

Evolutionary Potential is Maximized at Intermediate Diversity Levels

Bess L. Walker and Charles Ofria

Department of Computer Science and Engineering
BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48823
{blwalker, ofria}@msu.edu

Abstract

Diversity in a population is often cited as a major facilitator for the evolution of new complex features. The intuition behind this dynamic is that if a population is exploring multiple regions of a fitness landscape, more opportunities exist to find new functionality. We use the digital evolution software platform Avida to explore the effect of multiple limited resources on phenotypic Shannon diversity and, in turn, on evolvability of populations. We show that Shannon diversity peaks at intermediate levels of resource availability to the population, and we map the evolvability of a complex computational task on this availability-diversity gradient. While the evolvability of the complex task is highest at intermediate availabilities, it does not peak at the same resource inflow level as Shannon diversity, and it is more robust than diversity in its response to inflow level. These results indicate that while phenotypic Shannon diversity may play into the evolution of complex features, the selective pressures caused by diversity cannot be the only — or indeed even the main — pressures behind such evolution.

Introduction

Resource inflow and availability is a major factor affecting ecosystem diversity (Tilman, 1982; Chesson, 2000; Hall and Colegrave, 2007; Abrams et al., 2008; Cardinale et al., 2009). Diversity, in its turn, has been shown by the evolutionary computation community to encourage the evolution of solutions to complex problems through a more thorough exploration of the fitness landscape (Friedrich et al., 2009). Here, we explore the effect of the availability of multiple limited resources on phenotypic Shannon diversity, and use this availability-mediated diversity gradient to examine the relationship between Shannon diversity and the evolution of complex features.

Of the many types and measures of diversity, we choose to examine phenotypic Shannon diversity. We choose phenotypic over genotypic diversity because, of the two, phenotypic diversity is most easily manipulated with limited resources. We would also expect different drivers of genotypic diversity to have radically different results depending on whether different genotypes form a cloud in one area of the fitness landscape or are spread widely. Although similar issues can exist with phenotypic diversity, the range

of interesting phenotypes in these experiments is far more constrained than the range of interesting genotypes. Phenotypic diversity therefore provides a more fair treatment. We choose to measure phenotypic diversity as the Shannon entropy of the phenotypes in the population because Shannon entropy effectively balances the two main interesting qualities in diversity: the range of possible results and the evenness in their distribution.

Lenski et al. (2003) have investigated the evolutionary origin of complex features using Avida, using Boolean EQU as the specific complex task under study. This is the most complex of the one- and two-input Boolean operations to calculate, requiring at least five logical NAND operations. An Avidian organism requires at least 19 coordinated instructions to perform EQU, including at least five *nand* instructions. The ancestor starts out with none of these instructions in its genome; Lenski et al. found that in the 23 of 50 populations that evolved EQU in their experiments, EQU evolved in anywhere from 51 to 721 mutational steps.

In practice, the evolution of EQU is dependent on rewarding building blocks: the one- and two-input Boolean tasks of lower complexity. When Lenski et al. evolved populations in environments where only EQU was rewarded, none of the populations evolved EQU. However, they also found that the evolution of EQU does not depend on any particular building block or pair of building blocks. In fact, EQU can evolve in many different ways and is not dependent on any one thing; all 23 of Lenski et al.'s EQU-evolving populations evolved building blocks in different orders and organized them differently in their genomes.

Methods

Study System

We use the digital evolution software Avida (Ofria and Wilke, 2004), allowing precise manipulation of resource availability and a complete record of the course of evolution. The Avida system consists of a grid of digital organisms, each with a simple circular genome composed of instructions from an assembly-like Turing-complete instruction set. Time in Avida is measured in *updates*; each update

corresponds to an average number of 30 instruction executions per organism in the population. Organisms running quickly will execute more than 30 instructions per update, while slow organisms will execute fewer.

By executing its genome, each organism is capable of self-reproduction; during this process, copy mutations may be introduced into the offspring’s genome. Because the genetic instructions are drawn from a Turing-complete language, the organisms are also theoretically capable of any other Turing-computable task. The organisms have access to integers that they can manipulate; the researcher can choose to reward certain manipulations with additional CPU cycles. These additional CPU cycles allow the organism to execute its genome more quickly and thus increase fitness.

