

Evolving the “Feeling” of Time Through Sensory-Motor Coordination: A Robot Based Model

Elio Tuci, Vito Trianni, and Marco Dorigo

IRIDIA

Univeristé Libre de Bruxelles - Bruxelles - Belgium
{etuci,vtrianni,mdorigo}@ulb.ac.be

Abstract. In this paper, we aim to design decision-making mechanisms for an autonomous robot equipped with simple sensors, which integrates over time its perceptual experience in order to initiate a simple signalling response. Contrary to other similar studies, in this work the decision-making is uniquely controlled by the time-dependent structures of the agent’s controller, which in turn are tightly linked to the mechanisms for sensory-motor coordination. The results of this work show that a single dynamic neural network, shaped by evolution, makes an autonomous agent capable of “feeling” time through the flow of sensations determined by its actions.

1 Introduction

Animals that forage in patchy environments, and do not have any *a priori* knowledge concerning the quality of the patch, must decide when it is time to leave a patch to move to another one of potentially better quality. Optimal foraging theory models assume that the experience that the animals have of the patch during time has an incremental or a decremental effect on the animal tendency to remain in the patch. These models show that some animals behave as if they made their decision on information gained while foraging [1]. Artificial autonomous agents might face similar problems: they may be required to change their behaviour because of information gained through a repeated interaction with their environment. In this paper, we aim to design decision-making mechanisms for an autonomous robot equipped with simple sensors, which integrates over time its perceptual experience in order to initiate alternative actions. In other words, the behaviour of the agent should change as a consequence of its repeated interaction with particular environmental circumstances.

We are interested in exploiting an evolutionary biologically-inspired approach, based on the use of dynamical neural networks and genetic algorithms [2]. Generally speaking, the appeal of an evolutionary approach to robotics is twofold. Firstly, and most basically, it offers the possibility of automating a complex design task [3]. Secondly, since artificial evolution needs neither to understand, nor to decompose a problem in order to find a solution, it offers the

possibility of exploring regions of the design space that conventional design approaches are often constrained to ignore [4]. In our work, artificial evolution should tightly couple the agent’s decision-making mechanisms to the nature of the environment and to the sensory-motor capabilities of the agent.

The experiment, described in details in section 2, requires an autonomous agent to possess both navigational skills and decision-making mechanisms. That is, the agent should prove capable of navigating within a boundless arena in order to approach a light bulb positioned at a certain distance from its starting position. Moreover, it should prove capable of discriminating between two types of environment: one in which the light can be actually reached; the other one in which the light is surrounded by a “barrier” which prevents the agent from proceeding further toward its target. Due to the nature of the experimental setup, the agent can decide in which type of environment it is situated only if it proves capable of (a) moving coordinately in order to bring forth the perceptual experience required to discriminate between the two environments; (b) integrating over time its perceptual experience in order to initiate a signalling behaviour if situated in an environment in which the light cannot be reached.

The contribution of this paper consists in showing that a single dynamic neural network, shaped by evolution, make an autonomous agent capable of “feeling” time through the flow of sensations determined by its actions. In other words, the controller allows an agent to make coordinated movements which bring forth the perceptual experience necessary to discriminate between two different types of environment and thus to initiate a simple signalling behaviour. Low level “leaky-integrator” neurons, which constitute the elementary units of the robot’s controller, provide the agent with the required time-dependent structures. This is not the first experiment in which time-dependent structures are evolved to control the behaviour of agents required to make decision based on their experience (see, for example, [5–7]). However, in [7] and in [5] the task was simpler than the one described in here, because the controller was only in charge of making the decision, while the nature of the perceptual experience of the robot was determined by the experimenter. The work illustrated in [6] and the one described in this paper differ in term of the nature of the cue/s exploited by the agent to make the discrimination: in [6] the cues the agent exploits are environmental structures (regularities) which the agent’s controller has to extract; in our task, the cue the agent has to exploit concerns the persistence of a perceptual state, which is common to both types of environment.

