Executive Attention and Action Selection in a Neurally Controlled Simulated Robot

Jason Garforth[1], Sue McHale[2] and Anthony Meehan[1]

1. The Open University, Dept. of Computing, Walton Hall, MK7 6AA
2. Sheffield Hallam University, Dept. of Psychology, Pond Street, Sheffield, S1 1WB

jpg@janus.demon.co.uk; s.l.mchale@shu.ac.uk; a.s.meehan@open.ac.uk

Abstract

In this paper, we describe an integrated model for neural control of both routine and non-routine action selection. Functionally, the model is based upon an architecture originally proposed by Norman & Shallice (Norman and Shallice, 1986; Shallice 1988). This model includes a Contention Scheduler (CS) and a Supervisory Attentional System (SAS). The CS mechanism is based upon the CS architecture described by Presott, Redgrave and Gurney (1999). Little is known of the neural architecture of the prefrontal cortex (PFC) that might realise the SAS. Accordingly, we have developed a partial implementation in which the architecture is guided by an analysis of the functionality required by Shallice's SAS. The resulting model is used to control the behaviour of a simulated robot. This paper extends earlier work on normal executive control (Garforth, McHale, Meehan, 2003), describing lesion studies to explore the correspondence between the behaviour of the robot and behaviours seen in humans suffering lesions of the PFC. In doing so, we are able to assess the legitimacy of the SAS model as hypothesised and as implemented.

1 Introduction

When exhibiting high-level behaviours that require the sequencing of lower-level behaviours, performing the right action at the right time is important. Action selection has two manifestations that are believed to be governed by distinct systems. Routine action selection involves the unattended, 'automatic' selection of appropriate behaviour. Non-routine action selection occurs in situations that require 'willed' behaviour that require attentional resources. Non-routine action selection may be required in many circumstances, e.g. when executing a plan which requires significant variation in routine behaviour, trouble-shooting (dealing with minor novelty in the environment), or inhibition of a strongly triggered, but unintended, response.

Norman and Shallice (1986) and Shallice (1988) have proposed a functional model for the control of both routine and non-routine behaviour. Routine behaviour is believed to be managed by a contention scheduling mechanism. Non-routine behaviour is managed by a mechanism functionally labeled the Supervisory Attention System (Shallice, 1988), or the Supervisory System (Baddeley and Weiskrantz 1993), and is associated with an area of the prefrontal cortex (PFC).

For researchers within the connectionist/PDP paradigm exploring routine action selection, the focus of recent work has been on modelling neural structures for contention scheduling, i.e. structures capable of automatic selection of (sequences of) actions (Cooper and Shallice, 1997; 2000; Gilbert and Shallice, 2001). For some workers, the hypothesised association of the contention scheduling mechanism with the basal ganglia (Shallice, 1988) has led them to develop models guided/constrained by knowledge of its neural architecture (Houk, Davis and Beiser, 1995; Prescott, Redgrave and Gurney, 1999; Gurney, Prescott and Redgrave, 2001).

The problem of non-routine action selection has been explored to a more limited extent. Cohen, Dunbar and McClelland (1990) have developed a network in which task selection (between word reading and printed word colour naming) was governed by a 'context layer' that provided the supervisory input that determined which of the two tasks was performed. An exogenous control input is applied to determine which task is required. This system was later modified to perform other task

switching exercises (Cohen, Braver and O'Reilly, 1998).

Cooper (2003) has augmented an earlier system for managing routine action selection (Cooper and Shallice, 2000) to include rudimentary supervisory processes, including some monitoring and error recovery. The system is capable of generating sequences of basic actions associated with the high-level task of packing a lunch box. In the augmented model, behaviours are assigned pre- and post- conditions. A task's post-condition is evaluated once all of its compulsory sub-tasks have been attempted. If the post-condition fails, the plan is assumed to be incomplete and so remains active. This form of task monitoring and error correction may be realised either as a 'local' function of the CS mechanism in which failed sub-tasks continue to receive excitation from the parent task (and hence may be attempted again) or as external supervisory control which modulates the activation of (inappropriate) tasks and sub-tasks from outside the contention scheduling system.

