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

Moein Khajehnejad* MOEIN.KHAJEHNEJAD@MONASH.EDU
Turner Institute for Brain and Mental Health, Monash University, Clayton, VIC, AUS

Forough Habibollahi* FOROUGH@CORTICALLABS.COM
Alon Loeffler ALON@CORTICALLABS.COM
Brett J. Kagan BRETT@CORTICALLABS.COM
Cortical Labs, Melbourne, VIC, AUS

Adeel Razi ADEEL.RAZI@MONASH.EDU
Turner Institute for Brain and Mental Health, Monash University, Clayton, VIC, AUS
CIFAR Azrieli Global Scholars Program, CIFAR, Toronto, CA,
Wellcome Centre for Human Neuroimaging, University College London, UK

Editor: Sophia Sanborn, Christian Shewmake, Simone Azeglio, Nina Miolane

Abstract

In this study, we focus on characterizing the complex network dynamics of *in vitro* neuronal system of live biological cells during two distinct activity states: spontaneous rest state and engagement in a real-time (closed-loop) game environment. We use *DishBrain* which is a system that embodies *in vitro* neural networks with *in silico* computation using a high-density multi-electrode array. First, we embed the spiking activity of these channels in a lower-dimensional space using various representation learning methods. We then extract a subset of representative channels that are consistent across all of the neuronal preparations. Next, by analyzing these low-dimensional representations, we explore the patterns of macroscopic neuronal network dynamics during the learning process. 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 conditions. Furthermore, we characterise the evolving neuronal *connectivity* patterns within the *DishBrain* system over time during Gameplay in comparison to the Rest condition. Notably, our investigation shows dynamic changes in the overall connectivity within the same region and across multiple regions on the multi-electrode array only during Gameplay. These findings underscore the plasticity of these 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 or improved algorithms.

Keywords: In-vitro neuronal networks, DishBrain, Network dynamics, Low-dimensional representation

* Authors contributed equally to this manuscript.

© 2023 M. Khajehnejad, F. Habibollahi, A. Loeffler, B.J. Kagan & A. Razi.

1. Introduction

The *DishBrain* system introduces an innovative configuration that reveals elementary biological intelligence by leveraging adaptive and intrinsic neuronal traits. Within this construct, *in vitro* neuronal networks are seamlessly integrated with *in silico* computational elements utilizing high-density multi-electrode arrays (HD-MEAs). These cultivated neuronal ensembles demonstrate biologically-based adaptive intelligence, effectively emulated within a dynamic gaming environment through the implementation of closed-loop stimulation and concurrent recordings (Kagan et al., 2022, 2023). Specifically, the biological neural ensembles show self-organized adaptive electrophysiological dynamics, consistent with an innate ability for acquiring new knowledge and manifesting meaningful responses to constrained, albeit biologically plausible, external stimuli (Harrell et al., 2020). The empirical data is sourced from cortical cells derived from either embryonic rodent or human induced pluripotent stem cell (hiPSC) lineages. While synthetic biology methods demonstrate that *in vitro* biological networks of cortical cells are able to display real-time adaptive goal-directed learning in simulated environments (Kagan et al., 2022; Habibollahi et al., 2023, 2022), key underlying network dynamics associated with this learning remain unexplored.

Here, we study the spiking activity recorded from each channel of the HD-MEA to explore the dynamics of neuronal network structure and functional connectivity between the recorded units. Understanding the complex dynamics of neuronal networks is important to discover neural mechanisms of how learning occurs. Neurons in a complex network interact to process information and generate responses. Learning involves the modification of synaptic interactions, which affects the signal transmission within a neuronal network. Investigating the patterns of dynamic interactions between neurons provides insights into the mechanisms underlying learning. The temporal patterns and strength of these interactions represent the network's ability to encode, store, and retrieve information. By investigating the dynamic organisation of neural networks, we can discover the mechanisms driving synaptic modifications, leading to a deeper comprehension of the cellular and network-level processes intrinsic to learning. Such insights have crucial implications across disciplines, from neuroscience to artificial intelligence, potentially guiding the development of advanced learning algorithms and treatments for neurological disorders.

