

Evolving Mobile Robots Able to Display Collective Behaviours

Gianluca Baldassarre, Stefano Nolfi, Domenico Parisi

Institute of Cognitive Sciences and Technologies

National Research Council (ISTC-CNR)

Viale Marx 15, 00137, Rome, Italy

baldassarre@ip.rm.cnr.it, nolfi@ip.rm.cnr.it, parisi@ip.rm.cnr.it

Abstract

In this paper we present a set of experiments in which a group of simulated robots were evolved for the ability to aggregate and to move together toward a light target. Evolved individuals display interesting behavioural patterns in which groups of robots act as a single unit. Moreover, groups of identical evolved individuals display primitive forms of "situated" specializations in which different individuals play different behavioural functions according to the circumstances. Overall, the results presented in the paper demonstrates that the evolutionary method, by allowing the synthesis of behaviours that emerge from the interaction between the robots and the environment and between different individual robots, is a powerful method for synthesizing collective behaviour.

1. Introduction

The attempt to develop autonomous robots through a self-organization process based on artificial evolution is gathering increasing attention. The main advantage of this approach is that it is an ideal framework for synthesizing robots whose behaviour emerge from a large number of interactions among their constituent parts. This can be explained by considering that, in evolutionary experiments, robots are synthesized through a self-organization process based on random variation and selective reproduction where the selection process is based on the behaviours that emerge from the interactions among the robot's constituent elements and between these elements and the environment. This allows the evolutionary process to freely exploit interactions without the need to understand the relation between interactions and emerging properties as it is necessarily required in other approaches that rely more on explicit design.

On the basis of the same argument we should assume that the evolutionary approach can be successfully applied also to synthesize robots able to display collective behaviours. In this case evolving individuals might exploit not

only the properties that emerge from the interactions among the constituent elements of the robot and between the robot and the environment but also the interactions among different individual robots.

In this paper we present a set of experiments in which a group of simulated robots were evolved for the ability to aggregate and to move together toward a light target. As we will see, evolved individuals display interesting behavioural patterns in which groups of robots act as a single unit. Moreover we will see how groups of identical evolved individuals display primitive forms of "situated" specializations in which different individuals play different behavioural functions according to the circumstances.

In the next section we review the related literature. In section 3 and 4 we present our experimental framework and the obtained results. Finally, in section 5 we discuss the implications of the obtained results and the future directions of our work.

2. Related work

Despite the possible advantages mentioned above, the use of artificial evolution to synthesize robots able to display collective behaviour is still a rather unexplored area.

A pioneering work was conducted by Martinoli (1999) who used artificial evolution to synthesize the control system of a group of simulated Khepera robots (Mondada et al., 1993) that were asked to find "food items" randomly distributed on an arena. The robots were provided with 10 infrared sensors and two motors controlling the two corresponding wheels. Eight sensors were distributed around the robot body and were used to detect the other members of the group and 2 sensors were placed under the robot body and were used to detect food items (i.e. black painted disks scattered on the white ground). The controllers of the robot consisted of a simple neural network with 10 sensory neurons, that encoded the state of the infrared sensors, directly connected with two motor neurons, that controlled the two corresponding wheels. The genotype of the evolving individu-

als encoded the connection weights of the neural controllers and fitness consisted in the percentage of food items detected by a single individual (individual fitness) or by the entire group (group fitness). Unfortunately, a detailed and quantitative analysis of the obtained behaviours is missing. The author only reports that, in some cases, evolved individuals display interesting collective behaviours such as exploring the arena in couples.

Other experiments have been conducted by evolving groups of artificial creatures. Reynolds (1993) evolved the control system of a group of creatures placed in an environment with static obstacles and a manually programmed predator for the ability to avoid obstacles and predation. Despite the results described in the paper are rather preliminary, some evidences indicate that coordinated motion strategies begun to emerge. In the attempt to study the evolutionary origin of herding, Werner and Dyer (1993) co-evolved two populations of predators and prey creatures that were selected for the ability to catch prey and to find food and escape predators respectively. By analysing the result of a single run, the author observed that after some generations, during which predators evolved an ability to catch prey, creatures converged into small herds which were constantly spitting up and reforming. In a more recent work, Ward et al. (2001) evolved groups of artificial fish able to display schooling behaviours. Also in this case two populations of predator and prey creatures were evolved and placed in an 2D environment containing randomly distributed food elements. Creatures consisted of neural network controllers with two motor neurons controlling the speed and the direction of motion and two set of sensory neurons encoding: (a) the distance and the direction of the nearest prey, predator and food, and (b) the amount of changes in water pressure in the close proximity of the creature. The analysis of the distance between prey, prey and food, and predator and prey suggest that the emergence of the schooling behaviour is correlated with: (a) an advantage in the ability to find food clumps, and (b) a better protection from predation.