In this paper, we describe an integrated neural model for both routine and non-routine action selection. The model is based upon a functional architecture originally proposed by Norman & Shallice (Norman & Shallice, 1986; Shallice 1988). This model includes, as sub-components, a Contention Scheduler (CS) and a Supervisory Attentional System (SAS) that are associated with routine and non-routine action selection, respectively. We use the model to control a simulated robot, facilitating use of lesion studies to explore the correspondence between the behaviour of the robot and behaviours seen in humans suffering lesions of the PFC.

The remainder of this paper is organised as follows. Section 2 gives a fuller description of the Norman & Shallice model of executive control. Section 3 describes our implementation of this model, and especially of the SAS, in a simulated autonomous robot. Section 4 describes lesion studies of this architecture which induce a variety of behavioral pathologies in a robot that are analogous to those found in humans suffering damage to the PFC. Section 5 makes some concluding remarks.

2 Control of Action Selection

The Norman & Shallice model of action selection seeks to explain a number of behavioural pathologies observed in human patients. Examples of these behaviours include:

- 'capture errors' or utilisation behaviour - an inability to suppress a (strongly) triggered, but inappropriate behavior (Shallice, 1988); classic examples include William James dressing for bed instead of dressing for dinner and a patient reaching for a cup even though it has been explicitly agreed by the patient that they will not do when presented with it.
- the inability to act (akinesia) - attributable to the inability to resolve selection between competing behaviours (Robbins, 1991);
- persistence of a behaviour (stereotypy or perseveration) - a failure to notice significant cues (associated with successful completion of a task or sub-task) that should result in the expression of a different behaviour (Shallice, 1998).

These behaviours can be explained as a failure in error correction, i.e. the interruption of one behaviour and the willed initiation of another. In humans, they are associated with dysfunction of the pre-frontal cortex in an area which is functionally labeled the 'executive'. The executive initiates, monitors and modulates higher level behaviours (Parkin, 1996). Several accounts of the executive exist, notably those of Baddeley (Baddeley and Weiskrantz 1993) and Norman & Shallice (Shallice, 1998). Both the Baddeley and Norman & Shallice models contain an executive called the Supervisory System or Supervisory Attentional System (SAS), respectively.

Norman & Shallice's functional architecture for executive control of behaviour comprises several sub-components. A perceptual subsystem, via an associative mapping, causes a range of behaviours to '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). When two selected behaviours are incompatible, the CS is invoked. The contending behaviours are compared and adjusted dependent on the situation the agent finds itself in. 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.

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:

Monitor: the SAS must be able to compare the currently expressed action with an intended action (as formulated by the SAS or other 'planning' units). The monitor may be thought of as an 'arousal mechanism' that triggers the activation of the other attentional sub-units.

Modulate: when required, the SAS must provide a modulatory signal that attenuates the strength of triggering (salience) of inappropriate tasks and potentiates the salience of appropriate 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 a 'default' response;
- attenuate all intended behaviours for a given time, allowing the contention scheduler to express a behaviour governed by perception of the environment.

Generate: the SAS must create goals and strategies for solving novel problems.

The first two of these functions provide the basis of the SAS component in the integrated model we have developed.

3 Implementation

We have implemented the Norman & Shallice functional architecture as a large-scale, modular neural network controlling a simulated robot.

3.1 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).

3.2 The Model

The modular structure of the network corresponds, broadly, to the functional structure of the Norman & Shallice model as described. The network model is illustrated in Figure 2 (the size of the network, at ca. 700 neurons, requires a schematic representation.)

The large-scale network modules group clusters of highly interconnected neurons, most of which comprise four or eight input Elman or Jordan networks with up to three hidden layers. These clusters are individually trained using conventional training algorithms. In many cases, the pattern files are created by direct manipulation of the simulated robot in its environment. (All of the neurons used in the model have output values between 0.0 and 1.0 and are classed as 'active' above a value of 0.8. The weights on all connections between clusters are 1.0 so that weighting of inputs is determined within the cluster, thus reducing the burden of parameter setting.)

