

A Neural-Endocrine Architecture for Foraging in Swarm Robotic Systems

Jon Timmis and Lachlan Murray and Mark Neal

Abstract This paper presents the novel use of the Neural-endocrine architecture for swarm robotic systems. We make use of a number of behaviours to give rise to emergent swarm behaviour to allow a swarm of robots to collaborate in the task of foraging. Results show that the architecture is amenable to such a task, with the swarm being able to successfully complete the required task.

1 Introduction

Swarm robotic systems have many potential uses, ranging from the cleanup of hazardous waste or search and rescue operations at disaster sites that are often too dangerous for humans to respond effectively to or areas that need large coverage for monitoring (such as the ocean) and are simply too large a task for a single robot to cope. Good reviews of swarm robotics and associated issues can be found in [10, 7, 5]. However, in order to develop such systems, the task of *foraging* is used as a standard test arena for new approaches. Foraging is a popular task for mobile autonomous robots, both individual robots and swarms have been shown to successfully complete various types of foraging problem. The basic principles of foraging involve an agent collecting objects that are spread throughout the environment and returning them to some specified location. The task is completed once all of the

Jon Timmis

Department of Electronics and Department of Computer Science, University of York, Heslington, York. UK e-mail: jtimmis@cs.york.ac.uk

Lachlan Murray

Department of Electronics, University of York, Heslington, York. UK e-mail: ljm505@ohm.york.ac.uk

Mark Neal

Department of Computer Science, Aberystwyth University, Aberystwyth, Wales. UK. e-mail: mjn@aber.ac.uk

objects in the environment have been collected. Part of our on-going work is the development of a neural-endocrine architecture for deployment in ocean going robotic systems, and the eventual construction of a swarm of ocean going vessels that would be able to operate for prolonged periods of time.

This paper investigates and extends our previous work on a neural endocrine control architecture developed in [2, 3, 8, 6]. Until now its effectiveness at controlling a collection of robots has not been investigated, though work on using two robots has been undertaken in the context of task switching [9]. The addition of more robots brings added complexity to the system, it is necessary that a multi-robot control system not only encompasses the ability to control individual robots, but is also capable of appropriately handling the interactions with other robots. If we are to work towards developing an ocean going version of such a system then the understanding of the ability of our architecture to operate in a swarm of robots is essential. In order to assess the effectiveness of the system it was necessary to design a task for the robots performance to be measured on, the task chosen was a variant of foraging and was one of the most complicated tasks that the neural endocrine control architecture has been applied to. Specifically, in this paper we: investigate whether the neural endocrine control architecture is capable of controlling a multi-robot system; investigate how effective the architecture is at controlling a multi-robot system and investigate the capabilities of the architecture at a new and complex task.

2 Neural Endocrine Control Architecture

The neural endocrine control architecture of [2] is a combination of standard perceptron artificial neural networks, with a novel artificial endocrine system that has the ability to affect the weights of the neural networks, depending on external and internal factors. Here we review the basic neural endocrine architecture, for a more detailed description the reader is directed to [3, 6].

2.1 Artificial Endocrine Systems

The Artificial Endocrine System (AES) described here is based on the original design proposed by [2, 3] as well as subsequent modifications made by [6].

As is the case in the biological endocrine system, the two main components of an AES are glands and hormones. Artificial glands (g) release artificial hormones when they are stimulated. Stimulation can be caused by both the internal state of the system and external stimuli. In [6] *signal values* A_i were obtained by summing sensor inputs and similar gland *activation values* were calculated from the combination of sensor values and the internal state of the robot. The stimulation of a gland (R_g) as given by [6] is shown in equation 1 where α_g is the *stimulation rate*, that is the rate at which a hormone is released from a gland g .

$$R_g(t) = \alpha_g \sum_i A_i(t) \quad (1)$$

Our previous work, unpublished, investigated a second method of stimulation that also takes into account the current *concentration* of hormone $c_g(t)$ this is given by equation 2. As can be seen in equation 2 the amount of hormone released in this method is subject to a negative feedback mechanism, the reason for including this is to prevent the system from becoming over saturated with a particular type of hormone.

