

Necessary Conditions for Open-Ended Evolution Empirically Validated in Chromaria

L. B. Soros¹ and Kenneth O. Stanley²

¹Cross Labs, Cross Compass Ltd., Tokyo, Japan

²OpenAI, San Francisco, CA USA

lisa.soros@cross-compass.com

Introduction

The quest for open-ended evolution (OEE) in artificial life centers on constructing artificial evolutionary processes that make new discoveries indefinitely as evolution appears to do in nature. Though this phenomenon has been a longstanding topic of interest in the aife community, the field generally lacks consensus on its exact definition (Juric, 1994; Bedau et al., 1998; Channon, 2003, 2006; Maley, 1999). For example, OEE has been described as the continual production of either novel (Bedau et al., 1997; Standish, 2003; Lehman and Stanley, 2011; Nellis, 2012) or adaptive (Bedau et al., 1998) forms. Recognizing that these descriptions are not necessarily mutually exclusive, OEE researchers have moved towards a pluralism that admits different degrees of open-endedness (Taylor et al., 2016; Packard et al., 2019a,b). Nonetheless, the OEE community aims to replicate some dynamics of biological evolution (which is frequently interpreted as an effectively open-ended process) in the hope of creating *more* such open-ended processes.

The talk proposed in this abstract reports on new work completing a previous preliminary study by Soros and Stanley (Soros and Stanley, 2014), which introduced a set of four necessary conditions for OEE and also tested one of the conditions in an artificial life world called Chromaria (Figure 1) that was designed to empirically validate the four conditions. The full study is available in longer form as the dissertation of the first author (Soros, 2018). Now that it has been possible to test all four conditions, and not just one, this talk presents for the first time a full empirical study of a set of necessary conditions for open-endedness.

Methodology

The approach is to control for all four conditions proposed by Soros and Stanley (2014) and then evaluate the results using a diverse set of metrics (expanded from the preliminary experiments) over 20 runs of standard Chromaria and of each control setup through 500,000 individuals generated.

Condition 1: A rule should be enforced that individuals *must* meet some minimal criterion (MC) before they can reproduce, and that criterion must be nontrivial.




Figure 1: **Chromaria.** The first Chromarian is born at the center of the world and then must find an appropriate place to plant (satisfying a color matching function between the individual and its surrounding environment). Each successive Chromarian is then born within a fixed radius of its parent. The color-rich borders initially provide the only viable options, but more emerge as Chromarians continue to thrive in the environment (right). The world was intentionally designed to satisfy the four hypothesized necessary conditions for open-ended evolution described in this abstract and in prior publications (Soros and Stanley, 2014; Soros et al., 2016; Soros, 2018)

Control 1: All individuals are allowed to reproduce regardless of what they do in the world.

Condition 2: The evolution of new individuals should create novel opportunities for satisfying the MC.

Control 2: Individuals cannot see each other, which means that the evolution of new individuals has no impact on other individuals in the world. Note that this control was already demonstrated in short preliminary experiments by Soros and Stanley (2014).

Condition 3: Decisions about how and where individuals interact with the world should be made by the individuals themselves.

Control 3: Randomness is introduced into Chromarian interaction and reproduction via a two-step process. First, each new Chromarian is born in a random location within

a fixed radius of its parent. The new Chromarian is then given a random heading. This way, the individual's effective behavior (which determines whether or not an individual satisfies the MC) is determined entirely by a combination of chance and parent location.

Condition 4: The potential size and complexity of the individuals' phenotypes should be (in principle) unbounded.

Control 4: This condition is controlled for by setting an upper limit on the number of nodes and connections allowed in each individual's genome network (i.e. its CPPN), which in effect limits phenotypic complexity.

Metrics

The position advocated in this abstract, and supported by the "Barriers to Open-endedness" framework of Dolson et al. (2015), is that there is likely not just one key dimension of novelty or open-endedness; many factors contribute to a system's evolutionary dynamics. Thus the approach in this study is to paint a broad picture of evolution by collecting a variety of quantitative and qualitative data points that speak to the barriers described by Dolson et al. (2015).

Metric 1: Genome network size. The aim of this metric is to indicate the system's tendency to avoid the **complexity barrier**. At the genome level, complexity can be approximated by measuring the number of components in each individual's variable size genome. Once the genomes are decoded, the resulting phenotypes (each a morphology coupled with a behavioral controller) must be assessed independently. Genomic complexity is a good proxy for organismal complexity because the size of the genome represents a theoretical cap on the complexity of the organism, and genomes gradually increase in size through NEAT (the neuroevolution algorithm in Chromaria).

