THE NEW ATLANTIS
A JOURNAL OF TECHNOLOGY & SOCIETY

CORRESPONDENCE

The Logic of Science

In his article "Ghosts in the Evolutionary Machinery" [Fall 2007], *New Atlantis* contributing editor Steve Talbott takes to task scientists—including me—who use digital organisms to explore principles of evolution. This is not the first time that Talbott has decried modern science's materialist and reductionist approach to understanding the world around us; he has done so in previous essays elsewhere, as well as in these pages (see, for instance, "The Language of Nature" [Winter 2007]).

The central theme of Talbott's writings is that what he calls the "essence" of nature, its "material reality," its "vivid self-presentation"—in short, "the meaning" of life and the universe—can never be understood using the tools, language, and concepts of modern science. Laws, Talbott complains, may predict the behavior of an object, but not the object's uniqueness as compared to other perhaps similar objects in the world. According to Talbott, then, the things that matter are not the generalities but the idiosyncrasies, not the rules but the exceptions, not the classifications but the unclassifiable.

Such a point of view—besides being hopelessly romantic—is also metaphysical, if not mystical. It is based largely on a monumental category error because it assigns to physical objects attributes that only exist in the mind and that have no measurable correlate. For example, Talbott mentions the "vocal, full-bodied self-presentation of cloud, ocean, stone, and sparrow." None of these things has any objective existence. The words themselves do have an element of reality, and so do the feelings those words evoke about nature's beauty and bounty. But they are real by referring to each other: words in relationship to other words, feelings in relationship to other feelings or even to the absence of feeling. It is a category error to assign words and feelings to physical objects as if they represented ontological attributes.

Of course, insisting on physical rather than metaphysical characteristics is precisely what Talbott abhors, but the alternative to a materialistic description of nature—be it dualist or holistic—would entail a total abandonment of the idea of prediction as a measure of understanding. In science, the correctness of a generalization (or abstraction, or theory) is, as Karl Popper argued in *The Logic of Scientific Discovery*, determined by its ability to enable us to make predictions about phenomena we hitherto could not make predictions about, predictions that go beyond the facts that went into producing the theory. Thus, a scientific theory is much more than "a generalization [that] derives its only validity from [the] particulars," as Talbott would have it. Rather, a good theory creates understanding as measured by our ability to predict—to reduce the uncertainty that we have about our world—so that we can take advantage of the world's regularities. The recognition of these laws has allowed us to design new antibiotics, to manufacture eyeglasses and bifocals, to design airplanes and computers, and much more.

Talbott objects to research using digital organisms—self-replicating computer programs that autonomously mutate and evolve to adapt to the simulated world they inhabit—on the grounds that digital organisms are "immaterial" and do not have the "character of real things" (whatever that might be), or the benefit of "contextual reality." He dismisses this research because it represents something he disapproves of in science generally—namely abstraction, and ultimately prediction.

Talbott complains that digital organisms do not have bodies and therefore represent mere mental models that can be used to explore "certain possibilities of mathematical and algorithmic logic" but cannot "elucidate the actual character of real things." First, I note that it is not the business of science to elucidate the character of real things; that is the business of psychology or philosophy or religion. The business of science is to help us understand the world so that we can make predictions about things that, without a scientific understanding, would be unpredictable. Research with digital organisms has certainly had its share of successes—described below—when measured by this yardstick. Second, it is important to see that the focus of digital life research is the *process* of evolution, not any particular instantiation of evolution. That processes can be studied independently from the objects that participate in the processes is at the very heart of modern science; it has enabled a mode of analysis that is responsible for countless discoveries. Weather, for example, appears to be an enormously complex manifestation of the properties of a very large number of entities. If we only sought to understand weather in its glorious unity, we would never be able to predict whether or not we'll need an umbrella tomorrow. But by analyzing processes such as convection, turbulence, dissipation, and so forth—each of which can be studied in isolation, and independently of its substrate—we have made great advances in weather prediction.

