

# Alien life: how would we know?\*

Margaret A. Boden

Centre for Cognitive Science, University of Sussex, Sussex, UK  
e-mail: m.a.boden@sussex.ac.uk

**Abstract:** To recognize alien life, we would have to be clear about the defining criteria of ‘life’. Metabolism – in other words, biochemical fine-tuning – is one of these criteria. Three senses of metabolism are distinguished. The weakest allows strong artificial life (A-life): virtual creatures having physical existence in computer electronics, but not bodies, are classed as ‘alive’. The second excludes strong A-life but allows that some non-biochemical A-life robots could be classed as alive. The third, which stresses the body’s self-production by energy budgeting and self-equilibrating energy exchanges of some (necessary) complexity, excludes both strong A-life and living non-biochemical robots.

Received 1 June 2003, accepted 22 June 2003

**Key words:** autopoiesis, life, A-life, artificial life, metabolism, self-organization.

## Introduction

To recognize alien life, we would need to know what we mean by *alien*, and what we mean by *life*. Let us take the first thing first. We would see life as alien, I suggest, if it were discovered on a different planet, if it involved a fundamentally different biochemistry on planet Earth or if it was artificially generated by artificial intelligence/artificial life (AI/A-life) research (robots, perhaps?).

In particular, we would see it as alien if it consisted of purely *virtual* organisms: creatures existing only in computer memory, and manifested on the VDU screen. Most talk about life, alien or otherwise, assumes some physical object (some ‘body’) – perhaps microscopically small – existing as a material thing. But this is not applicable to purely virtual organisms. The claim that such cyber-creatures could properly be regarded as *alive* is the claim that ‘strong A-life’ is possible. (Strong A-life is so-called by analogy with strong AI (Searle 1980).) And that claim cannot be assessed without considering what, in general, we would count as *life*.

That is a notoriously difficult question. In this paper, I will concentrate on one of the commonly listed criteria of life: metabolism. As we shall see, metabolism (in the sense used by biologists) is a form of biochemical fine-tuning. It characterizes all known life, as a matter of fact – and, I will argue, as a matter of necessity too. If that is right, then something that does not metabolize cannot properly be regarded as being alive. So if virtual creatures are not fine-tuned in this sort of way, they cannot be counted as *living* things. Whatever marvels the realms of alien life may hold, cyber-organisms are not among them.

To address these questions, we must consider carefully just what is meant by ‘metabolism’. The first thing to note is that metabolism concerns *the role of matter/energy in organisms considered as physically existing things*. It is not an abstract functionalist concept, divorced from the specific material realities. By contrast, the other features typically mentioned in definitions of life – self-organization, emergence, autonomy, growth, development, reproduction, adaptation, responsiveness and (sometimes) evolution – can arguably be glossed in functionalist, informational terms.

The core concept of self-organization, for example, involves the emergence (and maintenance) of order, out of an origin that is ordered to a lesser degree. It concerns not mere superficial change, but fundamental structural development. The development is spontaneous, or autonomous, in that it results from the intrinsic character of the system (often in interaction with the environment), rather than being imposed on it by some external force or designer.

Similarly abstract definitions can be given of the other items on the list. Thus, emergence is the appearance of novel properties that seem (at least at first sight) to be inexplicable in terms of earlier stages or lower-level components. Growth is increase in quantity; development is autonomous structural change leading to a higher degree of order; adaptation is an improved response to the environment by means of structural and/or behavioural change (which may be heritable); reproduction is self-copying; and evolution is adaptive change by means of reproduction, heredity, variation and selection.

It is because no comparable definition can be given for metabolism that the concept is problematic for strong A-life. A-life in general is a functionalist enterprise; that is, A-life researchers typically think of vital phenomena in terms of information and computation, not matter or energy. For example, John von Neumann defined the general requirements

\* This paper is adapted from Boden, M.A. (1999). Is metabolism necessary? *Brit. J. Phil. Sci.* **50**, 231–248.

of reproduction in logical–computational terms, and pointed out that copying errors (an informational notion) could result in adaptive evolution (Burks 1966, 1970). Similarly, in the ‘Call for Papers’ for the first conference identifying ‘Artificial Life’ as a unitary project, Christopher Langton said: ‘The ultimate goal of A-life is to extract the logical form of living systems’ (Levy 1992, p. 113).

Of course, none of these A-life researchers doubts that living things are material entities of some sort. In other words, life is not pure information. Langton makes this explicit in his statement that life is ‘a property of the organization of matter, rather than a property of the matter which is so organized’ (Langton 1989, p. 2). So far, then, the question ‘Are matter and energy essential to life?’ seems to be answered with a guarded ‘yes’. Some matter is organized, somehow. But the nature of the material stuff is philosophically irrelevant to the status of the physical system as a living thing. It could, for example, be silicon. And nothing can (or need) be said about the general type of physicochemical processes that must be going on, except that they are organized in the relevant ways.

However, Langton (1986) also says: ‘The ultimate goal of the study of artificial life would be to create “life” in some other medium, ideally [sic] a virtual medium where the essence of life has been abstracted from the details of its implementation in any particular model’.

