

Using Bottom-Up Models to Investigate the Evolution of Life: Steps Towards an Improved Methodology

Tim Taylor

Department of Artificial Intelligence, University of Edinburgh
5 Forrest Hill, Edinburgh EH1 2QL, U.K.
timt@dai.ed.ac.uk

1 Introduction

Perhaps one of the few features shared by most artificial life approaches is that a phenomenon observed in biological life is studied by constructing a *bottom-up* model, in which a number of low-level components and interactions are explicitly encoded, and one or more higher-level phenomena are expected to emerge. While this is a perfectly valid approach, one has to be careful about how the model is constructed if it is to bring any scientific insight to bear on the phenomenon in question. Too many (but not all) studies of artificial life (my own included) have adopted a sloppy approach in the past, and this has meant that the field of artificial life has not contributed as much as it might have done to broader areas of scientific knowledge.

This paper highlights some areas of general methodology which should be carefully considered when designing a bottom-up simulation for scientific experimentation, and also suggests some considerations that are specifically relevant to A-life models designed to investigate the evolution of life. This is not, of course, the first time that concerns have been raised about the methodology of A-life, and very little is said here that has not been said before (see, for example, [4], [1], [5]). However, I believe that much current A-life work still suffers from poor methodology, and that it is therefore important to stress such issues at every available opportunity.

2 General Considerations

2.1 Explicit Assumptions and Predictions

The bottom-up approach to studying high-level phenomena is of scientific value only to the extent that the investigator has (a) made explicit exactly *what* high-level phenomenon he/she is trying to explain or investigate, and (b) explicitly enumerated a list of low-level phenomena (components and interactions) that he/she believes are necessary and sufficient to explain the high-level phenomenon.

If the assumptions and predictions have not been made explicit, then the output of the model will be able to tell us little of scientific value, no matter how surprising, interesting, or ‘life-like’ it may be. Although this point is fairly fundamental to scientific methodology in general and may seem so obvious that it is unnecessary to point it out, a quick skim through any A-life conference proceedings should be enough to demonstrate that these basic considerations are (very) often overlooked.

The number of assumptions (the low-level phenomena) that go into the model does not have to be large (e.g. they may be, say, (1) inert entities capable of (2) reproduction and (3) heritable

variation), and the high-level phenomenon under investigation does not have to be small (e.g. it could be, say, the evolution of life [but see Section 3]). However, the more explicit assumptions there are, and the more restricted the phenomenon to be explained, the more likely the model is to produce the desired results.

2.2 Minimal Models

Having devised an explicit list of low-level phenomena as a tentative reductive explanation for a specific high-level phenomenon, a model should be constructed that encapsulates these low-level phenomena *and nothing else*. In other words, it should be a minimal model. The model can then be run to see if it produces the expected results.

In practice, one generally has a choice of representations and algorithms that could be used to capture the low-level phenomena, and it may prove hard to be sure that no extra assumptions have crept into the model in the course of implementing it as a computer program (or as any other physical realization). However, as the list of low-level phenomena is explicit, the final implementation is open to testing, criticism and possible revision by others. David Marr essentially made the same point in his discussion of the three levels at which information-processing systems should be understood; he suggested that fields such as Artificial Intelligence were for too long hampered by a failure to recognize the theoretical distinction between *what* a system does (the ‘computational theory’ level), and *how* it does it (the ‘representation and algorithm’ and ‘hardware implementation’ levels) [3] (pp.19–29).

With the above in mind, once the model has been implemented, then if the expected results *are* observed, the model has demonstrated that the given assumptions are *sufficient* to explain the high-level phenomenon. To test whether all of the assumptions are *necessary*, further tests may be carried out in which assumptions are removed or relaxed one by one.

On the other hand, if the expected results are *not* observed, then the model has demonstrated that the assumptions are not sufficient. The model can then be revised by changing existing assumptions, or adding new ones.

Both cases can tell us something about the subject we are investigating, as we are always clear exactly *what* it is that we are trying to explain, and *how* we are trying to explain it. (It is much harder to conclude that assumptions are necessary to explain a given behaviour than it is to prove they are sufficient—indeed, we can *never* know that the behaviour may not also be achievable by completely different means. However, this problem is not specific to A-life, but is true of all science. All we can do is put forward our explanation as a possible model of the real world, and choose to accept the model that performs better (by some metric) than its competitors as our current ‘best guess’ on the matter [5].)

3 Specific Considerations for Models of the Evolution Of Life

3.1 The Low-Level Phenomena That Must Be Made Explicit

Darwinian (or, indeed, Lamarckian) evolution is a process of *change*. It tells us something about the *trajectory* of reproducing entities through their space of possible forms, and explains how reproducing entities become adapted to their environment. However, it *assumes* the existence of reproducing entities to begin with, and does not specify what sort of entities they should be, other than that they must be able to reproduce. Similarly, it does not specify that any particular *sort* of environment is necessary—evolution is a very general phenomenon.

