OPTIMIZING GENETICALLY-DRIVEN SYNAPTOGENESIS

Tommaso Boccato[1], **Matteo Ferrante**[1] **& Nicola Toschi**[1,2]
[1]Department of Biomedicine and Prevention, University of Rome Tor Vergata, Rome, Italy
[2]A.A. Martinos Center for Biomedical Imaging and Harvard Medical School, Boston, USA
tommaso.boccato@uniroma2.it

ABSTRACT

In this paper we introduce SynaptoGen, a novel framework that aims to bridge the gap between genetic manipulations and neuronal network behavior by simulating synaptogenesis and guiding the development of neuronal networks capable of solving predetermined computational tasks. Drawing inspiration from recent advancements in the field, we propose SynaptoGen as a bio-plausible approach to modeling synaptogenesis through differentiable functions. To validate SynaptoGen, we conduct a preliminary experiment using reinforcement learning as a benchmark learning framework, demonstrating its effectiveness in generating neuronal networks capable of solving the OpenAI Gym's Cart Pole task, compared to carefully designed baselines. The results highlight the potential of SynaptoGen to inspire further advancements in neuroscience and computational modeling, while also acknowledging the need for incorporating more realistic genetic rules and synaptic conductances in future research. Overall, SynaptoGen represents a promising avenue for exploring the intersection of genetics, neuroscience, and artificial intelligence.

1 INTRODUCTION

Let's imagine that, during brain development, **h1.** *we are able to manipulate, before synapses are formed, gene expression profiles of single neurons*. And let's imagine that **h2.** *we know how to act on these expression profiles in such a way as to guide synaptogenesis toward a specific neuronal network topology*. Maybe, **h3.** *we are also able to obtain the optimal computational graph, expressed as a composition of functions that represent the behaviour of neurons, required to solve a task of interest*. Small living organisms, or organoids (Bhaduri et al., 2020), could be, in principle, genetically programmed to fully develop with neuronal networks capable of solving pre-specified tasks. Such technology would lead to impressive applications – e.g., extreme low-power computing, micro-devices for the control of biological systems or therapies for disorders which are currently intractable. To date, hypothesis **h1.** seems to be verified (Nishikawa et al., 2014) while for **h3.** we can partially rely on spiking or non-spiking artificial neural networks and optimization techniques (Kingma & Ba, 2015; Graves, 2014).

In this work, we take a step toward the realization of the joint technology conceptualized in **h2.** and **h3.** by proposing SynaptoGen[1], a model that links, by means of differentiable functions, vector representations of gene expression profiles and genetic rules (i.e., interaction probabilities of protein pairs) with the average number of synaptic connections between pairs of neurons and their synaptic conductances. We substantiate our work through theoretical development which hinge on novel propositions and related mathematical proofs. SynaptoGen is compatible with backpropagation and can be inserted in learning frameworks where optimization is performed through gradient descent, enabling management of network sizes and task complexities beyond the capabilities of other optimization techniques. Finally, SynaptoGen is designed with flexibility in mind, allowing practitioners to choose which biological quantities to optimize (e.g., genetic rules, expression profiles or both).

[1]https://github.com/BoCtrl-C

Related Work. The resurgence of interest in organoid production for computational purposes gained momentum in 2022 with the unveiling of DishBrain (Kagan et al., 2022) by Cortical Labs. DishBrain is a system that integrates in-vitro neural networks derived from human or rodent sources with a simulated game world – "Pong" – through a high-density multielectrode array. The training approach employed by the DishBrain's authors follows the *free energy principle*, positing that neural networks learn to minimize the unpredictability of their sensory input by updating their beliefs and interacting with the environment. While the neuron cultures in DishBrain demonstrated statistically superior performance metrics in the game compared to controls defined by the authors, it remains challenging to assert that the neuronal networks fully mastered the assigned task.

