



A cognitive architecture that combines internal simulation with a global workspace

Murray Shanahan

Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK

Received 26 April 2005

Available online 27 December 2005

Abstract

This paper proposes a brain-inspired cognitive architecture that incorporates approximations to the concepts of consciousness, imagination, and emotion. To emulate the empirically established cognitive efficacy of conscious as opposed to non-conscious information processing in the mammalian brain, the architecture adopts a model of information flow from global workspace theory. Cognitive functions such as anticipation and planning are realised through internal simulation of interaction with the environment. Action selection, in both actual and internally simulated interaction with the environment, is mediated by affect. An implementation of the architecture is described which is based on weightless neurons and is used to control a simulated robot.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Global workspace theory; Simulation hypothesis; Models of consciousness

1. Introduction

Cotterill (1998, 2001) advances the proposal that thought is “internally simulated interaction with the environment,” and Hesslow (2002) argues that this “simulation hypothesis” can explain our experience of an inner world. However, while the simulation hypothesis has the potential to account for the content of conscious thought, it does not supply an answer to the question of what it is that distinguishes conscious from non-conscious activity in the brain. By contrast, global workspace theory can account for this distinction by appealing to an information processing architecture that features both competition among, and broadcast to, different brain processes (Baars, 1988, 1997).

The present article effects a marriage between these two proposals by presenting a neural-level cognitive architecture that realises an internal sensorimotor loop in which information passes through multiple competing cortical areas and a global workspace. This architecture, whose implementation is described here along with its application to the control of a simulated robot, serves to demonstrate that (i) the simulation hypothesis can be elegantly reconciled with global workspace theory and (ii) a robot controller which draws on these

E-mail address: m.shanahan@imperial.ac.uk.

contemporary ideas from the scientific study of consciousness is also viable from an engineering point of view.¹

For much of its history, mainstream cognitive science assumed language and reason to be the right conceptual foundations on which to build a scientific understanding of cognition. By contrast, the brain-inspired architecture described here, instead of manipulating declarative, language-like representations in the manner of classical AI and cognitive science, realises cognitive function through topographically organised maps of neurons, which can be thought of as a form of *analogical* (or *diagrammatic* or *iconic*) representation whose structure is close to that of the sensory input of the robot whose actions they mediate (Barsalou, 1999; Glasgow, Narayanan, & Chandrasekaran, 1995; Sloman, 1971).

Analogical representations are especially advantageous in the context of spatial cognition, which, though not the focus of the present paper, is a crucial capacity for both animals and robots. While common sense inferences about shape and space are notoriously difficult with traditional logic-based approaches (Shanahan, 2004), in an analogical representation basic spatial properties such as distance, size, shape, and location are inherent in the medium itself and require negligible computation to extract. Furthermore, traditional language-like representations bear a subtle and contentious relationship to the world they are supposed to represent, and raise difficult questions about intentionality and symbol grounding (Harnad, 1990; Shanahan, 2005a). With analogical representations, which closely resemble raw sensory input, this semantic gap is small and these questions are more easily answered.

In addition to these representational considerations, the design of the proposed architecture reflects the view, common among proponents of connectionism, that parallel computation should be embraced as a foundational concept rather than sidelined as a mere implementation issue. Specifically, the present paper advocates a computational architecture based on the *global workspace* model of information flow, in which a serial procession of states emerges from the interaction of many separate, parallel processes (Baars, 1988, 2002). This *serial* procession of states, which includes the unfolding of conscious content in human working memory (Baars & Franklin, 2003), facilitates anticipation and planning and enables a cognitively enhanced form of action selection. Yet the robustness and flexibility of these cognitive functions depends on the behind-the-scenes performance of extremely large numbers of *parallel* computations, only the most relevant of which end up making a contribution to the ongoing serial thread (Shanahan & Baars, 2005).

The architecture presented here is intended to be neurologically plausible at the level of large-scale neural assemblies, and contains analogues of a variety of brain structures and systems, including multiple motor-cortical populations (that compete for access to the global workspace), internal sensorimotor loops (capable of rehearsing trajectories through sensorimotor space), the basal ganglia (to carry out action selection), and the amygdala (to guide action selection through affect). But the central component is the global workspace itself, for which there are a number of candidate homologues in the vertebrate brain, including higher-order thalamocortical relays, and long-range corticocortical fibres.

In its overall conception, the architecture appeals to the notions of imagination and emotion as well as consciousness. Although perhaps only rough approximations to their humanly applicable counterparts, the way these concepts are deployed here is inspired by their increasingly important role in the brain sciences (Damasio, 2000). As such, the architecture described builds on the work of a number of other authors who have applied these ideas in the context of robotics or artificial intelligence.

- *Consciousness*. As already touched on, global workspace theory proposes a model of information flow in which conscious information processing is cognitively efficacious because it integrates the results of the brain's massively parallel computational resources (Baars, 1988, 2002). The global workspace architecture has previously been used in the design of software agents (Franklin, 2003; Franklin & Graesser, 1999), but its application to robotics has so far been neglected in the literature.

¹ The focus of the paper is the intersection of global workspace theory and the simulation hypothesis. It makes no claim to implement every aspect of the multi-faceted theory presented in Baars's, 1988 book. In particular, the global workspace architecture is not here deployed to assemble an effective reactive response to a novel situation, but only in the service of deliberation.

- *Imagination*. Although far from mainstream, the view that thought is internally simulated interaction with the environment or, to put it another way, the rehearsal of trajectories through sensorimotor space prior to their possible enactment, has influenced a number of robotics researchers in the biologically inspired tradition, including Chrisley (1990), Hoffmann and Möller (2004), Holland (2003), Stein (1995), Ziemke, Jirnhed, and Hesslow (2005).
- *Emotion*. Based on clinical studies, Damasio (1995) argued persuasively that the human capacity for rational deliberation is dependent on an intact affective system, and many other cognitive scientists subscribe to the view that affect addresses the problems of decision making and action selection (Picard, 1997; Sloman, 2001). It permits a number of factors to be blended together and brought to bear on the problem of contention for resources (i.e., muscles) by different brain processes. Neurologically plausible mechanisms of action selection compatible with this idea have been demonstrated in a robotics setting by Prescott, Redgrave, and Gurney (1999) and Cañamero (2003).

