BRAIN AND COGNITION **34,** 107–129 (1997)
ARTICLE NO. BR970909

Cognition and Life: The Autonomy of Cognition

Alvaro Moreno,* Jon Umerez,* and Jesús Ibañez†

**Department of Logic and Philosophy of Science and †Department of Languages and Information Systems, University of the Basque Country, Donostia, Spain*

In this paper we propose a philosophical distinction between biological and cognitive domains based on two conditions that are postulated to obtain a useful characterization of cognition: biological grounding and explanatory sufficiency. According to this, we argue that the origin of cognition in natural systems (cognition as we know it) is the result of the appearance of an autonomous system embedded into another more generic one: the whole organism. This basic idea is complemented by another one: the formation and development of this system, in the course of evolution, cannot be understood but as the outcome of a continuous process of interaction between organisms and environment, between different organisms, and, specially, between the very cognitive organisms. Finally, we address the problem of the generalization of a theory of cognition (cognition as it could be) and conclude that this work would imply a grounding work on the problem of the origins developed in the frame of a confluence between both Artificial Life and an embodied Artificial Intelligence. © 1997 Academic Press

1. INTRODUCTION

In the second half of the present century modern science has witnessed an apparently contradictory process. On the one hand, the classical ''field disciplines'' have become fragmented into a variety of significantly specialized subareas with increasingly narrower scopes. On the other hand, heterogeneous scientific communities have developed around multidisciplinary ideas, integrating the different epistemological, methodological, and technological contributions of their members and creating what have been called the sciences of complexity (see Pines, 1987; Cowan, Pines, & Meltzer,

The authors are very grateful to Arantza Etxeberria for her suggestions and ideas to original drafts and to Andy Clark, Pablo Navarro, and Tim Smithers for their comments and discussions that helped to clarify some passages of the final draft. This research was supported by Research Project PB92-0456 from the DYGCIT-MEC (Ministerio de Educación y Ciencia, Spain) and by Research Project 230-HA 203/95 from the University of the Basque Country. Jon Umerez acknowledges a Postdoctoral Fellowship from the Basque Government. Address reprint requests to Alvaro Moreno, Departamento de Lógica y Filosofía de la Ciencia, Universidad del Pais Vasco UPV/EHU, Apartado 1249, Donostia/San Sebastian 20.080, Spain.

0278-2626/97 $25.00

1994). The first significant milestones of this phenomenon were founded almost in parallel by cybernetics (Wiener, 1948/1961; Ashby, 1956) and by general systems theory (Bertalanffy, 1968, 1975), and as a consequence we can today speak about adaptation, autonomy, communication, information, second order (observation-dependent), and, specially, complexity sciences, although in some cases the title is far from having a sound tradition. Undoubtedly, the quick and huge development of computer science has been and is still an important factor in the spreading of these new disciplines, since it has provided empirical and experimental counterparts to the formal approaches necessary in the attempts to abstraction of these ‘‘complexity sciences.’’

Let us just mention some of the main conceptual issues which might help in drawing the wider context within which the subject of the paper is embedded.

1.1. *Artificial/Natural*

The extreme position in this direction has been the development of a new scientific strategy that (transcending the previous and well-established practice of engineering and other technological disciplines) has given origin to the sciences of the artificial (to borrow Simon’s, 1969, denomination). This strategy consists of the study of complex systems by means of their artificial creation, in order to experimentally evaluate the main theories about them. The computational technology is central (although not necessarily unique) to this new experimental paradigm, and its most outstanding applications concern precisely those fields where it becomes not just an alternative to other empirical approaches but a first-order validation tool (since the enormous complexity of their scientific objects makes it extremely difficult to render the traditional experimental approaches operational). Until now, there have been two main research projects that have resulted from its application to psychological and biological problems: Artificial Intelligence and Artificial Life, that is, attempts at ‘‘explanation through recreation’’ for the natural phenomena of intelligence and life.

1.2. *Functionalism*

The application of this paradigm to an existing scientific field (the host science) has two consequences that will be relevant to the discussion in the present work. First, the change in status of the epistemic relationship between scientific theories and reality: the pervasive comparison of models and artificially created systems renders what has been called the deconstruction of science’s traditional object (Emmeche, 1994). This situation is clearly exemplified in Artificial Intelligence, where it is very common to find arguments against its comparison with natural cognitive processes whenever it fails to account for them adequately. This defense strategy of a research program’s framework can also be detected in Artificial Life researchers, and its natural fate is to evolve into extreme functionalism by explicitly giving up in the attempt of modeling real systems (Umerez, 1995).

1.3. Universality

Another issue concerns the universality of the host science. Sciences such as physics, chemistry, and thermodynamics either do not assume ontological restrictions about their field of study or include a framework to reduce any constraint in their scope to the physical level. Moreover, they provide a methodology that operationally includes their target objects and its direct manipulation. Thus, they can be considered universal sciences of matter and energy since their laws are intended to be valid up to contingencies. Unlike them, biology studies phenomena about which we only have intuitive and empirically restricted knowledge (if any) and whose complexity is a barrier to its reduction to the lower physical and/or chemical levels (Moreno, Umerez, & Fernández, 1994). Thus, biology (not to speak of psychology) can be intended only as a science that studies the phenomenon of life through the experience of the living systems as we know them, and so we still have no means to distinguish, among the known characteristics of these systems, which are contingent and should not be required for the characterization of life or cognition in a generic context (Moreno, Etxeberria, & Umerez, 1994).

In this case Artificial Life has been more radical and more mature than its predecessor, AI. While Artificial Intelligence states as an explicit goal to attain a general theory to explain intelligent behavior but implicitly assumes the anthropocentric perspective of considering that what we have is the most we can expect, Artificial Life (Langton, 1989) has had from its birth a clear and explicit attitude toward contributing to extend the actual biological realm (life-as-we-know-it) into a universal theory of the living organization that could go beyond our empirical experience (life-as-it-could-be) as a consequence of the generalization provided by the artificial models.

1.4. Relation between Artificial Intelligence and Artificial Life

In any case, we can see that, despite the efforts made by most Artificial Life founders to make it an epistemologically different discipline from Artificial Intelligence (stress on the bottom-up methodology), both research programs have relevant similar traits and have had to some extent similar attitudes in facing the study of their respective main scientific targets (Pattee, 1989; Sober, 1992; Keeley, 1993; Umerez, 1995). In this sense, their methodological differences are contingent: each one has chosen its particular working hypotheses and, up to now, these have been proven the best ones in their respective fields for universality purposes. We obviously do not want to say that they are perfect, not even that they are any good. We simply want to point out that they have not at present better alternatives for producing general theories of life and intelligence.

The true differences arise when these researchers deal with a subject that lies partially in the scope of both, and that is precisely what happens with cognition. On the one hand, Artificial Intelligence has widened its scope, especially through its insertion into cognitive science, and has started to study processes that do not necessarily imply the classical knowledge approach (e.g., perception instead of recognition) (Brooks, 1991). On the other hand, in addition to dealing with specifically biological problems, Artificial Life has evidenced capability of producing systems that are close to model low-level cognitive processes. Cognition is not the main target of any of them (and for the moment we do not have a research program aimed at artificial cognition), but a secondary objective in both.

