Driven by Novelty? Integrating Executive Attention and Emotion for Autonomous Cognitive Development

Jason Garforth & Anthony Meehan

The Open University, Dept. of Computing, Walton Hall, MK7 6AA
jpg@janus.demon.co.uk; a.s.meehan@open.ac.uk

Abstract

In this paper, we describe a neural architecture and innate learning mechanism based upon executive attention that together provide a system capable of autonomous cognitive development expressed as behavioural adaptation. We suggest that such a mechanism arises from evolutionary pressure in favour of individuals who can reduce the burden on limited attentional capacity by learning new, unattended, reactive behaviours.

Functionally, the model is an elaboration of a neuropsychological model originally proposed by Norman & Shallice, and an extention to this model to include elements of emotion-based learning. This model includes a Supervisory Attentional System (SAS) capable of modulating purely reactive behaviour and generating novel behaviour. We show that deliberate attentional activity in the SAS provides the basis for a learning signal that leads to episodic and procedural learning.

A limited version of the model is used to control the behaviour of a simulated robot in which it is possible to observe both the behaviour of the machine in its environment and the neural activity levels of the SAS and other functional elements of the architecture. Such observations help asses the legitimacy of the model as hypothesised and as implemented.

1 Introduction

This paper builds upon previous work in which we have examined failures of attention in neurally controlled robots (Garforth et al. 2003; 2004). It presents a neural model for autonomous cognitive development based upon a 'learning signal' produced when executive attention is invoked in order to respond to new or unexpected events in the environment.

Learning the right thing to do at the right time is a matter of survival, especially for neonates. If we consider humans, and other 'higher' animals, we recognise that for the new-born there are many extrinsic signals available from the environment to reinforce the learning of what is, or is not, appropriate. It is clear that the neonate is capable of learning from the moment of birth, developing cognitive ability through interaction with its environment.

Neuroscience and neuropsychology have contributed enormously to our knowledge of the neural mechanisms underpinning the brain's structure and plasticity leading many to develop models which capture these mechanisms with varying degrees of neuroanatomical, neurophysiological or neuropsychological plausibility. Typically the models incorporate a dendro-axonal architecture, and a mechanism for modifying the synapses between neurons and/or the architecture itself[1].

Very recently, interest has developed in both attention (Taylor and Fragopanagos, 2004) and autonomous cognitive development (Weng et al 2000; Chen and Weng, 2004).

However, whilst there is knowledge about some of the architecture-level and control mechanisms that underpin learning, e.g., the role of the hippocampus in episodic learning, there is much less clarity as to the mechanism that might mediate experience of the world, translating experience into the invocation of learning. This paper outlines an architecture and a learning mechanism in which executive (deliberative) attention, plays a key role in promoting autonomous cognitive development.

The remainder of the paper is organised as follows. Section 2 seeks to establish a means by which attention-

[1] We use the terms mechanism very broadly and specifically include mechanisms that result in self-organisation, and genetic search based on 'natural' selection.

based learning might arise through evolutionary pressure. Section 3 outlines our architecture for attention-driven learning, and elaborates an innate learning mechanism that leads to autonomous cognitive development expressed as behavioural adaptation to novel events and problems. Section 4 illustrates the simulation used to investigate the model. Section 5 contains a discussion and account of current and future work.

2 Why Attention-based learning?

Executive attention is concerned with problem solving. It is invoked when there is a need to respond to a novel situation. The degree of novelty may be large or small; for the neonate everything is novel. Learning can be viewed as an adaptive strategy which diminishes attentional burden by progressively relegating that which is initially novel to the familiar or routine.

The benefits of this learning strategy flow from the well known constraints on attentional resources, which mean we are able to maintain some deliberate focus on about seven (plus or minus two) tasks or issues at any one time. In a dynamic and stochastic environment there is natural evolutionary pressure in favour of organisms that can distinguish between the novel and the familiar so that they can focus on the novel and treat the familiar in a routine, unattended, fashion.

