

Evolutionary Swarm Robotics: a theoretical and methodological itinerary from individual neuro-controllers to collective behaviours

Vito Trianni[†], Elio Tuci[†], Christos Ampatzis[‡] and Marco Dorigo^{*}

[†] LARAL-ISTC-CNR, Rome, Italy

[‡] ACT-ESA, ESTEC, Noordwijk, The Netherlands

^{*} IRIDIA-CoDE-ULB, Brussels, Belgium

1 Introduction

In the last decade, swarm robotics gathered much attention in the research community. By drawing inspiration from social insects and other self-organising systems, it focuses on large robot groups featuring distributed control, adaptation, high robustness and flexibility. Various reasons lay behind this interest in similar multi-robot systems. Above all, inspiration comes from the observation of social activities, which are based on concepts like division of labour, cooperation and communication. If societies are organised in such a way in order to be more efficient, then also robotic groups could benefit from similar paradigms.

Constructing tools from a collection of individuals is not a novel endeavor for man. A chain is a collection of links, a rake a collection of tines, and a broom a collection of bristles. Sweeping the sidewalk would certainly be difficult with a single or even a few bristles. Thus there must exist tasks that are easier to accomplish using a collection of robots, rather than just one [23].

A multi-robot approach can have many advantages over a single-robot system. First, a monolithic robot that could accomplish various tasks in varying environmental conditions is difficult to design. Moreover, the single-robot approach suffers from the problem that even small failures of the robotic unit may prevent the accomplishment of the whole task. On the contrary, a multi-robot approach can benefit from the parallelism of operation to be more efficient, from the versatility of its multiple, possibly heterogeneous, units and from the inherent redundancy given by the usage of multiple agents [21].

Swarm robotics pushes the cooperative approach to its extreme. It represents a theoretical and methodological approach to the design of “intelligent” multi-robot systems inspired by the efficiency and robustness observed in social insects in performing collective tasks [8]. Collective motion in fish, birds and mammals, as well as collective decisions, synchronisation and social differentiation are examples of collective responses observed in natural swarms (for some recent reviews, see [9, 15, 12, 30, 11]). In all these examples, the individual behaviour is relatively simple, but the global system behaviour presents complex features that result from the multiple interactions of the system components. Similarly, in a swarm robotics system, the complexity of the group behaviour should not reside in the individual controller, but in the interactions among the individuals. Thus, the main challenge in designing a swarm robotics system is represented by the need to identify suitable interaction rules among the individual robots. In other words, the challenge is designing the individual control rules that can lead to the desired global behaviour.

In the above perspective, self-organisation is the mechanism that can explain how complex collective behaviours can be obtained in a swarm robotics system from simple individual rules. In this context, a complex collective behaviour should be intended as some spatio-temporal organisation in a system that is brought forth through the interactions among the system components. Not every collective behaviour is self-organised, though [9]. The presence of a leader in the group, the presence of blueprints or recipes to be followed by the individual system components clashes with the concept of self-organisation, at least at the level of description in which leader or blueprints are involved. Another condition in which a collective behaviour cannot be considered self-organising is when environmental cues or heterogeneities are exploited to support the group organisation. For instance, animals that aggregate in a warm part of the environment following a temperature gradient do not self-organise. But animals that aggregate

to stay warm, and therefore create and support a temperature gradient in the environment, do self-organise. In both cases, the observer may recognise the presence of some structure (the aggregate) that correlates with the presence of an environmental heterogeneity (the temperature gradient). However, the two examples are radically different from the organisational point of view. Similar natural examples can be easily given also for the presence of leader or blueprints, to show that not every collective behaviour is self-organising [9]. Both the leader or the blueprint can be recognised as the place where the behavioural complexity of the group is centralised. In other words, the complexity of the group behaviour does not result from the multiple interactions among the individual behaviours. Rather, the group behaviour results from a **fixed pattern of interactions** among the system components that is either decided beforehand (in the case of a blueprint) or is centrally and/or continuously re-planned (in the case of a leader). In both cases, there is limited room for adaptiveness to unknown, unpredictable situations resulting from a highly dynamical environment, both physical and social.

The unpredictable nature of the (social) environment makes it difficult to predict in advance, and therefore design, the behavioural sequence and the pattern of interactions that would lead to a certain group behaviour. Moreover, “the adaptiveness of an autonomous multi-robot system is reduced if the circumstances an agent should take into account to make a decision concerning individual or collective behavior are defined by a set of a priori assumptions” [37]. This design problem can be bypassed by relying on Evolutionary Robotics (ER) techniques as an automatic methodology to synthesise the swarm behaviour [33]. In past researches conducted within the SWARM-BOTS project, we experimented with different tasks and defined a methodology that proved viable for the synthesis of self-organising systems. We focused on two particular kinds of self-organising systems: (i) systems that are able to achieve and maintain a certain organisation, and (ii) systems close to a bifurcation point, where robot-robot interactions and randomness lead to one or the other solution. In both cases, the problem is solved without placing any assumption on the kind of interaction pattern that would have been exploited to achieve a certain goal. Even more important, we have shown that determining a priori a certain form of interaction may result in worse performance with respect to an assumption-free setup.

In the rest of the chapter, we present the SWARM-BOTS project’s experience (Section 2), and we discuss in detail some examples of problems studied exploiting the ER approach (Section 3). Then, in Section 4 we speculate on the current limitations of the ER approach, and the future role of ER in the development of more complex behaviours and cognitive abilities for robotic swarms.

2 Swarm robotics and the *swarm-bots*

Even though research in swarm robotics is rather in its infancy, it is quickly developing thanks to the contribution of various pioneer studies [23, 6, 19, 24, 22]. A significant contribution to the field was given by the SWARM-BOTS project, which aimed at the design and development of an innovative swarm robotics platform, the *swarm-bot* [25, 14]. A *swarm-bot* is defined as a self-assembling, self-organising artifact formed by a number of independent robotic units, called *s-bots*. In the *swarm-bot* form, the *s-bots* become a single robotic system that can move and reconfigure. Physical connections between *s-bots* are essential for solving many collective tasks, such as the retrieval of a heavy object. Also, during navigation on rough terrain, physical links can serve as support if the *swarm-bot* has to pass over a hole wider than a single *s-bot*, or when it has to pass through a steep concave region. However, for tasks such as searching for a goal location or tracing an optimal path to a goal, a swarm of unconnected *s-bots* can be more efficient.

An *s-bot* is a small mobile autonomous robot with self-assembling capabilities, shown in Figure 1. It weighs 700 g and its main body has a diameter of about 12 cm. Its design is innovative concerning both sensors and actuators. The traction system is composed of both tracks and wheels—referred to as *treels*—that provide the *s-bot* with a differential drive motion. The wheels are connected to the chassis, which contains the batteries, some sensors and the corresponding electronics. The main body is a cylindrical turret mounted on the chassis by means of a motorised joint, that allows the relative rotation of the two parts. The gripper is mounted on the turret and can be used for connecting rigidly to other *s-bots* or to some objects. The shape of the gripper closely matches the T-shaped ring placed around the *s-bot*’s turret, so that a firm connection can be established. The gripper does not only open and close, but it also has a degree of freedom for lifting the grasped objects. The corresponding motor is powerful enough to lift another *s-bot*.

