

# Evolution of Implicit and Explicit Communication in Mobile Robots

Joachim de Greeff, Stefano Nolfi

Institute of Cognitive Sciences and Technologies, CNR, Roma, Italy  
stefano.nolfi@istc.cnr.it, joachim.degreeff@plymouth.ac.uk

**Abstract.** This work investigates the conditions in which a population of embodied agents evolved for the ability to display coordinated/cooperative skills can develop an ability to communicate, whether and to what extent the evolved communication system can complexify during the course of the evolutionary process, and how the characteristics of such communication system varies evolutionarily. The analysis of the obtained results indicates that evolving robots develop a capacity to access/generate information which has a communicative value, an ability to produce different signals encoding useful regularities, and an ability to react appropriately to explicit and implicit signals. The analysis of the obtained results allows us to formulate detailed hypothesis on the evolution of communication for what concern aspects such us: (i) how communication can emerge from a population of initially non-communicating agents, (ii) how communication systems can complexify, (iii) how signals/meanings can originate and how they can be grounded in agents sensory-motor states.

## 1 Introduction

The study of how populations of artificial agents that are embodied and situated can autonomously develop communication skills and a communication system while they interact with a physical and social environment presents two important advantages with respect to experimental methods: (a) it allows us to study how communication signals are grounded in agents' non-symbolic sensory-motor experiences, and (b) it allows us to come up with precise and operational models of how communication skills can originate and of how established communication systems can evolve and adapt to variations of the physical and social environment.

Existing models of the evolution of communication in agents that are embodied and situated often focus on specific aspects, such as: (i) how a shared vocabulary can emerge in a population of agents that interact on the basis of a pre-determined interaction protocol (e.g. [Cangelosi and Parisi, 1998](#); [Steels, 1999](#)), (ii) how a simple unstructured communication system can evolve into a structured communication system (e.g. [Cangelosi and Parisi, 2001](#)), (iii) how a given communication system can be acquired and eventually socially transmitted (e.g. [Billard and Dautenhahn, 1999](#); [Steels and Kaplan, 2001](#); [Sugita and](#)

Tani, 2004). In the work described in this chapter, instead, we focus on the more fundamental question of how a population of initially non-communicating robots might develop communication skills that are functional with respect to the task that the robots have to perform without being rewarded for communicating.

More precisely, we devised an experimental scenario that is simple enough to be systematically analysed but that at the same time includes all the elements that are necessary to investigate important questions concerning the evolution of communication such as: which are the conditions that might lead to the evolution of communication skills in a population of initially non-communicating individuals? What is the relation between agents' communicative and non-communicative behaviours and between different communication modalities (e.g. implicit and explicit communication)? How does the 'meaning' of signals originate and evolve and how is it grounded in agents' sensory experience? The key aspects of the chosen scenarios are: (i) the fact that the task/environment admits a variety of qualitatively different solutions, (ii) the fact that the robots are provided with a sensory-motor system that allows them to interact/communicate through different modalities, (iii) the fact that the evolving robots are not rewarded for communicating and are left free to determine the way in which they react to sensory states and sequences of sensory states.

The data collected through these synthetic experiments represent important evidence for understanding the evolution of animal and human communication given the paucity of empirical data. This shortage of empirical data is due to the impossibility of analysing the evolutionary process in action and the difficulty to reconstruct it from indirect evidence because communication and language do not leave traces in fossil records.

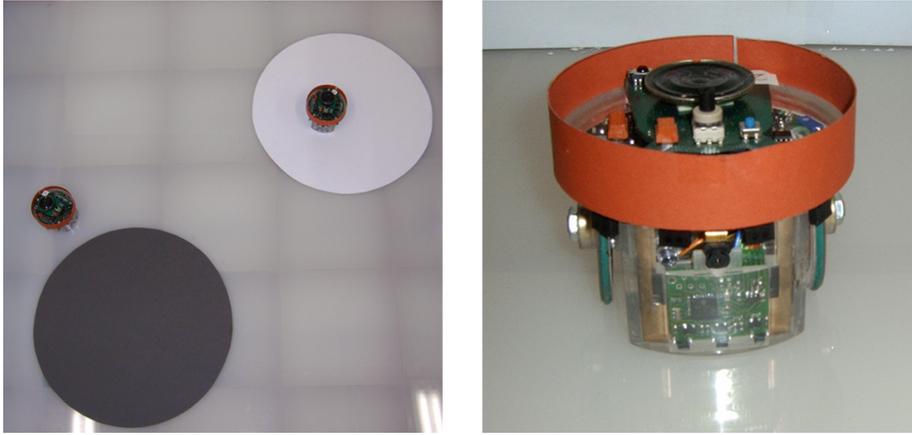
This rest of the chapter is organized as follows. In the next section we describe the experimental scenario. In section 3 and 4 we describe the obtained results and we discuss their implications. Finally, in section 5 we provide the details that allow interested readers to replicate the experiments.

## 2 Experimental setup

The experimental setup (Figure 1) involves two wheeled robots situated in an arena containing two target areas (one white and one black) that are evolved for being concurrently located in the two target areas and for switching areas as often as possible. The characteristics of the task/environment have been chosen so as to create a situation in which the robots should coordinate/cooperate to solve their adaptive problem. In the following sub-sections we describe the characteristics of the environment, of the robot's body and neural controller, and of the evolutionary algorithm.

### 2.1 The environment and the robots

The environment consists of an arena of either 110x110 or 150x150 cm surrounded by walls and containing two target areas with a diameter of 34 cm



**Fig. 1.** The environment and the robots. The two circular areas of the environment coloured in black and white represent the two target areas. Right: The e-puck robotic platform including the ground sensor board and a stripe of red paper around the top part of the body.

placed on two randomly selected but non-overlapping positions inside the arena. The floor of the arena and the walls are grey. The two circular portions of the arena corresponding to the two target areas are coloured black and white, respectively.

The robotic platforms consist of two e-Puck robots (Mondada and Bonani, 2007) provided with the ground sensor board extension. The robots, that have a diameter of 7.5 cm, are equipped with 2 motors that control the 2 corresponding wheels, 8 infrared proximity sensors located around the robot's body, 3 infrared sensors placed on the frontal side of the robot and oriented toward the ground, a VGA camera with a field of view of  $36^\circ$  pointing in the direction of forward motion, and a wireless Bluetooth interface that can be used to send and receive signals to and from other robots. The body of the robot has been covered with a circular stripe of red paper to allow robots to detect the presence of another robot in their field of view.

Signals consist of single floating point values ranging between  $[0.0, 1.0]$ , that are transmitted and received through the Bluetooth connection. Each time step both robots emit a signal and detect the signal produced by the other robot.

## 2.2 The neural controller

The neural controller of each robot is provided with 17 sensory neurons, 4 internal neurons with recurrent connections and 3 motor neurons. The internal neurons receive connections from the sensory neurons and from themselves. The motor neurons receive connections from both the sensory and the internal neurons (Figure 2).

The sensory layer includes 8 neurons that encode the state of the 8 corresponding infrared sensors, 3 neurons that encode whether the robot detects another robot in its field of view and the angular offset of the detected robot on the left or right side of the camera, 2 neurons that encode whether the ground sensor of the robot detects a white or black target area, 2 neurons that encode the previous state of the ground sensors and 2 signal sensors that encode the signal received from the other robot and the signal produced by the robot itself in the previous time step.

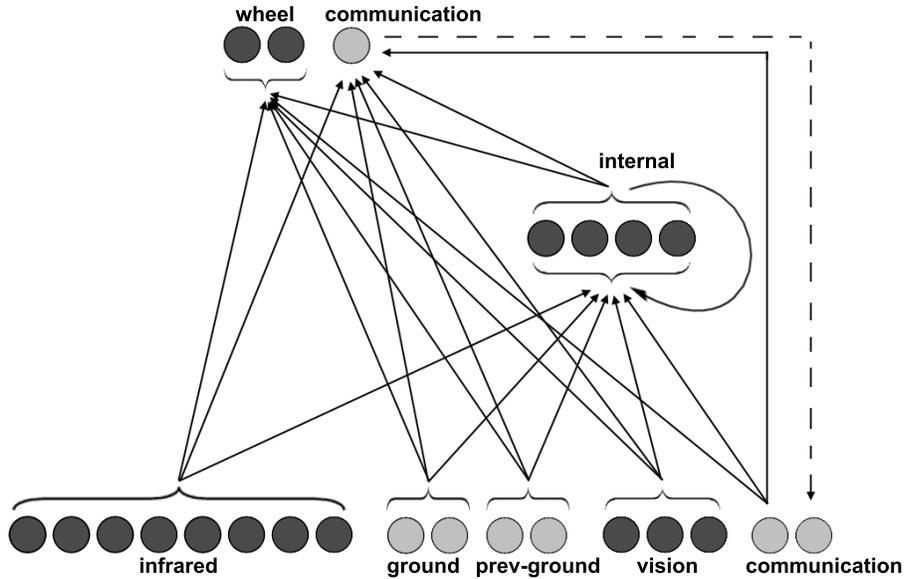
The motor layer includes 2 neurons that encode the desired speed of the 2 corresponding wheels and 1 neuron that encodes the value of the signal produced by the robot. The details on how this information is encoded in sensory and motor neurons are reported in section 5.2.

The state of sensory, internal and motor neurons are updated every 100ms (i.e. each time step lasts 100ms). The internal neurons consist of leaky integrator neurons that hold a certain amount of activation from the previous time step and in which the effect of the previous state on their current state is determined by their time constant parameter. The motor neurons consist of standard sigmoid units. More details are provided in section 5.2.

The type and number of sensors and actuators and the way in which the information extracted by the robots' sensors is encoded into the sensory neurons has been chosen to allow the robots to have a potentially rich interaction with their physical and social environment, while keeping the number of free parameters as low as possible.

More specifically, concerning communication, the possibility for the robots to perceive each other potentially allows the development of communication skills, i.e. the development of an ability to react to stimuli produced by other individuals in functional ways and/or the development of an ability to modify the perceptual environment of the other individuals in functional ways. Moreover, the possibility for robots to influence each other through different modalities (i.e. radio, vision, and infrared) potentially allows the robots to exploit both: (a) *implicit* communication forms, in which the robots develop an ability to react appropriately to the perceptual stimuli that are produced by the other individuals spontaneously, and (b) *explicit* communication forms, in which the robots develop an ability to functionally shape the way in which they affect the perceptual environment of the other individuals. For example, the possibility for the robots to visually detect the presence and the relative position of the other robot might allow the development of implicit communication forms which consist in reacting to such stimuli in a way which is functional with respect to the task/environment. On the other hand, the possibility for the robots to vary the radio signal produced in different robot/environmental circumstances allows them to develop explicit communication forms in which both the signal produced and detected and the reaction to such signals have been adapted.

For the sake of simplicity, in the rest of the chapter we will use the term *implicit signal* to indicate the signal that is generated by the actual physical position of the robots and that is detected by other robots through their visual



**Fig. 2.** The architecture of the robots' neural controller. The lower, middle and top layers indicate sensory, internal and motor neurons, respectively. Thin arrows indicate connections. The thick arrows indicates that the state of the communication motor neuron at time  $t$  is copied into the state of a sensory neuron at time  $t+1$ .

and infrared sensors and the term *explicit signals* to indicate the signals produced by a robot and received by other robots through the wireless connection (providing that robots do not always produce the same signal). This is justified by the fact that in this particular experimental setup, the latter stimuli, contrary to the former, can be shaped by the robot during the adaptive process. We should bear in mind, however, that there are no straightforward ways to formally distinguish between implicit and explicit signals. For example, as we will discuss below, also the perceptual stimuli generated by the physical position of the robots can be partially shaped in an adaptive way by the robots themselves through modification of the robots motor behaviours. Finally, the possibility to interact through different communication channels might lead to the development of communication forms that are based on a combination of implicit and explicit signals.

