Extended Homeostatic Adaptation: Improving the Link between Internal and Behavioural Stability

Hiroyuki Iizuka[1] and Ezequiel A. Di Paolo[2]

[1] Department of Media Architecture, Future University-Hakodate
116-2 Kamedanakano-cho, Hakodate, Hokkaido, 041-8655, Japan
[2] Centre for Computational Neuroscience and Robotics,
Department of Informatics, University of Sussex
Brighton, BN1 9QH, UK

Abstract. This study presents an extended model of homeostatic adaptation designed to exploit the internal dynamics of a neural network in the absence of sensory input. In order to avoid typical convergence to asymptotic states under these conditions plastic changes in the network are induced in evolved neurocontrollers leading to a renewal of dynamics that may favour sensorimotor adaptation. Other measures are taken to avoid loss of internal variability (as caused, for instance, by synaptic strength saturation). The method allows the generation of reliable adaptation to morphological disruptions in a simple simulated vehicle using a homeostatic neurocontroller that has been selected to behave homeostatically while performing the desired behaviour but non-homeostatically in other circumstances. The performance is compared with simple homeostatic neural controllers that have only been selected for a positive link between internal and behavioural stability. The extended homeostatic networks perform much better and are more adaptive to morphological disruptions that have never been experienced before by the agents.

1 Introduction

The use of homeostatic neurons as the basic building blocks for evolved neurocontrollers proposed in [1] provides a novel approach to modelling adaptivity in artificial systems in a way that resembles the adaptive dynamics of living organisms. The idea behind homeostatic adaptation is based on that of the ultrastable system proposed by Ashby [7]. This is a system – open to interaction with the environment – that will tend to change its own configuration plastically whenever stability is lost and until it finds a new internal dynamics which will make the system stable under the current conditions. Such systems are capable of remarkable adaptation and learning. They have been applied to legged robot locomotion [3], extended to different types of plastic functions [4], applied to the study of the minimal dynamics of behavioural preference [5], and as a model of perseverative reaching in infants (A-not-B error) [10].

In the original model, a neural controller inspired by this system was combined for the first time with the techniques of evolutionary robotics. In this model, local

M. Asada et al. (Eds.): SAB 2008, LNAI 5040, pp. 1–11, 2008.

plastic mechanisms change the incoming synaptic weights only when neural activations move out of a bounded region that is defined in advance by the designer. Plasticity keeps working until the neural activations return to the homeostatic region resulting in a "stable" configuration in the sense that the network weights do not change further as long the firing remains bounded. This mechanism was implemented in a simulated agent evolved with a fitness function simultaneously rewarding phototaxis and the maintenance of neural activations within the homeostatic region. The use of intermittent plasticity in combination with this dual selective pressure allows controllers to evolve where an association is created between internal homeostasis and the desired behaviour. This association is evolved to be positive: *high homeostasis goes together with good performance*. By selective design, once a neurocontroller gives rise to the right sensorimotor coordination within a given environmental situation in a way that results in internal stability, synaptic weight changes no longer happen and the agent behaves as desired. If the situation changes, such as in an inversion of the visual field or some other sensorimotor perturbation, this causes a breakdown of coordination. Under these circumstances some evolved agents also show a breakdown of internal homeostasis demonstrating that some agents evolve at least one negative association: *lack of phototaxis induces lack of homeostasis*. As this happens, the local adaptive mechanism activates until it finds a new synaptic configuration which can sustain the activations within the homeostatic region. In these conditions, some evolved agents are also able to re-form the behavioural coordination (even if they had not been trained to adapt to the induced perturbation). These agents are then able to re-create a positive association: *regaining homeostasis induces a recovery of the original behavioural performances*.

However, the original work has a problem in that these necessary further associations between internal and behavioural stability that allow adaptation to unseen perturbations are *contingent*. They may or they may not evolve in the original setup. This contingency is demonstrated by the high fitness sometimes achieved by solutions for which, under disruption of phototaxis, homeostasis remains unaffected. This problem with the method was first noted in [2] where an alternative model more closely resembling Ashby's homeostat was presented as a proof of concept. That model, however, was limited in that it used Braitenberg-style controllers. Improving the method to avoid contingent solutions using dynamical neural networks remains an important challenge if homeostatic adaptation is to be applied more widely in other areas of autonomous robotics. In this paper we move closer towards this aim. We propose an extended homeostatic neural controller where neurons are biased to have a strong resting membrane potential and an additional fitness condition rewarding not only a positive link between homeostasis and a desired behaviour but also a negative one between the breakdown of homeostasis and undesired behaviour [6].