An *s-bot* is provided with many sensory systems, useful for the perception of the surrounding environment or for proprioception. Infrared proximity sensors are distributed around the rotating turret. Four

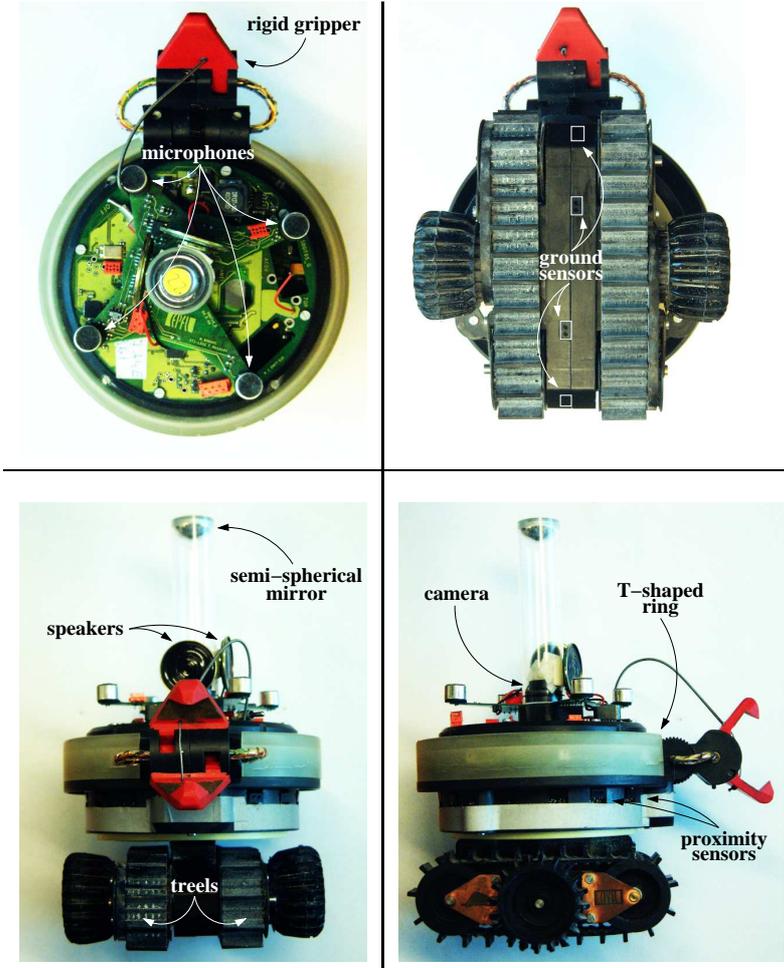


Figure 1: View of the *s-bot* from different sides. The main components are indicated (see text for more details).

proximity sensors placed under the chassis—referred to as *ground sensors*—can be used for perceiving holes or the terrain’s roughness (see Figure 1). Additionally, an *s-bot* is provided with eight light sensors uniformly distributed around the turret, two temperature/humidity sensors, a 3-axis accelerometer and incremental encoders on each degree of freedom.

Each robot is also equipped with sensors and devices to detect and communicate with other *s-bots*, such as an omni-directional camera, coloured LEDs around the *s-bots*’ turret, microphones and loudspeakers (see Figure 1). Eight groups of three coloured LEDs each—red, green and blue—are mounted around the turret. They can be used to emit a colour that can represent a particular internal state of the robot. The colour emitted by a robot can be detected by other *s-bots* using the omni-directional camera, which allows to grab panoramic views of the scene surrounding an *s-bot*. The loudspeaker can be used to emit a sound signal, which can be perceived by the microphones and processed by the on-board CPU. In addition to a large number of sensors for perceiving the environment, several sensors provide each *s-bot* with information about physical contacts, efforts, and reactions at the interconnection joints with other *s-bots*. These include torque sensors on most joints as well as a *traction sensor*, a sensor that detects the direction and the intensity of the pulling force that the turret exerts on the chassis resulting from the forces applied by other connected *s-bots*.

3 Experiments

By exploiting the *swarm-bot* robotic platform, we performed a series of experiments, all characterised by a coherent methodological approach. First of all, evolution was always performed in a simulated environment, which was designed to model the relevant features of the *s-bot*. When required by the experimental setup, the simulation exploited a full 3D physics simulation. This is the case for the

experiments presented in Section 3.1, in which pulling/pushing forces have a fundamental role in the *swarm-bot* behaviour. Otherwise, minimal simulations have been employed. In any case, the evolved controllers have been ported to reality to test the viability of the obtained controllers.

All evolutionary experiments share the same methodological approach as well. The algorithm is run for a fixed number of generation, and works on a single population of genotypes. Each genotype encodes the parameters of a single neural network controller. During evolution, a genotype is mapped into a control structure that is cloned and downloaded in all the *s-bots* taking part in the experiment (i.e., we make use of a homogeneous group of *s-bots*). Each genotype is evaluated over multiple trials. The fitness of a genotype is the average performance computed over the trials in which the corresponding neural controller is tested. The homogeneous group resulting from a single genotype allows to simplify the fitness assignment problem. In fact, a single controller is evaluated and selected for the group performance. This group selection also facilitates the evolution of cooperative strategies, given that there is no competition between different individuals in the group.

In the rest of this section, we present part of the experimental work performed within the SWARM-BOTS project exploiting the ER approach. We present four different experiments: coordinated motion in Section 3.1, synchronisation in Section 3.2, categorisation in Section 3.3 and self-assembly in Section 3.4. In all sections, we first introduce the scenario in which these experiments have been performed, we discuss the experimental setup and finally we draw some conclusions about the lesson learned from the study.

3.1 Coordinated motion and hole avoidance

The scenario For a *swarm-bot* to move coherently, *s-bots* need to negotiate a common direction of motion and maintain the group coordination against external disturbances. The coordinated motion of the assembled structure must take into account the variable number of assembled units, as well as a varying topology. Moreover, the *swarm-bot*'s navigation must be efficient with respect to obstacle and other hazards such as holes and rough terrain, which may be perceived only by a limited subset of the connected *s-bots*.

Coordinated motion has been widely studied in the literature [4, 16, 27, 28]. However, in the *swarm-bot* case, it takes a different flavour, due to the physical connections among the *s-bots*, which open the way to study novel interaction modalities that can be exploited for coordination. The experimental scenario can be summarised as follows: at the beginning of a trial, the *s-bots* start with their chassis oriented in a random direction. Their goal is to choose a common direction of motion on the basis of the only information provided by their traction sensor, and then to move as far as possible from the starting position [5]. In a different set of experiments, the experimental arena presents holes and open borders, in which a *swarm-bot* risks remaining trapped. In this case, *s-bots* must coordinate with the rest of the group to avoid falling [31]. Notice that this task is more difficult than it might appear at first sight. First, the group is not driven by a centralised controller (i.e., the control is distributed). Moreover, *s-bots* cannot use any type of landmark in the environment, such as light sources, or exploit predefined hierarchies between them to coordinate (i.e., there are no “leader robots” that decide and communicate to the other robots the direction of motion of the whole group). Finally, the *s-bots* do not have a predefined trajectory to follow, nor they are aware of their relative positions or about the structure of the *swarm-bot* in which they are assembled. As a consequence, the common direction of motion of the group should result from a self-organising process based on local interactions, which are shaped as traction forces. The problem of designing a controller capable of producing such a self-organised coordination is tackled using feed-forward neural networks synthesised by artificial evolution.

