Biosemiotics (2013) 6:379–391
DOI 10.1007/s12304-013-9174-8

ORIGINAL PAPER

Autopoiesis, Life, Mind and Cognition: Bases for a Proper Naturalistic Continuity

Mario Villalobos

Received: 5 October 2012 / Accepted: 17 December 2012 / Published online: 2 March 2013

Abstract The strong version of the life-mind continuity thesis claims that mind can be understood as an enriched version of *the same* functional and organizational properties of life. Contrary to this view, in this paper I argue that mental phenomena offer distinctive properties, such as intentionality or representational content, that have no counterpart in the phenomenon of life, and that must be explained by appealing to a different level of functional and organizational principles. As a strategy, and following Maturana's autopoietic theory of cognition, I introduce a conceptual distinction between mind and cognition. I argue that cognition corresponds to the natural behaviour that every living being exhibits in the realization of its existence, and that, viewed in that way, cognition is a dynamic process of structural coupling that, unlike mental phenomena, involves no representational contents. On the basis of this distinction, I try to show that while life suffices for cognition, it does not suffice for mind. That is, that the strong continuity is not between life and mind but between life and cognition.

Keywords Life-mind continuity thesis · Autopoiesis · Mind · Cognition

Introduction

The present work aims to show that though life may be considered a precondition for mind, mind cannot be understood simply as an enriched version of life. As opposed to the strong version of the life-mind continuity thesis (Godfrey-Smith 1996; Wheeler 1997, 2011; Swan and Goldberg 2010), which sees mind as an enriched version of *the same* functional and organizational properties of life, I claim that mental phenomena offer distinctive properties, such as intentionality or representational content, that (i) have no analog or counterpart in the phenomenon of life, and (ii) must be explained

Special Issue "Origins of Mind" edited by Liz Stillwaggon Swan and Andrew M. Winters

M. Villalobos (✉)
School of Philosophy, Psychology and Language Sciences, University of Edinburgh,
Edinburgh EH8 9AD, UK
e-mail: M.E.Villalobos@sms.ed.ac.uk

by appealing to a different level of functional and organizational principles. As I will try to show, this is so because the organization that defines living beings is autopoietic organization (Maturana 1975), which is realized physically in the metabolic network of living systems as a strictly mechanistic process without any trace of intentional or semantic aspects. In other words, I contend that life does not prefigure mind because it prefigures neither intentional nor representational phenomena.

Instead of searching for a direct (strong) continuity between life and mind, I will propose, though in a rather tentative way, an indirect or mediated (weak) continuity between life and mind. Embracing a sociolinguistic theory of mind, I will suggest that the key properties of mind are prefigured in language, i.e., that the mental domain emerges as an individual appropriation of the social linguistic domain. Following Maturana's biological theory, and introducing a more or less clear distinction between mind and cognition, I will argue that the strong continuity is not between life and mind but between life and cognition, and that mind emerges as a special kind of cognitive activity through the mediation of language.[1]

The structure of the paper is as follows. First I start by claiming that intentionality and representational content are distinctive marks of mental phenomena. Then I introduce the autopoietic theory of living systems as a preliminary framework to examine the nature of cognitive phenomena from a biological point of view. After that I present cognition as the basic know-how that every living system exhibits in the simple realization of its existence and try to show that, viewed in that way, cognition is a mechanistic process of structural coupling that neither involves representational contents nor exhibits intentionality. I do this in order to show that *while life suffices for cognition, it does not suffice for mind*, on the understanding that mind entails the presence of representational and intentional phenomena. Towards the end of the paper I propose some tentative ideas about the possible linguistic origins of the intentional and representational properties that characterize mental phenomena.

Intentionality and Representational Content: Essential Features of Mind

It is widely agreed that intentionality and representational content are distinctive marks of mental phenomena; i.e., properties without which we cannot identify a phenomenon as being mental. That a mental state is intentional means, in a philosophical sense, that it is *referred* to some content or that it is *about* something. For example, if I believe that Santa Claus does not exist, my belief is *about* the existence of Santa Claus. Every mental state seems to be *directed* (in a semantic or referential sense) to something, something that is the object of the intentional act (e.g., the object of my belief, the object of my desire, etc.). Another way of putting this is to say that mental states exhibit a determined representational content. Typically, mental events are viewed as phenomena that represent something, and, as a general conception, mind is viewed as something whose main function is to represent the world in a certain way.