Finally, Theraulaz and Bonabeau (1995) evolved a population of constructor agents who collectively build a nest structure by depositing bricks according to their perception of the local environment and to a set of behavioural rules (see also Bonabeau et al. 2000).

Overall, these experiments confirm that artificial evolution can be successfully applied to synthesize effective collective behaviours. Whether these results could be generalized to the synthesis of realistic creatures (robots), however, remains to be ascertained given that in these experiments creatures rely on sensory systems that provide “perfect sensory information” (e.g. information free from noise) and often “unrealistic sensory information” (e.g. distances of objects of different types).

3. The experimental framework

Within the E.U. project SWARM-BOTS whose goal is the development of swarms of self-assembling artefacts (IST-

2000-31010) we tried to evolve the control system of groups of mobile robots able to aggregate and to move together toward a light target.

The environment used in the experiments is a rectangular arena of 1×2 metres surrounded by walls and includes two halogen light bulbs of 230 watts that are located in the middle of the west and east walls at an height of 1.5 cm from the ground (Figure 1).

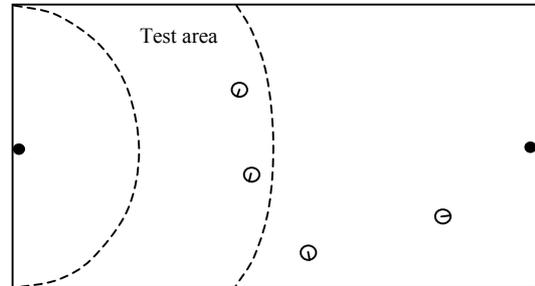


Figure 1: The environment and the robots. Lines represents the walls surrounding the arena. The full circles indicate the two light bulbs. The empty circles represent the robots. The dashed curves indicate the edges of the test area when the west light target is on (for an explanation see section 4).

Groups consists of four Khepera robots (Mondada et al., 1993) each provided with 8 infrared sensors able to detect walls and other robots up to a distance of 45 mms, 8 ambient light sensors, (i.e. the same infrared sensors used in a passive mode), and four “simulated” directional microphones. From the motor point of view, each robot is provided with two motors that control the speed of the two corresponding wheels and a simulated speaker that continuously emits a sound with a fixed amplitude and a frequency that randomly vary within a fixed range (Figure 2).

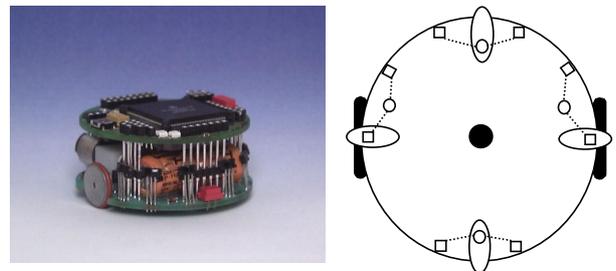


Figure 2: **Left:** The Khepera robot (sound speaker and microphones not present). **Right:** A schema of the position of sensors and motors within the robot body. Empty squares and ovals represent the 8 infrared sensors and the 4 directional microphones, respectively. The infrared sensors are also used to measure the ambient light. The four light sensors (represented by the empty circles) are obtained by averaging the activation value of two adjacent ambient light sensors. The two full segments and the full circle represent the two wheels controlled by two separate motors and the speaker.

Experiments were conducted in simulation by using an extended version of Evorobot (Nolfi, 2000). To simulate the

robot and the environment as accurately as possible, a sampling procedure was used to compute the activation state of the infrared and ambient light sensors (see Nolfi and Floreano, 2000). Walls and Khepera robots were sampled by placing one physical robot in front of them and by recording the state of the infrared sensors while the robot was turning 360 degrees at 20 different distances from of each object. These recorded values were used in simulation to set the activation states of the simulated infrared sensors on the basis of the current angle and distance of the robot with respect to the other robots and walls. The same procedure was used to sample the state of the ambient light sensors for different orientations and distances with respect to a light bulb. These recorded values together with a geometrical simulation of shadows were then used in simulation to set the activation state of the ambient light sensors.

