

Removing some ‘A’ from AI: Embodied Cultured Networks

Douglas J. Bakkum¹, Alexander C. Shkolnik², Guy Ben-Ary³, Phil Gamblen³, Thomas B. DeMarse⁴ and Steve M. Potter¹

¹Georgia Institute of Technology
bakkum@neuro.gatech.edu
steve.potter@bme.gatech.edu

²Massachusetts Institute of Technology
shkolnik@mit.edu

³University of Western Australia
guyba@cyllene.uwa.edu.au
pgamblen@hotmail.com

⁴University of Florida
tdemarse@bme.ufl.edu

Abstract. We embodied networks of cultured biological neurons in simulation and in robotics. This is a new research paradigm to study learning, memory, and information processing in real time: the Neurally-Controlled Animat. Neural activity was subject to detailed electrical and optical observation using multi-electrode arrays and microscopy in order to access the neural correlates of animat behavior. Neurobiology has given inspiration to AI since the advent of the perceptron and consequent artificial neural networks, developed using local properties of individual neurons. We wish to continue this trend by studying the network processing of ensembles of living neurons that lead to higher-level cognition and intelligent behavior.

1 Introduction

We present a new paradigm for studying the importance of interactions between an organism and its environment using a combination of biology and technology: embodying cultured living neurons via robotics. From this platform, explanations of the emergent neural network properties leading to cognition are sought through detailed optical and electrical observation of neural activity. A better understanding of the processes leading to biological cognition can, in turn, facilitate progress in understanding neural pathologies, designing neural prosthetics, and creating fundamentally different types of artificial intelligence. The Potter group is one of seven in the Laboratory for Neuroengineering (Neuro-lab¹) at the Georgia Institute of Technology, all working at the interface between neural

¹ <http://www.ece.gatech.edu/research/neuro/>

tissue and engineered systems. We envision a future in which mechanisms employed by brains to achieve intelligent behavior are also used in artificial systems; we overview three preliminary examples of the Neurally-Controlled Animats approach below. By using biology directly, we hope to remove some of the 'A' from AI.




Fig. 1. Connecting neurons to multi-electrode arrays. Left: Cells are plated inside a glass multi-electrode array culture dish such as this. Right: recorded voltage traces in the lighter boxes overlay a microscope image of the neuronal network growing on a 60-electrode array (electrode diameter, 30 μm). The thick lines are the electrode leads. The voltage spikes are neural signals.

No one would argue that environmental interaction, or embodiment, is unimportant in the wiring of the brain; no one is born with the innate ability to ride a bicycle or solve algebraic equations. Practice is needed. An individual's unique environmental interactions lead to a continuous 'experience-dependent' wiring of the brain [1]. This makes evolutionary sense as it is helpful to learn new abilities throughout life: if there are some advantageous features of an organism that can be attained through learning, then the ability to learn such features can be established through evolution (the Baldwin effect) [2]. Thus, the ability to learn is innate (learning usually being defined as the acquisition of novel behavior through experience [3]). We suggest that environmental interaction is needed to expose the underlying mechanisms for learning and intelligent behavior. Many researchers use *in vitro* models (brain slices or dissociated neural cell cultures) to study the basic mechanisms of neural plasticity underlying learning. We argue that because these systems are not embodied or situated, their applicability to learning *in vivo* is severely limited. We are developing systems to re-embody *in vitro* networks, and allow them to interact with an environment, so that we can watch the processes contributing to learning at the

cellular level *while they happen*.

We study networks of tens of thousands of brain cells in vitro (neurons and glia) on a scale of a few square millimeters. The cells in cortical tissue are separated using enzymes, and then cultured on a Petri dish with 60 electrodes embedded in the substrate, a multi-electrode array (MEA; from MultiChannel Systems) (Fig. 1) [4], [5]. The neurons in these cultures spontaneously branch out (Fig. 2). Even left to themselves without external input other than nutrients (cell culture media), they re-establish connections with their neighbors and begin communicating electrically and chemically within days, demonstrating an inherent goal to network; electrical and morphological observations suggest these cultures mature in about four weeks [6], [7], [8]. The neurons and supporting glia form a monolayer culture over the clear MEA substrate, amenable to optical imaging with conventional and two-photon microscopy [9], [10], [11]. With sub-micron resolution optical microscopy, we can observe learning-related changes in vitro with greater detail than is possible in living animals. The networks are also accessible to chemical or physical manipulation. We developed techniques to maintain neural cultures for up to two years, allowing for long-term continuous observation. For detailed methods, refer to [5].




Fig. 2. Microscope images of neurites (axons and dendrites) growing across a gap. The images were taken on three consecutive days beginning the second day after plating the cells. The *black circles* are the electrodes.

A multi-electrode array records extracellular neural signals fast enough to detect the firing of nearby neurons as voltage spikes (Fig. 1, right). Neurons detected by an electrode can be identified using spike-sorting algorithms [12]. Thus, the activity of multiple neurons can be observed in parallel and network phenomena can be studied. In addition to the expression of spontaneous activity, supplying electrical stimulation through the multiple electrodes induces neural activity; we have built custom circuitry to continuously stimulate the 60 electrodes [13]. The MEA forms a long-term non-destructive two-way interface to cultured neural tissue. The recorded signals can be used as motor commands, while the stimuli represent sensory inputs, in our embodied system. These techniques allow high resolution, long term, and continuous studies on the role of embodiment throughout the life of a cultured neural network.