$$R_g(t) = \frac{\alpha_g}{1 + c_g(t-1)} \sum_i A_i(t) \quad (2)$$

Every hormone has an associated decay rate (β_g) which takes a value from $[0, 1]$, this means that without stimulation the concentration of a hormone will eventually be reduced to an insignificant amount. The concentration of a particular hormone c_g at time $t + 1$ is given by equation 3.

$$c_g(t+1) = \beta_g c_g(t) + R_g(t+1) \quad (3)$$

2.2 Neural Endocrine Systems

Artificial hormones can only affect artificial neurons. In line with the biological endocrine system not all of the neurons in a system will be sensitive to all hormones, the sensitivity of a neuron i to the hormone released by a particular gland g is given by s_{ig} . The effect that hormones have on neurons can be calculated by equation 4 which takes into account the sensitivity of inputs to particular hormones and the concentration of those hormones using an artificial endocrine system with ng glands.

$$u = \sum_{i=0}^{nx} x_i \cdot w_i \sum_{g=0}^{ng} c_g \cdot s_{ig} \quad (4)$$

The most common form of coordination between networks is a cooperative approach whereby the outputs of each network are simply summed together. The resulting behaviour of a multi-network neural endocrine system is dependent on the current hormone levels of the system. High levels of a particular hormone will affect some networks more than others, giving these networks more or less influence over the global result when the network outputs are summed together.

3 System Design

3.1 Behaviours

In this work, we make use of eleven different behaviours, the majority of which can be categorised into the three different groups: taxes, reflexes and fixed-action patterns (FAP). One of the behaviours, *wander*, can not easily be classified by type. We also observe resultant emergent behaviours not programmed into the system.

Wander: A wander behaviour is necessary to ensure that robots keep exploring the environment even if none of their other behaviours are currently being stimulated, without a wander behaviour an unstimulated robot would just remain stationary. To implement a wander behaviour we take into account the current hormone levels of the system.

3.1.1 Reflexes

Reflexes are involuntary, spontaneous responses to stimuli, which last only as long as the stimulus that initiates them. The foraging task of this work requires only a single reflex behaviour. Because of their spontaneous and sporadic nature reflex behaviours do not require a neural endocrine control network, their response is simply tied directly to their stimulus.

Signal bin: As robots will have no awareness of the location of the bin, in order to improve their chances of finding it a signal bin behaviour is required, allowing robots to communicate the approximate location of the bin to others. In this case, robots signal that the bin is in their vicinity by the use of a light or beacon. The strength of the response should always be the same, i.e. the brightness of the light should not be effected by the closeness of the bin, it should either be on if the bin is in-sight, or off otherwise.

3.1.2 Taxes

Taxes are behavioural responses that cause agents to move towards, or away from certain stimuli. This work involves six taxes behaviours, two of which are repellent and four of which are attractive. Taxes behaviours are well suited to control using neural endocrine networks because both their inputs and outputs are continuous and should vary according to the current state of the system, i.e. the hormone levels. Robots have the capability of signalling via an LED, and observing that signal on other robots.

Obstacle avoid: Prevents robots from crashing into the walls of the environment, or obstacles within the environment. The response of an obstacle avoid behaviour should be proportional to the distance between a robot and its nearest obstacle, such that a robot responds more urgently to obstacles that are nearer. The inputs to the

network of an obstacle avoid behaviour come from a range finding sensor, for example a sonar.

Separation: Prevents robots from crashing into each other. The stimuli of a separation behaviour, also the inputs to the behaviour's network, are the locations of other robots, these can be determine using a camera device. In a similar manner to obstacle avoidance, the strength of a response should be proportional to the distance between a robot and its neighbours, such that the closer a fellow robot is, the faster the robot should retreat.

Cohesion: Attracts a robot its neighbours. As with separation, a cohesion behaviour is useful in the development of emergent global behaviours. The strength of the stimulus should have an effect on the strength of the response so that robots are less attracted to neighbours that are closer, reducing the chance of collisions. The inputs to a cohesion behaviour's network are similar to those of a separation behaviour and come from the positions of their neighbours via a camera device.

Seek rubbish: Robots should be stimulated by the presence of a piece of rubbish, which can be detected using a camera. Robots should be attracted to the location of the rubbish with a strength of response that is relative to the how far away the rubbish is, the further away, the stronger the attraction.

