

# Ghosts in the Evolutionary Machinery

*Steve Talbott*

A science straining toward what it imagines to be strictly material concepts will end up with abstract and general ones. That is, as the philologist and semantic historian Owen Barfield warned eighty years ago, our pursuit of materialism will paradoxically estrange us from concrete, material reality. The reason for this is that the world we know is a world of specific character, of particular, insistent presences, of expressive qualities—a world of smiling faces, fluttering leaves, stretching cats, billowing clouds. In turning away from these presences, from these qualities—in seeking the denatured, inert, non-experienceable stuff of the scientist's abstract imagination—we turn away from the one reality we are given. It is only natural, then, that direct and careful observation of the world's vivid, many-faceted character should yield more and more to one-dimensional measurement. As Barfield put it in 1963:

It must be admitted that the matter dealt with by the established sciences is coming to be composed less and less of actual observations, more and more of such things as pointer-readings on dials, the same pointer-readings arranged by electronic computers, inferences from inferences, higher mathematical formulae and other recondite abstractions. Yet modern science began with a turning away *from* abstract cerebration *to* objective observation!

It is hardly disputable that science has in fact listed heavily toward the imperceptible, immaterial, and abstract. The essence of modern science lies in the mathematization of reality, an exercise vastly more common than its necessary counterpart, which is the effort to characterize what sort of reality the mathematics refers to.

However, bloodless abstraction alone is impossible to sustain; we can't help wanting our equations to be statements about real things that we experience and understand. This is where models and machines come in handy: they enable scientists to re-embody their abstractions in a kind

---

*Steve Talbott*, a New Atlantis contributing editor, is a senior researcher at the Nature Institute and the editor of NetFuture, an online newsletter about technology and human responsibility ([netfuture.org](http://netfuture.org)). His new book is *Devices of the Soul: Battling for Our Selves in an Age of Machines* (O'Reilly). He can be reached at [stevet@oreilly.com](mailto:stevet@oreilly.com).

---

of alternate, well-understood, mechanistic reality, without being unduly bothered by the expressive presence of the original phenomena. The world is re-cast in the image of our own preferred habits of thought: orderly, understandable, manipulable. And now computer simulations are raising this modeling power almost to the level of pure conjuration, giving us programmed appearances that seem wonderfully real regardless of how thin the threads connecting them to physical reality. Some scientists seem to view the computer as a kind of Creation Machine; using it, they can watch the logic of their own thoughts, however arbitrary, congeal into convincing “material” phenomena.

So it is that, in the field of evolutionary studies, we now find strange séances in high-tech laboratories where biologists, bent over their computerized abstractions, struggle to bring them alive as mutating, replicating, evolving organisms.

### **Bringing Numbers Alive**

The scientists engaged in this occult-sounding project begin with a simple array of numbers—a data structure—where each number represents a low-level computer instruction. Perhaps, to begin with, the combined instructions don’t do anything sensible other than to direct the computer to replicate the entire data structure. However, you can arrange the supporting software so that every so often a “mutation” occurs, resulting, for example, in an altered, added, or deleted instruction. This might destroy the possibility for replication, but if you start with a considerable number of these data structures, an occasional mutation might yield a non-destructive new capability. For example, a mutated element in a data array might happen to be the computer’s instruction for performing a primitive logical operation, or, in combination with other, already present instructions, it might manage a more complex logical operation. Conceivably, a large number of such mutations could eventually produce a data structure capable of directing the computer to calculate the sum of any two supplied numbers, a feat that (depending on the computer) might require thirty-five or more instructions in a given sequence.

If you are like many scientists working in computational biology, you will be inclined to regard such data structures, along with the associated bits of programming, as “digital organisms”—organisms that, under the right software conditions, can be said to “reproduce,” “evolve,” and become “fitter.” Each data structure, with its array of computer instructions, can be thought of as a “genome.” When mutations produce a new logical capability,

the software “rewards” the organism with more opportunities to execute its instructions and to reproduce, whereas mutations that compromise its logical capabilities reduce those opportunities. The organism thus has a “metabolism”: it gains “energy” in the form of computer processing time whenever it is successful at logically processing numbers (“food”) received from its software environment and returning the numerical products of this activity back to the environment. This means that the more logically capable organisms tend to proliferate and are encouraged along their evolutionary path toward ever more sophisticated arithmetic prowess.

