quad V_i = V_m + V_e,$$

so the differences of V_e between neighbouring compartments drive axial current through the cell, which depolarises the membrane where it exits and hyperpolarises it where it enters; a spike is triggered when the depolarisation reaches threshold where sodium channels are present.

3.4.7. Stimulation protocol

The stimulation protocol reproduces the experimental amplitude and timing. The stimulus is a current-controlled, charge-balanced biphasic pulse of $\pm 50 \mu\text{A}$ per electrode ($150 \mu\text{A}$ in total for each column of three electrodes), anodic first, with two phases of opposite polarity of 0.25 ms each and no interphase interval, for a total duration of 0.5 ms . The phases are implemented with trapezoidal edges of $100 \mu\text{s}$ rather than ideal instantaneous steps, avoiding an unphysical discontinuity in the driving waveform.

The experimental stimulation frequency is $f = 0.2 \text{ Hz}$, corresponding to an inter-pulse interval $T_{\text{ISI}} = 5 \text{ s}$. The slowest relevant recovery processes in the model are much faster than this interval, with I_h evolving on the order of hundreds of milliseconds and calcium dynamics on a similar slow timescale. Consequently, successive pulses are treated as dynamically independent, which was checked by simulating pulse trains, and a single pulse is taken as representative of the individual stimulation events in the experimental train. Each simulated neuron therefore receives one pulse, and the full train need not be simulated for every neuron.

Each simulation starts from the settled resting state of the cell and consists of 5 ms before the pulse, the pulse (0.5 ms) and a post-pulse window of 1000 ms . The 5 ms before the pulse, the pulse and a 1 ms margin after it are integrated with a fixed time step of 0.025 ms ; the rest of the window is integrated with a variable time step (CVODE, absolute tolerance 10^{-6}). Simulations are run at 37°C . The total simulated time is therefore $5 + 0.25 + 0.25 + 1000 = 1005.5 \text{ ms}$.

3.4.8. Resting potential, leak tuning and field-to-membrane coupling

The resting potential is not imposed arbitrarily. The active model is first allowed to settle from an initial value of approximately -85 mV . A fast relaxation moves the membrane toward approximately -67 mV , followed by a slower relaxation toward a stable resting value near -74 mV , driven by the slow h -current kinetics. The settled resting potential is then used as the initial condition for stimulation simulations; in Full Tuned, the leak tuning imposes this rest on the whole arbour, so the initial state is a true equilibrium. This is important because starting directly from an arbitrary membrane potential would

introduce an artificial transient that could be incorrectly attributed to the extracellular pulse.

In Full Tuned the resting state is an equilibrium of the whole cell. The cell is first allowed to settle freely to its somatic resting potential, V_{rest} . Without tuning, the arbour is not isopotential at rest: different compartments rest at different potentials, so current flows between them even without stimulus. Therefore, with the whole arbour at V_{rest} , the leak reversal potential of every compartment is set to $e_{pas} = V_{rest} + i_{other}/g_{pas}$, where i_{other} is the sum of the active currents of that compartment at V_{rest} (sodium, potassium, calcium and h-currents). Each compartment then carries no net current at V_{rest} , the whole arbour rests at the same potential and no axial current flows. Because the unstimulated cell stays at rest, the change in somatic potential induced by the stimulus is measured from the resting potential:

$$\Delta V(t) = V_m(t) - V_{rest}$$

A zero-current run of every morphology and slab thickness confirmed that the unstimulated cell stays at V_{rest} (Results). In the preliminary Soma Only model, whose leak is not tuned, the cell drifts slowly from its initial condition; there the response was measured against such a zero-current run, with the same initial condition, time step and window.

3.4.8.1. Outcome classification

For each simulated neuron, the model receives the complete biphasic pulse, and the outcome is classified into four mutually exclusive classes.

The first criterion is whether the soma produces a spike, defined as the somatic membrane potential crossing 0 mV from below at least once at any time in the simulated window, including the post-pulse window. If it does, the neuron is classified as activated, whatever its potential at the end of the pulse. The spike is read at the soma because it provides a unique and experimentally comparable reference point.

If the soma does not spike, the change in membrane potential is evaluated at the end of the second phase, t_0 :

$$\Delta V(t) = V_m^{stim}(t) - V_{rest}$$

The end of the second phase is used, rather than the peak, because the two phases push the membrane in opposite directions: what matters is the state in which the pulse leaves the neuron, and the last phase dominates it. The subthreshold outcome is then:

$$\Delta V_{\text{end}} \geq +1 \text{ mV} \Rightarrow \text{depolarisation}, \quad \Delta V_{\text{end}} \leq -1 \text{ mV} \Rightarrow \text{hyperpolarisation},$$

and the neuron is classified as not affected otherwise. The raw output stores the value of ΔV_{end} and a sign-only label ($|\Delta V_{\text{end}}| > 10^{-6} \text{ mV}$); the 1 mV threshold is applied in the statistical analysis. The threshold is necessary because far from the electrodes ΔV_{end} is of the order of microvolts but still has a sign, so a sign-only classification would assign about half of the distant neurons to each subthreshold class. The value of 1 mV is a modelling choice. ΔV rather than the absolute membrane potential is used because the absolute value depends on each neuron's own resting potential and on the drift of the model, whereas the quantity of interest is the change produced by the stimulus.

3.4.9. From deterministic single-neuron response to probability

For a fixed morphology, position, orientation and layer, the model response is deterministic: a particular neuron either fires or does not fire, and its subthreshold outcome is fixed. The probability therefore does not belong to the individual neuron. It emerges from sampling a population of neurons with different morphologies, positions and orientations.

A sample is one random realisation of this population, made of independent neurons. For each sample, the morphologies are assigned to somas in equal proportions and in random order, the somas are placed uniformly in a square of $600 \times 600 \mu\text{m}$ centred on the centre of the stimulating dipole, and each soma receives a random orientation, uniform between 0° and 360° . The draws are generated from a seed and the sample index, so each sample is reproducible. The same realisation is then simulated at the three slab thicknesses. In the current design,

$$N_{\text{simulations}} = N_{\text{samples}} \times N_{\text{somas}} \times 3,$$

with 432 samples of 2000 somas, i.e. 864,000 neurons and 2,592,000 simulations. The morphologies are already represented among the somas of each sample; they are not multiplied by an additional factorial morphology term.

Each neuron is simulated alone: the neurons are independent and uncoupled, with no synapses and no effect on each other or on the field. This independence makes the outcomes of different neurons independent draws of the same random experiment, so that their fractions in each spatial bin are faithful estimates of the outcome probabilities. A sample is therefore a sampling unit, not a model of a culture or of tissue density: the 2000 somas are concentrated in the region where the stimulus acts, instead of spending most simulations on distant neurons with negligible responses. The per-neuron probabilities do not depend on this choice.

For a given spatial bin, the empirical probability of each outcome is the fraction of observations with that outcome, for instance

$$\hat{p}_{\text{act}} = \frac{n_{\text{fired}}}{n_{\text{observations}}}$$

and analogously for depolarisation and hyperpolarisation. The four classes are mutually exclusive, so the probabilities sum to one in each bin. Because the samples cover morphology, layer and orientation in fixed proportions, counting the fraction in a bin is equivalent to conditioning on position and marginalising over these variables:

$$P(Y \mid r, \theta) = \sum_m \sum_{\ell} p(m) p(\ell) \int_0^{2\pi} P(Y \mid r, \theta, \varphi, m, \ell) p(\varphi) d\varphi,$$

with $p(m) = 1/6$, $p(\ell) = 1/3$ and $p(\varphi) = 1/2\pi$, where Y is one of the outcomes. The variables are not discarded: their effect is what places the probability between 0 and 1. The number of observations per bin is not fixed.

The uncertainty of each probability is reported as a 95% Wilson confidence interval [E.B. Wilson, 1927]. The interval treats the observations as independent. Each neuron contributes three observations, one per slab thickness. They describe the precision of the estimate for the simulated population.

3.4.10. Spatial coordinates relative to the electrode array

A single distance from the array does not necessarily capture the geometry of a bipolar field. Two somas at the same radial distance can lie in different directions relative to the dipole and therefore experience different spatial field distributions.

The dipole reference frame is constructed from the centroid of the three anodes, $\mathbf{a}$, and the centroid of the three cathodes, $\mathbf{k}$:

$$\mathbf{c} = \frac{\mathbf{a} + \mathbf{k}}{2}, \quad \mathbf{d} = \frac{\mathbf{a} - \mathbf{k}}{\|\mathbf{a} - \mathbf{k}\|}.$$

Here $\mathbf{c}$ is the dipole centre and $\mathbf{d}$ is the unit vector pointing from the cathode side toward the anode side. For a soma at in-plane position $\mathbf{x}_i$, the in-plane offset is $\mathbf{r}_i = \mathbf{x}_i - \mathbf{c}$ and the distance used is $r_i = \sqrt{\|\mathbf{r}_i\|^2}$

The directional coordinate is the angle between the in-plane offset and the dipole axis:

$$\theta_i = \text{atan2}(\|\mathbf{r}_i \times \mathbf{d}\|, \mathbf{r}_i \cdot \mathbf{d}) \in [0, \pi].$$

The angle is deliberately kept unfolded over 0° – 180° , so that the two ends of the dipole are not merged. In the adopted convention, 0° points toward the anode and 180° toward the cathode. The resulting positional representation is therefore $P = P(r, \theta)$. Probabilities are estimated in bins of $5 \mu\text{m}$ in r and 5° in θ and for the maps in the plane of the array, in bins of $10 \times 10 \mu\text{m}$ in (x, y) .

The orientation of the neuron is a separate angle. It describes how the morphology is rotated around its own position. Because neuronal orientation is drawn at random for each neuron, it is not retained as a primary coordinate of the final probability; instead, its effect is integrated over by the population sampling.

Thus position is described by (r, θ) . All the other variables are averaged over: they influence the response of each individual neuron, but the probability at a given position is the average over the simulated population.

3.4.11. Activation region, simulated region and neuronal density

Three quantities are kept distinct. The activation region is the part of the plane in which somas can be directly activated by the pulse: in the final campaign, the probability of activation fell below 0.01 at 120 μm from the dipole centre and no neuron fired beyond 140 μm (Results). The simulated region is the square in which the somas of each sample are placed (half-side 300 μm , i.e. 0.36 mm^2 , centred on the dipole centre). It is about twice as wide as the activation region, so that it also contains the region in which subthreshold polarisation above the 1 mV threshold still occurs and all the effects of the stimulus fall inside it.

The number of somas per sample (2000) is a sampling choice, not a tissue density: it corresponds to about 5600 neurons/ mm^2 , about three times the density of 1700 neurons/ mm^2 assumed for the culture. Because each neuron is simulated alone, the density does not change the probability of each outcome at a given position; it only sets how many observations fall in each spatial bin and therefore the width of the confidence intervals.

3.4.12. Control simulations on the asymmetry between the two sides

Three sets of control simulations were run to identify the origin of the asymmetry between the anode and the cathode side (Results). They used one reduced morphology with the 80 μm slab and the same field, pulse amplitude and outcome classification as the campaign, with about 2000 soma positions and random orientations in a $300 \times 300 \mu\text{m}$ square centred on the dipole. In the first set, every neuron was stimulated with the campaign pulse and with the two phases in the opposite order, in the Full Tuned and in the Soma Only model. In the second, the resting potential of the Soma Only model was shifted to -74.3 mV , the value of Full Tuned on this morphology, by adding a constant to the leak reversal potential of all compartments, with no other change. In the third, somas were placed 20–35 μm or 60–80 μm from the nearest electrode of one column, on its

outer side, and the membrane potential of every compartment was recorded at the end of each phase, together with the time course at the soma.

3.4.13. Software and computation

Simulations were written in Python and run with NEURON through its Python interface; the reconstructions were read with *morphio*, and the analysis used numpy, pandas and matplotlib. The simulations were run on the davinci-1 high-performance computing facility as batch jobs (PBS) distributed over parallel workers, each keeping one live cell that is rebuilt when the morphology or the layer changes. A single configuration file defines every parameter of the pipeline; the model variant is recorded in every output row, results of different variants are written, merged and analysed separately, the draws are generated from a seed and the sample index, and identical duplicated observations are removed before the analysis.

Code availability. The code of the *in silico* pipeline and of the analyses presented in this thesis is publicly available at <https://github.com/matig331/ElectricalSimulation> under the MIT licence.

3.5. Statistical analysis

Statistical analyses were performed using GraphPad Prism (GraphPad Software, San Diego, CA, USA). The well was taken as the experimental unit, and each well was compared with itself before (PRE) and after (POST) electrical stimulation; clusters within a well were not treated as independent samples. Six wells from two MultiWell plates were analysed ($n = 6$). In one well (B2, MultiWell 2186) Super Bursts were abolished after stimulation, so the measures that are defined only when Super Bursts are present in both conditions could not be computed for it. These measures (Super Burst duration, inter-Super-Burst-event interval, trajectory speed, inter-fragment interval, generalised variance, negentropy and differential entropy) were therefore analysed in the remaining five wells ($n = 5$). All the other measures (Super Burst rate, fragment rate, cluster-to-network and cluster-to-cluster mutual information, variance explained by the principal components and effective dimensionality) were analysed in all six wells ($n = 6$). Because the comparisons are paired, the normality of the paired differences (POST – PRE) was first assessed with the Shapiro–Wilk test, which is recommended for small sample sizes, to choose the most appropriate statistical test. For normally distributed differences, PRE and POST were compared using the two-tailed paired t-test, which tests whether the mean of the paired differences is different from zero. For non-normally distributed differences, the comparison was performed using the Wilcoxon matched-pairs signed-rank test.

Data are presented as mean $\pm$ standard error of the mean (SEM). Statistical significance was considered for p -values < 0.05 and is indicated in the figures as follows: ns, not significant ($p \geq 0.05$); *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$. The significance of the mutual information against surrogate data and the confidence intervals of the simulated probabilities are described in the corresponding sections.

