On the appeals and dangers of synthetic reductionism

Eric W. Bonabeau

CNET Lannion B - OCM / TEP
Route de Trégastel
22301 Lannion Cédex
France
bonabeau@lannion.cnet.fr

Introduction

Common wisdom holds that being reductionist implies tracking the causes of macroscopic phenomena down to the level of elementary particles or more loosely down to the level of physical laws, which for instance any biological system ultimately obeys or seems to obey. There may be variations in this classical view of reductionism, since one can be ontologically, methodologically or epistemologically reductionist, depending on the extent to which one accepts and/or practices reductionist principles. But reductionism is always considered essentially analytic. This is a profound misconception: Science can be both reductionist and synthetic at the same time. Artificial Life (AL) is an example of that possible coexistence. More generally, the so-called "sciences of complex systems", in which AL is to some extent embedded, are based on synthetic-reductionist approaches to the study of natural systems. In this paper, we shall show how synthetic reductionism manifests itself in AL as well as within the larger domain of the sciences of complex systems, and highlight the appeals and dangers of such a way of practicing science, especially when it comes to the sciences of the living.

1 Reductionism and the nature of Artificial Life

Following [Wimsatt 86], we shall say that "...a reductionist is interested in understanding the character, properties, and behavior of the studied system in terms of the properties of its parts and their interrelations and interactions. This means that the reductionist is primarily interested in the entities and relations internal to the system of study". But Wimsatt added: "This is a sufficiently inclusive description that it probably captures any analytic method in general,..." unnecessarily restricting the scope of reductionism to the realm of analytic methodologies, while its definition does not refer to any kind of analysis. And from this definition, it is clear that AL, though synthetic, is 100% reductionist. That AL is a reductionist program has been obvious to its founders from the beginning: "The field of Artificial Life is unabashedly mechanistic and reductionist" [Langton 87]. But this very simple fact has been somewhat forgotten in the excitment. This is mainly due to the analytic view of reductionism which is most commonly shared. See for instance [Taylor 91]: "Reductionist science, ..., breaking things down (...) misses emergent properties". As was said in the introduction, there is no reason why reductionism should always imply analysis as is illustrated by the sciences of complex systems. If such sciences share some common properties with other classical reductionist sciences, they also exhibit some more specific features due to the mostly synthetic methods they resort to. And in the case of Artificial Life, its synthetic nature manifests itself in combination with its reductionist nature under some peculiar forms we shall describe in this paper. In particular, besides the usual flaws of reductionist methodologies when applied to the study of biological systems, the synthetic approach implies weakened explanatory status of models, huge spaces of exploration, absence of constraints (this is also caused by the life-as-it-could-be program), and the combination of reductionist approaches and synthetic methods leads to a lack of open-endedness, especially in the context of computational reductionism. While our object of study in this paper is AL, a lot of what is said can apply to a large extent to the emerging sciences of complex systems, which are also by nature mostly synthetic and reductionist.

Our point here is to discuss the advantages and the limits of this type of synthetic reductionist approach. The advantages are obvious: for instance AL constitutes an alternative to purely analytic natural sciences: "Alife represents a non-traditional approach to the study of emergent properties in nature and permits the study of issues that have otherwise been very difficult to address " [Taylor 91], and allows for a much wider spectrum of exploration compared to classical sciences which have well-defined objects of study - they start with the object of study and investigate its properties, while AL investigates some "dynamical" properties and finds objects. The limits must be studied in a more speculative manner simply because AL is a way of determining the extent they effectively constitute limits . We see mostly two limits to AL's approach: 1) AL's scope is too broad, and 2) AL's scope is not broad enough. It is easy to understand these seemingly incompatible facts:

1) By using synthetic exploration, AL deals with all the phenomena such an exploration may allow, thus the space of possibilities is too huge (otherwise stated, life-as-it-could-be is dramatically ill-defined). It seems that this sort of limit can be ignored by resorting to higher-level constraints, as we already claimed in [Bonabeau & al. 92]: AL must not fall into the trap of forgetting the higher-level sciences (see e.g "AI has for the most part neglected the fundamental biology of the nervous system. This criticism will undoubtedly be aimed at AL as well"[Pattee 89]), and must on the contrary accept the empirical constraints provided by the observations of these higher-level sciences, even if the ultimate hope is to go beyond the sciences of what exists toward the sciences of what could have existed -the latter providing (of course !) no observation at all. Moreover, how can we scientifically assess the validity of models without resorting to constraints ? If we resort to synthesis with only the goal of phenomenologically reproducing observed behaviors, it is hard to determine the extent to which a model explains the behavior it reproduces. Besides AL's program looks ambivalent since it needs to reduce the intractable number of possibilities offered by physics but not too much, so as to be capable of capturing even 'life-as-it-could-be'. Even more ambiguous is the way we judge AL's simulations; such a judgement is largely based on our intuitions, experience, and even emotions, which is in contradiction with AL's ambition to synthesize life-as-it-could-be: we judge simulations based on how well they meet our aesthetic requirements, which themselves rely on our experience of life-as-we-know-it (what other experience could we have ?). As a consequence, we will never be able to recognize or synthesize forms of life that are really far from life-as-we-know-it. Thus, instead of ambiguously and dangerously refusing constraints by defining a self-contradictory program, AL should make clear what constraints it chooses to be based upon. All this reminds us of an artistic approach: building an AL's creature, be it a cellular automaton, amounts to make some set of equations and our subconscious meet, like an artist. This parallel is not surprising if one remembers the importance of sensorial media (like videotapes or computer graphics) in AL demos.