2 Methods

Description of the task. At the start of each trial, a robot is positioned within a boundless arena, at about 100 cm west of the light, with a randomly determined orientation chosen between north-east and south-east (see Figure 1 left). The light is always turned on during the trial. The robot perceives the light through its ambient light sensors. The colour of the arena floor is white except for a circular band, centred around the lamp, within which the floor is in shades of grey. The circular band covers an area between 40 cm and 60




Fig. 1. Depiction of the task. The small black circles represent the robot at starting position. The small empty circles represent the light bulb. The arena floor is white everywhere except within a circular band surrounding the light. The *way in* zone corresponds to the sector of the band, indicated by dotted lines, in which the floor is white. In both pictures, the continuous arrows are examples of good navigational strategies; the dashed arrows are examples of forbidden trajectories. In *Env. B*, the continuous arrow gets thicker to indicate that the robot emits a sound after having made a loop around the light.

cm from the light; the floor is black at exactly 40 cm from the light; the grey level decreases linearly with the distance from the light. The robot perceives the colour of the floor through its floor sensor, positioned on its belly. The robot can freely move within the band, but it is not allowed to cross the black edge. The latter can be imagined as an obstacle, a trough, that prevents the robot to further approach the light (see dashed arrows in Figure 1). The area in shades of grey is meant to work as a warning signal which “tells” the robot how close it is to the danger—i.e., the black edge.

There are two types of environment. In one type, referred to as *Env. A*, the band presents a discontinuity (see Figure 1, left). This discontinuity, referred to as the *way in* zone, is a sector of the band in which the floor is white. In the other type, referred to as *Env. B*, the band completely surrounds the light (see Figure 1, right). The *way in* zone represents the path along which the robot is allowed to safely reach the light. A successful robot should prove to be capable of performing phototaxis as well as looking for the *way in* zone to avoid to cross the black edge of the band. Such a robot should always reach the light in *Env. A*, whereas, in *Env. B*, besides avoiding to cross the black edge, the robot should signal the absence of the *way in* zone by emitting a tone. How can the robot distinguish between *Env. A* and *Env. B*? The cue the agent should use is a temporal one: that is, the *Env. B* can be “recognised” by the persistence of the perception of the band for the amount of time that, given the trajectory and speed of the robot, corresponds to the time required to make a loop around the light. If the perception of the band persists long enough, this means that there is no *way in* zone, and a tone has to be emitted.

The difficulty of this experiment resides in synthesising a controller which, by solely integrating over time the perception of the colour of the floor under the robot’s belly, brings forth something similar to the “feeling” of being travelling within the band for the time required to complete a loop, so to “recognise” that there is no *way in* zone. The amount of time required for the robot to perform

a complete loop of the band depends on the dimensions of the band and on the way in which the robot moves within the band. The characteristics of the band are set by the experimenter and they do not change during the evolution. The way in which the robot moves within the band—e.g., its average speed and its trajectory—is determined by the robot’s controller, by directly setting the speed of the robot’s wheels. Thus, a successful controller should make the robot move in such a way that, if the perception of the band persists over a certain amount of time, the following conclusions can be drawn: (i) the band does not present any discontinuity; (ii) the sound signalling must be activated. Continuous time recurrent neural networks (CTRNNs), shaped by evolution, seem to be a suitable tool to obtain this kind of mix between mechanisms for sensory-motor coordination and time-dependent structures required to perform this task [5].

The simulation model. The robot and its world are simulated using a modified version of the “minimal simulation” technique described by Jakobi in [8]. Jakobi’s technique uses high levels of noise to guarantee that the simulated controller will transfer to a physically realised robot with no loss of performance. Our simulation models a Khepera robot, a 55 mm diameter cylindrical robot. This simulated robot is provided with two ambient light sensors, placed at 45 degrees (A_1) and -45 degrees (A_2) with respect to its heading, and a floor sensor positioned facing downward on the underside of the robot (F). The light sensors have an angle of acceptance of 120 degrees. Light levels change as a function of the robot’s distance from the lamp. The light sensor values are extrapolated from a look-up table which corresponds to the one provided with the Evorobot simulator (see [9]). The floor sensor can be conceived of as a proximity infra-red sensor capable of detecting the level of grey of the floor. It produces an output which is proportional to the level of grey, scaled between 0—when the robot is positioned over white floor—and 1—when it is over black floor.

