

# Ecological Theory Provides Insights about Evolutionary Computation

Emily Dolson (dolsonem@msu.edu), Wolfgang Banzhaf, and Charles Ofria

**Abstract**—Evolutionary algorithms often incorporate ecological concepts to help maintain diverse populations and drive continued innovation. However, while there is strong evidence for the value of ecological dynamics, a lack of overarching theoretical framework renders the precise mechanisms behind these results unclear. These gaps in our understanding make it challenging to predict which approaches will be most appropriate for a given problem. Biologists have been developing ecological theory for decades, but the resulting body of work has yet to be translated into an evolutionary computation context. This paper lays the groundwork for such a translation by applying ecological theory to three different selection mechanisms in evolutionary computation: fitness sharing, lexica selection, and Eco-EA. First, we use ecological ideas to establish a framework that clarifies how these selection schemes are alike and how they differ. We then build upon this framework by using metrics from ecology to gather empirical data about the underlying differences in the population dynamics that these approaches produce. Specifically, we measure interaction networks and phylogenetic diversity within the population to explore long-term stable coexistence. Notably, we find that selection methods affect phylogenetic diversity differently than phenotypic diversity. These results can inform parameter selection, choice of selection scheme, and the development of new selection schemes.

## I. INTRODUCTION

EVOLUTION and ecology are fully entwined in nature [1]. As such, evolutionary theory and ecological theory need to build upon each other to realize their full potential. Here, we argue that ecological theory is similarly important in the context of the theory of evolutionary algorithms. We can better understand much of what happens in evolutionary algorithms if we reformulate our analyses with ecological theory in mind. In particular, concepts such as diversity maintenance, species co-existence, niches, and spatial effects are suddenly at the forefront of concerns in such a perspective. While the No-Free-Lunch theorem states that there is no single best algorithm for solving all possible optimization problems [2], progress is possible because real-world problems have patterns that can be exploited, which tend to cluster into groups with similar properties. Algorithm can solve problems from across these classes by simultaneously exploring qualitatively different paths through the solution space. Such a technique effectively produces multiple “species” of solutions that coexist in an ecological community. Thus, by leveraging ecological theory, we should be able to use first principles to design evolutionary algorithms that take maximal advantage of this property and

mitigate the practical ramifications of the No-Free-Lunch theorem.

Evolutionary computation researchers have devoted substantial attention to understanding how to promote coexistence among lineages exploring different regions of a fitness landscape [3]–[6]. Promoting diversity in an evolving population is important for EC because it reduces premature convergence on suboptimal fitness peaks while still encouraging both exploration and exploitation. However, some types of diversity facilitate finding global optima better than other types. For example, a high mutation rate generates more new genotypes, but this increased exploration sacrifices exploitation of promising prospective solutions; the population plows ahead exploring rather than refining the solutions already found.

Even amongst more advanced diversity-promoting approaches, some produce forms of diversity that are spread across the fitness landscape in ways that appear to be more or less conducive to solving a given problem (see, for example, the difference between the probe and behavioral methods in [5], or the different evolutionary potential observed at equivalent diversity levels in [7]). Although we have a high-level idea of why many techniques promote diversity, the mechanistic details by which these populations spread across the fitness landscape are less well understood. For example, fitness sharing [3] promotes diversity via negative density dependence. Does this diversity represent qualitatively different portions of the fitness landscape than, for example, lexica selection [8]? The answer to this question depends on subtle differences in the evolutionary pressures that these different selection schemes place on the evolving population.

Because techniques for promoting diversity rely on creating fitness interactions between individuals in the population (beyond basic competition for space in the next generation), they are, by definition, creating simple ecologies. Ecologists have developed rigorous theory to predict how ecological communities change over time and are exploring their interactions with evolutionary dynamics. A particularly popular area of study concerns the conditions under which long-term stable coexistence between different species is possible [9]–[12]. Such a theoretical framework can be applied to understand stable coexistence in EC populations as well, ultimately allowing us to determine dynamic properties of the system. Of particular interest, such theory should facilitate the prediction of which pathways can simultaneously be traversed along a given fitness landscape. By understanding this interaction between diversity maintenance strategies and the fitness landscape, we can predict which algorithmic techniques will be the most effective *a priori*. These insights should apply both to choosing an

All authors are with the Department of Computer Science and Engineering and the BEACON Center for Evolution in Action at Michigan State University. ED and CO are also affiliated with the Ecology, Evolutionary Biology, and Behavior program at Michigan State University.

appropriate diversity maintenance technique and to setting its parameters. Additionally, an improved mechanistic understanding of why existing algorithms work should facilitate building more effective variations of those algorithms.

There are two main sets of tools from ecology that we expect can be helpful for analyzing evolving EC communities: mathematical theory and empirical techniques to evaluate that theory. Ecologists use mathematical theory to make *a priori* predictions about the fate of natural communities. However, translating these predictions to work with EC systems requires careful attention to implicit assumptions about nature that may or may not carry over. For example, ecologists can safely assume that organisms inhabit a finite three-dimensional Euclidean space. A more substantial problem is that most ecological theory assumes that evolution is too slow to be relevant and attempting to introduce evolution can create messy feedback loops (although eco-evolutionary dynamics research is bridging this gap [13]). Another important limitation is that equations in ecology generally calculate the average or expected behavior of a system. Such results can be misleading in highly contingent processes, such as evolution, where rare mutations can redirect a population into a new region of the fitness landscape. Nonetheless, ecological theory is a useful starting point for predictions. As such, we will lay groundwork here for using it to guide development of EC systems.

Empirical measurements complement ecological theory, allowing biologists to assess predictions and refine theoretical frameworks. Furthermore, empirical measurements allow us to uncover general patterns even in situations where theory is lacking [14]. To this end, ecologists have developed a toolbox of techniques to evaluate the diversity of a community. Some of these metrics, such as richness (number of unique species) and Shannon diversity (entropy) are already used in EC. Other valuable metrics have not yet been adopted in EC. Phylogenetic diversity [15], for example, can assess the extent to which an algorithm is maintaining independent subpopulations that are exploring distinct regions of the fitness landscape vs. organisms that only recently diverged from a common ancestor [14]. This distinction can impact how useful the population diversity is likely to be for adaptive evolution. We can also use empirical ecological measurements to assess hypothesized mechanisms for different diversity maintenance techniques. For example, ecologists often build graphs representing the pairwise interactions between community members; the topology of these graphs signals how the community is likely to change over time [16]. We can do the same for interacting communities in EC.

In the rest of this Introduction, we will provide additional background on ecological theory and the selection schemes that we will be examining. In the rest of this contribution, we will examine four illustrative selection schemes in the light of ecological theory. Our goal is to develop unifying principles for understanding the ecological dynamics inherent in an EC system and predicting its resulting adaptive potential. We intend for this theoretical framework to provide powerful insights about the dynamics in EC systems and lay the groundwork for easier application of ecological ideas in the future.

