Ecological Theory Provides Insights about Evolutionary Computation

Emily Dolson

BEACON Center for the Study of Evolution in Action, Department of Computer Science and Engineering, Program in Ecology, Evolutionary Biology and Behavior, Michigan State University East Lansing, Michigan dolsonem@msu.edu

CCS CONCEPTS

• Computing methodologies \rightarrow Artificial life; Genetic programming; • Applied computing \rightarrow Biological networks;

KEYWORDS

Diversity, diversity maintenance, fitness sharing, lexicase selection, Eco-EA, ecology, ecological theory, eco-evolutionary dynamics, phylogenetic analysis

ACM Reference Format:

Emily Dolson and Charles Ofria. 2018. Ecological Theory Provides Insights about Evolutionary Computation. In GECCO '18 Companion: Genetic and Evolutionary Computation Conference Companion, July 15–19, 2018, Kyoto, Japan. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3205651. 3205780

1 INTRODUCTION

Promoting diversity in an evolving population is important for Evolutionary Computation (EC) because it reduces premature convergence on suboptimal fitness peaks while still encouraging both exploration and exploitation [3]. However, some types of diversity facilitate finding global optima better than other types. For example, a high mutation rate may maintain high population-level diversity, but all of those genotypes are clustered in a local region of a fitness landscape. Fitness sharing [3], on the other hand, promotes diversity via negative density dependence forcing solutions apart. Lexicase selection [5] goes one step further, dynamically selecting for diverse phenotypic traits, encouraging solutions to actively represent many portions of the landscape.

Techniques for promoting diversity that create new interactions among individuals are, by definition, creating simple ecologies. Ecologists have developed rigorous theory to predict how ecological communities change over time [4]. If we can import this theory into EC, it can inform decisions about how each diversitymaintenance technique to will interact with a given problem. Additionally, an improved mechanistic understanding of how existing algorithms work should facilitate building more effective variations

GECCO '18 Companion, July 15-19, 2018, Kyoto, Japan © 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5764-7/18/07.

https://doi.org/10.1145/3205651.3205780

Charles Ofria

BEACON Center for the Study of Evolution in Action, Department of Computer Science and Engineering, Program in Ecology, Evolutionary Biology and Behavior, Michigan State University East Lansing, Michigan ofria@msu.edu

of those algorithms. Here, we begin to establish isomorphisms between natural ecological communities and communities created in EC.Additionally, we borrow some empirical techniques from ecology to analyze the ecological communities created by different EC selection schemes.

2 ECOLOGY IN EC SYSTEMS

Below, we discuss how three diversity-promoting techniques in EC can be described in a more ecological context.

In fitness sharing [3], an individual's fitness is reduced based on the number of other individuals that are similar to it. Likewise, in ecology, organisms compete more with others that they are most similar to. The closest ecological parallel to fitness sharing may be in environments defined by a resource that exhibits a continuous range of features. For example, seeds may vary in where they grow or how thick their shells are, so birds may evolve differently shaped beaks to limit their competition for the same seeds.

Mathematically, the equations governing stable coexistence (a prerequisite for diversity maintenance) in fitness sharing [1] turn out to be identical to those governing stable coexistence in natural ecologies [4]. Species can coexist if they limit their own growth more than they limit each other's growth. This scenario requires some combination of fitness similarity (equalization) and niche differences (stabilization).

In lexicase selection [5], individuals are selected for reproduction based on a large set of criteria. The entire population is evaluated on each criterion in random order and only the individuals that perform best on one criterion are evaluated on the next, until only one individual remains.

Lexicase selection is akin to a spatially heterogeneous environment in nature. Different species are the best at surviving in different regions. Each ordering of selection criteria corresponds to one of these regions. Species can coexist if they can simultaneously occupy enough space that stochastic extinction is rare.

In Eco-EA [2] limited resources are associated with simple problems, with a goal of generating subpopulations with building blocks that can be used to solve a more complex problem. 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 if they target the same resources. We know from ecological theory that this arrangement allows for stable coexistence

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored For all other uses, contact the owner/author(s).



