

Artificial Life as Theoretical Biology:
How to do real science with computer simulation

Geoffrey F. Miller

Cognitive and Computing Sciences

University of Sussex

Falmer, Brighton BN1 9QH, England

geoffm@cogs.susx.ac.uk

Abstract

Artificial Life (A-Life) research offers, among other things, a new style of computer simulation for understanding biological systems and processes. But most current A-Life work does not show

This romantic, colonialist view seems implicit in much of Artificial Life / A Life research. It maintains a kind of Orientalism of the Organic that views complex adaptive systems as newly discovered things and evolution, development, and learning as newly disco-

February 1994

designed through CS research. CS methods are much closer to analytical methods used by engineers such as finite element analysis, stability analysis, or perturbation analysis than to empirical methods used by scientists. This basic distinction leads to several other differences between RS and CS, which I will list with brazen simplicity and tongue slightly in cheek.

RS seeks knowledge of nature, whereas CS seeks profit through technical improvements in hardware and software. RS chooses problems for their theoretical and practical importance, whereas CS chooses problems for their economic relevance. RS analyzes existing natural systems through developing speculative theories subject to hypothesis testing by observation, experimentation, simulation, whereas CS builds new artificial systems through engineering and debugging. RS requires skills in observation, experimentation, comparative analysis, statistics, scholarship, interdisciplinary communication, and bold imagination, whereas CS requires skills in programming and mathematical analysis. RS rewards discoveries for their theoretical generality, experimental clarity, and practical applicability, whereas CS rewards inventions for their patentability, marketability, and complexity. RS advances through public, loosely organized collaborations called research areas by huge numbers of researchers working over many years, whereas CS advances through largely private, more tightly organized collaborations called project teams by smaller numbers of programmers working over

seve

whereas the phenomena studied by computer science (e.g. computation) depend on the science for their very existence. Computer science is more similar in nature and spirit to architecture and aeronautical engineering than it is to physics or biology.

Granted, computer science has been spectacularly successful as an engineering discipline; computer speeds and computer sales have grown exponentially over several decades, continuing to double every few years. But computer scientists run into trouble when they try to do real science, because they are simply not trained for it. Whenever computer scientists try to do real psychology for example, a boom and bust cycle results, as in artificial intelligence during the 1980s, cognitive science during the 1990s, and neural networks research during the 2000s. With suggestive initial results come wild promises of further progress and massive influxes of research funding; intellectual stagnation sets in as the promised conceptual breakthroughs remain elusive, but modestly useful real world applications keep the field limping along for another few years. I am afraid that A Life will repeat this same pattern.

Again and again, the same problems arise when computer scientists develop a new field that claims to be real science. The field shows historical amnesia, interdisciplinary blindness, and ignorance of current work in the relevant existing sciences. The field lacks explicit hypothesis testing, systematic observation, controlled experimentation, and statistical analysis. The field avoids recognizing or understanding its failures and overgeneralizes its successes in wildly inappropriate ways. The field jinks from one research fad to another, failing to replicate and extend its findings in ways that could lead to a conceptually integrated discipline. Many of these problems result from an engineering mentality that seeks to build impressively complicated masterpiece systems rather than to develop simple theories that explain complex phenomena. Masterpieces of computer programming such as those developed by

February 1994

Ph.D. students in artificial intelligence or cognitive science are almost always useless as scientific models because they usually require too much pre-processing of the input, too generous an interpretation of the output, and too many hidden assumptions and hacks.

These precedents suggest that unless A Life is very careful, it will become a historical curiosity in short order after being more or less successful in squandering many millions of dollars of public research funding that could have been better spent on a few bright evolutionary biologists and evolutionary psychologists. The remainder of this paper suggests some ways we can learn from the mistakes of the past, by identifying some methodological heuristics for doing A Life as good theoretical biology, and some areas of biology that may benefit especially from A Life simulations.

