

How the Strictness of the Minimal Criterion Impacts Open-Ended Evolution

L. B. Soros¹, Nick Cheney², and Kenneth O. Stanley¹

¹Department of Computer Science, University of Central Florida, Orlando, FL, USA

²Department of Biological Statistics and Computational Biology, Cornell University, Ithaca, NY, USA

lsoros@cs.ucf.edu, nac93@cornell.edu, kstanley@cs.ucf.edu

Abstract

Because the kind of open-ended complexity explosion seen on Earth remains beyond the observed dynamics of current artificial life worlds, it has become critical to isolate and investigate specific factors that may contribute to open-endedness. This paper focuses on one such factor that has previously received little attention in research on open-endedness: the *minimal criterion* (MC) for reproduction. Originally proposed as an enhancement to novelty search, the MC is in effect a different abstraction of evolution than the more conventional competition-focused fitness-based paradigm, instead focusing on the minimal task that must be completed for an organism to be allowed to produce offspring. The MC is interesting for studying open-endedness because in principle its *strictness* (i.e. how hard it is to satisfy) can be varied on a continuum to observe its effects. While in many artificial life worlds the MC strictness is implicit and therefore difficult to vary systematically, in the previously-introduced Chromaria world, the MC is designed to be set explicitly by the experimenter, making possible the systematic study of different levels of MC strictness in this paper. The main result, supported by visual, quantitative, and qualitative observations, is that the strictness of the MC can profoundly affect open-ended dynamics, ultimately deciding between complete stagnation (both with extreme strictness or complete relaxation) and orderly divergence. This result offers a lesson of particular importance to worlds whose MCs are not explicit by exposing an area of sensitivity within open-ended systems that is easy to overlook because of its implicit nature.

Introduction

Artificial life (alife) offers a unique opportunity to replay the tape of life (Gould, 1989) and thereby empirically test hypotheses about phenomena observed in nature. In particular, alife worlds offer an ideal platform for exploring and controlling for mechanisms of interest in the domain of open-ended evolution (OEE) (Channon and Damper, 2000; Lehman and Stanley, 2015; Miconi and Channon, 2005; Ofria and Wilke, 2004; Ray, 1992; Soros and Stanley, 2014; Spector et al., 2007; Yaeger, 1994). Though a universally satisfying definition of OEE has been elusive (Bedau et al., 1998; Channon, 2003, 2006; Juric, 1994; Maley, 1999), open-ended systems should *at minimum* not stagnate or converge. Biological evolution is widely considered an open-

ended process, exhibiting 3.8 billion years of diversification and increasing complexity that continues today (Miconi, 2008). A major challenge to investigating this phenomenon scientifically is of course that experiments on computers must be tractable within much shorter timespans. Given that experimental models necessarily entail such unnatural constraints on time (in addition to space), the need to induce a rate of evolution that achieves natural-seeming dynamics is paramount. However, little is yet known on how to press the delicate levers of evolution to control its speed and open-endedness.

One such lever is the *minimal criterion* (MC) for reproduction, which in effect means the minimal function an organism must perform for its lineage to continue. As an example, on Earth the MC is in effect to survive and physically create a copy of oneself. However, the MC on Earth is really only one of a vast spectrum of possibilities. In e.g. the recent Chromaria alife world (Soros and Stanley, 2014) the MC is simply to plant oneself somewhere that has colors similar to one's own morphology (which then will cause an offspring to be created without any need for an evolved reproductive apparatus). This general concept of MC began to appear in evolutionary computation literature in an algorithm called *minimal criterion novelty search* (MCNS) (Lehman and Stanley, 2010) as a low-level alternative abstraction to fitness-centric evolution. Its first appearance in OEE was in Chromaria.

The attractions of the MC for the purpose of developing a more fundamental understanding of OEE are threefold. First, it allows fitness, which is usually explicitly formalized as a score on a continuum, to become an implicit side effect of a deeper evolutionary principle. Second, by abstracting away the metric of fitness, it enables domains radically different from Earth to be tested in the context of OEE while still maintaining a parallel with natural evolution. Third, the MC becomes just the kind of lever that can be calibrated to alter the overall dynamics of long-term evolution. That is, it is possible to alter its *strictness*. On Earth, this notion of strictness is implicit in the difficulty of surviving to the point of physically making a copy of oneself, but in other conceiv-

able worlds, the difficulty of satisfying the MC can be made *explicit*, allowing the experimenter literally to tune it precisely and observe the consequence. Chromaria offers such an opportunity through its color-matching MC, the strictness of which can be tuned straightforwardly.

