Exploration and Exploitation of Computational Capabilities in In-Vitro Biological Neural Networks

1st Peter Aaser
Department of Computer Science
Norwegian University Of Science And Technology
Trondheim
peter.aaser@ntnu.no

2nd Ola Huse Ramstad
Department of Medicine and Health Sciences
Norwegian University Of Science And Technology
Trondheim
ola.h.ramstad@ntnu.no

3rd Gunnar Tufte
Department of Computer Science
Norwegian University Of Science And Technology
Trondheim
gunnar.tufte@ntnu.no

3rd Ioanna Sandvig
Department of Medicine and Health Sciences
Norwegian University Of Science And Technology
Trondheim
ioanna.sandvig@ntnu.no

***Abstract*—The human brain is a vastly parallel computing machine, expending as little as 20 watt, capable of sustaining severe damage without major loss of function, and the ability to quickly adapt and respond to novel situations. However, even after a century's worth of scientific inquiry into the inner workings of the brain, we know remarkably little about how computation in the brain functions. Common for every part of the brain is that its functionality arises as individual neurons form connections among themselves, thus studying the underlying principles that form these networks and how they compute is a much more tractable prospect than tackling computation on a cognitive level. In this paper we describe the cultivation of biological neural networks originating from human stem-cells *in-vitro* in special *micro-electrode arrays*, which allows the cultures to be interacted with through the medium of electricity. Next we describe the *Reservoir Computing* (RC) approach to harnessing the undirected computational capability of biological neural networks. Lastly we describe a proof of concept a closed loop system has been developed where the sensory input of a simulated robot is used to stimulate the neurons, and the resulting dynamics used to steer the robot using Reservoir Computing.**

***Index Terms*—cyborg, self-organization, neurons, electro-physiology, reservoir computing, bio-inspired artificial intelligence**

I. Introduction

The digital computer has enjoyed an exponential increase in processing power over the last 50 years thanks to the bi-annual increase of transistor density, dubbed Moore's law [1]. During this exponential growth it looked like there was no task that computers wouldn't eventually solve by just waiting a few years for transistor density to catch up, a sentiment best summed up by a quote from Marvin Minsky, one of the early AI pioneers, made in 1967: “Within a generation [...] the problem of creating 'artificial intelligence' will substantially be solved.” [2] However, as processors have conquered the world of digital logic where facts can be neatly represented by numbers and binary relations, they are still remarkably unsuited at interacting with the real world. Even as design complexity and power consumption has effectively ended Moore's law, it is hard to imagine how an increase in FLOPS would make tasks such navigating complex environments, understanding natural language, or acting independently any easier. Not incidentally, these are the tasks that the brain is extraordinarily well suited for, however the underlying mechanisms that allow the brain to effortlessly do what machines cannot remain as enigmatic as ever. Part of the reason for this is that nature follows a completely different approach to engineering structures. Humans prefer to design contraptions consisting of heterogeneous components that act together in a precise and well-defined manner according to some blueprint. In contrast, the brain emerges from a set of basic principles in a process of self-organization, creating reliable, adaptive and complex structures from (relatively) homogeneous neurons. As a result, no two brains are similar, not even those of identical twins which greatly complicates any attempt at explaining how the brain works. For instance, a description of how visual perception works must hold for every brain in order to be useful, not just for a single individual, since such a description is of no use in a different individual where the process of self-organization have opted for a different path.

Understanding how neurons self-organize may hold key insights for many different fields. As an example, computer scientists could use the principles employed by neurons to create more robust systems capable of solving real world problems, while neuroscientists could better understand why the brain often fails to heal from certain injuries like strokes. Despite the wildly diverging goals of these two fields, and many more, they all have a common interest in understanding the fundamental properties of neurons. To forward this goal the authors have constructed a system which combines neural tissue with a simulated machine to create a platform for further study of neurons.

