Evolving Minimally Creative Robots

Jon Bird and **Dustin Stokes** [1]

Abstract. We outline our motivations for using an evolutionary robotics approach to investigating creativity. We propose two necessary conditions that a robot must satisfy in order to demonstrate minimal creativity: agency; and novelty. Our preliminary results demonstrate that this methodology can produce simulated robots that mark their environments and interact with the lines that they have made. Our simulated robots satisfy a 'no strings attached' form of agency, but it is contentious whether their behaviour could be described as novel. Furthermore, it is an open question whether even as we incrementally increase the complexity of the robot controllers their behaviour will be classified as creative and their markings rudimentary drawings. However, we argue that a synthetic, bottom-up approach is a fruitful methodology for generating and testing hypotheses about creativity and generates concrete examples that can help to clarify the necessary and sufficient conditions for creative behaviour.

1 A MINIMAL ACCOUNT OF CREATIVITY

Going back to the ancients and early moderns, many philosophical theories of creativity go wrong right from the start, taking as their explanandum a radical form of creativity or genius ([18], [9], [21]). This is no less true of many contemporary philosophical and psychological theories of creativity ([8], [19], [10], [11], [12], [13], [7]). Such theories generally use a case study format, focussing on the introspective reports and historical observations of one or a few historically recognized geniuses. This loads the explanation in unhelpful ways, since both introspection and behavioural observation are fallible and, more importantly, geniuses provide undoubtedly exceptional examples of creativity. A theory would do well to start with something less complex, but nonetheless common to both radical and more mundane instances of creativity. What we are after then is a notion of minimal creativity.

1.1 Agency and autonomy

To what, at minimum, do we attribute creativity? Linguistic intuitions only get us so far, but they do suggest two necessary conditions. First, we attribute creativity only to things that result, in some non-trivial sense, from agency. Agency can be understood in radically different ways. Philosophers, for example, analyze what constitutes agency while researchers in artificial life tend to focus on the origins of agency. These differences in approach sometimes result in differences in concept use. However, the second, bottom-up approach might be understood instead as continuous with the first, that is, as an alternative way to clarify the constituents of agency. By default, we understand agency broadly. Some behaviour, artefact, or event x is the product of the agency of A only if x would not have occurred had A not acted in some autonomous way. We will not take 'autonomous' to imply 'intentional' or 'cognitive', but just to mean 'self-moving': for example, an autonomous robot displays behaviour that is not imposed by an external programmer.

The choice for construing autonomy and thus agency broadly is motivated by our underlying research project. Our ultimate goal is to see what can be learned about creativity and cognition using synthetic, bottom-up modelling techniques. It is an open question whether these techniques result in systems that possess or display robust autonomy, agency, or creativity. But the supposition that we can produce autonomous agents using a bottom-up, synthetic approach allows for fruitful methods of modelling and hypothesis generation. This, we take it, is a methodological assumption that we share with much of cognitive science.

1.2 Novelty

Linguistic intuitions also tell us that creativity requires novelty. Absolute novelty is novelty *simpliciter*, or perhaps slightly weakened, historical novelty, which is novelty given the history of ideas [1]. Boden contrasts this with psychological novelty, which is relative to some particular mind. As a conceptual decision, we opt for relative novelty. This decision invites an even trickier one, namely, selecting a suitable reference point. We are most interested in what we call *agent-relative novelty*, which is broadened from the purely psychological to include the bodily actions of a particular agent. A behaviour is novel, by this criterion, if it is new with respect to the previous behaviour of the agent. What sorts of changes underwrite this kind of novelty? Intuitions tell us that a behavioural change which depends not merely upon environmental change but upon some internal change is needed for novelty. This intuition, which may or may not be accurate, excludes purely reactive systems from acting in novel ways. We want the most for our penny, however, and so resist hasty dismissal of cheaper instances of novelty. In our initial experiments we start with reactive agents that do not undergo internal change, but that do influence aspects of their environment and in turn influence their own (purely reactive) behavior. Our initial question is whether such systems can display agent-relative novelty.