Experiments in this paper will reveal that the strictness of the MC in OEE has a similar effect as selection pressure in traditional evolutionary algorithms – the more strict, the more convergent evolution tends to become. At the other end of the spectrum, the MC can allow the world to degrade into chaos when not strict enough. However, the strictness of the MC interestingly does *not* correspond to the proportion of the population that is allowed to reproduce; in fact the MC can be very strict even as many still succeed in satisfying it (and thereby reproducing). At the same time, when set too strictly it becomes a force for convergence and hence an obstacle to OEE. As results in this study will show, the MC is instrumental in striking the right balance between order and chaos – a sweet spot for open-endedness made accessible to investigation by the explicit MC in Chromaria.

The implications of the results speak both to how to steer OEE towards desired dynamics and also to the delicacy of the balance that exists in open-ended systems such as nature where the MC is implicit. Stray too far from this golden balance, and the world can either stagnate or explode into meaningless chaos – a prescription worth considering for any future attempts to achieve OEE.

Background

To understand the significance of MC-based evolution, it is important first to understand the difference between *reproduction* (the passive generation of an offspring, usually as a reward for good performance) and *self-replication* (the active assembly of a copy of oneself). On Earth these processes are conflated because biological self-replication is the sole means of reproduction. However, in alife worlds, an evolutionary algorithm can allow an individual to reproduce even if the individual evolves no mechanism for self-replication. From an algorithmic point of view in such systems, traversing genotype space is all that really matters. Abstracted in this way, the MC can be unfastened from the familiar imperative of self-replication because the MC can in principle be to achieve something *other* than self-replication, which would still then lead to reproduction.

The idea of a MC distinct from self-replication first arose in the context of novelty search (Lehman and Stanley, 2011), an evolutionary algorithm that ignores fitness and instead searches only for behavioral novelty. While pure novelty search proved effective in closed domains with limited possible behavioral trajectories (such as for robots in an enclosed maze), Lehman and Stanley (2010) observed that the method gets lost in spaces that encompass vast regions of degenerate behaviors (such as when robots can wander outside the maze). To address this problem, they proposed adding a

MC that everyone in the population must satisfy by displaying a minimal level of competence (such as staying within the maze) to be even considered for further selection based on novelty. In this way, filtering selection through a MC allows the population to stay focused on the interesting part of the search space.

However, as emphasized by Lehman and Stanley (2010), the MC offers more than just practical utility. More significantly, it is an abstraction of a *fundamental aspect of evolution* that is not captured conceptually by traditional fitness-centric models. The shift in focus from performance-based selection in nature to surviving long enough literally to manufacture a rough copy of oneself transforms the view of evolution from emphasizing optimization instead into a picture of a kind of tinkerer that continually discovers new ways to do the same thing (i.e. satisfy the MC), reminiscent of the concepts of adaptive radiation (Schluter, 2000) and neutral networks (Wagner, 2011). This perspective merits further investigation because it provides a potentially fruitful alternative view of natural evolution largely orthogonal to the usual competition-centric and fitness-inspired interpretations, thereby opening up new avenues for investigation.

Of course, this alternate abstraction need not subsume the fitness-based view. Rather, it provides another angle for investigating evolutionary mechanics as demonstrated by experiments with the recent Chromaria alife world. The original paper introducing Chromaria (Soros and Stanley, 2014) hypothesizes that OEE requires a non-trivial MC and that self-replication is just one of many possibilities for such non-triviality. Chromaria becomes a test for this hypothesis by driving evolution with an explicit MC (planting in a region of similar color to oneself) that is *not* self-replication. Moreover, because the threshold for satisfying this MC (i.e. the degree of similarity required to plant) can be varied, Chromaria provided for the first time a world in which the strictness of the MC can be isolated and directly manipulated to see its effects. Presumably, the difficulty of satisfying the MC would have a profound effect upon the consequent evolutionary dynamics. Yet no studies thus far have explored how to adjust the difficulty of satisfying the MC to strike just the right balance between order and chaos in open-ended alife worlds. The consequent lack of understanding is problematic because, as the remainder of this section will demonstrate, the reproductive mechanisms in many existing alife worlds incorporate a form of MC even if it is not defined as such explicitly.

