

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)
15 Viale Marx, 00137 Rome, Italy
baldassarre@ip.rm.cnr.it, nolfi@ip.rm.cnr.it, parisi@ip.rm.cnr.it

We present a set of experiments in which simulated robots are evolved for the ability to aggregate and move together toward a light target. By developing and using quantitative indexes that capture the structural properties of the emerged formations, we show that evolved individuals display interesting behavioural patterns in which groups of robots act as a single unit. Moreover, evolved groups of robots with identical controllers display primitive forms of “situated specialization” and play different behavioural functions within the group according to the circumstances. Overall, the results presented in the paper demonstrate that evolutionary techniques, by exploiting the self-organising behavioural properties that emerge from the interactions between the robots and between the robots and the environment, are a powerful method for synthesizing collective behaviour.

The development of autonomous robots through a self-organization process based on artificial evolution is receiving increasing attention. The advantages of this approach become apparent when one wants to control a system governed by complex dynamics. An example of such a system is an autonomous robot interacting with a changing and noisy environment. Another example with even more complex dynamics is a group of robots interacting between them and with the environment. In these cases explicit design encounters two strong limitations. Firstly, the effects of the decisions taken at the level of control may be unpredictable and difficult to understand. Secondly, explicit design may be unable to anticipate and exploit the self-organising properties of the system that are potentially very useful for the solution of the task. Evolutionary methods synthesize behaviour by blindly and randomly varying controllers, and by selecting the successful ones *a posteriori*. Indeed, they select the controllers only after having tested them on the basis of their overall effects and behaviour. Furthermore, whenever the evolutionary process encounters self-organising properties of the system that might be useful to solve the task under focus, it incorporates them in the controllers. For this reason, evolutionary methods may have important advantages when applied to systems with complex dynamics.

In this paper we describe a set of experiments in which groups of simulated robots are evolved for the ability to aggregate and move together toward a light target. The robots use a sound emitter and directional microphones to aggregate and light sensors to locate the light target. As we will see, evolved individuals display interesting behavioural patterns in which groups of robots act as a single unit. Moreover, groups of robots with identical evolved controllers display primitive forms of “situated specialization” in which different individuals play different behavioural functions according to the circumstances.

In the next sections we review the related literature, we present the experimental framework used and the results obtained, and finally we discuss the implications of the results and some future directions of research.

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. Some research adopting evolutionary techniques to evolve collective behaviour has used “artificial creatures”. In his seminal work, Reynolds [8] evolved the control system of a group of simulated creatures placed in an environment with static obstacles and a manually programmed predator for the ability to avoid obstacles and predation. The results showed that coordinated motion strategies began to emerge. In an attempt to study the evolutionary origin of herding, Werner and Dyer [11] co-evolved two populations of predators and prey creatures. By analysing the result of a single run, the author observed that after some generations, during which predators evolved an ability to catch prey, both populations started to create small herds which repeatedly split and formed again. In more recent work, Ward et al. [10] evolved groups of artificial fish, guided by neural networks, which were able to display schooling behaviours. Two populations of predator and prey creatures were evolved in a 2D environment containing randomly distributed food elements. The analysis of the distances between prey, food, and predator suggests that the emergence of schooling behaviour is correlated with the advantage of finding food clumps and protection from predation. Finally, Theraulaz and Bonabeau [9] evolved a population of constructor agents who collectively built a nest structure by depositing bricks according to the perception of the local environment and to a set of behavioural rules (see also [2]). 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 more realistic creatures (robots), however, remains to be ascertained given that these experiments relied 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).

Some pioneering work with robots was conducted by Martinoli [3] who used artificial evolution to synthesize the neural control system of a group of simulated Khepera robots [4] that were asked to find “food items” randomly distributed in an arena. The robots were provided with 10 infrared sensors, two wheels, and two motors controlling them. In some cases the evolved individuals displayed interesting collective behaviours such as exploring the arena in couples. More recently, Quinn et al. [7] evolved simple mobile robots provided with infrared sensors for the ability to move in space while keeping close to each other. This work has been done in parallel with ours [1], and only recently the two research groups have discovered to be working on closely related problems. It is interesting that both research teams have independently converged on a form of genetic algorithm explained below involving team of homogenous robots. The results presented here and those of Quinn et al. demonstrate that, contrary to what claimed by Zaera et al. [12], 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 behaviours would be extremely difficult if not impossible.