### Box 1: Glossary

**Biological fitness ( $W$ ):** An individual's expected number of offspring.

**Equalizing forces:** Forces that increase the similarity of fitnesses in a population.

**Interaction network:** A weighted directed graph describing interactions between members of a population.

**Phylogenetic diversity:** The amount of diversity in ancestry represented within a population.

**Phylogeny:** The graph of parent-child relationships in a population.

**$R^*$ :** The lowest quantity of a limiting resource for which a given species' population growth rate is positive. In an EC context, can be thought of as the lowest quantity of a limiting resource for which attempting to use that resource increases fitness.

**Stable coexistence:** A scenario in which we expect a group of species to coexist indefinitely.

**Stabilizing forces:** Forces that increase the fitness of rare species.

### A. Ecological theory

One way that ecological theory can inform EC is to identify conditions under which different types of organisms can or cannot **stably coexist**. The ability to predict coexistence dynamics will inform us about types of lineages that can simultaneously explore the fitness landscape. As such, determining coexistence criteria is a primary focus of this paper.

Early theory on coexistence in ecology focused on competition for resources that are both *limited* in quantity and *limiting* of the growth rate of species that rely on them. The most important value for determining coexistence in this context is a species'  $R^*$ , the resource availability level at which the species' population growth rate is 0. If the current resource availability is less than a species'  $R^*$ , that population will shrink, reducing the utilization of the resource and (in the absence of other species) increasing its availability. Conversely, if the current resource availability is greater than the  $R^*$  for a species, its population will grow. Assuming no other ecological factors are affecting a species abundance, its population size will stabilize with a resource abundance equal to  $R^*$ .

In the simplest case of species competing for a single resource, the species with the lowest  $R^*$  should out-compete the others [17]. That species' population will continue to increase, depleting the resource until it reaches the species'  $R^*$ ; meanwhile, the resource availability will dip below the other species'  $R^*$ , causing their populations to decrease until they finally go extinct, having been out-competed. Adding additional resource types introduces the potential for stable coexistence among multiple species if each species is a better competitor for a different resource and consumes more of that resource [11] (summarized in [12]). Note that we use the word "species" here to be consistent with ecology; the same insights apply to other taxonomic units, such as phenotypes

or genotypes. In most cases, these are more appropriate taxonomic units in the context of EC.

This resource-mediated coexistence effect is an instance of a broader rule: species can coexist if individuals of each species compete with each other more than they compete with individuals of other species [10]. This rule works because it forces species to be self-limiting, creating negative frequency dependent dynamics where each additional member of a species reduces the competitive ability of other individuals of the same species [18]. The magnitude of difference between interspecific (between species) and intraspecific (within species) competition that is required to enable long-term stable coexistence is determined by the difference in the fitness of the two species. If one species is dramatically more fit, it can drive the other species to extinction even with limited competition [10]. Ecologists draw a distinction between “**stabilizing**” dynamics, which alter the ratio of interspecific competition and intraspecific competition, and “**equalizing**” dynamics which alter the difference in fitness between two species [18]. In the absence of any stabilizing dynamics, equalizing dynamics will lead to an unstable equilibrium; even if two species have identical fitness, one should eventually drift to extinction. Stabilizing dynamics, on the other hand, actively correct any deviation from equilibrium. Note that in biology, “**fitness**” is strictly a measurement of reproductive output. This definition is in contrast to the externally-defined “fitness functions” used in EC to determine competitive advantage.

If we consider more heterogeneous environments, there is another mechanism for species coexistence: spatial segregation. Such segregation can arise due to a variety of factors. The simplest of these factors is the presence of physical barriers that inhibit movement between regions. In biology, physical barriers play an important role in facilitating diversification [19], [20], and these benefits carry over to EC in the context of island models [21]. Spatial segregation can also be brought about by environmental conditions. For example, one region might require tolerance to extreme heat, while another might require tolerance to acidic soil. In this case, species would only compete when both could inhabit the same regions. The resulting traversal (over evolutionary time) of multiple environments has a profound effect on the traversal of the fitness landscape [22]. An even weaker form of spatial segregation occurs when species are capable of surviving in range of environments but are best optimized to a specific region [23].

What do these coexistence dynamics mean for EC? First, when choosing parameters, we should consider the circumstances under which they promote coexistence. When two independent lineages are traversing the fitness landscape, what genotypic or phenotypic differences are required for them to both stably persist? Second, when designing selection schemes, we should determine how many lineages we want to co-exist and how frequently we want lineages to turn over. If it is costly to stochastically lose established lineages, we should consider including stabilizing dynamics in addition to or in place of equalizing dynamics. Finally, we should give careful thought to the metaphors that we use to compare EC to biological populations; determining whether a given selection scheme is more akin to competition for resources, spatially




Fig. 1: Conceptual illustration of phylogenetic diversity metrics. The letters around the outside represent the full set of extant taxa and the circle in the middle indicates their most recent common ancestor. Branching points indicate the locations of intermediate taxa (note that in biology these have to be inferred but in EC we have perfect information). The three panels indicate the three different facets of phylogenetic diversity: richness, divergence, and regularity. Reproduced from [24].

segregated habitats, or something else all-together will make it much easier to draw parallels and gain a deeper understanding. If we can show that a system in EC is isomorphic to a system in biology, we can rapidly import the vast wealth of insights from biological research into how that system behaves.

## B. Empirical ecological techniques

Empirical ecologists seek to find general patterns in complex, messy data. While data from EC may be less noisy than field data, the interconnections within any ecological community are still complex. Ecologists have developed many tools for extracting meaning from this complexity, which may be equally helpful in trying to understand the fine-scale dynamics of EC systems. Here we focus on two approaches: phylogenetic analysis and interaction networks.

1) **Phylogenetic analysis:** Phylogenetic analysis refers to a suite of metrics that are used to quantify the topology of a populations’ **phylogeny** (ancestry tree) [15]. In biology, these trees generally need to be inferred from extant species, but in EC we can record the actual tree as it forms. Many of these methods measure the amount of evolutionary history discernible in a population. For example, evolving sub-populations that have stably coexisted for a long time will reflect deeper evolutionary history than a single population that is all descended from a recent common ancestor. Even if the latter population contains many unique phenotypes, they will all be relatively close to each other on the fitness landscape. As such, the population is likely exploring only one basin of attraction within the fitness landscape at a time, limiting the rate of adaptation. **Phylogenetic diversity** provides more direct evidence about the efficacy of diversity maintenance techniques than the more commonly used measures of genotypic and phenotypic diversity.

Techniques for analyzing the topology of phylogenies can be split into three categories [24] (see Figure 1 for an illustration of each):

- 1) **Richness:** the total quantity of evolutionary history represented

- 2) Divergence: how spread out the population is in phylogenetic space
- 3) Regularity: how evenly the population is divided across evolutionary space

Here, we focus on richness and divergence, as our goal is to analyze selection schemes, which regularity does not offer clear insight into. A number of metrics quantify phylogenetic richness and phylogenetic divergence (summarized in [15] and [24]). Here, we measure phylogenetic richness with the original phylogenetic diversity metric (simply referred to as “phylogenetic diversity”) [19], as the count of nodes in the minimal spanning tree connecting each extant genotype to the most recent common ancestor of all extant genotypes. We measure divergence as the mean pairwise distance between genotypes in the population [25].

