

Epigenetic adaptation in action selection environments with temporal dynamics

John Lones, Lola Cañamero, Matthew Lewis

Embodied Emotion, Cognition and (Inter-)Action Lab
School of Computer Science & STRI, University of Hertfordshire
College Lane, Hatfield, Herts, AL10 9AB, U.K.
Johnlones@gmail.com, L.Canamero@herts.ac.uk, matt-l@semiprime.com

Abstract

To operate in dynamic environments robots must be able to adapt their behaviour to meet the challenges that these pose while being constrained by their physical and computational limitation. In this paper we continue our study into using biologically inspired epigenetic adaptation through hormone modulation as a way to accommodate the needed flexibility in robots' behaviour, focusing on problems of temporal dynamics. We have specifically framed our study in three variants of dynamic three-resource action selection environment. The challenges posed by these environments include: moving resources, temporal and increasing unavailability of resources, and cyclic changes in type and availability of resources related to cyclic environmental changes.

Introduction

In autonomous robotics, there is still a trend to develop and tune controllers with certain explicit goals and environments in mind (see e.g., Suganol & Shirai, 2006; Krichmar, 2012 for an overview). This tuning can be either very direct such as pre-determining the weighting of environmental cues, or more subtle through the use of mechanisms such as reward feedback, fitness functions and activity functions (Krichmar, 2012; Lones & Cañamero, 2013).

However, even slight changes in the environment can lead to significant and often unpredictable changes in the trajectory of the same behaviour (Simon, 1969; Braitenberg, 1984; Steels, 1994, Maris & Boekhorst 1996). While environmental changes tend to modify the organism's behaviour in relation to the environmental change (see e.g., Clemens et al., 1978; Crew, 2010; Zhang & Ho, 2011), significant changes to the environment of robots possessing pre-programmed/determined adaptation mechanisms can lead to behaviours that are not only unsuitable but may render the robot inoperable (Tschacher & Dauwalder, 1999; Krichmar, 2012; Lones & Cañamero 2013).

Biological organisms are able to cope with environmental change through long-term evolutionary adaptation, more rapid ontogenetic adaptation, or through learning (Wilson et al, 1994; Cacioppo et al., 2002; Carere et al., 2005). In organisms, a form of epigenetic development occurs through

interactions with uncertain and dynamic environments (Jaenisch & Bird, 2003; Carere et al., 2005). These interactions can lead to changes in gene expression (Fowden & Forhead, 2011; Zhang & Ho, 2011) and subsequently to the appearance of new behaviours (Crews, 2011) adapted to a specific ecological niche (Narain, 2012). Recent studies have shown that hormones provide some of the signals needed to trigger the development of different aspects of the organism (Clemens et al., 1998; Crews, 2010; Fowden & Forhead, 2011).

In past experiments (Lones and Cañamero, 2013) we tested the viability of using epigenetic hormone modulation as a way to allow a robot to adapt to unknown environments. In that study, we placed the same architecture into various environments posing different challenges to the robot. For each experiment, we researched the ability of the epigenetic robot to develop unique behaviours in direct relation to the environmental challenges. In all cases, a significant increase in viability was noticed in the epigenetic model compared to an architecture lacking the epigenetic mechanism.

In the present study, we investigate the ability of a robot, endowed with the same architecture as in the above-mentioned study, to cope with environments posing different types of temporal dynamics problems. In our previous study, the environment we used, while possessing some dynamic qualities, were predominantly static. Changes in that environment occurred as a consequence of the robot's actions. However, in this study, each environment has its own dynamics. This creates an opportunity to examine the robot's behaviour when faced with constantly changing and potentially unpredictable environments.

Robotics model

The robot we have used in this study is the Koala II (www.k-team.com), a medium-sized wheeled robot. It is equipped with 16 infrared (IR) sensors placed around its body, and we use them as both proximity sensors and touch sensors. Proximity IR sensors are grouped to monitor the eight cardinal and ordinal directions surrounding the robot. In our case, this permits the detection of the direction that possesses the least resistance to movement. Touch IR sensors "extend" the

robot’s body by 1cm—what we refer to as “extended body”. Any encroachment of this area is categorised as contact, and the force of contact is dependent upon both velocity and persistence of the encroachment. Finally, we have fitted a webcam to the robot that, in combination with OpenCV, allows the robot to track specific coloured objects. For a more detailed overview of our setup please see (Lones and Cañamero, 2013).