Such ‘life’ would inhabit cyber-space, a virtual world of informational processes grounded in computers. The virtual creatures would be defined in purely informational terms, as strings of bits or computer instructions. But their activity (the execution of the instructions) – without which, they could not be regarded even as candidates for life – would require the computer. So, like biological creatures, they would have some physical existence: namely, the material ground, in computer memory, of the relevant information processing. The ‘matter which is organized’ would be the stuff of which the relevant computers are constructed – which might be almost anything. As Langton’s word ‘ideally’ makes clear, the molecules and physicochemical processes involved would be of no concern to the A-life functionalist. The only interesting properties of the virtual creatures, *qua* living things, would be abstract, informational ones.

This claim of Langton’s is disputed even by many A-life researchers, so his ‘ultimate goal’ cannot be ascribed to A-life work in general. (It follows that A-life as a whole could not be dismissed merely because one rejected strong A-life; similar remarks apply to weak and strong AI, as Searle allows.) Nevertheless, Langton is not alone in making such claims.

One of the A-life researchers who agree with him is Thomas Ray, an ecologist specializing in tropical forests. Indeed, Ray goes even further than Langton: he believes he has already implemented primitive forms of real – albeit virtual – life. His computer models of co-evolution in the virtual world ‘Tierra’ have led to the foundation of the ‘Digital Reserve’ (Ray 1992, 1994). This is a virtual memory space spread across a worldwide network of computers, which

allow their spare capacity to be used at idle times. Tierra is one example (another is described in the section on strong metabolism and strong A-life) of A-life work described by its proponent as the creation of actual, if primitive, life-forms.

The creatures (Ray’s word) inhabiting the Digital Reserve, like those within Tierra itself, are strings of self-replicating computer code. They can mate (exchange genetic instructions), mutate, compete and evolve. For example, some code-strings evolve which lack the instructions responsible for self-replication, but which can ‘parasitize’ the code of other creatures in order to replicate themselves. This is a successful evolutionary strategy because fitness is defined in terms of access to computer memory – and a ‘species’ with shorter strings can fit more individuals into a given memory space.

The creatures let loose in the Digital Reserve move from one computer to another in their search for unused memory space. (Because they are implemented in a virtual computer, simulated by some actual computer, the software creatures cannot ‘escape’ into computers not on the Reserve network, nor infest the everyday workings of those that are included.) Ray insists that the Digital Reserve is an experiment in the creation of new forms of life.

Those, like Langton and Ray, who regard strong A-life as a real possibility defend their counterintuitive view by making two interconnected claims. First, that the virtuality is limited: computers, after all, are material things, and need energy in order to function. Secondly, that the criteria for life are essentially abstract, or functionalist, saying nothing whatever about the nature of its (admittedly necessary) material grounding. To show that they are mistaken, one must show that at least one of these claims is false.

Since the first claim is indisputable, the focus falls on the second. I suggested, above, that all but one of the items on the typical list of vital properties can indeed be viewed as abstract, informational concepts. The one obvious exception is metabolism. The proponent of strong A-life must therefore show that virtual systems can genuinely metabolize. (The alternative strategy – dropping metabolism from the list of vital criteria – is discussed, and rejected, in a later section.)

In the next section, I distinguish three senses of metabolism. The first two (weaker) senses are found in the arguments of some proponents of strong A-life, for on each of these interpretations some A-life artefacts would count as genuinely alive. The third, strongest, sense is not. It is drawn rather from biology, and posits a form of bodily identity which (I shall argue in the section on strong metabolism and strong A-life) is not attained by virtual creatures.

Irrespective of questions concerning A-life, the strong sense of metabolism is more interesting than is sometimes thought. Besides referring to the biochemical processes (whatever they may be) that maintain the growth and function of an organism, it denotes various general properties that those processes *must necessarily* possess. Life on Mars or Alpha Centauri, then, would have those properties too.

### Three concepts of metabolism

What, exactly, is metabolism? It locates life in the physical world (no angels on pinheads). But it does not denote mere materiality. A volcano is a material thing, and so is a grain of sand, but neither of these metabolizes. Rather, metabolism – in the minimal sense of the term – denotes energy dependency, as a condition for the existence and persistence of the living thing.

If energy dependency were all there was to it, then strong A-life would be possible. For, as both Langton and Ray are quick to point out, virtual life satisfies this criterion. Strong A-life is utterly dependent on energy. Electrical power is needed to execute the information processes that define ‘this’ creature, or ‘that’ one. Pull the plugs on the computers, stop the electrons inside from jumping, and cyber-space is not merely emptied, but destroyed. Strong A-life, having once existed, would have died.

However, ‘metabolism’ is normally used to mean more than mere energy dependency. Two further senses of the term can be distinguished, each associated with notions of using, collecting, spending, storing and budgeting energy. These activities – a form of ‘fine-tuning’ – are characteristic of life. (Active volcanoes involve huge amounts of energy, without which they would not exist. But they do not use it, collect it, store it or even spend it, except in a weakly metaphorical sense – and they certainly do not budget it.)

A second (stronger) sense of metabolism supplements mere energy dependency with the idea of individual energy packets used to power the activities of the creature, its physical existence being taken for granted. Each living system has assigned to it, or collects for itself, a finite amount of energy. This is used up as it engages in its various activities. When the individual’s energy is spent, either because it is no longer available in the environment or because the system can no longer collect or use it, the energy-dependent behaviour must cease and the creature dies.

