Dorin, A., *Computer Based Life: Possibilities and Impossibilities*, in *ISIS: Information, Statistics and Induction in Science*, Dowe, Korb, and Oliver, Editors. 1996, World Scientific Press: Melbourne, Australia. p. 237-246.

Computer Based Life, Possibilities and Impossibilities.

Alan Dorin
Dept. Computer Science
Monash University
Clayton, Vic. 3168, Australia.

aland@cs.monash.edu.au

ph: +61-3-9905-5200
fax: +61-3-9905-5146

Abstract: The need for a concise definition for life has been accentuated by recent interest in computer based Artificial Life (A-Life). We attempt to apply conventional approaches to defining life in the domain of computer programs. Chemical autopoiesis in the physical space is accepted as necessary and sufficient for life. This forces us to make a distinction between physical and virtual entities. From this, we re-formulate the goal of creating algorithmic life, making allowances for the limitations of non-physical, virtual environments. We examine a number of potential 'virtual organisms' for possession of the necessary characteristics to determine in what sense they are living things. We are lead to conclude that non-physical entities, hence computer programs, cannot be living things. Some computer programs do however share characteristics with real life.

Keywords: Artificial Life, Autopoiesis, Virtual Environments, Virtual Life.

Area of Interest: Foundations of AI.

1 Introduction

"**Life**, n. a spiritual pickle preserving the body from decay." Ambrose Bierce, "The Devil's Dictionary", 1911.

A-Life research seeks to identify the underlying principles of life by examining life-like processes. The term 'Artificial Life', used loosely, refers to any model of an aspect of a living thing. Later we introduce a term 'Virtual Life' in order to remove some of the ambiguity surrounding the use of 'A-Life' to refer to a computer program.

We must continually re-assess our beliefs as we gain insight into an immature field such as Artificial Life. This paper re-examines recent results in the field, attempting to draw some bold claims into perspective. If we do not securely ground this field we risk repeating mistakes such as those made by many Artificial Intelligence researchers as they search for their own elusive Holy Grail.

Computer technology enables us to create and manipulate complex spaces which operate at a different level from the reality in which we, as carbon-based organisms, exist. These virtual spaces are represented by physical processes over which we have control. Is it meaningful to assign the same label to a virtual structure as we do to a structure with which we co-exist? This paper is intended to deal specifically with the question of life. When should a virtual construction be deemed 'alive'? Do we gain anything by making such a link between virtual and physical spaces?

In section 2 we introduce our stance. Section 3 discusses previous attempts to define life. Particularly we look at the use of a list of characteristics for defining life. We highlight the shortcomings of previous efforts in this vein. Section 4 defines autopoiesis and relates it to computer programs. Section 5 relates autopoiesis directly to life as we know it, allowing us to carry our ideas into an examination of computer based A-Life in sections 6 through 8. We conclude with section 9.

2 Virtual and physical environments

Let us touch briefly on the relationship between a program and the physical world to be elaborated in the remainder of the paper.

An entity's physical topology is defined by the closeness and connectedness in space of its components. Topology is retained through chemical and physical properties of matter.

Interactions between an entity and its environment are also determined by these properties and described by physical and chemical laws.

A computer program is a symbolic representation of an algorithm. It is realized physically in the state of switches within a computer. The topology of a computer program, the relation between the specific components that make it recognizable, is defined by the connectedness of its *instructions*. Although the instructions are realized in a physical medium, their connectivity is a *logical* relation established by rules for code interpretation. The *physical* relation between instructions is not relevant in determining the topology of a computer program. A computer program is therefore fundamentally different to a physical machine, even the computer on which it may run.

Programs resident in a digital silicon computer or a biological carbon based computer, operate at a different level to that of the physical universe. They require the matter of which the machine is built to behave according to the laws of physics, but the behaviour of the program as a whole is not reducible to these laws. (See Polyani [16], and Hull [6] for comment.)

The topology of a living system is realized through the chemical and physical laws of the universe. An organism may be viewed as a machine for computation with a program operating in some domain[1]. It is the machine itself which is the living system, not the program we interpret as running on it. A program running on any computer cannot be a living system due to its non-physicality. The remainder of this paper makes clear the reasons we believe physicality is a necessary characteristic of living systems.

3 Listing life's properties

Reductionism does not provide the answers to all questions. Analysis of biological components has contributed to our understanding of existing life but is yet to make clear what it means to live. We also lack a method for synthesizing arbitrary living systems. My conclusions rest partly on the belief that we may one day synthesize artificial organisms, even if the method escapes us at present.