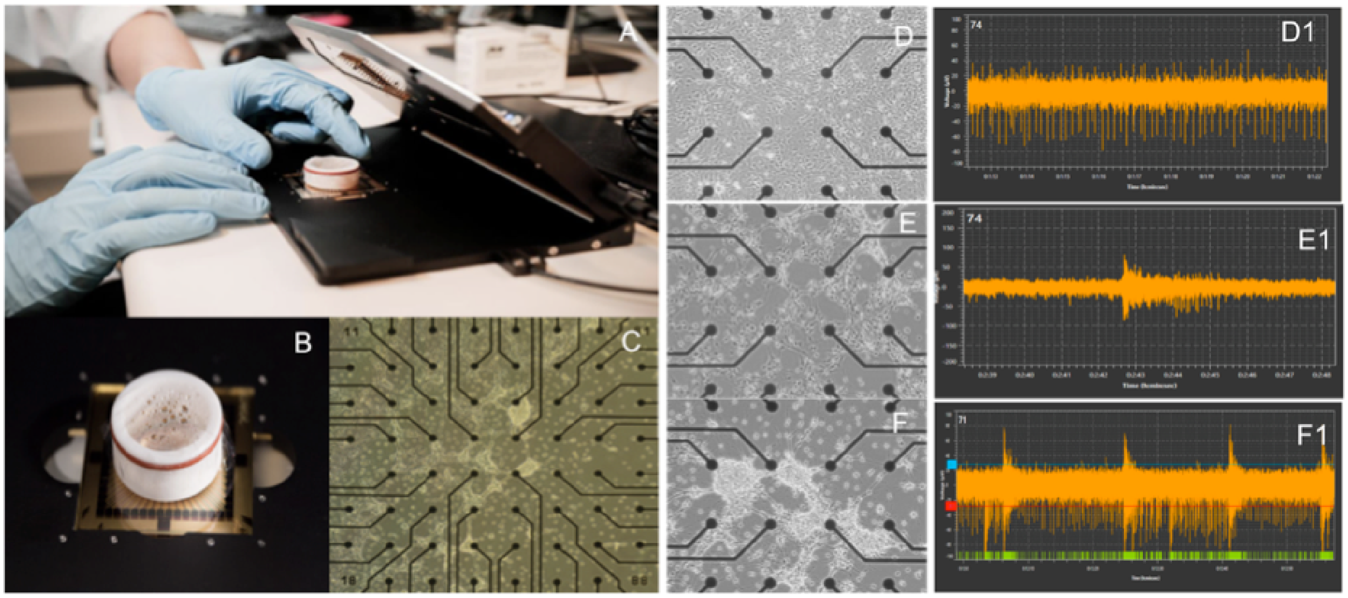

Fig. 1. A: The MEA2100 headstage with an MEA carrying a neural culture (B). C - F: Neural cultures growing on MEAs (B) viewed through a microscope. D1 - F1: Electrical activity (voltage) as a function of time measured on a single electrode from corresponding MEAs. D1 shows the electrical activity of a young network (D), whose spiking is frequent and seemingly random, whereas F1 shows more structured activity with large bursts typical of more mature networks.

A. Interacting With Neural Networks

In order to study the formation of structure and the emergence of function in neural networks it is necessary to be able to observe the physical structure and interact with the signaling process within the networks. This is achieved by growing neurons from stem cells in-vitro in *Micro Electrode Arrays* (MEA) as shown in figure 1 B. The MEA is lined with electrodes which can be viewed on a microscope as seen in figure 1 C, which enables real-time recording of electrical activity as seen in figure 1 D1 - F1. The measurements are done using a MEA2100 system supplied my multichannel systems GmbH, shown in 1 A.

B. A Proof Of Concept Cyborg

In order to provide a experiment-platform a proof of concept cyborg (cybernetic organism) has been implemented, based on the concept shown in figure 2 as part of an ongoing interdisciplinary project [3]. The proof of concept cyborg uses data from its sensors to stimulate a neural culture, interpreting the resulting *dynamics*, i.e the behaviour of the neural network to decide which direction it should move. This results in a cyborg that acts as a *closed loop system*, i.e one that operates without any outside (human or otherwise) correction or interference.

