

Evolving Social Behavior in Adverse Environments

Brian D. Connelly and Philip K. McKinley

Department of Computer Science and Engineering
Michigan State University
East Lansing, Michigan 48824, USA
{connel42, mckinley}@cse.msu.edu

Abstract. Cooperative behaviors are pervasive in the natural world. How organisms evolve stable cooperative strategies, specifically how selection can favor such costly behaviors, is a difficult problem for which several theories exist. In this work, we use digital evolution to explore the evolution of the production of a public resource that enables populations of organisms to survive in an adverse environment. Kin selection and limited dispersal are shown to promote cooperative acts, and evolved organisms stave off invasion by cheaters and survive in increasingly-adverse environments. Further, we observe how populations react to the disappearance and later re-emergence of adversity in the environment.

Keywords: digital evolution, cooperative behavior, kin selection, public resource.

1 Introduction

The evolution of stable cooperative acts, such as the sharing of food or resources [1], collaboration [2], and self-sacrifice [3], is difficult to explain, since the costs of these behaviors lower the fitness of the actors and make them exploitable by cheaters, who take advantage of the cooperative act, yet do not themselves contribute. One explanation is kin selection [4], whereby organisms maximize their *inclusive* fitness, which includes not only their individual fitness, but also the fitness of kin, who share either common ancestry or behavior. Because these recipients share genetic material, the reproductive successes of an organism's kin also benefit that organism. J.B.S. Haldane captured this idea in his oft-quoted statement, "I would lay down my life for two brothers or eight cousins." Organisms are able to target kin as recipients either through kin discrimination using expressed traits or implicitly if kin remain in close proximity.

Cooperative behaviors among microorganisms, in particular, have received considerable attention in recent years. Quorum sensing, or coordination mediated by chemicals deposited into the environment [5], is one such behavior. Other social acts, such as the production of iron-scavenging siderophores [6] or enzymes [7], have also been studied. Although researchers conducting work in these areas use sophisticated techniques, time often limits the number of generations which can be observed in wet lab experiments.

An alternate approach is to apply mathematical models. Game theory has been used to study the spread of behaviors in a population, most frequently with Prisoner's Dilemma or Snowdrift [8]. These models have provided insights into the dynamics of social behaviors, such as the effects of spatial structures [9], dynamic social ties [10],

and nonlinear benefits and resources [11]. One study using cellular automata [12] required agents to achieve a defined level of success in order to avoid being killed [13]. In these models, mutations altered the organism’s level of investment in the public good. A similar approach is used by Evolution Strategies (ES) [14], which seek to optimize some behavior tied to the mutated parameter. These approaches generally do not include environmental conditions or features aside from the cooperative act, however.

Artificial life tools such as digital evolution [15] provide a means to address these difficulties. Given complete control over environmental conditions, researchers are able to accurately model complex systems. These tools also offer direct access to organisms’ genomes, which greatly aids in understanding the behaviors observed and the conditions under which they are performed. For example, digital evolution has been used to examine group behaviors such as electing leaders, reaching consensus, population control, and foraging [16, 17]. Similar to the work described here, one study investigated kin discrimination in the altruistic suicide of colicinogenic bacteria [18], finding that cooperators were most successful against cheaters when their behaviors were more discriminatory.

In this paper, we describe a study using the Avida digital evolution platform to examine the evolution of cooperative behaviors in populations of organisms. These organisms were able to produce a resource that helped them to survive in adverse environments. When dispersal was limited, populations cooperated to prevent approximately 90% of their constituents from being killed. More well-mixed populations, while also successful, were not able to fare as well. Finally, we observe how populations reacted when the degree of adversity in the environment changed or when they were exposed to threats which they had not encountered in hundreds of generations.

2 The Avida Digital Evolution Platform

Avida is a software platform used to study the evolution of populations of self-replicating computer programs [15]. These computer programs, or “digital organisms,” compete for space by completing *tasks* in a user-defined environment. As depicted in Figure 1, each digital organism exists independently within its own *cell* in the environment. An organism’s behavior is defined by a circular list of instructions (its “genome”), which are executed sequentially on virtual hardware allocated to that organism. The virtual hardware comprises a CPU, three 32-bit registers, and two stacks. The CPU has four *heads*, which are used to control the execution flow of an organism’s genome and aid in self-replication. Organisms control these components using a Turing-complete instruction set. During every *update*, each organism is allotted a number of CPU cycles for executing instructions in its genome.

