What is (it like) to be a Hard Problem?

"Some books are important not because they solve a problem or even address it in a way that points to solution, but because they are symptomatic of the confusions of the time."
SEARLE[1]

Structure of the essay:

Abstract: In this essay I argue that a systemic perspective of cognition may be sufficient to explain whatever it must be necessary to explain about consciousness. By analysing Chalmers' diagnosis of the Hard Problem of consciousness we conclude that the only Hard Problem arises from the functionalist view of cognition. I argue that a functional explanation is not enough to explain consciousness (and that is why Chalmers' Hard Problem arises) and that an operational explanation is required. It follows that once we have specified the structure that makes us conscious then 'what phenomenal consciousness is' becomes a matter of *being* that structure and not something to be explained. Finally I argue that considering a system conscious depends on the operational conditions under which it is legitimate to describe an entity as conscious i.e. the necessary and sufficient operational conditions for a system to be conscious.

Key words: Consciousness, Cognitive Sciences, Cartesian Dualism, Explanatory gap, Dynamical approach, Phenomenal experience, Operational explanation.

1. INTRODUCTION

Since Descartes the relation between the phenomenological world (*res cogitans*) and the physical world (*res extensa*) occupies a privileged position between

[1] Searle, 1997, p.162. Searle's review on David Chalmers *The Conscious Mind; In Search of a Fundamental Theory* (Oxford University Press, 1996)

Xabier Barandiaran *barandi@euskalnet.net*

the unresolved philosophical questions of the Western history. If not so much Descartes' dualist ontology, its vocabulary and conceptual foundation remain alive in the contemporary debate (Searle, 1992) and, while consciousness is becoming an object of scientific study it's ontological and epistemological status is still in question. And it's not a trivial question since consciousness seems to us the highest human capacity, the more inaccessible domain, the most secret privacy, and the last hiding place of the individual against the objectivity. Can science grasp this mystery? Is there any mystery at all?

The current study of consciousness is characterised by a transdisciplinary, multidimensional and weakly co-ordinated approach. All sorts of theories and approaches inhabit the scene while they remain unconnected (at best) or incompatible. Moreover the claim of the impossibility of a scientific study of consciousness remains alive among some scientist and philosophers (Nagel and McGinn). In this context, as Searle's *Chinese Room Experiment* (1980) for the problem of intentionality, Chalmers' article *Facing Up to the Problem of Consciousness* (1995) became a reference point and focused the debate. Even if Chalmers' article has been considered a steep back in the debate (Dennett, 1996) I consider it symptomatic of a profound disagreement between different views in the field. I believe the really *Hard Problem* stands in the deep tension between confronted underlying assumptions in the current field (specially among Functionalists) and that we shall assume a biologically grounded operational perspective in order for the Hard Problem to vanish.

But what is the, so called, Hard Problem?

2. THE HARD PROBLEM

From Chalmers' article we will rescue the analysis of the so called Hard Problem (HP); a) because part of the current debate is focused in Chalmers diagnosis of the Hard Problem, b) because I consider (and I will try to argue) that Chalmers' mistake is already present in that diagnosis and that the latter development of his paper is a consequence of that mistake, and, c) because the claim of the HP, seems to me, is the point where consciousness, as a philosophical debate/problem, should arrive to an end, just because there is not such a problem (or at least the problem shows to be a philosophical problem in the more Wittgensteinian linguistic viewpoint).

At the beginning of the paper Chalmers divides the problems of the study of consciousness into the 'easy' problems and the 'hard' problem.

> "The easy problems of consciousness are those that seem directly susceptible to standard methods of cognitive science, whereby a phenomenon is explained in terms of computational or neural mechanisms. The hard problems are those that seem to resist those methods" (§ 3).

The easy problems: the easy problems are those concerning functional mechanisms:

- Ability to discriminate stimuli
- Integration of information
- Reportability of mental states
- Focus of attention
- Control of behaviour
- Difference between wakefulness and sleep

All those phenomena are associated to consciousness but they have a functional role in the cognitive processes of the cognitive agent; thus they can be explained in functional terms. However difficult they may turn to be in the future, Chalmers takes for grounded the conceptual frame on which they will be explained so that a good explanation is a matter of techno-scientific achievement but not one of conceptual re-formulation.

