

Life: Organisms with a plan [DRAFT]

Thomas Miconi
School of Computer Science
University of Birmingham
Edgbaston B15 2TT
Birmingham, UK

December 12, 2007

Abstract

The problem of defining “Life” is one of the recurring riddles of biology. We first point out that at least two concepts are denoted by the word ‘life’ namely the ensemble of all living things (‘collective’ life) and the property of being alive (‘individual’ life). We note that a convincing definition for the former exists: the so-called “NASA working definition of life.” However, the property of being alive cannot be covered by this definition, or any other which relies on reproduction, as shown by the example of sterile individuals and the theoretical possibility of *ab nihilo* biosynthesis. We then discuss the concept of an *organism*, that is, a set of interacting parts such that each part is constantly being produced and shaped as a result of the activity of others (and their own). We suggest that living beings are *organisms endowed with a plan*, that is, organisms in which the mutual production process is governed by a pre-established plan rather than mere contingency. We justify the position that the presence of a plan is indeed fundamental to our intuitive concept of being alive. As a case study, we apply our conceptual framework to the Gaia hypothesis. We conclude that our definition is indeed operative and useful for gaining insight into the nature of living and non-living complex entities.

[Main limitations of the current draft: the term “organism” is now entrenched as synonymous to living being, so another term (denoting self-construction by via the interactions of differentiated parts) should be used; also the notion of “plan” may not be explicit enough.]

1 Introduction

What is the definition of life? I remember a conference of the scientific elite that sought to answer that question. Is an enzyme alive? Is a virus alive? Is a cell alive? After many hours of launching promising balloons that defined life in a sentence, followed by equally conclusive punctures of these balloons, a solution seemed at hand: “The ability to reproduce—that is the essential characteristic of life,” said one statesman of science. Everyone nodded in agreement that the essential of a life was the ability to reproduce, until one small voice was heard. “Then one rabbit is dead. Two rabbits—a male and female—are alive but either one alone is dead.” At that point, we all became convinced that although everyone knows what life is there is no simple definition of life.

D.E. Koshland, *The Seven Pillars of Life*[11].

At first sight, life seems obvious. For just about any object of the world, we can classify it almost without thinking as either ‘living’ or ‘non-living’. Oaks, spiders, people are alive. Cars, houses, stars are not. So-called ‘borderline’ cases, quite often, simply aren’t - for example, viruses and prions are almost universally regarded as non-living in their free form, because of their utter absence of autonomous activity [19]. Of course, like any other rule in biology, this one has its exceptions: the exceptionally large, perhaps metabolically active mimivirus may well be one. But the overwhelming majority of entities fall squarely on one or the other side of the not-so-fine line which separates life from non-life.

It could be thought that our intuitive concept of ‘life’ is a modern construct, a product of recent scientific knowledge. Interestingly, it turns out not to be so. The unambiguous dichotomy between living and non-living creatures was apparent to the Ancients just as it is to us. Aristotle acknowledges the presence of life (and of some form of soul) in vegetals and animals, in opposition to non-living minerals. Since then, as scientists discovered quantities of new objects through their microscopes and telescopes, those were quickly identified as either living or non-living (with occasional delays due to insufficient evidence).

Despite this apparent self-evidence, the precise concept of life is notoriously difficult to pinpoint. Somewhat surprisingly, there seems to be no generally accepted definition of life. This does not mean that definitions of life are a scarce resource - quite the contrary, they abound, they thrive, they proliferate throughout the scholarsphere. But for some reason none has imposed itself upon the others.

This difficulty in providing an universally accepted definition of life has led some authors to question the very feasibility of the task. Cairns-Smith [2] posited that “Life is a fuzzy idea - and it is best left like that” (a view perhaps not entirely independent from his own objective of propounding a gradual origin of life). Cleland and Chyba (2002) argue that life, being philosophically a “natural kind,” is currently as difficult to define as, say, water was before modern chemical theories allowed us to define it unambiguously as H_2O : we lack such a theory for life, and must therefore be content with vague clusters of notions involving various properties of life, just like pre-modern thinkers had to be content with defining water by reference to its sensible properties. Another possible outcome is that, as skeptics such as Sober suspect, the question of defining life turns out

to have simply no interesting answer. While such objections may turn out to be valid, the utter decidability of the concept of life warrants at least an attempt at capturing its fundamental characteristic(s).

1.1 Common approaches to defining of life

Numerous definitions of life have been suggested from classical times to the present day. It is well-known that many of them can be divided among two broad categories, namely those which focus on the notions of *reproduction* and *evolution*, and those which favour the idea of *self-production* and *self-construction*.

As many authors have noticed, this division is not accidental: it actually follows a common conceptual division in biology, homologous to that which exist between heredity and self-maintenance, between reproduction and metabolism, between genotype and phenotype, between “a collective population of self-reproducing hereditary systems (life as evolution) and as individual self-maintaining dissipative units (life as metabolism)” (Ruiz-Mirazo and colleagues[18], paraphrasing Maynard Smith). In fact, as we shall see, the division actually lies within the word “life” itself: this term can actually designate two different concepts, each pertaining to one side of the divide.

2 The meaning(s) of ‘life’

The first point which we want to make is that the word ‘life’ is heavily *polysemic*: it has many different meanings. In particular, the word ‘life’ may denote at least two related, but distinct concepts:

1. Life as the sum total of all living creatures, seen as one enormous self-propagating chemical reaction. In this sense ‘Life’ (often capitalised) is essentially equivalent to ‘the living world’. This sense is apparent in such expressions as “the diversity of Life,” “wildlife,” etc. We may call this ‘collective life’, or ‘life-the-phenomenon’.
2. Life as the *property* of being alive, applying to individual creatures, as opposed to non-living or dead. This meaning is apparent in expressions such as “between life and death,” or “the breath of life.” We may call this ‘individual life’, or ‘life-the-property’.

The first of these concepts describes a tangible entity, or rather a large number of such entities, united by a common origin and an overarching process of reproduction and evolution. The second concept describes a property, which applies to certain entities and not to others, in a way which is intuitive to most of us, yet not easily expressed.

Such polysemy is quite common, especially with basic words of everyday language: for example, consider the word ‘glass’ which may denote a substance or an utensil. Perhaps more fittingly, consider the word ‘water’: this word obviously denotes a substance, but may also, in certain contexts, denote a large body of this substance (Coniston Water, etc. are lakes in the Lake District of England). Clearly the substance is not the same *thing* as the lake, though the relationship between the two is obvious. Similarly the population of living creatures is not the same thing as the property of being alive.

In addition, it would not be satisfying to simply define the former as the mere aggregation of the latter. Life as a collective phenomenon has fundamental properties of its own which are important but not reducible to the properties of individual living creatures. Collective life is a bit more than just a bunch of living creatures put together, much in the same way that the Lincoln cathedral is a bit more than a set of contiguous stones located at a certain location in the centre of Lincoln.

The dangers of overlooking polysemy in definitions are obvious. For example, it would not be possible to produce one all-encompassing definition for “glass”, which would cover both the utensil and the substance. Efforts at producing such a general definition would necessarily fail, since they would attempt to cover two fundamentally different things with one single definition. The same problem applies to definitions of life: because the same word, “life”, can mean different things, an all-encompassing definition of this word makes little sense.

