Evolution of Altruistic Robots

Dario Floreano[1], Sara Mitri[1], Andres Perez-Uribe[2], Laurent Keller[3]

[1]Laboratory of Intelligent Systems, EPFL, Lausanne, Switzerland
[2]University of Applied Sciences, Yverdon, Switzerland
[3]Department of Ecology and Evolution, University of Lausanne, Switzerland

Abstract. In this document we examine the evolutionary methods that may lead to the emergence of altruistic cooperation in robot collectives. We present four evolutionary algorithms that derive from biological theories on the evolution of altruism in nature and compare them systematically in two experimental scenarios where altruistic cooperation can lead to a performance increment. We discuss the relative merits and drawbacks of the four methods and provide recommendations for the choice of the most suitable method for evolving altruistic robots.

1 Altruistic Cooperation in Nature

The competition for survival and reproduction postulated by Darwin seems at odds with the observation that some organisms display cooperative behaviors. In order to understand the evolutionary conditions when cooperation can emerge, Lehmann and Keller [14] suggested to distinguish between two types of cooperation (figure 1), namely the situations where a cooperator does not pay a fitness cost from helping other individuals and the situations where a cooperator must pay a fitness cost for helping other individuals. Let us remember that in biology fitness benefits and costs translate into the number of genetic copies that an individual can produce or loose with respect to its baseline reproduction rate.

The situation where cooperation generates a fitness benefit without any cost to the cooperator is relatively common in nature. This situation can be further divided in two cases, when the benefit is immediate or direct and when the benefit is indirect. Examples of cooperation with direct benefits include nest building and group hunting. Whenever a cooperator obtains an immediate and direct benefit from helping another individual, cooperation will always evolve and remain stable, no matter whether the receiving individuals belong to another species or have never been seen before.

If the benefit is indirect, i.e., the act of helping is not immediately reciprocated or the benefit appears only in the long term, cooperation evolves only if individuals have an initial tendency to cooperate, interact together several times, and can both recognize the partner and remember the outcome of previous interactions. If these conditions are satisfied, cooperation will always evolve and remain stable even if cooperating individual belong to different species.

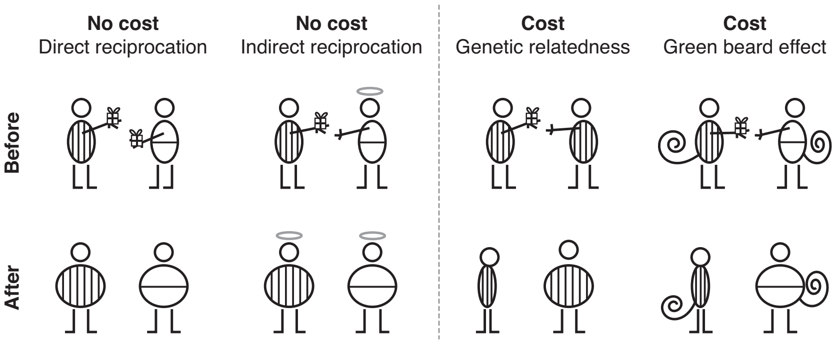

Fig. 1. Conditions for the evolution of cooperation according to the classification suggested by Lehmann and Keller [14]. When there is no cost for the cooperator, cooperation can evolve if there is direct reciprocation or indirect reciprocation (in the latter case, a reputation system may help). When there is a cost for the cooperator, cooperation can evolve if individuals have a high level of genetic relatedness or if they both have greenbeard genes. The pattern indicates the genetic similarity between individuals. The size change after cooperation indicates the cost or benefit of cooperation. Figure loosely inspired by figure 5.9 in [29].

It has also been shown that recognition of other individuals and memorization of the outcomes of the interactions is not necessary if there is a reputation system that informs how cooperative an individual is [21]. The way in which animals and people decide to cooperate has been studied extensively in game theory, notably within the framework of the Prisoner's Dilemma game.

On the other hand, the situation where cooperation implies a fitness cost for the cooperator is less common. Cooperation with a cost is also known as *altruism* because the cooperator helps other individuals at its own expense. Parental care is an instance of altruism directed towards offspring of the individual because it implies an energetic cost for the parent. The specialization of ant colonies into large numbers of sterile workers (for food collection, nest defense, rearing of the pupae of the queen, etc.) is yet another instance of altruistic cooperation where the helping workers incur the highest fitness cost because they cannot reproduce.

Building on earlier intuitions by Haldane [10], Hamilton [11] suggested that altruism can evolve if the cooperator is genetically related to the recipient of help. In this case, even if the cooperator cannot propagate its own genes to the next generation, its altruistic act will increase the probability that a large portion of those genes will be propagated through the reproduction of the recipient of the altruistic act. Hamilton [11] proposed the notion of *inclusive fitness*, which is the sum of the individual fitness and of the fitness effects caused by its own act on the portion of genes shared with other individuals. The portion of shared genes between two individuals is known as *genetic relatedness*. He [11] predicted that altruistic cooperation will evolve if the inclusive fitness of the helper is larger than zero

$$rb - c > 0 \tag{1}$$

where r is the coefficient of genetic relatedness, b is the fitness benefit of the recipient(s) of help, and c is the fitness cost of the helper. To use an example suggested by Haldane, in the case of brothers, where $r = 1/2$, an individual may be willing to sacrifice its own life and thus pay the maximum cost $c = 1$ if its act increases more than twice $b > 2$ the fitness of the brother. For cousins, where $r = 1/8$, an individual may be willing to pay the maximum cost if its act increases the fitness of the cousin more than eight times.