In contrast to DishBrain, our research endeavors to "train" networks of neurons by influencing synaptogenesis through genetic manipulations at the individual neuron level. The foundational principles of our work trace back to (Barabási & Czégel, 2021; Barabási & Barabási, 2020), which introduce methods for constructing networks based on genetic encodings inspired by the wiring rules of the brain. These methods were further elaborated in the Connectome Model (CM) (Kovács et al., 2020), where the authors decomposed the adjacency matrix of a connectome into the product of three matrices representing specific genetic quantities. Another development was presented in (Barabási et al., 2023), where the CM's matrix entries were treated as learnable parameters, resulting in the weight matrix of a Multilayer Perceptron (MLP) within the context of training neural networks. While this methodology has proven effective in producing parameter-efficient neural networks, it maintains a notable distance from the biological intricacies of real neuronal networks. A distinct generalization of the CM has also been proposed for the computational inference of synaptic polarities (Harris et al., 2022).

Similarly, our work draws inspiration from the CM but is geared towards a more bio-plausible computational modeling of synaptogenesis, with the novel elements extensively discussed in Section 2.

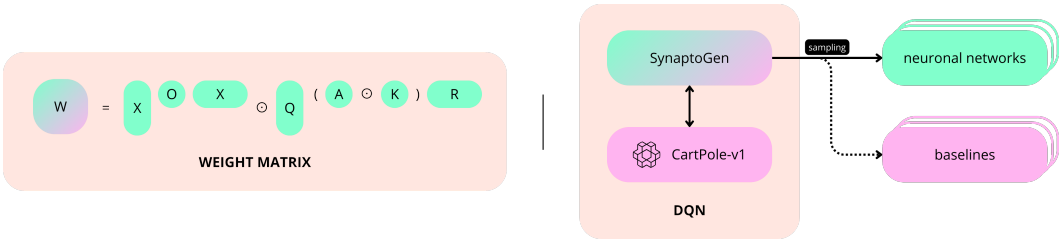

Figure 1: Overview of the SynaptoGen framework and related validation. **Left**: decomposition of the weighted connectome (W) introduced with the model's core equation. **Right**: validation performed on the Cart Pole reinforcement learning (RL) environment.

2 METHODS

In 2020, Kovács et al. proposed the CM, a novel strategy to link a brain's connectome (B) to the expression patterns of individual neurons (X) and existing biological mechanisms – or genetic rules – O:

$$B = XOX^T \tag{1}$$

In the CM's first interpretation, each row of X referred to a specific neuron while the i-th entry of the row described the binary expression (1 – "expressed" – or 0 – "not expressed") of gene i, one of the genes involved in synapse formation. Matrix O, instead, represented interaction compatibility for proteins translated from all gene pairs. Hence, $X \in \{0,1\}^{N\times G}$ and $O \in \{0,1\}^{G\times G}$ were defined as binary matrices while the entries of B, of shape $N \times N$, belonged to $\mathbb{Z}^+$; with N and G denoting the number of neurons and genes, respectively. When $G \ll N$ however, a very common scenario in nature (Koulakov et al., 2022), not all possible connectomes can be decomposed through equation 1. For this reason, the authors of the CM went on to relaxing the genetic rules matrix to $O \in [0,1]^{G\times G}$,

interpreting its entries as probabilities, and relying on the following approximation:

$$B \simeq XOX^T, \tag{2}$$

$$O = \arg\min_{O'} ||B - XO'X^T||^2 \tag{3}$$

where, in this context, $|| \bullet ||$ is intended as the Frobenius norm.

In this paper, we formulate a more general alternative to this framework. We build a model which takes into account synaptic conductances. While we start from equation 1, we design two novel interpretations tightly linked to the formalism with which the quantities of interest (i.e., the number of synapses between neurons and their conductances) have been represented. Our theoretical framework is as follows.

Let the number of synaptic connections between two neurons be represented by the following random variable:

$$\mathcal{B} = \sum_{i,j} \mathcal{B}^{ij} \tag{4}$$

where $\mathcal{B}^{ij}$ is a binomial random variable that expresses the contribution of the (i, j) gene pair to the total synaptic count:

$$\mathcal{B}^{ij} = \text{Bin}(n_{ij}, p_{ij}) \tag{5}$$

And let $\boldsymbol{x} \in \mathbb{R}^{+G}$ and $\boldsymbol{y} \in \mathbb{R}^{+G}$ be vector representations of gene expression in the pre- and post-synaptic neurons, respectively.