The experimental framework

The research presented here aimed at evolving the control system of groups of simulated robots (simbots) able to aggregate and move together toward a light target. The environment is a rectangular arena of 1×2 metres surrounded by walls. In the arena there are two halogen light bulbs of 230 watts each located in the middle of the west and east walls at an height of 1.5 cm from the ground (Fig. 1).

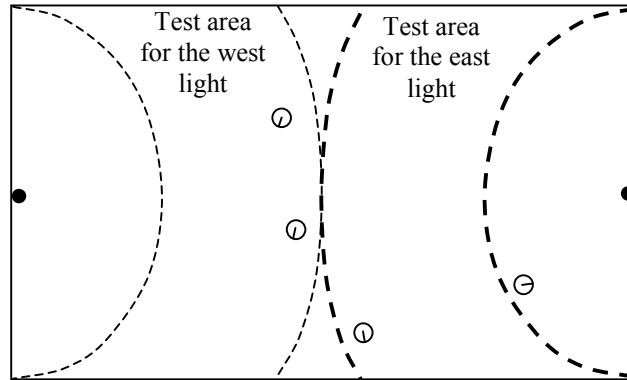


Fig. 1: The environment and the simbots. Lines represent the walls surrounding the arena. Filled circles indicate the two light bulbs. Empty circles represent the simbots. The thin dashed curves indicate the edges of the test area when the west light target is on, while the bold dashed curves indicate the test area when the east light target is on.

Groups consist of four simulated Khepera robots [4] each provided with 8 infrared sensors able to detect walls and other robots up to a distance of 45 mms, 8 ambient light sensors (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 varies within a fixed range (Fig. 2). Microphones and speakers allow simbots to detect each other at a much greater distance than infrared sensors.

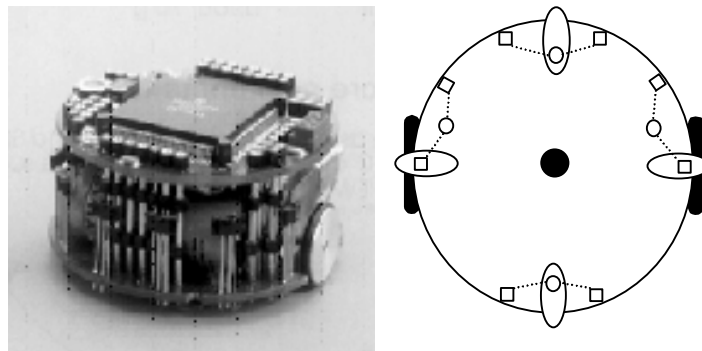


Fig. 2: **Left:** The Khepera robot (sound speaker and microphones not present). **Right:** A schema of the position of sensors and motors in the robot's body. Empty squares and ovals respectively represent the 8 infrared sensors and the 4 directional microphones. The infrared sensors are also used to 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 black filled segments represent the two wheels controlled by the two motors and the black centrally located circle represents the speaker.

Experiments were conducted in simulation by using an extended version of Evorobot [5]. 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 [6]. 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 either the wall or another robot. 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 walls and the other robots. The

same procedure was used to sample the state of the ambient light sensors for different orientations and distances with respect to one of the two light bulbs. These recorded values, together with a geometrical simulation of shadows, were used 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 guarantees 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 orientation and the sound source direction. AF is computed as follows:

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

where α ranges from 0 to 180 degrees and encodes the smaller angle between the orientation of the simulated microphone and the direction of the sound source, and 0.9 is a constant that guarantees 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 simbots, 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 contribution of the sound coming from simbot i of the $N-1$ mates and is 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. At each time step a random value with a uniform distribution that ranges over the interval [-0.05, 0.05] is added to all sensors to simulate sensor noise.

