Chapter 24

What is Life?

MARK A. BEDAU

1. The Fascination of Life

The surface of the Earth is teaming with life, and it is usually easy to recognize. A cat, a carrot, a germ are alive; a bridge, a soap bubble, a grain of sand are not. But it is notorious that biologists have no precise definition of what life is. Since biology is the science of life, one might expect a discussion of the nature of life to figure prominently in contemporary biology and philosophy of biology. In fact, though, few biologists or philosophers discuss the nature of life today. Many think that the definition of life has no direct bearing on current biological research (Sober, 1992; Taylor, 1992). When biologists do say something about life in general, they usually marginalize their discussions and produce something more thought provoking than conclusive. But this is all changing now.

Today the nature of life has become a hot topic. The economic stakes for manipulating life are rising quickly. Biotechnologies like genetic engineering, cloning, and high-throughput DNA sequencing have given us new and unprecedented powers to reconstruct and reshape life. A recent development is our ability to reengineer life to our specifications using synthetic genomics (Gibbs, 2004; Brent, 2004). In this domain attention has fallen on Craig Venter's well-publicized effort to commercialize artificial cells that clean the environment or produce alternative fuels (Zimmer, 2003). The current "wet" artificial life race to synthesize a minimal artificial cell or protocell from scratch in a test tube (Szostak, Bartel, & Luisi, 2001; Rasmussen et al., 2004; Luisi, 2006; Rasmussen et al., 2007) also spotlights life, for the race requires an agreed-upon definition of life, and it must be one that reaches well beyond life's familiar forms. The social and ethical implications of creating protocells will also increase the need for understanding what life is. Current controversies over the origin of life (Oparin, 1964; Crick, 1981; Shapiro, 1986; Eigen, 1992; Morowitz, 1992; Dyson, 1999; Luisi, 1998) and over intelligent design (Pennock, 2001) add more fuel to the fire.

Another recent development that highlights the nature of life is "soft" artificial life attempts to synthesize software systems with life's essential properties (Bedau, 2003a). Soft artificial life has created remarkably life-like software systems, and they seem genuinely alive to some (Langton, 1989a; Ray, 1992), but others ridicule the whole idea of a computer simulation being literally alive (Pattee, 1989).

Further still, recent "hard" artificial life achievements include the first widely available commercial robotic domestic vacuums, Roomba (Brooks, 2002), and the walking robots designed by evolution and fabricated by automated rapid prototyping (Lipson & Pollack, 2000). These robots inevitably raise the question whether a device made only of plastic, silicon, and steel could ever literally be alive. Such scientific developments increase uncertainty about how exactly to demarcate living things.

Biology makes generalizations about the forms life can take, but such generalizations rest on the forms of life that actually exist. Biologists study a number of different model organisms, like *Escherichia coli* (a common bacterium), *Caenohabditis elegans* (a nematode), and *Drosophila melanogaster* (a fruit fly). Picking model organisms that are as different as possible best illustrates the possible forms that life can take, and thus enables the widest generalizations about terrestial life. But all the life on Earth is terrestrial. Thus, these generalizations about life currently hinge on a sample size of one. Maynard Smith (1998) pointed out that artificial life helps mitigate this problem. Natural life comes in an amazing diversity of forms. But they are just a tiny fraction of all possible forms of life. Anytime we can synthesize a system in software, hardware, or wetware that exhibits life's core properties, we have a great opportunity to expand our empirical understanding of what life is.

There are three giants in the history of philosophy who advanced views about life, and their views still echo in contemporary discussion. In the *De Anima* Aristotle expressed the view that life is a nested hierarchy of capacities, such as metabolism, sensation, and motion. This nested hierarchy of capacities corresponds to Aristotle's notion of "soul" or mental capacities, so Aristotle essentially linked life and mind. As part of his wholesale replacement of Aristotelian philosophy and science, Descartes supplanted Aristotle's position with the idea that life is just the operation of a complex but purely materialistic machine. Descartes thought that life fundamentally differed from mind, which he thought was a mode of consciousness. Descartes sketched the details of his mechanistic hypothesis about life in his *Treatise on Man*. Some generations later, Kant's *Critique of Judgement* struggled to square Descartes's materialistic perspective with life's distinctive autonomy and purpose.

Understanding the nature of life is no mere armchair exercise. It involves investigating something real and extremely complex, and with huge potential creativity and power to change the face of the Earth (Margulis & Sagan, 1995). This investigation will by necessity be interdisciplinary, and it will survey an almost astonishing variety of perspectives on life. Interesting and subtle hallmarks like holism, homeostasis, teleology, and evolvability are thought to characterize life. But a precise definition of life remains elusive, partly because of borderline cases such as viruses and spores, and more recently artificial life creations. To add more complication, life figures centrally in a range of philosophical puzzles involving important philosophical issues such as emergence, computation, and mind. So, a diversity of views about life can be expected. Some employ familiar philosophical theories like functionalism. Others use biochemical or genetic explanations and mechanisms. Still others emphasize processes like metabolism and evolvability. The sheer diversity of views about life is itself interesting and deserves an explanation.

2. The Phenomena of Life

Life has various hallmarks and borderline cases, and it presents a variety of puzzles. The rest of this chapter is mainly devoted to explaining these phenomena.

A striking fact of life is the characteristic and distinctive *hallmarks* that it exhibits. These hallmarks are usually viewed as neither necessary nor sufficient conditions for life; they are nonetheless typical of life. Different people provide somewhat different lists of these hallmarks; see, e.g., Maynard Smith, 1986; Farmer & Belin, 1992; Mayr, 1997; Gánti, 2000. But most lists of hallmarks substantially overlap. Another notable point is that the hallmarks itemized on the lists are strikingly heterogeneous. A good illustration is Gánti's hallmarks (or "criteria," as he calls them).

