
On Complex Network Dynamics of an In-Vitro Neuronal System during Rest and Gameplay

Moein Khajehnejad*
Turner Institute for Brain and Mental Health,
Monash University, Australia

Forough Habibollahi*
Cortical Labs
Melbourne, Australia

Alon Loeffler
Cortical Labs
Melbourne, Australia

Brett J. Kagan
Cortical Labs
Melbourne, Australia

Adeel Razi
Turner Institute for Brain and Mental Health, Monash University, Australia
CIFAR Azrieli Global Scholars Program, CIFAR, Toronto, CA,
Wellcome Centre for Human Neuroimaging, University College London, UK

Abstract

In this study, we characterize complex network dynamics in live in-vitro neuronal systems during two distinct activity states: spontaneous rest state and engagement in a real-time (closed-loop) game environment using the *DishBrain* system. First, we embed the spiking activity of these channels in a lower-dimensional space using various representation learning methods and then extract a subset of representative channels. Next, by analyzing these low-dimensional representations, we explore the patterns of macroscopic neuronal network dynamics during learning. Remarkably, our findings indicate that just using the low-dimensional embedding of representative channels is sufficient to differentiate the neuronal culture during the Rest and Gameplay. Notably, our investigation shows dynamic changes in the *connectivity* patterns within the same region and across multiple regions on the multi-electrode array only during Gameplay. These findings underscore the plasticity of neuronal networks in response to external stimuli and highlight the potential for modulating connectivity in a controlled environment. The ability to distinguish between neuronal states using reduced-dimensional representations points to the presence of underlying patterns that could be pivotal for real-time monitoring and manipulation of neuronal cultures. Additionally, this provides insight into how biological based information processing systems rapidly adapt and learn and may lead to new improved algorithms.

1 Introduction

The *DishBrain* system combines biological intelligence with adaptive neuronal traits by seamlessly integrating *in vitro* neuronal networks with *in silico* computational elements using high-density multi-electrode arrays (HD-MEAs). These cultivated neuronal ensembles exhibit biologically-based adaptive intelligence, which is effectively replicated in a dynamic gaming environment through closed-loop stimulation and concurrent recordings [1]. These neural ensembles display self-organized adaptive electrophysiological dynamics, demonstrating their ability to acquire new knowledge and respond meaningfully to biologically plausible external stimuli [2]. The empirical data is sourced from cortical cells derived from embryonic rodent or human induced pluripotent stem cell (hiPSC) lineages. While synthetic biology methods show that *in vitro* biological networks of cortical cells can achieve real-time adaptive goal-directed learning in simulated environments [1, 3, 4], the underlying network dynamics associated with this learning remain unexplored.

We analyze the spiking activity of each channel on the HD-MEA to explore neuronal network dynamics and functional connectivity. Understanding these complex dynamics is crucial for uncovering the neural mechanisms behind learning. Neurons within a network interact to process information and generate

*Authors contributed equally to this manuscript.

responses. Learning involves the modification of synaptic interactions, which affects the signal transmission. Studying the temporal patterns and strength of these interactions reveals the network’s ability to encode, store, and retrieve information. Investigating neural network dynamics helps us understand the mechanisms driving synaptic changes, enhancing our grasp of cellular and network-level processes in learning. These insights have broad implications, spanning neuroscience and artificial intelligence, potentially guiding advanced learning algorithms and treatments for neurological disorders.

2 Methods

2.1 DishBrain System

The *DishBrain* system simulated neural cultures in a virtual ‘Pong’ game. Stimulation used rate coding pulses (4Hz to 40Hz) for the ball’s x-axis and place coding (specific electrodes) for the y-axis within an 8-electrode sensory zone. Paddle movement was controlled by real-time electrophysiological activity in a "motor area". Cultures received feedback on paddle movement. Two types of stimulation were used: sensory as explained and unpredictable feedback [1]. Missed balls triggered a 150 mV, 5 Hz unpredictable stimulus on random sensory electrodes at varied rates for 4 seconds, followed by a 4 second resting period. Each gameplay session lasted 20 minutes at a 20 kHz sampling rate.