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

Each genotype is used to build 4 identical neural controllers and each controller is used for one of the 4 simbots that together form a group (the unit of selection). Therefore, each group is made up of 4 “clones”. Each group is allowed to “live” for 4 “epochs” (each epoch consists of 1500 cycles and each cycle lasts 100ms). During each cycle, for each simbot: (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.

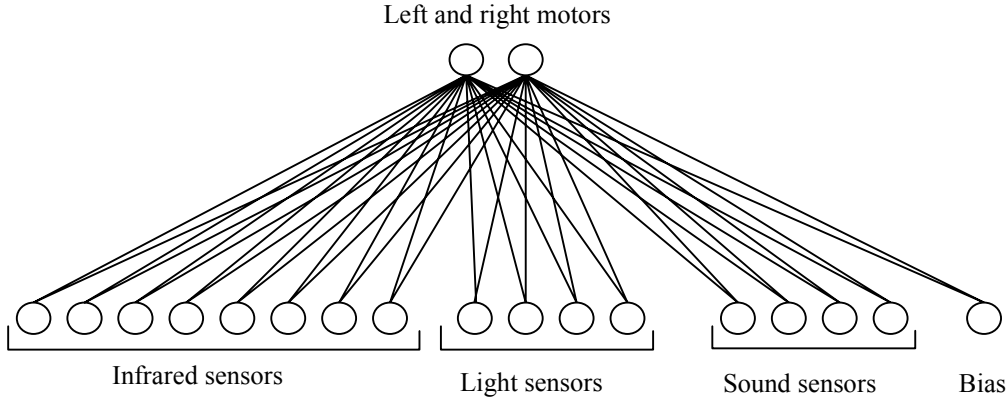


Fig. 3: The neural network controller.

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

To force the simbots to move together toward the light, we designed a fitness function with two components: a group compactness component (GCC) that rewards the simbots for clustering, and a group speed component (GSC) that rewards the simbots for moving fast towards the light target. The two components are computed as follows:

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

where SD_i is the distance of simbot 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 simbot or a group of simbots during one cycle (100 ms). Notice that if the group does not move GSC is equal to 0.5; if the group moves away from the light target GSC is between 0 and 0.5, while if the group approaches the light target GSC is between 0.5 and 1.0.

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

$$F = \sum_t^M [(GCC_t * W) + (GSC_t * (1 - W))] / M \quad (6)$$

where W is a constant in the interval [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 of each generation (each corresponding to the 4 identical neural controllers of one simbots' group) were reproduced by generating 5 copies of each genotype. Each bit of each copy was replaced, with a probability of 2%, by a randomly selected value

(mutation). The evolutionary process lasted 100 generations. The experiment was replicated 10 times by starting with different randomly generated genotype populations.

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. Fig. 4 shows the fitness of the best group of simbots of each replication. In all replications evolved simbots are able to spatially aggregate and to move toward the light target by keeping their aggregation. The sound emitter and sound microphones in each simbot are an important factor for achieving this result. Simulations run without these devices led to results where the simbots explored the arena randomly until they found a team mate and then stayed close to each other using the infrared sensors (results not shown). However, in these simulations the overall performance was rather poor since if subclusters of two simbots formed, the simbots of each subcluster had difficulties to move together to search for other simbots. Another important factor was the use of 4 directional microphones versus a single standard microphone. The single standard microphone forced the simbots to use the shallow sound gradient by moving in wide circles (results not shown), while directional microphones allowed the simbots to locate the direction of the “barycentre” of sounds emitted by the team mates, and to move toward such barycentre following a straight path.

With respect to the ability to move together toward the light target (GSC), significant differences were observed among different replications. The best performance was achieved in the first replication.