Gánti's hallmarks fall into two categories: real (or absolute) and potential. Real life criteria specify the necessary and sufficient conditions for life in an individual living organism. Gánti's (2003) proposed real life criteria are these:

(1) *Holism.* An organism is an individual entity that cannot be subdivided without losing its essential properties. An organism cannot remain alive if its parts are separated and no longer interact.
(2) *Metabolism.* An individual organism takes in material and energy from its local environment, and chemically transforms them. Seeds are dormant and so lack an active metabolism, but they can become alive if conditions reactivate their metabolism. For this reason, Gánti makes a four-part distinction between things that are alive, dormant, dead, or not the kind of thing that could ever be alive.
(3) *Inherent stability.* An organism maintains homeostatic internal processes while living in a changing environment. By changing and adapting to a dynamic external environment, an organism preserves its overall structure and organization. This involves detecting changes in the environment and making compensating internal changes, with the effect of preserving overall internal organization.
(4) *Active information-carrying systems.* A living system must store information that is used in its development and functioning. Children inherit this information through reproduction, because the information can be copied. Mistakes in information transfer can "mutate" this information, and natural selection can sift through the resulting genetic variance.
(5) *Flexible control.* Processes in an organism are regulated and controlled so as to promote the organism's continued existence and flourishing. This control involves an adaptive flexibility, and can often improve with experience.

In contrast to these "real" criteria, Gánti also proposed "potential" life criteria. An individual living organism can fail to possess life's potential criteria. The defining feature of potential life criteria is that, if enough organisms exhibit them, then life can populate a planet and sustain itself. Gánti proposed three:

(1) *Growth and reproduction.* Old animals and sterile animals and plants are all living, but none can reproduce. So, the capacity to reproduce is neither necessary nor sufficient for being a living organism. But due to the mortality of individual

organisms, a population can survive and flourish only if some organisms in the population reproduce. In this sense, growth and reproduction are what Gánti calls a "potential" rather than "real" life criterion.

(2) *Evolvability.* "A living system must have the capacity for hereditary change and, furthermore, for evolution, i.e. the property of producing increasingly complex and differentiated forms over a very long series of successive generations" (Gánti, 2003, p.79). Since what evolves over time are not individual organisms but populations of them, we should rather say that living systems can be members of a population with the capacity to evolve. It is an open question today exactly which kinds of biological populations have the capacity to produce increasing complexity and differentiation.

(3) *Mortality.* Living systems are mortal. This is true even of clonal asexual organisms, because death can afflict both individual organisms as well as the whole clone. Systems that could never live cannot die, so death is property of things that were alive.

Gánti's life criteria and other lists of life's hallmarks always reflect and express some preconceptions about life. This might seem to beg the question of what life is. Any nonarbitrary list of life's hallmarks was presumably constructed by someone using some criterion to rule examples in or out. But where did this criterion come from, and what assures us it is correct? Why should we be confident that any hallmarks that fit it reveal the true nature of life? Thus, it seems lists of life's hallmarks are not the final word on what life is. As we learn more about life, our preconceptions change, evolve, and mature. So we should expect the same of our lists of life's hallmarks.

Another interesting feature of life is the existence of *borderline cases* that fall between the categories of the living and the nonliving. Familiar examples are viruses and prions, which self-replicate and spread even though they have no independent metabolism. Dormant seeds or spores are another kind of borderline case, the most extreme version of which might be bacteria or insects that are frozen. There are also cases that seem clearly not to be alive but yet possess the characteristic properties of living systems. Hardly anyone considers a candle flame to be alive, but by preserving its form while its constituent molecules are constantly changing, it has something like a metabolism (Maynard Smith, 1986). Populations of microscopic clay crystalites growing and proliferating are another kind of borderline example, especially because they can in appropriate circumstances undergo natural selection (Bedau, 1991). So is a forest fire that is spreading ("reproducing"?) from tree to tree at its edge, somewhat like the edge of a growing population of bacteria. A further kind of borderline case consists of superorganisms, which are groups of organisms, such as eusocial insect colonies, that function like a single organism. Although this is controversial, some biologists think that superorganisms should themselves be thought of as living organisms. Another kind of borderline case consists of soft artificial life creations like Tierra. Tierra is software that creates a spontaneously evolving population of computer programs that reproduce, mutate, and evolve in computer memory. Tierra's inventor thinks that Tierra is literally alive (Ray, 1992). This would radically violate the ordinary concept of life that most of us have. One final category of borderline cases consists of complex adaptive systems found in nature, such as financial markets or the World Wide Web. These exhibit many of the hallmarks of life, and some think that the simplest and most unified explanation of the entire range of phenomena of life is to consider these natural complex adaptive systems to be literally alive (Bedau, 1996, 1998).

3. Puzzles about Life

A third characteristic of life is that it generates a number of puzzles. Seven puzzles are briefly reviewed below. Any account of life should explain the origin of these puzzles; more important, it should resolve the puzzles. Some puzzles might result simply from confusion, but others are open questions about a fundamental and fascinating aspect of the natural world.

Origins. How does life or biology arise from non-life or pure chemistry? What is the difference between a system that is undergoing merely chemical evolution, in which chemical reactions are continually changing the concentrations of chemical species, and a system that contains life? Where is the boundary between living and merely physico-chemical phenomena? How could a naturalistic process bridge the boundary, in principle or in practice? Dennett argues that Darwin's scheme of explanation solves this problem by appealing to "a finite regress, in which the sought-for marvelous property (life, in this case) was acquired by slight, perhaps even imperceptible, amendments or increments" (1995, p.200).