Obtained results As mentioned above, in order to move coordinately *s-bots* can rely only on the traction sensor information, which provides a coarse indication of the average direction of motion of the group. By physically integrating the pulling/pushing forces that the connected *s-bots* produce, the traction sensor provides a compact information that can be exploited for coordination. The problem is therefore designing a controller that would let the group self-organise by interacting through physical forces. The results obtained evolving coordinated motion are extremely interesting [5]. The evolved neural network encodes simple control rules that allow the robots to consistently achieve a common direction of motion in a very short time, and compensate possible misalignments during motion. In general terms, the evolved strategy is based on two feedback loops. Positive feedback makes robots match the average direction of motion of the group, as it is perceived through the traction sensor. Negative feedback makes robots persist in their own direction of motion, but when the traction and motion directions are opposite. Thus the positive feedback allows for a fast convergence towards a common direction of motion, that

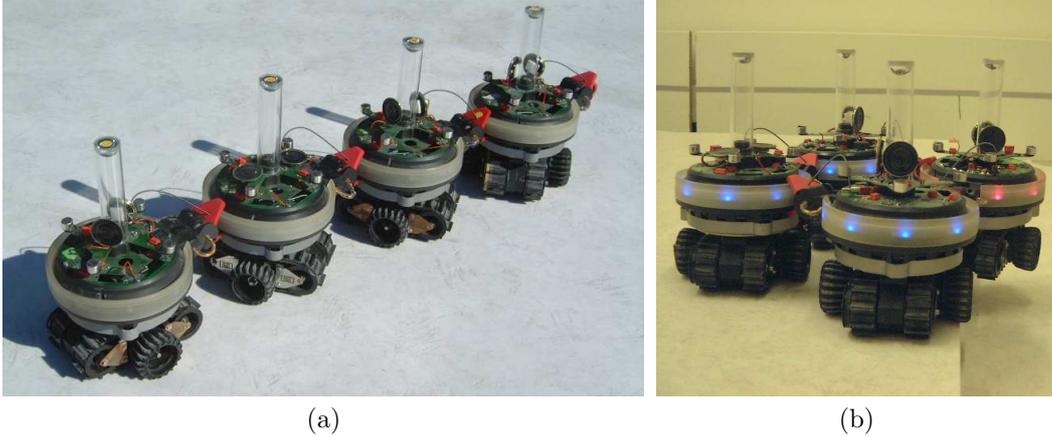


Figure 2: (a) Four real *s-bots* forming a linear *swarm-bot* during coordinated motion. (b) A physical *swarm-bot* while performing hole avoidance. Notice how physical connections among the *s-bots* can serve as support when a robot is suspended out of the arena, still allowing the whole system to work.

is stabilised by the negative feedback loop that avoids deadlock conditions. All this is synthesised in a simple neural network evolved in simulation and tested on real robots (see Figure 2a). The performance of the evolved controllers in terms of robustness, adaptation to varying environmental conditions and scalability to different number of robots and different topologies is striking, demonstrating how evolution synthesised a very efficient self-organising behaviour for coordinated motion [5].

Exploiting a similar setup, we also studied how a *swarm-bot* can navigate in an arena presenting holes or open borders in which the robots risk of remaining trapped [31]. In this case, we investigated how the *swarm-bot* can maintain coordination despite the presence of hazardous situations that are perceived only by a subset of the robots involved. To this purpose, some form of communication may be necessary to the group for a quick reaction. We tested three different communication modalities: (i) direct interactions (DI) through pulling/pushing forces, (ii) direct communication (DC), hand-crafted as a single tone signal emitted as a reflex to the perception of the hazard, and (iii) direct communication in which signalling was controlled by the evolved neural network (evolved communication, EC). In all cases, the motion of the *s-bots* was controlled by a simple perceptron network similar to the one used for coordinated motion. Additionally, *s-bots* could use their sensors for perceiving the presence of holes in the ground. In the DC and EC setups, *s-bots* could also communicate with each other through sound signalling [31].

The obtained results show that it is possible to evolve efficient navigation strategies with each communication paradigm we devised. In the DI setup, when only direct interactions are present, the pulling/pushing forces are sufficient to trigger collective hole avoidance. However, in some cases the *swarm-bot* is not able to avoid falling because the signal encoded in the traction force produced by the *s-bots* that perceive the hazard may not be strong enough to trigger the reaction of the whole group. A different situation can be observed in the DC and EC setup, in which direct communication allows a faster reaction of the whole group, as the emitted signal immediately reaches all the *s-bots*. Therefore, the use of direct communication among the *s-bots* is particularly beneficial in the case of hole avoidance. It is worth noting that direct communication acts here as a reinforcement of the direct interactions among the *s-bots*. In fact, *s-bots* react faster to the detection of the hole when they receive a sound signal, without waiting to perceive a traction strong enough to trigger the hole avoidance behaviour. However, traction is still necessary for avoiding the hole and coordinating the motion of the *swarm-bot* as a whole. We performed a statistical analysis to compare the three different setups we studied, and the results obtained showed that the completely evolved setup outperforms the setup in which direct communication is hand-crafted. This result is in our eyes particularly significant, because it shows how artificial evolution can synthesise solutions that would be very hard to design with conventional approaches. In fact, the most effective solutions discovered by evolution exploit some interesting mechanisms for the inhibition of communication that would have been difficult to devise without any *a priori* knowledge of the system's dynamics [31].

The lesson learned The experiments performed with coordinated motion and hole-avoidance revealed how direct interactions through pulling/pushing forces can be exploited to obtain robust coordination strategies in a *swarm-bot*. The connections among *s-bots* in fact represent an important mean of transfer-

ring information through physical forces. However, exploiting such information is not an easy endeavour if a precise model of the traction sensor is not available. In particular with respect to the synthesis of self-organising behaviours, the top-down approach runs into troubles due to the complex dynamical interactions among the system components that can hardly be predicted or modelled. The evolutionary approach, instead, does not need any precise model of the system. It is sufficient to test potential solutions and to compare their performance on the basis of a user-defined metric. With respect to hand-crafted solutions, the evolutionary approach can achieve a better performance as it can better exploit all system features, without being constrained by *a priori* assumptions. This is clear in the hole avoidance experiments, that show how the hand-crafted reflex signalling, which seemed perfectly reasonable at a first sight, is outperformed by the evolved signalling strategy, which could exploit self-inhibitory mechanisms that are counter-intuitive for a “naive” designer.

3.2 Synchronisation

The scenario An important feature of a swarm robotics system is the coordination of the activities through time. Normally, robots can be involved in different tasks, and higher efficiency may be achieved through the synchronisation of the activities within the swarm. Synchrony is a pervasive phenomenon: examples of synchronous behaviours can be found in the inanimate world as well as among living organisms [29]. The synchronisation behaviours observed in Nature can be a powerful source of inspiration for the design of swarm robotic systems, where emphasis is given to the emergence of coherent group behaviours from simple individual rules. Much work takes inspiration from the self-organised behaviour of fireflies or similar chorusing behaviours [18, 39, 10]. Here, we present a study of self-organising synchronisation in a group of robots based on minimal behavioural and communication strategies [32]. We follow the basic idea that if an individual displays a periodic behaviour, it can synchronise with other (nearly) identical individuals by temporarily modifying its behaviour in order to reduce the phase difference with the rest of the group. In this work, the period and the phase of the individual behaviour are defined by the sensory-motor coordination of the robot, that is, by the dynamical interactions with the environment that result from the robot embodiment. The studied task requires that each robot in the group displays a simple periodic behaviour, which should be entrained with the periodic behaviour of the other robots present in the arena. The individual periodic behaviour consists in oscillations along the y direction of a rectangular arena (see Figure 3). Oscillations are possible through the exploitation of a symmetric gradient in shades of grey painted on the ground. On the other hand, synchronisation of robots movements can be achieved by exploiting a *binary, global* communication: each robot can produce a continuous tone with fixed frequency and intensity. When a tone is emitted, it is perceived by every robot in the arena, including the signalling one. The tone is perceived in a binary way, that is, either there is someone signalling in the arena, or there is no one. This is a very minimal communication system for a swarm of robots, which carries no information about the number of signallers, nor about their position in the environment. No assumption is made on the way the robots should move on the arena, and on the way they should communicate. All the behavioural rules are designed by the evolution of feed-forward neural controllers.

