Chapter 3. What is life?

12 June 2005

One might think that the first task in discussing the potential for life to exist off of the Earth would be to make sure that we know what life is. After all, without a definition of life, how can we know what we're looking for or if we've found it? As we'll see, however, it's not so easy to come up with a definition that includes everything that we think of as being life, excludes everything that we think of as not being life, and also provides guidance as to what to look for on another planet. While it's tempting either to put off the discussion of the nature of life to the very end or to not discuss it at all, trying to address the question will allow us to put many of the issues that we saw in the previous chapter into proper context. It will also highlight the problem of determining which characteristics of life on Earth might be generic to life anywhere versus specific to terrestrial life.

Characteristics of life

We can begin a discussion of the definition of life by describing some of the characteristics that we usually think of as setting living things apart from non-living things. We often think of life as having order or structure that sets it apart from its surroundings, of being able to utilize energy, of taking in nutrients and giving off waste products, of being able to grow and develop, of carrying out known biochemical reactions, of reproducing, of responding to the environment in which it resides, and of being able to adapt to its surroundings by what we describe as Darwinian evolution. Let's discuss each of these briefly and see whether, together or separately, they can provide a useful definition of life.

Life has order or structure

Living organisms on Earth are not made up of a random assortment of the elements assembled together in a random fashion. Life is made up of certain key elements. Carbon, hydrogen, oxygen, and nitrogen are the most central, but another two dozen or so play important roles as well. These other elements include, for example, sulfur, phosphorous, calcium, iron, magnesium, and so on. But the list does not include, for example, uranium, beryllium, or titanium. Thus, at its most basic level, life has structure by having specific elemental components, and they are present in roughly consistent and uniform relative abundances.

At a somewhat larger scale, the atoms that are utilized are not arranged randomly, but are organized into specific molecules. Admittedly, there are a very large number of molecules that are important, but there is an even larger number that aren't present in terrestrial life. Some of the key molecules that are involved in life include ATP (adenosine triphosphate), nucleic acids, amino acids, and proteins (which are assembled from amino acids).

At an even larger scale, these molecules are themselves assembled into structures that are necessary components of terrestrial life. They include membranes that separate the inside of a cell from the outside, DNA and RNA molecules that contain genetic information and govern the reproduction within living things, and cellular structures such as chloroplasts (the photosynthetic organelle in plants) or mitochondria that can carry out specific chemical functions within a cell.

Having a well-defined order or structure alone cannot define life, however. Minerals and rocks have order and structure in their elemental composition, in the form of specific molecules that are made from these elements and in the details of macroscopic rocks. Similarly, large-scale geological features such as volcanoes, earthquake faults, and other geological constructs or edifices have order and structure, as well, and we recognize these entities as not being alive.

Taking in nutrients and giving off waste products

Terrestrial organisms take in nutrients and give off waste products. We're used to thinking about animals taking in oxygen, using it to in chemical reactions with ingested organic molecules (food), and giving off the carbon dioxide that forms, for example. Plants take in carbon dioxide as well as other elements and molecules from their surroundings (things that we feed our house plants as fertilizer, for instance), and give off oxygen. Certain kinds of bacteria take in iron atoms from the rock that surrounds them and give off a more-oxidized form of iron, and others do the same thing with manganese.

By itself, however, this characteristic wouldn't be a good discriminator of what is living, because there are lots of examples of non-living entities doing the same thing. Rocks can "take in nutrients" by absorbing oxygen from the atmosphere in chemical reactions that oxidize the minerals that make it up; these oxidized minerals can be "given off" as waste products in the form of mineral "weathering products". In fact, the net chemical reactions in this weathering are the same ones as can provide energy to some organisms, so even the specific chemical reactions cannot be used to discriminate the living from non-living. The Earth's atmosphere can exchange as well, taking in sunlight that powers photochemical reactions in the atmosphere, and giving off chemical byproducts that can precipitate into solid particles and be removed from the atmosphere by rain. Fire takes in nutrients (oxygen, fuel) and gives off waste products (soot, smoke, carbon monoxide, carbon dioxide, water).

Utilization of energy

The energy from chemical reactions between nutrients is utilized by living organisms to power other chemical reactions. The energy can be stored in the molecule ATP, which can be thought of as a storage battery for energy, and then released again when that energy is needed. The energy can be used to carry out specific functions such as to build molecules in growth of the organism (through chemical reactions that produce organic molecules as the products), or it can be utilized to carry out mechanical work (moving our muscles, for instance).

