

Evolution of Communication in Robots

Joachim de Greeff & Stefano Nolfi

Institute of Cognitive Sciences and Technologies, CNR
Via S. Martino della Battaglia, 44, 00141, Roma, Italy

1. Introduction

During the last ten years, the attempt to study the evolution of communication and language through computational and robotic models has attracted the attention of an increasing number of researchers (for a review see: Cangelosi and Parisi, 2002; Kirby, 2002; Steels, 2003; Wagner et al., 2003); Nolfi, 2005; Nolfi & Mirolli, in press). Indeed, 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 to study how communication signals are grounded in agents' non-symbolic sensory-motor experiences, and (b) it allows 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.

Within this area, evolutionary robotics can provide a key contribution because some of its foundational characteristics that differentiate it from other alternative learning methods; that is, the fact that fine-grained characteristics that regulate how the robots interact with the physical and social environment can be encoded into free parameters and the fact that variations can be retained or discarded on the basis of their affect at the level of the global behaviour exhibited by the robot/robots (Nolfi, in press). These features, in fact, allow the experimenter to reduce the number of characteristics that are predetermined and fixed to the minimum and leave the robots free to determine their way to solve the adaptive problem.

These characteristics also allow us to study whether and how communication can emerge in population of individuals which are not rewarded directly for communicating. Moreover, they allow us to study the role of the co-adaptation of behavioural and communication skills (Nolfi, 2005) which, as we will show in next sections, represents an essential pre-requisite for the emergence and complexification of robots' communication skills.

In this chapter we describe an experimental scenario which is simple enough to be analysed systematically but which at the same time includes all the elements that are necessary to investigate important questions concerning the evolution of communication such as: what are the conditions that might lead to the evolution of communication skills in a population of initially non-communicating robots? 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 this grounded (Harnad, 1990) in agents' sensory experience? The key aspects of the chosen scenarios are: (i) the fact that the task/environment allows qualitatively different solutions, (ii) the fact that the robots are provided with a sensory-motor system which 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 analysis of these experiments might allow us to generate new data which can partially compensate the paucity of empirical data caused by the fact that language and

communication do not leave direct traces in fossils records. More generally, as we will discuss in section 4, the analysis of these synthetic experiments provide hints for confirming or disconfirming existing theories on the evolution of communication and language as proposed by evolutionary biologists, as well as for formation of new explanations.

Finally, as we will also discuss on the concluding section, the analysis of these experiments provide interesting hints for understanding the role of social interactions in evolutionary robotics, with particular reference to the role that sociality might have in the manifestation of open-ended evolutionary processes.

This rest of the chapter is organized as follows. In the next section we describe the experimental scenario. In section 3 we describe the obtained results and in section 4 we discuss their implications.

2. Experimental setup

The experimental setup involves two wheeled robots situated in an arena containing two target areas (Figure 1) which 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 to identify 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 robots' body and neural controller, and of the evolutionary algorithm.