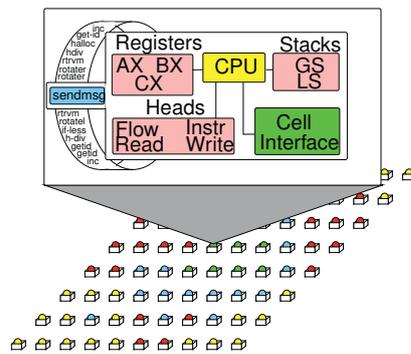


Fig. 1. The structure of a digital organism (top) in an Avida population (bottom) [19]

A population is seeded with an ancestral organism capable only of replication. All other behaviors exhibited by that organism's descendants must be evolved. As an organism replicates, it allocates additional space at the end of its genome into which it copies the instructions from its genome line by line. After this process has been completed, the genome is divided, and the new copy is placed into another cell in the environment, killing any organism at that location. During replication, mutations can cause the insertion, deletion, or changing of one or more instructions in the genome.

Organisms earn *merit* through the successful completion of *tasks*, which are defined by the user in terms of an organism's observable behaviors (i.e., its *phenotype*). Rewarded tasks may include performing a mathematical or logical operation [20], mitigating an attack [21], or cooperating to solve a distributed problem [17]. The number of CPU cycles an organism receives per update is directly proportional to its merit. This task mechanism creates competition in the population, as organisms with more merit are able to execute more quickly, and hence are likely to replicate more often and spread throughout the population. Tasks may involve the use of *resources*. By executing certain instructions, organisms are able to sense, consume, and produce resources. Resource levels in the environment can fluctuate over time through inflow, outflow, diffusion, decay, and consumption.

3 Experiments and Results

In this study, digital organisms evolved in environments in which a periodic event killed a portion of the population. An organism could avoid being killed if the amount of a resource in its cell was above a defined threshold. Although an organism could produce enough resource on its own to avoid being killed, that organism could also benefit from the resource production of its neighbors, due to the diffusion of the resource away from the cell in which it was created. Since the population exhibits limited dispersal, it is likely that neighbors are highly related, so kin selection [4] should enable this costly behavior to be maintained as long as it continues to provide a sufficient inclusive fitness benefit. However, since an organism can survive solely through the production of

Table 1. Logic tasks that could be completed by organisms. Upon completion, an organism's merit was multiplied by the reward listed. Organisms were rewarded once per task, but could continue to produce resource through multiple completions of OR NOT.

Task	Input	Output	Merit Bonus	Resource Produced
NOT	A	$\neg A$	2	0
NAND	A, B	$\neg(A \wedge B)$	2	0
AND	A, B	$A \wedge B$	4	0
OR	A, B	$A \vee B$	8	0
AND NOT	A, B	$A \wedge \neg B, \neg A \wedge B$	8	0
NOR	A, B	$\neg(A \vee B)$	16	0
XOR	A, B	$(A \wedge \neg B) \vee (\neg A \wedge B)$	16	0
EQU	A, B	$(A \wedge B) \vee (\neg A \wedge \neg B)$	32	0
OR NOT	A, B	$A \vee \neg B, \neg A \vee B$	0	1

resource by others, this environment also creates an opportunity for cheaters to exploit this public good.

Avida Configuration. For each experiment, multiple independent populations were evolved. Populations were seeded with a single ancestral organism that was placed into a cell in a 100x100-cell bounded grid environment; each interior cell had 8 neighboring cells. Two different replication methods were tested: limited dispersal, where offspring were placed in a cell neighboring their parent; and well-mixed, where offspring were placed in a random cell in the grid. During replication, an instruction mutated with probability 0.0075 during the copy phase, and an instruction was either added or removed with probability 0.05 as the organism divided. Each population started with a different random seed, so different evolutionary pathways were followed. These populations evolved for 50,000 updates, or approximately 7,000-10,000 generations.