2) Artificial Life's synthetic exploration procedure is based on the hope that simple (most often formal) elements in interactions will generate a sufficient richness of behaviors peculiar to life, but one may miss important phenomena because some external variables or conditions, accidental from the point of view of the model (i.e not taken into consideration by the model), may reveal crucial to the generation of behaviors constituting the

essence of life. These conditions, which are essential to the generation or the understanding of a particular phenomenon, are thereafter called "boundary conditions", and can be internal as well as external. In this paper, we shall describe at some length this notion of boundary conditions to show that Artificial Life relies on the hope that most "interesting" behaviors can be generated "internally" in some sense, that is with simple models of many non-linearly interacting elements. That second limitation has to do with the notion of open-endedness and is in close relationship with the criticisms made by systems science people [Cariani 89, 91, Kampis 91, Pattee 89, Rosen 78, 85]. This type of approach is common to almost all sciences of complex systems and raises the need for new epistemological concepts. That is why we will take some time to speak of complex systems in the next paragraph before coming back to Artificial Life in the last one.

2 The sciences of complex systems and the need for new epistemological concepts

Complex systems

The most commonly shared definition of a complex system states that it is a network of interacting objects, agents, elements or processes that exhibit a dynamic, aggregate behavior. The action of an object (possibly) affects subsequent actions of other objects in the network, so that the action of the whole is more than the simple sum of the actions of its parts. In other words, a system is complex if it is not reducible, in the following sense: since the beginning of time, science has been very busy dealing with either systems whose behaviors are reducible to a few degrees of freedom and thus can be characterized by low-dimensional deterministic equations, or systems with many degrees of freedom but whose behavior is reducible to a statistical description. A complex system has many degrees of freedom which strongly interact with each other, preventing from either of the two classical reductions. In a nutshell, it exhibits what Weaver called organized complexity [Weaver 68] (as opposed to organized simpicity and disorganized complexity). Besides, the complementary idea of chaos taught us that unpredictability can also arise in low-dimensional deterministic systems, showing that even "reducible" systems can be very hard to deal with. Systems such as those described as complex are by far the most numerous in nature. Since science has not been able to investigate (quantitatively and qualitatively) these systems in a satisfactory manner so far, it seems a good idea to look for laws, tools and methods originating from any field which has to deal with such systems, with no a priori restriction in the scope of research.

Why the sciences of complex systems did not emerge before has a simple reason: there was a sort of unexplored niche in the gigantic scientific ecology, because complex systems, be they low-dimensional chaotic, or high-dimensional with non-linearly coupled degrees of freedom, could not be studied before the last decades because they require high computational power -far beyond the (unaided) human brain's capabilities. **Synthesis has strong computational requirements.** Computer studies gave birth not only to quantitative results but also to theories. The best example is the scientific activity which has developed around dynamical systems, 10 to 15 years after the discovery *on a computer* of the notion of sensitivity to initial conditions in low-dimensional systems. Fractals are another example: we all know how to readily generate beautiful complex pictures with a computer, but what could we do without the computer ?

Disciplines adopting this type of approach can only benefit from others' experience. Transversal motion of concepts provides mutual enrichment to an extent never reached so far. Especially in the sciences of complex systems, because there seems to be a lot of common principles to be shared, almost all based on the popular notion of emergence. In fact, these principles can be considered **methodological** consequences of the very irreducibility of the objects of study of the sciences of complex systems: the synthetic nature of these sciences is explained by the fact that analytic methods fail to capture the essence of such "irreducible" objects by definition, because analytic methods apply efficiently only in those cases where some relevant simplifying tricks (reduction tools) allow for the transformation of the initial problem onto a tractable one with not too many variables. When reduction is not possible, analysis explodes (it can be for instance a combinatorial explosion), and must be replaced by synthesis, i.e emergence which then appears as a common principle while it is only a common methodology.

