

Chapter 10

Evolution of Signaling in a Multi-Robot System: Categorization and Communication

Christos Ampatzis, Elio Tuci, Vito Trianni,
and Marco Dorigo

Abstract We use Evolutionary Robotics to design robot controllers in which decision-making mechanisms to switch from solitary to social behavior are integrated with the mechanisms that underpin the sensory-motor repertoire of the robots. In particular, we study the evolution of behavioral and communicative skills in a categorization task. The individual decision-making structures are based on the integration over time of sensory information. The mechanisms for switching from solitary to social behavior and the ways in which the robots can affect each other's behavior are not predetermined by the experimenter, but are aspects of our model designed by artificial evolution. Our results show that evolved robots manage to cooperate and collectively discriminate between different environments by developing a simple communication protocol based on sound signaling. Communication emerges in the absence of explicit selective pressure coded in the fitness function. The evolution of communication is neither trivial nor obvious; for a meaningful signaling system to evolve, evolution must produce both appropriate signals and appropriate reactions to signals. The use of communication proves to be adaptive for the group, even if, in principle, non-cooperating robots can be equally successful with cooperating robots.

C. Ampatzis (✉)

Advanced Concepts Team, European Space Agency, Noordwijk, The Netherlands
e-mail: christos.ampatzis@esa.int

E. Tuci · V. Trianni

Institute of Cognitive Sciences and Technologies, CNR, Rome, Italy

E. Tuci

e-mail: elio.tuci@istc.cnr.it

V. Trianni

e-mail: vito.trianni@istc.cnr.it

M. Dorigo

IRIDIA-CoDE, ULB, Brussels, Belgium

e-mail: mdorigo@ulb.ac.be

1 Introduction

The work presented in this chapter is about the design of robot controllers in which decision-making mechanisms to switch from solitary to social behavior are integrated with the mechanisms that underpin the sensory-motor repertoire of the robots. In particular, we study the evolution of behavioral and communicative skills in a categorization task. Communication is the way in which the collective group response can be triggered, once one or more robots within the group take a decision. The individual decision-making structures are based on the integration over time of sensory information. The tool we use to implement such structures is the Continuous Time Recurrent Neural Network (hereafter CTRNN; see Beer and Gallagher 1992, for details). These structures should allow robots to initiate social behavior in response to the persistence of certain environmental stimuli.

As a design methodology, we use Evolutionary Robotics (ER; see Nolfi and Floreano 2000, and Chap. 7). The mechanisms for switching from solitary to social behavior and the ways in which the robots can affect each other's behavior (i.e., communication) are not predetermined by the experimenter, but are aspects of our model designed by artificial evolution. This approach is particularly suitable for our goal because it permits the co-evolution of communicative and non-communicative behavior; different strategies can co-adapt because selection depends only on an overall evaluation of the group (see Nolfi 2005).

Our results show that evolved robots manage to cooperate and collectively discriminate between different environments by developing a simple communication protocol based on sound signaling. Communication emerges in the absence of explicit selective pressure coded in the fitness function. The evolution of communication is neither trivial nor obvious; for a meaningful signaling system to evolve, evolution must produce both appropriate signals and appropriate reactions to signals. The use of communication proves to be adaptive for the group, even if, in principle, non-cooperating robots can be equally successful with cooperating robots.

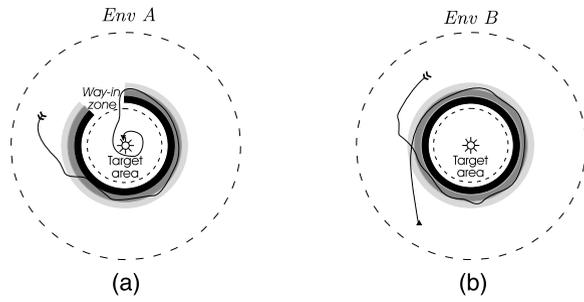
In Sect. 2, we introduce the task, the simulation model, the controller and the evolutionary algorithm and the fitness function employed to evolve the desired behavior. In Sect. 3, we present the results of the experiments we conducted and we discuss the adaptive significance of signaling. Finally, in Sect. 4 we draw conclusions.

2 Methods

2.1 Description of the Task

The task we consider is a categorization task in which two robots are required to discriminate between two different environments using temporal cues, that is, by integrating their perceptual inputs over time. At the beginning of each trial, two simulated robots are placed in a circular arena with a radius of 120 cm (see Fig. 1),