[1] This thesis presupposes a conceptual distinction between mind and cognition that is not usually made, but that proves to be, I believe, necessary and coherent within the theoretical framework that I present here.

Intentionality and representation are intimately linked. After all, every representation is always *about* something that is being represented, and the referential relation entailed in intentionality is, certainly, a *semantic like* relation. Some philosophers recognize this fact by using two different senses of intentionality: *referential* intentionality and *content* intentionality (Kim 2011). The first one concerns the *aboutness* of our mental states and the second one emphasizes the fact that our mental states have *meanings* or *semantic contents*.[2]

Intentionality and representation are viewed as distinctive marks of the mental because no physical state or event seems to possess such properties *as intrinsic aspects*. Physicochemical processes are what they are without being themselves, in the technical philosophical sense aforementioned, *about* other things. If the ambient temperature goes down to minus 3 °C, water will pass from a liquid state to a solid state. We can say that the falling of the ambient temperature is the *cause* of the freezing of water, but it would be odd to rephrase this causal relation by saying that the falling of the ambient temperature is also the meaning or representational content of the freezing of water. The freezing of water, by itself, has no meaning. A completely different thing is that I, as a representing system, may observe the solid state of water and take it as an indicator, a sign, or evidence that allows me to infer other facts, for example, that the ambient temperature must have descended to at least zero degrees. I can treat the causal connection between these physical events as a semantic relation and say, "The water is frozen. That *means* that the ambient temperature must have fallen down to at least zero degrees." But who establishes the referential connection here is me, not the physical states, and I can do that precisely because I have a mind.

The Autopoietic Organization of Living Beings

Maturana's biological theory is an attempt to answer the following question: What class of systems are living beings? The theory starts by establishing a basic distinction between the organization and the structure of a system (Maturana 1975, 1981). The organization refers to the set of relations that constitutes the system as such and that confers its class identity, whereas the structure denotes the actual instantiation of such relations, including the concrete components involved in such instantiation, their states and qualities. Maturana contends that what defines the class identity of a system is its organization, not its structure. A change in the structure *may or may not* lead to a change in the class identity of the system, whilst a change in the organization *always* leads to a change in the class identity of the system. That is, while every organizational change entails a structural change in the system, not every structural change in the system entails an organizational change. For example, a golden chair remains a chair as long as it maintains a particular organization among its pieces. While this organization is conserved, the chair can admit several structural changes without losing its class identity as a chair (e.g., we can change the shape of its legs, or replace them by wooden legs, etc.). Now, if we set the chair on fire, we will

[2] There are also some philosophers, such as Fodor (2009), who stress this semantic (language-like) feature by using the word intensionality (with an 's').

observe a sequence of structural changes (i.e., the melting of the gold) that lead finally to the disorganization of the chair as a chair. In the same way, if we disentangle the chair and reorganize its pieces to make a table, then the chair disappears as a system and a new system appears in its place. The chair exists as a chair as long as it maintains a certain organization that is recognizable for us, and ceases to exist when said organization changes or is lost. Its structure, on the contrary, can change within certain ranges without altering its class identity as a chair.

If organization is what defines the class identity of a system, claims Maturana, then the question about the class identity of living systems is a question about their organization. Transferred to the level of the basic living unity, the question is: What is the organization of cellular systems?

According to Maturana's definition, cellular systems are autopoietic machines materially realized in the molecular domain. The key word in the definition is "autopoietic". This neologism captures, according to Maturana, the proper organization of living systems. But before addressing this notion, someone might question the word "machines". What does it mean to say that living systems are machines? Well, nothing too special, really. 'Machine' is a term that cyberneticians use to designate any system in general, natural or artificial, whose changes of state follow a deterministic pattern. That is, a pattern in which, at every moment, the current state of the system is the result of the transformation of the previous state of the system. Thus, whether an entity "X" is a machine or not has to do with the kind of transitions by means of which its behaviour is generated, and not with the concrete characteristics of said behaviour (more about this in section "Structural determinism, operational closure and structural drift"). A cybernetic machine can be a formal system of algebraic transformations, a galaxy, a car industry, a neuron, etc. A machine is simply a state-determined system, or, for simplicity, a deterministic system.