Metric 2: Median reproduction number of last new species' emergence. The frequency at which new individuals appear directly measures the system's tendency to avoid the **change barrier** and **novelty barrier** by showing that the population has not converged to a single fixed point. Such nonconvergence is guaranteed if new species continue to emerge in the system.

Metric 3: Median total species count. The total number of unique species that emerge during a run quantify how much of the morphological search space is explored during a run. While this metric does not necessarily map to a specific barrier to OEE, it quantifies the degree of generativity of a system.

Metric 4: Planting attempt success rate. This metric, considered in tandem with with the time of last species emergence and total species count, helps create a holistic quantification of the system's evolutionary dynamics. Specifically,




Figure 2: **Snapshots of a single standard Chromaria run taken at each quarter of the run.** The final world state is shown at the bottom. These snapshots depict a clear progression of world states that has not stagnated and shows potential for continued evolution.

it can potentially quantify how difficult it is to satisfy the MC throughout a run.

In addition to the four metrics noted above, visual snapshots illuminate environmental change and the ability of the system to overcome the **ecological barrier**.

Results

Figure 2 depicts a sample standard Chromaria run at evenly-spaced intervals. In general, evolution tends to meander through color space, with new niches emerging as new ways of satisfying the MC become possible over time.

One unexpected outcome in the standard runs is that connection count in both types of genome networks actually *decreases* over time, which might initially seem to suggest a corresponding decrease in Chromarian complexity. However, this result does not necessarily indicate a lack of open-endedness in Chromaria, or even that individuals are not themselves becoming more complex over time. In fact, complex behavioral patterns can be observed at the ends of runs. The fact that connection counts generally decrease suggests that the genetic encoding in Chromaria (i.e. CPPNs) is sufficiently expressive that capping genome size (Control 4) would not have an impact on the complexity of observed behaviors. Hence, Control 4 is not performed.



Figure 3: **Representative end state snapshot, Control 1 (there is no MC for reproduction)**. In all 20 runs, the world becomes chaotic and partially gray partway through the run, and visual change stops occurring entirely. This result indicates that no planting attempts are made at all. (In fact, near the end of the runs most individuals simply run into walls repeatedly.)

Condition/Control 1 (Figure 3): Controller genome connection count decreases and does so at a much sharper rate than in standard runs. Morphology genome connection count decreases in both the standard and control setups, but the decrease is more dramatic in the control. The last new species is found on average relatively early in control runs compared to standard runs ($p=6.29 \times 10^{-8}$)¹, indicating premature convergence. The total species count is higher for control runs ($p=5.52 \times 10^{-8}$), which makes sense; more of the search space can be traversed if all regions are viable (because there is no MC for reproduction). However, the last successful planting attempt occurs early in the run, indicating a functional collapse of evolution not seen in standard runs.

Condition/Control 2 (Figure 4): Changes in genome size obey the same general patterns as the standard runs, but the decrease in controller connection count is not as dramatic in this control. However, both node and connection counts increase for morphology-encoding genomes in these control runs. The last new species emerges significantly earlier in control runs than in standard runs ($p=6.29 \times 10^{-8}$). It is clear from observing the visual snapshots that it is difficult, if not impossible, for evolution to traverse morphology space beyond the colors present in the initial state of the world. The inability for evolution to access most areas of morphology space is again evidenced by the significantly lower number of total species found (7) compared to the standard world (16; $p=5.77 \times 10^{-8}$). This low total species count combined with the relatively early emergence of the last new species indicates that evolution exhausts all available opportunities for innovation early on and then fails to discover or take advantage of new ones. Planting attempt success rates are lower when individuals cannot interact with each other compared to in the standard world.

¹Controls are compared to the standard world using the Wilcoxon rank-sum test.



Figure 4: **Representative end state snapshot, Control 2 (agents cannot interact with each other)**. In all 20 runs, the system fails to explore new niches beyond the initial color space.




Figure 5: **Representative end state snapshot, Control 3 (Chromarians cannot choose their own behaviors)**. In all 20 runs, the world becomes dominated by dark colors and appears relatively homogeneous compared to the standard runs.