Avida also supports a resource system, allowing task rewards to be tied to these resources. We accomplish resource manipulation in this system by manipulating the resource supply rate. Of course, precise manipulation of resource supply rate is possible in laboratory chemostat systems, but the use of a digital system allows us to know every detail of the population at any point in evolution, and to achieve very high generation counts over the course of just a few hours for each replicate population. Complete information about the population allows a precise calculation of diversity, which in this asexual system we define as the Shannon entropy of expressed resource-use phenotypes. It also allows a concrete definition of the complex feature we are examining; in this case, the Boolean EQU operation (Table 1).

| Function name | Boolean operation | Reward |
|---------------|--|--------------|
| NOT | $\neg A; \neg B$ | $\times 2^1$ |
| NAND | $\neg(A \wedge B)$ | $\times 2^1$ |
| AND | $A \wedge B$ | $\times 2^2$ |
| ORN | $(A \vee \neg B); (\neg A \vee B)$ | $\times 2^2$ |
| OR | $A \vee B$ | $\times 2^3$ |
| ANDN | $(A \wedge \neg B); (\neg A \wedge B)$ | $\times 2^3$ |
| NOR | $\neg A \wedge \neg B$ | $\times 2^4$ |
| XOR | $(A \wedge \neg B) \vee (\neg A \wedge B)$ | $\times 2^4$ |
| EQU | $(A \wedge B) \vee (\neg A \wedge \neg B)$ | $\times 2^5$ |

Table 1: NAND-count-based task rewards in Lenski et al. The symbol “ \neg ” denotes negation, while semicolons separate symmetrical functions. An organism which performs a task has its current execution rate multiplied by the amount of the task’s reward. Note that the EQU operation is sometimes known as XNOR.

Our experiments use a development version of Avida 2.12.3 with the default instruction set (inst-heads.cfg). The executable was compiled from publicly-available source code (at avida.devosoft.org); the specific git revision identifier of the code is e5ba9511df000bae780c8524abb6bd01987190a5.

We set the per-site copy mutation rate to .0025, while we

left the per-reproduction rates of insertion and deletion mutations at the Avida default value, .05.

The population structure is spatial; organisms reproduce into any of the nine cells surrounding and including the organism itself, preferring empty cells. The resource structure is non-spatial; all organisms access the same resource pools. Our world is a 60 x 60 toroidal grid, initially seeded with 3600 clones of an asexual ancestor organism capable only of reproduction. This ancestor is a modification of the default ancestor that ships with Avida, default-heads.org, to reduce its genotype from length 100 to length 50 by removing 50 lines of “blank tape” no-op instructions. Since the population experiences no bottlenecks, the entire world grid is populated throughout the experiments.

We used SciPy 0.10.1 to calculate statistics, and Matplotlib 1.1.0 to create graphs.

Configurations from Previous Experiments

In their investigation of the evolutionary origin of complex features, Lenski et al. rewarded digital organisms once for each distinct Boolean task performed. The value of each task corresponded to its complexity as approximated by the minimum number of Boolean NAND operations necessary for its performance (see Table 1).

Chow et al. (2004) investigated the relationship between resource inflow and diversity in Avida. They measured diversity as species richness; as the digital organisms are asexual, Chow et al. used a clustering algorithm based on phylogenetic distance to determine which genotypes belonged to the same “species”. Species richness in this system was the result of negative frequency-dependent selection due to multiple depletable resource pools. R_{INFLOW} units of resource flow into each resource pool at a constant rate over each update, and a percentage of each pool flows out, modeling a chemostat.

$$Inflow : R_{TASK} = R_{TASK} + R_{INFLOW} \quad (1)$$

$$Outflow : R_{TASK} = 0.01 * R_{TASK} \quad (2)$$

Chow et al. used the same set of Boolean computational tasks as Lenski et al., but linked each task to a separate resource pool. The amount of resource in a resource pool (R_{TASK}) determines the value of performing the associated task; the NAND-count is not considered. An individual organism depletes A_{TASK} units of resource from the task-linked pool when performing a Boolean task. This depletion results in negative frequency-dependent selection (Cooper and Ofria, 2002). Rewarding an organism for the performance of a task again consists of multiplying its current execution count by the amount of the reward.

$$A_{TASK} = 0.0025 * R_{TASK} \quad (3)$$

$$Depletion : R_{TASK} = R_{TASK} - A_{TASK} \quad (4)$$

$$Reward : \times 2^{A_{TASK}} \quad (5)$$

Limited-Resource Environment

Because the Lenski et al. environment determines task rewards purely by task complexity, it can be thought of as an environment with infinite resource inflow. Without the negative frequency-dependence of Chow et al.'s environments, populations converge to a single generalist genotype that performs all tasks. In the environments of Chow et al., high inflow rates result in populations that converge on a single genotype specialized on replication efficiency; these rarely perform more than one or two of the simpler Boolean tasks. This is because Chow et al. do not incorporate the difficulty of the task into the task's reward; at high resource abundance, there is little to no pressure to seek new resources, and thus no reason to do difficult tasks.

