The Evolved Radio and its Implications for Modelling the Evolution of Novel Sensors

Jon Bird[1] and Paul Layzell[2]

[1]Centre for Computational Neuroscience and Robotics
School of Biological Sciences
University of Sussex
Brighton BN1 9QG, UK
jonba@cogs.susx.ac.uk

[2]Hewlett-Packard Laboratories
Bristol BS34 8QZ, UK
paul_layzell@hp.com

Abstract – Sensor evolution research typically uses evolutionary algorithms (EAs) to generate sensors that near-optimally satisfy large numbers of constraints. This is qualitatively different from the phylogenetic process found in nature that has resulted, for example, in the mammalian auditory ossicles evolving from the jaw bones of amphibians and reptiles, that in turn had previously acted as gill arches in fish. This paper describes an evolvable hardware experiment that resulted in a network of transistors sensing and utilising the radio waves emanating from nearby PCs. We argue that this evolved 'radio' is only the second device ever whose sensors were constructed in a way that in key aspects is analogous to that found in nature. We highlight the advantages and disadvantages of this approach and show why it is practically impossible to implement a similar process in simulation.

I. INTRODUCTION

The molecular mechanisms underlying energy production and protein synthesis are virtually identical in all organisms. The diversity of species is evident in the multitude of different ways that organisms sense and protect themselves from changing conditions. About 5% of the molecular machinery in *E. Coli* is for sensing and motion, whereas in humans these processes constitute the majority of our bulk [1]. One of the key theoretical issues in sensor evolution research is to explain this increase in complexity: what processes lead to the development of novel sensors and effectors [2]? The hope is that theoretical insights might be applied to the engineering of robot sensors.

Sensor evolution research often uses evolutionary algorithms (EAs) to investigate the relationship between sensors and environmental conditions [2]. Typically, experiments simulate agents acting in environments and involve searching through large parameter sets to determine the near-optimal values that satisfy multiple constraints. For example, different researchers have coevolved robot sensor morphology and controllers [3,4] which can lead to insights into how particular environmental conditions affect sensory processing by agents [5]. However, the use of EAs has so far not lead to insights into how *novel* sensors evolve: sensors that transduce environmental stimuli not previously utilised by an organism.

Darwin was one of the first to observe that, "*throughout nature almost every part of each living being has probably served, in a slightly modified condition, for diverse purposes*" [6], a process exemplified by the changing role of the hyomandibular bone from a brachial structure in fish to part of the ear in mammals. This paper shows why it is practically impossible to implement an analogous process in a simulated environment. This is an important issue for two reasons. Firstly, it is necessary to be clear about the differences between EAs and natural evolution if the aim is a theoretical understanding of the evolution of novel sensors in organisms [7,8]. Secondly, from an engineering perspective, when designing robot sensors it is essential to be aware of the limitations, as well as the strengths, of particular methodologies.

In order to highlight some of the differences between EAs, as typically used in sensor evolution research, and phylogenetic processes, this paper describes a number of unconstrained hardware evolution (HE) experiments where circuits were evolved *intrinsically*: that is, their fitness was determined by instantiating and evaluating them in hardware, rather than in simulation. A key advantage of testing in real-world, physical environments is that the circuits are free to take advantage of a wide range of environmental invariants, none of which have to be specified by the experimenter at the outset. This range is further extended when the constraints adopted in conventional electronic engineering to ensure robust and predictable operation are relaxed [9]. Evolution is then free to explore very unusual designs: circuits with strange structures and intricate dynamical behaviours beyond the scope of conventional design. In unconstrained HE, the circuit primitives do not have their behaviour constrained within specific input and output ranges or by temporal coordination, nor are they restricted to playing specific functional roles. Consequently, the process of unconstrained intrinsic HE is more like tinkering than conventional engineering [10,11] and in some key aspects is analogous to natural evolution.