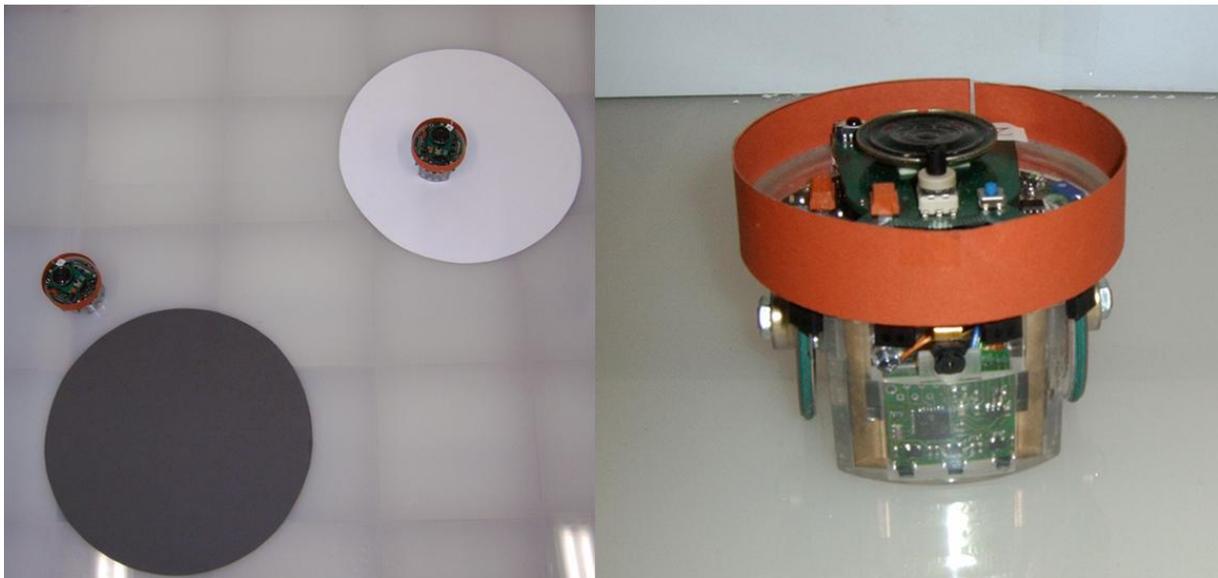


Figure 1. Left: The environment and the robots. The two circular areas of the environment coloured in black and white represent the two target areas. Right: The e-puck robotic platform including the ground sensor board and a stripe of red paper around the top part of the body, that allows for easier visual recognition.

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 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 platform consists of two e-Puck robots (Mondada & Bonani, 2007) equipped with the ground sensor board extension. The robots, which 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° 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. For more details, see De Greef and Nolfi (in press).

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 consists of 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 binary 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 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 a time constant parameter. The motor neurons consist of standard sigmoid units. For more details, see De Greef and Nolfi (in press).

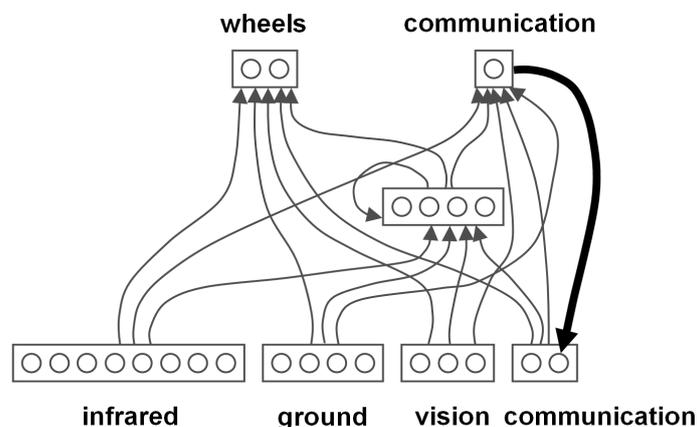


Figure 2. The architecture of the robots' neural controller. The lower, middle and top layers indicate sensory, internal and motor neurons, respectively. Thin arrows indicate connections. The thick arrows indicates that the state of the communication motor neuron at time t is copied into the state of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

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 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 distinguishing 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 (Nolfi and Marocco, 2001; Beer, 2003). 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 orientation of the robots are assigned randomly.

