Homeostasis and evolution together dealing with novelties and managing disruptions

Patricia A. Vargas
School of Maths and Computer Science, Heriot-Watt University, Edinburgh, UK
Renan C. Moioli
CCNR, Department of Informatics, University of Sussex, Brighton, UK
Fernando J. von Zuben
Laboratory of Bioinformatics and Bioinspired Computing, FEEC, University of Campinas, Campinas, Brazil, and
Phil Husbands
CCNR, Department of Informatics, University of Sussex, Brighton, UK

Received 30 November 2008
Revised 4 March 2009
Accepted 12 March 2009

Abstract

Purpose – The purpose of this paper is to present an artificial homeostatic system whose parameters are defined by means of an evolutionary process. The objective is to design a more biologically plausible system inspired by homeostatic regulations observed in nature, which is capable of exploring key issues in the context of robot behaviour adaptation and coordination.

Design/methodology/approach – The proposed system consists of an artificial endocrine system that coordinates two spatially unconstrained GasNet artificial neural network models, called non-spatial GasNets. Both systems are dedicated to the definition of control actions in autonomous navigation tasks via the use of an artificial hormone and a hormone receptor. A series of experiments are performed in a real and simulated scenario in order to investigate the performance of the system and its robustness to novel environmental conditions and internal sensory disruptions.

Findings – The designed system shows to be robust enough to self-adapt to a wider variety of disruptions and novel environments by making full use of its in-built homeostatic mechanisms. The system is also successfully tested on a real robot, indicating the viability of the proposed method for coping with the reality gap, a well-known issue for the evolutionary robotics community.

Originality/value – The proposed framework is inspired by the homeostatic regulations and gaseous neuro-modulation that are intrinsic to the human body. The incorporation of an artificial hormone receptor stands for the novelty of this paper. This hormone receptor proves to be vital to control the network's response to the signalling promoted by the presence of the artificial hormone. It is envisaged that the proposed framework is a step forward in the design of a generic model for coordinating many and more complex behaviours in simulated and real robots, employing multiple hormones and potentially coping with further severe disruptions.

Keywords Robotics, Artificial intelligence, Neural nets, Endocrine system

Paper type Research paper

Renan C. Moioli and Fernando J. von Zuben would like to thank CAPES and CNPq for their financial support. Patricia A. Vargas and Phil Husbands were supported by the Spatially Embedded Complex Systems Engineering Project, EPSRC Grant No. EP/C51632X/1.

International Journal of Intelligent Computing and Cybernetics
Vol. 2 No. 3, 2009
pp. 435-454

1756-378X
DOI 10.1108/17563780910982680

1. Introduction

The term homeostasis has its origins in the work of the French Physiologist Claude Bernard (1813-1878), who founded the principle of the internal environment, further expanded by Cannon (1929) as the process of homeostasis. Nonetheless, for Pfeifer and Scheier (1999), homeostasis was substantially defined by the English Cyberneticist and Psychiatrist Ashby (1952). For Ashby, the ability to adapt to a continuously changing and unpredictable environment (adaptivity) has a direct relation to intelligence. During the adaptive process, some variables need to be kept within predetermined limits, either by evolutionary changes, physiological reactions, sensory adjustment, or simply by learning novel behaviours. Therefore, with this regulatory task attributed to the homeostatic system, the organism or the artificial agent can operate and stay alive in a viability zone.

Basically, homeostasis can be considered paramount for the successful adaptation of the individual to dynamic environments, and thus is essential for survival. Moreover, Dyke and Harvey (2005, 2006) have pointed out that in order to understand real or artificial life, it is necessary to first understand the conceptual framework and basic mechanisms of homeostasis. In the human body, specific sensory receptors trigger specific responses in the nervous, immune and endocrine systems, which are the main systems directly related to the process of homeostasis (Besendovsky and Del Rey, 1996). Therefore, one can say that it is a consensus that homeostatic processes are strictly connected to the balance of any real or artificial life.

The theory presented by Ashby has motivated applications of homeostasis in the synthesis of autonomous systems in mobile robotics (Di Paolo, 2000; Neal and Timmis, 2003; Avila-García and Cañamero, 2004; Harvey, 2004; Hoinville and Henaff, 2004; Vargas *et al.*, 2005; Moioli *et al.*, 2008a, b). A possible approach for the development of such systems are found in the ideas presented by Alife researchers like Di Paolo (2000) and Hoinville and Hnaff (2004), who encompass homeostasis within one unique structure, i.e. an artificial neural network (ANN) capable of dynamically changing their connections by means of plasticity rules (hence modifying the response of the network).

Our present work approaches the issue in a different direction and is an extension of previous work by Vargas *et al.* (2005) and Moioli *et al.* (2008a, b). The framework further developed here, named evolutionary artificial homeostatic system (EAHS), is comprised by an evolved artificial endocrine system (AES) and two evolved spatially unconstrained recurrent ANN models, called non-spatial GasNet (NSGasNet) (Vargas *et al.*, 2007). The originality of this work comes from the fact that the EAHS now incorporates hormone receptors in an attempt to mediate the influence of the hormone released by the AES.

We believe that this is a step towards the development of a more biologically plausible system able to robustly coordinate multiple behaviours in a mobile robot exposed to a dynamic environment and susceptible to sensory disruptions. Experiments with a simulated and real robot are devised, proving the superior performance of the proposed system.

This paper is organized as follows: Section 2 presents the basis of the approach adopted in this work, together with the details of our proposal. Section 3 describes the experiments undertaken and their implementation procedures. Section 4 covers the simulation results of the artificial homeostatic system subject to a novel environment and some degrees of disruption, as well as its successful transfer to a real robot. Section 5 closes the paper with final remarks and some directions for future investigation.

2. Evolutionary artificial homeostatic systems

The framework proposed is particularly concerned and inspired by interactions between the nervous and endocrine systems in the human body. These neuro-endocrine interactions are employed to maintain homeostasis, metabolism and reproduction in the organism.

The endocrine system employs chemical signalling substances called hormones which also act as ligands (from the Latin word *ligare* – to bind) while binding to receptors. Receptors are protein molecules embedded in either the plasma membrane or cytoplasm of a cell. This binding is typically responsible for a cellular response. In this way, each cell has its own receptors which will mediate its response to specific hormones. The presence of receptors is vital to control these cellular responses to the signaling promoted by assorted types of hormones for hormones are released in the blood stream and therefore could reach literally any cell in the organism (Guyton and Hall, 1996). Moreover, it is known that the neuro-endocrine interaction is vital for the control and coordination of behaviour (Grillner, 2006).

Basically, on a typical neuro-endocrine interaction, hormones that are released by endocrine glands (EGs) can affect the nervous system, which in turn can transmit nerve impulses affecting the production and secretion of hormones, thus establishing a control loop mechanism. These are called positive and negative feedback mechanisms (Guyton and Hall, 1996). In our EAHS model, these mechanisms are represented by coupled difference equations as it will be described in the next section.

Two versions of the proposed EAHS are investigated in this work (Figure 1). One is composed of an evolved AES and two evolved NSGasNet models, both exhibiting discrete-time dynamics (Moioli *et al.*, 2008a). The novel model, named EAHS-R, derives from the former but incorporates a hormone receptor.

The incorporation of a hormone receptor was intended to enhance the system robustness in environments that differ from the ones employed during the evolution of the system. Next, we will describe both variants: first, the basic EAHS framework; and second, the EAHS-R, which includes the hormone receptor.

