

Four Puzzles About Life

Mark A. Bedau
Reed College
3203 SE Woodstock Blvd.
Portland, OR 97202
mab@reed.edu

Abstract To surmount the notorious difficulties of defining life, we should evaluate theories of life not by whether they provide necessary and sufficient conditions for our current preconceptions about life but by how well they explain living phenomena and how satisfactorily they resolve puzzles about life. On these grounds, the theory of life as supple adaptation gets support from the natural and compelling way it resolves the following four puzzles: (a) How are different forms of life at different levels of the vital hierarchy related? (b) Is there a continuum between life and nonlife? (c) Does life essentially concern a living entity's material composition or its form? (d) Are life and mind intrinsically connected?

Keywords

definition of life, theory of life, supple adaptation

1 What Explains the Phenomena of Life

Life seems to be one of the most basic kinds of actual natural phenomena. A bewildering variety of forms of life surrounds us, but we usually have no difficulty distinguishing the living from the nonliving. That flower, that mushroom, that worm, that bird are alive; that rock, that mountain, that river, that cloud are not. Just as any attempt to divide nature at its joints must account for mind and matter, so it must account for life.

Yet it is notoriously difficult to say what life is, exactly. Many have noted this (e.g., [59]); Farmer and Belin ([27], p. 818) put the point this way:

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 nonliving 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.

The fact today is that we know of no set of individually necessary and jointly sufficient conditions for life.

Nevertheless, there is broad agreement that life forms share certain distinctive hallmarks. Various hallmarks are discussed in the literature, and Mayr's list [47] is representative and influential:

1. All levels of living systems have an enormously complex and adaptive organization.
2. Living organisms are composed of a chemically unique set of macromolecules.
3. The important phenomena in living systems are predominantly qualitative, not quantitative.
4. All levels of living systems consist of highly variable groups of unique individuals.
5. All organisms possess historically evolved genetic programs which enable them to engage in teleonomic processes and activities.

6. Classes of living organisms are defined by historical connections of common descent.
7. Organisms are the product of natural selection.
8. Living processes are especially unpredictable.

I agree with Mayr ([47], p. 59) that the coexistence of these properties “make[s] it clear that a living system is something quite different from any inanimate object,” so I suspect that there is some unified explanation of vital phenomena. At the same time, it is puzzling that such heterogeneous properties characteristically coexist in nature, especially because each of the hallmarks can be possessed by nonliving things.

Appearances can be deceptive. Vital phenomena might have no unified explanation and life might not be a basic kind of natural phenomena. Skeptics such as Sober [58] think that the question of the nature of life, in general, has no interesting answer. But I suspect otherwise, along with those (e.g., [18, 54]) searching for extraterrestrial life; they are not searching just for extraterrestrial metabolizers and self-reproducers. Likewise for those searching for the origin of life on Earth (e.g., [17, 24]). Likewise for those in the field of artificial life attempting to synthesize life in artificial media [42]. This broadly based search for a unified theory of vital phenomena should retreat to skepticism, if at all, only as a last resort.

So, we face a quandary: We expect there is a unified theory of life but we doubt life has necessary and sufficient conditions. We can resolve this quandary if we reconceive our project in two ways. First, we must focus on the *phenomena* of life, not our current concept of life nor the current meaning of our word “life.” Physicists and chemists want to explain matter itself, not our current concept of matter or the current meaning of the word “matter.” I want to explain life itself; such an explanation is what I mean by a *theory* of life. It does not matter whether this theory supports our current preconceptions about life or fits the current meaning of our word “life.” Our current concept of life and the current meaning of our word “life” are contingent. They vary across space and time, changing with different human cultures at different places and in different ages. Our theories are connected with our concepts and words, of course, but the connection goes in the other direction, with our concepts and words following the lead of our currently best theories.

The second step in resolving our quandary is to shift our focus from living organisms to the process that *produces* organisms and other living phenomena. The generating process is primary and its products are secondary, for the process provides a unified explanation of the various products. Understanding how organisms and other living entities actually come into existence is the key to understanding what they are.

I believe that the process of supple adaptation is the primary form of life [9]. I defend this proposal here, on the grounds, not that supple adaptation is a necessary and sufficient condition for living organisms (it is not), nor that it matches our current concept of life or the current meaning of the word “life” (it might not), but that it provides the best unified explanation of the phenomena of life. Theories of life should be judged in part by how well they resolve basic puzzles about life. My specific concern here is how supple adaptation resolves four such puzzles. I propose no complete and final theory of life, nor definitive resolutions to the four puzzles. But I show that supple adaptation provides good explanations of the puzzles.

Can any rival theory explain the four puzzles as well? It is easy to dream up rival theories and to imagine that they have good explanations of the puzzles; it is another thing to support such dreams with substantial evidence. The theory of life as supple adaptation does not automatically fend off rival theories. Another theory that provided equally good explanations of the puzzles would be a serious contender. My goal here

is not to show that credible contenders are impossible but to establish what standard they must meet to be taken seriously.

2 The Theory of Life as Supple Adaptation

I propose that an automatic and continually creative evolutionary process of adapting to changing environments is the primary form of life. My proposal is broadly in the spirit of genetic definitions of life [57]; various similar proposals occur in the literature (e.g., [17, 45, 46]). From my perspective, what is distinctive of life is the way in which evolution automatically fashions and refashions appropriate strategies for coping as local contexts change.