The 4 sensors that encode both the current and the previous state of the ground sensors allow the robots to easily recognize whether they are or were recently located in one of the target areas.

Finally, leaky internal neurons with recurrent connections allow the evolving robots to integrate sensory-motor information through time (e.g. to detect

the duration of a given sensory state) and/or to remember and eventually communicate previously experienced sensory states (Beer, 2003; Nolfi and Marocco, 2001). In other words, the characteristics of the neural controllers potentially allow the robots to extract and communicate information that is not currently available through their sensors.

### 2.3 The evolutionary algorithm

An evolutionary technique is used to set the free parameters of the robots' neural controller (Nolfi and Floreano, 2000). The initial population consists of 100 randomly generated genotypes that encode the connection weights, the biases and time constants of 100 corresponding neural controllers (each parameter is encoded by 8 bits and normalized in the range  $[-5.0, +5.0]$  in the case of connection weights and biases and in the range  $[0.0, 1.0]$  in the case of time constants). Each genotype is translated into 2 identical neural controllers that are embodied in 2 corresponding robots situated in the environment (i.e. teams are homogeneous). The 20 best genotypes of each generation are allowed to reproduce by generating 5 copies each, with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasts 1000 generations (i.e. the process of testing, selecting and reproducing robots is iterated 1000 times). The experiment is replicated 10 times for each of the 2 experimental conditions (smaller and larger environment).

Each team of two robots is allowed to "live" for 20 trials, lasting 200 seconds each (i.e. 2000 time steps of 100 ms each). However, if a collision occurs a trial is immediately terminated. At the beginning of each trial the position of the two target areas and the position and the orientation of the robots are assigned randomly.

Each team of evolving robots is scored with 1 point every time the two robots occupy the two different target areas for the first time during a trial or after a switch (i.e. after the robot that previously occupied the white target area moved to the black target area and vice versa). The total performance of a team (fitness) consists of the average number of points scored during the 20 trials.

The robots' neural controllers have been evolved in simulation (see details in section 5.3). The best evolved neural controllers have been tested on hardware (i.e. have been downloaded on the two robots situated in the physical environment).

\* \* \*

Before concluding the description of the experimental setup it is important to clarify which are the characteristics that are pre-determined by the experimenter and fixed and which are those that are unspecified and that are left free to vary during the evolutionary process.

One aspect that is predetermined is the fact that the experimental scenario involves a full cooperative situation. This is due to the fact that the two robots forming the team have the same genetic characteristics and to the fact that selection operates on the basis of the performance of the team and not on the

performance of a single individuals (see Chapter ??). A second aspect that is pre-determined is constituted by the fitness function that is used to select the best individuals. The adaptive task faced by the robots, however, is only partially pre-determined since it depends on the fitness function and on the characteristics of the robots' bodies and of the environment, that are fixed, but also on the characteristics of the social environment (i.e. the other robots' behaviour), that is not predetermined and varies during the evolutionary process. In fact, the development of new behavioural and communication skills modifies the social environment of the robots themselves. These modifications, in turn, modify the adaptive task faced by the robots. Indeed, as we will discuss in section 4, modifications of robots' behaviour and communication skills might create the adaptive condition for the emergence of new skills.

The motor and communicative behaviours exhibited by the robots are not pre-determined since the way in which a robot reacts to any given sensory state or sequence of sensory states depends on the free parameters that are encoded in the genome of the population and are subjected to variations. Indeed, as we will see, evolving robots are left free to determine the characteristics of their behaviour within a large space of different behavioural solutions. More specifically, concerning robots' motor behaviour, the robots are free to determine the number and the type of elementary behaviours that they display and the way in which these behaviours are combined and arbitrated. Concerning their communicative behaviours, evolving robots are left free to determine how they will collect the information that has a communicative value from the physical and social environment, how many different signals they will produce, in which agent/environmental context each signal will be produced, and which will be the motor and communicative effects of the explicit and implicit signals that are detected. Finally, evolving robots are free to co-adapt their motor and communicative behaviours.

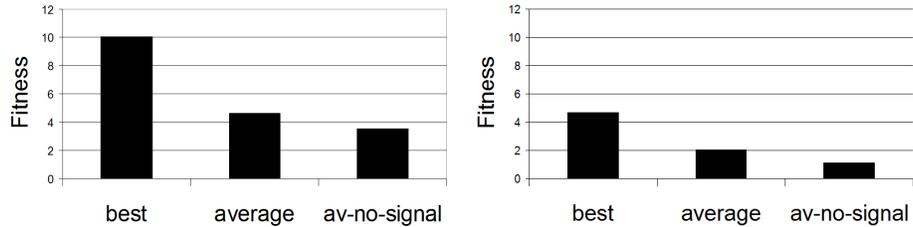
The theoretical approach and the methodology followed in this chapter is in line with the first Chapter of this section and with the work of [Baldassarre et al. \(2003\)](#); [Di Paolo \(1997, 2000\)](#); [Marocco and Nolfi \(2007\)](#); [Quinn \(2001\)](#); [Quinn et al. \(2003\)](#); [Trianni and Dorigo \(2006\)](#). However, the experimental scenario proposed here is more advanced than in the experimental works mentioned above with respect to the following aspects (or with respect to the possibility to study the following aspects in combination): (i) the complexity of the chosen task that allows us to study how several behavioural and communication skills are developed and co-adapted during the evolutionary process, (ii) the richness of the agents' sensory-motor system that supports, for example, the exploitation of both explicit and implicit communication, (iii) the validation of the results obtained in simulation in hardware.

### 3 Results

The analysis of the results obtained in different replications of the experiment and in different experimental conditions indicates that the robots solve the prob-

lem through qualitatively different strategies by exploiting the possibility to communicate through explicit and implicit signals. More precisely, the analysis of the performance of the best robots at the end of the evolutionary process, for the two different experimental conditions (i.e. 110x110 and 150x150 cm arenas) and for different replications of each of the two experiments, indicates that evolved robots display rather good performance in the case of the best replications and an ability to solve the task in all replications (see Figure 3, first two histograms of each graph). The comparison of the performance obtained in the normal condition and in a control condition in which the robots were not allowed to detect explicit signals indicates that the robots exploit explicit communication in order to coordinate and cooperate (see Figure 3, second and third histogram of each graph). Evolved robots also rely on implicit communication in most of the cases, as we will illustrate below.

By testing the robots evolved in simulation in a real environment (i.e. by embodying the neural controller on physical robots and by situating them in the physical environment shown in Figure 1) we observed that the behaviours exhibited in hardware are qualitatively very similar to those displayed in simulation. This can be clearly seen on the videos available at the following webpage <http://laral.istc.cnr.it/esm/evo-communication>.



**Fig. 3.** Fitness of the experiments performed in the 110x110 and 150x150 cm arenas, left and right picture respectively. Best indicates the performance of the best robots of the best replication. Average indicates the average performance of the best robots of all replications. Av-no-signal indicates the average performance of the best robots of all replications in a control experiment in which the robots are not allowed to detect the explicit signals produced by the other robots.

The visual inspection of the fittest evolved solutions indicates that they can be grouped in two qualitatively different strategies. In both strategies, the robots initially display an exploration behaviour that allows them to enter the two target areas (one robot per area) and then display a series of target switching behaviours in which each robot exploits the information provided by the other individual to navigate directly toward the other target area. The first strategy (that will be called *symmetrical strategy* from now on and that corresponds to

the strategy exhibited by the best robots of the best replication performed in the 110x110 arena) is characterized by a synchronized target switching behaviour in which the two robots, located in the two different target areas, simultaneously leave their current target area and move directly toward the other target area. The second strategy (that will be called *asymmetrical strategy* from now on and that corresponds to the strategy exhibited by the best robots of the best replication performed in the 150x150 arena) is characterized by a switching behaviour organized in two phases in which first a robot exits from its target area and travels toward the other target area containing the second robot and then the latter robot exits from its target area and travels directly toward the target area previously occupied by the former.

In sections 3.1 and 3.2 we will describe in details the behavioural and communication skills exhibited by evolved robots and the process through which such skills originate evolutionarily for the best replications of the experiment displaying the symmetrical and asymmetrical strategy, respectively.

The results of this analysis indicate that evolving robots display rather rich behavioural and communication skills, including: the ability to access/generate information that has a communicative value, the ability to produce different signals encoding useful regularities, and the ability to react appropriately to explicit and implicit signals by also regulating the reaction on the basis of the context in which signals are detected. This analysis provides a detailed account of how the meaning of signals is grounded in agents' sensory-motor experiences and on which are the relations between agents' behavioural and communicative skills. Moreover, the analysis of the evolutionary origins of robot' skills demonstrates how the co-development of behavioural and communicative skills might overcome the problems due to the need to develop two complementary abilities (i.e. produce useful signals and react to such signal appropriately). This might lead to an open-ended evolutionary process (within the limits imposed by the complexity of the task/scenario) in which agents' skills progressively complexify since the skills that have been developed at a certain stage tend to be preserved in successive generations and since the development of a new skills creates the adaptive conditions for the synthesis of further adaptive novelties.

Finally, in section 4, we will discuss the general implications of the obtained results related to our ability to model the evolution of communication in biological systems.

### 3.1 Symmetrical strategy

In this section we describe in detail the behaviour of the best evolved robots situated in the 110x110 cm arena (that display a symmetrical strategy). To do this, we will divide the overall behaviour exhibited by the robots into a set of elementary signalling and motor behaviours. These behaviours consist of sequence of movements that produce a given outcome (e.g. they allow the robots to avoid an obstacle or to move toward the other robot). For the purpose of explanation, behaviours having similar functions will be grouped together. For example, sequences of left-turning movements producing arc trajectories

over circles with a radius ranging from 7 to 15 cm might be grouped into the same type of elementary behaviour providing that different instances of this type of behaviour have similar, although not necessarily identical, functionalities. Similarly, sequences of signals varying within a given range, might be considered as instances of the same type of signal providing that different instances of the same signal have similar, although not necessarily identical, functionalities. In other words, the description in terms of classes of elementary behaviours that will be illustrated below does not capture the full story but only summarizes the most important characteristic of the evolved strategies.

