



UNIVERSITY OF GENOA

MASTER'S PROGRAM IN BIOENGINEERING

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

Characterising Network Burst Fragmentation in hiPSC-Derived Cortical Networks of Dravet Syndrome

Jacopo Canestrelli

October 2026

Thesis advisors: Prof. Michela Chiappalone
Prof. Monica Frega

Contents

Abstract	5
Chapter 1 Introduction	6
Chapter 2 State of the Art	10
2.1 Clinical Picture, Mechanisms, and Treatment.....	10
2.2 <i>In Vitro</i> Neural Models	12
2.3 hiPSC-Derived Neurons	13
2.4 Electrophysiological Assessment via Micro-Electrode Arrays (MEAs)	17
2.5 Neuronal Dynamics and Network Bursts.....	20
2.6 Algorithms for Network Burst Detection	25
Chapter 3 Materials and Methods	32
3.1 Cell Cultures, MEA Recordings and Inactive-Electrode Removal.....	33
3.1.1 Cell Lines	33
3.1.2 Differentiation and Plating	35
3.1.3 Dataset.....	36
3.1.4 Device, Acquisition, and Sampling	37
3.2 Data Analysis	38
3.2.1 Spike Detection.....	38
3.2.2 Single-Electrode Burst Detection.....	39
3.2.3 Network Burst Detection: Tampere (V1 and V2)	41
3.2.4 Network Burst Detection: Firing-Rate-Based	44
3.2.5 Fragmentation Detection	46
3.3 Spontaneous Activity Analysis.....	48
3.3.1 Activity Metrics.....	48
3.4 Frequency-Band Validation.....	50
3.4.1 Z-score and Peak Detection.....	50
3.4.2 Cimbalnik Amplitude	51
3.4.3 IN-NB and OUT-NB Event Classification.....	51
3.5 Statistical Procedures.....	52
3.5.1 Cross-Cell-Line Analysis (A1).....	53
3.5.2 Tampere vs Firing-Rate Detector Comparison (A2)	55

3.5.3 V1 vs V2 and Parameter-Set Comparison (A3).....	55
3.5.4 Spectral-Event Class Comparison (A4)	56
3.5.5 Visualisation Conventions	57
Chapter 4 Results	58
4.1 Spontaneous Neuronal Activity.....	58
4.1.1 Raster Plot Overview	58
4.1.2 Activity Metrics.....	60
4.1.3 Statistical Analysis.....	62
4.2 Network Burst Detection: Tampere vs Firing-Rate Detector	64
4.2.1 Raster-Based Comparison	64
4.2.2 Network Burst Rate and Duration.....	66
4.2.3 Inter-Burst Interval Regularity.....	69
4.2.4 Statistical Analysis.....	69
4.3 Fragmentation Analysis	72
4.3.1 Raster Plots with Smoothed Firing Rate.....	74
4.3.2 Fragment Occurrence and Count.....	76
4.3.3 Fragment Overlap between Algorithms	77
4.3.4 Statistical Analysis.....	79
4.4 Algorithm Optimisation and Parameter Sensitivity.....	80
4.4.1 Tampere V1 vs V2 and the ISI-N Threshold Difference	81
4.4.2 Parameter Set Comparison.....	82
4.4.3 Frequency-Band Validation.....	84
4.4.4 IN-NB and OUT-NB Event Classification.....	86
4.4.5 Statistical Analysis.....	88
Chapter 5 Discussion	90
5.1 A Delayed, Hypoactive Network Phenotype in Severe DS	90
5.2 NB Rate and Duration Depend on the Detector	92
5.3 The Fragmentation Phenotype and Its Dependence on the Detector	93
5.4 Robustness to Parameters and Spectral Correlates of Fragments	96
5.5 Convergent Evidence Supports NB Fragmentation as a DS Readout	97
Chapter 6 Conclusions.....	99
6.1 Limitations	100
6.2 Future Perspectives.....	102

Bibliography	106
Appendix A	119

Abstract

Dravet syndrome is a rare and severe genetic epilepsy of childhood, with frequent drug-resistant seizures and developmental delay. Neurons grown in the laboratory from patients' own reprogrammed cells allow the disease to be studied in human brain cells. Grown on devices carrying many small electrodes, they form networks that fire together in short, synchronised events called network bursts. In networks from patients with Dravet syndrome and other neurodevelopmental disorders, these bursts often break up into several consecutive sub-events, a pattern known as burst fragmentation, whose synaptic origin has recently been identified.

This thesis asked whether burst fragmentation is also present in a different model of the disease, networks of mixed excitatory and inhibitory cortical neurons derived from two patients with different clinical severity, and how its measurement depends on the method used to identify the bursts, a step for which no agreed standard exists.

The work was carried out within the Master's programme in Bioengineering of the University of Genoa, with experiments and analyses performed at Tampere University (Finland). Networks from the two patients and two healthy donors were recorded over several weeks of development. A fragmentation analysis was added to the laboratory's pipeline, and all recordings were analysed with two widely used burst-detection approaches, one based on the timing of spikes at each electrode, the other on the firing rate of the whole network.

The network from the more severely affected patient was the least active and showed the most fragmented bursts, especially early in development, in line with the clinical severity of the two patients. The two methods agreed on the average number of fragments per burst in most cell lines, but detected different numbers of bursts and fragments, particularly in the patient-derived networks. An analysis of the raw signal, independent of both methods, showed that fragments coincide with moments of especially strong synchronisation.

These results extend the fragmentation phenotype to a new model of Dravet syndrome and show that the number of fragments counted depends on the detection method: comparing fragmentation across studies, or using it to test drugs on patient-derived neurons, will require detection settings tuned to each dataset and validated by expert inspection of the recordings.

Chapter 1

Introduction

Dravet syndrome (DS) is a rare and severe developmental and epileptic encephalopathy (DEE), first described by Charlotte Dravet in 1978 (Dravet, 2011). Beyond drug-resistant seizures, DS carries a substantial burden of developmental slowing, intellectual disability, and one of the highest syndrome-specific mortality risks in paediatric epilepsy (Cooper et al., 2016; Nabbout et al., 2013), a combination that is described in full, together with its genetic and mechanistic basis, in Section 2.1. Pharmacological treatment remains notoriously difficult. Standard antiseizure medications (ASMs) achieve only partial control, several sodium-channel blockers are contraindicated because they can aggravate seizures, and even the more recently approved adjunctive therapies (stiripentol, cannabidiol, and fenfluramine) tend to reduce seizure frequency without abolishing it (Wirrell et al., 2022). Persistent seizures, together with cognitive and behavioural comorbidities that emerge largely independently of seizure control (Nabbout et al., 2013), are precisely what makes a deeper mechanistic understanding of DS needed. Current disease models, however, capture this human phenotype only partially: rodent models remain the standard for antiseizure drug discovery, but species differences in ion-channel properties, interneuron composition, and network organisation limit their translational value (Löscher, 2011), and findings from SCN1A mouse models have not always translated consistently to human neurons (Liu et al., 2013). This gap motivates the use of neurons derived from human induced pluripotent stem cells (hiPSCs), and specifically of network-level electrophysiological readouts obtained from micro-electrode array (MEA) recordings, as pursued in this thesis.

Approximately 70–85% of DS cases arise from pathogenic variants in SCN1A (Claes et al., 2001; Rampazzo et al., 2021): the gene encoding the alpha subunit of the voltage-gated sodium channel NaV1.1. The reported variants are diverse, including missense, nonsense, frameshift, insertion/deletion and larger structural changes. Truncating variants tend to cause more severe phenotypes (Rhodes et al., 2004; Scheffer & Nabbout, 2019). The traditional explanation for how this causes disease comes down to inhibition. Loss of

NaV1.1 function is thought to lower excitability preferentially in inhibitory GABAergic interneurons, essentially taking the brakes off from cortical circuits and pushing the network toward runaway excitation (Cheah et al., 2012; Han et al., 2012; Ogiwara et al., 2007; Yu et al., 2006). More recent evidence complicates that picture, implicating excitatory neurons and broader network-level changes that reach beyond the interneuron population (Bender et al., 2012; Mzezewa et al., 2025). The interneuron account carries weight for parvalbumin-positive cells, which help coordinate the timing of neuronal activity and generate the network oscillations that support cognition (Ogiwara et al., 2007). The disruption of these cells may therefore feed not only the seizures but also the cognitive and behavioural impairments that define the syndrome. This picture becomes more complex when data from hiPSC models are considered: the functional consequences of *SCN1A* variants can differ markedly between patient-derived lines (Mzezewa et al., 2025; Zayat et al., 2023). Human iPSC technology offers a powerful route to modelling DS in a human genetic background. Patient-derived lines can be differentiated into GABAergic, glutamatergic or heterogeneous cortical cultures, allowing disease-relevant phenotypes to be studied in the very cell types implicated in DS. Heterogeneous cultures are especially valuable here because they contain both excitatory and inhibitory neurons and therefore approximate cortical network organisation more closely than enriched monocultures. They have already revealed altered sodium currents, abnormal spontaneous activity, and variant-dependent differences in network bursting and synchronisation (Mzezewa et al., 2025; Zayat et al., 2023).

To capture and quantify this type of activity, MEAs are essential. They measure spontaneous network activity non-invasively from many electrodes at once and allow repeated measurements as the culture matures, which makes it possible to track how firing, bursting and network synchronisation develop over time (Cerina et al., 2023; Hyvärinen et al., 2019; Kapucu et al., 2022). On MEAs, hiPSC-derived neurons self-organise into network bursts (NBs), transient and synchronous increases in firing across the network, whose properties are widely used for phenotypic characterisation because they correlate with specific disease states (Mossink et al., 2021; Frega et al., 2019; Van Hugte et al., 2023).

In healthy networks the firing rate within an NB rises and then decays monotonically back to baseline, whereas in networks derived from patients it can fluctuate repeatedly within the same event, producing fragmented NBs. Fragmented NBs were the main phenotypic

feature of excitatory networks derived from patients carrying SCN1A variants across the GEFS+-to-DS spectrum (Van Hugte et al., 2023), and comparable within-burst fluctuations have been reported in models of Kleefstra (Frega et al., 2019), Kabuki (Gabriele et al., 2021) and Rett syndrome (Pradeepan et al., 2024); in dissociated rodent cultures similar events had long been described as reverberations or superbursts (Lau & Bi, 2005; Wagenaar et al., 2006). Doorn et al. (2024) combined a biophysically detailed *in silico* model with pharmacological experiments *in vitro* and showed that fragmentation arises from enhanced short-term synaptic depression together with persistent excitation, supported by asynchronous neurotransmitter release and NMDA-receptor currents. Fragmented NBs are therefore a recurring, mechanistically understood network phenotype of patient-derived neurons.

Quantifying fragmentation, however, first requires delimiting the NBs, and there is no agreed gold standard for this step. Some approaches reconstruct NBs from the inter-spike intervals recorded at each electrode, as in the pipeline used at NeuroGroup (Tampere), while others threshold the smoothed firing rate of the whole network, a widely used and computationally light approach, also adopted by Doorn et al. (2024). On the same recording, different methods can return substantially different NB counts and durations, and these differences can propagate to downstream measures such as the number of fragments (Bakkum et al., 2014; Chiappalone et al., 2005; Cotterill et al., 2016).

Two questions therefore motivate this work: whether the fragmented-NB phenotype is also present in heterogeneous hiPSC-derived cortical networks from DS patients (Mzezewa et al., 2025), which contain both excitatory and inhibitory neurons, and how the quantification of fragments depends on the approach used to detect the NBs.

The overall aim of this work is to investigate NB fragmentation, a phenotype repeatedly observed in patient-derived networks, in heterogeneous hiPSC-derived cortical networks from DS patients carrying different SCN1A variants (Mzezewa et al., 2025), and to assess how its quantification depends on the burst-detection method. The main hypothesis is that DS patient-derived networks display increased NB fragmentation compared with control networks, and that the number of fragments detected depends on the detection method employed. The work is structured into three aims, the first two methodological and the third biological:

1. **Aim 1:** to implement and apply a network-burst fragmentation analysis within the NeuroGroup MEA-analysis pipeline, and to test whether the fragmented-burst phenotype is present in Dravet syndrome patient-derived heterogeneous cortical networks;
2. **Aim 2:** to compare the burst-detection algorithms used (the Tampere log-ISI detector, in its V1 and V2 variants, and a widely used firing-rate-based detector), and to quantify the extent to which the electrophysiological metrics and the resulting phenotypic interpretation depend on the detector and its parameters;
3. **Aim 3:** to quantify spontaneous activity at the single-electrode and network levels across development, and to compare Dravet syndrome and control lines.

The remainder of this thesis is organised as follows. Chapter 2 reviews the biological and methodological background needed to interpret the results, covering Dravet syndrome, human stem-cell-derived neuronal models, micro-electrode array recordings, and existing network-burst detection methods. Chapter 3 describes the cell lines, the recordings, and the burst-detection and statistical pipelines used throughout this work. Chapter 4 presents the results: Aim 1 is addressed in Section 4.3, where the network-burst fragmentation analysis is implemented and applied; Aim 2 is addressed in Sections 4.2 and 4.4, where the two detection algorithms are compared and their sensitivity to detection parameters is assessed; and Aim 3 is addressed in Section 4.1, where spontaneous activity is characterised and compared between the Dravet syndrome and control lines. Chapter 5 discusses these results in relation to the existing literature, and Chapter 6 concludes the thesis with a summary of its main findings, its limitations, and its future perspectives.

Chapter 2

State of the Art

2.1 Clinical Picture, Mechanisms, and Treatment

DS is a rare and severe developmental and epileptic encephalopathy (DEE), first described by Charlotte Dravet in 1978 under the original name of severe myoclonic epilepsy of infancy (Dravet, 2011). Its estimated incidence ranges from roughly one in 15,700 to one in 40,000 live births (Wu et al., 2015), which places it among the most clinically severe monogenic epilepsies of childhood. Seizures typically begin within the first year of life in an otherwise healthy infant, often as prolonged, febrile, and frequently hemiclonic episodes, and the disorder then broadens into multiple seizure types, such as myoclonic, focal, atypical absence, together with recurrent episodes of status epilepticus (Dravet, 2011; Scheffer & Nabbout, 2019). What sets DS apart from more benign childhood epilepsies, however, is that the picture reaches well beyond the seizures themselves. From the second year of life onward, affected children show a marked slowing of developmental progress relative to age norms (Nabbout et al., 2013), together with intellectual disability, autistic-like behaviours, and ataxia and other motor impairments (Bender et al., 2012; Scheffer & Nabbout, 2019; Wirrell et al., 2022). DS also carries the highest syndrome-specific risk of sudden unexpected death in epilepsy (SUDEP) documented to date, with deaths occurring mainly in childhood (Cooper et al., 2016). This combination of drug-resistant seizures and a developmental trajectory that keeps worsening is what makes DS so difficult to manage clinically.

At the genetic level, the great majority of cases, approximately 70–85%, arise from pathogenic variants in the *SCN1A* gene, which encodes the alpha subunit of the voltage-gated sodium channel NaV1.1, and most of these occur *de novo* (Claes et al., 2001; Rampazzo et al., 2021). The reported variants are remarkably diverse, spanning missense,

nonsense, frameshift, insertion/deletion, and larger structural changes, with truncating and complete loss-of-function variants tending to produce the more severe phenotypes (Rhodes et al., 2004; Scheffer & Nabbout, 2019). Genotype, on its own, is nonetheless a poor predictor of clinical severity; in fact the same variant can yield very different outcomes across individuals, which points to a substantial contribution from genetic background and modifier genes (Cetica et al., 2017; Hammer et al., 2017).

The principal mechanistic account of DS is the interneuron hypothesis, which centres on impaired inhibition. Because NaV1.1 is expressed preferentially in inhibitory GABAergic interneurons, and most prominently in fast-spiking, parvalbumin-positive cells, loss of NaV1.1 function is thought to compromise the excitability of these neurons selectively, disinhibiting cortical and hippocampal circuits and tipping the network towards hyperexcitability (Cheah et al., 2012; Han et al., 2012; Ogiwara et al., 2007; Yu et al., 2006). This is a compelling model, partly because PV interneurons do more than just suppress firing (they also generate the rhythmic activity that underlies much of cognitive function), so their disruption could plausibly account not only for the seizures but also for the cognitive and behavioural comorbidities of the syndrome. More recent work, however, complicates this view. Evidence from mouse models and from hiPSC-derived neurons suggests that the consequences of SCN1A dysfunction reach beyond interneurons to involve excitatory neurons and broader, network-level reorganisation, with the relative contribution of these elements shifting over the course of development (Bender et al., 2012; Mzezewa et al., 2025). Rather than a clean 'interneuron problem', the picture is more likely one of a broader imbalance between excitation and inhibition that plays out differently at different stages of development.

Clinically, DS remains difficult to treat, and management is built around seizure control rather than cure. First-line therapy relies on broad-spectrum ASMs such as valproate and clobazam, yet control is usually only partial, and DS is regarded as a paradigmatically drug-resistant epilepsy (Wirrell et al., 2022). One point carries particular clinical weight: several sodium-channel-blocking drugs, like carbamazepine, lamotrigine, and phenytoin among them, are contraindicated because by further reducing sodium-channel function they tend to aggravate rather than relieve the seizures, exactly as the underlying NaV1.1 loss-of-function mechanism would predict. The more recently approved adjunctive therapies studied specifically in DS, namely stiripentol, cannabidiol, and fenfluramine, can bring seizure frequency down, but even these typically attenuate the seizures without abolishing

them (Wirrell et al., 2022). Moreover, the cognitive and behavioural impairments emerge largely independently of how well the seizures are controlled (Nabbout et al., 2013).

2.2 *In Vitro* Neural Models

Understanding how neurological disorders emerge at the cellular and network level calls for experimental systems that reproduce key features of human brain function under controlled conditions. The brain itself develops over years, through continuous molecular and structural change: gene expression is modulated, synaptic density is refined, and networks mature in stages (Budday et al., 2015). Studying these processes directly *in vivo* is difficult, partly for ethical reasons, partly because human neural tissue is largely inaccessible. *In vitro* neuronal cultures offer a practical way around this, and they remain among the most widely used tools in neuroscience for probing neuronal development, network dynamics, and disease-related phenotypes (Cerina et al., 2023; Mossink et al., 2021).

The conceptual roots of neuronal cell culture reach back to the early twentieth century, when Harrison kept embryonic nerve fibres alive and observable outside the organism using the hanging-drop technique (Harrison, 1910; Millet & Gillette, 2012). Over the following decades, rodent-derived preparations came to dominate. Neurons taken from embryonic or postnatal mice and rats can be grown as dissociated cells (Brewer et al., 1993) or maintained within organotypic brain slices (Gähwiler et al., 1997; Hutter-Schmid et al., 2015). The dissociated approach, in which cortical or hippocampal tissue is broken down enzymatically and mechanically, gives the experimenter close control over cell density and connectivity, and it supports high-resolution studies of single-cell physiology and network behaviour over long periods (Carter & Shieh, 2015; Gordon et al., 2013). Once plated on micro-electrode arrays (MEAs), these cultures begin to generate spontaneous activity, like spiking, bursting, and synchronised NBs, and such patterns have become functional benchmarks for pharmacological and disease studies (Charlesworth et al., 2015; Chiappalone et al., 2006; van Pelt et al., 2004). Organotypic slices make the opposite trade-off: they preserve the three-dimensional architecture and local circuitry of the source tissue, staying closer to the *in vivo* situation, but they are correspondingly harder to work with (Gähwiler et al., 1997).

Rodent preparations capture only part of what is specific to human neural function. Species differences in neuronal development, gene expression, ion-channel composition, and astrocyte morphology all constrain how well findings transfer to human pathophysiology (Hodge et al., 2019; Oberheim et al., 2009; Zeng et al., 2012). Human astrocytes for example, are considerably larger and more intricate than their rodent counterparts, with more elaborate processes and a greater capacity to shape synaptic transmission (Oberheim et al., 2009). The gap runs down to the transcriptome too, where even matched cortical cell types differ substantially between species (Hodge et al., 2019).

For genetic epilepsies such as DS the problem is sharper still, since pathogenic SCN1A variants act within a human-specific interneuron developmental context that mouse models cannot fully reproduce (Al-Jaberi et al., 2015; Mzezewa et al., 2025; Zayat et al., 2023). There is a practical cost as well, primary rodent cultures depend on fresh tissue for every experiment, which raises both ethical and reproducibility concerns.

In vitro models, however, cannot fully recapitulate the complexity of the *in vivo* brain environment (Cerina et al., 2023; Engle et al., 2018), and thus *in vivo* models for epilepsy research and drug development remain, despite their limited human relevance, the current gold standard (Löscher, 2011). This underscores the need to develop improved human neuron-based *in vitro* models that can better bridge the translational gap.

These limitations are a major reason why the field has pushed toward human-based models in the first place. This study employs two-dimensional (2D) hiPSC-derived neuronal cultures, a prominent approach in this field. They retain the patient's genetic background, allowing for the examination of patient-specific network phenotypes in their native genomic context (Frega et al., 2017; Mossink et al., 2021).

2.3 hiPSC-Derived Neurons

Human iPSCs are produced by reprogramming adult somatic cells, usually skin fibroblasts or peripheral blood mononuclear cells, through forced expression of a defined set of transcription factors i.e., Yamanaka factors (Takahashi & Yamanaka, 2006). Takahashi and Yamanaka (2006) first identified this set in mouse fibroblasts (OCT4, SOX2, KLF4, and c-MYC), and the approach was soon extended to human cells (Takahashi et al., 2007).

Reprogramming returns the cells to a pluripotent state, restoring indefinite self-renewal and the capacity to become almost any cell type, neurons and glia included (Figure 2.1).




Figure 2.1. Routes to hiPSC-derived neuronal models. Somatic cells from a human donor are reprogrammed into hiPSCs and then differentiated into hiPSC-derived neuronal cells, which can be maintained as two-dimensional cultures, three-dimensional cultures, or self-organising brain organoids. The two-dimensional cultures (top branch) are the configuration used in the present work. Adapted from Cerina et al. (2023).

The reprogrammed cells keep the donor's full genomic background. This makes it possible to grow patient-specific neuronal cultures that carry disease-associated variants in their native genetic context (Engle et al., 2018). For monogenic disorders such as DS the payoff is direct: patient-derived lines can be set against healthy controls, or against isogenic controls in which the pathogenic variant has been corrected by genome editing (Mzezewa et al., 2025; Zayat et al., 2023).

The earliest neuronal differentiation protocols leaned on spontaneous embryoid-body formation or on dual SMAD inhibition. Both tended to be slow and both yielded mixed neuronal populations of uneven maturity, although they follow the natural course of neurodevelopment more closely (Zhang et al., 2013). Transcription factor-based direct differentiation changed the picture, bringing real gains in speed, reproducibility, and cellular homogeneity. Forced expression of Neurogenin-2 (Ngn2), a proneural transcription factor, converts hiPSCs into excitatory glutamatergic cortical neurons within two to three weeks, skipping the prolonged progenitor stages altogether (Frega et al., 2017; Zhang et al., 2013). A parallel strategy uses Achaete-scute homolog 1 (Ascl1) together with forskolin to generate inhibitory GABAergic neurons with comparable efficiency (Shi et al., 2016). And where cortical GABAergic interneurons are specifically required, directed-differentiation protocols that recapitulate medial ganglionic eminence patterning offer yet

another route (Kim et al., 2014; Maroof et al., 2013; Nicholas et al., 2013). Together, these approaches allow the cell-type composition of the culture to be controlled, obtaining purely excitatory, purely inhibitory, or a defined mixture of the two.

Neurons rarely thrive on their own. Co-culturing hiPSC-derived neurons with astrocytes is important for survival, functional maturation, and the emergence of robust synaptic activity. Astrocytes supply nutrient support, keep extracellular ions and neurotransmitters in balance, and drive synaptogenesis through both secreted factors and direct cell-cell contact (Degl'Innocenti & Dell'Anno, 2023). Strip away the glial support and hiPSC-derived cultures become less viable, mature more slowly and show weaker network activity. For this reason, most established protocols add a defined proportion of astrocytes, rodent- or hiPSC-derived, to secure consistent network formation (Frega et al., 2017; Mossink et al., 2021). This addition is particularly necessary for the transcription-factor-based forced differentiation protocols described above: because they convert hiPSCs directly into a single targeted neuronal identity, skipping the progenitor stages of normal corticogenesis, they do not generate astrocytes as a matter of course. Dual-SMAD protocols, by more closely following the natural sequence of neurodevelopment, produce astrocytes spontaneously alongside neurons, although consistently only over extended culture periods (Chambers et al., 2009; Shi et al., 2012; Gunhanlar et al., 2018). Forced-differentiation cultures therefore rely on astrocytes added from an external source, commonly of rodent origin (Frega et al., 2017), which is itself a limitation revisited below. Human iPSC-derived networks spontaneously develop their electrophysiological activity that gradually increases over time. The firing is sparse and asynchronous at first; as synaptic connectivity matures, it gives way to organised bursting and then to synchronised network bursting (Heikkilä et al., 2009; Hyvärinen et al., 2019; Mossink et al., 2021). This trajectory echoes, at least in part, the maturation of cortical circuits *in vivo*, and it opens a moving window in which to watch how disease-associated variants reshape network formation over time. Recorded on multielectrode arrays, these networks return reproducible measures of spiking, bursting, and synchronisation that can be set side by side across neuronal cell lines and conditions (Hyvärinen et al., 2019; Kapucu et al., 2022; Mossink et al., 2021).