Lest you suspect me of being a little facetious in the foregoing, let me assure you that the biological terminology I’ve employed here is widely accepted and can be found, for example, in a 2003 article in the prestigious journal *Nature* detailing computer trials aimed at discovering how digital organisms can evolve the ability to add two numbers.\* “Digital organisms,” the authors tell us, “exist in a computational environment where copying [of data structures] is imperfect such that they mutate randomly and evolve spontaneously.” These organisms “compete for energy,” which they obtain by “performing logical functions.” Further, “reproduction is asexual,” and “each mutation alters the genome and may change an organism’s... replication efficiency, computational metabolism, and robustness.”

This work, which goes forward at places like the Digital Evolution Laboratory at Michigan State University and the Digital Life Laboratory at California Institute of Technology (Caltech), has many researchers galvanized. Robert Pennock, a philosopher at Michigan State and a co-author of the *Nature* article, finds it increasingly difficult to doubt that digital organisms are alive. “More and more of the features that biologists have said were necessary for life we can check off,” he says. “Does this, does that, does this. Metabolism? Maybe not quite yet, but getting pretty close.”

Acknowledging the discomfiture many people are likely to feel at this point, science writer Carl Zimmer concedes in a 2005 *Discover* magazine story that “it may seem strange to talk about a chunk of computer code in the same way you talk about a cherry tree or a dolphin.” But, he goes on, “the more biologists think about life, the more compelling the equation becomes”:

Computer programs and DNA are both sets of instructions. Computer programs tell a computer how to process information, while DNA instructs a cell how to assemble proteins.

---

\* Richard E. Lenski, Charles Ofria, Robert T. Pennock, and Christoph Adami, “The Evolutionary Origin of Complex Features,” *Nature* 423 (May 8, 2003): 139-144.

---

It's a breathtaking analogy, which Zimmer elaborates this way:

A cherry tree absorbs raw materials and turns them into useful things. In goes carbon dioxide, water, and nutrients. Out comes wood, cherries, and toxins to ward off insects. A computer program works the same. Consider a program that adds two numbers. The numbers go in like carbon dioxide and water, and the sum comes out like a cherry tree.

In this dizzying analogical world, the only thing required for drawing a profound comparison, it seems, is the common use of the two words “in” and “out”—never mind that there is no evident connection otherwise, or that even these two words are used in radically different senses in the two cases. If this sort of analogy is in fact what drives the work on digital organisms, one has to wonder: Does the actual material reality of the cherry tree—or bacterium, or elephant—figure at all in the researchers' theorizing about evolution? As we will see, the aggressive claim is being made, not only that it does not, but that it should not.

### Looking for a Body

When scientists tell us that an organism is reproducing, competing, mutating, metabolizing, and evolving, you might think they could point us to the creature doing these things—or at least that *something* physical exists whose evolution they are investigating. One cannot meaningfully say “compete” and “mutate” if these verbs are wandering around lost, without a subject.

Yet the digital organism is not a physical entity. We could write the relevant data structures on a blackboard, carrying out the computations of the software ourselves and showing how these structures change. But no one would then claim that we had discovered something important about the evolutionary behavior of certain traces of chalk or segments of blackboard. It's the data structures themselves, as immaterial, ideal constructions, that the researchers have in mind.

Similarly, nothing in the scientific literature suggests that anyone is thinking of bits of physical real estate in the computer when they imagine digital organisms performing feats of evolutionary transformation. It's a good thing, for given the design and operation of modern computers, an executing program with its data structures typically flits through memory banks and registers in a manner nearly impossible to trace. Even if we did go to the trouble of tracing the activity at a physical level, we would find no physically contiguous, let alone coherently organized, physical stuff we could

begin to think of as an evolving organism. And in any case, those discussing such organisms never concern themselves with physical manifestations. They are not talking about material forms, nor even looking for them.

By contrast, biologists who deal with actual organisms find themselves facing questions that can reasonably be termed evolutionary questions. The one-chambered heart in this lobster can be compared to the two-chambered heart in that fish, the three-chambered heart of a lizard, and the four-chambered heart of a human being. Setting a series of such organs side by side, we can, with an aesthetic eye, judge the relations between them as suggesting more or less strongly the possibility of a historical transformation of physical forms. We have physical phenomena to investigate.