Obtained results We performed 20 evolutionary replications, each resulting in the evolution of efficient synchronisation behaviours. The individual ability to perform oscillatory movements is based on the perception of the gradient painted on the arena floor, which gives information about the direction parallel to the y axis and about the point where to perform a U-turn and move back towards the x axis. The main role of the evolved communication strategy is to provide a coupling between the oscillating *s-bots*, in order to achieve synchronisation: we observed that *s-bots* change their behaviour in response to a perceived communication signal coming from other robots. Recall that the communication signal, being binary and global, does not carry information about either the sender or about its oscillation phase. The reaction to a perceived signal is therefore adapted by evolution to allow the robots reduce the phase difference between their oscillations, eventually achieving synchronous movements. In summary, the evolved synchronisation behaviours are the results of the dynamical relationship between the robot and the environment, modulated through the communicative interactions among robots. No further complexity is required at the level of the neural controller: simple and reactive behavioural and communication strategies are sufficient to implement effective synchronisation mechanisms. To better understand the dynamical relationship between individual sensory-motor coordination and communication, we introduced a dynamical system model of the robots interacting with the environment and among each other [32]. This model offers us the possibility to deeply understand the evolved behaviours, both at the individual

and collective level, by uncovering the mechanisms that artificial evolution synthesised to maximise the user-defined utility function. We assumed an idealised, noise-free and collision-free environment, and we modelled the *s-bot* individual behaviour as it is produced by the evolved neural network. By coupling the individual behaviours through the communication channel, we could study the effects of perturbations through sound signals over the robot oscillations. We analysed the different evolutionary runs performed, and we discovered two alternative mechanisms for synchronisation. With the *modulation mechanism*, *s-bots* synchronise by tuning their oscillatory frequency in response to the perceived communication signal coming from other robots, in order to match the other robots oscillations. This is basically performed by anticipating or delaying the U-turn. With the *reset mechanism*, *s-bots* “reset” their oscillation phase by moving to a particular position over the painted gradient, waiting for the other robots to reach a similar position. Qualitatively, similar mechanisms are also observed in biological oscillators. For instance, different species of fireflies present different synchronisation mechanisms, based on delayed or advanced phase responses.

Besides studying the synchronisation mechanisms, we performed a scalability analysis to test all evolved behaviours with varying group sizes. While scalability is ensured for small groups, we found that physical interactions may prevent the system from scaling to very large number of robots due to the higher probability of performing collision avoidance maneuvers. Still, the evolved synchronisation mechanism scales well if there are no physical interactions. We found that many controllers present perfect scalability, with only a slight decrease in performance due to the longer time required by larger groups to perfectly synchronise. Some controllers, however, present a communicative interference that prevent large groups from synchronising: the signals emitted by different *s-bots* overlap in time and are perceived as a fixed signalling pattern. If the perceived signal does not vary in time, it does not bring information to be exploited for synchronisation. This problem is mainly due to the global and binary communication form, in which the signal emitted by an *s-bot* is perceived by any other *s-bot* anywhere in the arena. Moreover, from the perception point of view, there is no difference between a single *s-bot* and a thousand signalling at the same time. In order to understand the conditions under which this communicative interference take place, we again exploited the mathematical model. We found that scalability can be predicted just by looking at the features of the individual behaviour: the synchronisation behaviour scales to any number of robots provided that an *s-bot* that perceives a communication signal never emits a signal itself. This is a very interesting result, as it directly relates the collective behaviour to the individual one, and indicates which are the building blocks for obtaining scalability in the system under study [32].

The lesson learned The synchronisation experiments show how temporal coordination can be achieved exploiting simple self-organising rules. To this purpose, it is not necessary to provide robots with complex behaviours and time-dependent structures. Instead, we show that a minimal complexity of the behavioural and communicative repertoire is sufficient to observe the onset of synchronisation. Robots can be described as *embodied oscillators*, their behaviour being characterised by a period and a phase. In this perspective, the movements of an *s-bot* correspond to advancements of its oscillation phase. Robots can modulate their oscillations simply by moving in the environment and by modifying their dynamical

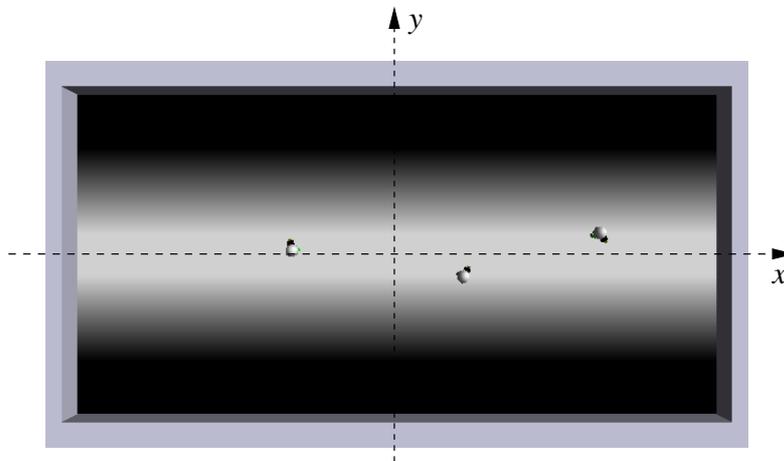


Figure 3: Snapshot of a simulation showing three robots in the experimental arena. The dashed lines indicate the reference frame used in the experiments.

relationship with it. Such modulations are brought forth in response to the perceived communication signals, which also depend on the dynamical relationship between the *s-bot* and the environment.

In this perspective, the dynamical system analysis proved very useful: we introduced a dynamical system model of the robots interacting with the environment and among each other. This model offered us the possibility to deeply understand the evolved behaviours, both at the individual and collective level, by uncovering the mechanisms that artificial evolution synthesised to maximise the user-defined utility function. Moreover, the developed model can be used to predict the ability of the evolved behaviour to efficiently scale with the group size. We believe that such predictions are of fundamental importance to quickly select or discard obtained solutions without performing a time-demanding scalability analysis, as well as to engineer swarm robotic systems that present the desired properties. For instance, the knowledge acquired through the performed analysis could be exploited to improve the experimental setup. We have found that the communicative interferences that prevent the group from synchronising are caused by a communication channel that is neither additive nor local. The locality of communication is for sure an important issue to take into account when studying a realistic experimental setup. Additivity, that is, the capability of perceiving the influence of multiple signals at the same time, is also crucial for self-organising behaviours. We tested the latter issue, and we discovered that it is sufficient to provide the robots with the average signalling activity of the group to systematically evolve scalable behaviours.

