

# Replication of the Hard Problem of Consciousness in AI and Bio-AI: An Early Conceptual Framework

Nicholas Boltuc <sup>1</sup>, Piotr Boltuc <sup>2</sup>,

<sup>1</sup> Bioengineering Department, University of Illinois Urbana Champaign

<sup>2</sup> Philosophy Department, University of Illinois Springfield

## Abstract

AI could in principle replicate consciousness (H-consciousness) in its first-person form (as described by Chalmers in *the hard problem of consciousness*.) If we can understand first-person consciousness in clear terms, we can provide an algorithm for it; if we have such algorithm, in principle we can build it. There are two questions that this argument opens. First, whether we *ever* will understand H-consciousness in clear terms. Second, whether we can build H-consciousness in inorganic substance. If organic substance is required, we would need to clearly grasp the difference between building a machine out of organic substance (Bio-AI) and *just* modifying a biological organism.

## The main thesis

Computers could in principle be equipped with Hard Consciousness that satisfies Chalmers' *hard problem*<sup>1</sup>. This thesis is to be called **the extra strong AI thesis**. In order to better understand the thesis, let us define the hard problem of consciousness. Conscious experience has a subjective quality, which can be thought of as *what it is like to be a given being*. Explaining this quality is termed **the hard problem of consciousness**. The hard problem is to be distinguished from the easy problem of consciousness, namely the issue of neural correlates of various conscious states [Chalmers 1995, 1997].

In order to present the extra strong AI thesis we start with the standard distinction between weak and strong AI. Within strong AI we distinguish between strong and extra

strong AI; for this we use a regulative definitions and clarify the distinction between consciousness and cognition.

Let us define Weak AI as design and construction of machines that simulate consciousness or cognitive processes closely associated with consciousness.

Let us define Strong AI as design and construction of conscious machines.

Note that conscious machines are those that reason in the way that conscious beings do, which requires them to satisfy the reductionist definition of conscious, but it does not require any first-person consciousness of the kind that many philosophers ascribe to living beings [Chalmers, Nagel, Harman] but others deny [Churchland, Dennett]. Hence, we should distinguish Strong AI from Extra Strong AI.

Let us define Extra Strong AI [ES-AI] as design and construction of conscious machines that satisfy Chalmers' hard problem of consciousness (H-consciousness).

Note that in order to address the issue of ES-AI one has to accept non-reductionistic views of consciousness. For consciousness-reductionists [Churchlands, Dennett] there is no distinction between Strong AI and Extra Strong AI since they deny the *differentia specifica* between the two, namely the occurrence of non-reductive first-person consciousness.

It is important to bring this largely philosophical issue to AI audience since different philosophers that charge into the realm of AI [Chalmers, Dennett, Churchland] import vastly different philosophical assumptions. Those assumptions have far reaching implications insofar as the formulation of strong AI goes.

## The argument

The thesis is that computers could in principle be equipped with Hard Consciousness that satisfies Chalmers' *hard problem*. This can be also formulated as the claim that extra strong AI (ES-AI) in principle can be built. This thesis relies on the following premises:

- a. If we understand consciousness (as in *the hard problem*) then we can present an algorithm<sup>2</sup>, *because* what

---

<sup>1</sup> Chalmers' approach, at different stages of its development, involves several detailed claims (for instance the separation of phenomenal consciousness from the functional roles, or a possible implication that we do not know what a solution of the hard problem of consciousness would look like) that I do not need to subscribe to in order to define my objective in this paper. Any non-reductionistic approach to phenomenal consciousness, such as Nagel's *first person point of view* or Harman's appropriation of the concept of *Das Ferstehen*, will do. To put it differently, the claim is that if we knew how exactly phenomenal consciousness works in a brain, we could, in principle, engineer one.

---

<sup>2</sup> By an algorithm we understand an executable function (both in terms of programming and an algorithmic description in natural science, e.g. chemistry or physics). This is more than a model; the

it is to understand something in scientific terms is to present its algorithm.

- b. If we have an algorithm that performs a certain function in a biological world, we can, in principle, engineer a machine such as to follow such algorithm and thereby perform such function<sup>3</sup>.

*Criterion b. satisfies at least the Weak AI Criterion*

- c. Consequently, if the function of a given algorithm is to produce consciousness (as in *the hard problem*) then if we inlay such an algorithm in matter it produces consciousness.

