

Ultrastable Neuroendocrine Robot Controller

Lenka Pitonakova

Note: A stylistically different version of this article was published in *Adaptive Behaviour* (2013), vol. 21, no. 1, p. 47–63

Abstract: The work of a simulated neuroendocrine controller with ultrastable neurons and glands is sketched and tested in terms of stability and adaptability. The artificial neurons connect to each other and to motors, while hormones produced by behaviour-related glands regulate their output. The ultrastable nature of the cells allows them to maintain their homeostasis by random reconfiguration of their connections and parameters without reference to the global goal of the system. Interactions of these ultrastable components cause individual robot behaviours to emerge to certain extents. The presented results show that the controller as a whole is capable of not only configuring itself to perform random walk, obstacle avoidance, mineral collection and recharging, but also to stay robust or adapt to a number of environmental and body perturbations without a need for a body model.

Keywords: distributed adaptation, ultrastability, hormonal modulation

1. Introduction

In the field of robotics, we continuously try to build better and more autonomous machines that could one day operate in unknown environments with no need for human intervention. Biological systems have developed mechanisms that support autonomy through maintenance of homeostasis. Minor behavioural alterations are often controlled by the amount of hormones in the body, whereas adaptation to rapid environmental changes involves alteration of survival strategies.

Biologically inspired engineering thus seems to be becoming more and more important when it comes to creation of autonomous agents. Two distinct groups of biologically inspired robot controllers are considered in this paper. On one hand, there are hormonally modulated controllers where different (but often pre-learned) behaviours are executed based on lev-

els of behaviourally dependent hormones (Neal and Timmis, 2003; Shen et al., 2004; French and namero, 2005; Vargas et al., 2005; Moili et al., 2008; Timmis et al., 2009; Sauzé and Neal, 2010; Liang et al., 2010; Schmickl et al., 2011), providing a robot with a need-based action selection. Work particularly relevant to this paper describes controllers where roles of artificial cells are adapted during the life-time of a robot based on hormonal levels (Hamann et al., 2010; Schmickl et al., 2011) and where hormonal release itself is adapted (Timmis et al., 2009).

On the other hand, there are ultrastable controllers that do not use action selection but attempt to maintain optimal levels of variables that are essential for their work. An ultrastable unit alters its own parameters in order to change the way actions are executed, providing adaptation driven by the goal of internal stability (Ashby, 1960; Manicka and Di Paolo, 2009; Portha et al., 2004).

This paper proposes an ultrastable neuroendocrine controller (UNEC) where both of these principles work in parallel, merging their advantages in a single simulated agent. The controller consists of ultrastable self-reconfiguring neurons and glands, commonly referred to as cells. The neurons are organised in a continuous-time self-recurrent neural network (CTRNN) (Beer, 1995) and are affected by hormones produced by glands, providing the robot with an action-selection mechanism. In the spirit of ultrastability, both cell types use a small Gaussian-random reconfiguration of their outgoing connections that is triggered each time a cell's activation is out of its 'comfort zone' for a certain amount of time. Stability of a connection increases with time elapsed since its last reconfiguration. All controller behaviours, including hormonal release itself, emerge to various extents from behaviours of the self-reconfiguring cells. Furthermore, the ultrastable nature of cells implies that there are no distinguished 'training' and 'trained behaviour' of the controller. Any behaviour can be changed at any point in time if constraints given to cells responsible for that behaviour are not satisfied. This paper approaches robot control in a novel way: the main question asked is not how to execute agent-level actions but why cells would need to execute them in the first place.

The agent is given a task of resource gathering and has to avoid obstacles, as well as maintain level of its energy by recharging near specified objects. It is hypothesised that behaviour of an UNEC-controlled agent will be robust to minor perturbations due to the elasticity that ultrastability affords. More severe perturbations are expected to degrade the controller performance but adaptation should return the performance close to normal, especially when body changes will occur while the environmental conditions will remain the same.

The following section explains the rationale behind hormonal modulation and ultrastability. Section 3 describes the controller mechanics, followed by Section 4 with preliminary behavioural and adaptation results. A full hormone-modulated controller is then tested with the methods described in Section 5 and the results are presented in Section 6. The nature of the proposed control mechanism, its limitations

and possible implications for robotics are discussed in Section 7.

2. Background

Despite our engineering efforts, the ability to withstand and adapt to various environmental conditions and changes to the body is still a unique trait of living organisms. For example, ants and bees sample food sources during foraging and abandon them if they are not suitable. They take into account not only the presence of needed substances and their balance in a food source (Josens and Rocas, 2000; Portha et al., 2004), but also their intake rate (O'Brien and Hooper-Bùi, 2005). A similar behaviour was recorded in domestic cats exposed to experimental diets (Hewson-Hughes et al., 2011). Furthermore, Lefebvre et al. (2007) showed that the patch-leaving behaviour of honeybees during nectar collection is related to quality of a visited patch as well as its distance from the nest. These and other experiments show that living organisms try to maintain the presence of essential elements in their bodies within a certain range (Simpson and Raubenheimer, 1993). At the same time, they indicate that biological entities are able to make effective trade-offs in food intake when different food sources provide varied amounts of required substances in order to satisfy as many of their needs as possible.

Biological adaptation even goes beyond short-term decisions and involves changing of behavioural strategies. For instance, Josens and Rocas (2000) showed that motivation to start foraging is largely affected by the nutritional state of an insect colony and that homeostatic mechanisms affect the type of behaviour that is executed. Their experiments with food-deprived ants revealed that higher motivation to eat suppressed behaviours such as carrying of food to the nest and scent marking of the food location.

There are various mechanisms in organic bodies that maintain the level of nutrients, water, temperature and other attributes in order to make cells function optimally. This phenomenon is referred to as homeostasis. Vertebrates evolved an especially complicated structure consisting of the neural, endocrine and immune system in order to main-

tain homeostasis (Neal and Timmis, 2003). Cells of the immune system protect the organism from antibodies, such as viruses and bacteria (Tsankova, 2009), and together with endocrine glands modulate the neural system through hormones (Simpson and Raubenheimer, 1993; Neal and Timmis, 2003). Hormones are produced when changes in the body are detected and their levels decay over time specific to a particular hormone (Neal and Timmis, 2003). For example, the information about level of nutrients in the body is carried through blood by hormones that affect taste sensors (Simpson and Raubenheimer, 1993) and by extension feeding. Neurons alter their synaptic strengths and consequently their spiking frequency (Surmeier et al., 2009) if they have appropriate hormone receptors and the detected level of the hormone reaches a specific threshold (Kaltenbach, 1988; Liang et al., 2010). Neurons thereby contribute to the homeostasis of the whole system by being homeostatic themselves, even if the system has no cognitive abilities (Sánchez-Lasheras et al., 2011).

From an engineering perspective, hormones found in bodily fluids maintain a global pool of information that the neural system reads from and is influenced by. As soon as activity of an organism makes its internal state change, glands seize production of a hormone. A basic experiment by Neal and Timmis (2003) showed that if an artificial hormone that is produced when obstacles are detected alters weights of a simple neural network responsible for random walk and obstacle avoidance, the behaviour can be more adaptive than if no hormonal control is used. The controller could thus produce both calm roaming in uncluttered environments and cautious obstacle approach in narrow corridors. A similar principle was used to drive energy replenishment through phototaxis (French and namero, 2005; Vargas et al., 2005; Moioli et al., 2008). Sauzé and Neal (2010) used several hormones in a sailing boat controller with pre-learned sensorimotor control to keep a boat in operation and drive it towards specified targets. Artificial hormones were also used to control pre-evolved GasNets that could avoid obstacles and steer towards or away from light sources based on the current homeostatic need (Moioli et al., 2009). The

Digital Hormone Model by Shen et al. (2004) used hormones during inter-robot communication in order to allow for collective seizing of a target, balanced spreading in space and even repairing swarm topological damage. This model is similar to the UNEC as it implies distributed control where behaviour of units is co-dependent, but is not unit-specific. Throughout all the above-mentioned experiments, hormonal modulation was used to temporarily change weights of neural connections in order to create an action-selection mechanism. The UNEC follows a similar approach but also considers adaptation of the hormonal modulation itself.

One kind of mechanism where action selection was adapted was implemented by Timmis et al. (2009), who used pre-evolved neural networks, outputs of which were affected by hormones and fed to a single-layer perceptron controlling a robot's motors. The controller was able to self-configure as well as adapt its hormonal secretion using a form of Hebbian learning. However, the reported time needed for the initial configuration was up to 45 minutes and not all of the adaptation experiments were successful due to uncontrolled connection weight increase. Another model proposed by Hamann et al. (2010), and also used by Schmickl et al. (2011), involved the use of artificial hormones that affected module actuators in a multi-modular robotic organism. The authors showed that continuous observing of hormonal gradients allowed the robotic organism to adapt to different body configurations, since actions of the individual modules could be executed based on positions of these modules within the body. Furthermore, a similar mechanism was applied to a swarm of robots that demonstrated successful collective obstacle avoidance and phototaxis (Schmickl et al., 2010).