In humans (and other higher animals exhibiting cognitive development and behavioural adaptation) the task of resolving what to do at any given time has two manifestations that are believed to be governed by largely distinct systems. Routine action selection involves the unattended, ‘automatic’ selection of appropriate behaviour through a tight coupling of perception to action, in a manner reproduced in many 'reactive' or 'situated' robots developed since the late 1980s. Non-routine action selection occurs in situations that require attentional resources and ‘willed’ behaviour. Non-routine action selection may be required in many circumstances, e.g. when executing a plan which requires significant variation in routine behaviour, or in trouble shooting (i.e. dealing with minor novelty in the environment), or when inhibition of a prepotent, strongly triggered, but unintended, response is appropriate (this includes suppression of reflex responses).

Norman and Shallice (1986) and Shallice (1988) have proposed a functional model for the control of both routine and non-routine behaviour. Non-routine behaviour is managed by a mechanism functionally labeled the Supervisory Attention System (SAS) associated with the prefrontal cortex (Shallice, 1988). Baddeley and Weiskrantz (1993) have proposed a broadly equivalent mechanism.

The architecture for autonomous cognitive development outlined in this paper is both an elaboration of Shallice's original SAS and an extension of that architecture to bring into scope mechanisms of emotion based learning associated with the limbic system. A distinctive element of our model is the use of attentional burden (SAS activity) as the basis of a 'learning signal' that leads to attended behaviour becoming routine, unattended behaviour.

Our architecture has been implemented as an integrated neural controller for a simulated robot, which facilitates experimentation and further refinement of the architecture and learning mechanism we propose.

3 Architecture of Attention-based Learning

Norman & Shallice's original functional architecture for executive control of behaviour comprises several subcomponents (Figure 1).

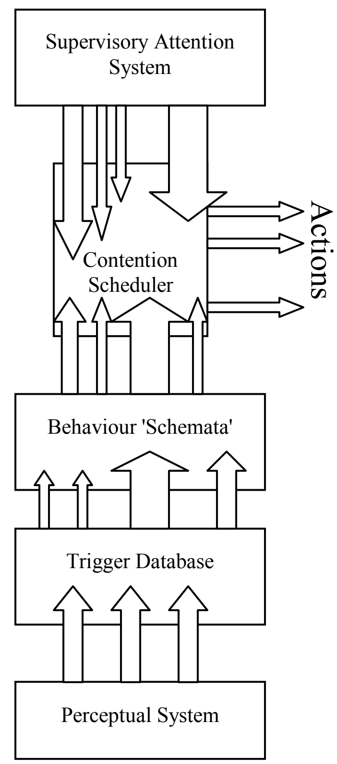

Figure 1. The Norman and Shallice model for willed and automatic control of behavior (after Shallice, 1988).

A perceptual subsystem, via an associative database, causes a range of behaviours (originally 'schemata') to be ‘triggered’ for possible expression. For each behaviour, the strength of the triggering depends upon the applicability of that behavior to the perceived state of the environment. The associative mapping takes account of the internal state of the agent and any goals that it has (as generated by cognitive subsystems). The CS resolves expression incompatible behaviours. A ‘willed’ action component is applied by a Supervisory Attention System (SAS) which modulates behaviour selection to correct errors and invoke actions to deal with novelty in the environment.

We have reinterpreted this functional architecture as a large scale, modular neural network. We have sought to maintain neuropsychological and neuroanatomical plausibility at the functional-structural level. Further, we have extended the model to include elements of the emotion (limbic) system associated with some forms of learning. Our model is illustrated in Figure 2.

The architecture is complex and the level of interconnection between functional modules makes it difficult to describe. Accordingly, only the connections that are central to our account of attention-based learning are featured. We outline the functional elements and operation in three steps:

- first we describe those elements which combine to provide for unattended, reactive behaviour in response to familiar (already learnt) situations;
- second we describe how executive attention directs action in unfamiliar or novel situations;
- finally we describe how this attentional mechanism provides the 'learning signal' that induces the reactive systems to learn the newly modified behaviour.

This sequence of elaboration provides a 'closed loop' in which there is no dependence, *ab initio*, on a body of learnt behaviors, other than the reflex behaviours present in any neonate.