The notion of propriety involved in supple adaptation is to be understood teleologically. A response is appropriate only if it promotes and furthers the adapting entity's intrinsic goals and purposes, where those goals and purposes are minimally to survive and, more generally, to flourish. For example, if a clam's shell becomes cracked, then an inappropriate response would be for the clam's soft tissue to ooze out the crack, and an appropriate response would be for the shell to be repaired. By contrast, although water flowing downhill automatically "adapts" to the local landscape's topography, the water has no intrinsic goals or purposes and flowing downhill serves no such goals or purposes. Similarly, a thermostat has no intrinsic goals or purposes, so its "adaptive responses" to local temperature changes can be considered appropriate only relative to the extrinsic goals or purposes that *we* have in using those artifacts. These teleological notions of intrinsic goals and purposes are certainly controversial, and I will not here rehearse my own attempts to resolve these controversies [5–7]. I trust that their connection to the relevant notion of adaptation is clear enough for present purposes.

My proposal is that the thread unifying the diversity of life is the *suppleness* of this process of producing adaptations—its ongoing and indefinitely creative production of significantly new kinds of adaptive responses to significantly new kinds of adaptive challenges and opportunities. A biological arms race [21] is one simple example of supple adaptation. By contrast, a thermostat's response to the ambient temperature is not "supple" in the relevant sense because it is the "same old" kind of response to the "same old" kind of temperature changes. Because the process of supple adaptation involves significantly new kinds of adaptive challenges and opportunities, those challenges and opportunities will be unanticipated by the adapting entities, and they will elicit an open-ended range of appropriate responses. Phrases such as "open-ended evolution" ([43], p. 310; [56], p. 372) or "perpetual novelty" ([38], p. 184) are sometimes used to refer to this process.

Supple adaptation is not to be equated with natural selection. For one thing, natural selection is not necessary for supple adaptation. Other adaptive mechanisms such as Lamarckian selection or Hebbian learning can produce supple adaptation. For another thing, natural selection is not sufficient for supple adaptation. Supple adaptation is the *ongoing* production of significant adaptive novelty. By contrast, the dynamics of natural selection often eventually stabilize in the long run, with the result that significantly new adaptations stop being produced. Even though new mutations continually occur, they yield at best only insignificantly different variants of familiar adaptations. So, natural selection produces supple adaptation only when it is continually creative. Adaptation cannot be continually creative without ongoing environmental change. One way to bring about ongoing environmental change is for the evolving system's own evolution to continually reshape the selection criteria or fitness function [51], perhaps through some mechanism like this: The organisms in the evolving system interact through their behavior. Each organism's environment consists to a large degree of its interactions with other organisms. So, if one organism evolves an innovative adaptive behavior,

this changes the environment of neighboring organisms. This environmental change in turn causes the neighboring organisms to evolve their own new adaptive behaviors, and this finally changes the environment of the original organism. In this way an organism's adaptive evolution ultimately changes the environment of that very organism. The net effect is that the population's adaptive evolution continually drives its own further adaptation.

I should call attention to the difference between a capacity and its exercise, because I hold that life involves the *exercise* of supple adaptation, not just the *capacity* to do so. For me, the key is not supple *adaptability* but actual supple *adaptation*. A system undergoing supple adaptation is not adapting at every moment, of course—the adaptation occurs in fits and starts. But the quiescent periods between adaptive events are transient; every quiescent period is followed by new adaptive events. If a system that *could* undergo supple adaptation never *does*, then by my lights it *could* be alive but never *is*.

Probably the most controversial feature of my theory of life is the claim that supple adaptation does not merely *produce* living entities: The *primary* forms of life are none other than the supplely adapting systems themselves. Other living entities are alive by virtue of bearing an appropriate relationship to a supplely adapting system; they are *secondary* forms of life. Different kinds of living entities (organisms, organs, cells, etc.) stand in different kinds of relationships to the supplely adapting system from which their life ultimately derives. In general, these relationships are ways in which the entity is created and sustained by a supplely adapting system. So, the general form of my theory of life can be captured by this definition:

X is living *iff*

1. *X* is a supplely adapting system, or
2. *X* is explained in the right way by a supplely adapting system.

The effect of this definition is to construe the primary form life as supplely adapting systems.

According to this definition, individual living organisms, organs, cells, and all the other living things count as alive because they are explained in the right way by a supplely adapting system. But the definition does not specify which kinds of explanations are the “right” ones. The explanations typically involve the way in which things are created and sustained, but it is not clear whether this is always true. Furthermore, some ways of being generated and sustained are clearly *not* what is intended by the definition, such as the way in which people create and sustain automobiles and garbage dumps, the way in which spiders create and sustain their webs and beavers their dams. I am leaving these details to be settled by whatever in the future best explains living phenomena, so I am not proposing a complete and final theory here. By claiming that the process of supple adaptation is the central explanatory factor underlying and unifying the various phenomena of life, the definition above delineates the central categories to be used in a final definition and proposes boundaries within which to seek that definition. My aim is not to give a particular definition but to set the stage for one to be produced in the future.

One important virtue of the theory of life as supple adaptation is its unified explanation for Mayr's hallmarks of life. The theory implies that we should expect those heterogeneous-seeming properties to coexist in nature. If life consists of supplely adapting systems and the entities they generate and sustain, we should expect life to involve the operation of natural selection producing complex adaptive organization in historically connected organisms with evolved genetic programs. Furthermore, the

random variation and historical contingency in supple adaptation explains why living phenomena are especially unpredictable and involve unique and variable individuals. Finally, if supple adaptation is produced by a branching process involving birth, reproduction, and death of individuals, such as natural selection, then we can understand why it would give rise to a wealth of qualitative phenomena characterized by frozen accidents like chemically unique macromolecules. The naturalness of all these explanations supports the theory of life as supple adaptation.

Another consideration in favor of the theory is its natural response to potential criticisms. For example, mules, the last living member of an about-to-be extinct species: neutered and spayed animals are all alive, but being infertile, such entities play no role in the supple adaptation of their lineages. However, infertile organisms exist only because of their connections with other fertile organisms that *do* play an active role in a supplely adapting biosphere, so they fall within the scope of my theory.