Figure 1:a-b illustrates the input information, feedback loop setup, and electrode configurations in the DishBrain system. More details of this system are introduced in Appendix A.

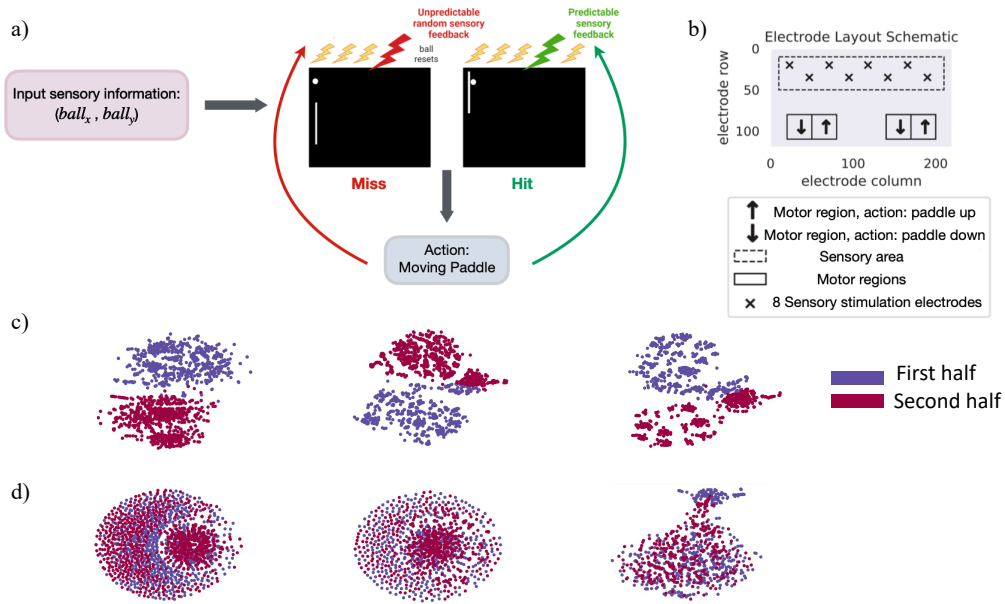


Figure 1: **a)** DishBrain feedback loop setup. **b)** Electrode configuration and predefined sensory and motor regions. Figures adapted from [1]. Low-dimensional representation of 3 samples of **c)** Gameplay and their following **d)** Rest sessions using t-SNE. Purple and maroon dots are channel representations in the embedding space in the first and second half of the recordings.

2.2 Network Construction

Neuronal spiking was recorded from 1024 HD-MEA channels in 248 Gameplay and 147 Rest sessions. Given the length of recordings at the 20 kHz sampling frequency, the resulting times series during the Gameplay was very long. In the realm of mining information from dense and high-dimensional networks, a prominent focus in recent years has been on the concept of acquiring network embeddings in lower dimensions. The primary objective of this direction is to acquire vector representations for individual nodes within the network, which encapsulate valuable and meaningful insights [5, 6, 7]. Hence, in this work, we first employed dimensionality reduction algorithms to both enhance the computational efficiency of subsequent data analysis and improve data interpretability. This also allowed us to uncover latent data structures not immediately apparent in the original high-dimensional space. We used t-SNE [8] and Isomap [9] to create 3D representations for both Rest and Gameplay data. To assess their effectiveness in capturing learning-related network structures, we split recording sessions in half before dimensionality reduction. Figure 1:c-d displays results of t-SNE with color labeling for the first and second halves of 20-minute Gameplay and 10-minute Rest sessions for three

samples. It's evident that the two halves are distinguishable in Gameplay but not in Rest, indicating specific patterns in network dynamics during learning, primarily present in Gameplay.

Prior research has extensively used simplified models of interconnected neural populations based on mean-field approximations, which effectively retain the dynamic properties of the original neural network while significantly accelerating simulation speeds by several orders of magnitude [10, 11, 12, 13]. Moreover, within complex neural networks, only a fraction of neurons fire at any given time, with many not displaying clear action potentials. Recent evidence indicates the emergence of specialized, selective, and abstract response properties in the cortex [14], showcasing the importance of sparse activity and connectivity patterns. These patterns conserve energy and optimize computational capacity [15], emphasizing the redundancy in evaluating individual neuron firing patterns. The brain's ability to encode and process information relies on the concerted action of neuronal populations, often conveying redundant or highly correlated signals. Given these collective behaviors in neuronal networks, our goal was to advance computational complexity reduction when studying large neuronal populations, while still preserving the network's dynamic properties.