Wilson [14] coined the term 'animat' (a computer simulated or robotic animal behaving in an environment) in his studies of intelligence in the interactions of artificial animals. Our interfacing of cultures to a simulated environment (described below) was the first Neurally-Controlled Animat (Fig. 3) [15], [16], [17]. For cultures interfaced to physical robots, we introduce the term 'hybrot' for hybrid biological robot. Mussa-Ivaldi's group created the first closed-loop hybrot by controlling a Khepera robot with a brain stem slice from a sea lamprey [18]. In a related approach, our Neurolab colleague Robert Butera studies detailed neural dynamics by coupling simulated neurons to real neurons using an artificial conductance circuit [19], [20]. Stephen DeWeerth's group in the Neurolab develops and studies, among others things, silicon model neurons interfaced with living mollusk and leech neurons [21].




Fig. 3. Hybrot (Hybrid living+robotic) setup. Optical and electrical data from neurons on an MEA are analyzed and used to control various robotic devices, while time-lapse imaging is carried out to make movies of neuronal plasticity.

Using simulated environments is a good first step and provides easier control and repeatability compared to robotics. However, a 'real' environment's great complexity provides two advantages. First, many seemingly complex behaviors of animals are emergent: simple behavioral rules applied in a complex environment produce complex and productive behavior [22], [23], [24]. Second, a complex environment produces a robust brain to take advantage of it: among other examples, this is evident in tool use [25] and in exploiting properties such as the biomechanics of muscle tissue in repositioning an arm without excessive vibrations. It is difficult to simulate a complex environment with realistic physics. If physics plays an important role in the complex behavior of intelligent systems, then by using robots in the real world, the researcher gets the physics "for free." We believe that this merging of artificial intelligence concepts (including robotics) into neurobiological experiments can inform future AI approaches, making AI a bit less artificial.

2 Examples: Three Embodied Neural Systems

Creating a neurally controlled robot that handles a specific task begins with a hypothesis of how information is encoded in the brain. Much remains to be determined, but numerous schemes have been proposed, most based on the quantity and/or relative timing of the firing of neural signals. A neural network may be considered as a type of processing unit with an input (synaptic or electrical stimulation patterns), and an output (neural firing patterns), which can perform interesting mappings to produce behavior. Below are overviews of three such systems. These examples could have been conducted with artificial neural networks. We use biological neural networks not as substitutes to artificial neural networks, but to tease out the intricacies of *biological* processing to inform future development of *artificial* processing. In particular, we analyzed how the properties of neurons lead to real-time control and adaptation to novel environments.

2.1 Living Neurons Control a Simulated Animal

The first Neurally-Controlled Animat [16] comprised a system for detecting spatio-temporal patterns of neural activity, which directed exploratory movement of a simulated animal in real time (Fig. 4). Neural firings were integrated over time to produce an activity vector every 200 ms, representing the current activity pattern, and recurring patterns were clustered in activity space. Each cluster was assigned a direction of movement (left, right, forward, backward). Proprioceptive and exteroceptive feedback via electrical stimulation was provided to the neural culture for each movement and for collisions with walls and barriers. The stimulation induced neural activity that, in turn, was detected through the activity vectors and used as commands for subsequent movements. We created the software and hardware necessary to enable a 15-ms sensory-motor feedback latency, since we feel it is important that a tight connection between the neural system and its environment is likely to be crucial to adaptive control and learning.

Within this real-time feedback loop, both spontaneous and stimulated neural activity patterns were observed. These patterns emerged over the course of the experiment, sometimes assembling into a recurrent sequence of patterns over several seconds, or the development of new patterns, as the system evolved. The overall effect of the feedback loop on neural activity was observed from the path of the animat's movement throughout its environment (Fig. 4). As the neural network moved its artificial body, it received feedback and in turn produced more movement. The behavioral output was a direct result of both spontaneous activity within the network as well as activity produced by feedback due to the networks interaction with its virtual environment. Hence the path of the animat was indicative of current activity as well as the effects of feedback. Analyzing the change in behavior of the neurally-controlled animat provided a simple behavioral tool to study shifts in the states of neural activity. However, this first Neurally-Controlled Animat did not

demonstrate noticeable goal-directed behavior, which the next example addresses explicitly.





Fig. 4. Animat setup and activity. Above: neural signals are used to control the movement of an animat, whose 'brain' is exposed to microscopic imaging; feedback from the environment determines subsequent electrical stimulation of the living neuronal network in an MEA. Below: One hour of the animat's path (*curved lines*), as it moves about within its environment under neural control, with feedback. The white boxes represent various environmental obstacles.