Chapter 4

Results

The results are presented in two parts. The first part concerns the *in vitro* experiments: the Super Bursts of the Dup15q networks before and after electrical stimulation, at the level of the whole well and of clusters of electrodes, the mutual information between the clusters and the rest of the network, and the population dynamics in the state space. The second part concerns the *in silico* model: the checks of the single-neuron model, and then the response of a simulated population of neurons to the stimulation pulse, expressed as spatial probabilities of activation, depolarisation and hyperpolarisation.

4.1. Whole-well Super Burst characterisation

To characterise the effect of electrical stimulation on the global network dynamics, Super Bursts were identified from the instantaneous firing rate (IFR) following the detection pipeline described in Chapter 3. Representative examples of the detection are shown in Figure 4.1A–D, where the whole-well IFR of a representative recording is displayed before and after stimulation together with the detected IFR peaks, Super Burst envelopes and the corresponding adaptive NIBI threshold.

4.1.1. Super Burst occurrence

The total number of Super Bursts detected in each recording was lower after stimulation in five of the six wells. This count is not directly comparable between conditions because the post-stimulation recording was considerably shorter than the pre-stimulation recording (9.55 min versus 20 min). Because a shorter window captures fewer events, the counts were normalised: the number of Super Bursts in each well was divided by its recording duration and expressed as a rate in events per minute (min^{-1}); the same normalisation was applied to the fragment count. Once expressed as a rate, the Super Burst frequency did not differ significantly between the two conditions (paired t-test, $p =$

0.65) and the response was heterogeneous across cultures, increasing in some wells and decreasing in others (Figure 4.1F). The reduction in the total count is therefore attributable to the shorter post-stimulation recording rather than to a change in the frequency of Super Bursts. In one well (B2, MultiWell 2186), Super Bursts disappeared entirely after stimulation; with a single well and no unstimulated control, this cannot be attributed to stimulation with certainty. The fragment rate, in contrast to the Super Burst rate, decreased in every well after stimulation (paired t-test, $p = 0.007$; Figure 4.1E), showing that, although Super Bursts occurred at a similar rate, each event contained fewer peaks after stimulation. Per-well values are reported in Tables A.1 and A.2.




Figure 4.1. Whole-well Super Burst detection and event rates before and after stimulation. (A, C) Whole-well instantaneous firing rate (IFR) of a representative recording (well A1, MultiWell 2186) before (A) and after (C) stimulation; the x-axis is time (s) and the y-axis the IFR (Hz). Red dots mark the detected IFR peaks and yellow shaded regions the intervals classified as Super Bursts. (B, D) Corresponding distributions of the inter-peak intervals (NIBI), on a logarithmic axis, with the adaptive threshold θ used to group consecutive peaks into the same Super Burst. The same detection procedure and parameters were applied to the two recordings. (E) Fragment rate and (F) Super Burst rate, in events per minute, before and after stimulation across the

six analysed wells (MultiWells 2186 and 2013). Bars: mean across wells; error bars: standard error of the mean (SEM). Differences between PRE and POST were assessed with two-tailed paired t-tests, each well being one paired observation (ns, $p \geq 0.05$; *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$).

4.2. Super Burst temporal organisation

To determine whether stimulation modified the temporal organisation of fragmented network activity, the duration of each Super Burst (SBD) and the interval separating consecutive Super Bursts (ISBEI) were quantified for each well. These two measures could be defined only in the five wells in which Super Bursts persisted after stimulation. Well B2 (MultiWell 2186) was excluded, because neither quantity exists in the absence of bursting. In every one of these five wells, stimulation shortened the mean Super Burst duration and lengthened the mean inter-Super-Burst-event interval, so that the direction of change was consistent across all cultures that retained bursting (Figure 4.2A, B; Table A.3). The decrease in Super Burst duration was statistically significant (paired t-test, $p = 0.049$, $n = 5$), whereas the increase in the inter-Super-Burst-event interval was not ($p = 0.12$), although it was observed in all five wells.




Figure 4.2. Super Burst duration and inter-Super-Burst-event interval before and after stimulation. (A) Mean Super Burst duration (SBD) and (B) mean inter-Super-Burst-event interval (ISBEI), in seconds, in the PRE and POST conditions, for the five

wells in which Super Bursts persisted after stimulation. Bars: mean across wells; error bars: standard error of the mean (SEM); points: individual wells. Differences between PRE and POST were assessed with two-tailed paired t-tests, each well being one paired observation (ns, $p \geq 0.05$; *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$).

4.2.1. Cluster-level Super Burst analysis

To establish whether a small group of electrodes carries the same description of the phenotype as the whole well, the Super Burst detection pipeline was applied without changes to each selected 3×3 cluster. Clusters were identified separately in the PRE and POST conditions, because cluster selection depends on the spatial distribution of active electrodes, which changes after stimulation. Clusters are therefore not paired observations, but local samples of the same network. Representative examples are shown in Figure 4.3D, E, and the position of the clusters on the array in Figure 4.3A.


The Super Burst features measured in the clusters reproduced those of the whole well, before and after stimulation (Figure 4.3B, C; Table 4.1). In the five wells in which Super Bursts persisted, the mean duration averaged over the clusters differed from the whole-well value by less than 1 s in every condition except two (well B3 after stimulation, 13.1 against 14.6 s; well B1 before stimulation, 50.9 against 45.6 s), and it decreased after stimulation in all five wells, as in the whole well. In well B2 Super Bursts were abolished after stimulation in every cluster, as in the whole well. The inter-Super-Burst-event interval followed the whole-well values in wells A1, B3, A2 and A3; in wells B2 and B1 it was shorter in the clusters than in the whole well.

Within a condition, the clusters of the same well gave very similar distributions (Figure 4.3D, E): in well A1 the per-cluster means of Super Burst duration ranged from 19.8 to 20.7 s before stimulation and from 9.2 to 10.2 s after it. This indicates that the clusters do not sample independent local events but record the same network-wide Super Bursts. The agreement was not equally close in all wells: in well B1 before stimulation the per-cluster means ranged from 17 to 84 s, with the number of detected Super Bursts varying from 14 to 53, because long events were split differently by the cluster-specific NIBI thresholds. For the purpose of this analysis the relevant result is that a group of nine electrodes carries, in most wells and conditions, the same temporal description of the phenotype as the whole well.


Table 4.1. Super Burst features in the whole well and in the 3×3 clusters. Mean Super Burst duration (SBD) and mean inter-Super-Burst-event interval (ISBEI) of the whole well, and the same quantities averaged across the clusters of the well (mean of the cluster means), in the PRE and POST conditions. In well B2 Super Bursts were abolished after stimulation (—).

MultiWell	Well	SBD PRE, whole well (s)	SBD PRE, clusters (s)	SBD POST, whole well (s)	SBD POST, clusters (s)	ISBEI PRE, whole well (s)	ISBEI PRE, clusters (s)	ISBEI POST, whole well (s)	ISBEI POST, clusters (s)
2186	A1	19.9	20.0	9.6	9.6	37.8	38.3	40.3	40.3
2186	B2	5.6	5.2	—	—	27.4	18.9	—	—
2186	B3	19.4	19.4	14.6	13.1	21.3	20.9	47.9	48.7
2013	A2	19.9	19.9	8.1	8.1	9.6	8.4	10.0	8.5
2013	A3	35.4	35.6	22.7	21.9	19.2	19.3	32.3	30.7
2013	B1	45.6	50.9	9.5	9.5	4.6	3.3	9.4	6.9


A Selected 3×3 clusters (well A1, PRE)




B Super Burst duration



C Inter-Super-Burst interval



D Cluster IFR and Super Bursts, PRE



E Cluster IFR and Super Bursts, POST




Figure 4.3. A 3×3 cluster reproduces the Super Burst features of the whole well. (A) Electrode grid of well A1 before stimulation with the ten selected 3×3 clusters (green: active electrodes; red: valid but not active; grey: excluded). (B) Mean Super Burst duration averaged over the clusters of a well against the whole-well value, one point per well and condition; dashed line: identity. (C) The same for the inter-Super-Burst-event interval. (D, E) Instantaneous firing rate of the ten clusters of well A1 before (D) and after (E) stimulation, first 220 s; red dots: detected peaks; shaded: Super Bursts.

4.3. Mutual Information

The Super Burst analysis shows that the clusters share the temporal features of the whole well. Mutual Information (MI) was used to establish whether the activity of each cluster is statistically dependent on that of the rest of the network (cluster-to-network MI) and on that of the other clusters (cluster-to-cluster MI).

The median cluster-to-network MI ranged from 0.58 to 0.88 nats across wells before stimulation and from 0.59 to 0.85 nats after it; the median cluster-to-cluster MI ranged from 0.44 to 0.65 and from 0.41 to 0.60 nats (Table A.4). After stimulation both medians were lower in three wells (A1, B2, A3) and higher in the other three (B3, A2, B1), so no common direction of change was observed (paired t-tests, $p = 0.95$ and $p = 0.87$) (Figure 4.4C, D). An example of the cluster-to-cluster values is shown in Figure 4.4A, B. Because the estimate depends on the properties of the signals, which change with the number and duration of the events, differences between conditions are reported descriptively and are not interpreted as changes in coupling.

4.3.1. Surrogate analysis

The cluster-to-network MI was compared with its distribution under the null hypothesis of independence, obtained from circular time shifts (Methods). With shifts of 0.2 to 2 s, the observed value exceeded every surrogate in all the clusters, 55 of 55 before stimulation and 58 of 58 after it: the mean of the surrogates was between 0.01 and 0.06 nats and their largest value 0.39 nats, against observed values between 0.39 and 0.99 nats, and in every cluster the observed value was at least 1.9 times the largest surrogate. With shifts drawn over the whole recording, which remove the alignment of long events as well, the surrogate values were close to zero (largest value 0.02 nats) and the result

was the same, 55 of 55 and 58 of 58 clusters. The activity of each cluster is therefore statistically dependent on the activity of the rest of the network, before and after stimulation. Together with the cluster-level Super Burst analysis, these results indicate that a 3×3 cluster can serve as a reference population for a network model of reduced size.




Figure 4.4. Mutual information between clusters and with the rest of the network.
(A, B) Normalised cluster-to-cluster mutual information in well A1 before (A) and after

(B) stimulation. (C) Median cluster-to-network mutual information of each well before and after stimulation. (D) The same for the cluster-to-cluster mutual information. (E, F) Representative surrogate analysis for one cluster (C1) of well A1 (MultiWell 2186), before (E) and after (F) stimulation, $k = 10$. Grey: distribution of the cluster-to-network mutual information obtained from 1000 circular-shift surrogates (shifts of 0.2–2 s; Gaussian kernel density estimate on a logarithmic axis); dashed green line: 95th percentile of the surrogate distribution ($\alpha = 0.05$); red line: observed mutual information. (C, D) Differences between PRE and POST were assessed with two-tailed paired t-tests, each well being one paired observation (ns, $p \geq 0.05$; *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$).

4.4. Neural Dynamics

In this section the activity of each well is described as a trajectory in the state space of the population. The principal component analysis and the effective dimensionality are reported first, then the speed of the trajectories, and finally the extent and the structure of the set of visited states.

4.4.1. Principal Component Analysis

The state-space analysis was restricted to the electrodes that were active in both the PRE and POST conditions, a smaller set than that used for the whole-well analysis (756 to 1232 electrodes per well). To verify that this restriction did not distort the population signal, Super Bursts were re-detected from the common-electrode instantaneous firing rate (IFR) using the same adaptive NIBI procedure. In the well shown, the common-electrode IFR retained the same bimodal NIBI distribution and adaptive threshold as the whole-well signal and yielded a number of Super Bursts consistent with the whole-well analysis. The population activity was dominated by the first principal component: PC1 captured on average 26% (PRE) and 28% (POST) of the total variance, whereas PC2 and PC3 each contributed a few per cent or less (Table A.5). The first three components together explained about 30% of the total variance. This does not describe a signal spread over three axes: PC1 is far larger than every other component, and the remaining variance is distributed over a long tail of minor directions. The variance structure was similar in the two conditions.

Because the first three components explain only about 30% of the variance, the number of components that carry population structure was estimated with Parallel Analysis against circular-shift surrogates (Methods). For each well, the spectrum of the recording was compared with the mean spectrum of 100 surrogates in which each electrode was circularly shifted by an independent random offset, which keeps the temporal statistics of each electrode and removes the correlations between electrodes. PC1 exceeded the null level by more than two orders of magnitude and PC2 by about one order of magnitude; from PC3 onward the spectrum of the recording approached the null band (Figure 4.5A). Parallel Analysis counted more components above the 95th percentile of the surrogates (7 to 28 before and 9 to 18 after stimulation, depending on the well): with about 60,000 time bins per recording even a small excess over the null is significant, but beyond PC2

the excess is small, as the effective dimensionality shows. After subtracting the null spectrum and renormalising, the effective dimensionality of the signal was 1.31 ± 0.13 before stimulation and 1.27 ± 0.08 after it (mean $\pm$ SD, $n = 6$; Table 4.2). The signal is therefore dominated by a single mode, with PC2 as a minor second axis. The quantitative trajectory analysis was carried out in the PC1–PC2 plane, which is the smallest space in which the geometry and the speed of a trajectory can be described.

The effective dimensionality was slightly lower after stimulation in five of the six wells, by 0.01 to 0.11; the difference was not statistically significant (paired t-test, $p = 0.079$) and is small compared with the variability between wells.


Table 4.2. Effective dimensionality of the population signal. Effective dimensionality after subtraction of the null spectrum (circular shifts, 100 surrogates).

