

Emergence of communication and language in evolving robots

Stefano Nolfi

Institute of Cognitive Sciences and Technologies

Modeling the evolution of communication and language is one of the most fascinating and challenging problems in science. Progressing toward this ambitious goal implies explaining how simple communication forms emerged in the first place and how they evolved into structured communication systems with the characteristics of human language. In this chapter, we will illustrate how communication systems originate and evolve in a population of robots that adapt to a given task/environment. The analysis of these synthetic experiments can help us to understand: (i) how communication can emerge in a population of initially noncommunicating individuals, (ii) what conditions represent a prerequisite for the emergence of a robust and stable communication system, (iii) how the communication system changes by eventually increasing in complexity, and (iv) how signals and meanings originate and how they are grounded in robots' sensorimotor states.

1. Introduction

Understanding the evolution of communication and language is one of the hardest problems in science and poses extremely difficult challenges (Christiansen & Kirby 2003). Indeed, language is one of the most complex human behaviors, the result of dynamic interactions between a population of individuals that in turn dynamically interact with their environment. Furthermore, communication and language continuously change over time by adapting to changes in the environment and/or in individuals.

One of the main difficulties in understanding the evolution of communication and language derives from the need to address highly interdisciplinary issues, such as how animal and human communication systems are structured and how they differ; how communication and language affect and are affected by the behavioral and cognitive capabilities of individuals; how humans acquire language during development; how evolution, learning, and culture interact; etc.

In this chapter, we will approach this issues from an artificial perspective (Cangelosi & Parisi 2002; Kirby 2002; Nolfi 2005; Nolfi & Mirolli 2010a; Steels 2003; Wagner, Reggia, Uriagereka & Wilkinson 2003); that is, we discuss how communication can emerge in populations of evolving robots that interact autonomously (i.e. without human intervention) with the physical world and among themselves.

Studying the evolution of communication by trying to reproduce this evolution in artificial embodied agents (robots) promises to fundamentally enhance our understanding of the evolution of real communication systems for at least three reasons.

The first reason has to do with *reproduction*, that is, with the synthetic approach. As is commonly recognized in an increasing number of scientific fields, computational and robotic models represent a fundamental new instrument for the development and expression of scientific theories that complements the other two classical forms of theorizing: mathematical equations and words. More specifically, this synthetic approach allows one to come up with precise and operational models of how communication originates and how established communication systems can evolve and adapt to variations in the physical and social environment.

The second reason has to do with the particular field of inquiry, that is, evolution, and in particular, the evolution of communication and language. The advantages of expressing theories by producing artifacts are particularly important in such a field, where it is difficult to do empirical experiments (because it is difficult to experiment on evolutionary phenomena), and in which the available empirical evidence is rather scanty and indirect (because communicative behaviors do not leave fossil records from which one might infer the history of the communication system).

The third reason has to do with embodiment, that is, with the fact that the agents of our models have a body and are situated in an environment with which they interact. This is important because it forces the researcher to take into account a number of factors that are likely to play an important role in the evolution of communication but that are typically ignored in non-embodied models, for example the fact that agents have access to information that is incomplete, noisy, and egocentric, or the fact that agents' sensory experiences are influenced by their actions. These aspects are crucial in modeling how "meanings" originate and vary and how communication signals are grounded in agents' nonsymbolic sensorimotor experiences.

More specifically, in this chapter we will review some recent studies of how relatively simple communication systems can emerge from "scratch" in a population of evolving robots. Other research addressing how a communication system with some of the features of human language can emerge in a population of robots playing language games (i.e. ritualized social interactions following a specific script) will be reviewed in Steels' chapter (this volume).

In Section 2, we will discuss whether and how honest communication can evolve despite the conflict of interests between individuals. In Section 3, we discuss whether and how communication can emerge and evolve despite the need to concurrently develop two interdependent skills at the same time (an ability to produce signals encoding useful information and an ability to react to these signals appropriately). Finally, in the fourth section, we will draw our conclusions and discuss the implications of these experiments for the study of the evolution of communication and language.

2. Evolution of a stable communication system

A key question in signaling theory, which has been subject to much debate, concerns the evolution of reliable signals and of a stable communication system when there are conflicts of interest between individuals (due, for example, to competition for limited resources) with particular reference to informative communication forms that are costly to the signaler and beneficial to the receiver.

As soon as a communication system is established (i.e. as soon as individuals produce a signal and react to that signal by modifying their behavior in a way that enhances their adaptive capability), the population will tend to be invaded by mutant deceivers that exploit the detected signal but stop producing signals or alter the signals produced so as to gain a relative advantage over deceived individuals (Dawkins & Krebs 1978).

Indeed, deceptive communication is widespread in nature even among conspecifics. Some animals, for example, produce false alarm calls to scare away conspecifics and gain access to an overcrowded food source (e.g. Møller 1988). Other examples include actively misleading conspecifics while they are searching for food (e.g. Bugnyar & Kotrschal 2004) and bluffing or exaggerating one's own strength or qualities (e.g. Adams & Caldwell 1990). For this reason, deceptive communication often results in a decline in performance with respect to a situation in which individuals do not communicate at all.

The emergence of stable communication can therefore only be explained by postulating other factors that eliminate or counterbalance the effects caused by the conflict of interests between individuals such as genetic relatedness between communicating individuals (Brown & Johnstone 2001; Diggle, Gardner, West & Griffin 2007; Keller & Surette 2006), group relatedness (i.e. the fact that selection might operate at the level of the group rather than at the level of the individuals; Frank 1998; Hamilton 1975), and the costs associated to the production of misleading signals (Grafen 1990; Zahavi 1975).

