Genetic Team Composition and Level of Selection in the Evolution of Cooperation

Markus Waibel, *Member, IEEE*, Laurent Keller, and Dario Floreano, *Senior Member, IEEE*

***Abstract*— In cooperative multi-agent systems agents interact to solve tasks. Global dynamics of multi-agent teams result from local agent interactions, and are complex and difficult to predict. Evolutionary computation has proven a promising approach to the design of such teams. The majority of current studies use teams composed of agents with identical control rules ("genetically homogeneous teams") and select behavior at the team level ("team-level selection"). Here we extend current approaches to include four combinations of genetic team composition and level of selection. We compare the performance of genetically homogeneous teams evolved with individual-level selection, genetically homogeneous teams evolved with team-level selection, genetically heterogeneous teams evolved with individual-level selection and genetically heterogeneous teams evolved with team-level selection. We use a simulated foraging task to show that the optimal combination depends on the amount of cooperation required by the task. Accordingly, we distinguish between three types of cooperative tasks and suggest guidelines for the optimal choice of genetic team composition and level of selection.**

***Index Terms*— Multi-agent systems, artificial evolution, evolutionary robotics, team composition, fitness allocation, cooperation, altruism.**

I. Introduction

MULTI-AGENT SYSTEMS (MAS) span a large number of research fields, from software agents to robotics, and play a key role in several industrial applications, such as ground and air vehicle control, supply chains or network routing. The design of control rules for multi-agent systems is challenging because agent behavior depends not only on interactions with the environment, but also on the behavior of other agents. As the number of interacting agents in a team grows, or when agent behaviors become more sophisticated, the design of suitable control rules rapidly becomes very complex. This is especially true when agents are expected to coordinate or cooperate to collectively achieve a desired task. Evolutionary computation has been advocated as an effective and promising strategy to generate control parameters and decision rules for collective agents [1], [2].

In addition to the methodological issues of evolving agents that operate in isolation [2], the evolution of agent teams must address two major issues: (1) It must determine optimal team composition. Agents of a team may either share control rules (genetically homogeneous teams) or employ different ones (genetically heterogeneous teams). (2) It requires a suitable method for selective reproduction of desired team behavior.

MW and DF are with the Laboratory of Intelligent Systems (LIS), Ecole Polytechnique Fédérale de Lausanne, Station 11, CH-1015 Lausanne, Switzerland

LK is with the Department of Ecology and Evolution (DEE), Biophore, University of Lausanne, CH-1015 Lausanne, Switzerland

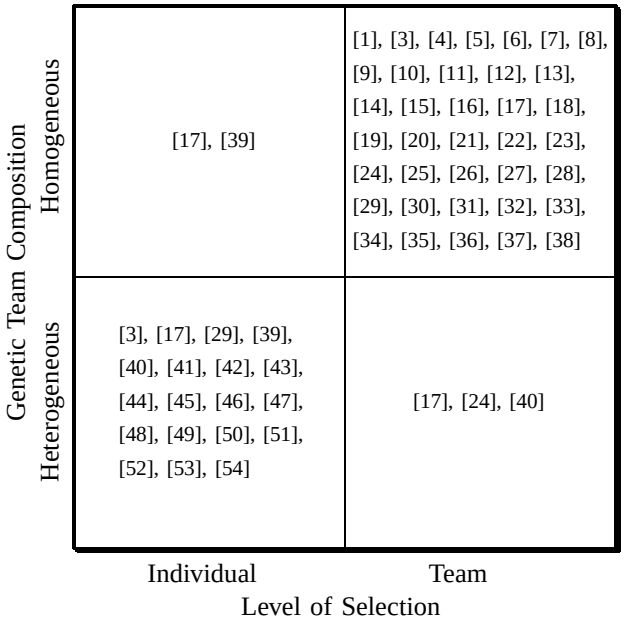

Fig. 1. A sample of approaches to the evolution of multi-agent teams. The majority of work uses genetically homogeneous teams, usually created from a cloned individual genome, with team selection. In some cases, authors created behaviorally heterogeneous agents out of a single team genome: Luke [25], [26] decoded team genomes into six separate sub-teams with one or two identical players each. Other authors [11], [20], [22], [28], [33] decoded one team genome into different single agent genomes. Yet another approach was taken by work using distributed, embodied evolution to evolve heterogeneous teams [43], [50]–[54]. In these cases selection and replication were entirely distributed among agents, with dynamics reminiscent of the replicator dynamics observed in bacterial evolution [55] and game theoretic models [56]. In some cases, teams were evolved using a continuously updated gene-pool rather than separate gene-pools for subsequent generations ("steady state evolution") [45]–[47]. Finally some authors have conducted more detailed comparisons of the influence of genetic team composition or level of selection alone: Martinoli [40] also considered more complex methods of selection. Stanley et al. [49] clustered genetically similar individuals into sub-teams that shared fitness, which resulted in partially heterogeneous teams. Mirolli et al. [39] also compared partially heterogeneous teams. Quinn [29] evaluated individuals in different heterogeneous teams to create robust homogeneous teams.

Selection may operate either on individuals (individual-level selection) or on teams (team-level selection). In the simplest case, one must decide between genetically homogeneous or heterogeneous teams, and between selecting agents at the individual or at the team level.

Figure 1 shows a sample of previous work on the evolution of multi-agent systems in robotics, combinatorial optimization, cellular automata, artificial life, genetic programming and others, plotted according to the chosen genetic team composition and level of selection. In addition to work cited in Figure 1, some authors have used cooperative co-evolutionary

algorithms (CCEAs, [57]) to evolve heterogeneous control rules for teams of agents [58]–[61]. CCEAs are applied by decomposing problem representations into subcomponents and then creating a separate population of individuals for each subcomponent. This approach allows teams to be composed of specialized sub-groups and corresponds to the biological co-evolution of multiple species. In their basic form, CCEAs require the designer to manually decompose the multi-agent task, and thus to solve part of the optimization problem beforehand. Work that used machine learning techniques other than evolutionary computation (e.g., reinforcement learning) was not considered in this review.

Figure 1 suggests that the majority of current approaches to the evolution of multi-agent systems use genetically homogeneous teams evolved with team-level selection. Where the reasons for the choice of genetically homogeneous teams are made explicit, it is argued that homogeneous teams are easy to use [8], [36], require fewer evaluations [25], [32], scale more easily [13] and are more robust against the failure of team members [13], [62] than heterogeneous teams. Many other approaches use genetically heterogeneous teams evolved with individual-level selection. Genetically heterogeneous teams are sometimes seen as providing more behavioral flexibility [25] and as providing advantages in tasks that require specialization [7], [25], [62], [63].

The terms “homogeneous team” and “heterogeneous team” used in the current literature cover many different aspects. It is important to note that while all agents in genetically homogeneous teams share the same genes, agents can nevertheless be behaviorally heterogeneous. This can happen when agents differentiate during their lifetime, for example due to varying initial conditions [30], or due to developmental processes or learning [64]. This can also happen when agents “activate” different parts of their genome, for example when each agent’s behavior is controlled by a different section of a single team genome [11], [22], [28], [33]. In this case, agents can specialize on different functions, yet be genetically identical, just like specialized cells in a biological organism. Conversely, it is important to note that genetically heterogeneous teams are those in which agents are, on average, not genetically more similar to team members than to agents in the rest of the population [65], [66]. This means that teams resulting from embodied evolution or common versions of steady state evolution are usually genetically heterogeneous although these algorithms often generate multiple offspring from a single parent, resulting in genetically similar (but not identical) agents. In some cases, teams consist of clonal subteams [25], [26] or of agents that share only part of their genome. Teams with agents that are, on average, genetically more similar (but not identical) to members of their team than to members of the rest of the population are termed “partially heterogeneous”. The effects of partial genetic heterogeneity on the evolution of multi-agent teams are not yet fully explored in evolutionary computation [39], but there is evidence that they can lead to improved specialization [25], [26]. These effects have been deeply studied in biology [67], [68].

The choice of level of selection is rarely discussed explicitly. Some research has addressed the related issue of credit assignment for the evolution of multi-agent systems [40], [69]. In the context of multi-agent systems, credit assignment is concerned with distributing fitness rewards among individual agents. Fitness distribution leads to credit assignment problems [70], [71] in many cooperative multi-agent tasks, because individual contributions to team performance are often difficult to estimate or difficult to monitor [72]. Selection is usually performed on the basis of accumulated individual or team fitness, which may be the result of many fitness rewards with different types of credit assignment. Therefore an optimal choice of level of selection is not only influenced by the type of task but also by the types of credit assignment used.