The Physiology of the robot consists of three survival-related homeostatic variables, which must be maintained within a pre-set boundary ($0 \leq h.var_i \leq 100$) for continued survival (see table 1). These three survival-related homeostatic variables are based upon plausible robotic needs in form of energy (E), physical condition (C) and temperature (T).

The robot’s *energy* depletes at a rate equivalent to a basal metabolic rate plus the energy cost of activating subsystems such as vision. Since these subsystems are always active in this implementation, energy decreases at a constant rate of δ per step. *Condition* represents a measure of health for the robot. Deficits occur in a semi-unpredictable manner from collisions. Both variables can be recovered by finding and consuming specific resources. Finally, *temperature* represents the internal heat level of the robot. The robot’s temperature rises as a function of a combination of the environment’s ambient temperature and the robot’s movement speed. Cooling down (dissipation of temperature) occurs at a constant rate. Assuming a moderate or rapid dissipation of excess heat, the robot is able to maintain a steady speed without running the risk of overheating. Table 1 provides an overview of the different internal variables.

H.Var	Ideal Value	Limit	Cause of deficit	Recover per step
E	100	0	$\delta = 0.1 \downarrow$	$1 \uparrow^*$
C	100	0	Contact \downarrow	$1 \uparrow^*$
T	0	100	Movement \uparrow	$0.8 \downarrow$

* The robot must be near the resource for recovery to commence

Table 1: The homeostatic variables of the robot. In this implementation, if Energy or Condition fall below 0, the robot “dies”. Temperature has an inverse effect.

These survival-related homeostatic variables give rise to a Viability zone (Physiological space), following Ashby (1952) and Avila-García and Cañamero (2004). The position in and management of the dynamics of the viability zone provide different ways to quantitatively measure the robot’s performance and wellbeing. Like Avila-García and Cañamero (2003, 2004) and our earlier paper Lones and Cañamero (2013) we have used this idea of the viability zone to create a performance indicator called “comfort”. Comfort provides a measure of the average homeostatic deficit at any time, and the “risk of death”, which indicates how close the internal state is from reaching lethal values. Comfort is calculated on a scale of 0 to 1; with a comfort level close 1 indicating homeostatic variables near their ideal levels. Whereas a comfort level near to 0 would indicate large homeostatic deficits and a high “risk of death”. Along with the comfort

level the standard deviation at specific points is also provided. This allows for a greater insight into the robot’s performance.

Hormones

Apart from providing a measure of wellbeing, the tendency to satisfy homeostatic needs provides part of the foundation for the formulation of motivations. Internal needs modelled as homeostatic variables have long been used to model motivations in robotics, providing efficient and understandable simple models that permit the generation of appropriate goal-oriented movements and behaviours (e.g., Cañamero, 1997; Breazeal & Scassellati, 1999; Arkin, 2003; Bach, 2011). However, in biological systems matters are more complex, as motivations do not come directly from homeostatic deficits. Rather, hormone secretion derived from homeostatic deficits (e.g., ghrelin in the case of hunger) are shown to be behind the formation of motivation (Wallen, 2001; Malik et al., 2008) and the motivational value of environmental cues (Wied, 1976; Martinez, 1981; Frijda 1986). The development of an organism’s hormonal gland activity (in the form of synthesis and release) as well as the development of receptor sensitivity are believed to be susceptible to both endogenous and exogenous environmental cues (Zhang & Ho, 2011). This would suggest that motivation is also in part affected by past experience.

An epigenetic hormonal motivation-like system could potentially provide an efficient method to allow robots to align their needs and goals with challenging environments on a more permanent basis, e.g., to “grow up” adapted to an environment presenting uneven opportunities to fulfil survival-related needs. This process would affect the tolerance to different homeostatic deficits and the priority with which they would be maintained as a function of the developmental environment. Through such an epigenetic process, during the earlier stages of the development of the robot, its hormone glands associated with underrepresented needs would become more sensitive. That is, smaller homeostatic deficits would trigger the same level of hormone secretion as we would see in robots that had “grown up” in a more balanced environment.

