

Cognition, Action Selection, and Inner Rehearsal

Murray Shanahan

Department of Electrical & Electronic Engineering,
Imperial College London,
Exhibition Rd., London SW7 2BT.

Abstract

This paper presents a large-scale model of the architecture of the mammalian brain, the core circuit of which carries out inner rehearsal of interaction with the environment to realise a form of cognitively mediated action selection. As it alternates between broadcast to and competition between its component neural assemblies, the core circuit exhibits an episodic dynamics suggestive of cortical processing in discrete frames. The implemented architecture is used to control a simulated robot, and a classic experimental paradigm in which rats performed apparently goal-directed action selection is emulated.

1 Introduction

In the 1940s, Tolman and Gletman used a classic experimental setup to demonstrate apparently goal-directed behaviour in rats (Tolman & Gletman, 1949). The rats were allowed to explore a T-maze containing a dark room on the left and a light room on the right (Fig. 1, left). Both rooms contained food. The rats were then placed in a separate enclosure resembling the dark room, and subjected to electric shocks through the feet. When reintroduced to the base of the T-maze, the rats always navigated directly to the light room, even though the actions of turning left and right had been equally reinforced.

The rat's ability to "think ahead" in this situation is hard to explain using reinforcement alone, and seems to require the inference of an indirect cause-and-effect relationship. However, Hesslow (2002) argues that the only extension to the paradigm of classical conditioning required to explain this sort of behaviour is a mechanism for inner rehearsal. Indeed, both Cotterill (1998) and Hesslow (2002) propose internally simulated interaction with the environment as the very basis of animal and human cognition.

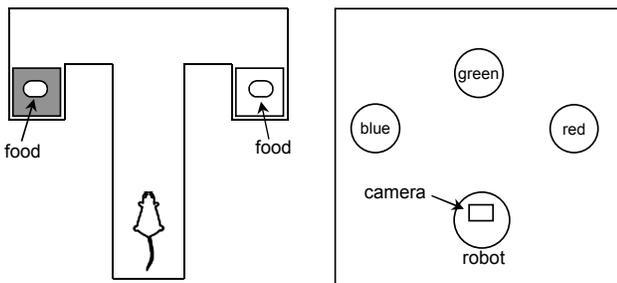


Fig 1: Rat and robot experiments

In pursuit of this suggestion, the present paper describes a large-scale, high-level neural model that realises goal-directed action selection for a simulated robot in an analogous experimental setup (Fig. 1, right). The model implements an architecture whose core circuit carries out inner rehearsal to anticipate the effects of currently executable actions, which are held on veto while these anticipated effects are evaluated by an affective system. This can bring about an increase or decrease in an action's salience, which in turn can result in the strengthening or weakening of its veto. When an action's salience exceeds a given threshold, its veto is released and the action is carried out.

The design of the core circuit facilitates the integration of the activities of multiple, parallel neural assemblies using a combination of competition and broadcast, and thereby realises a global workspace architecture (Baars, 1988; 2002). The dynamics of the core circuit exhibits a pattern of alternation between stability and rapid change, and is reminiscent of certain recent EEG findings suggestive of the idea that the cortex processes information in discrete frames (Freeman, 2003; 2004).

2 The Architecture of the Model

Fig. 2 shows a top-level schematic of the model's architecture. It can be thought of in terms of two interacting sub-systems. The first-order system is purely reactive, and determines an immediate motor response to the present situation without the intervention of cognition. But these unmediated motor responses are subject to a veto imposed by BG (the basal ganglia analogue). Through BG, which carries out salience-based action selection, the higher-order loop modulates the behaviour of the first-order system. It does this by adjusting the salience of currently executable actions. Sometimes this adjustment will result in a new action becoming the most salient, and sometimes it will boost an action's salience above the threshold required to release its veto, bringing about that action's execution.

The higher-order system determines these salience adjustments by carrying out off-line rehearsals of trajectories through (abstractions of) the robot's sensorimotor space. In this way – through the exercise of its "imagination" – the robot is able to anticipate and plan for potential rewards and threats without exhibiting overt behaviour. The first- and higher-order systems have the same basic components and structure. Both are sensorimotor loops. The key difference is that the first-order loop is closed through interaction with the world itself while the higher-order loop is closed internally. This internal closure is facilitated by AC, which simulates — or generates an abstraction of — the sensory