Hamilton's inequality applies to average genetic relatedness over the entire genotype and population, i.e. it is not restricted to the sharing of a specific set of genes. It also applies to the case where the act of cooperation benefits multiple individuals with various degrees of relatedness. The theory of *kin selection* [16], which developed from Hamilton's model, predicts that the ratio of altruistic individuals in a population is related to the degree of kinship, or genetic relatedness, among individuals. Although the theory is widely accepted, its quantitative validation in nature has not yet been done because it is difficult to precisely measure the values of the three variables in equation 1.

For evolution of altruism to occur, helping should be directed towards related individuals. This is more likely to happen when individuals share the same geographical space, such as a nest, for social activities. Indeed, most cases of altruistic cooperation are found in families of social insects [12]. Kin selection does not require that individuals recognize kin individuals or know their degree of genetic relatedness. As long as the act of altruism preferentially benefits genetically-related individuals, altruism will spread throughout the population and remain stable.

A particular case of altruism occurs when individuals share few specific genes that favor cooperating behaviors only between individuals having a specific phenotypic character, such as a green beard [7], and that express the same phenotypic character. However, altruism due to *greenbeard effects* can be disrupted if the linkage between the genes responsible for the green beard and the genes responsible for altruistic behavior is disrupted. For example, a mutant individual with a green beard but without the altruistic behavior will have larger inclusive fitness than individuals who have both types of genes; consequently, it will spread in the population and destroy altruistic cooperation [14].

The four conditions for the evolution of cooperation, direct or indirect reciprocity, genetic relatedness and greenbeard genes, which can all be included within a single model [14], hold only if cooperation brings a net fitness advantage to the individuals. In some societies, the actual values of benefits and costs are distorted by means of coercion and punishment to ensure maintenance of cooperative behavior.

Yet another explanation for the evolution of altruistic cooperation is provided by the theory of *levels of selection*, which argues that altruistic cooperation may also evolve in colonies of genetically unrelated individuals that are selected and reproduced all together at a higher rate than the single individuals composing the colony [31]. This could happen in situations where the synergetic effect of cooperation by different individuals provides a higher fitness to the group with respect to other competing groups.

However, the colony-level selection has been criticized because genetic mutations at the level of the individual are more likely and frequent than mutations at the level of the colony, thus creating stronger competition among individuals than among colonies. It has also been argued that the transition from uni-cellular to multi-cellular organisms can be explained by kin selection because all cells share the same genotype [30]. Although proponents of colony-level selection respond to these criticisms by pointing to evidence for the evolution of colony-level features that decrease individual conflict (such as a reduced mutation rate of individual organisms or cells that compose the colony), the theory of colony-level selection is still widely debated. Furthermore, colony-level selection may eventually lead to high genetic relatedness, thus making the disambiguation between the original driving forces that led to altruistic cooperation even more difficult.

2 Artificial Evolution of Cooperation

In robotics, the evolution of collective behaviors has been studied in several experiments, but often without attention to whether it involves only behavior coordination or also cooperation and whether cooperation involves a cost for the individuals. In those situations where cooperation is explicitly mentioned, it is described as a situation where robots obtain an advantage by working together rather than working in isolation.

When it comes to evolving teams of robots, the experimenter is presented with two design choices: 1) whether robots should be genetically identical or different; and 2) whether the fitness used for selection should take into account the performance of the entire group or only that of single individuals. These two choices are analogous to the issues of genetic relatedness and of level of selection that were discussed above in the context of the biological literature. If we consider only the extreme cases of each design choice, robots in a team can be genetically homogeneous (clones) or heterogeneous (they differ from each other); and the fitness can be computed at the level of the team (in which case, the entire team of individuals is reproduced) or at the level of the individual (in which case, only individuals of the team are selected for reproduction).

Biological theory tells us that the evolution of genetically related robots should lead to cooperative behaviors, but the question of the appropriate level of selection, or fitness computation, is still open for discussion. Furthermore, biological theory does not make any prediction on the comparative performances that we may expect from robots evolved under different conditions.

The majority of current approaches to the evolution of multi-agent systems use genetically homogeneous teams evolved with team-level selection (a comparative survey can be found in [27]). Where the reasons for the choice of genetically homogeneous teams are made explicit, it is

Fig. 2. A swarm-bot composed of four interconnected s-bots in chain formation.

argued that homogeneous teams are easy to use [3, 26], require fewer evaluations [15, 25], scale more easily [6], and are more robust against the failure of team members [6, 24] than heterogeneous teams.