Genetic team composition and level of selection have long been identified as two important factors for the evolution of biological multi-agent teams such as groups of genes, cells, individuals or other replicators [67], [73]. In particular the evolution of altruism [74], in which agents cooperate to increase team fitness in spite of an individual fitness cost to the cooperator, has received a lot of attention [68], [75]. Here we define cooperation as a behavior that increases the fitness of other agents, and altruistic cooperation (altruism) as a behavior that increases the fitness of other agents and decreases the cooperator’s fitness.

In this study, we focus on cooperative multi-agent tasks that do not require specialization. We compare the performance of robot teams evolved in four evolutionary conditions: 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. We evaluate the performance of robot teams evolved in these four evolutionary conditions for three classes of multi-robot tasks: a task that does not require cooperation; a task that requires cooperation but does not imply a cost for cooperators; and a task that requires altruistic cooperation, i.e., a task that implies an individual fitness cost for cooperators. Cooperative tasks that benefit from specialization were not considered in this study.

II. Evolutionary conditions

The four possible combinations of genetic team composition and level of selection were formalized into four evolutionary algorithms (Figure 2). For the remainder of the paper we will use the terms “homogeneous” and “heterogeneous” to designate genetically homogeneous and genetically heterogeneous teams, respectively, and the terms “individual selection” and “team selection” to designate teams evolved with individual-level selection and team-level selection, respectively. We considered populations composed of M teams, each composed of N individuals. Population size and team sizes were kept constant across generations. At each generation, the old population was entirely replaced by a new population of offspring. Individuals’ genomes were binary strings.

1) Algorithm 1 - Homogeneous teams, Individual selection: Each of the M teams at generation 0 was formed by generating one random genome and cloning it $N-1$ times to obtain N identical robot genomes (clones) per team. Teams were

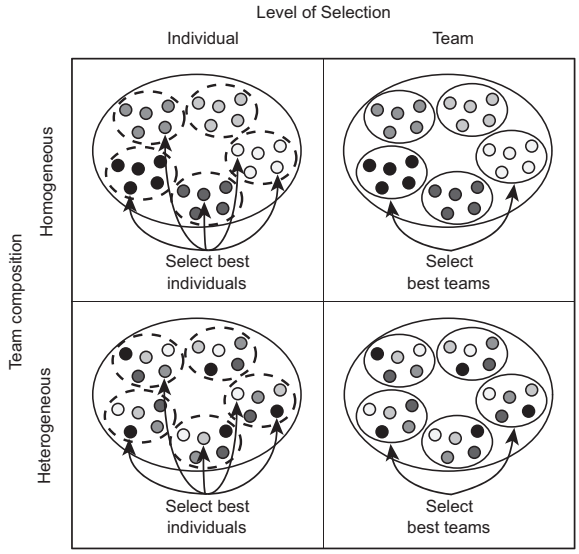

Fig. 2. The four evolutionary conditions. A population (large oval) was composed of several teams (medium ovals), each of which was composed of several robots (small circles) evaluated together. Genetic team composition was varied by either composing teams of robots with *identical genomes* (homogeneous, identical shading), or *different genomes* (heterogeneous, different shading). Level of selection was varied by selecting *teams* (team selection), or selecting *individuals*, independent of their team affiliation (individual selection).

Algorithm 1 Homogeneous teams, Individual selection

for each of M new teams **do**
  select two individuals from all old teams
  recombine their genomes to create one new genome
  mutate new genome
  clone new genome to obtain N genomes for new team
end for

evaluated in the task and an individual fitness determined for each of the N robots. For a new generation, each of the M new teams was created from two *individuals* selected among all individuals of all old teams in the population using roulette wheel selection. The two genomes of the selected individuals were recombined (one-point crossover, crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05 and then cloned $N-1$ times to generate the N robot genomes of the new team. Teams evolved using this evolutionary condition were thus genetically *homogeneous*.

Algorithm 2 Homogeneous teams, Team selection

for each of M new teams **do**
  select two old teams
  recombine their genomes to create one new genome
  mutate new genome
  clone new genome to obtain N genomes for new team
end for

2) Algorithm 2 - Homogeneous teams, Team selection: Each of the M teams at generation 0 was formed by generating one random genome and cloning it $N-1$ times to obtain N identical robot genomes (clones) per team. Teams were evaluated in the task, and for each team, a team fitness was determined as the sum of the individual fitnesses of all N robots. For a new generation, each of the M new teams was created from two old *teams* selected using roulette wheel selection. The two genomes of the selected teams were recombined (one-point crossover, crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05 and then cloned $N-1$ times to obtain the N robot genomes of the new team. Teams evolved using this evolutionary condition were thus genetically *homogeneous*.

Algorithm 3 Heterogeneous teams, Individual selection

for each of M new teams **do**
  for each of N new team members **do**
    select two individuals from all old teams
    recombine their genomes to create one new genome
    mutate new genome
    add new genome to new team
  end for
end for

3) Algorithm 3 - Heterogeneous teams, Individual selection: Each of the M teams at generation 0 was formed by generating N random genomes. Teams were evaluated in the task and an individual fitness determined for each of the N robots. For a new generation, each of the $N \times M$ new individuals was created from two *individuals* selected among all individuals of all old teams in the population using roulette wheel selection. The two genomes of the selected individuals were recombined (one-point crossover, crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05. This process was repeated $N \times M-1$ times to form M new teams of N individuals each. In this evolutionary condition robots were not, on average, genetically more similar to team members than to robots in the rest of the population, and thus teams were genetically *heterogeneous*.

Algorithm 4 Heterogeneous teams, Team selection

for each of M new teams **do**
  for each of N new team members **do**
    select two old teams
    randomly select two old team members
    recombine their genomes to create one new genome
    mutate new genome
    add new genome to new team
  end for
end for

4) Algorithm 4 - Heterogeneous teams, Team selection: Each of the M teams at generation 0 was formed by generating N random genomes. Teams were evaluated in the task, and for each team, a team fitness was determined as the sum of the individual fitnesses of all N robots. For a new generation, each of the $N \times M$ individuals was created from two old *teams* selected using roulette wheel selection. Two genomes, each

randomly selected among the members of a selected team, were recombined (one-point crossover, crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05. This process was repeated $N \times M - 1$ times to form M new teams of N individuals each. In this evolutionary condition robots were not, on average, genetically more similar to team members than to robots in the rest of the population, and thus teams were genetically *heterogeneous*.

III. Experimental Method

A. Scenario

The experimental setup (Figure 3) consisted of a 50×50 cm^2 arena with 10 micro-robots and two types of tokens, small and large. We chose to study a foraging task, because foraging combines several aspects of multi-agent tasks (distributed search, coordinated movement, transportation) and relates to many real-world problems [76], [77]. In addition, foraging is a wide-spread and well-studied behavior of many biological societies [78]–[80]. Experiments were conducted in simulation of micro-robots and evolved controllers were transferred to the real robots (see Section III-D).

Robots foraged tokens by transporting them into a 4 cm wide region at one side of the arena marked by a white wall. A single robot was sufficient to transport a small token. At least two robots were required to transport a large token, thus retrieval of large tokens required cooperation. Cooperating agents had to coordinate their behaviors to successfully align their positions before and during token transport.

The micro-robots [81] were small ($2 \times 2 \times 4$ cm^3), two-wheeled robots equipped with three infrared distance sensors at the front and one at the back, which could sense objects up to 3 cm away and allowed robots to distinguish between small and large tokens (Figure 4 left). An extension module with a fourth infrared distance sensor with a range of up to 6 cm and a linear camera were mounted higher on the robot, overlooking tokens but sensitive to other robots and walls.

B. Control and Genetic Architecture

Robots were controlled by a feed-forward neural network with a single layer of three hidden neurons (Figure 4 right) and a sigmoid activation function ($tanh$). The inputs were given by the activation values of five infrared sensors, two vision sensors, and a constant bias value of -1. Infrared sensor activation values were scaled in the range $[0; 1]$. Vision sensors were an average of three equi-distally spread camera pixels spanning a field of view of 18°, for the left or right side of the image, respectively. The averages were thresholded to yield 0 for a white or 1 for a black arena wall. Using the average value of three pixels rather than a single pixel allowed a robust detection of the white foraging target area in spite of the presence of other robots in the field of view. The two output units were used to control the left and right wheel motors. The activation values in the range $[-1; 1]$ were mapped into speeds in the range $[-4; 4]$ cm/s, with speeds in the interval of $[-2.5; 2.5]$ cm/s set to 0 because of unreliable motor response at low speeds.

