

Effects of Biased Group Selection on Cooperative Predation in Digital Organisms

Daniel J. Couvertier and Philip K. McKinley
BEACON Center for the Study of Evolution in Action
Department of Computer Science and Engineering
Michigan State University
East Lansing, Michigan 48824
{couverti,mckinley}@cse.msu.edu

ABSTRACT

A key issue in cooperative task completion is team composition. Prior studies have addressed two ends of a spectrum, with homogeneous teams on one end and heterogeneous teams on the other. In this paper we explore a space in between. In *biased group selection*, subpopulations compete against one another with respect to a cooperative task, but an external bias favors the genes of those individuals actually participating in the task. We evaluate this selection model on a cooperative predation task in digital organisms, where feasible solutions can be carried out by either homogeneous or heterogeneous teams. Our results show that, consistent with earlier studies, homogeneous teams tend to find better overall solutions than their heterogeneous counterparts. However, populations comprising teams with some degree of heterogeneity found solutions more frequently. Effectively, while evolution pushed heterogeneous teams toward functional homogeneity for this particular task, heterogeneity with a selection bias proved more effective at exploring the search space.

Categories and Subject Descriptors

I.2.8 [Computing Methodologies]: Artificial Intelligence—*Problem Solving, Control Methods, and Search*

General Terms

Experimentation

Keywords

Artificial life, digital evolution, multi-agent system, cooperative behavior, self-organization, group selection, predation.

1. INTRODUCTION

Cooperative behavior within and among natural organisms is one of the most important and pervasive phenomena found on Earth. Now, through evolutionary computation, we are “harnessing” the benefits of these behaviors to advance several emerging technologies, such as multi-agent systems and swarm robotics, that require cooperation among multiple, autonomous entities [2]. Like natural organisms, these systems need to adapt to dynamic and adverse conditions, conserve energy, and compensate for failures, all while cooperating to meet global objectives. A logical approach to developing these systems is to look to biology for selection methods

that can be codified and used to evolve cooperative behaviors *in silico*. An early explanation for the emergence of cooperation in natural systems was *group selection*, which proposed that a team comprising selfish, “cheating” individuals would die out due to the over-exploitation of resources, whereas one composed of cooperative members would constrain their behavior for the benefit and survival of the group as a whole [1]. While the concept is intriguing, later studies show that, in nature, groups do not go extinct quickly enough, and individuals move between groups at a sufficiently high rate, to render group selection an implausible explanation for the emergence of cooperation.

Fortunately, computational methods used to search for solutions to engineering problems are not limited to biological rules. Studies in evolutionary robotics and related areas have shown that group selection is indeed an effective way to obtain cooperative behaviors. Moreover, teams where all members are genetically identical have been shown to be highly effective in evolving cooperation, as the risk of cheating is diminished [4].

In this work, we investigate the role of team relatedness and individual selection in evolving cooperative predation behaviors in digital organisms. We focus on the relatively unexplored space of a spectrum where homogeneous teams reside at one end, and at the other, heterogeneous teams where intra-team genetic similarity is not necessarily greater, on average, than inter-team similarity. In between lie varying levels of group relatedness which are determined by user-defined parameters. In *biased group selection*, members of a group that contribute to the completion of tasks will have a higher fitness within the group, making them more likely to replicate. With this study, we hope to increase knowledge about the relationship between selection and team composition in computational evolution, and help to explain why it is so difficult to evolve heterogeneous teams with performance comparable to that of homogeneous teams.

2. METHODS

The task of cooperative predation, where multiple agents are needed to successfully hunt a prey, allows us to compare groups with varying levels of heterogeneity. Evolution could feasibly find solutions for this task using strategies that are evolved with either homogeneous or heterogeneous groups.

The Avida digital evolution platform [3] provided a framework suitable for carrying out the experiments in this study. Avida provides an environment through which digital organisms can move and interact. Each organism comprises a circular list of instructions (its *genome*), and a virtual CPU on which they are executed. We added new instructions for this study, including ATTACK, whereby a predator launches an attack on the prey; and HAIL, which en-

ables an organism to draw the attention of other organisms. When a hunter hails, all other predators within a certain range from the source will automatically turn their facing towards the hailer. This can be considered an instinctual behavior, as in animals that turn their heads towards the source of a loud sound.

