

MEART: The semi-living artist

Douglas J. Bakkum¹, Philip M. Gamblen², Guy Ben-Ary², Zenas C. Chao¹
and Steve M. Potter^{1,*}

1. Laboratory for Neuroengineering, School of Mechanical Engineering, Georgia Institute of Technology, USA

2. SymbioticA, School of Anatomy and Human Biology, University of Western Australia, Australia

Edited by: Frederic Kaplan, Ecole Polytechnique Federale de Lausanne, Switzerland

Reviewed by: Frederic Kaplan, Ecole Polytechnique Federale de Lausanne, Switzerland

Suguru N. Kudoh, National Institute of Advanced Industrial Science and Technology, Osaka, Japan

Sergio Martinoia, Department of Biophysical and Electronic Engineering (DIBE), University of Genova, Italy

Here, we and others describe an unusual neurobotic project, a merging of art and science called MEART, the semi-living artist. We built a pneumatically actuated robotic arm to create drawings, as controlled by a living network of neurons from rat cortex grown on a multi-electrode array (MEA). Such embodied cultured networks formed a real-time closed-loop system which could now behave and receive electrical stimulation as feedback on its behavior. We used MEART and simulated embodiments, or animats, to study the network mechanisms that produce adaptive, goal-directed behavior. This approach to neural interfacing will help instruct the design of other hybrid neural-robotic systems we call hybrot. The interfacing technologies and algorithms developed have potential applications in responsive deep brain stimulation systems and for motor prosthetics using sensory components. In a broader context, MEART educates the public about neuroscience, neural interfaces, and robotics. It has paved the way for critical discussions on the future of bio-art and of biotechnology.

Keywords: learning, embodiment, multi-electrode array, neural network, rat, art

INTRODUCTION

“The most beautiful thing we can experience is the mysterious. It is the source of all true art and all science.”—Albert Einstein, 1931

The emergence of the mind from its biological substrate is one of the greatest and most complex mysteries. We study the brain using a synthetic approach, building from scratch a simple artificial animal, a new type of model for studying the brain. To be useful and easy to control and study, a model necessarily is a simpler version of what it models. Although our approach is fairly reductionistic, we assume that the complexity found in living brain cells is crucial to their function, including their network dynamics. Thus, our synthetic model system incorporates living neuronal networks, and is therefore a cybernetic organism, or cyborg. To distance this approach the culturally loaded conception of a cyborg, we prefer to call simple hybrid neural-robotic systems used for neurobiology research “hybrot” (Potter, 2002).

We built a robotic drawing machine with two pneumatically actuated arms that move in concert to draw with ink markers on large sheets of paper (Figure 1) and designed software and hardware for it to converse with a network of rat cortical neurons grown in culture over a multi-electrode array (MEA, Figure 2) (Potter et al., 2006). The model system consisted of living neurons, growing in the laboratory for Neuroengineering at Georgia Tech, and connected by internet to the pen-wielding metal and plastic pair of arms behaving in gallery exhibitions around the world over

the past 5 years. The whole system was named MEART, an acronym derived from Multi-Electrode Array aRT. This geographically distributed, “semi-living artist” was one of the first closed-loop neurally controlled animats with a robotic body (Manson, 2004; Potter et al., 1997; Reger et al., 2000). Neuronal action potentials recorded by an MEA in Atlanta were processed in real-time and used to command movement at different exhibitions in Perth, Melbourne, Bilbao, New York, Moscow, Atlanta, and Shanghai (<http://www.fishandchips.uwa.edu.au/exhibitions.html>). Video images of the drawings in progress determined the subsequent feedback of electrical stimuli delivered to the neurons.

Artists in Perth and scientists in Atlanta collaborated to construct MEART, a concept originating from scientific inquiries into hybrid bio-robotic technology (DeMarse et al., 2001), and artistic expressions by SymbioticA, an art–science collaboratory in the School of Anatomy and Human Biology at the University of Western Australia. Our common interest was to explore the essence, or primordial substrates, of creativity and intelligence. Because MEAs are so much more accessible than brains in animals, they allow researchers to manipulate and quantify underlying neural mechanisms of small (a few thousand neurons) networks, including the physical manifestations of learning and memory (Jimbo et al., 1999; Potter et al., 2001).

The idea of controlling robots with brain cells removed from the body and interfaced to electronics went from the realm of science fiction to that of science when Mussa-Ivaldi and co-workers at Northwestern University interfaced the small, wheeled Khepera robot (K-Team) to a lamprey brainstem maintained *Ex vivo* (Reger et al., 2000), taking advantage of the existing vestibular circuitry in that part of the brain to study adaptation mechanisms. They re-mapped the lamprey brain’s circuitry to take input from the robot’s photosensors, and to control the motors with its vestibular response to this artificial input. “The semantics of the stimulus (gravity vs. light) is not likely to play any substantial role here,” they asserted (Karniel et al., 2005). This hybrot demonstrated phototaxis, and rudimentary learning, by changing its responses to light.

* Correspondence: S. M. Potter, Laboratory for Neuroengineering, Coulter Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0535, USA. e-mail: steve.potter@bme.gatech.edu

Received: 12 Sep. 2007; paper pending published: 9 Oct. 2007; accepted: 15 Oct. 2007; published online: 2 Nov. 2007

Full citation: *Frontiers in Neurorobotics* (2007) 1:5 doi: 10.3389/neuro.12/005.2007
Copyright: © 2007 Bakkum, Gamblen, Ben-Ary, Chao, Potter. This is an open-access article subject to an exclusive license agreement between the authors and the Frontiers Research Foundation, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are credited.




Figure 1. MEART's body. Two arms cooperated to grip a set of colored pens and move them across a sheet of paper, according to neural activity in a culture dish that was up to 12 000 miles away. A CCD camera aimed at the drawing provided sensory feedback to the neuronal network.

When cultured networks serve as the brain of a hybot, any intrinsic brain circuitry from the donor was lost during dissociation of the brain tissue during preparation of the cultures. A cortical culture lacks the 3D structure present in the brain and so lacks any computational advantages that this may have afforded. However, basic self-organizing principles and plasticity mechanisms such as spike timing-dependent plasticity (Bi and Poo, 1998) and homeostasis (Turrigiano and Nelson, 2000) persist and were the objects of our study. To what extent an organized network reforms *in vitro* is still up for debate. However, we and others have shown that even dissociated networks of neurons have the ability to produce complex, repeating patterns of activity (Rolston et al., 2007; Wagenaar et al., 2006a, 2006b). In 2002, we presented a poster describing a simple approach-avoidance task executed by a Khepera interfaced to a cultured cortical network (DeMarse et al., 2002). Others using hybots with cultured neurons as their brain include Kudoh and co-workers at the National Institute of Advanced Industrial Science and Technology in Japan (Kudoh and Taguchi, 2006) and Martinoia and co-workers at the University of Genoa in Italy (Martinoia et al., 2004; Novellino et al., 2007). Both of these groups also used the Khepera as the embodiment, in an obstacle-avoidance paradigm that included tetanic electrical stimulation to induce learning

“Certain types of feedback stimulation caused suppression of spontaneous network electrical activities and drastic re-organization of functional connections between neurons, when these activities are initially almost synchronized. The result suggests that neurons

in dissociated culture autonomously re-organized their functional neuronal networks [by interacting] with their environment. The spatio-temporal pattern of activity in the networks maybe a reflection of their external environment.” (Kudoh and Taguchi, 2006)

