Evolution of communication-based collaborative behavior in homogeneous robots

Onofrio Gigliotta[1] and Marco Mirolli[2]

[1]Natural and Artificial Cognition Lab, University of Naples Federico II, Napoli, Italy
[2]Institute of Cognitive Sciences and Technologies, CNR, Rome, Italy
onofrio.gigliotta@unina.it, marco.mirolli@istc.cnr.it

Abstract

In the field of collective robotics much research has been devoted to the study of coordinated and cooperative behaviours where typically all the robots play the same function. Much less attention has been devoted to the development of groups of robots that play different roles (robot teams), probably because evolutionary collective robotics tend to use groups of homogeneous robots, in which role differentiation poses difficult challenges. In the few Evolutionary robotics studies in which role differentiation has been demonstrated such differentiation depends exclusively on the robots physical interactions, making the solutions found by evolution quite fragile, in particular with respect to the number of robots that form the group. In this paper we apply a method for role differentiation developed in previous work to the evolution of teams of homogeneous robots in which role differentiation is based on a dedicated communication channel. Our evolved robots are able to negotiate their role through communication and perform very effectively their collaborative task, which requires that one robot is sent to a 'mission' away from the group while all other robots remain in a 'home'. Our simulations also show that the method proposed, based on the rewarding of communication-based role differentiation, is necessary for the evolution of the desired behaviour. Finally, we show that since role differentiation is based on communication and not only on robot physical interactions, evolved solutions are considerably robust with respect to the number of robots composing the group.

Introduction

Recently the study of collective robotics has been raising increasing interest in the Artificial Life community. In particular, Evolutionary Robotics techniques (Nolfi and Floreano, 2000; Harvey et al., 2005) have proven extremely successful to develop collective behaviours in group of robots (see, e.g. Baldassarre et al., 2003; Dorigo et al., 2004; Spector et al., 2005; Quinn et al., 2003; Brambilla et al., 2013). Evolutionary Robotics is a design methodology in which populations of robots are evolved for a specific task, defined by a *fitness function*, for a number of generations, until a satisfying solution to the task is found. The parameters that determine the behaviour of the robots (for example the connection weights of their neural networks) are encoded in their genomes and are selected on the basis only of the overall performance of the robots, defined by the fitness function. In this way, the evolutionary process can exploit the full power of self-organisation in which behaviours emerge from the interactions between the robots and their environment and, in the case of collective robotics, among the robots themselves (Nolfi, 1998).

According to Garnier et al. (2007) collective behaviours can be categorised in four kinds: 1) coordinated behaviours, where the group produces a specific spatio-temporal organization of the relative positions of its members and/or of the results of their activities which is functional with respect to a certain goal; 2) cooperative behaviours, where individuals combine their efforts to solve a problem that no one individual can solve by itself; 3) collective decision making, where a group of individuals collectively chooses one among several opportunities in order to maximize the performance with respect to a given problem; 4) collaborative behaviours, where a collective goal is achieved through different activities that are simultaneously performed by different individuals.

Most of the work that has used evolutionary robotics methods for developing collective behaviours has been devoted to coordinated or cooperative behaviours, in which the individuals do not need to differentiate their roles. This is probably due to the fact that evolutionary collective robotics typically works with groups of homogeneous robots, that is groups of robots with identical genomes and hence identical control systems. The reason is that interesting collective tasks tend to require cooperativeness and if the interacting agents are non-homogeneous the emergence of cooperation is extremely difficult due to the problem of altruism (Mirolli and Parisi, 2005 and Floreano et al., 2007 are two examples of works dedicated to the problem of altruism in groups of communicating agents). On the other hand, the use of homogenous robots makes it difficult the development of robot specialisation.

In this paper we are interested in the development of collective behaviours, where interactive robots form a *team*: i.e. (1) different individuals make different contributions to the

task, (2) cooperation is required as roles are interdependent, and (3) the group organisation persists over time (see the definition of team provided by Anderson and Frank, 2001 in the context of animal behavior).

Related work

A few previous works have been devoted to the evolution of role differentiation in groups of robots (Baldassarre et al., 2003; Quinn et al., 2003; Tuci et al., 2013). Baldassarre et al. (2003) evolved a group of robots for the ability to collectively navigate toward a light target. Beyond proximity and light sensors, each robot had a speaker producing a constant fixed sound and directional microphones for detecting the sound produced by the other robots. Evolved robots displayed different kinds of strategies, the most effective of which involved role specialisation: in particular, one individual tended to assume the frontal position and drive the others towards the light while the others stayed in the back and followed the 'leader'. Since the robots were homogeneous and their controllers were single layer perceptrons without any internal state, robot specialisation was completely situated, that is completely dependent on the different sensory stimulation that different robots received from the environment.

