The Caldron of Consciousness

Motivation, affect and self-organization

Advances in Consciousness Research

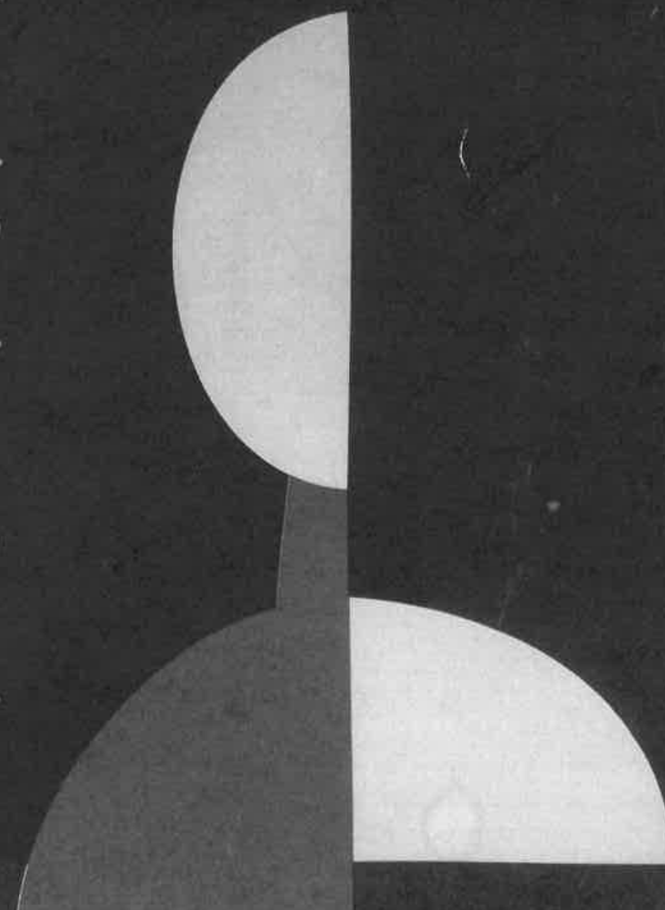

This is an offprint from:

Ralph D. Ellis and Natika Newton (eds.)
The Caldron of Consciousness
Motivation, affect and self-organization - An anthology
John Benjamins Publishing Company
Amsterdam/Philadelphia
2000
(Published as Vol. 16 of the series
ADVANCES IN CONSCIOUSNESS RESEARCH, ISSN 1381-589X)

ISBN 90 272 5136 3 (Eur.) / 1 55619 196 0 (US)

CHAPTER 9

Mind, Brain, and Chaos

Nicholas Georgalis
East Carolina University

Most work in cognitive science, whether computationally or biologically based, relies on the idea of 'unconscious representations.' I argue that there are no such things. In contrast, conscious states do represent objects outside of the agent. My second main thesis concerns the 'explanatory gap' between brain states and mental states. I propose an analogy to enhance our understanding of just what is required to close the explanatory gap; I argue that a new science is needed. There is third main thesis, related to the second. Of all the currently available approaches to the study of brain activity, only nonlinear dynamics, chaos theory in particular, has a unique characteristic sufficient to close the explanatory gap. Every other approach to the study of the brain appears doomed to failure in this regard. Identification of an additional feature of chaotic systems provides the framework for the explanation of another conundrum, mental causation. Furthermore, given the chaotic model of the dynamics of perception arising from Walter J. Freeman's work, it appears that emotion plays a central role in perception. The application of chaos theory to brain activity dovetails with my first thesis, as nonlinear dynamics does not require state transitions from representations to representations, yet it has the resources to account for the representational nature of conscious states.

1. Unconscious states are nonrepresentational

There is an innocent and straightforward sense of 'unconscious beliefs' such that no one, I take it, would deny that at any given moment, each of us has a vast number of them. I have in mind quite ordinary beliefs that each of us has but is not currently entertaining. All non-occurrent beliefs that one has are unconscious — beliefs that could be expressed by sentences such as 'Gold is a metal,' and 'Two is the only even prime number.' Just as clearly, any such belief is one that we could entertain and therefore, at times, be conscious of. Such beliefs easily pass from being unconscious to conscious, and back. Being unconscious in this sense does not make them deeply hidden; they are not Freudian.

Thus understood, there clearly are unconscious beliefs, and each *in some sense* is identical to its conscious manifestation. What identity there is between a conscious and an unconscious belief is secured by the *sentence* expressing the belief, which just means that the *sentence* expressing the belief is the same whether the belief is conscious or unconscious. But not everything we say of conscious beliefs is true of unconscious beliefs because their realizations are different.

Part of what is at issue here is just what it is to *have* a belief, and whether one *has* conscious and unconscious beliefs in the same way. I take it that when one "has" a belief, be it conscious or unconscious, it is realized in *some* way, be it neurobiological, computational, individuated somehow by its causal connections to the environment, or something else. We have every reason to expect that, however a belief is realized, its conscious and unconscious realizations are different. (I will here assume the realizations are neurobiological.)

An objection one might raise to my claim that conscious and unconscious realizations of the same belief are different might go as follows: although the realizations are unlikely to be identical, there is a "common component" to both, one that constitutes or corresponds to the belief itself. They are not quite identical, however, because a conscious belief obviously requires something additional. For example, if one's unconscious belief, **p,** is realized by some brain state, then one's conscious belief **p** is realized by that very same brain state plus others, or that brain state enters into different relations with other brain states when there is a conscious state; these other states or relations do not obtain when the belief is unconscious.

The main thing wrong with this kind of move is that it is committed to a "storehouse" view of unconscious beliefs, as if each and every belief one has, including the unconscious beliefs, is somehow stored intact in one's head. But at any given moment, one *has* indefinitely many, perhaps infinitely many,[1] unconscious beliefs and, according to this view, distinct states for each one must somehow be realized. As is well known, this is at best highly implausible. Better to reject the idea of a common realization, even partial, and think more along dispositional lines, in parallel with various dispositional physical properties. For example, that a glass *is fragile* is realized in its molecular structure. Importantly, the very same molecular realization of its fragility is *simultaneously* a realization of indefinitely many other dispositional properties. If the molecular structure that realizes a glass's fragility were exposed to extreme heat or an acid, rather than sharply struck, the molecular structure that would result, the realization of the new property, would be different from that which would result from its being struck. So the molecular structure that realizes the glass's fragility also realizes all its other dispositional properties. Countless dispositional properties are realized in the original molecular structure; which of these is actualized depends on what conditions that self-same molecular structure is subjected to.

The model for unconscious belief advanced here, more generally, unconscious mental states, is analogous: One and the same, say, brain state is *simultaneously* a realization of any number of unconscious mental states. Given that brain state, various conscious states may subsequently be realized; which is realized depends on how the agent is prompted. The parallel to dispositions is not exact but only a first step. Ultimately, the model for unconscious mental states is based on chaos theory and, therefore, unlike dispositions, is not stimulus driven. (See Section 3.)

The neuronal realizations of unconscious mental states should be further distinguished from the broader class of neuronal states that have nothing to do with cognition, those involved in, say, digestion or respiration. The value of categorizing some non-conscious states as 'unconscious *mental* states' is to indicate the existence of relations between them and various conscious mental states, relations that other non-conscious states do not share. In spite of these relations, my position is that *all* unconscious mental states are *non*-intentional, *non*-representational.