Utilization of energy occurs in living organisms but it also occurs in non-living things. Beaches utilize energy (pardon my anthropomorphizing an entity such as a beach by suggesting that they actually "do" something, but the point is the same) in the sense that wave energy can be used to move sand grains around to make larger structures (such as ripples on the sea floor) or to make additional grains of sand out of larger rocks and thereby create more beach.

In contrast, some things that we think of as being alive may not utilize energy at a given moment. Dormant spores or seeds don't take in nutrients or give off waste products, and they don't use energy for anything. Admittedly, they will take advantage of the right environmental conditions to do these things, but observing a spore or seed under the wrong conditions might lead one to conclude that they were inanimate and not alive.

Growth and development

When plants utilize energy, they create additional physical structure and get bigger. When animals utilize energy in their way, they also often get bigger. Microbes or individual cells within a multicellular organism can get bigger, but they also utilize energy to create the molecules that allow them to divide into two microbes or cells. Is this a defining characteristic of life?

Mountains also can grow, driven by the geological process of plate tectonics; rubble piles at the base of mountains grow as weathering breaks the mountains apart; sand dunes grow if there is a supply of sand, and the federal debt grows no matter what we do. Fire grows by taking in nutrients, with chemical reactions creating heat that causes expansion of the fire. As with many living things, growth of non-living things can continue as long as there is fuel to provide a source of energy.

Carrying out biochemical reactions

All living organisms on Earth carry out many of the same biochemical reactions. Every organism uses the same two dozen amino acids to construct proteins, uses the molecules ATP and ADP to store and release energy, and uses RNA and DNA molecules to store genetic information and to code for the construction of proteins. This list of reactions comes as close as may be possible to defining terrestrial life. If these molecules are present and their associated chemical reactions occur, then the entities in which they occur are alive; if they are not present, then the entity is not alive. However, it is not clear that these molecules are required by life in general. We don't know whether these specific biochemicals are the only possible solutions to the problem of storing energy and information. Would life elsewhere absolutely have to consist of the same molecules, or is it possible that other molecules, with either only a slightly different structure or possibly a completely different structure, could carry out these functions?

Molecules have recently been constructed in the laboratory that are theoretically able to carry out the same functions but that are made of different structures; they suggest that life elsewhere could have stumbled onto alternative molecules and that RNA, DNA, and ATP might not be universal indicators of life. In addition, we have hypothesized a form of life that doesn't use DNA and proteins. This life is based on only the RNA molecule, which could have carried out the DNA functions related to reproduction and the protein functions of catalyzing chemical reactions. This biological system is termed the "RNA World", and is thought to have predated the present "DNA World" here on Earth. Therefore, it seems likely that "life as we know it" does not have to be the only conceivable type of life and that other molecules could carry out the necessary functions.

Responding to its environment

Organisms respond to their environment. When we get cold, we put on a coat; if we don't have a coat, we shiver. Plants can turn to face in the direction of sunlight, so that they maximize their intake of energy for photosynthesis. Microbes can detect gradients in the composition of the medium in which they reside, and can move in the direction of either increased habitability (e.g., based on temperature, salinity, etc.) or greater supply of food.

But non-living things also respond to their environment in different ways. Most solids will expand in size or volume when heated or contract when cooled. In some cases, rocks can break ("reproduce") when heated rapidly. Rocky soils will overturn in response to repeated freeze-thaw cycles, "heave" rocks to the surface, and move them around to form polygonal patterns on the surface. Beaches or sand dunes will form different types of patterns on the surface in response to different intensities of waves or wind.

Reproduction

Living things reproduce. Some do so sexually, by combining the characteristics from two different organisms. Plants do it by producing multiple seeds that can each grow into a complete organism, or by allowing cuttings from an original plant to grow independently. Single-celled organisms do it asexually, typically by duplicating internal organelles, molecules, and structures and then splitting in half.

But some living things do not reproduce. Mules are a classic example; as the cross between a horse and a donkey, they are incapable of producing offspring. Or some individual organisms are incapable of reproducing due to genetic or physical defects. We don't imagine that these are non-living entities (imagine trying to tell your infertile cousin that they are therefore not alive), so we create a way around this problem — we suggest that the individual cells comprising a mule or an infertile being do reproduce and are alive, so therefore the entire organism must be alive.