### 3.1.1 Motor and communication behaviours repertoire

The analysis of the evolved robots indicates that they exhibit the following signalling and motor behaviours:

- A **signal-A** behaviour that consists in the emission of a signal in the range [0.7 to 0.9] that is produced by robots located outside target areas detecting a signal greater than 0.3. The fact that the production of this signal is influenced by the perceived signal implies that this signal is only produced when both robots are located outside target areas.
- A **signal-B** behaviour that consists in the emission of a signal in the range [0.4 to 0.6] that is produced by robots located in the white target area.
- A **signal-C** behaviour that consists in the emission of a signal in the range [0.0 to 0.3] that is produced by robots located in the black target area or by robot located near obstacles. The value of the signal produced by robots in the black target area increases from 0.0 to approximately 0.3 as the time spent by the robot in the area increases. Hence, through this mechanism the time spent on this target area is implicitly encoded.
- An **obstacle-avoidance** behaviour that consists of a sequence of right-turning movements performed near an obstacle constituted by a wall or by another robot through that the robot turns on the spot until the frontal side of the robot is free from obstacles. This behaviour is produced when the frontal-right infrared sensors of the robot are activated.
- A **move-straight** behaviour that consists of a sequence of move forward movements. This behaviour is produced by robots located outside target areas or by robots located in the white target area detecting signal A.
- A class of **turn-right** behaviours, that consists of a sequence of right-turning movements through which the robot moves forward and turns right by producing an arc trajectory over a circle with a diameter ranging from a few mm up to 100 cm. This type of behaviours is produced by robots detecting signal B or C or by robots located in the black target area detecting signal A. How much the robot turns during the execution of this behaviour depends on three factors (and on their combination): the value of the signal received, the location of the robot, and the time spent by the robot in the current location (for robots located in target areas). More precisely, the angle of turning is inversely proportional to the value of the received signal, and varies on the basis of the time spent by the robot in the area (it increases or decreases

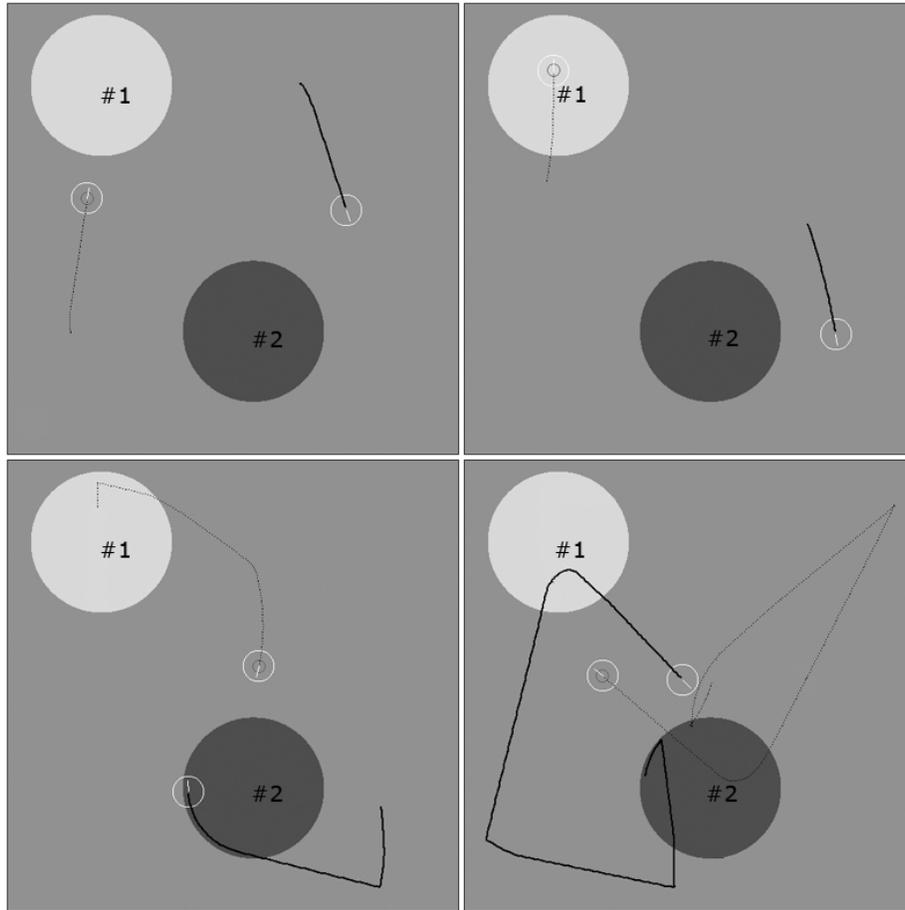
depending when the robot is in the white or black target area, respectively). For example, the arc trajectory produced by robots located outside target areas varies in the range [5, 100] cm for received signals varying in the range [0.0, 0.3]. Variation in the time spent by the robot on the target area results in the diameter of the arc trajectory produced by the robots varying over time (in the range [0, ~100] cm) for robots located on target areas and detecting signal 0.5.

- A **move-toward-robot** behaviour that consists of a sequence of move-forward and right-turning movements that allows a robot which visually perceives the other robot to move toward it and to keep it in its visual field for a certain period of time. This behaviour is produced by robots perceiving signal B or C and having the other robot in their field of view. Since the ability of the robots to maintain the other robot in their field of view is limited and since this behaviour is produced only when this condition holds, the execution of this behaviour might terminate after some time. The probability that this behaviour is terminated depends on the relative movement of the visually perceived robot. In fact, robots never lose visual contact with a robot that stays still or moves within a target area but quickly lose visual contact with a robot displaying exploration behaviour.

To illustrate how the combination of these behaviours allows the robots to solve their collective task, let us consider a typical trial in which the two robots are initially located outside target areas (Figure 4, top-left). In this situation the robots display a move-straight and signal-A behaviour when they are far from obstacles and an obstacle-avoidance and signal-C behaviour near obstacles. The combination of the move-straight and obstacle-avoidance behaviour allows the robots to explore the environment. The role of signal-A produced by each robot is to trigger the move-straight behaviour in the other robot. The obstacle-avoidance behaviour, instead, is triggered by the activation of the frontal-right infrared sensors independently from the detected signal (i.e. the obstacle-avoidance behaviour subsumes the move-straight behaviour). The signal C produced during the execution of obstacle-avoidance behaviour does not have an adaptive function. Indeed, a comparison of the fitness measured in a normal condition and in a control condition in which the robots are forced to produce signal-A rather than signal-C near obstacles did not produce significant differences (results not shown).

When a robot enters in the white target area it keeps producing the move-straight behaviour and modifies its signalling behaviour from A to B while the other robot keeps producing the move-straight and signal-A behaviour (Figure 4, top-right). The production of the signal-B behaviour during this phase therefore does not have any adaptive function.

When the first robot enters in the black target area and the second robot is still outside target areas (Figure 4, bottom-left), the first robot produces a turn-right and signal-C behaviour followed by normal exploration behaviour when it exits from the target area. The second robot, instead, produces a turn-right behaviour followed by a move-straight behaviour or followed by a move-toward-



**Fig. 4.** The behaviour produced by the best robots of the best replication of the experiment performed in the 110x110 cm arena. The 4 pictures show 4 snapshots of the trajectory produced by the robots during seconds 0-4, 4-9, 9-21 and 21-53 (from left to right and from top to bottom, respectively). The black and grey lines represent the trajectory of the two robots during the corresponding phases.

robot and move-straight behaviour (depending on whether the former robot switches back to signal  $\rightarrow A$  or not) The move-toward-robot behaviour, that is triggered when the latter robot visually detect the former robot, allows the latter robot to keep turning so to align toward the direction of the former robot (i.e. toward the direction of the black target area). When this happens, i.e. when the latter robot visually perceives the former robot, the combination of the behaviours described above produces a coordinated behaviour in which the two robots turn right, move toward each other by keeping the other robot on the

right side of their field of view, and then separate again by moving in opposite directions. Let us refer to this sequence of coordinated behaviours as position-exchange behaviour, since it leads to a state in which each robot roughly assumes the relative position and orientation previously assumed by the other robot.

The fact that the execution of a position-exchange behaviour results in the re-establishment of the conditions that triggered it in the first place implies that this behaviour tends to repeat itself periodically. The fact that the position and the orientation of the two robots before and after the execution of this behaviour is similar but not identical, ensures that the two robots keep moving in and out the black target area (thus preserving information on that location) while exploring different parts of the environment at the same time. This also allows the robots to identify the location of the white target area after some time.

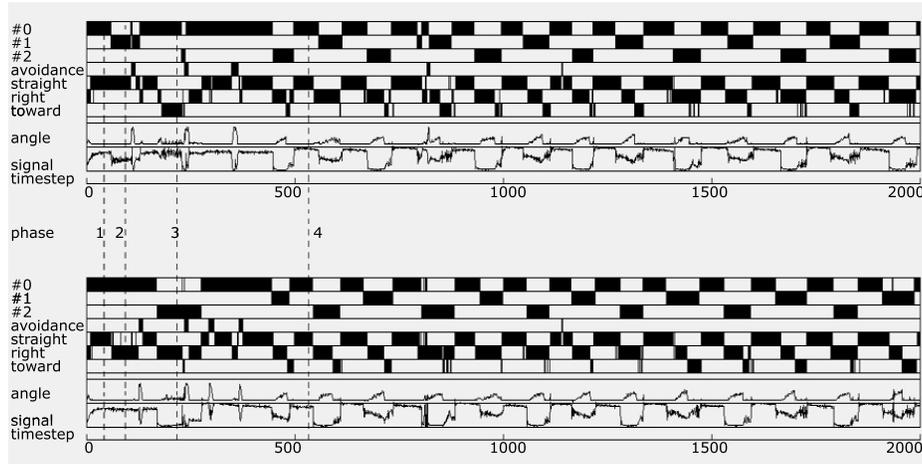
The execution of a position-exchange behaviour from a situation in which the two robots are concurrently located in the two target areas, instead, leads to the re-establishment of the same initial positions and orientations of the two robots. Hence, this leads to the production of a periodic target-switching behaviour in which each robot moves toward the target area previously occupied by the other robot.

What differentiates these two situations, namely the situation in which the two robots are initially located in the two target areas and the situation in which one robot is initially located in the black target area while the other robot is located outside target areas, consists in the fact that robots located in the white target area produce signal-B while robots located outside target area produce signal-A. This emitting of signal-B leads to a situation in which both robots turn toward each other and then visually perceive each other, that is facilitated by the fact that the detection of this signal makes the robot located in the black target area turn more quickly. The concurrent exhibition by the two robots of a move-toward-robot behaviour, in turn, ensures that they exchange their position by reassuming approximately the same position and orientation previously assumed by the other robot.

### 3.1.2 Communication system

The arbitration between different behaviours and the coordination between the robots is regulated by the social interaction between the two robots, mediated by the evolved communication system. This system includes the behaviours that enable the robots to access information that has a communicative value, the implicit and explicit signals produced by the robots, and the effects produced by the perception of these signals.

The behavioural skills that allow the robots to access useful information, in the case of this replication of the experiment, consist of an exploratory behaviour that allows the robots to find the two target areas. Since the robots keep moving in all environmental situations, the information concerning the location of a particular target area is accessed and communicated by the robot travelling over the area only for a limited period of time. However, when the two robots enter in



**Fig. 5.** Analysis of the behaviour displayed in Figure 4 throughout time. The 4 phases indicated by the dashed lines correspond to the 4 pictures in Figure 4. The top and the bottom graph show the data relative to the first and the second robot indicated with grey and black lines in Figure 4, respectively. For each graph, the top part indicates the location of each robot (#0 = outside target areas, #1 = in the white target area, #2 = in the black target area). The middle part indicates the elementary behaviour exhibited by the robot (avoidance = obstacle-avoidance, straight = move-straight, right = right turn, toward = move-toward-robot). The bottom part shows the amplitude of the angle of turn that characterizes different instances of behaviours belonging to the collection of turn-right behaviour and the value of the emitted signal. For more information on how data has been collected, see section 5.

a coordinated phase, the information on the location of previously visited areas is preserved in the relative position and orientation of the two robots. This means that evolved robots are also able to preserve and to implicitly communicate information about previously experienced sensory states.

The signals produced by the robots include three explicit signals (A, B, and C) and an implicit signal constituted by the presence or absence of one robot in the visual field of the other robot. More precisely, the signals A, B and C encode whether the robot emitting the signal is located outside target areas or in the white or black target area, respectively. The implicit signal provides information on the relative location of the robot. The combination of the implicit and the explicit signals B and C encode roughly the relative directions of the white and black target area. The signal A instead encodes whether both robots are located outside target areas. This information does not only concern the state of only the robot emitting the signal but the state of both robots. Indeed, the information

conveyed by this signal is generated through the communicative interaction of the two robots.

The signals effect is a modification of the robots' motor behaviour that is context dependent (i.e. the type of effects produced by a signal and/or whether or not a signal produces an effect depends on the state of the robot detecting the signal). More precisely:

The perception of signal-A and the absence of an implicit signal triggers a move-straight behaviour in robots located outside target areas and far from obstacles.

The perception of signal-B and the absence of an implicit signal: (1) a move-straight behaviour in robots located outside target areas and far from obstacles, and (2) a turn-right behaviour in robots located in the white target area.