Hence, once we understand consciousness (as in *the hard problem*) in scientific terms, we can in principle engineer such consciousness in a machine. Let us reemphasize that the issue is *obviously* not of interest, and in fact empty, to those who deny that there is an essential difference between computation and H-consciousness (or any phenomenal consciousness).

In the next two sections we discuss details of this argument. First, we examine premises a. and b., including some of the science behind them; more work needs to be done in this area. Next, we examine the distinction between stimulus-response, cognition, consciousness (in general) and H-consciousness. This last notion is given a closer look. We question its all or nothing character. We also point out to the distinction between consciousness (as in *the stream of consciousness*) and the objects of consciousness. In the last two sections we sketch out some potential objections and implications of this argument.

## Discussion of the premises and the science behind them

Premises a. and b. may require the further clarification.

In reference to premise a. let us say the following: Understanding any kind of events in scientific terms, strictly speaking, consists in providing an algorithm for those events; here we ascribe to van Fraassen's account of science without further proof [van Fraassen 1990]. Let us apply this definition of understanding to the issue of H-consciousness. While we do not have a generally accepted scientific theory of consciousness, and of H-consciousness in particular, there are a few competing theories of consciousness, which we mention for the sake of example so as to demonstrate that such an understanding could in principle be gained.

Let us visit the *dialogue of hemispheres* theory. Whatever its scientific value shall turn out to be it is one way of explaining how non-reductive theories *could be* and *may very well be indeed* compatible with non-reductive metaphysical views on consciousness. The focus of this approach is on the temporal difference between two functional systems in our brains. It is called **asynchrony of con-**

---

latter provides some level of understanding but would be insufficient to justify claim b.

<sup>3</sup> Our view does not imply fundamental computationalism since the algorithm can be understood in broadly operational terms.

**sciousness.** Essentially it states that the interaction between two parts of the brain, functionally or geographically defined, is what creates consciousness (as defined by Crick).

According to A. Bartels and S. Zeki "the visual brain ... shows that different attributes of the visual scene are consciously perceived at different times. This temporal asynchrony in visual perception reveals a plurality of visual consciousnesses that are asynchronous with respect to each other, reflecting the modular organization of the visual brain." They "further hypothesize that when two attributes (e.g. color and motion) are presented simultaneously, the activity of cells in a given processing system is sufficient to create a conscious experience of the corresponding attribute (e.g. color), without the necessity for interaction with the activities of cells in other processing systems (e.g. motion). Thus, any binding of the activity of cells in different systems should be more properly thought of as a binding of the conscious experiences generated in each system."

There are also other theories which try to establish a single part of the brain as the locus for consciousness. As early as the 1800s neurologists (John Hughlings Jackson) knew that removing the diencephalon (which includes the thalamus) from the brain results in a loss on consciousness. Others investigated recovery from a 'persistent vegetative state' (wakefulness without awareness). They found that overall cortical metabolism remained almost constant during recovery but that the metabolism in the prefrontal and association cortices became correlated with thalamic and precuneus activity [Laureys et. al 2002]. Without thalamic activity consciousness did not develop and cortical activity was reduced.

Another theory in consciousness formation is the reticulo-thalamo-cortical system [Herzog 1995a, b]. It states that consciousness is formed by the integration of three different stimulus processing activities. Other integration theories of consciousness include, but are not limited to, the global workspace theory [Baars 1994, 2005], information integration theory [Tononi 2005], and recurrent processing [Lamme 2004]. The details of these are still uncertain and being researched, however we are ever closer to the gold standard [Lamme 2006].

Any of the above theories of consciousness, if satisfactorily presented, would need to provide the equation describing the process of consciousness creation. For instance the theory that consciousness (H-consciousness) is created in the thalamus would have to present the exact consciousness-producing process that occurs in the thalamus and all relevant conditions for this process to occur. We are far from having such theory but two things seem obvious: first, that science progresses, though slowly, towards gaining an understanding of consciousness; second, that a scientific understanding of consciousness would require casting it in terms of equations (algorithm).

