

Evolution of Cooperation in Evolutionary Robotics : the Tradeoff between Evolvability and Efficiency

Arthur Bernard^{1,2}, Jean-Baptiste André³ and Nicolas Bredeche^{1,2}

¹Sorbonne Universités, UPMC Univ Paris 06, UMR 7222, ISIR, F-75005 Paris, France

²CNRS, UMR 7222, ISIR, F-75005, Paris, France

³Institut des Sciences de l'Evolution, Université de Montpellier, CNRS,
IRD, EPHE, CC065, Pl. E. Bataillon, 34095 Montpellier, France

Abstract

In this paper, we investigate the benefits and drawbacks of different approaches for solving a cooperative foraging task with two robots. We compare a classical clonal approach with an additional approach which favors the evolution of heterogeneous behaviors according to two defining criteria: the evolvability of the cooperative solution and the efficiency of the coordination behaviors evolved. Our results reveal a tradeoff between evolvability and efficiency: the clonal approach evolves cooperation with a higher probability than a non-clonal approach, but heterogeneous behaviors evolved with the non-clonal approach systematically show better fitness scores. We then propose to overcome this tradeoff and improve on both of these criteria for each approach. To this end, we investigate the use of incremental evolution to transfer coordination behaviors evolved in a simpler task. We show that this leads to a significant increase in evolvability for the non-clonal approach, while the clonal approach does not benefit from any gain in terms of efficiency.

Introduction

The evolution of cooperative actions in evolutionary robotics is as much a challenge as an interesting perspective for the design of complex collective systems (Doncieux et al., 2015). As such, it has been widely studied with very diverse approaches and objectives (Waibel et al., 2009; Hauert et al., 2014; Trianni et al., 2007; Lichocki et al., 2013). These works often use a clonal paradigm, where each robot has a copy of the same genome. This makes sense as this is the easiest way to ensure cooperation when individuals are expected to display similar behaviors. Moreover, using clones ensures maximal genetic relatedness between individuals, which is known to allow the evolution of altruism (Waibel et al., 2011; Montanier and Bredeche, 2011). As such, most research focus on increasing the probability for the cooperative solution to evolve.

In comparison, the nature of coordination behaviors and their influence on the quality of cooperation has yet to be thoroughly studied. In particular, interactions between clones in evolutionary robotics tend to produce homogeneous behaviors when most coordination tasks could benefit

from heterogeneous behaviors. This could be solved by using a non-clonal approach where paired individuals do not use the same genome, and could possibly evolve different behaviors more easily. However, a non-clonal approach may face a *chicken-and-egg* dilemma: multiple individuals need to behave in a particular fashion for cooperation to be rewarding, but no benefit can be extracted from this behavior unless all individuals cooperate. Therefore, without cooperating partners, those behaviors cannot be selected by the evolution as they do not benefit the individual. This is particularly problematic when a moderately rewarding solitary strategy overshadows a more rewarding, but also more challenging to evolve, cooperative strategy (Skyrms, 2004).

In this paper, we are interested in the comparison between clonal and non-clonal approaches on two different criteria:

- *Evolvability* of cooperation, which is the number of successful runs where cooperation evolved.
- *Efficiency* of cooperation. This criteria is focused on the quality of the evolved behaviors and is determined by the performance (w.r.t. fitness score) of the coordination strategies.

To that end, we design a foraging task where both cooperative and solitary strategies are possible but where cooperation provides the largest reward. This task is favored by the evolution of efficient cooperative behaviors and we compare different approaches on both criteria. The first approach is a straightforward implementation of the literature where interacting individuals are clones. In comparison, the second approach is a rather extreme implementation of a non-clonal approach: we use coevolution, where individuals are from two different populations, and where fitness scores are computed independently for each individual. While this scheme is typical of competitive coevolution (Floreano and Nolfi, 1997; Floreano et al., 1998; Panait and Luke, 2005), the nature of the task considered here makes cooperation more interesting, as both individuals can selfishly benefit from being cooperative.

