1 **Rapid #: -2727802** **Ariel IP: 128.205.111.1**

Status	Rapid Code	Branch Name	Start Date
Pending	UMC	Main Library	8/17/2009 1:12:14 PM

CALL #: **BF309 .J58**

LOCATION: **UMC :: Main Library :: UMCP McKeldin Library Periodical Stacks**

TYPE:	Article CC:CCG
JOURNAL TITLE:	Journal of consciousness studies : controversies in science & the humanities.
USER JOURNAL TITLE:	Journal of Consciousness Studies
UMC CATALOG TITLE:	Journal of consciousness studies :
ARTICLE TITLE:	Why axiomatic models of being conscious?
ARTICLE AUTHOR:	I. Aleksander, H. Morton
VOLUME:	14
ISSUE:	7
MONTH:	July
YEAR:	2007
PAGES:	0
ISSN:	
OCLC #:	32043673
CROSS REFERENCE ID:	515534
VERIFIED:	

BORROWER: **BUF :: Main Library**

PATRON: **Rapaport,William**

PATRON ID:	
PATRON ADDRESS:	
PATRON PHONE:	
PATRON FAX:	
PATRON E-MAIL:	-
PATRON DEPT:	Computer Science and Engineering
PATRON STATUS:	-
PATRON NOTES:	-

be defined or characterized; what different kinds of consciousness there are (core, phenomenal, access, functional, ...); how these different types relate to one another, and so on.

Finally, it is worth remembering that questions concerning artificial consciousness and artificial mentality more generally have important social and ethical ramifications, as more and more robots and other kinds of artificial humanlike agents are produced. Our predominating social and moral attitudes may be transformed as such technologies proliferate — indeed there may be tectonic shifts in prevailing notions of what 'society' is, and who 'we' are. So, in decades to come, developing work on machine consciousness may come strongly to affect how consciousness is seen, by both lay people and experts alike.

Acknowledgements

The editors thank Margarita Stapleton, whose organisational prowess was essential in running the Bristol AISB workshop; and Joel Parthemore for invaluable help at key stages in preparing of this special issue; also the conference chairs for the AISB-05 and AISB-06 conventions, and all who acted as anonymous referees, both for the workshops and for volume. Finally we thank Anthony Freeman for his patience and generous help at every stage in the production of ths issue.

References

Aleksander, I., & Dunmall, B. (2003), 'Axioms and tests for the presence of minimal consciousness in agents', *Journal of Consciousness Studies*, **10** (4), pp. 7–18.

Aleksander, I. (2007), 'Machine consciousness', in *The Blackwell Companion to Consciousness*, ed. Velmans and Schneider (Oxford: Blackwell), pp. 87–98.

Baars, B. (1988), *A Cognitive Theory of Consciousness* (Cambridge: CUP).

Clark, A. (2001). 'Visual experience and motor action: Are the bonds too tight?' *Philosophical Review*, **110** (4), pp. 495–519.

Chrisley, R., Clowes, R. and Torrance, S (ed. 2005), *Proceedings of the Symposium on Next-Generation Approaches to Machine Consciousness*, AISB-05, (Hatfield: University of Hertfordshire).

Clowes, R., Chrisley, R. and Torrance, S. (ed. 2006), *Proceedings of 2006 Symposium on Integrative Approaches to Machine Consciousness*. In T.Kovacs and J.Marshall (eds). *AISB'06: Adaptation in Artificial and Biological Systems, Vol 1*. (Bristol: University of Bristol), pp. 107–73.

Hesslow, G. (1994), 'Will neuroscience explain consciousness?' *Journal of Theoretical Biology*, **171**(1), pp. 29–39.

Holland, O. (ed. 2003), 'Editorial introduction', *Journal of Consciousness Studies*, **12** (4–5), Special issue on Machine Consciousness, pp. 1–6 .

Hurley, S. (1998), *Consciousness in Action* (Cambridge, MA: Harvard UP).

Maturana, H.R. & Varela, F.J. (1980), *Autopoiesis and Cognition* (Dordrecht: Reidel).

Metzinger, T. (2004), *Being No One* (Cambridge, MA: MIT Press).

O'Regan, J.K. & Noë, A. (2001), 'A sensorimotor account of vision and visual consciousness', *Behavioral and Brain Sciences*, **24** (5), pp. 939–73.

