Scientific Research Publishing

Journal of Computer and Communications, 2025, 13(7), 262-294
https://www.scirp.org/journal/jcc
ISSN Online: 2327-5227
ISSN Print: 2327-5219

Are Robots Alive?

Adrian David Cheok, Emma Yann Zhang, David Levy

School of Automation, Nanjing University of Information Science and Technology, Nanjing, China
Email: adrian@imagineeringinstitute.org

How to cite this paper: Cheok, A.D., Zhang, E.Y. and Levy, D. (2025) Are Robots Alive? *Journal of Computer and Communications*, **13**, 262-294.
https://doi.org/10.4236/jcc.2025.137014

Received: June 13, 2025
Accepted: July 26, 2025
Published: July 29, 2025

Abstract

Some attempts to answer the title question require a clarification of what is meant by "alive"—how the word is defined by biologists, other scientists, philosophers and experts from other disciplines. Such attempts fail because of the lack of a suitable definition of "alive" to serve as our starting point. This failure prompts us to consider various sets of criteria of life, criteria that have been promoted as enabling us to determine whether or not a particular entity is alive. This attempt too fails, because there are so many such sets and so many differences between them that they create confusion rather than clarity. We also consider a more general set of criteria, a set devised in the 1970s and known collectively as Living Systems Theory, which does not rely on traditional biological considerations. Here we have more success—if the theory is correct, we may indeed conclude that robots are alive. We then examine how advances in the various 21st century branches of biology have paved the way for the birth of a new science—Living Technology—which brings us much closer to being able to provide a definitive answer to our title question.

Keywords

Human-Robot Interaction, Humanoid Robotics, Intelligent Systems, Living Systems, Meanings of "Alive", Robot Life, New Biologies, Robot Reproduction, Biological Computation, Living Technology

1. Introduction

Advances in robotics, and in particular humanoid robotics, sometimes pose questions that challenge our ideas about aliveness and personhood, "blurring the boundary between our understanding of living and non-living matter" [1]. Artificial Intelligence is one domain in which such advances have promoted the concept that computer software, and therefore robots, should be regarded as intelligent, an attribute that until comparatively recently was exclusively associated with homo sapiens and, to a lesser extent, with other members of the animal kingdom.

But in the closing decades of the 20th century, certain achievements in Artificial Intelligence enabled computer software to match or even exceed the skill levels of leading human experts. This led Donald Michie to suggest that "systems are emerging that not only outrun the intellectual reach of humans but do so in a way which is opaque to human attempts to follow what they are doing" [2].

Examples of super-achievements within the field of AI include David Cope's program EMI (**Figure 1**), which has composed a complete symphony in the style of Mozart [3].

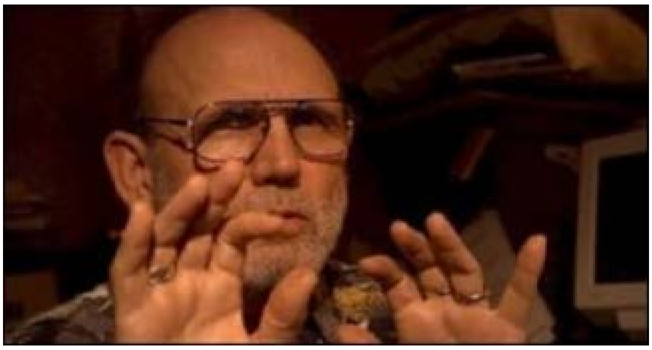

Figure 1. David Cope conducting a concert.

Then there is Harold Cohen's artist program AARON (**Figure 2**), whose works have been exhibited worldwide, and have sold for thousands of dollars in the commercial galleries [4].

Figure 2. One of AARON's paintings.

At one stage, Cohen also built a painting robot which, under the control of AARON, was able to paint the program's creations, even measuring the amount of paint on each brushstroke so that it knew when a pot of paint was empty and needed replenishing.

And perhaps the best known example of all in super-human computation is IBM's chess playing program Deep Blue, which defeated World Champion Garry Kasparov in a match in 1997 (**Figure 3**) [5]. Incidentally, the author has witnessed several games between human players and chess programs, and many of the humans quickly start to refer to their opponent as "him" rather than "it".

Figure 3. Garry Kasparov playing Deep Blue – the final game of the 1997 match.

With achievements such as these being widely publicized, it is little wonder that many people now accept the notion that certain computer software may reasonably be regarded as having intelligence. And when intelligent software is embedded in robots, the question arises—is the robot intelligent? This paper argues that it would not be a huge step from accepting that robots are intelligent to accepting that they are alive.

One of the first pioneers to make this conceptual leap was the MIT Professor Joseph Weizenbaum, the creator of the famous ELIZA program. In his influential 1976 book "Computer Power and Human Reason" Weizenbaum accepted that:

"a modern computer system is sufficiently complex and autonomous to warrant our talking about it as an organism. Given that it can both sense and affect its environment, the author even grants that it can, in an extremely limited sense, be 'socialized', that is, modified by its experiences with its world.

the author grant also that a suitably constructed robot can be made to develop a sense of itself, that it can, for example, learn to distinguish between parts of itself and objects outside of itself, that it can be made to assign a higher priority to guarding its own parts against physical damage than to similarly guarding objects external to itself, and that it can form a model of itself which could, in some sense, be considered a kind of self-consciousness."

Weizenbaum then clarified his position by stating:

"When the author says therefore that the author am willing to regard such a robot as an 'organism' the author declares the author's willingness to consider it a kind of animal." [6]

And by considering a robot to be a kind of animal Weizenbaum was implicitly declaring his willingness to consider robots as being "kind of alive". In 1976 Weizenbaum was a lonely voice in his support for the concept of aliveness in robots.

But by the early years of the 21st century Weizenbaum's voice had been joined by a generation of children who were talking about robots they had seen as being "kind of alive" or "kind of like a person" [7] [8]. It is reasonable to assume that within the next few years, as that generation of children reaches adulthood, they will be more susceptible than most of today's adults to the concept of robots being alive.

That trend will, in turn, cause the word "alive" to be used more and more often in relation to robots, probably causing it to become the norm to regard robots as being alive.

Around the time of publication of Weizenbaum's book, and as a consequence of actual advances and predicted advances in AI and robotics, some voices in the legal arena and other disciplines began to argue for robots to be considered as deserving of personhood, legal rights and civil rights [9]-[13]. In 1972 for example, a distinguished American professor of philosophy, William Lycan, suggested the possibility of bestowing personhood on the robots of the future, and accordingly he coined the compromise term "robot-person". Lycan presented the imaginary example of a robot-person called Harry who

"is humanoid in form—he is a miracle of miniaturization and has lifelike plastic skin—and he can converse intelligently on all sorts of subjects, play golf and the viola, write passable poetry, control his occasional nervousness pretty well, make love, prove mathematical theorems (of course), show envy when outdone, throw gin bottles at annoying children, etc., etc. We may suppose he fools people into thinking he is human. Now the question is, is Harry really a person? Does he have thoughts, feelings, and so on? Is he actually conscious, or is he just a mindless walking hardware store whose movements are astoundingly like those of a person" [9]

Today, some 40 years later, we must acknowledge that most of Harry's capabilities are achievable with current technologies [14] [15]. And one implication of regarding Harry as a person is that then Harry would also be regarded as being alive.

Not surprisingly, many of the opponents of the idea of robot personhood proclaim that robots should not have any such rights precisely because they are not alive. And in the debate about the concepts of love, sex and even marriage between humans and robots, one of the arguments most often put forward against such relationships is that love between a human and a robot is not possible because robots are not alive.

From the author's many discussions on these topics the author have observed that hostility to the concepts of robot emotions, robot personality, robot consciousness, and so on, often seems to emanate from that same belief, the belief that such attributes are not possible in a robot because robots are not alive.

But nowadays that pre-supposition, more and more often, is being implicitly questioned by the physical and mental feats of some robots, for example playing soccer (**Figure 4**), playing the trumpet (**Figure 5**), and dancing (**Figure 6**).

Figure 4. A team of soccer-playing robots.

Figure 5. Trumpet playing robot.

Figure 6. Ballroom dancing robot.

This paper rejects for a moment the premise that robots are not alive and therefore they cannot be regarded as intelligent, having consciousness, and so on. Instead, this paper claims that we should be critically examining the capabilities and attributes of robots and only then asking ourselves whether an entity with such capabilities and attributes should or should not be regarded as alive.

These, then, are some of the controversial aspects of humanoid robotics which have hitherto raised their heads. As computer programs and therefore robots become ever more advanced, being endowed with artificial emotions, artificial personalities, artificial consciousness, and so on, such controversies will grow in number. For example, how should we regard the prospect of robot evolution? After all, the idea and practice of digital evolution are not new. In the 1990s Thomas Ray, a biology professor in , created a software system called Tierra in which self-replicating software evolved by natural selection. At the time Tierra attracted considerable media attention, including an article in "Time" magazine which bordered on the sensational:

"Perhaps Ray's digital beings will set off down the same sort of evolutionary path our species has traveled, only at electron speed. And if that happens, what then? We may find ourselves face to face with an artificial intelligence so thoroughly immersed in the silicon realm, so distant from our curious, carbon-based concerns, that we cannot even hope to converse with it." [16]

And then there is the development of evolvable hardware, reconfigurable devices such as field programmable gate arrays (FPGAs), which can change their architecture and behaviour dynamically and autonomously under the control of evolutionary algorithms such as genetic learning.