Proposition 1. *If the product between the i-th entry of $\boldsymbol{x}$ and the j-th entry of $\boldsymbol{y}$ denotes the number of independent experiments that characterizes $\mathcal{B}^{ij}$ – i.e., $x_i y_j = n_{ij}$ – and entry O_{ij} corresponds to probability p_{ij}, then the expected number of synapses between two neurons can be calculated as:*

$$\mathbb{E}[\mathcal{B}] = \boldsymbol{x}^T O \boldsymbol{y} \tag{6}$$

Proof. From probability theory,

$$\mathbb{E}[\mathcal{B}^{ij}] = n_{ij} p_{ij}$$

and due to the linearity of expectation we have

$$\mathbb{E}[\mathcal{B}] = \sum_{i,j} n_{ij} p_{ij}$$

On the other hand,

$$\begin{aligned} \boldsymbol{x}^T O \boldsymbol{y} &= \boldsymbol{x}^T [\dots, \sum_j y_j O_{ij}, \dots]^T \\ &= \sum_i x_i \sum_j y_j O_{ij} \\ &= \sum_{i,j} x_i y_j O_{ij} \end{aligned}$$

Recalling that $x_i y_j = n_{ij}$ and $O_{ij} = p_{ij}$, the proof is concluded. □

In different terms, if the hypotheses of Proposition 1 are verified, gene expression in a pair of genes tells us how many attempts we can make to place a synapse between a pre- and a post-synaptic neuron; the genetic rule, instead, describes the probability of success, conditioned on the interaction between the proteins translated from the considered genes, of each attempt. It is worth noting that equation 6 represents the average number of links between two specific nodes of the connectome. Keeping in mind that genetic rules are shared across neurons, the equation can easily be generalized to the whole connectome:

$$\bar{B} = \mathbb{E}[B] = XOX^T \tag{7}$$

where X has been obtained by stacking the expression profiles of all neurons (e.g., $X^T = [\dots, \boldsymbol{x}, \dots, \boldsymbol{y}, \dots]$).

In order to model synaptic conductances, instead, a slightly more complex formalism is required. We restrict ourselves to chemical synapses, which are the result of the interplay between neurotransmitters released by pre-synaptic neurons and receptors in post-synaptic neurons. According to recent studies, a chemical synapse can also have an excitatory or inhibitory effect depending on the nature of the receptor that receives a specific neurotransmitter (Fenyves et al., 2020; Harris et al., 2022). The way in which synapses work in our framework is described by the following equation:

$$I_v = \sum_u G_{uv} V_u \tag{8}$$

where I_v is the current injected into post-synaptic neuron v while V_u is an input voltage from pre-synaptic neuron u; G_{uv}, is the equivalent conductance that takes care of all the synapses formed between u and v. To model the possibility of having the mentioned synapses characterized by different neurotransmitter-receptor pairs, we rely again on random variables as follows.

Let $\mathcal{T}$ be a multinomial random variable representing the process of randomly picking, from u, a synaptic vesicle filled with a specific neurotransmitter. And let $\mathcal{R}$ be a multinomial random variable representing the process of randomly selecting a specific receptor from the membrane of v. We define vectors $\boldsymbol{q} \in [0,1]^L$ and $\boldsymbol{r} \in [0,1]^M$ as the probability distributions associated to $\mathcal{T}$ and $\mathcal{R}$; where L denotes the total number of neurotransmitters while M the number of receptors. We also define $A \in \{-1,0,1\}^{L \times M}$ as the polarity matrix (for further details, refer to Appendix A) and $K \in \mathbb{R}^{+L \times M}$ as the conductance matrix. In detail, entry A_{ij} tells us the polarity of synapses derived from the interaction of the i-th neurotransmitter with the j-th receptor ($A_{ij} = 0$ if the considered neurotransmitter and receptor are not compatible) while K_{ij} stores its linked conductance. We finally set $\mathcal{G} = f(\mathcal{T}, \mathcal{R})$, with $f(i,j) = A_{ij}K_{ij}$. In other words, $\mathcal{G}$ represents the "signed" conductance of a synapse randomly selected from the ones which connect neurons u and v.