The Hard Problem: But… (Chalmers follows) if those phenomena (easy problems) are exhausted in their functional role… how is it possible that they give rise to phenomenal experience? And this is what Chalmers considers The HP. The HP, thus, is the problem of *experience* that carries on in the Mind-Body debate under different forms:

- Consciousness* (Harvey)
- First person ontology (Searle)
- Qualia
- Phenomenal consciousness (Ned Block)
- "What is it like" (Nagel)
- Explanatory gap (Joseph Levine)
- Knowledge argument (Frank Jackson)
- Phenomenal experience (Chalmers)
- Etc…

What makes the HP hard, for Chalmers, is that it goes beyond any performance of functions because "to explain a cognitive function we need only specify a mechanism that can perform that function" (§12) and after explaining all those mechanisms we still have something else to explain: consciousness is more than a mechanism. That's why, again on Chalmers' view, any attempt to explain consciousness in the current literature, doesn't work. We need an *extra ingredient*, a *something else* to fill the *explanatory gap* of *why* this functional or neural processes are *accompanied by experience* (the italic is mine to highlight Chalmers' most common expressions when describing the *problem*).

After analysing some case studies that fail to explain consciousness, Chalmers concludes

> At the end of the day, the same criticism applies to any purely physical account of consciousness. For any physical process we specify there will be an unanswered question: Why should this process give rise to experience? Given any such process, it is conceptually coherent that it could be instantiated in the absence of experience. It follows that no mere account of the physical process will tell us why experience arises. The emergence of experience goes beyond what can be derived from physical theory. (§ 43)

At this point, considering that any given cognitive process could exist without experience, Chalmers proposes to introduce *experience* as a fundamental feature of the world. Then he outlines a theory of consciousness whose central claim is 'The double-aspect theory of information' by which information is understood as being the basis of consciousness and the link with physics through the embeddedness of information in physical processes.

3. WHAT IS (IT LIKE) TO BE A HARD PROBLEM

3.1. The dissolution of the HP and the HP of traditional functionalism.

Explanatory gap, *extra ingredient*, *accompanied by experience*, *something else*, and so on seems to me dangerous expressions that cover a non-existing problem.

If we add to a functional explanation a structural bottom-up, operational[2] explanation, there is nothing else to be explained.

I don't deny conscious experience, the vivid sensation of perceiving a red apple or having an orgasm. The Knowledge Argument (Jackson, *Epiphenomenal Qualia* 1982, from Torrance, 1998) illustrates this point: it is still different to have an orgasm than knowing all the processes and biological structures involving the phenomenon. But the Knowledge Argument is not an argument against the explanation (as it has been used), it is not an argument that shows that any explanation of consciousness is not enough; it just shows that it is different to explain how a cognitive systems works than being a cognitive system (how are they going to be equal? They are not even in the same level to be compared!). Probably the narrow vision of the classical functionalist viewpoint of what a cognitive agent is makes impossible to imagine that experience *is* to *be* a cognitive system. If we take cognition (as traditional functionalists do) to be the computation of an algorithm (independently of the system that performs it) then experience (the inner vivid sensation of experience) seems to be something else and the HP arrives when we realise that phenomenal consciousness cannot be added to the list of algorithms that constitute cognition. That's why Chalmers considers that 'the something else' must be explained. But, from a systemic perspective, there is nothing to be explained about *being* a cognitive system because there is nothing on *the being* that can/should/must be explained. As well as there is nothing to be explained about "what it is like to be *solid*"; there is nothing to explain about *being* conscious. The only way of explaining solidity is specifying how the microstructure makes a solid macrostructure (what we have called structural[3] or operational explanation). If our theory of solidity was that of a *solidity algorithm* then a certain Zhalmers could argue that there is a HP of solidity since the computation of the *solidity algorithm* is not enough to explain solidity: the inner robustness of being solid. Zhalmers could argue that there could exist Zolids: entities that perform the *solidity algorithm* but are not solid. In this case we wouldn't

[2]Here operational vs. functional will be understood following (Di Paolo, 1999). By operational we mean an explanation which is "formulated in terms of a set of elements all pitched at a same descriptional level and also in terms of law-like realtionships between these elements so that an account can be given of how the phenomena are generated" (p.16) while by functional explanation we understand an explanation where "the terms of the reformulation are deemed to belong to a more encompassing context, in which te observer provides links and nexuses not supposed to operat in the domain in which the systems that generate the phenomena operate" (p.16). We intuitivelly understand operational explanations as specifying the structure of the system by establishing the elements and the law-like relations between the elements that constitute the system as such. While by funcional explanation we will understand particularly the kind of explanations of cognition held by Traditional AI where a cognitive agent can be defined solely in terms of the causal-computational relations between inputs and outputs requiring an external observer to specify them.

