

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

How to Train Your Organoid

Permalink

<https://escholarship.org/uc/item/0g88h0sz>

ISBN

9798241662620

Author

Robbins, Ash

Publication Date

2026-03-17

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

HOW TO TRAIN YOUR ORGANOID

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ELECTICAL AND COMPUTER ENGINEERING

by

Ash Robbins

March 2026

The Dissertation of Ash Robbins
is approved:

Professor Mircea Teodorescu, Chair

Professor David Haussler

Professor Marco Rolandi

Peter Biehl
Vice Provost and Dean of Graduate Studies

Copyright © by

Ash Robbins

2026

Table of Contents

List of Figures	vi
List of Tables	ix
Abstract	x
Dedication	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Background	5
1.3 Contribution to the Field	11
2 Hybrid Soft-Rigid Robotic Learning Environment	14
2.1 Motivation for learning to control biological-like systems	15
2.1.1 Design concept for hybrid joints	18
2.1.2 Control strategies	20
2.1.3 Network architecture	21
2.1.4 Reinforcement learning framework	22
2.2 Results	24
2.3 Discussion	28
3 Optogenetic Modulation of Epileptiform Activity in Human Brain Slices	29
3.1 Introduction	30
3.2 Human hippocampus slice culture and neuronal transduction	31
3.2.1 Design of optogenetic system	34
3.2.2 Optogenetic system software low-level functionality	35

3.2.3	Activity detection algorithm and closed-loop optical intervention	36
3.2.4	Optogenetic control of human hippocampal network	36
3.2.5	Optical HcKCR1-mediated inhibition of network activity .	37
3.2.6	Optogenetic suppression in provoked hyperexcitability . . .	39
3.2.7	Activity heatmaps	43
4	Internet-Connected Cortical Organoids for Project-Based Stem Cell and Neuroscience Education	46
4.1	Introduction	47
4.2	Methods	48
4.2.1	Ethics statement	48
4.2.2	Electrophysiology experiments	49
4.2.3	Data processing	50
4.2.4	Electrophysiology experimental design	50
4.2.5	Stimulation programming	51
4.3	Results	52
4.3.1	Electrophysiology software programs run on organoids in a neuroscience mathematics course	52
4.3.2	Student-designed electrophysiology experiments	60
4.3.3	Organoid experiments lead to higher interest in stem cell and neuroscience research in mathematics students	65
4.4	Discussion	68
5	Goal-Directed Learning in Cortical Organoids via Closed-Loop Electrophysiology	70
5.1	Introduction	71
5.2	Results	77
5.2.1	Multi-Phase Experimentation	77
5.2.2	Neural Configuration	80
5.2.3	Training Paradigm	83
5.2.4	Superior Performance of Adaptive Training Signals over Random and Null Conditions	85
5.2.5	Glutamatergic Transmission is Essential for Learning . . .	96
5.3	Discussion	99
5.4	Limitations	104
5.4.1	Experimental model and study participant details	105
5.4.2	Method details	110
5.4.3	Quantification and statistical analysis	123

6	Real-Time Spike Sorting for Closed-Loop Electrophysiology	145
6.1	Introduction	146
6.2	Results	147
6.2.1	Real-Time Spike Detection on Injected Spikes	147
6.2.2	Real-Time Spike Sorting on Ground Truth Datasets	151
6.2.3	Detection Latencies	155
6.2.4	RT-Sort Compared to Other Sorters	157
6.3	Discussion	162
6.4	Methods	165
6.4.1	Real-Time Spike Detection	165
7	BrainDance: A Framework for Closed-Loop Electrophysiology	167
7.1	Introduction	167
7.2	Architecture Overview	169
7.2.1	Electrical Stimulation	171
7.2.2	Phases	173
7.2.3	Hardware Interface	176
7.2.4	Experiment Management	178
7.2.5	Artifact Removal	179
7.2.6	Spike Detection	180
7.2.7	Real-Time Implementation	182
7.2.8	Coding Framework	183
7.3	Experimental Overview	184
7.3.1	Longitudinal Stimulation Frequency Sweeps	185
7.3.2	Cartpole Balancing Task	187
7.3.3	Rapid Spike Response Pairing	187
7.4	Discussion	189
8	Conclusion	192
9	Future Work	195
	Bibliography	198

List of Figures

2.1	Full assembly of the 3D model and physical robot.	17
2.2	Joints within molds and three experimental designs.	20
2.3	Task-based trajectories and errors.	25
3.1	HD-MEA recordings of human hippocampal slices.	33
3.2	Optogenetic device for MEA.	42
3.3	Optogenetic inhibition of human hippocampal activity.	44
4.1	Students' previous experience and reported interest in topics re- lated to mathematics, neuroscience and stem cell biology.	58
4.2	Mathematics self-concept of the students.	59
4.3	Student experiments explore programming neuronal circuits. . . .	62
4.4	Example of Mathematics of the Mind student-project applying tetanic pattern to cortical organoids.	64

4.5	Students report positive attitudes toward stem cells and neuro- science after organoid experiments.	67
5.1	Experimental overview for <i>in vitro</i> learning	76
5.2	Organoid generation and interaction with electrical interface . . .	79
5.3	Stimulation characterization and causal connectivity	82
5.4	Training paradigms enable task-specific learning	87
5.5	Continuous adaptive training reveals control policy development .	92
5.6	Statistical validation and neural correlates of learning	95
5.7	NBQX and APV abolish learning capabilities.	98
5.8	Additional immunohistochemistry validation of day 30 cortical organoids	135
5.9	Burst-evoking stimulation characterization	136
5.10	Autocorrelation of performance	137
5.11	Random control strategy performance	138
5.12	Evolution of the control policy manifold	139
5.13	Adaptive training reorganizes motor output activity rather than globally increasing firing rates	141
5.14	Spatial distribution of neural configurations across experiments . .	143
6.1	Action potential detection with high temporal accuracy using a CNN.	150
6.2	RT-Sort spike sorting performance on ground truth datasets. . . .	154

6.3	Real-time spike sorting with latencies in the range of synaptic trans-	
	mission.	157
6.4	RT-Sort detects spikes with consistent propagations and amplitudes.	161
7.1	Architecture of the BrainDance framework.	171
7.2	Longitudinal stimulation frequency sweep experiment (busy-bee).	191

List of Tables

5.1	Mixed Effects Model Summary of Drug Effects on Organoid Per-	
	formance	134

Abstract

How to train your organoid

by

Ash Robbins

Understanding how biological neural networks learn is both a fundamental scientific question and an increasingly urgent practical one. Despite decades of work demonstrating that *in vitro* neural cultures can be modified through electrical stimulation, the field lacks both a shared experimental framework for biological learning, as well as results mirroring those of artificial neural networks.

This dissertation presents the tools, experiments, and scientific findings developed while pursuing biological learning in a dish. The work begins with a hybrid soft-rigid robotic system where model-free reinforcement learning outperforms classical control, establishing a first learning environment framework and motivating the extension of reinforcement learning from artificial to biological substrates. Contributions to closed-loop optogenetic modulation of epileptiform activity in human brain slices and a cloud-connected platform for organoid-based neuroscience education provided foundational infrastructure and early validation that high-density microelectrode arrays could reliably interface with biological tissue for stimulation-driven experiments.

The central result is the first demonstration of goal-directed learning in brain organoids. Cortical organoids interfaced with high-density microelectrode arrays learned to balance an inverted pendulum (the “Cartpole” task) through closed-loop electrophysiology. Training signals were adaptively selected by a reinforcement learning algorithm and significantly outperformed both random stimulation and unstimulated controls. Pharmacological blockade of glutamatergic transmission abolished the learned improvements and washout restored them, confirming that the adaptation is biologically mediated and requires synaptic transmission. Causal connectivity analysis revealed that the strength of direct stimulus-evoked pathways predicts learning performance greater than traditional measures of functional connectivity, providing a basis for selecting electrode configurations prior to training.

Two final technical contributions provide a foundation for future exploration: RT-Sort, a real-time spike sorting algorithm, and BrainDance, an open-source Python framework for designing, running, and analyzing closed-loop electrophysiology experiments. BrainDance provides a modular architecture capable of interfacing with various hardware devices. By design, BrainDance facilitates accessible design and sharing of experiments. Together, these tools bring rapid-iteration to biological neural learning, allowing any laboratory with compatible hardware to run reproducible experiments on their own biology.

To my friends and family who have brought fun, support, and light-hearted times which have made everything that much more meaningful. And to Santa Cruz for providing the context to make this possible.

Chapter 1

Introduction

1.1 Motivation

Every form of learning that we've observed occur naturally has occurred in organisms which interact with their environment. Fundamentally, learning is the ability to adapt to an environment in order to achieve some desirable outcome. Behavioral studies have been fruitful in identifying how different forms of learning and adaption emerge in the brain of various species, from associative learning to goal-directed behavior. The human brain's capacity for adaptation comes from the coordinated plasticity of billions of interconnected neurons, yet we lack a clear understanding of how to directly guide the computation and learning in these networks. While this is so early, society is rapidly developing and deploying

invasive neural interfaces: brain-machine interfaces for motor restoration [1, 2], deep brain stimulation for neurological disorders, and closed-loop neuromodulation for psychiatric conditions [3]. Deepening our understanding of exactly how these algorithms can operate is now urgent.

We are at a technological turning point, high-density microelectrode arrays now provide the resolution to record from and stimulate individual neurons within complex networks [4] and are commercially available. Stem-cell-derived brain organoids offer three-dimensional, self-organized neural tissue that recapitulates key aspects of brain development, including diverse cell types and spontaneous electrophysiological activity [5, 6]. Together, these technologies enable systematic study of how living neural circuits process information, adapt to stimulation, and develop functional input-output mappings within the controlled setting of an *in-vitro* preparation.

Despite this technological maturity, the field of biological computing and *in vitro* learning suffers from a critical structural problem: many laboratories build their own experimental infrastructure from scratch. Hardware interfaces, stimulation protocols, analysis pipelines, and training paradigms are developed independently, yielding results that are difficult to compare, reproduce, or build upon. Decades of pioneering work have advanced our understanding, from Shahaf and Marom’s early demonstrations of learning in dissociated cultures [7], through

Potter’s embodied cortical networks [8, 9], to more recent efforts in reservoir computing and free-energy-based frameworks [10, 11]. Yet the lack of a shared experimental framework has fragmented progress. Each new group may invest years in infrastructure, preventing the cumulative advance and reproducibility which the field requires.

A similar issue occurred through the development of artificial reinforcement learning, and the introduction of OpenAI’s Gym [12] sparked rapid progress. Gym is a simple modular interface between learning agents and standardized environments, enabling artificial agents to be easily created and compared on common ground: DQN [13], PPO [14], SAC [15], and dozens of others evaluated on identical tasks from Atari games to continuous control problems. The resulting ecosystem of shared code, verifiable results, and consistent benchmarking transformed brought upon impressive progress on domains such as robotics, boardgames such as Go [16], video games like starcraft [17], and more. This same approach is needed for the upcoming wave of biological learning.

The need for such a paradigm shift is compounded by urgent societal considerations. Artificial intelligence has advanced rapidly without adequate regulation or public understanding, with large corporations steering deployment before society could establish guardrails [18]. Neural interface technologies are now following the same trajectory, progressing from academic research toward clinical and com-

mercial applications at an accelerating pace. Unlike conventional software, these systems operate directly on human cognition. Algorithms which determine how a brain is stimulated have the capacity to alter each part of the way we think, feel, and act. Having these algorithms developed behind closed doors (and optimized for commercial means) poses massive risks in ways that have only been touched on in science-fiction novels. The algorithms which determine how external stimulation shapes neural activity must therefore be studied transparently, with reproducible methods and openly shared tools, before they are commonly deployed in brain-machine interfaces.

In-vitro neural preparations, and brain organoids in particular, offer an effective platform for this work. They provide direct and easy access to living neural circuits without the ethical constraints of human or animal experimentation. Organoids can now be grown in large scales and take weeks to months to mature. Modern high-density electrode arrays support single-neuron-resolution for both recording and stimulation [4]. Organoids can serve as a testbed where the fundamental rules of biological neural learning can be examined through stimulus-driven experiments [19]. This creates a platform for exploring biological neural controllers.

The data generated by these high-density arrays (~ 1 GB/min) aligns with the strengths of current machine learning models. Virtual environments serve a dual

purpose: they provide an embodiment for the organoid while also enabling learning algorithms to operate around the biology. This opens the door for systematically searching across paradigms for encoding/decoding information to/from the organoid, as well as training signals. Virtual environments can help researchers leverage parameter sweeps and optimization to explore the vast parameter space for their desired learning tasks.

Realizing this potential requires a framework that abstracts away (1) the engineering complexity of hardware interfacing, (2) a suite of tasks and virtual environments, and (3) the necessary algorithmic tools for real-time closed-loop control. Without such a framework, each new experiment remains a bespoke engineering project, and the field cannot achieve the rapid, reproducible iteration that has characterized progress in computer science more broadly [12].

This dissertation addresses that need.

1.2 Background

The exploration of learning in biological neural networks *in vitro* has a short but rich history. Early work established that dissociated cortical neurons cultured on multi-electrode arrays (MEAs) exhibit spontaneous activity patterns, including synchronized bursting, and that these patterns can be modified through electrical stimulation. Maeda and colleagues demonstrated that strong localized

stimulation could potentiate network responses, increasing the probability of synchronized bursts and altering propagation patterns for periods exceeding twenty minutes [20]. Jimbo and colleagues showed that local tetanic stimulation induces pathway-specific changes in responsiveness lasting thirty minutes or more, with the direction of change (potentiation/depression) determined by the initial correlation structure between neurons rather than by properties of individual cells [21].

These foundational observations established that *in vitro* neural networks possess the synaptic plasticity machinery necessary for learning. The critical question then became whether this plasticity could be utilized for goal-directed behavior. Shahaf and Marom provided an early and influential demonstration [7], showing that cultured cortical networks could learn to increase their responsiveness to stimulation at a specific electrode when stimulation was removed depending on whether it achieved a target response rate. This work introduced a new paradigm, which looks at stimulation as a value which the network learns to reduce. This framing, termed the “stimulus regulation principle,” requires no separate reward circuitry.

The concept of embodiment, interfacing biological neural networks with external environments through closed-loop stimulation and recording, became popular due to the *coolness* factor of having brain-tissue controlling virtual or physical objects. Reger, Sanguineti, and Mussa-Ivaldi connected *in vitro* neural tissue from

lamprey brainstem to a mobile robot through a bidirectional interface, demonstrating that biological tissue could drive real-world behavior [22]. DeMarse, Wagenaar, and Potter developed the “neurally controlled animat” [8], embedding dissociated cortical cultures in simulated environments and demonstrating that network activity could be shaped through closed-loop interaction. Ruaro and colleagues showed that hippocampal cultures could perform a form of pattern recognition, with responses to spatial stimulation patterns selectively potentiated by tetanization at 100 Hz, showing that input-specific features could be extracted through plasticity [23]. DeMarse and Dockendorf extended the embodiment paradigm to adaptive flight control, using low- and high-frequency stimulation to enhance or depress network-wide evoked responses at separate sites, with evoked activity vectors mapped to pitch and roll control of a simulated aircraft [24]. Warwick and colleagues demonstrated closed-loop mobile robot control by identifying input-output electrode relationships based on functional connectivity and mapping neural activity to wheel velocity in a Braitenberg-inspired architecture [25]. Potter and colleagues showed that spatio-temporal electrical stimuli could shape embodied cortical network behavior in a simple goal-directed learning task, demonstrating that the specific pattern and timing of stimulation was important for inducing functional changes [9]. These studies collectively established the paradigm of embodied biological computation.

More recent work has expanded the scope of in vitro learning studies. Kagan and colleagues embedded dissociated cortical cultures in a simplified version of the game Pong, framing the system through the lens of the free-energy principle and demonstrating that cultures could develop responses consistent with game-play behavior [11]. Cai and colleagues demonstrated brain organoid reservoir computing, showing that organoids could serve as a computational substrate for pattern classification tasks through a reservoir computing framework [10]. Isomura and colleagues provided experimental validation of the free-energy principle with in vitro networks, showing that cultured neurons could perform blind source separation [26, 27]. Each of these efforts advanced the field, but most required substantial custom infrastructure and employed different paradigms, making direct comparison of results challenging.

In parallel, the technology for interfacing with neural tissue has advanced substantially. Modern high-density CMOS-based MEAs, such as the MaxWell system [4], provide 26,400 electrodes with the ability to simultaneously record from up to 1,024 selected channels and stimulate up to 32 electrodes at a 20 kHz sampling rate. This density enables subcellular-resolution recording, precise spatiotemporal spike signatures for identifying individual neurons, and targeted electrical stimulation of single neurons. Combined with real-time artifact removal algorithms such as SALPA [28] and emerging real-time spike sorting methods

such as RT-Sort [29], these arrays support the millisecond-precision closed-loop interaction that modern training paradigms demand.

Brain organoids have emerged as a particularly promising preparation for this work. Derived from stem cells, organoids self-organize into three-dimensional structures that recapitulate aspects of cortical development, including the generation of diverse neuronal and glial cell types, spontaneous electrophysiological maturation, and the formation of functional synaptic networks [5, 6]. Unlike dissociated cultures, organoids retain spatial organization and develop in three dimensions, providing a substrate that more closely resembles the architecture of developing cortical tissue. Mouse cortical organoids offer an attractive model for systematic investigation due to their relatively rapid developmental timeline, enabling higher experimental throughput than their human counterparts.

The field of artificial reinforcement learning offers a model for how easy-to-use software (much being based on python) can accelerate scientific progress. The introduction of OpenAI’s Gym [12] provided a standard interface between learning agents and environments, with consistent APIs for observation, action, and reward. The resulting ecosystem, with extensions for multi-agent settings, continuous control, and increasingly complex simulated environments, enabled reinforcement learning to grow to where it is currently dominant for training state-of-the-art AI models, due to its ability to optimize any task that can be benchmarked.

The framework also empowers users with the ability to design and share their own custom environments. This capability, wrapping a task into an environment, is what gives reinforcement learning its power. Compared to traditional machine learning where you need specific targets (e.g. classification, regression, etc.), reinforcement learning can optimize any task that can be benchmarked. The key finding is not that any single environment or algorithm was revolutionary, but that the experimental and training infrastructure enabled us to get where we are today, and it is a major driving force in current agentic AI.

Bringing this paradigm to biological neural learning requires addressing several challenges that do not arise in the purely computational setting. The biology introduces variability: each organoid develops somewhat differently, and network properties change over time. Hardware constraints impose real-time requirements on stimulation and recording, and other perturbations. The interface between computational algorithms and living tissue requires careful consideration of encoding schemes, timing, and the physiological effects of electrical stimulation. The experimental pipeline—from growing the organoid, to placing it on the array, to characterizing its activity, to running a training experiment—involves multiple stages that must be coordinated and standardized if results are to be reproducible across laboratories.

1.3 Contribution to the Field

This thesis presents contributions across experimental tools, algorithmic methods, and scientific findings that lay the groundwork for systematic investigation of learning in biological neural networks. The work spans several published and collaborative studies, outlined below.

1. **A hybrid soft-rigid robotic learning environment.** We designed compliant robotic joints with integrated elastic ligaments and demonstrated that model-free deep reinforcement learning outperforms classical control across multiple joint morphologies and tasks [30]. Published in *Robotics and Autonomous Systems* (2022), this work provides the first learning environment in this thesis and motivates extending reinforcement learning from artificial to biological substrates.
2. **Closed-loop optogenetic modulation of epileptiform activity in human brain slices.** We contributed to the first demonstration of AAV-mediated optogenetic reduction of network activity in human hippocampal slices recorded on high-density microelectrode arrays, published in *Nature Neuroscience* [31]. The closed-loop infrastructure was built on an early version of the BrainDance platform. A related follow-up characterizing microscale bursting dynamics in patient tissue was published in the *Journal of*

Neurophysiology [32].

3. Internet-connected cortical organoids for neuroscience education.

We co-led the design of a cloud-connected, project-based learning platform in which university students programmed electrophysiological stimulation experiments on live cortical organoids, published in *eNeuro* [33].

4. The first demonstration of goal-directed learning in brain organoids.

We show that cortical organoids interfaced with high-density microelectrode arrays can learn to balance an inverted pendulum through closed-loop electrophysiology, with training signals selected by a reinforcement learning algorithm. Learning was validated pharmacologically and across multiple organoids. Published in *Cell Reports* [34]. Related collaborative work on organoid self-organization [35] and interneuron fate plasticity [36] informed the biological substrate used in these experiments.

5. Real-time spike sorting for closed-loop electrophysiology. We contributed to RT-Sort, an action-potential propagation-based algorithm that assigns sorted spikes within approximately 7.5 ms of detection, enabling closed-loop experiments at single-neuron resolution, and becoming one of the core enabling technologies within the BrainDance platform. Published in *PLoS ONE* [29].

6. **BrainDance: an open-source framework for biological neural learning.** We developed a modular Python platform that provides standardized interfaces for designing, running, and analyzing closed-loop electrophysiology experiments—similar to OpenAI’s Gym for reinforcement learning but with biology-interfacing hardware. The framework integrates hardware control, a composable phase system for experimental protocols, real-time signal processing, and virtual environments for embodied biological experimentation. Collaborative work on IoT infrastructure for cloud biology [37] and automated organoid maintenance [38] contributed to the broader platform ecosystem.

Chapter 2

Hybrid Soft-Rigid Robotic Learning Environment

The following content in this Chapter is adapted from a journal article published in *Robotics and Autonomous Systems* by Elsevier, for which I served as first author alongside co-authors M. Ho and M. Teodorescu.

Biological systems are inherently complex, having nonlinear dynamics, stochasticity, and non-ideal actuators. One method to address these challenges is to learn how to control them with reinforcement learning, which requires a closed-loop interactive environment. This particular work explores both the design, characterization, environment, and control of a hybrid soft-rigid robotic joint. The adapted parts focus on the environment and control for the sake of brevity, but

more details can be found in the original paper [30].

2.1 Motivation for learning to control biological-like systems

Natural manipulators, such as the human arm, were crafted through the long-term optimization of evolution. They tend to be extremely versatile, having the dexterity to work with various objects and environments thus facilitating survival. The hybrid composition of rigid and soft components—including bones, muscles, and connective tissues—yield inherent compliance and flexibility. Biological joints often have passive stability and elasticity which create mechanical feedback that benefits disturbance responses [39]. Recent progress in the mechanical complexity of robots has popularized embedding intelligence within the system. Such robots may inherently dampen motion through elastic components [40] or enable complicated movements emerging from simple actuation—as with origami robots [41].

In contrast, traditional robot arms tend to feature rigid components that are susceptible to large moments propagating throughout the entire robot. This means the robot’s structural integrity can be compromised by a large unpredictable disturbance. Robotic manipulators involving rigid joints have strictly defined degrees

of freedom resulting from the mechanical design. These joints typically fit within three categories: prismatic joints (linear movement on an axis), revolute joints (rotational movement around an axis), or a combination of the two in order to achieve the desired degrees of freedom.

In this chapter, a methodology for the construction of novel robotic joints with a hybrid soft-rigid structure is presented as shown in Fig. 2.1 and model-free control strategies for this class of joints are evaluated.




Figure 2.1: Full assembly of the 3D model and physical robot. **A**, 3D model of the shoulder joint. **B**, 3D model of the elbow joint. **C**, 3D model of the full manipulator with motor mounts, cable routing pieces, and IMU mount. **D**, Physical prototype of the complete manipulator used for all experiments and analysis.

We present a modular approach to designing, implementing, and controlling custom hybrid robots. First, the joint design begins with a simplistic 3D printing approach enabling easy customization with a molding process that supports interfacing between soft and rigid parts. We demonstrate this by modifying the perforation of the rigid components, resulting in three joints (Perforated, Semi-perforated, and Solid) with distinct ligament meshes.

Secondly, we characterize the joints through spring and damper experiments to provide insight on how the modification changes the dynamics; these experiments expose significant nonlinearities in the system. In the full paper, we show how design parameters can modulate the spring-damper terms, leading to future work of programmable dynamics [30]. Measured parameter values inform a simulation in MuJoCo, enabling the use of Sim2real transfer learning.

Finally, we present results from two model-free control methodologies which require minimal user input: a proportional-integral-derivative (PID) controller and the Soft Actor-Critic (SAC) algorithm. Three tasks are designed to evaluate the accuracy of each controller: 1) Increasing steps, 2) Decreasing steps, and 3) Sinewave. The SAC policy yields greater precision through a wider range than

the tuned PID controller, and robustly solves the presented tasks generalizing through the different joint designs. In the case where SAC loses performance, a brief retraining period resolves the discrepancy. Our results suggest a method for designing soft-rigid robotic joints with parameterizable stiffness and damping, where model-free DRL compensates for variability and non-linearities.

2.1.1 Design concept for hybrid joints

We propose a process for constructing hybrid soft-rigid joints featuring a three-dimensional elastic mesh element which interweaves through rigid components. This approach effectively connects two rigid structures (such as the “ball” and “socket”) through a ligament-like elastic mesh that maintains compliance and promotes stable joint movement [42, 43] while allowing movement in major and minor degrees of freedom. We show an application of this design process by trying to achieve similar mechanical dexterity and robustness to a human arm; we design a manipulator using these hybrid joints combined with biologically inspired features. Carbon fiber rods take the place of bones while motors drive cables providing actuation in accordance to generalized muscle-tendon movement. This approach takes advantage of passive properties in the human arm that assist in joint motion and supports design modifications that result in a range of dynamic properties.

Fig. 2.2 shows our design of a cross-hatched network of holes within the joint bodies which the mesh engulfs. The two sets of through-holes lie in perpendicular planes and are aligned with the axis of intended motion. The combination of silicone surrounding the joint and within the through-holes creates a system of spring-dampers which can be modified by design parameters including hole quantity, outer-membrane width, and choice of flexible material. The human shoulder can be understood as a ball-and-socket joint and the elbow as a hinge joint [44], informing our design (Fig. 2.2A and B). We define the primary axes of rotation for the ball-and-socket joint as the y directions, and the hinge joint as the x direction. For both the shoulder and elbow joint, three designs involving differing perforations (Fig. 2.2C–E) were evaluated.




Figure 2.2: Joints within molds and three experimental designs. **A**, Shoulder joint components within mold. **B**, Elbow joint components within mold. **C–E**, Experimental designs for the shoulder ball component. Only one joint component is featured on the bottom to elucidate the quantity of through-holes. Adjusting perforation changes the amount of negative space, leading to changes in the amount of resulting silicone which then modulates dynamics.

2.1.2 Control strategies

Two methodologies for controlling the robotic arm were evaluated: 1) a proportional-integral-derivative (PID) controller fit via hyper-parameter optimization and 2) the model-free Soft Actor Critic (SAC) reinforcement learning algorithm. Model-based control techniques were not implemented due to the complex dynamics of

the system; however, we note that the system has the potential to be accurately modeled. Mean squared error (MSE) was used in both tuning the PID values and evaluating each control policy.

2.1.3 Network architecture

SAC was chosen as the algorithm to update network weights; the network architecture plays roles in the computational capability. Goal-conditioned algorithms allow for accomplishing a variety of tasks through policies that depend on specific goal inputs [45]. Naturally, goal-conditioned algorithms can be formulated similarly to controllers which require set points. Since there are delays in the hardware from sensing, communication to the main computer, and other unaccounted for attributes, the current state is non-Markovian. To computationally alleviate this issue, the ten previous observations were included. Both the simulation and physical prototype have a sampling rate of 100 Hz, permitting the history to account for 10 ms delays. One single observation (not including history) contains joint angles x, y recorded from the IMUs, joint velocities $\dot{x}, \dot{y}$ calculated from the angles, previous action a_{t-1} , and the goal joint angles g .

The policy network outputs actions for each motor parameterized as a Gaussian distribution, with outputs corresponding to the mean μ and log standard deviation $\log \sigma$. During training, actions are sampled from this distribution for

each motor to induce better exploration, but during evaluation the mean action is used. The policy output sends commands to each of the two motors which actuate over their corresponding dominant axis.

2.1.4 Reinforcement learning framework

The reinforcement learning problem is defined as maximizing the total discounted future reward $R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$, where γ is a discount factor and r_t is the reward given at time-step t . At each time-step, the agent uses a policy $\pi(s_t) = a_t$ to select an action a_t based on the current state s_t . The state is iterated forward and yields the reward via $(s_{t+1}, r_t) = f(s_t, a_t)$ based on the environment dynamics given by function f . The reinforcement learning algorithm attempts to find the optimal policy, $\pi^*(s_t) = a_t$, which returns the action that maximizes R_t .

The value function of a state is defined as $V(s) = \mathbb{E}(R_t | s_t = s)$ while the action-value function is $Q(s, a) = \mathbb{E}(R_t | s_t = s, a_t = a)$, where the probability distributions of actions are yielded from $\pi(s)$. Actor Critic algorithms are built to parameterize estimators for the value and policy functions [46]. For the deep reinforcement learning variant, the estimators used are feed-forward neural networks.

2.1.4.1 Soft Actor Critic algorithm

Recent model-free reinforcement learning techniques such as Proximal Policy Optimization (PPO) [47], Trust Region Policy Optimization (TRPO) [48], Asynchronous Advantage Actor Critic (A3C) [49] and Soft Actor Critic (SAC) [50] have the capacity to accomplish control tasks in both simulated and physical robotic control tasks [51, 52, 53, 54, 55, 56]. SAC is a sample-efficient and stable reinforcement learning algorithm which follows policy gradient methods, enabling use on continuous action spaces. An entropy maximization formulation encourages exploration and results in a policy that is not overly sensitive to hyper-parameters, making it a prime candidate for this robot.

SAC increases sample efficiency and stability of the training process by increased exploration through rewarding actions which result in higher variance of the policy [50]. This is given by the policy that solves the maximum entropy objective:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma (r_t + \alpha \mathcal{H}(\pi(\cdot|s_t))) \right] \quad (2.1)$$

where the expected value of $\mathcal{H}$ defines the entropy of the policy given the current state, and α is a temperature term that scales the impact of the entropy. The entropy term $\mathcal{H}$ is calculated as $\log \pi(a_t|s_t)$.

Feed-forward neural networks are optimized to minimize the objective functions for the value, Q-value, and policy functions. The corresponding objective

functions are:

$$\begin{aligned}
J_V &= \mathbb{E}_{s_t} \left[\frac{1}{2} (V(s_t) - \mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t) - \log \pi(a_t|s_t)])^2 \right] \\
J_Q &= \mathbb{E}_{(s_t, a_t)} \left[\frac{1}{2} (Q(s_t, a_t) - (r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} [V(s_{t+1})]))^2 \right] \\
J_\pi &= \mathbb{E}_{s_t} [\mathbb{E}_{a_t} [\alpha \log(\pi(a_t|s_t)) - Q(s_t, a_t)]]
\end{aligned}$$

which are alternately minimized until convergence. Further information on objective minimization and modifications are presented in [55]. The implementation used in this work comes from the Stable Baselines library [57].

2.2 Results

Three separate tasks were selected to test the effectiveness of the SAC and PID algorithms: (a) Increasing steps, (b) Decreasing steps, and (c) Sinewave trajectories. Using these experiments, we determine the ability of both control methods to handle smooth and discrete goal-position changes while compensating for gravity, mechanical hysteresis, and other non-linearities. Both joints were evaluated simultaneously, thus errors are dependent on each other. However, we assume the measurements are disentangled due to the orthogonality of the actuated axes.