The perception of signal-C and the absence of an implicit signal triggers a turn-right behaviour. The amplitude of turn of the turn-right behaviour varies depending on the location of the robot detecting the signal and on the time spent by the robot in that location (inside target areas).

The perception of the implicit-signal (i.e. the perception of the other robot in the visual field) triggers a move-toward-robot behaviour, independently from the type of explicit signals perceived. However, this move toward-robot behaviour is abandoned when the two robots get too close to each other (i.e. when the activation of the infrared sensors causes an obstacle-avoidance behaviour).

In most of the cases, the explicit signal produced by a robot encodes information that is currently available to the robot through its sensors. However, in some cases, signals also encode information concerning the sensory states previously experienced by the robot emitting the signal. This is the case for signal-C that encodes not only the fact that the robot is located in the black target area but also the time spent by the robot in that target area (the value of the signal increases from 0.0 to 0.3 as time passes). The information about the time spent in the black area conveyed by signal-C plays a functional role. Indeed, by testing the evolved individuals in a control condition in which robots located in the black target area were forced to produce a signal 0.0 independently from the time spent in the area, we observed a drop in performance from 9.673 to 4.189. The analysis of the behaviour displayed by the robots in the control condition indicates that the loss of performance is mainly due to the fact that, during the execution of the target switching behaviours, the constraint imposed on the signal produced by the robot located in the black target area makes the other robot turn too much toward it. This excessive turning causes the robot exiting from the white target area to move directly toward the robot exiting from the black target area so that the two robots are then forced to avoid each other and thus lose the information about the relative location of the target areas. For more on this aspect see the data provided in the following section.

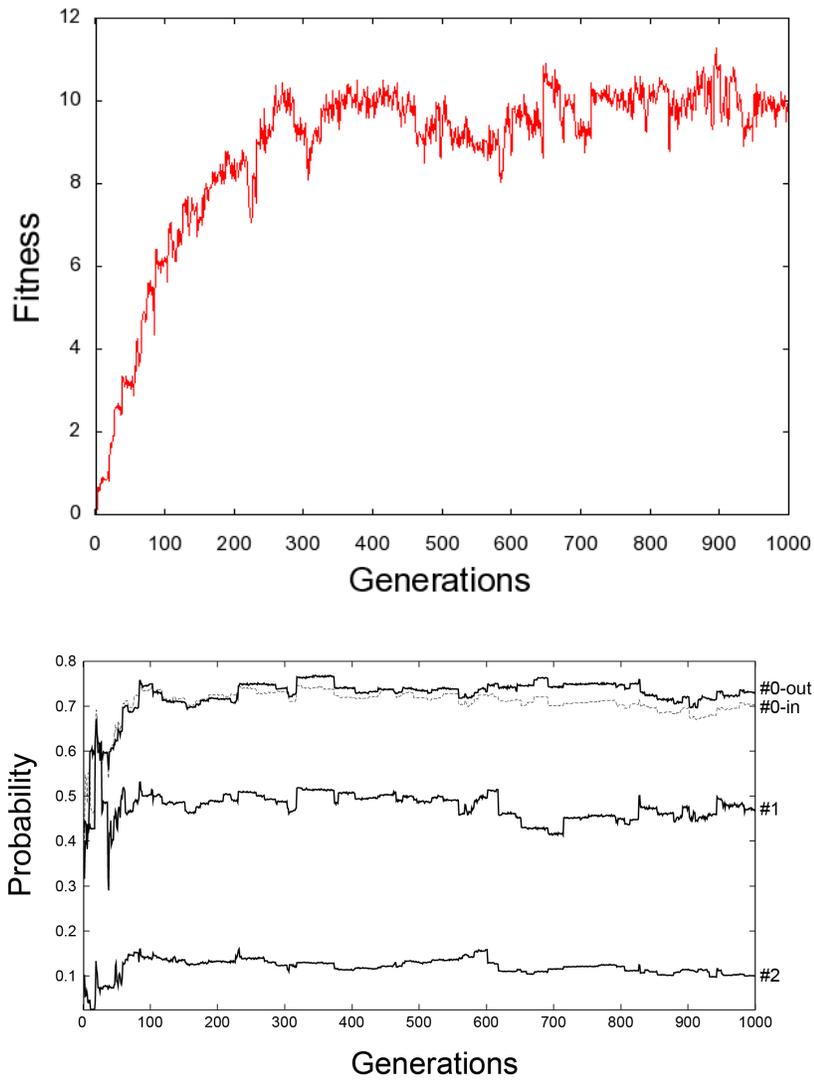
### **3.1.3 Evolutionary origin of robots' communicative and non-communicative skills**

In the case of this replication, the performance gradually increases during the

first 250 generations until a relatively stable state is reached both in terms of performance and in terms of the robots' behaviours (Figure 6, top). From the very first generation up to generation 20 the robots display an obstacle-avoidance behaviour close to obstacles, a move-straight/turn-right behaviour outside the white target area and a left-turning behaviour when located in the white target area. In this initial stage, the robots also produce a signal-C or a signal-B depending on whether they are located inside or outside the black target area (Figure 6, bottom). These behaviours allow the robots to explore the environment and occasionally reach the two target areas at the same time. The robots react to the implicit 'signal' constituted by the visual perception of the other robot by triggering a move-toward-robot behaviour and to 'signal-C' by triggering a turn-right behaviour characterized by an angle of turn that is inversely proportional to the value of the detected signal. At this stage, however, these interactions mediated by implicit and explicit signals do not lead to forms of coordinated behaviour with an adaptive function. We indicated these signals in quotes above since they do not yet play a functional role.

From generation 20 to generation 30, the turn-left behaviour originally displayed in the white target area is replaced with a move-straight behaviour and the move-straight behaviour originally displayed in the black target area is replaced with a turn-right behaviour. At the end of this phase the robots further differentiate their explicit signals by producing signal-B when they are located in the white target area and signal-A when they are located outside target areas (Figure 6, bottom). The number of signals produced by the robots and the value characterizing the different signals will remain rather stable from this phase on. As in the previous phase, the robots react to signal-C by triggering a turn-right behaviour and to the implicit signal constituted by the visual perception of the other robot by triggering a move-toward-robot behaviour. This allows the robots to occasionally switch their areas at least once without resorting in exploration behaviour. The turn-right behaviour produced by robots concurrently located in the two target areas (that is triggered by the state of the ground sensor for the robot located in the black target area and by the signal-C detected by the robot located in the white target area), increases the chance that the two robots will visually perceive each other and will trigger a move-toward-robot behaviour. However, target switching is achieved only very occasionally at this stage.

From generation 31 to generation 40, the robots improve their ability to exploit the available signals by modifying the way in which they react to detected signals. Indeed, from now on, the robots react to the detection of signal-A by producing a move-straight behaviour and to the detection of signal-B and signal-C by producing different types of turn-right behaviours. Moreover, robots located in the black target area differentiate their behaviour depending on whether the other robot is concurrently located in the white target area or not. Indeed, during this phase, robots located in the black target area develop an ability to turn-right sharply or to move straight depending on whether they detect signal B or A (i.e. depending on whether the other robot is currently located in the white target area or not, respectively). This in turn allows the robots to develop an ability



**Fig. 6.** Top: fitness through out generations. Data calculated by testing the ancestors of the best individuals of the last generation for 1000 trials. Bottom: Average value of the explicit signals produced by robots located in the white target area (#1), in the black target area (#2), and outside target areas (#0-in, #0-out). For the last case the two curves indicate the average signal produced by a robot when the other robot is located inside a target area (#0-in) or outside target areas (#0-out).

to occasionally switch areas by navigating directly toward the area previously occupied by the other robot.

From generation 41 on, the general strategy and the characteristics of the implicit and explicit signals remain substantially the same. However, the exact way in which the implicit and explicit signals regulate the behaviour of the robots in different conditions is progressively modified throughout generations so to maximize performance. These modifications affect in particular: (a) the way in which the turning angle that characterizes the turn-right behaviour varies as a function of the time spent by the robot on the target areas (as shown in Figure 7, top, for robots locate in the black target area), (b) the fine-grained characteristics of the move-toward-robot behaviour, and (c) the modulation of signal-C as a function of the time spent by the robot in the black target area (as shown in Figure 7, bottom). The analysis of the correlation between these variations and the value of the fitness throughout generations indicates that some of the variations are adaptive. This is the case, for example, for the significant variation in the turning angle that occurs at generation 66 and is later preserved in successive generations (Figure 7, top).

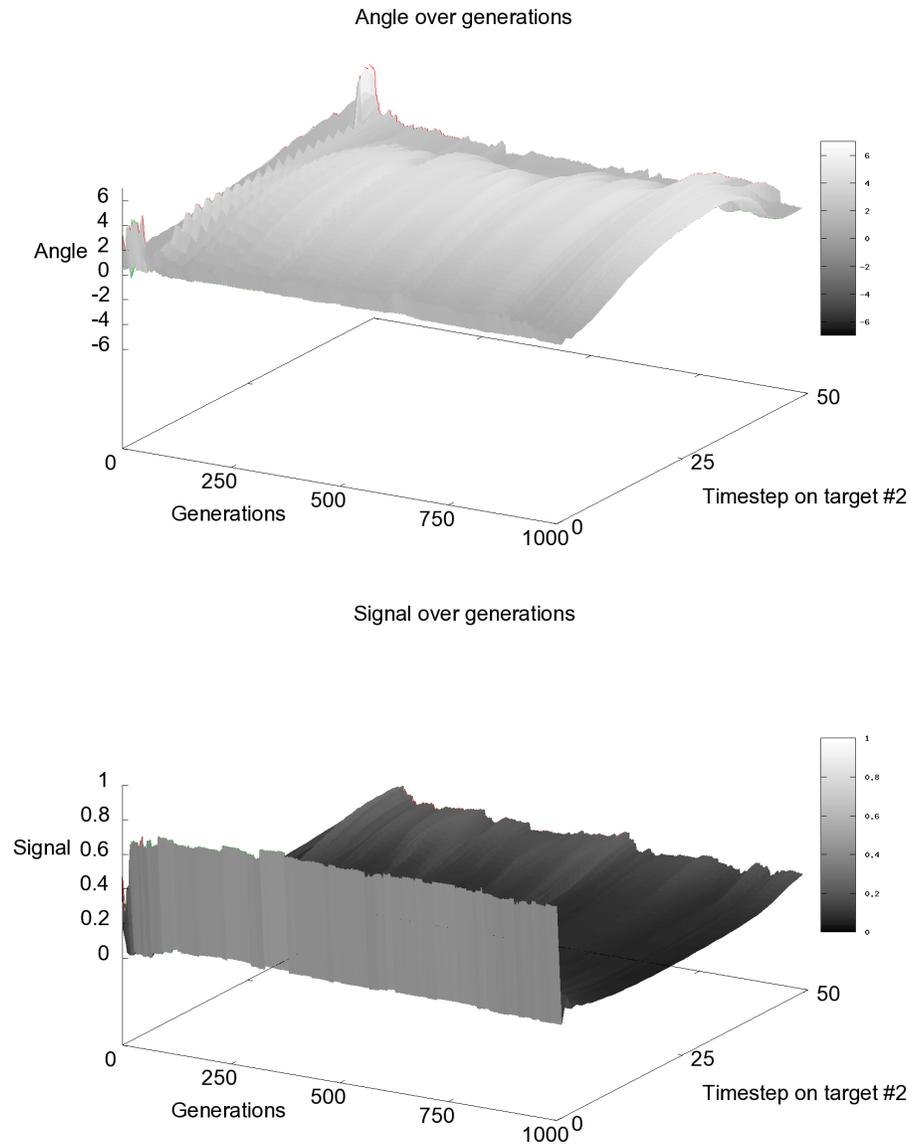
### 3.2 Asymmetrical strategy

In this section we describe the best evolved individual of the best replication of the experiment performed on a 150x150 cm environment (that displays an asymmetrical strategy).