Proposition 2. *If $\mathcal{T}$ and $\mathcal{R}$ are independent (i.e., the distribution of receptors in the post-synaptic neuron does not depend on the neurotransmitters synthesized by the pre-synaptic neuron), the expected "signed" conductance of a randomly picked synapse can be calculated as:*

$$\mathbb{E}[\mathcal{G}] = \boldsymbol{q}^T (A \odot K) \boldsymbol{r} \tag{9}$$

Proof. By expanding the matrix multiplications in equation 9, we have:

$$\begin{aligned} \boldsymbol{q}^T (A \odot K) \boldsymbol{r} &= \boldsymbol{q}^T [\dots, \sum_j r_j A_{ij} K_{ij}, \dots]^T \\ &= \sum_i q_i \sum_j r_j A_{ij} K_{ij} \\ &= \sum_{i,j} q_i r_j A_{ij} K_{ij} \end{aligned}$$

Thanks to the independence hypothesis:

$$\mathbb{P}[\mathcal{T} = i, \mathcal{R} = j] = q_i r_j$$

where $\mathbb{P}[\bullet]$ stands for "probability of". Hence,

$$\boldsymbol{q}^T (A \odot K) \boldsymbol{r} = \sum_{i,j} \mathbb{P}[\mathcal{T} = i, \mathcal{R} = j] f(i,j)$$

that corresponds exactly to the definition of $\mathbb{E}[\mathcal{G}]$. □

As for equation 7, also equation 9 can be generalized by stacking the neurotransmitter distributions of all pre-synaptic neurons in $Q = [\dots, \boldsymbol{q}, \dots]^T$ and the receptor distributions of post-synaptic neurons in $R = [\dots, \boldsymbol{r}, \dots]^T$:

$$\bar{G} = \mathbb{E}[G] = Q(A \odot K) R^T \tag{10}$$

As a next step, equation 7 and equation 10 can be inserted into the core equation of our model, which follows:

$$\begin{aligned} \bar{W} &= \bar{B} \odot \bar{G} \\ &= (XOX^T) \odot (Q(A \odot K) R^T) \end{aligned} \tag{11}$$

Summarizing, through equation 11 we are able to express the average equivalent conductance between all pairs of neurons as a differentiable function of their gene expression profiles and distributions of synthesized neurotransmitters and receptors, which, in turn, depend on gene expression. Furthermore, thanks to the adopted formalism, synaptogenesis can be simulated by sampling from the random variables introduced. For instance, the simplest approximation of synaptogenesis can be obtained as follows:

$$W = \tilde{B} \odot \bar{G} \tag{12}$$

with

$$\tilde{B} \sim B = \begin{bmatrix} \dots & \dots & \dots \\ \dots & \mathcal{B}_{uv} & \dots \\ \dots & \dots & \dots \end{bmatrix} \tag{13}$$

where $\sim$ stands for "sampled from".

3 EXPERIMENTS

To validate the proposed framework and assess its applicability in real-world scenarios, we conducted a preliminary experiment that involved simulating synaptogenesis in small populations of neurons (Figure 1). In simple terms, we simulated the formation of synapses in in-vitro neuron populations where genes and gene expression were manipulated at the level of individual neurons. These manipulations followed the genetic rules and expression profiles optimized by our model with the aim of enabling the resulting fully-developed neuronal networks to perform effectively in a pre-specified task.

Our approach began with a simplified setup, where a neuronal population comprised spatially-separated layers. We assumed that neurons in one layer could only attempt to form synapses with neurons in adjacent layers. This restriction allowed us to focus on multipartite network topologies and implement SynaptoGen as a customized MLP. The MLP's weight matrices have been indeed decomposed according to equation 11, with genetic rules shared across layers.

Regarding the task for our proof-of-concept experiments, we chose reinforcement learning (RL) as a bio-plausible benchmark. Our goal was to create a virtual neural agent capable of solving the control task defined by the `CartPole-v1` environment from the OpenAI Gym library (Brockman et al., 2016). In this environment, a pole is attached to a cart, which moves along a frictionless track. The objective is to balance the pole by applying forces to the left and right on the cart.

Initially, we trained[2] SynaptoGen (128 genes and 3 neurotransmitters), with a neuronal population size of $4 + 128 + 2$, on the Cart Pole environment (observation of dimension 4, 2 actions) using the DQN algorithm (Mnih et al., 2015). The training provided the genetic rules, gene expression profiles, neurotransmitter and receptor distributions, and synaptic conductances that enabled the agent to successfully solve the task with a mean reward of 500. Notably, this agent is built on the average equivalent conductances introduced in equation 10, and can be interpreted as an average agent which reflects the effects of the underlying genetically-derived quantities.