In particular, this paper details an unconstrained, intrinsic HE experiment where a network of transistors sensed and utilised the radio waves emanating from a nearby PC. Essentially, the EA led to the construction of a radio. This is, as far as the authors know, only the second example of a physical device whose sensors were constructed by a process analogous to that of phylogenetic change. We compare the circuit to the first device constructed in this way: Gordon Pask's electrochemical ear [12]. We argue that both of these devices display three key characteristics: they were constructed and tested in real environments; their basic primitives were not constrained to experimenter specified functional roles; and the primitives were sensitive to a wide range of environmental stimuli. We highlight the difficulties in implementing comparable processes in simulation and argue that only unconstrained physical systems situated in real-world environments can ever construct novel sensors in a way analogous to the phylogenetic process found in nature.

II. EVOLUTIONARY TINKERING

"*Evolution proceeds like a tinkerer who, during millions of years, has slowly modified his products, retouching, cutting, lengthening, using all opportunities to transform and create"* [11]. Through this process evolution has generated novel sensors "*often utilizing organs not originally 'intended' for the purpose they serve at present"* [2]. Formalizing this generative process is one of the key challenges in modelling the evolution of novel sensors.

A. Evolution of mammalian middle ear ossicles

It is instructive to trace the evolution of the auditory ossicles in the middle ear of vertebrates as this exemplifies how homologous structures play different functional roles over the course of evolutionary time.

In mammals, the function of the middle ear is to act as an impedance transformer between the low impedance tympanic membrane and the high impedance oval window of the cochlea. These two membranes are linked by the three middle ear ossicles: the malleus, incus and stapes. Without this impedance matching much of the sound energy arriving at the ear would be reflected back into the environment. The primary mechanism of impedance matching is that the area of the tympanic membrane is larger than that of the oval window: for example, in the cat it is about 35 times larger [13]. The pressure acting on the oval window is increased by the ratio of the two areas. The second mechanism is the lever action of the middle ear ossicles: the arm of the incus is shorter than that of the malleus and this causes an increased force on the stapes.

In fish, the homologue of the auditory ossicles is the hyomandibular, which was once part of the gill apparatus and then later functioned as a jaw prop [14]. In tetrapods, this bone functioned as a structural support and as a transmitter of vibrations (stapes). Gradually, the bone became finer and less attached and more and more suited to the task of vibration transmission. Mammals evolved a new joint system for the jaw and the older skeletal elements became the malleus and incus. The radical change in the function of the hyomandibular bone is a good illustration of the tinkering process of evolution; as Romer and Sturges memorably put it: "*Breathing aids have become feeding aids and finally hearing aids"* [14].

B. Contrast of Engineering and Tinkering

It is useful to compare the engineering and tinkering approaches to constructing objects. When EAs or other optimisation methods are applied to an engineering problem, such as component placement and routing, a clear goal is defined which is not necessarily constrained by previous solutions to design problems [11]. The design problem is represented by a set of alternatives (command variables in the terminology of Simon [15]) that have to adapt to a set of environmental parameters whose values are known with certainty or in terms of a probability distribution. The goal is then to find the values of the command variables that maximise the fitness (or other utility) function, given the values of the environmental parameters and any other constraints. The choice of command 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 the details of their implementation. For example, field-programmable gate arrays (FPGAs) implement Boolean logic using high-gain groups of analogue transistors that result in the output of each cell rapidly saturating high or low. In conventional electronic design the interaction of the cells and the overall behaviour of FPGAs is viewed at a functional, logic gate level, rather than in terms of transistor dynamics [16].

Engineering	Tinkering
Clear goal/plan	Often no goal/plan
Not necessarily dependent on previous designs	Uses whatever is to hand
Aims for best solution given constraints	Makes some kind of workable object
Insulates subsystems and minimises unforeseen side effects	Combines systems or transforms them for new uses

TABLE ONE – A COMPARISON OF ENGINEERING AND TINKERING

The design and implementation of computational models follows the engineering methodology. We outline some of the difficulties that this causes for simulating the evolution of novel sensors in the next section.