1.3 What else do we need?

There is more to be said about both agency and novelty, but we assume these two conditions as necessary for minimal creativity. Necessity aside, brief reflection reveals that satisfaction of these two conditions is likely insufficient for creativity. I can right now, in the privacy of my home, dance a little jig and sing the sentence 'Beware of the naked balloon guy' to the tune Twinkle, Twinkle Little Star. This is (and you must take this on faith) novel behaviour relative to me and is no doubt a product of my agency. But it is not obviously creative. So we will often need a third condition (or more) to suffice

[1] University of Sussex, United Kingdom, email: jonba@sussex.ac.uk

for creativity. The problem is that we do not obviously have consistent intuitions to take us beyond agency and novelty (assuming they even take us that far). In fact, it may well be that our concept of creativity is indeterminate and, more importantly, context-variant. Thus what completes the analysis depends upon certain, context-specific, theoretical decisions and commitments. We leave open what additional necessary conditions are conjointly sufficient for creativity (in some context or other). This should be no cause for concern: the two identified conditions are necessary and must be satisfied irrespective of any additional conditions one includes to complete an analysis of creativity. Thus to model creative behaviour, we must model something that satisfies, at minimum, the following conditions:

1. Agency: x is creative only if x is the product of agency;
2. Novelty: x is creative only if x is agent-relative novel.

2 DESIGNING SITUATED AGENTS IS HARD

The initial aim of our research is to produce a simple model of minimal creativity that satisfies the two necessary conditions discussed in the previous section: agency and novelty. As a first step we are focusing on the 'no strings attached' form of agency: a robotic model whose behavior is determined solely by its sensory-motor activity and is not controlled by an external observer. A robot that exhibits this minimal agency is necessarily situated. Given the difficulty of modelling robot-environment interactions, designing even simple, situated robots is a non-trivial task.

An illustration of this problem comes from Scutt [22] who carried out experiments with a Lego robot implementation of Braitenberg's Vehicle 2b [2]. The robot's simple sensory-motor system consisted of two wheels and two light sensors and its controller consisted of a connection from each sensor to the wheel motor on the opposite side. Each connection is configured so that the more its sensor is illuminated, the faster the connected wheel turns. If one of the sensors is nearer to the light than the other, then there will be a faster wheel speed on the less illuminated side and the robot will turn towards the source; if both sensors are equally stimulated then the robot will move straight ahead towards the light.

Braitenberg predicted this light seeking behaviour. Scutt found, however, that in a real world environment vehicle 2b will display more complex behaviour - avoiding any obstacles that are placed between it and the light source due to the way that shadows fall on the sensors. Scutt tested the robot with increasingly complex arrangements of obstacles until there was no light falling directly on the robot's sensors. One might reasonably predict that in this situation the robot will not move. However, in Scutt's experiments the robot moved away from the obstacles and along the wall of the testing arena, thereby avoiding all the intervening objects and eventually turning directly towards the light source. Neither Braitenberg nor Scutt predicted this behaviour, which occurred because the testing arena walls were painted white and the robot interacted with the light reflecting off them. This case illuminates both the challenge and promise of robotics: even rudimentary situated robots behave in unpredicted and interesting ways.

2.1 Evolutionary robotics (ER)

The methodology of evolutionary robotics has been successfully applied to the design of situated robots [16]. In current ER research, the fitness of candidate designs is tested either in simulation, in the real world or using a combination of the two. ER is a discovery methodology that is free to exploit any constraints arising from the interaction of components in the controller and between the robot and environment, even when the human experimenter is not aware of them. This can potentially produce simpler, more robust robots than conventional design ([15], [17]). It can also produce robots that exhibit unpredictable, and thus potentially novel, behaviour.

