Complexity Growth in Artificial Life

Russell K. Standish
High Performance Computing Support Unit
The University of New South Wales
Sydney, 2052, Australia
R.Standish@unsw.edu.au
http://parallel.acsu.unsw.edu.au/rks

Abstract

One of the striking properties of the evolution of life on earth is the almost monotonic growth in complexity from the primordial ooze to the space age. Yet it is equally striking that artificial life systems currently under study, such as *Tierra* or *Avida*, if anything, show the opposite trend toward simpler and more highly optimised creatures. This paper looks at where complexity might arise in complex systems, and proposes an artificial life experiment that might give rise to complexity growth.

Whose Complexity is it Anyway?

A major problem with examining the question of growth of complexity in evolution is that we do not have a clear handle on what we mean by complexity. Even when we can decide on a particular measure of complexity, do we then apply that to the most complex organism in any epoch, to an average[1] organism, or to the ecosystem as a whole. For example, in the current epoch, the average organism is a single cell bacterium, at least in terms of numbers, and the total biomass of bacteria is likely to outnumber the multicellular biomass by many orders of magnitude(7). It make a very real difference to which part of the ecosystem one applies the complexity measure. Gould(7) argues that the modal complexity[2] has not changed since the beginnings of life (corresponding to bacteria of minimal complexity). The maximally complex organism has become more complex over time by a purely random, undirected diffusion.

The few studies to date that have looked at these issues e.g. (5; 9) tend to focus on structural complexity, as this is the easiest to measure from the fossil record. One study(13) looked at trends in encephalization (braininess). It would probably be impossible to study the evolution of ecosystem structure, as there is no data in the fossil record for functional complexity, aside from extrapolation based on contemporary ecosystems. However, there is the real possibility that complexity evolution can be studied in artificial life ecosystems.

Then there is the issue of how we might measure complexity. Naïvely, we expect that complexity is somehow related to the *information content* of the object: perhaps it could be the size of the building plan, or the amount of work required to construct the system, or the number of parts and their interconnections. Clearly this introduces observer dependence. For example, the information content in an arrangement of objects clearly depends on whether one regards the positions of the objects as significant or not. Gell-Mann(6) gives a particularly clear discussion of the subtleties involved in developing a complexity measure. These essentially start off with the concept of Kolmogorov complexity, where the complexity of an object is related to the length of the smallest algorithm that generates it. The problem with this measure is that random sequences are incompressible, and so are of high Kolmogorov complexity. Gell-Mann introduces the concept of *effective complexity*, which is the length of schema that partially predicts the sequence. Sequences with high algorithmic information content have low effective complexity, because most of the behaviour is considered random by the observer. Perhaps the easiest way to picture this is to consider a hypothetical neural network that is being trained to recognise the complex objects. Bennett(3; 4) introduces the concept of *crypticity*, which is the length of time the neural network takes to settle down to its trained configuration, and *depth*, which is the length of time it takes for the neural network to recognise the object. Depth is roughly equivalent to effective complexity, and the depth of an object is relative to the crypticity of the observer.

Clearly complexity is a multidimensional quantity, as it is dependent on the framework of the observer.

---

[1] or median or modal

[2] the maximum point of the complexity distribution

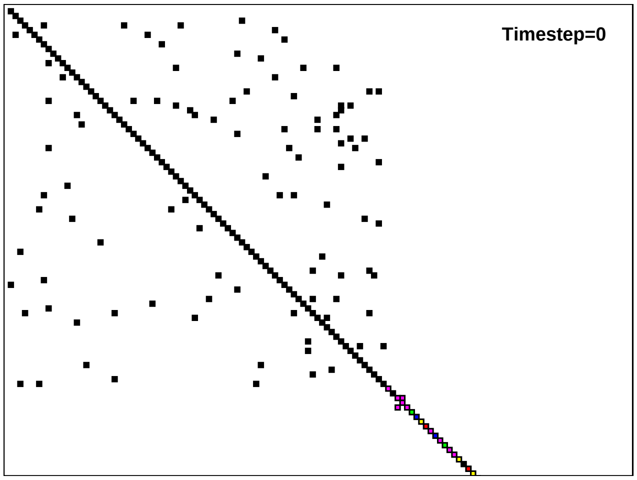

Figure 1: Ecolab Ecosystem Connectivity Plot

McShea(9) categorises four different types of complexity that relates to analysing the fossil record for evolutionary trends in complexity. He also examines evidence for trends in metazoa, or multicellular organisms, but concludes that whilst some trends are apparent, particularly in the early Phanerozoic, the data is too sparse and ambiguous to support or reject an increase in any type of complexity.

How do all these measures of complexity relate to growth of complexity in artificial life? The most pragmatic response is to assert that a system that "looks and smells like" it is increasing in complexity, most likely is. One can apply post-facto various measures of complexity to further refine exactly what is increasing, and to ascertain what features are the most important in that artificial system.