So cells are machines. But they are biological machines, living machines. In Maturana's words, they are *autopoietic* machines. What does that mean? Literally, it means that they are self-producing systems, i.e., systems that are organized as a network of productive processes that produce their own components. In the concept of autopoiesis the suffix 'poiesis' is used in its original Greek sense, meaning 'to make', 'to fabricate', or 'to build'. More specifically, the notion alludes to a process of 'synthesis' or 'composition' whereby a set of elements are assembled (combined under certain organization) to form a complex whole. Maturana wants to capture the permanent dynamic of *molecular synthesis* that takes place in the cell metabolism. The cell, basically, is viewed as a molecular factory that fabricates (synthesizes) the molecules that constitute it as such. What we must not forget, nonetheless, is that this molecular factory is, ultimately and despite its complexity, nothing more than a deterministic system of physicochemical transformations, i.e., a natural *machine*.

Cognition: The Praxis of Living

So far we have reviewed the autopoietic theory of Maturana with respect to the organization of living beings. But we have not yet said anything about the alleged relation, according to the theory, between life and cognition. Maturana contends that "livings systems are cognitive systems, and [that] living as a process is a process of cognition" (Maturana 1980). How should we understand this thesis?

First of all, we should start by recognizing that the concepts of knowledge and cognition are not uniform. We can distinguish different kinds of knowledge according to different criteria such as its source (reliable or unreliable, direct or indirect), its degree of certainty (probably or conclusively true), its accessibility (private or public), etc. For our purposes, nonetheless, there is an especially relevant distinction: the difference between declarative (discursive, theoretical) and performative (practical, behavioural) knowledge.

Probably the most familiar one for us humans is declarative knowledge, also called know-that. This is the kind of knowledge that is manifested through thoughts, propositions, judgments or sentences. For example, we know that the sun is a star, that Moscow is the capital of Russia, that two plus two equals four, etc. Basically, this is theoretical knowledge. Performative knowledge, on the contrary, has to do with actions, skills, with doing things instead of thinking or talking about them. For example, we know how to ride a bicycle, how to walk, how to prepare breakfast, etc. Basically, this is practical knowledge, also called know-how.

Peter, a student of physics, has general declarative knowledge about riding a horse—he knows that, to ride a horse, we need to keep certain postural balance and certain alignment in our body's centre of gravity. Yet Peter has never had contact with a real horse and does not even know how to mount one. Although he has some theoretical notions about riding a horse, he really does not know how to ride a horse.

Note that declarative knowledge seems to entail, in its turn, some kind of performative knowledge. After all, Peter needs to know how to compose propositions, how to connect concepts according to certain rules and so on in order to entertain certain thoughts and beliefs about riding a horse. Thinking or talking about riding a horse is not riding a horse, but thinking and talking are certainly actions in their own right. Every time we think or talk we are doing something, and that means that we know how to do it.

In that sense, it is interesting to note that performative knowledge seems to be much more fundamental and universal than declarative knowledge.[3] The latter requires the former, but not vice-versa. We learn to walk without having any theoretical notion about the complex biomechanic processes involved in the action of walking, and in the same way we learn to run, to climb, to speak and many other actions. The universality of performative knowledge is also easy to appreciate. Animals know how to do what they do (to swim, to hunt, to build a nest, etc.) without the mediation of any declarative or theoretical knowledge about their actions. When a spider is building a web, it is performing an action, it is doing something, and if the spider finishes its job in a successful way, there is a sense in which we can admit that the spider *knows how* to build a web. Yet it would be at least extravagant to say that the spider has a discursive knowledge about building a web.