This paper will present the extended model of homeostatic adaptation and compare it with more basic versions in terms of adaptivity and evolvability. It will be shown that agents evolved in the extended model are more adaptive against unexperienced morphological disruptions and random initial weight connections.

2 Model

Our proposed method is implemented in a simulated mobile agent with a plastic neural controller. The simulated agent is faced with a single light source. The task for the agent is to approach the light source. This task is deliberately simple in order to understand the basic interactions between neural mechanisms, plasticity, homeostasis and selection pressures before moving into more complex applications. Certain modifications are made to the original method for studying homeostatic adaptation [1] with the aim of improving the chances of internal homeostasis being linked with performance both in the positive and negative senses. In all cases, the set of initial weights for a neural controller is given randomly at the beginning of a trial. This modified setting makes the task more difficult than the original one. The agent is expected to adapt to a suitable weight set by the plasticity through interaction with the environment.

Agent. An agent is modelled as a simulated wheeled robot with a circular body of radius 4 and two diametrically opposed motors. The motors can drive the agent forwards in a 2-D unlimited plane. The agent has four light sensors mounted at angles $\pm\pi/4, \pm 3\pi/4$ radians to the forward direction. Light from point sources impinges on sensors with a local intensity proportional to the source intensity and the inverse of the distance from sensor to source. The model includes the shadows produced by the agent's body.

Plastic controller. A fully connected continuous-time recurrent neural network (CTRNN) [8] with 8 neurons is used as the agent's controller. The equations are modified from their ordinary form in order to enhance the conditions for homeostasis to be able to activate intermittent plasticity. In particular, a result to avoid is that of evolved agents capable of moving away from lights and remaining homeostatic. The lack on sensory activation (typically leading a CTRNN to converge to some form of asymptotic dynamics) should have as a consequence a breakdown of homeostasis. As before, the local homeostatic condition is fulfilled if a node is firing within a specified range. In the absence of sensory activation, certain parameters controlling resting potentials, synaptic strengths and size of the homeostatic region should be chosen as to enhance the chances for evolution to find solutions with the desired internal/interactive associations. To realize it, two parameters, α and β, are added in the typical CTRNN equations and another parameter γ controls the size of the homeostatic region.

In addition, a common problem in plastic neural networks is that of weight saturation. Typically, Hebbian-like rules, will tend to drive synaptic strengths either to their maximum or minimum values. This can sometimes be avoided with the use of directional damping factors [2]. But the problem is the more general one of loss of variability. In order to expand the range of possible perturbations that the system can adapt to, it is crucial that the ultrastable dynamics be provided with enough variability for its plastic reconfigurations. A possible way of achieving this is to map the weight values specified by the plasticity rules into a continuous but non-monotonic space. Since we are interested in some minimal

transformation that will solve this problem while still using simple plastic rules we eschew biological plausibility and make connection strengths a sinusoidal function of the weights. The time evolution of the states of neurons is then expressed as:

$$\tau_i \dot{y_i} = -(y_i - \beta) + \sum_{j=1}^{N} \alpha \sin(w_{ji}) z_j(y_j) + I_i, \tag{1}$$

$$z_i(x) = 1/(1 + e^{-x}), \tag{2}$$

where y_i represents the cell potential of neuron i, z_i is the firing rate, τ_i (range $[0.1, 10]$) is its time constant, I_i represents the sensory input, which is given to only sensory neurons. The sensory input is calculated by multiplying the local light intensity by a gain parameter (range $[1, 100]$), which is genetically encoded. There are two neurons for controlling each motor. The motor output is calculated from the ratio of firing rates of the neurons, which is mapped onto the range $[-5, 5]$ and is then multiplied by a gain parameter (range $[1, 10]$). This is to avoid that the agent cannot move when the dynamics converges to a small value in the absence of sensory stimulations. The *synaptic strength* for the connection from neuron j to i is determined $\alpha \sin(w_{ji})$, where α is a network constant given genetically that regulates how much the pre-synaptic neurons can affect post-synaptic neurons. The parameter β (resting potential) determines the equilibrium point for a node in the total absence of input. The balance between the two parameters, α and β, will become very important, because if α is too small the firing rates converge to their resting potentials (notice that the equation does not include bias terms), and if it is too big, β will not have a very significant effect on the dynamics. Therefore, if β is negative and with an appropriate α, firing rates will tend to converge to a small values unless there is enough stimulus coming from the sensors. The use of such a balanced combination, together with an appropriate choice of homeostatic region, will make it hard for networks to remain within homeostatic bounds in the absence of sensory input.

