

Requirements for Open-Ended Evolution in Natural and Artificial Systems

Tim Taylor

Department of Computer Science and York Centre for Complex Systems Analysis
University of York, United Kingdom
tim@tim-taylor.com

Abstract. Open-ended evolutionary dynamics remains an elusive goal for artificial evolutionary systems. Many ideas exist in the biological literature beyond the basic Darwinian requirements of variation, differential reproduction and inheritance. I argue that these ideas can be seen as aspects of five fundamental requirements for open-ended evolution: (1) robustly reproductive individuals, (2) a medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, (3) individuals capable of producing more complex offspring, (4) mutational pathways to other viable individuals, and (5) drive for continued evolution. I briefly discuss implications of this view for the design of artificial systems with greater evolutionary potential.

1 Introduction

If there is one lesson to be learned from the first 60 years of research into the evolution of digital organisms, it is that the classic Darwinian ingredients of *variation*, *differential reproduction* and *inheritance* are not, in themselves, sufficient for producing open-ended dynamics in which new, surprising, and sometimes more complex organisms continue to appear (Taylor et al., 2014).¹

Most evolutionary artificial life systems tend to rather quickly reach a quasi-stable state beyond which no qualitatively new innovations are seen to appear (Taylor, 2013). None has displayed dynamics which might be regarded as the holy grail of artificial life, where the continued evolution of novel forms is so interesting that the researcher is unwilling to press the “off” switch.

Various artificial life researchers have started to look at different aspects of the biological world for the missing ingredients. At the same time, our understanding of processes important in biological evolution has been greatly supplemented by new research in many areas, including epigenetics (Jablonka et al., 2005), non-coding regions of DNA (Comfort, 2015), neutral evolutionary networks (Wagner, 2011), facilitated variation (Gerhart and Kirschner, 2007), niche construction (Odling-Smee et al., 2003), and others.

¹ In this paper I will use an informal definition of open-ended evolution as “evolutionary dynamics in which new, surprising, and sometimes more complex organisms continue to appear.”

While these new research directions are exciting and promise new insights into the important ingredients of biological evolution, the underlying simplicity of the Darwinian picture of variation, differential reproduction and inheritance soon disappears in the panoply of new ideas. Of course, that might just be the price we have to pay for a deeper understanding of evolution—biology, unlike the physical sciences, is an historically contingent subject that can be fiercely resistant to Occam’s razor. On the other hand, it may be that these new ideas are all jigsaw pieces of a still simple, if somewhat expanded, framework in which we can understand biological evolution.

In the following section, I suggest that there are five fundamental requirements for a system to exhibit open-ended evolution. I show how the various ideas mentioned above fit into this picture, discuss how they relate to past work in artificial life, and suggest various directions that are indicated for future research.

2 Requirements

At a very general level, the following five features are necessary, and I claim sufficient, for a system to exhibit open-ended evolutionary dynamics:²

- Robustly reproductive individuals.
- A medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, at various levels of complexity.
- Individuals capable of producing more complex offspring.
- An evolutionary search space which typically offers mutational pathways from one viable individual to other viable (and potentially fitter) individuals.
- Drive for continued evolution.

Each of these features is discussed below.

2.1 Robustly reproductive individuals

The basic components of any evolutionary system are individual entities that can catalyse the production of (sometimes imperfect) copies of themselves. Successful individuals must be robust enough to survive in their environment until they have performed at least one reproduction. In order for an evolutionary process to be sustained, there must be at least some such robustly reproducing individuals in the population.³

While this may appear to be a fairly basic statement, the question of what are the appropriate ways to achieve robustness in artificial life systems has not often received the attention it deserves. Von Neumann’s self-reproducing automata

² This list is a refinement of the ideas presented in (Taylor, 2012).

³ Note that this requirement relates to the robustness of an *individual* to survive in its (potentially variable) environment. A separate consideration is the robustness of a *population* of individuals to cope with changing environments over evolutionary timescales; such population robustness is addressed in Section 2.4.