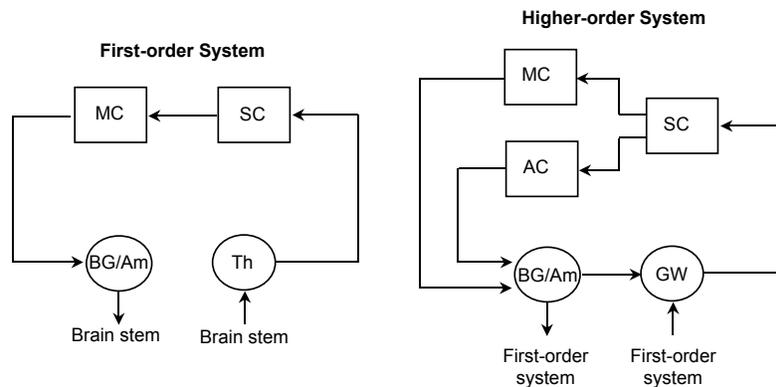


Fig. 2: A top-level schematic of the architecture. MC = motor cortex, SC = sensory cortex, AC = association cortex, BG = basal ganglia, Am = amygdala, Th = thalamus.

stimulus expected to follow from a given motor output, and fulfils a similar role to that of a *forward model* in the work of various authors (Demiris & Hayes, 2002; Hoffman & Möller, 2004; Grush, 2004; Ziemke, *et al.*, 2005). The cortical components of the higher-order system (SC, AC, and MC) correspond neurologically to regions of association cortex, including the prefrontal cortex which is implicated in planning and working memory (Fuster, 1997).

2.1 Affect and Action Selection

Analogues of various sub-cortical and limbic structures appear in both the first- and higher-order systems, namely the basal ganglia, the amygdala, and the thalamus. In both systems, the basal ganglia are implicated in action selection. Although, for ease of presentation, the schematic in Fig. 2 suggests that the final stage of motor output before the brain stem is the basal ganglia, the truth is more complicated in both the mammalian brain and the robot architecture under discussion.

In the mammalian brain, the pertinent class of basal ganglia circuits originate in cortex, then traverse a number of nuclei of the basal ganglia, and finally pass through the thalamus on their way back to the cortical site from which they originated. The projections up to cortex are thought to effect action selection by suppressing all motor output except for that having the highest salience, which thereby makes it directly to the brain stem and causes muscular movement (Redgrave, *et al.*, 1999). The basolateral nuclei of the amygdala are believed to modulate the affect-based salience information used by the basal ganglia through the association of cortically mediated stimuli with threat or reward (Baxter & Murray, 2002).

The robot architecture includes analogues of the basal ganglia and amygdala that function in a similar way. These operate in both the first- and higher-order systems. In the first-order system, the amygdala analogue associates patterns of cortical activation with either reward or punishment, and thereby modulates the salience attached to each currently executable action (Balkenius & Morén, 2001). The basal ganglia analogue adjudicates the competition between each executable action and, using a winner-takes-all strategy, selects the most salient for possible execution (Prescott,

et al., 1999). While the salience of the selected action falls below a given threshold it is held on veto, but as soon as its salience exceeds that threshold it is executed.

The roles of the basal ganglia and amygdala analogues in the higher-order system are similar, but not identical, to their roles in the first-order system (Cotterill, 2001). These structures are again responsible for action selection. However, action selection in the higher-order system does not determine overt behaviour but rather selects one path through the robot’s sensorimotor space for inner rehearsal in preference to all others. Moreover, as well as gating the output of motor association cortex (MC), the basal ganglia must gate the output of sensory association cortex (AC) accordingly, and thus determine the next hypothetical sensory state to be processed by the higher-order loop.

This distinction between first-order and higher-order functions within the basal ganglia is reflected in the relevant neuroanatomy. Distinct parallel circuits operate at each level (Nolte, 2002, p. 271). In the first-order circuit, sensorimotor cortex projects to the putamen (a basal ganglia input nucleus), and then to the globus pallidus (a basal ganglia output nucleus), which projects to the ventral lateral and ventral anterior nuclei of the thalamus, which in turn project back to sensorimotor cortex. In the higher-order circuit, association cortex projects to the caudate nucleus (a basal ganglia input structure), and then to the substantia nigra (a basal ganglia output nucleus), which projects to the mediodorsal nucleus of the thalamus, which in turn projects back to association cortex.

2.2 Global Workspace Theory

An important feature of the architecture, though not one that is explored fully in the present paper, is that it conforms to global workspace theory (Baars, 1988), which advances a model of information flow in which multiple, parallel, specialist processes compete and co-operate for access to a global workspace. Gaining access to the global workspace allows a winning coalition of processes to broadcast information back out to the entire set of specialists. Although the global workspace exhibits a serial procession of broadcast states, each successive state itself is the integrated product of parallel processing.