The fact that digital creatures are bereft of any tissues or organs—or even silicon, plastic, or copper—about which to raise questions of evolutionary development does not seem to worry anyone, or even to have been much noticed. For the authors of the article in *Nature*, the digital organism's phenotype—its observable physical qualities—has remarkably become nothing but a set of relations between mathematical formulae such as “replication efficiency” and “computational metabolism.” There is no material *thing* actually reproducing, or whose metabolism is hard at work.

### Is There Life Beyond Life?

The digital organism enthusiasts certainly think they are talking about *something* of real substance. The most direct conclusion to draw is that they have simply reified in their minds a set of computations and data structures. Having calculated certain ideal relationships expressed in program logic, they allow these relationships to condense, specter-like, into dim, vaguely imagined physical objects—objects that are, as a result, gratifyingly well-behaved in a logical sense. They then herald these ghostly compactions of logic as powerful demonstrations of how actual physical organisms evolve in obedience to the now perfectly displayed logic of their evolution.

There is something stunningly backward and tautological in all this. What these researchers are really doing is exploring certain possibilities of mathematical and algorithmic logic. It's a legitimate thing to do. Throughout the history of science the elaboration of mathematical formulae has often led, at least in the physical sciences, to subsequent discovery of application for these formulae. But this doesn't alter an obvious truth: *the discovery always needs to be made*, and it can be made only through observation of the world. Many of those who speak about artificial life seem strikingly casual about the role of observation.

The authors of the *Nature* article hint at this casual attitude when they say that the aim of their software project is “to shed light on principles relevant to any evolving system.” The key word is “any.” Writing in a different publication, one of the *Nature* co-authors makes the point more explicit: “Digital organisms provide a unique opportunity with which to study evolutionary biology in a form of life that shares no ancestry with carbon-based life forms, and hence to distinguish general principles of evolution from historical accidents that are particular to biochemical life.” In other words, understanding the “accidental” forms of life we actually can observe is not a central aim of the researchers.

Driving home the same point even more vigorously, the late Danish complexity theorist Per Bak opined that a general theory of life “cannot have any specific reference to actual species. The model may, perhaps, not even refer to basic chemical processes, or to the DNA molecules that are integral parts of any life form that we know.” After all, he wondered in his 1996 book *How Nature Works*, what might life forms on Mars be like?

We must learn to free ourselves from seeing things the way they are! A radical scientific view indeed! If, following traditional scientific methods, we concentrate on an accurate description of the details, we lose perspective. A theory of life is likely to be a theory of process, not a detailed account of utterly accidental details of that process, such as the emergence of humans.

When the only life forms we have known become dismissible details—when we are no longer bound by the onerous necessity to see things “the way they are”—then of course we are free to invent new forms of life in any way we wish, and to proclaim ourselves the discoverers of their wonderfully effective abstract laws—laws we no longer need to ground in biological phenomena. This experience in the projection of our own thoughts upon the world goes especially smoothly when, with the aid of the computer, we can simply allow our mathematical formulae to crystallize into “real” organisms of a new, post-biological sort. Such an exercise in creation by human design proves much easier than the empirically tedious task of demonstrating the spontaneous generation of life from rubbish heaps.

### Going Deep

The literature on artificial life and complexity contains countless references to the search for the simple and general. It’s a search that is supposed to yield deep explanatory principles precisely because these

---



principles are not tied to particular phenomena but instead generalize across broad fields of phenomena.

Christoph Adami, head of the Digital Life Laboratory at Caltech, acknowledges that “the worlds we’re dealing with here are extraordinarily simple compared with the real world.” The almost unfathomably complex processes by which DNA is transcribed into RNA and then translated into cell proteins, and by which those proteins contribute to particular traits of the organism, are intentionally left out of the picture. As he told a journalist writing for *PLoS Biology* in 2003, “We can’t see transcription and translation because we don’t have transcription and translation—we go right from sequence to function.” That is, the sequence of computer instructions in the digital organism’s “genome” is directly executed by the computer, and certain abstract features of this execution are considered to be the bodily functioning of the organism. “But,” Adami goes on,

the principles of evolutionary theory make such restrictions unimportant. Many of the [theory’s] predictions don’t depend on these little details of molecular biology. The principles are very, very general, and very simple, and in the end they are mostly responsible for the overall dynamics that you see in these simple systems.

