



The Neurally Controlled Animat: Biological Brains Acting with Simulated Bodies

THOMAS B. DEMARSE, DANIEL A. WAGENAAR, AXEL W. BLAU AND STEVE M. POTTER

Division of Biology 156-29, California Institute of Technology, Pasadena, CA, 91125, USA

spotter@gg.caltech.edu

Abstract. The brain is perhaps the most advanced and robust computation system known. We are creating a method to study how information is processed and encoded in living cultured neuronal networks by interfacing them to a computer-generated animal, the Neurally-Controlled Animat, within a virtual world. Cortical neurons from rats are dissociated and cultured on a surface containing a grid of electrodes (multi-electrode arrays, or MEAs) capable of both recording and stimulating neural activity. Distributed patterns of neural activity are used to control the behavior of the Animat in a simulated environment. The computer acts as its sensory system providing electrical feedback to the network about the Animat's movement within its environment. Changes in the Animat's behavior due to interaction with its surroundings are studied in concert with the biological processes (e.g., neural plasticity) that produced those changes, to understand how information is processed and encoded within a living neural network. Thus, we have created a hybrid real-time processing engine and control system that consists of living, electronic, and simulated components. Eventually this approach may be applied to controlling robotic devices, or lead to better real-time silicon-based information processing and control algorithms that are fault tolerant and can repair themselves.

Keywords: MEA, multi-electrode arrays, rat cortex, prosthetics, hybrid system, cybernetics

1. Introduction

The brain is arguably the most robust computational platform in existence. It is able to process complex information quickly, is fault tolerant, and can adapt to noisy inputs. However, studying how the brain processes and encodes information is often difficult because access to it is limited by skin, skull, and the sheer number of cells. We are developing a new technique, the Neurally-Controlled Animat approach, that uses multi-electrode array (MEA) culture dishes (Thomas et al., 1972; Gross, 1979; Pine, 1980) to record neural activity from populations of neurons and use that activity to control an artificial animal, also called an Animat (Wilson, 1985, 1987) (Fig. 1). Because the neurons are grown on a clear substrate, MEAs allow detailed (submicron) optical investigation of neural connectivity (e.g., using 2-photon time-lapse microscopy; Potter, 2000; Potter et al., 2001) or submillisecond

distributed activity (e.g., using voltage-sensitive dyes and a high-speed camera; Potter, 2001; Potter et al., 1997). The Neurally Controlled Animat approach allows correlations to be made between neural morphology, connectivity, and distributed activity, not presently feasible with in vivo neural interfaces (Georgopoulos et al., 1986; Nicolelis et al., 1998; Wessberg et al., 2000). By allowing the neuronal network to trigger its own stimulation, our approach incorporates real-time feedback not present in previous cultured neural interfaces (Jimbo et al., 1999; Tateno and Jimbo, 1999).

2. Design and Methods

2.1. Neural-Computer Interface

Cortical tissue (from day 18 embryonic rats) is dissociated (Potter and DeMarse, 2001) and cultured on a




Figure 1. Scheme for the Neurally Controlled Animat. A network of hundreds or thousands of dissociated mammalian cortical cells (neurons and glia) are cultured on a transparent multi-electrode array. Their activity is recorded extracellularly to control the behavior of an artificial animal (the Animat) within a simulated environment. Sensory input to the Animat is translated into patterns of electrical stimuli sent back into the network.

60 channel multi-electrode array (Multichannel Systems) shown in Fig. 2. Each electrode can detect the extracellular activity (action potentials) of several nearby neurons and can stimulate activity by passing a voltage or current through the electrode and across nearby cell membranes (e.g., ± 600 mV $400 \mu\text{s}$, biphasic pulses). Dissociated neurons begin forming connections within a few hours in culture, and within a few days establish an elaborate and spontaneously active living neural network. After one month in culture, development of these networks becomes relatively stable (Gross et al., 1993; Kamioka et al., 1996; Watanabe et al., 1996) and is characterized by spontaneous bursts of activity. This activity was measured in real-time and used to produce movements within a virtual environment.

2.2. Controlling the Animat in its Virtual Environment

One advantage of our cultured network approach, compared to studies using intact animals, is that virtually any mapping is possible between neural activity and various functions. For example, one potential application might be to use neural activity as a control system to guide a robotic device, or to control a robotic limb (e.g., Chapin et al., 1999; Wessberg et al., 2000).

In this preliminary experiment, we created a virtual environment, a simple room, in which a living neural network could initiate movement of a simulated body where the direction of movement was based on the spatio-temporal patterns of activity across the MEA. The room consisted of four walls and contained barrier objects. The Animat can move forward, back, left or right within its virtual environment.