Fig. 1 The task. (a) *Env A* is characterized by the way in zone. The target area is indicated by the dashed circle. (b) In *Env B* the target area cannot be reached. The continuous arrows are an example of a good navigation strategy for one robot



at the center of which a light bulb is always turned on. The robots are positioned randomly at a distance between 75 and 95 cm from the light, with a random orientation between -120° and $+120^\circ$ with respect to the light. The robots perceive the light through their ambient light sensors. The color of the arena floor is white except for a circular band, which is centered around the lamp and covers an area between 40 and 60 cm from it. The band is divided in three sub-zones of equal width but colored differently: light gray, dark gray, and black. Each robot perceives the color of the floor through its floor sensors, positioned under its chassis. Robots are not allowed to cross the black edge of the band close to the light. This black edge can be seen as a circular trough that prevents the robots from reaching the light. The colored zones can be seen as an indication of how close the robots are to the “danger.” There are two types of environment. In one type—referred to as *Env A*—the band has a gap, called the *way in zone*, where the floor is white (see Fig. 1(a)). In the other type, referred to as *Env B*, the band completely surrounds the light (see Fig. 1(b)). The *way in zone* represents the path along which the robots can safely reach the *target area* in *Env A*—an area of 25 cm around the light. In contrast, the robots cannot reach the proximity of the light in *Env B*, and in this situation their goal is to leave the band and reach a certain distance from the light source. Robots have to explore the arena, in order to get as close as possible to the light. If they encounter the circular band they have to start looking for the *way in zone* in order to continue approaching the light, and once they find it, they should get closer to the light and remain in its proximity for 30 s. After this time interval, the trial is successfully terminated. If there is no *way in zone* (i.e., the current environment is an *Env B*), the robots should be capable of “recognizing” the absence of the *way in zone* and leave the band by performing antiphototaxis.

Each robot is required to use a temporal cue in order to discriminate between *Env A* and *Env B*, as in Tuci et al. (2004). This discrimination is based on the persistence of the perception of a particular sensorial state (the floor, the light, or both) for the amount of time that, given the trajectory and speed of the robot, corresponds to the time required to make a loop around the light. The integration over time of the robots’ sensorial inputs is used to trigger antiphototaxis in *Env B*.

Robots are provided with a sound signaling system that can be used for communication. The emergence of a signaling convention by which the robots can affect each other’s behavior is entirely open to the dynamics of the evolutionary process.

The fitness function we use does not explicitly reward the use of signaling. Moreover, communication is not strictly required to solve the task considered, as robots that perform the discrimination individually and robots that perform it cooperatively can both be successful. However, our results show that the best evolved solutions to the problem make use of a simple communication system which enables the robots to cooperatively categorize different environmental situations and to display collective decision-making skills.

2.2 The Simulation Model

The controllers are evolved in a simulation environment which models some of the hardware characteristics of the *s-bots* (see Fig. 2(a)). The *s-bots* are wheeled cylindrical robots with a 5.8 cm radius, equipped with a variety of sensors, and whose mobility is provided by a differential drive system (see Mondada et al. 2004). In this work, we make use of four ambient light sensors, placed at -112.5° (L_1), -67.5° (L_2), 67.5° (L_3), and 112.5° (L_4) with respect to the *s-bot*'s heading, fifteen infra-red proximity sensors placed around the turret (P_1 to P_{15}), two floor sensors F_1 and F_2 positioned facing down on the underside of the robot with a distance of 4.5 cm between them, and an omni-directional sound sensor SI (see Fig. 2(b)). The motion of the robot implemented by the two wheel actuators (M_1 and M_2) is simulated by the differential drive kinematics equations, as presented in Dudek and Jenkin (2000), and a loudspeaker S is available for signaling. Light and proximity sensor values are simulated through a sampling technique (see Miglino et al. 1995). The robot floor sensors assume the following values: 0 if the sensor is positioned over white floor; $\frac{1}{3}$ if the sensor is positioned over light gray floor; $\frac{2}{3}$ if the sensor is positioned over dark gray floor; 1 if the sensor is positioned over black floor. The loudspeaker produces a binary output (on/off); the sound sensor detects the presence of sound in the environment (on/off). During evolution, 10% random noise was

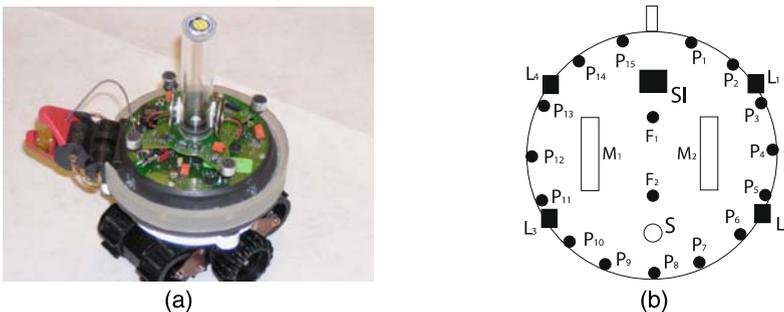


Fig. 2 (a) A picture of an *s-bot*. (b) Sensors and motors of the simulated robot. The robot is equipped with four ambient light sensors (L_1 to L_4), two floor sensors F_1 and F_2 , 15 proximity sensors (P_1 to P_{15}) and a binary sound sensor, called SI (see text for details). The wheel motors are indicated by M_1 and M_2 . S is the sound signaling system (loudspeaker)

added to the light and proximity sensor readings, the motor outputs and the position of the robot. We also added noise of 5% to the reading of the two floor sensors, by randomly flipping between the four aforementioned values. No noise was added to the sound sensor.

