

Kent Academic Repository

Full text document (pdf)

Citation for published version

Neal, Mark and Timmis, Jon (2005) Once More Unto the Breach: Towards Artificial Homeostasis?
In: de Castro, Leandro N. and von Zuben, F.J., eds. Recent Developments in Biologically Inspired Computing. Idea Group, pp. 340-365. ISBN 1-59140-313-8.

DOI

Link to record in KAR

<https://kar.kent.ac.uk/14382/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

ONCE MORE UNTO THE BREACH ...
TOWARDS ARTIFICIAL HOMEOSTASIS?

Mark Neal
Department of Computer Science
University of Wales, Aberystwyth. UK.
mjn@aber.ac.uk
Phone : +44 (0) 1970 622449

Jon Timmis
Computing Laboratory
University of Kent, Canterbury. UK
J.Timmis@kent.ac.uk
Phone: +44 (0) 1227 823636

ONCE MORE UNTO THE BREACH...
TOWARDS ARTIFICIAL HOMEOSTASIS?

ABSTRACT

The field of biologically inspired computing has generated many novel, interesting and useful computational systems. None of these systems alone is capable of approaching the level of behaviour for which the artificial intelligence and robotics communities strive. We suggest that it is now time to move on to integrating a number of these approaches in a biologically justifiable way. To this end we present a conceptual framework which integrates artificial neural networks, artificial immune systems and a novel artificial endocrine system. The natural counterparts of these three components are usually assumed to be the principal actors in maintaining homeostasis within biological systems. This chapter proposes a system, which promises to capitalise on the self-organising properties of these artificial systems to yield artificially homeostatic systems. The components develop in a common environment and interact in ways which draw heavily on their biological counterparts for inspiration. A case study is presented, in which aspects of the nervous and endocrine systems are exploited to create a simple robot controller. Mechanisms for the moderation of system growth using an artificial immune system are also presented.

KEYWORDS: Homeostasis, artificial immune systems, artificial neural networks, artificial endocrine systems, autonomous control, robotics

INTRODUCTION

The practice of drawing inspiration from biological systems for implementation in computing has a long and reasonably successful history. There remains however a wide gulf between the capabilities of computer systems and their biological counterparts. The variety of biological systems which have provided models is enormous as is the ingenuity of many implementations. Implementations range from hardware systems such as artificial retinas and neurons (Perrinet and Samuelides 2002), (Mead 1989) through software implementations of neural networks (Grossberg), (McClelland and Rumelhart 1986), genetic algorithms (Holland 1975), artificial immune systems (deCastro and Timmis 2002), cellular automata (Tomassini, Capcarrere et al. 1999) and a host of other techniques. Whilst each of these systems is undoubtedly extremely valuable in its own right, none has led to the type of behaviour which really warrants anything approaching the famous *Turing Test*, or is capable of life-like, long-term autonomous operation.

In this chapter we present a way forward which we believe represents an opportunity for biologically inspired computing in the current mould to break new ground in terms of generating complex, adaptive, autonomous and crucially: *self-organising* computational behaviour. We believe that all these properties are required for the implementation of systems capable of generating the type of behaviour sought by researchers in fields such as robotics, artificial intelligence and operating system design. With this in mind we wish to focus on one of the most impressive abilities of living organisms: their ability to ensure a reasonably stable internal state despite

wildly changing external environmental factors. This property, often termed homeostasis, is a major contributor to an organism's autonomy, and is the biological embodiment of the type of behaviour described above.

The investigation of animal behaviour by biologists has taken many forms (Aylett 1999), but the basic goal has been to understand the ways in which animals achieve this on-going autonomy of the individual. Of these approaches, the one which arguably most directly reflects the interest in on-going autonomy of a homeostatic nature is the dynamical systems approach. This considers the state of the individual in some state space which represents the state of the organism at any time. Homeostasis in such a state space is usually assumed to mean an orbit about some attractor which represents the “normal” condition for the organism. Clearly the presence of “attractive” values for particular variables will often lead to this type of cyclic path through state space, but should probably not be considered an immediate goal when constructing autonomous systems. This is for several reasons:

- The definition of such state spaces is fraught with problems such as: “What variables should be included?” and “How should behavioral attributes be represented?”
- The presence (or absence) of cyclical behaviour is often dependent on external factors such as the rising and setting of the sun. Thus, we need to define a set of circumstances under which a particular cyclic path will occur. Due to the unconstrained nature of the environment in which most workers wish their systems to operate, this is intractable;
- Once there are a significant number of interacting variables and control systems it is extremely hard to “design in” such cycles and to verify their presence and

robustness. Thus, systems designed in such a way tend to be of limited complexity and brittle.

We propose that in the spirit of biologically inspired computing we take one more lesson from biological systems. This lesson is that the existence of complex homeostatic systems is due to a series of selective pressures and to self-regulating growth and development of individual organisms. Thus we wish to create a system that is capable of developing robust homeostasis in a self-organising manner.

To achieve this we must consider the mechanisms usually associated with maintenance of homeostasis in organisms of a suitable level of complexity.

OVERVIEW OF THE APPROACH

As researchers we often (rather ambitiously!) wish to emulate the behaviour of animals such as sheep and dogs (Wilson and Neal 2001) and many other workers even more ambitiously wish to emulate the behaviour of humans (Fitzpatrick, Metta et al. 2003),(Adams, Breazeal et al. 2000), (Brooks, Breazeal et al. 1999)The mammalian body and its mechanisms for the maintenance of homeostasis are ripe for exploitation by computer scientists: nature itself can be perceived as performing computation and extracting metaphors from such systems has proved extremely useful in the past.

Selecting a level of granularity at which to represent the biological system is not a simple matter, and systems at both ends of the spectrum have been attempted: genetic algorithms model complete organisms as part of a population and systems such as artificial immune systems model interactions at a molecular and cellular level. In the case of homeostatic behaviour of organisms (which we believe to be a gross property

of organisms as a whole, rather than a reducible feature at a smaller scale) we have chosen an intermediate level of granularity for several reasons which will become apparent. There exists a body of literature which discusses the nature of homeostasis, autopoiesis, autonomy and life itself, which whilst interesting is largely of philosophical rather than practical use. Esoteric questions about the nature of boundaries will also for the moment be ignored and a pragmatic common-sense approach will be taken. At a suitable granularity which we believe to be of manageable complexity we can consider three major mechanisms for maintenance of homeostasis in mammals. These are the endocrine system, the immune system and the neural system. The current focus on identifying mechanisms appropriate for emulation is now we believe to some extent misguided as many of the artificial versions of various biological mechanisms perform very well (and to remarkably similar accuracies) in many situations. We believe that in order to move up to a different level of behaviour for our systems we must concentrate on the interactions between the various components that we have been broadly successful in reducing to their essences. We do not at this stage feel the need to introduce complex biologically plausible artificial neural networks, or very detailed models of the immune system, simple models which capture the essence of the systems will provide a suitable test-bed and sufficient complexity for a meaningful investigation. Thus we propose to consider the artificial counterparts of these three biological systems (neural, endocrine and immune systems) as a tightly coupled set of networks of interaction. Rather conveniently (and not coincidentally), useful models of two of these components are already well explored and sufficiently mature to consider as *off-the-shelf* components for our system: artificial neural networks and artificial immune systems. The third, the

artificial endocrine system, has been conceived as the final component for our system's development.

All animals have an immune system that is used to help keep them healthy. Artificial Immune Systems (AIS) have been researched for some years now (deCastro and Timmis 2002) and have been found to provide powerful and flexible tools for Computational Intelligence (CI) (deCastro and Timmis 2002), (Bradley and Tyrell 2002),(Forrest, Hofmeyr et al. 1997) to name a few. Initial work in AIS took inspiration from the basic functionality of the immune system (the ability to distinguish self cells from non-self cells) (Forrest, Perelson et al. 1994). However, the view that the immune system operates in such a manner is being questioned and alternative views are proposed: two such views are the one of self assertion (Varela, Coutinho et al. 1988),(Bersini 2002) and that of the danger theory (Matzinger 2002). With self assertion models, the system learns new senses of self over time, and will *assert* new cells to cover changes in the self. The danger model proposes that the immune system does not recognise self from non-self, but distinguishes danger from non-danger. Immunologists are debating this issue, and practitioners in the field of AIS are assessing the implications of self assertion models and danger models for artificial immune systems (Bersini 2002), (Aickelin and Cayzer 2002), (Secker, Freitas et al. 2003).