The controller and the evolutionary algorithm. Fully connected, eight neuron CTRNNs are used. All neurons are governed by the following state equation:

$$\frac{dy_i}{dt} = \frac{1}{\tau_i} \left(-y_i + \sum_{j=1}^8 \omega_{ji} \sigma(y_j + \beta_j) + gI_i \right) \quad \sigma(y_j + \beta_j) = \frac{1}{1 + e^{-x}} \quad (1)$$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate, ω_{ji} the strength of the synaptic connection from neuron j^{th} to neuron i^{th} , I_i the intensity of the sensory perturbation on sensory neuron i . Three neurons receive input (I_i) from the robot sensors: e.g., neuron N_1 takes input from A_1 , N_2 from A_2 , and N_3 from F . These input neurons receive a real value in the range [0,1], which is a simple linear scaling of the reading taken from its associated sensor. The other neurons do not receive any input from the robot’s sensors. The cell potential (y_i) of the 6th neuron, mapped into [0,1] by a sigmoid function (σ) and then set to 1 if bigger than 0.5 or 0 otherwise, can be used by the robot to control the sound signalling system. The cell potential (y_i) of

the 7th and the 8th neuron, mapped into $[0,1]$ by a sigmoid function (σ) and then linearly scaled into $[-10,10]$, set the robot motors output. The strength of synaptic connections ω_{ji} , the decay constant τ_i , the bias term β_j , and the gain factor g are genetically encoded parameters. Cell potentials are set to 0 any time the network is initialised or reset, and circuits are integrated using the forward Euler method with an integration step-size of 0.2.

A simple generational genetic algorithm (GA) is employed to set the parameters of the networks. The population contains 100 genotypes. Generations following the first one are produced by a combination of selection with elitism, recombination and mutation. For each new generation, the three highest scoring individuals (“the elite”) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness-proportional selection from the 70 best individuals of the old population. Each genotype is a vector comprising 81 real values (64 connections, 8 decay constants, 8 bias terms, and a gain factor). Initially, a random population of vectors is generated by initialising each component of each genotype to values chosen uniformly random from the range $[0,1]$. New genotypes, except “the elite”, are produced by applying recombination with a probability of 0.3 and mutation. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.15. The mean of the Gaussian is 0, and its standard deviation is 0.1. During evolution, all vector component values are constrained to remain within the range $[0,1]$. Genotype parameters are linearly mapped to produce CTRNN parameters with the following ranges: biases $\beta_j \in [-2,2]$, weights $\omega_{ji} \in [-6,6]$ and gain factor $g \in [1,12]$. Decay constants are firstly linearly mapped onto the range $[-0.7, 1.7]$ and then exponentially mapped into $\tau_i \in [10^{-0.7}, 10^{1.7}]$. The lower bound of τ_i corresponds to a value slightly smaller than the integration step-size used to update the controller; the upper bound corresponds to a value slightly bigger than the average time required by a robot to reach and to perform a complete loop of the band in shades of grey.

The evaluation function. During the evolution, each genotype is coded into a robot controller, and is evaluated 40 times—20 times in *Env. A* and 20 in *Env. B*. At the beginning of each trial, the neural network is reset—i.e., the activation value of each neuron is set to zero. Each trial differs from the others in the initialisation of the random number generator, which influences the robot starting position and orientation, the position and amplitude of the *way in* zone, and the noise added to motors and sensors. For each of the 20 trials in *Env. A*, the position of the *way in* zone is varied to facilitate the evolution of robust navigational strategies. Its amplitude varies within the interval $[\frac{\pi}{6}, \frac{\pi}{2}]$. Within a trial, the robot life-span is 80 s (400 simulation cycles). A trial is terminated earlier if either the robot crosses the black edge of the band (see dashed arrows in Figure 1) or because it reaches an Euclidean distance from the light higher than 120 cm. In each trial t , the robot is rewarded by an evaluation function f_t which corresponds to the sum of the following four components:

$$R_{\text{motion}} = \frac{d_f - d_n}{d_f} \quad R_{\text{error}} = -\frac{p_b}{t_b}$$

$$R_{\text{near}} = \begin{cases} p_c/t_c & \text{Env. A} \\ 0 & \text{Env. B} \end{cases} \quad R_{\text{signal}} = \begin{cases} 0 & \text{Env. A} \\ p_a/t_a & \text{Env. B} \end{cases}$$