2. Methods

2.1. DishBrain System

To assess the learning efficiency of cultured cortical networks during task engagement, we recorded neuronal cultures mounted onto a multi-electrode array (MEA) 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 (Ruaro et al., 2005) 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 (Kagan et al., 2022). 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 recording session lasted 20 minutes, with a sampling frequency set at 20kHz. More details of this system are introduced in Appendix A

Figure 1 illustrates the input information, feedback loop setup, and electrode configurations in the DishBrain system.

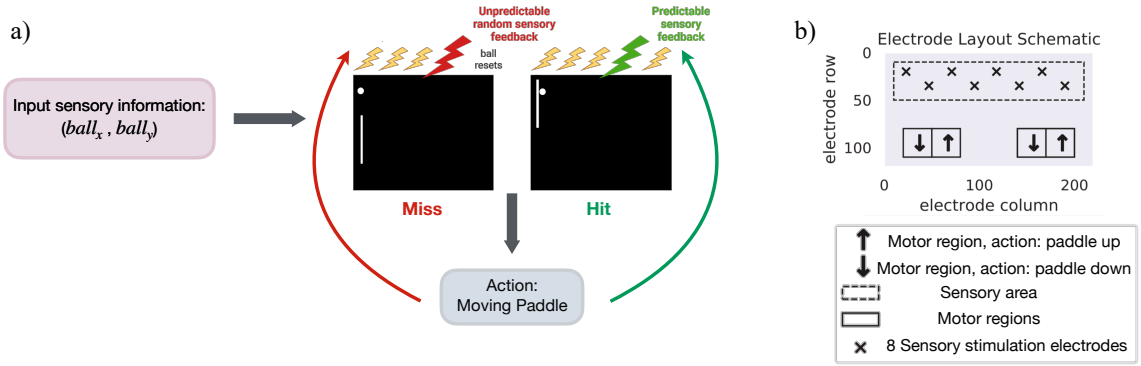

Figure 1: **a)** DishBrain feedback loop setup. **b)** Electrode configuration and predefined sensory and motor regions. Figures adapted and modified from (Kagan et al., 2022)

2.2. Network Construction

Neuronal spiking activity was measured from 1024 channels of HD-MEA in 248 Gameplay and 147 Rest session recordings. 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 (Perozzi et al., 2014; Tang et al., 2015; Khajehnejad, 2019). 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 (Van der Maaten and Hinton, 2008) and Isomap (Tenenbaum et al., 2000) to obtain 3-dimensional representations of both Rest and Gameplay recordings. To ascertain if these low-dimensional representations effectively capture latent network structures associated with learning during Gameplay, we divided all recording sessions in half prior to applying the dimensionality reduction techniques. Figure 2 illustrates the results after color labeling the first and second half of the recording sessions for three samples during Gameplay and Rest conditions, each of 10 and 5 minutes, respectively. The figure shows that while the two halves of recordings are easily distinguishable in the Gameplay sessions using either t-SNE or Isomap, the distinction is not as evident during the Rest sessions. This suggests the presence of specific patterns in the complex network dynamics during learning, captured only in the Gameplay.

Past studies have extensively utilized simplified models of interconnected neural populations based on mean-field approximations as an effective technique that maintains the dynamic properties of the neural network they are derived from, while also significantly accelerating simulation speeds by several orders of magnitude as well as allowing to study phase transitions (Renart et al., 2004; Baspinar et al., 2021; Bick et al., 2020; La Camera, 2021).