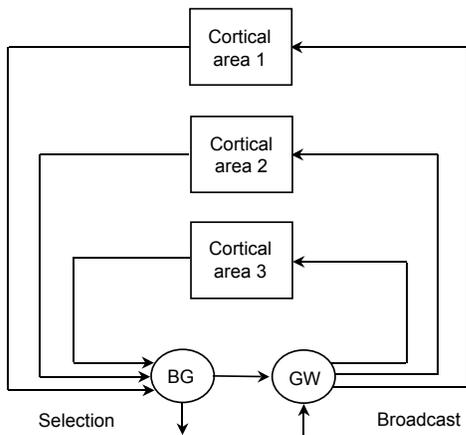


Fig 3: The fan-and-funnel model

According to global workspace theory, the mammalian brain instantiates this model of information flow, which permits a distinction to be drawn between conscious and unconscious information processing. Information that is broadcast via the global workspace is consciously processed while information processing that is confined to the specialists is unconscious. A considerable body of empirical evidence in favour of this distinction has accumulated in recent years (Baars, 2002). Although the topic of consciousness is orthogonal to the present paper, the combination of broadcast and competition that is the hallmark of the global workspace architecture is central to the action selection mechanism under investigation. During the process of internally exploring a space of possible sensorimotor trajectories, broadcast enables multiple branch points to be considered – in effect engaging many forward models simultaneously – while competition determines which of the candidate branches is actually explored next.

Moreover, the particular blend of serial and parallel computation favoured by global workspace theory suggests a way to address the frame problem – in the philosopher’s sense of that term (Fodor, 2000) – which in turn suggests that conscious information processing may be cognitively efficacious in a way that unconscious information processing is not (Shanahan & Baars, 2005). In particular, in the context of so-called informationally unencapsulated cognitive processes, it allows relevant information to be sifted from the irrelevant without incurring an impossible computational burden. More generally, broadcast interleaved with competition facilitates the integration of the activities of large numbers of specialist processes working separately. So the global workspace model can be thought of as one way to manage the massively parallel computational resources that surely underpin human and animal cognitive prowess.

The architecture of this paper conforms to the global workspace model of information flow by incorporating complementary mechanisms for the broadcast of information to multiple cortical areas and for selection between competing patterns of activation within those areas (Fig. 3). In Fig. 3, the locus of broadcast is denoted GW (for global workspace). Information fans out from GW to multiple cor-

tical sites (within which it may be subject to further local distribution). Conversely, information funnels back into GW, after competition within cortically localised regions, thanks to a process of selection between cortical sites realised by the basal ganglia.

A number of candidate structures exist in the brain that might fulfill the role of GW. For example, the first-order / higher-order distinction is preserved in the thalamus, which contains not only first-order relays that direct signals from the brain stem up to cortex (located, for example, in the lateral geniculate nucleus), but also higher-order relays that route cortical traffic back up to cortex (located, for example, in the pulvinar) (Sherman & Guillery, 2001). So the thalamus is one plausible candidate for a broadcast mechanism in the mammalian brain. But the same function could be realised by long-range corticocortical fibres, as proposed by Dehaene, *et al.* (2003), or indeed by some combination of thalamocortical and corticocortical communication.

Thankfully, there is no need to take a stand on this issue to supply an explanatory framework at an architectural level. What matters more in the present context is that the fan-and-funnel model of broadcast / distribution and competition / selection can be straightforwardly combined with the top-level schematic of Fig. 2, as is apparent from the diagrams. Indeed, the role of the BG component of the

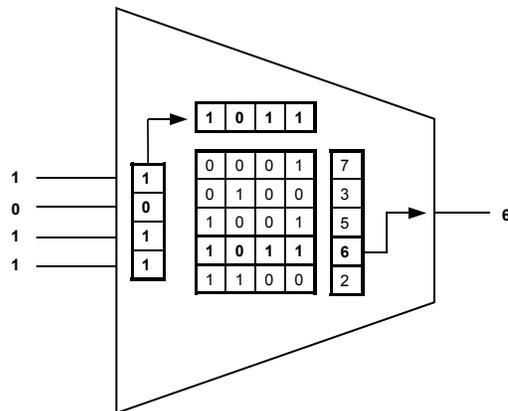


Fig 4: The G-RAM weightless neuron

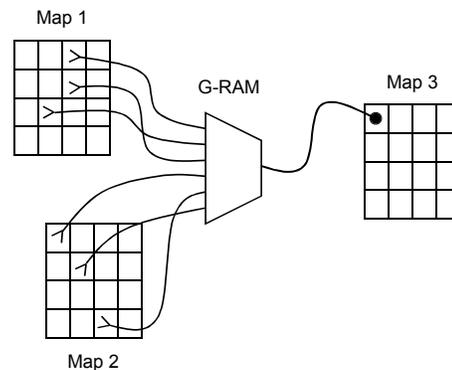


Fig 5: G-RAM maps and connections

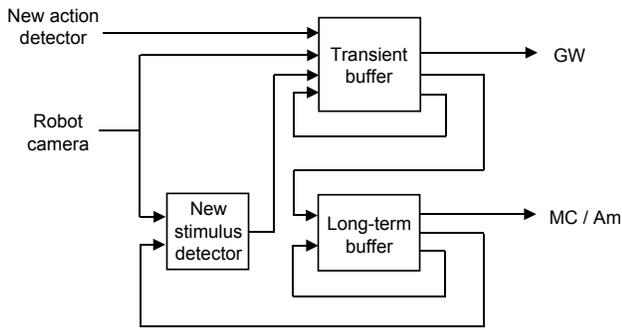


Fig. 6: Visual system circuitry (VC / IT). VC = visual cortex, IT = inferotemporal cortex.