2.2 Living Neurons Control a Mobile Robot

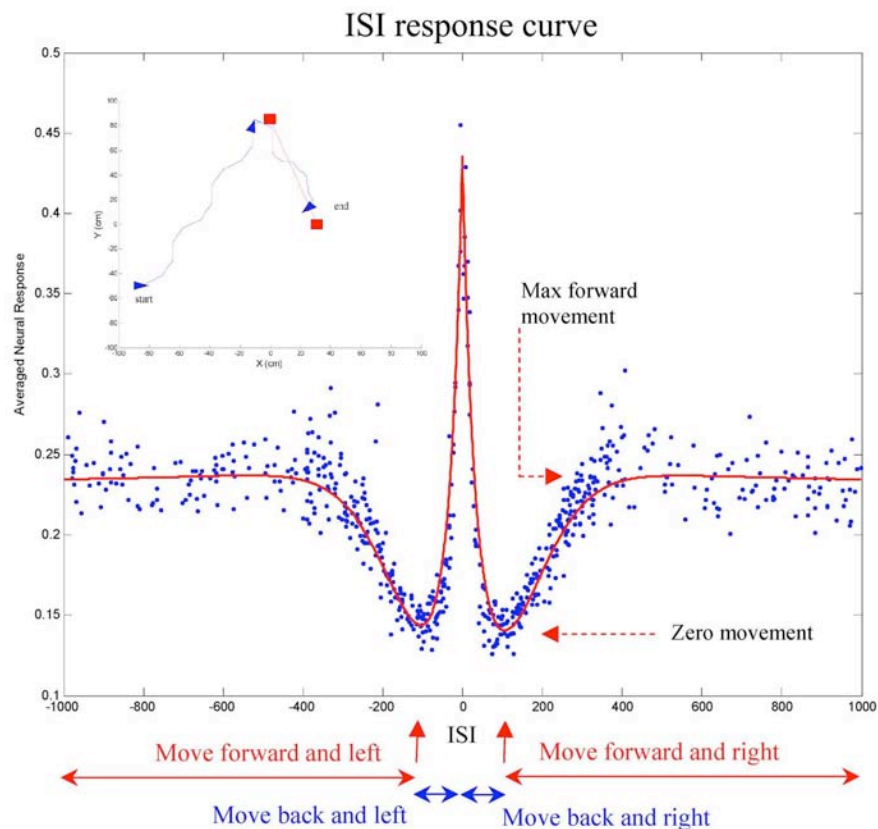


Fig. 5. Living neurons control a mobile robot. Neural firings in response to paired electrical stimulations at various inter-stimulus intervals (ISI) are plotted. In the experiments, the ISI was proportional to the distance between the neurally controlled approaching animat and its target object. It was considered positive if the target was located to the right of the animat and negative if left of the robot. The neural response determined the magnitude of subsequent animat movement; the direction of movement was determined from which quadrant the ISI fell into (see the arrows and movement key, bottom). Inset: the neurally controlled animat's trajectory (Koala robot, represented by the triangle). The target object (Khepera robot, represented by the square) was held stationary until the robot approached, and then it was moved continuously (down and to the right in the figure).

One of the simplest forms of 'intelligent' behavior is that of approach and avoidance. The goal of the second system was to create a neural interface between neuron and robot that

would approach a target object but not collide with it, maintaining a desired distance from the target. If a given neural reaction is repeatable with low variance, then the response may be used to control a robot to handle a specific task. Using one of these response properties, we created a system that could achieve the goal [26].

Networks stimulated with pairs of electrical stimuli applied at different electrodes reliably produce a nonlinear response, as a function of inter-stimulus interval (ISI). Figure 5 shows averaged firing rate over all 60 electrodes following two stimulations separated by a time interval. At short ISI's, the response of the network following stimulation was enhanced; at longer intervals, the response was depressed. Furthermore, the variance of the data for each ISI was relatively small, indicating the effect is robust and thus qualifies as a good candidate for an input/output mapping to perform computation.

By mapping the neural response to a given ISI as a transformation of distance to an object, we created a robot that reacts to environmental stimuli (in this case sensory information about distance from an object) by approaching and avoiding that target. To construct our "approach and follow" hybrot, sensory information (the location of a reference object with respect to the robot) was encoded in an ISI stimulation as follows: the closer the robot is to the object, the smaller the ISI. The response of the neurons to a stimulation pair, measured as an averaged firing rate across all electrodes for 100 ms after the second stimulus, was used to control the robot's movements: a larger neural response corresponded to a longer movement (either forward or backward) of the robot.

When the robot was far away from the reference object, the ISI of the stimulation pair was long, and the neural response was large, moving the robot towards the object (Fig. 5, right). As the robot moved closer to the object, the stimulation interval decreased until it reached 150 ms. At this point, the neural response was minimal, and no movement was commanded. In other words, the robot reached its desired location with respect to the reference object. If the robot was closer to the object, the neural reaction was larger (a very short ISI), this time driving the robot away from the object. We divided the input ISI into 4 quadrants (Fig. 5, left). Each of the 4 quadrants corresponded to a directional movement: forward/right, forward/left, backward/right, and backward/left. Then, a positive ISI caused movement in a direction opposite that for a negative ISI. Given the neural response to an ISI stimulation, we decoded which quadrant the response belonged to with good accuracy (>95%).