In the next section, we describe the methods and experimental setup used throughout our study. Then, we com-

pare the results of the two approaches on the cooperative task. This first experiment reveals that both approaches face a tradeoff between evolvability and efficiency, where neither one dominates the other on both criteria. We investigate in a second experiment the possibility to overcome this tradeoff for both approaches. To this end, we use incremental evolution (Harvey et al., 1994; Urzelai et al., 1998) and evolve coordination in a simpler task in order to improve both the evolvability and efficiency on the target task for each approach. Finally, we discuss the implication of our findings in the last section, in particular with respect to maximizing evolvability and efficiency alike.

Methods

Two robotic agents are placed in a 800 by 800 units square arena with four solid walls and emptied from any obstacle apart from the targets in the foraging task. Each circular-shaped agent, with a diameter of 20 units, has a collection of sensors divided between a 90 degrees front camera and 12 uniformly distributed proximity sensors. The camera is composed of 12 rays with infinite range which indicate the type (coded on 3 bits) and proximity (one value in R^n) of the nearest object or agent in their direction. Proximity sensors have a range of twice the agent body’s diameter and are used to get the distance to any obstacle nearby such as solid objects, the other agent or walls. The two agents always begin the simulation next to one another at one end of the arena, whereas all the objects’ initial positions are randomized.

Agents can move freely in the environment and are controlled by a fully connected multi-layer perceptron with a single hidden layer, the topology of which does not change during the evolution. Inputs of this neural network are fed with all the data extracted from the sensors: 48 neurons for the camera (4 neurons for each of the 12 rays) and 12 neurons for the proximity sensors. A bias neuron, whose value is always 1, brings the total number of input neurons to 61. The hidden layer is comprised of 8 neurons and the output layer of 2 neurons giving the speed of each of the agent’s wheels. The activation function used is a sigmoid.

In each experiment, individuals evolved during a fixed amount of evaluations thanks to an evolutionary algorithm. Their genome consists of a collection of the 506 connection weights (as real-values) of the neural network and is initially randomized for each individual in the population. Three evaluation setups are used to compare the different approaches of our experiment:

- In the *control* setup, each individual is evaluated against 5 other randomly chosen individuals in the population except itself. Therefore we ensure that there is no genetic relatedness between individuals in each pair. However, it is not clear how the evolutionary algorithm itself may impact the population’s diversity, especially because elitism is used;

- In the *clonal* setup, each individual is evaluated once against a clone of itself. This setup is used to study the results of the classical clonal approach (Waibel et al., 2009; Hauert et al., 2014; Trianni et al., 2007; Lichocki et al., 2013). While previous works have shown on multiple occasions that cooperation can evolve, it is not clear if individuals can take different roles during a cooperative interaction;
- In the *coevolution* setup, each of the two individuals comes from two different coevolved populations. In this setup, each individual from one population encounters 5 random individuals from the other population. As pairing considers individuals from two separate populations, the evolution of heterogeneous behaviors is theoretically easier. As a matter of fact, such a relation where two very different individuals find a selfish interest in mutual cooperation is actually quite common in nature (Connor, 1995).

A pair of individuals then interact in the arena described before for a fixed number of simulation steps called a trial. Each trial is conducted 5 times to account for the random initial positions of the objects and decrease the influence of the initial conditions on the individuals’ performance.

The selection method used in the evolutionary algorithm is an elitist (10+10)-ES where the 10 best individuals in the population are used to generate 10 offsprings for the next generation. We use no recombination and therefore each offspring is a mutated copy of its parent. Mutations were sampled according to a gaussian operator with a standard deviation of 1.10^{-2} and a gene’s mutation rate of 5.10^{-3} . Finally, population size was kept constant through the evolution with a number of 20 individuals. All experiments were done using the framework for evolutionary computation SFERESv2 (Mouret and Doncieux, 2010), which includes a fast 2D robotic simulator. The source code for reproducing the experiments is available for download at <http://pages.isir.upmc.fr/~bredeche/Experiments/ECAL2015-coop.tgz>.