And some non-living things reproduce. Sand ripples reproduce as a result of the physics of sand movement by the wind, for example, with one ripple spawning additional ripples downwind.

Many of the counter-examples that we've been discussing come from the world of geology, as that is the major category that's left on the Earth in the absence of biology. However, one can imagine additional problem entities. Could one construct a robot that contained the instructions to build an identical robot? Would this ability to reproduce make them alive?

Adaptation to its environment

The single most-significant unifying concept in biology is the idea of evolution by modification and natural selection, what we now call Darwinian evolution, and it is central to life on Earth. The concept of Darwinian evolution stems from a couple of simple concepts. First, there is always natural variation of characteristics within a population. These are driven in part by the occurrence of mutations due to errors in copying DNA and RNA molecules or to changes induced by cosmic-ray collisions that induce changes in atomic bonds within a molecule. These changes will affect the functioning of the molecule that was changed. This change in function will leave some organisms better able to survive within their environment and other organisms (most, actually) less able to survive. As these changes are incorporated into the genetic material of the organism, they will be passed on to their offspring.

Second, more offspring are produced than typically can be supported by the environment, so that many offspring do not survive to adulthood. Thus, there will be competition amongst the organisms in a given population. Those that are better able to survive will produce more offspring, and by virtue of retaining the genetic information from their parents they also will be better able to survive. Those that are less able to survive will not be able to compete in their environment as well and will give rise to fewer offspring of their own; their own functionality may be impaired so much by the mutation that they die without having been able to reproduce at all.

Over time, the positive changes will spread throughout the entire population. Changes that allow better competition can accumulate over time, giving rise to organisms that are increasingly better adapted to their environment. The accumulation of a sufficient number of changes eventually produces an organism that is so different that we consider it a new species. Over long periods of time, with multiple events that can split off groups of organisms from each other, this process created the multitude of different species that populate the Earth today.

This concept of change and natural selection was first described by Charles Darwin in his 1859 book *On the Origin of Species*, and we refer to it as Darwinian evolution. It has been such a unifying theme in biology that, as one biologist put it, nothing in biology makes any sense except in the context of Darwinian evolution.

It is so fundamental that one possible definition of life centers on it. Gerald Joyce, a molecular biologist working on biochemical reactions that can take place in the laboratory, put forward this definition in the introduction to a NASA report on life: Life is a self-contained chemical system capable of undergoing Darwinian evolution.

The guts of this characteristic are simple—Darwinian evolution is so fundamental to our understanding of life on Earth that it is impossible to imagine life without it. The other points are secondary to this basic point. The idea of "self-contained" precludes, for example, a laboratory environment in which a key step is supplied by outside agents (e.g., humans). And a "chemical system" precludes life based on a computer program that resides within a computer or our earlier hypothetical example of robots building robots.

The biggest problem with this definition is how to determine whether something is capable of Darwinian evolution. This definition applies to a population of organisms operating over long periods of time; it cannot apply to a single organism, nor can it apply to a population observed at a single instant.

Our example of a mule rears its ugly head again here, so to speak. It satisfies our intuitive definition of being alive, but because it cannot reproduce it is not capable even in theory of Darwinian evolution. Or is the problem moot because a single infertile individual is still a part of the population, and it is the population that is capable of Darwinian evolution?

In addition, there is the problem of "non-Darwinian evolution". Organisms are capable of exchanging DNA with other organisms. This idea of "lateral gene transfer" allows for sudden jumps in the genetic information of organisms and in their functional capability. Does this type of change lie outside the idea of Darwinian evolution, and does it change this definition of life?

Early in the history of life, it is possible that organisms were not independent entities each encased in their own membranes. The organic soup of biomolecules in liquid water, each reacting both with other molecules and with their surroundings, might have been the first life. In such a case, evolution might have operated at the molecular level rather than at the organism level. Does this process lie outside of this definition of life, even if we would recognize the entities as being alive?

Problems with definitions

Unfortunately, every characteristic that we have raised for possible inclusion in a definition of life has been found wanting. Each one has counterexamples that, while admittedly somewhat contrived, show the limitations of the definition. If we use these characteristics to construct a definition anyway, then we have tuned the resulting definition solely to be consistent with our preexisting biases toward what constitutes life on Earth. We've carefully defined life so that it includes all of the things here on Earth that we think should be included and excludes all of the things on Earth that we think should be excluded. In this sense, our definition is fundamentally no different from a long list that includes on it every terrestrial organism that we think is alive.