Some might object that an evolving system's supple adaptation has the wrong logical form to be the nature of life. Individual organisms are the paradigmatic living entities and an evolving population of organisms is of a different logical category than an individual organism. So, one might think that life cannot consist in a population-level property like supple adaptation. Now, individual organisms and populations of organisms are of different categories, to be sure, but phenomena from one category can explain phenomena from other categories. The theory of life as supple adaptation denies that individual organisms are the primary forms of life, but it does so consciously and deliberately, out of the conviction that the process of supple adaptation is our current best hope for unifying and explaining the phenomena of life. If the best explanation for life violates some of our currently dominant paradigms of life, so much the worse for those paradigms.

The possibility of an ecology that has reached a state of stable equilibrium and stopped adapting forever is a more direct challenge to my theory. After all, the organisms in such so-called "climax" ecosystems are certainly alive, yet the ecosystem containing them is not undergoing supple adaptation, so these organisms would seem to fall outside my theory. However, not only do climax ecosystems originate through a process of supple adaptation, but their quiescent periods are transitory. At least, that is the hypothesis behind the theory of life as supple adaptation. If this hypothesis is false and it turns out that climax ecosystems simply do not exhibit supple adaptation, then the theory of life as supple adaptation is also false. It is an empirical question whether the hypothesis is true. My theory implies not that the hypothesis is analytically true (it is not) nor that it is knowable a priori (it is not) but only that the nature of life, in fact, is supple adaptation. Being life's nature, it is an essential property of life and so holds necessarily, but it is a necessity that we learn about a posteriori through empirical science.

It is easy to *conceive* of circumstances that violate my account of life. Nothing prevents us from entertaining with Boden [16] the scientific fantasy of species that never evolve and adapt. For all I know, this is possible; that is, it is "epistemically" possible, as Kripke [41] might say. So is the possibility that there has been and ever will be only one living organism. So is the possibility that all organisms were created in seven days by an omnipotent, omniscient, and omnibenevolent deity. But these fantasies are just that—fantasies, with no bearing on the true nature of any form of life that we could discover or synthesize. My claim—a posteriori to be sure, but still true, I wager—is that all living organisms anywhere in the universe ultimately derive their existence and their characteristic lifelike features from having the right sort of explanatory connection with a system undergoing supple adaptation.

Are there not counterexamples of supplely adapting systems devoid of all life? Viruses are adapting against all our best efforts to eradicate them—the AIDS virus

evolves remarkably quickly—and viruses are a classic example of entities on some borderline between life and nonlife. Even less lifelike are populations of the tiny clay crystallites that make up mud, yet these seem to have the flexibility to adapt and evolve by natural selection [6, 17]. So do autocatalytic networks of chemical species [2], yet evolving populations of crystals or chemicals are ordinarily thought to involve no life whatsoever. Even more extreme examples are individual human mental activity and collective human intellectual, social, and economic activities; these all look like supple, open-ended capacities to adapt to unpredictably changing circumstances, yet none would ordinarily be called alive. Intellectual and economic activities are generated by living creatures, but the evolving intellectual or economic systems themselves are not thought to be alive. However, I am not offering supple adaptation as an explication of our current concept of life, so unintuitive classifications are no particular concern. These counterintuitive cases do not undermine the fact that supple adaptation is our best explanation of the phenomena of life. If life is supple adaptation, then virus and clay crystallite populations, autocatalytic chemical networks, and human intellectual and economic systems all deserve to be thought of as living if they exhibit supple adaptation. Our ordinary language may well reflect some linguistic pressure from this direction, because we speak of the vitality of such systems (though this might be only a metaphor, of course). If we seek to learn the true nature of the phenomena of life, we must be open to the possibility that life is quite unlike what we now suppose.

3 Four Puzzles and Proposed Solutions

We can evaluate a theory of life by how well it resolves persistent puzzles about life. In summary form, this is my present battery of puzzles, along with the resolutions implied by the theory of life as supple adaptation:

Puzzle 1: How are different forms of life at different levels of biological hierarchy related?

Solution: Life must exist at many levels of organization. Different levels involve different but related forms of life.

Puzzle 2: Is the distinction between life and nonlife dichotomous or continuous?

Solution: Various continua and dichotomies separate life and nonlife, but the primary distinction is continuous.

Puzzle 3: Does the essence of life involve matter or form?

Solution: Life is essentially a certain form of process. The suppleness of that form makes the process noncomputational, but a computer simulation of life can create real life.

Puzzle 4: Are life and mind intrinsically related?

Solution: Life and mind are expressions of essentially the same kind of process.

These puzzles are controversial and subtle. A compelling theory must not only resolve the puzzles; it should also explain why they arise in the first place. The theory of supple adaptation does all this.

3.1 Levels and Dependencies

Living phenomena fall into a complex hierarchy of levels—what I will call the *vital hierarchy*. Even broad brush strokes can distinguish at least eight levels in the vital hierarchy: ecosystems, which consist of communities, which consist of populations, which consist of organisms, which consist of organ systems (immune system, cardiovascular system), which consist of organs (heart, kidney, spleen), which consist of tissues, which consist of cells. Items at one level in the hierarchy constitute items at higher levels. For example, an individual population consists of a lineage of organisms that evolve over time. Individual organisms are born, live for a while, and then die. Taken together over time, these individuals constitute the evolving population. The vital hierarchy raises two basic kinds of questions about the nature of life. First, we may ask whether there is some inherent tendency for living systems to form hierarchies. Why are hierarchies so prevalent in the phenomena of life? The second question (really, set of questions) concerns the relationships among the kinds of life exhibited throughout the vital hierarchy. Are there different forms of life at different levels, and if so then how are these related? How are they similar and different? Which are prior and which posterior? What is the primary form of life? Haldane [35] and Mayr [47] are especially sensitive to these questions, although neither has a ready answer.

