Coordination and Behaviour Integration in Cooperating Simulated Robots

Gianluca Baldassarre **Domenico Parisi** **Stefano Nolfi**
Institute of Cognitive Sciences and Technologies,
National Research Council,
Viale Marx 15, 00137 Rome, Italy
baldassarre@ip.rm.cnr.it parisi@ip.rm.cnr.it nolfi@ip.rm.cnr.it

Abstract

This paper shows how a group of evolved physically-linked robots are able to display a variety of highly coordinated basic behaviours (coordinated motion, coordinated obstacle avoidance, coordinated light approaching) and to integrate such behaviours into a single coherent behaviour. In this way the group is capable of searching and approaching a light target in an environment scattered with obstacles, furrows, and holes and of dynamically changing its shape in order to pass through narrow passages. The paper analyses in detail the emerged basic behaviours and shows how the coordination of the group relies upon robust self-organising principles based on a traction sensor that allows the single robots to perceive the "average" direction of motion of the rest of the group.

1. Introduction

Consider the scenario shown in Figure 1 and Figure 8. A group of robots is placed in a maze with obstacles, furrows and holes. The mission of the robots is to explore the maze and search for a light target. Some robots form a *swarm-bot* (i.e. a group of physically linked robots with a particular topological structure) in order to pass over furrows and holes in which they would fall by moving alone. The challenges proposed by this scenario are several: how can the assembled robots move in the same direction? How can they avoid obstacles? How can they approach the light once it is in sight? To face these challenges, the robots should be able to display coordinated behaviours (e.g. coordinated movement, coordinated obstacle avoidance, and coordinated light approaching) and to integrate these behavioural capabilities into a single coherent behaviour. This paper will show how these problems can be solved (a) by suitably designing the hardware of the robots and by providing them with a traction sensor that allows them to feel the direction and the intensity of the traction caused by the group, and (b) by developing the controllers of the robots through an evolutionary method that allows evolving robots to exploit behavioural properties that emerge from the fine-grained interactions between the robots and between the robots and the environment.

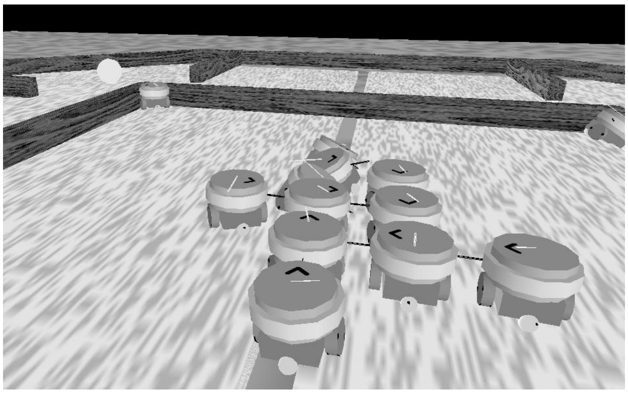

Figure 1: A group of eight assembled robots searches a light target (white sphere at the top left of the picture) in a maze with obstacles, furrows, and holes. Notice how single robots, as the three robots right behind the group, get stuck in holes and furrows.

Collective robotics is a growing research area within the broader field of robotics (Bonabeau et al., 1999; Grabowski et al., 2003). Part of this research area focuses on *distributed coordination* of groups of cooperating robots that have to accomplish common tasks (Baldassarre et al., 2002; Ijspeert et al., 2001). Distributed coordination implies that the characteristics of the collective behaviour exhibited by the group, such as the basic behaviours through which the task is accomplished, the roles played by the different robots, the synchronisation problems raised by their interactions, and so on, are not managed centrally by one or few "leader robots" but are the result of a self-organization process. This self-organization process is the result of the interactions between simple behavioural rules based on local sensory information, followed by the single robots. Depending on the specific tasks to be accomplished, distributed coordination might have some advantages compared to central coordination in terms of (a) resistance to failure of some robots, (b) reduced communication needs, and (c) simplicity of single robots in terms of sensors, actuators, and computational capabilities.

Coordination and behaviour integration is also a central topic in the study of collective behaviour in animals. In social insects, in particular, the interaction between simple behavioural rules followed by each individual might lead to rather complex collective behaviours such as chaining and formation of complex 3D structures (Anderson et al., 2002), collective nest building (Camazine et al., 2001), collective food retrieval (Kube and Bonabeau, 1998). These behaviours tend to rely upon self-organisation mechanisms (Camazine et al., 2001) such as positive feedbacks (e.g.: "do what the majority does") or negative feedbacks (e.g. "perform a given behaviour only while a certain stimulus is above a certain threshold"). Self-organization and behaviour coordination is also important in more complex animals such as non-human primates. For example baboons may use "move grunts" to coordinate the group's transition from resting to moving (Seyfarth and Cheney, 2003). Even humans, though endowed with sophisticated capacities such as language and planning to manage coordination, still heavily exploit distributed coordination. In these cases single individuals might even benefit from distributed coordination while being unaware of it (e.g. consider the coordination of individuals' actions through markets' prices: Von Hayek, 1945). Research studies such as those presented in this paper should produce insights on the general mechanisms of distributed coordination that might underlie all these types of behaviours from different realms.

As mentioned above, in our research evolutionary techniques are used to evolve the behaviours of the robots in simulation. Unfortunately, the time required to run evolutionary experiments can be prohibitively long, especially when detailed simulations of real robots and environments are used as it has been done here. To overcome this problem evolving robots were selected for their ability to solve critical aspects of the problem and then the final task was solved on the basis of the fact that they were able to generalize their abilities to new circumstances. This practical necessity, however, allowed shedding some light on what might be an important general property of nature, that is, searching for simple "building blocks" that can be re-arranged to generate new and more complex behaviours.

Section 2 presents the experimental setup. Section 3 illustrates the results in terms of overall behaviour. Section 4 and 5 respectively analyse the functioning of the basic behaviours and their integration. Section 6 draws conclusions.

2. Experimental setup

The scenario described in the introduction includes some of the problems that are being faced within a research project funded by the European Union ("SWARM-BOTS") which is developing *swarm-bots*, that is groups of fully autonomous robots able to physically connect/disconnect to form a single robotic system that can move and assume different physical shapes in order to solve problems that cannot be solved by single robots. Each robot is provided with its own neural controller and can only have access to local information, so connected or disconnected robots forming swarm-bots should coordinate in order to display coherent behaviours. This research focussed on the problem of how producing coordinated movements in a team of robots that are already assembled. Other research carried out within the project is tackling problems of how robots can self-assemble/disassemble.