The following well-known diagram

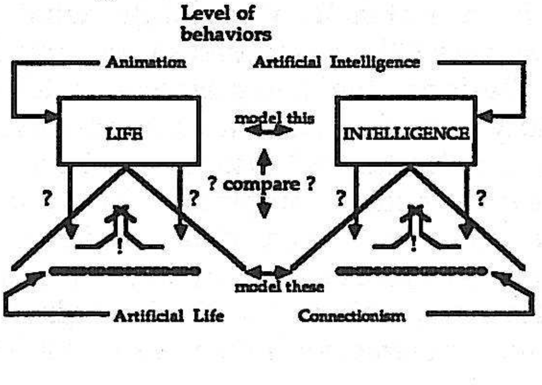

explains why theories of emergence constitute the essence of AL (as well as of most other sciences of complex systems): since it seems to be more difficult to start from manifestations of life and try to find its fundamental principles by top-down analysis, than to start from computational and physical simulations and try to synthesize more and more complex behaviors which in turn might capture the nature of some aspects of life, AL thus focuses on ways of achieving emergence to generate these more and more complex behaviors. "The most surprising lesson we have learned from simulating complex physical systems on computers is that complex behavior need not have complex roots. Indeed, tremendously interesting and beguilingly complex behavior can emerge from collections of extremely simple components" [Langton 89]. This also explains why People in Artificial Life as well as in the sciences of complex systems share an irrational faith in the power of emergence, though everybody acknowledges that "the concept of emergence in itself offers neither guidance on how to construct such a[*n emergent*] system nor insight into why it would work" [Hillis 88]: this makes AL and the sciences of complex systems even more reductionist than most classical reductionist sciences, because for them the laws of physics within a given system are almighty; not only must the system comply with physical rules, it is also defined by them since they generate sufficient 'boundary conditions' by themselves. This may lead to dangerous abuses: see e.g "Emergence offers a way to believe in physical causality while simultaneously maintaining the impossibility of a reductionist explanation of thought. For those who fear mechanistic explanations of the human mind, our ignorance of how local interactions produce emergent behavior offers a reassuring fog in which to hide the soul " [Hillis 88], or "if you take care of the computational setup, living behavior will emerge by itself" [Emmeche 91].

Let alone the danger of outrageously worshipping emergence, transversality of concepts can also be dangerous when it is not appropriately applied: for instance, synergetics has a particular way of applying transversality, "a direct comparison of physical and social systems on the phenomenological level can only lead to a superficial, short breathed analogy lacking structural depth", "deep and rather universal analogies between social and physical systems (...) reflect the fact that, due to the universal applicability of certain mathematical concepts to multi-component systems, all such systems exhibit an indirect similarity on the macroscopic collective level, which is independent of their possible comparability on the microscopic level" [Weidlich 91]. One can agree or not with the appropriateness of this method, but one fact cannot be questioned: all natural objects, be they physical, biological, social or else, are observed through systems, that is only a limited set of observables are chosen, and syntactic relationships are looked for between these observables to account for their (observed) behaviors. Two systems can share some similarities with respect to some set of observables, while they completely diverge when it comes to other observables. Thus, one must be very cautious when dealing with resemblances not to confuse these necessarily partial resemblances with global analogy at all levels of description and with respect to all possible sets of observables. In Artificial Life in particular, due to the lack of constraints, phenomenological relationships are almost the only criteria that can be used to judge simulations and models ("simulations are metaphorical, not litteral" [Pattee 88]), while theories usually "are judged by a much more comprehensive range of criteria, from the concrete test of how well they can predict specific values for the observables of the system being modeled, to abstract tests such as universality, conceptual coherence, simplicity and elegance" [Pattee 88]. Most of these criteria disappear when it comes to Artificial Life, which makes scientific explanation very weak in this case. One could argue that Artificial Life models are the simplest ones in some sense, since they rely on simple elements in interactions, and that it is epistemological common sense to start with simple models than with complicated ones. This is not completely true for at least two reasons: 1) one often starts with many more variables than necessary, and gradually simplifies the model to retain only relevant variables; 2) simplicity is not necessarily a quality as regards biological sciences, it is almost a counterintuitive criterion.

All these facts about the sciences of complex systems raise epistemological issues. For instance, there is an obvious mutation of the concept of scientific explanation. In the next paragraph, we are going to explore a sort of hierarchy of scientific explanations, which naturally will lead us to the notion of boundary conditions.

The need for a new epistemology

Putnam [Putnam 73] reached an interesting conclusion by making first a difference between 'to deduce' and 'to explain': being able to deduce the properties of a phenomenon from a set of causes is not equivalent to explaining this phenomenon, because only a few among the many possible causes may be relevant, "certain systems can have behaviors to which their microstructure is largely irrelevant" [Putnam 73]. Explaining the phenomenon amounts to determining what the relevant causes are. It may also simply be impossible to deduce the properties of a phenomenon from a set of causes originating from one single discipline: this is so because "the laws of the higher-level discipline are deducible from the laws of the lower-level discipline together with 'auxiliary hypotheses' which are accidental from the point of view of the lower-level discipline"[Putnam 73]. The laws of the higher-level discipline therefore depend on both the laws of the lower-level discipline and 'boundary conditions' which are "accidental from the point of view of physics but essential to the description of "[Putnam 73] the higher level. It is through the huge space of possibilities allowed by physics and through the many possible accidental causes that higher-level phenomena are somewhat autonomous relative to other levels. This is our first introduction of the concept of boundary conditions, to which we shall return later in this paper.