The rest of the paper is organised as follows. Section 2 presents a top-level schematic of the architecture, distinguishing first-order (external) from higher-order (internal) sensorimotor loops, and suggesting possible homologies between the components of the schematic and mammalian neuroanatomy. Section 3 introduces global workspace theory and shows how the top-level schematic can be made to conform to a global workspace architecture. Section 4 describes the implementation of the combined architecture at the level of its various neural assemblies and the circuitry that interconnects them. Section 5 outlines some experimental results obtained with the system, both with and without external connection to the robot. The concluding discussion addresses some of the methodological and philosophical issues the work brings up.

2. A top-level schematic

Fig. 1 shows a top-level schematic of the architecture. It can be thought of in terms of two interacting sensorimotor loops. The first-order loop, which is closed externally through the world, 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), which also carries out salience-based action selection. Through BG and Am (the amygdala analogue), 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 have no effect. But sometimes it will result in a new action becoming the most salient. And sometimes it will boost an action's salience above the threshold required to




Fig. 1. A top-level schematic of the architecture. MC, motor cortex; SC, sensory cortex; AC, association cortex; BG, basal ganglia; and Am, amygdala.

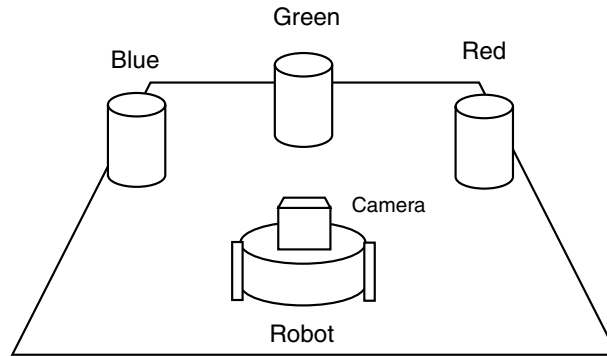


Fig. 2. The robot and its environment.

release its veto, bringing about that action's execution. The higher-order system computes these salience adjustments by carrying out off-line rehearsals of trajectories through (abstractions of) the robot's sensorimotor space. These rehearsals elicit an affective response in Am, which modulates the salience of the action currently selected by BG. 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.²

The first- and higher-order systems have the same basic components and organisation. The key difference is that, while the first-order loop is closed through interaction with the world itself, the higher-order loop is closed internally. This internal closure is facilitated by ACb, which simulates the sensory stimulus expected to follow from a given motor output, and fulfils a similar role to a *forward model* in the work of various authors (Demiris & Hayes, 2002; Grush, 2004; Wolpert, Doya, & Kawato, 2003). The first-order/higher-order distinction mirrors the corresponding distinction in neuroanatomy. The cortical components of the first-order system (SC and MC) correspond neurologically to primary sensory and motor areas, while those of the higher-order system (ACa and ACb) correspond to regions of association cortex whose major inputs arrive from other cortical areas. The distinction also reflects the functional roles of the relevant structures. In particular, the higher-order areas include the prefrontal cortex, which is implicated in planning and working memory (Fuster, 1997).

To get a better grasp of the interaction between the first- and higher-order systems, let's consider an example. Fig. 2 depicts the experimental setup used to test the ideas of this paper. It comprises a two-wheeled mobile robot located in a simple environment containing just three coloured cylinders. The robot is equipped with a camera, but the camera has a narrow field of view and can only see the cylinder ahead of it, not those to either side. In this situation, the robot has to choose one among three possible actions—to roll forwards, to rotate left, or to rotate right. Of particular interest here is the conflict between the immediate response of the reactive, first-order system and the cognitively mediated response of the higher-order system. The situation is analogous to that of a classic experimental paradigm in which rats act in a way that apparently requires the capacity to “think ahead” (Tolman & Gleitman, 1949), and is an example of what Beer (2000) calls *minimally cognitive behaviour*.

The MC component of the first-order, reactive loop immediately recommends a number of actions based on the current sensory input (mediated by SC). One of BG's duties is to select from these actions the one with the highest salience. In the present case, when SC has recognised a blue cylinder, suppose the most salient action recommended by MC is “rotate right.” If the first-order system were in total control, the robot would immediately execute this action. The relevant flow of information, initiated by the sensory stimulus, would pass uninterrupted through the structures of the first-order system, and result in a direct motor response.

² Franklin (2000) also describes the incorporation of deliberation into an implemented system based on a global workspace architecture. Franklin's aim is not to pursue the simulation hypothesis, nor is it to emulate neural architecture, except at a very abstract level. His implementation is accordingly very different from the one described here.

However, thanks to BG, the recommendations of the first-order system are intercepted before they issue in action. So the robot does not immediately rotate right. Instead, the rotate right action is held on veto, and the higher-order system kicks in.

The higher-order sensorimotor loop now simulates the effect of carrying out the “rotate right” action. The pattern of activation corresponding to the green cylinder stimulus is passed from SC to ACa. Assuming the system’s previously learned associations are right for the present environment, this will result in a pattern of activation appearing in ACb corresponding to the red cylinder image that would be expected if the currently selected action (rotate right) was carried out. The presence of the red cylinder image in ACb elicits a response from Am, which modulates the salience of the currently selected action. Suppose Am’s evaluation is negative—red cylinders are bad. The salience of the rotate right action is reduced accordingly, and a new action becomes the most salient, namely “rotate left.” The internal simulation process now restarts. The expected outcome of rotating left is a view of the blue cylinder. If Am’s evaluation of blue cylinders is sufficiently positive, this will increase the salience of the newly selected action above the threshold required for BG to release its veto. The robot will then rotate left until the blue cylinder comes into view.