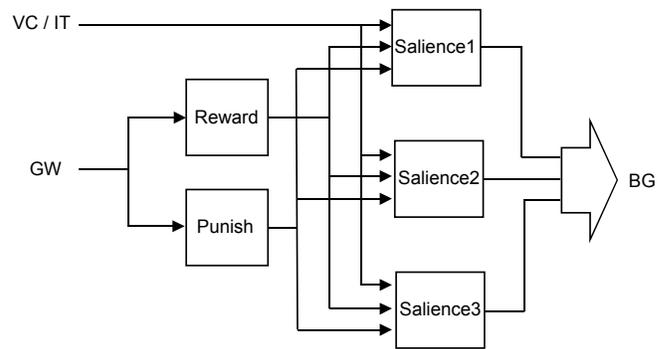


Fig. 7: Affect circuitry (Am)

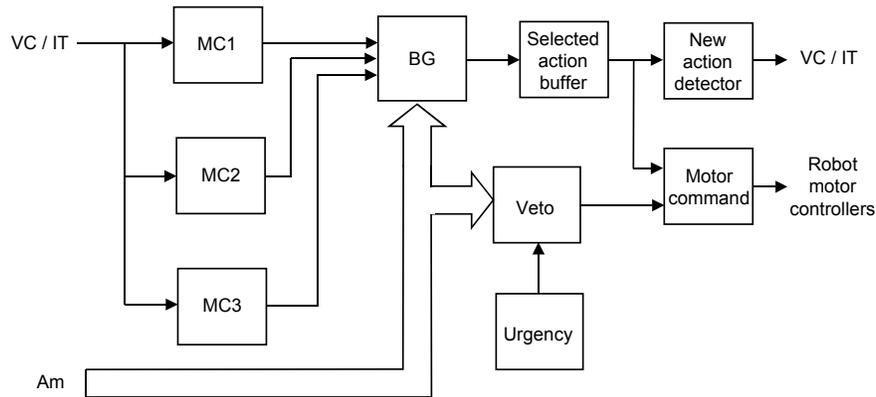


Fig. 8: Action selection circuitry (BG / MC)

higher-order loop introduced in Fig. 2 is precisely to effect the sort of selection between the outputs of multiple competing cortical areas shown in Fig. 3.

3 An Implementation

The brain-inspired architecture of the previous section has been implemented using NRM, a tool for building large-scale neural network models using G-RAMs (generalising random access memories) (Figs. 4 and 5). These are weightless neurons employing single-shot training whose update function can be rapidly computed (Aleksander, 1990), and which can be easily organised into attractor networks with similar properties to Hopfield nets (Lockwood & Aleksander, 2003).

The basic operation of a single G-RAM is illustrated in Fig. 4. The input vector is used to index a lookup table. In the example shown, the input vector of 1011 matches exactly with the fourth line of the table, which yields the output 6. When there is no exact match, the output is given by the line of the lookup table with the smallest Hamming distance from the input vector, so long as this exceeds a predefined threshold. In this example, if the input vector had been 1010, then none of the lines in the lookup table would yield an exact match. But the fourth line would again be the best

match, with a Hamming distance of 1, so the output would again be 6. If no line of the lookup table yields a sufficiently close match to the input vector the neuron outputs 0, which represents quiescence.

The core of the implementation, which comprises almost 40,000 neurons and over 3,000,000 connections, is a set of cascaded attractor networks corresponding to each of the components identified in the architectural blueprint of the previous section. The NRM model is interfaced to Webots, a commercial robot simulation environment. The simulated robot is a Khepera with a 64 × 64 pixel camera, and the simulated world contains cylindrical objects of various colours. The Khepera is programmed with a small suite of low-level actions including “rotate until an object is in the centre of the visual field” and “approach an object in the centre of the visual field”. These two actions alone are sufficient to permit simple exploration and navigation in the robot’s simple environment.