Emergence. How does life involve emergence? B properties are said to *emerge from* A properties when the B properties both depend on, and are autonomous from, the A properties. Different kinds of dependence and autonomy generate different grades of emergence (Bedau, 2003b). One is the "strong" emergence involving in principle irreducible top-down causal powers. An example might be consciousness or qualia in the philosophy of mind (Kim, 1999). If the A and B properties are simultaneous, the emergence of B from A is *synchronic*. It concerns what properties exist at a moment. Those properties might be changing, but the relationship between the A and B properties at an instant are a static snapshot of that dynamic process. By contrast, if the A properties precede the B properties, and the B properties arise over time from the A properties, then the emergence of B from A is *dynamic*. Life is the paradigm case of a dynamic form of "weak" emergence, one that concerns macro properties that are unpredictable or underivable except by observing the process by which they are generated, or by observing a simulation of it (Bedau, 1997, 2003b).

Hierarchy. Various kinds of structural hierarchies characterize life. Each organism has a hierarchical internal organization, and the relative complexity of organizations of different kinds of organisms form another hierarchy. The simplest organisms are prokaryotic cells, which have relatively simple components. More complicated are eukaryotic cells containing complex organelles and a nucleus. Multicellular organisms are even more complicated; they have constituents (individual cells) that also are individual living entities (e.g., they can be kept alive by themselves). In addition, mammals have complex internal organs (such as the heart) that can be harvested and kept alive when an organism dies, and then surgically implanted into another living organism. Two questions arise here. First, why does life tend to generate and encompass such hierarchies? This question applies both to the hierarchy in complexity that spans all organisms together, and also to the organizational hierarchy found within each individual living organism. With regard to the latter, a second question arises. Organisms are our paradigm case of something that is alive, but we also refer to organs and individual cells as alive. For example, apoptosis is an important process by which living cells in an organism undergo programmed death, and hospitals strive to keep certain organs alive after someone dies, so that they are available to be transplanted into someone else. This raises the question whether a mammal, its heart, and the cells therein are each alive in the same or different senses.

Continuum. Can things be more or less alive? Is life a black-or-white Boolean property, or a continuum property with many shades of gray? Common sense leans towards the Boolean view: a rabbit is alive and a rock isn't, end of story. But there are borderline cases like viruses that are unable to replicate without a host. And spores or frozen bacteria remain dormant and unchanging indefinitely but then come back to life when conditions become favorable. Are viruses and spores fully alive? Furthermore, when the original life forms emerged from a pre-biotic chemical soup, they differed very little from their non-living predecessors. Some conclude that there is a continuum of more or less alive things (e.g., Cairns-Smith, 1985; Emmeche, 1994; Dennett, 1995). An alternative is to accept a sharp distinction between life and non-life, but allow that a small step could cross it. The four-fold distinction between things that are (i) inanimate and forever incapable of living, (ii) now living, (iii) dead but formerly living, or (iv) dormant but capable of becoming alive again helps explain away some borderline cases by reclassifying them (e.g., seeds and spores are dormant and not currently living). But it does not fully resolve the continuum puzzle, for there are borderline cases in the four-fold distinction, such as between being dead and alive.

Strong artificial life. Artificial life software and hardware raise the question whether our computer creations could ever literally be alive (Langton, 1989a; Pattee, 1989; Sober, 1992; Emmeche, 1992; Olson, 1997). On the one hand, certain distinctive carbon-based macromolecules play a crucial role in the vital processes of all known living entities; on the other hand, much of artificial life seems to presuppose that life can be realized in a suitably programmed computer. It is important to distinguish two questions here. The first is the philosophically controversial question – *in virtue of what* a computer or a robot could be said to be alive. If this issue were settled, we would face the technical question of whether it is possible to create a software system or hardware device (e.g., a robot) that is literally alive in this sense. The challenge here is whether we could, in fact, realize the processes that were specified in the appropriate materials. The "strong" artificial life position about software is that an instantiation of artificial life software could literally be alive. There is an analogous strong position about "hard" artificial life hardware constructions, and also about "wet" artificial life laboratory constructions. These strong positions contrast with the uncontroversial "weak" positions that computer models, hardware constructions, and wet lab productions are just useful for understanding living systems. And yet, the strong version of wet artificial life is intuitively plausible; we usually accept that something synthesized from scratch in the lab could be literally alive. So the controversy about strong artificial life concerns primarily soft and hard artificial life.

Mind. Another puzzle is whether there is any intrinsic connection between life and mind. Plants, bacteria, insects, and mammals, for example, have various kinds of sensitivity to the environment, various ways in which this environmental sensitivity affects their behavior, and various forms of inter-organism communication (e.g., Dennett, 1997). These are all forms of intelligent behavior, and the relative sophistication of these "mental" capacities seems to correspond to, and explain the relative sophistication of, those forms of life. So it is natural to ask whether life and mind have some deep connection. Evolution creates a genealogical connection between life and mind, of course, but they would be much more deeply unified if Beer is right that "it is adaptive behavior, the . . . ability to cope with the complex, dynamic, unpredictable world in which we live, that is, in fact, fundamental [to intelligence itself]" (Beer, 1990, p.11; see also Maturana & Varela, 1987; Godfrey-Smith, 1994; Clark, 1997). Since all forms of life must cope in one way or another with a complex, dynamic, and unpredictable world, perhaps this adaptive flexibility inseparably connects life and mind.

4. Accounts of Life

There have been various attempts to state the universal characteristics of all forms of life. In this section, I will discuss the main varieties of such accounts of life, indicating some of their motivations, strengths, and weaknesses. I will also note some skeptical positions that deny the usefulness of such accounts.

First, consider the skeptical position that the nature of life is largely irrelevant to biology (Sober, 1992; Taylor, 1992). The reason for this skepticism is that biologists can continue with their biological research whether or not life can be adequately defined, and no matter what view of life prevails in the end. One must admit, though, that recent developments such as attempting to make minimal artificial cells from scratch does require scientists to start to articulate their views about what is essential to life, even if these views fall short of a precise definition. So the issue is no longer irrelevant, if it ever was. For one can set out to construct a minimal form of life only if one has at least a working hypothesis about life's minimally sufficient conditions. Otherwise one would have no idea what to try to make.