Some very early efforts in A-life (around mid-century) already involved the idea – and the reality – of individual energy packets. Grey Walter’s (1950, 1951) mechanical ‘tortoises’, Elmer and Elsie, were simple robots that used their energy to engage in physical behaviour. They moved around the floor by means of electric power, every so often abandoning their current activity in order to recharge their batteries. The second definition of metabolism would also cover those more recent A-life robots that are broadly comparable to Grey Walter’s tortoises, some of which even have distinct energy stores devoted to different types of activity. Such robots could, therefore, be termed alive insofar as this (second-sense) criterion is concerned.

But A-life robots are not germane to the question of whether strong A-life is possible. For ‘strong A-life’ does not refer indiscriminately to just any A-life artefacts, including robots and physical systems grounded in exotic biochemistries. Rather, it refers to virtual creatures inhabiting virtual worlds.

As remarked above, virtual creatures exist only in computer memory, manifested to the observer on the VDU screen. They ‘exist’ in the sense that they consist of a particular (perhaps continuously varying) distribution of electric charges at various (perhaps widely scattered) locations inside the machine (these locations may change, as the relevant instructions are swapped from one part of the machine to another for execution or storage). In this sense, then, they may be said to have physical existence. But that is not to say that they have bodies (see below). Nor is it to say (which is required for the second sense of metabolism) that they store and budget real energy so as to engage in their activities and to continue their physical existence.

Many virtual creatures are intended by their human creators as computer simulations of real life. That is, their manifest behaviour on the VDU screen (caused by the underlying electronic processes in computer memory) has some systematic relation to, or isomorphism with, certain features of living organisms. And some of these model metabolism (understood in the second sense), at least in a crude manner.

So examples abound of programs that simulate individual animals with distinct energy levels, raised by eating and rest, and reduced by activities such as food-seeking, fighting and mating. A few of these even assign different sub-packets of energy to various drives, so that at a particular time a creature might have energy available to mate, but not to fight. (For a very early example, where a simulated rat has to choose between seeking warmth and food, see Doran (1968).)

However, the ‘packets’ and ‘sub-packets’ here are not actual identifiable energy sources or energy stores, but mere simulations of these. At any given time, the program may dictate that the creature will seek food, but this merely means that some numerical variable has fallen below a threshold value, so triggering the food-seeking instructions. Certainly, energy is needed to execute the instructions. But this comes, via electric plug or battery, from the general undiscriminated energy source on which the whole program is passively dependent. If the program simulates more than one creature, this energy source is equally available to all, given the relevant program instructions. Metabolism in the first sense is achieved, but in the second sense it is merely modelled.

Suppose that separate energy sources, distinct real energy packets, were to be supplied (in the computer) for each simulated creature. What then? The second sense of metabolism would have been satisfied. If our concept of life involved this sense of the term, strong A-life would be conceivable.

However, one must note two important features of the second definition, as given above. First, it speaks of the creature’s ‘physical existence’, not of the creature’s ‘body’ – nor even of its being a ‘unitary’ physical system. Secondly, and crucially, it speaks of that physical existence being taken for granted.

Clearly, then, the second sense of metabolism is not the biologist’s concept of it. For no biologist ignores the

fact that an organism's physical existence is an integrated material system, or body. (Apparent exceptions include slime moulds, within whose life-cycle the multicellular organism splits into many unicellular 'amoebae', which later coalesce into a multicellular creature. But at every point, even at the amoebic stage, there is one or more integrated material system. Whether one chooses to call a reconstituted multicellular structure the 'same' organism, or body, is not important here.)

Furthermore, no biologist takes the existence of a creature's body for granted. On the contrary, one of the prime puzzles of biology is to explain how living bodies come into existence, and how they are maintained until the organism dies. We therefore need a third, still stronger, definition of metabolism if we are to capture what biologists normally mean by the term.

The third sense of metabolism refers to the use, and budgeting, of energy for bodily construction and maintenance, as well as for behaviour. Metabolism, in other words, is more than mere material self-organization. That occurs (for instance) in the Belousov–Zhabotinsky reaction, where mixing two liquids results in the spontaneous emergence of order (visible whorls and circles) – but no-one would speak of life here: too many of the other vital properties listed at the start of this paper are missing.

Rather, metabolism is a type of material self-organization which, unlike the Belousov–Zhabotinsky reaction, involves the autonomous use of matter and energy in building, growing, developing and maintaining the bodily fabric of a living thing. (For present purposes, we may apply the term 'body' to plants as well as animals.)

The matter is needed as the stuff of which the body is made. And the energy is needed to organize this matter, and new matter appropriated during the lifetime, into something that persists in its existence despite changes in external conditions. Metabolism, in this strong sense, both generates and maintains the distinction between the physical matter of the individual organism and that of other things, whether living or not.

Metabolism in this third sense *necessarily* involves closely interlocking biochemical processes. In other words, it necessarily involves a form of 'fine-tuning'.

There are several (inter-related) reasons for this. A multicellular organism must, and a unicellular organism may, sometimes grow. (I take it that multicellular organisms start off as unicellular ones; normally, this is a single spore or fertilized egg, but multicellular slime moulds 'grow' by the aggregation of many unicellular creatures.) And even a unicellular organism must (sometimes) repair damage. Since living matter cannot be created from nothing, growth and repair require that new molecules be synthesized by the organism – which molecules themselves make up the organism. Moreover, the living system (subject, like every physical thing, to the second law of thermodynamics) continuously tends to disorder and the dissipation of energy. Hence metabolism must involve continual energy intake from the environment.