3.3 Categorisation, integration over time and collective decisions

The scenario A general problem common to biology and robotics concerns the understanding of the mechanisms necessary to decide whether to pursue a particular activity or to give up and perform alternative behaviours. This problem is common to many activities that natural or artificial agents are required to carry out. Autonomous agents may be asked to change their behaviour in response to the information gained through repeated interactions with their environment. For example, after various unsuccessful attempts to retrieve a heavy prey, an ant may decide to give up and change its behaviour by either cutting the prey or recruiting some nest-mates for collective transport [13]. This example suggests that autonomous agents require adaptive mechanisms to decide whether it is better to pursue solitary actions or to initiate cooperative strategies.

We confronted with the decision-making problem by designing the experimental scenario depicted in Figure 4. Robots are positioned within a boundless arena containing a light source. Their goal is to reach a target area around the light sources. The colour of the arena floor is white except for a circular band, centred around the lamp, within which the floor is in shades of grey. The robots can freely move within the band, but they are not allowed to cross the black edge. The latter can be imagined as an obstacle or a trough, that prevents the robot from further approaching the light. The goal of the experiments is to show that the robots can learn to discriminate between two types of environments. In the first type—referred to as *Env. A*—the band presents a discontinuity (see Figure 4a). This discontinuity, referred to as the *way in zone*, is a sector of the band in which the floor is white. In the second type—referred to as *Env. B*—the band completely surrounds the light (see Figure 4b). The *way in zone* represents the path along which the robots are allowed to safely reach the light in *Env. A*. Successful robots should prove capable of performing phototaxis and of moving over the circular band in search for the *way in zone*, without crossing the black edge. When placed in *Env. A*, the robots should always reach the target area. On the contrary, in *Env. B* the robot should initiate an alternative action, such as signalling or moving away in order to search for other light sources.

Initial experimentation was performed using a single robot controlled by an evolved Continuous Time Recurrent Neural Network (CTRNN) [7]. The results revealed that decision-making can be performed by exploiting a temporal cue: the *Env. B* can be “recognised” by the persistence of a particular perceptual state for the amount of time necessary to discover that there is no *way in zone*. The flow of time, in turns, can be recognised through the integration of the perceptual information available to the robot. This means that the movements of the robot should bring forth the persistence of a certain perceptual condition, and the discrimination can be made only if the latter is maintained long enough.

We repeated the experiments using two robots having the same sensory-motor capabilities [1]. Additionally, robots are provided with a communication system similar to the one used in the synchronisation experiments: they can emit a single frequency tone that is perceived everywhere in the arena in a binary way. The experiments have been performed by varying the initial position of the two robots, and by rewarding them to perform antiphototaxis when placed in *Env. B*. However, no explicit reward was given for communication among the robots. In this way, we aimed at observing whether cooperative, communicative behaviour could emerge or not.

Obtained results Twenty evolutionary simulation runs, each using a different random initialisation, were run for 12000 generations. Thirteen evolutionary runs produced successful groups of robots: both robots approach the band and subsequently (i) reach the *target area* through the *way in zone* in *Env. A*; (ii) leave the band performing antiphototaxis in *Env. B*. The discrimination between the two environments is possible exploiting the integration over time ability of the leaky integrators that form the robot’s neural controller. While moving over the circular band, the *s-bot* accumulates evidence about the absence of the *way in zone*. If the latter is found, the integration over time is stopped and the robot continues in performing phototaxis. If, instead, the *way in zone* is not present, after approximately one loop, the robot leaves the band. This evolved behaviour closely resembles the one obtained with a single robot. However, a closer look reveals that among the thirteen successful groups, nine make use of sound signalling. In particular, signalling strongly characterises the behavioural strategies of the groups when they are located in *Env. B*. In *Env. A* signalling is, for all these groups, negligible. Note that the emission of sound is not demanded in order to navigate towards the target and discriminate *Env. A* from *Env. B*. Indeed, the task and the fitness function do not require the robots to display signalling behaviour. Mechanisms for phototaxis, antiphototaxis, and memory are sufficient for a robot to accomplish the task.

In order to reveal the adaptive significance of sound signalling, further tests have been performed. We looked at the behaviour of the robots that emit sound during a successful trial in each type of environment. We recorded the behaviour of the robots both in a normal condition and in a condition in which the robots cannot hear each other’s sound. In the normal condition we notice that, as soon as one of the robots starts signalling, both robots initiate an antiphototactic movement. On the contrary, when communication signals are blocked, we notice that each robot initiates antiphototaxis only at the time when it starts emitting its own sound. Sound signalling has therefore the function of stimulating antiphototaxis also for those robots that have not yet gathered enough evidence about the absence of the *way in zone*.

These results show that the majority of the successful strategies employ signalling behaviour and communication among the members of the groups. However, communication was not explicitly rewarded: communicating and non-communicating groups could in principle obtain equal fitness. This means that communication may have other functions that influence its adaptive significance. By looking at the behaviour of all successful groups, we discovered that whenever signalling is functionally relevant, it is employed by the robots in *Env. B* as a self-produced perceptual cue. This cue induces the emitter as well as the other robot of the group to change its behaviour from light-seeking to light-avoidance. This evidence constrains our investigation on the adaptive significance of sound signalling to two functions: on the one hand, sound is the means by which a robot emitter switches from phototaxis to antiphototaxis. We refer to this as the “solitary” function. On the other hand, sound is the means by which the robot emitter influences the behaviour of the other robot. We refer to this as the “social” function. From the data we gathered, it appears that signalling is beneficial mainly because of its “social” function.

The selective advantage of signalling groups is given by the beneficial effects of communication with respect to a robust disambiguation of *Env. A* from *Env. B*. The task in fact requires to find an optimal trade-off between speed and accuracy of the decision. The beneficial effect of communication corresponds to robust individual decision-making and faster group reaction, since signaller and hearer react at the

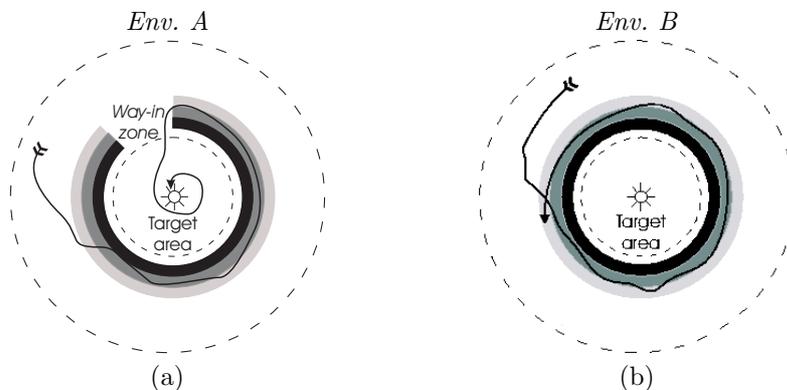


Figure 4: Depiction of the task. (a) *Env. A* is characterised by the way in zone. The target area, centred on the light, is indicated by the dashed circle. (b) In *Env. B* there is no way in zone and the target area cannot be reached. The continuous lines are an example of a good navigation strategy for one robot.