We believe that this problem has been the source of much unnecessary confusion: in effect, scientists who have tried to define “life”, without distinguishing the two meanings of the word, have put themselves in the position of philosophers attempting to define “glass” in general - an impossible proposition. By and large, two outcomes have occurred. On the one hand, some authors, having in mind one of these two concepts, have attempted to provide definitions which emphasised properties of this particular concept - only to be criticised by other authors who had the other concept in mind and (predictably) failed to recognise this other concept in the suggested definition. Of course, this outcome is manifest in the debate between those who emphasise reproduction and evolution (properties of collective life) and those who emphasise self-construction and self-regulation (hallmarks of individual life). On the other hand, some authors have attempted to capture the properties of both aspects of life into all-encompassing definitions, which could not but result in awkward constructions. Again, note the specific case of Ganti’s criteria, divided as they were into “real” life criteria and “potential” life criteria. We can now see that these two classes of criteria are actually both very real, but relate to two different things: the use of “real” and “potential” simply indicates an emphasis on individual life, that is, the property of being alive. The adjunction of the potential criteria reflects an attempt to include collective life into the definition, which is necessary for any attempt to cover “life” in general.

It turns out that we already possess a convincing definition for the first of these concepts, that is, for ‘collective’ Life: the so-called “NASA definition” (explained below), based both on self-sustenance and the capacity for Darwinian evolution. On the other hand the other concept associated with the word ‘life’, the property of being alive, has proven more elusive; however well-developed proposals exist, often centred on the self-construction.

3 “Collective” Life: reproduction and evolution

Reproduction, and its corrolary evolution, are often seen as characteristic features of living beings. Life thrives and multiplies everywhere it can. It also diversifies into a dazzling variety of forms, each subtly adapted to its own environment and lifestyle, sometimes through the effect of enormously complex adaptations.

This joint process of multiplication, diversification and adaptation is often seen as the defining feature of life, as the opening quotation of the present paper illustrates. Several definitions of life based on

This trend culminates in the recently proposed “NASA working definition of life” [3]:

“Life is a self-sustaining chemical system capable of Darwinian evolution.”

We believe that this definition is actually very appropriate for the concept of ‘collective’ life - life-the-phenomenon. The elegant formulation succinctly captures our intuitive notion of Life as a gigantic, granular chemical reaction, propagating and diversifying throughout the biosphere with remarkable success and awe-inspiring creativity.

However, definitions of life based on reproduction and evolution leave untouched the other concept known as ‘life’ - the property of being alive, the intangible difference between a living creature and a corpse. Clearly individuals don’t evolve. Many individuals cannot reproduce either - in fact, in sexual species, no isolated individual can reproduce at all (“one rabbit is dead; two rabbits are alive”). To push to the extreme, if reproduction is to be seen as the defining feature of life, then the philosopher Abelard produced some of his most enduring works after his “death.” Yet we would still call a single rabbit, or a mule, or an unfortunate philosopher, ‘alive.’

On the other hand there is a strong feeling that reproduction and evolution somehow play a role in our concept of being alive. To defend the importance of reproduction in the notion of life, it is sometimes observed that while all living creatures are not able to reproduce, all of them are the product of reproduction: all living creatures originate from another living creature. However this observation may be questioned: let us imagine a highly advanced civilisation where scientists would be able to create a living organism from scratch, without using any biotic material. This organism would not be the result of biological reproduction, but we would not call it any less alive for it. Thus we see that, while the notions of reproduction and evolution may well have some part in our concept of individual life, this cannot come from the physical act of reproduction itself.

4 The property of life

4.1 Autonomous growth

What does it mean, then, to be alive? What do we mean when we say of something that it has “life” in it? It turns out that this question also has an answer, although a more obscure one. In fact, we may turn back all the way to Aristotle, who disposes of the question with untypical conciseness and efficiency (On The Soul, Book II, chapt 1):

Of natural bodies some have life in them, others not; by life we mean self-nutrition and growth, with its correlative decay.

The capacity to absorb external matter and assimilate it, to turn non-self into self, is the basis of this notion of nutrition. It clearly corresponds to what we

now call metabolism. Importantly, the growth that Aristotle has in mind incurs much more than a mere increase in volume. After all, many phenomena incur the spontaneous assimilation of matter - fire and crystals being two obvious, oft-cited counter-examples. But as Aristotle pointed out, the growth of living beings is very different from that of fire, because it is regulated:

For while the growth of fire goes on without limit so long as there is a supply of fuel, in the case of all complex wholes formed in the course of nature there is a limit or ratio which determines their size and increase, and limit and ratio are marks of soul but not of fire. . .

Aristotle realised that biological growth is not a mere increase in volume: living systems grow in a *regulated* fashion, according to an inner directive force. Aristotle called this force the (vegetative) “soul”; we call it metabolism, self-organisation and genetic information. At any rate, this directive, regulating force is a fundamental difference with fire, which although it does grow, will happily eat up anything it can and grow itself into exhaustion as quickly as possible.

Furthermore, this regulation is not merely quantitative, far from it: it is actually formative, in that it moulds and shapes the growth of living beings with a creative power and a regularity which are simply inaccessible to non-living entities. In fact, it is on this basis that Aristotle rejected the possibility of a materialistic, objective biology, resulting from the blind interactions of brute matter, as put forward by Atomists such as Empedocles:

A difficulty presents itself: why should not nature work, not for the sake of something, nor because it is better so, but just as the sky rains, not in order to make the corn grow, but of necessity? What is drawn up must cool, and what has been cooled must become water and descend, the result of this being that the corn grows. Similarly if a man’s crop is spoiled on the threshing-floor, the rain did not fall for the sake of this-in order that the crop might be spoiled-but that result just followed. Why then should it not be the same with the parts in nature, e.g. that our teeth should come up of necessity-the front teeth sharp, fitted for tearing, the molars broad and useful for grinding down the food-since they did not arise for this end, but it was merely a coincident result; and so with all other parts in which we suppose that there is purpose? . . .

Yet it is impossible that this should be the true view. For teeth and all other natural things either invariably or normally come about in a given way; but of not one of the results of chance or spontaneity is this true. . . If then, it is agreed that things are either the result of coincidence or for an end, and these cannot be the result of coincidence or spontaneity, it follows that they must be for an end. . .

Living beings do not grow randomly and haphazardly; rather, they possess a certain well-defined *form*, which happens to be incredibly adapted to their lifestyle, and is highly unlikely to have arisen by ‘blind’ growth, solely guided by environmental and mechanical circumstances. Not only that, but this form is repeated identically in many individuals, proving beyond doubt that *something*

must be guiding this growth. For Aristotle, this ‘something’ is the “soul.” For us, the “soul” is basically biochemistry, and the direction that it follows have been arrived at by Darwinian evolution: we now know that neither of these two results (the self-construction of living beings and the obvious design apparent in their form) is beyond the capacities of a purely mechanical universe, and that in biology, final causes are indeed reducible to efficient causes. However, for our present purposes, the fundamental insight is not affected: living beings grow autonomously into complex forms, and this growth is not random and haphazard, but rather guided and directed.

4.2 The organism: mutual production and mutual causation

In everyday language, we often use the term “organism” as synonymous to living being. However, this is only a derived sense: originally, the term “organism” describes a certain property, that happens to be exhibited by all living beings - namely, being composed of mutually dependent, mutually producing parts.

Kant, in his *Critique of the Teleological Judgement*, devotes much thinking to the question of purpose and function in nature. In the course of his investigations, he attempts to capture the fundamental essence of living things, using the example of a tree:

In the first place, a tree generates another tree according to a known natural law. But the tree produced is of the same genus; and so it produces itself *generically*. On the one hand, as effect it is continually self-produced; on the other hand, as cause it continually produces itself, and so perpetuates itself generically.