I contrast my view with John Searle's, since he holds that unconscious beliefs are representational.[2] Searle's argument is straightforward; it relies on Leibniz's Law.[3] One of John's beliefs is that asparagus is a vegetable. It is the same belief that John has, whether John is consciously entertaining it or not, say, he is asleep, otherwise unconscious, or conscious but thinking about other things entirely. When in any of these latter states, it is correct to say he still has the self-same belief that asparagus is a vegetable. His conscious belief is intentional. Therefore, *this* belief when unconscious is intentional.

The appeal of this argument turns on the reification of beliefs. While it is correct to hold that one *has* the same belief, whether it is consciously or unconsciously had, it is not some identical *thing* that one has on either occasion. It is not some little 'nugget' tucked away somewhere in the brain. The dispositional model of beliefs advanced above undercuts treating belief as a thing. Thus applying Leibniz's Law as a basis for an objection to my view fails, since the realization of conscious and unconscious beliefs is different. For convenience we may speak of the 'same belief' consciously or unconsciously had on different occasions, but it is not the self-same *thing* in the agent. To say, as above, that John has the belief that asparagus is a vegetable when not consciously entertaining it is to make a pragmatic claim about what John would do or say when prompted in various ways. How the belief is "had" is explicated as how it is realized, and conscious and unconscious realizations may well have different properties.

Now the representational nature of conscious beliefs plays a clear role in the analysis of opacity of conscious beliefs. But it is not at all clear how representation could be involved in talk of *unconscious* beliefs realized in the brain.[4] Recall, on my view, a single sentence expresses both the conscious and the unconscious belief, and it is the sentence that bears the burden of the identity. Any opacity questions that may arise regarding unconscious mental states are reducible to the *sentences* that express those unconscious states; they just do not apply to the realizations of the unconscious mental states.

Even if Searle's argument were sound, however, it would provide no help as to *how* unconscious beliefs or mental states *could be* intentional. There appear to be enormous conceptual and methodological problems in answering this "how" question. Searle himself denies that the physiological facts are themselves sufficient to infer aspectual shape, even though on his view belief states are caused by and realized in brain states. (The aspectual shape of a mental state is very closely related to its representational nature. One might hold a belief under one representation, but disavow it under an extensionally equivalent representation.) He argues that no matter how complete the behavioral or even neurophysiological evidence, there would still be an inference from these facts to the aspectual facts. Behavioral or neurological facts cannot, he says, constitute aspectual nor intentional facts.[5] Given this, it would seem preferable to avoid this how question altogether. Obviously, the question is avoided on my account, since unconscious beliefs do not have aspectual shape, are not intentional, are not representational.

Why should we avoid this question? One reason relies on what Searle himself has argued, and just cited: behavioral or physiological facts cannot themselves constitute intentional facts. Unfortunately, this point wipes out too much, as it would seem to pose an equally formidable problem of explaining how *conscious* states could be intentional or representational. Though nothing I have yet said suggests a solution to this, I will address it later (Sections 2 and 3). It turns out that this is another version of the explanatory gap problem between brain and phenomenal states, discussed below. We will see that my application of a certain model of brain activity, supported by empirical evidence, provides the framework for answers to these questions. But even at this stage, there are grounds for treating conscious and unconscious states differently with regard to their representational nature.

There is strong evidence that some conscious states are representational, and there is no like evidence that any unconscious states are representational. Hardly anyone would dispute that at least some conscious mental states are representational. I submit that the reason this is so compelling is that it is evident from the first-person perspective that we each enjoy regarding some of our own states.[6] There is absolutely no corresponding datum, from neither the first- nor third-person perspective, regarding the representational nature of unconscious states. (Computational and biological models take for granted 'unconscious representations.' They do not offer direct evidence for such. To the extent that they are argued for at all, it is based on whatever success the models that employ them enjoy.) So representational unconscious states are at least initially problematic.

It is generally assumed in cognitive theories that a system's successful adaptive behavior is dependent on its having inner representations of its environment. Just how these inner representations are explicated varies considerably from theory to theory, but it is almost universally held that adaptive systems have representations and that they need not be conscious of them. Suppose, however, that the production of adaptive behavior can be explained by, say, nonrepresentational brain activity. This would undercut these appeals to a system's successful environmental behavior as support for unconscious representations. It would also be a neurophysiological basis supportive of my claim that unconscious beliefs are not intentional, not representational.

Such a neurophysiological model exists in the work of Walter Freeman.[7] Discussing that work, Christine Skarda argues that "... the patterns of neural activity responsible for behavior do not 'represent' anything, that brains do not 'read' them, and that 'neural representations' need not play a role in the production of behavior in animals." (1987: 189) She maintains that Walter Freeman's research on the olfactory bulb disconfirms the assumption that inner representations of the environment are required for adaptive behavior.[8] Surely, these are controversial claims; I will return to them in Section 3. My purpose in citing them is to remind us that there is an active empirical research program that denies the behavior of a system is dependent on the system's neuronal states *representing* the environment. So, though my contention that there are no unconscious representations is unorthodox, it is not without both philosophical and empirical support.

Consider the fact that we humans clearly do represent items in our environment. Do we not frequently do so unconsciously? For example, when we negotiate our environment to avoid obstacles without being conscious of them, does not this require *unconscious representations* of those obstacles and their relative positions to oneself?

Again, I can only suggest an answer here,[9] and I do so by appeal to the work of Kathleen Akins (1996). She distinguishes the *ontological* from the *sensory-motor project*. She argues that the former does involve representations of stable objects in the environment but the latter does not. After a careful examination of the neurophysiology involved in the sensory-motor project, she concludes that there is a rather large "... gap between the needs of the sensory-motor project and the demands of the ontological project ...". (1996: 370) The gap, she rightly insists, requires explanation. "... [H]ow exactly does the information provided by our sensory systems co-exist with, form a whole with, the ontology imposed by a representational system?" (1996: 370–371) She concludes: "Trace out the causal path between the object of perception, the stimulation of the receptors, and whatever neural events that thereafter eventuate [sensory-motor project] and this alone will not *explain*, in the required sense how *genuine representation* arises [ontological project]." (1996: 372, last emphasis added.)

The idea is that whatever occurs in the sensory-motor system, there is no straightforward interpretation of it such that it represents the stable objects in the environment, the objects of which we are conscious. There are no reliable correlations between types of physiological states and kinds of stimuli originating in the agent's environment. One might object that this does not show that the states of the sensory-motor system *themselves* are not representational. It could be maintained that these states *do* represent; it is just that *what* is represented is different from what is represented in the ontological project and at the conscious level. Of course, if Freeman is right, this objection is a non-starter. (As we will see in section 3, Freeman holds that the stimulus is "washed out".) But as I have yet to give sufficient details of this position, let us examine other ways of meeting this objection.

The objection itself is more powerful against me than it is against Akins, since all she argues is that certain traditional approaches to naturalizing aboutness that take the sensory-motor system as the basic aboutness relation will fail; whereas I maintain the stronger thesis that unconscious states of such a system simply do not represent. Period. Nevertheless, if one cannot get to the ontological project from the sensory-motor project, as Akins' work suggests, then at the very least, if "representations" are involved at the sensory-motor level, *they are of a very different sort than that at the conscious level*. And that may be all I need. Since what one obtains in pursuit of the sensory-motor project is not sufficient for the ontological project, it is not sufficient for genuine representation.

Now talk of 'genuine representation,' as Akins and I have, always raises suspicions as to what is so "genuine" about one's favorite analysis. But the point can be made innocently enough. What is important is not the word used, but the fact that "representations" in the sensory-motor project appear, at best, to be significantly different from those in the ontological project, and there seems to be no way of going from the former to the latter, for they are not reliably correlated. If the gap and the significant differences between the two senses are real, then the common practice of using the same term to cover both is a dubious one; so, the term 'representation' is better reserved for just one of these senses.