We propose a mechanism using a self assertion AIS approach, based on the immune network metaphor (Neal 2003). This can be used in conjunction with a novel artificial endocrine system (AES) to control the self-organising development of an ANN and AES by tracking the changes in self that occur during development of the system.

Whilst this approach, for the time-being, ignores the problem of using the AIS for identification of infective agents and their distinction from novel components of self, it is not seen as an insoluble problem and other work is tackling this problem of tracking changes in self whilst continuing to identify infective agents (Kim and Bentley 2002).

The use of artificial neural networks (ANNs) as a means of mapping inputs to outputs is very common in Computational Intelligence (CI) (Oyama, Chong et al. 2001). For example, in a mobile robot the ANN might connect sensors and actuators. Typically (after a learning period) the ANN defines a static, reactive response to a given input, but this is not the whole story in natural systems. In humans the endocrine system can affect the performance of the brain by means of various hormones, such as adrenaline (Besendovsky and del Ray 1996). Hormonal signals are also used as controls for many other processes that help to achieve homeostasis, so their use would appear to have significant potential for CI. We call such a system employing controlling hormones, an artificial endocrine system (AES).

Thus the development of the target system is a self-organising process that employs an artificial endocrine system and an artificial immune system to help shape and control an artificial neural network. The artificial neural network takes information from the perceptual channels and controls the state of any actuators that are present. The AIS and AES are essentially internal components, but in a similar way to its biological counterpart the AES can also respond to changes in perceptual state.

This chapter suggests a new method for combining ANN, AES and AIS in a single CI system, which develops an input-output model by means of positive and negative environmental cues. This chapter details the first steps to realising the proposed system. It first describes the background to each of these areas, explains how they might be combined, and suggests the properties that such a system might have and its advantages over existing systems. As an example, a case study is set out in which a robot has to achieve homeostasis, in a changing environment, whilst performing a simple object avoidance task and basking. It is envisaged that this will be expanded in the future to include tasks such as foraging or exploration to further test the ideas presented in this chapter. An embodied agent application is more demanding than a simulated agent and thus provides a more exacting test which is less prone to over-specificity and provides a richer problem space for the agent to explore.

We believe that any plausible test of the system will require an embodied agent in a non-trivial environment performing several tasks, the successful completion of at least one of which is required for survival.

There is a body of research that could be considered as potentially relevant to work in this chapter. To highlight one particular example, work in (Ogata and Sugano) attempts to create an emotional robot using “hormones”, but the mechanism differs dramatically from that proposed in this chapter. Interactive behaviour within humans is achieved by the robot through complex neural network architectures rather than hormone controlled responses. Work by (Gadanho and Canamero) also tackles some of the same problems as are presented here, but in very different ways. Work in the strong artificial life communities such as (Grand, Cliff et al. 1997), (Ray 1994) and (Terzopoulus 1994) attempts to create effects of emotions, and alludes to the idea of

hormone control, but again, no implementations like that proposed within this chapter were undertaken.

BIOLOGICAL MECHANISMS FOR HOMEOSTASIS

Homeostasis is the ability of an organism to achieve a steady state of internal body function in a varying environment (Besendovsky and del Ray 1996), (Varela 1981), (Vander, Sherman et al. 1990). This is achieved via complex interactions between a number of processes and systems within organisms. This section explores the biology behind the concept of homeostasis. In order to understand this process it is necessary to examine three of the major systems within organisms; the nervous system, the endocrine system and the immune system. By examining these systems and their interactions, it is possible to understand how organisms can achieve this state and therefore be exploited in the realm of CI. Work in (deCastro and Timmis 2002) describes at greater length the interactions of the biological systems, and the reader is directed to there for more detail.

The Nervous System

The nervous system (NS) is central to an organism's ability to process and act upon stimuli that it receives from an external source. Organisms ranging from slugs to humans are endowed with a nervous system which ranges in size, ability and function. This system will then develop over the lifetime of the organism, via processes such as growth, learning and memory (although not exclusively these).

Organisms are constantly being exposed to a vast number of stimuli, to which they must react. Stated simply, the NS takes sensory input and generates effector output. The sensory parts of the NS take input from vision, taste etc., which are stimuli for effector elements such as muscles. The NS consists of two types of cells: neurons and glial cells (we currently ignore the role of glial cells). Neurons are responsible for the firing of small electrical impulses in response to an input signal; neuroglia are cells which provide a type of support for neurons in the form of providing nutritional support, guiding development, the maintenance of the neuron environment and so on. Nearly half of the human nervous system is made from neurons, which are located in the brain. The basic components of a neuron can be considered as: (1) cell body and dendrites, (2) the axon and (3) axon terminals. Neurons will stimulate each other through the passing of an electrochemical signal from the axon of one neuron to the dendrite of another causing it to be stimulated. This in turn may cause the second neuron to send an activation (or possibly inhibitory) signal. The combined effect of the neurons in the network may ultimately stimulate an effector.

The Endocrine System

Within an organism, chemicals known as hormones implement a regulatory mechanism acting directly at an individual cell level. This system, the endocrine system, is responsible for the production and storage of these chemicals (Vander, Sherman et al. 1990). Hormones are also produced by neurons and immune cells such as T-cells, but for the current purposes these mechanisms will be ignored. These hormones have a great deal of influence over a large number of bodily functions and are key actors in the maintenance of homeostasis. Hormones have many functions

which affect behavior, assist growth, drive reproduction and so on. Typically, production of a hormone is in response to a change in state of the organism. Such changes are detected via the nervous system, immune system or by changes in other hormone or metabolite levels. Hormones are released into the blood or lymph system and are able to reach virtually all the tissues within the organism. It is quite possible (and normal) that there will be a number of different hormones present in the blood or lymph at any one time. However, not all cells will react to all hormones, as the response to hormones is highly specific: only certain cells are capable of responding to certain hormones. When a hormone locates its particular target cell, a binding takes place through specific receptors on the cells. Receptors on the target cell are usually located in one of two sites: within the cell nucleus (steroid hormone receptors) or in the plasma membrane (non-steroid hormone receptors, e.g., proteins, amines, and peptides). Non-steroid hormones decay and are ultimately removed from the organism at various rates. Built into the system is a mechanism by which hormones such as these will decay. This decay rate may well be a few minutes, but could potentially be a number of days. When a hormone binds with a receptor on the cell membrane, it stimulates internal signals to the appropriate sites within the cell, which in turn alter the cell's activity. Additionally, it can be noted that one hormone's presence can have an effect on another hormone.

Hormones are produced in various glands located around the body, including the ovaries, pancreas, adrenal glands, thymus, thyroid, hypothalamus and pituitary gland. Each of these glands produces hormones specific for certain tasks. For example, the hypothalamus produces hormones in response to signals from the nervous system and affects other glands within the body. The thymus is responsible for the development

and selection of T-cells within the immune system, hormones are produced here to stimulate the development of these T-cells and other lymph tissue. As can be seen, the endocrine system produces hormones that can affect body performance (such as steroids) but also hormones that interact with other body system to affect their performance.

The Immune System

The immune system is a remarkable, and complex, natural defence mechanism, which responds to foreign invaders called pathogens. Organisms typically have two lines of immunity, innate (inherited at birth) and adaptive (also known as acquired) which develops over the lifetime of the organism. However this is not the case for all organisms, such as the shark, which has a very powerful innate immune system and no acquired immune system. The innate immune system has first contact with any pathogenic substance and in a large amount of cases, this is all that is needed to remove the pathogenic material from the organism. However, there are many times the innate immune system is insufficient and cannot remove the infection. If this is the case, then the pathogen is passed over to the adaptive immune system.