A different adaptation mechanism first described by Ashby (1960) that builds on the principle of internal stability rather than on performance-driven behaviour (Portha et al., 2004) deals with adaptation directly. According to Ashby, an organism acts in order to maintain levels of its 'essential variables', such as energy, water, etc., in a (conceivably genetically specified) 'viability zone' and reconfigures randomly through small steps when conditions change and the levels of these variables cannot be maintained. While

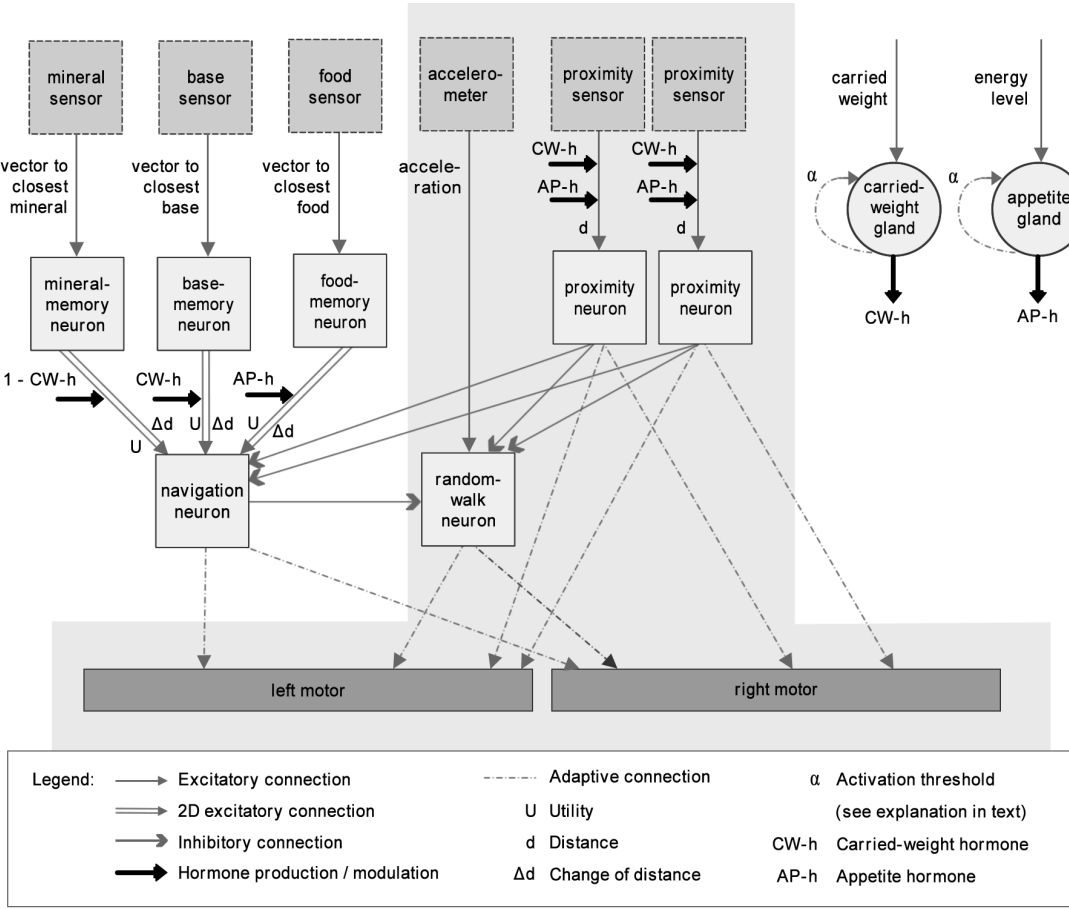


Figure 1: Schema of the ultrastable neuroendocrine controller. Only two out of eight proximity sensor-neuron pairs are shown for clarity. Schema of the initial controller described in Section 3.1 where hormonal modulation was not used is highlighted in the middle grey area.

interaction with the environment happens through sensorimotor feedback in the primary loop, the essential variables are observed by a slower secondary loop that can change the parameters of the system via random step functions. Ashby called this principle ultrastability and demonstrated its effectiveness by building Homeostat, a machine that was able to reach a stable state even after serious perturbations because of the distributed nature of its adaptation.

Di Paolo (2000) applied the principle of essential variables in a phototactic robot that was able

to adapt to sensory inversion and also worked on an ultrastable programmable spring machine (Manicka and Di Paolo, 2009). Montebelli et al. (2010) used a similar but non-self-reconfigurable principle to create a CTRNN driving an evolved ECOBOT controller that was able to collect water and food to sustain its energy. Despite the indications that ultrastability is beneficial for adaptation, its use in engineering is rather rare. According to Manicka and Di Paolo (2009), choosing an effective state change mechanism is a potential problem, since ran-

dom reconnection of elements could take a long time. In their paper, they used semi-random self-reconfiguration by identifying spring transitions that caused the system's behaviour to deviate from ideal and changed them until a desired behaviour was reached. The change was a product of the measured deviation, a mutation rate and a random number. The authors pointed out that maintenance of stability of found solutions was problematic, since introducing more randomness into their mutations caused the system to deviate from good solutions. Other authors implementing on-line learning (Iizuka and Di Paolo, 2008; Timmis et al., 2009; Santos et al., 2010) also altered weights in their neural networks by calculating the difference between actual and expected outcomes of the behaviour, but without introducing randomness.

On the other hand, the UNEC attempts to implement the Ashby's original concept by changing its weights in a Gaussian-random fashion, without any reference to a desired outcome. Stability is achieved by adding a secondary homeostatic loop as described by Ashby. Furthermore, plasticity of connections is based on the time elapsed since their last change, giving well-working configurations a higher chance to survive minor perturbations.

3. Ultrastable neuroendocrine framework

3.1. The basics: random walk and obstacle avoidance

The tested controller was situated inside of a simulated robot that was 60x60 units (pixels) large, had an octagonal shape and two wheels on its left- and right-hand side. All tests were performed in two-dimensional arenas 700x650 units large where obstacles, food and resources were placed. One time step of the simulation update loop was executed each 1/60 seconds. The robot movement was simulated as continuous (Bourg and Seeman, 2004) where the centre of the robot body was moved by a floating-point distance from range $\langle -2; 2 \rangle$ every time step, allowing the robot to move back- wards or forwards. The distance

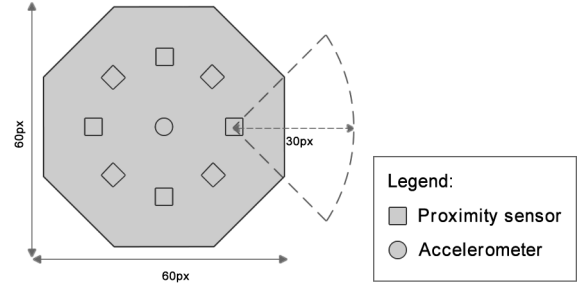


Figure 2: Position of proximity sensors and the accelerometer within the robot body.

moved was calculated based on positive or negative forces applied to the wheels. The effect of size and distance of the wheels was not simulated.

The controller consisted of several sensors, each of them having a corresponding neuron that handled the sensory data. The initial controller could configure itself to move randomly and avoid obstacles. It consisted of eight proximity sensor-neuron couples and one accelerometer - a random-walk neuron couple (Figure 1, highlighted part in the middle without hormonal modulation). Each neuron had a reconfigurable connection to each of the two motors. Both sensory and motor values were from range $\langle -1; 1 \rangle$, and a noise of uniformly distributed random values between $\langle -0.1; 0.1 \rangle$ was added to them.

The proximity sensors were spread around the robot's body (Figure 2) and their outputs (closeness of objects and self-movement for proximity and accelerometer sensors, respectively) were from the range $\langle 0; 1 \rangle$. Signals from the sensors caused corresponding continuous-time self-recurrent neurons n to change their activation level, AC , as shown in Equations (1) and (2):

$$AC(t+1)_n = AC(t)_n + \frac{1}{AC\theta_n}(-AC(t)_n + \lambda * compositeSigmoid(INPUT(t)_n, \beta, \alpha)) \quad (1)$$

$$\begin{aligned}
& \text{compositeSigmoid}(x, \beta, \alpha) \\
&= \begin{pmatrix} -1, & x = -1 \\ -\frac{1}{1+e^{7+\beta(x+\alpha)}}, & x > -1 \wedge x < 0 \\ 0, & x = 0 \\ \frac{1}{1+e^{7-\beta(x-\alpha)}}, & x > 0 \wedge x < 1 \\ 1 & x = 1 \end{pmatrix} \quad (2)
\end{aligned}$$

The parameters $AC\theta$ and λ were neuron-type specific and generally governed the sensitivity of activation in relation to a cell's input. The used composite-Sigmoid function consisted of two sigmoids inverted along the y-axis and its output value was either from range $\langle 0; 1 \rangle$ or $\langle -1; 0 \rangle$, depending on what range the input came from (Figure 3). This function could therefore modulate the input like a standard sigmoid function used in CTRNNs (Beer, 1995), amplifying extreme input values -1 , 0 and 1 .

Output O of a neuron was based on the activation level and a separate noisy connection signal S_{conn} was sent to each of the two motors in parallel with signals from other neurons:

$$O_n = \text{compositeSigmoid}(AC_n, 13, 0) \quad (3)$$

$$\begin{aligned}
S_{conn} &= (O_n * W_{conn}) + (\text{noise} * 2 * \\
&\text{rand}(0, 1) - \text{noise}); \quad \text{noise} = 0.1 \quad (4)
\end{aligned}$$

In Ashby's terms, the input-activation-output sequence provided the controller with a primary sensorimotor loop. A slower secondary loop that allowed for cell homeostasis maintenance was created by adding cell variable 'irritation' IR that observed activation and caused small random weight changes of the cell connections when its value was higher than a cell-type-specific irritation threshold $IR\theta$. Irritation (Equation (5)) started increasing rapidly when a cell's activation was out of its viability zone for a certain amount of time, that is, when the cell was not causing a behaviour that was required of it. For

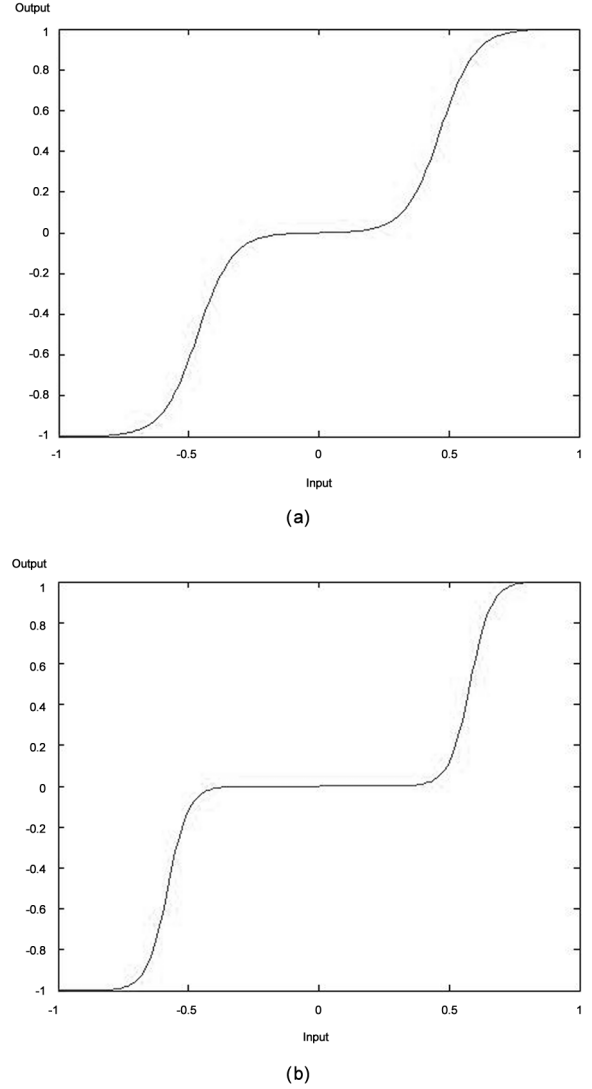


Figure 3: The compositeSigmoid function (Equation (2)) with (a) $\beta = 15$, $\alpha = 0$. (b) $\beta = 25$, $\alpha = 0.3$.

