Ad Hoc Teaming Through Evolution

Joshua Cook
Oregon State University
Corvallis, Oregon
cookjos@oregonstate.edu

Kagan Tumer
Oregon State University
Corvallis, Oregon
ktumer@oregonstate.edu

ABSTRACT

Cooperative Co-evolutionary Algorithms effectively train policies in multiagent systems with a single, statically defined team. However, many real-world problems, such as search and rescue, require agents to operate in multiple teams. When the structure of the team changes, these policies show reduced performance as they were trained to cooperate with only one team. In this work, we solve the cooperation problem by training agents to fill the needs of an arbitrary team, thereby gaining the ability to support a large variety of teams. We introduce Ad hoc Teaming Through Evolution (ATTE) which evolves a limited number of policy types using fitness aggregation across multiple teams. ATTE leverages agent types to reduce the dimensionality of the interaction search space, while fitness aggregation across teams selects for more adaptive policies. In a simulated multi-robot exploration task, ATTE is able to learn policies that are effective in a variety of teaming schemes, improving the performance of CCEA by a factor of up to five times.

CCS CONCEPTS

• **Computing methodologies** → **Genetic algorithms**; **Machine learning approaches**;

KEYWORDS

Multiagent, Co-evolution, Fitness Assignment

ACM Reference Format:
Joshua Cook and Kagan Tumer. 2021. Ad Hoc Teaming Through Evolution. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459560

1 INTRODUCTION

Cooperative co-evolutionary algorithms (CCEA) have proven effective in learning policies in multiagent systems with statically defined teams [6–8]. Often, real-world agents must work with teams other than a single static team. CCEA does not promote learning across multiple teams, leading to policies that do not transfer well to other teams. For example, CCEA could train a series of robots to perform search and rescue. However, if one of these robots needed to work with a team in another area, it would struggle as it does not have any experience working with this new team.

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07…$15.00
https://doi.org/10.1145/3449726.3459560

Developing this ability to work with a variety of teams is the goal of ad hoc teaming [9]. Learning ad hoc teaming policies is challenging as it compounds the number of interactions an agent needs to learn. Agent types reduce this interaction space by assuming that agents are not entirely unique, and instead fit the mold of one of a few agent types [2]. Rather than learning to coordinate with every agent separately, each agent learns to work with a limited number of agent types [3]. Current methods of learning ad hoc teaming policies include ATSIS and PLASTIC, but are not designed for domains where a team of agents must learn to cooperate in a continuous state-action space [4, 5].

In this work, we propose Ad hoc Teaming Through Evolution (ATTE) to train policies that operate effectively in multiple teams. This method first divides the problem into sub-tasks each agent may need to complete. Using an extension of CCEA, agent types are evolved to select these tasks in multiple teams. These agent types are continually randomly assigned to each agent in the system to simulate team variety. Then, fitness aggregation assigns fitnesses to promote better teamwork. Random agent type assignment and fitness aggregation were included to promote learning across different teams, thus generating more robust teaming policies.

The main contribution of this work is to learn multiple policies that enable agents to be members of diverse teams. Experimental results in an exploration domain show that ATTE produced policies which consistently scored up to five times higher than the baseline in terms of average performance across teams. These results confirm that ATTE learns adaptive policies that enable more robust teaming.

2 AD HOC TEAMING THROUGH EVOLUTION

We introduce Ad hoc Teaming Through Evolution (ATTE) as a method of learning robust cooperation across varying teams. ATTE evolves a group of agent types that are randomly distributed to the agents in the system. Through this random teaming and fitness aggregation, these agent types learn to coordinate with each other in multiple teams.

ATTE begins with an environment to receive fitnesses from, along with multiple populations of policies. Each population represents an agent type that can be copied to any agent in the system. At the beginning of each episode, we randomly assign an agent type to each agent to simulate team variety. Then, one policy is randomly selected from each population and copied to the agents of that type. Each policy is evaluated in the environment and the global fitness is added to the history of the policy being tested. This is repeated until all members of each policy have had one fitness added to each member's fitness history. Then the average of this history is the fitness assigned to each member of the population. This method of fitness aggregation reduces the fitness noise introduced by the random teaming combinations.

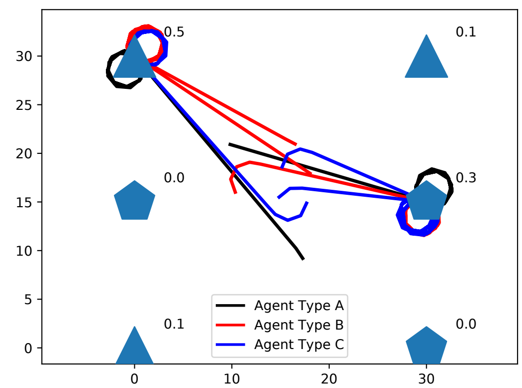

Figure 1: The paths of the 8 agents are shown, moving towards and viewing the two highest PoIs. The colors of the paths represent the agent type. The PoI shapes and values are also shown.