This *embodied cultured networks* approach is intended to bridge a large gap that exists between *in vivo* behavioral studies of learning and memory, and *in vitro* studies of cellular plasticity. With a hybot whose living brain can be easily probed and observed, behavior and learning can be observed in concert with the detailed and long-term multi-neuron electrophysiology available *in vitro* (Potter and DeMarse, 2001). We sought to find out whether MEART could learn something about the environment given to it, and whether a creative act could emerge from its interactions with this environment. We define learning in this context as a lasting change in behavior that results from experience. Here we present, along with artistic, philosophical, and scientific commentary, progress on engineering MEART's hardware, software, wetware, environment, and aesthetics. In experiments directed at making MEART learn, we applied patterned training stimuli (PTS) contingent on behavioral performance in order to achieve the goal-directed behavior of drawing geometrical shapes. Neural plasticity occurred, but successful learning did not. However, we modified the training algorithm using a living network connected to a simulated robot [an animat (Meyer and Guillot, 1994)]. Instead of a fixed transformation from sense data to stimuli, behavioral performance was used to continuously discover and refine effective sequences of PTSs, and in a preliminary experiment described below, an animat repeatedly learned to draw in different desired directions. By using more detailed sensation and motor output, we expect hybots to demonstrate increasingly complex and interesting behaviors. What questions would be posed if MEART was eventually deemed to show intelligent or creative behavior? What would be the implications for biotechnology if its drawings were considered aesthetically beautiful?

The unique nature of this art–science exploration in neurorobotics has stimulated wide-ranging discussion, about life, art, learning, embodiment, and other things, some of which is excerpted here. We hope that this discussion continues online via the Frontiers in NeuroRobotics web site.

METHODS: MAKING THE SEMI-LIVING ARTIST

MEART was comprised of living neurons, recording and stimulating electronics, robotic drawing arms, electronic control circuits for a pneumatic actuation system, a CCD camera to feedback images of drawings, and software communicating between the neurons and robot over the internet (Figure 3). The *simulated* animat was made of living neurons, recording and stimulating hardware, and a simple virtual embodiment on a computer. It was used to develop protocols in the




Figure 2. MEART's brain (above) and a MEA culture dish (below). A week-old culture of ~50 000 neurons and glia from embryonic rat cortex, growing in a MEA and forming a dense network 1–2 mm across. Fifty-nine 30 μm electrodes spaced at 200 μm intervals connect a few hundred of the network's neurons to the outside world, by allowing their activity to be extracellularly recorded or evoked by electrical stimulation.



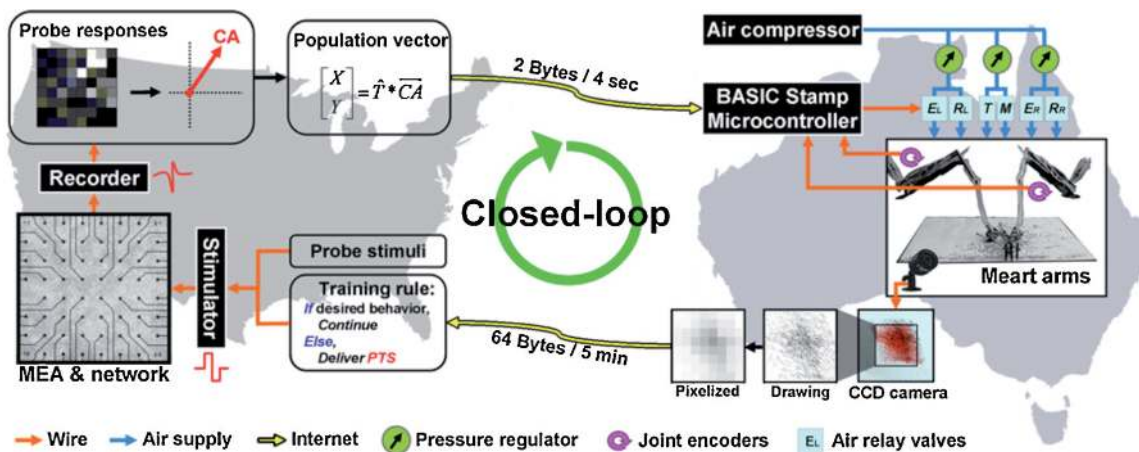


Figure 3. Schematic of the bio-robotic software algorithms and hardware, i.e., MEART's components. **Commanding movement:** The center of activity (CA) of neuronal action potentials was calculated from 100 ms of responses after a probe stimulation (8×8 box representing the MEA; increasing firing rate is black to white). Animat movement was instructed from a transformation (\hat{T}) of the CA into a population vector. The $[X, Y]$ movement command was sent over the internet (yellow arrows) to the robotic arms every 4 seconds. **Movement:** The robotic drawing machine consisted of two perpendicular arms actuated by braided pneumatic artificial muscles, allowing independent retraction (R) or extension (E) of the left (E_L/R_L) and right (E_R/R_R) arms within approximately a 30 cm by 30 cm workspace. Similarly, smaller muscles pressed the pens to the paper when at the target location (T), or optionally to trace movement trajectories (M). The supply line from an air compressor was split between three pressure regulators (green circles, one for each arm and one for the pens). 24 V AC pneumatic valves (light blue rectangles) controlled muscle air pressure. Joint encoders (purple arrows; 10 k potentiometers) tracked arm location, and a BASIC Stamp microcontroller (BS2SX-IC) modulated the relay valves to provide accurate movement as commanded by the neurons' activity. **Sensory feedback:** A CCD camera located above the workspace captured an image of accumulating markings every 5 minutes. The images were pixelated into 8 bit grayscale values (isomorphic to the electrodes on the MEA) and sent back over the internet to command feedback stimulation of the neurons. **Training:** Animat behavior was compared to the goal behavior to control training stimulation. Feedback stimuli could change neuronal activity, in turn varying subsequent animat movement and sensory feedback, thus forming a closed-loop system. TCP/IP sockets were used to communicate between the drawing robot and the neuronal network, which were often located on separate continents.

intervals between MEART exhibitions. Three major topics needed to be addressed to embody the cultured networks are as follows:

- The care and feeding of the biological brain;
- The hardware (or software) implementation of the body; and
- The sensory transformation, motor transformation, and training algorithms.

A Preparing and caring for MEART's brain

We have developed techniques to maintain neuronal cultures and conduct experiments for many months using MEAs (Potter and DeMarse, 2001). We describe these briefly, and refer the enthusiast to that paper for more details. Cells were obtained from embryonic-day-18 rat cortex according to protocols approved by the NIH and the Georgia Institute of Technology animal care and use committee. Brain tissue was dissociated with enzymes and mechanical trituration, to prepare a dense suspension of neurons and glia. A droplet of this suspension containing about 50 000 cells was pipetted into MEAs coated with polyethylene imine and laminin, and cultured at high density (~ 3000 cells/mm²) in serum-containing Dulbecco's Modified Eagle's Medium. The MEAs used were glass with silicon nitride insulation and 60 titanium nitride electrodes (multichannel systems). Neural activity was recorded using the MEA60 preamplifier and MCard analog-to-digital converter (multichannel systems) with each of 60 channels being digitized at 25 kHz. All cultures were allowed to grow 3 weeks prior to experimentation, with weekly medium replacement. Neurons spontaneously began communicating electrically and chemically within a week, demonstrating an inherent goal to form a functional network (Van Pelt et al., 2004; Wagenaar et al., 2006a) and distinct repeating patterns of activity (Rolston et al., 2007; Wagenaar et al., 2006b). Sensory input to the networks was delivered via the substrate electrodes as voltage-controlled pulses. These were biphasic pulses of 400 μ s duration and 500 mV magnitude per phase (Wagenaar et al., 2004) using a custom built all-electrode stimulator (Wagenaar and Potter, 2004). Data acquisition, visualization, artifact suppression, and spike detection were controlled using Meabench (Wagenaar et al., 2005a).




Figure 4. Life-support system for MEART's brain. The microscope used for observing neural cultures in long-term experiments was wrapped in insulation and outfitted with systems for control of temperature and carbon dioxide levels to maintain normal cell culturing conditions.

Experiments were conducted using sealed-lid MEAs (Potter and DeMarse, 2001) inside an environmentally controlled incubator built around an optical microscope (Figure 4), allowing us to monitor and stimulate the networks continuously for many days.

B MEART's body

Artistic design. The MEART data presented here were collected during the First Moscow Biennale of Contemporary Art at an exhibition entitled “art.digital.2004: I Click, Therefore I Am”, where MEART's goal of filling a square at the center of the drawing workspace was inspired by the Russian artist Kazimir Malevich's “Black Square” painting. From the art.digital.2004 program, “The action of MEART observing and drawing the Black Square explores the fundamentals of visual creativity and the way we communicate with the world through images, symbols, and their underlying meanings.” This goal behavior was a simplification of the mappings used during our previous exhibitions, to improve experimental controllability. In previous MEART exhibits, we added an element of interactivity by having MEART draw photographed faces of gallery attendees, entitled the “Portrait Series”. As with images from the drawing, the faces were pixilated and MEART's goal was to shade the drawing to match the grayscale pixel values. To give viewers a better understanding of MEART's brain, and the laboratory in which it was studied, live images from the laboratory, a close-up of the MEA, and a data display of neural activity were projected onto the exhibit walls. This, along with computers displaying the movement and feedback data streams, made the distributed nature of MEART more apparent.

MEART's body was designed to closely resemble organic forms in function and aesthetics. Shapes were based on bones [influenced by the photographer Andreas Feininger (Feininger and Schlatter, 2003)], and sanded Perspex offered an elegant look that referred to a skeletal structure. Similarly, the pneumatic muscles paralleled biological muscles. The design had no covering, never attempting to hide or deny the underlying technology. Analogously, the complex biology of the rat was reduced to a few thousand neurons and glia, grown *in vitro*. MEART was thus a symbol of the reductionist nature of science and of the stripping down to expose the physical substrates of the creative process.

Below, Emma McRae (2004) paints a verbal picture of MEART

“1. Introduction to a Cybernetic Entity

The soft popping sounds of air releasing, of the breaths taken between movements as the muscles contract and release on the mechanical structures at work on the table in the centre of the room, reach me first as I walk down the dark corridor in the Australian Centre for the Moving Image. I can see the plastic and metal arms and the tubes connected to two rows of valves—regular black garden hose valves—highlighted by a spotlight, that seem to create the movement of the arms. These arms (the creators call these structures arms, presumably because they hold pens and draw as human arms involved in drawing do) are busy drawing lines in apparently random directions with three different coloured pens on a large sheet of paper on the table. Behind the arms is a computer screen showing a photo of a man's face, a pixellated black and white image, a scrolling text box, and some graphs. The only other thing on the table is a camera which looks down over the arms at the picture they're drawing. A large screen on the wall behind the table shows a graph, a representation that looks like a glacial landscape and is constantly changing form, its peaks and troughs rising and falling in random motion, depicting varied intensities coloured in blue, yellow, white, and red. There are two smaller screens in the opposite corner of the room that intermittently display an image of a science laboratory, a close up of a petri-dish, a screen of 64 ECG-like blue tracking graphs, and a microscope view of cells.”

Movement. The drawing machine consisted of two perpendicular, rigid, jointed arms (aluminum and acrylic Perspex) fixed by hinges at their ends to a 3 m by 3 m table actuating the *X* and *Y* positions of a group of pens over a sheet of paper (Figures 1 and 5). Similar to biceps and triceps, McKibben braided pneumatic artificial muscles could contract individually, allowing independent flexion or extension of each arm within approximately a 30 cm by 30 cm workspace. Similarly, activation of smaller muscles pressed pens to the paper; a dark pen marked target




Figure 5. The body of MEART at the Moscow Biennale and drawings. (Top) Metal and plastic arms rested on a 3 m × 3 m table. Plastic tubes fed pressurized air to pneumatic muscles. A digital camera aimed at the paper captured images of the drawing as it progressed. (Bottom right) Development of MEART as reflected in its drawings: **New York** (July 2003). Video feedback was used for the first time to close the loop, but a “scribble” mode in effect randomized movement and pen placement. **Bilbao** (April 2004) Removing scribble demonstrated the arm moved between four points only, via eight movement directions corresponding to the possible combinations of muscle activation. Pen placement remained random. **Melbourne** (June 2004) Joint encoders were added to read in arm positions and command movement in a feedforward manner: Muscles were flexed for a duration proportional to the distance to reach the commanded location. Interior positions could be reached as in New York, however, accuracy was low. **Moscow** accuracy test (January 2005) A Basic STAMP microcontroller implemented feedback control of arm positions to achieve accurate movement. Outside pens were commanded down when at the target location. The middle pen was commanded down during arm movement.

locations, while an optional lighter colored pen traced the movement trajectories. The supply line from an air compressor was split between three pressure regulators, one for each arm and one for the pens, to isolate pressure fluctuations. Air pressure and thus arm and pen movement was controlled by opening and closing 24 V AC pneumatic valves. Pneumatic muscles, while offering a high power to weight ratio, produce nonlinear motion difficult to predict. Therefore, arm location was tracked using joint encoders (10 k potentiometers), and a BASIC Stamp microcontroller (BS2SX-IC) modulated valve opening to increase movement accuracy as commanded by the living network (Figure 6).

Sensory feedback. A digital camera located above the movement workspace captured images of the drawing in progress. Fluctuations in light from shadows and clouds could strongly influence the image quality. Therefore, ambient and natural light sources were reduced or eliminated except for bright spotlights on the drawing itself. Image inhomogeneity due to imperfect lighting was corrected by subtracting from the captured images an image of the sheet of paper when blank, prior to a drawing. The accumulation of markings was recorded every 5 minutes by retracting the






Figure 6. Accuracy test of the robotic drawing machine. Movements between seven locations were commanded 200 times in random order. A dark pen marked the target locations, while an offset lighter colored pen traced the movement trajectory. $3\text{ cm} \times 3\text{ cm}$ resolution targets could be reached within 4 second and a $1\text{ cm} \times 1\text{ cm}$ target around 10 second (not shown). A photograph of Malevich's "Black Square" painting can be seen projected on the gallery wall.

arms out of view and capturing an image, analogous to a painter stepping back from the canvas to check the work in progress.

Internet communication. TCP/IP sockets were used to send motor commands to the drawing machine and to return images of the progression of a drawing for feedback. To reduce internet bandwidth, 8 bit grayscale values of an 8×8 grid of pixels (isomorphic to the electrodes on the MEA) were transmitted over the internet and transformed into electrical stimulation feedback delivered to the neuronal network.

C Software development and experimental design

Motor transformation. For an animat to behave, sequences of neuronal action potentials need to be transformed into body movements, but understanding how such sequences might encode information is a subject of much scientific inquiry. Population vector coding is a candidate motor mapping found to occur in the motor cortex (Georgopoulos, 1994), premotor cortex (Caminiti et al., 1990), hippocampus (Wilson and McNaughton, 1993), and other cortical areas: the vector sum of firing rates of a group of broadly tuned neurons taken together provide a precisely tuned representation (e.g., to a preferred direction of arm movement).

We have used a new statistic, the center of neural activity (CA, analogous to the *center of mass*) to reliably quantify neuronal network plasticity on an MEA by including spatial information (Chao et al., 2007). Movement of MEART or a simulated animat was calculated from the CA of 100 ms of responses after each probe stimulus:

$$\text{Meart : } \begin{bmatrix} X \\ Y \end{bmatrix} = \hat{T} * \vec{CA} = \hat{T} * \frac{\sum_e (N_e \vec{W}_e)}{\sum_e N_e} \quad (1)$$

$$\text{Simulated animat : } \begin{bmatrix} dX \\ dY \end{bmatrix} = \hat{T} * \vec{CA} = \hat{T} * \frac{\sum_e (N_e \vec{W}_e)}{\sum_e N_e} \quad (2)$$

The CA is the vector summation of action potentials at each electrode e (N_e) weighted by the spatial location of the electrode, \vec{W}_e . The transformation, T , is a normalization matrix found prior to the closed-loop experiment to offset and scale the CAs (in electrode space) such that animat movement could produce a uniform distribution and the ability to place pen marks throughout the workspace (MEART) or move in any direction (simulated animat). Achieving a goal for either MEART or the animat required shifting the distribution of normalized CAs. Therefore, plasticity results were comparable. The responses to 1 Hz stimulation on a probe electrode were averaged between consecutive movements (every

4 second or 1/4 Hz) and used to command MEART pen location, while the responses to 1/4 Hz stimulation on a probe electrode were used to command the simulated animat. A single repeating probe electrode was used throughout an experiment.

Movement could be commanded by absolute location (MEART) or in relative increments (simulated animat). For each case, the activity was normalized to equally distribute the distribution of CAs prior to experiments. For absolute location, this set the possible pen locations to be distributed throughout the whole workspace. For incremental movement, this set the possible movement directions to be distributed throughout 360 degree. Absolute pen location was used with MEART to avoid movement exceeding the workspace, which would introduce discontinuities in behavior. Incremental movement (Equation 2) was later used for the simulated animat as workspace size was not physically limited, and we were more interested in direction of movement than position.

Training and sensory feedback. Previous MEART exhibits used a sensory mapping in which a camera's image, after reducing to 8×8 pixels, was directly mapped onto stimuli of the 8×8 grid of electrodes under the neuronal network. For the Moscow exhibit, the sensory system was simplified into a signal that merely indicated whether drawings were within a pre-defined square. Successful behavior was determined from comparisons between consecutive feedback images. If a larger proportion of markings occurred inside the target geometrical area than outside, behavior was considered successful. Otherwise, a change in the probe response was desired. For training, plasticity was induced by repetitive stimulation of paired electrodes, termed patterned training stimulation (PTS). A PTS was constructed by pairing the probe electrode with another active electrode (one that evokes network responses) 20 ms later, repetitively stimulated for 3 second with an inter-pair interval of 100 ms.

For the simulated animat (Bakkum et al., 2007), the training algorithm was modified in two ways. A pool of candidate PTSs was formed by pairing the probe electrode with other electrodes ($N_E = 58$) and inter-pulse intervals $\{-80, -40, -10, 10, 40, 80\text{ ms}\}$ ($N_{PTS} = 58 \times 6$). The probabilities of choosing a given PTS were initially uniform and increased or decreased based on whether subsequent animat performance was successful or not. This allowed an iterative search for an appropriate training "solution" to direct neuronal plasticity. Second, plasticity can arise from both the PTS stimuli and ongoing spontaneous activity occurring between probes. In a model network, a random stimulation stabilized neural synaptic weights (Chao et al., 2005). Therefore, when animat behavior was successful (no PTS application), a random background stimulation was used between probes such that the plasticity accumulated from a series of PTSs was maintained. The goal of the simulated animat was now to learn to move within ± 30 degree of a goal angle.

RESULTS

MEART was first exhibited in August 2002 at the Biennale of Electronic Arts Perth (BEAP). However, the precursor to MEART, Fish & Chips, was shown in 2001 at Ars Electronica in Austria. For this ground-breaking bio-art exhibit, SymbioticA Research Group created MEART's drawing arm and used it as the embodiment of a semi-living artist. This was called Fish & Chips because an acute goldfish brain slice was maintained and electrically interfaced on a silicon chip, and used as the controlling "brain" of the arm. From the collaboration between SymbioticA in Perth and the Potter laboratory in Atlanta, MEART was born: the first robot controlled by a network of neurons in a culture dish, with a two-way interface via a MEA. To the existing drawing arm, we added a sensory system, where images from a CCD camera were translated into electrical stimuli for the cultured network. It was also the first neurally controlled robot whose brain lay far away from its body, with the internet in some ways serving as a very long nerve connecting brain to body. It was the first physical embodiment for a cultured network that remained continuously connected for extended periods of several days, creating numerous drawings during exhibitions.

Early exhibitions were devoted to debugging the communication software and robot mechanics (Figure 5), and the most recent exhibitions allowed experimentation. We noticed early that continuous sensory input over the course of days tended to reduce the number of spontaneously occurring network-wide bursts. This led to a hypothesis that other types of bursting, such as epileptic seizures, might be treated by continuous multi-electrode stimulation. We quantified the short-term "quieting" effects of distributed multi-site stimulation on cortical cultures (Wagenaar et al., 2005b), and we are now pursuing the longer-term, or homeostatic effects of continuous stimulation that comes as a consequence of embodiment.

For the data presented here, MEART's behavioral goal was to draw a solid 12 cm × 12 cm square within the center of its 30 cm × 30 cm workspace. The simulated animat was used to test training algorithms between MEART exhibitions in order to improve behavioral performance. The simulated animat's behavioral goal was to incrementally move within ±30 degree of a desired angle (note that this differed from MEART's goal behavior of producing pen markings, commanded by absolute location). For both MEART and the simulated animat, the relationship between changes in behavior and the decision whether or not to apply feedback training stimulation were identical, and thus results about plasticity and learning were comparable.

Electrical stimulation can be an artificial inducer of neuronal plasticity, changing a network's input-output function. Bi and Poo found that for mono-synaptically connected cultured neurons firing within a few tens of milliseconds of each other, directional spike timing dependent synaptic plasticity occurred (Bi and Poo, 1998). Repetitive stimulation of pairs of electrodes in a PTS could therefore cause plasticity in shared pathways of neural activation.

For Meart, the transformation from visual sensation into the delivery of a PTS was fixed. For example, if previous movements occurred below the target area, the probe was paired with a predetermined electrode at the top of the MEA. Fetz and co-workers (Jackson et al., 2006) provided evidence in vivo of not only the induction of pathway plasticity, but of directional pathway plasticity: they repetitively stimulated a neuron in the primate motor cortex 5 milliseconds after the occurrence of an action potential on a different poly-synaptically connected neuron using a chronically implanted neural interface. After halting the stimulation, subsequent activity of the recorded neuron caused an increase in the firing rates in the vicinity of the stimulated neuron. In this manner, we hypothesized that the PTS would lead to potentiation of the probe response in the vicinity of the second paired electrode. In other words, a directional plasticity could arise during application of PTS, potentiating the pathway from the neurons evoked near the probe electrode to the neurons later evoked at the second paired electrode. This would modify subsequent CAs and population vectors in response to probe stimuli such that arm movements would approach the target area.

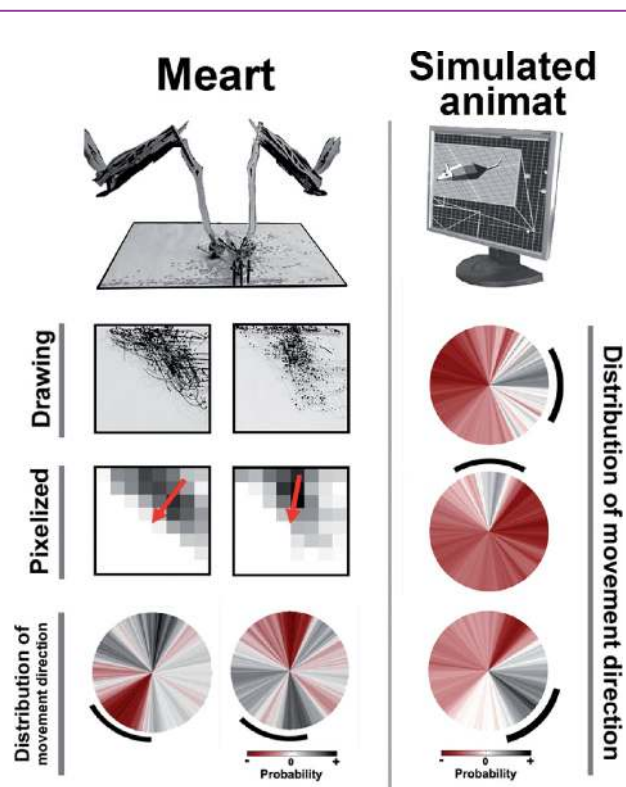


Figure 7. Plastic changes in MEART and animat behavior. *Unsuccessful and successful training of goal-directed animat behavior. MEART.* Training with predetermined PTS caused a shift in the probability distribution of commanded movement directions in two experiments (circles, bottom row), but in an uncontrolled manner. Marks first accumulated on a side of the drawing's workspace (CCD camera image of the drawing and pixelized feedback), but successful PTS training should shift the markings back toward the center (red arrow middle row; black arc bottom row). The probability distribution of movement directions during 10 minutes at the start of 2 hour experiments was subtracted from that during the final 10 minutes, thus allowing negative values (red). **Simulated animat.** Iteratively updating the probability of selecting a given PTS for training allowed an animat to learn to move in multiple directions (circles; see Methods: Making the Semi-Living Artist). Desired angles of 0, 90, and -45 degrees (black arcs) were applied in consecutive 2 hour periods. Successful behavior was considered to be movement within the desired angle ±30 degree. Notice the changes in probability distribution of movement direction were now more likely to be in the appropriate direction and more focused than for MEART.

While successful behavior did not occur (Figure 7), neural plasticity did (Figures 7 and 8), suggesting training stimuli had the potential to modify behavior. Normalized plasticity was defined as the difference in distribution of movement-controlling output (the CAs) in a given 10 minute period (CAPost) to those of the first 10 minutes (CAPre) as:

$$\text{Normalized change} = \frac{\text{Mean} \left\{ \left\| \bar{\vec{C}}_{\text{A}_{\text{Post}}} - \bar{\vec{C}}_{\text{A}_{\text{Pre}}} \right\|^2 \right\}}{\text{Variance} \left\{ \bar{\vec{C}}_{\text{A}_{\text{Pre}}} \right\}} = \frac{\sum_{\text{Post}} \left\| \bar{\vec{C}}_{\text{A}_{\text{Post}}} - \bar{\vec{C}}_{\text{A}_{\text{Pre}}} \right\|^2}{\sum_{\text{Pre}} \left\| \bar{\vec{C}}_{\text{A}_{\text{Pre}}} - \bar{\vec{C}}_{\text{A}_{\text{Pre}}} \right\|^2} \quad (3)$$

where $\bar{\vec{C}}_{\text{A}}$ is a mean of CA vectors. A value of 1 indicates no change.



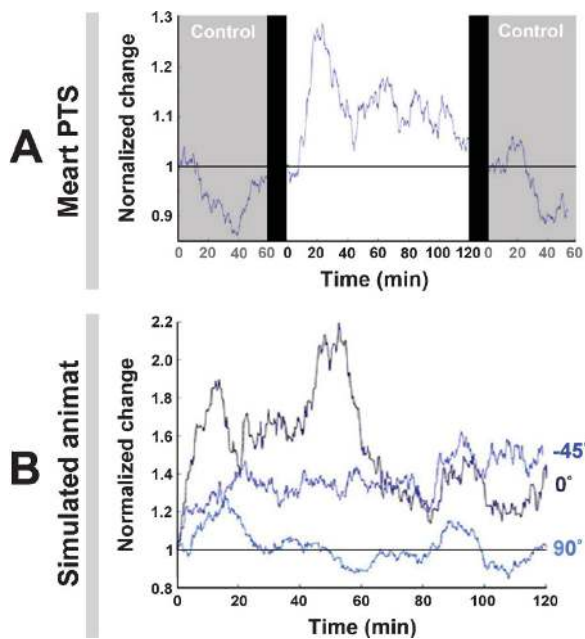


Figure 8. Neuronal plasticity. **A.** An experiment with MEART (data is the same as Figure 7, left) run for 2 hours and compared to 1 hour probe-only periods before and after. “Normalized change” is a comparison of the movement outputs (the CAs) in any 10 minute period to those of the first 10 minutes. At time=0, the same periods were compared, giving no change (a value of 1), and the 10-minute window for subsequent values was stepped by 1 minute. The drop below 1 in the control periods meant the variability in CAs decreased, possibly indicating a habituation to the stimulation. The addition of training stimuli caused plasticity, but not behavioral success (Figure 7). **B.** The experiments with the animat (data is the same as Figure 7, right) run for 2 hours. The adaptive training algorithm caused plasticity. For 90 degrees, change hovered around 1 because this was the direction of bias, a 60 degree/360 degree chance.

We concluded that since neurons at different electrodes are connected through multiple intermediate neurons and pathways, the effect of a given PTS cannot be predicted. By using feedback of behavioral performance to select and refine effective sequences of PTSs, instead of using MEART’s fixed PTSs, the simulated animat could now achieve its goal-directed behavior (Figure 7). To demonstrate that the successful behavior was a consequence of the biological changes in the neural network and not an artifact of the algorithms, the desired movement angle was switched between three angles every 2 hours. Even though movement was commanded by absolute location for MEART and incremental movement for the animat, training was intended to produce the same effect on neural activity: shift the distribution of CAs (and in turn movement angles) toward a desired goal direction.

The adaptive training algorithm allowed a search for “solutions” to achieve goal-directed behavior (Figures 7 and 8). Some PTSs may give desired neuronal plasticity while others may give the opposite or none. Furthermore, a neural network is continuously plastic, and the same PTS may have different effects at different times. The training algorithm commanded the application of a sequence of PTSs to produce the appropriate neural plasticity for successful adaptation. The learning curve in Figure 9 shows the percentage of successful movements in time; progressively fewer PTSs were needed to maintain the desired behavior, suggesting that the animat was learning the appropriate behavior.

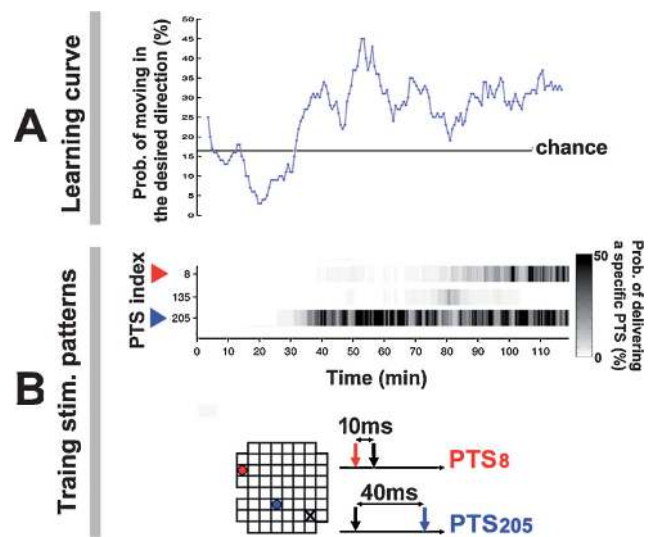


Figure 9. Training series and learning curve for the simulated animat. Animat learning curve and training history in living culture. (Data from the –45 degree desired angle trial in Figure 7). **A.** A greater portion of animat movement was toward the desired direction after 30 minute. An animat moving randomly would give a 16.67% chance that the movement was within the ± 30 degree range of the desired angle (60 degree/360 degree). **B.** Training was designed to select the PTS that induced appropriate neural plasticity as determined by subsequent animat behavior. The improved performance at 30 minute corresponded to an increase in the occurrence of PTS205, whose paired pulse pattern is shown below; its electrode location is shown in the 8×8 grid (blue dot; the probe electrode is a black X). A different PTS pattern increased the RBS occurrence at 80 minute (red).

DISCUSSION

“To view Meart is to witness a collage of contradictions. It offers us the actual biological substance of the thinking brain yet out of its biological context and system of developmental ordering. What is visible to us as Meart in the space of public display is a visualization of and/or window into ongoing experiments occurring thousands of miles away in a laboratory. The outcomes are neither pre-defined, nor are their meanings fully understood. Indeed, any of the aforementioned skeptical questions place us as viewers firmly in the midst of vigorous scientific debates—a fact underscored by the “real-time” nature of the Meart performance.

Like a work of science fiction, Meart stimulates broad inquiry into our own lived contexts. However, unlike sci-fi, it is not simply a representational text, but also an operational one. It cannot be dismissed as a mere illustrative flight of fancy, but must be interrogated as a concrete example. Meart is an ‘operational fiction’—a cyborg of representation and reality, art and science, and of course flesh and transistor.”

—Paul Vanouse, Excerpt from the Strange Attractors exhibition catalog, Zendai gallery, Shanghai, 2006.

Gallery visitors were first captivated by the aesthetics of the kinetic sculpture. MEART’s organic movement and the “breathing” sound of the pneumatic relay valves intermittently popping and hissing, not quite structured and not quite random, gave an intriguing sense of calm, maybe similar to watching trees sway in a gentle breeze. This hinted at the presence of an underlying natural process. A subtle curiosity to figure out what was happening turned into apprehension of the uniqueness of this semi-living artist, and then intense questions about the nature of the mind, the body, life, and about the artistic and scientific messages.

Art versus science

In our society, art and science are usually categorized into distinct disciplines. Humans are very adept at forming categories, and this is useful in making sense of the world, but convention is tailored by culture's current mood. The wide influence of 15th century artist and scientist Leonardo Da Vinci gives reason for pause and reminds us of the many connections between the artistic and scientific. After working on MEART, we have come to appreciate that both developing a work of art and making a scientific discovery require a curiosity and a passion to find new ideas, an ability to recognize a void in human understanding, and the creativity to form a solution. Does this comprise the "mysterious" in Einstein's quote?

Of course, tensions exist. The scientist needs to add precision and controllability to the project, then objectively document the results, constraints an artist may consider extraneous. In turn, the artist needs to conceptualize the project's importance and perfect its aesthetics, details a scientist may consider superficial. However, art and science also share the same goal: to expose new perspectives or forgotten truths about the world—to expand wisdom. Their presentation differs, but viewing an object of study from multiple angles broadens perspectives to new, possibly fertile ground. Exposure to the other's discourse can lead to a clash of cultures, but also a mirror to critically reassess one's own perspective.

If nothing else, MEART certainly got artists thinking more about science, and scientists thinking more about art. Since 2002, "MEART, the semi-living artist" has exhibited at galleries in Shanghai, Moscow, Atlanta, Melbourne, Bilbao, New York, and Perth (<http://www.fishandchips.uwa.edu.au/exhibitions.html>), often as part of larger exhibitions that focused on the use of new technology in art. The galleries became laboratories, as exhibitions were nearly the only time when experimentation was possible, and the scientific method became performance art. MEART has been presented at scientific conferences on artificial intelligence, neuroscience, and bioengineering in Switzerland (50th Anniversary Summit of Artificial Intelligence, Monte Verita, Switzerland, July 9–14, 2006), Germany (Embodied Artificial Intelligence, International Seminar, Dagstuhl Castle, Germany, July 7–11, 2003), Italy (European School of Neuro-IT and Neuroengineering, Genova, Italy, June 13–17, 2006), and France (29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Lyon, France, August 23–26, 2007), in addition to numerous other lectures to scientists and college students.

Embodiment and intelligence

The desire to breathe life into sculpted clay, or today into silicon microchips, has been around for thousands of years (Kac, 1997). This desire in part formed the scientific fields of artificial intelligence, cognitive science, and robotics. Their inquiries into the nature of intelligence began in the middle of the last century without a concern for its substrate: intelligent thought was considered the manipulation of abstract concepts. Digital computers have accomplished impressive feats, solving equations and defeating chess champions by relaying bits of information through discrete logic gates within nanoseconds. However, intelligence has not yet been attributed to computers or the robots they have been used to control. Tasks trivial to humans have proven difficult for computers such as adaptation, pattern recognition, fault tolerance, etc. This is likely due to significant differences in computational implementation, with brains using massively parallel processing, feedback loops on many scales, and components that learn and change function (Potter, 2007). Early predictions of how digital computers would change society were limited to things like calculators and the control of traffic lights. They did that, but obviously have embedded themselves in almost every aspect of our modern lives and technology. A better understanding of biological intelligence is expected to have its own presently unimaginable impact.

Now becoming more accepted by scientists is the hypothesis that intelligence is not disembodied, but intimately entwined with the

mechanics of the body and an interaction with the environment (Clark, 1997; Pfeifer and Bongard, 2007; Varela et al., 1993). The act of walking combines roles for neural signaling, proprioceptive feedback, the spring tension of muscles, the friction of shoes contacting pavement, and gravity to assist leg swing: both our brains and bodies were designed to take advantage of the physics in the world. With MEART and also biological movement, the presence of friction improved precision and stability by damping overshoot. MEART's muscles and other nonlinear components were not considered negatives, and our experiments tested the neuronal network's ability to learn the dynamics of its body to achieve goal-directed behaviors.

So MEART is embodied and *situated* in the real world. Does MEART manipulate abstract concepts of the external world in its small brain of a few thousand neurons? We doubt it, agreeing with the anti-representationalist stance of Neil Manson and his interpretation of our work, whether the cultured network is embodied in a simulated neurally controlled animat or an actual robot

"Anti-representationalist theorists propose an alternative model: an embodied agent conception of cognition (Clark, 1997; Franklin, 1995; Varela et al., 1991). On this conception the creature is viewed as part of the causal flux of its environment. Its success in satisfying its needs depends upon its competence in shaping its trajectory through the environment. Successful action requires creatures to use the information present in their environment (i.e., the causal regularities that actually obtain in their environment). This does not require the formation of an internal representation of the environment, it simply requires the creature to stand in the right kind of causal relations to its environment. Cognition on this view is an embodied, situated affair."

"The NCA [neurally-controlled animat] experiment has its background in this model of cognition. Earlier, I talked of the cognitive aspirations of the Potter Group. This can be read in two ways. If we assume the traditional model of cognition, the NCA methodology will only be of use for cognitive neuroscience if the cluster of neural cells gives rise to internal representations of the virtual environment. If we reject this model and situate the NCA methodology in its proper home—artificial life and embodied-agent AI—the cognitive aspirations look quite different. Some of the explananda of cognitive neuroscience (e.g., the brain's role in learning, adaptive behaviour, and linking perception and action) are amenable to embodied-agent modelling, and this is exactly what the Potter Group seem to be doing with the NCA experiment. On this second interpretation it need not be assumed that the neural cells subserve internal representations of the objects in the artificial world," (Manson, 2004).

A natural extension of embodied and situated AI is the use of external tools to scaffold intelligence (Clark, 1997). People have learned to extend memories with photographs, social networks with cell phones, vision with telescopes, and more. Ever since humans used sticks and stones to represent and keep track of things, we have been expanding our intellects with technology. The distinction between the technology and the biology that defines us as modern humans is becoming more ambiguous as some of this technology penetrates our skin (Clark, 2003). Many humans now live symbiotically with heart pacemakers and cochlear neural interfaces, and extend their life spans with medicine. MEART continues this conversation and further questions the body space of living agents by including the internet as part of its nervous system: its biological brain and artificial body were often located on different continents. This placed limitations on how "real-time" its responses to sensory input could be.

On the other hand, behavior is constrained by the limitations of the brain and the body. With MEART, movement was confined to a two-dimensional plane and constrained by the machine's speed and accuracy. The choice of how to map neuronal activity into motion and sensory






Figure 10. Does MEART create valuable art? MEART drawing and notes from an early accuracy test. This and four other drawings were purchased by a museum in Spain (MEIAC, Museu Iberoamericano de Arte Contemporânea) for their permanent collection.

feedback into electrical stimuli constrains which neuronal plasticity mechanisms could be observed behaviorally. This can be an advantage if investigating an individual mechanism or a disadvantage by limiting the available neuronal computational capacity. We might find that as we enhance the behavioral repertoire of MEART, we can study increasingly complex aspects of neural processing in its brain, perhaps eventually ones that underlie behaviors people regard as intelligent.

The nature of art and being an artist

MEART has many of the characteristics of a “real” artist. It lives, it dies, it leaves behind a body of work for others to contemplate, but can rat neurons and a mechanical body be labeled an artist? Maybe MEART is disqualified by being man-made. However, fillings for cavities in teeth and artificial hips make people part man-made, but no less human. MEART would have to be disqualified in some other sense. Does it possess sufficient creativity and intelligence to produce a work of art? Maybe not, but if so, would this suggest art is not solely a human endeavor; have we made an artist? If it possesses intention, maybe we have infringed on its intellectual property rights when drawings were purchased by a gallery (as discussed in Hughes, 2007) (Figure 10). Will the training algorithm enslave biology in order to steal from it? Or are such goals natural: does the body enslave the brain in order to live, by demanding it learn how to find and eat food?

Of course, MEART is a primitive construction, and much scientific/philosophical/artistic inquiry remains to be done. But the continued merging of biology and technology give substance to such questions. The answers given for the potential offspring of the MEART project maybe more controversial. For now, the tangible debate centers on what is the creative output: the drawings, the machine (if so then why not the brain?), a performance piece, conceptual art, or the system as a whole.

Fear and the future: Living with the semi-living

“Within thirty years, we will have the technological means to create superhuman intelligence. Shortly thereafter, the human era will be ended.”

Vernor Vinge—1993 essay “The Coming Technological Singularity”

After addressing viewer’s typical first questions during exhibitions: “Is it alive?”, “Is it thinking?”, “Is it creating art?” (“Partly.”, “That is the scientific question.”, “What do you think?”), a next question is often “Will this turn into Terminator II?”, a robotic harbinger of the apocalypse in a doomsday movie. One of the goals of MEART is to provide a public forum for education and dialog to address “fear of the unknown” and to critically examine the paths to be paved by biotechnology: we are more and more *living with* the semi-living as our artifacts become more life-like, and *becoming* the semi-living as we incorporate more technology into our bodies. Further understanding of biological intelligence is expected to improve artificial intelligence (Bakkum et al., 2004; Potter, 2007), but MEART remains rudimentary, and as mentioned above, digital computers and robotics lag behind the capabilities of biological agents.

The ethics of any technology lie not in the technology itself, but in how it is used. For example, nuclear energy can level cities and create a nuclear winter, or power cities and create life. Rats themselves, formerly plague-bearing and universally abhorred, have now become indispensable tools to advance science and medical technology (Burt, 2006). An understanding of how networks of neurons process information and how they can be best interfaced to achieve goal-directed behavior could influence future neural prosthetics for sensory deprived or paralyzed patients. Currently, prosthetics are being developed to restore hearing, vision, motion, and even anatomical parts of the brain itself (Berger and Glangzman, 2005). Will giving a bit back to those who have lost outweigh potential negatives and sacrificed animals?

More immediate are concerns about the continued melding of biology and technology and the role humans will have in creating life and death, especially if “semi-living” agents ever learn human qualities: intentionality, memory, irony, interpretation, creativity, etc. Moreover, the use of biology as an artistic pallet shifts art from imitation of nature to one that subsumes nature in its expression: partly alive artwork. MEART required constant care and attention. During the 2002 BEAP—Biofeel exhibition at PICA in Perth, MEART stopped moving when the neuronal culture died from insufficient environmental control (we since improved, see Methods: Making the Semi-Living Artist). The gallery went silent with the sudden realization that MEART had been somehow alive. The implications of such technology to manipulate life had been presented through the irony of a death, and

highlighted the need for compassion and a greater understanding of life (McRae, 2004). While MEART's arm can be re-animated by plugging in a new healthy neuronal network, we decided to permanently end MEART's intact-and-functioning existence, so that we could focus on developing the next semi-living artist. For MEART2, we intend to have more immediacy in the sensory-motor loop, so that gallery visitors can interact with it, and see by its behavior that they have become part of that loop, that they are an important part of the environment in which it is situated.

MEART and other hybrot provide a platform to continue philosophical inquiry and begin experimental inquiry into the fundamental makeup of intelligence, life, and existence.

CONFLICT OF INTEREST STATEMENT

We acknowledge the monetary value of MEART's creations and body, as works of art. We declare that the scientific research described was conducted in the absence of any other commercial or financial relationships that could be construed as a potential conflict of interest.

ACKNOWLEDGEMENTS

We thank all who helped make, improve, care for, and promote MEART, including Alec Shkolnik, Thomas DeMarse, Ryan Haynes, Radhika Madhavan, Peter Passaro, Boryanna Rossa, Emma McRae, SymbioticA residents and staff. We also thank ArtsWA, Australia Council for the Arts, the National Institutes of Health, and the National Science Foundation Center for Behavioral Neuroscience for funding.