Thus, if we summarize Putnam's ideas, we may state that (using Putnam's

terminology):
- a great number of laws of the lower-level are irrelevant to the understanding of higher-level phenomena;
- other laws originating from other perspectives are essential for the understanding of the higher-level phenomena, but purely accidental at the lower-level .

In the same spirit, one can see a major obstacle emerge: any kind of higher-level structure, be it self-generated, can be very hard to deal with, due to the fact that explanation is not transitive (i.e explanations at one level are not of the same nature as explanations at another level), which gives some unpleasant autonomy to higher levels relative to (explanations at) lower levels. If we say, as in [Weidlich 91] that "a level is a stratum of reality of a certain self-contained organization", that is with a "quasi autonomous" dynamics", then 1)- the immensity of the phase space allowed by the physics of level 1 can make the behavior of level 2 unpredictable, i.e it may be impossible to have any idea about shapes and structures appearing at the higher level given the laws of physics 2)- each passage from one level to another has its own boundary conditions, 3)- and external boundary conditions (external causes) are accidental [it is worth noticing that these external boundary conditions may also be generated by the higher levels in which the level under study is embedded]. As a consequence, it can be very hard to find tools to deal with higher levels when starting from a given level, and usually, "it appears that the lower level provides the constituent units for the next higher level only" [Weidlich 91].

Practicing science across levels or practicing synthetic science imply the redefinition of scientific explanation, falsifiability, etc... The main strategy of science until very recently has always been to reduce the understanding of a phenomenon to a (very) few essential degrees of freedom. There is an implicit 'simplicty hypothesis' underlying much of epistemology, and there is no reason why this implicit 'simplicity hypothesis' should hold forever, since, as we have argued before, irreducible systems are more often found in nature than reducible ones. In effect, when the complexity of a problem exceeds our power of analusis, we resort to idealizations and approximations, that is we reduce the complexity of the problem. But this is no longer possible when it comes to complex, irreducible systems. Thus, if scientific induction is superficially easy to understand in the sciences of complex systems (induction = synthetic exploration), deduction (and explanation) is becoming tricky. Classically, deduction and explanation were often confused (despite Putnam's warnings), simply because reduction was trivial (explanation can be seen as the projection of deduction on the space of relevant causes, thus if reduction to the relevant causes is trivial, explanation and deduction can be confused). Now, if reduction is no longer trivial or even possible, we have to define a hierarchy of explanations:

0- classical explanation, where a few relevant causes are pointed out;
1- since explanation is not transitive, explanation at one level is not of the same nature as explanation at another level;
2- logical deduction is different from explanation since among the many causes used to deduce an effect, only some of them may be relevant: this comes from the fact that in monotonic logic if A=>B then A&C=>B;
3- weak deduction can be defined as deduction using a real (as opposed to ideal) computer: on can use a computer to deduce (or derive) a result;
4- deduction, even if weak may take a long, long time, so that it is no longer possible to deduce "in practice";
5- deduction can be impossible even in principle (due to Gödel's theorem).

Points 0 to 3 are obviously positive definitions, and everyday language (in science) tends to reduce points 1, 2, and 3 to point 0. For instance, it is by now admitted that one can "explain" phenomenon using a computer. Note also that if this hierarchy does not completely illuminate the issue of the status of explanation based on phenomenological resemblances, defining this hierarchy was at least a prerequisite.

[The following excerpt from [Bonabeau & Bourgine 92] may help understand this hierarchy better: What does it mean that one cannot 'derive' the properties of a system at a particular level given a model of how a lower level behaves ? There are several possible (and not incompatible) answers to this question. For the sake of simplicity, we describe these answers in a formal setting: all processes we are talking about are computational. We shall see that the concept of detector is of importance.

- We can understand the notion of derivation in the context of Gödel's theorem. In this case, derivation is something very formal, it designates a way of proving new theorems from a set of axioms. Not being able to derive observed higher-level properties from the set of lower-level properties means not being able to prove theorems about higher-level entities while we know that these theorems are true. We know since Gödel that such a situation is possible. The only way of knowing the properties of the higher-level entities is to detect them, to be sensitive to them, since it is impossible, even 'in principle' to deduce them. The observed properties are of axiomatic nature.
- Another cause for the impossibility to derive the higher-level properties from lower-level considerations is the time -and possibly the 'memory' space- it would take to examine all the possible properties of the higher-level and determine those that are of interest. Time compression is allowed by computers, and in that sense we may speak of computational emergence since computers enable us to observe higher-level properties, indeed contained in the specification of the system from the beginning, but which are out of reach for a human being, because the space of possible properties is too huge. Computers might even not be sufficient, if the number of dimensions of the space of possibilities is 'immense' in Elsasser's sense [Elsasser 81], i.e a number that is 'not tractable and cannot be acted upon with present-day computers'[Kampis 91].
- There is yet another reason for which computers might not be the right tools to observe emergent properties: these properties may be logically deep, that is it might take a very, very long time even with the best program ever -not to speak of the difficulty of finding this

program- to compute these properties. It is not at all clear whether logical depth is due to the high dimensionality of the system under observation or to other factors intrinsic to the property (or the shape) which is the object of the computation. In this case, the only way of discovering a new 'deep' property is to have a detector which is at least as deep as the property, and having a high value of mutual information with that property due to a common history. The idea behind this is again that detectors coevolve with the shapes they can detect, or even if the shapes preexist in some way to the detectors, these latter have the same depth. It is interesting to note here the similarity between sophistication and logical depth: the depth of an object being roughly the length of the program producing it (without taking account of the input data), the deeper a property, the longer the program needed to 'compute' it.