#### 3.2.1 Motor and communication behaviours repertoire

The analysis of the evolved robots indicates that they exhibit the following behaviours:

- A **signal-A** behaviour that consists in the emission of a signal in the range [0.9 to 1.0]. This signal is always produced by robots located outside the black target area that are not detecting obstacles.
- A **signal-B** behaviour that consists in the emission of a signal in the range [0.0 to 0.6]. This signal is always produced by robots located in the black target area.
- An **obstacle-avoidance** behaviour that consists of a sequence of left-turning movements. This behaviour is always performed near an obstacle (a wall or another robot) when left, frontal, or right infrared sensors of the robot are activated, regardless of the signals perceived. The robot turns on the spot until the frontal side of the robot is free from obstacles.
- A **move-straight** behaviour that consists of a sequence of move forward movements. This behaviour is always produced by robots located outside target areas when no other robot is perceived visually and no obstacles are detected.
- A **follow-border** behaviour that consists of a combination of left-turning and move-forward movements that allow a robot to move counter-clockwise following the border of a target area. This behaviour is always produced by robots located in the white or black target area that do not visually perceive the other robot, regardless of the perceived signal.

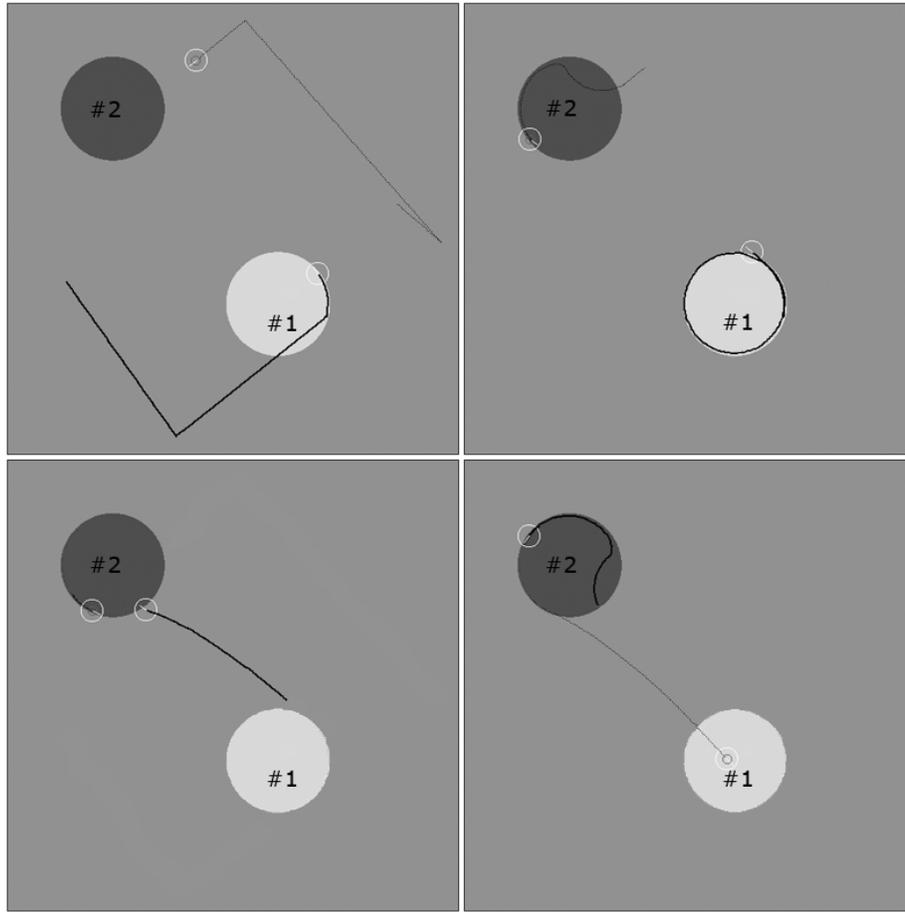


**Fig. 7.** Average turning angle (top) and signal (bottom) displayed by robots located in the black target area (target #2). Each picture shows how the corresponding value varies as the time spent by the robot in the area increases for individuals of successive generations. Data is computed by testing the ancestors of the best evolved individual of the last generation for 1000 trials.

- An **avoid-robot** behaviour that consists of a sequence of left-turning movements. This behaviour is produced by robots located outside areas that visually perceive the other robot in all cases beside the cases in which the additional conditions that trigger the execution of the move-toward-robot behaviour hold. The robot turns on the spot until the other robot exits from its field of view.
- A **move-toward-robot** behaviour that consists of a sequence of move forward and left-turning movements that allow a robot to move straight by slightly turning toward the direction of a visually perceived robot. This behaviour is always produced by robots that: are located outside target areas, previously visited the white target area, detect signal-B, and detect the other robot in their field of view.
- A **look-robot-and-follow-border** behaviour that consists of a combination of left-turning, right-turning and move-forward movements that allow the robot to maintain the other robot on the left-side of its field of view and to maintain the same relative position in a target area with respect to the other robot. This latter aspect is realized by remaining on the spot when the other visually-perceived robot is on the frontal or right side of the visual field and by moving counter-clockwise along the border of the area when the other robot is on the left side of the visual field. This behaviour is always produced by robots that are located in the black target area, perceive signal-A and visually perceive the other robot.
- An **exit-area-1** behaviour that consists of one or a few move-forward movements that allow a robot located in the white target area to exit from this area. This behaviour is always produced by robots located in the white target area that perceive signal-B and visually detect the other robot in the left part of their visual field.
- An **exit-area-2** behaviour that consists of one or a few move-forward movements that allow a robot located in the black target area to exit from this area. This behaviour is always produced by robots located in the black target area that perceive signal-B.

To illustrate how the combination of behaviours allows the robots to solve their task, let us consider a typical trial.

The two robots are initially located outside target areas (Figure 8, top-left). In this phase the robots display a move-straight behaviour when they are far from obstacles and do not visually perceive other robots, an obstacle-avoidance behaviour when they detect an obstacle with the infrared sensors, and an avoid-robot behaviour when they visually perceive the other robot. The combination of the move-straight and obstacle-avoidance behaviours allows the robots to explore the environment. The avoid-robot behaviour does not play a functional role when both robots are located outside target areas. Indeed, the performance measured in a normal condition does not significantly differ from the performance observed in a control condition in which the robots located outside target area were not allowed to visually detect the other robot. The signalling behaviours



**Fig. 8.** The behaviour displayed by the evolved robots of the best replication of the experiment performed in the 150x150 cm arena. The 4 pictures show 4 snapshots of the trajectory produced by the robots from 0 to 21, 21 to 38, 38 to 47 and 47 to 58 seconds (from left to right and from top to bottom, respectively). The black and the grey lines represent the trajectory of the two robots during the corresponding phases.

produced when both robots are located outside target areas do not alter the motor behaviour of the robots themselves and thus do not have any functionality.

When a robot enters in the white target area while the other robot is located outside target areas, it starts to produce a follow-border behaviour (Figure 8, top-left). This follow-border behaviour allows the robot to remain in the white target area until the other robot enters into the black target area. The signalling behaviour produced by the robot located in the white target area does not

have any functionality since it does not alter the motor behaviour of the other robot. The implicit signal produced by the robot located in the white target area triggers the avoid-robot behaviour in the other robot that plays an adaptive role in this circumstance. Indeed, the variance of the overall performance observed in a normal condition and in a control condition in which the robots located outside target areas were not allowed to visually detect robots located in the white target area is significant (average score of 4.723 and 3.941, respectively).

When a robot enters the black target area while the other robot is located outside target areas, it starts to produce a signal-B behaviour and a follow-border behaviour or a look-robot-and-follow-border behaviour, depending on whether or not it perceives the other robot visually. The function of the follow-border behaviour is to remain in the black target area and to look around in order to identify the relative position of the other robot. The look-robot-and-follow-border behaviour play several roles (that will be discussed below in more details): (i) it allows the robot to remain in the black target area, (ii) it allows the robot to assume a specific relative position in the target area with respect to the other robot that in turn provides for that robot an indication of the exact position of the black target area, (iii) it allows the robot to orient itself toward the centre of the white target area (as soon as the other robot enters that target area). Also in this case, the explicit signals produced by the two robots do not affect their motor behaviour and therefore do not have any functionality.

Finally, when the two robots are concurrently located in the two target areas they trigger a sequence of coordinated behaviours that is repeated over and over. This allows the two robots to quickly exchange their relative locations several times thus maximizing their fitness.

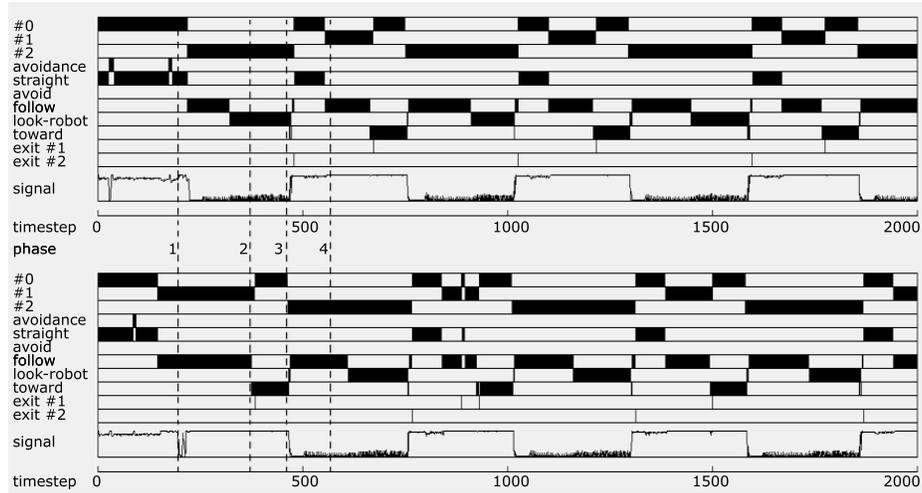
During the first phase of this sequence, the robot located in the black target area displays a follow-border behaviour or a look-robot-and-follow-border behaviour depending on whether it visually perceives the other robot or not. The robot located in the white target area displays a follow-border behaviour (Figure 8, top-right).

During the second phase, when both robots visually perceive each other on the left side of their field of view, the robot located in the white target area triggers and exit-area-1 behaviour that allows it to exit from the area and to initiate a move-straight behaviour toward the black target area (Figure 8, bottom-left).

During the third phase the robot that left the white target area displays a move-toward-robot behaviour through which it moves toward the direction of the other robot while the robot located in the black target area continues to look toward the approaching robot. The trajectory of the move-toward-robot behaviour allows the approaching robot to move approximately toward the centre of the black target area, thus maximizing the chance to enter this target area and avoiding the risk of obstructing the occupying robot. The look-robot-and-follow-border behaviour through which the occupying robot maintains the approaching robot on the left part of its visual field, allows the former robot to leave the black target area by being oriented toward the direction of the white target area.

During the fourth phase, as soon as the approaching robot enters the black target area and switches its signalling behaviour from A to B, the occupying robot leaves this target area by triggering an exit-area-2 behaviour and then a move-straight behaviour. The newly arrived robot triggers a follow-border and then a look-robot-and-follow-border behaviour (Figure 8, bottom-right). The orientation of the robot exiting from the black target area (that depends on the relative position assumed by the robot in the target area, the ability to keep the approaching robot on the left side of its visual field, and the ability of the approaching robot to move toward the centre of the area) ensures that the move-straight behaviour will bring this robot directly toward the centre of the white target area.

Finally during the fifth and last phase, the robot that left the black target area enters into the white target area. At this point the two robots are located again in the two target areas and the sequence of coordinated behaviours articulated in the five phases described above is repeated.