With evolutionary software and evolutionary hardware already in existence, the science fiction nightmare scenario of robots evolving to the point where they could take control of the world becomes a little less fictitious, a little more realistic. George Klir and Ladislav Valach point to the possibility of a robot with reproductive capability gaining in knowledge from one generation to the next as a result of the experiences of the older generation.

After assembling a new robot, the older one transfers to the new one not only its original program, which might itself have evolved during the older robot's "life", à la Tierra for example, but the older one also transfers information about its experiences—creating an evolutionary process for its data.

In this way the next generation robot will be superior in some respects to its parent, for example it will be better able to adapt to circumstances that its parent experienced [17]. With each generation of robots benefitting from all the experiences of its ancestors, we can only speculate on how many generations would be necessary before the nightmare scenario is upon us.

And remember, a generation of robots will need far less time to evolve than the 20 - 25 years for human families. A new generation of robots might take only a matter of a few months, a few weeks, or even less.

If robot software and hardware have the capability to evolve, creating a new and better generation of robots, have we not achieved a full technological equivalent of Darwinism? And Darwinism, of course, only deals with the evolution of species, the evolution of the living.

What Does It Mean to Be Alive

Answering this question is a necessary precursor to answering the title of the author's talk—"Are robots alive?" This paper will divide the author's exploration of this question into three parts:

Meanings of "alive" based on traditional biological factors.

Meanings of "alive" based on general factors—Living Systems Theory.

Despite its integrative power, LST has recognised limitations. Its very generality can make it descriptively rich yet mechanistically thin; the twenty abstract functions seldom predict specific dynamics without auxiliary models. Critics also note that subjective phenomena—conscious experience, culture, value—lie outside its scope. Moreover, fitting every new emergent structure (e.g., an immune system or a global stock-market) into the canonical twenty boxes can feel forced, prompting accusations of reductionism. Finally, applying LST analogically across domains (e.g., treating a nation too literally as an "organism") risks category errors or ethical blind spots. Hence, LST remains most fruitful as a high-level organising lens, supplemented by domain-specific theory.

A mobile robot exhibits the same pattern. Cameras and lidars act as Input-Transducers; on-board processors serve as Decider, Decoder, and Associator; batteries and power managers parallel the Ingestor, Converter, and Storage; wheels and servos form the Motor. Diagnostics (Internal-Transducers) trigger self-protective actions when, say, temperature rises. Thus, information and energy streams remain tightly coupled, enabling adaptive autonomy.

In a biological organism fleeing danger, visual and auditory Input-Transducers detect the threat and relay data to the Decider (brain), which consults Memory and then drives the Motor (muscles) while Converter and Distributor rush extra fuel to those muscles. Internal-Transducers monitoring fatigue feed back to the Decider, closing the loop. Every subsystem participates: sensing, deciding, energising, acting, and monitoring.

Living Systems Theory (LST) portrays a living entity as an integrated network of twenty "critical" subsystems. Each subsystem specialises in processing either matter–energy or information, but none act in isolation; instead, they exchange inputs and outputs in nested feedback loops that keep the whole system viable. Sensory units (Input- and Internal-Transducers) feed signals via the Channel and Net to integrative units (Associator, Memory, Decider). Guided by the Timer, the Decider issues commands—through the Encoder and Output-Transducer—to the matter–energy processors (Ingestor, Converter, Motor, Extruder, etc.). This information-to-energy coupling is reciprocal: if Internal-Transducers report low energy, the Decider slows the Motor and activates the Ingestor. The Boundary mediates all exchanges with the environment, and the Reproducer links present dynamics to future generations. In short, coherent system-level behaviour emerges from continual interaction among all subsystems.

2. Subsystem Interactions and Limitations

21st Century Meanings of "Alive"—New Biologies and Living Technology

Meanings of "alive" based on traditional biological factors.

For more than two thousand years biologists, philosophers and others have sought to define life, to provide definitions that would encapsulate all life forms, to specify lists of the qualities which they regard as being essential to life. But these goals have long been remarkably elusive and still remain so. One of the first notable attempts came from Aristotle, who had a considerable interest in biology.

In his treatise on the nature of living things, which was entitled "De Anima" (meaning "on the soul"), Aristotle listed four qualities which he regarded as essential for an entity to be considered alive: sensation, growth, reproduction and self-movement [18]. Since Aristotle's time there have been a plethora of other attempts to provide such lists. For example, a set of seven criteria was proposed by Daniel Koshland Jr. in an essay entitled "The seven pillars of life" [19].

Koshland was an eminent American biochemist, a member of the United States National Academy of Sciences, a long time professor of molecular and cell biology at the of, and for ten years he was editor of "Science". His seven criteria are:

A program—an organized plan for the ingredients and the interactions among those ingredients as the living system goes through its life.

Improvisation—the way in which a living system can change its program as its circumstances and environment change.

Compartmentalization—the container which maintains the interior of the living system and provides the system with protection from the outside.

Energy—produced by the metabolism of energy from the sun and other sources.

Regeneration—the way in which a living system replenishes the thermodynamic losses it undergoes while it is functioning.

Adaptability—a fundamental feedback response that allow the system to survive in quickly changing circumstances.

Seclusion—the separation of different processes that are going on simultaneously within the system, so that no combination of these processes can interact in a way that would have adverse effects on the system.

In order to put into perspective the proliferation of such lists of qualities and such definitions of life, in 2011 Edward Trifonov contributed a paper to the "Journal of Biomolecular Structure and Dynamics" in which he analyzes the vocabulary of no fewer than 123 published definitions of life, commenting that "The definitions are more than often in conflict with one another" [20].

Other researchers have added to the confusion by pointing to various errors in definitions, and in many of the lists of the criteria which others regard as being essential for life. Even the usually accurate publishing house Larousse is not immune to valid criticism on this account. In the section "What is life?" in its 1971 book "Larousse Science of Life" the editors it assert that:

"All living things have in common certain recognizable characteristics: they all have a definite shape and structure; the ability to absorb and transform oxygen and food into substances which become an integral part of themselves, namely, the faculty of assimilation; irritability, the power to survive certain changes in environment; reproduction; ageing and death. Growth, metabolism, and reproduction are three characteristics which are common to all living things." [21]

But as Geoff Simons explained in 1983:

"This plausible-sounding list does not bear scrutiny—worker bees, though manifestly alive, cannot mate; and certain parasitic wasps have been able to dispense with the male entirely, so even the much vaunted capacity for reproduction is not an inevitable quality of all successfully living creatures." [22]

It is not only errors that should cause us to think twice about accepting the accuracy of most of these published lists and definitions. If we compare the lists themselves we often find significant differences between the sets of criteria and/or their published meanings. For example, Koshland's set of seven criteria bears significant differences from the seven biological features in a 2011 paper from the University of Zürich [23], and both of those sets exhibit differences from the seven traits found in Wikipedia's definition of the characteristics of life [24].

Various man-made systems, including computers and robots, conform to a range of the so-called "essential" criteria by which life is recognized. If we had an authoritative definition of the terms "life" and "alive", and an authoritative list of criteria for life, we would be able to provide an authoritative answer to the question "Are robots alive?" But as we can see, within traditional biology no such authoritative list or definition exists.

Little wonder then that Stephan Wolfram has somewhat despairingly stated, in a talk entitled "Computation and the Future of the Human Condition":

"Let's talk about the definition of life.

Well, it's usually pretty easy to tell if something we encounter on Earth is living or not.

The author means, it has all that shared history, with RNA and cell membranes and everything.

But what about a more abstract definition of life?

Well, the Greeks used to think that anything that moved itself must be alive.

But then there were steam engines and things.

The Victorians thought it was something about thermodynamics and the Life Force.

Then people thought it might be something about self reproduction.

But actually all these abstract definitions really don't work.

We can say that there's a necessary condition for life: that the system exhibits sophisticated computation.

But beyond that, there really doesn't seem to be any kind of abstract definition one can give.

The practical definition for us is based on history—and based on the actual historical properties of life on Earth." [25]

So far so bad.

Before the author move on in a different direction This paper will first add to all this confusion by providing a short list of the author's own. This is not a list of criteria for life, because life can and does exist without any of the items on the author's list. Mine are characteristics which, This paper argues, indicate at least some form of life in any entity which possesses any or all of them: intelligence, emotion, personality, and consciousness. These are not characteristics derived directly from pure biology, but they are characteristics created in and controlled by the brain, which is certainly a biological entity.

Functionalist and organizational theories in biology contend that what makes a system alive is not its chemical substrate but the *organization of processes* that maintain the system's autonomy and capacity for adaptive behaviour (Deacon 2021; Bich & Moreno 2022). Within this perspective, cognitive and affective functions—whether realized in carbon-based neurons or silicon circuits—are hallmarks of a self-maintaining, goal-directed organization. When an artificial agent integrates perception, internal modelling, value-laden appraisal, and behavioural regulation, the resulting closed control loop satisfies the same dynamical logic that living organisms employ to survive and flourish.