3 Six Methodological Heuristics for A-Life

In no way is it intended to be a comprehensive list of heuristics for A-Life simulation.

This step is much harder than it sounds. Biology is a mature, successful science that has become quite sophisticated over the last few decades. One cannot just pick up a copy of

Dawkins (1982) or an introductory undergraduate biology textbook, find an interesting sounding issue, and forge ahead with a simulation. A basic maxim of modern science is that one should not attempt to simulate a complex system until one has a good understanding of the system's basic principles and processes.

February 1994

biological issues much better than A Life researchers do.

So, it is difficult to find a significant unsolved problem that can be addressed through A Life computer simulations, and that has not already been addressed by standard methods of theoretical biology e.g. verbal argument, formal population genetics models, optimality theory, evolutionary game theory

A Life as Theoretical Biology

Geoffrey F. Miller

February 2004

A Life as Theoretical Biology

Geoffrey F. Miller

February 1994

9

comparative method Harvti

A Life as Theoretical Biology

Geoffrey F. Miller

February 1994

1

uniform age and sex s

intellectually exciting. It is also more fun to speculate about simulation's metaphysics than its methodology. Philosophizing about A Life has its own seductions and snares that distract attention from A Life's real potentials and problems as science.

To ensure A Life's relevance as theoretical biology, we must develop better methods for measurement and experimentation in our simulations. Whereas graphics may briefly catch the eyes of theoretical biologists, only solid experiments will win over their hearts. Our observational and experimental methods will have to approach the sophistication of taxonomy, comparative biology, ethology, psychology, ecology, and evolutionary biology. And we will have to develop new methods, because simulation can yield data that empirical biology cannot.

Early A Life research consisted largely of proof of concept demonstrations: local interactions can lead to certain emergent effects suggestive of biological systems. Such results were important in establishing the concepts of self or

infer the dynamics of long term processes. Such processes could not be studied directly in replicable, controlled experiments that could yield large amounts of relevant data

Simulation allows access to much richer information as evo

February 1994

1

specific problems, it is not trivial to know whether the code is working appropriately. Making sure the code compiles and the program doesn't crash is just the first step. With humbling frequency, a simulation that seems to produce reasonable and interesting data under one set of conditions will produce, under slightly different conditions, weeks later and often immediately before a conference, anomalous output that highlights some hidden, tiny, critical bug that invalidates weeks of results. There are three strategies for a

biologists bother to read it

4 What distinguishes A-Life from other biological simulation?

Simulation is no stranger to biology. Theoretical biology papers that include simulations appear regularly in journals such as *Artificial Life*, *Journal of Theoretical Biology*, *Complexity*, and *Journal of Computational Biology*. Several journals are heavily biased towards biological simulation, such as *Artificial Life*, *Journal of Computational Biology*, *Journal of Theoretical Biology*, *Journal of Computational Intelligence and Applications*, *Journal of Computational Neuroscience*, and *Journal of Computational Biology*. What does A-Life offer that theoretical biology does not already have?

Traditional mathematics and simulations in theoretical biology try to capture self-organizational or evolutionary dynamics directly in equations or simple procedures that aim straight for the collective, emergent level. Such methods are only tractable when they implicitly represent the components of biological systems as simple, stable, homogeneous, and predictable. Equations don't generally allow surprising, emergent behaviour, and neither do simulations based on equations.

The key advance in A-Life has been to allow emergence by representing individual biological things explicitly as computational procedures, allowing self-organizational and evolutionary processes to emerge spontaneously from these things, and making observations and measurements about the resulting patterns and dynamics at the individual and collective levels (see Taylor & Jefferson, 2000). This advance has allowed A-Life to explore the interaction of many biological units (molecules, cells, organisms, or populations) at several levels of description over different time scales (behaviour, development, or evolution).