Vitalism, the belief in a 'life-force' peculiar to living things and different to all other physical forces, was proposed by Aristotle as a means of explaining the difference between living and non-living things. This theory went the way of Phlogiston and Ether. There is no experimental evidence to suggest it exists, nor is it necessary. Chemical and physical properties of matter seem sufficient to explain living phenomena.

Definitions for life often consist of a loose list of properties. The authors, for example Farmer and Belin [4], are frequently the first to admit that their lists are incomplete and inaccurate. The attraction to the list based approach may be due to the difficulty of describing a system of dynamic components. The tendency is to view the complex system's components in isolation, describe their properties, and hope vainly that this will shed light on the whole. Such attempts may be marginally helpful in the identification of Earthly life, even if they are not defining properties of life in general. More concise ways of describing living systems have been proposed.

Autonomy, self-production, dynamics, maintenance of identity under perturbation, these are important characteristics in any definition of life. Earthly life is characterized by continuous change within the body of an organism. This feature we take to be fundamental. There is no clear example of life as we know it for which this does not hold[2]. We therefore expect that a widely acceptable definition for life of any form, must account for this preconceived idea. A static structure we do not consider living, regardless of any other characteristics it possesses[3].

---

[1] Such as the domain of human thought.

[2] A virus is not dynamic. According to our criterion it is not alive. Other definitions for life may classify a virus along with living things.

[3] The author holds that intelligence may be possessed independently of life. Note that others, (for example Yaeger

We are interested in defining life as it applies to an individual organism. This is the level at which the word is used in everyday speech. It is therefore a criterion that any definition must meet if it is to be widely accepted. A definition that applies only to an entire geneology, or perhaps only to an entire ecosystem / planet, is not a definition for life. It is a definition for a new word.

A frequently recurring element of the list based approach to defining life is also the most clearly inaccurate element of the lists in which it appears. This planet's organisms can often be distinguished from the non-living by virtue of their evolvability and ability to reproduce. One may also witness the development of a phenotype (the organism and its behaviour) from a genotype (a DNA strand). It is not true that these properties are *necessary* features of life.

Both a mule and post-reproductive human are living systems lacking the ability to reproduce. The mule could never evolve. Could one not envisage a form of life, such as life created in-vitro, which was established by a process other than reproduction and evolution?

A line of argument proceeds: If we create life in a test-tube, then because that life is produced by us, it too has links to our evolutionary chain. This implies that all conceivable life is the product of an evolutionary chain.

Although (perhaps) minuscule, there is some measurable probability that a simple living thing, let's say a single cell, may 'spring to life' purely by chance combination of elements. Perhaps assembly would act under self-organizing principles.[4] We would have to consider the thing living if its current state was comparable to that of another cell that *was* produced by an evolutionary chain. That is, evolutionary history is not important in *defining* life, although knowing that something was produced by an evolutionary / reproductive chain, may guide us in its classification.

Once we have established that an organism may spring to life via some process other than reproduction and evolution, we may discard the idea that an organism (the phenotype) must develop from a genotype. An organism's developmental history is irrelevant when considering its status as a living system. We can decide the matter of life without reference to a system's past. Only an organism's current constitution determines whether or not it lives.

Related to reproduction and development is the storage of a self-representation such as a DNA strand. To say that an organism stores a 'self-representation' is misleading and unnecessary. All we are claiming is that the processes and components for topological maintenance are produced within the boundaries of the organism. The body arises out of the interaction between its components and the environment. These interactions may be explained by reference to the laws of physics / chemistry. An elephant arises out of the interactions between the bases which constitute elephant DNA in an environment suitable for development.

An observer labels a relatively simple structure as 'information' which, under suitable conditions, initiates development and maintenance of the organism. A clock does not have this characteristic of self-development or maintenance. It is not equipped to harness the laws of physics / chemistry in this manner.

Should we manufacture a living organism from scratch, there may be no need for it to possess a structure which can be viewed as a self-representation. A form may arise maintained only by the interactions of its compo-

[20]) may hold the view that an intelligent being is necessarily alive, regardless of its other properties. If we take the Turing test as a starting point for recognizing intelligence, then there is no reason to believe that a human can only attribute intelligence to a living thing.

[4]The *first* living thing must have arisen in some similar manner. It did not come about through reproduction because the first living thing must differ in some important way from the non-living things which preceded it. The non-living things at the time could therefore not be *reproducing* if they gave rise to a thing *substantially* different to themselves. Eg. If a robot gave rise to a caterpillar, it could not be said that the robot was *reproducing*.

nents with the environment and each other. No elementary information storage is necessary. Such a structure would be emergent from the environment. It would have no self-representation. Order arises from the laws which govern its surroundings and the interactions of its components.