Hence, we sampled 10 neural networks from the trained SynaptoGen model, simulating the development, based on the computed synaptogenesis rules, of multiple populations of in-vitro neurons into neuronal networks. We then measured their performance and compared the obtained metrics with those from two carefully designed baselines. The first baseline (`baseline 1`) involved randomly initializing[3] all matrices of SynaptoGen, akin to letting neurons randomly express genes with translated proteins following random interaction rules. The second baseline (`baseline 2`) leveraged the SynaptoGen's biology, utilizing the optimized genetic rules O and conductance matrix K, while still initializing gene expression profiles, neurotransmitter and receptor distributions randomly. We sampled 10 neuronal networks from each baseline.

Results, depicted in Figure 2, show the reward distribution for each investigated model family. As expected, networks sampled from the baselines maintained the pole upright for approximately 10 time units, coinciding with the time taken for the pole to exceed the Cart Pole termination angle without effective control. Conversely, networks directly sampled from the trained SynaptoGen

[2]Hyperparameters from `https://github.com/DLR-RM/rl-baselines3-zoo/blob/master/hyperparams/dqn.yml`.

[3]Initialization from (Barabási et al., 2023).

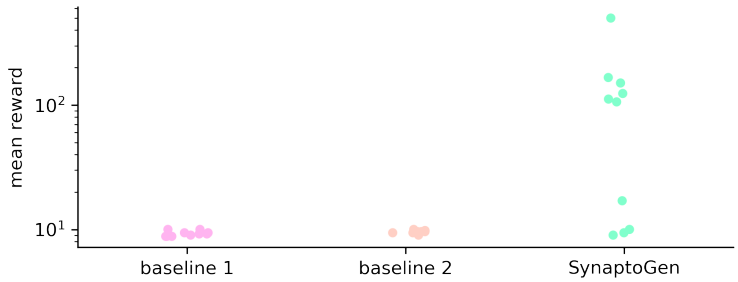

Figure 2: Mean reward distributions from the tested model families. Each *scatterplot* point represents the mean reward, averaged over 10 episodes, obtained by a specific agent. The y-axis is shown on a logarithmic scale.

exhibited significantly better performance, with the majority controlling the pole for a substantial duration. Remarkably, 1 out of 10 networks achieved perfect task resolution, maintaining the pole upright for the entire duration of the game (500 time units – Appendix B). This outcome suggests the successful generation of a fully functioning neuronal network after manipulating the genetics of just 10 neuronal populations.

4 Conclusions

In this paper we introduce SynaptoGen, a novel framework capable of simulating synaptogenesis and guiding a simulated neuronal population towards the formation of a neuronal network proficient in solving a predetermined computational task. Our approach involves framing synaptogenesis within the formalism of random variables and modeling their parameters through differentiable functions of matrices representing gene expression and protein interaction rules.

The validation of the framework was carried out by optimizing synapse formation in populations of $4 + 128 + 2$ neurons on a RL control task. During validation, the majority of networks exhibited a reduced loss of performance, with 1 out of 10 networks demonstrating an extraordinary absence of performance loss, perfectly solving the designated task.

Despite the success achieved in validation, it is crucial to acknowledge certain disparities with biological reality. Firstly, our model relies on a simplified conductance-based synapse model. Additionally, SynaptoGen has been implemented for multipartite networks consisting of standard artificial neurons. Regarding topology, this corresponds to genetically inhibiting, through an additional set of genes, all synapses that compatible neurons in non-adjacent layers could form, or employing an external posterior removal process.

Future plans involve extending the implementation of SynaptoGen to arbitrary feedforward networks by integrating its code with 4Ward (Boccato et al., 2024), a tool recently developed for converting arbitrary directed acyclic graphs into neural networks trainable with backpropagation. The neuron model used, instead, can be easily enhanced, for instance, by incorporating spiking neurons from the snnTorch library (Eshraghian et al., 2023) while maintaining overall compatibility with backpropagation.