For DS specifically, hiPSC-derived cultures make an unusually good model. Patient lines carrying pathogenic SCN1A variants can be pushed toward GABAergic interneurons, excitatory neurons, or mixed cortical cultures containing both, which lets the cell-type-

specific and network-level effects of NaV1.1 dysfunction be teased apart. The interneuron hypothesis of DS itself grew out of murine work showing that SCN1A haploinsufficiency hits GABAergic interneurons hardest (Yu et al., 2006; Ogiwara et al., 2007). Patient-derived hiPSC studies have since built on that foundation, reporting altered sodium current densities, elevated spontaneous firing, and disrupted network synchronisation relative to controls, though the magnitude, and even the direction, of these effects shifts from one line or protocol to the next (Higurashi et al., 2013; Jiao et al., 2013; Mzezewa et al., 2025; Sun et al., 2016). That variability reflects a biological reality: the consequences of an SCN1A variant depend on the individual's wider genetic background, and patient-specific hiPSC models are well placed to capture exactly this heterogeneity (Engle et al., 2018; Zayat et al., 2023).

Of course, hiPSC models are not a perfect solution either. Lines differ from one another, and independent differentiation batches differ too, which can undercut reproducibility, especially across laboratories running different protocols (Frega et al., 2017; Mossink et al., 2021). Maturation is a further concern: even after months in culture, hiPSC-derived neurons do not reach the developmental stage of adult human neurons, so their electrophysiology may reflect an immature, almost embryonic-like state (Mossink et al., 2021; Odawara et al., 2016). The rodent-derived astrocytes routinely used in co-culture add one more warning, convenient and well-characterised as they are, they introduce a cross-species element that may itself influence neuronal function and network dynamics (Degl'Innocenti & Dell'Anno, 2023).

Despite all of that, hiPSC-derived networks remain a valuable option for studying human-specific pathophysiology *in vitro*. They combine genetic fidelity with experimental tractability and functional relevance, a mix that suits them especially well to genetic epilepsies like DS, where differences in network activity between patients can be traced back to specific genetic variants and followed longitudinally with MEA recordings.

2.4 Electrophysiological Assessment via Micro-Electrode Arrays (MEAs)

Characterising *in vitro* networks functionally means recording from many neurons at once, with enough temporal resolution to catch fast spikes and enough spatial coverage to read network-level dynamics. MEA technology has become the standard answer to this need: it records non-invasively from many sites at once, tolerates long-term culture, and lets the same culture be revisited at successive developmental time points (Thomas et al., 1972). The idea behind MEA technology goes back to the early 1970s, when Thomas and colleagues showed that cultured cells could be recorded extracellularly on microfabricated electrode substrates (Thomas et al., 1972). Long-term recordings from dissociated cortical cultures soon proved practicable, opening the way to the systematic study of spontaneous network activity *in vitro*. The technology has come a long way since, with steady gains in electrode density, signal-to-noise ratio, substrate biocompatibility, and acquisition software. The result is the wide range of commercial platforms available today (Cerina et al., 2023; Jones et al., 2011).

Patch-clamp electrophysiology, by comparison, delivers single-cell resolution but at low throughput, and it cannot follow a culture over time. MEAs give up some of that resolution in exchange for the ability to watch spontaneous and evoked activity simultaneously across many electrodes spread over the culture. This trade-off is what makes MEAs valuable, you lose single-cell precision but gain the ability to watch how the whole network behaves, which is what matters when studying network-level activity and synchronisation (Cerina et al., 2023; Nam & Wheeler, 2011; Novellino et al., 2011).

A basic, planar single-well MEA usually carries multiple electrodes, such as 60 in an 8×8 grid in a model by Multichannel Systems, each electrode roughly 30 μm across and spaced 200 μm apart, as in the 60MEA200/30iR-Ti array (Multi Channel Systems MCS GmbH, Reutlingen, Germany). Devices of this kind have been used extensively to characterise spontaneous activity in dissociated cortical and hippocampal cultures, and they reinforce much of what is known about how network level activity develops, how pharmacology modulates network dynamics, and how synaptic plasticity manifests *in vitro* (Chiappalone et al., 2006; van Pelt et al., 2004; Wagenaar et al., 2006). Their drawback is plain enough:

one network per device. That becomes awkward the moment an experiment calls for several independent conditions or cell lines recorded in parallel. Multi-well MEA (mwMEA) platforms solve the problem by packing several recording wells, each with its own electrode array, into a single device that still fits standard incubator conditions. For disease-modelling work, where multiple cell lines, treatments, or time points have to be compared, this parallelism is exactly what is needed (Hyvärinen et al., 2019; Mossink et al., 2021).

Beyond these planar formats, the technology now spans several specialised architectures, summarised in Figure 2.2. Complementary metal-oxide semiconductor (CMOS) -based high-density MEAs (HD-MEAs) integrate thousands of tightly spaced micro-electrodes on a single chip, approaching near-cellular spatial resolution and allowing activity, and its propagation across a network, to be mapped at a far finer scale than conventional arrays. MEAs with non-planar, protruding micro-electrodes are designed to tighten the contact between electrode and cell membrane, raising the recorded signal amplitude and partly bridging the gap between extracellular and intracellular recording. 3D MEAs distribute electrodes through the depth of a culture rather than only across its base, matching the geometry of engineered 3D constructs and organoids in which activity is not confined to a single plane. Each of these designs trades spatial resolution, throughput, and experimental convenience against one another in its own way (Cerina et al., 2023; Jones et al., 2011).




Figure 2.2. Main families of micro-electrode array (MEA) devices: single-well MEAs (swMEAs), multi-well MEAs (mwMEAs), CMOS-based high-density MEAs (HD-MEAs), MEAs with non-planar micro-electrodes, and three-dimensional MEAs (3D MEAs). Adapted from Cerina et al. (2023).

MEA signal acquisition comprises three stages: recording, conditioning and digitisation of the extracellular electrical activity. This process relies on several components beyond the

MEA chip itself, with the final signal quality depending on the electrode design, the biological and instrumental noise, the noise-reduction strategy and the stability of the recording environment. The MEA plate is connected to a headstage that provides the first amplification stage (and stimulation, where required), placed close to the electrodes to limit interference. The extracellular potentials recorded at the electrodes are small, typically from about 10 μ V to a few hundred microvolts, and are amplified by a factor on the order of 10^4 – 10^5 . The amplified analog signal is then band-pass filtered to remove slow baseline drift and high-frequency interference, separating the higher-frequency spiking band from the lower-frequency local-field-potential band. Each channel is subsequently digitised by a high-resolution analog-to-digital converter, with a sampling rate high enough to resolve fast spikes while still capturing slower potentials (platform-dependent rates of 10–25 kHz are typical). The digitised, multiplexed data are then streamed to dedicated acquisition software for real-time visualisation and storage of the raw voltage traces. A central objective throughout is to keep the signal-to-noise ratio acceptable (commonly a target of at least five-fold): biological noise arises from distant action potentials, synaptic and subthreshold activity and local field potentials, whereas instrumental noise stems from the electronics, electromagnetic coupling, thermal and 1/f electrode–electrolyte processes, and mains interference at 50/60 Hz; these are mitigated by proper grounding, Faraday shielding and optimisation of the electrode impedance. Once digitised, an MEA recording yields a record of extracellular action potentials (APs) at each active electrode. Spikes are picked out by threshold-based detection on the filtered voltage trace, then grouped, first into single-electrode bursts, then into NBs, according to their timing and their co-occurrence across electrodes. From these arise the standard metrics: mean firing rate, burst rate, burst duration, NB rate, and inter-burst interval statistics, which together give a multidimensional picture of network activity sensitive to both developmental state and disease-related change (Cerina et al., 2023; Mossink et al., 2021). In DS, several of these metrics separate patient lines from controls in ways consistent with the network dysfunction expected from NaV1.1 haploinsufficiency, which supports MEA-based phenotyping as a functional readout in translational research (Mzezewa et al., 2025). The particular spike-, burst-, and NB-detection algorithms used here are described in full in Chapter 3.

2.5 Neuronal Dynamics and Network

Bursts

With the extracellular signals in hand, the spontaneous activity of an *in vitro* network can be described at several levels at once, from the firing of single neurons up to the coordinated behaviour of large populations. Read together, these levels give a functional summary of the culture's developmental state and connectivity and, indirectly, of certain aspects of its excitation–inhibition balance, but the E/I balance cannot be inferred from the MEA signal alone. Rather than a primary route, MEA electrophysiology is one of several complementary approaches for detecting disease-associated phenotypes in hiPSC-derived neuronal models, used alongside patch-clamp electrophysiology, morphological analysis and transcriptomics (Cerina et al., 2023; Mossink et al., 2021). At the single-cell level, neurons are excitable cells characterised by the presence of several types of voltage-dependent ion channels, primarily Na^+ , K^+ , and Ca^{2+} channels, which allow for rapid and transient changes in the membrane potential ($V_m \simeq -70 \text{ mV}$) (Chen & Lui, 2019; Hodgkin & Huxley, 1952). When a depolarising stimulus exceeds the activation threshold, it triggers the emission of an action potential (AP), followed by a refractory period during which the neuron is less excitable, due to a temporary increase in the threshold necessary to generate a second AP, a consequence of voltage-gated sodium channels remaining inactivated (Bean, 2007; Hodgkin & Huxley, 1952) (Figure 2.3).

Neuronal Action Potential




Figure 2.3. The neuronal AP. Phases of the AP (resting potential, threshold, depolarisation, repolarisation, refractory period, return to resting state) on the membrane-potential trace and at the level of voltage-gated Na⁺/K⁺ channels and the Na⁺/K⁺ pump. Adapted from Chen & Lui (2019).

The AP, or spike, is the most basic unit of activity recorded on an MEA. On an extracellular electrode it appears as a brief voltage deflection at a single site, usually one to three milliseconds long, and is detected by applying an amplitude threshold to the filtered signal (Maccione et al., 2009). Spikes and the bursting activity built from them, occupy the high-frequency component of the MEA signal, approximately 300–3000 Hz, which is isolated by band-pass filtering; the lower-frequency range (below ~300 Hz) carries the local field potentials, which are usually filtered out when the analysis focuses on spike-based metrics. Averaging spikes per second across all active electrodes gives the mean firing rate (MFR), a simple index of network excitability, and one of the most commonly reported measures of *in vitro* electrophysiology (Charlesworth et al., 2015; Kapucu et al., 2022; Novellino et al., 2011). In hiPSC-derived cultures the MFR climbs over weeks as neurons mature and

wire up, and it separates patient from control lines in disease models (Hyvärinen et al., 2019; Mossink et al., 2021; Mzezewa et al., 2025).

A burst is a rapid run of APs packed into a short window and set off from neighbouring events by longer silent stretches, the inter-burst intervals (IBIs). At a single electrode, bursting marks periods of heightened local excitability, and it reflects an interplay between fast sodium-driven spiking and the slower calcium-mediated depolarisations that keep firing going (Bologna et al., 2010; Chiappalone et al., 2006). Single-electrode bursts are usually detected through the inter-spike interval (ISI) distribution, which in a bursting neuron is typically bimodal: a first peak at short, intra-burst intervals and a second at longer, inter-burst ones. The valley between the two modes sets the threshold that divides spikes inside bursts from those outside (Pasquale et al., 2010). The analysis then yields several metrics such as burst rate, burst duration, and the percentage of random spikes (those belonging to no detected burst). A typical example of a single spike and burst from raw MEA recording is presented in Figure 2.4.




Figure 2.4. Hierarchical organisation of spontaneous MEA activity. Left: raw multi-electrode recording. Right: a magnified extracellular trace showing an isolated AP (spike, green) and a burst (red), a rapid run of spikes flanked by longer quiescent intervals. Spikes, single-electrode bursts, and NBs constitute the successive levels of description used throughout this work. Adapted from Cerina et al. (2023).

As connectivity grows, bursting at individual electrodes begins to line up across many of them, producing NBs: large synchronisation events in which a sizeable share of the active electrodes fire together or in quick succession (Eytan & Marom, 2006; van Pelt et al., 2004;

Wagenaar et al., 2006). NBs are a signature of functionally connected *in vitro* cultures, driven by recurrent excitatory connectivity and shaped by inhibitory interneurons. They are usually found by looking for windows of globally raised firing that involve some minimum fraction of the active electrodes, that can differ between implementations, commonly somewhere between 20% and 50% (e.g., Chiappalone et al., 2005; Mossink et al., 2021). The events are then summarised by metrics such as NB rate (NBR), NB duration (NBD), the coefficient of variation of the inter-burst interval (CV IBI), and the local variation (LvR), which tell us how regular the network-wide synchronisations are (Eisenman et al., 2015; Shinomoto et al., 2009; Kapucu et al., 2012). Examples of these metrics are shown in Figure 2.5.

High CV(IBM) values point to irregular bursting, a possible sign of immature connectivity or a skewed E/I balance; low values go with the stereotyped, almost clock-like bursting of highly synchronised networks (Eisenman et al., 2015).

One related phenomenon has drawn growing interest in disease modelling: NB fragmentation. Some NBs are not single, monolithic activations. Instead, the population firing rate within them rises and falls several times, producing distinct peaks separated by short dips in activity. Each such peak is called a fragment, and an NB made of several fragments is a fragmented NB; in practice, fragments are identified by peak detection on the smoothed population firing rate inside each detected NB, under minimum-height and minimum-prominence criteria, and their number per NB is counted (Doorn et al., 2024). The phenomenon is not new and has been given different names, including reverberations, minibursts, high-frequency bursts and superbursts (Doorn et al., 2024). In dissociated rodent cultures, reverberating activity following a burst was described long before hiPSC models became available (Lau & Bi, 2005; Wagenaar et al., 2006).

In hiPSC-derived networks, fragmented NBs emerged as the main phenotypic feature of excitatory networks from patients carrying SCN1A variants across the GEFS+-to-DS spectrum (Van Hugte et al., 2023), and they have been reported in models of other neurodevelopmental disorders, including Kleefstra, Kabuki and Rett syndrome (Frega et al., 2019; Gabriele et al., 2021; Pradeepan et al., 2024). In the SCN1A networks, antiseizure drugs that were effective in the donors reduced the number of fragments per NB, whereas drugs that worsened their clinical picture increased it (Van Hugte et al., 2023). Doorn and colleagues (2024) then clarified the mechanism: combining a biophysically detailed *in silico* model with pharmacological validation *in vitro*, they showed that fragmentation

requires substantial short-term synaptic depression, and that enhancing it with Dynasore was sufficient to fragment the NBs of healthy networks, while persistent excitation, sustained by asynchronous neurotransmitter release and NMDA-receptor currents, supports the generation of multiple fragments. Fragmentation therefore captures an aspect of network dysfunction that firing-rate or burst-rate measures alone can miss (Doorn et al., 2024; Mzezewa et al., 2025). A fragmented burst in a raw neuronal trace is indicated in Figure 2.6.




Figure 2.5. Activity metrics computed from MEA raster plots. Single-electrode metrics (left): mean firing rate (MFR) and inter-spike interval (ISI) at the spike level; mean burst rate, mean burst duration, mean ISI within a burst, and percentage of spikes within bursts at the single-electrode burst level; the percentage of random spikes denotes spikes outside any burst. Network-level metrics (right): NBR, NBD, mean network inter-burst interval (NIBI), and the coefficient of variation of the NIBI (CV_{NIBI}). The scheme is adapted to the present work, in which the network inter-burst interval and its coefficient of variation are written IBI and $CV(IBM)$ (Section 3.3.1) and the fragments-per-NB ratio (FragPerNB) is added as a further network-level descriptor. Adapted from Cerina et al. (2023).




Figure 2.6. NB fragmentation in DS. Representative extracellular trace (top) and smoothed population firing rate (bottom) for a healthy control hiPSC derived excitatory network (left) and a DS network (right). In the control the NB forms a single firing-rate peak, whereas in the DS culture the same event breaks into several distinct peaks (fragments, markers), illustrating the higher NB fragmentation reported for patient-derived cultures. Adapted from Doorn et al. (2024).

2.6 Algorithms for Network Burst

Detection

Detecting NBs from MEA recordings is a critical step, and far from a trivial one. A single-electrode burst (SEB) concerns one electrode's signal; a NB, by contrast, has to be picked out as a coordinated event spanning many electrodes at once, and exactly where its boundaries are placed propagates into every downstream metric analysed. Decades of work have not settled on a consensus standard. Different algorithms can return markedly different results on the very same recording, because they rest on different assumptions about signal shape, electrode participation, and where an event begins and ends (Bakkum et al., 2014; Chiappalone et al., 2005; Cotterill et al., 2016).

It is therefore important to specify precisely how each algorithm defines and detects bursts, since the choice of detector can systematically influence every metric derived from it.

An example of NB detection is shown in Figure 2.7.




Figure 2.7. Representative raster plot of spontaneous network activity recorded on a micro-electrode array. Each row corresponds to one electrode and each tick to a detected spike; the horizontal axis spans 150 s of recording. The vertical lines delimit the automatically detected network bursts (NBs): the green line marks the onset and the red line the offset of each event. Within these intervals firing is strongly synchronised across channels (dense black bands), whereas the inter-burst periods show only sparse, asynchronous spiking. Five NBs are visible over the displayed window. Adapted from Vinogradov et al. (2024).

The earliest and most widely used approaches apply a fixed amplitude threshold on the global firing-rate signal. Here a NB is any stretch in which the population firing rate, the summed spike count across all active electrodes in discrete time bins, rises above a preset level, usually expressed as a multiple of the recording's mean or baseline rate. Chiappalone and colleagues put forward an influential version of this idea, defining NBs as events where the global spike rate clears a fixed threshold and enough active electrodes take part; they used these criteria to track how spontaneous activity develops in dissociated cortical cultures over weeks *in vitro* (Chiappalone et al., 2005, 2006). Van Pelt and colleagues took a similar threshold-based route on MEA spike trains, charting the long-term dynamics of network bursting and showing systematic shifts in burst rate and duration across development (van Pelt et al., 2004). The downside is that if the baseline firing rate shifts between recordings, as cultures age or drugs are added, the fixed threshold stops making sense and often this type of detector needs manual tuning to stay consistent across conditions.

ISI-based methods were developed precisely to get around that parameter sensitivity, drawing burst boundaries from the statistical structure of the spike train rather than from absolute firing-rate values (Gourévitch & Eggermont, 2007; Legéndy & Salcman, 1985). The self-adapting algorithm of Pasquale and colleagues is the foundational example. It exploits the bimodal ISI distribution of bursting neurons by building a logarithmic ISI histogram (log-ISI) and then selecting, automatically, the threshold that best separates intra-burst from inter-burst intervals (Pasquale et al., 2010). Single-electrode bursts identified this way are then pooled, and NBs are defined from their cumulative occurrence across electrodes. This log-ISI logic is the conceptual basis of the NeuroGroup (Tampere) NB detector used in the present thesis. The detector comes in two variants, V1 and V2, that share the same N-th order log-ISI histogram but differ in how the threshold between intra- and inter-burst intervals is chosen; V2 is built to be more robust to ripples introduced by histogram smoothing. The full implementation, including the threshold-selection rules and parameter values, is set out in Chapter 3. The principle of log-ISI based burst-detection method is summarised in Figure 2.8.




Figure 2.8. Principle of the log-ISI burst-detection method. The inter-spike intervals of a spike train (single-electrode or pooled) are binned on a logarithmic scale; in a bursting train the histogram is bimodal, with a short intra-burst and a long inter-burst mode, and the threshold separating intra- from inter-burst intervals is placed at the valley between them. Runs of intervals below the threshold (with at least a minimum number of spikes) are accepted as bursts. This single-electrode principle is the basis of the network-level NeuroGroup (Tampere) detector used here (Chapter 3). After Pasquale et al. (2010).

Bakkum and colleagues added a complementary advance to ISI-based detection: a generalised framework built on the N-th ISI (ISI-N), in which the threshold comes from the distribution of N-th order ISIs pooled across all electrodes (Bakkum et al., 2014, Figure 2.9). It asks the user to set just one parameter ‘N’, the minimum number of spikes that defines the smallest burst worth counting, and it proved more robust and more sensitive than earlier algorithms across a wide range of activity levels. Because ISI-based methods, log-ISI and ISI-N alike, adapt to the statistics of each individual recording, they are particularly well suited to longitudinal studies in which firing rates change a lot from one developmental time point to the next.




Figure 2.9. Principle of the ISI-N (N-th order ISI) burst criterion. Rather than consecutive inter-spike intervals, the interval between each spike and the N-th following spike is computed; small ISI-N values correspond to locally high firing rates. A threshold on the ISI-N distribution delimits NBs as runs whose ISI-N falls below it, with N as the single free parameter. Adapted from Bakkum et al. (2014).

A third family, widely used for hiPSC-derived networks, works directly on the smoothed population firing rate, treating NBs as prominent peaks in the envelope of global activity. The recipe is to convolve the binned network spike rate with a smoothing kernel (often a Gaussian) and to flag events that cross thresholds defined relative to the maximum or mean of the smoothed signal. Where ISI-based methods read spike times electrode by electrode, firing-rate-based methods look at the collective dynamics of the whole population, and they are computationally light enough to be applied identically to recorded and simulated data. Doorn et al. (2024), for example, adopted it for both their in vitro recordings and their in silico simulations: the smoothed firing rate is computed over all active electrodes (MFR above 0.02 Hz), NB onset and offset are set by two amplitude thresholds (one-quarter and one-hundredth of the maximum firing rate), the rate must stay above the relevant threshold for at least 50 ms, and at least 30% of the active electrodes must take part; within each NB, fragments are the peaks of the smoothed rate whose height reaches one-sixteenth and whose prominence reaches one-tenth of the maximum firing rate. Because these thresholds are expressed relative to the data, their values are tuned to the recordings for which they were chosen. This smoothed firing-rate principle is illustrated in Figure 2.10.




Figure 2.10. Principle of smoothed firing-rate NB detection. Spikes pooled across electrodes are binned into a population firing rate and smoothed with a Gaussian kernel; NBs are the intervals where the smoothed rate exceeds a threshold, and sub-peaks within a burst are marked as fragments. After Doorn et al. (2024).

A separate approach that ignores spikes entirely and instead looks at the local field potential (LFP), that is the low-frequency part of the extracellular voltage, typically 1–300 Hz, for information that spike-based metrics alone do not capture. A genuine NB is accompanied by a transient, broadband increase in the power of the extracellular field, with prominent contributions in the higher oscillatory bands, beta, gamma and high-gamma. These spectral changes arise because the coordinated synaptic and spiking activity of the co-active neuronal population summates into larger extracellular field fluctuations, so the band-specific power increases index the synchrony of the population generating the event (Buzsáki & da Silva, 2012; Buzsáki et al., 2012). LFP-based methods are not used as primary burst detectors here. Frequency-band analysis is used, as an independent check: comparing band power inside and outside detected NBs helps establish that the events flagged by the spike-based algorithms are biologically plausible (Pré et al., 2026).

None of these approaches is clearly best; each makes different compromises between sensitivity, interpretability, and robustness to messy real-world recordings. ISI-based methods generally shrug off global firing-rate differences, but the shape of the ISI distribution can throw them when burst structure is complex or multimodal. Firing-rate-

based methods are good at catching population-level events, yet more exposed to noise and to artefactual swings in the smoothed signal. Fixed-threshold methods win on simplicity and transparency, and lose on the need for careful, condition-by-condition parameter tuning. Table 2.1 lays out the principal NB detection algorithmic families and their main features.