2.3 The Controller and the Evolutionary Algorithm

We use fully connected, thirteen neuron CTRNN (see Fig. 3 for a depiction of the network). All neurons are governed by the following state equation:

$$\frac{dy_i}{dt} = \frac{1}{\tau_i} \left(-y_i + \sum_{j=1}^{13} \omega_{ji} \sigma(y_j + \beta_j) + g I_i \right), \quad \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (1)$$

where, using terms derived from an analogy with real neurons, τ_i is the decay constant, y_i represents the cell potential, ω_{ji} the strength of the synaptic connection

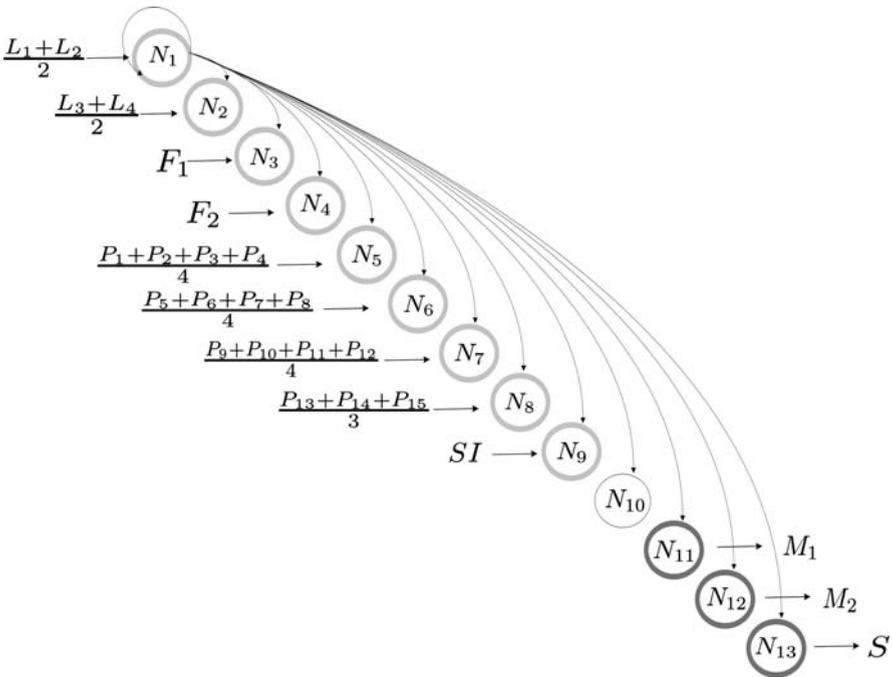


Fig. 3 The fully connected CTRNN architecture. Neurons are represented as circles. Circles with the light gray outline represent the input neurons, while circles with the heavy gray outline represent the output neurons. Only the efferent connections for N_1 are drawn; all other neurons are connected in the same way. We show for all input neurons the combination of sensors that serve as inputs, and for all output neurons the corresponding actuator. N_{10} is not connected to any sensor or actuator

from neuron j to neuron i , $\sigma(y_j + \beta_j)$ the firing rate, β_j the bias term, g the gain and I_i the intensity of the sensory perturbation on sensory neuron i . The connections of all neurons to sensors and actuators is shown in Fig. 3. Neurons N_1 to N_8 receive as input a real value in the range $[0, 1]$. Neuron N_1 takes as input $\frac{L_1+L_2}{2}$, $N_2 \leftarrow \frac{L_3+L_4}{2}$, $N_3 \leftarrow F_1$, $N_4 \leftarrow F_2$, $N_5 \leftarrow \frac{P_1+P_2+P_3+P_4}{4}$, $N_6 \leftarrow \frac{P_5+P_6+P_7+P_8}{4}$, $N_7 \leftarrow \frac{P_9+P_{10}+P_{11}+P_{12}}{4}$ and $N_8 \leftarrow \frac{P_{13}+P_{14}+P_{15}}{3}$. Neuron N_9 receives a binary input (i.e., 1 if a tone is emitted by any of the two agents, 0 otherwise) from the microphone SI . N_{10} does not receive input from any sensor and does not determine the output of any actuator. The cell potentials (y_i) of N_{11} and N_{12} , mapped into $[0, 1]$ by a sigmoid function (σ) and then linearly scaled into $[-4.0, 4.0]$, set the robot motors output. The cell potential of N_{13} , mapped into $[0, 1]$ by a sigmoid function (σ) is used by the robot to control the sound signaling system (the robot emits a sound if $y_{13} \geq 0.5$). The parameters ω_{ji} , τ_i , β_j , and g are genetically encoded. Cell potentials are set to 0 when the network is initialized or reset, and circuits are integrated using the forward Euler method with an integration step-size of 0.1. A simple generational genetic algorithm (GA) is employed to set the parameters of the networks (Goldberg 1989). The population contains 100 genotypes. Each genotype is a vector comprising 196 real values (169 connections, 13 decay constants, 13 bias terms, and a gain factor). Initially, a random population of vectors is generated by initializing each component of each genotype to values chosen uniformly random in the range $[0, 1]$. Subsequent generations are produced by a combination of selection with elitism, recombination and mutation. For each new generation, the three highest scoring individuals (“the elite”) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness-proportional selection from the 70 best individual genotypes of the old population. New genotypes, except “the elite,” are produced by applying recombination with a probability of 0.1 and mutation. Mutation entails a random Gaussian offset that is applied to each real-valued vector component encoded in the genotype, with a probability of 0.15. The mean of the Gaussian is 0, and its standard deviation is 0.1. During evolution, all vector component values are constrained within the range $[0, 1]$. Genotype parameters are linearly mapped to produce CTRNN parameters with the following ranges: biases $\beta_j \in [-2, 2]$, weights $\omega_{ji} \in [-6, 6]$, and gain factor $g \in [1, 12]$. Decay constants are firstly linearly mapped onto the range $[-0.7, 1.7]$ and then exponentially mapped into $\tau_i \in [10^{-0.7}, 10^{1.7}]$. The lower bound of τ_i corresponds to a value slightly smaller than the integration step-size used to update the controller; the upper bound corresponds to a value slightly bigger than the average time required for a robot to reach and perform a complete loop of the band in shades of gray.