For the cooperative predation task in the context of Avida, a group of organisms is placed in a two-dimensional toroidal world, where they have to locate and successfully attack a single, stationary prey. An attack is only successful if the number of predators in the prey vicinity is above a certain threshold, or *quorum*, when one of those predators executes the ATTACK instruction. The reward for a successful attack translates to energy that allows the group to replicate. The penalty for an unsuccessful attack is the “flight” of the prey to a new random location in the world, thereby wasting efforts put by the organisms to organize around it, forcing them to start anew. Scenarios depicting cases in which the attack would fail and succeed are demonstrated in Figure 1.




Figure 1: Attack scenarios. The prey is the 3×3 cell shaded area. The required quorum is 5. If an organism on the prey executes an ATTACK instruction in (a), the attack fails because quorum is not met. If the ATTACK is issued in (b), it succeeds because quorum is met.

We tested several types of selection, from homogeneous groups at one end, to heterogeneous (unbiased) groups at the other. In between these two extremes lie various degrees of biased group selection. This form of selection combines both group selection and individual selection. When the group completes the task, it is selected to replicate and thus outcompetes another group in the run. However, when composing the new group, there is competition *within* the group in the form of two levels of external bias. The first is an automatic guarantee of producing one offspring (which is exposed to mutations during the replication process) for those organisms that participated in the task. The remaining offspring are selected using the second level of bias, a user-defined probability that the parent is selected from among the participants. If the probability is 100%, then 100% of the remaining offspring population will be the progeny of participants. If the rate is 50%, then that fraction of the remaining offspring slots will be populated with organisms whose parents are participants, while the remaining slots will be occupied by the offspring of non-participants, and so on.

3. SAMPLE RESULTS

A total of 50 runs were performed for each of five treatments (Homogeneous, Biased-100%, Biased-50%, Biased-0%, and Unbiased). In each run, 400 groups of 10 organisms would compete during 75,000 Avidian time steps, or *updates*, with evolution enabled. Afterwards, an ecological period of 25,000 updates without evolution would realize competition of the remaining groups, allowing the fastest team to populate all 400 worlds.

The results were evaluated by comparing the number of suc-

cessful attacks (*kills*) completed per group per update within the best run in each treatment. We also considered the fraction of the 50 runs that successfully evolved a *persistent* solution within the treatment. We defined persistent solutions as those where, during the last 100 updates of the ecological period, the entire population of 400 groups carried out at least one successful attack. We also determined the composition of each team by calculating the average Levenshtein *distance* (minimum number of edits required to make two genomes equal) between all members in a group. A larger distance suggests greater heterogeneity among team member genomes. The resulting values for each of the treatments in a 30×30 cell environment are shown in Table 1.

Table 1: Treatment Comparison.

Treatment	Kills	Persistence	Distance
Homogeneous	0.53	48%	0.0%
Biased-100%	0.43	84%	4.6%
Biased-50%	0.39	64%	11.1%
Biased-0%	0.25	60%	18.0%
Unbiased	0.33	58%	2.2%

The Homogeneous treatment produced the overall best individual run. However, we also observed that the rate of persistent solutions was higher in all the Heterogeneous treatments than in the Homogeneous treatment, meaning they all produced more runs with persistent solutions, even though their individual performances were lower. Considering only the Biased treatments, a greater bias produced a more effective champion run and a higher rate of persistent solutions. The Unbiased treatment did not adhere to this general trend by providing an individual run that outperformed Biased-0%.

In general, we found that toward the homogeneous end of the spectrum, groups are very good at enhancing solutions, but compared to groups with heterogeneity, are not as effective at exploring the search space. On the other hand, heterogeneous treatments find more solutions, but those solutions are not as effective as those found by homogeneous groups. We also found that when presented with the choice of creating a homogeneous or heterogeneous group composition, evolution opts for homogeneity, as demonstrated by our results where the best-performing heterogeneous groups were effectively composed of functionally identical individuals.

4. ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation grants CNS-0915855, DBI-0939454, CCF-0820220, and CNS-0751155; and by U.S. Army Grant W911NF-08-1-0495.

5. REFERENCES

- [1] M. E. Borrello. The rise, fall and resurrection of group selection. *Endeavour*, 29(1):43–47, 2005.
- [2] P. K. McKinley, B. H. Cheng, C. Ofria, D. Knoester, B. Beckmann, and H. Goldsby. Harnessing digital evolution. *IEEE Computer*, 41(1), January 2008.
- [3] C. Ofria and C. O. Wilke. Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10:191–229, 2004.
- [4] M. Waibel, L. Keller, and D. Floreano. Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660, 2009.