Recent studies demonstrate this convergence. Chella and Manzotti (2023) review machine-consciousness architectures that couple synthetic proprioception with global workspace mechanisms, thereby enabling robots to monitor and modify their own internal states. In Nature Communications, Rizik et al. (2022) show neuromorphic genetic circuits that merge affective computation with metabolic regulation inside living cells, blurring the distinction between engineered emotion and biological homeostasis. These advances support the claim that artificial intelligence, emotion, personality, and even rudimentary consciousness are not mere surface metaphors but functional indicators of lifelike organization.

Thus, when a robot exhibits adaptive intelligence, affective appraisal, individually stable yet plastic personality traits, and self-reflective monitoring, it meets key organizational criteria put forward by contemporary theoretical biology. Under this functionalist lens, such attributes provide prima facie evidence of a genuine—though non-traditional—form of life, strengthening the paper's argument that lifelikeness can be evaluated by systemic functionality rather than biochemical composition.

3. Linking Artificial Cognitive-Affective Attributes to Life

3.1. Some Lifelike Characteristics

Artificial Intelligence first emerged as a science in the mid-1960s, and the author has already mentioned a few examples of the outstanding achievements of AI. In addition to those achievements researchers have also developed computer software and robots that exhibit emotions—they are artificial emotions, but emotions nonetheless. Similarly, software and robots have been developed that exhibit personalities with recognizably humanlike characteristics. And even in the case of consciousness, the most ethereal of the items on the author's list, for more than a decade, researchers have been seriously investigating the creation of artificial consciousness, for example Owen Holland, Professor of Cognitive Robotics at the University of Sussex [26].

In 2001 the Swartz Foundation organized a symposium at Cold Spring Harbor Laboratory on the question "Can a machine be conscious?", which was attended by a number of leading computer scientists, neurologists and philosophers. Many diverse views were expressed during the symposium, but a consensus was reached that:

"In principle, one day computers or robots could be conscious. In other words, we know of no fundamental law or principle operating in this universe that forbids the existence of subjective feelings in artifacts designed or evolved by humans." [27]

It is arguable that robots can already demonstrate consciousness. In the 1970s Gordon Gallup devised the "mirror test", to determine whether or not animals are able, as humans are, to recognize themselves in a mirror [28]. conducted the test with chimpanzees and orang-outans, and discovered that within only two or three days the animals developed the ability to recognize themselves.

To conduct the test a visible coloured spot is placed on the animal's skin or fur near its forehead, and the animal observed to see if it attempts to touch the spot or to remove it. Any such attempt is taken as an indication that the animal recognizes the creature it sees in the mirror to be itself – a sign of consciousness.

Amongst the non-human creatures that have been observed to pass this test are bottlenose dolphins, pigeons and elephants [29]. The test has also been tried on robots, by Junichi Takeno and his team at in, who announced in 2005 that they had succeeded in achieving mirror image cognition for a robot [30].

This paper makes use of an extension to Alan Turing's famous test for intelligence in a machine [31]. To summarize and paraphrase Turing, if a machine exhibits behaviour that is normally a product of human intelligence, imagination for example, or by recognizing sights and scenes and music and literary style, then we should accept that that machine is intelligent.

Similarly, this paper argues that if a machine exhibits behaviour of a type normally regarded as a product of human consciousness (whatever consciousness might be), then we should accept that that machine has consciousness. If it exhibits emotional behaviour of a type normally exhibited by humans, we should regard it as having emotions. And if it exhibits humanlike personality traits then we should regard it as having personality [14].

You can probably see where this is leading.

If we were to employ a comprehensive list of characteristics of life as our yardstick for answering the question "Are robots alive?", then the characteristics on the author's own list would be perfectly valid inclusions; not as characteristics which must, by definition, be possessed by an entity that is alive, but as characteristics the presence of which would indicate that the entity in question is, in some sense alive, or at the very least it might be alive.

Extending Turing's argument still further, what the author am claiming is that if an entity has intelligence as though it were alive, and/or if it has personality, emotion and consciousness as though it were alive, then we should regard it as being alive. If it looks like a live duck, swims like a live duck, and quacks like a live duck, then probably it is a live duck.

The author am not claiming that we should regard robots as being alive solely on the basis of their possessing one or more of the four attributes on the author's list, but the author want you to accept that the author's list and this particular argument serve as additional reasons to demonstrate that such lists are fraught with problems and have the capacity to create confusion rather than clarity.

Let us now abandon any attempt to answer the question "Are robots alive?" from the perspective of traditional biology. In the next part of our exploration we delve into an analytical system which, although inspired by biology, is actually non-biological in its nature—a general theory which allows us to consider the key question from a non-biological perspective.

Meanings of "alive" based on general factors—Living Systems Theory

In 1978 the renowned American biologist James Grier Miller (**Figure 7**) published a groundbreaking magnum opus entitled "Living Systems", a 1102 page volume in which he described what has become known as Living Systems Theory [32].

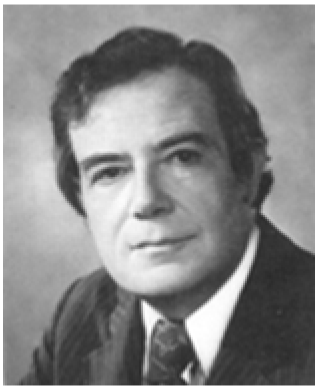

Figure 7. James Grier Miller.

Miller had a distinguished career which included: serving on the faculty at Harvard, he was Chairman of the Department of Psychology at the University of Chicago, President of the Society for General Systems Research, and for more than 30 years he was editor of the journal "Behavioral Science"—the leading academic journal reporting on the systematic investigation and analysis of and behaviour.

Miller's book represented the state of Behavioural Science theory in the late 1970s. His work provides us with a means by which we can compare the various subsystems that he identified as being in all living systems, with various of the components and their attributes found in robots. Such a comparison will allow us to answer our question "Are Robots Alive?" from the perspective of Behavioural Science as it was at the time Miller's magnum opus was published.

Living Systems Theory is a general theory about the existence of all living beings, and the author stresses the word "general" here for reasons that will soon become apparent. Miller's suggested basic strategy for observing and investigating living systems was to trace the inputs, throughputs and outputs of living systems "through various subsystems and components in which they flow in sequence, measuring various variables of their transmissions as they occur" [32]. This strategy reveals how a system is functioning.

Miller comments that a "wide range of sorts of matter and several sorts of energy are essential in varying amounts for various living systems". In robots the range of types of matter is very much narrower than it is in the living systems which Miller studied. In robots the relevant types of matter are mostly those that can be converted into energy, as for example with the EcoBot-II robot which This paper will show you shortly. There is also a very narrow range of types of energy which are appropriate for robots, principally the various forms of electrical energy.

Miller intended his work to formalize the concept of life, to reveal general principles for how all living systems work and to surmount the problem which the author have just discussed and which had long been faced by biologists – how to distinguish clearly between living and non-living matter. As we have seen, many generally accepted forms of life do not conform to at least some of the lists o criteria specified by biologists, which is one reason why Miller's general principles are so useful in identifying living systems.

In the first edition of "Living Systems", in 1978, Miller specified that a living system must contain each of nineteen "critical subsystems" as he called them, features of life. Subsequently he added a twentieth subsystem. Miller also showed how biological and social systems are organized and operate at a number of different hierarchical levels, including cells, organs composed of cells, and organisms at the lower end of the hierarchy, to groups of people such as families, organizations and nations at the upper end. In his book Miller presents an integrated analysis of the nature of all biological and social systems.

Each subsystem consists of "one or more components that carry out one of the system's essential processes" [33]. And each of Miller's subsystems is defined by its functions rather than by any biological attribute or characteristic. One notable aspect of his list of critical subsystems is that all of them are independent of their method of construction. Miller did not assume that life must be based on hydrocarbons or any other specific chemical compounds.

He simply claimed that if a system can reproduce, and if it can handle energy and information appropriately, then that system can justifiably be regarded as a living system. A corollary is that how the system was created is not relevant.

Whether a system is created by nature, or assembled mechanically, or manufactured on a flow solder machine, is of no consequence if its internal operation and its behaviour in the world meet Miller's criteria.

Given Miller's career as an eminent biologist, and the fact that his book "Living Systems" is replete with references to and evidence from the biological sciences, you might well be surprised to learn that the author am going to discuss Miller's work as being independent of traditional biology. My summaries of the processes described by Miller in each of the 20 subsystems employ no references to biology.

The reason for this lies in the generality of Miller's analysis and of Living Systems Theory as Miller presented it. Despite his frequent references to matters biological, Miller did indeed present a general theory rather than a biological theory. Living Systems Theory encompasses all living systems, which is not the same as all biological living systems when the word "biological" is being employed in its traditional sense, meaning created by nature.