Seek power: Robots should be attracted to charging stations. Inputs are provided in the same manner as the seek rubbish behaviour, using a camera device, and the strength of the response is once again relative to the distance of the stimulus.

Seek bin : A seek bin behaviour is very similar to both the seek power and seek rubbish behaviours, however in this case robots should be attracted to the bin. The stimulus is the presence of the bin, and the strength of response is relative to the distance between the robot and the bin.

3.1.3 Fixed-action-patterns

Fixed Action Patterns (FAP) are behaviours that continue even if the stimulus that triggered them is not present, usually they run uninterrupted until completion. Their response is always identical and so they are not suitable for control using neural endocrine networks, like reflexes they can be implemented by directly tying stimulus to response.

Pickup rubbish: A pickup rubbish behaviour should be stimulated when a robot is close enough to a piece of rubbish and is not already carrying some. The behaviour should involve the robot moving towards the piece of rubbish and either successfully or unsuccessfully picking it up, both of which should result in the end of the pattern, however if the pickup is unsuccessful it is possible that the behaviour will be re-stimulated immediately.

Drop Rubbish: If a robot is carrying a piece of rubbish and is close enough to the bin, the drop rubbish pattern should be stimulated. The pattern starts with the robot approaching the bin and continues until the robot has either successfully or unsuccessfully dropped the rubbish into the bin.

Recharge: A recharge behaviour should be stimulated when a robot is close enough to a charging station and its internal state dictates that it needs to recharge. The behaviour should begin with the robot moving towards the charging station and attempting to dock with it, if the robot fails to dock, the pattern should end, if the robot successfully docks the pattern should continue until the robot is fully charged.

3.2 Neural-Endocrine Design

In section 2, it was noted that not every hormone in a system must affect every neuron. In all previous work the approach has been to make all the neurons of a single network sensitive to the same hormones, for example in a system with two hormones h_a and h_b and two networks N_a and N_b , a possible configuration would be that all the neurons of N_a are sensitive to h_a and all the neurons of N_b are sensitive to h_b , this is shown in figure 1. The alternative is to make different neurons of the same network sensitive to different hormones, for example in a system with two hormones h_a and h_b and a single network of seven nodes $\{n_1, n_2, \dots, n_7\}$, nodes $n_1 - n_4$ might be sensitive to h_a and nodes $n_5 - n_7$ might be sensitive to h_b , this is shown in figure 2. Since each network in a system corresponds to a single behaviour, it seems sensible that, as is the case in the previous approach, each network should be affected by the same hormones. For simplicity, here each network is only associated with a single gland-hormone pair. The sensitivity of a neuron i to a particular gland g is denoted s_{ig} , in theory s_{ig} can take any value, however in this work s_{ig} only takes the value 1 or 0, representing full or no sensitivity of i to g .




Fig. 1 Two networks, the neurons of which are all sensitive to the same hormone

Having decided that all neurons in a network will be sensitive to the same hormones, that each network will only respond to one hormone and that sensitivity is only ever 1 or 0 it is possible to refine equation 4 from section 2. Removing the sensitivity and multiplicity of 4 leaves 5 where c_g is the concentration of the network's only associated hormone.

$$u = \sum_{i=0}^{nx} x_i \cdot w_i \cdot c_g \quad (5)$$




Fig. 2 A single network with different neurons sensitive to different hormones

The activation of a gland can be calculated from a combination of both internal and external properties of the system. Each gland is associated with a single activation parameter which changes over time according to a dedicated function and is represented here as a_g . The stimulation of a gland, as was seen in section 2 can be calculated in one of two ways. Early implementations had shown success with using a negative feedback mechanism, we therefore adopted that approach in this work. The stimulation of a gland is calculated using 6, which is a slightly adjusted version of 2 in order to take into account the new representation of activation.

$$R_g(t) = \frac{\alpha_g \cdot a_g(t)}{1 + c_g(t-1)} \quad (6)$$

The final consideration with ANN-AES integration, is what values of stimulation (α_g) and decay (β_g) rate are used by the networks. The stimulation rate helps determine the amount of hormone released by a gland at a particular time-step and the decay rate determines how long the hormone remains in the system, hence they both have a big influence on the behavioural response. Values of α_g and β_g can vary widely between different networks. These values were chosen experimentally, however an automated learning process could be adopted.