The adaptive immune system primarily consists of B- and T-lymphocytes (cells). Through receptors on the cell, they are capable of binding with pathogenic material (antigens). Binding will occur between the receptors (paratopes) and antigen receptors (epitopes) if the affinity between the two is above a certain threshold. If a T-cell successfully binds an antigen this will cause the T-cell to stimulate B-cells through the emission of lymphokines. Additionally, B-cells can also bind with antigens, and

therefore a notion of antigenic affinity is created. The B-cells receive stimulation from this interaction with the antigen. Through the combination of these two interactions (antigens and T-cells) a B-cell then becomes stimulated and reaches a threshold at which it transforms into a *blast cell*. These blast cells then produce large amounts of clones (in proportion to antigenic affinity: the higher the affinity, the larger the number of clones produced) and also a large number of free antibodies, which undergo somatic hypermutation to increase the diversity of the immune response. This whole process is known as affinity maturation and is part of the *clonal selection theory* (Burnet 1959), which is the term used to identify the process described above. These antibodies (with the assistance of killer T-cells) will remove the antigen from the system. The immune system maintains an *immune memory* of cells, so that when exposed to the same (or slightly different) antigen, a quicker secondary response can be elicited which results in quicker removal of the infection.

The immune system remembers encounters with antigenic material (Tizzard 1988). There are a number of theories on how the immune system remembers encounters with antigenic material, with the most favoured view being that of clonal selection and memory cells (Burnet 1959). However, a theory first proposed in (Jerne 1974) suggested an idiotypic network and the immune network theory. Although not widely accepted, this theory is interesting especially for computer scientists and is the model which we choose to exploit. The idiotypic network was devised to explain the stimulation of B-cells in the absence of antigens. This is achieved by stimulation and suppression between cells via a network communicating via idiotypes on paratopes. The network acts as a self-organising and self-regulatory mechanism that captures antigenic information. Notable work in (Farmer, Packard et al. 1986) further explored

the immune network theory and created a simple model of the idiotypic network, which was further extended by (Perelson 1989). It can be noted that such a self-regulated system is akin to a homeostatic system, i.e. is capable of maintaining its own internal steady state.

Interactions between Biological Systems

So far, attention has been given to three systems within an organism: the nervous system, endocrine system and immune system. These systems do not act independently but as one large complex system.

First, immune, neural and endocrine cells can express receptors for each other. This allows interaction and communication between cells and molecules all three ways. Secondly, it appears that products from immune and neural systems can exist in lymphoid, endocrine and neural tissue at the same time. This indicates that there is a bi-directional link between the nervous system and immune system. Third, it would seem that both endocrine and neural systems can affect the immune system. There is evidence to suggest that by stimulating areas of the brain it is possible to affect certain immune responses, and also that stress (which is regulated by the endocrine system) can suppress immune responses: this is also reciprocal in that immune cells can affect endocrine and neural systems. The action of various endocrine products on the neural system is accepted to be an important stimulus of a wide variety of behaviours. These range from behaviours such as flight and sexual activity to sleeping and eating.

The primary function of the immune system is to defend the body against foreign invaders and malfunctioning cells. There are a wide variety of components that are used to achieve this, ranging from the bone marrow to lymph nodes. We currently propose to exploit the ability of the immune system to eliminate cells from the body. This will not exclude our exploitation of the other mechanisms that the immune system fulfils at a later stage in the development of the system. The immune system displays a number of interactions with other biological systems including the following: immune cell populations have receptor profiles for modulators such as neurotransmitters and endocrine hormones; and immune products also exist in neuroendocrine tissues (deCastro and Timmis 2002).

The nervous system's functions are the reception of stimuli, with the transmission of nerve impulses and activation of muscle (or effector) mechanisms. The nervous system has a number of interactions, which can be summarised as follows. Neural cells express receptors for cytokines, hormones and neuro-transmitters. The brain can stimulate defense mechanisms against infection, thus engaging the immune system. The hypothalamus within the brain, controls the pituitary and other endocrine glands and it is known that neural products coexist in immune and endocrine tissues (deCastro and Timmis 2002).

Finally, the endocrine system's function is to secrete hormones into the blood and other body fluids, with the aim being to regulate metabolism, growth etc. There are a large number of components that make up the system including glands such as the thyroid, pineal and the thymus. These glands are closely related to three fundamental activities in which we are interested: growth, release of hormones to the brain, and

immune system development. There are a number of interactions that the endocrine system is involved with: endocrine cells express receptors for cytokines, hormones, and neuro-transmitters; hormones provide feedback to the brain that affect neural processing; hormones including the reproductive hormones also affect the development of the nervous system. Again, endocrine products also exist in both immune and nervous tissue (deCastro and Timmis 2002). A good example of the close coupling between the neural and endocrine systems is the hypothalamus:

The hypothalamus is an excellent example of the interactions between these three major physiologic systems of the human body. Anatomically, the hypothalamus is part of the brain; it is located beneath the thalamus in the diencephalon. Signals from the limbic system are the primary neural trigger for the hypothalamus. Electrochemical signals from the hypothalamus trigger the auto-nomic nervous system as well as the pituitary. Nevertheless, the hypothalamus also produces a variety of hormones that are conveyed through a group of blood vessels to the pituitary, triggering the release or inhibition of the corresponding pituitary hormones. Furthermore the hypothalamus is an integral part of a series of feedback loops which not only regulate many systemic physiologic processes, but also adjust those processes to deal with environmental or internal changes and/or threats to the organism. As part of this feedback system, the hypothalamus senses the amount of certain hormones in the blood, the amount of neural stimulation in the limbic system, and the amount of certain thymic hormones. This information is then processed by the hypothalamus and adjustments in both neural and hormonal secretions are accomplished. The adjustment can be either to restore homeostasis or to move in either direction from it,

depending upon the result of the combined information processed by the hypothalamus. (deCastro and Timmis 2002) pp 176.

A FRAMEWORK FOR ARTIFICIAL HOMEOSTASIS

We now wish to propose a *framework* for creating systems based on analogues of the biological systems described above. The concept of a framework for CI is not new, and amongst many others work in (deCastro and Timmis 2002) argued that a framework would consist of:

- a representation of the components of the system;
- mechanisms by which to evaluate interactions of these components;
- procedures for adaptation.

Under such a conceptualisation, it is easier to discuss how such systems may be combined to form a more complex system. Table 1 captures the salient features of our proposed system in this context.

Framework Element	ANN	AES	AIS
Representation	Neuron	Endocrine gland	Lymphocyte
Evaluations	Activation function	Time delayed activation function	Affinity measures
Adaptation	Learning algorithms	Hormone structure update	Immune algorithms

Table 1- ANN, AES and AIS in a simple framework

Artificial Counterparts of the Biological Systems

Significant work has been done in extracting useful metaphors from the nervous system for the creation of artificial neural networks of a huge variety of types and functionalities (Haykin 1999). Work is now emerging in the field of artificial immune systems (deCastro and Timmis 2002), but little work has been done on artificial endocrine systems (AES). This section will discuss ANN and AES and postulate that through the combination of these approaches and an AIS it may be possible to create an artificial system capable of developing homeostatic behaviour in a manner analogous to biological systems.

Neural Networks

Artificial Neural Networks (ANN) are parallel distributed processing systems that are constructed via the connection of simple processing elements known as artificial neurons (McClelland and Rumelhart 1986). ANN have been applied to a vast array of problem areas such as machine vision e.g. (Sandini, Bosero et al. 1989) and robot control (Oyama, Chong et al. 2001). Figure 1 is a graphical depiction of a simple artificial neuron. In order to be of any practical use, individual neurons are connected together to form networks. These networks are *trained* in order to be able to perform an input-output mapping, of an input (x) through the constant adjusting of the weights (w_i) until the ANN produces some output (y). The weights can be adjusted via a number of *learning algorithms*, backpropagation being one of the most popular. An artificial neuron can be represented mathematically as:

$$u = \sum_{i=0}^{nx} w_i \cdot x_i$$

Equation 1

Where n is the number of weights and w_i is the weight for a given input x_i

Once the summing of the inputs has taken place, the neuron will fire, depending on the activation function $f(u)$, in this work a standard sigmoidal activation function has been employed,

$$f(u) = \frac{1}{1 + e^{-u}}$$

Equation 2

Initially we propose to use a relatively standard neural network system augmented with the AES and AIS components.