The overall system can be divided into four separate modules – the visual system (Fig. 6), the affective system (Fig. 7), the action selection system (Fig. 8), and the broadcast / inner rehearsal system (Fig. 9). Each box in these figures denotes a layer of neurons and each path denotes a bundle of connections. If a path connects a layer A to an $n \times n$ layer B then it comprises n^2 separate pathways – one for

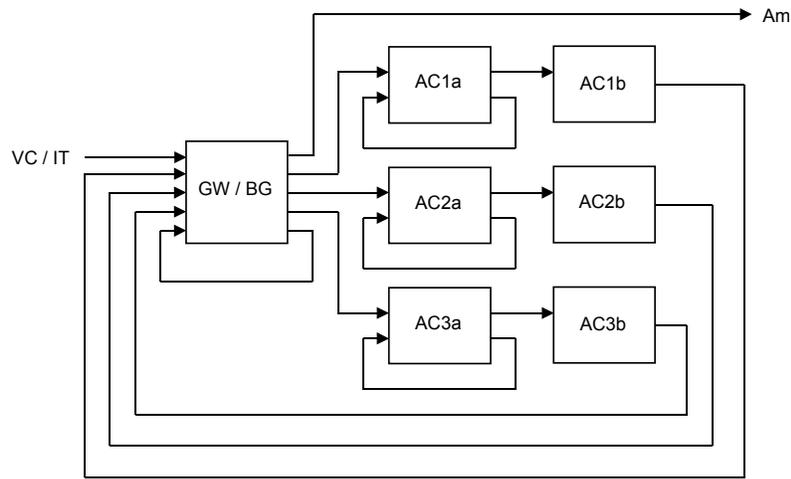


Fig. 9: Circuitry for broadcast and inner rehearsal (GW / BG / AC). GW = global workspace.

each of the neurons in B – each of which itself consist of m input connections originating in a randomly assigned subset of the neurons in A (Fig. 5). For the majority of visual maps m is set to 32.

The two buffers in the visual system comprise 64×64 topographically organised neurons (Fig. 6). These are both attractor networks, a property indicated by the presence of a local feedback path. The transient buffer is activated by the presence of a new visual stimulus. The hallmark of a new stimulus is that it can jog the long-term visual buffer out of one attractor and into another. The GW component of the inner rehearsal system is loaded from the transient visual buffer, whose contents rapidly fade allowing the dynamics of inner rehearsal to be temporarily dominated by intrinsic activity rather than sensory input.

The contents of the long-term visual buffer are fed to three competing motor-cortical areas, MC1 to MC3 (Fig. 8), each of which responds either with inactivity or with a recommended motor response to the current stimulus. Each recommended response has an associated salience (Fig. 7). This is used by the action selection system to determine the currently most salient action, which is loaded into the “selected action buffer” (Fig. 8). But the currently selected action is subject to a veto. Only if its salience is sufficiently high does it get loaded into the “motor command” buffer, whose contents is forwarded to the robot’s motor controllers for immediate execution.

So far the mechanism described is little different from a standard behaviour-based robot control architecture. What sets it apart from a purely reactive system is its capacity for inner rehearsal. This is realised by the core circuit depicted in Fig. 9, which is similar in both structure and function to the recurrent neural network of Tani (1996). When a new visual stimulus arrives, it overwrites the present contents of GW, and is thereby broadcast to the three cortical association areas AC1a to AC3a. The contents of these areas stimulates the association areas AC1b to AC3b to take on patterns of activation corresponding to the expected out-

comes of the actions recommended by their motor-cortical counterparts. These patterns are fed back to GW / BG, leading to further associations corresponding to the outcomes of later hypothetical actions. By following chains of associations in this way, the system can explore the potential consequences of its actions prior to their performance, enabling it to anticipate and plan ahead.

But for this capacity to be useful, the system needs to be able to *evaluate* hypothetical futures as it discovers them. So as a result of inner rehearsal, the salience of the currently selected action becomes modulated according to the affective value of the situations to which it might lead (Fig. 7). If the currently selected action potentially leads to a desirable situation, a small population of “reward” neurons becomes active, causing an increase in the salience of that action. This in turn may be sufficient to trigger the release of its veto, bringing about its execution. Conversely, if the currently selected action potentially leads to an undesirable situation, a small population of “punish” neurons becomes active. The resulting decrease in the salience of that action may cause a new action to become the most salient. In this case, the transient visual buffer is reloaded, its contents is passed on to GW, and the process of inner rehearsal is restarted. This is, in effect, a form of backtracking, allowing the system to perform a limited search of the space of possible courses of action.

To ensure that the system never gets stuck in a “thinking rut”, endlessly pondering the possible consequences of its actions instead of actually doing something, a small population of neurons acts as an indicator of the urgency with which the robot should act (Fig. 8). At the onset of a new stimulus, this neural population becomes quiescent, reflecting a lack of urgency, holding the currently selected action on veto and giving the inner rehearsal system time to work. But its level of activity grows with time, reflecting an increasing sense of urgency, and the need to act soon. The veto on the execution of the currently most favoured action is thereby gradually weakened, and eventually this action