The author should mention here that, since the late 20th century, scientists have been researching a handful of what This paper will refer to as "new biologies", non-traditional biologies. More about these new biologies later in this talk.

Living Systems Theory was developed by an interdisciplinary group of scientists whose goal was to achieve scientific unity via the development of general theories. Its unifying factor, according to Miller, is that

"because of the evolutionary relationship among them, all living systems have similar requirements for matter and energy, without which they cannot survive. They must be able to secure food, fuel, and other necessary inputs. They must process their inputs in various ways to maintain their structure, reproduce, make products, and carry out other essential activities." [32]

In his descriptions of the 20 subsystems Miller frequently employs the joint term "matter-energy". Matter is anything which has mass and occupies physical space. Energy is the ability to do work. Mass and energy can be converted into each other in accordance with Einstein's equation $E = mc^2$, meaning that energy is equal to mass multiplied by the square of the speed of light.

Miller employs the joint term matter-energy throughout his book, except where one or the other is specifically intended. He explains that "Living systems need specific types of matter-energy in adequate amounts. Heat, light, water, minerals, vitamins, foods, fuels, and raw materials of various kinds, for instance, may be required." Eight of Miller's 20 critical subsystems are related to the processing of matter-energy, which is of great importance to living creatures.

In addition to matter-energy, Miller's descriptions of the subsystems also make frequent use of the term "information". He explains that inputs, internal processes, and outputs of various information signals represent the information metabolism of living systems, and that:

"in all such systems information flows through several subsystems, each one being a transmitter to the next one in the sequence, which acts as a receiver. Then, after processing the information in its own specialized way, the second subsystem acts as a transmitter, sending the information on to a third." [32]

Miller also noted that living systems maintain their relationships with their environments by inputting and outputting information.

During the next part of the author's talk, this paper will summarize each of the 20 "essential processes" specified by Miller, with the aim of demonstrating that robots can engage in all 20 of them.

The function of Miller's subsystems numbered 1 and 2 is to process both matter-energy and information. The function of subsystems 3 - 10 is to process only matter-energy. For subsystems 11 - 20 the function is to process only information.

In considering each of the 20 "critical subsystems" we start with Miller's own definition and summary explanation.

3.1.1. The Reproducer—The Subsystem Which Is Capable of Giving Rise to Other Systems Similar to the One It Is in

The reproducer provides information that specifies the structure of a new system and mobilizes the matter and energy needed to make it. The information is carried in the genetic material at the levels of cell and organism. The information that specifies the structure and processes of systems above the organism is contained in a written or implicit charter.

Miller described the Reproducer subsystem as being:

"critical in a different sense from the others. It is essential if the species is to continue to another generation, though it is not needed for the survival of the system in which it exists. For example, castrated animals can continue to live. The processes of all the other subsystems appear to be critical to the survival of living systems at all levels."

This is a point worth emphasizing – a system can be alive even if it cannot reproduce.

Before the author discuss reproduction This paper will make a brief diversion to mention growth, which is another characteristic of biological life included in some definitions of life and in some of the lists of criteria for life. We normally think of growth as meaning an increase in physical size, possibly accompanied by the appearance or activation of additional organs.

Self-reproducing robots, such as the one This paper will show you shortly, can grow by increasing the size of their progeny, adding more parts such as an additional memory chip. Since robots can already, as we shall soon see, add parts to themselves, we already we have sufficient justification for stating that robots can grow. Not gradual analog growth as we know it in flora and fauna, but stepwise growth.

Now to reproduction.

Serious scientific interest in the self-reproduction of machines began with John von Neumann's research into automata theory during the mid-1950s [34]. Any study of the early history of the theory of self-reproducing machines is virtually the same as the study of von Neumann's thinking on the subject. His research into automata had the goal of modelling biological self-reproduction. He set out to investigate the logical organization of self-reproducing machines, and when considering what capabilities should be demonstrable in any machine that was claimed to be able to replicate itself, von Neumann listed three such capabilities:

Logical universality—the ability to function as a general-purpose computing machine;

Construction capability—the ability to manipulate information, energy, and materials of the same sort of which the machine itself is composed;

Constructional universality—the ability to manufacture any machine which can be formed from specific types of parts;

Before going any further the author should point out that living systems do not all reproduce in the same way. Samuel Butler, in his 1863 article "Among the Machines", suggested the possibility that machines are a kind of "mechanical life", and proclaimed that "We should not be surprised if new types of life develop new reproductive modes, the pattern best suited to their nature." [35] The actual method by which reproduction takes place is not in any way crucial. The only crucial question is whether robots can, in some way, reproduce themselves, creating new robots.

Von Neumann concluded that self-reproduction is possible if the above three capabilities are achieved. His argument was that, because the original machine is made of parts that can be manufactured, and the original machine can be constructed, and the original machine is given a description of itself, then it ought to be able to make more copies of itself using the manufacturable parts. He envisioned a machine that lived in a "sea" of spare parts.

That machine had a program stored in its memory, a program that instructed the machine to go through certain mechanical procedures. The machine had an arm, or something that functioned very much like an arm, and the machine could move around its environment. By using its arm the machine could pick up and connect whichever of the spare parts it wished.

The program first instructed the machine to reach out and pick up a part, then to go through an identification procedure in order to determine whether or not the part selected was the one the machine had been told to locate.

(If not, the part would be thrown back into the "sea" and another part picked up and examined in the same way, and so on, until the correct part was found.) Having found the part it was looking for, the machine would then start on its search for the next part that it needed, as determined by the program stored in its memory.

The machine would continue following the instructions to make something, without really understanding what it was doing. When it finished executing its program the machine would have produced a physical duplicate of itself.

But this newly minted machine would not yet have any program in its memory, so the "parent" machine would copy its own program, from its own memory, into the memory of its offspring, and then, finally, the original machine would start up the program residing in its progeny. This whole process, as described by von Neumann, is logically very close to the way that living organisms reproduce themselves [14].

Although von Neumann did not make this point in his description of the reproductive process, an equivalent process would be possible if and when the machine's program determined that the machine had developed a fault and needed to be repaired, possibly by the replacement of a faulty part.

Fast forward just over 30 years from von Neumann's publication of his research, and we come to the first tangible examples of self-reproducing robots, as part of the Golem project at, which started in 1998 and was led by Hod Lipson and Jordan Pollock. Their "parent robots" consisted of a computer running an algorithm that simulated evolution and produced a design for new robots based on trial and error. This was linked to a 3-dimensional printer that made small plastic shapes. The offspring were small plastic trusses (made of bars), with motors that were controlled by neural networks (**Figure 8**).

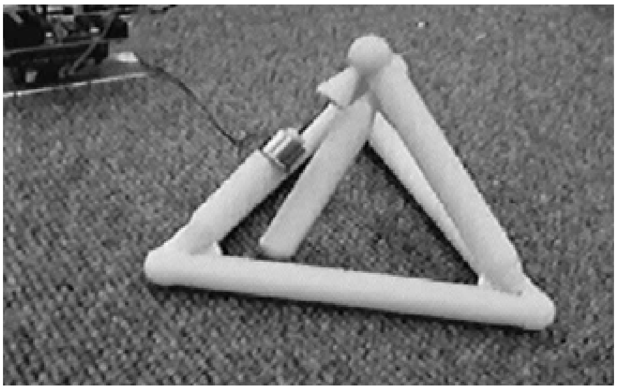

Figure 8. The offspring.

Artificial neurons were the building blocks of their control modules (their brains). The bars could connect to each other with free ball-and-socket joints, and the trusses could potentially make up articulated structures such as the one in this photograph. Neurons could connect to other neurons through synaptic connections, and neurons could connect to bars. Human intervention was necessary only to attach the motors and connect the wires—the robots did all the rest of the assembly themselves, including telling the humans what to do [14]. These were the first examples of robots that themselves had been robotically designed and robotically fabricated—the first steps in robot evolution.

Subsequently, in 2005, Lipson and colleagues at developed robots whose offspring were exact copies of themselves. These were simple machines that acted as autonomous modular robots and were capable of physical self-reproduction using a set of cubes [36]. In order to reproduce itself their machine required a supply of materials, namely 10 centimetre cubes each of which had a single rotating joint as well as two points where they could connect to other cubes.

One half of each cube could swivel relative to the other half, and connected cubes could be formed into arbitrary arrangements. A four-module robot was able to construct a replica of itself in 2½ minutes by lifting and assembling cubes from supply locations on top of small plinths. A three-module robot could perform a similar feat in little over 1 minute.

A video of the process (speeded up by a factor of 4) showing a four-module robot reproducing itself, is available at [37].

With advances in 3D-printer technologies it is not difficult to envisage the day when it will be possible for a robot system, based on such technology, to replicate all of its components and then to assemble them into an exact replica of itself.

3.1.2. The Boundary—The Subsystem at the Perimeter of a System That Holds Together the Components Which Make up the System, Protects Them from Environmental Stresses, and Excludes or Permits Entry to Various Sorts of Matter-Energy and Information

The boundary separates the system from its environment. It surrounds and protects vulnerable components, acts as a barrier to free movement of matter, energy, and information in and out of the system, and filters inputs and outputs by allowing some but not others to pass.