We used the Koala and Khepera robots (manufactured by K-Team) to embody the cultured network, and to provide an environment with a moving object. The Koala robot was used as the neurally controlled robot, while the Khepera served as the reference object, moving at random under computer control. Under neural control, the Koala successfully approached the Khepera and maintained a distance from it, moving forward if the Khepera moved away, or backing up if the Khepera approached.

In addition to demonstrating the computational capacity inherent in cultured neurons, this hybrot can be used to study learning in cultured neural networks. In this case, learning would be manifested through changes in the neural activity and changes at the behavioral level of the robot. Preliminary studies indicate that quantifiable behavioral traits,

such as the speed with which the hybrot approaches the object, may be manipulated through mechanisms of neural plasticity.

2.3 Living Neurons Control a Drawing Arm




Fig. 6. Meart—The Semi-Living Artist. Left: Meart’s arms used markers to draw on a piece of paper, under live neural control. In the background was a projection of the MEA and cultured net, Meart’s ‘brain’. Right: one drawing created by Meart in an exhibition.

Meart (Multi-Electrode Array art) was a hybrot born from collaboration with the SymbioticA Research Group². The ‘brain’ of dissociated rat neurons in culture was grown on an MEA in our lab in Atlanta while the geographically detached ‘body’ resided in Perth. The body consisted of pneumatically actuated robotic arms moving pens on a piece of paper (Fig. 6). A camera located above the workspace captured the progress of drawings created by the neurally-controlled movement of the arms. The visual data then instructed stimulation frequencies for the 60 electrodes on the MEA. The brain and body interacted through the internet (TCP/IP) in real time providing closed loop communication for a neurally controlled ‘semi-living artist’. We see this as a medium from which to address various scientific, philosophical, and artistic questions.

Meart has brought neurobiology research to two artistic events: *Biennale of Electronic Arts Perth* and most recently at *Artbots: the Robot Talent Show* in New York. The robotic arm and video sensors were shipped to New York while the living neurons sent and re-

² SymbioticA: the Art and Science Collaborative Research Laboratory (<http://www.fishandchips.uwa.edu.au/>), based in the School of Anatomy and Human Biology at the University of Western Australia in Perth.

ceived signals from Atlanta. An overview of how Meart worked may best be described by the artistic conception behind the Artbots presentation: portrait drawing. First, a blank piece of paper was placed beneath the arm's end-effector and a digital photograph was taken of an audience member. Then, communication between the arm and the neurons was begun. The neural stimulation via the MEA was determined by a comparison of the actual drawing, found using a video camera taking images of the drawing paper, to the target image of a person's photograph. Both the actual image and the target image were reduced to 60 pixels, corresponding to the MEA electrodes, and the gray scale intensity of each pixel was found. Similar to how an artist continually compares her work to her subject, the gray scale percentages for corresponding pixels on the two images were continuously compared, in this case subtracted to produce a matrix of error values. The 60 error values determined in real-time the stimulation frequency per electrode using a custom stimulation circuit built by Thomas DeMarse. Arm movement was determined by the recorded neural activity, using averaged firing rates of the induced and spontaneous activity per stimulation. Stimulation affected this neural activity, and so the communication formed a loop, with a loop time of approximately one second.

In the prior example, the sensory-motor mappings used a stable neural property to reliably control the robot. With Meart, the sensory-motor mappings are less well defined, in the hope of demonstrating a micro-scale version of the brain's creative processes. The behavioral response of the robot sheds light on the properties of the neural network and directs further encoding refinements. Thus, Meart is a 'work in progress' with the sensory-motor encoding continuously being improved to demonstrate learning processes. An example drawing is shown in Figure 6. The drawings changed throughout the life of cultures (and were different for different cultures) demonstrating neural plasticity, however, the mechanisms are still under investigation.

3 Discussion

3.1 Embodying Cultures: Theory

A Blank Slate. Since the cultured neurons were first separated and allowed to settle onto the MEA at random, they start from a 'blank slate'. Neural structure is lost and the function of neural activity is no longer obvious, yet neural network processing remains, evidenced by the complex activity patterns we have observed. For traditional in vitro neural models, function is cloudy since activity no longer relates to or causes behavioral states or actions. One cannot say 'this neuron is involved in color perception' or 'this neural structure helps to coordinate balance' as could be said for in vivo experiments. Artificially embodying and situating cultured neurons redefines their behavioral function concretely.

The structure of neuronal networks is likely to be important in neuronal processing, and changes in structure are likely to underlie learning and memory [27]. Our cultured neurons

form two-dimensional monolayers; functional importance may lie in the affordances given by the three-dimensional layered nature of the cortex. We and others in the Neurolab are pursuing the construction of 3D MEAs to support three-dimensional cultures, as part of an NIH Bioengineering Research Partnership [28], [29]. However, even cultured cortical monolayers (without 3D structure nor sub-cortical regions) have demonstrated an ability to adapt following stimulation via potentiation and/or depression [30], [31], [32], [33]. We are exploring using these plasticity mechanisms as a means to shape the network during development, within the Neurally-Controlled Animat paradigm, so it is no longer a blank slate.