Task Environment. Organisms evolved to complete one- and two-input logic tasks, listed in Table 1. When an organism executed an IO instruction, randomly-generated numbers were placed into its registers. The organism then had to perform the operation, place the result in the correct output register, and issue an additional IO instruction. If the output value matched the result of one of these tasks, the organism was rewarded.

Unlike other tasks, the completion of OR NOT did not yield any merit reward. Instead, it resulted in 1 unit of an extracellular resource being deposited into the environment at the organism's location. This resource both diffused and decayed at a rate of 1% per update and represented a beneficial product to the organism, akin to extracellular polymeric substance (EPS) [22], siderophores [6], or enzymes [23] in microorganisms. The amount of this resource present in a cell was used to determine whether or not an organism residing there was killed. Although this resource was beneficial to the organisms as described, it was not required for the completion of any other tasks.

Initial Experiments. Because each task could potentially serve as either a building block or a hindrance to completing another task, as in [20], we evolved populations in an environment without the periodic kill event in order to determine how many OR NOT tasks would be completed (and consequently how much resource would be produced) when this behavior was neither rewarded nor necessary for survival. These runs produced a mean per-cell resource level of 0.37 units. By using a significantly-higher threshold, we can infer that the success of organisms at staving off periodic killing was not simply a by-product of completing rewarded tasks.

Random Attacks. In this experiment, 121 cells (1.21% of the population) were randomly chosen at each update, and an organism living in those cells was killed if the level of resource in its cell was below 2 units. By the end of the run, organisms had evolved the production of the resource at a level that prevented approximately 66% of organisms selected from being killed. From this experiment alone, however, it is unclear whether the resource production was simply self-preserving or mutually-beneficial [24]. Since only 1.21% of organisms were subject to being killed at each update, the probability of an organism and all of its kin being killed was very small. Indeed, organisms in several runs preferred to play these odds and stopped producing resource, focusing entirely on completing the rewarded tasks.

Localized Attacks. In the next experiment, the kill event was localized, increasing the likelihood of one organism and its kin being killed by a single event. At each update, one cell was chosen at random, and any organism within a 5-cell radius (121 cells total) whose cell contained less than 2 units of resource was killed. The behaviors seen in these populations are plotted in Figure 2. At the end of these runs, an average of approximately 53% of organisms were cooperators (resource producers). Cheaters, which focused solely on the completion of one or more rewarded task but did not produce the resource, accounted for 23% of living organisms. Approximately 24% were unable to complete any tasks, most likely due to deleterious mutations.

We also tracked the level of resource in each cell as the runs progressed. One might predict this level to converge to the threshold amount of 2 units; however, modest stability was reached at approximately 9.5 units. This surplus can be viewed as cooperative: Cooperators produced enough to allow themselves to survive and to help their neighbors survive as well. This asymmetry follows the *Tragedy of the Commune* [25], where levels of investment in public goods may not be uniform. Figure 3(a) shows the distribution of this resource at the end of a typical run. Although only about half of the living organisms produced resource, it was sufficient to prevent 90% of organisms in the population from being killed as shown in Figure 3(b).