Cooperative Foraging Task

In this first experiment, we investigate the evolution of cooperation in a foraging task. The environment is filled with 18 solid targets that the agents can collect. To collect a target, an agent has to stay close to this object for a fixed amount of simulation steps (800). After this duration, the target disappears and any agent close to it is rewarded with its value. Targets are of two types. Green targets always reward 50 when collected whereas purple ones reward 250 only when the agents collect it together (Table 1). If a solitary agent collects a purple target, it disappears and rewards nothing. Consequently, there is both an incentive and a risk to cooperate as cooperation is dependent on successful coordination. This setup is a robotic implementation of a well-known

problem in game theory for studying the evolution of mutualistic cooperation: the *Stag Hunt* (Skyrms, 2004).

The fitness score (F) of an individual is the average reward per trial:

$$F = \frac{1}{N * M} \sum_{i=1}^N \sum_{j=1}^M f_{ij}$$

Where N is the number of individuals encountered (5 in the control and coevolution setups, 1 in the clonal setup), M the number of trials (5) and f_{ij} the rewards obtained at trial j with individual i .

When a target is collected, another target of the same type is then placed at a random position in the arena to keep a constant ratio between green and purple targets. Each evaluation lasted 20000 simulation steps and 60 independent runs were conducted for each experimental setup, each one lasting 40000 evaluations.

| Target | | Reward |
|--------|--------------|--------|
| Green | <i>alone</i> | 50 |
| | <i>coop</i> | 50 |
| Purple | <i>alone</i> | 0 |
| | <i>coop</i> | 250 |

Table 1: Rewards for the foraging of the different targets, depending on whether they were collected alone or cooperatively.

Results

| Setting | # Coop. | # Solitary | Total |
|--------------------|---------|------------|-----------|
| <i>Control</i> | 10 | 50 | 60 |
| <i>Clonal</i> | 28 | 32 | 60 |
| <i>Coevolution</i> | 14 | 46 | 60 |

Table 2: Number of simulations where the best individual evolved a cooperative strategy (collecting purple targets) or a solitary strategy (collecting green targets) for each setup in the foraging task.

We are interested in the number of simulations where cooperation evolved (i.e. the *evolvability* of each approach), which means simulations where the best individual in the population evolved the cooperative foraging of the purple targets (i.e. more than 50% of the collected targets are purple). Results for the three setups are displayed in Table 2. As could be expected from the literature, the clonal setup displays a greater evolvability w.r.t. evolving cooperation (28/60), whereas coevolution (14/60) is on par with the control setup (10/60). It is also apparent that cooperation is still difficult to evolve as in the best case (clonal), no more than

half the simulations display the evolution of cooperative behaviors.




Figure 1: Median fitness score of the best individuals in each of the runs where cooperation evolved for each setup over time. The fitness score of an individual is computed as the average reward the individual earned per trial by foraging targets. The colored areas around the medians represent the first and third quartiles.

However, cooperative individuals do not perform with the same *efficiency* from one setup to another. We show in Figure 1 the median fitness score of the best individuals in each independent run where cooperation evolved over time and for each setup. Fitness scores are significantly different in each setup with the best score obtained in the coevolution setup and the worst in the control setup (Mann-Whitney U-test on the fitness score of the best individuals at the last evaluation, p -value < 0.001).

These differences in efficiency can be explained by looking at the nature of the cooperative behaviors evolved, which reveals two types of behaviors: *turning* and *leader/follower*.

Individuals adopting the turning strategy turn around one another so that they always see the other individual as well as stay close to it (Figure 2(a)). This allows the two individuals to approach simultaneously a same target and therefore forage it in a cooperative fashion. In this strategy, both individuals have a similar behavior and no role division is necessary for their successful cooperation.

In comparison, individuals which evolve a leader/follower strategy adopt a differentiation between two roles: *leader* and *follower* (Figure 2(b)). The individual we call leader always goes first on a target whereas the follower always arrives second on the same target. We observe that the follower's behavior consists in staying close to the leader and always keeping it in front of itself. In comparison the leader

shows a lesser interest in the presence of its follower and rarely checks on its position.