The choice of level of selection is rarely discussed explicitly despite the fact that fitness distribution leads to credit assignment problems [9, 19] in many cooperative multi-agent tasks because individual contributions to team performance are often difficult to estimate or difficult to monitor [23].

Let us consider the case of evolving control systems for a population of identical robots, the s-bots shown in figure 2, which can self-connect to form a swarm-bot [20]. In a simple case, a swarm-bot of four s-bots assembled in chain formation were evolved for the ability to move coordinately on a flat terrain. Each s-bot was provided with a neural controller where sensory neurons were directly connected to the motors neurons that controlled the desired speed of the tracks. The sensory neurons received information from distance sensors around the body of the robot and from a torque sensor that measured the amount of torsional force exerted by other robots. In this case, all s-bots in the swarm-bot were genetically identical and the fitness measured the progress of the entire swarm-bot on the ground. Evolved controllers were also capable of producing coordinated movement also when the swarm-bot was augmented by additional s-bots and re-organized in different shapes. Swarm-bots also dynamically rearranged their shape so as to effectively negotiate narrow passages and were capable of moving on rough terrains over holes or slopes that could not be passed by a single robot. Such robots also collectively avoided obstacles and coordinated to transport heavy objects [1, 2, 26].

The choice of team-level selection in this case was imposed by the difficulty to assign fitness values to individual s-bots that composed the swarm-bot. However, the choice of genetically related teams was not

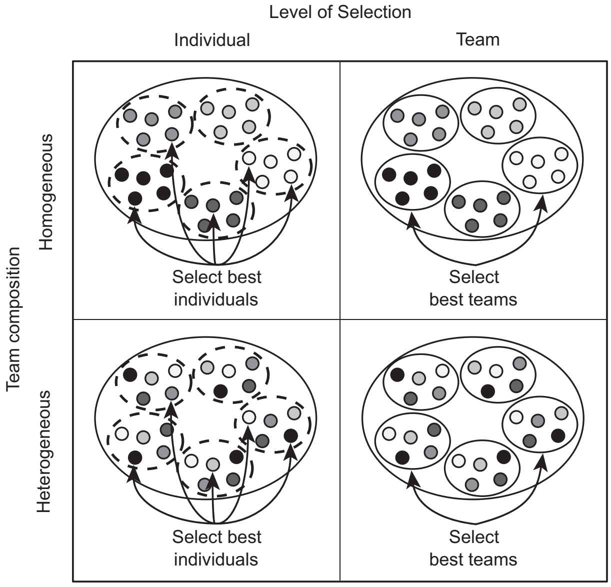

Fig. 3. Four conditions for the evolution of robot collectives. A population (large oval) is composed of several teams (medium ovals), each of which is composed of several robots (small circles). Genetic team composition is varied by either composing teams of robots with identical genomes (homogeneous, identical shading), or different genomes (heterogeneous, different shading). The level of selection is varied by either measuring team performance and selecting teams (team-level selection) or measuring individual performance and selecting individuals independently of their team affiliation (individual-level selection).

duly justified because it may have prevented the emergence of specialized individuals.

The question therefore remains of what is the best performing set of choices for tasks that benefit from cooperative behaviors when there is both a choice between genetic relatedness and level of selection. In the remainder of this chapter, we will describe the systematic comparison of these design choices for two sets of experiments that can benefit from the evolution of altruistic cooperation.

2.1 Evolutionary Conditions

We compared four evolutionary conditions (figure 3): genetically homogeneous teams evolved with team-level selection; genetically homogeneous teams evolved with individual-level selection; genetically heterogeneous teams evolved with team-level selection; and genetically heterogeneous teams evolved with individual-level selection. Team-level selection (akin to colony-level selection) consisted of computing the fitness of the team and reproducing the robots in the best teams to create a new population

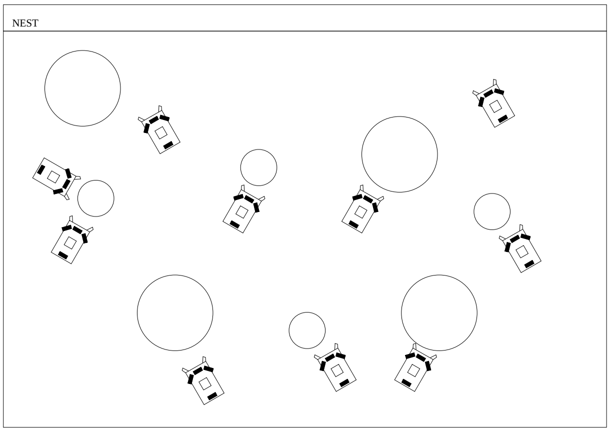

Fig. 4. A team of artificial ants is foraging for food tokens. Small food tokens can be transported by a single ant and are consumed by that ant when it manages to get to the nest. Large food tokens require the cooperation of two ants to be transported to the nest, but they are shared by the entire team. However, the share of a large food token provides less food intake to each individual than a small token. For the sake of simplicity, in this figure we are only showing 10 artificial ants.

of robot teams. Individual-level selection instead consisted of computing the fitness of individual robots (notice that even robots with identical genomes can obtain different fitness because they are exposed to different situations) and reproducing the best ones independently of their team affiliation to recreate new teams.
The comparisons were carried out in situations where both selfish and altruistic behaviors could produce fitness increments over generations, but altruistic behavior corresponded to larger fitness increments, that is to a larger quantity of work accomplished by the team of robots. In a first set of experiments, we resorted to simplified behaviors and simulated environments in order to disentangle fitness differences due to the effects of the evolvability of control systems in situated environments from the effects of the four evolutionary conditions. In a second set of experiments, we resorted to neural controllers in real and simulated robots.

