

Communication based dynamic role allocation in a group of homogeneous robots

Onofrio Gigliotta · Marco Mirolli · Stefano Nolfi

Published online: 10 July 2014
© Springer Science+Business Media Dordrecht 2014

Abstract The field of collective robotics has been raising increasing interest in the last few years. In the vast majority of works devoted to collective robotics all interacting robots play always the same function, while less attention has been paid to groups of collaborating robots in which different robots play different roles. In this paper we evolve a population of homogeneous robots for dynamically allocating roles through communicative interactions. In particular, we focus on the development of a team of robots in which one and only one individual (the 'leader') must differentiate its communicative behaviour from that of all the others ('non-leaders'). Evolved solutions prove to be very robust with respect to changes in the size of the group. Furthermore, both behavioural analyses and a comparison with a control condition in which robots are not allowed to move demonstrate the importance of co-adapting communicative and non-communicative behaviours, and, in particular, of being allowed to dynamically change the topology of communicative interactions. Finally, we show how the same method can be used for solving other kinds of role-allocation tasks. The general idea proposed in this paper might be used in the future for evolving general, robust, and scalable role differentiation mechanisms which can be exploited to develop non-communicative

collaborative behaviours that require specialisation of roles within groups of homogeneous individuals.

Keywords Evolutionary robotics · Dynamic role allocation · Neural controllers

1 Introduction

Adaptive technique such as artificial evolution in which a robot or a group of robots develop the ability to solve a certain problem autonomously in interaction with the environment (Harvey et al. 2005; Nolfi and Floreano (2000) have been effectively applied to the synthesis of collective behaviour in groups of robots (Baldassarre et al. 2003; Dorigo et al. 2004; Spector et al. 2005; Quinn et al. (2003). Indeed, these methods are particularly useful for developing collective behaviours in decentralized systems in which each robot act independently and has access to local information only (Nolfi 2006). In this case the global behaviour is the result of a self-organization process that emerges from the numerous interactions taking place between each robot and the environment and between the robots. Systems that rely on self-organization are highly desirable since they tend to be characterized by important properties such as robustness, flexibility and scalability.

Following Garnier et al. (2007) collective behaviours can be categorized in four types: coordinated, cooperative, collective decision-making and collaborative behaviours. Coordinated behaviour consists in the ability of a group of individuals to produce a specific spatio-temporal organization of their relative position and/or of the results of their activities which is functional with respect to a certain goal. Cooperative behaviours are those in which individuals must combine their efforts to solve a problem that goes

O. Gigliotta (✉)
Nac Lab, Department of Humanities, University of Naples
Federico II, Naples, Italy
e-mail: onofrio.gigliotta@unina.it

M. Mirolli · S. Nolfi
Institute of Cognitive Sciences and Technologies, Rome, Italy
e-mail: marco.mirolli@istc.cnr.it

S. Nolfi
e-mail: stefano.nolfi@istc.cnr.it

beyond their individual abilities. Collective decision making behaviours are behaviours in which a group of individuals faces several opportunities and collectively chooses the opportunity to maximize the performance with respect to a given problem. Finally, collaborative behaviours are behaviours in which different activities should be simultaneously performed by a group of specialized individuals.

For the moment, most of the work that has applied evolutionary robotics methods to the synthesis of collective behaviour has studied coordinated or cooperative behaviour in which the individuals do not need to differentiate their roles. A possible reason for this is that typically Evolutionary Robotics experiments use groups of homogeneous robots in order to avoid problems related to altruistic behaviours. In fact, if interacting agents are non-homogeneous, then the problem of altruism immediately arises, making the emergence of collective (cooperative) behaviours extremely difficult (two examples of works devoted to the problem of altruism in groups of communicating agents are Mirolli and Parisi 2005 and Floreano et al. 2007).

According to the categorization provided above, a group of collaborating robots constitutes what in the animal behavior literature has been defined a ‘team’ (Anderson and Franks 2001), that is a group of agents in which (1) different individuals make different contributions to the success of the task, (2) roles are interdependent thus requiring cooperation, and (3) organization persists over time. For the reason exposed above, the development collaborative behaviours in robot teams through Evolutionary Robotics techniques seems to pose a difficult challenge since it is not clear how homogeneous individuals might be able to assume different roles in a persistent manner.

In this paper we propose a possible solution to this problem which is based on endowing robots with communication capabilities, so that role allocation might be negotiated through the exchange of signals. In recent years several studies have demonstrated the possibility to evolve communication in homogeneous robots so to accomplish collective tasks (e.g. Di Paolo 2000; Marocco and Nolfi 2007; Quinn 2001). In this work we evolve groups of robot to differentiate themselves by simply sending different communication signals. The rationale behind this idea is that if we can evolve groups of homogeneous robots which are able to negotiate their roles through the exchange of signals, then this ability might be exploited in the future for the development of non-communicative collaborative behaviours requiring role specialization. More precisely, we will focus on the development of a team of robots in which one and only one individual should differentiate from the others by sending a different signal.

The obtained results illustrate how the method proposed allows evolving robots to solve the specialization problem

and to come up with solutions which are robust and scalable with respect to the number of interacting robots. In addition we demonstrate how one key aspect of the evolved solutions consists in exploiting the robot’s motor behaviour so to dynamically change the topology which determines who interact with whom, and we show how the proposed solution can be generalized to different kinds of role allocation tasks.

The rest of the paper is structured as follows. In Sect. 2 we describe the relation between the model presented in this paper with the state of the art. In Sect. 3 we present the experimental set-up. Section 4, which presents the results of our experiments, is divided in several subsections: in Sect. 4.1 we present the general results; in Sect. 4.2 we test the robustness of the best evolved solution with respect to the number of interacting robots; in Sect. 4.3 we perform a behavioural analysis of the best individual in order to understand the mechanisms underlying the evolved solution; in Sect. 4.4 we compare the results of our experiment with a control condition in which the robots are not allowed to move (i.e. are not allowed to modify the topology of the interactions); in Sect. 4.5 we demonstrate how the same method can be applied to solve a different role allocation task. Finally, in Sect. 5 we discuss the significance of the proposed method and of the obtained results and the implications for future works.