Each robot (Mondada et al., in press) has a cylindrical body with a diameter of 116 mm and consists of a mobile base ("chassis") with two drive mechanisms each controlling a track and a teethed wheel (Figure 2), and a main body ("turret") provided with two grippers, one rigid and one flexible, that allow the robots to self-assemble and to grasp objects. Only the first gripper was used in the experiments presented below. The turret has a motor with which it can rotate with respect to the chassis. Each robot is provided with a number of different sensors (Mondada et al., in press), but only the traction sensor described below has been used in the experiments reported in this paper.

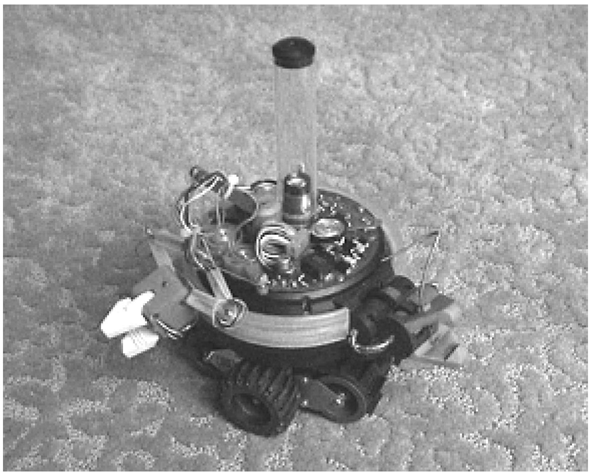

Figure 2: The hardware prototype of an individual robot.

Given that the robots of the project are still under production, a simulator software based on the SDK Vortex™ toolkit (Critical Mass Labs, Canada), which allows programming realistic simulations of dynamics and collision of rigid bodies in 3D, was designed. Given the high computation costs of simulations, the evolutionary experiments were speeded up as follows: (a) only few relevant characteristics of the sensors, actuators and body of the robot were simulated; (b) the size of the robots and the gravitational acceleration coefficient were reduced to have the possibility of increasing the simulation time step without having instabilities; (c) the controller of the robots was evolved in a simplified environment (see Figure 3) and then tested in the more complex environment shown in Figure 1 (for a comparison of the behaviour of simulated and real robots, and for more details on the simulator, see Mondada et al., in press).

The motor system of a robot was modelled by four wheels: two lateral motorised wheels that model the external wheels of the real robot and two spherical passive wheels placed at the front and at the back that stabilise the robot. The four wheels were connected to the chassis, which underpins the rotating turret modelled as a cylinder (Figure 3). The turret was endowed with a simplified gripper which was modelled by creating a physical joint between two robots when needed (Figure 3). The active and passive wheels had a radius of 1.15 and 0.575 cm, respectively. The turret had a radius of 2.9 cm and a height of 2.3 cm.

During evolution, spherical collision models were used for all the wheels and for the chassis, as they required less computations, but equivalent results were obtained by testing evolved controllers with the collision models shown in Figure 3. The gravitational acceleration coefficient was set at 9.8 cm/s^2. This low value, that caused a low friction of the wheels on the ground, was compensated for by setting the maximum torque of the motors at a low value, 70 dynes-centimetre. The friction coefficient was set at 0.6 (Vortex simulates friction according to the Coulomb friction model). The desired speed of the wheels was allowed to vary within the range [–5, +5] radians per second. The desired speed applied to the turret-chassis motor was equal to the difference between the desired speed of the left and right wheel times 0.26. The effect of this was that when the chassis turned the turret turned in the opposite direction, so that its orientation did not change with respect to the environment. This greatly helped the robots to turn the chassis when attached to other robots. The state of the sensors and the motors, and the differential equations used by Vortex to simulate the bodies' dynamics, were updated every 100 ms.

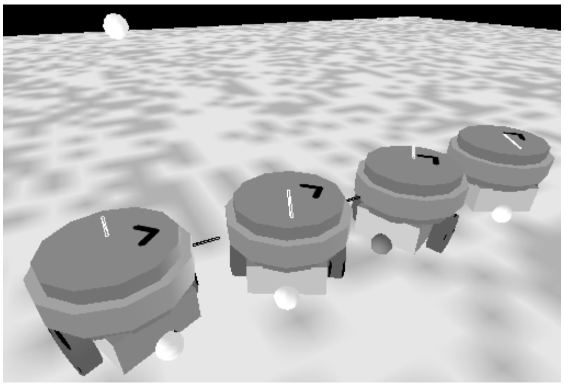

Figure 3: Four simulated robots linked up to form a linear structure. For each robot, the cylinders and the spheres respectively represent the motorised and the passive wheels. The cylindrical structure represents the turret. The arrow indicates the orientation of the turret. The black segment between two robots represents the physical link between them. The white line above each robot, that goes from the robot's centre towards a point on its perimeter, indicates the direction of traction and, with its size, the intensity of traction. The environment consists of a flat terrain and a light source (large white sphere at the top left corner of the picture).

Each robot was provided with a "traction sensor", placed at the turret-chassis junction (Figure 4). This sensor returned the direction (angle with respect to the chassis' orientation) and the intensity of the force of traction (henceforth called "traction") that the turret exerted on the chassis. Traction was caused by the movements of both the connected robots and the robot's own chassis. Notice that, by being assembled, the turrets of the robots *physically integrated* the forces produced by all robots. As a consequence, the traction sensor measured the mismatch between the directions toward which the entire group and the robot's chassis were trying to move. The intensity of the traction measured the size of this mismatch. As it will be shown below, given that robots are physically connected the traction sensor informs robots on what the other robots are doing and thus implements a sort of implicit communication. Traction, seen as a vector, was affected by a 2D noise with a size of ±5% of its maximum length.

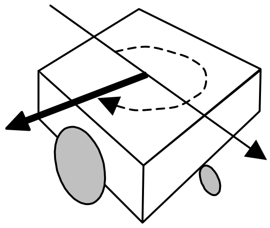

Figure 4: Traction force detected by the robot traction sensor. The parallelepiped represents the chassis. The turret has not been drawn for clarity. The large and small gray circles respectively represents the right motorised wheel and the front passive wheel. The thin arrow indicates the orientation of the chassis, the bold arrow indicates the vector of the traction force that the turret exerts on the chassis, the dotted arrow indicates the angle of the traction.