One more point needs to be made in regard to premise a. There are independent reasons within neuroscience to doubt the whole neural correlates approach and hence Chalmers' distinction between the hard and the easy problem of consciousness, which is relevant for the argument in the next

section. However, this does not seem to be the dominant view at the moment. Many are convinced that there are clear neural correlates of consciousness (NCC) as defined: An NCC (for content) is a minimal neural representational system N such that representation of a content in N is sufficient, under conditions C, for representation of that content in consciousness [Chalmers, 2000]. Francis Crick, the father of modern DNA research goes so far as to say "we...need to discover the neural correlates of consciousness.... For this task the primate visual system seems especially attractive.... No longer need one spend time attempting ... to endure the tedium of philosophers perpetually disagreeing with each other. Consciousness is now largely a scientific problem" [Crick, 1996]. However despite research on NCCs [Baars 2002, Blake 2001] there are still doubts about the matching-content doctrine and NCCs. There is no reason to think that the neural states that have been shown experimentally to be correlated with conscious visual experiences match those experiences in content" [A. Noe and E. Thompson 2004 p. 4]. Hence, the issue is open to further challenges; it is not easy to predict whether such a challenge, if successful, would have us reformulate or abandon the thesis of this paper.

In reference to premises b. let us say the following: If we have a model or any algorithm that performs a certain function we can instantiate it in various substances (this is the multiple instantiation thesis). Hence, we can, in principle, engineer a machine such as to follow this algorithm and thereby to perform such function. Hence, if we have an algorithm for the process of H-consciousness it becomes a technical issue in what substances such algorithm can be engineered. There are some reasons to believe that the process such as H-consciousness could be engineered in non-organic media, such as silicone based computers. There are also reasons why this may be possible solely in organic substance. This may be true since certain reactions must involve organic carbon compounds, and they might indeed form an important basis in consciousness creation.

Should the organic substance be required in order to build h-consciousness, we would need to clearly grasp the difference between building a machine out of organic substance (Bio-AI) and *just* modifying or growing a biological organism. This interesting problem goes beyond the scope of the present paper.

Instantiability only in organic media would not be destructive to the extra strong AI thesis, unless this required bio-modification, since bio-computing, and bio-AI, are just one more engineering alternative; we say a little more about it when addressing objection B. below. We will be clearer on whether bioengineering is required for extra strong AI only when we are closer to getting the H-consciousness equation.

### Clarifications of certain terms and the philosophy behind them

The standard distinction between consciousness and cognition [Crick] helps clarify the distinction between weak and

strong AI; by adding H-consciousness we distinguish between strong and extra strong AI. This section is devoted to a slower examination of those concepts.

Weak AI systems may have cognitive functions (cognition) but only Strong AI systems have conscious functions; extra strong AI systems would have first-person consciousness (H-consciousness). Hence, we shall define cognition, consciousness and various levels of H-consciousness.

Cognition can be defined as a set of mental functions such as the ability to think, reason, and remember. Cognition can also be defined, more broadly, without reference to mental functions, as processing of sensory input and its assimilation into existing schemes. The latter definition is common to human cognition and to AI.

Applied to organisms (biological and robotic) cognition can be defined as "processes by which an organism gains knowledge of or becomes aware of events or objects in its environment and uses that knowledge for comprehension and problem-solving."<sup>4</sup> Even in this somewhat *zoocentric* definition cognition does not presuppose consciousness for the following reasons: 1. It makes sense to talk about unconscious cognitive processes (the way psychoanalysis does); 2. It is natural to ask whether a human, or a simpler animal, is conscious while engaging in those processes; the alternative is a *zombie* hypothesis. This is a different question from whether this organism is learning through cognition since the latter process may not involve consciousness. Going down the evolutionary chain, at some point, probably at the level of roaches, that question becomes meaningless since we seem to have no reason to presume consciousness apart from those processes.

**Cognition**, or a set of cognitive functions, must be distinguished from **consciousness** on one side (since not all cognitive processes are conscious) and from a **stimulus response** (input, transmission, integration, response) on the other side. Cognition engages previously learned and adapted patterns of behavior specific for an individual organism, while a stimulus response is only a reaction to the current input, and does not involve previous conditioning. Consciousness is a high cognitive function, which involves visualization and other processes. We should distinguish reductive and non-reductive explanations of consciousness. Importantly, many scientists do not ascribe to the reductive camp, for instance, according to Francis Crick, one of the fathers of contemporary neuroscience (and molecular biology), consciousness is "**seeing with the mind's eye**".