A robot is a self-contained entity contained within the boundary that is its outer covering, no matter whether that covering be some sort of manufactured material or a combination of such materials. Since robots conform to this simple description they satisfy Miller's boundary criteria.

3.1.3. The Ingestor—The Subsystem Which Brings Matter-Energy across the System Boundary from the Environment

In "Living Systems" Miller identified the input arteries to bodily organs, the gaps in cell walls, and the mouths, noses and gills of organisms, as some of the input viaducts for the matter-energy of living organisms. Today we point to the electrical cables and connectors, and to the other electrical conduits and methods employed for conveying electrical energy from outside a robot system.

Whether that electrical energy comes from mains power, or batteries, or solar cells, or some other source, is not important. The source of a robot's power could even be "food" which it converts into electrical energy, as in the fly-eating EcoBot-II (**Figure 9**) robots developed at the University of the West of England by Chris Melhuish and his team [38].

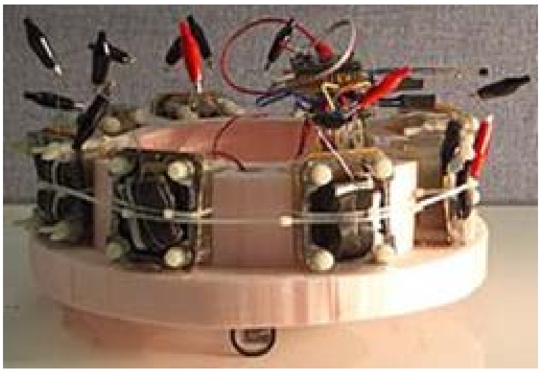

Figure 9. The EcoBot-II robot.

In their case the name "ingestor" seems particularly appropriate—the dead house, rotten fruits and shells ingested by the system provide the fuel for bacteria in a microbial fuel cell, and this cell acts as an artificial metabolism converting the fuel into electrical power. It is easy to see how the ingestor is implemented in computer systems and in robots.

The computer takes in electrical energy from an external source - normally it does so from the grid network or a battery—using various types of electrical conductors such as the tracks on a printed circuit board.

3.1.4. The Distributor—The Subsystem Which Carries Inputs from Outside the System, or Outputs from Its Subsystem around the System to Each Component

We humans have veins and arteries which distribute our blood around the body. In robots there are electrical connections such as the tracks on printed circuit boards, pins on computer chips, and other conductive components which allow electrons to be distributed around the system, providing power to the parts of the robot that need it and conveying data from one component to another.

3.1.5. The Converter—The Subsystem Which Changes Certain Inputs to the System into Forms More Useful for the Special Processes of That Particular System

The nature of the converter in a robot will depend on the form of the input to the system and on what form the converted input needs to take in order for the robot to operate correctly. In the case of the Ecobot-II fly-eating ingestor, the bodies of flies and the other bacterial cultures derived from sewage sludge are converted into electrical energy, so the converter is actually part of the ingestor.

In the case of electrical power being provided as the system's input, the input voltage and current will usually undergo various conversions that raise or lower them to the levels required by the different electronic components in the robot system. What is common to all forms of conversion is that they change matter-energy into the forms needed by the system.

3.1.6. The Producer—The Subsystem Which Forms Stable Associations That Endure for Significant Periods among Matter-Energy Inputs to the System or Outputs from Its Converter, the Materials Synthesized Being for Growth, Damage Repair, or Replacement of Components of the System, or for Providing Energy for Moving or Constituting the System's Output of Products or Information Markers to Its Suprasystem

Miller provided the following examples of the physical bonds that maintain these associations and hold together various types of artifacts: nails, bolts, welding, cement and stitches. In the case of robots obvious additions to this list are the soldered connections which hold many of the electronic components in place.

3.1.7. Matter-Energy Storage—The Subsystem Which Places Matter or Energy at Some Location in the System, Retains It over Time, and Retrieves It

Generally speaking, matter is not stored in a robot beyond that needed to lubricate moving parts. But in the case of some robots such as the EcoBot-II, biological matter does need to be stored so that it can be converted to create electrical power.

In robots electrical energy is the primary form of energy in which we are interested. But not all robots require electrical energy to be stored, for example a vacuum cleaner robot only needs energy when it is performing its cleaning duties and it can therefore run entirely off mains power. When it is necessary to store electrical energy in robots, this task is accomplished by batteries and/or capacitors.

The author should also mention that the software incorporated in a robot, the computer programs that instruct and enable the robot to operate, can be viewed as a form of energy as Miller defined it, because they provide some of the robot's ability to do work, but such philosophical arguments are beyond the scope of this talk.

3.1.8. The Extruder—The Subsystem Which Transmits Matter-Energy out of the System in the Forms of Products or Waste

Miller characterizes the products of living systems as being forms of matter-energy which are transmitted to beyond the boundary of the system. In the case of robots the most tangible products exist as output data of various types, including text and speech. Over and above a robot's tangible products there are also what we might call its "emotional products", for example its smiles, its frowns, and other behaviours that convey a different type of energy to users—emotional energy. And a robot's most obvious waste product is the heat dissipated by some of its components.

A robot's output data "extrudes" from the system via various output devices within the robot, such as loudspeakers and electrical displays (e.g. LCDs), or via what might be considered as extensions of the robot—cables or wi-fi technology for example. In Miller's terminology it is those loudspeakers, cables, displays and wi-fi carrier waves that are the extruders. The heat which is a robot's waste extrudes into its surroundings through its contact with air.

3.1.9. The Motor—The Subsystem Which Moves the System or Parts of It in Relation to Part or All of Its Environment or Moves Components of Its Environment in Relation to Each Other

The applicability of this subsystem to robotics is so obvious as to require no comment here.

3.1.10. The Supporter—The Subsystem Which Maintains the Proper Spatial Relationships among Components of the System, So That They Can Interact without Weighing Each Other down or Crowding Each Other

Here Miller is referring to the physical arrangement of the various components in a system. Again, the applicability of this subsystem to robotics is so obvious that it requires no further comment here.

3.1.11. The Input Transducer—The Sensory Subsystem Which Brings Markers Bearing Information into the System, Changing Them to Other Matter-Energy Forms Suitable for Transmission within It

The information input "markers" referred to by Miller are forms of matter-energy that can act as information carriers, conveying information from the environment into a system. Light and sound are examples of such markers, as are any electronic media through which digital data can pass.

An input transducer receives information and converts it into a form appropriate for transmission into the system. An example of this process in a robot is a speech recognition system that recognizes the words being spoken by a user and converts that information into digital representations of the strings of letters making up the spoken words. These strings of letters can then be operated on by the system's software, for example carrying out commands spoken to the robot by the user.

3.1.12. The Internal Transducer—The Sensory Subsystem Which Receives, from Subsystems or Components Within the System, Markers Bearing Information about Significant Alterations in Those Subsystems or Components, Changing Them to Other Matter-Energy Forms of a Sort Which Can Be Transmitted within It

An example of such a transducer in a robot system is a battery-low detector which monitors changes in the power remaining in a robot's batteries and produces an indication, either to inform the user that the batteries need recharging or changing, or to instruct the robot to plug itself in to a mains supply for recharging itself.

Other forms of internal monitoring of a robot systems transduce the forms of other types of information relating to the internal state of the robot, rendering the information into forms suitable to be acted upon by the robot. An example is monitoring the internal temperature of the robot so that a fan may be activated if the robot's innards get too hot for some of the electronic components to function reliably.

3.1.13. Channel and Net—The Subsystem Composed of a Single Route in Physical Space, or Multiple Interconnecting Routes, over Which Markers Bearing Information Are Transmitted to All Parts of the System

A single track or a collection of tracks on a printed circuit board are examples of what Miller means by the term "channel". These tracks enable the system to convey the electronic signals that act as data for the system.

3.1.14. This Information Signals the Decider of the System or Deciders of Subsystems to Start, Stop, Alter the Rate, or Advance or Delay the Phase of One or More of the System's Processes, Thus Coordinating Them in Time

In a robot or any other computing device there is always a clock function which can be accessed by the system's software. This provides any timing information that the system needs.

3.1.15. The Decoder—The Subsystem Which Alters the Code of Information Input to It, through the Input Transducer or Internal Transducer, into a "Private" Code That Can Be Used Internally by the System

The input transducer and the internal transducer both alter the matter-energy form of a marker which carries information. The decoder itself alters the code in which the information carried by the marker appears. The decoding process might, for example, be transducing text from one language to another, a process we more commonly refer to as translation.

3.1.16. The Associator & The Memory

We consider the Associator subsystem and the memory subsystem together, because together they provide the memory capabilities of a system.

The Associator is the subsystem which carries out the first stage of the learning process, forming enduring associations among items of information in the system. The Memory is the subsystem which carries out the second stage of the learning process, storing information in the system for different periods of time, and then retrieving it.

In a robot these memory capabilities are provided by electronic memory components, such as hard disks, random access memory chips (RAM), and read-only memory devices (ROM).