Figure 1 - A simple artificial neuron

The case study presented in this chapter shows how we propose to augment this basic artificial neuron, with interactions from an artificial endocrine system.

We expect to be able to use a standard learning rule (such as backpropagation) to perform weight-updates in the standard way, and to initially make no explicit allowance for interactions from the AES. This is not to say that they will not affect

how the weights alter during the learning process: clearly with an algorithm such as backpropagation any alteration in activity of neurons (such as those caused by AES interactions) will affect the weight update mechanism. We expect this to have significant effects at times, and also envisage that hormones affecting synaptic plasticity (and thus learning rate) will also be easy to incorporate. Clearly there are many complex interactions to be studied in this area once the initial mechanisms have been shown to be of value.

Artificial Endocrine Systems

Work in (Neal and Timmis 2003) proposes a new biologically inspired technique known as an Artificial Endocrine System (AES). The role of the AES is to provide a medium-term regulatory control mechanism for the behaviour of the system. The AES proposed consists of *gland cells* which produce and secrete *hormones* in response to external stimuli. The amount of hormone secreted is expressed as r_g for a gland g :

$$r_g = \alpha_g \sum_{i=0}^{nx} x_i$$

Equation 3

where α_g is the rate at which hormones are released for a particular gland g , x_i is the input to that gland and n is the number of inputs to that gland. This is a similar mechanism to that employed in neural network models.

The level of hormone is subject to geometric decay:

$$c(t+1)_g = c(t)_g \cdot \beta$$

Equation 4

where $c(t)_g$ is the hormone concentration at a time t for a gland g and β is the decay constant.

Membrane receptors located on artificial neurons and artificial immune system components are sensitive to hormones, thus providing a mechanism for the regulation of the ANN and AIS by the AES. Gland cells produce, secrete and record the concentration of hormones present in the system. Each gland cell secretes a specific hormone, represented by a simple string of bits. Within the integrated AES-ANN-AIS the hormone sensitive membranes of cells simply have a list of hormone receptors (again, represented as bit patterns) to which hormones are matched and a cell-specific action associated with each receptor. At present, perfect matches of hormone to receptor are considered (though this is not necessarily required: imperfect matches should generate lesser reactions). In the natural endocrine system, hormones are transported throughout the body: the same effect is achieved in the artificial endocrine system through the matching of each hormone secreted to the receptors on each cell's membrane in turn. A record of the current concentration of a hormone is maintained in the gland cell which secretes the hormone, and is then used to moderate the strength of reaction.

True to the analogy with the biological endocrine system, different cell types react to particular hormones, in different ways. The actions which are triggered in individual cells can vary according to four factors: the hormone which is detected, its concentration, the type of receiving cell and the individual cell's make-up. The former two of these factors are explained above, but the latter require further explanation. The type of cell receiving the hormone signal will clearly dictate what actions it is

capable of performing. For example, a neural cell may lower (or raise) its threshold value or increase (or decrease) its sensitivity to one or many of its inputs; a gland cell may increase (or decrease) secretion rate of a hormone; and an AIS cell (such as a B-cell) may increase (or decrease) its affinity threshold. The precise make-up of cells is fixed when they are added to the system. This may include variations in membrane characteristics (abilities to receive hormone signals), the effects that those signals have within the cell and other cell-type-specific characteristics such as connectivity pattern of a neuron etc. Cells are integrated into the neural, endocrine or immune system as appropriate and may then be culled by the immune system if the immune system is sufficiently stimulated to do so. This will depend on how different the cell is from those already present (which the AIS will tolerate) and on the state of the AES at the time. Of course the AIS will also be gradually extending the range of cells that it tolerates, and thus after a period of time new cells will be added to its repertoire and thus be safe from attack by the AIS.

In order AES-ANN interactions to be useful, the hormone levels affect the input weights in the ANN. Figure 2 provides a simple graphical representation of how this is achieved. Here the recorded hormone level affects each input weight on a particular neuron. It is easier to see this when these interactions are described mathematically as

$$u = \sum_{i=0}^{nx} w_i \cdot x_i \cdot \prod_{j=0}^{ng} C_j \cdot S_{ij} \cdot M_{ij}$$

Equation 5

where in this case x_i and w_i are the same as Equation 1 and ng is the number of glands in the system, C is the concentration of hormone, S is the sensitivity of the connection for receptor i to hormone j and M is the match between the receptor i and hormone j and is defined as:

$$M = \frac{1}{1 + dis(i, j)}$$

Equation 6

where *dis* is a distance measure function.




Figure 2 - The effect of endocrine interaction on the artificial neuron

Equation 1 and Equation 5 can now be compared. It should be noted that the new equation for the AES-ANN interaction is simple augmentation of the original equation, with the application of hormone levels applied to each input weight in the neuron. It should also be noted that this new AES augmented neuron bears a passing resemblance to the *Sigma Pi Neurons* (McClelland and Rumelhart 1986), however it is fundamentally different upon further examination.

Artificial Immune Systems

True to the analogy with the immune system, ultimately we propose to utilise the immune system metaphor in two ways. The first is for the AIS to act as a growth regulator for cells within the artificial system. Within the system, cells correspond to B-cells (AIS cells), neurons, endocrine glands and connections between these. As it is

proposed that this system will develop over time, some mechanism for controlling growth is required. The role of the AIS will be to remove cells and or connections that have a detrimental impact on the functioning of the system. The other role of the AIS will be to act in the more traditional manner of responding to pathogenic agents (which could be environmental changes, inputs from sensors, malfunctioning parts of the robot) and affect the ANN and AES accordingly. The AIS will be implemented using an immune network algorithm.

The immune network theory proposes that the B-cells in the body interact with each other to maintain the immune memory. The mechanism proposed is that B-cells which are capable of recognising similar (but not necessarily identical) patterns are also capable of recognising and stimulating each other (Jerne 1974). Thus a dynamic feedback mechanism can maintain parts of the immunological memory which are not frequently stimulated. Clearly however not all B-cells have sufficient stimulation to survive indefinitely and thus some will die out. In the human immune system T-cells both perform a surveillance role and interact with B-cells which complicates the mechanism somewhat. In our artificial immune system the role of T-cells is currently ignored. In the real immune system there are very large numbers of identical B-cells to deal with each type of infection. In an artificial system such repetition can be coded without representing all the identical cells individually. Fortunately the concept of a *recognition ball* which represents a region of antigen space that is covered by a particular type of B-cell can replace the repetition of individuals (Perelson 1989). So our AIS consists of a network of artificial recognition balls (ARB) which are linked together if they are close to each other in antigen space (see (Neal 2003) for a more detailed description of the algorithm). Items to be tested (such as cells recently

added to the system) can be considered to be points in this antigen space, and thus proximity can be defined as a simple distance function. The proximity required to evoke a reaction from an ARB is called its *affinity threshold* and determines whether an item is recognised by the AIS or not. Items falling outside this range evoke a primary immune response which attempts to extend the capability of the AIS to deal with the new item. If the hormones controlling the activity of the AIS (specifically the fitness hormone) indicate that the system is performing poorly, then the item will be removed rapidly. If however the system is performing well, then the killer immune response will be suppressed and the item (such as a beneficial new cell or connection) will survive and the repertoire of the AIS will be extended to cover it. This is effectively adding the new cell to the AIS definition of *self*.

PROPOSED MECHANISMS FOR ARTIFICIAL HOMEOSTASIS

The control mechanism proposed, unites in a single framework the techniques presented above. It is argued that a system that allows the low-level interaction of (computational) neural, immune and endocrine systems provides much scope for the construction and self-organising development of highly functional computational intelligences. The structure proposes to follow to some degree the biological realities of cell proliferation, differentiation and apoptosis. Cells that are added to the system are differentiated on a very restrictive epigenetic landscape. Components added will become one of: neurons, synapses, B cells, hormone producers or connections to hormone producers. These types have been chosen as they are usually seen as being fundamental to control of behaviour and maintenance of homeostasis in mammals.