2.1 Emergence of communication and information suppression in evolving robots with conflicting interests

Recent experiments performed with evolving robots provide a way to test the theories reviewed above and suggest more detailed hypotheses concerning the characteristics that enable the emergence of a stable communication system in individuals with conflicting interests. In the following sections, we will review a series of experiments performed at the EPFL in Lausanne by Sara Mitri, Dario Floreano, Laurent Keller, and Stephane Magnenat (Floreano, Mitri, Magnenat & Keller 2007; Mitri 2009; Mitri, Floreano & Keller 2009).

In these experiments, groups of 10 robots situated in a square arena surrounded by walls evolved the ability to forage for food while avoiding poison sources (i.e. for the ability to reach and remain close to the foraging area while avoiding the poison area, see Figure 1).




Figure 1. Left: The s-bot robot (Mondada et al. 2004). The bottom part in black is the track system. The light plastic ring at the center of the robot represents the LED ring. The cylindrical protuberance on the center of the top is the omnidirectional camera. Right: Representation of a group of 10 robots situated in their environment. The small circles indicate the ten robots (the blue and white colors are used to indicate whether the blue LEDs are turned on or off). The two red circles labeled with the letters “F” and “P” represent the food and poison feeding areas. The light gray and dark gray circular areas around the feeding areas represent the color of the ground around the food and poison areas, respectively.

Adapted from Mitri (2009, Figure 2.1)

2.2 Experimental setup

Both food and poison areas constantly emit red light. The floor around the food and poison sources is covered by a circular strip of light gray and dark gray paper, respectively, which cannot be detected until the robots approach the corresponding area. The robots are provided with a track system that allows them to move in the environment,

an LED ring that allows them to produce or not produce blue light, an omnidirectional camera that allows them to detect at a distance the red light emitted by food or poison sources and the blue light emitted by the other robots, and a ground sensor that allows them to detect the color of the ground under them (which can be used to discriminate whether they are located far from foraging areas, near a food area, or near a poison area).

The robots are provided with an extremely simple brain (Figure 2) consisting of a neural network with two sensory neurons that encode whether the ground sensor detects a light gray or dark gray color, eight sensory neurons that encode the average of red and blue light perceived over the frontal-left, frontal-right, rear-left, and rear-right sectors of the omnidirectional camera, and a bias unit that is always set to 1.0. Sensory neurons are directly connected to the two motor neurons that control the desired speed of the left and right wheels of the track and one motor neuron that controls whether or not the robot emits blue light.




Figure 2. The robots' neural controller. The 11 and 3 circles represent the sensory and motor neurons, respectively. The lines represent the connections between neurons. The figure at the bottom shows how the circular visual field is divided into four sectors.

Adapted from Mitri (2009, Figure 2.2)

The way a robot responds to environmental stimuli depends on the values of the synaptic weights that are encoded in artificial genomes and adapted through evolutionary methods (Nolfi & Floreano 2000). More specifically, the evolutionary process is conducted as follows. At the beginning, an initial population of 500 genotypes, each

formed from 36 genes that encode the value of the 33 connection weights of the robots' neural network, is randomly generated. Then all genotypes are translated into neural controllers that are embodied in robots that are situated in the environment in groups of 10 and evaluated for their ability to reach and remain close to food elements. Once all individuals in the population have been evaluated, the genotypes corresponding to the best individuals are reproduced by generating copies of them with the addition of variations introduced through mutations and recombination. The process is repeated for several generations. For a review of alternative methods in which the genotype does not directly specify the connection strengths but rather characteristics that influence how the connection weights change while the robots interact with the environment, see Nolfi and Floreano (1999).

The robots' performance is calculated by increasing or decreasing a robot's fitness by one point for every time step spent in the food or poison area, respectively. Robots are not directly rewarded for communicating.

Individual robots compete with their conspecifics to reproduce (since only the fittest are allowed to reproduce) and to forage (since a maximum of eight robots can fit around the food source and forage at the same time). This means that there are conflicts of interest between individuals.

To investigate the effects of level of selection and degree of genetic relatedness, the authors ran four series of experiments that varied with respect to two binary conditions: individual versus group selection, and related versus unrelated groups. In the group selection condition, all the individuals belonging to the best groups were selected. In the individual selection condition, the best robots were selected independently of the group to which they belonged. In the related condition, groups were formed by closely related (identical) individuals that were obtained by cloning each reproduced genotype 10 times. In the unrelated condition, groups consisted of unrelated individuals. Each series of experiments was replicated 10 times, starting with populations of robots with different, randomly generated, genotypes.

Finally each of the four experiments, corresponding to the four experimental conditions, was replicated in a normal and in a control condition in which robots were or were not allowed to produce blue light. For more details, see Floreano et al. (2007), Mitri et al. (2009), and Mitri (2009).

2.3 Results

The analysis of the results indicated that, in three experimental conditions (group-selection/related, group-selection/unrelated, individual-selection/related), a stable communication system had evolved. The visual inspection of the behavior exhibited by evolved robots indicated that in all replications the robots moved toward the red light (i.e. the food or poison area) and then remained near the area or moved away

when they detected the light gray or dark gray circular area, respectively. Moreover, in all replications, evolved robots used the possibility to communicate (i.e. to turn their blue LEDs on and off depending on the circumstances and to vary their behavior appropriately when they detected a blue light emitted by another robot). In the majority of cases, the evolutionary process converged on a “signal-food” strategy in which the robots turned their blue LEDs on near the food source and reacted to the detection of blue light by moving toward it. In the remaining cases, the evolutionary process converged on a “signal-poison” strategy in which the robots turned their blue LEDs on near a poison source and reacted to the detection of blue light by avoiding it.

A comparison of performance in the normal and control conditions in which the robots could or could not communicate (were or were not allowed to turn their blue LEDs on) indicates that communication provided an adaptive advantage (i.e. led to better performance). And a comparison of the performance of individuals relying on the signal-food strategy versus the signal-poison strategy indicates that the former was more effective than the latter.