Each robot was also endowed with four light sensors, positioned on the perimeter of the turret, simulated by using a sampling procedure (see Miglino et al., 1995). A noise of ±5% of the maximum intensity was added to the sensors. Shadows were simulated by computing the geometrical projections of obstacles. In order to provide information about the light gradient with respect to the orientation of the chassis (this greatly eased control since the wheels were connected to the chassis) four virtual light sensors were simulated. These virtual sensors were activated on the basis of the weighted average of the activation of the two closest light sensors, with weights proportional to the angular distance of them from the virtual sensor (Figure 5).

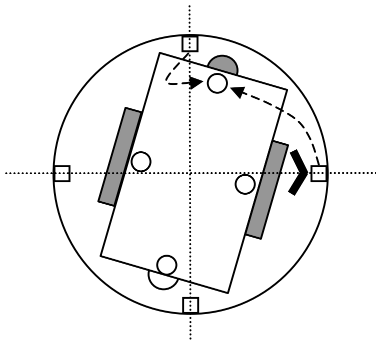

Figure 5: The four empty squares represent the light sensors placed on the turret. The four empty circles represent the virtual light sensors located on the chassis. As an example, the dotted arrows indicate the sensors that were used to compute the activation of the virtual sensor located on the top side of the picture.

Each robot's controller (Figure 6) consisted of a neural network with nine sensory neurons directly connected to two motor neurons. The first four sensory neurons encoded the intensity of the traction from four different preferential orientations with respect to the chassis (rear, left, front and right). Each sensory neuron had an activation proportional to the cosine of the angle between the sensor's preferential orientation and the traction direction when the angle was within [–90, +90] degrees, and zero otherwise. This activation was then scaled by the traction intensity. The next four sensory neurons were activated by the four virtual light sensors described above. The last sensor was the bias unit that was always activated with 1.0. The activation state of the two motor neurons was normalized within [–5, +5] rad/s and was used to set the desired speed of the two corresponding wheels and of the turret-chassis motor.

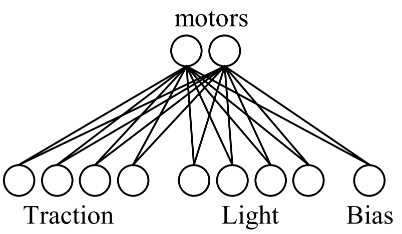

Figure 6: The neural controller of each robot consisted of a two-layer neural network with nine sensory neurons directly connected to two motor neurons.

The connection weights of the neural controllers were evolved (Nolfi & Floreano, 2000). The initial population consisted of 100 randomly generated genotypes that encoded the connection weights of 100 corresponding neural controllers. Each connection weight was represented in the genotype by eight bits that were transformed into a number in the interval [–10, +10]. Each genotype encoded the connection weights of four identical neural controllers, that were used to control the four robots linked up to form the swarm-bot shown in Figure 3. The case of groups of robots with different neural controllers in which different robots might specialize by assuming different roles was not studied here. Although investigating this aspect would have been interesting, the fact that near-optimal performance was obtained without specialization indicates that specialization is not needed to solve the problem discussed here.

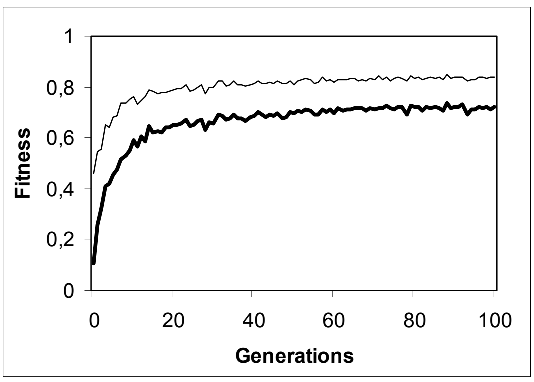

Figure 7: Fitness throughout 100 generations. Thin line: fitness of the best swarm-bot of each generation. Thick line: average fitness of the population. Each line indicates the average performance over 20 replications of evolution.

Each swarm-bot was tested in six "epochs", each lasting 150 time steps of 100 ms each. The light source was present only in three of the six epochs. At the beginning of each epoch the orientations of the chassis of the four robots were randomly assigned and, when present, the light source was placed at a random selected position at a distance of 100 cm from the swarm-bot. The 20 best genotypes of each generation were allowed to reproduce by generating five copies each, with 3% of their bits replaced by a new randomly selected value. The evolutionary process lasted 100 generations. The evolution was replicated 20 times by starting with different initial random genotype populations.

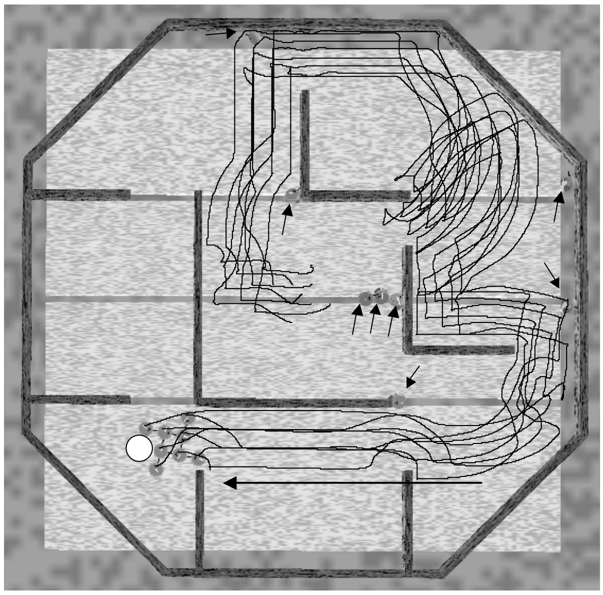

Figure 8: The trajectory produced by a star-shaped swarm-bot in the environment shown in Figure 1. The shape of the swarm-bot is depicted in its final position near the light target represented by the white sphere. The black irregular lines indicate the trajectories followed by the eight robots forming the swarm-bot. While the eight single robots (indicated by short arrows) get stuck in furrows, the swarm-bot passed over them, succeeds to free its robots that fall in the holes near the walls, and searches and finds the light that was not visible from the starting position (centre of graph). As soon as the light is in sight, the swarm-bot reaches it by following a quite direct path (long arrow).

To develop swarm-bots able to explore the environment and to approach the light target when it was in sight, the swarm-bots were selected for the ability to move as fast and as straight as possible when the light was off, and to move toward the light when the light was on. More specifically, the fitness of a swarm-bot was computed by summing the Euclidean distance between the centre of mass of the swarm-bot at the beginning and at the end of the epoch in the three epochs in which the light was off, and the Euclidean distance travelled by the swarm-bot toward the light in the three epochs in which the light was on. To normalize the value of the fitness between [0.0, 1.0], the total fitness of one swarm-bot was computed by dividing the average distance travelled during one epoch by the maximum distance travelled by a single robot moving straight at maximum speed for 150 steps. Notice how the few and short *epochs focussed* only on the two critical aspects of the final task, namely coordinated motion and light approaching, since an evolution directly tackling the final task (see Figure 8) would have required a prohibitive amount of time.