Each team of evolving robots scores 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 are evolved in simulation and 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 characteristics 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 pre-determined aspect is that the experimental scenario involves a full cooperative situation. This is due to the fact that the two robots forming a team have the same genetic characteristics and that selection operates on the basis of the performance of the team and not on the performance of a single individuals (Floreano et al, 2007). A second pre-determined aspect 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 on the one hand it depends on fixed aspects like the fitness function and characteristics of the robots' body and of the environment, but on the other hand also on the characteristics of the social environment (i.e. the other robots' behaviour), that is not pre-determined and varies during the evolutionary process. 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 may 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 robots' communicative behaviours, evolving robots are left free to determine how to use 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 work of Di Paolo (1997, 2000), Quinn (2001), Quinn et al. (2003), Baldassarre et al. (2003), Trianni and Dorigo (2006), Marocco and Nolfi (2007), Williams, Beer, Gasser (2008). 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 problem through qualitatively different strategies by exploiting the possibility to communicate through explicit and implicit signals (section 3.1). The analysis of the best solutions indicates that evolving robots display rather rich behavioural and communication skills, including: the ability to access/generate information which 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 (section 3.2). The analysis of the evolutionary development of the best replications sheds light on how signals/meanings originate and how robots behavioural and communication skills progressively complexify as a result of an incremental process. New skills are added on top of previously developed skills which thus represent a pre-requisite for the development and the exhibition of the new skills and which are retained during the successive course of the evolutionary process (section 3.2.4 and 4). The analysis of evolved behavioural and communication skills and of their origin also provides insights on the mechanisms which allow evolving robots to solve the problem resulting from the need to develop an ability to produce useful signals and to react to signals appropriately at the same time (section 4).

3.1 Performance and evolved strategies

By analysing the results obtained at the end of the evolutionary process for different replications of the experiment and for different experimental conditions (i.e. 110x110 and 150x150 cm arenas) we observed that evolved robots display an ability to be concurrently located in the two areas and to switch area several times in the case of the best replications and only a few times in the case of the worst replications. More precisely, the mean number of switches made (+ 1 point for finding the two target areas for the first time) is 10.035 and 4.680 for the best individuals evolved and tested in the 110x110 and 150x150 environments, respectively.

Evolved robots exploit the possibility to communicate through explicit signals in most of the replications. Indeed, by analysing the variance of the performance obtained in a standard condition and a control condition in which the evolving robots are forced to always produce a 0.0 signal, we observed that the ability to vary the explicit signals significantly affects the overall performance of the robots (Kruskal-Wallis test, $p < .0005$ for both the experiments performed in the small and large arenas). Evolved robots also rely on implicit communications in most of the cases, as we will illustrate below.

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 robot.

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) we observed that the behaviours exhibited in hardware are qualitatively very similar to those shown in simulation. Examples of the best evolved behaviours both in simulation and in the real environment can be seen at the following webpage <http://laral.istc.cnr.it/esm/evo-communication>).

3.2 Detailed analysis of an exemplar solution (asymmetrical strategy)

In this section we describe in details the behavioural and communication skills of the best evolved robots of the best replication of the experiment performed in the 150x150cm arena, that display an asymmetrical strategy. Moreover, we describe the origin of such skills by analysing how the behavioural and communication skills exhibited by robots of succeeding generations vary over the course of evolution.

To perform this analysis we divided the overall behaviour exhibited by the robots into a list of selected elementary motor and communicative behaviours corresponding to sequences of sensory-motor interactions that produce a given functionality (e.g. that allow a robot to avoid an obstacle, or to move toward the other robot located in the other target area, or to produce a signal that allows the other robot to exit from its current area when the two robots are concurrently located in the two areas). The division of the robots' overall individual behaviour into individual elementary behaviour has been realized through the use of mutual exclusive conditions (for details see De Greeff and Nolfi, in press). For example, the sequences of robot/environmental interactions in which the robots' infrared sensors are activated above a given threshold are classified as an obstacle-avoidance behaviour (until the infrared sensors are no more activated). Similarly, the sequences of robot/environmental interactions in which a robot is located on the border or a target area and in which this robot moves forward (and turns slightly left or right) are classified as follow-border behaviour (provided that the robots' infrared sensors are not activated).

For reasons of clarity, behaviours having similar functions and/or constituted by sequence of similar, but not necessarily identical, sensory-motor interactions are grouped into the same elementary behaviours. For example, sequences of sensory-motor interactions in which the agents produce similar, although not identical, explicit signals are grouped into the same elementary signalling behaviour provided that the effect of the signals produced have a similar effect on the other robot.