2 Related work

The issue of dynamical role allocation in evolving robots has already been tackled by Baldassarre et al. (2003), Quinn et al. (2003), Ampatzis et al. (2009) and Tuci et al. (2013). Baldassarre et al. (2003) evolved a group of robots for the ability to aggregate and collectively navigate toward a light target. The robots were equipped with infrared sensors, ambient light sensors, a speaker and directional microphones which were used to detect the intensity of the sounds produced by the other robots (a constant sound with fixed intensity). Evolved robots of different replications displayed three different families of strategies. The most effective of these strategies involved a collaborative behaviour in which different individuals played different roles (i.e. in which some individuals assume and maintain the frontal position with respect to the group and drive the group toward the light while other individuals assume the rear position with respect to the group and follow the individuals located at the front). Since the groups were formed by homogeneous individuals, and since robots’ controllers were formed by simple perceptrons and hence did not have any internal state, robot’s specialization was *situated*, in the sense that it completely depended on the different input patterns that robots received from the environment.

In a similar work, Quinn et al. (2003) evolved a team of three homogeneous robots able to dynamically allocating their roles in order to navigate as a group. In this experiment, robots equipment was really minimal: each robot had just four infrared sensors and two motor-driven wheels. As in the case of the previous experiment, the robots were asked to move together. The analysis of evolved robots' behaviour showed that the task was completed by relying on two phases: during the first phase robots organize themselves into a line formation, while in the second phase the robots start to move swinging clockwise and anticlockwise while maintaining their relative positions.

Ampatzis et al. (2009) evolved a team of two wheeled robots provided with a gripper to show a self-assembly behaviour. In particular, robots shared the same neural controller and the role allocation (two role were allowed: s-bot -gripper and s-bot -grippee) was accomplished by a dynamical physical interaction between the robots without an explicit communication channel.

In a more recent work, Tuci et al. (2013), evolved a group of robots to show two behaviours: foraging and nest patrolling. The authors explicitly rewarded the ability of robots to switch their role in different scenarios. In particular the nest and the foraging areas were indicated by a green and red light respectively. Furthermore the robots were not equipped with a specific communication system.

In the first three works, the robots were not directly evolved for the ability to assume different roles. Thus, the emergence of an ability to assume different roles (when present) can be explained by considering that it represents a prerequisite for the development of effective behaviour (i.e. collaborative behaviours). Moreover, in the above mentioned works, the authors did not analyse the robustness of the solutions involving role specialization with respect to the number of interacting robots (indeed, the solution found by the robots of Quinn and colleagues did demonstrate not to be robust to the lack of an individual: if one of the three robots is removed from the formation the remaining pair maintain the same configuration as when in full formation but ceasing the forward movement). In the present work we follow a complementary approach. We evolve the robots directly for the ability to assume different roles and we look to whether evolving robots come up with solutions which are robust with respect to the number of robots composing the group. This latter objective with a in deep analysis of the internal dynamics that subserves the role allocation behaviour though a specific communication system differentiates the current study from the work of Tuci and colleagues 2013. In fact, our objective is to come up with general, robust, and scalable role differentiation mechanisms which can be later exploited to develop collaborative behaviours.



Fig. 1 The experimental set up. *Left* the e-puck robot. *Right* the arena with ten simulated robots

3 Experimental setup

The experimental set-up, implemented using Evorobot* (Nolfi and Gigliotta 2010), consists in a group of ten identical E-puck (Fig. 1, left) robots placed in a box shaped arena of 600×600 mm (Fig. 1, right).

Robots can move in the arena by sending commands to their two wheels and can exchange signals between themselves through a dedicated communication channel. Communication is local as each robot perceives only the higher signal emitted by surrounding robots within the distance of 150 mm. Robots' signals are not only used for communication: they also represent the role of the signalling robot. We evolve our robots for their ability to differentiate their roles through the differentiation of their signals: one of the robots must maximise the value of its communicative output, while all other robots must minimize the values of their signals. (For simplicity, in the remainder of the paper we will call the individuals that send high value signals 'leader' and those that send low value signals as non-leaders.) More concretely, we calculate the fitness of a group of robot in the following way. For each cycle, we take the average of the differences between the communicative output of the current 'leader' (i.e. the robot with maximal communicative output) and the communicative outputs of all other robots. The fitness is the average of this value for all the cycles of all the trials. Formally, this is how fitness is calculated:

$$F = \frac{\sum_j^C \sum_i^N Max - O_i}{C(N - 1)}, \quad (1)$$

where N is the number of robots in the group (i.e. 10), C is the total number of life-cycles of each individual (i.e. 1,000 cycles times 40 trials = 40,000), Max is the value of the signal of the current leader and O_i is the value of the signal of robot i . Each robot is controlled by a neural network with a fixed architecture shown in Fig. 2. There are ten sensory units: 8 input units encode the state of the 8 infrared sensors, which are placed around the robot's body;

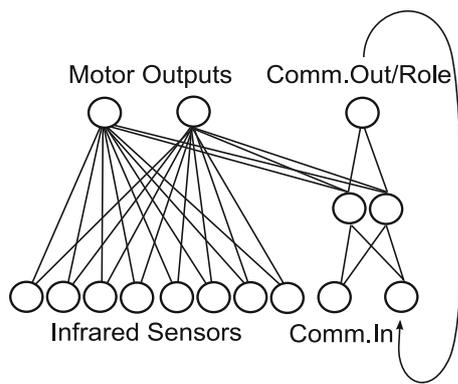


Fig. 2 The neural controller

one input unit encodes the higher signal emitted by surrounding robots; and the last input unit encodes the activation of the same robot's communicative output unit during the previous cycle. All sensory units (but the one encoding the same robot's communicative output) are noisy: in each cycle we add a random value uniformly distributed in the range $[-0.05, 0.05]$ to the estimated value of each sensor. All the infrared sensor units send connections to two motor units, while communication input units send connections to two hidden units, which are leaky integrators with evolvable time constant. These two hidden units send connections to both the communicative output unit and the two motor output units, which are used to control the motors of the two wheels.