Additionally, our goal in defining life is not to be able to categorize living things here on Earth. We don't need a definition to allow us to do so. Rather, we wish to examine samples from another planet and to determine whether they contain evidence for life within them. By having tuned our definition to terrestrial life, we are *a priori* limiting ourselves to identifying or finding life that is very similar to terrestrial life. Life that might have a different biochemistry, multiply or reproduce on very different time scales, or in a snapshot analysis be indistinguishable from natural, non-living entities, could not be identified.

Another way of saying this was described recently by Carol Cleland, a philosopher of science at the University of Colorado with an interest in the attempts to define life. She pointed out that, with a single example of life — that on Earth — we have the same problem as scientists did in trying to define water prior to the existence of molecular theory. We know today that water is uniquely and absolutely defined as the molecule described by the formula H_2O. Prior to understanding the nature of molecules, water was defined based on what it could or could not do — how it would interact with other chemicals, what it looked like or tasted like, and so on. Trying to use these characteristics as a way of defining life clearly is inadequate, and this approach is remarkably reminiscent of trying to use the characteristics of life described above to arrive at a unique definition of life today.

Without having a sample of life that has an independent origin from that on Earth, and that might show a wider range of characteristics, molecular structures, or biochemical compositions and reactions, it is not possible to determine what characteristics of terrestrial life are specific to Earth life and which are general to all life. Thus, we are absolutely precluded from coming up with a singular and unique definition of life.

Origin of life — the same problem revisited

We have the same problem in identifying a unique boundary between living and non-living entities when we consider the origin of life on Earth some four billion years ago. The rock record is too sparse to allow us to determine the processes that were associated with the origin of life. However, let us imagine that we could have watched the origin taking place over many hundreds of millions of years. Early on, we would have seen a system that included the surface of the Earth, the oceans, and the atmosphere. Undoubtedly, there would have been significant chemical reactions taking place, given the wide variety of molecules that would have been supplied from space during the Earth's formation, from the interior of the Earth as a result of the differentiation and outgassing that took place during its earliest history, and from chemical reactions that were taking place in diverse chemical environments. Places of interest might have been the oceans themselves, the "warm little pond" that Darwin referred to where chemicals eventually might have come together to create life, or hydrothermal systems where heat from volcanoes or from large asteroid impacts would have driven circulation of water and created environments where disequilibrium chemical processes might have been important. This environment clearly would have been dominated by chemistry and geochemistry, with a complete absence of life. [If my description of the ongoing chemical reactions sounded too biological and you have doubt that this time was one at which there was no life, simply back up another 100 million years; continue to back up until you're convinced that the system was non-biological.]

If we jump ahead a half-billion years or so, to perhaps 3.5 billion years ago, we find convincing evidence in the rock record that life existed at that time. We see structures in certain kinds of rocks that are interpreted as fossilized cells, and we see larger-scale structures that appear to have been made in environments that consisted of colonies of fossilized cells. In addition, there is isotopic evidence suggesting that life was present at this time. Questions have been raised about each of these lines of evidence, though. [If you have doubts as to whether life existed at this time, simply jump forward in increments of a half-billion years or so until your doubts disappear; doing so will not affect our argument.] At this time, then, there is evidence that everybody would accept for biological activity (even if we have to move far ahead in time to, say, only 2 billion years ago).

The environment on Earth clearly underwent a transition from a time at which only geochemical processes were operating to one in which biochemical processes were operating as well. Was there a distinct moment when we went from having no life to having life, as if changed by a switch?

The answer is, probably not. We think that the origin of life involved the creation of an "RNA World". In such a world, molecules composed of a string of nucleic acids, much like a modern RNA molecule, would have been able to carry out the functions both of carrying genetic information and of catalyzing chemical reactions. That is, one molecule, similar to RNA, could have carried out the functions that today are performed by both DNA and proteins. It's possible to imagine how we could have gotten from an RNA world to the present DNA world. And it's also possible to imagine how we could have gone from a "geochemical world" to an RNA world.