3. The evolved behaviour

Figure 7 shows how the fitness of the best swarm-bot and the average fitness of the population increase throughout the generations. By testing for 100 epochs the swarm-bots of the last generations of the 20 replications of the evolution, it is find that the performance of the best swarm-bot is 0.87 while the average performance is 0.78 with a standard error of 0.07. This means that all evolved robots are able to coordinate so as to move effectively toward the light when the light is on and to move straight when the light is off.

By testing evolved swarm-bots in new conditions (the controllers used are *identical* to those evolved with groups of four robots forming a linear structure) it can be seen that they are able to generalize their coordinated motion and coordinated light approaching ability in rather different circumstances and also to display new interesting behaviours (cf. Baldassarre et al., 2003). In particular, by placing the evolved swarm-bots in the environment shown in Figure 1, it can be seen that they are able to explore the environment avoiding walls and to display a coordinated light approaching behaviour as soon as the light is not shadowed by walls.

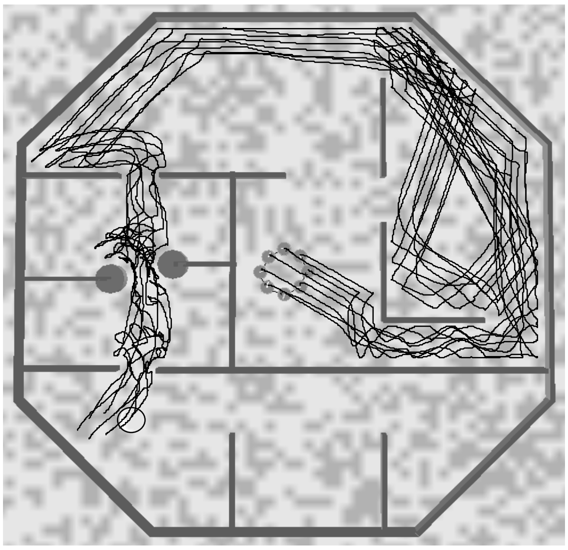

Figure 9: A swarm-bot with a circular shape formed by robots connected through flexible links. The irregular lines, that indicate the trajectories of the single robots, indirectly provide an indication of how the shape of swarm-bot changes while it is moving. With respect to the environment shown in Figure 8, this environment includes additional obstacles and a narrow passage formed by two close cylindrical objects.

Figure 8 shows how the same neural controllers (i.e. the controllers evolved in a team of four robots forming a linear structure) are able to generalize their ability to swarm-bots made up by a *larger number* of robots (eight robots instead of the four used during the evolutionary process) and assembled to form a *different shape* (a star shape instead of a linear shape). This graph and direct observation of behaviour indicate that the robots forming the swarm-bot, whose chassis initially have different orientations, quickly negotiate a common direction of motion and move along such direction by compensating further mismatches arising during the movement. Later on, when one or more of the robots forming the swarm-bot hit a wall, the swarm-bot changes direction by displaying a very effective and coordinated obstacle avoidance behaviour. After avoiding one wall, the swarm-bot keeps moving straight in the arena until it hits another obstacle. The combination of coordinated motion and collective obstacle avoidance behaviour allows the swarm-bot to explore the environment. When finally the swarm-bot sees the light, it reaches it following a straight path.

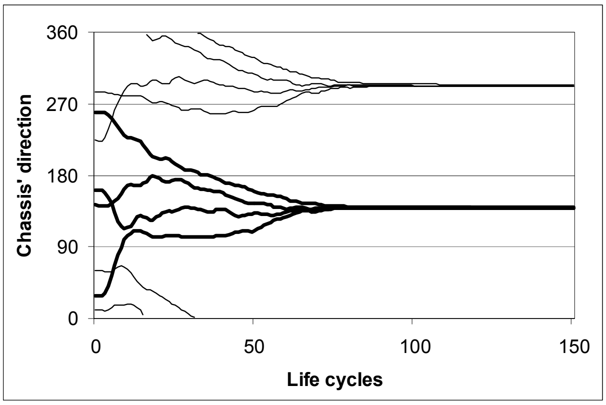

Figure 10: The orientation of the chassis of four robots assembled to form a linear structure placed in the simple environment shown in Figure 3 with the light turned off. The thick and thin lines represent data obtained in two independent tests in which the robots' chassis are initially assigned different random orientations.

Figure 9 shows how the same neural controllers (i.e. the controllers evolved with a group of four robots forming a linear structure through rigid links) are able to generalize their ability to swarm-bots having a circular shape and in which robots are connected through *flexible links* (i.e. links that, at their centre, have a hinge joint with a passive degree of freedom around a vertical axis). Swarm-bots formed by robots connected through flexible links might modify their shape while moving. Interestingly, as shown in Figure 9, the evolved neural controllers display an ability to dynamically adapt the shape of the swarm-bot to the local environmental configuration in order, for example, to pass through narrow passages. As explained in detail below, this result can be explained by considering that the robots that see the light move toward it and, by doing so, produce a traction in that direction that is felt by the other robots that react accordingly. The deformation of the shape of the swarm-bot that allows the swarm-bot itself to pass the narrow passage can be explained by considering the fact that flexible links can bend thus allowing the swarm-bot to change its shape as a result of the movements of the different robots and of collisions. The generalization ability with respect to the shape of the swarm-bot, that in turn depends on how the turrets of the robots are assembled and on the position of the links (in the case flexible links), can be explained by considering that the control system receives only sensory information with respect to the orientation of the chassis.