A Life has used emergence in two main ways⁴ as a proof of concept to show that certain biological phenomena can arise from distributed interactions among many local components, or as an extension to make current theoretical biology models more complete and realistic. The first way has led to some messianic predictions that theoretical biology will be revolutionized, perhaps with emergence replacing evolution as the central explanatory principle of life. The second way is more conservative and, I think, more useful⁵ it embraces emergence without getting obsessed with it. Kauffman's⁶ research, despite its Emergentist rhetoric, seems a powerful example of this second strategy.

A Life simulations have other advantages. Conceptually, the requirements of programming force researchers to make assumptions explicit and processes computable; formal equations are actually quite ambiguous compared to computer programs. Computationally, the speed of modern hardware allows biological systems to be simulated at levels of complexity unimaginable only a couple of decades ago. Experimentally, simulation offers a high degree of control, flexibility, and replicability. Analytically, simulation allows accurate measurement of very large amounts of data, and automatic statistical analysis and visualization of that data. Socially, simulation code can be shared over computer networks, promoting easy replication and extension of results.

There is a continuum between A Life as high level theoretical biology, studying general processes and patterns of evolution, and A Life as empirical biology, modelling specific data from certain taxa. Some A Life models the mechanisms or effects of a certain set of behaviours from a single species; but many biologists already develop such models routinely in their empirical work. More usefully, A Life can model more general classes of behaviours shared across many species, such as courtship, cooperation, pursuit and evasion, communication, collective behaviour, or flocking. Models that predict different varieties of the behaviour for different species under different conditions would prove especially valuable to empirical biologists.

5 Open questions in theoretical biology that A-Life might usefully address

Given these strengths of A-Life, we can turn now to ask what open questions in theoretical biology might be especially amenable to A-Life simulation. The following areas seem promising: evolutionary innovations, interactions between different adaptive processes and different selective forces, origins and effects of mental and behavioural adaptations, and life as it could be – logically and extra-terrestrially.

Theoretical biology cannot yet explain major evolutionary innovations such as the evolution of life (Eigen & Schuster, 1989; 45; sex (Williams, 1966; 45; Margulis & Sagan, 1990; ; Maynard Smith, 1982; ; Michod & Levin, 1992) – multi-cellular bodies (Buss, 1988) – or the human brain (Miller, 1999; ; Ridley & Maynard Smith, 1981; Mara

A Life as Theoretical Biology

and comparative psychology only reveal the outcomes of neural and cognitive evolution, and fossils do not provide details of neural circuits during evolution. A Life systems that explicitly simulate the evolution of nervous systems interacting with each other and with complex environments may be our only hope for constructing theories of mental and behavioural evolution.

Once behavioural and cognitive adaptations have evolved, they can influence the further course of evolution within and across species. The role of mind in guiding evolution has usually been overlooked entirely, or conflated with a mystical, progressivist, animism as in the work of Herbert Spencer (1859); see Godfrey Smith (1999) and William McDougall (1909); see Boden (1990). Only a few biologists developed a Darwinian view of minds as selective forces (Morgan (1902); Thompson (1908); see Richards (1987)). But very recently, several theoretical biologists have begun to recognize the importance of perception and cognition as selective forces in the evolution of diverse phenomena such as camouflage, mimicry, warning colouration, sexual ornaments, flowers, fruits,

6 Does Strong A-Life allow stronger theoretical biology research than Weak A-Life?

The debate over strong A-Life (computer processes as realizations of living systems) versus weak A-Life (computer processes as simulations of living systems) can be extended in a methodological direction by asking: would it make any difference to theoretical biology if an A-Life system were construed as a realization rather than a simulation? Clearly, empirical biology would be affected: we would have to add a sixth kingdom of life to the current five (see Margulis & Schwartz, 1998), and I suppose that databases of biological phylogenies would have to be updated every time a new Ph.D. thesis in A-Life was written. Also granted is that acceptance of strong A-Life would imply that life-like mind is a functional, emergent property of certain systems with certain internal relations, external relations, and evolutionary histories (see Millikan, 1999). But my focus is on biological theory: what could we learn about life and evolution from strong A-Life that we could not learn from doing weak A-Life?

