 Mind & Matter Vol. 4(1), pp. 45–68

(Proto-) Consciousness as a Contextually Emergent Property of Self-Sustaining Systems

J. Scott Jordan
Department of Psychology
Illinois State University, Normal, USA
and
Marcello Ghin
Faculty of Cultural Sciences
University of Paderborn, Germany

Abstract

The concept of contextual emergence has been introduced as a specific kind of emergence in which some, but not all of the conditions for a higher-level phenomenon exist at a lower level. Further conditions exist in contingent contexts that provide stability conditions at the lower level, which in turn afford the emergence of novelty at the higher level. The purpose of the present paper is to propose that (proto-) consciousness is a contextually emergent property of self-sustaining systems. The core assumption is that living organisms constitute self-sustaining embodiments of the contingent contexts that afford their emergence. We propose that the emergence of such systems constitutes the emergence of *content-bearing* systems because the lower-level processes of such systems give rise to and sustain the macro-level whole (i.e., body) in which they are nested, while the emergent macro-level whole constitutes the context in which the lower-level processes can be for something (i.e., be functional). Such embodied functionality is necessarily and naturally *about* the contexts that it has embodied. It is this notion of self-sustaining embodied *aboutness* that we propose to represent a type of content capable of evolving into consciousness.

1. Introduction

Many current theories of consciousness make use of the concept emergence. Although they differ when it comes to the details, all such theories assume consciousness emerges out of the properties of some physical system, be it a brain region (Crick 1994), a synchronized pattern of activity across multiple brain regions (Edelman 1989, Engel and Singer 2001), or a dynamic synergistic nesting of brain, body and world (Clark 1997, Myin and O'Regan 2002, O'Regan and Nöe 2001, Thompson and Varela 2001, Varela *et al.* 1991). Having framed the issue this way, the tendency is to then work toward determining the necessary and sufficient conditions underlying such emergence.

While, on the one hand, framing the problem of consciousness in terms of emergence from physical systems seems straightforward, there is some uncertainty in the field about the meaning of the distinction. Specifically, there seems to be a number of meanings to the concept *physical*. Atmanspacher and Kronz (1998), for example, describe a range of meanings of *physical* (as well as its synonym *material*) that extends from its ontological utilization in *physicalism*, "... the idea that the basis of reality consists of the material world alone; anything like qualia, consciousness, psyche, mind, or spirit is based on the material elements and fundamental laws of physics" (p. 282), to its epistemological utilization in *methodological dualism*, which "... utilizes the mind-matter distinction as a basic, but maybe not the only possible methodological tool to inquire into the structure of the world" (p. 283). To be sure, most theories utilize the concepts *physical, mental* and *emergence* in the ontological sense – that is, they assume there exist real, observer-independent physical "things", and real conscious mental states that emerge from such "things". Atmanspacher and Kronz (1998) describe this practice in terms of ontic and epistemic descriptions, and assert that making ontic assertions about lower-level phenomena so as to get an epistemic grip on higher-level phenomena is a successful strategy in science.

While we agree that the onticity of an ontic description must be specified as *relative to* another, higher-level description (ontic relativity, compare Quine's (1969) notion of relativistic ontology), we also believe that ontic assumptions about the physical, in emergence-based theories of consciousness, can lead one to implicitly take the stability of the physical as given, much in the way one assumes stability at the chemical level when making ontic assumptions about chemistry so as to get an epistemic grip on biology. Doing so may lead one to inadvertently overlook the possibility that stability itself constitutes an important property of consciousness – not in terms of stable physical properties *giving rise to* consciousness but, rather, in terms of dynamic, self-sustaining processes *being* consciousness. While this may smack of reductionism and identity theory – positions that emergence was meant to discount – we propose it constitutes a case of *contextual emergence* (Atmanspacher and Bishop 2006) in which some, but not all of the conditions for a higher-level phenomenon exist at a lower level. Further conditions exist due to contingent contexts that provide stability criteria at the lower level which simultaneously afford the emergence of novelty at the higher level. Conceptualized this way, consciousness is not so much a higher-level phenomenon, which emerges from lower-level stability, as it is a self-referential, self-sustaining embodiment of the higher-level contingent contexts a system has to address (i.e., embody) in order to sustain itself. In short, we propose that consciousness constitutes a contextually emergent property of self-sustaining systems.

2. The Contextually Emergent Nature of Self-Sustaining Systems

In his classic work on autocatalytic systems, Kauffman (1995) proposed that living systems emerged by a phase transition in the pre-biotic soup from systems of chemical reactions unable to sustain themselves to chemical systems capable of self-sustainment. Self-sustainment was possible because such systems were autocatalytic; the reactions in the system produced their own catalysts. Thus, as long as raw materials were available, autocatalytic networks of chemical reactions were able to sustain themselves.

The emergence of autocatalytic systems constituted a case of *contextual emergence* in the sense of Atmanspacher and Bishop (2006). While certain necessary conditions existed at the lower level (i.e., the catalytic properties of various chemicals), the emergence of an autocatalytic network required additional contingent contexts – specifically, the ratio of the number of possible reactions between chemicals to the number of diverse chemical types available in the pre-biotic soup. According to Kauffman (1995), once this ratio reached a value $r = 0.5$, chemical systems underwent a phase transition from non-autocatalytic to autocatalytic. This constitutes emergence because the macro-structure (e.g., an autocatalytic network as a whole) comes to be, spontaneously, out of the catalytic nature of certain chemical reactions. It constitutes contextual emergence because such macro-structures are only able to emerge within the larger-scale contingent context of a specific ratio of possible reactions and diverse chemical types. Thus, while the catalytic properties of certain chemical reactions are necessary, they are not sufficient. Only within the contingent context of a critical ratio r can autocatalytic networks emerge and sustain themselves.

Once such self-sustaining systems emerged, they eventually gave rise to a new context that afforded the emergence of more sophisticated self-sustaining systems. Specifically, as more and more self-sustaining systems came to be (i.e., single cell organisms) the chemical energy encapsulated within them provided a potential fuel source for any system capable of capturing such energy and using it to sustain itself. In this example of contextual emergence, the existence of single-cell organisms constitutes the lower-level necessary conditions, while the large-scale availability of such systems constitutes the contingent context that affords the emergence of larger-scale organisms capable of sustaining themselves on the energy encapsulated in single-cell systems. Large-scale availability represents a contingent context because it provides the stability conditions under which a single-cell consumer can emerge and sustain itself. Without this contingent context (i.e., stable fuel supply) a single-cell consumer is not possible.