R_{motion} rewards movements toward the light bulb: d_f and d_n represent respectively the furthest and the nearest Euclidean distance between the robot and the light bulb. In particular, d_f is updated whenever the robot increases its maximum distance from the light bulb. At the beginning of the trial, d_n is fixed as equal to d_f , and it is subsequently updated every time step when (i) the robot gets closer to the light bulb; (ii) d_f is updated. In this latter case, d_n is set equal to the new d_f . In *Env. A*, d_n is set to 0 if the robot is less than 7.5 cm away from the light bulb. In *Env. B*, d_n is set to 0 if the robot makes a complete loop around the light bulb while remaining within the circular band. R_{error} is negative to penalise the robot for (i) signalling in *Env. A*, and (ii) signalling in *Env. B* before having made a loop around the light: p_b is the number of simulation cycles during which the robot has erroneously emitted a tone, and t_b is the number of simulation cycles during which the robot was not required to signal. R_{near} rewards movements for remaining close to the light bulb: p_c is the number of simulation cycles during which the robot was no further than 7.5 cm away from the light bulb in *Env. A*, and t_c is the robot life-span. In *Env. B* the robot cannot get closer than 40 cm to the light, therefore, this component is equal to 0. R_{signal} rewards signalling in *Env. B*: p_a is the number of simulation cycles during which the robot has emitted a tone after having made a loop around the light, and t_a is the number of simulation cycles during which the robot was required to emit a tone. In *Env. A*, this component is always set to zero. Recall that the robot is also penalised for crossing the black edge of the band and for reaching a distance from the light higher than 120 cm. In these cases, the trial is ended and the robot's fitness is computed by considering the current state of the system.

3 Results

Twenty evolutionary simulations, each using a different random initialisation, were run for 6000 generations. We examined the best individual of the final generation from each of these runs in order to establish whether they evolved the required behaviour. During re-evaluation, each of the twenty best evolved controllers was subjected to a set of 100 trials in *Env. A* and a set of 100 trials in *Env. B*. At the beginning of each re-evaluation trial, the controllers are reset. Each trial has a different initialisation (see section 2 for details). During re-evaluation, the robot life-span is 120 s (600 simulation cycles).

Firstly, we analyse the navigational abilities of the best evolved robot in an *Env. A*. Recall that, in these circumstances, a successful robot should reach the light bulb going through the *way in zone*, without signalling. Figure 2 shows the percentage of time each of these robots spent in an area located in the proximity of the light bulb during the 100 re-evaluation trials in *Env. A*, as computed by the fitness component R_{near} . These results prove that almost all the best

evolved robots employ successful navigational strategies which allow them to find the *way in zone*, and to spend between 40% and 80% of their life-time close to the target. In both types of environment, the run n. 1 and n. 8 are slightly less successful than the others. We observed that their failure was caused by a tendency to cross the black edge of the band. A qualitative analysis of the robots’ behaviour shows that, when the best evolved robots are situated in an *Env. B*, their navigational strategies allow them (i) to approach the light as much as possible without crossing the black edge of the band, and (ii) to make a loop around the light, between 40 cm and 60 cm from the light, following a trajectory nearly circular. The agents are not evolved just to navigate properly toward the light, but also for accurately discriminating between the two types of environment. Recall that the agents are required to make their choice by emitting a tone only if they “feel” they have been situated in an *Env. B*. We observed that none of the best evolved robots emits a tone if situated in *Env. A*. Table 1 shows data concerning the signalling behaviour of the best evolved robots in *Env. B*. In particular, the column (*Succ.*) shows the percentage of successful trials at the discrimination task, for each of these robots, during the 100 re-evaluation trials. We can notice that eleven out of twenty robots never missed to emit a tone if situated in *Env. B* (i.e., run n. 2, 3, 6, 7, 10, 13, 14, 16, 17, 18, 20). The robot run n. 4 shows a 96% success rate. The other eight robots did not emit any tone during each of the 100 re-evaluation trials in *Env. B* (i.e., run n. 1, 5, 8, 9, 11, 12, 15, 19).