[3] The term *structural* or *structure* won't be used in this essay as oposed to *operational* (Varela) but rather as oposed to *functional*.

have a HP of solidity at all but a wrong explanation of solidity that becomes problematic. But *being* solid or experiencing orgasms does not enter the scientific discourse. And this is mostly because of the very meaning of 'explanation' in science, which was (fortunately) moved from Aristotelian 'intrinsic property of an entity' to the Galilean law-like relations between entities (Rorty, 1994).

Lets go back to the beginning of this section; now we can see how terms like *explanatory gap*, *extra ingredient*, *accompanied by experience*, *something else*, and so on, are absolutely mistaken. There is no *explanatory gap* because there are not two objects to be linked. From our perspective *accompanied by experience* is completely nonsensical, the perceptive process of perceiving red IS the experience, thus, there is no *extra ingredient* to be added. But Chalmers mistake is still worst since the assumption of any extra ingredient falls under the Hard Problem again (the extended HP, Torrance, 1998), and no matter how many ingredients we add we will always need another one. At this point Zhalmers could argue against Chalmers with his own arguments: "why does any aspect of information give rise to experience? We need a third extra ingredient since it is, still, conceptually coherent to imagine any physical process + informational process without *experience*." and so on, and so on *ad infinitum*.

Chalmers' Hard Problem itself is a problem of functionalism (whose view of the mental and cognition is disembodied and functional and never operational), not a problem of Cognitive Sciences. As Jackendorff pointed out: if consciousness has no causal effect (functional role) then 'it is useless' (Jackendorff, 1987 p.26, from Varela et al. 1991. P. 82). But all there is in the domain of the mental is not functional. Chalmers functionalist view of the mind (only considering the causal connections between representations) is the *real* Hard Problem. As Searle points out (Searle, 1997) Chalmers mistake stands on trying to hold both functionalism and property dualism or irreducibility of consciousness; and it just doesn't work. As Chalmers itself sees, functionalism is not enough to account for consciousness. And that is because functionalism has an horizontal concept of causation. Functionalism studies relations between representations (propositional attitudes) and this is not enough to account for cognition; on the lower boundary of cognition to account for the symbol grounding problem (Harnard, 1990), embodied situated cognition (Brooks, 1991) etc.; on the upper boundary for such non-functional 'phenomena'[4] as the self (Varela et al. 1991), consciousness, etc.

[4] I quoted 'phenomena' because the self and consciousness cannot be properly called phenomena because they are prior to any phenomenon as such. In fact, this is the whole point of the essay.

3.2. What is (it like) to be a cognitive system

If we now consider a dynamical-embodied perspective of cognition, the functionalist HP comes to an end: our experience is embodied (Varela et al. 1991, Varela 1996) thus if we don't want to fall into explanatory gap problems we have to consider cognition as embodied, as realised (and only realisable?) in a biological-dynamical structure. And then take (conscious) experience as *being* this embodied dynamical structure. Them, ones we assume this viewpoint, what makes us being as we are, which are the concrete dynamics that constitute a conscious being, will become a scientific task to be resolved, an operational description to be made.

At this point I would like to point to a dynamic-systemic Artificial Life (A-Life) approach to cognition from which I believe the problem could be addressed in a fruitful way (as this new approach to cognitive phenomena has showed with many other problems of Cognitive Science). Within this view cognition as a process can be understood as the structural coupling between an agent and its environment and a cognitive agent can be studied as a dynamical system. The dynamical approach (van Gelder, 1993, Van Gelder and Port 1995) does not entail, necessarily, the exclusion of symbolic/computational explanations of some cognitive processes because any computational process can (in principle) be explained by dynamical system theory (even if the concrete mechanisms require a research effort not yet resolved –Crutchfield, 1998). Thus any informational and functional account of consciousness (the easy problems) is not rejected but subsumed in an embodied dynamical bottom-up explanation of cognition (at the same time an embodied perspective can solve the symbol-grounding problem, the intentionality (Searle) of some cognitive processes[5]). As Di Paolo points out "the purpose of an operational explanation is not necessarily to act as a replacement of a functional one, but rather to act as a constraint to the possible functional interpretations that we may need to provide for pragmatic purposes." (Di Paolo, 2000:4). The other way round works as well so that a functional explanation cannot replace an operational one and if we are to take experience as real then we must avoid any functional explanation of it because we know that experience cannot be an added function to the purely cognitive ones. What, in fact, is proposed here is that an operational explanation must be necessary and sufficient to *explain* whatever could be and must be explained about inner experience, consciousness*,