MultiWell	Well	Electrodes	Effective dimensionality PRE	Effective dimensionality POST
2186	A1	1232	1.15	1.17
2186	B2	756	1.44	1.33
2186	B3	919	1.29	1.28
2013	A2	965	1.31	1.25
2013	A3	1173	1.22	1.20
2013	B1	1025	1.48	1.39
	Mean $\pm$ SD		1.31 ± 0.13	1.27 ± 0.08


Because the trajectories are expressed in a single PRE-derived coordinate system, the comparison is meaningful only if the dominant modes are preserved across conditions. Two checks indicate that they are. First, the magnitudes of the PC1 loadings were strongly correlated between conditions in every well (Spearman's $\rho = 0.94 \pm 0.03$; range 0.88–0.96; Figure 4.5C), showing that the same electrodes drive the dominant mode before and after stimulation. Second, the principal angles between the PRE and POST subspaces were small for the leading components (PC1, $9.1^\circ \pm 0.8^\circ$; PC2, $17.9^\circ \pm 1.7^\circ$; mean $\pm$ SEM) and larger and more variable for the third (PC3, $49.5^\circ \pm 11.5^\circ$, range 27–89°; Figure 4.5B). The PC1–PC2 plane is therefore closely preserved, whereas the poor preservation of PC3 is a further reason to restrict the quantitative analysis to PC1–PC2. This check used all six wells, including well B2, whose continuous trajectory exists although its Super Bursts are abolished after stimulation. The comparisons that follow use the five wells in which Super Bursts persisted.

Projected onto the PRE-derived components, each event becomes a path of the population state through the reduced space. At the Super Burst scale, every event is an excursion that leaves the baseline region near the origin, extends along PC1 and returns, so that the events form a compact bundle (Figure 4.6A). At the Network Burst scale, each event traces a closed loop, leaving the baseline along PC1 and returning by a slightly different path, so that together the events outline a ring-shaped region with an empty centre (Figure 4.6B). Both geometries are present before and after stimulation. These three-dimensional overlays are shown for illustration only; the quantitative measures were computed in the PC1–PC2 plane.

A Spectrum of the recording and null spectrum (well A1)



B Preservation of the subspace



C PC1 loadings, PRE against POST (A1)






Figure 4.5. Dimensionality and preservation of the population subspace. (A) Variance explained by the principal components of the recording and null spectrum from

circular shifts (95th percentile), well A1, before and after stimulation. (B) Principal angles between the PRE and POST subspaces (bars: mean; points: wells). (C) Magnitude of the PC1 loading of each electrode before and after stimulation (well A1).

A Super Burst trajectories (well A1): PRE blue, POST red



B Network Burst trajectories (well A1): PRE blue, POST red





Figure 4.6. Population trajectories in the reduced state space (well A1). (A) Super Burst trajectories and (B) Network Burst trajectories, projected on the first three components of the PRE-derived basis; blue: before stimulation; red: after stimulation. Illustration only: the measures are computed in the PC1–PC2 plane.

4.4.1.1. Trajectory speed and state-space extent

Each event traces a path of the population state in the PC1–PC2 plane, and the speed at which the state moves along this path is the magnitude of the time derivative of the trajectory (Methods).

A single fragment is an out-and-back excursion of the population state. Starting from the baseline region near the origin, the state moves along PC1 to the apex of its displacement, where it reverses, and then returns towards baseline. Such a path produces a two-peaked speed profile: the speed is high during the outward movement, falls to a minimum at the apex, where the state reverses, and rises again during the return (Figure 4.7A, B). A Super Burst is a train of such fragments, so at the Super Burst scale this two-peaked profile repeats once for every fragment of the envelope (Figure 4.7D, E).

The two legs of the excursion correspond to the rise and the fall of the population activity. The global instantaneous firing rate, summed across the common electrodes and computed directly from the spike trains, increased during the outward leg, reached its maximum at the apex of the excursion and decreased during the return leg, in all five wells and both conditions (Figure 4.7C). This correspondence is expected and is not an independent result: the first principal component follows the overall amplitude of the population activity, so the speed along it is close to the rate of change of the population firing rate. The trajectory speed is therefore a description, in the state space, of how fast the population is recruited and how fast its activity terminates: its two maxima mark the phases of fastest change, and its minimum the moment of maximal activity.

The peak trajectory speed (v_{peak}) and the scale-normalised speed ($\tilde{v}$) were quantified in the five wells in which Super Bursts persisted after stimulation; well B2 (MultiWell 2186) was excluded. After stimulation the peak speed was higher at the Super Burst scale in all five wells (by 3 to 19%; paired t-test, $p = 0.012$; Figure 4.7F, G; Table A.6) and at the fragment scale in four of the five (well A2 decreased; $p = 0.13$). The scale-normalised speed changed little: at the Super Burst scale it increased in three wells and decreased in two ($p = 0.72$), and at the fragment scale it decreased in four wells by 1 to 8% ($p = 0.24$) (Figure 4.7H, I; Table A.7).

These values must be read with one limitation of the analysis in mind. The activity of each electrode was standardised within each condition, with its own mean and standard deviation, before projection. The peak speed is therefore expressed in units that depend

on the variability of the activity in that condition: if the events occupy a smaller fraction of the recording after stimulation, as the shorter Super Bursts imply, the same event becomes larger in standardised units. The increase in peak speed cannot therefore be attributed to a larger excursion of the population in absolute terms. The scale-normalised speed, which is divided by the size of the excursion of the same event, is not affected by this limitation. When the analysis was repeated with a common standardisation (PRE mean and standard deviation applied to both conditions), the peak speed of the Super Bursts was again higher after stimulation in all five wells (mean from 597 to 672 a.u./s), but the difference was not statistically significant (paired t-test, $p = 0.076$; Table 4.3).

The speed profile also resolves the fragmented structure of the events: each fragment produces one out-and-back excursion, so the temporal spacing of the fragments is the interval between successive fragment peaks. Each Super Burst is composed of P peaks occurring at times $t_1 < t_2 < \dots < t_P$, the same peaks used to identify the Super Burst. The inter-fragment interval is the time between two consecutive peaks,

$$IFI_j = t_{j+1} - t_j, \quad j = 1, \dots, P - 1,$$

and each event was summarised by the median of its inter-fragment intervals; its reciprocal is the inter-fragment frequency. After stimulation the inter-fragment interval was longer in three wells (B3, A2 and B1, by 7 to 31%) and unchanged in two (A1 and A3, change of less than 2%; paired t-test, $p = 0.10$; Figure 4.7J; Table A.8). Where it changed, the fragments recurred less frequently within the Super Burst.




Figure 4.7. Trajectory speed. (A) One fragment as a trajectory in the PC1–PC2 plane lifted by its speed. (B) Speed of the same fragment against time. (C) Trajectory speed (blue) and population firing rate (orange) of one fragment; dashed line: apex of the excursion. (D, E) Speed during one Super Burst before (D) and after (E) stimulation; black dots: peaks; note the different time axes. (F, G) Peak speed at the fragment and Super Burst scales, one point per well. (H, I) Scale-normalised speed. (J) Inter-fragment interval. a.u.: units standardised within each condition. (F–J) Differences between PRE and POST were assessed with two-tailed paired t-tests, each well being one paired observation (ns, $p \geq 0.05$; *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$). For the peak speed of the Super Bursts (G), the symbol refers to the control with a common standardisation ($p = 0.076$).

4.4.1.2. Entropy and volume

Beyond the timing of the events, the analysis asked whether the set of states visited by the network differed after stimulation. At every instant the activity of the well is summarised by its position in the PC1–PC2 plane, so that the instants belonging to the events form a cloud of points (Figure 4.8A). Two properties of this cloud were measured:

its extent, with the generalised variance V (the area of the covariance ellipse), and its structure, with the negentropy ΔH (how far the cloud departs from a Gaussian cloud with the same covariance). The differential entropy H , in nats (natural logarithm), is reported as well, but it is not a third independent quantity: by definition $H = \log V + \text{constant} - \Delta H$, so the entropy increases when the cloud becomes larger or less structured. For the network bursts, for example, the increase in entropy (+0.27 nats) is the sum of the increase in log-volume (+0.12) and of the decrease in negentropy (0.14). The analysis was restricted to the five wells in which Super Bursts persisted, with the well as the unit of observation.

The generalised variance was higher after stimulation in all five wells for both event types: for the network bursts by 3 to 23% (mean from 123 to 139, in squared standardised units; paired t-test, $p = 0.013$) and for the Super Bursts by 5 to 20% (mean from 59 to 66; $p = 0.032$; Figure 4.8B, E; Table A.9).

The negentropy of the network bursts was lower after stimulation in four of the five wells, by 0.10 to 0.38 nats, and unchanged in well A3 (+0.02 nats); its mean went from 1.85 to 1.71 nats (paired t-test, $p = 0.095$; Figure 4.8C, F; Table A.10). The cloud of the network bursts therefore tended to be less structured, closer to a Gaussian cloud of the same size, although the change was not significant. For the Super Bursts the negentropy decreased in three wells and increased in two, with an unchanged mean (2.04 and 2.09 nats; $p = 0.61$).

The entropy follows from these two changes. For the network bursts it was higher after stimulation in all five wells, by 0.05 to 0.42 nats (mean +0.27 nats; paired t-test, $p = 0.011$; Figure 4.8D, G; Table A.11), the smallest change being in well A3. For the Super Bursts it increased in three wells and decreased in two (mean +0.05 nats; $p = 0.57$).

Table 4.3. Summary of the state-space measures before and after stimulation. Mean $\pm$ SEM across the five wells in which Super Bursts persisted (well B2 excluded), and p values of the two-tailed paired t-test between PRE and POST, with the standardisation applied to each condition separately (as for the values in the table) and, as a control, with a common standardisation (PRE mean and standard deviation). †: Wilcoxon matched-pairs signed-rank test (differences not normally distributed). –: not recomputed (the measure does not depend on the scale). The values of each well are in Tables A.6–A.11. a.u.: units standardised within each condition.

Measure	Events	PRE	POST	p
Peak speed (a.u./s)	Fragments	831.9 $\pm$ 104.8	875.7 $\pm$ 113.0	0.13
Peak speed (a.u./s)	Super Bursts	597.2 $\pm$ 51.8	666.8 $\pm$ 60.0	0.012
Scale-normalised speed (1/s)	Fragments	40.66 $\pm$ 2.02	39.75 $\pm$ 2.27	0.24
Scale-normalised speed (1/s)	Super Bursts	39.35 $\pm$ 1.07	39.67 $\pm$ 1.51	0.72
Inter-fragment interval (s)	Super Bursts	0.397 $\pm$ 0.016	0.444 $\pm$ 0.024	0.10
Generalised variance (a.u. ²)	Network bursts	123.4 $\pm$ 11.0	139.3 $\pm$ 11.6	0.013
Generalised variance (a.u. ²)	Super Bursts	59.4 $\pm$ 4.4	66.0 $\pm$ 5.8	0.032
Negentropy (nats)	Network bursts	1.854 $\pm$ 0.150	1.711 $\pm$ 0.185	0.095
Negentropy (nats)	Super Bursts	2.041 $\pm$ 0.134	2.090 $\pm$ 0.130	0.61
Differential entropy (nats)	Network bursts	5.782 $\pm$ 0.069	6.050 $\pm$ 0.110	0.011
Differential entropy (nats)	Super Bursts	4.871 $\pm$ 0.145	4.923 $\pm$ 0.176	0.57

In summary (Table 4.3), in the state space the events kept the same geometry after stimulation, and the scale-normalised speed and the effective dimensionality changed little. The cloud of states visited during network bursts was larger after stimulation and tended to be less structured: its generalised variance and its entropy increased in all five wells ($p = 0.013$ and $p = 0.011$), while its negentropy decreased in four of the five ($p = 0.095$). At the Super Burst scale the peak speed and the generalised variance also increased in all five wells ($p = 0.012$ and $p = 0.032$).




Figure 4.8. Repertoire of visited states. (A) States visited during network bursts in the PC1–PC2 plane (well A1) with the covariance ellipse, before (blue) and after (red) stimulation. (B–D) Network bursts: generalised variance, negentropy and differential entropy of each well before and after stimulation. (E–G) The same for Super Bursts. (B–G) Differences between PRE and POST were assessed with two-tailed paired t-tests, each well being one paired observation (ns, $p \geq 0.05$; *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$). For the generalised variance (B, E) and the entropy of the network bursts (D), the symbols refer to the control with a common standardisation ($p = 0.13, 0.11$ and 0.060).

4.5. *In silico* model

This section first reports the checks of the single-neuron model: morphological reduction, extracellular field, resting state, intrinsic properties, axon and single-pulse approximation. It then reports the response of the simulated population: direct activation, subthreshold polarisation, their spatial structure, and their dependence on slab thickness, orientation and morphology. It ends with the comparison with the preliminary campaign and with one sample of 2000 neurons taken as an example.

4.5.1. Morphological reduction and numerical resolution

The six reconstructed human L2/3 pyramidal neurons were reduced from their original three-dimensional geometry to a quasi-planar representation compatible with the geometry of the MEA culture (Figure 4.9A). Three slab thicknesses were considered: 40, 80 and 120 μm . For each morphology, a dendritic section was retained when at least 80% of its arc length lay within the slab and its parent section was also retained, preserving connectivity with the soma (Figure 4.9B).


Increasing the slab thickness progressively increased the retained dendritic arbour while preserving the sections already present at smaller thicknesses. The three reductions therefore formed nested representations of the same reconstructed neuron, providing controlled levels of dendritic complexity rather than independent morphological realisations. This allowed the effect of the uncertain out-of-plane extent of the cultured neurons to be evaluated without changing the underlying morphology. The reduction acts only along the vertical axis, so the retained arbour can extend far from the soma in the plane: for specimen 60308, for example, the retained dendrites reach up to 267 μm from the soma at 40 μm and up to 774 μm at 80 and 120 μm .

The reduced morphologies were discretised into isopotential compartments according to the d_λ rule (Figure 4.9C–D).