2.4 The Fitness Function

During evolution, each individual genotype is coded into a CTRNN controller, and is evaluated for ten trials, five in each environment. Both robots in the ten trials have the same controller, that is, the group of robots is homogeneous. The group of

two robots is evaluated at the end of each trial as a whole. Notice that in Chap. 8 it is demonstrated that using a homogeneous control system favors the emergence of cooperation and communication, while if robots are evolved under individual selection, deceptive communication emerges. The sequence order of environments within the ten trials does not influence the overall performance of the group since each robot controller is reset at the beginning of each trial. Each trial differs from the others in the initialization of the random number generator, which influences the robots' starting positions and orientation, the position and amplitude of the *way in zone* (between 45° to 81°), and the noise added to motors and sensors. Within a trial, the robot life-span is 100 s (1000 simulation cycles). The final fitness attributed to each genotype is the average fitness score of the ten trials. In each trial, the fitness function E is given by the following formula:

$$E = \frac{E_1 + E_2}{2 \times (n_c + 1)},$$

where n_c is the number of (virtual) collisions in a trial, that is the number of times the robots get closer than 2.5 cm to each other (if $n_c > 3$, the trial is terminated) and E_i , $i = 1, 2$, is the fitness score of robot i , calculated as follows:

- If the trial is in *Env A*, or the robot in either environment has not yet touched the band in shades of gray or crossed the black edge of the band, then its fitness score is given by $E_i = \frac{d_i - d_f}{d_i}$. If *robot_i* ends up in the *target area* in *Env A*, we set $E_i = 2$.
- Otherwise, that is if the band is reached in *Env B*, $E_i = 1 + \frac{d_f - 40}{d_{\max} - 40}$. If a robot ends up 120 cm from the light ($d_f = 120$), $E_i = 2$.

d_i is the initial distance of the robot to the light, d_f is the distance of the robot to the light at the end of the trial and $d_{\max} = 120$ cm is the maximum possible distance of a robot from the light. If both robots are successful, the genotype corresponding to their neuro-controller gets the maximum score of 2. An important feature of this fitness function is that it rewards agents that develop successful discrimination strategies and end up doing the correct action in each environment, regardless of any use of sound signaling. That is, a genotype that controls a group that solves the task without any signaling or communication gets the same fitness as one that makes use of communication.

3 Results

In this section, we present a series of post-evaluation tests concerning simulated robots. In particular, in Sect. 3.1, we select and re-evaluate the best evolved strategies of a series of twenty evolutionary simulations. In Sect. 3.2, we show that sound signaling is a functional element of the behavioral strategies in the majority of successful groups of robots. In Sect. 3.3, we run further post-evaluation tests aimed at unveiling the adaptive significance of sound signaling behavior. Notice that one

successful controller using communication has been transferred to real hardware (the *s-bot*) (see Ampatzis et al. 2008, for details).

3.1 A First Series of Post-evaluation Tests

Twenty evolutionary simulation runs, each using a different random initialization, were run for 12,000 generations. Thirteen evolutionary runs produced successful groups of robots. Note that a group is successful if both robots approach the band and subsequently (i) reach the *target area* through the *way in* zone in *Env A*, (ii) leave the band performing antiphototaxis in *Env B*. We arbitrarily demand that the successful accomplishment of this task corresponds to an average fitness score $F \geq 1.8$. In the seven evolutionary runs considered not successful, the fitness score recorded during the evolutionary phase by the best groups at each generation was always lower than 1.8. For each successful run, we chose to post-evaluate the best group of each generation whose fitness score was higher than 1.8.

