

The MODES Toolbox: Measurements of Open-ended Dynamics in Evolving Systems

Emily L. Dolson^{1,2,3}, Anya E. Vostinar⁴, Michael J. Wiser^{1,3}, and Charles Ofria^{1,2,3}

¹BEACON Center for the Study of Evolution in Action

²Department of Computer Science and Engineering, Michigan State University, MI, 48823

³Program in Ecology, Evolutionary Biology, and Behavior, Michigan State University, MI, 48823

⁴Grinnell College, Grinnell, IA, 50112

dolsonem@msu.edu

Abstract

Building more open-ended evolutionary systems can simultaneously advance our understanding of biology, artificial life, and evolutionary computation. In order to do so, however, we need a way to determine when we are moving closer to this goal. We propose a set of metrics that allow us to measure a system's ability to produce commonly-agreed-upon hallmarks of open-ended evolution: change potential, novelty potential, complexity potential, and ecological potential. Our goal is to make these metrics easy to incorporate into a system, and comparable across systems so that we can make coherent progress as a field. To this end, we provide detailed algorithms (including C++ implementations) for these metrics that should be easy to incorporate into existing artificial life systems. Furthermore, we expect this toolbox to continue to grow as researchers implement these metrics in new languages and as the community reaches consensus about additional hallmarks of open-ended evolution. For example, we would welcome a measurement of a system's potential to produce major transitions in individuality. To confirm that our metrics accurately measure the hallmarks we are interested in, we test them on two very different experimental systems: NK Landscapes and the Avida Digital Evolution Platform. We find that our observed results are consistent with our prior knowledge about these systems, suggesting that our proposed metrics are effective and should generalize to other systems.

Introduction

A central goal of the field of artificial life is to build evolving systems that capture the full range of dynamics found in natural systems. Such systems should be capable of producing evolutionary outcomes such as sophisticated navigation behaviors, novel cooperative strategies, complex ecosystems, and major evolutionary transitions, to name but a few. Researchers seek such "open-ended" systems for a number of reasons:

1. For biologists, access to systems exhibiting complex and nuanced evolutionary processes allows rapid experimen-

tation and facilitates developing a deep intuition for underlying mechanisms (Tenailon et al., 2016).

2. For evolutionary computation researchers, insights from open-ended evolving systems will allow researchers to break complexity barriers, expanding the classes of engineering problems that evolutionary algorithms can solve (Hara and Nagao, 1999; Potter and Jong, 2000) and producing more general forms of evolved intelligence.
3. For artificial life researchers, it is concerning that there may be dynamics of fundamental importance to biology that artificial life systems do not exhibit. The existence of such dynamics suggests that we are not building evolving systems as innovative as those found in nature, be it due to limited memory, limited time, or simply an insufficient understanding of the necessary components. Identifying these missing factors should allow us to better understand life as it is and to better explore life as it could be.

While various artificial life systems have reproduced individual dynamics – such as the evolution of complex traits (Lenski et al., 2003b), cooperative behaviors (Goldsby et al., 2012), and coexistence of diverse ecotypes (Cooper and Ofria, 2003) – these accomplishments have been in highly controlled circumstances. The overarching goal of open-ended evolution research is to create a system where all of these dynamics emerge more organically, as in nature. Additionally, replicating this process would provide substantial insights into our own origins, including the evolution of human brains. Indeed, harnessing a more open-ended set of evolutionary dynamics could help us spur breakthroughs in the evolution of general artificial intelligence.

Open-ended evolution is a many-faceted concept. A number of patterns are considered to be hallmarks of open-ended evolution (Taylor et al., 2016), most notably the continual production of novelty (Lehman and Stanley, 2011; Banzhaf et al., 2016), unconstrained increases in diversity (Bedau and Bahm, 1994), ongoing increases in complexity (Lenski et al., 2003b; Korb and Dorin, 2011), and shifts in individuality such as those often associated with major transitions in evolution (Smith and Szathmary, 1997). There is a growing

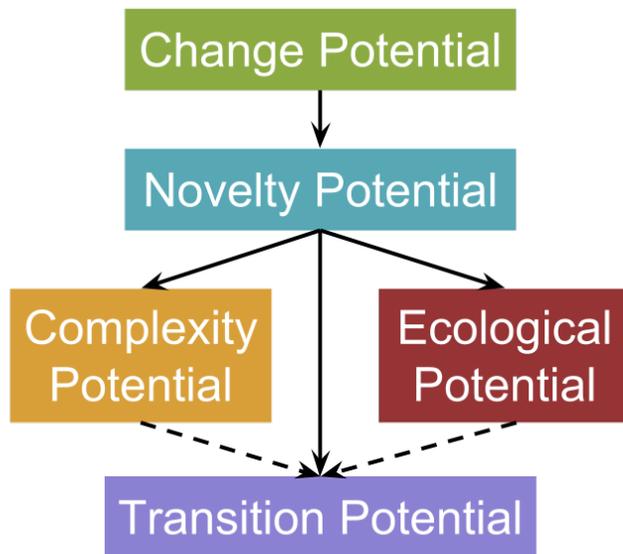


Figure 1: **Relationships between the metrics.** Originally published in (Dolson et al., 2015). Solid lines with arrows indicate metrics which are prerequisites for other metrics.

consensus in the field that all of these dynamics are important pieces of the open-ended evolution puzzle (Taylor et al., 2016). In addition, we have previously suggested that there is a fifth necessary and even simpler dynamic: continuous change in the information content of genomes in the population (Dolson et al., 2015).

These five properties of a system fit into a hierarchy, as shown in Figure 1. For novelty to exist, there must be some degree of change in the information within a population. While this observation is trivially true, many evolutionary algorithms suffer from premature convergence, the absence of non-trivial change. Thus, it remains an important prerequisite to define and explicitly address. Similarly, complexity and diversity can only increase indefinitely if novel members of the population continue to be generated. Finally, transitions in individuality typically involve multiple organisms coming together into a single individual, building from complex and diverse progenitors. All of these dynamics capture different subsets of interesting behavior that evolving systems might exhibit and we propose they are all necessary (but perhaps not sufficient) in a fully open-ended system.

To draw conclusions about what factors of a system promote or inhibit these dynamics, we need methods for measuring the extent to which each dynamic is present. Importantly, these methods must be applicable across a wide variety of systems. Some progress has been made toward this end with evolutionary activity statistics (Bedau et al., 1998a; Channon, 2001), an approach to isolating and quantifying the adaptive component of an evolving system, sepa-

rating out the non-adaptive dynamics. Evolutionary activity statistics require that the user decide on two things ahead of time: a definition for “components” (meaningful individual pieces of a system) and a way of filtering noise out of the system (typically by contrasting with a shadow population that evolves with selective pressures turned off).

Thus far, components have needed to be defined for each system on a case-by-case basis. In artificial life systems, alleles or genotypes are typically used as components, while in the fossil record, whole species were used as components (Bedau et al., 1998b). This flexibility to choose different components is valuable, as it allows for the study of open-ended evolution at different scales of organization. However, it also means that care must be taken when comparing evolutionary activity statistics across systems. Here, we suggest a component definition that should work for any system in which genomes are composed of elements that collectively determine fitness (see Identifying Meaningful Sites in the Genome).

Due to the critical role of stochasticity in evolution, most evolving systems are noisy. In order to make behavioral generalizations, we need a way to distinguish evolutionary signal from this noise. In the original description of evolutionary activity statistics, a specific method was proposed for doing so: for each run of a system, there should be a corresponding “shadow” run in which any outcome of selection is replaced with a random choice. Dynamics observed in the shadow run can then be subtracted out from those in the main run. While this control can be highly informative, it is challenging to implement in many systems and requires researchers to be able to isolate all selective events in the system. For example, when evolutionary activity statistics were applied to the fossil record, a different filter had to be used: the assumption that any species that was successful enough to have made it into the fossil record was probably evolutionarily successful for a substantial amount of time. In this paper, we build on this idea to propose a filter for evolutionary activity that can be more easily implemented across a variety of systems (see Filtering Out Noise).

Evolutionary activity statistics classify evolving systems based on how open-ended they are. However, it is relatively easy to create a system that falls into the most open-ended class while still failing to further our goals for open-ended evolution research or to match our subjective understanding of what we would expect from a truly open-ended system (Maley (1999)). Indeed, there is debate over whether “open-endedness” is even quantifiable (Stanley and Soros, 2016). Moreover, it is our opinion that most efforts to define systems as either open-ended or not have largely been unproductive; open-endedness is likely better thought of as a continuum than as a binary. While there is much debate over what would constitute a fully open-ended system, there is consensus in the field that we are not particularly close

to building such a system¹. Our goal in this paper is to extend evolutionary activity statistics into easy-to-use diagnostic criteria that quantitatively measure key hallmarks of open-ended evolution. We want researchers to be able to isolate the effects of experimental settings on these hallmarks, keeping such results relevant across experimental platforms. As such, we hope to spur a more consistent and comparable march toward true open-endedness, adding new metrics to this toolbox as the community reaches a consensus on the features that we should promote.