The schematic’s suggestive use of labelling—BG for the basal ganglia, Am for the amygdala, and so on—reflects the fact that these components in the schematic are functionally analogous to real neuroanatomical structures. 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 (Mink, 1996; Redgrave, Prescott, & Gurney, 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; Cardinal, Parkinson, Hall, & Everitt, 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 thalamocortical 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.³

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 analogue 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.

³ A related action selection mechanism in combination with a global workspace architecture is described by Negatu and Franklin (2002).

3. Global workspace theory

A major architectural feature left out of the schematic of the previous section is parallelism. Yet the massively parallel computational resources of the biological brain undoubtedly underpin human cognitive prowess. Global workspace theory proposes an architecture for marshalling and managing these resources, and purports to account for the distinction between conscious and unconscious information processing within that architecture (Baars, 1988, 1997).

At the heart of the theory is 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 (Fig. 3). Although the global workspace exhibits a serial procession of broadcast states, each successive state itself is the integrated product of parallel processing. According to global workspace theory, the mammalian brain instantiates this model of information flow, which permits a distinction to be drawn between conscious and non-conscious information processing. Information that is broadcast via the global workspace is consciously processed while information processing that is confined to the specialists is non-conscious. A considerable body of empirical evidence in favour of this distinction has accumulated in recent years (Baars, 2002).

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 non-conscious 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.

The architecture of this paper conforms to the global workspace model of information flow insofar as the inner rehearsal mechanism incorporates complementary mechanisms for the broadcast of information to multiple cortical areas and for selection between competing patterns of activation within those areas (Fig. 4). As shown in Fig. 4, information fans out from the global workspace (GW) to multiple cortical sites (within which it may be subject to further local distribution). Conversely, information funnels back into the global workspace, possibly after competition within cortically localised regions (Koch, 2004, ch. 2), thanks to a process of selection between cortical sites realised by the basal ganglia.

Although the global workspace is depicted in Figs. 3 and 4 as a distinct component of the architecture, it should not be thought of as a functional unit that stores or processes information. Instead, the global workspace can be likened to the infrastructure of a communications network—the means by which information is gathered from and disseminated to multiple brain regions. Likewise, we should not expect it to be identifiable with any anatomically localised brain region. In neurological terms, a global workspace might be realised in a

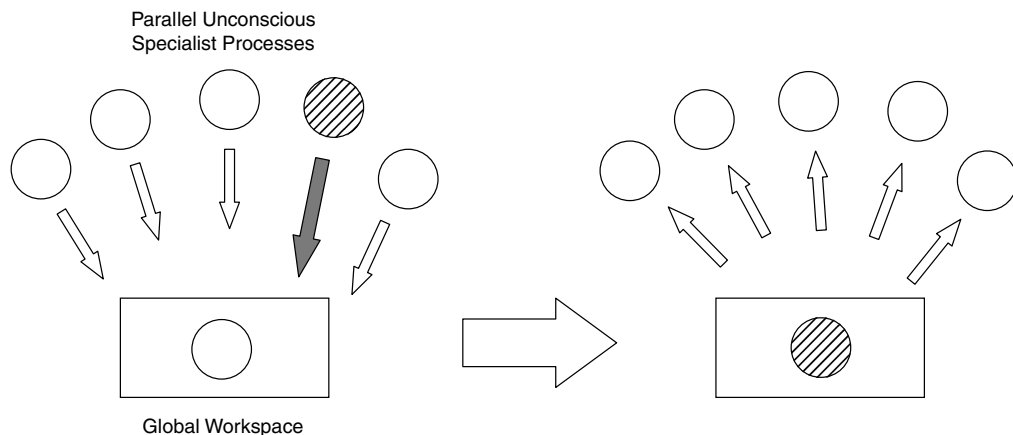


Fig. 3. The global workspace architecture.

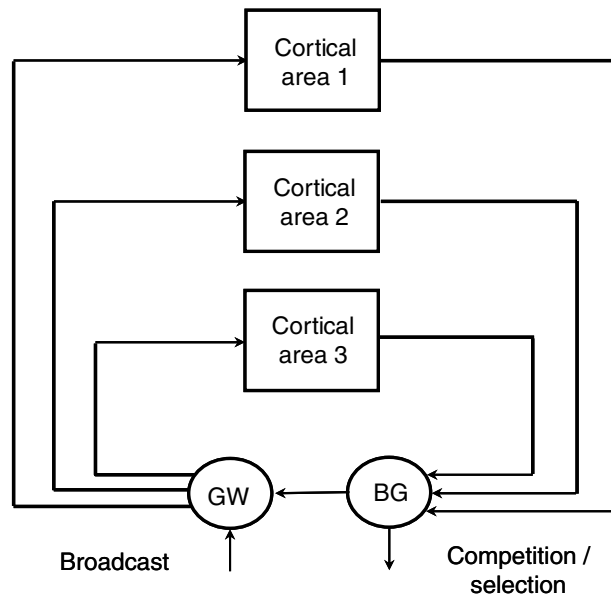


Fig. 4. The fan-and-funnel model.

variety of ways, such as by ascending thalamocortical fibres, or by long-range corticocortical fibres, or by some combination of both.

The first possibility—the thalamocortical workspace hypothesis—is supported by the fact that 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, 2002). However, while Sherman and Guillery have suggested a dominant role for thalamocortical relays in cortical communication, others have challenged this position. According to the opposing view, direct corticocortical projections are the main bearers of information from one cortical region to another (Van Essen, 2005). In line with this, Dehaene and his colleagues, who have also implemented a neural-level computer simulation of the global workspace architecture, favour the corticocortical workspace hypothesis (Dehaene & Changeaux, 2005; Dehaene, Sergent, & Changeaux, 2003).

