Hyb-CCEA: Cooperative Coevolution of Hybrid Teams

Jorge Gomes
BioMachines Lab &
Instituto de Telecomunicações
& Faculdade de Ciências,
Universidade de Lisboa, BioISI
Lisbon, Portugal
jgomes@di.fc.ul.pt

Pedro Mariano
Faculdade de Ciências,
Universidade de Lisboa, BioISI
Lisbon, Portugal
plmariano@fc.ul.pt

Anders Lyhne Christensen
BioMachines Lab &
Instituto de Telecomunicações
& Instituto Universitário de
Lisboa (ISCTE-IUL)
Lisbon, Portugal
anders.christensen@iscte.pt

1. INTRODUCTION

Cooperative coevolution algorithms (CCEAs) evolve solutions that consist of interacting, coadapted components [5]. CCEAs are capable of evolving a heterogeneous set of cooperating agent behaviours, where each agent can have a specialised behaviour. CCEAs are, however, associated with inherent scalability issues [4], since each agent behaviour typically evolves in a separate population. The computational complexity therefore increases at least linearly with the number of agents. Moreover, credit assignment issues might arise in large multiagent systems: it can be hard to assess the contribution of each individual agent to the performance of the team as a whole.

In multiagent systems with large numbers of agents, successful solutions may contain agents with similar behaviours [3]. Nonetheless, if each population is isolated, as it is typically the case in CCEAs, similar agent behaviours might have to be learned multiple times in different populations. One way to increase the scalability of multiagent learning is through the reduction of the heterogeneity in the system [4]. Partially heterogeneous multiagent systems, also known as *hybrids*, are composed of multiple homogeneous sub-teams — sub-teams in which all agents have identical controllers. By allowing partial heterogeneity, the number of agent controllers that need to be evolved is reduced, improving the scalability of the learning process. Most previous studies on the evolution of partially heterogeneous teams are focused on *team learning*, where one single genome encodes the behaviour and/or the composition of the whole team [1]. The cooperative coevolution of hybrid multiagent systems is still unexplored.

This paper is a summary of the following work:
J. Gomes, P. Mariano, and A. L. Christensen. Cooperative Coevolution of Partially Heterogeneous Multiagent Systems. In *International Conference on Autononomous Agents & Multiagent Systems*, pages 297–305. IFAAMAS, 2015.

GECCO '15 Companion, July 11 - 15, 2015, Madrid, Spain

DOI: http://dx.doi.org/10.1145/2739482.2768495

2. HYB-CCEA

We propose *Hyb-CCEA*, an extension of a CCEA to evolve agent controllers for physically homogeneous, behaviourally heterogeneous multiagent systems. In traditional applications of CCEAs in multiagent systems, there is a one-to-one relation between agents and populations. Our approach departs from this concept: we allow population individuals to encode a controller that can be used by multiple agents. Each population thus becomes responsible for the evolution of a homogeneous sub-team inside the larger heterogeneous team, not just one specific agent.

In *Hyb-CCEA*, each population is assigned to a subset of agents (see Figure 1). The number of population individuals is constant across all populations, and two different populations cannot be assigned to the same agent. The rest of the coevolutionary evaluation operates the same way as a traditional CCEA [5]: individuals are joined with representative individuals from the other populations for evaluation, and the individual being evaluated receives the fitness score that the team as a whole obtained.

The distinctive aspect of *Hyb-CCEA* is that it does not assume that the optimal number of sub-teams and their composition are known beforehand: we extend the CCEA so that the number of composition of the sub-teams is also under evolutionary control. The evolutionary process can start either fully homogeneous (one population assigned to all agents) or fully heterogeneous (one population for each agent). Different levels of heterogeneity can then be explored throughout the evolutionary process. To this end, we propose: (i) a procedure for merging two populations, thus decreasing the heterogeneity of the system (Figure 2); and (ii) a procedure for splitting a population, thus increasing heterogeneity (Figure 3). Merging and splitting operations can occur at any time during the evolutionary process. Both the split and merge procedures are implemented so that they have a minimal immediate impact in the performance of the teams, thus avoiding major evolutionary disruptions.

The proposed merging procedure allows agents evolving in separate populations to become genetically homogeneous. If two separate populations are evolving similar agent behaviours, they are merged into one population. To identify behaviour similarities between populations, the behaviour of each agent is characterised during the evaluation process, when cooperating in the team with which it was evaluated. The *agent behaviour characterisation* [2] should be provided by the experimenter. If two populations have a large overlap in agent behaviour space (above a fixed threshold), these populations are merged.

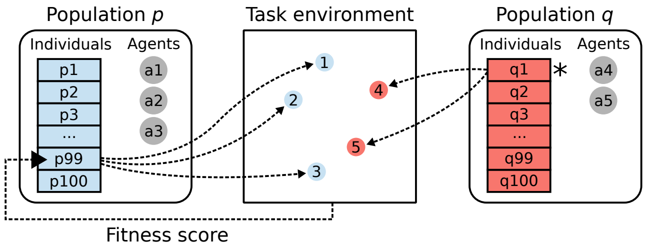

Figure 1: Evaluation phase: the individual under evaluation ($p99$) is joined with the representatives from the other populations. Each individual is assigned to the respective agents. The fitness of the whole team is assigned to individual being evaluated.