Having briefly reviewed this basic distinction, one might say that every living being, in its continuous doing, reveals a certain know-how that is congruent with its particular form of existence. From this point of view, it seems natural to admit that every living being is a cognitive system insofar as it exhibits certain know-how in the

[3] For a different interpretation see Fodor (2008).

art of living. That is the sense of the autopoietic aphorism “to live is to know”, which means that cognition, in its most basic and embracing sense, corresponds to the *praxis of living*.[4]

The interesting point is that, viewed in that way, life appears to be a sufficient condition for cognition, and that appears to be the case in the individual behaviour of every organism. If we accept (i) that practical knowledge (know-how) is a form of cognition, and (ii) that every organism, in its natural behaviour, reveals a particular know-how in the art of living, it seems to me that we are entitled to conclude the following: if something is a living system, then it is also a cognitive system. This conclusion, bold as it may sound, is entirely valid (i.e., internally consistent with the premises of the argument). What someone legitimately might object to are rather the premises of the argument (i.e., show that the argument, though valid, is not sound). Someone could demonstrate that practical knowledge is not really a form of knowledge, or that, even if it is a genuine form of cognition, the natural behaviour of living beings does not fall within that category. All this also depends on how one defines cognition and the kind of relation that one establishes between mind and cognition. Usually, though not always, cognitive scientists and philosophers understand mind and cognition as relatively coextensive terms: where there is a mind there is cognition, and where there is cognition there is a mind. Here, on the contrary, we are working under the assumption that mind and cognition are different categories.

Let us recall that our hypothesis is that life suffices for cognition but not for mind. We have said that the presence of mind entails the presence of representational and intentional phenomena, and that simple cognition, understood as the basic know-how exhibited by living beings, does not involve representations or intentionality. Nonetheless, we still have not shown that cognition, so understood, does not involve such properties. It might well be the case that the practical cognition exhibited by living beings is in fact a representational phenomenon, or at least a phenomenon in which some intentional aspects are involved. So our next task is to address the phenomenon of cognition from a biological point of view, and try to show that in that phenomenon there is no room for either representations or intentionality.

Structural Determinism, Operational Closure and Structural Drift

In the previous section we said that cognition corresponds to the natural behaviour of living beings, and in section “The autopoietic organization of living beings” we defined living beings as autopoietic physical machines. Now is time to link these ideas in order to understand cognition as a natural phenomenon. If cognition is the natural behaviour of living beings, and if living beings are self-producing physical machines, then cognition is simply the behaviour of such self-producing physical

[4] Note that the aphorism is unidirectional and does not entail any ontological parity between the terms. It just says that to live is to know, but it does not entail the reciprocal “to know is to live”, nor does it present ‘cognition’ as having the same ontological (i.e., natural) status as life. Reading the aphorism as entailing identity (life = cognition), or as implying that ‘cognition’ is a natural category similar to the category of ‘living beings’, leads to a series of confusions or pseudo-problems that, though interesting, we cannot address in this paper. Authors who have, I think, fallen in this kind of pseudo-problem are, for example, Bourgine and Stewart (2004), Bitbol and Luisi (2004).

machines. In other words, cognition is nothing more than the behaviour (i.e., the trajectory of changes of state) of a particular class of deterministic physical machines.

To understand cognition as a natural phenomenon we have to recall that the peculiarity of living beings lies uniquely in their organization as autopoietic systems, and not in the structural logic that generates their changes of state (i.e., their behaviour). The know-how exhibited by living beings is an expression of both their organization and their structure. Their organization is unique (autopoietic), but their structural logic is not. *The way* in which a living being undergoes its structural changes is indistinguishable from *the way* in which a non-living system undergoes its structural changes. What is different—actually remarkably different—are the behaviours generated in each system. Yet this difference has to do with the way in which each system is structurally organized and not with their respective logics of structural change. For example, it is evident that a drifting boat and a dolphin behave in very different ways. What is not so evident is that this behavioural difference has to do with the way in which these systems are organized, and not with their respective logics of structural change. What separates the dolphin from the boat is its organization as an autopoietic system; the fact that he is in a constant process of self-production. It is this structural organization that generates a peculiar kind of behaviour in the dolphin; a kind of behaviour that we usually call "cognitive" or "intelligent". Nonetheless, both the dolphin and the boat are physical machines whose structural logic is strictly deterministic. Knowing this, usually we would not describe the boat's behaviour in semantic terms, appealing to intentional notions or alleged representational mechanisms. Why, apart from purely pragmatic reasons, should we do it with the dolphin? Maturana's position is that we should respect the mechanistic ontology of living beings and describe their behaviour in simple structural terms, without invoking semantic notions (Maturana 2002). That is the main message entailed in the definition of living beings as autopoietic machines, wherein the notion of machine is the key. Let us review the argumentation behind this thesis.