Hormones are however not limited to motivations. In our previous study (Lones and Cañamero 2013), we showed how an epigenetic hormone-like system can give rise to diverse behaviours tailored to different environments. While hormone-modulated behaviours had already been successfully modelled in the past, what sets our model apart from others such as (Avila-García and Cañamero, 2004; and Krichmar, 2012), is that: (a) instead of having a limited number of pre-set behaviours, behaviours emerge from the combination of the hormone-activated sub-systems within the robot; and (b) that due to the epigenetic nature of the hormone glands, this means that two robots with the same motivational tendency but with different developmental histories may behave in different ways.

The Action selection mechanism

The ASM incorporated a “voting-based” (VB) policy based upon ideas presented by Tyrell (1993). By using the VB architecture, actions selected by the robot will be those that

provide the greatest overall benefit. In comparison, a “winner-takes-all” (WTA) policy would lead to the selection of the actions that satisfy the current greatest need. Although Avila-García et al. (2003) found that a WTA outperformed the VB architectures in dynamic environments, in their environments the dynamics was introduced by the presence of predators, thus posing very different challenges. Using our model, in preliminary experiments we found that the VB architecture performed better, as shown in figure 1. These preliminary experiments consisted of five 5-minute runs of each architecture type. Performance was measured using comfort as an indication of the robot’s wellbeing.

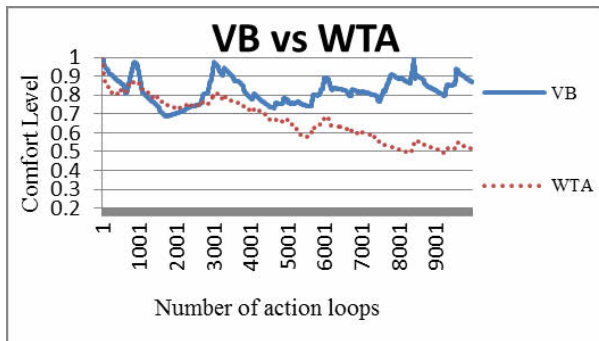


Figure 1: Comparison of the performance of VB and WTA architectures.

Hormone System

At the core of the VB ASM lies a hormone-like system influenced by models developed by Avila-García & Cañamero (2004) and Krichmar (2012). In a new development, we have implemented two different types of hormone, which are classed as either endocrine hormones (Eh) or neurohormones (Nh) (see Table 2). Drawing on biological systems, our Eh-like implementation consists of hormones with the primary purpose to try to maintain homeostasis (Murphy & Bloom 2006). The Eh group is made up of three hormones: one associated with each homeostatic variable. For each hormone, h , secretion occurs via a gland, g_h , and the rate of secretion, s_h , depends upon the current homeostatic deficit, d_h , and the activity level of the gland, Ψ_h ,

$$s_h = k(\Psi_h^{d_h}) \tag{1}$$

where k is a constant that scales the size of secretion. Once released, each secretion persists in the system for a random number of action loops (within a fixed range) before decay of that particular secretion occurs. The larger the secretion, the longer it will take to fully decay. The concentration ($conc_h$) of these hormones is thus determined by the total sum of each active secretion.

The second group of hormones, Nh, contains only one hormone, D1. This hormone facilitates what can be described as “dominant” or potential “aggressive” behaviour. This is achieved by having the hormone suppresses environmental cues that are associated with negative stimuli. For example a robot with a high D1 level that detects a desired resource will move towards it directly at a high speed pushing aside any obstacles, disregarding the potential of damage from

collisions. In contrast a robot in the same situation but with a low D1 level would instead move around obstacles to reach the desired location.

Rather than being triggered by internal deficits, as with Eh hormones, Nh secretion is linked to the mean of the external environmental cues (ec_{dt} where d is the direction of the cue and t the type e.g. energy or repair source). Therefore the $conc_h$ of the neurohormone is determined by

$$conc_h = (conc_h\beta) + \alpha\bar{e}\Psi_h \tag{2}$$

where α is a predetermined weighting factor ($0 \leq \alpha \leq 1$) and β is the disperses rate of the hormone which is set to 0.9 (leading to a 10% disperse rate each loop) during these experiments.