Searle, J.R. (1980). 'Minds, brains, and programs.' *Behavioral and Brain Sciences*, **3**(3), pp. 417–57.

Varela, F., Thompson, E. & Rosch, E. (1991), *The Embodied Mind*. (Cambridge, MA: MIT Press).

Igor Aleksander and Helen Morton

Why Axiomatic Models of Being Conscious?

Abstract: *This paper looks closely at previously enunciated axioms that specifically include phenomenology as the sense of a self in a perceptual world. This, we suggest, is an appropriate way of doing science on a first-person phenomenon. The axioms break consciousness down into five key components: presence, imagination, attention, volition and emotions. The paper examines anew the mechanism of each and how they interact to give a single sensation. An abstract architecture, the Kernel Architecture, is introduced as a starting point for building computational models. The thrust of the paper is to relate the axioms to the kernel architecture and indicate that this opens a way of discussing some first-person issues: tests for consciousness, animal consciousness and Higher Order Thought.*

Introduction and Overview

It is now commonplace that to attempt a logical analysis of consciousness is controversial due to the first-person or phenomenological character of the target topic. This paper argues that it is possible to design models based on introspection as a way of including the first person in the model itself. Our motivation within machine consciousness comes not so much from a desire to create artificially conscious artefacts as from a desire to pin down the necessary properties of mechanisms that support consciousness. We use a *constructive* approach: understanding by synthesis. Since early discussions of machine models that specifically attempt to capture elements of consciousness there

Correspondence:
Igor Aleksander, Department of Electrical & Electronic Engineering, Imperial College, London SW7 2BT. *i.aleksander@imperial.ac.uk*
Helen Morton, School of Social Sciences, Brunel University, Uxbridge, UB8 3PH, U.K. *Helen.morton@brunel.ac.uk* (also Imperial College).

has been an apparent spectrum from the functional to the phenomenological. In functional models the presence of consciousness is judged on behaviour (e.g. Franklin, 2003). At the other end of the spectrum, modelling involves hypotheses of what is necessary in a mechanism to support internal sensations akin to our personal experience. Our work falls into this category and is distinguished by a decomposition of sensations into five classes which we call axioms.

Taking the phenomenological stance, we see consciousness as a definite product of the brain and seek to reconcile what is known of the neurology of the brain with the axiomatic features mentioned in the Abstract. By 'brain', however, we refer to a general class of machine that could be described without reference to detailed physical makeup, but which is characterized by its 'state dynamics' and the dependence of such dynamics on broad architectural characteristics of the mechanisms. This is the classical methodology of 'automata theory', which is the *lingua franca* of our approach (see Aleksander, 1996). In particular we address and argue against the notion of an unassailable 'hard' problem, which implies that a brain may become totally bereft of its consciousness (not just in sleep) without any physically measurable change in its neurology and the states available to such neurology.

Our study of consciousness starts with an introspective decomposition of different aspects of being conscious into five major axiomatic sensations (presence, imagination, attention, planning and emotions) so as to make the modelling task tractable. These axioms are mapped into mechanisms that form an interlocking ensemble called a kernel architecture (KA) as it forms the basis of many computational models implemented to date. Also, in developing the mechanisms that support the axioms, we introduce a central mechanism: *depiction*. This involves not only sensory pathways in the brain, but also signals from the musculature of the body: it is the mechanism that is the basis of the phenomenology of our models.

Then, to illustrate our computational methodology and the role of the KA, we show that it leads to tests for claims that an object is conscious, suggests that most animals are conscious and that Higher Order Thought is a 'content' issue that requires no machinery over and above the KA.

A Comment on the History of Machine Consciousness

At a small workshop at the Cold Spring Harbor Laboratory, the paradigm of machine consciousness appears to have been established

among international workers in 2001. Organized by Koch, Chalmers, Goodman and Holland, a mixed group of neurologists, computer scientists and philosophers were reported to have concluded:

> The only (near) universal consensus at the workshop was that, in principle, one day computers or robots could be conscious. In other words, that we know of no fundamental law or principle operating in this universe that forbids the existence of subjective feelings in artifacts designed or evolved by humans. (Koch, 2001).

This led to the publication of a special issue of the *Journal of Consciousness Studies* on machine consciousness, edited by Owen Holland (2003). Here contrasting styles of machine modelling of consciousness appeared: functional (where an appearance of being conscious would be judged from the behaviour of the system) and 'phenomenological' approaches. In the latter category, and the same issue of *JCS*, we first indicated the axiomatic approach (Aleksander and Dunmall, 2003). This approach is elaborated in this paper and will be shown to link to phenomenology through a property called 'depiction', while the five axioms, derived from introspection, are used to develop computational mechanisms.