3.2.1 Network size and Weights

ANNs can be defined by four properties: the number of hidden layers, the number of nodes in each of the hidden layers, the number of nodes in the input layer and the number of nodes in the output layer. For more information on neural networks, the reader is directed to [1]. The number of nodes in the input layer of a network are determined by the number of sensor values needed to define the stimulus of that behaviour, for example in the case of obstacle avoidance which is stimulated by the presence of nearby objects, the number of sonar devices (two in this piece of work) determines the number of input nodes. The number of output nodes is determined by the actuator that the response affects, in most cases, where the response affects the locomotion of the robot, it is the number of inputs to the motors that decides the number of output nodes (which again in this study is two).

The number of hidden layers and the number of hidden layer nodes is less dependent on the behaviour, and puts more pressure on the designer to choose sensible values. It was known from our previous work that the networks required would be relatively simple, consequently we only include one hidden layer.

With regards to setting the weights, we used a combination of determining the weights by hand and back-propagation [1].

3.2.2 Coordination of different behaviour types

Behaviours that are encapsulated as neural endocrine networks are coordinated by summing their outputs. We have not discussed how these behaviours are coordinated with the other types of behaviour, such as the fixed-action-patterns and reflexes. The signal bin behaviour is a reflex, it does not affect anything other than the state of the robot's beacon and so it does not need to be coordinated with the other behaviours. In terms of the FAPs, when stimulated, these will always take complete control of the robot's motors, inhibiting any of the suggested commands from the other behaviours. It is very rare for conflicts to arise between different FAPs since it is never the case that a robot will want to both drop and pickup rubbish at the same time and because the bin and charging posts are positioned far apart (in the experiments carried out in this work) there will never be a conflict between wanting to charge and wanting to drop rubbish. However, we recognise that this is an avenue for further exploration.

3.2.3 Environments

In order to test the adaptability of the system it was necessary to test the performance of the robots in two different environments. Both of the environments were designed with the capabilities of the robots in mind, for example, it was known that because the robots had only two sonar sensors, both of which were located at the front, they would struggle to find their way out of concave obstacles with small internal angles. When faced with concave obstacles robots can be indecisive about which way to turn and in the end may either end up stalling or crashing into the obstacle. Another deficiency caused by the poor sonar coverage is that if an obstacle is too small (smaller than the width of the robot) when a robot approaches it head on, its sonar devices will not recognise it and the robot will crash. Due to these problems, both of the environments were designed to contain no concave obstacles (with small internal angles) and no obstacles smaller than the width of a robot.

The first environment, referred to as *world 1*, can be seen in figure 3, it contains a single bin, shown by the large square; three charging stations, represented by the circles; and twenty pieces of rubbish, depicted as very small squares, the robots are the squares located by the bin. The world was made deliberately challenging by placing the bin in the centre of the environment and surrounding it with obstacles.

The reason for placing the bin in a difficult position, was the expectation that to reach it, robots would fair better if they cooperated with each other, for example by signalling and flocking. A second world was used, but space restricts the inclusion of those results.

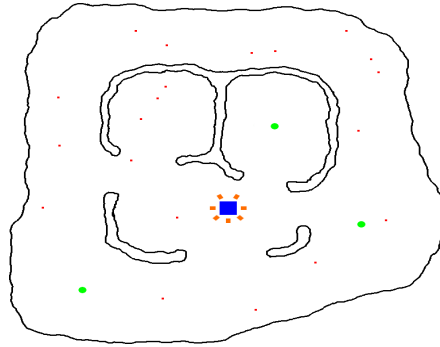


Fig. 3 World 1

4 Experiments

All experiments were carried out in the Player/Stage environment [4], running on Linux. We simulated Pioneer mobile robots, containing sonars, a camera, a gripper and a beacon. All code is available on request. The variant of foraging that was chosen for this project is known as rubbish or garbage collection. The task of rubbish collection used here involves a group of robots collecting pieces of rubbish that are randomly distributed throughout the environment and returning them to a bin. In order to make the task slightly more complex and to model the real world closer robots are required to monitor their power levels and when they are running low find a charging station at which to recharge.