example, activation of a proximity neuron remained high while there was a high input from a proximity sensor, that is, when an object was not avoided:

$$\begin{aligned}
IR(t+1)_n &= IR(t)_n + \frac{1}{IR\theta_n} * \\
&(\text{brokenLine}(AC\vartheta_n, |AC(t)_n|)) \quad (5)
\end{aligned}$$

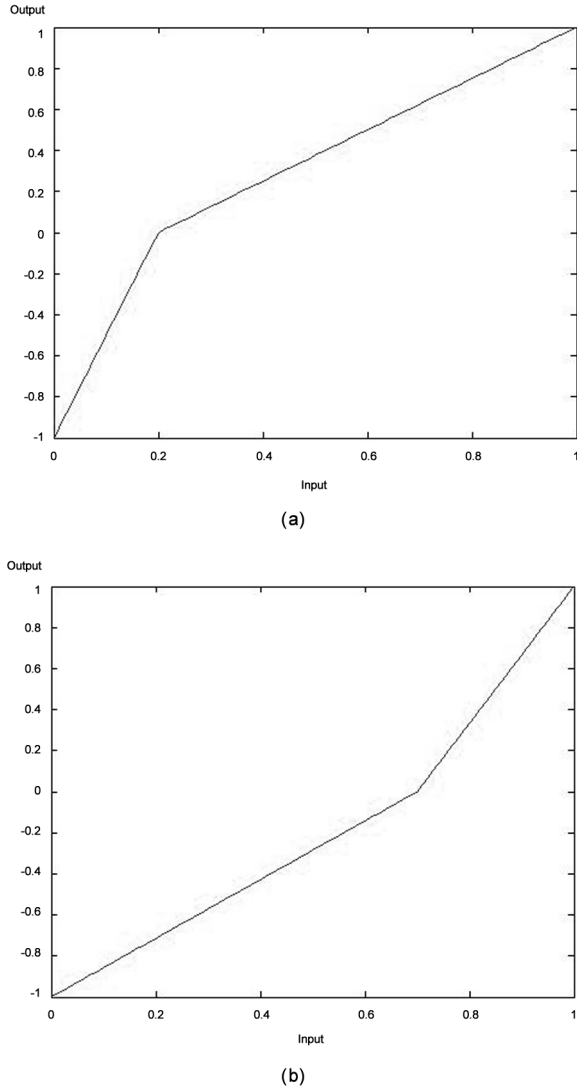


Figure 4: The brokenLine function (Equation (6)) with (a) $\vartheta = 0.2$ (b) $\vartheta = 0.7$.

$$brokenLine(\vartheta, x) = \begin{cases} \frac{x}{\vartheta} - 1, & x < \vartheta \\ \frac{x - \vartheta}{1 - \vartheta}, & x \geq \vartheta \end{cases} \quad (6)$$

Note that only the amount by which activation deviated from a neuron resting state equivalent to

$AC = 0$ was important for calculation of irritation, while the sign of activation was irrelevant. The speed of irritation change was controlled by the brokenLine function (Equation (6)), more specifically by a neuron-type-based activation threshold $AC\vartheta$. BrokenLine was a linear function that returned output from range $\langle 0; 1 \rangle$ but had two different gradients based on the input value (Figure 4). The $AC\vartheta$ parameter thus defined ranges of activation that caused irritation to change at different speeds, specifying when a neuron was more sensitive towards its homeostasis.

At the point of reconnection, the irritation was reset to zero and a new weight was chosen separately for each connection using random Gaussian function with the mean in the current connection weight and a variance VAR based on a maximum pre-defined connection change variance $CCVAR$ of a cell and the cell's current activation (Equation (7)). The probability of connection change (*probChange*) depended on the current connection age with respect to $maxConnAge \approx 167$ seconds and was at least $MINPROB = 0.2$ (Equation (8)). The connection age was reset to 0 when its weight changed. The value of $maxConnAge$ represented 10,000 update loops and was chosen from a range of possible values found experimentally that allowed for both long-term stability and effective adaptation. The connection weights were initialised at random and kept within the range $\langle -1; 1 \rangle$:

$$VAR = CCVAR_n * AC_n \quad (7)$$

$$probChange = (1 - MINPROB) * \left(1 - \frac{connAge}{maxConnAge}\right) + MINPROB \quad (8)$$

Movement of the robot was generally governed by the frequently reconfiguring random-walk neuron $RW-n$ ($\lambda_{RW-n} = 0.7$, $AC\theta_{RW-n} = 20$, $AC\vartheta_{RW-n} = 0.65$, $IR\theta_{RW-n} = 40$, $IR\vartheta_{RW-n} = 0.2$, $CCVAR_{RW-n} = 1.5$) that tried to maintain the robot's speed based on the input from the accelerometer. The proximity neurons inhibited $RW-n$ when the proximity sensors detected objects nearby, interrupting random walk and triggering obstacle avoid-

ance. The resulting behaviour and adaptation obtained with this controller are presented in Section 4.1.

3.2. Extended behaviours and hormonal action selection

An extension of the controller involved adding of carried-weight and energy sensors and their corresponding glands in order to allow for recharging and mineral collection. Furthermore, base-, mineral- and food-memory neurons, carried-weight and appetite glands, as well as a navigation neuron, were added (Figure 1). The memory neurons maintained relative vectors towards items of interest and they were connected to the navigation neuron N-n that affected motors and inhibited the random-walk neuron. The outputs of memory neurons were from range $\langle -1; 1 \rangle$ and expressed magnitude change of the remembered vectors, producing high positive output when distance to the object significantly increased and high negative input when it significantly decreased. The activation of N-n was thus smaller when it correctly navigated towards an item, keeping its irritation low and preventing the neuron's connections to motors from changing. This mechanism caused the robot to move towards the item but change its course quickly when it started drifting away. N-n considered only a signal from a memory neuron with the highest utility each update loop, giving the food-memory neuron precedence in case of a tie (Moioli et al., 2009).

Utility U of a memory neuron was carried along its output in a two-dimensional signal to N-n. Its value was mostly based on a corresponding hormonal concentration C_h , but it was also high when an object of interest was very close to the robot, giving the robot a slightly opportunistic behaviour (Equation (9)):

$$U_n = \text{MAX}(0.9 * f(C_h), \text{distanceComponent}) \quad (9)$$

The carried-weight gland CW-g produced high amounts of carried-weight hormone CW-h when there was little or no weight carried, causing high utility of the mineral-memory neuron MM-n. The gland ceased the hormonal production when the

robot carried a certain amount of weight, which decreased utility of MM-n and increased utility of base-memory neuron BM-n. Similarly, the appetite gland produced appetite hormone AP-h when the level of energy started getting low and caused the utility of the food-memory neuron to become high and take over navigation. The resulting behaviour during a gathering recharging cycle is presented in Section 4.2.

All hormones were produced by a hormone repository based on signals S from glands g , a hormone decay rate ρ_h and a hormone growth multiplier γ_h , as shown in Equation (10). The glands were ultrastable and had their own activation and irritation parameters. However, in contrast with neurons, they were more effective by adapting the activation thresholds so that governed sensitivity of activation in respect to input (Equation (1) and (2)), rather than the connection weights to the repository:

$$C_h = C_h + \gamma_h * \rho_h * |S_{conn}| - \rho_h \quad (10)$$

Apart from affecting the utility of memory neurons, the carried-weight and appetite hormones inverted the signs of inputs from the proximity sensors to proximity neurons when objects such as food, minerals or base were detected, effectively making the robot steer towards the objects rather than away from them when obstacle avoidance would normally take place. A similar type of hormonal modulation where hormones affected inputs of neurons was implemented by Neal and Timmis (2003), Timmis et al. (2009) and Hamann et al. (2010).

4. Initial results

4.1. Behaviour and adaptation of the initial randomly walking robot

Figures 5 and 6 show an exemplary run with the simple obstacle-avoiding controller described in Section 3.1. The run consisted of Trial 0, where the initial configuration of the proximity neurons occurred, and Trial 1, where the resultant controller from Trial 0 was used in the same arena. A significant amount

of obstacle hitting was noted during the first two thirds of Trial 0 (Figure 5(a)), accompanied by frequent weight changes of proximity neuron connections (Figure 6(a)). The reconnection stabilised between the 110th and 120th seconds into a configuration where the collective inputs from the neurons to motors allowed the robot to steer away from obstacles, keeping the neural activations and thus connection reconfigurations to minimum. Two minor reconnections occurred at around the 135th and 155th seconds, caused by an obstacle hit in the first case and a close encounter with an obstacle in the second case. There was only one obstacle hit in the subsequent Trial 1, which caused minor reconnections on several neurons at around the 10th second. Small reconnections also occurred later in Trial 1 due to longer near encounters with obstacles, although the obstacles were successfully avoided.

It is important to point out that the controller never completely stops reconfiguring. If near an obstacle for long enough, a cell's irritation reaches a high level and causes a connection weight to change, although the probability of that happening decreases for well-working connections. Furthermore, small and isolated alterations usually do not impact the overall obstacle avoidance performance and obstacle avoidance can rarely fail after the initial reconfiguration. Figure 7 shows average obstacle hitting times based on 50 runs, sampled every 20 seconds during Trials 0 and 1. Encounters with obstacles gradually decreased during Trial 0, from initial 20% to approximately 5%. The random reconfigurations of proximity neurons had on average the most significant impact between the 40th and 60th seconds when the obstacle hitting time dropped from 19.15% to 11.27%. A slightly less significant decrease was noted between the 100th and 120th seconds, from 10.85% to 8.02%, after which the performance slowly improved by no more than 2.5% during one 20-second window. The average obstacle hitting time during Trial 1 was 3.19% and did not exceed 5.62%.