Note that in most cases biologists implicitly assume that trees are rooted at the most recent common ancestor of all taxonomic units being compared. This assumption is unavoidable in the context of biology, because phylogeny reconstruction techniques cannot make inferences about anything preceding the most recent common ancestor. In EC, we have the full history. However, including it would not add additional information; phylogenetic diversity would increase by a constant for each member of the population and mean pairwise distance would not change. Note that in this paper we calculate these metrics on a per-genotype basis, but they can also be calculated per-individual or per-phenotype.

2) *Interaction networks*: Since ecological communities are collections of interacting organisms, drawing a graph representing the network of pairwise interactions is a useful technique for understanding them [16]. The topology of this graph illustrates the selective pressures that organisms place on each other (see Figure 2). Ecologists often generate these networks from knowledge of the species involved such as the fact that two species use the same limited resource and thus must compete with each other. Of course, in EC we can directly measure the fitness of each member of the population with and without each other member present to identify interactions.




Fig. 2: Example **interaction network** with three nodes. Red arrows denote harmful interactions and blue arrows denote beneficial interactions. Line width and color darkness indicate magnitude of effect. In this example, A reduces B's fitness a lot, B reduces A's fitness a little, and B reduces C's fitness a lot. A increases C's fitness a lot. A plausible mechanism for this positive interaction is that, by harming B, A reduces B's harmful effect on C.

### C. Selection schemes

For the purposes of this paper, we have chosen a subset of selection schemes from the vast range of diversity maintenance techniques: fitness sharing [3], lexica selection [8], and Eco-EA [26]. We compare these to each other, and to standard tournament selection, as a control. Specifically, we have selected a set of approaches that clearly map onto ecology. In the future, we plan to expand this framework to include a wider variety of diversity maintenance techniques. All of the selection schemes summarized below have their population size and mutation parameters held constant. Here we identify parameters unique to each system.

1) *Tournament selection*: As tournament selection is one of the most popular selection techniques and does not create an ecology, we use it as a control. In its simplest form, tournament selection has one distinct parameter,  $T$ , the number of individuals to randomly pick from the population (with replacement) for each tournament. The fittest individual in the tournament is then chosen to reproduce. An independent tournament is run to fill each position in the population.

2) *Fitness sharing*: Fitness sharing was an early use of ecology to promote diversity in EC [3]. In fitness sharing, the fitness of every member of the population is reduced in proportion to the density of similar individuals (i.e. they “share” their fitness). Specifically, the extent to which each organism needs to share its fitness is calculated by summing the sharing equation,  $sh(d)$ , over the population:

$$sh(d) = \begin{cases} 1 - (\frac{d}{\sigma_{share}})^\alpha & d < \sigma_{share} \\ 0 & d \geq \sigma_{share} \end{cases} \quad (1)$$

In this equation,  $d$  is the distance (defined in terms of genotype or phenotype) between the individual that is having its fitness calculated and the other member of the population that it is being compared to.  $\alpha$  is a parameter that tunes the relationship between the similarity of individuals and how much they compete with each other.  $\sigma_{share}$  (the sharing threshold) specifies how similar individuals need to be to compete with each other at all. For each individual,  $Sh(d)$  is summed across the population and that individual's fitness is divided by the sum.

3) *Lexica selection*: In lexica selection, solutions are evaluated on a large number of criteria [8]. Traditionally, these criteria are test cases, but other types of fitness functions are also effective [27]. To choose an individual to reproduce, the selection criteria are applied to the population in a random order. Only the individuals that perform best on each criterion move on to be evaluated on the next, until only one individual remains. In case of a perfect tie, the winner is selected randomly from the remaining options.

Lexica selection is highly effective at solving challenging problems in genetic programming [28], [29] while maintaining diverse populations [30]. Prior work on the population dynamics of lexica selection has noted the surprising prevalence of “hyperselection events,” occasions in which a single individual is selected as a parent for the vast majority of the next generation [31]. Given lexica selection's success at maintaining diversity, the fact that such events occurred at all



was unexpected. Further research, however, has suggested that these events are not important to lexibase selection's success at solving challenging problems and maintaining diversity [32]. Additional case studies of phylogenetic trees generated by lexibase selection suggest that lexibase selection promotes the existence of sub-populations that are focused on a specific selection criterion [33].

4) *Eco-EA*: Eco-EA is an approach to solving complex problems by associating limited resources with simpler challenges that can be used as stepping stones [26], [34], [35]. These challenges can be sub-problems of the larger problem, individual test cases, or other related tasks that may be used as components of a global solution. Resources remain plentiful until a solution to the corresponding sub-problem is discovered. When a population first solves a sub-problem, the large amount of resource dramatically boosts the fitness of individuals that use it. As the number of individuals using the resource increases, it become less valuable. As long as there is a cost to attempting to use each resource, negative frequency dependence fosters small, but stable sub-populations capable of solving each sub-problem. For problems without such tradeoffs, an explicit cost can be introduced [27]. These sub-population are available to be co-opted into solutions for harder challenges, including the global problem.

## II. ECOLOGY IN EC SYSTEMS

In this section, we discuss how each of our chosen EC selection schemes can be intuitively mapped onto ecology. We believe that these metaphors will facilitate a two-way flow of ideas between fields. With the aid of these metaphors, we consider how to calculate two ecologically-important factors: an individual's expected number of offspring (i.e. biological fitness,  $W$ ) and the conditions under which stable coexistence is possible. The former is critical for calculating interaction networks within ecological communities (see Section 3.1) and the latter provides insight into not only the amount of diversity that a selection scheme promotes, but how the population is situated across the fitness landscape.

### A. Tournament selection

Individuals in tournament selection have no interaction beyond competition for space in the next population. Since this competition affects all individuals equally, tournament selection approximates a condition where there is no ecological interaction. Nature lacks comparable situations, as ecology is nearly ubiquitous outside of carefully controlled experimental environments (although evolutionary theorists do often make this simplifying assumption). The absence of ecology in tournament selection means that stable coexistence is impossible in the long term.

Calculating the biological fitness,  $W$ , of an individual in tournament selection requires determining the proportion of the population with a lower fitness score than that individual's,  $p_{\text{less}}$ . Based on this number, we can calculate the probability of the focal individual winning a tournament. We must also account for ties based on  $p_{\text{equal}}$ , the proportion of the entire population with a fitness score equal to that of the focal

individual (including that individual itself, since we select tournament members with replacement). Thus, the fitness of an individual is the number of tournaments it is expected to compete in ( $T$ ) times the probability of that individual winning a tournament:

$$W = T \times \sum_{i=1}^T \left( \frac{p_{\text{equal}}^{i-1}}{i} \times p_{\text{less}}^{T-i} \right) \quad (2)$$

Note that the summation steps through all possible  $i$ , the number of individuals in the tournament with fitness equal to that of the focal individual (at minimum 1, the focal individual itself, and at maximum the entire tournament). Thus, we enumerate all possible combinations of individuals with less or equal fitness. Note that any tournament won by the focal individual must not contain individuals with a higher fitness.