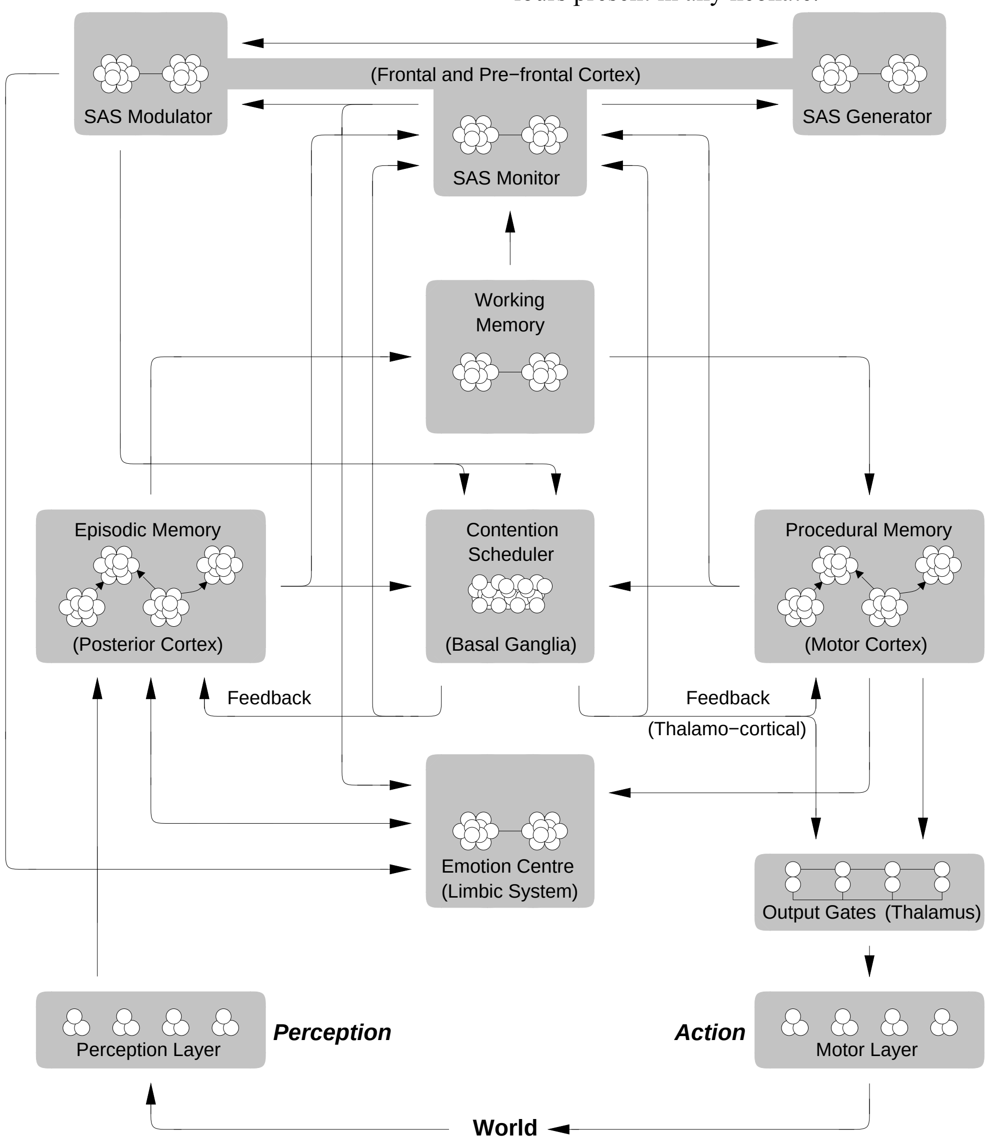

Figure 2. Architecture for autonomous cognitive development based upon the integration of systems for executive attention and emotion.

3.1 Unattended action

The following elements provide the basis for reactive, unattended behaviour:

Perception Layer: responsible for encoding low level perception mapping this into episodic memory.

Episodic memory (EM): EM (Baddeley et al., 2004) relates the perceived environment to previously learnt/familiar narrative structures or 'episodes'. These may be interpreted as possible plans or behavioural schemas (Schank, 1982). At an unattended level, EM seeks to express the salience of episodes. Learning in EM is known to be influenced by elements of the emotion (limbic) system, e.g. the hippocampus.

Procedural Memory (PM): encodes procedural motor skills. EM and PM are heavily interconnected (not shown), providing the connectivity for the tight coupling of perception to action. PM includes central pattern generators and motor pattern generators for "binding" motor action to perception.