[5] It is not a coincidence that what Chalmers calls the Hard Problem is strongly related to the problem of semantic content and intentionality, which is, at the same time, one of the major problems of functionalism (are consciousness and intentionality very far from being the same problem?).

phenomenal consciousness or qualia. The *rest* is not a rest, an object, nor a function, nor information; the phenomenal experience "is not a 'thing that exists' in the sense that matter and charge, indeed tables, and apples, exist." (Harvey:13), but the mere *being* that operationally described system, the *being* an observer: "Phenomenal experience is primary, and through our experience we construct matter, charge, tables, and apples as objects that allow us to make sense of the world" (Harvey).

Up to now we have been addressing the explanation of consciousness to a structural-operational study. But how does consciousness emerge from the dynamics of the brain? There is an increasing work on the foundations of biology, A-life and complex dynamical systems (Varela et al., 1991), where the concept of emergence plays a central role. Collier (Collier, 1998) argues that emergent properties entail cohesion, where “cohesion represents those factors that causally bind the components of something through space and time, so it acts coherently and resists which internal and external fluctuations”. But cohesion, as causal condition for the emergence of a property can be understood in terms of transfer of information (according to Collier). Thus, Chalmers was probably not completely wrong with his double aspect theory of information after all? Well it depends on what we understand by completely wrong, which is clear is that Collier’s argument does not support any interpretation of a phenomenal side of information. On the contrary I suggest that consciousness could be understood as an emergent property of informational processes happening in the brain[6]; where conscious experience, the phenomenal side of consciousness, is just *being* this process. Nevertheless, the emergence of consciousness is a matter of science. I only wanted to point to this coarse proposal as a coherent alternative to the Hard Problem of functionalism. But I shall point out a major problem of my last proposal which is that it requires to built an operational explanation on the top of a functional one that takes cognitive information (which may considered a functional term), for grounded. I address this issue to Moreno's perspective in cognition and the operational closure of the nervous system in informational terms (Moreno et al., 1997)

4. DISTINGUISHING A CONSCIOUS BEING

Up to now we have dealt with what it could be defined as an explanatory problem. But immersed in the literature around the topic of consciousness we find another sort of problem which is strongly linked to the above one not necessarily

[6] In this sense my position could be compared with Searle’s *biological naturalism.*

determining it (as it happens for some authors) , i.e. the problem of *distinguishing* a conscious being from a non-conscious one. I find the problem quite similar to the Turing test and all the ongoing functionalist problems of Strong and Weak AI (Searle, 1997). If we reduce consciousness (the cognitive and the phenomenal side) to a functional explanation, thus to an observable behaviour, then the Hard Problem arises and Zombies enter the scene (as well as if we reduce cognition to functional computation the 'symbol grounding problem' arises and 'Chinese room' kind of arguments enter the scene).

But nor functionalism is enough to explain cognition nor we do discriminate solely by behavioural observations. Until the sciences of the artificial, and robotics arose, for a certain behaviour there was usually the same kind of physical structures performing it (i.e. a human body for linguistic behaviour). That is why we, humans, take for grounded that for a certain behaviour there is a corresponding structure performing it, with its evolutionary history, structural causality and so on. But if, by chance, a plastic ball gets out of the window, nobody will attribute to that behaviour any intentionality of willing to suicide, nor any other intentional instance, because we know that the causal relations that make that ball going out of the window are structurally different from ours.