The simplest conceivable living things might take their energy directly from the environment whenever they needed it. (They would satisfy only the first sense of metabolism, not the third.) Perhaps the very earliest organisms actually did this. But this would leave them vulnerable to situations in which no 'new' energy was immediately available. (Analogously, computers that rely on the plug in the wall are vulnerable to power cuts.) If, by chance, the organism became able to store even small amounts of excess energy for later use, its viability – and Darwinian fitness – would be enormously increased. Once evolution got started, this fact would be reflected in the evolution of metabolism.

Inevitably, then, all metabolic systems (other than the very earliest, perhaps) must not only exchange energy with the outside world but also perform internal energy budgeting. Excess energy is stored, so that reliance on direct energy collection is avoided. If (which is likely) the inputted energy cannot be conveniently stored in its initial form, it must be changed into some other form. In other words, living organisms must convert external energy into some substance ('currency') that can be used to provide energy for any of the many different processes going on inside the organism. This is 'the first fundamental law of bio-energetics'.

Apparently, only three convertible energy currencies – one of which is ATP, or adenosine tri-phosphate – are used by terrestrial life (Moran *et al.* 1997, paragraph 4.2). This *may* be an evolutionary accident, ATP having turned up (due to mutation) early in the game and having been retained ever since. In other words, it is not known (so my informants tell me) whether some other substance *could*, in principle, fulfil the same role. If not, then 'alien' life on Mars would be ATP-based too. If so, then perhaps it would not. It would show metabolic fine-tuning, to be sure. But the basic chemicals playing the core biophysical tunes would be different.

Additional, purely internal, energy exchanges are required as the collected energy is first converted into substances suitable for storage and then, on the breakdown of those substances, released for use. Very likely, these processes will produce waste materials, which have to be neutralized and/or excreted by still other processes. In short, metabolism necessarily involves a nice equilibrium between anabolism and catabolism, requiring a complex biochemistry to effect these vital functions.

Bodily maintenance is normally continuous. But the underlying metabolic processes are more active at some times – of the day, year and life-cycle – than at others. Sometimes, they are drastically slowed down, or (perhaps) even temporarily suspended. In hibernating animals, for instance, metabolism is kept to a minimum: respiration and excretion occur at a very low rate. Even in the case of seeds or spores frozen, or entombed, for centuries, some minimal metabolic activity may have been going on.

But what if it has not? It is not clear that this strong concept of metabolism assumes that active self-maintenance must be absolutely continuous, allowing no interruptions whatsoever. If biochemical research were to show that metabolism is occasionally interrupted, in highly abnormal

conditions (such as freezing), so be it. Indeed, we already speak of ‘suspended animation’: a spore may be currently inactive, but if it retains the potential to metabolize in suitable conditions we do not regard it as ‘dead’.

### Strong metabolism and strong A-life

The previous section showed that the first sense of metabolism is satisfied by all A-life systems, and that the second could conceivably be satisfied by certain types of A-life simulation. But what of the third, strongest, sense? Could this be found in any A-life creatures, so allowing us to regard them as living things? If so, would these creatures necessarily be robots, or could they also include virtual life?

A-life robots as currently envisaged do not fit the bill. These are typically ‘situated’ robots, engineered (or evolved) to respond directly to environmental cues. Some do not look at all life-like (Cliff *et al.* 1993). Others resemble insects in their physical form, and may have control systems closely modelled on insect neuroanatomy (Brooks 1986, 1991; Beer 1990).

Certainly, such robots are in a significant sense autonomous, especially if they have been automatically evolved over many thousands of generations (Boden 2000a). And they undoubtedly consume real energy as they make their way around their physical environment. Unlike classical robots, they are embedded in the world, in the sense that they react directly to it rather than by means of a complex internal world model. But being embedded does not necessitate being (truly) embodied. Earlier I argued that a body is not a mere lump of matter, but the physical aspect of a living system, created and maintained as a functional unity by an autonomous metabolism. If that is right, then these robots do not have bodies.

Conceivably, some future A-life robots might be self-regulating material systems, based on some familiar or exotic biochemistry. Just how exotic their biochemistry might be is unclear, however.

In principle, it need not even be carbon-based. It may be the case, however, that carbon is the only element capable of forming the wide range of stable yet complex molecular structures that seem to be necessary for life. And Eric Drexler (1989) has argued that even utterly alien (non-carbon) biochemistries would have to share certain *relational* properties with ours. They would have to employ general diffusion, not channels devoted to specific molecules; molecular shape-matching, not assembly by precise positioning; topological, not geometric, structures; and adaptive, not inert, components. In effect, Drexler is offering a functionalist characterization of biochemistry (the chemistry of metabolism), which can perhaps be instantiated in many different ways. Metabolism has also been characterized in even more abstract, thermodynamic, terms (Moreno & Ruiz 1999).