For example, in Avida (Ofria and Wilke, 2004), individuals are encoded as software programs and the MC, as on Earth, is to self-replicate. However, Avida additionally rewards individuals that evolve to perform certain computational tasks by allowing them to loop through their evolved programs (and thereby potentially reproduce) more frequently than others can. When an individual reproduces in this way, its offspring replaces a preexisting individual




Figure 1: **Chromaria**. The initial Chromarian is born at the center of the world and then must find an appropriate place to plant. Subsequent individuals are born wherever their parents planted. The color-rich borders initially provide the only viable options, but more emerge as Chromarians continue to thrive in the environment.

in the population. Thus the MC of self-replication is made more difficult by introducing the threat of effectively random replacement.

Polyworld (Yaeger, 1994) also implements an Earth-like MC by removing individuals from the population if they fail to survive long enough to mate. In contrast with Avida, Polyworld does not award fitness bonuses. However, Polyworld creatures are given a primitive fighting behavior that can be used to kill other creatures and thereby generate food. In this way, selection incentivizes the evolution of adversarial behaviors, which in turn increases the difficulty of meeting the survival-based MC.

The reproductive mechanisms implemented in Avida and Polyworld are typical of those in other alife worlds, which frequently include similar nature-inspired mechanisms that effectively increase the difficulty of satisfying the MC. However, alife is not bound by nature's constraints and thus admits a variety of alternate MCs that can be more explicitly quantified and manipulated. As a result, alife opens up the opportunity to learn via experimentation the largely unexplored implications of varying MC strictness. The next section describes the alife world of Chromaria, which serves as the domain in this paper for experiments varying the MC strictness.

Chromaria setup

Chromaria¹ (Soros and Stanley, 2014) is an alife world explicitly designed for testing theories about the necessary conditions for OEE (figure 1). In deliberate contrast with many other alife worlds, Chromaria is designed without an explicit notion of competition or relative fitness. Instead, Chromarians qualify for reproduction (via an evolutionary algorithm) by satisfying the unique MC of navigating the world and finding an appropriate place to plant themselves.

Each Chromarian's morphology is a two-dimensional image composed of RGB pixels. To facilitate intelligent nav-

¹Source code is available at <http://github.com/lsoros/chromaria>.

igation of the world, each Chromarian is equipped with a rectangular visual field consisting of 100 RGB sensors centered at the forefront of the Chromarian's body. Half of these sensors fall underneath the body and the rest extend in front of the creature. The exact resolution of the visual field depends on the creature's morphology; as body length and width increase, the distance between neighboring sensors grows. The visual field is complemented by a heading-sensitive *compass* consisting of 8 pie slice sensors. All sensor values are scaled and then input to an evolved multimodal neural controller (Figure 2). The output layer of this network has four effector nodes corresponding to the Chromarian's rotation (L and R), speed (S), and inclination to plant itself (P). If the P node's activation level exceeds a constant threshold, the Chromarian is immobilized and the function described in Figure 3 calculates a *matching value* quantifying how closely the creature's morphology matches the colors already on the ground at the requested planting location. The attempt succeeds (and the MC is thereby considered met) if the matching value exceeds a configurable *threshold* (τ), which in effect thereby controls the strictness of the MC. Successful planting thus requires evolving a synergistic combination of morphology and behavior.




Figure 2: **Behavioral controller**. All sensor values are scaled and then input to an evolved multimodal neural controller. Each plane represents an array of sensors or neurons. Arrows between planes in this schematic denote *sets* of connections between one plane and another. The input layer contains three individual color fields, and an additional set of heading inputs. The four output nodes control movement and planting behaviors (which enable the Chromarian to satisfy the MC). The maximum number of connections in this network (evolved by HyperNEAT) is 30,448.

In the unconventional MC-driven main loop in Chromaria, the Chromarians that have successfully planted most recently are kept in a *parent queue* capped at 100 individuals. The parent queue initially contains only one initial seed Chromarian, which is found by preliminary search (Soros and Stanley, 2014) and thereby guaranteed to successfully