Figure 1: Interaction networks for the same community under three different selection schemes: lexicase selection, Eco-EA, and fitness sharing. Red edges indicate actively competitive interactions, blue edges indicate facilitative (beneficial) interactions. Edge width denotes interaction strength.

when species use different resources or sets of resources from each other [4].

3 EMPIRICAL RESULTS

3.1 Interaction networks

We can understand the ecological community created by a selection scheme by drawing a graph representing the interactions between individuals in the population. Different selection schemes create strikingly different interaction networks (see Figure 1). Notably, lexicase selection's rigid population structure leads to fewer and more negative interactions than under other schemes. Eco-EA creates far more interactions, because individuals can use resources to different extents; importantly, many of these interactions are beneficial. Lastly, in fitness sharing, most individuals harm each other approximately the same amount.

3.2 Phylogenetic analysis

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. We focus on phylogenetic diversity (mean distance in a phylogeny between all pairs of members of a population) and phenotypic diversity (Shannon diversity of phenotypes in the population). The dynamics of a population depend on the part of the fitness landscape that it is currently exploring, regardless of the selection scheme. To assess the effect of differences in fitness landscapes, we analyze three different genetic programming problems believed to have qualitatively different fitness landscapes: squaring the input, calculating the Collatz sequence, and the Dow chemical symbolic regression benchmark problem [6]. We evolved 30 populations of linear genetic programs for each selection scheme and problem for 1000 generations. All code for generating and analyzing the data presented in this abstract is open source and available at https:// $github.com/emilydolson/ecology_of_evolutionary_computation.$

Phylogenetic diversity is low for all selection schemes on the square problem (see Figure 2), due to the population's rapid convergence on a perfect solution. Results from the Collatz problem support our hypothesis that selection schemes with more restrictions on which individuals compete promote phylogenetic diversity (see Figure 2). Interestingly, phenotypic diversity does not correlate especially closely with phylogenetic diversity (see Figure 2).



Figure 2: Phentypic and phylogenetic diversity over time for each problem. Shaded area is the bootstrapped 95% confidence interval around the mean.

Results from the Dow problem (see Figure 2) illustrate the strong effect that 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.

ACKNOWLEDGMENTS

This research was supported by the National Science Foundation (NSF) BEACON Center under Cooperative Agreement DBI-0939454, an NSF Graduate Research Fellowship to ED under Grant No. DGE-1424871, and Michigan State University's 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] Kalyanmoy Deb and David E. Goldberg. 1989. An Investigation of Niche and Species Formation in Genetic Function Optimization. In Proceedings of the 3rd International Conference on Genetic Algorithms. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 42–50. http://dl.acm.org/citation.cfm?id=645512.657099
- [2] S. Goings and C. Ofria. 2009. Ecological approaches to diversity maintenance in evolutionary algorithms. In *IEEE Symposium on Artificial Life*, 2009. ALife '09. 124–130. https://doi.org/10.1109/ALIFE.2009.4937703
- [3] David E. Goldberg and Jon Richardson. 1987. 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, 41–49.
- [4] Andrew D. Letten, Po-Ju Ke, and Tadashi Fukami. 2017. Linking modern coexistence theory and contemporary niche theory. *Ecological Monographs* 87, 2 (May 2017), 161–177. https://doi.org/10.1002/ecm.1242
- [5] Lee Spector. 2012. Assessment of problem modality by differential performance of lexicase selection in genetic programming: a preliminary report. In Proceedings of the 14th annual conference companion on Genetic and evolutionary computation. ACM, 401–408. http://dl.acm.org/citation.cfm?id=2330846
- [6] David R. White, James McDermott, Mauro Castelli, Luca Manzoni, Brian W. Goldman, Gabriel Kronberger, Wojciech JaÅŹkowski, Una-May OåÅŹReilly, and Sean Luke. 2013. Better GP benchmarks: community survey results and proposals. *Genetic Programming and Evolvable Machines* 14, 1 (March 2013), 3–29. https://doi.org/10.1007/s10710-012-9177-2