III. MODELS OF SENSOR EVOLUTION

Sensor evolution research investigates the relationship between sensors and environmental conditions [2]. One of the outstanding questions in this field is how to use EAs to construct sensors that are, "*able to tap new information channels in simulated and real-world (hardware) environments"* [17].

Sensor evolution research using EAs tends to use *static* fitness functions; this is very much the engineering methodology where the goal is to find a near-optimal solution to a well-defined problem. This is clearly *not* analogous to natural evolution, where the fitness landscape is dynamic and there is no clearly defined goal [18]. Some sensor evolution research tries to make the search less constrained by allowing limited changes in the dimensionality of the search space. For example, the number of sensors and/or the size of the controller, that maps sensor states to behaviour, are varied [3,5,19]. Menczer and Belew [20] argue that fitness functions should be *implicit*, for example, based on energy levels, in order to allow "*creative, 'open-ended' evolution*". However, these approaches cannot overcome a fundamental constraint in simulating sensor evolution: the experimenter sets a *bound* on the possible interactions between the agent and the environment. This is a direct consequence of the simulation process: firstly, the experimenter has to model *explicitly* how different environmental stimuli change the state of the sensors; secondly, experimenters only simulate those aspects of the environment that they think are relevant to their experiment, otherwise the simulation would become computationally intractable. These constraints make it very difficult to see how there can be a simulation of the evolution of *novel* sensors, as the possible sensor/environment interactions are prespecified and cannot vary: an external observer can model the system deterministically [21]. It might be argued that a simulation can model the evolution of a novel sensor from an *agent's* perspective. However, constructing a novel sensor does not involve selecting which environmental stimulus to utilise from a prespecified finite list. Lewontin [22] points out that the world can be partitioned *a priori* into an infinite number of ecological niches but that we can only know which of these partitions *are* niches by the presence of an organism. The same argument holds for environmental stimuli, which can only be defined by reference to an organism. Novel sensors are constructed when a device, rather than an experimenter, determines which of the infinite number of environmental perturbations act as useful stimuli.

The next section details some hardware evolution experiments which demonstrate the conditions under which novel sensors *can* be constructed by an EA.

IV. UNCONSTRAINED INTRINSIC HARDWARE EVOLUTION

Unconstrained intrinsic HE design usually comprises a computer running an EA and a reconfigurable device, such as an FPGA, on which individual genotypes are instantiated as physical electronic circuits. The fitness of a given circuit is determined solely by its real time behaviour and other factors, such as topology, are not considered. For example, Thompson [9] evolved a circuit on a small corner of a Xilinx XC6216 FPGA that was able to discriminate between two square wave inputs of 1 kHz and 10 kHz without using any of the counters/timers or RC networks that conventional design would require for this task. The evolved circuit contained several continuous-time recurrent loops and the timing mechanism relied on a subtle analogue property - possibly parasitic capacitance - which affected delays in the internal signal paths according to the input frequency [23]. Both the loops and the timing mechanism would have been forbidden under conventional design procedure, but the evolved circuit made more parsimonious use of the silicon.

Unconstrained, intrinsic HE therefore shows potential for the design of analogue dynamical systems that may prove more successful for certain tasks than conventional design. This approach may also lead to the discovery of novel electronic 'tricks' not yet exploited by conventional design. Layzell [24] 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 where he intrinsically evolved the first oscillators to reach their target frequency.