Since 2003, the concern with phenomenology in machine consciousness studies (under the title of *Synthetic Phenomenology*) can be noted (for example, Chrisley, in this issue, and Aleksander and Morton, 2006).

The Challenge of the First Person: Is there a Hard Problem?

The word 'phenomenology' is generally used for philosophical concerns that include, if not begin with, personal, internal feelings of being conscious. From a modelling standpoint this implies that introspection is a suitable starting point for the model. The implication for computational studies is to establish a link between experienced conscious states and the observable states of some underlying physical computational mechanism. First, the word 'computational' needs to be qualified. Conventionally this suggests a process with a defined algorithm that runs on a von Neumann type of architecture. Here we do not mean this at all. Computation is taken to be the time development of the inner states of an architecture that is controlled by and in touch with an external world. The prototype model for such an architecture is the living brain. This is an evolved system the fitness function of which is to keep the organism informed of the reality of the world and the organism's own potential in this world. To approach the problem computationally we need to assert:

> *Assertion 1: It is necessary to abstract the physical/informational properties of the brain that determine conscious states and study these as structures that are* ***virtual*** *on a conventional computer architecture.*

In other words, the algorithms needed to create the virtual architectures are unimportant. Brief mention must be made on how the above assertion relates to Chalmers' (1996) 'hard problem'. He allows that while the states of mind we actually sense may (and do) have a correlation with physical events that may be measured in the brain, these physical events do not, in a logically guaranteed way, lead to the presence of phenomenological personal events. In other words, this opens the possibility of finding 'zombies': organisms that are entirely physically equivalent to conscious ones except for the lack of phenomenology. Therefore, while work on the physical may be very interesting, it misses precisely the source of the inner phenomenological sensation. We take the view that if one is going to understand phenomenology it is necessary to start with an organism that is undoubtedly phenomenologically conscious. Then it is possible to discuss the mechanistic necessities of such an organism. But the only organism that we each know is undoubtedly conscious is our very self, and it is for this reason that the axioms developed below are based on introspection.

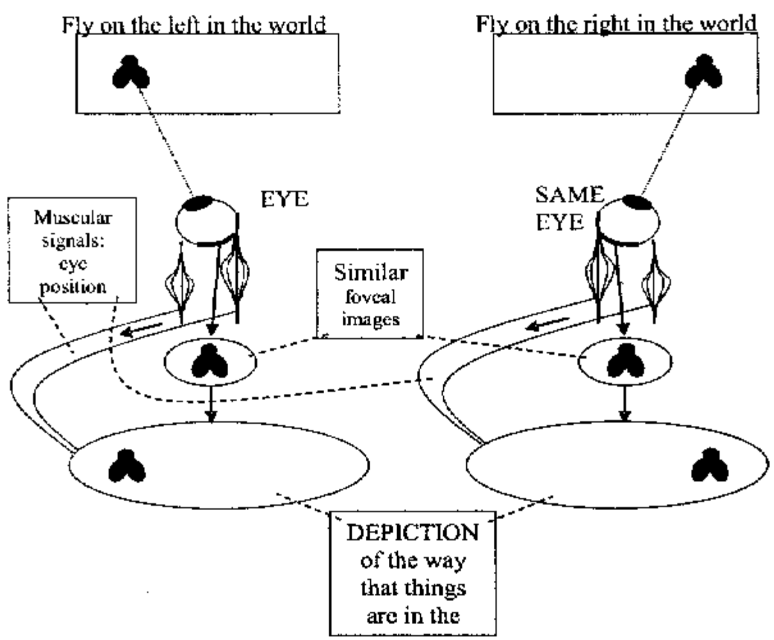

Figure 1. An example of a depictive process which involves muscle action (reproduced from Aleksander, 2005).

Depiction: The Key Mechanism for Conscious Representation

Depiction is a direct representation of *where* elements of the world are. It uses the efforts of the mechanism to attend to such elements. An example of depiction appears in Figure 1.