2.1 The EAHS framework

The EAHS framework encompasses one AES and two NSGasNets. The AES consists of three main modules: hormone level (HL), hormone production controller (HPC) and EG. The HL has a record of the level of hormone in the organism; the HPC is responsible for controlling the production of hormones in response to variations in the internal state of the organism and to external stimulation; and the EG receives inputs from the HPC, being responsible for producing and secreting hormones when required. The hormone production HP is updated as follows:

$$\mathrm{HP}(t+1) = \begin{cases} 0, & \text{if} \quad \mathrm{IS} < \theta \\ (100 - \%\mathrm{ES}) \times \alpha(\mathrm{Max}(\mathrm{HL}) - \mathrm{HL}(t)), & \text{otherwise} \end{cases} \tag{1}$$

where θ is the target threshold of the internal state IS; ES is the external stimulus; α is the scaling factor; HL is the hormone level; and t is the discrete time index. If the internal state IS is greater than or equal to a target threshold, then hormone will be produced at a rate that will depend upon the level of the external stimulus received and the level of hormone already present within the artificial organism. Otherwise,

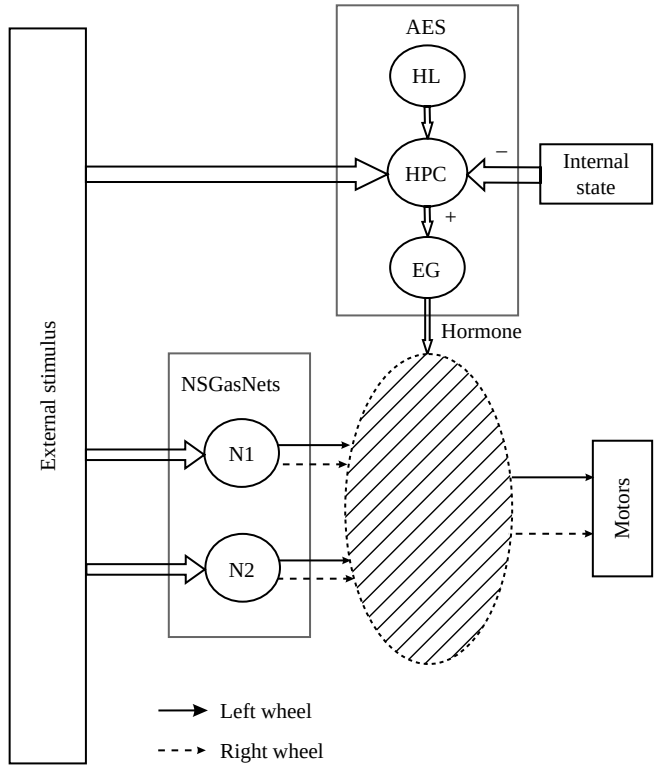

Figure 1.
The basic framework of the EAHS

Note: The shaded area represents the hormone influence which depends on the modulation employed, i.e. with or without hormone receptors

hormone production will cease. The internal state IS is governed by:

$$\mathrm{IS}(t+1) = \begin{cases} 0, & \text{if} \quad (\mathrm{ES} \geq \lambda) \text{ and } (\mathrm{HL} \geq \omega) \\ \mathrm{IS}(t) + \beta(\mathrm{Max}(\mathrm{IS}) - \mathrm{IS}(t)), & \text{otherwise} \end{cases} \tag{2}$$

where λ and ω are thresholds associated with ES and HL, respectively; and β is a gain value for the rate of change of the internal state. The hormone level HL represents the amount of hormone stimulating the ANN, and is a function of its current value and of the amount of hormone produced:

$$\mathrm{HL}(t+1) = \mathrm{HL}(t) \times \mathrm{e}^{-1/T} + \mathrm{HP}(t) \tag{3}$$

where T is the half-life constant.

The two NSGasNets (N1 and N2 in Figure 1) are previously and separately evolved to accomplish two distinct and possibly conflicting behaviours. The NSGasNet is a spatially unconstrained model, which was shown to present superior performance when compared to the original GasNet model on a pattern generation task and on a

delayed response robot task (Vargas *et al.*, 2007, 2008). Details about the NSGasNet model are presented in Section 3.1.

For the EAHS, the outputs of the NSGasNets are modulated by the hormone level HL (equation (3)), giving rise to the dynamical coordination of both behaviours. Specifically, for N1 the modulation factor is given by (1-HL) and for N2 by (HL) (Figure 2(a)). The modulation factor that differentiates the EAHS from the EAHS-R is explained in the following section.

2.2 The EAHS-R framework

The EAHS-R is an extended version of the EAHS and represents the originality of this work. It incorporates the idea of hormone receptors to the former framework and involves a more intricate modulation of the network outputs.

In this version, the modulation factor, which is also dependent on the HLs, affects the NSGasNet outputs (Section 2) by multiplying it by a δ value (Figure 2(b)). This value represents the hormone receptor sensitivity dictated by the following equation:

$$\delta = 1 - \frac{[\text{Min}(\text{lightSensorReadings}) - \text{Min}\,R]^3}{(\text{Max}\,R - \text{Min}\,R)^3} \tag{4}$$

where min(lightSensorReadings) is the most stimulated light sensor value. The parameters Min R and Max R assume the values 65 and 450, respectively. In this first attempt to include a receptor on the EAHS, the parameters Min R and Max R of equation (4) were devised empirically and were specific for the chosen tasks. Figure 3 shows the value of variable δ according to the external stimulus (ES) (i.e. light sensor readings).

It is important to highlight that the variable δ, representing the hormone receptor sensitivity, is driven by the light source radiance, thus regulating the agent's response to the presence of the hormone. One can understand this sensitivityto the light source as if a fictitious hypothalamus releases a neurotransmitter while sensing the presence of light. This neurotransmitter would then bind to the receptor changing the agent's sensitivity to the HL mediated by the proximity of the light source.

Given the original architecture of the AES, one might enquire: "Why not evolving the rules that dictate the AES coupling behaviour?" Similar to the laws of physics,

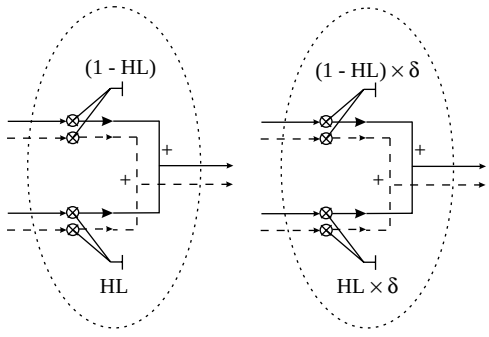

Notes: (a) EAHS without the hormone receptor; and
(b) EAHS-R with hormone receptor sensitivity value δ

Figure 2.
Detailed hormone influence for each model

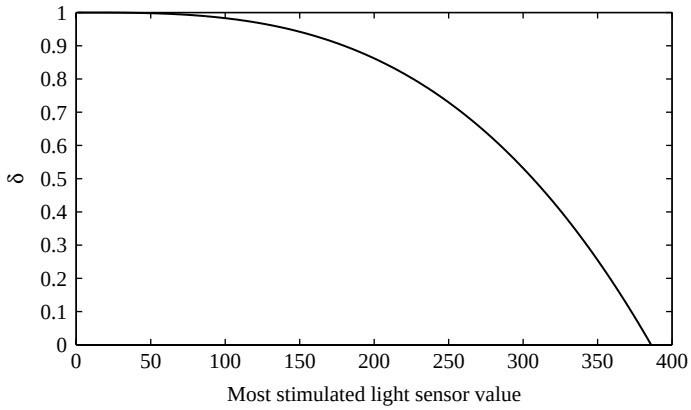

Figure 3.
δ variation over the most stimulated light sensor readings

which dictate the nature of the physical coupling between a cell and its exterior environment (Di Paolo, 2005), the proposed rules will govern the artificial coupling between the AES, the NSGasNets, the environment and the artificial agent (equations (1)-(3)). In this sense, our new AES will evolve only the parameters of these rules.