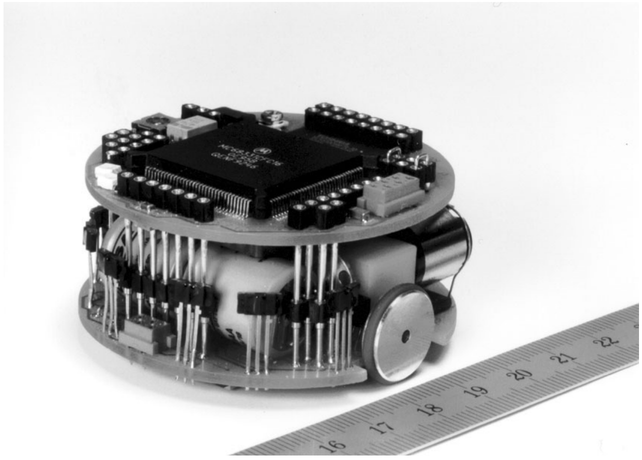

Figure 1. Khepera robot

Another feature of ER suitable to our minimal approach to creativity is that it provides a means of generating novel designs that can, to some extent, overcome inductive bias [20]. This is the phenomenon where the explicit and implicit biases of an experimenter constrain the possible space of designs that is explored. By artificially evolving control architectures from suitably low level primitives, the final controller "need not be tightly restricted by human designers' prejudices" [3, p.83]. ER therefore has the potential to produce models of minimal creativity not dominated by our preconceptions, case studies, or, perhaps mistaken, theories of creativity.

Finally, artificially evolving neural networks as robot controllers allows for open-ended evolution, since their architecture can be incrementally increased in complexity by adding processing units and connections ([6], [4]). Thus, even if robot behavior does not exhibit relative novelty at early stages in experimentation, incremental increases in neural complexity may well make the relevant difference.

2.2 Simulation

The aim of the Drawbots project is to carry out experiments on physical robots. However, there are two reasons why initially the behaviour of the robots is tested in simulation, rather than in the real world:

1. the testing procedure takes a long time and can be slowed down further due to the power requirements of robots which might require batteries changing regularly;
2. robot hardware is generally not robust enough for long periods of testing and at initial stages of the artificial evolutionary process the robots can behave in ways that damage themselves.

There are difficulties involved in using ER simulations:

1. simulating physical environments is a non-trivial task and the process of abstraction may introduce unintentional artifacts that are

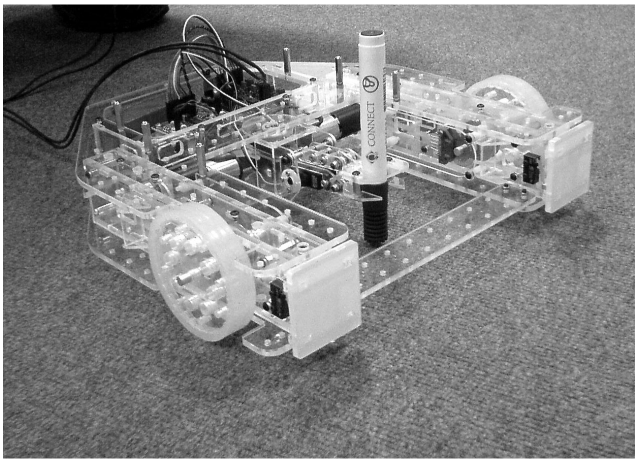

Figure 2. Prototype *Drawbot* that will used to test evolved controllers in the real world. To give a sense of scale, note that the marker pen, which can be raised and lowered, is approximately 15cm in length.

exploited by the GA, but do not hold in the real world, therefore preventing the evolved controllers from successfully transferring to real robots;

2. getting the right level of noise in a simulation is critical, as if there is a discrepancy between noise in the simulation and the real world, again, the evolved controllers will not transfer: too much can lead to the GA exploiting stochastic resonance effects (that is, weak signals are boosted); too little can make a task easier in simulation than in the real world.

The simulator used in the Drawbots experiments has been adapted from software that has been successfully used in previous ER experiments to evolve robots that successfully transferred from simulation onto physical robots ([15]. We are currently developing a custom robot platform for future real world experiments (Figure 2.).