The neural network connection weights were in the range of $[-2; 2]$ and coded on 8 bits. The genome of one individual was thus 8x32 bits long.

C. Collective Tasks

We devised three types of foraging tasks that differed in the amount of cooperation required from agents.

1) Task 1 - Individual Foraging: The arena contained 6 small tokens, which each awarded 1 fitness point to the foraging robot. This task did not require cooperation, because a single agent was sufficient to transport a small token.

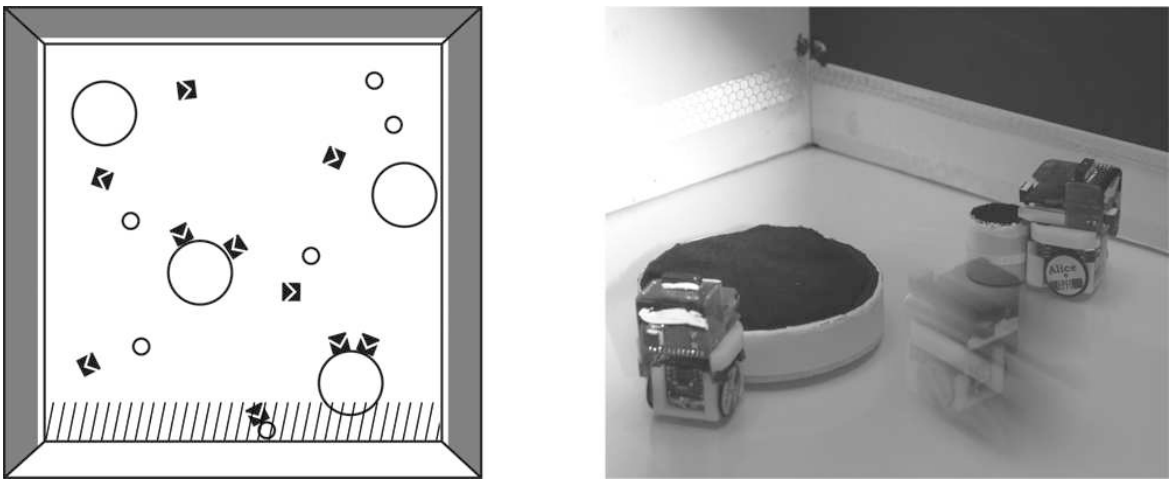

Fig. 3. Left: The experimental setup for task 3, the altruistic cooperative foraging task. Ten micro-robots (black squares with arrows) searched for small and large tokens and transported them to the target area (hatched area at bottom) under the white wall (the other three walls were black). An identical setup was used in the other two tasks, except that the arena contained either only small tokens in task 1, or only large tokens in task 2. Right: Three micro-robots in task 3, the altruistic cooperative foraging task. The robot in the background could transport the small token by itself. The robot at the left could not transport the large token by itself and needed to wait for the arrival of a second robot (blurred in the picture due to its rapid movement towards the large token).

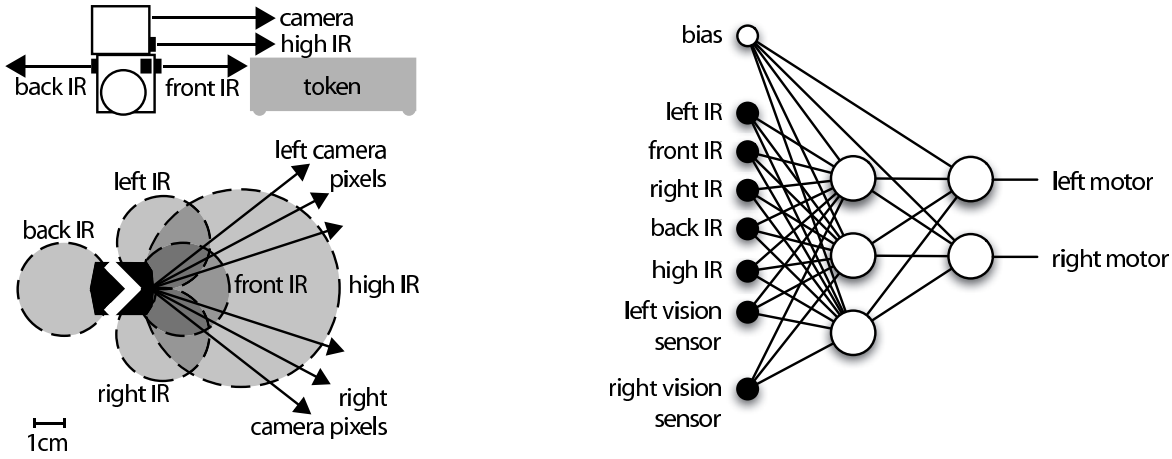

Fig. 4. Left: Side and top-view schematics of a simulated micro-robot. The robot was equipped with four infrared (IR) distance sensors (three at the front, one at the back) to detect tokens, and a camera to identify the target area. A fifth infrared distance sensor (high IR) was mounted higher on the robot and thus overlooked tokens. This allowed robots to distinguish tokens from walls and other robots. Right: The neural network architecture, a feed-forward neural network with a single layer of three hidden neurons. Inputs were given by the activation values of five infrared (IR) sensors and two vision sensors with activation values computed from left and right camera pixels (see text).

2) Task 2 - Cooperative Foraging: The arena contained 4 large tokens, which each awarded 1 fitness point to each team member, irrespective of its participation in the token foraging. This corresponded to a situation where the individual contributions to team performance were not known, i.e., a situation with credit assignment problems [70], [71], which is the case for many cooperative multi-agent tasks [72]. This task required cooperation because it could not be accomplished by a single agent.

3) Task 3 - Altruistic Cooperative Foraging: The arena contained 6 small and 4 large tokens. Small tokens each awarded 1 fitness point to the foraging robot and large tokens each awarded 1 fitness point to each team member, irrespective of their participation in the token foraging. In this task cooperation was costly for individuals, because individuals that did not cooperate always had higher fitness than their cooperating team mates. This meant that cooperators suffered a relative individual fitness cost and therefore this task required altruistic cooperation [68].

D. Evolutionary Experiments

Due to the large number of evaluations required for the evolution of robot behaviors, all evolutionary experiments were conducted using a physics-based 2D simulator [82], which is available as part of an open evolutionary framework [83]. All simulation parameters, including robot size, shape, speed and weight, as well as collision dynamics, friction forces and sensor and actuator modalities, were based on the micro-robots described in Section III-A.

We evolved teams of robots under the four evolutionary conditions separately for each of the three tasks, making a total of 12 experimental lines. Evolutionary experiments lasted for 300 generations. Twenty independent runs were performed for each experimental line. Populations consisted of 100 teams of 10 agents each. Each team was evaluated 10 times for three minutes with random token and robot starting positions and orientations. Fitness was averaged over the 10 evaluations.

To compare the efficiency of the four evolutionary conditions, we re-evaluated the best teams at generation 300 for 1000 times and compared their team fitness. Fitness values were analyzed using Wilcoxon rank sum tests. All fitness values were normalized for each task, with 0 being the minimal possible fitness and 1 the theoretical maximum value.

IV. Results

A. Task 1 - Individual Foraging

Successful foraging behavior evolved for all four evolutionary conditions (Figure 5). After 300 generations of artificial evolution, heterogeneous teams evolved with individual selection collected all 10 tokens in most evaluations and achieved fitness values close to the maximum value achievable. These fitness values were higher than those of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection (Wilcoxon rank sum test, $df = 38$, $P < 0.001$ and $P < 0.006$, respectively). A possible reason could be the disparities in genome evaluation in homogeneous and heterogeneous teams. For a team size of N agents, heterogeneous teams evaluated N times more genomes than homogeneous teams. This was because each heterogeneous team consisted of N different genomes, whereas homogeneous teams consisted of N identical genomes. On the other hand, homogeneous teams evaluated each genome N times more often than heterogeneous teams. This was because each team evaluation evaluated an identical genome N times. Our results suggest that higher evaluation accuracy may have been less important than a larger number of different genomes in this task. The larger number of genomes may have allowed heterogeneous teams to discover solutions faster than homogeneous teams, which could explain the steep initial fitness increase. It may also have allowed heterogeneous teams to discover better solutions than homogeneous teams, which could explain the higher final fitness obtained with this evolutionary condition. To test whether these disparities in genome evaluation caused the high team performance of heterogeneous teams evolved