3. Methods

A simulated and a real mobile robot equipped with an internal battery meter has to perform two coupled but distinct tasks: to explore the environment while avoiding collisions and to search for a light source when its battery level is low (the light source indicates the location of the battery charger). This experiment was first proposed by Vargas *et al.* (2005) to assess the performance of the artificial homeostatic system that has served as an inspiration to the current proposal.

The robot is a Khepera-II model (Figure 4(a)) (KTEAM SA, 2009) with two wheels which are responsible for its motion (each wheel has an independent electric motor). The robot has eight infrared sensors. The sensors measure the environmental luminosity (ranging from 50 to 500 – 50 being the highest luminosity that can be sensed) and the obstacle distance (ranging from 0 to 1,023 – the latter value represents the closest distance to an object). The range of the obstacle sensors is approximately 10 cm. The simulations were carried out using a robot simulator named KiKS (Storm, 2004). It reproduces both sensory and motor behaviours of the real Khepera-II robot.

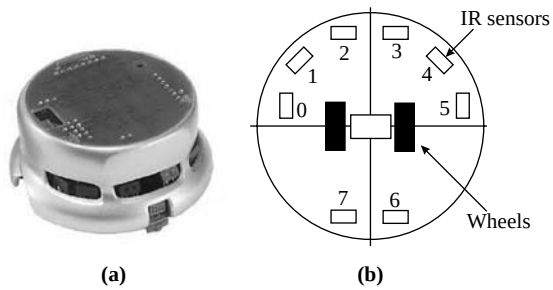

Figure 4.
Real Khepera-II robot (a) and its schematic representation, including the IR sensors (b)

The evolution of the system was divided in two steps. First, the two NSGasNets were evolved independently employing a distributed genetic algorithm (Hillis, 1990; Collins and Jefferson, 1991) (one NSGasNet evolved for each task). The AES was evolved thereafter as a coordination module responsible for the swapping of behaviours between the NSGasNets. No crossover is employed. A generation is defined as 25 breeding events, and the evolutionary algorithm runs for a maximum of 50 generations. The fitness criteria are specific for each task. There are two mutation operators applied to 10 per cent of the genes. The first operator is only for continuous variables. It produces a change at each *locus* by an amount within the $[-10, +10]$ range. For the second mutation operator, designed to deal with discrete variables, a randomly chosen gene *locus* is replaced with a new value that can be any value within the [0, 99] range, in a uniform distribution. For further details about the application of the genetic algorithm to the evolution of GasNet models, the reader should refer to Husbands *et al.* (1998).

The first experiment assesses the performance of the EAHS without any internal disruption (i.e. there is no change in the values of the AES parameters after evolution). The second and third experiments focus on the analysis of the parameters of the AES under internal disruptions, a constant disruption and a variable disruption, respectively.

The extended version of the EAHS, the EAHS-R (Section 2.2), which includes a hormone receptor, is compared with the standard EAHS in a forth experiment, in which both systems are validated on a simulated and real robot.

All experiments are an attempt to verify the performance of the EAHS in terms of the homeostatic regulation process and robustness to changes in environmental conditions. It is expected that such an endeavour will promote the automatic adjustment of variables under regulatory control.

3.1 Evolving the NSGasNets

GasNets are a class of ANNs that makes use of an analogue of volume signalling, whereby neurotransmitters freely diffuse into a relatively large volume around a nerve cell, potentially affecting many other neurons (Gally *et al.*, 1990). They are essentially recurrent ANNs augmented by a chemical signaling system comprising a diffusing virtual gas (or gases) which can modulate the response of other neurons.

A number of GasNet variants, inspired by different aspects of real nervous systems, have been explored in evolutionary robotics context as controllers for mobile autonomous robots. They have been shown to be significantly more evolvable, in terms of speed of evolution, than other forms of neural networks for a variety of robot tasks and behaviours (Husbands *et al.*, 1998; Philippides *et al.*, 2005; Vargas *et al.*, 2008). The model used in the work reported here is the NSGasNet (Vargas *et al.*, 2007).

The transfer function of the node i in the network is given by equation (5):

$$O_i(t) = \tanh\left[K_i(t)\left(\sum_{j\in C_i} w_{ji}O_j(t-1) + I_i(t)\right) + b_i\right] \tag{5}$$

where C_i is the set of nodes with connections to node i, w_{ji} is the connection weight value (ranging from -1 to $+1$), $O_j(t-1)$ is the previous output of neuron j, $I_i(t)$ is the external input to neuron i at time t, if the node has external inputs, b_i is the bias of the neuron and $K_i(t)$ represents the modulation of the transfer function caused by the gases.

The K_i^t parameter has its value determined from the set of equations (6)-(9):

$$K_i^t = P[D_i^t] \tag{6}$$

$$P = \{-4.0, -2.0, -1.0, -0.5, -0.25, -0.125, 0.0, 0.125, 0.25, 0.5, 1.0, 2.0, 4.0\} \tag{7}$$

$$D_i^t = f\left(D_i^0 + \frac{C_1^t}{C_0 \times K}\left(N - D_i^0\right) - \frac{C_2^t}{C_0 \times K} D_i^0\right) \tag{8}$$

$$f(x) = \begin{cases} 0, & x \leq 0 \\ \lfloor x \rfloor, & 0 < x < N \\ N - 1, & \text{otherwise} \end{cases} \tag{9}$$

where $P[i]$ is the set of values K_i^t can assume in the array of N positions, D_i^0 is the genetically defined value of D_i^t, C_1^t and C_2^t are the concentrations of gases 1 and 2 at time t, respectively. The transfer function K_i^t is increased by the presence of gas 1 and decreased by the presence of gas 2. C_0 and K are global constants.