2.3 Artificial neural networks

We have initially used artificial neural network (ANN) controllers for our robots, based on Nolfi's emergent modularity architecture [14]. The low level primitives consist of the neurons, or processing units, and the connections between them. The emergent modularity architecture has been successfully evolved to control complex robot behaviours. When evolving garbage collection robots, the experimenters demonstrated that this architecture performed better than hand-designed controllers and argued "the engineering oriented approach based on decomposition and integration can have serious limitations in the case of behavioural systems (such as mobile robots) where the observed behavior is the result of the dynamical interaction between the robot and the environment" [16, p.134].

In our experiment each robot controller consists of seven sensors (six IR and one line detector) and six motor neurons (a pair of left motor neurons, a pair of right motor neurons and a pair of pen motor neurons). At each sensory-motor cycle, the most strongly activated neuron out of each pair of motor neurons is selected to control the appropriate motor. Each sensor connects to every motor neuron, giving 42 connections in the network.

All six motor neurons in the controller are governed by:

$$z_i^t = \Phi(\sum_j w_{ij} z_j^t + b_i) \quad (1)$$

where Φ represents a sigmoid function, restricting the neuron activation values to the range $[0, 1]$, z_i^t represents the activation of neuron i at time t, w_{ij} the strength of the synaptic connection from neuron j to neuron i (range $[-10, 10]$) and b_i the bias (range $[-10, 10]$) of neuron i.

The activation of each of the six IR sensors is in the range $[0, 1.0]$ and is calculated from look up tables generated by sampling the responses of a real Khepera's sensors placed at a range of distances and different orientations from a wall. A small amount of noise is added to the measured value (range $[0, 0.05]$). The line sensor is 2mm $\times$2mm and positioned under the front of the robot. It the has two states: 0 indicating that there is no mark under any part of the sensor, 1 indicating the presence of a mark under the sensor.

A right and left integer wheel speed (range $[-10, 10]$) is generated from the corresponding motor neuron using the following formula (which shows the calculation of the right wheel speed):

$$o_r^t = (m_r^t \times 20) - 10 \quad (2)$$

where o_r^t represents the right motor signal at time t and m_r^t is the activation of the right motor neuron at time t (range $[0, 1.0]$).

If the pen motor neuron activation is greater than 0.75 then the pen is in the down position and the robot marks the arena floor when it moves.

2.4 Genetic algorithms (GAs)

Our initial experiments use a genetic algorithm (GA) a technique inspired by natural evolution and introduced by Holland ([5]) that efficiently searches through a parameter set that defines a particular optimization problem in order to find an acceptable solution. A population of trial solutions (phenotypes) is encoded as a string of symbols (genotypes). The initial population is randomly generated, with symbols on the genotype representing different values of the parameters that define the problem. In our experiments, the genotypes encode neural network parameters (the biases of the neurons and the strengths of the connections between them) which are decoded into robot controllers. Each phenotype in the population is decoded in turn from its genotype and tested and assigned a fitness. This is usually done automatically by a fitness function specified by the programmer. A new generation of solutions is generated by randomly selecting genotypes, with a bias towards the fitter ones, and carrying out various operations on their data that are inspired by evolution. The major operators are random mutation of the symbols in the string and crossover of symbols between two strings. This process of generate and test is repeated until an acceptable solution is found to the problem.

In our preliminary experiment there was a mutation rate of 0.01 per allele and we did not use crossover. Each of the robot controller parameters (42 connection weights and 6 motor neuron biases) was encoded as an 8 bit integer-valued vector (range $[0, 255]$) and mutation consisted of flipping one of the 8 bits.The population size was 100 and the experiments were run for 600 generations.