Figure 2.3: Task-based trajectories and errors. Evaluation of control algorithms on three separate tasks with the shoulder joint (left) and elbow joint (right). The goal and experimental trajectories of multiple trials are shown on top, while the corresponding mean error is shown on bottom with an orange tint. **A**, Increasing steps task used to train SAC, which is most similar to the PID tuning task. **B**, Decreasing steps task highlighting asymmetry of joint dynamics. **C**, Sine wave task including a SAC-retrained agent that has been fine-tuned to properly capture dynamic goals.

To see how well each model adapted to new, unfamiliar static and dynamic tasks, we trained the SAC algorithm and tuned the PID controller on a static task (Increasing steps) shown in Fig. 2.3A. For the dynamic task (Sinewave), a decrease in SAC performance was resolved with a brief retraining period which compensated for the unencountered domain. The PID controller did not require re-tuning because it was tuned about a single point within the range of motion and does not differentiate between static and dynamic tasks.

Overall, the SAC agent showed lower error for both the elbow and shoulder joints during the Increasing steps task as shown in Fig. 2.3A. The PID controller was only able to perform comparably when operating near the position the control coefficients were tuned for (35°), otherwise having higher errors. The PID controller performs worse on the shoulder joint than the elbow, likely due to uncompensated gravitational forces and the bidirectional mechanical hysteresis (depicted in the full paper [30]). In contrast, the SAC agent performs consistently throughout each joint’s entire range of motion, indicating that the system’s transient dynamics have been characterized by the policy network.

Fig. 2.3B shows the Decreasing steps task, where the SAC agent continues to generalize well. The PID controller failed to minimize the steady state error throughout the range of the experiment, and we see that it performs better on the elbow joint compared to the shoulder joint indicating that the elbow is a simpler

control task. In both cases, SAC keeps reducing in error and the PID controller does not.

In Fig. 2.3C, the shape of the error graph shows a “bouncing” behavior, where the minimal error points represent the time where the true position crosses over the goal trajectory. The PID generalized better than the original SAC for the elbow joint and comparably for the shoulder, explained by its inability to distinguish between task types and more direct control of the elbow, yet by design it is not able to compensate for the lag involved in the task. We see that the original SAC algorithm cannot solve this task as well, having much larger error magnitudes indicating failure of generalization. Since the task involves a dynamic trajectory rather than static goal positions, a retrained version of the same algorithm was tested to see if performance improved. The SAC-retrained agent performed significantly better, outperforming the other control methods and inherently developing lag compensation.

Overall, the SAC agents perform best in 23/27 cases compared to the PID controller. The results table can be found in the full paper [30]. For the most important measurement—the Combined error—SAC outperforms PID control in 9/9 cases. Note that the SAC results include the SAC-retrained.

2.3 Discussion

We propose a new step in flexible robotic joint design, where elastic ligament meshes interweave traditional robotic joints to introduce compliance while maintaining stability. This novel class of robotic joints was used to construct a humanoid arm, demonstrating how the ligament mesh design can be manipulated to adjust dynamic characteristics; thus, three designs were constructed. Traditional robotic joint components (i.e. ball-and-socket and hinge joint components) and molds were designed to accommodate ligaments and 3D printed with negative space to cast a surrounding and interweaving elastic membrane. The simple casting process allows for rapid, low-cost prototyping, opening an avenue for creating systems with dynamic properties through simple mechanical modifications in 3D CAD. The variable nature of this system was accounted for through model-free controllers which successfully control the manipulator across joint designs and tasks.

Chapter 3

Optogenetic Modulation of Epileptiform Activity in Human Brain Slices

The following content in this Chapter is adapted from a journal article published in *Nature Neuroscience* by Nature Publishing Group [58], for which I served as a contributing author alongside co-authors John P. Andrews, Jinghui Geng, Kateryna Voitiuk, Matthew A.T. Elliott, David Shin, Alex Spaeth, Albert Wang, Lin Li, Daniel Solis, and others. My contributions included helping develop the base software for closed-loop electrophysiology as well as analysis relating to the activity and heatmaps used in the experiments. The analysis and closed-loop

control infrastructure utilized in this work was built on an early version of the BrainDance platform (<https://braingeneers.github.io/braindance>) which enabled the optogenetic intervention experiments described herein.

3.1 Introduction

Epilepsy is a disease of aberrant neuronal activity that results from an imbalance of excitation and inhibition [59, 60, 61]. Francis Crick famously envisioned the potential of using light to precisely control neuronal populations in the human brain [62, 63, 64]. Since their discovery, light-responsive channelrhodopsins [65] have become a cornerstone of modern neuroscience studies [66, 67]. Although optogenetic interventions have shown tremendous promise in non-human disease models of epilepsy [68, 69], whether such tools will translate into human neural circuits, and under what conditions, remains unclear.

Although direct application of optogenetics in the human brain is not currently possible, it has been translated for use in human retina [70, 71]. The use of optogenetics in research on human brain tissue has the potential to uncover key parameters for evaluating the potential efficacy of alternative therapeutic strategies. Development of any such therapies involving vector delivery of activity-modulating genes would benefit from data in a directly relevant model, namely human brain tissue.

Human brain tissue slices have been used for decades to probe mechanisms of epilepsy [72, 73, 74, 75], and single-neuron patch-clamping data exist showing that cells transduced with an activating channelrhodopsin in human brain slice can be depolarized with light [76]. However, optogenetics has never been used to modulate network activity in a human brain slice model. Activation of inhibitory GABAergic neurons [68, 77] and silencing of glutamatergic neurons [78, 79] have both shown promise in animal models of epilepsy. Still, neither of these approaches has been demonstrated in a human brain slice model.

3.2 Human hippocampus slice culture and neuronal transduction

To model human epileptiform activity, we established human organotypic tissue slices from hippocampus tissue (Fig. 3.1a) donated to research by patients with drug-refractory epilepsy. Tissues were obtained from patients both with and without hippocampal sclerosis as defined by International League Against Epilepsy (ILAE) criteria [80]. Resected tissue was sliced to 300 μm and cultured on cell inserts at the air-liquid interface in serum-free media as previously described [74, 81]. Slices were transduced (Methods) with AAV9 carrying an HcKCR1 [82] transgene driven by a CAMK2A promoter and a fluorescent tag (enhanced yellow

fluorescent protein (eYFP)). Fluorescent reporter expression was observed in live imaging with an epi-fluorescent microscope by day 4–6 (transduced on day 0). On the day of experiments, the slice was plated with neuron-dense areas of interest positioned over the high-density microelectrode array (HD-MEA) recording surface (Fig. 3.1b). A scan of the recording area electrodes was performed immediately after plating the slice to observe areas of spontaneous activity (Fig. 3.1c). After experiments, slices were fixed, and reporter expression was confirmed by immunohistochemistry (Fig. 3.1d). Reporter expression was enriched in neurons of the granule cell layer and the dentate gyrus as well as pyramidal cells (Fig. 3.1e). The proportion of neurons expressing eYFP in areas of high transduction ranged from 10% to 54%. This variability may be the result of focal micropipette injection. Slices that did not show detectable eYFP were excluded from the analysis of transduction rates (Methods). Together, these results indicate that we established organotypic tissue slice cultures of the human hippocampus and reproducibly delivered genetic payloads into hippocampal CAMK2A+ neurons using a combination of AAV9 and a CAMK2A promoter-driven transcript.




Figure 3.1: HD-MEA recordings of human hippocampal slices. **a**, Workflow of hippocampal slice recordings of HD-MEAs. Hippocampus specimens were collected on postoperative day 0 (POD0, day of surgery), sliced and plated on cell culture inserts. Transducing viral vectors was done on the same day as surgery. After 5–8 d in culture, slices were plated on HD-MEAs for recording experiments. **b**, NeuN staining of a slice with a black rectangle overlaying the HD-MEA recording area (left) and higher magnification of the same image with black rectangle denoting the HD-MEA recording surface (right). **c**, Activity scan showing amplitude of spike activity in the area denoted by the rectangle in **b**. **d**, Co-localization of eYFP with neuron-dense areas of the same hippocampal slice from **b** (slice 10F), one of eight slices transduced with AAV9-HcKCR1 with correlative electrophysiologic data. **e**, Left, insets of areas shown in **d**, correlating with the GCL of the dentate gyrus. Right, area of CA1 showing areas of pyramidal (PYR) cell morphology. **f**, Top, raster plot of unit activity with average unit firing rate (Hz) overlaid in red, demonstrating a sample of rhythmic bursting activity after adding KA to a hippocampal slice (slice 10F). Bottom, local field potential from select electrodes from recording above, showing discrete increases in local field potential frequency bands strongest in the theta range. Scale bars are shown in microns.

3.2.1 Design of optogenetic system

The purpose of the designed hardware platform is to trigger an optogenetic stimulation pulse and observe the response of neurons on a high-density electrode array. The optogenetic system was designed to be easy to replicate and does not require custom printed circuit boards (PCBs) or soldering (Assembly instructions are provided in the full paper [58], as well as the bill of materials and relevant software).

To run an experiment, the user must specify a calibration CSV file, which contains data on an LED's power output (see 'Optogenetic power density measurement and calculation' subsection in the full paper [58]). The user also connects to the Arduino over a USB serial port. Afterwards, the program creates a stimulation log CSV file into which the stimulation events with their parameters are recorded. Now, the user can construct their own stimulation sequences and use several helper functions with configurable stimulation pattern templates (Fig. 3.2f).

As the user runs the experiment, stimulation events are logged in both the stimulation log CSV file and the MaxWell voltage recording .hdf5 file. The MaxWell .hdf5 file has a field that stores a single bit value representing the GPIO pin signal ('1' for on, '0' for off) for every frame of the recording (every 0.05 ms assuming 20-kHz sampling). Therefore, we can extract at which recording frame the light

turned on and at which recording frame the light went back to being off. The CSV log file stores a timestamp of each event as well as the intensity of light used and the duration of ON and OFF requested by the user.

3.2.2 Optogenetic system software low-level functionality

The computer sends the Arduino an integer value over USB serial connection corresponding to the desired intensity. Although the intensity is represented to the user as a decimal number between 0 and 1, it is converted to the nearest proportional integer in the range of the 12-bit digital-to-analog converter (DAC) (0–4,095). The range 0–4,095 was chosen instead of the full 12-bit integer range of 0–4,096 because it more closely corresponded to 0–5-V DAC output when measured on an oscilloscope.

The computer also sends the electrophysiology recording system a signal to set GPIO high. GPIO high is +2.0–3.3 V, and anything less than or close to zero is considered GPIO low. When the GPIO signal is high, the Arduino sends a value to the DAC. Thus, to turn the light on, the computer needs to send a non-zero intensity value to the Arduino and also trigger the recording system to set GPIO high.

3.2.3 Activity detection algorithm and closed-loop optical intervention

We implemented an algorithm that monitors spike activity in real time, using data from the HD-MEA (Fig. 3.3). This algorithm is meant to demonstrate a therapeutic use-case scenario. When a predefined spike threshold is detected, the algorithm triggers illumination of the slice. Spike activity was measured by counting the total number of 5-root-mean-square (rms) spikes that occurred in 10 s across all the electrodes on the HD-MEA.

Experimentally, this threshold was to be 2–4 times above the average baseline 5-rms spike activity for 1 min before application of bicuculline. After the drug was applied, activity quickly surpassed this threshold. After the threshold is surpassed for a 10-s window, the algorithm triggers 10 s of illumination. Afterwards, the algorithm goes back to monitoring activity. This algorithm is intentionally simple in nature, to provide a proof of concept that user-defined signals can be detected in real time using HD-MEA interface with this *ex vivo* model, in a closed-loop process.

3.2.4 Optogenetic control of human hippocampal network

HcKCR1 encodes a kalium channelrhodopsin—a potassium-selective, light-sensitive ion channel that hyperpolarizes the neuronal membrane, thus reduc-

ing the probability of spiking when activated by 530-nm light [82, 83]. Unlike a more traditional Channelrhodopsin-2 (ChR2), which encodes a sodium-selective ion channel, inhibitory channelrhodopsins have never been applied to human tissue. To confirm that HcKCR1 illumination drives membrane hyperpolarization, intracellular voltage-clamp recordings confirmed hyperpolarizing currents in human hippocampal neurons, consistent with the expected channel physiology [58]. To assess whether AAV-based transduction of human neurons with optogenetic constructs is specific to HcKCR1, we transduced additional slices with an excitatory channelrhodopsin (AAV9-CAMK2A-ChR2-eYFP) packaged in the same viral vector. In this case, we observed a similar extent of expression, suggesting that our methods are not limited to HcKCR1. Together, these results demonstrate the feasibility of applying inhibitory optogenetics to human neurons.

3.2.5 Optical HcKCR1-mediated inhibition of network activity

HcKCR1 activation suppressed activity in human hippocampal neurons across all experimental conditions of spontaneous and provoked activity (Fig. 3.3). Illumination with 530-nm light [82] resulted in rapid silencing of hippocampal activity.

Intensity scans across the full range of light intensity produced by the LED driver were carried out to determine the dynamic range of the control of neuronal

firing rates as a function of illumination power. The intensity threshold was typically between 26.2 mW mm^{-2} and 31.2 mW mm^{-2} , correlating to 40–50% of the possible output from our drivers. For experiments, we used one step above the minimum threshold (that is, 50–60% intensity, translating to $31.2\text{--}35.8 \text{ mW mm}^{-2}$).

Decreases in neuronal firing rates during HcKCR1 optogenetic activation were obtained from $n = 8$ ex vivo hippocampal slices. We additionally carried out recordings in a hippocampal slice transduced with an identical AAV vector driving ChR2 expression (AAV9–CAMK2A–ChR2–eYFP), which demonstrated an increase in unit activity with 470-nm light illumination. This result confirms that the specificity of light-mediated change in neuronal activity is determined by the channelrhodopsin-specific physiology and not an artifact of light illumination.

AAV mediated HcKCR1 expression is sufficient to enable optogenetic modulation of network-wide spontaneous activity despite incomplete transduction. Spontaneous activity arises from networks perched in a balanced regime that is prone to neither silence nor saturation [84]. As a result, relatively small increases in inhibitory tone at a cellular level may be insufficient to offset the imbalance necessary to generate epileptiform activity at a network level.

3.2.6 Optogenetic suppression in provoked hyperexcitability

Next, we tested HcKCR1-mediated optical inhibition in the context of hyperexcitable conditions.

We used magnesium-free (0-Mg) media, a long-standing approach used to increase neuronal excitability and induce epileptiform activity in hippocampal preparations [72, 85, 86]. Clear, spontaneous bursts of activity were observed in two of four slices bathed with 0-Mg media (Fig. 3.3c). These bursts of activity were less than 1 s in duration and several minutes apart, reminiscent of interictal epileptiform spikes [87]. During LED light illumination of slices bathed in 0-Mg media, mean firing rates were reduced compared to the rates in the absence of light illumination, although not all units responded uniformly. On a unit-level analysis, 15–40% of units in these slices had a decrease in firing rate of at least 90% during light-ON conditions. Spike events during 0-Mg recording conditions were not completely absent during optical inhibition of two separate slices (Fig. 3.3c). The spikes were far enough apart in time to prevent strong conclusions about the effects of optical inhibition on these events.

Addition of kainic acid (KA) has been used in various epilepsy models [88, 89, 90]. Based on other groups reporting the effectiveness of combining multiple proconvulsant pharmacologic interventions to induce epileptiform activity [88, 91], we

tested whether adding kainate to 0-Mg bathed slices would lead to epileptiform activity. The addition of KA resulted in acute elevations in firing rate, and, in two of four slices recorded with KA, rhythmic bursts of activity occurred (Fig. 3.3d). Slices exhibiting this rhythmic bursting of unit activity also showed increased coherence across frequency power bands, most clearly in the theta band (Fig. 3.3f,g). Repeated 10-s periods of illumination reduced firing rates during light-ON conditions (Fig. 3.3d), but a causal effect of illumination on the coordinated bursts is less clear in the absence of more comprehensive characterization of the typical effects and expected frequency of rhythmic bursting with addition of KA.

To explore this platform’s ability to perform closed-loop illumination, we used a GABAAR blockade via the application of bicuculline to elevate firing rates above their baseline levels (Fig. 3.3j,k). We used the reproducible increases in firing rate with GABAAR blockade to detect and disrupt elevations in network firing. We used an automated, real-time spike event detection algorithm to trigger LED illumination when a firing rate threshold is crossed (Fig. 3.3h–k). The closed-loop illumination software is user-activated after bicuculline application, and when firing rates surpass a predefined threshold, it triggers 10 s of continuous illumination. This markedly reduced the firing of recorded units (Fig. 3.3h–k). This was replicated in $n = 3$ HcKCR1-expressing slices derived from $n = 2$ patients. These results demonstrate a potential benefit of optogenetic interventions

in epilepsy: modulation of brain activity restricted precisely to time and anatomy in which pathologic elevations in neuronal firing occur.

In these three distinct methods of provoking a hyperexcitable state (0-Mg, 0-Mg + KA and GABAAR blockade), HcKCR1 activation in a subset of CAMK2A-expressing neurons reduced network firing rates. The magnitude of firing rate reductions in 0-Mg and 0-Mg + KA conditions was generally less than that observed in physiologic media and GABAergic blockade conditions. Although several variables may contribute to this, HcKCR1-mediated inhibition of pre-synaptic CAMK2A-expressing neurons may be partially offset by post-synaptic glutamatergic receptor activation in 0-Mg and 0-Mg + KA conditions. Whether HcKCR1 activation of CAMK2A-expressing neurons effectively interrupts provoked coherence seen in 0-Mg + KA conditions will require additional experiments to thoroughly characterize the effects of the rhythmic burst activity induced with 0-Mg + KA.




Figure 3.2: Optogenetic device for MEA. **a**, Hardware components: the MaxWell CMOS HD-MEA system is augmented with an optogenetic platform, consisting of an Arduino with a DAC controlling an LED driver to modulate light timing and intensity delivered to the biological sample on the HD-MEA. The optogenetic platform uses off-the-shelf optoelectronic equipment and 3D printable components (the lid and insert) for reproducibility. The MaxWell natively supports electrical stimulation. **b**, Cross-section of optogenetic well insert showing main components. **c**, Close-up of the optical fiber creating illumination. **d**, Software modules: code and data flow supporting optogenetic and electrical stimulation wrapped into Python libraries. Additional files store calibration for the expected LED power outputs, a log of stimulations delivered during an experiment with respective parameters and MaxWell voltage recording data. Cyan elements are specific to optogenetic stimulation. Green elements are specific for electrical stimulation. All other software elements are shared. **e**, Experimental setup inside the incubator showing the MaxWell chip with the optogenetic insert during an optical stimulation event. **f**, Configurable stimulation templates are part of the software modules.

3.2.7 Activity heatmaps

Activity heatmaps in Figs. 3.1 and 3.3 show the spatial distribution of significant voltage events. The MaxWell software computes a moving rms value for each electrode, and we used this to determine events that crossed a threshold set at $5\times$ the rms value of each respective electrode. A 2D spatial grid of the counts per second was constructed, and then a 2D Gaussian blur with radius, r , was used to enable better visual continuity. Each position in the grid was divided by $2\pi r^2$ so the values displayed would correspond to events per second, making it so the Gaussian blur does not significantly alter the values of each data point. These values were then plotted as the activity heatmaps, where warmer colors represent areas of higher event frequency, and darker colors denote regions with less activity.




Figure 3.3: Optogenetic inhibition of human hippocampal activity. **a,b**, Stacked raster and electrode heat maps (slice 3C, HcKCR1) across five 10-s trials (pre/ON/OFF); mean firing rate (red); black: HcKCR1-eYFP. **c,d**, Rasters from slice 10F in 0-Mg (**c**) and after KA (**d**). **e–g**, Expanded burst epoch (**e**),

LFP spectrum (**f**), and theta-band coherence (pink; KLD) with firing rate (**g**). **h,i**, Rose plots of theta phase: non-random in **g**(I), random in **g**(II). **j,k**, Raster and heat maps showing optogenetic inhibition after bicuculline blockade (slice 5C); black: HcKCR1-eYFP. **l,m**, Closed-loop illumination schematic (**l**) and silencing of bicuculline-provoked activity (**m**). Shading: SEM. fr, firing rate.

Chapter 4

Internet-Connected Cortical Organoids for Project-Based Stem Cell and Neuroscience Education

This chapter has been adapted from a journal article published in *eNeuro* by the Society for Neuroscience, for which I served as first author alongside co-first authors Matthew A.T. Elliott and Hunter E. Schweiger. The adapted sections focus on the portion of the study related to the internet-based stimulation and recording of organoids. For information on the other course, see the full paper [19].

This methodology used an early version of the BrainDance software, providing a study on the usability and outcomes with undergraduate students.

4.1 Introduction

The use of cortical organoids in academic research and biotechnology has drastically increased in recent years [92]. For example, cortical organoids are used as tools for drug discovery [92], to study genetic mutations underlying disease [93, 94], and to study infectious diseases that affect the brain, such as Zika and COVID-19 [95, 96]. Given their applications and the trends in the use of these models in research and development, there is a growing need for students to be trained in stem cell culture and cortical organoid generation, manipulation, and analysis.

Internet-enabled technologies can be a powerful tool to overcome this gap. When laboratory equipment is connected to the cloud, scientists and students can easily access and manipulate these technologies remotely [97, 98, 99, 100]. Moreover, this approach allows exceptional platform scalability, allowing dozens or hundreds of users simultaneously [101]. In the classroom, Internet-enabled technologies have been used to drive project-based learning (PBL) approaches in a variety of settings, including the use of microscopes for remote observations of behaviors of model organisms and cell cultures [102, 97, 101], the use of sensors to monitor environmental conditions [103, 104], and even the use of lab-on-a-chip

microfluidic devices to perform bacterial detection in water samples [105].

We take advantage of Internet-enabled technologies to introduce cortical organoid PBL in two different university courses: an introductory tissue culture course, where students monitor the growth of organoids in the presence of different drugs, and an advanced neuroscience mathematics course, where students stimulate and record from cortical organoids using multielectrode arrays (MEAs). In both approaches, the students can perform PBL over several days. We showed that Internet-enabled PBL leads to a high level of interest in neuroscience and stem cell biology in the students in both courses. Altogether we propose that Internet-enabled PBL can become an effective approach to deliver complex concepts in stem cell and organoid culture. Here we focus primarily on the internet-based stimulation and recording of organoids. For information on the other course, see the full paper [19].

4.2 Methods

4.2.1 Ethics statement

The University of California Santa Cruz Institutional Review Board (IRB) reviewed this work at the proposal stage and determined that it did not constitute human subject study. Therefore, the work was IRB exempt.

The students who performed the electrophysiology experiments were enrolled in the BME118: Mathematics of the Mind course at the University of California Santa Cruz. A total of 24 students was enrolled in this course. We performed two surveys: one before the experiment, and one after the experiment. For the pre-experiment survey, 18 students responded. For the postexperiment survey, all 24 students responded.

4.2.2 Electrophysiology experiments

Using the Maxwell Biosystems software, activity scans were performed every 2–3 days. We sampled signals from 1024 of the $\sim 26,000$ electrodes at a time in a sweeping formation across the MEA at an interval of 30–45 s for each configuration. A base level of activity was measured by electrical events which crossed the 5-root mean square (rms) noise threshold. If activity was detected, 5- to 12-min recordings were taken using Maxwell Biosystems’ electrode selection algorithm which maximizes clusters of electrodes around hotspots of high relative firing rate. Neurons were identified by analyzing the neuronal footprint, the signal averaged waveform across a patch of multiple electrodes surrounding the potential neuron. Neurons chosen to be stimulated required detection of at least 10 threshold-crossing events. We chose stimulation channels as the channels that recorded the highest amplitude signals from the neuron to have the highest prob-

ability of evoking action potentials [106].

4.2.3 Data processing

Raw electrical data from the MEA was saved during each of the experiments in the HDF5 format. These data were spike-sorted into individual units (putative neurons) using Kilosort 2 [107]. This process filters the data, whitens [107], then clusters neurons based on matching spike-waveform templates. For each identified neuron, this process outputs a list of spike times together with the spatial location on the MEA. To validate the biological plausibility of the identified neurons, expert human curation was conducted through the open-source software Phy (<https://github.com/cortex-lab/phy>). Putative neural units are retained if they have key biological features such as a biologically consistent shape and duration for their action potential waveforms. Each experiment recording was individually processed with the same set of spike sorting parameters. Electric stimulation artifacts and noise were removed during the curation. These data were delivered to the students in Python NumPy format allowing easy subsampling.

4.2.4 Electrophysiology experimental design

Mathematics of the Mind students were introduced to the idea of analyzing electrophysiological measurements by conducting analysis from a previous set of

published electrophysiological recordings in human organoids [108]. The students were taught to use measures including spike-time tiling [109], correlation, interspike intervals, and latency distributions to characterize neural circuit behavior. After this initial homework assignment, students were asked to propose a circuit perturbation experiment using electrode-supplied stimulation on live tissue to augment the underlying neural circuitry, where they would analyze spontaneous activity before and after the perturbation. The setup of the experiment involved a 5-min baseline recording, a 5-min recording with student-designed stimulation, and a 15-min poststimulation recording. The students were specifically tasked with designing and coding their own stimulation paradigm to be executed during their experiment, and proposing a hypothesis for changes that would occur between the baseline and poststimulation recordings. They were encouraged to be creative.

4.2.5 Stimulation programming

To make the task of stimulation paradigm design feasible and straightforward for the Mathematics of the Mind students, a simple application programming interface (API) was designed. The students were told that they could stimulate three neurons in any manner that the API enabled, and they could select the approximate distance of the three neurons as a proxy for connectivity. At a high

level, the students craft sequences of stimulations using the API, and choose a frequency to iterate through these stimulation sequences. The API enables the creation of stimulation paradigms by three simple commands: “stim,” “delay,” and “next.” The “stim” command has three parameters: (1) the list of neurons to stimulate, (2) the amplitude in millivolts to stimulate, (3) the width of the biphasic pulse. Recommended values of 150 mV for the amplitude and 200 μ s per phase for the biphasic pulse were given to students. The “delay” command has one parameter: the time to delay in milliseconds. It pauses the stimulation. The “next” command effectively ends the current stimulation sequence.

4.3 Results

4.3.1 Electrophysiology software programs run on organoids in a neuroscience mathematics course

A landmark in neuronal maturation is the acquisition of electrophysiological properties [110]. These properties follow a stereotypical cell type-specific identity [111, 112, 113, 114]. Given the great diversity of neurons in the cerebral cortex, it is expected that complex models, such as organoids, will have different electrophysiological profiles representative of their cell type composition [115]. Indeed, the emergence of complex networks in cortical organoids is an active area of research

that is expected to continue growing over the next few years [116, 110, 108, 117].

One of the most exciting aspects of the emergence of neuronal circuits is that it provides scientists with the ability to study how neurons in organoids “learn” by adapting to stimuli provided to them [118]. Neuronal plasticity can be evoked in cerebral organoids by sending electrophysiological stimulus patterns [118]. These patterns are written using computer code. Therefore, an organoid’s neuronal circuitry can be “programmed” to produce different types of predictable neuronal responses.

To date, PBL-based teaching of electrophysiology concepts in the classroom have been limited to either simulation of single neurons using web-based apps [119, 120, 121, 122] or the use of simple devices to activate and record neuromuscular electrophysiological signals in either invertebrates or the students’ own muscles [123, 124, 125, 126]. While remarkable, these experiments do not offer the capability to ask complex questions in circuit assembly and brain function. Specifically, students do not have the ability to program their own experiments, where they see how neuronal circuitry responds to their code. This is made possible using high-density multielectrode arrays (HD MEAs), in which thousands of electrodes can stimulate and record from a single circuit [127]. For example, the MaxOne system from MaxWell Biosystems has been used to record circuit activity in cortical organoids through its 26,400 electrodes [94, 108]. Yet, because

of the high costs and experimental complexity associated with these experiments, HD MEAs have not been introduced in the classroom. We have previously developed an Internet-connected electrophysiology platform that allows the users to record from MEAs remotely [128]. This technology could become key for cloud-enabled PBL [101], in which students use experimental platforms that are not located at or owned by their own institutions.

To deploy this technology for neuroscience education, we integrated it into the Mathematics of the Mind course at the University of California Santa Cruz. This course is a mixed upper-level undergraduate and beginning graduate class where most enrolled students are from quantitative fields such as computer science, physics, applied mathematics, and engineering. Students were taught concepts and given supplementary readings within topics such as tetanus-induced potentiation dependent on connectivity [129], Hebbian theory [130], spike-time dependent plasticity [131] and homeostatic regulation [132]. For the students to properly construct experiments, they were introduced to applied electrophysiology techniques like spike sorting, a neurophysiology technique that allows neuronal spikes to be clustered based on similarities [133]. To understand how performing live electrophysiological experiments impacted the students, we first surveyed them after the introduction of these concepts but before performing the experiments (Fig. 4.1). Of the 24 students enrolled in the class, 18 responded to the

precourse survey. We found that the large majority (83.7%) of the students had previous experience in mathematics and computer programming (Fig. 4.1A,B). In contrast, 66.6% of the students reported having little to no previous stem cell biology experience (Fig. 4.1C) and 94% of the students reported having little to no previous experience with neuroscience (Fig. 4.1D).

We then asked the students questions related to their interest in the class and the topics discussed up to this point, and before the execution of live experiments. We found that 94.4% of the students reported that they are interested in the applications of mathematics to biological processes (Fig. 4.1E), and all of them (100%) wanted to learn about Internet-of-Things technologies (Fig. 4.1F). Similarly, 94.4% of the students reported that this course required them to think at a deep level or use critical thinking (Fig. 4.1G). In addition, 83.3% of the students thought that this course allowed them to examine questions that matter beyond the classroom (Fig. 4.1H). When asked about the biological side of the class, 77.7% of the students were interested in learning about stem cells (Fig. 4.1I) and 94.4% of them were interested in learning about neuroscience (Fig. 4.1J). Similarly, 88.8% of the students claimed to be interested in learning about “connectoids” (Fig. 4.1K), long-range connected organoids [134, 135, 136]. When asked about the technologies to be used in the experiments, 66.6% of the students were comfortable using WetAI (Fig. 4.1L), our custom interface with the experiment [101],

and 72.2% of the students claimed to be comfortable using Jupyter notebooks to perform experiments (Fig. 4.1M). We found that 61.1% of the students said they were comfortable reading neuroscience literature (Fig. 4.1N), while 44.4% of the students felt comfortable performing high-density MEA experiments at this time (Fig. 4.1O). In addition, 66.6% of the students reported that they would be comfortable identifying spike burst signals in high-density MEA data (Fig. 4.1P). Altogether, we observed that, consistent with their academic background, most students are more comfortable with the quantitative aspect of the class than with the biological section, although they demonstrate a high interest in learning topics in stem cell and neuroscience.




Figure 4.1: Students’ previous experience and reported interest in topics related to mathematics, neuroscience and stem cell biology. **A – P** , Students’ answers to a survey assessing their level of experience and interest in the topics covered in the Mathematics of the Mind course. $n = 18$ students. A total of 14 of the 16 survey questions were statistically significant, using a one-sample two-tailed t test (p -value ≤ 0.05). **C** and **P** were not significant. y -axis represents percent of students.