Table 3 shows the distribution of cooperative strategies for all three setups. Whereas the control and clonal setups always resulted in turning strategies (resp. 10/10 and 28/28), the coevolution setup always displayed the evolution of a leader/follower strategy (14). We observe that this latter strategy leads to more efficient cooperation. Indeed, individuals adopting the turning strategy are forced to check constantly on the other individual’s position. Consequently, they cannot be as fast as individuals with a leader/follower strategy where they move to the target in a straight line under the leader’s direction. Moreover, due to the random proximity of the targets, the turning strategy is prone to errors. Namely, they often get to another target than that of their partner whenever two targets are too close to each other.

A possible explanation as to why no leader/follower strategy could evolve in the control and clonal setups may be because of the need to differentiate between the two roles. Indeed, there needs to be the existence of an asymmetry between the two individuals for this phenomenon to appear. With coevolved populations, this asymmetry is deliberately created by the separation between the two populations. Indeed, we observe that one population exclusively contains leaders while the other exclusively contains followers.

The two other setups fail to evolve heterogeneous behaviors. In the control setup, this may be due to the evolutionary algorithm used, especially with elitism enforcing the homogenization of the population throughout the course of evolution (as hinted in the Methods Section). Then, the clonal setup introduces yet another challenge as switching to a particular role can only be done during evaluation as both individuals are by definition genetically similar.

| Setting | # Leader/Follower Strategy | # Turning Strategy | Total |
|--------------------|----------------------------|--------------------|-----------|
| <i>Control</i> | 0 | 10 | 10 |
| <i>Clonal</i> | 0 | 28 | 28 |
| <i>Coevolution</i> | 14 | 0 | 14 |

Table 3: Repartition of the different strategies evolved in each of the runs where cooperation evolved for each setup in the foraging task. We indicate in each cell the number of simulations where a particular strategy evolved.

Going Beyond the Evolvability vs. Efficiency Tradeoff using Incremental Evolution

The previous section revealed a tradeoff between evolvability and efficiency. In the clonal setup, cooperation evolves more often than with other setups. However, the coevolution setup yields cooperative behaviors which are more efficient, with paired individuals displaying asymmetrical behaviors.

In this section, we address the following question: is it possible to benefit from both evolvability *and* efficiency with the clonal and/or the coevolution setups? In other words, we explore (1) whether the clonal setup can be used to evolve pairs with heterogeneous behaviors, and (2) whether the coevolution setup can be improved in terms of number of runs where cooperation evolved.

In order to address this question, we use incremental evolution, a rather common method in evolutionary robotics for solving challenging problems (Dorigo and Colombetti, 1994; Saksida et al., 1997; Bongard, 2008; Doncieux, 2013). The main principle is to ease the learning of a complex task by splitting it into simpler sub-tasks (Perkins and Hayes, 1996).

In the following, we introduce an additional task, the *waypoints crossing* task, which requires the evolution of coordination behaviors, and is simpler to address than the previous task. Individuals evolved in this first task are then used as starting point for the original task described earlier, hoping that cooperative behavior will be *recycled* from the first task to the second task.

Waypoints Crossing Task

We consider a task where robotic agents have to cross randomly positioned waypoints. As such, these round waypoints do not act as obstacles and have a diameter of 30 units. As soon as an agent goes through a waypoint, it can not be seen by this agent anymore. All 18 waypoints have the same color and can be crossed in any order. The fitness score (F) of each individual is defined as the average longest sequence of waypoints shared by both agents per trial:

$$F = \frac{1}{N * M} \sum_{i=1}^N \sum_{j=1}^M l_{max_{ij}}$$

Where N is the number of individuals encountered (5 in the control and coevolution setups, 1 in the clonal setup), M the number of trials (5) and $l_{max_{ij}}$ the longest sequence of waypoints shared by both individuals at trial j with individual i .

This implies that the two individuals are rewarded when crossing waypoints in the same order as well as maximizing the number of waypoints crossed. Each evaluation lasted 10000 simulation steps and 60 independent runs were conducted for each experimental setup, each one lasting 40000 evaluations.

All simulations showed an increase in fitness score for each of the three setups (cf. Figure 3). This was expected as this task does not represent a particular challenge for the individuals: it simply needs the evolution of a successful coordination strategy. However, whereas the coevolution and clonal setups performed equally, they both surpassed the performance of individuals from the control setup (Mann-Whitney, p -value < 0.001).