A third possibility is that both hypotheses are true. The tendency for complex biological systems such as the brain to exhibit degeneracy (Edelman & Gally, 2001) suggests that both thalamocortical and corticocortical connections might independently be capable of supporting a global workspace. Indeed, the replication principle of Shipp (2003), according to which the connectivity of indirect transthalamic cortical pathways mimics that of direct corticocortical pathways, is supportive of this possibility. However, the present architecture is neutral with respect to this whole issue, as it does not attempt to model the brain at a sufficiently low level of detail for the distinction between the thalamocortical and corticocortical hypotheses to make sense.

The fan-and-funnel model of broadcast/distribution and competition/selection depicted in Fig. 4 can be straightforwardly combined with the top-level schematic of Fig. 1, as can be seen from the diagrams. In particular, the higher-order sensorimotor loop of Fig. 1 can be parallelised according to the model of Fig. 4, thus reconciling the simulation hypothesis with global workspace theory. The role of the BG component of Fig. 1's higher-order loop is to effect a selection between the outputs of multiple competing cortical areas, as shown in Fig. 4. The arrow connecting BG to GW in Fig. 4 is the path of internal closure for the higher-order loop of Fig. 1.

As we shall see shortly, the behaviour of the resulting system is best understood in dynamical system terms, and in this sense the outlook of the present paper accords with the methodology advocated by authors such as Van Gelder (1997), Beer (2000), and Wheeler (2005). The global workspace and the competing cortical assemblies each define an attractor landscape. Perceptual categories become attractors in a state space whose struc-

ture mirrors that of raw sensory input. Predictions are made by advancing the higher-order sensorimotor system along a simulated trajectory through (an abstraction of) that state space, enabling the global workspace to visit a sequence of attractors. But the various cortical assemblies are coupled in such a way that the vector of motion through this state space is determined by associations between one attractor and another, and different associations compete (and sometimes combine) to influence this vector of motion.

4. An implementation

The brain-inspired architecture of Sections 2 and 3 has been implemented using NRM (Dunmall, 2000), a tool for building large-scale neural network models using G-RAMs (generalising random access memories) (Figs. 5 and 6). These are weightless neurons employing single-shot training whose update function can be rapidly computed (Aleksander, 1990).

The basic operation of a single G-RAM is illustrated in Fig. 5. 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




Fig. 5. The G-RAM weightless neuron.

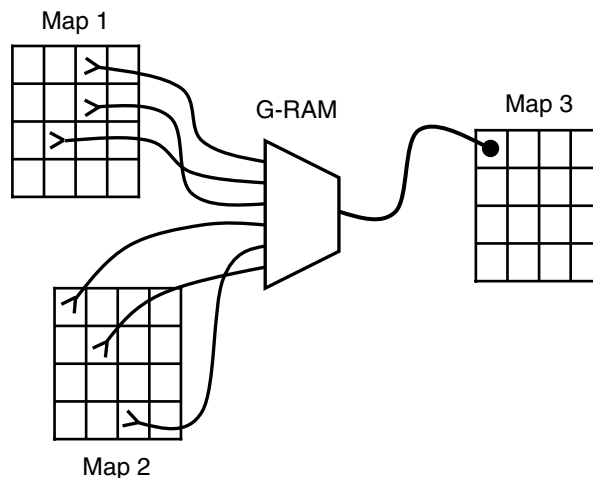


Fig. 6. G-RAM maps and connections.

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 does not exceed 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 implemented system exploits the fact that G-RAMs can be easily organised into attractor networks with similar properties to Hopfield nets (Lockwood & Aleksander, 2003). 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 Sections 2 and 3.

The NRM model is interfaced to Webots, a commercial robot simulation environment (Michel, 2004). 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 navigation in the robot’s simple environment.

The overall system can be divided into four separate modules—the visual system (Fig. 7), the affective system (Fig. 8), the action selection system (Fig. 9), and the broadcast/inner rehearsal system (Fig. 10). 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 each of the neurons in *B*—each of which itself consists of *m* input connections originating in a randomly assigned subset of the neurons in *A* (Fig. 6). 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. 7). 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-

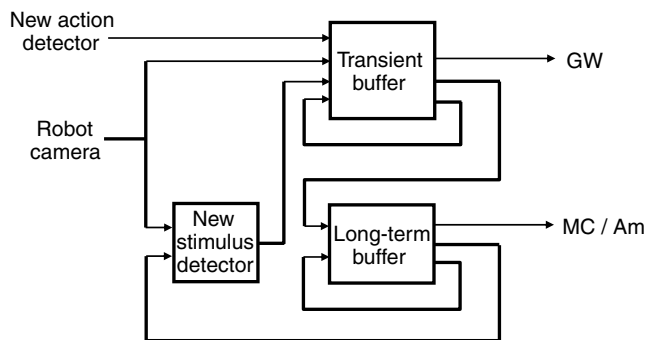


Fig. 7. Visual system circuitry (VC/IT). VC, visual cortex; IT, inferotemporal cortex.

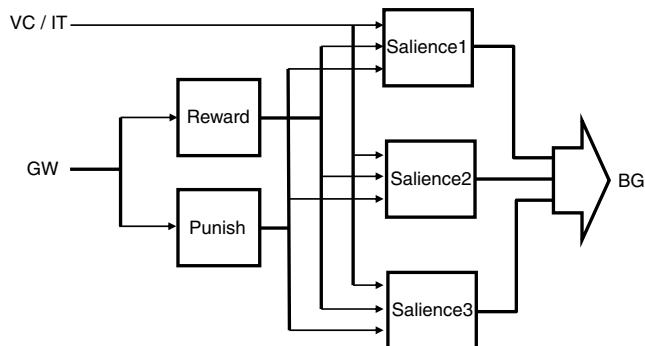


Fig. 8. Affect circuitry (Am).




Fig. 9. Action selection circuitry (BG/MC).