Darwin's description of the process of evolution was, in many ways, abstract; he did not know about DNA, proteins, or even the nature of mutations. And indeed, we do not have to make reference to any of these instantiations or implementations when we study evolution, so long as information is stored in a physical substrate. The substrate could be nucleic acids, or characters sketched in crayon, or voltage differences representing bits in a computer. The process of evolution affects all of these equally, so long as the three elements of the Darwinian triad are present: variation (for example, by mutation), inheritance (for example, by self-replication), and selection (through the differential survival of variants). The study of evolution is the examination of this process within a complex, ever-changing, messy, and unpredictable world. The point of digital life research is that this study can be pursued even if evolution is abstracted all the way down to the level of information processing.

Indeed, in studying the past evolutionary history of biological life, we can reconstruct full lines of descent from species' genomes—that is, the informational abstractions of species. This is precisely the reason why bodies are not necessary in the study of the process of evolution. However, while bodies are optional, the *information* that evolves must be real, and there is no doubting the reality of the persistent ones and zeros that constitute the digital organism's genomes. They are as real as the viruses that once in a while attack our host computers.

If our work on digital life were only useful in making predictions about other digital life forms, this endeavor would be a poor one indeed. But just as the great French molecular biologist Jacques Monod exclaimed that "What's true for *E. coli* is true for the elephant—only more so," we have found that insights from digital organisms have analogues in the biological world. Our discovery of the "survival of the flattest" effect in digital organisms (published in *Nature* in 2001) was soon followed by the discovery of the same effect in viroids, several viruses, and microRNA. Our identification of highly epistatic (that is, interacting) mutations in the evolution of complexity (published in a 2003 paper in *Nature*, the very one that Talbott critically cites) was recently discovered in the evolution of steroid receptors by researchers at the University of Oregon. These insights and several others vindicate our approach to abstraction in evolution.

But Talbott's gripe isn't really with digital organisms; he only sees them as a conspicuous perfidy of science. His grievance is with scientific abstraction in general, but this objection is entirely fallacious because it stems from confounding laws with predictions. A law can be fundamental, abstract, and general, and yet still make highly detailed predictions about particulars because a prediction is obtained by plugging initial conditions into the law and deducing the predictions by application of the law. This way the law of gravitational attraction, for example, can be completely general, yet make highly detailed predictions, from the collapse of molecular clouds in the formation of galaxies to the interaction of an atomic force microscope tip with a surface. So clearly a law that is "true of everything" can, contra Talbott, be used to make exquisite predictions about particulars.

It is interesting to note as a final thought that Darwin's theory was so original and unusual in science precisely because of its emphasis on *variation* as the essence of the living world. That centrality of diversity also holds for the critters in our computer-based evolutionary machinery, but this essence can be measured and analyzed. In our computers just as in the biosphere at large, to quote Darwin's final phrase in his *Origin of Species*, "from so simple a beginning endless forms most beautiful and most wonderful have been, and are being evolved."

CHRISTOPH ADAMI
Keck Graduate Institute of
Applied Life Sciences
Claremont, Cal.

STEVE TALBOTT responds: Christoph Adami seems perfectly happy to confirm one of my central contentions: he has little interest in characterizing real entities, biological or otherwise. While he does believe himself to be learning about the evolution of living creatures, he does not pretend that the parameters of his digital organisms correspond either to the identifiable features of known physical life forms, or to any coherently describable and evolving physical part of the computer. His fascination is with neatly computable patterns of logic as such—programs and "information." This fascination is not matched by any evident appreciation of the disciplined observational and experimental work required in order to discover whether and how the finely spun logic of the programmer's closed thought-world bears on real things.

Adami claims that it is not the business of science to elucidate the *character* of real things but rather to seek the kind of understanding, rooted in abstraction, that enables us to make *predictions* about things.

This point of view, which infects much of science today, does indeed go to the heart of the issues between us. Adami is wrong, however, in thinking that I have no use for abstraction or the power of prediction it brings us. I worry, not about abstraction as such, but about the consequences for our understanding when we become so enamored of our abstractions—for example, our ability to abstract measures from observable phenomena—that we forget the phenomena themselves, which alone give us sensible interpretations (meanings) for the abstractions.