3.2.1 Motor and communication behaviours repertoire

In this section we describe the elementary motor and communicative behaviour exhibited by the best robots of the last generation. For each elementary behaviour we briefly describe the functionality of the behaviour (with respect to the task), the conditions in which it is executed, and the actions which are produced during its execution.

- A **signal-A** behaviour that consists of 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 of 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 by following the border of an area. The follow-border behaviour is always produced by robots located in the black area that do not visually perceive the other robot, regardless of any perceived signal. This behaviour originates evolutionarily from the modification of a *remain-on-black-area behaviour* that allows the robot to remain on the target area by producing circular trajectories without necessarily moving along the border independently from whether the other robot is visually perceived or not (see section 3.2.3).
- An **avoid-robot** behaviour that consists of a sequence of left-turning movements that make the robot turn on the spot until the other robot exits from its field of view. This behaviour is produced by robots located outside areas that visually perceive the other robot in all cases except cases in which additional conditions trigger the execution of the move-toward-robot behaviour.
- 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 remain on the border of the area while maintaining the other robot on the left-side of its field of view. This behaviour also allows the robot to reach a particular location in its target area with respect to the other robot located in the other target area and hence with respect to this

other target area (see below). 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-from-white-area** 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-from-black-area** 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.

The identification of the robot's elementary behaviours, in this case, is simplified by the fact that this robot displays a reactive behaviour, i.e. always reacts in the same way to the same sensory states. For an analysis of other individual solutions in which the internal dynamic occurring within the agents' control system play a significant role, see De Greeff & Nolfi (in press).

3.2.2 Arbitration and combination of the elementary behaviours

To illustrate how the elementary behaviours described above are combined and arbitrated to solve the robots' adaptive task we will describe a typical trial (see the videos available from <http://laral.istc.cnr.it/esm/evo-communication>).

At the beginning of a trial the two robots are located outside target areas. 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 through 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 in a normal condition does not significantly differ from the performance in a control condition in which the robots located outside target areas were not allowed to visually detect the other robot. The signalling behaviours 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. 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 (mean 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 that 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.

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 an exit-area-1 behaviour that allows it to exit from the area and to initiate a move-straight behaviour toward the black target area.

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 former 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 while 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. The orientation of the robot exiting from the black target area (that depends on the relative position assumed by the robot in this 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.

3.2.3 Communication system

In this section we focus on the communication system possessed by evolved robots and on the relation between robots' behavioural and communication skills. More precisely, we will describe the motor behaviours that allow the robots to access the information that has a communicative value, the explicit and implicit signals produced, and the (context dependent) effect of the detected signals.

The elementary behaviours that allow the robots to access and to generate information that has a communicative value include: an exploration behaviour (a combination of an obstacle-avoidance and a move-forward behaviour) that allows the robots to identify the location of the two target areas, a follow-border behaviour that allows the robots to maintain this 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 constituted by the body of a robot itself that can be visually detected by the other robot and that provides an indication of its relative position. 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).

3.2.4 Evolutionary origin of robots' motor and communicative skills

The analysis of individuals of successive generations indicates that the behavioural and communication repertoire exhibited by the robots progressively complexifies throughout generations as described below.

- 1) In the very first generations the robots develop an *exploration* behaviour that consists of the combination of a *move-forward* and *obstacle-avoidance* behaviour. The exhibition of these behaviours allows the robots to occasionally score 1 point when they happen to transit over the two target areas at the same time.
- 2) During generations 5-10 the robots develop a *remain-on-black-area* behaviour that allows them to remain in the black area when they enter into it. The exhibition of this new behaviour increases the probability that the two robot happen to be concurrently located in the two areas since it eliminates a situation in which the latter robot enters the white area, while the former robot already has abandoned the black area.