### B. Fitness sharing

Fitness sharing operates on the basic ecological assumption that individuals compete more against others that they are more similar to. The most intuitive ecological scenario to compare fitness sharing to is one where a population consumes a single, continuously-varying resource. In cases where the distance function is calculated over a discrete number of dimensions (e.g. Euclidean distance between two vectors of equal length), the resource can be thought of as varying over the same number of dimensions. An analogous situation in ecology is that of Darwin's finches [36]. Many of these finches eat seeds, which vary along dimensions such as size and shape. Each species of finch has a beak that is best adapted to eat seeds near a target size and shape. Over evolutionary time, the finch species have partitioned the space of possible beak morphologies into stable niches specialized on different seed types. In the context of fitness sharing, we expect the population to partition the space of possible phenotypes into stable niches. The theory of limiting similarity suggests that these niches should be somewhat separated from each other in genotypic/phenotypic space [9]. Fitness sharing results suggest that these niches are often associated with peaks in the fitness landscape [3].

Deb and Goldberg have already established the criteria for stable coexistence in fitness sharing [37], which we find to be mathematically identical to Chesson's predictions under modern coexistence theory [10], [38]. For the algebra demonstrating this equivalence, see supplemental information. This coexistence criterion can be summarized by the following equation:

$$\rho \leq \frac{k_1}{k_2} \leq \frac{1}{\rho} \quad (3)$$

where  $k_1$  and  $k_2$  are the fitness function scores of the genotypes being evaluated and  $\rho$  is the niche count between these two genotypes, as calculated by Equation 1. We chose these symbols to be consistent with Chesson's [10]. Note that in Chesson's framework,  $\rho$  generalizes to be the amount of niche overlap between the two individuals. The intensity of stabilizing dynamics (i.e. negative frequency dependence) can

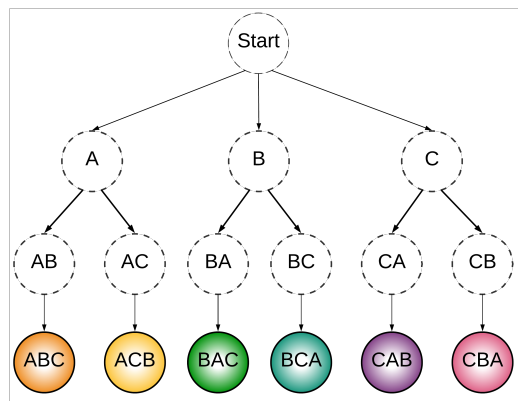


Fig. 3: Depiction of a Lexicase decision tree. A, B, and C are test-cases, and the internal nodes represent the order in which they are picked. Each possible path through the tree leads to a different leaf node ("island"). The individuals on that "island" are the ones that have the potential to be selected for reproduction.

be calculated as  $1 - \rho$ . The smaller  $\rho$  is, the greater the difference in fitness values it is capable of stabilizing.

The biological fitness,  $W$ , of individuals in fitness sharing can be calculated by modifying fitness scores based on Equation 1 and then applying Equation 2 based on the adjusted scores.

### C. Lexicase selection

In ecological terms, lexicase selection creates a vast number of niches nested within each other. While it is tempting to use a resource metaphor to understand the competition within these niches, the concept of resources does not clearly map onto lexicase selection. Whereas using a resource harms all other individuals that use that resource, improving on a selection criterion in lexicase selection only harms a (usually small) subset of the population. Instead, we argue that population structure is a more apt metaphor.

Imagine  $N!$  islands of equal size, where  $N$  is the number of selection criteria. Each island corresponds to a single potential ordering of selection criteria and can only be inhabited by the individuals that are best at that ordering (see Figure 3). This arrangement is analogous to situations in nature where niches are defined exclusively by an organism's ability to survive a set of harsh abiotic conditions. Being better able to survive these conditions increases the number of offspring an individual can have. Over time, genotypes that are better at surviving in a given set of conditions competitively exclude those that are worse at surviving there. This competitive exclusion happens more rapidly in lexicase selection than it usually would in biology, but the principle is the same.

While multiple genotypes can theoretically inhabit the same island if they are phenotypically identical, this coexistence will be unstable due to the lack of stabilizing dynamics to increase




Fig. 4: The probability of long-term survival under lexicase selection across different population sizes ( $S$ ) and lengths of evolutionary time ( $G$ ).

the populations of rare genotypes on an island. Eventually, all but one should stochastically go extinct. Furthermore, much as an island in nature can experience a random catastrophic event, (e.g., a volcanic eruption), the stochastic nature of lexicase selection means that in every generation there is a chance that an island will get unlucky and not be selected. The probability of such an event progressively increases as the proportion of the population living on each island ( $\frac{\text{population size}}{N!}$ ) decreases. When the total population size is less than the number of islands, genotypes must inhabit multiple islands to survive. Thus, stable coexistence in lexicase selection depends on whether a genotype's range (i.e. the proportion of islands it can occupy) is large enough to survive in the long term.

The proportion of islands a genotype can occupy is equivalent to its probability of selection. Just as species in nature face greater risk of stochastic extinction if their range (the area they inhabit) is too small, so too do genotypes in lexicase selection. Details for calculating this quantity are described in equation 4 of [39]. To convert the probability of selecting an individual under lexicase selection to biological fitness, we simply need to multiply it by the number of selection events that will occur:

$$W = S \times P \quad (4)$$

where  $S$  is population size (the number of selection events per generation) and  $P$  is the proportion of islands occupied. What does this condition mean for stable coexistence in lexicase selection? The chances of a genotype with fitness  $W$  surviving for  $G$  generations in a population of size  $S$  are given by the equation:

$$P(\text{survival}) = (1 - (1 - W)^S)^G \quad (5)$$

Plotting this function for various parameter values, we see that there is generally a cut-off where the probability of survival rises abruptly from approximately 0 to approximately 1 (see Figure 4). The value at which the transition occurs is effectively the minimal percentage of islands that a genotype must occupy in order to be expected to survive in the long term. This value is closely related to the concept of a minimum viable population in biology, and is useful to consider when making decisions about how large of a population size to use.

#### D. Eco-EA

Eco-EA is analogous to a traditional resource competition scenario in ecology. All individuals occupy the same region of space and compete with each other by targeting the same resources. We can calculate the threshold amount of a resource in the environment needed by each individual to benefit from using it:

$$R^* = \frac{\text{cost}}{C_f * \text{score}} \quad (6)$$

where *cost* is the cost of performing the corresponding task,  $C_f$  is the fraction of the available resource consumed, and *score* is how well the individual performed the task (normalized to 1). In terms of resource competition theory, this value is effectively that individual's  $R^*$  for that resource.<sup>1</sup> Note that a higher *score* reduces  $R^*$ , meaning that the individual can benefit from the task even if less of the resource is available, making it a better competitor for that resource. Thus, as in natural ecosystems, the individual with the lowest  $R^*$  can out-compete individuals with a higher  $R^*$ .