Associations. The biological brain makes associations between different phenomena observed through sensation, whether between various external stimuli or between the actions of a body and their consequences, and then commands movement accordingly. Our methods have been developed to study these processes in real time with enough resolution to capture the dynamics of these interactions. These processes can be expressed using dynamical systems theory (DST), a mathematical framework to describe systems that change in time. For example, the formation of certain functional structures (ocular dominance columns) in the visual cortex has been described using Alan Turing's reaction-diffusion equations [34]. Kuniyoshi and his group explore DST to connect sensory-motor control to the cognitive level [35]. As applied to cognition [34], DST describes the mind with a set of complex, recursive filters. This opposes the classical cognitive concept of neural processing being analogous to a digital computer, containing distinct storage and processing of symbols [36], [37]. DST contends that multiple feedback loops and transmission delays, both of which are widespread in the brain, provide a time dimension to allow higher-level cognition to emerge without the need for symbolic processing [38]. DST is a framework compatible with embodied perspectives. The dynamical systems perspective has too often been neglected in neurobiology and cognitive sciences.

In contrast to an intact brain in an animal, cultures of neurons are isolated because they do not contain the afferent sensory inputs or efferent motor outputs a body would provide and therefore no longer have a world with which to reference their activity. Under these conditions, what associations can the network make, and what would those associations mean? Moreover, what symbols are operated on? Because of this, any associations that are made must consequently be self-referential or circular and neural activity may be misleading. The network as a set of complex, recursive filters has no external signals to filter, possibly leading to the abnormal barrage activity described below. To address this major shortcoming of in vitro systems, our neural cultures are embodied with sensory feedback systems, motor systems, and situated in an environment, providing a new frame of reference. New findings about the dynamics of living neural networks might be used to design more biological, less artificial AI.

Intelligence and Meaning. By embodying cultured neurons, the 'meaning' of neural activity emerges, since this activity affects subsequent stimulation. Now the network has a body behaving and producing experiences, allowing for the study of concepts such as in-

telligence. We will take a behavioral definition of intelligence as our start: Rodney Brooks describes intelligence in terms of how successfully an agent interacts with its world to achieve goal directed behavior [39]. William James states, "Intelligent beings find a way to reach their goal, even if circuitous," [40]. Neurons have inherent local goals (to transmit signals, integrate synaptic input, optimize synaptic strengths, and much more) that provide the foundation to intelligently achieve meaningful behavioral goals. No doubt the basis for intelligence is inherent at birth, but an interaction with a sufficiently complex environment (learning) is needed to develop it.

In our cultured networks, the local goals of neural interaction are subject to detailed optical and electrical observation, while the execution of higher-level behavioral goals are observed through the activities of the robotic body. (Note that the behavioral goals are artificially constrained by the stimulation and recording transformations chosen.) We hope this combination will lead to a clearer definition and a better understanding of the neurological basis of intelligence, in addition to explanations of other psychological terms: learning, memory, creativity, etc. Neurobiology has given inspiration to AI since the advent of the perceptron and consequent artificial neural networks, which are based on the local properties (goals) of individual neurons. We wish to continue this trend by finding the principles of network processing by multiple neurons that lead to higher-level goals.

Network-wide Bursting. The activity of cultured neurons tends towards the formation of dish-wide global bursts (barrages) [8]: sweeps of fast, multiple neural firings throughout the network lasting between hundreds of milliseconds to seconds in duration. These barrages have been observed often in cultured neurons [41] but also in cortical slices [42] and in computer models [43]. Barrages of activity are reported in the cortex in vivo during early development, during epileptic seizures, while asleep, and when under anesthesia. These in vivo examples of barrages occur over finite periods of time. In contrast, barrages in vitro are continuous over the life of the culture. We consider the possibility that at some stage, dish-wide barrages of spiking activity are abnormal, a consequence of 'sensory deprivation' (manuscript in preparation), or a sign of arrested development [44].

For both a model system [43] and for cultured mouse spinal neurons [45], if more than 30% of the neurons are endogenously active, the neurons fire at a low steady rate of 1 to 5 Hz per neuron, while a reduction in the fraction of endogenously active cells leads to barrage activity. Endogenous activity is functionally similar to activity induced by afferent input, suggesting embodiment would lead to low steady firing rates. The hypothesis is then that the barrage activity may be due to the lack of an external environment with which to interact. We are developing animat mappings in which continuous sensory input quiets barrages, bringing the networks to a less 'sensory-deprived' state that allows more complex, localized activity patterns.

3.2 The Importance of Embodiment

The World and the Brain. Environmental deprivation leads to abnormal brain structure and function, and environmental exposure shapes neural development. Similarly, patterned stimulation supplied to cultured neurons may lead to more robust network structure and functioning than with trivial or no stimulation. The most dramatic examples of the importance of embodiment come from studies during development, when the brain is most malleable. Cognitive tests were performed on institutionalized children in Romania, children typically deprived of proper environmental and social interaction early in life [46], [47]. Compared to peers, the children showed severe developmental impairment that improved, however, after transplantation to a stable family. Those adopted prior to 6 months of age achieved nearly complete cognitive catch-up to similarly aged children, while those adopted after 6 months of age had significant but incomplete catch-up. Likewise, laboratory rats raised in environments with mazes and varied visual stimuli had 30% greater cortical synaptic density than those raised in minimalist environments, and performed better in various cognitive experiments [48], [49]. Synaptic morphology in adults [1] and adult neurogenesis is dependent on external cues [50] demonstrating that environmental interaction is important throughout life.