2.2 Altruistic Foraging

In the first set of experiments, we used an agent-based model of a team of artificial ants performing a foraging task (figure 4). The agents or artificial ants (e.g., robots) are supposed to look for food items randomly scattered in a foraging area. There are two kinds of food items, small food items which can be transported by single agents to the nest, and large food items, which can only be transported if two ants cooperate. When a cooperative foraging ant happens to find a large food item, it sends a local message asking for help. Given the local nature of the help message, another cooperative individual will only be able to help the first one if it happens to be close to it and hear its message. For sake of simplicity large food items can only be transported by a pair of ants and we have not included a pheromone-like communication among ants.

Each ant is endowed with a set of three genes encoding three threshold values that are used to determine if one or more predefined behaviors (b_0, b_1 or b_2) are activated at each step of a foraging trial, as shown in the table.

b_0	b_1	b_2	**Behavioral strategies**
0	0	0	do nothing
1	0	0	if a small food item is found, bring it to the nest, ignore large food items, and do not help other ants
0	1	0	if a large food item is found, stay and ask for help, ignore small food items, and do not help other ants
0	0	1	if a help message is perceived, go and help, ignore small and large food items
1	1	0	if a small food item is found, bring it to the nest, if a large food item is found ask for help, but do not help other ants
1	0	1	if a small food item is found, bring it to the nest, help other ants, but ignore large food items
0	1	1	if a large food item is found, stay and ask for help, ignore small food items, and help other ants
1	1	1	if a small food item is found, bring it to the nest, if a large food item is found, stay and ask for help, and help other ants

The expression of a given behavior b_i depends on the number of foragers already engaged in that behavior and is mediated by the thresholds values that are genetically encoded, as suggested by the response threshold value of [4]. For example, if the proportion of members of the team having activated a given behavior j is smaller than the corresponding threshold of ant k, behavior b_j^k is set to '1' (i.e., it is activated).

The agents were not physically simulated; the model assumed a random walk and took into account the probability of finding a food token at each time step, which decreased in proportion to the number of token collected by the agents. The model also included a probabilistic function of perception and action.

We used 20 agents foraging for 4 large food tokens and 4 small food tokens. The performance of the robot teams was measured using the average score obtained during 20 foraging trials. The small food items provided a score of 1.0 to the single ant who transported it to the nest, while the large food items provided a total score of 16.0. However, since the large food items were shared with the whole team, each individual obtained a score of 0.8 for any large food item taken to the nest.

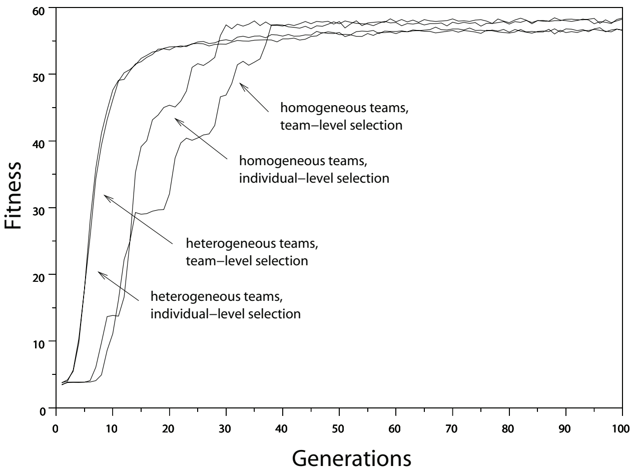

Fig. 5. Evolution of the mean performance of homogeneous and heterogeneous colonies under individual and team-level selection (each curve is the average over 10 different evolutionary runs of mean population fitness).

According to these payoffs, all individuals, including those that do not cooperate, can get 0.8 points for every large food item transported by other individuals of the team, whereas the individuals that cooperate in foraging for large food items, pay a cost of 0.2 points compared to the score 1.0 that they would made if they foraged on small food items. The total performance of the team, or total energy brought to the nest, was highest when individuals were altruist rather than selfish.

Performance differences appeared to be cause mainly by genetic relatedness (figure 5). Homogeneous colonies displayed significantly higher mean fitness than heterogeneous colonies. The difference between homogeneous and heterogeneous fitness depends on the relative cost and benefit ratios, as postulated by Hamilton's inequality. However, there was no significant difference between the mean performance of homogeneous colonies evolved using team-level selection and mean performance of homogeneous colonies evolved using individual-level selection.