A representation of self appears only to be useful in reproduction, development and regeneration. We do not require a living thing to reproduce or develop from a genotype. Nor do we require an organism to possess a representation for self maintenance. Therefore we do not require possession of a model of self.

Some authors, notably Kauffman [7] and Prigogine [17], claim organisms are chemical autocatalytic systems. Living systems are necessarily autocatalytic but chemical autocatalysis is not sufficient for life. If it were a sufficient condition, we must include chemical dissipative structures and other ‘far from equilibrium’ systems in our list of organisms. These things are not normally considered to live. In addition to autocatalysis, we require that the organism be responsible for the maintenance and construction of its own topology. Chemical autocatalytic systems do not all meet this requirement. The following section explains what we mean by maintenance of topology.

4 Autopoiesis

Having discarded a few of the conventional list items as being unsatisfactory in helping us define life, let us now examine autopoiesis. It is argued by Maturana and Varela [13] that autopoiesis in the physical space, in the domain of chemical reactions, is necessary and sufficient to characterize life. We do not elaborate their argument here as we believe it to be adequately dealt with in the original publication.

The precursors of Maturana and Varela are briefly described by Zeleny [24]. What follows is a brief discussion of the definition and its terms by way of a summary of the early descriptions given by Maturana and Varela [13] and Zeleny [25].

"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 realise 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 realisation as such a network." [13]

4.1 Organization and structure

An observer distinguishes a physical or non-physical unity from its background by perceiving a boundary around it. A unity is fully characterized by its *organization* and *structure*. We shall now examine these terms as they apply to physical entities.

We may view a mechanical clock as a *simple* unity, a unity without constituent parts. We can only characterize this clock by the properties we assign it in our perception of it as separate from its environment. If we distinguish between the components of the clock, we are viewing it as a *composite* unity. We can now also characterize the clock by the relations between its parts. It is the relationship between the components allowing us to classify the unity as a clock, that we label the machine’s *organization*. The clock's components themselves play no part in the organization of the machine, only the relations which they must satisfy for the unity to be labelled a clock.

The perceived function of a unity plays no part in characterising its organization, although it may be a convenient way to communicate an intended organization. That a machine's purpose is to allow us to read the time for example, is not a feature of the organization of a clock. It is a feature of the context in which we view the machine i.e. the environment in which the machine is likely to act including the participation of a clock-reading observer.

The particular components that realise an instance of a unity, and the precise relations between them, constitute the unity's *structure*. A unity's specific organizational relations are an aspect of its structure.

If the organization of a clock shifts just outside the acceptable limit we label the clock 'broken'. We still find it believable that the unity was, or will be, a clock. Should we radically alter the organization of the machine we may well not recognize the unity as clock-like at all.

Varying the structure in such a way as to maintain the necessary organization does not change the class of the unity. It does allow us to distinguish between different instances of the unity. Such changes allow us to distinguish between a digital clock and a mechanical clock, or between a particular grandfather clock and a different grandfather clock.

4.2 Organization and structure of computer programs

Above we have seen how a physical object may be classified by its organization and recognized by its structure. We now clarify notions of structure and organization of non-physical entities such as computer programs.

Instructions and values contained in a set of memory locations constitute the structure of a computer program. If we view a computer program as a composite unity, the components may be the procedure calls, data structures, instructions or switching patterns which realise the algorithm in the physical space. In this environment physical relations play no part in the organization, hence the classification of a unity.

The organization of a program is based upon the order of access to its operations. If we move a program from one set of physical memory locations to another, we are changing its structure (not its organization). Two copies of a program in memory may be distinguished because of their different structures.

A sorting program is classified by the relations between its operations. Within the class of programs for sorting, an infinite number of different structures are acceptable. So long as the correct relationship between the parts holds, the program's organization is that of a sorting algorithm.

A sorting program containing a 'bug' may still be recognized as a 'broken' sorting algorithm. Repairing the bug will ensure the algorithm's organization satisfies the criteria required of a sorting algorithm.

We may alter the structure of a sorting algorithm by substituting equivalent logical instructions, switching the order of commutative operations, or even altering the entire set of operations, without destroying its organization. For example, a merge sort and a bubble sort appear radically different and yet both are classified as sorting programs. The two algorithms could be realised on two different architectures, in widely different computer languages, yet the organization of the unities that allows us to classify them both as sorting algorithms is invariant. We can distinguish the unities as separate because of their different structures.