A disembodied neural culture, whose activity never influences future stimulation, will not develop meaningful associations to an input. In the brain, if a sensation is not useful in influencing future behavior (no association is made between the two) the percept of the sensation fades. The environment triggers an enormous number of sensory signals, and the brain develops to filter out the excess while perceiving the behaviorally relevant. All one-month-old infants can distinguish between the English L and R sounds. Five months later, Japanese infants lose the ability while American infants maintain it, because the distinction is not needed to understand the Japanese language [51]. Japanese adults consequently have great difficulty distinguishing these sounds, but perception of the distinction can be learned through targeted instruction. These studies further demonstrate how brain (re)wiring depends on environmental context and occurs throughout life: the brain focuses on perceiving the portions of the environment relevant to produce a meaningful interaction.

The Body and the Brain. The choice of how to instantiate an animat or hybrid is important to processing in cultured neural networks. For example, the body, with its various sensory apparatus and motor output, is what detects and interacts with the environment. In addition to how different environments cause differences in the brain, differences in the body will have analogous effects on the brain. Changes in the frequency or type of sensory input via practice or surgical manipulation of the body causes gross shifts in the functional organization of corresponding cortical areas (the somatotopic maps) [52]. Amputation causes a sudden change to a body, and amputees later report having at times a sensation or impression that the limb is still attached. The impression lasts for days or weeks in most cases (years or decades in other cases) and then gradually fades from consciousness [53]. These false 'phantom limb' sensations arise because the brain has wired

itself for a given body that has now changed. This discrepancy further suggests the body and its interaction with the environment influence brain wiring and cognitive function. Neurally-Controlled Animats allow an unlimited variety of bodies to be studied; their structure and operating parameters can be easily varied to test effects on brain-body interactions.

4 Summary: Integrating Brain, Body, and Environment.

The above paragraphs were worded as if the entities brain, body, and environment are independent. Finding physical boundaries between the three is easy, but since the brain is so enmeshed in the states of the body (influencing mood, attention, and more), which in turn are so enmeshed in the body's interaction with its environment, finding functional boundaries between the three is difficult, if possible at all [25], [54], [55]. Damasio contends that the mind depends on the complex interplay of the brain and the body, and consequently emotions and rationality cannot be segregated [56].

We have integrated our hybros' brain (cultured network), body (robot or simulated animat), and environment (simulation, lab, or gallery) into a functional whole, even while the parts are sometimes 12,000 miles apart. Our experiments with these Neurally-Controlled Animats so far are rudimentary: we are still setting up the microscopic imaging systems to allow us to make correlations between changes in behavior and changes in neuron or network structure; we have not yet developed sensory-motor mappings that reliably result in learning. But in the process of creating this new research paradigm of embodied, situated cultured networks, we have already sparked a philosophical debate about the epistemological status of such semi-living systems,³ and have raised a number of issues about the validity of traditional (disembodied) in vitro neural research. We hope that others will make use of the tools we have developed such as our MeaBench software,⁴ sealed-dish culture system [5], and multi-site stimulation tools [57], to pursue a wide variety of questions about how neural systems function. We expect that these inquiries will lead to fundamentally different, more capable, and less artificial forms of AI.

Acknowledgments

We thank the NIH (NINDS, NIBIB), the Whitaker Foundation, the NSF Center for Behavioral Neuroscience, and Arts Western Australia for funding. We thank Daniel Wagenaar, Radhika Madhavan, John Brumfield, Zenas Chao, Eno Ekong, Gustavo Prado, Bryan Williams, Peter Passaro, and Ian Sweetman for their many contributions to this work.

³ Manson, N (2004) "Brains, vats, and neurally-controlled animats," in *Studies in the History and Philosophy of Biology and the Biomedical Sciences*, special issue on "The Brain in a Vat."