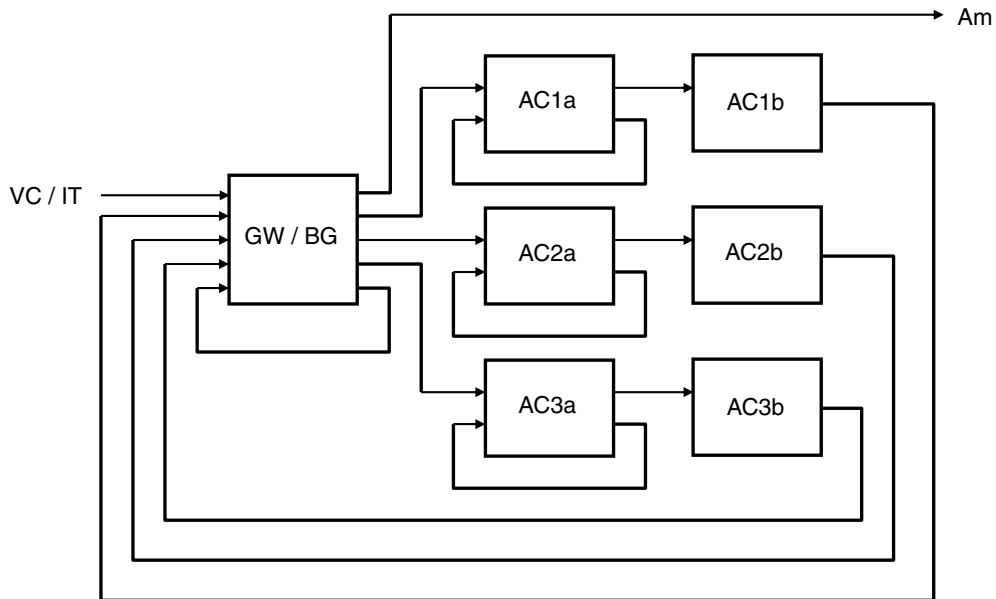


Fig. 10. Circuitry for broadcast and inner rehearsal (GW/BG/AC). GW, global workspace.

term visual buffer out of one attractor and into another. The global workspace component of the inner rehearsal system is loaded from the transient visual buffer, whose contents rapidly fade allowing the dynamics of the higher-order sensorimotor loop 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–MC3 (Fig. 9), 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. 8). This is used by the action selection system to determine the currently most salient action, which is loaded into the “selected action buffer” (Fig. 9). 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 are 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 (Brooks, 1986). 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. 10.⁴ When a new visual stimulus arrives, it overwrites the present contents of GW, and is thereby broadcast to the three cortical association areas AC1a–AC3a. The contents of these areas stimulate the association areas AC1b–AC3b to take on patterns of activation corresponding to the expected outcomes 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. 8). 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 salience of that action may cause a new action to become the most salient. In this case, the transient visual buffer is reloaded from the long-term buffer (Fig. 7), its contents is passed on to GW, and the process of inner rehearsal is restarted, permitting the system to explore a different possible future.

5. Experimental results

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 750 ms. A large proportion of this time is taken up by internal communication and graphics processing.

Fig. 11 illustrates an interesting property of the circuit of Fig. 10. The graph plots the percentage of neurons in the four maps GW and AC1a–AC3a that changed state from one time step to the next during a typical run in which the system was given a set of arbitrary associations between abstract coloured shapes but no external sensory input was presented to the circuit. (A similar pattern is typically produced soon after the initial presentation of an external stimulus.) 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 phenomenon of aperiodic alternation between pan-cortical coherent and decoherent EEG activity reported by various authors (Freeman & Rogers, 2003; Rodriguez et al., 1999). Freeman interprets these results to suggest that 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).

By computing the mutual information between GW and areas AC1a–AC3a (Fig. 12), it can be shown that the periods of stability depicted in the graph occur when the contents of GW are being successfully broadcast to those three cortical regions. This is supportive of Freeman’s hypothesis and, given the framework of global workspace theory, is suggestive of the possibility that conscious information processing itself is not continuous but quantised (VanRullen & Koch, 2003).⁵ Each peak in Fig. 12 denotes that the corresponding region of association cortex has fallen into a similar attractor to GW.⁶ (The absence of simultaneous peaks for all three

⁴ Similar networks are deployed by Tani (1996), based on the work of Elman (1990). These networks also features recurrent pathways capable of reproducing sequences of previously encountered states. The chief distinction in the present circuit is the use of local feedback connections—absent in Elman’s and Tani’s networks—which make the circuit of Fig. 10 into a recurrent cascade of attractor nets and give it a distinctive dynamics.

⁵ See also (Llinás, Ribary, Contreras, & Pedroarena, 1998, p. 1845), where it is conjectured that “consciousness is a non-continuous event determined by synchronous activity in the thalamocortical system.”

⁶ The notion of “similarity” makes sense here because all the maps in question are 2D arrays of neurons whose organisation matches that of the system’s raw visual input. In the context of such topographic mappings between areas, the similarity of their patterns of activation can be assessed using techniques from computer vision. The information-theoretic measure used here is adapted from (Viola & Wells, 1997).




Fig. 11. Cycles of system-wide stability and instability.




Fig. 12. Mutual information as an index of broadcast.

association areas is due to the differences in their attractor landscapes. When an association area lacks an attractor close to the pattern being broadcast from GW it becomes quiescent.) Conversely, the spikes of instability in Fig. 11 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–AC3b. The resulting new contents of GW are then broadcast to AC1a–AC3a, causing new activation patterns to form in AC1b–AC3b, which in turn give rise to a renewed competition for access to GW.

This tendency to chain a series of associations together accounts for the system's ability to look several possible actions into the future. In the way it migrates from one attractor to another, the dynamics of the system

can be characterised as *metastable* (Bressler & Kelso, 2001), and its behaviour also bears an interesting resemblance to *chaotic itinerancy* (Freeman, 2003; Tsuda, 2001), although in the context of discrete computation the concept is not strictly applicable without some modification.