Throughout the 50 runs, a few controllers were found that learned obstacle avoidance poorly relative to the other controllers. Figure 8 shows the dependence of the amount of obstacle hitting during Trial 1 (y-axis) on Trial 0 obstacle hitting (x-axis). A high

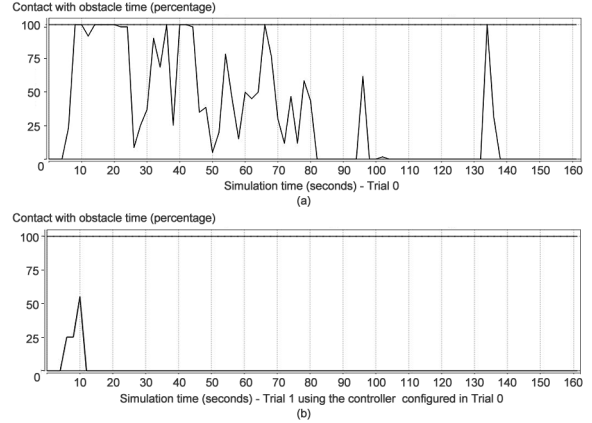


Figure 5: Example of obstacle avoidance learning. Percentage of time during which the robot was hitting an obstacle was sampled for each second during (a) the training Trial 0 and (b) a subsequent Trial 1 with a single controller.

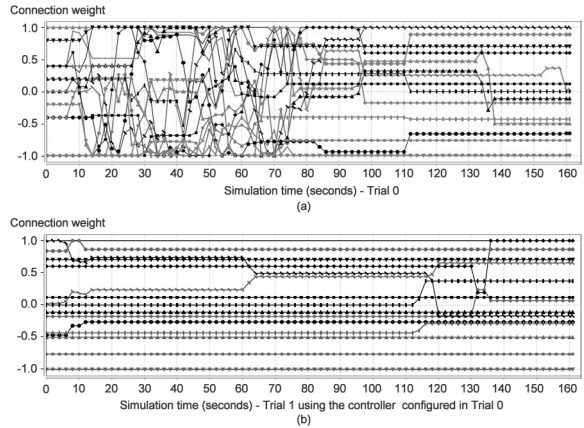


Figure 6: Changes in connection weights during obstacle avoidance learning presented in Figure 5. Weights of connections from the eight proximity sensors to two motors were sampled for each second during (a) Trial 0 and (b) subsequent Trial 1.

amount of obstacle hitting during Trial 1 correlated with low obstacle hitting times in Trial 0. This was an expected result, since obstacle avoidance could not be properly learned when a small number of obstacles were encountered in the initial trial.

In addition to initial self-configuration, the controller was also able to adapt to perturbations includ-

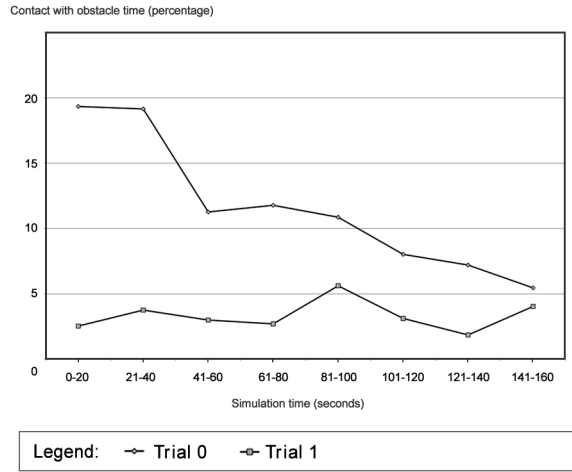


Figure 7: Obstacle hitting throughout initial configuration tests. The shown values represent average percentage of time during which the robot was hitting an obstacle sampled between two data points, each 20 seconds apart.

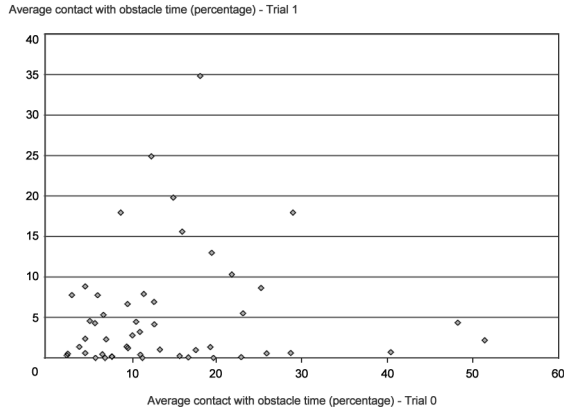


Figure 8: Relationship between obstacle hitting time in Trials 0 and 1.

ing higher sensory noise of up to 0.4 and inversion of motor rotation directions that caused the robot to steer in opposite ways than it originally configured itself for. The robot was initially configured for 120 seconds under normal conditions (Trial 0), after which a perturbation was introduced for 360 seconds

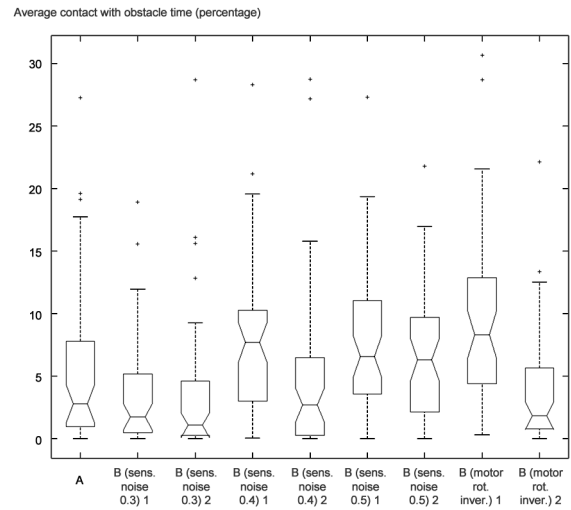


Figure 9: Analysis of variance of initial adaptation results (MatlabA). Average obstacle hitting time of Controller A that did not experience perturbations is compared to Controllers B that did experience them in Trials 1 and 2. The used annotation pattern is [controller][perturbation] [trial no.]. Medians are shown as lines in the middle of each box, with the 25th and 75th percentiles represented by tops and bottoms of each box respectively. The whiskers represent interquartile ranges of the samples and observations beyond the whisker length are marked as outliers (MatlabB).

(Trial 1) and then again for 360 seconds (Trial 2). The robot's position and the cell activations and irritations were reset at the beginning of each trial. The duration of Trial 0 was chosen based on the fact that no significant improvements in obstacle avoidance were noted after the 120th second with an unperturbed controller. The performances from Trials 1 and 2 were compared with the performance of Controller A, which did not experience perturbations and ran for 360 seconds after Trial 0. Controller A achieved a median obstacle hitting time of 2.79%.

Figure 9 reveals that the obstacle avoidance behaviour was robust and even slightly improved for sensory noise of 0.3 (1.75% median hitting time in Trial 1 and 1.09% in Trial 2). The hitting time was significantly higher in Trial 1 with both sensory noise of 0.4 (median hitting time 7.72%) and 0.5 (6.57%) and only improved in Trial 2 with noise

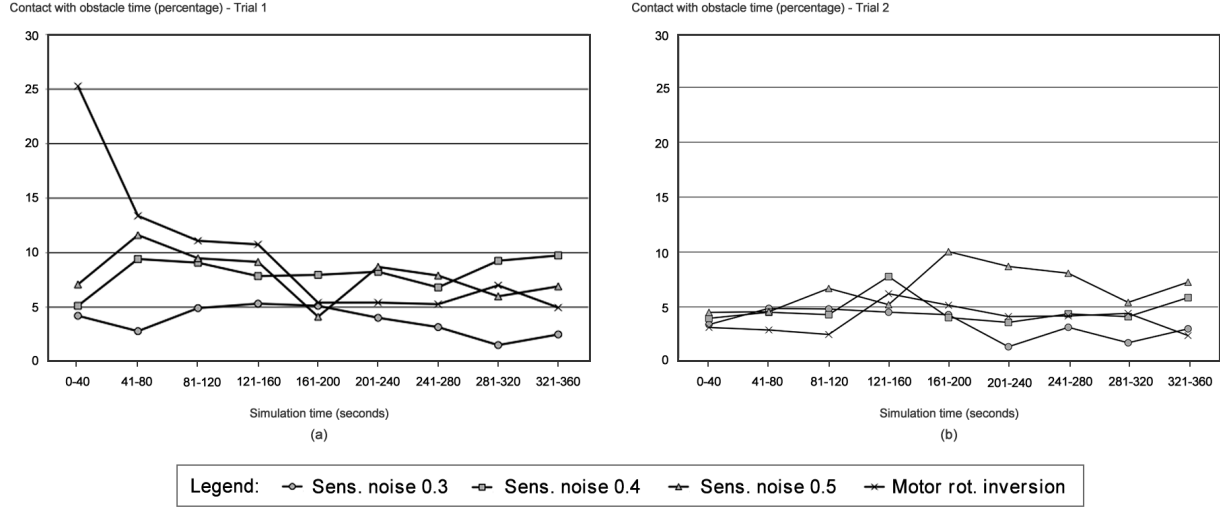


Figure 10: Obstacle hitting throughout initial adaptation tests. The shown values represent average percentage of time during which the robot was hitting an obstacle sampled between two data points, each 40 seconds apart. Results from tests with Controller B during Trials (a) 1 and (b) 2 are shown.

of 0.4 (2.65%). The performance during Trial 2 with noise of 0.5 did not improve (median hitting time 6.32%), as effective reconnection was not possible due to constantly high activation of the proximity neurons. These results are also reflected in Figure 10, which shows the average obstacle hitting time throughout Trials 1 (a) and 2 (b) sampled for each 40 seconds. It is interesting to note that when sensory noise of 0.4 was used, the performance did not improve throughout Trial 1, but was on average significantly better in Trial 2. It seems that there were frequent encounters with obstacles between the 280th and 360th seconds in Trial 1 and between the 120th and 160th seconds in Trial 2 that probably caused final correct configurations to be found. Also, note the instability of performance when sensory noise of 0.5 occurred. A significant performance improvement was recorded on average between the 160th and 200th seconds of Trial 1, but the configurations often broke down and the performance degraded again.

The impact of motor rotation direction change was immediate and high at the beginning of Trial 1 (Figure 9, median hitting time 8.32%). Adaptation was most significant during the first 200 seconds of

Trial 1 (Figure 10) and returned the controller to its normal obstacle avoidance performance in Trial 2, slightly outperforming Controller A (median hitting time 1.83%).

Similarly to the results from the initial controller configuration, the success of adaptation depended on how often obstacles were encountered during Trial 1, giving controllers that experienced more hits a better chance to adapt.