Also different to the Eh model, Ψ_h is not a set value. Instead, the activity level of the gland is stimulated by the mean concentration of the Eh hormones ($conc_{EH}$) in the body, similar to tropic hormones. Where in biological systems, these hormones have been demonstrated to cause/increase the secretion/production of other hormones (Sherwood, 2003):

$$\Psi_h = e^{-4e^{-0.5conc_{EH}}} \tag{3}$$

Ψ_{D1} will lie between a value of 0 and 1, with a value of 1 indicating the gland is fully active, and 0 signifying that the gland is inactive. The final part of the Nh equation models neuroreceptor sensitivity (sen_h).

$$streng_h = e^{de \times e^{sen_h conc_h}} \tag{4}$$

where $streng_h$ the cumulative effect of the neurohormone on the system once the concentration and sensitivity to it are taken into account, de is the minimum stimulation needed for activation of the receptor, and sen_h the sensitivity of the receptor to the hormone.

H. Name	H. Type	Trigger	Ψ_h
E1	Eh	Energy deficit	0.09
C1	Eh	Cond deficit	
T1	Eh	Temp deficit	
D1	Nh	Visual cues	Varies

Table 2: Robotic hormones

Hormones and the ASM

The VB ASM consists of a two-step computation (see figure 2). The first step calculates the current homeostatic motivations or drives (m_{dt}) (see Table 3). Although three drives are present, we only need to directly calculate the intensity of hunger and damage. The hyperthermia drive, which can be satisfied by reduced or no movement, instead suppresses other drives. In addition to the internal state, motivations are influenced by environmental cues,

$$m_{dt} = \begin{cases} \frac{conc_h(t)ec_{dt}}{1+conc_{T1}} & \text{if } strength_{D1} > 0.75 \times conc_{T1} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

The perceived environmental cue in the forward direction, ec_{2t} , is given an additional +1 score to simulate a restless mechanism and allow forward movement without external stimuli. To further reduce excessive switching of motivations ec_{2t} is given a 10% bonus to its value as a form of “hysteresis”.

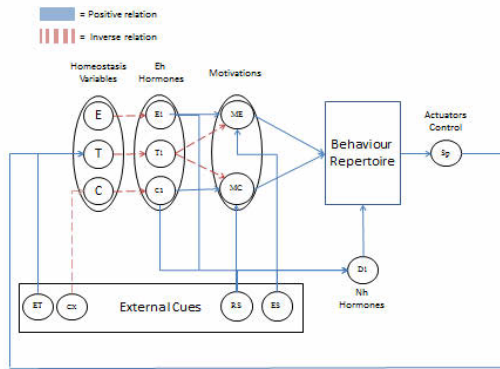


Figure 2: Hormone-based Architecture.

Motivation	Drive	ext Stimuli	Suppressed by
Hunger (ME)	E1	Eng source	T1
Damage (MC)	C1	Rep source	T1
Hyperthermia	T1	Climate	E1 & C1

Table 3: Motivations of the robot

The second step in the ASM calculates the behaviour to execute given the current motivational state and environmental conditions. Unlike previous hormone-based architectures such as Avila-García & Cañamero (2004) and Krichmar (2012) no explicit behaviours have been modelled. In our case, behaviours occur from dynamic combinations of different systems with no pre-set physiological cost or gain. The cost or gain of behaviour execution results from the sum of physiological changes that occurred during the action.

One of these subsystems is the robot’s personal space (Ps) (see Hall, 1966), an area that the robot will treat almost as an extension of its own body. Using a similar technique as with the “extended body” (the IR-based touch sensors around the robot’s body), the robot will normally maintain the Ps free from other objects. The radius of the Ps zone is determined by the current C1 hormone concentration. Encroachment will lead to attempts to re-establish a space by moving along the path of least resistance (d_n), with a slight preference to going forward. D1 counteracts the tendency to keep the Ps empty, allowing objects within the Ps while trying to satiate drives. At high levels D1 will facilitate physical contact, allowing the robot to push or “attack” anything standing between itself and its target, the size of the Ps at any given time is show in equation 6.

$$Ps = \left(ns + ns \frac{conc_{C1}}{conc_{C1\{max\}}} \right) (1 - strength_{D1}) \quad (6)$$

where ns is the normal, unadjusted size of personal space, and $conc_{C1\{max\}}$ the maximum potential concentration of the hormone.