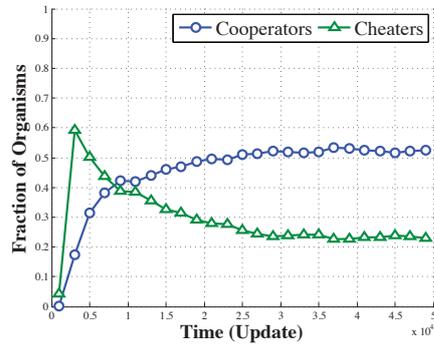


Fig. 2. Mean distribution of behaviors among 40 populations

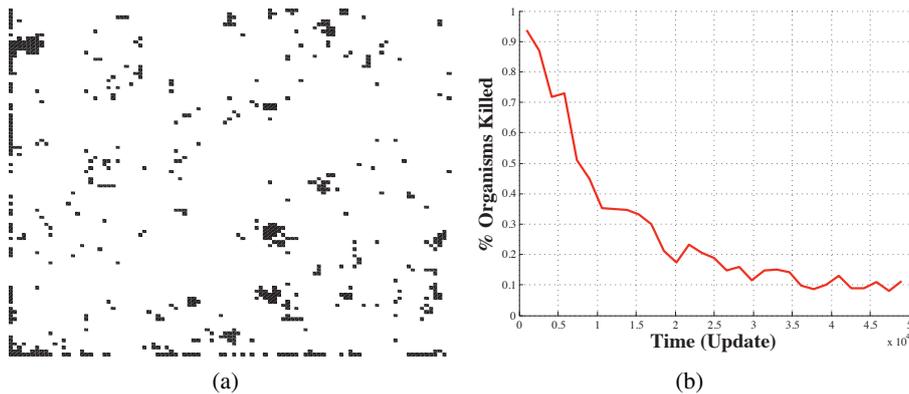


Fig. 3. Resource distribution and organisms killed: (a) Distribution of resource levels in a typical environment. White indicates an above-threshold level, while black indicates a level below threshold. (b) Mean fraction of organisms killed within the target region.

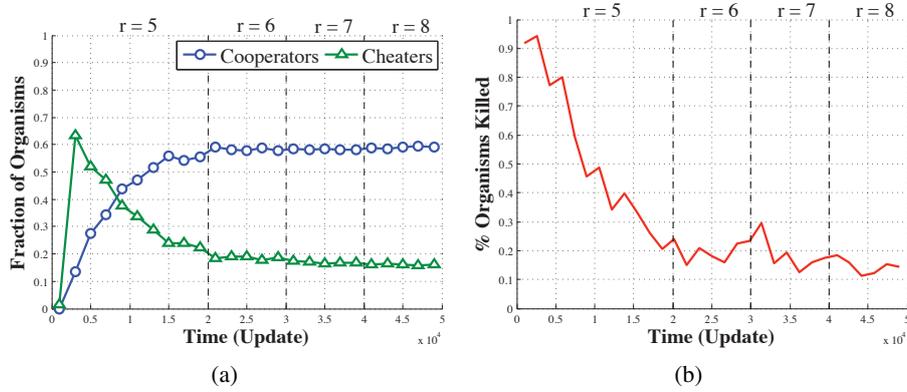


Fig. 4. Results for an increasingly-adverse environment: (a) Mean distribution of behaviors in 20 populations (b) Mean fraction of organisms within the target region killed during each kill event

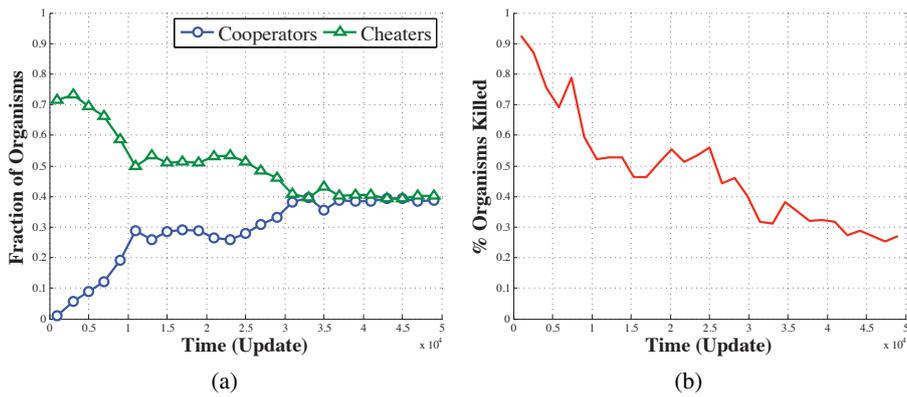


Fig. 5. Results in a well-mixed environment: (a) Mean distribution of cooperators and cheaters in 20 populations (b) Mean fraction of organisms killed within a 5-cell radius

Greater Adversity. To investigate how populations would evolve in more adverse environments, we repeated these experiments using kill radii up to 13 cells and a resource threshold of 3 units. These environments proved to be too adverse, and populations were not able to persist with kill radii above 5 cells. However, when the kill radius was expanded incrementally from 5 cells to 8, more than doubling the number of organisms at risk, the populations adapted and produced enough resource to survive these events. The distribution of resource producers and cheaters is shown in Figure 4(a). Here, organisms produced a mean resource level of 7.9 units per cell, enabling 85% of organisms to avoid being killed, as shown in Figure 4(b).