The analysis of the evolutionary process indicates that, in most cases, it led to a stable communication system (i.e. the ability to communicate, the situation in which signals are produced, and the general way in which the detection of a signal affects the robots’ behavior were preserved over successive generations). In a few cases, a communication system based on the signal-poison strategy might be later replaced by a communication system based on the signal-food strategy – which, however, then remained stable.

In contrast, in the individual-selection/unrelated condition, the evolutionary process converged on a form of deceptive communication in which the robots tended to emit blue lights when they were far from the food sources and to react to blue light by approaching rather than avoiding it. Our comparison of performance in the normal and control conditions in which the robots were or were not allowed to communicate (i.e. were or were not allowed to turn their blue LEDs on) indicates that the possibility of communicating led to lower performance with respect to the control condition in which the robots were not allowed to communicate.

The analysis of the evolutionary process indicates that this suboptimal behavior can be explained by the evolutionary process tending to converge on a limit cycle dynamics (see Figure 3). More specifically:

1. During the first phase of this process, the robots develop the same individual skills that we described above: an ability to move toward red light and remain near the food source (when they detect the light gray ground color) and avoid the poison source (when they detect the dark gray ground color).
2. Then, in phase 2 the robots start to develop a communication skill that consists in producing the blue light near food, and

3. In reacting to the detection of blue light by approaching it.
4. In phase 3, however, variations that lead to a deceptive form of communication in which individuals preferentially turn on their blue lights far from the food become adaptively advantageous and tend to invade the population.
5. The proliferation of individuals that deceive their conspecifics creates the adaptive conditions for retaining mutant individuals with a reduced tendency to approach blue light. This leads to phase 5, in which the robots are attracted to blue light only weakly or not at all. However, the absent or reduced tendency to approach blue light recreates a condition similar to the one that triggered phase 2, thus leading to a limit cycle dynamics in which robots tend to keep changing their behavior according to phases 2, 3, 4, and 5, over and over again.

The fact that the evolutionary process is constrained by this limit cycle dynamics prevents convergence on a truly stable communication system and instead causes convergence on a semistable equilibrium in which, on average, the robots preferentially tend to turn their blue LEDs on far from food sources and react weakly to the detection of blue light; however, this state is characterized by high intra- and intergenerational variability.

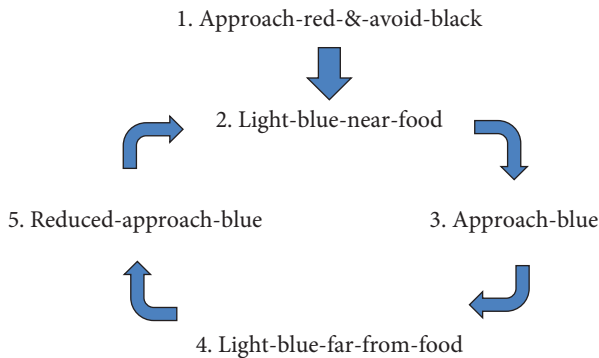


Figure 3. A schematization of the limit cycle dynamics that characterize the evolutionary process.

Adapted from Mitri (2009, Figure 6.1)

2.4 Discussion

As the authors pointed out, the results of these studies confirmed the crucial role of genetic relatedness and/or group selection for the evolution of honest communication. More specifically, they confirmed the hypothesis (Diggle et al. 2007; Keller & Surette 2006) that individuals sharing identical genes would evolve an honest communication system since this allows them to achieve higher inclusive fitness, thus making them more likely to be selected for subsequent generations (Floreano et al. 2007). Moreover,

they confirmed the hypothesis that evolution would lead to the emergence of an honest communication system when selection operates at the group level (Frank 1998), since the likelihood of survival depends not only on one's own performance but also on the performance of other group members (Floreano et al. 2007). They also allowed us to make more precise quantitative predictions about the relative roles of the different factors (see Floreano et al. 2007).

Moreover, as reported by the authors (Mitri et al. 2009), these experiments demonstrate that, when conflicts of interest prevent the evolution of honest communication, deceptive communication (Dawkins & Krebs 1978) or information suppression (in which the inadvertently provided information is eliminated; Hauser 1992) are not the only possible outcomes. Indeed, as we have seen, the evolutionary process can lead to limit cycle dynamics, resulting in a semistable suboptimal equilibrium. Alternatively, as observed in the artificial experiments reported in Mirolli and Parisi (2008), it might lead to a limit cycle dynamics in which communication is periodically discovered or rediscovered and then lost again.

Finally, as reported by the authors, another interesting and surprising result is that a complete suppression of information never occurred. Instead, the system stabilized at an equilibrium, characterized by high variation in both signaling and response strategies, where little information was transmitted and response to the information was weak (Mitri et al. 2009). Such a suboptimal equilibrium results from the complex interplay between variation introduced through mutations and weak selection pressure against the suppression of inadvertent signaling and against the reaction to deceptive signals. As the authors hypothesized, this may provide a possible explanation for the long-standing question of high intrapopulation polymorphism observed in signaling systems where the signalers' and receivers' interests conflict (Bond 2007; Bond & Kamil 2002; Dewitt, Sih & Hucko 1999; Poulton 1890; Whiteley, Owen & Smith 1997).

3. How forms of communication originate and evolve

The emergence of communication represents an evolutionary paradox, since it is unclear why an individual should evolve the ability to signal if no receivers had evolved a response to the signal. Conversely, if no signal exists, there is nothing for individuals to evolve a response to (Maynard Smith 1997). A first open question thus is how does communication emerge even though it requires the concurrent development of two interdependent skills that are adaptively neutral in isolation? A second related question is how do established communication systems evolve (i.e. how does evolution mold established communicative forms) by eventually complexifying?