Given that we do not have physical sound sensors available yet, we simulated the sound amplitude recorded by a directional microphone as follows:

$$A = (1 / (1 + (D^2/1000^2))) * AF \quad (1)$$

where A is the amplitude, D is the distance from the sound source in mm, 1000 is a constant that assures that the maximum amplitude is 0.5 after 1 m, AF is an attenuation factor that is proportional to the angular difference between the microphone and the sound source. AF is computed as follows:

$$AF = 1 - 0.9 * (\alpha/180) \quad (2)$$

where α ranges from 0 to 180 degrees and encodes the angle of the simulated microphone with respect to the sound source, and 0.9 is a constant that assures that the range of the attenuation factor is [0.1, 1.0].

The total perceived amplitude (TPA), i.e. the result of the contribution of different sound sources corresponding to different robots, is computed according to the following function:

$$TPA = (2 / (1 + \exp(-\sum_i^{N-1}[A_i]))) - 1 \quad (3)$$

where A_i is the contributions of the sound coming from the other $N-1$ robots and computed according to (1). It should be noted that sound waves can sum up or cancel out on the basis of their frequency and phase. The co-existence of different sound sources with slightly different frequencies and phases that originate from different positions, therefore, tends to produce a semi-irregular wave with picks and valleys caused by summation or deletion of waves with different amplitudes and frequencies. For this reason we assumed that the signal recorded by each microphone is passed to a filter that returns a moving average of the maximum picks recorded within a given time window. Therefore, in equation 3, different sound sources sum up less than linearly. Finally, a random value with a uniform distribution within the range [-0.05, 0.05] is added to all sensors each time step.

The initial population consists of 100 randomly generated genotype strings that encode the connection weights of 100 corresponding neural controllers (Figure 3). Each controller is made up of a neural network with 17 sensory neurons, that encodes the state of the infrared, light, sound sensors, and a bias unit (i.e. a unit whose activation state is always 1.0), directly connected with 2 motor neurons, that control the speed of the two wheels. Each connection weight is represented in the genotype with 8 bits that are transformed in a number in the interval [-10, +10]. Therefore, the total length of the genotype is $34 * 8 = 272$ bits.

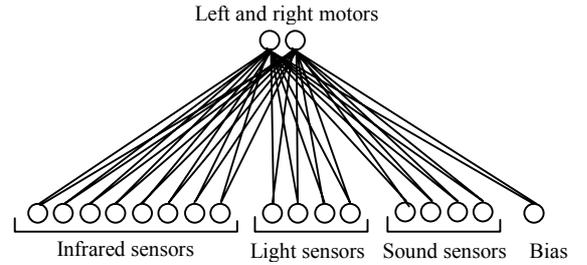


Figure 3: The neural network controller.

Each genotype is translated into 4 identical neural controllers corresponding to a group of 4 robots. The group is allowed to “live” for 4 “epochs” (each epoch consists of 1500 cycles and each cycle lasts 100ms). During each cycle, for each robot: (1) the activation state of the sensors is set according to the procedure described above; (2) the activation state of the two motor neurons is computed according to the standard logistic function; (3) the desired speed of the two wheels is set according the activation states of the motor units.

At the beginning of each epoch the four robots are placed in randomly selected positions and orientations within the arena and only the east light is turned on. However, in order to push robots to continuously move toward the light, when the barycentre of the robots’ group (i.e. the average of their x and y coordinates) reach a distance lower than 30 cm from the light that is currently on, the light that is on is turned off and the other light is turned on.

To force the robots to move together toward the light, we designed a fitness function with two components: a group’s compactness component (GCC), and a group’s speed component (GSC) that are computed as follows:

$$GCC = \sum_i^N [1 - (SD_i / 300)] / N \quad (4)$$

where SD_i is the distance of the robot i from the group’s barycentre. SD_i are thresholded to 300mm so that distances larger than 300 mm do not contribute to the fitness of the group;

$$GSC = (1 + (-\Delta GD / 7)) / 2 \quad (5)$$

where ΔGD is the variation of the distance between the group’s barycentre and the light target, and 7 is a constant

value that corresponds to the maximum possible advancement in mm of a robot or a group or robots during a cycle (100 ms). Notice that if the group does not move GSC is equal to 0.5.