In justifying their use of a simple and therefore tractable computational model, the authors of the *Nature* article refer to philosopher Daniel Dennett’s remark in his 1995 book *Darwin’s Dangerous Idea* that “evolution will occur whenever and wherever three conditions are met: replication, variation (mutation), and differential fitness (competition).” These conditions, Dennett claims, manifest in an algorithmic (recipe-like) procedure, but one that isn’t a procedure *of* anything in particular. That is, the procedure is “substrate-neutral”: its power “is due to its *logical* structure, not the causal powers of the materials used in the instantiation, just so long as those causal powers permit the prescribed steps to be followed exactly” (emphasis in original).

In other words, evolution is like a computer program that can run on many different kinds of computers. Dennett’s illustration of an algorithm is illuminating. He discusses a computer program called the “Game of Life,” originally devised by mathematician John Horton Conway in 1970. The program divides your computer screen into a fine-meshed rectangular grid, wherein each tiny cell can be either bright or dark, on or off, “alive” or “dead.” The idea is to start with an initial configuration of alive cells and then, with each tick of the clock, see how the configuration changes as these simple rules are applied:

- If exactly two of a cell's eight immediate neighbors are alive at the clock tick ending one interval, the cell will remain in its current state (alive or dead) during the next interval.
- If exactly three of a cell's immediate neighbors are alive, the cell will be alive during the next interval regardless of its current state.
- And in all other cases—that is, if less than two or more than three of the neighbors are alive—the cell will be dead during the next interval.

You can, then (as the usual advice goes), think of a cell as dying from loneliness if too few of its neighbors are alive, and dying from overcrowding if too many of them are alive.

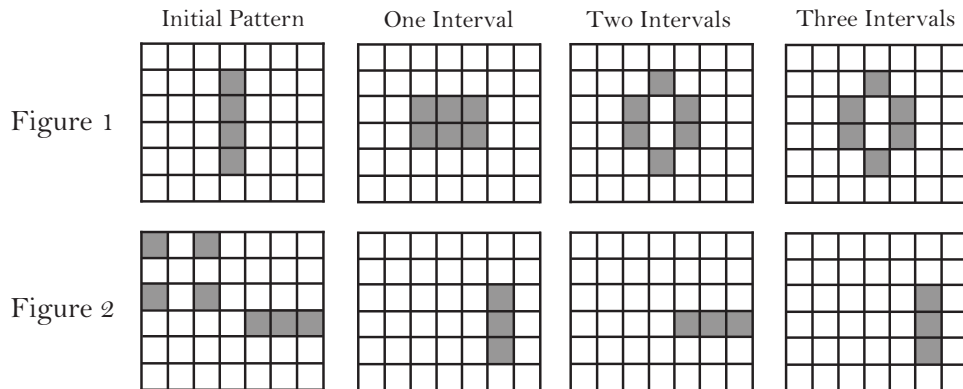
What intrigues many researchers is the fact that, given well-designed initial configurations, fascinating patterns are produced as the program unfolds. Some of these patterns remain stable or even reproduce themselves endlessly. Investigations of such behavior have led to the new discipline known as “artificial life.”

Referring to the Game of Life and the three-part rule governing its performance, Dennett remarks that “the entire physics of the Life world is captured in that single, unexceptioned law.” As a result, in the Life world “our powers of prediction are perfect: there is no [statistical] noise, no uncertainty, no probability less than one.” The Life world “perfectly instantiates the determinism made famous by [French mathematician Pierre-Simon] Laplace: if we are given the state description of this world at an instant, we observers can perfectly predict the future instants by the simple application of our one law of physics.”