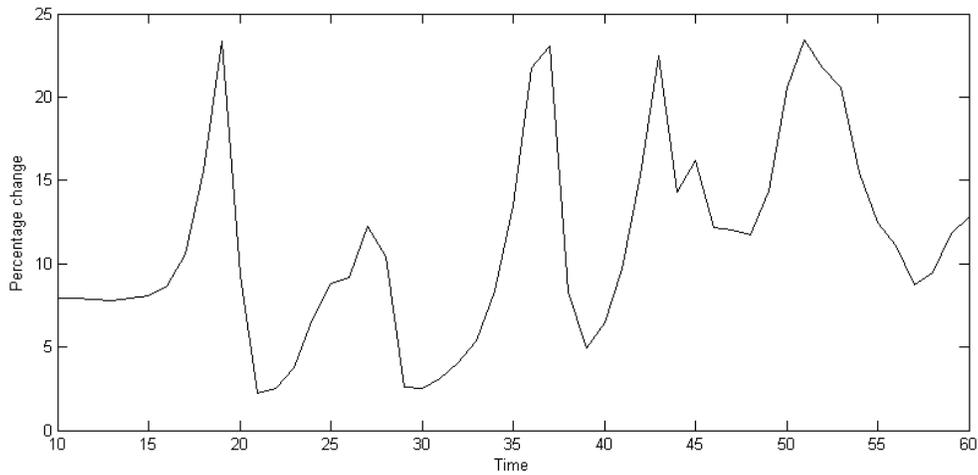


Fig. 10: Cycles of stability and instability

will be executed regardless of ongoing rehearsal. In this way, a balance is struck between reactivity and cognitively mediated, deliberative behaviour.

4 Results and Discussion

The implemented system currently runs on a 2.5 GHz Pentium 4 machine. Both Webots and NRM are run on the same machine, and the two systems communicate through an internal TCP socket. Under these somewhat unfavourable circumstances, each update cycle for the whole set of neurons takes approximately 750ms. A large proportion of this time is taken up by internal communication and graphics processing.

In each of the following experiments, the system runs a predefined training script prior to exhibiting the behaviour reported. Running this script sets up associations between patterns of visual input (VC / IT) and, for a subset of the three motor-neuronal assemblies (MC1 to MC3), corresponding recommended actions (Fig. 8) and their saliences (Fig. 7). This is analogous to reinforcement learning, acquiring a number of preferred immediate responses to an ongoing situation. In addition, the training script sets up associations between the current contents of GW and the punishment / reward neurons of Fig. 7. These permit the inner rehearsal mechanism, via the amygdala (Am), to exercise its influence on action selection. Producing similar results with a less supervised form of learning is an obvious theme for future research.

Fig. 10 illustrates an interesting property of the circuit of Fig. 9. The graph plots the percentage of neurons in the four maps GW and AC1a to AC3a that changed state from one time step to the next (where a time step corresponds to one complete cycle of updates to all the neurons in the system) during a typical run in which no external sensory input was presented to the robot. (A similar pattern is typically produced soon after the initial presentation of an external stimulus.) In order to study long chains of associations, a set of images of abstract coloured shapes (lozenges, stars, and

so on) was used as a training set, rather than images obtained from the Webots simulator. But the same effect is apparent with images obtained directly from the simulated robot's camera. Specifically, the graph shows that the system of inner rehearsal exhibits a procession of stable states punctuated by episodes of instability, a pattern which is reminiscent of the recently reported phenomenon of aperiodic alternation between pan-cortical coherent and decoherent EEG activity (Freeman & Rogers, 2003; Freeman, 2004). According to Freeman, these results suggest that the cortex processes information in a series of movie-like frames corresponding to "recurring episodes of exchange and sharing of perceptual information among multiple sensory cortices" (Freeman, 2004, p. 2077).

In a similar vein, the periods of stability depicted in the graph occur when the contents of GW is being successfully broadcast to the three cortical regions, while the spikes of instability indicate that GW is being nudged out of its previous attractor and is starting to fall into a new one. The new attractor will be the outcome of a competition between AC1b to AC3b. The resulting new contents of GW is then broadcast to AC1a to AC3a, causing new activation patterns to form in AC1b to AC3b, which in turn give rise to a renewed competition for access to GW. This tendency to chain a series of associations together is what gives the system its ability to look several actions ahead.

Tables 1 and 2 summarise episodes within two typical runs of the system, corresponding respectively to the with-out-aversion and with-aversion conditions in the classic experiment of Tolman & Gleitman (1949) described in the introduction (Fig. 1, left). Each episode starts with the initial presentation of a new stimulus, and ends with the robot's first action. Under both conditions, the robot's environment contained just three cylinders – one green, one red, and one blue (Fig. 1, right). Area MC1 of the motor-cortical system was trained to recommend "rotate right" (RR) when presented with a green cylinder, while area MC2 was trained to recommend "rotate left" (RL).