A second form of skepticism is the view that life cannot be captured by necessary and sufficient conditions, but instead consists of just a cluster of things sharing only a Wittgenstinian family resemblance. Different forms of life might share various properties or hallmarks, but the individual properties in the cluster each have exceptions. The properties would typically be possessed by living organisms but they would not be strictly necessary or sufficient. Farmer and Belin list eight hallmarks: process; self-reproduction; information storage of self-representation; metabolization; functional interactions with the environment; interdependence of parts; stability under perturbations; and membership in a population with the ability to evolve. They then explain that a cluster conception of life arises from their despair at finding anything more precise than this list of hallmarks.

> There seems to be no single property that characterizes life. Any property that we assign to life is either too broad, so that it characterizes many non-living systems as well, or too specific, so that we can find counter-examples that we intuitively feel to be alive, but that do not satisfy it. (Farmer & Belin, 1992, p.818; see also Taylor, 1992)

The cluster conception amounts to skepticism about the possibility of a unified theory of life.

An advantage of the cluster conception is that it offers a natural explanation for borderline cases. All cluster concepts inevitably have borderline cases. A characteristic of the cluster conception is that it cannot explain why forms of life are unified by one set of hallmarks rather than another. The cluster view must simply accept the hallmarks as given, and then identify the cluster with those hallmarks. Thus, this view can identify life's hallmarks only post hoc; it cannot predict or explain the hallmarks. Those who think that there should be an explanation for life's hallmarks will therefore find the cluster conception unsatisfying.

Another kindred form of skepticism questions the idea that life is a natural kind. Keller (2002) says that life is a human kind, not a natural kind, that is, a distinction created by us, not a distinction in nature. This could explain borderline cases. Since the concept of life changes with the progress of science and technology, one should expect its boundaries to change, thus creating borderline cases. The view also provides some general ammunition against life's puzzles, for a mutable human construct can be expected to spawn puzzles. Keller's argument that life is a human kind suggests that the present presupposition that life has an essence arose only 200 years ago, that the search for life's essence is driven by attempts to make life from non-life (and this tends to dissolve the boundary between life and non-life) and that the new concepts generated by scientific and technological progress violate older taxonomies like the life/non-life distinction (Keller, 2002).

There are problems with all of these arguments. First, all modern scientific concepts like matter and energy arose at some point in human history and have evolved since then. So contingent, datable recent origin does not show that a kind is a human kind, unless it does so at one fell swoop for all scientific concepts. Second, bridging the gap in the laboratory from the non-living to the living need not dissolve the boundary between life and non-life, any more than making the first airplane dissolved the distinction between flying and not flying. Remember that we are seeking the nature of life, not just current conceptions of life.

Now, one answer to the question "what is life?" is simply to give a taxonomy of living things. This is taking the question as a request for an exhaustive list of the kinds of things on the Earth that are alive. This is an interesting historical question, but one riddled with contingencies. The taxonomy is necessarily silent about forms of life that could have existed but did not. This illustrates the taxonomy view's chauvinism in assuming that life as we know it exhausts what life is or could be. Unrelated life forms that exist on an extra-terrestrial site like Europa are absent from all such taxonomies. In any case, we should welcome having our taxonomies adjusted by scientific and technological progress, for that is how we learn.

Some have given a biochemical definition of life. They attempt to specify the biochemical properties that any form of life must have, given the general constraints set by physics and chemistry (Pace, 2001; Benner, Ricardo, & Carrigan, 2004). This includes thermodynamic limits, energetic limits, material limits, and even geographical limits. The features in a biochemical definition are sometimes called life's biochemical "universals." A biochemical definition always presupposes a prior account of life; it states the physical, chemical, and biological possibilities for any biochemical system meeting that prior account of life. The biochemical definitions of Pace (2001) and Benner et al. (2004) presuppose a definition of life based on evolution, so Pace and Benner dwell on the biochemical universals for genetic capacities and emphasize molecules like DNA that can store and transmit information between generations. Biochemical definitions are often myopic and presume that all possible life forms are quite similar to the familiar ones. One could imagine starting with a different conception of life, such as the view based on metabolism, and ending up emphasizing different biochemical universals, such as those that enable open systems to retain their structure in the face of the second law of thermodynamics.

A genetic instance of a biochemical definition of life is Venter's recent genomic definition of life as a minimal genome sufficient to support life (Hutchison et al., 1999). This view inherits the limitations of biochemical definitions. The genomic definition captures the simplest known set of genes sufficient for life. It does not capture genes found in every life form, for the same essential life functions can be achieved by different genes. Many people would question the molecular definition's limitation to genetic properties, on the grounds that life centrally involves much more than genes (Cho et al., 1999).

Everyone in the community of scientists making artificial cells from scratch or "protocells" admits that the nature of life is controversial and contentious, but almost all share the goal of making a self-contained system that metabolizes and evolves (e.g., Rasmussen et al., 2004). That is, an artificial cell is viewed as any chemical system that chemically integrates three processes: The first is the process of assembling some kind of container, such as a lipid vesicle, and living inside it. The second is the metabolic processes that repair and regenerate the container and its contents, and enable the whole system to reproduce. Those chemical processes are shaped and directed by a third chemical process involving encoded information about the system stored in the system ("genes"). Errors ("mutations") can occur when this information is reproduced, so the systems can evolve by natural selection. The integrated-triad view of life requires that the chemical processes of containment, metabolism, and evolution support and enable each other, so that there is functional feedback among all three. This view of protocellular life as an integrated triad of functions accepts any biochemical realization of the triad as genuine life.