There is a vast literature of the existence of cells that only react to sensory signals if addressed by the appropriate muscular activity (e.g. called 'Gaze-Locked' cells in vision, Galletti and Battaglini, 1989). This involves a variety of inputs from musculature, including head, body and even arm movement. For example, if one points at an object, this helps to depict that object in the brain. So if there is an overall mechanism that is essential for a machine to support conscious sensation it is depiction. This leads to the second assertion.

> *Assertion 2: The parts of a mechanism that sustain conscious experience can only do so if they are the product of a depictive process.*

This needs elaboration.

Current Perspectives on Five Axioms

Following the introspective route, we identify how being conscious may be broken down into distinct experiences. Five of these were identified some years ago (Aleksander and Dunmall, 2003), and the list has not changed. Clearly being conscious of being at a concert at the Albert Hall is a different experience from waking up in the morning and trying to work out what to do during the day. Here we introduce this intuitive way of dividing conscious experience as brief statements of the following five axioms (stated in the first person to underscore their introspective nature), leaving it until later to elaborate their mechanisms:

1. Presence: I feel that I am an entity in the world that is outside of me.

This appears to be the most fundamental conscious experience one can have: the feeling of being an empowered, active agent in a real sensory world. In previous descriptions of the axioms we referred to this sensation as *Perception*, but this is now seen as an understatement. Presence in the world is perception plus a sense of *being*. Although the visual sensory modality is often used to illustrate this sensation, hearing and touch too are implicated. Lack of one or two of these can be compensated by the remaining modalities.

2. Imagination: I can recall previous sensory experience as a more or less degraded version of that experience. Driven by language I can imagine experiences I never had.

Sometimes this is called memory. But that is an understatement of the sensation which allows me to have a rich imagery of the world as once experienced through axiom 1. It also allows me to create scenarios that might have been experienced or even ones that are not close to reality (say the world of Harry Potter).

3. Attention: I am selectively conscious of the world outside of me and can select sensory events I wish to imagine.

Attention has conscious and unconscious components. Sometimes it feels as if what I choose to experience is automatic. For example a bright flash will attract my attention leading to the subsequent experience of the firework patterns that go with it. A sound bang can do the same. At other times it feels as if I set the target of my attention more purposefully. For example if, on seeing a motorcar, I wish to identify its make, I purposefully shift my gaze to the point on the bonnet, the wheels or the rear where an emblem is expected to be found.

4. Volition: I can imagine the results of taking actions and select an action I wish to take.

It feels as if I consider options for making a choice sequence. Previously this axiom was called 'planning'. So here too we have chosen a more appropriate term so as to avoid the implication that the mechanism is one resulting from a rule-based planning process as found in AI. This sequential traversal of options must eventually alight on a decision based on likes and dislikes, that is, it involves emotions.

5. Emotion: I evaluate events and the expected results of actions according to criteria usually called emotions.

As suggested above, what we call emotions guide our decision making. It can also be the case that, sometimes when it is difficult to make a selection, I recognize that an arbitrary decision is made. In the next section we give an example of decision making in a restaurant which stresses the mechanistic link between emotion and volition.

Mechanistic Use of the Axioms: The Kernel Architecture

So far we have merely identified what seem to be some important ways in which we can highlight sensations that appear to be different in kind. But this is merely the first step towards a more important aim: to have an understanding of distinct mechanisms that may be

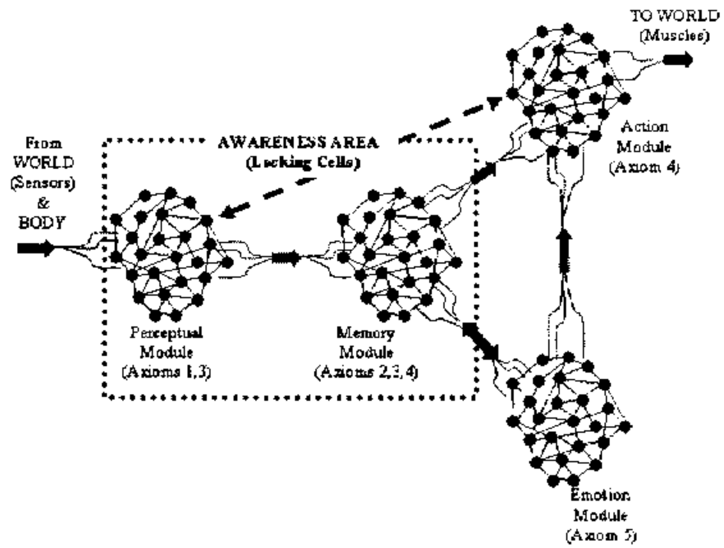

Figure 2. The Kernel Architecture.

necessary to support these distinct sensations. Once such mechanisms have been derived, we ask how they *interlock* to provide a sensation of a coherent consciousness. To this end, we suggest a kernel architecture (Figure 2, first introduced in Aleksander, 2005). We extend this here to discuss the way that the axiomatic mechanisms interact during specific conscious acts.