In studying the effect of resource supply on both phenotypic Shannon diversity *and* the evolution of complex features, it is useful to create environments in which both the difficulty of the task and its rarity in the population (via the availability of its associated resource) contribute to the reward an organism receives for performing that task. To that end, we have devised a limited-resource environment starting with Lenski et al.'s reward scheme, but where a linked resource pool mediates the amount of the reward as in Chow et al.; Table 2 describes this hybrid reward scheme.

| Function name | # NAND | Depletion | Reward |
|---------------|--------|------------|-------------------------------|
| NOT | 1 | A_{NOT} | $\times 2^{1 \cdot A_{NOT}}$ |
| NAND | 1 | A_{NAND} | $\times 2^{1 \cdot A_{NAND}}$ |
| AND | 2 | A_{AND} | $\times 2^{2 \cdot A_{AND}}$ |
| ORN | 2 | A_{ORN} | $\times 2^{2 \cdot A_{ORN}}$ |
| OR | 3 | A_{OR} | $\times 2^{3 \cdot A_{OR}}$ |
| ANDN | 3 | A_{ANDN} | $\times 2^{3 \cdot A_{ANDN}}$ |
| NOR | 4 | A_{NOR} | $\times 2^{4 \cdot A_{NOR}}$ |
| XOR | 4 | A_{XOR} | $\times 2^{4 \cdot A_{XOR}}$ |
| EQU | 5 | A_{EQU} | $\times 2^{5 \cdot A_{EQU}}$ |

Table 2: Hybrid task rewards, based both on task complexity and resource availability (A_{TASK} denotes the number of resource units an organism uses from the $TASK$'s pool). An organism that performs a task has its current execution rate multiplied by the amount of the task's reward.

Results and Discussion

Diversity Peaks at Intermediate Productivity

Of the inflow rates we examined, the intermediate R_{INFLOW} of 10 (Figure 1) had the highest diversity; observing the highly unimodal trend of this data, we conclude that diversity in this system peaks somewhere between an R_{INFLOW} of 3 and 30. At lower inflow rates, Shannon diversity drops off quickly; with too-low resource levels, each pool supports too few organisms to make any substantial impact on diversity. The number of phenotypes may remain

high, but the Shannon entropy of the population as a whole is low. At higher inflow rates, diversity drops more slowly as resources become so plentiful they might as well be unlimited. Indeed, at inflow levels of 1000 and above, negative frequency-dependent pressures are effectively removed. This result corresponds to the results in other studies of the effects of resource supply on diversity (e.g. Kassen et al., 2000; Chow et al., 2004; Hall and Colegrave, 2007)

Diversity Distributions over Resource Inflow Rates

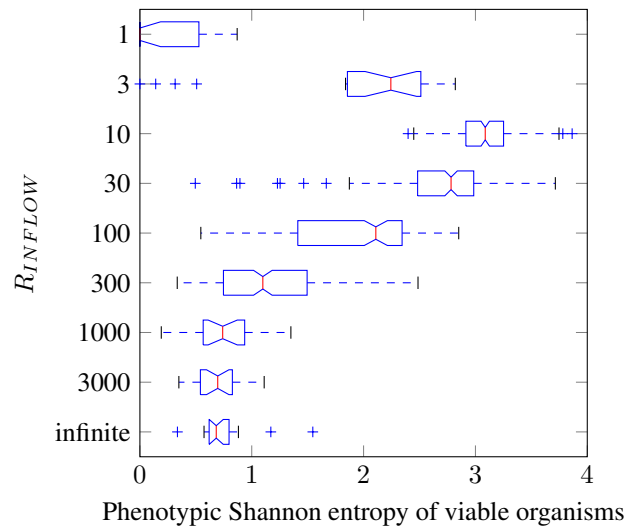


Figure 1: Diversity distributions across inflow rates, measured as the phenotypic Shannon entropy of all viable organisms in the population. Data for inflows 1, 3, 1000, 3000, and infinite are drawn from 20 populations; inflows 10, 30, 100, and 300 from 200 populations.