The modulator bias (Mbias) is responsible for dictating to what extent the node could be affected by the gases emitted by all the other nodes, ranging from 0 to 1. A "0" value means that the neuron is not affected by the specified emitting neuron. A value above "0" means that the neuron will be affected by the specified emitting neuron, at a rate proportional to the stimulation level. Equation (10) defines the concentration of gas at the neuron:

$$C(t) = \text{Mbias} \times T(t) \tag{10}$$

Functions $T(t)$ and $H(x)$ (equations (11) and (12)) model the increase and decay of the gases; t_e and t_s are the last time the neuron started and ceased the emission of gas, respectively; s is a constant related to the build up and decay of the gas emission at each time step t:

$$T(t) = \begin{cases} H\left(\frac{t - t_e}{s}\right), & \text{emitting} \\ H\left(H\left(\frac{t_s - t_e}{s}\right) - H\left(\frac{t - t_s}{s}\right)\right), & \text{not emitting} \end{cases} \tag{11}$$

$$H(x) = \begin{cases} 0, & x \leq 0 \\ x, & 0 < x < 1 \\ 1, & \text{otherwise} \end{cases} \tag{12}$$

The network genotype consists of an array of integer variables lying in the range [0, 99] (each variable occupies a gene *locus*). The decoding from genotype to phenotype adopted is the same as the original model (Husbands *et al.*, 1998). The NSGasNet model has six variables associated with each node plus one modulator bias (Mbias_{ij}) for each node, plus task-dependent parameters. Therefore, networks N1 and N2 (Figure 1) will

have six Mbias for each node (for each network has six nodes as it will be described later in this section):

$$< \text{genotype} >::< \text{rec} >< \text{TE} >< \text{CE} >< D_i^0 >< b_i >< s > \\ < \text{Mbias}_{i1} > \cdots < \text{Mbias}_{(ij)} >$$

where the rec parameter determines if the recurrent connection is excitatory, inhibitory, or inexistent; TE stands for the circumstances under which the neuron will emit a gas: if its electrical activity exceeds a predetermined threshold, if the concentration of gas 1 exceeds a predetermined threshold, if the concentration of gas 2 exceeds a predetermined threshold, or if the neuron does not emit gases under any circumstance; CE specifies which gas the neuron emits, gas 1 or gas 2; D_i^0, b_i and s are referred in equations (8), (5) and (11), respectively, and the Mbias parameter is the modulation bias of the node. For a more detailed explanation of the mechanisms of GasNets and NSGasNets, the reader should refer to Husbands *et al.* (1998) and Vargas *et al.* (2007).

Network N1 (Figure 1), which is responsible for the straight-line motion with obstacle avoidance behaviour, has four inputs: the most stimulated left, right, front and back distance sensors (Figure 4(b)). Two additional neurons were considered to be output neurons, so the network consists of six neurons. The output neurons correspond to the motor neurons.

The fitness function (equation (13)) and the training scenario (Figure 5(a)) were inspired by the work of Nolfi and Floreano (2004):

$$\phi = V(1 - \sqrt{\Delta v})(1 - i) \tag{13}$$

where V is the sum of the instantaneous rotation speed of the wheels (stimulating high speeds), Δv the absolute value of the algebraic difference between the speeds of the wheels (stimulating forward movement), and i is the normalized value of the distance sensor of highest activation (stimulating obstacle avoidance).

A trial is considered to be 2,000 iterations of the control algorithm. At the end of each trial, the robot is randomly repositioned in the environment.

The structure of network N2 (Figure 1), which is responsible for the phototaxis behaviour, is similar to the obstacle avoidance network. Only, the distance sensors

(a) **(b)** **(c)**

Notes: (a) Obstacle avoidance; (b) phototaxis; and (c) for the AES

Figure 5.
Scenarios used for the evolution of the NSGasNets

were replaced by the luminosity sensors (Figure 4(b)). The environment consists of a square arena, where the robot has an initial fixed position at the beginning of each trial (Figure 5(b)). Each trial corresponds to 2,000 simulation steps. The fitness function is given by equation (14):

$$\phi = V(1 - i) \tag{14}$$

The parameter i (referring to sensory activation) is minimized when the robot is near the light, due to the sensory structure of the robot.

3.2 Evolving the AES

The genotype of the AES consists of four parameters: ω, θ, α and T (equations (1)-(3)). A trial is considered to be 800 iterations of the control algorithm. The parameters λ and ω stand for minimum light intensity and HL, respectively. β is the internal state (IS) growing rate. In our model, as in Moioli *et al.* (2008a), the IS of the artificial agent (equation (2)) stands forth inverse of the battery meter reading. It implies that the lower the battery level, the higher the IS. θ is the IS level threshold, above which hormone production starts at a rate dictated by α. T is simply the half-life of the hormone. As the battery is always discharging (given that the robot is turned on), β should have a predefined value associated with it. Similarly, the minimum light intensity above which the robot could recharge should also be predefined. The presence of light here indicates a recharging area.

The evolution starts with a robot exploring the arena (Figure 5(c)), controlled by the obstacle avoidance network. The AES was designed to sense the internal state of the robot. If the internal state goes above 90, on a 0-100 scale, the robot is considered to be "dead." To obtain a successful performance the robot should be able to efficiently switch between exploration behaviour and phototaxis behaviour. This switching is expected to be due to the production of the hormone related to the decrease of battery level. After battery recharging (associated with being close to the light), and the consequent decrease in the related HL, the robot should switch to its original exploratory behaviour. Equation (15) shows the fitness function adopted for this task:

$$\phi = \frac{V(1 - i)t}{M} \tag{15}$$

where V is the absolute value of the sum of the instantaneous rotation speed of the wheels (stimulating forward movement), i is the normalized value of the distance sensor of highest activation (stimulating obstacle avoidance), t is the number of iterations in which the robot remains alive, and M is the maximum number of iterations a trial can have. Thus, an acceptable performance would consist of adjusting the hormone production thresholds and growth rate in order to allow maximum exploration with regular recharging events. Owing to the environment set-up (Figure 5(c)), the robot could not stay close to the light sources when performing exploration of the environment for the light sources are themselves located close to the wall.

Figure 6 shows the fitness value (minimum, average and maximum of three experiments) throughout the evolutionary process, for the NSGasNets, (a) and (b), and the whole system (c). The final network configuration indicates that network N1 (for obstacle avoidance) has evolved a more symmetrical architecture in terms of connections and gas emission.

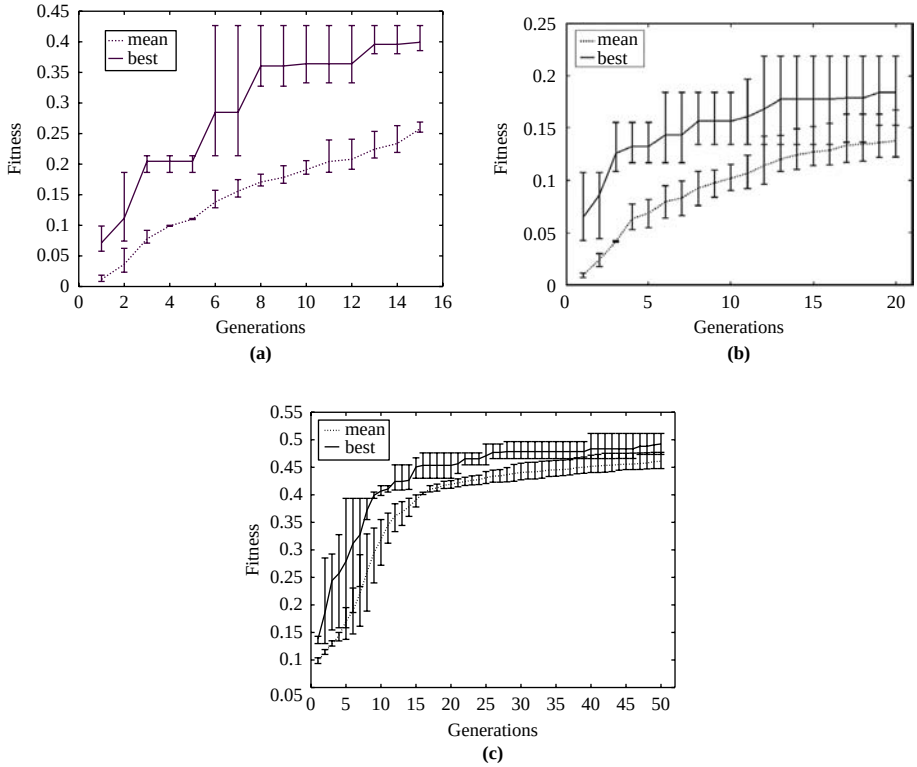

Notes: Mean values refer to a set of 3 experiments; (a) obstacle avoidance; (b) phototaxis; and (c) EAHS

Figure 6.
Fitness evolution

4. Results

The following experiments include previous results obtained with the original EAHS (Moioli *et al.*, 2008a) plus the performance comparison experiment between the former and the novel EAHS-R.