These are startlingly errant statements from one of the most influential philosophers of our day. The three-part rule, after all, is not a law of physics, and not even, in any relevant sense, a law of the computer on which it is executing. It defines an algorithm, and its deterministic, Laplacian perfection holds true only so long as we remain within the perfectly abstract realm of the algorithm's crystal-line logical structure. Try to embody this structure in any particular stuff of the world, and its perfection suddenly vanishes. For example, if you execute it in a running computer, you can be absolutely sure that the algorithm will fail at some point, if not because of spilled coffee or a power failure or an operating system glitch, then because of normal wear and tear on the computer over time. Contrary to Dennett's

---





*Examples of the “Game of Life.”* In Figure 1, the column of live cells changes shape twice; after the second interval, it has reached a stable arrangement. In Figure 2, the four live cells in the upper left all die after one interval; the other block of cells will switch back and forth from horizontal to vertical with each interval.

claim, you will find in every physical implementation of this algorithm that there is noise, no certainty, and no probability equal to one.

### An Irrepressible World

Digital evolution where nothing, no *thing*, is evolving; a physics that is not the physics of anything; an understanding that tries to “free itself from seeing things the way they are”—there is in all this a strange recoiling from the actual manifestation of the world with its insistent character, as if the investigators do not want to meet whatever is there, as if they prefer the clarity and the narcissistic pleasures of their own cleanly articulated, one-dimensional thoughts—so easily reflected by the computer—to the vocal, full-bodied self-presentation of cloud, ocean, stone, and sparrow.

Yet we are always coming up against this presentation in one way or another. Even if our habit is immediately to imprison it in reified formulae, the world itself forever finds new ways to surprise us. We have long possessed rigorous mathematical statements of the fundamental physical laws, and yet the new environments we explore, from the quasars and pulsars of deep space to the ethereal domain of “elves,” “red sprites,” and “blue jets” above the cloud tops to the thermal vents on ocean floors, repeatedly present us with *kinds of phenomena*, eloquent contexts, never before imagined, let alone predicted. Every year brings a string of such discoveries, as if the law-abiding cosmos were determined to express itself in a qualitative variety as endless and irrepressible as that of life on earth.

Every phenomenon manifests its law-abiding nature in its own way,

according to a characteristic artfulness of its own being that expresses much more than a mere submission to general law. It's not as if physical laws represent an iron necessity external to the things themselves. The laws are what emerge as a common grammar of the language that every phenomenon speaks in its own unique voice. And so, if you walk through the New England landscape on a bright spring day, for instance, and if you are observantly open to the character of your surroundings, you can't help experiencing how, for example, every tree's vivid self-presentation virtually shouts at you, "Recognize me! I am my own kind of being, completely, utterly, exuberantly different from all those other species!" In this striking distinctiveness of expression, recognizable to a child, in the indisputably aesthetic character of everything we encounter in nature, we come up against the essence of the problem of the presence of real stuff—the problem our science has done its best to ignore these past several hundred years.

If we took our "deep" generalizations with the seriousness they deserve, they would all point us to this problem of real, expressive content, because they all presuppose such content. But the habit of inattention to the world—a habit that comes to such striking expression in the work on digital organisms—runs deep. The unspoken assumption underlying this habit seems to be that the given substance of the world can simply be taken for granted, since its entire being must consist of nothing but its behavior according to law. What sort of thing might be doing the behaving—what sort of concrete manifestation it is that can respect and fulfill law—never comes up for question, just as the body of the digital organism never comes up for question. If the law is everything, then who needs actual beings? They can only threaten to do something unexpected as they go about their lawful business; better to dispense with them.

So long as we ignore the problem of real content, we will still bring the material world into our science—we cannot help doing so—but we will do it in a largely unconscious and highly distorted manner, sacrificing the world as much as we can to abstract mental models and above all to that miracle of condensed logic, the computer.

### **Machines, Design, and the World**

There is one distinction I have thus far glossed over. While the mathematically rigorous laws of physics can contribute in a real and profound way to our understanding of the physical world, the logical syntax of a computer does not in the same way contribute to our understanding of the physical machine. The law of gravity is a native law of copper, glass,

and silicon in a way that the computer's program logic is not. Rather, the program logic relates primarily to the way we have articulated the physical parts one with another so as to create a humanly useful mechanism. The computer's logic is a function of design activity external to the materials themselves—an activity imposed from without—whereas the law of gravity arises from what matter and space *are*. Remove the program from the computer, or disassemble the physical machine, and there is no loss to the nature of copper, glass, and silicon; but you cannot remove gravitation without losing the materials themselves—their very substance is in part a “gravitational way of being.”