One possible answer to these questions is that communicative forms might originate from inadvertent cues. Or more specifically, as we will demonstrate below, they

might emerge from the exploitation of cues produced as a side effect of behaviors that had (and possibly still have) independent noncommunicative functions. These cues create the adaptive conditions for the development of appropriate responses, thus leading to the establishment of a communicative form.

This is indeed the mechanism that explains how communication originates in the experiments reviewed in Section 2. In fact, the robots' individual foraging behavior (which consists in moving toward red targets and then remaining close to the food area while abandoning the poison area) produces, as a side effect, a concentration of robots producing blue light near the food area even if robots turned the blue light on in a quasirandom way (Mitri 2009). The development of such individual behavior thus confers an information value on the blue light that creates the adaptive conditions for the development of an ability to react to the blue light by approaching it. Finally, the establishment of such communicative behavior creates the conditions for improving the quality of the signal through the development of an ability to turn on the blue light only near food.

In this respect, it is important to notice that the inadvertent production of cues might arise as a result of qualitatively different processes. In some cases, as in the case we just mentioned, the cues are the inevitable consequence of behaviors that have a function for the individual performing them. In other cases, cues are constituted by behaviors that do not represent inevitable consequences of the execution of functional behaviors but that nevertheless tend to be produced spontaneously and to correlate with the production of functional behaviors. The former case is well documented in animal communication studies. In many species, individuals have been shown to monitor each other to decide how to behave (Bshary & Grutter 2006; Dall, Giraldeau, Olsson, McNamara & Stephens 2005; Franks 1999; Giraldeau 1997; Shuker & West 2004; Tibbetts & Dale 2004; West, Herre & Sheldon 2000). For example, during foraging activities, individuals observe their conspecifics' behavior to gain indirect information about food locations (Buckley 1997; Galef & Giraldeau 2001). An interesting piece of evidence supporting the latter case comes from an experiment reported in Hihara, Yamada, Iriki, and Okanoya (2002) in which monkeys were trained to use a rake-shaped tool to retrieve distant food placed on a table. The monkeys began to spontaneously vocalize during the tool training and then modified their vocalization behavior once they noticed that it tended to trigger a specific effect on the experimenter (the experimenter passed the tool to the monkey after hearing a specific coo-call). This led to the establishment of a communicative form between the monkey and the experimenter in which the monkey reliably produced a specific coo-call to request the tool and the experimenter responded by passing the tool (Hihara et al. 2002). Additional evidence demonstrating how communication behaviors can originate from spontaneous behaviors correlated with the execution of functional

behaviors comes from artificial experiments reported by Cangelosi and Parisi (1998) and Mirolli and Parisi (2008).

A second factor that might explain the origin of communication is the so-called “receiver biases hypothesis” (Endler & Basolo 1998; Maynard Smith & Harper 2003), which hypothesizes that communication behaviors might originate from the exploitation of the sensory biases of individuals, that is, from the tendency of individuals to react to certain sensory states by exhibiting specific behaviors. The presence of these biases creates the adaptive conditions for the development of an ability to produce signals, similar to the environmental states that trigger specific behavioral responses, thereby triggering the execution of those behaviors by other individuals when appropriate (Guilford & Dawkins 1991; Maynard-Smith & Harper 2003).

In the following section, we will review another experiment that provides some suggestive evidence about how communication emerges (i.e. how signals and their associated meaning originate) and how the established communication system complexifies to enable additional functions.

3.1 Origins and evolution of communication in cooperating robots

In a series of experiments carried out at the Italian National Research Council by de Greeff and Nolfi (2010), pairs of robots situated in a square arena surrounded by walls evolved the ability to perform a coordinated behavior. This behavior consisted in reaching the black and white circular areas and then switching areas as quickly as possible (Figure 4).




Figure 4. Left: The environment and the robots. The two circular areas of the environment colored in black and white represent the target areas. Right: The e-puck robotic platform (Mondada & Bonani 2007) including the ground sensory board and a strip of red paper around the top of the body

3.2 Experimental setup

The robots were 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 two motors that control the two corresponding wheels, eight infrared proximity sensors uniformly distributed around the robot's body, one infrared sensor 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. Each robot's body is covered with a circular strip 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] and [1.0] that are transmitted and received through the Bluetooth connection. At each time step, both robots emit a signal and detect the signal produced by the other robot.

Each robot's neural controller (Figure 5) has 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. The sensory layer consists of eight neurons that encode the state of the eight corresponding infrared sensors; three neurons that encode whether the robot detects another robot in the frontal-left, frontal, or frontal right side of its field of view; two neurons that encode in binary whether the ground sensor of the robot detects a white or black target area; two neurons that encode the previous state of the ground sensors; and two signal sensors that encode the signal received from the other robot and the signal produced by the robot itself. The motor layer includes two neurons that encode the desired speed of the two corresponding wheels and one neuron that encodes the value of the signal produced by the robot.




Figure 5. The architecture of the robots' neural controller

To favor the emergence of cooperation, the group is composed of genetically identical individuals (i.e. clones) and selection operates at the level of the group. More specifically, the group scores one point every time the two robots are concurrently located in the two areas (one in each area) for the first time or after a switch. As in the previous experiment, the robots are not rewarded for communicating. The way the robots behave, the signals they produce, if any, and how they react to signals and other sensory states depends on the connection weights that are encoded in genotypes and evolved. The initial population consists of 100 randomly generated genotypes that encode the connection weights, biases, and time constants of 100 corresponding neural controllers. Each genotype is translated into two identical neural controllers that are embodied in two robots. The 20 best genotypes in each generation are allowed to reproduce by generating five copies each, with 2% of their bits replaced with a new randomly selected value. For more details, see de Greeff and Nolfi (2010).