If a group of individuals use completely different resource types from each other, there will be a high degree of stabilization, meaning they should coexist under a wide range of fitness differences. The more complex coexistence scenario occurs when two individuals compete for the same resources. Ecologists distinguish resources that can be used interchangeably with each other ("substitutable resources") from resources that cannot be used interchangeably ("essential resources"). Resources in Eco-EA are substitutable, because an equivalent fitness gain could be achieved using either one (up to a point); one just requires more resource to do so. In this scenario, two individuals can stably coexist as long as they have lower  $R^*$ s for opposite resources, and each one consumes more of the resource it has a lower  $R^*$  for. Because both  $R^*$  and the amount of resource consumed are determined by the individual's score on the task associated with the resource, the latter criterion will always be met. Thus, as long as neither individual has the lowest  $R^*$  for every resource, coexistence should be possible. In EC terms, this requirement boils down to each individual being Pareto dominant. As before, coexistence also requires that the individuals have somewhat similar fitnesses. The amount of stabilization (and thus the maximum fitness difference) can be calculated based on the resource consumption of each species [12].

Biological fitness,  $W$ , can be determined in the same way as in fitness sharing: first, calculate the adjusted fitnesses based on resource use, then use Equation 2.

#### E. Summary

As we would expect, selection schemes that successfully maintain diversity allow for long-term stable coexistence. For fitness sharing and Eco-EA, the criteria are described by

<sup>1</sup>This definition is subtly different from the traditional definition of  $R^*$ , which states that  $R^*$  is the lowest value of resource for which a species' per capita population growth is not negative [11]. We frame  $R^*$  in terms of benefit rather than population growth because population growth in EC is generally restricted by competition for space in the population.

Chesson's coexistence theory [10], [12]; to stably coexist, each species/genotype/phenotype must limit itself more than it limits others (as formalized in Equation 3). In lexicase selection, biological fitness (i.e. the proportion of "islands" the species dominates) must be above a cutoff (described in Figure 4).

Although Eco-EA and fitness sharing have the same coexistence criteria, they have an important difference from each other: in Eco-EA, competition occurs along multiple, potentially orthogonal dimensions, whereas in fitness sharing competition is mediated through a single function that summarizes all aspects of an individual. This difference should result in fewer, more intense pairwise competitive interaction in Eco-EA than in fitness sharing (although still not as few, or as intense, as those in lexicase selection). We predict that the more focused interactions in Eco-EA and lexicase selection will promote forms of evolutionary divergence that are more useful for adaptive evolution. The requirement in Eco-EA that coexisting individuals must be non-dominated suggests that individuals in Eco-EA should fall roughly along a Pareto front, similarly to individuals in lexicase selection. An important distinction between these two selection schemes, however, is that lexicase selection's rigid population structure produces strong pressure for specialists, whereas generalists can be successful in Eco-EA.

### III. EMPIRICAL METHODS

To confirm and extend the intuition we developed in the previous section, we now empirically investigate these selection mechanisms in the context of actual evolving populations.

#### A. Interaction networks

In Section II, we predicted the competitive pressures that different selection schemes exert on populations. Here, we assess these predictions by drawing interaction networks for real populations. For simplicity, we use a population containing 10 individuals with 5 integer traits selected from a geometric distribution. Each trait represents a niche that these individuals are competing to occupy, with higher numbers corresponding to higher competitive ability. For lexicase selection, each trait is a selection criterion and the individual with the highest value there wins. In Eco-EA, we added a resource associated with each trait and an individual's value for that trait defines its ability to use that resource. In fitness sharing, distance is measured as euclidean distance between sets of traits. The fitness landscape is otherwise flat. We compare the same population across all selection schemes. To calculate the effect that a given individual,  $A$  has on another individual,  $B$ , we first calculate the fitness of  $B$  in the presence of the whole population. Then we remove  $A$  from the population and recalculate the fitness of  $B$ . The difference between these two fitnesses is the effect of  $A$  on  $B$ . Note that these fitness values are biological fitness, i.e. the expected number of offspring, rather than the fitness produced by the fitness function. For fitness sharing and Eco-EA, we assume a tournament size of 2 when making this calculation.

### B. Phylogenetic analysis of evolved populations

What is the long-term effect of these different interaction network topologies? We arrive at a first-order approximation by analyzing the phylogenies of populations evolved under each selection scheme. Across selection schemes, the behavior of a population depends on the part of the fitness landscape that it is currently exploring. If the entire population is climbing a steep hill, all forms of diversity should be low, due to frequent selective sweeps. To assess the effect of such differences in fitness landscapes, we compare phylogenetic diversity across three different genetic programming problems believed to have qualitatively different fitness landscapes. The first task, chosen to be quickly-solvable, is calculating the square of a number. In contrast, the second task is a math problem known to be challenging: calculating numbers in the Collatz sequence [28]. The third task is the Dow chemical challenge, a real-world symbolic regression problem [40]. Using a simple linear genetic programming representation (described in [41]), we evolve linear genetic programs for 1000 generations. We use 11 test cases for the squares problem and 100 test cases each for the Collatz problem and Dow chemical challenge. Phenotypes are the vectors of inverse error for each test case (so higher scores are better). Fitness is calculated as the sum of these vectors. For lexica selection and Eco-EA, each test case corresponds to a selection criterion or resource.

We evolved 30 populations for each selection scheme and each problem. Because of the profound effect of the sharing threshold parameter on the behavior of fitness sharing, we also performed 30 runs each of fitness sharing with five different sharing thresholds. Statistics across selection schemes were calculated using the sharing threshold with the highest phylogenetic diversity. For each run, we calculated a variety of metrics, including phenotypic diversity (measured as Shannon entropy), genotypic diversity, phylogenetic diversity [19], and mean pairwise distance of genotypes in the phylogeny [25]. Across all runs, we used a genome length of 200 and a population size of 500. We applied mutations to every new individual by randomly changing the code at up to three sites in the genome (the specific number was selected from a uniform distribution). To simplify calculation of the phylogenetic metrics, we seeded each population with a single individual that served as the common ancestor to all others. For Eco-EA, we used a resource inflow rate of 100, a cost of 1, and a consumption fraction for 0.0025.

### C. Statistical methods

That statistical significance of all comparisons of metrics across conditions was assessed using a Kruskal-Wallis test. Differences between specific conditions were assessed with a post-hoc pairwise Wilcoxon rank-sum test accompanied by a Bonferroni correction for multiple comparisons.