Condition/Control 3 (Figures 5,6): Controller genome size becomes minimal and controller-encoding networks become *completely disconnected* when controllers no longer serve the purpose of deciding behaviors. Morphological genomes become larger than in the standard runs. The last species emerges at median reproduction 238,000 when individuals cannot decide for themselves where and how to interact with the environment, compared to reproduction 330,750 in the standard world ($p=6.29 \times 10^{-8}$). The number of species found in this world (18) is significantly different but only slightly greater than in the standard world (16; $p=0.01017$). Considering this result in tandem with the time of last new species emergence indicates that a higher volume of species are found earlier in these runs compared to the standard world. However, this elevated degree of generativity is not sustainable. Planting rates are atypically high in the control case because every individual is forced to attempt to plant once its location has been (semi-randomly) determined.

Conclusion

The experiments described in this abstract validated three of the four conditions hypothesized to be necessary for open-ended evolution: that a nontrivial minimal criterion is necessary for selective reproduction, that new ways of satisfying this minimal criterion should become possible as new indi-




Figure 6: **Representative end state snapshots, standard Chromaria and Control 3.** End states of sample standard runs (left) are presented side by side with end states of sample Control 3 runs (right). The visual disparity between these two sets of end states underscores the relative homogeneity that is generated when individuals cannot make decisions about their own actions.

viduals evolve, and that individuals should control their own interactions with the environment.

However, the results also showed that Chromaria cannot offer empirical evidence to support the hypothesis that individual complexity must be able to increase without bound. Still, validating even one of the hypothesized necessary conditions is a useful step towards a comprehensive theory of OEE because of the lack of any systematic empirical study verifying hypothesized conditions. Furthermore, the knowledge gained from this study can provide concrete suggestions for increasing open-endedness in artificial evolutionary systems writ broadly.

Acknowledgements

This work was supported in part by the National Science Foundation under grant no. IIS-1421925. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Bedau, M. A., Snyder, E., and Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and the biosphere. In *Proceedings of the Fourth European Con-*

ference on Artificial Life, pages 125–134, Cambridge, MA, MIT Press.

Bedau, M. A., Snyder, E., and Packard, N. H. (1998). A classification of longterm evolutionary dynamics. In *Proc. of Artificial Life VI*, pages 189–198, Cambridge, MA, MIT Press.

Channon, A. (2003). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in geb as unbounded. In *Proc. of Artificial Life VIII*, pages 173–181, Cambridge, MA, MIT Press.

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.

Dolson, E., Vostinar, A., and Ofria, C. (2015). What’s holding artificial life back from open-ended evolution? *The Winnower*.

Juric, M. (1994). An anti-adaptationist approach to genetic algorithms. In *Proc. of First IEEE Conf. on Evolutionary Computation*, volume 2, pages 619–623. IEEE.

Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.

Maley, C. C. (1999). Four steps toward open-ended evolution. In *Proc. of the Genetic and Evolutionary Computation Conf. (GECCO '99)*, volume 2, pages 1336–1343. Morgan Kaufmann.

Nellis, A. (2012). *Towards meta-evolution via embodiment in artificial chemistries*. PhD thesis, University of York.

Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K., and Taylor, T. (2019a). Open-Ended Evolution and Open-Endedness: Editorial Introduction to the Open-Ended Evolution I Special Issue. *Artificial Life*, 25(1):1–3.

Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., and Taylor, T. (2019b). An Overview of Open-Ended Evolution: Editorial Introduction to the Open-Ended Evolution II Special Issue. *Artificial Life*, 25(2):93–103.

Soros, L., Cheney, N., and Stanley, K. O. (2016). How the strictness of the minimal criterion impacts open-ended evolution. In *Artificial Life 2016 (Proceedings of the Fifteenth International Conference on the Simulation and Synthesis of Living Systems)*, pages 208–215, Cambridge, MA, MIT Press.

Soros, L. B. (2018). *Necessary Conditions for Open-Ended Evolution*. PhD thesis, University of Central Florida.

Soros, L. B. and Stanley, K. O. (2014). Identifying minimal conditions for open-ended evolution through the artificial life world of chromaria. In *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 793–800, Cambridge, MA, MIT Press.

Standish, R. K. (2003). Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, 3(02):167–175.

Taylor, T., Bedau, M., and Channon, A. (2016). Open-ended evolution: Perspectives from the oee workshop in york. *Artificial Life*, 22:408–423.