In the rest of this paper we will introduce the Measurements of Open-ended Dynamics in Evolving Systems (MODES) toolbox and explore the behavior of the metrics it contains in the context of two evolving systems: NK-Landscapes (Kauffman and Levin, 1987) and the Avida Digital Evolution Platform (Ofria and Wilke, 2004).

Background

Evolutionary Activity

Evolutionary activity statistics attempt to quantify the extent to which adaptive dynamics are occurring in a population. In most applications, evolutionary activity has been measured as the length of time that components exist in the population beyond what would be expected in the absence of selection (Bedau et al., 1997, 1998a; Channon, 2003). This value was chosen because it translates easily across systems and represents a universal measure of evolutionary “success”. In earlier work, a measure of selective sweeps was used instead of component existence time (Bedau and Packard, 1992), but this metric could not be easily generalized across systems.

Multiple facets of evolutionary activity are used in the interpretation of evolutionary activity statistics: the activity of new components (A_{new}), the mean (or median) cumulative activity of components in the population (\bar{A}_{cum}), and the diversity of components in the population (D). Based on the long-term behavior of these values, systems that exhibit qualitatively similar dynamics have been grouped together into a class of evolutionary dynamics. Initially, three possible classes were described: no evolutionary activity, bounded evolutionary activity, and unbounded evolutionary activity. Over time, additional classes have been added to more precisely reflect the types of systems observed. For ease of referring to these classes, table 1 merges together all prior additions to the original classification system of which we are aware.

¹This line of thought originally lead us to conceptualize the metrics described here in terms of possible barriers a system might encounter that would prevent it from being open-ended (Dolson et al., 2015). However, our attempts to measure these barriers align closely with dynamics that have since been identified as hallmarks of open-ended evolution. Ultimately, these perspectives are two sides of the same coin and both are useful frames through which to view open-endedness. For simplicity, we phrase this paper in terms of hallmarks rather than barriers.

According to the original formulation of evolutionary activity statistics, in order for a system to be categorized among the most open-ended systems (originally class 3, now class 4) it must exhibit unbounded growth in summed evolutionary activity across all components in the population (Bedau et al., 1997) (see Table 1). Technically, this growth could happen either because of an unbounded increase in the number of components (diversity) or because of an unbounded increase in the average evolutionary activity of components in the population. The latter case was originally thought to not occur (Bedau et al., 1998a), however when such a case was observed, Channon suggested that class 3 open-ended dynamics should be broken up into three subcategories. These subcategories depend on whether the growth in evolutionary activity was driven by diversity, per-component evolutionary activity, or a combination of both (Channon, 2001) (see Table 1).

In parallel, Skusa and Bedau refined the classification in a different way (Skusa and Bedau, 2003), inserting a new second class in which evolutionary activity was unbounded but no novel components came into being (see Table 1). Such a situation would describe purely ecological dynamics. This observation may seem surprising at first - shouldn't unbounded evolutionary activity involve adaptation? However, when evolutionary activity is measured as the existence time of a component, evolutionary activity statistics draw no distinction between stabilizing selection and selection favoring changes to the status quo (Channon, 2003). Thus, pressure for multiple eco-types to continue existing in their current form (i.e., ecology) will show up as as evolutionary activity above and beyond what is observed in the shadow run.

In fact, the presence of a single component under stabilizing selection will trivially cause the mean evolutionary activity to increase indefinitely; such a component will sit in the population, increasing the population's activity counter despite being quickly lost from the shadow population. This behavior casts doubt on how we should interpret class 4b, as well. To remedy this concern, Channon suggested that we should look for unbounded growth in median (rather than mean) per-component evolutionary activity (Channon, 2003). This adjustment is a drastic improvement, but it still does not eliminate the possibility that systems exhibiting class 4b evolutionary dynamics are not doing quite what we would expect. If at least 51% of the components in the population are under stabilizing selection – as would be expected in an ecological system – the rest of the population could still be behaving like a class 3 system. While such a system would still be interesting for ecological studies, our understanding of it would not be well-served by conflating it with

Class	Median ² Evolutionary Activity (\tilde{A}_{cum})	Change	Novelty A_{new}	Diversity (D)	Complexity	Type of evolutionary dynamics	Described in
1	zero	?	zero	bounded	bounded	None	Bedau et al, 1997
2	unbounded	?	zero	bounded	bounded	Uncreative	Skusa and Bedau, 2003
3	bounded	positive	positive	bounded	bounded	Bounded	Bedau et al, 1997
4a	bounded	positive	positive	unbounded	bounded	Unbounded	Bedau et al, 1997
4b	unbounded	positive	positive	bounded	?	Unbounded	Channon, 2001
4c	unbounded	positive	positive	unbounded	?	Unbounded	Channon, 2001

Table 1: Table combining all previously described classes of evolutionary dynamics as measured with evolutionary activity statistics. For each class, we show the response of all quantities measured for evolutionary activity statistics and in our proposed metrics (novelty and diversity should behave equivalently between the two systems). Note that we expect bounded evolutionary activity to imply bounded complexity, as any scenario in which complexity is growing without bound should imply that evolutionary activity is too. Question marks indicate that the value of a given metric is not specified in the description for a class of evolutionary activity. Higher-numbered classes are generally believed to fall further along the continuum of open-endedness than lower-numbered classes.

systems that are exhibiting open-ended adaptive evolution.

How can we know whether evolutionary activity is driven by stabilizing selection rather than more interesting dynamics? If every component is experiencing directional (as opposed to stabilizing) selection, the change metric we propose here should theoretically be comparable to the number of components³. In contrast, if most of the population is under stabilizing selection, the change metric should be very low.

Ultimately, our change metric (described in the next section) is in keeping with the original evolutionary activity measurement, which sought to quantify the acquisition of new genetic information (Bedau and Packard, 1992). For this reason, in our suite of metrics, we replace the concept of evolutionary activity with change. We believe that this framing will be easier to measure and interpret with little loss of information (although of course we encourage the use of other measures of evolutionary activity where appropriate). Our change metric does have the downside of not being possible to usefully classify as bounded or unbounded. Because we seek only to compare systems and identify progress toward higher levels of open-endedness, this limitation should not be a problem for us.

Prior work using MODES

Soros (2018) used a preliminary version of our framework (Dolson et al., 2015) to study open-ended evolution in the

²The original formulation of evolutionary activity statistics used mean rather than median, but Channon (2003) makes a compelling argument for using median instead. Using median rather than mean does not change any of the intuitions for how we expect this metric to behave and reduces the risk of non-intuitive behavior due to outliers.

³However, the change metric may often be lower than the number of components because not every component will change during every measurement period.

artificial life system Chromaria. Agents in Chromaria are colorful circles controlled by CPPNs that must find a region of the world that matches their color in order to reproduce. These agents can be classified into species based on their patterns of coloration, and change and novelty can be assessed by measuring the emergence of new species. Ecological interactions in Chromaria occur as a result of individuals planting themselves in the world, which alters the color environment that subsequent agents must navigate. Thus, Soros was able to measure the ecology of Chromaria through a series of visual snapshots of the world, as well as by measuring the number of species that occur over the course of a run. Lastly, she measured complexity in terms of the number of elements in the CPPNs controlling the agents.

Using these MODES-inspired metrics, Soros (2018) investigated three hypothesized necessary conditions for open-ended evolution: 1) some sort of minimal criterion (Soros and Stanley, 2014) must be met before reproduction, 2) when new types of individuals evolve, it should create new ways to satisfy the minimal criterion, and 3) individuals should be responsible for making decisions about how they interact with the world. By measuring hallmarks of open-ended evolution under various controls that removed these conditions, Soros (2018) found strong evidence that all of the conditions are indeed necessary for change and novelty (let alone ecology and complexity) in Chromaria. These experiments perfectly illustrate the kind of hypothesis-driven research that we hope a further formalization of our metrics will enable. Additionally, they serve as an example of the range of approaches that can be taken to translate these concepts between systems.

