

## Evolution as a Random Walk on a High-Dimensional Manifold Defined by Physical Law: Implications for Open-Ended Artificial Life

Andrew L. Nelson<sup>1</sup> and Brenae L. Bailey<sup>2</sup>

<sup>1</sup>Androtics LLC, Tucson, Arizona 85733, USA

<sup>2</sup>Department of Mathematics, University of Arizona, Tucson, Arizona 85721, USA  
[alnelson@ieee.org](mailto:alnelson@ieee.org)

### Extended Abstract

This short article presents a discussion of the underlying conditions under which natural evolution of life occurs and how these natural conditions may be extremely difficult to implement in artificial life (ALife) systems. In particular, the Darwinian concept of adaptation via natural selection may not have a complete or functional macro-level description that could be used to build any evolutionary environment that is defined at the agent-environment interaction level.

Here we are specifically addressing open-ended evolution (Ruiz-Mirazo et al., 2004) in artificial systems (Standish, 2003; Nolfi, 2012; Mouret and Doncieux, 2012). One of the goals of ALife is to generate systems capable of sustained evolution of life-like complexity. Such systems, although artificial, could then be considered to produce real life of a kind (Pattee, 1987; Ray, 1993), as opposed to being just simulations of life.

For evolutionary computing applications aimed at solving particular problems, a well-defined goal that is separate from survival in and of itself can be formalized into a selection criterion and used to evolve solutions (Oduguwa et al., 2005). However, in the case of open-ended evolution, many researchers now accept that agent-level definitions of fitness are unsuitable to drive differential selection and replication (Lynch, 2007; McShea, 1991; Lehman and Stanley, 2011). This includes even the most unbiased and high-level implicit fitness criteria in which replication is seemingly made to be a direct result of agent interaction with the environment (see for example Yaeger, 1994). Interestingly, concerns about the adequacy of neo-Darwinian and Darwinian theory to fully describe the evolution of life have come from several ALife researchers, often after attempting to implement evolving systems (Mitchell and Forrest, 1994; Lehman and Stanley, 2011; Watson, 2012; Nolfi, 2012). Ray, for example, indicates that there is something “oddly self-referential” about evolution (Ray, 1993). Dawkins describes how his views of natural biological evolution changed after playing around with ALife simulations (Dawkins, 2003).

An underlying tenet of science is that all observable natural phenomena result from a fundamental set of physical laws and that physical law is essentially unchanging (Feynman, 1967; Zilsel et al., 2003). Such a set of laws, although not yet fully elucidated by physicists, is presumed to exist. If this were not the case, some fundamental cornerstones of science such as

repeatability of experiments, as well as a host of epistemological underpinnings, would not hold. A consequence of the existence of such a set of elemental physical laws is that fundamental driving forces producing change in natural evolution result from or reflect the topology of a *static* space defined solely by unchanging physical law (Ray, 1993). In this sense (and noting that physics is thought to have an intrinsic stochastic aspect), evolution can be described as a random walk on a static manifold, one of extremely high dimensionality. The only fundamental non-random “force” driving change in nature is imparted by the underlying topology of this static extremely low-level and high-dimensional landscape. Furthermore, this low-level view of the universe is not mediated by a replication cycle per se.

The discussion above implies that a system defined only in terms of a suitable set of elemental rules might in theory support open-ended evolution, and that our natural universe is an example of such a system. This raises the possibility that high-level representations (including the differential survival and replication paradigm upon which Darwinian evolution is based), while describing evolution sufficiently to generate simulations, might not fully functionally specify evolution to the degree needed to generate artificial realizations of evolution. (See Pattee (1987) for a discussion of the distinction between simulation and realization.) Below, we loosely summarize an argument that implies that high-level descriptions of complex systems are likely to be functionally incomplete.

When complex systems with a high level of granularity are converted to lower levels of resolution, information is usually lost, even if overall patterns are seemingly more evident (Katsoulakis and Trashorras, 2006). Hence, if the behavior of a complex system is fully described (but not over-specified) at one level, it is in fact not likely to be fully described at a reduced level of resolution. The implication is that macro-level traits in biological systems, being essentially extremely low-resolution views of matter/energy configurations, do not contain sufficient information to fully predict replication efficiency distributions (as generalizations of adaptive fitness landscapes (Wright, 1932) might have suggested).

Relating variation in macro-level traits to replication efficiency would then not fully define the underlying forces driving evolution, not even in theory. In this case, the

paradigm of natural selection would still be useful in a retrospective sense for summarizing some high-level relationships between observed phenotype and reproductive efficiency (Kauffman, 1993), but would not be sufficient to functionally define or drive open-ended evolution of such complexity (Lynch, 2007; McShea, 1991).