The formation of RNA-like molecules has been demonstrated in the laboratory using minerals as a catalytic surface on which chemical reactions could take place, and it has been hypothesized as possibly taking place on a surface composed of layers of organic molecules. Longer and longer chains could have been built up, with each having some catalytic capability, until one became long enough and complex enough that it was capable of catalyzing its own reproduction independent of the mineral surface catalyst. While this change may sound like a switch-on of life, crossing over a boundary, this is not necessarily so. For example, one could imagine that the fidelity of copying could have gotten better as longer chains were formed. Or that a chain could only reproduce for a certain number of generations before using up the chemical ability to induce additions of new bases at one end of a copied chain, and that a longer chain could have gone on for many generations before petering out; it might be only a matter of semantics as to how many generations of reproduction need to occur before one thinks of it as living.

In this case, there is not necessarily a sharp boundary between living and non-living at the time of the origin of life. The extremes are easy to identify and characterize — early on, with no life, and later on, with life — but the intermediate stages cannot be uniquely or readily categorized. In essence, there could have been a smooth, gradational boundary in the creation of life, with no place along the transition being a clearly definable moment at which life first existed.

Gradational boundary for life—definition by consensus

In the same sense, the present-day boundary between living and non-living may be inherently impossible to define. There may easily be intermediate forms, such as viruses, that just aren't amenable to a unique categorization. We can then arrive at a gradational definition for life. Think of the half-dozen characteristics with which we began this chapter. If an entity satisfied few or none of these criteria, everybody would agree that it was non-living. Things that might fall within this category might include the wind, the atmosphere, rocks, mountains, sand dunes, or fire.

If an entity satisfied most or all of these criteria, then most everybody would agree that it was alive. Certainly, any example that we would universally agree to be living appears to satisfy most of these characteristics (perhaps failing to satisfy no more than perhaps one of them). This conclusion applies to single-celled organisms such as an amoeba or the *e. coli* bacterium, to macroscopic multicellular organisms such as flies, rabbits, or elephants, and to humans.

In the middle, it's a gray area. The more characteristics that are satisfied, the more likely it is that we would agree that it is living, or the greater the fraction of knowledgeable people would accept it as living.

We can see that this concept gives us the possibility of describing how we might identify a newly discovered entity and determine if it alive. By examining the available criteria, is there a consensus as to whether it meets most or all of the characteristics of life? For those in which it doesn't, do we think that the failure is "fatal" to the idea of it being alive, or is it only incidental?

As we will see later, much of science is done by consensus anyway, so that a consensus means to determine whether something is alive might actually make sense. And, this approach provides some pragmatic advice on how to interpret observations of potentially living organisms found on other planets.

The proof is in the pudding: Real-world applications

This approach to determining whether something is alive makes sense only if it can provide some guidance in real-world situations. We'll look at three examples — determining whether viruses are alive or not, determining whether ancient terrestrial fossils were alive, and determining whether features identified in rocks that came from Mars were alive. We'll also discuss the (at present) hypothetical case of looking for life on Mars either *in situ* or by examining samples that we collect there and return to the Earth.

Modern life—what about viruses?

Viruses are, in essence, extracellular nucleic acids that are encased in a protein coating. The nucleic acids can be either DNA or RNA. By themselves, they do not have the machinery to replicate, and so cannot reproduce independently; instead, they infect biological hosts and utilize their biochemical machinery. The host organism provides the biochemical machinery to make copies of the DNA or RNA and thereby make multiple copies of the original virus. The host cell bursts and releases these multiple new copies into the environment, where they are able to infect additional cells, create new copies of itself again, and so on.

On the one hand, viruses can be considered as being alive: Where initially there was a single virus, multiple copies are created. They are capable of competing via Darwinian evolution, with mutations creating changes in the genome of the virus and with competition allowing those that are better able to reproduce to out-compete the others. And they contain RNA and/or DNA and thereby satisfy the closest thing we have to a litmus test for terrestrial life. On the other hand, they are not able to reproduce without borrowing the biochemical machinery of other organisms, and they cannot carry out their complete life cycle alone. In this sense, they fail Gerald Joyce's attempt to define life by not being a "self-contained" chemical system.