Family	Representative method	Input signal	Free parameters	Strengths	Limitations
Log-ISI histogram (NeuroGroup detector, V1 & V2)	Pasquale et al. (2010); NeuroGroup (Tampere) implementation. V1 : first local minimum of the log-ISI histogram after the global maximum. V2 : global minimum within the valley between the two leading prominent peaks (height- and prominence-filtered).	Per-electrode log-ISI histogram; NB aggregated from SEBs and refined by channel-participation criteria	ISI-N histogram parameters (N, span, bins/decade, smoother); minimum-height/prominence (V2); merging, minimum duration, minimum channels, minimum dominant channels, minimum channel contribution	Self-adapting to recording statistics; explicit two-stage logic (single-electrode → network); V2 robust to smoothing ripples	Requires reasonably bimodal log-ISI distribution; V1 can latch onto spurious local minima; V2 fails when fewer than two prominent peaks are present
Smoothed firing-rate with locally-adaptive thresholds	Widely used in hiPSC-MEA studies; parameters used here follow Doorn et al. (2024)	Smoothed population firing rate (Gaussian-smoothed across all active electrodes)	Start/stop thresholds relative to the maximum firing rate (1/4 and 1/100 in Doorn et al., 2024); 50 ms minimum duration; 30% active-electrode participation; peak-detection parameters for fragments	Captures population-level dynamics; computationally light; fragment detection runs on the same smoothed signal	Sensitive to smoothing kernel choice; thresholds normalised to the maximum firing rate may be affected by outliers and need tuning to each dataset
Fixed amplitude threshold on global firing rate	Chiappalone et al. (2006); van Pelt et al. (2004)	Binned population firing rate	Absolute or baseline-scaled rate threshold; minimum fraction of active electrodes	Conceptually simple and transparent; computationally inexpensive	Sensitive to global firing-rate variability across recordings; often requires manual tuning across conditions
Pooled ISI-N	Bakkum et al. (2014)	Network-wide pooled spike train	Single parameter: minimum number of spikes N defining the smallest burst	Single free parameter; not biased towards large bursts; precise boundary assignment	Choice of N influences sensitivity; pooled analysis discards per-electrode structure
LFP / frequency-band analysis	Buzsáki et al. (2012)	Local field potential (1-300 Hz); beta, gamma, high-gamma band power	Frequency bands; spectral estimation window	Independent confirmation of population events; sensitive to oscillatory phenomena not visible in spike rate	Not used as a primary burst detector in this work; serves as an orthogonal validation

Table 2.1. Principal families of NB detection algorithms. The first two rows correspond to the algorithms compared in the present study; the remaining rows summarise representative alternatives discussed in the text.

Chapter 3

Materials and Methods

An overview of the full analytical pipeline is given in Figure 3.1.

Starting from the raw extracellular signals recorded on the MEA, spikes were first detected on each electrode and inactive electrodes were removed; the retained spike trains were then used to detect SEBs and, by pooling activity across electrodes, NBs. NBs were identified with three detectors applied in parallel — the NeuroGroup (Tampere) log-ISI-N detector in two threshold variants (V1 and V2) and a widely used smoothed firing-rate detector (FR detector) with parameters adapted from Doorn et al. (2024) — and each accepted NB was further characterised by an integrated fragmentation analysis. From the spike trains, the SEBs and the NBs, a set of recording-level activity metrics was then computed for every well and time point. These metrics were compared statistically across cell lines and detectors using linear mixed-effects models (Analyses A1–A3), and the NB and fragment detection was cross-validated against an independent spectral high-frequency-event analysis (Cimbalnik framework). Each stage is described in detail in the sections that follow.




Figure 3.1. Overview of the analytical pipeline. From raw MEA signals to statistical comparison.

3.1 Cell Cultures, MEA Recordings and Inactive-Electrode Removal

3.1.1 Cell Lines

In this study, four hiPSC lines were analysed: two patient-derived lines from individuals diagnosed with DS (DD1C and DD5A) and two healthy control lines (TUBA and C10902). All were obtained from consenting donors, and their use in neuronal research is covered by

an ethical statement from the Ethics Committee of the Expert Responsibility Area of Tampere University Hospital (R20159).

The two healthy control lines were C10902 (10902.EURCCs), from a 78-year-old female donor, and TUBA (AICS-0012), commercially available from the Coriell Institute (USA) and generated from a 30-year-old male donor; TUBA carries a mono-allelic mEGFP knock-in at the endogenous TUBA1B locus. The two patient-derived lines, DD1C and DD5A, came from female donors diagnosed with DS and carry distinct pathogenic variants in SCN1A. DD1C, from a 17-year-old donor, carries a *de novo* frameshift variant (c.5502_5509dup; p.Pro1837Argfs*24), predicted to result in complete loss of function of the encoded NaV1.1 channel. DD5A, from a 9-year-old donor, instead carries a pathogenic missense variant (c.651C>G; p.Thr217Arg) located in domain I, segment 4 of NaV1.1 (the region encoding the voltage sensor of the channel) and is predicted to impair channel gating (Schuster et al., 2019; Mzezewa et al., 2025). The four lines and their relevant characteristics are summarised in Table 3.1.

Table 3.1. Summary of the four hiPSC lines used in this work, with donor information, group assignment, SCN1A variant in HGVS notation, predicted functional effect, and source. *TUBA1B (AICS-0012) is a commercial reporter line obtained from the Coriell Institute (Camden, NJ, USA).

Cell line	Donor sex / age	Group	SCN1A variant (HGVS)	Predicted effect	Source / reference
DD1C	F, 17 y	DS (severe)	<i>De novo</i> frameshift: c.5502_5509dupGCTTGAAC; p.Pro1837Argfs*24	Complete loss of function of the NaV1.1 channel	Schuster et al. (2019); Mzezewa et al. (2025)
DD5A	F, 9 y	DS (mild)	Missense: c.651C>G; p.Thr217Arg (domain I, segment 4: voltage sensor of NaV1.1)	Predicted to impair channel gating	Schuster et al. (2019); Mzezewa et al. (2025)
TUBA1B (AICS-0012)*	M, 30 y	Healthy control	No SCN1A variant; mono-allelic mEGFP knock-in at TUBA1B (WTC parental)	Reference / fluorescent control line	Kreitzer et al. (2013); Coriell Institute (AICS-0012)
10902.EURCCs (C10902)	F, 78 y	Healthy control	No SCN1A variant	Reference control line	Kiamehr et al. (2019); Mzezewa et al. (2025)

3.1.2 Differentiation and Plating

The four hiPSC lines were differentiated into cortical neurons following a previously published protocol (Hyvärinen et al., 2019; Mzezewa et al., 2025). In brief, cells were taken through staged neural induction, proliferation, and maturation over 32 days *in vitro* (DIV), then enzymatically dissociated and plated onto MEA devices in neural maturation medium (NMM) for electrophysiological recording. Throughout this work DIV refers to days *in vitro*; cells were therefore plated on MEA at DIV 32, which corresponds to day 1 on MEA (DOM 1).

Before plating, CytoView MEA 48-well plates (Axion Biosystems, Atlanta, GA, USA) were coated with 0.1% poly(ethyleneimine) (PEI; Sigma-Aldrich, Stockholm, Sweden) in 0.1 M borate buffer and then with 50 µg/ml human recombinant laminin (LN521; BioLamina, Sundbyberg, Sweden). Neurons were plated at 80,000 cells per well, corresponding to roughly 635,000 cells/cm². On DOM 1 the medium was refreshed with a 1:1 mixture of NMM and BrainPhys™ maturation medium (BMM; STEMCELL Technologies, Vancouver, BC, Canada); from DOM 2 onwards the cultures were kept exclusively in BMM. Medium was changed three times per week throughout the culture

period, and cultures were maintained at 37 °C in a humidified atmosphere of 5% CO₂ (Hyvärinen et al., 2019; Mzezewa et al., 2025).

3.1.3 Dataset

Spontaneous electrophysiological activity was recorded at five developmental time points: DIV 53, 60, 67, 83, and 90. Two of these are nominal labels: the time points reported throughout as DIV 60 and DIV 90 correspond to recordings acquired at DIV 62 and DIV 88 respectively; the trajectory figures are labelled with the nominal values, whereas the raster figures are annotated with the acquisition DIV. The first three (DIV 53, 60, and 67) sample the early-to-mid network consolidation phase, while DIV 83 and 90 sample the later functional-stabilisation phase.

The dataset comprised 136 wells distributed across 5 independent MEA plates, with each well contributing a recording at every time point, for a total of 680 individual recordings. The two patient lines, DD1C and DD5A, were each recorded across three independent experiments (DD1C: 16, 16, and 12 wells; DD5A: 12, 12, and 16 wells; 44 and 40 wells in total). The two healthy control lines, TUBA and C10902, were each recorded across two experiments i.e., two MEA plates (TUBA: 8 and 16 wells; C10902: 12 and 16 wells; 24 and 28 wells in total). The five plates were shared across cell lines rather than dedicated to single lines (plate N2894 carried DD1C and TUBA wells; N2897 carried DD5A; N2905 carried DD5A and TUBA; N2912 carried DD1C and C10902; N2942 carried DD1C, DD5A and C10902), which is why plate identity is modelled as a random intercept in the statistical analysis (Section 3.5). The complete breakdown is given in Table 3.2.

Cell line	Group	Plates (n)	Wells per plate	Total wells	Recordings (wells × 5 DIVs)
DD1C	DS (severe)	3	16, 16, 12	44	220
DD5A	DS (mild)	3	12, 12, 16	40	200
TUBA1B	Healthy control	2	8, 16	24	120
10902.EURCCs	Healthy control	2	12, 16	28	140
Total		5		136	680

Table 3.2. Breakdown of the experimental dataset by cell line, number of independent plates, wells per plate, total wells, and total recordings (wells $\times$ 5 DIV time points). The per-line plate counts sum to ten because three of the five plates carried wells of more than one cell line; the total row therefore reports the five independent plates.

3.1.4 Device, Acquisition, and Sampling

Spontaneous activity was recorded on the Maestro multi-well MEA platform (Axion Biosystems, Atlanta, GA, USA). Recordings used CytoView MEA 48-well plates, in which each well has 16 extracellular electrodes in a 4×4 grid, with an electrode diameter of $30 \mu\text{m}$ and an inter-electrode spacing of $350 \mu\text{m}$ (Figure 3.2).




Figure 3.2. The MEA system used in this work. CytoView MEA 48-well plate (Axion Biosystems): wells are arranged in a 6×8 layout (rows A-F, columns 1-8), and each well contains 16 extracellular electrodes in a 4×4 grid (enlargement), labelled 11-44 (rows-columns). Spontaneous activity was recorded on the Maestro multi-well platform at 12.5 kHz per channel. Adapted from Kapucu et al. (2022).

Extracellular signals were sampled continuously at 12.5 kHz per channel, enough temporal resolution to resolve both individual APs and population-level synchronisation events. Acquisition and real-time visualisation used AxIS Navigator software (Axion Biosystems), following previously published protocols with the same platform and cell lines (Hyvärinen et al., 2019; Mzezewa et al., 2025). Each recording lasted 600 s (10 min), a standardised window for comparing spontaneous network dynamics across cell lines and developmental time points.

3.2 Data Analysis

3.2.1 Spike Detection

Spike detection relied on the in-house MATLAB pipeline described in Hyvärinen et al. (2019), combining SWTTEO with ABS (Lieb et al., 2017; Mayer et al., 2018). SWTTEO decomposes the signal with a stationary wavelet transform, which resamples the wavelet filters at each scale rather than down-sampling the coefficients as in the standard discrete wavelet transform, thereby preserving a time-invariant representation suited to non-stationary MEA recordings; the Teager energy operator is then applied to the resulting sub-band coefficients, nonlinearly emphasising the energy concentrated around action-potential waveforms, and each sub-band is thresholded with an SD-based rule before the sub-bands are recombined into a single detection trace (Lieb et al., 2017; Mayer et al., 2018). After ABS-based detection, SWTTEO was configured to return the same number of spikes as ABS. Amplitude thresholding used a threshold of $4.5 \times$ the estimated noise standard deviation (Quiroga et al., 2004; Mzezewa et al., 2025), and a spike was accepted only if both algorithms detected it. Requiring agreement between the two algorithms served as a cross-validation step: because ABS and SWTTEO rely on different signal properties, raw amplitude versus wavelet-domain energy, an event flagged by only one of them more likely reflects noise or a threshold artefact than a genuine action potential. Constraining SWTTEO to return the same spike count as ABS therefore did not simply equalise two arbitrary numbers, but forced the two independent criteria to converge on the same underlying set of events, so that only spikes consistent with both detection principles were retained for the downstream burst and network-burst analyses. Electrodes whose MFR exceeded 0.2 Hz over the 10-minute recording were defined as active, and only active electrodes were retained for analysis; channels below this threshold were treated as inactive and excluded. The spike-detection pipeline is shown in Figure 3.3.




Figure 3.3. Spike-detection pipeline. The extracellular signal of each electrode (12.5 kHz) is processed in parallel by the ABS detector (amplitude threshold at $4.5 \times$ the estimated noise SD) and the SWTTEO detector (tuned to return the same number of spikes as ABS); a spike is retained only if both algorithms detect it (coincidence acceptance). Adapted from Hyvärinen et al. (2019).

3.2.2 Single-Electrode Burst Detection

Single-electrode bursts were detected from each active electrode's spike train using a self-adapting algorithm based on the log-ISI histogram (Pasquale et al., 2010), implemented in MATLAB within the in-house developed pipeline (Kapucu et al., 2012). For each electrode, inter-spike intervals were binned on a logarithmic scale and the histogram was smoothed with a lowess smoother. Where the histogram showed two well-separated modes (judged by a void-parameter criterion), the threshold dividing intra-burst from inter-burst intervals was placed automatically at the valley minimum between them; where the bimodality criterion was not met, a fixed default ISI threshold was applied instead. Bursts were defined as runs of at least the minimum-spikes threshold with successive inter-spike intervals below the identified threshold, and an upper-bound ISI threshold guarded against unrealistically long bursts; burst parameters were computed only for burst-detecting electrodes. Parameter values for the four sets appear in Table 3.3, and the reference analyses used Set 2 (BinsPerDec = 20, VoidThr = 0.99, DefaultThr = 170 ms, MaxThr = 1000 ms, MinSpikes = 10, Span = 4, MPD = 2).

The parameters of the detector are illustrated in Figure 3.4.




Figure 3.4. Parameters of the log-ISI single-electrode burst detector. Left: a burst requires a minimum number of spikes (minSpikes) whose successive inter-spike intervals stay below the intra-burst threshold (maxISI), subject to a minimum burst duration (minDur), with a minimum inter-burst interval (minIBI) separating consecutive bursts. Right: the burst- and network-level metrics derived from the detected events on the raw signal. Adapted from Cerina et al. (2023).

Parameter	Set 1 (baseline)	Set 2 (reference)	Set 3 (strict)	Set 4 (strictest)	Notes
Bins per decade	20	20	25	25	Logarithmic histogram resolution
Void parameter	0.99	0.99	0.995	0.995	Bimodality criterion; higher = stricter
Default ISI thr. (ms)	170	170	150	130	Fallback threshold when bimodality criterion is not met
Max ISI thr. (ms)	1000	1000	800	600	Upper bound for histogram-derived threshold
Minimum spikes	8	10	12	15	Minimum spikes per burst
Smoother	lowess	lowess	lowess	lowess	Fixed across sets
Span	5	4	3	3	Smoother span (bins)
Min peak distance	2 bins	2 bins	2 bins	3 bins	Minimum separation between histogram peaks

Table 3.3. Parameters of the single-electrode burst detector (log-ISI histogram method) across the four parameter sets. Set 2 (yellow) was used as the reference set. The smoother type (lowess) is fixed across all sets; bins per decade, void threshold, ISI thresholds, minimum spikes, and smoothing span increase in strictness from Set 1 to Set 4.

3.2.3 Network Burst Detection: Tampere (V1 and V2)

NB detection used two independent algorithms applied in parallel to all recordings: the NeuroGroup (Tampere) log-ISI detector, referred to here as the Tampere detector and implemented in two variants (V1 and V2), and a firing-rate-based detector (FR detector) with parameters adapted from Doorn et al. (2024). Detection was run on spike trains exported from AxIS Navigator, in MATLAB R2025b (The MathWorks, Natick, MA, USA). The two algorithms take different approaches: the Tampere detector aggregates SEBs using the log-ISI-N distribution, while the FR detector looks for peaks in the smoothed population firing rate using locally adaptive thresholds. They were applied to the same dataset so that their outputs could be compared systematically in terms of detected NBR, duration, and fragmentation.

The Tampere detector exploits the bimodal structure of the inter-spike interval (ISI) distribution at the network level. For each recording, spike times pooled across active electrodes were used to build the N-th order ISI histogram (ISI-N) on a logarithmic scale, where N is the order of the interval computed for each spike, the interval between spike i and spike $i+N$ rather than the immediately following one. The value of N ($N = 20$) was chosen empirically by inspecting the burst structure and activity level of the recordings, and kept fixed across all recordings and parameter sets. The two variants share the same histogram; the only thing that differs between them is how they pick the threshold that separates intra-burst from inter-burst intervals.

In V1, the ISI-N threshold was set at the first local minimum of the smoothed log-ISI-N histogram occurring after the global maximum (the intra-burst peak). This works well in recordings with a clearly bimodal ISI distribution, but it can be thrown by local irregularities or by secondary peaks introduced through histogram smoothing. When no local minimum was found after the global maximum, the algorithm returned no NBs for that recording.

In V2, the threshold was derived in two steps. Candidate histogram peaks were first selected with a peak-detection routine requiring both a minimum peak height and a minimum peak prominence of 10% of the histogram maximum. When at least two such peaks were found, the threshold was placed at the global absolute minimum within the

valley between the two leading peaks, rather than at the first local minimum after the global maximum. When fewer than two peaks met the height-and-prominence criteria, or when the two leading peaks were adjacent (separation ≤ 1 bin), the algorithm returned no NBs for that recording. This makes V2 much less sensitive to smoothing artefacts and the thresholds it picks tend to be more consistent across recordings, especially ones with disrupted activity patterns.

Once the threshold was set, candidate NBs were assembled by applying the N-th order ISI criterion of Bakkum et al. (2014) to the pooled network-level spike train. The candidate NBs were then passed through a sequence of quality filters, applied identically in V1 and V2:

(1) Merging step. Closely spaced candidate NBs likely to represent a single fragmented event were combined, using either a fixed inter-burst interval threshold or an adaptive strategy based on the local inter-burst interval distribution. (2) Minimum-duration filter. NBs shorter than a minimum-duration criterion were rejected. (3) Channel-participation filter. Only NBs involving at least a minimum number of unique electrodes were retained. (4) Dominant-channel filter. Within each surviving NB, every participating electrode had to contribute at least a minimum fraction of the mean per-electrode spike count of the NB, and only NBs with at least a minimum number of such dominant electrodes were accepted. The Tampere log-ISI-N NB detection pipeline is summarised in Figure 3.5.




Figure 3.5. The Tampere log-ISI-N NB detector. Spike trains pooled across active electrodes build an N-th order log-ISI histogram ($N = 20$, lowess-smoothed). The two variants differ only in the ISI-N threshold rule: V1 takes the first local minimum after the global maximum, V2 the global minimum in the valley between the two leading peaks. The ISI-N burst criterion (Bakkum et al., 2014) yields candidate NBs, refined by quality filters (merging, minimum duration, minimum electrodes, dominant-channel criterion) into the accepted NBs. Adapted from Pasquale et al. (2010) and Bakkum et al. (2014).

Four NB parameter sets were evaluated: Set 1 (baseline), Set 2 (reference), Set 3 (strict), and Set 4 (strictest). Set 2 served as the reference configuration for all primary statistical analyses in Chapter 4, while the other three were examined in the parameter-sensitivity analysis (Chapter 4, Section 4.4). Set 1 keeps the original baseline kernel (kernel size 6 bins, $\sigma = 3$ bins), which violates the standard $6\sigma+1$ support rule and effectively covers only ± 2.5 bins rather than the intended $\pm 3\sigma$; it is included to quantify the effect of this historical implementation error. Sets 2-4 all use correctly sized kernels and differ systematically in their smoothing scale (75, 150, and 250 ms) and in the stringency of their NB acceptance filters, going from the most permissive settings to the most restrictive, step by step. Table 3.4 lists all the parameter values used.

Parameter	Set 1 (baseline)	Set 2 (reference)	Set 3 (strict)	Set 4 (strictest)	Notes
ISI-N order N	20	20	20	20	N-th order ISI for the log-ISI-N histogram; fixed across sets
Bins per decade	20	22	25	30	Logarithmic histogram resolution
Histogram span	3	3	2	2	Lowess-smoother span (bins)
Smoother	lowess	lowess	lowess	lowess	Local-regression smoother; fixed across sets
Gaussian σ (ms)	75 ms [†]	75 ms	150 ms	250 ms	Kernel SD; scales the NB time-scale assumed by the detector
Kernel size (bins)	6 [†]	19	37	61	[†] Set 1: buggy value ($6\sigma+1$ rule violated; covers only ± 2.5 bins); Sets 2-4: correct $6\sigma+1$
MinCount	8	10	12	15	Minimum candidate NBs per recording
MinDuration (s)	0.35	0.40	0.50	0.60	Minimum accepted NBD
MergingType	Adp.	Adp.	Adp.	Adp.	Adaptive merging; fixed across sets
MergingFrac	0.20	0.20	0.30	0.40	Adaptive merging threshold parameter
MinChannels	3	4	5	6	Minimum unique electrodes per NB
MinDomChannels	3	4	5	6	Minimum dominant electrodes (after contribution filter)
MinChannelContrib	0.20	0.25	0.30	0.40	Minimum per-electrode spike fraction relative to NB mean

Table 3.4. Parameters of the Tampere detector (V1 and V2) across the four parameter sets evaluated in this work. Set 2 (yellow) was used as the reference set for the principal analyses of Chapter 4. Set 1 (orange) retains the original baseline kernel (size 6 bins, $\sigma = 3$ bins), which violates the $6\sigma+1$ rule and covers only ± 2.5 bins; it is included to quantify the impact of this historical bug. All sets share the ISI-N order $N = 20$, the lowess smoother, and an adaptive merging strategy. Both V1 and V2 use the same parameter values and differ only in the threshold-selection criterion on the log-ISI-N histogram.

3.2.4 Network Burst Detection: Firing-Rate-Based

The firing-rate-based detector (hereafter FR detector) follows the smoothed population-firing-rate approach described in Section 2.6, with parameters adapted from Doorn et al. (2024). For each recording, spike counts were binned across all active electrodes in 25 ms time bins and the resulting population firing rate was smoothed with a Gaussian kernel (reference Set 2: $\sigma = 75$ ms (3 bins), kernel size 19 bins; the four parameter sets are listed in Table 3.5).

NBs were delineated directly on the smoothed signal using thresholds derived for each candidate region from the local maximum of the smoothed signal within a moving 30 s window (reference Set 2; Table 3.5). Onset required the smoothed rate to stay above 25% of the local maximum for at least 50 ms, with at least 30% of active electrodes co-active over the same window. The NB was extended as long as the smoothed rate stayed above an offset threshold of 2% of the local maximum within a sliding 100 ms (2×50 ms) check window, and its end was placed at the bin where this condition first failed; the original implementation uses 1/100 of the maximum firing rate instead (Doorn et al., 2024). Because the thresholds are computed locally, the detector adjusts to the activity level of each recording. The fractional threshold values, however, were taken from Doorn et al. (2024) and not re-optimised for the present dataset: in the absence of an expert-annotated reference, there was no objective criterion for choosing among alternative settings. Only the smoothing kernel and the moving window were therefore varied across parameter sets (Table 3.5; Section 5.2). Separately, peak detection on the smoothed signal identified the prominent peaks of global activity, requiring a minimum peak height of 1/16 and a minimum peak prominence of 1/10 of the maximum global firing rate. These peaks play no part in delineating the NBs themselves; instead, the number of them falling inside each detected NB gives its fragment count (Section 3.2.5). The whole analysis pipeline is shown in Figure 3.6, while the full set of parameter values used here is reported in Table 3.5.




Figure 3.6. The firing-rate-based (FR) NB detector. Spikes pooled across active electrodes (mean firing rate ≥ 0.02 Hz) are binned at 25 ms into a population firing rate and smoothed with a Gaussian kernel ($\sigma = 75$ ms, reference set). NBs are delimited by locally adaptive onset/offset thresholds (onset: rate $> 25\%$ of the local maximum for ≥ 50 ms, with $\geq 30\%$ of active electrodes participating; offset: rate above 2% of the local maximum within a sliding 100 ms window). Parameters adapted from Doorn et al. (2024).

Parameter	Set 1 (baseline [†])	Set 2 (reference [*])	Set 3 (strict)	Set 4 (strictest)	Notes
Time bin	25 ms	25 ms	25 ms	25 ms	Fixed across sets
Gaussian σ (ms)	75 ms [†]	75 ms	150 ms	250 ms	Kernel SD; matched to Tampere sets
Kernel size (bins)	6 [†]	19	37	61	[†] Set 1: same buggy kernel as Tampere Set 1
Moving window (s)	30	30	25	20	Window for local maximum driving the adaptive thresholds
Active electrode thr.	0.02 Hz	0.02 Hz	0.02 Hz	0.02 Hz	Fixed across sets
Onset threshold	25%	25%	25%	25%	Fraction of local max; fixed across sets
Onset duration	50 ms	50 ms	50 ms	50 ms	Minimum duration above onset threshold; fixed
Offset threshold	2%	2%	2%	2%	Fraction of local max; fixed across sets
Electrode participation	30%	30%	30%	30%	Fixed across sets
MinPeakHeight	max/16	max/16	max/16	max/16	Fixed across sets (fragmentation)
MinPeakProminence	max/10	max/10	max/10	max/10	Fixed across sets (fragmentation)

Table 3.5. Parameters of the FR detector across the four parameter sets. *Set 2 (yellow) was used as the reference set. [†]Set 1 (orange) uses the same buggy kernel as the Tampere baseline. The onset/offset thresholds, electrode-participation criterion, and fragment-detection thresholds are fixed across all four sets; only the Gaussian smoothing kernel and the moving-window duration vary.