Whatever the details, artefacts grounded in exotic biochemistries might well merit the ascription of life: not strong A-life confined to cyber-space, but real, metabolizing, life. There is nothing in A-life at present that promises such alien

creatures. (It is conceivable, however, that human biochemists have already created artificial life-forms – though not robots – without realizing it, by unwittingly ‘creating the conditions under which [metabolizing systems] form themselves’ (Zeleny 1977, p. 27).) In any event, such artefacts are irrelevant to our main question. If novel robots and biochemistries were to be engineered or artificially evolved, they would count as *successful* A-life rather than *strong* A-life. The question thus remains as to whether the third sense of metabolism rules out strong A-life.

Metabolism in this strong sense, as we have seen, involves material embodiment – embodiment, not mere physical existence. It also requires a complex equilibrium of biochemical processes of certain definable types. It cannot be adequately modelled by a system freely helping itself to electricity by plug or battery, or even by assigning notional ‘parcels’ of computer power to distinct functions within the program. Virtual creatures might have individual energy packets, and some form of energy budgeting, but these would be pale simulations of the real thing. Even ‘biochemical’ A-life models are excluded from the realm of the living, if they are confined to cyber-space.

This forbids us from regarding as truly living things a ‘species’ of A-life that has recently attracted considerable attention – and whose main designer Steve Grand insists that its virtual denizens are primitive forms of life (Grand, personal communication). I am thinking of the cyber-beings conjured up by running ‘Creatures’, a computer game, or more accurately a computer world, built by the use of A-life techniques (Grand *et al.* 1996). It is a far richer virtual world than that of other computerized ‘pets’ – such as ‘Dogz’, ‘Catz’ and the electronic Tamagochi chick that the user must rest, exercise and clean. What is of special interest here is that Creatures includes a (crude) model of metabolism, as well as of behaviour.

The human user of Creatures can hatch, nurture, aid, teach and evolve apparently cuddly little VDU creatures called norns. Up to ten norns can co-exist in the virtual world (future increases in computer power will make larger populations possible), but even one solitary individual will keep the person quite busy.

One of the user’s tasks is to ensure that all the norns can find food when they are hungry, and to help them learn to eat the right food and avoid poisons. Another is to teach them to respond to simple linguistic inputs (proper names, categories and commands), with different norns receiving different lessons. Yet another is to help them learn to cooperate in various simple ways. In addition, the user must protect them – and teach them to protect themselves – from grendels, predatory creatures also present in the virtual world. The human can evolve new norns likely to combine preferred features of appearance and behaviour, since mating two individuals results in (random) recombinations of their ‘genes’.

A norn’s genes determine its outward appearance and the initial state of its unique neural-network ‘brain’ (at birth, 1000 neurones and 5000 synapses), with the specific

connection weights changing with the individual's experience. The genes also determine its idiosyncratic 'metabolism'. Each creature's behaviour is significantly influenced by its (simulated) biochemistry. This models global features such as widespread information flow in the brain, hormonal modulations within the body, the norn's basic metabolism and the state of its immune system.

The virtual biochemistry is defined in terms of four types of biochemical object. First, there are 255 different 'chemicals', each of which can be present in differing concentrations. (These are not identified with specific biochemical molecules: the functions of the 255 substances are assigned randomly.) Secondly, various biochemical 'reactions' are represented. These include fusion, transformation, exponential decay and catalysis (of transformation and of breakdown). Thirdly and fourthly, there are a number of 'emitter' and 'receptor' chemicals, representing various processes in the brain and body (for example, activity in the sense organs). Taken together, these (abstract, functional) biochemical categories are used to build feedback paths modelling phenomena such as reinforcement learning, drive reduction, synaptic atrophy, glucose metabolism, toxins and the production of antibodies.

This general architecture offers significant potential for theoretically interesting advances in A-life modelling. Its largely untapped complexity, including its ability to model global features of information processing, makes it a promising test-bed. It could be developed, for example, by incorporating recent AI ideas on the computational architecture underlying motivation and emotion (Sloman 1990; Beaudoin 1994; Wright *et al.* 1996), which as yet have been modelled only in very preliminary ways (Wright 1997).

Even now, without such additions, *Creatures* is undeniably seductive. All but the most hard-headed of users spontaneously address the norns as though they were alive, and some mourn the demise of individuals (each of whose 'life history' is unique) despite being able to hatch others at the touch of a button.

For all that, *Creatures* is a simulation of life, not a realization of it. There is no actual glucose, and no actual chemical transformation; the system is not even a chemically plausible model of specific molecular processes. Moreover, the simulated metabolism is concerned with controlling the norns' behaviour, not with building or maintaining its 'bodily fabric'. (Still less does it regulate the VDU creature's underlying, electronic, physical existence.)

Admittedly, the 'foods' and 'poisons' are associated with simulated metabolites and metabolic processes. At present, however, these affect the norns' behavioural, not bodily, integrity. They do not froth at the mouth when ingesting poison; and they do not have 'hearts' that stop beating, or 'flesh' that rots without oxygen. Certainly, some future development of *Creatures* might include a much richer metabolic simulation. The user might even be able to help a favourite norn to acquire a suntan, or to feed and exercise so as to develop its 'biceps'. Nevertheless, there would be no real metabolism, no real body – and no real life.

What if the 'foods' were to be associated with real energy, which was used only to run the electronic processes underlying the VDU manifestation of the individual norn?