In other words, we cannot think of the logic or mathematics of gravity in relation to the physical world the way we think of program syntax in relation to a computer. The importance of this can hardly be overestimated at a time when the lawfulness of the universe is increasingly conceived as a kind of software governing a world-machine.

Here, incidentally, we can recognize the common ground shared by the advocates of Intelligent Design and their conventional opponents: both view the universe as a grand machine. This groundless assumption is the explicit foundation equally of the case for intelligent design (“the machine requires a Designer”) and the case for a materialistic, mindless universe (“a machine is *merely* a machine—and we learned long ago simply to ignore the question of a Designer or First Cause, or to conceal it behind the obscurity of the Big Bang”). The theists correctly understand that a machine requires an intelligent designer, whether we acknowledge this fact as such or attempt to smuggle the designer into our thinking by obscure bits and pieces. The materialists, in turn, see well enough that a machine-world is no suitable habitation for a human soul and spirit.

One way out of the ill-tempered and lightless debate between the two sides is to recognize that the intelligence we see in the world is not imposed from the outside upon pre-existing material, in the way we impose our design upon a machine. The intelligence in nature works always from within. In the world's phenomena we see intelligence embodying itself in that visible, significant, aesthetically compelling speech we can't help recognizing everywhere around us. The one thing we can be certain of is that whatever—or whoever—speaks through these phenomena is not doing so in the way we speak through the design of our machines. It is the height of hubris to think that we have become creators in that fundamental sense. Our design of machines does not bring material reality itself into existence as the embodiment of our own expressive powers. It is not both the lawfulness and the substance of things.

---

## Failed Science

Against this backdrop we can now usefully summarize and amplify the web of confusions ensnaring the work on digital organisms.

To begin with—and insofar as the computer programs are thought to provide actual instances of life and not merely simulations of it—the problem couldn't be more striking: we are given presumed physical organisms without identifiable or even coherently conceivable bodies. It is strange to find scientists with a declared commitment to strictly material concepts hailing “organisms” composed of disembodied quantities as if they were discoveries at the frontier of science. The fact that the missing bodies haven't even provoked discussion in the literature indicates the unconscious force of the tendency to replace the world with abstraction. No bodily subject is required because the research isn't designed to elucidate the actual character of real things.

A second confusion arises from the belief, shared by most if not all the digital organism researchers, that physical organisms are machines whose behavior, development, and evolution are determined by something like software—for example, by rules programmed into DNA and by a larger evolutionary algorithm programmed into the world as a whole. The problem here is that software is never intrinsic to the machine on which it executes. If we want to know the native lawfulness of the living organism, we cannot conceive it in the manner of a computer program. To devise software imagined as the operating system of a machine-organism is to turn away from any revelation of the inherent potentials of the organism as a creature of nature.

However, basic physical laws *do* give us a kind of mathematical regularity that is intrinsic to the physical world, and there is presumably *some* sort of distinctive biological lawfulness at work in living organisms. Digital software, it might be claimed, can simulate or model the laws of evolving organisms in the way computers now model, for example, the weather. And just as the models enable forecasters to predict the weather (approximately and in the short term), so, too, the digital organism software might enable at least approximate prediction of patterns of evolution.

But—leaving aside any discussion of biological lawfulness and of the limitations of models in general (a discussion that badly needs to be advanced)—no one will doubt that the value of a model depends on its fidelity to the facts of the case. The digital organism researchers seem remarkably unconcerned about the fact that the numerical parameters of their software bear little relation to any measured parameters derived

from living organisms. We have already seen a primary reason for their unconcern: they don't feel bound by what is known about existing, carbon-based life forms. By remaining at a general level, they hope to discover deeper, more universal truths. As the authors of the *Nature* article put it, "By using this tractable system, we aim to shed light on principles relevant to any evolving system."

"Any" evolving system? This is hard to fathom. There is, after all, a tension between generality and depth. If you want the most general truth of all, then just say for every organism and thing in the universe, "X exists"—or, in mathematical terms, "X = 1." But what knowledge of anything at all does this buy us? A universal law, precisely because it tells us what is true of everything, cannot tell us much about anything in particular. It cannot distinguish one thing from another. To say that you've got some "deep" principles applying not only to biological life as we know it but also to all sorts of other forms of life not yet even conceived—and in saying this to willfully ignore whatever we do know about the life forms we study on earth—is to give up saying anything at all revelatory about this life.