- 3) The development of a capacity to remain on the black area, however, also has an additional function; it allows the robot located in the black area to access information which is potentially useful for the other robot when it reaches the white area (information that, as we will see, may allow the other robot to decide whether it should remain or exit from its area). This creates the conditions for the development of an ability to communicate to the other robot whether a robot is located in a black area or not through the production of two different signals (A and B).
- 4) The development of the *signal-A* and *signal-B* communication behaviours described above does not lead to an improvement in performance in itself, but creates the adaptive conditions in the next generations for the development of an *exit-from-black-area* behaviour which is executed by robots located in the black area detecting the signal-B produced by another robot also located in the black area. This new behaviour allows the robots to occasionally exchange areas in the following situation: a robot that visited the white area enters the black area which already contains the other robot; through the *exit-from-black-area* behaviour this other robot exits from the black area and later reaches the white area.
- 5) The development of the behavioural and communication skills described above, in turn, creates the conditions for the development of a *remain-on-the-white-area* behaviour that allows the robots to remain in the white area until they do not detect the signal that indicates that the other robot is located in the black area. The development of this new behavioural skill eliminates the problem caused by the fact that while the second robot enters in the black area the first robot has already exited the white area.

At this stage of the evolutionary process, the robots are able to reach the two areas through a exploration behaviour, to remain on the black area unless the other robot also enters the black area. On the basis of these skills, they are able to be located in the two areas for the first time in most of the trials but they are able to switch areas only occasionally. In many cases the trial ends before the robot exiting from the black area succeeds in finding the white area, because, after exiting from the black area, it resumes on a simple but time consuming exploration behaviour.

During the next phase of the evolutionary process, however, the robots manage to develop new additional skills that allow them to switch areas several times in most of the trials:

- 6) Around generation 205, the robots develop a *move-toward-robot* behaviour that allows the robot exiting from the white area to navigate toward the robot emitting signal-B and therefore directly toward the black area. This new behavioural skill drastically reduces the time needed by the robot located in the white area to reach the black area.
- 7) Finally, after a long substantially stable phase, at generation 814, the robots develop a new way to remain in the black area that consists in remaining in the border of the area itself while looking to the other robot (i.e. the *look-robot-and-follow-border* behaviour). Also in this case, this new behaviour is realized through the exploitation of the explicit and implicit signals produced by the robots located outside the black area which in turn are based on the ability to produce the behavioural skills previously developed. The function of this new behaviour is to ensure that the robot located in the black area positions itself toward the other robot located in the white area (or travelling from the white to the black area) and therefore toward the white area. This, in turn, ensures that when the robot exits from the black area, it will travel directly toward the white area (i.e. toward the direction previously occupied by the other robot). The development of this new behaviour also creates the adaptive conditions for a further improvement of the *move-toward-robot* behaviour previously developed. Indeed, the exhibition of *look-robot-and-follow-border* behaviour implies that the robot located in the black area assumes a

specific relative position with respect to the other robot located in the white area (i.e. the left side of the area with respect to the other robot). The fact that the robot located in the black area now assumes such specific position allows the robot travelling from the white to the black area through the *move-toward-robot* behaviour to position itself toward the centre of the black area, thus minimising the risk to miss this target. This ability is refined in the following generations.

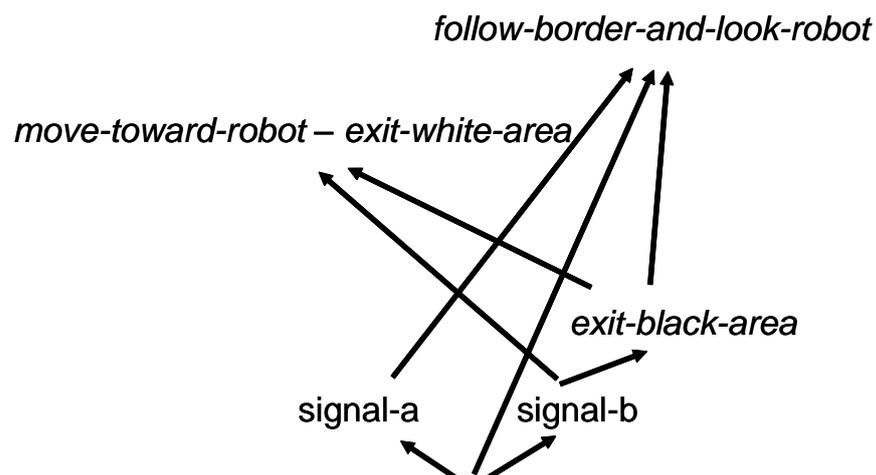


Figure 3. Schematic representation of the relations between different behavioural and communication skills. The Y-axis indicates the course of the evolutionary process and the order in which skills are developed. The arrows summarize the most important dependencies between the different elementary skills. More precisely, the arrows pointing to a given capacity indicate that the development of the skills at the bottom of the arrow created the adaptive conditions for the development of the new capacity and the fact that the new capacity is based (i.e. depends) on the previously developed skills.

Overall this analysis shows how the behavioural and communication skills developed by the robots at a certain stage of the adaptive process often create the conditions for the development of further skills with additional functionalities, are based on previously developed skills. With the sentence “create the conditions for the development of further skills” we mean that the skills that have been developed after would not have been developed (or would have had a lower probability to be developed) without the previously developed skills. With the sentence “..are based on previously developed skills” we mean that the newly developed skills require the previously developed skills and that the eventual lost of one or more of the previously developed skills would also imply the lost of the newly developed skills.

The development of new skills that are based on previously developed skills implies that old skills tend to assume additional functions (i.e. support of these newly developed skills). The creation of these chains of dependencies explains why the adaptive processes observed in this and other replications of the experiments can be described fundamentally as an incremental process in which new skills are often developed on top of previously developed skills and in which previously developed skills tend to be preserved in successive generations.

Moreover, the observation that new behavioural and communication skills are often based on simpler previously developed skills implies that the signals which are produced and exploited by the robots are not simply ‘grounded’ on robots’ sensory-motor states but also on robots’ behaviours.

4. Discussion

We believe that the experimental scenario illustrated in this chapter represents a minimalistic model which allows us to study how communication can evolve in a population of initially non-communicating robots and how robots’ communication skills can progressively complexify as they 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?* The evolution of a communication skill, in fact, requires the development of two complementary but interdependent 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 the same signals, but not to both, are adaptively neutral unless the two abilities are developed at the same time. This aspect seems to indicate that the evolution of communication would be an extremely unlikely event, a consideration which 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 playing 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 consists of the ability to produce an action which could potentially assume a signalling value (as proposed by Konrad Lorenz and other earlier ethologists) or in the tendency to react in a certain way to signals which could potentially assume a communicative value (Maynard Smith and Harper, 2003). 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). Evidence supporting the latter hypothesis is constituted by the colourful phenotype of *Papilio memnon* which 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 paper confirm that indeed, communication can emerge despite the fact that the traits which are necessary for its emergence – namely, an ability to produce useful signals and an ability to react to signals appropriately – taken in isolation are adaptively neutral in that they 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 which are adaptively neutral in isolation is solved. Indeed, the analysis reported in section 3.2.4 indicates that the evolution of communication skills occurs through the exploitation of traits which originally did not served a communicative function or which did not originally play any functionality.

An example of a case in which pre-existing signal acquires a communication 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 3.2.4). Up to this point the signal-B, which is produced by robots located in the black area, triggers an exit behaviour in robot located in the white area (thus allowing the robots to eventually reach the other area later on). From generation 210 on, the same signal plays an additional functionality which is realized by triggering a move-toward-robot behaviour in robots that previously visited the white 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 conditions in which it is produced.

An example 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 which triggers the exit-black-area behaviour. The tendency to react to this signal by exiting from the black area, in fact, is displayed already from generation 10 on (which was be observed in a control condition in which one robot is forced by the experimenter to produce the signal-B while the other robot is located in the black area). The ability to produce signal-B in the black area is developed several generations after. The presence of a trait that makes the robot exit from the black area in reaction to signal-B is thus exploited by developing an ability to produce the signal-B in a situation in which the *exit-black-area* behaviour has a functionality.