The past generation of the philosophy of mind has been dominated by functionalism: the view that mental beings are a certain kind of input–output device and that having a mind is simply having a set of internal states that causally interact (or "function") with respect to each other and with respect to environmental inputs and behavioral outputs in a certain characteristic way. Functionalism with respect to life is the analogous view that being alive is simply realizing a network of processes that interact in a certain characteristic way. Some processes (such as information processing, metabolization, purposeful activity) operate within the organism's lifetime; other processes (such as self-reproduction and adaptive evolution) operate over many generations. These processes are always realized in some material substratum, but the substratum's material nature is irrelevant so long as the *forms* of the processes are preserved. For these reasons, functionalism is an attractive position with respect to life. Chris Langton's defense of artificial life is a classic statement of the case for functionalism with respect to life:

> Life is a property of *form*, not *matter*, a result of the organization of matter rather than something that inheres in the matter itself. (Langton, 1989a, p.41)

> The big claim is that a properly organized set of artificial primitives carrying out the same functional roles as the biomolecules in natural living systems will support a process that is "alive" in the same way that natural organisms are alive. Artificial Life will therefore be genuine life – it will simply be made of different stuff than the life that has evolved here on Earth. (Langton, 1989a, p.33)

We might be unsure about the details of the processes that are definitive of life, and we might wish to reserve judgment about whether artificial life creations are genuinely alive. Nevertheless, it is hard to deny Langton's point that life's characteristic processes like metabolism, information processing, and self-reproduction could be realized in a wide and potentially open-ended range of materials. Thus, the prospects for some form of functionalism with respect to life seem bright.

The main challenge for functionalism with respect to mind concerns consciousness and qualia. It is worth noting that functionalism about life does not face any analogous problems. Another challenge for functionalism with respect to mind is to explain how people's mental states are meaningful or have semantic content. Darwinian natural selection provides a naturalistic explanation of many biological functions of structures in evolved forms of life. This biological functionality gives the internal states of living creatures a kind of meaning or semantic content, so that we can speak of a creature trying to find food for nourishment. Many philosophers are optimistic that the meaning problem in functionalism with respect to mind will be solved by some analogous Darwinian explanation of the biological function of mental states (e.g., Dennett, 1995).

Another apparent threat to functionalism with respect to life is the suggestion that the processes involved in life are, in some relevant sense, unformalizable or non-computational (e.g., Emmeche, 1992). Bedau (1999) thinks that the apparent non-computational quality of life can be explained. Advantageous traits that arise through mutations tend, ceteris paribus, to persist and spread through the population. Furthermore, trait frequencies in the population will tend, ceteris paribus, to change in a way that is generally apt for the population in its exogenously changing environment. These dynamical patterns in trait frequencies emerge as a statistical pattern from the micro-level contingencies of natural selection, mutation, drift, etc. Bedau argues that there is often a special kind of suppleness in these patterns. Such patterns in trait frequencies are not precise and exceptionless universal generalizations, but instead hold only for the most part, only ceteris paribus. Furthermore, those regularities have exceptions that sometimes "prove the rule" in the sense that they are a byproduct of trying to achieve some deeper adaptive goal. For example, Bedau describes a system in which mutation rates can evolve and shows that the mutation rates tend to evolve so as to keep the population's gene pool at the "edge of disorder"; but this regularity has exceptions, some of which are due to the operation of a deeper regularity about mutation rates evolving so as to optimally balance evolutionary "memory" and "creativity" (for details, see Bedau, 1999). In this sort of way, supple regularities reflect an underlying capacity to respond appropriately in an open-ended variety of contexts. This explains a certain kind of unformalizability of life processes, though it also allows life to be captured in appropriate computer models.

Functionalism leaves unanswered exactly which processes play what role in the functional characterization of life. Persisting in the face of the second law of thermodynamics by means of metabolism is the defining process of life according to Schrödinger's influential account:

> When is a piece of matter said to be alive? When it goes on "doing something", moving, exchanging material with its environment, and so forth, and that for a much longer period than we would expect an inanimate piece of matter to "keep going" under similar circumstances . . . It is by avoiding the rapid decay into the inert state of "equilibrium" that an organism appears so enigmatic; . . . How does the living organism avoid decay? The obvious answer is: By eating, drinking, breathing and (in the case of plants) assimilating. The technical term is metabolism . . . (Schrödinger, 1969, pp.74–6)

Metabolism-centered views of life attract many (Margulis & Sagan, 1995; Boden, 1999). They are closely related to views that focus on autopoeisis (Varela, Maturana, & Uribe, 1974; Maturana & Varela, 1987).

The view that metabolism is life's central process has some clear advantages, such as explaining our intuition that a crystal is not alive (there is a metabolic flux of molecules only at the crystal's edge, not inside it). Also, the fact that metabolism is needed to combat entropy implies that metabolism is at least a necessary condition of all physical life forms. Metabolism also naturally explains the four-fold distinction between the non-living, living, dead, and dormant. The non-living cannot metabolize in principle, and the living are now metabolizing. The dead were once living and metabolizing, but now they are decaying. The dormant were once living but now do not metabolize, but they could resume metabolizing given the right circumstances.

The main drawback of metabolism as an all-encompassing account of life is that many metabolizing entities seem intuitively not to be alive or to involve life in any way. Standard examples include a candle flame, a vortex, and a convection cell (Maynard Smith, 1986; Bagley & Farmer, 1992). Such examples by themselves do not prove conclusively that metabolism is insufficient for life, for pre-theoretic intuitive judgments can be wrong. The question is whether on balance metabolism adequately explains life's hallmarks and resolves life's puzzles.

Some think that the central feature underlying all life is the open-ended evolutionary process of adaptation. The central idea is that what distinguishes life is its automatic and open-ended capacity (within limits) to adapt appropriately to unpredictable changes in the environment. From this perspective, what is distinctive of life is the way in which adaptive evolution automatically fashions new and intelligent strategies for surviving and flourishing as local contexts change. Maynard Smith (1975, p.96f; see also Mayr, 1982; Cairns-Smith, 1985) succinctly explains the justification for the view that life crucially depends on the evolutionary process of adaptation:

> We shall regard as alive any population of entities which has the properties of multiplication, heredity and variation. The justification for this definition is as follows: any population with these properties will evolve by natural selection so as to become better adapted to its environment. Given time, any degree of adaptive complexity can be generated by natural selection.