The Grand View

The various components of the system will be implemented using relatively standard machine learning components for the neural and immune systems. This is not possible for the endocrine system as there is not an established literature concerning the implementation of artificial endocrine systems and therefore, this chapter has proposed such a mechanism that will be employed.

The functions performed by the natural endocrine system are both diverse and pervasive. The release of hormones into the bodily fluids and blood-stream permits their rapid transport throughout the body. An artificial endocrine system (AES) must be capable of such global activity. The natural endocrine system plays a pivotal role in initiation and regulation of a huge variety of homeostatic functions. The artificial system must be capable of such interactions. The innate endocrine components outlined in this chapter (control of growth, epigenesis and apoptosis) are only examples of what is a much more general mechanism. The growth of the artificial endocrine system is an on-going process which allows the development (via cell reproduction and mutation) of new hormones and new hormone release mechanisms. The hormones released by the AES will be globally applied to the cells in the system which react dependent upon their type and sensitivity. Novel reactions to hormones may be generated by the component cells of the different parts of the system (also via cell reproduction and mutation). The sensitivity to hormones of all cells in the system will be defined by a membrane definition associated with each cell. Membranes provide various binding sites for hormones generated and are capable of inhibition and excitation of the cell in question. Binding sites on the membranes are defined as

bit strings which are related to hormones in circulation by checking the hamming distance between the binding site pattern and the hormone molecule pattern.

Concentrations of hormones are increased by endocrine cells which release their particular hormone molecule into the system. The concentrations are decreased by the continuous breakdown of the molecules in circulation. This breakdown rate varies between hormones. The release of a particular hormone molecule by a particular gland cell may be triggered or suppressed by several mechanisms. These are:

- A particular pattern of activity on a perceptual channel (or channels)
- A particular pattern of activity in the artificial neural system
- A particular pattern of activity in the artificial immune system
- The presence of a hormone (or hormones) at a particular concentration

Thus a gland cell must be allowed to monitor any of several state indicators internal and external to the system. In order to achieve this it must be permitted to "connect" in various ways. These must include:

- Connections directly to perceptual channels
- Connections to the outputs of artificial neurons
- Connections to monitor the stimulation level of artificial immune system components

In addition to these connections, gland cells must have a membrane definition as described above. The combination of these sources of stimulation provides a suitable complexity of stimuli for the gland cells and provides scope for the evolution of

useful and interesting functionality in the AES. A large degree of commonality exists between the various cell types of the system. The functionality of the various cell types can be summarised as follows.

The addition of cells to the system is currently envisaged as stochastic in nature and occurs at a rate determined by a combination of the growth hormone and epigenetic hormone concentrations. At any given time the state of the system will be monitored by the AIS, and thus there will be knowledge of "self" encoded which represents the current state. Clearly when a cell is added to the system this state will alter and the AIS will identify the new neuron as "non-self". This will result in the AIS beginning the process of killing the cell at a particular rate. The rate of this process is however controlled (as are most other processes in the system) by the concentrations of various hormones. Of these hormones, the fitness hormone (the generation of which is inhibited and stimulated by negative and positive reinforcement stimuli respectively) is very important. If the fitness hormone is being secreted in large quantities then the time taken for the new cell to elicit a full reaction from the AIS is drastically increased. This allows sufficient time for the AIS to extend its definition of self to include the new neuron. As the processes of cloning and mutation are continuously ongoing in the AIS, this may occur and prevent the culling of the cell without any further intervention from other mechanisms. If however the fitness hormone concentration is low then the process will proceed apace, and the cell will be destroyed rapidly. Thus if, for example, a neuron is added which increases the fitness it is likely to survive due to its influence in increasing the fitness hormone concentration. Neurons which decrease the fitness will have the opposite effect and are thus more likely to be culled rapidly. Clearly this type of simple selection requires

a low rate of change of the system in order to ensure that the mechanisms have time to take effect and in order to ensure that changes caused by cells added at the same time do not always cancel each other out.

The system as a whole thus has the ability to add cells to each of its three internal components in ways suitable to each. The three components have the ability to interact in complex networks and to develop feedback control in an ongoing way under pressure of positive and negative reinforcement. Over time the system will add cells to its three components and we expect it to develop complex control strategies in a self-organising fashion. We hope that such controllers will display stability of external and internal behaviour that can be described as emergent homeostasis. Table 2 summarises the basic functionality of the three main cell types contained within the system.

Cell Type	Hormone Sensitive Membrane	Connectivity			Comments
		Perceptual Channels	Neural Outputs	Actuators	
Endocrine	Yes	Input	Input	None	Produce hormones to regulate internal processes
Neural	Yes	Input	Input	Output	Read perceptual channels and active actuators
Immune	Yes	None	None	None	Monitor internal state and cull cells that inhibit performance

Table 2- Summary of Cell Type Functionality

Components for Artificial Homeostasis

There are four innate components at the beginning of development of the system. These are engineered to ensure a suitable starting point for development to begin. There are three sets of cells that constitute embryonic starting points for the development of the three components described above, namely: the neural system, the immune system and the endocrine system. In addition to these there is the external boundary of the system which provides access to sensory apparatus and actuators. Each of these innate components will now be described, illustrating the necessary starting points required for each component and how each of these components will develop over time.

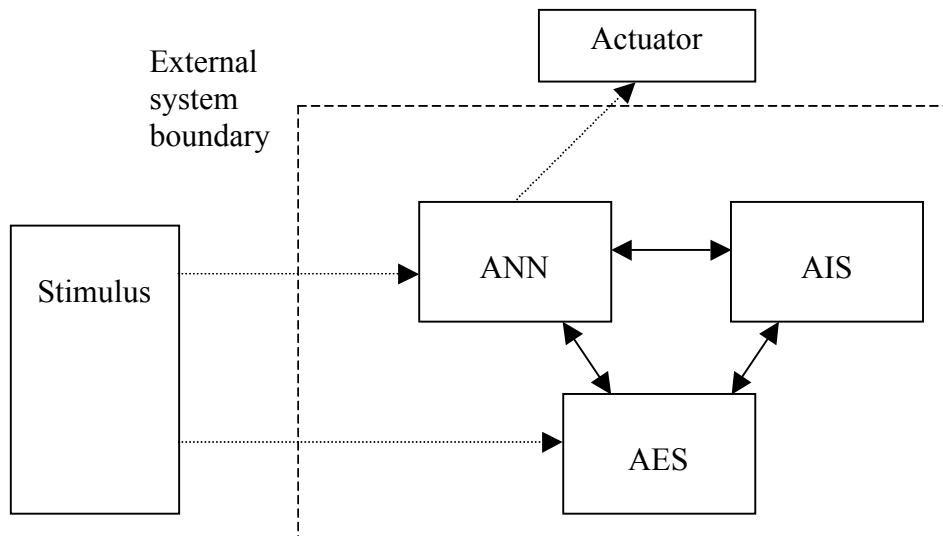


Figure 3- Overall System View of an Integrated Artificial Homeostatic Controller. All three components interact at various levels to achieve artificial homeostasis

The Innate Neural System and its Development

At the outset, a minimal neural network is constructed linking a number of the perceptual units and actuators which are available to the system. This minimal system may be of arbitrary design, but it is to be expected that the initial network will be hard-wired to generate basic reflexive innate behaviour, which can be then used as a bootstrap for development. The artificial neural system fulfils the same role as that played by the natural neural system in that it will connect to both the sensory apparatus and the actuators (effectors) of the system. This will in turn, allow reactive behaviour to be elicited. Cells created by the system, which differentiate to become neurons will be inserted into the network using a suitable insertion mechanism and will be evaluated with respect to the overall fitness of the system: thus allowing the ANN to "grow" over time. If there is a sustained drop in the overall system fitness after the addition of a particular neuron to the network then that node may be culled. This culling process will be triggered by one of the modes of operation of the AIS that is built into the system. Neurons that exhibit other behaviour, such as firing continuously or for a very large proportion of the time, may also be culled by the AIS as these will not be regarded as normal behaviour, as they could be potentially damaging to the system overall. Clearly neural cells may monitor perceptual channels and stimulate external activity directly. This neural activity may be moderated by hormones released by the artificial endocrine system, but only in a relatively homogeneous fashion. That is to say that functionally identical neurons will be affected identically by hormones at any one time. The neurons will be standard summation and squashing function units as are typically used in many ANN systems.