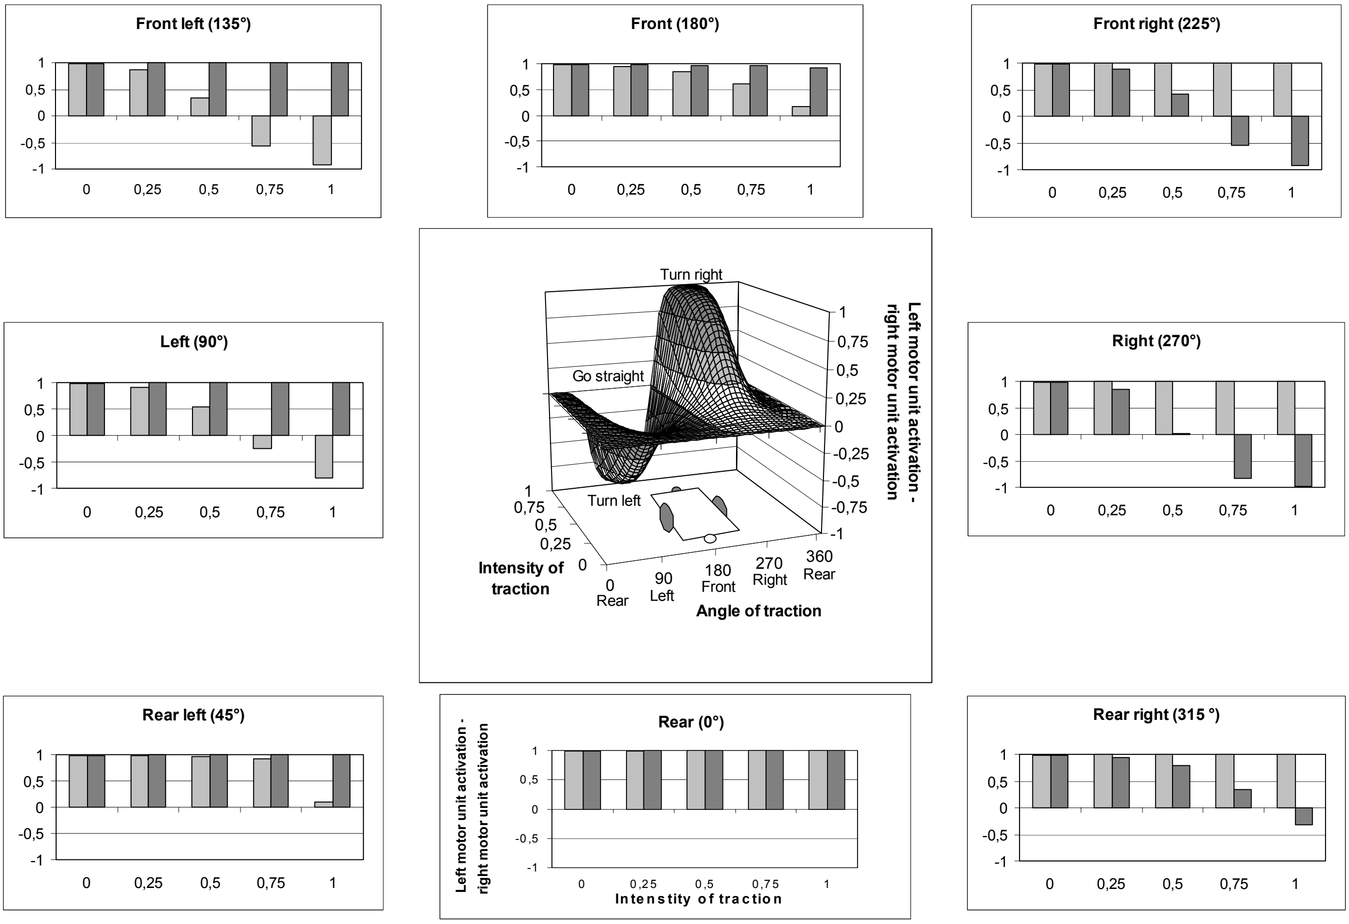

Figure 11: The picture in the centre shows how a robot reacts to a traction with different angles and intensities. The vertical axis indicates the difference between the left and the right motor neurons that, in turn, indicates whether the robot goes straight, turns left or turns right, with different speed. The schematic little picture represent a chassis and should aid the "visualization" of the direction of traction with respect to the chassis itself: the white little wheel represents the rear of the chassis and corresponds to an angle of traction of 0°, measured anticlockwise. The histograms show the desired speed of the left and right wheel as indicated by the activation state of the left and right motor neurons (respectively light grey and dark grey columns of the histograms, measured on the y-axis) for a traction with different angles (see the title of each histogram) and intensities (histograms' x-axis shows four different intensities of traction normalized in [0.0, 1.0]). Data of histograms were obtained with a typical evolved neural controller by manually setting the activation state of the sensory neurons and by measuring the corresponding activation state of the motor neurons (the activation state of the light sensors has been always set to zero).

Above it has been mentioned that, for efficiency reasons, evolution was focussed on some critical aspects (coordinated motion and light approaching) of the final task. When the controllers obtained in *different replications* of the evolution are tested in the final task (see Figure 8 and Figure 9) they display different behaviours. For example, some swarm-bots are more stubborn than others in passing through narrow corridors, some explore the arena very efficiently by doing a sort of collective wall following, some sometimes get stuck in narrows passages when the light is in sight, some other rotate on the spot with some particular initial chassis' orientations (i.e. they reach a stable local minimum in terms of the minimisation of the traction intensity). Such a variety of behaviours produces the possibility for the experimenter to choose the controllers of those replications that better satisfy the requirements of the final task.

To analyse more in detail how the evolved swarm-bots are able to display such complex behaviours and to generalize their ability to such different circumstances the next two sections will describe how different basic behaviours are produced and how they are integrated.

4. Analysis of the basic behaviours

First, it is analysed how evolved swarm-bots perform coordinated motion when the light is off or in shadow. In this condition the robots (a) start to pull/push in different directions, (b) orient their chassis in the direction where the majority of the other robots are pulling, and (c) move straight along the direction that emerges from the initial negotiation by compensating successive mismatches that arise while moving. As shown in Figure 10, the absolute direction that emerges from the robots' negotiation changes in different tests depending on the initial orientation of the robots, but the robots always converge toward a single direction.

By analysing how evolved robots react to traction with different angles and intensities (Figure 11) and by observing the behaviour of the corresponding swarm-bots, the control strategy of a robot might be described as follows:

1) When the traction comes from the front (about 180°), the robot is oriented toward a direction that is close to the mean direction of the other robots. In this situation, when the intensity of the traction is low the robot moves straight. When the intensity is higher (above 0.25, see Figure 11), the robot tends to turn left. However, this latter condition tends to not take place given that when the robots are oriented toward similar directions, the intensity of the traction tends to be low.
2) When the traction comes from the left or the right side (i.e. the angle of the traction significantly differs from 0° and 180°) there is a significant mismatch between the orientation of the robot and the mean orientation of the other robots. In this condition the robot turns toward the direction of traction, that is towards the mean direction of the other robots, by turning left when the traction comes from the left side and right otherwise. The speed of turn is proportional to the intensity of the traction.
3) When the traction comes from the rear (about 0^o) the robot goes straight at maximum speed independently of the intensity of the traction and even if its orientation might significantly differ from the mean orientation of the other robots.