1. Presence

The key necessary mechanism here is *depiction* which has been discussed earlier. That is, the output of a neuron that is going to be representative of where an element of the sensory world is located must not only react to that element as a sensory signal, but must also be indexed by signals that position the element with respect to some frame of reference. Such an indexed neuron becomes intimately connected with the event in the world in relation to other such neurons while its absolute position in the 'brain' network becomes less important.

For example, it is known that an elemental event in the world such as a small moving green triangle stimulates representational neurons in different parts of the visual cortex: a colour area, a motion area and a shape area. How this *integrates* to provide a coherent sensation has always been a problem for neurophysiologists. This is the 'binding' problem. In accepting that neurons become depictive, that is, indexed by world events, the binding problem disappears. The three separate neurons, being indexed to the same world event will depict colour, motion and shape. But they each position their contribution in the same element of the world and hence overlap in consciousness.

In the KA it is the 'perceptual' module that implements the support mechanisms for Axiom 1, while depiction is obtained through the dotted connection to the action module, hence this *interaction* is essential for a coherent sensation.

2. Imagination

The simplest imaginational act is to look at an object on one side of the room (say the left), then switch attention, by moving one's head, to another object somewhere else in the room (say the right), and then close one's eyes. It becomes possible to recall, say, the object on the left (call it A), then the one on the right, (B). The ability to recall suggests the presence of a dynamic neural system for which A and B are different attractors. In the KA this is labelled 'The Memory Module'. During perception, the depiction in the perceptual module acts as input to the memory module. A version of this input *Ai'* also becomes a state of the memory module through a process known as Iconic learning (see Aleksander, 1996). In simplified terms, iconic learning is the transfer of input pattern *Ai* to the state variables to become *Ai'* through the following learning sequence. The left of the equations below is the input to the memory module and the right is the resulting or taught state.

The format therefore is *(state, input → next state)*.

$$(R, Ai(t)) \rightarrow Ai'(t+1)$$

where R is a random initial state of the memory module

$$(Ai'(t+1), Ai(t)) \rightarrow Ai'(t+2).$$

In a further step

$$(Ai'(t+2), S) \rightarrow Ai'(t+3)$$

where S is a return of the input to some neutral, 'eyes closed' state.

Ai and *Ai'* are a depictive like the 'fly on the left' depiction in Figure 1. *Ai'* becomes the depictive attractor-memory of *Ai*. The same occurs for *Bi*.

It should be noted that *Ai* and *Ai'* need not be exactly the same, but are expected to be similar. After looking at something and closing our eyes, we commonly experience memory that is less precise than the perception. This depends on the properties of the network. We note that sequences of *Ai* can lead to sequence attractor trajectories in the memory network. So, in general terms, the *interaction* of the memory

and perception networks in the KA form a system for laying down experience. This experience has a state structure (i.e. state graphs) in the state space of the memory automaton.

3. Attention

The lowest level of attention is not shown in the KA. This is largely the unconscious process that causes eyes (or hands, say in the case of the blind) to move to salient parts of the world as mediated by sensory input. In the case of vision in the brain, the key mechanism is that of the superior colliculus which moves the eyes on the basis of its own representation of the retinal image. However, in the above example of looking left and right, to remember what is on the left and right is an instance of *purposeful* attention. How does the mechanism of 'thinking' about what is on the left and what is on the right work in the KA?

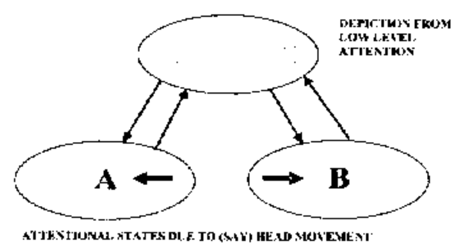

Figure 3. Attention for indexing imagined events.