Well-Mixed Environment. The previous experiment demonstrated the benefits of kin selection to cooperative behaviors in populations with limited dispersal. Our next set of experiments examined how cooperation is affected when dispersal is increased. To

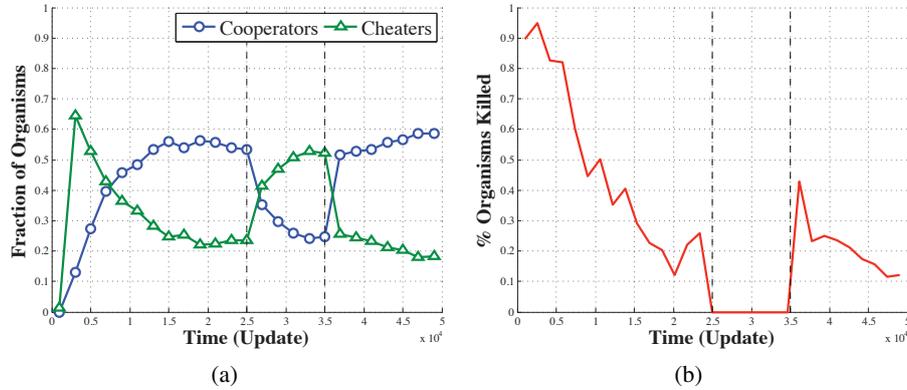


Fig. 6. Results when the kill event was suspended between updates 25,000 and 35,000: (a) Mean distribution of behaviors (b) Mean fraction of organisms killed within the target region

achieve this, offspring were no longer placed into a neighboring cell, but instead into a random cell in the environment, greatly reducing the possibility of remaining near relatives. In these runs, cooperation emerged and persisted, resulting in approximately 40% of the population producing resource (Figure 5(a)). Interestingly, cheaters made up the same portion of the population. These results match those modeled with Snowdrift [8], where one's best strategy is to cooperate if their opponent cheats, and vice versa. These strategies enabled approximately 72% of organisms in kill regions to be spared, as shown in Figure 5(b).

Dynamic Conditions. To determine how populations would react to the absence and return of adversity, we evolved organisms in an environment where the kill event was suspended during updates 25,000 through 35,000. In these runs, offspring were once again placed into neighboring cells, and the kill event used a 5-cell radius with a 3-unit threshold. Figures 6(a) and 6(b) plot the strategies used by organisms and the fraction of organisms killed, respectively. During the respite, mean cellular resource levels fell below the kill threshold, leaving organisms vulnerable. Upon the return of the kill action, cooperative resource production re-evolved in 19 of the 20 populations.

4 Conclusions

This work has demonstrated that cooperative behaviors can evolve in populations exposed to adverse environments. Specifically, organisms evolved to complete tasks that did not provide direct bonuses as did other tasks, but instead produced a resource that helped prevent that organism and its neighbors from being killed. Kin selection was observed to provide incentive for developing such strategies when dispersal is limited. The increase in resource production by such populations indicates that this behavior is not simply selfish; rather, the behavior is mutually beneficial. We have also seen that

populations can quickly re-gain cooperative strategies in reaction to the return of adverse conditions after a calm period. Insights into cooperative behaviors offer potential applications as treatments for infectious disease, motivating future research in this area.

Acknowledgments

The authors gratefully acknowledge the contributions of Ben Beckmann, Art Covert, Tracy Teal and the MSU Devolab. This work was supported by NSF Grants CCF-0750787, CNS-0751155, CCF-0820220; by U.S. Army Grant W911NF-08-1-0495; and by a Quality Fund Grant from Michigan State University.