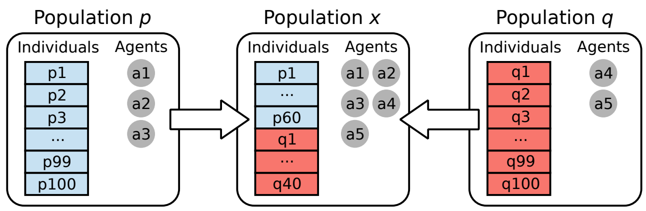

Figure 2: Merge procedure: the new population x replaces the two parents, p and q. The population x is formed by a subset of their individuals, and is assigned to all the parents' agents.

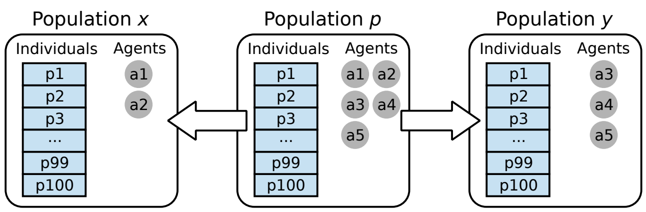

Figure 3: Split procedure: two new populations, x and y, replace the parent p. The populations x and y are copies of p, but each one is assigned to a disjoint set of agents.

The splitting mechanism is stochastic: since all agents assigned to a given population have a copy of the same controller, they will most likely display very similar behaviours, which makes it unfeasible to rely on behaviour characterisations to regulate splits. We resort to a stochastic approach where the chances of splitting a given population increase with the population age and the number of agents assigned to it. Such *blind* splits are feasible because they can later be reverted by the *merge procedure*. Both the merge and split procedures are regulated by a single time threshold that influences the frequency of such procedures over the evolutionary run.

3. RESULTS

We study the proposed approach in a simulated herding task, where each agent is controlled by a neural network. In this task, a group of physically homogeneous shepherds must corral one or more sheep. Additionally, one to three foxes are present, which try to capture the sheep, and must be kept away by the shepherds. Only the controllers for the shepherds are evolved. We explore multiple versions of the herding task, that involve different numbers of shepherds (from 5 to 10), where a single shepherd is sufficient to keep one fox away (W), and where multiple shepherds are necessary for each fox (E). The different task versions therefore require the evolution of different specialisations and levels of heterogeneity.

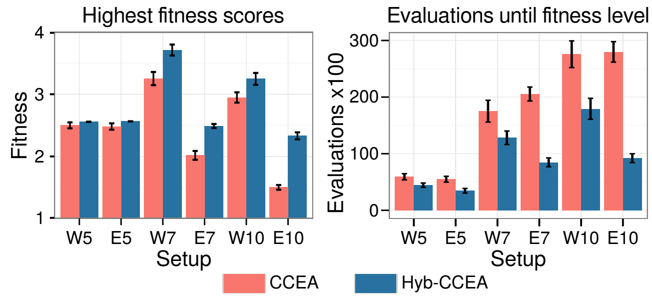

Figure 4: Left: Highest fitness scores achieved in each run. Right: Number of evaluations to reach the following fitness levels – W5/E5: 2.5; W7: 3.0; E7: 1.75; W10: 2.75; E10: 1.5.

We evaluate the performance improvements of *Hyb-CCEA* over a fully-heterogeneous CCEA. Figure 4 summarises the performance of each method. Our results were generally consistent across all task setups: *Hyb-CCEA* could reach good solutions for all tasks in significantly fewer evaluations, and it could also achieve significantly higher fitness scores in the end. We found that *Hyb-CCEA* can adapt the level of heterogeneity in the system to the task at hand, and that it is especially effective in task variants that require the formation of sub-teams.

Our experiments revealed that the performance gains of *Hyb-CCEA* become higher as the number of agents increases. In ongoing work, we are evaluating the proposed approach in additional tasks, requiring a higher number of agents. To the best of our knowledge, *Hyb-CCEA* is the first cooperative coevolution algorithm that allows for the emergence of partially heterogeneous teams. *Hyb-CCEA* reduced the number of coevolving populations without sacrificing solution quality. Our study opens new interesting avenues of research, since it brings together the emergence of team compositions and the advantages of cooperative coevolution.

Acknowledgements: This research is supported by Fundação para a Ciência e Tecnologia (FCT), with grant SFRH/BD/89095/2012 and projects UID/EEA/50008/2013, UID/Multi/04046/2013, and EXPL/EEI-AUT/0329/2013.

4. REFERENCES

[1] J. C. Bongard. The legion system: A novel approach to evolving heterogeneity for collective problem solving. In *Genetic Programming*, volume 1802 of *LNCS*, pages 16–28. Springer, 2000.

[2] J. Gomes, P. Mariano, and A. L. Christensen. Systematic derivation of behaviour characterisations in evolutionary robotics. In *International Conference on the Synthesis and Simulation of Living Systems (ALife)*, pages 212–219. MIT Press, 2014.

[3] G. Nitschke. Behavioral heterogeneity, cooperation, and collective construction. In *Congress on Evolutionary Computation (CEC)*, pages 1–8. IEEE Press, 2012.

[4] L. Panait and S. Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents & Multi-Agent Systems*, 11(3):387–434, 2005.

[5] M. A. Potter and K. A. D. Jong. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1):1–29, 2000.