Maturana contends that living beings are, like any other physical system, *structurally determined systems* (Maturana 1987), and that their trajectory of changes of state, their behaviour, occurs as a simple *structural drift*. The principle of structural determinism, as I read it, says two things:

a) Every structural change that takes place in a system occurs because the structure of the system admits such change (otherwise it could not take place).
b) Every structural change undergone by the system as a result of its interaction with the environment is always specified (determined) by the structural state of the system itself, and not by the structural conditions of the environment.

The structural changes that occur in the system may be the result of its own internal dynamic or may be triggered by the action of some external factor. The principle of structural determinism states that in both cases, internal and external, the system always follows its own structural logic. This is important to bear in mind especially in the case of living beings, since they have a rich internal dynamic. But before addressing the particular case of living beings, let us see this idea through a general example.

A person presses a button on a laptop with his finger and as a consequence of this the laptop turns on. After a couple of minutes the same finger presses the same button on the same laptop and, now the laptop turns off. What has happened? We have the same elements interacting in the same way but ending in different results. Well, there is nothing mysterious about that. Although the finger is pressing the same button with the same force, the current structural state of the laptop is different in each case and, consequently, so too is the structural change that takes place. The mechanical interaction with the finger triggers the change of state that is possible in every moment according to the current structural state of the system, but *it does not specify* the nature of such change. The laptop in its turn only reacts in the way in which its actual structure allows. By pressing the button it may turn on or turn off, but not dance, cry, or cook a pizza. The point is not that it cannot react in these ways but that for doing such things it would need a different structure.

Every system responds as it responds, reacts as it reacts, and does what it does always following its own structural legality. Every system is a structurally autonomous and independent system. Autonomy, so viewed, is a trivial property of every physical system and not a distinctive mark of living beings. Every living being reacts as it reacts and does what it does according to what its current structural state specifies in every moment. Its internal dynamic, its operations, remain always within its own structural logic. This idea does not mean, as it might seem, that the external factors do not play any role in the chain of changes of state of the system. It just means that they cannot specify the logic of such structural changes. Actually, since the system is always in contact with the environment, and since this interaction is able to trigger certain changes of state in the system, the result is that none of these interactions are, after all, trivial for the structural trajectory of the system. In our example, that the finger, through mechanical interaction with a certain button, can trigger only the changes of states specified by the laptop, has to do with the structural logic of the laptop. That is true. But the fact that the laptop, staying previously off, is now on, has to do with the changes of state triggered by the action of some external agent (that may be the finger), and in fact we cannot explain such change without considering the triggering action of the external agent.

Living beings are in permanent interaction with the different structural states of the environment, and nothing of what occurs in such interactions is trivial for them (Maturana 1987). External events are constantly triggering certain changes of state in living systems, thus modulating their structural trajectories. We use the word "modulating" because the dynamic of structural changes in the environment can just impinge on, *but not specify*, the internal structural dynamic of living beings. This means, from the *operational* point of view, that the structural dynamic of living beings constitutes a closed network of operations, in the sense that all that happens in the network are transitions of state defined by the structure and organization of the system itself.

In order to avoid misunderstandings with this idea, we have to highlight the following: we are talking about the operational logic of the systems in general (alive or not), not about the concrete material conditions under which they conserve their organization. We know that living beings, as dissipative systems, incorporate matter and energy from the outside and release in turn matter and energy to the outside. From the material and energetic point of view, living beings are essentially open systems. Nonetheless, from the operational point of view they are closed or auto-defined systems. As analogy, and just as analogy, we can take a dictionary and follow its operational logic as a lexical network. No matter in what point of the network we start the navigation, it will always send us to another item within the same network, which in its turn will send us to another item within the same network, and so on in an infinite auto-referential loop. In a similar sense, claims Maturana, living systems constitute closed domains of structural transformations.