The theory of life as supple adaptation involves a two-tier picture with connected but different forms of life. The first tier consists of the primary form of life—the supplely adapting systems themselves. A supplely adapting system is an evolving population of organisms, or a whole evolving ecosystem of many populations, or, in the final analysis, a whole evolving biosphere with many interacting ecosystems. At the second tier, entities that are suitably generated and sustained by such a supplely adapting system branch off as different but connected secondary forms of life. These secondary forms of life include organisms, organs, and cells. So the idea that various forms of life are found at various levels of the biological hierarchy follows from the very structure of the theory of life.

Notice also that the very notion of a supplely adapting system implies simultaneous multiple levels of activity. Adaptive evolution involves the interaction between phenomena at a variety of levels, including at least genes and individual organisms and populations, so the process implies a system with activity at macro, meso, and micro levels. Thus, the theory of life as supple adaptation explains why life involves multiple levels of living phenomena. The agents constituting a supplely adapting population are not in every instance themselves alive. The simplest kind of supplely adapting systems seem to be something like an autocatalytic network of chemical species, such as those hypothesized to be involved in the origin of life [2, 28], and it is implausible to attribute life to the chemical species that constitute these supplely adapting systems. Nevertheless, the agents in most supplely adapting populations are alive; organisms are the paradigm case of this.

There is another more indirect and much more controversial way in which supple adaptation might explain why there is a vital hierarchy. No one doubts that organisms have parts that function to ensure the organism's survival and reproduction, and no one doubts that in some cases these parts themselves have a complex hierarchical structure (think of the immune system or the brain). The progression of evolution in our biosphere seems to show a remarkable overall increase in complexity, from simple prokaryotic one-celled life to eukaryotic cellular life forms with a nucleus and numerous other cytoplasmic structures, then to life forms composed of a multiplicity of cells, then to large-bodied vertebrate creatures with sophisticated sensory processing capacities, and ultimately to highly intelligent creatures that use language and develop sophisticated technology. This evidence is consistent with the hypothesis that open-ended evolutionary processes have an inherent, lawlike tendency to create creatures

with increasingly complicated functional organization. Just as the arrow of entropy in the second law of thermodynamics asserts that the entropy in all physical systems has a general tendency to increase with time, the hypothesis of the arrow of complexity asserts that the complex functional organization of the most complex products of open-ended evolutionary systems has a general tendency to increase with time. Make no mistake: The arrow of complexity hypothesis is far from settled. Some biologists are sympathetic but plenty are skeptical; see, for example, Gould [33, 34], Maynard Smith and Szathmáry [46], and McShea [48], as well as many of the chapters in [49] and [4]. I am not trying to resolve this controversy here. In fact, I think we have no compelling evidence either for or against the hypothesis right now [11]. My point here is that, *if* the arrow of complexity hypothesis is true, then supplely adapting systems have an inherent, internal tendency to produce entities with a complex, hierarchical structure, and so the theory of life as supple adaptation has a deep explanation of the vital hierarchy.

Whether or not the arrow of complexity hypothesis proves true, the theory of life as supple adaptation resolves the puzzle about the levels of life in a way that provides a natural explanation for why this puzzle arises in the first place.

3.2 Continuum or Dichotomy

Can things be more or less alive? Serious reflection about life quickly raises the question whether life is a Boolean (black-or-white) property, as it seems at first blush, or whether it is a continuum property, coming in many shades of gray. Common sense leans toward the Boolean view: A rabbit is alive and a rock is not, and there is little apparent sense in the idea of something falling in between these two states, being partly but not fully alive. But the common sense view is put under stress by various borderline cases such as viruses that are unable to replicate without a host and spores or frozen sperm that remain dormant and unchanging indefinitely but then “come back to life” when conditions become suitable. Furthermore, we all agree that the original life forms somehow emerged from a prebiotic chemical soup, and this suggests that there is very little, if any, principled distinction between life and nonlife. Many have concluded this implies that there is an ineliminable continuum of things being more or less alive [2, 17, 22, 26]. But is this right?

If life is viewed as supple adaptation, then the most important life/nonlife distinction involves a continuum because the activity of supple adaptability itself comes in degrees. Different systems can exhibit supple adaptation to different degrees, and a given system’s level of supple adaptation can fluctuate over time. A system’s level of supple adaptation can smoothly drop to nothing or smoothly rise from nothing. It is obvious enough that evolving systems’ level of supple adaptation can rise or fall continuously. In fact, there are methods for quantifying various aspects of an evolving system’s level of supple adaptation [8, 12], and this enables the dynamics of supple adaptation in artificial and natural systems to be compared directly [13, 14]. Thus, if we view life as supple adaptation, then being alive is a matter of degree. In addition to asking whether something is alive, we can also ask about the extent of its life; indeed, its life might vary along more than one dimension.

It is possible, of course, to define various sharp, Boolean distinctions on top of the continuum of the activity of supple adaptability. One natural distinction is whether a system’s level of supple adaptation is positive; this dichotomy marks whether or not a system is alive. But it must be admitted that any such Boolean distinction involves an unmistakable element of arbitrariness; we could just as well focus on whether or not a system’s level of supple adaptation exceeds 17 or 3.14159. Furthermore, such dichotomies would be defined in terms of a prior and more fundamental continuum of levels of supple adaptation; a system’s level of supple adaptation could be arbitrarily

close to our chosen cut-off point. Thus, the continuum is the truth underlying the dichotomies that it can be used to define.