3.2.5 Fragmentation Detection

Both NB detectors include an integrated fragmentation analysis that characterises the internal temporal structure of each accepted NB. The peak-detection routine was applied to the smoothed population firing rate within the boundaries of each accepted NB, requiring a minimum peak height of 1/16 and a minimum peak prominence of 1/10 of the maximum

global firing rate. Each peak detected inside an NB was treated as a fragment, and NBs containing two or more such peaks were classified as fragmented NBs. For each recording and each detector, the fragmentation metrics computed were: the proportion of NBs classified as fragmented relative to the total number of accepted NBs; and the mean number of fragments per NB, computed over NBs containing at least one detected peak (FragPerNB, Section 3.3.1). A fragment-overlap analysis between the Tampere and FR detectors was also carried out to gauge how far the two agreed in identifying fragmented events within the same recording (Section 4.3.3). The pipeline for NB fragmentation analysis is shown in Figure 3.7.




Figure 3.7. Fragmentation analysis. Within each accepted NB (Tampere or FR), peak detection is run on the smoothed population firing rate (height $\geq 1/16$, prominence $\geq 1/10$ of the maximum); each intra-NB peak counts as one fragment, and an NB with more than one fragment is a fragmented NB. The fragmentation metrics (percentage of fragmented NBs, fragments per NB) follow from these counts. Adapted from Doorn et al. (2024).

3.3 Spontaneous Activity Analysis

3.3.1 Activity Metrics

From the spike trains, the SEBs (Section 3.2.2), and the NBs returned by the detectors of Section 3.2, a set of recording-level activity metrics was computed. The specific subset of metrics carried into each statistical analysis (A1–A3) is given in Section 3.5; the remaining metrics provide complementary descriptive information used in Chapter 4. These metrics are summarised schematically in Figure 2.5 (Section 2.5).

Mean firing rate (MFR). For each recording, the MFR was taken as the total number of detected spikes in the well, pooled across all electrodes, divided by the recording duration (600 s), giving a recording-level firing rate in spikes per second (Hz). Recordings with no detected spikes were assigned $\text{MFR} = \text{NaN}$ and excluded from the corresponding statistical comparisons.

Single-electrode burst rate (SEB rate). An electrode's burst rate was the number of SEBs detected on it (Section 3.2.2), divided by the recording duration in minutes (10 min); the recording-level SEB rate was the mean across the well's active, burst-detecting electrodes.

Single-electrode burst duration (SEB duration). Each SEB was timed from its first to its last spike; the recording-level SEB duration was the mean burst duration across all bursts detected in the recording.

Percentage of random spikes (%RS). For each recording, %RS was the proportion of spikes falling outside any detected SEB, expressed as a percentage of the total spike count. It gives a complementary view of how strongly the spiking is organised into bursts.

Network burst rate (NBR). For each recording and each detector (Tampere V1, Tampere V2, FR), the number of accepted NBs was divided by the recording duration in minutes (10 min), giving the NBR in bursts per minute.

Network burst duration (NBD). Each accepted NB was timed between its onset and offset, as returned by the detector; the recording-level NBD was the mean NBD across all accepted NBs of the recording.

Coefficient of variation of the inter-burst interval (CV(IBM)). Within each recording, the inter-burst intervals (IBIs) between consecutive accepted NBs were extracted, and CV(IBM)

was computed as the ratio of their standard deviation to their mean. Low CV(IBM) values indicate regular NB occurrence, while higher values indicate more irregular network bursting (Chiappalone et al., 2006).

Local variation of the inter-burst interval (LvR). As a complementary descriptor of NB regularity, the revised local variation (LvR; Shinomoto et al., 2009) was computed, separately for each detector, from the sequence of inter-burst intervals between consecutive accepted NBs. Unlike the coefficient of variation, LvR is based on the stepwise relationship between consecutive inter-burst intervals and exhibits enhanced invariance to slow rate fluctuations. For a sequence of n inter-burst intervals, it is defined as $LvR = [3/(n-1)] \sum_i \{ [1 - 4 I_i I_{i+1} / (I_i + I_{i+1})^2] \cdot [1 + 4R/(I_i + I_{i+1})] \}$ with $(i = 1, \dots, n-1)$ and R is a refractoriness constant (set to 5 ms in this work). Because $R = 5$ ms is negligible relative to the inter-burst intervals (on the order of seconds), LvR effectively reduces to the local variation Lv applied to the NB sequence (Shinomoto et al., 2009; Sanlı & Lambiotte, 2015). The metric distinguishes three canonical firing regimes according to the relative displacement of consecutive spikes: $LvR \approx 0.5$ corresponds to a regular pattern, $LvR \approx 1$ to a random (Poisson-like) pattern and $LvR \approx 1.5$ to a bursty pattern (Shinomoto et al., 2009; Figure 3.8).




Figure 3.8. Spike-generation regimes distinguished by the local variation LvR. (A) LvR distribution across the whole neuronal population, with the regular ($LvR \approx 0.5$, blue), random ($LvR \approx 1$, green) and bursty ($LvR \approx 1.5$, red) ranges highlighted. (B) Representative single-neuron spike trains and corresponding LvR distributions for the three regimes. (C) LvR distributions for individual cortical areas, illustrating how the metric discriminates regional firing patterns. Reproduced from Shinomoto et al. (2009).

Fragments per NB (FragPerNB). Within each recording, the fragmentation analysis (Section 3.2.5) returns the number of internal fragments for every accepted NB. The recording-level FragPerNB was defined as the mean number of fragments per NB, restricted to NBs containing at least one fragment; recordings with no fragmented NBs were assigned $\text{FragPerNB} = \text{NaN}$.

3.4 Frequency-Band Validation

To validate the Tampere V1 NB and fragment detector independently of the spike-based pipeline, the spectral event-detection framework introduced by Cimbalnik et al. (2018a, 2018b) for detecting high-frequency oscillations (HFOs) in intracranial EEG was adapted and applied to the raw MEA recordings. In intracranial recordings, ripples and fast ripples are established electrophysiological biomarkers of epileptogenic tissue (Jacobs et al., 2008; Worrell et al., 2004, 2008). The framework was chosen because it provides a validated event detector that runs entirely on the raw filtered signal, independently of the spike detection underlying the NB detector. The point of using this framework is that it is independent of the spike-detection and NB-detection algorithms, though not of spike energy in the raw signal itself, since the Fast Ripple band (250-600 Hz) overlaps the spike band: if the NB detector correctly marks windows of genuine network synchronisation, then spectral events detected independently from the raw signal should cluster within NBs, and more tightly still within fragment windows and stay sparse outside them.

3.4.1 Z-score and Peak Detection

Following Cimbalnik et al. (2018b), the raw signal from each electrode was band-pass filtered (8th-order Butterworth, zero-phase via `filtfilt`) into three frequency bands: High Gamma (65-100 Hz), Ripple (100-250 Hz), and Fast Ripple (250-600 Hz). For each electrode, the Hilbert envelope was computed and z-scored against the whole recording (baseline = full recording duration per electrode). To capture network-level co-activation rather than local single-electrode events, a cross-electrode maximum z-score trace was built by taking, at each time point, the highest z-score across all 16 electrodes. This inverts the logic of the original iEEG framework, where simultaneous multi-channel events are rejected as artefacts because physiological HFOs are spatially confined to a small patch of

cortical tissue (Blanco et al., 2010); by contrast, within *in vitro* MEA data cross-electrode synchrony is the signal of interest rather than an artefact, consistent with the NB paradigm and with HFO observations in MEA cultures (Hales et al., 2012). Two features of the pipeline keep this aggregation from being driven by common-mode interference: each electrode is z-scored against its own full-recording baseline, so any persistent shared signal (for example mains interference) is absorbed into the baseline rather than producing transient supra-threshold events; as reported in Section 4.4, detected events concentrate within NBs rather than being uniformly distributed across the recording, which a continuous artefact would not be. Spectral events were defined as continuous segments exceeding $z = 3.0$ for longer than three oscillation cycles at the lower cutoff frequency of each band, following the parameters of Cimbalnik et al. (2018b).

3.4.2 Cimbalnik Amplitude

For each detected event, the mean amplitude of the Hilbert envelope was computed as the key spectral feature, following the feature analysis of Cimbalnik et al. (2018a). In that work the amplitude feature separates pathological from physiological HFOs (Cimbalnik et al., 2018a, 2018b): pathHFO show a pronounced long-tailed distribution toward higher amplitudes relative to physHFO (their Figure 4). Here it is used not as a biomarker of epileptogenicity, a meaning specific to epileptic iEEG that does not transfer to healthy cultures, but as an index of how intense the network synchronisation associated with each event is. The peak z-score of each event was also retained as a complementary measure of event strength.

3.4.3 IN-NB and OUT-NB Event Classification

Each detected spectral event was assigned to one of three mutually exclusive classes, according to whether its peak time fell within (1) a fragment window (defined as ± 25 ms around each fragment peak time; IN-FRAG), (2) a NB but outside any fragment (IN-NB-OUT), or (3) outside any NB (OUT-NB). This three-level stratification lets NB detection and fragment detection be assessed separately: events should be strong inside NBs relative to outside and stronger inside fragments relative to the interstitial within-NB period. The validation metrics derived from this classification were the peak z-score distribution across

the three classes; the high-gamma event rate normalised by the cumulative time of each class; the NB and fragment hit-rate (the percentage of NBs or fragments containing at least one high-gamma event); and the mean envelope amplitude distribution across classes. The fragment-window half-width was fixed at 25 ms. The underlying spectral event-detection framework is the previously validated detector of Cimbalnik et al. (2018a, 2018b), whose parameters were established and characterised in that work; here it is adopted and simply adapted to MEA recordings as an independent validation tool, so its robustness was not re-assessed in the present study. It is worth being clear about what this validation can and cannot tell us: since both methods are ultimately picking up the same thing (coordinated firing), agreement between them is reassuring but does not prove either one is correct. The corresponding results are reported in Section 4.4.

3.5 Statistical Procedures

Four families of statistical analyses were run: Analyses 1 to 3 in MATLAB R2025b (The MathWorks, Natick, MA, USA) and Analysis 4 in Python. Analysis 1 (A1) compared the four cell lines longitudinally for each core metric using the Tampere V1 detector; Analysis 2 (A2) compared the Tampere V1 detector against the FR detector within each cell line; and Analysis 3 (A3) compared the two variants of the Tampere detector (V1 vs V2) within each cell line and also compared the four parameter sets of the Tampere V1 detector. Analysis 4 (A4) compared the three spectral-event classes and is described separately in Section 3.5.4. Analyses A1 to A3 are built on a linear mixed-effects (LME) framework with the plate of origin as a random intercept, to absorb between-plate biological variability across independent differentiation batches, complemented by per-DIV pairwise comparisons; three of the four network-burst timing metrics of A2 are assessed by the per-DIV comparisons alone (Section 3.5.2), and A4 follows a non-parametric path. An $\alpha = 0.05$ was used throughout, and corrected for multiple comparisons using the Benjamini-Hochberg FDR procedure within each analysis (FDR-BH; Benjamini & Hochberg, 1995). The overall statistical analysis framework is described in Figure 3.9.




Figure 3.9. Decision flow of the statistical analysis. After assumption checks (Shapiro-Wilk for normality, F-test for equal variance), each metric follows a parametric path (linear mixed-effects model with plate as a random intercept; ANOVA / Welch t as fallback) when the assumptions hold, or a non-parametric path (Kruskal-Wallis for ≥ 3 groups, Mann-Whitney U for two) otherwise; raw p-values are corrected with the Benjamini-Hochberg FDR procedure and reported with Cohen's d effect sizes (and 95% CIs on descriptives). The same workflow underlies Analyses A1-A3 (Sections 3.5.1-3.5.3).

3.5.1 Cross-Cell-Line Analysis (A1)

For each metric (MFR, SEB rate, SEB duration, %RS, network synchrony, NBR, NBD, CV(IBM), LvR, FragPerNB, and fragment count), the well-level values from the Tampere V1 detector were modelled as a function of cell line, DIV, and their interaction, with plate as a random intercept (formula: $\text{metric} \sim \text{CellLine} \times \text{DIV} + (1 | \text{Plate})$). The model was fitted with MATLAB's `fitlme` function. The overall p-value reported for each metric corresponds to the interaction term of the omnibus ANOVA of the fitted LME, and the adjusted R^2 of the fitted model is reported as a model-level effect-size summary. Residual normality was checked with the Shapiro-Wilk test, falling back to Lilliefors when needed; if residuals were strongly non-normal ($p < 0.01$), the LME result was kept but flagged as approximate rather than dropping it entirely, because no standard non-parametric test preserves the plate-level random-intercept structure and the F-tests of mixed models are

known to be robust to moderate departures from residual normality; given the limited number of recordings, retaining the (flagged) LME was preferable to discarding the random-effects structure. Pairwise post-hoc comparisons between cell lines were computed at each DIV with an adaptive test: when both groups met the Shapiro-Wilk normality criterion ($p > 0.05$), a two-sample t-test was used (Welch correction when variances differed at the F-test level $p < 0.05$); otherwise a Mann-Whitney U test was used. The resulting raw p-values were corrected within each metric using the Benjamini-Hochberg false discovery rate procedure (FDR-BH; Benjamini & Hochberg, 1995). Pairwise contrasts derived directly from the LME coefficients were also computed via `coefTest` (with FDR-BH correction) and exported for verification; because those contrasts are built from the cell-line main-effect coefficients of a model that also contains the $\text{CellLine} \times \text{DIV}$ interaction, they estimate the between-line difference at the reference time point (DIV 53) rather than a difference averaged over DIVs, and for this reason they are not reported in Chapter 4.

In addition to the per-DIV post-hoc comparisons, an *overall* between-line test was computed for each metric by pooling all recordings of a given well across time points; this is what the significance brackets on the right-hand margin of the trajectory line plots report. The test is intended as a summary indication of whether two lines differ globally in a given feature, complementary to — and not a substitute for — the per-DIV comparisons, which remain the confirmatory basis of the developmental claims made in this work. In this pooled test the unit of analysis is the **well $\times$ DIV observation** rather than the well, so that a well recorded at four time points contributes four rows ($n = 299\text{--}541$ observations across 105–127 wells, depending on the metric and the detector). Groups were compared with the same adaptive scheme used for the per-DIV post-hocs (Shapiro-Wilk at $\alpha = 0.05$, followed by Welch’s t-test or the Mann-Whitney U test), and raw p-values were corrected across the six pairwise line contrasts within each metric with the Benjamini-Hochberg procedure.

Because repeated measurements of the same well are not fully independent, this pooling inflates the nominal sample size and yields anti-conservative p-values. The magnitude of that inflation was quantified rather than assumed: variance-component models of the form $\text{metric} \sim \text{CellLine} + \text{DIV} + (1 \mid \text{Plate}) + (1 \mid \text{Well})$ were fitted to each metric, and the resulting intraclass correlation ρ was used to compute the design effect, $\text{DEFF} = 1 + (m - 1)\rho$, where m is the mean number of observations per well. Across all metrics and detectors DEFF ranged from 1.00 to 2.04 ($m = 2.2\text{--}4.3$), corresponding to $\rho < 0.32$ and to an effective

sample size of roughly 50–100% of the nominal one. The design is moreover balanced in time, each line contributing the same number of wells at every time point (mean DIV 70.2–70.6 across the four lines), so pooling does not confound line identity with developmental stage and leaves the point estimates and effect sizes unbiased. Two sensitivity analyses were performed accordingly (Section 4.1.3 and Appendix A): in the first, each well was collapsed to the mean of its time points and the identical adaptive test was re-applied, so that only the experimental unit changed; in the second, contrasts were derived from a linear mixed-effects model of the form $metric \sim CellLine + DIV + (1 | Well)$ fitted on the log1p scale, which uses all observations while estimating the within-well correlation from the data.

3.5.2 Tampere vs Firing-Rate Detector Comparison (A2)

Within each cell line, the well-level values of NBR and FragPerNB from the Tampere V1 detector and from the FR detector were compared with an LME (formula: $metric \sim Detector + (1 | Plate)$). The omnibus effect of Detector was taken from the ANOVA of the fitted LME. When the LME could not be fitted (insufficient data, $n < 6$), the analysis fell back to a Mann-Whitney U test. Pairwise comparisons between the two detectors were also computed at each DIV using the same adaptive test as in A1, and raw p-values were FDR-BH corrected within each metric. The remaining network-burst timing metrics — NBD, CV(IBM) and LvR, together with NBR — were compared between the two detectors at each DIV with the same adaptive test and the same FDR-BH correction, but without an omnibus mixed-effects model; the significance heat-maps for these three metrics therefore report per-DIV contrasts only.

3.5.3 V1 vs V2 and Parameter-Set Comparison (A3)

The comparison between Tampere V1 and Tampere V2 followed the same procedure as A2: an LME of the form $metric \sim Detector + (1 | Plate)$ within each cell line for NBR and

FragPerNB, complemented by per-DIV adaptive pairwise tests, with FDR-BH correction within each metric.

As a further component of A3, the sensitivity of the detectors to their parameter configuration was assessed by comparing the four parameter sets of Table 3.4 (Sets 1–4). This comparison used the same LME framework as A2 and A3: within each cell line, the well-level metrics were modelled as $\text{metric} \sim \text{ParamSet} + (1 \mid \text{Plate})$, the omnibus effect of the parameter set being taken from the ANOVA of the fitted model (with a Kruskal–Wallis fallback when the residuals were not normally distributed).

3.5.4 Spectral-Event Class Comparison (A4)

The fourth analysis (A4) tested whether the three spectral-event classes defined in Section 3.4.3 — events inside fragments (IN-FRAG), events in the within-NB background (IN-NB-OUT) and out-of-NB events (OUT-NB) — differed in their spectral properties, separately for the High Gamma, Ripple and Fast Ripple bands. Because each recording contributed hundreds to thousands of mutually correlated events, the unit of analysis was the cell-line $\times$ DIV block rather than the individual event: for every block the median of each measure (peak z-score and mean Hilbert-envelope amplitude) was computed within each class, and only blocks in which all three classes were present were retained. The three class medians were then compared with a non-parametric repeated-measures design — a Friedman omnibus test across the three classes, followed, when the omnibus was significant, by pairwise Wilcoxon signed-rank tests with Holm step-down correction over the three comparisons. Holm’s step-down procedure was preferred here over a false-discovery-rate (Benjamini–Hochberg) correction because each band involves only a small, closed family of three pre-specified pairwise contrasts on paired within-block data, a setting in which strong control of the family-wise error rate is both feasible and more appropriate than the FDR control used for the large parallel comparison families elsewhere in this work. Significance is annotated directly on Figures 4.19 and 4.20 (*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$; ns, not significant); the analogous test on the High-Gamma event rate is reported in Section 4.4.3. These analyses were carried out in Python using the SciPy library.

3.5.5 Visualisation Conventions

Results are reported with three complementary visualisations. (1) Line plots of metric means across DIVs, with shaded bands giving the 95% confidence interval of the mean (Student's t-distribution). These bands were computed after removing outliers within each DIV using the interquartile-range (IQR) rule (within each DIV, missing (NaN) values were first removed and then, as long as at least four valid values were left, values were discarded below $Q1 - 1.5 \times IQR$ or above $Q3 + 1.5 \times IQR$, where $Q1$ and $Q3$ are the 25th and 75th percentiles and $IQR = Q3 - Q1$). The same interquartile-range outlier removal was applied to all of the computed metrics when building the descriptive line plots; for the inferential statistics, by contrast, the full dataset was retained: given the limited number of recordings, no data were discarded, so the LME models and per-DIV tests described above operate on the full data. Significance brackets drawn in the right-hand margin of the trajectory line plots (* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$) report the overall between-line test described in Section 3.5.1, and are shown whenever that contrast is significant after FDR-BH correction, independently of the outcome of the omnibus LME. (2) Pairwise heat-maps of $-\log_{10}(\text{adjusted } p\text{-value})$ for each pairwise comparison at each DIV; these are shown even when the overall LME is not significant, with white cells indicating comparisons that were not tested, light grey cells comparisons that were tested but not significant (labelled "ns"), and progressively darker blue shades significant comparisons, with asterisk annotations (*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$) in each significant cell. The cell colour thus indicates significance and the annotation its direction (based on the sign of the mean difference); the standardised effect size (Cohen's d), although computed and stored in the output tables for each comparison, is not mapped onto the heat-maps. (3) Boxplots of well-level distributions per cell line and DIV, used where helpful to convey within-group dispersion.

Chapter 4

Results

4.1 Spontaneous Neuronal Activity

Spontaneous activity was characterised across the five recording time points (DIV 53, 60, 67, 83, and 90), both at the level of single-electrode spiking and bursting and at the level of network-wide synchronisation. Section 4.1.1 shows the qualitative differences in single-electrode activity through representative raster plots and multivariate activity profiles; Section 4.1.2 reports the quantitative activity metrics; and Section 4.1.3 reports the statistical analysis.

4.1.1 Raster Plot Overview

To illustrate spontaneous activity development, representative single-electrode raster plots are presented for the cell lines at the first (DIV 53) and last (DIV 90) recorded time points (Figure 4.1). Each raster contains rows representing electrodes, vertical ticks indicating spikes and red rectangles denoting single-electrode bursts (SEBs). Comparing these two time points reveals a progressive organisation of spiking into bursts over development; looking across the lines demonstrates a gradient from the disrupted activity of the severe DS line to the well-organised activity of the controls.




Figure 4.1. Representative single-electrode raster plots (with SEB overlay) for the four cell lines at DIV 53 and DIV 88. Rows are electrodes, vertical ticks are spikes and red rectangles denote SEBs. The control lines already show dense, well-spread bursting at DIV 53, whereas the severe DS line DD1C shows sparse activity concentrated on a few electrodes.

Visually, the control lines already show dense, well-spread bursting at DIV 53 and by DIV 88 it is even more pronounced. The severe DS line DD1C shows sparse activity concentrated on a few electrodes and develops only weakly over time; the mild DS line DD5A falls somewhere in between.

Figure 4.2 gives a complementary overview of the activity profile of each cell line at all five DIV time points, summarising six core metrics simultaneously on a single radar chart per DIV (Tampere V1, raw values). This gives a broad view of how each line develops before the individual metrics are examined below. The six axes represent: NB rate, single-electrode burst rate (SB Rate), mean firing rate (MFR), coefficient of variation of the inter-burst interval (CV IBI), NB duration (NBD), and fragmentation prevalence (Frag. Prev.). A bigger area generally means more or longer activity, with two exceptions: CV IBI (smaller = more regular bursting) and Frag. Prev. (which runs 0-1).




Figure 4.2. Multivariate activity profiles (radar charts) of the four cell lines, overlaid within each panel, at five developmental time points (Tampere V1, raw values, no normalisation): (a) DIV 53, (b) DIV 60, (c) DIV 67, (d) DIV 83, (e) DIV 90. The six axes in every panel are NB rate, single-electrode burst rate (SEB rate), MFR, CV(IBM), NB duration and fragmentation prevalence (Frag. Prev.).

Across all DIVs, the two control lines (TUBA, C10902) cover a substantially larger area than the DS lines on NB rate, SEB rate, and MFR; the gap is already visible at DIV 53 (panel a) and persists through DIV 90 (panel e). The severe line DD1C shows the smallest profile area throughout, and the mild line DD5A sits in between. CV IBM runs the other way: DD1C keeps a comparatively high CV IBM (a sign of less regular bursting), while NB duration stands out for the large extent of DD5A at the intermediate time points.

4.1.2 Activity Metrics

The quantitative single-electrode and network activity metrics are summarised together in Figure 4.3, which shows the developmental trajectories of mean firing rate (MFR), single-electrode burst rate (SEB rate), single-electrode burst duration (SEB duration), network synchrony (active electrodes per NB) and the percentage of random, non-burst spikes (%RS) for the four cell lines across DIVs.




Figure 4.3. Developmental trajectories of the activity metrics across DIVs for the four cell lines (mean $\pm$ 95% CI; IQR outlier removal, $k = 1.5$): (A) mean firing rate, (B) single-electrode burst rate, (C) single-electrode burst duration, (D) network synchrony (active electrodes per NB) and (E) percentage of random spikes.

Mean firing rate (MFR). How the MFR developed differed markedly between lines (Figure 4.3A). The severe DS line DD1C stayed essentially flat across all time points, at a low firing rate. DD5A, the mild line, followed a more typical trajectory, rising up to DIV 83 and then plateauing between DIV 83 and 90. TUBA traced the same developmental shape but reached higher firing rates, climbing steeply from DIV 67 onward. C10902 behaved in the opposite way, starting high at DIV 53 and decreasing over time.

Single-electrode burst rate (SEB rate). Figure 4.3B shows the SEB rate rising over time for the TUBA, DD1C and DD5A electrodes, suggesting a developmental shift from random spiking to organised bursting at the electrode level. TUBA achieved the highest SEB rate by DIV 90. C10902 behaved non-monotonically, falling and then rising before dropping again at the final time point.