The use of pre-defined behaviors allowed us to precisely measure the amount of altruistic individuals in the evolving teams in each of the four evolutionary conditions (figure 6). We considered an individual to be "altruistic" when it expressed behaviors that did not "pay attention" to small food items and concentrated only on large food items, either by searching for large food items or by helping other individuals to transport large food items (see table above).

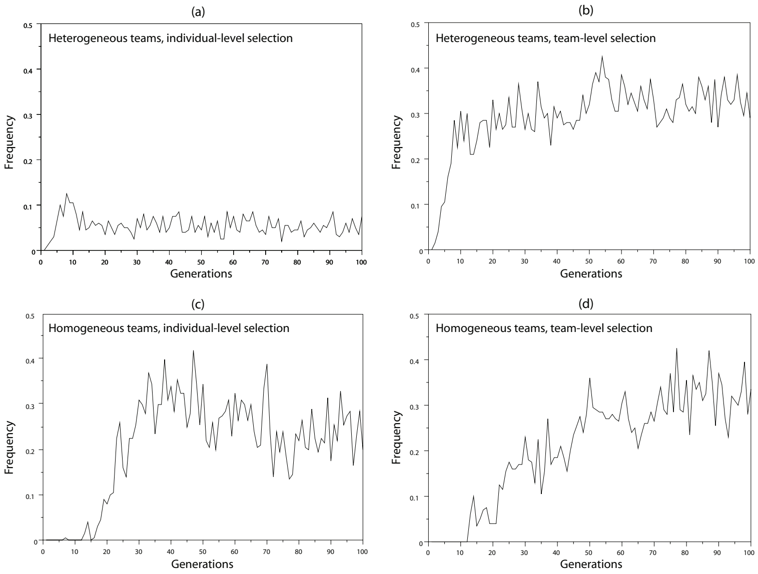

Fig. 6. Evolution of the frequency of altruistic individuals in the simulated ant populations (average of 10 runs) given the following experimental setups: (a) Heterogeneous teams, individual-level selection, (b) Heterogeneous teams, team-level selection, (c) Homogeneous teams, individual-level selection, and (d) Homogeneous teams, team-level selection.

As expected, the frequency of altruistic individuals within populations of heterogeneous teams evolved using individual-level selection remained below 10%. However, in all other three conditions we observed a gradual dominance of altruistic individuals in the population. In particular, the resulting number of altruistic individuals is higher when using team-level selection (Figure 6b and Figure 6d). This is understandable because team-level selection favors the individuals that work for the team and not the ones that specialize in the foraging of small food items for their own benefit.

This set of experiments indicated that homogeneous teams were conducive to higher performances in a scenario that could benefit from altruistic behavior and that team-level selection tended to produce more altruistic individuals than individual-level selection. Therefore, it came with no surprise that teams of heterogeneous individuals evolved with individual-level selection produced very few altruistic individuals and obtained lower fitness. The question however remained of why heterogeneous teams evolved with team-level selection produced a majority of altruistic agents, but did not result in better fitness than heterogeneous teams evolved with individual-level selection.

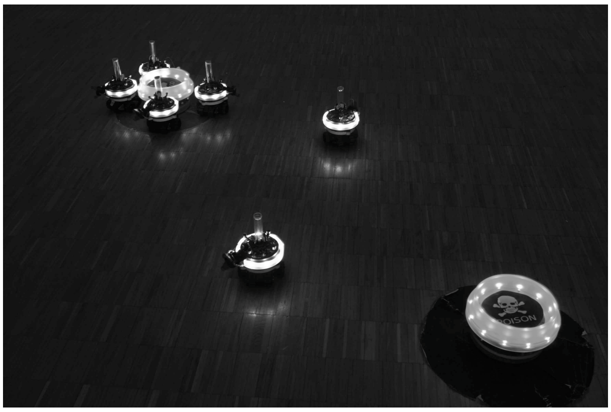

Fig. 7. A team of s-bots engaged in cooperative communication. A team of four s-bots feed on the food objects while they are lit up in blue color. Two s-bots in white color are attracted by the blue signal and move away from the poison object.

We will get back to this issue in the next set of experiments where we repeated our comparison of the four evolutionary conditions in a more realistic scenario both with physics-based robot simulations and with real robots.

2.3 Altruistic Communication

The evolution of communication is a particularly challenging problem both in biological and in robotic systems because efficient communication requires tight co-evolution between the signal emitted and the response elicited [17]. Furthermore, most communication systems are also costly because of the energy required for signal production [32] and/or increased competition for resources resulting from the transmitted information. For example, if organisms decide to communicate the location of a limited food source, individuals may pay a cost due to decreased food intake. In these situations, communication is another example of altruism and its evolvability and efficiency may depend on the four evolutionary conditions mentioned above.