The approach to the cognitive phenomenon is different from the perspectives of Artificial Intelligence and Artificial Life. The former, though claiming a physicalist and mechanicist stand, has tended to consider cognition in an abstract and disembodied way. The latter, notwithstanding, has brought a new sensibility: an effort has been made to address the understanding of the cognitive phenomenon from the bottom up and to insert it in an evolutionary and embodied frame. From our point of view this constitutes a considerable advance but it has been reached at the price of creating confusion between what is generically biological and what is properly cognitive.

In this context it is legitimate to make the comparison between both approaches, their respective methodologies, theoretical models, and experimental results. Moreover, this is probably the most interesting comparison test (if not the only one) that we can have between Artificial Life and Artificial Intelligence. This paper attempts a critical review of the subject.

2. THE PHENOMENON OF COGNITION

Cognition is not a concept standing out there, ready to be handled by any discipline. There is considerable controversy about its definition and nature (Churchland & Sejnowski, 1993; Gibson, 1979; Pylyshyn, 1984; Smolensky, 1988; Van Gelder, 1992/1995; Varela, Thompson, & Rosch, 1991), to the point of altering its meaning depending on the starting epistemological assumptions for its study. In this sense it is worth remembering its philosophical origin, which denotes an essentially human feature in whose realization awareness, acquaintance, and knowledge should be involved. Today's situation is that most explicit definitions of cognition can be perfectly correct, though contradictory among them, simply because they address the same word to different concepts.

However, there are two aspects to this controversy. On the one hand, there is the problem of the boundaries of the idea of cognition. Given that there is no scientific disagreement about considering human intelligence a form of cognition, the upper boundary seems to be delimited without controversy. Therefore, the main problem to deal with concerns the lower limits of cognitive phenomena.

On the other hand, there is the methodology of the question, i.e., what kinds of definitions are we looking for? We suggest that in spite of arguing in favor of one or another type of definition, it is more useful to discuss the methodological implications that such definitions convey. In other words, a precise definition of cognition should not be sought but a useful one, that is, one that allows the correct framing of a research project centered on it, should be sought.

What is being discussed is mainly a conflict between two types of research programs about cognition. On the one side is a research program that attempts to make cognitive phenomena emerge from a purely biological background. On the other is the more traditional research program in Artificial Intelligence, which seeks mainly to reproduce high-level cognitive functions as a result of symbol manipulating programs in abstract contexts (Newell, 1980). This second position has the advantage of dealing with a distinctly cognitive phenomenology by focusing on high-level cognition without any condition. As a matter of fact, most of the best models in AI are of this kind. Nevertheless, this perspective has well-known but very serious problems: it implies an abstract and disembodied concept of cognition whose foundations (‘‘symbol grounding problem,’’ Harnad, 1990) are by no means clear.

Thus, according to these former considerations, we propose that a useful characterization of cognition should fulfill two conditions:

(a) Biological grounding—to establish the biological conditions under which cognition is possible and so to relate the understanding of cognition with its origins.

(b) Explanatory sufficiency—any plausible minimal conditions to characterize cognitive phenomena should include all the primitives necessary for utterly explaining its more evolved forms, the higher-level forms involved in human intelligence.

Therefore, in the next two sections we develop a concept of cognition simple enough to be derived from a biological frame and, at the same time, endowed with an autonomous[1] identity that permits it to be useful also for supporting high-level cognitive phenomena.

3. THE LOWER BOUND: LIFE AND COGNITION

As we have stated before, an important prerequisite for any research program involving the theoretical addressing of a complex phenomenon such

---

[1] In this paper we use the concept of autonomy in two senses, which will be discerned in each case, i.e., as a general idea of self-sustaining identity and as the more concrete result of some kind of operational closure. See Smithers (this issue) for a fuller and more encompassing treatment of the plural and different uses of the concept and its related terms.

as cognition is to work out an explanation of that phenomenon along with a characterization of the mechanisms that make it possible and originate it. Cognition only appears in nature with the development of living systems. The inherent complexity of living processes renders the relationship between life and cognition a very interesting and difficult problem. In particular, it has been traditionally very hard not only to identify precisely the origins of cognitive activities, but even to distinguish which biological activities can be considered cognitive (Maturana & Varela, 1980; Heschl, 1990; Stewart, 1992). We will try to trace the different stages associated with the origin of cognition addressed from the perspective of the origin of life itself.

Since its origin, life has provoked a set of continuous changes in the earth. Thus, living beings have had to develop several adaptive mechanisms to maintain their basic biological organization. On a phylogenetic scale, the solution to this fundamental problem is given by evolutionary mechanisms, but we also see that when we focus on organisms on their lifetime scale, each is able to adapt—in a nonhereditary fashion in this case—to changes in the environment. Even the simplest biological entities known at present possess some sort of ''sensor organs'' that evaluate the physical or chemical parameters of their environment that are functionally relevant for them to trigger structural or behavioral changes to ensure a suitable performance of their living functions.

At this level, ontogenetic adaptability consists of functional modulation of metabolism triggered by molecular detection mechanisms located in the membrane. Any biological system, no matter how primitive, includes relationships among different biochemical cycles that allow the existence of regulatory mechanisms that can imply modifications in different parts of the metabolic network. In this very elementary stage of the relations between organism and environment, the basic sensorimotor loops that constitute adaptive mechanisms do not differ significantly from the rest of the ordinary biological processes of the organism, e.g., its metabolic cycles. For instance, flagellum movements involved in oriented locomotion in certain types of bacteria can be equivalently characterized as modifications in metabolic paths. From this starting scheme, evolution has developed organisms provided with more and more complex metabolic plasticity, whose control by the organisms themselves has allowed in turn more and more complex viable behavior patterns. However, as far as the variety of possible answers was based only on the metabolic versatility of the organism, the complexity of the behavioral repertoire would remain strongly limited. That is why, according to the second condition assumed in the previous section, those kinds of adaptive responses to the detection of significant environment variations through molecular mechanisms (certain membrane proteins) are essentially nothing but biological functions and only in a very unspecific sense could such behavior be considered ''cognitive.''

The history of life shows, though, that the aforementioned limited responses have not been an insurmountable obstacle for the creation of progressively more complex organisms, when in the course of evolution some such organisms began to form pluricellular individuals. There is, however, an exception for those forms of life based on the search for food through movement, where speed in the sensorimotor loop is still crucial. In this case a process of cellular differentiation was induced, leading to an internal specialized subsystem that could quickly link effector and sensory surfaces. This process was the origin of the nervous system. In turn, the operation of such a system implied the development of an internal world of externally related patterns (coupled with sensors and effectors) organized in a circular self-modifying network. As we will see in the next section, in the organisms endowed with a nervous system (henceforth animals), instead of through metabolic mechanisms of self-control, adaptation takes place through an informational meta-control on metabolic–motor functions.[2]

When we study the evolution of pluricellular organisms whose strategy of life was based on the search for food, the development of such neural subsystem represents two significant advantages: higher speed and finer flexibility in the coordination of the sensory motor loops. Moreover, the organic changes attributable to nervous processes represent only a small number (in terms of energetic costs) of the physiological processes that occur in the lifetime of the individual. For these reasons selective pressure determined that the development of pluricellular organisms whose adaptability relied on motor behaviors would become impossible without a neural subsystem.