The quality of the signalling behaviour can be established with reference to the amount of error of type I (*Err. I*) and error of type II (*Err. II*) made by the successful robots. The *Err. I* refers to those cases in which the robot emits a tone **before** having made a loop around the light. The *Err. II* refers to those cases in which the robot emits a tone **after** having completed the loop. *Err. I* can be considered as a false positive error—i.e., signalling that there is no *way in zone* when there may be one. *Err. II* can be considered as a false negative error—i.e., not accurately signalling that there is no *way in zone*. Both types of error are calculated with respect to the angular displacement of the robot




Fig. 2. Box-and-whisker plot visualising the fitness component R_{near} of the best evolved strategies for each run, computed over 100 trials in *Env. A*. The box comprises observations ranging from the first to the third quartile. The median is indicated by a horizontal bar. The whiskers extend to the most extreme data point which is no more than 1.5 times the interquartile range. Outliers are indicated as circles.




Fig. 3. The graph shows the average and standard deviation of Err I—i.e., negative angles—and Err II—positive angles—for each environmental condition defined by the distance between the black edge of the band and the light. The values are averaged over 100 re-evaluation trials.

around the light from the starting position—the position at the time when the robot enters into the circular band—to the signalling position—the position at the time when the robot starts signalling. If the robot makes no errors, this angle is 2π . The bigger is the angle, the less reliable is the signalling mechanism. However, we should notice that, due to the nature of the task, it is very difficult to make no errors—i.e., emitting a tone precisely at the time in which an entire loop around the light is made. For our purpose, we consider successful an agent that, in order to signal the absence of the *way in zone*, manages to reduce the amount of errors of both types. As shown in column “*Err. I—Avg.*” and in column “*Err. II—Avg.*” of table 1, the robots evolved in runs n. 17, 18 and 20 manage to keep their average error of both types smaller than an angle of 10 degrees. The robots evolved in runs n. 3 and n. 7 are also quite accurate, with both types of error smaller than 16 degrees. All the other “signalling” robots are less accurate, with average errors bigger than 20 degrees.

Table 1. Quality of the performance of the twenty best evolved robots in *Env. B* during the 100 re-evaluation trials. The table shows the average angle (degrees), the standard deviation, and the number of times the error was made for both *Err. I* and *Err. II*. The success rate (%) at the discrimination task is indicated by *Succ.*

run	Err. I			Err. II			Succ.	run	Err. I			Err. II			Succ.
	Avg.	Std	n.	Avg.	Std	n.	(%)		Avg.	Std	n.	Avg.	Std	n.	(%)
n. 1	-	-	-	-	-	-	0	n. 11	-	-	-	-	-	-	0
n. 2	18.55	5.05	3	99.96	24.45	97	100	n. 12	-	-	-	-	-	-	0
n. 3	8.50	5.67	36	11.70	8.68	64	100	n. 13	7.88	8.53	14	34.84	23.29	86	100
n. 4	0	0	0	64.65	13.88	96	96	n. 14	4.05	4.40	11	25.30	21.62	89	100
n. 5	-	-	-	-	-	-	0	n. 15	-	-	-	-	-	-	0
n. 6	5.14	3.26	13	30.66	20.92	87	100	n. 16	6.41	4.94	37	22.69	38.80	63	100
n. 7	2.86	4.13	9	15.14	9.02	91	100	n. 17	6.67	3.36	96	1.88	1.44	4	100
n. 8	-	-	-	-	-	-	0	n. 18	7.65	5.01	70	5.89	4.42	30	100
n. 9	-	-	-	-	-	-	0	n. 19	-	-	-	-	-	-	0
n. 10	0	0	0	44.63	15.48	100	100	n. 20	6.11	3.92	59	9.02	6.88	41	100