The free parameters of the robots' neural controllers (i.e. the connection weights, the biases, and time constants of the two leaky neurons) are evolved (Nolfi and Floreano 2000). Each parameter is encoded as an 8 bits string, whose value is then uniformly mapped in the range $[5.0, +5.0]$ for weights and biases and in the range $[0, 1]$ for time constants. The initial population consists of 100 randomly generated genotypes. Each genotype is tested for 40 trials, lasting 1,000 cycles each. At the beginning of each trial, the genotype is translated into a corresponding neural controller which is duplicated ten times and embodied into ten robots (the group of robots is homogeneous). The robots are randomly placed into the arena and are left free to move and to communicate between each other for all the rest of the trail. After all teams of robots have been tested, the 20 best genotypes of each generation are allowed to reproduce by generating five copies each, with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasts 250 generations (i.e. the process of testing, selecting and reproducing robots is iterated 250 times). The experiment was replicated 10 times by starting with differently randomly generated genotypes.

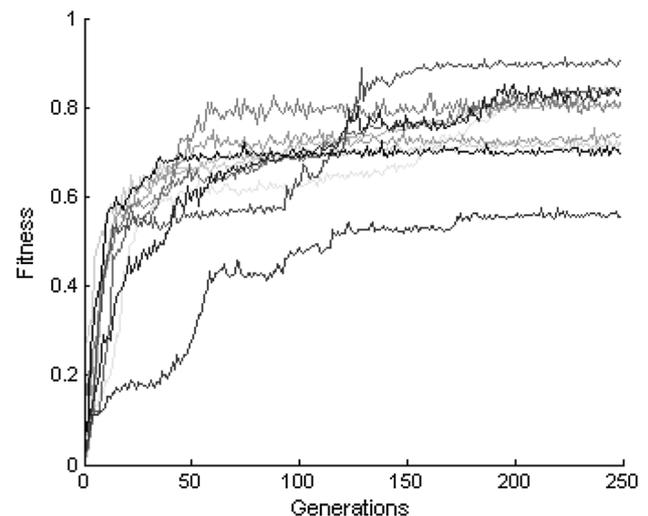


Fig. 3 Fitness of the best individuals of the ten replications of the experiment

4 Results and analyses

4.1 General results

By analyzing the obtained results we observed in five out of ten replications evolved robots display an ability to solve the role allocation task after 250 generations (Fig. 3). Since we were interested in investigating whether role differentiation could be evolved in groups of homogeneous robots, we discarded all the replications in which good role allocation did not emerge. Within the successful replications we observed different strategies, which can be divided in two types depending on whether the roles towards which individuals converge are either reversible or irreversible. In reversible strategies, the robots which assume a non-leader role can maintain their non-leader role or turn to a leader role depending on whether they keep interacting with leader robots or not. On the contrary, in irreversible strategies, once a robot has assumed a non-leader role, it remains a non-leader for the rest of the trial independently from its interactions with other robots. Both kinds of strategies have advantages and disadvantages, and they might be appropriate for supporting different kinds of non-communicative collective behaviors requiring specialization (see the discussion). However, for reasons of space in the rest of this paper we will focus our analyses only on the best evolved individual, which belongs to the irreversible strategy type.

4.2 Robustness

An important advantage of solutions based on distributed mechanisms consists in the fact that these solutions might

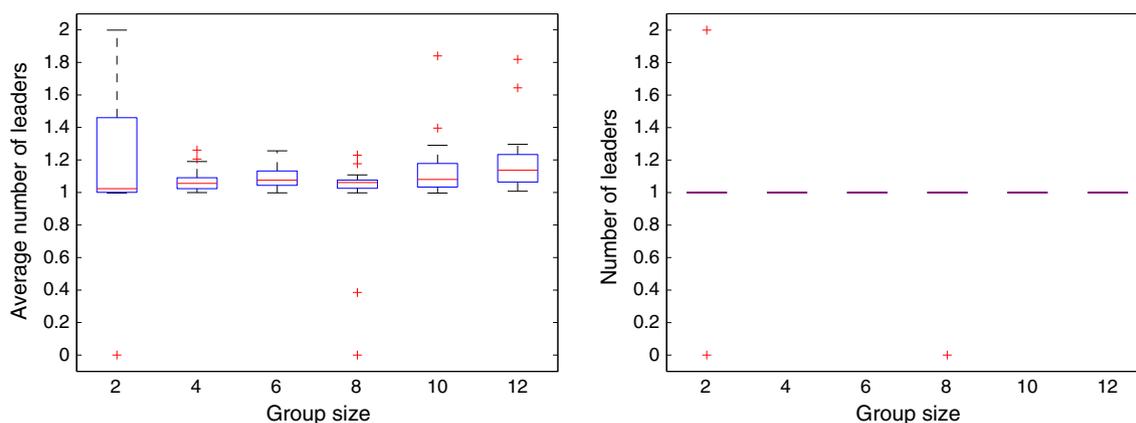


Fig. 4 Average number of leaders over 25 trials of 3000 cycles (*left*) and number of leaders recorded on the last lifecycle of 25 trials (*right*) for the best evolved controller embodied in groups of 2, 4, 6, 8, 10, and 12 robots

generalize with respect to the number of robots (Baldassarre et al. 2006; Trianni and Nolfi 2007). This generalization ability is important in collective robots since the number of robots available might vary due to variations of the environmental circumstances or due to the malfunctioning of some of the robots. In order to verify whether our method produced robust solutions we tested the best individual of the best evolutionary run using different group sizes: in particular, we tested groups of 2, 4, 6, 8, 10, and 12 robots (remember that during the evolutionary process groups of robots were always composed of 10 individuals). For each group size we made 25 trials, each of which lasted 3,000 life cycles. For each life cycle we recorded the number of leaders within the group. This latter measure was computed assigning the leader role to the robots with the role/communication output higher than 0.5 and the non-leader role to the robots with the role/communication output lower than 0.5. The average (over the 25 trials) number of leaders and the number of leaders on the last life cycle are shown in the boxplots reported in Fig. 4, left and 4, right, respectively.

The results reported in the figure show that evolved robots do generalise with respect to the number of robots. In fact, there are not significant difference between group size whether we consider the average number of leaders over all life cycles in tested trials (Friedman test, $p = 0,22493$), or the number of leaders in the last life cycle, where a convergence is supposed to be reached (Friedman test, $p = 0,072$). The fact that different number of robots in the same arena produce basically the same number of leaders demonstrates that, within reasonable limits, also the size of the arena does not matter much for role allocation, since by changing the number of robots it changes the crowding of the arena. On the other hand, if we increase the number of robots too much while maintaining the size of the arena constant, performance starts to decrease significantly just

because the arena becomes overcrowded and the robots are not able to move and meet each other in order to appropriately negotiate their roles. For this reason, in order to test whether the evolved solution was effective even with larger group sizes, we had to increase the size of the arena. In particular, we tested in larger arenas (see measures between parentheses) groups formed of 20 (800 × 800 mm), 30 (900 × 900 mm), and 40 (1,000 × 1,000 mm) agents. Results shown in Fig. 5, left indicate that notwithstanding increase in the arena size performance gradually decrease as the group size increases. In order to test whether the problem of crowding with many robots might be overcome not only through larger arenas but also through the extension of the communication range of the robots, we re-tested groups of 20, 30, 40, and 50 with a communication range of 400 mm instead of 150 mm. Figure 5, right shows how extending the communication range from 150 to 400 mm is able to improve the performance of the system.

4.3 Behavioural analysis

To understand the strategy used by the evolved robots to solve the role allocation task we tested the best individual with respect to both its communicative and non-communicative behaviour. With respect to communicative behaviour, we analysed a situation in which the team is composed by two robots initially located at a distance which is higher than their communicative range and facing toward each other. In this situation the robots move forward toward each other, thus reducing their relative distance up to the point in which they start to detect their respective communication outputs. The communicative output of the two robots during the first 50 cycles of this test is shown in Fig. 6, left. As can be seen from the figure, when the robots do not detect any signal, they quickly increase their communication output up to a value of 0.955

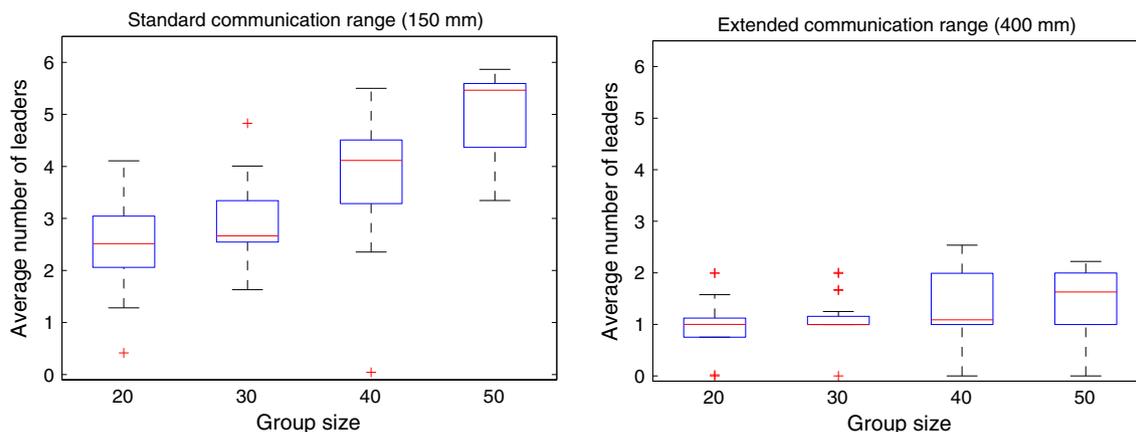


Fig. 5 Average number of leaders over 25 trials of 3,000 cycles recorded on for the best evolved controller embodied in groups of 20, 30, 40, and 50 robots. *Left* with a communication range of 150 mm. *Right* with an extend communication range of 400 mm

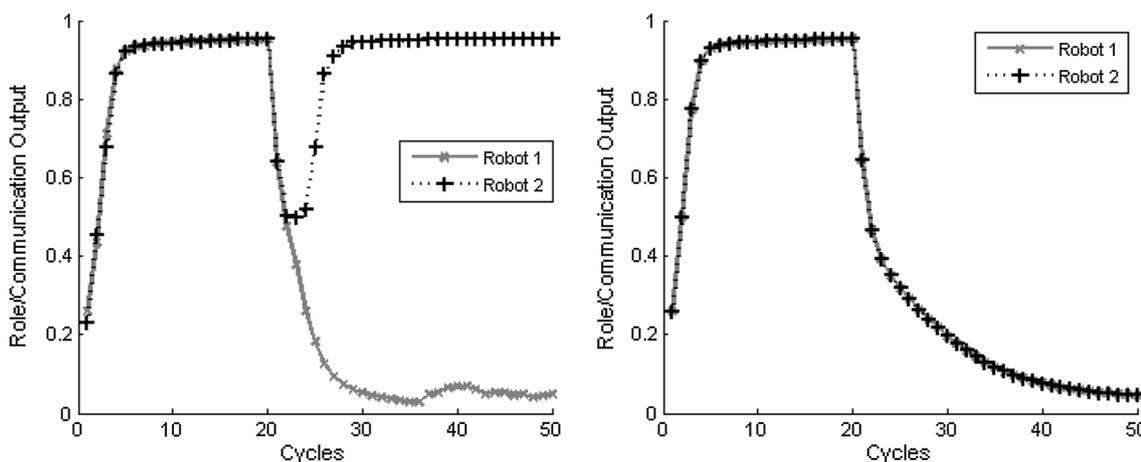


Fig. 6 Communicative outputs of two interacting robots over 50 life-cycles with (*left*) and without (*right*) noise in the communicative input channel

(i.e. they become leaders). Later on, when the two leader robots detect the signal of the other robot which indicates that also the other robot is a leader (cycle 19 of Fig. 6, left), they progressively and concurrently reduce their communication outputs until a point (cycle 22) in which the two communication outputs start to differentiate. As soon as the symmetry between the states of the two robots is broken, the difference between the two states tends to be amplified. As a consequence, the robot with the highest communication output progressively increases its communication output further thus taking the role of leader while the other robot keeps on decreasing its communication output further up to a value of 0.036 thus becoming a non-leader. In order to break the symmetry between the two communication outputs the two robots exploit the noise on the sensory state (which is simulated by adding a random uniformly distributed noise in the range $[-0.05, 0.05]$ to the estimated value of the sensor). Indeed, by

repeating the experiment in a test in which the noise has been removed (Fig. 6, right), the robots fail to differentiate their state and both converge on the non-leader role. Overall, the differentiation mechanism can be explained by the exploitation of (a) noise, and (b) a bifurcation point in the communication output (and in the internal state) of the robot located between the two attractors points corresponding to a leader and non-leader role and corresponding to a value of about 0.5.

What about non-communicative behaviour? Is it affected by the role assumed by a robot? Since the two hidden units which determine communicative behaviour project also to the two motor units controlling the two robot wheels it is likely that the role assumed by a robot will modulate also non-communicative behaviour. This is indeed the case: if left free to move all alone in the arena leader robots tend to move straight and just perform obstacle avoidance when encountering walls, while non-leader robots tend to

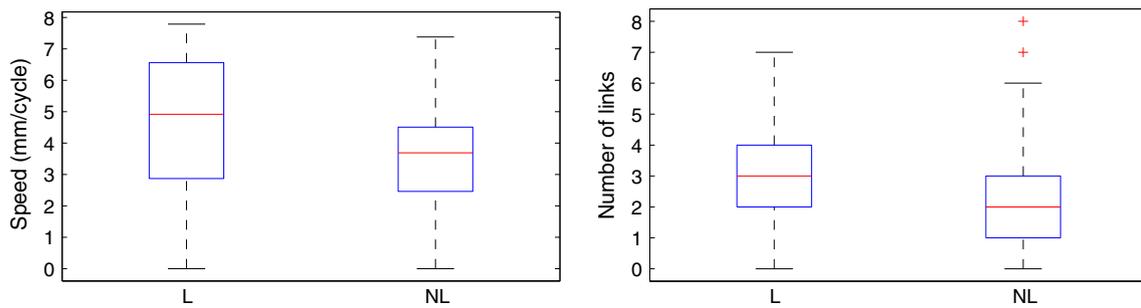


Fig. 7 Average speed (*left*) and number of links (*right*) for Leaders (L) and Non-Leaders (NL). Results are averaged over 100 test trials lasting 3,000 cycles each

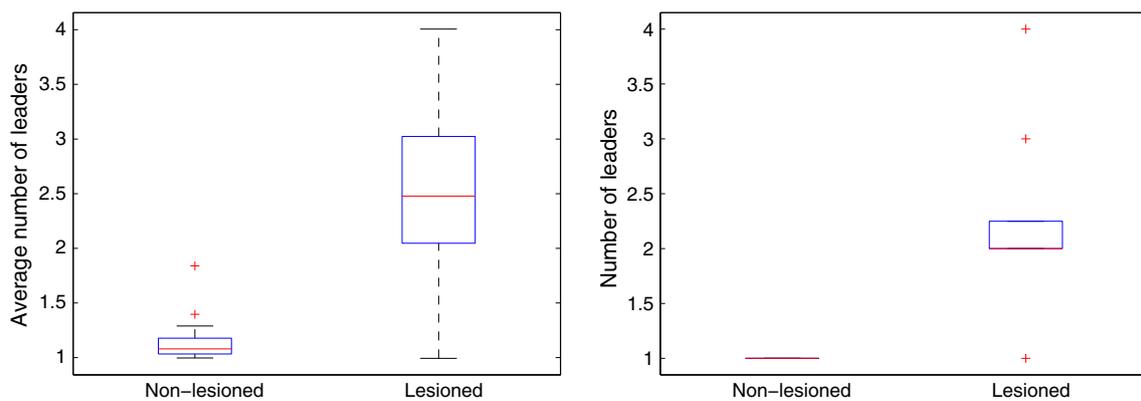


Fig. 8 Average number of leaders over ten trials during the 3,000 cycles of a trial (*left*) and on the last cycle of each trial (*right*) for the best evolved robots (non-lesioned) and robots with the same neural

network but with the links connecting the internal units to the motor units lesioned (i.e. set to 0)

have a wide circular motor behaviour. But is this difference also present in robot’s *ecological* conditions (i.e. when the arena is crowded by ten interacting robots)? If so, what, if any, is the adaptive function of this behavioural differentiation? In order to answer these questions we re-run 100 test trials lasting 3,000 cycles each with groups of 10 robots while measuring, for each cycle and for each robot, three values: the communicative output, the instantaneous speed, and the number of interaction links (i.e. the number of robots within the communication range). Figure 7 shows the average speed (left) and number of links (right) for robots with a communicative output higher than 0.5 (leaders) and lower than 0.5 (non-leaders). It turns out that even in ecological conditions leaders tend to move faster than non-leaders, and this has the result of slightly increasing their average number of communicative interactions (though quantitatively small, both the differences are statistically significant—Mann–Whitney U , $p < 0.001$ —and correspond to the qualitatively very different behaviours explained above).

In order to ascertain whether this behavioural difference is epiphenomenal or has an adaptive function, we ran other two tests of ten trials of 3,000 cycles each comparing the

performance (number of leaders) in two conditions: ten robots with the best evolved controller versus ten robots with the same controller but with the 4 connections linking the two internal units with the motor output units lesioned (i.e. set to 0). Hence, in the latter condition the role that the robot has assumed can not influence the robot non-communicative behaviour. In this condition all the robots (both leaders and non-leaders) produce the wide circular motor behaviours produced by non-leaders in normal conditions. Consequently, it is harder for leaders to meet each other, resulting in a significantly higher number of leaders in the group (Fig. 8, in both the comparisons, Mann–Whitney U , $p < 10^{-7}$).

The results presented in this section demonstrate that evolved robots have co-adapted their communicative and non-communicative behaviour so to maximise the efficiency of the role allocation task: on the one hand, the robots allocate their roles on the basis of their encounters due to their motor behaviour; on the other hand, their roles influence their motor behaviour so maximise the probability that two leaders meet each other and re-negotiate their roles so that only one leader remain in the group as required by the fitness function.

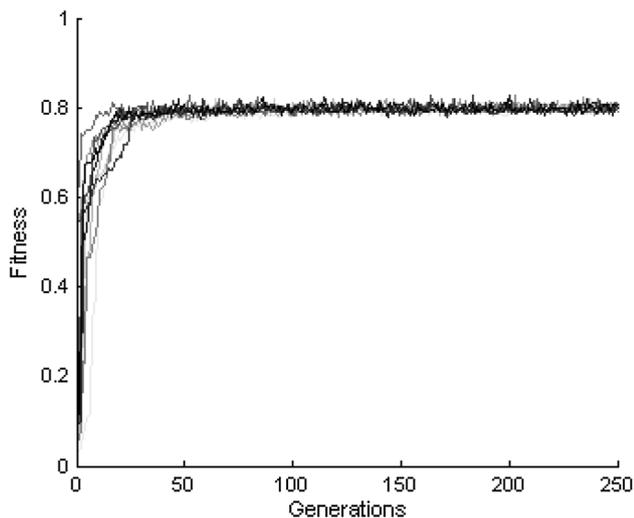


Fig. 9 Fitness of the best individuals of the ten replications of the experiment with fixed topology

4.4 Role of topology

In the previous section we showed that evolved robots not only move in the environment, but also differentiate their non-communicative behaviour according to their role so to maximise their fitness. In this section we investigate whether the possibility to move so to change the topology of interactions is really necessary for solving the role-allocation task. To do that, we run ten replications of a new evolutionary experiment in which the ten interacting robots have a configuration which is fixed throughout the whole trial, in that robots are not allowed to move. For each trial, the fixed configuration is chosen randomly by placing the first robot at the centre of the arena and then positioning the other robots one by one with the constraint that each newly introduced robot must be within the communication range of at least one of the robots which are already present in the arena (this constrain has been introduced in order to assure that in the resulting topology there is never a robot or a group of robots which is completely disconnected from the others, which would have prevented the possibility of finding a solution to the role-allocation task). Figure 9 shows the fitness of the best individual of each generation for the ten replications of this new experiment. In this condition all the evolutionary processes converge in less than 50 generations to a stable value of about 0.8, which is significantly lower than the best solution found in the basic experiment (about 0.9). The fact that in this condition evolution is much faster and more reliable than in the normal case is not surprising since in the fixed-topology condition robots have only to evolve communicative behavior, while in the normal condition they have also to

evolve motor behaviours such as obstacle avoidance and exploring behavior and to co-adapt them with their communicative behavior.

In order to compare the best solutions of the two conditions (the normal condition and the one with fixed topology) we performed the robustness test presented in Sect. 4.2 with the best evolved individual of the fixed topology condition: 25 trials of 3,000 cycles each for six different group sizes (2, 4, 6, 8, 10, and 12). The results of these tests, shown in Fig. 10, clearly demonstrate the advantage of having the possibility to co-adapt communicative and non-communicative behavior (all the difference between group of the same size resulted statistically significant Mann–Whitney U , $p < 0.01$; only in the case of two robots the fixed topology presents better results). Since they cannot rely on dynamically changing the topology of the interactions through displacing in the environment, robots of the fixed topology condition cannot but rapidly converge on suboptimal solutions. In fact, with groups composed by 4 to 12 robots performance is not only suboptimal (Fig. 10, right), but the groups do not converge on having just one leader. On the contrary, the higher is the number of interacting robots, the higher is the average number of leaders of the groups (Fig. 10, left). Only in groups composed by just two robots the fixed-topology solution stably converges on having only one leader and outperforms the condition in which robots are free to move. And this is not a problem of robustness: even in the condition for which robots have been evolved (i.e. groups of 10) the fixed-topology condition produce groups of robots that have on average about 3 leaders instead of 1. The difference between groups of different size in the fixed topology condition (reported in Fig. 10) thus is widely significant (Friedman test, $p = 0.00000$ for the average number of leaders and number of leaders in the last lifecycle).

4.5 Task generalization

So far we have demonstrated that the approach proposed in this paper of evolving groups for robots for dynamically allocating roles through local communicative interactions is feasible and promising at least for the case in which one of the robots must assume the role of the leader and all the other robots must assume the role of non-leaders. But since nothing in this approach is strictly related to the emergence of a (single) leader, the same idea might be exploited also for developing robots able to dynamically allocate different kinds and numbers of roles. In order to demonstrate this, we applied the same method for evolving groups which have to split in two evenly distributed subgroups, with half of the robots sending high-value signals and the other half

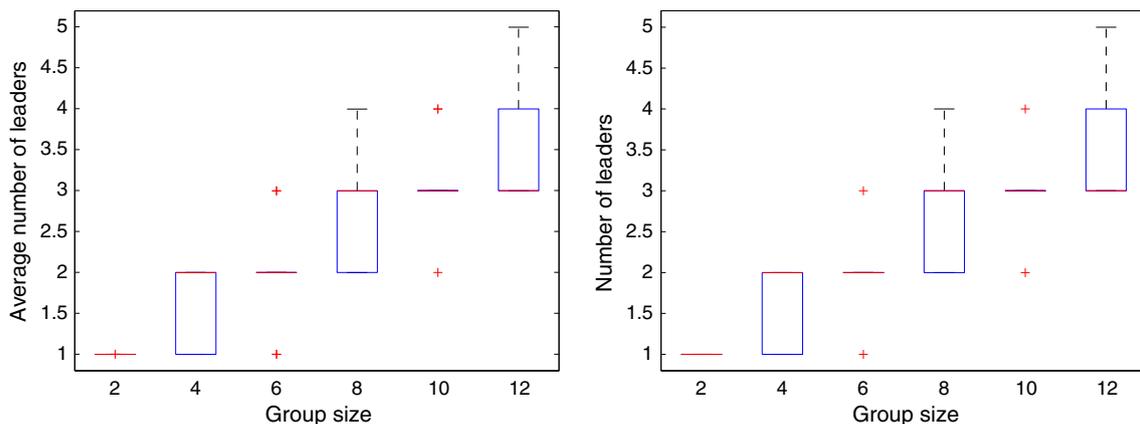


Fig. 10 Average number of leaders over 25 trials of 3,000 cycles (right) and number of leaders recorded on the last lifecycle of 25 trials (left) for the best evolved controller with fixed topology embodied in groups of 2, 4, 6, 8, 10, and 12 robots

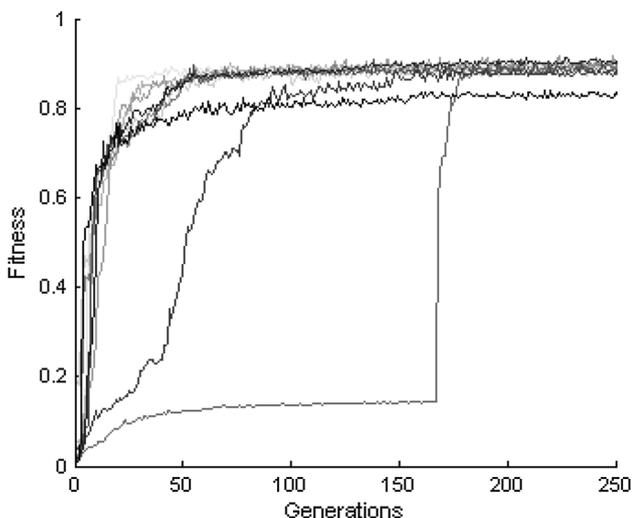


Fig. 11 Fitness of the best individuals of the ten replications of the experiment in which robots are evolved for their ability to split in two sub-groups. See text for details

sending low value signals. The new experimental set-up is identical to the previous one (described in Sect. 3) but for the fitness function, which is calculated as follows. During each cycle, we sort robots by their role/communication output activation from the highest to the lowest. After that, we split the group of N robots in two sub-groups composed of $N/2$ robots each, according to the robots’ communicative output (the 50 % with the higher and the 50% with the lower output value). Then, we compute the absolute difference between the means of the communicative output within the two groups. The fitness value is the average of this difference over the total number of life-cycles (i.e. 1,000 cycles times 40 trials = 40,000). This is the fitness formula:

$$F = \frac{\sum_j^C \left| \sum_{i=1}^{N/2} \frac{O_i}{N/2} - \sum_{i=N/2+1}^N \frac{O_i}{N/2} \right|}{C}, \tag{2}$$

where C is the total number of cycles, N is the number of robots and O_i is the communicative output of robot i . Figure 11 shows the fitness of the best individual of each generation for the ten replications of the same experiment with different random initial conditions. As it is clear from the figure, evolution is able to find very effective solutions also to this new dynamic role allocation task. Indeed, this task seems to be much easier to solve than the previous one since nine out of ten evolutionary runs have reached a stable fitness value of about 0.9 (which is the fitness reached by the best replication of the previous experiment), and the worst seed reached a fitness of 0.8.

5 Discussion

In this paper we presented evolutionary robotics experiments in which groups of homogeneous robots are evolved for their ability to dynamically allocate their roles through their communicative and non-communicative interactions. Each individual robot is an independent entity which only has access to local information. As a consequence, the allocation of the different roles is the results of a dynamical process in which the robots negotiate their roles on the fly (for alternative approaches in which the role allocation process is based on global information shared by the individuals see, Mataric 1995 and Stone and Veloso 1999). In particular, we analyzed the case in which a single robot should assume the role of the leader (i.e. should turn their communication output on) and all other robots should assume the role of non-leader (i.e. should turn their communication output off). The fact that the robots are

homogeneous implies that each individual can potentially assume any possible role.

The analysis of the obtained results indicates that the method proposed leads to the development of effective solutions which scale well with respect to the number of interacting robots. In fact, the best evolved solution demonstrated to be very robust to changes in the number of robots composing the group, which might be higher or lower with respect to the number of robots experienced during the evolutionary learning process. This aspect is particularly important in the case of mobile robots in which the number of interacting robots might vary depending on the robot relative position.

A detailed behavioural analysis showed that evolved solutions are based on the combination of two dynamical mechanisms which (1) regulate how the state encoding the role of a robot changes while the robot interacts with other individuals located nearby, and (2) regulate the network of interacting robots (i.e. who interact with whom).

The first mechanism concerns the communicative behaviour of the robots. When a robot does not interact with other robots, it tends to increase its communicative output so to progressively assume the role of leader. However, when two leader robots interact, they initiate a mutual inhibition phase which finally ends in a situation in which one robot acts as a leader and the other robot acts as a non-leader. This result is achieved through a dynamical process involving three phases: (a) an initial phase in which the two robots inhibit each other so to reach an intermediate state representing a bifurcation point in robots' internal dynamics (i.e. a situation from which the state of the robot can easily move toward a leader or non-leader role corresponding to a fully activated or fully deactivated state); (b) a transition phase in which the symmetry between the states of the two robots is broken by exploiting small differences in robots' sensors state due to noise; and (c) a final phase in which these small differences are progressively amplified thanks to the mutual interactions between the robots and a positive feedback mechanism which tends to reduce the discrepancy between the current state of a robot and one of the two fixed attractor states corresponding to the fully activated and the fully deactivated states.

The second mechanism concerns the motor behaviours of the robots which influence the topology of who interacts with whom. Despite the selection criterion does not directly reward the robots for their ability to move and modify their interaction topology, evolved robots do indeed exploit this possibility in order to increase their ability to efficiently differentiate their roles. By comparing the results obtained in the standard condition with a control experiment in which the robots are not allowed to move, we observed that the robots in the former condition

significantly outperform the robots in the latter condition. Moreover, we observed that evolved solutions also exploit the possibility to regulate the robots motor behaviour depending on the current robots role to increase their ability to solve the role allocation problem. In fact, leader robots tend to move more than non-leader robots in order to increase the probability that two leader robots meet and so to reduce the risk that the same group include more than one leader.

Finally, by carrying out another experiment in which exactly the same methodology was applied to the solution of a different role allocation task we also demonstrated how the proposed methodology generalizes to different types of problems.

Since our ultimate goal is to study the emergence of teams of robots in which different individuals play different roles and collaborate for the solution of a collective task, the principal line for future research consists in finding a way to exploit the dynamic role allocation system presented in this paper for the accomplishment of non-communicative collaborative tasks, in particular, tasks which require the presence of a leader. In fact, while several swarm-like behaviours might be successfully accomplished by groups of robots without any significant distinction between the behaviours of the members of the group, there are many cases in which the presence of a leader might significantly improve the performance of the group (see Anderson and Franks 2001 for examples in the animal kingdom and Baldassarre et al. 2003; Gigliotta et al. 2007 for examples within the artificial life community).

The solution proposed in the present paper seems promising in this respect: since role differentiation depends on communicative behaviour and not on different robots having different control systems, we can solve two potentially conflicting problems. On the one hand, once the robots have assumed different roles through the signals they emit, they can differentiate also their non-communicative behaviour depending on their roles in order to accomplish a non-communicative task that requires role specialisation (in fact, the results presented in Sect. 4.3 show that a differentiation in non-communicative behaviours already happens in the present simulations because it helps the maintenance of a single leader). On the other hand, the fact that interacting robots share the same genes (are clones) assures that the problem of altruism cannot raise and that robots with different specialisations would behave collaboratively so to maximise their fitness.