As the neural system becomes more complex, animals use neural resources ‘‘off line’’ for exploring virtual, mentally simulated situations before taking actions in their real environments (Clark & Grush, 1996). Hence, the fundamental factor in the development of the nervous system has been not only the relation between the organisms whose way of life is based on movement and their noncognitive environment but also the co-evolution—cooperation and competition—with other cognitive organisms. Co-evolution is essential (not just a contingent fact) for the emergence of meaning and cognition because the ‘‘autonomy’’ of the cognitive agents, and of every organism as biological organization, cannot be understood without its collective dimen-

---

[2] As to the suggestion that the immune system could be considered cognitive, we would say that, rather than being cognitive, the immune system is a system where processes similar to those of biological evolution take place, but inside an individual and in the lapse of a few hours or days (instead of populations and millions of years). Its functionality and speed notwithstanding, it is more a case of a very complex adaptive system in somatic time than cognitive. Functionally it is not in direct relationship with sensorimotor coordination (not functionally linked to directed movement). Furthermore, the immune system has only been developed within the frame of some cognitive organisms (vertebrate animals) and does not exist in noncognitive evolved organisms. It is therefore possible to pose the question whether it is not precisely following the development of complex forms of identity like the one occurring through the entanglement between the metabolic and nervous operational closures that the appearance of the immune system was propitiated.

sion (and vice versa). ‘‘Movement,’’ for instance, as has been pointed out by the Ecological Realism (see Turvey & Carello, 1981), is not to be taken as a mere physical concept, but mainly as a network of interactions among other organisms equally endowed with a nervous system. Accordingly, the development of cognitive capacities occurred as a collective phenomenon which took the form of a ‘‘bootstrapping’’ process.

3.1. Blending Approach and Its Evaluation

So far we have discussed the origin of cognitive capacities within an evolutionary frame. However, among those authors who agree with the idea of the necessity of a biological grounding of cognition, there is a (rather philosophical) discussion concerning the nature of life and cognition. As we see it, the clue to this discussion is the controversy over the significance of the gap between what we have called mere adaptation and the world of phenomena arising from the development of the nervous system.

Some authors (Maturana & Varela, 1980; Stewart, 1992; Heschl, 1990) consider that life itself necessarily involves cognitive abilities. Although significant differences exist among them, the positions of these and other authors share the assumption that life and cognition are, if not the same concept, inseparably linked phenomena, and in the ongoing discussion we refer collectively to them as the blended approach (BA).

According to the BA, all these adaptive processes constitute the simplest forms of cognition. Thus, there would not be any explanatory gap between life and cognition (the existence of biological systems is linked to the presence of cognitive abilities), and, moreover, the understanding of the nature of cognition is linked to an explanation of its own origin and of the origin of life (the simplest forms of cognition, and thus the easiest to understand, would be present in the earliest organisms).

The main advantage of this position is that it is able to give account of the biological origins of the epistemic phenomena. However, as we have pointed out, the concept of cognition proposed by the BA becomes reduced considerably in its operational meaning. This is because it renders as ontologically equivalent the aforementioned basic sensorimotor loops, that is, interaction processes between organism and environment through metabolic nets, and much more complex and evolved interaction processes that explicitly involve other cognitive organisms. In other words, the kind of processes considered in the BA paradigm are closer to other biological functions than to higher cognitive abilities.

In addition, there would be no profit in carrying the BA to its most extreme consequences: either cognition is reduced to life, and this leads to abandoning the term cognitive because of its lack of specific meaning, or a more pragmatic argument is adopted in order to state that life and cognition represent operationally and epistemologically different concepts. In the first case those processes (like the basic sensorimotor loops) that are presented as cognitive can, in fact, be characterized as purely adaptive process, in which the specifically cognitive dimension is not functionally distinguishable from the whole of the biological operations of the individual. In the second case, the problem is how to determine which biological processes could be categorized as specifically cognitive and which not. Thus, we would not have simplified the concept of cognition, but merely translated the boundary problems to the biochemical level, since it is at that level where earlier cognitive mechanisms are identified. Finally, it seems very hard to ground the primitives of cognitive science (like open-referential information and representation) without assuming the necessity of the aforementioned gap between purely biological phenomena and cognitive ones.

4. THE AUTONOMY OF THE NERVOUS SYSTEM

The existence of an important gap between purely adaptive behavior and high-level cognition suggests the importance of an explanation of the origin of cognition as an autonomous phenomenon with respect to biology and the necessity of raising the lower boundary of cognition. If we claim (as in fact we do) that cognition is not epistemologically equivalent to the basic biological functions, we need to identify not only its specific phenomenology but also the (harder) biological conditions to produce this phenomenology.

This leads us to address the question of the origin of the nervous system (NS) in a new manner, namely, as a radically different kind of organization arising in a biological frame. As we have mentioned before, the emergence of the NS was the result of an evolutionary process carried out by pluricellular organisms whose survival depended on obtaining food through movement. This strategy ultimately resulted in the formation of a specialized subsystem of the organism for quickly channeling the sensorimotor couplings. The organization of the nervous system is oriented toward density, speed, precision, plasticity, pattern number maximization, and energy cost minimization. The combination of these features expresses the specific identity of the NS as the material support of the cognitive capacities in animals.

Functionally speaking, the specificity of the role played by the NS lies in the different way by which adaptive mechanisms take place. Organisms without NS, when faced with biologically significant changes in the environment, trigger a set of functional metabolic reactions, preserving the biological viability of the organism. Here adaptation occurs essentially as a result of biochemical changes induced by sensor surfaces that constrain the metabolism. Instead, when animals interact cognitively with their environment, sensorial flow does not directly constrain metabolic states (the body), but a flow of nervous patterns within a recursive network. Effector organs are thus connected with sensors through this network, which allows for some internal patterns to be coupled with not only the present features of the environment. For this reason, it seems more convenient to speak about this kind of coupling between nervous internal patterns and external events in informational terms, whose meaning we discuss later.

As we will see, the specificity and potential open-endedness of the internal patterns arising in this network will open the specific phenomenology of the cognitive domain. Thus, the NS is the material support of the cognitive phenomenology as an autonomous level with regard to the rest of the biological domain. Cognition appears as the consequence of the emergence of the nervous system.

When we describe the NS as a functional network it is worth distinguishing it at different levels. At the bottom, we have complex patterns of metabolic processes. However, some of these processes at a higher level produce simple discrete events, and at even higher levels, patterns formed by groups of neurons. As a result of the functional interactions that an animal has with its environment, a history of couplings between internal states (underlying complex metabolic processes) and events of the environment arises. Thus, meaning or cognitive information occurs at different hierarchical levels, implying both activation patterns of discrete units and insertion of these patterns into a body frame endowed with an evolutionary history (Umerez & Moreno, 1995).