- The issue of undecidability is also of interest here: there are some 'axioms' that are undecidable, that is their being true or false does not change anything to the set of axioms describing the system. But some theorems may be very long to prove without these axioms, while they are very short with them. The link between explanation and undecidability is that some higher-level property can be very difficult to derive (with the different meanings previously exposed) without a particular 'undecidable' device, namely a detector.
- Finally, we can come back to Putnam's distinction between 'to deduce' and 'to explain' [Putnam 73]

3 Artificial Life and Boundary Conditions

AL lost in immensity

Studying life-as-it-could-be (a sort of side effect of the synthetic, emergence-based nature of AL) may constitute an intractable task, all the more as 'real life' already covers a large spectrum of possible behaviors: "Life is self-organizing in the sense that it leads to very special forms, i.e from a wide basin of attraction it leads to a much narrower set of meaningful states. But this alone would not yet be surprising: the surprising aspect is that this attraction is not at all rigid. Although the attractor is very small compared to full phase space, it is still huge and it therefore allows for a wide spectrum of behaviors "[Grassberger 89]. That is why we should follow Sober's suggestion to approach the general questions of the 'nature of mind or the nature of life' by "focusing on more specific psychological and biological properties... this strategy makes the general questions more tractable" [Sober 91]. We propose to focus on the notion of autonomous system, while accepting the existence of empirical constraints, because if "to understand mind and life we must abstract away from physical details", "the problem is to do this without going too far" [Sober 91]. Using Putnam's words [Putnam 88], the only way for AL not to be "*one damned thing after another*" is to have accept empirical constraints and eventually have a 'Master Program' (why not autonomous systems ?), otherwise AL reasearchers would be tinkers -like evolution-, and the number of "damned things" we may think of may be *astronomical*. There are other sciences dealing with life, adaptation, evolution, ..., which can provide sufficient constraints. The best bottom-up approach needs some kind of validation by top-down data. This data constitute phenomenological and empirical "boundary conditions" which must be satisfied. In the next section, we deal with a completely different type of boundary conditions.

Boundary conditions

Coming back to Wimsatt's general definition, we see that being reductionist leads to a particular interest in the "entities and relations internal to the systems": this is the crux of the original properties of Artificial Life. In a previous article [Bonabeau & al. 1992], the notion of 'self-generated boundary condition' was proposed as the basic, implicit principle motivating most Artificial Life's research. The idea behind boundary conditions [Polanyi 68] is that the functioning of a system (say a machine) is highly underdetermined by lower-level considerations (say physical laws): it is the design of the system that constitutes the right level of description in order to understand how this system works, not the level of physical laws. It is true that the higher level has a behavior which is compatible with physical laws, but physical laws alone are unspecific, they cannot determine the behavior of the higher level: boundary conditions make the link between the two levels by 'directing lower-level processes to definite channels' [Kampis 91]. Vitalistic conclusions may easily be drawn from these considerations, if one believes that irreducible boundary conditions underlie the appearance of life: in effect, among the **immense** [Elsasser 81] number of possible states of the world allowed by physics, only a few are compatible with life, and such compatibility may not be deducible from the laws of physics. The idea of boundary condition is closely linked to Elsasser's immensity [Elsasser 81], to Pattee's non-holonomic constraints and to Rössler's privileged zero property, nicely summarized in [Kampis 91]. The notion of **self-generated** boundary condition is by now easy to grasp: we use this terminology to describe the property of some systems which generate boundary conditions from inside (when nonlinear laws of interaction are present), i.e which exhibit a highly specific (in some sense to be defined) behavior without the help of any exogeneous phenomenon. Such a phenomenon would be 'purely accidental' from the point of view of the internal dynamics [Putnam 73].

Reductionism and external boundary conditions

- Reductionists do not like environments

As emphasized in [Wimsatt 86], reductionists usually look for internal explanations (intrasystemic mechanisms) rather than for external causes (intersystemic mechanisms), and in any case internal mechanisms are almost always considred more "fundamental". Reductionists "simplify the description of the environment before simplifying the description of the system", "construct experimental arrangements so as to keep environment variables constant", in a nutshell, they "ignore, oversimplify, or otherwise underestimate the importance of the context of the system under study" [Wimsatt 86]. But, for example, "evolution depends on a result of microstructure (variation in genotype) but it also depends on conditions (presence of oxygen) which are accidental from the point of view of physics and chemistry" [Putnam 73]. This last remark in particular reminds us of the multitude of "frozen accidents" that have occurred during evolution: these frozen accidents were mainly caused by external conditions (external relative to a given system's laws of functioning). Thus, the task of reproducing evolution (i.e to synthesize artificial life) by purely self-generated boundary conditions seems hopeless, since at certain points in evolution, external causes have produced relevant changes.