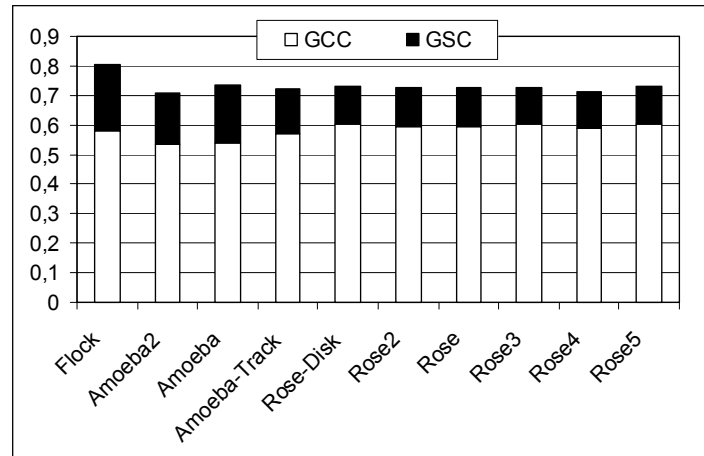


Fig. 4. Performance of the 10 best simbots of the 10 replications of the simulation after 100 generations. Data obtained by testing each individual for 10 epochs. The white and black parts of the histogram respectively represent the contribution of the compactness and speed components. Notice that the relative differences between the two components is mainly due to the value of the W parameter. The names indicate 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 simbots (detailed results not shown), we observed that all emerged formations and strategies could be categorized in three qualitatively different classes that we named *Flock*, *Amoeba*, and *Rose* (Fig. 5).

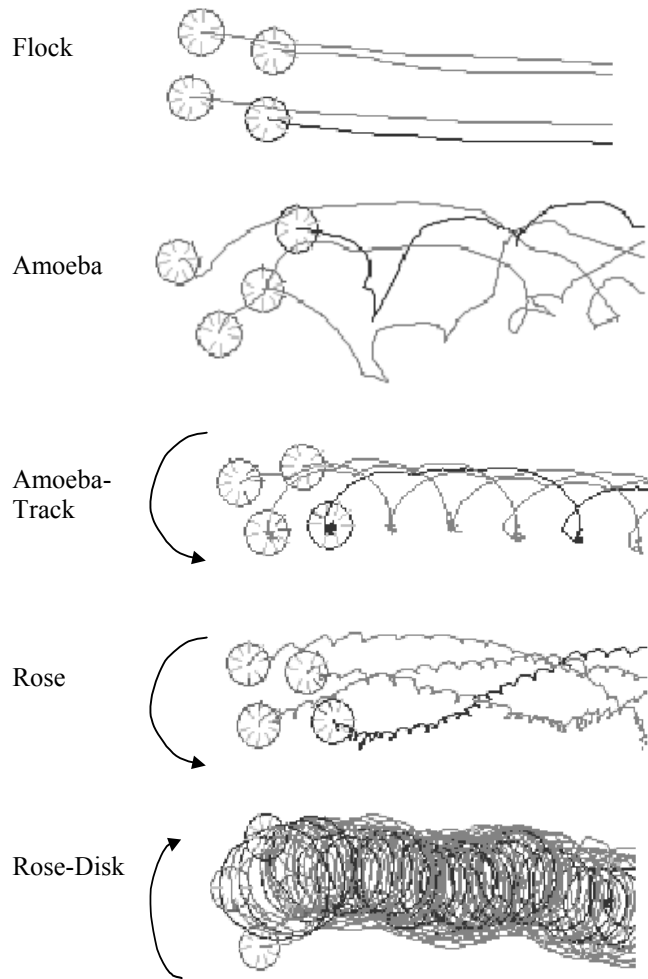


Fig. 5: Three prototypical formations: Flock, Amoeba, and Rose, and some of their variants, Amoeba-Track and the Rose-Disk. The four circles represent the simbots. The lines represent the trajectories of the simbots while the group is approaching the light target (for clarity, the trajectory of one of the four simbots 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. Simbots form a rather compact group that moves straight toward the light target. While the group approaches the target, individual simbots tend to maintain their relative position with respect to each other and with respect to the target (e.g., a simbot that happens to occupy the right-front position in the group at the beginning of the light approaching phase tends to remain in that position). This strategy is the most effective, in particular with regards to the GSC component (the group approaches the light target with maximum speed).