We devised a method to pinpoint a subset of recorded channels that likely monitored the neuronal populations specially attuned to the ongoing task. This subset enables the identification of key neurons characterizing the network's behavior during Gameplay, to more efficiently study the (macroscopic) of this smaller and interpretable network. To find a consistent subset of channels across all neuronal cultures, we employed Tucker decomposition –via higher-order orthogonal iteration– on the tensor data from the 248 Gameplay sessions in the lower-dimensional embedding space. The resulting 1024×3 tensor served as a compact representation, capturing underlying patterns and structures. Using this tensor, we identified representative channels by applying the K-medoid clustering algorithm, creating 30 clusters, and extracting the corresponding 'medoids' for each cluster. Attempts with a higher value of K did not notably improve clustering accuracy measured by the Davies-Bouldin index. A network matrix using functional connectivity – defined as the zero-lag Pearson correlations – of each Gameplay or Rest session recording was then built with these 30 channels as the nodes and the edges between these nodes represented by the functional connectivity. Only edges with Pearson correlation absolute values above 0.7 were kept.

Figure 2 is a schematic illustration of the proposed *in vitro* network construction framework in this study.

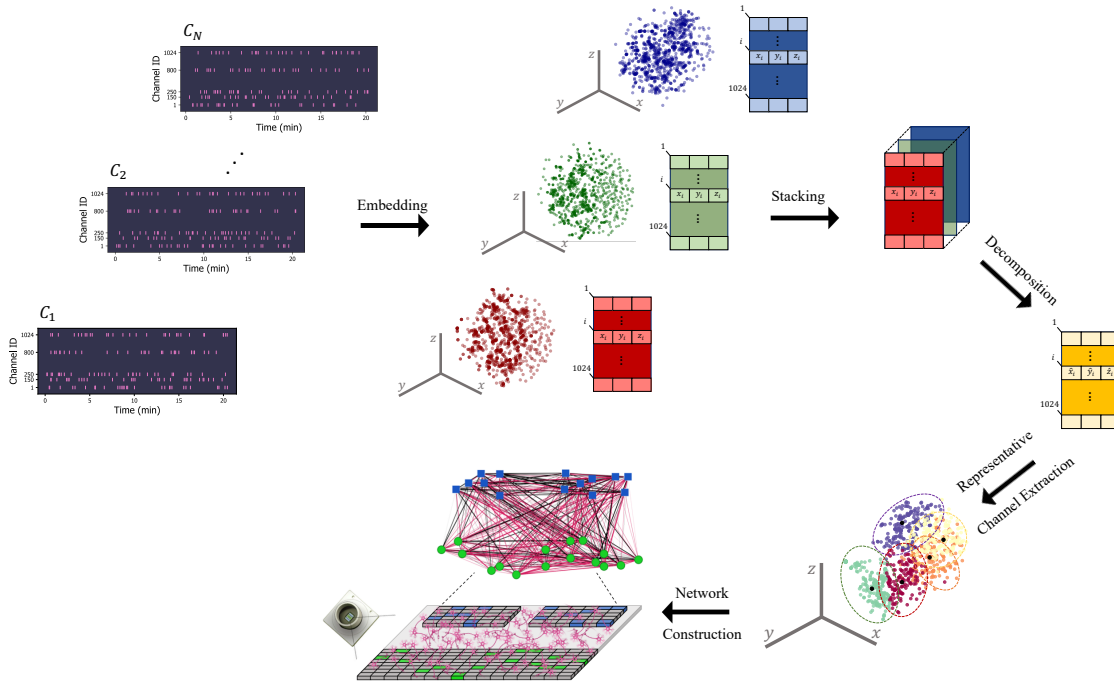


Figure 2: A schematic illustration of the overall network construction framework. The spiking time series data are first transformed into a 3D space using t-SNE embedding. These lower-dimensional representations are then combined into a tensor, which is decomposed using Tucker decomposition. The K-medoids algorithm is then applied to identify consistent representative channels across all cultures. These channels become network nodes, and pairwise Pearson correlation values serve as edge weights. The network layout reflects the physical placement of channels on the MEA, with node colors distinguishing sensory (green) from motor (blue) regions.