We therefore set up an experimental scenario for comparing the four evolutionary conditions where communication provides both benefits and costs [8]. We used teams of 10 s-bots that could forage in an environment containing a food and a poison source that both emitted red light (figure 7). Under such circumstances, foraging efficiency could potentially be increased if robots transmitted information on food and poison location. However, such communication also incurred direct costs to the signaler because it resulted in higher robot density and increased competition and interference nearby the food (i.e., spatial constraints around the food source allowed a maximum of 8 robots out of 10 to feed simultaneously and resulted in robots sometimes pushing each other away from the food). Thus, while beneficial to other team members, signaling of a food location effectively constituted a costly act because it decreased the food intake of signaling robots. This setting thus mimics the natural situation where communicating almost invariably incurs costs in terms of signal production or increased competition for resources.

The experiments were conducted multiple times using a physics-based simulator which accurately models the dynamical properties of the s-bots. The results were then verified by running a single evolutionary experiment for each of the four conditions with the physical robots. The robots had a translucent ring around the body that could emit blue light and a 360° vision system that could detect the amount and intensity of red and blue light. A circular piece of gray paper was placed under the food source and a similar black paper under the poison source. These paper circles could be detected by infrared ground sensors located between the tracks underneath the robot and thus allowed discrimination of food and poison.

The robots were equipped with a neural network to process the visual information and ground sensor input in order to set the direction and speed of the two tracks and control the emission of blue light accordingly every $50ms$ cycle. During each cycle, a robot gained one performance unit if it detected food with its ground sensors and lost one performance unit if it detected poison. The performance of each robot at the end of a trial was computed as the sum of performance units obtained during that trial (1200 sensory motor cycles of $50ms$) and the robot performance was quantified as the sum of performance units over all 10 trials. Team performance was equal to the average performance of all robots in the team.

The feed-forward neural controller had 10 input and 3 output neurons (figure 8). Once a robot had detected the food or poison source, the corresponding neuron was set to 1. This value decayed to 0 by a factor of 0.95 every $50ms$, thereby providing a short-term memory even after the robot's sensors were no longer in contact with the gray and black paper circles placed below the food and poison. The remaining 8 neurons were used to encode the 360° visual input image, which was divided into four sections of 90° each. For each section, the average of the blue and red channels was calculated and normalized within the range of 0 and 1, such that one neural input was used for the blue and one for the red value. The activation of each of the output neurons was computed as the sum of all inputs multiplied by the weight of the connection and passed through the continuous $\tanh(x)$ function (i.e., their output was between -1 and 1). Two of the three output neurons were used to control the two tracks, where the output value of each neuron gave the direction of rotation (forward if > 0 and backward if < 0) and velocity (the absolute value) of one of the two tracks. The third output neuron determined whether to emit blue light, which was the case if the output was greater than 0. The genotype of an individual encoded the synaptic weights of the neural network in a bit string. Each synaptic weight was encoded in 8 bits, giving 256 values that were mapped onto the interval $[-1, 1]$.

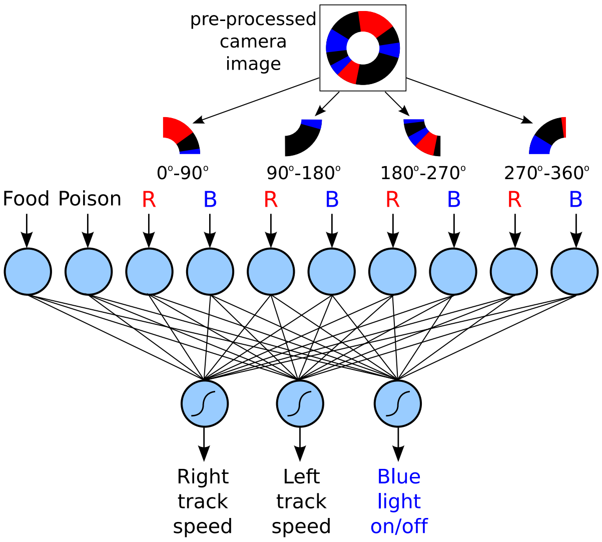

Fig. 8. The neural network architecture used in the experiments on communication.

For each of the four conditions, we ran 20 independent evolutionary experiments with 100 colonies of 10 robots. Furthermore, as a control situation, we repeated all experiments (4 times 20 runs) by disabling the light ring of the robots, but the neural architecture and genotype were the same as in the normal condition.

To compare team performance between treatments, we calculated the average performance of the 100 colonies over the last 50 generations for each of the 20 experiments per condition (figure 9). In evolving teams where robots could produce blue light, foraging efficiency greatly increased over generations and was significantly greater compared to control experiments for all evolutionary conditions, except for the condition of heterogeneous teams under individual-level selection. An analysis of the robot behavior revealed that this performance increment in the three conditions of genetic relatedness or team-level selection was associated with the evolution of effective systems of communication [8].