In answering this question, we must appreciate that many theoretical biologists view instances of real terrestrial life as little more than the outcomes of simulations by Nature to inform them about how evolution works. That is, empirical biology constrains biological theory in almost the same way that simulation does, so for some theoretical biologists, even real terrestrial biology could be considered weak A-Life. Life forms that result from artificial selection by human breeders or genetic engineering also blur the distinction between realization and simulation. If one makes a strong division between Nature and Culture, such life forms are experimental simulations of what would happen if a lineage were subjected to some selective pressure or mutation in Nature; from a more integrated perspective, such life forms are simply the outcome of a thoroughly Natural process that happens to include humans as selective forces. Likewise, experimental biology research that records animal behaviour in unnatural laboratory conditions could be viewed either as realizations of behaviour *n*

February 1994

or simulations of natural behaviour. Does it make any difference to theoretical biology either way?

Consider an issue in evolutionary theory that might be solved in two ways. Ms. Goodmaths develops a mathematical population genetics model that represents changes in gene frequencies using differential equations, whereas Mr. Badmaths programs an A Life model that represents genes themselves in a genetic algorithm. In the latter case, one might argue that the genes are "alive" in the limited sense of replicating because they really are copied within computer data structures. Suppose the two models are both good and give the same answer. The "aliveness" of the genes in the A Life model is simply irrelevant to the theoretical biology. If the models are formally equivalent, one can move smoothly from the differential equations through discrete iterative approximations to genetic algorithms without affecting the results in the slightest. The evolutionary dynamics have multiple realizability in formal equations, computational approximations, and procedural simulations. Only if Mr. Badmaths' simulation goes beyond the complexity that Ms. Goodmaths' equations can represent, is there any point in doing the simulation.

An analogy to high energy physics may also help here. The events that occur within colliders at CERN have an ambiguous status. Empirical physicists treat them as "real physical events" that reveal forces operating outside colliders. But for theoretical physicists, colliders may as well be viewed as very special, very expensive computers that simulate the physics of the very early universe, shortly after the Big Bang. It does not really matter for theoretical physics whether collider events are viewed as "realizations of current physics" or "simulations of early universe physics," because in both cases the problem remains of how to generalize to processes outside the collider.

February 1994

Todd [unclear]. Such simulations will probably become the most important scientific tool for understanding evolution since Darwin [unclear] first de

Acknowledgments

This research was supported by NSF NATO Post Doctoral Research Fellowship RCD 445 and NSF Research Grant INT 990909. Thanks to Maggie Boden, Dave Cliff, Inman Harvey, Peter Todd, and Michael Wheeler for useful discussions.

References

- Ackley, D., & Littman, M. (1999). Interactions between learning and evolution. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 45-60). New York: Addison Wesley.
- Ackley, D., & Littman, M. (1999). A case for Lamarckian evolution. In C. Langton (Ed.), *Artificial life III* (pp. 1-10). New York: Addison Wesley.
- Alcock, J. (1999). *An introduction to behavioral ecology* (4th Ed.). Sunderland, MA: Sinauer Associates.
- Andersson, M. B. (1999). *Sexual selection*. Princeton U. Press.
- Barkow, J. H., Cosmides, L., & Tooby, J. (Eds.). (1999). *The cognitive niche: Social intelligence, language, and human evolution*. Oxford U. Press.
- Barth, F. G. (1999). *In search of a naturalistic theory of evolution*. Princeton U. Press.
- Bateson, P. (1999). The active role of behavior in evolution. In M. W. Ho & S. W. Fox (Eds.), *Evolution and behavior* (pp. 1-10). New York: John Wiley.
- Bedau, M., & Packard, N. (1999). Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 1-10). New York: Addison Wesley.
- Boden, M. (1999). *Artificial intelligence: Foundations and applications*. Harvard U. Press.