Secondly, a tree produces itself as an *individual*. This kind of effect no doubt we call growth; but it is quite different from any increase according to mechanical laws, and is to be reckoned as generation, though under another name. The matter that the tree incorporates it previously works up into a specifically peculiar quality, which natural mechanism external to it cannot supply; and thus it develops itself by aid of a material which, as compounded, is its own product...

Thirdly, each part of a tree generates itself in such a way that *the maintenance of any one part depends reciprocally on the maintenance of the rest* [emphasis added]. A bud of one tree engrafted on the twig of another produces in the alien stock a plant of its own kind, and so also a scion engrafted on a foreign stem. Hence we may regard each twig or leaf of the same tree as merely engrafted or inoculated into it, and so as an independent tree attached to another and parasitically nourished by it. At the same time, while the leaves are products of the tree they also in turn give support to it; for the repeated defoliation of a tree kills it, and its growth thus depends on the action of the leaves upon the stem.

The first point indicates that living beings do not grow at random, but follow a certain form, associated with their genus; this echoes Aristotle’s insistence regarding the directedness of biological growth. The second part also

echoes Aristotle in pointing out the capacity of living beings for nutrition and growth, that is, for turning non-self into self - not just quantitatively, but also qualitatively.

The third point, however, is fundamentally novel. Instead of ascribing the self-production of living beings to a non-mechanistic “soul”, Kant notices that this self-production takes place in a very specific manner. a living being is a machine, an *organised* system, with a very specific type of organisation: it is composed of *parts which produce each other*. Each part in the system performs a certain function, upon which other parts (and the whole system) are dependent: remove the leaves from a tree, and the tree dies. But conversely, their appearance, growth and continued existence of each part is itself dependent on the activity of the rest of the system: a leave, removed from the tree, will quickly disappear. Thus, we see that living beings are composed of parts which recursively produce each other, and are both cause and effect of each other and of the whole. We summarise this by saying that living beings are *organised*, and call them *organisms*.

This remarkable aspect of living beings was already pointed out by Aristotle, but has been especially emphasised by Kant, who also used the term ‘organised’ for this concept: “An organised product of nature is one in which every part is reciprocally purpose, [end] and means.” (paragraph 66) (Critique of Judgement, Part II, Paragraph 64-66). As we shall explain in the following section, the concept of an organism, in and by itself, has important properties, such as the spontaneous emergence of purpose and function and the possibility for (some) self-regulation. The concept of organism has been fertile in representations, even to the modern day with Rosen’s notion of closure under efficient causation, or Kauffman’s autocatalytic sets.

4.3 Autopoiesis

The organismal view of life culminates in the much more formal concept of autopoiesis. Indeed, autopoiesis captures much more than just the organismal concept of mutual construction. First, this mutual construction must include a boundary between the organism and the outer world, which defines its physical unity. Second, the autopoietic machine must actively maintain its organisation, in a form of homeostasis which has its own configuration as a target:

An autopoietic machine is a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components that produces the components which: (i) through their interactions and transformations continuously regenerate and realize the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in the space in which they (the components) exist by specifying the topological domain of its realization as such a network. . . Therefore, an autopoietic machine is an homeostatic (or rather a relations-static) system which has its own organization (defining network of relations) as the fundamental variable that it maintains constant.

This definition leads to interesting questions. Can we really say that the parts of living beings “continuously regenerate and realize the network of processes. . . that produced them,” or (a fortiori) that they have their own “networks

of relations” as a fundamental invariant? In fact, in many cases, living beings seem to adopt pretty much the opposite behaviour: interacting organs may not only alter the graph of relationships, but even actively destroy each other in complex (yet remarkably well-ordered) processes of organismal alteration. A part may very well conspire in the destruction of the parts which gave rise to it. Change, sometimes in massive proportions, seems to be the rule more than the exception. Spectacular examples of discrete transformation abound (the metamorphosis of insects being a striking example), but many species show continuous, deliberate alteration of some or all of their parts. This culminates in situations where the whole organism may commit suicide altogether, or at least deliberately put itself at risk (as in apoptosis, or in some cases of altruism): in this case it is not just its own organisation that the system attacks, but its very existence, that is, the very fact of being organised at all. It is not easy to maintain that the defining feature of living creatures is that they somehow strive to maintain their own “networks of relations.”

Or should we understand that “their organisation” does not refer to their actual, current configuration, but precisely to the fact of being thus organised in a mutually producing manner - that what autopoietic machines homeostatically maintain is autopoietic organisation itself? However, the fact is that living beings do much more than just maintaining their own “organism-ness”. While living systems change and alter themselves over time, they are clearly neither amorphous blobs, nor chaotic shapeshifters. they possess a form which identifies them, and that they strive to reach. Also, the problem of suicidal or risky behaviours remains: when the neural impulses of a doe lead her to confront a predator in order to save her fawn, she is clearly violating the homeostatic maintenance of her organisation, in any sense. Should we call her any less ‘alive’ for it?

Additionally, autopoiesis explicitly does not consider any link between the organism and the reproductive-evolutionary process: it is an uncompromisingly organismal view of life. Some authors have expressed reservations about this absolute disconnection between the property of life and the concepts of reproduction and evolution [18].

4.4 An a minima concept of the organism

At any rate, for our present purposes, we do not need the full machinery of autopoietic discourse. Instead, we choose to use an *a minima* notion of organism, one which (as we will see) is clearly insufficient to capture the notion of ‘being a living creature’, but can nonetheless be used as a basis for it.

An *organism* is a self-enclosing set of interacting parts, such that the existence and activity of each part is dependent on the existence and activity of others.

This simple, minimal concept of an organism is deliberately similar to Kant’s or Aristotle’s, except for the ‘self-enclosing’ aspect which is clearly derived from autopoiesis.

Notice the use of passive voice and indirection. This is not an accident: the indirect wording of this definition is meant to avoid the suggestion of any functional or purposeful behaviour, that is, the idea that parts should somehow strive to ensure each other’s continued existence and maintenance in their

current shape. This is why we do not simply state that parts ‘produce’ each other: production might imply directed effort, which we want to avoid. Let us take an example to illustrate this point: it is certainly reasonable to say that a population of foxes is being constantly produced as a result of the activity of rabbits, as well as their own (imagine that rabbits stop eating and mating, and as a result all die off: clearly the production of foxes would greatly suffer). But it would not be quite as reasonable to say that the rabbits ‘produce’ the foxes. The reason is that the rabbits do not go out of their way to ensure the existence of the foxes - quite the contrary, we can objectively observe that their actions (such as running away from foxes) are actually in opposition to the well-being and multiplication of foxes. This seems to be incompatible with the common idea of ‘producing’ something. In living beings, we often see parts acting for the good of each other; but, as explained above, we also see the exact opposite, namely, mutual destruction. The above definition is meant to accommodate for such behaviours. If one part stops functioning, or even if it is destroyed altogether by the action of other parts, then so be it: as long as other parts maintain a mutually productive configuration, the whole is still an organism.

4.5 Spontaneously occurring organisms

‘Organisms’ in this minimal sense need not be the result of a priori design: they can occur spontaneously. Let us consider a system in which, initially, a certain type of element grows spontaneously, extracting matter and energy from their environment. Now let us imagine that a new type of element appears, growing by feeding on the previous element. Let us imagine further that many such types of entities appear, or disappear, each being somehow dependent on other entities, extracting matter and energy either from the outer environment or from other elements within the system. In the course of this process many new entities appear, while others disappear, and most of them change over time, constantly creating new relationships and dependencies.