4.1 Experiment 1 – EAHS performance under normal conditions

After the evolution of the AES, the phenotype of the best individual was: $\alpha = 0.0099$; $T = 11.1$; $\omega = 50.5$ and $\theta = 52.0$. Parameters $\beta = 0.01$ and $\lambda = 103$ were defined empirically and kept fixed.

When the HL increases above ω, the robot stops exploring the environment and starts chasing the light. This confirms the influence of the HL over the robot's autonomous behaviour. During the experiment, the robot traverses the arena at maximum speed, only adjusting its speed when avoiding collision courses or when changing the behaviour to phototaxis.

4.2 Experiment 2 – EAHS performance under constant disruptions

This experiment aims at submitting the evolved artificial homeostatic system to perturbations in order to analyse its performance when dealing with situations not faced during evolution. The task consists of changing the value of the parameter β, thus simulating a faster or slower battery discharge. Figure 7 shows the trajectory and

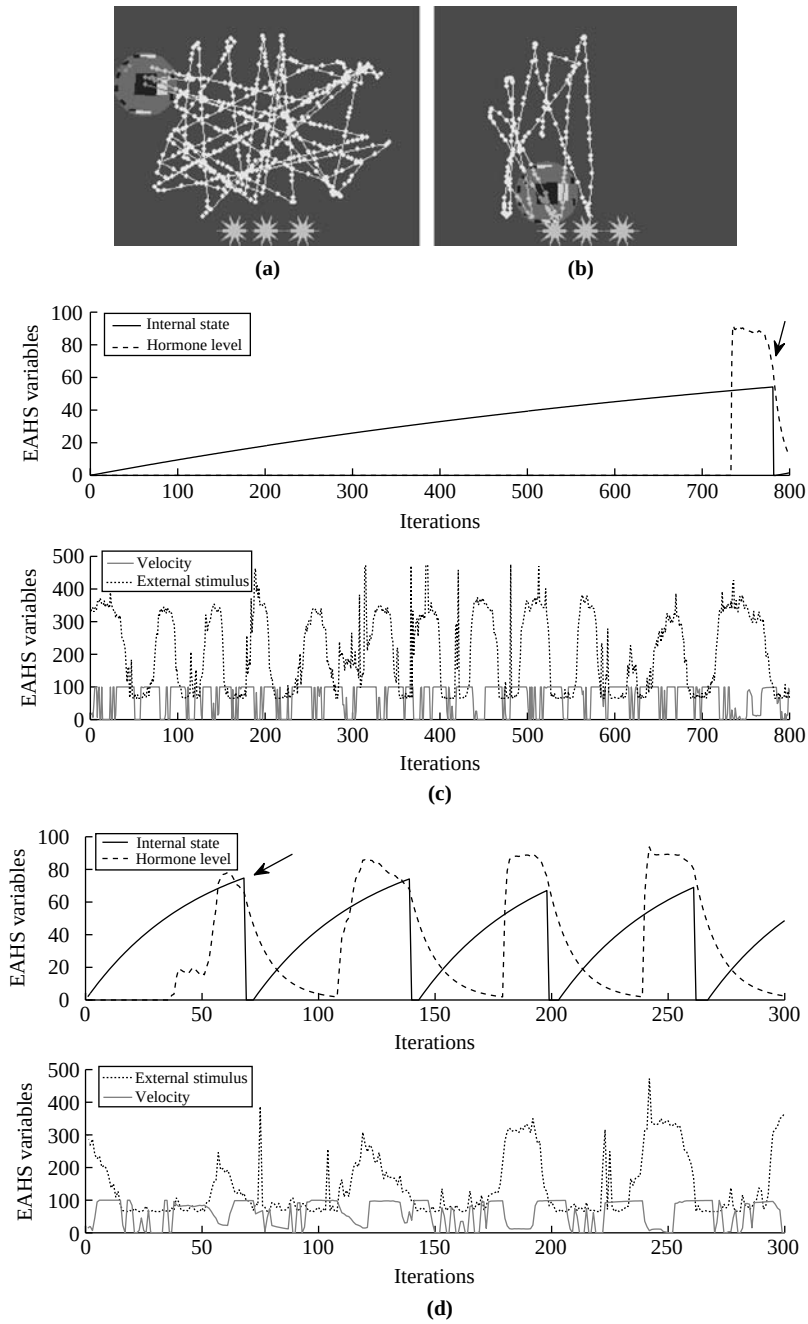

Figure 7.
Trajectories and the behaviour of EAHS variables for $\beta = 0.001$ ((a) and (c)) and for $\beta = 0.02$ ((b) and (d))

Note: The black arrow indicates the end of the phototaxis behaviour

the parameter analysis for two different values of β. Remember that, during evolution, β was kept fixed at the value 0.01. Figure 7(a) and (c) shows the results for $\beta = 0.001$, i.e. a slower discharge rate of the battery, and Figure 7(b) and (d) shows the results for $\beta = 0.02$, simulating a faster discharge rate.

In the first case (Figure 7(a)), the robot's IS grows slowly enough to allow the robot to explore the whole arena before recharging is needed (around iteration 800). However, when the discharge rate is increased, the robot has a smaller period of time to explore the arena between consecutive recharging (Figure 7(b)). The robot is very good in timing the switch of behaviour, avoiding getting too far from the battery charger and eventually "dying."

The behaviour of the velocity curve also indicates a quantitative change of behaviour. With a lower β, the robot only needs to turn (and consequently reduce its velocity) when facing a wall and sometimes when adjusting its direction. At higher values, β forces the robot to turn and seek the light more frequently, which requires operation at a lower velocity. Hence, on average, when the exploratory behaviour is enforced, the robot tends to move much faster.

4.3 Experiment 3 – EAHS performance under variable disruptions

In this experiment, the parameter β varies within a single trial and the speed of the robot is disrupted.

First, the absolute rotation speed of the wheels of the robot is linked to a greater energy consumption. In this way, the greater the rotation, the greater the value of β. β is allowed to vary linearly between 0.0015 and 0.015. Figure 8(a) shows the behaviour of the EAHS variables, including the β parameter. The robot was able to cope with variable disruption by successfully navigating and exploring the environment, chasing the light when recharge is necessary, and consequently maintaining its integrity.

Second, a disruption in the speed of the robot is simulated: the velocity value command is multiplied by 2 before reaching the robot wheels. But this information is not directly available to the robot or the control system. Figure 8(b) shows the experimental results. Once more the system was able to self-adapt to this unpredicted and unknown variable disruption, being robust enough to perform accordingly.