Applying MODES to biology

Since many hypotheses about open-ended evolution involve comparisons to the biosphere, it is critical that MODES met-

rics are applicable not only to digital systems, but are also relevant to experimental biological systems. To confirm that they are, we consider how we would apply them to a well-studied wet-lab experiment. The Long-Term Evolution Experiment (LTEE) (Lenski et al., 1991) is an exemplar of experimental evolution, consisting of 12 populations of the bacteria *E. coli*, which have been evolving independently for more than 60,000 generations (Good et al., 2017). As detailed in (Taylor et al., 2016), the LTEE exhibits many hallmarks of open-ended evolution, including the criteria we propose here. Because fitness within the LTEE is best described by an unbounded power law function (Wiser et al., 2013; Lenski et al., 2015), the system meets the change metric: populations continue to change in non-trivial ways over time. Further, studies of individual populations within the LTEE have shown numerous examples of the generation of novelty, including exploration of new areas of the fitness landscape (Tenaillon et al., 2016), repeated selective sweeps (Maddamsetti et al., 2015), and new diversity arising after such sweeps (Blount et al., 2012). Toward the ecology metric, several populations within the LTEE demonstrate frequency-dependent fitness dynamics (Ribeck and Lenski, 2015; Rozen and Lenski, 2000; Le Gac et al., 2012; Maddamsetti et al., 2015), which are necessarily cases of ecological interactions. Included in these cases of frequency dependence is a special case (Blount et al., 2008, 2012; Turner et al., 2015a) driven by cross-feeding and specialization on different resources (Turner et al., 2015b). Because each population in the LTEE descends from a single ancestor present at the start of the experiment, all ecological complexity in any population must have arisen during the course of the experiment, and thus satisfies the ecological metric. The complexity metric is inherently harder to quantify in a biological system than in a computational one, but recent large-scale genome sequencing from the LTEE (Tenaillon et al., 2016) offers the promise of being able to measure complexity at the genome level over the course of the experiment. Because our metrics can theoretically be applied to a well-studied and open-ended biological system, they can be used to compare dynamics in a broad range of systems and enable the field of artificial life to move forward in quantifiable steps to open-ended evolution.

Metrics

Overarching techniques

We use two broad techniques to ensure that our metrics can focus on the most relevant and meaningful information in an evolving population. Additionally, we describe a technique for determining whether a metric is bounded or unbounded.

Filtering out noise In any evolving population, mutations continually produce new maladapted genomes that are then purged from the population via natural selection. All of the MODES metrics assume that some form of filtering has been

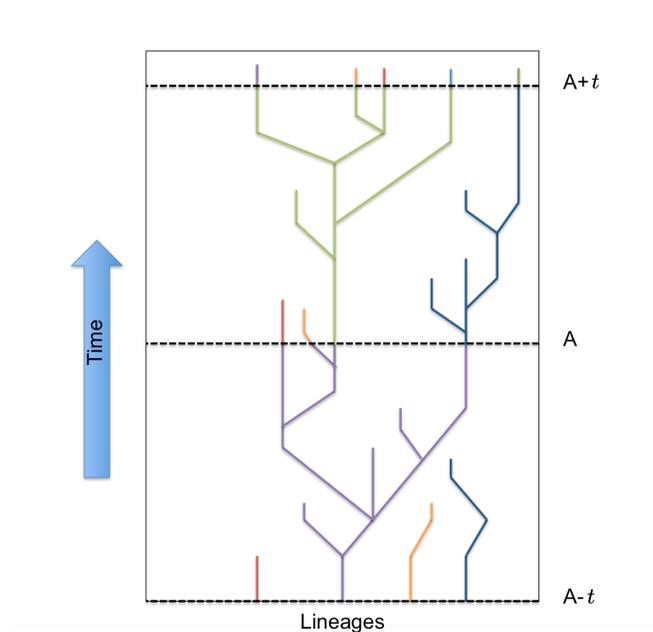


Figure 2: **An illustrative example of how we filter genomes for persistent lineages.** At time point A , the purple lineage has proven to be persistent and therefore the original genome from $A-t$ will be considered meaningful. Similarly, the green and blue lineages persist to time point $A+t$ and so the original green and blue genomes will be considered meaningful as they existed at time point A .

applied before-hand, to prevent maladapted genomes from overwhelming the hallmarks that we are measuring. Here, we describe a persistence filter that we use in our experiments. However, shadow runs are also a viable filter option, and there are likely further useful filtering techniques that have not yet been invented.

To focus only on the adaptive products of evolution, we limit our analysis to those genomes whose descendants persist for a substantial number of generations. We refer to this technique as a persistence filter. We mark each organism with a lineage ID at a given time point A , as demonstrated in Figure 2 (where color indicates lineage ID). The lineage IDs are passed on to offspring for the next t generations, where t is a pre-determined number of generations indicating the length of our filtering process (hereafter referred to as filter length). At time point $A+t$, we determine which genomes from the population at A have descendants at $A+t$. At this point, those genomes are considered **persistent**; in the example in Figure 2 the individuals at the bases of the green and blue lineages are considered persistent at time point $A+t$. These are the individuals that would be evaluated in the MODES metrics. This filtering leads to a delay in counting a genome in a metric until t generations later, but enables us to avoid an apparent increase in metrics

due to drift via mutation. For example, the red and orange genomes from time point A would not be considered in our metrics because their lineages do not persist to time point $A + t$.

How large should t be? The correct value depends on our goals. If we are interested in evolution on a shorter time scale, we may only want to filter out deleterious mutants, which will likely only survive a few generations. In this case, a relatively small value of t should suffice. Indeed, prior open-ended evolution research has used what is effectively a persistence filter with $t = 1$ as a supplemental filtering technique (Channon, 2003). If we are interested in a broader time scale, however, we may want to filter out neutral mutants too and measure only adaptive evolution. In this case, coalescence theory can inform our choice of t . In an asexual population without diversity-preserving forces, the population will periodically “coalesce”, i.e. neutral clades will die out resulting in a new most-recent-common ancestor of the current population. If we take a snapshot of such a population at any given point in time and let the population continue evolving for long enough, a single individual from the snapshot will eventually be a common ancestor of the entire extant population. We refer to coalescence time here as the amount of time that this process takes, although it should be noted that coalescence time is more commonly thought of retrospectively.

If we want to filter out all neutral mutants that do not go on to play a critical role in evolution, it would be ideal to choose a value for t that falls above the expected distribution for coalescence times. If we did so, then we could be confident that any individuals that made it past the filter represented a meaningful part of the evolutionary history of the population. If only a single individual makes it through the filter, that individual must be along the line of descent for the entire population. Multiple individuals making it through the filter would be evidence of ecological dynamics promoting their coexistence.

The median coalescence time for a well-mixed asexual population of N haploid individuals under no selective pressure is $2N$ generations (Fu and Li, 1999). Unfortunately, the expected distribution of coalescence time is exponential, meaning that we would have to choose a potentially impractically large value for t if we want to guarantee that it is rare to get through the filter by chance. However, the presence of selective pressure dramatically reduces expected coalescence time. Since most systems in which people study open-ended evolution do have selective pressure of some form, in practice relatively low values of t are still effective filters.

For a meaningful comparison across populations, we must filter them using consistent values of t . We always expect filters with lower t values to let more individuals through, and it is challenging to separate this effect from changes in the underlying dynamics of the system. Additionally, t must be measured in generations to ensure con-

sistency in the amount of filtering that occurs. Researchers studying systems that use a different time scale need only calculate the average generations within the population to measure t .

In evaluating results, the field should strive to use consistent values of t relative to population size and be aware that, all else being equal, increasing selective pressure will reduce the number of taxa that get through the filter.

This effect brings up an important distinction between this filtering technique and the shadow run traditionally used with evolutionary activity statistics. Whereas shadow runs filter out the effect of neutral processes, the persistence filter does not entirely. We view this reduced filtering primarily as an advantage – drift can be an important part of the evolutionary process – but there may also be situations where it is undesirable. Our metrics are unable to distinguish between class 1 and 2 dynamics or between class 3 and 4b dynamics (see Table 1), although they are able to distinguish between useful subcategories within those classes (as discussed in the Change Metric section below).

Identifying meaningful sites in genomes While a genome may have descendants in t generations, if t is relatively small this persistent genome may not be phenotypically different from another persistent genome in the population. To ensure that we are not separately counting genomes that differ only in non-coding regions, we use an additional filter in which we determine which sites in the genome contain information about the environment. In calculating all of the following metrics, we first reduce the genome to its informative sites.

This approach can easily be extended to any system in which the genome is made up of a set of elements that collectively determine fitness. Whether or not a genomic position is informative can be approximated by measuring the overall fitness effect of either removing it or changing it to a null alternative that is known to not contribute information. A null alternative should be used in cases where changing the structure of the genome changes the meaning of other sites. For example, in *Avida* it is critical that we replace instructions with nulls rather than completely removing them because information can be encoded in the number of instructions between two other instructions. A more accurate technique would be to examine the fitness effect of substituting all possible alternative elements and calculate the potential entropy at that site. When null substitutions are not possible, this technique is an effective method. A caveat to this technique is that genomes that achieve a given result in an excessively fragile manner may appear more complex than more robustly built genomes. To mitigate this issue, the combined fitness effect of eliminating multiple genome elements at once can be measured.

Note that, although identifying informative sites can be computationally intensive, we would need to do so anyway

to calculate the complexity metric. As such, this additional layer of filtering is effectively free.