The fitness F of a robots' group is the average of the weighted sum of the two components during the individuals' lifetime and is computed as follows:

$$F = \frac{1}{M} \sum_{t=1}^M [(GCC_t * W) + (GSC_t * (1 - W))] \quad (6)$$

where W is a constant within $[0.0$ and $1.0]$ that determines the relative weights of the two components of the fitness, and t is the current lifecycle ranging from 1 to M ($M=6000$).

The best 20 genotypes (corresponding to a group of four identical neural controllers) of each generation were allowed to reproduce by generating 5 copies of their genotype with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasted 100 generations. The experiment was replicated 10 times by starting with different randomly generated initial populations.

4. Results and discussion

By running a set of experiments in which W was set to 0.75 we observed the emergence of an interesting variety of strategies. Figure 4 shows the fitness of the best group of individuals of each replication. In all replications evolved individuals are able to form groups, keep all together, and move toward the light target. By looking at the contribution of the two components with respect to the overall fitness, we see how, with respect to the ability to form a group and keep all together (GCC), rather good performance are obtained in all replications. On the contrary, with respect to the ability to move together toward the light target (GSC), significant differences are observed within different replications and best performance are achieved in the first replication (notice that the absolute differences between the two components is due to the value of the W parameter).

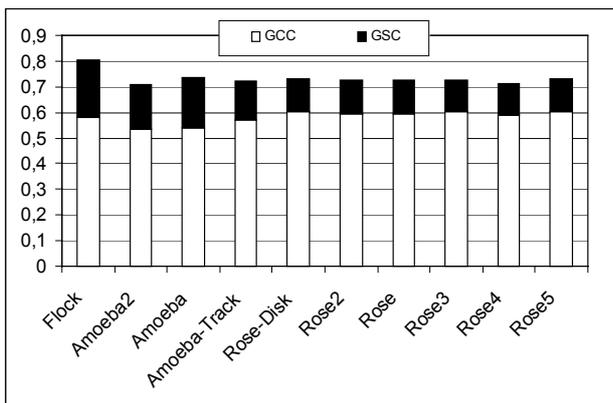


Figure 4. Performance of the best group of robots of each replication after 100 generations. Data obtained by testing each group of individuals for 10 epochs. The white and black parts of each histogram represents the contribution of the compact and speed components, respectively. The names at the bottom indicates a qualitative

description of the type of strategy emerged in the corresponding replication (see below).

By analysing this set of experiments and by running other experiments with W set to 0.5 or 0.8 and/or with groups of 2 or 7 robots (detailed results not shown), we observed that all emerged formations and strategies can be categorized within three qualitatively different classes that we named *Flock*, *Amoeba*, and *Rose* (Figure 5).

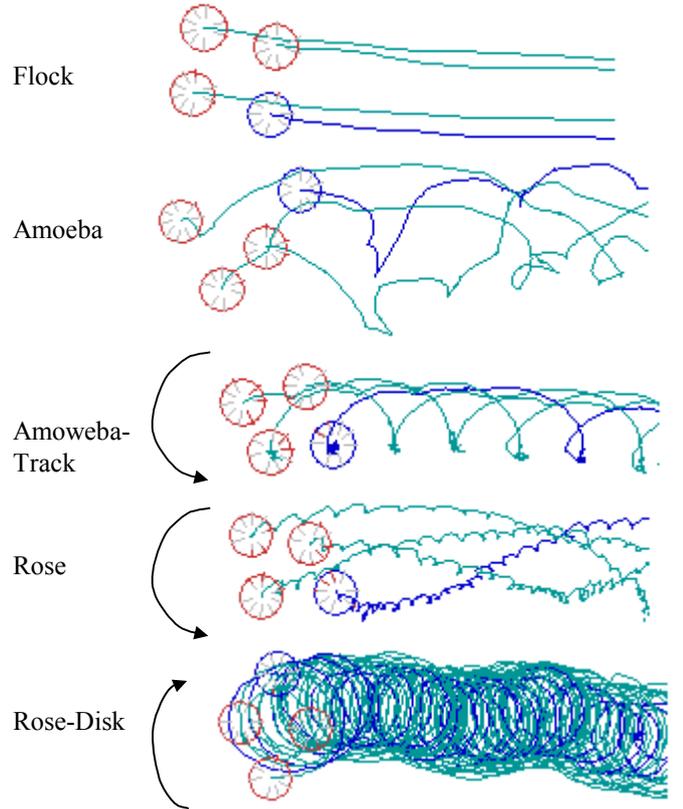


Figure 5: Three prototypical formations, *Flock*, *Amoeba* and *Rose* and some of their variants, *Amoeba-Track* and the *Rose-Disk*. The four circles represents the robots. The lines represents the trajectory of the robots while the group is approaching the west light (for clarity, the trajectory of one of the four robots is displayed with a darker colour). The arrows indicate the direction of the rotation of the group with respect to the barycentre of the group itself.