3.3 Results

The analysis of the results obtained in different replications of the experiment indicates that the robots solved the problem by exploiting the possibility to communicate. Visual inspection of the best solutions indicates that they can be grouped into two qualitatively different classes of strategies. In both types, the robots initially displayed an exploration behavior that allowed them to enter the two target areas (one robot per area) and then displayed a series of target switching behaviors in which each robot exploited the information provided by the other individual to navigate directly toward the other target area. The first strategy involved a synchronized target-switching behavior in which the two robots, located in the two different target areas, simultaneously left their current area and moved directly toward the other target area. The second strategy type was characterized by a switching behavior organized in two phases in which first a robot exited its target area and traveled toward the other target area containing the other robot and then the latter robot exited its target area and traveled directly toward the target area previously occupied by the former robot (videos with examples of the behavior can be seen at <http://laral.istc.cnr.it/esm/evo-communication/>).

In the rest of this section, we summarize how the robots' behavior varied across generations in a replication of the experiment that converged on the two-phase switching behavior strategy. We will summarize the main changes occurring during the seven successive phases, which extended over about 1,800 generations (Figure 6), and finally led to the near-optimal solution described above and below.

Phase 1. During the very first generations, the robots developed two simple individual behaviors that consisted in moving forward when the infrared sensors were off, and avoiding obstacles when the infrared sensors were on. The combination of these

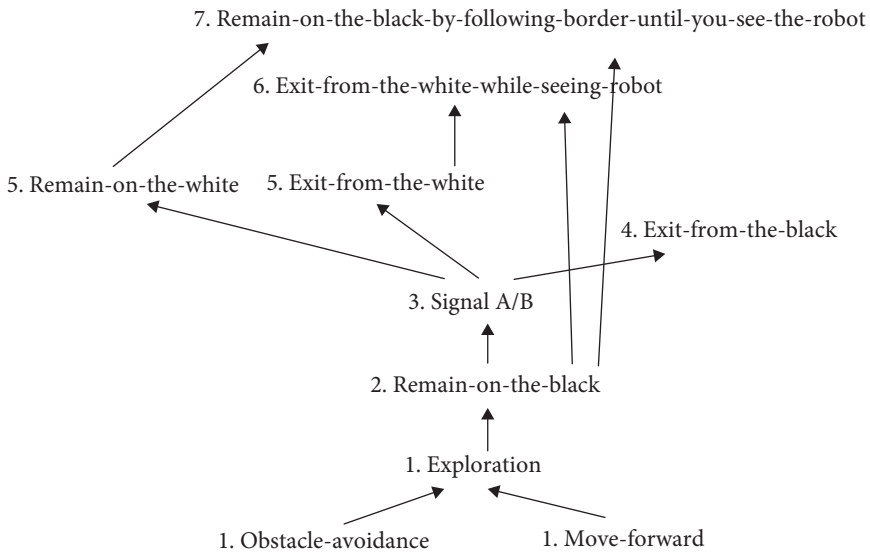


Figure 6. Schematization of the most important changes occurring during the evolutionary process. Labels indicate behavioral and/or communication skills. The number preceding the label indicates the phase in which the skill was developed. The arrows indicate (some of) the dependencies, namely the pre-existing behaviors that created the adaptive conditions for the development of a new behavior

two behaviors allowed the robots to avoid crashing into obstacles and to explore the environment, thereby occasionally scoring a point when the two robots happened to cross over the two areas at the same time.

Phase 2. The robots developed a new remain-on-the-black behavior that consisted in turning when the ground sensors detected a black area so as to remain inside it. This new behavior improved the robots' performance since, as soon as the other robot happened to cross the white target area, both robots were rewarded. At this stage of the evolutionary process, the behavior of the robots was not affected significantly by the other individuals. Thus, each robot behaved individually.

Phase 3. The remain-on-black-area behavior not only resulted in performance improvement but also allowed the robots to access information that had a communicative value (the knowledge of being in the black area and the knowledge of the position of the area). Moreover, as a side effect, it caused the robots to start to produce two differentiated signals (which we will call signals A and B) outside and inside the black area, respectively. This is an example of the production of a cue associated with the execution of a functional behavior (i.e. remain-on-the-black) that does not constitute an inevitable consequence of the execution of that behavior. Nevertheless, a cue correlating with the execution of that behavior tends to be spontaneously produced. This can be explained by considering that the robots' internal neurons tend to assume

different states depending on whether or not they are executing the remain-on-the-black behavior and that the internal neurons affect either the motor neurons controlling wheel speed or the motor neuron controlling signal production. This effect should be expected in any system in which the neural area controlling the actuators responsible for the production of a specific functional behavior is not fully separated from neural areas controlling other actuators (e.g. the actuators controlling facial expression and/or vocalization).

Phase 4. The availability of a cue that provides information on whether or not the other individual is located on the black area created the adaptive conditions for the development of an ability to exit-from-the-black-area when another robot was inside the same area – a new capacity that led to improved performance. This means that the cues originally produced as a side effect of the execution of the remain-on-the-black-area behavior later became a signal (signal B) that elicited a functionally appropriate response in the robots detecting it.

Phase 5. Signal B also created the adaptive conditions for the development of an ability to remain-on-the-white-area, when the other robot was not in the black area, and to exit-from-the-white-area, when the other robot was on the black area.

In sum, these skills allowed the group of robots to explore the environment, find the areas, and remain on their area until the two robots were concurrently located in the two areas. Moreover, they allowed the robots to switch areas through a two-phase strategy enabled by the fact that the robot located in the white area exited the area as soon as the other robot was in the black area, and the robot located in the black area exited the area as soon as the other robot entered its area. At this stage, however, the robots relied on a time-consuming exploration strategy to switch areas since they were still not able to travel directly toward the opposite area. On the other hand, these established behavioral and communication skills created the conditions for the emergence of new skills that would allow the robots to further improve their performance.