The Evolution of EQU is Common in Intermediate Productivities

We are now equipped to examine the evolvability of complex features on this resource inflow gradient, and to observe how it relates to the corresponding Shannon diversity gradient. In this case, we measure the evolvability of complex features by the proportion of populations that have evolved EQU by the end of 100,000 updates. At 20 populations per treatment (Figure 2), it is clear that the evolvability of EQU is highest at intermediate productivities. Indeed, between the intermediate inflow levels of 10 and 300 units per resource per update, the evolvability of EQU seems robust to increasing resource supply and decreasing phenotypic Shannon diversity.

To determine whether only complex tasks are sensitive to resource supply levels, we also examined the evolvability of the other 8 tasks rewarded in this environment (Figure 3). As a general trend, these tasks indicate that more complex

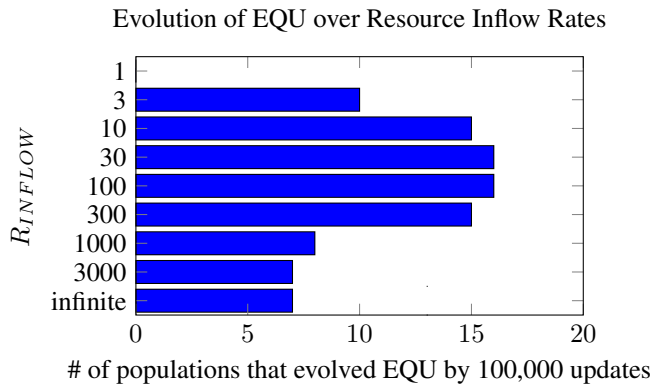


Figure 2: Evolvability of EQU across inflow rates, measured as the number of populations for which some genotype in the final population can perform EQU. Data for all inflows is drawn from 20 populations.

tasks are more sensitive to the resource supply level. Further investigation on this added axis of task complexity was beyond the scope of this paper.

We focused on the inflow rates in which populations were most successful in evolving EQU (10, 30, 100, and 300), and performed 10 times as many experimental runs at each to gain a higher resolution (Table 3). At this resolution, we saw that the evolvability of EQU is not truly unaffected by the variation of resource supply and Shannon diversity in this inflow range. The number of populations evolving EQU by the end of 100,000 updates is significantly higher at the 100 unit inflow rate than at the 10, 30, or 300 unit inflow rates. While these data do not indicate the precise R_{INFLOW} at which the evolutionary potential of EQU peaks, it is clearly a different — and greater — R_{INFLOW} than that at which phenotypic diversity reaches its peak.

| R_{INFLOW} | 10 | 30 | 100 | 300 |
|--------------|----------|--------|-----|--------|
| #pops/200 | 141 | 152 | 171 | 152 |
| p-value | <0.00001 | <0.045 | N/A | <0.045 |

Table 3: Number of populations out of 200 that evolved EQU at intermediate inflow rates. We performed a chi-squared test to determine if the evolvability of EQU for at least one inflow rate differed significantly from the rest ($p < .005$, $\chi^2 = 13.156$, 3 degrees of freedom). With this confirmed, we calculated the significance of each ratio's difference from 171/200 with Fisher's exact test, two-tailed, and corrected with the sequential Bonferroni correction; the $n=2$ correction was applied to both the $R_{INFLOW} = 30$ and $R_{INFLOW} = 300$ data, since they can be ordered arbitrarily.




Figure 3: Evolvability of tasks requiring fewer NAND operations than the EQU task, measured as the number of populations for which some genotype in the final population can perform the task of interest. The first seven tasks (only NOT is shown) showed evolvability across resource inflows that was qualitatively similar to the NOT task shown here, with low evolvability at $R_{INFLOW} = 1$ and very high evolvability for all other inflow rates. The XOR task shows evolvability results qualitatively similar to the EQU task, with XOR being more evolvable at intermediate values of R_{INFLOW} . Data for all inflows and tasks is drawn from 20 populations.