A. *The Evolvable Motherboard (EM)*

The evolvable motherboard is essentially a triangular matrix of analogue switches, into which daughterboards containing the desired circuit primitives for evolution can be inserted. Any component from transistors and operational amplifiers to function-level integrated circuits may be used. Each daughterboard takes up to 8 lines on the switch matrix, plus a further 8 connections to allow for various power lines and I/O which may be required by certain components. The matrix is designed to provide the minimum number of switches necessary so that every combination of interconnection between primitives can be configured. By the

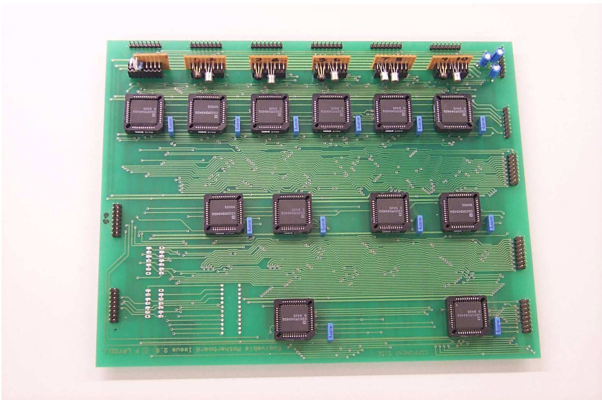

Figure 1: the evolvable motherboard (EM)

appropriate choice of genotype to phenotype mapping, more or less restrictive interconnection architectures can be investigated. The analogue switches are configured via an interface card plugged into a host PC's internal I/O ports, enabling genotypes to be instantiated in less than 1ms.

The analogue switches are themselves 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 play an active part in any evolving circuit.

B. Oscillator Experiments

There are established techniques for designing oscillators. In conventional circuits the necessary timing is supplied by a capacitor whose charge release is controlled by a resistor; this combination of components is known as an RC time constant. As the desired frequency decreases, the value of the RC product increases. Large value capacitors are difficult to implement in VLSI and are generally provided externally, at some expense. The motivation was to evolve an oscillator of a precise frequency *without* using capacitors. The tone discriminator experiment discussed above had demonstrated that evolution can make use of parasitic properties to form suitable time constants. However, oscillator evolution is a difficult task when the basic components are transistors. Whereas oscillation is the likely outcome of recurrent loops of digital gates or operational amplifiers, 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 where 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 evolutionary process.

The experiment used 10 bipolar transistors as the circuit primitives. A generational GA was used, with single point crossover, rank-based selection and elitism.

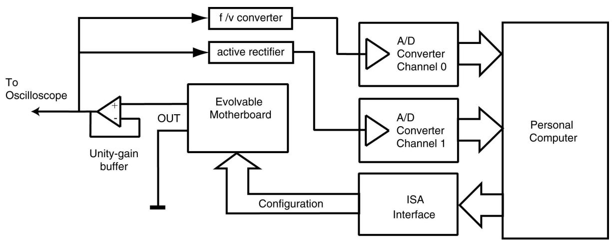

Figure 2: the oscillator experimental setup

The output of the candidate circuits was sampled directly using a hardware frequency to voltage (f/v) converter placed between the EM and an a/d converter on the host PC. This arrangement prevented aliasing errors. If a signal of amplitude greater than 10 mV is present at its input, the f/v converter outputs a d.c. voltage in the range [0, 6V] which is equal to the input frequency multiplied by a constant, *k*, whose value is determined by the midpoint of the f/v converter's range. The fitness function was as follows:

$$\text{fitness} = \begin{cases} \bar{a} + k\dfrac{f_{\min}}{f_{\max}}\left(f_{\text{target}} - \left| f_{\text{target}} - \bar{f} \right| \right), \text{if } \bar{f} > 60\,\text{Hz} \\ \bar{a}, \text{if } \bar{f} <= 60\,\text{Hz} \end{cases} \quad (1)$$

$\bar{a}$ and $\bar{f}$ represent respectively output amplitude and frequency, averaged over 20 samples, each taken at 100 µs intervals. $f_{\min}$ and $f_{\max}$ are the minimum and maximum of the 20 frequencies sampled. f_{target} is the target frequency. The ratio of the minimum and maximum frequencies rewards constant output frequency. The f/v converter's time constant was configured so that the target frequency corresponded to 3V output, ensuring that the function was smooth for frequencies above 60 Hz. This lower bound was chosen to ensure that the f/v converter was detecting oscillation and not mains hum, which is 50 Hz in the UK. The target frequency was 25 kHz.