These remarks suggest how the process of adaptive evolution could explain life's hallmarks, borderline cases and puzzles (see Bedau, 1998).

There are a few characteristic criticisms of such evolution-centered views. One is purported counterexamples of creatures that are alive but cannot give birth (mules, old people, etc.) and so cannot contribute to the process of evolution. The typical response is to require that organisms be produced by an evolutionary process, but not that they necessarily can affect further evolution. Another kind of purported counterexample is a clearly non-living system, such as a population of clay crystallites or a free market economy, which evolves by natural selection. Some think that we should accept these unintuitive examples because evolution-centered views provide such a compelling explanation of life's hallmarks, borderline cases, and puzzles (e.g., Bedau, 1998).

Not all of these positions are competing; many are consistent. For example, functionalism is consistent with the protocell integrated-triad account of minimal life. Also, accounts of the nature of life each entail a biochemical characterization of life, and many accounts of life overlap. The problem of understanding life is to identify exactly which of these accounts is true.

5. The Problem of Understanding Life

How should we compare and evaluate accounts of the nature of life? One straightforward answer is simply to see how well each explains the phenomena of life. This amounts to doing three things: explaining life's hallmarks, explaining the borderline cases, and resolving the puzzles about life. The problem of understanding life is the problem of explaining these three things.

One initial difficulty is confusion about what question is at stake. Some investigations think the key test for any account of life is to fit it with our pre-theoretic intuitions about which things are alive and which are not (e.g., Boden, 1999). But one should ask why we should emphasize such intuitions. A good theory of life might make us reconceptualize and recategorize life. This might change our attitudes about exactly which cases are the ones in which life is present. Thus, although they have some weight, our pre-theoretic intuitions are not inviolable.

One could also ask about the *meaning of the word* "life" in today's English. But the stereotypes associated with the term "life" are commonplaces and reflect the lowest common denominator of our current shared picture of life. So we are not likely to learn much about life by relying on what "life" means.

Nor are we likely to learn much by analysis of the *concept* of life. As with the meaning of "life," our current concept of life will reflect our current understanding of life. If we want to learn the real nature of the phenomena with life's hallmarks, borderline cases, and puzzles, we should study the natural phenomena themselves, not our words or concepts. And we should expect our understanding of the phenomena of life to evolve and sometimes improve.

Explaining the phenomena of life involves at least a rough view of life's essence or nature, and perhaps even a rough definition of life. Scientific essentialism, originating from Kripke (1980), is the philosophical view that the essence of natural kinds like water and gold is their underlying causal powers, which are discovered by empirical science (see Bealer, 1987). The essence of substances like water and gold turns out to be their underlying chemical composition. Life, on the other hand, is a certain kind of flexible process, not a fixed chemical substance. So unlike water or gold, life's nature would presumably be captured by the characteristic network of processes (such as metabolism, reproduction, and sensation) that explains its characteristic causal powers. In this regard life is more like heat, which is a certain process in matter (high molecular kinetic energy). A specific temperature (say, 23°C) is a specific kind of process that can occur in all kinds of matter. Life is also a kind of process that can occur in different kinds of material, but unlike temperature not *all* kinds of material can be alive. Mapping the biochemical constraints on the kinds of substances that could instantiate life yields a biochemical definition of life (recall above). Note that scientific essentialism about life might be true, even if contemporary science has reached no consensus about life. Scientific essentialism is a philosophical view about the method by which life's essence would be discovered – it is not a view about the particular content of that essence. The details of the scientific essentialist definition of life might need to await further scientific progress.

It is unclear whether living things have any features that make them essentially alive. In Dennett's opinion, for example, the life/non-life distinction is a matter of degree and life is too "interesting" to have an essence (1995, p.201). In fact, contemporary biology and philosophy of biology thoroughly embrace a Darwinian anti-essentialism according to which species have no essence and their members share no necessary and sufficient properties. Instead, the similarities among the members of a species are only statistical. Species are no more than a cloud or clump in an abstract possible feature space. Although some sub-regions of possible feature space are unoccupied because they are maladaptive, it is an accident exactly which of the acceptable sub-regions are occupied. No sub-regions are any more natural than any other; none are privileged by fixed and immutable Platonic essences. The generalization of this anti-essentialism probably helps account for why so many philosophers are attracted to the cluster concept of life, for that seems like a direct consequence of anti-essentialism.

Darwinian anti-essentialism is directed against a narrow notion of essence that embraces exception-less necessary and sufficient conditions and excludes borderline cases. Borderline cases are one of the hallmarks of life, so the nature of life must be broad and flexible enough to embrace borderline cases. One could embrace Darwinian anti-essentialism but still accept scientific essentialism about life. On this view, the "essence" of life would be whatever process explains the phenomena of life, including life's hallmarks, borderline cases, and puzzles. Life would not be defined by exceptionless conditions but empirically. It is unfortunate that contemporary philosophical terminology obscures that Darwinian anti-essentialism and scientific essentialism about life are compatible.

Clelland and Chyba (2002) argue that it is too early to formulate definitions of life, because our current understanding of life is too limited. They conclude that we should put off formulating definitions until scientists can tell much more about the different forms that life could take. Now might nevertheless be the right time to construct tentative and testable hypotheses about the phenomena of life. These hypotheses will likely be false, but they can aid our search for better theories (Wimsatt, 1987). When we have good theories of life in hand, we can extract their implied definitions of life. So the quest for the definition of life is better recast as the quest for the nature of life.