It is, after all, the characterizable phenomena that distinguish a work of science from a strictly mathematical or logical (or metaphysical) text. It's remarkable how quickly we have forgotten that the whole point of the scientific revolution was to bring the free flight of medieval intellection down to earth by means of careful observation of real things.

If the language of scientific discovery and prediction were an algebra pure and simple, with no reference to real things, then we could content ourselves with celebrating the joys of abstraction. But, fortunately, the scientist always does refer to real things, and therefore always makes the "monumental category error" of "assign[ing] to physical objects attributes…that have no measurable correlate." It's not at all clear how the rest of us are to make sense of Adami's scientific papers if, as he would have it, we are not allowed to construe his words as referring meaningfully to things.

When we pretend not to be characterizing things, our characterizations simply drop out of consciousness and thereby escape critical attention. We then all too easily begin to imagine some vague sort of inert, Cartesian, machine-like "stuff" whose sole mission, conveniently, is to perfectly instantiate the machine-like rules or algorithms we have come to love—this despite the fact that the physicist gives us nothing remotely like such stuff to work with. (And, in fact, the physicist may nowadays be found reflecting upon how consciousness figures in the science of matter—as blatant an instance of the "monumental category error" as one could find.)

Adami hopes that, if only digital organisms are programmed to exhibit variation, inheritance, and selection—the three ingredients of his evolutionary "algorithm"—then they will tell him what he wants to know about the evolution of living creatures. But an algorithm—a computer algorithm, for example—obtains its algorithmic reliability from the precision of its conception and the uncompromised rigor of its mapping to the minute, painstakingly laid out, and perfectly defined structures of a computer's logical apparatus. Twiddle a zero or one here or there, and the whole thing falls to pieces. Adami and his coworkers do not pretend to provide even the faintest hint of a scheme for mapping their computer algorithms to the hugely complex, ever-changing, mutually interacting, "bit-level" details of organisms, real or imagined. But unless actual organisms can be shown to follow these algorithms, what insights do we gain?

To speak only of variation, inheritance, and selection is unexceptionable, but vacuous. All the worlds of likelihood or unlikelihood, of possibility or impossibility, of development or dead-end extinction, turn upon the details of *how* things vary (historically and exactly), *how* traits are inherited, and *how* selection occurs. This reality scarcely concerns the computational biologist, who enjoys the enviable knowledge that, if he wants to kill off all his "organisms" or make them thrive beyond hope, he needs only to tweak a software parameter or two. Without an effort to match these parameters conceptually or quantitatively to the goings-on in living organisms, what opportunity do we give reality to constrain the programmer's untethered freedom at the keyboard?

According to Adami, the ones and zeroes of his digital organisms are "as real as" computer viruses. It's a useful comparison. Certainly we are referring to *something* when we speak of computer viruses, and this something has to do with the physical states of real machines. But if you want to say anything profound about the "evolution of viruses," then you had better be able to specify what, exactly, you are talking about, whether it's the evolving technical know-how and ethical stance of malicious programmers, or the changing capabilities of anti-virus software, or the continual innovations in computer architecture, or particular characteristics of the prevailing networks, or the education and libidinal desires of computer users, or.... And when you have mastered the intricate complexity of the relevant factors, one of the many lessons you will have learned is the futility of any attempted reduction of the whole to a precise algorithm.

The instinct of the digital-organism enthusiasts when facing such complexity is to "go general." Forget the details; find a truth that applies to any conceivable situation. But, as I pointed out in my article, to go general in an abstract manner generally means to go superficial. Yes, you may arrive at some predictions that, as long as you love generality, you will take some satisfaction in. But they are not likely to be of much help in understanding and working with the objective world. Once you think you have devised an algorithm sufficiently general to say something valid about the evolution of viruses, ask yourself how much predictive power it will give you in the face of whatever network changes or operating system patches or new viruses may show up next week. Of course, we no doubt *can* arrive at good and useful generalizations about viruses, but they will necessarily precipitate out of the detailed sort of understanding mentioned above, and they will not be reducible to simple algorithms.