**Fig. 9.** Analysis of the behaviour displayed in Figure 8 throughout time, the 4 phases indicated by dashed lines correspond to the 4 pictures in Figure 8. The top and the bottom graph show the data relative to the first and the second robot indicated with grey and black lines in Figure 8, respectively. For each graph, the top part indicates the location of each robot (#0 = outside target areas, #1 = in the white target area, #2 = in the black target area). The middle part indicates the behaviour exhibited by the robot (avoidance = obstacle-avoidance, straight = move-straight, avoid = avoid-robot, follow = follow-border, look = look-robot-and-follow-border, toward = move-toward-robot, exit #1 = exit-area-1 and exit #2 = exit-area-2). The bottom part shows the explicit signals produced by the robots.

### 3.2.2 Communication system

As it is the case of the robots described above displaying the symmetrical strategy, the arbitration between different behaviours and the coordination between robots is regulated by the social interaction between the two robots mediated by the evolved communication system.

The behavioural skills that allow the robots to access and to generate information that has a communicative value include: an exploratory behaviour that allows the robots to identify the location of the two target areas, a follow-border behaviour that allows the robots to maintain information over time, and a look-robot-and-follow-border behaviour that allows the robots to identify and assume a specific position in a target area with respect to the location of the other robot. Interestingly, part of the information conveyed through implicit and explicit signals is not simply extracted from the environment but is generated through the behavioural and communicative interaction between the two robots. For example, information that encodes the location of the centre of the two target areas (that cannot be detected directly by a single robot) is extracted by the two robots through a coordinated behaviour that allows the robots to assume a precise relative position in the target area with respect to the other robot.

The signals produced by the robots include two explicit signals (A and B) that encode whether a robot is located outside or inside the black target area, respectively, and an implicit signal that consists of the relative position of the visually perceived robots. The fact that the explicit signals do not differentiate the white target area from the regions outside target areas does not constitute a source of ambiguity since this information is exploited only by robots currently located in target areas and because robots never occupy the same target area.

The effects of implicit and explicit signals consist in a modification of the robots' motor behaviour that is context dependent (i.e. the type of effect produced and/or whether or not the effect will be produced depends on the state of the robot detecting the signal). More precisely:

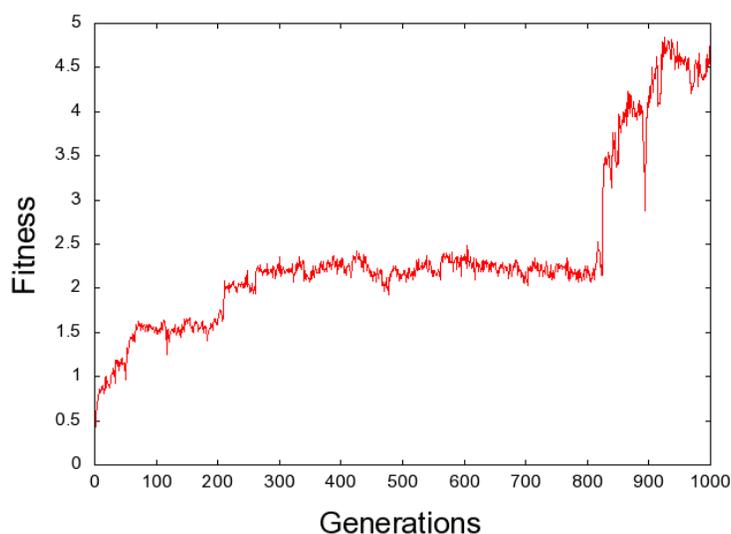
- the perception of signal-B always triggers an exit behaviour in robots located in the black target area;
- the perception of signal-B in combination with an implicit signal constituted by the visual perception of the other robot on the left side of the visual field always triggers an exit behaviour in robots located in the white target area;
- the perception of signal-B in combination with an implicit signal constituted by the visual perception of the other robot triggers a move-toward-robot behaviour in robots located outside target areas that previously visited the white target area.
- the perception of the implicit signal always triggers an avoid-robot behaviour in robots located outside target areas (with the exception of the case reported above that triggers the execution of the move-toward-robot behaviour).

### 3.2.3 Evolutionary origin of robots communicative and non-communicative skills

By looking at how performance varies throughout generations, we can observe

the alternation of periods of rapid improvements and periods of stasis (see Figure 10).

During the first 10 generations the robots already display an ability to concurrently reach the two target areas at least once in most of the trials, but not yet an ability to switch between target areas. This is achieved through the exhibition of an exploration behaviour outside target areas (consisting of a combination of a move-straight/turn-right and obstacle avoidance behaviour) and of a remain-on-target behaviour (in the white target area during the first generations and in the black target area during the successive generations). The remain-on-target behaviour ensures that as soon as the other robot enters in the other area, the two robots are concurrently located in the two target areas as requested. These behavioural skills are preserved in the following phases. The remain-on-target area behaviour is realized by individuals of successive generations through a follow-border behaviour or through a turns-on-the-spot behaviour (that is realized by moving the two wheels at approximately the same speed in opposite directions). The follow-border behaviour, however, represents a crucial pre-requisite for the emergence of the move-toward-robot and look-robot-and-follow-border behaviours as described below. Therefore, when these new behavioural skills are developed, the turn-on-the-spot behaviour extinguishes and the follow-border behaviour fixates in the evolving population. During this phase the robots also develop the avoid-robot behaviour that allows the robot located outside target areas to avoid moving in the direction of the area already occupied by the other robot.



**Fig. 10.** Fitness throughout generations. Data are calculated by testing the ancestors of the best individuals of the last generation for 1000 trials.

From generation 10 to 209, the robots develop a signal-A behaviour (that is produced by robots located outside target areas), a signal-B behaviour (that is produced by robots located in the black target area) and an exit-target-2 behaviour (that is produced by robots located in the black target area detecting signal-B). These new behaviours, that are preserved in succeeding generations, allow the robots to occasionally switch target areas and thus gather additional fitness points. The robots exiting from target areas, however, are not yet able to navigate directly toward the other target area (i.e. they exhibit an exploration behaviour independently from whether they previously visited the two target areas or not).

At generation 210 the robots located in the white target area develop an ability to remain in the area by exhibiting a follow-border behaviour even when they detect signal-B until they do not visually perceive the other robot (Figure 11). This new behaviour allows them to leave when they are oriented toward the direction of the black target area and to navigate directly toward that area through a move-toward-robot behaviour (i.e. without the need to resort on exploration behaviour). The acquisition of this new behavioural skill is reflected by the fact that the percentage of time spent by the robot on the border of an area by moving forward at a significant speed (see section 5.4) significantly increases from generation 209 to generation 210 (Figure 11, frame A). The fact that the exhibition of this behaviour plays an adaptive role in this stage is reflected by the corresponding increase of performance observed during the same generations (Figure 10).

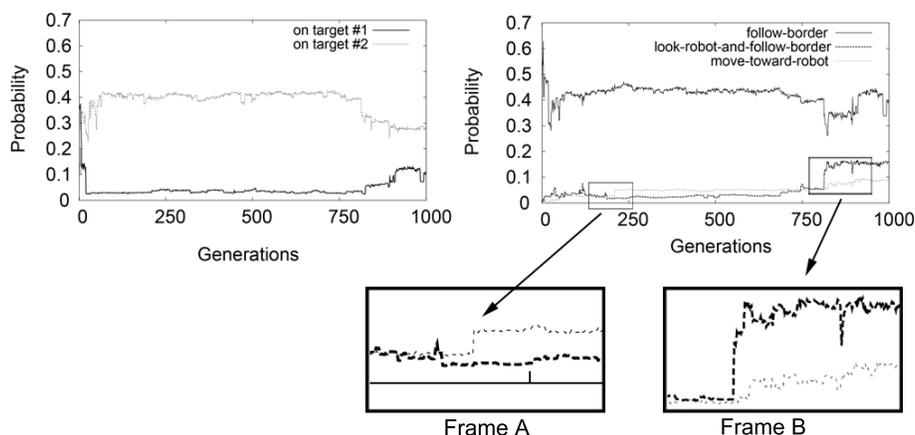
During the long evolutionary phase from 210 to 813 the move-toward-robot behaviour is slowly optimized with respect to the ability to avoid approaching the robot located in the black target area too much (thus triggering an obstacle avoidance behaviour by the robot located in that area) while minimizing the risk of missing the target area itself.

From generation 813 to 820 the robots develop a look-robot-and-follow-border behaviour that allows the robot exiting from the black target area to navigate directly toward the white target area (Figure 11, frame B). This new behaviour, that arises suddenly at generation 813 as a result of a single point mutation and is then refined in the few successive generations, plays a significant adaptive role, as is reflected by the concurrent increase in performance. The development of this new behaviour is reflected by the fact that the percentage of time spent by the robot on the border of the area while being oriented toward the other robot significantly increases from generation 812 to generation 813 (Figure 11, frame B). The fact that the robot occasionally finds itself in this condition also before generation 813 can be explained by considering that the condition used to identify the look-robot-and-follow-border behaviour (see section 5) might occasionally occur also during the execution of the follow-border behaviour.

The development of the look-robot-and-follow-border behaviour represents a remarkable innovation that, in addition to having the functionality described above, allows the robot located in the black target area to reach and remain

in a specific position of the area with respect to the other robot. This creates the basis for the further improvement the efficiency of the move-toward-robot behaviour that allows the robot exiting from the white target area to navigate directly toward the black target area.

From generation 820 on, in fact, the robots develop an ability to navigate exactly toward the centre of the black target area, thus avoiding the risk of missing this target area.



**Fig. 11.** Probability of executing different types of behaviour for the best individuals of successive generations. Data are computed by testing each individual for 1000 trials. For the sake of clarity, the pictures only show the data for a few selected behaviours. The left picture shows the data for the remain-on-target-#1 and remain-on-target-#2 behaviours. The right picture shows the data for the follow-border, the look-robot-and-follow-border, and the move-toward-robot behaviours. The zoomed information displays the significant increase at generation 210 of the probability of executing the move-toward-robot behaviour (Frame A) and the significant increase at generation 813 of the probability of executing the look-robot-and-follow-border behaviour (Frame B). The description of the way in which the probability of executing different behaviours has been computed is provided in section 5.4

## 4 Discussion

The experimental scenario illustrated in this chapter represents a minimal model that allows us to study how communication can evolve in a population of initially non-communicating agents and how agents' communication skills can progressively complexify as the agents adapt to their task/environment. In this section

we will discuss how the obtained results can help us to better answer the general questions identified at the beginning of this chapter.

The first issue that we want to consider is the question of *whether and how communication can evolve in the first place*. As also discussed in Chapter ?? and ??, the evolution of a communication skill requires the development of two complementary but independent abilities: an ability to produce signals that are useful (from the point of view of the signaller, the receiver, or both) and an ability to react to signals in a way that is useful (from the point of view of the signaller, the receiver, or both). As Maynard Smith puts it: “It’s no good making a signal unless it is understood, and a signal will not be understood the first time it is made” (Maynard Smith, 1997). From the point of view of the evolution of explicit signalling capabilities, this implies that variations that lead to the production of a useful signal will tend to be retained only if agents already have the complementary ability to react to that signal in an appropriate way, or vice versa, variations that lead to an ability to react to signals in a useful way tend to be retained only if agents already have the complementary ability to produce the corresponding signal. This means that adaptive variations that lead to the production of useful signals or to the exploitation of signals that are adaptively neutral (and therefore might be lost) until the corresponding complementary condition is met. This aspect seems to indicate that the evolution of communication would be an extremely unlikely event, a consideration that is in contrast to experimental evidence.