Hormone-Signalled Epigenetics

The final and following aspect of the model introduces an epigenetic adaptation mechanism into the architecture. Taking inspiration from recent biological studies (see Crews, 2008 & 2010, Fowden & Forhead, 2011 for an overview) hormones trigger epigenetic changes in the robot. In our robot, hormone levels both indirectly and directly provides a fairly accurate measure of current conditions in the environment and level of situatedness. For instance, the current level of the E1 hormone is an indication of how well the robot is managing its need for energy. Combined with the concentration of D1, it is possible to determine the root of the imbalance, as either issue of scarcity, or difficulty of access to the resources.

These hormones can thus act as signals for epigenetic adaptation, whereby development of the glands that secrete hormones and receptors that receive them are influenced by the external environment. For example, an autonomous robot that is often low on condition/health will have a high concentration of the C1 hormone within its system. The high concentration will lead to a long-term increase in the activity level (Ψ) of the gland that secretes C1. This will mean that sub systems such as the desire to maintain a degree of personal space or find repair resources will be much more prevalent within the model. Formula 7 shows method used to facilitate the epigenetic change in activity levels (Ψ) of the gland for hormones in the eh group.

$$\Psi_{h+} = \frac{conc_h}{l} \quad (7)$$

where l is a constant to regulate the speed of epigenetic change.

Formula 8 shows the method in which epigenetic change can occur to the sensitivity of neurohormone receptors for the hormones in the Nh group

$$sen_h = \frac{(strength_h - 0.05)}{j} \quad (8)$$

where j is a constant to regulate the speed of epigenetic change.

Drawing on the notion of critical periods in biological organisms, the epigenetic process above is active during the early period of the robot’s life. This critical period represents a window frame when organisms are most susceptible to the influences of external perturbations (Winks & Berthouze, 2008), mediated via hormone modulation (Crews, 2010), among other things.

The temporal three-resource problem

The architecture described here has been tested in a temporal three-resource action selection problem framework, in which a robot needs to timely and appropriately select among and satisfy three needs using resources available in the environment in order to survive (remain operational or “alive”). Our experimental design included three different sets of experiments corresponding to three variants of an environment that pose different challenges arising from the temporal dynamics of the resources. Each set took place within a 2mx2m bordered environment inhabited by a single robot. Within each environment a number of energy and repair resources were available to allow the robot to replenish homeostatic deficits. These resources were represented by two different coloured sets of balls. The environments also contain an ambient temperature that is sensed internally by the robot.

Scenario one consists of the base environment with one of each resource moving in a continuous pattern at a constant speed, slightly faster than the robot’s average, around the arena, see figure 3. At the end of each movement path (represented by a letter) the resource would pause for a period of 2 seconds. In cases where the robot was in the direct path of a resource, the resource would be manoeuvred around the robot using the shortest path before returning to its original trajectory. In the case where the resource was pinned or the movement was blocked by the robot, no attempts were made to push the robot aside. Instead, movement of the resource was halted until the robot moved away and a viable path was visible. At the start of each run the resources started at a different opposite points, e.g. A and E.




Figure 3: The pathways and start points of the resources.

Scenario two was again based on the base environment. However, in this scenario the energy resource appears at set points within the environment once every minute, the period during which it is available reduces over time, i.e., it becomes decreasingly available. For the first five runs, the energy source would remain for 30 seconds before being removed. In the second five sets the duration was reduced to 20 seconds and the final five saw the resource only accessible for 10 seconds of every minute. The set points are the same as the start of pathways as seen in figure 3. In order to avoid biases, the order of set points where the resource would appear was

predetermined randomly before each run. The choice to have the temporal properties apply to only the energy source was done to examine the robot’s ability to deal with the increasing disparities between the availability of the repair and energy sources.

It is worth noting that the robot has no capacity to monitor time. Therefore, there is no facility to try to directly predict when the resource will appear. Rather, over time the robot will adapt to the scarcity and rarity of the resource. The use of a strict time period was to ensure each robot had same constraints and opportunities.

Scenario three examines the ability of the robot to adapt to the effects of dynamic climatic changes. In this experiment the standard base set up of the environment was used with one of each resource available at all times. However the ambient temperature of the environment would increase and decrease over time, simulating a day-and-night temperature cycle. The entire cycle lasts for four minutes, as can be seen in figure 4. To simplify the model, ambient temperature ranked between 0 (cold) and 10 (scorching heat).

In order to increase the dynamics of the environment, temperature was allowed to fluctuate by up 2 points to simulate potential meteorological phenomena. The fluctuations were calculated at start of each 10-second period and lasted until the next period.