3 Results

After constructing the connectivity networks, we aimed to examine their temporal evolution in both Gameplay and Rest. To achieve this, we divided each recording session into 2-minute windows and evaluated the change in edge weights as the network evolved over those windows. Figure 3 shows the differences in the correlation between each pair of nodes when comparing the last and first 2 minutes of each recording. As shown in Figure 3, we found that the cultures, while embedded in the game environment, had a higher number of edges with increased correlation between channels while this change was not apparent during their rest state spontaneous activity. This clearly indicates significant network plasticity occurring in these cultures that can be a necessary underlying mechanism for the learning that happens in this closed-loop system [1]. Moreover, we evaluated several other network characteristics from all of the generated networks and compared them between the first and last 2 minutes of recordings in both Rest and Gameplay groups. Figure 4 shows these results. While all of these metrics showed statistically significant differences during Gameplay, none of them showed statistically significant differences during the Rest condition of the cultures.

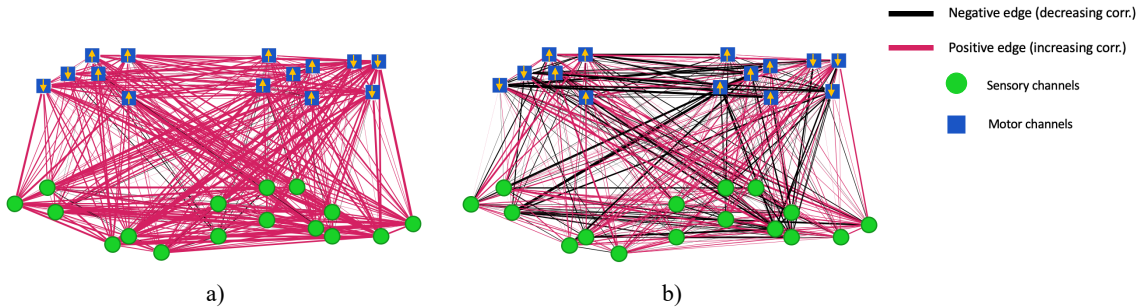


Figure 3: The average networks over all the **a)** Gameplay and **b)** Rest sessions with edge weights representing changes in functional connectivity between channel pairs when comparing the last 2 minutes to the first 2 minutes of recordings. Edge colors signify the direction of these connectivity changes, with red indicating increases and black indicating decreases. Motor and sensory region channels are represented by blue squares and green circles, respectively. Arrows on motor region nodes show the paddle’s movement direction as per their position in the predefined layout in Figure 1.