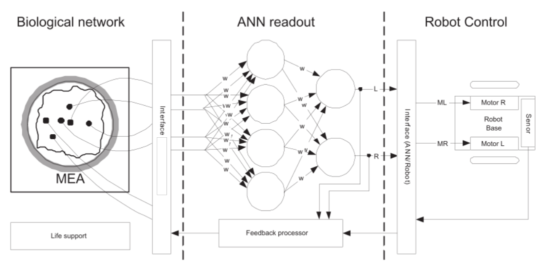

Fig. 2. The concept for the cyborg. Electrical activity from a neural network in an MEA is measured and used as input to an artificial neural network which steers a robot either left or right. The sensory data captured by the cyborg is transformed into an electrical signal which applied to the neural culture, effectively creating a feedback loop between the behavior of the robot and the behavior of the neural culture.

II. Cultivating Neural Cultures

Cultured networks consisted of dopaminergic (DA) neurons seeded at a density of 100 000 cells per MEA (Seen in 1 C). The DA neurons were derived from induced pluripotent stem cells (iPSCs) through exogenous addition of timed mesencephalic patterning factors over a period of 16 days allowing partial recapitulation of developmental processes. DA neuron morphology was confirmed by positive expression of molecular marker neuronal marker beta-III tubulin (Tuj-1), and tyrosine hydroxylase. MEAs were coated with poly-L-ornithine (PLO) and laminin. The cultures allowed to mature for 4 weeks prior to electrophysiological recording.

III. Interpreting Neural Dynamics

Harnessing the computational power of a neural network is no simple task. No two networks are the same, and the same network may exhibit different behavior as it grows and matures. As a result, attempting to understand the detailed workings of a single culture is nearly insurmountable, the chosen approach must hold for every single network. When approached from the framework of artificial neural network it may be tempting to attempt to shape the networks through Hebbian learning [4], unfortunately this approach does not scale beyond the ability to imprint very simple spiking behavior. Furthermore, attempting to impose desired functionality on a network is ignorant to the fact that the network has already been organized into a computing network by the neurons themselves, it does not need to be ”programmed” as a computer would. Instead we apply the *Reservoir Computing* (RC) where we view a neural network as an unsupervised classifier. The general principle of RC is shown in figure 3: In order to classify input is used to perturb a *reservoir*, and the resulting *dynamics* are classified using a linear classifier called a readout layer. An important detail is that only a small portion of the dynamics of the reservoir is measured, as sampling the state of the entire reservoir is impossible. In terms of neural cultures, the dynamics of the reservoir encompasses chemical gradients, structures and differences in voltage, however the *measured dynamics* is restricted to only 60 electrodes, greatly limiting what we can know of the reservoirs inner state. Although interacting with neural cultures using reservoir computing is currently unproven there are several reasons to suspect it to be a good fit: Firstly because it has been been successfully applied to digital reservoirs which have taken inspiration from living neural networks, such as Echo state networks [5] and recurrent neural networks [6]. Secondly, because it is has been shown to work on a physical system in the form of an analog circuit [7]

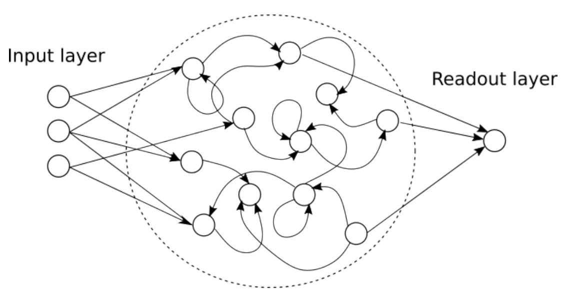

Fig. 3. A conceptual reservoir computing setup. The reservoir can be a physical or virtual system as long as its dynamics responds in a nonlinear fashion when it is perturbed. In order to approximates some function $f(x)y$, input (x) is translated to a perturbation suitable for the chosen reservoir, while the resulting dynamics is interpreted to the resulting value (y) by the readout layer.