Additionally, within complex neural networks, only a fraction of neurons fire at any given time, while the majority do not show clear action potentials. Growing evidence suggests the emergence of specialized, selective, and abstract response properties within the cortex (Wolfe et al., 2010). Such sparse activity and connectivity patterns conserve energy and optimize computational capacity (Olshausen and Field, 2004). Yet, this sparsity highlights the redundancy in evaluating every neuron's individual firing patterns. The neuronal network's remarkable ability to encode and process information relies on the concerted action of neuronal populations, with individual neurons often conveying redundant or highly correlated signals.

Motivated by this emergent and collective behaviour of neuronal networks, we aimed to advance the reduction of computational complexity, when studying large neuronal populations while simultaneously preserving the dynamical properties of the network. To do so, we developed a method to identify a subset of recorded channels that likely monitored the neuronal populations specially attuned to the ongoing task. Such a subset allows for specification of neurons that characterise the network's behavior during Gameplay to more efficiently study the (macroscopic) dynamics of this smaller and interpretable network of neurons. To identify a consistent subset of channels across all neuronal cultures, we first used Tucker decomposition – via higher-order orthogonal iteration – on the tensor data from the 248 Gameplay sessions in the lower-dimensional embedding space. The resultant

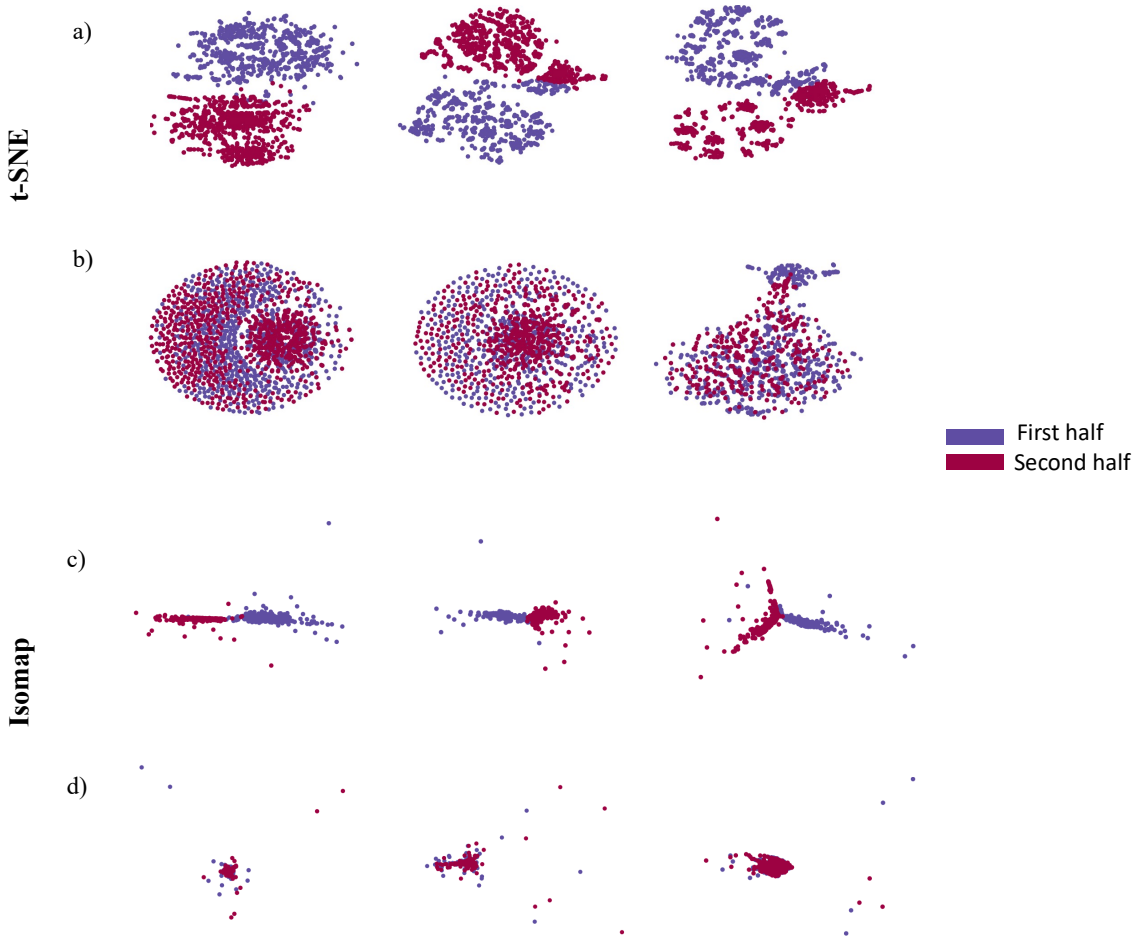

Figure 2: Low-dimensional representation of 3 samples of **a)** Gameplay sessions and their following **b)** Rest sessions using t-SNE as well as **c)** Gameplay sessions and their following **d)** Rest sessions using Isomap. The purple and maroon dots are the channel representations in the embedding space in the first and second halves of the recordings, respectively. Both dimensionality reduction algorithms were able to distinguish between the two halves of recording during Gameplay but not during Rest sessions.

tensor with dimensions of 1024×3 acts as the compact representation of the original data capturing its underlying patterns and structures. Using this tensor, we identified the set of representative channels by employing the K-medoid clustering algorithm to partition the data into $K = 30$ clusters and extract the corresponding 'medoids' of each cluster. Selecting $K > 30$ did not significantly improve the clustering accuracy measured by the red 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 3 is a schematic illustration of the proposed *in vitro* network construction framework in this study.