February 1994

Boekhorst, I. J. A. te., Hogeweg, P. 1993. Effects of tree size on travelband formation in *Orania*. Data analysis suggested by a model study. In R. A. Brooks & P. Maes, Eds., *Artificial Life II*, pp. 45-52. MIT Press Bradford Books.

Boyd, R., Richerson, P. J. 1985. *The Logic of Cultural Evolution*. U. Chicago Press.

Buss, D. M. 1988. *The Evolution of Human Intrasexual Competition*. New York: Basic Books.

Buss, L. W. 1988. *Sexual Selection: A Comprehensive Theory*. Princeton U. Press.

Byrne, R., Whiten, A. Eds. 1989. *Intentional Cognition: Animal Cognition and the Role of the Mind*. Oxford U. Press.

Cariani, P. 1993. Emergence and artificial life. In C. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen, Eds., *Artificial Life II*, pp. 45-52. New York: Addison Wesley.

Cavalli Sforza, L. L., Feldman, M. W. 1981. *Evolution and the Human Genome*. Princeton U. Press.

Cliff, D. 1993. Computational neuroethology: A provisional manifesto. In J. A. Meyer & S. W. Wilson, Eds., *Artificial Life III*, pp. 1-10. MIT Press.

Cliff, D., Miller, G. F. in press. Co Evolution of pursuit and evasion II: Simulation methods and results. For *Artificial Life III*.

Clutton Brock, T. H. 1991. *Evolutionary Ecology*. Princeton U. Press.

Collins, R. J., Jefferson, D. R. 1993. The evolution of sexual selection and female choice. In F. J. Varela and P. Bourguin, Eds., *Artificial Life III*, pp. 1-10. MIT Press.

Cronin, H. 1991. *Evolutionary Ecology*. Cambridge U. Press.

Daly, M., Wilson, M. 1988. *Evolutionary Psychology*. Boston: Willard Grant Press.

A Life as Theoretical Biology

Geoffrey F. Miller

February 1994

Daly

February 1999 45

Eldredge, N. 1985. *Evolution of the eye*. Oxford U. Press

Eldredge, N. 1995. *Developmental constraints*. New York: McGraw Hill

Endler, J. A. 1999. Signals, signal conditions, and the direction of evolution. *American Naturalist* 154: 46-45

Ewald, P. 1999. *Evolution of the brain*. Oxford U. Press

Fagen, R. 1991. *An evolutionary approach to design*. Oxford U. Press

Fisher, R. A. 1930. *The genetical theory of natural selection*. Clarendon Press

French, R. M., Messinger, A. 1999. Genes, phenes, and the Baldwin Effect: Learning and evolution in a simulated population. In R. A. Brooks, P. Maes, Eds. *Artificial Intelligence and Evolutionary Computing*. MIT Press Bradford Books

Foley, R. 1995. *Animal communication*. Harlow, Essex, England: Longman Scientific Technical

Futuyama, D. J. 1986. *Evolutionary biology*. 2nd Ed. Sunderland, MA: Sinauer

Futuyama, D. J., Slatkin, M. Eds. 1983. *Evolutionary biology*. Sunderland, MA: Sinauer

Gale, J. S. 1999. *Evolution of the eye*. London: Unwin Hyman

Gallistel, C. R. 1999. *Learning and cognition*. MIT Press

Godfrey Smith, P. 1999. Spencer and Dewey on life and mind. In R. A. Brooks, P. Maes, Eds. *Artificial Intelligence and Evolutionary Computing*. MIT Press Bradford Books

Guilford, T., Dawkins, M. S. 1991. Receiver psychology and the evolution of animal signals. *Animal Behaviour* 41: 4-1991

Harvey, P. H., Pagel, M. D. 1991. *The evolution of the eye*. Oxford U. Press

Hirschfeld, L. A., & Gelman, S. A. Eds. 1994. *Evolution and the Origin of Species*. Cambridge U. Press.