In these and other respects (such as its use of re-entrant circuitry), the system also echoes Edelman and Tononi's *dynamic core* hypothesis (Edelman, 2003; Edelman & Tononi, 2000; Seth & Baars, 2005), according to which consciousness is realised in the brain by a shifting subset of its neurons (the dynamic core), whose composition is subject to constant and rapid change, but always comprises a set of distributed neural groups whose membership is tightly internally coupled while being functionally separated from the rest of the brain. Similarly, within the present simulation, competing areas of association cortex take turns to dominate the dynamics of the system through access to the global workspace from where their patterns of activation can achieve widespread influence. At any given time, the combination of GW and the areas of association cortex that are migrating to new attractors under its influence (such as AC2a and AC3a at time 20 in Fig. 12) can be likened to the dynamic core in Edelman and Tononi's sense.

Table 1 presents an illustrative sequence of events that occurred in a typical run of the whole system in which the ability to look ahead is put to good use. The episode described starts with the initial presentation of a new stimulus to the robot's camera, and ends with the robot's first action. The time is given in perception-update-action cycles, so the overall time between stimulus and response is around 17 s. This suggests that real-time performance would be attainable with current technology using a higher-end platform, assuming the Webots simulator is run on a different machine.

Prior to each such experimental run, the system executes a predefined training script. Running this script defines the input–output relation of each of the system's component neural areas. First, the script sets up associations between patterns of visual input (VC/IT) and, for a subset of the three motor-neuronal assemblies (MC1–MC3), corresponding recommended actions (Fig. 9) and their saliences (Fig. 8). This is analogous to reinforcement learning, acquiring a number of preferred immediate responses to an ongoing situation. Second, for the inner rehearsal mechanism (AC), the script sets up associations between various input stimuli and the expected sensory outcome of performing the actions that are salient for those stimuli (Fig. 10). Third, the training script sets up associations between the contents of GW and the punishment/reward neurons of Fig. 8, permitting 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.

A subset of the training data used in this experiment is summarised in Table 2. For the run presented, the robot's environment contained just three cylinders—one green, one red, and one blue (Fig. 2). 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). MC1's recommendation has the higher initial salience, and in a purely reactive system this action would be executed straight away. But thanks to the imposition of a veto, the inner rehearsal system gets a chance to anticipate the outcome of the recommend-

Table 1
An episode in a typical run

| Time | Events |
|------|--|
| 0 | Green cylinder comes into view |
| 4 | 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 |
| 7 | Green cylinder image in GW and broadcast to AC1a–AC3a. AC1b has association with red cylinder, AC2b has association with blue cylinder |
| 8 | Associated red cylinder image now in GW |
| 11 | “Punish” neurons active, salience of RR going down |
| 13 | Salience of RR very low. RL becomes currently selected action |
| 14 | Transient visual buffer reloaded with green cylinder image |
| 16 | Green cylinder image in GW and broadcast to AC1a–AC3a |
| 20 | Associated blue cylinder image now in GW. “Reward” neurons active. Salience of RL going up |
| 22 | Salience of RL very high. Veto released |
| 23 | RL passed on to motor command area. Robot rotates left until blue cylinder in view |

Table 2
A subset of the training data

| Area | Input pattern | Output pattern |
|-------------|----------------|--------------------|
| MC1 | Green cylinder | Rotate right |
| MC2 | Green cylinder | Rotate left |
| Saliency 1 | Green cylinder | Moderately salient |
| Saliency 2 | Green cylinder | Slightly salient |
| AC1b | Green cylinder | Red cylinder |
| AC2b | Green cylinder | Blue cylinder |
| Punish (Am) | Red cylinder | Active |
| Reward (Am) | Blue cylinder | Active |

ed action. This turns out to be an undesirable view of a red cylinder. So the system considers an alternative action (RR). Rehearsing the RR action yields a preferable expected outcome, namely a view of the blue cylinder, and the RR is therefore duly executed.

6. Discussion

Although only a prototype, the implemented system has demonstrated the viability of the proposed architecture. As the episode in Table 1 illustrates, a system conforming to the architecture is capable of generating a cognitively enhanced motor response to an ongoing situation. The design methodology used is, of course, quite different to that currently favoured by researchers in mainstream cognitive robotics (Lespérance et al., 1994), and is more closely allied to the research programme hinted at by Clark and Grush (1999). In place of viewpoint-free propositional representations, the present system employs viewer-centred analogical representations, and in place of symbolic reasoning it deploys a recurrent cascade of attractor networks. But compared with related products of the classical approach, the current implementation inherits several well-known disadvantages.

- While traditional propositional representations possess a compositional structure, and therefore comply with Fodor and Pylyshyn's *systematicity* constraint (Fodor & Pylyshyn, 1988), this is not true of the patterns of neuronal activity in the present system.
- Traditional propositional representations are adept at coping with *incomplete information* using disjunction and existential quantification. The present system can only deal with alternatives by using competitive parallelism and by exploring different threads of possibility at different times.
- Traditional planning systems are typically capable of effecting a *complete search* of the space of possible plans, while the presently implemented system of inner rehearsal ignores large tracts of search space and is only capable of a very crude form of backtracking.

Each of these issues is the subject of ongoing research, and a variety of techniques can be used to address them. For example, Tani has demonstrated that networks similar to the core circuit of Fig. 10 can embed combinatorial structure (Tani & Fukumura, 1995), and are capable of effective search (Tani, 1996), in both cases by exploiting chaotic dynamics. But brain-inspired cognitive architectures are relatively unexplored in artificial intelligence, and much work needs to be done before they can offer a viable alternative to the classical methodology in the domain of cognition.