This would be an example of the type of A-life system discussed above (in relation to the second sense of metabolism), in which the creature's continuing physical existence depends upon its being able to commandeer specific packets of real energy. In such a case, since the norns can evolve, they might even evolve new ways of attracting real energy and of using it (for instance) to repair their electronic grounding when damaged. Nevertheless, the points remarked above still stand: this imaginary scenario concerns the creature's physical existence, not its metabolically integrated body, and it takes that physical existence for granted. The construction of the computer, and of the parts/processes within it that constitute the norn's material being, was effected by artificial construction, not by autonomous metabolism.

In short, if we regard metabolism (in the third, biological, sense) as – literally – vital, we must reject the claim that norns, and their cyber-cousins, are simple forms of life. Even energy-gobbling and self-repairing norns, evolved without human direction, would not metabolize in this strong sense.

### Can we drop metabolism?

Someone might suggest at this point that we adopt a weaker sense of metabolism when defining life, or that we drop the criterion of metabolism altogether. In that event, some of the virtual artefacts envisaged by Langton, Ray or Grand *could* properly be regarded as alive. Such suggestions cannot be instantly dismissed as absurd. For one cannot define life, define metabolism and conclude that strong A-life is – or is not – possible in a way that will immediately convince everyone. In contrast, the concept of life is negotiable.

There are two reasons for this. First, there is no universally agreed definition of life. It is not even obvious that what one should do, in this situation, is to *try* to justify (*a priori*) a list of necessary and sufficient conditions, since our everyday concept may not name what philosophers call a 'natural kind'. I noted one example of definitional disagreement in the introduction, where I remarked that evolution is 'sometimes' – which is to say, not always – added to the typical list of vital properties. Indeed, it is regarded as 'the' fundamental criterion by many biologists, and by some philosophers – such as Mark Bedau (1996). But there are problems lurking here.

Taking evolution (or, in Bedau's terminology, 'supple adaptation') to be essential has several philosophical difficulties, as Bedau himself admits. One is that creationist biology becomes logically incoherent, not just empirically false. Another is that evolving populations, rather than individual organisms, must be taken as the paradigm case of life. This conflicts with ordinary usage, where we say that the lion – not the lion lineage – is 'alive'. It also sits uneasily with the concept of metabolism: we saw earlier that even the weakest sense of this term is defined with reference to the physical maintenance of individual things. (By the same token,

including metabolism in the list of vital criteria underscores our usual assumption that individual organisms are paradigms of life.) Nevertheless, Bedau argues that evolution is so important in theoretical biology that it should be regarded as the very essence of life. Others, by contrast, argue that evolution – and reproduction, too – is merely a secondary feature of life, and that one can envisage living things incapable of either (see below). Secondly, the concept is negotiable because *even if* everyone today defined ‘life’ in the same way, they might tomorrow have good reason for defining it differently. Scientific discoveries might lead to an (*a posteriori*) theoretical identification of the real essence of life, and hence to a change in the way that non-scientists use the term.

The suggestion that evolution be taken as essential, for example, dates from the development of modern biology. Before Darwin’s theoretical work, perhaps even before the twentieth-century synthesis of Darwinism and genetics, it would have been unreasonable to propose this – *even though* many of his predecessors believed that, as a matter of fact (not explanatory theory), living things had somehow evolved. Again, one of the research aims of A-life is to study ‘life as it could be’, not merely ‘life as we know it’ (Langton 1989, p. 2). This might eventually lead to a different, more inclusive, definition. Indeed, one new ‘essential’ vital property has already been suggested: Langton (1990, 1992) conjectures that all living things satisfy a narrow range of numerical values of the ‘lambda parameter’, a simple statistical measure of the degree of order and novelty in a system. It is not obvious that this sort of discovery is impossible. In short, with scientific advance, the list of vital properties can change.

It might appear, then, that the possibility of strong A-life hangs on mere definitional fiat. Given that there are several senses of metabolism, why not simply choose the weakest, or the strongest, so as to allow or disallow strong A-life respectively? More radically, why not drop metabolism entirely? If we can consider adding evolution, surely we can consider dropping metabolism? We could retain a commitment to physicalism. And metabolism would still be recognized as a universal characteristic of the sort of (biological) life we happen to know about. But it would no longer be seen as essential.

To see the situation in this way is to confuse fiat with negotiation. I said, above, that the concept of life is negotiable, not that it can be defined just anyhow. Both scientific and philosophical judgment must be involved in favouring one definition rather than another. And both types of judgment imply that to drop metabolism from the concept of life would not be a sensible move. That is, the analogy we are asked to draw here – between adding evolution and dropping metabolism – is too weak to be persuasive.

There are strong scientific reasons for adding evolution to the definition of life, even for making it the most fundamental criterion. Specifically, evolutionary theory has enormous explanatory and integrative power, interconnecting all (or most) biological phenomena. Even in molecular biology and genetics, evolutionary explanations provide many insights.

And most biologists who resist the reductionist approach of molecular biology, taking the form of whole organs and organisms as their explanandum, see it as not merely universal, but fundamental.

A minority do not. For instance, Brian Goodwin (Goodwin 1990; Webster & Goodwin 1996, part 2) and Stuart Kauffman (1992) argue that biological self-organization is a more fundamental explanatory concept than evolution – and that the two processes can sometimes pull in different directions (see also Wheeler 1997). But even these theoretical mavericks allow that Darwinian evolution selects, and so (superficially) shapes, the range of living things that survive, given the (deeper, wider) potentialities afforded by self-organization. In short, all serious biologists – I do not include creationists – acknowledge that evolution has considerable explanatory force. This is why Bedau is willing to accept the admittedly counter-intuitive implications of taking evolution to be necessary.