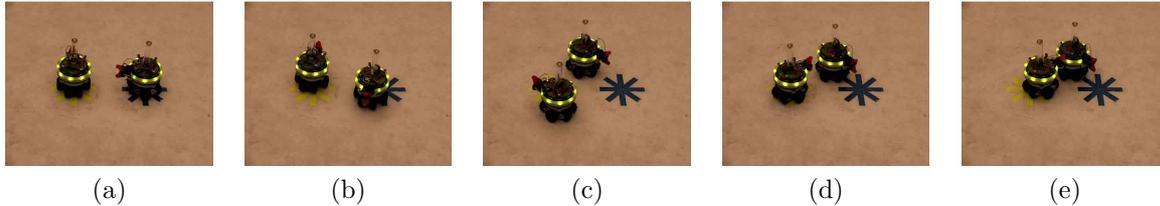


Figure 5: Snapshots from a successful trial. (a) Initial configuration (b) Starting phase (c) Role allocation phase (d) Gripping phase (e) Success (grip)

same time. In fact, a robust individual decision requires longer time spent over the circular band to accumulate evidence of the absence of the *way in zone*, due to the environmental noise that influences the sensors and to the uncertainty of the action outcomes. In total, in those groups in which antiphototaxis is triggered by the perception of sound, a robot which by itself is not ready to make a decision concerning the nature of the environment can rely on the decision taken by the other robot of the group. In average, communication allows the group to accomplish the task earlier, and more reliably. In this way, signalling groups are better adapted to the “danger” of discrimination mistakes in *Env. A* than are non-signalling groups, and thus “early” signalling seems to be an issue that has been taken care of by evolution. In fact, once signalling groups evolve, their signalling behaviour is *refined* by categorising the world later than in the case of non-signalling groups. This happens in order to ensure that the chances of a potential disadvantage resulting from social behaviour are minimised. In other words, the use of communication in a system can also affect aspects of the behaviour not directly related to communication (i.e., the process of integration of inputs over time).

The lesson learned The experiments presented in this section show how individual decision-making and group behaviour can be co-evolved to obtain a robust and efficient system. The need to perform a decision on the basis of information accumulated over time creates a natural trade-off between speed and accuracy. Each *s-bot* has to solve a dilemma: to continue searching for the *way in zone*, or to leave for good? The solution of this dilemma, under normal evolutionary pressures, would tune the individual behaviour to limit the time spent searching to the minimum. However, the introduction of other robots contemporaneously solving the same task, and the possibility of communication, changes the evolutionary dynamics. By exploiting the information gathered by other robots, it is possible to improve the accuracy of the group decision without reducing the decision speed. This is a relevant fact, which justifies the usage of a collective robotics setup even for those conditions in which it is not explicitly required. Additionally, the exploitation of communicative strategies allows to spread acquired information to the group, and to share information retrieval duties among group members: in fact, as soon as communication is in place, the individual behaviour can be refined to exploit the redundancy of the system to the maximum.

3.4 Self-assembly and autonomous role allocation

The scenario Self-assembly is an ubiquitous process in Nature. According to [38], it is defined as “the autonomous organisation of components into patterns or structures without human intervention”. At the nano- or microscopic scale, the interaction among components is essentially stochastic and depends on their shape, structure or chemical nature. Nature also provides many examples of self-assembly at the macroscopic scale, the most striking being animals forming collective structures by connecting to one another. Individuals of various ant, bee and wasp species self-assemble and manage to build complex structures such as bivouacs and ladders [3, 20].

As mentioned in Section 1, the robotics community has been largely inspired from cooperative behaviour in animal societies when designing controllers for groups of robots that have to accomplish a given task. In particular, self-assembly provides a novel way of cooperation in groups of robots. However, it is important to notice that some characteristics of the hardware may impose important constraints on the control of the modules of a self-assembling system. As argued in [36], some hardware platforms consist of morphologically heterogeneous modules, that can only play a predefined role in the assembly process. In others, the hardware design does not allow, for example, the assembly of more than two modules, or requires extremely precise alignment during the connection phase, that is, it requires a great accuracy. The *swarm-bot* platform, thanks to its sensors and actuators and its connection apparatus, does not severely constrain the design of control mechanisms for self-assembly. The lack of hardware constraints

and the homogeneity of the robots requires that self-assembly must be achieved through a differentiation of roles, resulting in the definition of a *s-bot-gripper* (i.e., the robot that makes the action of gripping) and a *s-bot-grippee* (i.e., the robot that is gripped). In work carried out within the SWARM-BOTS project by using other control design techniques than ER, the *s-bot-gripper/s-bot-grippee* differentiation was either predefined (see [17]), or it was based on stochastic events and a complex communication protocol (see [26]). Thanks to the use of ER we designed control strategies for real assembling robots that are not constrained by either morphological or behavioural heterogeneities introduced by the hardware and control method, respectively (see [2] for details). Instead of a priori defining the mechanisms leading to role allocation and self-assembly, ER allowed us to let behavioural heterogeneity emerge from the interaction among the system’s homogeneous components. Moreover, coordination and cooperation in self-assembly between physical robots is achieved without requiring explicit signalling of internal states, as assumed, for example, in [17].

Self-assembly is studied in a scenario in which two *s-bots* are positioned in a boundless arena at a distance randomly generated in the interval [25 cm, 30 cm], and with predefined initial orientations. The robots are required to get closer to each other and to physically assemble through the gripper. The agents perceive each other through their omni-directional camera mounted on the turret, which returns rough information about robot distance and orientation. We also make use of the an optical barrier mounted on the gripper, which informs a robot about the presence of an object between the gripper claws. The agent controller is composed of a continuous time recurrent neural network (CTRNN), whose control parameters are evolved through a rank-based evolutionary algorithm.

Obtained results The results of this work prove that dynamical neural networks shaped by evolutionary computation techniques directly controlling the robots’ actuators can provide physical robots all the required mechanisms to autonomously perform self-assembly. Owing to the ER approach, the assembly is initiated and regulated by perceptual cues that are brought forth by the homogeneous robots through their dynamical interactions. Moreover, in spite of the system being homogeneous, role allocation—i.e., who is the *s-bot-gripper* and who is the *s-bot-grippee*—is successfully accomplished by the robots through an autonomous negotiation phase between the two *s-bots*, as confirmed by our behavioural analyses (see Figure 5). We observed that the role allocation unfolds in time during the entire duration of a trial. Whenever the two robots have different initial perceptions, the role that each *s-bot* assumes can be predicted knowing the combination of the initial relative orientations of the robots. In other words, the combination of relative orientations leads to a pattern of interactions among the robots with a predictable outcome, from the observer point of view. However, a robot has no such information. Perceiving the other robot at a specific distance and orientation does not inform a robot about the role it will assume at the end of the trial. In summary, whenever the initial orientations are asymmetrical, robots engage in a role negotiation phase, and the dynamical system composed of the two interacting robots almost always converge to the same final condition, which depends only on the initial conditions.

In those cases in which the robots start with an identical perception, symmetry does not hinder the robots from autonomously allocating different roles to successfully accomplish their goal. The robots engage in a dynamical interaction which eventually leads to a role assignment. However, in this case it is not possible to predict the outcome of the role allocation process: both robots have 50% probability of assuming the *s-bot-gripper* or the *s-bot-grippee* role. Post-evaluation tests have shown that the random noise inherent in the system is the causal factor that drives the system through sequences of actions that turn out to be successful. In other words, the dynamical system composed by the two interacting robots starts from an unstable equilibrium point, from which it can converge either to one or to the other stable conditions, that is, one of the two alternative role allocations. It is important to notice that the symmetry breaking is performed by exploiting randomness present in the system, which is amplified by the neural controllers as a result of the evolutionary optimisation.