Figure 3: **Boolean MC function in Chromaria.** Each pixel of both the Chromarian’s morphology and sensor field (an area of ground overlapping the front of the Chromarian) is placed into one of eight bins: black, white, red, green, blue, yellow, cyan, or magenta. Here, a simple morphology is shown to the left of its binned equivalent and a histogram of the bins. The bins are defined by halving the ranges $[0,255]$ that the R, G, and B component values can take. For instance, any pixel with $R \in [0, \frac{255}{2}]$ (more non-red than red), $G \in [0, \frac{255}{2}]$ (more non-green than green), and $B \in [0, \frac{255}{2}]$ (more non-blue than blue) falls into the black bin because black has values $R,G,B = 0,0,0$. Once every pixel is binned in this way, color ratios are calculated for each bin by dividing the bin size by the total pixel count. Ratios are recorded for both the morphology and sensor field. The differences between these ratios for each color are summed to get a *matching value*. If this value is less than a configurable threshold τ (effectively the threshold for satisfying the MC), the function is satisfied.

plant. The initial seed individual is born at the center of the world and navigates to plant on the color-rich border. Subsequent offspring are born on top of their parents and then must similarly move away from their spawn point to find an appropriate place to plant. To prevent individuals who attempt trivially to plant without moving from reproducing, planting attempts made during the first 25 ticks of an individual’s life are automatically invalidated. If an individual either makes an invalid planting attempt or fails to plant altogether, the Chromarian is removed from the world and does not generate any children. However, if a Chromarian succeeds in moving to an appropriate location and then requests to plant, its body remains frozen and thereby supercedes all other pixels at that location. Its offspring is then generated (because anyone who satisfies the MC is given an offspring) based on the reproductive mechanics of HyperNEAT (Stanley et al., 2009). Accordingly, the genotype used in this system is a compositional pattern producing network (CPPN; Stanley 2007), an indirect encoding that generates patterns with regularities seen in nature such as symmetry, repetition, and repetition with variation. Separate CPPNs are used to encode the morphology (Figure 4) and neural controller. Parameters for evolution are included with the source code.

Note that reproductive dynamics in Chromaria are unlike those in many other alife worlds. To both allow the steady state evolutionary algorithm to explore multiple lineages simultaneously and to guarantee that *all* Chromarians who successfully plant eventually get to reproduce, offspring who satisfy the MC are placed at the end of the parent




Figure 4: **Morphology-encoding CPPN and decoding process.** The CPPN on the left is decoded by iteratively activating the network with pairs of polar coordinates r and θ as input and using the resulting outputs to draw the morphology pixel by pixel. Upon activation, the CPPN returns an r_{max} for each value of θ , which determines how far the perimeter of the Chromarian’s body extends at that angle. The CPPN is then activated again for every r on the interior of this border to get the corresponding RGB values. In this way the CPPN determines both the shape (via the r_{max} output) and internal color (via the RGB outputs) of the Chromarian. These characteristics ultimately determine where the Chromarian can successfully plant and thereby satisfy the MC.

queue and only reproduce once they reach the front of the queue. (The oldest parent in the queue is removed to make room for the new offspring if the queue is at maximum capacity.) After reproducing, if it is not the oldest, the Chromarian at the front of the queue is sent to the back again. Thus the individual active in the world is not always the child of the individual that was active just before it. Though potentially unintuitive, this process forces the system to allow every preexisting member of the population to reproduce before a newcomer. As a result, the explicit competition so central to many alife worlds is intentionally absent from Chromaria. Note that if offspring of enough parents in the queue fail to plant, some parents may receive more than one opportunity to reproduce when they advance again to the front of the line.

Chromaria is chosen as the domain for experiments in this paper precisely because it allows the MC to be studied in isolation from competition and other selective pressures (everyone who satisfies the MC is guaranteed an offspring). Additionally, the difficulty of satisfying the MC is easily adjusted by simply setting a different τ for the planting function described in Figure 3. An interesting question that will be clarified by these experiments is what features of open-endedness can be achieved when more conventional competitive mechanisms are absent. However, the hope is that the results will inform the design of systems that *do* involve more intricate evolutionary pressures.

Experiment

One aim of this study is to develop an intuition for how the MC affects the balance of order and chaos in potentially open-ended evolutionary systems. For this reason, experimental runs are carried out through 350,000 reproductions, each lasting 5 to 7 days (wall clock time). Fifty runs were performed in total, each starting from the same initial seed and varying only τ . Twenty baseline runs were performed with a **moderate** τ of 87.5%, which is the same matching threshold implemented in previous experiments by Soros and Stanley (2014). Additionally, the MC was made more **strict** by increasing τ to 95% and made more **forgiving** by decreasing it to 75%. Ten runs were performed with each of these thresholds.