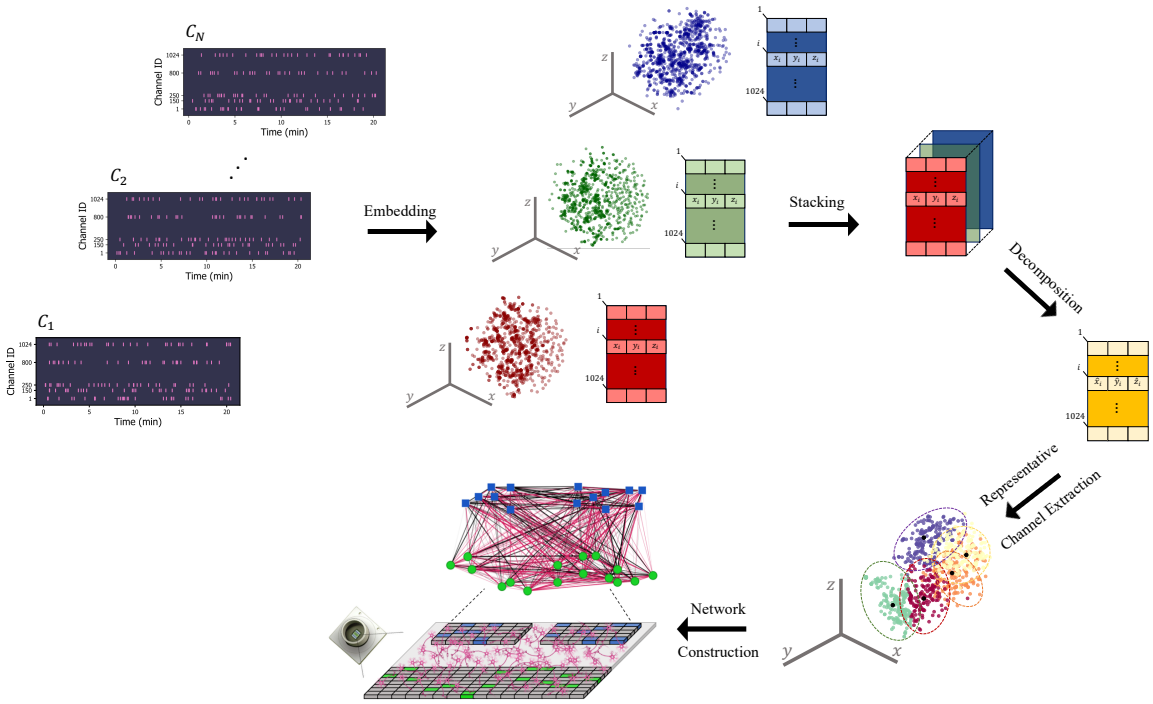

Figure 3: A schematic illustration of the overall network construction framework. The spiking time series of all channels from each recorded neuronal culture are first mapped to a 3-dimensional space using the t-SNE embedding algorithm. These lower-dimensional representations of all the recorded cultures are then stacked to form a tensor. Using tucker decomposition, a tensor is extracted and using K-medoids algorithm, the representative channels are identified which are consistent across all neuronal cultures. These channels are then used as the network nodes and the pairwise Pearson correlation values are used as edge weights for these networks. The network layout matches the actual location of the selected channels on the MEA and the node colors indicate whether these nodes belong to the predefined sensory (green) or motor (blue) regions.

3. Results

After constructing the connectivity networks corresponding to each recording session, we aimed to examine the temporal evolution of these networks in both Gameplay and Rest conditions. 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 4 shows the differences in the correlation between each pair of nodes when comparing the

last and first 2 minutes of each recording. This figure shows the average networks over all the Gameplay or Rest sessions with red and black colors indicating increased and decreased correlations, respectively. The edge weights are proportional to the absolute value of these differences in functional connectivity.

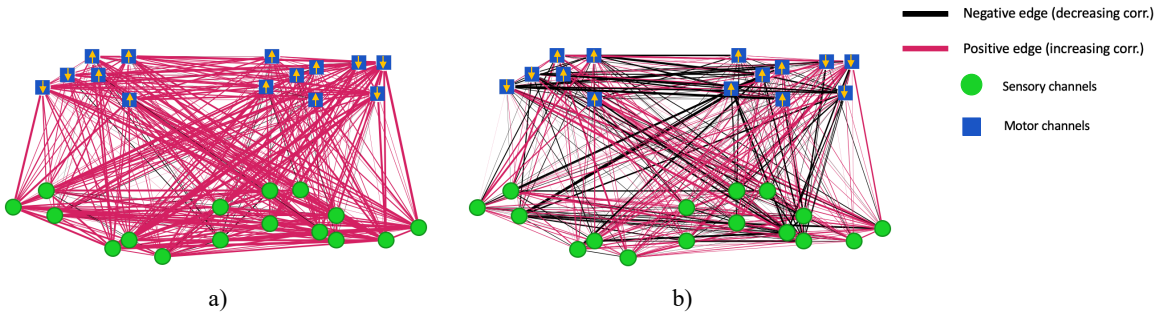

Figure 4: The average networks over all the **a)** Gameplay and **b)** Rest sessions with edge weights corresponding to the difference of correlation values, representing differences in functional connectivity, between each pair of channels when comparing the last 2 minutes to the first 2 minutes of the recordings. The edge colors indicate the valence of these differences in functional connectivity with red/black coloured edges denoting increased/decreased connectivity, respectively. Blue square and green circle nodes represent channels selected from motor and sensory regions, respectively. The arrows on the motor region nodes indicate the direction of paddle movement of which those channels are representative in the predefined layout shown in Figure 1.

As shown in Figure 4, 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 (Kagan et al., 2022). 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 5 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.

The increased average weight of the networks in the Gameplay sessions confirms the same patterns also observed in Figure 4. 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.

Figure 5: Comparing the network summary statistics between the first and last 2 minutes of recording for **top)** Gameplay and **bottom)** Rest. While all the evaluated network metrics show statistically significant differences during Gameplay, no statistically significant difference is detected during spontaneous rest state period (Rest). One-way ANOVA test, $^{***}p < 10^{-3}$. Box plots show interquartile range, with bars demonstrating 1.5X interquartile range, the line marks the median and the black triangle marks the mean. Error bands, 1 SE.