4.3 Processes for Autopoiesis

Now that we have clarified the terms organization and structure, we may proceed to apply them to autopoiesis. According to Maturana and Varela's definition outlined above, an autopoietic machine must support special processes acting on its components. The processes of production, transformation and destruction of components must produce the components which regenerate the same set of processes and the organization from which they themselves were produced.

An autopoietic machine must be in a state of flux. For a unity to be autopoietic it must undergo continual structural change by the production of new components which realise its own organization. If the processes are in some way interrupted or suspended, the unity is not, whilst the suspension continues, autopoietic.

The autopoietic system is a machine whose organization is defined by the relations between processes for the construction of those same processes. If this ceases the machine is said to be allopoietic, the result of its operation is something other than its own organization. It is true that there are a large number of manufactured unities which maintain dynamic relations amongst their components. These machines do not specify their own organization by this process. They are not autopoietic.

The process carried out by a mechanical clock is one of transformation of stored energy into the motion of gears and the hands on the clock face. Such a clock is allopoietic. The organization of the clock is not responsible for the creation of new components, nor processes for the construction of components. It is not specifying its own organization.

A sorting algorithm is allopoietic in the same way as a clock. The processes of operation of a sorting algorithm are in no way responsible for creating afresh the components and organization of that same algorithm.

The second requirement of processes in an autopoietic machine is that they realise a particular instance of the unity in the space in which the components exist. This is achieved by specifying the topology in which the processes act. The boundary between the unity and its environment is defined by the space in which the component producing processes act.

The production of a clock is not the result of that same clock's operation. A machine quite separate from the unity we label 'clock' is responsible for establishing the topology of the clock. The same can trivially be said of sorting programs.

4.4 On the topology of computer programs

A physical body possesses a topology determined by the spatial relations and connectedness of its parts. We have seen that this is determined by the physical properties of matter and laws governing its behaviour. Interactions between components are causally connected. Causality is a property of physical matter.

As discussed above, physical connectedness of a computer program is irrelevant to its organization. This is also true of the program's topology. A program's topology is established by rules governing the environment in which it exists. These rules are logical rules for code interpretation and order of access to the components of an algorithm. The interactions of a program's components are not causally connected. Only the behaviour of a physical machine (such as the computer in which the program resides) is causally determined.

In many cases, the state of a computer program will be represented visually using computer graphics. The use of such representations simplifies the act of distinguishing a unity from its background. This is especially helpful where the unity in question exists only as a sequence of invisible states in an enormous set of electronic switches.

We may choose to view the visual representation of the computer's internal state as the topology of the unity represented by that state. We should not forget that this is not a 'concrete' topology. The form we see is one of an infinite number of arbitrary ways to represent a pattern which exists in the memory of the computer. The apparent physical connectedness of lit areas on a computer monitor does not directly reflect, nor have any bearing on, the underlying physical connectedness of the algorithm's components.

We may also choose to view the program's topology as a set of data structures. This too is an arbitrary way of representing the program's current state. It too is unrelated to physical proximity of the components.

For a (composite) unity to be autopoietic, it must exist in an environment complete with laws establishing consistent interactions between its components. The laws must support the establishment and maintenance of a topology. This implies that the laws must support some temporal equivalent.

A unity must produce particular components at a location and rate suitable for autopoiesis. The environment must operate under a set of laws which differentiate between the behaviour exhibited by different types of materials and their interactions. It is possible for such laws to be created in the virtual space. As we shall see this does not necessarily help us to create life.

5 Autopoiesis and life

Maturana and Varela's claim that a physico-chemical, autopoietic machine <u>is</u> a living entity is different to conventional approaches for defining life. The components of a living thing are examined in relation to the organization they must maintain under perturbation from the environment, rather than any individual properties they require.

We can examine the components of an autopoietic unity as if they had an input, output and a function, that is, as allopoietic machines. Doing so for all the components will not reveal the whole's autopoietic nature. Autopoiesis emerges from the dynamic interaction of processes.

A still image of an organism is insufficient to demonstrate its autopoiesis. An image does not capture the continuing processes of production and transformation required for autopoiesis. The autopoietic organization of a machine is defined by the relations between dynamic processes, not by spatial relations existing between static components.