Here is a refresher on a familiar distinction, followed by some critical discussion. Following Chalmers, we distinguish between two different problems of consciousness: the easy and the hard one. Chalmers argues that neural correlates of the content of consciousness should constitute **the easy problem of consciousness** [Chalmers 1995, Chalmers 2002] where "a phenomenon is explained in terms of computational or neural mechanisms" [Chalmers 1995].

---

<sup>4</sup> The Society for Neuroscience, [www.sfn.org/index.cfm?page\\_name=brainBackgrounders\\_theShortAnswer](http://www.sfn.org/index.cfm?page_name=brainBackgrounders_theShortAnswer)

The hard problem of consciousness is what fails to be grasped by this approach: the subjective, first-person conscious experience. Chalmers characterizes **the hard problem of consciousness** [Chalmers 1995] as "the problem of experience". Those experiences include "the sound of a clarinet, the smell of mothballs ... bodily sensations, from pains to orgasms; mental images that are conjured up internally; the felt quality of emotion, and the experience of conscious thought". [loc. cit.] Chalmers claims that "what unites all these states is that there is something it is like to be in them" [loc. cit.] But is this subjective quality of experience everything? For instance, are mental images properly and exhaustively described as solely states of experience? We pose that it is not the case, by far.

The issue seems to be (mis-)informed by the remnants of Locke's primary- secondary qualities distinction. If I see the shape of a rose, it is supposed to be *objective* (a primary quality) while its color seems *subjective* (a secondary quality). Yet, a simple experiment of moving a magnifying glass towards an object and back shows how shape can be seen as a secondary quality, quite dependent on the senses of the observer. /Saul Kripke pointed out to this fact/. More broadly, any inductive knowledge of the kind people acquire in the third person perspective derives from first-person observations (either directly, of the objects involved, or of some readings of the measuring equipment). There is an aspect of first person experiences that is "subjective" and can be acquired only by acquaintance, but much of the content of knowledge by acquaintance goes beyond the subjective character of experience. It is actually *the content* thanks to which knowledge by acquaintance is a kind of knowledge.

Human beings have sensory access to various objects: some of those sensory experiences are corrigible more directly from the first- than third-person perspective than others; also, some are outside of our bodies while others are inside. There is a valid epistemological distinction involved in this, but the problem of experience invaded ontology due to an unhappy confusion. The confusion pertains to dodging the difference between consciousness (awareness) and the objects of consciousness (phenomenal content)<sup>5</sup>.

Another problem in the understanding of consciousness comes from the tendency to view it as an all or nothing event [Ledwig]. Consciousness is a gradual process, which is another argument for a naturalistic non-reductionism. We can show how different brains, for instance of sick or injured people, have more or less conscious functions. The same gradualization of consciousness can be shown on various animals depending on their level of conscious activity (normal or abnormal to their species), although – due to alternative areas of focus of the activity of various brains – it is sometimes hard to establish a clear hierarchy of dif-

<sup>5</sup> C. O. Evans was one of the few analytic philosophers to avoid the confusion.

ferent levels of consciousness among different species.<sup>6</sup> Those activities are conscious if they are geared through the center of consciousness as opposed to just organized through non-conscious cognitive functions. As Ingarden observed already in the 1960s, the distinction between the functions that come through consciousness and those that do not is anything but sharp, with many cognitive functions just touching on consciousness (and the same is true for the distinction between consciousness and H-consciousness). Let us propose the lightened candle analogy for the functioning of phenomenal consciousness. Consciousness can be weaker or stronger (a rat versus a human being) just like a candle can have a stronger or weaker flame, but there is also a difference between the states in which a candle is lit and extinguished which is analogous to the all or nothing difference between a given organism being (medically) conscious or unconscious.

One more clarification is in order. Non-reductive naturalism Chalmers style seems more robust than epiphenomenalism or double aspect theories; it is in fact not a materialism strictly speaking, since Chalmers postulates consciousness as one of the substances – the building blocks of the world. However, as clarified in ref. 1, our argument does not hinge on details of Chalmers' or any other view in philosophy of consciousness. EH-AI is compatible with any non-reductionism of consciousness.

## Potential objections

We shall list, and attempt to address, the main potential objections to the thesis and argument of this paper, based in part on the past discussions.<sup>8</sup> The readers are encouraged to present the authors with further objections.