At this point the problem can be understood as an attitude problem:

> So if one accepts that the creation of robots with consciousness1-3 offers merely 'easy' problems (...), the additional magic ingredient for consciousness* is merely a change of attitude in us, the observers. Such a change of attitude cannot be achieved arbitrarily; the right conditions of complexity of behaviour, of similarity to humans, are required first
> (Harvey:10)

But an attitude problem is not merely an 'attitude' problem because my attitude with my teddy bear does not make the teddy bear conscious. The problem must be addressed as under what condition it is legitimate the attitude of considering or not a certain entity to be conscious. If phenomenal consciousness is referred to that *being* a cognitive system independently of the performing functional behaviour of the moment, the attribution of consciousness to a certain entity cannot be a matter of behaviour but a matter of the operational organisation of the entity that performs such behaviour. This way the matter of 'attitude' becomes an epistemological matter of establishing the operational conditions under which it is legitimate to describe an entity as conscious i.e. the necessary and sufficient operational conditions for a system to be conscious.

5. CONCLUSION: CONSCIOUS EXPERIENCE AND SCIENTIFIC STUDY:

By the claim that being conscious is no more (nor less) than *being* a certain kind of cognitive system (with the appropriate structure and processes from where consciousness emerges) and thus that the HP does not entail any problem at all (but addresses a HP inside the functionalist assumptions), I don't mean that we have no inner experience. What I suggest is that the *existence* of experience, of phenomenal consciousness, does not require any intrinsic explanation at all: because it is impossible a priori (by definition of the term explanation), and because the concept of intrinsic explanation entails the same problem ones again ad infinitum. Neither do I mean that there is no place in science for conscious experience. I suggest that the place of experience should be methodological rather than ontological. In this sense Varela's proposal for a neurophenomenological framework (Varela, 1996) seems to me completely coherent with my argument. If there is anything to be explained this must be in functional and structural-operational terms, as proposed above.

But this claim does not narrow the methodological scope, it just opens it ones we realise that there is no extra ingredient nor explanatory gap. In this context, doubtless, phenomenology stands for one of the most powerful tools to take into account. Just because been cognitive systems put us in a privileged position to know how 3 trillion neurones work (and how the embodied study of those 3 trillion neurones work, as well). After all, the problem, seems to me, is more methodological than ontological. We know that something special (cognition, intentionality, consciousness) is going on in our brains, because we experience it[7], because we *are it*: a complex self-organising activation patterns and processes which can be (however difficult it must turn to be), in principle, studied by neurophysiology, computer simulations, dynamical system theory etc. The study of this *going on in the brain* should provably be divided into different emergent levels, it will probably require the appropriate mathematical tools to understand emergence (Collier, 1998), as well as an account of how parallel processes have a unitary functionality[8], it will probably be challenged by pluralist approaches (Torrance, 1996) and so on. But this is all it is to be explained. That is what is (it like) *to be* a Hard Problem

[7] I want to point to the contradiction involved in this expresion, namely that we cannot experience consciousness because consciousness is not an object to be experienced or hadled by a subject but the very fact of been a subject. Language makes the whole subject/object dicotomy hard to solve since the very structure of language involves such dualism.

[8] Probably explanable by the 40 Hz hypothesys (Crick and Koch).

6. DISCUSSION: POINTING TO A HARD PROBLEM AND A CRUCIAL GAP

I wish to finish this essay by pointing to a further discussion about some related issues that I find specially important but far from most of the efforts in the field. I will briefly note them.

In what accounts for consciousness we are dealing as well with the notion of the self (and still *worst,* with the *experience* of the self) considered as the unifying substrate of experience. I consider that the underlying problem of the self has great possibilities of becoming a real hard problem since it is one of the constitutive notions of western civilisation on which science is immersed.

On the other hand (but somehow linked with the former problem) I find the problem of the gap between dynamic experiences (understood as personal/individual experiences) and intellectual experiences (purely symbolic/abstract experiences). Following the work by Varela et al. (1991) I believe that working in this two issues is fundamental if we want knowledge to serve human purposes and not *vice versa*.

7. ACKNOWLEDGEMENTS

To my father for his help on rigour, for his patience, for being always there. To Steve Torrance and Alvaro Moreno for guiding me first steeps into Cognitive Sciences. To Alfredo for technical support on *What Macintosh still can't do* (on decompressing on-line articles).

8. BIBLIOGRAPHY and REFERENCES

BROOKS, R. (1991) *Intelligence without representation.* Artificial Intelligence **47** (1991), 139-159.