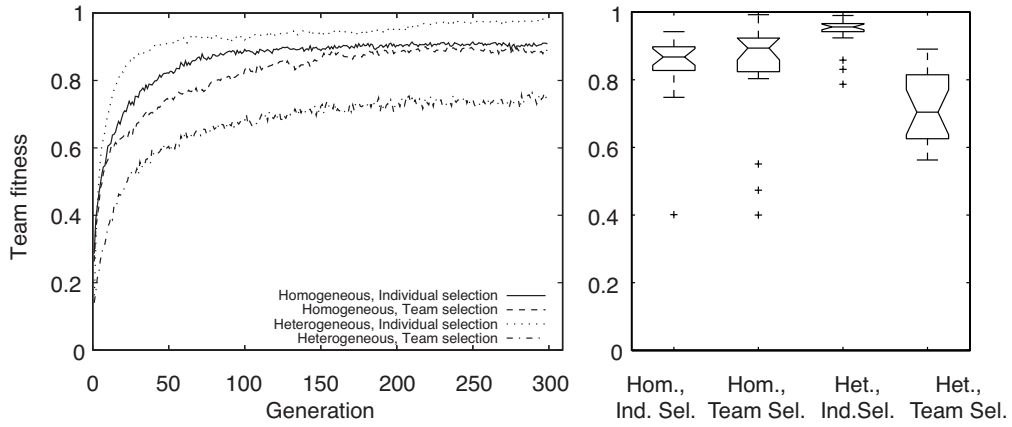

Fig. 5. Task 1 - Individual Foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. The mid line in the box is the median, while the box represents the upper and lower quartile above and below the median. The bars outside the box generally represent the max and min values, except when there are outliers, which are shown as crosses. We define outliers as data points which differ more than 1.5 times the interquartile range from the border of the box. The notches represent the uncertainty in the difference of the medians for box-to-box comparison. Boxes whose notches do not overlap indicate that the medians differ at the 5 % significance level [84]. In this task, which did not require cooperation, heterogeneous teams evolved with individual selection performed best, followed by homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection. Heterogeneous teams evolved with team selection performed significantly worse than all other evolutionary conditions.

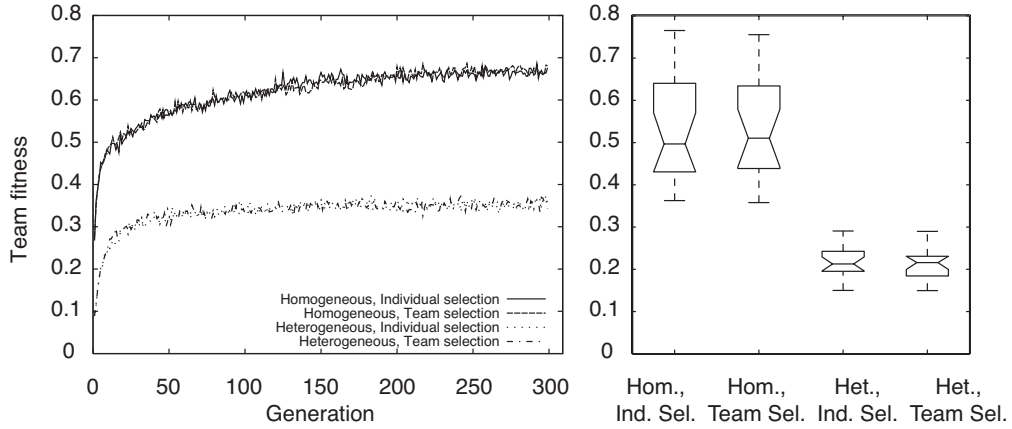

Fig. 6. Task 2 - Cooperative Foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. Homogeneous teams performed significantly better than heterogeneous teams. Boxplot explanations see Figure 5.

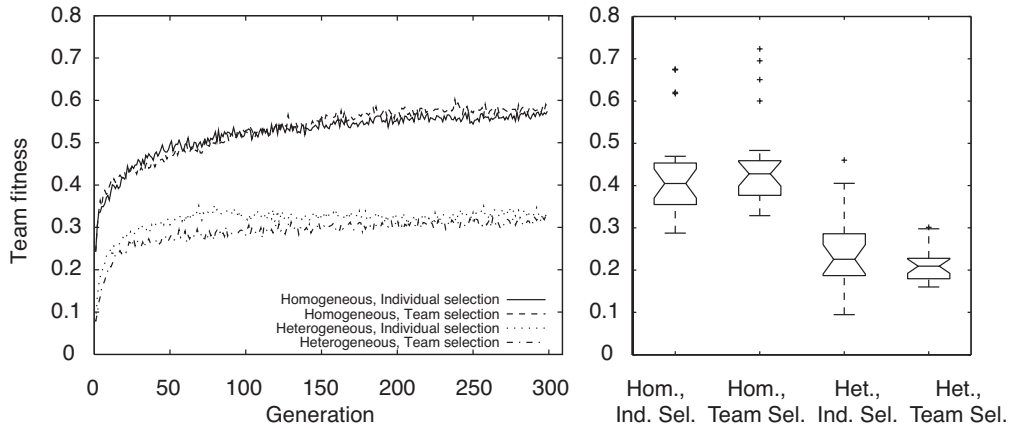

Fig. 7. Task 3 - Altruistic Cooperative Foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. Homogeneous teams performed significantly better than heterogeneous teams. Boxplot explanations see Figure 5.

with individual selection we performed a set of additional experiments (see additional experiments without these disparities in the next section).

Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ($P = 0.337$). This was because with roulette wheel selection, the probability of a team to be selected was the same as the sum of the probabilities of each individual team member to be selected. Since all team members of homogeneous teams shared the same genome, selection probabilities for a given genome were equal for both homogeneous evolutionary conditions. It should be noted, however, that this is not necessarily true for other types of selection. Selection mechanisms where the fitness of a genome is not directly proportional to its probability to be selected (e.g., truncation or rank-based selection) may lead to differences in the number of selected individuals with a given genotype and consequently affect the relative performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection. In these cases individual selection may select for genomes that lead to higher maximum but lower average individual performance. However, additional investigations using truncation selection (selection of best 30% of the population; all other experimental parameters identical) did not find such performance differences in any of the three types of task ($P = 0.350/0.903/0.394$ for tasks $1/2/3$ respectively; see Figure S4 in the supplementary material[1]).

Heterogeneous teams evolved with team selection performed significantly worse than all other evolutionary conditions (all three $P < 0.002$). This was because, unlike all other three evolutionary conditions, this evolutionary condition did not allow a direct link between the performance of a genome and its probability to be selected. Instead, selection of good genomes could only happen indirectly, by selecting those teams that contained a better mix of genomes than other teams. Since good genomes could be part of bad teams or bad genomes part of good teams, selection for good individual genomes was inefficient. This explains the slow initial fitness increase and the lowest final fitness of heterogeneous teams evolved with team selection.

B. Task 1 - Individual Foraging: Disparities in genome evaluation and credit assignment

To test the hypothesis that the high team performance of heterogeneous teams evolved with individual selection was caused by disparities in genome evaluation, we performed a set of additional experiments. First, we evolved homogeneous teams in the same task, but used only 1 evaluation per team rather than 10 evaluations (Figure 8 (a)). Second, we evolved heterogeneous teams in the same task, but used only 100 agents per population rather than 1000 agents (Figure 8 (b)). In this set of experiments homogeneous and heterogeneous teams therefore evaluated the same number of genomes and had the same number of evaluations per genome.

[1]An electronic supplement for this paper is available online at `http://lis.epfl.ch/documentation.php`

Without disparities in genome evaluation, heterogeneous teams evolved with individual selection performed similarly to homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection (all three $P > 0.597$). Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions (all three $P < 0.001$), because the efficiency of selection was not affected by the changes in genome evaluation.

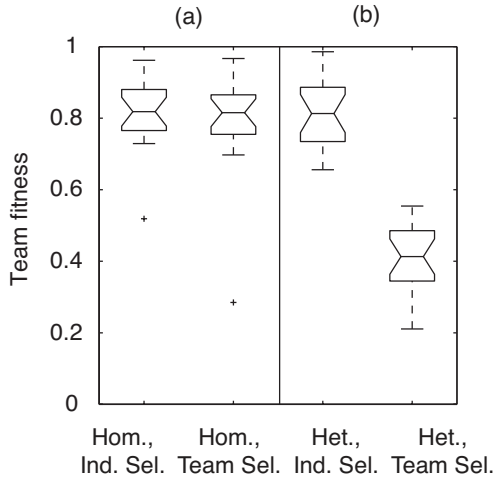

Fig. 8. Task 1 - Individual Foraging without disparities in genome evaluation. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). Heterogeneous teams evolved with individual selection performed similarly to homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection in this task. Boxplot explanations see Figure 5.