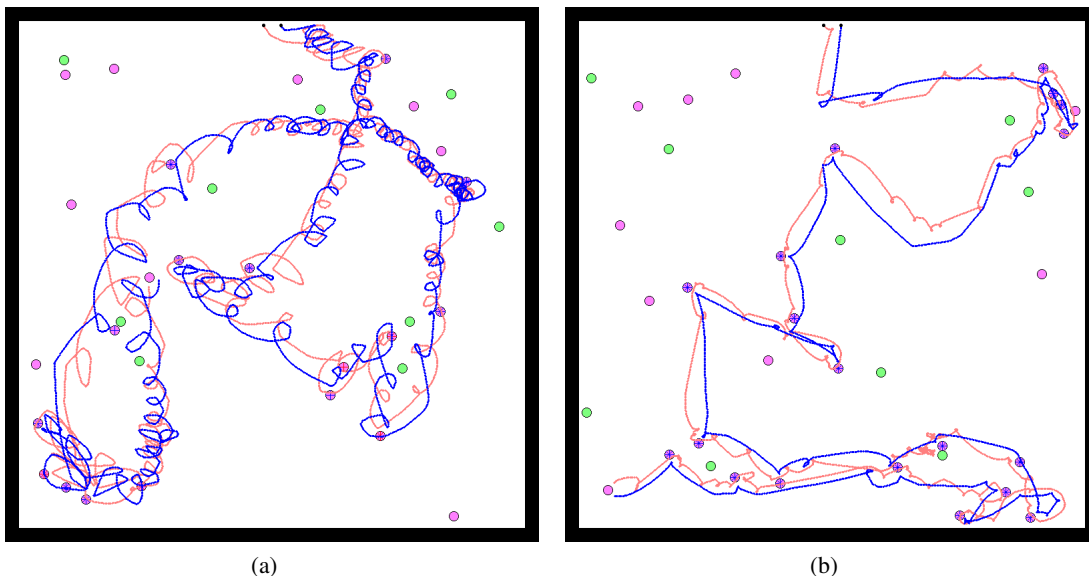


Figure 2: Snapshots of the simulation after an entire trial in the foraging task. The path of each robotic agent from their initial positions (black dots) is represented in red and blue. The green and purple discs represent the 18 targets in the environment. When a target is foraged by the two agents, a red cross (resp. blue) is drawn on the target if the red agent (resp. blue) arrived on it first. Each snapshot corresponds to a trial where agents adopted a different behavior: (a) turning or (b) leader/follower.

As with the previous foraging task, we can hypothesize that these differences in fitness scores are due to differences in the behaviors evolved. Table 4 gives a classification of the cooperative behaviors for each setup. They are similar to those in the previous task with the addition of a third rare strategy: the *wall-following* strategy (which is regrouped in “Other”). Wall-followers simply follow the walls around the arena and cross any waypoints close to the wall they are adjacent to. As such, this is a far less efficient strategy than the two others.

| Setting | # Lead. | # Turn. | # Other | Total |
|-------------|---------|---------|---------|-----------|
| Control | 19 | 37 | 4 | 60 |
| Clonal | 23 | 31 | 6 | 60 |
| Coevolution | 59 | 1 | 0 | 60 |

Table 4: Repartition of the different strategies evolved in each of the 60 independent runs for each setup in the waypoints task. We indicate in each cell the number of simulations where a particular strategy evolved: *Leader/follower* (Lead.), *Turning* (Turn.) or *Other*. “Other” regroups wall-following strategies or simulations where no recognizable strategy evolved.

In the coevolution setup, nearly all runs (59/60) evolved a leader/follower strategy. Interestingly, although fitness scores in the clonal and control setups are significantly different, this behavior evolved in roughly one third of the runs for both setups. To explain the difference in fitness scores, we must take into account the quality of the leader/follower

strategy in each setup. We measure the proportion of leadership in each interaction, which is computed as the proportion of waypoints crossed by both individuals for which the leader arrived first. Figure 4 displays the boxplots of the proportion of leadership for the best individuals in each setup and only for the simulations where a successful leader/follower strategy evolved (a minimal threshold of 0.75 is chosen to consider only the best performing runs). We show that the proportion of leadership is greater in the clonal and coevolution setups than in the control one (Mann-Whitney, p -value < 0.005). These differences means that the individuals are more efficient in their leader/follower strategy in the clonal and coevolution settings than in the control one. This explains the differences in fitness scores observed in Figure 3.