After the genotypes had been instantiated as circuits, there was a 5 ms delay to allow the f/v converter and rectifier to stabilise. From 20 runs, 10 resulted in successful oscillation, attaining the target frequency within 1% and with minimum amplitude of 100 mV. These represent the first intrinsically evolved oscillators to reach their target frequencies.

It has proved difficult to clarify exactly how these circuits work. Probing a typical one with an oscilloscope has shown that it does not use beat frequencies to achieve the target frequency. If the transistors are swapped for nominally identical ones, then the output frequency changes by as much as 30%. A simulation was created that incorporated all the parasitic capacitance expected to exist within the physical circuit, but the simulated circuits failed to oscillate. The programmable switches almost certainly play an important role in the behaviour of the circuit and it is only possible to probe their input and output connections and not the circuitry in which they are embedded.

C. The Evolved Radio

Some of the circuits achieved high fitness, but when they were examined with an oscilloscope they did not oscillate stably: the signals were of the order of 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 need an aerial and an extremely high input impedance. This was achieved by using as an input the printed circuit board tracks on the EM connected to an open programmable switch whose impedance is at least 100 MΩ. The high impedance was confirmed by an electrometer behaviour observed in many of the non-oscillating circuits: if a person's hand was brought close to the circuit, then the d.c. output voltage rose; if the person remained there, the output voltage remained high, falling if the person was earthed. The evolutionary process had utilised not only the EM's transistors, but also the analogue switches and the printed circuit to which they were connected.

D. Other Environmental Effects

In earlier experiments Layzell [25] found that circuits utilised the oscilloscope used to measure their behaviour as a path to 0V, via the 10 MΩ impedance of the oscilloscope. If the oscilloscope was unplugged, the circuit did not work. In a SPICE simulation where the oscilloscope was represented by a resistance, the circuit worked, confirming its functional role.

Some of the evolved oscillators worked successfully until a soldering iron on a nearby workbench was disconnected from the mains, at which point oscillation ceased. This occurred despite high quality laboratory power supplies and extensive mains filtering. The circuit was apparently sensitive to tiny transients in its voltage supply. The circuit worked if it was reinstantiated on the EM, regardless of whether the soldering iron was on or off. However, tests showed that it failed to oscillate if during instantiation the programmable switches were set in a different order to that used originally. It seems that the circuit was dependent on some initial condition, such as charge, that only occurred if the switches were set in a particular sequence.

These results demonstrate that unconstrained, intrinsic HE will potentially exploit *any* physical characteristic that can influence circuit behaviour, and that these characteristics are present in the *entire* evolutionary environment. The fact that the circuits sometimes utilise very particular environmental conditions and component properties does mean that they do not always generalise well. This is also the case with many organisms that live in environments of low variability, as these niches can be effectively exploited by efficient specialisations; general solutions are only found in organisms that inhabit high variability environments [26]. If we constrain the evolutionary process then we can make the circuits more transparent, but we also lose any possible advantages of unconventional design, one of which is the construction of novel sensors.

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 which enable the construction of novel sensors.

V. PASK'S ELECTROCHEMICAL EAR

In 1958 Gordon Pask demonstrated a number of remarkable mechanisms that were able to construct novel sensors and thereby determine the relations between their own states and the environment. In other words, these devices were able to generate and explore their own state space. Any observer trying to model the behaviour of these devices would be forced to *change* the dimensionality of their model over time as the devices can transform the underlying generative system.

A. Description of the Mechanism

The devices are electrochemical assemblages consisting of a number of small platinum electrodes that are inserted in a dish of ferrous sulphate solution and connected to a current limited electrical source. Depending on the activity of the system, these electrodes can act as sinks or 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 and so current will 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 fundamentally consists of two opposing processes: one which 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 neighbouring threads and also by previously developed structures. Slender branches extend from a thread in many directions and most of these dissolve except for the one following the path of maximum current. If there is an ambiguous path then a thread can bifurcate. As the total current entering the system is restricted, threads compete for resources. However, when there are a number of neighbouring 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*.