I have argued for an analysis of unconscious beliefs that allows them no intentionality. This stands apart from my more recent claim as to the nonexistence of unconscious representations. I have only offered considerations on behalf of this bolder claim here. But If I am right on all this, it would go far in explaining why there has been so much success in studies of cognition that ignore consciousness, even if consciousness is necessary for representation. These studies succeed when they focus on unconscious processes. It is a grave error, however, and one that is frequently made, to infer from such successes that results pertaining to unconscious "beliefs" or "representations" are transferable to conscious beliefs and representations, or that consciousness itself is epiphenomenal and not central to representation and cognition. Such errors stem from the deeply flawed, but widely held, assumption that the sense of *belief* and *representation* employed in discussions of conscious and unconscious states is the same. If I am right, this assumption is false.

2. Physicalism and the explanatory gap

Physicalism comes in a variety of forms. Common to all such forms is the view that everything is physical, and there is a physical explanation of every fact. Just what 'physical' means in such formulations varies. I will recommend below one account that draws its inspiration from a well-known episode in the history of physics.

It is widely agreed that our phenomenal states pose a particularly "hard problem" for physicalism, however construed, since there seems to be an "explanatory gap"[10] between, say, brain states and our phenomenal experiences. I will discuss an important solved analog to the explanatory gap problem. It provides us with a concrete understanding of just what it is we *ought to seek*, when we seek to close the explanatory gap. This case also demonstrates that an entailment relation is not necessary and so, contrary to some, need not be sought.

My account of physicalism and how the gap problem should be closed preserves the explanatory gap as the serious problem that it is, but shows that it poses no serious threat to physicalism, properly understood. Still, I will show that physicalism must avoid certain evasive, "ostrich" strategies, and that any appeal to fundamental mysteries also should be avoided. I also generalize the problem in that there is a gap not only between brain and phenomenal states, but also between brain and contentful mental states such as beliefs and desires.

What is the close analogy to the explanatory gap problem that has been solved? Consider an investigator at a relatively early stage of chemistry, perhaps a fictitious stage: Suppose it is known that H_2O molecules constitute water but little is known about intermolecular bonds. Molecules are (roughly) *discrete entities*, yet water presents itself as a *continuous and flowing stuff*, a liquid. When I use the term 'liquid' here, I do not mean it in its technical, physio-chemical sense; rather, it should be understood as a straightforward observation term. It is what a speaker of English would understand by the term, even if entirely ignorant of chemistry and physics. It is this sense of the term that is relevant, if the analogy is to work.[11] That is, the gap in the water case is between its discrete molecular constitution and its ordinarily observed continuous flow. Being so disparate in nature, how could the former give rise to the latter? We can understand this once we know that H_2O molecules can form weak intermolecular bonds, such that the bonds between large clusters of them can easily break and reform with other large clusters, the clusters themselves changing over time, thereby "sliding," as it were, across one another.[12] Furthermore, it is a familiar fact that small discrete entities in near proximity to one another may give the appearance of continuity, especially when viewed from a distance. Thus is dispelled the "mystery" that justifiably exists concerning the gap between the liquidity of water, its observed continuity and ability to flow, on the one hand, and its discrete molecular constitution, on the other.

Note that the actual theory of molecular bonding is quite complicated, the success of the explanation, however, does not turn on any arcane knowledge of that theory. *Only some of the simplest general features of the theory,* coupled with continuity considerations, *are needed to provide an adequate account of the liquidity of water*. With that slim knowledge, what is mysterious becomes *comprehensible to us*.

We should also note that this account is *not* tantamount to simply saying that liquidity *supervenes* on a certain type of molecular bonding, though it may be doing that too. Mere talk of supervenience without appeal to, say, intermolecular bonds and continuity considerations would be *just* to name the mystery; it would not provide a plausible "mechanism" for the supervening property.[13] While many different properties may 'supervene' on various sets of base properties, mere statement of a supervenience relation between the sets of properties is nearly vacuous. One must also state how the base properties *determine* the supervenient ones. The importance of specifying the dependence relation becomes all the more evident once it is realized these dependence relations will certainly vary from case to case. Thus, though weak intermolecular bonds may constitute the dependence relation in the molecule/liquid case, no one would expect that same dependence relation operative in the brain/mental state case.[14]

The gap problem between brain and mental states is so difficult not only because no one has identified any plausible mechanism connecting them that enables us to understand how the one *could* give rise to the other but, relatedly, because the features of the former are so radically different from the latter. A solution to the gap problem requires the *specification* of some "mechanism," something that could play a similar role to the intermolecular bonds in the molecule/liquid case. Merely saying that mental states supervene on brain states is almost a vacuous claim; it certainly advances no understanding.

Hypothesis: some unknown entities, states, processes, or features, **X**, of the brain are to phenomenal or mental states as intermolecular bonds are to liquidity. There may be no such **X**, but any physicalist research program, one of whose goals is to solve the explanatory gap problem, must seek such an **X**.

For definiteness, I will assume throughout that a solution lies in neuroscience. Whether current neurophysiological theories and approaches are sufficient to uncover such an **X**, if it exists, is an empirical question. With this caveat in mind, there seem to be three options:

1. Current neurophysiological resources are generally sufficient, though certain relevant details or connections among them have yet to be determined.
2. The current resources are not adequate — a scientific revolution in neurophysiology, of the kind Kuhn has described in other fields, is required.
3. There is no such neurophysiological **X**.

The current intractableness of the problem of the explanatory gap, while not conclusively establishing that option 1 is a dead end, does make it implausible. The analogy provides the reason for saying this: the solution to the molecule/liquidity gap did not rely on any technical or detailed knowledge of the character of intermolecular bonds, *but only on some of their simple most general aspects*. But while we have a rather good understanding of the firing rates of neurons and the various chemical neurotransmitters, we still *do not have a clue how such features could give rise to phenomenal qualities*. If specialized knowledge is not required and we already have the essentials of what we need, the gap problem should have been solved already, or at least the outline of its solution should be at hand. But it is not, and we do not even have a plausible line on it. Some exploit this recalcitrant fact to the extent of claiming that it shows physicalism is false, that is, they endorse option 3, or some generalized version of it, maintaining that it is an unsolvable mystery, beyond our capabilities. We shall see that such a conclusion is premature. For although option 1 currently seems unlikely, that is not the only way physicalism could be saved and the explanatory gap problem solved. There is still option 2–we need a new science.[15] For the reasons cited above, I reject option 1. (Later, I will give a more conclusive reason why option 1 is not tenable.) I further hold that, given a plausible understanding of physicalism, we must reject option 3. To see this consider first a simple-minded view of Physicalism, one no one would take seriously today. It holds that all things, events, states, or processes are explainable within a Newtonian framework. Clearly, physicalism thus understood is false. No contemporary would be tempted by such a narrow physicalism, even so, such a view would have been, indeed was, held by many in the late eighteenth and first half of the nineteenth century.

The history of this case is well known. A relevant problem then was to account for electromagnetic phenomena in Newtonian terms. (I note that I am not advancing this as a gap problem, though one might; I exploit it for its bearing on the concept of physicalism.) All such efforts failed. The Newtonian laws required supplementation with those of Maxwell. Was this a refutation of physicalism? No one, I think, would say so, not even then. Though it was a defeat for a physicalism narrowly understood in Newtonian terms, the concept was enlarged. To my mind, this was a rational reaction. Newtonian mechanics enjoyed huge successes and held out tremendous promise in its early years, but subsequent sustained scientific investigation determined that more had been promised than could be fulfilled; furthermore, the Maxwellian enhancement, which solved the target problems, was good science by all relevant standards. Subsequent scientific developments, the special and general theories of relativity and quantum electrodynamics, expanded our conception of the physical even more radically.