Determining boundedness In the design of these metrics, we have primarily focused on determining the effect that small changes to a system have on the extent to which that system exhibits hallmarks. However, they can also be used to classify systems in much the same way that evolutionary activity statistics do. As described in Table 1, this classification requires determining whether diversity is increasing without bound. In addition, it would be informative to determine whether complexity is growing without bound. In previous work, the definition of boundedness in this context has been stated in terms of the limit of the supremum of diversity as time goes to infinity (Bedau et al., 1998a). While this is an excellent theoretical definition, taking limits of empirical data as time goes to infinity is generally not practical. Previous applications of evolutionary activity statistics seem to determine boundedness based on whether or not a line on a graph appears to be plateauing. This technique has the potential to be misleading (Wiser et al., 2018).

Instead, we advocate the use of statistics to determine what mathematical model best fits the observed data. We can then classify the pattern as bounded or unbounded based on the limit of the best-fitting mathematical model. Such an approach has previously been used to demonstrate that fitness is following an unbounded growth pattern in a long-term wet lab evolution experiment with *E. coli* (Wiser et al., 2013; Lenski et al., 2015).

MODES Metrics

Change Metric Our first metric focuses on whether the genetic makeup of the population is changing in a non-trivial way. This metric will be above zero during adaptive evolution, including situations where the population is returning to previous states, perhaps due to environmental cycling. In the work presented here, we use a persistence filter (explained above) to ensure that we mark a genome as new only if its lineage full t generations. However, a different filtering technique (e.g. a shadow run) could be used instead. For this comparison, we first find the genomes from persistent lineages from generation A by determining which genomes have descendants in generation A + t . In the example shown in Figure 2, the genomes at the roots of the green and blue lineages would count as persistent. We then compare these genomes to those found to have been from persistent lineages in the previous time point (e.g. we would compare the roots of the blue and green lineages to the root of the purple lineage in Figure 2). In this way, we create a sliding window to observe change in the population. Note that the example in the figure assumes the resolution at which data are collected (i.e. the number of generations between time points) is equivalent to the value of t , but this does not need to be the case. It may be desirable to have a very long length of

t but still gather data frequently. In such a case, each time point is individually filtered by looking ahead t generations, but change is calculated by comparing the set of persistent taxa in the current time point to the set of persistent taxa in the previous time point.

While there is no change metric in the original conception of evolutionary activity statistics, we expect that it will provide similar information to cumulative evolutionary activity (Bedau et al., 1997). Change must be positive in systems exhibiting class 3 or higher evolutionary dynamics, as these systems must all exhibit positive novelty. Class 1 systems may or may not exhibit change; an evolving system that stagnates (e.g. many genetic algorithms) would have zero change, whereas a completely neutral system where all change was caused by drift would sometimes have a non-zero amount of change (depending on the value of t). Class 2 systems would have non-zero change if they were cycling between fixed states, but not if they were purely the result of stabilizing selection.

Novelty Metric The novelty metric measures how many genomes have evolved in the population that have never been seen previously in the experiment. For this metric we again filter out genomes that do not have descendants in t generations, enabling us to focus on meaningful novelty. As with change, we could have used a different filtering technique instead. To measure novelty, we simply count how many genomes from persistent lineages have never been in a previous time points persistent genome pool. It is possible with this metric for a genome to evolve, but not persist, and therefore not be recorded in the permanent history, but then evolve and persist at a later point and be counted as novel. Once a genome has been counted as novel, however, it is part of the permanent history and will never be counted in the novelty metric again. Thus, while a genome could be delayed in being counted as novel, or not counted if it never persists, it will not be counted twice. Our novelty metric is functionally equivalent to A_{new} in evolutionary activity statistics (Bedau et al., 1997).

Complexity Metric The complexity metric measures the maximum complexity (informative sites) of any organism found in the entire population. We recommend the approach described in the section on “Identifying meaningful sites in genomes” above. Once the meaningful sites have been identified, they can be counted to get a measurement of complexity; the value of the complexity metric at a given time point is the highest observed count of informative sites across all taxa in the population that make it through the filter. Further, complexity can be measured even more accurately by using information-theoretic techniques where all possible mutations are considered at each site, and ideally some epistatic interactions.

There is no equivalent to the complexity metric in evolu-

tionary activity statistics. However, as many believe growth in complexity to be an important hallmark of open-ended evolution (Taylor et al., 2016), we feel it is a critical addition. In particular, it would be interesting to find non-trivial systems that exhibit unbounded growth in complexity. We suspect that such growth could only occur in systems exhibiting class 4b or 4c evolutionary dynamics, as bounded evolutionary activity should imply bounded complexity (although the converse is not true).

Ecological Metric The ecological metric measures the amount of information in the population as a whole. While organisms may not contain increasing amounts of information in their individual genomes (as measured by the complexity metric), they could still be increasingly diverse and therefore contain increased information collectively in the population. Ideally, we would measure this collective information by tracking the origin of each piece of information across all genomes in the population and summing up the number of unique pieces of information. Unfortunately, this approach is not computationally practical for many systems. As a proxy, we can look at the diversity of post-filter genotypes reduced to informative sites. Complex ecologies in which multiple subsets of the population are using different information about the environment to survive are likely to be characterized by a relatively balanced distribution of individuals across the various successful phenotypes. Thus, we use Shannon entropy, a popular metric of diversity that also measures evenness, to measure the diversity of the persistent genotypes and calculate the ecological metric. This metric is equivalent to D in evolutionary activity statistics (Bedau et al., 1997), assuming we chose genotypes reduced to informative sites as components.

Experimental Systems

We used two radically different experimental systems in order to ensure both that these metrics can be broadly applied and that they produce meaningfully consistent results.

NK Landscape

To begin a systematic examination of MODES metrics, we used a simple NK model (Kauffman and Levin, 1987). An NK model uses two parameters, N and K , to randomly generate a fitness landscape. N specifies the number of sites in the genome, each of which is a 0 or a 1. The fitness landscape specifies the effect of a given value at a given site on the fitness of the bit-string organism. This fitness effect depends on the values at the K subsequent adjacent sites. As such, K tunes the ruggedness of the landscape; low values of K produce smooth landscapes with few peaks, whereas high values produce landscapes with many peaks. We chose to use NK models because they are a well-understood system for studying general questions about evolutionary dynamics.

Experimental Treatments Our basic treatment used $N = 20$ (i.e., 20 bits in an individual) and $K = 3$ (the fitness contribution of each bit was influenced by three other bits). We used a population size of 200 and a per site mutation rate of .05, with tournament selection and a tournament size of 2. In addition to this baseline treatment, we tested the effects of eight experimental treatments: **High K** tests the effect of a highly rugged landscape ($K=10$) where fitness is effectively randomized whenever a mutation occurs. **High N** tests the effect of longer bit-string genomes ($N=100$; mutation rate was adjusted to 0.01 to keep the whole-genome mutation rate consistent with the base condition), allowing for a higher potential complexity. **Low Mutation** and **High Mutation** test the effects of more extreme mutation rates (0.005 and 0.1 respectively); we expect the mutation rate to be important for finding new areas of the fitness landscape and thus our novelty metric. **Small Pop** and **Large Pop** vary the population size (to 20 and 1000 respectively); in small populations we expect more drift in the population, allowing more change, while in a large population we expect stronger selection and consequently that a higher percentage of changes along the line of descent are beneficial. Finally, we included two special treatments: in **Oscillating Environment**, the fitness function was toggled between two pre-defined NK Landscapes every 500 generations, allowing us to see the effect of changing selective pressures where the populations was not able to stay on a single peak. In **Fitness Sharing** organisms that were too similar to each other detracted from each other's fitness, creating a pressure to explore multiple portions of the landscape at the same time and, ideally, maintain a high diversity. We used the fitness sharing equations described by Goldberg and Richardson (1987), with a sharing threshold of 50 and an α of 1. For all experiments, we used a filter length (t) equal to the population size.

Avida

The Avida Digital Evolution platform is a popular artificial life system for studying evolutionary dynamics (Ofria and Wilke, 2004). Avida consists of a population of self-replicating digital organisms with circular genomes composed of assembly-code-like instructions. Over the course of their lifetimes, organisms in Avida execute the code in their genome. The population is initially seeded with a single hand-coded organism that inefficiently copies itself and does nothing else. Each organism lives in its own cell in a toroidal grid. When an organism copies itself, its offspring is placed in a different cell, overwriting any previous occupant of that cell. Thus, there is pressure for individuals to reproduce quickly, before others copy over them. During the replication process, mutations are probabilistically introduced. Thus, the system contains inheritance, variation, and selection, causing evolution by natural selection to occur. Optionally, "tasks" can be added to the environment

in Avida. These are computational problems that organisms can perform for a reward in the form of additional CPU cycles that allow organisms to execute their code faster.

Experimental Treatments To understand how MODES metrics will behave in a full-featured artificial life system, we tested them in Avida under a variety of scenarios. For all experiments, we used a well-mixed population in order to speed up the expected rate of coalescence. All other parameters in Avida were left at their default values. We ran experiments in two different environments. The **empty** environment has no tasks - all evolution is focused entirely on optimizing the efficiency with which organisms can self-replicate. The **logic-9** environment, which has been used in many prior experiments (e.g. (Lenski et al., 2003a)), contains tasks for all non-trivial one- and two-input boolean functions.