Figure 4: An example of an average weather cycle with meteorological phenomena. The periods between 6 and 18 or minute 2 and 3 are analogue to daytime, with the highest temperature occurring midday equivalent to the sun at its peak in a natural environment.

Experiments and Results

The robot was tested over a total of 35 runs split in 10/15/10 runs amongst the three previously described scenarios. Each run lasted a maximum of 10,000 steps around 10 minutes 40 seconds per run. The epigenetic system was active during the first 3 minutes (2880 steps). A second set of runs was conducted in the same manner for a robot without the epigenetic mechanism to serve as a basis for comparison. The viability of both architectures was assessed using the previously discussed *Comfort* measure and standard deviation as well as visible observation. In cases where a robot died before the end of a run, a comfort value of 0 would be recorded for any remaining loops.

Scenario One

This environment provides the robot with two distinct challenges. The first and most obvious was the need to develop a consumption behaviour suitable for moving resources. Secondly, this environment presents the first situation where the robot can be damaged by other elements (objects or organisms) of the environment. While as previously stated resources will move around the robot if it is directly in their path, they will still move close enough to encroach upon the extended body, causing damage. Therefore, the robot will also need to adapt to co-exist with the resources, not just how to exploit them. The results of the first experiments can be seen below in figure 5.




Figure 5: The combined results for Scenario One.

As can be seen in figure 5, the epigenetic robot performed at a higher level overall, but more interestingly had a much lower standard deviation of 0.05 compared to 0.17. The differences in standard deviation can be attributed to the dynamic nature of the resources. In some situations the robot was positioned in the ideal location to catch and consume resources as they passed. This led to timely management of the robot’s homeostatic needs. However, in other cases the robot would need to actively move across the arena and chase a resource. Since the resources moved slightly faster than the robot’s average speed, the motivation to consume the resource had to outweigh the motivation to limit speed in order maintain a low temperature.

Distinctive behaviour developed for each of the architectures in this environment. The epigenetic model would develop an “ambush-like strategy”: the robot would remain sedentary until an energy source passed closely, at which point the robot would give chase at full speed often pinning the resource to a wall until it had finished consuming it.

In contrast, the non-epigenetic model would engage in “drawn-out chases”. As the motivation to consume the resource allowed it to generate the speed needed to catch up, excess heat was generated. This heat generation led to premature end of the chase on a number of occasions. Finally the epigenetic robots displayed more adaptive behaviour at avoiding unnecessary collisions with resources, and almost no unwanted collisions occurred after the early periods.

Scenario Two

In scenario two, we tested the ability of the robot to deal with resources only available for limited periods of time. The 15

runs were divided into 3 groups of increasingly challenging runs with the resource present for 30/20/10 seconds of every minute, challenging the robot to act in a timely manner when the opportunity to recover from homeostatic deficits was present. This temporal quality only applied to the energy resource. This further challenged the robot to overcome the “distraction” of the more readily available repair resource. The results of this scenario can be seen below in figure 6.




Figure 6: The combined results for Scenario Two.

Both robots performed at a similar level during the first five runs with 30 second window of opportunity. While the epigenetic robot moved more promptly to resources when they appeared, neither robot ever was in any real danger. However, as the window of opportunity shrunk, the differences between the two models became very apparent, as can be seen in figure 6.

As the point where the resource would appear next was unknown to the robot, it was inevitable that both architectures would miss some opportunities to replenish. However, the epigenetic model was generally quicker to find any resource due to the development of the E1 and D1 glands, thus giving the robot a greater chance of survival also when the opportunities were missed.

Finally, due to missed opportunities to fully recover deficits, both robots often contained significant level of the D1 hormone. This in turn resulted in higher occurrences of collision in later runs, subsequently increasing the need for repair resources. In multiple cases this lead to similar levels of need for both the energy and repair resource. This resulted in the non-epigenetic robot sometimes going to the readily available repair source during the limited periods when the energy source was present and seen. This occurred on some occasions even when condition deficits were not significant. In contrast, the epigenetic model had adapted to the rarity of the resource. It only missed the opportunity to replenish energy once. This occurred when its condition levels were critical. In total, 7 of the non-epigenetic robot runs ended prematurely compared to a single death in the epigenetic model. Due to the high level of fatalities, the hormone-only model actually had a lower standard deviation of 0.03 in contrast to 0.08 in the epigenetic model.