Kauffman, S. 1995. *At Home in the Universe: The Search for Life's Self-Organizing Principle*. Oxford U. Press.

Kohn, D. Ed. 1994. *Complexity: The Emerging Science at the Edge of Order and Chaos*. Princeton U. Press.

Krebs, J. R., & N. B. Davies Eds. 1997. *Foraging Ecology and Optimal Decision Making*. 2nd Ed. Oxford Blackwell Scientific.

Langton, C. 1995. Artificial life. In Langton, C. Ed. *Artificial Life*, pp. 1-10. New York Addison Wesley.

Lewin, R. 1999. *The Core of Life*. 2nd Ed. Boston Blackwell Scientific.

Mangel, M., & Clark, C. W. 1989. *Dynamic Modeling in Plant Ecology*. Cambridge U. Press.

February 1994

Miller, G. F. (1994). Two dynamic criteria for validating claims of optimality. *Journal of Theoretical Biology*, 166, 41-59.

Miller, G. F. (1999). *The Mating Game: How Sexual Selection Shaped the Evolution of Human Intrasexual Competition*. Ph.D. thesis, Stanford University Psychology Department. To be published as a book by MIT Press Bradford Books in 1994.

Miller, G. F. (1999). Exploiting mate choice in evolutionary computation. Sexual selection as a process of search, optimization, and diversification. In T. C. Fogarty (Ed.), *Artificial Evolution: Proceedings of the First International Conference on Artificial Evolution* (pp. 4-19). Springer Verlag.

Miller, G. F., & Cliff, D. (1999). Protean behavior in dynamic games. Arguments for the co-evolution of pursuit evasion tactics in simulated robots. In D. Cliff, P. Husbands, J. A. Meyer, & S. Wilson (Eds.), *An Artificially Evolving World: Proceedings of the First International Conference on Artificial Evolution* (pp. 1-10). MIT Press.

Miller, G. F., & Freyd, J. J. (1999). *Evolutionary Psychology: The Origin of Mind*. Cognitive Science Research Paper CSRP 9, University of Sussex.

Miller, G. F., & Todd, P. M. (1999). Exploring adaptive agency I. Theory and methods for simulating the evolution of learning. In D. S. Touretsky, J. L. Elman, T. J. Sejnowski, & G. E. Hinton (Eds.), *Evolutionary Learning: Proceedings of the First International Conference on Evolutionary Learning* (pp. 4-15). San Mateo, CA: Morgan Kaufmann.

Miller, G. F., & Todd, P. M. (1999). Let evolution take care of its own. *Journal of Theoretical Biology*, 211, 1-10.

Miller, G. F., & Todd, P. M. (1999). Evolutionary wanderlust. Sexual selection with directional mate preferences. In J. A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), *An Artificially Evolving World: Proceedings of the First International Conference on Artificial Evolution* (pp. 1-10). MIT Press.

February 1994

1994

Sigmund, K. 1994. *The Philosophy of Language*. Oxford U. Press.

Simpson, G. 1994. *The Philosophy of Language*. Columbia U. Press.

Sober, E. 1994. Learning from functionalism: Prospects for strong artificial life. In C. Langton, C.

Taylor, J. D. Farmer. S. R.

February 1999 45

Webb, B. (1999). Robotic experiments in cricket phonotaxis. In D. Cliff, P. Husbands, J. A. Meyer, S. Wilson, & An. An. In *Artificial Intelligence and Robotics*. MIT Press.

Weismann, A. (1984). The selection theory. In *The Philosophy of Language*. New York: Boni and Liveright.

Williams, G. C. (1966). *Adaptation and Natural Selection*. Princeton U. Press.

Williams, G. C. (1975). *Sex and Selection*. Princeton U. Press.

Wilson, E. O. (1975). *Social Insects*. Harvard U. Press.

Wright, S. (1978). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Evolution and the Genetics of Populations*. Princeton U. Press.