A related aim of this study is to investigate one of the hypothesized necessary conditions for OEE proposed by Soros and Stanley (2014). Though artificial life provides a promising platform for answering the question of exactly what is necessary for OEE, surprisingly few empirical studies have been performed to test the validity of the few frameworks proposed thus far (i.e. by Waddington (1969) and Taylor (2004)). According to the hypothesis that the MC is *necessary* for OEE (Soros and Stanley, 2014), evolving systems should become degenerate if every individual is allowed to reproduce. This hypothesis in effect addresses the extreme case wherein the MC is so relaxed that there is **no MC (the control)**. To this end, an additional set of ten control runs are performed wherein every individual generates an offspring regardless of what it does during its lifetime. Note that this control setup is slightly more complicated than simply setting τ to 0% because planting attempts during the first 25 ticks of an individual's life are normally invalidated, but this restriction is removed in the control case. Individuals that *never* ask to plant also generate offspring (after timing out) in the control case, though only the bodies of Chromarians who actively ask to plant remain visible in the world after their lifetime ends.

Results

Chromaria's visual design allows evolutionary dynamics to be observed in an intuitive, engaging manner. Harnessing this capability, sequences of screenshots up to 350,000 reproductions demonstrate typical effects of varying τ to produce different MC strictness levels (Figure 5). The story told by these sequences illustrates the profound effect of MC strictness on OEE. At one end of the spectrum, when the MC is strict, the world only slightly changes over hundreds of thousands of reproductions. Interestingly, the extreme opposite end of the spectrum, when there is no MC at all, while initially chaotic, also ends up descending into stagnation because the entire breeding population eventually loses the ability even to request to plant. It is only at the middle ranges that consistent and coherent change (both in terms of morphology and behavior) is readily observable.

Of course, screenshots offer only a static portrait of life in Chromaria and cannot capture the diversity and complexity of individual Chromarians' behaviors. For this reason, videos (and additional screenshots) are available at <http://eplex.cs.ucf.edu/chromaria-resources/>. Observing these dynamic worlds reveals that the most interesting behaviors are expressed when the environment contains a variety of simple RGB signals. When the MC is strict, little deviation from the simple dichromatic initial seed is tolerated, and as a result the world becomes filled with strong magenta, blue, and white signals. Individuals evolve behaviors such as turning left when a magenta/blue edge appears and then planting when a patch of whiteness (such as at the frontier created by new planters) is encountered. When the MC is moderate, environmental complexity increases as a result of Chromarians with more complex morphologies successfully planting. However, the color-based signals in the environment become harder to differentiate and individuals begin to rely on heading (which provides a more reliable signal) to affect planting behaviors. This phenomenon is readily observed in the later stages of the runs with moderate MC; though some individuals visibly change direction based on environmental signals, others ignore the RGB inputs and simply move in smooth arcs and then plant once a desired heading has been reached. Runs with the forgiving MC exhibit the widest diversity of behaviors, reflecting the clear yet diverse RGB signals that emerge in the environment as individuals successfully plant. Finally, when the MC is completely absent, individuals tend to follow simple trajectories with no inherent purpose. Eventually, planting behaviors end up disappearing completely in this case, and senselessly crashing into walls becomes a popular behavior. Recall that the second goal of this experiment is to test the hypothesis that a MC is necessary for OEE. If this hypothesis is true, then evolution should stagnate when the MC is absent. In fact such stagnation (as depicted in Figure 5) *does* occur in every control run.

The qualitative view is strengthened by complementary quantitative analysis. Of course, the quantification of open-endedness is a contentious topic, reflecting the general lack of consensus regarding the phenomenon's salient features. Activity statistics (Bedau et al., 1998; Channon, 2003), which measure the persistence of advantageous genotypes over the course of evolution, present one possibility for quantification. However, alternate metrics are devised for the experiments in this paper so that the control runs (which completely eliminate genotypic advantage by giving every individual an offspring) can be included in comparisons.