It is clear that we have just described a very simple kind of abstract ecosystem. Furthermore, a system of this kind may well be self-enclosing: it suffices that one or more of the entities thus produced somehow encircle or contain the entire system (in the real world, we have the examples of atolls or oases, or even simply forests with a well-defined border). Such a mutually dependent, self-enclosed system fulfills the definition of an organism: the maintenance of each part is dependent on the continued existence and activity of others. Thus, an ecosystem of this kind does qualify as an ‘organism’ in our minimal sense. This property of ecosystems may give rise to complex dynamics. However it is immediately obvious that such a system would exhibit important differences with living beings, as we shall explain below.

4.6 Spontaneous emergence of purpose and function in organisms

In an organism, each element is both cause and effect of each other; so are the whole and each individual part. This immediately creates a notion of purpose, or *function*, in the individual behaviour of these parts. If an element simply produces another, then all we have is merely a straightforward production process. But if the activity of each part is essential for the continued existence

of other parts (and of the whole), then this activity becomes, from this very fact, a function. The dependency of the whole upon the proper behaviour of individual elements induces a natural notion of functionality and purpose. Kant has brought attention to this, noting that “in such a product of nature every part not only exists by means of the other parts, but is thought as existing *for the sake* of the others and the whole, that is as an (organic) instrument.” A set of mutual efficient causes justifies a discourse based on final causes, escaping the Panglossian trap¹. Kauffman makes the same point in the context of autocatalytic systems, with typical eloquence [9] (p. 388):

Part of the interest in models of autocatalytic polymer system. . . lies in the fact that such systems afford a crystalline founding example of functional wholeness, hence functional integration. Given the underlying model of chemical interactions, once an autocatalytic set of polymers emerges, it is a coherent whole by virtue of achieving catalytic closure. Given the underlying model chemistry and catalytic closure, the functional role of each polymer or monomer in the continued existence and proliferation of the autocatalytic set is clear. Note that we here feel impelled, almost required, to begin to use functional language. This requirement reflects the fact that such a self-reproducing system allows a natural definition of the “purpose” of any polymer part, a purpose which is subservient to the overarching purpose of the autocatalytic set, which is of course, abetted by natural selection, to persist and prevail.

In short, an organism lends itself to analysis in terms of *physiology*: the study of the role and function of various parts within an organised system.

However it is important to note that this spontaneous emergence of function is dependent on certain conditions, perhaps most importantly on the stability of the overall organisation. If the mutual production process is too unstable, if parts emerge, grow and disappear arbitrarily, without any noticeable pattern, then no part can ever be seen to fulfill any particular function, and the impression of purpose is lost. In fact, it appears that purpose and functionality do not arise from the mutual production process, but from the mutual *causation* aspect: only when a reasonably stable set of mutual effective causes emerges (which presumes that the parts last long enough than some of them can actually be seen as both cause and consequence of each other’s activity), can we begin to ascribe function and purposes to each particular activity.

4.7 Intrinsic properties of organisms

Due to their constant self-construction, organisms automatically exhibit certain important properties. The most obvious one is a capacity for spontaneous self-repair in the face of moderate quantitative damage. If one part has somehow lost some of its material, the constant activity of other parts may regenerate the damaged part. Of course this may not happen if the damage is too important (to the point of compromising the whole organisation) or if the part being damaged

¹In the first chapter of Voltaire’s *Candide*, Doctor Pangloss produces several instances of dubious reasoning based on final causes, such as: “Observe, for instance, the nose is formed for spectacles, therefore we wear spectacles”

is itself highly organised and thus cannot be easily repaired. For example, in humans, the loss of a portion of liver, or of a fingertip (above the nail line) is easily repaired; the loss of a portion of skin may be repaired with varying success depending on the location and extent of the damage; and the loss of a limb can usually not be repaired at all.

Additionally, under certain conditions, the mutual dependency of construction processes may spontaneously give rise to basic forms of self-regulation, that is, to the maintenance of certain quantitative variables. This is especially the case when the mutual dependency cycle between various parts takes the form of direct mutual consumption: a food web in which elements are being consumed by each other. This gives rise to negative feedback relationships of the form “the more there is of me, the less there is of you,” which are inherently stabilising in the face of fluctuation. Paraphrasing Kauffman’s “order for free” concept, we might call this “self-regulation for free.”

However, this spontaneous self-regulation is merely a possibility, pertaining to certain types of dependency; it does not mechanically follow from the notion of an organism. In particular, it may not apply when the relationships are of a different type. For example, the mutual dependency may take the form of a producer-consumer relationship in which elements produce compounds which are consumed by others in their own formation and activity. In this case, if one component happened to produce more than its usual level, this would not mechanically create a negative, “corrective” force on its level of production.

Furthermore, these spontaneous instances of self-regulation can only be of a very basic form: while they may regulate small quantitative variations, they may not be able to deal with qualitative alterations or other, more serious conditions. This does not mean that organisms are fundamentally unable to deal with such problems - merely that being an organism, in and by itself, does not warrant any such high-level self-preservation mechanisms. Those may be present, or not. Additionally, any spontaneously occurring feedback mechanism may just as well take the form of a *positive* feedback, which hamper regulation and control rather than assisting them.

4.8 “Organisms” do not cover “living beings”

An organism simply describes a certain type of relationships within a set of elements. This notion is clearly important in and by itself: organisms are likely to exhibit interesting properties, independently of other considerations. Furthermore, the notion of an organism (and the properties it entails) is likely to play an important role in our vision of life. Can we say, however, that this concept is sufficient to capture the intuitive notion of being “alive”? It does not seem to be so.

Let us consider again the abstract ecosystem which we introduced in section 4.5. This self-enclosing system of continually co-evolving entities does fulfill the requirements of an organism. On the other hand, if left to itself, the structure of this system would be quite instable over large timescales: various types of entities would appear and disappear without any pattern, in an apparently chaotic manner, guided only by the whim of opportunistic evolution within the constraints posed by other entities and the outer environment. No general structure would be able to endure for any significant period of time. The continual turnover, not just of the individual entities themselves, but also of the types of

these entities, would prevent the appearance of any stable organisation.

In fact, it is probable that we would sooner or later observe the emergence of ‘selfish’ or ‘polluting’ entities, which would somehow grow at the expense of other entities (either by depleting resources, or by releasing toxic byproducts). These would correspond to cancers in multicellular organisms. The noxious effects of such entities might cause their own demise in the end, but they could severely damage and impoverish the entire ecosystem in the meantime, perhaps even beyond recovery.

4.9 From organisms to living beings: the notion of plan

Living organisms are not quite so chaotic. Somehow, they grow in a well-ordered manner. A young tree is not going to grow into a random patchwork of mutually constructing tissues - it is going to grow into a mature tree. A larva will grow into an insect, and a human infant will grow into a human adult. We feel that there is a very strong amount of non-randomness in the way living beings grow, which is not incidental, but rather fundamental to their very nature. Living beings are endowed with an inherent body-plan, an *imago*, which means that they grow into a certain pre-defined form rather than into some unpredictable assemblage of randomly mutating parts. Although their form changes over time, sometimes in a spectacular manner (as in insect metamorphosis), this change is not chaotic to the point of blurring any notion of body-plan altogether: a larva is different from an insect, but both are well-defined forms which are being actively built and maintained by the living being at any time. More importantly, the change itself is not an effect of pure contingency, but instead occurs in a well-ordered manner as part of an overarching, pre-defined mechanism. The change from one body-plan to the other is itself part of a plan.