Also of utmost importance is the fact that the complexity of an organism is often believed to reflect the complexity of its environment, at least to some extent. The now popular idea of enaction [Bourgine & Varela 92] is based on the statement that an organism and its environment are mutually defined. Even if one does not believe in complete mutual specification, it raises the issue of evaluating the influence of environmental structures on an organism's structures. How can one hope to do so without embodying artificial creatures in somewhat realistic, varying environments ? Emergence is not necessarily self-generated, it can also come from both internal and external interactions. Moreover, this has to do with the always implicitly assumed existence of useful/relevant simplifications, whereby one can transform a complex problem about the environment into "a simpler problem, the answer to which is usually a reliable guide to the answer to the complex problem" [Wimsatt 86]. This assumption can no longer hold since it implies a possible reduction of the system under study, which is not permitted in the case of complex systems.

- Function as a side effect of structure ?

The concept of function is a behavioral one. One can say that a given structure realizes a particular function if there is an observable effect on the environment or in the system when the structure is applied to or plunged into a given environment. Thus, a function is detected in association with the structure which implements it because of the modifications the structure induces in the 'behavior' of the environment (which in turn may influence the structure's behavior) when 'dropped' into this environment. This is where accidental causes are of great importance to understand the difficulty of predicting what function a given structure can implement. In effect, shapes are much more intrinsic to systems than functions, the latter being defined only behaviorally, i.e with respect to an environment and all the laws of this environment, which are accidental from the point of view of the structure. Functions usually depend on many more parameters than structures (we must make a distinction between a geometrical description of a structure, which may be simple, and a functional description which in turn is highly complex, much more than the functions it "implements"). Thus we externally recognize structures more easily than functions, but the importance of some structures comes from the functions they have been able to implement, and it is thanks to these functions that we have evolved detectors to recognize the corresponding structures. "Mind and world in short have evolved together, and in consequence are something of a mutual fit. (...) That is to say that our various ways of feeling and thinking have grown to be what they are because of their utility in shaping our reactions on the outer world" [James 1879].

The reductionist bias highlighted in the previous section leads to a focus of interest on the notion of structure while functional aspects are quite often neglected. Sometimes, these two aspects are confused with one another, because reductionists do not see the purpose of making separate studies on structure and on function, the latter being considered a side effect of the former (in effect, a function is just the consequence of plunging a structure into an environment, and since environments are not considered important, why should functions ?). But, this side effect can have dramatic consequences: "any adaptation has systematically specifiable conditions under which its employment will actually decrease the fitness of the organism" [Wimsatt 86]. Let us clarify the differences between structure and function (not to be taken litterally): a function F is specified by its effects on a given (finite) subset of environmental variables, and a structure S is functionally defined by its effects (when plunged into a given environment) with respect to all possible environmental variables in their whole ranges. That's the difference. More precisely, let (Xi) be all possible environmental variables (i can be a continuous indice, but that's not important to get the idea), and let $E=\{X_1, ..., X_n\}$ be a subset of these variables, acted upon by F: $F(X_1, ..., X_n)=(Y_1, ..., Y_n)$. For instance, F can represent the modification of the states of some variables in time: in continuous time $(dX_1/dt, ..., dX_n/dt)=F(X_1, ..., X_n)$, or in discrete time, $F(X1[t], ..., Xn[t])=(X1[t+1], ..., Xn[t+1])$. F is a function of these n variables. Then S is said to implement function F in environment {Xi} iff $S(X_1, ..., X_n, \text{and all other } X_i) = F(X_1, ..., X_n)$. What is usually assumed by reductionists is $\{ \text{all other } X_i\}=\text{constant}$, which is incorrectly derived into $\partial S/\partial X_i=0$ for $i\neq 1 ... n$. This can be true for some X_i and for some range of values, but this is generally false. What we can see here is that many different structures can implement the same function, and that the same structure can implement many different functions (even functions defined on the same set of variables) if different values of "irrelevant" variables are assumed. Moreover, the function implemented by S in a given environment depends on what variables we have chosen to look at: S can also have an effect on other variables, accidental from the point of view of the chosen variables.

This focus on emerging structure rather than on emerging function is a particular feature of synthetic reductionism as applied to the biological sciences. Any model of a natural system is partial and exists at a given level of description, therefore only some aspects of the system can be captured by the model -this is a common feature of all sciences having to deal with natural systems. But, even if AL is such a science, even if it had some successes in applying its methods, the very fact that its objects of study are "bio-logical" systems in a wide sense, makes it sensitive to the structure/function issue, which is not the case for other sciences of complex systems.