It is important to note that our framework would benefit from a more realistic validation incorporating genetic rules that closely mirror those guiding synaptogenesis in nature. In this regard, also injecting synaptic conductances resulting from specific neurotransmitter-receptor interactions based on experimental data could be beneficial. Although data currently available for integration into our framework is limited, numerous research groups have made significant strides in developing methods aligning with this direction (Taylor et al., 2021; Kovács et al., 2020; Fenyves et al., 2020; Harris et al., 2022).

Despite the mentioned limitations, we firmly believe that SynaptoGen will serve as inspiration for novel methodologies and experiments, propelling the neuroscience community towards the creation of biological neuronal networks ready for deployment in a variety of cutting-edge applications.

ACKNOWLEDGMENTS

This work is supported and funded by: #NEXTGENERATIONEU (NGEU); the Ministry of University and Research (MUR); the National Recovery and Resilience Plan (NRRP); project MNESYS (PE0000006, to NT) - *A Multiscale integrated approach to the study of the nervous system in health and disease* (DN. 1553 11.10.2022); the MUR-PNRR M4C2I1.3 PE6 project PE00000019 Heal Italia (to NT); the NATIONAL CENTRE FOR HPC, BIG DATA AND QUANTUM COMPUTING, within the spoke *"Multiscale Modeling and Engineering Applications"* (to NT); the European Innovation Council (Project CROSSBRAIN - Grant Agreement 101070908, Project BRAINSTORM - Grant Agreement 101099355); the Horizon 2020 research and innovation Programme (Project EXPERIENCE - Grant Agreement 101017727). Tommaso Boccato is a PhD student enrolled in the National PhD in Artificial Intelligence, XXXVII cycle, course on Health and Life Sciences, organized by Università Campus Bio-Medico di Roma.

REFERENCES

Dániel L. Barabási and Albert-László Barabási. A genetic model of the connectome. *Neuron*, 105(3):435–445.e5, 2020. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2019.10.031. URL https://www.sciencedirect.com/science/article/pii/S0896627319309262.

Dániel L. Barabási and Dániel Czégel. Constructing graphs from genetic encodings. *Scientific Reports*, 11(1):13270, December 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-92577-2. URL http://www.nature.com/articles/s41598-021-92577-2.

Dániel L. Barabási, Taliesin Beynon, Ádám Katona, and Nicolas Perez-Nieves. Complex computation from developmental priors. *Nature Communications*, 14(1):2226, April 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-37980-1. URL https://www.nature.com/articles/s41467-023-37980-1.

Aparna Bhaduri, Madeline G. Andrews, Arnold R. Kriegstein, and Tomasz J. Nowakowski. Are organoids ready for prime time? *Cell Stem Cell*, 27(3):361–365, 2020. ISSN 1934-5909. doi: https://doi.org/10.1016/j.stem.2020.08.013. URL https://www.sciencedirect.com/science/article/pii/S1934590920304094.

Tommaso Boccato, Matteo Ferrante, Andrea Duggento, and Nicola Toschi. 4ward: A relayering strategy for efficient training of arbitrarily complex directed acyclic graphs. *Neurocomputing*, 568:127058, 2024. ISSN 0925-2312. doi: https://doi.org/10.1016/j.neucom.2023.127058. URL https://www.sciencedirect.com/science/article/pii/S0925231223011815.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

Jason K Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bennamoun, Doo Seok Jeong, and Wei D Lu. Training spiking neural networks using lessons from deep learning. *Proceedings of the IEEE*, 111(9):1016–1054, 2023.

Bánk G. Fenyves, Gábor S. Szilágyi, Zsolt Vassy, Csaba Sőti, and Peter Csermely. Synaptic polarity and sign-balance prediction using gene expression data in the caenorhabditis elegans chemical synapse neuronal connectome network. *PLOS Computational Biology*, 16(12):1–19, 12 2020. doi: 10.1371/journal.pcbi.1007974. URL https://doi.org/10.1371/journal.pcbi.1007974.

Alex Graves. Generating sequences with recurrent neural networks, 2014.

Michael R. Harris, Thomas P. Wytock, and István A. Kovács. Computational inference of synaptic polarities in neuronal networks. *Advanced Science*, 9(16):2104906, 2022. doi: https://doi.org/10.1002/advs.202104906. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/advs.202104906.