Interestingly, whereas in the foraging task no leader/follower strategy could evolve in the control and clonal setups, this strategy did evolve in one third of the simulations for this task. This could mean that these individuals use information in the environment to adopt one role or the other. Indeed, we observe that this is achieved by exploiting the differences in the initial starting positions, with one individual on the left and the other on the right. They both turn to the same direction (left or right, depending on the runs) at the beginning of the simulation which results in one individual (the leader) turning its back to the other, while the second individual (the follower) looking at its partner.

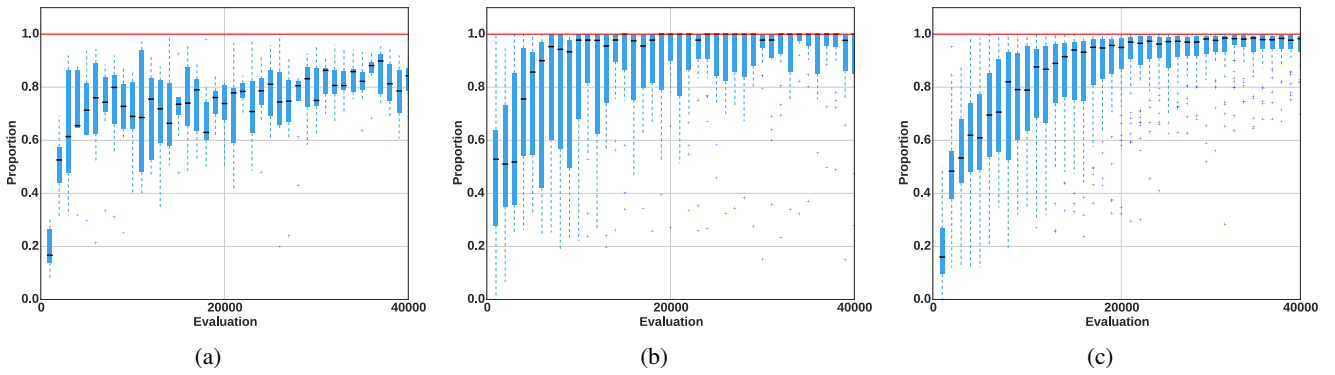


Figure 4: Boxplots of the proportion of leadership over time for the best individuals in each runs where the proportion at the last evaluation was greater than 0.75 in the (a) control, (b) clonal or (c) coevolution setup. This value represents the proportion of waypoints crossed by both individuals for which the leader arrived first.




Figure 3: Median fitness score of the best individuals in each of the 60 independent runs and for each setup over time. Fitness score is computed as the average longest sequence of waypoints shared by both agents per trial. The colored areas around the medians represent the first and third quartiles.

Recycling Cooperative Behaviors in the Foraging Task

Coming back to the initial foraging task, we perform the exact same experiment described at the beginning of this paper, with one notable exception: the initial population is initialized with genomes evolved for solving the waypoint task. This implies that coordination is possible starting from the very first generation of each setup. Given that we have already shown that such coordination is a desirable feature, the question is: will it be possible to retain cooperative behaviors in order to solve the foraging task?

Table 5 gives the results in terms of evolved behaviors

| Setting | # Coop. | | # Solitary | Total |
|--------------------|---------|---------|------------|-----------|
| | # Lead. | # Turn. | | |
| <i>Control</i> | 0 | 20 | 40 | 60 |
| <i>Clonal</i> | 0 | 24 | 36 | 60 |
| <i>Coevolution</i> | 28 | 0 | 32 | 60 |

Table 5: Proportion of the 60 independent simulations where the best individual evolved a cooperative strategy (collecting purple targets) or a solitary strategy (collecting green targets) for each setup in the foraging task when individuals are previously evolved in the waypoints task. In addition, the repartition of the different strategies is indicated when cooperation evolved: *Leader/Follower* (Lead.) or *Turning* (Turn.).

from the 60 independent runs for each setup. The coevolution setup evolves cooperation slightly more often (28/60) than both the control (20/60) and the clonal (24/60) setups. A first remark is that the number of occurrences of cooperation for the coevolution and control setups have actually doubled compared to previous results without incremental evolution (see Table 2). This is not the case for the clonal setup, which does not appear to benefit from incremental evolution.