4.2. Gathering-recharging cycle of the hormone-controlled robot

The following paragraph and Figure 11 describe normal behaviour of a fully implemented hormone-modulated UNEC as described in Section 3.2. The robot started off near a base carrying zero weight and having full energy. Since there initially was no weight carried, the amount of carried-weight hormone CW-h was low, causing high utility of the mineral-memory neuron MM-n. However, because there was no mineral in the memory yet, the output of MM-n was still 0, RW-n was not suppressed and the robot moved randomly (Figure 11, trajectories between Points 1 and 2). When a mineral was detected, the output

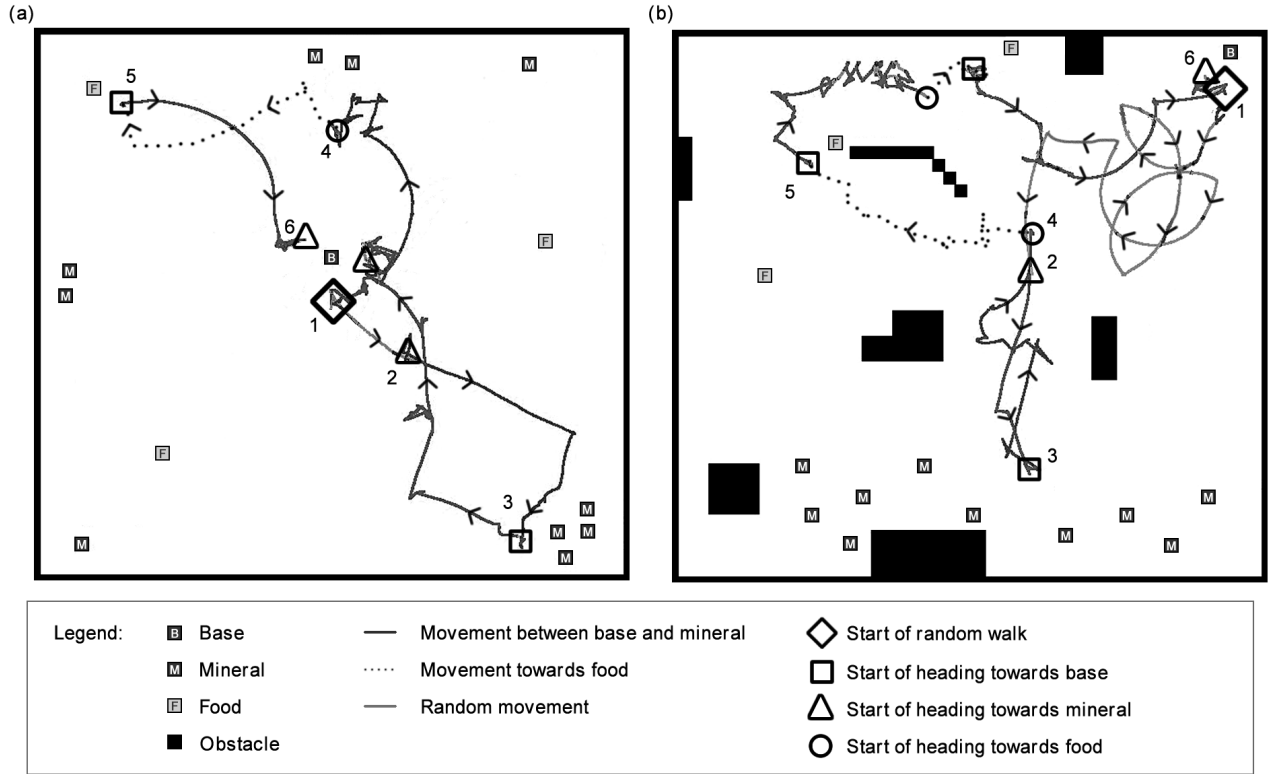


Figure 11: Examples of the robot performing gathering and recharging. Trajectories from the first mineral gathering - feeding cycles only are shown.

transferred from MM-n to N-n became high, causing N-n to suppress RW-n and the robot to steer towards the mineral (Figure 11, Points 2). The level of CW-h became high when there was a certain amount of weight carried (Figure 11, Points 3). This hormonal change caused utility of the base-memory neuron BM-n to become the highest, thus making the N-n start reconnecting in order to minimise the distance between the robot and the base. The location of the base was always known as it was the starting position of the robot. The mineral collection loop kept repeating while there were minerals in the memory (Figure 11(a), unlabelled triangle point near the base). Otherwise, the robot resumed random walk as RW-n was not suppressed anymore, which continued until another mineral source was found.

The appetite gland AP-g started producing the ap-

petite hormone AP-h when the energy level started getting low, causing high utility of a food-memory neuron FM-n. Similar to the case with minerals, behaviour of the N-n based on the input from FM-n caused the robot to minimise its distance towards food or to perform random walk, depending on whether a food item was in the memory. This meant that the robot started searching for food rather than continuing to collect minerals when the energy was low but no food was found so far. Food following or search could be triggered at any time during the gathering cycle (Figure 11, Points 4) and the gathering cycle resumed once the robot's energy was high again (Figure 11, Points 5). The whole process repeated once the base was reached and minerals were unloaded (Figure 11, Points 6).

5. Adaptation evaluation methods

The navigation abilities of the robot were evaluated on their own, as well as alongside mineral gathering and recharging, by comparing the performance of Controller A, which did not experience perturbations, and Controller B, which did. The individually introduced perturbations included increased sensory noise, increased motor noise and inversion of motor rotation direction.

Trial 0, which lasted 120 seconds under normal conditions with sensory and motor noise of 0.1, was executed in an arena with no resources or base in order to allow for initial obstacle avoidance configuration. Subsequent Trials 1 and 2 with Controller B were performed for 360 seconds with the individual perturbations. Controller A ran for 360 seconds after Trial 0 under normal conditions.

The following hypotheses were tested:

1. The performance of Controller B in both trials will not significantly differ from Trial 1 of Controller A under perturbations up to a certain value, that is, Controller B will stay robust to small perturbations.
2. If degradation of behaviour of Controller B in Trial 1 occurs, adaptation will cause the performance in Trial 2 to be better.
3. in the case of body changes, the performance of Controller B in Trial 2 will be similar to that of Controller A in Trial 1, as the environmental conditions for executing a tested behaviour will not change and adaptation will be able to return the controller to its normal working state.

The controller performances were compared using balanced one-way analysis of variance (MatlabA) and Dunn Sidak adjusted comparison tests (MatlabC).

The testing was performed in three arenas, labelled GAT1, GAT2 and BIG, containing three food items each holding one floating-point unit of energy, one base and eleven mineral items each containing 2.0 units of minerals. Food items started recharging

after 50 seconds of being depleted and their recharging took approximately 13 seconds. While resources were placed in between obstacles in the first two arenas, there were no obstacles in the BIG arena, apart from its surrounding walls. The following three test cases consisting of 50 runs each were executed in order to gradually test the performance of:

- Test Case 1: obstacle avoidance
- Test Case 2: obstacle avoidance and mineral gathering
- Test Case 3: obstacle avoidance, mineral gathering and recharging

The robot's energy level was fixed to 1.0 in Test Cases 1 and 2. Because random configuration of the appetite gland could potentially make the robot deplete its energy before the gland could function properly, an emergency energy supply was used for both Controllers A and B in Test Case 3 during a separate configuration Trial 0b that lasted 360 seconds and preceded Trial 1. The energy level decreased by approximately 0.9% per second, meaning that the energy was completely depleted in approximately 110 seconds, allowing for short contingences in cases when food needed to be searched for by random walk. The robot's speed linearly decreased with the amount of carried minerals, which introduced a penalty to robots that tried to cheat by not returning to the base to unload the minerals at reasonable intervals.

6. Adaptation results

6.1. Sensory perturbations

The robot behaviour was evaluated with sensory noise of 0.3 and 0.4 that simulated worse visibility, for example rainy or foggy conditions. The navigation abilities of Controller B in Test Case 1 were not affected by noise of 0.3 and the controller was able to adapt to sensory noise of 0.4 in all three testing arenas (data not shown), mirroring results of the initial controller without glands and confirming Hypotheses 1 and 2. A higher input noise of 0.5 caused high