A) Consciousness may be created by interactions between brain hemispheres and therefore its creation would involve re-creation of the whole brain.

If we understand the brain we can re-create its multi-hemisphere system; this is true by an argument analogous to the one presented for extra hard AI.

B. Perhaps there are some features of the organic matter that can not be recreated in inorganic matter.

When we consider the creation of consciousness in machines we must make a point that a machine can be

<sup>6</sup> Such inter-species hierarchy naturally branches, with dolphins being more conscious of echolocation-positioning than apes, but apes with some abilities of using tools that, as far as we know, superior to those of other animals. Still, there are clear.

<sup>7</sup> Some observations indicate that even this is not quite the case; the issue goes beyond the scope of this paper.

<sup>8</sup> The philosophical part of this paper has been presented at the: APA Eastern Division, NYC, December 2006, section on Philosophy and Computers. A recent draft was presented at The University of Illinois at Springfield in September 2007. We are grateful to all involved, especially to John Barker, Marvin Croy, Gordana Dodig-Crnkovich and Keith Miller for helpful comments.



of two varieties: organic and inorganic. Organic machines count as having artificial consciousness as long there was no previous non-artificial consciousness among their components. Eligible machines could in principle include not just chips with phospholipid bilayers and computers with bone structural components, but even machines with completely organic tissues engineered to have consciousness.

C) Does the paper cover the cyborgs?

Adding cyborg elements to an already conscious being is interesting but not for the current project. This is because the issue at hand is not whether we could put the part of the brain that is likely to produce consciousness in a different environment (e.g. in a machine). The issue is rather whether we can build a machine that works as a generator of H-consciousness (and what the implications of its successful implementation are).

D) Perhaps it is essentially impossible to ever find the algorithm that describes consciousness *as in the hard problem*. This could be for at least three reasons: 1. because the problem is too hard for human beings [Nagel, McGinn 1991]; 2. Consciousness may be located at the quantum level and therefore we are unable to engineer it. 3. Because the brain does not contain consciousness, but only links to it [Eccless].

The problem may be hard but past progress of science gives strong support to the heuristic hypothesis that we should be able to gain this knowledge in a long run. If McGinn is right and we shall never understand the nature of consciousness as in Chalmers' 'hard problem' then the argument of this paper is empty-satisfied.

While contemporary so-called quantum computing may be using sub-atomic particles, not strictly quantum particles, there is no reason to believe that this will not be done. The history of technology demonstrates that all such artificial thresholds are always crossed eventually.

The hypothesis that the brain does not produce consciousness is no longer represented, aside from people motivated by non-scientific reasons. (This hypothesis needs to be kept in the back-burner in case new facts confirm it in some version).

E) The reader is encouraged to raise further challenges.

### The importance of this finding

The main thesis has theoretical and indirect practical implications. The theoretical implications come from the understanding that the hard problem of consciousness is not applicable solely to natural biological brains.

Its practical implications for ethics are potentially enormous and deserve a different article. In particular, there are

good reasons to believe that ascription of moral value should follow, in part, subjects of (advanced) *consciousness as in the hard problem* [Boltuc 2007]. While AI ethics tackles largely the issue of responsibility of humans for creating robots, or ascription of moral responsibility for actions committed by robotic devices, the issue of inherent moral value of robots as moral patients is rarely addressed. Its main presentation is in the theory that ascribes moral value to creating anti-entropy [Floridi 2007]. However, creation of robots with H-consciousness would make it more persuasive to ascribe the inherent value to robots as moral patients since we tend to ascribe such value to conscious intelligent beings. Therefore, it follows that inherent moral value should be ascribed to organic (human and non-human) as well as non-organic sources of intelligent consciousness.

There are also potential practical implications. It is likely that first-person consciousness (H-consciousness) allows for certain cognitive functions, such as self-reflection, that are not possible for unconscious (zombie-like) creatures.<sup>9</sup> Alternatively, H-consciousness may allow for a simpler execution of certain cognitive functions, which are possible, but more complex, for unconscious creatures<sup>10</sup>. In both instances it would, in principle, make sense to consider engineering H-consciousness in thinking machines.