Contention Scheduler (CS): unattended selection of routine actions results in a number of contending actions of varying salience seeking expression at the effector level. The primary function of the CS is to select which behaviors are granted expression (hence becoming observable behaviour). (It is interesting to note that the CS architecture allows more than one behaviour to be expressed as long as there is no contradiction at the effector level.)

The CS output provides positive feedback to the PM (Houston and Sumida, 1985) to reinforce the persistence (McFarland, 1989), or perseverance (Shallice 1988), of the current behaviour so that minor fluctuations in perception do not result in rapid behaviour switching.

Output Gates: behaviours selected for expression by the CS are 'gated' to effector systems (in the thalamus). By default, the gating system *inhibits* the expression of highly salient behaviours (from PM). The CS *disinhibits* those behaviours selected for expression. (It is interesting to note that this mechanism enables an organism to 'always do something'; evolution seems to have selected for organisms for which 'doing nothing is not an option').

3.2 Attended action

The following elements facilitate attended, or deliberate action.

Working Memory (WM): provides the attention system with access to current goals and intentions and salient biographical 'episodes' from EM. (Connections from EM and PM converge in pre-frontal cortex (Fuster, 1995), and this may also provide the attentional system with a means to combine episodes and representations of motor functions which can be 'played forward', in the absence of behavioural expression, thus providing a means to predict or 'look ahead'.

Supervisory Attention System (SAS): In order to correct errors and determine non-routine courses of action, a supervisory system requires a number of distinct sub-functions: Shallice (1988) distinguishes some of them as follows:

SAS Monitor: The SAS monitor can be thought of as a novelty detection system. The SAS Monitor is connected into WM and detects novelty, defined as departure from expectation in relation to the perceived world (from EM), intended action (WM), expressed action (PM and CS), and the outcome of expressed action perceived through the changed state of the world (EM and WM). The monitor may be thought of as an 'arousal mechanism' that triggers the activation of the other attentional SAS sub-units.

SAS Modulator: when a novel or unexpected situation arises, the SAS must provide a modulatory signal that both attenuates the salience (strength of triggering) of inappropriate actions and potentiates the salience of alternative, attended tasks. Shallice suggests three possible modulatory responses:

- attenuate the currently expressed behaviour for a given time and potentiate an intended behaviour;
- attenuate the active behaviour for a given time and potentiate some default, 'try something', response;
- attenuate *all* intended behaviours for a given time, allowing the expression of a purely reactive behaviour governed by perception of the environment ('try anything').

SAS Generator: if modulatory responses described above prove inadequate to the novelty of a situation, the SAS Generator must be able to produce novel strategies for solving new problems. Such strategies might include altering a sequence of actions in a plan or creating novel sequences of intended actions (plans).

3.3 Attention-based Learning

Earlier in the paper, we argued that the need for learning arose from the evolutionary advantage to be gained from releasing attentional resources from dealing with the familiar in order for them to be available to deal with the demands of the unfamiliar. Thus, we hypothesise that the level of activity in the SAS (part of prefrontal cortex) is the basis of a reinforcing signal that promotes the learning of currently attended perception and action so that it becomes increasingly unattended perception and action.

It is known that the emotion system (limbic system, and especially the hippocampus) plays a role in promoting episodic learning. The limbic system is tightly connected to the prefrontal cortex, providing a material basis for our hypothesis.

The learning of new procedural skills does not directly depend upon the hippocampus. Thus we need a second mechanism to support this form of learning. We have already seen that the SAS modulates the expression of behaviour by potentiating and/or attenuating the salience of behaviours contending for expression by the CS. This, in itself does not appear to constitute a learning signal. However, once behaviour is expressed by the CS, reinforcing (thalamo-cortical) feedback to the PM seeks to promote the persistence of the expression of this behaviour. (This 'self-priming' mechanism serves to avoid constant switching of behaviour arising from very minor change in the perceived environment.) This feedback mechanism is a plausible basis for a reinforcement signal which results in the learning of the attended response.

4 Implementation

We have begun to implement the architecture described above as a large-scale, modular neural network controlling a simulated robot. This allows us to illustrate the operation of the architecture and explore its properties.

