

Open-Ended Evolution and Open-Endedness: Editorial Introduction to the Open-Ended Evolution I Special Issue

Abstract Nature’s spectacular inventiveness, reflected in the enormous diversity of form and function displayed by the biosphere, is a feature of life that distinguishes *living* most strongly from *nonliving*. It is, therefore, not surprising that this aspect of life should become a central focus of artificial life. We have known since Darwin that the diversity is produced dynamically, through the process of evolution; this has led life’s creative productivity to be called *Open-Ended Evolution* (OEE) in the field. This article introduces the first of two special issues on current research on OEE and on the more general concept of open-endedness. Most of the papers presented in these special issues are elaborations of work presented at the Third Workshop on Open-Ended Evolution, held in Tokyo as part of the 2018 Conference on Artificial Life.

Norman Packard*

Protolife, Inc.
norman.packard@protolife.com

Mark A. Bedau
Reed College

Alastair Channon
Keele University

Takashi Ikegami
University of Tokyo

Steen Rasmussen
University of Southern Denmark
Santa Fe Institute

Kenneth Stanley
Uber AI Labs
University of Central Florida

Tim Taylor
Monash University

Keywords

Open-ended evolution, open-endedness,
novelty, innovation, complexity growth

The variety of organisms produced by biological evolution is staggering. Recent estimates based on scaling arguments estimate the number of microbial species alone to be $\sim 10^{12}$ [5]. In a very real sense, the entire biosphere has been produced by (or one might say “invented by” [9]) the process of biological evolution. This creative productivity is one of the most striking features of life, and so it has not surprisingly become a focus of artificial life research and is considered one of the grand challenges in the field [2]. Decades of research have explored evolutionary algorithms of various sorts, but so far a full understanding of the creative productivity of evolution (biological or non-biological) has remained out of reach, both in terms of understanding how it has worked to produce the biosphere, and understanding what principles might be used or instantiated outside of biology (the purview of artificial life) to achieve comparable levels of creative productivity.

Nature’s creative productivity is generally assumed to be ongoing, not only at work over the past 4 billion years, but also at work presently in the biosphere. As far as we know, the same physical, chemical, and biological mechanisms remain at work now as they have in the past, leading most scientists to believe that biological evolution is *open-ended*. This has historically led us to call the study of the creative production of ongoing novelty *open-ended evolution*. As this field of study has

* Corresponding author.

developed, it has become clear that processes having the property of ongoing creative productivity are not necessarily biological, chemical, or even evolutionary (at least in a Darwinian sense). This realization leads us to generalize the description of the endeavor to include *open-endedness* in any form.

These special issues on open-ended evolution and open-endedness are the product of an ongoing collaboration within the artificial life community. An important benchmark in this collaboration was the report from the first Workshop on Open Ended Evolution in York [10]. One of the report's main conclusions was that open-endedness is a variegated concept; there is not a simple single litmus for the phenomenon, but instead it comes in a variety of forms, and a given system might exhibit one form but not another. This produces many varieties of open-endedness. The different kinds of open-ended evolution listed in the report from York are these:

1. *Ongoing generation of adaptive novelty:*
 - (a) *Ongoing generation of new adaptations*
 - (b) *Ongoing generation of new kinds of entities*
 - (c) *Emergence of a dynamical hierarchy and major transitions*
 - (d) *Evolution of evolvability*
2. *Ongoing growth of complexity:*
 - (a) *Entity complexity*
 - (b) *Interaction complexity*

The main lesson to learn from this list is that open-ended evolution is not a monolithic topic; to forestall confusion we must distinguish different kinds of OEE and give them distinct names. We should expect the phrase “open-ended evolution” to elicit different ideas in different people, so we should be quite clear at all times about which kinds of open-endedness we are discussing. In this note we describe OEE using the York categories.

The latest and most innovative work on open-ended evolution is well-represented by these special issues on OEE. The papers discuss each of the kinds of OEE listed above, and the number of papers is so large that it spans the first two issues of this journal's current volume.