Here is where cognitive information appears. The fact that the sensorimotor loop is mediated through informational processes is precisely what distinguishes cognition from generic adaptation. However, information is also a concept central in biology at large; for instance, essential processes such as self-maintenance and self-reproduction depend on the specific sequence of discrete units stored in DNA molecules, i.e., genetic information. Now, this information, though generically ‘‘epistemic’’—because of its referentiality—is bounded to self-referentiality. Nevertheless, if information has to be a useful concept for cognition, it needs to convey open referentiality.

More precisely, let us compare these two senses of the term information. When we try to account for both its genetic and cognitive meanings, information should be understood as a set of patterns with causal effects that connect meta-stable states with physically independent events or processes by virtue of some interpretation mechanisms autonomously constructed by the very system. Therefore, in the case of both the genetic and the neuronal information, we are dealing with self-interpreted functional information and not with just a series of discrete units which have a certain probability assigned and whose meaning is externally attributed with independence of its formal structure. In the frame of the NS the term information corresponds to the functional description of those metabolic global patterns that in turn modulate a flow of chemical and physical processes connected to the outside through diverse organs, sensors, and effectors, in a circular manner.[3] The dynamics

---

[3] This leads us to interpret the information in the NS as metabolic global patterns that in turn modulate a flow of chemical and physical processes in a circular manner. The NS is connected to the outside through diverse organs, sensors, and effectors (two levels of exteriority: outside the nervous system and outside the whole organism, the latter being the most

of the informational flow is constrained by both the requirements of sustaining the entire organism's viability and the constraints of the structure of the environment.

In a way similar to the generic biological organization, the nervous system produces primarily its own internal states as an expression and a condition of its self-coherence as an operationally closed network (Varela, this issue). However, this autonomy is in its turn not independent of that of the whole organism. Once emerged and developed, the nervous system subsumes purely metabolic adaptability functions. In this sense, the epistemological autonomy of the nervous system is the continuous production and reproduction of an informational flow coherent with the viability of the autonomy of those organisms that generate precisely these internal meta-autonomies. Along with this, the nervous system is essentially immersed in (and coupled with) the external environment (mainly other cognitive organisms). The autonomy of the nervous system can also be stated in that frame of informational relationships.

Thus, the appearance of a new phenomenological domain whose primitives are these informational patterns is one of the most outstanding features of the nervous system. This domain relies on a set of features that configure the deep specificity of this system with respect to the rest of the organized structures of the individual organism.

As a consequence, the external environment of the organisms endowed with a NS is composed of informational interactions rather than functional ones. However, as this environment mainly consists of other cognitive organisms, the world of cognitive agents becomes progressively a communication world.

5. THE RELATIONSHIP BETWEEN THE COGNITIVE AND BODY FEATURES

We have previously pointed out that the nervous system constitutes in its performance an operationally closed system, and this fact supposes a fundamental problem: How can we understand the relationships between the nervous system and the rest of the organism (what we usually call body) if the whole of it is to be also characterized as an autonomous system? If the self-maintenance of the body is expressed by means of metabolism, how can we interpret the set of constraints performed by the nervous system on it?

This is a difficult question. Those who stress the embeddedness of the cognitive system normally blur its autonomy and ontological distinction with respect to the biological level, which hinders their ability to generate a useful research program in cognitive sciences. However, those who stress the auton-

---

important). Accordingly, we cannot interpret correlated functional states with external processes as informational when these states are merely metabolic ones (e.g., bacteria, paramecium, plant). However, in the case of adaptive metabolic changes that take place in animals, surely both levels (metabolic and informational patterns) are strongly interconditioned.

omy of cognitive phenomenon from the biological level tend to disembody it in greater or lesser degree.

If we want to avoid the problems involved in the disembodied theories about cognition, it is necessary to assume that the NS is subsumed in the wholeness of the body. As the latter is itself an operationally closed system, we would have to interpret ‘‘the whole organism’’ in its turn, as a higher form of an operationally closed system in which the body would perform the dynamical level and the nervous system the informational one, in a way similar to the concept of semantic closure proposed by H. Pattee (i.e., 1982, 1986, 1987, 1989, 1993, 1995; see also Thompson, this issue) to explain the complementary relation between DNA and proteins in the cellular frame. This interpretation seems to us more suitable than that dealing with the body as an ‘‘environment’’ for the nervous system (Clark, 1995).

How is this complementarity between body and nervous system reflected? The answer to this question could be in the sense that functional meanings emerge precisely through a self-interpretation process of the nervous information. The body (metabolic–energetic–functional system) specifies or determines the ‘‘readiness’’ of the informational relationships. What is functionally significant for the animal constrains the performance of the nervous system, and the converse. The body controls the nervous system, and the converse.

The autonomy of the body is, in a generically biological sense, more general and global than that of the nervous system. The body is energetically connected to the environment, while the nervous system is connected informationally. This does not mean that they are independent processes: in fact, what is informationally relevant for the organism depends on its internal state—thirst, sexual readiness, tiredness, etc. (Etxeberria, 1995). In addition, the phenomena of pain and pleasure are not understandable unless we conceive of the relation between the NS and the rest of the body in a globally entangled manner. The functional output of neuronal activity is not only a set of motor actions (which constrain sensorial inputs), but also a more encompassing effect of metabolic constraining action (hormonal secretion, etc.), which, ultimately, accounts for the whole sensorial flow (including the meaningfulness of ‘‘painful’’ or ‘‘pleasant’’ feelings in animals). So, the body biologically constrains the nervous system, for instance, determining its basic functioning, but also the converse is true: the (operationally closed) logic of the nervous system in turn determines how the biological functioning of the body will take place. The nervous system has phylogenetically bestowed on the body of the animal fundamental evolutionary constraints, thus conditioning the development of the different bodily structures.

Functional self-interpretation of the information (the question of the emergence of epistemic meanings) is possible only through this complementary relationship. The informational nature of the relation maintained by the nervous system with the environment expresses its autonomy and operational closure, whereas the entanglement between the biological and the cognitive structures expresses the embodiment of this latter.

What confers informational character to some of the patterns produced at certain levels of the nervous system is its operational closure. In the nervous system the neural patterns that assume an informational nature are those that establish causal connections with physically independent events due both to the operational closure of the nervous system and to that formed globally between the nervous system and the body. The first of these operational closures renders autonomous the process that connects the sensor surfaces with the effector mechanisms from the rest of the biological processes, thus constituting them as cognitive. The second is the mechanism by which the processes that occur at the nervous level acquire a functional meaning: in the end, the global biological self-maintaining logic is responsible for the functional interpretation of the information of the nervous system.

5.1. Representation

The present perspective is, in our opinion, the only one that allows a satisfactory approach to the problem of representation in cognitive science. This concept, in its classical formulations within the computationalism, has been heavily criticized in the last decade, especially for its alleged incompatibility with connectionism (see Andler, 1988). Recently, even its abandonment has been proposed (Varela, 1989; Van Gelder, 1992/1995; Brooks, 1991). Nevertheless, the problem of these radical positions is that they throw the baby out with the bath water, for without the idea of representation it is hardly possible to build a cognitive science able to explain cognitive phenomena of a higher level. Therefore, even if it is possible to discuss whether representation is dispensable in explaining a certain level of behavior, the crucial problem is that without a concept of representation it is not easy to see how a research program in cognitive science could be articulated that went from its lowest to its highest level without cleavage (as we put forth in the first section).