What bearing does all this have on ruling out option 3? The point is that any attempt to limit the concept of physicalism to our *current* successful scientific theories is a mistake. For among other reasons, it unjustifiably assumes that our current science has "got it essentially right." Not only is such a belief contrary to the inherent tentative nature of science, it is inductively unsound, given the various dramatic changes and upheavals in scientific theories over the years that has resulted in our expanded view of the physical.

Physicalism understood as restricted to current science is most probably false. If one adopts this restriction, then given the apparent hopelessness of closing the gap with the resources of current science, one is pushed toward option 3. I have attempted to show that such a restriction, however, is unjustified. If the exact character of physicalism is understood as characterized by our best science, not taken as an essentially finished product, but as ongoing and open-ended, then physicalism is very probably true. For thus understood, it is simply the stricture that we will not count any thing, process, state, or feature as explained until we have a scientific account of it, *whatever that is*. This makes physicalism a rather uncontroversial, indeed, trivial thesis, since all that is ruled out is non-scientific accounts, say, supernaturalism.

Given this conception of physicalism, endorsement of option 3, in its general interpretation, would be tantamount to giving up and declaring the explanatory gap to be an unsolve*able* mystery, as some have done. Such a stance bears a trace of arrogance in that it unjustifiably presupposes foreknowledge of what future science will bring, more accurately, what it will *not* bring. We *may* never have the requisite science to close the gap (though I think it is at hand, as I will explain in the next section); still, it is difficult to see why, at this stage of infancy of brain science we should draw the pessimistic conclusion that closing the gap is forever beyond our reach.[16]

A conception of physicalism constrained to some particular stage of scientific development, while more dramatic, is also far more problematic. Certainly, the restricted version to current science is quite dubious, since there is not even a single general and basic paradigm in neuroscience. Neuroscience has yet to produce its Newton or Einstein. What could decide in favor of such a narrow conception of Physicalism? Nothing short of its actual fulfillment and this, certainly, is not on the horizon.

Physicalism conceived as broadly as I have recommended, though so uncontroversial as to be trivial, has the virtue of shifting the attention to more productive discussions of the explanatory gap. Frequently physicalists and anti-physicalists alike view the explanatory gap as a threat to physicalism. The gap is only a threat if one adopts an unjustifiably narrow view of physicalism. Taking this narrower view, coupled with the current (apparent) hopelessness of the situation, results in a pernicious reaction to the explanatory gap by both sides. Physicalists tend to argue that the gap is not that big, that serious or even that it is an illusion, while antiphysicalists declare it is not solvable. If I am right, these defensive strategies on the part of physicalists, one might call them "ostrich strategies," are as misguided as the "mysterian" strategies of the antiphysicalist, for the narrow view of physicalism presupposed by both sides is unwarranted.

The explanatory gap is serious and — *at this time* — quite mysterious. That said, there is no reason to hide our heads in the sand or to be unduly and prematurely pessimistic. Brain science is, after all, still in its infancy. We have no reason to think that there can be no *new* understanding of the way the brain works that will close the gap between brain activity and mental states in a way analogous to how our understanding of intermolecular bonds closed the gap between discrete molecules and the fluidity of liquids. In short, to avoid 'ostrich' or "mysterian" strategies, we should endorse the wider conception of physicalism.

Some have rejected this line. For example, David Chalmers, has argued against the idea advanced here that we need a new science. (See his 1996: 118. See also Churchland 1996.) Chalmers' argument is that all physical theories come down to two basic elements: structure and dynamics of physical processes. "But from structure and dynamics, we can only get more structure and dynamics." (1996: 118) He concludes that new structure and dynamics will not help in closing the explanatory gap, since it is just more of the same.

Chalmers supplements this argument with two other requirements that he thinks advocates of a new science cannot fulfill. He says that it is difficult to evaluate the claim that it will take a new science in the abstract, "[o]ne would at least like to see an example of how such a new physics [science] might *possibly* go" (1996: 118). Furthermore, he holds that "[n]o set of facts about physical structure and dynamics can add up to a fact about phenomenology" (1996: 118).

A fair interpretation of the locution 'add up' is that of entailment, since on the very same page he raises the rhetorical question: "How could a theory that is recognizably a physical theory entail the existence of consciousness?" So the two further requirements are: A. a concrete proposal, or an indication of how the new science might go. B. The new science entails consciousness.

As to requirement A, I will present a specific proposal later. For now, I note that I have already indicated how a new science *could possibly go*. I did so by reminding us how science has gone with respect to a similar problem; namely, it uncovered an informative connection between the diverse elements of discrete molecules and continuous liquidity. But if one accepts requirement B, then one might protest that the analogy did not even accomplish that. For it does not close its own gap, since the structure and dynamics of molecules does not *entail* liquidity.

I am inclined to agree that the entailment relation is lacking in the analogy, though it may well be that there are a number of plausible missing premises that could be uncovered that would turn it into an entailment. Still, I am not tempted to mount such a defense, as I think the entailment requirement is too strong. Rather, I offer the analogy as *evidence* that the entailment relation is simply not necessary: we *do* see how understanding the loose intermolecular bonds render the gap between the disparate items *understandable to us without an entailment relation.* If an entailment relation is not necessary to close this particular gap, why should we require it *a priori* in the brain/mental states gap? Where is the argument for the claim that the relation that closes the gap between diverse entities must be entailment?

Let us suppose, as seems true, that the molecule/liquid gap is closed without an entailment relationship between the structure and dynamics of molecules and liquidity. What seems to be operative is some "plausibility relation" between the disparate items. Given these disparate items and the weak intermolecular bonds, it is *comprehensible to us* that such structure and dynamics could give rise to the observed liquidity.

The problem of the explanatory gap is just such a "comprehensible to us" problem. How can we understand how brain states, apparently so different from mental states, nevertheless, give rise to them? Requiring that the relevant scientific theories entail that there are mental states is, to my mind, to set too high a standard. I, for one, would think that the explanatory gap between these disparate items closed, if one could give a comparable account of their relations to that given in the case of molecules and liquidity. *If we could only just see how something new and different could arise out of the brain states*. Of course, none of this is to eschew entailment relations. All the better if they are produced. My point is simply that they are not necessary for closing the explanatory gap.[17]

There remains Chalmers' main objection. According to him, what is essential to a physical theory is structure and dynamics, and a new physical theory will just give us more of the same; so, a new physical theory will not help in closing the explanatory gap. But Chalmers casts his argument at far too general a level. The *particular character* of the structure and dynamics is what makes a difference. Just as I earlier argued that talk of supervenience is empty without specification of the relevant determination relation, so too is general talk of structure and dynamics. In the molecule/liquid analogy, it is the *particular character* (weak intermolecular bonds) of the *structure and dynamics of molecules* that enables us to understand how liquidity can arise out of such qualitatively different items. So, too, we may expect that the *particular* character of the new dynamics is what will enable us to close the explanatory gap between brain and mental states. Additionally, the analogy leads us to the hopeful expectation that detailed or highly sophisticated knowledge of the new dynamics may not be necessary to close the gap.