C. Task 2 - Cooperative Foraging

Successful foraging behavior evolved for all four evolutionary conditions (Figure 6). The experiments with a cooperative task led to a change in the relative performance of the four evolutionary conditions. Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection was significantly higher than that of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection (all four $P < 0.001$), with the best fitness values in homogeneous teams up to 70% higher than those in heterogeneous teams. One possible reason are disparities in genome evaluation between homogeneous and heterogeneous teams (see next section). Another possible reason is that selection of good genomes could only happen indirectly in this task, which may have led to inefficient selection just as in heterogeneous teams evolved with team selection in task 1 (Section IV-B). This could have been because fitness in this task was assigned to all team members, irrespective of their participation in the token foraging.

Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ($P = 0.839$). This was because fitness in this task was assigned to all team members, irrespective of their participation in the token foraging. For the

same reason, performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ($P = 0.365$).

D. Task 2 - Cooperative Foraging: Disparities in genome evaluation and credit assignment

To test the hypothesis that the differences in performance of heterogeneous teams evolved with individual selection and homogeneous teams evolved with individual selection and team selection were caused by disparities in genome evaluation or by the fitness assignment to all team members, we performed two sets of additional experiments. First, we again corrected for the disparities in genome evaluation. However, correcting for this factor alone did not eliminate the performance differences (see Figure S2 in the supplementary material[1]). Second, we performed experiments where we again corrected for the disparities in genome evaluation and where fitness was only assigned to team members that participated in the token foraging. In these experiments, each of the 4 large tokens awarded 5 fitness points to each of the two transporting robots, rather than 1 fitness point to each of the 10 team members. This second additional set of experiments therefore corresponded to a situation where the individual contributions to team performance were known, i.e., a situation without credit assignment problems.

Without the disparities in genome evaluation and without credit assignment problems, heterogeneous teams evolved with individual selection outperformed heterogeneous teams evolved with team selection ($P < 0.001$). This was because selection of good genomes could now happen directly, which allowed for efficient selection. However, the performance of heterogeneous teams evolved with individual selection remained lower than that of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection ($P < 0.001$ and $P < 0.002$, respectively, Figure 9). A possible reason is that heterogeneous teams had to solve a more complex optimization task than homogeneous teams. Successful cooperation in heterogeneous teams required individuals to evolve behaviors to coordinate their actions with $N - 1$ different team members, while individuals in homogeneous teams only had to evolve behaviors to coordinate with a single type of team member. In other words, homogeneous teams led to a smaller search space because all team members were per definition identical, and thus only a subset of the total number of possible team compositions was considered in these teams. Furthermore, individuals in heterogeneous teams were not just different in a team, but team members changed from one generation to the next. Both factors may have hindered the evolution of cooperative behavior in heterogeneous teams.

The performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ($P = 0.441$) in this second additional set of experiments.

Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions due to inefficient selection (all three $P < 0.001$).

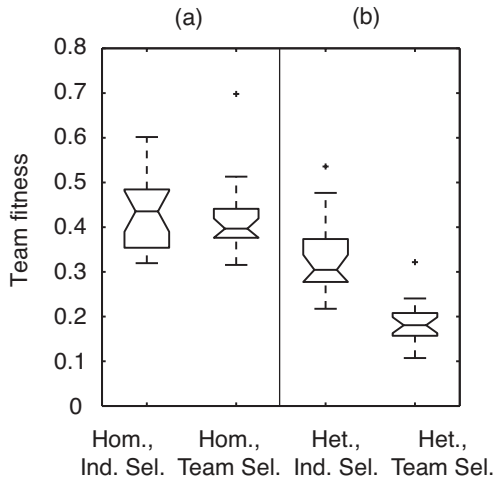

Fig. 9. Task 2 - Cooperative Foraging without disparities in genome evaluation and without credit assignment problems. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). The performance of heterogeneous teams evolved with individual selection was higher than the performance of heterogeneous teams evolved with team selection, but did not reach that of homogeneous teams in this task. Boxplot explanations see Figure 5.

E. Task 3 - Altruistic Cooperative Foraging

Successful foraging behaviors evolved for all four evolutionary conditions (Figure 7). Team performance in the altruistic cooperative foraging task was systematically lower than in the cooperative foraging task. This may seem surprising because the larger number of tokens in the arena increased the total number of fitness points available. A possible reason is that the increased number of tokens led to more clutter in the arena, which made successful token transport more difficult (see video supplied with supplementary material[1]).

Homogeneous teams achieved significantly higher fitness values than heterogeneous teams (all four $P < 0.001$). Possible reasons are disparities in genome evaluation and inefficient selection for the foraging of large tokens because fitness points gained from large tokens were assigned to all team members, irrespective of their participation in the token foraging (see next section).

Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ($P = 0.310$). This was because selection probabilities for a given genome were again equal for both homogeneous evolutionary conditions.

Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ($P = 0.490$). However, the four evolutionary conditions resulted in different foraging strategies in this task (Figure 10): While homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection as well as heterogeneous teams evolved with team selection collected a significantly higher proportion of large tokens than small tokens (all three $P < 0.001$), heterogeneous teams evolved with individual selection collected a significantly higher proportion of small tokens than large tokens ($P < 0.001$). In comparison to the other

three evolutionary conditions, heterogeneous teams evolved with individual selection collected the significantly highest proportion of small tokens (all three $P < 0.001$), but the significantly lowest proportion of large tokens of all four evolutionary conditions (all three $P < 0.003$).

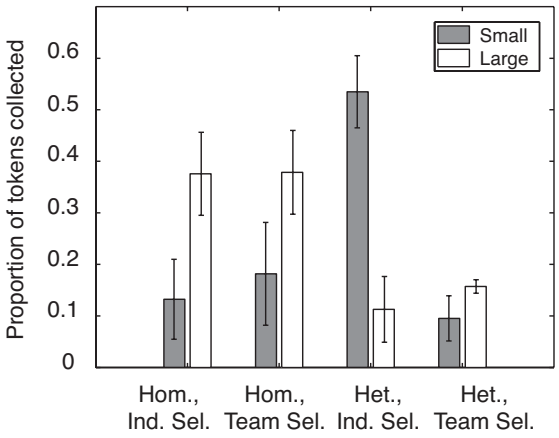

Fig. 10. Task 3 - Altruistic Cooperative Foraging. The plot shows the average proportion of the six small tokens and four large tokens collected by the best teams at generation 300 for each of the 20 independent experiments and for each of the four evolutionary conditions. Heterogeneous teams evolved with individual selection pursued a different foraging strategy than teams of the other three evolutionary conditions, collecting very few large tokens but most small tokens.

F. Task 3 - Altruistic Cooperative Foraging: Disparities in genome evaluation and credit assignment

To test the hypothesis that the differences in performance of heterogeneous teams evolved with individual selection were caused by disparities in genome evaluation or by the fitness assignment to all team members, we performed two sets of additional experiments similar to those described in Section IV-D. First, we again corrected for the disparities in genome evaluation. However, correcting for this factor alone did not eliminate the performance differences (see Figure S3 in the supplementary material[1]). Second, we again performed experiments where we corrected for the disparities in genome evaluation and for credit assignment problems.

Without the disparities in genome evaluation and without credit assignment problems heterogeneous teams evolved with individual selection outperformed heterogeneous teams evolved with team selection ($P < 0.001$). This was because selection of good genomes could again happen directly, which allowed for efficient selection. However, the performance of heterogeneous teams evolved with individual selection remained lower than that of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection ($P < 0.015$ and $P < 0.003$, respectively, Figure 11). This may have been because heterogeneous teams had to solve a more complex optimization task.

In this second additional set of experiments, the performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ ($P = 0.133$).

Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions due to inefficient selection (all three $P < 0.001$).

