Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Neural Networks 20 (2007) 929–931
Neural Networks
www.elsevier.com/locate/neunet

Introduction

Introduction to the special issue on 'Brain and Consciousness'

It is now ten years since two of us (WF and JGT) edited a Special Issue of this journal on 'Neural Network Models of Consciousness' (**Freeman & Taylor, 1997).** Much has happened in the intervening ten years in this area of study, caused both by the enormous increase in experimental results relevant to any study of brain and mind and also by the creation and further development of more complex models claiming to be able to explain how consciousness arises through suitable brain activity. One question to ask is therefore: how much have we really moved forward in our study of consciousness? It is clear that we have recognised ever more fully as time has passed how consciousness presents a very deep and difficult problem. The earlier optimistic view presented in the papers in (**Freeman & Taylor, 1997**) has been replaced by a more cautious approach. That is clear from the papers of this Special Issue.

We start the Special Issue with papers of a more theoretical nature. The first paper (**Aleksander & Morton, 2007**) sets out to formulate some deep principles for the creation of consciousness in an information processing system. It develops further the Axiomatic Consciousness Theory (ACT), already published by the authors elsewhere, to construct an approach which is later in the paper implemented by a simulation. The basic axioms of ACT involve inclusion of the five components of presence, imagination, attention, volition and emotion. These are implemented in an iconic neural state machine, which is activated by external stimuli in a manner similar to the external stimulus. Further components are considered as part of the simulation of the development of sensations associated with eye movements.

We continue with an approach to consciousness through control theory of (**Sanz, López, Rodríguez, & Hernández, 2007**). A cognitive system is defined in the paper as one which exploits internal models of other systems. This leads to system awareness, defined as the continuous updating of internal models of other systems. Consciousness is then suggested as arising in a given system through the updates of internal models that the system has created of itself. This leads to five axioms for the properties of the system enabling it to be expected to possess consciousness, these axioms being very close to those presented in the previous paper, the ACT axioms.

The third paper of this special issue, by (**Coward & Sun, 2007**), is a valuable and salutary approach to consciousness. It considers the general nature of scientific theories and in particular emphasises the hierarchical nature of the most effective scientific theory we have presently, that of the physical world. The authors then argue that the brain itself has a hierarchical character in its processing, so should be able to be handled in a similar manner. After a discussion of possibilities in approaching consciousness in this way, critiques from this approach are made of several current approaches to consciousness.

A paper is then presented (**Baars & Franklin, 2007**) in which the influential 'Global Workspace' (GW) idea presented earlier by one of the authors (BB) is put into a simulation of commercial value. This is the IDA system, an interesting workspace for testing the GW hypothesis. IDA tackles the problem of assessing sailors entering the navy to determine what sort of job they would be most suited to. It is used as a platform for the various mechanisms that have been proposed as components of the GW approach. As such the IDA system is a valuable foundational neural network simulator enabling further ideas on consciousness and more general brain processing to be tested out by detailed simulation.

The fifth paper in this Special Issue (**Rolls, 2007**) brings us to a more directly brain-based approach to consciousness, at a global level. The paper begins with a careful discussion of the nature of coding in temporal cortex (which shows that synchronisation of neuronal activity is not necessarily involved) and that the threshold for awareness of a given stimulus in temporal cortical neurons is higher than that for activity outside awareness (to be regarded as possible noise). These results and others are then used to support a higher-order syntactic thought (HOST) approach to consciousness. Such a higher-order approach is seen as natural when a multi-step plan needs to be corrected; the emergence of qualia is seen as secondary to such a higher-order correction system.

The next paper (**Taylor, 2007a**) also uses a brain-based hierarchical approach, and also employs engineering control ideas related to attention (relating back to Sanz et al.). The construction of the model (termed CODAM) uses results of brain imaging and single cell experiments on attention. A detailed model of attention is built extending to attention control models of motor response as well as ballistic models of attention used by modellers. The extra feature in CODAM is the COrollary Discharge of Attention Movement (so the

0893-6080/$ - see front matter © 2007 Published by Elsevier Ltd
doi:10.1016/j.neunet.2007.09.010

acronym). CODAM allows exploration of dynamics to see how consciousness might arise from attention movement, and resulting properties of such consciousness are explored. In particular an explanation is given of the important 'immunity to error through misidentification of the first person pronoun'; explained by an 'ownership' signal (the corollary discharge) creating the sense of 'I'.