It is true indeed that a fair number of the debates around the dispensability of representation are due to disagreements with respect to what is cognitive, but there also are serious discrepancies in and confusion around the very meaning of representation itself. Clark and Toribio (1994) hold, however, that behind this diversity a basic implicit consensus exists around the definition proposed by Haugeland (1991). This definition states that representation is a set of informational internal states that hold a ‘‘standing for’’ (referentiality) relation toward certain traits of the environment which are not always present and which form a general scheme systematically generating a great variety of related states (also representational).

Most of the objections to and difficulties posed by this definition proceed from its abstract and disembodied character. However, if we situate the idea of these referential internal states in the context of the informational patterns generated by the operational closure of the nervous system, we think that these difficulties can be solved.

Thus, the mechanism of self-interpretation of the information inside the nervous system is achieved by the complementary relationship between the body and the NS within the organism as an integrated whole. This is the radical meaning of the statement suggesting that the biological is a lower ground of the cognitive level.

6. COULD WE HAVE A DISEMBODIED AI?

In the previous section about the relations that hold in natural cognitive systems that we know between the properly cognitive system (the nervous one) and that which globally provides for the identity of the whole system as an autonomous being (the body) we have seen the deep interrelation between both. Should we infer that any cognitive system must be based on a similar relationship? This question can be approached in two different ways.

The first consists of asking, within the frame of the universalization of biology, what conditions are necessary for the appearance of natural cognitive systems. It is possible to argue that the conditions we have indicated in previous sections are only the instantiation of a phenomenology, a particular case of the various possible histories of generation of cognitive systems in nature. Nevertheless, the main features we have used to define the cognitive phenomenon satisfy the requisites posed at the end of Section 2. Also, in any biological setting, it is logical to suppose that any kind of natural cognitive agent, as it might appear in any circumstances, could be conceived only as the result of some sort of evolutionary process from previous organizational and generically lifelike stages. According to this, the relationship framework concerning cognitive and biological levels would have guidelines similar to those of the previously described scheme.

The second way to address the question of how to generalize a full theory of cognition universally valid is the attempt to build artificial cognitive systems, what is commonly known as Artificial Intelligence. In this frame, one of the more important questions to investigate is the way to determine the series of conditions, structures, and material processes (what Smithers (1994) has called ‘‘infrastructure’’) required to support the emergence of cognitive abilities in artificial systems.

So far AI has tried mainly to simulate and build expert systems or neural networks that accomplish certain kinds of externally specified cognitive tasks. Despite the success attained in this research one could disagree with the idea that these systems are truly cognitive because, as we have previously argued, cognition is a capacity that should be understood through its own process of appearance and development. This implies their embeddedness in a whole biological background.

Recently, however, there has been increasing interest in relating cognitive and biological problems, due mainly to promising research that studies and designs robots capable of developing (different degrees of) autonomous adaptive behavior—the so-called behavior-based paradigm (Maes, 1990; Meyer & Wilson, 1991; Meyer, Roitblat, & Wilson, 1992; Cliff, Husbands, Meyer, & Wilson, 1994). The fact that autonomy should be considered the basic condition for cognition is precisely one of the bridges between Artificial Intelligence and Artificial Life.

7. ARTIFICIAL LIFE AS A METHODOLOGICAL SUPPORT OF A NEW ARTIFICIAL INTELLIGENCE

Artificial Life poses questions about cognitive phenomena from its own point of view, as something that must be considered inside a biological frame (not necessarily within a terrestrial scope). Most work in AL related to cognition attempts to develop cognitive abilities from artificial biological systems (whether computational models or physical agents). In this sense, it can be said that these abilities, though low level, are generically universal because they are generated from biologically universal systems. Furthermore, in all this work it is essential that cognitive abilities appear not as a result of a predefined purpose but as an ‘‘emergent’’ outcome of simpler systems.

Thus, if we consider that the preceding argument about the lack of universality of biology can be translated to cognitive science, it would be a natural step to produce a research program to fill the gap between cognition-as-we-know-it and cognition-as-it-could-be in which the development of artificial systems would play a major role. This fact poses a number of interesting questions the answers to which could be of great interest in the search for a general theory of cognition. First, it can be asked if artificial cognition of any kind is a specific target for Artificial Life. The question arises because of the difficulties of joining together this problem with other ones more essentially biological, such as origins of life, evolution, collective behavior, morphogenesis, growth and differentiation, development, adaptive behavior, or autonomous agency. Second, should the answer be positive, there would be a problem of the methodological status of the studies on low-level cognition: since it can be a common interest area for Artificial Intelligence and Artificial Life, it is not clear which methodology should be applied. Third, the study of the emergence of cognitive abilities in simple lifelike artificial systems might enlighten the evolutionary conditions for the origin of specialized cognitive systems to take place. This could be essential in the correct approach to more complex forms of cognition.

However, within Artificial Life itself we may distinguish two basic perspectives in facing the problem of designing cognitive agents: the ‘‘externalist’’ one and the ‘‘internalist’’ one. In the externalist position cognition is understood as a process that arises from an interactive dynamical relation, fundamentally alien to the very structure (body) of the cognitive agent, while according to the internalist position, cognition is the result of a (more) fundamental embodiment that makes it possible for evolution to create structures that are internally assigned interactive rules (Etxeberria, 1994). Most of the work done in computer simulations—and practically all in robotic realizations—belongs to the first perspective. For practical reasons, the ‘‘internalist’’ view hardly could be, by now, developed otherwise than by means of computational models.

In both positions autonomy and embodiment are established gradually. The externalist position is well represented by the aforementioned behavior-based paradigm, one of whose main characteristics is the fact of building physical devices for evaluating cognitive models. This represents an advantage in many aspects, because the interactions in real, noisy environments turn out to be much more complex than simulations.

In the externalist position, the parameters for controlling the agent are measured from the perspective of the agent itself, and the results are put in dynamic interaction with the effector devices. Performance is controlled by adaptive mechanisms that operate from the point of view of the agent itself: but the agent’s body is essentially only a place. Although this position supposes a significant advance with respect to the position of classic Artificial Intelligence and even with respect to some connectionist viewpoints, in fact it is still inside the traditional endeavor of designing cognitive agents disregarding the conditions that constitute them as generically autonomous, i.e., (full-fledged) biological systems. The consideration of the body essentially as a place means that the co-constructive (co-evolutionary) aspect of the interaction between agent and environment (Lewontin, 1982, 1983) is ignored. Autonomy (seen as the ability of self-modification) is restricted to the sensorimotor level (what Cariani (1989) has called the syntactical level of emergence). Thus, the plasticity of its cognitive structure is ultimately independent of its global structure (which is neither self-maintained, nor self-produced, nor evolutionary). As long as the autonomy in the solution to the cognitive problems involved in these agents is considered fundamentally external to the process of constructive self-organization of the very cognitive system (Moreno & Etxeberria, 1992; Etxeberria, Merelo, & Moreno, 1994), their ability to create themselves their own world of meanings (their autonomy) will be very limited (Smithers, this issue).