The connection weights between neurons, w_{ij}, are randomly determined at the beginning of a trial and a plastic mechanism allows for the lifetime modification of the connections. A homeostatic region is described as the finite zero-value set of a plasticity function of the post-synaptic firing rate. This function is modulated by parameter γ (range $[0, 0.5]$), which is genetically determined. Weights from neuron i to j are updated according to :

$$\Delta w_{ij} = \eta_{ij}(1 - z_i)p(z_j), \tag{3}$$

$$p(x) = \begin{cases} 0 & x > \gamma \\ 1 - x/\gamma & \text{else} \end{cases} \tag{4}$$

where z_i and z_j are the firing rates of pre- and post-synaptic neurons, respectively, Δw_{ij} is the change per unit of time to w_{ij}, η_{ij} is a rate of change (range $[-1, 1]$), which is genetically set for each connection, and $p(x)$ is the plastic function that defines the homeostatic region. The reason why this is called homeostatic is that if z_j is more than γ, the weight connection does not change.

Otherwise, the plasticity works and the weight connection keeps changing until z_j is stabilized in the homeostatic region (more than γ). Here we can see that the effect of the balance between α and β can be (if α is small enough and β negative enough) to land the dynamics into the zone of active plasticity. The parameter, γ, is also evolved, however, the trivial solution of evolving $\gamma = 0$ so that neurons are always homeostatic is prevented by the need to use plasticity to organize the weight configuration at the beginning of a trial. The only way this can happen is for firing rates to move out of the homeostatic regions, hence gamma may evolve to be small, but not zero.

3 Evolutionary Setup

A population of agents is evolved using a rank-based genetic algorithm with elitism. All fixed network parameters, τ_i, η_{ij}, α, γ and the gains are represented by a real-valued vector ([0,1]) which is decoded linearly to the range corresponding to the parameters (with the exception of gain values which are exponentially scaled). Crossover and vector mutation operators, which adds a small random vector to the real-valued genotype, are used [9].

In the extension of the original method presented here, half of trials during the evaluation process correspond to the presence of light and the other half are carried out in the dark. The light condition consists of the serial presentation of 8 distant light sources that the agent must approach and remain close to. A single source is presented at a time for a relative long time period of 1000. In the dark condition, there is no light in the arena. Consequently, the network receives no sensory stimulus. The agent can move freely in the unlimited arena for the same period as in the light condition. There will be a selective pressure to evolve homeostatic dynamics in the light condition but to *avoid* homeostasis in the dark condition. The scheme is expected to evolve networks for which lack of sensory stimulation leads to non-homeostasis. Therefore, the agents are evaluated by measuring three factors: the proportion of time that the agent spends near the light source (at a distance less than 20 in this paper), f_s, the time-average of the proportion of neurons that have behaved homeostatically in the light condition, f_h, and that have *not* behaved homeostatically in the dark condition, f_{Nh}. The fitness function is given by this, $F = \zeta * f_s * f_h + (1 - \zeta) * avg.(f_s) * f_{Nh}$, where ζ decides the weighted selective pressure between light and dark conditions.

4 Results

The performance of the extended homeostatic neural controller is compared to that of more basic forms of homeostatic neural controllers not using the biased resting potential and that have only been evolved only in the light condition. This basic homeostatic neural controller can be described in our formulations with the two parameters, β and ζ, which are set to 0 and 1.0, respectively. For the extended version, β and ζ are set to -5 and 0.5, respectively. In order to confirm whether the basic homeostatic neural network with the evaluation of light/dark

conditions can evolve agents properly to link internal and behavioural stability, a controller with $(\beta, \zeta) = (0, 0.5)$ is also tested. It should be noted that this basic homeostatic neural network has a different form to that of the original homeostatic neural network proposed by Di Paolo [1] but follows the same basic selection scheme.

4.1 Evolvability and Adaptivity

All homeostatic neural controllers are evolved for 4000 generations. Ten independent runs were made for each of the three conditions. Extending the number of generations to 10000 has not produced any observable increase in fitness.

The average of the best fitness values across 10 runs are shown in Fig. 1. It is clear that the extended homeostatic networks evolve better than the other conditions which become saturated at fitness ceilings of about 0.2 and 0.3 in each case. The main observed difference between these conditions and the extended one $(\beta, \zeta) = (-5, 0.5)$ is that agents are not able to sustain a configuration that allows them to perform phototaxis reliably.