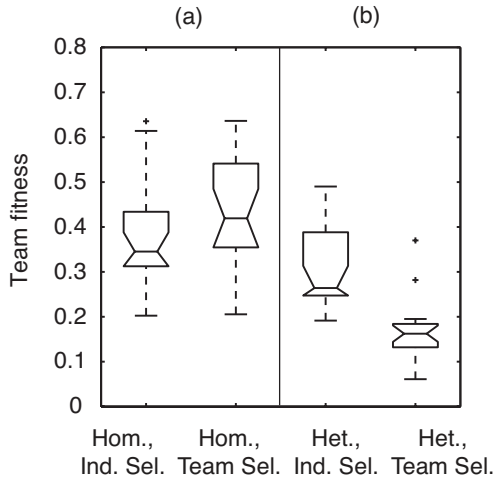

Fig. 11. Task 3 - Altruistic Cooperative Foraging without disparities in genome evaluation and without credit assignment problems. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). The performance of heterogeneous teams evolved with individual selection was higher than the performance of heterogeneous teams evolved with team selection, but did not reach that of homogeneous teams. Boxplot explanations see Figure 5.

Importantly, the altruistic cooperative foraging task led to the evolution of a different foraging strategy in heterogeneous teams evolved with individual selection than in the other three evolutionary conditions (Figure 10). A possible reason is that cooperation to collect large tokens now implied a cost for individuals. To test this hypothesis we performed additional experiments with this evolutionary condition. First we repeated the experiments with a setup identical to that of task 3, i.e., with 1000 agents per population and 10 evaluations per team, but with known individual contributions to large token foraging, i.e., a situation without credit assignment problems. Each of the 4 large tokens awarded 5 fitness points to each of the two transporting robots, rather than 1 fitness point to each of the 10 team members. Then, at generation 300, we changed the fitness assignment and assumed unknown individual contributions to large token foraging, i.e., a situation with credit assignment problems. Each of the 4 large tokens awarded 1 fitness point to each team member, irrespective of their participation in token foraging.

This change in fitness assignment resulted in a drastic change in foraging strategy (Figure 12). While at generation 300 heterogeneous teams evolved with individual selection collected a significantly higher proportion of large tokens than small tokens ($P < 0.001$), at generation 500 they collected a significantly lower proportion of large than small tokens ($P < 0.001$). As a direct result of this change, team performance decreased significantly ($P < 0.001$) between generation 300 and generation 500. This was because after the introduction of credit assignment problems, fitness points gained from large

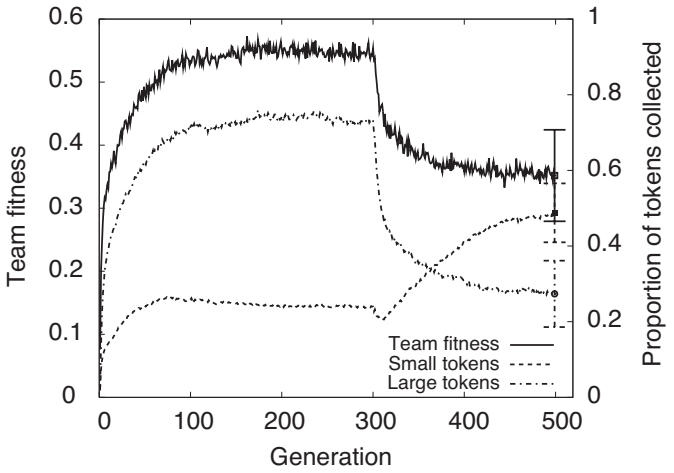

Fig. 12. Task 3 - Altruistic Cooperative Foraging in heterogeneous teams evolved with individual selection. For the first 300 generations, individual contributions to the cooperative foraging of large tokens were known (no credit assignment problems). From generation 300 onward individual contributions to the foraging of large tokens were presumed unknown (credit assignment problems). The introduction of credit assignment problems led to the rapid collapse of cooperation and a decrease in team fitness.

tokens were assigned to all team members, and therefore individuals collecting small tokens gained a fitness advantage over their team mates. This led to the selection of individuals that foraged for small tokens and resulted in fewer and fewer individuals foraging for large tokens. The observed drop in team fitness also implies a drop in average individual fitness. This illustrates that fitness is a relative measure of performance and therefore evolution selects for performance increase relative to the performance of competitors, rather than for absolute performance. The simplicity of the neural network controllers did not allow individuals to accurately discriminate large and small tokens, which explains the incomplete collapse of large token foraging.

In contrast, the foraging strategy in homogeneous teams evolved with individual selection and in homogeneous teams evolved with team selection was not affected by the costs implied in large token foraging (see Figures S5 and S6 in the supplementary material[1]). This was because relative fitness differences between team members could not have an influence on the selection of genomes when individuals were genetically identical.

Foraging strategy in heterogeneous teams evolved with team selection was not affected by the costs implied in large token foraging (see Figures S5 and S6 in the supplementary material[1]). This was because relative fitness differences between team members did not have an influence on selection of genomes when selection acted at the level of the team.

V. Conclusion

This study provides an experimental demonstration of how the choice of genetic team composition and level of selection influences the performance of multi-agent systems in tasks with varying levels of cooperation that do not provide a benefit for specialization. We have identified three different types of multi-agent tasks depending on the amount of cooperation required between team members. Our results demonstrate that different combinations of genetic team composition and level of selection lead to significant performance differences. No combination achieved optimal performance in all three types of task.

We have identified and studied three different types of multi-agents tasks depending on the amount of cooperation required between team members.

In tasks that did not require cooperation, heterogeneous teams evolved with individual level selection achieved the highest team performance. Team heterogeneity allowed to evaluate a high number of different genomes in parallel, and individual selection allowed efficient selection of good genomes. However, these teams performed poorly in tasks that required cooperation and in tasks with credit assignment problems.

For multi-agent tasks that required cooperation, the highest team performance was achieved by homogeneous teams. These teams led to efficient cooperation between team members and they were not affected by credit assignment problems or costs associated with cooperation. Our results suggest that homogeneous teams are a safe choice in tasks that do not benefit from specialization when the requirements for agent cooperation are difficult to estimate. Compared to heterogeneous teams, homogeneous teams evaluate less genomes, which may result in premature convergence to sub-optimal solutions. Our experimental results indicate that a simple way to prevent this problem is to use populations made of a large number of homogeneous teams.

Heterogeneous teams evolved with team selection were inefficient at selecting for good performance in all three types of tasks studied here and and therefore cannot be recommended for cooperative tasks that do not require specialization.

Converging evidence towards these guidelines was recently found in a study on the evolutionary conditions for the emergence of communication in societies of robots with different morphology and sensing abilities than those described in this paper [17].

However, it should be noted that evidence from both studies has two notable limitations. First, it did not address tasks that benefit from specialization in addition to cooperation. There is evidence that behavioral heterogeneity can lead to significant performance advantages for such tasks [29], [63], [85]–[88]. Second, it does not consider teams with intermediate genetic similarity. Biological research has shown that such teams can overcome individual fitness costs of cooperation [55], [89], thus combining the best of both worlds, enhanced genetic diversity with readiness to adopt altruistic behaviors. A good understanding of those conditions will require significant further research.

Acknowledgments

The authors thank Martijn Bosch, Peter Dürr, Claudio Mattiussi, Sara Mitri and Danesh Tarapore for insightful discussion as well as three anonymous referees for useful comments. This research was sponsored by the Swiss National Science Foundation and the European Commission (FP6, IST-FET projects ECAgents and Swarmanoids).

REFERENCES

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

[2] S. Nolfi and D. Floreano, *Evolutionary Robotics*. Bradford, MIT Press, 2000.

[3] A. Agogino, K. Stanley, and R. Miikkulainen, "Online interactive neuro-evolution," *Neural Processing Letters*, vol. 11, no. 1, pp. 29–38, 2000.

[4] D. Andre, "The automatic programming of agents that learn mental models and create simple plans of action," *IJCAI-95: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, vol. 1, pp. 741–747, 1995.

[5] D. Andre and A. Teller, *RoboCup-98: Robot Soccer World Cup II*. Springer Berlin / Heidelberg, 1999, vol. 1604, ch. Evolving Team Darwin United, pp. 346–351.

[6] G. Baldassarre, S. Nolfi, and D. Parisi, "Evolution of collective behavior in a team of physically linked robots," in *Applications in Evolutionary Computing*, S. E. A. Cagnoni, Ed. Berlin: Springer Verlag, 2003, pp. 581–592.

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

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

[9] ——, "Effects of population size upon emergent group behavior," *Complexity International*, vol. 6, 1998.

[10] J. K. Bassett and K. A. De Jong, "Evolving behaviors for cooperating agents," in *Foundations of Intelligent Systems: 12th International Symposium, ISMIS 2000, Proceedings.*, Z. W. Ras and S. Ohsuga, Eds. Springer Berlin / Heidelberg, 2000, pp. 157–165.