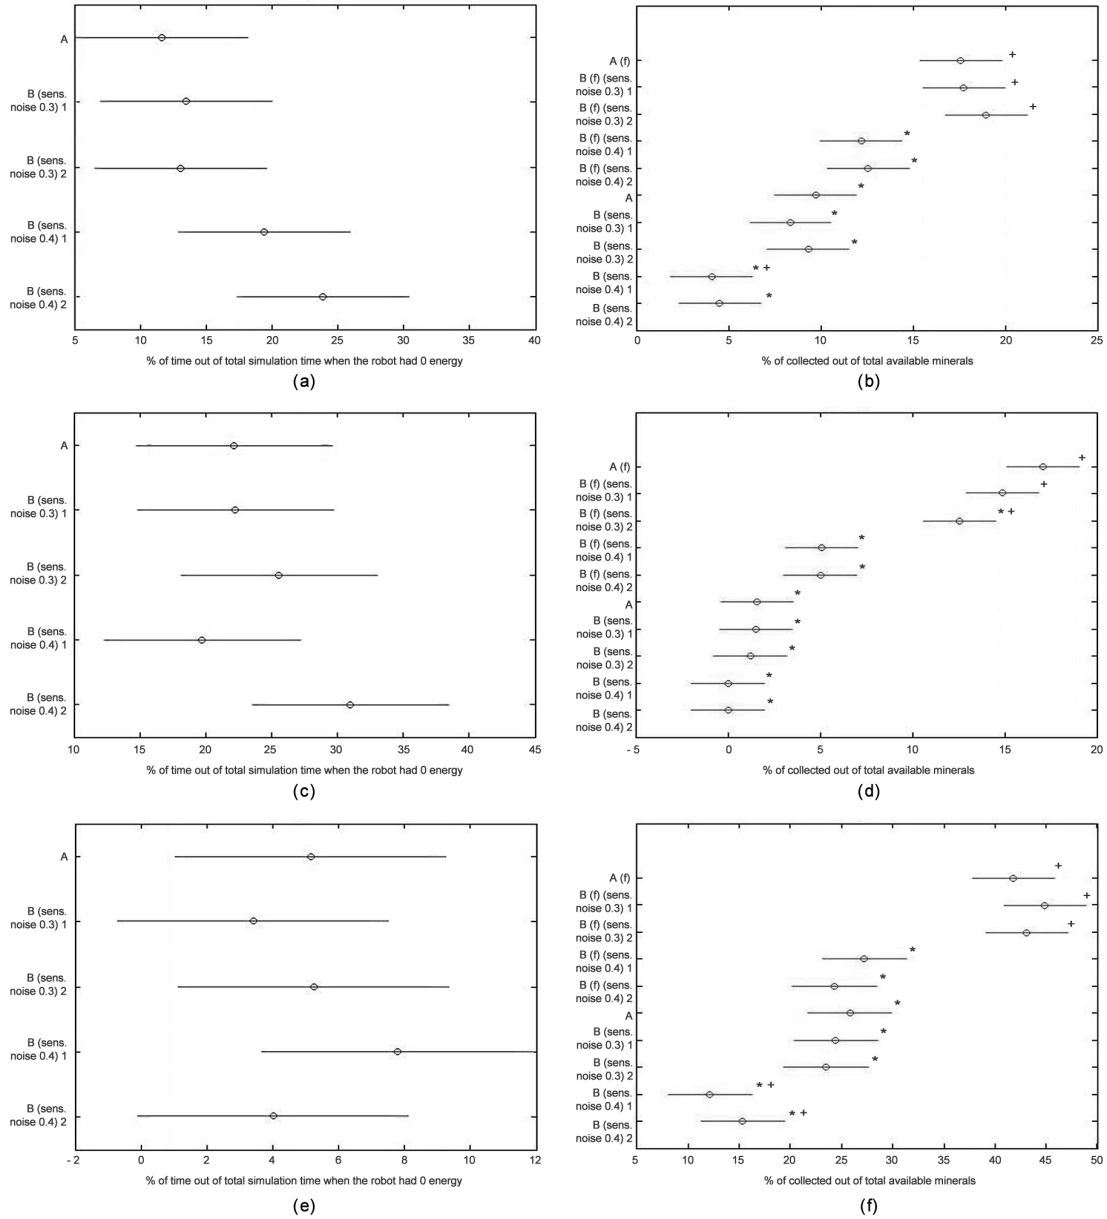


Figure 12: Dunn Sidak adjusted comparison of gathering and recharging controller performances during sensory perturbations. The used annotation pattern is [controller]([perturbation]) [trial no.]. The controllers from Test case 2 have (f) appended to their names. Group means are represented by circles and group intervals are shown as whiskers (MatlabC). Group means significantly different ($p < 0.05$) from the normal Controller A(f) of Test Case 2 are marked with *. Group means significantly different from the normal Controller A of Test Case 3 are marked with +. (a) Arena GAT1, 0-energy time. (b) GAT1, gathered minerals. (c) GAT2, 0-energy time. (d) GAT2, gathered minerals. (e) BIG, 0-energy time. (f) BIG, gathered minerals.

activation of the proximity neurons and spontaneous reconnections that prevented adaptation.

Results from Test Cases 2 and 3 are presented in Figure 12. Gathering was mostly not affected by sensory noise of 0.3 in Test Case 2, confirming Hypothesis 1 (median amount of gathered minerals during Trial 1 17.72% compared to Controller A s 17.58% in GAT1 and 44.91% compared to 41.73% in BIG). However, this did not occur in GAT2, where gathering performance dropped in Trial 2 (12.51% compared to Controller A s 17.03%). Gathering also degraded with sensory noise of 0.4 (median gathered obstacles during Trial 1: 12.20% in GAT1, 5.05% in GAT2 and 27.20% in BIG) and was not improved by adaptation (Trial 2 results: 12.56% in GAT1, 4.99% in GAT2 and 24.11% in BIG).

Similar results were achieved during Test Case 3. Hypothesis 1 was confirmed in all arenas during tests with sensory noise of 0.3, where gathering remained robust, although the amount of gathered resources was smaller due to the robot s need to recharge (Trial 1 results: 8.38% compared to Controller A s 9.71% in GAT1, 1.54% compared to 1.58% in GAT2, 24.39% compared to 25.81% in BIG). Adaptation with sensory noise of 0.4 was not successful in any arena, although the amount of gathered minerals slightly improved in BIG (12.11% in Trial 1, 15.34% in Trial 2). Recharging in Test Case 3 was robust to sensory noise of 0.3 (Trial 1 0-energy time 13.60% compared to Controller A s 11.60% in GAT1, 22.3% for both controllers in GAT2 and 3.42% compared to 5.08% in BIG) and broke down with sensory noise of 0.4 in arenas GAT1 (19.41% in Trial 1 and 23.83% in Trial 2) and GAT2 (19.62% in Trial 1 and 30.89% in Trial 2), although the results were not statistically significant. Hypothesis 2 was only confirmed in the BIG arena with noise of 0.4, where the recharging was worse in Trial 1 (0-energy time 7.66%) but improved in Trial 2 (3.98%).

6.2. Motor perturbations

Higher motor noise was introduced in order to test how the controller could adapt to changed environmental conditions such as rougher terrain. Furthermore, a change in the robot body was simulated by inverting motor rotation directions, which caused the

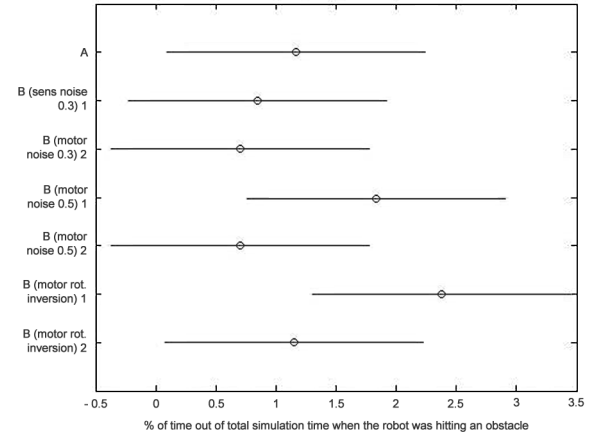


Figure 13: Obstacle hitting time in the BIG arena during Test Case 1 with motor perturbations. The used annotation pattern for the shown groups is [controller]([perturbation]) [trial no.]. Group means are represented by circles and group intervals are shown as whiskers (MatlabC). None of the groups has a significantly different mean than the original Controller A.

robot to steer in opposite ways than it initially configured itself for. Obstacle avoidance of Controller B was robust during tests with motor noise of 0.3, confirming Hypothesis 1. The performance with other perturbations was worse in Trial 1, but the controller adapted successfully and its behaviour returned to normal in Trial 2, confirming Hypotheses 2 and 3. Figure 13 shows the results from arena BIG, where the differences between Trials 1 and 2 were the most significant (motor noise of 0.5: median obstacle hitting time 1.85% in Trial 1 and 0.67% in Trial 2, motor rotation inversion: 2.36% in Trial 1 and 1.08% in Trial 2).

Results from Test Cases 2 and 3 are shown in Figure 14. Similar to the tests with sensory noise, the gathering performance of Controller B remained robust with motor noise of 0.3 during Test Case 2 (median amount of gathered minerals in Trial 1: 18.18% compared to Controller A s 18.28% in GAT1, 17.62% compared to 17.13% in GAT2 and 43.87% compared to 41.78% in BIG). This result was also achieved when the motor rotation direction was changed in