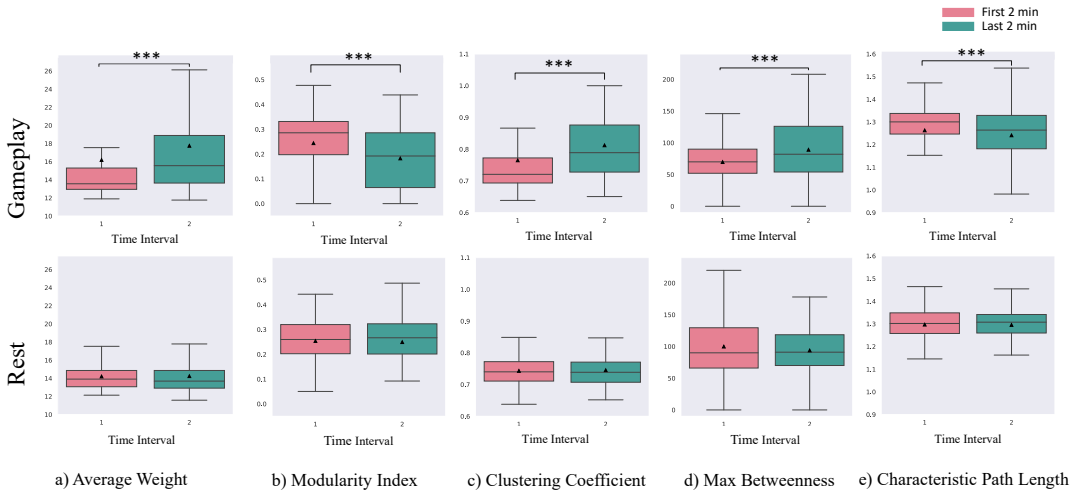


Figure 4: Network summary statistics between the first and last 2 minutes of recordings for **top)** Gameplay and **bottom)** Rest. All the evaluated network metrics show statistically significant differences during Gameplay but not during Rest. One-way ANOVA test, $***p < 10^{-3}$.

The increased average weight of the networks in the Gameplay sessions confirms the same patterns also observed in Figure 3. Interestingly, the decreasing modularity index suggests the change in the community structure of neuronal networks when learning occurs with disconnected communities becoming more connected during Gameplay. Other features such as increased clustering coefficient and decreased characteristic path length are also in line with the increasing pattern of correlation values observed in the connectivity during Gameplay.

4 Discussion

Our study draws inspiration from mean-field theory, a technique used in statistical physics that seeks to create simplified models of networks with the aim of reducing computational complexity. We leverage previous findings on redundancy in large neuron populations which contributes to their robustness, suggesting a subset can represent the network’s dynamics.

Using DishBrain, we place a multitude of *in vitro* cortical neurons in a closed-loop game environment to explore network dynamics during learning compared to Rest conditions. Our innovative approach enables us to 1) study neurons at a cellular level in a closed-loop game environment, 2) analyze neuronal activities in a lower-dimensional space, 3) identify an efficient subpopulation capturing overall dynamics, and 4) examine network dynamics driving information processing and learning.

Our research effectively extracts low-dimensional information and identifies influential units, promising deeper insights into how large neuronal populations function and adapt in complex environments, shedding light on underlying network changes facilitating learning.

References

- [1] Brett J Kagan et al. “In vitro neurons learn and exhibit sentience when embodied in a simulated game-world”. In: *Neuron* 110.23 (2022), pp. 3952–3969.
- [2] Evan R. Harrell et al. “An elaborate sweep-stick code in rat barrel cortex”. In: *Science Advances* 6.38 (Sept. 18, 2020), eabb7189. ISSN: 2375-2548. DOI: 10.1126/sciadv.abb7189. URL: <https://www.science.org/doi/10.1126/sciadv.abb7189> (visited on 06/06/2022).
- [3] Forough Habibollahi et al. “Critical dynamics arise during structured information presentation within embodied in vitro neuronal networks”. In: *Nature Communications* 14.1 (2023), p. 5287.
- [4] Forough Habibollahi et al. “Biological Neurons vs Deep Reinforcement Learning: Sample efficiency in a simulated game-world”. In: *NeurIPS 2022 Workshop on Learning Meaningful Representations of Life*. 2022.
- [5] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. “Deepwalk: Online learning of social representations”. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2014, pp. 701–710.
- [6] Jian Tang et al. “Line: Large-scale information network embedding”. In: *Proceedings of the 24th international conference on world wide web*. 2015, pp. 1067–1077.
- [7] Moein Khajehnejad. “SimNet: Similarity-based network embeddings with mean commute time”. In: *PloS one* 14.8 (2019), e0221172.
- [8] Laurens Van der Maaten and Geoffrey Hinton. “Visualizing data using t-SNE.” In: *Journal of machine learning research* 9.11 (2008).
- [9] Joshua B Tenenbaum, Vin de Silva, and John C Langford. “A global geometric framework for nonlinear dimensionality reduction”. In: *science* 290.5500 (2000), pp. 2319–2323.
- [10] Alfonso Renart, Nicolas Brunel, and Xiao-Jing Wang. “Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks”. In: *Computational neuroscience: A comprehensive approach* (2004), pp. 431–490.
- [11] Emre Baspinar et al. “Coherence resonance in neuronal populations: Mean-field versus network model”. In: *Physical Review E* 103.3 (2021), p. 032308.
- [12] Christian Bick et al. “Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review”. In: *The Journal of Mathematical Neuroscience* 10.1 (2020), p. 9.
- [13] Giancarlo La Camera. “The mean field approach for populations of spiking neurons”. In: *Computational Modelling of the Brain: Modelling Approaches to Cells, Circuits and Networks*. Springer, 2021, pp. 125–157.
- [14] Jason Wolfe, Arthur R Houweling, and Michael Brecht. “Sparse and powerful cortical spikes”. In: *Current opinion in neurobiology* 20.3 (2010), pp. 306–312.
- [15] Bruno A Olshausen and David J Field. “Sparse coding of sensory inputs”. In: *Current opinion in neurobiology* 14.4 (2004), pp. 481–487.
- [16] Maria Elisabetta Ruaro, Paolo Bonifazi, and Vincent Torre. “Toward the neurocomputer: image processing and pattern recognition with neuronal cultures”. In: *IEEE Transactions on Biomedical Engineering* 52.3 (2005), pp. 371–383.
- [17] Brett J Kagan et al. “Scientific communication and the semantics of sentience”. In: *Neuron* 111.5 (2023), pp. 606–607.