Table 1: Without aversion to red cylinders

Time	Events
0	Green cylinder comes into view.
2	Green cylinder image in both visual buffers. MC1 recommends RR, MC2 recommends RL. RR has higher salience and is currently selected action. Veto is on.
3	Green cylinder image in GW and broadcast to AC1a to AC3a. AC1b has association with red cylinder, AC2b has association with blue cylinder.
6	Associated red cylinder image in GW.
8	Affective system quiescent, but urgency increasing.
19	Urgency very high. Veto released.
20	RR passed on to motor command area. Robot rotates right until red cylinder in view.

Table 2: With aversion to red cylinders

Time	Events
0	Green cylinder comes into view.
2	Green cylinder image in both visual buffers. MC1 recommends RR, MC2 recommends RL. RR has higher salience and is currently selected action. Veto is on.
3	Green cylinder image in GW and broadcast to AC1a to AC3a. AC1b has association with red cylinder, AC2b has association with blue cylinder.
5	Associated red cylinder image in GW.
6	“Punish” neurons active, salience of RR going down.
9	Salience of RR very low. RL becomes currently selected action.
10	Transient visual buffer reloaded with green cylinder image.
14	Green cylinder image in GW and broadcast to AC1a to AC3a.
15	Associated blue cylinder image in GW. “Reward” neurons active. Salience of RL going up.
16	Salience of RL very high. Veto released.
17	RL passed on to motor command area. Robot rotates left until blue cylinder in view.

The action selection networks were trained in such a way that MC1’s recommendation (rotate right) had the higher initial salience, and in a purely reactive system this action would have been immediately executed under both the without- and with-aversion conditions. But thanks to the imposition of a veto, the inner rehearsal system had a chance to anticipate the outcome of the recommended action, giving rise to contrasting behaviours in the two experimental conditions, as in Tolman and Gleitman’s rat experiments. The inner rehearsal system was trained, using a predefined script matching the experimental setup, to associate 1) the RR action and the image of the green cylinder

with the subsequent presentation of the red cylinder, and 2) the RL cylinder and the image of the green cylinder with the subsequent presentation of the blue cylinder.

To emulate the without-aversion condition, the affective system was trained so that neither its “reward” nor its “punishment” neurons fired when GW contained the image of a red cylinder. Under this condition, the robot’s behaviour is the result of pure reinforcement. As Table 1 shows, this brought about the execution of RR – the system’s immediately preferred, reactive response – as soon as the combination of urgency and salience exceeded the threshold required to release the veto on that action.

By contrast, to emulate the with-aversion condition, the “punish” neurons were trained to fire when GW contained the image of the red cylinder. As Table 2 shows, this led the system to reduce the salience of its initially preferred action (RR) following a period of inner rehearsal that revealed its unpleasant expected consequences. The inner rehearsal system then explored the consequences of the alternative RL action. When these turned out to be more palatable, the salience of the RL action increased until its veto was eventually released, the RL command was forwarded to the motor output area, and the robot finally rotated to face the blue cylinder.

As all of this took place, urgency was increasing, but not fast enough to outpace the process of rehearsal and prevent it from influencing the selected action. The upper row of Table 3 summarises the results of eight further trials under the with-aversion condition, using the same training script but with a different randomly generated network configuration for each trial. The RL action is selected on each occasion, with some variation in timing.

Table 3: Sample runs with aversion

	Time to first action / action taken							
	1	2	3	4	5	6	7	8
□=8	17	16	15	16	15	15	20	15
	RL	RL	RL	RL	RL	RL	RL	RL
□=24	6	9	3	15	2	14	17	3
	RR	RL	RR	RL	RR	RL	RL	RR

The behaviour the system exhibits under these two experimental conditions demonstrates that the architecture is capable of an elementary form of cognitively mediated action selection similar to that first reported by Tolman and Gleitman (1949). Moreover, the architecture is broadly consistent with contemporary high-level neuroanatomy, and it conforms to the theoretical proposals of both Baars (1998) and Hesslow (2003). In addition, the episodic dynamics of its core circuit is supportive of Freeman’s interpretation of recent EEG findings in terms of discrete frames of cortical processing (Freeman & Rogers, 2003; Freeman, 2004). Neither the architecture nor the current implementation is confined to the simple experimental setup described in this paper, and their use in richer environments is the subject of ongoing work.

For example, by varying the system’s baseline level of urgency (□), it is possible to adjust the trade-off between

deliberation and reactivity – a high baseline level of urgency results in a tendency to act quickly but “unthinkingly” (Table 3, lower row), while a low baseline level of urgency results in slower but sometimes more effective action selection (Table 3, upper row). Preliminary experimentation also suggests that it may be possible to reproduce the behavioural phenomenon of “microchoices” reported by Brown (1992), wherein rats make tentative small explorations of arms of a star-maze before eventually making an apparently goal-directed choice. Using the mechanisms described here, a similar effect can be had by selecting a baseline level of urgency that allows for some anticipation of the consequences of actions, but only enough to look a very few actions ahead. The long-term hope is that, through experiments such as this, the conceptual framework and architecture of the present paper will help to further our understanding of the basis of cognition in both animals and machines.