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 also build upon previous findings that have underscored the existence of redundancy within large populations of neurons, which contributes to the robustness of neural networks. It has been suggested that only a subpopulation of neurons is needed to capture the comprehensive dynamics of the entire neuronal network.

In our investigation, we work with a large population of *in vitro* cortical neurons, placing them within a closed-loop game environment in the DishBrain system. Our primary objective is to analyze the network dynamics of these neurons during learning, comparing their complex network dynamics in the experimental setting to conditions of Rest. To tackle this challenge, we introduce an innovative approach using DishBrain that allows us 1) to study neuronal populations on a cellular level in a closed-loop game environment, 2) to study neuronal activities in a lower-dimensional space and 2) to identify a specific subset of the population that can efficiently encapsulate the complete population dynamics, and 4) to study the underlying network dynamics leading to the emergence of information processing and learning.

Our research outcomes demonstrate the effectiveness of our framework in extracting this low-dimensional information, along with the identification of a small group of the most

influential recorded units. These findings convincingly indicate that our approach is highly efficient in capturing variations in the network's macroscopic properties. This represents a promising step toward a deeper understanding of how large neuronal populations function and adapt in complex environments and what could be the underlying network changes that give rise to learning.

References

Emre Baspinar, Leonhard Schülen, Simona Olmi, and Anna Zakharova. Coherence resonance in neuronal populations: Mean-field versus network model. *Physical Review E*, 103 (3):032308, 2021.

Christian Bick, Marc Goodfellow, Carlo R Laing, and Erik A Martens. Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review. *The Journal of Mathematical Neuroscience*, 10(1):9, 2020.

Forough Habibollahi, Moein Khajehnejad, Amitesh Gaurav, and Brett Joseph Kagan. Biological neurons vs deep reinforcement learning: Sample efficiency in a simulated game-world. In *NeurIPS 2022 Workshop on Learning Meaningful Representations of Life*, 2022.

Forough Habibollahi, Brett J Kagan, Anthony N Burkitt, and Chris French. Critical dynamics arise during structured information presentation within embodied in vitro neuronal networks. *Nature Communications*, 14(1):5287, 2023.

Evan R. Harrell, Matías A. Goldin, Brice Bathellier, and Daniel E. Shulz. An elaborate sweep-stick code in rat barrel cortex. 6(38):eabb7189, 2020. ISSN 2375-2548.

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, 2022.

Brett J Kagan, Adeel Razi, Anjali Bhat, Andy C Kitchen, Nhi T Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J Parker, Ben Rollo, and Karl J Friston. Scientific communication and the semantics of sentience. *Neuron*, 111(5):606–607, 2023.

Moein Khajehnejad. Simnet: Similarity-based network embeddings with mean commute time. *PloS one*, 14(8):e0221172, 2019.

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*, pages 125–157. Springer, 2021.

Bruno A Olshausen and David J Field. Sparse coding of sensory inputs. *Current opinion in neurobiology*, 14(4):481–487, 2004.

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*, pages 701–710, 2014.

Alfonso Renart, Nicolas Brunel, and Xiao-Jing Wang. Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. *Computational neuroscience: A comprehensive approach*, pages 431–490, 2004.

Maria Elisabetta Ruaro, Paolo Bonifazi, and Vincent Torre. Toward the neurocomputer: image processing and pattern recognition with neuronal cultures. *IEEE Transactions on Biomedical Engineering*, 52(3):371–383, 2005.

Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th international conference on world wide web*, pages 1067–1077, 2015.

Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Jason Wolfe, Arthur R Houweling, and Michael Brecht. Sparse and powerful cortical spikes. *Current opinion in neurobiology*, 20(3):306–312, 2010.