Given the qualities of these students, we then asked them questions related to their mathematics self-concept, as this metric has been previously shown to be a predictor of performance [137, 138, 139]. To do so, we asked the students their level of agreement with different statements. We found that 61.1% of the students agreed with the statement “I am capable and skillful at Mathematics” (Fig. 4.2A), while 44.4% agreed with the statement “Being a good mathematics student makes me feel that my classmates and teachers think more of me” (Fig. 4.2B). These results are similar to previous reports of mathematics self-concept in university students [139]. However, unlike previous reports, we find that half of the students (50%) agree with the statement “In the mathematics examinations I feel unsure, desperate, and nervous” (Fig. 4.2C), while the minority of the students (44.4%) agree with the statement “My performance in mathematics largely depends on the methodology and empathy with teachers” (Fig. 4.2D). This self-concept profile is consistent with university students who are high achievers in online mathematics training [139], suggesting that our students were likely to succeed in a remote PBL experiment. Finally, when asked whether they agreed with the statement

“Mathematics is useful and necessary in all areas of life,” we found that 83.3% of the students agreed (Fig. 4.2E), while 88.8% of the students agreed with the statement “Mathematics is useful and necessary for a career in Biology” (Fig. 4.2F). The combination of the students’ quantitative skills and interests, as well as their self-concept consistent with high achievers in mathematics, allowed us to conclude that this cohort of students is ideal for implementing remote PBL experiments using high-density MEAs.




Figure 4.2: Mathematics self-concept of the students. A – F , Students’ answers to a survey assessing their mathematics self-concept. $n = 18$ students. Four of the six survey questions were statistically significant, using a one-sample two-tailed t test (p -value ≤ 0.05). C and D were not significant. y -axis represents percent of students.

4.3.2 Student-designed electrophysiology experiments

As part of the class, we asked students to work in groups to design and program stimulation patterns to be given to cortical organoids [140]. They were asked to propose a hypothesis on what changes they expected to see in network behavior following stimulation. The students then analyzed the results of their stimulation patterns and presented their findings to the class. Each group was assigned two organoids: one to be used as an experimental organoid, and one control (Fig. 4.3A,B). To help the students conduct their electrophysiology experiment, we designed a simple application programming interface (API). Each team of students had the opportunity to watch their stimulation experiment happen on the organoid live, in real time. The experiments were hosted via Zoom, with the class TA presenting the changes in neuronal activity that occurred while the organoids were being stimulated.

The API mentioned above provided students with considerable freedom in the types of experiment they could construct. One notable group, for example, translated the rhythm of a minimalist song: Steve Reich’s Clapping Music (Fig. 4.3C,D). In this song, two artists clap at rhythms that synchronize and desynchronize throughout the song [141, 142]. The students wanted to see whether this syncing of stimulation patterns might yield a Hebbian-like learning outcome. This rhythm has been previously used in mathematics problems in combinatorics and

group theory [141, 142]. The students translated each of the clapping sequences to an electrode stimulation pattern using the API (Fig. 4.3D). These patterns were then applied to the experimental organoid (Fig. 4.3E) and the students found that the stimulation pattern had a measurable effect on neuronal behavior (Fig. 4.3F).




Figure 4.3: Student experiments explore programming neuronal circuits. **A** , Students were assigned two cerebral organoids (test and control) placed on an HD MEA, which delivered their electrophysiology stimulus pattern. **B** , Students used online software to explore data, including neuron locations across both organoids, stimulation sites, and individual spike rasters. **C** , One team encoded the rhythm of Steve Reich’s duet, Clapping Music, sending each performer’s pattern to different neurons. Image from Reich’s music score. **D** , The online software was used to verify the stimulation pattern matched the musical rhythm; red and blue represent patterns sent to two different neuronal sites. **E** , Students analyzed spike raster data from before, during, and after stimulation to determine whether their stimulus evoked the predicted neuronal response. **F** , Example computational analysis: one team applied a statistical test to the cosine similarity score to assess whether stimulation altered neuronal activity, comparing spontaneous firing immediately after stimulation to that 10 min later and finding a significant difference. MCMC fit = Markov-Chain Monte Carlo fit.

The other groups created experiments of similar rigor and creativity (Fig. 4.4). Three of the five groups stimulated organoids with a high frequency tetanic pattern [143]. This experiment was popular because it was explained to have a high chance of prominent changes. Each group approached the problem using a unique protocol they designed. One group used the freedom that the stimulation API provided to design an experiment which generated a random walk over stimulation amplitudes, producing a unique sequence during each generation. This group was interested in whether these stimulation sequences produced unique patterns between neurons connected in a circuit.

The students were told to rigorously analyze the resulting data from their experiment to determine whether the experiment’s hypothesis was correct through techniques they were taught throughout the class: correlation matrix, spike time

tiling coefficient, interspike intervals, and interneuronal spike latencies. Most importantly, the students had to implement some form of original analysis not taught during the class. This could either be a method learned from reading literature, or a novel technique they design themselves. Most groups implemented either a basic statistical test to verify their hypothesis, or an unsupervised machine learning technique to help interpret results. However, some groups developed novel analysis methods. Of note, one group implemented the underlying algorithm used in principal components analysis (PCA) but using the spike time tiling matrix as the input into the algorithm (Fig. 4.4A–D). The group showed that this new technique was able to discern the test from control organoid based on neural signals (Fig. 4.4D).




Figure 4.4: Example of Mathematics of the Mind student-project applying tetanic pattern to cortical organoids. **A** , Spike time tiling matrix for the electrophysiology data recorded from cortical organoids. **B** , Eigendecomposition on the spike time tiling matrix to extract its eigenvalues and eigenvectors. For illustration, the first four eigenvectors are plotted. **C** , Comparison of eigendecomposition on the spike time tiling matrix to other popular correlation techniques. Specifically, they considered how the method's reconstruction error compared with that of performing principal components analysis on the correlation matrix. **D** , Superimposition of the first eigenvector's values (represented as small circles in the blue-yellow range) on top of their corresponding neural unit allowed the students to differentiate which neurons came from which organoid.

4.3.3 Organoid experiments lead to higher interest in stem cell and neuroscience research in mathematics students

To understand how the integration of HD MEA recording and stimulation of neuronal cell cultures in the classroom affects the students' interest in neuroscience and stem cell research, we surveyed the Mathematics of the Mind students immediately after their presentations. All 24 students enrolled in the class responded to the survey. We found that 95.8% of the students said that they enjoyed learning about organoids (Fig. 4.5A), and 91.6% enjoyed performing these experiments (Fig. 4.5B). In addition, we found that 91.6% of the students thought that performing remote electrophysiology experiments was interesting (Fig. 4.5C), while 95.8% of the students thought that the experiment selected helped solidify the concepts discussed in the class (Fig. 4.5D).

We then asked the students for their perceptions after performing remote electrophysiology experiments. We found that 83.3% of the students thought that performing these experiments allowed them to perform novel and complex experiments (Fig. 4.5E), while 87.5% of the students thought that this approach allowed them to do experiments that required complex training (Fig. 4.5F). 91.6% of the students thought that these experiments allowed them to study new areas beyond their academic program (Fig. 4.5G). A total of 95.8% of the students consid-

ered that performing remote experiments allowed them to do projects that are not available to most students around the world (Fig. 4.5H), while 62.5% of the students thought that performing remote experiments allowed them to multitask more than in person labs (Fig. 4.5I).

Then, we focused on understanding how performing these experiments influenced the students' interest in neuroscience and stem cell research. We found that 83.3% of the students thought that performing these experiments increased their interest in stem cells (Fig. 4.5J), while 87.5% reported an increased interest in organoids (Fig. 4.5K). Remarkably, we also find that after this course most of the students report a higher comfort in reading and discussing scientific literature in neuroscience (91.6%; Fig. 4.5L), stem cells (79.1%; Fig. 4.5M), and mathematics (66.6%; Fig. 4.5N).

Finally, we asked the students how this work affected their career prospects. We found that 87.5% of the students would consider continuing further education in stem cells and neuroscience (Fig. 4.5O) and 79.1% would consider applying to jobs analyzing stem cell and neuroscience data (Fig. 4.5P). Altogether we conclude that our approach positively affected the students' interest in stem cell and neuroscience.




Figure 4.5: Students report positive attitudes toward stem cells and neuroscience after organoid experiments. **A – P** , Postexperiment survey results for the Mathematics of the Mind course ($n = 24$). All 11 questions were statistically significant (one-sample two-tailed t test, $p \leq 0.05$). y -axis: percent of students.

4.4 Discussion

The scalability of electrophysiology projects would follow a different approach. While the HD MEAs used in these experiments have a high cost (\$200/chip) and require expensive capital equipment, there are $\sim 26,000$ electrodes in the HD MEAs, which can be distributed to lower the per user costs. Specifically, users can read from 1024 electrodes simultaneously and of those they can stimulate from [140]. We can switch to a new set of 1024 in seconds. This allows us to multiplex the system in various ways: by assigning different sets of electrodes for recording and stimulation to different subgroups within the class or individually to each student, depending on class size. For extended experiments, each student or group can design a stimulation pattern and our interface can apply these extended patterns to the same or different chips sequentially. In addition to the multiplexing of experiments, we have developed a strategy to scale the use of the software and lecture materials to reach more students. The coding lectures and homework were aggregated into a single public repository for future education (<https://github.com/Braingeneers-Education>). We integrated a newly de-

veloped cloud environment, GitHub Codespaces, to allow anyone to access course tutorials. The software is web based, requiring no downloads, and can be run from any computer or tablet.

Chapter 5

Goal-Directed Learning in Cortical Organoids via Closed-Loop Electrophysiology

This chapter has been adapted from a journal article published in *Cell Reports* [34], for which I served as first author. The work presents the central experimental contribution of this dissertation: the first demonstration of goal-directed learning in brain organoids. Through closed-loop electrophysiology, cortical organoids interfaced with high-density microelectrode arrays learned to balance an inverted pendulum. We used adaptively selected training signals guided by reinforcement learning to modify the biological circuitry. The experimental framework developed

here builds on the tooling and infrastructure described throughout this thesis, primarily the BrainDance platform (Chapter 7).

5.1 Introduction

The brain has an unmatched capability for intelligent and dynamic information processing—deep artificial neural networks require over five layers to achieve an adequate approximation of a single biological neuron [144]—yet biological neural networks operate at a fraction of the power consumption of artificial systems [145]. Understanding how they process and dynamically adapt to information has profound implications, from developing targeted stimulation protocols for neural rehabilitation to creating energy-efficient bio-electronic computation [146]. Reinforcement learning concepts provide a high-level blueprint for learning *in vivo* [147, 148, 2, 149], and low-level learning rules including Hebbian plasticity and synaptic dynamics have been studied extensively [150, 151, 152, 153, 154, 155, 156]. However, despite extensive research at both scales, we lack a clear understanding of how low-level synaptic modifications give rise to goal-directed learning. This gap limits our ability to harness biological learning.

Studies of neural adaptation *in vitro* could reveal these mesoscale learning principles. Pluripotent stem cell-derived brain organoids are an ideal model for *in vitro* learning, as they recapitulate key aspects of neuronal development, including di-

verse cell types and complex electrophysiological circuitry [5, 157]. Although these organoids can model many aspects of development, they are typically grown without key inputs from other brain regions or sensory organs [158, 6], limiting their ability to model *in vivo* neurodevelopmental processes that depend on such interactions. High-density microelectrode arrays (HD-MEAs) can both record from and stimulate organoids at single-cell resolution [159], enabling precise control of the electrophysiological aspects of an organoid’s interaction with its surroundings. This fine-grained control allows us to systematically perturb neuronal networks within organoids and longitudinally observe their responses, providing a platform for iterative investigation of learning at the circuit-level.

Within the last three decades, labs have used varying methodologies of electrophysiological stimulation in efforts to harness biological computation through closed-loop interaction [9, 11, 10]. Pioneering work repurposed vestibular circuits of a lamprey to control a robot [22], leading to numerous studies embodying cultures into closed-loop tasks [8, 24, 160, 25, 161, 162, 163, 164]. Subsequently, different frameworks for *in vitro* neural adaptation have emerged, influenced by factors like species, brain regions, neuron distributions, and network architectures. Early work accomplished supervised learning through slow frequency stimulation which evoked network bursts, ceasing stimulations when the desired effect was achieved [7, 165], formalized as learning by stimulation avoidance [166]. Many

approaches instead use tetanic (high-frequency) stimulation patterns as training signals [151, 9]. These tetanic electrical stimulations (> 5 Hz) have been shown to induce functional and stimulus-evoked changes for pattern recognition [23] and burst information [151], and have been explored with respect to LTP/LTD [167, 21, 168, 169], and embodied tasks [24, 160, 161]. Tetanic stimulation is thought to induce changes through associative mechanisms due to increased neural activity.

More recent work on reservoir computing with organoids, which leverages machine learning readout layers, has demonstrated success in unsupervised tasks like predicting dynamical equations and classifying speech snippets [10], as well as a control task [163]. The free energy principle has been proposed as a theoretical framework to explain success in blind source separation tasks [170, 171] and control in a simplified “Pong” environment [11], drawing parallels to artificial reinforcement learning approaches to curiosity [172]. However, other work has shown that even inanimate hydrogels, when electrically embodied in a “Pong” task space, can demonstrate improvement [173], highlighting the importance of appropriate controls and careful interpretation and investigation of apparent learning in these systems. A key insight from previous studies is that repeated and formulaic high-frequency stimulation patterns can induce persistent shifts in network dynamics [21, 9]. These shifts effectively reconfigure the network’s input-output mapping,

potentially enabling real-time adaptation. Building on this principle, we hypothesized that by systematically selecting which neural pairs to stimulate based on task performance, we could guide these network reconfigurations toward improved control policies.

While previous work has demonstrated that stimulation can modify synaptic plasticity, key questions remain about the optimal selection of target neurons, stimulation frequencies, and timing. In this paper, we focus on feedback-driven neural plasticity—on the scale of seconds to tens of minutes—as a form of short-term learning. We propose a systematic approach to characterizing and training biological controllers using high-frequency stimulation within a task-based framework (Fig. 5.1). We embody cortical organoids into an environment that simulates the classic problem of inverted pendulum stabilization [174] (similar to balancing a ruler vertically on your hand). In this problem, commonly known as “Cartpole” [175, 176], a simulated two dimensional cart must be moved along the horizontal axis to balance a vertical pole which is attached by a hinge. Unlike pattern recognition tasks, Cartpole requires continuous, active stabilization of an inherently unstable system where small lapses in control lead to failure—making it particularly suitable for assessing both real-time processing and short-term adaptation. We demonstrate that organoids can significantly improve their performance on this task through optimized training patterns, and show that this form of learn-

ing depends on glutamatergic neurotransmission, as it ceases when administering pharmacological blockers of AMPA and NMDA receptors.




Figure 5.1: Experimental overview for *in vitro* learning

a Schematic of multiphase experimental design. The *Record* phase uses a spontaneous recording to locate and characterize putative neural units. The *Stimulate* phase uses electrical stimulation on each of these units to measure stimulus-evoked activity through different temporal ranges. Human experimenters then select putative neuron roles from the causal connectivity analysis. The *Train* phase consists of repeated interactions with the simulated dynamical environment, organized into episodes. **b** Flowchart of each episode of the training loop. The virtual environment is simulated at a fixed frame rate while interacting with the organoid in a closed loop (yellow box). The episode is terminated when the pole falls into an unrecoverable position. Finally, depending on the training paradigm, a pattern of pulses may be delivered.

5.2 Results

5.2.1 Multi-Phase Experimentation

We developed a framework consisting of three key experimental phases to achieve goal-directed learning in cortical organoids: 1) network characterization through spontaneous recording, 2) stimulus-response mapping through targeted electrical stimulation, and 3) closed-loop training in the dynamical task (Figure 5.1a). Each phase builds upon automated analysis from the previous phase to systematically identify and interface with relevant neural circuits. This approach enables us to first characterize the network’s causal connectivity before attempting to dynamically modify the stimulus-evoked responses through training. The framework provides millisecond-precision control to minimize latency between the neural culture and virtual environment, and it also supports reproducibility through an automated analysis pipeline.

To implement this framework, we generated mouse cortical organoids to serve as a biological substrate for learning. Functional neural networks were developed through directed patterning and self-organization (Figure 5.2a) of three-dimensional aggregates from embryonic stem cells into structured neural tissue (Figure 5.2b) that recapitulates key features of cortical architecture [177]. By day 10, the networks show dorsal forebrain specified radial glia (Pax6) and medial gan-

glionic eminence progenitors (Nkx2.1) (Figure 5.2c). The tissue markers of proper differentiation at day 30 express cortical subtype-specific markers including upper (Satb2 and Brn2) (Fig. 5.8a) and deep (TLe4 and Ctip2) (Fig. 5.8b) layer excitatory projection neurons along with inhibitory neurons (Sst) (Fig. 5.8c), dendritic processes (Map2), and astrocytes (Gfap) (Fig. 5.8d). We specifically chose cortical patterning due to the cortex’s well-established role in adaptive information processing and its ability to encode, decode, and modify responses to novel inputs [1]. The organoids were interfaced with high-density microelectrode arrays (HD-MEAs) [178, 179, 58] (Figure 5.2d-f), providing precise spatio-temporal control over the culture with a high number of putative neuronal units for potential computation.




Figure 5.2: Organoid generation and interaction with electrical interface

a Schematic of the protocol for generating mouse embryonic stem cell (ESC) derived cortical organoids. **b** Brightfield images of organoids over development; images taken at varying magnifications. **c** Immunohistochemistry of day 10 cortical organoids show the presence of Pax6 (radial glia), Nkx2.1 (medial ganglionic eminence). Scale bar, 150 μ m. **d** Illustration of organoid plated on HD-MEA chip. Scale bar, 200 μ m. **e** Heatmap of activity from electrode configuration overlaid on the HD-MEA and a representative plot of the electrode configuration. Scale bar, 200 μ m. **f** Representative plot of the waveforms from individual neuronal units overlaid on electrode configuration

5.2.2 Neural Configuration

Our targeted approach to cortical organoid computation focuses on characterizing capabilities within small sub-circuits. To do this, we require a few separate components. First, we begin by capturing spontaneous neural activity in the *Record* phase, using it to identify locations of putative neurons and the best electrodes to stimulate them. As shown in a previous study, electrically stimulating the axon initial segment yields the best chance at triggering action potentials, with electrode locations showing higher recorded amplitudes providing better stimulation efficacy [180]. The spontaneous activity recording is used to generate a spatial map of putative neural unit locations (Methods: Spontaneous recording and neural unit identification) through a metric incorporating firing rate and action potential amplitude (Figure 5.2e,f). We then perform signal-averaging triggered by the peak amplitude of detected action potentials at these locations to extract spatio-temporal footprints and remove duplicates for each unit. Larger amplitudes yield neurons which are easier to identify in real-time experiments. This analysis enables us to reliably detect neural activity on distinct electrodes throughout the experiment and identify optimal stimulation sites based on peak amplitude locations.

To map causal information flow through these neural circuits, we characterized stimulus-response relationships during the *Stimulate* phase by delivering bi-phasic

pulses to each identified neural unit one-by-one for 50 repetitions (delivered at 2Hz) as seen in Figure 5.3a and c. We used automated analysis to quantify three major temporal response signatures: first-order causal connectivity representing direct neural pathways (Figure 5.3b), multi-order causal connectivity showing network-mediated responses (Figure 5.3d), and probability of evoking network-wide bursts (Fig. 5.9). Based on these metrics, we created a neural configuration defining units for encoding (input), decoding (output), and training. We select two encoding and two decoding units, prioritizing first-order connectivity from encoding to decoding units, and used multi-order connectivity as a secondary preference (Figure 5.3e and f). Units exhibiting frequent network bursting responses were deemed less suitable as encoding units since widespread activation could interfere with more fine-grained control. After encoding and decoding units were selected, between 5–12 training units were selected independent of connectivity patterns to explore optimal training stimulations.




Figure 5.3: Stimulation characterization and causal connectivity
a Stimulus-evoked responses: (top) 50 voltage traces from unit 7 after stimulating unit 6; (bottom) spike latency from stimulation raster. Red marks the first detected spike in each trial. **b** First-order causal connectivity showing the probability that stimulation of one unit (x-axis) evokes a response in another unit (y-axis) within 18ms. **c** As in a, showing a multi-order bursting response. **d** Multi-order causal connectivity showing mean response count within 200ms

when each unit is stimulated (x-axis) and responses are measured in other units (y-axis). **e** Network graph of first-order causal connectivity with assigned neural roles. Edge weights indicate probability of evoking a direct response within 18ms. Node positions are schematic. **f** As in e, but for multi-order connectivity. Edge weights indicate mean spike count within 200ms, capturing network-mediated responses that may include polysynaptic pathways through intermediate neurons.

5.2.3 Training Paradigm

Using our characterized neural configuration, in the *Training* phase we trained organoids to balance a pole on the cartpole task, where an inverted pendulum attached to a cart must be balanced by applying horizontal forces (between -10 N (Newtons) and 10 N) to the cart (Fig. 5.1). Each episode represents a single attempt at balancing the pole until failure (defined as the pole exceeding a terminal angle of $\pm 16^\circ$, which represents a nearly unrecoverable state), and typically lasts for tens of seconds depending on performance. The electrophysiological interaction between the environment simulation and the organoid was specified by a rate-coding scheme [181, 182, 149] for both encoding pole angle (stimulating the input units) and decoding horizontal forces (reading smoothed spike counts of the output units) to apply to the cart. This scheme is employed for each discrete timestep within an episode (Methods: Rate-coding scheme for encoding and decoding). Each episode ran from a random initialization until failure. A series of consecutive episodes is called a training cycle. We evaluate the performance within a training cycle as the amount of time balanced per episode. A result

indicating learning within a training cycle shows an increase in the duration of episodes. To allow for both learning and recovery, experiments were organized into typically 15-minute training cycles with 45-minute rest periods. The organoid receives real-time sensory information about pole angle and provides motor commands throughout each episode. However, training signals—the high-frequency stimulation patterns that modify neural dynamics—are delivered only at episode completion and only when a performance criteria is met. This design ensures that training signals cannot directly influence task performance through immediate artifacts or interference, but instead must modify the organoid’s neural dynamics and information processing during subsequent episodes. Specifically, training signals were administered whenever the 5-episode mean performance did not exceed the 20-episode mean, allowing the system to focus on larger-scale adaptations while filtering out the impact of individual episodes of poor performance. If the time balanced never surpassed 10 s within the first training cycle, different encoding/decoding units were selected. See Methods: Cartpole task environment for further details on training and Fig. 5.14 for configuration layouts.

Our approach leverages a fundamental property of neural networks: high-frequency stimulation can induce shifts in network dynamics that persist beyond the stimulation period [21, 151]. By delivering specific paired-pulse patterns to designated training neurons at episode completion, we aim to perturb the network

state in ways that modify the input-output transformation between encoding and decoding units. The adaptive selection of these perturbations—choosing which neural pairs to co-activate based on recent performance—allows us to explore different network configurations, effectively searching for control policies that improve task performance. Each stimulation pattern acts as a targeted perturbation that can shift the network’s dynamical landscape, and by tracking which perturbations lead to better balancing, we guide the system toward more effective configurations. Therefore, we call these network perturbations training patterns, and we optimize the training pattern selection using artificial reinforcement learning principles.

5.2.4 Superior Performance of Adaptive Training Signals over Random and Null Conditions

To investigate whether training signals sent to the training neurons could modify information processing between input and output neurons, we compared three distinct experimental conditions: 1) “*Null*” without stimulation, 2) “*Random*” with stimulation of 5 randomly chosen training units (neurons), and 3) “*Adaptive*” with stimulation of a pair of training units (neurons) selected through reinforcement learning with eligibility traces—a computational method that tracks which neuron pairs have previously led to performance improvements and preferentially

selects those pairs for future training, while gradually forgetting older associations (Figure 5.4a,b). We refer to either a 5-tuple or a pair of units selected for training stimulation feedback as a *training pattern*. We chose to stimulate random-ordered 5-tuples to replicate similar conditions from previous studies [9, 11]. Pairs instead of 5-tuples were used in the Adaptive Stimulation paradigm to both reduce the total number of possible training signals and to test whether a more fundamental pulse would enable modifications. Refer to Methods Section: Training signal paradigms for a detailed explanation of the training options tested. In each experiment, we tested two or more of these conditions in sequential 15-minute cycles (e.g., Null $\rightarrow$ Random $\rightarrow$ Adaptive) with one condition per cycle. Figure 5.4c showcases a representative experiment where the adaptive condition (green) repeatedly achieved superior performance, improving from a baseline of 10 seconds to over 60 seconds of balanced control across multiple cycles (see Supplementary Video 1-3 for examples of each condition). This improvement is quantified in the cycle-averaged performance metrics (Figure 5.4d). The emergence of effective control behavior becomes evident when examining the pole angle trajectories over time (Figure 5.4e). The effects of training stimulation delivered when the 5-episode mean performance dropped below the 20-episode mean is illustrated for three example cycles (Figure 5.4f).




Figure 5.4: Training paradigms enable task-specific learning

a Square wave biphasic pulse, 400 μV peak-to-peak, 400 μs period. Patterns consist of multi-channel pulses spaced 10ms apart, repeated at 10Hz. **b** Three training paradigms: Null (no stimulation), Random (5-pulse patterns with randomly ordered channels), and Adaptive (2-pulse patterns selected based on performance-attributed improvements). **c** Alternating paradigm experiment showing performance (balance time in seconds) by color (blue-null, red-random, green-adaptive), with raw (transparent) and smoothed (solid) traces. Cycles last 15 minutes with 45-minute rests; 21 cumulative hours shown. **d** Mean and inter-quartile range of performance per paradigm for the experiment in c. **e** Overlaid pole-angle

trajectories per episode within selected cycles. X's mark terminations; color indicates temporal progression (blue = early, pink = late). **f** Individual cycles with episodes as scatter points and training delivery times marked. Dark line is the 5-episode trailing mean; training was delivered when performance was not improving.

We evaluated 19 cortical organoids from 8 batches across 41 experiments (38 cartpole training, 3 pharmacological intervention), totaling over 125 hours of recorded activity. For cartpole training experiments, organoids underwent 1–6 experiments each (median: 2) spanning 3–75 cycles (median: 27). The most extensive experiment ran continuously over 48 hours. To quantify consistent performance, we used the 90th percentile of episode duration (time balanced) within each cycle as our primary metric. Organoids were deemed proficient if time balanced exceeded three standard deviations above the mean performance of the best of five *in silico* purely stochastic neural firing patterns (Fig. 5.11).

We compared distributions of the 90th percentile performance from each cycle of these training paradigms (Null: $n = 130$, Random: $n = 67$, Adaptive: $n = 91$). We found the adaptive training significantly outperformed both random and null cases ($p = 0.0081$ and $p = 1.87 \times 10^{-7}$ respectively, Dunn's test with Holm correction). Random stimulation did not significantly outperform the null case ($p = 0.079$). While 23.1% of cycles reached proficiency under adaptive training, only 4.5% did so with random stimulation, and 2.3% with no stimulation.

Given the superior performance of adaptive training, we conducted extended

experiments using only the adaptive paradigm without cycling between conditions. We will refer to this training strategy as “*Continuous Adaptive*” training. Figure 5.5a demonstrates sustained learning across multiple hours, with performance consistently exceeding the proficiency threshold. Notably, while organoids demonstrated clear performance improvements within individual training cycles, these improvements did not consistently persist across the 45-minute rest periods between cycles. The temporal pattern of performance peaks during training periods shows clear autocorrelation (Fig. 5.10), suggesting underlying state-dependent changes occurring over multi-hour periods. Post-hoc analysis revealed certain pulse combinations yielded consistently higher improvement metrics (Figure 5.5b), measured as the cumulative change in time balanced following each training signal delivery (see Fig. 5.4.3). These highly effective patterns—identified by improvement exceeding random walk bounds—often shared common training units.

To elucidate how the adaptive training paradigm optimizes its 2-unit stimulation patterns, Figure 5.5c visualizes the dynamic tracking of how the algorithm weights each possible training pattern during an experiment with multiple cycles. Unlike the random paradigm which always samples 5-unit patterns randomly, the adaptive method uses eligibility trace-based temporal difference reinforcement learning described in Methods: Training signal paradigms to assign higher weights to training patterns that have historically contributed to performance im-

provements. The inset highlights how individual pulse patterns can drive either improvement or deterioration at different timepoints, likely reflecting differing network states, emphasizing the importance of adaptive training signal selection.

We visualize a representative proficient organoid’s control strategy in two complementary ways. A simplified sigmoid policy estimation (Figure 5.5e) shows the emergence of structured control centered around the vertical pole position, while the complete neural response dynamics showing how the organoid’s motor output patterns relate to sensory input states (Figure 5.5f,g) and reveal richer dynamics. These flow fields—which capture how neural responses relate to the pole’s state through time—show adaptation towards an off-center balancing point that incorporates both angle and angular velocity information. This suggests the organoid’s control strategy becomes sensitive to the pole’s dynamic state rather than responding solely to instantaneous position (see Fig. 5.12 and Supplementary Video 4). Early episodes show scattered, inconsistent responses, but late episodes demonstrate coherent control strategies with multiple stable oscillation patterns and increased activity density near this preferred state. In late episodes, the organoid’s control policy becomes more structured and consistent, with episodes converging toward consistent initial starting conditions (Figure 5.5f,g). The input-output mappings show that neural responses develop stable trajectories that reflect systematic control strategies[183], with well-trained organoids repeatedly cy-

cling through these consistent response patterns during extended balancing periods. Adaptive training did not simply increase decode unit firing rates (Fig. 5.13), but reorganized them in a performance-dependent manner.




Figure 5.5: Continuous adaptive training reveals control policy development

a Performance across all adaptive training cycles. Dashed threshold at 20.5s denotes *proficient* performance (Fig. 5.11). **b** Improvement metric per training pattern; line and circle colors denote the first and second neurons, respectively. Grey envelope indicates Brownian motion baseline (see Fig. 5.4.3). Roman numerals mark peak improvements. **c** Real-time adaptive pattern selection. Top: performance over the experiment. Bottom: value estimates for all 2-unit patterns, with color intensity indicating selection probability. Black markers show delivery times; red boxes correspond to panel d. **d** Value estimate adaptation for two example patterns (purple, orange). Purple values rise after yielding improvements; orange values rise then fall following performance loss, illustrating the algorithm's sensitivity to outcomes. **e** Sigmoid fits of the organoid's control policy within one cycle. Early episodes lack structure; late episodes converge toward a sigmoid centered at 0° . **f,g** Phase-plane dynamics of spike count difference vs. input frequency across early (f) and late (g) episodes. Late episodes show more round trips, higher density near 0° , and less variability in initial responses.