In addition to its potential engineering application, the architecture presented here can be construed as a concrete statement of a specific hypothesis about human brain function. In line with the methodological stance outlined in Section 1, this hypothesis ascribes the capacity for high-level cognition to the interplay of consciousness, emotion, and imagination. Building a computer model and using it to control a robot is one way to give a clear interpretation to these concepts and to make precise their hypothesised role in mediating behaviour. But to conclude, let's consider the extent to which these philosophically difficult concepts—consciousness, emotion, and imagination—can legitimately be applied to artefacts conforming to the architectural blueprint of the present paper, such as the robot controller described in the previous section.

Let's begin with the concept of consciousness. The architecture respects all five of the “axioms of consciousness” proposed by Aleksander and Dunmall (2003). However, the present paper draws more heavily on the empirically grounded distinction between conscious and non-conscious information processing hypothesised by global workspace theory (Baars, 1988, 2002). This carries over straightforwardly to the broadcast/inner rehearsal system of Fig. 10. The processing of activation patterns that appear in GW and are subsequently successfully broadcast to cortex can be considered “conscious,” while all other information processing that goes on in the system is “non-conscious.” In accordance with global workspace theory, information that has been thus processed “consciously” integrates the contributions of many parallel processes, although the parallelism is very small scale in the implemented robot controller described here.

Similar considerations apply to the concepts of emotion and imagination. The functional role of the affective and inner rehearsal systems in the present architecture is identical to that proposed for emotion and imagination by many authors for the human case (Damasio, 1995, 2000; Harris, 2000). The argument, in a nutshell, is that “human beings have evolved a planning system in which felt emotion plays a critical role. By imagining what we might do, we can trigger in an anticipatory fashion the emotions that we would feel were we to actually do it” (Harris, 2000, p. 88). In much the same vein, the higher-order loop of Fig. 10 “imagines” what the robot might do, and this triggers an “emotional” response in the affective system of Fig 8.

However, the liberal use of scare quotes in the above paragraphs remains appropriate. There are many possible objections to the literal application of concepts such as consciousness and emotion to a robot such as the one described here. Prominent among these is the sheer poverty of the robot's external environment, the consequent poverty of its control system's internal dynamics, and the limited range of behaviour it can exhibit as a result. But consider a future humanoid robot in an unconstrained natural environment, equipped with a control system conforming to the proposed architecture. Suppose the robot's broadcast/inner rehearsal system comprised not six cortical regions but 100,000. Perhaps it would be harder to rein in the use of these concepts in such a case. But for now this remains pure science fiction.⁷

Acknowledgments

For help, discussion, and inspiration thanks to Igor Aleksander, Bernie Baars, Lola Cañamero, Ron Chrisley, Rodney Cotterill, Yiannis Demiris, Barry Dunmall, Ray Guillery, Gerry Hesslow, Owen Holland, Mercedes Lahnstein, Pete Redgrave, S.Murray Sherman, and Stewart Shipp. Thanks also to two anonymous reviewers for suggesting many improvements to the paper.

References

- Aleksander, I. (1990). Neural systems engineering: Towards a unified design discipline? *Computing and Control Engineering Journal*, 1(6), 259–265.
- Aleksander, I., & Dunmall, B. (2003). Axioms and tests for the presence of minimal consciousness in agents. *Journal of Consciousness Studies*, 10(4–5), 7–18.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge, MA: Cambridge University Press.
- Baars, B. J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford: Oxford University Press.
- Baars, B. J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Science*, 6(1), 47–52.
- Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Science*, 7(4), 166–172.
- Balkenius, C., & Morén, J. (2001). Emotional learning: A computational model of the amygdala. *Cybernetics and Systems*, 32(6), 611–636.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Baxter, M. G., & Murray, E. A. (2002). The amygdala and reward. *Nature Reviews Neuroscience*, 3, 563–573.
- Beer, R. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Science*, 4(3), 91–99.
- Bressler, S. L., & Kelso, J. A. S. (2001). Cortical co-ordination dynamics and cognition. *Trends in Cognitive Science*, 5(1), 26–36.
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2, 14–23.
- Cañamero, L. D. (2003). Designing emotions for activity selection in autonomous agents. In R. Trappé, P. Petta, & S. Payr (Eds.), *Emotions in humans and artifacts* (pp. 115–148). Cambridge, MA: MIT Press.

⁷ This is, of course, a very large and subtle topic, touching on contemporary debates about the so-called hard and easy problems of consciousness, the nature of the self, the prerequisites for subjectivity, and so on. The present author's views on these matters can be consulted in (Shanahan, 2005b).