That is not to say that everyone will judge the strong reasons for adding evolution to the definition to be strong enough. In particular, those who stress metabolism as a criterion are likely to insist that we should continue to take individual creatures, not evolved species, as the paradigm of life.

Consider, for example, the argument of the biologists Humberto Maturana and Francisco Varela (1980, pp. 105–107). Their definition of life as ‘autopoiesis in the physical space’ is broadly equivalent to the third sense of metabolism defined previously (broadly, but not exactly: see Boden 2000b). They remark that the concept of evolution *logically presupposes* the existence of some identifiable unity – that is, of a living thing self-generated and self-sustained by autopoiesis. Evolution, therefore, cannot be a defining criterion of life.

Their refusal to regard evolution as essential is not a merely semantic point, following trivially from their preferred definition of life. Rather, it is a biological hypothesis. They point out that a living, self-organizing, cell could conceivably be incapable of reproduction. Even if it could be split (either accidentally or autonomously) into two autopoietic halves, there might be no self-copying involved. Self-copying requires some relation of particulate heredity between the mother and daughter systems. Furthermore, without such (digital) heredity, there can be no evolution (Maynard Smith 1966, p. 117). So the first living things might not have been capable of evolution.

My own view is that to regard evolution as an essential criterion of life is unwise. For the reasons I have outlined above, it would be better regarded as a universal characteristic, though admittedly one offering enormous explanatory power. Because of this explanatory power, it is not surprising that many biologists take evolution to be a defining property. But this definition, interpreted strictly, generates too many counterintuitive – and biologically paradoxical – implications. That is, I do not find Bedau’s arguments compelling. Even so, one must allow that he and others like him have a respectable case to make.

The same cannot be said of someone who proposes to drop metabolism as a defining criterion of life. There is no persuasive argument for rejecting our intuitions about its necessity. We have just seen that metabolism is even more fundamental than evolution, since non-reproducing organisms are conceivable and may once have lived. And the second section showed that metabolism, in the third sense, is essential for self-organizing bodily creatures that take in energy from their environment. Or rather, it is essential if that energy is not always immediately available, and it is useful if the energy is not always immediately needed. As for explanatory power, metabolism provides this. Biochemists have identified a host of specific molecular reactions involving general types of metabolic relation (such as breakdown and catalysis), and satisfying general principles concerning the storage and budgeting of energy (the ‘laws of bio-energetics’ mentioned when discussing concepts of metabolism).

In short, scientific advance in biology and biochemistry reinforces our everyday assumption that metabolism is crucial, while also enriching the concept considerably.

To outweigh this combination of scientific theory and everyday usage, powerful countervailing considerations would be needed. But none exists. The only reason for proposing that we drop metabolism from our concept of life is to allow a strictly functionalist-informational account of life in general, and A-life in particular. The same applies in respect of suggestions that we weaken the notion of metabolism, abandoning the third interpretation and substituting mere energy dependency (with or without individual energy packets). The only purpose of this recommendation is to allow virtual beings, which have physical existence but no body, to count as life. These question-begging proposals have no independent grounds to buttress them.

Significantly, it is even difficult to imagine what such independent grounds could be like. Perhaps some future science might discover strange wispy clouds, distributed over a large space yet somehow identifiable as (one or more) unitary individuals, and having causal properties analogous to those of living things – but lacking metabolism? In that case, we would have to think again. The concept of life remains negotiable, after all.

However, this futuristic scenario is well-nigh unintelligible. What are these ‘causal properties analogous to those of living things’ that do not require bodily unity? And how, in the absence of metabolism, could the clouds satisfy any self-organizing principle of living unity? The fact that science fiction writers (including the cosmologist Fred Hoyle) have sometimes asked us to consider such ideas does not prove that, carefully considered, they make sense.

Similar remarks apply to the speculative idea of a ‘cosmic computer’ (or ‘computers’) distributed across the atmosphere, supposedly supporting information-processes that evolve and adapt much as Ray’s virtual creatures do. Many philosophers argue that life is a necessary ground of cognition. If that is so, then nothing can be regarded as intelligent

which is not also alive. And if life requires some metabolizing bodily unity, then the ‘cosmic computer’ is irredeemably suspect.

The argument of this paper suggests that such science-fictional ideas are not just implausible, but irredeemably incoherent. Without independent grounds for doing so, we should not drop metabolism from the concept of life. Nor should we weaken our (third) interpretation of it. On the contrary, we should acknowledge it as a fundamental requisite of the sort of fine-tuned self-organization that is characteristic of – indeed, necessary for – life.

In summary, metabolism is necessary, so strong A-life is impossible. Alien organisms, whatever their habitats may be, do not exist in cyber-space.