The neural activity on the MEA was recorded by amplifiers and A/D converters producing a real-time data stream of 2.95 MB per second (60 channels sampled at 25 KHZ). The computer analyzed the stream in real time to detect spikes (action potentials) produced by neurons firing near an electrode (Fig. 2). A clustering algorithm was trained to recognize spatio-temporal patterns in the spike train, to use as motor commands for the Animat, as follows: Activity on each channel was integrated and decayed following each spike, i , by:

$$A_n(t_i) = A_n(t_{i-1})e^{-\beta(t_i-t_{i-1})} + 1$$

where n is the MEA channel number from 1 to 60 (top left to bottom right in Fig. 2), t_i is the time of the current spike, t_{i-1} is the time of the previous spike on that channel n , β is a decay constant, $\beta = 1\text{s}^{-1}$. This produces a vector, \mathbf{A} , representing the current spatial

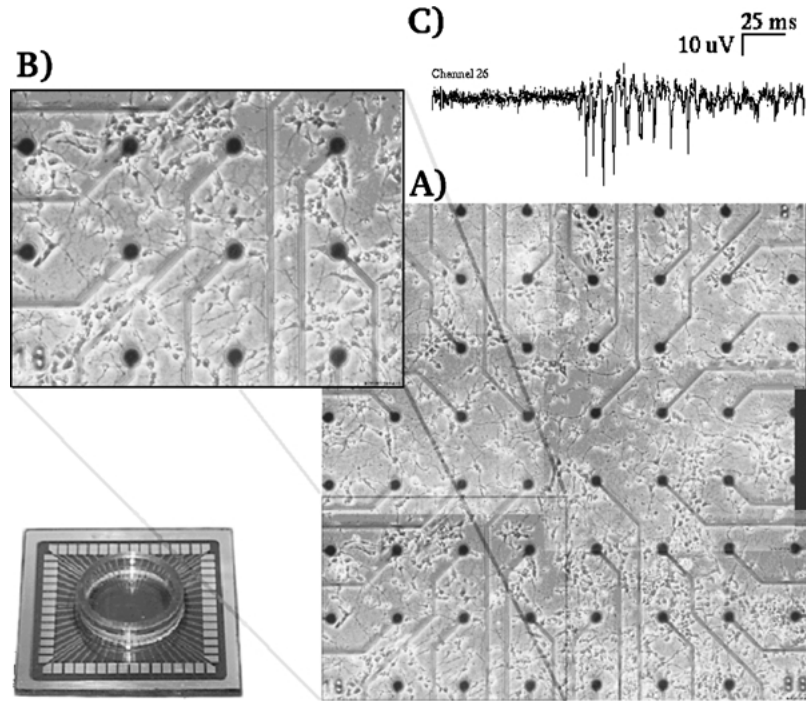


Figure 2. (A) A culture of live neurons after 3 days in culture on a 8×8 grid of 60- 10 μm -electrodes separated by 200 μm . (B) Expanded view of the lower left hand corner of the array showing the neurons and axons that form the living neural network. The large circular pads are the recording/stimulating electrodes that detect the extracellular neural activity of neurons located on or near each electrode. Insulated traces (dark lines in the figure) interface each electrode to amplifiers and A/D converters. MEAs and recording electronics are from MultiChannel Systems. (C) Raw signal from one electrode showing action potentials and background noise (signal to noise of 20 or more, typically). Spikes (action potentials) are detected in real-time as voltage peaks above a threshold of $5 \times$ standard deviation of the signal amplitude, calculated independently for each channel.

activity on the MEA. The activity vector \mathbf{A} was then sent through a squashing function:

$$P_n(t_i) = \tanh(\delta A_n(t_i))$$

where $\delta = 0.1$. This normalization was used to improve the dynamic range of the algorithm by limiting the contribution of channels with extremely high spike rates.

Every 200 ms, the computer sampled the current pattern of normalized activity \mathbf{P} , and clustered it according to the following: A pattern was grouped with the nearest cluster \mathbf{M}_k if the Euclidean distance to that cluster was less than a threshold Δ , where $\Delta = 0.9$. Clusters were allowed to shift their centers towards repeated matches by:

$$\mathbf{M}_k \Leftarrow \left(\frac{N_k}{N_k + 1} \right) \mathbf{M}_k + \left(\frac{\mathbf{P}}{N_k + 1} \right)$$