The second issue that we want to consider is the question of *how and to what extent the evolved communication system can complexify?* 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 which has a communicative value and which 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 (Mirolli & Nolfi, in press).

The analysis of the evolutionary process described above 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 different replications of the experiments also shows how solutions which 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 the case of the two best replications of the experiments performed in large and small arenas, evolved individuals display a rich behavioural and communicative repertoire which 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 robots to access, generate, and elaborate communicative information, in most of the cases explicit signals encode non-abstract information which is directly and currently available through the sensors of the robots and which 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 symmetrical strategy, however, explicit signals encoding abstract information are also observed (see the analysis reported at <http://laral.istc.cnr.it/esm/evo-communication>).

Implicit signals and combination of implicit and explicit signals, on the other hand, 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 states and motor actions previously perceived and performed by the robot itself. The need to extract and communicate information about previous experienced sensory states therefore is solved by selecting behavioural skills which 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 a robot located in the black area which allows the other robot to infer the direction in which it should navigate to reach the central part of the black area -- information which is not directly available from the state of the robot's sensors and which reflects the effects of the previous sensory-motor interactions between the robot and the environment (see section 3.2.4).

All used signals are deictic (i.e. they provide information which is dependent from the current context experience by the sender (Hockett, 1960). Displaced signals (i.e. signals providing information which is independent from the current context of the sender (Hockett, 1960) are not observed. 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 few cases, however, relational signals are also observed. By relational signals we mean signals that are generated through a communicative interaction and which allow a group of individuals to perform collaborative task which cannot be accomplished by a single individuals (i.e. signals analogous to the vocal duetting produced by several species which allow them to establish and maintain a pair bond, Bailey, 2003; Farabaugh, 1982; Haimoff, 1986). An example of relational signal is observed in the experiment displaying the symmetrical strategy. The signalling interaction occurring between the two robots allows the robots to produce two different signals which are produced, respectively, when only one robot or both robots are located outside target areas (see De Greef and Nolfi, in press). On the evolution of relation signals in a similar experimental setting see also Marocco and Nolfi (2007).

The third issue that we want to consider is *the role of innovations in the evolutionary process and the relation between social/communicative interaction and open-ended evolution*. 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 very initial generations as a result of the random generation of the genome of the initial population and as a result of the high stochasticity which characterizes the selection process in the very first generations. On the other hand, the analysis of the most successful replications of the experiment also demonstrates how robots' skills can be progressively transformed and how the number and the complexity of the robots' elementary skills can increase during the adaptive process until optimal or close to optimal solutions are discovered. As we mentioned

in section 3.2.4, such progressive complexification of robots' skills seem to occur as a result of an incremental process in which the development of new skills often creates the adaptive condition for the development of further skills and in which previously developed skills tend to be retained.

The tendency to preserve previously developed skills can be explained by considering that new skills often exploit (are based on) previously developed skills. For example, in the case of the experiment displaying the asymmetrical strategy described above, the *move-toward-robot behaviour* that allows the robots located in the white area to navigate directly toward the black area depends on the follow-border behaviour exhibited by robots remaining on the black area that was initially developed to allow the robots to reach the two areas for the first time and not for switching areas. Moreover, the *move-toward-robot behaviour* also depends on the *signal-B* behaviour that was previously developed to allow one robot to exit from the black area when the other robot is also located in the same area. This means that the *move-toward-robot behaviour* is based on the other two previously developed behavioural skills and that the development of the *move-toward-robot behaviour* causes the *follow-border* and *signal-B* behaviours to acquire additional functionality – that of supporting the *move-toward-robot behaviour*.