Of course, this aspect has its own problems — no organism can be entirely self contained. In the simplest sense, all organisms interact with their environment and are not self-contained. Some organisms, however, are symbiotic with other organisms, and rely on them for very specific functions. Some bacteria take in chemicals that are produced as waste products by other organisms, for example with the bacteria that live within tube worms at mid-ocean spreading centers. Other bacteria live within the gut of animals, utilizing the animal's partially digested food as its own and simultaneously helping the animal to digest its food; they are present within both humans and cows, for example. Other bacteria have become so intimately integrated with their symbiont that they no longer can live alone; chloroplasts and mitochondria, for example, are thought to have once been independent organisms and are now integral parts of other cells.

Viruses occurring within each of the three major branches of life (archaea, bacteria, eukarya) share common characteristics, suggesting that they may predate the ancient split into these branches. Other characteristics have been interpreted as indicating an existence prior to the origin of life. It remains to be seen whether viruses played an important role in the origin of life, originated as miscopied pieces of the DNA of ancient organisms, or represent an independent origin of life here on Earth.

Viruses are often singled out as being non-living, but it isn't obvious from this discussion that this conclusion is correct. They meet some of the requirements of our definition, yet do not meet others. There is no consensus as to whether they are living, and they appear to fall in the middle of our sliding scale. Interestingly, there are bacteria that are a little bit more complex than viruses and require some of the biochemical machinery of host organisms in order to function. *Rickettsias* and *chlamydias* fall within this category. And, the smallest bacterial genomes are only a little bit larger than the largest viral genomes. There also are entities known as viroids, that consist essentially of naked RNA molecules. Viroids are simpler than viruses, lacking the protein coat and some of the enzymes that often accompany the DNA and RNA in the viral interior. They also are capable of infecting organisms, with the biochemical machinery of the host organism carrying out all aspects of their reproduction.

It seems plausible that, from viroid to virus to the smallest bacteria to other bacteria, we are seeing a smooth transition in capability and functionality in addition to a transition from non-living to living.

Ancient terrestrial fossils?

There's been controversy recently over the oldest fossil evidence for life here on Earth. Work done during the past several decades, led in large part by Bill Schopf, a paleontologist from UCLA, has pushed back the oldest fossil evidence for life to about 3.5 billion years ago. The evidence took the form of fossil entities found in rocks that looked like fossilized single-cell organisms. The appearance was not overwhelmingly convincing by itself, however, because the fossils had been severely degraded due to the heat, pressure, and other geological processes that had taken place during the billions of years since they formed. These entities formed a continuous "series" in terms of their physical characteristics that seemed to fall in place with entities in younger rocks that looked much more like fossils. That is, the younger ones were convincing, and the older ones looked enough like the younger ones given their substantial degradation that it was straightforward to conclude that they were biological.

However, a recent reanalysis of these same fossils has called their biological nature into question. It turns out that the fossils are found not in the rocks that are dated at 3.5 billion years, but in a cross-cutting rock that is, of necessity, younger by some unknown amount. In addition, some of the characteristics of the putative fossils do not match up cleanly with those that would be expected of living things. Some have branching extensions, for example, that are not found in the biological world, or are at one end of a continuum of characteristics (size, shape) for which the other end is not thought to be biological. In addition, laboratory experiments have demonstrated that similar-looking features can be formed by non-biological processes.

Are these features biological or not? Right now, there is no consensus in the scientific community. Rather than jumping to a rash conclusion, the community has responded by trying to get more data that might shed light on the issue. Without a consensus, and without more data that would lead to a consensus, it's hard to believe either extreme camp—with one group saying that these features are clearly biological and the other that they are non-biological. In this instance, the entities in question either were or were not alive; the only issue is our ability to decipher the evidence in a unique, convincing, or compelling way. And, the scientific process is following along our view of a "consensus" definition of life — as more data is obtained, a consensus ultimately will develop as to what the best interpretation of these features is.

Fossil life in a meteorite from Mars?

Meteorites have been found on the Earth that are convincingly thought to have come from Mars based on their ages, oxygen isotopic composition, and the similarity of trapped gases to the composition of the martian atmosphere. To date, more than thirty such rocks have been found. One of them, known as ALH84001 and collected in Antarctica in 1984, has been controversial in the discussion of possible martian life. This one rock is 4.5 billion years old, which means that it was present on Mars from about the time that the planet formed, and that it was present when the surface environment might have been more conducive to life. In addition, it contains deposits of carbonate minerals filling voids within the rock, which requires that liquid water flowed through the rock and that dissolved minerals precipitated within the rock; these processes occur very commonly here on Earth. The evidence for the prior occurrence of liquid water made it an interesting rock in which to look for evidence of life.