Artificial life systems necessarily have constraints on the amount of time and memory we can give them. It is important in open-ended evolution research to determine whether these constraints are imposing practical limitations on the dynamics the system exhibits (Zaman, 2018). To do so, we ran experiments in each environment at three different population sizes: 500, 1000, and 2000. In each condition, we ran 30 replicate runs of Avida.

Additionally, to understand how sensitive our metrics are to the choice of the filter length (t) we conducted some additional experiments in the empty environment in which we varied t . In general, since our Avida runs are so long, we aim to filter neutral mutants out with our persistence filter, rather than just deleterious mutants. At each of the three population sizes, we tried t values of 500, 1000, and 2000. To ensure that we always have data from a filter length larger than population size, we also included a condition with a population size of 2000 and a t of 4000.

Implementation Details

If not implemented with care, these metrics can become computationally intractable in the context of the long experiments that open-ended evolution research often entails. In particular, RAM requirements can become prohibitive. We provide a few high-level approaches to mitigating these issues.

The largest memory cost is imposed by the novelty metric's requirement that we keep track of every taxon that has ever passed the persistence filter. Because we only need to know when we encounter a repeat taxon (rather than storing an archive of all taxa we have encountered), we can dramatically reduce this cost by using a Bloom filter (Bloom, 1970). Although this approach does introduce a (tunable) risk of false negatives (i.e. miss-classifying a novel taxon as not novel), this risk only makes the metrics more conservative.

The next largest cost is imposed by needing to keep track

of the phylogeny over time. In addition to standard phylogenetic pruning techniques (such as removing all taxa that do not have extant descendants), we can safely remove all taxa that died before the current generation minus t^4 . This optimization prevents the tree from growing without bound over the course of the experiment.

Lastly, it is helpful to be aware that increasing t will reduce computational demands by increasing the percentage of taxa that will be filtered out. With these optimizations, MODES metrics can be implemented with minimal overhead.

Statistical Methods

We assessed significance using Kruskal-Wallis tests followed by post-hoc Wilcoxon tests comparing each treatment to the baseline condition. To correct for multiple comparisons, we used a Bonferroni correction. Effect size measurements determine whether a treatment has a meaningful impact on a variable. Because standard deviations varied wildly among conditions, we used Glass's Δ as our measure of effect size (Hedges, 1981). As a general guideline for interpreting Δ , a value of 0.2 is generally thought to be low while a value of 0.8 is generally thought to be high (although this guideline is context dependent). All analyses and statistics for this paper were conducted using the R Statistical Computing Language version 3.4.4 (R Core Team, 2017) and the ggplot2 plotting library (Wickham, 2016). Distributions of final metric values are visualized using rain-cloud plots (Allen et al., 2018). Statistical code and supplemental statistical information is freely available (Dolson, 2018).

Code Availability

A C++ implementation of the MODES toolbox is available as part of the Empirical library (Ofria et al., 2018). The library is header only, and designed to be as easy to integrate into existing systems as possible. Code reliability is ensured with a suite of unit tests automatically run when code is added. As a proof of this concept, the Avida experiments presented in this paper were carried out using a lightly modified version of Avida that incorporated this implementation of our metrics (Bryson et al., 2018). All code used in this paper is open source and freely available (Dolson, 2018).

Results and Discussion

To ensure that these metrics are capturing the dynamics that we want them to, we tested them on a range of variants of our basic NK model and a range of conditions in Avida. The preliminary results for each metric are presented here.

⁴An important caveat is that this approach will only work with a strictly increasing unit of time. In many systems (including Avida) the average generation is not guaranteed to consistently increase. To support such systems, our implementation of the metrics allows for time to be tracked using two units at once, one corresponding to generation, and one that is guaranteed to be strictly increasing.

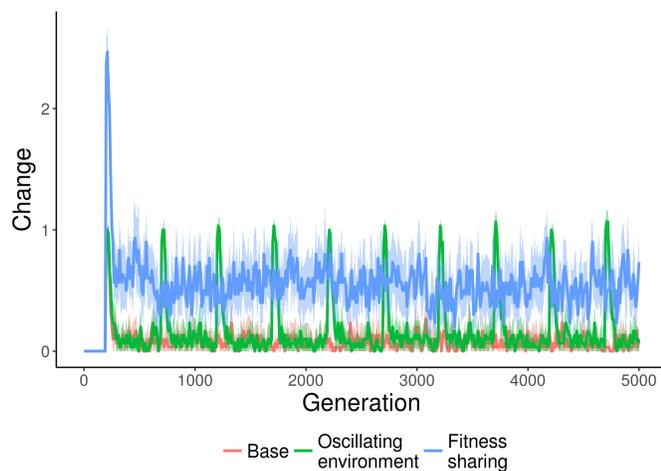


Figure 3: Amount of change at over time in varying NK Landscape environments. Fitness sharing increases the amount of change in the population over time. Conversely, a routinely changing environment leads to spikes in change that quickly drop as the population converges again. Shaded region represents a bootstrapped 95% confidence interval around the mean.

Change Metric

In the baseline and low mutation rate conditions for the NK Landscape, change is close to 0 (see Figures 3 and 4), indicating that our metrics are capable of detecting the stagnation typical of many genetic algorithms. As shown in Figure 3, several environmental changes increase the amount of meaningful change found in the NK Landscape populations over time. When organisms are forced to share fitness between others with the same genotype, the amount of change increases and remains higher than the baseline over time. Conversely, when the environment changes frequently, there is an initial spike of increased change that quickly drops back down to the baseline value.

The majority of environmental conditions we tested in the NK Landscape system produced dynamics over time qualitatively similar to the baseline treatment. In Figure 4 we show the amount of meaningful change in populations at the final time point in more environmental conditions. A higher mutation rate leads to increased meaningful change ($p < 0.0001$, Wilcoxon test; $\Delta = 0.80$) because mutations are necessary to create any meaningful change in this system. A smaller population size produces more meaningful change ($p = 0.004$, Wilcoxon test; $\Delta = 0.33$) because a small population cannot hold as many different genomes at one time and therefore there are more genomes that can arise that are different than what is in the previous population. Finally, fitness sharing produces increased meaningful change ($p < 0.0001$, Wilcoxon test; $\Delta = 0.63$) because it creates a constant pressure for the population to adapt away from

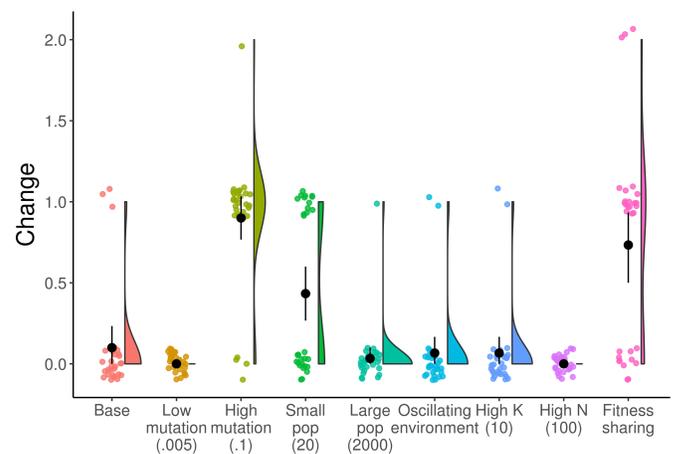


Figure 4: Rain-cloud plot of change at final generation across NK Landscape treatments. Environmental conditions that increase the amount of change at the final time point include: increasing the mutation rate, decreasing the population size, and implementing fitness sharing. Black circle and line indicate mean and bootstrapped 95% confidence interval.

whatever is the current dominant genotype.

In Avida, there is always at least a little change (see Figure 5). This observation is consistent with previous findings that fitness in Avida increases indefinitely (Wiser, 2015), as an increase in fitness implies both change and novelty. Based on coalescence theory, we would expect change in the empty environment to usually be less than or equal to one, because it is a single-niche environment. During each interval, either a fitter genotype will arise and sweep the population or the current fittest genotype will remain dominant. Because our value of t is not higher than the maximum expected coalescence interval, we should also expect to see the occasional time point with change greater than one. Our data are roughly consistent with this expectation. In addition, there is a subtle downward trend in the change data, likely due to the progressively increasing difficulty of finding beneficial mutations.