3.2.1 Perception and Associative Layers

The Perception Layer processes and fuses sensor signals to produce a representation of the environment. The output is distributed to the Associative Layer which maps the perceived state of the world to individual behaviours in the Behaviour Layer. Feedback from the Behaviour Layer into the Associative Layer enables the behaviours to provide excitation for the 'priming' of other relevant behaviours, e.g. a task, 'satisfy hunger' would stimulate associative memory clusters which propagate this raised salience forward to other food-related behaviours.

3.2.2 Behaviour Layer

The Behaviour Layer groups a number of clusters which exhibit basic behaviours such as 'wander safely', 'aggregate', 'disperse' (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 "schemas" in the Norman & Shallice model.)

The behaviour clusters 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 the Associative Layer, other behaviour clusters in the Behaviour Layer, and 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.). If a behaviour is granted expression by the CS, the behaviour experiences feedback from the CS designed to enhance persistence (Alexander, 1995).

3.2.3 Contention Scheduler

Prescott et al (1999) have developed a model of the CS based on the computational properties of the basal ganglia. We have developed an independent implementation of the CS following the principles established by Prescott et al. The contention scheduler takes input from behaviours in the behaviour network. The inputs represent the strength of 'requests' for access to the effector systems. The primary function of the CS is to select which behaviors are given that access to effectors, and hence expression in observable behaviour.

The contention scheduler is a layered network in which the layering reflects 'adjacency' of effector systems. Thus, localised lateral inhibition in the CS serves to prevent contradictory behavioural requests being made on a single effector system. This architecture allows the CS to resolve conflicting behavioural requests. For example, if, in the robot, the behaviours 'move_forward' and 'turn_left' both seek expression, then the regions of the CS controlling access to the left and right motors will select one behaviour for expression. (It is interesting to note that this CS architecture allows both behaviours to achieve expression at the right-hand drive wheel, but one must be excluded at the left.)

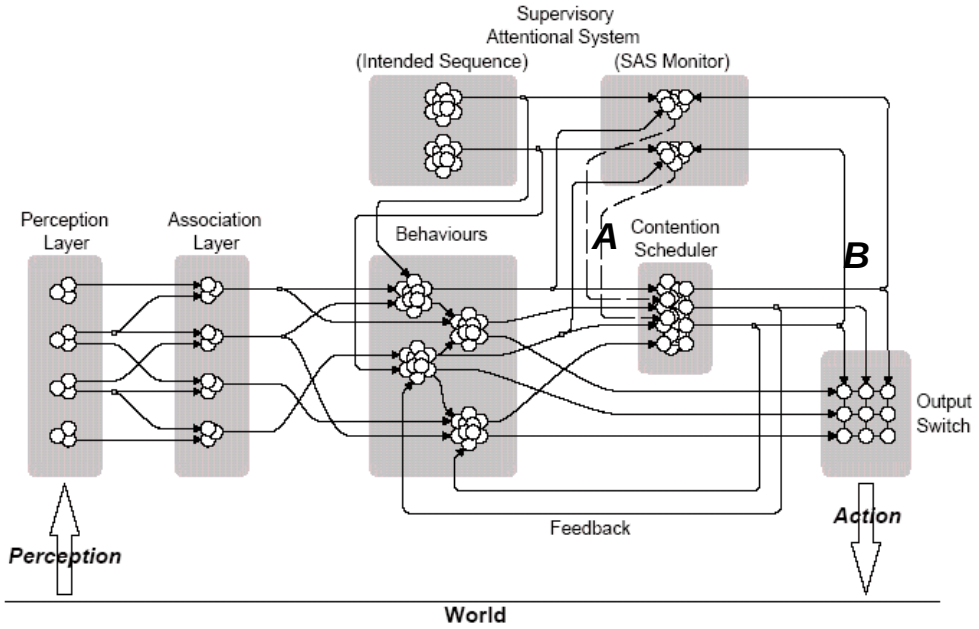

Figure 1. Implementation of the Norman & Shallice model of executive attentional control showing the major sub-components and how they interact.