In a similar work, Quinn et al. (2003) evolved a team of three homogeneous robots for the ability to move together in the environment. The peculiarity of this experiment was the minimalism of robot equipment: each robot had just four infrared sensors and two motor-driven wheels. Evolved robots were able to navigate together by relying on a strategy that had two phases: in the first phase the robots negotiated their roles until they reached a line formation; in the second phase they started to move by swinging clockwise and anti-clockwise while maintaining their relative positions.

Finally, Tuci et al. (2013) evolved a group of homogeneous robots for their ability to perform two behaviours: some robots had to remain and patrol the 'nest', while other robots had to leave the nest and go foraging in a food area. Furthermore, the role allocation strategy depended on the context: in one context the number of robot patrolling the nest had to be higher than the number of robots foraging, while in a second context the opposite holded.

Notwithstanding the importance of these works, they developed solutions that are not general but task-specific: the dynamic role allocation demonstrated by these robots can be used only for the tasks for which the robots were evolved and cannot be exploited for any other purposes. Furthermore, those solution do not seem to have the robustness that is typically assured by the use of homogeneous robots. As correctly discussed by Quinn and colleagues, the use of homogeneous robots has the potential to assure high robustness: since all robots are identical and hence equally able to play any role, teams of homogeneous robots should be in principle able to cope with the loss of individual members

or even to the addition of more robots. However, this advantage is not demonstrated in any of the works discussed above: Baldassarre et al. did not touch the problem of robustness, and it is not clear whether their robots' flocking behaviour would generalise with respect to the number of robots in the group; Quinn et al. report that their robots were not able to cope with the lack of one individual; Tuci et al. report that their evolved strategies "seem quite fragile with respect to various sources of variability, such as the cardinality of the team...". A possible reason for this fragility is that the evolution of task-specific role allocation will tend to produce task-specific solutions which deeply rely on the specific conditions under which evolution took place.

In order to solve this problem, in previous work (Gigliotta et al., 2009) we have proposed a simple mechanism for evolving dynamic role allocation in groups of robots that is based on a dedicated communication channel. In particular, robots were evolved for producing different communication signals: for example, one task consisted in having one robot (the 'leader') producing one signal and all other robots (the 'followers') producing another one (another task was to have half of the robots producing one signals and the other half the other signal). We found that evolved solutions not only involved the co-adaptation of communicative and non-communicative behaviours (Nolfi, 2005; Mirolli and Nolfi, 2010), but, most importantly, that they were very robust with respect to the number of robots involved. However, the limit of that work was that role-allocation was the only behaviour that the robots were required to show: in other words, fitness depended only on the signals emitted by the robot and role differentiation had no other adaptive function.

In this paper we apply the idea of communication-based dynamic role allocation to a simple task in which one robot (the 'leader') has to leave the group to perform its own task, while all other robots (the 'followers') have to remain together in their home area. We compare a condition in which the fitness function rewards only the non-communicative behaviour with one in which we add to this fitness function a component, inherited by our previous work, which also rewards communicative role-allocation. In this way we demonstrate that our idea of communication-based role-allocation can indeed be exploited to develop non-communicative team behaviour in homogeneous robots. Furthermore, since this kind of role allocation does not depend on the task nor on the number of robots, we also show that the resulting behaviour is robust with respect to the number of robots constituting the team.

Experimental set-up

Robot and task

We want to demonstrate that a group of homogeneous robots can dynamically and autonomously learn to allocate different roles (i.e. 'leader' and 'followers') and behave according to these roles. Furthermore, we want to demonstrate that the
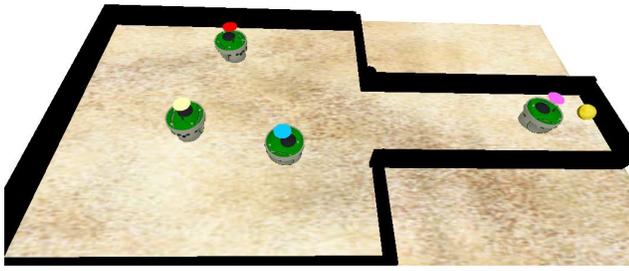
Figure 1: The set-up

Figure 2: The neural controller

robot form a real team, where the group organisation persists over time (see above) and is not strictly dependent on the continuous physical interactions between the robots. For these reasons the metaphor that guided us for developing our experimental set-up was that of a group of robots which had to send one and only one robot away from the group in order to accomplish a 'mission' (e.g. finding resources, exploring some other place, rescuing somebody...).