Phase 6. After about 150 generations during which performance remained rather stable, the robots modified their exit-from-the-white-area behavior by exiting the area not only when they detected signal B (i.e. when the other robot was on the black area) but also when they visually detected the other robot (i.e. when the robot in the black area was located in front of them). This new behavior was triggered by two cues that were originally produced as a side effect of the development of the remain-on-black-area behavior. The first cue is signal B, whose origins have been discussed above. The second cue is the position of the other robot during the execution of the remain-on-the-black-area behavior, which implicitly and inevitably indicates to the robot the relative position of the black area.

Phase 7. Once the position of the robot during the execution of the remain-on-the-black-area behavior became a signal triggering a functionally appropriate response (as described above), the way the robots displayed the remain-on-the-black-area behavior became subject to adaptive pressures. The positions that the robot assumed

within the black area codetermined whether or not the robot coming from the white area would successfully reach the black area. This favored the retention of a series of variations that finally led to the development of a smart solution that consisted in remaining on the black area by moving counterclockwise along the border of the area and stopping when the other robot was visually detected. The execution of this follow-the-border-until-you-see-the-robot behavior allowed the robot located in the black area to assume a precise relative position with respect to the center of the area and with respect to the other robot. This, in turn, allowed the robot exiting the white area to orient precisely toward the center of the black area, thus avoiding the risk of missing it (Figure 7). Moreover, the development of this behavior allowed the robot located in the black area to remain oriented toward the other robot (i.e. toward the white area) so that, when it left the black area, it could navigate directly toward the white area (Figure 7).




Figure 7. Schematization of the target-switching behavior produced at the end of the evolutionary process. The dashed black and white lines represent the trajectory of the robots located in the black and white areas, respectively. The white circle represents the position in which the robot located in the black area stops while looking toward the robot located in the white area. The dashed black line indicates, first, the phase in which the robot remains in the white area by moving counterclockwise along the border, and, second, the phase in which it exits from the area (when it detects signal B and sees the other robot). The dotted white line indicates the trajectory produced by the robot that exits from the black area as soon as the other robot also enters the black area. Notice how the relative positions and orientations that the two robots assume before exiting from their areas allow them to travel toward the center of their destination areas despite their sensory system's inability to discriminate between different parts of the areas

Overall, these behavioral and communication skills allowed the robots to explore the environment, reach the two target areas, and reliably switch areas several times by navigating directly toward the appropriate destination area.

3.4 Discussion

The analysis of the process whereby signals and meanings originate and evolve demonstrates a strict interdependence between the evolution of behavioral and of communication skills (Nolfi 2005). In particular, the development of functional behavioral

skills tends, as a side effect, to produce cues that have a communicative value. Such cues might be the inevitable results of the execution of the corresponding behavior (e.g. they might consist of the position of the robot during the execution of a behavior that requires it to assume that position) or they might consist of spontaneously produced cues that correlate with the execution of functional behavior (e.g. they might be constituted by originally meaningless vocalizations that vary during the production of different functional behaviors). The availability of these cues creates the adaptive conditions for the development of adaptive responses that transform them into signals encoding a given meaning and triggering a functionally appropriate response in the individuals that perceive them. As soon as cues become signals, either the traits underlying their production and the traits underlying the response they trigger become subject to adaptive pressures. The adaptive pressure on the communicative form established in this way typically leads to a transformation of the form and of the meaning of the signals that enhances the functionality of the communicative form. This kind of transformation is achieved through behavioral modifications and, more specifically, through modifications of the behavior enabling signal production and of the behavioral responses to detected signals.

The analysis of the experiments reported on here also demonstrates how the development of action skills (e.g. remaining in the black area) creates the adaptive conditions for the development of communication skills (e.g. producing signals A and B) and how the development of these communication skills (e.g. signals A and B) then creates the adaptive conditions for the development of action skills (e.g. exiting-from-the-black-area). This process can potentially lead to long-sustained periods of innovation in which new capacities create the adaptive conditions for development of further new capacities. Moreover, this process can lead incrementally to the progressive complexification of action and communication skills, in which new capacities tend to be developed on top of established capacities and old established capacities tends to assume additional functionalities and thus to be retained.

Overall, this implies that signal and meaning are not grounded (Harnad 1990) directly in sensorimotor states but rather in behaviors that are ultimately grounded in sensorimotor states.

4. Conclusion

In this paper, we reviewed a series of experiments that demonstrate how communication can emerge in an evolving population of initially noncommunicating robots. We studied the conditions under which a stable communication system can evolve. We showed which mechanisms overcome the problems of developing two interdependent abilities, and how established communication systems can emerge. These experiments

provide a way to test theories on the origin of communication and provide a more detailed account of how communication emerges and evolves. In some cases, they also suggest new hypotheses. For example, the experiments reviewed in Section 2 indicate that conflicts of interests between individuals do not necessarily lead to deceptive communication or information suppression, as usually hypothesized, but might also lead to limit cycle dynamics, resulting in a semistable suboptimal equilibrium state or in an unstable state in which communication is periodically rediscovered and lost. The experiments in Section 3 indicate that the development of new behavioral skills often creates the conditions for the development of communication skills, which in turn might create the conditions for the development of new behavioral skills, thus leading to a progressive expansion of the individuals' behavioral and communicative repertoire.

Whether these models can be extended to the study of the evolution of communication systems with the characteristics of human language still represents an open issue (for a discussion, see Mirolli & Nolfi 2010; Nolfi & Mirolli 2010b; Parisi 2010). For an example of how this methodology can be used to study the evolution of displaced signals (Hockett 1960), namely signals providing information about regularities that are extracted from previous sensorimotor experiences rather than from the current context experienced by the agent emitting the signal, see Williams, Beer, and Gasser (2008).