4.4 Experiment 4 – comparing the EAHS with the EAHS-R

4.4.1 Simulation. The experiments in this section are central for this work. The objective is to explore the relevance of the addition of a receptor to the EAHS. For the sake of simplicity, the EAHS with receptor was named EAHS-R (Section 2.2).

In order to compare the performance of both systems, one robot (named Robot1) controlled solely by the EAHS and another (named Robot2) controlled solely by the EAHS-R are placed on a larger arena which differs from the arenas used to evolve the entire system. They are validated separately for 400 iterations each or until the robot "dies."

Figure 9 shows Robot1's trajectory and Figure 10 shows the behaviour of the EAHS variables throughout the experiment. Note that Robot1, which is controlled solely by the EAHS, is not capable of dealing with the novel environment and "dies." This happens when the battery level falls beyond a threshold (at the instant represented by letter A in Figure 9 and iteration number 69 in Figure 10) and therefore the hormone was released signaling the moment to swap the behaviours from exploration to phototaxis.

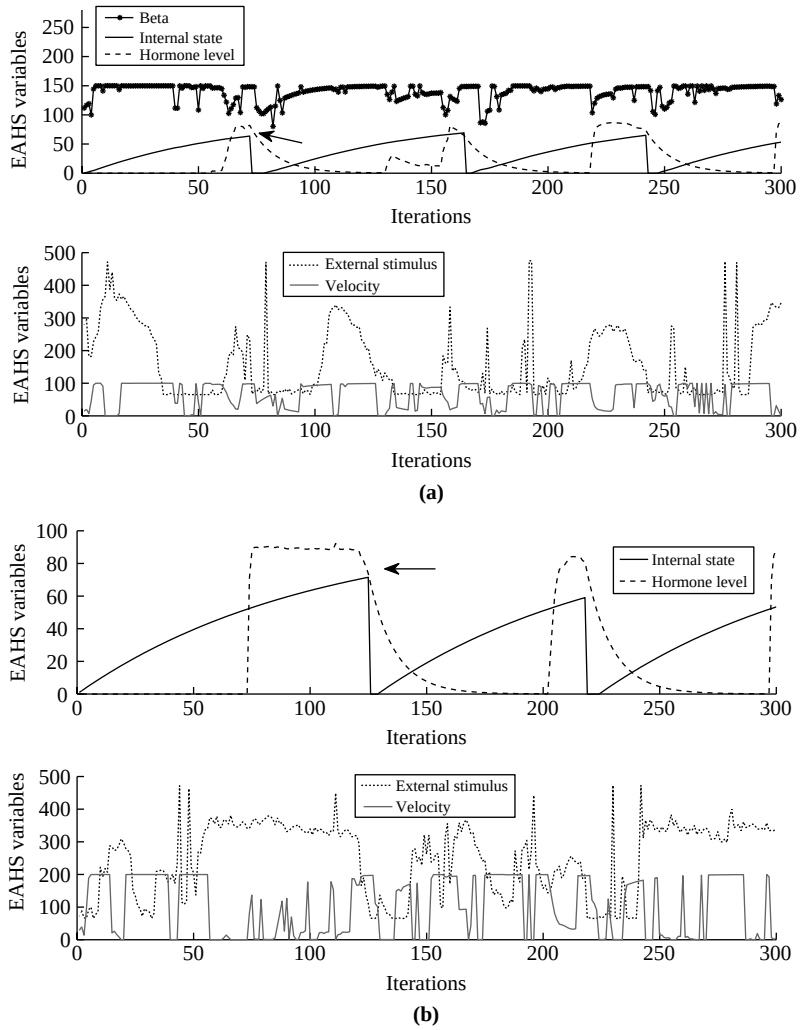

Figure 8.
EAHS variables for $0.005 \leq \beta \leq 0.015$ and for double velocity disruption, respectively

Note: The black arrow indicates the end of the phototaxis behaviour

However, as the robot sensors could not detect the light source (i.e. the battery charger), the robot starts to rotate until its battery is completely discharged.

Conversely, Robot2's trajectory (Figure 11) shows that it could slowly swap between both behaviours. It could take into account not only the HL but also the proximity of the light source. This particular instant is represented by letter B in Figure 11.

This superior performance can also be verified in Figure 12 which portrays the EAHS-R variables throughout the experiment. Note that the first time the hormone was released the robot was near the light source (ES) and therefore the receptor was active, allowing the direct switch to the phototaxis behaviour. Nonetheless, the second time this

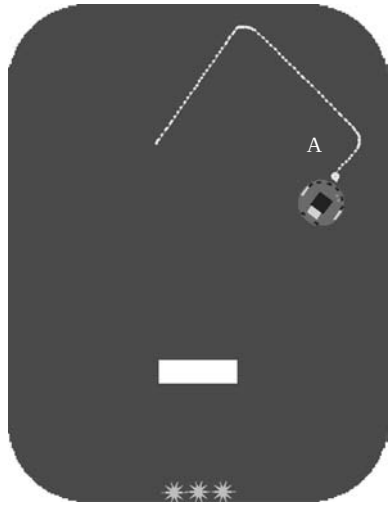

Figure 9.
Trajectory of the robot controlled by the EAHS

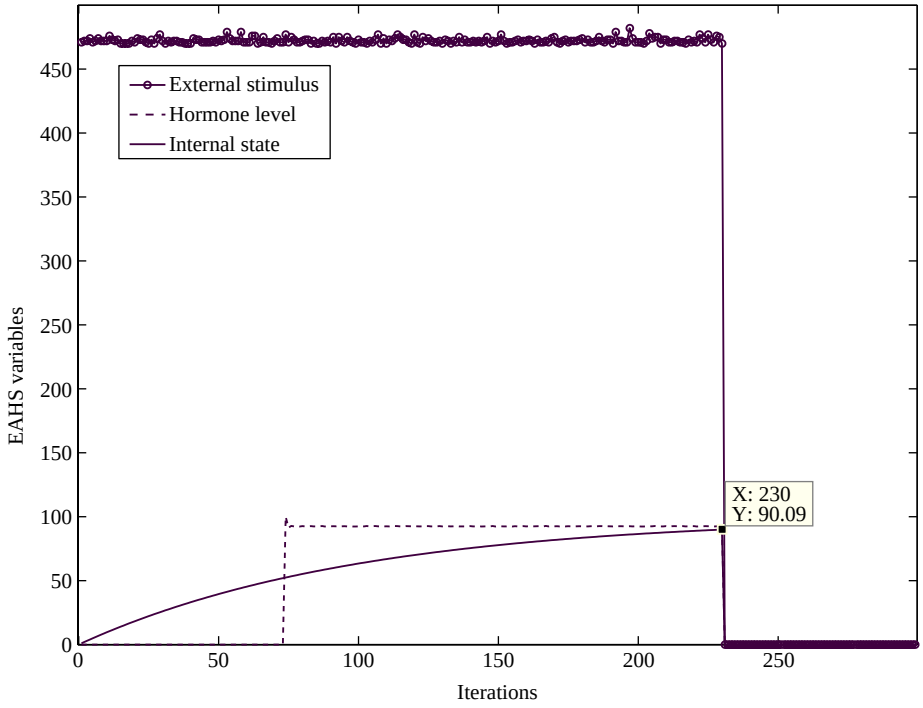

Figure 10.
EAHS variables HL, ES and IS throughout the experiment

happened, the robot was far from the light source and thus the receptor was inactive. Therefore, the switching to the phototaxis behaviour only happened when the robot achieved a region close to the light source (ES). In other words, although the HL was high, the robot only switched the behaviour when the receptor was active, thus allowing the hormone to be sensed by the ANNs.