CHALMERS, D. (1995). Facing Up to the Problem of Consciousness. Journal of Consciousness Studies. **2**:200-220.
[I used the HTML version on
http://cogprints.soton.ac.uk/archives/phil/papers/199806/199806022/…/cosnciousness.html]

COLLIER, J. D. (1998). *The dynamical basis of emergence in natural hierarchies*. From George Farre and Tarko Oksala (eds*) Emergence, Complexity, Hierarchy and Organization, Selected and Edited Papers from the ECHO III Conference, Acta Polytechnica Scandinavica, MA19* (Finish Academy of Technology, Espoo, 1998.)

CRUTCHFIELD, J P. (1998) *Dynamical Embodiment of Computation in Cognitive Processes*. Submited as Open Peer Commentary on T. van Gelder (1998) The Dynamical Hypothesis in Cognitive Science, BBS to appear.
[Taken from the internet http://www.santafe.edu/jpc]

DENNETT, D. (1996). *Facing Backwards on the Problem of Consciousness*. Published in Journal of Consciousness Studies, vol.3, no.1, 1996, 4-6.
[I used the HTML version on http://ase.tufts.edu/cogstud/papers/chalmers.htm]

DI PAOLO, E. (1999). *On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical*. DPhil Thesis, School of Cognitive and Computing Sciences, University of Sussex.
[I used the internet version from www.cogs.susx.ac.uk/users/ezequiel/thesis.html]

DI PAOLO, E. (2000). *Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents*. Adaptive Behavior 8:1. 25-46. Special issue on Simulation Models of Social Agents. K. Dautenhahn (guest ed.)

GUTTENPLAN, S. (editor). *A companion to the philosophy of mind*. Blackwell, 1998.

HARNAD, S. (1990) The Symbol Grounding Problem. Physica D42: 335-346
[I used the HTML version on http://cogprints.soton.ac.uk/archives/psyc/papers/199803/199803014/doc.html/The_Symbol_Grounding_Problem.html]

HARVEY, I. *Evolving Robot Consciousness: The Easy Problems and the Rest*. To appear in Evolving Consciousness, G. Mulhauser (ed.), Advances in Consciousness Research Series, John Benjamins, Amsterdam. In preparation.
[I used the internte version at ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/consc.ps.gz]

MORENO, A., UMEREZ, J. & IBAÑEZ, J. (1997) *Cognition and Life. The Autonomy of Cognition.* Brain & Cognition **34 (1)** Special Issue Academic Press pp. 107-129.

RORTY, R. (1994).*Consciousness, Intentionality, and the Philosophy of Mind*. From *The Mind-Body Problem (A guide to the current debate)*. Edited by Richard Wagner and Tadesz Szubka. Blackwell, 1994. P. 121-127.

SEARLE, J. (1992) *What's Wrong With the Philosophy of Mind?*. 1st chapter of 'The rediscovery of the mind' (MIT press, 1992). Taken from *The Mind-Body Problem (A guide to the current debate)*. Edited by Richard Wagner and Tadesz Szubka. Blackwell, 1994. p. 277-298.

SEARLE, J. (1997). *The Mystery of Consciousness*. Granta Books 1998.

TORRANCE, S. (1996). *Real world: embedding and traditional AI.*

TORRANCE, S. (1998). *The Taste of Lemons: A New Twist to the Cosnciousness Debate.* (Talk delivered in Psychology Group, Middlesex University, November 1998)

Van GELDER, T. (1993) What can cognition be if not computation? From *III International Workshop on Artificial Life and Arificial Intelligence*, Workshop Notes, second edition, UPV. San Sebastian, 1995.

Van GELDER, T. (1995). *Mind as Motion: Explorations in the Dynamics of cognition.* MIT press.

VARELA, THOMPSON AND ROSCH, (1991). *The Embodied Mind: Cognitive Science and Human Experience*, Cambridge MA:MIT Press. (I used the Spanish edition: Editorial Gedisa, 1997)

VARELA, F. (1996). *Neurophenomenology: A Methodological Remedy for the Hard Problem.* In Journal of Consciousness Studies, "Special Issues on the Hard Problems", J.Shear (Ed.) June 1996.

[I used the HTML version on http://www.ccr.jussieu.fr/varela/
human_consciousness/article01.html]