The action of the CS is to *dis-inhibit* active behaviours which are otherwise inhibited by an effector gateway (also located in the thalamus).

We have already mentioned that the CS provides feedback to the Behaviour Layer to reinforce the currently expressed behaviour(s).

3.2.4 Supervisory Attention System

Although the SAS has several functions, including the generation of novel behaviours, only two functions are implemented currently; these are Monitor and Modulate (see Section 2, above). The full SAS has a Generate function to create novel plans. We have not yet implemented this function (neural systems for dynamic planning are at early stages of development).

Currently, we represent the 'result' of dynamic planning as 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 Monitor network clusters have three inputs: the environmentally induced behaviours (from the Behaviour Layer), the 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 Modulation network in the SAS.

The Modulation clusters generate outputs that modulates the signals from the Behaviours Layer into the CS so that intended behaviour is potentiated and the 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.

4 Lesion Studies

This section illustrates the operation of the simulated robot. First, we illustrate normal functioning of a robot that locates a food source, picks it up, takes it to a 'home', and puts it down (a behaviour called foraging). In the first instance, there are no distractions and the action sequence is entirely routine; it does not need the intervention of the SAS. We then demonstrate the ability of the SAS to correct the behaviour of the robot that is momentarily 'distracted' by the introduction of additional food before it has taken the original food item home. Second, we demonstrate the effects of two types of lesion to the SAS network.

4.1 Normal Behaviour

The normal functioning of the robot, with a non-lesioned controller is illustrated in Figure 3a. In this figure, the robot (labeled Penny), starting at the top left corner of the world, has detected food in a region near the bottom of the world, oriented and then moved towards that food, collected it, and is proceeding to take it towards 'home'.

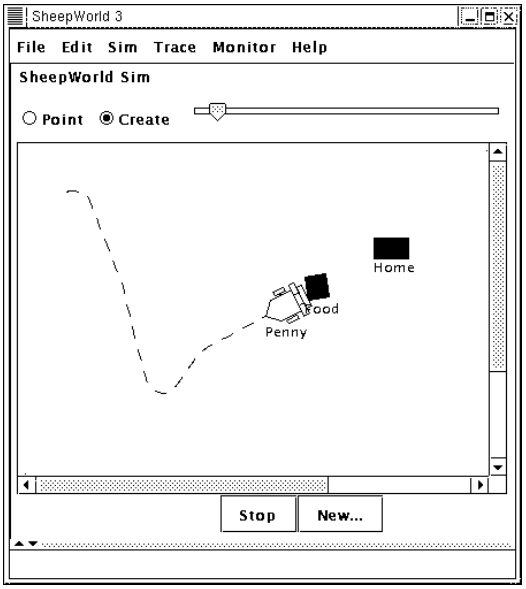

Figure 3(a) Routine foraging behaviour.

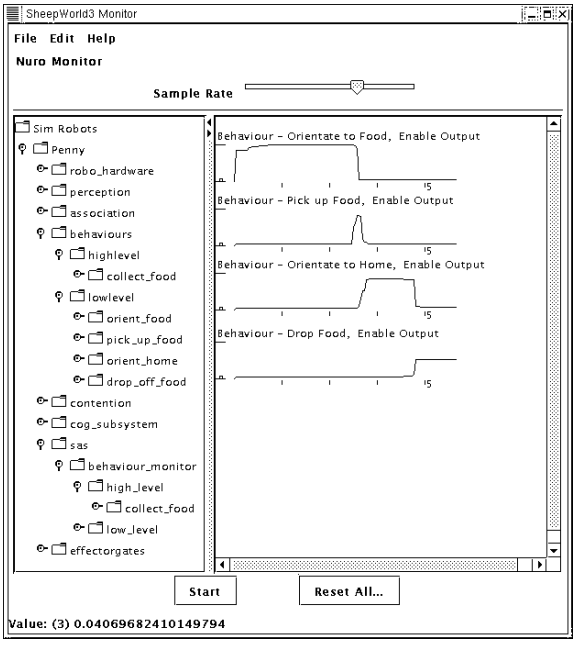

Figure 3(b). Selected traces from the neural network behaviour clusters which contribute to foraging behaviour (see text).