A The six reconstructed human L2/3 pyramidal neurons



B Reduction to slabs of 40, 80 and 120 μm (top: 3D; bottom: plane of the array)



C d_λ rule on one branch



D All dendritic compartments





Figure 4.9. From the reconstructions to the simulated cells. (A) The six reconstructed human L2/3 pyramidal neurons. (B) Reduction of one neuron (specimen 60308) to slabs of 40, 80 and 120 μm centred on the soma: three-dimensional view (top) and view in the

plane of the array (bottom); orange: apical dendrite; green: basal dendrites. (C) The d_λ rule on one dendritic branch: physical length, length constant at 100 Hz and compartments. (D) Compartment length relative to the length constant for all the dendritic compartments of one reduced cell; dashed line: the 0.1 criterion.


4.5.1.1. Extracellular stimulation field

The extracellular potential was computed with a quasi-static point-source formulation in a homogeneous conductive half-space, summing the contributions of the six electrodes (Figure 4.10A, B). The implementation reproduced the expected inverse-distance dependence for a single source, $V_e \propto 1/r$, and the potential of the full array coincided with the sum of the six individual contributions, confirming linear superposition. The biphasic stimulus was charge-balanced and the field vanished when the current was removed (Figure 4.10C). These analyses establish the internal consistency of the implemented field model and numerical solver.

A Extracellular potential at the height of the somata, first and second phase



B Polarity of the two phases



C Current waveform




Figure 4.10. Stimulus. (A) Extracellular potential at the height of the somas at the peak of the first (left) and of the second phase (right); red: anodes, blue: cathodes; dashed square: region in which somas were placed in the preliminary tests. (B) Polarity of the two phases. (C) Current waveform with 100 μ s ramps and its time derivative.

4.5.1.2. Resting state and leak tuning

The resting membrane potential was not imposed as a fixed value. Each morphology–layer combination was allowed to settle without stimulation: the somatic potential relaxed quickly from the initial value and then slowly, with the kinetics of the human I_h , whose density increases along the apical dendrite (Figure 4.11B), towards about -74 mV (Figure 4.12A). In the Full Tuned model used for the final campaign, the leak reversal


potential was then set compartment by compartment so that this rest held uniformly across the arbour.

To check that the unstimulated cell stays at rest, every morphology and slab thickness was also simulated with zero current. In the untuned Soma Only model the drift at t_0 was +0.081 mV on average (0.054 to 0.108 mV, one value per morphology and thickness); in the Full Tuned campaign it was below 10^{-4} mV (mean -3.7×10^{-5} mV, largest absolute value 6.5×10^{-5} mV). In both cases it is far below the 1 mV classification threshold. In Full Tuned the response can therefore be measured from the resting potential; in Soma Only the zero-current run was subtracted. The somatic resting potential was -73.8 mV on average (-74.8 to -73.1 mV) in Full Tuned and -83.4 mV (-83.7 to -82.6 mV) in Soma Only (Figure 4.12B).

4.5.2. Intrinsic properties of the neuron model

The membrane model combines the human conductances [S. Rich, 2021], with their kinetics and spatial distributions, and the Eyal morphologies [G. Eyal, 2016; G. Eyal, 2018] (Figure 4.11A). The implemented conductances and their distributions were compared with the source model. The intrinsic behaviour was then assessed independently of extracellular stimulation by somatic current injection. The response depended on the morphology: specimen 60303 fired regularly at all the currents tested, whereas other morphologies entered depolarisation block. Specimen 60308, for example, fired at the onset of the current step but sustained firing only at 0.2 nA, and its potential settled on a depolarised plateau from 0.1 nA. Perturbation analyses tested whether the block arose from the channel set, the AIS, the morphological reduction or the activation shifts: none of these removed it, and the block was therefore interpreted as morphology-dependent rather than as a defect of the channel set. These tests were run on the Full Active model. The Full Tuned model, with its passive axon, also fired repetitively under somatic current injection (4, 16 and 26 spikes in 500 ms at 0.1, 0.2 and 0.4 nA).

A Where the channels are (specimen 60308, 80 μm slab; grey: absent)



B Rules that depend on the distance from the soma: Ca hot spot (left), I_h gradient (right)




Figure 4.11. Placement of the ionic channels. (A) Maximal conductance of each channel on specimen 60308 (80 μm slab); grey: channel absent. (B) The two rules that depend on the path distance from the soma: calcium hot spot between 360 and 600 μm (left) and exponential gradient of the h-current along the apical dendrite (right).

4.5.2.1. Channel dynamics

The ionic conductances were followed before, during and after stimulation. At rest, the potential is set by the passive and ionic currents, with I_h partly active and the persistent sodium current contributing to the balance. During a suprathreshold response, the fast sodium conductance drives the upstroke and the potassium currents the repolarisation; calcium conductances open around the spike peak and activate SK_{E2} . The key observation was that I_h remained essentially unchanged during the 0.5 ms pulse: its slow kinetics prevent it from being recruited during the stimulus. After the pulse, the fast conductances recover within milliseconds, whereas calcium handling and I_h recover more slowly.

4.5.3. Axon initial segment and stylised axon

In cortical pyramidal neurons, action potentials evoked by synaptic or somatic input initiate in the axon initial segment (AIS) [G. Stuart, 1997; L.M. Palmer, 2006; M.H.P. Kole, 2012], and extracellular stimulation preferentially activates axons rather than somas [A.S. Aberra, 2018; L.G. Nowak, 1998; C.C. McIntyre, 1999], with the lowest thresholds near the AIS in cultured neurons on HD-MEAs [M. Radivojevic, 2016]. The Rich model leaves the axon passive, which is adequate for somatic current injection. The Full Active variant therefore added sodium and potassium channels to the AIS of the stylised axon. In that model, under sustained somatic current injection, setting the sodium and potassium conductances of the AIS to zero did not remove the depolarisation block, so the AIS was not its origin.

The contribution of the axon to the campaign outcome was measured directly on the six morphologies (80 μm slab), each at the same 30 random placements and orientations within 120 μm of the dipole centre (180 paired simulations per model, 50 μA). The same placements were simulated with the Full Tuned and the Full Active models, recording the membrane potential and the sodium current of every compartment. The site of initiation was defined as the region (soma, apical dendrite, basal dendrite, axon) whose inward sodium current first reached 10% of the peak total sodium current of the cell.

The two models gave the same fire/no-fire outcome in 178 of the 180 paired simulations (43 firing neurons in Full Tuned, 45 in Full Active). The two disagreements were at the same placement, 6 μm from the dipole centre on the midline between the electrode

columns, in two morphologies (60308 and 130305): the neuron fired only in Full Active. In those two cells Full Active also starts from a slightly less negative potential (-73.9 mV against -74.0 and -74.3 mV), so these borderline cases cannot be attributed to the initial segment alone. In both models, the spike initiated at the soma in every firing neuron and in every morphology (Table 4.4; Figure 4.12C). In Full Tuned the soma carried 99.3% of the sodium charge of the spike; in Full Active the axon carried 25% (18–32% across neurons), but its current never preceded the somatic one.

Table 4.4. Site of initiation and effect of the axon (six morphologies $\times$ 30 paired placements, 80 μ m slab; axon-length test: one morphology).

	Full Tuned (passive axon)	Full Active (active AIS)
Firing neurons	43 / 180	45 / 180 (the same 43, plus 2)
Initiation at the soma	43 / 43	45 / 45
Sodium charge: soma	99.3%	74.5%
Sodium charge: axon	0%	25.0%
Sodium charge: apical dendrite	0.7%	0.5%
Soma crosses 0 mV in phase 1 / phase 2 / after the pulse	20 / 18 / 5	20 / 18 / 7

Axon length (Full Tuned, 150 placements within 200 μ m)	6 mm	600 μ m	60 μ m
Somatic rest	-74.28 mV	-74.22 mV	-74.01 mV
Same fire/no-fire outcome as 6 mm	–	150 / 150	150 / 150
Same subthreshold class as 6 mm (non-firing)	–	124 / 124	121 / 124
Largest change in ΔV_{end} (non-firing)	–	0.24 mV	0.49 mV

The stylised axon is retained in all model variants, even where it is passive: it loads the soma, it is polarised by the field like any other neurite, and it keeps the geometry identical across variants, so that they differ only in the placement of the voltage-gated channels.

The stylised axon was also tested for sensitivity to its length: shortening it from 6 mm to 600 or 60 μ m left the fire/no-fire outcome unchanged in all placements and changed the subthreshold response by less than 0.5 mV (Table 4.4; Figure 4.12E). In this model the outcome is therefore decided by the soma and the dendrites, and the stylised axon contributes little to it.

In the final Full Tuned campaign the axon is passive, as in the Rich model, and spikes initiate at the soma, where the model concentrates the excitability. Using the passive

axon instead of the active initial segment of Full Active changes the outcome in about 1% of the neurons near the electrodes (2 of 180).

4.5.4. Single-pulse approximation

The experimental train delivers biphasic pulses of 0.5 ms at 0.2 Hz, i.e. every 5 s. The channel dynamics showed a clear separation between the pulse and the slow recovery processes: I_h is essentially unchanged during the pulse, and the slow calcium-dependent and I_h -dependent recovery is complete well within the 5 s interval. In a simulated train of three pulses 800 ms apart, an interval much shorter than the experimental 5 s, the three somatic responses overlapped, with a maximum difference of 2.4 mV between pulses. Each pulse was therefore treated as an independent perturbation, and the final campaign simulated one pulse per neuron rather than the full 36-pulse train.




Figure 4.12. Checks of the single-neuron model. (A) Settling of the somatic potential and of the net membrane current without stimulation (specimen 60308, model with active initial segment). (B) Change of the somatic potential of the unstimulated cell at the end of the pulse window in the Soma Only and Full Tuned campaigns; each bar is one morphology and slab thickness. (C) Share of the sodium charge of the spike carried by soma, dendrites and axon with a passive axon and with an active initial segment (six morphologies); the soma fired first in all firing neurons. (D) Number of firing neurons (of 150) when the discretisation of dendrites and axon, or of the soma, is refined (one morphology). (E) Probability of firing against distance (left) and against amplitude

(right) for three lengths of the stylised axon (6 mm, 1000 μm and 500 μm ; the three curves coincide; specimen 60308, model with active initial segment; the test at 600 and 60 μm is in Table 4.4).

4.5.5. Population response to stimulation

The final campaign simulated 432 samples of 2000 independent neurons, each simulated at the three slab thicknesses: 864,000 neurons and 2,592,000 single-cell simulations, with the six morphologies in equal proportions. The somas were placed uniformly in a $600 \times 600 \mu\text{m}$ square centred on the stimulating dipole. The spatial sampling was uniform: each $10 \times 10 \mu\text{m}$ bin contained about 720 observations (range about 560–920), and in the polar bins ($5 \mu\text{m} \times 5^\circ$) the number of observations grew with distance, from a few tens near the centre to about 2500 at the corners of the square. No duplicated observations were found.

Over the simulated region, 4.68% of the observations were activations (95% Wilson interval 4.66–4.71%), 7.22% depolarisations (7.19–7.25%) and 6.36% hyperpolarisations (6.33–6.39%); the remaining 81.74% showed no effect above threshold.

4.5.5.1. Direct activation

The probability of direct activation (Figure 4.13B, E) reached 0.83 (95% Wilson interval 0.82–0.84) at 45–50 μm from the dipole centre, fell to 0.5 at 77 μm , to 0.1 at 99 μm and to 0.01 at 120 μm ; no neuron fired beyond 140 μm . Right at the dipole centre it was lower, about 0.5, because the centre lies on the midline between the electrode columns. In the plane, activation followed the electrode layout: two regions over the anode and cathode columns, larger and more probable over the cathodes, with a dip on the midline between them (Figure 4.13F). Within 150 μm of the centre, 28.6% of the neurons on the cathode side fired, against 19.4% on the anode side.

4.5.5.2. Subthreshold polarisation

Among the neurons that did not fire, the median stimulus-induced change at the soma (Figure 4.13G) was several millivolts within about 100 μm , fell below 1 mV at about 150 μm and to about 0.01 mV at 400 μm . Far from the electrodes the change is therefore smaller but still has a sign: a classification by sign alone assigned about half of these neurons to each subthreshold class, giving a uniform probability of about 0.5 far from the

array (Figure 4.14F). With the 1 mV threshold, a neutral class appears and the subthreshold probabilities fall below 0.01 at 209 μm and vanish beyond 220 μm .

With this threshold, the probability of depolarisation (Figure 4.13C, E) peaked at 0.45 at about 100 μm and that of hyperpolarisation at 0.42 at about 115 μm . Subthreshold polarisation therefore extends about twice as far as direct activation, and over the simulated region about three times more neurons were polarised below threshold than activated (13.58% against 4.68% of the observations, a ratio of 2.9).

The extent of the subthreshold response depends only weakly on the threshold: halving or doubling it moved the distance at which the subthreshold probability falls below 0.01 by about 20 μm (Table A.12; Figure 4.13H).




Figure 4.13. Response of the simulated population to one stimulation pulse (final campaign, 2,592,000 simulations). (A) Extracellular potential at the height of the somas during the second phase of the pulse (anodes negative, cathodes positive); squares mark

the six electrodes. (B–D) Probability that a neuron is activated, depolarised or hyperpolarised as a function of the position of its soma ($10 \times 10 \mu\text{m}$ bins). (E) Probability of the three outcomes against the distance from the dipole centre, with 95% Wilson intervals. (F) Profile of the three probabilities along x , for somas within $90 \mu\text{m}$ of the array axis; A and C mark the anode and cathode columns. (G) Absolute change in somatic potential at the end of the pulse in non-firing neurons (median, 25–75% and 5–95%); dashed line: the 1 mV threshold. (H) Probability of a subthreshold outcome with thresholds of 0.5, 1 and 2 mV.

