Balancing Selection Pressures, Multiple Objectives, and Neural Modularity to Coevolve Cooperative Agent Behavior

Alex C. Rollins
Dept. of Mathematics and Computer Science
Southwestern University
Georgetown, Texas 78626 USA
rollinsa@southwestern.edu

Jacob Schrum
Dept. of Mathematics and Computer Science
Southwestern University
Georgetown, Texas 78626 USA
schrum2@southwestern.edu

ABSTRACT

Previous research using evolutionary computation in Multi-Agent Systems indicates that assigning fitness based on team vs. individual behavior has a strong impact on the ability of evolved teams of artificial agents to exhibit teamwork in challenging tasks. Such research made use of single-objective evolution, but when multiobjective evolution is used, populations can be subject to individual-level objectives, team-level objectives, or combinations of the two. This paper explores the performance of cooperatively coevolved teams of agents controlled by artificial neural networks subject to these types of objectives. Because of the tension between individual and team behaviors, multiple modes of behavior can be useful, so the effect of modular neural networks is also explored. Results demonstrate that fitness rewarding individual behavior is superior to fitness rewarding team behavior, despite being applied to a cooperative task. However, networks with multiple modules can discover intelligent behavior, regardless of which type of objectives are used.

CCS CONCEPTS

•**Computing methodologies** → **Artificial life;** *Cooperation and coordination; Neural networks;*

KEYWORDS

Artificial life, Co-evolution, Multi-agent systems, Neural networks

ACM Reference format:
Alex C. Rollins and Jacob Schrum. 2017. Balancing Selection Pressures, Multiple Objectives, and Neural Modularity to Coevolve Cooperative Agent Behavior. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3067695.3075979

1 INTRODUCTION

Past research has studied the effects of selection pressures [9], co-evolution [2], modular neural networks [5], and multiple objectives [4] in the evolution of complex agent behavior, but none of this research studies all at once. This paper explores how these concepts work in tandem. Combinations of different types of multiobjective selection and numbers of network output modules show how these components interact in the evolution of cooperative behavior.

GECCO '17 Companion, Berlin, Germany

DOI: http://dx.doi.org/10.1145/3067695.3075979

The concept of selection pressures stems from research on the rewarding of individual vs. team behavior [9]; does one agent get credit for good outcomes, or does the whole team? Rather than have only one objective assigning credit in these ways, one can apply Pareto-based multiobjective optimization [1]. Multiobjective optimization is natural in domains that would otherwise use a single complex objective consisting of several components. Multiple objectives also make it easy to go beyond a simple either/or comparison, and also study the combination of individual and team objectives.

Skilled cooperation depends on balancing selfish and cooperative actions, both across different team members and by individuals. Evolved neural networks control agents in this paper, and such agents can more easily exhibit multimodal behavior using modular network structures [6]. Past research has demonstrated the ability of modular networks to succeed in domains with homogeneous teams [5], but this paper utilizes cooperative coevolution with separate and distinct sub-populations for each evolved team member. This approach allows individuals to develop specialized behaviors.

2 PREDATOR/PREY EXPERIMENT

These issues are studied in a torus-shaped predator/prey grid world, variants of which have been used by researchers in many ways [2, 8]. All experiments feature a team of three evolved predators trying to catch two scripted prey agents that flee the nearest predator.

Predators must work together in order to herd and capture prey. Selfishly chasing the prey generally leads to all agents going in circles around the torus. The success of each individual at least partially relies on the success of the team. Consequently, predators develop specializations as valuable members of the team. Some common roles that emerge in successful teams are *blocker*, *herder*, and *aggressor*. The blockers do not move very much but align themselves at a distance with the side to side movement of the more aggressive predators so that they can force the prey to run toward the blocker. The herders work to keep the prey in front of the aggressors by running parallel to the prey's direction of movement, so that it does not slip by to one side. The job of the aggressor is to simply close the gap on the prey as quickly as it can.

Predator agents were evolved using Modular Multiobjective Neuro-Evolution of Augmenting Topologies (MM-NEAT [6][1]), which combines the multiobjective evolutionary algorithm Non-Dominated Sorting Genetic Algorithm-II (NSGA-II [1]) and standard NEAT [7]. MM-NEAT also allows for the evolution of networks with multiple output modules. MM-NEAT has been extended in this paper to support cooperative coevolution of separate sub-populations.

[1] Download at http://nn.cs.utexas.edu/?mm-neat

	IndCatch	IndDist	TeamCatch	TeamDist
Individual Selection	1	2	0	0
Team Selection	0	0	1	2
Both Selection	1	2	1	2

Table 1: Objectives For Each Sub-population. Table shows the number of fitness functions for each sub-population in each experiment. Numbers are the same whether networks have one or two modules. Ind stands for Individual Selection, and Team stands for Team Selection. Catch indicates the maximization of the number of prey caught. Dist indicates the minimization of distances between predators and prey (distinct fitness functions measure distance to each prey agent).