In order to develop groups of homogeneous robots able to efficiently accomplish collaborative tasks requiring role specialisation (in particular the presence of a leader) we envisage three possibilities. The first possibility is to seed the task-specific evolutionary search with our evolved robots, which are already able to dynamically negotiate

their roles. In this respect it is interesting to note that our evolutionary experiments resulted in two types of strategies. The first strategy is characterized by robots which converge toward a certain role allocation state which cannot be modified further. In particular, in this case robots which became non-leaders will never become leaders later on. The second strategy is characterized by robots which tend to converge toward a certain role but which can always change their role later on when the appropriate conditions are met. In this case, a non-leader robot which does not meet a leader for a certain amount of time will always tend to assume the leader role. These two kinds of basic strategies might be useful for different kinds of collaborative tasks. For example, if a task requires that one and only one robot (the leader) have to leave the group while all the other robots have to remain together, then the first strategy is to be preferred since non-leader robots will never change their role and leave the group. On the other hand, if the task is such that if the leader robot leaves the group it must be replaced (or if the role functionality is context-dependent and hence roles need to be re-arranged if the context changes, as in Baldassarre et al. 2003, Quinn et al. 2003), then the second, more flexible, role-allocation strategy will probably be more effective.

The second possible way of using the ideas proposed in this paper for designing teams of robots able to accomplish non-communicative collaborative tasks consists in designing multi-objective fitness functions in which one component of the fitness is specific to the collaborative task while a second component rewards the robots for their ability to dynamically allocate their roles. This might lead to the co-evolution of behavioural skills and rule allocation mechanisms which suit each other while maintaining the possibility to force the evolutionary process to develop a suitable role allocation mechanism toward a direct reward.

A third possibility which will release the designer from the need to specify the number and the distribution of roles within the group would consist in using information theoretic measures such as Shannon entropy (Shannon 1948) to reward robots for assuming different roles in different contexts while leaving the robots free to determine the right number and distribution of roles within the group. Finally, aside theoretical aspects and last but not least, in future work we will port the simulated setting into real robots by using a specific turret developed for the e-puck robots. In particular we will implement our communication system by using the e-RandB turret developed by Robolabo, Iridia and RBZ Robot Design. Robots equipped with this turret are able to exchange data within a specified range through an infrared protocol that allow them to exchange messages containing each a 16bit information packet. A capability that is pretty sufficient for our needs.

References

- Ampatzis C, Tuci E, Trianni V, Christensen AL, Dorigo M (2009) Evolving self-assembly in autonomous homogeneous robots: experiments with two physical robots. *Artif Life* 15(4): 465–484
- Anderson C, Franks NR (2001) Teams in animal societies. *Behav Ecol* 12(5):534–540. doi:10.1093/be-hecco/12.5.534
- Baldassarre G, Nolfi S, Parisi D (2003) Evolving mobile robots able to display collective behaviors. *Artif Life* 9(3):255–267. doi:10.1162/10645460322392460
- Baldassarre G, Parisi D, Nolfi S (2006) Distributed coordination of simulated robots based on self-organization. *Artif Life* 12(3):289–311. doi:10.1162/artl.2006.12.3.289
- Di Paolo EA (2000) Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adapt Behav* 8(1):27–48
- Dorigo M, Trianni V, Sahin E, Gross R, Labella TH, Baldassarre G, Gambardella LM (2004) Evolving self-organizing behaviors for a swarmbot. *Auton Robots* 17(2–3):223–245
- Floreano D, Mitri S, Magnenat S, Keller L (2007) Evolutionary conditions for the emergence of communication in robots. *Curr Biol* 17(6):514–519
- Garnier S, Gautrais J, Theraulaz G (2007) The biological principles of swarm intelligence. *Swarm Intell* 1:3–31. doi:10.1007/s11721-007-0004-y
- Gigliotta O, Miglino O, Parisi D (2007) Groups of agents with a leader. *J Artif Soc Soc Simul* 10(4):1
- Harvey I, Paolo ED, Wood R, Quinn M, Tuci E (2005) Evolutionary robotics: a 15 new scientific tool for studying cognition. *Artif Life* 11(1–2):79–98. doi:10.1162/1064546053278991
- Marocco D, Nolfi S (2007) Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connect Sci* 19:53–74. doi:10.1080/09540090601015067
- Mataric MJ (1995) Designing and understanding adaptive group behavior. *Adapt Behav* 4:51–80
- Mirolli M, Parisi D (2005) How we can explain the emergence of a language which benefits the hearer but not the speaker? *Connect Sci* 17(34):325–341
- Nolfi S (2006) Behaviour as a complex adaptive system: on the role of self-organization in the development of individual and collective behaviour. *ComplexUs* 2(3–4):195–203
- Nolfi S, Floreano D (2000) *Evolutionary robotics*. MIT Press, Cambridge
- Nolfi S, Gigliotta O (2010) Evorobot*. In: Nolfi S, Mirolli M (eds) *Evolution of communication and language in embodied agents*. Springer, New York
- Quinn M (2001) Evolving communication without dedicated communication channels. In Kelemen J, Sosik P (eds.) *Advances in artificial life: sixth european conference on artificial life (ecal'01)*. Springer, Berlin. vol 2159, pp 357–366. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.28.5890>
- Quinn M, Smith L, Mayley G, Husbands P (2003) Evolving controllers for a homogeneous system of physical robots: structured cooperation with minimal sensors. *Philos Trans R Soc Lond Ser A* 361:2321–2344
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27(379–423):623–656
- Spector L, Klein J, Perry C, Feinstein M (2005) Emergence of collective behavior in evolving populations of flying agents. *Genet Program Evolv Mach* 6:111–125. doi:10.1007/s10710-005-7620-3

- Stone P, Veloso M (1999) Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artif Intell* 110(2):241–273
- Trianni V, Nolfi S. (2007) Minimal communication strategies for self-organising synchronisation behaviours. In: *Proceedings of the first IEEE symposium on artificial life*, pp 199–206
- Tuci E, Mitavskiy B, Benedettini S, Francesca G (2013) On the evolution of self-organised role-allocation and role-switching behaviour in swarm robotics: a case study. *Adv Artif Life Ecol* 12:379–386