It is possible to associate some of the electrodes with output devices that enable the behaviour of the system to be assessed by a user. 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 leads to current flowing in such a way that the user rewards the system. Importantly, the reward is simply an increased capacity for growth and there is not any specification of what form it should take.

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 grow structures that are sensitive to different environmental stimuli. Pask was able to train an assemblage to act as an 'ear' that could discriminate between a 50 Hz and 100 Hz tone in about half a day. He was also able to grow a system that could detect magnetism and one that was sensitive to pH differences. The development of sensors constitutes a change in the state space of the assemblage that was not specified by a designer explicitly.

VI. DISCUSSION

We have described an unconstrained, intrinsic HE experiment that resulted in the construction of a *novel* radio wave sensor. The EM is the second ever experimental system to construct novel sensors, unconstrained by prespecified sensor/environment channels. Like Pask's ear, the evolved radio determined the nature of its relation to, and knowledge of, the world. Both of these devices are *epistemically autonomous*: they are not restricted to experimenter specified information channels [27]. By using a process analogous to the tinkering of natural evolution, epistemically autonomous devices alter their relationship with the environment depending on whether a particular configuration generates rewarded behaviour.

We have argued that there are three key properties that devices must embody in order for selection pressure to form them into novel sensors:

- they are situated in the physical world;
- they consist of primitives with no fixed functional roles;
- and the primitives are sensitive to a wide range of environmental stimuli.

In 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 releasing electronic components from the constraints of their conventional operating ranges.

We argue that devices such as this are useful for highlighting the practical impossibility of *simulating* the evolution of novel sensors: programming a simulation necessarily involves prespecifying the possible sensor/environment interactions. Novel sensors are constructed when a device, rather than an experimenter, determines which of the infinite number of environmental perturbations act as useful stimuli. Unconstrained, intrinsic HE has provided a concrete example of such a device and is potentially a powerful approach to designing robot sensors as it enables circuits to exploit the rich dynamics of semiconductor physics and thereby explore regions of design space that are inaccessible to the conventional engineering approach.

References

[1] D.S. Goodsell. *The Machinery of Life*. Copernicus, Springer-Verlag, 1998.

[2] K. Dautenhahn, D. Polani and T. Uthmann. Guest Editors' Introduction: Special Issue on Sensor Evolution. *Artificial Life* 7: 95-97, 2001.

[3] D. Cliff, I. Harvey and P. Husbands. Explorations in Evolutionary Robotics. *Adaptive Behavior*, Vol. 2, 1, 73 – 110.

[4] P. Husbands, T.M.C. Smith, N. Jakobi and M. O'Shea. Better Living through Chemistry: Evolving Gasnets for Robot Control. *Connection Science*, 10 (3-4), 185 – 210.

[5] K. Balakrishnan and V. Honovar. On Sensor Evolution in Robotics. In J. Koza, D. Goldberg, D. Fogel and R.Riolo (Eds.), *Proceedings of the First International Conference on Genetic Programming (GP-96)*, MIT Press, 455 – 460, 1996.

[6] C. Darwin. *The Various Contrivances by which Orchids are Fertilised by Insects*. D. Appleton, 1886.

[7] E.A. Di Paolo, J. Noble and S. Bullock. Simulation Models as Opaque Thought Experiments. In M.A. Bedau, J.S. McCaskill, N.H. Packard and S. Rasmussen (Eds.), *Artificial Life VII: The Seventh International Conference on the Simulation and Synthesis of Living Systems, Reed College, Portland, Oregon, USA, 1 - 6 August*, 497 – 506, 2000.