Simplifying things to the extreme, Figure 3 shows the states of the perceptual automaton entered as a result of muscular action. The top state showing both A and B is a somewhat hazy depiction that might be created by unconscious eye movement, but one that requires a more direct look to identify the objects at the left and then at the right. The arrow leading to the A state is caused by purposeful movement, say, of the head (what causes it is not important here). Being depictive, the neurons that record this focus-on-A state will also be indexed by the left head movement. This has just been indicated as a thick arrow. Similar events take place on the B side. This exact state structure, including the motion sensations, is iconically transferred to memory automaton which, due to the generalization mentioned above, can explore it *internally* without input from the perception module. Were the system to be endowed with language, the memory automaton has all the necessary information to drive language that says: 'if I attend to the left object, I clearly remember an A'. This, then, is an example of inner attention shown to be dependent on the *interaction* of the perception, memory and action modules.

4. Volition and 5. Emotion

While these axioms involve two separate mechanisms, it is helpful to see them as they work together. All four modules of the Kernel Architecture interact to create a state structure as in Figure 4. This represents state activity as seen from the perspective of the memory module and we refer to a very simple scenario of being in a restaurant, looking at the menu.

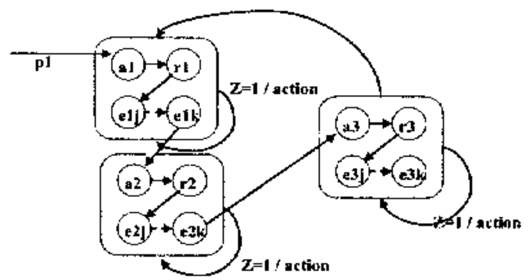

Figure 4. A state diagram for an example of volition and emotion.

The state of the Ax1 module is labelled *p1* and it represents a depiction of the menu. This acts as an input to the memory module where it causes first a memory of the first item *a1* (say a pizza). This, in turn, recalls the experience of eating the pizza *(r1)* and the emotions that accompanied the experience (iconically absorbed from previous experience in the emotion module *(e1j,e2k . . .)* where these may be akin to pleasure in the taste of the food, guilt (for eating fattening food) etc. Having dealt with the pizza, the memory module repeats the exercise for other items on the menu in some kind of succession. It is noted that no action need occur during this contemplative activity. What is changing, however, is some cumulative 'wantedness' value in the emotion module related to each of the felt emotions. This value could be positive or negative, but it is assumed that the emotion

module is capable of accumulating a wantedness value for each of the contemplated food items. If this value exceeds some threshold it triggers a signal Z from 0 to 1 and this does two things: first it discontinues the search through the items and causes the wanted event to be held in the memory module; second, it activates the action module that causes whatever action needs to be performed (calling the waiter, say) to obtain the wanted food.

Of course it is quite possible that the Z parameter is never triggered. This, in the strongest version, would be described as a pathological 'freezing' of the organism. Weaker versions might be described just as 'not being able to make up one's mind'. However, it is quite possible, certainly in principle, for there to be some source of noise which is additive to the wantedness value and which causes Z to be triggered. Introspectively this corresponds to a feeling of 'I'll make an arbitrary choice'.

KA Implication 1: Is a Given Organism Conscious?

Here we begin to look at the KA and the axioms as tools with which to enter important discussions about consciousness. The first is as a set of tests of consciousness.

> *Assertion 3: There is a* prima facie *case for an organism to be conscious if it has a system that is structurally isomorphic with the Kernel Architecture to the extent that it can support the axiomatic mechanisms.*

Before applying this, it is most important to stress that this is a *necessary* condition for consciousness to be supported. It is not sufficient — the KA does not disappear if the organism becomes unconscious by going to sleep. It also provides no assessment of *what* the organism might be conscious. Clearly a virtual machine based on the kernel architecture and I might in abstract principle share some mechanical properties that support our conscious experiences, but such experiences need not be the same. So we here separate the content of consciousness from the mechanisms that are necessary for there to be any content at all. Therefore, the test suggested by Assertion 3 looks for the potential support of consciousness without reference to content.