As expected, increasing the filter length, t , decreases the amount of change observed because fewer taxa are able to get through the filter (see Figure 5). In general, using a value of t equal to population size seems to be an adequate filter. The confidence interval for the mean of these conditions always overlaps 1, indicating that a substantial amount of filtering is occurring. Using lower values of t begins to lead to substantial increases in the variance of observed change. Using a higher value of t does further reduce noise, but with diminishing returns. In the empty environment, population size does not appear to have much impact on change, implying (unsurprisingly) that population size does not exert

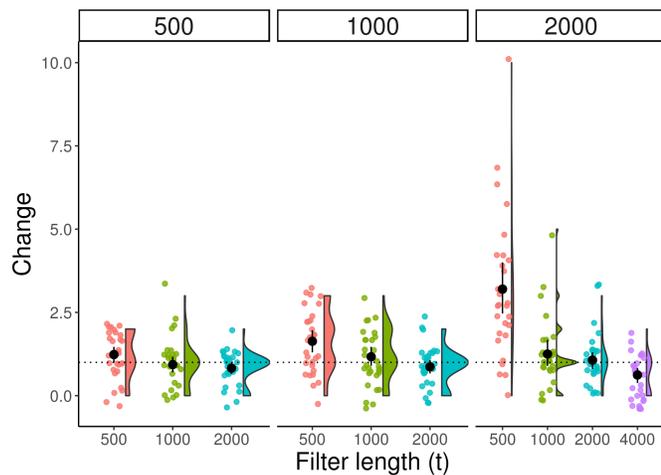


Figure 5: **Rain-cloud plot of change at final generation across multiple population sizes and filter lengths in Avida in the empty environment.** Labels along the top indicate population size. Black circle and lines indicate mean and bootstrapped 95% confidence interval. Horizontal bar indicates change = 1, the expected average change in the empty environment.

pressure on evolutionary dynamics in this environment.

In the logic-9 environment, however, there is a slight increase in change as population size increases, particularly early in the experiment (see Figure 6). Additionally, change is generally a little higher in the logic-9 environment than the empty environment. This distinction is an unexpected benefit of using a value of t too low to guarantee coalescence. Logic-9 is a single-niche environment, so if we chose a large enough value of t , we would expect change to always be less than or equal to 1. However, logic-9 is a more complex environment than the empty environment, which increases the odds that multiple lineages will be able to keep evolutionary pace with other for a substantial amount of time. As such, at the values of t we used, our change metric is able to reflect the fact that more is going on in the logic-9 environment than the empty environment.

While change is a metric often not considered in discussions of open-ended evolution, these results show that the amount of meaningful change can reflect differences in the environment and evolution of the populations and is likely a necessary dynamic for open-ended evolution. Our change metric responds in intuitive ways to variations in parameter settings, suggesting that it is a reliable indicator of the dynamics we designed it to capture.

Novelty Metric

As shown in Fig 7, a higher mutation rate increases the amount of novelty generated by the NK Landscape system. This result is to be expected, because more mutations make

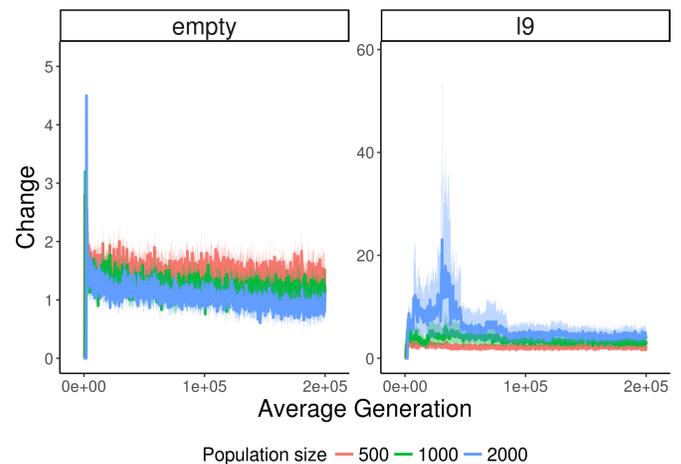


Figure 6: **Change over time across different environments and population sizes in Avida.** Note that y axes have different scales. In general, change is much higher in the Logic-9 environment. Filter length, t , is equal to population size.

it easier to cross fitness valleys and otherwise traverse the fitness landscape. Even at a high mutation rate, novelty does start to decrease over time as the search space is explored.

We again found that the majority of treatments had a qualitatively similar trajectory over time and therefore in Figure 8 we show only the final novelty value. As predicted, the baseline treatment has, with the exception of 1 replicate, stopped producing meaningful novelty by the final time point. Indeed, the only environment still reliably producing meaningful novelty is the high mutation rate condition ($p = 0.003$, Wilcoxon test; $\Delta = 0.3$). High mutation rates produce ongoing novelty by shifting the mutation-selection balance such that drift is able to preserve novel lineages for longer than t . Notably, fitness sharing does not increase the final novelty ($p = 0.333$, Wilcoxon test), presumably because it is promoting cycling among previously-discovered genomes.

Like NK Landscapes, Avida shows gradually declining novelty over time (although novelty in Avida declines more slowly) (Dolson, 2018). This trend presumably reflects the declining availability of beneficial mutations. Interestingly, the novelty and change graphs from Avida look almost identical, implying that there is little cycling among previously discovered solutions. This result, too, is consistent with the previously observed indefinite fitness increases in Avida (Wiser, 2015) - if the population is always getting fitter, genotypes from earlier in the run would be unlikely to survive in the present.

These results highlight the power of the novelty metric to identify environments and populations that have the potential to be open-ended due to the high number of new geno-

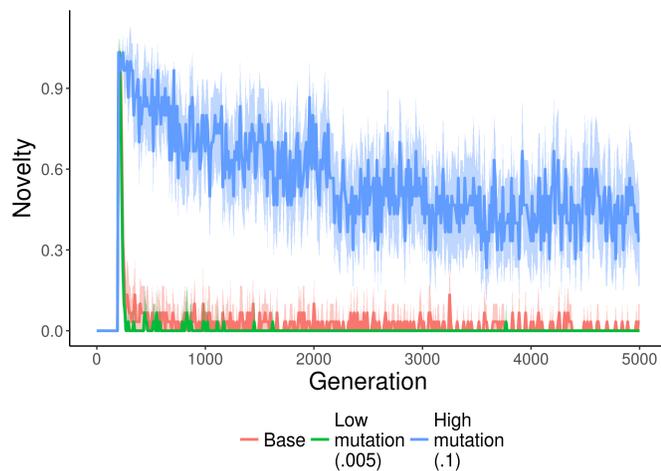


Figure 7: Amount of novelty over time in NK Landscape with varying mutation rate. The novelty metric measures the number of completely new meaningful genomes that have lineages that persisted since the previous time point. As the mutation rate increases, more novelty is continuously produced. However, at all mutation rates the novelty decreases over time. Shaded region represents a bootstrapped 95% confidence interval around the mean.

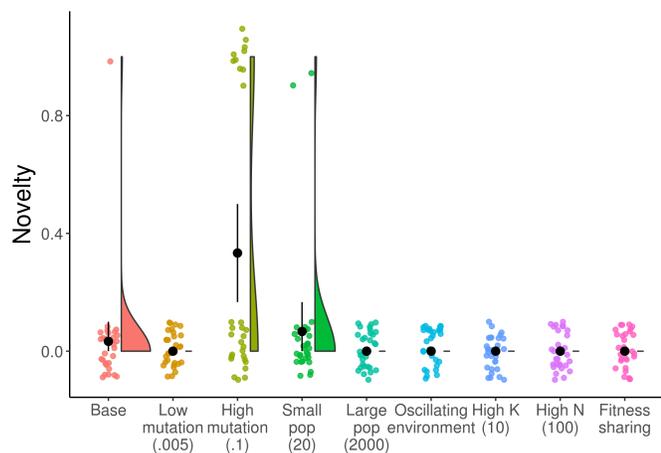


Figure 8: Rain-cloud plot of novelty at final generation across NK Landscape treatments. Black circle and lines indicate mean and bootstrapped 95% confidence interval. At the final time point, little meaningful novelty is found in our baseline populations. Only increasing the mutation rate increases novelty, implying that other conditions with high change are simply promoting cycling between a fixed set of genotypes.

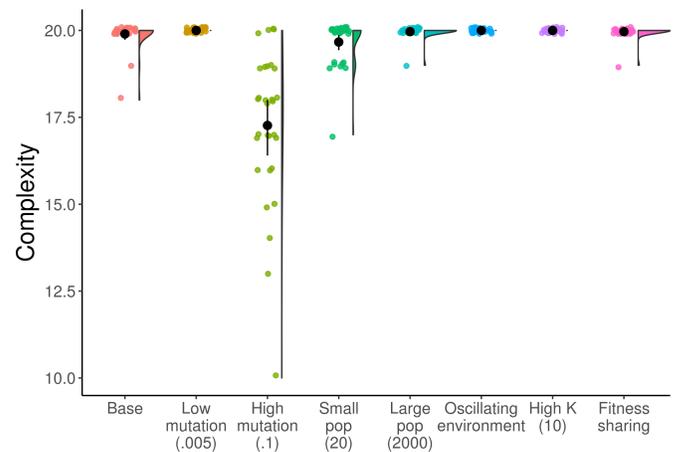


Figure 9: Rain-cloud plot of complexity at final generation across NK Landscape treatments. Black circle and lines indicate mean and bootstrapped 95% confidence interval. Note that we have excluded the high N condition from this graph because it throws off the axes. Most of the treatments reach the maximum complexity allowed by the genome length (20 or 100) and cannot continue to increase. High mutation rate and smaller populations decrease the final complexity achieved by the populations on average.

types being consistently discovered. Novelty is likely necessary, but not sufficient, for open-ended evolution because if nothing new is being produced by a population, neither the complexity nor the ecological metric can be non-zero.