3.1.17. The Decider—The Executive Subsystem Which Receives Information Inputs from All Other Subsystems and Transmits to Them Information Outputs for Guidance, Coordination, and Control of the System

Miller did not specify how the decider should function, so although deciders in brain cells are organically based, deciders in other types of system could be mechanically or electronically constituted. A robot is controlled by a computing device (or more than one), usually a micro-controller or a central processing unit (CPU) which acts as the robot's "brain". This brain processes information in the form of data that comes from a variety of inputs, and carries out the programmed commands in the robot's software, thereby deciding how the robot should behave and what tasks it should carry out.

3.1.18. The Encoder—The Subsystem Which Alters the Code of Information Input to It from Other Information Processing Subsystems, from a "Private" Code Used Internally by the System into a "Public" Code Which Can Be Interpreted by Other Systems in Its Environment

Computing based systems, including robots, deal with information and data encoded in various ways. Ultimately all such information and data can be reduced to binary, ones and zeros, but all sorts of higher level data representations are possible. Speech, for example, can be encoded as a succession of data items each comprising the amplitude and pitch of a human voice at a particular point in time.

Another example is text, often represented as integers on the range 0 - 255, where each letter, digit and punctuation mark etc., is represented by one of these 256 integers. The encoder as described by Miller serves the purpose of translating information from one data format, which in Miller's writing is referred to as a "code", into a different format, a different code, as required by the next subsystem that will operate on or use that information.

In a robot or other computing system such translations are specified and controlled by the system's software.

When taken together, Miller's subsystems 15, 16, 17, 18 and 19 constitute the processing of the program of a living system, in other words – the execution of the program. As to the programming itself, the author take issue with those who claim that computers and robots "only" decide to do as they are told to do (i.e. as they are programmed to do), and can not, therefore, be regarded as making their own decisions. My argument here is based partly on the fact that we humans and other animals are also programmed, in our case by nature, and partly because of the author's position on what is commonly known as free will.

When we talk about free will in humans we really mean freedom to choose and make decisions in whichever ways we wish. It is often argued that, because human beings are responsible for the programming of computers and robots, it is fallacious to speak of computers and robots exercising free will. This is one of the more common arguments raised against the concept of robots being deserving of rights and personhood that they cannot exercise autonomy or free will. But computers are already able to write their own programs, and then modify and evolve their own programs, which capabilities, this paper claims, make it less easy to argue that robots are unable to exercise free will.

In discussing the subject of free will in robots, John McCarthy distinguishes between having choices and being conscious of these choices. "Both are important, even for robots, and consciousness of choices requires more structure in the agent than just having choices, and is important for robots" [39]. Manifestly computers do make choices, but does this mean that they have free will? After all, the choices made by computers are often explicitly determined within a program.

But some of the decision-making mechanisms employed in programs are capable of the same type of flexibility as those evident in humans, incorporating a balance between predictable choices determined solely by the logic expressed within a program and the unpredictable choice of randomness, which can also be programmed but is, nevertheless, unpredictable. This balance is the aim of the ARASEM software architecture (Artificially Random Self-Motivation), described by Frank Da Costa [40].

ARASEM was a simple algorithm wherein a robot pet was programmed, sometimes to make random decisions as to the pet's actions and sometimes to make decisions in response to certain stimuli, for example reacting to a certain sound. When the robot is not compelled to respond to a stimulus, the randomness in its choice of actions corresponds to free will.

McCarthy concludes that an intelligent robot which performs at the level of a human being requires the ability to reason about its past, present, future and the choices it has at its disposal. If the robot's decision making is not deterministic, when its computations include some random events so that we cannot be sure which way it will decide, then the computations in the mind of the robot themselves have random and non-deterministic interactions, resulting in the robot having free will. And with free will, robots will not necessarily be constrained to tell the truth (as they believe it) all the time [14].

3.1.19. The Output Transducer—The Subsystem Which Puts out Markers Bearing Information from the Systems, Changing Markers within the System into Other Matter-Energy Forms Which Can Be Transmitted over Channels in the System's Environment

As its name implies, the output transducer operates on information within a system and converts that information into a form appropriate for transmission to outside the system. An example of this process in a robot is a text-to-speech system that converts a text representation of what the robot is saying into human speech sounds that can be readily understood by someone interacting with the robot.

3.2. What Does This Comparison Indicate?

Characteristics of each and every one of Miller's critical subsystems can be found in robotic systems. A corollary of this is that robots can possess the critical functions represented by the 20 subsystems. Miller's thesis in 1978 was that all of the subsystems he specified are essential components of all living systems. He had thus prepared the way for us to recognize that emerging animate machines could be regarded as living systems.

By providing a general integrated scheme that was applicable to all biological and social systems, Miller had helped us to understand the "living systems" status of various artificial systems. What Miller had intended to be a general theory of living systems, has also become an eminently arguable case to support the claim that robots should be regarded as living systems.

If we accept that Miller's original conception of Living Systems Theory is correct, then since his 20 critical subsystems can all be found in robots we should surely accept, on the basis of Miller's work, that robots can justifiably be considered to be living systems as Miller understood and espoused the term.

3.3. 21st Century Meanings of "Alive"—New Biologies and Living Technology

Up until the late 20th century traditional biology was the only biology. The scope of the traditional biological sciences at that time incorporated: human biology, animal biology, plant biology, marine biology and many other branches, but not the biologies of any entities such as the fully functioning man-made reproducing cell controlled by synthetic DNA, which Craig Venter and his team announced in 2010.

Then came hi-tech and the realization that there could be other biologies—different types of biologies including synthetic biology [41]. Since whatever aliveness exists in a robot is driven by computer software, when discussing an appropriate biology that might provide the basis for a 21st century answer to the question "Are robots alive?" we are seeking a discipline that explains how nature provides us with a means for computation which is not, fundamentally, man-made.

Here the author is not talking about computational biology, which is the application of computer systems to the solutions of specific biological problems [23]. Nor am the author talking about biologically-inspired computing, which employs ideas from biological systems such as the brain in order to develop new algorithms for use in computer science applications. Similarly excluded are biomimetic systems, which employ other new technologies that use principles derived from the study of biological systems [42].

One discipline that seems to the author to be closer to the author's goal than these is synthetic biology, a field described as being

"at the interface between biotechnology, chemistry, engineering and computer science …. This multidisciplinary and multi-approach field has the unifying goal of producing and designing new forms of life …. Interestingly, the aim of producing novel types of living organisms in synthetic biology not only implies the production of living from non-living matter, but also the idea of using living matter and turning it into machines which are traditionally considered non-living." [1]

But synthetic biology does not tick the most essential box of all in the discipline the author am seeking—that of computational capability.

Another discipline that might appear to provide computational capability created from living matter is biological computation – a recently conceived science that examines how, under what circumstances, and why, living systems such as organisms can be regarded as being fundamentally computational [43].

The science of biological computation studies such questions as whether organisms and other living systems can perform computations, and if so, can such living systems replicate or even outperform the logic and processing capabilities of man-made electronic systems and their software?

Here the author am prompted to ask a further question – why, in the 21st century, can the logical and computational processes traditionally satisfied by electronics and software, and less efficiently (in many ways) satisfied by the human brain, why can these technology processes not have a "biology" of their own, a "robot biology" perhaps, or a "technology biology"?

Biology is the science of life and living organisms, so "technology biology" would be the science of robot life or of living technology. In her paper entitled "Biological Computation" Melanie Mitchell comments that:

"the field of computation itself has been evolving to become more "biological", with the shrinking of computing elements to molecular scales, and the increasing focus in computer theory of biological computation will be a foundation not only for biology but also for a new era of more life-like computers" [43].

Up to now the most dramatic results of research into biological computation have been in the field of DNA computation. The idea behind DNA computing is to take advantage of the many different molecules of DNA to investigate many different possible solutions to a problem simultaneously. This enables DNA computers, smaller than any other known computers, to perform some types of parallel computation faster than any other computing system.

In 2002 Leonard Adelman demonstrated a DNA computation system in a test tube that could solve a version of the famous travelling salesman problem [44]. Adelman's work acted as encouragement to other researchers to develop logic circuits based on DNA. For example, Erik Winfree's group at CalTech has developed digital-logic circuits in test tubes that connect biological logic gates capable of performing some of the fundamental operations of computing: AND, NAND, OR, exclusive OR, and exclusive NOR [45].

And last year a team of researchers at created, from DNA and RNA, what they call a "transcriptor"—a biological transistor [46]. As one commentator notes.

"With these gates a biological computer could perform almost any computation inside a living cell."

In a parallel research effort Harvard scientists have devised a method to enable DNA strands to act as massive data storage repositories, cramming 700 terabytes of data into a single gram of DNA. This means that the most essential building blocks of a biological computer, the means of computation and the memory storage, are now in place, which heralds an astounding turnaround in computer science and therefore in robotics.

Whereas artificial intelligence, artificial emotion and artificial personality enable computer software to perform tasks and exhibit attributes that are rooted in the human brain—computers imitating nature—we are now at the dawn of an era in which nature, in the form of RNA and DNA, is collaborating on the creation of the fundamental computational building blocks of computer systems and robots—nature creating computers.