What is worse, we cannot even arrive at our generalizations about earthly life and other possible life without first thoroughly understanding the things from which we are trying to generalize. Far from freely contradicting all the particulars underlying it, a generalization derives its only validity from those particulars. By disavowing the importance of the "accident" of life as we know it, we forgo all possibility of "deep" generalization.

### **A Proper Eye for Evolution**

Finally, even if we observed precisely definable, universal laws in living organisms (where the generally recognized truth is that we have no biological laws of the sort we have in physics), and even if the software perfectly modeled these laws, we would still have no answers to the questions the digital organism researchers are asking.

The *Nature* article was intended to explore whether "complex functions can originate by random mutation and natural selection." By playing with philosopher Dennett's three conditions—replication, mutation, and competition—abstractly and outside any actual biological context, the digital organism researchers claim to have answered their question positively: they say they have shown that complex features—classic examples are the wing of a bird or the eye of an insect or mammal—could have appeared in earth history as a result of known and understood evolutionary mechanisms.

The conclusion is supposed to follow from the fact that their software data structures “evolved” the complex ability to add two numbers.

But this is to leave out of consideration almost everything. It is to ignore the entire character of the material and contextual reality expressing itself in evolution. To take just one of Dennett’s conditions: do the researchers really believe that what we have already discovered and may in the future discover about *mutation*—for example, about the newly emerging field of adaptive mutations, or about the whole domain of epigenetic variation, or about the validity or invalidity of the problematic notion of “chance” mutation—has no bearing on our judgment about the possibilities for evolution of the mammalian eye? And if these things do have a bearing, then what is the use of an investigation that leaves not only them, but all other concrete considerations of biological reality out of the picture?

Or, again, if we can scarcely begin to define what makes for “fitness” in real biological contexts—and who would be bold enough to suggest that we yet have much of a clue about it?—how can we possibly program fitness into a computer algorithm that will then tell us anything serious about evolutionary processes?

It is of decisive importance for the scientist to recognize that fixed laws, computer algorithms, and logic cannot by themselves give us a scientifically understood content of the world. A leaf may obey the same gravitational law and other physical laws as a stone, but a leaf is not capable of going through anything like the kind of “life history” a stone may go through. If we know only the common lawfulness they display but do not know the different things themselves—well, then, we don’t know the things themselves. We don’t know what they are capable of. Laws do not magically “add up” to real substance; they cannot give us the totality or the character of whatever has the power of asserting its existence in its own, lawfully respectful way.

And what is true of a leaf and stone is true of everything else—for example, DNA. We can know what sort of life history it is capable of within its natural context only by observing its actual substance and behavioral tendencies—and not, as was done following Crick and Watson’s elucidation of certain structural features of DNA, by projecting upon it a simplistic logic imagined to capture the entire significance and tendency of this particular material way of being. How different the history of biology would have been if in Crick and Watson’s time we had known what we should have known then and what is now being forced upon us by a dizzying and accelerating series of discoveries—namely that in the DNA “book of life,” as in every production giving us various possibilities of



logical analysis, context counts for nearly everything, and the play of logic always derives from the character of the larger expressive goings-on.

When we are willing to see the mammalian eye with all our powers of observation, we can't help acknowledging that nothing in the digital organism experiments tells us anything at all about whether and how this eye could have evolved upon earth. We can be quite sure that it has in fact evolved, for here it is—but how it might have done so remains largely hidden from us. And those experimenting with digital organisms seem determined to flee all possibility for real answers. Transfixed by the intrinsic force of their own logic, they have lost their investigative anchor in the world's sense-perceptible phenomena. The world has become in their imagination a mere crystallization of their own logic, a process greatly helped by the false conviction that the world can be understood the way we understand the humanly imposed logic of a machine. It's as if the only task of all material substance were to put the logic on display—which is much like saying that the only task of speech is to pronounce whatever logic or grammar we can extract from the speech. But just as speech always has a content setting the terms for any further play of its logic, so, too, the world has a content giving direction to the play of the laws we discover at work in it. Investigating “organisms” without bodies isn't a productive way to explore this content.