The notion of boundary conditions is now easier to understand as an analogy with the boundary conditions needed to specify the solution of e.g a partial differential equation: in order for a model of a structure to explain the function(s) implemented by this structure in a given environment, it may have to contain a lot of variables which do not seem directly relevant to the function(s), but are crucial for the function(s) to be realized. Some models contain enough specifications to effectively realize or explain the function(s) -they generate boundary conditions internally-, while some models may require additional constraints, or external boundary conditions. These conditions restricts the space of behaviors to the relevant ones.

Computational reductionism

It is true that AL as well as the sciences of complex systems have greatly benefited from the advances of computers in the last decades: these advances have enabled a "**time compression**" allowing for the simulation of processes that would otherwise have taken years and years. But this time compression is the only (though colossal) **known** advantage of computers. Thus, the questions are to know 1) - whether time compression is powerful enough to explore all possible behaviors (including interesting ones) of a formally defined system, and 2) - whether finite specifications can lead to open-endedness. Computational reductionism stands on the -true- idea that any phenomenon that obeys the **laws** of physics can be simulated on a computer (a sort of weak physical version of Church's thesis). Thus, classical reductionism conjectures the reducibility of any biological process to the laws of physics, while computational reductionism goes further by 'transitivity of reduction': any biological process can be simulated on a computer. This leads to three remarks:

1)- We are able to simulate some low-dimensional models of some aspects of the living, but it is not sure that we are going to be able to do more with a computer. This is a major problem: models scale very poorly, and toy models cannot give birth to well-formed theories. What if life can be 'explained' only by an immense-dimensional model, such that the number of relevant degrees of freedom itself is not even tractable and cannot be acted upon with present-day computers ? Practical computers are not Turing machines, and human programmers usually do not live more than a billion years.

2)- Also, the algorithmically-based notion of logical depth, although it may be of limited scope in this context (for instance it is always assumed that a finite number of operations take place at each step), showed us that some (very deep) objects can be simulated only by themselves, in the sense that there exists no shortcut to generating them. Thus, if evolution is depth-generating, it may be very hard to reproduce its latest products on a computer by using a synthetic procedure very similar to an "artificial evolution" (this is a philosophical objection which does not jeopardize such simulations - they will obviously teach us something- but which questions their scope).

3)- It must be clearly understood that arguments against computational-synthetic reductionism are NOT related to arguments against any kind of weak physical version of Church-Turing thesis. In effect, if it is obviously true that any process that obeys physical laws can be simulated (on a Turing machine, i.e provided enough space -memory- and time are available), nothing can be said about synthesis, because simulation and synthesis have two very different status. The laws of physics themselves have finite specifications, they are defined with respect to a set of chosen observables (properties of the object) which are transformed into variables to form a system together with relationships between them. It is thus trivial to express the above mentioned sentence about the simulation of processes that obey physical laws. But such processes can be simulated once the relevant laws have been found, i.e given one phenomenon, we can look for laws governing its behavior, and once this is done, the phenomenon can be simulated. Now, we synthesize some behavior with a computer, it this behavior will be doomed to obey the physical laws, possibly specific, which can be formally derived from the specifications of the system. It is now a completely unresolved question to know whether or not these derivable behaviors are open endedly diversified.

Close in spirit to these issues is the question of understanding the influence exerted by the medium of "implementation" through the boundary conditions it provides to the "simulated" process. These boundary conditions are difficult to deal with, because they may be essential to the implementation without being clearly taken into account in the model, or simply because they are hard to track down due to the high complexity of the medium.

Conclusion

In conclusion, we should first remember that while top-down approaches usually forget to obey lower-level constraints and laws, purely bottom-up approaches usually forget to look at higher-level constraints, and this leads in both cases to considerable flaws. Artificial Life, being 'very bottom-up', needs constraints. We propose that both empirical constraints originating from biology and other natural sciences, and pragmatic constraints oriented by the design of [useful] lifelike systems should be taken into account. Besides, external boundary conditions should not be ignored, because many phenomena in nature certainly occur with the help of "accidental causes". A way of taking external causes into account is by making embodiment a clear goal of all AL's theories and simulations. Embodying an artificial creature in some kind of environment (with the ultimate goal of plunging it into a real one) implies making a thorough investigation of the notion of external boundary conditions. AL may well constitute a first step toward this goal, in the sense that it is an attempt to delimit the power of self-generated boundary conditions and therefore to locate the frontiers beyond which it is the realm of accidental causes. Then, once internal boundary conditions are well understood, embodiment -that is taking into account environmental variables- will be allowed to begin.

Acknowledgements

Many thanks to Paul Bourgine, Chris Langton and Guy Theraulaz for many exciting debates and discussions, to Mark Bedau, Peter Cariani, Claus Emmeche, George Kampis, and Bill Wimsatt for fruitful intellectual exchanges.