Flock. Robots form a rather compact group that then move straight toward the light target. While the group approaches a target, individual robots tend to maintain their initial relative position with respect to the rest of the group and with respect to the target (e.g. a robot that happens to occupies the right-front position in the group at the beginning of an approaching phase tends to remain in that position). This strategy is the most effective overall and is particularly effective from the point of view of the GSC component (i.e. from the point of view of maximizing the speed of the group's movement toward the light target).

Amoeba. Robots form a group and then move toward the light target by varying their relative position. This strategy is

not very effective from the point of view of the GCC component but is quite effective from the point of view of the GSC component (although slightly less effective than the flock strategy given that individual robots do not move fully straight toward the light). Robots do not keep the same relative positions while approaching the light target. For example the robot that is currently at the rear of the group might speed up and reach the front of the group while the other robots tend to stay still or to rotate on themselves. *Amoeba-Track* is a variation of the strategy described above. The group is more compact in this case and the relative movements of the individuals of the group are more regular so that the movement of the whole group resembles that of the track of a caterpillar seen from a side.

Rose. Robots form a very compact group and rotates on themselves, and eventually over the barycentre of the group, while moving toward the light. The shape of the group tend to vary for groups of different numbers (i.e. typically 4 robots form a rhombus while 7 robots form a rose with 6 robots forming a circle around 1 robot located in the centre of the circle itself). *Rose-Disk* is a variant of the rose in which robots rotate around the barycentre of the group but not on themselves. Given that these formations are very compact and stable in time, this strategy is the most effective from the point of view of the GCC component. However, it is also very ineffective from the point of view of the GSC component. Indeed, compactness is achieved through motor behaviour (e.g. turning on itself, and/or turning around the group's barycentre, and/or oscillating back and fourth) that interfere with the ability to fast approach the light.

To analyse these strategies also quantitatively we designed few statistical indexes that allow to better characterize the behaviour of single robots and of the whole group that are described below.

The first index, named *group stability index* (GSI), measures how much a formation is stable in time or, in other words, how much the relative position of each robot with respect to the other robots does not change in time. For example, if the group of 4 robots tend to form a square formation, the index will indicate how long the robots maintain the same geometrical formation and how stable the dimension of the square is in time. More precisely, the stability index (GSI) at cycle t is computed as follows:

$$D_{ij}^t = \text{dist}[(x_i^t, y_i^t), (x_j^t, y_j^t)] \quad (7)$$

$$\text{GSI}^t = 1 - \left(\frac{\sum_i^N \left[\sum_{j=i+1}^N [|D_{ij}^t - D_{ij}^{t-1}|] / 7 \right]}{(N * (N-1) / 2)} \right)$$

where D_{ij}^t is the Cartesian distance between each couple of robot i and j at time t , N is the number of robots in the group (Figure 6, top), $(N * (N-1) / 2)$ is the number of distances between the robots, and 7 is the maximum possible displacement during a cycle. Notice that GSI index usually ranges between the interval $[0, 1]$ and only occasionally

becomes negative (e.g. when robots go away from each other at maximum speed).

The second index, named *group role index* (GRI), measures how much individual robots within the group tend to keep the same relative position with respect to the other robots and the light target. The index corresponds to the average variation between two successive lifecycles of the x and y coordinates of each robot with respect to a Cartesian reference system in which the origin corresponds to the barycentre of the group and the y axis correspond to the line that connect the barycentre to the light target (see Figure 6, bottom). More precisely, the index at cycle t is computed as follows:

$$\text{GRI}^t = 1 - \left(\frac{\sum_i^N [\text{dist}[(x_i^t, y_i^t), (x_i^{t-1}, y_i^{t-1})] / 7]}{N} \right) \quad (8)$$

where $\text{dist}[\]$ is the distance between the x and y position of robot i at time t and $t-1$ (notice that x and y coordinates refer to the current Cartesian reference system), and 7 is the maximum possible displacement during a cycle. Notice that RI index typically ranges between the interval $[0, 1]$ and only occasionally becomes negative (e.g. when one light is turned on and the other is turned off).