### D. Code availability

Data for the interaction network analysis was generated using a simulation written with Python 3.6.3. Graphs were visualized using the networkx package [42]. The

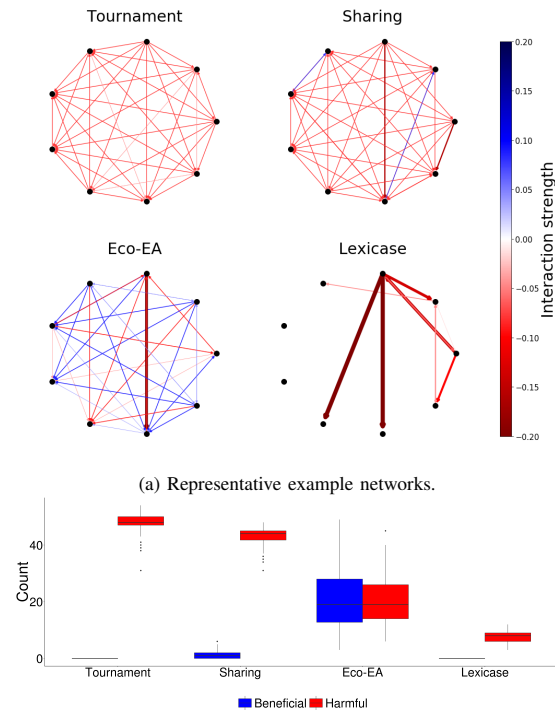


Fig. 5: Interaction networks for the same community under four different selection schemes: tournament selection, fitness sharing, Eco-EA, and lexica selection. a) Shows the networks for a single population. Red edges indicate harmful interactions, blue edges indicate beneficial interactions. Edge width denotes interaction strength. An interactive version of this paper is available [45]. b) Shows boxplots summarizing the number of beneficial and harmful interactions observed under each selection scheme across 100 randomly generated populations.

code for empirical analysis of evolved populations was written in C++ using the Empirical library, available at <https://github.com/devosoft/Empirical>. Statistical analysis was performed using the R statistical computing language version 3.4.3 [43] and graphs were made with the ggplot2 package [44]. All code for generating and analyzing the data presented in this paper is open source and available [45].

## IV. RESULTS AND DISCUSSION

### A. Interaction networks

Different selection schemes create strikingly different interaction networks (see Figure 5). Notably, lexica selection's rigid population structure leads to far fewer interactions of both types<sup>2</sup> than under other schemes (pairwise Wilcoxon rank-sum

<sup>2</sup>With one exception: we observed no beneficial interactions in lexica selection, so these groups were not different from each other.






Fig. 6: Phenotypic and phylogenetic diversity over time for each problem. Shaded area is the bootstrapped 95% confidence interval around the mean.

test,  $p < .0001$ ), and ones that are predominantly negative. Eco-EA creates far more interactions, because individuals can use different resources to different extents. Importantly, there are many beneficial interactions; if  $A$  competes with  $B$  and  $B$  competes with  $C$ ,  $A$  can help  $C$  by suppressing  $B$ . These interactions imply a stabilizing dynamic between  $A$  and  $C$ , because they indicate that  $A$  is competing with  $C$  less than with itself. Additionally, the Eco-EA community has many weak interactions, which can promote stabilizing dynamics community-wide [46]. Lastly, in fitness sharing, most individuals harm each other approximately the same amount. The sharing threshold subtly affects interaction strength.

#### B. Phylogenetic analysis

As hypothesized, phylogenetic diversity is low for all selection schemes on the square problem (see Figure 6). This result is presumably due to fitness differences so large that stabilizing effects are insufficient to overcome them. In the context of such a steep evolutionary gradient, such behavior is expected and often desirable. Due to negative frequency dependence, phenotypic diversity remains high for fitness sharing and Eco-EA, even after most populations solved the problem (generally between 200 and 500 generations in). For lexicase selection, on the other hand, it initially increases and then drops rapidly as the populations converge on the solution.

Results from the Collatz problem support our hypothesis that selection schemes with more restrictions on which individuals compete promote phylogenetic diversity (see Figure 6). Lexicase selection and Eco-EA did not have significantly different final phylogenetic diversity (Wilcoxon rank-sum test,  $p=0.31$ ), but all other pairs of selection schemes did (Wilcoxon rank-sum tests,  $p < .05$ ). Results for mean pairwise distance




Fig. 7: Phenotypic and phylogenetic diversities in fitness sharing across various values of sharing threshold ( $\sigma_{\text{share}}$ ) in the context of the Collatz problem. Shaded areas are bootstrapped 95% confidence intervals around the mean.

were similar, suggesting that lexicase selection and Eco-EA (and to a lesser extent fitness sharing) do promote the coexistence of divergent branches.

Interestingly, phenotypic diversity does not correlate especially closely with phylogenetic diversity (see Figure 6). In particular, Eco-EA has relatively low phenotypic diversity, despite its high phylogenetic diversity. Conversely, fitness sharing has relatively high phenotypic diversity despite its mid-range phylogenetic diversity. This result suggests that whereas lexicase selection and fitness sharing allow similar phenotypes to coexist, Eco-EA forces them to converge. A potential explanation for this difference is that Eco-EA rewards generalists substantially more than lexicase selection.

This discrepancy is even more pronounced in the context of selecting a sharing threshold for fitness sharing (see Figure 7). Phenotypic diversity is maximized at  $\sigma_{\text{share}} = 1$  (Wilcoxon rank-sum tests,  $p < .005$ ), whereas mean pairwise distance and phylogenetic diversity are highest at  $\sigma_{\text{share}} = 10$  (Wilcoxon rank-sum tests,  $p < .005$ ). This result emphasizes the fact that phylogenetic diversity is meaningfully different from phenotypic and genotypic diversity in ways that can affect choice of parameter values.

Results from the Dow problem (see Figure 6) illustrate the strong effect that the topology of the fitness landscape has on which techniques are most effective at maintaining phylogenetic diversity. In contrast to its high efficacy on the Collatz problem, Eco-EA maintains no more phylogenetic diversity than fitness sharing and tournament selection on the Dow chemical problem (Wilcoxon rank-sum tests,  $p = 1$  for all). Lexicase selection, on the other hand, maintains phylogenetic diversity on both of these problems (Wilcoxon rank-sum tests,  $p < .005$ ). Understanding what properties of the fitness landscape account for this difference is an important area of future research.

#### V. CONCLUSIONS

In this paper, we have taken a high-level tour of tools that ecology can contribute to EC. Using metaphors and mathematical theory from ecology, we have identified important differences across four selection schemes. Importantly, we

have shown mathematical equivalence between rules governing coexistence in ecology and in EC. This framework should allow us to directly translate insights between these fields.

Building on this conceptual understanding, we empirically measured the interactions that describe the communities created by different selection schemes. The topology of these interaction networks supports our theoretical predictions, and provides insight into the mechanisms by which these techniques maintain diversity. A more systematic study of interaction networks under various selection schemes will further improve our understanding. We gained further insight by exploring the long-term effects of these selection schemes with phylogenetic metrics. Notably, we demonstrated that phylogenetic analysis metrics describe aspects of the underlying dynamics of diversification that are distinct from those described by genotypic and phenotypic diversity. This result holds across different parameter choices for the same selection scheme, and across different selection schemes. Analyzing the phylogenies of populations evolving on a wider range of fitness landscapes has the potential to help predict which parameters and selection schemes are most appropriate for which problems.