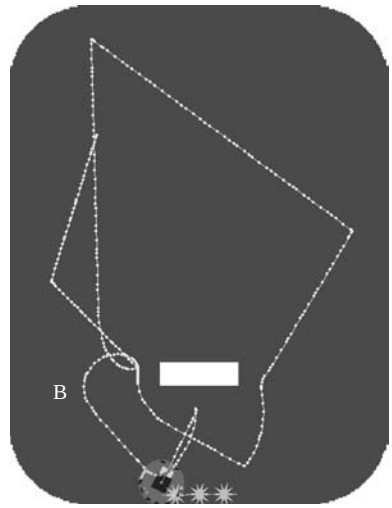

Figure 11.
Trajectory of the robot controlled by the EAHS-R

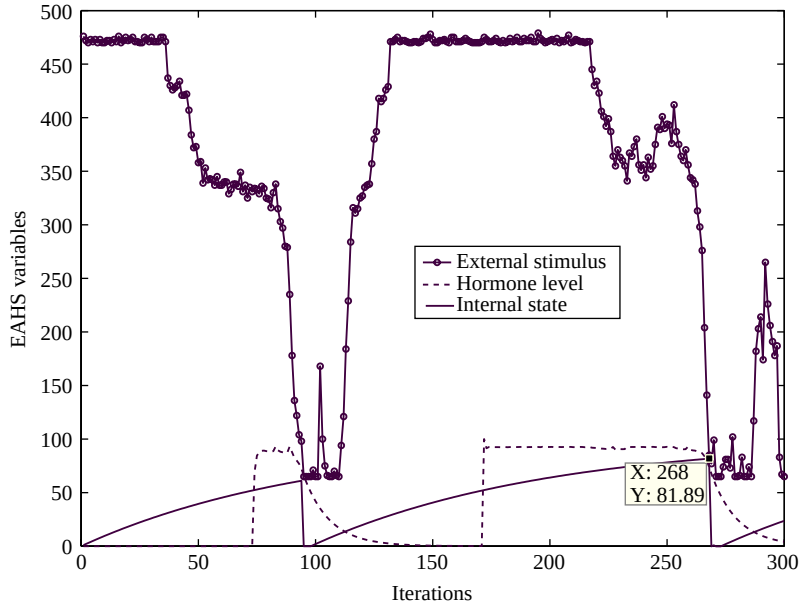

Figure 12.
EAHS-R variables HL, ES and IS throughout the experiment

4.4.2 Real robot. In order to verify whether the same performance could be observed in the real robot, an additional comparison was performed using the Khepera-II robot.

Figures 13 and 14 show both trajectories for Robot1 and Robot2, respectively. The letter A represents the instant when the robot needs to recharge its battery and letter B represents the moment when the battery is recharged.

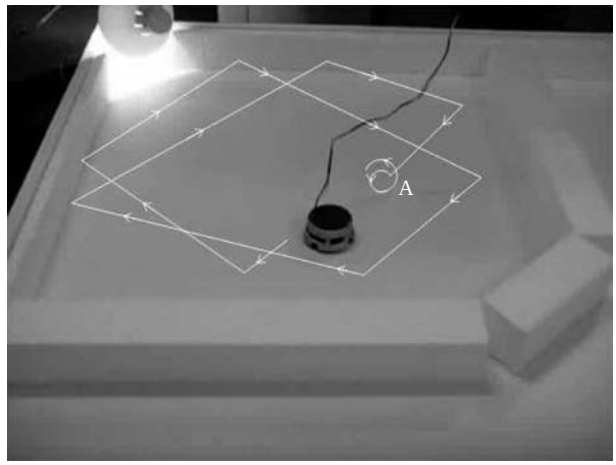

Figure 13.
Trajectory of the real robot controlled by the EAHS

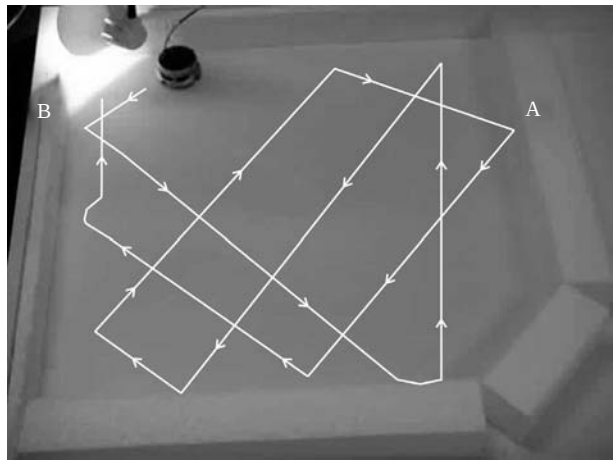

Figure 14.
Trajectory of the real robot controlled by the EAHS-R

Once more, Robot1 failed to cope with the novel arena, and hence "died." This instant is represented by letter A in Figure 13. This happened for the robot entered yet again a rotational trajectory until its battery was completely discharged.

On the other hand, Robot2 could successfully accomplish its task by managing to maintain its internal variables within acceptable limits, hence, ensuring its survival (Figure 14).

5. Discussion and future work

This work is a step forward in the design of an EAHS. Towards the goal of creating an even more robust system, we have introduced an artificial homeostatic system whose parameters are defined by means of an evolutionary process, and most importantly incorporates a hormone receptor.

The EAHS consists of two evolved ANNs coordinated by an AES. The ANNs followed the model proposed by Vargas *et al.* (2007), drawing inspirations from gaseous neuro-modulation in real brains. The objective was to design a more biologically plausible system inspired by homeostatic regulations pervasive in the human body, which is capable of exploring key issues in the context of behaviour adaptation and coordination.

This work deals with a flexible behaviour-based approach in the sense that coordination between modules of behaviour is evolved, not predesigned. A series of experiments were performed in order to investigate the performance of the system and its robustness to novel environmental conditions and internal sensory disruptions. The designed system showed to be robust enough to self-adapt to a wider variety of disruptions and novel environments by making full use of its in-built homeostatic mechanisms.

The incorporation of a receptor proved to be vital to cope with novel environments, which differ from the ones used during evolution in terms of dimension and target positioning. This fact could be observed on the simulated and real robots. That indicates the viability of the proposed method for coping with the reality gap, a well-known issue for the evolutionary robotics community (Nolfi and Floreano, 2004; Vargas *et al.*, 2010).

Another important outcome is illustrated by the precise coordination when β increases taking into account the severe changes on the environment and physical conditions of the robot. Recall from the experiments that β is the internal state gain factor (Section 2). When $\beta = 0.02$, the EAHS was able to adapt to the disruption, albeit reducing the exploration ratio. When employing only the exploration behaviour, if $\beta = 0.02$, the robot's internal state could eventually rise above 90, meaning its "death". This aspect reinforces the homeostatic regulatory behaviour of the proposed system.

Future work would include a deeper analysis of the neuro-endocrine interactions, eventually proposing mechanisms for the coordination of more complex sensorimotor behaviours. For instance, coordination of more than two behaviours and/or the implementation of more than one hormone, including multiple hormone receptors, are going to be investigated. Moreover, the automatic establishment of some of the predefined mechanisms of the current work (e.g. hormone-receptor interactions) should also be investigated.