This example suggests an algorithm for the recursive scaling-up of contextually emergent self-sustaining systems. The lower-level necessary conditions are to be found in the self-sustaining properties of certain types of *work* (i.e., energy transformation). At the single-cell level, self-sustaining *work* refers to the autocatalytic properties of certain chemical systems. At the level of the single-cell consumer, self-sustaining work refers to the energy-transformations the system undergoes (e.g., swimming, chewing and digesting) to capture and release the energy entailed in the single-cell system. At both levels of scale (i.e., the chemical and the biological) the work produces a product (i.e., a catalyst in the single-cell case, and the release of encapsulated energy in the single-cell-consumer case) that feeds back into and sustains the work that produced the product. Once a plethora of such lower-level self-sustaining work is available, this wide-scale availability constitutes a contingent context providing stability conditions (i.e., a stable fuel source) for the emergence of even more sophisticated self-sustaining systems. Thus, the self-sustaining work of plants (i.e., the conversion of electromagnetic radiation into chemical energy) provides lower-level conditions, given by the contingent context of the wide-scale availability of encapsulated plant energy, for the emergence of herbivores and, subsequently, carnivores.

To be sure, different scholars have alluded to this recursive scaling-up of self-sustaining energy-transformation systems in different ways for different reasons. Boltzmann (1905) and Lotka (1945) cast evolution as the struggle for available energy. Wiener (1948) conceptualized organisms as thermodynamically open energy transformation systems in order to conceptualize purposive systems (versus non-purposive systems), and Schrödinger (1945) did so while attempting to conceptualize life. Recently, Odum (1988) and Vandervert (1995) conceptualized nature as a self-organizing energy-transformation hierarchy, and even applied this notion to the emergence of human knowledge.

Given that all these theories converge on the notion of an energy-transformation hierarchy, one might question our focus on contextual emergence and claim that it is more appropriate to model the entire energy-transformation hierarchy as one large self-sustaining system that is ultimately reducible to the lower-level properties of autocatalytic systems. Odum (1988), for example, models nature in such a way in order to quantify and model the nature of the energy exchange between different levels within the system. We believe, however, that there is something important to be gained by focusing on contextual emergence and self-sustainment at multiple nested levels. For, while the influences exerted on a self-sustaining system by its local context may be describable in terms of low-level local effects, the self-sustaining system simultaneously embodies the contingent context of the large-scale availability of its fuel source. It does not embody this context in the sense that it "knows" of such availability in any way, or has maps of such availability within itself. Rather, its ability to sustain itself, *in and of itself*, is testament to (i.e., embodies) the availability of its fuel source.

A single cell, for example, is a form of self-sustaining work that sustains itself as a whole on a time-scale that emerges from and is larger than the time scale of the chemical systems of which it is constituted. The dynamics at this larger time scale (e.g., swimming and tumbling) differ from the dynamics of the lower time-scale (i.e., chemical reactions). And the sustainability of the larger global context (i.e., the single cell as a whole) across a larger time scale is testament to (i.e., embodies) the wide-scale availability of its fuel source (which collectively exists at a larger time scale than that of individual chemicals). It is our position that this emphasis on contextual emergence and the idea that organisms constitute *self-sustaining embodiments of the contingent contexts that afforded their emergence* provides a point of entry into the consciousness debate that is not available by treating the entire energy-transformation hierarchy as one system.

3. Self-Sustainment, Contextual Emergence, and Consciousness

In Sec. 1, we mentioned that most emergence-based approaches to consciousness make ontic assertions in which conscious mental properties are assumed to emerge from physical properties. We further claimed that framing the relationship between the concepts *physical, mental* and *emergence* in this way might lead one to take the stability of the lower level (i.e., the physical in this case) as given and, as a result, overlook the possibility that stability itself might constitute an essential property of consciousness – not in terms of stable physical properties *giving rise* to consciousness but, rather, in terms of self-sustaining processes *being* consciousness. In what follows we clarify why we think this is the case.

3.1 Function, Content, Aboutness, and Process

When addressing the relationship between internal states and consciousness, many theories focus on the *function* of internal states (Davidson 1987, Millikan 1984, Adams and Dietrich 2004). Most theories of such function have teleological features and assume that internal states gain their status as having function because of the evolutionary and developmental history of the system bearing the internal states. According to Davidson (1987), however, these teleological theories ultimately lead to non-functional internal states because they leave open the logical possibility of an organism that spontaneously emerges from a convergence of molecules in the air. Given the organism's lack of evolutionary or developmental history, its internal states would not qualify as functional. In light of this problem, Bickhard (2001) defines *function* in terms of self-sustaining systems. That is, the internal states of self-sustaining systems have function precisely because they give rise to and sustain the macro-level whole in which they are nested (i.e., the organism). Having defined function in this manner, it is logically impossible for the internal states of a self-sustaining system to be non-functional.

We agree with Bickhard's approach to function. Within our framework of contextually emergent self-sustaining systems, we propose that the emergence of such functional internal states constitutes the advent of content-bearing systems or, as we refer to them, systems with *proto-consciousness* (see Bickhard (2001) for a variant of this approach to content and self-sustainment). To be sure, there are many different approaches to the notion of content. What we mean by proto-consciousness is that the content of self-sustaining systems constitutes a type of content that is ultimately able to evolve into consciousness.

The basis for this claim is that, within such contextually emergent self-sustaining systems, the micro-level work (i.e., chemical reactions) gives rise to (i.e., is *for*) the macro-level whole (i.e., the network as a whole) which, synergistically, provides a sustained context in which the micro-level work can continue. Within the context of such *self-sustaining recursive reciprocity* a process (i.e., work) constitutes *content*. This emphasis on self-sustaining recursive reciprocity and micro-level work being for the macro-level whole that it gives rise to is similar to Boden's assertion that "... metabolism is a type of material self-organization which, unlike the Belousov-Zhabotinsky reaction, involves the autonomous use of matter and energy in building, growing, developing, *and maintaining the bodily fabric of a living thing*" (Boden 1999, p. 237, italics added). Content emerges and sustains itself in the midst of self-sustaining metabolism.

Given this approach to content, we further propose that the *meaning* of such content derives from the contingent contexts that a self-sustaining system has to embody in order to sustain itself. An autocatalytic network, for example, embodies the larger-scale contingent context (i.e., the critical ratio between possible chemical reactions and the number of diverse chemicals) in which it emerges and sustains itself. In this sense, its internal states are *about* that larger-scale contingent context in a very real way. In our view, this "about" does not refer to "messages" about external states carried by internal states. Rather, it is identical with self-sustaining processes of being that emerge from context, depend on context, and sustain themselves within context. In other words, *aboutness* in our approach refers to embedded, embodied, self-sustaining context (Jordan 2003).