The experimental set-up (fig. 1) consists in four simulated e-puck robots placed in an arena formed by a 'home' square (side = 600 mm) and a rectangular corridor (500 x 200 mm). At the end of the corridor there is a small light (yellow sphere in fig 1). The task for our group of robots is to send one of the robots into the corridor to get close to the light while all other robots remain in the home area.

Neural controller

The controller of the robots is the neural network depicted in fig 2. The sensory system is composed by: 8 infrared sensors, placed around the robot's body, which provide a noisy indication of the proximity of an object (another robot or a wall); 8 light sensors, also placed around the robot's body, which provide a noisy indication of the proximity of the light (if within sight, i.e. at about 300 mm); one communication sensor, which encodes the value of the highest signal emitted by the other robots (within a distance of 1000 mm, meaning that basically each robot can hear each other within the arena); and a unit which encodes the signal emitted by the same robot in the previous time step (for the importance of talking-to-oneself in artificial life and robotics, see Mirolli and Parisi, 2005, 2011). All sensory units are directly connected both to the output units and to a group of 8 fully recurrent internal units, which in turn send connections to the output units. The group of output units is composed by the two motor units, encoding the speed of the two wheels, and one communication unit, which encodes the signal emitted by the robot (in $[0, 1]$). Output units are simple logistic units, whose activation o_i is given by

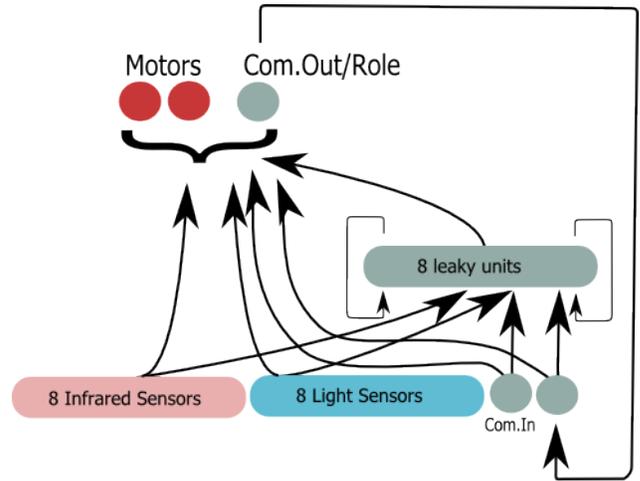$$o_i = \phi(I_i), \phi(x) = (1 + e^x)^{-1}, I_i = \sum_j^N a_j \times w_{ij} - b_i \quad (1)$$

where a_j is the activation of the unit j that sends connection w_{ij} to output i and b_i is the bias of unit i. The activation of each motor unit (in $[0, 1]$) is remapped in $[-8, +8]$ and is sent to the motors of the robot. Internal units are leaky integrator units, whose activation is given by

$$o_i^t = (1 - \tau_i) \times \phi(I_i^t) + \tau_i \times o_i^{t-1} \quad (2)$$

where τ_i is the time constant (in $[0, 1]$) of unit i.

Evolutionary algorithm

The genome of a robot encodes the values of the connection weights, the biases and the time constants of the internal units. Each parameter is encoded as an 8 bits string, whose value is uniformly mapped in the range $[-5.0, +5.0]$ for weights and biases and in the range $[0, 1]$ for time constants. The first population consists of 100 randomly generated genomes. Each genome is tested for 15 trials, each lasting 2000 time steps. At the beginning of each trial the genome is translated in the corresponding neural controller, and the same controller is embedded in each of the four robots that constitute the group. The four robots are randomly placed into the home square of the arena and let free to move and communicate for the rest of the trail. After the fitness of each genome in a generation has been calculated (see below), the 20 best genomes are copied 5 times each (reproduction) and 1% of the bits of each new genome is flipped (mutation). The evolutionary process lasts 1000 generations.

We run two kinds of experiments which differ only for the fitness function employed. In the first experiment, which we will call 'base-line', the fitness function is purely behavioural, rewarding only the ability of the robots to stay in the configuration that we are aiming for, that is having one and only one robot in the corridor close to the light while all other robots remain in the home square. In particular,

the fitness of the base-line experiment is composed by two components, rewarding separately a) the behaviour of the 'leader' (defined as the robot who is, in each time step, most close to the light), and b) the behaviour of the 'followers' (i.e. all other robots). In particular, the first behavioural fitness component (BFC1) measures how close the leader is to the light, normalised in $[0, 1]$:

$$BFC1(g) = \frac{\sum_t^T \frac{max(0,(M-d(L,light)))}{M}}{T} \qquad (3)$$

where M is a maximal distance, set to 900 mm and $d(L, light)$ is the distance between the leader and the light, and T are all the time steps of all the trials of a genome's 'life' (i.e. 1000 time steps x 10 trials = 10000).

The second component (BFC2) measures the average distance of the followers from the light, again normalised in $[0, 1]$:

$$BFC2(g) = \frac{\sum_t^T \sum_i^F \frac{min(1,d(F_i,light))}{M}}{T \times F} \qquad (4)$$

where F_i is follower i and F is the number of followers, i.e. 3. Since the behaviour of the leader is more important than that of the follower, the global behavioural fitness (BF) gives more weight to F1 than to F2, in particular:

$$BF(g) = 0.75 \times BFC1(g) + 0.25 \times BFC2(g) \qquad (5)$$

The second experiment, which we will call the 'communication rewarded' experiment is identical to the first one (the ratio between the leader and the followers fitness components is kept constant) but for the addiction of another component to the fitness, which is identical to the fitness that we used in our previous work (Gigliotta et al., 2009) for evolving role-allocation based on communication. In particular, this component rewards the genome for having one robot sending a high value signal and all the other sending a low value signal. More precisely, the communication fitness component (CFC) is measured as follows:

$$CFC(g) = \frac{\sum_t^T \sum_i^N O_{max} - O_i}{T \times (N-1)} \qquad (6)$$

where O_{max} is the highest signal value, O_i is the value of the signal of robot i, and N is the number of robots in the group, i.e. 4. Finally, the global fitness of the communication rewarded simulation is

$$CRF(g) = 0.8 \times BF(g) + 0.2 \times CFC(g) \qquad (7)$$
$$= 0.2 \times BFC1(g) + 0.6 \times BFC2(g) + 0.2 \times CFC(g) \qquad (8)$$

Each experiment is replicated 20 times by starting with different random conditions.

Results

Figure 3 shows the evolution of the fitness of the best individual of each generation for the two experiments: both the global value and the three different components are shown (average results of 20 replications). By looking at the global fitness (that determining the evolutionary process: fig 3a), we can see that the communication rewarded condition is much better than the base-line one: in the former we can clearly see an evolutionary improvement (from 0.3 to 0.75); in the latter fitness is almost flat (from 0.3 to 0.4). It is already clear that this difference can not be explained only by the communicative fitness component, which influences the fitness of the communication rewarded condition only for one fifth (0.2). Indeed, if we look at the fitness related to the behaviour of the leader (BFC1, fig. 3b) we see that there is a big difference between the two conditions. At the end of the evolution, in the communication rewarded condition the best genomes get on average a score of more than 0.5, meaning that the leaders stay on average at a distance of less than 450 mm form the light, which is *inside* the corridor, while in the base-line condition they score 0.2. Even more interesting is the evolution of the component related to the behaviour of the followers (BFC2, fig. 3c): while in the communication rewarded condition this component is stable to 0.9, in the base-line condition it *decreases* during evolution. Such a decrease is concomitant with the increase in the first behavioural fitness component, meaning that as the leader learns to get closer to the light, so do also the followers. Finally, the plot of the communication component, which is present only in the communication rewarded condition, shows that the evolution of this component is very fast: it reaches a steady state of about 0.95 after about 50 generations. In particular, this fitness component evolves well *before* than the (leader) behavioural component. This, together with the fact that the absence of this component prevents the evolution of the desired behaviour, shows that it is just the ability to dynamically allocate tasks through communication, evolved thanks to the communication fitness component, that allows the emergence of the ability to solve the behavioural task.

This conclusion is further corroborated by looking at the communicative behaviour of the two conditions. In particular, we measured how much the signals of the robots in a group are differentiated by running, for each of the best evolved individuals of the two conditions, 100 test trials. During these tests we recorded, for each timestep, the robot communicative behaviour by binarizing the communication output: we assign 1 to robots whose communication unit exceeds 0.5 and 0 otherwise. Hence, for a group formed by 4 robots we have 16 (2^4) possible configurations. The frequencies of such states are reported in fig 4. Furthermore, we computed the entropy (Shannon, 1948) of the two distri-

(a) Global fitness
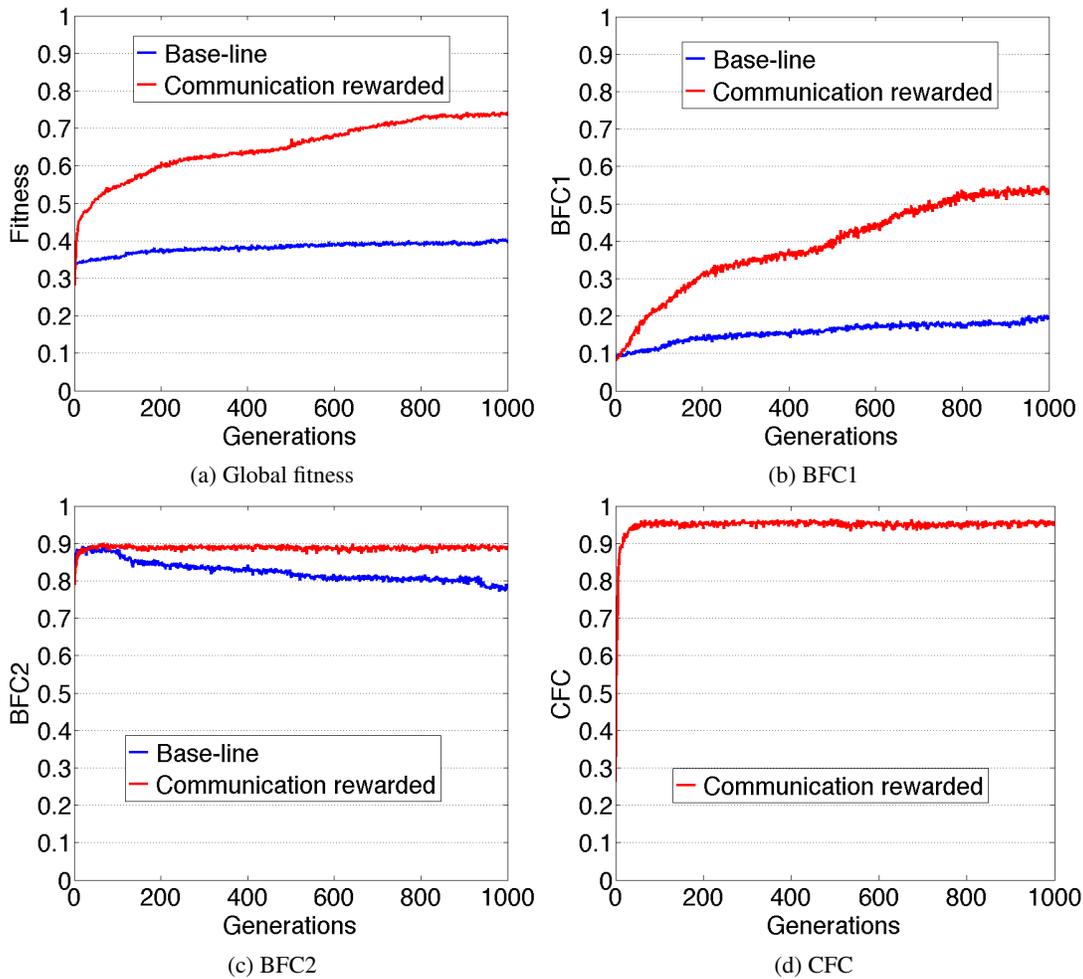

(b) BFC1

(c) BFC2

(d) CFC

Figure 3: Evolution of fitness. All data refer to the averages of the best individuals of the 20 replications. BFC1 = first behavioural fitness component (related to the leader). BFC2 = second behavioural fitness component (related to the followers). CFC = communication fitness component (only present in the communication-rewarded condition.

butions (of the two conditions):

$$H = -\sum_{i=1}^{16} p_i \log_2 p_i \qquad (9)$$

where p_i is the experimental probability (i.e. the relative frequency) of the i_{th} configuration.

While in the communication rewarded condition the signals are differentiated (since, as requested by the communication fitness component, one robot sends a high value signal while the others send a low value one) presenting an entropy equal to 2.23, in the base-line condition the entropy is very close to 0 (0.00036), meaning that all robots send always the same signal (that is 0 in this case). This means that the robots of the base-line condition do not use communication at all. The inability of the base-line condition to evolve an appropriate communication system with which to communicate the different roles to one another is the reason for
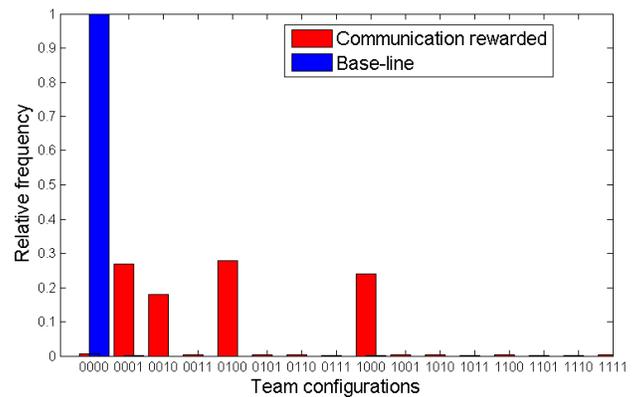