This apparent paradox can be solved by hypothesizing that: (a) originally neutral traits can later acquire a communicative function, and (b) traits originally having a certain function can later be exapted (Gould, 1977) to play an additional communicative function. These general hypothesis can be further articulated into two cases depending on whether the pre-existing trait consisted in the ability to produce an action that could potentially assume a communicative value (as proposed by Konrad Lorenz and other earlier ethologists) or in the tendency to react in a certain way to signals that could potentially assume a communicative value (Maynard Smith and Harper, 2003). An evidence supporting the former hypothesis is constituted by the observation that the beak-wiping behaviour serving a preening function displayed by several species of grass finches, in some species plays the role of a courtship signal (Morris, 1958). The results of the simulation experiments reported in Chapter ?? (see also Mirolli and Parisi, 2008) constitute an additional source of evidence for this hypothesis. An piece of evidence supporting the latter hypothesis is constituted by the colourful phenotype of *Papilio memnon* that increases the chances of survival of this species by exploiting the tendency of its predator to avoid distasteful insects characterized by a colourful phenotype (Maynard Smith and Harper, 2003).

The results obtained through the synthetic experiments presented in this chapter confirm that, indeed, communication can emerge despite the traits that are necessary for its emergence — namely, an ability to produce useful signals and an ability to react to signals appropriately — do not, per se, increase the reproductive chances of individuals that possess them. Moreover, the possibility to

analyse the course of the evolutionary process in detail, thanks to the synthetic nature of these experiments, allows us to identify how the problem consisting in the need to develop two interdependent traits that are adaptively neutral in isolation is solved. Indeed, the analysis reported in section 3.1.3 and 3.2.3 indicates that the evolution of communication skills occurs through the exploitation of traits that originally did not serve a communicative function or that did not originally have any functionality.

As an exemplification of a case in which adaptively neutral traits later acquire a communicative functionality we can consider the ability of the robots to produce signal-C or a signal-B depending on whether they are located inside or outside the black target area, the ability to react to the perception of the implicit signal by triggering a move-toward-robot behaviour, and the tendency to react to signal-C by triggering turn-right behaviours (see section 3.1.3). These behavioural traits, that do not play any significant functionality during the first 20 generations, later assume a functional role as a consequence of retained variations that alter the behavioural and communicative skills exhibited by the robot in other circumstances.

As an exemplification of cases in which pre-existing signalling capability acquire a communicative functionality through a variation in the way in which agents react to the signals (and not through a variation of the signal themselves), we can consider the variations occurring at generations 31-40 that makes the robots located in the black target area turn sharply or move forward depending on whether they detect signal B or A, respectively (section 3.1.3). These variations indeed, confer to the two pre-existing signals a communicative function that allows the robots located in the two target areas to switch areas.

An exemplification of a case in which a pre-existing functional signal acquires an additional functionality through a variation in the way in which agents react to the signal (and not through a variation of the signal itself) is constituted by the variations occurring from generation 210 on that lead to the development of the move-toward-robot behaviour (section 3.2.3). Up to this point the signal B, that is produced by robots located in the black target area, triggers an exit behaviour in robots located in the white or black target area (thus allowing the robots to eventually reach the other area later on). From generation 210 on, the same signals play an additional functionality which is realized by triggering a move-toward-robot behaviour in robot located outside target area that previously visited the white target area. The new functionality is achieved through a variation that modifies the way in which the robots react to the signal but not the signal itself or the condition in which the signal is produced.

An exemplification of a case in which a pre-existing ability to react to signals in a specific way acquires a functionality through a variation of the signal produced, but not of the way in which the robots react to the signal, is constituted by the development of the signal-B that triggers the exit-area-2 behaviour. The tendency to react to this signal by exiting from the black target area, in fact, is displayed from generation 10 on. The ability to produce signal-B in the black target area is developed several generations after. The ability to produce this

signal thus immediately assumes a functional role thanks to the exploitation of the pre-existing tendency to react to this signal in that way.

The second issue that we want to consider is the question *how and to what extent robots' behavioural and communication skills increase in complexity*. Complexity can be measured along different dimensions. A first dimension concerns the number of different elementary behaviours produced by the agents. A second dimension concerns the number of signals or combination of signals serving a communicative function which co-determine the expressive power of the communication system. A third dimension concerns the diversification of the effects that each signal produces depending on the context in which the signal is detected. A fourth dimension concerns the ability of the agents to access and to generate information that has a communicative value and that can then be conveyed through communication signals. Finally, a fifth dimension concerns the nature of signals developed, namely whether a signal encodes information directly available through the agents' sensors or more complex re-elaborated information (see below). We discuss the complexification of behavioural and communication skills together since, as we have seen, these skills are strongly interdependent.

The analysis of the evolutionary process indicates that improvements in terms of performance are often correlated with a complexification of agents' skills with respect to one or more of the dimensions described above. The comparative analysis of the two different replications of the experiments (reported in section 3.1 and 3.2) also shows how solutions that are comparable in terms of performance and in terms of overall complexity of the evolved strategy can differ significantly with respect to the complexity along different dimensions.

In general terms, in the case of the two best replications of the experiments described above, evolved individuals display a rich behavioural and communicative repertoire that includes 7-10 different elementary behaviours and 4-6 signals (constituted by different explicit signals or combination of implicit and explicit signals) each producing 1-3 different effects depending on the context in which signals (or combination of signals) are experienced.

With respect to the ability of the agent to access, generate, and elaborate communicative information, in most of the cases explicit signals encode non-abstract information that is directly and currently available through the sensors of the robots and that is accessed through the exhibition of simple behaviours (i.e. an exploration and/or a remain-on-target behaviour). Non-abstract signals of this form do not involve a significant re-elaboration of the sensory information and/or the integration of sensory-motor information through time (Hauser, 1996; Rendall et al., 1999). In the case of the experiment reported in section 3.1, however, we also observed the development of explicit signals that encode abstract information (e.g. the time spent by the robot within a target area).

Moreover, we observed that the implicit signals, and combinations of implicit and explicit signals, also often encode abstract information. This can be explained by considering that the implicit signal constituted by the actual position of a robot often implicitly encodes useful information concerning the sensory and motor states previously sensed by the robot itself. The need to extract and

communicate information about previous experienced sensory states is therefore solved by selecting behavioural skills that allow the robots to integrate and elaborate information by acting in the environment rather than by performing internal operations. An example of an abstract signal is constituted by the combination of the implicit signal and the explicit signal-B produced by the robot located in the black target area (section 3.2) that allows the other robot to infer the direction toward which it should navigate in order to reach the central part of the black target area — information that is not directly available from the state of the robot’s sensors and that reflects the effects of the previous sensory-motor interactions between the robot and the environment (see section 3.2).

All used signals are deictic (i.e. they provide information that is dependent from the current position/state of the sender [Hockett, 1960](#)). Displaced signals (i.e. signals providing information that is independent from the current context of the sender [Hockett, 1960](#)) are not observed. For a simple experimental scenario in which simulated agents develop an ability to communicate the location of spatially distant objects, see [Williams et al. \(2008\)](#). Finally, most of the used signals are informative/manipulative (i.e. they convey information possessed by one of the individuals or in which one individual manipulates a second individual so as to accomplish a certain adaptive function). In a few cases, however, relational signals are also observed. By relational signals we mean signals that are generated through a communicative interaction and that allow a group of individuals to perform collaborative task that cannot be accomplished by a single individual (i.e. signals analogous to the vocal duetting produced by several species that allow them to establish and maintain a pair bond, [Bailey, 2003](#); [Farabaugh, 1982](#); [Haimoff, 1986](#)). An example of relational signal is constituted by the family of signal-A produced by robots outside target areas that assume different values (within the [0.7 to 0.9]) depending on whether the other robot is located outside target areas or not (section 3.1). For other examples see [Marocco and Nolfi \(2007\)](#).

The third issue that we want to consider is *the role of innovations in the evolutionary process and the extent to which the proposed scenario can be considered open-ended*. The analysis of the origins of robots’ behavioural and communicative skills demonstrates how the evolutionary process is strongly influenced by the initial capabilities of the robots. These, in turn, depend on the family of strategies on which the evolutionary process converges in the initial generations as a result of the random generation of the genome of the initial population and as a result of the high stochasticity that characterizes the first phase of the evolutionary process. On the other hand, the analysis of the most successful replications of the experiment also demonstrates how agents’ skills can be progressively transformed and complexified during the adaptive process until optimal or close to optimal solutions are discovered. This remarkable result is achieved because of the ability to preserve previously developed functional traits and the ability to generate new skills by exploiting previously developed skills.

The tendency to preserve previously acquired functional traits can be explained by considering that new skills are often based on previously developed

skills. For example, in the case of the asymmetrical strategy described in section 3.2, the move-toward-robot behaviour that allows robots located in the white target area to navigate directly toward the black target area depends on the exit-area-1 behaviour and the remain-on-target-1 behaviour (that is achieved through the exhibition of a follow-border behaviour) developed previously for other functionalities. In turn, the exit-area-1 and remain-on-area-1 behaviours depend on the signal-B behaviour developed previously for other functionalities. Moreover, the signal-B behaviour depends on the remain-on-area-2 behaviour developed previously for other functionalities. Finally, the remain-on-area-2 behaviour depends on the exploration behaviour developed previously for other functionalities. The establishment of this cascade of dependencies prevents the possibility that variations affecting previously developed skills can be retained. It thus leads to an incremental process in which new skills are added on top of previously developed skills while old skills are preserved.

The ability to generate the required new skills can be explained by considering the possibility to exploit previously developed skills. The development of new skills, in fact, does not only lead to an improvement of agents' performance but often also leads to the establishment of the adaptive condition that enables the development of further and more complex skills. For example, the development of an ability to remain in the black target area creates the basis for the development of a signalling capability that allows other robots to navigate directly toward this target area. Similarly, the ability to remain in the black target area displaying a look-robot-and-follow-border behaviour (by assuming a precise position in the target area with respect to the other robot) creates the condition for the development of an ability to navigate directly toward the centre of this target area. More generally, concerning communicative behaviour, the development of an ability to signal relevant information enriches the perceptual environment of the robots, thus creating the adaptive conditions for the emergence of new skills that are based on the signalled information. Similarly, concerning motor skills, the development of a given behavioural skill often creates the basis for the development of further skills through the synthesis of new traits that, in combination with previously developed traits, lead to the exhibition of new functional skills. In other words, the generation of novelties often creates the adaptive conditions for the development of additional novelties, thus producing an evolutionary process that is open-ended (within the limits imposed by the complexity of the task/scenario).

The combination of these two aspects, i.e. the possibility to generate new skills by exploiting previously developed skills and the tendency to preserve previously developed skills, leads to the progressive complexification of agents' behavioural and communicative skills observed in the experiments reported in this chapter.

Before concluding, we want to briefly discuss the extent to which the contributions provided in this chapter can be generalized to different setups involving artificial agents. We believe, in fact, that although the specific characteristics of the robot certainly affect the obtained results, the dynamics that characterize

how communication skills originate and complexify should generalize to different experimental setups. For example, the origin of communication skills should be observed independently of whether the robot can produce and detect radio, or sound, or light signals (i.e. independently from the communication channel available) provided that robots have the possibility to vary the signal produced and to vary the way in which they react to detected signals. Similarly, the origin of implicit communication forms should be observed independently from how robots get information from the other individuals and from whether the implicit signals encode information about the relative position of the other individuals or about the action performed by these individuals. Similarly, we hypothesize that the emergence and the progressive complexification of behavioural and communication skills can be observed independently from the task and the characteristics of the evolutionary algorithm provided that: (i) the task requires an ability of the robots to coordinate and cooperate and permits a variety of partial solutions that represent stepping stones toward achieving the final solution (ii) the variations introduced during the evolutionary process affect the rules that regulate the fine-grained interaction between the robots and the environment that in turn affects the behavioural and communication behaviours exhibited by the robots. Moreover, variations are retained or discarded on the basis of their effect at the level of the global behaviour exhibited by the robots.