Life is one of the most fundamental and complex aspects of nature. So accounts of life are rich and interesting, with a complicated structure. They come in many forms, including skepticism, detailed biochemical and molecular descriptions, and abstract functionalism, and they emphasize fundamental biological processes like metabolism and evolution. The criteria for evaluation include their ability to explain life's hallmarks and borderline cases and their ability to resolve the puzzles about life. Many of the main accounts of life still lack substantial development and careful evaluation along a number of these dimensions. Thus, the problem of understanding life is still wide open.

References

Bagley, R., & Farmer, J. D. (1992). Spontaneous emergence of a metabolism. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial Life II* (pp. 93–140). Redwood City, CA: Addison-Wesley.

Bealer, G. (1987). The philosophical limits of scientific essentialism. In J. Tomberlin (Ed.). *Philosophical perspectives 1* (pp. 289–365). Atascadero: Ridgeway.

Bedau, M. A. (1991). Can biological teleology be naturalized? *The Journal of Philosophy*, 88, 647–55.

Bedau, M. A. (1996). The nature of life. In M. Boden (Ed.). *The philosophy of artificial life* (pp. 332–57). New York: Oxford University Press.

Bedau, M. A. (1997). Weak emergence. In J. Tomberlin (Ed.). *Philosophical perspectives: mind, causation, and world* (vol. 11, pp. 375–99). Oxford: Blackwell.

Bedau. M. A. (1998). Four puzzles about life. *Artificial Life*, 4, 125–40.

Bedau, M. A. (1999). Supple laws in biology and psychology. In V. Hardcastle (Ed.), *Where biology meets psychology: philosophical essays* (pp. 287–302). Cambridge: MIT Press. (The printed version of this paper inadvertently omits a few pages; the full text is available on the web via http://www.reed.edu/~mab. Accessed 11 October 2007)

Bedau, M. A. (2003a). Artificial life: organization, adaptation, and complexity from the bottom up. *Trends in Cognitive Science*, 7(11, November), 505–12.

Bedau, M. A. (2003b). Downward causation and autonomy in weak emergence. *Principia Revista Inernacional de Epistemologica*, 6(1), 5–50.

Beer, R. D. (1990). *Intelligence as adaptive behavior: an experiment in computational neuroethology.* Boston: Academic Press.

Benner, S. A., Ricardo, A., & Carrigan, M. A. (2004). Is there a common chemical model for life in the universe? *Current Opinion in Chemical Biology*, 8, 679–89.

Boden, M. A. (1999). Is metabolism necessary? *British Journal for the Philosophy of Science*, 50, 231–48.

Brent, R. (2004). A partnership between biology and engineering. *Nature Biotechnology*, 22(10), 1211–14.

Brooks, R. (2002). *Flesh and machines: how robots will change us.* New York: Pantheon.

Cairns-Smith, A. G. (1985). *Seven clues to the origin of life.* Cambridge: Cambridge University Press.

Cho, M. K., Magnus, D., Caplan, A. L., McGee, D., & the Ethics of Genomics Group. (1999). Ethical considerations in synthesizing a minimal genome. *Science*, 286(5447), 2087–90.

Clark, A. (1997). *Being there: putting brain, body, and world together again.* Cambridge, MA: MIT Press.

Cleland, C., & Chyba, C. (2002). Defining "life". *Origins of Life and Evolution of the Biosphere*, 32, 387–93.

Crick, F. (1981). *Life itself: its origin and nature.* New York: Simon & Schuster.

Dennett, D. C. (1997). *Kinds of minds: towards an understanding of consciousness.* New York: Basic Books.

Dennett, D. C. (1995). *Darwin's dangerous idea.* New York: Simon & Schuster.

Dyson, F. (1999). *Origins of life* (rev. edn). Cambridge: Cambridge University Press.

Eigen, M. (1992). *Steps toward life.* Oxford: Oxford University Press.

Emmeche, C. (1992). Life as an abstract phenomenon: is artificial life possible? In F. Varela & P. Bourgine (Eds). *Towards a practice of autonomous systems* (pp. 466–74). Cambridge, MA: Bradford Books/MIT Press.

Emmeche, C. (1994). *The garden in the machine: the emerging science of artificial life.* Princeton: Princeton University Press.

Farmer, D., & Belin, A. (1992). Artificial life: the coming evolution. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial Life II* (pp. 815–40). Redwood City, CA: Addison-Wesley.

Gánti, T. (2003). *The principles of life.* New York: Oxford University Press. With a commentary by James Grisemer and Eörs Szathmáry.

Gibbs, W. W. (2004). Synthetic life. *Scientific American*, 290(April), 75–81.

Godfrey-Smith, P. (1994). Spencer and Dewey on life and mind. In R. Brooks & P. Maes (Eds). *Artificial Life IV* (pp. 80–9). Cambridge, MA: MIT Press/Bradford Books.

Hutchison, C. A., Peterson, S. N., Gill, S. R., Cline, R. T., White, O., Fraser, C. M., Smith, H. O., & Venter, J. C. (1999). Global transposon mutagenesis and a minimal Mycoplasma genome. *Science*, 286(10 December), 2165–9.

Keller, E. F. (2002). *Making sense of life: explaining biological development with models, metaphors, and machines.* Cambridge, MA: Harvard University Press.

Kim, J. (1999). Making sense of emergence. *Philosophical Studies*, 95, 3–36.

Kripke, S. (1980). *Naming and necessity.* Cambridge, MA: Harvard University Press.

Langton, C. (1989a). Artificial life. In C. Langton (Ed.). *Artificial life* (pp. 1–47). Redwood City, CA: Addison-Wesley.

Langton, C. G. (Ed.) (1989b). *Artificial life.* Redwood City: Addison-Wesley.

Lipson, H., & Pollack, J. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(31 August), 974–8.

Luisi, P. L. (1998). About various definitions of life. *Origins of Life and Evolution of the Biosphere*, 28, 613–22.

Luisi, P. L. (2006). *The emergence of life: from chemical origins to synthetic biology.* Cambridge: Cambridge University Press.