At first glance, this position is difficult to accept because most theories treat *aboutness* or *content* as a property a system has *in addition* to its physical properties. We believe this is because most emergence theories treat the concept physical as ontic, take stability for granted, and assume that mental properties emerge from these stable physical properties. We propose, however, that at the ontic level these systems constitute *self-sustaining processes* emerging from the context of processes in which they are embedded. This reveals how such processes can simultaneously be self-referential (i.e., be *about*) and have function (i.e., be *for*). *Content* and *aboutness*, in our view, are not mental properties emergent from physical systems. Rather, they are holistic, ontic properties of embedded, embodied self-sustaining process. The concepts *physical* and *mental*, if used at all in our framework, are utilized in a methodologically epistemic way, as a means of illustrating the apparent epistemic divide between processes that *feel* as though they depend on me for their existence (i.e., what those who take the Cartesian divide as ontic refer to as *mental* properties) and processes that don't feel as though they depend on me for their existence (i.e., what those who take the Cartesian divide as ontic refer to as *physical* properties).

Atmanspacher (2006) takes a somewhat similar approach when attempting to relate particular *complexity* measures to the notion of *meaning*. Specifically, he bases this approach on Peirce's (1958) semiotic theory, in which signs have meaning because they are embedded in a genuinely triadic system of relations between the sign, its object, and its interpreter. He then argues that Peirce's focus on the notion of genuine embedded relations leads to an ultimately holistic theory of meaning. Atmanspacher refers to this as "implicit meaning" and proposes that perhaps the concepts *complexity* and *meaning* can be considered as complementary notions. "A complementarity relation between two (or more) concepts typically indicates that the respective concepts share important features at a level of description underlying that at which the complementarity relation applies" (Atmanspacher 2006, p. 87). According to this idea, both meaning and complexity are implicit at the holistic level at which they are embedded, and from which they emerge as explicitly different concepts when one attempts to get an epistemic grip on the phenomena distinguishing the complexity of physical properties from the meaning of mental properties.

As stated above, within our approach to *content* and *aboutness* the distinction between physical and mental properties constitutes a methodological tool rather than ontological categories (Jordan 1998a, 2000a). Theories assuming an ontic Cartesian distinction have difficultly addressing *meaning* and *aboutness* properly because, after dividing the world into two kinds of properties, they have to fit one back into the other, or have one emerge from the other. In many variants of corresponding maneuvers, the concept *causality* is paired with the physical, and mental properties are rendered *epiphenomenal*. For instance, Chalmers (1996, p. 150) writes:

> The physical world is more or less causally closed, in that for any given physical event, it seems that there is a physical explanation (modulo a small amount of quantum indeterminacy). This implies that there is not room for a non-physical consciousness to do any independent causal work. It seems to be a mere epiphenomenon, hanging off the engine of physical causation, but making no difference in the physical world.

Avoiding the temptation of an ontic Cartesian divide we allow for *aboutness* and *content* to constitute properties of the natural world. They are not properties of physical "things", however, and they do not emerge from physical "things". Rather, they are holistic embedded properties of self-sustaining processes which, themselves, are holistic and embedded. As regards whether or not such properties are causally connected, the logic of the theory leads to the possibility that the concept of causality may, itself, be epistemic rather than ontic. That is, it may be the case that causality is a methodological conceptual tool (Hume 1977) that we utilize when we attempt to understand the phenomena in which we find ourselves embedded. Thus, the realist ideal of doing science in order to reveal the true causal relations of things might actually be somewhat misguided (Manicas and Secord 1983).

Though this may seem far-fetched, it is actually consistent with Atmanspacher and Kronz's (1998) assertion that ontic statements about "things" constitute category mistakes insofar as "things" do only exist in contexts, and not as context-independent entities. Since "things" cannot be ontic in a context-independent fashion, statements about their causal relations are also context-dependent, and there is no such thing as "true" properties plus *ceteris paribus*. Seen in this light, the issue of whether or not *content* and *aboutness* are causal may simply be misconstrued. *Content* and *aboutness* are holistic embedded self-sustaining processes that constitute self-sustaining embodiments of the contingent contexts that afford their emergence.

3.2 Scaling-up of Content and Aboutness

Having defined *content* and *aboutness* in this fashion, the next step is to see how this framework gets us to *consciousness*. The mechanism for scaling-up from *content* to *consciousness* has already been described. Specifically, it is a form of contextual emergence in which a given level of self-sustaining work constitutes necessary lower-level conditions, while the widespread availability of such systems constitutes the contingent context that provides the lower-level stability conditions from which a higher-level self-sustaining system can emerge. Plant life constitutes a necessary condition for the emergence of herbivores, while the widespread availability of plant life constitutes a contingent context that affords their emergence. Herbivores therefore constitute self-sustaining embodiments of the contingent contexts that have to be addressed in order for a system to sustain itself off of the energy encapsulated in plants. For example, the teeth of a rabbit embody the constraints that have to be addressed to release the energy entailed in plants. The rabbit's digestive system embodies the chemical dynamics of released plant energy. Its neuromuscular architecture embodies the geographic distribution of plant energy. And its circulatory system embodies the energy consumption required to sustain the neuromuscular architecture needed to capture plant energy. A rabbit, therefore, constitutes a self-sustaining embodiment of the contingent contexts that afford its emergence.

During the recursive contextual emergence of carnivores, one of the contingent contexts that has to be addressed and embodied is the fact that the carnivore's fuel source (i.e., the herbivore) is mobile. To overcome this context, carnivores had to propel themselves, as a whole, along *anticipatory* pursuit curves. Data indicate that certain aspects of this contingent context are embodied in cerebral-cerebellar loops that allow organisms to embody regularities in motor-command-feedback cycles (Kawato *et al.* 1987). By embedding command-feedback cycles in cerebral-cerebellar loops, organisms are able to control their locomotion at time scales faster than that afforded by feedback garnered via kinesthetic sensors. This is because kinesthetic feedback requires roughly 120 milliseconds to reach motor cortex, while cerebral-cerebellar loops require only 10–20 milliseconds. Clark (1997) and Grush (2004) refer to these latter loops as *virtual feedback* because they reflect regularities in command-feedback cycles that can *stand in for* regular feedback and afford locomotion at faster time scales.

For the purposes of the present paper, virtual feedback is important because it represents a case in which a contingent context (i.e., the mobility of a fuel source) affords the emergence of virtual content. Virtual, in this case, does not refer to being unreal, or not existing. It means that the content is *about* a state that does not exist in the present context (i.e., the relationship between the predator's current body-world state and the *future* location of the prey). In addition, it has *function* and *content* precisely because it is necessary for the sustainment of the predator system, and it is *about* the contingent context (i.e., capturing moving prey) in which it emerges, respectively.

An interesting and important aspect of virtual content is that it is inherently *social*. That is, it is not about command-feedback patterns *per se* but about command-feedback patterns and their relationship to the patterns of *others*. This is the case because others (i.e., prey, predators and con-specifics) constitute an aspect of the contingent context that affords its emergence. In addition to cerebro-cerebellar loops, this assertion is supported by the discovery of areas in the brain (i.e., mirror neurons) that are active both when one plans a goal-related action and when one observes another execute such an action (Rizzolatti *et al.* 2002). What this means, in short, is that when a person produces a goal-directed action, he/she simultaneously puts an observer's brain in a planning state for the same goal-related action. These planning states are, by definition, about the future, and are thus inherently virtual. Given the inherently social nature of virtual content, its phylogenetic appearance constitutes the contextual emergence of anticipation-laden virtual social worlds in which prey, predators and con-specifics reciprocally constrain and contextualize each other's planning content.

Once such self-sustaining, inherently social, virtual content emerges, it constitutes a contingent context that affords the emergence of more abstract embodied content. That is, the virtual content embodied in a neural network is, itself, available for "capture" in the sense that emerging neural circuitry can tap into it and use it to sustain itself (Grush 2004). Emerging neural circuitry sustains itself via Hebb's (1949) notion of the cell assembly. This notion recognizes self-sustainment at the level of single neurons in a neural network: The work of being a neuron (i.e., producing action potentials) sustains itself due to the reciprocal influence of connected neurons. Edelman (1989) recognized the autocatalytic, self-sustaining nature of neural networks within the developing brain as a whole and referred to it as *neural Darwinism*. Thus, emerging neural circuitry, just like any contextually emergent self-sustaining system, embodies the constraints that afford its emergence. The content of these newly emerging systems would be *about* the content already embodied in the neural systems (i.e., contingent contexts) that afford their emergence.

In short, this describes the emergence of virtual meaning *about* virtual meaning. An example would be the neural circuitry that generates and sustains the concept "buffalo". Within our framework, the concept "buffalo" constitutes a self-sustaining embodiment of the contingent contexts addressed while a system works to sustain organism-buffalo relationships – a phenomenon usually referred to as perception (see Barsalou 1999). The "work" in these neural circuits (usually referred to as concept rehearsal, underlying our thoughts) is self-sustaining in the sense that the work allows the organism to generate organism-buffalo simulations (i.e., virtual-content sustainment) that allow it to capture more energy for less work. That is, the energy expended in running simulations is much less than the energy expended actually acting out alternatives. Once these regularities have been embedded, they too constitute a contingent context that affords the emergence of more recursively nested self-sustaining networks. As a consequence, the organism-environment coordinations an organism can sustain become increasingly virtual and recursively bootstrap themselves to the point at which systems emerge that are capable of sustaining relationships with virtual contexts such as "tomorrow" and "next month".

At some point it actually becomes more energy efficient to change the environment according to one's needs by "unpacking" tailored regularities into the environment. As early humans used rocks to scratch an image of a buffalo onto a cave wall, they were unpacking embedded regularities. And the "work" of the structures enabling such unpacking autocatalytically paid for itself in different ways. For example, the resultant unpacked regularity (i.e., a symbol) allowed individuals to create and manipulate complex external virtual structures, thus saving the high amount of energy it takes to sustain these structures internally – Clark (2001) refers to this as *cognitive off-loading.* In addition, the unpacked regularities allowed groups to share virtual content. As a result, each member was afforded the opportunity to generate virtual content within the context of a shared, unpacked simulation, coordinate individual actions around the same virtual content (i.e., a group plan) and, ultimately, capture more energy per unit work than was possible either alone or without the benefit of shared virtual content.

Sharing such virtual content is possible because it is inherently social in the first place. That is, it is about sustaining oneself in a world of prey, predators and con-specifics. What originally emerged as embodied (i.e., internalized), *other*-relative virtual content, autocatalytically bootstrapped itself, via contextual emergence, into externalized *we*-relative virtual content that afforded the sustainment of increasingly distal joint futures. This idea of a shared, we-relative virtual framework based on simulation seems similar to Gallese's notion of a shared manifold of intersubjectivity that "... relies on a specific functional mechanism, which is probably also a basic feature of how our brain/body system models its interactions with the world: *embodied simulation* ..." (Gallese 2003, p. 517, italics added).

Eventually, these inherently social, unpacked embodiments of content (i.e., symbols), which autocatalytically emerged from a shared manifold of intersubjectivity, constitute a contingent context that affords the emergence of more energy-efficient unpacking systems such as written and verbal language (Vandervert 1995). They also afford the emergence of more sophisticated forms of what Metzinger (2003) refers to as a *phenomenal self-model* (PSM). According to Metzinger, the PSM is almost exclusively generated by internal input from proprioceptive systems and constitutes the ongoing self-sustaining dynamic process of establishing a stable, centered, emotional embodied being. The PSM is integrated into a global model of the world that results in a *virtual*, self-sustaining self-world border, which enables the system to build representations of the relations between itself and the world, itself and specific objects, itself and other selves, and itself and itself.

The emergence of such content constitutes a necessary condition for the emergence of phenomena such as *mineness* and *selfhood* (Ghin 2005). And according to our perspective, the higher-level contingent context that affords its emergence is the wide-spread availability of externalized virtual content. For, in order to intake externalized virtual content, a system has to be able to distinguish virtual content generated *within* the system from virtual content generated *outside* the system. Thus, just as the proteins that make-up the wall of a single cell are inherently *about* those aspects of the contingent context that have to be discriminated (i.e., toxins) in order for the system to sustain itself at the chemical level, PSMs are likewise *about* those aspects of the contingent context (i.e., virtual content externalized by others) that have to be discriminated in order for the system to sustain itself at the social level.

Here then, is our take on how self-sustaining systems recursively bootstrap themselves from content-bearing single-cell systems to full-blown phenomenal selves. At every step in this recursive process, self-sustaining embodiments emerge, constitute necessary conditions for further emergence, become widely available, and give rise to the contingent contexts that afford the emergence of another recursion on the theme of self-sustaining process. In short, the stairway from content to consciousness is an inherently social, self-referential process of abstraction from less to more abstract forms of self-sustaining embodied meaning.

Within this framework one can conceptualize consciousness as self-sustaining aboutness and assert that all self-sustaining systems constitute consciousness to some extent. Or, one can define consciousness in a particular way (e.g., as the having of qualia, in opposition to "unconsciousness", or as "reportable" experience) and then assert that the defined form of "consciousness" exists at a particular level within a self-sustaining system's structural hierarchy (i.e., within a particular level, or perhaps across several levels). Given our assertion that content exists at all levels within self-sustaining systems, we are more inclined to go with the former over the latter option. By doing so, however, we are not denying the utility of the latter. To the contrary, our framework provides an ontic-epistemic space in which one can compare the utility of the two options.

4. Implications of Contextually Emergent Self-Sustaining Systems

The notion of contextually emergent self-sustaining systems holds many implications for consciousness studies specifically, and cognitive science in general (Jordan 1998a,b, Jordan 2003). In the following we examine a few such points of contact, and limit ourselves to those that seem most likely to clarify the position that we present in this contribution.

4.1 Information Processing Theory

Common to both the present position and information processing theory (IPT) is the notion of content-bearing internal states. While IPT tends to refer to such states as representations, we refer to them as self-sustaining embodiments. As mentioned above, the advantage of conceptualizing internal states in terms of self-sustaining embodiments is that this provides an account of *function* (Bickhard 2001) and *content* (Jordan 2003, Jordan and Ghin 2006) which avoids problems encountered by most teleological accounts. It does so because the internal states are enmeshed in a self-sustaining recursive reciprocity in which internal processes give rise to a macro-level whole that recursively constitutes a context in which the internal work sustains itself. And the meaning of the content derives from the contexts the system has to embody in order to sustain itself.

Self-sustaining embodiment (versus representation) also provides a different account of how internal states are coupled with each other and the external environment. According to IPT, *mental representations* constitute a distinct level in a system's structural hierarchy. Due to vertical separation (among levels in the hierarchy) and loose horizontal coupling (between component processes at the same level), representations are assumed to entail measurable causal properties that play a role both horizontally (i.e., between representation-processing components) and vertically (i.e., between sensory inputs and representations-*perceptions*, and between representations and outputs-*actions*). The assertion of measurable causal properties is based on the further assumption that the level at which mental representations express their causal properties is effectively isolated from higher and lower levels such that "... changes on different time-scales may be separated in terms of their causal implications – we may isolate causal properties on different timescales" (van Orden and Holden 2002, p. 91).

By contrast, the notion of self-sustaining embodiment assumes that, while levels of embodiment are vertically separated (e.g., single neurons, neural networks and the brain as a whole), they interact in a synergistic, interdependent fashion. This is because the higher levels emerge from, and are sustained by, the work of the lower levels, while the higher levels synergistically constrain (i.e., stabilize) the lower levels from which they emerge. Within this framework, *sustainment* and/or *stability* do not refer to static states, but rather to phenomena as dynamic as the sustainment of spatio-temporal coherence among the levels of work within the system. Nikolaev *et al.* (2005) for example, measured EEG phase synchronizations in the alpha band across different electrode sites (e.g., Cz and Pz) and found that the duration of moments of synchrony and de-synchrony between pairs of electrodes varied in a highly-complex, ordered (i.e., non-random) fashion that was constant across all subjects. Such multi-scale interdependence is characteristic for systems that exhibit dynamical co-operative behavior and flexibility.

While the differences between these two approaches to vertical separation and horizontal coupling might at first seem purely semantic, the highly recursive nature of neural interaction (Edelman 1989) and its continuous modulation by the environment and the organism itself makes it increasingly difficult to sustain the notion of a distinct hierarchical level, i.e. a cognitive band (Clark 2001), containing causally discrete representations – be they symbolic or distributed – whose content is clearly representative of discrete bodily or environmental events. To be sure, it is the case that the activity in certain brain areas is highly correlated with specific types of bodily and environmental events. But it does not follow from these correlations that those areas "represent" those events in the form of discrete, efficacious codes. According to the notion of self-sustaining embodiments, such correlations are perhaps better conceptualized as *resonances* (Jordan 2004) among the multiple levels of self-sustaining systems that constitute organisms (i.e., neurons, neural networks, the brain, and the body) and the environmental context in which the system is embedded.

The advantages of the latter option can perhaps be clarified by examining the ontic assumptions each position makes about internal states. IPT grants ontic status to mental representations, while the notion of self-sustaining systems grants ontic status to self-sustaining process. Having granted ontic status to mental representations, IPT then has to account for how such representations garner their status as being functional as well as how they come to entail content. As stated earlier, taking self-sustaining embodied processes as ontic, as opposed to representation, solves these problems. In addition, it has the major payoff of overcoming the reference problem (Kronfeld 1990). This is because the internal states of self-sustaining systems constitute embodiments (i.e., internalizations) of the contingent contexts in which they are embedded and must sustain themselves. In short, the *inside* is an embodiment of the *outside*. Thus, there is no need to ask how internal states map onto external states. Diachronically, they are embodiments of those external states. And any moment-to-moment synchronies between internal and external events constitute resonances, not representations.

To be sure, most cognitive scientists use the concept representation in an epistemic (versus an ontic) way. That is, they use the concept as a means of getting a grip on how phenomena we call "mental" are related to phenomena we call "physical". As long as this epistemic framework proves empirically useful, its use will sustain itself. The framework gets into trouble, however, when it treats representation in an ontic fashion, because this brings about the problems of function, content and reference. Given that the notion of self-sustaining embodiment can address these issues while simultaneously providing a useful epistemic grip on the role of internal states, parsimony and coherence may ultimately favor its use as a preferred framework for cognitive science.

4.2 Embodied Cognition

While IPT theorists tend to frame the issue of content in terms of the correspondence between internal and external states (Adams and Dietrich 2004), researchers espousing the embodied cognition perspective prefer to do so in terms of on-going dynamic interactions between brain, body and world (Clark 1997, Myin and O'Regan 2002, O'Regan and Nöe 2001, Thompson and Varela 2001, Varela *et al.* 1991). A major issue facing the embodiment position is that it does not explain why consciousness necessarily accompanies dynamic interactions among brain, body and world. In other words, it is logically possible for all the dynamic interactions to go on without any accompanying consciousness – the zombie problem (Chalmers 1996).

Our approach based on self-sustaining systems overcomes this problem and grounds the notion of embodiment. Once content and consciousness are conceptualized in terms of self-sustaining, self-referential embodied processes, it becomes logically impossible for an exact copy of such a system to lack content. In addition, it becomes clear why grounding consciousness in dynamic interactions between brain, body and world is a good direction to take. It is so because brains and bodies constitute self-sustaining embodiments of those aspects of the world (i.e., contingent contexts) that afford their emergence. In a sense, one could say that brains and bodies constitute *embodied world*.

4.3 Autocatalysis and Autopoiesis

As an early proponent of the embodiment position, Varela proposed autopoesis (versus autocatalysis) as the dynamic underlying the emergence of living systems (Ruiz-Mirazo and Moreno 2004). According to his view (Varela and Bourgine 2002, p. 5),

> ... an autopoietic system is organized (defined as a unity) as a network of processes of production (synthesis and destruction) of components such that these components (i) continuously regenerate and realize the network that produces them, and (ii) constitute the system as a distinguishable unity in the domain in which they exist.