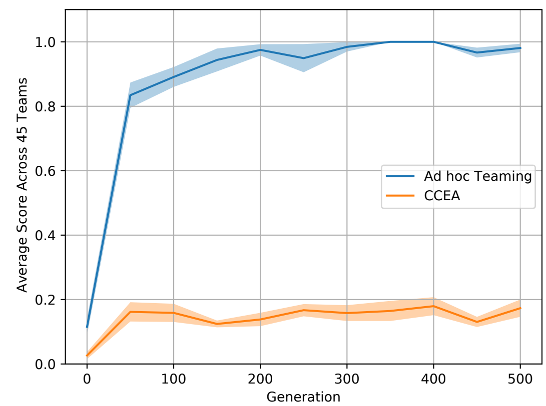

Figure 2: Fitness curves for teams of 8 agents, averaged across 45 team combinations. ATTE was compared to CCEA, with ATTE consistently achieving a max fitness across multiple teams.

At the end of an episode, the best policies are selected and copied. The copies mutate and their fitness histories are reset as their previous history no longer an accurate representation of their performance.

3 RESULTS

We test our algorithm on the Rover Domain, a multi-robot exploration domain [1]. In the basic formulation of this problem, a group of rovers must explore a desired two-dimensional location and successfully observe various points of interest (PoI). These PoI are represented as n-sided polygons that require observation of each side by a separate rover. Rovers are only evaluated for each PoI that is fully observed. Each PoI contains a value proportional to its importance. To solve this task, agents must form groups around the PoIs and observe them, while prioritizing the higher valued locations.

Each robot in this domain contains two types of continuous sensors: PoI sensors and rover sensors. Each sensor is capable of viewing a 90-degree area from the robot. To fully view their surroundings, each robot contains four of each type of sensor, spaced evenly around the robot, for a total of eight sensors. Robots can take continuous actions consisting of linear and angular movement along the plane of the world.

In our experiments, we employ a modified version of the Rover Domain in which the PoIs have an additional shape attribute. The robots themselves only have the ability to observe a PoI of a single specific shape. The robots are also aware of which shape of PoI they can view.

Performance in each experiment is defined as the average global fitness received by all combinations of three agent types distributed amongst eight agents. The resulting score represents the agents' ability to work well in every possible team configuration. Each performance score is the average of eight statistical runs with error shading representing the standard error. This performance metric is scaled to a maximum value of 1.0.

In this experiment, six PoIs were divided into two shapes. Half the agents could view one shape and the other half could view the other shape. Each PoI needed three agents for successful observation. To complete this task, each agent needed to move towards the highest valued PoI with a shape they could observe, shown in figure 1.

ATTE is compared to a baseline CCEA algorithm, each with a population size of 50. Figure 2 shows the learned performance over time. ATTE significantly outperforms CCEA, reaching an optimal score. The learned policies were able to achieve the maximum score across all 45 teaming combinations, while CCEA struggled. This shows ATTE's ability to produce policies that perform well in multiple teams.

ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation under grant No. IIS-1815886 and the Air Force Office of Scientific Research under grant No. FA9550-19-1-0195.

REFERENCES

[1] Adrian K Agogino and Kagan Tumer. 2008. Analyzing and visualizing multiagent rewards in dynamic and stochastic domains. *Autonomous Agents and Multi-Agent Systems* 17, 2 (2008), 320–338.
[2] Stefano V Albrecht and Peter Stone. 2018. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* 258 (2018), 66–95.
[3] Nolan Bard, Michael Johanson, Neil Burch, and Michael Bowling. 2013. Online implicit agent modelling. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. 255–262.
[4] Samuel Barrett and Peter Stone. 2015. Cooperating with Unknown Teammates in Complex Domains: A Robot Soccer Case Study of Ad Hoc Teamwork.. In *AAAI*, Vol. 15. Citeseer, 2010–2016.
[5] Shuo Chen, Ewa Andrejczuk, Athirai Aravazhi Irissappane, and Jie Zhang. 2019. ATSIS: achieving the ad hoc teamwork by sub-task inference and selection. (2019).
[6] Mitchell Colby, Matt Knudson, and Kagan Tumer. 2014. Multiagent flight control in dynamic environments with cooperative coevolutionary algorithms. In *2014 AAAI Spring Symposium Series*.
[7] Atil Iscen, Adrian Agogino, Vytas SunSpiral, and Kagan Tumer. 2014. Flop and roll: Learning robust goal-directed locomotion for a tensegrity robot. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2236–2243.
[8] Atil Iscen, Ken Caluwaerts, Jonathan Bruce, Adrian Agogino, Vytas SunSpiral, and Kagan Tumer. 2015. Learning tensegrity locomotion using open-loop control signals and coevolutionary algorithms. *Artificial life* 21, 2 (2015), 119–140.
[9] Peter Stone, Gal A Kaminka, Sarit Kraus, Jeffrey S Rosenschein, et al. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination.