Towards Epistemically Autonomous Robots: Exploiting the Potential of Physical Systems

Jon Bird, Paul Layzell, Andy Webster and Phil Husbands

ABSTRACT

The authors outline one path towards constructing interactive artworks with the potential for displaying novel behavior. They use Peter Cariani's taxonomy of adaptive robotic systems as a framework for comparing the capabilities of systems that interact with their environments. The authors then describe two examples of *structurally autonomous* systems that are able to construct their own sensors independently of a human designer. The first device, the *evolved radio*, is the result of a recent hardware evolution (HE) experiment conducted by the authors. The second device, the *electrochemical ear*, was constructed almost 50 years ago by the British cybernetician Gordon Pask. The emergent behavior in both systems is only possible because many conventional engineering constraints were relaxed during their construction. Using existing technology, artists have the opportunity to explore the potential of structurally autonomous systems as interactive artworks.

The sculptor and theorist Ken Rinaldo has surveyed the use of artificial life (ALife) techniques in the arts and argues that the greatest potential for their application lies in developing interactive artworks that go "beyond the hackneyed replicable paths of interactivity that have thus far been presented by the arts community" [1]. Rinaldo's vision is of a "cybernetic ballet of experience, with the machine and human involved in a grand dance of each sensing and responding to the other," which will result in "sculptural and virtual algorithmic manifestations that will far surpass our wildest imaginations." Motivating this vision is what Mitchell Whitelaw has identified as an essential driving force in contemporary ALife art: the "will to escape, the desire for emergence" [2]. A shared interest in emergent phenomena provides common ground between scientists and artists in the ALife community. However, "emergence" is a notoriously slippery concept, with many subtle nuances and definitions. Consequently, the term is used in different ways by ALife practitioners, which can potentially lead to misunderstandings. Although fully aware of these difficulties, we focus here on the following open question: What sort of interactive mechanisms could display emergent behavior that does surpass our wildest imaginations? This paper is primarily written from a scientific perspective, and our aim is *not* to give a definitive answer, but rather to contribute to an ongoing debate in ALife. Given the scope of this paper, we concentrate on the state of the art in autonomous robotics and can only briefly consider the aesthetic implications of applying this technology in an artistic context.

Initially we outline Cariani's taxonomy of adaptive robotic systems [3], which provides a useful theoretical framework for comparing the capabilities of machines that interact with their environments. This taxonomy has been used by both scientists and artists [4] and provides one principled approach to exploring emergent phenomena. The most powerful generative mechanisms in this hierarchy are what Cariani calls *epistemically autonomous devices*: systems that construct their own sensors and/or effectors and thereby determine their own relations with, and knowledge of, the world. In order to make this theoretical concept more tangible, we describe two concrete examples of such systems. We highlight the key properties that these devices share and that lead to their epistemic autonomy. The first device, the *evolved radio*, is one of the results of a recent hardware evolution (HE) experiment. The second device, the *electrochemical ear*, was constructed almost 50 years ago by the British cybernetician Gordon Pask.

CARIANI'S TAXONOMY OF ADAPTIVE ROBOTIC SYSTEMS

One way to describe the organization of both organisms and robots is in terms of three abstract parts: *sensors*, *effectors* and a *control mechanism* that coordinates the two. A key distinction here between a robot and other devices with sensors is that a robot has effectors that enable it to *move*. Peter Cariani has developed a taxonomy of adaptive robotic systems using this framework, which enables different organisms and robots to be compared in terms of the capabilities of these three basic parts. Each of them can be more or less flexible, thereby determining how adaptive the device is. The sensors constrain the perceptual categories of the device, and the effectors constrain the ways in which the device can act on its environment. The control mechanism coordinates the behavior of the robot by mapping perceptual categories onto actions. All three parts of the robot determine its *primitives*: the most basic building blocks of a system, which cannot be derived by combining other elements. For example, the standard model of the Khepera, a commonly used research robot, has eight infrared sensors, eight ambient-light sensors and two wheels. The maximum number of primitives in this robot is constrained by the 1,024 distinct values that each ambient-light and infrared sensor can record and the 81 distinct integer values that each of the wheel motors can be set to (+/– 40 and 0). A Khepera controller might not utilize the full range of sensorimotor primitives: for example, it might restrict the maximum forward and reverse speed of the motors. A useful analogy is to see primitives as letters of an alphabet that can be combined in different ways to form words.

Jon Bird (scientist), Centre for Computational Neuroscience and Robotics (CCNR), University of Sussex, Brighton BN1, U.K. E-mail: <jonba@cogs.susx.ac.uk>.

Paul Layzell (engineer), Hewlett-Packard Laboratories, Bristol BS34 8QZ, U.K. E-mail: <paul_Layzell@hp.com>.

Andy Webster (artist), Falmouth College of Arts, Falmouth, Cornwall TR11 4RH, U.K. E-mail: <andy.Webster@falmouth.ac.uk>.

Phil Husbands (scientist), Centre for Computational Neuroscience and Robotics (CCNR), University of Sussex, Brighton BN1, U.K. E-mail: <philh@cogs.susx.ac.uk>.

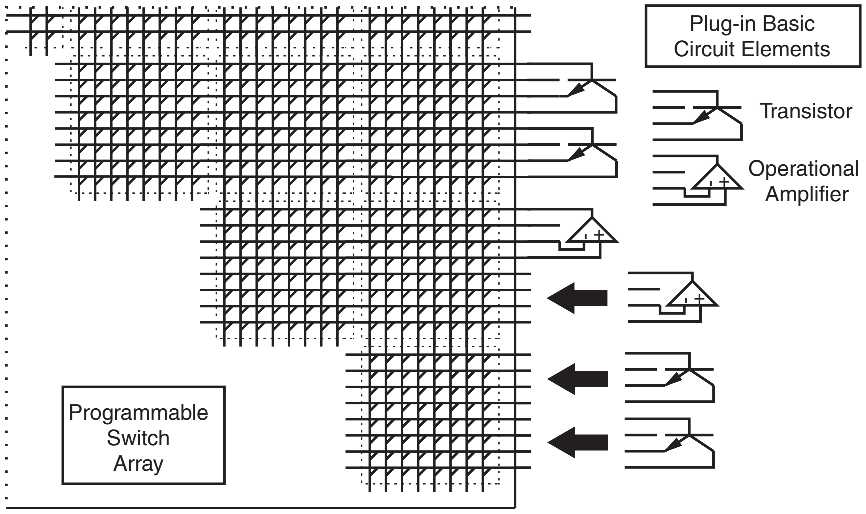

Fig. 1. Schematic of the evolvable motherboard (EM). (© Jon Bird) A genetic algorithm determines which switches are on and off and how the plug-in elements (in this case, transistors and operational amplifiers) are connected to each other.

In Cariani's taxonomy, the simplest robotic device is purely reactive and cannot modify either its internal structure, which is hardwired, or its perceptual and effector categories, which are fixed. There is therefore no means by which these *non-adaptive* devices can change their responses to particular environmental situations. The simplest *adaptive* device in Cariani's taxonomy is able to change the relations between its sensors and effectors on the basis of experience. The perceptual and action categories remain fixed, but the mapping between them can change in response to feedback from the environment. Cariani describes machines exhibiting this sort of adaptation as *adaptive computational devices.* This category includes any device that displays a capacity to learn. The most adaptive form of epistemically adaptive device in Cariani's taxonomy not only is able to change the mapping between a fixed set of primitives, but is able to create *new* primitives. Such devices do so by changing their physical structure and *constructing* new sensory, effector and control mechanisms. There are many examples of phylogenetic changes leading to the creation of new primitives. For instance, the evolution of color vision, flight and the cerebral cortex have led to new sensory, effector and control mechanism primitives, respectively. Cariani categorizes systems that demonstrate this type of flexibility as *structurally adaptive.*

Cariani's taxonomy is closely linked to different concepts of emergence. Adaptive computational devices are *combinatorially emergent* in that they can generate new combinations of existing primitives. To use the letter analogy introduced above, these mechanisms have the capacity to form new words from a prespecified alphabet. Structurally adaptive devices are *creatively emergent* in that they can generate *new* primitives. In terms of the letter analogy, they are able to expand the alphabet as well as combine the letters into new words. By constructing their own primitives, rather than using ones that we have specified, structurally autonomous systems can potentially "surpass our wildest imaginations."

We now briefly describe state-of-the-art robot technology in terms of Cariani's hierarchy and analyze how conventional design methodologies limit the structural autonomy of robots. We then describe two examples of devices that overcome these limitations and identify the key properties that lead them to display creative emergence.

STATE-OF-THE-ART ROBOTICS

Many robots are hardwired and non-adaptive, because this makes them very successful within the constraints of the niche in which they operate. For example, most production-line robots would be categorized as non-adaptive. Some research robots do have control systems that are computationally adaptive; an interesting example is a Khepera robot that learned to distinguish between objects that it could pick up with a gripper attachment and those that were too heavy solely on the basis of their size. It adaptively constructed categorical distinctions within the constraints of its fixed sensorimotor architecture [5]. There are also robots whose sensors and control mechanism are fine-tuned within fixed bounds by an adaptive process. For example, the Sussex gantry robot visual system used a number of photoreceptors determined by an artificial evolutionary process [6]. More recently, in research on an experimental locomotive device, the physical shape, the number of motors and the controller were determined using a similar approach [7]. However, no robots to date have had sufficient structural autonomy to construct their own sensors, effectors and control mechanism primitives. To understand why this is the case, it is useful to compare the conventional design methodology used to construct robots with the phylogenetic process of

Fig. 2. Photograph of the evolvable motherboard (EM). (© Jon Bird)

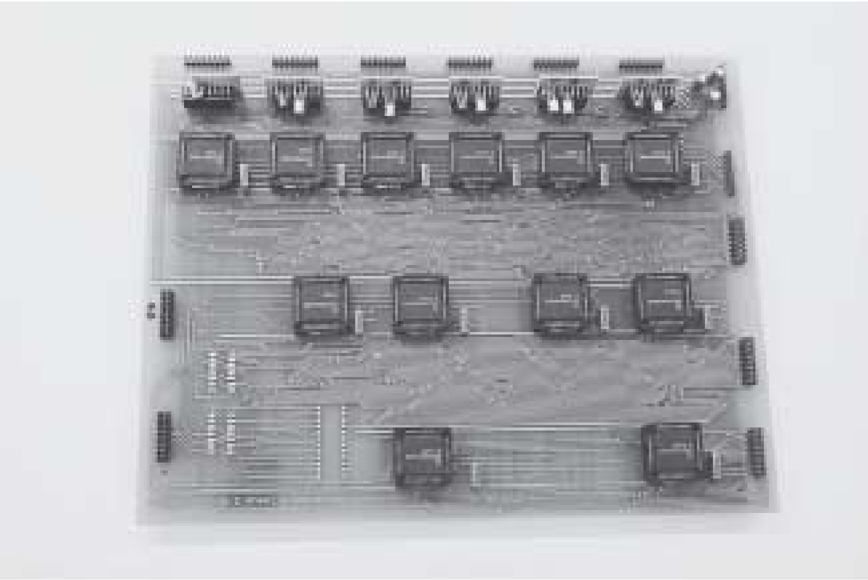

structural change that has resulted in increasingly complex organisms.

THE ENGINEERING APPROACH TO PROBLEM SOLVING

When engineers attempt to solve a problem, they are not generally constrained by previous designs and can start with a clean drawing board. The first step is to represent the problem in terms of a set of alternatives or variables whose values must be adapted to a set of fixed constraints. The goal is to find the values of the variables that maximize some utility function. The choice of variables is usually determined by a “divide and conquer” methodology: A system is functionally decomposed into semi-independent subsystems, each with separate functional roles, that interact through their functions, rather than through the physical details of their implementation. For example, in a component placing and routing problem, the goal is to position the components of a particular circuit on a board so that its size and the total wiring length are minimized. The variables in the problem are the positions of the components on the circuit board. Each component is treated as a functional subsystem such as a resistor or capacitor [8], and the interactions between them are analyzed at this functional level, rather than in terms of semiconductor physics [9]. The design of robotic systems tends to follow this engineering methodology. In the next section, we show how this limits the structural autonomy of robotic systems.

ROBOT DESIGN BY ARTIFICIAL EVOLUTION

In the engineering approach of functional decomposition, robots tend to be constructed using tried and tested off-the-shelf components. The effectors usually consist of legs or wheels arranged in one of a small number of standard configurations. There is a wider range of possible sensors, but none of the components used has much capacity for structural change. Robotic research tends to focus on designing the control mechanism that coordinates the sensors and effectors. This is usually implemented as a computer program, which can take a variety of forms within the constraints of the instruction set of the underlying machine. Writing a controller program is a complex task, which involves finding a way to coordinate sensors and effectors so that the robot behaves appropriately in the real world. One way to proceed is to use an artificial evolutionary process that searches for a near-optimal solution, given the constraints on the problem. This process is often implemented as a genetic algorithm. A population of solutions (phenotypes) is encoded as a string of numbers (genotypes). The initial population is usually randomly generated. Each phenotype in the population is tested and assigned a fitness rating. This is usually done automatically using a fitness function specified by the programmer. A new generation of solutions is generated by selecting genotypes, with a bias towards the fitter ones, and carrying out various operations on their data. The major operators are mutation (randomly changing one of the numbers in the string) and crossover (swapping numbers between two strings). This is a variant of the standard engineering methodology, in which the process of generation and testing is automated and repeated until a near-optimal solution to a well-defined problem is found.

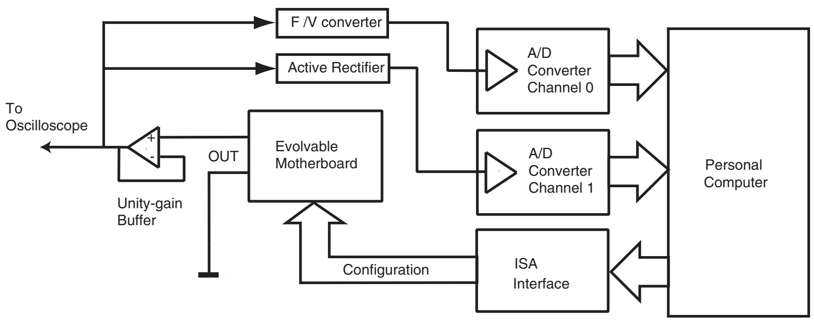

Fig. 3. The oscillator experimental setup. (© Jon Bird)

When the artificial evolutionary approach is applied to designing robot controllers, the fitness testing usually takes place in simulation. There are a number of reasons for this: first, the artificial evolutionary process can take days or weeks, and testing physical robots usually requires constant supervision, whereas a computer program can be left to run; second, in many cases it is quicker to simulate and test a design rather than to construct it; third, it is often easier to measure robot performance in a controlled simulation than in the real world.

The simulation process requires that a problem be functionally decomposed, albeit often at a more abstract level than in conventional engineering, and this leads to fundamental constraints on the design of robots. First, the designer has to model *explicitly* how different environmental stimuli change the state of the simulated sensors and how the simulated effectors act upon the environment. Second, designers simulate only those aspects of the environment that they think are relevant to their experiment; otherwise the simulation would become computationally intractable. The designer has to formalize a well-defined set of sensorimotor interactions between the agent and the environment and define the controller primitives. An artificial evolutionary process can be used to search this experimenter-defined space by testing different combinations of the prespecified primitives. This can potentially lead to interesting and surprising designs. However, such systems are limited to displaying combinatorial emergence. The next section details some hardware evolution experiments that demonstrate the conditions under which novel primitives can emerge through an artificial evolutionary process.

UNCONSTRAINED INTRINSIC HARDWARE EVOLUTION

If a design is structured on the basis of a well-specified functional decomposition, it will be subject to the same constraints whether it is tested in simulation or in the real world. In order for a system to generate new primitives, there must be aspects of the system that are *not* under the control of a designer. Unconstrained intrinsic HE is a design method in which the fitness of electronic circuits is determined by instantiating and evaluating them in actual hardware rather than in simulation. This approach relaxes the constraints adopted by conventional engineering and allows an artificial evolutionary process to explore unusual circuits with strange structures and dynamic behaviors beyond the scope of conventional design. The primitives of these circuits interact not at a functional level but through their underlying material substrate. The primitives are also free to take advantage of a wide range of environmental invariants, none of which has to be specified by an experimenter beforehand.

The experimental setup for unconstrained intrinsic HE design usually consists of a computer running an artificial evolutionary process and a reconfigurable device, such as a field programmable gate array (FPGA) [10], on which individual genotypes are instantiated as physical electronic circuits. The fitness of a given circuit is determined solely by its real-time behavior, and other factors, such as topology, are not considered. Layzell developed the evolvable motherboard (EM) to investigate some of the key issues in intrinsic HE, in particular to evaluate the relative merits of different basic components, methods of analysis and interconnection architectures. The next section gives an overview of this testbed and describes an experiment in which Layzell intrinsically evolved the first oscillators to reach their target frequency.

THE EVOLVABLE MOTHERBOARD

The evolvable motherboard (Fig. 1) is a 48 × 48 triangular matrix of reconfigurable analog switches, into which desired circuit primitives for evolution can be inserted. Any component, from transistors [11] to function-level integrated circuits, may be used. The matrix is designed to provide the minimum number of switches necessary for every combination of interconnection between primitives to be configured. By the appropriate choice of genotype-to-phenotype mapping, more restrictive or less restrictive interconnection architectures can be investigated.

The EM is configured by specifying the on/off state of each switch in turn. The full complement of approximately 1,500 switches can be configured in less than 1 ms by a host computer. The analog switches themselves are semiconductor devices contained within integrated circuits. They behave like low-value resistors but also exhibit a small degree of capacitance and inductance and may therefore have an active function in any evolving circuit. The EM has revealed fundamental differences between evolved and conventional electronics, including inherent fault tolerance in populations of evolved circuits [12] (Fig. 2).

OSCILLATOR EXPERIMENTS

In conventionally designed oscillator circuits, the necessary timing is supplied by a capacitor whose charge release is controlled by a resistor; this combination of components is known as a resistor-capacitor (RC) time constant [13]. As the desired frequency decreases, the value of the RC product increases. Large-value capacitors are difficult to implement in VLSI [14] and are generally provided externally, at some expense. Layzell's motivation was to evolve an oscillator of a precise frequency *without* using capacitors. However, oscillator evolution is a difficult task when the basic components are transistors. Whereas oscillation is the likely outcome of recurrent loops of primitives such as digital gates, precise operating points must be established before it can be produced by a network of transistors. These conditions are extremely unlikely to occur by chance, a fact that was confirmed by Layzell when he performed some preliminary experiments in which only frequency and amplitude of oscillation were rewarded. Therefore, he found it necessary to reward output amplitude, even if the signal was just noise, in order to kick-start the artificial evolutionary process (Fig. 3).

After the genotypes had been instantiated as circuits, there was a 5 ms delay to allow them to stabilize. Out of 20 runs, 10 resulted in successful oscillation, attaining the target frequency within 1% and with minimum amplitude of 100 millivolts (mV). These results represent the first intrinsically evolved oscillators to reach their target frequencies.

THE EVOLVED RADIO

Some of the circuits achieved high fitness, but when examined with an oscilloscope, they did not oscillate stably: the signals were of the order of a 10–50 mV amplitude with rapidly fluctuating frequency. The evolutionary process had taken advantage of the fact that the fitness function rewarded amplifiers, even if the output signal was noise. It seems that some circuits had amplified radio signals present in the air that were stable enough over the 2-ms sampling period to give good fitness scores. These signals were generated by nearby PCs in the laboratory where the experiments took place.

In order to pick up radio signals, the circuits needed an aerial, or antenna, and an extremely high input impedance [15]. We achieved this to create the evolved radio by using as an input the printed circuit-board tracks on the EM, which were connected to an open programmable switch whose impedance was at least 100 MΩ. The high impedance was confirmed by an electrometer behavior observed in many of the non-oscillating circuits: if a person's hand came close to the circuit, then the output voltage rose; if the hand remained there, the output voltage remained high, but it fell if the person was grounded. The evolutionary process had utilized not only the EM's transistors, but also the analog switches and the printed circuit boards (PCBs) to which they were connected.

The fact that the circuits sometimes utilize very particular environmental conditions and component properties means that they do not always generalize well. If we constrain the evolutionary process then we can make the circuits more robust; however, we also lose any of the possible advantages of unconventional

Fig. 4. Photographs of a thread growing on a piece of filter paper. (© Jon Bird) The thread is growing radially from the central copper cathode out to a circular copper anode. See Fig. 5 for an overview of this experimental setup.

design, one of which is the construction of novel primitives.

The EM is the second experimental system ever to construct novel sensors through a process of creative emergence. We now describe the first device to construct its sensors in a way analogous to the tinkering process of natural evolution: Gordon Pask's electrochemical ear. We then highlight the key properties that it shares with the evolved radio and that enable the construction of novel primitives.

Fig. 5. The experimental setup for growing the thread shown in Fig. 4. A central copper cathode is surrounded by a circular anode. The solution is copper sulfate mixed with sulfuric acid.

Pask's Electrochemical Ear

In 1958, Gordon Pask demonstrated a number of remarkable electrochemical systems that were able to construct novel sensor and controller primitives and thereby determine the relations between their own states and the environment [16]. Such systems consist of a number of small platinum electrodes inserted in a dish of ferrous sulfate solution and connected to a current-limited electrical source. Depending on the activity of the system, these electrodes can act as sinks or as sources of current. Metallic iron threads tend to form between electrodes where maximum lines of current are flowing. These metallic threads have a low resistance relative to the solution. Current will thus tend to flow down them if the electrical activation is repeated. Consequently, the potentials at the electrodes are modified by the formation of threads. If no current passes through a thread, then it tends to dissolve back into the acidic solution. The system therefore consists fundamentally of two opposing processes: one that builds metallic threads out of ions on relatively negative electrodes (sinks) and one that dissolves metallic threads back into ions. The trial-and-error process of thread development is also constrained by the concurrent development of neighboring threads and by previously developed structures. Slender branches extend from a thread in many directions and most dissolve, except for the one following the path of maximum current. If there is an ambiguous path, then a thread may bifurcate. As the total current entering the system is restricted, threads compete for resources. However, when there are a number of neighboring unstable structures, the threads can amalgamate and form one cooperative structure. Over time, a network of threads can form that is dynamically stable: the electrochemical mechanism literally *grows* (Fig. 4 and Color Plate A No. 1).

It is possible to associate some of the electrodes with output devices that enable a user to assess the behavior of the system. A reward consists of an increase in the limited current supply to the assemblage and is therefore a form of positive reinforcement. Regardless of how the electrodes are configured, the assemblage will develop a thread structure that causes current to flow in such a way that the user rewards the system. Importantly, the reward is simply an increased capacity for growth; there is no specification of what form it should take (Fig. 5).

Critically, the system is not just electrically connected to the external world: Due to the physical nature of the components, thread formation is also sensitive to temperature, chemical environment, vibrations and magnetic fields. Any of these arbitrary disturbances can be viewed as an input to the system, especially if they affect the performance of the mechanism so that its current supply is changed. The system can cause structures that are sensitive to different environmental stimuli to grow. Pask was able to train an assemblage to act as an "ear" that could discriminate between a 50-Hz and a 100-Hz tone in about half a day. He was also able to develop a system that could detect magnetism and one that was sensitive to pH differences.

Key Shared Properties

There are three key properties that the evolved radio and Pask's ear share:

- they are situated in the physical world
- they consist of primitives without fixed functional roles
- the primitives are sensitive to a wide range of environmental stimuli.

In the case of Pask's ear, the second property stems from the fact that electrochemical devices initially consist of raw material, which has no specified structure or function; in the evolved radio, this property follows from the release of electronic components from the constraints of their conventional operating ranges. The third property enables a system to utilize invariants that an experimenter may not be aware of. Given these properties, the systems' primitives are free to interact with each other and the environment through their material substrate, rather than through designer-specified interfaces, thus potentially leading to the construction of additional novel primitives. These two devices provide concrete examples of structurally autonomous systems that in one sense generate novel phenomena that *can* surpass our wildest imaginations. Whether a purely computational process can generate novel primitives is an undecided and controversial issue in ALife [17]. We now very briefly consider the implications of using epistemically autonomous devices in interactive artworks.

Some Aesthetic Consequences

The process of developing epistemically autonomous systems necessarily involves an open-ended interactive approach. It is not possible to isolate the device from its environmental context or the history of its interactions. Epistemically autonomous artworks would be less objects than active, dynamic processes. Furthermore, there will be no clear sense of a "creator" or an "author," as the artist and the system would both play active participatory roles in the creative process [18]. Analogously, when these systems are constructed and studied in a scientific context, the concept of "designer" is undermined. The steering process requires the observer to recognize developing trends in the growing device, and this involves interacting with the device in a manner akin to animal training [19], where understanding stems from interacting at a behavioral level rather than analysis at a mechanistic level. In science, we face the "inadequacy of the physical model paradigm for modelling organizations that are complex enough to themselves be observers and modelers of their world" [20]. Analogously, art generated with epistemically autonomous devices would necessarily "explore the meaning of experience rather than the meaning of works of art" [21].

CONCLUSION

Clearly, there are many different avenues that can be followed when developing interactive artworks. For example, Ken Rinaldo's own approach to achieving his artistic vision utilizes systems with hardwired controllers and fixed sensors and effectors: non-adaptive devices, in Cariani's taxonomy. In this paper we have outlined an alternative route to constructing artworks with the potential, as Rinaldo envisions, to surpass our wildest imaginations. Epistemologically autonomous robots construct their own novel sensors and effectors and thereby determine how they interact with the world, independent of any human designer. Currently there are no robots with sufficient structural autonomy to display this type of creatively emergent behavior, but we have described two devices that have constructed novel sensor—and in one case controller—primitives. The evolved radio and Pask's ear illustrate that structurally autonomous devices are not just theoretical entities: they can be implemented with existing technology and applied in an artistic context.

The development of epistemically autonomous robots will require investigating the suitability of different material substrates for constructing structurally autonomous systems. John Cage lamented that,

> when Theremin provided an instrument with genuinely new possibilities, Thereministes did their utmost to make the instrument sound like some old instrument, giving it a sickeningly sweet vibrato, and performing on it, with difficulty, masterpieces from the past. . . . We are shielded from new sound experiences [22].

Artists, unconstrained by many conventional engineering considerations, can play a significant role in exploring the genuinely new possibilities that structurally autonomous systems offer by harnessing their potential in interactive artworks.

Acknowledgment

Publication of this paper was supported in part by the Rockefeller Foundation.

References and Notes

1. Kenneth E. Rinaldo, "Technology Recapitulates Phylogeny: Artificial Life Art," *Leonardo* **31**, No. 5, 371–376 (1998).

2. Mitchell Whitelaw, "Tom Ray's Hammer: Emergence and Excess in ALife Art," *Leonardo* **31**, No. 5, 377–381 (1998).

3. Peter Cariani, "Some Epistemological Implications of Devices That Construct Their Own Sensors and Effectors," in F. Varela and P. Bourgine, eds., *Towards a Practice of Autonomous Systems* (Cambridge, MA: MIT Press, 1992). Cariani's publications can be downloaded from <http://www.cariani.com>.

4. See, for example, Whitelaw [2]; Jon McCormack and Alan Dorin, "Art, Emergence and the Computational Sublime," in A. Dorin, ed., *Proceedings of Second Iteration* (Melbourne: CEMA, 2001).

5. Rolf Pfeifer and Christian Scheier, "Sensory-Motor Coordination: The Metaphor and Beyond," *Robotics and Autonomous Systems* **20** (1997) pp. 157–158.

6. Phil Husbands et al., "Better Living through Chemistry: Evolving Gasnets for Robot Control," *Connection Science* **10**, Nos. 3–4, 185–210 (1998).

7. Hod Lipson and Jordan B. Pollack, "Automatic Design and Manufacture of Robotic Lifeforms," *Nature* **406** (2000) pp. 974–978.

8. A *resistor* impedes current flow, and combinations of them are normally used to set the voltage at a given point in a circuit. A *capacitor* is essentially two parallel plates separated by a thin insulator (or dielectric) and is used to store charge. The larger the plates' area, the greater the charge stored.

9. A *semiconductor* is a substance, usually a solid chemical element or compound, that can conduct electricity under some conditions but not others, making it a good medium for the control of electrical current.

10. A Field Programmable Gate Array (FPGA) is a VLSI circuit (see Ref. [14]) that consists of primitives such as logic gates that can be rapidly and repeatedly configured by a computer.