Figures 1–4 show the ecosystem connectivity plot for a typical run of the author's Ecolab system(18; 14; 17). The full sequence may be viewed as a GIF89 movie on http://parallel.acsu.unsw.edu.au/rks/ecolab-snaps/complex-anim.gif. Each species has a row and a column, and a coloured square corresponds to a connection between the species that own the row and the column that contain the square. The individual species are sorted and coloured to indicate the connected subecologies. The system is started from a random configuration, which is highly connected, but unstable. It rapidly collapses to a number of isolated simple ecosystems. Over time, the ecosystem complexity grows as more species are added to the system, and isolated ecosystems are connected together. Note that I haven't applied any quantitative complexity measures to the ecosystem, but it is clear that most of the relevant proposed complexity measures will show an increase.

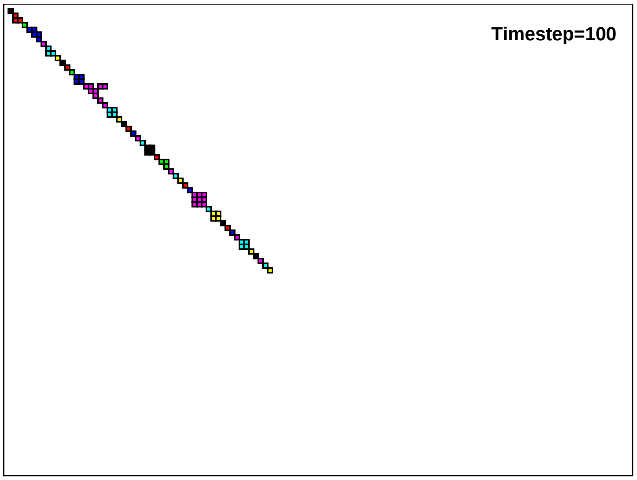

Figure 2: Ecolab Ecosystem Connectivity Plot

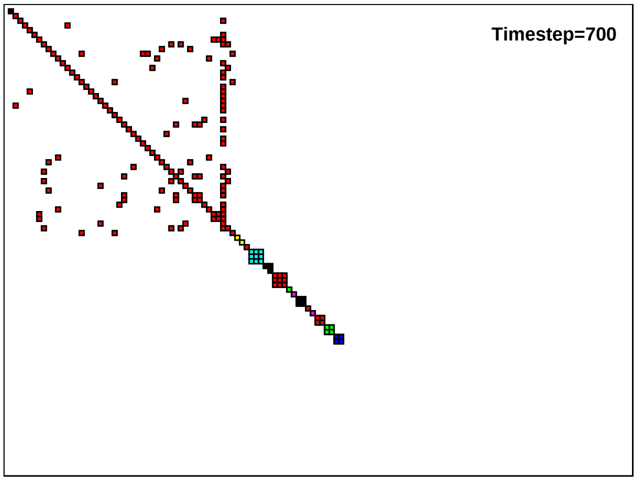

Figure 3: Ecolab Ecosystem Connectivity Plot

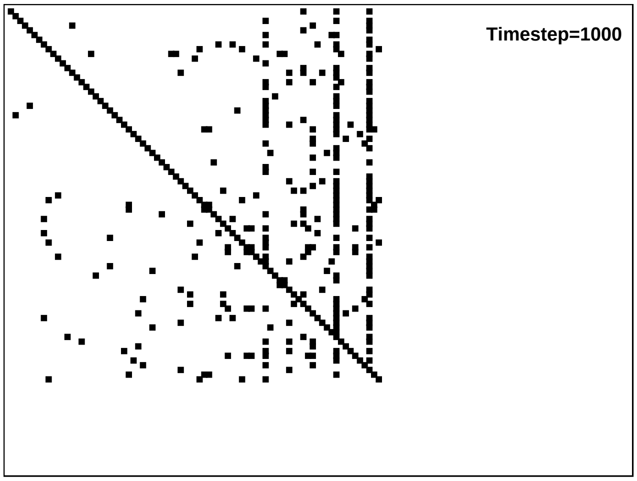

Figure 4: Ecolab Ecosystem Connectivity Plot

Growth of Complexity in Artificial Life

What lessons can we draw from the Ecolab observation? One of the first is the importance of networks, which Green(8) observed to be a recurring feature of complex systems. The second is to move from the organismal layer to the ecosystem layer. There are certainly examples from biology where more complex systems appear to have arisen from a symbiosis of two or more simpler systems; the mitochondria within eukaryotic cells for example are thought to have evolved from prokaryotic cells living symbiotically with the eukaryotic cell.

Taking this back to an artificial life system like *Tierra*(10; 12; 11), connections between organisms arise through the template matching process. One could generate ecosystem connectivity plots like figures 1–4 to study the evolution of the ecosystem. One problem that needs to be addressed is the large amount of neutral variation between Tierran organisms(15) — you must categorise organisms by phenotypic behaviour, not genotypes as is usual in Tierra.