The first metric, which is an application of the graph theoretic notion of *connected components* (Hopcroft and Tarjan, 1973), interprets the RGB phenotype space as a three-dimensional cube-shaped graph (Figure 6). Every 250 reproductions, each individual in the parent queue is placed into one of 216 bins based on the R, G, and B ratios in its mor-

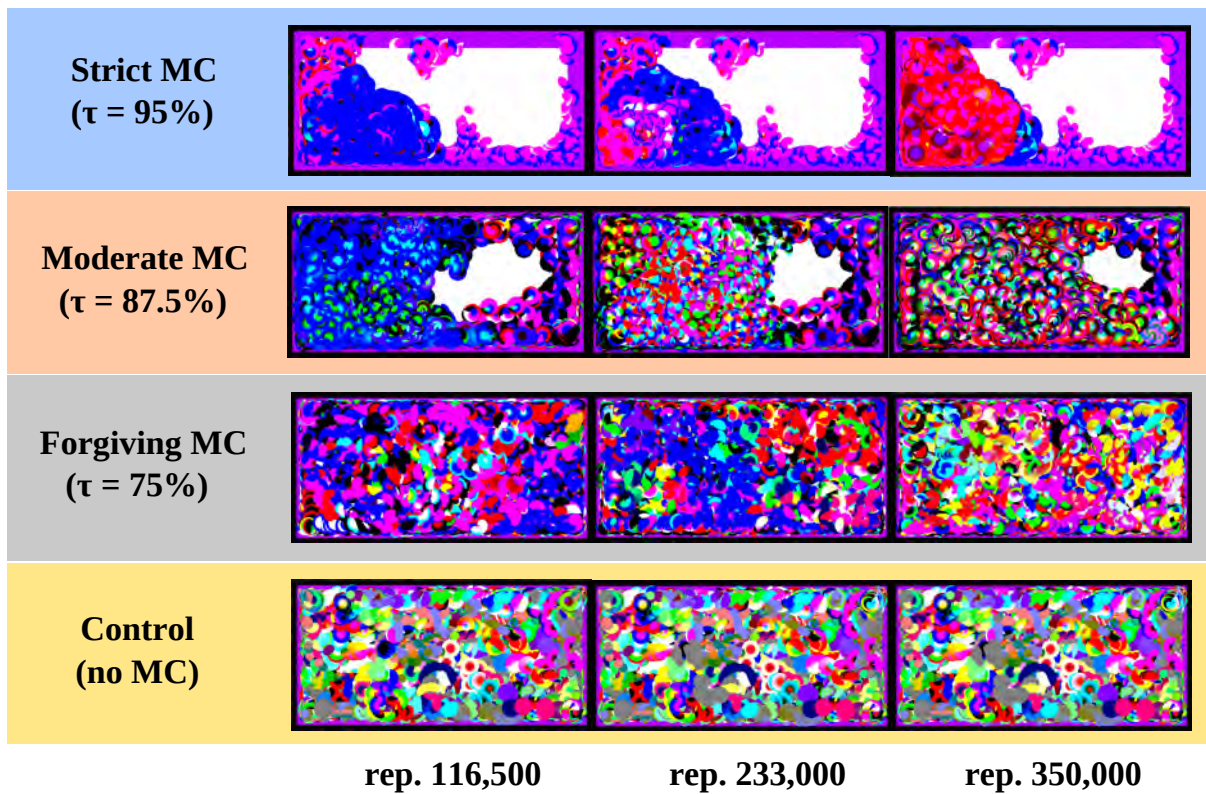


Figure 5: **Representative worlds over time.** Worlds with different MC strictness are depicted after the specified numbers of reproductions have occurred. The variable τ quantifies the difficulty of the MC (detailed in the Experiment section). When the MC is strict, evolution exhibits little progression and the breeding population lacks diversity (an undesirable result for OEE). As the MC is relaxed, diversity increases but stability eventually plunges into chaos (eventually leading back to stagnation). Videos demonstrating the complexity of behavior in additional runs are available at <http://eplex.cs.ucf.edu/chromaria-resources/>.

phology. The individual bins can then be treated as nodes in a graph, with connections existing between adjacent nodes. Evolutionary divergence (or lack thereof) can then be approximated by the average number of *disjoint components* (\overline{dc}) in the RGB cube at any given time in the run. Interestingly, the minimal value of 1 therefore jointly indicates that evolution is *either* maximally ordered (when few bins are filled) or maximally chaotic (when all bins are filled). Between these extremes, \overline{dc} increases as evolution diverges. However, this metric (which indicates divergence only every 250 reproductions) does not tell the full story because it does not reveal how much of the space was eventually explored over the *entire* run. For this reason, the average total number of bins filled over the course of an entire run (\overline{bf}) is also recorded for each MC level.