The process that leads to the final communication system, the evolution of previously developed communication skills, and the role of implicit and explicit communication, on the other hand, might differ significantly in cases in which there are conflicts of interests between the individuals, and/or in cases in which the characteristics of the interacting individuals differ significantly (due to differences in their genetic characteristics or due to variations occurred during their lifetime). The simple experimental scenarios described in Chapters ?? and ?? provide an analysis of these issues concerning the origin of simple communication skills. However, further work is necessary in order to analyse the impact of these aspects on how simple communication skills might vary and complexify during the successive course of the evolutionary process.

Finally, an important issue concerns the potential impact that this type of synthetic experiments might have on the modelling of the evolution of animal communication. We believe that in addition to providing evidence that might support or disconfirm existing theories (as discussed above in this section), the analysis of the results obtained in these experiments might lead to the identification of general mechanisms that could generalize to natural systems and that might be tested through empirical studies. In particular, in the case of the experiments reported in this chapter, we believe that they demonstrate that signals and signals meaning are grounded (Harnad, 1990) not only in agents' sensory states but also in agents behavioural skills and that the co-adaptation of agents behavioural and communicative skills represent a crucial prerequisites for the emergence and progressive complexification of agents' communicative skills.

## 5 Appendix

This appendix describes in details the sensory and motor system of the robots (section 5.1), the way in which neurons are updated (section 5.2), the way in which the robot/environmental interactions have been simulated during the evolutionary process (section 5.3), and the criteria used to analyze the different behaviours exhibited by the robots (section 5.4).

### 5.1 Sensors and actuators

- The 8 infrared sensory neurons encode the activation state of the 8 corresponding infrared sensors of the robots normalized in the range  $[0.0, 1.0]$ . To compensate for the significant differences in the resting level of different individual sensors, a constant value corresponding to the resting level of each sensor measured far from obstacles is subtracted from the current state of each sensor.
- The 3 visual sensory neurons encode whether another robot is detected through the camera (that has a view angle of  $36^\circ$ ) or not and also the relative angular position of this detected robot. This information is computed by elaborating on an image of  $60 \times 15$  pixels. In particular, the activation of the first neuron is set on or off depending on whether two or more pixels located in different columns of the picture have a predominant red colour. The state of the 2 other sensory neurons encode the offset on the left and right side, respectively, of the barycentre of the red pixels corresponding to the visually detected robot normalized in the range  $[0.0, 1.0]$ . For example, when the camera does not have any predominant red pixel, the state of these two sensors is set to  $[0.0, 0.0]$ . When the barycentre of the detected red pixels is at  $0^\circ$ , the state of the two sensors is set to  $[1.0, 1.0]$ . When the barycentre of the detected red pixels is at  $9^\circ$  on the right side, the state of the two sensors is set to  $[0.0, 0.5]$ . When the barycentre of the detected pixels is at  $6^\circ$  on the left side the state of the two sensors is set to  $[0.66, 0.0]$ .
- The 2 ground sensory neurons binarily encode whether the central ground sensor located on the frontal side of the robot currently detects a white or a black colour (i.e. the white or black target area, respectively). The 2 ground sensory neurons are turned on when the activation of the ground sensors is in the range  $[415, 900]$  or  $[0, 414]$ , respectively.
- The 2 additional ground sensory neurons encode the previous state of the 2 corresponding ground sensory neurons. In one set of experiments the state of these neurons is set equal to the state of the ground sensor at the previous time step. In a second set of experiments the state of these neurons was updated on the basis of the following equation:

$$O_i(t) = 0.99G_i(t-1) + 0.01G_i(t) \quad (1)$$

where  $O_i(t)$  represents the output of the neuron  $i^{th}$  at time  $t$ ,  $G_i(t-1)$  represents the output of the corresponding ground sensor at time  $t-1$ , and

$G_i(t)$  represents the output of the corresponding ground sensor at time  $t$ . No significant differences were observed between the two sets of experiments based on the two different encodings.

- The 2 signal sensory neurons encode the output of the communication motor neuron of the other robot (see below) and the output of the communication motor neuron of the same robot in the previous time step.
- The first 2 motor neurons encode the desired speed of the two wheels of the robot. The activation of each neuron is linearly scaled in the range  $[-2\pi, 2\pi]$  rad/sec, and used to set the desired angular speed of the corresponding motor.
- The communication motor neuron encodes the signal that is transmitted to the other robot through the wireless Bluetooth connection.

## 5.2 Update functions of the neurons

- The output of the motor neurons at time  $t$  is computed as the weighted sum of all inputs units and bias, filtered through a sigmoid functions:

$$O_j(t) = \sigma \left( \sum_i w_{ij} I_i(t) + \beta_j \right), \quad \sigma(z) = \frac{1}{1 + e^{-z}} \quad (2)$$

where  $I_i(t)$  corresponds to the activation of the  $i^{th}$  neuron at time  $t$ ,  $w_{ij}$  is the weight of the synaptic connection between the input neuron  $i^{th}$  neuron and the current neuron  $j$ , and  $\beta_j$  is a bias term.

- The output of the internal neurons at time  $t$  is computed on the basis of the following equation:

$$O_i(t) = \tau O_i(t-1) + (1 - \tau) O_i(t) \quad (3)$$

Where  $O_i(t-1)$  represents the output of the neuron at time  $t-1$ , and  $O_i(t)$  represents the weighted sum of all input units and bias filtered through a sigmoid function (see above),  $\tau$  represents a time constant ranging between  $[0.0, 1.0]$ .

## 5.3 Simulation

The movement of the robots in simulation has been computed by using a simple kinematic model (i.e. the role of friction and inertia has been simulated in consideration of the light weight and the limited maximum speed of the e-puck robots). The state of the infrared sensors has been computed by using a sampling technique (Miglino et al., 1995). The shape of the robot body and the position of the sensors have been simulated with an accuracy of floating point precision. The state of the ground sensors has been computed on the basis of the current position of the ground sensor on the simulated environment. The state of the vision sensors has been computed on the basis of the relative angle between the barycentre of the two robots bodies and of their directions. The state

of the communication sensors has been updated on the basis of the state of the communication motor neuron of the other robot and of the robot itself at the previous time step. The position of the robot and the state of the sensors have been updated every 100ms (i.e. at the same time rate of the neural controller).

#### 5.4 Criteria used to identify the behavior exhibited by the robots analyzed in section 3.2

To identify the behavior exhibited by the robot during the robot/environmental interaction we identified the conditions (reported in the table below) that are mutually exclusive and that can be used to estimate automatically the behavior exhibited by the robot in any time step (i.e. every 100ms) within a list of possible behaviors. The probability of exhibiting each behavior of the list is later determined by calculating the percentage of time steps that have been assigned to the corresponding behavior. See Table 1

### Supplementary Data

Supplemental data for this article including movies of the behaviours displayed by evolved robots of different replications of the experiment can be found at the following address: <http://laral.istc.cnr.it/esm/evo-communication>. An open-software for replicating the experiments in simulation and hardware including the source codes, a manual, a tutorial, and the sample files of the experiment can be downloaded from <http://laral.istc.cnr.it/evorobotstar> (see also Chapter ?? in the Appendix of this book)

### Acknowledgments

The authors thank Davide Marocco for his important suggestions on the definition of the experimental scenario and on the analysis of the obtained results.

**Table 1. XXXXXX**

Behaviour	Condition
Obstacle-avoidance	1. infrared sensors are active
avoid-robot	1. infrared sensors not active 2. not on target areas 3. robot visually detected at least once in the last 5 time steps 4. turning angle over preceding 9 time steps $> 2^\circ$
follow-border	1. infrared sensors not active 2. ground sensor located on the border of a target area 3. average distance traveled over 5 time steps $> 2\text{mm}$
look-robot-and-follow-border	1. infrared sensors not active 2. ground sensor indicates target border 3. robot visually detected at least once in the last 5 time steps
move-toward-robot	1. infrared sensors not active 2. ground sensor not active 3. robot visually detected at least in the last 5 time steps 4. average distance traveled over 5 time steps $> 2\text{mm}$ 5. turning angle over preceding 9 time steps $< 2^\circ$
exit-area-1	1. infrared sensors not active 2. ground sensor detected the white target area at time t-1 but does not detect the white target area at time t.
exit-area-2	1. infrared sensors not active 2. ground sensor detected the black target area at time t-1 but does not detect the black target area at time t.
move-straight	1. infrared sensors not active 2. no other of the conditions apply (i.e. default behaviour when infrared sensors are not active).

## Bibliography

- Bailey, W. J. (2003). Insect duets: underlying mechanisms and their evolution. *Physiological Entomology*, 28:157–174.
- Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9:255–267.
- Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent (with commentary and response). *Adaptive Behavior*, 11(4):209–243.
- Billard, A. and Dautenhahn, K. (1999). Experiments in learning by imitation grounding and the use of communication in robotic agents. *Adaptive Behavior*, 7(3/4):415–434.
- Cangelosi, A. and Parisi, D. (1998). The emergence of a language in an evolving population of neural networks. *Connection Science*, 10:83–97.
- Cangelosi, A. and Parisi, D. (2001). How nouns and verbs differentially affect the behavior of artificial organisms. In Moore, D. and Stenning, K., editors, *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pages 170–175. London: Lawrence Erlbaum Associates.
- Di Paolo, E. (1997). An investigation into the evolution of communication. *Adaptive Behavior*, 6(2):285–324.
- Di Paolo, E. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):25–46.
- Farabaugh, S. M. (1982). The ecological and social significance of duetting. In Kroodsma, D. E. and Miller, E. H., editors, *Acoustic Communication in Birds*, pages 85–124. New York: Academic Press.
- Gould, S. (1977). *Ontogeny and Phylogeny*. Harvard University Press: Cambridge, MA.
- Haimoff, E. (1986). Convergence in the duetting of monogamous old world primates. *J. Hum. Evol.*, 15:767–782.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.
- Hauser, M. D. (1996). *The Evolution of Communication*. MIT Press / Bradford Books, Cambridge, MA.
- Hockett, C. (1960). The origin of speech. *Scientific American*, 203:88–96.
- Marocco, D. and Nolfi, S. (2007). Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Con-*

- nection Science*, 19(1):53–74.
- Maynard Smith, J. (1997). *The Theory of Evolution*. C.U.P, 3rd edition.
- Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press.
- Miglino, O., Lund, H., and Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, 4(2):417–434.
- Mirolli, M. and Parisi, D. (2008). How producer biases can favor the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16(1):27–52.
- Mondada, F. and Bonani, M. (2007). The e-puck education robot.
- Morris, D. (1958). The comparative ethology of grass-finches (erythrurae) and mannikins (amadinae). In *Proc. Zool. Soc. London*, volume 131, pages 389–439.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books.
- Nolfi, S. and Marocco, D. (2001). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences*, 120:287–310.
- Quinn, M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life: Sixth European Conference on Artificial Life (ECAL 2001)*. Berlin: Springer Verlag.
- Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences*, 361:2321–2344.
- Rendall, D., Cheney, D., Seyfarth, R., and Owren, M. (1999). The meaning and function of grunt variants in baboons. *Animal Behaviour*, 57:583–592.
- Steels, L. (1999). *The Talking Heads Experiment*. Antwerpen, Laboratorium.
- Steels, L. and Kaplan, F. (2001). Aibo’s first words: The social learning of language and meaning. *Evolution of Communication*, 4:3–32.
- Sugita, Y. and Tani, J. (2004). A connectionist approach to learn association between sentences and behavioral patterns of a robot. In Schaal, S., LjSpeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J., editors, *Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB04)*. Cambridge, MA: The MIT Press.
- Trianni, V. and Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231.

Williams, P., Beer, R., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*. Cambridge, MA: MIT Press.