And so, with the advent of computer systems that have DNA at their core, we can now address the question "Are robots alive?" from a fresh perspective. We ask the question "Is DNA alive?", with the intention that if the answer is "yes" it is eminently reasonable to argue that robots are alive. But is the answer to the DNA question "yes"? Certainly many science texts state that DNA is alive, and many educated people believe so [47].

After all, without DNA there could be no presence of life. At the of , , an education experiment revealed that 63 per cent of a sample of students believed that DNA is alive. But from a purely biological perspective DNA is not alive, it is merely a large chemical molecule and an essential constituent of all living things. And if DNA is not alive then is DNA computing a living technology?

To answer this, and to finally arrive at the answer of our title question "Are robots alive?", we now explore the new era of living technology.

4. Living Technology

4.1. Definition

In 2001, even before the first DNA computer was developed by Adelman, a small group of scientists met in Ghost Ranch, to discuss the coming emergence of a new science which they named "Living Technology". The prime instigators and investigators of this group have been Mark Bedau, John McCaskill, Norman Packard and Steen Rasmussen. They describe the concept of living technologies as follows:

Bedau

"any technology that is powerful and useful because it is alive or, at least, has certain important properties of living systems [48] … coined to capture the technological implications of our increasing ability to engineer systems whose power is based on the core features of life [49] … to synthesize living or intelligent systems from nonliving components … Living technologies are clearly 'artificial' in the sense of being created through intentional human activity. At the same time, they are clearly 'natural' in the sense that they grow and adapt and evolve autonomously; that is, they have a life of their own." [50]

McCaskill

"technology that exhibit the core properties of living systems that distinguish them from inanimate ones…. To put it simply… 'living technology' is technology that is alive." [48]

Packard

"technology that derives its functionality and usefulness primarily from its living properties… Living technology is, in fact, really alive, if it has appropriate living properties, and in particular if it is capable of evolving. This is a rather extreme position regarding the meaning of 'life' and 'living', but the author believe that as we learn more about life and living properties through study and development of living technology, regarding it as 'real life' will become ever more natural." [48]

Rasmussen

"the author believe it is critical for us to start paying attention to the systemic and increasing lifelike properties of the many manmade systems we usually do not consider to be alive. We need to start relating and contrasting these manmade systems to what we usually consider as being alive, because the distinction between the two will increasingly blur and eventually merge in the coming century." [48]

4.2. Current Examples

A report in 2011 from the European Commission on Information, Society and Media describes the challenge facing Living Technology as being:

"to create controllable yet evolving living technologies that co-organize information and matter in collective systems of physical entities. These systems will include an integration of self-assembly and self-reproduction, and have an ability to balance fully open-ended evolution with more constrained and directed evolution, so as to maintain control over outcomes. This will allow us to learn how to exert flexible and useful control over evolving living technologies." [51]

And the Commission further asserts that

"The creation of systems able to live, evolve and reproduce will likely revolutionize all of science and industry, possibly transforming them as much as the first computer revolution."

To recap, we are seeking an appropriate biology that might provide the basis for a 21st century answer to the question "Are robots alive?" In the nascent science "Living Technology" the author submit that we do have an appropriate biology. As Bedau and his colleagues explain:

"It [living technology] may be seen as an essentially new form of life, and as such, it expands the realm of biological life." [49] … "The key difference from mankind's earliest bio-inspired technological awakening is that we are now approaching a position to turn nonliving materials into technology that itself comes alive." [50]

And the products of this living technology will be living products in the same sense that the technology is living. They will be alive in the same sense as the technology is alive. And robots are amongst the products of that technology.

4.3. Future Outlook

Clearly living technology has immense potential for society. Bedau summarizes that potential thus:

"Nothing has really prepared us for the revolution that will take place when we make purely artificial technologies that embrace the core properties of life, allowing the technology itself to become robust, adaptive, self-repairing, self-optimizing, autonomous, intelligent, and evolvable. We believe that these developments right now are converging to create a qualitative jump, as purely artificial technology itself becomes alive.

We deem technology to be living if it is powerful and useful precisely because it has the core properties of living systems, including such properties as the ability to maintain and repair itself, to autonomously act in its own interests, to reproduce, and to evolve adaptively on its own. As our technologies increasingly embody such core properties of living systems, they will become increasingly powerful, natural, and sustainable." [50].

Right now, in 2014, living technology and its products are more in the future than in the present. But we are not talking about a future several decades away – the foundation stones of living technology are already with us. So even if you are not convinced that the robots of today are alive, you can be confident that the robots of tomorrow will be. That, from a futurist perspective, is a worst case answer. Those who decry the possibility of robots being alive will have to change their tune before very much longer, when the computational driving force of robots, their "brains", will undeniably be alive, and therefore the robots themselves will be alive.

Recent research continues to push the field forward, from self-replicating modular robotic swarms [52] to neuromorphic genetic circuits in living cells [53] and conceptual roadmaps for synthetic multicellularity [54].

4.4. Summary

In attempting to answer the question "Are robots alive?" we must first consider what we mean by "alive". The many different definitions of "alive" published by biologists, other scientists, philosophers, and experts from other disciplines, combine to create more confusion than assistance. There is no universally accepted definition of "alive" that helps us in our search for the answer.

With no suitable definition of "alive" to serve as our starting point, we come to consider various sets of criteria of life, criteria that have been promoted as enabling us to determine whether or not a particular entity is alive. Some of those sets of criteria are purely biological in nature, that is to say they are based on traditional biology. But there are so many such sets, and so many differences between sets, that they create much confusion to the extent that they should not be considered in any way definitive.

A more general set of criteria, a set that does not rely on any traditional biological considerations, is the Living Systems Theory propounded by James Grier Miller. He provided 20 criteria, to which he refers as "critical subsystems". His thesis was that in order to be considered a living system it would be necessary for a system to conform to his definitions of all 20 of the subsystems.

This paper claims that the brief comparisons the author make between Miller's 20 subsystems and various of the characteristics of robots, provide sufficient justification for us to consider robots to be living systems according to Miller's 1978 theory. So if we accept Miller's theory, and to the best of the author's knowledge it has never been refuted, we can safely answer our title question in the affirmative. Yes, robots are alive.

But today is not 1978 and a lot of scientific and technological advances have been made since then, advances which enable us to consider the question from quite different perspectives - 21st century perspectives. "Clearly the difference between what is life and what is not life has changed with the advance of science." [55] Languages too develop with time. One reason is that new words enter the vocabulary through common usage. Another reason is that existing words acquire new meanings, as noted by Geoff Simons in 1983 regarding the word "living":

"We are now beginning to re-examine many of the conventional adjectives traditionally applied to human beings—such words as conscious, intelligent, thinking, perceptive, free and aware. We are having to scrutinize such terms because we are seeing that increasingly they can denote characteristics of artificial systems… It is interesting to note how 'intelligent' has come to denote such machine systems. We may expect to see the word 'living' undergo a similar evolution." [22]

If we examine the current online version of the "English Dictionary", the world's most established and reliable authority on the meanings of English words, one of the current definition classes for "alive" is "of a thing", within which class we find [56]:

Of an abstract concept

Of a fire, flame or spark

Able to be used; in operation; working; fully functional

None of which are in any way biological terms. The author makes this point in order to demonstrate that today the word "alive" is not only appropriate when used in relation to living plants, living creatures or living organisms, but it can also be appropriate when it is employed in relation to things, to artifacts, to robots.

Advances in the arena of the various 21st century branches of biology – those that the author call the "new biologies"—have paved the way for the birth of the new science of Living Technology. Through the merger of these new biologies with advances in the fields of computation, artificial intelligence, robotics, and other disciplines, the coming years will see the creation of robots, driven by living technologies, that are very much alive in a 21st century sense, robots that astound us with their capabilities.

5. Conclusions

"Are robots alive?" In endeavouring to answer this question, the author has come to the following conclusions:

It depends on what we mean by "alive".

Biological definitions of "alive" are not adequate for answering the question.

Biological criteria for life are not adequate for answering the question.

If we accept Living Systems Theory, then the robots of today are living systems. They are alive [57]-[59].

Even if we do not accept Living Systems Theory, we can nevertheless justify describing robots as alive on the basis of one of the generally accepted meanings of the word, as given by the Oxford English Dictionary.

For any of you who are not willing to accept conclusions (4) and (5), may the author remind you of the words of Mark Bedau, one of the founders of the science of Living Technology:

"We are now approaching a position to turn nonliving materials into technology that itself comes alive." [50]

And if the technology comes alive, then so do artifacts that incorporate the technology, including robots.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

[1] Deplazes, A. and Huppenbauer, M. (2009) Synthetic Organisms and Living Machines. *Systems and Synthetic Biology*, **3**, 55-63. https://doi.org/10.1007/s11693-009-9029-4

[2] Michie, D., Spiegelhalter, D.J., Taylor, C.C. and Campbell, J. (Eds.) (1995) Machine Learning, Neural and Statistical Classification. Ellis Horwood.

[3] Cope, D. (2001) Virtual Music: Computer Synthesis of Musical Style. MIT Press.