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 which enables the development of further and more complex skills. For example, the ability to remain on the black area by displaying a *look-robot-and-follow-border* behaviour (i.e. 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 leave the white area by navigating toward the centre of the black area. More generally, concerning communication and social interaction, 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 which exploit information encoded in detected signals. In other words, the generation of novelties often creates the adaptive conditions for the development of additional novelties, thus producing an evolutionary process which is open-ended (within the limits imposed by the complexity of the task/scenario).

Supplementary Data

Supplemental data for this chapter 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>. 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/>

References

- Bailey, W. J. (2003). Insect duets: underlying mechanisms and their evolution. *Physiological Entomology*, 28: 157-174.
- Baldassarre G., Nolfi S., Parisi D. (2003). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9: 255-267.
- Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent (with commentary and response). *Adaptive Behavior* 11(4):209-243.

- Cangelosi, A. and Parisi, D. (2002) Computer Simulation: A New Scientific Approach to the Study of Language Evolution. In Angelo Cangelosi and Domenico Parisi, editors, *Simulating the Evolution of Language*, pages 3--28. London: Springer Verlag.
- De Greeff J. & Nolfi S. (in press). Evolution of Implicit and Explicit Communication in a Group of Mobile Robots. In Nolfi S. & Mirolli M. (Eds.), *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.
- Di Paolo E.A. (1997). An investigation into the evolution of communication, *Adaptive Behavior* 6 (2): 285-324.
- Di Paolo E.A. (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 D. E. Kroodsma & E. H. Miller, *Acoustic Communication in Birds* (Eds.), pp. 85-124. New York: Academic Press.
- Floreano, D., Mitri, S., Magnenat, S. and Keller, L. (2007). Evolutionary Conditions for the Emergence of Communication in Robots. *Current Biology*, 17 pp. 514-519.
- Gould, S.J. (1977). *Ontogeny and Phylogeny*. Harvard University Press: Cambridge, MA.
- Haimoff E.H. (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.F. (1960), The origin of Speech, *Scientific American* 203, 88-96.
- 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.
- Marocco D. & Nolfi S. (2007). Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connection Science*, 19 (1): 53-74.
- Mirolli M. & Nolfi S. (in press). Evolving communication in embodied agents: Theory, methods, and evaluation. In S. Nolfi & M. Mirolli (Eds.), *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.
- Mondada F., Bonani M. (2007). The e-puck education robot. <http://www.e-puck.org/>.
- Morris, D (1958). The comparative ethology of grass-finches (*Erythrurae*) and mannikins (*Amadinae*). *Proc. Zool. Soc. London* 131: 389-439
- Nolfi S. & Floreano D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books.
- Nolfi S. & Marocco D. (2001). Evolving robots able to integrate sensory-motor information over time, *Theory in Biosciences*, 120:287-310.
- Nolfi, S. (2005) Emergence of communication in embodied agents: co-adapting communicative and non-communicative behaviours. *Connection Science*, 17(3-4):231-248.
- Nolfi S. (in press). Behavior and cognition as a complex adaptive system: Insights from robotic experiments. In C Hooker (Ed.), *Philosophy of Complex Systems, Handbook on Foundational/Philosophical Issues for Complex Systems in Science*. Elsevier.
- Nolfi S., Mirolli M. (in press). *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.
- Quinn M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosik, P. (Eds.) *Advances in Artificial Life: Sixth European Conference on Artificial Life (ECAL 2001)*. Berlin: Springer Verlag.
- Quinn M., Smith L., Mayley G. & 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.L., Seyfarth R.M. and Owren, M.J. (1999). The meaning and function of grunt variants in baboons. *Animal Behaviour*, 57: 583-592.
- Steels, L. (2003) Evolving Grounded Communication for Robots. *Trends in Cognitive Sciences*, 7 (7):308-312.
- Trianni V., Dorigo M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, volume 95, pages 213-231, 2006.
- Wagner, K., Reggia, J. A., Uriagereka, J., and Wilkinson, G. S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37-69.
- Williams, P.L., Beer, R.D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In S. Bullock, J. Noble, R. Watson and M.A. Bedau (Eds.), *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*. Cambridge, MA: MIT Press.