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, 2022. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2022.09.001. URL https://www.sciencedirect.com/science/article/pii/S0896627322008066.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Alexei Koulakov, Sergey Shuvaev, Divyansha Lachi, and Anthony Zador. Encoding innate ability through a genomic bottleneck. *bioRxiv*, 2022. doi: 10.1101/2021.03.16.435261. URL https://www.biorxiv.org/content/early/2022/05/26/2021.03.16.435261.

István A. Kovács, Dániel L. Barabási, and Albert-László Barabási. Uncovering the genetic blueprint of the ¡i¿c. elegans¡/i¿ nervous system. *Proceedings of the National Academy of Sciences*, 117 (52):33570–33577, 2020. doi: 10.1073/pnas.2009093117. URL https://www.pnas.org/doi/abs/10.1073/pnas.2009093117.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015. ISSN 1476-4687. doi: 10.1038/nature14236. URL https://doi.org/10.1038/nature14236.

Shin Nishikawa, Naohide Hirashima, and Masahiko Tanaka. Optimization of Single-Cell Electroporation Protocol for Forced Gene Expression in Primary Neuronal Cultures. *Molecular Biotechnology*, 56(9):824–832, September 2014. ISSN 1559-0305. doi: 10.1007/s12033-014-9761-1. URL https://doi.org/10.1007/s12033-014-9761-1.

Seth R. Taylor, Gabriel Santpere, Alexis Weinreb, Alec Barrett, Molly B. Reilly, Chuan Xu, Erdem Varol, Panos Oikonomou, Lori Glenwinkel, Rebecca McWhirter, Abigail Poff, Manasa Basavaraju, Ibnul Rafi, Eviatar Yemini, Steven J. Cook, Alexander Abrams, Berta Vidal, Cyril Cros, Saeed Tavazoie, Nenad Sestan, Marc Hammarlund, Oliver Hobert, and David M. Miller. Molecular topography of an entire nervous system. *Cell*, 184(16):4329–4347.e23, 2021. ISSN 0092-8674. doi: https://doi.org/10.1016/j.cell.2021.06.023. URL https://www.sciencedirect.com/science/article/pii/S0092867421007583.

A THE POLARITY MATRIX

As outlined in (Fenyves et al., 2020), synapse polarity in *C. elegans*, a well-studied small nematode, is elucidated by the interplay among 3 neurotransmitters – glutamate, acetylcholine, and GABA – and their corresponding receptors. Specifically, each neurotransmitter can be associated with receptors capable of exerting both excitatory and inhibitory effects on synaptic connections. This relationship can be represented, abstractly, through a $3 \times (2 \cdot 3)$ polarity matrix:

$$A = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix} \tag{14}$$

Here, each neurotransmitter synthesized in pre-synaptic neurons can be bind to either a positive (+) or negative (-) receptor in post-synaptic neurons. The 0s in the matrix signify that receptors attuned to a specific neurotransmitter are incapable of receiving different ones. This formalism readily extends to accommodate an arbitrary number of neurotransmitters by setting $M = 2L$ and expanding A to an $L \times 2L$ block diagonal matrix, where each block is represented as $[1, -1]$. It is worth noting that, in the experiments of Section 3, the entries of A do not belong to the set of learnable parameters.

B A CART POLE'S EPISODE

We show in Figure 3 some frames captured from an example episode played by the best agent generated with SynaptoGen. In the reported case, the agent obtained a final reward of 500, and the pole was kept in balance throughout the entire duration of the episode.

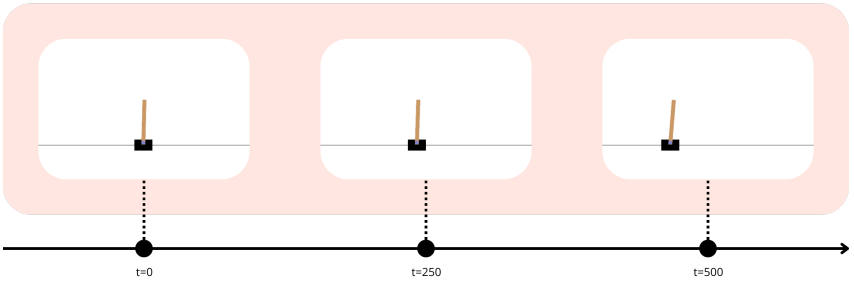

Figure 3: Frames captured from a Cart Pole episode. Frames are arranged in chronological order.