Single-electrode burst duration (SEB duration). Burst duration followed line-specific temporal trajectories (Figure 4.3C). DD1C stayed essentially flat, with only a slight decrease from DIV 53 to DIV 90. DD5A rose sharply up to DIV 67 and then fell, consistent

with a network that breaks into shorter bursts at later stages. TUBA increased overall, though along an up-and-down path, while C10902 decreased over time, in line with a relative shift back towards random spiking.

Network synchrony (active channels per NB). The mean number of electrodes taking part in each NB indexes network-wide synchronisation (Figure 4.3D). DD1C, the severe DS line, showed the lowest synchrony throughout development, as expected. DD5A, the mild line, was intermediate: above DD1C but below the controls. Both control lines (TUBA, C10902) held high synchrony across the entire period.

Percentage of random spikes (%RS). This metric complements network synchrony, measuring the fraction of spikes that belong to no detected burst (Figure 4.3E). The percentage decreased during development in every line, marking the progressive organisation of spiking into bursts as the networks matured. At every time point the two DS lines carried the highest %RS relative to the controls, DD1C consistently the highest, which confirms the reduced bursting organisation of the disease lines.

4.1.3 Statistical Analysis

The four cell lines were compared longitudinally with the cross-cell-line analysis (Analysis 1; Section 3.5.1): for each metric, a linear mixed-effects (LME) model with cell line, DIV, and their interaction as fixed effects and plate as a random intercept, followed by per-DIV pairwise comparisons corrected with the Benjamini-Hochberg false discovery rate (FDR). The statistical results for all activity metrics are summarised in the pairwise significance heat-maps of Figure 4.4 (colours match the previous section: DD1C red; DD5A orange; TUBA dark blue; C10902 light blue).




Figure 4.4. Analysis 1 (cross-cell-line LME with Benjamini-Hochberg FDR correction): pairwise significance heat-maps ($-\log_{10}$ adjusted p-value) for each cell-line pair at each DIV, for MFR, SEB rate, SEB duration, %RS and network synchrony. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

The MFR gave the clearest and most consistent separation between lines (Figure 4.4). The severe DS line DD1C had a significantly lower MFR than all three other lines at essentially every time point: significantly below DD5A ($p < 0.05$ - 0.01 across DIVs), and strongly below both controls (DD1C < TUBA and DD1C < C10902, $p < 0.001$ at all DIVs). DD5A was significantly lower than TUBA at most time points. This ordering, DD1C < DD5A < controls, recapitulates the clinical severity gradient and was the most robust cross-line effect seen in this study.

The remaining activity metrics (SEB rate, SEB duration, %RS and network synchrony) followed the same disease-control gradient (Figure 4.4): the disease lines, and DD1C in particular, separated significantly from the controls at the majority of time points, while the two control lines were largely indistinguishable from each other. The matching statistics for the NB rate and the fragments-per-NB ratio (FragPerNB) appear alongside the detector comparison in Sections 4.2.4 and 4.3.4, respectively.

Because the overall between-line brackets shown on the trajectory plots pool the recordings of each well across time points (Section 3.5.1), two sensitivity analyses were run to verify that this pooling does not drive the reported differences (Appendix A). Collapsing each

well to its DIV-averaged value and re-applying the identical test altered the significance status of 12 unique contrasts out of 168: ten were lost and two gained. Six of the ten losses were patient-versus-patient contrasts (DD1C vs DD5A), which do not bear on any conclusion drawn here, and two were control-versus-control contrasts, whose loss is consistent with the two control lines behaving similarly. Only two disease-versus-control contrasts were affected — NB rate for DD1C vs C10902 and NB duration for DD5A vs C10902, both with the Tampere V1 detector — while one further disease-versus-control contrast, FragPerNB for DD1C vs TUBA with Tampere V2, gained significance. The mixed-effects contrasts reproduced the pooled results in the large majority of metric-detector combinations. The disease–control gradient described above is therefore not an artefact of the unit of analysis.

Taken together, the spontaneous-activity metrics describe a consistent disease–control gradient across development: the control lines (TUBA, C10902) build up denser, more synchronised and more regular bursting, whereas the DS lines — the severe DD1C in particular — remain sparser and flatter, with the mean firing rate providing the single most discriminating metric between disease and control.

4.2 Network Burst Detection:

Tampere vs Firing-Rate Detector

This section compares the two NB detection algorithms, the Tampere log-ISI detector (V1) and the FR detector, applied to the same recordings, in terms of NB rate and duration (Section 4.2.2) and inter-burst-interval regularity (Section 4.2.3). All results in this section use the reference parameter set (Set 2; see Section 3.2.3 and Chapter 3, Table 3.4).

4.2.1 Raster-Based Comparison

To show qualitatively how the two detectors delineate NBs, representative rasters of the same well are displayed with the NBs found by the FR detector (panel A, blue rectangles) and by the Tampere detector (panel B, red rectangles), at the first (DIV 53) and last (DIV 90) time points. The disease lines (DD1C, DD5A) are shown first in Figure 4.5 and the

control lines (TUBA, C10902) in Figure 4.6. These panels visualise how the detected events evolve from DIV 53 to DIV 90 and how the two algorithms can diverge on a given well.




Figure 4.5. Representative NB rasters for the disease lines, severe Dravet DD1C and mild Dravet DD5A, at DIV 53 and DIV 88: (A) FR detector (blue rectangles) and (B) Tampere (red rectangles).




Figure 4.6. Representative NB rasters for the control lines TUBA and C10902 at DIV 53 and DIV 88: (A) FR detector (blue rectangles) and (B) Tampere (red rectangles).

4.2.2 Network Burst Rate and Duration

The developmental trajectories of the NB metrics for both detectors are shown together in Figure 4.7 (NB rate, NB duration, CV(IBM) and LvR). Network burst rate (NBR). With the Tampere detector, all four lines showed the expected developmental increase in NBR: TUBA, DD1C and DD5A rose up to DIV 83 and then dipped slightly at the final time point, whereas C10902 kept increasing up to DIV 90, with TUBA leading from the start. Overall the control lines showed higher NB activity than the disease lines. Within the Dravet pair, DD1C started below DD5A at DIV 53, in line with the clinical severity of the donors, but the two trajectories crossed soon afterwards and DD5A remained the lower of the two for the rest of the recording period. The FR detector gave broadly similar results for the disease–control comparison but showed the opposite behaviour within the Dravet pair: DD1C started above DD5A and the two trajectories crossed after DIV 67, from which point

DD1C < DD5A, the ordering matching the clinical phenotype; the trajectories of the two control lines also crossed over after DIV 67. For the other NB metrics (NB duration, CV(IBM) and LvR), the ordering of the lines was broadly consistent with clinical severity under both detectors.

Network burst duration (NBD). NBD followed line-specific temporal trajectories (Figure 4.7). For DD1C it decreased after DIV 67, which may reflect disruption of sustained network activity; DD5A peaked later, falling after DIV 83; the two control lines stayed comparatively flat, indicating a more stable NBD across development. The FR detector produced trajectories broadly similar to the Tampere ones, supporting the robustness of the NB-duration trends to the choice of detector. The corresponding per-DIV statistical comparison between the two detectors is reported in the statistical analysis (Section 4.2.4).




Figure 4.7. Tampere (red) vs FR (blue) developmental trajectories across DIVs for the four cell lines (mean $\pm$ 95% CI), for the NB metrics: (A-B) NB rate, (C-D) NB duration, (E-F) CV(IBM) of the inter-burst interval, and (G-H) LvR. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

Direct comparison per DIV. Across most lines and time points the Tampere detector returned a higher NBR than the FR detector (Figure 4.8, Figure 4.9), i.e. it delimited more, and shorter, events; which of the two partitions is closer to the actual burst structure cannot

be established without an expert-annotated reference. The statistical analysis is reported in Section 4.2.4.

4.2.3 Inter-Burst Interval Regularity

The regularity of network bursting was assessed with two complementary measures of the inter-burst interval (IBI): the coefficient of variation (CV(IBI)) and the local variation LvR (Figure 4.7). For the CV, a value of 1 corresponds to a Poisson (memoryless) process, in which the standard deviation of the IBI distribution equals its mean, whereas a value of 0 corresponds to a perfectly regular sequence of intervals. The LvR (computed here with a refractoriness constant $R = 5$ ms) is a local measure of regularity, more robust to slow rate fluctuations than the global CV.

CV(IBI) (Tampere). The per-DIV CV(IBI) distributions show C10902 holding the lowest, most regular values at the earlier time points, while the DS lines tended towards higher and more variable CV values, consistent with less regular network bursting. LvR (Tampere). The LvR trajectories revealed line-specific developmental patterns: the healthy control TUBA decreased over development; the severe DS line DD1C decreased even more sharply, a loss of irregularity that may reflect the loss of inhibitory modulation associated with SCN1A loss-of-function; the mild line DD5A showed a U-shaped trajectory, decreasing up to DIV 67 and then rising, possibly signalling a late compensatory mechanism in keeping with its milder clinical phenotype; and C10902 increased slightly while staying very low throughout, consistent with a line whose network activity is dense and highly synchronised from the outset.

CV(IBI) and LvR (FR detector). With the FR detector, the LvR followed a different pattern from the Tampere one, underscoring that IBI-regularity metrics are sensitive to the choice of detection algorithm; this dependence is examined further in the statistical analysis (Section 4.2.4) and in the parameter-sensitivity experiments (Section 4.4).

4.2.4 Statistical Analysis

The NB metrics were assessed statistically in two ways: first across the four cell lines with the Tampere V1 detector (Analysis 1; Section 3.5.1), and then between the Tampere V1 and FR detectors within each cell line (Analysis 2; Section 3.5.2), in both cases with per-

DIV pairwise comparisons corrected with the Benjamini-Hochberg FDR, complemented by an LME model (detector or cell line as fixed effect, plate as a random intercept) for those metrics for which one was fitted (Sections 3.5.1 and 3.5.2). The statistical results for NB rate, NB duration, CV(IBM) and LvR are summarised in the pairwise significance heat-maps of Figure 4.8 and Figure 4.9.




Figure 4.8. Detailed Tampere-vs-FR pairwise significance heat-maps (panels A-H) for the NB metrics across cell lines and DIVs (Analysis 2; per-DIV pairwise tests with FDR correction). Coloured cells mark statistically significant detector differences and their direction. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.




Figure 4.9. Per-DIV statistical comparison of the NB metrics (NBR rate, NBD duration, CV(IBI), LvR) between the Tampere and FR detectors within each cell line (Analysis 2; per-DIV pairwise tests with Benjamini-Hochberg FDR correction), shown as a significance heat-map. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

Cross-cell-line comparison (Analysis 1). Across the four cell lines, the NBR computed with the Tampere V1 detector showed a developmentally varying pattern. The strongest and most stable effect was $DD5A < TUBA$, significant from DIV 60 onward; DD5A also fell significantly below C10902 at several time points. DD1C differed significantly from C10902 at DIV 53 ($DD1C < C10902$, $p < 0.001$) and showed rare significant differences from DD5A and TUBA at intermediate time points. By the latest time points several pairwise differences had dropped out of significance, reflecting the convergence and increased dispersion of the NBR trajectories at DIV 83-90.

Tampere versus FR detector (Analysis 2). Shifting to the detector comparison, the Tampere V1 and FR detectors were compared within each cell line (Analysis 2; Section 3.5.2), using an LME model with detector as fixed effect and plate as a random intercept, with per-DIV

pairwise tests and FDR correction (Figure 4.8 and Figure 4.9; Tampere V1 in red, FR in blue). The detector effect on NBR was significant in three of the four cell lines but did not point the same way in all of them. In DD1C and TUBA the Tampere detector returned a significantly higher NBR than the FR detector (DD1C: $p < 0.001$, Cohen's $d = 0.55$; TUBA: $p < 0.001$, $d = 0.43$). In DD5A the two detectors did not differ significantly (ns), and in C10902 the relationship reversed, the FR detector returning the higher rate ($p < 0.05$, $d = -0.35$). The per-DIV comparison localises the effect to the earlier time points of DD1C (DIV 53-67, Tampere $>$ FR), with only isolated significant cells elsewhere. The statistical analysis was extended to all NB metrics (NB duration, CV(IBM) and LVR) for both detectors, summarised in the same heat-maps.

Overall, the two detectors broadly agree on the disease–control ordering of the cell lines, and for NB duration, CV(IBM) and LVR the lines follow the clinical severity with both; for NBR, however, the relative position of the two DS lines depends on both the detector and the time point. The detectors also disagree on the absolute values of the metrics: the Tampere V1 detector systematically reports more and shorter network bursts than the FR detector, so the choice of detector affects the magnitude of the NB metrics while leaving the disease–control comparison qualitatively unchanged.

4.3 Fragmentation Analysis

This section examines NB fragmentation, a phenotype reported in several patient-derived network models (Van Hugte et al., 2023; Doorn et al., 2024). Fragmentation was quantified in two complementary ways: the total fragment count per well (the absolute number of fragments detected across all NBs of a recording) and the fragments-per-NB ratio (FragPerNB; the mean number of fragments per NB, computed over NBs containing at least one detected peak). Both were computed for the Tampere and the FR detectors.

Direct comparison per DIV. To move from line-level trends to a direct, pooled algorithm comparison, the total fragment count (Figure 4.10) and the fragments-per-NB ratio (Figure 4.11) were compared between the two detectors at each time point with well-level boxplots and pairwise significance brackets (Tampere in red, FR in blue).




Figure 4.10. Total fragments per well, Tampere (red) vs FR (blue), by cell line and DIV (boxplots; one data point = one well). Pairwise comparisons used an adaptive test (Shapiro-Wilk normality check, then a Student's/Welch t-test or Mann-Whitney U test) with Benjamini-Hochberg FDR correction. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.




Figure 4.11. Fragments per NB (FragPerNB), Tampere (red) vs FR (blue), by cell line and DIV (boxplots; one data point = one well). Pairwise comparisons used an adaptive test (Shapiro-Wilk normality check, then a Student's/Welch t-test or Mann-Whitney U test) with Benjamini-Hochberg FDR correction. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

4.3.1 Raster Plots with Smoothed Firing Rate

To visualise how fragments are detected, representative recordings are shown with the spike raster overlaid on the z-scored smoothed population firing rate (blue dashed line), together with the detected NBs (blue rectangles) and the fragment peaks within them (magenta dots). The Tampere and FR detectors appear side by side for the same well (Figure 4.12), with the number of NBs in the displayed window and the mean fragment count annotated.




Figure 4.12. Smoothed-firing-rate rasters: spike raster overlaid on the z-scored smoothed population firing rate (blue dashed line), with detected NBs (blue rectangles) and fragment peaks (magenta dots), for the (A) Tampere and (B) FR detectors on the same well.

4.3.2 Fragment Occurrence and Count

The developmental trajectories of the total fragment count per well and the fragments-per-NB ratio (FragPerNB), for both detectors, are shown in Figure 4.13. With the Tampere detector the total fragment count per well followed line-specific temporal patterns: DD1C traced a U-shaped trajectory with a minimum around DIV 67, while DD5A showed a less pronounced, oppositely oriented trend; the two control lines stayed comparatively stable. With the FR detector all four lines increased their total fragment count over development, with the controls reaching higher counts than the disease lines, the opposite ordering to what a disease-specific fragmentation phenotype would predict, a first sign that the absolute fragment count is strongly detector-dependent.

The FragPerNB ratio computed with the Tampere detector showed the expected disease-related pattern: DD1C had a higher ratio than the other three lines at every time point, while the other three stayed close to 1 (essentially unfragmented NBs). This is the single clearest fragmentation signal in the dataset and it isolates the severe DS line DD1C. With the FR detector the FragPerNB ratio behaved differently, with the control line TUBA showing elevated fragment ratios across the first three time points. The per-DIV statistical comparison of the fragmentation metrics between the two detectors is reported in the statistical analysis (Section 4.3.4).




Figure 4.13. Developmental trajectories of the total fragment count per well and the fragments-per-NB ratio (FragPerNB) across DIVs for the four cell lines, Tampere and FR detectors (mean $\pm$ 95% CI).

4.3.3 Fragment Overlap between Algorithms

To quantify how far the two detectors agree at the level of individual fragments, a fragment-overlap analysis was performed: a fragment detected by one algorithm counted as shared if a fragment from the other algorithm fell within a tolerance window of ± 0.05 s. The result is summarised as a Venn diagram per cell line (Figure 4.14). The diagrams shown here are for DIV 83 (the same analysis was run at every time point), reporting, for each line, the number of fragments unique to Tampere, unique to the FR detector, and shared between the two.




Figure 4.14. Fragment overlap between detectors at DIV 83 (tolerance ± 0.05 s), shown as Venn diagrams per cell line: number of fragments unique to Tampere, unique to the FR detector, and shared.

For the controls, the two detectors were largely picking up the same fragments: for TUBA the great majority of fragments were shared, and for C10902 almost every Tampere fragment coincided with an FR fragment, as expected from the well-defined, regular NBs of the controls. The DS lines agreed less, and in DD1C in particular there was a much larger number of FR-unique fragments, showing that the two detectors diverge most precisely where the network activity is most disrupted. This points clearly to fragment-level quantification being sensitive to which detection algorithm is used, particularly in the disease lines.

4.3.4 Statistical Analysis

The fragmentation phenotype was assessed statistically in two ways: first across the four cell lines with the Tampere V1 detector (Analysis 1, Section 3.5.1), and then between the Tampere V1 and FR detectors (Analysis 2, Section 3.5.2). Both analyses used an LME model with the relevant grouping factor as fixed effect and plate as a random intercept, followed by per-DIV pairwise comparisons with Benjamini-Hochberg FDR correction; the Analysis 2 results are summarised in the pairwise significance heat-map of Figure 4.15.




Figure 4.15. Per-DIV statistical comparison of the fragmentation metrics (total fragment count and FragPerNB) between the Tampere and FR detectors within each cell line (Analysis 2; per-DIV pairwise tests with Benjamini-Hochberg FDR correction), shown as a significance heat-map. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

Cross-cell-line comparison (Analysis 1). Across the four cell lines, the fragments-per-NB ratio computed with the Tampere V1 detector showed that the severe line DD1C had a significantly higher fragmentation ratio than the others, mostly at the earlier time points: DD1C > C10902 ($p < 0.001$ at DIV 53, 60, 67), DD1C > TUBA (significant at DIV 53, 67, 83), and DD1C > DD5A (significant at DIV 60, 67). The two control lines also differed

from each other early on (TUBA > C10902 at DIV 53-67). By DIV 83-90 most pairwise differences in FragPerNB had dropped. The early DD1C-dominated signal suggests that the phenotype is easiest to detect at the earlier time points.

Tampere versus FR detector (Analysis 2). Unlike the NBR, the FragPerNB ratio was statistically indistinguishable between the two detectors in three of the four lines (DD1C, DD5A, and C10902: ns; Figure 4.15 and Figure 4.11). The exception was the control line TUBA, where the pooled detector effect was significant ($p < 0.05$, Cohen's $d = -0.38$, FR slightly higher), although no individual per-DIV comparison reached significance. Across all lines, the two detectors return essentially the same average degree of internal fragmentation per detected NB, even though they disagree, in a line-dependent way, on how many NBs there are.

In summary, the fragments-per-NB ratio carries a genuine, statistically significant disease signal across cell lines (Analysis 1), while the absolute fragment counts and the line-level orderings depend strongly on the detection algorithm (Section 4.3.2 and the overlap analysis of Section 4.3.3). The FragPerNB ratio itself is essentially consistent between the two detectors (Analysis 2). On the strength of the cross-line signal of Analysis 1, NB fragmentation (FragPerNB) stands as a candidate phenotypic marker of the NB network phenotype, isolating the severe line DD1C. It is worth being precise about what Analysis 2 does and does not establish: it shows that the two detectors return essentially the same average value of the FragPerNB ratio (Figure 4.15), not that the between-line differences of Analysis 1 recur when the FR output is substituted for the Tampere V1 output. Because the cross-line comparison was run on the Tampere V1 detector alone, the detector-independence of the fragmentation phenotype itself was not tested here.

4.4 Algorithm Optimisation and Parameter Sensitivity

This section examines how robust the NB and fragmentation results are to two methodological choices and reports an independent spectral validation of the detector.

4.4.1 Tampere V1 vs V2 and the ISI-N Threshold Difference

The two variants of the Tampere detector differ only in how they pick the ISI-N threshold on the log-ISI-N histogram: V1 takes the first local minimum after the global maximum, whereas V2 takes the global minimum of the valley between the two leading prominent peaks (Chapter 3, Section 3.2.3). To make the difference concrete, the per-well log-ISI-N histograms are shown with the V1 (red) and V2 (green) thresholds overlaid (Figure 4.16). In many wells the two thresholds nearly coincide; in others they diverge substantially, for instance when a secondary peak or a smoothing ripple creates an early local minimum that V1 selects while V2 places the threshold in the deeper valley between the two dominant modes. This is the mechanism by which the two variants can produce different burst boundaries on the same recording.




Figure 4.16. Per-well log-ISI-N histograms with the V1 threshold (red) and V2 threshold (green) overlaid. In many wells the two thresholds nearly coincide; in others they diverge, the mechanism by which the two variants can produce different burst boundaries on the same recording.

The effect of this threshold choice on the detected metrics is shown in Figure 4.17, which overlays the V1 and V2 developmental trajectories of NB rate and fragment rate for the four cell lines. The two variants stay closely matched across lines and time points; their formal statistical comparison (Analysis 3) is reported in Section 4.4.5.




Figure 4.17. Tampere V1 (red) vs V2 (green) developmental trajectories of NB rate and fragment rate for the four cell lines (mean $\pm$ 95% CI).

4.4.2 Parameter Set Comparison

How robust the results are to the detection parameters was tested by re-running the full pipeline with the four parameter sets defined in Chapter 3 (Set 1 baseline, Set 2 reference, Set 3 strict, Set 4 strictest; Tables 3.3-3.5). For each cell line, the four sets were compared on the core metrics (NBR for both detectors, NBD, CV(IBM), and FragPerNB) using the same LME framework, with pairwise tests between sets where an overall difference appeared (Figure 4.18). Across the four cell lines, NBR (both detectors), NBD and CV(IBM) were essentially unaffected by the parameter set, with no significant omnibus effect in any line (the sole exception a marginal NBD effect in DD1C, $p = 0.044$, with no pairwise contrast surviving FDR correction).

The one responsive metric was the fragmentation ratio (FragPerNB), which differed across sets in DD1C ($p = 3.4 \times 10^{-6}$) and TUBA ($p = 1.0 \times 10^{-5}$). These differences were driven entirely by Set 1: in both lines Set 1 differed from Sets 2 and 4, whereas Sets 2–4 were mutually indistinguishable (Figure 4.18). As Set 1 retains the historical baseline kernel (size 6 bins, $\sigma = 3$, violating the 6- $\sigma + 1$ rule; Table 3.4), this is expected: an under-sized smoothing kernel reshapes the sub-peaks on which fragmentation is computed. Excluding Set 1, the results are robust within the explored range. Overall, the detection pipeline proved robust to the choice of parameter set: none of the NB metrics (NBR, NBD, CV(IBM)) changed significantly across sets in any line, and the only sensitivity — on the fragmentation ratio (FragPerNB) — was confined to the legacy Set 1 kernel, while the three principled sets (Sets 2–4) gave statistically indistinguishable results, supporting Set 2 as the reference configuration.




Figure 4.18. Parameter-set comparison for the four cell lines. (A–D) Boxplots of the core metrics (NB rate, NB duration, CV(IBM), FragPerNB) across the four parameter sets (Set 1 baseline, Set 2

reference, Set 3 strict, Set 4 strictest; one data point = one well; IQR $k = 1.5$) for DD1C, DD5A, TUBA, and C10902 respectively, compared with an LME model. (E) Pairwise tests with Benjamini-Hochberg FDR correction. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

4.4.3 Frequency-Band Validation

The Tampere V1 NB and fragment detector was validated independently with the adapted Cimbalnik et al. (2018a, 2018b) spectral framework (Chapter 3, Section 3.4). Each spectral event detected on the raw filtered signal was assigned to one of three classes by its temporal relationship to the spike-based detections: inside a fragment window (IN-FRAG), inside a NB but outside any fragment (IN-NB-OUT), or outside any NB (OUT-NB). A correct detector predicts a monotonic ordering $\text{IN-FRAG} > \text{IN-NB-OUT} > \text{OUT-NB}$ in both the strength and the rate of high-frequency events.

Event strength. The peak z-score of the detected events followed the expected ordering, most clearly in the High Gamma band: events inside fragments reached the highest z-scores, those in the within-NB background were intermediate, and out-of-NB events were lowest, with all three classes lying above the $z = 3$ detection threshold by construction. The amplitude distribution of IN-FRAG events showed a pronounced long tail toward high z-scores, echoing the long-tailed pathological-HFO distributions reported by Cimbalnik et al. (2018a) for epileptogenic tissue, reinterpreted here as reflecting the elevated synchrony of NB sub-units rather than any pathological process. The High Gamma band is the most informative of the three for this validation: in dissociated cultures genuine spontaneous high-frequency network synchronisation is concentrated in the gamma range, whereas true ripple and fast-ripple oscillations *in vitro* are largely stimulus-evoked rather than spontaneously generated (Hales et al., 2012), so the High Gamma band carries most of the physiological signal while the higher bands contribute comparatively little (Figure 4.19).