There is a pragmatic dimension of the issue whether life at bottom is Boolean or continuous. If we quantify a system's level of supple adaptation in the way that Norman Packard and I have proposed [8, 9, 12, 14], then one needs a certain amount of data, and so a certain amount of time to gather the data, to determine (to within a certain level of statistical confidence) whether a system has a given level of supple adaptation. So, a system exhibiting very little supple adaptation will take a long time to generate enough data to distinguish it from the null hypothesis. But on that same time scale the system could exhaust some essential resource and perish. Thus, it might be impossible in practice to detect supple adaptation below a certain level on a certain time scale, and this would create a dichotomy separating detectable life from everything else. Still, this would not lessen the fact that in principle a continuum underlies this dichotomy.

The theory that life is supple adaptation, at least as I construe it, holds that life is the *activity* of supple adaptation, not merely the capacity for it. But the existence of this capacity is more basic than the extent to which it is exercised; the capacity is prior to its exercise. So we might ask whether this capacity is a Boolean property. Even if we do not know exactly what it takes for a system to have this capacity, it might seem that a system either has or lacks that capacity; it might look as if a system either can or cannot undergo open-ended evolution. But the truth seems more complicated. Supple adaptation is the process of producing *significantly new* kinds of adaptive responses to *significantly new* kinds of adaptive challenges and opportunities. Because it is dubious that there is a sharp divide between those challenges and responses that are significantly new and those that are not, the property of having the capacity for life seems to be a matter of degree.

So far we have focused on the supplely adapting system itself, as well we should if supple adaptation is the primary form of life, as I have been urging. But other things, such as individual organisms, individual organs, and individual cells, are also alive, if only secondarily, and we should ask whether their life is a matter of degree. Intuitively one would think that a flea or paramecium is no less alive than a cow or human being; likewise, my heart is no less alive than a flea's heart, and a cell in my body is no less alive than a flea's cell. The theory of life as supple adaptation supports these intuitions. The theory attributes different derivative forms of life to entities that have the right connections with a supplely adapting system and in general it is an equally determinate and dichotomous matter for humans and fleas whether such connections obtain. When something definitely does or does not satisfy the conditions of derivative life, it definitely is or is not alive. There still are the familiar borderline cases, though, such as viruses, frozen sperm, and dormant spores. But notice that these are precisely those cases in which connections with the supplely adapting system deviate from the norm. The derivative form of reproduction of viruses makes their participation in the supplely adapting system less autonomous than other organisms. Frozen sperm and dormant spores have become disconnected from the supplely adapting system but when those connections are reestablished they are brought "back to life." In this sort of way the theory of life as supple adaptation offers a natural explanation for why borderline cases *are* borderline cases.

If the theory of life as supple adaptation is right, then both continuous scales and dichotomous divisions separate the living and the nonliving. Given this complexity, it is no wonder that we are puzzled about whether there is a continuum between life and nonlife.

3.3 Matter or Form

The advent of the field of artificial life has focused attention on a set of questions about the role of matter and form in life [25, 26, 42, 53, 58]. 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, life seems more like a kind of process than a kind of substance. Furthermore, much of the practice of artificial life research seems to presuppose that life can be realized in a suitably programmed computer (see [50] for a good recent discussion of this). This raises a number of related questions: Does the essence of life concern the material out of which something is composed or the form in which that material is arranged? If the latter, is that form static or is it a process? If the latter, is that process computational? Is the property of being alive a functional property? Is it realizable in an indefinitely long list of different material substrata? Could a computer simulation of a living process ever be a realization of life, that is, literally be alive?

Supple adaptation is a kind of process, not a kind of stuff. Although this process cannot occur unless it is realized in some material, and although it cannot be realized in just *any* kind of material, the range of materials that *can* realize it seems quite open ended. After all, even economic or intellectual systems can exhibit supple adaptation. So, supple adaptation is multiply realizable. What is essential to supple adaptation is the *form* of interactions among the components, not the stuff those components are made from. Thus, what determines whether something is an instance of the process of supple adaptation is whether the right sort of functional structure is present. In other words, the process of supple adaptation has a functional definition.

Of course, the theory of life as supple adaptation leaves room for secondary life forms, which would be delineated by a more specific form of the second clause of the definition in Section 2 above. But it would seem that the clauses in such a definition will also specify structural, causal, or functional conditions and relationships, and these will also be multiply realizable. So the theory of life as supple adaptation construes life entirely as a functional property. So, on this theory, functionalism captures the truth about life. Furthermore, there is no evident reason why the functional structure specified in the theory could not be realized in a suitably structured computational medium. If so, then a computer “simulation” of life could in principle create a real, literally living entity.

A seductive misunderstanding arises at this juncture. In claiming that supple adaptation can be realized in a computational medium I am *not* claiming that the process of supple adaptation corresponds to a fixed algorithm. What blocks this is supple adaptation’s *suppleness*—its ability to respond appropriately to an open-ended and unanticipatable range of contingencies. The history of the so-called “frame problem” in artificial intelligence illustrates the problem (see, e.g., the chapters in [55]). One could try to embody a supple process in a fixed algorithm, along the lines of traditional artificial intelligence’s use of heuristics, expert systems, and the rest. But the empirical fact is that these algorithms do not supplely respond to an open-ended variety of contingencies (see, e.g., [23, 36, 37]). Their behavior is brittle, lacking the supple sensitivity to context distinctive of intelligence. For the same reason, the suppleness of supple adaptation cannot be captured in a fixed algorithm.