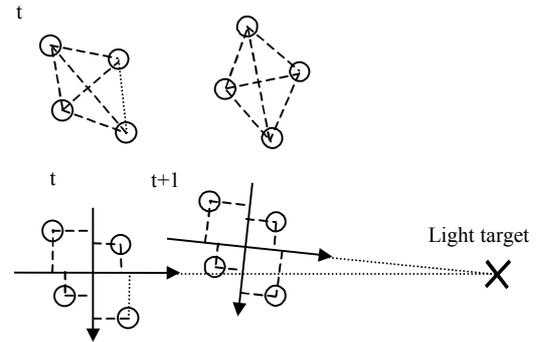


Figure 6: **Top**: an example of how the compact index is computed. The circles represent the robots, and the dotted lines indicate the distances between each couple of robot. The index measures the average difference in distances at time t and $t+1$. **Bottom**: an example of how the role index is computed. The circles and the arrows represent the position of the robots and the Cartesian reference system at lifecycle t and $t+1$. The dotted lines indicate the x and y co-ordinates of each robot with respect to the reference system.

The third and last index is the *rotational index* (RI), that measures how much individual robots rotate on themselves. This index at time t is computed on the basis of the different speed of the left and right wheel:

$$\text{RI}^t = \frac{\sum_i^N [|(RW_i - LW_i) / 7|]}{N} \quad (8)$$

where RW_i and LW_i are the distances covered by the right and left wheel respectively of the robot i in one cycle.

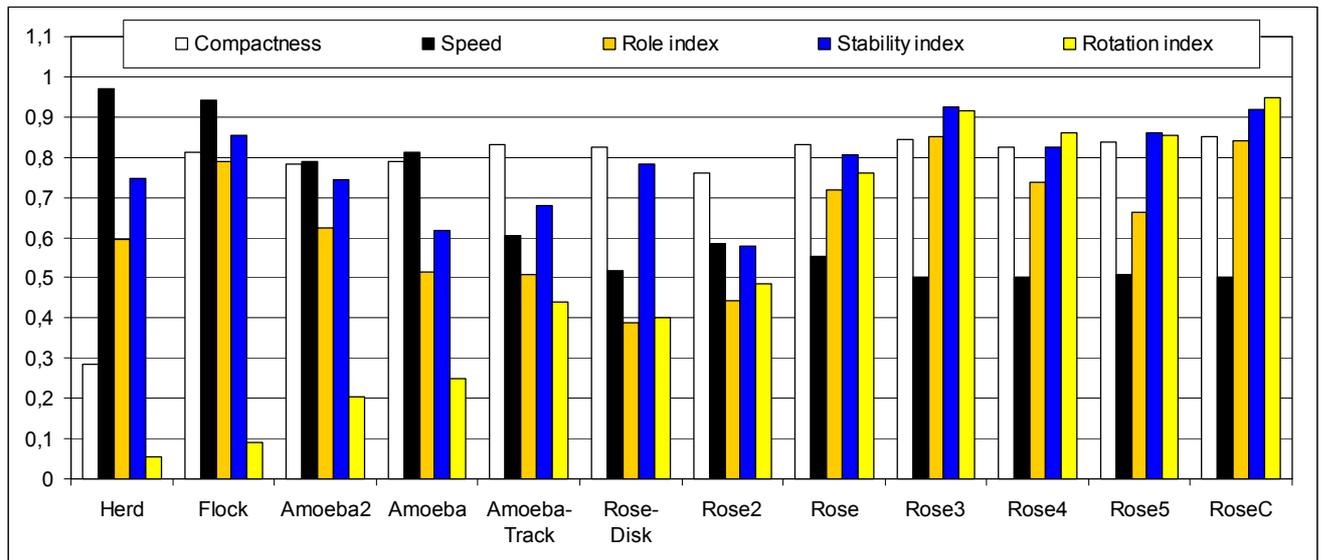


Figure 7: Fitness components and behavioural indexes obtained in the case of the 10 replication of the experiments also described in Figure 4 and in the case of two control experiments named Herd and RoseC in which the W parameter was set to 0.0 and 1.0 respectively.

Figure 7: shows the value of the three indexes and of the two components of the fitness for the best individuals of the 10 replications of the experiment reported in Figure 4 and of two control experiments named *Herd* and *RoseC* in which W was respectively set to 0 (i.e. the fitness has only the speed component) and to 1 (the fitness has only the compactness component). It should be noted that in the first control experiment (*Herd*), the robots go towards the light target without trying to keep close to each other, while in the second control experiment (*RoseC*), robots form a very stable, compact and motionless rose. The value of the indexes and fitness components has been computed only within the test area shown in Figure 1 during five approaching phase of the west light target.