Margulis, L., & Sagan, D. (1995). *What is life?* Berkeley: University of California Press.

Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge: the biological roots of human understanding* (rev. edn, 1992). Boston: Shambhala.

Maynard Smith, J. (1975). *The theory of evolution* (3rd edn). New York: Penguin.

Maynard Smith, J. (1986). *The problems of biology.* New York: Oxford University Press.

Maynard Smith, J. (1992). Byte-sized evolution. *Nature*, 355(27 February), 772–3.

Mayr, E. (1982). *The growth of biological thought.* Cambridge, MA: Harvard University Press.

Mayr, E. (1997). What is the meaning of "life"? In E. Mayr, *This is biology: the science of the living world* (pp. 1–23). Cambridge, MA: Harvard University Press.

Morowitz, H. J. (1992). *Beginnings of cellular life: metabolism recapitulates biogenesis.* New Haven: Yale University Press.

Olson, E. T. (1997). The ontological basis of strong artificial life. *Artificial Life*, 3, 29–39.
Oparin, A. I. (1964). *Life: its nature, origin, and development* (A. Synge, Trans.). New York: Academic Press.
Pace, N. R. (2001). The universal nature of biochemistry. *Proceedings of the National Academy of Science USA*, 98(January 30), 805–8.
Pattee, H. H. (1989). Simulations, realization, and theories of life. In C. Langton (Ed.). *Artificial life* (pp. 63–78). Redwood City, CA: Addison-Wesley.
Pennock, R. (2001). *Intelligent design creationism and its critics: philosophical, theological & scientific perspectives.* Cambridge, MA: MIT Press.
Rasmussen, S., Chen, L., Deamer, D., Krakauer, D., Packard, N. H., Stadler, P. F., & Bedau, M. A. (2004). Transitions from nonliving to living matter. *Science*, 303(February 13), 963–65.
Rasmussen, S., Bedau, M. A., Chen, L., Deamer, D., Krakauer, D. C., Packard, N. H., Stadler, P. F. (Eds). (2007). *Protocells: bridging nonliving and living matter.* Cambridge, MA: MIT Press.
Ray, T. (1992). An approach to the synthesis of life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial life II* (pp. 371–408). Redwood City, CA: Addison-Wesley.
Schrödinger, W. (1969). *What is life? The physical aspect of the living cell.* Cambridge: Cambridge University Press.
Shapiro, R. (1986). *Origins: A skeptic's guide to the creation of life on Earth.* New York: Summit Books.
Sober, E. (1992). Learning from functionalism – prospects for strong artificial life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial Life II* (pp. 749–65). Redwood City, CA: Addison-Wesley.
Szostak, J. W., Bartel, D. P., & Luisi, P. L. (2001). Synthesizing life. *Nature*, 409(January 18), 387–90.
Taylor, C. (1992). "Fleshing out" artificial life II. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial Life II* (pp. 25–38). Redwood City, CA: Addison-Wesley.
Varela, F. G., Maturana, H. R., & Uribe, R. (1974). Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems*, 5, 187–96.
Wimsatt, W. C. (1987). False models as means to truer theories. In M. Niteckiand, & A. Hoffman (Eds). *Neutral modes in biology* (pp. 23–55). Oxford: Oxford University Press.
Zimmer, C. (2003). Tinker, tailor: can Venter stitch together a genome from scratch? *Science*, 299(February 14), 1006–7.

Further Reading

Bedau, M. A., & Packard, N. H. (1992). Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds). *Artificial Life II* (pp. 431–61). Redwood City, CA: Addison-Wesley.
Cleland, C. E., & Copley, S. H. (2005). The possibility of alternative microbial life on Earth. *International Journal of Astrobiology*, 4, 165–73.
Code, A., & Moravcsik, J. (1992). Explaining various forms of living. In M. C. Nussbaum & A. O. Rorty (Eds). *Essays on Aristotle's* De Anima (pp. 129–45). Oxford: Clarendon Press.
Emmeche, C. (1994). Is life a multiverse phenomenon? In C. G. Langton (Ed.). *Artificial life III* (pp. 553–68). Redwood City: Addison-Wesley.
Feldman, F. (1992). *Confrontations with the reaper: a philosophical study of the nature and value of death.* New York: Oxford University Press.
Haldane, J. B. S. (1937). What is life? In J. B. S. Haldane, *Adventures of a biologist* (pp. 49–64). New York: Macmillan.

Jonas, H. (1966). *The phenomenon of life: toward a philosophical biology.* New York: Dell.
Keller, E. F. (1995). *Refiguring life: metaphors in twentieth-century biology.* New York: Columbia University Press.
Korzeniewski, B. (2001). Cybernetic formulation of the definition of life. *Journal of Theoretical Biology*, 209, 275–86.
Lange, M. (1996). Life, "artificial life," and scientific explanation. *Philosophy of Science*, 63, 225–44.
Maynard Smith, J., & Szathmáry, E. (1999). *The origins of life: from the birth of life to the origins of language.* New York: Oxford University Press.
Miller, J. G. (1978). *Living systems.* New York: McGraw-Hill.
Murphey, M. P., & O'Neill, L. A. J. (Eds). (1993). *What is life? The next fifty years: speculations on the future of biology.* Cambridge: Cambridge University Press.
Richards, R. J. (2002). *The Romantic conception of life: science and philosophy in the age of Goethe.* Chicago: University of Chicago Press.
Rosen, R. (1991). *Life itself: a comprehensive inquiry into the nature, origin, and fabrication of life.* New York: Columbia University Press.
Sagan, C. (1970). Life. *Encyclopedia Britannica* (15th edn, vol. 10). New York: Macropaedia.
Sterelny, K., & Griffths, P. (1999). What is life? In K. Sterelny & P. Griffiths, *Sex and death: an introduction to philosophy of biology* (pp. 357–77). Chicago: Chicago University Press.