[11] J. C. Bongard, *Genetic Programming*. Springer Berlin / Heidelberg, 2000, vol. 1802/2004, ch. The Legion System: A Novel Approach to Evolving Heterogeneity for Collective Problem Solving, pp. 16–28.

[12] H. Botee and E. Bonabeau, "Evolving ant colony optimization," *Advances in Complex Systems*, vol. 1, no. 2/3, pp. 149–159, 1998.

[13] B. Bryant and R. Miikkulainen, "Neuroevolution for adaptive teams," in *Proceedings of the 2003 Congress on Evolutionary Computation, CEC '03*, vol. 3, 2003, pp. 2194–2201.

[14] L. Bull and O. Holland, "Evolutionary computing in multiagent environments: Eusociality," *Proceedings of the Annual Conference On Genetic Programming*, pp. 347–352, 1997.

[15] R. J. Collins and D. R. Jefferson, "Antfarm: Towards simulated evolution," in *Artificial Life II*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, Eds. Redwood City, CA: Addison-Wesley, 1992, pp. 579–601.

[16] M. Dorigo, V. Trianni, E. Scedilahin, R. Gross, T. H. Labella, G. Baldassarre, S. Nolfi, J. L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella, "Evolving self-organizing behaviors for a swarm-bot," *Autonomous Robots*, vol. 17, no. 2-3, pp. 223–245, 2004.

[17] D. Floreano, S. Mitri, S. Magnenat, and L. Keller, "Evolutionary Conditions for the Emergence of Communication in Robots," *Current Biology*, vol. 17, pp. 514–519, 2007.

[18] R. Gross and M. Dorigo, "Cooperative transport of objects of different shapes and sizes," *Lecture Notes in Computer Science*, vol. 3172, pp. 106–117, 2004.

[19] ——, "Evolving a cooperative transport behavior for two simple robots," *Lecture Notes in Computer Science*, vol. 2936, pp. 305–316, 2004.

[20] A. Hara and T. Nagao, "Emergence of cooperative behavior using ADG; Automatically Defined Groups," in *Proceedings of the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, Eds. Morgan Kaufmann, 1999, pp. 1038–1046.

[21] T. Haynes, R. Wainwright, S. Sen, and D. Schoenefeld, "Strongly typed genetic programming in evolving cooperation strategies," in *Proceedings of the Sixth International Conference on Genetic Algorithms*, S. Forrest, Ed. Morgan Kaufman, 1995, pp. 271–278.

[22] T. Haynes and S. Sen, "Crossover operators for evolving a team," in *Genetic Programming 1997: Proceedings of the Second Annual Conference*, J. R. Koza, K. Deb, M. Dorigo, D. B. Fogel, H. I. M. Garzon, and R. L. Riolo, Eds. Morgan Kaufmann: Stanford University, CA, USA, 1997, pp. 162–167.

[23] J. Lohn and J. Reggia, "Discovery of self-replicating structures using a genetic algorithm," in *Evolutionary Computation, 1995., IEEE International Conference on*, vol. 2, 1995, pp. 678–683.

[24] S. Luke and L. Spector, "Evolving teamwork and coordination with genetic programming," in *Proceedings of Genetic Programming 1996 (GP96), Stanford*, 1996, pp. 150–156.

[25] 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*. Springer-Verlag, 1997, pp. 214–222.

[26] S. Luke, "Genetic programming produced competitive soccer softbot teams for robocup97," in *Genetic Programming 1998: Proceedings of the Third Annual Conference*, W. Banzhaf, K. Chellapilla, K. Deb, M. Dorigo, D. Fogel, M. Garzon, D. Goldberg, I. H., J. Koza, and R. Riolo, Eds. Morgan Kaufmann, 1998, pp. 214–222.

[27] D. Marocco and S. Nolfi, "Self-organization of communication in evolving robots," in *Proceeding of the Tenth International Conference On Artificial Life: AlifeX*, L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone, and A. Vespignani, Eds. Boomington: MIT Press, 2006, pp. 199–205.

[28] T. Miconi, "When evolving populations is better than coevolving individuals: The blind mice problem," in *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, G. Gottlob and T. Walsh, Eds. Morgan Kaufmann, 2003, pp. 647–652.

[29] M. Quinn, "A comparison of approaches to the evolution of homogeneous multi-robot teams," in *Proceedings of the 2001 Congress On Evolutionary Computation: CEC2001*. Seoul and Korea: IEEE, 2001, pp. 128–135.

[30] M. Quinn, L. Smith, G. Mayley, and P. Husbands, "Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors," *Philosophical Transactions of the Royal Society of London and Series A: Mathematical and Physical and Engineering Sciences*, vol. 361, pp. 2321–2344, 2003.

[31] S. Raik and B. Durnota, "The evolution of sporting strategies," in *Complex Systems: Mechanisms of Adaption*, R. J. Stonier and X. H. Yu, Eds. Amsterdam, Netherlands: IOS Press, 1994, pp. 85–92.

[32] M. Richards, D. Whitley, J. 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*, 2005, pp. 1721–1728.

[33] A. Robinson and L. Spector, "Using genetic programming with multiple data types and automatic modularization to evolve decentralized and coordinated navigation in multi-agent systems," in *Late-Breaking Papers of the Genetic and Evolutionary Computation Conference (GECCO-2002)*. The International Society for Genetic and Evolutionary Computation, 2002.

[34] G. Saunders and J. Pollack, "The evolution of communication schemes over continuous channels," in *From Animals to Animats 4 - Proceedings of the Fourth International Conference on Adaptive Behaviour*, 1996.

[35] V. Trianni, R. Gross, T. H. Labella, E. Sahin, and M. Dorigo, "Evolving aggregation behaviors in a swarm of robots," *Lecture Notes in Computer Science*, vol. 2801, pp. 865–874, 2003.

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

[37] A. S. Wu, A. C. Schultz, and A. Agah, "Evolving control for distributed micro air vehicles," in *IEEE Conference on Computational Intelligence in Robotics*, 1999, pp. 174–179.

[38] N. Zaera, C. Cliff, and J. Bruten, "(Not) evolving collective behaviours in synthetic fish," in *Proceedings of the Forth International Conference On Simulation of Adaptive Behavior (SAB'96)*, vol. 4, 1996, pp. 635–644.

[39] M. Mirolli and D. Parisi, "How can we explain the emergence of a language that benefits the hearer but not the speaker," *Connection Science*, vol. 17, pp. 307–324, 2005.

[40] A. Martinoli, "Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies," Ph.D. dissertation, EPFL, 1999.

[41] A. Agah and G. A. Bekey, "Phylogenetic and ontogenetic learning in a colony of interacting robots," *Autonomous Robots*, vol. 4, pp. 85–100, 1997.

[42] A. Agah and K. Tanie, "Robots playing to win: evolutionary soccer strategies," *Robotics and Automation, 1997. Proceedings., 1997 IEEE International Conference on*, vol. 1, pp. 632–637, 1997.

[43] R. Bianco and S. Nolfi, "Toward open-ended evolutionary robotics: Evolving elementary robotic units able to self-assemble and self-reproduce," *Connection Science*, vol. 16, no. 4, pp. 227–248, 2004.

[44] A. Cangelosi and D. Parisi, "The emergence of a 'language' in an evolving population of neural networks," *Connection Science*, vol. 10, pp. 83–97, 1998.