Complexity Metric

In an NK Landscape populations, organisms cannot evolve to be more complex than N . As a result, the complexity increases over time and then saturates. High mutation rate ($p < 0.0001$, Wilcoxon test; $\Delta = -2.6$) and small population size ($p < 0.0001$, Wilcoxon test; $\Delta = -0.23$) reduce complexity because they shift mutation-selection balance to make staying on a fitness peak more challenging (see Figure 9). All other treatments were able to achieve maximal complexity fairly reliably. Note that, due to the rough nature of our complexity metric, this merely indicates the presence of an individual on a fitness peak. We cannot distinguish between higher and lower peaks. The high N treatment was, unsurprisingly, able to attain a drastically higher complexity due to its increased upper bound on complexity.

The NK Landscape results demonstrate that the complexity metric correctly identifies a system that is not able to continuously produce more complex solutions. Once the maximum complexity allowed by the genome length is reached, no higher value is possible.

In Avida, the complexity metric reveals a stark difference between the empty and logic-9 environments. In the empty

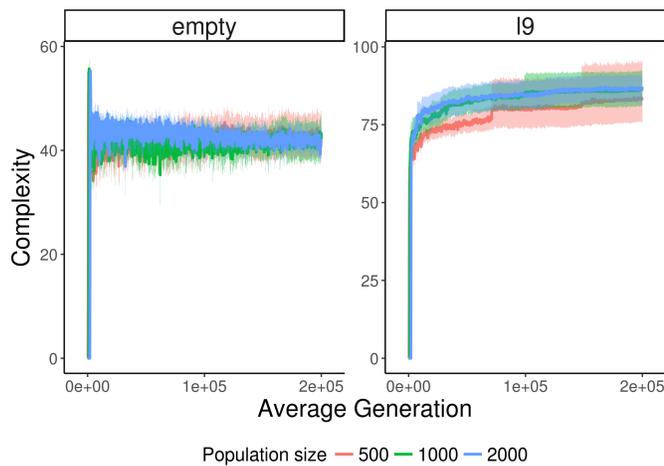


Figure 10: Complexity in Avida over time across different environments and population sizes. Note that y axes have different scales. In general, complexity appears to continue increasing in the Logic-9 environment, whereas it drops and then stabilizes in the empty environment. Filter length, t , is equal to population size.

environment, there is a rapid rise in complexity followed by a decrease and leveling out. This behavior is likely due to the strong pressure in Avida to become a more efficient self-replicator by optimizing code. Over time, simpler solutions are selected for, all else being equal. In the logic-9 environment, on the other hand, there is an ongoing upward trajectory in complexity. While logic-9 still rewards efficiency, algorithms that can make maximally efficient use of the tasks are complex. These results are consistent with other measurements of complexity in Avida over time (Adami et al., 2000).

In the biosphere, complexity appears to be growing without bound (Korb and Dorin, 2011), although there is debate over the mechanisms behind this process. As such, building a non-trivial system that exhibits such behavior is a worthwhile goal for open-ended evolution research. Unbounded growth in complexity is only possible in a system with a sufficiently complex environment such that there is always new information to be integrated into the genome.

Ecology Metric

Across both of our systems, the only condition that creates a multi-niche environment is the fitness sharing condition in NK Landscapes. Accordingly, that is the only condition in which we observe ecology significantly above the baseline ($p < 0.0001$, Wilcoxon test; $\Delta = 0.91$) (see Figure 11). Because fitness sharing specifically rewards organisms with less common genotypes, it promotes a stably high ecology value over time. This result demonstrates the trade-off inherent in fitness sharing because it leads to higher ecology at

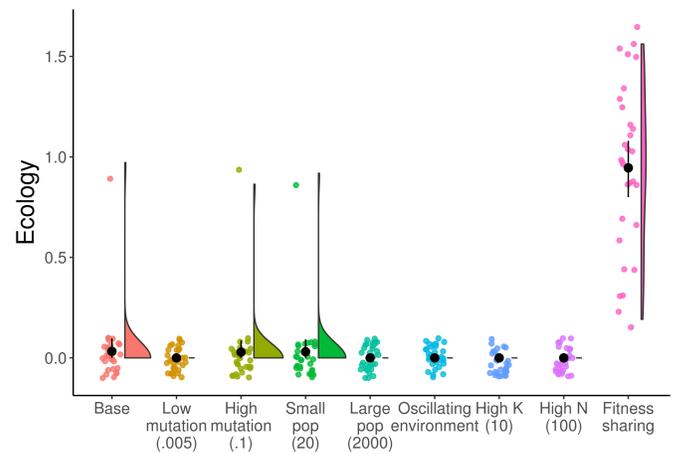


Figure 11: Rain-cloud plot of ecology at final generation across NK Landscape treatments. Black circle and lines indicate mean and bootstrapped 95% confidence interval. Fitness sharing is the only condition that reliably produces ecology.

the expense of lower complexity. Interestingly, in the Avida data, we see a consistent low level of ecology across all conditions (Dolson, 2018). We hypothesize that this slight increase over NK landscapes is due to the noise introduced by Avida's variable generation times.

Ecology is an incredibly powerful force in nature, leading to feedback cycles of ever-increasing diversity. Like complexity, diversity seems to be growing without bound in the biosphere (Harmon and Harrison, 2015). Thus, it will be important to see what mechanisms are important to promoting it in artificial life systems, too.

Conclusions

We have proposed a suite of metrics that quantify the presence of four generally-accepted hallmarks of evolution. These metrics build on prior work with evolutionary activity statistics and are largely compatible with them. Additionally, we have proposed techniques for reducing noise in these statistics. By testing them on two very different well-understood evolutionary systems, we have demonstrated that our metrics respond in an intuitive way to the dynamics these systems exhibit. Thus, these metrics should also be useful in understanding the extent to which novel systems exhibit hallmarks of open-ended evolution. Moreover, we can use them to understand the impact of incremental changes to a system. By breaking the seemingly monolithic problem of designing an open-ended evolutionary system into smaller, measurable pieces, we facilitate improved use of the scientific method. One of the primary goals in building an open-ended evolutionary system is to understand the underlying components that are necessary to do so. By measuring the

effects that controlled changes to a system have on this suite of metrics, we can more productively work toward these goals.

Going forward, it will be interesting to see how a wider variety of artificial life systems respond to the metrics in the MODES toolbox. In particular, further investigation on the differences between using a shadow run as a filter and using the persistence filter described here would be worthwhile. Ultimately, these two techniques capture sufficiently different information that it may be valuable to use each in turn.

For a long time, the field of open-ended evolution has been plagued by a lack of data that can be meaningfully compared across systems. We believe that the MODES toolbox will help remedy this problem by making useful metrics easily accessible. As new hallmarks of open-ended evolution are identified and new techniques of reducing noise are developed, we encourage contributions to the toolbox. By working together as a community of users, researchers, and developers we can dramatically increase the rate at which open-ended evolution research progresses.

Acknowledgments

We thank members of the Artificial Life community for feedback on the initial ideas for the MODES toolbox. We further thank the MSU Digital Evolution Lab for helpful feedback on this manuscript. This research was supported by the National Science Foundation (NSF) through the BEACON Center (Cooperative Agreement DBI-0939454), a Graduate Research Fellowship to ED (Grant No. DGE-1424871), and NSF Grant No. DEB-1655715 to CO. Michigan State University provided computational resources through the Institute for Cyber-Enabled Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or MSU.