This non-randomness is not only seen in the overall form of living beings. We may actually find its traces by observing certain specific features, which reveal the existence of a priori design. One obvious tell-tale of this ordered behaviour is the presence of repeated parts: if the elements of an organism randomly grew out of their local conditions, we would not expect to see almost exactly similar, complex parts popping up in different places - the probability of just one such replication would be astronomically small, and it is estimated that a mature tree has about 200000 leaves. Additionally, the growth of living organisms exhibits uncanny *prescience*. This is particularly striking in the existence of parts which apparently do not benefit the whole in any way in day-to-day existence - until a certain specific situation occurs, in which case these parts become vital to the whole, in that their specific features (which, importantly, would *not* be expected to have occurred by chance) prevent the demise of the whole by specifically addressing this situation (again, with an efficiency that could not be the result of plain luck). Platelet cells are a typical example.

This prescience is also seen in the development itself, especially in its high degree of synchronicity: different parts grow simultaneously to produce a harmonious whole. Note that this prescience is quite different from mere mutual adaptation. Even in a randomly-growing organism such as an ecosystem, we may expect that at any given time, existing parts will be reasonably well-adapted to each other - otherwise they would not have arisen, or would have been

wiped out². But in living organisms, this adaptation actually occurs *a priori*: the parts exhibit features or behaviours which benefit the whole in conditions which have not yet been encountered. Also we observe that the majority of feedback processes are regulative and homeostatic in nature, which again would not be expected by chance. When positive feedbacks do occur (as in the case of immunatory response), they are almost always beneficial to the maintenance of the whole, rather than destabilising.

Similarly, living organisms are striking by the lack of *selfishness* of their various parts: in general, parts do not grow selfishly at the ruthless expense of others, but rather maintain a very urbane co-existence. As Aristotle pointed out (see section 4.1), they restrict their own growth to fit a certain well-defined shape and organisation. This unselfishness is actually so strong as to result in generalised *altruism*: parts produce certain elements, at a cost to themselves (in terms of growth), which benefit other parts. Indeed, in some cases, parts actively destroy themselves for the good of the whole. When selfishness occurs in living beings, in the shape of ‘rogue’ growing parts (for example, in multicellular organisms, as cancers), the striking difference between the behaviour of those selfish elements and other, normal parts highlights the extreme ‘politeness’ of the latter.

Unguided mutual construction, on its own, is insufficient to explain these remarkable features which we deeply associate with the living condition. These aspects of life are immediately explained, however, by considering a universal feature of living creatures: the fact that they are all endowed with a *plan*.

5 Living beings: organisms endowed with a plan

5.1 The plan as a fundamental feature of life

We suggest that living beings are essentially *organisms endowed with a plan*, that is, self-enclosed sets of interacting parts which constantly produce and shape each other through their activities, such that the behaviour of each part (instead of being left to local contingency) follows a well-defined, pre-established plan and results in a well-defined, pre-established general form - their imago. This imago can change over time (in contrast to autopoietic machines), but never so chaotically as to vanish into indefiniteness, and most importantly, this change itself is specified by the plan.

That the existence of a plan is a defining feature of living beings was already suspected by Aristotle, as we mentioned earlier. Although the passage quoted in section 4.1 only mentions quantitative self-regulation, the entire book from which it is taken (as well as significant portions of ‘On the soul’) is devoted to the examination of qualitative aspects of this planning. It is important to note that Aristotle *did* consider the possibility of spontaneous growth generated by purely local material causes, as put forward by the Atomists - and then rejected it in the face of overwhelming contrary evidence (see part 8 of book II of Physics). The regularity of repeated parts, the uncanny prescience exhibited by living processes, betray the existence of an overarching plan because they cannot be

²Thus, perhaps counter-intuitively, mutual adaptation among parts is not, in and by itself, a sign of a priori design.

explained by “chance”³.

As Aristotle explains, fire has no plan, and as a result will gladly consume itself into exhaustion. Living beings regulate their growth, not only with regard to size, but also with regard to form and substance, according to a well-defined architecture. Rather than being entirely dependent on external conditions and contingency to shape their overall forms, they actually impose their own form (up to a point) to the external environment. We suggest a metaphor borrowed from linguistics: living beings are, in this sense, *strong*, very much in the manner of a ‘strong verb’ which is able to impose its own specific form to the conjugation system rather than being shaped by it (weak verbs, such as ‘love’ and ‘talk’, become ‘loved’ and ‘talked’, following an external pattern proper to the tense; by contrast, strong verbs such as ‘grow’ or ‘sing’ become ‘grew’ or ‘sang’, imposing their own alternate root to the tense in question).

Some authors have gone so far as to posit that planned autonomous construction was, in itself, the essence of life. But such processes also occur in non-living beings, such as crystals. The growth of crystals certainly does follow a well-defined plan. Furthermore, this growth can even result in ‘reproduction’ when the crystal breaks into two or more parts, which keep growing independently. Monod, in particular, found no other difference between living beings and crystals than the sheer volume of information necessary to describe (and enforce) their overall organisation[17]⁴.

However, from our viewpoint, there is an important difference between crystals and living beings, namely the fact that crystals are not organisms. We see no process of continual mutual production and maintenance in crystals. Crystalline growth occurs as a straightforward production process (by catalysed accretion), rather than mutual production. Particles are included in the growth layer of the crystal, catalyse the inclusion of other particles, and then are sequestered within the crystal where they remain, indefinitely inert - effectively dead. Indeed, if we draw the graph of interactions within (or rather ‘on the top layer of’) a crystal, we will see no cycle. This absence of loops in the dependency graph of crystals demonstrates their non-organismal nature. This shows that both the plan *and* the organismal process are necessary to capture the essence of living beings.

5.2 Planning permits a priori design

The most important consequence of the existence of a plan is that it opens the possibility of *design*. Instead of being subject to the vagaries and uncertainties of purely opportunistic production processes, organisms endowed with a plan have the possibility (depending on the contents of the plan) of constraining these processes in arbitrary ways. These constraints are fundamentally different from the purely local constraints imparted by contingency and the laws of physics. This is because, through the medium of the plan, the self-construction process

³For Aristotle, both the process of self-construction and the planning of this process are imparted by the “soul.” For us, the “soul” is basically biochemistry, and the contents of the plan have been arrived at by evolution: we now know that neither of these two processes is beyond the capacities of a purely mechanical universe, and that in biology, final causes are indeed reducible to efficient causes.

⁴We may also note that the plan of most crystals cannot be modified, and thus is not really a ‘plan’ in our sense, since it cannot be a basis for design. However, some authors have argued that certain crystalline structures can indeed evolve in a certain sense[2].

can be influenced by information *which it has not yet encountered itself*. The consequences of this simple fact are tremendous.

5.3 Design permits arbitrary self-regulation

We saw that organisms, by their very nature, could sometimes exhibit a certain amount of basic self-regulation: for example, certain types of mutual dependencies automatically produce negative feedback relations. However, these fortuitous feedback processes are necessarily limited in their regulative power, and may just as readily take the form of destabilising positive feedbacks. By contrast, in living beings, the pre-defined plan permits *arbitrary* self-regulation to take place. This includes self-regulation and self-maintenance processes which cannot emerge spontaneously, but must have been put in place through a design process. Immune systems, cancer-fighting mechanisms, endocrine systems, inflammatory responses, all these ‘prescient’ processes (which only help the whole in situations which have not yet been encountered at the time of their initial production) are made possible by the plan.

Again, we stress that these additional capacities are not just related to defence against external aggression. They include the capacity to protect the whole against potentially devastating instabilities of unguided mutual production. Organisms endowed with a plan are no longer utterly defenseless against runaway positive feedbacks, rogue (‘selfish’) processes, etc. Cancer-fighting mechanisms can be put in place. Hormones and glands can put the production of various parts in check. Design offers some protection, not just against the external environment, but also against the fundamentally unstable nature of self-constructing processes when left to themselves.