## References

- Beaudoin, L.P. (1994). Goal-processing in autonomous agents. *PhD thesis*, School of Computer Science, University of Birmingham.
- Bedau, M.A. (1996). The nature of life. In *The Philosophy of Artificial Life*, ed. Boden, M.A., pp. 332–357. Oxford University Press, Oxford.
- Beer, R.D. (1990). *Intelligence as Adaptive Behavior: an Experiment in Computational Neuroethology*. Academic Press, New York.
- Boden, M.A. (2000a). Life and cognition. In *The Foundations of Cognitive Science at the End of the Century*, ed. Branquinho, J., pp. 11–22. Oxford University Press, Oxford.
- Boden, M.A. (2000b). Autopoiesis and Life. *Cognitive Sci. Q.* **1**, 115–143.
- Brooks, R.A. (1986). A robust layered control system for a mobile robot. *IEEE J. Robotics Automation* **2**, 14–23.
- Brooks, R.A. (1991). Intelligence without representation. *Artificial Intell.* **47**, 139–159.
- Burks, A.W. (1966). *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, IL.
- Burks, A.W. (1970). *Essays on Cellular Automata*. University of Illinois Press, Urbana, IL.
- Cliff, D., Harvey, I. & Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior* **2**, 71–108.
- Doran, J.E. (1968). Experiments with a pleasure-seeking automaton. In *Machine Intelligence III*, ed. Michie, D., pp. 195–216. Edinburgh University Press, Edinburgh.
- Drexler, K.E. (1989). Biological and nanomechanical systems: contrasts in evolutionary complexity. In *Artificial Life*, ed. Langton, C.G., pp. 501–519. Addison-Wesley, Redwood City, CA.
- Goodwin, B.C. (1990). Structuralism in biology. *Sci. Progr.* **74**, 227–244.
- Grand, S., Cliff, D. & Malhotra, A. (1996). Creatures: artificial life autonomous software agents for home entertainment. *Research report CSRP 434*. University of Sussex School of Cognitive and Computing Sciences, Brighton.
- Grey Walter, W. (1950). An imitation of life. *Sci. Am.* **182**, 42–45.
- Grey Walter, W. (1951). A machine that learns. *Sci. Am.* **185**, 60–63.
- Kauffman, S.A. (1992). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford.
- Langton, C. (1986). Studying artificial life with cellular automata. *Physica D* **22**, 1120–1149.
- Langton, C.G. (1989). Artificial life. In *Artificial Life*, pp. 1–47. Addison-Wesley, Redwood City, CA. (Reprinted, with revisions, in Boden, M.A. (ed.), *The Philosophy of Artificial Life*, pp. 39–94. Oxford University Press, Oxford.)
- Langton, C.G. (1990). Computation at the edge of chaos: phase-transitions and emergent computation. *Physica D* **42**, 12–37.
- Langton, C.G. (1992). Life at the edge of chaos. In *Artificial Life II*, eds Langton, C.G., Taylor, C., Farmer, J.D. & Rasmussen, S., pp. 41–91. Addison-Wesley, Redwood City, CA.



- Levy, S. (1992). *Artificial Life: the Quest for a New Creation*. Pantheon, New York.
- Maturana, H.R. & Varela, F.J. (1980). *Autopoiesis and Cognition: the Realization of the Living*. Reidel, London.
- Maynard Smith, J. (1996). Evolution – natural and artificial. In *The Philosophy of Artificial Life*, Boden, M.A., pp. 173–178. Oxford University Press, Oxford.
- Moran, F., Moreno, A., Montero, F. & Minch, E. (1997). Further steps towards a realistic description of the essence of life. In *Artificial Life V: Proc. 5th Int. Workshop on the Synthesis and Simulation of Living Systems*, pp. 255–263. MIT Press Cambridge, MA.
- Moreno, A. & Ruiz, K. (1999). Metabolism and the problem of its universalization. *BioSystems* **49**, 45–61.
- Ray, T.S. (1992). An approach to the synthesis of life. In *Artificial Life II*, eds Langton, C.G., Taylor, C., Farmer, J.D. & Rasmussen, S., pp. 371–408. Addison-Wesley, Redwood City, CA. (Reprinted in Boden, M.A. (ed.), *The Philosophy of Artificial Life*, Oxford: Oxford University Press, pp. 111–145.)
- Ray, T.S. (1994). An evolutionary approach to synthetic biology: zen and the art of creating life. *Artif. Life* **1**, 179–210.
- Searle, J.R. (1980). Minds, brains, programs. *Behav. Brain Sci.* **3**, 417–457.
- Sloman, A. (1990). Motives, mechanisms, emotions. In *The Philosophy of Artificial Intelligence*, ed. Boden, M.A., pp. 231–247. Oxford University Press, Oxford.
- Webster, G. & Goodwin, B.C. (1996). *Form and Transformation: Generative and Relational Principles in Biology*. Cambridge University Press, Cambridge.
- Wheeler, M. (1997). Cognition's coming home. In *4th Eur. Conf. on Artificial Life*, eds Husbands, P. & Harvey, I., pp. 10–19. MIT Press, Cambridge, MA.
- Wright, I.P. (1997). Emotional agents. *PhD thesis*. School of Computer Science, University of Birmingham.
- Wright, I.P., Sloman, A. & Beaudoin, L.P. (1996). Towards a design-based analysis of emotional episodes. *Philos. Psych. Psychol.* **3**, 101–137.
- Zeleny, M. (1977). Self-organization of living systems: a formal model of autopoiesis. *Int. J. Gen. Syst.* **4**, 13–22.