Figure 4: Frequency distribution of team configurations for the best individual of the base-line and the communication rewarded condition.

the inability to evolve the appropriate non-communicative behaviour. Indeed, given that all robots in a group have the same body and the same control system, in order to behave differently they need something that might differentiate the two behaviours of a) going in the corridor and towards the light (leader) and b) remaining outside the corridor in the home (followers). The only thing that seems to be able to permit this differentiation is communication. In fact, the signals have a twofold function. The first is social, that is to make sure that in the group there is one and only one leader. The second is individual, that is to permit the differentiation of individual behaviours: the robot that sends a high signal is the one that looks for the corridor, enters it, and is attracted by the light; the robots that send low signals remain in the home.

The difference between the results of the base-line and the communication rewarded conditions can be best appreciated if we measure the performance of the groups of robots on the basis of the task that we wanted them to accomplish in the first place. Our goal was to have one robot to go 'on mission' inside the corridor and all the others to remain in the home. For this reason we measured performance as the proportion of time steps in which the robot are in the desired configuration, i.e. one robot in the corridor and all the other outside. Furthermore, since it necessarily takes some time for the group to get to the desired state (as roles needs to be negotiated and the leader has to find the corridor and enter it), we take the performance measure only in the last 100 time steps of a trial (choosing a different interval like the second half of the trial or just the last time step gives almost the same results). The average results of the 20 replications of the two conditions in the 100 test trials are shown in fig. 5. The results speak for themselves. The median of the communication rewarded condition is about 0.9, meaning that in most of the replications the robots solve the task almost always; there are also some 0s, as in a few replications the desired behaviour has not evolved (yet?). In striking contrast, the performance for all the best evolved genomes of the base-line condition is 0, meaning that no group of that condition has ever achieved the desired configuration. This result is related to the fact, shown above, that as the fitness component rewarding the behaviour of the leader increases that of the behaviour of the follower decreases. Given that in the base-line condition robots never learn to differentiate their communicative behaviour and hence their roles, in those replications in which the behaviour of entering the corridor and approaching the light evolves, all the robots of the group exhibit that same behaviour, resulting in a configuration where all the four robots queue inside the corridor towards the light. Given that the behaviour of the leader is weighted more, in the fitness, than the average behaviour of the followers (eq. 5), that different robots enter the corridor at different times, and that for physical reasons the last robot to enter the corridor remain distant from the light be-
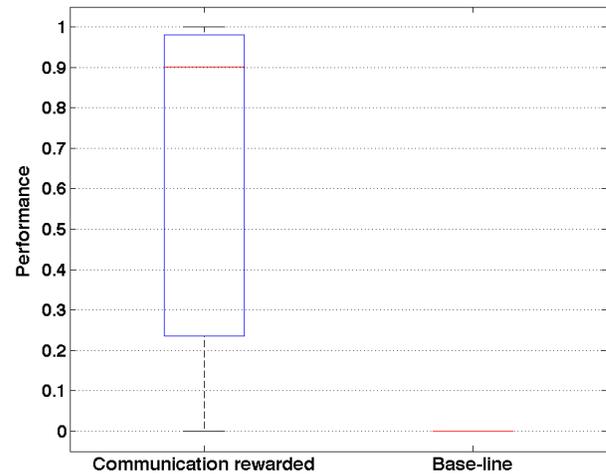

Figure 5: Post-evaluation results for all the 20 best individuals of the base-line and the communication rewarded conditions. Performance is calculated as the proportion of time steps, among the the last 100 of a trial, in which the group of robots accomplishes the task, meaning that one and only one robot is inside the corridor. Each box represents the inter-quartile range of the data, while the red line inside the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box.

cause they have other robots in front of them, this kind of behaviour tends to provide a higher fitness than the behaviour of remaining all the robots in the home, but nonetheless, it does not solve the task, and results in a performance of 0.