Experiments re-enacting the evolution of communication in artificial agents can help us to identify universal properties that might be at the basis of the evolution of all types of communication systems, including human language. One of these fundamental properties might be the strict interdependence between action and communicative behavior. The hypothesis that there might be such a strict interdependence in language evolution is also supported by the abundant neural and psychological evidence of integration between action and language in humans (Buccino et al. 2005; Cappa & Perani 2003; Gallese 2008; Glenberg & Kaschak 2002; Pulvermuller 2003; Rizzolatti & Arbib 1998). Moreover, it is supported by recent neurorobotic studies showing how the codevelopment of language and behavioral skills can lead to a form of compositionality that allows the robot to appropriately interpret new sentences by recombining the meaning of their constituent words (Arie et al. 2010; Sugita & Tani 2005, 2008; Tuci, Ferrauto, Massera & Nolfi 2010a, 2010b).

Finally, the analysis of how primitive forms of communication evolve by increasing in complexity might lead to the identification of new factors that promote language evolution. The results reported in Section 3, which reveals a cocausal effect between the development of action and communication skills due to an expansion of the adaptive needs and the possibility of developing new skills by reusing previously acquired skills, suggest that an expansion of the individuals' behavioral skills might have been one of the main factors that triggered the origins of language. The expansion of individuals'

behavioral repertoires might have led to a parallel expansion of their communication system as well as of their cognitive and social skills (e.g. selective attention, memory, sequential control, social learning, and imitation) that might in turn have created an adaptive need for more sophisticated communication systems (e.g. “protolanguages” with compositional features) whose development might have been favored by the opportunities for skill reuse as well.

References

- Adams, E.S., & Caldwell, R.L. (1990). Deceptive communication in asymmetric fights of the stomatopod crustacean *Gonodactylus bredini*. *Animal Behaviour*, 39, 706–716.
- Arie H., Endo T., Jeong S., Lee M., Sugano S., & Tani J. (2010). Integrative learning between language and action: A neuro-robotics experiment. In *Proceedings of the 20th International Conference on Neural Networks*. Berlin: Springer.
- Bond, A.B. (2007). The evolution of color polymorphism: Crypticity, searching images, and apostatic selection. *Annual Review of Ecology, Evolution, and Systematics*, 38, 489–514.
- Bond, A.B., & Kamil, A.C. (2002). Visual predators select for crypticity and polymorphism in virtual prey. *Nature*, 415, 609–613.
- Brown, S.P., & Johnstone, R.A. (2001). Cooperation in the dark: Signalling and collective action in quorum-sensing bacteria. *Proceedings of the Royal Society B: Biological Sciences*, 268, 961–965.
- Bshary, R., & Grutter, A.S. (2006). Image scoring and cooperation in a cleaner fish mutualism. *Nature*, 441, 975–978.
- Buccino, G., Riggio, T., Mellia, G., Binkofski, F., Gallese, V., & Rizzolatti, G. (2005). Listening to action-related sentences modulates the activity of the motor system: A combined TMS and behavioral study. *Cognitive Brain Research*, 24, 355–363.
- Buckley, N.J. (1997). Spatial concentration effects and the importance of local enhancement in the evolution of colonial breeding in seabirds. *The American Naturalist*, 149, 1091–1112.
- Bugnyar, T., & Kotrschal, K. (2004). Leading a conspecific away from food in ravens (*Corvus corax*)? *Animal Cognition*, 7, 69–76.
- Cangelosi, A., & Parisi, D. (1998). The emergence of a “language” in an evolving population of neural networks. *Connection Science*, 10, 83–97.
- Cangelosi, A., & Parisi, D. (2002). Computer simulation: A new scientific approach to the study of language evolution. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 3–28). London: Springer.
- Cappa, S., & Perani, D. (2003). The neural correlates of noun and verb processing. *Journal of Neurolinguistics*, 16, 183–189.
- Christiansen, M.H., & Kirby, S. (2003). Language evolution: Consensus and controversies. *Trends in Cognitive Sciences*, 7, 300–307.
- Dall, S.R. X., Giraldeau, L.A., Olsson, O., McNamara, J.M., & Stephens, D.W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology and Evolution*, 20, 187–193.
- Dawkins, R., & Krebs, J.R. (1978). *Animal signals: Information or manipulation?* Oxford: Blackwell.