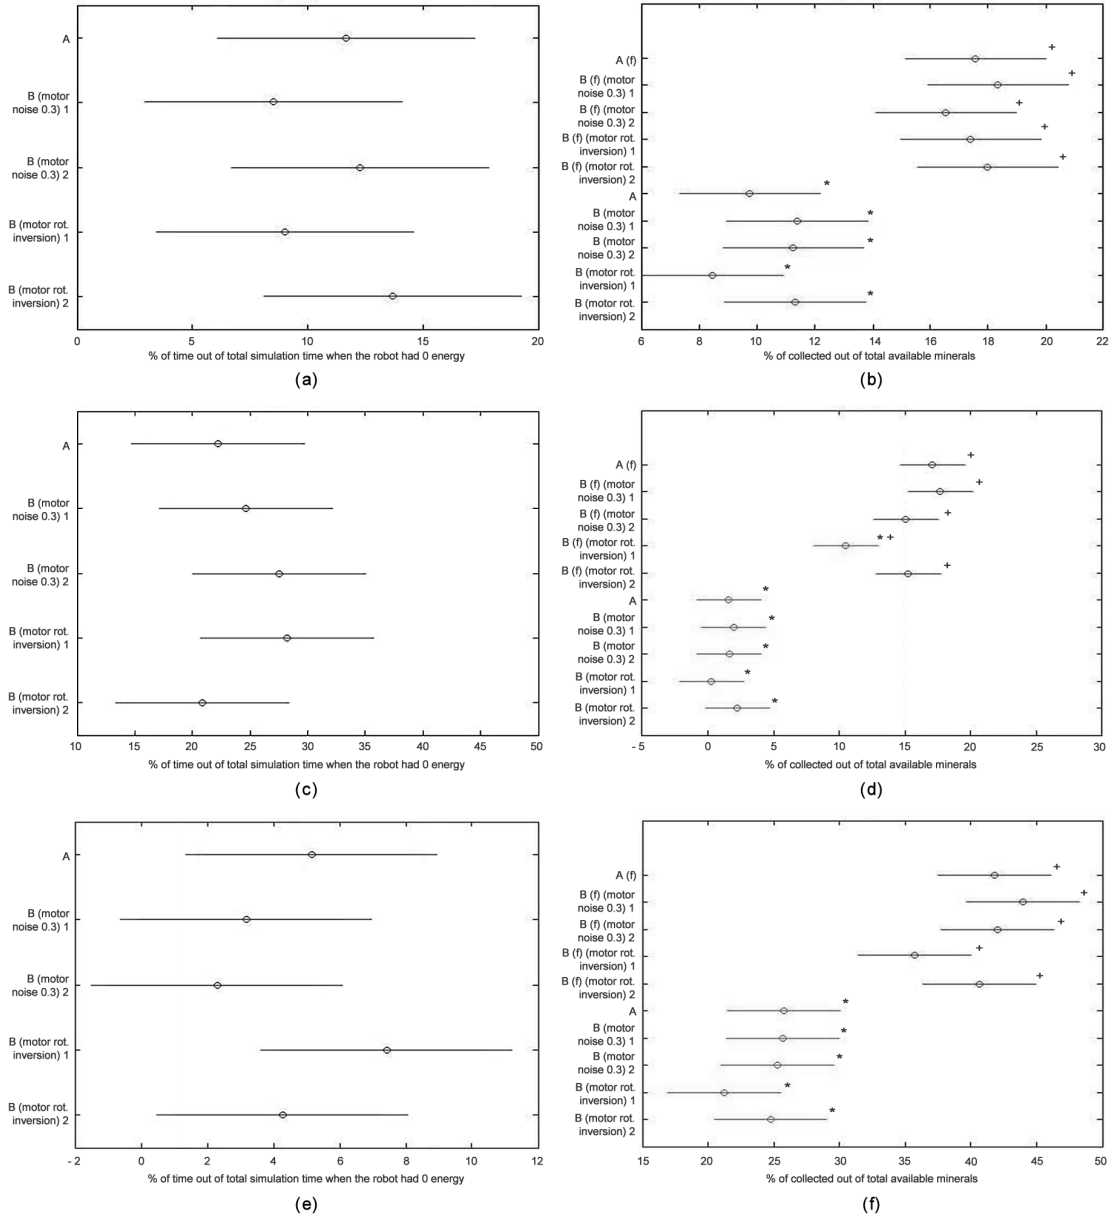


Figure 14: Dunn Sidak adjusted comparison of gathering and recharging controller performances during sensory perturbations. The used annotation pattern is [controller]([perturbation]) [trial no.]. The controllers from Test case 2 have (f) appended to their names. Group means are represented by circles and group intervals are shown as whiskers (MatlabC). Group means significantly different ($p < 0.05$) from the normal Controller A(f) of Test Case 2 are marked with *. Group means significantly different from the normal Controller A of Test Case 3 are marked with +. (a) Arena GAT1, 0-energy time. (b) GAT1, gathered minerals. (c) GAT2, 0-energy time. (d) GAT2, gathered minerals. (e) BIG, 0-energy time. (f) BIG, gathered minerals.

GAT1 (17.40%). Hypotheses 2 and 3 were confirmed when the motor rotation inversion was introduced in GAT2 and BIG, where the amount of gathered minerals dropped in Trial 1 (10.42% in GAT 2 and 35.45% in BIG) but increased in Trial 2 (15.11% in GAT 2 and 40.51% in BIG), getting closer to the performance of Controller A.

Similar results were found when looking at gathering during Test Case 3, where motor noise of 0.3 did not affect the performance significantly (Trial 1 results: 11.41% compared to 9.69% in GAT1, 1.91% compared to 1.58% in GAT2 and 25.78% compared to 21.18% in BIG) The controller's performance dropped in Trial 1 when the motor rotation inversion was introduced (8.40% in GAT1, 0.25% in GAT2 and 21.11% in BIG), but the performance matched or was better than Controller A's in Trial 2 (11.20% in GAT1, 2.26% in GAT2 and 24.72% in BIG).

Although the gathering performance results obtained during motor perturbation tests were not as statistically different as with sensory perturbations, ANOVA evaluation (results not shown) revealed that the adapted gathering performance was generally similar or better and more stable during Trial 2 compared to Trial 1, implying that adaptation did play its role.

Finally, recharging did not seem to be significantly affected when motor noise was increased, although the results varied by arena, and the performance was not as consistent as the gathering performance. Hypotheses 2 and 3 were confirmed when motor rotation was inverted in GAT2 (0-energy time 28.20% during Trial 1 compared to Controller A's 22.21% and 20.08% during Trial 2) and BIG (7.38% during Trial 1 compared to 5.12% and 4.26% during Trial 2). The results from GAT1 were not statistically significant.

6.3. The impact of random walk

It was apparent especially during tests with higher sensory noise that the effectiveness of navigation based on randomly reconfiguring cells was relatively low, causing multiple encounters with the same object when the agent tried to steer away from it and then resume its original path. This negatively affected feeding and gathering in arenas GAT1 and

GAT2, where obstacles were often in the way between the robot and its item of interest.

Furthermore, the search efficiency of random walk was also limited. Even in small testing arenas, food or minerals were sometimes not found and there was no search behaviour to execute since there was no memory of visited places. Adding more sophisticated motor control along with visited landmarks memory neurons could potentially support path planning and help with these issues.

7. Discussion

Work of a simulated UNEC where cells reconfigure in order to maintain their homeostasis was sketched and tested in terms of obstacle avoidance, alteration between recharging and gathering and adaptability to environmental (sensory and motor noise) and body (motor rotation inversion) perturbations. It was hypothesised that:

1. The controller will stay robust during minor perturbations.
2. If degradation of behaviour due to a perturbation occurs, adaptation will cause the performance to improve after a certain time.
3. In the case of body changes, the adapted performance will be similar to the original performance without a perturbation, as the general conditions for executing a tested behaviour will not change and adaptation will be able to return the controller to its normal working state.

The results presented in Sections 4 and 6 show that the controller was mostly capable of remaining robust to minor environmental and body perturbations, confirming Hypothesis 1. Adaptation to high sensory noise was only possible for obstacle avoidance, although terrain roughness and changes in the robot's motors were handled more appropriately and hormone-modulated behaviours were adapted as well. Hypothesis 2 was therefore mainly satisfied for navigation. Hypothesis 3 expressed the intuition that it is not only the adaptation mechanism

itself but also the range of possibilities afforded by a changed environment that affect a robot's performance. This was confirmed, as the adapted behaviour of a controller that experienced motor rotation inversion was either similar to or more effective than behaviour of the original controller that did not experience any perturbations.

The fact that the same mechanism was applied to both the controller configuration and adaptation and to both action execution (obstacle avoidance, gathering, etc.) and action selection (working of glands) indicates that the properties of the presented ultra-stable cells are not limited to a specific cell type or functionality and are applicable to a larger variety of problems than those targeted by the particular simulations developed for this paper. The controller components utilising ultrastability had different roles and operated on different time scales, while their main common property was that they interacted with each other through their environment, that is, the internal and external state of the robot. Rather than using action selection to achieve ultrastability on the robot level, ultrastability was used on the cellular level to configure behaviours including action selection. This ties in more closely with biological neurons (Sánchez-Lasheras et al., 2011).

The nature of the UNEC allowed for the following advantages.

1. Robustness

It was shown in Section 4 that cell connections stabilised at the point when a behaviour started working effectively, that is, when activation of a cell was kept within acceptable bounds. Furthermore, a connection was stable under minor activation perturbations due to a time gap between high activation and reconnection provided by irritation, as well as due to the fact that there was a smaller probability for older connections to reconfigure. The UNEC thus provides a possible solution to the problem of achieving robustness, which is one of the reasons why ultrastability tends not to be widely used in engineering (Manicka and Di Paolo, 2009).

2. Reduced design requirements

Adaptation is random and parallel and there is no predefined correct configuration of the controller because of the fact that cells behaviourally depend on each other. Tolley and Lipson (2011) advocate stochasticity in motion planning and locomotion and argue that it reduces design requirements. It is certainly the case in the UNEC, where the robot behaviour is not dependent on exact cell parameter values and it was often found that values within the same or very similar order of magnitude delivered statistically similar results.

3. Hormonal modulation learned through ultrastability

Most of today's work with hormonal modulation (French and namero, 2005; Vargas et al., 2005; Liang et al., 2010; Moiola et al., 2008, 2009; Sauzé and Neal, 2010; Schmickl et al., 2011; Timmis et al., 2009) involves (simulated) robots with pre-learned behaviours that are switched on or off depending on levels of hormones but are not adaptive themselves. Also, the way hormonal modulation works is often evolved and predefined for life-time of a robot (Hamann et al., 2010; Moiola et al., 2009), although some authors implemented learning algorithms for on-line adjustment of hormonal release (Iizuka and Di Paolo, 2008; Santos et al., 2010; Timmis et al., 2009). Walker and Wilson (2011) developed on-line adaptation of hormone-to-task sensitivity based on a robot's effectiveness at a task, showing that heterogeneous robots pre-evolved to be better at different tasks could pick those tasks when given a choice.

Similarly, the hormonal release in the UNEC is adaptable, although task performance is evaluated indirectly through the irritation level of neurons responsible for a particular behaviour. Furthermore, on the contrary to the previous work discussed above, there is not a set of pre-configured behaviours and even the most basic random walk and obstacle avoidance are constantly able to change, giving an UNEC-operated robot the ability to perform hormone-selected behaviours under various environmental conditions. More hormone-related perturbations need to be tested in order to confirm this intuition.

4. On-line adaption without a body model module

There is a relatively new branch of robotics interested in adaptation through body and forward action-reaction modelling advocated, for example, by Bongard et al. (2006), Moriguchi and Lipson (2011) and Hild et al. (2011). The tests with the UNEC indicate the possibility that there is no need for such a model, at least in a sense that there is no need for a specific place where the model would be stored or calculated and that it is possible to approach behavioural adaptation as a property of a system where adjustments to changes are distributed between the system's parts.

The following problems with the UNEC were noted.

1. No time autonomy

The time gap between high activation and reconfiguration is governed by currently hard-coded irritation parameters, preventing the activation from being time autonomous. This problem was apparent especially during adaptation experiments with the appetite gland, which became unstable in cases when food was not available for a prolonged amount of time. It was also problematic to stabilise a proximity neuron when its activation rose too quickly, for example during tests with sensory noise of 0.5. This was observed across different cell parameter ranges. A similar time-dependency problem was found by Timmis et al. (2009). It is possible that there is a more fundamental version of ultrastability that could deal with this issue.

2. Target of cell adaptation

A human designer needs to specify the source of cell activation, that is, what the cell is trying to optimise itself against. Therefore, artificial evolution would probably be needed in order to create more complex controllers.

It is apparent from the above points that apart from relatively poor navigation, which could very probably be improved by adding more complexity,

the main problems of the controller were related to designer-specified constants that governed behaviour of the artificial cells. This fact only underlines the importance of ultrastability as a principle that should be applied to as many parameters as possible in order to minimise human-imposed restrictions and increase autonomy.

The mechanics based on the type of presented ultrastability tie in with the argument of Haselhager (2005) that artificial cells should not act because humans programmed them to do so but because they are trying to do something beneficial for themselves, such as to maintain their survival or survival of the artificial creature. Arguably, robots that are hard-coded to blindly execute pre-programmed behaviours have limited life-like behaviour, no self-purpose and a very limited autonomy. It is also questionable whether learning where body and forward models are calculated in a black-box mind is not equally blind. Even if such learning can make a machine walk autonomously, autonomy and adaptation will most probably be limited to what a program given by the designer can model and what conditions were assumed when it was designed. This became apparent when most of the problems of the UNEC were related to values of designer-specified parameters.