We can observe the output histories (traces) of selected network clusters during the experiment (Figure 3b). In this case, the traces illustrate the activities of behaviours selected for expression during forage: 'orient to food', 'pickup food', 'orient to home' and 'drop food'. The top trace shows the orient to food behaviour requesting expression; the second trace shows the spike associated with the short event of picking up the food in the gripper; the third trace represents expression of orient to home the final trace represents the dropping of the food at the home location.

In Figure 3c we illustrate normal operation of the SAS. The procedure is as for the first experiment. However, on this occasion, a distracting food source is introduced (2 seconds into this trace) before the robot has dropped the food at home.

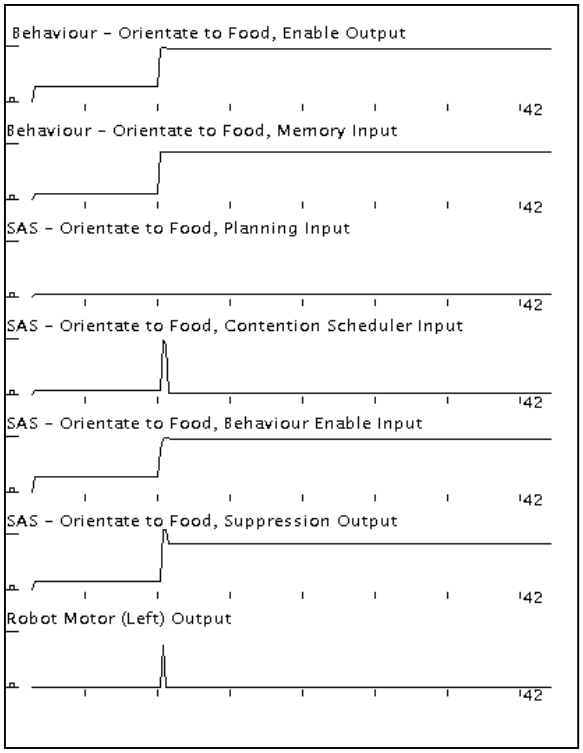

Figure 3c. Successful suppression of an inappropriate behaviour by the SAS.

Food is detected (trace 1) and 'orient to food' is triggered (trace 2). The SAS is not expecting this behaviour to be active at this point in its plan to take food 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 detects this (trace 5), and generates 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.

4.2 Lesion Study 1

In this experiment, the executive attentional control network is lesioned by blocking the modulatory signal from the SAS (Figure 2, marked *A*).

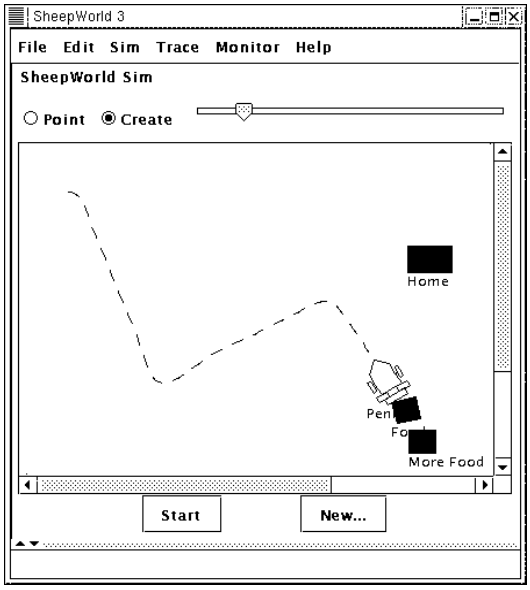

Figure 4a Robot exhibiting distracted behaviour.