Amoeba. Simbots 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 less effective than the flock strategy given that individual simbots do not move fully straight toward the light). Simbots do not keep the same relative positions while approaching the light target. For example, a simbot that was initially at the rear of the group might speed up and reach the front of the group while the other simbots tend to stay still or to rotate. *Amoeba-Track* is a variation of the amoeba strategy. In this case the group is more compact and the relative movements of the individuals of the group are more regular so that the movement of the whole group resembles that of a track of a caterpillar.

Rose. Simbots form a very compact group and rotate (and eventually they revolve around the group's barycentre) while moving toward the light. The shape of the group tends to vary for groups of different numbers (i.e., typically 4 simbots form a rhombus while 7 simbots form a rose with 6 simbots forming a circle around 1 simbot located in the centre of the circle). *Rose-Disk* is a variant of the rose in which simbots revolve 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 a motor behaviour (turning on itself, and/or turning around the group's barycentre) that interferes with the ability of fast approaching the light.

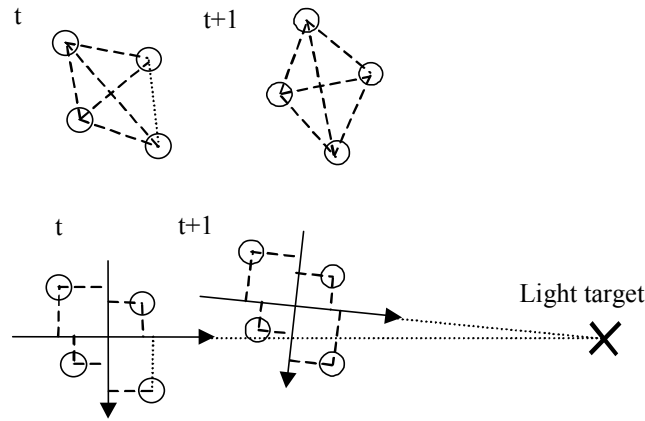


Fig. 6: **Top**: an example of how the group stability index (GSI) is computed. The circles represent the simbots and the dotted lines indicate the distances between each couple of simbots. The index is based on the average variation of distances between time t and $t+1$. **Bottom**: an example of how the role index (GRI) is computed. The circles and the arrows represent the position of the simbots and the Cartesian reference system at lifecycles t and $t+1$. The dotted lines indicate the x and y co-ordinates of each simbot with respect to the reference system. The index is based on the average variation of the simbots' positions indicated by those co-ordinates between time t and $t+1$.

To analyse these strategies quantitatively we designed two statistical indexes that make it possible to characterize the behaviour of single simbots and of simbot groups. The first index, named *group stability index*, measures how much a formation is stable in time, i.e., how much the relative position of each simbot with respect to the other simbots remains identical across time. For example, if the group of 4 simbots forms a square formation, the index will indicate how long the simbots maintain the same geometrical formation and how stable the size of the square remains in time (see Fig. 6, top). The group stability index (GSI) at cycle t is computed as follows:

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

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

where D_{ij}^t is the Cartesian distance between each couple of simbots i and j at time t , N is the number of simbots in the group, $(N * (N-1) / 2)$ is the number of distances between the simbots taken into account by the index, $| |$ is the operator that returns the absolute value of the argument, and 7 is the maximum distance that a simbot can cover in one cycle. Notice that GSI index usually ranges over the interval $[0, 1]$ and only occasionally becomes negative (e.g. when simbots move away from each other at maximum speed).