A second remark is that cooperation in the coevolution setup systematically corresponds to a leader/follower strategy, which is *never* the case with the two other setups. This has a significant, though expected, impact on fitness scores, as shown in Figure 5. Cooperation evolved with the coevolution setup leads to significantly greater fitness scores (Mann-Whitney, p -value < 0.001).

Results from this experiment make it possible to revise our initial statement. Using pre-trained individuals strongly benefits the coevolution setup in terms of evolvability. But this is not the case with the clonal setup, for which using pre-trained individuals improves neither evolvability nor ef-

efficiency. Therefore, we may face a tradeoff which does not concern evolvability and efficiency, but one that implies computational cost: the coevolution setup outperforms the clonal setup on both evolvability and efficiency *at the cost of additional computational effort*.

The control and clonal setups completely failed to maintain a leader/follower strategy, even though such strategy originally evolved. An explanation is provided by considering the difference between the waypoints task, where leader/follower evolved, and the current foraging task. In the waypoints task, symmetry breaking could be achieved at the beginning of the evaluation (as explained earlier), and could be retained afterwards as the follower was always behind the leader. However, the current foraging setup requires that the two robots display the same behavior to cooperatively collect a target (ie. both robots have to touch the target), which implies that leader/follower roles are lost, as they depend on the relative position of robots with one another.




Figure 5: Median fitness score of the best individuals in each of the runs where cooperation evolved for each setup over time. The fitness score of an individual is computed as the average reward the individual earned per trial by foraging targets. The colored areas around the medians represent the first and third quartiles.

Discussion and Conclusion

In this paper, we considered several approaches for the evolution of cooperation in evolutionary robotics: a clonal approach, where all individuals in a group share the same genome, and a non-clonal approach, where individuals are independent from one another, but may share a common interest in cooperating.

We first showed that there exists a tradeoff between evolvability and efficiency. On the one hand, the clonal approach

evolves cooperative behaviors on a more frequent basis than with the other approach. On the other hand, the non-clonal approach, which is implemented using a coevolution setup, results in more efficient behaviors in terms of pure performance whenever cooperation evolved. The non-clonal approach actually enables the evolution of asymmetric behaviors, such as a leader/follower strategy.

We then used incremental evolution to evolve coordination behaviors using a simpler task in order to overcome this tradeoff and improve both evolvability and efficiency in each setup. We showed that while no improvement was observed in the clonal setup on either criteria, the outcome is very different for the coevolution setup: the probability of evolving cooperation actually increases, and the evolved cooperative solutions remain the most efficient.

This work raises several questions. Firstly, heterogeneous behaviors were obtained with coevolution, a rather radical way to enable asymmetrical behaviors during cooperation. However, the waypoints task revealed that breaking symmetry can also be done with identical individuals using environmental feedback, even though such cooperation is difficult to obtain. As a consequence, we intend to investigate the evolution of cooperation with heterogeneous behavior without resorting to coevolution. In particular, we will study how more elaborated neural architectures (e.g. using plasticity) can switch to a particular persistent regime depending on environmental cues available at the beginning of the evaluation.

Secondly, incremental evolution requires an added computational cost in order to increase evolvability in the non-clonal approach. However, it may be possible to avoid this extra cost by considering other evolutionary methods. In particular, we intend to explore how a multiobjective approach which considers both performance and *diversity* could improve the optimization process (Lehman and Stanley, 2008; Doncieux and Mouret, 2014). Though this approach looks promising, it is not clear yet how diversity should be implemented in the context of cooperative problem solving.

Acknowledgements

This work is supported by the ANR CreAdapt project (ANR-12-JS03-0009) and the European Union’s Horizon 2020 research and innovation programme under grand agreement No 640891. Experiments presented in this paper were carried out using the Grid’5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>).