The mechanisms that the successful robots employ to solve the discrimination task are tuned to those environmental conditions experienced during evolution. As expected, they do not properly work if the environment changes. For example, we observed that both the reduction and the increment of the distance between the black edge of the band and the light disrupt the robots’ performance: the smaller is the distance, the bigger is the *Err. II*—i.e., signalling after having made a loop around the light; the higher is the distance, the bigger is *Err. I*—i.e., signalling before having made a loop around the light. However, as far as it concerns the robot evolved in run n. 17, we observed that in particular circumstances it is capable of adjusting its behaviour to the altered environmental conditions. As shown in Figure 3, when the black edge of the band is 34 cm away from the light, the performance of the robot is comparable to the one recorded in those conditions experienced during evolution. Moreover, when this distance is 54 cm, the *Err I* gets smaller. How do we explain the robustness of this behaviour? An explanation could be that the robot is taking into account (i.e., it integrates over time) both the perception of the floor and the intensity of the light. Obviously, in the altered environmental conditions, the perception of the floor is not disrupted. That is, the robot can still freely move within the band in order to bring forth the perception of the colour of the floor that it was used to experience during evolution. However, for a given perception of the floor, the corresponding light intensity is altered by the fact that the black edge of the band is not at the same distance from the light, as during evolution. It is reasonable to conclude that, for an agent that integrates both the perception of the floor and the intensity of the light, the relationship between these two sensory inputs might have a bearing on the emission of the tone. For example, for a given level of grey, the higher/lower is the intensity of the light the shorter/longer is the time it takes to the robot to emit a tone. This mechanism improves the robustness and the adaptiveness of the agent’s controller in particular environments that differ from the one experienced during evolution. Interestingly, the evolution of this mechanism has not been favoured by any selective pressures explicitly introduced to encourage the emergence of robust controllers. Serendipitously the artificial neural networks turned out to be capable of tracking significant variations in environmental conditions—i.e., the relationship between the intensity of the light and levels of grey of the floor.

4 Conclusions

In this paper, we have shown that a single dynamic neural network can be synthesised by evolution to allow an autonomous agent to make coordinated movements that bring forth the perceptual experience necessary to discriminate between two types of environments. The results illustrated here are of particular interest because, contrary to other previous similar studies, in this work the decision-making is uniquely controlled by the time-dependent structures of the agent controller, which in turn, are tightly linked to the mechanisms for sensory-motor coordination. Unexpectedly, one of the evolved controllers proved robust to variation in the environmental conditions without being explicitly evolved for

this. Based on this preliminary but encouraging results, in future works, we will consider more challenging experimental setups. In particular, the integration of the agent perception over time will not be solely finalised to a simple signalling response, but it will trigger effective alternative actions as it is the case for animal species making decisions about the quality of foraging sites.

Acknowledgement

This work was supported by the “SWARM-BOTS” project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-2000-31010. The authors thank Christos Ampatzis, Paul Graham, Roderich Gross, Thomas Halva Labella, Shervin Nouyan, Peter Todd, and Andreas Wilke for useful comments. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate, through the grant “Virtual Swarm-bots”, contract no. 9.4515.03, and from the “ANTS” project, an “Action de Recherche Concertée” funded by the Scientific Research Directorate of the French Community of Belgium. The information provided is the sole responsibility of the authors and does not reflect the Community’s opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

References

1. Nonacs, P.: State dependent behavior and the marginal value theorem. *Behavioral Ecology* **12** (2003) 71–83
2. Beer, R.D.: A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* **72** (1995) 173–215
3. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA (2000)
4. Harvey, I., Husbands, P., Cliff, D.: Issues in evolutionary robotics. In Meyer, J.A., Roitblat, H., Wilson, S., eds.: *Proc. of the 2nd Int. Conf. on Simulation of Adaptive Behavior*, MIT Press, Cambridge, MA (1992) 364–373
5. Beer, R.D., Gallagher, J.C.: Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior* **1** (1992) 91–122
6. Nolfi, S.: Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science* **14** (2002) 231–244
7. Tuci, E., Harvey, I., Todd, P.M.: Using a net to catch a mate: Evolving CTRNNs for the dowry problem. In Hallam, B., Floreano, D., Hallam, J., Hayes, G., Meyer, J.A., eds.: *Proceedings of SAB’02*, MIT press, Cambridge, MA (2002)
8. Jakobi, N.: Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior* **6** (1997) 325–368
9. Nolfi, S.: *EvoRob 1.1 User Manual*. Institute of Psychology, National Research Council (CNR). (2000) Available at <http://gral.ip.rm.cnr.it/evorobot/simulator.html>.