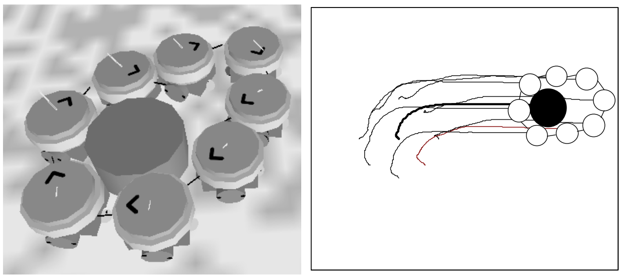

Figure 12: *Left*: Eight robots connected around an object that can move freely. *Right*: the irregular thin and bold lines respectively represent the trajectories left by the robots and the object. The white and black circles respectively represent the final position of the robots and of the object.

Summarising, the ability to display coordinated motion is the result of a *conformist tendency*, that is, a tendency to follow the direction of the traction that provides an indication of the average direction of the other robots (see previous points 1 and 2). Coordinated motion is also a result of a *stubborn tendency*, that is, a tendency to persevere in one direction independently of the intensity of the traction. The stubborn tendency is due to the following factors: (a) the fact that turning toward the direction of the traction takes time (i.e. the mismatch cannot be compensated for instantaneously), (b) the tendency to go straight when the intensity of the traction is low, independently of its direction, and (c) the tendency to go straight when the traction comes from the rear (i.e. from around 0^o, see point 3 above). The stubborn tendency is likely to play an important role in the ability of the robots to keep the equilibrium state once a common direction of movement has emerged from the negotiation and to avoid never-ending negotiation phases in which robots keep changing orientation in order to eliminate small mismatches (in this case the swarm-bot would enter into a limit cycle dynamics without ever reaching a stable state).

Another effect of the stubborn tendency, with particular reference to the tendency to move straight when the traction come from the rear, is that evolved robots spontaneously display coordinated object pushing/pulling behaviours when linked to or around an external object. Figure 12 shows how, after coordinating toward a single direction, robots start to push/pull the object in the direction that emerged from the negotiation notwithstanding the fact that the friction with the ground of the object produces a traction in the opposite direction with respect to the direction of the group's motion.

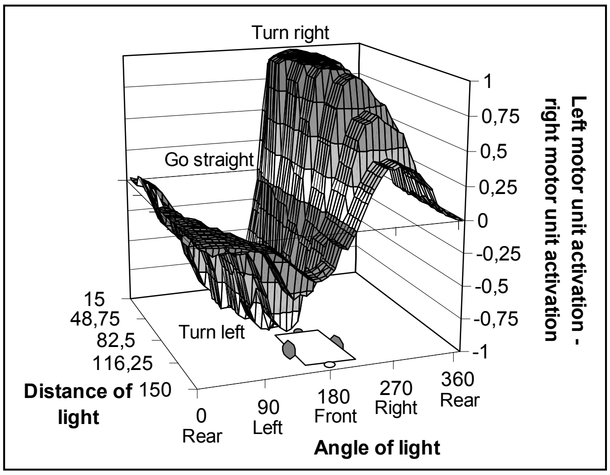

Figure 13: This graph shows how a robot reacts to different orientations and distances with respect to the light. The vertical axis indicates the difference between the left and the right motor neurons that, in turn, indicates whether the robot goes straight, turns left or turns right with different speed. Data have been obtained for a typical evolved neural controller by placing a single robot at different orientations and distances with respect to the light and by recording the corresponding activation state of the motor neurons. The activation state of the traction sensors was always set to zero.

Now a description of how evolved robots spontaneously display an obstacle avoidance behaviour and how swarm-bots consisting of several assembled robots display a coordinated obstacle avoidance behaviour will be presented. When a single robot hits an obstacle, the collision produces a traction with a direction that points away from the obstacle. The robot follows this traction and so avoids the obstacle. When a swarm-bot consisting of several assembled robots hits an obstacle, only a single or few robots collide with it. However, the resistance of the obstacle will propagate through the turrets of these robots to the turrets of the other robots via the links. As a consequence, all the robots will start to turn away from the obstacle, eventually tending to select slightly different directions. Traction will immediately average between this tendencies and will guarantee a synchronous well coordinated turning of the whole swarm-bot. Indeed, direct observation of behaviour shows how the single robots perform a surprisingly highly synchronised turning when the swarm-bot hits an obstacle.

Finally, how evolved swarm-bots display a collective light approaching behaviour is now explained. The fact that individual robots display such behaviour can be explained by analysing how evolved neural controllers react to light sources located at different angles and distances with respect to the robot. As shown in Figure 13, when the traction is null, evolved robots follow the light gradient by turning left or right when the light is respectively located at the left or at the right side of the robot's chassis.

The ability to display a coordinated light approaching behaviour despite the fact that different individual robots forming a swarm-bot have different sensory information (because of their different relative positions and because of shadows) can be explained by considering, once more, the ability to coordinate through the traction sensor and the effects of the motor behaviour of each individual robot on the other assembled robots. Indeed, as soon as a single robot starts to perceive and follow a light gradient turning in the direction of the light, it creates a traction force that is felt by the other robots which, as a consequence, will also turn in the direction of the light. To confirm this, Figure 14 shows an experiment where a single robot provided with light sensors is able to drive a swarm-bot composed of up to twelve robots deprived of light sensors towards a light target. Although performance decreases with the size of the swarm-bot, the experiment shows the power of traction in supporting coordination.

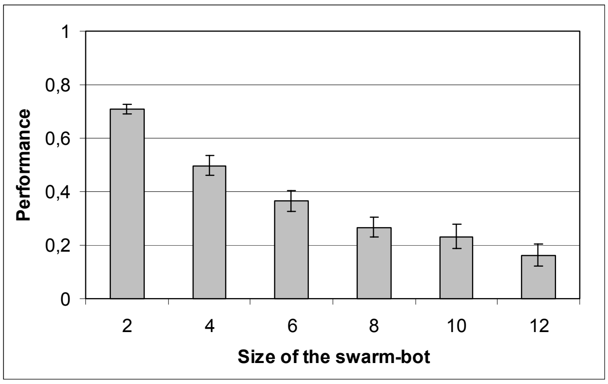

Figure 14: Performance of swarm-bots composed of two to twelve robots forming a linear structure in the simple environment shown in Figure 3 with the light always on. Only one robot forming the swarm-bots is provided with the light sensors while the light sensors of the other robots are always set to zero. Data have been obtained by using the usual neural controller evolved with four robots forming a linear swarm-bot. Columns and bars respectively indicate averages and standard errors over 60 epochs.