(von Neumann, 1966), and other systems of self-reproduction in 2D cellular automata, are generally not robust: they do not engage in self-maintenance and self-repair, and are susceptible to disruptive perturbations from neighbouring individuals. Hence, while these systems might possess some desirable theoretical evolutionary capacity (see Section 2.3), in practice they are evolutionary non-starters.

Digital organism systems such as Tierra and Avida hard-wire robustness into the system by not granting individuals write-access to other parts of memory (except in the special case where some new memory has been allocated for reproduction). This was a critical design decision that allowed prolonged evolution to happen in these systems, in contrast to predecessors such as Core War, where individuals could overwrite each other with no such restrictions (Ray, 1991). However, by hard-wiring write protection into the system, programs in Tierra and Avida become relatively isolated from each other, with consequences for what kinds of interactions are possible.⁴

Biological organisms need to actively maintain their organisation against the disruptive pull of the second law of thermodynamics. Concerns of entropy increase are not immediately applicable to digital organisms, unless entropy is intentionally built into the digital physics of the system.⁵ If entropy *was* built into an artificial life system, it would mean that the digital organisms would have to concern themselves with self-maintenance, and that most structures would naturally decay without the need for arbitrary mechanisms like reaper queues. This would entail the organisation of digital organisms more closely corresponding to the characterisation of living organisation as self-building, self-maintaining and self-reproducing systems, e.g. (Varela et al., 1974; Rosen, 1991; Gánti et al., 2003).

An significant open question for artificial life research is understanding the importance of topics such as entropy and self-maintenance for open-ended evolution.

2.2 A medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, at various levels of complexity

A clear requirement for open-ended evolution is that many different types of organism must be conceivable within the system. The medium in which the evolutionary process is unfolding must allow the possibility of a practically unlimited diversity of organism organisations, processes and interactions.

Much previous work within artificial life has concentrated on the ability of organisms to evolve complex computational and information processing capabilities, such as the ability of digital organisms in Avida to solve logic functions

⁴ This therefore also has consequences for the degree of drive for continued evolution (see Section 2.5).

⁵ Although note that in some physically-inspired models of computation such as *conservative logic*, there are more clearly defined analogies of heat dissipation and entropy (Fredkin and Toffoli, 1982).

(Lenski et al., 2003) or the evolution of complex neural network-driven behaviour in systems such as Polyworld (Yaeger et al., 2010) and Geb (Channon, 2006).

However, it is restrictive to only consider the evolution of information processing capabilities. Some of the most remarkable events in biological evolutionary history have involved the evolution of new ways of interacting with the environment via new sensors and effectors. The geochemical-physical medium in which biological evolution unfolds offers an enormously rich source of complex dynamics, across many different modalities of phenomena, that may potentially be exploited by organisms to promote their survival and reproduction.

The need for complex environments for the production of interesting evolution in artificial life systems has been recognised right back to the earliest work in the area. Barricelli (1963) spoke in terms of adding “toy bricks” to the environment to allow his digital organisms to evolve interesting behaviours.

In addition, the major transitions in evolution identified by Maynard Smith and Szathmáry (1995) involve changes in the organisation of individuals over evolutionary time. Hence, open-ended artificial life systems should allow the organisation of individual organisms to evolve as well.

Many issues arise when designing complex virtual environments in which organisms can evolve to access and exploit that complexity for their own ends. These include questions such as whether the medium should have “messy” processes with side effects, to allow for the serendipity often apparent in biological evolution, and “matter of degree” rather than “all or nothing” processes to allow for gradual evolution—described by Dennett (2013) as “sorta” evolution. Further issues concern the origin of signs and signals, i.e. biosemiosis (Hoffmeyer, 2007), and the representational relationship between organisms and environment such that aspects such as new sensors and effectors can evolve without being “programmed in” by the designer (Taylor, 2004).