References

1. Kaplan, H., Hill, K.: Food sharing among ache foragers: Tests of explanatory hypotheses. *Current Anthropology* 26(2), 223 (1985)
2. Symonds, M.R., Elgar, M.A.: The evolution of pheromone diversity. *Trends in Ecology & Evolution* 23(4), 220–228 (2008)
3. Tofilski, A., et al.: Preemptive defensive self-sacrifice by ant workers. *The American Naturalist* 172(5), E239–E243 (2008)
4. Hamilton, W.D.: The genetical evolution of social behaviour. I.II. *Journal of Theoretical Biology* 7(1), 1–52 (1964)
5. Williams, P., et al.: Quorum sensing and the population-dependent control of virulence. *Philosophical Transactions of the Royal Society B: Biological Sciences* 355(1397), 667–680 (2000)
6. Griffin, A.S., West, S.A., Buckling, A.: Cooperation and competition in pathogenic bacteria. *Nature* 430, 1024–1027 (2004)
7. Burnham, J., Collart, S., Daft, M.: Myxococcal predation of the cyanobacterium *Phormidium luridum* in aqueous environments. *Archives of Microbiology* 137(3), 220–225 (1984)
8. Doebeli, M., Hauert, C.: Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game. *Ecol. Lett.* 8(7), 748–766 (2005)
9. Nowak, M.A., May, R.M.: Evolutionary games and spatial chaos. *Nature* 359, 826–829 (1992)
10. Pestelacci, E., Tomassini, M.: Hawks and doves in an artificial dynamically structured society. In: *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pp. 466–473 (2008)
11. Gore, J., Youk, H., van Oudenaarden, A.: Snowdrift game dynamics and facultative cheating in yeast. *Nature*, 253–256 (2009)
12. Neumann, J.V.: *Theory of self-reproducing automata*. University of Illinois Press, US (1966)
13. Alonso, J., Fernández, A., Fort, H.: Prisoner's dilemma cellular automata revisited: Evolution of cooperation under environmental pressure. *Journal of Statistical Mechanics: Theory and Experiment* 2006(06), P06013 (2006)
14. Rechenberg, I.: *Evolutionsstrategie: Optimierung technischer systeme nach prinzipien der biologischen evolution*. Frommann-Holzboog (1973)
15. Ofria, C., Wilke, C.: Avida: A software platform for research in computational evolutionary biology. *Artificial Life* 10(2), 191–229 (2004)
16. McKinley, P.K., Cheng, B., Ofria, C., Knoester, D., Beckmann, B.E., Goldsby, H.: Harnessing digital evolution. *IEEE Computer* 41(1), 54–63 (2008)

17. Connelly, B.D., McKinley, P.K., Beckmann, B.E.: Evolving cooperative pheromone usage in digital organisms. In: Proceedings of IEEE Symposium on Artificial Life, pp. 184–191 (2009)
18. Goings, S., Clune, J., Ofria, C., Pennock, R.T.: Kin-selection: The rise and fall of kin-cheaters. In: Proceedings of the Ninth International Conference on Artificial Life, pp. 303–308 (2004)
19. Knoester, D.B., McKinley, P.K.: Cooperative network construction using digital germlines. In: Proceedings of the Genetic and Evolutionary Computation Conference, Atlanta, Georgia (2008)
20. Lenski, R.E., Ofria, C., Pennock, R.T., Adami, C.: The evolutionary origin of complex features. *Nature* 423, 139–144 (2003)
21. Beckmann, B.E., McKinley, P.K.: Evolution of adaptive population control in multi-agent systems. In: Proceedings of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems, Venice, Italy, pp. 181–190 (2008)
22. Rainey, P.B., Rainey, K.: Evolution of cooperation and conflict in experimental bacterial populations. *Nature* 425, 72–74 (2003)
23. Carlson, M., Botstein, D.: Two differentially regulated mRNAs with different 5' ends encode secreted and intracellular forms of yeast invertase. *Cell* 28(1), 145–154 (1982)
24. West, S.A., Griffin, A.S., Gardner, A., Diggle, S.P.: Social evolution theory for microorganisms. *Nature Reviews: Microbiology* 4, 597–607 (2006)
25. Doebeli, M., Hauert, C., Killingback, T.: The evolutionary origin of cooperators and defectors. *Science* 306, 859–862 (2004)