In order to study the homeostatic controllers evolved by the three conditions, we select one successful agent at the 4000th generation from each run and investigate them in terms of their adaptivity.

First, we study adaptivity against variation of the initial weights by measuring the performances by the best agents for different initial weight configurations. The performances for phototaxis and neural homeostasis in 100 independent trials are shown in Fig. 2.

The best agents evolved using $\beta = 0$ can in some cases establish phototaxis but fail to adapt against the many initial weight configurations. It can also be

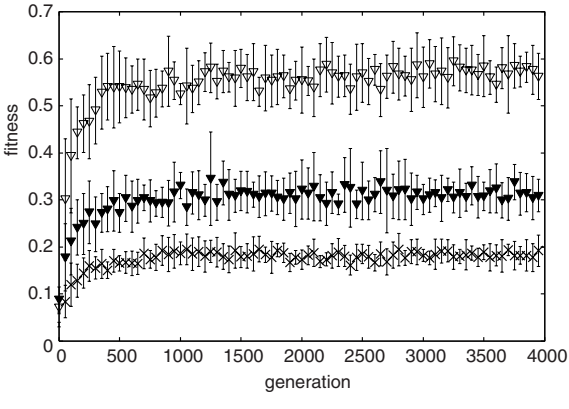

Fig. 1. Average fitness values of the best agents at each GA generation over 10 independent runs evolved with three kinds of the evolutionary regimes. The empty triangles are the results of agents with the extended method, $(\beta, \zeta) = (-5, 0.5)$, and the filled ones are for homeostatic networks without the biased resting potential nor evaluation in the dark condition, $(\beta, \zeta) = (0, 1)$. The crosses are for the one without the biased resting potential and with evaluation of both conditions, $(\beta, \zeta) = (0, 0.5)$.

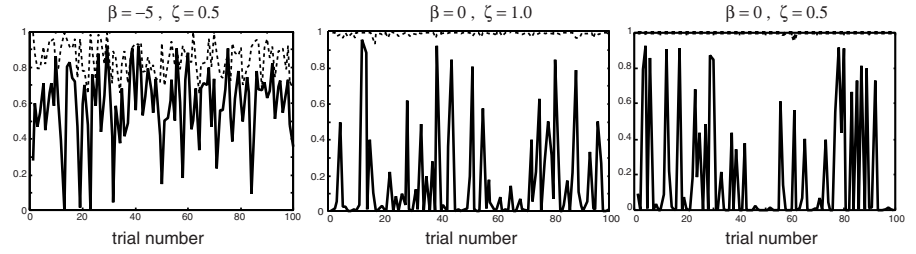

Fig. 2. A comparison of the performances for random initial weights of the best agents evolved with different evolutionary regimes, i.e. our proposed method (left), the basic homeostatic network (center), and the basic one with light/dark evaluation (right). The solid lines show the proportion of time that the agent spends within 30-units distance to light sources (initial distances to the light when it appears are 60). The dashed lines represent the time-averages of the proportion of neurons that have behaved homeostatically in each trial. A single trial consists of the serial presentation of 50 distant light sources. Trial numbers indicate different random seeds for weight connections.

observed that the time-averages of the homeostatic neural firing are always high regardless of the behaviour. This exemplifies the contingency of the link between the two selected evolutionary aims: homeostasis remains unaffected even under disruption of phototaxis since the two are realized independently. On the other hand, the extended condition adapts more widely to the initial random configuration. Depending on the initial weight values, the network has a good chance of establishing an appropriate sensorimotor coordination that results in phototaxis. Notice how the level of homeostasis correlates with phototaxis success. This indicates that agents that do not perform the correct behaviour will also tend to be non-homeostatic.

In the original homeostatic adaptation model agents show adaptation to radical unexperienced perturbations such as swapping sensor positions left and right. Although it is not always guaranteed that such perturbations can be adapted to, the extended model should be able to adapt more reliably than the basic homeostatic controllers because it decreases the chances of internal and behavioural stability being independent of each other. Adaptation to sensory inversion is measured by checking how much the agent can return to phototactic behaviour in the new condition after having adapted to the random initial weights. A recovery of phototaxis is defined as approaching at least 10 lights in sequence at some point after the sensor inversion.

