A Discussion of the Use of Artificial Life Models to Evaluate Gould's Hypothesis about Progress in Evolution

Paul Domjan

The Plan II Honors Program
The University of Texas at Austin
WCH 4.104
Austin, Texas 78712-1105
pdomjan@mail.utexas.edu

Abstract

Stephen Gould has introduced the hypothesis that progress in evolutionary history is due to contingency. Daniel Dennett has further suggested that this hypothesis could be confirmed or denied by artificial life models, in which multiple evolutionary histories can be produced and the role of contingency thus evaluated. While existing models do not allow for this evaluation, Calabretta *et al.* (1998) suggest a general form to be used in developing artificial life models to answer questions about biology. Using this form, a model is outlined which may be able to test this hypothesis.

Gould's System

In 1979, Stephen Gould and Richard Lewontin published their seminal article "The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme." This article argued that, although we have traditionally viewed all variation in traits over time as adaptive, many traits actually result from structural necessity or chance. In Gould and Lewontin's view, it is inappropriate to tell an adaptationist story for each individual trait of the organism, rather we should look at traits as largely the result of how the organism uses and changes what originally appears as a result of its form.

This simple critique is extended in Gould's 1989 discussion of the Burgess Shale, *Wonderful Life*. In *Wonderful Life,* Gould argues that, if we were to begin with the extreme diversity of multi-cellular life present at the time of the Burgess Shale and replay the tape over again and again, we would get a new collection of surviving phyla, and, thus, life would take a radically different course. Contingency plays a major role in the process of evolution, which selected the Burgess creatures that would survive. Nevertheless, one is still able to tell an adaptationist story about why certain types of creatures survived. Under this view of the Burgess shale, the appearance of intelligence and consciousness are the result of contingency in the history of life.

However, contingency and adaptationist selection are not incompatible. Rather, the issue at stake is to what degree each has contributed to the create the current collection of forms of life. The adaptationists argue that life appears in its current form entirely, or at least primarily, as a result of the specific adaptive advantage of each particular trait. Gould and Lewontin, on the other hand, argue that, while some traits may evolve for adaptive reasons, contingency is the most important factor in determining survival and many traits appear because of contingency.

In Gould's 1996 work, *Full House*, he argues that the maximum complexity, which is closely related to maximum organism size, of organisms has increased not as a result of some essential direction in evolution. This claim has two parts: (1) If size has no adaptive value, but there is a minimum viable size, then a tail of high complexity organisms will develop, and (2) in the biosphere size has no adaptive value. Organism size has a minimum, namely the smallest a prokaryote can be, but no maximum. Thus, without any particular tendency towards increased complexity, random variation will tend to produce larger, more complex organisms. However, the mode of organism complexity remains the prokaryote, who lie very close to the minimum organism size. Furthermore, there is no direct line of increase in complexity. Rather, largely unrelated organisms occasionally stumble onto the large end of the complexity scale. As such, any replay of the tape would lead to a vastly different collection of complex organisms, though not necessarily a different distribution of sizes. Thus, thus "progress," that is continual increase in organism complexity, is not a result of the adaptive value of complexity, but of the neutral value of complexity.

Dennett's Critique of Gould

Daniel Dennett (1995) has criticized Gould for not using a computer, and, as a result of his technological incompetence, ignoring the possibility that computer models of evolution produced by the field of Artificial Life might actually allow one to replay the tape of life again and again, seeing which direction evolution will take. Interestingly, Dennett does not do this either. However, were one to do the experiment, either each run will produce vastly different results, suggesting that Gould is correct, or, as Dennett theorizes, each run will find the same "Good Moves in Design Space" each time it is replayed, producing similar results each time. Thus, Dennett believes that every time the tape is replayed, certain useful traits, like consciousness, will be arrived at in some way by some species. This simple criticism opens the possibility that Artificial Life modeling could develop real instantiations of evolution, which could then be studied to evaluate Gould's criticism.

Modeling Evolution: Replaying the Tape

Artificial Life is the study of synthetic, computer based, systems which are actual instantiations of some properties of real living organisms. Thus, while 'Boids' (Reynolds 1987), an artificial model of flocking behavior, does not include any real birds, the 'boids' in it are actually flocking. That is, each 'boid,' based on its surroundings, is making decisions about what direction to move, and, as such, the 'boids' move together as a flock. This sort of model could be used to analyze and categorize flocking behavior.