What is needed is some way of recognising when a group of Tierran organisms is behaving like a single organism. One approach is to add the ability for organisms to build cell membranes. These would take the form of additional NOP instructions that act as barriers to template matching searches. It is rather unlikely that this can be done meaningfully in the 1-D diffusion space[3] that is Tierra's soup. More likely, a 2-D space is required, such as the *Avida* system(1). Then we would look for self reproducing systems enclosed with a membrane. Clearly efficiency of the template matching operation is important to a successful system(16).

Another potential difficulty limiting growth of complexity is the homogeneity of resources within Tierra. Each organism is given a fixed time slice. Ray has performed some experiments with "artificial chlorophyll", but no results have been published on this to the author's awareness.

Evolutionary Physics and Chemistry

What has been discussed so far is a proposal to tinker with the instruction sets of existing artificial life systems. I want to end this paper with another perspective and a note of caution. It has generally been assumed that there are a plethora of artificial "physics and chemistries" that can support life. Usually, the ability to build a Turing machine in the alife instruction set is taken to be sufficient for life to evolve. Not so strongly advocated, but widely believed, is that growth in complexity will arise naturally given enough simulation time and space, once the system has sufficient richness to support a Turing machine.

However, the real world physics and chemistry is actually extremely fine tuned for life as we know it(2). There is a very real prospect that the Universe as a whole has undergone some evolutionary process to give rise to a physics that generates complexity(19). Perhaps we will not see genuine growth of complexity in alife until we can allow the instruction set itself to evolve.

---

[3] template searches can be considered as the program counter diffusing through the soup of instructions until a matching template is found

References

[1] Chris Adami and C. Titus Brown. Evolutionary learning in the 2d artificial life system "avida". In R. Brooks and P. Maes, editors, *Artificial Life IV*. MIT Press, 1994.

[2] J. D. Barrow and F. J. Tipler. *The Anthropic Cosmological Principle*. Clarendon, Oxford, 1986.

[3] Charles H. Bennett. Dissipation, information, computational complexity and the definition of organization. In David Pines, editor, *Emerging Syntheses in Science*, volume 1, page 215. Addison-Wesley, Reading, Mass., 1988.

[4] Charles H. Bennett. How to define complexity in physics and why. In Wojciech Zurek, editor, *Complexity, Entropy and the Physics of Information*, page 137. Addison-Wesley, Reading, Mass., 1990.

[5] J. T. Bonner. *The Evolution of Complexity*. Princeton UP, 1988.

[6] M. Gell-Mann. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. Freeman, 1994.

[7] Stephen Jay Gould. *Full House: The Spread of Excellence from Plato to Darwin*. Harmony, 1996.

[8] David G. Green. Towards a mathematics of complexity. In R. Stocker, H. Jelinek, B. Durnota, and T. Bossomeier, editors, *Complex Systems: From Local Interactions to Global Phenomena*, pages 98–105. IOS, Amsterdam, 1996. also *Complexity International*, vol. 3, http://www.csu.edu.au/ci.

[9] Daniel W. McShea. Metazoan complexity and evolution: Is there a trend? *Evolution*, 50:477–492, 1996.

[10] T. S. Ray. An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, page 371. Addison-Wesley, New York, 1991.

[11] T. S. Ray. Evolution, complexity, entropy and artificial reality. *Physica D*, 75:239–263, 1994.

[12] T. S. Ray. An evolutionary approach to synthetic biology, zen and the art of creating life. *Artificial Life*, 1:195, 1994.

[13] D. A. Russell. *Adv. Space Res.*, 3:95, 1983.

[14] R. K. Standish. Ecolab: Where to now? In R. Stocker, H. Jelinek, B. Durnota, and T. Bossomeier, editors, *Complex Systems: From Local Interaction to Global Phenomena*, pages 263–271. IOS, 1996. also *Complexity International*, vol. 3, http://www.csu.edu.au/ci.

[15] R. K. Standish. Embryology in tierra: A study of a genotype to phenotype map. *Complexity International*, 1997.

[16] R. K. Standish. On an efficient implementation of tierra. *Complexity International*, 4, 1997. http://www.csu.edu.au/ci.

[17] Russell K. Standish. Ecolab documentation. Available at http://parallel.acsu.unsw.edu.au/rks/ecolab.html.

[18] Russell K. Standish. Population models with random embryologies as a paradigm for evolution. In Russel J. Stonier and Xing Huo Yu, editors, *Complex Systems: Mechanism of Adaption.* IOS Press, Amsterdam, 1994. also *Complexity International*, vol. 2, http://www.csu.edu.au/ci.

[19] Walter Thirring. Do the laws of nature evolve? In Michael P. Murphy and Luke A. J. O'Neill, editors, *What is Life? The Next Fifty Years*, page 131. Cambridge UP, Cambridge, 1995.