Nevertheless, there is no evident reason why the process of supple adaptation cannot be realized in a computational medium, provided there is a suitably supple mechanism for changing the algorithms involved. This is one of the first important lessons of the field of artificial life. Vital processes typically are supple; think of metabolism or the process of adaptation itself. Successful adaptation depends on the ability to explore an appropriate number of viable evolutionary alternatives; too many or too few can make adaptation difficult or even impossible. In other words, success requires striking a balance between the competing demands for “creativity” (trying new alternatives)

and “memory” (retaining what has proved successful). Furthermore, as the context for evolution changes, the appropriate balance between creativity and memory can shift in a way that resists precise and exceptionless formulation. Still, artificial life models can show a supple flexibility in how they balance creativity and novelty [10] because the underlying algorithmic behavior is supplely shaped and reshaped through the process of evolution. The key feature behind the supple vital dynamics produced by genetic algorithms [38] and other supple mechanisms that underlie artificial life models is their “bottom-up” architecture [42]. The supple dynamics is the emergent macro-level effect of a context-dependent competition for influence in a population of micro-level entities in the model. The micro-level models are precise and fixed algorithmic objects, of course, but their emergent macro-level supple dynamics are not. For this reason, supple adaptation can be realized as a nonalgorithmic emergent macro-level effect of an algorithmic micro-level process. Although the multiple realization of supple adaptation implies that life has a functional definition, the suppleness of this functional structure implies that the process of life is not a fixed algorithm. I have elsewhere called this special form of functionalism *emergent functionalism* [10].

This line of thought identifies three factors that fuel the puzzle about whether life depends on form or matter. One is the inherent subtlety of the relationship functionalism implies between form and matter; what is essential to supple adaptation is a certain form of process, but this form of process cannot exist without being embodied in some matter. No doubt the mechanistic, reductionistic thrust of molecular biology, fueled by the celebrated discovery of DNA’s double helix and recently re-energized by the cloning of an adult sheep, also contributes to the puzzle about whether life is form or matter. The mistaken equation of functionalism and computationalism is a third cause of the puzzle. All of this helps explain why the puzzle about whether life involves form or matter is so animated.

3.4 Life and Mind

A fourth 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 interorganism communication. Thus, various kinds of what one could, broadly speaking, call “mental” capacities are present throughout the biosphere. Furthermore, 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. The process of evolution establishes a genealogical connection between life and mind, of course, but life and mind might be much more deeply unified. For example, life and mind would be strikingly unified if Beer ([15], p. 11) 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]” (see also [1, 19, 44, 52, 59]). Because 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. Resolving how, if at all, life and mind are connected is one of the basic puzzles about life.

If mental capacities are adaptations produced by the process of evolution, then the theory of life as supple adaptation implies that mental capacities are produced by life itself. Some view the evolution of the mind as an entirely unpredictable historical accident [33, 34]; or as a plausible adaptation to environmental complexity [32]; or as an almost inevitable consequence of the evolutionary process—what Dennett calls a “forced move” [22]. All such views agree that the mind is at most just one adaptation among many. Thus, this line of thought implies that there is a genealogical connection between life and mind but it is not unique, so life and mind have no intrinsic unity.

This contrasts with Aristotle's idea that there is a deep unity between life and mind. Code and Moravcsik ([20], p. 130) explain Aristotle's position as follows:

In the case of a living thing, . . . its "psychological" activity is the exercise . . . of the various capacities and potentialities . . . assigned to its soul. . . *[F]or a living thing its natural/physical activity just is its psychological activity.* [emphasis added]

An analogously direct connection between life and mind can be grounded on the theory of life as supple adaptation, for one can view the mind as an expression of essentially the same underlying capacity for supple adaptation. It is well known that the emergent dynamical patterns among our mental states are especially difficult to describe and explain. An ineliminable open-ended list of exceptions seems to infect descriptions of all mental patterns, for which reason these patterns are sometimes called "fluid" [36] or "soft" [40]. But there are different kinds of fluidity and softness. Fodor [30] and others have emphasized the functionalist point that softness can result from malfunctions in the material and processes implementing mental phenomena. Horgan and Tienson [39, 40] have emphasized the softness that results from the indeterminate results of competition among a potentially open-ended range of conflicting desires. But what is most relevant here are specifically those exceptions to the rule that reflect our *ability to act appropriately* in the face of an open-ended range of contextual contingencies. These exceptions occur when we make *appropriate* adjustment to contingencies. Some people conclude that this supple capacity for adaptive behavior is the defining feature of intelligence or mind [1, 10, 11, 15, 44, 52, 60].

This quasi-Aristotelian view construes the mind as essentially the expression of a form of supple adaptation. Natural selection is not necessarily involved, for Lamarckian selection or some other adaptive process might do the trick. Rather, leaving aside the mechanisms of adaptation, my claim is that the process of having a mind is something quite like the process of being alive. So, the theory of life as supple adaptation is naturally allied with the theory of mind as supple adaptation. Just as the essence of life is the process that generates the phenomena of life, for the same reason the essence of mind is the process that generates intelligent behavior. If life and mind are both produced by basically the same process of supple adaptation, then the mind is not just one adaptation among many. Rather, an essential feature of the mind is involved in the explanation of all other local adaptations, so life and mind could hardly fail to coexist. They exhibit the strong continuity [31] characterized by both exhibiting the abstract pattern of supple adaptation. From the perspective of the theory of life as supple adaptation and the quasi-Aristotelian approach to the mind, it is no wonder that people think that life and mind are deeply connected.

A complete solution to the puzzle about the connection between life and mind should also explain why this connection is largely ignored today, especially among philosophers. The theory of life as supple adaptation combined with the quasi-Aristotelian approach to mind can blame this on Descartes. Contemporary philosophy of mind is a culture deeply influenced by Descartes. Descartes rejected the then orthodox scholastic Aristotelian framework in favor of the view that living substances have no essential connection with mental substances (except for the unmediated causal connection unifying each person's mind and body). Descartes focused on the intrinsic nature of isolated living and mental substances, ignoring the processes that created and sustained them, and concluding that living substances were purely material mechanisms while mental substances are essentially immaterial and spiritual consciousness. Today, even contemporary philosophy of mind that rejects Descartes's dualism of body and mind typically embraces consciousness as the essence of mind and shares Descartes's

unconcern about how living and mental substances are produced. One testament to Descartes's persistent influence is the present difficulty of initially motivating the puzzle about how life and mind are connected.