A group led by David McKay and Everett Gibson at NASA's Johnson Space Center examined this rock carefully to see if there was evidence for fossil life within it, and published their results in 1996. They identified five characteristics that they thought were consistent with the presence of life. These included the presence of a specific type of organic molecules that could be either decay products or precursors of life, the presence of minerals within the voids that were adjacent to each other yet out of chemical equilibrium with each other, the presence of magnetite grains that have sizes and shapes very similar to those thought to be produced only by bacteria here on Earth, and the presence of morphological shapes reminiscent in appearance to terrestrial microbes. While none of them could be attributed uniquely to biological activity, they argued that, taken together, they made a strong case.

The announcement of possible fossil martian life, in addition to creating a major public stir, triggered intense scientific analysis of these samples. Over a period of a half-dozen years, a couple of dozen scientific teams examined the rock in more detail, and were able to provide a much clearer view of its history and characteristics. Today, each of the original characteristics has been called into question, either as an artifact of the handling, processing, and analysis or as something that could be produced by non-biological processes. The discovery of modern terrestrial microbes living as a contaminant within the rock, for example, has cast doubt on any possible conclusions about extraterrestrial origin of biological features.

In this case, the consensus that has emerged is that none of the features identified in the rock requires the existence of martian biota, and that all of them are likely explainable by non-biological processes. This consensus does not mean that this perspective is unanimous, in that many respectable scientists have come down on the "biology" side of the features. However, it does mean that, at best, a conclusion that this rock contains evidence for martian biota is premature. We need more evidence in order to reach a stronger conclusion about whether these features were formed from biological entities. (Keep in mind that this consensus is based on available evidence, that new evidence continually comes into play, and that the nature of consensus can change over time.)

Searching for life elsewhere

So if we have trouble identifying evidence of living organisms in terrestrial rocks, even after more than thirty years of analysis, and if we can be so misled by analysis of extraterrestrial rocks without the tremendous effort that was subsequently put into analyzing a single rock, how are we ever going to identify life elsewhere with any certainty? What would it take to be sufficient or convincing evidence?

One view is that a search for life on Mars, for instance, has two possible outcomes. One is that we might identify features that are immediately accepted by knowledgeable scientists as convincing evidence for martian life. This evidence could take the form of obvious cells, with easily identifiable well membranes, interior structures, and so on, or of cells caught in the act of multiplying, etc. Such obvious cells could be either living (extant life) or fossil (past life). A second possibility is that the evidence might be less compelling or ambiguous, and that it would engender years of debate and uncertainty. This case would be analogous to the debate over fossil life in ALH84001. Over time, a consensus might or might not develop but, as with ALH84001, even a consensus probably would not be unanimous and would be accompanied by never-ending debate over whether life was present. (A third outcome is possible, of course, in which little or no evidence indicative of life would be found. In this case, an immediate consensus might even form that there was no life in the sample, but this result is different from what would be required to reach a conclusion that there was no life on Mars.)

A second view is that it is not possible to predict what we will find or how we will react to specific discoveries. Finding entities within a martian sample that meet some of the criteria for life would not necessarily provide convincing evidence for the presence of life. Finding chemical or mineralogical structures and order would not necessarily be proof of life, but it would point to places that are worth examining in more detail. Further examination might identify features that could point in a convincing way to either a biological or a non-biological origin for them. This approach, championed by Ken Nealson of USC among others, suggests that we can't actually know in advance what it might take to convince us that we have found life, but that we may know it when we see it; at the least, we can hope for seeing something that would require additional investigation.

Where does that leave us?

Do we need a definition of life? Does it help us in understanding the origin and evolution of life on Earth or the environmental conditions required to support life? Does it help us in determining whether there is life elsewhere?

It seems clear that we do not have enough information to arrive at a unique definition of life. At the same time, attempting to define life allows us to focus attention on some of the issues in understanding life's characteristics. And it will be of help as we explore the ancient fossil record on Earth or the fossil or possibly biological record on other planets. But, in the end, without a unique definition, it will be by consensus within the knowledgeable scientific community that we are able to reach any conclusions.

Or, we can always fall back on Mark Twain's attempt to define life: "Life is just one damn thing after another".