The living things about us are autopoietic and die if their autopoiesis ceases. Their autopoiesis exists in the domain of chemical reactions. That is, the processes in an organism are chemical ones. They are processes of production, transformation and destruction of molecules. Systems that meet these criteria include complex animals, plants, Protozoa and Fungi.

Our argument does not require the sufficiency of chemical autopoiesis for life. It is enough here that we accept that it is at least necessary for life.

Our view of life is based on a single class of examples, the class of earthly carbon based life. We must expect that a broad definition will require us to make some adjustments to our ideas on life. Taking autopoiesis as a necessary condition for life may require us to change our way of thinking. In particular, we question the legitimacy of referring to computer programs as living.

In the use of language, a definition comes to represent a border between objects of one class and other objects. There is usually a *grey* area in the middle where things do not clearly fit into one category or the other. We should be satisfied with a definition for life if commonly encountered organisms will continue to be called living, whilst clocks, sorting algorithms and such things remain outside the class of living things.

We must be prepared for the shifting of the grey area between the categories. In this case, the grey area does not refer to any problem distinguishing between autopoietic and allopoietic entities, rather that certain things once considered alive may no longer be defined as such. Things once thought not to be alive may now fit the definition for life. A virus may be seen to fall within the grey area. Changes such as these are encountered daily as the environment in which we use language drifts.

6 Candidates for computer life

Computer based A-Life has demonstrated its usefulness in furthering our understanding of life on Earth [9,10,11,21]. There are many noteworthy examples of A-Life 'organisms' which exhibit behaviour normally associated with life (for example, Reynolds [19]). Pattee [15] and many others have made the point that a simulation remains just that. These A-Life organisms are not living things. We have expounded a similar view above although we have argued along a different line to many.

An autopoietic simulation realised on a computer is not potentially a living thing. It may at least have some limited claim to 'life' within its simulated environment. If we take MacLennan's [12] view of a computer as a machine for manipulating matter / electrons, simulated autopoiesis may be realised in the physical world. The computer acts as the intermediary between our controlling program and the matter with which we are accustomed to dealing. Whether or not the process simulating the autopoiesis is itself autopoietic is a separate issue.

It is possible to simulate autopoiesis in the virtual environment whilst the matter of which the simulation is maintained does not constitute an autopoietic system. In such a case the organism *may* only be said to 'live' in the virtual world, not in the physical world. It may well be nonsensical to refer to it as a living organism at all. We discuss this shortly. Maturana and Varela's claim extends only to autopoiesis in the physical world.

6.1 PolyWorld & Tierra

Yaeger's PolyWorld creatures [22] and the Tierrans of Ray [18] provide fascinating insights into the behaviour of artificial systems. Inhabitants of these worlds simulate many characteristics of complex organisms. These include behaviours such as grazing, flocking, self-defence, parasitism and its resistance as well as reproduction and evolvability. The creatures of PolyWorld and Tierra however do not satisfy our proposed necessary criterion for life. Even within their own space they are not autopoietic systems.

It is commendable that Yaeger compares his creature's properties to those suggested by Farmer and Belin [4] as having some bearing on life. Farmer and Belin point out the unsatisfactory nature of their list. Its problems are noted by Yaeger who concludes in reference to his PolyWorld creations, "So with the above caveats, questions, arguments and counter arguments, it would appear that the organisms of PolyWorld come surprisingly close to fulfilling Farmer and Belin's set of criteria; indeed, they *may* do so entirely." (Yaeger's emphasis) Perhaps I am misinterpreting Yaeger when I say that he appears incredulous. He is not letting go of his experience when his conclusions seem to be counter-intuitive. Although this is often a mistake in dealing with complex systems, he is right to question the status of his creations as living things.

The creatures of PolyWorld are static data structures specified in advance by Yaeger in the construction of the simulation. The environment and its laws are also specified by the world's ruler via a computer program which operates on the data structures according to built in rules. These rules also fully define the topology, if one may call it that, of the individual creatures which inhabit the world. The topology and components of PolyWorld creatures are not generated by any process of which they are a part. Regardless of their other characteristics, the creatures of PolyWorld are not autopoietic, even within their virtual realm. As established, autopoiesis appears at least to be necessary for life. We therefore claim that the creatures of PolyWorld do not live in the space in which their topology is defined.

The unities of interest in the Tierra simulation consist entirely of blocks of machine code operating on virtual CPU's. Tierran programs replicate by constructing a copy of their own code. Evolution is supported by introducing a preset failure rate for the execution of a Tierran instruction. This may occur during the production of a copy of an organism (replication), or in the execution of other Tierran machine code operations.