4 Open Questions and Conclusions

I offer no final and complete theory of life and no final and complete solution to the four puzzles about life, but I do defend the general form of the theory of life as supple adaptation. My defense consists of showing the theory's promising and illuminating solutions to four puzzles about life.

This defense highlights three open questions. The first is to determine what, in the end, is the best explanation of the salient phenomena and puzzles concerning life. Even if supple adaptation provides good explanations of these matters, this leaves room for other theories to provide better explanations. Our final understanding of what life really is will turn on which theory in the end provides the best explanations.

When we try to settle exactly how well supple adaptation explains these matters, two more questions arise. First, one thing, this theory is no clearer than the notion of supple adaptation itself, and there is still much to learn about supple adaptation. For example, not a single artificial evolutionary model has unambiguously shown the sort of continually growing supple adaptation evident in the biosphere [13, 14], not even those models with unpredictably changing selection criteria and an infinite space of genetic possibilities, such as John Holland's Echo [38], Kristian Lindgren's evolving strategies for infinite prisoner's dilemmas [43], and Tom Ray's Tierra [56]. The problem seems to be that no existing model creates a continually unfolding accessible space of new kinds of adaptive innovations. Synthesizing even one demonstrable instance of continually growing supple adaptation would profoundly advance our understanding of this process. The task of producing and certifying such a model falls squarely to the field of artificial life. If life is supple adaptation, finding such a model is one of the field's most pressing current challenges.

Finally, even if our understanding of supple adaptation were complete, we still would need to settle how best to use it to define life. For example, we need to determine the different ways in which different forms of life can be explained by a supplely adapting system. These details will replace the place-holding expression "explained in the right way" in the definition given above. We also need to decide what weight to place on different mechanisms for producing supple adaptation. Natural selection is one such mechanism, but there is an open-ended variety of others (Lamarckian selection, etc.). Once we have delineated all those mechanisms, we will be faced with a choice: Is the primary form of life supplely adapting systems produced by any mechanism? Only by natural selection? The way to settle this question, in the end, is to determine which choice provides the most illuminating understanding of the phenomena and puzzles surrounding life.

I intend the present discussion to establish two main conclusions about the theory of life. The first is methodological: The search for a theory of life is more productive if it focuses on the best explanation of life, including deep and persistent puzzles about life. This methodology frees us from many traditional worries caused by our current preconceptions about life, including worrying about necessary and sufficient conditions for all and only living organisms. My second conclusion is substantive: The theory of life as supple adaptation deserves our serious consideration. To be sure, the theory generates tension with our present preconceptions of life, but this is no strike against the theory. Rival theories are credible contenders only if they explain living phenomena and resolve the four puzzles at least as well as the theory of supple adaptation.

Acknowledgments

For helpful discussion of these topics or comments on previous drafts, thanks to Hugo Bedau, Maggie Boden, Andy Clark, Peter Godfrey-Smith, David McFarland, Dan McShea, David Reeve, Francisco Varela, Bill Wimsatt, Marty Zwick, the spring 1997 audience at the conference on the Philosophy of Artificial Life, *PAL97*, Christ Church, Oxford University, the reviewers and summer 1997 audience at the Fourth European Conference on Artificial Life, *ECAL97*, Brighton, England, the winter 1998 audience at my Systems Science colloquium at Portland State University, and the anonymous reviewers for this journal. Thanks to the Santa Fe Institute for support and hospitality during the visits that started and sustained this work. Special thanks to Norman Packard for years of collaboration that have profoundly influenced my thinking about this topic.

References

1. Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
2. 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.
3. Banzhaf, W., & Eeckman, F. (Eds.). (1995). *Evolution and biocomputation—Computational models of evolution*. Berlin: Springer-Verlag.
4. Barlow, C. (1995). *Evolution extended: Biological debates about the meanings of life*. Cambridge, MA: MIT Press.
5. Bedau, M. A. (1990). Against mentalism in teleology. *American Philosophical Quarterly*, 27, 61–70.
6. Bedau, M. A. (1991). Can biological teleology be naturalized? *The Journal of Philosophy*, 88, 647–655.
7. Bedau, M. A. (1992). Where's the good in teleology? *Philosophy and Phenomenological Research*, 52, 781–805.
8. Bedau, M. A. (1995). Three illustrations of artificial life's working hypothesis. In W. Banzhaf & F. Eeckman (Eds.), *Evolution and biocomputation—Computational models of evolution* (pp. 53–68). Berlin: Springer-Verlag.
9. Bedau, M. A. (1996). The Nature of life. In M. Boden (Ed.), *The philosophy of artificial life* (pp. 332–357). New York: Oxford University Press.
10. Bedau, M. A. (1997). Emergent models of supple dynamics in life and mind. *Brain and Cognition*, 34, 5–27.
11. Bedau, M. A. (1997). Philosophical content and method in artificial life. In T. W. Bynam & J. H. Moor (Eds.), *The digital phoenix: How computers are changing philosophy* (pp. 135–152). New York: Blackwell.
12. 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. 431–461). Redwood City, CA: Addison-Wesley.
13. Bedau, M. A., Snyder, E., Brown, C. T., & Packard, N. H. (1997). A comparison of evolutionary activity in artificial systems and in the biosphere. In P. Husbands & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life ECAL 97* (pp. 125–134). Cambridge, MA: MIT Press/Bradford Books.
14. Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Artificial life VI* (pp. 228–237). Cambridge, MA: MIT Press/Bradford Books.
15. Beer, R. D. (1990). *Intelligence as adaptive behavior: An experiment in computational neuroethology*. Boston: Academic Press.