From our analysis thus far, we can conclude that lexica selection and Eco-EA rely on distinct underlying mechanisms to promote the evolution and maintenance of solutions that are evolutionarily divergent. Lexica selection creates a small number of intense competitive interactions, whereas Eco-EA creates a larger number of weaker interactions that can be either competitive or facilitative. As a result, lexica selection generates many specialist phenotypes, while Eco-EA supports fewer, more generalist phenotypes. Both of these selection schemes produce populations representing a wider diversity of evolutionary history than those produced by fitness sharing.

Ecological techniques have the power to revolutionize our understanding of diversity in EC, and our work here has barely started. In the future, we plan to delve deeper into these approaches, to import more ideas and statistical techniques from ecology, and to apply them to an even wider range of selection schemes.

#### ACKNOWLEDGMENT

We thank Alex Lalejini, Matthew Andres Moreno, and Michael Wiser for their helpful comments and suggestions on early drafts of this manuscript. This research has been supported by the National Science Foundation (NSF) BEACON Center under Cooperative Agreement DBI-0939454, by an NSF Graduate Research Fellowship to ED under Grant No. DGE-1424871, and by Michigan State University through computational resources provided by 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.

#### REFERENCES

- [1] H. Kokko, A. Chaturvedi, D. Croll, M. C. Fischer, F. Guillaume, S. Karrenberg, B. Kerr, G. Rolshausen, and J. Stapley, "Can Evolution Supply What Ecology Demands?" *Trends in Ecology & Evolution*, vol. 32, no. 3, pp. 187–197, Mar. 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0169534716302385>
- [2] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 67–82, 1997.
- [3] D. E. Goldberg and J. Richardson, "Genetic algorithms with sharing for multimodal function optimization," in *Genetic algorithms and their applications: Proceedings of the Second International Conference on Genetic Algorithms*. Hillsdale, NJ: Lawrence Erlbaum, 1987, pp. 41–49.
- [4] S. W. Mahfoud, "Niching methods for genetic algorithms," *Urbana*, vol. 51, no. 95001, 1995.
- [5] J.-B. Mouret and S. Doncieux, "Using Behavioral Exploration Objectives to Solve Deceptive Problems in Neuro-evolution," in *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '09. New York, NY, USA: ACM, 2009, pp. 627–634. [Online]. Available: <http://doi.acm.org/10.1145/1569901.1569988>
- [6] J. K. Pugh, L. B. Soros, P. A. Szerlip, and K. O. Stanley, "Confronting the Challenge of Quality Diversity," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '15. New York, NY, USA: ACM, 2015, pp. 967–974. [Online]. Available: <http://doi.acm.org/10.1145/2739480.2754664>
- [7] B. L. Walker and C. Ofria, "Evolutionary Potential is Maximized at Intermediate Diversity Levels," in *Proceedings of the 13th International Conference for Artificial Life*. East Lansing, MI, USA: MIT Press, Jul. 2012, pp. 116–120. [Online]. Available: <http://mitpress.mit.edu/sites/default/files/titles/content/alife13/ch017.html>
- [8] L. Spector, "Assessment of problem modality by differential performance of lexica selection in genetic programming: a preliminary report," in *Proceedings of the 14th annual conference companion on Genetic and evolutionary computation*. ACM, 2012, pp. 401–408. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2330846>
- [9] S. W. Pacala and D. Tilman, "Limiting Similarity in Mechanistic and Spatial Models of Plant Competition in Heterogeneous Environments," *The American Naturalist*, vol. 143, no. 2, pp. 222–257, 1994. [Online]. Available: <https://www.jstor.org/stable/2462642>
- [10] P. Chesson, "Mechanisms of Maintenance of Species Diversity," *Annual Review of Ecology and Systematics*, vol. 31, pp. 343–366, 2000.
- [11] J. M. Chase and M. A. Leibold, *Ecological Niches: Linking Classical and Contemporary Approaches*. University of Chicago Press, Jul. 2003, google-Books-ID: Ssmcl\_ubQUQC.
- [12] A. D. Letten, P.-J. Ke, and T. Fukami, "Linking modern coexistence theory and contemporary niche theory," *Ecological Monographs*, vol. 87, no. 2, pp. 161–177, May 2017. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1002/ecm.1242/abstract>
- [13] A. P. Hendry, *Eco-evolutionary Dynamics*. Princeton University Press, Nov. 2016, google-Books-ID: QMlDAAQAQBAJ.
- [14] E. Dolson, A. Lalejini, S. Jorgensen, and C. Ofria, "Quantifying the Tape of Life: Ancestry-based Metrics Provide Insights and Intuition about Evolutionary Dynamics," in *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*. Tokyo, Japan: MIT Press, 2018, pp. 75–82. [Online]. Available: [https://www.mitpressjournals.org/doi/abs/10.1162/isal\\_a\\_00020](https://www.mitpressjournals.org/doi/abs/10.1162/isal_a_00020)
- [15] M. Winter, V. Devictor, and O. Schweiger, "Phylogenetic diversity and nature conservation: where are we?" *Trends in Ecology & Evolution*, vol. 28, no. 4, pp. 199–204, Apr. 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0169534712002881>
- [16] C. Fontaine, P. R. Guimares, S. Kfi, N. Loeuille, J. Memmott, W. H. van der Putten, F. J. F. van Veen, and E. Thibault, "The ecological and evolutionary implications of merging different types of networks," *Ecology Letters*, vol. 14, no. 11, pp. 1170–1181, Nov. 2011. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1111/j.1461-0248.2011.01688.x/abstract>
- [17] J. P. Grover, *Resource Competition*. Springer Science & Business Media, Jul. 1997.
- [18] P. B. Adler, J. HilleRisLambers, and J. M. Levine, "A niche for neutrality," *Ecology Letters*, vol. 10, no. 2, pp. 95–104, Feb. 2007. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1111/j.1461-0248.2006.00996.x/abstract>
- [19] D. P. Faith, "Conservation evaluation and phylogenetic diversity," *Biological Conservation*, vol. 61, no. 1, pp. 1–10, Jan. 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0006320792912013>
- [20] E. Dolson, M. J. Wiser, and C. A. Ofria, "The Effects of Evolution and Spatial Structure on Diversity in Biological Reserves," in *Artificial Life XV: Proceedings of the Fifteenth International Conference on Artificial*