3. Chaos

It is the fact that mental states are so qualitatively different from brain states that gives rise to the conundrum of the explanatory gap. In short, we seem to get something totally new and different out of the activity of neurons and neurotransmitters. A significant contributing factor to this mystifying situation is the implicit assumption that the relevant mathematics is linear. For within such models, it is true that 'the whole is the sum of its parts.' You do not get more, and you certainly do not get anything different from the linear combination of the parts.[18] Mental states being so evidently different from brain states, showing how they could arise out of the latter seems to be an impossible, even logically impossible, undertaking, so long as inquiries are restricted to linear models. Until recently, nonlinear models were eschewed in scientific investigations because of their complexity, while linear models applied to the activity of neurons and neurotransmitters, for the reasons given, appear *logically inadequate* to close the gap. The situation, thus, seems desperate and is a compelling reason to reject option 1, listed earlier. In consequence, if one's sights are limited to linear models, option 3 appears inescapable.

Prospects are brighter when *self-organized* systems governed by nonlinear dynamics are considered.[19] These are open systems, some of whose properties are not completely explainable by properties of parts of the system. Consider a simple example. When oil is heated uniformly from below, the heat is randomly dissipated, and there is no qualitative change in the properties of the liquid. However, a qualitative change occurs once a certain temperature is reached: the liquid rapidly assumes a lattice structure of hexagonal convection cells (Bénard cells). This new behavior is not merely driven by input and initial conditions; the molecules collectively organize themselves. This collective behavior of large populations of elements of the system constrain the behavior of individual elements so that there is a kind of "top-down causality," contrary to the usual patterns of explanations. Such self-organized behavior governs the dappling patterns of animals' coats, hurricanes, cellular slime molds, and even in the formation of the visual cortex (Kelso 1997: 6–15).

From such cases, we may extract a unique and peculiar attribute of nonlinear interactions of elements that provides for at least the possibility of the closing the explanatory gap. It is the **X** of Section 2:

(i) The dynamics of self-organized systems described by nonlinear equations can yield properties of the system that are qualitatively different from any linear combination of its variables.

Nonlinear dynamics is the only known candidate possessing this feature, and this is exactly the kind of feature that can resolve the basic conundrum over the explanatory gap. For the closing of the gap seems to require getting something more and different out of interacting elements.[20] Freed from the constraint of linear models and given (i), what appeared to be a recalcitrant mystery is transformed into a manageable, if still formidable, problem. The enormous task remains to characterize the nonlinear equations of brain state variables so that something more and different, mental states, can arise out of their interactions. Viewing the brain as a self-organized system governed by nonlinear dynamics is the general core of a new paradigm of brain behavior that at least provides the framework for a solution, a solution that appears impossible against a background of linear models.

Noticing that (i) is exactly what is *needed* to do the job falls far short of actually producing the relevant nonlinear mathematics. Still, it is gratifying to realize the abstract point that there is this unusual and powerful feature of nonlinear dynamics that provides a *basis for our comprehension* of how brain states can give rise to mental states (even though they are so radically different from one another). This is rather like the situation with our earlier analogy. Knowledge of the general character of weak intermolecular bonds is sufficient for us to comprehend how liquidity can arise from a multitude of discrete entities; no sophisticated knowledge of the theory of molecular bonds is required. It is no small point that, in a similar way, a most serious logical obstacle to closing the explanatory gap between the disparate items of mental and brain states is swept away. Furthermore, we will see that there is some cause for optimism with regard to the existence of nonlinear models of brain activity.

In the discussion of the heated oil, we observed that the emergent property of the system, hexagonal convection cells, exerted a kind of top-down causality, properties of large collections of molecules constrain the behavior of individual molecules. This is a second powerful and unique feature of self-organized systems that provides the basis for removal of another major logical conundrum between brain and mind, mental causation. Here, certain high-level mental states, such as beliefs and desires, appear to activate certain microstates, the activation of motor neurons, which results in the agent moving to fulfill her desire. We may express the feature at issue as follows:

(ii) The dynamics of self-organized systems described by nonlinear equations yield global properties of the system that causally constrain microelements of the system.

Assume, for the moment, that conscious thoughts and perceptions are just these sorts of emergent properties of self-organized brains governed by nonlinear dynamics. (Two prominent neuroscientists, Earl MacCormac (1996: 151) and Walter J. Freeman (1991 and 1995, e.g., have so hypothesized, respectively). The idea that conscious thoughts or perceptions actually bring about certain behaviors by means of certain sensory motor responses, is *demystified* on this assumption, for we have simple analogical illustrations (Bénard cells discussed above, and the other examples cited) of how such emergent and qualitatively different states can constrain microstates.

So, a second apparently logically impossible realization of a relation between brain and mind becomes *comprehensible to us, given the new paradigm.* Again, this only holds out the abstract possibility, one whose promise could only be fulfilled after completion of enormous scientific work, work demonstrating the relevant nonlinear equations that support such speculations. But, once again, it is a most gratifying abstract possibility, as it transforms a seemingly logically impossible situation into one that is within our grasp.

Thus, I suggest that the conundrums of the explanatory gap and mental causation may well be simply artifacts of methodologies based in linear micro-reduction models.

Walter J. Freeman has successfully applied chaos theory to neurophysiological systems.[21] He did extensive work on the olfactory system of rabbits, which I will focus on, but he has found similar results in rats and cats, and in the visual system of monkeys (Freeman and Dijk 1987), cats (Eckhorn *et al.* 1988; Gray *et al.* 1989), and humans (Schippers 1990). Somewhat more tentatively, similar results were found in the somatosensory system of humans (Freeman and Mauer 1989). I must add that there is no way I can do justice to the extensive important work that Freeman has done in this area. I will only touch on a few highlights of his varied experimental and mathematical support for his model and strongly encourage the reader to "follow his nose" in pursuing Freeman's extensive research.[22]

There is a significant split among brain researchers as to the importance of relatively local behavior of neurons (feature neurons or "grandmother cells"), on the one hand, or the global activity of large populations of neurons, on the other, for understanding perception and cognition. Freeman apparently started in the former camp, but twelve years of experimentation using traditional methods on the olfactory system of the rabbit transformed him into a leading exponent of the importance of global activity and the application of chaos theory. His attempts to find significant relations between odorants and relatively local neural populations failed.[23]

The initial detection of an odor is by a spatial sheet of receptors in the nose. These receptors signal through the olfactory nerve another sheet of cells in the olfactory bulb. Although single neurons do fire in response to specific odors, the same cell fires in response to different odors and the same odor can excite many different cells. Freeman used electroencephalograms (EEGs) to record the activity. What he found was that the patterns of two successive sniffs of the same odor were as different from one another as were the patterns between two sniffs of different odors. The spatial patterns of the EEGs did not reliably correlate with the input stimuli. They changed not only when stimuli changed but when anything else changed. Change even occurred on reinforcement of an odor not previously reinforced. This change in reinforcement pattern for a particular odor had a ripple effect in that the spatial patterns for all the other odors also changed, indicating associative memory.

In brief, the receptors in the nose activate certain neurons in the olfactory bulb. The bulb then constructs a global spatial pattern that is dependent on past experience in that which neurons fire synchronously is a function of the specific past activations of neurons. (We will see below that the entorhinal cortex, thought to be involved with both memory and emotion, also plays a significant role in the construction of this pattern.) It is the global pattern of cooperative activity of neurons that is transmitted to the olfactory cortex and other areas. Freeman allows that the sheets of receptors in the nose and at the initial input to the olfactory bulb are feature detectors that refine sensory input, but they do not constitute perception. Importantly, the raw sensory data recorded in these sheets is "washed out," not transmitted as such (though it can be "recovered" if the need arises). Stimuli destabilize neuronal populations, causing them to construct patterns that express the *significance* of the stimuli to the animal based on its experience; it does *not* express the stimulus *per se*. (Compare Akins' (1996) discussion of heat, where no single correlation is found between the heat stimulus and the various resultant changes in physiological states.)