5.4 On the meanings of the word ‘plan’

Ambiguity pervades human languages, and the word ‘plan’ is no exception. The Oxford English Dictionary tells us that this word can be understood in several ways, including the following two:

1. “An organized (and usually detailed) proposal according to which something is to be done; a scheme of action; a strategy; a programme, schedule.”
2. “A design according to which elements of something are arranged; a scheme of organization; a configuration, arrangement, or type of structure (esp. regarded as the product of design).”

In other words, a ‘plan’ can specify either a *procedure*, or an *architecture*. Importantly, in our discussion, both of these senses are important, though is the latter (the a priori specification of a general configuration) is in a sense primary. In their process of self-construction, living beings strive to reach a certain overall form (an imago), and in order to reach this particular form, they need to follow a certain pre-defined procedure. The imago being built may change over time, but this change in configuration is well-ordered rather than chaotic, and at any rate is itself specified and enforced by the procedure. The change of body-plan is itself part of the plan.

In this regard it may be useful to quote one of the examples chosen by the Dictionary to illustrate the second sense of the word: “A phylum is a group of creatures which share a distinct body plan.” The body-plan of phyla exists at a high level of organisation; however, each species has its own plan, which is precisely what makes us call it a ‘species’ and allows us to identify it as such. A young oak is not going to grow into a rose or a cauliflower - it is going to grow into an oak. A dog grows into a dog, and a man into a man. The existence of an imago (a body plan) is a fundamental component of our concept of living beings.

As it turns out, in earthly living beings, we observe that the physical support of the plan (DNA) actually encodes a sequence of instructions, that is, a procedure (“in these or those conditions, synthesise this or that protein”). However, when considering examples of “life as it could be,” this is not a logical necessity. In fact it is not even necessary that the plan should physically be embedded in a separate component at all. This does not mean that the question has no importance, far from it: for example, it has often been pointed out that reproduction using physically embedded instructions has different properties than reproduction by self-inspection[16].

5.5 The plan as the link between collective and individual life

The fact that living creatures are universally endowed with a plan offers a natural link between the individual property of being alive, and the visible aspect of collective life: the concept of a plan naturally suggests the concepts of reproduction and evolution. A plan can be instantiated many times, which implies reproduction. A plan can be altered and modified to suit certain criteria, which suggests evolution.

Notice the symmetry between ‘collective’ life (as per the NASA definition) and individual life (as per ours). The NASA definition of life requires that the chemical system be “self-sustaining” (a notion associated with self-construction), while the definition of individual life requires a “plan” (a notion associated with reproduction and evolution). Thus, both of these definitions include requirements which are supplied by properties of the other: self-sustenance is ensured by the self-construction of individually living creatures, while the plan is automatically present in a reproduction process.

This symmetry shows that, although ‘collective’ life and ‘individual’ life do denote different concepts, those two concepts are nonetheless deeply intertwined: one naturally calls for the other. It is indeed conceivable that one might somehow be brought to existence without the other⁵; but we can hardly think of any way in which either could arise without the other *spontaneously* (that is, without human intervention). On the other hand, their properties support each other’s emergence in conjunction: self-constructing entities endowed with a plan are most likely to emerge out of some early reproductive

⁵Let us recall the example of scientists creating a cell from scratch, in isolation from any reproductive-evolutionary process. Conversely, we may one day build robots which could ‘reproduce’ by self-inspecting and producing copies of themselves by extracting, refining and assembling material from the environment; such robots could be called a self-sustaining chemical system capable of Darwinian evolution, but they would not qualify as organisms due to the absence of self-construction.

process, initially kept going by external factors (such a process is precisely suggested by Fernando and Rowe [6]); conversely, as soon as these entities become truly self-constructing, and therefore self-sustaining, so does the reproductive-evolutionary process. Collective and individual life, in their early intermediate stages, stand out as natural ‘scaffolds’ for each other.

As a consequence, the fact that, on Earth, they have both emerged from one single chemical reaction (being, in a sense, the large scale and the small-scale view of it) may not be a mere accident, but rather an indication of their mutual suitability as a support for each other. If this interpretation is correct, it is likely that if life exists elsewhere in the universe, it will also occur under the form of one enormous, ‘granular’ chemical reaction (collective life) generating, and composed of, many organisms endowed with a plan (individual life). The unity of collective and individual life, more than a statement of observed chemical fact, may well be an expected consequence of their nature.

5.6 Is the plan decidable?

It may be objected that the notion of a plan is very much an observational property: the plan is in the eye of the beholder. The individual part doesn’t ‘know’ that it is following a plan: from its own ‘viewpoint’, it is merely reacting to constantly changing local conditions and stimuli. What the plan does is precisely to ensure that these conditions change in a very specific manner, so that the behaviour of the part is constrained at the service of the whole, but the individual part cannot ‘see’ this. If we adopt the organ’s eye-view, there is no possibility to distinguish between purely reactive and designed growth.

However, we are not limited to the organ’s eye-view: we can, and should, consider the whole organism. When this is done, several objective factors may be assessed in order to detect the presence of a plan. The features discussed in section 4.9 are the obvious candidates: within an organism, repetition of similar parts, obvious cases of prescience, and (perhaps most importantly) unselfishness, are strong indications that a certain pre-established plan drives the process. Of course, other indicators exist. For example, the presence of many similar organisms, such that their similarity cannot be explained merely by chance or by environmental influences, would suggest the existence of a generic plan in and by itself.

[More !]

6 A case study: Gaia

6.1 The Gaia hypothesis: is the Earth “alive” ?

[Fill up the introduction]

The Gaia hypothesis, initially proposed in the late seventies by Lovelock [13], states that the biosphere (living beings and their immediate geological environment) behaves in certain respects and to a certain extent like a living organism. Just in which respects, and to which extent, has been the subject of intense debate, through which the original hypothesis has been thoroughly refined. Let us consider the following formulation, taken from the most recent treatment of the hypothesis by its originator [14]:

Gaia is a thin spherical shell of matter that surrounds the incandescent interior; it begins where the crustal rocks meet the magma of the Earth's hot interior, about 100 miles below the surface, and proceeds another 100 miles outwards through the ocean and air to the even hotter thermosphere at the edge of space. It includes the biosphere and is a dynamic physiological system that has kept our planet fit for life for over three billion years. I call Gaia a physiological system because it appears to have the unconscious goal of regulating the climate and the chemistry at a comfortable state for life. Its goals are not set points but adjustable for whatever is the current environment and adaptable to whatever form of life it carries.

The central tenet of the Gaia hypothesis is that living organisms collectively create large-scale feedback mechanisms, which maintain certain quantities (temperature, gas concentrations, oceanic salinity, etc.) within tight bounds. Most controversially, certain interpretations of the Gaia hypothesis posit that some of these feedback systems are being altruistically maintained by living organisms: certain groups generate a certain product, at a cost to themselves, for the benefit of others. Can we use our conceptual framework about the nature of life to gain insight into the matter?

6.2 Gaia is an 'organism'

First of all, Gaia is clearly an 'organism' according to our own limited definition. Living organisms and their immediate environment (atmospheric gases, minerals, etc.) are indeed caught up in an enormous dependency cycle, in which the existence and maintenance of each individual entity is dependent on the maintenance of other entities within the cycle. This is a natural consequence from the fact that living beings disturb their environment: by so doing, they create new challenges and new opportunities for other living beings. Life constantly evolves to fit the conditions that it creates. In addition to this, some of these conditions (such as the presence of oxygen in the atmosphere) have been maintained with notable consistence over time. While the non-living, mineral environment is clearly not evolving in a Darwinian sense (which requires multiplication, heredity and selection), it is very much part of the 'organismic' cycle of mutual dependency and production, and is involved in the patterns which emerge within it. Examples include the quantities of oxygen and carbon dioxide in the atmosphere, which are maintained at reasonably stable values (for sufficiently short geological timescales) in the face of massive biological turnover⁶.