References

Ashby, W.R. (1952), *Design for a Brain: The Origin of Adaptive Behaviour*, Chapman & Hall, London.

Avila-García, O. and Cañamero, L. (2004), "Using hormonal feedback to modulate action selection in a competitive scenario", in Schaal, S., Ijspeert, A., Billard, S. and Vijayakumar, J. (Eds), *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behaviour*, MIT Press, Cambridge, MA, pp. 243-52.

Besendovsky, H.O. and Del Rey, A. (1996), "Immune-neuro-endocrine interactions: facts and hypotheses", *Endocrine Reviews*, Vol. 17, pp. 64-102.

Cannon, W.B. (1929), "Organization for physiological homeostasis", *Physiological Review*, Vol. 9, pp. 399-431.

Collins, R. and Jefferson, D. (1991), "Selection in massively parallel genetic algorithms", *Proceedings of the 4th International Conference on Genetic Algorithms, ICGA-91*, Morgan Kaufmann, San Fransisco, CA, pp. 249-56.

Di Paolo, E.A. (2000), "Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions", *From Animals to Animals: Proceedings of the 6th International Conference on the Simulation of Adaptive Behavior*, MIT Press, Cambridge, MA, pp. 440-9.

Di Paolo, E.A. (2005), "Autopoiesis, adaptivity, teleology, agency", *Phenomenology and the Cognitive Sciences*, Vol. 4 No. 4, pp. 429-52.

Dyke, J.G. and Harvey, I.R. (2005), "Hysteresis and the limits of homeostasis: from daisyworld to phototaxis", in Capcarrere, M., Freitas, A., Bentley, J., Johnson, C. and Timmis, J. (Eds), *Advances in Artificial Life: 8th European Conference, ECAL 2005, Canterbury*, Springer, Berlin, 5-9 September, pp. 332-42.

Dyke, J.G. and Harvey, I.R. (2006), "Pushing up the daisies", *Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems*, MIT Press, Cambridge, MA, pp. 426-31.

Gally, J., Montague, P., Reeke, G. and Edelman, G. (1990), "The NO hypothesis: possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system", *Proceedings of the National Academy of Sciences USA*, Vol. 87, pp. 3547-51.

Grillner, S. (2006), "Biological pattern generation: the cellular and computational logic of networks in motion", *Neuron*, Vol. 52, pp. 751-66.

Guyton, A.C. and Hall, J.E. (1996), *Textbook of Medical Physiology*, 9th ed., W.B. Saunders Company, Philadelphia, PA.

Harvey, I. (2004), "Homeostasis and rein control: from daisyworld to active perception", *Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems, ALIFE9*, MIT Press, Cambridge, MA, pp. 309-14.

Hillis, W.D. (1990), "Co-evolving parasites improve simulated evolution as an optimization procedure", *Physica D*, Vol. 42 Nos 1/3, pp. 228-34.

Hoinville, T. and Henaff, P. (2004), "Comparative study of two homeostatic mechanisms in evolved neural controllers for legged locomotion", *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 3, pp. 2624-9.

Husbands, P., Smith, T., Jakobi, N. and O'Shea, M. (1998), "Better living through chemistry: evolving GasNets for robot control", *Connection Science*, Vol. 10 No. 4, pp. 185-210.

KTEAM, S.A. (2009), available at: www.k-team.com

Moioli, R.C., Vargas, P.A., Zuben, F.J.V. and Husbands, P. (2008a), "Evolving an artificial homeostatic system", *2008 SBIA Joint Conference (SBIA 2008)*, pp. 278-88.

Moioli, R.C., Vargas, P.A., Zuben, F.J.V. and Husbands, P. (2008b), "Towards the evolution of an artificial homeostatic system", *2008 IEEE Congress on Evolutionary Computation (CEC 2008)*, pp. 4024-31.

Neal, M. and Timmis, J. (2003), "Timidity: a useful mechanism for robot control", *Informatica*, Vol. 7, pp. 197-203.

Nolfi, S. and Floreano, D. (2004), *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-organizing Machines*, MIT Press, Cambridge, MA.

Pfeifer, R. and Scheier, C. (1999), *Understanding Intelligence*, MIT Press, Cambridge, MA.

Philippides, A., Husbands, P., Smith, T. and O'Shea, M. (2005), "Flexible couplings: diffusing neuromodulators and adaptive robotics", *Artificial Life*, Vol. 11 Nos 1/2, pp. 139-60.

Storm, T. (2004), KiKS, a Khepera simulator for Matlab 5.3 and 6.0, available at: http://theodor.zoomin.se/index/2866.html

Vargas, P.A., Di Paolo, E.A. and Husbands, P. (2007), "Preliminary investigations on the evolvability of a non-spatial GasNet model", *Proceedings of the 9th European Conference on Artificial life ECAL 2007*, Springer, Berlin, pp. 966-75.

Vargas, P.A., Di Paolo, E.A. and Husbands, P. (2008), "A study of GasNet spatial embedding in a delayed-response task", *Proceedings of the XIth International Conference on the Simulation and Synthesis of Living Systems, ALIFE-XI, Winchester, 5-8 August*, p. 640.

Vargas, P.A., Di Paolo, E.A., Harvey, I. and Husbands, P. (2010), *The Horizons of Evolutionary Robotics*, The MIT Press, Cambridge, MA (in press).

Vargas, P.A., Moioli, R.C., Castro, L.N., Timmis, J., Neal, M. and von Zuben, F. (2005), "Artificial homeostatic system: a novel approach", *Proceedings of the VIIIth European Conference on Artificial Life*, pp. 754-64.

About the authors

Patricia A. Vargas received her PhD in Computer Engineering from the University of Campinas, Unicamp (Brazil) in 2005. She is currently a Research Fellow at the School of Maths and Computer Science, Heriot-Watt University (Edinburgh, Scotland, UK). She was a post-doc at the Centre for Computational Neuroscience and Robotics, University of Sussex (England, UK) for the past three years. Her research interests include evolutionary robotics, ANNs, artificial immune and endocrine systems, learning classifier systems, biologically-inspired algorithms and robot autonomous navigation and control. Patricia A. Vargas is the corresponding author and can be contacted at: p.a.vargas@hw.ac.uk

Renan C. Moioli received his BSc in Electrical Engineering in 2006 and MSc in Computer Engineering in 2008 from the University of Campinas, Unicamp (Brazil). He is currently a PhD student at the Centre for Computational Neuroscience and Robotics at the University of Sussex (England, UK). His research interests include computational intelligence, evolutionary robotics, computational neuroscience and complex adaptive systems.

Fernando J. von Zuben is currently the Head of the Laboratory of Bioinformatics and Bioinspired Computing, and he is an Associate Professor at the Department of Computer Engineering and Industrial Automation, University of Campinas, Unicamp (Brazil). His research interests are computational intelligence, bioinspired computing, multivariate data analysis and machine learning. He coordinates open-ended research projects in these topics, tackling real-world problems in the areas of decision making, pattern recognition and discrete and continuous optimization.

Phil Husbands is a Professor of Artificial Intelligence and Co-director of the Centre for Computational Neuroscience and Robotics, University of Sussex (England, UK). His research interests include biologically inspired adaptive robotics, evolutionary systems, computational neuroscience and creative systems.