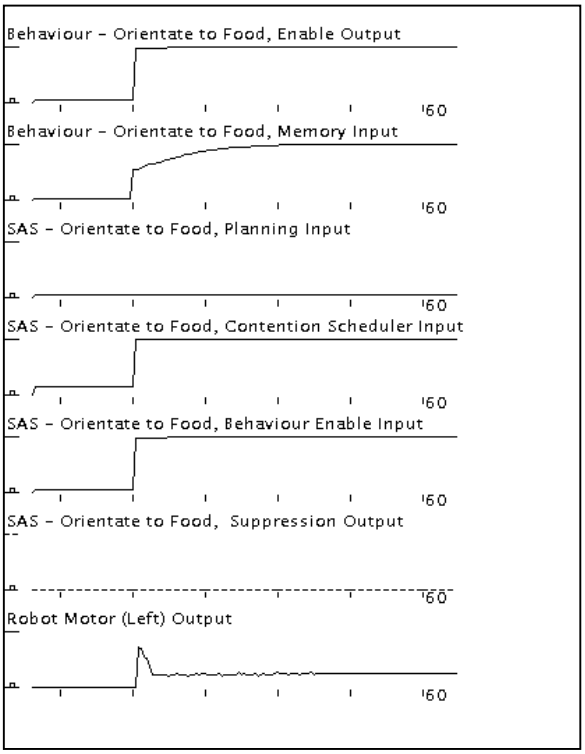

Figure 4b Robot exhibiting distracted behaviour resulting from failure to suppress inappropriate behaviour as a result of a lesion to the SAS.

Figure 4a illustrates the resulting robot behaviour; the robot responds to the new food by orienting towards it before it has taken the first food item home. In the corresponding activity traces of Figure 4b the first four traces are equivalent to those in Figure 3b: the food is detected (trace 1), 'orient to food' is triggered (trace 2), the SAS is not expecting this behaviour at this time (trace 3), the strength of the orient to food signal leads the CS to select (inappropriately) the 'orient to food' behaviour (rising slope in trace 4). However, on this occasion, and in keeping with the lesion introduced, the SAS fails to produce the modulatory signal (trace 6) and so there is no modulation of the inappropriately exhibited behaviour (no falling slope in trace 4) and the motor signal which results in the robot moving to the food (trace 7) is not completely suppressed.

4.3 Lesion Study 2

In this second experiment, the executive attentional control network is lesioned by preventing the SAS from monitoring the level of activity of a behavior prior to its expression (Figure 2, marked *B*). In this case, the SAS only detects expression of an inappropriate behavior after it is inappropriately expressed at the effector level (i.e. via perception of the changed environment).

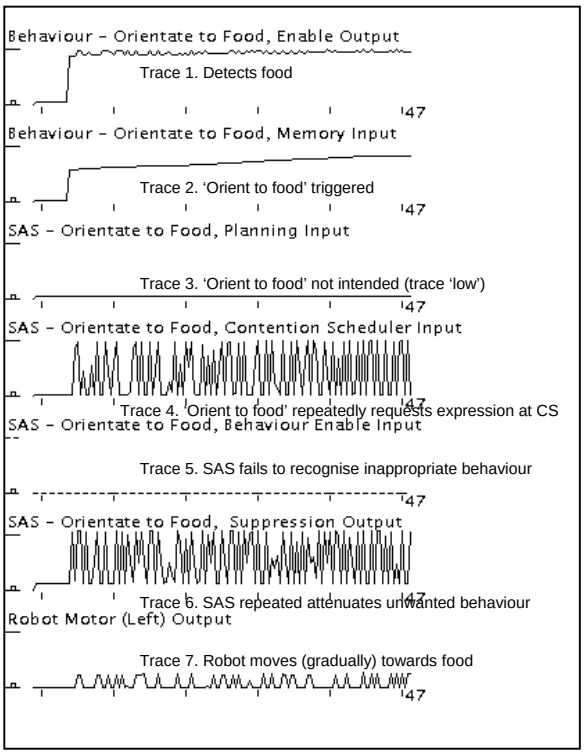

Figure 5 Robot exhibiting distracted behaviour resulting from failure to suppress inappropriate behaviour as a result of a lesion to the SAS.