4.5.5.3. *Spatial structure: anisotropy and symmetry*

Depolarisation prevailed on the cathode side and hyperpolarisation on the anode side (Figure 4.13C, D; Figure 4.14A–D): along the cathode-to-anode axis, the probability of depolarisation increased towards the cathodes and that of hyperpolarisation towards the anodes. On the outer flanks the two maps were mirror images, with the same peak of about 0.385 at $x = +85 \mu\text{m}$ for depolarisation and $x = -85 \mu\text{m}$ for hyperpolarisation, as expected from a field that is negative on one side and positive on the other (Figure 4.13A). Close to the electrodes the symmetry was broken: around the anodes, which are negative at the end of the pulse, a ring of depolarised neurons appeared (peaks of about 0.245 and 0.22 along x), whereas around the cathodes the corresponding hyperpolarised ring was weaker (about 0.17 and 0.10), largely replaced by the wider activated region. Because firing neurons are removed from the subthreshold classes, the three probabilities and the probability of no effect sum to one in each bin.

The asymmetry was quantified by comparing the map of depolarisation with the mirror image of the map of hyperpolarisation (Figure 4.15D–F). The mean absolute difference was 0.105 within $100 \mu\text{m}$ of the dipole centre and 0.008 beyond it: the asymmetry is confined to the region where neurons fire.

The counts on the two sides explain most of it (Table 4.5; Figure 4.15G). On the outer flanks the two sides mirror each other (43.4% depolarised on the cathode side, 43.4% hyperpolarised on the anode side). Near the electrodes, the neurons on the cathode side that fire (28.6%) or end hyperpolarised (11.0%) together amount to 39.6%; on the anode side, those that fire (19.4%) or end depolarised (19.8%) amount to 39.2%. The strongly driven population is thus the same on the two sides, but a larger share of it fires on the

cathode side, where the first phase is depolarising, and these neurons are counted as activated rather than as hyperpolarised.

Table 4.5. Outcomes within 150 μm of the dipole centre, by side.

	Activated	Depolarised	Hyperpolarised
Full Tuned, anode side ($x < 0$)	19.4%	19.8%	43.4%
Full Tuned, cathode side ($x > 0$)	28.6%	43.4%	11.0%
Soma Only, anode side ($x < 0$)	14.1%	24.5%	44.6%
Soma Only, cathode side ($x > 0$)	20.6%	44.9%	17.6%




Figure 4.14. Dependence on the side of the array and on the orientation of the neuron. (A–C) Probability of activation, depolarisation and hyperpolarisation as a function of the distance from the dipole centre and of the position angle θ (0° : anode side; 180° : cathode side). (D) Probability of the three outcomes against the position angle, for somas within 150 μm of the dipole centre. (E) The same against the orientation of the neuron: no dependence. (F) Probability of depolarisation and hyperpolarisation

classified by sign only (dashed) and with the 1 mV threshold (solid): without a threshold, about half of the distant neurons fall in each class.

4.5.5.4. Origin of the asymmetry near the electrodes

The two columns of electrodes deliver opposite pulses: the anodes inject +50 μ A in the first phase and -50 μ A in the second, and the cathodes the reverse. A soma close to a cathode is therefore depolarised in the first phase and hyperpolarised in the second, whereas a soma close to an anode is hyperpolarised first and depolarised second. Since the outcome is read at the end of the pulse, the sign of a subthreshold response is the one left by the second phase.

This order explains the two rings. In the final campaign, among the neurons with the soma 20–35 μ m from a cathode, on the outer side of the column, about 86% fired, 8% ended hyperpolarised and 4% depolarised: most of them reach threshold during the first phase and are counted as activated, so few remain to form a hyperpolarised ring. At the same distance from an anode, about 42% fired and 54% ended depolarised: here the depolarisation arrives in the second phase and starts from a hyperpolarised soma, so about half of the neurons stay below threshold and form the depolarised ring. Further out, at 50–100 μ m, fewer than 1% of the neurons fired and the two sides were mirror images, with 67% of the neurons hyperpolarised on the anode side and 67% depolarised on the cathode side.

The role of the phase order was tested on one morphology (Methods). With the campaign pulse, 17.9% of the neurons on the anode side and 27.5% on the cathode side fired, 19.9% and 46.2% were depolarised, and 47.1% and 10.3% hyperpolarised, in agreement with Table 4.5. With the two phases in the opposite order the two sides were exchanged: 25.5% and 17.3% fired, 46.6% and 19.3% were depolarised, and 13.7% and 47.7% hyperpolarised; the largest difference from an exact exchange was 3.4 percentage points. The same exchange was obtained with the Soma Only model (2.9 percentage points).

The two models differ in the size of the effect. Around the cathodes, about 57% of the Soma Only neurons fired and 36% ended hyperpolarised, against 86% and 8% in Full Tuned; around the anodes, 18% fired and 77% ended depolarised, against 42% and 54%. The hyperpolarised ring around the cathodes is thus visible in the Soma Only map (Figure 4.16C) and almost absent in the Full Tuned one (Figure 4.13D). The difference follows from the resting potential: Soma Only neurons rest about 10 mV further from

threshold (Table 4.6), so fewer of them fire and more remain in the subthreshold classes. When the resting potential of the Soma Only model was shifted to -74.3 mV on the test morphology, 71.4% of the neurons around the cathodes fired and 13.0% ended hyperpolarised, close to Full Tuned on the same morphology (69.6% and 15.2%) and far from the unshifted Soma Only model (33.8% and 49.4%). Beyond $50\text{ }\mu\text{m}$ the two models gave the same probabilities.

The opposite sign of the ring and of the outer band on the same side follows from the way the field polarises an extended neuron. Below threshold the response is passive: the field drives axial currents along the neurites, towards the regions where the extracellular potential is lower, so that the part of the neuron closest to a negative electrode is depolarised and the rest is hyperpolarised [F. Rattay, 1986; F. Rattay, 1999]. In the test morphology, the change in dendritic membrane potential had the sign imposed directly by the electrode within about $30\text{ }\mu\text{m}$ of it and the opposite sign between 40 and $100\text{ }\mu\text{m}$, and the two signs were exchanged between the first and the second phase. When the soma is the part of the neuron closest to the electrode, axial current converges on it: in the non-firing Full Tuned neurons $20\text{--}35\text{ }\mu\text{m}$ from a cathode, the median somatic change was $+12$ mV at the end of the first phase and $+0.5$ mV at the end of the pulse ($+18$ and -7 mV in Soma Only), and $20\text{--}35\text{ }\mu\text{m}$ from an anode -23 mV and $+11$ mV. When the soma is $60\text{--}80\text{ }\mu\text{m}$ away, the dendrites that reach towards the electrode are the closest part and the soma takes the opposite sign: -2.6 mV and then $+5.7$ mV on the cathode side, $+1.7$ mV and then -4.2 mV on the anode side. The neurons that fired near the cathodes crossed 0 mV about 0.275 ms after the onset of the pulse, just after the end of the first phase, whereas those near the anodes fired after the end of the pulse, at about 0.575 ms.

4.5.5.5. Slab thickness, orientation and morphology

The slab thickness had a small effect. The 80 and $120\text{ }\mu\text{m}$ slabs were indistinguishable, whereas the $40\text{ }\mu\text{m}$ slab gave a slightly lower activation (peak 0.80 against 0.84 ; Table A.13; Figure 4.15A). Neuron by neuron, the outcome was the same at the three thicknesses for 97.9% of the 864,000 neurons, and the fire/no-fire outcome for 99.6%. This held although the lateral reach of the retained arbour changed several-fold between 40 and $80\text{ }\mu\text{m}$. The probabilities did not depend on the orientation of the neuron: within $150\text{ }\mu\text{m}$ of the dipole centre they were the same at every orientation (Figure 4.14E).

The morphology changed the peak probability (from 0.75 to 0.89) and the distance at which the probability falls to 0.5 (from 67 to 83 μm), whereas the outer limit of activation was nearly the same for all six morphologies (0.1 between 96 and 105 μm ; Table A.14; Figure 4.15B). The two most excitable specimens are also those with the least negative resting potential (-73.2 and -73.4 mV); whether this accounts for the difference was not tested. In this model, the response is therefore determined mainly by the position of the soma relative to the electrodes and modulated by the morphology.




Figure 4.15. Robustness of the response and symmetry of the maps. (A) Probability of activation against distance for the three slab thicknesses. (B) The same for the six morphologies. (C) Probability of the three outcomes in the preliminary Soma Only model (dashed) and in the final Full Tuned model (solid); colours as in Figure 4.13. (D) Probability of depolarisation at (x, y). (E) Probability of hyperpolarisation at the mirror

position $(-x, y)$. (F) Difference between D and E: zero on the outer flanks, positive around the anode column. (G) Outcomes of the neurons within 150 μm of the dipole centre, on the anode and on the cathode side.

4.5.6. Comparison with the preliminary campaign

The preliminary campaign used the Soma Only model (720 samples of 1700 neurons on 1 mm^2 , 3,672,000 simulations). The two campaigns were analysed with the same bins and the same 1 mV threshold (Table 4.6; Figure 4.15C; Figure 4.16). Because the simulated areas differ, the overall fractions are not comparable; the curves against distance and the ratio of subthreshold to activated neurons are.

Direct activation was stronger and reached further in the Full Tuned model: the peak probability rose from 0.67 to 0.83 and the distance at which it falls to 0.5 from 61 to 77 μm (Figure 4.16D, E). The subthreshold response beyond the activation region was instead the same in the two models: the curves of depolarisation and hyperpolarisation coincide beyond about 100 μm and fall below 0.01 at 209 μm in both. The asymmetry near the electrodes had a similar overall size (mean absolute difference within 100 μm of the dipole centre: 0.091 in Soma Only, 0.105 in Full Tuned), although the hyperpolarised ring around the cathodes was more evident in Soma Only, and the maps of the three outcomes had the same spatial structure in the two models (Figure 4.16A–C): activation over the two electrode columns, depolarisation on the cathode flank and around the anodes, hyperpolarisation on the anode flank. In the preliminary campaign the response of the non-firing neurons also fell below 1 mV at about 150 μm (Figure 4.16F). The far-field subthreshold response is therefore set by the passive properties and the geometry, not by the dendritic conductances.

The difference in activation cannot be attributed to the active dendrites alone. The two models also differ in the leak tuning and in resting potential (-83.4 mV in Soma Only against -73.8 mV in Full Tuned), so that the Soma Only neurons start about 10 mV further from threshold. The comparison therefore shows how the results change between the preliminary and the final model, not the contribution of a single mechanism.

Table 4.6. Preliminary and final campaigns.

	Soma Only	Full Tuned
Simulations (neurons)	3,672,000 (1,224,000)	2,592,000 (864,000)
Simulated area	1000 × 1000 μm	600 × 600 μm
Pulses per simulation	1	1
Somatic rest	−83.4 mV	−73.8 mV
Drift without stimulus at t_0	+0.081 mV	< 10 ^{−4} mV
Peak P(activation)	0.67 at 27.5 μm	0.83 at 47.5 μm
P(activation) = 0.5 at	61 μm	77 μm
P(activation) = 0.1 at	92 μm	99 μm
P(activation) = 0.01 at	102 μm	120 μm
Subthreshold probability < 0.01 beyond	209 μm	209 μm
Subthreshold / activated	4.4	2.9
Mirror asymmetry within 100 μm	0.091	0.105
Mirror asymmetry beyond 100 μm	0.009	0.008
Same outcome at the three thicknesses	99.3%	97.9%
Expected in a 1 mm ² well (1700 neurons): activated / subthreshold	21 / 91	29 / 83




Figure 4.16. Preliminary campaign with the Soma Only model (3,672,000 simulations). (A–C) Probability that a neuron is activated, depolarised or hyperpolarised as a function of the position of its soma. (D) Probability of the three outcomes against the distance from the dipole centre, with 95% Wilson intervals. (E) Profile of the three probabilities along x, for somas within 90 μm of the array axis (A: anode column; C: cathode column). (F) Absolute change in somatic potential at the end of the pulse in non-firing neurons (median, 25–75% and 5–95%); dashed line: the 1 mV threshold. Same bins, same threshold and same colours as in Figure 4.13.

4.5.7. One sample of 2000 neurons

In a single sample of 2000 neurons (80 μm slab), 104 neurons fired, 143 were depolarised and 148 hyperpolarised (Figure 4.17A–C). The neurons were arranged as in the probability maps: activated neurons over the electrodes, depolarised neurons mainly on the cathode side, hyperpolarised neurons mainly on the anode side. Each neuron responds deterministically; the probability maps describe how often each response occurs at each position.




Figure 4.17. One sample of 2000 neurons. Somas of one sample of 2000 independent neurons (80 μm slab; grey: no effect above threshold). (A) Activated neurons. (B) Neurons depolarised at the end of the pulse, coloured by the size of the change. (C) Hyperpolarised neurons. (D) Number of neurons in each class in this sample (solid bars) and mean number per sample over the whole campaign, i.e. 2000 times the overall probability of each outcome (pale bars; whiskers: ± 2 binomial standard deviations).

Chapter 5

Discussion

The question addressed in this thesis is what an electrical stimulus delivered through a multi-electrode array does, at the level of the single neuron, to a culture of human derived neurons. A patient-derived network on an array is a biological twin whose activity can be recorded and perturbed, but a change in its activity after stimulation does not identify the mechanism behind it [N. Doorn, 2024; N. Doorn, 2025; L.S. Fekonja, 2024]. To address this, I built an *in silico* pipeline that couples reconstructed human pyramidal morphologies and human membrane biophysics to the extracellular field of the stimulating electrodes, simulates large samples of independent neurons in which position, orientation, morphology and dendritic complexity vary, and expresses the responses as probabilities of activation, of subthreshold depolarisation and of subthreshold hyperpolarisation as a function of position. I also asked how much these probabilities depend on the features of the model that are least constrained by data: the thickness of the culture, the morphology and the properties of the axon.