Scenario three

In the final scenario we tested the ability of the two robot architectures to deal with cyclical climates, with the cycle of change in ambient temperature previously shown in figure 4. Like scenario two, this environment challenged the robot’s

ability to take advantage of limited windows of opportunity. During the periods where ambient temperature reached its peak, even limited movement soon led to overheating. Two of each of the resources, spread evenly in each corner, were constantly available in the environment. The results for this experiment can be seen in figure 7.




Figure 7: The combined results for Scenario Three.

As can be seen, the epigenetic robot had much greater success. After the initial 3 or 4 cycles the robot's hormone glands had developed in such a way that, during periods with the highest ambient temperature, virtually all actions would be suspended. As soon as the ambient temperature dropped, the robot would move to replenish any deficits. The epigenetic robots developed two contrasting behaviours in order to survive the periods of high ambient temperature. One group simply over consumed and in effect "hibernated". The second group would instead stay near the energy source at all times apart from the occasional need to repair, allowing itself to consumer energy during the increased climate with only very limited movement needed.

In contrast, the non-epigenetic model often ran low on energy during the day cycle. This resulted in the robot being forced to move to energy sources, generating significant overheating, which led to the death of the robot on 3 occasions.

Conclusion

In our past study (Lones and Cañamero 2013) we have shown how epigenetic changes through hormone modulation increase the adaptability of a robot. Specifically we demonstrated how this process leads to behaviours tailored to specific environmental niche. These robots were placed into different environments with exactly the same starting architecture. However, through epigenetic processes, the robots developed distinct traits and behaviours depending on the environment in which they developed.

In the study presented in this paper, we have investigated the same architecture under new criteria. Specifically, we focused on the ability of the robot to adapt to environments that presented temporal dynamics challenges. In the first experiment, the robot needed to adapt to fast-moving resources. While the robot could simply have "chased after" the resource at top speed, this would lead to unwanted overheating and would not guarantee appropriate satisfaction of its homeostatic needs. Instead, the robot developed what could be considered equivalent to an "ambush-like hunting

tactic". In the second experiment, the robot was challenged to adapt to limited windows of opportunity to satisfy a homeostatic need, all the while needing to adapt and disregard opportunities offered by more easily available resources that permitted to satisfy other needs. Needing to find a balance between maintaining the different homeostatic needs, the robot was able to respond in a timely manner to rare occurrences while still finding time to satisfy the other needs. In the final experiment, we examined the ability of the robot to adapt to cyclical events. Under this scenario, the robot needed to fully utilise the cooler periods of the day, which allowed it to be in a position to survive hotter periods when most actions would need to be suspended. This experiment marked the first time we saw the epigenetic model divide into two distinct groups. Each group developed a different method to deal with the debilitating temperature.

As we have shown, epigenetic adaption though hormone modulation potentially offers a suitable method to allow a base architecture to develop behaviours to adapt to environments presenting different temporal challenges.

Acknowledgements

John Lones is supported by a research studentship of the University of Hertfordshire. Partial support was also provided by the FP7 project ALIZ-E. We would like to thank the other members of our lab for their feedback. John is also grateful for additional feedback and advice from Zornitsa Trendafilova.

References

- Arkin, R (2003). Moving up the food chain: motivation and emotion in behavior based robots. In *Who Needs Emotions: The Brain Meets the Robot*, Fellous J., Arbib M., editors. Oxford University Press, p.p 245–270.
- Ashby, W. R (1952). *Design for a Brain*, Chapman and Hall, London
- Avila-García, O., Cañamero, L. and Boekhorst, R. (2003). Analyzing the performance of 'Winner-Take-All' and 'Voting-Based' action selection policies within the two-resource-problem. In *Proceedings of seventh European conference in artificial life (ECAL03)* p.p 733–742
- Avila-García, O. and Cañamero, L. (2004). Using hormonal feedback to modulate action selection in a competitive scenario. In: *From Animals to Animats: Proceedings of the 8th International Conference of Adaptive Behavior (SAB'04)*, Cambridge, MA, MIT Press, p.p 243–252.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology* Cambridge, MA:MIT press/Bradford Book
- Breazeal, C. and Scassellati, B. (1999). "A context-dependent attention system for a social robot," in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI 99)*. Stockholm, Sweden, 1146–1151.
- Cacioppo, J., Brenston, G., Taylor, S. and Schacter, D. (2002). *Foundation in Social Neuroscience* Cambridge, MA:MIT press
- Cañamero, L. (1997). Modeling motivation and emotions as a basis for intelligent behaviour. In W.I Johnson editor *proceeding of the first int conf on autonomous agents (agent 97)*, pages 148-155. New York: ACM Press
- Carere, C., Drent, P., Koolhaas J. and Groothuis, T. (2005). Epigenetic effects on personality traits: early food provisioning and sibling competition *Behaviour* volume 142 number 9-10, p.p 1329-1355
- Clemens, L., Gladue, B. and Coniglio, L. (1978). Prenatal endogenous androgenic influences on masculine sexual behaviour and genital morphology in male and female rats. *Hormones and Behavior* volume 10 issue 1, p.p 40-53