- De Greeff, J., & Nolfi, S. (2010). Evolution of implicit and explicit communication in a group of mobile robots. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 179–214). Berlin: Springer.
- Dewitt, T.J., Sih, A., & Hucks, J.A. (1999). Trait compensation and cospecialization in a freshwater snail: Size, shape and antipredator behaviour. *Animal Behaviour*, 58, 397–407.
- Diggle, S.P., Gardner, A., West, S.A., & Griffin, A.S. (2007). Evolutionary theory of bacterial quorum sensing: When is a signal not a signal? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 1241–1249.
- Endler, J.A., & Basolo, A.L. (1998). Sensory ecology, receiver biases and sexual selection. *Trends in Ecology and Evolution*, 13, 415–420.
- Floreano, D., Mitri, S., Magnenat, S., & Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17, 514–519.
- Frank, S.A. (1998). *Foundations of social evolution*. Princeton, NJ: Princeton University Press.
- Franks, N.R. (1999). Information flow in the social domain: How individuals decide what to do next. In C. Detrain, J.L. Deneubourg, & J. M. Pasteels (Eds.), *Information processing in social insects* (pp. 101–112). Basel: Birkhäuser Verlag.
- Galef, B.G., & Giraldeau, L.A. (2001). Social influences on foraging in vertebrates: Causal mechanisms and adaptive functions. *Animal Behaviour*, 61, 3–15.
- Gallese, V. (2008). Mirror neurons and the social nature of language: The neural exploitation hypothesis. *Social Neuroscience*, 3, 317–333.
- Giraldeau, L.A. (1997). The ecology of information use. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural ecology: An evolutionary approach* (pp. 42–68). Oxford: Blackwell.
- Glenberg, A., & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review*, 9, 558–565.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144, 517–546.
- Guilford, T., & Dawkins, M.S. (1991). Receiver psychology and the evolution of animal signals. *Animal Behavior*, 42, 1–14.
- Hamilton, W.D. (1975). *Innate social aptitudes of man: An approach from evolutionary genetics*. London: Malaby Press.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Hauser, M.D. (1992). Costs of deception: Cheaters are punished in rhesus monkeys (*Macaca mulatta*). *Proceedings of the National Academy of Sciences of the United States of America*, 89, 12137–12139.
- Hihara, S., Yamada, H., Iriki, A., & Okanoya, K. (2003). Spontaneous vocal differentiation of coo-calls for tools and food in Japanese monkeys. *Neuroscience Research*, 45, 383–389.
- Hockett, C.F. (1960). The origin of speech. *Scientific American*, 203, 88–96.
- Keller, L., & Surette, M.G. (2006). Communication in bacteria: An ecological and evolutionary perspective. *Nature Reviews Microbiology*, 4, 249–258.
- Kirby, S. (2002). Natural language from artificial life. *Artificial Life*, 8, 185–215.
- Maynard Smith, J. (1997). *The theory of evolution* (3rd ed.). Cambridge, UK: Cambridge University Press.
- Maynard Smith, J., & Harper, D. (2003). *Animal signals*. Oxford: Oxford University Press.
- Mirolli, M., & Nolfi, S. (2010). Evolving communication in embodied agents: Theory, methods, and evaluation. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 105–121). Berlin: Springer.
- Mirolli, M., & Parisi, D. (2008). How producer biases can favor the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16, 27–52.

- Mitri, S. (2009). *The evolution of communication in robot societies*. Ph.D. thesis, EPFL, Lausanne.
- Mitri, S., Floreano, D., & Keller, L. (2009). The evolution of information suppression in communicating robots with conflicting interests. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 15786–15790.
- Mondada, F., & Bonani, M. (2007). The e-puck education robot. (<http://www.e-puck.org/>)
- Mondada, F., Pettinaro, G.C., Guignard, A., Kwee, I.V., Floreano, D., Deneubourg, J.-L., et al. (2004). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17, 193–221.
- Møller, A. (1988). False alarm calls as a means of resource usurpation in the great tit *Parus major*. *Ethology*, 79, 25–30.
- Nolfi, S. (2005). Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. *Connection Science*, 17, 231–248.
- Nolfi, S., & Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, 7, 89–113.
- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: The MIT Press/Bradford Books.
- Nolfi, S., & Mirolli, M. (Eds.). (2010a). *Evolution of communication and language in embodied agents*. Berlin: Springer.
- Nolfi, S., & Mirolli, M. (2010b). Evolving communication in embodied agents: Assessment and open challenges. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 215–220). Berlin: Springer.
- Parisi, D. (2010). Artificial organisms with human language. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 13–35). Berlin: Springer.
- Poulton, E.B. (1890). *The colours of animals: Their meaning and use, especially considered in the case of insects*. New York, NY: Appleton.
- Pulvermuller, F. (2003). *The neuroscience of language. On brain circuits of words and serial order*. Cambridge, UK: Cambridge University Press.
- Rizzolatti, G., & Arbib, M. (1998). Language within our grasp. *Trends in Neuroscience*, 21, 188–194.
- Shuker, D.M., & West, S.A. (2004). Information constraints and the precision of adaptation: Sex ratio manipulation in wasps. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 10363–10367.
- Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7, 308–312.
- Sugita, Y., & Tani, J. (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior*, 13, 33–52.
- Sugita, Y., & Tani, J. (2008). Acquiring a functionally compositional system of goal-directed actions of a simulated agent. In M. Asada, J. Hallam, J.-A. Meyer, & J. Tani (Eds.), *Proceedings of the 10th International Conference on Simulation of Adaptive Behavior (SAB2008)* (pp. 331–341). Berlin: Springer.
- Tibbetts, E.A., & Dale, J. (2004). A socially enforced signal of quality in a paper wasp. *Nature*, 432, 218–222.
- Tuci, E., Ferrauto, T., Massera, G., & Nolfi, S. (2010a). The evolution of behavioural and linguistic skills to execute and generate two-word instructions in agents controlled by dynamical neural networks. In H. Fellersmann et al. (Eds.), *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems (ALife XII)*. Cambridge, MA: The MIT Press.

- Tuci, E., Ferrauto, T., Massera, G., & Nolfi, S. (2010b). Co-development of linguistic and behavioural skills: compositional semantics and behaviour generalisation. In S. Doncieux et al. (Eds.), *Proceedings of the 11th International Conference on Simulation of Adaptive Behavior (SAB2010)* (pp. 523–532). Heidelberg: Springer.
- Wagner, K., Reggia, J.A., Uriagereka, J., & Wilkinson, G.S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11, 37–69.
- West, S.A., Herre, E.A., & Sheldon, B.C. (2000). The benefits of allocating sex. *Science*, 290, 288–290.
- Whiteley, D.A. A., Owen, D.F., & Smith, D.A.S. (1997). Massive polymorphism and natural selection in *Donacilla cornea* (Poli, 1791) (Bivalvia: Mesodesmatidae). *Biological Journal of the Linnean Society*, 62, 475–494.
- Williams, P., Beer, R., & Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In S. Bullock, J. Noble, R. Watson, & M. A. Bedau (Eds.), *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems* (pp. 702–709). Cambridge, MA: The MIT Press.
- Zahavi, A. (1975). Mate selection – A selection for a handicap. *Journal of Theoretical Biology*, 53, 205–214.