

Investigating Open-Ended Coevolution in Digital Organisms

Luis Zaman

University of Michigan, Ann Arbor, MI 48109
zamanlh@umich.edu

One of the most striking features of life on earth is the dramatic diversity and creativity produced by Darwinian evolution. As the field of artificial life has matured, interest has grown in recapitulating similar levels of diversity *in silico*. It is immediately obvious that artificial life has not yet reached this goal (Bedau et al., 1998). On the other hand, there is no reason to think that artificial life systems are missing some fundamental component of the evolutionary process.

Emily Dolson and her coauthors have worked towards formalizing an approach focused on identifying (thus potentially overcoming) *barriers* to open-ended evolution (Dolson et al., 2015; Taylor et al., 2016). The goal of this work is principally to operationalize part of this framework to study whether or not coevolving host-parasite communities in Avida exhibit ongoing open-ended evolution. While doing so, we will introduce a few quantitative measures of evolutionary change that are rooted in traditional ecological theory and are suitable for coevolving communities.

Avida

Briefly, in Avida hosts and parasites are self-replicating computer programs that execute CPU cycles to copy their genomes and thus reproduce. Copying instructions is not perfect, so mutations can affect offspring genomes (Ofria and Wilke, 2004).

Host-Parasite Coevolution

In these experiments, parasites were able to infect hosts whenever they performed at least one task the hosts did not also perform (Zaman et al., 2014). We seeded the populations with a single host unable to perform any tasks. After a brief period of time (2000 updates, or about 20 generations), we infected at most 400 hosts with a parasite capable of only performing the simplest logic function, NOT. These evolution experiments only ran for a total of 100,000 updates (about 1,000 generations) in order to maintain tractability when using extremely large populations.

In order to vary the number of tasks available to digital organisms, we first ran experiments where hosts and parasites had available all 77 one, two, and three input logic functions. We sorted the frequency of evolved tasks, and selected

the 25 most frequent preserving their rank ordering. Thus, when we ran experiments with only three tasks, the most frequently evolved three were used. We ran 20 replicates for each of the task treatments. For two task treatments (7 and 15), we additionally ran coevolution experiments where we increased or decreased the population size.

Measures of Evolution

Shannon Diversity of Edges A classic metric of diversity commonly used across disciplines is Shannon’s Diversity Index (H). It combines information about the number and the evenness of the distribution of species in a community. The equation for H is

$$H = - \sum_{i=1}^X x_i \log_2(x_i), \quad (1)$$

where X is the number of unique species, and x_i is the proportion of the community composed of species i .

Here, we used Shannon’s index to measure the diversity of unique **edges** in the interaction network between digital hosts and parasite by identifying which host and parasite phenotypes interacted and their frequency in the community.

Morisita-Horn Community Overlap Index Another classic, though much more obscure, ecological metric is Morisita’s overlap index (Horn, 1966). Intuitively, the Morisita-Horn extension spans the range $[0, 1]$, where $M_H = 0$ indicates totally distinct communities and $M_H = 1$ indicates two communities composed of the same species with the same abundances. The Morisita-Horn overlap is defined as

$$M_H = \frac{2 \sum_{i=1}^S x_i y_i}{\left(\frac{\sum_{i=1}^S x_i^2}{X^2} + \frac{\sum_{i=1}^S y_i^2}{Y^2} \right) XY}, \quad (2)$$

where X and Y are the number of species in community X and Y respectively, $x_i(y_i)$ is the proportion of community

$X(Y)$ composed of species i , and S is the total number of unique species.

Here, we used this metric to measure the overlap *between* independently coevolved replicate communities within the same treatment as well as the overlap *within* the same community at different points in time. The latter measurement allowed us to quantify how much churn coevolving communities experienced through time, while the former gave us a metric of coevolution's repeatability. We again used the edges in the interaction network rather than the individual species in each community as the focal unit (see **Shannon Diversity of Edges**).

Results and Discussion

We observe an increase in the diversity of edges in coevolved communities as we increase the number of potential tasks organisms have available to them (Fig. 1). Increasing the number of available tasks is akin to increasing the complexity of the phenotypic space hosts and parasites can occupy. Interestingly, the diversity levels off after only about 7 tasks (Fig. 1, green line). When we increase the size of the world, we see a continued increase in the level of diversity as the number of tasks increases. Thus, we conclude that the diversity of the community is limited by the number of organisms we are able to keep *alive*.

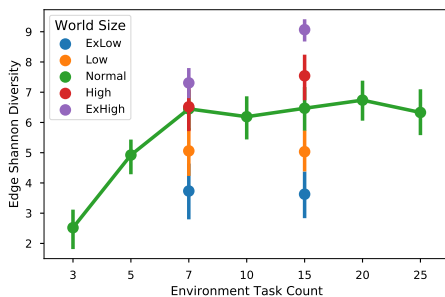


Figure 1: Shannon diversity of host-parasite edges. The connected green line depicts the Normal world size data, while the two extra sets of dots indicate increased and decreased world sizes (ExLow - 900; Low - 3,600; Normal - 14,400; High - 57,600; ExHigh - 230,400).

We see a hump-shaped curve when measuring the total community change as the number of available tasks increases. We measured this by summing the difference between the maximum similarity possible and the measured Morisita-Horn index at each time point (Fig 2). While increasing the world size has a substantial effect when there are 7 tasks available, the effect is muted in the more complex 15-task environment. One potential explanation is that community change is smoother in large phenotypic spaces, leading to fewer spikes of activity.




Figure 2: Sum of community change over time. The same treatments are depicted as in Fig 1.

Finally, we observe that coevolution with limited phenotypic space available leads to similar communities, while increasing the number of tasks in the environment quickly leads to distinct host-parasite networks (Fig 3). Interestingly, increasing the population size slightly increases the similarity between replicates. Perhaps this is because larger populations are able to exploit more phenotypic space.

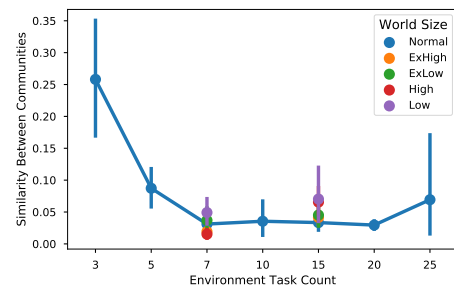


Figure 3: Morisita-Horn overlap between replicate communities within treatments.

References

- Bedau, M., Snyder, E., and Packard, N. (1998). A classification of long-term evolutionary dynamics. *Artificial life VI*, pages 228–237.
- Dolson, E., Vostinar, A., and Ofria, C. (2015). What's holding artificial life back from open-ended evolution? *The Winnower*.
- Horn, H. S. (1966). Measurement of "Overlap" in Comparative Ecological Studies. *The American Naturalist*, 914:419–424.
- Ofria, C. and Wilke, C. O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10:191–229.
- Taylor, T. et al. (2016). Open-Ended Evolution: Perspectives from the OEE Workshop in York. *Artificial Life*, 22:408–423.
- Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., and Ofria, C. (2014). Coevolution Drives the Emergence of Complex Traits and Promotes Evolvability. *PLoS Biology*, 12:e1002023.