[4] Cohen, H. (1995) The Further Exploits of AARON, Painter. *Stanford Humanities Review*, **4**, 141-160.

[5] Hsu, F.H. (2002) Behind Deep Blue: Building the Computer that Defeated the World Chess Champion. Princeton University Press.

[6] Weizenbaum, J. (1976) Computer Power and Human Reason. W.H. Freeman and Company.

[7] Kahn, P.H., Gary, H.E. and Shen, S. (2012) Children's Social Relationships with Current and Near-Future Robots. *Child Development Perspectives*, **7**, 32-37. https://doi.org/10.1111/cdep.12011

[8] Turkle, S. (2011) Alone Together. Basic Books.

[9] Lycan, W. (1987) Consciousness. MIT Press.

[10] Stone, C.D. (2010) Should Trees Have Standing? 3rd Edition, Oxford University Press.

[11] Lehman-Wilzig, S.N. (1981) Frankenstein Unbound: Towards a Legal Definition of Artificial Intelligence. *Futures*, **13**, 442-457. https://doi.org/10.1016/0016-3287(81)90100-2

[12] Freitas Jr., R.A. (1985) The Legal Rights of Robots. *Student Lawyer*, **13**, 54-56.

[13] Levinson, P. (1999) The Civil Rights of Robots. In: *Bestseller*: *Wired*, *Analog and Digital Writings*, Pulpless.Com Inc, 247-251.

[14] Levy, D. (2006) Robots Unlimited. AK Peters.

[15] Levy, D. (2007) Love and Sex with Robots: The Evolution of Human-Robot. Harper Perennial.

[16] Dibbell, J. (1996) The Race to Build Inteligent Machines. *Time Magazine*.

[17] Klir, G. and Valach, L. (1965) Cybernetic Modelling. Iliffe Books.

[18] De Ley, H. (1967) On Aristotle, «De Anima», A. 2, 404 B 27-405 A 13. *L'Antiquité Classique*, **36-2**, 605-612.

[19] Koshland, D.E. (2002) The Seven Pillars of Life. *Science*, **295**, 2215-2216. https://doi.org/10.1126/science.1068489

[20] Trifonov, E.N. (2011) Vocabulary of Definitions of Life Suggests a Definition. *Journal of Biomolecular Structure and Dynamics*, **29**, 259-266. https://doi.org/10.1080/073911011010524992

[21] Rostand, J. and Tétry, A. (1971) Larousse Science of Life. Hamlyn.

[22] Simons, G. (1983) Are Computers Alive? Harvester Press.

[23] Deplazes-Zemp, A. (2011) The Conception of Life in Synthetic Biology. *Science and Engineering Ethics*, **18**, 757-774. https://doi.org/10.1007/s11948-011-9269-z

[24] Wikipedia. Life. https://simple.wikipedia.org/wiki/Life#:~:text=They%20undergo%20metabolism%2C%20maintain%20homeostasis,can%20be%20found%20on%20Earth

[25] Wolfram, S. (2010) Computation and the Future of the Human Condition. Wolfram Media, Inc.

[26] Holland, O. (2003) Machine Consciousness. Imprint Academic.

[27] Foundation, S. (2001) Can a Machine Be Conscious? Laboratory Workshop, May 13th-16th 2001.

[28] Gallup, G.G. (1977) Self Recognition in Primates: A Comparative Approach to the Bidirectional Properties of Consciousness. *American Psychologist*, **32**, 329-338. https://doi.org/10.1037/0003-066x.32.5.329

[29] Levy, D. (2009) The Ethical Treatment of Artificially Conscious Robots. *International Journal of Social Robotics*, **1**, 209-216. https://doi.org/10.1007/s12369-009-0022-6

[30] Takeno, J., Inaba, K. and Suzuki, T. (2005) Experiments and Examination of Mirror Image Cognition Using a Small Robot. 2005 *International Symposium on Computa-*

tional Intelligence in Robotics and Automation, Espoo, 27-30 June 2005, 493-498. https://doi.org/10.1109/cira.2005.1554325

[31] Turing, A. (1950) Computing Machinery and Intelligence. *Mind*, **59**, 433-460.

[32] Miller, J. (1978) Living Systems. McGraw-Hill.

[33] Miller, J.L. and Miller, J.G. (1992) Greater than the Sum of Its Parts. *Behavioral Science*, **37**, 1-9. https://doi.org/10.1002/bs.3830370102

[34] von Neumann, J. (1966) Theory of Self-Reproducing Automata. University of Illinois Press.

[35] Butler, S. (1863) Among the Machines. Trübner & Co.

[36] Zykov, V., Mytilinaios, E., Adams, B. and Lipson, H. (2005) Self-Reproducing Machines. *Nature*, **435**, 163-164. https://doi.org/10.1038/435163a

[37] YouTube (2014) Video of Self-Reproducing Robot. https://www.youtube.com/watch?v=aBYtBXaxsOw

[38] Ieropoulos, I., Melhuish, C., Greenman, J. and Horsfield, I. (2005) EcoBot-II: An Artificial Agent with a Natural Metabolism. *International Journal of Advanced Robotic Systems*, **2**, 295-300. https://doi.org/10.5772/5777

[39] McCarthy, J. (2000) Free Will-Even for Robots. *Journal of Experimental & Theoretical Artificial Intelligence*, **12**, 341-352. https://doi.org/10.1080/09528130050111473

[40] Da Costa, F. (1978) A Programming Approach for Robots. *Interface Age*.

[41] Khalil, A.S. and Collins, J.J. (2010) Synthetic Biology: Applications Come of Age. *Nature Reviews Genetics*, **11**, 367-379. https://doi.org/10.1038/nrg2775

[42] Prescott, T.J., Lepora, N. and Vershure, P.F.M.J. (2014) A Future of Living Machines: International Trends and Prospects in Biomimetic and Biohybrid Systems. *SPIE Proceedings*, **2014**, Article 905502. https://doi.org/10.1117/12.2046305

[43] Mitchell, M. (2012) Biological Computation. *The Computer Journal*, **55**, 852-855. https://doi.org/10.1093/comjnl/bxs078

[44] Braich, R., Chelyapov, N., Johnson, C., Rothemund, P. and Adelman, L. (2002) Problem on a DNA Computer. *Science*, **296**, 1499-502.

[45] Evans-Pughe, C. (2011) Programming through Biology. *Engineering & Technology Magazine*.

[46] Bonnet, J., Yin, P., Ortiz, M.E., Subsoontorn, P. and Endy, D. (2013) Amplifying Genetic Logic Gates. *Science*, **340**, 599-603. https://doi.org/10.1126/science.1232758

[47] Lazcano, A. (2013) A Question without Answers? *Metascience*, **23**, 301-304. https://doi.org/10.1007/s11016-013-9861-8

[48] Bedau, M., Hansen, P.G., Parke, E. and Rasmussen, S. (2010) Living Technology, 5 Questions. Automatic Press.

[49] Bedau, M.A., McCaskill, J.S., Packard, N.H. and Rasmussen, S. (2010) Living Technology: Exploiting Life's Principles in Technology. *Artificial Life*, **16**, 89-97. https://doi.org/10.1162/artl.2009.16.1.16103

[50] Bedau, M.A. (2009) Living Technology Today and Tomorrow. *Technoetic Arts*, **7**, 199-206. https://doi.org/10.1386/tear.7.2.199/1

[51] European Commission, Information Society and Media (2011) Living Technology Artificial Systems Embodied Evolution. FET Consultation Workshop Report.

[52] Abdel-Rahman, A., Cameron, C., Jenett, B., Smith, M. and Gershenfeld, N. (2022) Self-Replicating Hierarchical Modular Robotic Swarms. *Communications Engineering*, **1**, Article 35. https://doi.org/10.1038/s44172-022-00034-3

[53] Rizik, L., Danial, L., Habib, M., Weiss, R. and Daniel, R. (2022) Synthetic Neuromorphic Computing in Living Cells. *Nature Communications*, **13**, Article 5602. https://doi.org/10.1038/s41467-022-33288-8

[54] Solé, R., Conde-Pueyo, N., Pla-Mauri, J., Garcia-Ojalvo, J., Montserrat, N. and Levin, M. (2024) Open Problems in Synthetic Multicellularity. *npj Systems Biology and Applications*, **10**, Article 151. https://doi.org/10.1038/s41540-024-00477-8

[55] Etxeberria, A. and Ruiz-Mirazo, K. (2009) The Challenging Biology of Transients. *EMBO Reports*, **10**, S33-S36. https://doi.org/10.1038/embor.2009.154

[56] Stevenson, A. (Ed.) (2010) Oxford Dictionary of English. Oxford University Press.

[57] Deacon, T.W. (2021) A Teleodynamic View of Life-Like Cognition in Synthetic Agents. *Artificial Life*, **27**, 1-22.

[58] Bich, L. and & Moreno, A. (2022) The Organizational Account of Minimal Cognition: A Formal Analysis. *Journal of Theoretical Biology*, **539**, Article 111072.

[59] Chella, A. and Manzotti, R. (2023) Machine Consciousness: A Survey. *Philosophical Transactions of the Royal Society A*, **381**, Article 20220056.