Although the Tierran programs exhibit many characteristics one would find in a list of lifelike qualities, they are not autopoietic. A program's organization is established by human hand or by a pre-existing parent program. A Tierran does not specify its own organization and is not responsible for the maintenance of its topology within the Tierran space. A Tierran program is not dynamic. Its structure remains constant throughout its identifiable existence. A

Tierran program's organization is not defined by a set of processes for maintenance of those same processes.

Yaeger [22] reports that Ray defines a living system as any system capable of replication and open ended evolution [18], although in personal communication also reported in [22], Ray seems to have changed his mind. This is healthy speculation. As we have seen, replication and evolvability are neither necessary nor sufficient criteria for living systems. The Tierran programs may examine themselves or other programs in order to manufacture a copy of their own code. This does not contribute to autopoiesis. A Tierran is not autopoietic. It does not live even in its own environment.

6.2 Cellular Automata

Conway's Game of Life [1], Zeleny's simplified model of an autopoietic system [24,25,26] and many other systems are known collectively as Cellular Automata. CA's provide a rich environment for studying emergent behaviour and A-Life whilst avoiding the unnecessary baggage of evolution, reproduction, the maintenance of a model of self etc.

The laws of the universe in Conway's Game of Life govern relations between adjacent cells in a two dimensional grid over discrete time intervals. All unities are visually represented by adjacent cells on this grid being lit in a pattern distinguishable from the background. Many stable patterns emerge under the rules specified by Conway. We shall examine a stable structure known as the *spinner* in relation to our definition of autopoiesis.

A spinner is identifiable as a composite unity consisting of three collinear, neighbouring, lit cells. Two neighbouring unlit cells forming a copy of the structure rotated by ninety degrees with the same lit center cell, are guaranteed to be part of the spinner at the next time step (fig. 1). The spinner's cycle arises through processes of production, transformation and destruction of components at a rate and position specified by the logical laws of interaction of the space.

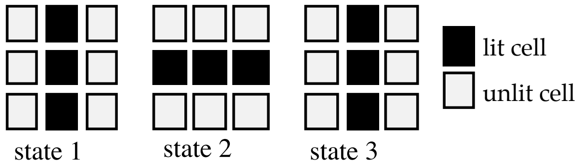

Figure 1: States of a Spinner

The cycle continues to regenerate its own components, allowing the continuation of the same process which created them. The organization of the unity is responsible for its own maintenance and the maintenance of its topology. In this world, virtual grid topology is the be all and end all of existence. The processes of production, transformation and destruction of components occur in a logical or virtual space governed by the rules of a computer program.

Should any cell adjacent to a spinner become lit, the interaction between it and the spinner will destroy the spinner's organization rendering it unidentifiable. The cells we associate with the topology of the spinner do not prevent interference between the spinner and any other lit cells in the environment. Interference from outside is fatal to the maintenance of the spinner's virtual autopoiesis.

The only way in which a unity can be identified in the Game of Life is through recognition of its topology. Topology emerges in this world through the dynamic interaction of the same cells which are its components. A recognizable and stable unity in this environment cannot fail to be autopoietic. The Game of Life does not support the conservation of virtual matter as the physical world conserves matter (and energy), so it does not support the existence of static structures. All structures must be maintained within the limitations of the dynamics of the simulation as governed by Conway's laws.

It is possible that an autopoietic unity may act as a machine for the production of other autopoietic unities. A 'glider gun' is an example of such a machine. Production of gliders from the gun occurs as a part of the autopoiesis of the gun. The glider gun is a stable but dynamic collection of cells which produces a stream of unities which themselves maintain a recognizable form. The gliders produced by the gun are, like the spinner, cyclic in their topology. Unlike the spinners, they maintain a form which is in constant motion across the grid world.

7 Autopoiesis and simulation

Supposing we accept that physico-chemical autopoiesis is at least necessary for life, let us turn our attention to determining whether or not Conway's CA organisms are potentially alive.

In a previous section we discussed the topology of computer programs and found them to be non-physical structures. No organism on Conway's grid can be autopoietic in the physical space. No such structure possesses the necessary property of life.

What if we attempt to physically implement a spinner? We build a dedicated automata grid of physically independent automata connected by wires and acting under Conway's rules. A unity on this machine can be recognized by the relation between the automata whose states continually re-establish its pattern. The arrangement of switching patterns in this case is <u>not</u> autopoietic in the physical space. Its topology in the physical space is determined by the manufacturer of the switches and the wire through which the electrons flow. This is not connected to the state of the spinner / switching pattern.