We employed the average fitness score F over a set of 500 trials in each type of environment as a quantitative measure of the effectiveness of the evolved groups' strategy. Table 1 shows, for each successful evolutionary run (i), the results of the best group among those chosen for post-evaluation. These groups are referred to as g_i . We can notice that all these groups achieve an average fitness score in each

Table 1 Results of post-evaluation tests showing for each best evolved successful group of each evolutionary run (g_i): the average and standard deviation of the fitness over 500 trials in *Env A* and in *Env B*; the average and standard deviation of the percentage of timesteps sound was emitted by either robot over 500 trials in *Env A* and in *Env B*

Group	<i>Env A</i>				<i>Env B</i>			
	Fitness		Signaling (%)		Fitness		Signaling (%)	
	Mean	Sd	Mean	Sd	Mean	Sd	Mean	Sd
g_1	1.92	0.31	0.00	0.00	1.98	0.13	17.39	0.30
g_2	1.94	0.28	0.72	3.72	1.99	0.00	18.22	1.36
g_5	1.99	0.10	0.00	0.00	1.98	0.10	13.36	1.58
g_6	1.96	0.21	0.00	0.00	1.99	0.11	16.47	2.38
g_7	1.99	0.11	0.00	0.00	1.95	0.21	15.06	2.82
g_8	1.96	0.25	0.00	0.00	1.99	0.02	16.47	2.08
g_9	1.99	0.12	0.00	0.00	1.97	0.16	16.38	2.62
g_{10}	1.91	0.31	0.00	0.00	1.91	0.36	0.00	0.00
g_{13}	1.87	0.43	1.72	8.14	1.95	0.09	20.88	2.44
g_{14}	1.96	0.17	0.00	0.00	1.98	0.17	0.00	0.00
g_{16}	1.89	0.33	0.00	0.00	1.94	0.27	0.00	0.00
g_{18}	1.81	0.45	0.00	0.00	1.87	0.16	0.00	0.00
g_{19}	1.91	0.27	0.00	0.00	1.98	0.06	12.65	0.99

environment higher than 1.8 (see Table 1). Thus, they proved to be particularly successful in performing the task. The post-evaluation tests also reveal that among the successful groups, nine groups ($g_1, g_2, g_5, g_6, g_7, g_8, g_9, g_{13}, g_{19}$) make use of sound signaling. In particular, the use of sound strongly characterizes the behavioral strategies of the groups when they are located in *Env B*. In *Env A* signaling is, for all these groups, rather negligible—see Table 1. In groups $g_{10}, g_{14}, g_{16}, g_{18}$, the robots do not emit sound during post-evaluation in either environment.

3.2 Sound Signaling and Communication

In this section we show the results of further post-evaluation tests on the groups in which the robots emit sound during the accomplishment of the task. These tests aim to determine whether sound has a functional significance within the behavioral strategies of the groups and, if the answer is positive, to identify the adaptive function of sound use.

3.2.1 Behavioral Features and Mechanisms

We looked at the behavior of the robots that emit sound during a successful trial in each type of environment. During each trial, we recorded for each robot of a group the distance to the light and the change over time of the sound output (i.e., cell potential of neuron N_{13} mapped into $[0.0, 1.0]$ by a sigmoid function σ). These two variables are recorded both in a *normal* condition and in a condition in which the robots can not hear each other's sound (i.e., the *not-other-sound* condition). In the latter circumstances, the input of neuron N_9 of each robot controller is set to 1 only if the sound in the environment is produced by the robot itself. Figure 4 shows the results of the tests for robots of group g_2 in *Env B* only. We do not show the results of the tests in *Env A* because they are less relevant to the issue of sound, as signaling in *Env A* is rather negligible. We show only the results of one signaling group (i.e., g_2) since it turned out that the groups that emit sound in *Env B* share similar behavioral strategies. Therefore, everything that is said for group g_2 with respect to sound signaling, qualitatively applies to groups $g_1, g_5, g_6, g_7, g_8, g_9, g_{13}, g_{19}$.

In Figs. 4(a) and 4(b), we plot the robot-light distances in the *normal* and the *not-other-sound* condition. In both figures, the areas in shades of gray represent the circular band. From these figures, we can recognize three phases in the behavior of the robots. In the first phase, the robot-light distance initially decreases for both robots (phototaxis phase). When the robots touch the band, the distance to the light remains quite constant as the robots circle around the band trying to find the *way in zone* (integration over time phase). In the third phase the robot-light distances increase and reach their maximum at the end of the trial (antiphototaxis phase). We immediately notice that the behavior of the robots in the *normal* condition (see