⁴ <http://www.its.caltech.edu/~pinelab/wagenaar/meabench.html>

References

1. Weiler, I. J., Hawrylak, N. & Greenough, W. T.: Morphogenesis in Memory Formation - Synaptic and Cellular Mechanisms. *Behavioural Brain Research* 66 (1995) 1-6
2. Dennett, D. C.: *Consciousness Explained*. Little, Brown and Co., Boston, (1991)
3. Morris, C. G.: *Psychology: An Introduction*. Appleton-Century-Crofts, New York, (1973)
4. Potter, S. M.: Distributed processing in cultured neuronal networks. In: Nicolelis, M. A. L. (ed.^(eds.): *Progress In Brain Research: Advances in Neural Population Coding*. Elsevier, Amsterdam, (2001) 49-62
5. Potter, S. M. & DeMarse, T. B.: A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods* 110 (2001) 17-24
6. Watanabe, S., Jimbo, Y., Kamioka, H., Kirino, Y. & Kawana, A.: Development of low magnesium-induced spontaneous synchronized bursting and GABAergic modulation in cultured rat neocortical neurons. *Neuroscience Letters* 210 (1996) 41-44
7. Gross, G. W., Rhoades, B. K. & Kowalski, J. K.: Dynamics of burst patterns generated by monolayer networks in culture. In: Bothe, H. W., Samii, M. & Eckmiller, R. (ed.^(eds.): *Neuro-bionics: An Interdisciplinary Approach to Substitute Impaired Functions of the Human Nervous System*. North-Holland, Amsterdam, (1993) 89-121
8. Kamioka, H., Maeda, E., Jimbo, Y., Robinson, H. P. C. & Kawana, A.: Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neuroscience Letters* 206 (1996) 109-112
9. Potter, S. M.: Two-Photon Microscopy for 4D Imaging of Living Neurons. In: Yuste, R., Lanni, F. & Konnerth, A. (ed.^(eds.): *Imaging Neurons: A Laboratory Manual*. CSHL Press, Cold Spring Harbor, (2000) 20.1-20.16
10. Potter, S. M.: Vital imaging: Two photons are better than one. *Current Biology* 6 (1996) 1595-1598
11. Potter, S. M., Lukina, N., Longmuir, K. J. & Wu, Y.: Multi-site two-photon imaging of neurons on multi-electrode arrays. *SPIE Proceedings* 4262 (2001) 104-110
12. Wheeler, B. C.: Automatic discrimination of single units. In: Nicolelis, M. A. L. (ed.^(eds.): *Methods for Neural Ensemble Recordings*. CRC Press, Boca Raton, (1999) 61-77
13. Wagenaar, D. A. & Potter, S. M.: A versatile all-channel stimulator for electrode arrays, with real-time control. *Journal of Neural Engineering* 1 (2004) 1-7
14. Meyer, J. A. & Wilson, S. W.: *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, (1991)
15. Potter, S. M., Fraser, S. E. & Pine, J.: Animat in a Petri Dish: Cultured Neural Networks for Studying Neural Computation. *Proc. 4th Joint Symposium on Neural Computation, UCSD* (1997) 167-174
16. DeMarse, T. B., Wagenaar, D. A., Blau, A. W. & Potter, S. M.: The Neurally Controlled Animat: Biological Brains Acting with Simulated Bodies. *Autonomous Robots* 11 (2001) 305-310
17. DeMarse, T. B., Wagenaar, D. A. & Potter, S. M.: The neurally-controlled artificial animal: A neural-computer interface between cultured neural networks and a robotic body. *Society for Neuroscience Abstracts* 28 (2002) 347.1
18. Reger, B. D., Fleming, K. M., Sanguinetti, V., Alford, S. & Mussa-Ivaldi, F. A. Connecting brains to robots: The development of a hybrid system for the study of learning in neural tissues. In: *Proc. of the VIIth Intl. Conf. on Artificial Life*. (2000) 263-272

19. Sharp, A. A., Abbott, L. F. & al, e.: The Dynamic Clamp: Computer-generated conductances in real neurons. Pre-print (1992)
20. Butera, R. J., Wilson, C. G., DelNegro, C. A. & Smith, J. C.: A methodology for achieving high-speed rates for artificial conductance injection in electrically excitable biological cells. *Ieee Transactions on Biomedical Engineering* 48 (2001) 1460-1470
21. Simoni, M., Cymbaluyk, G., Sorensen, M., Calabrese, R. & DeWeerth, S.: Development of Hybrid Systems: Interfacing a Silicon Neuron to a Leech Heart Interneuron. In: Leen, T. K., Dietterich, T. G. & Tresp, V. (ed.^(eds.): *Advances in Neural Information Processing Systems* 13, NIPS 2000. MIT Press, Boston, (2001) 173-179
22. Braitenberg, V.: *Vehicles, experiments in synthetic psychology*. MIT Press, Cambridge, Mass., (1984)
23. Arkin, R. C.: *Behavior-Based Robotics*. MIT Press, Cambridge, (1999)
24. Brooks, R. A.: *Cambrian Intelligence: The Early History of the New AI*. MIT Press, Cambridge, MA, (1999)
25. Clark, A.: *Being There: Putting Brain, Body, and the World Together Again*. MIT Press, Cambridge, (1997)
26. Shkolnik, A. C. *Neurally Controlled Simulated Robot: Applying Cultured Neurons to Handle and Approach/Avoidance Task in Real Time, and a Framework for Studying Learning In Vitro*. In: Potter, S. M. & Lu, J.: Dept. of Mathematics and Computer Science. Emory University, Atlanta (2003)
27. Engert, F. & Bonhoeffer, T.: Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 399 (1999) 66-70
28. Choi, Y. et al. High aspect ratio SU-8 structures for 3-D culturing of neurons. In: *ASME International Mechanical Engineering Congress and RD&D Expo*. Washington, D. C. (2003)
29. Blum, R. A. et al.: A custom multielectrode array with integrated low-noise preamplifiers. In: Akay, M. (ed.^(eds.): *Proceedings of the IEEE Engineering in Medicine and Biology Conference*. (2003) 3396-3399
30. Jimbo, Y., Tateno, T. & Robinson, H. P. C.: Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophysical Journal* 76 (1999) 670-678
31. Tateno, T. & Jimbo, Y.: Activity-dependent enhancement in the reliability of correlated spike timings in cultured cortical neurons. *Biological Cybernetics* 80 (1999) 45-55
32. Marom, S. & Shahaf, G.: Development, learning and memory in large random networks of cortical neurons: Lessons beyond anatomy. *Quarterly Reviews of Biophysics* 35 (2002) 63-87
33. Eytan, D., Brenner, N. & Marom, S.: Selective adaptation in networks of cortical neurons. *Journal of Neuroscience* 23 (2003) 9349-9356
34. Turing, A. M.: The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B* 237 (1953) 37-72
35. Yamamoto, T. & Kuniyoshi, Y. Stability and controllability in a rising motion: a global dynamics approach. In: *International Conference on Intelligent Robots and Systems (IROS)*. Lausanne, Switzerland (2002) 2467-2472
36. Fodor, J. A.: Methodological Solipsism Considered as a Research Strategy in Cognitive-Psychology. *Behavioral and Brain Sciences* 3 (1980) 63-73
37. Vera, A. H. & Simon, H. A.: Situated Action - a Symbolic Interpretation. *Cognitive Science* 17 (1993) 7-48