Once having adapted to the initial weights, each pair of the diagonal sensors are swapped left and right ($\pm\pi/4 \leftrightarrow \mp 3\pi/4$). If the agent is able to recover phototactic behaviour by the 300th light source presented, it is counted as a success. The success rate is calculated over 100 trials after having adapted successfully to the initial weights. Results are shown in Fig. 3. All controllers evolved with the extended method can return to the phototactic behaviour with high probabilities. Having evolved under a pressure to behave non-homeostatically in the absence of sensor input has induced a condition where plasticity remains active

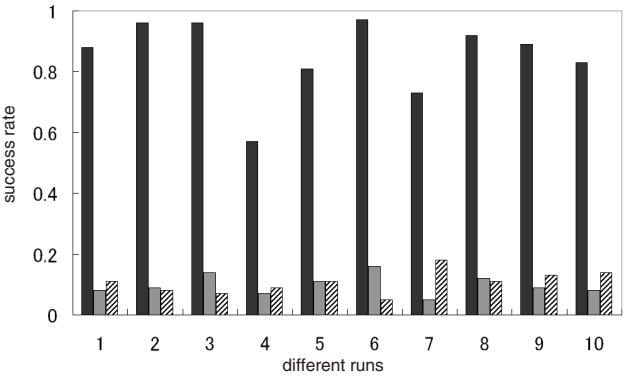

Fig. 3. A comparison of adaptivity to a new sensor configuration after the best agents have established the normal sensorimotor coupling. The new condition is achieved by swapping each pair of diagonal sensors ($\pm\pi/4 \leftrightarrow \mp 3\pi/4$). The vertical axis shows the rate of each best agent successfully showing the phototactic behaviour after the disruption. From the left, the bars means the results of the best agent using the homeostatic neural controller with ($\beta = -5, \zeta = 0.5$), ($\beta = 0, \zeta = 1.0$), and ($\beta = 0, \zeta = 0.5$), respectively. See text for details.

until sensors become active again. The system behaves in the desired ultrastable manner and plastic changes stop only as phototaxis is recovered.

4.2 Homeostatic Adaptation at Work

An example of the lifetime adaptation to the initial weight configuration by the best agent with the extended method is shown in Fig. 4. At the beginning of the trial, weight connections are randomly set so that the resulting behaviour cannot be approaching the light source. If too little light stimulus is provided the neurons cannot maintain the neural dynamics in the homeostatic region. This is the combined effect of using a biased resting potential and the extra fitness constraint for the dark condition. Following the plastic rules, the network starts changing the network structures that can lead to homeostasis and phototaxis at the same time. Regardless of the initial weight values, the best agent can successfully establish both phototactic behaviour and homeostasis. Converged weight values are very different each time the agent is re-initialized since the established sensorimotor coupling is dynamically constructed through the interaction.

The best agents also show adaptivity to inversion of visual field. After having adapted the normal positions of sensors, the left-right swapping is applied. Figure 5(left) shows the distances from the agent to the light sources. The light sources appear at a new place in every 1000 steps and the sensors are swapped when the 13th light source appears. Adaptation to the initial weight set happens before swapping sensors. When the sensors are swapped, the agent moves predictably away from the light source. This causes a breakdown of internal homeostasis and synaptic plasticity is turned on. After a period of adaptation

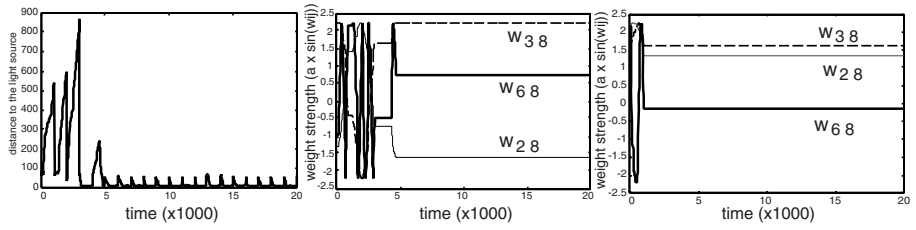

Fig. 4. Left: Distance from robot to sources. Each source lasts for 1000 time steps. Center: Change of synaptic weights corresponding to the same run of the left. For the clarity, only three of weights are shown. Right: Another example of synaptic changes of the same agent starting from different initial configurations.