This section can be concluded by observing that the coordinated motion behaviour, by relying upon the conformist tendency, is a nice example of coordination based on a mechanism of self-organisation named "positive feedback". This mechanism allows groups of individuals to converge on the same selection, among some possible alternatives, without a centralised decision maker (cf. Camazine et al., 2001).

5. Integration of behaviours

The ability to integrate the different behaviours described in the previous section into a single coherent behaviour results from the combination of three mechanisms.

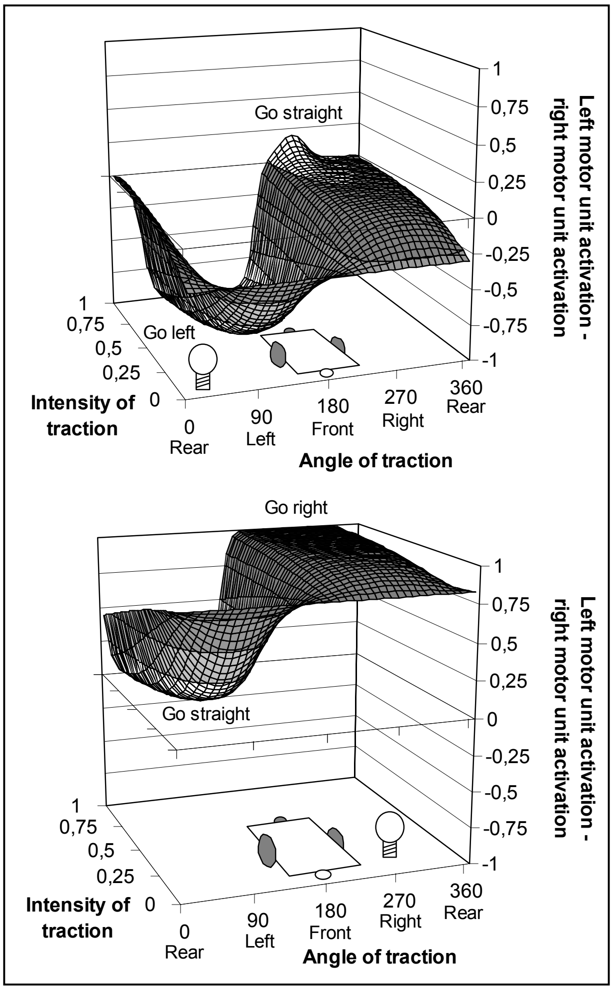

Figure 15: The pictures show how a robot reacts to traction with different direction and intensities when it also perceives a light at a distance of 100 cm coming from the left (90°) or from the right (270°) side (top and bottom pictures, respectively). The y-axis indicates the difference between the left and the right motor neuron that, in turn, indicates whether the robot goes straight or turns left or right with different speed.

As the reader might have already realised from the analyses presented above, one important mechanism is the ability to coordinate through the traction sensor and through the effects that arise from the fact that the robots are physically linked. In fact, as it has been shown previously, the ability to coordinate through the traction sensor and the fact that robots are physically assembled do not only allow swarm-bots to display coordinated motion but also play a crucial role in the ability of the swarm-bot to display coordinated obstacle avoidance and coordinated light approaching behaviours. In fact, the tendency to coordinate through the traction sensors and links between the robots assure that the swarm-bot produces a coherent and coordinated behaviour even when the individual robots have different, incomplete or noisy perceptions. In other words, the ability to coordinate through traction constitutes an important behavioural building block also for the ability to display other more complex behaviours.

The second mechanism, already mentioned above, consists in the fact that the control strategies responsible for the ability to display the different basic behaviours described in the previous section are independent of the particular shape of the swarm-bot. This is due to the fact that the direction of the traction returned by the traction sensor corresponds to the direction of the force that the turret exerts on the chassis *independently* of the particular *turret's orientation*. This characteristic is important not only because it allows the swarm-bot to behave robustly independently of the number of robots, of the way in which robots are assembled, and of possible damages that might affect some individual robots, but also because it allows the swarm-bot to behave robustly even when its shape changes dynamically by adapting to the current situation of the environment.

A third important aspect consists in the fact that the ability to coordinate through the traction sensor and the ability to approach the light and to avoid obstacles integrate without interfering with each other. This can be seen by examining, for example, how the evolved robots react when the information coming from the traction and the light sensors are consistent or inconsistent.

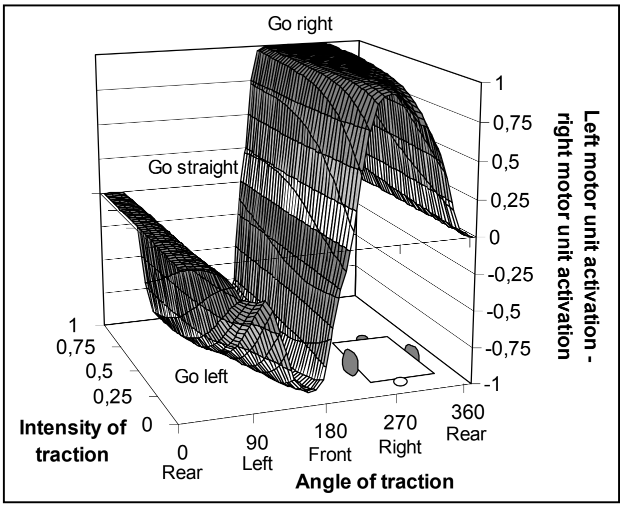

Figure 16: The graph shows the behaviour of a robot when it perceives a traction with different angles and intensities and a light with the same angle as the traction and a fixed distance of one metre.

Figure 15 shows how an evolved robot reacts to traction with different directions and intensities when it also perceives a light at a fixed distance of 100 cm coming from a fixed position. When the light comes from the left side (Figure 15, top graph) the robot turns mostly left by varying the speed of turning on the basis of the intensity of the traction. On the contrary, when the light comes from the right side (Figure 15, bottom graph) the robot turns mostly right by varying the speed of turning on the basis of the intensity of the traction. The combination of these two tendencies allows the robots to approach the light and to maintain the group's coherence at the same time. Figure 15, however, only shows how the control mechanisms responsible for the two behaviours interact from a static point of view. The interaction from a dynamical perspective is even more interesting.