The Innate Artificial Endocrine System

It is proposed that the initial population of gland cells will be small and consist of cells that secrete a limited number of hormones that are key to the development and control of the initial cell population in the system. Three fundamental factors are required however. These are growth hormone, fitness hormone and an epigenetic control hormone. The growth hormone will be secreted by a gland cell which monitors the size of the current population of cells in the system and via a negative feedback mechanism (more cells implies less hormone release) achieve a steady state size of cell population. The fitness hormone will be secreted each time an improvement in performance is detected in order to suppress apoptosis of new cells which have recently been added to the system. Here, performance can be considered a measure of items such as battery consumption, distance covered, possible faults occurring in mechanical components and so on. We expect that this will encourage retention of novel cells which improve performance. The epigenetic control hormone will be released by a cell which monitors the relative sizes of the populations in each of the parts of the system, and via a negative feedback system maintain a dynamic equilibrium between the components. Clearly these hormones must act directly on cells which are added to the system and affect the cells which are produced. Cells in the artificial endocrine system can monitor the internal state of the system as well as the perceptual channels to regulate the release of their individual hormones. The hormones can only indirectly (by affecting the AIS and neural system) affect the external behaviour of the system.

The Innate Artificial Immune System

The initial immune system will be based on an idiotypic network model (Neal 2003), and will contain a small number of components which recognise the cells that are present within the innate parts of the other components of the system (the endocrine cells and the neural cells) as self. Thus any new cells generated that vary significantly in their properties from the initial sets may be recognised as non-self and as such will possibly elicit a destructive response from the AIS. In order to select those cells that are possibly of benefit to the overall performance of the system, when an improvement in performance is detected a suppression of the immune system's destructive power will be required accompanied by an increase in its ability to expand its definition of self, via the mechanisms of self assertion inherent within the immune network model. This suppression will be elicited by the release of a hormone from a cell in the innate artificial endocrine system which is sensitive to improvements in performance (as described above). In general, cells in the AIS will be sensitive to hormone concentration and will be capable of being suppressed or stimulated by them. The cells in the AIS will only be capable of detecting and destroying cells within the system that are persistently classified as non-self. This labelling is however continuously modified as the idiotypic network of the AIS evolves. The cells in the AIS will not directly affect the external behaviour of the system via its actuators, or monitor directly the perceptual channels. The AIS in this case interacts with both the neural and endocrine systems. With the former, removing redundant or useless neurons or promoting the inclusion of good neurons. With the latter, the AIS responds to the level of the fitness hormone in the system: for example, with higher levels of fitness hormones decreasing the culling mechanisms of the AIS.

External Boundary

The external boundary of the system provides an interface to the world that receives sensory input via perceptual channels which in the first instance will be considered to simply be streams of real-valued numbers. These perceptual channels can only be monitored by connections to cells in the neural system and endocrine system. In addition to the perceptual channels the boundary provides actuators which can be activated by neural outputs only.

THE CASE STUDY

Preliminary work to test the framework outlined above has begun, and is showing great promise for the future (see Neal and Mendao 2003 and Neal and Timmis 2003). Whilst this work does not demonstrate the techniques in their entirety it does demonstrate the efficacy of some of the mechanisms proposed. From initial implementations and some (prolonged) initial experimentation it seems that the mechanisms for control of neural networks via hormone suppression and excitation have effects within the bounds of our expectations.

In this case study we first present an outline of the controller, which has already been implemented and tested the innate neural and endocrine system, as discussed in the previous section. An outline of a more ambitious implementation extending this initial work is then proposed. All of this work was carried out using an ActivMedia 2DX mobile robot platform equipped with sixteen ultrasound range-sensors arranged

around its perimeter, a camera equipped with a panoramic mirror to provide a 360° field of view and an internal battery charge state sensor (although this has not yet been exploited). Whilst throughout this case study we will use this robot architecture we feel that there is nothing restricting the controller architecture to this particular robot or type of robot. Indeed we intend to extend this work using a variety of different platforms. Locomotion is achieved by driving the wheels independently at varying speeds and currently these motors are the only actuators available to the controller.

Initial work in (Neal, M. and J. Timmis (2003)) demonstrated that hormone mechanisms similar to those proposed here could lead to interesting and potentially useful medium-term (in the range of a few minutes to several hours) changes in the behaviour of robot controllers. Experiments described in (Neal and Mendao 2003) used a more general version of this artificial neuro-endocrine interaction mechanism to elicit cyclic behaviour using two independent neural networks and a pair of endocrine cells producing and releasing hormones to control them.

The controller contained two neural networks which were independently trained to perform two different tasks. The first was trained to move the robot to a position on a black placard placed on the floor, and the second was trained to move the robot to a position on top of a white placard placed nearby. The background around the placards was a grey carpet. The neural networks used a simple 16 input (plus one bias), 4 hidden, 2 output architecture, and took a simple greyscale camera image which was divided into 16 segments and averaged as input. The outputs were simply the motor speeds that were required in order to drive the robot toward the targets as required.

The training of the networks used the standard backpropagation weight update mechanism, and they were given a fixed concentration of the hormone to which they were sensitive. This ensured that the networks were effectively trained to carry out their tasks. These were tested independently and functioned well: as would normally be expected. This formed the neural part of the controller.

The endocrine part of the controller consisted of two endocrine cells which were sensitive to the maximum brightness present in the image. These cells produced hormone at a rate controlled proportionally by this intensity level, as well as having a “background” production rate which did not vary. Early versions of the endocrine elements released their hormone directly to the neural system, but later experimentation showed that in order to avoid heavy damping of the dynamics of the system a “store and release” mechanism was required. Whilst this is an important detail (and well worthy of further experimentation) we do not feel that such technical details are at stake here (see Neal and Mendao 2003). One of these cells had a negatively weighted input, and the other positive, thus one produced more hormone when the white placard was visible and the other produced more when the black placard was visible. The hormone which was produced when white was visible stimulated the network trained to seek the white placard, and the hormone produced when black was visible stimulated the network trained to seek the black placard. The endocrine cells were functionally very similar to the neural network elements, and it would certainly be possible to use weight update mechanisms such as backpropagation to train them to release hormone in the same way as neural networks. It is worth reiterating the point that although the endocrine cells look very much like neural network elements, they are different in that their effects are potentially global

within the system and that they have very different time dynamics (the hormone is produced, released and decayed in very different ways from the activation of neural network elements).

The controller constructed this way showed three notable effects:

- 1) Neural networks which are “over-driven” and “under-driven” through the endocrine mechanisms presented here still exhibit useful behaviour.
- 2) It is possible to use this endocrine mechanism to perform a type of sub-symbolic action selection.
- 3) The overall dynamics of such a system can be interesting: in this case we observed reliably periodic behaviour.

The first two of these effects came as no great surprise, as we had already seen hints of this in (Neal and Timmis 2003), but the latter was what we were hoping to achieve. The robot was seen to repeatedly move from the white placard to the black placard, and back again, pausing for a few minutes on each. In addition to the externally observable cyclic behaviour, the internal state of the endocrine system showed interesting periodic effects (see figure 5). Whilst this is in itself not terribly complex or difficult to achieve using standard symbolic approaches, we believe that the ability to combine behaviour from a number of neural networks in a way which is governed by the environment opens the way to much more interesting subsymbolic control mechanisms which are capable of developing themselves whilst responding in appropriate ways to their environment. Clearly this has not yet been shown, but we

believe that a mechanism such as the AIS mechanism proposed will provide sufficient self-organizational properties and stability to complete the picture presented here.

The next controller under consideration is intended to “roam” in a varying environment avoiding collisions with objects (including people and furniture), seek a charging station when required and periodically seek the white and black placards.