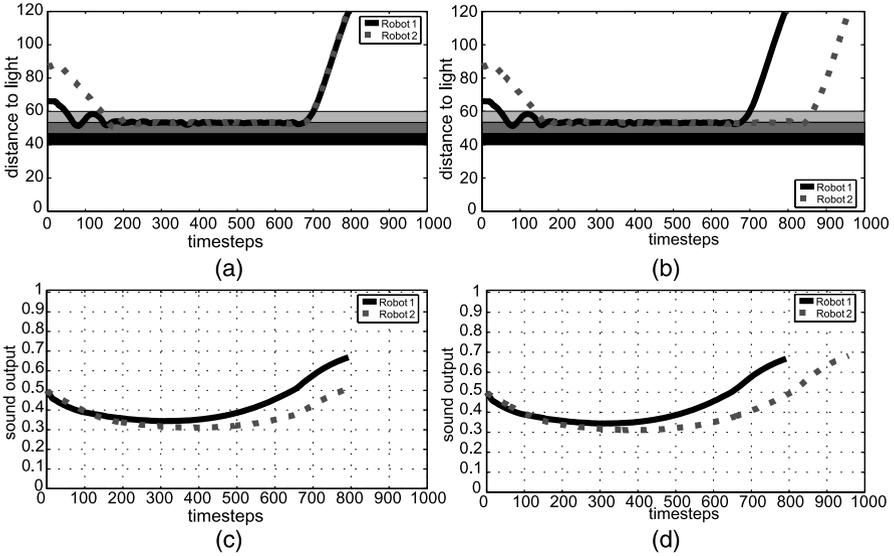


Fig. 4 The graphs show some features of the behavior of the group of robots g_2 at each timestep of a successful trial in *Env B*. Graphs (a) and (b) show the robots' distance to light, in the *normal* and *not-other-sound* condition, respectively. The areas in shades of gray represent the circular band. Graphs (c) and (d) show the cell potential of neuron N_{13} mapped into $[0.0, 1.0]$ by a sigmoid function σ (i.e., the sound output) of each robot controller, in the *normal* and *not-other-sound* condition, respectively. Robot 1—see *continuous lines*—is always initialized closer to the light than Robot 2—see *dashed lines*

Fig. 4(a) only slightly differs from what can be observed in the *not-other-sound* condition (see Fig. 4(b)). The only difference concerns the third phase. In particular, while in the *normal* condition both robots begin to move away from the light at the same time, in the *not-other-sound* condition Robot 2 initiates the antiphototactic behavior after Robot 1. If observed with respect to how the robots' sound output unfolds in time, this small behavioral difference turns out to be an extremely indicative cue as to the function of sound signaling.

Figures 4(c) and 4(d) show that for both robots the sound output changes smoothly and in the same way in both conditions. During the phototaxis phase, the sound output decreases. During the integration over time phase, this trend is reversed. The sound output starts to increase up to the point at which its value rises over the threshold of 0.5. The increment seems to be induced by the persistence of a particular sensory state corresponding to the robot moving around the band. Once the sound output of a robot increases over the threshold set to 0.5, that robot starts emitting a tone. In the *normal* condition we notice that, as soon as the sound output of Robot 1 rises over the threshold of 0.5 (see *continuous line* in Fig. 4(c) around timestep 650) both robots initiate an antiphototactic movement. Robot 2 leaves the band the moment Robot 1 emits a signal, despite the fact that its own sound output is not yet over the threshold of 0.5. Contrary to this, in the *not-other-sound* condition we notice that Robot 2 does not leave the band at the same

time as Robot 1, but it initiates antiphototaxis only at the time when it starts emitting its own sound (see dashed line in Fig. 4(d) around timestep 830).

3.2.2 The Role of Sound

The way in which the distance to the light and the sound output of each robot change over time in the two experimental conditions suggests that the sound is functionally relevant to the accomplishment of the task. In particular, signaling behavior seems to be strongly linked to mechanisms for environmental categorization. As long as the latter mechanisms work properly, the emission of sound after approximately one loop around the light becomes a perceptual cue that reliably indicates to a robot the necessity to move away from the light. Moreover, sound has a communicative function: that is, once broadcast into the environment by one robot (e.g., Robot 1 in *normal* condition), it changes the behavior of the other robot (i.e., Robot 2 in *normal* condition) which stops circling around the light and initiates antiphototaxis (see Figs. 4(a) and 4(b)). To further test the causal relationship between the emission of sound and the switch from phototaxis to antiphototaxis, we performed further post-evaluation tests. In these tests, we post-evaluated group g_2 for 500 trials in *Env A* and 500 trials in *Env B*, in conditions in which the robots are not capable of perceiving sound. That is, their sound input is set to 0 regardless of whether any agent emits a signal. We refer to this condition as the *deaf* setup. We remind the reader that similar phenomena to the one concerning g_2 and illustrated in Table 2, have been observed for all the other signaling groups. As far as *Env A* is concerned, the average fitness of the group does not differ much from the average fitness obtained in the *normal* setup (see Table 2). Concerning *Env B*, the average fitness of the group is lower than the average fitness recorded in the *normal* setup (see Table 2). Moreover, the robots' average final distance to the light is only about the same as the radius of the outer edge of the band (i.e., 60 cm to the light; see Table 2). Given that the robots never collided, the decrease of the average fitness recorded in *Env B* in the *deaf* setup can only be attributed to the fact that the robots do not perform antiphototaxis. This confirms that, in conditions in which the robots can not hear any sound, they do not switch from phototaxis to antiphototaxis. The role of sound is indeed to trigger antiphototaxis in both the emitter and the robot that is not emitting a tone yet.