2.3 Individuals capable of producing more complex offspring

Beyond having a medium in which a wide variety of organism designs could possibly exist, in order for complex adaptations to *evolve* from simple progenitors, it must be *possible* for an individual (or multiple individuals) to produce offspring that are more complicated than their parent(s).

There are (at least) two ways in which this may occur:

- A single individual is capable of producing an offspring of greater complexity than itself.
- Two or more individuals are jointly capable of producing an offspring of greater complexity than any one of its parents.

The first solution is exactly the issue addressed by von Neumann (1966) in his *Theory of Self-Reproducing Automata*. The fundamental requirement identified by von Neumann is that the inherited information-bearing structures must be involved in two distinct processes: (1) they are *interpreted* by the phenotype’s

machinery as instructions to guide the construction of an individual, and (2) they are copied *uninterpreted* from parent to offspring.

Seen in this general light, we can say that von Neumann’s requirements are satisfied by biological cells (in 3D), by his proposed self-reproducing cellular automata (in 2D), and by digital organisms such as those in Tierra (in 1D). Note, however, that in the case of Tierra, the interpretation machinery is hard-coded into an organism’s “virtual CPU” is therefore not evolvable. In addition, it is also desirable to allow for the evolution of other aspects of the evolutionary process itself, such as allowing new forms of genetic transmission, evolution of the organisation of the genome, evolution of mutation rates, etc. (Hindré et al., 2012). Hence, issues of explicit versus implicit encoded, embeddedness in the medium, etc., are also important concerns here (Taylor, 2013).

Biological examples of the second solution include horizontal gene transfer (HGT) and symbiogenesis. These processes are much less well explored in the artificial life literature, despite their significance in biological evolution and the fact that they provide a feasible complementary (or alternative) route to increased complexity.

2.4 Mutational pathways to other viable individuals

For an open-ended evolutionary process, it is insufficient for individuals to have the theoretical capacity for producing more complicated offspring. The fitness landscape of the system must be such that there are often viable mutational pathways open to individuals, leading to different individuals that are of roughly the same fitness, or of higher fitness, than their parents. That is, there must often be the opportunity for adaptive, or at least neutral, evolution. Otherwise, the evolutionary process will often get stuck in local optima (dead ends) beyond which no further change is possible.

While this has been understood for a long time—e.g. Rensch (1947) discussed the need for “improvements allowing further improvement”—the task of understanding the requirements for a fitness landscape to have this property is now a very active area of research.

A wide variety of work can be seen as contributing to this topic, including Wagner (2011)’s work on evolutionary innovations and neutral networks, a wide range of work on the evolution of evolution, e.g. (Hindré et al., 2012), evolvable genotype–phenotype mappings, e.g. (Gerhart and Kirschner, 2007; Wagner and Altenberg, 1996; Wills, 2014), and major transitions, e.g. (Maynard Smith and Szathmáry, 1995). Also relevant is work on understanding how complex structures can evolve from simpler components in modular, hierarchical and nearly-decomposable systems, e.g. (Simon, 1962; Watson, 2006; Calcott, 2008), and related work on semiosis in the origin of modular and loosely coupled systems, e.g. (Auletta et al., 2008). Conrad (1990) has also argued that redundant, loosely coupled systems can aid evolvability by creating “extradimensional by-passes” that prevent evolution from getting stuck in local optima.

The importance of exaptation—where an existing phenotypic structure becomes selected for a different function—is well recognised in biology (Gould and

Vrba, 1982; Whitacre, 2010). A challenge for achieving open-ended evolution in artificial systems is to work with structures that potentially have multi-functional properties, perhaps in different domains of interaction (Taylor, 2013).

All of the topics mentioned here (and many others too) provide us with ideas of how to create artificial evolutionary systems in which individuals have *room to move* as they explore the evolutionary landscape.

2.5 Drive for continued evolution

Even with the first four requirements in place, a continued drive is required to force the system to explore new states.