[45] C. Reynolds, "An evolved, vision-based behavioral model of coordinated group motion," in *From Animals to Animats 2*. MIT Press, 1993, pp. 384–392.
[46] G. M. Werner and M. G. Dyer, "Evolution of herding behavior in artificial animals," in *Proceedings of the Second International Conference On From Animals to Animats 2 : Simulation of Adaptive Behavior*. MIT Press Cambridge, MA, USA, 1992, pp. 393–399.
[47] ——, "Evolution of communication in artificial organisms," in *Artificial Life II*, C. Langton, C. Taylor, D. Farmer, and S. Rasmussen, Eds. Redwood City, CA: Addison-Wesley Pub., 1992, pp. 659–687.
[48] C. R. Ward, F. Gobet, and G. Kendall, "Evolving collective behavior in an artificial ecology," *Artificial Life*, vol. 7, pp. 191–209, 2001.
[49] K. Stanley, B. Bryant, and R. Miikkulainen, "Real-time neuroevolution in the NERO video game," *Evolutionary Computation, IEEE Transactions on*, vol. 9, pp. 653–668, 2005.
[50] S. G. Ficici, R. A. Watson, and J. B. Pollack, "Embodied evolution: A response to challenges in evolutionary robotics," in *Eighth European Workshop On Learning*, 1999, pp. 14–22.
[51] E. D. V. Simoes and D. A. C. Barone, "Predation: An approach to improving the evolution of real robots with a distributed evolutionary controller," in *Proceedings of the 2002 IEEE International Conference On Robotics and Automation*, 2002, pp. 664–669.
[52] L. Spector, J. Klein, C. Perry, and M. Feinstein, "Emergence of collective behavior in evolving populations of flying agents," *Lecture Notes in Computer Science*, vol. 2723, pp. 61–73, 2003.
[53] ——, "Emergence of collective behavior in evolving populations of flying agents," *Genetic Programming and Evolvable Machines*, vol. 6, pp. 111–125, 2005.
[54] R. A. Watson, S. G. Ficici, and J. B. Pollack, "Embodied evolution: Distributing an evolutionary algorithm in a population of robots," *Robotics and Autonomous Systems*, vol. 39, pp. 1–18, 2002.
[55] A. S. Griffin, S. A. West, and A. Buckling, "Cooperation and competition in pathogenic bacteria," *Nature*, vol. 430, pp. 1024–1027, 2004.
[56] C. Hauert, S. De Monte, J. Hofbauer, and K. Sigmund, "Replicator dynamics for optional public good games," *Journal of Theoretical Biology*, vol. 218, pp. 187–194, 2002.
[57] M. A. Potter and D. V. De Jong, "A cooperative coevolutionary approach to function optimization," *Lecture Notes in Computer Science*, vol. 866, pp. 249–257, 1994.
[58] H. J. Blumenthal and G. B. Parker, "Co-evolving team capture strategies for dissimilar robots," in *Proceedings of the AAAI 2004 Symposium On Artificial Multiagent Learning, Washington.*, 2004.
[59] M. A. Potter and D. V. De Jong, "Cooperative coevolution: An architecture for evolving coadapted subcomponents," *Evolutionary Computation*, vol. 8, pp. 1–29, 2000.
[60] R. P. Wiegand, W. C. Liles, and D. V. DeJong, "An empirical analysis of collaboration methods in cooperative coevolutionary algorithms," *Proceedings of the Genetic and Evolutionary Computation Conference 2001*, 2001.
[61] K.-C. Jim and C. L. Giles, "Talking helps: Evolving communicating agents for the predator-prey pursuit problem," *Artificial Life*, vol. 6, pp. 237–254, 2001.
[62] 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*. MIT Press, 2002, pp. 302–311.
[63] J. C. Bongard, "Reducing collective behavioural complexity through heterogeneity," in *Artificial Life VII: Proceedings of the Seventh International Conference*, M. B. et al, Ed., 2000, pp. 327–336.
[64] X. Yao, "Evolving artificial neural networks," *Proceedings of the IEEE*, vol. 87, pp. 1423–1447, 1999.
[65] D. C. Queller, "Genetic relatedness in viscous populations," *Evolutionary Ecology*, vol. 8, pp. 70–73, 1994.
[66] S. A. West, M. G. Murray, C. A. Machado, A. S. Griffin, and E. A. Herre, "Testing Hamilton's rule with competition between relatives," *Nature*, vol. 409, pp. 510–513, 2001.
[67] W. D. Hamilton, "The Genetical Evolution of Social Behavior I+II," *Journal of Theoretical Biology*, vol. 7, pp. 1–52, 1964.
[68] L. Lehmann and L. Keller, "The evolution of cooperation and altruism: A general framework and a classification of models." *Journal of Evolutionary Biology*, vol. 19, pp. 1365–1376, 2006.
[69] A. Agogino and K. Tumer, "Efficient evaluation functions for multi-rover systems," *Lecture Notes in Computer Science*, vol. 3102, pp. 1–11, 2004.
[70] J. Grefenstette, "Credit assignment in rule discovery systems based on genetic algorithms," *Machine Learning*, vol. 3, no. 2, pp. 225–245, 1988.
[71] M. Minsky, "Steps toward artificial intelligence," *Proceedings of the Institute of Radio Engineers*, vol. 49, pp. 8–30, 1961, reprinted in Computers and Thought, McGraw-Hill, 1963.
[72] L. Panait and S. Luke, "Cooperative multi-agent learning: The state of the art," *Autonomous Agents and Multi-Agent Systems*, vol. 11, pp. 387–434, 2005, dOI 10.1007/s10458-005-2631-2.
[73] L. Keller, Ed., *Levels of Selection in Evolution*. Princeton University Press, 1999.
[74] S. West, A. Gardner, and A. Griffin, "Altruism," *Current Biology*, vol. 16, no. 13, pp. 482–483, 2006.
[75] W. D. Hamilton, "Altruism and related phenomena, mainly in social insects," *Annual Review of Ecology and Systematics*, vol. 3, pp. 193–232, 1972.
[76] T. Balch, "The impact of diversity on performance in multi-robot foraging," *International Conference On Autonomous Agents, Seattle, Washington, United States*, pp. 92–99, 1999.
[77] Y. U. Cao, A. S. Fukunage, and A. B. Kahng, "Cooperative mobile robotics: Antecedents and directions," in *Expanded Version of 1995 IEEE/RSJ IROS Conference Proceedings*, 1997, pp. 7–27.
[78] M. J. B. Krieger, J. B. Billeter, and L. Keller, "Ant-like task allocation and recruitment in cooperative robots," *Nature*, vol. 406, pp. 992–995, 2000.
[79] J. L. Reyes Lopez, "Optimal foraging in seed-harvester ants: Computer-aided simulation," *Ecology*, vol. 68, pp. 1630–1633, 1987.
[80] J. F. A. Traniello, "Foraging strategies of ants," *Annual Review Entomology*, vol. 34, pp. 191–210, 1989.
[81] G. Caprari, "Autonomous micro-robots: Applications and limitations," Ph.D. dissertation, EPFL, 2003.
[82] S. Magnenat and M. Waibel, "Enki, a fast physics-based 2D simulator," http://teem.epfl.ch, 2003-2007.
[83] S. Magnenat, M. Waibel, and A. Beyeler, "Teem, an open evolutionary framework," http://teem.epfl.ch, 2003-2007.
[84] R. McGill, J. Tukey, and W. Larsen, "Variations of box plots," *The American Statistician*, vol. 32, no. 1, pp. 12–16, 1978.
[85] T. Balch, "Behavioral diversity in learning robot teams," Ph.D. dissertation, Georgia Institute of Technology, 1998.
[86] M. Waibel, D. Floreano, S. Magnenat, and L. Keller, "Division of labour and colony efficiency in social insects: Effects of interactions between genetic architecture, colony kin structure and rate of perturbations." *Proceedings of the Royal Society B*, vol. 273, pp. 1815–23, 2006.
[87] M. A. Potter, L. A. Meeden, and A. C. Schultz, "Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists," in *Proceedings of The Seventeenth International Conference on Artificial Intelligence*, 2001, pp. 1337–1343.
[88] D. Tarapore, D. Floreano, and L. Keller, "Influence of the level of polyandry and genetic architecture on division of labour," in *The Tenth International Conference on the Simulation and Synthesis of Living Systems*, 2006, pp. 358–364.
[89] A. S. Griffin and S. A. West, "Kin selection: Fact and fiction," *Trends Ecol. Evol.*, vol. 17, no. 1, pp. 15–21, 2002.

Markus Waibel received a Masters in Physics from the Technical University of Vienna in 2003 and a Ph. D. from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland in 2007. He is interested in the evolution of cooperation in both natural and artificial systems. He coordinated the popular Talking Robots and Robots podcasts and is a regular contributor to the IEEE's robotics blog Automaton.

Laurent Keller is Professor of Ecology and Evolution, and Head of the Department of Ecology and Evolution, at the University of Lausanne. He is interested in the principles governing the evolution of animal societies and the ecological and evolutionary consequences of social life. As well as numerous research papers, he has edited the books Queen Number and Sociality in Insects (1993) and Levels of Selection in Evolution (1999). In 2005 he was awarded the E. O. Wilson Naturalist Award.

Dario Floreano (SM'05) is Associate Professor and Director of the Laboratory of Intelligent Systems in the School of Engineering of the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. He is interested in the analysis and synthesis of bio-inspired adaptive systems, such as bio-mimetic robotics, neuromorphic engineering, and artificial evolution. He published two books on those topics: Evolutionary Robotics in 2000 with Stefano Nolfi by MIT Press and Bio-inspired Artificial Intelligence in 2008 with Claudio Mattiussi by MIT Press too.