Appendix A. DishBrain System

A.1. 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 differentation protocols as previously described (Kagan et al., 2022). 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 culture development, for NGN2 cultures this occured at DIV 3.

A.2. 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 10^6 cells were plated on MEA after preparation as per (Kagan et al., 2022). 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 free media. Media changes always occurred after all recording sessions.

A.3. 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 visualiseris 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 μs. 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. Figures 6 and 7 illustrate schematic views of Software components and data flow in the DishBrain closed loop system.

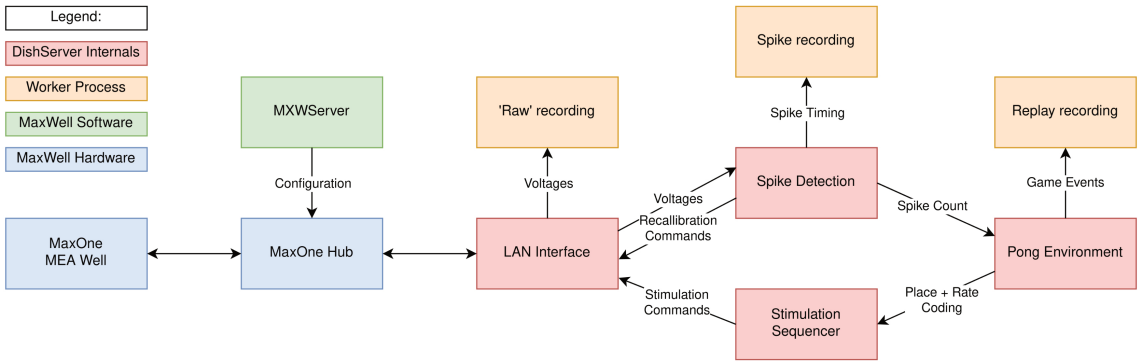

Figure 6: Software components and data flow in the DishBrain closed loop system. Voltage samples flow from the MEA to the 'Pong' environment, and sensory information flows from the 'Pong' environment back to the MEA, forming a closed loop. The blue rectangles mark proprietary pieces of hardware from 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. Figures adapted from (Kagan et al., 2022).

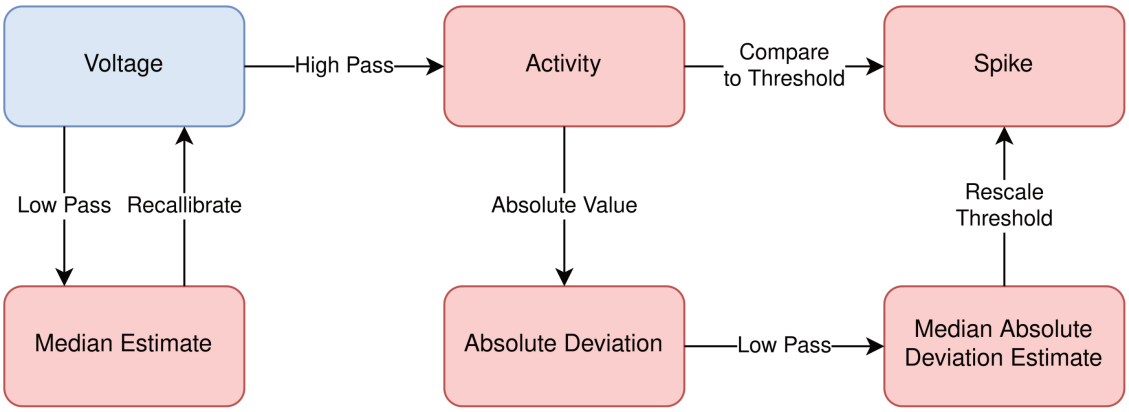

Figure 7: Numeric operations in the real-time spike detection component of the DishBrain closed loop system, including multiple IIR filters. Running a virtual environment in a closed loop imposes strict performance requirements, and digital signal processing is the main bottleneck of this system, with 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 (Kagan et al., 2022).