The second index, named *group role index*, measures how much individual simbots within the group tend to keep the same position and orientation with respect to the other simbots *and* the light target. The index corresponds to the average variation between two successive lifecycles of the x and y coordinates of each simbot with respect to a Cartesian reference system in which the origin corresponds to the barycentre of the group and the y axis corresponds to the line that connects the barycentre to the light target (see Fig. 6, bottom). Moreover, the index is lower if the simbots rotate because rotation implies a change in the orientation with respect to the light target. The group role index (GRI) at cycle t is computed as follows:

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

where $\text{dist}[\]$ is the distance between the x and y position of simbot i at time t and t-1 (notice that x and y coordinates refer to the Cartesian reference system at time t), and 7 is the maximum possible distance covered by one simbot during one cycle. RI^t measures the rotation of the simbots, i.e., how much simbots rotate. It is computed on the basis of the different speed of the left and right wheels:

$$RI^t = \sum_i^N [(RW_i - LW_i) / 7] / N \quad (8)$$

where RW_i and LW_i are the distances respectively covered by the simbot's right and left wheel in one cycle. Notice that the GRI index typically ranges over the interval [0, 1] and only occasionally becomes negative (e.g., when one light is turned off and the other is turned on).

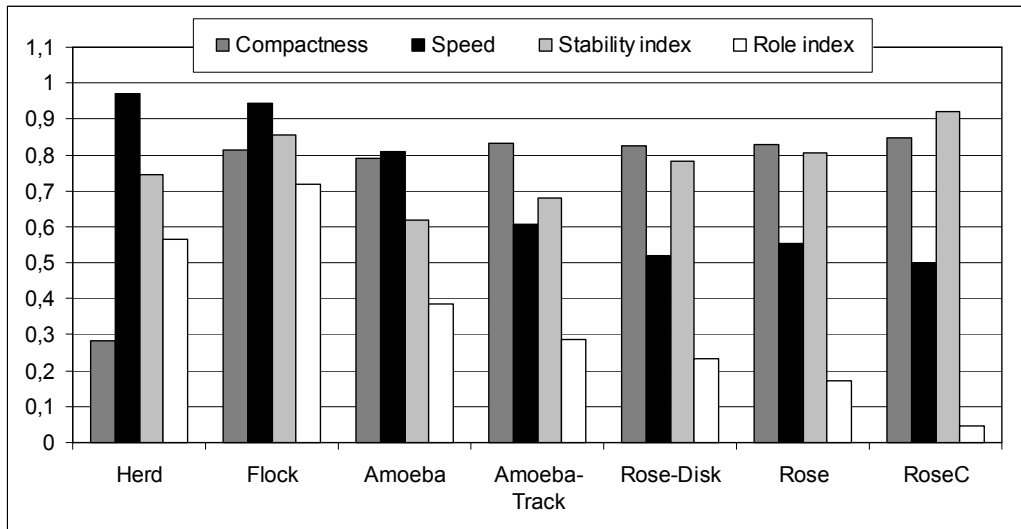
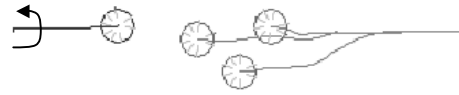


Fig. 7: Fitness components and behavioural indexes obtained in the 10 replications of the experiment described in Fig. 4 and in two control experiments named *Herd* and *RoseC* in which the W parameter was respectively set to 0.0 and 1.0. With regards to speed, a value of 0.5 implies a motionless group, a value below 0.5 corresponds to a group that moves away from the light, while a value above 0.5 corresponds to a group that approaches the light.

Fig. 7 shows the value of the three indexes and of the two fitness components for the best individuals of the strategies reported in Fig. 5 and for two control experiments named *Herd* and *RoseC* in which W was respectively set to 0 (only fitness speed component) and to 1 (only fitness compactness component). In the *Herd* experiment the simbots go toward the

light target without forming a group, while in the RoseC experiment the simbots form a stable, compact but motionless rose. The values reported have been computed within the test area (see Fig. 1; average of 5 passages through the area).