Unfortunately, open ended evolution is much more complicated than flocking. The time scale of flocking is very short: either the organisms move in a flock or they do not. Furthermore, flocking is a property of small collections of organisms, not a property of a larger, time-indexed system formed by the totality of all the organisms. Evolution, however, takes place across time and is a property of the evolving system, not any smaller collection of individuals in the system. There are three steps which must be completed in order to evaluate Gould's hypothesis. First: a model of open ended evolution must be developed. Second: the model must be tested to determine the role of contingency in its evolution. Third: the model must be tested to determine whether it approximates evolution in the biosphere. We will find that, while artificial life models exhibit one sort of evolution, this is not the same sort of evolution as in the biosphere. Thus, while they tell us something about the nature of evolution, they do not help us evaluate Gould's hypothesis about the natural of evolution in the biosphere.

Modeling Evolution

Tierra (Ray 1992) is probably the most famous artificial evolving system ever developed. The system has an ancestor organism, a simple computer program, which has a start sequence, a reproductive loop, and a copy sequence. The system evolves for a period (a couple of hours at most on a relatively fast computer), develops a relatively complicated ecosystem, which includes parasitism and, occasionally, social behavior, and then stops changing significantly. While there is still some genomic change in the system, this change can all be attributed to genetic drift rather than adaptation.

Ray and Hart (1998) have developed an extended, multi-threaded (multi-cellular) version of Tierra, Network Tierra, which incorporates two different types of cells. Their initial data seems to suggest that Network Tierra suffers from the same problem as Tierra: the environment presents a problem, which the organisms then proceed to solve. Once the organisms have developed an efficient solution to the problem, adaptive evolution stops (Bedau *et al.* 1997a).

Although there are many artificial life models, some of the major ones being Echo (Holland 1994, 1995), Bugs (Bedau and Packard 1992), and Tierra derivatives, like COSMOS (Taylor 1997) and Avida (Adami and Brown 1994), all of these models are similar in character to Tierra. They all continue evolving until they solve the problem presented by their environment. Thus, in order use these models to evaluate Gould's hypothesis, we must determine, among other things, whether a model which only evolves to a point can tell us about evolution in the biosphere.

Determining the Role of Contingency in Evolution

In 1992, Ray formulated the following agenda for artificial life:

> Because biology is based on a sample size of one, we cannot know what features of life are peculiar to earth, and what features are general, characteristic of all life.... A practical alternative to an inter-planetary or mythical biology is to create synthetic life in a computer. "Evolution in a bottle" provides a valuable tool for the experimental study of evolution and ecology. (371-2)

While evolution on earth cannot be replayed, computer evolution can be replayed many times. However, artificial life has only recently been used to address this question. Taylor and Hallam (1998) attempt to further illuminate this issue by looking at multiple runs of COSMOS (Taylor 1997) to determine what role contingency plays in COSMOS. The authors find that, regardless of the seed, the simulation produces similar results, thus, they conclude that, in artificial life models, while contingency effects the specifics of the outcome, the general result is determined by the problems presented by the environment. In Dennett's (1995) formulation, while Stephen Gould's existence may be due to contingency, the fact that someone would inhabit his office at Harvard is not. Taylor and Hallam's results point in this general direction. However, Taylor and Hallam realize that these initial results do not shed much light on Gould's hypothesis.

Nevertheless, it may be that there are central differences between the current generation of artificial life systems and the biosphere which cause multiple runs of an artificial life model to evolve in basically the same direction, making such systems unfit to evaluate Gould's claim. In Gould's (1996) view, increased body size is the result of drift away from a left wall with little or no pressure to evolve in either direction. Thus, given time, complexity will continue to drift away from the simplicity of the modal bacter. Artificial life models, however, usually force organisms to compete for clock cycles on the computer. As such, larger organisms are significantly less fit than their better optimized cousins, because they require more clock cycles to reproduce, and, thus, these models are not a system in with size has no adaptive value, but there is a minimum size. Because these models do not fulfill the conditions of the disjunctive first part of Gould's claim, they cannot be used to evaluate this claim, and they certainly cannot be used to evaluate the empirical hypothesis of the second part of his claim. Therefore, they cannot be used to evaluate the claim. It may be that, while contingency essentially plays a major role in the development of a tail of complex organisms, leading away from the modally dominant bacterium is the biosphere, such a tail is unable to develop in an artificial model, because large organisms are significantly maladapted.

Evaluating the Models

One should be hesitant to use artificial life models to make claims about the nature of the biosphere, although, as Calabretta *et al.* (1998) show, evolutionary models can be very useful in investigating a simple property of a system. Firstly, one needs to isolate the feature. Secondly, with a sufficient understanding of the feature to begin with, one must build an artificial life model which incorporates the aspects of the feature pertinent to the problem in question. Finally, this artificial life model can be used to directly manipulate the feature and evaluate the consequences.