References

Bongard, J. (2008). Behavior Chaining : Incremental Behavior Integration for Evolutionary Robotics. In *Artificial Life XI*:

- Proceedings of the 11th International Conference on the Simulation and Synthesis of Living Systems*, pages 64–71. MIT Press.
- Connor, R. C. (1995). The Benefits of Mutualism: A Conceptual Framework. *Biological Reviews*, 70(3):427–457.
- Doncieux, S. (2013). Transfer Learning for Direct Policy Search: A Reward Shaping Approach. In *Proceedings of the International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 1–6. IEEE.
- Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. E. (2015). Evolutionary Robotics: What, Why, and Where to. *Frontiers in Robotics and AI*, 2(4).
- Doncieux, S. and Mouret, J.-B. (2014). Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, pages 1–18.
- Dorigo, M. and Colombetti, M. (1994). Robot Shaping : Developing Autonomous Agents through Learning. *Artificial Intelligence*, 70(2):321–370.
- Floreano, D. and Nolfi, S. (1997). God Save the Red Queen! Competition in Co-Evolutionary Robotics. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 398–406.
- Floreano, D., Nolfi, S., and Mondada, F. (1998). Competitive Co-Evolutionary Robotics: From Theory to Practice. In *From Animals to Animats 5: Proceedings of the 5th International Conference on Simulation of Adaptive Behavior*, pages 515–524.
- Harvey, I., Husbands, P., and Cliff, D. (1994). Seeing The Light : Artificial Evolution, Real Vision. In *From Animals to Animats 3: Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior*, volume 1994, pages 392–401.
- Hauert, S., Mitri, S., Keller, L., and Floreano, D. (2014). Evolving Cooperation : From Biology to Engineering. In *The Horizons of Evolutionary Robotics*, pages 203–217. MIT Press.
- Lehman, J. and Stanley, K. O. (2008). Exploiting Open-Endedness to Solve Problems Through the Search for Novelty. *Artificial Life XI*, pages 329–336.
- Lichocki, P., Wischmann, S., Keller, L., and Floreano, D. (2013). Evolving team compositions by agent swapping. *IEEE Transactions on Evolutionary Computation*, 17(2):282–298.
- Montanier, J.-M. and Bredeche, N. (2011). Surviving the Tragedy of Commons : Emergence of Altruism in a Population of Evolving Autonomous Agents. In *Proceedings of the 11th International Conference on Artificial Life (ECAL'11)*, pages 550–557.
- Mouret, J.-B. and Doncieux, S. (2010). SFERESv2: Evolving in the Multi-Core World. In *Proceedings of Congress on Evolutionary Computation (CEC)*, pages 4079–4086. IEEE.
- Panait, L. and Luke, S. (2005). Cooperative Multi-Agent Learning: The State of the Art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434.
- Perkins, S. and Hayes, G. (1996). Robot Shaping - Principles, Methods and Architectures. Technical report, University of Edinburgh, Edinburgh, GB.
- Saksida, L. M., Raymond, S. M., and Touretzky, D. S. (1997). Shaping Robot Behavior using Principles from Instrumental Conditioning. *Robotics and Autonomous Systems*, 22:231–249.
- Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.
- Trianni, V., Ampatzis, C., Christensen, A., Tuci, E., Dorigo, M., and Nolfi, S. (2007). From Solitary to Collective Behaviours: Decision Making and Cooperation. In Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., and Coutinho, A., editors, *Advances in Artificial Life*, volume 4648 of *Lecture Notes in Computer Science*, pages 575–584. Springer Berlin Heidelberg.
- Urzelai, J., Floreano, D., Dorigo, M., and Colombetti, M. (1998). Incremental Robot Shaping. *Connection Science*, 10(3-4):341–360.
- Waibel, M., Floreano, D., and Keller, L. (2011). A Quantitative Test of Hamilton’s Rule for the Evolution of Altruism. *PLoS biology*, 9(5):e1000615.
- Waibel, M., Keller, L., Floreano, D., and Member, S. (2009). Genetic Team Composition and Level of Selection in the Evolution of Cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660.