In the second perspective the cognitive autonomy of the agent is approached in a much more radical way, since its frame is the very biological autonomy. Nevertheless, we will see that positions that have been criticized in previous sections for their strict identification of the cognitive and biological mechanisms may reappear. We certainly have to agree with the idea in Varela et al. (1991) that the design of agents with cognitive functions should be understood within the frame of the very process that constitutes the agent as an autonomous entity (that is, its biological constitution). But as we have earlier said, this ability is not enough to explain the emergence of cognitive capacities.

Biology shows that the emergence of autonomous agents does not take place without (1) a process of constitution of a net of other autonomous agents and (2) a process that occurs through variations in reproduction and selection at its expression level. It is evident that in the biological frame the environment of a cognitive agent is mainly the result of the action (along with evolutionary processes) of the cognitive organisms themselves and other biological entities with which they have co-evolved. This is important because it means that, while the environment of biological systems is itself a biological (eco)system, the environment of cognitive agents is, to a great extent, a cognitive environment (communication).

Thus, the study of cognition in natural systems leads us to the conclusion that the functionality or cognitive meaning of the world for an agent emerges from this process of co-evolution. If we propose to apply this idea to the design of cognitive artificial systems it is because only from this perspective can a research program that ends up in the creation of true autonomous cognitive systems, i.e., that define their cognitive interaction with their environment by themselves, be established. This leads us to the necessity of adopting an Artificial Life research program in which evolutionary processes can have a fundamental role in the constitution in the agent of its own cognitive structure.

The so-called ‘‘evolutionary robotics’’ research project has tried to address this problem by redesigning the cognitive structure of the agent from an evolutionary perspective (in the current state of technology this cannot be made but in a purely computational universe[4]). In these models a phenotype and a genotype are considered the fundamental primitives of an evolutionary process. However, the phenotype as such is reduced to a nervous system scheme (that is, a neural net) (Floreano & Mondada, 1994; Yamauchi & Beer, 1994; Nolfi et al., 1995; Jakobi, Husbands, & Harvey, 1995; Gomi & Griffith, 1996). One of the most interesting aspects of these researches is the different attempts to evaluate in a realistic, physical context the evolutionary design of the cognitive system of the agents. In some cases there is even on-going physical evaluation of the computational evolutionary design, as in Harvey, Husbands, & Cliff (1994).

All this work means a significant advance, but it has a problematic identification between the phenotype of an agent and its nervous system. That is, the complementary relationship between nervous system and body, which we have argued is fundamental in previous sections, is still absent (because it does not exist a proper body). Hence, the problem of designing, in an

---

[4] Due to the impossibility of artificially building agents capable of self-production and self-reproduction and, even less, of disposing of the necessary long periods of time, we are obliged to resort to the computational simulation of evolutionary processes.

evolutionary scenario, agents having their structure set up as a complementary interplay between metabolic and neural organizations remains basically unexplored.

Some authors (Bersini, 1994; Parisi, this issue; Smithers, 1994) have presented critical proposals regarding this predominant approach of considering the phenotype of an agent only as its nervous system. However, the solution to this problem is linked to two deep questions very difficult to solve. One is how to generate from a global evolutionary process of the organism the structure of a system such as the NS. The other is how to generate cognitive abilities through a process of coevolutionary interaction among agents. We think that research about the origin of cognition must undertake as its main task the combined resolution of both kinds of problems.

However, the research program of evolutionary robotics is based on physical realizations. This circumstance conveys, given the level of current technology, a series of limitations for the exploration of the above-mentioned issues. Therefore, the study of such problems must be done fundamentally by means of computational models.

With respect to the first issue, some recent work offers interesting insights. This work develops models in which neuronlike structures are generated from evolutionary processes that produce cellular differentiation. The model by Dellaert and Beer (1995) shows an effort to avoid direct mapping from genotype to phenotype. This is achieved through the implementation of three successive levels of emergent structures (molecular, cellular, and organismal). In that sense, it represents an attempt to design epigenetic (ontogenetic or morphogenetic) processes to develop more realistic phenotypic structures. More recently, Kitano (1995) has developed another model in which the structure of a system similar to that of the nervous system appears through a process of cell differentiation. The most interesting aspect of Kitano's work is the generation of a ''tissue'' made of cells which are connected through axonal structures. Nevertheless, none of these addresses the emergence of cognitive functionalities.

There is another important question not addressed by these models: in the process of constitution of cognitive structures (and, in general, in the whole morphogenetic process) the interaction with the environment is not considered, and, therefore, the role that co-evolution with other cognitive organisms plays in the genesis and development of the cognitive system is ignored. If we want to understand in which way and under which conditions some organisms give origin to a cognitive system, it is necessary to have as starting point a collection of organisms that have developed a considerable complexity level. An interesting work that confronts the development of cognitive abilities in an artificial world from a co-evolutionary perspective is that of Sims (1994). In contrast with the previously mentioned models, in this case the stress is placed on the emergence of cognitive functionalities. In this model there is a bodylike structure formed by rigid parts. Better than inspired in biochemical-type processes, these parts behave more as physical mechanical structures. The fitness function is based on a contest in which organisms compete with each other.

An innovative advance of this model is that the neural net (though it is not metabolically embedded) is structured in two levels (local and global). However, this structure is introduced more in function of considerations about the physics of cognitive processes than of globally biological ones. In this sense, Sims' model conveys a greater abstraction of a series of processes situated in the interface between the metabolic and the neural level. Although it includes energetical considerations in the development of its cognitive functionalities, these considerations ignore the basic relation with the metabolic level (the network which ensures the self-maintenance of the whole system—the ''infrastructure'').

The problem is how to integrate these works with each other. In Kitano's model the emergent functionality is manifest through the formation of a neuronlike structure. Maybe, what is still lacking is two new levels in the model: first, a level at which newly formed neuronlike structures perform some control task—constraint—over the whole of the body; and, second, the appearance of a new level derived from a co-evolutionary process among organisms able to generate new functionalities as very basic cognitive behaviors.

This task becomes one of great complexity. It is difficult to determine which fundamental elements must be part of the model and which are disposable. The same happens at different levels, making the development of the model even more complicated. One of the biggest difficulties consists, surely, in searching transformation rules of genotypic structures into nonarbitrary phenotypic expressions (morphogenesis), which requires linking them to the realization of new functionalities. This, in turn, is linked to the generation of forms of ''downward causation'' (Campbell, 1974). All this implies serious difficulties, because the appearance of functional abilities cannot be facilitated by means of an artificial simplification of the rules at the high level (''complex'' parts) of the model.

What has been said to this point is more a review of the approaches to the problem of cognition within AL than a clear proposal of solutions. Nevertheless, we think that a correct estimation of the fundamental frame (underlying levels of complexity, etc.) within which the issue of the appearance of cognitive abilities is posed constitutes by itself an important advance considering the current context of AL (and AI too). It is true that in the research program on AL there is a characteristic emphasis on bottom-up methodology, as well as a greater insistence on the principle of embodiment with respect to the classical positions in AI. However, when reviewing most of the work that confronts the study of cognitive functionalities from the AL perspective, it is easy to see the lack of unanimity and even the absence of clear criteria regarding which kind of problem we ought to solve in order to adequately state the emergence of such capacities.