References

- Adami, C., Ofria, C., and Collier, T. C. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, 97(9):4463–4468.
- Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., and Kievit, R. (2018). Raincloud plots: a multi-platform tool for robust data visualization. Technical Report e27137v1, PeerJ Inc.
- Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Foster, J. A., McMullin, B., Melo, V. V. d., Miconi, T., Spector, L., Stepney, S., and White, R. (2016). Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences*, 135(3):131–161.
- Bedau, M. A. and Bahm, A. (1994). Bifurcation structure in diversity dynamics. In *Artificial life IV*, pages 258–268.
- Bedau, M. A. and Packard, N. H. (1992). Measurement of Evolutionary Activity, Teleology, and Life. In *Artificial Life II*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 431–461. Addison-Wesley, Redwood City, CA.
- Bedau, M. A., Snyder, E., Brown, C. T., and Packard, N. H. (1997). A Comparison of Evolutionary Activity in Artificial Evolving Systems and in the Biosphere. In *Proceedings of the Fourth European Conference on Artificial Life, ECAL97*, pages 125–134, Cambridge. MIT Press.
- Bedau, M. A., Snyder, E., and Packard, N. H. (1998a). A classification of long-term evolutionary dynamics. In *Artificial life VI*, pages 228–237. Cambridge: MIT Press.
- Bedau, M. A., Snyder, E., and Packard, N. H. (1998b). A classification of long-term evolutionary dynamics. In *Artificial life VI*, pages 228–237. Cambridge: MIT Press.
- Bloom, B. H. (1970). Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426.
- Blount, Z. D., Barrick, J. E., Davidson, C. J., and Lenski, R. E. (2012). Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature*, 489(7417):513–518.
- Blount, Z. D., Borland, C. Z., and Lenski, R. E. (2008). Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 105(23):7899–7906.
- Bryson, D., Baer, B., Ofria, C., Barrick, J., Vostinar, A., Goldsby, H., Zaman, L., Chandler, C., Goings, S., Dolson, E., Vogel, M., Covert, A., Wright, G., Nahum, J., Misevic, D., Wagner, A., Rupp, M., Yilmaz, E., Pakanati, A., and Biswas, R. (2018). emilydolson/avida-empirical: Code for OEE Avida experiments at the time of submission. DOI: 10.5281/zenodo.1439479.
- Channon, A. (2001). Passing the ALife test: Activity statistics classify evolution in Geb as unbounded. In *Advances in Artificial Life*, pages 417–426. Springer.
- Channon, A. (2003). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. *Proceedings of Artificial Life VIII, Sydney, RK Standish, MA Bedau, and HA Abbass, (eds.), MIT Press: Cambridge, MA*, pages 173–181.

- Cooper, T. F. and Ofria, C. (2003). Evolution of stable ecosystems in populations of digital organisms. In *Artificial Life VIII: Proceedings of the Eighth International Conference on Artificial Life*, pages 227–232.
- Dolson, E. (2018). Supplemental information for The MODES Toolbox: Measurements of Open-ended Dynamics in Evolving Systems. DOI: 10.5281/zenodo.1442906.
- Dolson, E., Vostinar, A. E., and Ofria, C. (2015). What's holding artificial life back from open-ended evolution?
- Fu, Y.-X. and Li, W.-H. (1999). Coalescing into the 21st Century: An Overview and Prospects of Coalescent Theory. *Theoretical Population Biology*, 56(1):1–10.
- Goldberg, D. E. and Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization. In *Genetic algorithms and their applications: Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49. Hillsdale, NJ: Lawrence Erlbaum.
- Goldsby, H. J., Dornhaus, A., Kerr, B., and Ofria, C. (2012). Task-switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences*, 109(34):13686–13691.
- Good, B. H., McDonald, M. J., Barrick, J. E., Lenski, R. E., and Desai, M. M. (2017). The Dynamics of Molecular Evolution Over 60,000 Generations. *Nature*, 551(7678):45–50.
- Hara, A. and Nagao, T. (1999). Emergence of cooperative behavior using ADG; Automatically Defined Groups. In *Proceedings of the 1999 Genetic and Evolutionary Computation Conference*.
- Harmon, L. J. and Harrison, S. (2015). Species Diversity Is Dynamic and Unbounded at Local and Continental Scales. *The American Naturalist*, 185(5):584–593.
- Hedges, L. V. (1981). Distribution theory for glass's estimator of effect size and related estimators. *Journal of Educational Statistics*, 6(2):107–128.
- Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45.
- Korb, K. B. and Dorin, A. (2011). Evolution unbound: releasing the arrow of complexity. *Biology & Philosophy*, 26(3):317–338.
- Le Gac, M., Plucain, J., Hindr, T., Lenski, R. E., and Schneider, D. (2012). Ecological and evolutionary dynamics of coexisting lineages during a long-term experiment with *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 109(24):9487–9492.
- Lehman, J. and Stanley, K. O. (2011). Abandoning Objectives: Evolution Through the Search for Novelty Alone. *Evol. Comput.*, 19(2):189–223.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003a). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003b). The evolutionary origin of complex features. *Nature*, 423:139–144.
- Lenski, R. E., Rose, M. R., Simpson, S. C., and Tadler, S. C. (1991). Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *The American Naturalist*, 138(6):1315–1341.
- Lenski, R. E., Wisser, M. J., Ribeck, N., Blount, Z. D., Nahum, J. R., Morris, J. J., Zaman, L., Turner, C. B., Wade, B. D., Maddamsetti, R., Burmeister, A. R., Baird, E. J., Bundy, J., Grant, N. A., Card, K. J., Rowles, M., Weatherspoon, K., Papoulis, S. E., Sullivan, R., Clark, C., Mulka, J. S., and Hajela, N. (2015). Sustained fitness gains and variability in fitness trajectories in the long-term evolution experiment with *Escherichia coli*. *Proceedings of the Royal Society of London B: Biological Sciences*, 282(1821).
- Maddamsetti, R., Lenski, R. E., and Barrick, J. E. (2015). Adaptation, Clonal Interference, and Frequency-Dependent Interactions in a Long-Term Evolution Experiment with *Escherichia coli*. *Genetics*, 200(2):619–631.
- Maley, C. C. (1999). Four Steps Toward Open-ended Evolution. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 2*, GECCO'99, pages 1336–1343, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Ofria, C., Dolson, E., Lalejini, A., Fenton, J., Jorgensen, S., Miller, R., Moreno, M. A., Stredwick, J., Zaman, L., Schossau, J., Gillespie, L., CG, N., and Vostinar, A. (2018). emilydolson/Empirical: Code for Open-Ended Evolution at the time of submission. DOI: 10.5281/zenodo.1439475.
- Ofria, C. and Wilke, C. O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10(2):191–229.
- Potter, M. A. and Jong, K. A. D. (2000). Cooperative Coevolution: An Architecture for Evolving Coadapted Subcomponents. *Evolutionary Computation*, 8(1):1–29.

- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ribeck, N. and Lenski, R. E. (2015). Modeling and quantifying frequency-dependent fitness in microbial populations with cross-feeding interactions. *Evolution*, 69(5):1313–1320.
- Rozen, D. E. and Lenski, R. E. (2000). Long-Term Experimental Evolution in *Escherichia coli*. VIII. Dynamics of a Balanced Polymorphism. *The American Naturalist*, 155(1):24–35.
- Skusa, A. and Bedau, M. A. (2003). Towards a comparison of evolutionary creativity in biological and cultural evolution. *Artificial Life*, 8:233–242.
- Smith, J. M. and Szathmari, E. (1997). *The major transitions in evolution*. Oxford University Press.
- Soros, L. (2018). Necessary Conditions for Open-Ended Evolution. *Electronic Theses and Dissertations*.
- Soros, L. and Stanley, K. (2014). Identifying Necessary Conditions for Open-Ended Evolution through the Artificial Life World of Chromaria. In *ALIFE 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 793–800, New York. The MIT Press.
- Stanley, K. O. and Soros, L. B. (2016). The Role of Subjectivity in the Evaluation of Open-Endedness. In *OEE2: The second open-ended evolution workshop at Artificial Life XV*, Cancun, Mexico.
- 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., and Wisner, M. (2016). Open-Ended Evolution: Perspectives from the OEE Workshop in York. *Artificial Life*, 22(3):408–423.
- Tenaillon, O., Barrick, J. E., Ribeck, N., Deatherage, D. E., Blanchard, J. L., Dasgupta, A., Wu, G. C., Wielgoss, S., Cruveiller, S., Medigue, C., Schneider, D., and Lenski, R. E. (2016). Tempo and mode of genome evolution in a 50,000-generation experiment. *bioRxiv*.
- Turner, C. B., Blount, Z. D., and Lenski, R. E. (2015a). Replaying Evolution to Test the Cause of Extinction of One Ecotype in an Experimentally Evolved Population. *PLoS ONE*, 10(11):e0142050.
- Turner, C. B., Blount, Z. D., Mitchell, D. H., and Lenski, R. E. (2015b). Evolution and coexistence in response to a key innovation in a long-term evolution experiment with *Escherichia coli*. *bioRxiv*.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.
- Wiser, M. J. (2015). *An analysis of fitness in long-term asexual evolution experiments*. PhD thesis, Michigan State University.
- Wiser, M. J., Dolson, E. L., Vostinar, A., Lenski, R. E., and Ofria, C. (2018). The Boundedness Illusion: Asymptotic projections from early evolution underestimate evolutionary potential. *PeerJ Preprints*, (e27246v1). DOI: 10.7287/peerj.preprints.27246v1.
- Wiser, M. J., Ribeck, N., and Lenski, R. E. (2013). Long-term dynamics of adaptation in asexual populations. *Science*, 342(6164):1364–1367.
- Zaman, L. (2018). Investigating Open-Ended Coevolution in Digital Organisms. In *The 2018 Conference on Artificial Life*, pages 258–259, Tokyo, Japan. MIT Press.