The design proposed contains a variety of elements which can be summarised as follows:

Black and white seeking ANN modules:

These will be trained and function exactly as described above.

An ANN *obstacle avoidance* module:

This module is a neural network with seventeen input nodes (one connected to each ultrasound range sensor and one bias input), two hidden nodes and two output nodes. All layers are fully connected. The two output nodes directly drive the wheel motors (one left and the other right). The weights in the network are set so as to generate turns away from any obstacle into unobstructed areas. This is relatively easy to achieve in a manner similar to Braitenberg’s vehicles (Braitenberg 1984).

An ANN *charge seeking* module:

This module consists of 17 input nodes (sixteen connected to the camera in the same way as described for the black seeking and white seeking above, and the other a bias node) connected to four hidden nodes, which in turn connect to the same two output

nodes as the obstacle avoidance neural network module. This network will be trained to move towards a bright red light mounted on top of the charging station.

Black seek promoting and *white seek promoting* endocrine cells:

These will function exactly as described above for the initial implementation.

A *caution promoting* endocrine cell:

This endocrine cell is connected directly to all of the ultrasound range sensors, and is stimulated to release hormone when they are registering obstacles close by. As suggested earlier this hormone release is proportionate to this stimulation (i.e. closer objects cause more hormone to be released). The hormone released has a stimulatory effect on all of the synapses in the “obstacle avoidance” neural network module, and thus leads to more sensitive and more rapid detection and avoidance of objects.

A *tiredness* endocrine cell:

This endocrine cell is connected only to the battery charge state sensor and is stimulated to release more hormone as the battery voltage decreases. This hormone has an inhibitory effect on the synapses connecting to the output nodes of the *obstacle avoidance* neural network module and an excitatory effect on the synapses in the *charge seeking module*. This will perform a similar action selection role as was observed for the black and white seeking modules described above.

These elements are intended to demonstrate a variety of features:

ANN interactions between *modules*:

Clearly there is nothing novel about interaction between neural network elements. However we feel that it is worth highlighting the highly “connectable” nature of ANNs and the fact that pre-trained (or manually set) networks can be combined intelligently. This implies that partially designed systems can be built (using combinations of pre-fabricated components) and used or subsequently trained further in slightly different configurations. Clearly such interactions would be capable of producing surprising results, but that does not preclude the successful use of either manual or automated methods for introducing links between previously trained neural systems, or indeed other components.

Endocrine mechanisms and selective sensitivity:

The use of several different endocrine cells which release different hormones highlights the fact that this mechanism provides immense flexibility for the suppression and excitation of large numbers of neurons without the use of overly connected neural network architectures or unpredictable recurrent architectures. The fact that the hormones affect subsets of the synapses in the neural system and that their effects can be combined further highlights this flexibility. This is most obvious when considering the output nodes which drive the motors. These node's synapses are stimulated by the *caution promoting* hormone, but are inhibited by the *tiredness* hormone. Clearly there will at times be conflicting signals, but this will often lead to a useful compromise as it will in this case. If the battery voltage is a little low and the robot is not currently being “threatened” by close by obstacles the motor speeds will

be reduced to save energy; however if close-by obstacles are encountered the level of *caution promoting* hormone will increase leading to more rapid motor movement to try to remove the threat. Clearly a trade-off between obstacle proximity and power usage will be made via this simple mechanism. The time period required for the decay of the hormone levels will help to eliminate problems with rapid oscillations in behaviour patterns which tend to occur in some other types of controller, but interestingly seem absent in most animal behaviour studies (McFarland 1999).

The potential for modularity:

Clearly the design employed here is modular in nature and we would argue that each of the components in this example is easily comprehensible. This goes some way to addressing the frequently highlighted problems with sub-symbolic systems (and especially ANNs): namely that they are opaque to examination, and that they cannot take advantage of previously existing knowledge. We would argue that these complaints are most obvious when dealing with large fully-connected networks and that the types of structure presented here will allow some degree of examination and certainly some “engineering-in” of domain specific knowledge. Even if subsequent automated development of the system leads to significant additions and alterations, it *is* possible to engineer-in some domain specific knowledge, and it is likely that structures that are engineered-in at the start will remain recognisable for a significant period.

The potential for partial design and (automated) refinement:

We propose that automated stochastic development could be used to develop the types of controller under consideration here. Possible mechanisms for this include:

connection of new inputs to neurons, connection of new inputs to endocrine cells, addition of neurons, addition of endocrine cells, addition of receptors for a particular hormone to a neuron or endocrine cell, and of course removal of any of these. It is easy to imagine both successful and unsuccessful changes of these types. We assume that we have some measure of success or failure in order to assess such changes. In this example we may wish to choose some function of the battery voltage and proximity to obstacles. This function could be used to ascertain if performance had improved or worsened after a particular change to the controller. An example of an advantageous change might be the addition of inhibitory receptors sensitive to the tiredness hormone to synapses in the black seeking and white seeking neural networks. This should reduce the likelihood of the robot being “distracted” by white or black seeking activity when the battery charge state is low. This may result in a higher battery charge state on average and thus positive reinforcement and retention of such an addition. An example of a disadvantageous change might be the addition of an excitatory neural connection from the input node in the charge seeking network directly to the output node responsible for driving the left wheel. This would result in a continual propensity to turn to the right which is likely to both reduce the effectiveness of the obstacle avoidance network and to consume more battery power in the process. It is easy to imagine mechanisms which allow the rolling back or retention of such changes dependent on performance.




Figure 4 - The Artificial Neural Network augmented with the Artificial Endocrine Gland. There are 16 sensor inputs, 2 hidden layer nodes and 2 output nodes. Each output node controls either the left motor or right motor control. The shaded area indicates that the hormonal gland influences the neurons in the network.

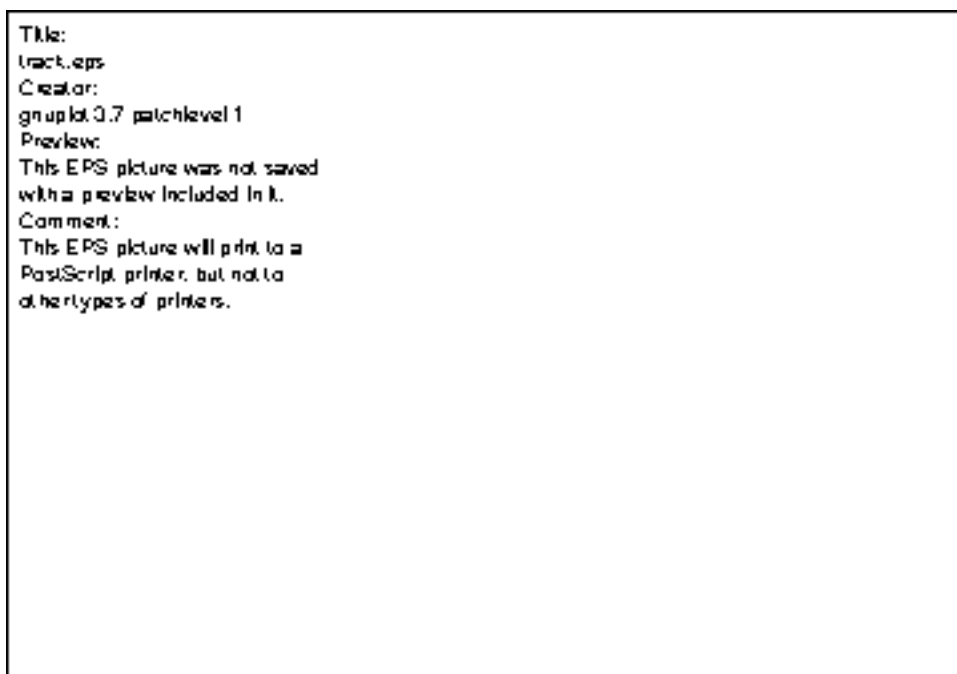


Figure 5 – Oscillations in behaviour observed in the black-seeking/white-seeking robot controller. The line shows the path followed by the robot over a period of 22 minutes.