In Calabretta *et al.*'s case, modularity of structures and recursive design was investigated by developing neural networks, some with modules and some with the possibility of developing modules. Both sets of networks were presented with a problem, and those with modules where able to solve it faster. However, upon examining the networks which began without modules, it was found that they developed much more specific modularity, thus allowing the researchers to see how modules could be seen as functional units. This research was successful, in large part, because the structure the researchers were investigating was simple and they were able to easily isolate the aspects of it which were pertinent to their problem.

Bedau *et al.* (1997a, b) has done extensive work in classifying evolutionary system, with the aim of showing that artificial life systems are qualitatively different than the biosphere. Bedau's work is based on using a neutral variant of a simulation, in which genotype does not affect survival, to determine how long genotypes persist in the absence of selection. This data is then used to determine whether genotypes, in the presence of natural selection, persist longer than they would in the neutral model, indicating that they are persisting because they are better adapted. Using this data, Bedau *et al.* (1998) find that artificial life models either never engage in adaptive evolution, or stop changing after a point.

It is assumed that:

> [T]he mere fact that a [taxonomic] family appears in the fossil record is good evidence that its persistence reflects its adaptive significance. Significantly maladaptive taxonomic families would likely go extinct before leaving a trace in the fossil record. (Bedau *et al.* 1998: 229)

This begs the question by assuming that genomes which appear in the fossil record are evolutionarily significant. Because it is impossible to create a neutral model for the biosphere, it may be that the biosphere, like the artificial life models, also has solved the problem presented by its environment, and, now, is simply drifting through evolutionary space relative to the basic structure of its environment, meaning that, at an early stage in evolution, organisms developed solutions which allowed them to survive in the biosphere (Domjan 1999). This is not to say that there are no further evolutionary pressures, on the contrary constant threats, being eaten by tigers for example, may appear to threaten organisms. However, being eaten by tigers is not a problem of the structural features of one's environment, but a problem created by the other organisms in one's environment. Both the tiger and the animal it is trying to eat have solved the basic problem of being alive in the biosphere.

This would even be consistent with Bedau *et al.*'s assumption about taxonomic families. That is, if the biosphere has solved the problem presented by its environment, none of the surviving taxonomic families will be "significantly maladaptive." In Tierra, for example, once adaptive evolution stops, random genomic change persists, and the new genotypes which survive are not significantly maladaptive. Rather, they are not significantly better adapted than the original ancestor organism. Furthermore, the fact that the majority of organisms are bacteria further supports this possibility. If the environment of the biosphere presented a problem which required an increase in complexity to solve, bacteria would have been superseded by larger organisms. However, because this has not happened, it stands to reason that a bacteria is capable of solving the problems presented by our biosphere. We cannot dismiss Gould's claim on the basis of artificial life evidence produced by existing models because they may not be sufficiently analogous to the biosphere, either because a structural feature of their environments (i.e. competition for clock cycles) imposes significant downward pressure on organism size, or because, unlike the biosphere, they do not engage in long run adaptation. These problems show us that, using our current generation of artificial life models, we are not able, as Calabretta *et al.* (1998) did, to isolate and implement those features of evolution (i.e. the first half of the conditional part of Gould's claim) which are pertinent to the question at hand.

Towards Testing Gould's Hypothesis

As Calabretta *et al.* (1998) show, models can be useful to investigate problems in biology. In order to do this, one must isolate the problem in question. I believe that this could be done in the case of Gould's hypothesis about progress in evolution. The question to be answered is as follows:

> In a system where there is no evolutionary pressure in terms of size, a system where size does not significantly affect survival, will organisms continue to grow in size, developing the tail that Gould (1996) describes, while the mode of life remains next to the wall of minimum complexity?

Tierra is not such a system, because large size is maladaptive. However, if such a model were developed, and the answer to the above question was yes, the conditional first portion of Gould's claim would be largely confirmed. The only remaining problem would be to determine whether size has a positive adaptive value in the biosphere. However, if the species at the end of the tail do not fall into a direct lineage (Gould 1996: 171-2), it would seem that this model does approximate our biosphere. Furthermore, once such a model was developed, its taxonomic lineage could be compared to that of the biosphere through a taxonomic analysis. McShea (1994, 1996) provides such an analysis of portions of the fossil record based on large-scale evolutionary trends. If the taxonomic lineage of produced by this next generation artificial life model was statically similar to that of the biosphere, as evaluated by McShea, this would add significant support to the later portion of Gould's claim.