2.5 Fitness function

In keeping with our approach to reducing inductive bias, the fitness function used in our preliminary experiments does not specify the type of marks the robots should make. Instead, it rewards correlated changes in the state of the robot's line sensor (line detected or no line detected) and the pen position (up or down) within a short time window (two sensory-motor cycles). For example, if a line is detected and in the next sensory-motor cycle the pen is either raised or lowered, the individual accumulates fitness. The fitness function also rewards robots that make marks that are in separated regions of the arena. The fitness function consists of 2 weighted elements, each element in the range $[0, 1]$ and both elements summing to 1.

$$\text{fitness} \quad = \quad \alpha \frac{1}{T} \sum_{t=0}^{T} f^{t}_{corr} \times f^{t}_{movement} + \beta f_{spread} \qquad (3)$$

where α and β are weighting values (0.4 and 0.6 respectively), T is the total number of time steps in the trial, f_{corr} is determined by the correlation between changes in the state of the light sensor and changes in the state of the pen position and calculated as follows :

$$f_{corr} \quad = \quad L^{t}_{change} \times P^{t+\Delta t}_{change} \qquad (4)$$

where

$$L^{t}_{change} \quad = \quad \begin{cases} 1 & : \text{ line sensor changes state at time } t \\ 0 & : \text{ otherwise} \end{cases} \qquad (5)$$

$$P^{t+\Delta t}_{change} \quad = \quad \begin{cases} 1 & : \text{ pen changes state during } t+\Delta t \\ 0 & : \text{ otherwise} \end{cases} \qquad (6)$$

where $\Delta t = 1$, f_{spread} is the largest distance between any two of the marks made by the robot in a given trial and $f_{movement}$ is calculated over a moving time window Δt_{move} (10 sensory-motor cycles).

$$f_{movement} \quad = \quad \begin{cases} 1 & : \text{ robot moves } > R_{min} \text{ in } \Delta t_{move} \\ 0.1 & : \text{ otherwise} \end{cases} \qquad (7)$$

where $R_{min} = 5$mm. If the robot crashes into an arena wall then the trial is stopped but the fitness accumulated up to that point is still averaged over the total number of time steps. This penalizes robots that do not avoid obstacles.

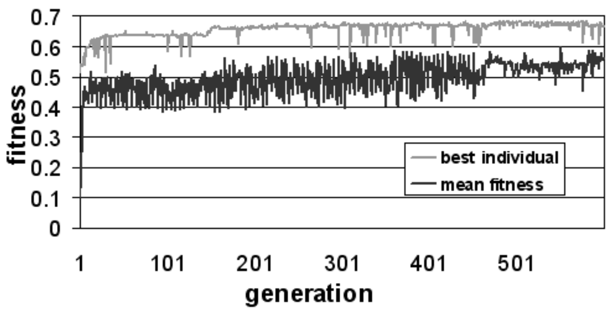

Figure 3. Typical increase in best and mean fitness over 600 generations: a rapid increase in fitness is followed by a slow increase that plateaus before another small increase occurs

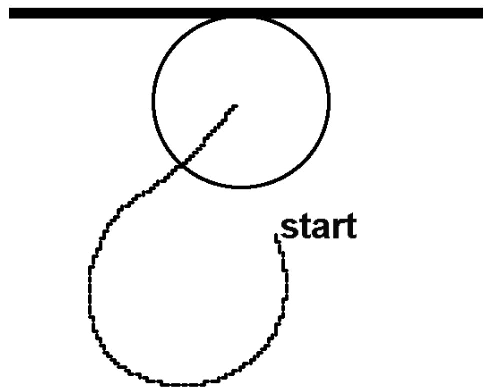

Figure 4. An example of low-fitness behaviour typical of a robot in the first generation. It makes a continuous line up to the point that it crashes into the arena wall. The robot robot does not change its pen state during the whole trial and only gains a small amount of fitness based on the maximum distance between points it has marked on the arena floor.

3 INITIAL EXPERIMENTS

Each genotoype was instantiated as robot controller in a walled arena (650mm $\times 350$mm) and each generation every individual was tested over 10 trials where they were placed in a a random position and orientation and tested for 200 (± 20) sensory-motor cycles. The robots were all tested on the same series of initial positions and orientations each generation, and these changed every generation. The pen was always placed in the down state at the beginning of a trial.