Figure 7 depicts \overline{dc} over time for each run. As expected, when the MC is strict ($\tau = 95\%$) the average number of disjoint components and number of bins filled per run are both low ($\overline{dc} = 1.34$, $\sigma = 0.59$; $\overline{bf} = 51/216$), indicating limited evolutionary divergence. Interestingly, when the MC is relaxed to the more moderate baseline value ($\tau = 87.5\%$), \overline{dc} increases only slightly to 2.24 ($\sigma = 0.42$), but \overline{bf} sharply




Figure 6: **Discretized RGB space.** Every Chromarian in the parent queue is placed into 1 of 216 bins based on the R, G, and B ratios in its morphology. The number of connected components in the color space approximates the amount of correlation in the parent queue. When all bins are filled (as is the case when the system is completely chaotic) as shown above, the number of connected components is 1. It is also 1 when the system is completely converged.

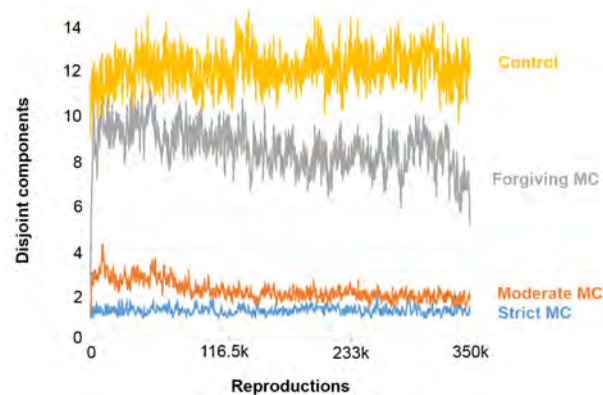


Figure 7: **Average number of disjoint components in the parent queue over time.** A minimal value of 1 indicates that evolution is either maximally ordered (when few bins are filled) or maximally chaotic (when all bins are filled). However, *high* values may also indicate a tendency towards chaos. For this experiment, samples are taken every 250 reproductions. Student’s t-tests with p-values $p < 0.05$ indicate significant differences between strict versus moderate runs for 97.21% of parent queue samples, for 100% of the samples between moderate versus forgiving runs, and for 98.64% of samples between forgiving and control (no MC) runs.




Figure 8: **Average planting success rates over time.** The number of successful planters increases steadily and eventually approximately doubles over the course of the non-control runs, indicating that evolution reliably discovers nontrivial behaviors. However, when there is no MC (the control) evolution degenerates so rapidly into triviality that evolution halts altogether. This result confirms the hypothesis that a minimal criterion is necessary for OEE.

increases to 178.5/216 ($\sigma = 17.56$), signifying a similar degree of order as for the strict MC, but much faster traversal of phenotype space over time. As the MC becomes easier to satisfy, the runs display an increased tendency towards divergence and chaos. In the forgiving case ($\tau=75\%$), \overline{dc} increases dramatically to 8.59 ($\sigma = 0.97$) and eventually fills

all 216 bins. The \overline{dc} measure remains high (12.16, $\sigma = 0.79$) and all bins are still filled when the MC is completely absent, but recall that the world nevertheless actually eventually stops changing entirely without the MC because of the eventual loss of the ability to plant.

The catastrophic effects of removing the MC are clearly reflected also in the immediate and inevitable decline in planting success rates (Figure 8). In contrast, when there is an MC of *any* strictness, planting success rates steadily *increase* over time, suggesting a nontrivial evolutionary process leading to gradual behavioral improvements.

Discussion

Combining the number of connected components with the overall number of bins filled paints a quantitative picture of evolution’s trajectory. When the MC is too strict, evolution proceeds so slowly as to become stagnant. Unsurprisingly, there is little variation between runs in terms of color space coverage because search does not progress far beyond the initial magenta and blue subspace. The MC in this case encourages genetic homogeneity and leads to an effective equilibrium. At the other extreme, when the minimal criterion is absent (the control case), the system initially exists so far from equilibrium that it verges on chaos. While such a system may avoid attractors and thus maximize coverage of the search space, it may also lack the qualities of cohesion and correlation that differentiate evolution from random search. Most problematic, though, is the proliferation of degenerate behaviors because of the lack of pressure to satisfy the MC. The lack of new planters even in the control when τ is 0 (evidenced by eventual planting success rates of 0) indicates that *nobody is requesting to plant* even though any such request would be granted regardless of the colors in the world. This result may at first seem surprising, but in fact strongly supports the hypothesis that maintaining a minimal level of behavioral complexity is necessary for OEE.