Notice that the increase in the percentage of signaling in *Env B* observed in the *deaf* setup (51.13% vs. 18.22% in the *normal* setup) can be attributed to the fact that robots in this condition keep on signaling while failing to perform antiphototaxis until the trial ends. Moreover, for the sake of clarity, we should say that, when signaling groups are located in *Env A*, the robots' sound output undergoes a trend similar to the one shown in Fig. 4(c). That is, it decreases during the initial phototactic phase and starts rising during the integration over time phase. However, when the robots are placed in *Env A*, the increment of their sound output is interrupted by the encounter of the *way in* zone. As soon as the robot gets closer to the light via the *way in* zone, the sound output begins to decrease. This process has been shaped by evolution in such a way that, in order for the sound output to rise over the threshold

Table 2 Comparison between the *deaf* and *normal* setups. We show for group g_2 the average and standard deviation of the fitness over 500 trials in *Env A* and in *Env B*; the average and standard deviation of the percentage of timesteps the sound was on by either robot over 500 trials in *Env A* and in *Env B*; the average and standard deviation of the final distance (d_f) of each robot to the light in *Env B*. The row in gray shows again the result of group g_2 in the *normal* condition, with no disruptions applied to the propagation of sound signals

Group g_2											
<i>Env A</i>				<i>Env B</i>							
Fitness		Signaling (%)		Fitness		Signaling (%)		Robot 1 (d_f)		Robot 2 (d_f)	
Mean	Sd	Mean	Sd	Mean	Sd	Mean	Sd	Mean	Sd	Mean	Sd
1.97	0.16	1.35	7.03	1.26	0.09	51.13	4.35	66.52	14.46	54.90	3.12
1.94	0.28	0.72	3.72	1.99	0.00	18.22	1.36	119.65	0.20	119.64	0.20

of 0.5, it must be the case that no *way in* zone has been encountered by the robots. In other words, it takes more or less the time to make a loop around the light while moving on the circular band for a robot's sound output to rise over the threshold. Consequently, when the robot is located in *Env A*, no sound is emitted. Those post-evaluation trials in which sound has been recorded in *Env A* in signaling groups (see Table 1, groups g_2 and g_{13}) were probably due to atypical navigation trajectories which caused the sound output of either robot to rise above the threshold.

Finally, we should say that for all the best-evolved groups of robots, we found that there is a neuron other than the sound output neuron (either a neuron that receives input from the sensors or N_{10}) whose firing rate behaves similarly to neuron N_{13} of the robots in group g_2 . That is, there is a neuron whose firing rate increases in response to the persistence of the sensory states associated with moving around the light on the band. For groups that never emit sound (i.e., g_{10} , g_{14} , g_{16} , g_{18}), if this increase is not interrupted by the encounter of the *way in* zone, it eventually induces antiphototaxis.¹ For groups that emit sound (i.e., g_1 , g_2 , g_5 , g_6 , g_7 , g_8 , g_9 , g_{13} , g_{19}), this mechanism is linked to the behavior of neuron N_{13} as shown in Fig. 4c. The relationship between mechanisms for integration of time and neuron N_{13} is the basic difference between signaling and non-signaling groups.

3.3 On the Adaptive Significance of Signaling

In this section, we provide evidence that there are selective pressures which favor signaling over non-signaling groups, that is, that communication has a selective advantage. More in detail, we demonstrate that this adaptive advantage consists in (i) triggering antiphototaxis faster by exploiting sound signals, and (ii) a more robust and more reliable discrimination between the two environments.

¹See <http://iridia.ulb.ac.be/supp/IridiaSupp2006-007> for supplementary graphs showing the behavior of all neurons and a lesion analysis aimed to prove the functionality of each neuron.

3.3.1 Functions of Sound Signaling

We started our analysis by trying to understand whether during evolution sound had fulfilled functions other than the one we observed in the best evolved groups of robots during the post-evaluation tests shown in Sect. 3.2. To do this, we post-evaluated (500 times in each type of environment) all the best groups at each generation (1 to 12000) of all the successful evolutionary runs. During this post-evaluation we recorded the average fitness in each environment and the average percentage of time per environment either robot emits a signal during a trial. After post-evaluating these groups, we isolated those whose average fitness was higher than 1.8. We noticed that after having excluded (i) those groups that signal throughout the entire duration of a trial in both environments,² (ii) those groups that never signal in a trial in both environments, and (iii) those groups in which sound was not functionally relevant for their behavioral strategies, we were left with groups that signal only in *Env B* for an average time of about one fourth of the duration of a trial. Further investigation on the behavior of these groups revealed that in all of them sound was fulfilling one and only one function: triggering antiphototaxis in *Env B*.

In other words, looking at the behavior of all successful signaling groups of any evolutionary simulation run we discovered that whenever signaling is functionally relevant to the success of the group, it is employed by the robots in *Env B* as a self-produced perceptual cue. This cue induces the emitter as well as the other robot of the group to change its behavior from light-seeking to light-avoidance. This evidence constrains our investigation on the adaptive significance of sound signaling to only a specific case in which we can arbitrarily associate to sound two functionalities: on the one hand, sound is the means by which a robot emitter switches from phototaxis to antiphototaxis. We refer to this as the “solitary” function. On the other hand, sound is the means by which the robot emitter influences the behavior of the other robot. In fact, the perception of the sound triggers antiphototaxis in the emitter as well as in the robot that is not yet emitting a tone (see Figs. 4(a) and 4(c)). We refer to this as the “social” function. In the following, we illustrate the results of post-evaluations that prove and explain why it is the latter functionality which makes a group of signaling robots better adapted than other group types.