Since Gaia is an organism, it is not surprising that feedbacks between various entities should spontaneously emerge. This is certainly evident between living creatures: the well-known Lotka-Volterra cycle is an obvious example of a direct negative feedback. Similar negative feedbacks also involve mineral components.

⁶This 'organismal' view of Gaia is supported by Lynn Margulis, co-founder of the Gaia hypothesis: "Selection pressures, the insistence of all organisms that they grow and reproduce, favor certain types of life under certain specified conditions. These grow, expand, remove waste, and recycle. As they do, they place enormous selection pressures on still different types of life. The result is Gaia."

For example, several cycles involving carbon dioxide and temperature are known to produce negative feedbacks (such as the silicate weathering ‘thermostat’ [1]). These cycles are thought to be important factors in climate regulation.

Additionally, the notions of purpose and function, which as we have seen spontaneously emerge in (sufficiently stable) organisms, have long been recognised in natural ecosystems: we talk routinely of ‘primary producers’, ‘decomposers’, ‘top predators’, and other ecological roles associated with certain functions within an ecosystem. The Gaia hypothesis suggests that we should extend this recognition to the global ecosystem, and include mineral components within the picture. This is not unreasonable, considering that certain biotic groups do play a global role (the production of oxygen by plants and algae, the sequestration of carbon by plants, etc.). In this limited sense, and with due precautions, it is not exaggerated to regard Gaia as a “pysiological system.”

6.3 Gaia is not a living being

However, while Gaia is an organism in our ‘weak’ sense, it is clearly not a living being according to our definition. The reason is that the global ecosystem does not seem to follow any pre-established plan. Rather, it develops spontaneously as the result of the constant, autonomous evolution of lineages. This simple fact happens to have important implications regarding the limits of the Gaia hypothesis.

6.3.1 Fragility of regulation

The first consequence of this lack of plan is rather obvious: if one lineage suddenly decides to make a living at the expense of all others, nothing will prevent it from doing so. The most obvious example is the massive release of free oxygen in the atmosphere by early photosynthesisers. Living organisms at the time were not prepared to deal with highly reactive oxygen, and as a result many disappeared or were banished to the anoxic fringes of the environment. Margulis [15] has called this event the “Oxygen Holocaust,” a rather unambiguous term. Thankfully, the remarkable adaptivity of bacteria allowed them to overcome this crisis and turn the previously noxious oxygen to their advantage, harnessing its chemical power to their profit: this was the birth of modern respiration. However, it is not easy to see how this process could be described as an adjustment of the geochemical environment “to the form of life it carries.” More recently, another group has begun to grow at the expense of other lifeforms with devastating efficiency, to the point of starting not only global climactic changes but also a mass extinction event. This group, of course, is *Homo Sapiens*, the Great Extincter. More generally, the point is that the growth of entities within the Gaian dependency cycle is manifestly not subordinated to the overall well-being of Gaia itself, for any reasonable definition of such ‘well-being’. Parts of the Gaian organism can, and will, grow at the expense of others if they find a way to do it. Gaia, having no plan, is devoid of the targeted mechanisms which could keep inner aggressions in check.

Another obvious consequence of this lack of plan is visible in the nature of the many feedback mechanisms which affect the Gaian environment. As we have seen, negative feedbacks exist, as could be expected in any set of interdependent growing parts. However, Gaia seems to contain just as many *positive* feedbacks

as negative ones. Again, this is what would be expected from a spontaneous, unguided organism, but stands in contrast with the remarkably coordinated homeostatic mechanisms of living beings. In fact, some of these feedbacks may lead to an extreme *impoverishment* of the biosphere, sometimes of absolutely cataclysmic proportions. One striking example is the occurrence of runaway glaciations (the “Snowball Earth” events [8]) : growing ice caps reflect ever more solar radiation, leading to a further fall in temperature and increased ice formation. Eventually continental ice caps reach the equator and the Earth is covered in ice. Unsurprisingly such conditions have dramatic consequences for life⁷. More generally, Kirchner [10] reviews evidence from the literature which tends to indicate that the overall effect of biological feedbacks is actually destabilising: imbalances in quantities of interest (carbon dioxide, methane, sulfur compounds, etc.) seem to be mostly aggravated, rather than mostly corrected, by biological response. At any rate, the important point is that Gaia is not ruled by a complex network of highly efficient, biologically-enforced, negative feedbacks - it is not *homeostatic*: biogeochemical feedbacks may be negative or positive, as would be expected from a spontaneous, unplanned organism. This is in contrast with the exquisitely coordinated regulation mechanisms which exist in living organisms, which are a benefit of a priori design⁸.

6.3.2 Altruism

More generally, the stronger interpretations of the Gaia hypothesis are imbued with the idea that the behaviour of individual entities is somehow curtailed in favour of the whole: entities supposedly perform certain task, at their own cost, for the benefit of all, and subsequently reap an indirect reward from the enhanced health of the Gaian community. This idea is perhaps the crucial factor which distinguishes the widely accepted versions of the hypothesis from the more controversial ones.

Let us consider one recent example: it has been suggested ([14], p. 19) that the excretion of urea by animals is actually a costly process, wasting both water and the energy contained in the expanded urea. Supposedly, breaking down urea into easily disposable nitrogen, carbon dioxide and water would be more efficient. Why, then, do we not do it? The proposed reason is that urea is a fertiliser: it favours the growth of certain plants (usually after being degraded back into ammonia by bacterial symbionts of these plants). More plants means more food for animals. As a result, urea production is indirectly beneficial to animals despite its immediate costs, and (in the author’s opinion) should therefore be favoured by evolution.

Unfortunately this ‘altruistic’ explanation flies in the face of modern evolutionary theory. If, as the scenario suggests, degrading urea into nitrogen and

⁷Interestingly, biological causes have been suggested for such Snowball Earth events[12, 8], though the more usual explanations rely on geological factors such as the presence of unified supercontinents.

⁸Such complex feedback mechanisms were actually purported in early expositions of the Gaia hypothesis. In one example [13], a highly engineered mechanism for oxygen control was suggested, involving (among others) two fine-grained, mutually counter-balancing feedback systems based on methane and nitrous oxide. The author argued (p. 76): “We may be sure that the efficient biosphere is unlikely to squander the energy required in making this odd gas [nitrous oxide] unless it has some useful function.” These are precisely the type of mechanisms (and the type of explanations) which are highly unlikely to emerge without a priori design, which implies the presence of a plan.

carbon dioxide is both easy and beneficial (if it is not, we need not wonder why it failed to evolve), then a population entirely composed of altruistic urea excreters will invariably be invaded by ‘selfish’ urea degraders. The reason is that the advantages of this selfishness benefit only the selfish individuals, while its costs (reduced plant growth) are equally borne by all. Even in ideal conditions, the process would at the very least produce a certain balance of ‘selfish’ urea degraders and ‘altruistic’ urea excreters.

As it turns out, we do not see any such thing in the world. Rather, we see that the excretion of urea, or uric acid (which is more costly, but conserves water) is universal among land-dwelling and flying vertebrates. This universality is simply not compatible with the idea that urea degradation is both easy to obtain and selectively advantageous. Rather, the most plausible position is that urea degradation is either difficult to evolve, or is just not quite as beneficial as the author suggests. Terrestrial animals excrete urea or uric acid because they have to. Other organisms (mostly bacteria, but also certain fungi and plants) have then evolved ways of degrading this urea into a usable form - namely, ammonia.