$$N_k \Leftarrow N_k + 1$$

where N_k is the number of occasions this pattern has matched so far, gradually freezing \mathbf{M}_k as N_k becomes large. If \mathbf{P} is not close to any cluster \mathbf{M}_k , a new cluster is constructed at position \mathbf{P} . A direction of movement within the virtual environment was then arbitrarily assigned to be triggered by that pattern. No movement was produced when the neural network was inactive by assigning \mathbf{M}_0 as the null cluster (all elements were set to zero). The choice of a 200 ms integrating window allowed the computer to process the raw data stream from 60 channels, perform online spike detection, integrate activity, and categorize the pattern sampled during that window. The decay constant allowed the pattern detection algorithm to detect emergent patterns (e.g., sequences of bursts) that extended beyond the window and reduced the artifact produced by sampling midway within bursts since activity was carried forward into the next window. In this manner, real-time spatio-temporal patterns of neural activity were categorized across the MEA and based on these patterns, a movement was produced within the virtual environment.

2.3. Sensory Information and Feedback

Proprioceptive feedback was provided for each movement within the virtual world as well as for the effects of those movements from collisions with walls or barriers. Feedback into the neuronal network was accomplished by inducing neural activity near one of five possible electrodes using custom hardware that delivered four ± 400 mV, 200 μ s pulses. Stimulated, or 'sensory' channels were chosen to be spatially distributed across the MEA, and capable of eliciting a reproducible response (action potentials) when stimulated. The stimulus strength was chosen to produce approximately half-maximal response from the network. Feedback stimuli typically occurred within 100 ms after pattern detection, often producing bursts that would propagate from the stimulating channel along multi-synaptic pathways. Similar stimulation pulses have been shown to produce changes in probability and latency to firing by neurons that are pathway-specific to the channel stimulated (Bi and Poo, 1999; Jimbo et al., 1999; Tateno and Jimbo, 1999). The Animat's movement resulted in feedback to one of the five assigned channels: either collision detection (one channel), or kinesthetic, (four channels, based on the four possible movements). Hence, each movement resulted in rapid feedback and that feedback could potentially change the activity of the network, and therefore, the behavior of the Animat.

3. The Activity of the Animat

Training of the neural pattern recognizer (clustering algorithm) began by interfacing a neural culture to the system for approximately 8 minutes without feedback, allowing the clustering algorithm to discover some of the recurring activity patterns. The feedback system was then activated and the Animat was run for approximately 50 minutes (but could have been run for hours, or perhaps days). During the time the Animat was online, it produced movements through much of its environment. The top left panel of Fig. 3 shows the trajectory of the Animat's steps within this artificial world. Movements of the Animat often occurred in rapid succession, in most cases as a result of periods where the culture exhibited bursts of neural activity.

Over the course of the run many different patterns of neural activity emerged. The bottom right panel of Fig. 3 shows the total number of patterns detected as the session progressed. Over the first few minutes the clustering algorithm quickly learned to recognize many

of the patterns of activity occurring in the MEA and, after 8 minutes, the number appears to stabilize. At this time the feedback system was activated and the number of patterns began to grow. The top right panel shows the frequency at which various patterns would recur throughout the session, in the order in which they were initially detected. Many of the patterns that were detected recurred frequently. For example, patterns 6, 9, 12, and 20 (shown in the lower left panel) appeared often while some patterns appeared only occasionally. Interestingly, these patterns would often occur in sequences of two or three that would repeat for a few seconds and then would reassemble into new combinations as the session progressed. The differences among these patterns, and the frequencies at which they occur, represent the current state of the living neural network as those patterns evolve either naturally, or as a result of the feedback we provided within the virtual environment.

4. Summary and Conclusions

The goal of the Animat project is to create a neurally-controlled artificial animal with which we can study learning in-vitro. This preliminary work has shown that it is possible to construct a system that can respond to and provide feedback in real-time to a living neural network. We do not yet know in detail how the complex patterns of activity were affected by the stimulation we provided, nor what changes within the network are responsible for producing the different patterns. Our current efforts are focused on clarifying these issues, as well as on understanding how these patterns are encoded within the network at a cellular level. We are also investigating how changes in the networks' activity patterns, whether spontaneous or as a result of Animat- or experimenter-initiated stimuli, can be mapped onto different Animat behaviors. Increased understanding of how feedback changes network activity, connectivity and cell morphology, might enable the development of more robust biologically-based artificial animals and control systems, and shed light on the neural codes within these networks. With this information we could create artificial animals as a control system to solve a wide variety of tasks, or map the neural processing power to perform calculations, pattern recognition, or process sensory input. Moreover, because the control system is biologically based, these artificial animals possess many potential advantages over their silicon counterparts. For example, they might exhibit both