Structural Coupling and Structural Drift

Strictly speaking, the interaction between a living being and its environment is the interaction between two structurally determined systems whose dynamics, though independent and operationally closed (they follow their own structural legalities), remain coupled and coevolving. With Maturana (1975), we may say that a living being and its environment remain in *structural coupling*. This coupling process is a process of *structural* co-evolution, not a phenomenon of semantic correspondence. The environment triggers certain structural changes in the living being, at the same time that the living being triggers certain structural changes in the environment. We use the word "trigger" because, being structurally determined systems, they can only initiate, but never instruct or specify, one or more structural changes in each other. This distinction is subtle but important. It means that the internal dynamic of living beings, due to their structural determination, can be perturbed or affected but never instructed or "informed" by the environmental factors. Living beings, as operationally closed systems, are not "open" to informational inputs. Nonetheless, they co-evolve with their environment in a joint history of structural coupling. This coupling process is like a continuous structural dance in which nobody is the leader, nobody is controlling the trajectory of movements, and nobody is giving orders or instructions. It is a dance without a preset choreography, but in which one observes a fine structural coherence. And it is this structural coherence, when interpreted from the point of view of the organism, that one connotes with notions such as "cognition," "adaptation" or "intelligent behavior." The behavioral congruence that one observes between the organism and its environment (the fact that the organism seems to "know" what to do in every situation) has to do with this shared history of structural changes, not with alleged informational inputs.

Usually, it is said that living beings are "intelligent" systems because they are capable of controlling their behavior. Nonetheless, here we are saying that the "intelligent behavior" of living beings is the spontaneous result of an uninterrupted history of structural coupling rather than the result of alleged control mechanisms. We reject the notion of control because in a structurally determined system such a notion is, to a large extent, meaningless. The behavior of any natural machine, alive or not, is generated under the following condition: given the current structural configuration of the system and the current structural configuration of its environment, the next state of the system is fully determined. This amounts to saying that in the structural dynamic of the living systems there are no possibilities of action at all; there are no alternatives to select, no options for making a choice. 'Possibilities', strictly speaking, are a descriptive function introduced by the observer, not an existential condition of the observed system. Under a regime like this, it seems clear that the system cannot perform any control or regulation over the trajectory of its structural changes.

Maturana, recognizing this fact, offers the notion of "structural drift". Like a drifting boat in the sea, without a helmsman, without control or regulation, every system, alive or not, follows a deterministic trajectory of structural changes that, in spite of having no direction or purpose, proves congruent with the trajectory of structural changes in its environment. This congruency is not an exceptional situation but a trivial property of the interaction of every system with its environment. This is so because, as we already know, system and environment are nothing more than two systems co-evolving in a joint structural dance.

If we have assimilated the idea that living beings are autopoietic physical *machines*, we should also be able to accept the idea of *structural drift* without too much trouble, that is, without misinterpreting the analogy. The analogy does not say that a drifting boat and a living being behave in the same way. What it says is that their behaviors, though different, exhibit the same structural logic (recall the example of the boat and the dolphin). The natural behavior of living beings, according to this view, turns out to be just a particular version of structural drift.

Cognition, Representation and Intentionality

There are many theoretical approaches that see the intelligent behaviour of living beings as a phenomenon mediated by internal representations. I will point out in general terms the kind of assumptions on which these theories build their interpretations, showing that such assumptions are, according to all we have said here, wrong or incompatible with the operational-structural logic of living beings.

Most of the representational cognitive theories subscribe to one or more of the following assumptions: a) that biological systems in general, or certain subsystems like the nervous system in particular, pick up or collect information from the environment, b) that this information is used for building internal representations, c) that living beings are goal-oriented or teleological systems, and d) that biological systems in general, or certain subsystems like the nervous system in particular, are systems of control and regulation of behaviour.

The informational assumption is that living beings are, from the cognitive point of view, systems that consume information (*informavore* systems, in the words of Pylyshyn 1986). This information is in the environment and living systems pick it up through their sensorial channels. As with nutrients, the ingestion of information is vital for the organisms. Why? Because, following the reasoning of these theories, without the acquisition of some minimal informational content about the environment, the organisms cannot elaborate any internal representation about the world, and without a minimal map of the world they cannot know how to behave in that world. Insofar as informational content appears as the raw material for building internal representations, these representationalist theories assume what we could call the "informational prerequisite", which says: "there is no representation without information". Internal representations are *made of* information. Being a necessary condition for building representations, it follows that if a system is unable to acquire or "ingest" information, it will not be able to build any kind of internal representation. Now we have to ask ourselves: are living systems really 'informavore' systems?