We focused on autocatalysis instead of autopoiesis because the notion of autocatalysis allows us to flesh out both what it means for a system to be embodied and why such embodiment constitutes content (and ultimately consciousness). In addition, when we use the abstract notion of autopoiesis, we have to say that a civilization is an autopoietic system. One way to overcome this problem would be to require that the system be embodied and continuously reiterate the processes that sustain its existence. But then we would need a notion of embodiment that is independent of the notion of autopoiesis. Otherwise, we end up with a circular definition of autopoiesis. One way to avoid this problem would be to add either the notion of *strong metabolism* as proposed by Boden (1999), or the notion of autocatalysis as proposed by Kauffmann (1995). Both stress the importance of the work done by self-sustaining systems. But the advantage of Kauffmann's analysis is that it shows why we get something like strong metabolism in the first place.

In addition, our focus on autocatalysis as self-sustaining work sets us up to look for a means of scaling-up that is not conceptually grounded in biochemistry, per se. That is, what is common to single-cell organisms, rabbits, and PSMs is that they constitute self-sustaining work embodying the contingent contexts that afford its emergence. Thus, by focusing on autocatalysis and the notion of self-sustaining work, versus autopoiesis and the notion of living systems, we are able to ground embodiment and content in a way that clarifies the homological relationship between the nested levels of meaningful embodiment that constitute living systems.

4.4 Types of Content

One thing we have not done in the present paper is to distinguish different types of content. That is, we have not clarified how mental content, intentional content and phenomenal content differ (see Metzinger 2003). We have not done so for two reasons. First, our goal was to develop a framework that naturalizes all the above-mentioned types of content. Given we define content as self-sustaining embodied *aboutness*, and then propose that content scales-up via self-referential recursions, just as self-sustaining systems do, it seems to be the case that these various types of content also emerge in a self-referential recursive fashion. Thus, within our framework, they do not seem to constitute different types of content, as much as they constitute *levels* of content that range on a gradient from the proto-content of the first autocatalytic systems to the full-blown phenomenal content of self-sustaining *selves.*

Second, while we agree with the distinctions between mental, intentional and phenomenal content, we take some issue with the labels. Specifically, in our framework, all content is *intentional* in the sense that it is *about* the sustainment of the system in which it is embedded. In fact, Jordan (2000b, 2003) has used this idea to argue that self-sustaining systems constitute *embodied intentionality.* Thus, we agree that content becomes recursively abstract as self-sustaining systems phylogenetically bootstrap themselves from having *body-world* models to having *self-self* models (Metzinger 2003). But we do not see the recursions being distinguished in terms of mental, intentional and phenomenal – they are all intentional. And what distinguishes them is the degree of recursion they constitute.

To be sure, there are some forms of content that seem difficult to fit into this notion of self-sustainment (e.g., the smell of a rose, or the belief that $2 + 2 = 4$). The point of our thesis is not to account for the content of all possible embodiments, or to assume that every possible embodiment must, in some way, be traceable to immediate system sustainment. As a matter of fact, once somatic markers emerge, i.e. emotions (Bechara and Damasio 2005, Damasio 1999, Ghin 2005), they create a context in which virtual content, i.e. thoughts and memories, can be generated and sustained simply for emotional reasons. That is, one can engage in thoughts and memories simply because it feels good.

With this context in place, content is able to emerge that, *at first glance*, might seem un-related to immediate sustainment, and some of it, in fact, may not be (e.g., suicidal thoughts, the technology of nuclear power). Certain forms of virtual content (e.g., religion, sciences, arts), however, produced and continue to produce products that have had such a profound impact on our ability to sustain ourselves in terms of social organization, technology and entertainment, respectively, that they have basically "paid" for themselves as well as for *the entire enterprise of generating and sustaining virtual content.* Thus, while some forms of virtual content may seem irrelevant to immediate sustainment, the self-sustaining processes that make virtual content possible, are anything but.

4.5 Living Systems Only?

According to our position, a micro-macro synergy entails*content* if it constitutes a self-sustaining embodiment of the contingent contexts that afford its emergence. This is because the emergent self-sustaining process is necessarily *about* the context of processes from which it emerges and within which it sustains itself. Within this framework, *content* and *phenomenal selves* constitute extremes on a continuum of embodied meaning (i.e., self-sustaining embodied aboutness).

An immediate implication of our position is that all self-sustaining systems, be they living or non-living, entail content. We think doing so requires to clearly distinguish *self-sustaining* systems from *self-organizing* systems and *self-regulating* systems. The state of any one molecule in a self-organizing system such as a convection roll is *about* the state of all the other molecules in the system as well as the macro-level convection roll that emerges from the interactions among molecules. In this sense, *about* refers to continuous, reciprocal causality. Self-organizing systems, therefore, can be said to entail *aboutness*.

In micro-macro synergies that are self-*sustaining* (as opposed to simply self-*organizing*), the micro-macro transformations entail *aboutness* in both the causal sense and an additional sense. Specifically, they are about the contingent contexts they have to embody in order to sustain themselves. Again, this means that the micro-level work gives rise to (i.e., is *for*) the macro-level whole that synergistically provides a sustained *context* in which the micro-level work can continue. It is within the context of such *self-sustaining recursive reciprocity* that a process (i.e., work) constitutes *content*.

Thus, within our framework thermostats, convection rolls and mountains have *aboutness* but lack*content* because they lack the property of self-*sustainment*. The micro-level transformations in a convection roll are not *for* the convection roll because their work is not *about* the sustainment of the convection roll. Rather, the micro-macro synergies between the component molecules and the macro structure of the convection rolls emerge and disappear with the flow of energy through the system. The system dissipates energy, but does not entail dynamics that can keep it far from thermodynamic equilibrium. Turn off the heat and the convection roll disappears.

Another difference between *self-sustaining* and *self-organizing* systems is that the former are *self-regulating*. The convection rolls that emerge in heated fluids are not self-regulating because the convection rolls themselves have no means of offsetting perturbations to the state of the convection roll. Convection-roll dynamics do not regulate. The autocatalytic micro-macro synergy of a single cell (as well as the micro-macro synergies of all self-sustaining systems), however, is capable of offsetting perturbations and is, therefore, self-regulating. Specifically, the micro- and macro-states of the cell (i.e., the molecular structures comprising the whole and the cell as a whole, respectively) are synergistically enmeshed in such a way that certain transformations in the former give rise to phase transitions in the latter (i.e., abrupt shifts in macro-level organization) that actually serve to keep the former in particular states of transformation. The micro-level cell-wall proteins that allow food particles into the cell, for example, are kept in a state of transformation we refer to as food-intake via the macro-level transformations of the whole cell – what we refer to as swimming or tumbling, depending upon the species. These macro-level phase transitions emerge out of the micro-level transformation states that comprise the cell as a whole. As a result of this micro-macro synergy, the cell is capable of regulating the inputs to its cell-wall proteins.