Only fairly recently, neuroscientists started to lean towards the argument that a self-model of a human probably exists as a configuration of the nervous system itself (Metzinger, 2008) and that the feeling of self is its secondary property (Wegner, 2003; Metzinger, 2006, 2008). From this perspective, imposing of self-modelling as a separate function of an artificial body represents a step back towards the Good Old Fashioned AI, when reasoning modules did all sensorimotor calculations *in vacuo*. This approach was proven limited some time ago (Dreyfus, 1979; Brooks, 1991). On the other hand, it is highly probable that a truly space- and time- autonomous robot will plan and feel with its whole body and reconfigure not only its behaviour, but also behavioural planning and execution in order to survive and satisfy its needs. We perhaps need to start looking at approaches to robot design where the overall behaviour itself is not as important as the behaviour of a robot's interacting artificial cells, appreciating the complex-

ity of their interactions as well as the emergence of phenomena these interactions result in.

7.0.1 Acknowledgements

This work was performed during my Masters program at the University of Sussex, UK. I would like to thank Professor Phil Husbands and Dr Luc Berthouze for inspiration and feedback. I would also like to thank Dr Inman Harvey and Dr Andy Philippides for their useful advice during writing of this article.

7.0.2 Funding

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

References

- Matlab a: anova 1. <http://www.mathworks.co.uk/help/stats/anova1.html>. [Online; accessed September 2012].
- Matlab b: box plots. <http://www.mathworks.co.uk/help/stats/boxplot.html>. [Online; accessed September 2012].
- Matlab c: multcompare. <http://www.mathworks.co.uk/help/stats/multcompare.html>. [Online; accessed September 2012].
- Ashby, W. R. (1960). *Design for a Brain: the origin of adaptive behaviour*. London, UK: Chapman and Hall, 2nd edition.
- Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behaviour*, (3):469–509.
- Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, (314):1118–1121.
- Bourg, D. M. and Seeman, G. (2004). *AI for game developers*, pages 16–19. Sebastopol: O’Reilly Media, 1st edition.
- Brooks, R. A. (1991). Intelligence without reason. In *Computers and Thought, IJCAI- 91, April 1991*.
- Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S. W., editors, *From Animals to Animats, Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, pages 440–449. Cambridge, MA: MIT Press.
- Dreyfus, H. L. (1979). *What computers can’t do: The limits of artificial intelligence*. New York: Harper and Row, revised edition.
- French, R. L. B. and namero, L. C. (2005). Introducing neuromodulation to a braitenberg vehicle. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 4188–4193.
- Hamann, H., Strander, J., Schmickl, T., and Crailsheim, K. (2010). A hormone-based controller for evolutionary modular robotics: from single modules to gait learning. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 244–251. Piscataway, NJ: IEEE Press.
- Haselhager, W. F. G. (2005). Robotics, philosophy and the problems of autonomy. *Pragmatics & Cognition*, (13):515–532.
- Hewson-Hughes, A. K., Hewson-Hughes, V. L., Miller, A. T., Hall, S. R., Simpson, S. J., and Raubenheimer, D. (2011). Geometric analysis of macronutrient selection in the adult domestic cat, *felis catus*. *Journal of Experimental Biology*, (214):1039–1051.
- Hild, M., Kubisch, M., and Höfer, S. (2011). Using quadric-representing neurons (qrens) for real-time learning of an implicit body model. *Robotica* 2011.
- Iizuka, H. and Di Paolo, E. A. (2008). Extended homeostatic adaptation: Improving the link between internal and behavioural stability. In Asada, M., Hallam, J. C., Meyer, J., and Tani, J., editors, *From Animats to Animals 10, the Tenth International Conference on the Simulation of Adaptive Behavior*, pages 1–11. Berlin: Springer.
- Josens, R. B. and Roces, F. (2000). Foraging in the ant *camponotus mus*: nectar-intake rate and crop filling depend on colony starvation. *Journal of Insect Physiology*, (46):1103–1110.
- Kaltenbach, J. C. (1988). Endocrine aspects of homeostasis. *American Zoologist*, (28):761–773.
- Lefebvre, D., Pierre, J., Outreman, Y., and Pierre, J.-S. (2007). Patch departure rules in bumblebees: evidence of a decremental motivational mechanism. *Behavioral Ecology and Sociobiology*, (61):1707–1715.
- Liang, J., You, H., and Wang, X. (2010). A hormone-modulated emotional model. In *2010 2nd International Conference on Computer Engineering and Technology*. [http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5485816 accessed September 2012].
- Manicka, S. and Di Paolo, E. A. (2009). Local ultrastability in a real system based on programmable springs. In Kampis, G., Karsai, I., and Szathmari, E., editors, *Advances in Artificial Life, Proceedings of the 10th European Conference on Artificial Life, ECAL09*, number LNAI 5777, pages 87–94. Berlin: Springer.

- Metzinger, T. (2006). Conscious volition and mental representation: towards a more fine-grained analysis. In Sebanz, N. and Prinz, W., editors, *Disorders of volition*, pages 19–48. Cambridge, MA: MIT Press.
- Metzinger, T. (2008). Empirical perspectives from the self-model theory of subjectivity: a brief summary with examples. *Progress in Brain Research*, (168):215–244.
- Moioli, R. C., Vargas, P. A., and Husbands, P. (2009). A multiple hormone approach to the homeostatic control of conflicting behaviours in an autonomous mobile robot. In Tyrrell, A., editor, *Proceedings of IEEE Congress on Evolutionary Computation*, pages 47–54. Piscataway, NJ: IEEE Press.
- Moioli, R. C., Vargas, P. A., Von Zuben, F. J., and Husbands, P. (2008). Evolving an artificial homeostatic system. *Lecture Notes in Artificial Intelligence*, (5249):278–288.
- Montebelli, A., Lowe, R., Ieropoulos, I., Melhuish, C., Greenman, J., and Ziemke, T. (2010). Microbial fuel cell driven behavioral dynamics in robot simulations. In Fellerman, H., 2, A., 3, A., and 4, A., editors, *Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 749–756. Cambridge, MA: MIT Press.
- Moriguchi, H. and Lipson, H. (2011). Learning symbolic forward models for robotic motion planning and control. In *Proceedings of European Conference of Artificial Life (ECAL 2011)*, pages 558–564. Cambridge, MA: MIT Press.
- Neal, M. and Timmis, J. (2003). Timidity: a useful mechanism for robot control? *Informatica*, (4):197–204.
- O’Brien, K. S. and Hooper-Bùi, L. M. (2005). Sucrose intake and percent effort by the red imported fire ant, *solenopsis invicta* buren (hymenoptera: Formicidae). In Lee, C.-Y. and Robinson, W. H., editors, *Proceedings of the 5th International Conference on Urban Pests*, pages 135–139. Malaysia: Perniagaan Ph’ng.
- Portha, S., Deneubourg, J.-L., and Detrain, C. (2004). How food type and brood influence foraging decisions of *lasius niger* scouts. *Animal Behaviour*, (68):115–122.
- Sánchez-Lasheras, C., Könnner, A. C., and Brüning, J. C. (2011). Integrative neurobiology of energy homeostasis - neurocircuits, signals and mediators. *Frontiers in Neuroendocrinology*, (31):4–15.
- Santos, B. A., Husbands, P., and Froese, T. (2010). Accommodating homeostatically stable dynamical regimes to cope with different environmental conditions. In Fellerman, H., Dörr, M., Hanczyc, M. M., and Laursen, L. L., editors, *Proceedings of the Alife XII Conference*, pages 395–402. Cambridge, MA: MIT Press.
- Sauzé, C. and Neal, M. (2010). A neuro-endocrine inspired approach to long term energy autonomy in sailing robots. In *Proceedings of TAROS (Towards Autonomous Robotic Systems) 2010, August 31st–September 2nd 2010*, pages 255–262.
- Schmickl, T., Hamann, H., Strander, J., Mayet, R., and Crailsheim, K. (2010). Complex taxis-behaviour in a novel bio-inspired robot controller. In Fellerman, H., 2, A., 3, A., and 4, A., editors, *Proceedings of the Alife XII Conference*, pages 648–655. Cambridge, MA: MIT Press.
- Schmickl, T., Thenius, R., Strander, J., Hamann, H., and Crailsheim, K. (2011). Robotic organisms - artificial homeostatic hormone system and virtual embryogenesis as examples for adaptive reaction-diffusion controllers. In Kernbach, S. and Fitch, R., editors, *Reconfigurable Modular Robotics: Challenges of Mechatronic and Bio-Chemo-Hybrid Systems (IROS 2011 Workshop)*.
- Shen, W., Will, P., and Galstyan, A. (2004). Hormone-inspired self-organization and distributed control of robotic swarms. *Autonomous Robots*, (17):93–105.
- Simpson, S. J. and Raubenheimer, D. (1993). A multi-level analysis of feeding behaviour: the geometry of nutritional decisions. *Philosophical Transactions of the Royal Society*, (342):381–402.
- Surmeier, D. J., Plotkin, J., and Shen, W. (2009). Dopamine and synaptic plasticity in dorsal striatal circuits controlling action selection. *Current Opinion in Neurobiology*, (19):621–628.
- Timmis, J., Neal, M., and Thorniley, J. (2009). An adaptive neuro-endocrine system for robotic systems. In *Workshop on Robotic Intelligence in Informationally Structured Space, (RIISS ’09)*, pages 129–136. Piscataway, NJ: IEEE Press.
- Tolley, M. T. and Lipson, H. (2011). On-line assembly planning for stochastically reconfigurable systems. *International Journal of Robotic Research*, (30):1566–1584.
- Tsankova, D. D. (2009). Emotional intervention on an action selection mechanism based on artificial immune networks for navigation of autonomous agents. *Adaptive Behaviour*, (17):135–152.
- Vargas, P., Moioli, R., de Castro, L. N., Timmis, J., Neal, M., and Von Zuben, F. J. (2005). Artificial homeostatic system: a novel approach. In Capcarrere, M., 2, A., 3, A., and 4, A., editors, *ECAL 2005*, number LNAI 3630, pages 754–764. Berlin: Springer.
- Walker, J. H. and Wilson, M. S. (2011). Task allocation for robots using inspiration from hormones. *Adaptive Behaviour*, (19):208–224.
- Wegner, D. M. (2003). *The illusion of conscious will*. Cambridge, MA: MIT Press.