A model in which a population of integers reproduce with occasional mutation, and differential survival based, perhaps, upon how large the integer is, will exhibit evolution, but it will never produce anything more than just integers. To take a more familiar example, genetic algorithms (see, e.g., [2]) satisfy the basic requirements for the evolution of the individual ‘chromosomes’, but all that is generally evolving is the encoded solution to some predetermined problem. Thus it is clear that if we are interested in modelling the evolution of *life*, we must (a) have a clear idea of what sorts of functions or roles a reproducing entity must fulfil if we are to consider it alive, i.e. a definition of life (this does not, of course, have to be universally agreed upon, but it *does* have to be explicitly stated), and (b) include in our model explicit components and interactions not only to allow for an (open-ended) evolutionary process to emerge, but also to allow for the existence of entities that fulfil any other functions or roles that we have specified as necessary for life.

In other words, evolution is not sufficient to explain life; we also require a theory of living organisation, and of the sorts of worlds which are able to support the emergence and evolution of such organisations. We must incorporate all of these into any A-life model designed to investigate the emergence and evolution of life.

3.2 A Definition of Life

If we are to build models to investigate or explain the emergence and evolution of life, we therefore need an explicit definition of life. That is, we need to be clear about exactly *what* we are trying to explain (as pointed out in Section 2). A number of definitions may be found in the literature, e.g. Maturana and Varela’s notion of autopoiesis (see, e.g., [6]). It is emphasised that any definition adopted does not have to be universally agreed upon (although it would obviously be desirable if it were widely accepted), but it does have to be explicitly stated if we are hoping that the model will be able to tell us anything of scientific value about the evolution of life (rather than just evolution in general). Many existing A-life models that claim to have been designed to investigate the evolution of life are accompanied with no explicit statement of exactly what they are trying to demonstrate (and often also have no explicit list of the assumptions and theory involved in the construction of the model), so it is impossible to judge whether they have succeeded or failed and they can therefore tell us little of any interest.

3.3 Ecological Considerations

An aspect of biological life that seems to be particularly overlooked in many A-life models is that biological organisms are dissipative organisations that participate in exchanges of energy and matter with their biotic and abiotic environment. Perhaps more accurately, most A-life models tend to focus upon *either* evolutionary *or* ecological aspects of life, but few consider both equally. Any acceptable definition of living organisation is likely to concentrate on an organism’s capability of self-maintenance in the face of environmental perturbations (caused by biotic or abiotic factors). It therefore seems probable that any A-life model of the sort we are considering will have to make explicit assumptions about the sorts of ecological interactions that are necessary and sufficient, as well as what sorts of organisations should be classified as living, and by what mechanisms they may evolve. A model that contains all of these things, and is capable of supporting a large population of organisms, may turn out to be prohibitively large for most computers at present (but maybe not). However, these are the design criteria we should move towards if models of this sort are to make significant contributions to the more general study of living systems.

4 Summary

It has been suggested in this paper that too much of the current research being done in A-life still suffers from a poor methodological approach. Specific recommendations are given to improve the situation; these basically boil down to having an explicit high-level natural phenomenon to be explained, and proposing an explicit list of low-level phenomena as a tentative reductive explanation. In Section 3 specific weaknesses are identified in the particular area of current A-life research into the evolution of life. It is suggested that such studies require a definition of living organisation, together with consideration for the sorts of environment which can support such organisation, as well as a mechanism for open-ended evolution.

Acknowledgements

Thanks to John Hallam, John Demiris and the other members of the Mobile Robots Group in the Department of Artificial Intelligence at Edinburgh University for their interesting discussions and suggestions on this topic.

References

- [1] Ezequiel A. Di Paolo. Some false starts in the construction of a research methodology for artificial life. In J. Noble and S. Parsowith, editors, *The 9th White House Papers: Graduate Research in the Cognitive and Computing Sciences at Sussex*. School of Cognitive and Computing Sciences, University of Sussex, 1996.
- [2] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
- [3] David Marr. *Vision*. W.H. Freeman, New York, 1982.
- [4] Geoffrey F. Miller. Artificial life as theoretical biology: How to do real science with computer simulation. Cognitive Science Research Paper 378, School of Cognitive and Computing Sciences, University of Sussex, 1995.
- [5] Jason Noble. The scientific status of artificial life. Poster presented at the Fourth European Conference on Artificial Life (ECAL97), Brighton, U.K., 1997.
- [6] Francisco J. Varela. *Principles of Biological Autonomy*. North Holland, Amsterdam, 1979.