Within our framework, self-regulating servo-mechanisms such as thermostats entail *aboutness*, but not*content*. This is because we characterize content as a property of self-sustaining work, in which self-sustaining internal states are *for* the macro-whole they give rise to and sustain. The transformations taking place inside a thermostat, or a digital computer for that matter, do not give rise to and sustain the macro-whole in which they work. They do regulate certain of their own states, but they do not do so via self-sustaining recursive reciprocity. As a result, the work taking place in the system is not for the system itself and, therefore, does not constitute *content*.

There is another type of self-organizing system that needs to be addressed; specifically human-made neural nets. In this context, Atlan (1991) defines *meaning* in terms of self-organization and then distinguishes between self-organization in the weak sense, in which a goal is set-up outside the system, and self-organization in the strong sense, in which the goal is an emergent property within the system itself. Atlan uses this distinction to argue that weakly self-organizing systems, such as human-made neural nets, entail *meaning* (i.e., content) because their structure is *functional* (i.e., is good for something). Meaning with true intentionality, however, according to Atlan, emerges in systems entailing strong self-organization.

It is tempting to relate Atlan's approach with its distinction between weak and strong self-organization to our distinction between self-organizing and self-sustaining systems. For in the latter, the goal is actually the macro-level whole that the nested micro-level work sustains. Thus, our position grants what Atlan refers to as *meaning*, and what we refer to as*content*, to a system that is self-organizing in the strong sense as long as the self-emergent goal of the system is self-sustainment. Human-made neural networks, from our perspective, do not yet entail content because the work going on in such systems is not yet self-sustaining. When the work taking place in a human-made system constitutes a functional architecture whose work (i.e., energy expenditure) sustains itself, our position will have to grant it content.

Boden (1999) makes a similar argument when she claims that artificial-life systems such as genetic algorithms running inside a computer or robots negotiating an environment do not constitute life. Her basis for this claim is their lack of a self-sustaining metabolism. Thus, even though such systems are capable of doing things that living systems can do, such as learn (i.e., internalize the context in which they sustain themselves), self-replicate, mutate, compete and evolve, Boden argues they are still not living because self-sustaining metabolism constitutes too fundamental a property of living systems for those who lack it to be granted the status of living.

We agree with Boden's emphasis on metabolism and conceptualize it in terms of self-sustaining work. We do so because this allows us to scale up from the bio-chemical self-sustaining work of single cells to the psychosocial self-sustaining work of social interaction. And, whereas Boden claims systems that lack such self-sustaining work are not alive, we would argue they also lack content.

4.6 Future Recursions?

In addition to providing an alternative approach to living systems and consciousness, the notion of self-sustaining contextual emergence provides a slightly changed view on evolution. Our position implies that self-sustaining systems came into being not because they were "selected", but because they were able to sustain themselves. That is, their work (i.e., the energy they expended) paid for itself by producing products that sustained the work. As a result, there are many varieties of body, brain and world that have phylogenetically emerged and sustained themselves. Ants and bees, for example, also constitute self-sustaining embodiments of the contexts that afford their emergence.

Unlike ants and bees, however, humans have the ability to embody and externalize *virtual* content. As a result, they have been able to affect massive transformations of the contexts in which they sustain themselves and, for the most part, prolong their sustainment. Odum (1988) and Vandervert (1999) model the emergence of minds and cultures in such terms. As an example, they argue that the "higher-quality" energy embodied in the work of a teacher increases the quality of the work embodied in a student, while the work of the student, via feedback (i.e., educational outcomes), reciprocally increases the quality of the teacher's work. And all of this reciprocal work sustains itself because it produces products (i.e., embodiments of higher-quality energy, more sophisticated minds) that sustain the work – more sophisticated minds become participants in culture.

Conceptualizing the evolution of such systems in terms of "sustainment" versus "selection" has the advantage of consistently reminding us of the precarious, recursively dependent nature of being a self-sustaining system. In addition, it reminds us that the process is still ongoing.

5. Conclusions

In conclusion, the concept of contextual emergence (Atmanspacher and Bishop 2006) allows us to conceptualize content as the self-sustaining embodiment of contingent contexts, i.e. as self-sustaining aboutness. This framework addresses many problems encountered by contemporary theories of consciousness. It provides a meaningful framework for representationalism and the reference problem by showing that internal states constitute embodiments of external states as contingent contexts. It grounds the idea of consciousness as a property of ongoing dynamic interactions between brain, body and world by demonstrating that brains and bodies constitute embodiments of world as contingent contexts. And it overcomes the zombie problem by demonstrating that content, defined as self-sustaining embodiment, constitutes a property of self-sustaining systems. Thus, any exact copy of such a system will necessarily have content.

Perhaps most importantly, conceptualizing content in terms of the self-sustaining embodiment of contingent contexts may provide the field of consciousness studies a way to conceptualize the Cartesian divide epistemically, as a methodological tool, rather than ontically, as a division of reality into physical and mental domains. This is because, within our framework, we take *process*, and not physical and mental properties, to be ontic. This leaves us with a holistic account of meaning in which self-sustaining embodied process constitutes content because it emerges from, embodies, and sustains itself within the context of processes in which it is embedded. This allows one to see the distinction between the *physical* and the *mental* as a methodological tools. In addition, it shifts the focus of a science of consciousness from a search for causal neural correlates to a search for the laws of dynamic self-sustainment.

Acknowledgments

This work was partially supported by a grant to the first author from the Zentrum für interdisziplinäre Forschung (ZiF) der Universität Bielefeld. We thank Harald Atmanspacher, Michael Spivey, Cees van Leeuwen and Andrew Bailey for their insightful and constructive comments on an earlier version of this paper. We also thank the members of the ZiF research group *Embodied Communication in Humans and Machines* – Ipke Wachsmuth, Manuela Lenzen, Günter Knoblich, Jens Alwood, Elizabeth Ahlséen, Thomas Metzinger, Karl Grammer and Jürgen Streeck – for many informative conversations regarding the present work.

References

Adams F. and Dietrich L. (2004): Swampman's revenge: Squabbles among the representationalists. *Philosophical Psychology* **17**(3), 323–340.

Atlan H. (1991): Intentional self–organization in nature and the origin of meaning. In *Ecological Physical Chemistry*, ed. by C. Rossi and E. Tiezzi, Elsevier, Amsterdam, pp. 311–331.

Atmanspacher H. (2006): A semiotic approach to complex systems. In *Aspects of Automatic Text Analysis*, ed. by A. Mehler and R. Kohler, Springer, Berlin, pp. 79–91.

Atmanspacher H. and Bishop R. (2006): Stability conditions in contextual emergence. *Chaos and Complexity Letters*. In press.