A. Establishing A Common Language

Reservoir computing provides a nice way of circumventing the need for a deep understanding of the inner workings of the dynamics of the reservoir substrate [6], i.e. neural network culture. However, it is still necessary to translate electrical activity to a format that can be interpreted by the readout layer, and from sensory input in the robot to electrical stimuli. Figure 4 shows the first step of this process: Analog electrical activity recorded from the MEA is processed by a spike filter, reducing it to a digital spiking/not spiking representation. This representation is then further condensed to a vector representing the frequency of spikes for some interval of time which can then be fed to the readout layer. In the opposite direction, shown in figure 5, the distances measured by the robots sensors is transformed to a frequency where a short distance translates to a high frequency and vice versa. Stimulus is defined only as a frequency, i.e the period between application of a square pulse (although before an experiment parameters such as amplitude and duration for this pulse, or even its shape can be altered). These frequencies are then transformed into analog waveforms resembling spikes which are then applied at different electrodes. The feedback system is highly configurable, allowing for many different mappings between sensory input and resulting stiumulus, with the default configuration shown in figure 5 where the distance perceived by each ”eye” on the robot corresponds to the stimulus frequency on a given electrode.

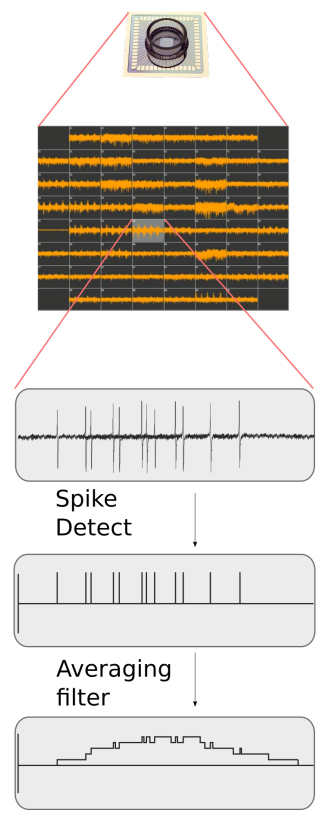

Fig. 4. Analog electrical activity is first transformed to a binary spiking/not spiking representation. Next the spikes are aggregated in an averaging filter, resulting in a vector which can be input to the readout layer.

B. Architecture

An overview of the final cyborg implementation is shown in figure 6. In the overview the data-flow between neural culture and robot control is shown, highlighting the fact that the robot control does not have to be located near the neurons since the information is transmitted over network. In the top

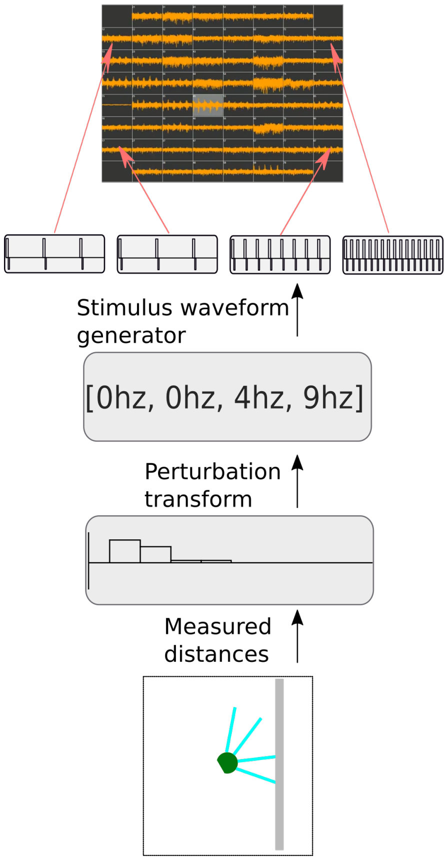

Fig. 5. Sensory data captured by the simulated robot is transformed into a set of frequencies where a short distance to a wall translates to a high frequency. These frequencies are then used to create analog waveforms resembling the spiking of neurons. The perturbation transform can be configured on a per-experiment basis, ranging from a binary on/off transform where seeing a wall elicits maximum allowed stimuli, to more sophisticated transforms such as exponential falloff.