These improvements in control strategy were reflected in overall performance metrics—continuous adaptive experiments achieved proficiency in 46.4% of cycles, substantially higher than the 23.1% rate observed in adaptive alternating experiments (Figure 5.6a,b). When comparing distributions of the 90th percentile performance from each cycle, continuous adaptive significantly outperformed all cases ($p = 7.07 \times 10^{-20}$, $p = 2.08 \times 10^{-8}$, and $p = 8.11 \times 10^{-3}$ for null, random, and adaptive respectively, Dunn’s test with Holm correction). We analyzed correlations between the best 90th percentile performance for each experiment and a range of connectivity metrics derived from both spontaneous and stimulus-evoked activity (see Methods: Stimulus-response characterization). Our first-order causal connectivity metric proved especially predictive of performance outcomes ($R^2 = 0.446$, $p = 6.0 \times 10^{-4}$, $n = 30$), substantially outperforming traditional functional connectivity as calculated by spike-time tiling coefficient within 20 ms [184] ($R^2 = 0.288$, $p = 0.004$, $n = 32$) (Figure 5.6c,d). This advantage was most pronounced in proficient experiments, where causal connectivity showed remarkably strong correlation with learning performance ($R^2 = 0.59$, $p < 0.01$, $n = 12$), whereas functional connectivity showed a weaker correlation ($R^2 = 0.20$, $p = 0.14$, $n = 12$). The strength of first-order causal connections, which guided our neural configuration selection, emerged as a key predictor of learning capability (Figure 5.6e). Output units’ ability to evoke multi-order

responses ($p < 0.01$) and network-wide bursts ($p < 0.01$) correlated with performance, suggesting that output units' capacity to recruit broader network activity may facilitate adaptive control. Other connectivity metrics, including burst probability and functional coupling between non-input/output units, showed weaker or non-significant correlations, further supporting the importance of first-order causal pathways in enabling successful learning.

These results demonstrate both the effectiveness of adaptive training and the importance of proper neural configuration selection. While high-frequency stimulation generally enhanced network performance over baseline, adaptive selection of training signals led to significantly better outcomes than random selection, with proficiency rates doubling in continuous experiments. The state-dependent nature of each cycle's 90th percentile performance suggests persistent network states conducive to adaptation, rather than independent and isolated initializations.




Figure 5.6: Statistical validation and neural correlates of learning

a Box plots (IQR) of 90th-percentile cycle performance per condition (Null: $n = 130$; Random: $n = 67$; Adaptive: $n = 91$; Continuous Adaptive: $n = 138$; 7–12 organoids each). Adaptive stimulation outperforms both baselines (Dunn's test, Holm correction: adaptive vs null $***P = 1.87 \times 10^{-7}$, vs random $**P = 8.11 \times 10^{-3}$; continuous adaptive vs null $***P = 7.07 \times 10^{-20}$; random vs null $P = 0.079$). Red line: *proficient* threshold. **b** Proficiency rates: Continuous

Adaptive 46.4%, Adaptive 23.1%, Random 4.5%, Null 2.3%. **c,d** Peak performance predicted by baseline functional connectivity (c; $R^2 = 0.288$, $P = 0.004$, $n = 32$) and first-order causal connectivity (d; $R^2 = 0.446$, $P = 6.0 \times 10^{-4}$, $n = 30$). **e** Feature correlations with peak performance ($n = 30$): direct causal connections ($***P = 2.0 \times 10^{-4}$), evoked response ($**P = 2.0 \times 10^{-3}$), output-unit burst rate ($**P = 8.0 \times 10^{-3}$). *i/o*: input/output units. $*P < 0.05$, $**P < 0.01$, $***P < 0.001$.

5.2.5 Glutamatergic Transmission is Essential for Learning

To investigate the underlying mechanisms of the observed learning capabilities, we performed pharmacological intervention experiments using glutamatergic receptor antagonists. Application of AMPA/kainate receptor antagonist NBQX (20 μM) and NMDA receptor antagonist APV (100 μM) [147, 29] significantly impaired the ability of organoids to balance the pole (Figure 5.7a). Mixed effects modeling across three organoids revealed that drug application reduced performance by 10.38 seconds (95% CI: -13.52 to -7.25, $p < 0.001$), representing a 64% decrease from baseline performance regardless of pre-drug performance (Figure 5.7b; Table S1). Following washout of the antagonists, performance substantially recovered, with only a non-significant residual deficit of 1.51 seconds (95% CI: -4.48 to 1.45, $p = 0.318$) compared to baseline. These results demonstrate that glutamatergic transmission through both AMPA/kainate and NMDA receptors is necessary for the goal-directed learning observed in our experiments, and

that this effect is reversible through washout.




Figure 5.7: NBQX and APV abolish learning capabilities.

a Three organoid experiments (top, middle, bottom panels) demonstrating the effects of glutamatergic receptor antagonists on task performance. Performance is measured as time balanced (seconds) over experimental time (hours). All organoids show proficient performance prior to drug application (Pre), complete loss of learning during NBQX (20 μ M) and NMDA receptor antagonist APV (100 μ M) treatment (orange shaded area), and recovery of baseline performance after washout (Post). Circles indicate the 90th percentile performance in each cycle. **b** 90th percentile performance normalized to pre-drug baseline. Each color represents a different organoid as shown in panel a (purple, green, gray). Vertical dashed lines demarcate experimental phases (Pre, During, Post) and are colored with their respective organoid color. During treatment, 90th percentile performance decreases to less than 30% of baseline performance, and recovers after washout over varying timescales across organoids. Markers denote the 90th percentile values for each cycle (n=38 cycles from 3 organoids from 1 batch (3 experiments): Pre n=3, During n=6, Post n=29; total 3193 episodes).

5.3 Discussion

Our results demonstrate goal-directed learning in brain organoids using intentionally selected circuits, establishing a new paradigm for biological computation beyond the successful unsupervised reservoir learning approach demonstrated in [10]. Organoids naturally recapitulate neural development and, to some extent, circuit architecture in three-dimensional space, providing a highly parallel substrate for spatial computation. HD-MEAs enable millisecond-precision closed-loop control, allowing us to map specific computational roles to individual neurons despite capturing only a subset of the total network activity. Using a novel multi-phase experimental approach, we designed a rapid analysis pipeline that enables iterative experimentation, characterizing neural responses and adjusting parameters within minutes between experimental phases. This capability for systematic investigation mirrors the rapid prototyping that accelerated artificial neural network development. Mouse cortical organoids, with their quick development cycle [19, 185], provide an ideal testbed for this engineering approach to biological computation. This framework opens new possibilities for both fundamental neuroscience and hybrid bio-electronic systems.

Unlike previous *in vitro* learning studies focused on pattern recognition or fixed behaviors [10, 9], our framework tackles a continuously unstable control problem requiring active state maintenance, demonstrating that biological circuits

can optimize complex dynamical systems even without canonical reward circuitry. Our experimental design—delivering training patterns only at episode completion and selecting training units distinct from input/output units—isolates the learning effects. Random stimulation did not significantly improve performance over baseline. However, our adaptive training paradigm, which optimized training signal selection based on performance changes, significantly outperformed both random and null conditions, achieving proficiency in 46.5% of continuous cycles. This demonstrates that biological neural networks can be systematically modified through precise electronic control.

Our adaptive approach represents a hybrid form of reinforcement learning that bridges artificial and biological systems. We use artificial reinforcement learning algorithms with eligibility traces [46] to optimize the selection of training signals delivered to biological neural networks. In this framework, the reinforcement learning algorithm learns which pairs of stimulation sites, when activated together, often improve task performance over time. Unlike canonical biological reinforcement learning, which relies on neuromodulatory signals like dopamine [186], our system delivers direct electrical stimulation to selected neurons based solely on task performance metrics. The organoid’s neural dynamics adapt in response to these externally-imposed training patterns, improving its input-output mappings for pole balancing. This adaptation likely occurs through stimulation-induced

shifts in network dynamics where each training pattern effectively explores a different configuration of the network’s dynamical landscape, with successful patterns reinforcing beneficial network states [9, 151]. Interestingly, extensive work has shown that cortical neurons, particularly in prefrontal and orbitofrontal cortex, can represent reward expectation and value signals [187, 188], suggesting that cortical circuits may possess intrinsic mechanisms for performance-based optimization. Our results demonstrate that such circuits can be shaped toward specific behavioral goals through systematic electrical stimulation, potentially mimicking endogenous processes by which cortical networks identify and strengthen task-relevant connections [189].

The effectiveness of training signals showed clear state-dependence, building on observations from [9] that different stimulation patterns succeed at different times. This suggests an interaction between training signals and recent activity history, providing an opportunity to explore fundamental stimulus-based learning rules at the circuit level. We observed rapid modulation of control policies through these brief training cycles while also discovering temporal correlation in the network’s ability to achieve high performance over multi-hour periods, revealing distinct timescales of network modulation. Such adaptation across multiple timescales demonstrates a key advantage of biological systems for hybrid bio-electronic applications: while our training signals enabled rapid tuning of task-specific responses,

the observed long-term state dependencies suggest potential for shaping underlying network dynamics to support and enhance these rapid adaptations. Training adaptations in different organoids lead to similar control strategies that effectively model inverted pendulum dynamics, even though the details of the neural connectivity in each organoid differ greatly. This convergence toward functionally similar input-output mappings demonstrates the organoids' ability to develop equivalent control policies through different underlying neural pathways.

Causal connectivity analysis reveals how interfacing with the network influences learning capability, providing a framework for predicting and optimizing performance. The strong predictive power of first-order causal connectivity ($R^2 = 0.59$ in proficient cycles) compared to traditional functional connectivity metrics demonstrates the importance of direct, stimulus-evoked pathways in architecting successful learning. The correlation between output units' ability to evoke network-wide responses and performance suggests that effective neural interfaces benefit from neurons capable of recruiting broader network activity. This methodology of causal characterization could inform both therapeutic interventions and the design of future biological computing systems. By isolating fundamental stimulus-based functionality in simplified circuits, we can systematically identify and explore learning rules that may scale to complex bio-electronic systems.

To test the biological basis of learning, we blocked glutamatergic transmission

with AMPA/kainate receptor antagonist NBQX and NMDA receptor antagonist APV. Application of these antagonists reduced performance by 64%, which recovered after washout, confirming that learning depends on intact glutamatergic transmission. This establishes a synaptic basis for the observed plasticity. Adaptive training modifies biologically-mediated neural dynamics rather than exploiting artifacts of the interface or task structure. Both AMPA and NMDA receptor-mediated neurotransmission are necessary for the adaptive responses. The rapid impairment and recovery suggest that our training paradigm induces changes in neural dynamics that require active glutamatergic signaling, though the precise cellular mechanisms underlying these stimulus-evoked changes remain to be characterized.

Our open-source Python-based platform **BrainDance** (<https://braingeneers.github.io/braindance>) enables exploration of biological learning mechanisms through flexible and accessible experimental design, following the rapid iteration paradigm that accelerated machine learning development. By isolating fundamental stimulus-based functionality in simplified circuits and providing tools for systematic investigation, this framework enables identification of learning rules that may scale from organoid models to therapeutic neural interfaces and hybrid bio-electronic computation systems.

5.4 Limitations

Our experimental approach has several technical constraints that affect interpretation and generalizability. Planar HD-MEA recordings capture only activity from the organoid surface in contact with the array, missing the three-dimensional network structure above. Single-electrode thresholding for spike detection likely includes multi-unit activity, which hinders understanding connectivity on a per-neuron level.

The observed performance improvements did not persist through the 45-minute rest periods between training cycle, requiring relearning in subsequent cycles. This suggests our training paradigm induces short-term plasticity operating on timescales distinct from canonical long-term potentiation. Whether this reflects fundamental constraints of the approach or could be overcome through modified training protocols remains unclear.

Learning capability proved configuration-dependent. Even within the same organoid, different selections of input/output neurons yielded varying success rates. Only 22% of adaptive training cycles (45% for continuous adaptive) achieved proficiency, indicating that successful learning under this paradigm requires specific network architecture features we do not yet fully understand. Our manual selection of neural configurations based on causal connectivity, while informed by data, lacks systematic optimization and likely misses better alternatives.

The rate-coding scheme and reinforcement learning parameters (eligibility trace decay, learning rate) were hand-tuned rather than systematically optimized. More sophisticated approaches—adaptive encoding schemes, latent space representations of organoid state, or advanced RL algorithms that incorporate task structure—could improve performance. Integration of local field potentials might provide richer and more stable network state information beyond spike counts.

Future investigations should examine how learning capability changes across organoid developmental stages and with different regional specifications, test systematic methods for neural configuration selection, explore how stimulation timing and pattern structure affect plasticity induction, and identify the cellular mechanisms underlying the glutamatergic-dependent adaptation we observe.

5.4.1 Experimental model and study participant details

5.4.1.1 Mouse embryonic stem cell line

All experiments were performed in the adapted BRUCE-4 mouse embryonic stem cell (ESC) line (Millipore Sigma #SF-CMTI-2). This line is derived from a male of the C57/BL6J mouse strain. Mycoplasma testing confirmed lack of contamination.

ESCs were maintained on Recombinant Human Protein Vitronectin (Thermo Fisher Scientific #A14700) coated plates using mESC maintenance media contain-

ing Glasgow Minimum Essential Medium (Thermo Fisher Scientific #11710035), Embryonic Stem Cell-Qualified Fetal Bovine Serum (Thermo Fisher Scientific #10439001), 0.1 mM MEM Non-Essential Amino Acids (Thermo Fisher Scientific #11140050), 1 mM Sodium Pyruvate (Millipore Sigma #S8636), 2 mM Glutamax supplement (Thermo Fisher Scientific #35050061), 0.1 mM 2-Mercaptoethanol (Millipore Sigma #M3148), and 0.05 mg/ml Primocin (Invitrogen #ant-pm-05). mESC maintenance media was supplemented with 1,000 units/mL of Recombinant Mouse Leukemia Inhibitory Factor (Millipore Sigma #ESG1107). Media was changed every day.

Vitronectin coating was incubated for 15 min at a concentration of 0.5 µg/mL dissolved in 1X PBS. Dissociation and cell passages were done using ReLeSR passaging reagent (Stem Cell Technologies #05872) according to the manufacturer's instructions. Cell freezing was done in mFreSR cryopreservation medium (Stem Cell Technologies #05855) according to the manufacturer's instructions.

5.4.1.2 Mouse cortical organoid generation

Mouse cortical organoids were grown as previously described by our group [190, 19] with some modifications. To generate cortical organoids we single cell dissociated ESCs using TrypLE Express Enzyme (ThermoFisher Scientific #12604021) for 5 minutes at 37°C and re-aggregated in lipidure-coated 96-well V-bottom plates at a density of 3,000 cells per aggregate, in 100 µL of mESC maintenance media

supplemented with Rho Kinase Inhibitor (Y-27632, 10 μ M, Tocris #1254) and 1,000 units/mL of Recombinant Mouse Leukemia Inhibitory Factor (Millipore Sigma #ESG1107) (Day -1).

After one day (Day 0), we replaced the medium with cortical differentiation medium containing Glasgow Minimum Essential Medium (Thermo Fisher Scientific # 11710035), 10% Knockout Serum Replacement (Thermo Fisher Scientific # 10828028), 0.1 mM MEM Non-Essential Amino Acids (Thermo Fisher Scientific # 11140050), 1 mM Sodium Pyruvate (Millipore Sigma # S8636), 2 mM Glutamax supplement (Thermo Fisher Scientific # 35050061), 0.1 mM 2-Mercaptoethanol (Millipore Sigma # M3148), and 0.05 mg/ml Primocin (Invitrogen # ant-pm-05). Cortical differentiation medium was supplemented with Rho Kinase Inhibitor (Y-27632, 20 μ M # 1254), WNT inhibitor (IWR1- ϵ , 3 μ M, Cayman Chemical # 13659), and TGF-Beta inhibitor (SB431542, Tocris # 1614, 5 μ M). Days 0-5 media were changed every day.

On day 5, organoids were transferred to ultra-low adhesion plates (Millipore Sigma # CLS3471) where media was aspirated and replaced with fresh neuronal differentiation media. The plate with organoids was put on an orbital shaker at 60 revolutions per minute. Neuronal differentiation medium contained Dulbecco's Modified Eagle Medium: Nutrient Mixture F-12 with GlutaMAX supplement (Thermo Fisher Scientific # 10565018), 1X N-2 Supplement (Thermo

Fisher Scientific # 17502048), 1X B-27 Supplement without vitamin a (Thermo Fisher Scientific # 12587010) 1X Chemically Defined Lipid Concentrate (Thermo Fisher Scientific # 11905031) and 0.05 mg/ml Primocin (Invitrogen # ant-pm-05). Organoids were grown under 5% CO₂ conditions. The medium was changed every 2-3 days.

On day 14 and onward, we transferred the organoids to neuronal maturation media containing BrainPhys Neuronal Medium (Stem Cell Technologies # 05790), 1X N-2 Supplement, 1X Chemically Defined Lipid Concentrate (Thermo Fisher Scientific # 11905031), 1X B-27 Supplement serum free (Thermo Fisher Scientific # 17504044), 0.05 mg/ml Primocin (Invitrogen # ant-pm-05) and 0.5% v/v Matrigel Growth Factor Reduced (GFR) Basement Membrane Matrix, LDEV-free.

5.4.1.3 HD-MEA plating and coating

Mouse cortical organoids were plated at day 25 on MaxOne high-density multielectrode arrays (Maxwell Biosystems) as previously described by our group [19, 180]. Prior to organoid plating, the multielectrode arrays were coated in two steps: First, we performed an overnight coating with 0.01% Poly-L-ornithine (Millipore Sigma #P4957) at 37°C. Plates were then washed 3 times with 1X Phosphate-buffered saline (PBS) pH 7.4 (Thermo Fisher Scientific #70011044). We then performed an overnight coating with 5 µg/mL mouse Laminin (Fisher Scientific #CB40232) and 5 µg/mL human Fibronectin (Fisher Scientific #CB40008)

at 37°C.

After coating, we placed the organoids on the chip and removed excess media. The organoids were then incubated at 37°C for 10 minutes to promote attachment. We then added pre-warmed neuronal maturation media and changed the media every 2-3 days.

5.4.1.4 Experimental design and sample

Experiments were performed on 19 mouse cortical organoids derived from 8 independent differentiation batches of BRUCE-4 embryonic stem cells. Each organoid was plated on a separate HD-MEA chip at day 25-30 of differentiation.

Biological replicates: Each organoid represents an independent biological replicate. Organoids from the same batch plated on different days are considered separate biological replicates due to variability in differentiation and plating outcomes.

Technical replicates: Multiple experiments on the same organoid represent technical replicates since we did not observe significant long-term modification. Multiple cycles within an experiment represent repeated measurements.

Sample sizes across all experiments:

- 19 organoids from 8 independent batches
- 16 organoids from 7 independent batches for cartpole training

- 432 total training cycles across cycled and continuous conditions
- Null: n=131 cycles from 12 organoids from 5 batches (29 experiments)
- Random: n=68 cycles from 9 organoids from 5 batches (20 experiments)
- Adaptive (cycled): n=92 cycles from 12 organoids from 5 batches (26 experiments)
- Continuous Adaptive: n=141 cycles from 7 organoids from 4 batches (10 experiments)
- NBQX + APV experiments: n=38 cycles from 3 organoids from 1 batch (3 experiments)
- >125 hours cumulative recording time

Three organoids were tested in both cycled and continuous adaptive paradigms in separate experiments.

5.4.2 Method details

5.4.2.1 High-density microelectrode array recording

Extracellular signals were obtained through the MaxWell MaxOne system (MaxOne, Maxwell Biosystems) [4], using custom experimental software for precise stimulation setup and timing [179]. Signals were recorded at a sampling

rate of 20 kHz/channel with a 1 Hz hardware filter, for up to 1024 channels. At most 32 stimulation electrodes could be selected at a time. Cultures were maintained within incubators full-time during recordings. Experiments were run with the BrainDance (<https://braingeneers.github.io/braindance>) Python library and longitudinally scheduled using an Internet of Things framework [191].

5.4.2.2 Spontaneous recording and neural unit identification

To rapidly determine the location and footprint of neurons, we performed spontaneous recordings capturing baseline neural activity. Automatic analysis located neurons based on quantity and magnitude of action potentials above threshold. Spike detection threshold was set at 5 times the root-mean-square amplitude of the channel signal in MaxLab Live software (MaxWell Biosystems). This conservative threshold ensured selection of the most relevant units, and was higher compared to later experimental phases.

Channels with activity were combined into putative neural unit footprints using spike-triggered averaging of electrical signals across all channels. Channels with peak-to-peak variation above an empirically chosen threshold in their spike-triggered average were considered part of the same footprint, forming an ad-hoc test for channels significantly likely to demonstrate simultaneous activity.

Redundant footprints were eliminated by identifying pairs sharing at least 50% of spike times and retaining the footprint with larger peak-to-peak signal

amplitude in its spike-triggered average.

From these footprints, up to 32 putative units were selected for stimulation based on a custom metric η designed to identify electrodes where spikes could be detected most consistently. This metric combined firing rate and mean spike amplitude μ_{amp} :

$$\eta = (1 + \hat{r}) \cdot (1 + 0.1|\mu_{\text{amp}}|) \quad (5.1)$$

where the normalized log spike rate $\hat{r}$ is defined with respect to minimum and maximum spike rates across all electrodes:

$$\hat{r} = \frac{\ln(1 + r) - \min(\ln(1 + r))}{\max(\ln(1 + r)) - \min(\ln(1 + r))} \quad (5.2)$$

5.4.2.3 Stimulation waveform parameters

All electrical stimulations used charge-balanced biphasic square wave pulses with 400 μV peak-to-peak amplitude and 200 μs per phase (400 μs total duration), positive phase first. For encoding sensory information, stimulation frequencies were dynamically adjusted according to the rate-coding equations described above. For training signals, patterns consisted of sequential biphasic pulses delivered with 10 ms inter-pulse intervals, repeated at 10 Hz for the duration of the training period (200-300 ms depending on the experiment).

5.4.2.4 Stimulus-response characterization

After identifying putative neural units, we characterize their stimulus-response properties by delivering biphasic pulses to each unit and recording the network’s response. Previous studies have identified multiple different response timescales [192], so we consider two separate connectivity metrics: the first-order connectivity C_1 , and the multi-order connectivity C_m .

For a window of stimulus-relative time, we define a response tensor indexed by the stimulus repetition k , the stimulus electrode i , the recording channel j :

$$\mathbf{R}_k^{i,j} \in \mathbb{N}^{n_{\text{reps}} \times n_{\text{stim}} \times n_{\text{channels}}} \quad (5.3)$$

where n_{reps} is the number of stimulus repetitions (50), n_{stim} is the number of stimulation electrodes corresponding to putative neurons, and n_{channels} is the number of recording channels specifically under putative neurons. The value of the tensor is the number of spikes in an interval of time relative to each stimulus, with spikes identified by thresholding at 3 standard deviations. The 3σ threshold is applied after SALPA artifact removal [28], which substantially reduces artifacts from electrical stimulation, and maintains real-time performance.

First-order connectivity C_1 is the observed probability of evoking at least one spike within 10ms:

$$C_1(i, j) = \frac{1}{n_{\text{reps}}} \sum_{k=1}^{n_{\text{reps}}} \mathbb{I}(\mathbf{R}_k^{i,j}{}_{[0,10\text{ms}]} > 0) \quad (5.4)$$

Multi-order connectivity C_m is the average number of spikes observed in a longer window:

$$C_m(i, j) = \frac{1}{n_{\text{reps}}} \sum_{k=1}^{n_{\text{reps}}} \mathbf{R}_k^{i,j}{}_{[10\text{ms}, 200\text{ms}]} \quad (5.5)$$

To identify network-wide bursts of the units, we compute the total spike count across all channels for each stimulation. A burst is detected when this count exceeds the median count plus three median absolute deviations (MAD). Bursts are excluded from the multi-order connectivity calculation to focus on specific neural pathways rather than larger network activations.

These metrics enable identification of both direct and network-mediated connections between neural units.

5.4.2.5 Neural configuration selection

Neural configurations were chosen following a three-part systematic approach. First, we identified encoding units with selective connectivity patterns, characterized by columns in the first-order connectivity matrix where only a small fraction of units showed high response probabilities. We prioritized encoding unit pairs (e.g., A→B and C→D) that maintained independence, avoiding configurations

where stimulation of one encoding unit would strongly activate the other encoding unit or the opposite decoding unit (i.e., avoiding strong $A \rightarrow C$ or $A \rightarrow D$ connections). When suitable first-order connections were sparse, we used the multi-order connectivity matrix as a secondary criterion. Finally, we excluded units that consistently evoked network-wide bursts (burst probability > 0.5), as widespread activation interfered with precise information transmission on the task timescale. Configurations were tested for one training cycle, and if the 90th percentile time balanced did not exceed 10 s, an alternative configuration was selected.

We did not enforce a minimum distance between stimulation electrodes. To avoid redundant detection of the same neuron from multiple electrodes, we employed a conservative footprint deduplication method during the Record phase. Putative spatio-temporal footprints were extracted for each active channel, and pairs sharing more than 50% of their constituent electrodes were identified as redundant. For each redundant pair, we retained only the footprint with the larger peak-to-peak signal amplitude in its spike-triggered average. Representative spatial distributions of neural configurations are shown in Supplementary Figure 7.

5.4.2.6 Cartpole task environment

Cartpole (also known as inverted pendulum) is a traditional control problem which tests an agent’s ability to balance an unstable dynamical system. The task is typically restricted to 2D, involving a cart that can move horizontally balancing

a pole vertically that is free to rotate around the attachment joint. The goal is to keep the pole upright by applying horizontal forces to the cart. This benchmark is valuable for its simple formulation yet inherent instability, real-world parallels, and widespread use in both control theory and reinforcement learning [175, 12].

The system state is fully described by four variables: cart position x , cart velocity $\dot{x}$, pole angle θ , and angular velocity $\dot{\theta}$. The agent can apply a force $F \in \{-10\text{N}, 10\text{N}\}$ to the cart at each timestep. The full equations of motion describing the nonlinear dynamics can be found in [193]. Episodes begin with small random perturbations to θ and $\dot{\theta}$ and proceed in discrete timesteps until $|\theta| > 16^\circ$, at which point the pole is considered to be in an unrecoverable state. We chose to remove traditional horizontal position bounds to focus solely on the pole balancing aspect of the task.

Each timestep involves decoding spike rates from the culture to determine the applied force, updating the virtual environment state with said force, and adjusting the encoding stimulation rates corresponding to the current virtual state.

5.4.2.7 Rate-coding scheme for encoding and decoding

Throughout this study, we use rate coding for both input/encoding and output/decoding signals, which provides sufficient information transmission, especially when focusing on small numbers of neurons [194, 2, 195]. The virtual state is encoded to the culture through two input neurons receiving stimulation fre-

quencies determined by the pole angle θ :

$$\begin{aligned} f_1 &= a \cdot (b - \sin(\theta))^2 \\ f_2 &= a \cdot (b + \sin(\theta))^2 \end{aligned} \tag{5.6}$$

where $a = 7$ and $b = 0.15$ are scaling factors that maintain stimulation frequencies in a biologically relevant range (both neurons receive ≈ 1.1 Hz when the pole is vertical, with frequencies diverging to ≈ 0.8 Hz and ≈ 8.9 Hz at terminal angles) while still being able to represent the sign of the angle. This functional form was chosen because the sinusoidal encoding ensures that the neurons receive opposing signals based on the pole's deviation from vertical, while the exponent enables the frequency to change faster as it nears the terminal angle.

For decoding signal from the culture, spikes are detected by a threshold of 3 standard deviations. Each output neuron's activity is then converted to a smoothed firing rate by an exponential moving average:

$$r_t = \alpha r_{t-1} + (1 - \alpha) c_t \tag{5.7}$$

where r_t is the smoothed firing rate at time t , c_t is the raw spike count in the current window, and $\alpha = 0.2$ balances immediate responses with firing rate stability.

The force F applied to the cart is determined by the difference between the

two output units’ smoothed firing rates, which are clipped to the range from -1 to 1 and then multiplied by 10 N .

5.4.2.8 Real-time closed-loop implementation

The closed-loop system operates through discrete timesteps involving three key components: the organoid’s neural activity, the simulated environment, and the training signals. Each timestep consists of:

1. Read Period (200 ms):
 - Process raw signals using SALPA (Subtraction of Artifacts by Local Polynomial Approximation) [28] for artifact removal
 - Detect spikes via thresholding at 3 standard deviations post-SALPA
 - Monitor motor neuron activity via smoothed firing rates r_t (Eq. 5.7)
2. Environment Update:
 - Decode motor signals $r_{1,t}$ and $r_{2,t}$ to output force F
 - Update cartpole state according to physics
 - Encode new state to stimulation frequencies (Eq. 5.6)
3. Training Period (300 ms, conditional):
 - Skip unless episode completed

- Deliver training patterns if 5-episode mean performance below 20-episode mean
- For adaptive protocol (see below), update value estimates based on performance

Each phase requires precise timing to maintain real-time control. A longer read phase yields reliable spike counts and smoother actions, whereas a shorter read phase leads to more responsive, yet noisy control. Experiments used either 200 ms or 300 ms read phases. For direct comparison, performance metrics from 200 ms experiments were scaled by a factor of 1.5 to normalize to a 300 ms time base.

5.4.2.9 Training signal paradigms

Training signals were delivered in a manner to induce a shift in how the network processed inputs. The training patterns were delivered only in response to two conditions being met: 1) termination of the task (pole drops below a terminal angle), and 2) the performance of the 5-episode mean is below the performance of the 20-episode mean. This second criterion results in the system focusing on larger-scale adaptations while ignoring singular episodes with poor performance. When these both are met, we enter the training period, where this is the only stimulation being delivered. Three experimental conditions were tested: a null

case without stimulation, and two stimulation paradigms, random and adaptive, with signals administered to the corresponding training neurons. The pool from which individual training neurons and corresponding electrodes ("training electrodes") were selected to receive feedback stimulation ($n = 8-15$) was created in the Record phase based on the activity metric described in Methods: Spontaneous recording and neural unit identification.

The random paradigm delivered sequential biphasic pulses (5 ms inter-pulse interval) through 5 training electrodes selected randomly without replacement from the pool of possible training electrodes. Each stimulation epoch consisted of one complete 5-pulse sequence delivered at 10 Hz for 200 ms.

In contrast to the random paradigm, the adaptive paradigm utilized a weighted sampling approach for paired-pulse patterns (5 ms inter-pulse interval). We construct a pool of all possible training patterns (e.g., if electrodes A, B, C, D, E are available, patterns include (A,B), (A,C), (A,D), ..., (D,E)), and sample based on a its historical contribution to performance improvements. Each training pattern i had a value estimate V_i updated at the end of each episode via:

$$V_{i,t+1} = V_{i,t} + \alpha(R_t - V_{i,t})E_{i,t} \quad (5.8)$$

where R_t denotes time balanced (our performance metric, episode duration),

α (0.3) is the learning rate, and $E_{i,t}$ represents the eligibility trace:

$$E_{i,t} = \gamma E_{i,t-1} + \begin{cases} 1 & \text{for selected training pattern } i \\ 0 & \text{for all other training patterns} \end{cases} \quad (5.9)$$

where γ (0.3) is the decay factor that diminishes the influence of past selections. The eligibility trace enables temporal credit assignment by storing a decaying record of recent stimulation patterns, weighted by their previous success in improving performance. The eligibility trace keeps track of which previous training patterns contributed to the current received performance.

Value estimates were not allowed to decrease between the minimum possible episode reward of 10. Selection probabilities were computed as:

$$P(i) = \frac{V_i}{\sum_j V_j} \quad (5.10)$$

Training patterns were sampled according to these probabilities and delivered as paired pulses at 10 Hz for 300 ms.

5.4.2.10 Pharmacological inhibition of glutamate receptors

Drug concentrations were selected based on established effective doses from previous studies [196, 197]. NBQX (Abcam #ab120045) was used at 20 μ M and APV at 100 μ M. Stock solutions were prepared to enable 1:1000 dilution (1 μ L per

1 mL medium), with NBQX dissolved in DMSO and APV in water. After drug administration, organoids were incubated for at least 30 minutes before undergoing drug-modulated training cycles.

Immunohistochemistry and confocal imaging

Organoids were collected and fixed in room temperature 4% Paraformaldehyde (PFA) (Thermo Fisher Scientific # 28908) and cryopreserved in 30% Sucrose (Millipore Sigma # S8501). They were then embedded in a solution containing 50% of Tissue-Tek O.C.T. Compound (Sakura # 4583) and 50% of 30% sucrose dissolved in 1X Phosphate-buffered saline (PBS) pH 7.4 (Thermo Fisher Scientific # 70011044). Organoids were then sectioned to 20 μ m using a cryostat (Leica Biosystems # CM3050) directly onto glass slides. After 2 washes of 5 minutes in 1X PBS and 1 wash in deionized water (Chem world #CW-DW-2G), the sections were incubated in blocking solution 5% v/v donkey serum (Millipore Sigma # D9663), and 0.1% Triton X-100 (Millipore Sigma # X100) for 1 hour. The sections were then incubated in primary antibodies overnight at 4°C. They were then washed 3 times for 10 minutes in PBS and incubated in secondary antibodies diluted in blocking solution for 90 minutes at room temperature. They were then washed 3 times for 10 minutes in PBS and coverslipped with Fluoromount-G Mounting Medium (Thermo Fisher Scientific # 00-4958-02).

Primary antibodies used were: mouse anti Gfap (Thermo Fisher Scientific

#G6171, RRID: AB_1840893, 1:100); rabbit anti Map2 (Proteintech #17490-1-AP, RRID: AB_2137880, 1:2000); mouse anti Sst (Santa Cruz Biotechnology #sc-55565, RRID: AB_831726, 1:100); mouse anti Satb2 (Abcam #ab51502, RRID: AB_882455, 1:100); mouse anti Pax6 (BD Biosciences #561462, RRID: AB_10715442, 1:100). rabbit anti Brn2 (Thermo Fisher #PA530124, RRID: AB_2547598, 1:100); rabbit anti Tle4 (Thermo Fisher #PA530192, RRID: AB_2547666, 1:100); rat anti Ctip2 (Abcam #ab18465, RRID: AB_2064130, 1:100), rabbit anti Nkx2.1 (Abcam #ab76013, RRID: AB_1310784, 1:400).

Secondary antibodies were of the Alexa series (Thermo Fisher Scientific), used at a concentration of 1:750. Nuclear counterstain was performed using 300 nM DAPI (4',6-Diamidino-2-Phenylindole, Dihydrochloride) (Thermo Fisher # D1306).

Imaging was done using an inverted confocal microscope (Zeiss 880) and Zen Blue software (Zeiss). Images were processed using Zen Black (Zeiss) and Fiji/ImageJ software (NIH) [198].

5.4.3 Quantification and statistical analysis

All statistical analyses were performed in Python using `scipy.stats` and `statsmodels` libraries. Prior to statistical testing, we assessed normality of the 90th percentile performance distributions using Shapiro-Wilk tests, which revealed signif-

icant deviations from normality across all conditions ($p < 0.001$). Sample sizes varied between experimental conditions and correlation analyses due to technical factors including intermittent signal loss during stimulation and occasional hardware issues during longitudinal experiments that prevented complete data collection.

Given these non-normal distributions, we employed the non-parametric Kruskal-Wallis H-test for overall differences between conditions ($H(3) = 91.69$, $p = 9.52 \times 10^{-20}$) and Dunn's post hoc test for subsequent pairwise comparisons between conditions (null, random, adaptive, and continuous adaptive), with Holm-Bonferroni correction for multiple comparisons. All statistical tests used an alpha level of 0.05, with p -values reported where appropriate. Organoids were considered proficient if their 90th percentile performance exceeded a threshold defined by random controller simulations (see Supplemental Information: Proficiency Threshold for detailed methodology). Sample sizes for each analysis are reported as: $n =$ [cycles] from [organoids] from [batches]. Statistical parameters (exact value of n , what n represents, SEM, SD, etc.) are reported in figure legends and results text. Statistical significance is denoted as $*P < 0.05$, $**P < 0.01$, $***P < 0.001$.

Supplemental Methods: Detailed methods for threshold determination, pharmacology, and data analysis

Proficiency Threshold

The performance threshold above which an organoid was considered “proficient” was based on significantly exceeding the performance (> 3 SD) of simulated random controllers. A cycle was considered “proficient” if its 90th percentile time balanced exceeded this threshold. To establish a performance baseline for the cartpole task, we simulated various random control strategies, each representing different approaches to generating control signals similar to potential biological noise patterns. Typical spike count from the experiments were used to inform the parameters chosen for each simulation. We implemented five distinct methods.

Brownian Motion Controller

This controller generates actions through a random walk process, where each new action is the previous action plus a normally distributed random step (step size = 0.3). This creates temporally correlated control signals that change smoothly over time, similar to gradual shifts in neural activity.

Independent Gaussian Controller

Actions are sampled independently from a Gaussian distribution ($\mu = 0$, $\sigma = 1$) at each timestep. This models uncorrelated random fluctuations around a neutral position, representing instantaneous neural noise without temporal dependency.

Uniform Random Controller

Actions are sampled uniformly from the interval $[-3, 3]$, representing an unbiased exploration of the action space with equal probability across all possible forces. This provides a baseline for completely random control without any structural assumptions.

Noisy Two-Neuron Network

This controller simulates two independent neurons with random firing rates drawn from a normal distribution ($\sigma = 3$). The difference between these rates determines the action, mimicking our experimental setup’s output decoding scheme. This represents instantaneous neural noise in a two-neuron circuit.

Brownian Two-Neuron Network

Similar to the noisy network, but each neuron’s firing rate follows a Brownian motion process (step size = 1). This combines temporal correlation with the

two-neuron architecture, representing gradual changes in neural firing rates.

Proficiency Threshold Selection

For each strategy, we conducted 50,000 episodes with a maximum duration of 1000 timesteps. All generated actions were clipped to $[-1, 1]$ action output or $[-10, 10]$ N to match the constraints of our biological experiments. The independent Gaussian controller achieved the highest mean performance (8.64 ± 4.00 seconds), likely due to its ability to rapidly change without being too similar with previous actions. We defined our proficiency threshold at three standard deviations above the mean (20.64 seconds) of the Gaussian controller’s performance distribution. This threshold provides a rigorous benchmark for distinguishing genuine learning from random success.

Pharmacological Intervention Experiments

To investigate the neural mechanisms underlying the observed learning capabilities, we performed pharmacological intervention experiments using glutamatergic receptor antagonists. We applied 2-amino-3-(5-methyl-3-oxo-1,2-oxazol-4-yl)propanoic acid (AMPA)/kainate receptor antagonist NBQX ($20 \mu M$) and N-methyl-D-aspartate (NMDA) receptor antagonist APV ($100 \mu M$) between training cycles.

Each experiment consisted of three parts: (1) a pre-drug baseline cycle where organoids demonstrated learning on the cartpole task, (2) drug application where NBQX and APV were added to the medium for 30 minutes, along with at least two attempted training cycles, and (3) a post-period where drugs were washed out and recovery was monitored over repeated cycles. We compared performance between these cycles and assessed recovery capabilities.

In all three organoids tested, application of NBQX and APV completely abolished the ability to balance the pole beyond baseline levels. Mixed effects modeling across 3193 episodes revealed that drug application significantly reduced performance by 10.38 seconds (95% CI = $[-13.52, -7.25]$, $p = 8.23 \times 10^{-11}$), representing a 64% decrease from baseline performance of 16.15 seconds. This effect was consistent despite varying pre-drug performance levels across organoids. After washout, performance substantially recovered, with only a non-significant residual deficit of 1.51 seconds (95% CI = $[-4.48, 1.45]$, $p = 0.318$) compared to baseline.

The complete elimination of learning capabilities during glutamatergic receptor blockade suggests that both AMPA/kainate and NMDA receptor-mediated synaptic transmission are essential for the adaptive responses observed in our experiments. This finding supports that glutamatergic synaptic transmission is necessary for goal-directed learning demonstrated by the organoids in the cartpole

task.

Autocorrelation Analysis

To investigate long temporal structure in organoid performance beyond individual episodes, we analyzed autocorrelations in performance (time balanced) over multi-hour periods. For each experiment, we calculated the normalized autocorrelation:

$$R(\tau) = \frac{\sum_{k=-\infty}^{\infty} x_k x_{k+\tau}}{R_0}$$

where x_k is the time balanced in episode k , and $R_0 = \sum x_k^2$ is the unnormalized autocorrelation at zero lag (the signal’s maximum correlation with itself). The correlation is computed using all possible overlaps, then normalized to set $R(0) = 1$. To establish statistical significance, we generated null distributions by temporally shuffling the time balanced values ($n = 10$ shuffles per experiment) while preserving their value distribution. This shuffling breaks any temporal dependencies while maintaining the overall performance statistics. Comparing the original autocorrelation to the shuffled distributions reveals temporal structure in learning beyond what would be expected by chance.

Burst Detection

To identify network-wide bursts during the stimulation phase, we computed spike count matrix C where $C_{i,j}$ represents the total spikes across all recording electrodes for stimulation electrode i and repetition j . A single burst threshold was established using the median absolute deviation across all counts:

$$\text{threshold} = \text{median}(C) + 3 \cdot \text{median}(|C - \text{median}(C)|)$$

where the median operations are applied to all elements in C . The burst probability for each stimulation electrode i was calculated as the fraction of repetitions where C_i exceeded this threshold. Bursts were excluded from multi-order connectivity calculations to isolate specific neural pathways from widespread network activations. In neural configuration selection, electrodes with burst probabilities exceeding 0.5 were generally excluded as input units to maintain precise information transmission.

Improvement Metric

To evaluate the effectiveness of different training signals, we defined an improvement metric that quantifies the cumulative change in performance following each signal delivery. For a given training pattern, its improvement value started at

0 and accumulated the direct performance changes (measured as time balanced) observed for each delivered pattern. To establish a statistical baseline for this metric, we modeled the null hypothesis as a Brownian motion process with parameters matching our experimental timescale. Specifically, we simulated 1000 random trajectories of values of a training pattern with deliveries occurring at intervals drawn from a normal distribution with mean equal to the number of patterns and $\sigma = 10$ episodes, along with step magnitudes drawn from a normal distribution ($\mu = 0$, $\sigma = 10$). This models the expected random fluctuations in performance that might occur if patterns were selected uniformly at random. We then identified training patterns that produced improvement trajectories exceeding five standard deviations above the mean of these simulated trajectories, indicating highly statistically significant positive effects on performance.

Firing Rate Analysis

To address whether performance improvements reflect increased network activity rather than reorganized control, we analyzed firing rates at output recording sites across conditions (Supplementary Figure 6). Baseline firing rate distributions were nearly identical (Null: 12.3 ± 11.7 Hz, Random: 12.0 ± 11.5 Hz, Adaptive: 12.8 ± 12.2 Hz; Cohen’s $d < 0.07$ for all pairwise comparisons), demonstrating that adaptive training does not cause generalized network excitation. However, when

examining firing rates at matched performance levels, we observed performance-dependent differences. At low performance (< 10 s), adaptive conditions showed reduced activity (median 16.4 Hz) compared to null and random (23.5–24.7 Hz), while at high performance (> 30 s), adaptive maintained elevated activity (9.7 Hz) compared to other conditions (6.1–6.4 Hz). Effect sizes increased systematically with performance (Cohen’s d : $0.04 \rightarrow 0.52$ across performance bins), indicating that adaptive training reorganizes motor output toward task-relevant patterns rather than simply increasing overall activity.

Policy Manifold Estimation

To visualize how the organoid’s control strategy evolves, we mapped the relationship between state variables (pole angle and angular velocity) and motor output (spike count difference between decoding neurons). We compared early (first quarter) and late (final quarter) episodes within a cycle. The manifolds were constructed using a grid interpolation method. For a grid of size 100×100 spanning the observed ranges of pole angles and angular velocities, we estimated the manifold by linearly interpolating each datapoint onto the grid. The resulting surface was spatially smoothed using a Gaussian filter ($\sigma = 1$) to reduce noise while preserving major features. Notably, while only pole angle was directly encoded to the organoid through stimulation frequencies, some of the successful

manifolds reveal structure along the angular velocity dimension, indicating that the organoid developed sensitivity to hidden state variables.

Table 5.1: Mixed Effects Model Summary of Drug Effects on Organoid Performance

Metric	Estimate	95% CI	p-value
Intercept (Pre)	16.15	(3.05, 29.26)	N/A
During Drug Effect	-10.38	(-13.52, -7.25)	$8.23 \times 10^{-11}^*$
Post Drug Effect	-1.51	(-4.48, 1.45)	0.3176




Figure 5.8: Additional immunohistochemistry validation of day 30 cortical organoids **a** Expression of upper layer cortical subtype markers showing DAPI (blue), Satb2 (green), Brn2 (red), and merged image. **b** Expression of deep layer cortical subtype markers showing DAPI (blue), Ctip2 (green), Tle4 (red), and merged image. **c** Expression of interneuron marker showing DAPI (blue), Sst (red), and merged image. **d** Neuronal and astrocytic differentiation marker showing DAPI (blue), Map2 (green), Gfap (red), and merged image. Scale bar for all images: 50 μm.




Figure 5.9: Burst-evoking stimulation characterization

a Matrix **C** showing total spike counts across all recording electrodes for each stimulation electrode (rows) and repetition (columns). Red X's indicate counts exceeding the global MAD-based threshold. **b** Representative raster plot from stimulation electrode 3 showing burst responses across multiple recording electrodes, with $> 50\%$ of stimulations evoking network-wide bursts. **c** Representative raster plot showing typical non-burst responses with more selective activation patterns.




Figure 5.10: Autocorrelation of performance

a i. Autocorrelation of a showcase experiment compared to a shuffled baseline. ii. Difference between autocorrelation and autocorrelation of the shuffled baseline, with colored highlights showing the polarity of the differences. iii. Performance analysis across episodes for the same experiment. **b** Extended autocorrelation analysis over larger time scales, with the red dashed line indicating an 8-hour boundary. The color gradient represents the best cycle each experiment obtained measured by 90th percentile within the cycle.




Figure 5.11: Random control strategy performance

a Performance distributions for five random control strategies implemented as baselines. Each violin plot shows the distribution of episode durations across 50,000 trials. **b** Mean performance $\pm 3\sigma$ for each control strategy. The independent Gaussian controller achieved the highest mean performance (8.64 ± 4.00 seconds), informing our proficiency threshold of 20.64 seconds.




Figure 5.12: Evolution of the control policy manifold a,b Early and late episode manifolds under random training showing spike count difference as a function of pole angle and angular velocity. c,d Under adaptive training, manifolds evolve from scattered early responses to a structured policy with clear

velocity-dependent control. **e,f** Null condition manifolds showing minimal policy refinement over time. Colors indicate spike count difference between left and right decoding neurons (range: -10 to 10 spikes). Positive differences (red) drive the cart rightward while negative (blue) drive leftward.




Figure 5.13: Adaptive training reorganizes motor output activity rather than globally increasing firing rates

a Violin plots of output firing rates across all episodes. Overall distributions are similar (Null: 12.3 ± 11.7 Hz; Random: 12.0 ± 11.5 Hz; Adaptive: 12.8 ± 12.2 Hz; Kruskal-Wallis $p < 0.001$), but effect sizes are negligible (Cohen's $d < 0.07$), indicating the significance is driven by large sample size. **b** Firing rates binned by episode duration reveal performance-dependent reorganization. At low performance (< 10 s) adaptive shows lower activity than null and random; at high performance (> 30 s) it maintains elevated activity. Effect sizes increase with performance (Cohen's d adaptive vs. null: 0.04, 0.07, 0.31, 0.52 for < 10 s, 10–20s, 20–30s, > 30 s). **c** Cumulative distributions confirm similar overall firing rate profiles across conditions. **d** Firing rate evolution over the first 10 minutes. All

conditions show an initial decline (Null: -29.3% ; Random: -28.2% ; Adaptive: -17.7%), with adaptive decaying most slowly. Multi-unit activity summed across both output sites; error bars are IQR.




Figure 5.14: Spatial distribution of neural configurations across experiments Representative neural configurations from 32 experiments showing

spatial locations of encode (red), decode (blue), and training (blue) overlaid on a representation of all electrode sites (grey) on the MEA array. Each panel represents one experiment's configuration.

Chapter 6

Real-Time Spike Sorting for Closed-Loop Electrophysiology

The content within this Chapter is adapted from a journal article published in *PLoS ONE* by Public Library of Science, for which I served as contributing author alongside co-authors Tjitse van der Molen, Max Lim, Julian Bartram, Zhuowei Cheng, Ash Robbins, David F. Parks, Linda R. Petzold, Andreas Hierlemann, David Haussler, Paul K. Hansma, Kenneth R. Tovar, and Kenneth S. Kosik. My contributions included assisting in the methodological design, testing early implementations of the axonal propagation algorithm, writing and editing.

In this chapter, we will discuss the development and evaluation of RT-Sort, a spike sorting algorithm that has the capability to sort spikes in real time through

short-latency axonal propagations. This work is part of an effort to build BrainDance, enabling closed-loop electrophysiology with rapid processing of neural data.

6.1 Introduction

We set out to develop RT-Sort (Real Time Sorting), a spike detection and sorting algorithm that overcomes the issues related to overlapping spikes, changes in waveform shapes, and the inability to perform spike sorting in real time. Instead of relying on waveform shape similarity, we based our sorting method on a more biologically plausible phenomenon: the fact that an action potential propagates along a neural axon and thus gets detected with very high fidelity by a sequence of electrodes, in the same order, with sub-millisecond but nonzero time delays [199]. Interestingly, all real-time spike sorting methods described above still apply some form of template matching, where the shape of the spike waveform recorded on one or multiple electrodes directly determines the unit cluster to which the spike is assigned. For the first time, we used action potential propagation detection instead of spike waveform template matching to identify individual neurons in our recordings. This idea has already been acknowledged and utilized in the classification of neuron types based on high-density recordings [200], but has not yet explicitly been applied to enhance spike sorting routines.

In this work, we present RT-Sort, an algorithm that can sort spikes in real time based on their action potential propagation trajectory. In doing so, we show that RT-Sort performs at the level of a broad series of conventionally used template-matching-based spike sorting algorithms [201, 202, 203, 204, 205, 206], while reporting latencies for detection and sorting of 7.5 ± 1.5 ms (mean $\pm$ STD) from the trough of the waveform on replayed recordings. Furthermore, RT-Sort can be applied to any system with sufficiently high electrode density ($\leq 50 \mu\text{m}$ pitch), as exemplified by our implementations on some of the most commonly used recording devices for both in vivo [207] and in vitro [208] experiments.

6.2 Results

6.2.1 Real-Time Spike Detection on Injected Spikes

A CNN with 4 layers (Fig. 6.1A) was trained to detect spikes in recording-device-specific noise. Separate models were trained for MaxWell MEA recordings performed on human brain organoids and Neuropixels 1.0 recordings performed on mice. The training was performed using averaged waveform shapes of electrodes selected along the whole recorded axon, the trajectory along which the action potential propagates (Fig. 6.1B). These averaged waveform shapes were then pasted into recording-device-specific noise fragments, which enabled artifi-

cial amplification of the training dataset (Fig. 6.1C) and the generation of training cases of overlapping waveform shapes with ground truth (Fig. 6.1D). The small differences between the training loss and testing loss indicated that our models are neither overfitting nor underfitting the data (Fig. 6.1E). See the corresponding sections in Methods for a more detailed description of the model training and evaluation.

When testing the trained models on the artificially generated data from the held-out organoid recordings, using the detection threshold yielding an optimal F1 score, we obtained a recall score of $89.1\% \pm 5.59\%$ and a precision score of $91.5\% \pm 3.23\%$ (mean $\pm$ STD, Fig. 6.1F, G). Similarly, the mouse Neuropixels recordings yielded a recall score of $84.7\% \pm 5.81\%$ and a precision score of $86.4\% \pm 7.90\%$ (mean $\pm$ STD). In both cases, the obtained recall was significantly larger than applying a 5 RMS threshold to the injected spikes (organoids: 5 RMS recall mean $\pm$ STD = $63.6\% \pm 7.36\%$, two-sided paired t -test $P = 1.25 \times 10^{-7}$, $n = 6$, Fig. 6.1F; mouse: 5 RMS recall mean $\pm$ STD = $65.1\% \pm 3.37\%$, two-sided paired t -test $P = 5.19 \times 10^{-5}$, $n = 6$). Meanwhile, the precision was not significantly different between the CNN detection model and the 5 RMS threshold (organoids: 5 RMS precision mean $\pm$ STD = $95.0\% \pm 3.36\%$, two-sided paired t -test $P = 0.13$, $n = 6$, Fig. 6.1F; mouse: mean $\pm$ STD = $79.7\% \pm 4.60\%$ (smaller than RT-Sort CNN model), two-sided paired t -test $P = 0.15$, $n = 6$). The significantly

higher recall scores were due to the fact that the model could detect spikes with amplitudes below 5 RMS, with the 5th percentile of the amplitude distribution over all detected waveforms at 4.08 RMS for the MEA model and 3.67 RMS for the Neuropixels model. Meanwhile, the models hardly missed spikes with amplitudes above 5 RMS (Fig. 6.1H).

Adjusting the detection threshold for the organoid model from its optimal value allowed for enhanced precision scores at the cost of a slightly lower recall score. This enabled the use of both a stringent detection score at the optimal F1 score, obtained with a detection threshold of 27.5%, as well as a loose detection threshold at 10% (Fig. 6.1I; 17.5% and 7.5% for the Neuropixels model). Even though we considered a detection to be a true positive if it was detected within $400\ \mu\text{s}$ of the actual waveform trough, the average absolute deviation from the correct trough time was only $13\ \mu\text{s}$ and $14\ \mu\text{s}$ for the organoid and mouse recordings respectively (Fig. 6.1J), reflecting the high temporal sensitivity of the model detections.




Figure 6.1: Action potential detection with high temporal accuracy using a CNN. **A**, Architecture of the 4-layer CNN (200-node input, 10 ms window; 120-node output, 2–8 ms). **B**, Averaged waveform footprint from Kilosort2; purple traces selected for training, highlighting the variety of waveform shapes along the axon. **C**, Training sample creation: noise fragment (i), waveform selection (ii), injection with ground truth trough (iii), and CNN prediction (iv). **D**, Same as C with overlapping waveforms. **E**, Training and validation loss across epochs; small gap indicates no over/underfitting. **F–G**, Recall and precision of the CNN vs. 5RMS threshold on held-out organoid recordings; CNN achieves significantly higher recall ($P = 1.25 \times 10^{-7}$) with no significant difference in precision ($P = 0.13$). **H**, Amplitude distribution of detected vs. missed spikes, showing sensitivity below 5 RMS. **I**, F1, precision, and recall as a function of detection threshold; loose (L) and stringent (S) thresholds marked. **J**, Distribution of absolute trough-detection deviations (mean $13 \mu s$, 0.26 frames).

6.2.2 Real-Time Spike Sorting on Ground Truth Datasets

Subsequently, we utilized the CNN outputs to perform spike sorting in real time using two different ground truth datasets. A detailed explanation of the spike sorting algorithm can be found in Methods: Propagation Sequence Detection and following sections. In summary, within a short pre-recording made right before the experiment, recurring sequential detections on multiple electrodes with short but nonzero temporal interelectrode intervals were detected, reflecting the propagation of an action potential. We refer to these detections as propagation signals. Detections were based on high-fidelity interelectrode intervals and compared to the highest-amplitude electrode in the sequence. Sequences that were highly similar in their propagation pattern were merged into single propagation signals, yielding the final set of detected propagation signal sequences. In the subsequent experiment, sequential detections were compared to the pre-detected propagation signal sequences and, if there was sufficient overlap, the sequential detection was counted as a spike for that propagation signal sequence.

The first ground truth dataset consisted of recordings made of rat primary neurons plated on MaxWell HD-MEAs while one of the neurons was simultaneously recorded by patch clamp (Methods: Patch-MEA Ground Truth Recording). Applying RT-Sort to this recording yielded the detection of 49 different units, including the neuron recorded using patch clamp (Fig. 6.2A–C).

Because our CNN model was trained to detect general waveform shapes in sample-specific noise, which are then sorted based on their action potential propagation patterns, RT-Sort is well capable of detecting all spikes from the same neuron despite changes in waveform shape (Fig. 6.2D, E). Furthermore, since our CNN was trained on multiple waveforms injected into the same detection window, RT-Sort is well suited to detect overlapping waveforms in both injected (Fig. 6.1D) and actual recordings (Fig. 6.2F). Performing the ground truth tests on $n = 4$ neurons from 2 different MEAs using RT-Sort in offline mode resulted in a precision of $97.5\% \pm 4.4\%$ and a recall of $90.9\% \pm 5.57\%$ (mean $\pm$ STD, Fig. 6.2G). For the array with a pre-recording, RT-Sort was also applied in online mode on the $n = 2$ neurons on this array, yielding a precision of $98.2\% \pm 1.78\%$ and a recall of $97.8\% \pm 1.60\%$ (mean $\pm$ STD, Fig. 6.2H).

We further tested RT-Sort on a different ground truth dataset consisting of a simulated mouse in vivo Neuropixels dataset used in [209]. This simulated dataset consisted of 250 biophysically detailed neurons that each exhibit independent Poisson firing of spikes with added Gaussian noise. To first test the ability of RT-Sort to assign spikes to the correct sequences, sequence metrics were initially obtained using the ground truth spike times from the modelled recording. Subsequently assigning spikes to these sequences yielded 142 units detected on the minimally required 4 electrodes to measure the propagating action potential. Over these 142

units, a precision of $98.8\% \pm 12.4\%$ (mean $\pm$ STD) was obtained, reflecting a low false positive detection rate on spike detection (Fig. 6.2I). Subsequently, RT-Sort was run in offline mode on the first 5 minutes of the simulated dataset to detect propagation sequences, yielding 120 detectable action potential propagation sequences. Over these 120 units, a precision of $97.3\% \pm 13.4\%$ (mean $\pm$ STD) was obtained, again reflecting a low false positive rate (Fig. 6.2J), in accordance with the patch-MEA results.

These results indicate that RT-Sort is capable of detecting a subset of units with low false positive rates while being unable to detect others. This is likely due to the relatively stringent detection constraints imposed on RT-Sort in order to reliably detect propagating action potentials. Unlike the patch-MEA recordings, RT-Sort did miss some true positive spikes in the simulated dataset. This resulted in an overlap score per unit of 0.65 ± 0.24 (mean $\pm$ STD, Fig. 6.2K). The missed true positive spikes could partially be due to the fact that the spike detection model was trained on real waveforms, while the biophysically detailed neurons did not always yield realistically shaped waveforms. RT-Sort is more sensitive to unrealistically shaped waveforms than template matching algorithms, which makes it less likely to produce false positive spike detections.




Figure 6.2: RT-Sort spike sorting performance on ground truth datasets. A–B, Averaged waveform footprint (A) and CNN detection scores (B) for a patch-clamped neuron on HD-MEA; root electrode marked with red star, detected loose electrodes in bold, trace color encodes propagation delay. C, Raw MEA trace, CNN scores, RT-Sort detections (yellow dots), and simultaneous patch recording (blue) for an example period. D–E, Single-spike footprints showing waveform shape variability across spikes. F, Footprint of a spike overlapping with another unit; RT-Sort correctly assigns it. G–H, Precision and recall on 4 patch-MEA neurons in offline mode (97.5%/90.9%, G) and 2 neurons in online mode (98.2%/97.8%, H). I–J, Per-unit precision on the simulated Neuropixels dataset using ground truth sequences (98.8% $\pm$ 12.4%, I) and detected sequences (97.3% $\pm$ 13.4%, J), with corresponding waveform amplitudes. K, Per-unit overlap scores for the simulated dataset (mean $\pm$ STD = 0.65 $\pm$ 0.24).

6.2.3 Detection Latencies

For RT-Sort to be used in real-time experiments, it is important to minimize the latency between the occurrence of the action potential and its detection. In our case, the latency between the waveform trough and its sorted detection can be ascribed to three different factors: biological constraints, detection speed, and sorting speed (Fig. 6.3A). The biological constraints consist of the duration of the remainder of the spike waveform and propagation occurring after the waveform trough. For spike detection, using a single NVIDIA RTX A5000 GPU allowed us to perform a forward pass through the MaxWell MEA spike detection model for 1020 electrodes simultaneously (the total number of available recording electrodes for MaxWell Biosystems HD-MEAs) in 1.12 ms, and a forward pass through the Neuropixels spike detection model for 384 electrodes simultaneously (the total number of available recording electrodes on a Neuropixels probe) in 0.92 ms. Adding additional electrodes led only to a linear increase in the required computation times (Fig. 6.3B). A first-order linear regression model fitted to the forward pass duration yielded:

$$\text{duration (ms)} = 0.155 + 0.001 \times N_{\text{elec}} \quad (R^2 = 0.997, P = 4.14 \times 10^{-15}) \quad (6.1)$$

for the MEA model and

$$\text{duration (ms)} = 0.155 + 0.002 \times N_{\text{elec}} \quad (R^2 = 0.999, P = 4.09 \times 10^{-19}) \quad (6.2)$$

for the Neuropixels model. The steeper slope for the Neuropixels model is due to its higher sampling rate of 30 kHz compared to 20 kHz.




Figure 6.3: Real-time spike sorting with latencies in the range of synaptic transmission. **A**, Schematic decomposing total latency into biological constraints (waveform duration + 0.5 ms propagation), detection speed, and sorting speed. **B**, Forward pass computation time vs. number of electrodes for MEA (purple) and Neuropixels (cyan) models; both scale linearly ($R^2 > 0.997$). **C–D**, Sorting duration distributions for MEA (C) and Neuropixels (D) recordings. **E–F**, Total detection+sorting latency from waveform trough for MEA (E, 7.24 ± 1.49 ms) and Neuropixels (F, 7.57 ± 1.57 ms). **G**, Offline sequence detection and spike sorting time per minute of recording for patch-MEA (purple) and simulated Neuropixels pre-recording (cyan).

6.2.4 RT-Sort Compared to Other Sorters

We next compared RT-Sort to 6 conventionally used sorting algorithms (SpyKING Circus, HDSort, HerdingSpikes2, Kilosort2, IronClust, and Tridesclous [201, 202, 203, 204, 205, 206]) that were all applied to the same mouse Neuropixels

recording [210]. RT-Sort detected a total of 168 units in the Neuropixels recording, which is fewer than any of the other sorters, in agreement with the results observed in the simulated dataset. The units consisted of detected propagations of $40\text{ }\mu\text{m}$ to $150\text{ }\mu\text{m}$ in length and up to $500\text{ }\mu\text{s}$ in duration.

We sought to quantify the quality of the detected units by comparing them to the other sorters. A previous comparison between the six other sorters on this dataset showed only very limited overlap in the detected units between different sorters [209] (see Methods: Overlap Score for how the overlap was assessed). We hypothesize that the low overlap between units could be due to false positive spike detections of different sorters, which could make the spike trains of different sorters so dissimilar that they are not recognised as coming from the same unit. This would mean that good performance of a certain sorter could potentially be obscured by bad performance of a different sorter (Fig. 6.4A). As a result, for comparing detections between sorters, we quantified how similar the spikes detected by either RT-Sort or a comparison sorter alone were compared to the spikes that both sorters agreed on.

The distribution of the interelectrode intervals between each loose electrode in the sequence and the root electrode was constructed for “matched” spikes, “RT-Sort only” spikes, and “other sorter only” spikes (Fig. 6.4B). For every spike in each of these three groups, the interelectrode interval difference between the spike

and the distribution mean of the matched spikes was used as a measure of how different the propagation detected on that loose electrode was compared to the spikes agreed upon by RT-Sort and the other sorter, with absolute differences clipped at a maximum of 7 frames (Fig. 6.4C). An interelectrode interval difference score per unit was then defined as the average absolute difference over all spikes on all loose electrodes. To account for variability intrinsic to the variable action potential propagation speed, the average absolute difference over all matched spikes was subtracted from the average absolute difference of both the “RT-Sort only” and “other sorter only” spikes to yield the final interelectrode interval difference score for both groups.

A paired one-sided *t*-test showed that the interelectrode interval difference score was significantly larger for each of the other sorters compared to the RT-Sort only spikes, taken over all units with an overlap score of ≥ 0.1 where both sorters detected at least 10 unique spikes (Fig. 6.4D, E; SpyKING Circus: $P = 3.50 \times 10^{-7}$, $n = 115$ units; Kilosort2: $P = 1.49 \times 10^{-8}$, $n = 119$ units; HerdingSpikes2: $P = 9.76 \times 10^{-5}$, $n = 78$ units; Tridesclous: $P = 5.87 \times 10^{-10}$, $n = 104$ units; HDSort: $P = 2.36 \times 10^{-13}$, $n = 108$ units; IronClust: $P = 2.43 \times 10^{-4}$, $n = 106$ units). Note that n is notably lower than the number of units detected by these sorters, due to RT-Sort’s lower unit detection recall arising from more stringent detection thresholds.

We applied the same procedure to detected spike amplitude differences instead of interelectrode interval differences. Here too, we found a significantly larger difference score for the “other sorter only” spikes compared to the “RT-Sort only” spikes, except for HDSort (Fig. 6.4F, G; SpyKING Circus: $P = 1.15 \times 10^{-8}$; Kilosort2: $P = 2.49 \times 10^{-8}$; HerdingSpikes2: $P = 5.83 \times 10^{-4}$; Tridesclous: $P = 5.55 \times 10^{-7}$; HDSort: $P = 0.0789$; IronClust: $P = 8.24 \times 10^{-6}$). However, HDSort had the largest increase in the interelectrode interval difference score relative to RT-Sort of all the other sorters.

Despite the absence of a ground truth for these recordings, these results indicate that the RT-Sort only detections were more similar to the matched detections than those made exclusively by the other sorter. This makes it more likely that the RT-Sort only detections reflect false negatives of the other sorter, while the other-sorter-only detections may reflect false positives of that sorter. We therefore conclude that RT-Sort detects a lower number of units, but that the units detected are of relatively high quality. Most importantly, these results were obtained by detecting spikes in online mode, where it took 7.62 ± 1.58 ms (mean $\pm$ STD) between the trough of the waveform and its sorted detection (Fig. 6.4H).




Figure 6.4: RT-Sort detects spikes with consistent propagations and amplitudes. **A**, Averaged waveform footprints for an example unit: spikes matched between RT-Sort and SpyKing Circus (left), RT-Sort only (middle), and SpyKing Circus only (right). **B–C**, Interelectrode interval distributions (**B**) and per-spike interval differences from the matched mean (**C**) for each spike group; SpyKing Circus-only spikes show a wider distribution indicating unit contamination. **D–E**, Interelectrode interval difference scores across all matching units for SpyKing Circus (**D**) and all six sorters (**E**); other-sorter-only spikes are significantly more different than RT-Sort-only spikes ($P < 0.05$, one-sided paired t -test). **F–G**, Same analysis for spike amplitude differences; all sorters significant except HD-Sort ($P = 0.079$). **H**, Total detection+sorting latency for real-time sorting on the Neuropixels dataset (7.62 ± 1.58 ms).

Altogether, our results indicate that RT-Sort performs at least at a comparable level to currently available spike sorting algorithms, albeit with a lower unit recall. However, unlike all other spike sorting algorithms, RT-Sort achieves

this performance in real time, with sorting occurring within milliseconds of spike occurrence while the experiment is still ongoing. Thus it is the first algorithm that allows for closed-loop feedback experiments with single-neuron resolution and mono-synaptic delay times.

6.3 Discussion

The technological advancements in hardware used for extracellular neural recordings have yielded dense multi-electrode recording devices for both in vivo [207, 211] and in vitro [208] recordings with interelectrode pitches smaller than $20\text{ }\mu\text{m}$. These devices allow for the recording of the activity of a single neuron by multiple adjacent electrodes. With the high temporal resolution of these recording devices—in the range of 20–30 thousand frames per second—consistent time delays can be observed among electrodes that detect the action potential of a single neuron, resulting from the time it takes for this action potential to propagate along the axon [33]. This action potential propagation detection has been used to enhance the classification of neuron types based on high-density recordings [200]. In this work, we presented RT-Sort (Real Time Sorting) that for the first time utilizes action potential propagation detection instead of spike waveform template matching as a spike sorting methodology to identify individual neurons in high-density extracellular field recordings, enabling millisecond latencies from action

potential occurrence to sorted detection.

To obtain optimal sorting performance, it is important to detect the propagating action potential on as many electrodes as possible, with high temporal precision. Therefore, we developed a convolutional neural network (CNN) that detects waveform shapes in noisy traces on single electrodes (Fig. 6.1A–D), with a significantly higher recall than 5 RMS-based spike detection (Fig. 6.1G). When training our model, we ensured that it is sensitive to the wide variety of waveform shapes that can occur while the action potential propagates along various parts of a neuron. As a result, even the in vivo recordings from Neuropixels probes yielded propagating action potentials over $150\text{ }\mu\text{m}$, taking up to $500\text{ }\mu\text{s}$ for their propagation.

Despite leaving space for improvements, we think that in its current form RT-Sort will already provide significant value to researchers conducting neural extracellular field recordings with dense multi-electrode recording devices, both in vivo and in vitro. RT-Sort enables, for the first time, an easily implementable spike sorting routine with time latencies for sorted spike detection that are comparable to single synaptic delay times. This important feature opens up interesting opportunities for closed-loop experiments, where a neural system receives external inputs based on preceding activity recorded from the same system. These types of experiments have been performed in various forms but always utilize multi-unit

activity as a read-out of the neural system [212, 213, 214, 215]. With RT-Sort, closed-loop experiments can now be conducted with real-time single-unit read-out accuracy, which enables studying neural circuit dynamics in response to external inputs with greatly enhanced granularity. Similarly, we anticipate that this single-unit real-time detection granularity for large populations of recorded neurons can greatly enhance the performance of brain-machine interfaces utilized in real time.

In summary, RT-Sort exemplifies a novel approach to spike sorting based on action potential propagation, which comprises high-fidelity sequential detection of a propagating action potential on multiple adjacent electrodes with sub-millisecond but nonzero time delays. This approach lies closer to the biology that underlies action potentials than template matching, the current norm in spike sorting algorithms. As a result, RT-Sort is more robust against waveform shape changes and overlapping waveforms than conventional methods, yielding a reduced false positive rate in detected units and spikes. Importantly, RT-Sort is able to detect and sort spikes within less than 10 ms after their occurrence, on the order of single synaptic delay times. This makes it a useful tool for closed-loop experiments and brain-machine interfaces, as well as various other implementations. We hope that the broader community will take interest in action-potential-propagation-based spike sorting and apply this new paradigm to continuously iterate on better-performing real-time spike detection algorithms.

6.4 Methods

6.4.1 Real-Time Spike Detection

RT-Sort uses action potential propagation detection to identify individual neurons in recordings with dense electrode configurations of $\leq 50\mu\text{m}$ interelectrode pitch. This requires sensitive detection of the action potential on multiple electrodes followed by a sorting procedure that assigns the action potentials to the correct neurons. All of this needs to happen on millisecond timescales to enable RT-Sort to be used in real time. For this, we developed a spike detector that can simultaneously detect spikes on over a thousand different electrodes within millisecond timescales, while being more sensitive than 5 RMS threshold crossings and able to detect overlapping waveforms.

A convolutional neural network (CNN) was trained to detect neural waveform shapes in a noisy signal. Due to variability in the type of noise produced by different recording devices, it is important to train a specific model per recording device type. In this work, separate models are trained for human brain organoids recorded with MaxWell Biosystems multi-electrode arrays [3] and for intact mouse brains recorded with Neuropixels probes [1]. The data used for the training and validation process come from [210, 157] respectively and can be found at <https://datadryad.org/stash/dataset/doi:10.25349/D9031Z> and <https://>

`buzsakilab.nyumc.org/datasets/SiegleJ/AllenInstitute_744912849/session_766640955/`. All training and validation is performed using an NVIDIA RTX A5000 GPU.

For more information on the architecture, algorithm, or other experiments and analysis conducted, please see the full paper [29].

Chapter 7

BrainDance: A Framework for Closed-Loop Electrophysiology

This chapter describes the BrainDance framework for closed-loop electrophysiology experimentation. The work fully encapsulates the culmination of the tooling and infrastructure developed throughout this thesis. The current version will remain up-to-date and live at <https://braingeneers.github.io/braindance>.

7.1 Introduction

Biological computing is a growing domain using electrode arrays and methods invoking activity in neurons such as electrical stimulation or optogenetics [216, 217, 218]. This has been explored from the fundamental level [219, 220, 221] to

basic learning on small-counted arrays through pivotal works by Potter [222, 223, 224] and Jimbo [225, 226], to more recent works involving high-density CMOS arrays [227, 226, 228, 229]. Although many labs have been advancing the field, many develop their own frameworks, use varying arrays as the technology progresses, and there appears to be a lack of proper reproducibility.

However, when looking at the similar field of reinforcement learning, we see that there was a catalyst of community efforts surrounding an extremely useful open-source interface which modularly enables the interaction of intelligent agents and environments—OpenAI’s Gym [12]. This sparked researchers to create, evaluate, and test various RL algorithms including DQN, SAC, and PPO on environments including Atari games, basic control problems, 3D simulations, and more. There have been extensions for networking [230], many more games, multiple agents [231], and so on. The nature of these interactive environments first and foremost enables a simple way to have an agent interact with highly complex environments. This type of framework permits easy sharing of code, verifiable results, consistent benchmarking, and general accessibility. The impact these environments have had on advancing research is undeniable.

Throughout this chapter, we introduce BrainDance, a modular interface for exploring, training, and evaluating biological *in vitro* neuronal experiments. We begin with a hardware interface module which interacts directly with the experi-

mental hardware, enabling recording and interaction with the biological cultures. The primary use case is through electrical stimulation and recording, however other forms of interaction are possible to automate and control (optogenetics, pharmacology, etc.). We provide three methods of interacting with the cultures. First is through pre-built phases which facilitate experiments with automated phase-specific analysis. Examples of phases include: a raw recording which yields neural footprints, repetitive stimulation yielding evoked responses, or gamified closed-loop tasks which yield performance metrics. The second interaction method is through direct interaction, where electrical reading occurs as fast as hardware allows and custom methods can be created for real-time processing. The third interaction is through a easy-to-use GUI which allows for piecing phases together in order to run experiments and analysis.

7.2 Architecture Overview

BrainDance is a modular, extensible framework for designing, running, and analyzing neural electrophysiology experiments. The architecture follows a hierarchical design with three levels of abstraction:

- At the lowest level, direct hardware interfaces for recording and stimulating with precise timing.

- The middle level consists of a phase system which has parameterizable common experimental procedures as pre-defined phases.
- The top level provides a GUI for easy construction of phases and analysis.

This approach enables all levels of accessibility from high-school students to expert researchers while maintaining reproducibility and standardization. Figure 7.1 illustrates the overall architecture. At the physical level (Figure 7.1a), analog neural signals are converted to digital signals by the recording hardware and streamed to the host computer, which in turn dispatches stimulation commands back to the device. The software stack (Figure 7.1b) is organized into loosely coupled optional modules including feature engineering, spike sorting, experiment management, real-time phases, analysis phases, and a hardware interface. They communicate through a shared data context. This modular design allows individual components to be swapped, extended, or reused across experiments without altering the rest of the pipeline.




Figure 7.1: Architecture of the BrainDance framework.

a Data flow between the biological culture, recording hardware, and host computer. The hardware converts raw neural signals to digital signals and streams them to the computer, which issues hardware commands and stimulation patterns back to the device. **b** Software architecture of BrainDance. Incoming data passes through a *Feature Engineering* module (artifact removal, filtering, LFP band extraction) and a *Spike Sorter* (RT-Sort or thresholding). The *Experimenter* manages session control, data saving, and timing. Pre-built *Real-time Phases* implement closed-loop tasks and stimulation logic, while *Analysis Phases* provide post-hoc plot generation and connectivity analysis. A *Hardware Interface* layer handles low-level command generation, waveform construction, and precise timing. All modules communicate through a shared data context and are coordinated by the experiment manager.

7.2.1 Electrical Stimulation

BrainDance provides a flexible interface for delivering electrical stimulation through the array's digital-to-analog converter (DAC). The fundamental stim-

ulation primitive is a charge-balanced, negative-first biphasic square wave pulse. Each pulse consists of a cathodic (negative) phase followed by an anodic (positive) phase of equal amplitude and duration, returning to a neutral baseline afterward. The biphasic waveform ensures charge balance across the electrode–electrolyte interface, preventing electrode degradation and tissue damage from net charge accumulation [216, 217, 232, 233].

The pulse is parameterized by two values: the amplitude A (in mV), which sets the peak voltage of each phase, and the phase duration τ (in μs), which sets the length of each half of the waveform. The amplitude is converted from millivolts to DAC units (least significant bits) via the hardware-specific scaling factor, and the phase duration is converted to sample counts at the corresponding sampling rate.

These pulses can be delivered individually or composed into sequences for more complex stimulation paradigms. The sequence interface supports three command types:

stim: Delivers a biphasic pulse to a specified set of electrodes with given amplitude and phase duration. Multiple electrodes can be stimulated in a single command by powering up each corresponding stimulation unit before emitting the pulse.

delay: Inserts a programmable delay (in sample counts) between successive com-

mands, enabling precise control of inter-pulse intervals.

next: Acts as a delimiter, marking the boundary of a stimulation pattern when sequences are triggered periodically or at externally determined timepoints.

By composing these commands, experimenters can construct arbitrary stimulation protocols starting from single-electrode probing pulses and paired-pulse paradigms to multi-electrode sequential patterns delivered at controlled frequencies. BrainDance also supports alternative waveforms including sinusoidal stimulation (parameterized by amplitude and frequency) and direct DC voltage control through the DAC, although the biphasic pulse remains the default for its well-characterized ability to evoke action potentials.

7.2.2 Phases

The phase system forms the backbone for building intricate experiments with BrainDance. Phases are self-contained, reusable modules which perform specific experimental features. They have a standard interface for initialization, execution, data collection, and cleanup.

Each phase inherits either (a) the **DataPhase** class which interfaces with experimental portions that occur on one of the supported devices and yields raw data or (b) the **AnalysisPhase** class which provides analysis that builds off of recorded data and yields derived data.

7.2.2.1 Pre-built Phases

RecordPhase: Allows for a timed recording of neural activity without any stimulation. Typically used for baseline measurements and initial characterization. *Provides:* `baseline_data`.

NeuralSweepPhase: Systematically stimulates each set stimulation electrode at a chosen frequency. A static amplitude can be used, or a range can be provided. The order of electrodes, amplitudes, and repetitions can be configured. *Requires:* `stim_electrodes`. *Provides:* `neuralsweep_data`, `stim_log`.

FrequencyStimPhase: Delivers stimulations at specified frequencies. The stimulations can either be single-pulse stimulations on any number of stimulation electrodes, or complex patterns as defined in Section 7.2.1. *Requires:* `stim_electrodes`. *Provides:* `stim_data`, `stim_log`.

ActivityPhase: Identifies a per-electrode activity metric based on detected spikes via thresholding. The default metric is $(1 + \text{firing_rate})(1 + \text{amplitude})$, or a custom metric can be specified. Calculates local maxima to identify activity hotspots. *Requires:* `baseline_data`. *Provides:* `selected_electrodes`.

FootprintPhase: Determines spatial footprints of neurons across electrodes from `selected_electrodes`. Removes duplicates which are part of the same footprint in favor of the largest amplitude, which indicates the location with

the best chance of evoking a spike [234, 232]. *Requires:* `baseline_data`, `selected_electrodes`. *Provides:* `stim_electrodes`.

CausalAnalysisPhase: Analyzes evoked responses by identifying all stimulation patterns in a file, creating windows post-stimulation, and removing the stimulation artifact (see Section 7.2.5). For spike detection, parameters allow using either thresholding or RT-Sort (Section 7.2.6.2). For an experiment with 10 stimulation patterns (or electrodes if using single-pulse, single-electrode), 20 reaction channels, 100 repetitions each, and 200 ms (4000 frames at a sampling rate of 20 frames/ms), we compute:

- `clean_data` — shape (stim_repetitions, stim_patterns, react_channels, timesteps) or (100, 10, 20, 4000): an `np.ndarray` of data with stimulation artifacts removed.
- `spikes` — shape (100, 10, 20): an `np.ndarray` of lists, where indexing yields a list of spike times for that repetition, stimulation pattern, and reaction channel.

Requires: `stim_data`, `stim_electrodes`, `react_channels`. *Provides:* `clean_data`, `spikes`.

RTSortPhase: Runs RT-Sort on baseline data, creates an RT-Sort object, and spike sorts the data file (see Section 7.2.6.2). This RT-Sort object can be

used later for spike sorting in real-time experiments if a sufficient GPU is available. *Requires:* `baseline_data`. *Provides:* `rt_sort_object`, `spike_data`.

7.2.3 Hardware Interface

BrainDance abstracts the interface between experimental logic and recording hardware through a two-layer environment design. At the base layer, a **BaseEnv** class defines the minimal contract that any hardware backend must satisfy: a `step()` method that returns observations and optionally delivers stimulation, a `reset()` method for reinitializing state between experimental runs, internal time management, and a maximum experiment duration after which the environment terminates. This base class tracks wall-clock time so that all higher-level code can query elapsed time and inter-step intervals without depending on hardware-specific timing. This base class is extensible to support other hardware platforms.

The primary concrete implementation is **MaxwellEnv**, which interfaces with the Maxwell CMOS microelectrode array. On initialization, it loads an electrode configuration file that specifies which of the array's electrodes are routed for recording and which are designated for stimulation (up to 32). The environment then initializes the hardware: setting amplifier gain, configuring the on-chip high-pass filter, connecting stimulation units to their respective electrodes, power-cycling the stimulation buffers, and computing voltage offsets. A ZeroMQ subscriber

connects to the MaxWell data server to receive streaming frames of voltage data and spike events at the 20 kHz sampling rate.

Each call to `step()` receives the latest data packet from the hardware, parses spike events from the binary stream, and optionally delivers a stimulation command. Stimulation commands are constructed as sequences of DAC operations as described in Section 7.2.1, supporting single biphasic pulses, multi-electrode pulse sequences, or manually constructed waveforms. All stimulations are timestamped and logged to a CSV file for post-hoc analysis. The environment supports two observation modes: (1) a `spikes` mode that returns parsed spike events (filtered to reject frames where artifact-like activity spans more than 15% of channels), and (2) a `raw` mode that returns the raw voltage trace across all routed channels, optionally buffered across multiple frames.

For experiments requiring low-latency stimulation in parallel with data acquisition, BrainDance supports a multiprocess architecture where data reception and stimulation delivery run in separate processes communicating through shared memory and event flags. This decoupling prevents stimulation timing from being blocked by data processing and vice versa. A dummy server mode is also provided, which can either replay recorded data or generate synthetic signals over the same ZeroMQ interface, acting as a test harness and enabling experimental development without needing the physical hardware.

7.2.4 Experiment Management

The experiment manager orchestrates multi-phase experimental protocols with automatic data flow between phases. An **Experiment** object is initialized with a name, a parameter dictionary (or path to a JSON file), and a save directory. Phases are added sequentially, and the manager validates the dependency chain before execution: each phase declares the data keys it *requires* and the keys it *provides*, and the validator ensures that every requirement is satisfied by a preceding phase or by pre-loaded data.

During execution, the manager iterates through phases in order. For each phase, it configures the hardware environment (creating or reusing a **BaseEnv** as needed), inputs the references to the shared data context, runs the phase, and merges the returned data into the shared context. Phase groups allow multiple phases to share a single hardware environment instance and save-file. All results (including timing, success, or errors) are recorded and saved as a JSON summary alongside the data files.

The data context acts similarly to a python dictionary that persists across phases and supports saving to and loading from disk. This enables experiments to be resumed from intermediate checkpoints: on initialization, the manager can automatically detect and reload previously saved data, allowing new phases to build on prior results without re-running earlier steps. Data from separate exper-

iments can also be loaded, enabling cross-experiment analysis pipelines.

Custom phases can be defined either by subclassing the phase base class or by using a decorator that enables easy in-line phase creation. This lightweight mechanism allows experimenters to inject custom analysis or selection logic between pre-built phases without modifying the framework. Listing 7.1 demonstrates a complete closed-loop experiment defined in this manner.

Listing 7.1: A closed-loop plasticity experiment constructed from pre-built and custom phases. The `@phase` decorator wraps a function into a phase object. Each phase’s outputs are automatically available to subsequent phases through the shared data context.

```
1 exp = Experiment("closed-loop", params=params)
2
3 @phase("SelectRandomPair")
4 def select_random_pair(exp):
5     C = exp.data.connectivity_matrix
6     for _ in range(500):
7         i, j = np.random.randint(0, C.shape[0], 2)
8         if 0.3 < C[i, j] < 0.6:
9             return {'selected_pair': (i, j)}
10    raise ValueError("No suitable pair found")
11
12 exp.add_phase(RecordPhase(duration=300))
13 exp.add_phase(RTSortPhase(sorter='rt_sort'))
14 exp.add_phase(ConnectivityPhase(window_ms=20.0))
15 exp.add_phase(select_random_pair)
16 exp.add_phase(ClosedLoopPhase(duration=300))
17 exp.add_phase(RecordPhase(duration=300))
18
19 exp.run()
```

7.2.5 Artifact Removal

Using electrical stimulations with CMOS arrays generates electrical artifacts that fall off spatially with distance from the electrode. We specifically show electrodes labeled “close,” “near,” and “far.” These artifacts pose a problem when

analyzing data after experiments and especially during experiments when timing is important for closed-loop paradigms. This can be seen in deep brain stimulation using template subtraction [235] or principal component projections post hoc [236]. In order to remove these artifacts, we implemented real-time artifact removal using the SALPA algorithm [237]. This algorithm fits a third-order Taylor polynomial and has a recursive formulation which is conducive to real-time processing. This algorithm enables us to read spikes within 2 ms from the electrode stimulation post-saturation. We include an implementation which translates Python to optimized machine code via Numba [238]. This could only handle a few hundred channels, so we extended it to a simpler method which iteratively estimates a linear fit over a sufficiently large time window, yielding similar results with significantly higher throughput.

7.2.6 Spike Detection

7.2.6.1 Thresholding

BrainDance implements flexible thresholding approaches for spike detection which do not require extra hardware and can be tailored to experimental needs. The framework supports both fixed voltage thresholds and adaptive thresholds calculated from baseline activity on a per-channel basis. For adaptive thresholding, users can specify multipliers of the standard deviation (σ) of the baseline

signal, with thresholds typically set between -3σ and -5σ for reliable spike detection. Importantly, threshold scaling can be applied either before or after artifact removal using the SALPA algorithm (Section 7.2.5), allowing experimenters to optimize detection sensitivity in stimulation experiments where artifacts may affect signal statistics.

7.2.6.2 RT-Sort

BrainDance natively supports RT-Sort [29], a novel spike sorting algorithm that enables real-time detection and sorting of action potentials based on their propagation patterns rather than template matching. RT-Sort requires a sufficient GPU, which is model-dependent. RT-Sort achieves sorting latencies of ~ 7.5 ms from spike detection to sorted assignment, making it ideal for closed-loop experiments. The algorithm uses a convolutional neural network for initial spike detection followed by tracking of consistent spatiotemporal propagation patterns across multiple electrodes.

Key features of RT-Sort in BrainDance include:

- Real-time processing of streaming data from multiple hardware platforms (MaxWell, Neuropixels)
- Built-in tools for training custom spike detection models optimized for specific experimental setups

- Integration with both online and offline analysis pipelines
- High precision in spike assignment ($> 97\%$) while maintaining millisecond-scale latencies
- Robust handling of overlapping spikes and waveform variations

RT-Sort complements BrainDance’s broader capabilities by enabling true closed-loop experiments with single-neuron resolution and synaptic-scale temporal precision. This allows for sophisticated experimental paradigms like neural feedback control and adaptive stimulation based on specific cellular activity patterns.

7.2.7 Real-Time Implementation

The difficulties of real-time dense CMOS electrophysiology lie in two separate parts. First, the hardware chosen has set limitations. For our primary use case, we work with the MaxWell MaxOne. Thus, we have to build around the current hardware and software bottlenecks, although others have created systems with shorter steps—the time between reading data in and enacting a stimulation. Other similar programs have had less than 20 ms lag time [224], or even lower when using FPGAs. However, for our experiments, less than 15 ms should be sufficient.

The second difficulty is spike detection and data processing. Since we are receiving up to 1024 channels of data sampled at 20 kHz, the storage, manipulation, and processing of data needs to be heavily optimized, as any computationally

heavy process compounds. As a primary goal of this software is ease of use, we keep the implementations in Python where possible. Our implementation uses multiple threads for layers of processing.

For managing the data, we first read the data off of the MaxWell server and fill a queue. Next, we pull the data from the queue and move it into a circular buffer. A second thread is then launched, implementing the SALPA algorithm (Section 7.2.5) to act as an artifact filter. A third thread thresholds the activity and maintains a spike rate for each channel.

7.2.8 Coding Framework

We build and release this software to enable similar capabilities across labs. This software provides a framework for conducting experiments of arbitrary complexity, where experimental algorithms and analyses can be shared between experimenters. We also aim to enhance the replicability of experiments by providing straightforward mechanisms for sharing experimental protocols. Scheduling of experiments provides a means to conduct longitudinal experiments as well.

Since the experiments are designed through code and not through expensive proprietary software, they can be created anywhere. This allows for remote experimentation which can provide accessible experiments, enabling uses in remote learning or at locations which do not have the funding or resources to run exper-

iments on their own.

7.3 Experimental Overview

We demonstrate the utility of BrainDance for running *in vitro* neural experiments through three experimental paradigms, showcasing the ability to craft complicated experimental designs at a high level. The experiments are as follows:

1. Longitudinal Stimulation Frequency Sweeps (“busy-bee”)

- Type: Open-loop. Phases: recording, causal stimulation.
- Demonstrates alternating between spontaneous recording phases and neural stimulation phases. The stimulation phases explore how neurons respond to differing stimulation frequencies, while the baseline recordings are used to interpret temporal consistency.

2. Cartpole Balancing Task (“cartpole”)

- Type: Closed-loop. Phases: recording, heatmap, causal, causal connectivity, cartpole.
- Uses phases to sequentially characterize putative neurons, firing rates, and stimulus responses in order to select units to interface with for the traditional control task of balancing an inverted pendulum on a moving cart.

3. Rapid Spike Response Pairing (“rapid”)

- Type: Closed-loop. Phases: recording, RT-Sort baseline, functional connectivity, rapid response.
- Employs RT-Sort to first identify spikes and create an RT-sorter, then identifies a unit pair with a desired functional connectivity. A pair of neurons A and B are identified. Whenever A fires, a stimulation is delivered to B . Finally, we observe how the functional connectivity changes in a concluding spontaneous recording.

7.3.1 Longitudinal Stimulation Frequency Sweeps

This experimental design can be seen in Figure 7.2a, where recording phases and stimulation phases alternate over the course of the experiment. The recording phases and stimulation phases have set durations, with stimulus rotating between 0.5, 1, 2, 4, and 8 Hz, repeating throughout the length of the experiment. The stimulation pulses were set to the default stimulation pulse. From a first recording phase, neurons were identified to stimulate. We prepared the stimulation phases to deliver the stimulations in the following order: neuron 1 repetition 1, neuron 1 repetition 2, ..., neuron N repetition 1, neuron N repetition 2, ..., neuron N repetition R . The working code for constructing this experiment is shown in Figure 7.2a, illustrating how BrainDance’s phase system enables concise specification

of complex longitudinal protocols.

We first examine the evoked responses at the single-trial level. Figure 7.2b shows a spike raster from a single recording during 0.5 Hz stimulation (electrode 11565, neuron 73), revealing short-latency responses concentrated within approximately 25 ms of stimulus onset. To assess whether these evoked responses remain stable across the longitudinal experiment, we condense all repetitions from each recording into a peri-stimulus time histogram (PSTH). Figure 7.2c displays this PSTH across all 13 recordings at 0.5 Hz for the same neuron, where each row corresponds to one recording session. Crucially, these 13 recordings are not consecutive—they are interspersed throughout the full experiment, separated by spontaneous recording phases and stimulation phases at other frequencies (1, 2, 4, and 8 Hz). Despite this interleaving, the evoked response timing remains consistent across recordings, while overall excitability varies, indicating that causal pathways are preserved even as network-wide activity fluctuates.

Beyond single-neuron spike responses, we observe that stimulation synchronizes network-level burst activity. Figure 7.2d shows burst events from an input electrode over 100 stimulation repetitions, demonstrating that bursts become entrained to the stimulus timing across trials. To examine whether this burst entrainment persists throughout the experiment, Figure 7.2e presents a burst activity density map at 1 Hz stimulation, where each row represents one recording

session. The consistent peak in burst density near stimulus onset across all sessions confirms temporal stability of the evoked burst response over the longitudinal timescale.

7.3.2 Cartpole Balancing Task

The inverted pendulum control task, denoted “cartpole,” serves as a closed-loop demonstration of BrainDance’s phase system. The experiment uses sequential phases—baseline recording, heatmap analysis, causal stimulation, neural configuration selection, and closed-loop training—to automatically characterize the network and select encoding, decoding, and training units without human input. Each phase’s outputs feed directly into subsequent phases through the shared data context, and all parameters (encoding/decoding method, training signal selection, time windows) are configurable, enabling easy and automated iterative experimentation. The full experimental design, results, and in-depth analysis of the cartpole task are presented in Chapter 5.

7.3.3 Rapid Spike Response Pairing

RT-Sort enables higher-quality experimentation due to its ability to provide real-time spike identities [29]. This is a substantial improvement over thresholded spike events, yet requires a sufficient GPU to run. First, we run a baseline record-

ing phase, then an RT-Sort baseline phase to instantiate our RT-sorter object. Next, an RT-Sort offline analysis phase is used on the baseline recording to identify spike trains, and a functional connectivity phase is used to determine the spike-time tiling coefficient.