We believe this view has relevance to the long-term success of ALife. A guiding approach employed in much of ALife is that of setting up an environment in which agents defined at the macro level compete to survive and replicate. This approach to producing systems capable of supporting open-ended complexity may be fundamentally flawed, even when implemented without overt bias and with asynchronous local reproduction and careful attention to definitions of fitness as highlighted in Lichocki et al. (2012).

Although obtaining an explicit description of a given complex phenomenon solely in terms of elemental physical law may be intractable, we maintain that it is not theoretically impossible. Further, this may be the only level at which complex biological phenomena are completely causally described.

A system defined by elemental physical laws is clearly sufficient to produce open-ended evolution, as this describes life in our own universe. However, in order to generate an artificial system capable of the open-ended evolution of complex agents, it may be not merely sufficient but necessary to define environments in terms of elemental rules (Ray, 1993). Agents in such a system must either be constructed using only these rules, or perhaps arise through abiogenesis, as natural life did. If complex self-replicators and their environment are constructed from a single set of consistent rules, the system could be considered to contain *endogenous* ALife. Currently such systems remain beyond the state of the art (Nelson, 2013), but recent work in artificial chemistry and soft ALife have made considerable advances in this direction (Joachimczak et al., 2012; Fontana, 2010).

Many questions remain. For example, at what level must in silico ALife environments be specified? Can any system defined at the macro/agent level be considered to be free of implicit fitness functions? Does the increase in complexity observed in, e.g., vertebrate evolution represent a general aspect of possible life, or is it just an artifact of life on Earth?

To summarize, in this short paper we have argued that macro-level concepts of natural selection cannot be used to define systems capable of supporting the open-ended evolution of complex life-like self-replicators. Further, the only fully explanatory driving force behind the evolution of natural life is imparted by the topology of a static low-level landscape defined by unchanging physical law. Higher-level descriptions that include differential survival explicitly linked to replication cycles are only adequate to generate simulations of evolution, not realizations of evolution in which complex agents actually arise.

## References

Dawkins, R. (2003). The evolution of evolvability. In *On Growth, Form and Computers*, pages 239–255.

Feynman, R. (1967). The character of physical law. Vol. 66. MIT Press.

Fontana, A. (2010). Devo co-evolution of shape and metabolism for an artificial organ. In *Artificial Life XII*, pages 19–23.

Joachimczak, M., Kowaliw, T., Doursat, R., and Wrobel, B. (2012). Brainless bodies: controlling the development and behavior of multicellular animats by gene regulation and diffusive signals. In *Artificial Life 13*, pages 349–356.

Katsoulakis, M. A. and Trashorras, J. (2006). Information loss in coarse-graining of stochastic particle dynamics. *Journal of Statistical Physics*, 122(1):115–135.

Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.

Lehman, J. and Stanley, K. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19.2 (2011):189–223.

Lichocki, P., Keller, L., and Floreano, D. (2012). Differences in the concept of fitness between artificial evolution and natural selection. In *Artificial Life 13*, pages 530–531.

Lynch, M. (2007). The frailty of adaptive hypotheses for the origins of organismal complexity. *PNAS*, 104:8597–8604.

McShea, D. W. (1991). Complexity and evolution: what everybody knows. *Biology and Philosophy*, 6(3):303–324.

Mitchell, M. and Forrest, S. (1994). Genetic algorithms and artificial life. *Artificial Life*, 1(3):267–289.

Mouret, J-B. and Doncieux, S. (2012). Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evolutionary Computation*, 20(1):91–133.

Nelson, A. L. (2013). Artificial life and machine consciousness. In *2013 AAAI Fall Symposium Series: AAAI Technical Report FS-13-02*, pages 52–57.

Nolfi, S. (2012). Co-evolving predator and prey robots. *Adaptive Behavior*, 20(1):10–15.

Oduguwa, V., Tiwari, A., and Roy, R. (2005). Evolutionary computing in manufacturing industry: an overview of recent applications. *Applied Soft Computing*, 5(3):281–299.

Pattee, H. H. (1987). Simulations, realizations, and theories of life. In *Artificial Life*, pages 63–78.

Ray, T. S. (1993). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1(1):179–209.

Ruiz-Mirazo, K., Peretó, J., and Moreno, A. (2004). A universal definition of life: autonomy and open-ended evolution. *Origins of Life and Evolution of the Biosphere*, 34(3):323–346.

Standish, R. K. (2003). Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, 3(2):167–175.

Watson, R. A. (2012). Is evolution by natural selection the algorithm of biological evolution? In *Artificial Life 13*, pages 121–128.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress on Genetics*, 1(6):356–366.

Yaeger, L. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or PolyWorld: life in a new context. In *Artificial Life III*, pages 263–298.

Zilsel, E., Raven, D., Krohn, W., and Cohen, R. S. (2003). The genesis of the concept of physical law. In *The Social Origins of Modern Science*, Boston Studies in the Philosophy of Science 200:96–122, Springer Netherlands.