Figure 4.19. Peak z-score of spectral events by class (IN-FRAG, IN-NB-OUT, OUT-NB) for the High Gamma, Ripple and Fast Ripple bands (Tampere V1). (A) High Gamma, (B) Ripple, (C) Fast Ripple. Events inside fragments reach the highest z-scores, within-NB background events are intermediate, and out-of-NB events are lowest; the separation is clearest in the High Gamma band. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

Mean envelope amplitude. The mean Hilbert-envelope amplitude, the key feature of the Cimbalnik framework, likewise followed the IN-FRAG $\geq$ IN-NB-OUT $>$ OUT-NB ordering, most clearly in the High Gamma band. This is the measure that links most directly to Cimbalnik et al. (2018a), where envelope amplitude separates pathological from physiological HFOs; here it serves as an index of network-synchronisation intensity, not as a marker of epileptogenicity. Taken together across these measures, the High Gamma band (65–100 Hz) carried by far the clearest separation between the three event classes and is therefore the most informative band for this validation; in these dissociated cultures the ripple and fast-ripple bands, dominated by noise rather than genuine network oscillations, added little discriminative power (Figure 4.20).




Figure 4.20. Mean Hilbert-envelope amplitude of spectral events by class for the High Gamma, Ripple and Fast Ripple bands (Tampere V1). (A) High Gamma, (B) Ripple, (C) Fast Ripple. Amplitude follows the IN-FRAG $\geq$ IN-NB-OUT $>$ OUT-NB ordering, most clearly in the High Gamma band. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns, not significant.

4.4.4 IN-NB and OUT-NB Event Classification

The hit-rate, the percentage of NBs (or fragments) containing at least one high-gamma event, measures how reliably the detected events coincide with spike-based detections, and how that coincidence develops with network maturation. The NB hit-rate rose markedly with DIV, exceeding 75% at the mature time points for the control lines and approaching 100% by the latest ones, confirming that the Tampere V1 detector reliably identifies windows of elevated broadband activity. The fragment hit-rate was lower by construction: each fragment window spans only about 50 ms, far less than a whole NB, so its chance of containing a high-gamma event is correspondingly smaller, but it consistently beat chance and rose with maturation, mirroring the developmental trajectory of network connectivity. As in the event-rate analysis, DD1C showed the lowest hit-rates, in keeping with its sparse, weakly synchronised activity.

In representative NBs (for example in the control line C10902), the peak of the cross-electrode multi-unit activity (MUA) envelope coincided with the spike-based NB window, and the fragment peak times fell on sub-peaks of that envelope, supporting the idea that

fragments pinpoint the most intense synchronisation windows within a burst. High-frequency events detected independently by the spectral pipeline clustered inside the NB window and were densest around the fragment peaks (IN-FRAG) relative to the sparse out-of-NB events (OUT-NB), in line with the quantitative results above (Figure 4.21).




Figure 4.21. Illustrative single network bursts for the four cell lines (Tampere V1). (A) DD1C, (B) DD5A, (C) TUBA, (D) C10902. Each panel shows, from top to bottom, the cross-electrode multi-unit activity (MUA) envelope, the High Gamma, Ripple and Fast Ripple band traces, and the per-event peak z-score heat-map; the spike-based NB window and the fragment peaks align with the maxima of the MUA envelope and of the high-frequency activity.

The frequency-band validation backs this up from a largely independent angle: without relying on the spike-based detector at all, it points to the same windows of high-frequency network synchronisation that V1 flags, and shows that fragments correspond to sub-units of especially intense synchronisation within NBs. As discussed in Chapter 3 (Section 3.4), this is a concurrent validation rather than a gold-standard one: it confirms temporal concordance between two independent measures of coordinated firing, but does not by itself validate the absolute NB or duration estimates.

4.4.5 Statistical Analysis

The two Tampere variants (V1 and V2) were compared statistically within each cell line (Analysis 3, Section 3.5.3), using an LME model with detector variant as fixed effect and plate as a random intercept, for NBR and FragPerNB, with per-DIV pairwise tests and FDR correction (the descriptive trajectories are shown in Figure 4.17; V1 in red, V2 in green).

Fragments per NB. Descriptively, V1 tended to return a higher fragment ratio than V2, visible for the severe line DD1C, where the V1 trajectory lies above V2 with a peak around DIV 60. This fits the design of the two variants, V2 being the more conservative criterion. The LME analysis, however, found no significant difference between V1 and V2 in FragPerNB in any cell line at any time point: once between-plate variability is accounted for, the apparent V1-V2 gap does not reach significance.

Network burst rate. The NBR likewise did not differ significantly between V1 and V2 in any cell line at any time point. Overall, despite the conceptual gap between V1 and V2 and the per-well threshold divergences visible in Figure 4.16, the choice of variant does not materially change the quantitative conclusions, while V2 keeps the practical advantage of being more robust to smoothing-induced ripples (Chapter 3).

The detailed V1-versus-V2 statistics (Analysis 3) confirm that neither NBR nor FragPerNB differed significantly between the two variants in any cell line at any time point. For the parameter-set comparison (Section 4.4.2), NBR, NBD and CV(IBM) were unaffected by the parameter set in all lines, the only responsive metric being FragPerNB in DD1C and TUBA, where the difference was driven entirely by the baseline Set 1 (Sets 2-4 mutually indistinguishable); excluding Set 1, the main results are robust to the detection parameters within the explored range. The frequency-band validation (Sections 4.4.3-4.4.4) is descriptive and convergent in nature: every metric checked (event strength, rate, envelope amplitude, hit-rate) lined up in the expected order, IN-FRAG on top, then IN-NB-OUT, then OUT-NB.

Spectral-event classes (A4). The three spectral-event classes (IN-FRAG, IN-NB-OUT, OUT-NB) were compared within each frequency band on the cell-line $\times$ DIV block medians (Analysis 4, Section 3.5.4), using a Friedman omnibus test followed by Holm-corrected pairwise Wilcoxon signed-rank tests. For the peak z-score, the class effect was highly significant in the High Gamma band (Friedman $\chi^2 = 48.7$, $p = 2.7 \times 10^{-11}$, $n = 40$ blocks), with all three pairwise comparisons surviving Holm correction (IN-FRAG > IN-

NB-OUT > OUT-NB; all Holm-adjusted $p < 10^{-4}$); the Fast Ripple band was also significant ($\chi^2 = 21.9$, $p = 1.8 \times 10^{-5}$, $n = 22$; all pairs $p < 0.05$), whereas the Ripple band showed no class effect ($\chi^2 = 2.5$, $p = 0.29$, $n = 13$). For the mean Hilbert-envelope amplitude, the High Gamma band again showed a significant class effect ($\chi^2 = 20.2$, $p = 4.2 \times 10^{-5}$, $n = 40$ blocks), with all three pairwise comparisons significant after Holm correction (Holm-adjusted $p \leq 0.004$); the Fast Ripple omnibus was significant ($\chi^2 = 9.8$, $p = 7.4 \times 10^{-3}$, $n = 22$) but no individual pair survived correction, and the Ripple band showed no effect ($\chi^2 = 0.5$, $p = 0.79$, $n = 13$).

Taken together, the optimisation analyses indicate that the core results are robust to the methodological choices examined, the V1/V2 threshold rule and the parameter set leave the NB metrics essentially unchanged, while the independent spectral (Cimbalnik) validation confirms, without relying on spike detection, that the detected fragments mark sub-units of genuine high-frequency network synchronisation.

Chapter 5

Discussion

The work reported here set out to investigate NB fragmentation in heterogeneous hiPSC-derived cortical networks recorded on MEAs, and to examine how its quantification depends on NB detection, using the same four-cell-line longitudinal dataset established for this DS model system (Mzezewa et al., 2025): the severe SCN1A line DD1C, the milder line DD5A, and two unrelated control lines (TUBA, C10902). Three questions organised the analysis: whether the four lines separate along disease-relevant activity metrics, whether the choice of NB detector alters that picture, and whether the NB fragmentation metric behaves as a candidate phenotypic marker.

The specific contribution made here is twofold: the fragmented-NB phenotype, so far characterised mainly in excitatory patient-derived networks, was examined in heterogeneous cortical networks from DS patients, and the NeuroGroup (Tampere) pipeline was extended with an integrated fragmentation analysis whose output was compared with that of a widely used firing-rate-based detector, across parameter sets, and against a spectral measure computed independently of spike detection. Fragmentation as a phenomenon is well established (Van Hugte et al., 2023; Doorn et al., 2024); what this thesis examines is how its quantification behaves when the NBs are delimited in different ways.

5.1 A Delayed, Hypoactive Network Phenotype in Severe DS

DS is a severe, SCN1A-linked developmental and epileptic encephalopathy, and hiPSC-derived neuronal networks are increasingly used to characterise its cellular and network-level phenotype in vitro, complementing the murine models that have dominated the field to date (Higurashi et al., 2013; Sun et al., 2016). MEA recordings of such networks provide

a longitudinal, non-invasive readout of population activity, and the four-line panel analysed here is the same one used to first characterise the functional phenotype of this DS model system (Mzezewa et al., 2025). Among all the spontaneous-activity measures analysed, the MFR distinguished the genotypes most sharply.

The severe line DD1C showed significantly lower firing than every other line at almost all time points, and the ordering that emerged, DD1C below DD5A below the two control lines, mirrors the clinical severity gradient of the donors. A pronounced reduction in baseline activity is consistent with the cellular logic of SCN1A haploinsufficiency, since NaV1.1 is preferentially expressed in GABAergic interneurons and the loss of one functional allele should attenuate inhibitory drive (Yu et al., 2006; Ogiwara et al., 2007). The network-level effect is not simply disinhibition; it is a flatter, less synchronised regime, observable in the disease lines, corresponding to the interneuron-centred channel dysfunction described in hiPSC DS models (Higurashi et al., 2013; Sun et al., 2016). The remaining metrics, namely the SEB rate, the %RS, and network synchrony, all follow the same disease-control split, which reinforces the significance of the gradient.

The longitudinal design also allowed the developmental trajectory of each line to be tracked directly. This proved as informative as the between-genotype comparison itself: as the cultures matured, activity in the control lines increased and consolidated more markedly than in the disease lines. TUBA and C10902 progressed toward more regular and synchronised firing, consistent with the developmental profile previously reported for these lines (Mzezewa et al., 2025), while the disease lines, DD1C in particular, lagged behind and settled onto a lower plateau. This pattern also matches the classical maturation signature described for dissociated and hiPSC-derived cultures on MEAs, a staged transition from sparse, uncorrelated spiking toward organised network activity (van Pelt et al., 2004; Wagenaar et al., 2006; Odawara et al., 2016).

Taken together, these observations suggest that DS networks are not simply hypoactive, but instead follow a delayed developmental trajectory. This reading is reinforced by a further finding: the sharpest between-line contrasts appeared at the early and intermediate time points and then gradually narrowed later in the recording window. The convergence of the trajectories, together with the later widening of their dispersion, explains why the phenotype behaves as a windowed phenomenon, most evident while the control networks are still maturing ahead of the disease lines (Mzezewa et al., 2025).

5.2 NB Rate and Duration Depend on the Detector

The comparison between the Tampere detector and the FR detector showed that results vary with both the metric considered and the detector employed, with points of agreement and disagreement across measures. On the overall disease–control ordering of the lines the two detectors broadly agreed, and for NB duration, CV(IBM) and LvR the lines followed the clinical severity with both. For NBR, however, the relative position of the two DS lines depended on the detector: DD1C fell below DD5A, as expected from the clinical severity, only at the first time point with the Tampere detector and only after DIV 67 with the FR detector. On the absolute values of NBR and NBD the two detectors did not agree. The Tampere V1 detector, which thresholds the log-ISI histogram following the self-adapting logic of Pasquale et al. (2010), systematically reported more numerous and shorter bursts than the FR detector, and the direction of the detector effect even reversed across lines.

This divergence reflects a structural difference between the two algorithmic families rather than a flaw specific to either implementation. ISI-based detectors read burst structure from the statistical properties of individual spike trains and are largely insensitive to global firing-rate differences, whereas firing-rate-based detectors such as the FR detector read a population-level envelope and are comparatively more exposed to noise and to the choice of amplitude normalisation (Bakkum et al., 2014; Chiappalone et al., 2005, 2006). Systematic comparisons of burst-detection methods have repeatedly shown that both the count and the boundaries of the detected events depend strongly on the algorithm and its parameters (Cotterill et al., 2016), and the present results reproduce this pattern at the network-burst level.

Neither family is intrinsically more biologically meaningful. The log-ISI approach is well suited when firing rates vary widely across development or genotype, as in this longitudinal dataset, while firing-rate-based detection more directly captures population-level envelope dynamics and is the natural substrate for the sub-peak fragment criterion used by Doorn et al. (2024) and in this work. It is for this reason, rather than because one algorithm is more valid than the other, that the Tampere detector was retained as the primary NB definition for the cross-line comparisons in this thesis, while the FR detector was used to examine

how the fragmentation readout changes with a different algorithmic family. The results must therefore always be read together with the detector that produced them.

Two further considerations qualify the comparison. First, the thresholds of the FR detector are data-driven: in the original work they were chosen for excitatory networks and a different recording and spike-detection setup (Doorn et al., 2024), whereas here they were kept fixed, with only the smoothing kernel and the moving window varied across parameter sets (Table 3.5). Part of the disagreement between the two detectors may therefore reflect thresholds that are not optimal for heterogeneous cortical networks recorded on the Axion platform, and a fair comparison would require re-tuning them for this dataset. Second, neither detector can be declared closer to reality on the present data, because no ground truth was available. Deciding which quantification better reflects the actual burst and fragment structure would require recordings annotated by expert observers, against which each detector, and each parameter setting, could be scored.

5.3 The Fragmentation Phenotype and Its Dependence on the Detector

This work implemented the fragmented-NB analysis within the NeuroGroup (Tampere) pipeline and quantified fragmentation across all four studied lines. The concept itself draws on a growing body of hiPSC-MEA work: fragmented NBs emerged as the main phenotypic driver in excitatory networks from GEFS+, DS, and febrile-seizure patients carrying SCN1A variants (Van Hugte et al., 2023), and comparable within-burst fluctuations, sometimes termed reverberating bursts, mini-bursts, or superbusts depending on the study, have been reported in models of Kleefstra syndrome (Frega et al., 2019), Kabuki syndrome (Gabriele et al., 2021), and Rett syndrome (Pradeepan et al., 2024). Doorn et al. (2024) synthesised this evidence and, using a biophysically detailed in silico model together with pharmacological validation, linked fragmentation mechanistically to enhanced short-term synaptic depression, with asynchronous neurotransmitter release and NMDA-receptor currents identified as contributing factors; blockers of NMDA receptors and of asynchronous release had independently been shown to suppress fragments in the Kleefstra and Rett models, respectively (Frega et al., 2019; Pradeepan et al., 2024). The present

results are broadly consistent with this picture: DD1C, the severe SCN1A line, fragmented its bursts more than any other line, and the signal was strongest on the earliest recording days. What this thesis adds is the observation of the phenotype in heterogeneous cortical networks from DS patients, together with a comparison of how two detection approaches, several parameter sets and a spike-independent spectral measure quantify it.

The key methodological finding is that not all fragmentation-related quantities are equally robust. FragPerNB, the mean number of fragments per NB, conveyed a genuine cross-line signal and was statistically indistinguishable between the Tampere and FR detectors in three of the four lines, whereas the absolute fragment count diverged between detectors in every line. This is expected: the two algorithms partition the same recording into different burst ensembles even though the underlying fragment-detection routine is identical, so any count that is not normalised by the number of fragmented NBs inherits the detector's own segmentation choices. The ratio, by contrast, is comparatively insensitive to how many NBs are called in the first place, which is why, in this dataset, it was less sensitive than the raw count to the choice of detector; whether this holds for other detectors and datasets was not tested here.

It is important, however, to be precise about what this detector-agreement result does and does not show. The finding that FragPerNB values are similar between detectors was established at the level of the ratio itself, pooled across the dataset; the cross-line, DD1C-specific phenotype was instead formally tested only on the Tampere V1 output. Whether the same DD1C-versus-control contrast recurs, and reaches significance, under the FR detector was not directly tested here, so detector-independence of the disease phenotype, as distinct from detector-agreement of the underlying ratio, remains an open question. A fragmentation phenotype reported without its detection method would be ambiguous; reported as a ratio, and with the method stated, it is less ambiguous, and it is in this qualified sense that fragmentation qualifies here as a candidate, and not an established, marker of the severe DS network phenotype. This distinction carries a direct implication for reproducibility: a future study switching to a different NB detector should not assume that the DD1C-versus-control contrast itself will reproduce merely because the FragPerNB ratio does: detector-agreement of the metric and detector-independence of the disease phenotype are not the same claim.

A mechanistic interpretation can be drawn from previous work. In purely excitatory SCN1A networks, fragmentation emerges without any interneurons (Van Hugte et al.,

2023), and it requires enhanced short-term synaptic depression together with persistent excitation sustained by asynchronous release and NMDA-receptor currents (Doorn et al., 2024). The same synaptic mechanisms could underlie the stronger fragmentation of DD1C in the heterogeneous networks studied here, possibly modulated by the altered inhibitory drive expected from NaV1.1 loss of function. This explanation is plausible and consistent with the direction of the present results, but it should be stated as an interpretation rather than as a direct finding: synaptic dynamics and E/I balance were not measured in this work, and the fragmentation phenotype was inferred purely from spike-timing statistics.

A related confound deserves explicit treatment. DD1C is simultaneously the least active and the most fragmented line, and the sub-peak criterion used to split an NB into fragments applies thresholds defined relative to the peak firing rate of that same NB, so the number of fragments counted in a sparser, more irregular population could be inflated relative to a denser one, a rate-dependent quantification bias of the kind systematically documented across burst-detection methods, including for low-activity spike trains, in independent benchmarking work (Cotterill et al., 2016). An exploratory check using the frequency-band validation data, in which the density of detected high-frequency events inside NBs served as a proxy for local firing intensity, found a consistently negative association with FragPerNB in every line, controls included. This analysis rests on a proxy variable rather than on spike trains and is not conclusive, but it means that a contribution of firing sparsity to the magnitude of the fragmentation readout cannot be excluded on the present data; disentangling the two would require a rate-matched or surrogate-spike-train control.

The fragmentation phenotype was also more clearly expressed at the earliest recording days than later in the developmental window, mirroring the pattern already noted for the spontaneous-activity metrics (Section 5.1). The most cautious reading is that fragmentation is easier to detect against a background of low, still-developing activity, rather than that the underlying network pathology itself fades with maturation; disentangling these two possibilities would require a design that explicitly decouples chronological age from network maturational stage. Taken together, the fragmentation data support FragPerNB as a promising, detection-sensitive candidate readout of the DS network phenotype, whose between-line contrast should be confirmed with dataset-specific detection settings, expert-validated fragment annotation, and a larger, ideally isogenic cohort.

5.4 Robustness to Parameters and Spectral Correlates of Fragments

Two further checks were carried out to test how far the fragmentation results depend on incidental methodological choices: whether the specific threshold rule used within the Tampere log-ISI detector matters, and whether the results hinge on a particular smoothing-kernel parameter set.

The comparison between the two Tampere threshold variants, V1 and V2, showed no significant difference in either NBR or FragPerNB in any line. This suggests that the conceptual distinction between selecting the first local minimum and selecting the deepest valley of the log-ISI histogram does not affect the quantitative results, although V2 offers a practical, by-design advantage in resisting smoothing-induced ripples, even though this did not translate into a measurable difference in NB metrics here. Re-running the entire pipeline with four different parameter sets likewise produced similar results. NB metrics remained largely unchanged across all sets, and the only responsive quantity, FragPerNB in DD1C and TUBA, shifted only because of the legacy Set 1 kernel, whose undersized smoothing window distorts the sub-peaks used to compute fragmentation. Once this legacy setting is set aside, the three principled sets (Sets 2–4) yield statistically comparable results, which supports Set 2 as the reference parameterisation and indicates that the headline findings do not hinge on a specific choice of smoothing parameters.

The most independent line of evidence came from the spectral validation, which addresses the same recordings from an entirely different direction, using a framework based on spectral events rather than spike trains. Adapting the high-frequency-oscillation (HFO) framework of Cimbalka et al. (2018a, 2018b), spectral events were detected on the filtered raw signal and classified according to their position relative to the spike-based NB and fragment boundaries. In the High Gamma band, every measure examined, including peak z-score and envelope amplitude, followed the predicted order: strongest and densest inside the fragments, intermediate across the rest of the NB, and weakest outside any burst. The higher bands added little, the Ripple band showing no class effect at all; this is consistent with previous work showing that genuine spontaneous high-frequency synchronisation in dissociated cultures concentrates in the gamma range, whereas true ripple and fast-ripple

oscillations *in vitro* are largely stimulus-evoked rather than spontaneously generated (Hales et al., 2012; Buzsáki & da Silva, 2012).

The envelope-amplitude feature, originally developed to flag pathological high-frequency events in clinical intracranial recordings (Cimbalnik et al., 2018a), is repurposed here to index the intensity of physiological network synchronisation *in vitro*. This is a deliberate reinterpretation of the metric, not a claim of pathology in the recorded material, since the present dataset consists of dissociated healthy and patient-derived hiPSC-neuronal networks rather than clinical or *ex vivo* tissue samples. The validation this analysis provides is concurrent rather than gold-standard: it demonstrates that two independent measures of coordinated firing, spike-based and spectral, occur within the same time windows, without certifying the absolute burst or fragment boundaries produced by either method. Within these constraints, the spectral pipeline, computed independently of spike detection, supports the interpretation that the fragments identified by the spike-based analysis coincide with sub-units of especially intense network synchronisation.

5.5 Convergent Evidence Supports NB Fragmentation as a DS Readout

Read together, the results converge on a central theme: the clearest disease signal lies in the MFR and the FragPerNB ratio, both of which isolate the severe line DD1C. Two NB detectors built on different principles return a comparable FragPerNB ratio in three of the four lines; two control lines of independent origin behave consistently relative to the patient lines; and a spectral measure that does not use the spike-based detector at all shows that fragments coincide with the most intense moments of network synchronisation. A convergence of independent lines of evidence of this kind is harder to explain through a single source of bias than any individual comparison (Munafò & Davey Smith, 2018). Together with the mechanistic evidence already available for fragmentation (Doorn et al., 2024), this supports NB fragmentation as a promising readout for *in vitro* network phenotyping of DS. At the same time, the number of fragments detected changed with the detection method, particularly in the patient-derived networks, so fragmentation values can

be compared across studies only when the detection settings have been adapted to each dataset and checked against expert inspection of the recordings.

More broadly, these results illustrate both the potential and the practical challenges of hiPSC-based disease modelling on MEAs. Patient-derived networks reproduce genotype-linked functional differences that murine models capture only partially, and this is the central rationale for the approach used throughout this work (Engle et al., 2018; Cerina et al., 2023). At the same time, the same recordings show how heavily the readout depends on the analytical choices applied to the raw signal, from the burst-detection algorithm to the smoothing kernel. The convergence-based strategy adopted here, testing each finding against an independent detector, parameter set, or modality, is one practical way to probe which network-level readouts are robust enough to serve as candidate phenotypic markers, although it cannot replace an expert-annotated ground truth.

Chapter 6

Conclusions

This thesis investigated network burst fragmentation, a network phenotype repeatedly observed in patient-derived neurons, in heterogeneous cortical networks derived from patients with Dravet syndrome, and examined how its measurement depends on the method used to identify the bursts in the first place.

Across four patient- and control-derived cell lines followed over several weeks of maturation, the number of fragments per network burst emerged, together with the mean firing rate, as the feature that most clearly isolated the network derived from the most severely affected patient, especially early in the recording period. The phenotype described in excitatory patient-derived networks is therefore also visible in heterogeneous cortical networks containing inhibitory neurons, and it follows the clinical severity of the donors. When the same recordings were analysed with two widely used detection approaches, the number of bursts and the absolute number of fragments differed between them, especially in the patient-derived networks, whereas the number of fragments per burst remained comparable in three of the four lines. A spectral analysis of the raw signal, independent of both detectors, showed that the detected fragments coincide with the most intense moments of network synchronisation.

The main message is therefore twofold. Biologically, burst fragmentation, expressed as fragments per burst, is a promising marker of Dravet syndrome severity in this model. Methodologically, how many fragments are counted depends on how the bursts are delimited: the between-line comparison rests on a single detector, and neither detector could be shown to be closer to reality, because no expert-annotated reference was available. Before fragmentation values are compared across detectors, laboratories or studies, detection thresholds should be tuned to each dataset and validated against expert inspection of the recordings.