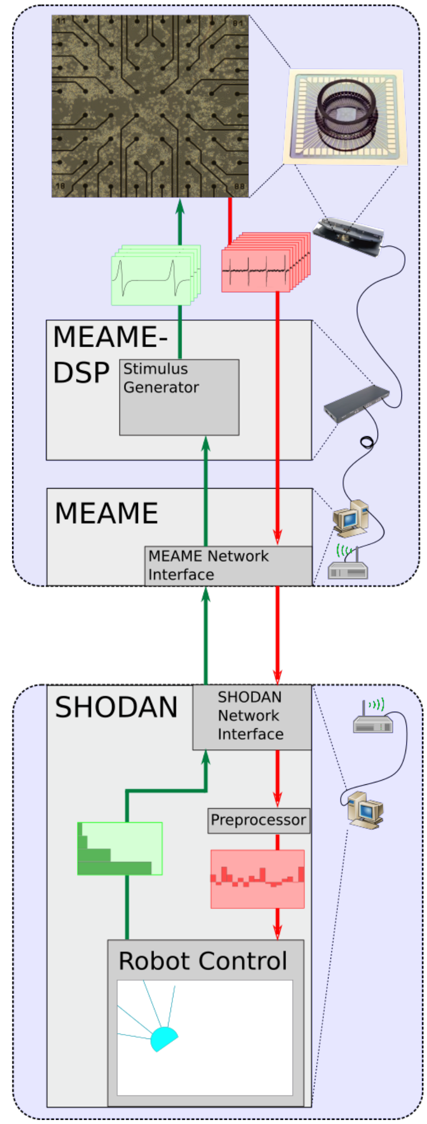

Fig. 6. A condensed overview of the final system. Data going from the neural culture to the robot control unit is colored red (also indicated with arrowheads for the colorblind) while sensory data going from the robot control unit to the neural culture is colored in green. The figure also highlights the shape of the data, showing how neural wave-forms from several electrodes are transformed to a vector of scalars as input for the robot control, and vice versa for stimulus. The flow of data going from the reservoir to the robot controller (red) corresponds to figure 4, while the flow going from robot control to the neural culture (green) corresponds to figure 5.

right part of the figure the MEA2100 lab equipment used to take measurements can be seen. An MEA holding a culture is inserted into the *Headstage* which can measure and apply stimuli at high precision. The headstage is connected to an interface board which hosts a user-programmable digital signal processor which is being used to apply periodic stimuli via the MEAME-DSP module. The interface board is connected to a lab-computer which hosts the MEAME module, responsible for broadcasting neural recordings in real time and hosting a stimulus request service. At the other side of the network gap the SHODAN system which consists of two main components: The mundane infrastructure parts, data filtering and such, and the more specialized robot controller which uses reservoir computing to control a simulated robot.

C. Robot Controller

Owing to the mercurial nature of neural cultures it is necessary for the reservoir computer to be able to reconfigure the readout layer during an experiment since there is no way to reset a neural culture to a previous state. As shown in figure 7 the robot is faced with several challenges and awarded a score based on how well it avoids the wall. The data acquired from each separate run is then used by an optimizing algorithm which runs concurrently with the experiment in order to optimize the readout layer to avoid colliding with the wall such that at the conclusion of each individual run the experiment runner can then fetch the best readout layer currently found.