Finally, we tested the robustness of the best solution to the variation in the number of robots of the group. We took the best individual out of all the replications of the communication rewarded simulations and we tested its performance with a number of robots that goes from 2 to 10. In particular, for each test we run 100 trials where we measure, as before, the proportion of time steps in which one and only one robot is in the corridor, during the last 100 time steps of each trial. The results (fig. 6) show that the evolved solution is extremely robust to the number of robots in the team: from 2 to 7 robots performance is above 0.9, and even with 10 robots we obtain a quite high performance of about 0.65. Furthermore, by observing the behaviour of the team with 10 robots, we can see that even when the robots fail to accomplish the task, this is not due to a disruption of the group behaviour, but rather to the fact when there are so many robots the arena is overcrowded and it becomes difficult for the leader to find the corridor and enter it. If we would just prolong the duration of the trial or enlarge the arena the performance would return to optimal values even with 10 robots.
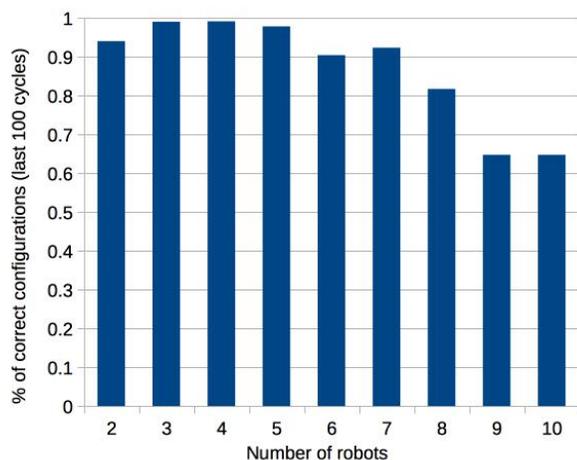
Figure 6: Generalisation tests of the best evolved individual of the 'communication rewarded' condition. The proportion of time steps (among the last 100 of a trial) in which one and only one robot is in the corridor is plotted against the number of robots in the team.

Discussion and conclusion

In this paper we have demonstrated that it is possible to evolve a team of homogeneous robots in which the robots autonomously negotiate different roles among themselves and behave according to the allocated roles. In particular, our robots are able to elect a 'leader' that is sent to a 'mission' outside the home area, while all other 'follower' robots remain in the home.

In contrast to all the analogous previous works which we are aware of, where the evolved behaviours were very fragile, in particular with respect to the number of robots, our system proved to be extremely robust, both to the addiction and to the subtraction of robots in the group. We maintain that this high robustness depends on the presence of a communication channel that is used to allocate the roles. If, as in previous works (Baldassarre et al., 2003; Quinn et al., 2003; Tuci et al., 2013) role allocation depends only on robots behaviour, it will tend to be situated, that is dependent on the context. This means that as the context changes, for example because the groups contains a different number of robots, then the behaviour will tend to stop functioning. On the contrary, if roles are allocated and maintained through a dedicated communication channel, they will tend to depend much less from the context, which renders the system much more robust.

The present work demonstrates that the method we had proposed for communication-based role allocation in previous work (Gigliotta et al., 2009) can be effectively applied to develop robots able to accomplish non-communicative collaborative task. Furthermore, our simulations show also that, in some circumstances (like the ones of our scenario), di-

rectly rewarding communication-based role allocation can be necessary for the non-communicative team behaviour to evolve. The fact that if communication-based role allocation is not rewarded through the fitness function the robots cannot solve the task suggests that it is too difficult to co-evolve the necessary communicative and non-communicative behaviours from scratch. Consider that the evolution of communicative behaviours per se can be difficult as the traits which are necessary for its emergence – namely, sending good signals and responding appropriately to received signals – taken in isolation do not increase the reproductive chances of the individuals that possess them (Maynard-Smith, 1997; Mirolli and Parisi, 2008). Furthermore, our robots have to develop not only their communicative behaviours, but these must be co-adapted to the appropriate non-communicative behaviours, that is going 'on mission' for the leader and remaining in the home for the followers. Even if this is surely possible in principle, our simulation show that it is too difficult in practice. On the contrary, if communication-based role allocation is rewarded through the fitness function the evolution of the desired behaviour becomes feasible, and even easy (remember that the median performance in the tests of the best individuals of 20 replications of the experiments was 0.9). The reason is that, if explicitly rewarded, the evolution of communication-based role allocation is straightforward, and once role are allocated the evolution of the behaviours that are appropriate to the roles is not so difficult.