Figure 6 – Oscillations in concentration of black-seek promoting hormone for the same experiment as presented in figure 5. The plot only shows the first half of the run for clarity. Note that the rate at which hormone is produced and stored (“Hormone held within endocrine cell”) varies, this is dependent on the perceptual state of the robot at the time. Also note that when hormone is released from the cell the free hormone concentration increases accordingly (“free concentration of hormone”).

CONCLUSIONS

This chapter has shown what we believe to be a viable route to implementing computational intelligence which surpasses current biologically inspired approaches. We believe that the current range of biologically inspired methods is approaching the limits of the capabilities of systems which will be of use as “one size fits all” stand-alone techniques. We envisage the future of biologically inspired computing to be in the integration of sets of these techniques into more structured and more complex architectures which have a better chance of achieving the levels of intelligence to which we aspire as robotics and AI researchers. Whilst the approach taken here represents a relatively naive view of the systems which are responsible for

homeostasis in mammals, we believe that such caricatures of biological systems have proved to be successful in the past in fields such as artificial neural networks, and are still showing benefits in emerging fields such as artificial immune systems. We believe that simplistic approaches such as that presented here should at least be eliminated before moving on to more complex, more biologically plausible techniques which may be required to fully realise our goal of artificial homeostasis.

Our system has taken a small subset of the currently viable biologically motivated computation approaches and using suitable biologically motivated interactions between them has set down a framework within which we believe artificial homeostasis at the level of complete organisms is achievable. Whilst this remains an open question we are working towards this goal, and are not yet prepared to abandon the biologically motivated approach to robotics and AI: we will attempt to imitate the action of the tiger for a little longer.

Once more unto the breach, dear friends, once more,

Or close the wall up with our English dead!

In peace there 's nothing so becomes a man

As modest stillness and humility;

But when the blast of war blows in our ears,

Then imitate the action of the tiger:

Stiffen the sinews, summon up the blood. (Shakespeare)

Adams, B., C. Breazeal, et al. (2000). "Humanoid Robots: A New Kind of Tool." IEEE Intelligent Systems.

Aickelin, U. and S. Cayzer (2002). The Danger Theory and Its Application to Artificial Immune Systems. International Conference on Artificial Immune Systems (ICARIS 2002), University of Kent, Canterbury, UK.

Aylett, R. (1999). Emotion in Behavioural Architectures. Workshop on Emotion-based Agents, Seattle, USA.

Bersini, H. (2002). Self-assertion versus Self-recognition: A Tribute to Francisco Varela. 1st International Conference on Artificial Immune Systems (ICARIS 2002), Canterbury, UK.

Besendorfsky, H. O. and A. del Ray (1996). "Immune-Neural-Endocrine Interactions: Facts and Hypotheses." Nature **249**: 356-358.

Bradley, D. and A. Tyrell (2002). A Hardware Immune System for Benchmark State Machine Error Detection. World Congress on Computational Intelligence (WCCI), HI. USA.

Braitenberg, V. (1984). Vehicles -- experiments in synthetic psychology. MIT Press.

Brooks, R., C. Breazeal, et al. (1999). The Cog Project: Building a Humanoid Robot. Lecture Notes in Computer Science, Springer. **1562**: 52-87.

Burnet, F. (1959). The Clonal Selection Theory of Acquired Immunity, Cambridge University Press.

deCastro, L. N. and J. Timmis (2002). An Artificial Immune Network for Multi-Modal Function Optimisation. World Congress on Computational Intelligence (WCCI), HI. ISA., IEEE.

deCastro, L. N. and J. Timmis (2002). Artificial Immune Systems: A New Computational Intelligence Approach, Springer.

Farmer, J. D., N. H. Packard, et al. (1986). "The immune System, Adaptation and machine learning." Physica **22D**: 187-204.

Fitzpatrick, P., G. Metta, et al. (2003). What am i? Initial Steps Towards Artificial Cognition. IEEE International Conference on Robotics and Automation.

Forrest, S., S. Hofmeyr, et al. (1997). "Computer Immunology." Communications of the ACM **40**(10): 88-96.

Forrest, S., A. S. Perelson, et al. (1994). Self-Nonself Discrimination in a Computer. IEEE Symposium on Research in Security and Privacy, Los Alamos, USA, IEEE Computer Society Press (1994).

Grand, S., D. Cliff, et al. (1997). Creatures: Artificial Life Autonomous Software Agents for Home Entertainment. ICAA.

Grossberg, S. "Nonlinear neural networks: Principles, mechanisms and architectures." Neural Networks **1**: 17-61.

Haykin, S. (1999). Neural Networks: A comprehensive foundation, Prentice Hall.

Holland, J. (1975). Adaptation in Natural and Artificial Systems, University of Michigan Press.

Jerne, N. K. (1974). "Towards a network theory of the immune system." Annals of Immunology **125**(C): 373-389.

Kim, J. and P. Bentley (2002). Immune Memory in the Dynamic Clonal Selection Algorithm. 1st International Conference on Artificial Immune Systems (CICARIS), Canterbury. UK., UKC Press.

Matzinger, P. (2002). "The Danger Model: A Renewed Sense of Self." Science **296**: 301-305.

McClelland, M. and D. E. Rumelhart, Eds. (1986). Parallel Distributed Processing, MIT Press.

McFarland, D. (1999). Animal Behavior, Pearson.

Mead, C. (1989). Analog VLSI and neural systems, Addison Wesley.

Neal, M. (2003). Meta Stable Memory in an Artificial Immune Network. 2nd International Conference on Artificial Immune Systems (ICARIS). J. Timmis and P. J. Bentley, Springer. **2787**: 168-180.

Neal, M. and J. Timmis (2003). "Timidity: A Useful Mechanism for Robot Control." Informatica: Special Issue on Perception and Emotion **7**: 197-203.

Ogata, T. and S. Sugano "Emotional Communication Robot: WAMOEBA-2R - Emotion Model and Evaluation Experiments."

Oyama, E., N. Y. Chong, et al. (2001). Inverse Kinematics Learning by Modular Architecture Neural Networks with Performance Prediction Networks. IEEE International Conference on Robotics and Automation, IEEE.

Perelson, A. S. (1989). "Immune Network Theory." Immunological Review **110**: 5-36.

Perrinet, L. and M. Samuelides (2002). Sparse image coding using an asynchronous spiking neural network. ESANN.

Ray, T. S. (1994). "An Evolutionary Approach to Synthetica Biology, Zen and the Art of Creating Life." Artificial Life **1**(1/2): 179-209.

Sandini, G., F. Boserio, et al. (1989). "The Use of Anthropomorphic Visual Sensor for Motion Estimation of Object Tracking." IEEE Transactions on Neural Networks **1**(1): 28-43.

Secker, A., A. A. Freitas, et al. (2003). A Danger Theory Inspired Approach to Web Mining. The Proceedings of Second International Conference on Artificial Immune Systems (ICARIS 2003). J. Timmis, P. Bentley and E. Hart, Springer. **2787**: 156-167.

Shakespeare, W. King Henry V. Act iii. Sc. 1.

Terzopoulos, P. e. a. (1994). "Artificial Fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world." Artificial Life **1**(4): 327-351.

Tizzard, I. (1988). The Response of B-cells to Antigen. Immunology: An Introduction, Saunders College Publishing: 199-223.

Tommassini, M., M. Capcarrere, et al. (1999). "A statistical study of a class of cellular evolutionary algorithms." Evolutionary Computation **7**: 255-274.

Vander, A. J., J. Sherman, et al. (1990). Human Physiology: The Mechanisms of Body Function, McGraw-Hill.

Varela, F. (1981). Autonomy and Autopoiesis. Self Organising Systems, New York Campus Press: 14-23.

Varela, F., A. Countinho, et al. (1988). "Cognitive Networks: Immune, Neural and Otherwise." Theoretical Immunology **2**: 359-375.

Wilson, M. and M. Neal (2001). "Diminishing returns of engineering effort in telerobotic systems." IEEE Transactions on Systems, Man and Cybernetics - Part A: Systems and Humans, Special Issue on Socially Intelligent Agents - The Human in the Loop **31**: 459-465.