Conclusion

Although there is no existing model which allows us to evaluate Gould's hypothesis, it might be possible to develop one. If Gould is correct, this model would also provide us with the first model of evolution which evolves complexity, providing a basis for a great deal of other research in artificial life and evolutionary biology. However, because of the time the biosphere required to evolve the complexity it currently exhibits, it is unlikely that even a correct model will provide the complexity necessary to study other problems in artificial life and evolutionary biology in the near future. Nevertheless, this appears to be a fruitful direction for both artificial life and evolutionary biology to begin to move in.

Acknowledgments

My research is supported by grants from the Dean of Liberal Arts, Plan II Honors Program, and Department of Philosophy at the University of Texas at Austin.

References

Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E. (1998) (eds.), *Artificial Life VI* (UCLA, Proceedings; Cambridge, Mass.: MIT Press).

———, Brown, C. T. (1994), 'Evolutionary learning in the 2D artificial life system "avida"', in Brooks, R. and Maes, P. (1994), 377-381.

Bedau, M. A., Synder, E., Packard, N. H. (1998), 'A Classification of Long-Term Evolutionary Dynamics', in Adami *et al.* (1998), 228-237.

——— Synder, E., Brown, C. T., and Packard, N. H. (1997a), 'A Comparison of Evolutionary Activity in Artificial Evolving Systems and in the Biosphere', in Husbands and Harvey (1997), 293-301.

——— and Packard, N. H. (1992), 'Measurement of Evolutionary Activity, Teleology, and Life', in Langton *et al.* (1992), 431-61.

——— and Brown, C. T. (1997b), 'Visualizing Evolutionary Activity of Genotypes', *Adaptive Behavior*. Forthcoming.

Brooks, R. and Maes, P. (1994), *Artificial Life IV* (Cambridge, Mass.: MIT Press).

Calabretta, R., Nolfi, S., Parisi, D., Wagner, G. P. (1998), 'A Case Study of the Evolution of Modularity: Towards a Bridge between Evolutionary Biology, Artificial Life, Neuro- and Cognitive Science', in Adami *et al.* (1998), 275-284.

Cowen, G. A., Pines, D., Meltzer, D. (1994) (eds.), *Complexity: Metaphors, Models and Reality* (Redwood City, Calif.: Addison-Wesley).

Dennett, D. C. (1995), 'Tinker to Evers to Chance: The Burgess Shale Double-Play Mystery', *Darwin's Dangerous Idea* (New York: Simon & Schuster), 299-312.

Domjan, P (1999), 'Are Romance Novels Really Alive? A Discussion of the Supple Adaptation View of Life', in Floreano *et al.* (1999), 21-25.

Floreano, D. *et al.* (1999) (eds.), *Advances in Artificial Life* (Berlin: Springer-Verlag).

Gould, S. J. (1992), *Full House: The spread of excellence from Plato to Darwin* (New York: Harmony Books).

——— and Lewontin, R. C. (1979), 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme', *Proceedings of the Royal Society* B, 205, 581-598.

———. (1989), *Wonderful Life: The Burgess Shale and the Nature of History* (New York: W.W. Norton).

Holland, J. H. (1994), 'Echoing emergence: objectives, rough definitions, and speculations for echo-class models', in Cowen, G. A. *et al.* (1994).

———. (1995), *Hidden Order: How Adaptation Builds Complexity* (Helix Books).

Husbands, P., and Harvey, I. (1997) (eds.), *Proceedings of the Fourth European Conference on Artificial Life* (Cambridge, Mass.: MIT Press).

Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S. (1992) (eds.), *Artificial Life II* (Santa Fe Institute Studies in the Science of Complexity, Proceedings, 10; Redwood City, Calif.: Addison-Wesley).

McShea, D. W. (1994), 'Mechanisms of large-scale evolutionary trends', *Evolution* 48: 1747-1763.

———. (1996), 'Metazoan complexity and evolution: is there a trend?', *Evolution* 50: 477-492.

Ray, T. S. (1992), 'An Approach to the Synthesis of Life', in Langton *et al.* (1992), 371 408.

——— and Hart, J. (1998), 'Evolution of Differentiated Multi-threaded Digital Organisms', in Adami *et al.* (1998), 295-304.

Reynolds, C. W. (1987), 'Flocks, Herds, and Schools: A Distributed Behavioral Model', Proceedings of SIGGRAPH '87, *Computer Graphics* V 21/4: 25-34.

Taylor, T. (1997), "The COSMOS artificial life system', Working Paper 263, Department of Artificial Intelligence, University of Edinburgh. Available from http://www.dai.ed.ac.uk/daidb/people/homes/timt/papers/.

——— and Hallam, J. (1998), 'Replaying the Tape: An Investigation into the Role of Contingency in Evolution', in Adami *et al.* (1998), 256-265.