References

[Bedau 92] - M.Bedau, "*Philosophical Aspects of Artificial Life*", Proceedings of ECAL 91, MIT Press, 1992

[Bonabeau 91] - E.Bonabeau, "*Flowers and beautiful women abound*", ENST Paris, 1991

[Bonabeau & Bourgine 92] - "*Artificial Life as it could be*", in Artificial Life , ed. C.Langton, Addison-Wesley, 1993.

[Bourgine & Varela 92] - P.Bourgine, F.Varela, "*Towards a practice of autonomous systems*", Proceedings of ECAL 91, MIT Press, 1992

[Cariani 89] - P.Cariani, "*On the design of devices with emergent semantic functions*", Ph.D dissertation, State University of New York at Binghamton, 1989

[Cariani 91] - P.Cariani, "*Adaptivity and emergence in Organisms and Devices*", in World Futures, vlo.31, pp.49-70, Gordon and Breach Science Publishers S.A, 1991

[Emmeche 92] - C.Emmeche, "*Life as an Abstract Phenomenon : Is AL Possible ?*", Proceedings of ECAL 1991, MIT Press 1992

[Elsasser 81] - W.Elsasser, "*Principles of a New Biological Theory: A Summary*", J.Theor.Biol., 89, 131-150, 1981

[Farmer & Belin 91] - J.D.Farmer, A.d'A.Belin, "*Artificial Life: The Coming Evolution* ", in Artificial Life II, Addison Wesley, 1991

[Forrest 90] - S.Forrest, "Emergent Computation", Physica D 42, North-Holland, 1990

[Grassberger 89] - P.Grassberger, "*Problems in quantifying self-generated complexity*", Helvetica Physica Acta, 62, 489-511, 1989

[Hillis 88] - D.Hillis, "*Intelligence as an Emergent Behavior; or, The Songs of Eden*", in The Artificial Intelligence Debate, S.R.Graubard ed., MIT Press, 1988

[Kampis 91] - G.Kampis, "*Emergent Computations, Life and Cognition*", in World Futures, vol.31, pp 33-48, Gordon and Breach Science Publishers S.A, 1991

[Küppers 91] - B.-O.Küppers, "*Information and the origin of life*", MIT Press, 1991

[Langton 86] - C.Langton, "*Studying Artificial Life with cellular automata*", Physica D, 22, North-Holland, 1986

[Langton 89] - C.Langton, ed., "*Artificial Life*", Addison-Wesley, 1989

[Langton 91] - C.Langton, C.Taylor, D.Farmer, S.Rasmussen, eds., "*Artificial Life II*", Addison-Wesley, 1991

[Lewes 1874] - G.-H.Lewes, in "*Emergence*", Dictionnaire de la Langue Philosophique, Foulquié.

[Nicolis & Prigogine 89] - G.Nicolis, I.Prigogine, "*Exploring Complexity: An Introduction*", R.Piper GmbH & Co. KG Verlag, 1989

[Pattee 1989] - H.Pattee, "*Simulations, realizations, and theories of life*", in 'Artificial Life', Addison-Wesley, 1989

[Polanyi 68] - M.Polanyi, "*Life's irreducible structure*", Science 160, 1308-1312, 1968

[Popper 83] - K.Popper, "*Der Materialismus überwindet sich selbst*", in Das Ich und sein Gehirn, Ed. K.Popper & J.C.Eccles, Münich, 1983

[Popper 87] - K.Popper, "*Natural selection and the emergence of mind*", in Evolutionary epistemology, rationality and the sociology of knowledge, Open Court, La Salle, 1987

[Putnam 73] - H.Putnam, "*Reductionism and the nature of psychology*", Cognition, 2, 131-146, 1973

[Putnam 88] - H.Putnam, "*Much Ado About Not Very Much*", in The Artificial Intelligence Debate, S.R.Graubard ed., MIT Press, 1988

[Rosen 78] - R.Rosen, "*Fundamentals of measurement and representation of natural systems*", North Holland, New York, 1978

[Rosen 85] - R.Rosen, "*Anticipatory Systems*", Pergamon Press, New York, 1985

[Sober 91] - E.Sober, "*Learning from functionalism - Prospects for strong AL*", "-Artificial Life II, Addison-Wesley, 1991

[Taylor 91] - R.Taylor, "*Fleshing out Artificial Life II*", in Artificial Life II, Addison Wesley, 1991

[Weaver 68] - W.Weaver, "*Science and Complexity*", Am. Scientist, 36, 536-544, 1968

[Weidlich 91] - W.Weidlich, "*Physics and Social science - the approach of synergetics*", Physics Reports, vol. 204, 1, 1 - 163, 1991.

[Wimsatt 86] - W.C. Wimsatt, "*Heuristics and the Study of Human Behavior*", Ch. 13 in W. Fiske & K. Scwheder eds., Chicago: University of Chicago Press, 1986.

[Weiss 70] - P.Weiss, "*Das lebende System: Ein Beispiel für den Schichtendeterminismus*", in Das Neue Menschenbild, Ed. A.Koestler & J.R.Smythies, Vienna, 1970