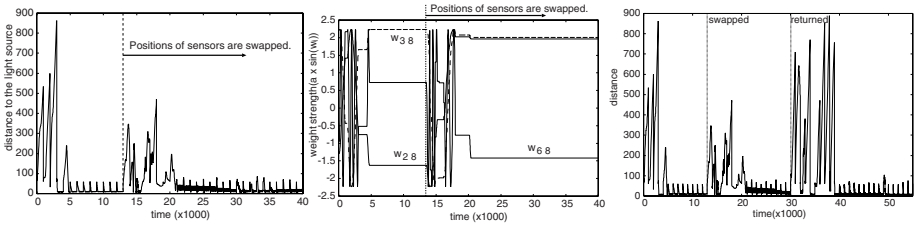

Fig. 5. Left: Distance from robot to sources. The vertical dashed line shows onset of swapping diagonal sensors ($\pm\pi/4 \leftrightarrow \mp 3\pi/4$). All initial configurations are same as Fig. 4. Center: Synaptic changes over time. Right: Re-adaptation to swapped sensor positions. The sensors are swapped at the appearance of 13th light source and are returned to original configuration when the 30th light appears.

the network finds a new homeostatic state that can approach the light. Once an appropriate new sensorimotor coupling has been established, weights becomes stable again (Fig. 5(center)). Re-adaptation is also tested. The measures that enhance dynamical variability in the extended model allow the agent to re-adapt when sensors are returned to the original configuration, something that was not observed in the original model (Fig. 5(right)).

5 Conclusions

This paper presents an extended homeostatic adaptation method capable of improved adaptive performance. The method works by addressing the problem of contingent associations between internal and behavioural stability in the original homeostatic adaptation model. The contingent link between these two requirements could always lead to solutions where behavioural disruption leaves internal stability unaffected. This situation, presumably does not occur so easily in living systems that must behave in a goal-directed manner in most cases *in order to* guarantee continued inner stability. Our proposed method moves closer to this natural situation by explicitly selecting an association between the undesired

behaviour and lack of homeostasis and using a strong negative resting potential for neurons that contributes to this association. It also moves closer to the original Ashbyan idea of ultrastability by preventing loss of variability during the adaptation process (through avoidance of synaptic strength saturation). As a result, evolved agents show improved adaptivity to previously unseen radical perturbations such as sensor inversions.

This result leads to wider lessons despite the simplicity of the task. In conventional evolutionary robotics models, sensorimotor connections as low level descriptions are tightly related to the behavioural performances as macro observations. Even in those cases where sensorimotor coordination is allowed to change plastically, the appropriate relation is established by an external process of artificial selection. In the homeostatic adaptation models, the sensorimotor coupling is intermittently reconfigured by the system to maintain the internal requirement of homeostasis, which is associated with the desired behaviour by the external selective process. Therefore, the bottom-up construction from the sensorimotor dynamics and the top-down regulation from the behavioural performances are mutually coupled processes *intrinsic* to the system. It is their 'high-level' link that is constructed by artificial evolution. The careful construction of this link can produce systems that are more adaptive and more lifelike.

Acknowledgements

This research was partially supported by the Japanese Ministry of Education, Science, Sports and Culture, Grant-in-Aid for JSPS Fellows, 17-04443.

References

1. Di Paolo, E.A.: Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In: From Animals to Animats VI: Proceedings of the 6th International Conference on Simulation of Adaptive Behavior, pp. 440–449 (2000)
2. Di Paolo, E.A.: Spike timing dependent plasticity for evolved robots. Adaptive Behavior 10(3/4), 243–263 (2002)
3. Hoinville, T., Henaff, P.: Evolving plastic neural controllers stabilized by homeostatic mechanisms for adaptation to a perturbation. In: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems, pp. 81–87 (2004)
4. Williams, H.: Homeostatic plasticity in recurrent neural networks. In: From Animals to Animats 8: Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior, pp. 344–353 (2004)
5. Iizuka, H., Di Paolo, E.A.: Toward Spinozist robotics: Exploring the minimal dynamics of behavioural preference. Adaptive Behavior 15(4), 359–376 (2007)
6. Iizuka, H.: Evolving homeostatic neural controller without depending on initial weights. In: Proceedings of 13th International Symposium of Artificial Life and Robotics (2008)
7. Ashby, W.R.: Design for a brain: The origin of adaptive behaviour, 2nd edn. Chapman and Hall, London (1960)

8. Beer, R.D.: Intelligence as adaptive behavior: An experiment in computational neuroscience. Academic Press, San Diego (1990)
9. Beer, R.D.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In: From Animals to Animats 4: Proceedings of the 4th International Conference on Simulation of Adaptive Behavior, pp. 421–429 (1996)
10. Wood, R., Di Paolo, E.A.: New models for old questions: Evolutionary robotics and the 'A Not B' error. In: Proceedings of 9th European Conference on Artificial Life, pp. 1141–1150 (2007)