What we have shown in the previous section is precisely that living systems, due to their operational closure and structural determinism, do not interact with the environment through informational exchanges. We have seen that, though from the thermodynamic point of view living systems are open systems, from the informational point of view they are closed systems. Their sensorial surfaces are not "windows" through which alleged informational contents are received but simple points of structural contact and coupling. Consequently, having no possibility of picking up information, living systems cannot build any internal representation of their environment.

What about action? Is not action mediated or facilitated by some kind of (at least minimal or weak) internal representation? Here the main assumption is that the "intelligent" behaviour of living beings is the result of a set of internal processes that control, regulate and guide the behaviour. The nervous system is viewed as a control system that commands the actions of the organism through a series of instructions and informational messages. These neural mechanisms, in their turn, are viewed as internal representations oriented to actions, as mechanisms that assist the generation of intelligent (flexible, adaptive) behaviour. Notions like action-oriented representations (Clark 1997; Wheeler 2005) or anticipatory models like the "emulator" (Grush 2004; Clark and Grush 1999) are good examples of this way of reasoning.

Nevertheless, we saw in the previous section that living beings are just a particular kind of natural machine, and that as deterministic systems they cannot perform any control or regulation over their behaviours. Their behaviours, though flexible and "intelligent" according to our appraisals, are just the result of a permanent process of structural drift wherein there are neither possibilities of action nor goals to follow. As genuine drifting systems, living beings neither foresee the consequences of their actions nor perform anticipatory mechanisms. Living systems in general, and nervous system in particular, are not control systems. If this is so, the idea that intelligent behaviour is controlled by means of internal representation proves to be not only misleading but incompatible with the ontology of living beings as autopoietic machines.

Are Actions Intentional Phenomena?

If the reader recalls the way in which we characterized the relationship between the living being and its environment, it should be apparent that the mechanism of structural coupling does not entail any intentional relation between them. To put it briefly, structural relations are not semantic relations. The organism and its medium work on the basis of a continuous process of mutual perturbation, not on the basis of referential relations. The structural states of the organism are not *about* the structural states of the environment (and vice-versa). That which we call cognition, i.e., the know-how of living beings, simply corresponds to the trajectory of structural changes that the organisms exhibit in congruency with their environment. In Maturana's words (1987), cognition "takes place in a living system as it operates in its domain of perturbations, and as such it has no content and is not "about" anything". If that is the case, it seems natural to conclude that the behaviour of living beings is not an intentional phenomenon.

Final Comments: Towards a Sociolinguistic Hypothesis of Mind

In this paper I have tried to show, through an extensive analysis of the structural constitution of living beings and their operational logic, that the cognitive behaviour of living beings neither exhibits intentional directedness nor is mediated by internal representations. At the same time, following (Maturana 1980), I have claimed that every living being is a cognitive system and that living as a process is a process of cognition. In claiming this, I have established a strong continuity between life and cognition, but not between life and mind.[5] Why? The reason is plain and the attentive reader, at this point, should be able to provide the answer. Representational content and intentionality are fundamental properties of mind, and life as a cognitive phenomenon simply does not exhibit such properties. Consequently, mind cannot be considered just an enriched version of *the same* properties of life. To be a minded system requires something more than being alive. What more is needed?

My hypothesis is that while the root of cognition is biological and individual, the origin of mental phenomena is essentially linguistic and social. In a nutshell, the idea is that to be a minded system the organism needs to develop through its social relations a special kind of communicative behaviour, rich and complex enough to generate the semantic and representational phenomena that later will become the core features of its mental life. In other words, I think that to be a minded system the organism needs to develop language. My suggestion, consequently, is that to understand mindedness we should explore and develop a sociolinguistic hypothesis of mind, which means expanding the focus of our analysis and turning our attention to what happens when two or more living systems start to interact in a recurrent manner constituting new domains of structural coupling. That is, we should study the communicative dynamics developed by living beings and the emergence of language as a special kind of communicative pattern within that domain.