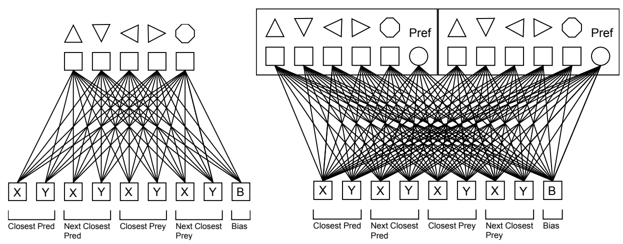

(a) One Module Network **(b) Two Module Network**

Figure 1: Starting Network Configurations. New populations start with no hidden neurons, but each output is fully connected to all inputs. (a) Networks with one module have outputs for moving up, down, left, and right, and staying still. Inputs are x/y offsets to each other agent, followed by a constant bias of 1.0. The agent inputs are grouped into predators and prey, and sorted according to proximity in terms of Manhattan Distance. (b) Networks with two modules use the same inputs, but have two distinct output modules. Each module has all of the outputs possessed by the one module network, as well as a preference neuron. For each set of inputs, the two module network will pick the action from the module whose preference neuron output is higher.

Full details of the experiments are discussed in an associated technical report [3]. The main focus of the experiments is in studying different combinations of fitness functions in conjunction with different numbers of network output modules. The individual fitness functions were combined in three ways, as summarized in Table 1. The three specific groups of fitness functions used focus either entirely on individual selection, entirely on team selection, or on both. Evolved networks could consist of either one or two preference modules (Figure 1). The experimental runs have the labels Individual1M, Individual2M, Team1M, Team2M, Both1M, and Both2M.

3 RESULTS

The results show that two modules are better than one module and that individual selection and combination setups are better than pure team selection. Fitness plots of the average number of prey caught by the champion team of each generation across 30 runs for each method are shown in Figure 2. Statistical comparisons indicate that Team1M is significantly inferior to other methods ($p < 0.05$). Further discussion of behaviors and results is available in the technical report [3], and videos of representative behaviors can be seen at southwestern.edu/~schrum2/SCOPE/predprey.html.

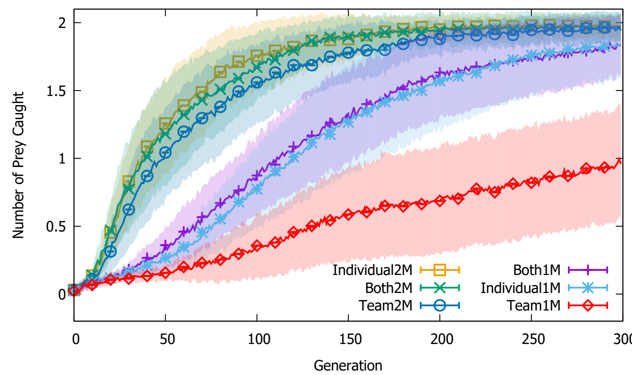

Figure 2: Average Number of Prey Caught For Each Approach. Average prey caught by champion teams across 30 runs of each method are plotted by generation with 95% confidence intervals shown. All 2M variants are superior to their 1M counterparts. Among 1M configurations, Both1M and Individual1M are superior to Team1M.

4 CONCLUSION

The predator/prey task is an interesting domain requiring teamwork and specialization. Results demonstrate that multimodal networks are extremely helpful, and that individual selection can be superior to team selection in a cooperative task when coevolution across distinct sub-populations is used. The combination of both types of selection pressures via multiobjective optimization was also effective. Using multiobjective optimization with multimodal networks could produce interesting cooperative behavior in more complex domains as well.

ACKNOWLEDGMENTS

This work is supported in part by the Howard Hughes Medical Institute through the Undergraduate Science Education Initiative Program under Grant No.: 52007558.

REFERENCES

[1] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2002. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6 (2002), 182–197.
[2] Aditya Rawal, Padmini Rajagopalan, and Risto Miikkulainen. 2010. Constructing Competitive and Cooperative Agent Behavior Using Coevolution. In *Conference on Computational Intelligence and Games*. IEEE.
[3] Alex C. Rollins and Jacob Schrum. 2017. Balancing Selection Pressures, Multiple Objectives, and Neural Modularity to Coevolve Cooperative Agent Behavior. (2017). arXiv:1703.08577
[4] Jacob Schrum and Risto Miikkulainen. 2008. Constructing Complex NPC Behavior via Multi-Objective Neuroevolution. In *Artificial Intelligence and Interactive Digital Entertainment*. 108–113.
[5] Jacob Schrum and Risto Miikkulainen. 2012. Evolving Multimodal Networks for Multitask Games. *IEEE Transactions on Computational Intelligence and AI in Games* 4, 2 (2012), 94–111.
[6] Jacob Schrum and Risto Miikkulainen. 2016. Discovering Multimodal Behavior in Ms. Pac-Man through Evolution of Modular Neural Networks. *IEEE Transactions on Computational Intelligence and AI in Games* 8, 1 (2016), 67–81.
[7] Kenneth O. Stanley and Risto Miikkulainen. 2002. Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation* 10 (2002), 99–127.
[8] Ming Tan. 1993. Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents. In *International Conference on Machine Learning*. 330–337.
[9] Markus Waibel, Laurent Keller, and Dario Floreano. 2009. Genetic Team Composition and Level of Selection in the Evolution of Multi-Agent Systems. *IEEE Transactions on Evolutionary Computation* 13, 3 (2009), 648–660.