In teams of genetically related robots with team-level selection, two distinct communication strategies evolved. In 12 of the 20 evolutionary experiments, robots preferentially produced light in the vicinity of the food and were attracted by blue light (figure 10, left). Instead, in the other 8 evolutionary experiments, robots tended to emit light near the poison and were repulsed by blue light (figure 10, right). Teams of robots that signaled food resulted in higher team performance. Interestingly, once one type of communication was well established, there was no transition to the alternate and more efficient strategy. This was because a change in either the signaling or response strategy would completely destroy

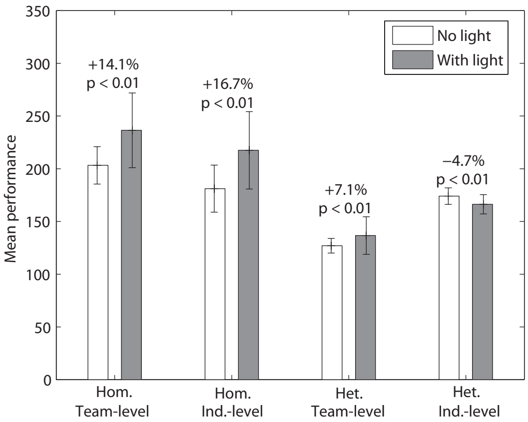

Fig. 9. Mean (+ S.D.) performance of robots during the last 50 generations for each condition when robots could versus could not emit blue light (20 experiments per condition).

the communication system and result in a performance decrease. Thus, each communication strategy effectively constituted an adaptive peak separated by a valley with lower performance values.

Heterogeneous teams evolved with team-level selection reliably established communication protocols and displayed increased performance with respect to the control situation. However, their performance was similar to that of heterogeneous teams evolved with individual-level selection, who did not communicate. This result was analogous to the previous example where heterogeneous team evolved with team-level selection displayed a high number of altruistic foragers, but their performance was similar to that of heterogeneous teams evolved with team-level selection, who had very few altruistic foragers.

3 Conclusion

We have presented four algorithms for evolving robot collectives that are presented with situations where altruistic cooperation can lead to a performance increment. Only three of the four algorithms lead to altruistic cooperation, as predicted by kin selection and levels of selection. Heterogeneous teams of robots evolved with individual-level selection do not display altruistic cooperation and consequently result in lower fitness values in tasks that require altruistic cooperation.

Heterogeneous teams evolved with team-level selection represent a special case because in both examples they did evolve stable altruistic cooperators, but their fitness was lower than that of homogeneous teams.

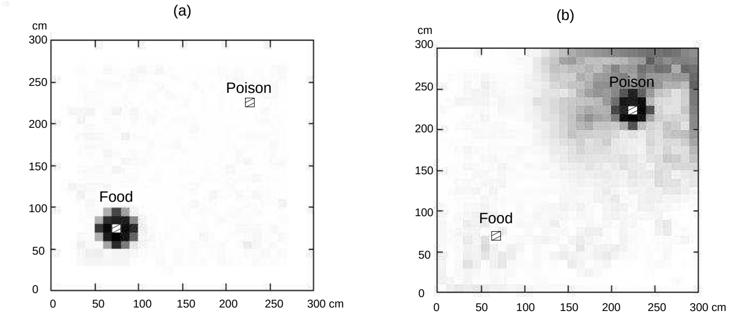

Fig. 10. Signaling frequency measured in each area of the arena for robots from two different evolved teams. a) The team was one where robots signal the presence of food. b) In this team robots signal the presence of poison. The darkness of each square is proportional to the amount of signaling in that area of the arena. From [8].

We think that this was due to the fact that after making copies of the individuals belonging to the best teams, those individuals were mated with individuals from other teams and randomly re-grouped in new teams. Although this was biologically plausible and necessary to prevent the genetic convergence of inbreeding teams, which would have rapidly led to homogeneous teams and thus confused the experimental design, it resulted in sub-optimal performance because combinations of well-integrated diverse individuals were disrupted at every generation.

From a practical perspective, homogeneous teams evolved with team-level selection are recommended for tasks that can benefit from altruistic cooperation. Not only do they bring together both conditions for the emergence of reliable altruism and thus result in higher performance, but they also do not require the need for separately computing the individual performance of each individual in a team. This is particularly useful in robotic tasks where only the resulting work of the team is known, but not what each robot in the team did and how.

We would like to emphasize that the results described in this chapter are specific to the case where there is an opportunity for altruistic cooperation and where altruistic cooperation results in higher fitness. We are currently expanding this line of investigation into three directions. First, we systematically compare the four evolutionary conditions described in this paper across experimental scenarios that require different degrees of cooperation, ranging from simple coordination to cooperation without a cost all the way to altruistic cooperation. Second, we compare these evolutionary conditions with other evolutionary methods in tasks that can benefit from non-trivial division of labor. Third, we compare the four evolutionary conditions in situations where the individuals in the team have a specific identity and can recognize each other, which was not the case in these experiments.

The study of the evolution of robotic collectives is not only promising for developing efficient control systems and testing biological hypotheses, but may also have an impact in a larger number of areas that require an optimal trade-off between the good of the individual and that of the society, such as internet agents, plant optimization, logistics, and economics.