A pair of neurons A and B are then selected based on a target connectivity range (spike-time tiling coefficient between 0.3 and 0.6), chosen to identify pairs with moderate but not saturated functional connectivity. During the closed-loop phase, whenever neuron A fires, a stimulation is delivered to neuron B as fast as possible given the hardware and RT-sort latency. This spike-triggered stimulation paradigm is designed to strengthen the functional connection from A to B through repeated temporally precise pairing similar to the fundamental principles of spike-timing-dependent plasticity [220, 221]. A concluding spontaneous recording phase is then used to assess whether the functional connectivity between the pair has changed as a result of the pairing. The full working code for this experiment is shown in Listing 7.1, demonstrating how a multi-phase closed-loop plasticity experiment can be specified in under 20 lines using BrainDance’s phase system and custom phase decorator.

7.4 Discussion

BrainDance addresses a structural problem in the field of biological neural learning: the absence of a shared experimental framework. By providing a modular architecture that can integrate with varying hardware interfaces, reusable experimental phases, and an experiment manager with automatic data management, the framework turns complicated experiments into a short script. The three experiments presented in this chapter illustrate this concretely. The longitudinal frequency sweep, cartpole task, and rapid spike response pairing each combine multiple phases in distinct configurations, yet all are built from similar base components and run through the same experiment manager.

The phase system's requires/provides contract enforces that data dependencies are satisfied before execution, catching configuration errors early rather than mid-experiment when the data is being accessed. This is particularly valuable for longitudinal or overnight experiments where failures ruin both experiments and potentially even the biology. The shared data context and checkpoint system further allow experiments to be paused and resumed, or for analysis phases to be re-run on previously collected data.

There are clear limitations. The hardware interface is currently implemented stably only for the MaxWell CMOS array, and extending to other platforms requires writing a new environment class. The GUI is functional but early in de-

velopment. RT-Sort requires a dedicated GPU, which limits its accessibility. The framework also inherits the constraints of the hardware: latency, channel count, and stimulation electrode limits are set by the array, not by the software.

Despite these limitations, BrainDance provides a foundation that can grow with the field. Experimental protocols built with the framework are shareable as code, enabling other labs to reproduce experiments or build on them directly. As the field moves toward larger-scale studies with more replicates and longer time horizons, this kind of infrastructure will be necessary to make progress cumulative.




Figure 7.2: Longitudinal stimulation frequency sweep experiment (busy-bee).

a Experimental workflow and working code for constructing the busy-bee experiment. After a baseline recording to locate neurons, the experiment alternates between spontaneous recording phases and stimulation phases at increasing frequencies (0.5, 1, 2, 4, and 8 Hz), repeating this cycle over the full duration. **b** Evoked spike raster from a single recording during 0.5 Hz stimulation (electrode 11565, neuron 73), showing short-latency responses concentrated within approximately 25 ms of stimulus onset. **c** Peri-stimulus time histogram (PSTH) across all 13 recordings at 0.5 Hz for the same neuron. Each row condenses all repetitions from one recording session into a single firing-rate estimate. These recordings are not consecutive but are interspersed throughout the experiment, separated by spontaneous and other-frequency stimulation phases. Evoked response timing remains consistent across recordings while overall excitability varies. **d** Burst raster from an input electrode (electrode 11565) over 100 stimulation repetitions, showing synchronization of burst events to the stimulation input across trials. Bin: 50 ms. **e** Burst activity density over time at 1 Hz stimulation, where each row represents one recording session. The persistent peak in burst density near stimulus onset demonstrates temporal consistency of the evoked burst response across the longitudinal experiment.

Chapter 8

Conclusion

This dissertation asked whether cortical organoids can be systematically trained to perform goal-directed tasks, and what infrastructure is needed to make that investigation reproducible and accessible. The work began with a hybrid soft-rigid robotic system where model-free reinforcement learning outperformed classical control on complex, nonlinear dynamics, establishing a software interface for mapping reinforcement learning to a "messy" system. Next, we progressed toward applying closed-loop learning paradigms to living neural tissue. Contributions to closed-loop optogenetic modulation of epileptiform activity in human brain slices and an internet-connected platform for organoid-based neuroscience education provided foundational software and early validation that high-density microelectrode arrays, combined with real-time control software, could accessibly

and reliably interface with biological tissue.

The central result is the first demonstration of goal-directed learning in brain organoids. We embodied cortical organoids in the inverted pendulum task “Cart-pole” through closed-loop electrophysiology and showed that an adaptive training paradigm, in which reinforcement learning optimizes the selection of electrical training signals, significantly outperforms both random stimulation and unstimulated controls. Pharmacological blockade of glutamatergic transmission abolished the learned improvements, and washout restored them, confirming that the adaptation depends on synaptic transmission rather than artifacts of electrical stimulation. Causal connectivity analysis revealed that the strength of direct stimulus-evoked pathways predicts learning performance, providing a basis for selecting electrode configurations before training.

For experiments requiring real-time spike sorting, we contributed to RT-Sort, which uses machine learning detection to achieve sorted spike assignment within approximately 7.5 ms, enabling closed-loop paradigms at single-neuron resolution. As a final contribution, we built BrainDance, an open-source Python framework to orchestrate reproducible biological learning experiments. We demonstrate the framework through three example experiments: (1) a longitudinal stimulation frequency sweep, (2) the cartpole balancing task, and (3) a rapid spike response pairing task. Together, these tools bring the rapid-iteration paradigm of OpenAI’s

Gym to biological neural learning. Any laboratory with compatible hardware (or the will to write a driver for their hardware) can design, run, and share reproducible experiments without rebuilding infrastructure from scratch. This work is available at <https://braingeneers.github.io/braindance>.

Chapter 9

Future Work

The work presented in this dissertation should be seen as a foundation, not a final answer. Translating these results into a deeper understanding of biological learning requires effort across several directions.

The most immediate priority is expanding BrainDance. The framework needs a comprehensive test harness and support for hardware modalities beyond electrical stimulation, including optogenetics, pharmacological delivery, and imaging. The setup to make these modalities available within the existing design is there, but the implementation needs to be completed. This will let experimenters combine perturbation types within a single protocol, enabling substantially richer experimental designs.

The encoding and decoding schemes used in this work, rate coding and simple

thresholding, represent a small region of the design space. Future work should move toward learned input-output mappings. The space of possible strategies is too large to search exhaustively through reinforcement learning alone, and pre-trained or foundational models that capture the structure of neural data could provide a scaffold for faster convergence on effective mappings for the chosen task.

The volume of data generated by high-density arrays and longitudinal experiments is well suited to large-scale machine learning. Models trained on neural dynamics could capture long-term adaptations, identify latent structure in network activity, and help design experiments by acting as a surrogate (or simulation) of the biological network.

The broader goal is to understand what kinds of computation different biological substrates are best equipped to perform. Systematic studies across organoid developmental stages, regional specifications, cell-type compositions, and activity profiles could reveal which tissue properties have the capacity to solve a given task. BrainDance makes this exploration feasible by making it easier to build a field-wide library of tasks. This will result in easier sharing of tissue properties, the methods for encoding/decoding information, and the methods for training. We can take insights from behavioral neuroscience and begin exploring the broad space of tasks that biological neural networks can learn. BrainDance is built to

make that accessible.

Bibliography

- [1] Jose M Carmena, Mikhail A Lebedev, Roy E Crist, Joseph E O’Doherty, David M Santucci, Dragan F Dimitrov, Parag G Patil, Craig S Henriquez, and Miguel A L Nicolelis. Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS biology*, 1(2):e42, 2003.
- [2] J. K. Chapin, K. A. Moxon, R. S. Markowitz, and M. A. Nicolelis. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2(7):664–670, July 1999.
- [3] Andy Zhou, Benjamin C Johnson, and Rikky Muller. Toward true closed-loop neuromodulation: artifact-free recording during stimulation. *Current Opinion in Neurobiology*, 50:119–127, June 2018.
- [4] Jan Müller, Marco Ballini, Paolo Livi, Yihui Chen, Milos Radivojevic, Amir Shadmani, Vijay Viswam, Ian L Jones, Michele Fiscella, Roland Diggelmann, et al. High-resolution cmos mea platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13):2767–2780, 2015.
- [5] Cleber A. Trujillo, Richard Gao, Priscilla D. Negraes, Jing Gu, Justin Buchanan, Sebastian Preissl, Allen Wang, Wei Wu, Gabriel G. Haddad, Isaac A. Chaim, Alain Domissy, Matthieu Vandenberghe, Anna Devor, Gene W. Yeo, Bradley Voytek, and Alysson R. Muotri. Complex oscillatory waves emerging from cortical organoids model early human brain network development. *Cell stem cell*, 25(4):558–569.e7, October 2019.
- [6] Silvia Velasco, Amanda J Kedaigle, Sean K Simmons, Allison Nash, Marina Rocha, Giorgia Quadrato, Bruna Paulsen, Lan Nguyen, Xian Adiconis, Aviv Regev, et al. Individual brain organoids reproducibly form cell diversity of the human cerebral cortex. *Nature*, 570(7762):523–527, 2019.
- [7] G. Shahaf and S. Marom. Learning in networks of cortical neurons. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 21(22):8782–8788, November 2001.

- [8] Thomas B. DeMarse, Daniel A. Wagenaar, Axel W. Blau, and Steve M. Potter. The Neurally Controlled Animat: Biological Brains Acting with Simulated Bodies. *Autonomous Robots*, 11(3):305–310, November 2001.
- [9] Douglas J. Bakkum, Zenas C. Chao, and Steve M. Potter. Spatio-temporal electrical stimuli shape behavior of an embodied cortical network in a goal-directed learning task. *Journal of Neural Engineering*, 5(3):310–323, August 2008.
- [10] Hongwei Cai, Zheng Ao, Chunhui Tian, Zhuohao Wu, Hongcheng Liu, Jason Tchieu, Mingxia Gu, Ken Mackie, and Feng Guo. Brain organoid reservoir computing for artificial intelligence. *Nature Electronics*, 6(12):1032–1039, December 2023.
- [11] Brett J. Kagan, Andy C. Kitchen, Nhi T. Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J. Parker, Anjali Bhat, Ben Rollo, Adeel Razi, and Karl J. Friston. *In vitro* neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969.e8, December 2022.
- [12] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- [13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, and others. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015. Publisher: Nature Publishing Group.
- [14] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [15] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and others. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- [16] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016. Publisher: Nature Publishing Group.

- [17] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, November 2019. Publisher: Nature Publishing Group.
- [18] Luciano Floridi, Josh Cowls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, Burkhard Schafer, Peggy Valcke, and Effy Vayena. AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4):689–707, December 2018.
- [19] Matthew AT Elliott, Hunter E Schweiger, Ash Robbins, Samira Vera-Choqueccota, Drew Ehrlich, Sebastian Hernandez, Kateryna Voitiuk, Jinghui Geng, Jess L Sevetson, Cordero Core, et al. Internet-connected cortical organoids for project-based stem cell and neuroscience education. *eneuro*, 10(12):ENEURO.0308–23.2023, 2023.
- [20] Eisaku Maeda, Yoichiro Kuroda, Hugh P. C. Robinson, and Akio Kawana. Modification of parallel activity elicited by propagating bursts in developing networks of rat cortical neurones. *European Journal of Neuroscience*, 10(2):488–496, 1998.
- [21] Y. Jimbo, T. Tateno, and H. P. C. Robinson. Simultaneous Induction of Pathway-Specific Potentiation and Depression in Networks of Cortical Neurons. *Biophysical Journal*, 76(2):670–678, February 1999.
- [22] Bernard D. Reger, Karen M. Fleming, Vittorio Sanguineti, Simon Alford, and Ferdinando A. Mussa-Ivaldi. Connecting Brains to Robots: An Artificial Body for Studying the Computational Properties of Neural Tissues. *Artificial Life*, 6(4):307–324, October 2000.
- [23] Maria Elisabetta Ruaro, Paolo Bonifazi, and Vincent Torre. Toward the neurocomputer: image processing and pattern recognition with neuronal cul-

- tures. *IEEE transactions on bio-medical engineering*, 52(3):371–383, March 2005.
- [24] T.B. DeMarse and K.P. Dockendorf. Adaptive flight control with living neuronal networks on microelectrode arrays. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 3, pages 1548–1551 vol. 3, July 2005. ISSN: 2161-4407.
 - [25] Kevin Warwick, Dimitris Xydias, Slawomir J. Nasuto, Victor M. Becerra, Mark W. Hammond, Julia Downes, Simon Marshall, and Benjamin J. Whalley. Controlling a mobile robot with a biological brain. *Defence Science Journal*, 60(1):5–14, January 2010.
 - [26] Takuya Isomura, Kiyoshi Kotani, Yasuhiko Jimbo, and Karl J. Friston. Experimental validation of the free-energy principle with in vitro neural networks. *Nature Communications*, 14(1):4547, August 2023.
 - [27] Takuya Isomura, Kiyoshi Kotani, and Yasuhiko Jimbo. Cultured Cortical Neurons Can Perform Blind Source Separation According to the Free-Energy Principle. *PLoS computational biology*, 11(12):e1004643, December 2015.
 - [28] Daniel A Wagenaar and Steve M Potter. Real-time multi-channel stimulus artifact suppression by local curve fitting. *Journal of Neuroscience Methods*, 120(2):113–120, October 2002.
 - [29] Tjitse van der Molen, Max Lim, Julian Bartram, Zhuowei Cheng, Ash Robbins, David F Parks, Linda R Petzold, Andreas Hierlemann, David Haussler, Paul K Hansma, et al. RT-Sort: an action potential propagation-based algorithm for real time spike detection and sorting with millisecond latencies. *PloS one*, 19(12):e0312438, 2024.
 - [30] A. S. Robbins, M. Ho, and M. Teodorescu. Model-free dynamic control of robotic joints with integrated elastic ligaments. *Robotics and Autonomous Systems*, 155:104150, September 2022.
 - [31] John P. Andrews, Jinghui Geng, Kateryna Voitiuk, Matthew A. T. Elliott, David Shin, Ash Robbins, Alex Spaeth, Albert Wang, Lin Li, Daniel Solis, Matthew G. Keefe, Jessica L. Sevetson, Julio A. Rivera de Jesús, Kevin C. Donohue, H. Hanh Larson, Drew Ehrlich, Kurtis I. Augustine, Sofie Salama, Vikaas Sohal, Tal Sharf, David Haussler, Cathryn R. Cadwell, David V. Schaffer, Edward F. Chang, Mircea Teodorescu, and Tomasz Jan Nowakowski. Multimodal evaluation of network activity and optogenetic interventions in human hippocampal slices. *Nature Neuroscience*, 27(12):2487–2499, December 2024. Publisher: Nature Publishing Group.

- [32] Matthew A. T. Elliott, John P. Andrews, Tjitse van der Molen, Jinghui Geng, Alex Spaeth, Anna Toledo, Kateryna Voitiuk, Cordero Core, Thomas Gillespie, Ari Sinervo, David F. Parks, Ash Robbins, Daniel Solís, Edward F. Chang, Tomasz Jan Nowakowski, Mircea Teodorescu, David Haussler, and Tal Sharf. Microscale maps of bursting dynamics across human hippocampal slices from patients with epilepsy. *Journal of Neurophysiology*, 135(1):273–285, January 2026.
- [33] Matthew A. T. Elliott, Hunter E. Schweiger, Ash Robbins, Samira Vera-Choqueccota, Drew Ehrlich, Sebastian Hernandez, Kateryna Voitiuk, Jinghui Geng, Jess L. Sevetson, Cordero Core, Yohei M. Rosen, Mircea Teodorescu, Nico O. Wagner, David Haussler, and Mohammed A. Mostajo-Radji. Internet-Connected Cortical Organoids for Project-Based Stem Cell and Neuroscience Education. *eNeuro*, 10(12), December 2023. Publisher: Society for Neuroscience Section: Research Article: New Research.
- [34] Ash Robbins, Hunter E. Schweiger, Sebastian Hernandez, Alex Spaeth, Kateryna Voitiuk, David F. Parks, Tjitse van der Molen, Jinghui Geng, Isabel Cline, Kenneth S. Kosik, Sofie R. Salama, Tal Sharf, Mohammed A. Mostajo-Radji, David Haussler, and Mircea Teodorescu. Goal-directed learning in cortical organoids. *Cell Reports*, 0(0), February 2026. Publisher: Elsevier.
- [35] Sebastian Hernandez, Hunter E. Schweiger, Isabel Cline, Gregory A. Kaurala, Ash Robbins, Daniel Solis, Jinghui Geng, Tjitse van der Molen, Francisco Reyes, Chinweike Norman Asogwa, Kateryna Voitiuk, Mattia Chini, Marco Rolandi, Sofie R. Salama, Bradley M. Colquitt, Tal Sharf, David Haussler, Mircea Teodorescu, and Mohammed A. Mostajo-Radji. Self-Organizing Neural Networks in Organoids Reveal Principles of Forebrain Circuit Assembly. *bioRxiv*, page 2025.05.01.651773, May 2025.
- [36] Mohammed A. Mostajo-Radji, Walter R. Mancia Leon, Arnar Breevoort, Jesus Gonzalez-Ferrer, Hunter E. Schweiger, Julian Lehrer, Li Zhou, Matthew T. Schmitz, Yonatan Perez, and Tanzila Mukhtar. Fate plasticity of interneuron specification. *Isience*, 28(4), 2025. Publisher: Elsevier.
- [37] David F. Parks, Kateryna Voitiuk, Jinghui Geng, Matthew A. T. Elliott, Matthew G. Keefe, Erik A. Jung, Ash Robbins, Pierre V. Baudin, Victoria T. Ly, Nico Hawthorne, Dylan Yong, Sebastian E. Sanso, Nick Rezaee, Jess L. Sevetson, Spencer T. Seiler, Rob Currie, Alex A. Pollen, Keith B. Hengen, Tomasz J. Nowakowski, Mohammed A. Mostajo-Radji, Sofie R. Salama, Mircea Teodorescu, and David Haussler. IoT cloud laboratory:

Internet of Things architecture for cellular biology. *Internet of Things*, 20:100618, November 2022.

- [38] Kateryna Voitiuk, Spencer T. Seiler, Mirella Pessoa de Melo, Jinghui Geng, Tjitse van der Molen, Sebastian Hernandez, Hunter E. Schweiger, Jess L. Sevetson, David F. Parks, Ash Robbins, Sebastian Torres-Montoya, Drew Ehrlich, Matthew A. T. Elliott, Tal Sharf, David Haussler, Mohammed A. Mostajo-Radji, Sofie R. Salama, and Mircea Teodorescu. A feedback-driven brain organoid platform enables automated maintenance and high-resolution neural activity monitoring. *Internet of Things*, 33:101671, September 2025.
- [39] Michael H Dickinson, Claire T Farley, Robert J Full, MAR Koehl, Rodger Kram, and Steven Lehman. How animals move: an integrative view. *science*, 288(5463):100–106, 2000.
- [40] Maryam Tebyani, Ash Robbins, William Asper, Sri Kurniawan, Mircea Teodorescu, Zhongkui Wang, and Shinichi Hirai. 3d printing an assembled biomimetic robotic finger. In *2020 17th International Conference on Ubiquitous Robots (UR)*, pages 526–532. IEEE, 2020.
- [41] Shuhei Miyashita, Steven Guitron, Marvin Ludersdorfer, Cynthia R Sung, and Daniela Rus. An untethered miniature origami robot that self-folds, walks, swims, and degrades. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1490–1496. IEEE, 2015.
- [42] Septika Pedi Artati, Patrick Van der Smagt, Dipl-Ing Nadine Krüger, and José Ma Benitez Baena. *Calculation of Human Arm Stiffness using a Biomechanical Model*. PhD thesis, Technische Universität München, 2008.
- [43] Stefan Fornalski, Ranjan Gupta, and Thay Q Lee. Anatomy and biomechanics of the elbow joint. *Sports medicine and arthroscopy review*, 11(1):1–9, 2003.
- [44] Ge Wu, Frans CT Van der Helm, HEJ DirkJan Veeger, Mohsen Makhsous, Peter Van Roy, Carolyn Anglin, Jochem Nagels, Andrew R Karduna, Kevin McQuade, Xuguang Wang, et al. Isb recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion—part ii: shoulder, elbow, wrist and hand. *Journal of biomechanics*, 38(5):981–992, 2005.
- [45] Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, pages 9191–9200, 2018.

- [46] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, Cambridge, MA, 2nd edition, 2018.
- [47] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [48] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897, 2015.
- [49] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- [50] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [51] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [52] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.
- [53] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1–8. IEEE, 2018.
- [54] Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332*, 2018.
- [55] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

- [56] Sarthak Bhagat, Hritwick Banerjee, Zion Tsz Ho Tse, and Hongliang Ren. Deep reinforcement learning for soft, flexible robots: brief review with impending challenges. *Robotics*, 8(1):4, 2019.
- [57] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- [58] John P Andrews, Jinghui Geng, Kateryna Voitiuk, Matthew AT Elliott, David Shin, Ash Robbins, Alex Spaeth, Albert Wang, Lin Li, Daniel Solis, et al. Multimodal evaluation of network activity and optogenetic interventions in human hippocampal slices. *Nature Neuroscience*, 27:2487–2499, 2024.
- [59] K. Staley. Molecular mechanisms of epilepsy. *Nature Neuroscience*, 18:367–372, 2015.
- [60] K. J. Staley and F. E. Dudek. Interictal spikes and epileptogenesis. *Epilepsy Currents*, 6:199–202, 2006.
- [61] J. T. Paz and J. R. Huguenard. Microcircuits and their interactions in epilepsy: is the focus out of focus? *Nature Neuroscience*, 18:351–359, 2015.
- [62] F. H. Crick. Thinking about the brain. *Scientific American*, 241:219–233, 1979.
- [63] E. S. Boyden. Optogenetics and the future of neuroscience. *Nature Neuroscience*, 18:1200–1201, 2015.
- [64] S.-C. Lin, K. Deisseroth, and J. M. Henderson. Optogenetics: background and concepts for neurosurgery. *Neurosurgery*, 69:1–3, 2011.
- [65] E. S. Boyden, F. Zhang, E. Bamberg, G. Nagel, and K. Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8:1263–1268, 2005.
- [66] V. Gradinaru et al. Molecular and cellular approaches for diversifying and extending optogenetics. *Cell*, 141:154–165, 2010.
- [67] F. Zhang et al. Optogenetic interrogation of neural circuits: technology for probing mammalian brain structures. *Nature Protocols*, 5:439–456, 2010.

- [68] E. Krook-Magnuson, C. Armstrong, M. Oijala, and I. Soltesz. On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy. *Nature Communications*, 4:1376, 2013.
- [69] E. Krook-Magnuson et al. In vivo evaluation of the dentate gate theory in epilepsy. *Journal of Physiology*, 593:2379–2388, 2015.
- [70] M. Fradot et al. Gene therapy in ophthalmology: validation on cultured retinal cells and explants from postmortem human eyes. *Human Gene Therapy*, 22:587–593, 2011.
- [71] J.-A. Sahel et al. Partial recovery of visual function in a blind patient after optogenetic therapy. *Nature Medicine*, 27:1223–1229, 2021.
- [72] M. Avoli et al. Epileptiform activity induced by low extracellular magnesium in the human cortex maintained in vitro. *Annals of Neurology*, 30:589–596, 1991.
- [73] L. M. Masukawa, M. Higashima, J. H. Kim, and D. D. Spencer. Epileptiform discharges evoked in hippocampal brain slices from epileptic patients. *Brain Research*, 493:168–174, 1989.
- [74] I. Cohen, V. Navarro, S. Clemenceau, M. Baulac, and R. Miles. On the origin of interictal activity in human temporal lobe epilepsy in vitro. *Science*, 298:1418–1421, 2002.
- [75] R. Köhling and M. Avoli. Methodological approaches to exploring epileptic disorders in the human brain in vitro. *Journal of Neuroscience Methods*, 155:1–19, 2006.
- [76] M. Andersson et al. Optogenetic control of human neurons in organotypic brain cultures. *Scientific Reports*, 6:24818, 2016.
- [77] M. Ledri, M. Andersson, J. Wickham, and M. Kokaia. Optogenetics for controlling seizure circuits for translational approaches. *Neurobiology of Disease*, 184:106234, 2023.
- [78] R. C. Wykes et al. Optogenetic and potassium channel gene therapy in a rodent model of focal neocortical epilepsy. *Science Translational Medicine*, 4:161ra152, 2012.
- [79] R. C. Wykes, D. M. Kullmann, I. Pavlov, and V. Magloire. Optogenetic approaches to treat epilepsy. *Journal of Neuroscience Methods*, 260:215–220, 2016.

- [80] I. Blümcke et al. International consensus classification of hippocampal sclerosis in temporal lobe epilepsy: a task force report from the ilae commission on diagnostic methods. *Epilepsia*, 54:1315–1329, 2013.
- [81] R. S. Jones, A. B. da Silva, R. G. Whittaker, G. L. Woodhall, and M. O. Cunningham. Human brain slices for epilepsy research: pitfalls, solutions and future challenges. *Journal of Neuroscience Methods*, 260:221–232, 2016.
- [82] E. G. Govorunova et al. Kalium channelrhodopsins are natural light-gated potassium channels that mediate optogenetic inhibition. *Nature Neuroscience*, 25:967–974, 2022.
- [83] J. Vierock et al. WiChR, a highly potassium-selective channelrhodopsin for low-light one-and two-photon inhibition of excitable cells. *Science Advances*, 8:eadd7729, 2022.
- [84] Z. Ma, G. G. Turrigiano, R. Wessel, and K. B. Hengen. Cortical circuit dynamics are homeostatically tuned to criticality in vivo. *Neuron*, 104:655–664, 2019.
- [85] M. Mele et al. Transient incubation of cultured hippocampal neurons in the absence of magnesium induces rhythmic and synchronized epileptiform-like activity. *Scientific Reports*, 11:11374, 2021.
- [86] I. Mody, J. Lambert, and U. Heinemann. Low extracellular magnesium induces epileptiform activity and spreading depression in rat hippocampal slices. *Journal of Neurophysiology*, 57:869–888, 1987.
- [87] C. J. Keller et al. Heterogeneous neuronal firing patterns during interictal epileptiform discharges in the human cortex. *Brain*, 133:1668–1681, 2010.
- [88] M. Lévesque and M. Avoli. The kainic acid model of temporal lobe epilepsy. *Neuroscience & Biobehavioral Reviews*, 37:2887–2899, 2013.
- [89] A. Vezzani et al. Interleukin-1 β immunoreactivity and microglia are enhanced in the rat hippocampus by focal kainate application: functional evidence for enhancement of electrographic seizures. *Journal of Neuroscience*, 19:5054–5065, 1999.
- [90] E. M. Jimenez-Mateos et al. Silencing microRNA-134 produces neuroprotective and prolonged seizure-suppressive effects. *Nature Medicine*, 18:1087–1094, 2012.

- [91] G. Huberfeld et al. Glutamatergic pre-ictal discharges emerge at the transition to seizure in human epilepsy. *Nature Neuroscience*, 14:627–634, 2011.
- [92] M. R. Salick, E. Lubeck, A. Riesselman, and A. Kaykas. The future of cerebral organoids in drug discovery. *Seminars in Cell & Developmental Biology*, 111:67–73, 2021.
- [93] H. Wang. Modeling neurological diseases with human brain organoids. *Frontiers in Synaptic Neuroscience*, 10:15, 2018.
- [94] B. Paulsen et al. Autism genes converge on asynchronous development of shared neuron classes. *Nature*, 602:268–273, 2022.
- [95] P. P. Garcez, E. C. Loiola, R. Madeiro da Costa, L. M. Higa, P. Trindade, R. Delvecchio, J. M. Nascimento, R. Brindeiro, A. Tanuri, and S. K. Rehen. Zika virus impairs growth in human neurospheres and brain organoids. *Science*, 352:816–818, 2016.
- [96] M. G. Andrews et al. Tropism of SARS-CoV-2 for human cortical astrocytes. *Proceedings of the National Academy of Sciences of the United States of America*, 119:e2122236119, 2022.
- [97] V. T. Ly, P. V. Baudin, P. Pansodtee, E. A. Jung, K. Voitiuk, Y. M. Rosen, H. R. Willsey, G. L. Mantalas, S. T. Seiler, J. A. Selberg, S. A. Cordero, J. M. Ross, M. Rolandi, A. A. Pollen, T. J. Nowakowski, D. Haussler, M. A. Mostajo-Radji, S. R. Salama, and M. Teodorescu. Picroscope: low-cost system for simultaneous longitudinal biological imaging. *Communications Biology*, 4:1261, 2021.
- [98] P. V. Baudin, V. T. Ly, P. Pansodtee, E. A. Jung, R. Currie, R. Hoffman, H. R. Willsey, A. A. Pollen, T. J. Nowakowski, D. Haussler, M. A. Mostajo-Radji, S. R. Salama, and M. Teodorescu. Low cost cloud based remote microscopy for biological sciences. *Internet of Things*, 18:100454, 2022.
- [99] S. B. Nagarathna, A. Gehlot, M. Tiwari, T. Tiwari, M. K. Chakravarthi, and D. Verma. A review of bio-cell culture processes in real-time monitoring approach with cloud computing techniques. In *2022 2nd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)*, pages 618–621. IEEE, 2022.
- [100] D. F. Parks et al. IoT cloud laboratory: internet of things architecture for cellular biology. *Internet of Things*, 20:100618, 2022.