As shown in the Figure, all strategies, and especially Flock and the Rose, are very effective in terms of compactness (GCC) as in the *RoseC* control condition, and are much better than the *Herd* control condition. In general, a high compactness (GCC) is associated with a high stability (GSI), as shown by the stability index. This suggests that, by keeping stable spatial relationships, individual robots can reduce their reciprocal distances.

More significant differences are observed in the capacity to move towards the light target (GSC). From this point of view the Flock strategy significantly outperforms all other strategies and achieves a performance similar to the *Herd* control experiment in which the robots are only asked to approach the light target. Amoeba strategies achieve intermediate results, and Rose strategies rather poor results. Finally it should be noted that the GSC component is correlated with a high group role index (GRI) and a low rotational index (RI). This can be explained by considering that, to maximize the movement toward the light, robots should turn toward the light and start to move straight toward it.

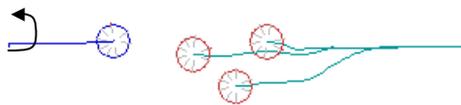
This implies that the two fitness components (GCC and GSC), by requiring the robots to turn toward the other robots and toward the light might interfere between themselves. Apparently, the only way to resolve the interference between the two sub-goals consists in strategies where different individuals play and maintain different functions.

This is what happens in the case of the Flock strategy. In this case, in fact, individuals that are located in the frontal side of the group with respect to the light target do not turn toward the rest of the group but keep their orientation toward the light eventually moving backward to avoid losing the rest of the group. On the contrary individuals located behind tend to turn and move toward the other members of the group. Moreover, once individuals form a compact group and start to move toward the light, each individual try to maintain its current role. The final result is that, once the group has been formed, the individual on the frontal side drive the whole group toward the light and wait if the other individuals get away, while the individuals behind tend to be in shadow and only try to keep close to the individuals at the front so that the whole group remains compact.

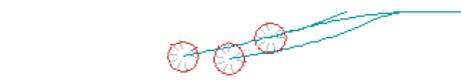
Figure 8 shows how the individuals that display the Flock strategy play different functions in different circumstances. Figure 8a shows how the individual that is closer to the light target assumes and maintains the function of “leader”. Indeed, this individual turns toward the light and waits the rest of the group before driving the entire group toward the light target. After turning toward the light, it also moves backward to speed up the formation of a compact group but, as soon as the rest of the group gets closer, it starts to move toward the light target thus keeping the frontal position with respect to the rest of the group. Figure 8b shows another situation in which the individual that is closer to the light target does not turn toward the rest of the group but keeps its relative position by waiting the rest of the group and by starting to move toward the light as soon as the

rest of the group approach it. Figure 8c shows that individuals that are shadowed (in this case the second robot from left to right) turns and move toward the rest of the group. Finally, Figure 8d shows that couple of individuals located in similar conditions with respect to the light target and to the rest of the group might both assume and maintain the role of leaders or followers.

a) (1000, 500, 180) (1500, 500, 0) (1600, 500, 0) (1700, 500, 0)



b) (1000, 500, 0) (1500, 800, 0) (1600, 800, 0) (1700, 800, 0)



c) (700, 500, 0) (1000, 500, 0) (1300, 500, 0) (1600, 500, 0)



d) (1000, 300, 0) (1100, 300, 0) (1000, 700, 0) (1100, 700, 0)

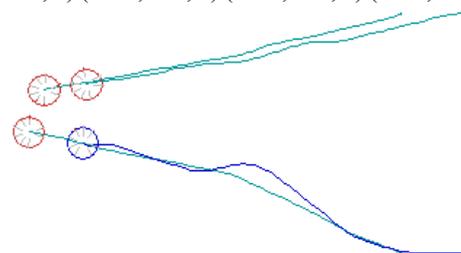


Figure 8: Behaviour displayed by the individuals relying on a Flock strategy initially located in four selected starting positions and orientations. In all cases the light target is located on the west side. The number in parenthesis indicate the initial x and y position and the orientation (0 = west, 180 = east) of each robot. The lines represent the trajectories of the four robot and the circles represent the final position of the robot after a given amount of time. The arrows indicate quick changes of the orientation of individual robots.