4.1 Results for Neuro-endocrine Swarms

The success of the system is measured in terms of the *amount* of rubbish that was collected. Graphs are presented to show how the success of the system changed as the number of robots was varied. The total amount of rubbish collected by the group as a whole, as well as the number of pieces collected per robot are analysed.

4.2 Results

Figures 4(a) and 4(b) show the success of the robots after periods of 300 and 1200 seconds respectively. Each boxplot shows the results of ten different runs with ten different starting positions for the rubbish. Both the graphs show a strong positive correlation between the number of robots and the number of pieces of rubbish collected, until the case where five robots were used, at which the performance starts to level out and even drops in figure 4(b). The levelling out is expected in 4(b) since the maximum number of pieces that can be collected is twenty, but the fact that it is observed in 4(a) and that the performance drops in 4(b) indicates that interference starts to have an effect after five robots. The case with five robots also had the smallest interquartile range showing that five robots not only performed the best, but did so consistently.

The first outlier in 4(b), where the number of robots was three and the number of pieces picked up was six, was caused by one robot crashing, and the other robots crashing into the obstruction formed by the other robots, which emphasises the importance of redundancy in multi-robot systems. The outlier where the number of robots was five and the number of pieces collected was sixteen can be attributed, at least partly, to the simulator and the way the bin is represented. Since robots cannot see the inside of the bin from the outside, there is always the danger that collisions can occur as one robot travels into and one robot travels out of the bin, this is what happened in case of this outlier, two robots crashed whilst entering and leaving the bin which meant that when other robots came to drop rubbish there was a pileup effect. Only one other crash at the bin was observed in the seventy experiments of world 1, again for an experiment involving five robots however in this case it did not involve all of the robots and two were able to continue functioning, resulting in nineteen pieces being collected.

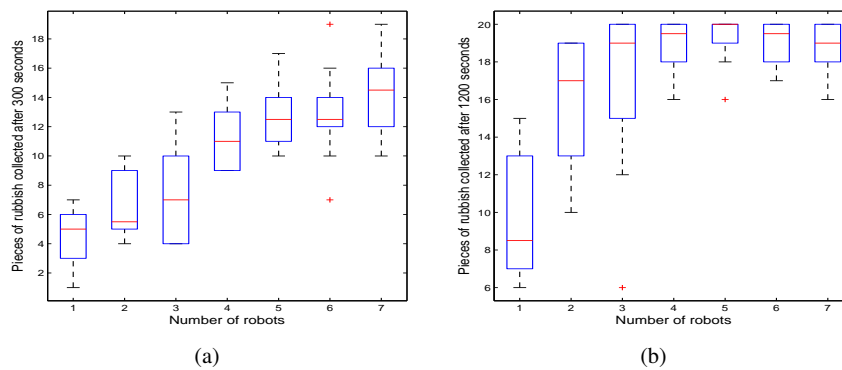


Fig. 4 Graphs showing the number of pieces of rubbish collected over periods of 300 (a) and 1200 (b) seconds, with varying numbers of robots between one and seven: World 1

Figure 5 shows the number of pieces of rubbish collected per robot after 300 and 1200 seconds, as to be expected, in both graphs the number of pieces drops as more robots are added. What is interesting about figure 5(b) is that the smallest interquartile range is observed for the case where there were five robots, showing that a group of five robots is most consistent on an individual level as well as a group level as indicated by figure 4(b). The outliers in figure 5(b) relate to the same runs as in figure 4(b).