3.1 Preliminary results

Initially, the robots either were unable to move or crashed into the arena walls (Figure 4.). However, within 10 generations the majority of the population were able to avoid obstacles and starting to raise and lower their pen. As fitness increased, robots started to move backwards and forwards over marks, co-ordinating the raising and lower of their pens with the activity of the line sensor. Figure 5 shows the marks resulting from this behaviour. After 500 generations, the fittest individuals followed the walls leaving a long interrupted line. When they had completed a circuit of the arena and sensed the lines they had previously drawn they began to swerve left and right over the line, raising and lowering their pen and leaving marks parallel to the line marked on the initial circuit (Figure 6.).

4 CONCLUSION

Our project is in the early stages of conceptual and experimental development and it is a contentious issue whether our preliminary results demonstrate agent-relative novelty. The ANN controllers are feed forward and the connection strengths are fixed: the robots are purely reactive and their behaviour is a consequence of the current sensory-motor activity. Any novelty is therefore driven by changes in the environment. However, what makes the experimental results interesting conceptually is that the there is reciprocal feedback between the robot and the environment, which is changed by the mark making behaviour of the robot. Moreover, if we hold the starting conditions and testing arena constant, different agents (individuated by

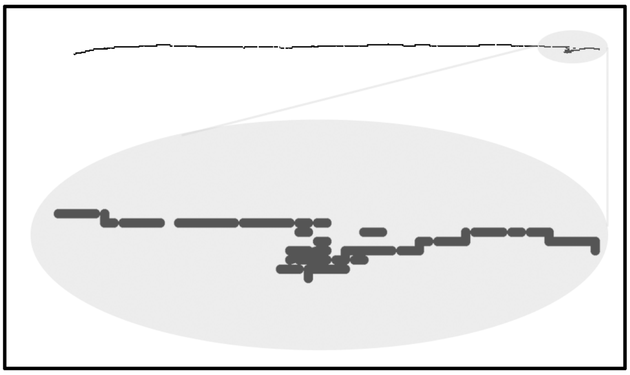

Figure 5. Performance of a mid fitness robot in an early generation. The image shows the line made in the arena - the robot started near the top left hand corner and moved across to the right top hand corner, raising and lowering the pen during this movement. The grey region is magnified to show the more complex marks the robot made at the end of the trial period: it made short line segments and moved forwards and backwards over them. The robot thereby gains fitness for correlated activity between changes in its pen state ('up to down' or 'down to up') and changes in its line sensor state ('on to off' or 'off to on').

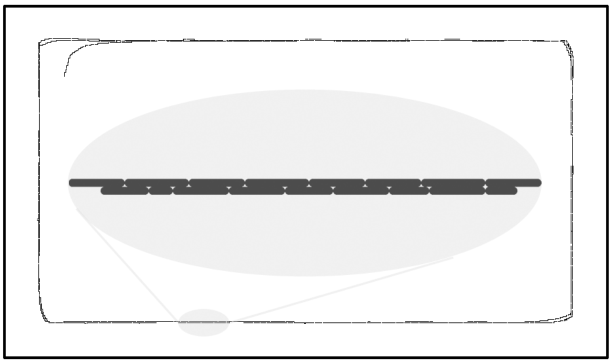

Figure 6. Performance of a high fitness individual after 500 generations. The robot completes one circuit of the arena with its pen down by following the walls. On its second circuit it sweeps left and right over the line, marking line segments parallel to the line it had previously left.

their ANN controllers) respond differently. And some, we are inclined to say, act in unexpected ways which are novel relative to other agents in the population (demarcated by the set of starting and testing conditions in question). Sometimes these behaviours enhance fitness, sometimes not. Our results are promising at such an early stage, and with incremental increases in complexity of the ANNs (for example, using recurrent and plastic connections), the behavior is likely to become less and less predictable and, we hope, more and more creative.