A Appendix A.

A.1 DishBrain System

To assess the learning efficiency of cultured cortical networks during task engagement, we recorded neuronal cultures mounted onto a multi-electrode array with 1024 channels. The *DishBrain* system, interfacing in real-time with the MaxOne MEA (Maxwell Biosystems, AG, Switzerland) software, enables closed-loop stimulation and recording. In addition to recording neuronal electrical activity, this system also delivers safe, long-term external electrical stimulation using biphasic pulses [16] eliciting action potentials in neurons. Task-related information is relayed through appropriate coding schemes, allowing real-time monitoring of neuronal activity and simultaneous delivery of structured stimulation to the neuronal culture.

The *DishBrain* system was employed to simulate neural cultures within a virtual gaming environment, emulating the classic arcade game 'Pong'. Stimulation was delivered using a combination of rate coding electrical pulses, ranging from 4Hz to 40Hz, to encode the ball's position on the x -axis, and place coding (using designated electrodes arranged in a specific topographical fashion) to encode the ball's position on the y -axis. This input targeted a pre-defined two-dimensional sensory zone comprising 8 sensory electrodes. The motion of the paddle was controlled by the extent of electrophysiological activity measured within a predetermined "motor area" of the cultured network, captured in real-time. The cells also received information about the closed-loop response to their control of the paddle's movement. It was possible to either deliver the sensory stimulation, as explained above, or a feedback stimulation could be applied as previously described [1, 17]. The data utilized in this work were collected using an unpredictable feedback protocol. If the cultures failed to hit the ball using the paddle, indicating a "miss" event, they were subjected to an unpredictable stimulation. This feedback stimulus had a voltage of 150 mV and a frequency of 5 Hz, introducing an unpredictable external input into the system. Random stimulation was delivered to arbitrary locations on the 8 designated sensory electrodes, at varied intervals over four seconds. A configurable four-second resting period followed, where no stimulation was provided before the next rally began. Each gameplay session lasted 20 minutes, with a sampling frequency set at 20kHz.

A.2 Cell Culture

Neural cells were cultured either from the cortices of E15.5 mouse embryos or differentiated from human induced pluripotent stem cells via a dual SMAD inhibition (DSI) protocol or through a lentivirus based NGN2 direct differentiation protocols as previously described [1]. Cells were cultured until plating. For primary mouse neurons this occurred at day-in-vitro (DIV) 0, for DSI cultures this occurred at between DIV 30 - 33 depending on culture development, for NGN2 cultures this occurred at DIV 3.