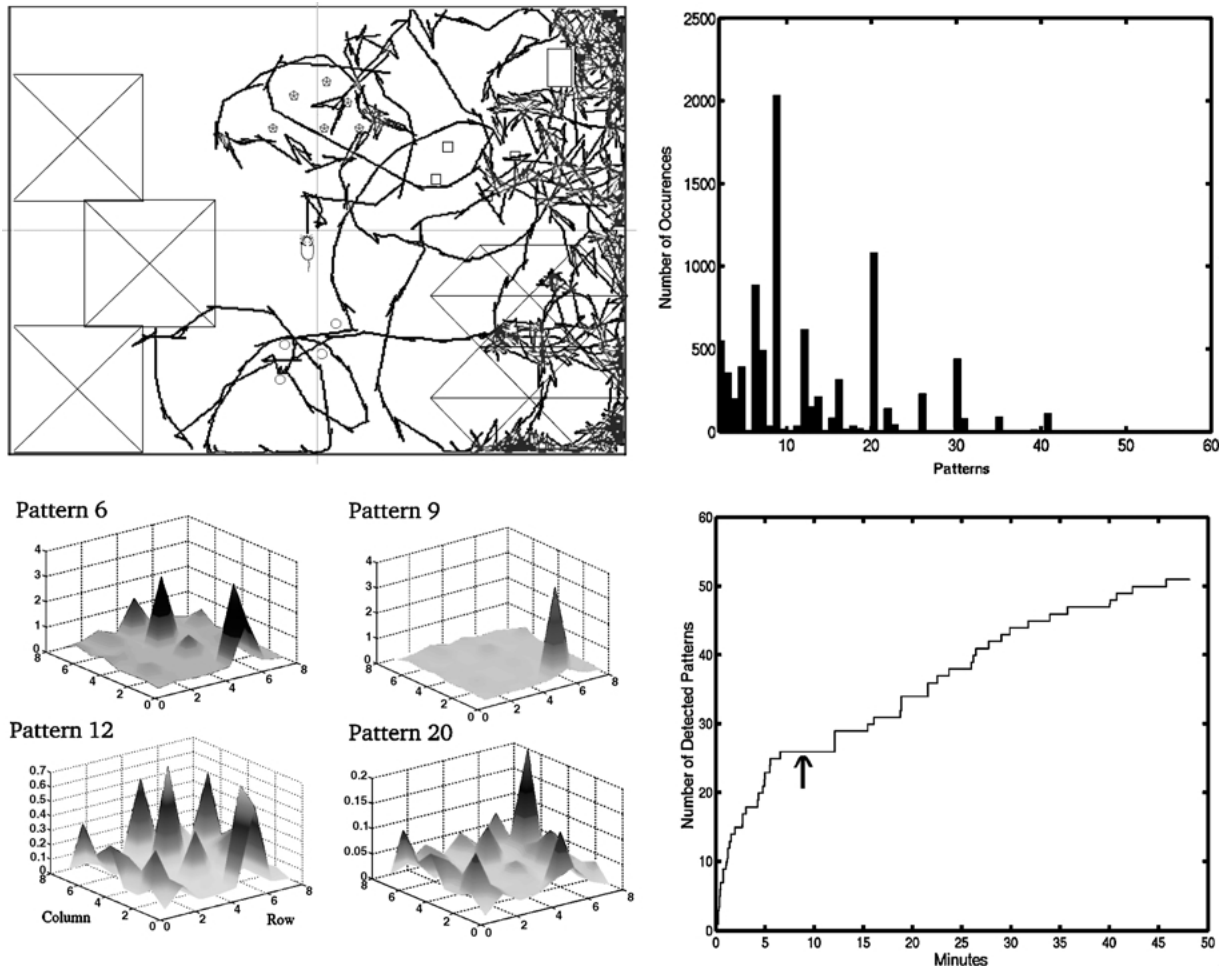


Figure 3. Movement and results of the session in which the Animat was online. Top left: Trajectory of movement of the Animat in a simulated room. Top right: The frequency that different patterns would recur. Bottom left: Examples of the four most common patterns. Bottom right: Number of novel patterns that the clustering algorithm detected as the session progressed (total 51). Bottom left: Examples of the four most common patterns.

complex and adaptive behavior as well as an intrinsic fault tolerance, i.e., the living network could potentially repair itself following damage. Studying how these changes occur as a result of feedback at the cellular and population levels will perhaps one day lead to more robust computing methods and control systems, and may provide information about how we process and encode information in the real world.