The traces of Figure 5 illustrate this failure. New food is introduced (0.5 seconds) as a distraction whilst the robot is taking food it already holds to the home location. The

'orient food' behaviour is activated (traces 1 and 2). Again, the robot is not intending to exhibit this behaviour at this time (trace 3). As in normal operation (Figure 3c) the CS expresses the highly active behaviour (first rising spike of trace 4) and the SAS monitor detects that a modulatory signal is required to suppress this behaviour (first rising slope in trace 6). This suppresses the inappropriate behaviour which is then no longer selected by the CS (first falling slope in trace 4). However, the lesion introduced to the network prevents the SAS from recognising that the inappropriate behaviour is still strongly active, accordingly, when that behaviour is no longer enabled by the CS, the SAS ceases its modulation of the signal. This results in the inappropriate behaviour again achieving expression via the contention scheduler and the oscillatory cycle of expression and modulation is established (traces 4 and 6). The resulting behaviour of the robot (as indicated by trace 7) is that it repeatedly switches between two behaviours ('orient food' and 'orient home') moving very slowly in fits and starts.

5 Discussion

In the work described in this paper a functional model of executive attention, due to Norman and Shallice, was used as the basis of an implementation of a modular neural control architecture capable of willed behaviour. The experiments described demonstrate that the integrated architecture for the control of routine and non-routine actions is capable of sustaining normal behaviour in respect of simple tasks in a simulated robot.

The two lesion studies have resulted in distinct behavioural errors. Seeking to interpret the errors arising from the lesions, we suggest that the first lesion, blocking modulatory output from the SAS, provides an example of utilisation behaviour: i.e. despite an intention (maintained in WM) not to respond to the powerful stimulus of additional food, the robot is unable to suppress this inappropriate behaviour.

The behaviour seen in the second lesion study appears to correspond to a form of akinesia – the robot movement is very indecisive. Alternatively, it might be argued that it provides an example of perseverative behaviour in that, despite the fact that the robot already holds some food, it persists (in an oscillatory fashion) to exhibit the desire to move towards food.

The continued usefulness of the Norman & Shallice model in designing neural controllers for robots is supported by these studies.

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.

Cohen, J.D., Dunbar, K., McClelland, J.L. (1990) On the control of automatic processes: a parallel distributed processing account of the Stroop effect. Psychological Review, 97(3): 332-361.

Cohen, J.D., Braver, R., O'Reilly, A. (1998) A computational approach to prefrontal cortex, cognitive control and schizophrenia: recent developments and current challenges. In Roberts, A.C., Robbins, T.W., Weiskrantz, L. (eds.) The Prefrontal Cortex. Oxford University Press, 1998, 195-220.

Cooper, R.P. (2003) Organisation and disorganisation of complex sequential activities. In Detje, F., Dörner, D. & Schaub, H. (eds). Proceedings of the Fifth International Conference on Cognitive Modelling. Universitäts-Verlag, Bamberg, Germany, 51–56.

Cooper, R.P and Shallice, T. (1997) Modelling the Selection of Routine Action: Exploring the Criticality of Parameter Values. Proc. Cognitive Science Society, 1997, LEA, 131-136.

Cooper, R.P and Shallice, T. (2000) Contention Scheduling and the control of routine activities. Cognitive Neuropsychology, 17, 297–338.

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

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.

Gilbert, S.J. and Shallice, T. (2001) Task Switching: a PDP model. Cognitive Psychology, AP, Elsevier Science, 2001.

Gurney, K., Prescott T.J., Redgrave, P. (2001) A computational model of action selection in the basal ganglia (i), a new functional anatomy. Biological Cybernetics, 84: 401-410.

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

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

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.

Parkin A. J. 1996. Explorations in Cognitive Neuropsychology. Psychology Press, 1996.

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

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