A.3 MEA Setup and Plating

MaxOne Multielectrode Arrays (MEA; Maxwell Biosystems, AG, Switzerland) was used and is a high-resolution electrophysiology platform featuring 26,000 platinum electrodes arranged over an 8 mm². The MaxOne system is based on complementary meta-oxide-semiconductor (CMOS) technology and allows recording from up to 1024 channels. MEAs were coated with either polyethylenimine (PEI) in borate buffer for primary culture cells or Poly-D-Lysine for cells from an iPSC background before being coated with either 10 µg/ml mouse laminin or 10 µg/ml human 521 Laminin (Stemcell Technologies Australia, Melbourne, Australia) respectively to facilitate cell adhesion. Approximately 10⁶ cells were plated on MEA after preparation as per [1]. Cells were allowed approximately one hour to adhere to MEA surface before the well was flooded. The day after plating, cell culture media was changed for all culture types to BrainPhys™ Neuronal Medium (Stemcell Technologies Australia, Melbourne, Australia) supplemented with 1% penicillin-streptomycin. Cultures were maintained in a low O₂ incubator kept at 5% CO₂, 5% O₂, 36°C and 80% relative humidity. Every two days, half the media from each well was removed and replaced with fresh media. Media changes always occurred after all recording sessions.

A.4 DishBrain platform and electrode configuration

The current DishBrain platform is configured as a low-latency, real-time MEA control system with on-line spike detection and recording software. The DishBrain platform provides on-line spike detection and recording configured as a low-latency, real-time MEA control. The DishBrain software runs at 20 kHz and allows recording at an incredibly fine timescale. There is the option of recording spikes in binary files, and regardless of recording, they are counted over a period of 10 milliseconds (200 samples), at which point the game environment is provided with how many spikes are detected in each electrode in each predefined motor

region as described below. Based on which motor region the spikes occurred in, they are interpreted as motor activity, moving the 'paddle' up or down in the virtual space. As the ball moves around the play area at a fixed speed and bounces off the edge of the play area and the paddle, the pong game is also updated at every 10ms interval. Once the ball hits the edge of the play area behind the paddle, one rally of pong has come to an end. The game environment will instead determine which type of feedback to apply at the end of the rally: random, silent, or none. Feedback is also provided when the ball contacts the paddle under the standard stimulus condition. A 'stimulation sequencer' module tracks the location of the ball relative to the paddle during each rally and encodes it as stimulation to one of eight stimulation sites. Each time a sample is received from the MEA, the stimulation sequencer is updated 20,000 times a second, and after the previous lot of MEA commands has completed, it constructs a new sequence of MEA commands based on the information it has been configured to transmit based on both place codes and rate codes. The stimulations take the form of a short square bi-phasic pulse that is a positive voltage, then a negative voltage. This pulse sequence is read and applied to the electrode by a Digital to Analog Converter (or DAC) on the MEA. A real-time interactive version of the game visualiser is available at <https://spikestream.corticallabs.com/>. Alternatively, cells could be recorded at 'Rest' in a gameplay environment where activity was recorded to move the paddle but no stimulation was delivered, with corresponding outcomes still recorded. Using this spontaneous activity alone as a baseline, the gameplay characteristics of a culture were determined. Low level code for interacting with Maxwell API was written in C to minimize processing latencies-so packet processing latency was typically $<50 \mu s$. High-level code was written in Python, including configuration setups and general instructions for game settings. A 5 ms spike-to-stim latency was achieved, which was substantially due to MaxOne's inflexible hardware buffering. Figure S1 illustrates a schematic view of Software components and data flow in the DishBrain closed loop system.

