

Artificial Life: A Feast for the Imagination

DANIEL C. DENNETT

*Center for Cognitive Studies
Tufts University
Medford, MA 02155
U.S.A.*

A Review of Christopher G. Langton (ed.), *Artificial Life*, Addison-Wesley, Redwood City, CA, 1989, 564 pp. plus annotated bibliography and index (655 pp. total) (Volume VI of the Santa Fe Studies in the Sciences of Complexity), \$21.50.

In 1956 history was made at Dartmouth, at a summer conference organized by John McCarthy which inaugurated (and baptized) the field of Artificial Intelligence (AI). The fathers of the field were all there: McCarthy, Minsky, Samuel, Simon, Selfridge, and they took their inspiration from Alan Turing and John von Neumann, whose combined genius had given the world not only the serial computer, but also a strikingly clear vision of its *cognitive* powers. Thirty-one years later, in September of 1987, a meeting was held at Los Alamos which may in time overshadow the Dartmouth meeting and its aftermath: the first international Artificial Life (AL) conference, drawing together thinkers from around the world whose work might be said to owe its inspiration to one of John von Neumann's *other* great ideas: the idea of self-replicating cellular automata.

Not that all workers in AL concern themselves directly with either self-replication or cellular automata. Indeed one of the most striking features of the volume under review, the proceedings of the first Artificial Life conference, is the great diversity of approaches, goals, and grain-sizes of the simulations presented and discussed. This is an extraordinarily rich volume, exciting in tone, kaleidoscopic in its perspectives and contributions.

AI is known for its simulations of cognitive competences modeled "from the top down" and perhaps the most basic difference between AI and AL is that most AL simulations are adamantly bottom up in design philosophy. While some of the virtues of bottom-up modeling are now familiar, especially in the manifestos of the connectionists, broader, deeper reasons are offered here. In *Vehicles* (MIT Press, 1984), Valentino Braitenberg put forward his "law of uphill analysis, downhill synthesis" which made the compelling point that it is simply much harder to figure out the inner machinery of something whose behavior you have data on

than to figu. But the behavior of something whose innards you synthesize (or just posit). Several authors in the present volume press a more radical epistemological point: there is a fundamental difference between the architectures and powers of systems designed in the two different ways, due in part to the epistemological limitations on analyst-designers. Different versions of this claim are developed by Christopher Langton in the introduction, and by Pauline Hogeweg. As Francis Crick has often noted, “Orgel’s Second Law” is that Mother Nature is smarter than you are. Human designers, being far-sighted but blinkered, tend to find their designs *thwarted* by unforeseen side-effects and interactions, so they try to guard against them; Mother Nature, myopic, but with lots of time and the broadest of perspectives, is equipped to *take advantage of* side-effects and interactions. The result is a very different style of system, whose salencies are about as often emergent properties as, shall we say, skeletal properties. And to think that only a few years ago “emergence” was such a loaded term that I felt I had to coin the softener “innocently emergent properties” to forestall premature dismissal!

The ideology of AL is in most other regards in harmony with the ideology of mainstream AI, and I found myself wondering with some amusement whether modern-day vitalists will arise to do battle against “strong AL,” shoulder to shoulder with John Searle. The gauntlet is laid down in a ringing declaration in the Introduction:

Life is a property of *form*, not *matter*, a result of the organization of matter rather than something that inheres in the matter itself. Neither nucleotides nor amino acids nor any other carbon-chain molecule is alive — yet put them together in the right way, and the dynamic behavior that emerges out of their interactions is what we call life. It is effects, not things, upon which life is based — life is a kind of behavior, not a kind of stuff — and as such, it is constituted of simpler behaviors, not simpler stuff. (P. 41)

A common design theme encountered in various guises here is organisms and other systems getting along without benefit of *headquarters* — or at least bosses. Even when entities have heads (and hence, in one sense, headquarters), their headquarters emerge on closer inspection to be unoccupied. The homunculi have all gone away, leaving nothing but machinery behind. This is the essence of von Neumann’s great idea, and hence it grows on one, but initially it is an alien and vaguely repellent idea. Consider, for instance, this simple but eerie description of von Neumann’s basic idea:

Then the parent organism, making use of a memory-stored description of itself, directs the construction arm to move from cell to cell and to emit pulses so as to produce a configuration of cells that is identical to that of the original organism (though initially lacking the “genetic” memory contents of the original). The parent organism then reads its memory a second time and loads a copy of the “genome” information into the memory unit of the offspring organism. The parent organism then injects a pulse to turn

on the new organism and withdraws its constructing arm. This completes the artificial organism's self-reproduction. (Richard Laing, 51).