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)

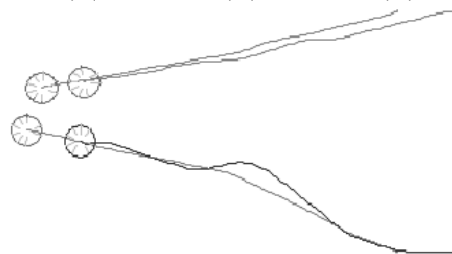


Fig. 8: Behaviour displayed by the simbots relying on a Flock strategy, located in four different starting positions and assigned particular orientations. In all cases the light target is located on the west side. The numbers in parentheses indicate the initial x and y position and the orientation (0 = west, 180 = east) of each simbot. The lines represent the trajectories of the four simbots and the circles represent the final position occupied by the simbots after some time elapsed. The arrows indicate quick changes of the orientation of individual simbots.

Fig. 7 shows that all strategies, and especially Flock and Rose, are as effective, in terms of compactness (GCC), as the RoseC control condition and much better than the Herd control condition. In general, a high compactness (GCC) is associated with a high stability, as shown by the stability index (GSI: in particular, compare the flock and rose versus the amoeba). This suggests that simbots can reduce their reciprocal distances by keeping stable spatial relationships.

Significant differences are observed in the capacity to move towards the light target (GSC) however. The Flock strategy significantly outperforms all other strategies and achieves a performance similar to the Herd control experiment in which the simbots can focus on speed.

Amoeba strategies achieve intermediate results, and Rose strategies rather poor results. The GSC correlates with a high group role index (GRI). In fact to maximize the movement toward the light, simbots have to turn toward the light and move straight toward it.

These results imply that the two fitness components (GCC and GSC), by requiring the simbots to turn both toward the other simbots and toward the light, might interfere with each other. Apparently, the only way to resolve this interference consists in strategies where different individuals play different functions, notwithstanding they have identical control and body, on the basis of the different input pattern they receive from the environment (“situated specialisation”). This is exactly what happens in the case of the Flock strategy. In this case, individuals located in the front of the group with respect to the light (they perceive the sound only from the back) do not turn toward the rest of the group but keep their orientation toward the light, possibly moving backward without turning to avoid losing the rest of the group. On the contrary, individuals in the rear part of the group (that tend to be in the shadow with respect to the light and to perceive the sound with the frontal microphones) turn and move toward the other members of the group. Once simbots form a compact group, and start to move toward the light, each individual tries to maintain its current role: the simbots in front guide the group while the simbots in the back follow those in front while everybody is maintaining a proper distance and orientation with respects to everybody else. When the light target is reached and it is turned off while the other light target is turned on, the formation is disrupted but in few cycles the simbots closer to the light assume the guide role while the simbots behind other simbots with regards to the light assume the follower role (cf. [7] on “role allocation” in a group of robots with identical neural controllers with internal states).

Fig. 8 shows how the simbots that display the Flock strategy play different roles in different circumstances. Fig. 8a shows how the simbot that is closer to the light target assumes and maintains the function of “leader”. Indeed, this simbot 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, the “leader” may occasionally move backward to speed up or maintain 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. Fig. 8b shows another situation in which the simbot that is closer to the light target does not turn toward the rest of the group but keeps its position, waits the rest of the group, and starts to move toward the light as soon as the rest of the group gets closer. Fig. 8c shows that individuals that are in the shadow (in this case the second simbot from left) turn and move toward the rest of the group. Finally, Fig. 8d shows that simbots in couples assume and maintain the roles of leader or follower while approaching the rest of the group.

The positive consequence of “situated specialization” 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 assume and maintain qualitatively different functions may also explain why this strategy, that clearly outperforms the Amoeba and Rose strategies (cf. also Fig. 4), emerges only in one out of 10 experiments. Indeed, we can hypothesize that the Amoeba and Rose strategies, being more simple, are also easier to evolve. To facilitate the emergence of the flock strategy, W should be lowered in order to give a higher weight to the speed component of the fitness. Indeed, in a simulation where W was set to 0.5, the flock strategy emerged 6 times and the amoeba 4 times, out of 10 replications; in a simulation where W was set to 0.8 the rose strategy emerged 7 times, the rose-disk 2 times, and the amoeba 1 time, out of 10 replications.