Finally, tests with real robots revealed that the evolved mechanisms proved to be robust with respect to changes in the colour of the light displayed by the LEDs. Furthermore, the self-assembling robotic system designed by using ER techniques exhibits recovery capabilities that could not be observed during the artificial evolution and that were not coded or foreseen by the experimenter [2]. Such a feature in our case comes for free, while in the case of [17] a recovery mechanism had to be designed as a specific behavioural module to be activated every time the robots failed to achieve assembly.

The lesson learned The main contribution of this work lies in the design of control strategies for real assembling robots that are not constrained by either morphological or behavioural heterogeneities introduced by the hardware and control method, respectively. Contrary to the modular or hand-coded

controllers described in [17, 26], the evolutionary robotics approach did not require the experimenter to make any a priori assumption concerning the roles of the robots during self-assembly (i.e., either *s-bot-gripper* or *s-bot-grippee*) or about their status (e.g., either capable of moving or required not to move). We showed with physical robots that coordination and cooperation in self-assembly do not require explicit signalling of internal states, as assumed, for example, by [17]. In other words, we present a setup that requires minimal cognitive and communicative capacities on behalf of the robots. The absence of a priori assumptions allows evolution to exploit the dynamical interaction among the robots to produce an autonomous role allocation mechanism. This can be considered an example of a self-organising system close to a bifurcation point, in which the random fluctuations of the system are amplified to let the system overcome the impasse given by symmetric starting conditions and converge towards a desired solution.

4 Discussion

The experiments presented in Section 3 are representative of a coherent theoretical and methodological approach to the synthesis of self-organising behaviours for a swarm robotics system. What are the limits of this approach? The main problem to deal with is the evolvability of the system related to the scaling in complexity of the collective behaviour. By practising with evolutionary swarm robotics, it appears rather easy to evolve self-organising behaviours in which the system achieves and maintains a certain spatio-temporal pattern. For instance, coordinated motion of the *swarm-bot* or synchronisation are not particularly difficult to evolve (e.g. they require few generations, and successful controllers are almost always obtained), once a suitable experimental setup has been defined (see Section 3.1 and Section 3.2). On the one hand, this is justified by the simplicity of the neural controller and the rather limited number of free parameters that need to be optimised by the evolutionary machinery. On the other hand, the quality of the interactions among the robots contains in itself part of the solution to the self-organisation problem. In the whole, simple controllers and well-defined interactions represent a perfect starting point for the evolution of self-organising behaviour. As a matter of fact, in similar conditions successful behaviours are systematically obtained in all evolutionary runs.

The situation though is slightly different when evolution must produce self-organising systems close to a bifurcation point, in which multiple solutions are possible as a result of the interactions, feedback loops and randomness of the system. This is the case of the categorisation experiment, in which robots had to take a collective decision (Section 3.3), and of the self-assembly experiment, in which complementary roles need to emerge from the robot-robot interactions and the amplification of random fluctuations of the system (Section 3.4). In similar conditions, evolvability is limited by the need to contemporaneously evolve different behavioural traits, and by the presence of multiple stable conditions which create local optima in which evolution may remain trapped. In the experiments we performed, a very large number of generation was necessary to find a suitable solution. Also, the success rate was never close to 100%, and some evolutionary runs resulted in partial solutions of the problem. A similar problem is experienced with the evolution of communication, which requires the evolution of both the signal and the response to the signal, which individually may be counter-adaptive or neutral with respect to the devised fitness function (see Section 3.3).

The experiments presented in Section 3.3 are interesting also from a different point of view, that is, the influence that the individual behaviour has on the evolution of the group behaviour. Here, we can distinguish between two different organisational levels: (i) the individual level, in which sensory-motor coordination and integration over time support the decision-making, and (ii) the collective level, in which information spreading through communication leads to increased group efficiency. We believe that future directions in evolutionary swarm robotics should focus on systems characterised by *multiple levels of organisation*. More complex self-organising behaviours can be obtained through a layered evolution, that proceeds through individual sensory-motor coordination, individual categorisation abilities, communication and exploitation of the social environment, aiming at some collective intelligence. As experienced in our experiments, each different level of organisation is supported by the lower levels, and in turns influences their dynamics. In a swarm robotics scenario, the influences of higher organisational level on the lower ones could be exploited to simplify the individual behaviour in favour of more robust, collective solutions. Brought to the limit, each robot in the swarm could behave as a neuron-like device that can move in the environment and interact, physically or through communication, with neighbouring robots, while the swarm brings forth complex processes as a whole. In this respect, we believe that the cognitive abilities of swarms should be studied and compared with those observed in the vertebrate brain, in the attempt to find the common mechanisms that underly cognition. In this respect, robotics models of

swarm behaviour may represent extremely powerful tools for the study of swarm cognition [34].

Another possible direction in the study of evolutionary swarm robotics concerns the exploitation of *heterogeneous swarms*, in which different types of robots are organised in swarms, which cooperate for a collective goal. We are currently investigating swarms of heterogeneous robots within the project Swarmanoid,¹ in which three types of robots are studied: *eye-bots*, *foot-bots* and *hand-bots*. Eye-bots are robots specialised in sensing and analysing the environment from a high position to provide an overview that *foot-bots* or *hand-bots* cannot have. Eye-bots fly or are attached to the ceiling. Hand-bots are specialised in moving and acting in a space zone between the one covered by the foot-bots (the ground) and the one covered by the eye-bots (the ceiling). Hand-bots can climb vertical surfaces. Foot-bots are specialised in moving on rough terrain and transporting either objects or other robots. They are based on the *s-bot* platform, and extend it with novel functionalities. The combination of these three types of autonomous agents form an heterogeneous swarm robotic system that is capable of operating in a 3D space.

Generally speaking, dealing with heterogeneity in a collective robotics setup often leads to specialisation and team work: the task is decomposed on the basis of the different robots available, and roles are assigned correspondingly. With heterogeneous swarms, the redundancy of the system opens the way to various scenarios. On one extreme, the classical scenario accounts for different swarms that specialise to particular sub-tasks, and that are loosely coupled. For instance, a swarm of eye-bots is responsible of locating areas of particular interest, for instance areas that contain objects to be retrieved. The eye-bots direct the action of a swarm of foot-bots that collectively retrieve such objects. On the other extreme, robots can form a swarm of homogeneous entities, where each entity is a small heterogeneous, tightly cooperating team. For instance, two or three foot-bots can self-assemble to transport a single hand-bot, therefore creating a small team, which can coordinate its activities within a swarm of similar foot-bot/hand-bot teams. Between these two extreme scenarios, there can be an infinite blend of possibilities for cooperating heterogeneous swarms. In this respect, ER can give a strong contribution to define the individual behaviours, and shape the self-organisation of the heterogeneous swarm. In particular, ER can be exploited to define the behaviour of the heterogeneous robots by evolving one controller for each robot type. An alternative, interesting scenario consists in synthesising homogeneous controllers for heterogeneous robots, in which the controller adapts to the dynamics of the robot on which it is downloaded without *a priori* knowledge of its type. We recently performed preliminary studies by evolving controllers for a heterogeneous group of three simulated robots [35]. The agents are required to cooperate in order to approach a light source avoiding collisions. The robots are morphologically different: two of them are equipped with infrared sensors, one with light sensors. Thus, the two morphologically identical robots should take care of obstacle avoidance, while the other one should take care of phototaxis. Since all the agents can emit and perceive sound, the group's coordination of actions is based on acoustic communication. The results of this study are a "proof-of-concept": they show that dynamic artificial neural networks can be successfully synthesised by artificial evolution to design the neural mechanisms required to underpin the behavioural strategies and adaptive communication capabilities demanded by this task. Thus, ER represents a promising method that should be considered in future research works dealing with the design of homogeneous controllers for groups of heterogeneous cooperating and communicating robots.