[8] J. Bird. Can Robots Illuminate Adaptive Behaviour? In J. Kelemen and P. Sosík, (Eds.), *Advances in Artificial Life: Proceedings of the Sixth European Conference on Artificial Life*, Springer, 495 - 498, 2001.

[9] A. Thompson. An evolved circuit, intrinsic in silicon, entwined with physics. *International Conference on Evolvable Systems (ICES96)*, Vol. 1259, Springer LNCS, 640 – 656, 1997.

[10] D. C. Dennett. Cognitive Science as Reverse Engineering: Several Meanings of "Top-Down" and "Bottom-Up". In D. Prawitz, B. Skyrms and D. Westerstahl (Eds.), *Logic, Methodology, and Philosophy of Science IX*, Elsevier Science, 679 – 689, 1994.

[11] F. Jacob. *The Possible and the Actual*. Penguin Books, 1989.

[12] G. 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*, Vol. II H.M.S.O., 879 – 928, 1958.

[13] J. O. Pickles. *An Introduction to the Physiology of Hearing*. Second Edition. Academic Press, 1988.

[14] A. S. Romer and T.S. Sturges. *The Vertebrate Body*. Sixth Edition. Saunders College, 1986.

[15] H.A. Simon. *The Sciences of the Artificial*. Third Edition. MIT Press, 1996.

[16] A. Thompson. Artificial Evolution in the Physical World. In T. Gomi (Ed.), *Evolutionary Robotics: From Intelligent Robots to Artificial Life (ER'97)*, AAI Books, 101 – 125, 1997.

[17] A. Liese, D. Polani and T. Uthmann. A Study of the Simulated Evolution of the Spectral Sensitivity of Visual Agent Receptors. *Artificial Life* 7: 99 – 124, 2001.

[18] I. Harvey. Artificial Evolution: A Continuing SAGA. In T. Gomi (Ed), *Evolutionary Robotics: From Intelligent Robots to Artificial Life*, Proc. of 8th Intl. Symposium on Evolutionary Robotics (ER2001), 2217 LNCS, Springer, 2001.

[19] M. D. Bugajska and A.C. Schultz. Co-Evolution of Form and Function in the Design of Autonomous Agents: Micro Air Vehicle Project. *GECCO-2000 Workshop on Evolution of Sensors in Nature, Hardware and Simulation*, 240 – 244, 2000.

[20] F. Menczer and R.K. Belew. Evolving sensors in environments of controlled complexity. In P. Maes and R.A. Brooks (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems.* MIT press, 210 – 221, 1994.

[21] P. Cariani. Emergence and Artificial Life. In C.G. Langton, C. Taylor, J.D. Farmer and S. Rasmussen (Eds.), *Artificial Life II*, Santa Fe Studies in the Sciences of Complexity, Proceedings, 10, Addison-Wesley, 775-797,1991.

[22] R. Lewontin. Organism and Environment. In H.C. Plotkin (Ed.), *Learning, Development and Culture: Essays in Evolutionary Epistemology*, Wiley, 1982.

[23] A. Thompson and P. Layzell. Analysis of Unconventional Evolved Electronics. *Communications of the ACM*, 42 (4), 71 – 79, 1999.

[24] P. Layzell. A New Research Tool for Intrinsic Hardware Evolution. In M. Sipper, D. Mange and A. Perez-Uribe (Eds.) *Proceedings of the Second International Conference On Evolvable Systems (ICES98)*, 1478 LNCS, Springer, 47-56, 1998.

[25] P. Layzell. *Hardware Evolution: On the Nature of Artificially Evolved Electronic Circuits*. DPhil Thesis, School of Cognitive and Computing Sciences, University of Sussex, 2001.

[26] M. Conrad. *Adaptability: The significance of variability from molecule to ecosystem.* Plenum, 1983.

[27] P. Cariani. To evolve an ear: Epistemological implications of Gordon Pask's electrochemical devices. *Systems Research*, 10 (3), 19 – 33, 1993.