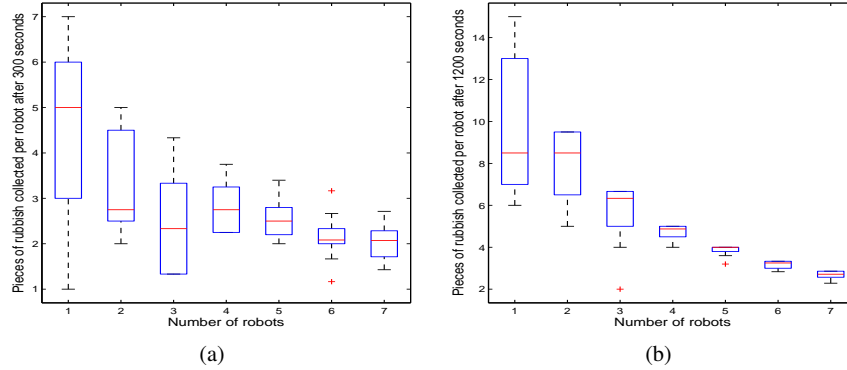


Fig. 5 Graphs showing the number of pieces of rubbish collected per robot over periods of 300 (a) and 1200 (b) seconds, with varying numbers of robots between one and seven: World 1

What is interesting to note from the observation of the experimental runs are the *emergence* of certain types of behaviour: specifically flocking of robots and dispersion of robots. Flocking emerges from the combination of obstacle avoidance, seek bin, signal bin, separation and cohesion and dispersion emerges from the combination of obstacle avoidance and separation, simply stated it is the spreading out of robots over the environment to ensure the greatest amount of coverage. Robots recharge when necessary, and collaborate together, through flocking etc. to remove as much garbage as possible from the environment. We have not undertaken a comparison between other approaches as yet, this would be outside the scope of a conference paper. However, we have investigated the efficiency and the effect of speed up on the swarm (how does adding more swarm members effect the overall performance), but have not room to report those results here. In summary, however, we have been able to show that there is an optimal number of robots for each world to achieve the best performance in garbage collection.

5 Conclusions

This work has adapted the neural-endocrine architecture for the development of swarm robotic systems. An architecture has been proposed for the task of foraging

and has been showed to allow for good collection of garbage over two basic environments. The work has also shown us that the simple neural-endocrine approach can easily be used for the development of such swarm systems. We observe that too many robots in the environment causes a potential problem (to be expected) for the optimal collection of garbage. The work presented in this paper is also the most complex task that the neural-endocrine approach has been used for to date. This gives us confidence in our approach and further work will investigate the actual role of each behaviour, and its importance to the overall performance of the system, and developing neural-endocrine systems on an ocean-going platform.

Acknowledgements This work is funded by EOARD, grant number FA-8655-07-3061

References

1. Simon Haykin. *Neural Networks - A Comprehensive Foundation*. Prentice-Hall, 1999.
2. Mark Neal and Jon Timmis. Timidity: A useful emotional mechanism for robot control? *Informatica*, 27(2):197–204, 2003.
3. Mark Neal and Jon Timmis. Once more unto the breach: Towards artificial homeostasis? In *Recent Developments in Biologically Inspired Computing*, pages 340–365. Idea Group, January 2005.
4. Player. The Player Project, 2009. Accessed online at 23-Apr 2009.
5. Erol Şahin and Alan Winfield. Special issues on swarm robotics. *Swarm Intelligence*, 2(2-4):69–72, 2008.
6. Jon Timmis, Mark Neal, and James Thorniley. An adaptive neuro-endocrine system for robotic systems. In *IEEE Workshop on Robotic Intelligence in Informationally Structured Space. Part of IEEE Workshops on Computational Intelligence*, pages 129–136, 2009.
7. Ali E. Turgut, Hande Çelikkanat, Fatih Gökçe, and Erol Şahin. Self-organized flocking in mobile robot swarms. *Swarm Intelligence*, 2(2-4):97–120, 2008.
8. Patrícia Vargas, Renan Moiolli, Leandro N. de Castro, Jon Timmis, Mark Neal, and Fernando J. Von Zuben. Artificial homeostatic system: A novel approach. In *Proceedings of the European Conference on Artificial Life*, volume 3630 of *LNAI*, pages 293–306. Springer, 2005.
9. Joanne Walker and Myra Wilson. A performance sensitive hormone-inspired system for task distribution amongst evolving robots. In *Proceedings of IEEE/RSJ 2008 International Conference on Intelligent Robots and Systems*, 2008.
10. Alan FT Winfield and Julien Nembrini. Safety in numbers: Fault tolerance in robot swarms. *International Journal of Modelling, Identification and Control*, 1(1):3–37, 2006.