- Crews, D. (2008). Epigenetics and its implications for behavioural neuroendocrinology. *Frontiers in Neuroendocrinology* 29 p.p 344-357
- Crews, D. (2010) Epigenetics, brain, behavior, and the environment *HORMONES*, 9(1): p.g 41-50
- Fowden, A. L. and Forhead, A. J. (2009). Hormones as epigenetic signal in developmental programming, *Exp Physiol.* June; p.p 607-25.
- Frijda, N. (1986). *The Emotions*, Cambridge University Press
- Jaenisch, R. and Bird, A. (2003). Epigenetic regulation of gene expression: how the genome integrates intrinsic and environmental signals. *Nat Genet* 33(Suppl): p.p 245–254
- Krichmar, J. L. (2012). “A biologically inspired action selection algorithm based on principles of neuromodulation” in *The 2012 International Joint Conference on neural networks* p.p 1-8
- Lones, J. and Cañamero, L. (2013). Epigenetic adaptation through hormone modulation in autonomous robots. Submitted for publication
- Malik, S., McGlone, F., Bedrossian, D. and Dagher, A. (2008). Ghrelin modulates brain activity in areas that control appetitive behavior. *Cell Metab* 7:400–409
- Maris, M. and Te Boekhorst, R. (1996). Exploiting physical constraints: heap formation through behavioral error in a group of robots. In *Proc. IROS'96, IEEE/RSJ International Conference on Intelligent* p.p 1655-1660
- Martinez, J. L. (1981). *Endogenous Peptides and Learning and Memory Processes* Academic Press Inc
- Murphy, K. G. and Bloom, S. R. (2006). Gut hormones and the regulation of energy homeostasis. *Nature* 444 p.p 854-859
- Narain, C. (2012). Changing behaviour with epigenetics, *Natural Neuroscience* 15, p.p 1329
- Sherwood, L. (2003). *Human Physiology: From Cells to Systems* Wadsworth Publishing Co Inc; 5th Revised edition
- Simon, H. A. (1969). *The Sciences of the artificial*. Cambridge, Mass.: MIT Press.
- Steels, L. (1994). A case study in the behaviour oriented design of autonomous agents
- Sugano, S. and Shirai, Y. (2006). Robot design and environment design Waseda Robot House Project”. In *Proc. SICE-ICASE 31–34. Busan*
- Tschacher, W. and Dauwalder, J P. (1999). *Dynamics, synergetics, autonomous agents—nonlinear systems approaches to cognitive psychology and cognitive science* World Scientific Publishing Co Pte Ltd
- Tyrrell, T. (1993). *Computational Mechanisms for Action Selection* PhD thesis University of Edinburgh
- Wallen, K. (2001). Sex and context: hormones and primate sexual motivation. *Horm. Behav.* 40, 339–357
- Wied, D. (1976). Hormonal influences on motivation, learning, and memory processes. *Hosp Practice* 11:123–131
- Sloan Wilson, D., Clark, A. B., Coleman, K. and Dearstyne, T. (1994). Shyness and boldness in humans and other animals *Trends Ecol Evol* 9 p.p 442-446
- Winks, O. and Berthouze, L. (2008). Predicting the timing of critical periods in development In *Proceedings of the 8th International Conference on Epigenetic Robotics*, pp. 127-134
- Zhang, X. and Ho, S. M. (2011). Epigenetics meets endocrinology *Journal of Molecular Endocrinology* 46, p.p 11–32