38. Edelman, G. M. & Tononi, G.: A universe of consciousness: how matter becomes imagination. Basic Books, New York, (2000)
39. Brooks, R. A.: Intelligence without representation. *Artificial Intelligence* 47 (1991) 139-159
40. James, W.: The principles of psychology. H. Holt, New York, (1890)
41. Nakanishi, K. & Kukita, F.: Functional synapses in synchronized bursting of neocortical neurons in culture. *Brain Research* 795 (1998) 137-146
42. Corner, M. A. & Ramakers, G. J.: Spontaneous bioelectric activity as both dependent and independent variable in cortical maturation. Chronic tetrodotoxin versus picrotoxin effects on spike-train patterns in developing rat neocortex neurons during long-term culture. *Ann N Y Acad Sci* 627 (1991) 349-53
43. Latham, P. E., Richmond, B. J., Nelson, P. G. & Nirenberg, S.: Intrinsic dynamics in neuronal networks. I. Theory. *Journal of Neurophysiology* 83 (2000) 808-827
44. Tabak, J. & Latham, P. E.: Analysis of spontaneous bursting activity in random neural networks. *Neuroreport* 14 (2003) 1445-1449
45. Latham, P. E., Richmond, B. J., Nirenberg, S. & Nelson, P. G.: Intrinsic dynamics in neuronal networks. II. Experiment. *Journal of Neurophysiology* 83 (2000) 828-835
46. Rutter, M.: Developmental catch-up, and deficit, following adoption after severe global early privation. *Journal of Child Psychology and Psychiatry and Allied Disciplines* 39 (1998) 465-476
47. O'Connor, T. G., Rutter, M., Beckett, C., Keaveney, L. & Kreppner, J. M.: The effects of global severe privation on cognitive competence: Extension and longitudinal follow-up. *Child Development* 71 (2000) 376-390
48. Black, J. E., Isaacs, K. R., Anderson, B. J., Alcantara, A. A. & Greenough, W. T.: Learning Causes Synaptogenesis, Whereas Motor-Activity Causes Angiogenesis, in Cerebellar Cortex of Adult-Rats. *Proceedings of the National Academy of Sciences of the United States of America* 87 (1990) 5568-5572
49. Diamond, M.: Morphological cortical changes as a consequence of learning and experience. In: Scheibel, A. B. & Wechsler, A. F. (ed.^(eds.): *Neurobiology of Higher Cognitive Function*. Guilford Press, New York, (1990) 370
50. Gross, C. G.: Neurogenesis in the adult brain: death of a dogma. *Nature Reviews Neuroscience* 1 (2000) 67-73
51. Kuhl, P. K. et al.: Cross-language analysis of phonetic units in language addressed to infants. *Science* 277 (1997) 684-686
52. Buonomano, D. V. & Merzenich, M. M.: Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience* 21 (1998) 149-186
53. Ramachandran, V. S. & Hirstein, W.: The perception of phantom limbs - The D.O. Hebb lecture. *Brain* 121 (1998) 1603-1630
54. Varela, F. J., Thompson, E. & Rosch, E.: *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge, Mass., (1993)
55. Pfeifer, R. & Scheier, C.: *Understanding Intelligence*. The MIT Press, Cambridge, Massachusetts, (1999)
56. Damasio, A. R.: *Descartes' Error: Emotion, Reason, and the Human Brain*. Gosset/Putnam Press, New York, (1994)
57. Wagenaar, D. A. & Potter, S. M.: Real-time multi-channel stimulus artifact suppression by local curve fitting. *J. Neurosci. Methods* 120 (2002) 113-120