### Bibliography

- Baars, B. J. & Newman, J. 1994. A neurobiological interpretation of the Global Workspace theory of consciousness. In (A. Revonsuo & M. Kamppinen, eds) *Consciousness in Philosophy and Cognitive Neuroscience*. Lawrence Erlbaum
- Baars, B. J. 1997. In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies* 4:292-309
- Baars, B. J. 2005. Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in brain research*. 150:45-53.
- Blake, R. 2001. 'A primer on binocular rivalry, including current controversies', *Brain and Mind*, 2, pp. 5-38.
- Boltuc P. 2008. Solidarity and The Shape of Moral Space Forthcoming.
- Chalmers, D.J. 1990, Consciousness and Cognition unpublished.

<sup>9</sup> H-consciousness may be creating some kind of a *global workspace* that requires conscious access in order to function at all.

<sup>10</sup> Here H-consciousness may be needed to enhance the functioning of such global workspace. The last process would work like color tagging; colorless tags would probably work too but the use of color tags makes it more convenient (simpler) to recognize the relevant differences among the objects tagged.

- Chalmers, D.J. 1995. 'The puzzle of consciousness' *Scientific American* :92-100.
- Chalmers, D.J. 1997. Moving Forward on the Problem of Consciousness *Journal of Consciousness Studies* 4(1):3-46
- Chalmers, D.J. 2003. 'Consciousness and its place in nature' in S. Stich & F., Warfield, eds, *Blackwell Guide to Philosophy of Mind*. Blackwell, 2003.
- Churchland, P. M. & Churchland, P. S. 1981. Functionalism, qualia and intentionality. *Philosophical Topics* 12:121-32
- Crick, F. 1996. 'Visual perception: Rivalry and consciousness', *Nature*, 379, pp. 485–6.
- Crick, F. & Koch, C. 1990. Toward a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2:263-275
- Crick, F. and Koch, C. 1995. 'Are we aware of neural activity in primal visual cortex?', *Nature*, 375, pp. 121–3
- Dennett, D. 1991. *Consciousness Explained*, Penguin
- Eccles, J. 1985. *The Brain and the Unity of Conscious Experience*. Cambridge University Press.
- Engelmore, R., and Morgan, A. eds. 1986. *Blackboard Systems*. Reading, Mass.: Addison-Wesley.
- Evans, C. O. 1970. *The Subject of Consciousness* George Allen & Unwin Ltd
- Floridi, L.; Sanders, J. W. 2002: Mapping the Foundationalist Debate in Computer Ethics in: *Ethics and Information Technology* 4/2002 pp. 1-9
- Harman G. 2007. Explaining an Explanatory Gap *APA Newsletter on Philosophy and Computers* 06/02
- Hernegger, R. 1995. Changes of paradigm in consciousness research. <http://chaos.ip-technik.de/Hernegger/consciousness.html>
- Ingarden, R. 1983. *Man and Value/Roman Ingarden* Catholic University of America Press
- Lamme, V.A.F. 2004. Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Netw.* 17, 861–872
- Lamme, V.A.F. 2006. Towards a true neural stance on consciousness *Trends in Cognitive Sciences* 10/11, 494-501
- Laureys, S. , Owen, A. M. , & Schiff, N. D. 2004. Brain function in coma, vegetative state, and related disorders. *The Lancet Neurology* 3:537-546.
- Laureys, S. et al. 2002. Brain function in the vegetative state. *Acta neurol. Belg.*, 2002, 102, 177-185
- Ledwig, M. 2007. To Understand the Understanding. *APA Newsletter on Philosophy and Computers* 2/2007
- McGinn, C. 1991. *The Problem of Consciousness: Essays Toward a Resolution*. Blackwell
- McGinn, C. 2004. *Consciousness and Its Objects* Oxford University Press University Press
- Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* 83:435-50
- Nagel, Thomas 2002. *The Psychophysical Nexus in his: Concealment and Exposure and Other Essays* New York, Oxford University Press, 2002 pp. 194-235
- Noe, A, Thompson, E 2004. Are there neural correlates of consciousness?. *Journal of Consciousness Studies* 11: 3 28.
- Tononi, G. 2004. An Information Integration Theory of Consciousness. *BMC Neuroscience* 5: 42.
- van Fraassen, B. 1990. *The Scientific Image*. Oxford University Press
- van Fraassen, B. 2002. *The Empirical Stance*. Yale University Press
- Zeki, S. & Bartels, A. 1999. Toward a theory of visual consciousness. *Consciousness and Cognition* 8:225-59