The EEGs for a particular odorant indicate an attractor. Although different sniffs of the same odorant produce differences in EEG as great as do different odors, the spatial patterns of EEG for a given odorant are similar, as remarked above. Each inhalation causes a burst of bulbar activity. Although each EEG tracing varies, a common waveform or carrier wave is embedded in the tracings, and the identity of the odorant is reliably identified in the bulbwide spatial pattern of the carrier wave amplitude. The aperiodic common carrier wave is everywhere in the bulb. It occurs not only during bursts but also between bursts, when there is no extractable stimulus. As Freeman says: "The carrier wave is not imposed from outside the olfactory system by the receptors or by other parts of the brain. It emerges as a cooperative activity that is self-organized by the neural masses." (1992: 468) This lack of external driving of the activity is what indicates that it is self-organized which, in turn, is characteristic of chaotic systems.

The varying set of different receptor cells in the nose that cause the global activity in the olfactory bulb to exhibit this spatial pattern is its basin of attraction. Thus, one attractor can be entered from a variety of starting points. I emphasize that the particular array of receptor cells excited varies considerably from one sniff to another, even when of the same odorant, but each odorant has its own attractor and basin. Input to the basal chaotic attractor causes a burst to the appropriate chaotic attractor for the input. To Freeman, chaos is essential, as it is the only way to account for the observed rapid shifts from one brain state to another.[24]

According to Freeman, the general dynamics of perception is as follows. The brain seeks information by directing an individual to sense. Self-organizing activity in the limbic system (a part of the brain that includes the entorhinal cortex and is thought to be involved in emotion and memory) directs the search by transmitting a search command to the motor systems and simultaneously sends what is called a reafference message to the sensory systems. The reafference message directs the sensory systems to prepare to respond to new information. The sensory systems strongly respond with a burst, every neuron in a given region participating in a collective activity. Resulting synchronous activity from these systems is then transmitted back to the limbic system, where they are combined and form what Freeman calls a 'gestalt'. In a fraction of a second, the process repeats. Freeman concludes: "... an act of perception is not the copying of an incoming stimulus. It is a step in a trajectory by which brains grows, reorganize themselves and reach into their environment to change it to their own advantage" (Freeman 1991: 85).

Generalizing on some earlier work of others (Helmholtz, von Holst, Mittlestadt and Sperry), Freeman sees all goal-directed movement as initiated by the limbic system, with the entorhinal cortex playing a central role. All sensory systems converge in the entorhinal cortex, and the entorhinal cortex is involved in memory and emotion. It was already noted that the pattern of the burst activity associated with a sniff was a function of experience and did not correlate with the external stimulus itself. What all this suggests is that the experience that shapes the burst activity is not simply a recording of what has happened, but includes its emotional value or significance to the animal.

One would expect that the significance of the smell of a fox would be different for a rabbit and a dog; for the one, memories of chase and fear, for the other, food and expectation of a meal. According to Freeman, the bursts correlated with such smells are unique to the individual. It is a commonplace that one's emotions can affect one's perceptions and thoughts. The model of perception provided by Freeman would seem to provide some physiological basis for such a commonplace, as the reafferent messages transmitted to the sensory cortex from the entorhinal cortex apparently reflect the significance of the stimulus to the particular animal based on its experience.

In concluding, let me remind the reader that in Section 2. I argued that a new paradigm of brain science is needed to close the explanatory gap. I argued that what was important in the new science is its *particular* structure and dynamics. I illustrated this in my recommendation of an analogy that provides a model of what we need to close the explanatory gap. We saw that the *particular* structure and dynamics must be such that they can make *comprehensible to us* how the interaction of certain elements and their properties can give rise to qualitatively different properties.

The new science that appears to provide a basis for closing the explanatory gap between qualitatively different mental and brain states is that of chaos theory applied to brain activity. The new *particular* feature that this theory provides, and which is unique to it, is presented above as (i). We found yet another special feature of this theory, (ii), which may well provide the basis to make comprehensible to us how there can be mental causation. Based on how the molecule/liquid analogy turned out, we expected and hoped that detailed knowledge of the new theory would not be required to make comprehensible to us how it can resolve the conundrums we face. Apparently, our hopes and expectations in this regard have been fulfilled. We may take further gratification in the fact that not only does chaos theory do all this, which is at a rather abstract level, but also it already has been successfully applied to concrete questions of brain activity.

Some have criticized the application of chaos theory to brain activity because unlike, say, computational models, there is no clear or significant place for representations in chaotic dynamics, and it is a widely held assumption that representations (both conscious and unconscious) play a pervasive role in cognition. The point of my first section is to diffuse this concern by arguing that there simply are no unconscious representations, though many take them for granted. Representations are central to *conscious* thoughts — but only there.[25] Where they are operative, they may well be manifested as bursts to attractor states. Such bursts constrain the firing of individual neurons in a top-down fashion; this is a manifestation of feature (ii). Furthermore, we know that these bursts are self-organized and are not driven by external stimulation. By feature (i), we may expect that the resulting burst states are significantly and qualitatively different from their antecedent states. This is exactly what we should expect if only conscious states are representational.

Appendix

A general and brief explanation of some of central concepts of chaos theory are here provided for the convenience of the reader. The relevant form of chaos is deterministic as opposed to entropic. Entropic chaos simply moves to increasing disorder without emerging patterns of activity in the system. Where there is deterministic chaos, similar patterns appear and disappear periodically in the dynamical activity of a system, but they do so with a complexity that makes it difficult to detect and, therefore, appear to be entropic. The complexity is due to the iterative nature of the equations governing the behavior. In iterative equations, the outcome of each computation is the input for the next computation. A characteristic feature of chaotic systems is that very slight differences in the starting point or weak input can lead to huge differences in eventual outcome. This extreme sensitivity to initial conditions is referred to as the 'butterfly effect.' This results in the extraordinary feature that the outcomes are not predictable, even though the equations are often quite simple and deterministic. Superficially, at least, this feature would seem to go far in explaining why, although we can frequently give an ad hoc explanation of someone's behavior, we are not as successful in predicting what someone will do, despite similar initial conditions.

A key concept required for our purposes is that of an *attractor*. When a dynamical system settles into a certain pattern, that pattern is known as an attractor; there are several kinds of attractors. The simplest attractor is the *single-point attractor*. The point at which a pendulum subject to friction eventually stops typically illustrates it. When a pendulum is not subject to friction, it will continue in exactly the same pattern of motion forever. Such repeating patterns are known as *limit-cycle attractors*. There are a number of other kinds of attractors, which like the one just defined, continue to repeat the same pattern of motion. Earlier I stated that similar patterns are repeated in chaotic systems. Characteristic of these systems, unlike the ones just discussed, is that they *never* traverse the same path, but the paths are similar and do trace recognizable patterns. Because of the uniqueness of each path, such patterns are known as *strange attractors*.

The set of initial values that terminate in an attractor is known as its *basin of attraction*. In chaotic systems, there are rapid *bursts* between strange attractors as well as back to the basal state. These bursts are called *bifurcations*; they constitute significant phase transitions in the system. The transition from rest to walking, or from walking to a trot, or in fluid dynamics, the shift from laminar flow to turbulence, are all common illustrations of bifurcations.