Acknowledgments

The authors gratefully acknowledge support from the Swiss National Science Foundation and from the European projects IST-FET ECAgents and IST-FET Swarmanoids.

Bibliography

[1] G. Baldassarre, S. Nolfi, and D. Parisi. Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9:255–267, 2003.

[2] G. Baldassarre, D. Parisi, and S. Nolfi. Coordination and behavior integration in cooperating simulated robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J-A. Meyer, editors, *From Animals to Animats 8: Proceedings of the VIII International Conference on Simulation of Adaptive Behavior*, Cambridge (MA), 2003. MIT Press.

[3] C. Baray. Evolving cooperation via communication in homogeneous multi-agent systems. *Intelligent Information Systems, 1997. IIS'97. Proceedings*, pages 204–208, 1997.

[4] E. Bonabeau, G. Theraulaz, B. Schatz, and J.-L. Deneubourg. Response threshold model of division of labour in a ponerine ant. *Proc. Roy. Soc. London B*, 265:327–335, 1998.

[5] A. F. G. Bourke and N. R. Franks. *Social Evolution in Ants.* Princeton University Press, Princeton, NJ, 1995.

[6] B.D. Bryant and R. Miikkulainen. Neuroevolution for adaptive teams. In *Proceedings of the 2003 Congress on Evolutionary Computation, CEC '03*, volume 3, pages 2194–2201, 2003.

[7] R. Dawkins. *The Selfish Gene.* Oxford University Press, Oxford, 1976.

[8] D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519, 2007.

[9] J.J. Grefenstette. Credit assignment in rule discovery systems based on genetic algorithms. *Machine Learning*, 3(2):225–245, 1988.

[10] J. B. S. Haldane. Population genetics. *New Biology*, 18:34–51, 1955.

[11] W. D. Hamilton. The genetical evolution of social behavior, I and II. *Journal of Theoretical Biology*, 7:1–52, 1964.

[12] L. Keller and M. Chapuisat. Eusociality and cooperation. In *Encyclopedia of Life Sciences*, pages 1–8. MacMillan, London, 2002.

[13] L. Keller and M. G. Surette. Communication in bacteria. *Nature Reviews Microbiology*, 4:249–258, 2006.

[14] L. Lehmann and L. Keller. The evolution of cooperation and altruism – a general framework and a classification of models. *Journal of Evolutionary Biology*, 19:1365–1376, 2006.

[15] S. Luke, C. Hohn, J. Farris, G. Jackson, and J. Hendler. Co-evolving soccer softbot team coordination with genetic programming. In *Proceedings of The First International Workshop On RoboCup, IJCAI-97, Nagoya*, pages 214–222. Springer-Verlag, 1997.

[16] J. Maynard-Smith. Group selection and kin selection. *Nature*, 201:1145–1147, 1964.

[17] J. Maynard-Smith and D. Harper. *Animal Signals.* Oxford University Press, Oxford, UK, 2003.

[18] R. E. Michod. *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality.* Princeton University Press, Princeton, NJ, 1999.

[19] M. Minsky. Steps toward artificial intelligence. *Proceedings of the Institute of Radio Engineers*, 49:8–30, 1961. Reprinted in Computers and Thought, McGraw-Hill, 1963.

[20] F. Mondada, G. Pettinaro, A. Guignard, I. Kwee, D. Floreano, J-L. Deneubourg, S. Nolfi, L.M. Gambardella, and M. Dorigo. Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, 17:193–221, 2004.

[21] M.A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577, 1998.

[22] M.A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437:1291–1298, 2005.

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

[24] M. Quinn, L. Smith, G. Mayley, and P. Husbands. Evolving teamwork and role-allocation with real robots. In *Artificial Life VIII: Proceedings of the Eighth International Conference*, pages 302–311. MIT Press, 2002.

[25] M.D. Richards, D. Whitley, J.R. Beveridge, T. Mytkowicz, D. Nguyen, and D. Rome. Evolving cooperative strategies for UAV teams. In *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, pages 1721–1728, 2005.

[26] V. Trianni, S. Nolfi, and M. Dorigo. Cooperative hole-avoidance in a swarm-bot. *Robotics Autonomous Systems*, 54:97–103, 2006.

[27] M. Waibel, L. Keller, and D. Floreano. Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions in Evolutionary Computation*, 2008. under review.

[28] E. O. Wilson. *The Insect Societies.* The Belknap Press of Harvard University Press, Cambridge, MA, 1971.

[29] E. O. Wilson. *Sociobiology: The New Synthesis.* The Belknap Press of Harvard University Press, Cambridge, MA, 2000. Twenty-Fifth Anniversary Edition.

[30] L. Wolpert and E. Szathmary. Multicellularity: Evolution and the egg. *Nature*, 420:745–745, 2002.

[31] V. C. Wynne-Edwards. *Evolution Through Group Selection.* Blackwell, Palo Alto, CA, 1986.

[32] A. Zahavi and A. Zahavi. *The Handicap Principle. A Missing Piece of Darwin's Puzzle.* Oxford University Press, New York, 1997.