4.1 The Robot

The robot has two, forward facing sonar sensors and eight olfactory sensors that allow it to sense the presence of obstacles or objects of interest such as food, nesting materials and other robots. Its effectors are two independent drive wheels and a gripper for picking up objects of interest. The dynamics of the robot motion and the sensor behaviors are modelled on the techniques prescribed in Dudek and Jenkin (2000).

4.2 The Attention-based Neural Controller

The modular structure of the network corresponds to the functional structure of Figure 2. The large-scale modules group clusters of highly interconnected neurons, most of which comprise four or eight input recurrent (Elman or Jordan) networks with up to three hidden layers. The Perception Layer processes and fuses sensor signals to produce a representation of the environment to an associative layer (EM) which maps the perceived state of the world to behaviours in the PM.

PM groups neural clusters which exhibit a small number of basis behaviours (c.f. Mataric, 1996). Basis behaviours are low-level behaviours that may be combined to provide higher-level behaviours. The basis behaviours, and higher-level behaviours arising from them, serve the same role as "schemata" in the Norman & Shallice model. Feedback from PM to EM enables the behaviours to provide excitation for the 'priming' of other relevant associations.

The procedural behaviour clusters in PM are layered in a tree like structure; clusters at the bottom of the tree correspond to primitive actions, and those further up represent either composite (complementary or parallel) or sequenced (conflicting and therefore sequential) behaviours. Exciting a composite or sequenced behaviour cluster causes that behaviour to excite (in parallel or in sequence, respectively) clusters representing each sub-behaviour. Every behaviour cluster has inputs from EM, other behaviour clusters in PM, the CS and SAS (see below). The strength of the output to the CS represents a 'request' for expression of the behaviour at the robot effectors (wheels, gripper, etc.).

The CS is based on the computational properties of the basal ganglia (Alexander, 1995; Houk et al., 1995) and is an independent implementation of the CS described by Prescott et al. (1999). The CS *disinhibits* active behaviours which are otherwise inhibited by the effector gateway.

We have already mentioned that the CS provides feedback to PM to reinforce the perseverance of currently expressed behaviour(s).

Although the SAS has several functions, including the generation of novel behaviours, only two functions are implemented currently; these are Monitor and Modulate (as described above). (The full SAS has a Generate function to create novel plans. We have not yet implemented this function.)

Currently, we represent simulate the 'result' of dynamic planning as an encoding of sequences of intended behaviours held in working memory (WM). As the excited behaviour is expressed (via the CS) the WM sequence primes the next behaviour in the sequence so that it will be more readily triggered when (if) the prerequisite change in the environment occurs through expression of the current behaviour.

The SAS Monitor network clusters have three inputs: the environmentally induced behavioural associations from EM, the currently intended behaviours (from WM), and the behaviours expressed by the CS. If the currently expressed behaviour is not strongly triggered, or if it is not intended (or both), the Monitor generates an 'arousal' stimulus to the SAS Modulator.

The Modulation clusters generate outputs that modulates the signals from PM into the CS so that intended behaviour is potentiated and other behaviours are attenuated. (It is important to recognise that this does not *guarantee* the selection of the intended behaviour, as this risks overriding behaviours strongly and appropriately triggered by the environment, e.g., those designed to prevent collisions.)

The SAS Modulator provides an output signal, which reflects its level of activity, and this provides the basis for a signal to the Emotion Centre. The emotion centre presents this signal to EM as a 'learning stimulus'.

5 Illustration

This section illustrates the use of the simulated robot to observe the operation of the control network. We take as an example the need for an 'infant' robot to learn to avoid highly tempting, prepotent stimuli in order to complete an intended task. Before learning this response, the infant is readily distracted by a strong stimulus. In order to avoid distraction and complete its original goal the SAS is required to suppress the distracted behaviour. In doing so, the SAS activity provides the learning signal which reinforces the attended behaviour ('avoid the distraction') so that this becomes routine. Thereafter, introduction of the distracting stimulus has no significant effect and the SAS remains quiescent.

The scenario (Figure 3a) involves an infant robot (Penny) collecting 'food' (a stimulus to which it responds enthusiastically) with a goal of and taking it 'home'. When the distraction ('more food') is introduced, the robot orients towards the stimulus even though it has not taken the food it already has home.