Acknowledgments

Thanks to Igor Aleksander, Bernie Baars, Rodney Cotterill, Barry Dunmall, Gerry Hesslow, and to the workshop’s anonymous reviewers.

References

- Aleksander, I. (1990). Neural Systems Engineering: Towards a Unified Design Discipline? *Computing and Control Engineering Journal* 1(6), 259–265.
- Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Baars, B.J. (2002). The Conscious Access Hypothesis: Origins and Recent Evidence. *Trends in Cognitive Science* 6 (1), 47–52.
- Balkenius, C. & Morén, J. (2001). Emotional Learning: A Computational Model of the Amygdala. *Cybernetics and Systems* 32 (6), 611–636.
- Baxter, M.G. & Murray, E.A. (2002). The Amygdala and Reward. *Nature Reviews Neuroscience* 3, 563–573.
- Brown, M.F. (1992). Does a Cognitive Map Guide Choices in the Radial Arm Maze? *Journal of Experimental Psychology: Animal Behavior Processes* 18 (1), 56–66.
- Cotterill, R. (1998). *Enchanted Looms: Conscious Networks in Brains and Computers*. Cambridge University Press.
- Cotterill, R. (2001). Cooperation of the Basal Ganglia, Cerebellum, Sensory Cerebrum and Hippocampus: Possible Implications for Cognition, Consciousness, Intelligence and Creativity. *Progress in Neurobiology* 64, 1–33.
- Dehaene, S., Sergent, C. & Changeux, J.-P. (2003). A Neuronal Network Model Linking Subjective Reports and Objective Physiological Data During Conscious Perception. *Proceedings of the National Academy of Science* 100 (14), 8520–8525.
- Demiris, Y. & Hayes, G. (2002). Imitation as a Dual-Route Process Featuring Predictive and Learning Components: a Biologically-Plausible Computational Model. In K.Dautenhahn & C.Nehaniv (eds.), *Imitation in Animals and Artifacts*, MIT Press, pp. 327–361.
- Fodor, J.A. (2000). *The Mind Doesn't Work That Way*. MIT Press.
- Freeman, W.J. & Rogers, L.J. (2003). A Neurobiological Theory of Meaning in Perception Part V: Multicortical Patterns of Phase Modulation in Gamma EEG. *International Journal of Bifurcation and Chaos* 13 (10), 2867–2887.
- Freeman, W.J. (2004). Origin, Structure, and Role of Background EEG Activity. Part 1. Analytic Amplitude. *Clinical Neurophysiology* 115 (9), 2077–2088.
- Fuster, J.M. (1997). *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*. Lippincott-Raven.
- Grush, R. (2004). The Emulation Theory of Representation: Motor Control, Imagery, and Perception. *Behavioral and Brain Sciences* 27, 377–396.
- Hesslow, G. (2002). Conscious Thought as Simulation of Behaviour and Perception. *Trends in Cognitive Science* 6 (6), 242–247.
- Hoffmann, H. & Möller, R. (2004). Action Selection and Mental Transformation Based on a Chain of Forward Models. In *Proc. 8th International Conference on the Simulation of Behaviour (SAB 04)*, pp. 213–222.
- Lockwood, G.G. & Aleksander, I. (2003). Predicting the Behaviour of G-RAM Networks. *Neural Networks* 16, 91–100.
- Nolte, J. (2002). *The Human Brain: An Introduction to its Functional Anatomy*. Mosby.
- Prescott, T.J., Redgrave, P. & Gurney, K. (1999). Layered Control Architectures in Robots and Vertebrates. *Adaptive Behavior* 7 (1), 99–127.
- Redgrave, P., Prescott, T.J. & Gurney, K. (1999). The Basal Ganglia: A Vertebrate Solution to the Selection Problem. *Neuroscience* 89 (4), 1009–1023.
- Shanahan, M.P. & Baars, B. (2005). Applying Global Workspace Theory to the Frame Problem. *Cognition*, in press.
- Sherman, S.M. & Guillery, R.W. (2001). *Exploring the Thalamus*. Academic Press.
- Tani, J. (1996). Model-Based Learning for Mobile Robot Navigation from the Dynamical Systems Perspective. *IEEE Transactions on Systems, Man, and Cybernetics B* 26 (3), 421–436.
- Tolman, E.C. & Gleitman, H. (1949). Studies in Learning and Motivation: I. Equal Reinforcements in Both End-boxes, Followed by Shock in One End-box. *Journal of Experimental Psychology* 39, 810–819.
- Ziemke, T., Jirnhed, D.-A. & Hesslow, G. (2005). Internal Simulation of Perception: A Minimal Neuro-robotic Model. *Neurocomputing*, in press.