Notes

1. E.g., for each particular number, I unconsciously believe that it is a natural number.
2. Let me make it clear that I think Searle is, for the most part, profoundly right about intentionality, and though I think he is wrong in holding that unconscious beliefs are intentional, abandoning this does not necessitate wholesale changes in his view.
3. In arguing for the connection principle (1992: 156–62), Searle advances his claim that unconscious beliefs have aspectual shape in some detail. He has confirmed (in conversation) that something like the short argument here presented captures the thrust of the point contained in the larger argument of the broader conclusion.
4. The above would seem to undermine Searle's *connection principle:* "all unconscious intentional states are in principle accessible to consciousness ..." (1992: 156); there are no unconscious intentional states on my view. Still, a modified form of the connection principle survives; one that I believe preserves the spirit of Searle's principle. Space prohibits elaboration.
5. He thinks we would also need some law-like connections between the levels, as he says, "... we would still have to have some law-like connection that would enable us to infer from our observations of the neural architecture and neuron firings that they were realizations [for example] of the desire for water not of the desire for H2O," (1992: 158)
6. See my (1990) where I argue that certain third-person accounts fail in providing for this first-person access, and (1999) where I argue that Burge's original and influential thought experiment in favor of anti-individualism fails. My positive account is in my (1994) and (1996).
7. I discuss Freeman's work in some detail in section 3.
8. More fully, she says:"The patterns of neural activity that Freeman records in the EEG is the result of a self-organized, acquired tendency to behave in a certain way given the goals of the system considered as a whole ... The deciding factor in brain function is not neural activity patterns per se, ... not the input pattern and its transformation within the animal, but the internally generated neural dynamics created by the system itself." (1987: 198)
9. I argue for this at length in "Represent and Information Bearers", unpublished.
10. It is well known that David Chalmers (1996) is credited with coining the former expression and Joe Levine (1983) the latter.
11. It is important to stress this ordinary observational sense of the term, as David Chalmers objected to this analogy, used in a distant ancestor of this paper, based on a different interpretation. He says, "... the water case is not a useful analogy for you, because in that case it's clear that what needs to be explained is structure and function — we explain the liquidity, transparency, etc. and it's obvious that those are the sort of structural/functional matters that are wide-open to physical explanation." (12/13/98) For it to be "obvious" that these properties are open to a structural/functional explanation, one must presume, at the least, that these properties are understood in a physio-chemical way, rather than in the way I recommend here, a way that

corresponds to our ordinary observation of liquids. I will have more to say about the role of structure and function in these matters latter in this section.

12. Similar remarks apply to the solidity of ice and water vapor, where the strength of the intermolecular bonds are much stronger or weaker, respectively.
13. I have discussed this in my (1995).
14. I raised these points in criticism of John Heil's treatment of supervenience in my (1995) review of his (1992). Apparently, he now agrees, for in the conclusion of his (1998: 153–154) he states, "I have argued that the concept of supervenience as standardly formulated provides little in the way of ontological illumination ... [S]upervenience is not explanatory ... Supervenience claims hold, when they do, because the world is a particular way. What we need to be clear about is what that way is. Different cases yield different results."
15. As to the competition between options 1 and 2, one is reminded of Kuhn's remarks pertaining to the differing reactions of scientists to a paradigm under stress.
16. Colin McGinn (1989) has argued for such a pessimistic conclusion. Space limitations do not permit discussion here of his subtle arguments.
17. One reason someone might think they are necessary is if the explanatory gap problem is confused with a reduction of phenomenal states to brain states and one accepts the classical model of reduction of one kind of entity to another. I think it is a mistake to view the explanatory gap problem as a reduction problem, regardless of the model of reduction.
18. True, the whole may function differently than the parts individually, but the whole is not qualitatively different from its parts. Most functionalists (not without exception) admit their inability to reduce qualia, phenomenal states, to functional states. In my 1994 and 1996, I argue that they are also unsuccessful in reducing propositional contents to functional states.
19. Self-organization is a characteristic of chaotic systems. See Appendix.
20. J. A. Scott Kelso (1997) lists among the conditions for a self-organized system that there must be a large number of elements with nonlinear interactions. He points out that this requirement constitutes a major break with Sir Isaac Newton, whom he quotes as holding that "[t]he motion of the whole is the sum of the motion of all the parts." (Definition II, Principia) In contrast, Kelso holds "[f]or us, the motion of the whole is not only greater than, but different than the sum of the motions of the parts, due to nonlinear interactions among the parts or between the parts and the environment." (Kelso 1997: 16)
21. Detailed discussion of Freeman's work supporting the role of chaos in brain processing of information is in the papers by Chris King, Carl M. Anderson and Arnold J. Mandell, Earl R. MacCormac, and David M. Alexander and Gordon G. Globus in Stamenov and MacCormac (1996). Others also successfully apply chaos theory to the brain. For a fascinating and clear account of other successful applications of chaos theory to the brain, see Kelso (1997).
22. See, for example, his 1996, 1991, 1992 1995, a sequence that goes, roughly, from simpler to more technical works.
23. It is worthwhile to note that Freeman rejects the "binding problem" as a psuedo-problem, for it is based on the idea of feature detector neurons in each of the sensory cortexes, which he rejects because it is observer relative and what is significant for the brain is the global activity of neurons. (The binding problem is how the outputs from the different feature detectors are connected.) A proposed solution, one which Freeman himself had a hand in, is that feature neurons fire synchronously (40 Hz.). Aside from the difficulties with feature detector neurons,

this particular solution is now rejected by him since the frequency distribution in the visual and olfactory systems is not synchronized but broad spectrum, as Freeman has demonstrated.

24. A further fact that supports the contention that the system is chaotic is that the bulb and the olfactory cortex excite each other so that neither settles down, nor do they agree on a common oscillation frequency. (If the connection is severed between the two, they become stable and quiet.) This competition increases sensitivity and instability and contributes to chaos.

25. I regret that space limitations do not allow for further discussion of the implications of my thesis that there are no unconscious representations for Searle's Connection Principle.

References

Akins, Kathleen. 1996. "Of Sensory Systems and the 'Aboutness' of Mental States". *The Journal of Philosophy XCIII*: 337–372.

Alexander, David M., and Gordon G. Globus. 1996. "Edge-of-Chaos Dynamics in Recursively Organized Neural Systems". In MacCormac, Earl, and Maxim I. Stamenov, (eds), *Fractals of Brain, Fractals of Mind.* Amsterdam and Philadelphia: John Benjamins, 31–74.

Anderson, Carl M., and Arnold J. Mandell. 1996. "Fractal Time and the Foundations of Consciousness: Vertical Convergence of 1/f Phenomena From Ion Channels to Behavioral States". In MacCormac, Earl, and Maxim I. Stamenov, (eds.), *Fractals of Brain, Fractals of Mind.* Amsterdam and Philadelphia: John Benjamins, 75–126.

Chalmers, David. 1996. *The Conscious Mind.* New York: Oxford University Press.

Churchland, Patricia Smith. 1996. "The Hornswoggle Problem". http://www.merlin.com.au/brain_proj/psch_2.htm.

Eckhorn, R., R. Bauer, W. Jordan, M. Brosch, W. Kruse, M. Munk, and H. J. Reitboeck. 1988. "Coherent Oscillations: A Mechanism of Feature Linking in Visual Cortex?" *Biological Cybernetics 60*: 121–130.

Freeman, Walter. 1996. "Interview with Jean Burns: Societies of Brains". Journal of Consciousness Studies 3 (2):, 172–180.

Freeman, Walter J. 1995. *Societies of Brains: A Study in the Neuroscience of Love and Hate*. New Jersey: Lawrence Erlbaum Associates.

Freeman, Walter J. 1992. "Tutorial on Neurobiology: From Single Neurons to Brain Chaos". *International Journal of Bifurcation and Chaos* 2(3): 451–482.

Freeman, Walter J. 1991. "The Physiology of Perception". *Scientific American 264 (2)*: 78–85.

Freeman, W. J., and K. Maurer. 1989. "Advances in Brain Theory Give New Directions to the Use of the Technologies of Brain Mapping in Behavioral Studies". In Maurer, K (ed.), *Proceedings, Conference on Topographic Brain Mapping*. Berlin: Springer-Verlag, 118–126.