In conclusion, based on the results obtained in past research and on the perspective of future achievements, we believe that the bidirectional influence arrow connecting ER and swarm robotics can be enforced in both directions. ER can offer to swarm robotics a bias-free method to automatically obtain robust and sophisticated control structures that exploits aspects of the experimental setup not always *a priori* evident to the experimenter. Equally, swarm robotics can broaden the horizons of ER beyond the current limits. In our opinion, the swarm cognition approach, as well as studies with heterogeneous swarms, are two of the most promising directions.

References

- [1] C. Ampatzis, E-Tuci, V. Trianni, and M. Dorigo. Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behaviour*, 16(1):5–26, 2008.
- [2] C. Ampatzis, E. Tuci, V. Trianni, A.L. Christensen, and M. Dorigo. Evolving autonomous self-assembly in homogeneous robots. *Artificial Life*, 15(4):465–484, 2009.

¹Swarmanoid is a project funded by the Future and Emerging Technologies program of the European Community, under grant IST-022888. See <http://www.swarmanoid.org>

- [3] C. Anderson, G. Theraulaz, and J.-L. Deneubourg. Self-assemblages in insect societies. *Insectes Sociaux*, 49(2):99–110, 2002.
- [4] T. Balch and R. C. Arkin. Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, 1998.
- [5] G. Baldassarre, V. Trianni, M. Bonani, F. Mondada, M. Dorigo, and S. Nolfi. Self-organised coordinated motion in groups of physically connected robots. *IEEE Transactions on Systems, Man and Cybernetics - Part B: Cybernetics*, 37(1):224–239, 2007.
- [6] R. Beckers, O. E. Holland, and J.-L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. A. Brooks and P. Maes, editors, *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems (Artificial Life IV)*, pages 181–189. MIT Press, Cambridge, MA, 1994.
- [7] R. D. Beer. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215, 1995.
- [8] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY, 1999.
- [9] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- [10] A. L. Christensen, R. O’Grady, and M. Dorigo. From fireflies to fault-tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation, Special Issue on Swarm Intelligence*, 13(4):754–766, 2009.
- [11] I.D. Couzin. Collective minds. *Nature*, 455(7129):715–715, 2007.
- [12] I.D. Couzin and J. Krause. Self-organization and collective behavior of vertebrates. *Advances in the Study of Behavior*, 32:1–75, 2003.
- [13] C. Detrain and J.-L. Deneubourg. Scavenging by *pheidole pallidula*: a key for understanding decision-making systems in ants. *Animal Behaviour*, 53:537–547, 1997.
- [14] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a *swarm-bot*. *Autonomous Robots*, 17(2–3):223–245, 2004.
- [15] N.R. Franks, S.C. Pratt, E.B. Mallon, N.F. Britton, and D.J.T. Sumpter. Information flow, opinion-polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357(1427):1567–1583, 2002.
- [16] J. Fredslund and M. J. Matarić. A general algorithm for robot formations using local sensing and minimal communication. *IEEE Transactions on Robotics and Automation*, 18(5):837–846, 2002.
- [17] R. Groß, M. Bonani, F. Mondada, and M. Dorigo. Autonomous self-assembly in swarm-bots. *IEEE Transactions on Robotics*, 22(6):1115–1130, 2006.
- [18] O. Holland and C. Melhuish. An interactive method for controlling group size in multiple mobile robot systems. In *Proceedings of the 8th International Conference on Advanced Robotics (ICAR ’97)*, pages 201–206. IEEE Press, Piscataway, NJ, 1997.
- [19] O. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5(2):173–202, 1999.
- [20] B. Hölldobler and E. O. Wilson. The multiple recruitment systems of the african weaver ant, (*Ecophylla longinoda* (latreille) (hymenoptera: formicidae). *Behavioural Ecology and Sociobiology*, 3:19–60, 1978.
- [21] C. V. Jones and M. J. Matarić. Behavior-based coordination in multi-robot systems. In S. S. Ge and F. L. Lewis, editors, *Autonomous Mobile Robots: Sensing, Control, Decision-Making, and Applications*, pages 549–569. CRC Press, Boca Raton, FL, 2006.

- [22] M. J. B. Krieger, J.-B. Billeter, and L. Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature*, 406:992–995, 2000.
- [23] C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–219, 1993.
- [24] A. Martinoli, A. J. Ijspeert, and F. Mondada. Understanding collective aggregation mechanisms: From probabilistic modeling to experiments with real robots. *Robotics and Autonomous Systems*, 29:51–63, 1999.
- [25] F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221, 2004.
- [26] R. O’Grady, R. Groß, F. Mondada, M. Bonani, and M. Dorigo. Self-assembly on demand in a group of physical autonomous mobile robots navigating rough terrain. In *Proceedings of the 8th European Conference on Artificial Life (ECAL’05)*, number 3630 in LNCS, pages 272–281. Springer Verlag, Berlin, Germany, 2005.
- [27] M. Quinn, L. Smith, G. Mayley, and P. Husbands. Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences*, 361(1811):2321–2343, 2003.
- [28] L. Spector, J. Klein, C. Perry, and M. Feinstein. Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, 6(1):111–125, 2005.
- [29] S. H. Strogatz. *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press, New York, NY, 2003.
- [30] D.J.T. Sumpter. The principles of collective animal behaviour. *Philosophical Transactions of the Royal Society of London: Series B*, 361:5–22, 2006.
- [31] V. Trianni and M. Dorigo. Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231, 2006.
- [32] V. Trianni and S. Nolfi. Self-organising sync in a robotic swarm. a dynamical system view. *IEEE Transactions on Evolutionary Computation, Special Issue on Swarm Intelligennce*, 13(4):722–741, 2009.
- [33] V. Trianni, S. Nolfi, and M. Dorigo. Evolution, self-organisation and swarm robotics. In C. Blum and D. Merkle, editors, *Swarm Intelligence. Introduction and Applications*, Natural Computing Series, pages 163–192. Springer Verlag, Berlin, Germany, 2008.
- [34] V. Trianni and E. Tuci. Swarm cognition and artificial life. In *Advances in Artificial Life. Proceedings of the 10th European Conference on Artificial Life (ECAL 2009)*. Springer Verlag, Berlin, Germany, in press.
- [35] E. Tuci, C. Ampatzis, F. Vicentini, and M. Dorigo. Evolving homogeneous neurocontrollers for a group of heterogeneous robots: Coordinated motion, cooperation, and acoustic communication. *Artificial Life*, 14(2):157–178, 2008.
- [36] E. Tuci, R. Groß, V. Trianni, M. Bonani, F. Mondada, and M. Dorigo. Cooperation through self-assembling in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):115–150, 2006.
- [37] E. Tuci, R. Groß, V. Trianni, F. Mondada, M. Bonani, and M. Dorigo. Cooperation through self-assembling in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):115–150, 2006.
- [38] G. Whitesides and B. Grzybowski. Self-assembly at all scales. *Science*, 295:2418 – 2421, March 2002.
- [39] S. Wischmann, M. Huelse, J. F. Knabe, and F. Pasemann. Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14(2):117–127, 2006.