Conclusions

We have described 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. A simple form of sound communication, used by the robots to locate the team mates, and an algorithm inspired by kin selection that facilitated the evolution of cooperative behaviours, were two key factors for allowing the robots to aggregate and move in a coherent fashion. The results obtained demonstrate that evolutionary techniques, which exploit the self-organising behavioural patterns that emerge from the interaction between the robots and between the robots and the environment, are a powerful method for synthesizing collective behaviour. Indeed, evolved individuals displayed a 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 interactions among rather simple individual behaviours.

By analysing the different types 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 robots with identical evolved controllers play different behavioural functions according to the circumstances. Such forms of functional specialization seem 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 in order to aggregate and maintain aggregation and the need to move toward the light target. In future research we plan to investigate more deeply how functional specializations emerge 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 experience.

Acknowledgements

This research has been supported by the SWARM-BOT project funded by the Future and Emerging Technologies program (IST-FET) of the European Community under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that may be made of data appearing in this publication. The Swiss participants of the project are supported under grant 01.0012 by the Swiss Government. The authors thank Federico Cecconi (ISTC-CNR, Italy) and Francesco Mondada (EPFL, Switzerland) for helpful discussions on how to simulate sound.

References

- [1] Baldassarre, G., Nolfi, S., & Parisi, D. (2002). Evolving mobile robots able to display collective behaviour. In C. K. Hemelrijk (Ed.), *International Workshop on Self-Organisation and Evolution of Social Behaviour* (pp. 11-22). Zurich: Swiss Federal Institute of Technology.
- [2] Bonabeau, E., Guerin, S., Snyers, D., Kuntz, P., & Theraulaz, G. (2000). Three dimensional architectures grown by simple stigmergic agents. *Biosystems*, 56, 13-32.
- [3] 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, Ecole Polytechnique Fédérale de Lausanne.
- [4] Mondada, F., Franzi, E., & Ienne, P. (1993). Mobile robot miniaturisation: a tool for investigation in control algorithms. In T. Yoshikawa & F. Miyazaki (Eds.), *Proceedings of the Third International Symposium on Experimental Robotics*. Berlin: Springer Verlag.
- [5] Nolfi, S. (2000). *Evorobot 1.1 User Manual*. Technical Report. Rome: Institute of Cognitive Sciences and Technologies, National Research Council of Italy.
- [6] Nolfi, S., & Floreano, D. (2000). *Evolutionary Robotics*. Cambridge, Mass.: The MIT Press.
- [7] Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2002). Evolving teamwork and role allocation with real robots. In R. K. Standish, M. A. Bedau & H. A. Abbass (Eds.), *Proceedings of the 8th International Conference on Artificial Life* (pp. 302-311). Cambridge, Mass.: The MIT Press.
- [8] Reynolds, C. W. (1993). An evolved, vision-based behavioral model of coordinated group motion. In J.-A. Meyer, H. L. Roitblat & S. W. Wilson (Eds.), *From Animals to Animals 2: Proceedings of the second international conference on the simulation of adaptive behaviour* (pp. 384-392). Cambridge, Mass.: The MIT Press.
- [9] Theraulaz, G., & Bonabeau, E. (1995). Coordination in Distributed Building. *Science*, 269, 686-688.
- [10] Ward, C. R., Gobet, F., & Kendall, G. (2001). Evolving collective behavior in an artificial ecology. *Artificial Life*, 7, 191-209.
- [11] Werner, G. M., & Dyer, M.G. (1993). Evolution of herding behavior in artificial animals. In J.-A. Meyer, H.L. Roitblat & S.W. Wilson (Eds.), *From Animals to Animals 2: Proceedings of the second international conference on the simulation of adaptive behaviour* (pp. 393-399). Cambridge, Mass.: The MIT Press.
- [12] Zaera, N., Cliff, D., & Bruten, J. (1996). *(Not) evolving collective behaviours in synthetic fish*. Technical Report HPL-96-04. Palo Alto, Ca.: Hewlett-Packard Laboratories.