REFERENCES

- Bakkum, D. J., Shkolnik, A. C., Ben-Ary, G., Gamblen, P., DeMarse, T. B., and Potter, S. M. (2004). Removing some 'A' from AI: embodied cultured networks. In *Embodied Artificial Intelligence*, F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi, eds. (New York, Springer), pp. 130–145.
- Bakkum, D. J., Chao, Z. C., and Potter, S. M. (2007). Adaptive goal-directed behavior in embodied cultured networks: living neuronal networks and a simulated model. Paper presented at the 3rd International IEEE EMBS Conference on Neural Engineering, Kohala.
- Berger, T. W., and Glanzman, D. L. (2005). *Toward replacement parts for the brain: implantable biomimetic electronics as neural prostheses* (Boston, MIT Press).
- Bi, G., and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Burt, J. (2006). *Rat. Reaktion Books*.
- Caminiti, R., Johnson, P. B., Burnod, Y., Galli, C., and Ferraina, S. (1990). Shift of preferred directions of premotor cortical cells with arm movements performed across the workspace. *Exp. Brain Res.* 83(1), 228–232.
- Chao, Z. C., Bakkum, D. J., Potter, S. M. (2005a). A more animal-like in vitro model for the study of learning and embodiment. Biomedical Engineering Society Conference, Chicago.
- Chao, Z. C., Bakkum, D. J., and Potter, S. M. (2007). Region-specific network plasticity in simulated and living cortical networks: comparison of the center of activity trajectory (CAT) with other statistics. *J. Neural Eng.* 4(3).
- Chao, Z. C., Bakkum, D. J., Wagenaar, D. A., and Potter, S. M. (2005b). Effects of random external background stimulation on network synaptic stability after tetanization: a modeling study. *Neuroinformatics* 3(3), 263–280.
- Clark, A. (1997). *Being there: putting brain, body, and the world together again* (Cambridge, MIT Press).
- Clark, A. (2003). *Natural-born cyborgs: minds, technologies, and the future of human intelligence* (New York, USA, Oxford University Press).
- Cozzi, L., Angelo, P. D., and Sanguineti, V. (2006). Encoding of time-varying stimuli in populations of cultured neurons. *Biol. Cybern.* 94(5), 335–349.
- DeMarse, T. B., Wagenaar, D. A., Blau, A. W., and Potter, S. M. (2001). The neurally controlled animat: biological brains acting with simulated bodies. *Auton. Robots* 11, 305–310.
- DeMarse, T. B., Wagenaar, D. A., and Potter, S. M. (2002). The neurally-controlled artificial animal: a neural-computer interface between cultured neural networks and a robotic body. *Soc. Neurosci. Abstr.* 28, 347.1.
- Feininger, A., and Schlatter, N. E. (2003). "Structures of nature: photographs by Andreas Feininger," University of Richmond Museums.
- Franklin, S. (1995). *Artificial minds* (Cambridge, MA, MIT Press).
- Georgopoulos, A. (1994). Population activity in the control of movement. *Selectionism and the brain* (San Diego, Academic Press), pp. 103–119.
- Hughes, R. (2007). The semi-living author: post-human creative agency. In *Architecture and Authorship*, T. Anstey, K. Grillner, and R. Hughes, eds. (London, Black Dog Publishing).
- Jackson, A., Mavoori, J., and Fetzi, E. E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. *Nature* 444(7115), 56–60.
- Jimbo, Y., Tateno, T., and Robinson, H. P. C. (1999). Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys. J.* 76, 670–678.
- Kac, E. (1997). *Art Journal*, Vol. 56, N. 3, Digital reflections: the dialogue of art and technology, special issue on electronic art, J. Drucker, ed. (New York, CAA), pp. 60–67.
- Karniel, A., Kositsky, M., Fleming, K. M., Chiappalone, M., Sanguineti, V., Alford, S. T., and Mussa-Ivaldi, F. A. (2005). Computational analysis in vitro: dynamics and plasticity of a neuro-robotic system. *J. Neural Eng.* 2, S250–S265.
- Kudoh, S. N., and Taguchi, T. (2006). Interaction and intelligence in living neuronal networks interfaced with moving robot. *Proceedings of the SPIE: BioMEMS and Nanotechnology II*, 6036, 60360S.
- Manson, N. C. (2004). Brains, vats, and neurally-controlled animats. *Stud. Hist. Philos. Biol. Biomed. Sci.* 35, 249–268.
- Martinoia, S., Sanguineti, V., Cozzi, L., Berdondini, L., Van Pelt, J., Thomas, J., Le Masson, G., and Davide, F. A. (2004). Towards an embodied in-vitro electrophysiology: the NeuroBIT project. *Neurocomputing* 58–60: 1065–1072.
- McRae, E. (2004). Department of History and Philosophy of Science, University of Melbourne, <http://www.fishandchips.uwa.edu.au/project/publications.html>.
- Meyer, J.-A., and Guillot, A. (1994). From SAB90 to SAB94: four years of animat research. In D. H. Cliff, P. Meyer, J.-A. Wilson, S. W., eds. (Cambridge, MIT Press), pp. 2–11.
- Mussa-Ivaldi, S. (2000). Real brains for real robots. *Nature* 408(6810), 305–306.
- Novellino, A., D'Angelo, P., Cozzi, L., Chiappalone, M., Sanguineti, V., and Martinoia, S. (2007). Connecting neurons to a mobile robot: an in vitro bidirectional neural interface. *Computational Intelligence and Neuroscience*, doi:10.1155/2007/12725, 1–13.
- Pfeifer, R., and Bongard, J. (2007). *How the body shapes the way we think: a new view of intelligence* (Boston, MIT Press).
- Potter, S. M., Fraser, S. E., and Pine, J. (1997). Animat in a petri dish: cultured neural networks for studying neural computation. *Proceedings of 4th Joint Symposium on Neural Computation*, UCSD, 167–174.
- Potter, S. M., and DeMarse, T. B. (2001). A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods* 110(1–2), 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. *SPIE Proc.* 4262, 104–110.
- Potter, S. M. (2002). Hybrot: hybrid systems of cultured neurons + robots, for studying dynamic computation and learning. Paper presented at the *Simulation of Adaptive Behavior 7: Workshop on Motor Control in Humans and Robots-On the interplay of real brains and artificial devices*, Edinburgh, Scotland.
- Potter, S. M., Wagenaar, D. A., and DeMarse, T. B. (2006). Closing the loop: stimulation feedback systems for embodied MEA cultures. In *Advances in Network Electrophysiology using Multi-Electrode Arrays*, M. Taketani, and M. Baudry, eds. (New York, Springer), pp. 215–242.
- Potter, S. M. (2007). What can artificial intelligence get from neuroscience? In *Artificial Intelligence: The Next 50 years*, M. Lungarella, ed. (New York, Springer).
- Reger, B. D., Fleming, K. M., Sanguineti, V., Alford, S., and Mussa-Ivaldi, F. A. (2000). Connecting brains to robots: an artificial body for studying the computational properties of neural tissues. *Artif. Life* 6, 307–324.
- Rolston, J. D., Wagenaar, D. A., and Potter, S. M. (2007). Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience* 148, 294–303.
- Turrigiano, G. G., and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* 10(3), 358–364.
- Van Pelt, J., Corner, M. A., Wolters, P. S., Rutten, W. L. C., and Ramakers, G. J. A. (2004). Long-term stability and developmental changes in spontaneous network burst firing patterns in dissociated rat cerebral cortex cell cultures on multielectrode arrays. *Neurosci. Lett.* 361, 86–89.
- Varela, F. J., Thompson, E., and Rosch, E. (1993). *The embodied mind: cognitive science and human experience* (Cambridge, Massachusetts, MIT Press).
- Wagenaar, D. A., Pine, J., and Potter, S. M. (2004). Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J. Neurosci. Methods* 138, 27–37.
- Wagenaar, D. A., and Potter, S. M. (2004). A versatile all-channel stimulator for electrode arrays, with real-time control. *J. Neural Eng.* 1(1), 39–45.
- Wagenaar, D. A., DeMarse, T. B., and Potter, S. M. (2005a). MeaBench: a toolset for multi-electrode data acquisition and on-line analysis. 2nd Intl. IEEE EMBS Conference on Neural Engineering, 518–521.
- Wagenaar, D. A., Madhavan, R., Pine, J., and Potter, S. M. (2005b). Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *J. Neurosci.* 25, 680–688.
- Wagenaar, D. A., Pine, J., and Potter, S. M. (2006a). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci.* 7, 11.
- Wagenaar, D. A., Nadasdy, Z., and Potter, S. M. (2006b). Persistent dynamic attractors in activity patterns of cultured neuronal networks. *Phys. Rev. E* 73, 51907.1–51907.8.
- Wilson, M. A., and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* 261(5124), 1055–1058.