Adami cites the law of gravitational attraction as an example of a general rule that allows for highly detailed predictions. Understood rightly, this is certainly true. We can be quite confident today that everything we encounter—subatomic particles, light, rocks on earth's hard surface, fish in the sea, solar atmospheres, plasmas, quasars, black holes—will respect a properly stated law of gravity. But in the centuries since Newton, have we been spared the arduous task of acquainting ourselves through direct observation with the radically different characters of these various environments? Laws manifest themselves with different "emphasis" in different contexts; they can be understood more as expressions of the character of such contexts than as one-sided determiners of them. Nothing delivers us from the necessity of learning what sort of context we are dealing with.

When molecular biologists discovered a structure and code for DNA in the 1950s and 1960s, they were given wonderfully neat, clean, and computable abstractions to play with in their minds, and they immediately set off upon an orgy of explanation and expectation based on these abstractions. The "keys of life" were in their hands, and it remained only to work out the details in accordance with a straightforward logic. It has taken these subsequent several decades for it finally to be borne in upon biologists that the logic was horribly simplistic and needed to be radically revised based on the observable details—details that are turning out to be almost unfathomably complex, with arrows of cause and effect running in every possible direction. The researcher has continually had to notice the larger context of the organism in order to get a realistic and more healthily organic picture of the once-unproblematic "mechanisms" of DNA.

As a result of this healthier picture, the concept of the gene has become so obscure that philosopher of science Philip Kitcher could mischievously remark, "A gene is anything a competent biologist has chosen to call a gene." And geneticist William Gelbart adds more seriously that "we may well have come to the point where the use of the term 'gene' is of limited value and might in fact be a hindrance to our understanding of the genome. Although this may sound heretical, especially coming from a card-carrying geneticist, it reflects the fact that, unlike chromosomes, genes are not physical objects but are merely concepts that have acquired a great deal of historic baggage over the past decades." The problem is precisely that the concept was driven too much by dreams of logic and code, and therefore was dissevered from the observations—of the unified, organic, and contextual character of cells and organisms—that alone could discipline the concept and give it a proper meaning. How can you make valid generalizations about the gene and its role in ontogeny and phylogeny without first having a full understanding of the phenomena about which you are trying to generalize?

The simple, numeric gene of the digital-organism researchers, freely available for manipulation by the transcendent programmer and unencumbered by the burdens of physical reality, makes the inflated gene of mid-twentieth-century molecular biology look like the very model of disciplined scientific observation.

All this, I hope, highlights the irony in Adami's characterization of my view, rather than his own, as "metaphysical." He sees me as a hopeless romantic who is interested not in generalities, rules, and classifications, but in idiosyncrasies, exceptions, and the unclassifiable. The truth of the matter is that I highly value generalities, rules, and classification. But I also know that when our fascination with logical structures runs ahead of our disciplined observation, then science really does stand at risk of mysticism, romanticism, and all the rest. How else to explain those occult denizens of the laboratory known as "digital organisms"?

Biodiversity and the Bible

As an Orthodox Christian, I was deeply disappointed in S. M. Hutchens's review of E. O. Wilson's book, *The Creation: An Appeal to Save Life on Earth* ["The Evangelical Ecologist," Fall 2007]. Mr. Hutchens claims that "Christianity's principal reason for the earth's existence is to serve the cause of human redemption"; that the earth exists primarily to be consumed by man; and that the preservation of biodiversity is not "in itself a matter of the highest order." In contrast, the fourth century Church Father Basil the Great prayed,

> O God, enlarge within us the sense of fellowship with all living things, our little brothers the animals, to whom Thou hast given this earth as their home in common with us. We remember with shame that in the past we have exercised the high dominion of man with ruthless cruelty, so that the voice of this earth which should have gone up to Thee in song has been a groan