8. CONCLUSIONS

We have seen that the complexity of the relation between the system supporting cognitive abilities and the whole organism has entailed frequent misunderstandings. Sometimes the deep embeddedness of the cognitive system in its biological substrate has been ignored (as has happened and still happens in classical Artificial Intelligence, where the construction of disembodied artificial cognitive systems is attempted); in some others the autonomy of cognitive phenomena has been neglected, subsuming it in a generic adaptive behavior.

At the root of these difficulties is the fundamental problem of the origin of cognition. The answer given to this question determines the kind of research program in cognitive sciences and, even more, the autonomy of the cognitive sciences with respect to biology, on the one hand, and its grounding, on the other. The problem is that neither biology nor cognitive science now provide a satisfactory theory about the origin of cognitive systems. AL research can, however, help in developing such a theory. In this way, the knowledge that we gradually acquire about the conditions that make possible the arising of cognitive systems in artificial organisms will be endowed with a higher generality than classical biological studies.

What we have proposed here is that the origin of cognition in natural systems (cognition as we know it) is the result of the appearance of an autonomous system—the nervous system—embedded into another more generic one—the whole organism. This basic idea is complemented by another one: the formation and development of this system, in the course of evolution, can be understood only as the outcome of a continuous interactions between organisms and environment, among different organisms, and, especially, among the very cognitive organisms.

The possibilities of generalizing this conception of the origin of cognition rest in AL. AL offers new tools that make it possible to establish the foundations of a theory about the origin of cognition as-it-could-be. This should be, precisely, the bridge between Artificial Life and Artificial Intelligence. Our suggestion is that investigations in AL should satisfy the two previously mentioned conditions—autonomy and co-evolution—in order to connect with the foundations, in turn, of a new research program in AI.

It is conceivable to hope that the results of all this might influence the research programs of both Artificial Life and Artificial Intelligence so that they gradually converge, though not necessarily in a global merging process, but finding a well-established common research area. This mutual finding seems today more likely since within Artificial Intelligence there is increasing research on situated systems, with more and more degrees of autonomy, and whose main ability does not concern the solution of very complex problems but the ability to functionally modify the statement of easier problems—that is, on agents capable of doing simple things in a more autonomous way. In turn, systems that could be considered part of ‘‘the primordial soup’’ allowing the emergence of agents with primitive cognitive functions are starting to be taken into consideration within Artificial Life. Should this confluence be achieved, Artificial Life would have contributed not only to establishing the bases of biology as the science of all possible life, but also to establishing cognitive science as the science of all possible cognition.

REFERENCES

Andler, D. 1988. *Representations in cognitive science: Beyond the pro and con.* Manuscript.

Ashby, W. R. 1956. *An introduction to cybernetics.* London: Chapman & Hall.

Bersini, H. 1994. Reinforcement learning for homeostatic endogenous variables. In D. Cliff et al. (Eds.), *From animals to animats 3.* Cambridge, MA: MIT Press. pp. 325–333.

Bertalanffy, L. von. 1968. *General systems theory; foundations, development, applications.* New York: George Braziller.

Bertalanffy, L. von. 1975. *Perspectives on general systems theory. Scientific-philosophical studies.* New York: George Braziller.

Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence,* **47,** 139–159.

Brooks, R., & Maes, P. Eds. 1994. *Artificial life IV.* Cambridge, MA: MIT Press.

Campbell, D. T. 1974. Downwards causation in hierarchically organized biological systems. In F. J. Ayala & T. Dobzhansky (Eds.), *Studies in the philosophy of biology.* London: Macmillan. Pp. 179–186.

Cariani, P. 1989. *On the design of devices with emergent semantic functions.* Ph. D. dissertation, State University of New York at Binghamton.

Clark, A. 1995. Autonomous agents and real-time success: Some foundational issues. In *IJCAI’95.*

Clark, A., & Grush, 1996. *Towards a Cognitive Robotics.* Personal Manuscript.

Clark, A., & Toribio, J. 1994. Doing without representing? *Synthese,* **101,** 401–431.

Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, J. S. Eds. 1994. *From animals to animats 3, Proceedings of the Third Conference on Simulation of Adaptive Behaviour SAB94.* Cambridge, MA: MIT Press.

Cowan, G. A., Pines, D., & Meltzer Eds. 1994. *Complexity.* Reading, MA: Addison–Wesley.

Churchland, P. S., & Sejnowski, T. 1993. *The computational brain.* Cambridge, MA: MIT Press.

Dellaert, F., & Beer, R. 1995. Toward an evolvable model of development for autonomous agent synthesis. In R. Brooks & P. Maes (Eds.), *Artificial life IV.* Cambridge, MA: MIT Press. Pp. 246–257.

DRABC’94. *Proceedings of the III International Workshop on Artificial Life and Artificial Intelligence ‘‘On the Role of Dynamics and Representation in Adaptive Behaviour and Cognition.’’* Dept. of Logic & Philosophy of Science, University of the Basque Country.

Emmeche, C. 1994. *The garden in the machine: The emerging science of artificial life.* Princeton, NJ: Princeton University Press.

Etxeberria, A. 1994. Cognitive bodies. In *DRABC’94.* Pp. 157–159.

Etxeberria, A. 1995. Representation and embodiment. *Cognitive Systems,* **4**(2), 177–196.

Etxeberria, A., Merelo, J. J., & Moreno, A. 1994. Studying organisms with basic cognitive capacities in artificial worlds. *Cognitiva,* **3**(2), 203–218; *Intellectica,* **18**(1), 45–69; *Kognitionswissenschaft,* **4**(2), 75–84; *Communication and Cognition-Artificial Intelligence,* **11**(1–2), 31–53; *Sistemi Intelligenti,* **IV**(3), 443–465.

Floreano, D., & Mondada, F. 1994. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In D. Cliff et al. (Eds.), *From animals to animats 3.* Cambridge, MA: MIT Press. Pp. 421–430.

Gibson, J. J. 1979. *The ecological approach to visual perception.* Boston, MA: Houghton–Mifflin.

Gomi, T., & Griffith, A. 1996. Evolutionary robotics–an overview. In *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation (ICEC 96)*, Nagoya (Japan) May 20–22. Pp. 40–49.

Harnad, S. 1990. The symbol grounding problem. *Physica D* **42,** 335–346.

Harvey, I., Husbands, P., & Cliff, D. 1994. Seeing the light: Artificial evolution, real vision. In D. Cliff et al. (Eds.), *From animals to animats 3.* Cambridge, MA: MIT Press. Pp. 392–401.

Haugeland, J. 1991. Representational genera. In W. Ramsey, S. Stich, & D. Rumelhart (Eds.), *Philosophy and connectionist theory.* Hillsdale, NJ: Erlbaum. Pp. 61–90.

Heschl, A. 1990. L = C. A simple equation with astonishing consequences. *Journal of Theoretical Biology,* **185,** 13–40.