The six papers published in this first set focus mostly on Type 1 OEE (ongoing generation of adaptive novelty): different papers focus on different kinds, and many papers discuss more than one kind. Type 1a OEE (ongoing generation of new adaptations) is one focus of most of the six papers. Harrington and Pollack [4] and Moran and Pollack [6] both consider evolving populations of game players as entities that evolve new strategic adaptations. Rasmussen and Sibani [8] make a distinction between two different modes of evolution at work in (presumably) any version of open-ended evolution—optimization and expansion—and both modes refer (sometimes indirectly) to ongoing generation of new adaptations. The MODES toolbox of Dolson et al. [3] quantifies several aspects of evolution producing new adaptations (e.g., change potential, novelty potential), and is hence relevant to Type 1a OEE. Dolson et al.'s MODES toolbox also addresses Type 1b OEE (ongoing generation of new adaptations), which is also the focus of an empirical analysis of technology evolution by Bedau et al. [1]. Pattee and Sayama [7] discuss how open-endedness might itself be a product of evolution, which concerns Type 1d OEE (evolution of evolvability).

Type 2 OEE (ongoing growth of complexity) is discussed in the papers by Harrington and Pollack [4], who present a version of Type 2a OEE (ongoing increase in entity complexity), as quantified by memory length escalation in an evolving population of competing entities, and by Moran and Pollack [6], whose game-playing agents display an ongoing increase in strategy complexity (both Type 2a and Type 2b OEE). Ongoing growth of complexity (Type 2 OEE) is also treated by Dolson et al. [3], who propose metrics to characterize both growth of entity complexity (change potential, novelty potential, complexity potential) and interaction complexity (ecological potential).

The subsequent issue of this journal will contain the second set of new papers on OEE, and it will include an overview of all the papers in the special issues. It will also address whether and how to revise the list of kinds of open-ended evolution in the report from York [10]. The list from York is not chiseled in stone; rather, it is designed to grow and be modified as we learn more about open-endedness and the mechanisms behind it. An excellent source of inspiration for this kind of incremental improvement are the papers in this special issue.

References

1. Bedau, M. A., Gigliotti, N., Janssen, T., Kosik, A., Nambiar, A., & Packard, N. (2019). Open-ended technological innovation. *Artificial Life*, 25(1), 33–49.
2. Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., & Ray, T. S. (2000). Open problems in artificial life. *Artificial Life*, 6(4), 363–376.
3. Dolson, E. L., Vostinar, A. E., Wisner, M. J., & Ofria, C. (2019). The MODES toolbox: Measurements of open-ended dynamics in evolving systems. *Artificial Life*, 25(1), 50–73.
4. Harrington, K., & Pollack, J. (2019). Escalation of memory length in finite populations. *Artificial Life*, 25(1), 22–32.
5. Locey, K. J., & Lennon, J. T. (2016). Scaling laws predict global microbial diversity. *Proceedings of the National Academy of Sciences*, 113(21), 5970–5975.
6. Moran, N., & Pollack, J. (2019). Evolving complexity in prediction games. *Artificial Life*, 25(1), 74–91.
7. Pattee, H. H., & Sayama, H. (2019). Evolved open-endedness, not open-ended evolution. *Artificial Life*, 25(1), 4–8.
8. Rasmussen, S., & Sibani, P. (2019). Two modes of evolution: Optimization and expansion. *Artificial Life*, 25(1), 9–21.
9. Stanley, K., Lehman, J., & Soros, L. (2017). Open-endedness: The last grand challenge you've never heard of. <https://www.oreilly.com/ideas/open-endedness-the-last-grand-challenge-youve-never-heard-of>.
10. Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., McMullin, B., Packard, N., Rasmussen, S., Virgo, N., Agmon, E., Clark, E., McGregor, S., Ofria, C., Ropella, G., Spector, L., Stanley, K. O., Stanton, A., Timperley, C., Vostinar, A., & Wisner, M. (2016). Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life*, 22(3), 408–423.