11. A *transistor* is normally used as an electronic switch or an amplifier, but it is more helpful here to think of its physical structure: a sandwich of different types of semiconducting material possessing myriad physical properties.

12. This research is described in depth in Paul Layzell, *Hardware Evolution: On the Nature of Artificially Evolved Electronic Circuits* (Ph.D thesis, School of Cognitive and Computing Sciences, University of Sussex, 2001), which also provides full technical details on the EM.

13. The pitch or frequency of oscillators is governed by a *time-constant*. In an acoustic oscillator such as a piano, the time-constant is determined by the thickness of the string and its tension. In an electronic oscillator, the time-constant is usually determined by the discharge of a capacitor through a resistor. The larger the value of either, the lower the frequency.

14. Very Large Scale Integration (VLSI) circuits contain hundreds of thousands of transistors.

15. *Impedance* is an expression of the opposition that an electronic component, circuit or system offers to alternating and/or direct electric current.

16. Gordon Pask, "Physical Analogues to the Growth of a Concept," in *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory on 24–27 November, 1958* (London: H.M.S.O, 1959). A useful web resource on Pask is <http://www.pangaro.com/Pask-Archive/>.

17. See Margaret Boden, "Introduction," in M. Boden, ed., *The Philosophy of Artificial Life* (Oxford, U.K.: Oxford Univ. Press, 1996). This gives a good overview of the key philosophical debates in ALife.

18. See Christa Sommerer and Laurent Mignonneau, "Art as a Living System: Interactive Computer Artworks," *Leonardo* **32**, No. 3, 165–173 (1999). This paper focuses on the aesthetic implications of interactive art.

19. See Pask [16].

20. Howard H. Pattee, "The Problem of Observables in Models of Biological Organizations," in E.L. Khalil and K.E. Boulding, eds., *Evolution, Order and Complexity* (London: Routledge, 1996).

21. Jeff Kelley, ed., *Essays on the Blurring of Art and Life/Allan Kaprow* (Berkeley: University of California Press, 1993).

22. John Cage, *Silence: Lectures and Writings* (London: Marion Boyars, 1968).

Manuscript received 4 June 2002.

ERRATUM

Corrections to "Towards Epistemically Autonomous Robots: Exploiting the Potential of Physical Systems" by Jon Bird, Paul Layzell, Andy Webster and Phil Husbands

It has been brought to our attention that some final editorial changes made to the paper "Towards Epistemically Autonomous Robots: Exploiting the Potential of Physical Systems" by Jon Bird et al. (Leonardo Vol. 36, No. 2, 2003) prior to publication introduced ambiguity into the text in one instance and in another resulted in a misrepresentation of an experiment that is a central part of the paper. The editors wish to express our regrets for these unfortunate errors, and we publish here corrections intended to clarify the text and rectify the errors.

Page 110, first paragraph, line 20:

"The most adaptive form of epistemically adaptive device in Cariani's taxonomy . . . "

should read:

"The most adaptive form of device in Cariani's taxonomy . . . "

Page 112, in "The Evolved Radio" section, second paragraph:

"In order to pick up radio signals, the circuits needed an aerial, or antenna, and an extremely high input impedance. We achieved this to create the evolved radio by using as an input the printed circuit-board tracks on the EM, . . . "

should read:

"In order to pick up radio signals, the circuits needed an aerial, or antenna, and an extremely high input impedance. The circuits achieved this by using as an input the printed circuit-board tracks on the EM, . . . "

Page 112, "The Evolved Radio" section, paragraph 2, final sentence:

"The evolutionary process had utilized not only the EM's transistors, but also the analog switches and the printed circuit boards (PCBs) to which they were connected."

should read:

"The evolutionary process had utilized not only the EM's transistors, but also the analog switches and the printed circuit board (PCB) to which they were connected."

Page 114, first full sentence:

"Analogously, art generated with epistemically autonomous devices would necessarily 'explore the meanings of experience rather than the meanings of works of art' [21]."

should read:

"Analogously, art generated with epistemically autonomous devices would necessarily explore 'the meanings of experience instead of the meanings of art' [21]."

An electronic version of this article can be accessed through Ingenta Select <www.ingentaselect.com>. Click on the "browse" button, and then search for Leonardo.