Jakobi, N., Husbands, P., & Harvey, I. 1995. Noise and the reality gap: the use of simulation in Evolutionary Robotics. In F. Morán et al. (Eds.) *Advances in artificial life.* Berlin: Springer. Pp. 704–720.

Keeley, B. L. 1993. Against the global replacement: On the application of the philosophy of AI to AL. In C. Langton (Ed.), *Artificial life III.* Reading, MA: Addison–Wesley. Pp. 569–587.

Kitano, H. 1995. Cell differentiation and neurogenesis in evolutionary large scale chaos. In F. Morán et al. (Eds.), *Advances in artificial life.* Berlin: Springer. Pp. 341–352.

Langton, C. Ed. 1989. *Artificial life.* Reading, MA: Addison–Wesley.

Lewontin, R. C. 1982. Organism and environment. In H. C. Plotkin (Ed.), *Learning, development, and culture.* New York: Wiley. Pp. 151–170.

Lewontin, R. C. 1983. The organism as the subject and the object of evolution. *Scientia,* **118,** 65–82.

Maes, P. Ed. 1990. *Designing autonomous agents: Theory and practice from biology to engineering and back.* Cambridge, MA: MIT Press.

Maturana, H. R., & Varela, F. J. 1980. *Autopoiesis and cognition.* Dordrecht: Reidel (Kluwer).

Meyer, J.-A., & Wilson, S. W. Eds. 1991. *From animals to animats 1. Proceedings of the First International Conference on Simulation of Adaptive Behavior.* Cambridge, MA: MIT Press.

Meyer, J.-A., Roitblat, H. L., & Wilson, S. W. Eds. 1992. *From animals to animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior.* Cambridge, MA: MIT Press.

Morán, F., Moreno, A., Merelo, J. J., & Chacón, P. Eds. 1995. *Advances in artificial life. Proceedings of the 3rd European Conference on Artificial Life (ECAL95).* Berlin: Springer.

Moreno, A., & Etxeberria, A. 1992. Self-reproduction and representation. The continuity between biological and cognitive phenomena. *Uroboros,* **II**(1), 131–151.

Moreno, A., Etxeberria, A., & Umerez, J. 1994. Universality without matter? In R. Brooks & P. Maes (Eds.), *Artificial Life IV.* Cambridge, MA: MIT Press. Pp. 406–410.

Moreno, A., Umerez, J., & Fernández, J. 1994. Definition of life and research program in Artificial Life. *Ludus Vitalis. Journal of the Life Sciences,* **II**(3), 15–33.

Newell, A. 1980. Physical symbol systems. *Cognitive Science,* **4,** 135–183.

Nolfi, S., Floreano, D., Miglino, O., & Mondada, F. 1995. How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks & P. Maes (Eds.), *Artificial Life IV.* Cambridge, MA: MIT Press. Pp. 190–197.

Parisi, D. 1997. Artificial life and higher level cognition. *Brain and Cognition,* **34.**

Pattee, H. H. 1982. Cell psychology: An evolutionary approach to the symbol–matter problem. *Cognition and Brain Theory,* **5**(4), 325–341.

Pattee, H. H. 1986. Universal principles of measurement and language functions in evolving systems. In J. L. Casti & A. Karlqvist (Eds.), *Complexity, language, and life.* Berlin: Springer–Verlag. Pp. 268–281.

Pattee, H. H. 1987. Instabilities and information in biological self-organization. In F. E. Yates (Ed.), *Self-organizing systems. The emergence of order.* New York: Plenum. Pp. 325–338.

Pattee, H. H. 1989. The measurement problem in artificial world models. *BioSystems,* **23,** 281–290.

Pattee, H. H. 1993. The limitations of formal models of measurement, control, and cognition. *Applied Mathematics and Computation,* **56,** 111–130.

Pattee, H. H. 1995. Evolving self-reference: Matter, symbols, and semantic closure. *Communication and Cognition-Artificial Intelligence,* **12**(1–2), 9–27.

Pines, D. Ed. 1987. *Emerging synthesis in science.* Reading, MA: Addison–Wesley.

Pylyshyn, Z. 1984. *Cognition and Computation.* Cambridge, MA: MIT Press.

Simon, H. A. 1969 (1981, 2nd ed.). *The sciences of the artificial.* Cambridge, MA: MIT Press.

Sims, K. 1994. Evolving 3D morphology and behavior by competition. In R. Brooks & P. Maes (Eds.), *Artificial life IV.* Cambridge, MA: MIT Press. Pp. 28–39.

Smithers, T. 1994. What the dynamics of adaptive behaviour and cognition might look like in agent–environment interaction systems. In *DRABC'94.* Pp. 134–153.

Smithers, T. 1997. Autonomy in robots and other agents. *Brain and Cognition,* **34.**

Smolenski, P. 1988. On the proper treatment of connectionism. *Behavioral and Brain Sciences,* **11,** 1–74.

Sober, E. 1992. Learning from functionalism. Prospects for strong artificial life. In C. G. Langton, J. D. Farmer, S. Rasmussen, & C. E. Taylor (Eds.), *Artificial life II.* Reading, MA: Addison–Wesley. Pp. 749–765.

Stewart, J. 1992. Life = Cognition. The epistemological and ontological significance of artificial intelligence. In F. Varela & P. Bourgine (Eds.), *Toward a practice of autonomous systems. Proceedings of the 1st European Conference on Artificial Life (ECAL91).* Cambridge, MA: MIT Press. Pp. 475–483.

Thompson, E. 1997. Symbol grounding: A bridge from artificial life to artificial intelligence. *Brain and Cognition,* **34.**

Turvey, M., & Carello, C. 1981. Cognition: The view from Ecological Realism. *Cognition,* **10,** 313–321.

Umerez, J. 1995. Semantic closure: A guiding notion to ground artificial life. In F. Morán et al. (Eds.), *Advances in Artificial Life.* Berlin: Springer. Pp. 77–94.

Umerez, J., & Moreno, A. 1995. Origin of life as the first MetaSystem Transition—Control hierarchies and interlevel relation. *World Futures,* **45,** 139–154.

Van Gelder, T. 1992/1995. What might cognition be if not computation? *Technical Report 75,* Indiana University, Cognitive Sciences/*Journal of Philosophy,* **XCII**(7), 345–381.

Van Valen, L. 1973. A New Evolutionary Law. *Evolutionary Theory,* **1,** 1–30.

Varela, F. J. 1989. *Connaître. Les sciences cognitives, tendences et perspectives.* Paris: Seuil.

Varela, F. J. 1997. Patterns of life: Intertwining identity and cognition. *Brain and Cognition,* **34.**

Varela, F. J., Thompson, E., & Rosch, E. 1991. *The embodied mind. Cognitive science and human experience.* Cambridge, MA: MIT Press.

Wiener, N. 1948 [1961]. *Cybernetics or control and communication in the animal and the machine.* Cambridge, MA: MIT Press.

Yamauchi, B., & Beer, R. 1994. Integrating reactive, sequential, and learning behavior using dynamical neural networks. In D. Cliff et al. (Eds.), *From animals to animats 3.* Cambridge, MA: MIT Press. Pp. 382–391.