- Cardinal, R. N., Parkinson, J. A., Hall, J., & Everitt, B. J. (2002). Emotion and motivation: The role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience and Biobehavioral Reviews*, 26, 321–352.
- Chrisley, R. (1990). Cognitive map construction and use: A parallel distributed processing approach. In D. Touretzky, J. Elman, G. Hinton, & T. Sejnowski (Eds.), *Connectionist models: Proceedings of the 1990 summer school* (pp. 287–302). Los Altos, CA: Morgan Kaufman.
- Clark, A., & Grush, R. (1999). Towards a cognitive robotics. *Adaptive Behavior*, 7(1), 5–16.
- Cotterill, R. (1998). *Enchanted looms: Conscious networks in brains and computers*. Cambridge, MA: 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.
- Damasio, A. R. (1995). *Descartes' error: Emotion, reason and the human brain*. Picador.
- Damasio, A. R. (2000). *The feeling of what happens: Body, emotion and the making of consciousness*. NY: Vintage.
- Dehaene, S., & Changeaux, J.-P. (2005). Ongoing spontaneous activity controls access to consciousness: A neuronal model for inattention blindness. *Public Library of Science Biology*, 3(5), e141.
- 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 Sciences of the United States of America*, 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* (pp. 327–361). Cambridge, MA: MIT Press.
- Dunmall, B. (2000). *Representing the sensed world in a non-biological neural system*. Imperial College London: Department of Electrical & Electronic Engineering.
- Edelman, G. M. (2003). Naturalizing consciousness: A theoretical framework. *Proceedings of the National Academy of Sciences of United States of America*, 100(9), 5520–5524.
- Edelman, G. M., & Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences of United States of America*, 98(24), 13763–13768.
- Edelman, G. M., & Tononi, G. (2000). *A universe of consciousness: How matter becomes imagination*. London: Basic Books.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.
- Fodor, J. A. (2000). *The mind doesn't work that way*. Cambridge, MA: MIT Press.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critique. *Cognition*, 28, 3–71.
- Franklin, S. (2000). Deliberation and voluntary action in “conscious” software agents. *Neural Network World*, 10, 505–521.
- Franklin, S. (2003). IDA: A conscious artifact? *Journal of Consciousness Studies*, 10(4–5), 47–66.
- Franklin, S., & Graesser, A. (1999). A software agent model of consciousness. *Consciousness and Cognition*, 8, 285–301.
- Freeman, W. J. (2003). Evidence from human scalp EEG of global chaotic itinerancy. *Chaos*, 13(3), 1–11.
- Freeman, W. J. (2004). Origin, structure, and role of background EEG activity. Part 1. Analytic amplitude. *Clinical Neurophysiology*, 115(9), 2077–2088.
- 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.
- Fuster, J. M. (1997). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe*. Lippincott-Raven.
- Glasgow, J., Narayanan, N. H., & Chandrasekaran, B. (1995). *Diagrammatic reasoning: Cognitive and computational perspectives*. Cambridge, MA: MIT Press.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377–396.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Harris, P. L. (2000). *The work of the imagination*. Oxford: Blackwell.
- 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: *Proceedings of the eighth international conference on the simulation of behaviour* (pp. 213–222). SAB 04.
- Holland, O. (2003). Robots with internal models. *Journal of Consciousness Studies*, 10(4–5), 77–109.
- Koch, C. (2004). *The quest for consciousness*. Roberts and Company.
- Lespérance, Y., Levesque, H. J., Lin, F., Marcu, D., Reiter, R., & Scherl, R. B. (1994). A logical approach to high-level robot programming: A progress report. In: B. Kuipers (Ed.), *Control of the physical world by intelligent systems: Papers from the 1994 AAAI fall symposium* (pp. 79–85).
- Linás, R., Ribary, U., Contreras, D., & Pedroarena, C. (1998). The neuronal basis for consciousness. *Philosophical Transactions of the Royal Society B*, 353, 1841–1849.
- Lockwood, G. G., & Aleksander, I. (2003). Predicting the behaviour of G-RAM networks. *Neural Networks*, 16, 91–100.
- Michel, O. (2004). Webots: Professional mobile robot simulation. *International Journal of Advanced Robotics Systems*, 1(1), 39–42.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50, 381–425.
- Negatu, A., & Franklin, S. (2002). An action selection mechanism for “conscious” software agents. *Cognitive Science Quarterly*, 2, 363–386.
- Nolte, J. (2002). *The human brain: An introduction to its functional anatomy*. St. Louis, MO: Mosby.
- Picard, R. (1997). *Affective computing*. Cambridge, MA: MIT Press.
- Prescott, T. J., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7, 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.
- Rodriguez, E., George, N., Lachaux, J.-P., Martinerie, J., Renault, B., & Varela, F. (1999). Perception's shadow: Long-distance synchronization of human brain activity. *Nature*, 397, 430–433.
- Seth, A. K., & Baars, B. J. (2005). Neural darwinism and consciousness. *Consciousness and Cognition*, 14, 140–168.
- Shanahan, M. P. (2004). An attempt to formalise a non-trivial benchmark problem in common sense reasoning. *Artificial Intelligence*, 153, 141–165.
- Shanahan, M. P. (2005a). Perception as abduction: Turning sensor data into meaningful representation. *Cognitive Science*, 29, 109–140.
- Shanahan, M. P. (2005b). Global access, embodiment, and the conscious subject. *Journal of Consciousness Studies*, in press.
- Shanahan, M. P., & Baars, B. J. (2005). Applying global workspace theory to the frame problem. *Cognition*, 98(2), 157–176.
- Sherman, S. M., & Guillery, R. W. (2001). *Exploring the thalamus*. New York: Academic Press.
- Sherman, S. M., & Guillery, R. W. (2002). The role of thalamus in the flow of information to cortex. *Philosophical Transactions of the Royal Society B*, 357, 1695–1708.
- Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philosophical Transactions of the Royal Society B*, 358, 1605–1624.
- Sloman, A. (1971). Interactions between philosophy and artificial intelligence: The role of intuition and non-logical reasoning in intelligence. *Artificial Intelligence*, 2, 209–225.
- Sloman, A. (2001). Beyond shallow models of emotion. *Cognitive Processing*, 2(1), 177–198.
- Stein, L. A. (1995). Imagination and situated cognition. In K. M. Ford, C. Glymour, & P. J. Hayes (Eds.), *Android epistemology* (pp. 167–182). Cambridge, MA: MIT 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.
- Tani, J., & Fukumura, N. (1995). Embedding a grammatical description in deterministic chaos. *Biological Cybernetics*, 72, 365–370.
- 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.
- Tsuda, I. (2001). Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behavioral and Brain Sciences*, 24, 793–810.
- Van Essen, D. C. (2005). Cortico-cortical and thalamo-cortical information flow in the primate visual system. *Progress in Brain Research*, 149, 173–185.
- Van Gelder, T. (1997). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628.
- VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Science*, 7(5), 207–213.
- Viola, P., & Wells, W. III, (1997). Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2), 137–154.
- Wheeler, M. (2005). *Reconstructing the cognitive world*. Cambridge, MA: MIT Press.
- Wolpert, D. M., Doya, K., & Kawato, M. (2003). A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society B*, 358, 593–602.
- Ziemke, T., Jirnhed, D.-A., & Hesselow, G. (2005). Internal simulation of perception: A minimal neuro-robotic model. *Neurocomputing*, 68, 85–104.