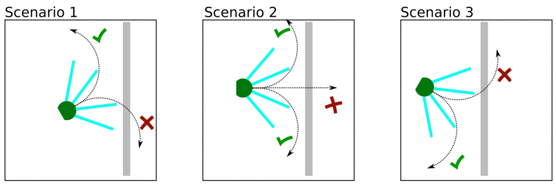

Fig. 7. The goal of the cyborg is to not collide into walls in a simple maze-environment. Four distance sensors (blue) radiate from the body (green) which allows the cyborg to sense the presence of an obstruction. Since the optimal strategy of avoiding walls is simply running in a circle, the faces scenarios where it is on collision course forcing it to respond.

IV. Outlook

The described system for interfacing neural in-vitro cultures to the digital domain can be a starting point for hybrid biological-digital systems. However, the project main purpose is to elucidate neural responses to perturbation, such as the computational efficiency of Parkinsonian or axotomized (severed) neural network compared to a healthy network. Further, the project aims to use bio-inspired methods toward introducing the underlying self-organizing process of learning into artificial systems.

Dynamics is the common property of neural culture's spiking behaviour and the reservoir's change of state in reservoir computation. Knowledge of how neural systems adapts it's dynamics by self organization include both exploration and stability this is features hardly found in artificial systems. Artificial Neural Networks (ANNs) are trained to one task and needs to be retrained to solve a new task, i.e. Catastrophic interference [8], instead of adapting to new information. By introducing similar dynamic adaption properties in reservoirs by mimicking the neural adaption by tuning physical resevoirs, e.g. nanomagnets [9] or carbon nanotubes [10], artificial systems that may learn (adapt) by self organization instead of being trained by an 2external" global system.

V. Conclusion

A complete system for a closed-loop cyborg has been implemented, serving as a platform for performing experiments on the computational capability of neural cultures. This entails a pipeline for broadcasting measured analog neural activity to a network, a flexible pipeline for processing analog neural measurements to digital spike data, an online re-configurable reservoir computing setup, and a pipeline for transforming digital problem representations to analog spiking as stimuli. Cultures specifically intended for use in the proof of concept cyborg are currently being grown in the lab, providing the final piece.

References

[1] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, 1965.

[2] H. L. Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press, 1992.

[3] P. Aaser, M. Knudsen, O. H. Ramstad, R. van de Wijdeven, S. Nichele, I. Sandvig, G. Tufte, U. S. Bauer, O. Halaas, S. Hendseth, A. Sandvig, and V. Valderhaug, "Towards making a cyborg: A closed-loop reservoir-neuro system," in *ECAL*, 2017.

[4] A. Odawara, H. Katoh, N. Matsuda, and I. Suzuki, "Induction of long-term potentiation and depression phenomena in human induced pluripotent stem cell-derived cortical neurons," *Biochem. Biophys. Res. Commun.*, vol. 469, no. 4, pp. 856–862, Jan 2016.

[5] W. Maass, T. Natschlager, and T. Markram, "Real-time computing without stable states: A new framework for neural computation based on perturbations," *Neural Computation*, vol. 14, no. 11, pp. 2531 – 2560, 2009.

[6] M. Lukosevicius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review*, vol. 3, no. 3, pp. 127 – 149, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1574013709000173

[7] J. H. Jensen and G. Tufte, "Reservoir computing with a chaotic circuit," *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, no. 29, pp. 222–229, 2017. [Online]. Available: https://www.mitpressjournals.org/doi/abs/10.1162/isal_a_039

[8] M. McCloskey and N. J. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," ser. Psychology of Learning and Motivation, G. H. Bower, Ed. Academic Press, 1989, vol. 24, pp. 109 – 165. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0079742108605368

[9] J. H. Jensen, E. Folven, and G. Tufte, "Computation in artificial spin ice," *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, no. 30, pp. 15–22, 2018. [Online]. Available: https://www.mitpressjournals.org/doi/abs/10.1162/isal_a_00011

[10] M. Dale, J. F. Miller, S. Stepney, and M. A. Trefzer, "Evolving carbon nanotube reservoir computers," in *Unconventional Computation and Natural Computation*, M. Amos and A. CONDON, Eds. Cham: Springer International Publishing, 2016, pp. 49–61.