The starting point was the existing literature, which provided the components but not their combination. Detailed models of cortical neurons, including human ones, supply the morphologies and the conductances [E. Hay, 2011; S. Rich, 2021; H. Mohan, 2015; G. Eyal, 2016; G. Eyal, 2018], but they were developed to study the intrinsic properties of the neuron and not its response to an extracellular field. The theory of extracellular stimulation, and its application to realistic morphologies, explains how a field polarises the membrane and where excitation starts [F. Rattay, 1986; F. Rattay, 1999; A.S. Aberra, 2018; G.-S. Yi, 2017], but it considers one neuron at a time. And stimulation of cultures on multi-electrode arrays has been characterised experimentally and modelled at the level of the field and of the network [D.A. Wagenaar, 2004; J. Wen, 2022; G. Zanini, 2023; G. Parodi, 2024; S. Joucla, 2014], without a neuron-level description of how the stimulus acts across the variability of a culture. The pipeline presented here joins these three lines. The first result is that direct activation is local. The probability of firing is high only close to the electrodes and falls to zero within a short distance from them, so that only a small fraction of the neurons of the culture fires in direct response to the pulse. This is

consistent with the classical decrease of activation with distance from the electrode [E.J. Tehovnik, 2006]. In vivo, however, microstimulation recruits sparse and spatially distributed populations, largely through axons passing near the electrode [M.H. Histed, 2009]; this mode of activation is absent from the model, whose axons are stylised, so the activation maps describe the direct recruitment of somas and proximal neurites only.

The second result is that stimulation is not all-or-none. Around the neurons that fire there is a larger population that is polarised without firing, and this population extends further from the electrodes than activation does. The sign of the polarisation depends on the side of the electrode group: at the end of the pulse, neurons on the side of the cathodes are mostly depolarised and neurons on the side of the anodes mostly hyperpolarised, as expected from the sign of the field on the two sides. The neurons of this outer population have the sign opposite to that of the neurons next to the electrodes on the same side, because below threshold the field acts through axial currents along the neurites: the part of the neuron closest to a negative electrode is depolarised and the rest is hyperpolarised, so a distant soma is polarised in the direction opposite to its dendrites that reach towards the electrodes. Subthreshold polarisation by electric fields, and its dependence on the morphology and orientation of the cell, is well documented for uniform fields [G.-S. Yi, 2017], and regions of hyperpolarisation flanking the region of excitation are a known feature of extracellular stimulation [F. Rattay, 1986; F. Rattay, 1999]. What this work adds is their extent and probability for the electrode configuration and the pulse used in the laboratory. A single pulse therefore does three things at once: it makes a few neurons fire, it brings others closer to threshold and it moves others away from it.

Close to the electrodes the two sides are not mirror images: more neurons fire on the side of the cathodes than on the side of the anodes, and a ring of depolarised neurons surrounds the anodes without an equivalent hyperpolarised ring around the cathodes. This asymmetry comes from the order of the two phases of the pulse, in line with the known dependence of the recruited neurons on the waveform of a biphasic stimulus [D.R. Merrill, 2005]. Near the cathodes the depolarising phase comes first, so most neurons fire and leave the subthreshold classes; near the anodes it comes second, after a hyperpolarisation, and about half of the neurons stay depolarised below threshold. With the two phases in the opposite order the two sides were exchanged, in both models. The size of the effect depends on the resting potential: the preliminary model rests about 10 mV further from threshold, fewer of its neurons fire near the cathodes and the

hyperpolarised ring remains visible, and shifting its resting potential to that of the final model removed the difference. The active dendritic conductances are therefore not required. These control simulations were run on one morphology.

The third result is that the outcome depends on the position of the neuron more than on the details of the cell. The thickness of the slab, which sets how much of the dendritic tree is retained, left the outcome almost unchanged; the six morphologies gave similar curves, with moderate differences in their peak and in the distance at which activation is lost; and the properties of the axon, whether passive or with an active initial segment, long or short, changed the outcome in a negligible number of cases. This robustness is what makes the reduction to probabilities meaningful: the map is determined mainly by the geometry of the stimulation, which is known, and only weakly by the features of the neuron that were varied here (slab thickness, morphology, axon); the resting potential and the discretisation of the soma, by contrast, changed it appreciably.

One result differs from the literature. In the model the spike starts at the soma, whereas experimental and modelling studies place the site of lowest threshold in the axon, at the initial segment or along the fibre [F. Rattay, 1999; A.S. Aberra, 2018; M.H. Histed, 2009; L.G. Nowak, 1998], and recordings on high-density arrays show the same for cultured neurons [M. Radivojevic, 2016]. This is most likely a property of the model and not a finding about human neurons: in the model the axon is stylised and not reconstructed, the axon is passive in the final model (and in the Full Active variant the sodium conductance of its initial segment is much lower than that of the soma), and the soma lies immediately above the plane of the electrodes, where the field is strongest. The conclusion that the axon has little influence on the map therefore holds for this model and should not be generalised.

Other limits define how far the results can be taken: the soma is a single compartment, the neurons are adult cortical neurons and not hiPSC-derived ones, the field is that of point sources in a homogeneous medium, only one protocol was simulated, and the map has not yet been compared with recordings (Chapter 7). What the stimulus model can hand to a network model is therefore, for each neuron, according to its position, a probability of firing and a probability of being left depolarised or hyperpolarised. In the laboratory, stimulation changes the fragmented bursts of the patient-derived networks (Chapter 4); since the pulse directly reaches only the neurons near the electrodes, a change at the level of the whole network must be produced by the network itself. Placing

the map on a network model calibrated on the recordings, for instance with simulation-based inference [N. Doorn, 2024; N. Doorn, 2025], would allow one to ask which network mechanisms must be engaged for the phenotype to change.

In conclusion, the direct effect of extracellular stimulation on a culture of human neurons is local, graded and set mainly by position. A stimulus on a multi-electrode array therefore cannot be represented only by the neurons it makes fire, because for each of them several others are polarised without firing. The probability maps obtained here are the first component of a digital twin in which the stimulus is a parameter that can be explored *in silico*: the biological twin shows what changes after stimulation, and the digital twin, once it includes the network, can be used to ask why.

Chapter 6

Future perspectives

The pipeline developed in this thesis is a model of the stimulus: it describes what a pulse does to each neuron, depending on its position relative to the stimulating electrodes. Four directions follow from this result: bringing the single-neuron model closer to the biological twin, validating the predicted maps against MEA recordings, coupling the stimulus model with a network model calibrated on patient data by simulation-based inference (SBI), and using the coupled model to design stimulation protocols in silico.

6.1. Towards a digital twin of the biological twin

The digital twin should reproduce the system recorded on the MEA, rather than the adult brain. The current model combines reconstructed adult human layer 2/3 pyramidal morphologies with a human layer 5 channel model [S. Rich, 2021], whereas the biological twin consists of excitatory neurons derived from hiPSCs through Ngn2 overexpression, which are functional at DIV 35 but immature compared with adult cortical neurons. Three refinements would reduce this gap.

The adult reconstructions can be replaced with reconstructions of Ngn2 neurons grown in the same conditions as the biological twin, obtained by sparse labelling, confocal imaging and three-dimensional tracing. The slab-based reduction can then be applied unchanged, with the slab thickness measured directly in the images and the slab defined from the substrate upwards, so that no branch lies below the electrode plane.

Channel densities can be constrained with patch-clamp recordings from the same neurons (resting potential, input resistance, rheobase, action-potential shape, sodium and potassium currents), separately for control and Dup15q lines, which would make the single-neuron model patient-specific.

hiPSC-derived neurons are cultured with astrocytes, which may make the conductivity of the medium non-uniform and shape network activity after stimulation through potassium

buffering and glutamate uptake. The first effect can be included in the field computation, the second in the network model.

6.2. Experimental validation on HD-MEA

The probability maps are predictions that can be tested directly on the biological twin. Stimulation is delivered with the same protocol and electrode configuration as in the model, and excitatory synaptic transmission is blocked with CNQX (AMPA receptor antagonist) and APV (NMDA receptor antagonist), so that the short-latency spikes that remain after the stimulus artefact reflect direct activation only. For each recording electrode, the experimental response probability is the fraction of pulses followed by at least one spike within the post-stimulus window.

Soma positions are not known experimentally, but each electrode detects spikes from neurons within a detection radius of a few tens of micrometres. The corresponding prediction is the probability that at least one neuron within this radius is directly activated, obtained from the activation map and the neuronal density under the assumption of random placement. Experimental and predicted probabilities are then compared electrode by electrode, as a function of distance and angle from the centre of the stimulating electrodes. Since direct activation is predicted to vanish within about 140 μm of the stimulating electrodes, only the few tens of electrodes within a few hundred micrometres are informative; more distant electrodes serve as negative controls, and those adjacent to the stimulating ones are excluded because of the stimulus artefact. The detection radius is an assumption and should be varied, and the density should be measured independently. Subthreshold depolarisation and hyperpolarisation are not visible in extracellular recordings and could only be assessed indirectly, for instance as changes in spontaneous firing probability after each pulse.

Extending the dataset. The comparison between conditions in this thesis rests on six wells. Stimulating control networks, and recording unstimulated Dup15q wells over the same interval, in more wells and plates, would establish whether the changes are specific to the stimulus and to the genotype. Recording at several times after the protocol would distinguish an acute from a lasting effect.

6.3. Coupling with an SBI-calibrated network digital twin

The stimulus model describes what a pulse does to each neuron, but not how the network responds; network models of hiPSC-derived neuronal networks on MEA describe how the network is built and how it evolves spontaneously, but not how an extracellular stimulus acts on it. Coupling the two is the step required to address the question raised by the biological twin: which network mechanisms explain the shorter Super Bursts and the longer intervals between them observed after stimulation in Dup15q networks.

A biophysical network model of hiPSC-derived excitatory networks on MEA has been developed by Doorn et al. [N. Doorn, 2024; N. Doorn, 2023]: Hodgkin–Huxley neurons connected by AMPA and NMDA receptor-mediated synapses with short-term synaptic depression [H. Markram, 1998], placed on a grid with virtual electrodes that mimic the MEA recordings. Its parameters cannot be measured directly, and many different combinations can produce similar activity. SBI estimates their posterior distribution given a recording without evaluating the likelihood [K. Cranmer, 2020; M. Deistler, 2025], by training a neural density estimator on pairs of parameters and simulated features [G. Papamakarios, 2016; P.J. Gonçalves, 2020]. With this approach, Doorn et al. [N. Doorn, 2025] identified parameter changes consistent with pharmacological manipulations and with known disease mechanisms in patient-derived networks.

The network model is calibrated with SBI on spontaneous recordings of Dup15q and control networks. Summary features must be computed identically on recorded and simulated data: since the model contains at most one hundred virtual electrodes and the HD-MEA records from 2304, the recorded electrodes must be selected or pooled to match the virtual layout [N. Doorn, 2025], and the features must capture Super Burst structure, for instance the number of fragments per burst and the Super Burst duration and interval. The cluster analysis of this thesis is a first step in this direction: a 3×3 cluster reproduced the Super Burst features of the whole well, so the clusters are candidates to provide the reduced set of signals to which the virtual electrodes are matched. The state-space measures can be added to the summary features or kept for the posterior predictive check.

The six stimulating electrodes are placed at their physical positions within the virtual electrode grid, and each network neuron is assigned its coordinates relative to them. If the neuronal density of the model is comparable to that of the culture, the outcome of each neuron at each pulse is drawn directly from the maps; if it is lower, each network neuron represents a region of the culture and receives the average probability over that region. An activated neuron fires a spike at the pulse, a depolarised or hyperpolarised neuron receives a membrane-potential offset drawn from the simulated distribution in its bin, and a neutral neuron is unaffected.

The calibrated network is then simulated with the laboratory protocol, 36 biphasic pulses at 0.2 Hz, with parameters sampled from the posterior inferred from the pre-stimulation recordings, and its features are compared with those of the post-stimulation recordings [M. Deistler, 2025]. The test succeeds if the model reproduces the shorter Super Bursts and longer inter-Super-Burst intervals in Dup15q networks. If the effect is acute, it may emerge from the interaction between the stimulus and unchanged network parameters, for instance through depletion of synaptic resources, and the mechanism can be identified by removing components of the model, such as short-term depression or adaptation. If it persists after stimulation has ended, the network itself has changed, and posteriors inferred separately from pre- and post-stimulation recordings identify the parameters affected by stimulation [N. Doorn, 2025].

Methodological requirements. Three recommendations from Doorn et al. [N. Doorn, 2025] apply directly: posteriors should be compared only between networks cultured and recorded on the same MEA, parameter changes should be considered only if consistent across independent batches, and SBI can only identify changes that are reflected in the chosen features. Model misspecification must also be checked, with prior and posterior predictive checks and simulation-based calibration [M. Deistler, 2025].

6.4. In silico design of stimulation protocols

Once validated, the coupled model can be used to explore stimulation protocols before testing them *in vitro*: electrode configuration and number, amplitude, waveform, inter-pulse interval and train duration. Each protocol defines new probability maps, which the modular pipeline recomputes by changing only the stimulation block. Since the network

parameters are known only through a posterior distribution, the choice of protocol can be framed as Bayesian decision making [M. Deistler, 2025]: the best protocol is the one that, on average over the posterior, brings the features of the stimulated Dup15q network closest to those of control networks, possibly with a penalty on the total charge delivered. The most promising protocols can then be tested on the biological twin, closing the loop between the digital and the biological twin.

Chapter 7

Limitations

This chapter summarises the limitations of the work. The axon is stylised and not reconstructed, because the reconstructions contain only a short fragment of it (Methods); since axonal morphology strongly shapes the response to extracellular stimulation [A.S. Abera, 2018], the model does not describe axonal initiation, and a neuron whose axon passes close to an electrode may fire at a distance at which the model predicts no activation. The morphologies are those of adult human cortical neurons and the membrane model is that of a human layer-5 pyramidal neuron [S. Rich, 2021], with its distance-dependent channel distributions transferred to the six morphologies; neither is fitted to the hiPSC-derived neurons of the culture, which are smaller and less mature. The reduction of the three-dimensional morphology to a slab is a modelling rule, and the three thicknesses tested are a sensitivity analysis and not a measurement of the culture. The field is that of point sources in a homogeneous medium, with a compact arrangement of three anodes and three cathodes, whereas the electrodes selected in the experiment depend on the activity and may be arranged differently; the neurons do not feed back on the field. The soma is a single compartment and the tests of axon length and of spatial convergence, like the control simulations on the asymmetry between the two sides, were run on one morphology.