These experiments illustrate the complex manner in which the strictness of the explicit MC impacts evolution in Chromaria. However, the results are important more broadly for their potential to generalize to any open-ended alife world that limits which individuals are allowed to reproduce. In other words, modifying the intrinsic strictness of the *implicit* MC in other worlds may have similarly dramatic implications for their degree of open-endedness. This issue is nontrivial because the MC prunes the space being explored by evolution. If there were no MC on Earth, for example, every organism would reproduce regardless of its viability. As a result, degeneracy would thrive whereas otherwise it could not. If the MC were too difficult (in this case, anything requiring functionality beyond that of a single cell), foundational building blocks would be discarded. Thus setting the wrong MC can have disastrous implications for a system’s ability to uncover the maximal amount of interest in a search space. While a designer may be able to

compensate for an overly strict MC by increasing run time if the run time is capped, it may be desirable to relax the MC through whatever means available to increase coverage of the search space in the allotted time.

An emphasis on optimization in evolutionary computation has influenced the primacy of *competition* in the design of alife worlds. Unlike competition-based selection, the MC admits innovations that offer no discernible benefit for an individual (which would otherwise have been discarded by optimization-oriented evolution). For this reason, the MC allows evolution to maximize its potential as a creative force.

However, competition is not the only overly restrictive mechanism. Any selective force that arbitrarily imposes a strict limit on reproduction (i.e. by only allowing a certain percentage of the population to generate offspring) will suffer the same inhibited creative output, including on Earth. Many mechanisms implemented in artificial life systems can be reduced to a MC-based interpretation in this way. By abstracting away the system-specific details and working at the most general level, we can develop a model of open-ended evolution with maximum applicability.

Conclusion

Life strikes a delicate balance between order and chaos. In alife worlds, this balance is moderated at least in part by an often-implicitly-implemented minimal criterion (MC). As results in this paper have shown, tuning this mechanism incorrectly can have disastrous implications for a system's open-endedness. However, if configured correctly, the MC can create a dynamic environment in which evolved life flourishes.

Acknowledgements

This work was supported 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. (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.

Channon, A. D. and Damper, R. I. (2000). Towards the evolutionary emergence of increasingly complex advantageous behaviours. *International Journal of Systems Science*, 31(7):843–860.

Gould, S. J. (1989). *Wonderful Life: The Burgess Shale and the Nature of History*. W. W. Norton and Company.

Hopcroft, J. and Tarjan, R. (1973). Algorithm 447: Efficient algorithms for graph manipulation. *Communications of the ACM*, 16(6):372–378.

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. (2010). Revising the evolutionary computation abstraction: minimal criteria novelty search. In *Proc. of the 12th annual conf. on Genetic and evolutionary computation*, pages 103–110. ACM.

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

Lehman, J. and Stanley, K. O. (2015). Investigating biological assumptions through radical reimplementations. *Artificial Life*, 21(1):21–46.

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

Miconi, T. (2008). *The road to everywhere: Evolution, complexity and progress in natural and artificial systems*. PhD thesis, University of Birmingham.

Miconi, T. and Channon, A. (2005). A virtual creatures model for studies in artificial evolution. In *The 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 565–572. IEEE.

Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial life*, 10(2):191–229.

Ray, T. S. (1992). An approach to the synthesis of life. In *Proc. of Artificial Life II*, pages 371–408.

Schluter, D. (2000). *The Ecology of Adaptive Radiation*. Oxford University Press.

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.

Spector, L., Klein, J., and Feinstein, M. (2007). Division blocks and the open-ended evolution of development, form, and behavior. In *Proc. of the 9th annual conf. on Genetic and evolutionary computation*, pages 316–323. ACM.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162.

Stanley, K. O., D'Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.

Taylor, T. (2004). Niche construction and the evolution of complexity. In *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Artificial Life*, pages 375–380, Cambridge, MA. MIT Press.

Waddington, C. H. (1969). Paradigm for an evolutionary process. In Waddington, C., editor, *Towards a Theoretical Biology, Volume 2*. Edinburgh University Press, Edinburgh, Scotland.

Wagner, A. (2011). *The Origins of Evolutionary Innovations: A Theory of Transformative Change in Living Systems*. Oxford University Press.

Yaeger, L. (1994). PolyWorld: Life in a new context. *Proc. Artificial Life*, 3:263–263.