Even if we are enthusiastic about this idea in some applications, we want to exempt ourselves (we want to exempt our *selves*). The problem is that it *seems* as if *we* at least are very different: *we* are top-down, centered, globally directing. We feel like von Neumann machines, not parallel processors. That is why we have such a dizzy sensation when we encounter "the alternative", why we impute a soul to a termite colony, a behind-the-scenes presence to be the Central Processing Unit. Our (apparent) centeredness must be explained, and the more we learn about biological mechanisms, the more our (apparent) centeredness appears unique in the biosphere. Elsewhere there is Nobody Home. Reflection may convince us that we already know that our own sense of centeredness *must* be something of an illusion, that ultimately we will be explained in all our privacy by "more of the same" — the same kind of modeling exhibited in this volume — but the imaginative chasm is not yet bridged by this observation.¹

The thought-experimental explorations of design-space presented in this volume illuminate the nature of biological possibility, the deepest principles of organization. Reproduction, growth, and self-repair are talents just as exquisite and initially awe-inspiring as cognition, and thanks to these pioneers, we are now beginning to see how they are possible in a "mere mechanism". Philosophers with an occupational bias in favor of cognition should also note that both reproduction and self-repair involve *as a matter of logic* some reflexivity, some self-examination, and thus invoke themes dear to our hearts.

"Bottom up" means *really* bottom up in this book. Neurons are vastly too complicated as starting points for some of the investigations, which deal with such Urquestions as: is it possible for random shuffling to get to the threshold for self-replicating DNA (Rasmussen)?; how does asymmetry in morphogenesis get started from a symmetrical gamete (Turing, initially, but followed up in several ways in this volume)? These are explorations in what we might call proto-biology, but there are also simulations of whole creatures, whole colonies, whole ecosystems, modeled at different levels, investigating different questions.

There are ingenious new toy problems, many of them strongly reminiscent of the "blocks worlds" of AI, but with novel twists. They strike me as equally deep — sometimes deeper than the toy problems of AI — even while they avail themselves of the same sort of audacious oversimplifications, ideas such as the entire economy of nutrition and metabolism boiled down to a "tax" on moving, decrementing a single variable incremented according to an "intake" schedule.

RAM, a simulation platform developed at UCLA, is clearly a magnifi-

cent vehicle for exploration, but so is MIRROR, a rather differently inspired system from Holland. Norman Packard's "bugs" in an artificial ecosystem can be very illuminatingly compared to Richard Dawkins' "Blind Watchmaker" program (complete with a crystal-clear account of the Pascal code behind the scenes). One of the deliberate simplifications of Dawkins' system is what Packard calls *extrinsic adaptation*, and one can discern in the differences between Packard's and Dawkins' approaches the substance of the issues that still divide Dawkins from such anti-adaptationists as Lewontin and Gould.

The simulations span the space from ultra-abstract ("computational metabolism", "typogenetics" and "Petworld") to ultra-concrete (LEGO robots, the prospects of nanotechnological robots that might "live" in your bloodstream, making repairs and performing other services). The color plates of various simulated life forms from the molecular to the macroscopic are gorgeous. There is a good mix of straightforward presentation of running simulations on the one hand and speculation and moral-drawing on the other. In addition to the authors mentioned above, I particularly recommend K. Eric Drexler's observations on "M systems" versus "O systems", which explore the fundamental differences between designed and evolved systems, the relative powers of diffusive versus channelled transport, matching versus positional attachment, and the reasons why (engineer-) designed systems can't evolve.

One can see that AL has learned from the errors and controversies of AI — but not enough! Some of the rhapsodic advertisements for particular systems are all too reminiscent of the overselling of AI, but the spirit of constructive criticism is refreshing, and there is very little badmouthing. Will it come later? I hope not. Will AL revolutionize biology the way AI has revolutionized psychology? I think there is a good chance that it will, and for similar reasons. Philosophers, in the meantime, will find in this opening manifesto a feast for the imagination.

NOTE

¹ Nicholas Humphrey and I sketch a model for the process of self-creation in 'Speaking for Our Selves: An Examination of Multiple Personality Disorder', *Raritan*, Summer, 1989.