Selfishness, opportunism and invasion are not the imaginary creations of “reductionists,” supposedly blind to the interconnectedness of nature. They are very real, powerful forces in evolution, which can only produce altruistic behaviours in specific circumstances. Certainly some degree of local altruism can, and does, appear: symbiosis is clearly frequent in the world. But, as modern evolutionary theory has painfully realised [20, 5, 4], explanations based on ‘the good of the population’ are extremely fragile against the deeply opportunistic nature of Darwinian evolution. Living beings are able to enforce this altruism among their own parts, by subjecting their growth to a pre-established plan. Gaia, on the other hand, cannot rely on any such device. As a result, the global altruism which pervades the boldest accounts of the Gaia hypothesis is hardly tenable.

6.3.3 Reproduction

One final aspect of the lack of plan in Gaia concerns reproduction. Early objectors to the Gaia hypothesis pointed out that reproduction is usually an essential criterion for life, and “planets have no kids” [7]. However several authors, including Margulis, argue that on the contrary, biospheres may in principle reproduce. The reasoning is that some living creatures will eventually colonise outer space and settle previously lifeless planets; eventually, these creatures may initiate the creation of entire biospheres on those other planets, which will affect (and depend on) their environment just as much as on our own. A common example is that of humans terraforming other planets, but bacterial spores transported between planets by meteorites are not a material impossibility.

However, it is highly unlikely that those extra-terrestrial biospheres would resemble their ‘ancestors’; in fact, they would probably look very different. The reason is that they would evolve according to the basic geophysical conditions of their new host planet, which would create entirely different evolutionary challenges and opportunities. While life may well eventually flourish on Titan, it is quite unreasonable to expect that the same chemical processes (water-breaking photosynthesis, carbon dioxide fixation, oxydising atmosphere, etc.) would evolve. A fortiori, the lifeforms which would eventually emerge would

probably not be similar to Earthly organisms. Gaia, in other words, is not *strong* - its morphology is deeply dependent on its external environment, and two different environments are likely to produce very different biospheres. The fact that it is itself able to alter this same environment is important, but does not change the fact that ‘biosphere reproduction’ is nothing like the real reproduction of living beings. Again, this is a consequence of the fact that Gaia’s growth is not bound by a pre-established plan.

6.4 Conclusion

In summary, we see that Gaia is a spontaneous organism; as such, it exhibits multiple interactions which may take the form of complex feedbacks, and (at a certain scale) justifies the use of physiological metaphors. However, it is not a living being. As a result, it is necessarily limited in its powers to regulate itself, as a whole. We note that these limitations can usually be traced to the absence of a plan, and of the a priori design made possible by it.

Importantly, we are not just stating that ‘Gaia has no plan, therefore is not alive’. This would be a mere assertion of our definition. Rather, we observe that, precisely because Gaia lacks a plan (and is thus not ‘alive’ according to our definition), it is simply unable to do what its most enthusiastic proponents suggest that it does (and which happens to include those features which we associate most deeply with the nature of living beings). We believe this supports the importance of the plan (together with organism status) as a defining characteristic of life, rather than organisation alone.

In retrospect, it is easy to see how the realisation of Gaia’s organismal nature (a realisation to which the Gaia hypothesis itself contributed enormously) may have led to enthusiastic comparisons with living beings. A lack of distinction between organisms and living beings may have contributed to some confusion regarding the inner workings of Gaia.

7 Conclusion

[TBW]

[In summary:

1. ”Life” means (at least) two different things, namely the gigantic chemical reaction consisting of all living beings, and the property of being alive.
2. The NASA definition of life is very convincing for the first of these meanings. However, it cannot cover the second - neither can any definition (including cluster definitions) which include reproduction or evolution.
3. Distinction between organisms and living beings. An organism is a self-enclosing set of active parts, such that the continued existence and activity of each part is dependent on the activity of others.
4. A living being is an organism endowed with a plan, that is, an organism in which mutual production, growth, alterations and general behaviour are not the contingent result of chance, but are being guided by a pre-existing plan, to reach a certain pre-defined design. This design may change over time (development), but not so chaotically as to vanish altogether into

indefiniteness - and most importantly, these changes are themselves commanded by the plan. The existence of a plan is indeed fundamental to our intuitive concept of a living being.

5. Not all organisms have plans: some spontaneously occurring organisms may well exist as the result of unguided processes. But because they have a plan, living beings can benefit from a priori design. This is what permits the dazzling complexity and adaptation of living beings.
6. As an example, Gaia is indeed an organism. However, claims that Gaia does exhibit such lifelike adaptive complexity cannot be valid. Gaia, having no pre-defined plan, is not a living being - and, as a result, is unable to produce the level of adaptive complexity which we associate with living beings.

Suggestion: there are probably other entities, similar to Gaia and ecosystems, which are organisms without being living beings.]

References

- [1] P. Brady and S. Carroll. Direct effects of co₂ and temperature on silicate weathering: Possible implications for climate control. *Geochimica et Cosmochimica Acta*, 58(7), 1994.
- [2] A. Cairns-Smith. *Seven clues to the origin of life*. Cambridge University Press, Cambridge, 1985.
- [3] C. Cleland and C. Chyba. Defining life. *Origins of Life and Evolution of the Biosphere*, 32(4):387–93, 2002.
- [4] R. Dawkins. *The extended phenotype: the gene as the unit of selection*. Freeman, 1982.
- [5] R. Dawkins. *The selfish gene*. Oxford University Press, 1989.
- [6] C. Fernando and J. Rowe. Natural selection in chemical evolution. *Journal of Theoretical Biology*, 2006. in press.
- [7] S. J. Gould. Self-help for a hedgehog stuck on a molehill. *evolution*, 51(3), 1997.
- [8] P. Hoffman, A. Kaufman, G. Halverson, and D. Schrag. A neoproterozoic snowball earth. *Science*, 281(5381):1342, 1998.
- [9] S. Kauffman. *The Origins of Order*. Oxford University Press, 1993.
- [10] J. Kirchner. The gaia hypothesis: Conjectures and refutations. *Climatic Change*, 58(1):21–45, 2003.
- [11] D. Koshland. The seven pillars of life. *Science*, 295(22):2215–2216, 2002.
- [12] T. Lenton, H. Schellnhuber, and E. Szathmari. Climbing the co-evolution ladder. *Nature*, 431(7011):913, 2004.

- [13] J. Lovelock. *Gaia: A New Look at Life on Earth*. Oxford University Press, New York, 1979.
- [14] J. Lovelock. *The Revenge of Gaia: Why the Earth is Fighting Back - and how We Can Still Save Humanity*. Allen Lane, 2006.
- [15] L. Margulis. *Symbiotic planet: a new look at evolution*. Basic Books, 1998.
- [16] B. McMullin. John von neumann and the evolutionary growth of complexity. *Artificial Life*, 6(4):347–361, 2000.
- [17] J. Monod. *Chance and Necessity*. William Collins Sons & Co Ltd, 1972.
- [18] K. Ruiz-Mirazo, J. Peretó, and A. Moreno. A universal definition of life: Autonomy and open-ended evolution. *Origins of Life and Evolution of the Biosphere*, 34(3), 2004.
- [19] L. Villarreal. Are viruses alive? *Scientific American*, 291(6):100–105, 2004.
- [20] G. Williams. *Adaptation and natural selection*. Princeton University Press, Princeton, 1966.