One first important aspect to notice is that the three basic behaviours described above are different forms of *taxis* or *anti-taxis* (i.e. behaviours through which a robot should approach or avoid a stimulation source; notice how the surface graphs of Figure 11 and Figure 13 have similar shapes), and as a consequence in some conditions they can suitably integrate and sum up. A second important aspect is that the behaviour based on traction is a special type of taxis in that it is not related to a stimulus anchored to the environment, as the light pursuing behaviour, but it is related to the average motion direction of the group. This direction tends to change dynamically on the horizontal plane (due to the mechanisms described above) so as to eventually coincide with the direction of the light. When the two directions match completely, the tendency to move toward the average direction of the group and toward the light gradient tend to sum up and to amplify each other as shown in Figure 16.

The possibility of integrating the coordination behaviour based on traction with other behaviours eliminates the need to use behavioural selection mechanisms such as the subsumption architectures (Brooks, 1986) in which only one control mechanism among many is activated at any one time. The possibility of avoiding behaviour selection mechanisms, in turn, eliminates the need of evolving arbitration mechanisms able to select the right behaviour at the right time (e.g. coordinated motion or obstacle avoidance or light approaching). Moreover, it allows evolution to exploit the synergies that might emerge from the interplay between different basic behaviours.

It should be noted that the fact that the control mechanisms which are responsible for the three different basic behaviours cooperate and never interfere with each other is not only due to the characteristics of the basic behaviours themselves but also to the way in which sensory information is encoded in the sensors, to the organization of the robots' body, and to the interaction of these characteristics with the environment. For instance, the possibility of integrating obstacle avoidance and coordinated motion behaviours is due to the fact that robots collide with obstacles with the turret (since the turret is bigger than the chassis) and this generates a traction with a direction opposite with respect to the collision. A second example refers to the obstacle avoidance and the light approaching behaviours. The fact that obstacles shadow the light has the effect that the anti-taxis behaviour related to obstacles never interferes with the taxis behaviour related to light. This means that the possibility of finding simple and robust solutions, from the point of view of the control system of the robots, crucially depends on the way in which sensory and motor information are encoded, on the structure of the robot's body, and on the relation of these with the environment's properties.

6. Conclusions and future work

The paper has described how an evolved group of physically assembled robots can solve problems that could not be solved by single robots, on the basis of simple control mechanisms, a suitable integration of them, and robust coordination mechanisms. In particular, the analysis of the evolved robots indicates that the ability to coordinate through a "traction sensor", that in turn allows exploiting the physical interaction between assembled robots, not only is at the basis of the ability of the robots to display coordinated motion but it also constitutes an important building block for the ability to display other more complex behavioural capabilities.

The possibility of achieving these results can be ascribed to (a) the use an evolutionary technique, and (b) a careful design of the robots' hardware structure. The importance of the robots' hardware structure can be explained by considering that the particular shape of the robots' body, the type and positions of the sensors, and the way in which sensory information is encoded in the neural controllers might crucially affect the complexity of sensory-motor mapping that should be produced by the control systems of the robots and the possibility of integrating different behavioural abilities. These characteristics have been carefully designed and often re-designed on the basis of the results obtained by exploratory experiments performed in simulation. The importance of the latter point can be explained by considering that, as discussed elsewhere (Nolfi and Floreano, 2000; Baldassarre, Nolfi, and Parisi, 2003), artificial evolution allows evolved robots to find solutions that exploit useful behavioural characteristics emerging from the interaction between the control systems of the robots, the structure of their bodies, and the external environment, including the *social* environment made up by other robots.

This work has focussed on the coordination problems faced by some robots that are manually assembled by the experimenter before being tested. Future work will focus on the interesting issue of self-assembling robots, eventually triggered by the challenges of tasks that can be suitably tackled only by assembled robots, for example passing over furrows, passing over high obstacles, linking up around a heavy objects in order to drag them, forming a supporting structure for other robots.

Acknowledgements

This research has been supported by the "SWARM-BOTS" project founded by the Future and Emerging Technologies program (IST-FET) of the European Commission under grant IST-2000-31010.

References

Anderson C., Theraulaz G., Deneubourg J.-L. (2002). Self-assemblages in insect societies. *Insectes Sociaux*, vol. 49, pp. 1-12.

Baldassarre G., Nolfi S., Parisi D. (2002). Evolving mobile robots able to display collective behaviour. In Hemelrijk C. K. (ed.), *International Workshop on Self-Organisation and Evolution of Social Behaviour*, pp. 11-22. Zurich: Swiss Federal Institute of Technology.

Baldassarre G., Nolfi S., Parisi D. (2003). Evolution of collective behaviour in a group of physically linked robots. In Raidl G., Guillot A., Meyer J.-A. (eds.), *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics*, pp. 581-592. Berlin: Springer Verlag.

Bonabeau E., Dorigo M., Theraulaz G. (1999). *Swarm intelligence*. Oxford: Oxford University Press.

Brooks R.A. (1986). A robust layered control system for a mobile robot. *Journal of Robotics and Automation*, vol. 2, p. 14-23.

Camazine S., Deneubourg J.L., Franks N.R., Sneyd J., Theraulaz G., Bonabeau E. (2001). *Self-organization in biological systems*. Princeton, NJ.: Princeton University Press.

Dorigo M., Trianni V., Sahin E., Labella T.H., Gross R., Baldassarre G., Nolfi S., Deneubourg J.-L., Mondada F., Floreano D., Gambardella L.M. (submitted). Evolving self-organising behaviour for a swarm-bot. *Autonomous Robots (Special issue on Swarm intelligence)*.

Grabowski R., Navarro-Serment L.E., Khosla P.K. (2003). An army of small robots. *Scientific American*, November, pp. 42-47.

Ijspeert A.J., Martinoli A., Billard A., Gambardella L.M. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, vol. 11, pp. 149-171.

Kube R.C., Bonabeau E. (1998). Cooperative transport by ants and robots. *Robotics and autonomous systems*, vol. 30, pp. 85-101.

Miglino O., Lund H.H., Nolfi S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, vol. 4, pp. 417-434.

Mondada F., Pettinaro G.C., Guignard A., Ivo K.W., Floreano D., Deneubourg J.-L., Nolfi S., Dorigo M., Gambardella L.M. (in press). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots (Special issue on Swarm intelligence)*.

Nolfi S., Floreano D. (2000). *Evolutionary Robotics*. Cambridge, Mass.: The MIT Press.

Seyfarth R.M., Cheney D.L. (2003). Signallers and receivers in animal communication. *Annual Review of Psychology*, vol. 54, pp. 145-173.

Von Hayek F. (1945). *American Economic Review*, vol. 35, pp. 519-30.