6.1 Limitations

The results presented in this thesis should be read in the light of several limitations, most of which stem from the intrinsic constraints of in vitro disease modelling based on patient-derived cells.

The comparison analysed four hiPSC lines carrying different SCN1A mutations: a severe DS line (DD1C), a milder DS line (DD5A), and two control lines. With a single line per genotype, it is impossible to separate the effect of the mutation from donor- or clone-specific biology, since any difference attributed to SCN1A haploinsufficiency could in principle reflect the idiosyncrasies of a particular clone rather than the mutation itself (Engle et al., 2018). The study is therefore exploratory rather than confirmatory, and a definitive attribution would require isogenic controls, such as a gene-corrected counterpart of each patient line (Sun et al., 2016), which were not available in this work.

The two control lines derive from different sources and were not generated together with the patient lines, so the control baseline is not a single, internally consistent reference. Because the cultures were recorded longitudinally across DIV, maturational timing and genotype are partly confounded: a difference observed at a given time point could reflect a mismatch of the developmental clock between the lines (van Pelt et al., 2004; Odawara et al., 2016) rather than a genuine genotype effect. Following whole trajectories rather than single snapshots mitigates this confound, without eliminating it entirely.

NB fragmentation is treated throughout this work as a candidate rather than an established phenotypic marker, and its quantification is detector-dependent: the Tampere and FR detectors did not agree on the absolute number of fragments per NB in this dataset, consistent with the broader observation that burst-detection outputs are algorithm-sensitive (Cotterill et al., 2016), so the magnitude of any fragmentation phenotype reported here is sensitive to the detection pipeline used. Only the FragPerNB ratio proved comparable between the two detectors, and the between-line comparison itself was performed on the Tampere V1 output alone, so the detector-independence of the between-line contrast has not been established. The thresholds of the FR detector, moreover, were adopted from a different dataset and kept fixed rather than re-tuned for these recordings, and no expert-annotated reference was available to judge which detector better captures the actual burst and fragment structure. As the present dataset overlaps in part with that used in the original

characterisation of this model system, the limitations discussed by Mzezewa et al. (2025) regarding line number and donor variability apply equally here.

A further caveat concerns the fragmentation metric itself. DD1C is at once the least active and the most fragmented line, and the sub-peak criterion used to split an NB into fragments applies thresholds defined relative to the maximum of that NB, so fragment counts in a sparser population could be inflated. An exploratory check on the frequency-band validation data, using the density of detected high-frequency events inside NBs as a proxy for firing intensity, found a consistently negative association with FragPerNB within every line, controls included, consistent with the rate-dependent detection bias documented for burst-detection methods more generally (Cotterill et al., 2016). That analysis rests on proxy variables rather than on spike trains and is not conclusive, but it means a contribution of firing sparsity to the fragmentation readout cannot be excluded on the present data; a rate-matched or surrogate-train control would be required to separate the two. The magnitude of the fragmentation readout should therefore be interpreted with caution, and the between-line contrast confirmed in larger and ideally isogenic cohorts.

The number of plates (two or three per cell line, five in total) and the small number of lines limit statistical power, particularly for the mixed-effects models that include plate as a random intercept. Effect sizes (Cohen's *d*) are therefore reported alongside *p*-values throughout the text, since with a small sample the magnitude of an effect is more informative than its nominal significance (Cohen, 1988). The spectral validation, finally, is a concurrent validation and not a gold standard, as stated earlier: both the NB detector and the spectral detector reflect the same underlying phenomenon, coordinated neuronal firing, through different processing routes, so their agreement is encouraging but cannot certify that either one correctly captures the NB boundaries. The envelope-amplitude feature was used in a way it was not originally designed for, serving as an approximate index of synchronisation intensity in healthy and patient-derived neuronal networks *in vitro*, rather than as a clinical biomarker of pathological activity.

Two further features of the design constrain the statistical treatment of the data. The cell lines were not fully crossed with the MEA plates: TUBA was recorded only on plates N2894 and N2905 and C10902 only on N2912 and N2942, with no plate shared between them, so line identity and plate identity are partially non-identifiable and models adjusting for plate can return unstable estimates for that contrast. In addition, the overall between-line tests pool repeated recordings of the same well; although the measured design effect

is modest (1.00–2.04) and the time-point structure is balanced across the lines, the resulting p-values remain anti-conservative and are best read as summary indications rather than independent confirmatory evidence (Section 3.5.1 and Appendix A). The limitations reported by Mzezewa et al. (2025) for this dataset apply here as well: isogenic lines would be required to attribute the functional differences directly to the disease-causing variants, and a network-level readout cannot resolve which neuronal subtype drives the altered activity. Both issues would be resolved by a design in which every line is represented on every plate and independent differentiation batches, rather than wells, provide the biological replicates.

6.2 Future Perspectives

The relevance of this work extends beyond the four cell lines studied here. Patient-derived neuronal networks are increasingly used to model genetic epilepsies and neurodevelopmental disorders, because they carry the patient's own genetic background and can be studied directly in human cells, narrowing the translational gap that limits rodent models (Löscher, 2011). MEA recordings make these networks particularly attractive, since the same culture can be followed non-invasively over weeks and many cultures can be recorded in parallel (Hyvärinen et al., 2019; Kapucu et al., 2022). For such platforms to be used to compare laboratories, to stratify patients, or to test drugs, however, their readouts must be reliable. The central methodological finding of this thesis, that the apparent size of a disease phenotype can change with the burst-detection algorithm alone, therefore concerns any study that relies on network-burst analysis, and points to the need for standardised and cross-validated analysis pipelines (Cotterill et al., 2016).

Three broader directions of impact follow from this. First, a validated fragmentation readout could support drug testing in patient-derived networks. In DS, several commonly used antiseizure medications can aggravate seizures and even the approved therapies rarely abolish them (Wirrell et al., 2022), so a human in vitro assay able to show whether a candidate compound normalises network behaviour in cells from a given patient would be of direct clinical interest, in line with the pharmacological targets already identified by Doorn et al. (2024). Second, because the analysis developed here does not depend on the specific mutation, the same framework could be applied to other genetic epilepsies and

neurodevelopmental conditions in which altered network synchronisation is expected, providing a common and comparable readout across disease models. Third, on the methodological side, the dependence of fragment counts on the detection method shown here argues for reporting the detector and its parameters alongside any network-level metric, and for building shared, expert-annotated reference recordings, a prerequisite for comparing and pooling data across laboratories. The paragraphs below outline the concrete steps needed to move in these directions.

The most direct extension of this work is to close the genetic-attribution gap identified above and discussed for this same model system by Mzezewa et al. (2025). Pairing each patient line with a CRISPR-corrected isogenic counterpart would allow the fragmentation and firing-rate phenotypes to be tested against the SCN1A mutation itself, rather than against a confounded genetic background (Sun et al., 2016). Given that DS is caused by more than 1,800 distinct SCN1A variants (Ding et al., 2021), however, a purely mutation-specific isogenic strategy may not by itself be the most informative route. Van Hugte et al. (2023), studying the same GEFS⁺-to-DS clinical spectrum with unrelated healthy controls, found that the excitatory network phenotype was mutation-specific rather than shared uniformly across the spectrum, which argues for combining several complementary strategies: unrelated healthy controls, as used here, isogenic gene-corrected lines, and familial (parental or sibling) controls of the kind used by Xie et al. (2020), who paired an unaffected-sibling control with an isogenic corrected line within the same SCN1A/GEFS⁺ family to separate mutation effects from genetic background, so that the *in vitro* phenotype can be characterised both at the level of a specific variant and at the more general level of SCN1A haploinsufficiency, following best-practice recommendations for control design in hiPSC disease modelling (Engle et al., 2018). A design of this kind, expanding the panel to several clones per genotype, would also average out donor-specific idiosyncrasies and increase the statistical power that the current plate design necessarily limits.

Methodologically, the detector-dependence highlighted throughout this thesis suggests a more systematic approach for future work: rather than relying on a single algorithm, subsequent analyses could evaluate fragmentation using a panel of detectors, drawing on the broader family of NB-detection methods reviewed in Chapter 2 (Pasquale et al., 2010; Bakkum et al., 2014; Doorn et al., 2024), and synthesise their results explicitly. This would turn the convergence argument used informally in this thesis into an explicit, quantitative analysis criterion; standardised benchmarking of burst- and NB-detection methods on

shared datasets would support exactly this goal (Cotterill et al., 2016), provided that the thresholds of each detector are tuned to the data and that a subset of recordings is annotated by expert observers to serve as ground truth. Extending the longitudinal sampling and aligning recordings to a maturational rather than a chronological index, for instance anchoring each line to the point at which it first reaches stable NB activity, or to a composite network-maturation score, rather than to absolute DIV, would also help separate developmental timing from genotype, a confound that the trajectory analysis performed here only partly resolves.

The spectral validation also opens a route from concurrent to mechanistic evidence. The High Gamma signature that tracks the fragment sub-units in this work could be combined with pharmacological manipulation of synaptic dynamics, following the approach already used by Doorn et al. (2024), who showed that pharmacologically enhancing short-term synaptic depression (with the dynamin inhibitor Dynasore) was sufficient to induce fragmented NBs in healthy networks, and that NMDA-receptor and asynchronous-release blockers could suppress them. Applying a comparable pharmacological framework to the DS lines studied here would test directly whether the same synaptic mechanisms underlie the fragmentation phenotype observed in DD1C.

More broadly, complementary single-cell or transcriptomic phenotyping, combined with the electrophysiological readouts used here, is likely to be the most informative route forward: multimodal approaches of this kind have already proven valuable for linking network-level MEA phenotypes to cellular and molecular mechanisms in other hiPSC neurodevelopmental disease models (Verboven et al., 2026). Such an approach would link the network-level readouts developed here to the cellular deficits of SCN1A haploinsufficiency, and would help clarify whether NB fragmentation, once validated and standardised, could mature from a candidate marker into a quantitative assay for screening candidate anti-seizure medications (ASMs) in patient-derived networks, in the spirit of the pharmacological targets already proposed by Doorn et al. (2024).

A final direction concerns the fragmentation metric itself. Because FragPerNB behaved as the most robust disease-linked quantity, refining its definition deserves dedicated attention. Systematically varying the fragment window, the smoothing kernel, and the sub-peak criterion would map how strongly the fragment count depends on each of these choices. Coupling such a sensitivity analysis to the spectral validation used here as an independent test would allow each candidate definition to be evaluated not only by its statistical

separation between the lines, but also by its agreement with an independent physical signal, providing an objective basis for choosing the fragmentation estimator rather than fixing it by convention. A fragmentation readout refined in this way, and tested through cross-detector and, where feasible, cross-laboratory comparison, would be a step towards a more standardised functional readout for DS and related SCN1A-related disorders.

In summary, combining patient-derived neuronal networks with multi-electrode recordings offers a scalable way to study human brain disorders in human cells, to follow them over time, and potentially to reduce the reliance on animal models. Making such platforms trustworthy, reproducible across laboratories and clearly linked to the underlying biology is as important a contribution as identifying any single new phenotype. The present work contributes to this goal on two fronts: biologically, by showing that burst fragmentation tracks Dravet syndrome severity in heterogeneous cortical networks, and methodologically, by showing that the number of fragments counted depends on the detection method, so that dataset-specific tuning and expert validation are needed before such counts are compared.

Bibliography

- Al-Jaberi, N., Lindsay, S., Sarma, S., Bayatti, N., & Clowry, G. J. (2015). The early fetal development of human neocortical GABAergic interneurons. *Cerebral Cortex*, 25(3), 631–645. <https://doi.org/10.1093/cercor/bht254>
- Bakkum, D. J., Radivojevic, M., Frey, U., Franke, F., Hierlemann, A., & Takahashi, H. (2014). Parameters for burst detection. *Frontiers in Computational Neuroscience*, 7, 193. <https://doi.org/10.3389/fncom.2013.00193>
- Bean, B. P. (2007). The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, 8(6), 451–465. <https://doi.org/10.1038/nrn2148>
- Bender, A. C., Morse, R. P., Scott, R. C., Holmes, G. L., & Lenck-Santini, P. P. (2012). SCN1A mutations in Dravet syndrome: Impact of interneuron dysfunction on neural networks and cognitive outcome. *Epilepsy & Behavior*, 23(3), 177–186. <https://doi.org/10.1016/j.yebeh.2011.11.022>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Blanco, J. A., Stead, M., Krieger, A., Viventi, J., Marsh, W. R., Lee, K. H., Worrell, G. A., & Litt, B. (2010). Unsupervised classification of high-frequency oscillations in human neocortical epilepsy and control patients. *Journal of Neurophysiology*, 104(5), 2900–2912. <https://doi.org/10.1152/jn.01082.2009>
- Bologna, L. L., Nieuwenhuis, T., Tedesco, M., Chiappalone, M., Benfenati, F., & Martinoia, S. (2010). Low-frequency stimulation enhances burst activity in cortical cultures during development. *Neuroscience*, 165(3), 692–704. <https://doi.org/10.1016/j.neuroscience.2009.11.018>
- Brewer, G. J., Torricelli, J. R., Evege, E. K., & Price, P. J. (1993). Optimized survival of hippocampal neurons in B27-supplemented Neurobasal, a new serum-free medium combination. *Journal of Neuroscience Research*, 35(5), 567–576. <https://doi.org/10.1002/jnr.490350513>
- Budday, S., Steinmann, P., & Kuhl, E. (2015). Physical biology of human brain development. *Frontiers in Cellular Neuroscience*, 9, 257. <https://doi.org/10.3389/fncel.2015.00257>

- Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), 407–420. <https://doi.org/10.1038/nrn3241>
- Buzsáki, G., & da Silva, F. L. (2012). High frequency oscillations in the intact brain. *Progress in Neurobiology*, 98(3), 241–249. <https://doi.org/10.1016/j.pneurobio.2012.02.004>
- Carter, M., & Shieh, J. C. (2015). *Guide to research techniques in neuroscience* (2nd ed.). Elsevier/Academic Press.
- Cerina, M., Piastra, M. C., & Frega, M. (2023). The potential of *in vitro* neuronal networks cultured on micro electrode arrays for biomedical research. *Progress in Biomedical Engineering*, 5(3), 032002. <https://doi.org/10.1088/2516-1091/acce12>
- Cetica, V., Chiari, S., Mei, D., Parrini, E., Grisotto, L., Marini, C., Pucatti, D., Ferrari, A., Sicca, F., Specchio, N., Trivisano, M., Battaglia, D., Contaldo, I., Zamponi, N., Petrelli, C., Granata, T., Ragona, F., Avanzini, G., & Guerrini, R. (2017). Clinical and genetic factors predicting Dravet syndrome in infants with SCN1A mutations. *Neurology*, 88(11), 1037–1044. <https://doi.org/10.1212/WNL.0000000000003716>
- Chambers, S. M., Fasano, C. A., Papapetrou, E. P., Tomishima, M., Sadelain, M., & Studer, L. (2009). Highly efficient neural conversion of human ES and iPS cells by dual inhibition of SMAD signaling. *Nature Biotechnology*, 27(3), 275–280. <https://doi.org/10.1038/nbt.1529>
- Charlesworth, P., Cotterill, E., Morton, A., Grant, S. G. N., & Eglen, S. J. (2015). Quantitative differences in developmental profiles of spontaneous activity in cortical and hippocampal cultures. *Neural Development*, 10, 1. <https://doi.org/10.1186/s13064-014-0028-0>
- Cheah, C. S., Yu, F. H., Westenbroek, R. E., Kalume, F. K., Oakley, J. C., Potter, G. B., Rubenstein, J. L. R., & Catterall, W. A. (2012). Specific deletion of NaV1.1 sodium channels in inhibitory interneurons causes seizures and premature death in a mouse model of Dravet syndrome. *Proceedings of the National Academy of Sciences*, 109(36), 14646–14651. <https://doi.org/10.1073/pnas.1211591109>
- Chen, I., & Lui, F. (2019). Neuroanatomy, neuron action potential. In *StatPearls*. StatPearls Publishing. <https://www.ncbi.nlm.nih.gov/books/NBK546639/>
- Chiappalone, M., Bove, M., Vato, A., Tedesco, M., & Martinoia, S. (2006). Dissociated cortical networks show spontaneously correlated activity patterns during *in vitro*

- development. *Brain Research*, 1093(1), 41–53.
<https://doi.org/10.1016/j.brainres.2006.03.049>
- Chiappalone, M., Novellino, A., Vajda, I., Vato, A., Martinoia, S., & van Pelt, J. (2005). Burst detection algorithms for the analysis of spatio-temporal patterns in cortical networks of neurons. *Neurocomputing*, 65–66, 653–662.
<https://doi.org/10.1016/j.neucom.2004.10.094>
- Cimbalnik, J., Brinkmann, B., Kremen, V., Jurak, P., Berry, B., Van Gompel, J., Stead, M., & Worrell, G. (2018a). Physiological and pathological high frequency oscillations in focal epilepsy. *Annals of Clinical and Translational Neurology*, 5(9), 1062–1076. <https://doi.org/10.1002/acn3.618>
- Cimbalnik, J., Hewitt, A., Worrell, G., & Stead, M. (2018b). The CS algorithm: A novel method for high frequency oscillation detection in EEG. *Journal of Neuroscience Methods*, 293, 6–16. <https://doi.org/10.1016/j.jneumeth.2017.08.023>
- Claes, L., Del-Favero, J., Ceulemans, B., Lagae, L., Van Broeckhoven, C., & De Jonghe, P. (2001). *De novo* mutations in the sodium-channel gene *SCN1A* cause severe myoclonic epilepsy of infancy. *American Journal of Human Genetics*, 68(6), 1327–1332. <https://doi.org/10.1086/320609>
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Cooper, M. S., McIntosh, A., Crompton, D. E., McMahon, J. M., Schneider, A., Farrell, K., Ganesan, V., Gill, D., Kivity, S., Lerman-Sagie, T., McLellan, A., Pelekanos, J., Ramesh, V., Sadleir, L., Wirrell, E., & Scheffer, I. E. (2016). Mortality in Dravet syndrome. *Epilepsy Research*, 128, 43–47.
<https://doi.org/10.1016/j.epilepsyres.2016.10.006>
- Cotterill, E., Charlesworth, P., Thomas, C. W., Paulsen, O., & Eglen, S. J. (2016). A comparison of computational methods for detecting bursts in neuronal spike trains and their application to human stem cell-derived neuronal networks. *Journal of Neurophysiology*, 116(2), 306–321. <https://doi.org/10.1152/jn.00093.2016>
- Degl’Innocenti, E., & Dell’Anno, M. T. (2023). Human and mouse cortical astrocytes: A comparative view from development to morphological and functional characterization. *Frontiers in Neuroanatomy*, 17, 1130729.
<https://doi.org/10.3389/fnana.2023.1130729>
- Ding, J., Li, X., Tian, H., Wang, L., Guo, B., Wang, Y., Li, W., Wang, F., & Sun, T. (2021). SCN1A mutation—beyond Dravet syndrome: A systematic review and

- narrative synthesis. *Frontiers in Neurology*, 12, 743726.
<https://doi.org/10.3389/fneur.2021.743726>
- Doorn, N., Voogd, E. J. H. F., Levers, M. R., van Putten, M. J. A. M., & Frega, M. (2024). Breaking the burst: Unveiling mechanisms behind fragmented network bursts in patient-derived neurons. *Stem Cell Reports*, 19(11), 1583–1597.
<https://doi.org/10.1016/j.stemcr.2024.09.001>
- Dravet, C. (2011). Dravet syndrome history. *Developmental Medicine & Child Neurology*, 53(Suppl. 2), 1–6. <https://doi.org/10.1111/j.1469-8749.2011.03964.x>
- Eisenman, L. N., Emnett, C. M., Mohan, J., Zorumski, C. F., & Mennerick, S. (2015). Quantification of bursting and synchrony in cultured hippocampal neurons. *Journal of Neurophysiology*, 114(2), 1059–1071.
<https://doi.org/10.1152/jn.00079.2015>
- Engle, S. J., Blaha, L., & Kleiman, R. J. (2018). Best practices for translational disease modeling using human iPSC-derived neurons. *Neuron*, 100(4), 783–797.
<https://doi.org/10.1016/j.neuron.2018.10.033>
- Eytan, D., & Marom, S. (2006). Dynamics and effective topology underlying synchronization in networks of cortical neurons. *Journal of Neuroscience*, 26(33), 8465–8476. <https://doi.org/10.1523/JNEUROSCI.1627-06.2006>
- Frega, M., Linda, K., Keller, J. M., Gümüş-Akay, G., Mossink, B., van Rhijn, J. R., Negwer, M., Klein Gunnewiek, T., Foreman, K., Kompier, N., Schoenmaker, C., van den Akker, W., van der Werf, I., Oudakker, A., Zhou, H., Kleefstra, T., Schubert, D., van Bokhoven, H., & Nadif Kasri, N. (2019). Neuronal network dysfunction in a model for Kleefstra syndrome mediated by enhanced NMDAR signaling. *Nature Communications*, 10, Article 4928.
<https://doi.org/10.1038/s41467-019-12947-3>
- Frega, M., van Gestel, S. H. C., Linda, K., van der Raadt, J., Keller, J., Van Rhijn, J.-R., Schubert, D., Albers, C. A., & Nadif Kasri, N. (2017). Rapid neuronal differentiation of induced pluripotent stem cells for measuring network activity on micro-electrode arrays. *Journal of Visualized Experiments*, 119, 54900.
<https://doi.org/10.3791/54900>
- Gabriele, M., Vitriolo, A., Cuvertino, S., Pereira, M. F., Franconi, C., Germain, P.-L., et al. (2021). KMT2D haploinsufficiency in Kabuki syndrome disrupts neuronal function through transcriptional and chromatin rewiring independent of H3K4-monomethylation. *bioRxiv*. <https://doi.org/10.1101/2021.04.22.440945>

- Gähwiler, B. H., Capogna, M., Debanne, D., McKinney, R. A., & Thompson, S. M. (1997). Organotypic slice cultures: A technique has come of age. *Trends in Neurosciences*, 20(10), 471–477. [https://doi.org/10.1016/s0166-2236\(97\)01122-3](https://doi.org/10.1016/s0166-2236(97)01122-3)
- Gordon, J., Amini, S., & White, M. K. (2013). General overview of neuronal cell culture. *Methods in Molecular Biology*, 1078, 1–8. https://doi.org/10.1007/978-1-62703-640-5_1
- Gourévitch, B., & Eggermont, J. J. (2007). A nonparametric approach for detection of bursts in spike trains. *Journal of Neuroscience Methods*, 160(2), 349–358. <https://doi.org/10.1016/j.jneumeth.2006.09.024>
- Gunhanlar, N., Shpak, G., van der Kroeg, M., Gouty-Colomer, L. A., Munshi, S. T., Lendemeijer, B., Ghazvini, M., Dupont, C., Hoogendijk, W. J. G., Gribnau, J., de Vrij, F. M. S., & Kushner, S. A. (2018). A simplified protocol for differentiation of electrophysiologically mature neuronal networks from human induced pluripotent stem cells. *Molecular Psychiatry*, 23(5), 1336–1344. <https://doi.org/10.1038/mp.2017.56>
- Hales, C. M., Zeller-Townson, R., Newman, J. P., Shoemaker, J. T., Killian, N. J., & Potter, S. M. (2012). Stimulus-evoked high frequency oscillations are present in neuronal networks on microelectrode arrays. *Frontiers in Neural Circuits*, 6, 29. <https://doi.org/10.3389/fncir.2012.00029>
- Hammer, M. F., Ishii, A., Johnstone, L., Tchourbanov, A., Lau, B., Sprissler, R., Hallmark, B., Zhang, M., Zhou, J., Watkins, J., & Hirose, S. (2017). Rare variants of small effect size in neuronal excitability genes influence clinical outcome in Japanese cases of SCN1A truncation-positive Dravet syndrome. *PLoS ONE*, 12(7), e0180485. <https://doi.org/10.1371/journal.pone.0180485>
- Han, S., Tai, C., Westenbroek, R. E., Yu, F. H., Cheah, C. S., Potter, G. B., Rubenstein, J. L., Scheuer, T., de la Iglesia, H. O., & Catterall, W. A. (2012). Autistic-like behaviour in *Scn1a*^{+/-} mice and rescue by enhanced GABA-mediated neurotransmission. *Nature*, 489(7416), 385–390. <https://doi.org/10.1038/nature11356>
- Harrison, R. G. (1910). The outgrowth of the nerve fiber as a mode of protoplasmic movement. *Journal of Experimental Zoology*, 9(4), 787–846. <https://doi.org/10.1002/jez.1400090405>
- Heikkilä, T. J., Ylä-Outinen, L., Tanskanen, J. M. A., Lappalainen, R. S., Skottman, H., Suuronen, R., Mikkonen, J. E., Hyttinen, J. A. K., & Narkilahti, S. (2009). Human