A sociolinguistic hypothesis of mind is naturalistic in character. It does not appeal to any mysterious or supernatural order of phenomena (unless someone thinks that communicative behaviours and language are supernatural phenomena). Both language and social behaviours are natural phenomena easily observable in a wide range of ecological conditions. To point to them as the possible roots of mental phenomena is, accordingly, to offer a naturalistic hypothesis about the origins of mind. As I see it, the virtue of such a sociolinguistic approach is that it offers a naturalistic (though not individualistic) hypothesis about the origins of mind without forcing us to extrapolate our mental properties into the non-linguistic biological levels. That is, it offers a hypothesis that shows mental properties as belonging to the natural world, but that does not assume that such properties must be instantiated (or proto-instantiated) in all forms of life.

To be clear; linguistic phenomena constitute a subclass of communicative phenomena, and communicative phenomena, in their turn, a subclass of social phenomena. But since social phenomena presuppose the existence of individual biological unities, ultimately, neither social nor linguistic phenomena exist outside the biological

[5] As it is clear now, Maturana's autopoietic theory only establishes a strong continuity between life and cognition. Yet it is a common mistake to interpret this thesis as implying a strong continuity between life and mind. As an example of this misinterpretation see Godfrey-Smith (1996), chapter 3.

domain, broadly understood. In this general sense, linguistic phenomena are continuous with the rest of nature. If mental phenomena are conceived as emerging from linguistic phenomena, then they are conceived as continuous with the rest of nature too, but this in the general (or weak) sense that their specific properties (representational content and intentionality) emerge from other phenomena that are in their turn continuous with the rest of nature, not in the narrow (strong) sense that such properties are ubiquitous in the biological world.

References

Bitbol, M., & Luisi, P. L. (2004). Autopoiesis with or without cognition: defining life at its edge. *Journal of the Royal Society Interface, 1*, 99–107.

Bourgine, P., & Stewart, J. (2004). Autopoiesis and cognition. *Artificial Life, 10*(3), 327–345.

Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, Mass.: MIT Press.

Clark, A., & Grush, R. (1999). Towards a cognitive robotics. *Adaptive Behavior, 7*, 5–16.

Fodor, J. (2008). *LOT 2: The language of thought revisited*. Oxford: Oxford University Press.

Fodor, J. (2009). Where is my mind? *London Review of Books, 31*(3), 13–15.

Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. NY: Cambridge University Press.

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

Kim, J. (2011). *Philosophy of mind*. CO: Westview Press.

Maturana, H. (1975). The organization of the living: a theory of the living organization. *International Journal of Man–machine studies, 7*, 313–332.

Maturana, H. (1980). Biology of cognition. In H. Maturana & F. Varela (Eds.), *Autopoiesis and cognition: The realization of the living* (pp. 5–56). Dordrecht, Holland: Kluwer Academic Publishers.

Maturana, H. (1981). Autopoiesis. In M. Zeleny (Ed.), *Autopoiesis: A theory of living organization* (pp. 21–33). New York; Oxford: North Holland.

Maturana, H. R. (1987). Ontology of observing: The biological foundations of self consciousness and the physical domain of existence. In E. Caianiello (Ed.), *Physics of cognitive processes* (pp. 324–379). Singapore: World Scientific.

Maturana, H. (2002). Autopoiesis, structural coupling and cognition: a history of these and other notions in the biology of cognition. *Cybernetics & Human Knowing, 9*(3–4), 5–34.

Pylyshyn, Z. (1986). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, Mass.: MIT Press.

Swan, L. S., & Goldberg, L. J. (2010). Biosymbols: symbols in life and mind. *Biosemiotics, 3*(1), 17–31.

Wheeler, M. (1997). Cognition's coming home: the reunion of life and mind. In I. Harvey & P. Husbands (Eds.), *Proceedings of European conference on artificial life IV* (pp. 10–19). London: MIT.

Wheeler, M. (2005). *Reconstructing the cognitive world*. Cambridge, Mass.: MIT Press.

Wheeler, M. (2011). Mind in life or life in mind? Making sense of deep continuity. *Journal of Consciousness Studies, 18*(5–6), 148–168.