The probabilities are specific to the protocol: they hold for the amplitude, the pulse duration, the waveform and the electrode configuration that were simulated, and for pulses far enough apart to be treated as independent (0.2 Hz). A different protocol requires new simulations. The maps have not yet been compared with recordings, and only the probability of direct activation can be tested on a multi-electrode array, because subthreshold polarisation cannot be observed (Chapter 6).

The experimental analyses are exploratory. They rest on six wells of one patient-derived line, without stimulated control networks or unstimulated wells recorded over the same interval; the changes after stimulation should therefore be read as indications and not as established effects.

Chapter 8

Appendix

Table A.1. Whole-well Super Burst rate before and after stimulation. Super Burst rate (min^{-1}), computed as the number of Super Bursts detected in each well divided by its recording duration, before (PRE) and after (POST) electrical stimulation for MultiWells 2186 and 2013. Each well is identified by its plate number.

MultiWell	Well	Super Burst rate PRE (min^{-1})	Super Burst rate POST (min^{-1})
2186	A1	1.00	1.26
2186	B2	1.55	0.00
2186	B3	1.50	0.84
2013	A2	2.05	3.46
2013	A3	1.10	1.15
2013	B1	1.20	3.25

Table A.2. Whole-well fragment rate before and after stimulation. Fragment rate (min^{-1}), computed as the total number of intra-envelope IFR peaks (fragments) detected within Super Bursts in each well divided by its recording duration, before (PRE) and after (POST) electrical stimulation for MultiWells 2186 and 2013. Each well is identified by its plate number.

MultiWell	Well	Fragment rate PRE (min^{-1})	Fragment rate POST (min^{-1})
2186	A1	60.85	36.02
2186	B2	15.65	0.00
2186	B3	73.35	27.64
2013	A2	84.30	53.51
2013	A3	95.85	63.77
2013	B1	133.45	58.95

Table A.3. Whole-well temporal Super Burst features. Mean Super Burst duration (SBD) and mean inter-Super-Burst-event interval (ISBEI) in each analysed well during PRE and POST electrical stimulation for MultiWells 2186 and 2013. Values are reported for the five wells in which Super Bursts persisted after stimulation; in well B2 (MultiWell 2186) Super Bursts were abolished after stimulation, so SBD and ISBEI are

not defined in the POST condition (indicated by "—") and this well is excluded from the corresponding comparisons. Each well is identified by its plate number.

MultiWell	Well	SBD PRE (s)	SBD POST (s)	ISBEI PRE (s)	ISBEI POST (s)
2186	A1	19.9	9.6	37.8	40.3
2186	B2	5.6	—	27.4	—
2186	B3	19.4	14.6	21.3	47.9
2013	A2	19.9	8.1	9.6	10.0
2013	A3	35.4	22.7	19.2	32.3
2013	B1	45.6	9.5	4.6	9.4

Table A.4. Median mutual information for each well (nats, k = 10).

MultiWell	Well	Clusters (PRE / POST)	Cluster-to-network PRE	Cluster-to-network POST	Cluster-to-cluster PRE	Cluster-to-cluster POST
2186	A1	10 / 10	0.83	0.75	0.61	0.53
2186	B2	5 / 8	0.65	0.59	0.45	0.41
2186	B3	10 / 10	0.58	0.64	0.44	0.46
2013	A2	10 / 10	0.70	0.79	0.47	0.54
2013	A3	10 / 10	0.88	0.72	0.65	0.52
2013	B1	10 / 10	0.72	0.85	0.48	0.60

Table A.5. Variance explained by the first three principal components. Percentage of the total variance captured by PC1, PC2 and PC3, and their cumulative value, for each well before and after stimulation.

MultiWell	Well	PC1 PRE (%)	PC1 POST (%)	PC2 PRE (%)	PC2 POST (%)	PC3 PRE (%)	PC3 POST (%)	PC1–3 PRE (%)	PC1–3 POST (%)
2186	A1	38.0	37.3	2.34	2.26	0.56	0.52	40.9	40.1
2186	B2	19.2	24.4	2.59	2.68	0.60	0.55	22.4	27.6
2186	B3	26.3	28.4	2.85	2.55	0.65	0.58	29.7	31.5
2013	A2	23.6	24.1	2.85	2.16	0.66	0.40	27.1	26.7
2013	A3	31.2	31.0	2.55	2.21	0.60	0.53	34.4	33.8
2013	B1	19.6	22.2	2.82	2.08	0.79	0.39	23.3	24.7
	Mean ± SD	26.3 ± 7.1	27.9 ± 5.6	2.67 ± 0.21	2.32 ± 0.24	0.64 ± 0.08	0.49 ± 0.08	29.6 ± 6.9	30.7 ± 5.6

Table A.6. Peak trajectory speed (standardised units per second) at the fragment and Super Burst scales, before and after stimulation. Each value is the mean over the events of the well. Well B2 is excluded.

MultiWell	Well	Fragment PRE	Fragment POST	Super Burst PRE	Super Burst POST
2186	A1	1190.5	1268.4	731.0	840.5
2186	B3	814.7	833.9	575.6	593.3
2013	A2	680.7	648.1	509.3	562.6
2013	A3	894.3	955.4	701.5	781.8
2013	B1	579.2	672.9	468.4	555.6
	Mean $\pm$ SEM	831.9 $\pm$ 104.8	875.7 $\pm$ 113.0	597.2 $\pm$ 51.8	666.8 $\pm$ 60.0

Table A.7. Scale-normalised trajectory speed (per second) at the fragment and Super Burst scales. Each value is the mean over the events of the well. Well B2 is excluded.

MultiWell	Well	Fragment PRE	Fragment POST	Super Burst PRE	Super Burst POST
2186	A1	47.04	46.62	37.10	37.41
2186	B3	43.10	41.18	40.08	37.87
2013	A2	36.14	33.41	36.61	37.17
2013	A3	40.18	41.24	42.06	45.13
2013	B1	36.85	36.32	40.91	40.79
	Mean $\pm$ SEM	40.66 $\pm$ 2.02	39.75 $\pm$ 2.27	39.35 $\pm$ 1.07	39.67 $\pm$ 1.51

Table A.8. Inter-fragment interval (IFI) before and after stimulation. The IFI of each event is the median of its inter-fragment intervals; the value of each well is the mean over its events. Well B2 is excluded.

MultiWell	Well	IFI PRE (s)	IFI POST (s)
2186	A1	0.376	0.381
2186	B3	0.370	0.449
2013	A2	0.460	0.493
2013	A3	0.397	0.398
2013	B1	0.381	0.499
	Mean $\pm$ SEM	0.397 $\pm$ 0.016	0.444 $\pm$ 0.024

Table A.9. Generalised variance (squared standardised units) of the cloud of states.

MultiWell	Well	Network bursts PRE	Network bursts POST	Super Bursts PRE	Super Bursts POST
2186	A1	150.4	175.3	72.0	86.0
2186	B3	116.5	136.1	54.2	62.2
2013	A2	118.9	123.1	56.6	60.9
2013	A3	142.8	153.2	66.8	70.0
2013	B1	88.1	108.7	47.6	51.0
	Mean $\pm$ SEM	123.4 $\pm$ 11.0	139.3 $\pm$ 11.6	59.4 $\pm$ 4.4	66.0 $\pm$ 5.8

Table A.10. Negentropy (nats) of the cloud of states.

MultiWell	Well	Network bursts PRE	Network bursts POST	Super Bursts PRE	Super Bursts POST
2186	A1	2.241	2.120	1.730	1.691
2186	B3	1.937	1.834	2.076	2.394
2013	A2	1.684	1.300	2.342	2.162
2013	A3	2.036	2.053	2.319	2.300
2013	B1	1.371	1.247	1.740	1.901
	Mean $\pm$ SEM	1.854 $\pm$ 0.150	1.711 $\pm$ 0.185	2.041 $\pm$ 0.134	2.090 $\pm$ 0.130

Table A.11. Differential entropy (nats) of the cloud of states.

MultiWell	Well	Network bursts PRE	Network bursts POST	Super Bursts PRE	Super Bursts POST
2186	A1	5.610	5.885	5.385	5.601
2186	B3	5.659	5.917	4.755	4.574
2013	A2	5.933	6.351	4.532	4.784
2013	A3	5.763	5.816	4.721	4.786
2013	B1	5.945	6.279	4.961	4.869
	Mean $\pm$ SEM	5.782 $\pm$ 0.069	6.050 $\pm$ 0.110	4.871 $\pm$ 0.145	4.923 $\pm$ 0.176

Table A.12. Sensitivity to the subthreshold threshold (Full Tuned).

Threshold	Depolarised	Hyperpolarised	Subthreshold probability < 0.01 beyond
0.5 mV	9.63%	8.77%	230 μ m
1 mV	7.22%	6.36%	209 μ m
2 mV	5.16%	4.31%	190 μ m

Table A.13. Activation by slab thickness (Full Tuned).

Thickness	Overall P(activation)	Peak	P = 0.5 at	P = 0.1 at
40 μ m	4.52%	0.803	75.5 μ m	97.8 μ m
80 μ m	4.75%	0.841	78.0 μ m	99.4 μ m
120 μ m	4.78%	0.844	78.1 μ m	99.6 μ m

Table A.14. Activation by morphology (Full Tuned).

Specimen	Rest (mV)	Overall P(activation)	Peak	P = 0.5 at	P = 0.1 at
60303	-74.0	4.22%	0.846	67.5 μ m	95.7 μ m
60308	-74.1	4.51%	0.818	75.8 μ m	97.6 μ m
60311	-73.4	5.25%	0.887	82.8 μ m	102.5 μ m
130303	-73.2	5.46%	0.885	80.6 μ m	105.2 μ m
130305	-74.3	4.46%	0.808	74.6 μ m	98.5 μ m
130306	-74.0	4.19%	0.751	72.9 μ m	95.8 μ m

Bibliography

- [A.S. Aberra, 2018] A.S. Aberra, A.V. Peterchev, W.M. Grill. Biophysically realistic neuron models for simulation of cortical stimulation. *Journal of Neural Engineering*, 15(6):066023, 2018.
- [H. Amin, 2016] H. Amin, A. Maccione, F. Marinaro, S. Zordan, T. Nieuws, L. Berdondini. Electrical responses and spontaneous activity of human iPS-derived neuronal networks characterized for 3-month culture with 4096-electrode arrays. *Frontiers in Neuroscience*, 10:121, 2016.
- [L. Berdondini, 2009] L. Berdondini, K. Imfeld, A. Maccione, M. Tedesco, S. Neukom, M. Koudelka-Hep, S. Martinoia. Active pixel sensor array for high spatio-temporal resolution electrophysiological recordings from single cell to large scale neuronal networks. *Lab on a Chip*, 9(18):2644–2651, 2009.
- [S.M. Chambers, 2009] S.M. Chambers, C.A. Fasano, E.P. Papapetrou, M. Tomishima, M. Sadelain, L. Studer. Highly efficient neural conversion of human ES and iPS cells by dual inhibition of SMAD signaling. *Nature Biotechnology*, 27(3):275–280, 2009.
- [T.M. Cover, 2006] T.M. Cover, J.A. Thomas. *Elements of information theory*. 2nd ed. Wiley-Interscience, Hoboken, NJ, 2006.
- [K. Cranmer, 2020] K. Cranmer, J. Brehmer, G. Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences of the United States of America*, 117(48):30055–30062, 2020.
- [J.P. Cunningham, 2014] J.P. Cunningham, B.M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014.
- [Deciphering Developmental Disorders Study, 2015] Deciphering Developmental Disorders Study. Large-scale discovery of novel genetic causes of developmental disorders. *Nature*, 519(7542):223–228, 2015.
- [M. Deistler, 2025] M. Deistler, J. Boelts, P. Steinbach, G. Moss, T. Moreau, M. Gloeckler, P.L.C. Rodrigues, J. Linhart, J.K. Lappalainen, B.K. Miller, et al. Simulation-based inference: a practical guide. *arXiv preprint arXiv:2508.12939*, 2025.
- [N. Doorn, 2023] N. Doorn, E.J.H. van Hughte, U. Ciptasari, A. Mordelt, H.G.E. Meijer, D. Schubert, M. Frega, N. Nadif Kasri, M.J.A.M. van Putten. An in silico and in vitro

human neuronal network model reveals cellular mechanisms beyond Nav1.1 underlying Dravet syndrome. *Stem Cell Reports*, 18(8):1686–1700, 2023.

[N. Doorn, 2024] N. Doorn, E.J.H.F. Voogd, M.R. Levers, M.J.A.M. van Putten, M. Frega. Breaking the burst: unveiling mechanisms behind fragmented network bursts in patient-derived neurons. *Stem Cell Reports*, 19(11):1583–1597, 2024.

[N. Doorn, 2025] N. Doorn, M.J.A.M. van Putten, M. Frega. Automated inference of disease mechanisms in patient-hiPSC-derived neuronal networks. *Communications Biology*, 8(1):768, 2025.

[M. Elamin, 2023] M. Elamin, A. Dumarchey, C. Stoddard, T.M. Robinson, C. Cowie, D. Gorka, S.J. Chamberlain, E.S. Levine. The role of UBE3A in the autism and epilepsy-related Dup15q syndrome using patient-derived, CRISPR-corrected neurons. *Stem Cell Reports*, 18(4):884–898, 2023.

[G. Eyal, 2016] G. Eyal, M.B. Verhoog, G. Testa-Silva, Y. Deitcher, J.C. Lodder, R. Benavides-Piccione, J. Morales, J. DeFelipe, C.P.J. de Kock, H.D. Mansvelder, et al. Unique membrane properties and enhanced signal processing in human neocortical neurons. *eLife*, 5:e16553, 2016.