Freeman, W. J., and B.Van Dijk. 1987. "Spatial Patterns of Visual Cortical Fast EEG During Conditioned Reflex in a Rhesus Monkey. *Brain Research 422*: 267–276.

Georgalis, Nicholas. 1999. "Rethinking Burge's Thought Experiment". *Synthese* 18: 145–64.

Georgalis, Nicholas. 1996. "Awareness, Understanding, and Functionalism". *Erkenntnis 44*: 225–256.

Georgalis, Nicholas. 1995. "Review of The Nature of True Minds". *Philosophical Psychology 8 (2)*: 189–193.

Georgalis, Nicholas. 1990. "No Access for the Externalist: Discussion of Heil's 'Privileged Access'". *Mind 99*: 101–108.

Georgalis, Nicholas. 1994. "Asymmetry of Access to Intentional States". *Erkenntnis 40*: 185- 211.

Gray, C. M., P. Koenig, K. A. Engel, and W. Singer. 1989. "Oscillatory Responses in Cat Visual Cortex Exhibit Intercolumnar Synchronization Which Reflects Global Stimulus Properties". *Nature 338*: 334–337.

Heil, John. 1998. "Supervenience Deconstructed". *European Journal of Philosophy* 6 (2): 146- 155.

Heil, John. 1992. *The Nature of True Minds*. Cambridge: Cambridge University Press.

Kelso, J. A. Scott. 1997. *Dynamical Patterns*. Cambridge: MIT Press.

King, Chris. 1996. "Fractal Neurodynamics and Quantum Chaos: Resolving the Mind-Brain Paradox Through Novel Biophysics". In MacCormac, Earl, and Maxim I. Stamenov (eds), *Fractals of Brain, Fractals of Mind,* 179–234. Amsterdam and Philadelphia: John Benjamins.

Levine, Joseph. 1983. "Materialism and Qualia: The Explanatory Gap". *Pacific Philosophical Quarterly 64*: 354–61.

MacCormac, Earl R. 1996. "Fractal Thinking: Self-Organizing Brain Processing". In MacCormac, Earl, and Maxim I. Stamenov (Amsterdam and Philadelphia: John Benjamins, 127–154.

McGinn, Colin. 1989. "Can We Solve the Mind-Body Problem?" Mind 98:, 349–66.

Schippers, B. 1990. "Spatial Patterns of High-Frequency Visual Cortical Activity During Conditioned Reflex in Man". Masters Thesis, Laboratory for Medical Physics, University of Amsterdam, Netherlands.

Searle, John. 1992. *Rediscovery of the Mind*. Cambridge: MIT Press.

Skarda, Christine. 1987. "Explaining Behavior: Bringing the Brain Back". *Inquiry 29*: 187–202.

In the series ADVANCES IN CONSCIOUSNESS RESEARCH (AiCR) the following titles have been published thus far or are scheduled for publication:

1. GLOBUS, Gordon G.: *The Postmodern Brain.* 1995.
2. ELLIS, Ralph D.: *Questioning Consciousness. The interplay of imagery, cognition, and emotion in the human brain.* 1995.
3. JIBU, Mari and Kunio YASUE: *Quantum Brain Dynamics and Consciousness. An introduction.* 1995.
4. HARDCASTLE, Valerie Gray: *Locating Consciousness.* 1995.
5. STUBENBERG, Leopold: *Consciousness and Qualia.* 1998.
6. GENNARO, Rocco J.: *Consciousness and Self-Consciousness. A defense of the higher-order thought theory of consciousness.* 1996.
7. MAC CORMAC, Earl and Maxim I. STAMENOV (eds): *Fractals of Brain, Fractals of Mind. In search of a symmetry bond.* 1996.
8. GROSSENBACHER, Peter G. (ed.): *Finding Consciousness in the Brain. A neuro-cognitive approach.* n.y.p.
9. Ó NUALLÁIN, Seán, Paul MC KEVITT and Eoghan MAC AOGÁIN (eds): *Two Sciences of Mind. Readings in cognitive science and consciousness.* 1997.
10. NEWTON, Natika: *Foundations of Understanding.* 1996.
11. PYLKKÖ, Pauli: *The Aconceptual Mind. Heideggerian themes in holistic naturalism.* 1998.
12. STAMENOV, Maxim I. (ed.): *Language Structure, Discourse and the Access to Consciousness.* 1997.
13. VELMANS, Max (ed.): *Investigating Phenomenal Consciousness. Methodologies and Maps.* n.y.p.
14. SHEETS-JOHNSTONE, Maxine: *The Primacy of Movement.* 1999.
15. CHALLIS, Bradford H. and Boris M. VELICHKOVSKY (eds.): *Stratification in Cognition and Consciousness.* 1999.
16. ELLIS, Ralph D. and Natika NEWTON (eds.): *The Caldron of Consciousness. Motivation, affect and self-organization - An anthology.* 2000.
17. HUTTO, Daniel D.: *The Presence of Mind.* 1999.
18. PALMER, Gary B. and Debra J. OCCHI (eds.): *Languages of Sentiment. Cultural constructions of emotional substrates.* 1999.
19. DAUTENHAHN, Kerstin (ed.): *Human Cognition and Social Agent Technology.* 2000.
20. KUNZENDORF, Robert G. and Benjamin WALLACE (eds.): *Individual Differences in Conscious Experience.* 2000.
21. HUTTO, Daniel D.: *Beyond Physicalism.* 2000.
22. ROSSETTI, Yves and Antti REVONSUO (eds.): *Beyond Dissociation. Interaction between dissociated implicit and explicit processing.* n.y.p.
23. ZAHAVI, Dan (ed.): *Exploring the Self. Philosophical and psychopathological perspectives on self-experience.* 2000.
24. ROVEE-COLLIER, Carolyn, Harlene HAYNE and Michael COLOMBO: *The Development of Implicit and Explicit Memory.* n.y.p.
25. BACHMANN, Talis: *Microgenetic Approach to the Conscious Mind.* n.y.p.
26. Ó NUALLÁIN, Seán (ed.): *Spatial Cognition. Selected papers from Mind III, Annual Conference of the Cognitive Science Society of Ireland, 1998.* n.y.p.
27. McMILLAN, John and Grant R. GILLETT: *Consciousness and Intentionality.* n.y.p.
28. ZACHAR, Peter: *Psychological Concepts and Biological Psychiatry. A philosophical analysis.* n.y.p.
29. VAN LOOCKE, Philip (ed.): *The Physical Nature of Consciousness.* n.y.p.

These new studies by prominent neuroscientists, psychologists and philosophers work toward a coherent framework for understanding emotion and its contribution to the functioning of consciousness in general, as an aspect of self-organizing, embodied subjects. Distinguishing consciousness from unconscious information processing hinges on the role of motivating emotions in all conscious modalities, and how emotional brain processes interact with those traditionally associated with cognitive function. Computationally registering/processing sensory signals (e.g. in the occipital lobe or area V4) by itself does not result in perceptual *consciousness*, which requires subcortical structures such as amygdala, hypothalamus, and brain stem. This interdisciplinary anthology attempts to understand the complexity of emotional intentionality; *why* the role of motivation in self-organizing processes is crucial in distinguishing conscious from unconscious processes; how emotions account for 'agency'; and how an adequate approach to emotion–motivation can address the traditional mind–body problem through a holistic understanding of the conscious, behaving organism.

John Benjamins Publishing Company
ISBN 90 272 5136 3 (Eur) / 1 55619 196 0 (US)