- [101] P. V. Baudin et al. Cloud-controlled microscopy enables remote project-based biology education in underserved Latinx communities. *Heliyon*, 8:e11596, 2022.
- [102] Z. Hossain, E. W. Bumbacher, A. M. Chung, H. Kim, C. Litton, A. D. Walter, S. N. Pradhan, K. Jona, P. Blikstein, and I. H. Riedel-Kruse. Interactive and scalable biology cloud experimentation for scientific inquiry and education. *Nature Biotechnology*, 34:1293–1298, 2016.
- [103] D. Tsybulsky and E. Sinai. IoT in project-based biology learning: students’ experiences and skill development. *Journal of Science Education and Technology*, 31:542–553, 2022.
- [104] B. Tabuenca, J. L. Moreno-Sancho, J. Arquero-Gallego, W. Greller, and D. Hernández-Leo. Generating an environmental awareness system for learning using IoT technology. *Internet of Things*, 22:100756, 2023.
- [105] T. Sano, M. J. N. Sampad, J. Gonzalez-Ferrer, S. Hernandez, S. Vera-Choqueccota, P. A. Vargas, R. Urcuyo, N. M. Duran, M. Teodorescu, D. Haussler, H. Schmidt, and M. A. Mostajo-Radji. Open-loop lab-on-a-chip technology enables remote computer science training in Latinx life sciences students. *bioRxiv*, page 538776, 2023.
- [106] M. Radivojevic, D. Jäckel, M. Altermatt, J. Müller, V. Viswam, A. Hierlemann, and D. J. Bakkum. Electrical identification and selective microstimulation of neuronal compartments based on features of extracellular action potentials. *Scientific Reports*, 6:31332, 2016.
- [107] M. Pachitariu, S. Sridhar, and C. Stringer. Solving the spike sorting problem with Kilosort. *bioRxiv*, page 523036, 2023.
- [108] T. Sharf, T. van der Molen, S. M. K. Glasauer, E. Guzman, A. P. Buccino, G. Luna, Z. Cheng, M. Audouard, K. G. Ranasinghe, K. Kudo, S. S. Nagarajan, K. R. Tovar, L. R. Petzold, A. Hierlemann, P. K. Hansma, and K. S. Kosik. Functional neuronal circuitry and oscillatory dynamics in human brain organoids. *Nature Communications*, 13:4403, 2022.
- [109] C. S. Cutts and S. J. Eglén. Detecting pairwise correlations in spike trains: an objective comparison of methods and application to the study of retinal waves. *Journal of Neuroscience*, 34:14288–14303, 2014.
- [110] S. R. Fair, D. Julian, A. M. Hartlaub, S. T. Pusuluri, G. Malik, T. L. Summerfield, G. Zhao, A. B. Hester, W. E. Ackerman, E. W. Hollingsworth,

- M. Ali, C. A. McElroy, I. A. Buhimschi, J. Imitola, N. L. Maitre, T. A. Bedrosian, and M. E. Hester. Electrophysiological maturation of cerebral organoids correlates with dynamic morphological and cellular development. *Stem Cell Reports*, 15:855–868, 2020.
- [111] Z. Ye, M. A. Mostajo-Radji, J. R. Brown, C. Rouaux, G. S. Tomassy, T. K. Hensch, and P. Arlotta. Instructing perisomatic inhibition by direct lineage reprogramming of neocortical projection neurons. *Neuron*, 88:475–483, 2015.
- [112] C. R. Cadwell, A. Palasantza, X. Jiang, P. Berens, Q. Deng, M. Yilmaz, J. Reimer, S. Shen, M. Bethge, K. F. Tolias, R. Sandberg, and A. S. Tolias. Electrophysiological, transcriptomic and morphologic profiling of single neurons using Patch-seq. *Nature Biotechnology*, 34:199–203, 2016.
- [113] H. Zeng and J. R. Sanes. Neuronal cell-type classification: challenges, opportunities and the path forward. *Nature Reviews Neuroscience*, 18:530–546, 2017.
- [114] N. W. Gouwens et al. Classification of electrophysiological and morphological neuron types in the mouse visual cortex. *Nature Neuroscience*, 22:1182–1195, 2019.
- [115] A. P. Passaro and S. L. Stice. Electrophysiological analysis of brain organoids: current approaches and advancements. *Frontiers in Neuroscience*, 14:622137, 2021.
- [116] C. A. Trujillo, R. Gao, P. D. Negraes, J. Gu, J. Buchanan, S. Preissl, A. Wang, W. Wu, G. G. Haddad, I. A. Chaim, A. Domissy, M. Vandenberghe, A. Devor, G. W. Yeo, B. Voytek, and A. R. Muotri. Complex oscillatory waves emerging from cortical organoids model early human brain network development. *Cell Stem Cell*, 25:558–569.e7, 2019.
- [117] H. Cai, Z. Ao, C. Tian, Z. Wu, H. Liu, J. Tchieu, M. Gu, K. Mackie, and F. Guo. Brain organoid computing for artificial intelligence. *bioRxiv*, page 530502, 2023.
- [118] M. P. Zafeiriou, G. Bao, J. Hudson, R. Halder, A. Blenkle, M. K. Schreiber, A. Fischer, D. Schild, and W. H. Zimmermann. Developmental GABA polarity switch and neuronal plasticity in bioengineered neuronal organoids. *Nature Communications*, 11:3791, 2020.
- [119] S. S. Demir. Interactive cell modeling web-resource, iCell, as a simulation-based teaching and learning tool to supplement electrophysiology education. *Annals of Biomedical Engineering*, 34:1077–1087, 2006.

- [120] R. C. Cannon and C. Hammond. Virtual electrophysiology with SPatch [1] applications in undergraduate education. *Brains Minds Media*, 3:bmm1403, 2008.
- [121] S. Formella-Zimmermann, M. W. Kleespies, and P. W. Dierkes. Motivation and technology acceptance in a neuroscience student lab—an empirical comparison between virtual experiments with and without hands-on elements. *Frontiers in Education*, 7:89, 2022.
- [122] T. Yamamoto, A. Weitemier, and M. Kurokawa. Smartphone-enabled web-based simulation of cellular neurophysiology for laboratory course and its effectiveness. *Journal of Undergraduate Neuroscience Education*, 21:A151–A158, 2023.
- [123] T. C. Marzullo and G. J. Gage. The SpikerBox: a low cost, open-source bioamplifier for increasing public participation in neuroscience inquiry. *PLoS One*, 7:e30837, 2012.
- [124] L. M. R. Ferreira, G. A. Carosso, N. Montellano Duran, S. V. Bohorquez-Massud, G. Vaca-Diez, L. I. Rivera-Betancourt, Y. Rodriguez, D. G. Ordonez, D. K. Alatraste-Gonzalez, A. Vacaflares, L. Gonzalez Auza, C. Schuetz, L. E. Alvarado-Arnez, C. V. Alexander-Savino, O. Gandarilla, and M. A. Mostajo-Radji. Effective participatory science education in a diverse Latin American population. *Palgrave Communications*, 5:63, 2019.
- [125] C. Hanzlick-Burton, J. Ciric, M. Diaz-Rios, 3rd Colgan, W., and G. J. Gage. Developing and implementing low-cost remote laboratories for undergraduate biology and neuroscience courses. *Journal of Undergraduate Neuroscience Education*, 19:A118–A123, 2020.
- [126] B. Ramadan and U. M. Ricoy. The NEURON program: utilizing low-cost neuroscience for remote education outreach. *Journal of Undergraduate Neuroscience Education*, 21:A126–A132, 2023.
- [127] B. Miccoli, C. M. Lopez, E. Goikoetxea, J. Putzeys, M. Sekeri, O. Krylychkin, S. W. Chang, A. Firrincieli, A. Andrei, V. Reumers, and D. Braeken. High-density electrical recording and impedance imaging with a multi-modal CMOS multi-electrode array chip. *Frontiers in Neuroscience*, 13:641, 2019.
- [128] K. Voitiuk, J. Geng, M. G. Keefe, D. F. Parks, S. E. Sanso, N. Hawthorne, D. B. Freeman, R. Currie, M. A. Mostajo-Radji, A. A. Pollen, T. J. Nowakowski, S. R. Salama, M. Teodorescu, and D. Haussler. Light-weight electrophysiology hardware and software platform for cloud-based neural recording experiments. *Journal of Neural Engineering*, 18:066004, 2021.

- [129] Y. Jimbo, T. Tatenno, and H. P. C. Robinson. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophysical Journal*, 76:670–678, 1999.
- [130] D. O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press, London, 2005.
- [131] N. Caporale and Y. Dan. Spike timing-dependent plasticity: a Hebbian learning rule. *Annual Review of Neuroscience*, 31:25–46, 2008.
- [132] D. Harnack, M. Pelko, A. Chaillet, Y. Chitour, and M. C. van Rossum. Stability of neuronal networks with homeostatic regulation. *PLoS Computational Biology*, 11:e1004357, 2015.
- [133] R. Q. Quiroga. Spike sorting. *Current Biology*, 22:R45–R46, 2012.
- [134] D. K. Cullen, W. J. Gordián-Vélez, L. A. Struzyna, D. Jgamadze, J. Lim, K. L. Wofford, K. D. Browne, and H. I. Chen. Bundled three-dimensional human axon tracts derived from brain organoids. *iScience*, 21:57–67, 2019.
- [135] T. Kirihaara, Z. Luo, S. Y. A. Chow, R. Misawa, J. Kawada, S. Shibata, F. Khoiratee, C. A. Vollette, V. Volz, T. Levi, T. Fujii, and Y. Ikeuchi. A human induced pluripotent stem cell-derived tissue model of a cerebral tract connecting two cortical regions. *iScience*, 14:301–311, 2019.
- [136] D. O. Adewole, L. A. Struzyna, J. C. Burrell, J. P. Harris, A. D. Nemes, D. Petrov, R. H. Kraft, H. I. Chen, M. D. Serruya, J. A. Wolf, and D. K. Cullen. Development of optically controlled “living electrodes” with long-projecting axon tracts for a synaptic brain-machine interface. *Science Advances*, 7:eay5347, 2021.
- [137] S. Colmar, G. A. D. Liem, J. Connor, and A. J. Martin. Exploring the relationships between academic buoyancy, academic self-concept, and academic performance: a study of mathematics and reading among primary school students. *Educational Psychology*, 39:1068–1089, 2019.
- [138] J. Cribbs, X. Huang, and K. Piatek-Jimenez. Relations of mathematics mindset, mathematics anxiety, mathematics identity, and mathematics self-efficacy to STEM career choice: a structural equation modeling approach. *School Science and Mathematics*, 121:275–287, 2021.
- [139] K. L. Rueda-Gómez, L. J. Rodríguez-Muñiz, and L. Muñiz-Rodríguez. Performance and mathematical self-concept in university students using Khan Academy. *Heliyon*, 9:e15441, 2023.

- [140] S. Ronchi, M. Fiscella, C. Marchetti, V. Viswam, J. Müller, U. Frey, and A. Hierlemann. Single-cell electrical stimulation using CMOS-based high-density microelectrode arrays. *Frontiers in Neuroscience*, 13:208, 2019.
- [141] J. K. Haack. Clapping music—a combinatorial problem. *College Mathematics Journal*, 22:224–227, 1991.
- [142] J. K. Haack. The mathematics of Steve Reich’s “clapping music”. In *Bridges: Mathematical Connections in Art, Music, and Science*, pages 87–92, 1998.
- [143] Y. Tamura, K. Kouzaki, T. Kotani, and K. Nakazato. Electrically stimulated contractile activity-induced transcriptomic responses and metabolic remodeling in C2C12 myotubes: twitch vs. tetanic contractions. *American Journal of Physiology—Cell Physiology*, 319:C1029–C1044, 2020.
- [144] David Beniaguev, Idan Segev, and Michael London. Single cortical neurons as deep artificial neural networks. *Neuron*, 109(17):2727–2739.e3, September 2021.
- [145] Danijela Marković, Alice Mizrahi, Damien Querlioz, and Julie Grollier. Physics for neuromorphic computing. *Nature Reviews Physics*, 2(9):499–510, 2020.
- [146] Lena Smirnova, Brian S. Caffo, David H. Gracias, Qi Huang, Itzy E. Morales Pantoja, Bohao Tang, Donald J. Zack, Cynthia A. Berlinicke, J. Lomax Boyd, Timothy D. Harris, Erik C. Johnson, Brett J. Kagan, Jeffrey Kahn, Alysson R. Muotri, Barton L. Paulhamus, Jens C. Schwamborn, Jesse Plotkin, Alexander S. Szalay, Joshua T. Vogelstein, Paul F. Worley, and Thomas Hartung. Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish. *Frontiers in Science*, 1:1017235, February 2023.
- [147] Jeffrey R. Hollerman and Wolfram Schultz. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature neuroscience*, 1(4):304–309, 1998.
- [148] Wolfram Schultz. Getting Formal with Dopamine and Reward. *Neuron*, 36(2):241–263, October 2002.
- [149] Eric E. Thomson, Rafael Carra, and Miguel A. L. Nicolelis. Perceiving invisible light through a somatosensory cortical prosthesis. *Nature Communications*, 4(1):1482, February 2013.

- [150] D. O. Hebb. *The organization of behavior; a neuropsychological theory*. The organization of behavior; a neuropsychological theory. Wiley, Oxford, England, 1949.
- [151] Radhika Madhavan, Zenas C Chao, and Steve M Potter. Plasticity of recurring spatiotemporal activity patterns in cortical networks. *Physical biology*, 4(3):181–193, October 2007.
- [152] Verena Pawlak, Jeffery Wickens, Alfredo Kirkwood, and Jason Kerr. Timing is not Everything: Neuromodulation Opens the STDP Gate. *Frontiers in Synaptic Neuroscience*, 2:146, 2010.
- [153] Harel Shouval, Samuel Wang, and Gayle Wittenberg. Spike Timing Dependent Plasticity: A Consequence of More Fundamental Learning Rules. *Frontiers in Computational Neuroscience*, 4:19, 2010.
- [154] Stephanie C. Seeman, Brian J. Mogen, Eberhard E. Fetz, and Steve I. Perlmuter. Paired Stimulation for Spike-Timing-Dependent Plasticity in Primate Sensorimotor Cortex. *Journal of Neuroscience*, 37(7):1935–1949, February 2017.
- [155] Pieter R. Roelfsema and Anthony Holtmaat. Control of synaptic plasticity in deep cortical networks. *Nature Reviews. Neuroscience*, 19(3):166–180, February 2018.
- [156] Luis Carrillo-Reid and Rafael Yuste. Playing the piano with the cortex: role of neuronal ensembles and pattern completion in perception and behavior. *Current Opinion in Neurobiology*, 64:89–95, October 2020.
- [157] Tal Sharf, Tjitse van der Molen, Stella MK Glasauer, Elmer Guzman, Alessio P Buccino, Gabriel Luna, Zhuowei Cheng, Morgane Audouard, Kamalini G Ranasinghe, Kiwamu Kudo, et al. Functional neuronal circuitry and oscillatory dynamics in human brain organoids. *Nature communications*, 13(1):4403, 2022.
- [158] Steven A Sloan, Jimena Andersen, Anca M Paşca, Fikri Birey, and Sergiu P Paşca. Generation and assembly of human brain region-specific three-dimensional cultures. *Nature protocols*, 13(9):2062–2085, 2018.
- [159] Silvia Ronchi, Michele Fiscella, Camilla Marchetti, Vijay Viswam, Jan Müller, Urs Frey, and Andreas Hierlemann. Single-Cell Electrical Stimulation Using CMOS-Based High-Density Microelectrode Arrays. *Frontiers in Neuroscience*, 13:208, 2019.

- [160] A. Novellino, P. D’Angelo, L. Cozzi, M. Chiappalone, V. Sanguineti, and S. Martinoia. Connecting Neurons to a Mobile Robot: An In Vitro Bidirectional Neural Interface. *Computational Intelligence and Neuroscience*, 2007(1):012725, 2007.
- [161] Jacopo Tessadori, Marta Bisio, Sergio Martinoia, and Michela Chiappalone. Modular Neuronal Assemblies Embodied in a Closed-Loop Environment: Toward Future Integration of Brains and Machines. *Frontiers in Neural Circuits*, 6:99, 2012.
- [162] Yongcheng Li, Rong Sun, Bin Zhang, Yuechao Wang, and Hongyi Li. Application of Hierarchical Dissociated Neural Network in Closed-Loop Hybrid System Integrating Biological and Mechanical Intelligence. *PLOS ONE*, 10(5):e0127452, May 2015.
- [163] Yuichiro Yada, Shusaku Yasuda, and Hirokazu Takahashi. Physical reservoir computing with FORCE learning in a living neuronal culture. *Applied Physics Letters*, 119(17):173701, October 2021.
- [164] Hiroyuki Tetsuka, Samuele Gobbi, Takaaki Hatanaka, Lorenzo Pirrami, and Su Ryon Shin. Wirelessly steerable bioelectronic neuromuscular robots adapting neurocardiac junctions. *Science Robotics*, 9(94):eado0051, September 2024.
- [165] Joost le Feber, Jan Stegenga, and Wim L. C. Rutten. The effect of slow electrical stimuli to achieve learning in cultured networks of rat cortical neurons. *PloS One*, 5(1):e8871, January 2010.
- [166] Lana Sinapayen, Atsushi Masumori, and Takashi Ikegami. Learning by stimulation avoidance: A principle to control spiking neural networks dynamics. *PLOS ONE*, 12(2):e0170388, February 2017.
- [167] W R Chen, S Lee, K Kato, D D Spencer, G M Shepherd, and A Williamson. Long-term modifications of synaptic efficacy in the human inferior and middle temporal cortex. *Proceedings of the National Academy of Sciences*, 93(15):8011–8015, July 1996.
- [168] Therese Abrahamsson, Txomin Lalanne, Alanna J. Watt, and P. Jesper Sjöström. In Vitro Investigation of Synaptic Plasticity. *Cold Spring Harbor protocols*, 2016(6):pdb.top087262, June 2016.
- [169] A. Odawara, H. Katoh, N. Matsuda, and I. Suzuki. Induction of long-term potentiation and depression phenomena in human induced pluripotent

- stem cell-derived cortical neurons. *Biochemical and Biophysical Research Communications*, 469(4):856–862, January 2016.
- [170] Takuya Isomura and Karl Friston. In vitro neural networks minimise variational free energy. *Scientific Reports*, 8(1):16926, November 2018.
 - [171] Takuya Isomura, Kiyoshi Kotani, Yasuhiko Jimbo, and Karl J. Friston. Experimental validation of the free-energy principle with in vitro neural networks. *Nature Communications*, 14(1):4547, August 2023.
 - [172] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR, 2017.
 - [173] Vincent Strong, William Holderbaum, and Yoshikatsu Hayashi. Electroactive polymer hydrogels exhibit emergent memory when embodied in a simulated game environment. *Cell Reports Physical Science*, 5(9):102151, September 2024.
 - [174] Andrew Stephenson. Xx. on induced stability. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 15(86):233–236, 1908.
 - [175] Andrew G Barto, Richard S Sutton, and Charles W Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics*, SMC-13(5):834–846, 1983.
 - [176] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.
 - [177] Mototsugu Eiraku, Kiichi Watanabe, Mami Matsuo-Takasaki, Masako Kawada, Shigenobu Yonemura, Michiru Matsumura, Takafumi Wataya, Ayaka Nishiyama, Keiko Muguruma, and Yoshiki Sasai. Self-organized formation of polarized cortical tissues from escs and its active manipulation by extrinsic signals. *Cell stem cell*, 3(5):519–532, 2008.
 - [178] Jinghui Geng, Kateryna Voitiuk, David F Parks, Ash Robbins, Alex Spaeth, Jessica L Sevetson, Sebastian Hernandez, Hunter E Schweiger, John P Andrews, Spencer T Seiler, et al. Multiscale cloud-based pipeline for neuronal electrophysiology analysis and visualization. *bioRxiv*, page 2024.11.14.623530, 2024.

- [179] Kateryna Voitiuk, Spencer T Seiler, Mirella Pessoa de Melo, Jinghui Geng, Sebastian Hernandez, Hunter E Schweiger, Jess L Sevetson, David F Parks, Ash Robbins, Sebastian Torres-Montoya, et al. A feedback-driven iot microfluidic, electrophysiology, and imaging platform for brain organoid studies. *bioRxiv*, 2024.
- [180] Milos Radivojevic, David Jäckel, Michael Altermatt, Jan Müller, Vijay Viswam, Andreas Hierlemann, and Douglas J. Bakkum. Electrical Identification and Selective Microstimulation of Neuronal Compartments Based on Features of Extracellular Action Potentials. *Scientific Reports*, 6(1):31332, August 2016.
- [181] Ranulfo Romo, Adrián Hernández, Anótonio Zainos, and Emilio Salinas. Somatosensory discrimination based on cortical microstimulation. *Nature*, 392(6674):387–390, March 1998.
- [182] Wilder Penfield and Edwin Boldrey. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60(4):389–443, December 1937.
- [183] Juan A. Gallego, Matthew G. Perich, Lee E. Miller, and Sara A. Solla. Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984, June 2017. Publisher: Elsevier.
- [184] Catherine S. Cutts and Stephen J. Eglen. Detecting Pairwise Correlations in Spike Trains: An Objective Comparison of Methods and Application to the Study of Retinal Waves. *The Journal of Neuroscience*, 34(43):14288–14303, October 2014.
- [185] Mohammed A Mostajo-Radji, Walter R Mancia Leon, Arnar Breevoort, Jesus Gonzalez-Ferrer, Hunter E Schweiger, Julian Lehrer, Li Zhou, Matthew T Schmitz, Yonatan Perez, Tanzila Mukhtar, et al. Fate plasticity of interneuron specification. *bioRxiv*, page 2024.10.02.614266, 2024.
- [186] Wolfram Schultz. Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, 80(1):1–27, 1998.
- [187] Masataka Watanabe. Reward expectancy in primate prefrontal neurons. *Nature*, 382(6592):629–632, 1996.
- [188] Camillo Padoa-Schioppa and John A. Assad. Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090):223–226, May 2006.

- [189] Pieter R Roelfsema and Anthony Holtmaat. Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*, 19(3):166–180, 2018.
- [190] Yunjeong Park, Sebastian Hernandez, Cristian O Hernandez, Hunter E Schweiger, Houpu Li, Kateryna Voitiuk, Harika Dechiraju, Nico Hawthorne, Elana M Muzzy, John A Selberg, et al. Modulation of neuronal activity in cortical organoids with bioelectronic delivery of ions and neurotransmitters. *Cell Reports Methods*, 4(1):100686, 2024.
- [191] David F Parks, Kateryna Voitiuk, Jinghui Geng, Matthew AT Elliott, Matthew G Keefe, Erik A Jung, Ash Robbins, Pierre V Baudin, Victoria T Ly, Nico Hawthorne, et al. Iot cloud laboratory: Internet of things architecture for cellular biology. *Internet of Things*, 20:100618, 2022.
- [192] Daniel A. Wagenaar, Jerome Pine, and Steve M. Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *Journal of Neuroscience Methods*, 138(1):27–37, September 2004.
- [193] Razvan V Florian. Correct equations for the dynamics of the cart-pole system, 2005.
- [194] E. D. Adrian and Yngve Zotterman. The impulses produced by sensory nerve endings. *The Journal of Physiology*, 61(4):465–483, August 1926.
- [195] Zenas C. Chao, Douglas J. Bakkum, and Steve M. Potter. Region-specific network plasticity in simulated and living cortical networks: comparison of the center of activity trajectory (CAT) with other statistics. *Journal of Neural Engineering*, 4(3):294, July 2007.
- [196] Simone Mayer, Jiadong Chen, Dmitry Velmeshev, Andreas Mayer, Ugomma C Eze, Aparna Bhaduri, Carlos E Cunha, Diane Jung, Arpana Arjun, Emmy Li, et al. Multimodal single-cell analysis reveals physiological maturation in the developing human neocortex. *Neuron*, 102(1):143–158, 2019.
- [197] Tjitse van der Molen, Alex Spaeth, Mattia Chini, Sebastian Hernandez, Gregory A. Kaurala, Hunter E. Schweiger, Cole Duncan, Sawyer McKenna, Jinghui Geng, Max Lim, Julian Bartram, Aditya Dendukuri, Zongren Zhang, Jesus Gonzalez-Ferrer, Kiran Bhaskaran-Nair, Lon J. Blauvelt, Cole R. K. Harder, Linda R. Petzold, Dowlette-Mary Alam El Din, Jason Laird, Maren Schenke, Lena Smirnova, Bradley M. Colquitt, Mohammed A. Mostajo-Radji, Paul K. Hansma, Mircea Teodorescu, Andreas Hierlemann, Keith B. Hengen, Ileana L. Hanganu-Opatz, Kenneth S. Kosik, and Tal

Sharf. Protosequences in brain organoids model intrinsic brain states, January 2025. Pages: 2023.12.29.573646 Section: New Results.

- [198] Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. Fiji: an open-source platform for biological-image analysis. *Nature methods*, 9(7):676–682, 2012.
- [199] Kenneth R Tovar, Daniel C Bridges, Bian Wu, Clayton Randall, Morgane Audouard, Jayashree Bhatt, et al. Action potential propagation recorded from single axonal arbors using multielectrode arrays. *Journal of Neurophysiology*, 120(7):306–320, 2018.
- [200] Xiaoxuan Jia, Joshua H Siegle, Corbett Bennett, Samuel D Gale, Daniel J Denman, Christof Koch, et al. High-density extracellular probes reveal dendritic backpropagation and facilitate neuron classification. *Journal of Neurophysiology*, 121(5):1831–1847, 2019.
- [201] Gerrit Hilgen, Martino Sorbaro, Sahar Pirmoradian, Jens-Oliver Muthmann, Ibolya Edit Kepiro, Simona Ullo, et al. Unsupervised spike sorting for large-scale, high-density multielectrode arrays. *Cell Reports*, 18(10):2521–2532, 2017.
- [202] Marius Pachitariu, Nicholas Steinmetz, and Jennifer Colonell. Kilosort2. <https://github.com/MouseLand/Kilosort2>, 2018.
- [203] James J Jun, Catalin Mitelut, Chongxi Lai, Sergey L Gratiy, Costas A Anastassiou, and Timothy D Harris. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. *bioRxiv*, page 101030, 2017.
- [204] Pierre Yger, Giulia Lia Spampinato, Elric Esposito, Baptiste Lefebvre, Stéphane Deny, Christophe Gardella, et al. A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife*, 7:e34518, 2018.
- [205] Samuel Garcia and Christophe Pouzat. Tridesclous. <https://github.com/tridesclous/tridesclous>, 2015.
- [206] Roland Diggelmann, Michele Fiscella, Andreas Hierlemann, and Felix Franke. Automatic spike sorting for high-density microelectrode arrays. *Journal of Neurophysiology*, 120(6):3155–3171, 2018.

- [207] James J Jun, Nicholas A Steinmetz, Joshua H Siegle, Daniel J Denman, Marius Bauza, Brian Barbarits, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017.
- [208] Marco Ballini, Jan Müller, Paolo Livi, Yihui Chen, Urs Frey, Alexander Stettler, et al. A 1024-channel CMOS microelectrode array with 26,400 electrodes for recording and stimulation of electrogenic cells in vitro. *IEEE Journal of Solid-State Circuits*, 49(11):2705–2719, 2014.
- [209] Alessio Paolo Buccino, Cole L Hurwitz, Samuel Garcia, Jeremy Magland, Joshua H Siegle, Roger Hurwitz, et al. SpikeInterface, a unified framework for spike sorting. *eLife*, 9:e61834, 2020.
- [210] Joshua H Siegle, Xiaoxuan Jia, Séverine Durand, Sam Gale, Corbett Bennett, Nile Graddis, et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, 592(7852):86–92, 2021.
- [211] Zhiwen Ye, Andrew M Shelton, Jawad R Shaker, Julien Boussard, Jennifer Colonell, Sahar Manavi, et al. Ultra-high density electrodes improve detection, yield, and cell type specificity of brain recordings. *bioRxiv*, 2023.
- [212] Daniel A Wagenaar, Radhika Madhavan, Jerome Pine, and Steve M Potter. Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *Journal of Neuroscience*, 25(3):680–688, 2005.
- [213] Thomas B DeMarse, Daniel A Wagenaar, Axel W Blau, and Steve M Potter. The neurally controlled animat: biological brains acting with simulated bodies. *Autonomous Robots*, 11:305–310, 2001.
- [214] Brett J Kagan, Andy C Kitchen, Nhi T Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J Parker, et al. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969, 2022.
- [215] Douglas J Bakkum, Zenas C Chao, and Steve M Potter. Adaptive goal-directed behavior in embodied cultured networks: living neuronal networks and a simulated model. In *2007 3rd International IEEE/EMBS Conference on Neural Engineering*, pages 46–49. IEEE, 2007.
- [216] Daniel A. Wagenaar, Jerome Pine, and Steve M. Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *Journal of Neuroscience Methods*, 138(1–2):27–37, September 2004.

- [217] Chadwick M. Hales, John D. Rolston, and Steve M. Potter. How to culture, record and stimulate neuronal networks on micro-electrode arrays (meas). *Journal of Visualized Experiments*, (39), May 2010.
- [218] Edward S Boyden, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9):1263–1268, August 2005.
- [219] Donald O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York, 1949.
- [220] Natalia Caporale and Yang Dan. Spike timing-dependent plasticity: A hebbian learning rule. *Annual Review of Neuroscience*, 31(1):25–46, July 2008.
- [221] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, 18(24):10464–10472, December 1998.
- [222] Thomas B. DeMarse, Daniel A. Wagenaar, Axel W. Blau, and Steve M. Potter. The neurally controlled animat: Biological brains acting with simulated bodies. *Autonomous Robots*, 11(3):305–310, November 2001.
- [223] Douglas J Bakkum, Zenas C Chao, and Steve M Potter. Spatio-temporal electrical stimuli shape behavior of an embodied cortical network in a goal-directed learning task. *Journal of Neural Engineering*, 5(3):310–323, August 2008.
- [224] Daniel A. Wagenaar, Thomas B. DeMarse, and Steve M. Potter. MeaBench: A toolset for multi-electrode data acquisition and on-line analysis. In *2nd International IEEE EMBS Conference on Neural Engineering*, pages 518–521, 2005.
- [225] Y. Jimbo, T. Tateno, and H.P.C. Robinson. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophysical Journal*, 76(2):670–678, February 1999.
- [226] Takuya Isomura, Kiyoshi Kotani, Yasuhiko Jimbo, and Karl J. Friston. Experimental validation of the free-energy principle with in vitro neural networks. *Nature Communications*, 14(1), August 2023.
- [227] Brett J. Kagan, Andy C. Kitchen, Nhi T. Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J. Parker, Anjali Bhat, Ben Rollo, Adeel Razi, and

- Karl J. Friston. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969, December 2022.
- [228] Hongwei Cai, Zheng Ao, Chunhui Tian, Zhuohao Wu, Hongcheng Liu, Jason Tchieu, Mingxia Gu, Ken Mackie, and Feng Guo. Brain organoid reservoir computing for artificial intelligence. *Nature Electronics*, 6(12):1032–1039, December 2023.
- [229] Silvia Ronchi, Michele Fiscella, Camilla Marchetti, Vijay Viswam, Jan Muller, Urs Frey, and Andreas Hierlemann. Single-cell electrical stimulation using cmos-based high-density microelectrode arrays. *Frontiers in Neuroscience*, 13, March 2019.
- [230] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Perolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, Daniel Hennes, Dustin Morrill, Paul Muller, Timo Ewalds, Ryan Faulkner, Janos Kramar, Bart De Vylder, Brennan Saeta, James Bradbury, David Ding, Sebastian Borgeaud, Matthew Ghosh, Eugene Lockhart, Kory McKee, Yee Whye Teh, Georgios Piliouras, and Thore Graepel. OpenSpiel: A framework for reinforcement learning in games, 2019.
- [231] J. K. Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S. Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, Niall L. Williams, Yashas Lokesh, and Praveen Ravi. PettingZoo: Gym for multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [232] Milos Radivojevic, David Jackel, Michael Altermatt, Jan Muller, Vijay Viswam, Andreas Hierlemann, and Douglas J. Bakkum. Electrical identification and selective microstimulation of neuronal compartments based on features of extracellular action potentials. *Scientific Reports*, 6(1), August 2016.
- [233] Daniel R. Merrill, Marom Bikson, and John G.R. Jefferys. Electrical stimulation of excitable tissue: design of efficacious and safe protocols. *Journal of Neuroscience Methods*, 141(2):171–198, February 2005.
- [234] Douglas J. Bakkum, Urs Frey, Milos Radivojevic, Thomas L. Russell, Jan Muller, Michele Fiscella, Hirokazu Takahashi, and Andreas Hierlemann. Tracking axonal action potential propagation on a high-density microelectrode array across hundreds of sites. *Nature Communications*, 4(1), July 2013.

- [235] Marco Emilio Vazquez, Enrico Opri, Brandon Parks, and Aysegul Gunduz. DBS artifact removal through template subtraction. *UF Journal of Undergraduate Research*, 21(2), 2019.
- [236] Daniel J. O’Shea and Krishna V. Shenoy. ERAASR: An algorithm for removing electrical stimulation artifacts from multielectrode array recordings. *Journal of Neural Engineering*, 15(2):026020, April 2018.
- [237] Daniel A Wagenaar and Steve M Potter. Real-time multi-channel stimulus artifact suppression by local curve fitting. *Journal of Neuroscience Methods*, 120(2):113–120, October 2002.
- [238] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A LLVM-based Python JIT compiler, 2015.