- Life*. Cancun, Mexico: MIT Press, 2016, pp. 434–440. [Online]. Available: <http://www.biorxiv.org/content/early/2016/03/10/043083>
- [21] M. Tomassini, *Spatially Structured Evolutionary Algorithms: Artificial Evolution in Space and Time*, ser. Natural Computing Series. Berlin, Heidelberg: Springer Berlin Heidelberg, Jan. 2005.
- [22] E. Dolson and C. Ofria, “Spatial resource heterogeneity creates local hotspots of evolutionary potential,” in *Proceedings of the European Conference on Artificial Life*. Lyon, France: MIT Press, 2017, pp. 122–129. [Online]. Available: [https://www.mitpressjournals.org/doi/abs/10.1162/ecal\\_a\\_0023\\_14](https://www.mitpressjournals.org/doi/abs/10.1162/ecal_a_0023_14)
- [23] E. Dolson, S. Perez, R. Olson, and C. Ofria, “Spatial resource heterogeneity increases diversity and evolutionary potential,” *bioRxiv*, 2017. [Online]. Available: <http://biorxiv.org/content/early/2017/06/12/148973>
- [24] C. M. Tucker, M. W. Cadotte, S. B. Carvalho, T. J. Davies, S. Ferrier, S. A. Fritz, R. Grenyer, M. R. Helmus, L. S. Jin, A. O. Mooers, S. Pavoine, O. Purschke, D. W. Redding, D. F. Rosauer, M. Winter, and F. Mazel, “A guide to phylogenetic metrics for conservation, community ecology and macroecology,” *Biological Reviews*, vol. 92, no. 2, pp. 698–715, May 2017. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1111/brv.12252/abstract>
- [25] C. O. Webb and A. E. J. B. Losos, “Exploring the Phylogenetic Structure of Ecological Communities: An Example for Rain Forest Trees,” *The American Naturalist*, vol. 156, no. 2, pp. 145–155, 2000. [Online]. Available: <http://www.jstor.org/stable/10.1086/303378>
- [26] S. Goings and C. Ofria, “Ecological approaches to diversity maintenance in evolutionary algorithms,” in *IEEE Symposium on Artificial Life, 2009. ALife '09*, Mar. 2009, pp. 124–130.
- [27] E. Dolson, W. Banzhaf, and C. Ofria, “Applying Ecological Principles to Genetic Programming,” in *Genetic Programming Theory and Practice XV*, W. Banzhaf, R. S. Olson, W. Tozier, and R. Riolo, Eds. Cham: Springer International Publishing, 2018, pp. 73–88.
- [28] T. Helmuth and L. Spector, “General Program Synthesis Benchmark Suite,” in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '15. New York, NY, USA: ACM, 2015, pp. 1039–1046. [Online]. Available: <http://doi.acm.org/10.1145/2739480.2754769>
- [29] T. Helmuth, L. Spector, and J. Matheson, “Solving Uncompromising Problems With Lexicase Selection,” *IEEE Transactions on Evolutionary Computation*, vol. 19, no. 5, pp. 630–643, Oct. 2015.
- [30] T. Helmuth, N. F. McPhee, and L. Spector, “Effects of Lexicase and Tournament Selection on Diversity Recovery and Maintenance,” *ACM Press*, 2016, pp. 983–990. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2908961.2931657>
- [31] N. F. McPhee, D. Donatucci, and T. Helmuth, “Using Graph Databases to Explore the Dynamics of Genetic Programming Runs,” in *Genetic Programming Theory and Practice XIII*, ser. Genetic and Evolutionary Computation. Springer, Cham, 2016, pp. 185–201.
- [32] T. Helmuth, N. F. McPhee, and L. Spector, “The Impact of Hyperselection on Lexicase Selection,” *ACM Press*, 2016, pp. 717–724. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2908812.2908851>
- [33] N. F. McPhee, M. M. Casale, M. Finzel, T. Helmuth, and L. Spector, “Visualizing Genetic Programming Ancestries,” in *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion*, ser. GECCO '16 Companion. New York, NY, USA: ACM, 2016, pp. 1419–1426. [Online]. Available: <http://doi.acm.org/10.1145/2908961.2931741>
- [34] S. Goings, “Natural niching: Applying ecological principles to evolutionary computation,” Ph.D., Michigan State University, United States – Michigan, 2010. [Online]. Available: <http://search.proquest.com/docview/748246734/abstract/B48FAA8216574AE5PQ/1>
- [35] S. Goings, H. J. Goldsby, B. H. Cheng, and C. Ofria, “An ecology-based evolutionary algorithm to evolve solutions to complex problems,” *Artificial Life*, vol. 13, pp. 171–177, 2012. [Online]. Available: <http://www.cse.msu.edu/~ofria/pubs/2012GoingsEtAl.pdf>
- [36] D. Schluter, T. D. Price, and P. R. Grant, “Ecological Character Displacement in Darwin’s Finches,” *Science*, vol. 227, no. 4690, pp. 1056–1059, 1985. [Online]. Available: <http://www.jstor.org/stable/1694816>
- [37] K. Deb and D. E. Goldberg, “An Investigation of Niche and Species Formation in Genetic Function Optimization,” in *Proceedings of the 3rd International Conference on Genetic Algorithms*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1989, pp. 42–50. [Online]. Available: <http://dl.acm.org/citation.cfm?id=645512.657099>
- [38] P. Chesson and J. J. Kuang, “The interaction between predation and competition,” *Nature*, vol. 456, no. 7219, pp. 235–238, Nov. 2008. [Online]. Available: <http://www.nature.com/articles/nature07248>
- [39] W. La Cava, T. Helmuth, L. Spector, and J. H. Moore, “A probabilistic and multi-objective analysis of lexicase selection and -lexicase selection,” *Evolutionary Computation*, vol. 0, no. 0, pp. 1–26, 0, pMID: 29746157. [Online]. Available: [https://doi.org/10.1162/evco\\_a\\_00224](https://doi.org/10.1162/evco_a_00224)
- [40] D. R. White, J. McDermott, M. Castelli, L. Manzoni, B. W. Goldman, G. Kronberger, W. Jakowski, U.-M. O’Reilly, and S. Luke, “Better GP benchmarks: community survey results and proposals,” *Genetic Programming and Evolvable Machines*, vol. 14, no. 1, pp. 3–29, Mar. 2013. [Online]. Available: <https://link.springer.com/article/10.1007/s10710-012-9177-2>
- [41] E. Dolson, A. Lalejini, and C. Ofria, “Exploring genetic programming systems with MAP-Elites,” *PeerJ Preprints*, vol. 6, p. e27154v1, Aug. 2018. [Online]. Available: <https://doi.org/10.7287/peerj.preprints.27154v1>
- [42] A. A. Hagberg, D. A. Schult, and P. J. Swart, “Exploring network structure, dynamics, and function using NetworkX,” in *Proceedings of the 7th Python in Science Conference (SciPy2008)*, Pasadena, CA USA, Aug. 2008, pp. 11–15.
- [43] R Core Team, *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, 2017. [Online]. Available: <https://www.R-project.org/>
- [44] H. Wickham, *ggplot2: elegant graphics for data analysis*. Springer New York, 2009. [Online]. Available: <http://had.co.nz/ggplot2/book>
- [45] E. Dolson, “Code for Ecological Theory Provides Insights into Evolutionary Computation at the time of preprint/submission to journal,” Nov. 2018. [Online]. Available: <https://doi.org/10.5281/zenodo.1476709>
- [46] K. S. McCann, “The diversitystability debate,” *Nature*, vol. 405, no. 6783, pp. 228–233, May 2000. [Online]. Available: <https://www.nature.com/nature/journal/v405/n6783/abs/405228a0.html>