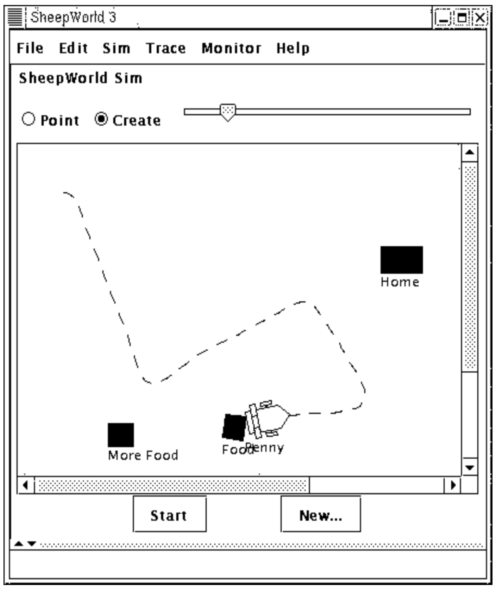

Figure 3a. Behaviour in an 'infant' robot that has yet to learn how to avoid inappropriate distraction.

We can observe the output histories (traces) of selected network clusters during the experiment (Figure 3b) and thus observe operation of the SAS. The distracting food source is introduced (2 seconds into this trace) and before the robot has dropped the food it already has at home. Food is detected (trace 1) and the behaviour 'orient to food' is triggered (trace 2). This is not an expected event in the plan to take the food it already has its to home (permanently low intention for this behaviour in trace 3). The strength of the orient to food response leads the CS to select (inappropriately) the 'orient to food' behaviour (rising spike of trace 4). The SAS monitor detects this unexpected response (trace 5), and produces a modulatory signal to suppress this behaviour as seen by the CS (trace 6). This results in the falling spike of trace 4. (Trace 7 illustrates a momentary (<0.1s.) expression of the inappropriate behaviour at one of the motors.)

The SAS activity provides the reinforcing learning signal which leads EM to acquire a new 'episode' in which presence of the distraction becomes 'familiar' and the mapping to PM is modified so that the appropriate behaviour (continue towards home in the presence of the distraction) is made routine (unattended).

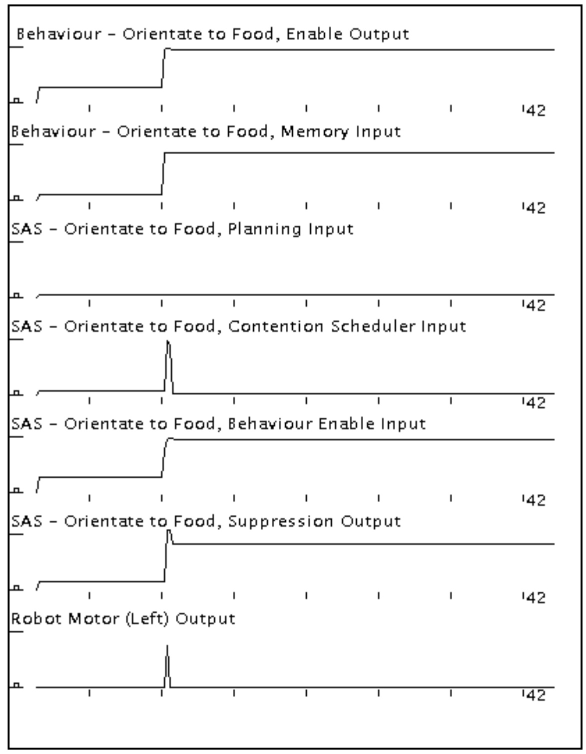

Figure 3b. Suppression of distracted behaviour by the SAS.

5 Discussion

In the work described in this paper a functional model of executive attention was used as the basis of an implementation of a modular neural control architecture capable of learning new unattended behaviour on the basis of an attention-based learning signal. The paper illustrated how the integrated architecture is capable of inducing learning in episodic memory and procedural memory so that initially novel perceptual patterns become more familiar, and the associated actions become capable of unattended expression. It was suggested that such a mechanism might arise from evolutionary pressure in favour of reducing the burden on limited attentional resources.