Some researchers in the Evolutionary Robotics community may not like the explicit rewarding of the communicative behaviour of our robots, as this might sound like 'cheating', in the sense of providing too much information and canalising the evolutionary process. Even though we sympathise with the efforts to let evolution as free as possible to find his own solution, we also agree with the maxim ascribed to Albert Einstein that "everything should be made as simple as possible, but no simpler". In our case, the fact that rewarding communication leads to effective behaviour but avoiding to do so leads just to failure widely justifies the use of such a procedure. Of course, the impact of such a choice depends on one's goals. If one's goals are technological, i.e. to develop teams of autonomous robots that are able to allocate different roles and behave accordingly for practical purposes, there is no reason to avoid the use any kind of expedient, let alone one that proves to be very effective. If one's goal is scientific, that is if one is interested in understanding role differentiation in real organism, the use of fitness components that directly reward communicative behaviours is more problematic, as it is not clear how such communicative behaviours per se may lead to higher chances of survival and reproduction. However, if one is not interested in understanding the *evolution* of role differentiation but rather how role differentiation can be realised in groups of homogeneous individuals, the way in which we

get to our role-differentiating robots is not relevant and ours is just an effective way to obtain the behaviour that we want to investigate which otherwise could not be studied (because without rewarding communication role differentiation could not evolve). The detailed analysis of how our robots manage to solve the task is the first line of future research.

A second possible line of future research consists in trying to evolve other collaborative behaviours which have different kinds and numbers of roles. Our approach is not strictly related to the development of a group in which there is one individual that plays one role (the 'leader') and all the others play another role (the 'followers'). In fact, in previous work we have already shown that other ways of allocating roles are possible: for example, we were able to evolve groups which had to split in two evenly distributed subgroups, with half of the robots sending high-value signals and the other half sending low value signals. It will be interesting to demonstrate that other interesting non-communicative collaborative behaviours can be developed that require this or other kinds of role allocation.

Finally, another possible development consists in trying to release the designer from the need to specify the number and distribution of roles within the group by using information theoretic measures like Shannon entropy (Shannon, 1948) to reward robots for assuming different roles in different contexts while leaving the robots free to autonomously determine the right number and distribution of roles depending on the task at hand.

Acknowledgements

References

Anderson, C. and Frank, N. (2001). Teams in animal societies. *Behavioral Ecology*, 12:534–540.

Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9:255–267.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. pages 1–41.

Dorigo, M., Trianni, V., Sahin, E., Gross, R., Labella, T., Baldassarre, G., Nolfi, S., Deneubourg, J.-L., Mondada, F., Floreano, D., and Gambardella, L. (2004). Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2-3):223–245.

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

Garnier, S., Gautrais, J., and Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intelligence*, 1:3–31.

Gigliotta, O., Mirolli, M., and Nolfi, S. (2009). Who is the leader? dynamic role allocation through communication in a population of homogeneous robots. In Serra, R., Villani, M., and Poli, I., editors, *Artificial Life and Evolutionary Computation (Proceedings of Wivace 2008)*, pages 167–177. World Scientific Publishing.

Gigliotta, O., Mirolli, M., and Nolfi, S. (Subm). Communication based dynamic role allocation in a group of homogeneous robots. *Natural Computing*.

Harvey, I., Di Paolo, Ezequiel andTuci, E., Quinn, M., and Wood, R. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1–2):79–98.

Maynard-Smith, J. (1997). *The theory of evolution*. Cambridge University Press, Cambridge.

Mirolli, M. and Nolfi, S. (2010). Evolving communication in embodied agents: Theory, methods, and evaluation. In Nolfi, S. and Mirolli, M., editors, *Evolution of Communication and Language in Embodied and Situated Agents*, pages 105–121. Springer, Berlin.

Mirolli, M. and Parisi, D. (2005). How can we explain the emergence of a language which benefits the hearer but not the speaker? *Connection Science*, 17(3-4):325–341.

Mirolli, M. and Parisi, D. (2008). How producer biases can favour the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*.

Mirolli, M. and Parisi, D. (2011). Towards a vygotskyan cognitive robotics: The role of language as a cognitive tool. *New Ideas in Psychology*, 29:298–311.

Nolfi, S. (1998). Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science*, 10(3-4):167–183.

Nolfi, S. (2005). Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. *Connection Science*, 17(3-4):231–248.

Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA.

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

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.

Spector, L., Klein, J., Perry, C., and Feinstein, M. (2005). Emergence of collective behavior in evolving populations of flying agents. In *Genetic Programming and Evolvable Machines*, 6, pages 111–125.

Tuci, E., Mitavskiy, B., and Francesca, G. (2013). On the evolution of self-organised role-allocation and role-switching behaviour in swarm robotics: a case study. In *Advances in Artificial Life (Proceedings of ECAL2013)*, pages 379–386.