Conclusions

We have seen that, for the inflow rates we tested, the phenotypic Shannon diversity of populations is highest at the 10 unit inflow rate (so likely peaks between R_{INFLOW} of 3 and 30). On the other hand, for the same set of inflow rates, the evolvability of this complex feature is highest at the 100 unit inflow rate (so likely peaks between R_{INFLOW} of 30 and 300). These ranges do not overlap; this difference indicates that diversity cannot be the only driver of the evolution of complex features, which is not unexpected. While the evolvability of complex features is indeed high at peak Shannon diversity, it seems that complex features may require more productive environments to evolve most often. We speculate that this greater resource availability and lesser

phenotypic diversity represent environments where more-abundant resources allow the desperate scramble for survival to relax slightly, allowing organisms to accumulate a collection of building blocks necessary for complex tasks.

These results indicate that evolutionary theory still has a great deal of work to do in tracking down the pressures responsible for the evolution of complex features. However, we have seen in this paper that the evolution of complex features is relatively robust, suggesting that the search for such pressures will not be akin to seeking a needle in a haystack — complex features evolve at a high rate at a large range of diversities in these experiments, and the number of times that EQU successfully evolved displays a decidedly unimodal nature. It is therefore likely to be similarly easy to track down the point of peak evolvability of complex features for other hypothesized pressures.

Future Work

In this paper, we have investigated only phenotypic Shannon diversity as caused by resource-based negative frequency-dependent selection. Negative frequency-dependent selection allows adaptive radiation in the Avida system's homogeneous environment, but it is not the only driver of diversity in nature. The relationship between the evolvability of complex features and diversity as driven by other factors (e.g. spatial structure, heterogeneous environments, or parasite pressures) certainly deserves investigation. Other measures of diversity ought also to be considered. Further, examination of the relationship between diversity and the evolvability of complex features only begins to explore the possible pressures driving the evolution of complex features. Although the mechanisms allowing complex features to evolve have been the subject of much investigation and debate (see Gregory, 2008, for an excellent overview), the exploration of pressures involved in the evolution of complex features has only begun.

Acknowledgements

We would like to thank Richard Lenski for helpful discussions about the relationship between diversity and the evolution of complex features, Heather Goldsby for her comments on the introductory section of this paper, and Michael Wisner for his comments and careful copyediting of the final draft.

This material is based in part upon work supported by the National Science Foundation under Grant No. CCF-0643952 and Cooperative Agreement No. DBI-0939454. 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 National Science Foundation.

References

- Abrams, P. A., Rueffler, C., and Kim, G. (2008). Determinants of the strength of disruptive and/or divergent selection arising from resource competition. *Evolution*, 62(7):1571–1586.
- Cardinale, B. J., Hillebrand, H., Harpole, W. S., Gross, K., and Ptacnik, R. (2009). Separating the influence of resource 'availability' from resource 'imbalance' on productivity–diversity relationships. *Ecology Letters*, 12(6):475–487.
- Chesson, P. (2000). Mechanisms of maintenance of species diversity. *Annual Review of Ecology and Systematics*, 31:343–366.
- Chow, S. S., Claus, W. O., Ofria, C., Lenski, R. E., and Adami, C. (2004). Adaptive radiation from resource competition in digital organisms. *Science*, 305(5680):84–86.
- Cooper, T. F. and Ofria, C. (2002). Evolution of stable ecosystems in populations of digital organisms. In *Proceedings of VIII International Conference on Artificial Life*, pages 227–232.
- Friedrich, T., Oliveto, P. S., Sudholt, D., and Witt, C. (2009). Analysis of diversity-preserving mechanisms for global exploration. *Evolutionary Computation, 2003. CEC'03. The 2003 Congress on*, 17(4):455–476.
- Gregory, T. R. (2008). The evolution of complex organs. *Evolution: Education and Outreach*, 1:358–389.
- Hall, A. R. and Colegrave, N. (2007). How does resource supply affect evolutionary diversification? *Proceedings of the Royal Society B: Biological Sciences*, 274(1606):73–78.
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*, 9(3):90–95.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001–). SciPy: Open source scientific tools for Python.
- Kassen, R., Buckling, A., Bell, G., and Rainey, P. B. (2000). Diversity peaks at intermediate productivity in a laboratory microcosm. *Nature*, 406(6795):508–512.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.
- Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2):191–229.
- Tilman, D. (1982). *Resource Competition and Community Structure*. Princeton University Press, Princeton, NJ.