The seventh paper (**Taylor & Fragopanagos, 2007**) addresses the problem of the relation between attention and consciousness, which is presently somewhat controversial. Various successful simulations are presented of the results of critical paradigms (such as the attentional blink and change blindness) and others are discussed which have been claimed to indicate the independent nature of attention and consciousness. The simulations are created inside the CODAM attention control framework of the previous paper, and thereby imply that the claimed separation of attention and consciousness need not be well-founded. Attention is implied, from this study, still to be the gateway to consciousness.

The following paper (**LaBerge & Kasevich, 2007**) continues with the direct brain-based approach, but now looking more closely at the substrate used in brain activity, that of the neurons in the cortex. The paper presents a detailed analysis of the structures and dynamics of cortical neurons, and proposes that it is in the neural activity in the apical dendrites of cortical neurons that consciousness finds its crucial dynamics. The flow of activity vertical to the cortical surface (in thalamocortical loops) is proposed as being suitably continued in time so as to support consciousness creation. Such creation is further aided by the associated electromagnetic fields with this activity flow, so leading to a field-theoretic basis for consciousness (see also ideas presented in the following paper).

A further direct brain-based approach is continued in the paper of (**Freeman, 2007**), which is based on his pioneering measurements of cortical activity in rabbits, cats and man using multiple surface electrode recordings. He considers how a general dynamical systems approach to brain activity especially that associated with motion towards attractors using synchronised neural activity, allows a detailed grounding of the reflex arc (at the base of all responses) in the brain. This dynamical systems approach leads to a picture of consciousness as a field phenomenon which is suggested even to extend both sub-atomically and socially to interactions between individuals.

Paper ten (**Cleeremans, Timmermans & Pasquali, 2007**) starts from the Higher-Order Thought approach to consciousness, and considers possible computational mechanisms that can implement it. The simulations are performed of two interacting networks: one, a first-order network trained to perform a simple categorization task being input to a second-order network trained as an encoder, which thus 'observes' the states of the first-order network and reproduces these states on its output units. This is proposed as the beginning of a computational mechanism to account for mental attitudes, that is, an understanding by a cognitive system of the manner in which its first-order knowledge is held. Consciousness, in this approach, is knowledge of the geography of one's own internal representations, learned over time in terms of the relative importance of its components to the agent embedded in a physical world.

The following paper (**Grossberg, 2007**) returns to a brain-based approach by way of Adaptive Resonance Theory (ART) and the conjecture that "all conscious states are resonant states". The article reviews theoretical considerations of this approach and some of the rapidly growing body of behavioral and brain data supporting it. The article summarizes ART models that predict functional roles for identified cells in laminar thalamocortical circuits. These predictions include explanations of how slow perceptual learning can occur without conscious awareness, and why oscillation frequencies in the lower layers of neocortex are sometimes slower beta oscillations, rather than the higher-frequency gamma oscillations that occur more frequently in superficial cortical layers. ART traces these properties to the existence of intracortical feedback loops, and to various detailed reset mechanisms.

The final paper (**Herzog, Esfeld, & Gerstner, 2007**) 'sets the cat among the pigeons', so to speak. It indicates that most of the present neural models of consciousness on the market suffer the fate of the 'small network' paradox: they all imply that consciousness should be present in small networks (composed of no more than ten or twenty neurons). Such a result can be taken to read that there is either something missing from all of these models (so that large number of neurons are needed to implement them) or that these small network systems are indeed conscious. We refer to the commentary of (Taylor) in this issue for a further discussion of the implications of this result.

Given this set of papers, and especially that of (Herzog et al.), we need to take stock of where the study of consciousness as a brain-based phenomenon is going. Of course the set of papers presented here is only a small indication of the overall strength of the subject. For example we direct any interested reader to the proceedings of the annual Association for the Scientific Study of Consciousness conference, and to the journals such as the 'Journal of Consciousness Studies' and 'Consciousness and Cognition' to get a broader perspective on present-day consciousness studies. However the present papers provide, we claim, a good perspective on the possible neural network-based models that are 'cutting edge' in consciousness studies. They cover a broad range of the extant models and indicate the progress that modellers are making to face up to the challenge of consciousness for computational neuroscientists. So how far do they go to indicating real progress is being made? Are there any real breakthroughs?

We have given separately a commentary of the paper of (Herzog et al., 2007) due to its surprising and important result; it clearly requires careful treatment, as the commentary (Taylor, 2007) implies. Let us turn therefore to address the question just raised: are there any breakthroughs on the modelling of consciousness since our 1997 Special Issue? We have to admit that there are not, but the papers in this Special Issue show that steady progress has occurred in the subject over the last decade. The experimental side has provided greater clarity as to the various sorts of brain activity that are concomitant with the creation or otherwise of consciousness by different inputs. It is these that must be used to test any proposed neural model of consciousness. This experimentally-based modelling approach is well-reflected especially in the papers of (Freeman, 2007; LaBerge, 2007: Rolls, 2007; Taylor & Fragopanagos, 2007) in this issue. The theoretical analysis of models of consciousness and of consciousness itself, have also advanced to greater depths as evinced by the papers of (Coward & Sun, 2007; Aleksander & Morton, 2007; Herzog et al., 2007).