KA Implication 2: Animals and Higher Order Thought

Assertion 3 above allows us to approach the question of animal consciousness. This arises not only in the context of animal rights and ethics, but also in the scientific context of exercising emerging theories of consciousness. We do not intend to make judgments about ethics, but merely to show how some puzzling issues might be clarified. First, the axiomatic/kernel approach aligns closely with the central features of the work of Crick and Koch (1998) in the sense that they have argued a separation between unconscious input sections of the visual system from 'exrastriate' areas that contribute to consciousness due to their connections to motor regions. Of importance is the fact that the brain maps used for the Crick/Koch work were those of a macaque monkey. It would be inconsistent then to argue that if an animal's brain tells us about the neural correlates of consciousness, that such an animal is not conscious. This is extended to the suggestion that finding kernel architectures in brains is concomitant with being potentially conscious, as this is what kernel architectures do. It is of some interest that kernel architectures may clearly be seen in invertebrates such as bees, but perhaps not in amoebae.

Another argument sometimes used against animal consciousness is that of Higher Order Thought — HOT (Rosenthal, 1993). According to this, a sensation becomes conscious only if it is accompanied by a higher level awareness that the perception is happening. That is, in HOT theory an organism has conscious thoughts only if they occur at both these two levels. Some have argued that animals do not need or in any way exhibit this distinct higher level of thought and are therefore not conscious.

Briefly, to see how axiomatic theory addresses the HOT issue, we need to understand how 'I see a chair' may be distinguished from 'I am conscious that I am seeing a chair'. While in this paper we have not concerned ourselves with language, it is possible to imagine, in using the KA to model the human case, that there needs to be a decoder of activity in the KA that constructs appropriate utterances in language machinery. In uttering 'I see a chair' the system would be signalling that perception machinery and memory machinery are depicting a chair (say labelled as such in the memory machinery). In the case of the more complex sentence, 'I am conscious that . . .' may be taken to be a decoding of activity in the entire KA in addition to a time-separated decoding of the content of the two awareness areas. In animals, both local and global KA activities are bound to be ongoing, but without language these would not be expressed or differentiated. Taken to its conclusion, this supports the notion that HOT is an artefact of language and not a definition of consciousness. The nature of language machinery necessary to express the contents of the KA is the subject of present studies.

Conclusion: The Legitimacy of First-Person Modelling

The content of this paper may be summarized as confirming that a modelling process can start with an introspectively derived set of axioms and finish with an implemented model that creates for itself depictions of self in the real world. These have the same role in the model as have the first-person sensations to which we refer when we say we are conscious. As an aside, this allows us to compare computer-displayed depictions with our own and discuss the sensitivity of conscious experience to architectural parameters in a 'Kernel Architecture'. This provides tests for the consciousness of an organism which here have allowed us to assert that animals are conscious in virtue of being equipped with a Kernel Architecture. Higher-Order Thought theories then turn out to be about our ability to attend and translate into language the activity of different parts of the KA without requiring additional higher-order machinery.

References

Aleksander, I. (1996), *Impossible Minds: My Consciousness, My Neurons, My Consciousness* (London: Imperial College Press).

Aleksander, I. (2005), *The World in My Mind My Mind in the World: Key Mechanisms* (Exeter: Imprint Academic).

Aleksander, I. and Dunmall, B. (2003), 'Axioms and tests for minimal consciousness in agents', *Journal of Consciousness Studies*, **10** (4–5), pp. 7–19.

Aleksander, I. and Morton, H.B. (2006), 'Phenomenology in computational models of consciousness', *International Journal of Pattern Recognition and Artificial Intelligence*, **20** (6), pp. 785–96.

Chalmers, D.J. (1996), *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).

Crick, F. and Koch, C. (1998), 'Consciousness and neuroscience', *Cereb. Cortex*, **8** (2), pp. 97–107.

Franklin, S. (2003), 'IDA, a conscious artefact?', *Journal of Consciousness Studies*, **10** (4–5), pp. 47–66.

Galletti, C. and Battaglini, P. (1989), 'Gaze-dependent visual neurons in area V3A of monkey prestriate cortex', *Journal of Neuroscience*, **6**, pp. 1112–25.

Holland, O. (ed. 2003), *Machine Consciousness*, Special issue of the Journal of Consciousness Studies, **10** (4–5) (Exeter: Imprint Academic).

Koch, C. (2001), 'Final report of the workshop "Can a machine be conscious?" ', The Banbury Center, Cold Spring Harbor Laboratory, 13–16 May 2001 (http://www.swartzneuro.org/abstracts/2001_summary.asp)

Rosenthal, D. (1993), 'Thinking that one thinks', in *Consciousness*, ed. M. Davies and G. Humphreys (Oxford: Blackwell), pp. 197–223.