In the context of autonomous cognitive development, it is important to recognise that the adaptation arises from an innate mechanism (there is no 'teacher') and that the mechanism assumes only initial (basis) behaviours at 'birth'. We might note in passing that there is no 'store' of 'world knowledge'; the world is its own model (c.f. Chen and Weng, 2004).

It is possible to see in this model the basis upon which neophilic behaviour (characterised by traits of 'adventurousness' and 'curiosity') might arise. As each initially novel circumstance is encountered it demands limited attentional resources. The application of attention (in our case, to action) induces learning that needs reduced, eventually zero, attentional resource. The freed resource can then be applied to the next novel circumstance.

Further, the inclusion of the emotion (limbic) system would seem to link learning with a sense of pleasure. If this system were to be extended to include elements of the endocrine system (which is closely associated with limbic structures) then it may even be possible to establish an addictive underpinning of learning, even under stress. Such an organism would be 'driven by novelty'.

Future work aims to extend the SAS to include the SAS generator so that significant new behaviours, as opposed to modified existing behaviours, can be learnt.

References

Alexander, G. A., (1995) Basal Ganglia. In Arbib M. A, The Handbook of Brain Theory and Neural Networks, MIT Press 1995, 139-144.

Baddeley, A. and Weiskrantz, L. (1993) Attention: Selection, Awareness, & Control. Oxford University Press, 1993, 152-170.

Baddeley, Conway and Aggleton (2002) Episodic Memory. Oxford University Press 2002.

Chen, Y. and Weng, J. (2004) A Case Study of Developmental Robotics in Understanding "Object Permanence". In L.S. Smith, A. Hussein and I. Aleksander (eds.) Brian Inspired Cognitive Systems (BICS) 2004, 29 August - 1 September 2004, University of Stirling, Scotland, 2004, CD-ROM, BIS 1.7 1-7

Dudek, G. and Jenkin, M. (2000) Computational Principles of Mobile Robotics. Cambridge University Press, 2000.

Fuster, J.M. (1995) Memory in the Cerebral Cortex. MIT Press, Cambridge MA, 1995.

Garforth, J, McHale, S, Meehan, A, (2003) Neural Executive Attentional Control in Robots. International Joint Conference on Artificial Intelligence (IJCAI-03), Acapulco, Mexico, August 2003.

Garforth, J, McHale, S, Meehan, A, (2004) Executive Attention and Action Selection in a Neurally Controlled Simulated Robot. In L.S. Smith, A. Hussein and I. Aleksander (eds.) Brian Inspired Cognitive Systems (BICS) 2004, 29 August - 1 September 2004, University of Stirling, Scotland, CD-ROM, BIS 1.6 1-7

Houk, J., Davis J., Beiser, D. (1995) Models of information processing in the basal ganglia. MIT Press, Cambridge MA, 1995.

Houston, A. and Sumida, B. (1985) A positive feedback model for swithching between two activities. Animal Behaviour 33(Feb): 315-325.

Mataric, M. J. (1996) Designing and Understanding Adaptive Group Behaviour. Adaptive Behaviour, 4:1, 51-80.

McFarland, D. (1989) Problems of Animal Behaviour. Harlow, UK, Longman, 1989.

Norman, D.A. and Shallice, T. (1986) Attention to action: willed and automatic control of behaviour. In R. Davidson, G. Schwartz and D. Shapiro (eds), Conciousness and Self Regulation: Advances in Theory and Research, 4, Academic Press, 1986, pages 515-549.

Prescott T.J., Redgrave, P., Gurney, K. (1999) Layered control architectures in robots and vertebrates. Adaptive Behaviour 1999, 7(1) 99-127.

Schank, R. (1982) Dynamic Memory. Cambridge: Cambridge University Press, 1982.

Shallice, T. (1988) From Neuropsychology to Mental Structure. Cambridge University Press, 1988.

Taylor, J.G. and Fragopanagos, N. (2004) Modelling Interaction of Attention and Emotion. In L.S. Smith, A. Hussein and I. Aleksander (eds.) Brian Inspired Cognitive Systems (BICS) 2004, 29 August - 1 September 2004, University of Stirling, Scotland, CD-ROM, CNS 1.2 1-4

Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur M., and Thelen, E. (2000) Autonomous Mental Development by Robots and Animals, Science, vol. 291, no. 5504, pp. 599 - 600.