[G. Eyal, 2018] G. Eyal, M.B. Verhoog, G. Testa-Silva, Y. Deitcher, R. Benavides-Piccione, J. DeFelipe, C.P.J. de Kock, H.D. Mansvelder, I. Segev. Human cortical pyramidal neurons: from spines to spikes via models. *Frontiers in Cellular Neuroscience*, 12:181, 2018.

[L.S. Fekonja, 2024] L.S. Fekonja, R. Schenk, E. Schröder, R. Tomasello, S. Tomšič, T. Picht. The digital twin in neuroscience: from theory to tailored therapy. *Frontiers in Neuroscience*, 18:1454856, 2024.

[J.J. Fink, 2021] J.J. Fink, J.D. Schreiner, J.E. Bloom, J. James, D.S. Baker, T.M. Robinson, R. Lieberman, L.M. Loew, S.J. Chamberlain, E.S. Levine. Hyperexcitable phenotypes in induced pluripotent stem cell-derived neurons from patients with 15q11-q13 duplication syndrome, a genetic form of autism. *Biological Psychiatry*, 90(11):756–765, 2021.

[M. Frega, 2017] M. Frega, S.H.C. van Gestel, K. Linda, J. van der Raadt, J. Keller, J.R. van Rhijn, D. Schubert, C.A. Albers, N. Nadif Kasri. Rapid neuronal differentiation of induced pluripotent stem cells for measuring network activity on micro-electrode arrays. *Journal of Visualized Experiments*, (119):e54900, 2017.

- [M. Frega, 2019] M. Frega, K. Linda, J.M. Keller, G. Gümüş-Akay, B. Mossink, J.R. van Rhijn, M. Negwer, T. Klein Gunnewiek, K. Foreman, N. Kompier, et al. Neuronal network dysfunction in a model for Kleefstra syndrome mediated by enhanced NMDAR signaling. *Nature Communications*, 10(1):4928, 2019.
- [J. Frohlich, 2016] J. Frohlich, D. Senturk, V. Saravanapandian, P. Golshani, L.T. Reiter, R. Sankar, R.L. Thibert, C. DiStefano, S. Huberty, E.H. Cook, et al. A quantitative electrophysiological biomarker of duplication 15q11.2-q13.1 syndrome. *PLoS ONE*, 11(12):e0167179, 2016.
- [J. Frohlich, 2019] J. Frohlich, L.T. Reiter, V. Saravanapandian, C. DiStefano, S. Huberty, C. Hyde, S. Chamberlain, C.E. Bearden, P. Golshani, A. Irimia, et al. Mechanisms underlying the EEG biomarker in Dup15q syndrome. *Molecular Autism*, 10(1):29, 2019.
- [J.A. Gallego, 2017] J.A. Gallego, M.G. Perich, L.E. Miller, S.A. Solla. Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, 2017.
- [J.A. Gallego, 2018] J.A. Gallego, M.G. Perich, S.N. Naufel, C. Ethier, S.A. Solla, L.E. Miller. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, 9(1):4233, 2018.
- [P. Gao, 2017] P. Gao, E. Trautmann, B.M. Yu, G. Santhanam, S. Ryu, K. Shenoy, S. Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 214262, 2017.
- [P.L. Germain, 2017] P.L. Germain, G. Testa. Taming human genetic variability: transcriptomic meta-analysis guides the experimental design and interpretation of iPSC-based disease modeling. *Stem Cell Reports*, 8(6):1784–1796, 2017.
- [P.J. Gonçalves, 2020] P.J. Gonçalves, J.-M. Lueckmann, M. Deistler, M. Nonnenmacher, K. Öcal, G. Bassetto, C. Chintaluri, W.F. Podlaski, S.A. Haddad, T.P. Vogels, et al. Training deep neural density estimators to identify mechanistic models of neural dynamics. *eLife*, 9:e56261, 2020.
- [E. Hay, 2011] E. Hay, S. Hill, F. Schürmann, H. Markram, I. Segev. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS Computational Biology*, 7(7):e1002107, 2011.
- [H.O. Heyne, 2018] H.O. Heyne, T. Singh, H. Stamberger, R. Abou Jamra, H. Caglayan, D. Craiu, P. De Jonghe, R. Guerrini, K.L. Helbig, B.P.C. Koeleman, et al. De novo

variants in neurodevelopmental disorders with epilepsy. *Nature Genetics*, 50(7):1048–1053, 2018.

[M.L. Hines, 2001] M.L. Hines, N.T. Carnevale. NEURON: a tool for neuroscientists. *The Neuroscientist*, 7(2):123–135, 2001.

[M.H. Histed, 2009] M.H. Histed, V. Bonin, R.C. Reid. Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. *Neuron*, 63(4):508–522, 2009.

[J.L. Horn, 1965] J.L. Horn. A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185, 1965.

[A. Hyvärinen, 2000] A. Hyvärinen, E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4–5):411–430, 2000.

[S. Joucla, 2014] S. Joucla, A. Glière, B. Yvert. Current approaches to model extracellular electrical neural microstimulation. *Frontiers in Computational Neuroscience*, 8:13, 2014.

[L. Kalsner, 2015] L. Kalsner, S.J. Chamberlain. Prader-Willi, Angelman, and 15q11-q13 duplication syndromes. *Pediatric Clinics of North America*, 62(3):587–606, 2015.

[M.H.P. Kole, 2012] M.H.P. Kole, G.J. Stuart. Signal processing in the axon initial segment. *Neuron*, 73(2):235–247, 2012.

[M. Komarov, 2019] M. Komarov, P. Malerba, R. Golden, P. Nunez, E. Halgren, M. Bazhenov. Selective recruitment of cortical neurons by electrical stimulation. *PLoS Computational Biology*, 15(8):e1007277, 2019.

[L.F. Kozachenko, 1987] L.F. Kozachenko, N.N. Leonenko. Sample estimate of the entropy of a random vector. *Problems of Information Transmission*, 23(2):95–101, 1987.

[A. Kraskov, 2004] A. Kraskov, H. Stögbauer, P. Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, 2004.

[A. Maccione, 2009] A. Maccione, M. Gandolfo, P. Massobrio, A. Novellino, S. Martinoia, M. Chiappalone. A novel algorithm for precise identification of spikes in extracellularly recorded neuronal signals. *Journal of Neuroscience Methods*, 177(1):241–249, 2009.

[H. Markram, 1998] H. Markram, Y. Wang, M. Tsodyks. Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences of the United States of America*, 95(9):5323–5328, 1998.

- [M. Mastrangelo, 2026] M. Mastrangelo, C. Di Bonaventura, G. Lentini, F. Pisani. An updated overview on clinical and developmental trajectories of developmental and epileptic encephalopathies: from childhood to adulthood. *Seizure*, 134:190–197, 2026.
- [C.C. McIntyre, 1999] C.C. McIntyre, W.M. Grill. Excitation of central nervous system neurons by nonuniform electric fields. *Biophysical Journal*, 76(2):878–888, 1999.
- [D.R. Merrill, 2005] D.R. Merrill, M. Bikson, J.G.R. Jefferys. Electrical stimulation of excitable tissue: design of efficacious and safe protocols. *Journal of Neuroscience Methods*, 141(2):171–198, 2005.
- [H. Mohan, 2015] H. Mohan, M.B. Verhoog, K.K. Doreswamy, G. Eyal, R. Aardse, B.N. Lodder, N.A. Goriounova, B. Asamoah, A.B.C.B. Brakspear, C. Groot, et al. Dendritic and axonal architecture of individual pyramidal neurons across layers of adult human neocortex. *Cerebral Cortex*, 25(12):4839–4853, 2015.
- [B. Mossink, 2021] B. Mossink, A.H.A. Verboven, E.J.H. van Hugte, T.M. Klein Gunnewiek, G. Parodi, K. Linda, C. Schoenmaker, T. Kleefstra, T. Kozicz, H. van Bokhoven, et al. Human neuronal networks on micro-electrode arrays are a highly robust tool to study disease-specific genotype-phenotype correlations in vitro. *Stem Cell Reports*, 16(9):2182–2196, 2021.
- [L.G. Nowak, 1998] L.G. Nowak, J. Bullier. Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter I. Evidence from chronaxie measurements. *Experimental Brain Research*, 118(4):477–488, 1998.
- [L.M. Palmer, 2006] L.M. Palmer, G.J. Stuart. Site of action potential initiation in layer 5 pyramidal neurons. *Journal of Neuroscience*, 26(6):1854–1863, 2006.
- [G. Papamakarios, 2016] G. Papamakarios, I. Murray. Fast ϵ -free inference of simulation models with Bayesian conditional density estimation. *Advances in Neural Information Processing Systems*, 29:1028–1036, 2016.
- [G. Parodi, 2024] G. Parodi, G. Zanini, M. Chiappalone, S. Martinoia. Electrical and chemical modulation of homogeneous and heterogeneous human-iPSCs-derived neuronal networks on high density arrays. *Frontiers in Molecular Neuroscience*, 17:1304507, 2024.
- [V. Pasquale, 2010] V. Pasquale, S. Martinoia, M. Chiappalone. A self-adapting approach for the detection of bursts and network bursts in neuronal cultures. *Journal of Computational Neuroscience*, 29(1–2):213–229, 2010.

- [Y. Perez, 2025] Y. Perez, D. Velmeshev, L. Wang, M.L. White, C. Siebert, J. Baltazar, G. Zuo, J.A. Moriano, S. Chen, D.M. Steffen, et al. Single-cell analysis of dup15q syndrome reveals developmental and postnatal molecular changes in autism. *Nature Communications*, 16(1):6177, 2025.
- [K.S. Pradeepan, 2024] K.S. Pradeepan, F.P. McCready, W. Wei, M. Khaki, W. Zhang, M.W. Salter, J. Ellis, J. Martinez-Trujillo. Calcium-dependent hyperexcitability in human stem cell–derived Rett syndrome neuronal networks. *Biological Psychiatry: Global Open Science*, 4(2):100290, 2024.
- [M. Radivojevic, 2016] M. Radivojevic, D. Jäckel, M. Altermatt, J. Müller, V. Viswam, A. Hierlemann, D.J. Bakkum. Electrical identification and selective microstimulation of neuronal compartments based on features of extracellular action potentials. *Scientific Reports*, 6:31332, 2016.
- [F. Rattay, 1986] F. Rattay. Analysis of models for external stimulation of axons. *IEEE Transactions on Biomedical Engineering*, BME-33(10):974–977, 1986.
- [F. Rattay, 1999] F. Rattay. The basic mechanism for the electrical stimulation of the nervous system. *Neuroscience*, 89(2):335–346, 1999.
- [S. Rich, 2021] S. Rich, H. Moradi Chameh, V. Sekulic, T.A. Valiante, F.K. Skinner. Modeling reveals human–rodent differences in H-current kinetics influencing resonance in cortical layer 5 neurons. *Cerebral Cortex*, 31(2):845–872, 2021.
- [S. Ronchi, 2019] S. Ronchi, M. Fiscella, C. Marchetti, V. Viswam, J. Müller, U. Frey, A. Hierlemann. Single-cell electrical stimulation using CMOS-based high-density microelectrode arrays. *Frontiers in Neuroscience*, 13:208, 2019.
- [M. Schröter, 2025] M. Schröter, F. Cardes, C.V.H. Bui, L.D. Dodi, T. Gänswain, J. Bartram, L. Sadiraj, P. Hornauer, S. Kumar, M. Pascual-Garcia, et al. Advances in large-scale electrophysiology with high-density microelectrode arrays. *Lab on a Chip*, 25(19):4844–4885, 2025.
- [G. Stuart, 1997] G. Stuart, N. Spruston, B. Sakmann, M. Häusser. Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends in Neurosciences*, 20(3):125–131, 1997.
- [K. Takahashi, 2006] K. Takahashi, S. Yamanaka. Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell*, 126(4):663–676, 2006.

- [K. Takahashi, 2007] K. Takahashi, K. Tanabe, M. Ohnuki, M. Narita, T. Ichisaka, K. Tomoda, S. Yamanaka. Induction of pluripotent stem cells from adult human fibroblasts by defined factors. *Cell*, 131(5):861–872, 2007.
- [E.J. Tehovnik, 2006] E.J. Tehovnik, A.S. Tolias, F. Sultan, W.M. Slocum, N.K. Logothetis. Direct and indirect activation of cortical neurons by electrical microstimulation. *Journal of Neurophysiology*, 96(2):512–521, 2006.
- [E.J.H. van Hugte, 2023] E.J.H. van Hugte, E.I. Lewerissa, K.M. Wu, N. Scheefhals, G. Parodi, T.W. van Voorst, S. Puvogel, N. Kogo, J.M. Keller, M. Frega, et al. SCN1A-deficient excitatory neuronal networks display mutation-specific phenotypes. *Brain*, 146(12):5153–5167, 2023.
- [D.A. Wagenaar, 2004] D.A. Wagenaar, J. Pine, S.M. Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *Journal of Neuroscience Methods*, 138(1–2):27–37, 2004.
- [J. Wen, 2022] J. Wen, M. Peitz, O. Brüstle. A defined human-specific platform for modeling neuronal network stimulation in vitro and in silico. *Journal of Neuroscience Methods*, 373:109562, 2022.
- [E.B. Wilson, 1927] E.B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927.
- [G.-S. Yi, 2017] G.-S. Yi, J. Wang, B. Deng, X.-L. Wei. Morphology controls how hippocampal CA1 pyramidal neuron responds to uniform electric fields: a biophysical modeling study. *Scientific Reports*, 7(1):3210, 2017.
- [G. Zanini, 2023] G. Zanini, G. Parodi, M. Chiappalone, S. Martinoia. Investigating the reliability of the evoked response in human iPSCs-derived neuronal networks coupled to micro-electrode arrays. *APL Bioengineering*, 7(4):046121, 2023.
- [Y. Zhang, 2013] Y. Zhang, C. Pak, Y. Han, H. Ahlenius, Z. Zhang, S. Chanda, S. Marro, C. Patzke, C. Acuna, J. Covy, et al. Rapid single-step induction of functional neurons from human pluripotent stem cells. *Neuron*, 78(5):785–798, 2013.