Acknowledgments

We thank C. Michael Atkin, Gray Rybka, and Samuel Thompson for early programming on the Animat and neural interface and MultiChannel Systems (<http://www.multichannelsystems.com>) for their

gracious technical support; Sami Barghshoon and Sheri McKinney for help with cell culture; Jerome Pine and Scott E. Fraser for support, advice, and infrastructure; and Mary Flowers, Shannan Boss, and Vanna Santoro for ordering and lab management. This research was supported by a grant from the National Institute of Neurological Disorders and Stroke, RO1 NS38628 (SMP) and by the Burroughs-Wellcome/Caltech Computational Molecular Biology fund (DAW).

References

Bi, G.Q. and Poo, M.M. 1999. Distributed synaptic modification in neural networks induced by patterned stimulation. *Letters to Nature*, 401:792–796.

- Chapin, J.K., Moxon, K.A., Markowitz, R.S., and Nicolelis, M.A.L. 1999. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2:664–670.
- Georgopoulos, A.P., Schwartz, A.B., and Kettner, R.E. 1986. Neuronal population coding of movement direction. *Science*, 233:1416–1419.
- Gross, G.W. 1979. Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multimicroelectrode surface. *IEEE Transactions on Biomedical Engineering*, 26:273–279.
- Gross, G.W., Rhoades, B.K., and Kowalski, J.K. 1993. Dynamics of burst patterns generated by monolayer networks in culture. *Neurobiology: An Interdisciplinary Approach to Substitute Impaired Functions of the Human Nervous System*, H.W. Bothe, M. Samii, and R. Eckmiller (Eds.), Amsterdam: North-Holland, pp. 89–121.
- Jimbo, Y., Tateno, T., and Robinson, H.P.C. 1999. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophysical Journal*, 76:670–678.
- Kamioka, H., Maeda, E., Jimbo, Y., Robinson, H.P.C., and Kawana, A. 1996. Spontaneous periodic synchronized bursting during the formation of mature patterns of connections in cortical neurons. *Neuroscience Letters*, 206:109–112.
- Nicolelis, M.A.L., Ghazanfar, A.A., Stambaugh, C.R., Oliveira, L.M.O., Laubach, M., Chapin, J.K., Nelson, R.J., and Kaas, J.H. 1998. Simultaneous encoding of tactile information by three primate cortical areas. *Nature Neuroscience*, 1:621–630.
- Pine, J. 1980. Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *Journal of Neuroscience Methods*, 2:19–31.
- Potter, S.M. 2000. Two-photon microscopy for 4D imaging of living neurons. In *Imaging Neurons: A Laboratory Manual*, R. Yuste, F. Lanni, and A. Konnerth (Eds.), CSHL Press: Cold Spring Harbor, pp. 20.1–20.16.A.
- Potter, S.M. 2001. Distributed processing in cultured neuronal networks. In *Progress in Brain Research: Advances in Neural Population Coding*, Nicolelis, M.A.L. (Ed.), Vol. 130, Amsterdam: Elsevier, pp. 49–62.
- Potter, S.M. and DeMarse, T.B. 2001. A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods*, 110:17–24.
- Potter, S.M., Lukina, N., Longmuir, K.J., and Wu, Y. 2001. Multi-site two-photon imaging of neurons on multi-electrode arrays. In *SPIE Proceedings*, 4262:104–110.
- Potter, S.M., Mart, A.N., and Pine, J. 1997. High-speed CCD movie camera with random pixel selection, for neurobiology research. In *SPIE Proceedings*, 2869:243–253.
- Tateno, T. and Jimbo, Y. 1999. Activity-dependent enhancement in the reliability of correlated spike timings in cultured cortical neurons. *Biological Cybernetics*, 80:45–55.
- Thomas, C.A., Springer, P.A., Loeb, G.E., Berwald-Netter, Y., and Okun, L.M. 1972. A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp. Cell Res.*, 74:61–66.
- Watanabe, S., Jimbo, Y., Kamioka, H., Kirino, Y., and Kawana, A. 1996. Development of low magnesium-induced spontaneous synchronized bursting and GABAergic modulation in cultured rat neocortical neurons. *Neuroscience Letters*, 210:41–44.
- Wessberg, J., Stambaugh, C.R., Kralik, J.D., Beck, P.D., Laubach, M., Chapin, J.K., Kim, J., Biggs, J., Srinivasan, M.A., and Nicolelis, M.A.L. 2000. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408:361–365.
- Wilson, S.W. 1985. Knowledge growth in an artificial animal. In *Proceedings of the First International Conference on Genetic Algorithms and their Applications*, Grefenstette (Ed.), Lawrence Erlbaum Associates: Hillsdale, NJ. pp. 16–23.
- Wilson, S.W. 1987. Classifier systems and the animat problem. *Machine Learning*, 2:199–228.