By our earlier claim, the spinner is not alive. Unless we can realise the Game of Life in the physical world, where its organization is determined by the physical laws governing matter, where the interactions of its components are causally connected, its inhabitants do not possess the necessary characteristic of life.

The organization of a computer program is not governed by chemistry or physics. Although it is realised under these laws, the laws are not important in determining the class of computer program. The components of the computer program do not exist in the physical world, they are algorithmic constructs defined in the virtual space. A virtual unity is not even potentially a living thing.

8 How close can we get?

Strictly speaking the autopoietic structures in the Game of Life are no more alive than those of Ray and Yaeger. <u>If</u> we accept the following criteria, then a virtual autopoietic unity may be considered to live in its own space:

(i) The virtual environment is a valid place for an organism to live.

(ii) Autopoiesis (in some domain) in a given space is sufficient for life in that space.

By convention the virtual world is not a place in which an organism may live. We do acknowledge that the virtual arena provides an environment useful for research into A-Life. Although the virtual environment cannot support life, not even A-Life, it does allow us to model many aspects of living things.

Whether or not we accept autopoiesis in the domain of chemistry is sufficient for life is outside the scope of this paper. The author is satisfied with the notion, the reader is referred to the original paper by Maturana & Varela [13] and subsequent papers for further discussion. If we accept this idea, then we have more to consider before we accept the second proposition.

A whirlwind is an autopoietic system. Its organization is defined by a set of processes for the production of components which satisfy Maturana and Varela's definition above. The whirlpool exists in the physical space but the processes of production, transformation and destruction of components are not *chemical* , they do not involve the production etc. of *molecules*. A whirlpool is therefore not potentially a living system by Maturana and Varela's definition. This means of course that autopoiesis in a space is not sufficient for life in that space. In the physical space we have both physical and chemical autopoiesis, only one of which we claim gives rise to a living system.

What then of other spaces? Perhaps the notion of 'life' should best be left as an organization arising in the physical space. Having said that, it is pertinent to point out that there is no common analogy at present to distinguish between a virtual chemistry and a virtual physics.

We have agreed upon the boundary between the domains of physics and chemistry. We can therefore claim that life is a particular kind of organization concerning the interaction of molecules. These interactions fall within the domain of chemistry. Seeing as all interactions in virtual space are at present 'virtual interactions', perhaps virtual autopoiesis may be as near as we can get to life in the virtual space. There is no other domain in the virtual space in which things can interact.

If we do not accept (i) or (ii) we should at least recognize the special properties of a virtual autopoietic unity. We may do this by maintaining the possibility that they be classified as examples of *Virtual Life*. Such creations could not be considered living organisms and therefore not even examples of A-Life in the strict sense of the term. If we accept that autopoiesis in a space is at least necessary for life in that space, then a virtual autopoietic system has the potential to fit into this new class of computer programs.

We may refer to an example of Virtual Life as a Virtual Organism (or just a *Vorganism* to carry on the A-Life tradition of inventing buzz-words).

A Virtual Organism is a unity whose components maintain the topology of the whole only within the virtual world, but which satisfies the other necessary characteristics of life. One of these characteristics (possibly the only one) being autopoiesis in some as yet unspecified subset of the virtual domain.

We have now steered the way towards a reasonable goal. A computer, in the current sense of the word, cannot support Artificial Life. Instead we make a distinction between Virtual Life and Artificial Life. We propose that the goal of computer programmers be the creation of Virtual Life. As a means of understanding the behaviour and organization of organisms, we may study Virtual Organisms which possess some important traits of living things.

9 Conclusions

We have highlighted a number of misconceptions regarding the use of computers to create life. A computer program is fundamentally different to a living thing. A computer program's topology is logical, not physical. The interactions of its components are not causally connected. This on its own is sufficient reason to disqualify a program of any sort from being classified as a living system.

We have distinguished the creation of Virtual Life from the general study of Artificial Life by admitting that even a non-physical program may satisfy other criteria necessary for life including autopoiesis in some subset of the virtual/logical space.

It is hoped that this distinction will allow us to further discuss the lifelike properties of virtual organisms, without having to lay extravagant claims regarding their status as living organisms.

Acknowledgments

Thanks to Kevin Korb for giving me a kick and for hours of thought provoking discussion.

Bibliography

[1] Conway, J.H. “What Is Life?”, Winning Ways for Your Mathematical Plays, Berlekamp et al (eds), Vol. 2, chap. 25, Academic Press, 1982.