16. Boden, M. (1998). *Is metabolism necessary?* (Cognitive Science Research Paper 482). University of Sussex.
17. Cairns-Smith, A. G. (1985). *Seven clues to the origin of life*. Cambridge, U.K.: Cambridge University Press.
18. Chyba, C. F., & McDonald, G. D. (1995). The origin of life in the solar system: Current issues. *Annual Review of Earth and Planetary Sciences*, 23, 215–249.
19. Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
20. 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. 129–145). Oxford: Clarendon Press.
21. Dawkins, R., & Krebs, J. R. (1978). Arms races between and within species. *Proceedings of the Royal Society of London, B* 205, 489–511.
22. Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon and Schuster.
23. Dreyfus, H. (1979). *What computers cannot do* (2nd ed.). New York: Harper and Row.
24. Eigen, M. (1992). *Steps toward life*. Oxford: Oxford University Press.
25. 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. 466–474). Cambridge, MA: MIT Press/Bradford Books.
26. Emmeche, C. (1994). *The garden in the machine: The emerging science of artificial life*. Princeton, NJ: Princeton University Press.
27. 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. 815–840). Redwood City, CA: Addison-Wesley.
28. Farmer, J. D., Kauffman, S. A., & Packard, N. H. (1986). Autocatalytic replication of polymers. In J. D. Farmer, A. Lapedes, N. Packard, & B. Wendroff (Eds.), *Evolution, games, and learning: Models for adaptation for machines and nature* (pp. 50–67). Amsterdam: North Holland.
29. Farmer, J. D., Lapedes, A., Packard, N., & Wendroff, B. (Eds.). (1986). *Evolution, games, and learning: Models for adaptation for machines and nature*. Amsterdam: North Holland.
30. Fodor, J. A. (1981). *Representations*. Cambridge, MA: MIT Press/Bradford Books.
31. Godfrey-Smith, P. (1994). Spencer and Dewey on life and mind. In R. Brooks & P. Maes (Eds.), *Artificial life IV* (pp. 80–89). Cambridge, MA: MIT Press/Bradford Books.
32. Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. Cambridge, U.K.: Cambridge University Press.
33. Gould, S. J. (1989). *Wonderful life: The Burgess shale and the nature of history*. New York: Norton.
34. Gould, S. J. (1996). *Full house: The spread of excellence from Plato to Darwin*. New York: Harmony Books.
35. Haldane, J. B. S. (1937). *Adventures of a biologist*. New York: Macmillan.
36. Hofstadter, D. R. (1985). Waking up from the Boolean dream, or, subcognition as computation. In *Metamagical themas: Questing for the essence of mind and pattern* (pp. 631–665). New York: Basic Books.
37. Holland, J. H. (1986). Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning II* (pp. 593–623). Los Altos, CA: Morgan Kaufmann.

38. Holland, J. H. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: MIT Press/Bradford Books.
39. Horgan, T., & Tienson, J. (1989). Representation without rules. *Philosophical Topics*, 17, 147–174.
40. Horgan, T., & Tienson, J. (1990). Soft laws. *Midwest Studies in Philosophy*, 15, 256–279.
41. Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
42. Langton, C. (1989). Artificial life. In C. Langton (Ed.), *Artificial life* (pp. 1–47). Redwood City, CA: Addison-Wesley.
43. Lindgren, K. (1992). Evolutionary phenomena in simple dynamics. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 295–312). Redwood City, CA: Addison-Wesley.
44. Maturana, H. R., and Varela, F. J. (1987). *The tree of knowledge: The biological roots of human understanding* (Rev. ed.). (1992) Boston: Shambhala.
45. Maynard Smith, J. (1975). *The theory of evolution* (3rd ed.). New York: Penguin.
46. Maynard Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. New York: Freeman.
47. Mayr, E. (1982). *The growth of biological thought*. Cambridge, MA: Harvard University Press.
48. McShea, D. W. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution*, 50, 477–492.
49. Nitecki, M. H. (1988). *Evolutionary progress*. Chicago: University of Chicago Press.
50. Olson, E. T. (1997). The ontological basis of strong artificial life. *Artificial Life*, 3, 29–39.
51. Packard, N. (1989). Intrinsic adaptation in a simple model for evolution. In C. Langton (Ed.), *Artificial life* (pp. 141–155). Redwood City, CA: Addison-Wesley.
52. Parisi, D., Nolfi, N., & Cecconi, F. (1992). Learning, behavior, and evolution. In F. Varela & P. Bourguin (Eds.), *Towards a practice of autonomous systems* (pp. 207–216). Cambridge, MA: MIT Press/Bradford Books.
53. Pattee, H. H. (1989). Simulations, realization, and theories of life. In C. Langton (Ed.), *Artificial life* (pp. 63–78). Redwood City, CA: Addison-Wesley.
54. Pirie, N. W. (1972). On recognizing life. In D. L. Rohlfsing & A. I. Oparin (Eds.), *Molecular evolution: Prebiological and biological* (pp. 67–76). New York: Plenum Press.
55. Pylyshyn, Z. W. (1987). *The robots dilemma: The frame problem in artificial intelligence*. Norwood, NJ: Ablex.
56. 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.
57. Sagan, C. (1998). Life. In *The New Encyclopaedia Britannica* (15th ed. Macropaedia, pp. 964–981). Chicago: Encyclopaedia Britannica, Inc.
58. 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. 749–765). Redwood City, CA: Addison-Wesley.
59. Taylor, C. (1992). Fleshing out. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 25–38). Redwood City, CA: Addison-Wesley.
60. Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press/Bradford Books.