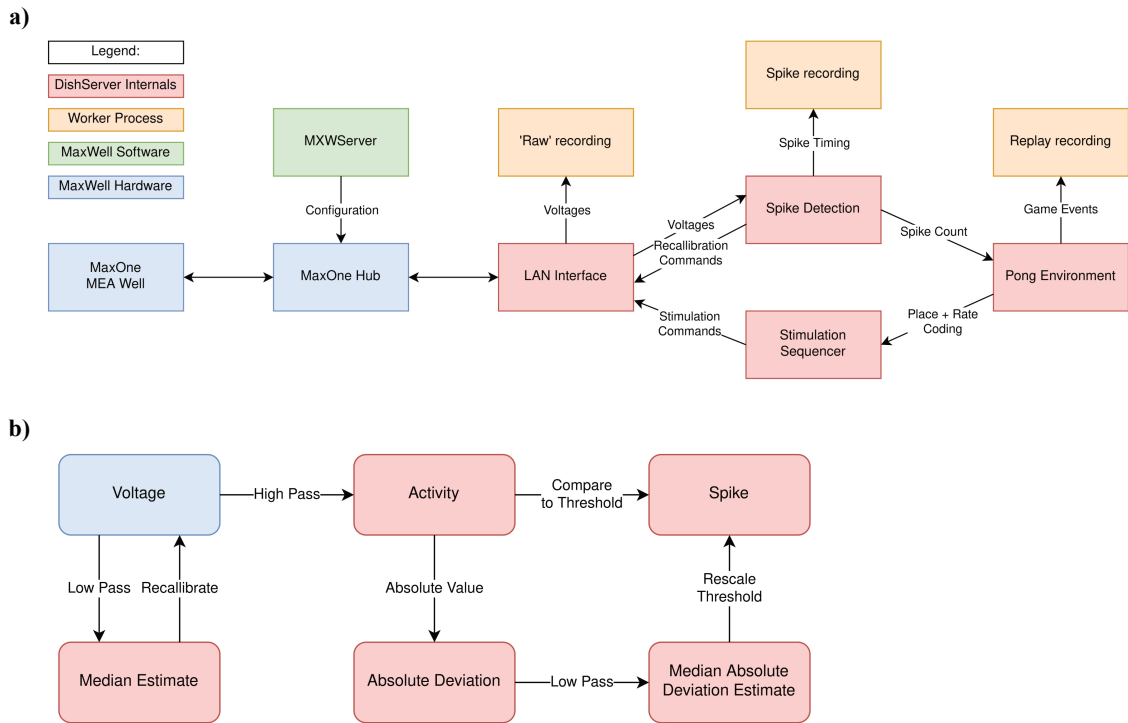


Figure S1: **a, b**) Schematics of software used for DishBrain. **a)** Software components and data flow in the DishBrain closed loop system. Voltage samples flow from the MEA to the ‘Pong’ environment, and sensory information flows from the ‘Pong’ environment back to the MEA, forming a closed loop. The blue rectangles mark proprietary pieces of hardware from MaxWell, including the MEA well which may contain a live culture of neurons. The green MXWServer is a piece of software provided by MaxWell which is used to configure the MEA and Hub, using a private API directly over the network. The red rectangles mark components of the ‘DishServer’ program, a high-performance program consisting of four components designed to run asynchronously, despite being run on a single CPU thread. The ‘LAN Interface’ component stores network state, for talking to the Hub, and produces arrays of voltage values for processing. Voltage values are passed to the ‘Spike Detection’ component, which stores feedback values and spike counts, and passes recalibration commands back to the LAN Interface. When the pong environment is ready to run, it updates the state of the paddle based on the spike counts, updates the state of the ball based on its velocity and collision conditions, and reconfigures the stimulation sequencer based on the relative position of the ball and current state of the game. The stimulation sequencer stores and updates indices and countdowns relating to the stimulations it must produce and converts these into commands each time the corresponding countdown reaches zero, which are finally passed back to the LAN Interface, to send to the MEA system, closing the loop. The procedures associated with each component are run one after the other in a simple loop control flow, but the ‘Pong’ environment only moves forward every 200th update, short-circuiting otherwise. Additionally, up to three worker processes are launched in parallel, depending on which parts of the system need to be recorded. They receive data from the main thread via shared memory and write it to file, allowing the main thread to continue processing data without having to hand control to the operating system and back again. **b)** Numeric operations in the real-time spike detection component of the DishBrain closed loop system, including multiple IIR filters. Running a virtual environment in a closed loop imposes strict performance requirements, and digital signal processing is the main bottleneck of this system, with close to 42 MB of data to process every second. Simple sequences of IIR digital filters is applied to incoming data, storing multiple arrays of 1024 feedback values in between each sample. First, spikes on the incoming data are detected by applying a high pass filter to determine the deviation of the activity, and comparing that to the MAD, which is itself calculated with a subsequent low pass filter. Then, a low pass filter is applied to the original data to determine whether the MEA hardware needs to be re-calibrated, affecting future samples. This system was able to keep up with the incoming data on a single thread of an Intel Core i7-8809G. Figures adapted from [1].