The research project can be seen as an 'embodied thought experiment' and the resulting models may not uncontroversially demonstrate the necessary conditions for creativity. The project will still be successful, however, if it helps to clarify how those conditions can be incorporated into future models. The general methods of evolutionary robotics and our particular hybrid application of simulation and real world situated agency provide a promising way to incorporate such conditions and thus model minimally creative processes and behaviour. As some have it, this is 'blue sky science', full of open-ended possibilities and requiring patient theorization. We are confident that, given time, it will reveal a number of interesting features of creativity.

ACKNOWLEDGEMENTS

The *Computational Intelligence, Creativity and Cognition* project is funded by the AHRC and led by Paul Brown in collaboration with Phil Husbands, Margaret Boden and Charlie Gere.

REFERENCES

[1] M. A. Boden, *The Creative Mind*, Routledge, London, second edn., 2004.
[2] V. Braitenberg, *Vehicles: Experiments in Synthetic Psychology*, The MIT Press, Cambridge, MA, 1987.
[3] D. Cliff, I. Harvey, and P. Husbands, 'Explorations in evolutionary robotics', *Adaptive Behavior*, **2**, 73–110, (1993).
[4] I. Harvey, *The Artificial Evolution of Adaptive Behaviour*, Ph.D. dissertation, School of Cognitive and Computing Sciences, University of Sussex, 1995.
[5] J. Holland, *Adaptation in Natural and Artificial Systems*, MIT Press, Cambridge: MA, 1975.
[6] P. Husbands, I. Harvey, D. Cliff, and G. Miller, 'Artificial evolution: A new path for artificial intelligence', *Brain and cognition*, **34**, 130 – 159, (1997).
[7] P. Kivy, *The Possessor and the Possessed*, Yale University Press, New Haven, 2001.
[8] A. Koestler, *The Act of Creation*, Macmillan, New York, 1964.
[9] Longinus, *On Great Writing: On the Sublime*, Hackett Publishing, Cambridge, 1991. trans. G.M.Grube.
[10] C. Martindale, 'Creativity, consciousness and cortical arousal', *Journal of Altered States of Consciousness*, **3**, 69–87, (1977).
[11] C. Martindale, *Cognition and Consciousness*, Dorsey Press, Homewood: IL, 1981.
[12] C. Martindale, 'Personality, situation and creativity', in *Handbook of Creativity*, eds., J.Glover, R.Ronning, and C.Reynolds, chapter Personality, situation and creativity, Plenum, New York, (1989).
[13] C. Martindale, 'Creativity and connectionism', in *Creative Cognition Approach*, eds., S. Smith, T.B. Ward, and R. Finke, MIT Press, Cambridge: MA, (1995).
[14] S. Nolfi, 'Using emergent modularity to develop control system for mobile robots', *Adaptive Behavior*, **5**, 343–363, (1997).
[15] S. Nolfi, *Evorobot 1.1*, Institute of Cognitive Sciences and Technologies, 2000. http://gral.ip.rm.cnr.it/evorobot/simulator.html.
[16] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*, The MIT Press, Cambridge, MA, 2000.
[17] R. Pfeifer and C. Scheier, *Understanding Intelligence*, The MIT Press, Cambridge, MA, 1999.
[18] Plato, *Plato: Collected Dialogues*, Princeton University Press, New Jersey, 1989.
[19] H. Poincaré, *Foundations of Science*, University Press of America, Washington, D.C., 1982.
[20] J.B. Pollack, H. Lipson, G. Hornby, and P. Funes, 'Three generations of automatically designed robots', *Artificial Life*, **7**, 215 – 223, (2001).
[21] A. Schopenhauer, *The World as Will and Representation*, Falcon's Wing Press, Indian Hills: CO, 1958. trans. E.F.J. Payne.
[22] T. Scutt, 'The five neuron trick: Using classical conditioning to learn how to seek light', in *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, eds., D. Cliff, P. Husbands, J.-A. Meyer, and S.W. Wilson, pp. 364 – 370, Cambridge: MA, (1994). MIT Press.