To create *any* drive in the system at all, selection pressure must exist. In general, this can be brought about by competition for some kind of limited resource (which may be matter, energy, space), or through environmental conditions, etc. Selection creates an adaptive landscape in which some variations of organism are favoured over others.

In order to achieve *continued* drive, the individuals must experience a *changing* adaptive landscape (Waddington, 1969). In biological populations this is brought about by other individuals being part of the ecological environment—those individuals are also evolving, and can alter the fitness landscape by direct interaction, e.g. co-evolution (Thompson, 1994), or indirectly through their actions, e.g. ecosystem engineering (Jones et al., 1997) and niche construction (Odling-Smee et al., 2003). Changes can also come about through (passive or active) diffusion of species to new environments (e.g. migration).

A changing adaptive landscape also has bearing on the available mutational pathways of the system (Section 2.4), as it will have consequences for what set of mutational neighbours of an individual are now viable.⁶

Some kinds of drive will push the system towards higher complexity (e.g. co-evolutionary arms races), whereas others will lead to change but not necessarily higher complexity. Whether the latter counts as “open-ended evolution” depends on one’s definition.

Many artificial evolutionary systems lack the rich *connectedness* of individuals brought about by ecological interactions, niche construction, etc., and this is no doubt part of the explanation of why open-endedness remains elusive in those systems (Taylor, 2013). In addition, if we wish to engineer artificial evolutionary systems aimed at solving particular problems, an important question is how to appropriately introduce some kind of extrinsic selection (e.g. fitness function), rather than relying on purely intrinsic natural selection, while still retaining an effective drive at each step of the process.

3 Conclusion

While it is clear that the requirements for open-ended evolution extend far beyond the basic Darwinian demands of variation, differential reproduction and

⁶ Indeed, this process has been observed experimentally in studies of virus–bacteria coevolution (Meyer et al., 2012; Thompson, 2012).

inheritance, I have argued that these additional ideas can be seen as aspects of five basic requirements: (1) robustly reproductive individuals, (2) a medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, (3) individuals capable of producing more complex offspring, (4) mutational pathways to other viable individuals, and (5) drive for continued evolution.

While advances in the evolutionary potential of artificial systems can come about by careful consideration of the details of all of the topics and theories discussed, it is useful to consider these five basic features as the foundation upon which open-ended evolution can be achieved.

References

- Auletta, G., Ellis, G., and Jaeger, L. (2008). Top-down causation by information control: from a philosophical problem to a scientific research programme. *Journal of The Royal Society Interface*, 5(27):1159–1172.
- Barricelli, N. A. (1963). Numerical testing of evolution theories. Part II. Preliminary tests of performance. Symbiogenesis and terrestrial life. *Acta Biotheoretica*, XVI(3/4):99–126.
- Calcott, B. (2008). The other cooperation problem: generating benefit. *Biology and Philosophy*, 23(2):179–203.
- Channon, A. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7(3):253–281.
- Comfort, N. (2015). Genetics: We are the 98%. *Nature*, 520(7549):615–616.
- Conrad, M. (1990). The geometry of evolution. *BioSystems*, 24(1):61–81.
- Dennett, D. (2013). Turing’s “strange inversion of reasoning”. In Cooper, S. B. and van Leeuwen, J., editors, *Alan Turing: His Work and Impact*, pages 569–573. Elsevier.
- Fredkin, E. and Toffoli, T. (1982). Conservative logic. *International Journal of Theoretical Physics*, 21(3-4):219–253.
- Gánti, T., Griesemer, J., and Szathmáry, E. (2003). *The principles of life*. Oxford University Press.
- Gerhart, J. and Kirschner, M. (2007). The theory of facilitated variation. *Proceedings of the National Academy of Sciences*, 104(suppl 1):8582–8589.
- Gould, S. J. and Vrba, E. S. (1982). Exaptation—a missing term in the science of form. *Paleobiology*, 8(1):4–15.
- Hindré, T., Knibbe, C., Beslon, G., and Schneider, D. (2012). New insights into bacterial adaptation through in vivo and in silico experimental evolution. *Nature Reviews Microbiology*, 10(5):352–365.
- Hoffmeyer, J. (2007). Semiotic scaffolding of living systems. In Barbieri, M., editor, *Introduction to Biosemiotics*, pages 149–166. Springer.
- Jablonka, E., Lamb, M. J., and Zeligowski, A. (2005). *Evolution in four dimensions: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT Press.
- Jones, C. G., Lawton, J. H., and Shachak, M. (1997). Positive and negative effects of organisms as physical ecosystem engineers. *Ecology*, 78(7):1946–1957.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423:139–144.

- Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. W.H. Freeman, Oxford.
- Meyer, J. R., Dobias, D. T., Weitz, J. S., Barrick, J. E., Quick, R. T., and Lenski, R. E. (2012). Repeatability and contingency in the evolution of a key innovation in phage lambda. *Science*, 335(6067):428–432.
- Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.
- Ray, T. S. (1991). An approach to the synthesis of life. In Langton, C., Taylor, C., Farmer, J., and Rasmussen, S., editors, *Proceedings of Artificial Life II*, pages 371–408. Addison-Wesley.
- Rensch, B. (1947). *Neuere Probleme der Abstammungslehre*. Enke. In German. English translation published as *Evolution above the Species Level*, 1959, Columbia University Press.
- Rosen, R. (1991). *Life itself: a comprehensive inquiry into the nature, origin, and fabrication of life*. Columbia University Press.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482.
- Taylor, T. (2004). Redrawing the boundary between organism and environment. In Pollack, J., Bedau, M., Husbands, P., Ikehami, T., and Watson, R., editors, *Proceedings of Artificial Life IX*. MIT Press.
- Taylor, T. (2012). Exploring the concept of open-ended evolution. In Adami, C., Bryson, D., Ofria, C., and Pennock, R., editors, *Proceedings of the 13th International Conference on Artificial Life*, pages 540–541. MIT Press.
- Taylor, T. (2013). Evolution in virtual worlds. In Grimshaw, M., editor, *The Oxford Handbook of Virtuality*, chapter 32. Oxford University Press.
- Taylor, T., Dorin, A., and Korb, K. (2014). Digital genesis: Computers, evolution and artificial life. Presented at the 7th Munich-Sydney-Tilburg Philosophy of Science Conference: Evolutionary Thinking, University of Sydney.
- Thompson, J. N. (1994). *The Coevolutionary Process*. University of Chicago Press.
- Thompson, J. N. (2012). The role of coevolution. *Science*, 335:410–411.
- Varela, F. G., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems*, 5(4):187–196.
- von Neumann, J. (1966). *The Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, Ill. Editor: A.W. Burks.
- Waddington, C. (1969). Paradigm for an evolutionary process. In Waddington, C., editor, *Towards a Theoretical Biology*, volume 2, pages 106–128. Edinburgh University Press.
- Wagner, A. (2011). *The origins of evolutionary innovations: a theory of transformative change in living systems*. Oxford University Press.
- Wagner, G. P. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976.
- Watson, R. A. (2006). *Compositional evolution: the impact of sex, symbiosis and modularity on the gradualist framework of evolution*. MIT Press.
- Whitacre, J. M. (2010). Degeneracy: a link between evolvability, robustness and complexity in biological systems. *Theoretical Biology and Medical Modelling*, 7(6):1–17.
- Wills, P. R. (2014). Genetic information, physical interpreters and thermodynamics; the material-informatic basis of biosemiosis. *Biosemitotics*, 7:141–165.
- Yaeger, L., Sporns, O., Williams, S., Shuai, X., and Dougherty, S. (2010). Evolutionary selection of network structure and function. In Fellermann, H., Dörr, M., Hanczy, M. M., Laursen, L. L., Maurer, S., Merkle, D., Monnard, P.-A., Støy, K., and Rasmussen, S., editors, *Proceedings of Artificial Life XII*, pages 313–320. MIT Press.