The overall result of the ability to display and maintain “situated” specializations is that individuals can quickly form a compact group and then move straight toward the light target. The fact that the Flock strategy requires that

different identical individuals are able to assume and maintain qualitatively different functions may also explain why this strategy, that clearly outperforms Amoeba and Rose strategies (cf. also Figure 4), emerges only in one out of 10 experiments. Indeed we can hypothesize that the Amoeba and Rose strategies, being more simple, might be easier to evolve.

5. Conclusions

We presented a set of experiments in which a group of simulated robots endowed with simple reactive controllers were evolved for the ability to aggregate and to move together toward a light target. The obtained results demonstrate that the evolutionary method, by allowing the synthesis of behaviours that emerge from the interaction between the robots and the environment and between different robots, is a powerful method for synthesizing collective behaviour. Indeed, evolved individuals display a large variety of interesting behavioural patterns in which a group of robots seem to act as a single unit (e.g. by forming a circular structure that rotates around the barycentre of the group and moves toward the light target, or a stable rectangular formation that moves straight toward the light target) and in which the behaviour of the whole group emerges from the interaction between rather simple individual behaviours. The fact that successfully results can be obtained without fine tuning the parameters and/or the fitness function demonstrates that, contrary to what claimed by Zaera et al. (1996), artificial evolution is an effective method for automating the design process of robots able to exhibit collective behaviours. Indeed, explicitly programming the control systems of robots able to show such types of behaviours would be extremely difficult, if not impossible.

By analysing the type of strategies that emerged in different replications of the experiment we observed that all strategies can be categorized within three qualitatively different classes and that the most effective strategy includes primitive forms of “situated” specializations in which identical evolved individuals played different behavioural functions according to the circumstances. Such forms of functional specialization seems to be due to the need to reduce the interference between conflicting sub-goals such as the need to turn and move toward the rest of the group and toward the light target. In future research we plan to investigate more deeply how functional specialization emerges in these experiments and whether more complex forms of specializations can be observed in individuals that are also provided with internal neurons and/or are able to modify their behaviour on the basis of their previous experience.

Acknowledgements

This research has been supported by the European Communities, Information Society Technologies Program (IST),

Future and Emerging Technology (FET), IST-2000-31010. The authors thank Federico Cecconi (National Research Council, ISTC-CNR) and Francesco Mondada (EPFL, Switzerland) for the helpful discussions on how to simulate sounds.

References

- Bonabeau E. et al. (2000). Three-dimensional architecture grown by simple 'stigmergic' agents. *Biosystems*, 56:130-32.
- Martinoli A. (1999). *Swarm intelligence in autonomous Collective robotics: from tools to the analysis and synthesis of distributed control strategies*. PhD Thesis. Lausanne: Computer Science Department, EPFL.
- Mondada F., Franzi, E., Ienne P. (1993). Mobile Robot Miniaturisation: A Tool For Investigation In Control Algorithms. *Proceedings of the Third International Symposium on Experimental Robotics*.
- Nolfi S. (2000). *Evorobot 1.1 User Manual*. Technical Report. Rome: Institute of Psychology (available at <http://gral.ip.rm.cnr.it/evorobot/simulator.html>).
- Nolfi S., Floreano D. (2000). *Evolutionary Robotics*. Cambridge, Mass.: The MIT Press.
- Reynolds C.W. (1993). An evolved, vision-based behavioral model of coordinated group motion. In Meyer J.A., Roitblat H.L., Wilson S.W. (eds.), *From Animals to Animats 2: Proceedings of the second international conference on the simulation of adaptive behaviour*. Cambridge, Mass.: The MIT Press.
- Theraulaz G., and Bonabeau E. (1995). Coordination in Distributed Building. *Science*, 269: 686-688.
- Ward C.R., Gobet F., Kendall G. (2001). Evolving collective behavior in an artificial ecology. *Artificial Life*. Vol. 7, pp. 191-209.
- Werner G.M., Dyer M.G. (1993). Evolution of herding behavior in artificial animals. In Meyer J.A., Roitblat H.L., Wilson S.W. (eds.), *From Animals to Animats 2: Proceedings of the second international conference on the simulation of adaptive behaviour*. Cambridge, Mass.: The MIT Press.
- Zaera N., Cliff D., Bruten J. (1996). *(Not) evolving collective behaviours in synthetic fish*. Technical Report HPL-96-04. Hewlett-Packard Laboratories.