There are two commentaries to conclude the Special Issue. The first **(Cooper, 2007)** considers the problems presented by consciousness when looked at through the eyes of someone trained in the physical world. He eloquently points out the need to explore first the possible (but very subtle) mechanisms that could be used by suitable aggregates of neurons to create conscious experience. Only if this attempt to give a neural basis for consciousness fails should we then consider the second step, that of assuming a separate consciousness 'component' to matter, distinct from the quarks and other components which are now known to make it up. He finishes with a call for leaps of imagination to make the second step superfluous.

The second commentary **(Taylor, 2007b)** discusses in more depth the implications of the small network argument of (Herzog et al.) in this Issue. These implications are discussed from various angles, especially an evolutionary one. There is also the need for lower level encoding of input stimuli, thus increasing the number of neurons needed for any consciousness system. However if this encoding and the associated consciousness networks are sparse then it is expected that there may indeed be conscious experience, but of a very 'thin' character.

Where do we go from here? More of the same or do we see the need for a completely new direction? This covers a new type of criterion applied to models, as given, for example, by the 'small network' argument of (Herzog et al., 2007). It would also cover a new type of mechanism for creating consciousness from brain activity. Whilst we do not have any such completely new contributions the theoretical analyses contained in our special issue provide useful constraints on the form any such new development would need to take.

We seem to be in the middle of the longer haul, gradually climbing up to the peak of the mountain of consciousness. We are not there yet, but the contributions in this issue add to the general feeling that real progress is occurring and that we are mounting ever higher towards the peak through our endeavours.

References

Aleksander, I., & Morton, H. (2007). Phenomenology and digital neural architectures. *Neural Networks*, *20*(9), 932–937.

Baars, B. J., & Franklin, S. (2007). An architectural model of conscious and unconscious brain functions: Global Workspace Theory and IDA. *Neural Networks*, *20*(9), 955–961.

Cleeremans, A., Timmermans, B., & Pasquali, A. (2007). Consciousness and metarepresentation: A computational sketch. *Neural Networks*, *20*(9), 1032–1039.

Cooper, L. N. (2007). On the problem of consciousness. *Neural Networks*, *20*(9), 1057–1058.

Coward, L. A., & Sun, R. (2007). Hierarchical approaches to understanding consciousness. *Neural Networks*, *20*(9), 947–954.

Freeman, W. J. (2007). Indirect biological measures of consciousness from field studies of brains as dynamical systems. *Neural Networks*, *20*(9), 1021–1031.

Grossberg, S. (2007). Consciousness CLEARS the mind. *Neural Networks*, *20*(9), 1040–1053.

Herzog, M. H., Esfeld, M., & Gerstner, W. (2007). Consciousness & the small network argument. *Neural Networks*, *20*(9), 1054–1056.

LaBerge, D., & Kasevich, R. (2007). The apical dendrite theory of consciousness. *Neural Networks*, *20*(9), 1004–1020.

Rolls, E. T. (2007). A computational neuroscience approach to consciousness. *Neural Networks*, *20*(9), 962–982.

Sanz, R., López, I., Rodríguez, M., & Hernández, C. (2007). Principles for consciousness in integrated cognitive control. *Neural Networks*, *20*(9), 938–946.

Taylor, J. G. (2007a). CODAM: A neural network model of consciousness. *Neural Networks*, *20*(9), 983–992.

Taylor, J. G. (2007b). Commentary on the 'small network' argument. *Neural Networks*, *20*(9), 1059–1060.

Taylor, J. G., & Fragopanagos, N. (2007). Resolving some confusions over attention and consciousness. *Neural Networks*, *20*(9), 993–1003.

J.G. Taylor*
Department of Mathematics, King's College, Strand, London, UK
E-mail address: john.g.taylor@kcl.ac.uk.

W. Freeman
Department of Molecular & Cell Biology, University of California at Berkeley, Berkeley, California, USA

A. Cleeremans
Cognitive Science Research Unit, Universite Libre de Bruxelles, Brussels, Belgium

* Corresponding editor. Tel.: +44 0207 848 2214; fax: +44 0207 848 2017.