Atmanspacher H. and Kronz F. (1998): Many realisms. In *Modeling Consciousness Across the Disciplines*, ed. by J.S. Jordan, University Press of America, Baltimore MD, pp. 281–305.

Barsalou L. (1999): Perceptual symbol systems. *Behavioral and Brain Science* **22**, 577–660.

Bechara A. and Damasio A. (2005): The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior* **52**, 336–372.

Bickhard M.H. (2001): The emergence of contentful experience. In *What Should be Computed to Understand and Model Brain Function?*, ed. by T. Kitamura, World Scientific, Singapore, pp. 217–237.

Boden M.A. (1999): Is metabolism necessary? *British Journal for the Philosophy of Science* **50**, 231–238.

Boltzmann L. (1905): *The Second Law of Thermodynamics*, Reidel, Dordrecht.

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

Clark A. (1997): *Being There: Putting Brain, Body, and World Together Again*, MIT Press, Cambridge MA.

Clark A. (2001): *Mindware. An Introduction to the Philosophy of Cognitive Science*, Oxford University Press, New York.

Crick F. (1994): *The Astonishing Hypothesis: The Scientific Search for the Soul*, Simon and Schuster, New York.

Damasio A.R. (1999): *Descartes' Error: Emotion, Reason and the Human Brain*, Grosset/Putman, New York.

Davidson D. (1987): Knowing one's own mind. *Proceedings of the American Philosophical Association* **60**, 441–458.

Edelman G.M. (1989): *Neural Darwinism: The Theory of Group Neuronal Selection*, Oxford University Press, Oxford.

Engel A.K. and Singer W. (2001): Temporal binding and the neural correlates of sensory awareness. *Trends in Cognitive Sciences* **5**, 16–25.

Gallese V. (2003). The manifold nature of interpersonal relations: The quest for a common mechanism. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* **358**, 517–528.

Ghin M. (2005): What a Self could be. *Psyche* **11**(5), 1–10.

Grush R. (2004): The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences* **27**, 377–442.

Hebb D.O. (1949): *The Organization of Behavior: A Neuropsychological Theory*, Wiley, New York.

Hume D. (1748/1977): *An Inquiry Concerning Human Understanding*, The Bobbs-Merrill Company, Indianapolis.

Jordan J.S. (1998a): Intentionality, perception and autocatalytic closure: A potential means of repaying psychology's conceptual debt. In *Systems Theories and a Priori Aspects of Perception*, ed. by J.S. Jordan, Elsevier, Amsterdam, pp. 181–208.

Jordan J.S. (1998b): Recasting Dewey's critique of the reflex–arc concept via a theory of anticipatory consciousness: Implications for theories of perception. *New Ideas in Psychology* **16**(3), 165–187.

Jordan J.S. (2000a). The world in the organism: Living systems are knowledge. *Psycoloquy* **11**, #113.

Jordan J.S. (2000b): The role of "control" in an embodied cognition. *Philosophical Psychology* **13**, 233–237.

Jordan J.S. (2003): The embodiment of intentionality. In *Dynamical Systems Approaches to Embodied Cognition*, ed. by W. Tschacher and J. Dauwalder, Springer, Berlin, pp. 201–228.

Jordan J.S. (2004): The role of 'pre–specification' in an embodied cognition. *Behavioral and Brain Sciences* **27**(3), 409–409.

Jordan J.S. and Ghin M. (2006): Born to be wild: Grounding embodiment and content in selfsustaining systems. Submitted manuscript, available from the authors.

Kauffman S. (1995): *At Home in the Universe*, Oxford University Press, New York.

Kawato M., Furukawa K. and Suzuki R. (1987): A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics* **57**, 169–185.

Kronfeld A. (1990): *Reference and Computation*, Cambridge University Press, Cambridge.

Lotka A.J. (1945): The law of evolution as a maximal principle. *Human Biology* **17**, 167–194.

Manicas P.T. and Secord P.F. (1983): Implications for psychology of the new philosophy of science. *American Psychologist* **38**, 399–413.

Metzinger T. (2003): *Being No One. The Self-Model Theory of Subjectivity*, MIT Press, Cambridge MA.

Millikan R.G. (1984): *Language, Thought, and Other Biological Categories*, MIT Press, Cambridge.

Myin E. and O'Regan J.K. (2002): Perceptual consciousness, access to modality and skill theories: A way to naturalise phenomenology? *Journal of Consciousness Studies* **9**(1), 27–45.

Nikolaev A.R,. Gong P., and van Leeuwen C. (2005): Dynamical properties of global phase synchronization patterns in human EEG. Manuscript available from the authors.

Odum H.T. (1988): Self-organization, transformity, and information. *Science* **242**, 1132–1139.

O'Regan J.K. and Nöe A. (2001): A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* **24**(5), 939–1011.

Peirce C.S. (1958): *The Collected Papers*, Vol. 1-6 ed. by C. Hartshorne and P. Weiss, Vol. 7-8 ed. by A.W. Burks, Harvard University Press, Cambridge.

Quine W. (1969): *Ontological Relativity and other Essays*, Columbia University Press, New York.

Rizzolatti G., Fadiga L., Fogassi L., and Gallese V. (2002): From mirror neurons to imitation: Facts and speculations. In *The Imitative Mind*, ed. by A. Meltzoff and W. Prinz, Oxford University Press, New York, pp. 247–266.

Ruiz-Mirazo K. and Moreno A. (2004): Basic autonomy as a fundamental step in the synthesis of life. *Artificial Life* **10**(3), 235–260.

Schrödinger E. (1945): *What is Life?* Macmillan, New York.

Thompson E. and Varela F. (2001): Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Science* **5**, 418–25.

Vandervert L.R. (1995): Chaos theory and the evolution of consciousness and mind: A thermodynamic-holographic resolution to the mind-body problem. *New Ideas in Psychology* **13**(2), 107–127.

Vandervert L.R. (1999): Maximizing consciousness across the disciplines: Mechanisms of information growth in general education. In *Modeling Consciousness Across the Disciplines*, ed. by J.S. Jordan, University Press of America, Lanham MD, pp. 3–25.

Van Orden G.C. and Holden J.G. (2002): Intentional contents and self–control. *Ecological Psychology* **14**(1–2), 87–109.

Varela F. and Bourgine P. (2002): *Toward a Practice of Autonomous Systems*, MIT Press, Cambridge MA.

Varela F., Thompson E., and Rosch E. (1991): *The Embodied Mind: Cognitive Science and Human Experience*, MIT Press, Cambridge.

Wiener N. (1948): *Cybernetics: Control and Communication in the Animal and the Machine*, Wiley, New York.

Received: 13 November 2005
Revised: 20 March 2006
Accepted: 20 April 2006

Reviewed by Andrew Bailey, Cees van Leeuwen, and Michael Spivey