[2] Dawkins, R. “The Evolution of Evolvability”, *Artificial Life*, SFI Studies in the Sciences of Complexity, Vol VI, Langton (ed), Addison-Wesley 1988, pp201-220.

[3] Dorin, A., Martin J. “A Model Of Protozoan Movement For Artificial Life”, *Computer Graphics International ‘94* Conference Proceedings, (in press).

[4] Farmer, J., Belin, A. “Artificial Life: The Coming Evolution”, *Artificial Life II*, SFI Studies in the Sciences of Complexity, Vol X, Langton et al (eds), Addison-Wesley, 1991, pp815-840

[5] Holland, J. "Adaptation In Natural And Artificial Systems", MIT Press, 1992.

[6] Hull, D. “Philosophy of Biological Science” Prentice-Hall, 1974.

[7] Kauffman, S.A. “The Origins Of Order”, Oxford University Press, 1993.

[8] Koza, J. R. “Genetic Programming”, MIT Press 1992.

[9] Langton, C.G. (ed) “Artificial Life”, SFI Studies in the Sciences of Complexity, Vol VI, Addison-Wesley, 1988.

[10] Langton, C.G., Taylor, C., Farmer, J.D., Rasmussen, S. (eds) “Artificial Life II”, SFI Studies in the Sciences of Complexity, Vol X, Addison-Wesley, 1991.

[11] Langton, C.G. (ed) “Artificial Life III”, SFI Studies in the Sciences of Complexity, Vol XVII, Addison-Wesley, 1994.

[12] MacLennan, B. “Synthetic Ethology: An Approach to the Study of Communication”, *Artificial Life II*, SFI Studies in the Sciences of Complexity, Langton et al (eds), Vol X, Addison-Wesley 1992, pp631-658.

[13] Maturana, H. Varela, F. “Autopoiesis, The Organization of the Living”, in “Autopoiesis and Cognition, The Realization of the Living”, Reidel, 1980, pp73-140.

[14] van de Panne, M., Fiume, E. “Sensor-Actuator Networks”, *SIGGRAPH 93* Conference Proceedings, ACM Press 1993, pp335-343.

[15] Pattee, H. “Simulations, Realizations, and Theories of Life”, *Artificial Life*, SFI Studies in the Sciences of Complexity, Langton (ed), Vol VI, Addison-Wesley, 1988, pp63-77.

[16] Polyani, M. “Life’s Irreducible Structure”, Science, Vol 160, 21 June 1968, pp1308-1312.

[17] Prigogine, I., Stengers, I. “Order Out Of Chaos”, Flamingo/Harper-Collins, 1985.

[18] Ray, T.S. “An Approach to the Synthesis of Life”, *Artificial Life II*, SFI Studies in the Sciences of Complexity, Langton et al (eds), Vol X, Addison-Wesley 1992, pp371-408.

[19] Reynolds, C. W. “Flocks, Herds and Schools: A Distributed Behavioural Model”, *SIGGRAPH 87* Conference Proceedings, ACM Press 1987, pp25-35.

[20] Sims, K. “Evolving Virtual Creatures”, *SIGGRAPH 94* Conference Proceedings, ACM Press 1994, pp15-22.

[21] Varela, F.J., Bourgine, P. (eds) “Toward a Practice of Autonomous Systems”, Proceedings of the First European Conference on Artificial Life, MIT Press, 1992.

[22] Yaeger, L. “Computational Genetics, Physiology, metabolism, Neural Systems, Learning, Vision and Behaviour or PolyWorld: Life in a New Context”, *Artificial Life III*, SFI Studies in the Sciences of Complexity, Vol XVII, Langton (ed), Addison-Wesley, 1994, pp263-298.

[23] Zeleny M, (ed) “Autopoiesis, Dissipative Structures and Spontaneous Social Orders” Westview Press, AAAS Selected Symposia Series 1980.

[24] Zeleny, M. “Autopoiesis, A Paradigm Lost?”, *Autopoiesis, Dissipative Structures and Spontaneous Social Orders*, Zeleny (ed), Westview Press, 1980, pp3-45.

[25] Zeleny, M. “What Is Autopoiesis?”, *Autopoiesis, A Theory Of Living Organization*, Zeleny (ed), North Holland, 1981, p4-17

[26] Zeleny, M., Klir, G.J., Hufford, K.D. “Precipitation Membranes, Osmotic Growths, and Synthetic Biology”, *Artificial Life*, SFI Studies in the Sciences of Complexity, Langton (ed), Addison-Wesley, 1988, pp125-139.