- embryonic stem cell-derived neuronal cells form spontaneously active neuronal networks *in vitro*. *Experimental Neurology*, 218(1), 109–116.
<https://doi.org/10.1016/j.expneurol.2009.04.011>
- Higurashi, N., Uchida, T., Lossin, C., Misumi, Y., Okada, Y., Akamatsu, W., Imaizumi, Y., Zhang, B., Nabeshima, K., Mori, M. X., Katsurabayashi, S., Shirasaka, Y., Okano, H., & Hirose, S. (2013). A human Dravet syndrome model from patient induced pluripotent stem cells. *Molecular Brain*, 6, 19.
<https://doi.org/10.1186/1756-6606-6-19>
- Hodge, R. D., Bakken, T. E., Miller, J. A., Smith, K. A., Barkan, E. R., Graybuck, L. T., Close, J. L., Long, B., Johansen, N., Penn, O., Yao, Z., Eggermont, J., Höllt, T., Levi, B. P., Shehata, S. I., Aevermann, B., Beller, A., Bertagnolli, D., Brouner, K., Casper, T., ... Lein, E. S. (2019). Conserved cell types with divergent features in human versus mouse cortex. *Nature*, 573(7772), 61–68.
<https://doi.org/10.1038/s41586-019-1506-7>
- Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), 500–544. <https://doi.org/10.1113/jphysiol.1952.sp004764>
- Hutter-Schmid, B., Kniewallner, K. M., & Humpel, C. (2015). Organotypic brain slice cultures as a model to study angiogenesis of brain vessels. *Frontiers in Cell and Developmental Biology*, 3, 52. <https://doi.org/10.3389/fcell.2015.00052>
- Hyvärinen, T., Hyysalo, A., Kapucu, F. E., Aarnos, L., Vinogradov, A., Eglen, S. J., Ylä-Outinen, L., & Narkilahti, S. (2019). Functional characterization of human pluripotent stem cell-derived cortical networks differentiated on laminin-521 substrate: Comparison to rat cortical cultures. *Scientific Reports*, 9(1), 17125.
<https://doi.org/10.1038/s41598-019-53647-8>
- Jacobs, J., LeVan, P., Chander, R., Hall, J., Dubeau, F., & Gotman, J. (2008). Interictal high-frequency oscillations (80–500 Hz) are an indicator of seizure onset areas independent of spikes in the human epileptic brain. *Epilepsia*, 49(11), 1893–1907.
<https://doi.org/10.1111/j.1528-1167.2008.01656.x>
- Jiao, J., Yang, Y., Shi, Y., Chen, J., Gao, R., Fan, Y., Yao, H., Liao, W., Sun, X.-F., & Gao, S. (2013). Modeling Dravet syndrome using induced pluripotent stem cells (iPSCs) and directly converted neurons. *Human Molecular Genetics*, 22(21), 4241–4252. <https://doi.org/10.1093/hmg/ddt275>

- Jones, I. L., Livi, P., Lewandowska, M. K., Fiscella, M., Roscic, B., & Hierlemann, A. (2011). The potential of microelectrode arrays and microelectronics for biomedical research and diagnostics. *Analytical and Bioanalytical Chemistry*, 399(7), 2313–2329. <https://doi.org/10.1007/s00216-010-3968-1>
- Kapucu, F. E., Tanskanen, J. M., Mikkonen, J. E., Ylä-Outinen, L., Narkilahti, S., & Hyttinen, J. A. (2012). Burst analysis tool for developing neuronal networks exhibiting highly varying action potential dynamics. *Frontiers in computational neuroscience*, 6, 38. <https://doi.org/10.3389/fncom.2012.00038>
- Kapucu, F., Vinogradov, A., Hyvärinen, T., Ylä-Outinen, L., & Narkilahti, S. (2022). Comparative microelectrode array data of the functional development of hPSC-derived and rat neuronal networks. *Scientific Data*, 9, 120. <https://doi.org/10.1038/s41597-022-01242-4>
- Kiamehr, M., Klettner, A., Richert, E., Koskela, A., Koistinen, A., Skottman, H., Kaarniranta, K., Aalto-Setälä, K., & Juuti-Uusitalo, K. (2019). Compromised barrier function in human induced pluripotent stem-cell-derived retinal pigment epithelial cells from type 2 diabetic patients. *International Journal of Molecular Sciences*, 20(15), 3773. <https://doi.org/10.3390/ijms20153773>
- Kim, T.-G., Yao, R., Monnell, T., Cho, J.-H., Vasudevan, A., Koh, A., Peeyush, K. T., Moon, M., Datta, D., Bolshakov, V. Y., Kim, K.-S., & Chung, S. (2014). Efficient specification of interneurons from human pluripotent stem cells by dorsoventral and rostrocaudal modulation. *Stem Cells*, 32(7), 1789–1804. <https://doi.org/10.1002/stem.1704>
- Kreitzer, F. R., Salomonis, N., Sheehan, A., Huang, M., Park, J. S., Spindler, M. J., Lizarraga, P., Weiss, W. A., So, P.-L., & Conklin, B. R. (2013). A robust method to derive functional neural crest cells from human pluripotent stem cells. *American Journal of Stem Cells*, 2(2), 119–131.
- Lau, P.-M., & Bi, G.-Q. (2005). Synaptic mechanisms of persistent reverberatory activity in neuronal networks. *Proceedings of the National Academy of Sciences*, 102(29), 10333–10338. <https://doi.org/10.1073/pnas.0500717102>
- Legéndy, C. R., & Salcman, M. (1985). Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *Journal of Neurophysiology*, 53(4), 926–939. <https://doi.org/10.1152/jn.1985.53.4.926>
- Lieb, F., Stark, H.-G., & Thielemann, C. (2017). A stationary wavelet transform and a time-frequency based spike detection algorithm for extracellular recorded data.

Journal of Neural Engineering, 14(3), 036013. <https://doi.org/10.1088/1741-2552/aa654b>

- Liu, Y., Lopez-Santiago, L. F., Yuan, Y., Jones, J. M., Zhang, H., O'Malley, H. A., Patino, G. A., O'Brien, J. E., Rusconi, R., Gupta, A., Thompson, R. C., Natowicz, M. R., Meisler, M. H., Isom, L. L., & Parent, J. M. (2013). Dravet syndrome patient-derived neurons suggest a novel epilepsy mechanism. *Annals of Neurology*, 74(1), 128–139. <https://doi.org/10.1002/ana.23897>
- Löscher, W. (2011). Critical review of current animal models of seizures and epilepsy used in the discovery and development of new antiepileptic drugs. *Seizure*, 20(5), 359–368. <https://doi.org/10.1016/j.seizure.2011.01.003>
- Maccione, A., Gandolfo, M., Massobrio, P., Novellino, A., Martinoia, S., & Chiappalone, M. (2009). A novel algorithm for precise identification of spikes in extracellularly recorded neuronal signals. *Journal of Neuroscience Methods*, 177(1), 241–249. <https://doi.org/10.1016/j.jneumeth.2008.09.026>
- Maroof, A. M., Keros, S., Tyson, J. A., Ying, S.-W., Ganat, Y. M., Merkle, F. T., Liu, B., Goulburn, A., Stanley, E. G., Elefanty, A. G., Widmer, H. R., Eggen, K., Goldstein, P. A., Anderson, S. A., & Studer, L. (2013). Directed differentiation and functional maturation of cortical interneurons from human embryonic stem cells. *Cell Stem Cell*, 12(5), 559–572. <https://doi.org/10.1016/j.stem.2013.04.008>
- Mayer, M., Arrizabalaga, O., Lieb, F., Ciba, M., Ritter, S., & Thielemann, C. (2018). Electrophysiological investigation of human embryonic stem cell derived neurospheres using a novel spike detection algorithm. *Biosensors and Bioelectronics*, 100, 462–468. <https://doi.org/10.1016/j.bios.2017.09.034>
- Millet, L. J., & Gillette, M. U. (2012). Over a century of neuron culture: From the hanging drop to microfluidic devices. *The Yale Journal of Biology and Medicine*, 85(4), 501–521.
- Mossink, B., Verboven, A. H. A., van Hugte, E. J. H., Klein Gunnewiek, T. M., Parodi, G., Linda, K., Schoenmaker, C., Kleefstra, T., Kozicz, T., van Bokhoven, H., Schubert, D., Nadif Kasri, N., & Frega, M. (2021). 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. <https://doi.org/10.1016/j.stemcr.2021.07.001>

- Munafò, M. R., & Davey Smith, G. (2018). Robust research needs many lines of evidence. *Nature*, 553(7689), 399–401. <https://doi.org/10.1038/d41586-018-01023-3>
- Mzezewa, R., Hyvärinen, T., Kulta, O., Vinogradov, A., Pesu, E., Isosaari, L., Vuolanto, V., Kapucu, F. E., Schuster, J., Dahl, N., & Narkilahti, S. (2025). Abnormalities in the functional activity of neural networks in a human iPSC model of Dravet syndrome. *Neuroscience Research*, 220, Article 104958. <https://doi.org/10.1016/j.neures.2025.104958>
- Nabbout, R., Chemaly, N., Chipaux, M., Barcia, G., Bouis, C., Dubouch, C., Leunen, D., Jambaqué, I., Dulac, O., Dellatolas, G., & Chiron, C. (2013). Encephalopathy in children with Dravet syndrome is not a pure consequence of epilepsy. *Orphanet Journal of Rare Diseases*, 8, 176. <https://doi.org/10.1186/1750-1172-8-176>
- Nam, Y., & Wheeler, B. C. (2011). *In vitro* microelectrode array technology and neural recordings. *Critical Reviews in Biomedical Engineering*, 39(1), 45–62. <https://doi.org/10.1615/CritRevBiomedEng.v39.i1.40>
- Nicholas, C. R., Chen, J., Tang, Y., Southwell, D. G., Chalmers, N., Vogt, D., Arnold, C. M., Chen, Y.-J. J., Stanley, E. G., Elefanty, A. G., Sasai, Y., Alvarez-Buylla, A., Rubenstein, J. L. R., & Kriegstein, A. R. (2013). Functional maturation of hPSC-derived forebrain interneurons requires an extended timeline and mimics human neural development. *Cell Stem Cell*, 12(5), 573–586. <https://doi.org/10.1016/j.stem.2013.04.005>
- Novellino, A., Scelfo, B., Palosaari, T., Price, A., Sobanski, T., Shafer, T. J., Johnstone, A. F. M., Gross, G. W., Gramowski, A., Schroeder, O., Jügel, K., Chiappalone, M., Benfenati, F., Martinoia, S., Tedesco, M. T., Defranchi, E., D'Angelo, P., & Whelan, M. (2011). Development of micro-electrode array based tests for neurotoxicity: Assessment of interlaboratory reproducibility with neuroactive chemicals. *Frontiers in Neuroengineering*, 4, 4. <https://doi.org/10.3389/fneng.2011.00004>
- Oberheim, N. A., Takano, T., Han, X., He, W., Lin, J. H., Wang, F., Xu, Q., Wyatt, J. D., Pilcher, W., Ojemann, J. G., Ransom, B. R., Goldman, S. A., & Nedergaard, M. (2009). Uniquely hominid features of adult human astrocytes. *The Journal of Neuroscience*, 29(10), 3276–3287. <https://doi.org/10.1523/JNEUROSCI.4707-08.2009>

- Odawara, A., Katoh, H., Matsuda, N., & Suzuki, I. (2016). Physiological maturation and drug responses of human induced pluripotent stem cell-derived cortical neuronal networks in long-term culture. *Scientific Reports*, 6, 26181. <https://doi.org/10.1038/srep26181>
- Ogiwara, I., Miyamoto, H., Morita, N., Atapour, N., Mazaki, E., Inoue, I., Takeuchi, T., Itohara, S., Yanagawa, Y., Obata, K., Furuichi, T., Hensch, T. K., & Yamakawa, K. (2007). NaV1.1 localizes to axons of parvalbumin-positive inhibitory interneurons: A circuit basis for epileptic seizures in mice carrying an *Scn1a* gene mutation. *Journal of Neuroscience*, 27(22), 5903–5914. <https://doi.org/10.1523/JNEUROSCI.5270-06.2007>
- Pasquale, V., Martinoia, S., & Chiappalone, M. (2010). A self-adapting approach for the detection of bursts and network bursts in neuronal cultures. *Journal of Computational Neuroscience*, 29(1), 213–229. <https://doi.org/10.1007/s10827-009-0175-1>
- Pradeepan, K. S., McCready, F. P., Wei, W., Khaki, M., Zhang, W., Salter, M. W., Ellis, J., & Martinez-Trujillo, J. (2024). Calcium-dependent hyperexcitability in human stem cell-derived Rett syndrome neuronal networks. *Biological Psychiatry Global Open Science*, 4(2), Article 100290. <https://doi.org/10.1016/j.bpsgos.2024.100290>
- Pré, D., Cazares, C., Wooten, A. T., Zhou, H., Onofre, I., Neil, A., Logan, T., Hu, R., Lui, J. H., Voytek, B., & Bang, A. G. (2026). Pharmacological manipulation of nested oscillations in human iPSC-derived 2D neuronal networks. *Neurobiology of Disease*, 220, Article 107281. <https://doi.org/10.1016/j.nbd.2026.107281>
- Quiroga, R. Q., Nadasdy, Z., & Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Computation*, 16(8), 1661–1687. <https://doi.org/10.1162/089976604774201631>
- Rampazzo, A. C. M., Dos Santos, R. R. P., Maluf, F. A., Simm, R. F., Marson, F. A. L., Ortega, M. M., & de Aguiar, P. H. P. (2021). Dravet syndrome and Dravet syndrome-like phenotype: A systematic review of the *SCN1A* and *PCDH19* variants. *Neurogenetics*, 22(2), 105–115. <https://doi.org/10.1007/s10048-021-00644-7>
- Rhodes, T. H., Lossin, C., Vanoye, C. G., Wang, D. W., & George, A. L., Jr. (2004). Noninactivating voltage-gated sodium channels in severe myoclonic epilepsy of infancy. *Proceedings of the National Academy of Sciences*, 101(30), 11147–11152. <https://doi.org/10.1073/pnas.0402482101>

- Sanli, C., & Lambiotte, R. (2015). Local variation of hashtag spike trains and popularity in Twitter. *PLOS ONE*, *10*(7), Article e0131704. <https://doi.org/10.1371/journal.pone.0131704>
- Scheffer, I. E., & Nabbout, R. (2019). SCN1A-related phenotypes: Epilepsy and beyond. *Epilepsia*, *60*(S3), S17–S24. <https://doi.org/10.1111/epi.16386>
- Schuster, J., Fatima, A., Sobol, M., Norradin, F. H., Laan, L., & Dahl, N. (2019). Generation of three human induced pluripotent stem cell (iPSC) lines from three patients with Dravet syndrome carrying distinct SCN1A gene mutations. *Stem Cell Research*, *39*, 101523. <https://doi.org/10.1016/j.scr.2019.101523>
- Shi, Y., Kirwan, P., & Livesey, F. J. (2012). Directed differentiation of human pluripotent stem cells to cerebral cortex neurons and neural networks. *Nature Protocols*, *7*(10), 1836–1846. <https://doi.org/10.1038/nprot.2012.116>
- Shi, Z., Zhang, J., Chen, S., Li, Y., Lei, X., Qiao, H., Zhu, Q., Hu, B., Zhou, Q., & Jiao, J. (2016). Conversion of fibroblasts to parvalbumin neurons by one transcription factor, Ascl1, and the chemical compound forskolin. *Journal of Biological Chemistry*, *291*(26), 13560–13570. <https://doi.org/10.1074/jbc.M115.709808>
- Shinomoto, S., Kim, H., Shimokawa, T., Matsuno, N., Funahashi, S., Shima, K., Fujita, I., Tamura, H., Doi, T., Kawano, K., Inaba, N., Fukushima, K., Kurkin, S., Kurata, K., Taira, M., Tsutsui, K.-I., Komatsu, H., Ogawa, T., Koida, K., ... Toyama, K. (2009). Relating neuronal firing patterns to functional differentiation of cerebral cortex. *PLoS Computational Biology*, *5*(7), e1000433. <https://doi.org/10.1371/journal.pcbi.1000433>
- Sun, Y., Paşca, S. P., Portmann, T., Goold, C., Worringer, K. A., Guan, W., Chan, K. C. C., Gai, H., Vogt, D., Chen, Y.-J. J., Mao, R., Chan, K., Rubenstein, J. L. R., Dixon, D. M., & Dolmetsch, R. E. (2016). A deleterious Nav1.1 mutation selectively impairs telencephalic inhibitory neurons derived from Dravet syndrome patients. *eLife*, *5*, e13073. <https://doi.org/10.7554/eLife.13073>
- Takahashi, K., & Yamanaka, S. (2006). Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell*, *126*(4), 663–676. <https://doi.org/10.1016/j.cell.2006.07.024>
- Takahashi, K., Tanabe, K., Ohnuki, M., Narita, M., Ichisaka, T., Tomoda, K., & Yamanaka, S. (2007). Induction of pluripotent stem cells from adult human fibroblasts by defined factors. *Cell*, *131*(5), 861–872. <https://doi.org/10.1016/j.cell.2007.11.019>

- Thomas, C. A., Springer, P. A., Loeb, G. E., Berwald-Netter, Y., & Okun, L. M. (1972). A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Experimental Cell Research*, 74(1), 61–66. [https://doi.org/10.1016/0014-4827\(72\)90481-8](https://doi.org/10.1016/0014-4827(72)90481-8)
- Van Hugte, E. J. H., Lewerissa, E. I., Wu, K. M., Scheefhals, N., Parodi, G., van Voorst, T. W., Puvogel, S., Kogo, N., Keller, J. M., Frega, M., Schubert, D., Schelhaas, H. J., Verhoeven, J., Majoie, M., van Bokhoven, H., & Nadif Kasri, N. (2023). SCN1A-deficient excitatory neuronal networks display mutation-specific phenotypes. *Brain*, 146(12), 5153–5167. <https://doi.org/10.1093/brain/awad245>
- van Pelt, J., Wolters, P. S., Corner, M. A., Rutten, W. L. C., & Ramakers, G. J. A. (2004). Long-term characterization of firing dynamics of spontaneous bursts in cultured neural networks. *IEEE Transactions on Biomedical Engineering*, 51(11), 2051–2062. <https://doi.org/10.1109/TBME.2004.827936>
- Verboven, A. H. A., Puvogel, S., Latour, B. L., Kolsters, N., Lewerissa, E. I., Mijdam, R., Dyke, E., Duineveld, D., Zoodsma, M., Claassen, W. J. J., Oudakker, A., Schoenmaker, C., Koolen, D. A., de Vries, B. B. A., 't Hoen, P. A. C., Linda, K., & Nadif Kasri, N. (2026). Integrative transcriptomics and electrophysiological profiling of hiPSC-derived neurons identifies novel druggable pathways in Koolen-de Vries Syndrome. *Molecular Psychiatry*, 31(6), 3558–3575. <https://doi.org/10.1038/s41380-026-03482-x>
- Vinogradov, A., Kapucu, E. F., & Narkilahti, S. (2024). Exploring kainic acid-induced alterations in circular tripartite networks with advanced analysis tools. *eNeuro*, 11, ENEURO.0035-24.2024. <https://doi.org/10.1523/ENEURO.0035-24.2024>
- Wagenaar, D. A., Pine, J., & Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience*, 7(1), 11. <https://doi.org/10.1186/1471-2202-7-11>
- Wirrell, E. C., Hood, V., Knupp, K. G., Meskis, M. A., Nabbout, R., Scheffer, I. E., Wilmschurst, J., & Sullivan, J. (2022). International consensus on diagnosis and management of Dravet syndrome. *Epilepsia*, 63(7), 1761–1777. <https://doi.org/10.1111/epi.17274>
- Worrell, G. A., Gardner, A. B., Stead, S. M., Hu, S., Goerss, S., Cascino, G. J., Meyer, F. B., Marsh, R., & Litt, B. (2008). High-frequency oscillations in human temporal lobe: Simultaneous microwire and clinical macroelectrode recordings. *Brain*, 131(4), 928–937. <https://doi.org/10.1093/brain/awn006>

- Worrell, G. A., Parish, L., Cranstoun, S. D., Jonas, R., Baltuch, G., & Litt, B. (2004). High-frequency oscillations and seizure generation in neocortical epilepsy. *Brain*, 127(7), 1496–1506. <https://doi.org/10.1093/brain/awh149>
- Wu, Y. W., Sullivan, J., McDaniel, S. S., Meisler, M. H., Walsh, E. M., Li, S. X., & Kuzniewicz, M. W. (2015). Incidence of Dravet syndrome in a US population. *Pediatrics*, 136(5), e1310–e1315. <https://doi.org/10.1542/peds.2015-1807>
- Xie, Y., Ng, N. N., Safrina, O. S., Ramos, C. M., Ess, K. C., Schwartz, P. H., Smith, M. A., & O'Dowd, D. K. (2020). Comparisons of dual isogenic human iPSC pairs identify functional alterations directly caused by an epilepsy associated SCN1A mutation. *Neurobiology of Disease*, 134, Article 104627. <https://doi.org/10.1016/j.nbd.2019.104627>
- Yu, F. H., Mantegazza, M., Westenbroek, R. E., Robbins, C. A., Kalume, F., Burton, K. A., Spain, W. J., McKnight, G. S., Scheuer, T., & Catterall, W. A. (2006). Reduced sodium current in GABAergic interneurons in a mouse model of severe myoclonic epilepsy in infancy. *Nature Neuroscience*, 9(9), 1142–1149. <https://doi.org/10.1038/nn1754>
- Zayat, V., Kuczynska, Z., Liput, M., Metin, E., Rzonca-Niewczas, S., Smyk, M., Mazurczak, T., Goszczanska-Ciuchta, A., Leszczynski, P., Hoffman-Zacharska, D., & Buzanska, L. (2023). The generation of human iPSC lines from three individuals with Dravet syndrome and characterization of neural differentiation markers in iPSC-derived ventral forebrain organoid model. *Cells*, 12(2), 339. <https://doi.org/10.3390/cells12020339>
- Zeng, H., Shen, E. H., Hohmann, J. G., Oh, S. W., Bernard, A., Royall, J. J., Glattfelder, K. J., Sunkin, S. M., Morris, J. A., Guillozet-Bongaarts, A. L., Smith, K. A., Ebbert, A. J., Swanson, B., Kuan, L., Page, D. T., Overly, C. C., Lein, E. S., Hawrylycz, M. J., Hof, P. R., Hyde, T. M., ... Jones, A. R. (2012). Large-scale cellular-resolution gene profiling in human neocortex reveals species-specific molecular signatures. *Cell*, 149(2), 483–496. <https://doi.org/10.1016/j.cell.2012.02.052>
- Zhang, Y., Pak, C., Han, Y., Ahlenius, H., Zhang, Z., Chanda, S., Marro, S., Patzke, C., Acuna, C., Covy, J., Xu, W., Yang, N., Danko, T., Chen, L., Wernig, M., & Südhof, T. C. (2013). Rapid single-step induction of functional neurons from human pluripotent stem cells. *Neuron*, 78(5), 785–798. <https://doi.org/10.1016/j.neuron.2013.05.029>

Appendix A

The overall between-line comparisons reported on the trajectory line plots pool the recordings of each well across developmental time points (Section 3.5.1). To verify that this choice does not drive the reported differences, each well was collapsed to the mean of its time points and the identical adaptive test was re-applied, so that only the experimental unit changed. Of 168 pairwise contrasts, 15 changed significance status; after removing contrasts duplicated across analysis blocks, 12 remain, of which ten were lost and two gained. Six are patient–patient contrasts, three are control–control and three are patient–control. They are listed in Table A1.

Detector	Metric	Contrast	Contrast type	p (pooled)	p (well)	Outcome
Tampere V1	MFR	DD1C vs DD5A	patient–patient	6.9e-07	0.094	lost
Tampere V1	SEB duration	DD1C vs DD5A	patient–patient	9.4e-10	0.056	lost
Tampere V1	Synchrony	DD1C vs DD5A	patient–patient	0.028	0.724	lost
Tampere V1	CV(IBM)	DD1C vs DD5A	patient–patient	0.044	0.094	lost
Tampere V1	NBR	DD1C vs C10902	patient–control	0.020	0.238	lost
Tampere V1	NBD	DD5A vs C10902	patient–control	2.8e-06	0.366	lost
Tampere V1	NBD	TUBA vs C10902	control–control	2.8e-06	0.092	lost
Tampere V1	%RS	TUBA vs C10902	control–control	0.131	0.024	gained
FR	NBR	DD1C vs DD5A	patient–patient	0.028	0.483	lost
FR	NBD	DD1C vs DD5A	patient–patient	1.8e-05	0.268	lost
FR	NBD	TUBA vs C10902	control–control	0.0011	0.268	lost
Tampere V2	FragPerNB	DD1C vs TUBA	patient–control	0.062	0.0016	gained

Table A1. Pairwise contrasts between cell lines whose significance status changed when the experimental unit was moved from the well $\times$ DIV observation to the well. The statistical test is identical in the two p-value columns; only the unit of analysis differs. p-values are Benjamini-Hochberg adjusted within metric, $\alpha = 0.05$. Contrasts duplicated across analysis blocks are listed once.