3.3.2 The Social Function of Sound Signaling as a Means to Obtain Robustness

The statistics shown in Table 3 refer to a series of tests in which we post-evaluated (500 times in each environment) 100 different groups of robots of five different evolutionary runs (runs 2, 10, 14, 16, 18), chosen among the best of each generation whose average fitness was higher than 1.8. As far as it concerns run 2, we post

²We do not further analyze the cases in which the robots signal throughout the entire duration of a trial since we consider it obvious that in these cases the sound-emitting behavior does not serve any specific function.

Table 3 The table shows the statistics of post-evaluation tests in which 100 different groups of robots of five different evolutionary runs (runs 2, 10, 14, 16, 18), chosen among the best of each generation whose average fitness was higher than 1.8. For run 2, we post evaluated: (i) 100 groups that use sound signaling in the *normal* setup (see row “sig”) and in the *not-other-sound* setup (see row “not-other”); (ii) 100 groups that do not use sound signaling (see row “non-sig”)

Run	Groups	Mean	Sd	Lower quartile	Median	Upper quartile
<i>Env B</i>						
2	sig	1.989	0.082	1.995	1.996	1.997
	non-sig	1.923	0.261	1.964	1.995	1.997
	not-other	1.747	0.268	1.589	1.760	1.982
10	non-sig	1.905	0.308	1.966	1.995	1.997
14	non-sig	1.943	0.226	1.993	1.996	1.997
16	non-sig	1.945	0.210	1.992	1.995	1.997
18	non-sig	1.880	0.326	1.918	1.995	1.997

evaluated: (i) 100 groups that use sound signaling in the *normal* setup (see Table 3 second row “sig”) and in the *not-other-sound* setup (see Table 3 fourth row “not-other”); (ii) 100 groups that do not use sound signaling (see Table 3 third row “non-sig”). Recall that the *not-other-sound* setup refers to the case in which the robots do not hear each other’s sound (see also Sect. 3.2). The 100 non-signaling groups of robots of evolutionary run 2 are “predecessors” of the signaling one. That is, they were the best groups some generations before the evolution of successful signaling groups. By looking at the statistics shown in Table 3 we notice that: (a) the fitness of signaling groups (run 2) is significantly higher than the fitness of any of the non-signaling groups (run 2 “not-sig”, 10, 14, 16, and 18, pairwise Wilcoxon test with 99% confidence interval); (b) the standard deviation of the fitness of signaling groups (run 2) is smaller than the standard deviation of the fitness of any of the non-signaling groups (run 2 “not-sig”, 10, 14, 16, and 18); (c) the fitness of signaling groups (run 2) recorded in the *not-other-sound* condition is significantly smaller than the fitness of any of the non-signaling groups (run 2 “not-sig,” 10, 14, 16, and 18, pairwise Wilcoxon test with 99% confidence interval). We consider (a) and (b) evidence which suggests that indeed signaling groups are on average better than non-signaling groups. Notice that, although the difference among the groups is small, during evolution it may have influenced the distribution of genetic material and consequently the emergence of the behavioral strategies. For the sake of completeness, we also show the lower and upper quartile and the median of the distributions. These data confirm that the difference in performance between the two groups seems to lie in the fact that non-signaling groups display a slightly worse performance than signaling groups in a few cases (see lower quartiles, run 2 “sig” and “non-sig”). We consider (c) as evidence suggesting that the beneficial effect of signaling is not linked to the “solitary” function, since if we prevent signaling robots from hearing each other’s sound (i.e., the *not-other-sound* setup) the “solitary” function is not by itself sufficient to make the robots on average better than those that do not use signaling at all. Consequently, it appears that groups of robots that

use sound signaling have a selective advantage over other types of groups, due to the “social” function of signaling.

In particular, we believe that the selective advantage of signaling groups is given by the beneficial effects of communication with respect to a robust disambiguation of *Env A* from *Env B*. The beneficial effect corresponds to robust individual decision-making and faster group reaction, since signaler and hearer react at the same time. Moreover, the effectiveness of the mechanisms which integrate sensory information over time in order to produce the categorization of the environment is disrupted by the random noise explicitly injected into the simulated world, which strongly affects the sensors’ reading and the outcome of any “planned” action. However, by communicating the outcome of their decision about the state of the environment, signaling groups, contrary to other types of group, might exploit social behavior to counterbalance the disruptive effect of noise on individual mechanisms for environmental discrimination. In total, in those groups in which antiphototaxis is triggered by the perception of sound rather than by an internal state of the controller, a robot which by itself is not capable or not ready yet to make a decision concerning the nature of the environment can rely on the decision taken by the other robot of the group. Therefore, by reacting to the sound signal emitted by the group mate, a robot initiates an action (i.e., antiphototaxis) which it may not have been capable of, or ready to perform, otherwise. If a robot that reacts to the “non-self” produced sound could not have exploited the signal emitted by the other member of its group, it would have wasted precious time orbiting around the light. Eventually, it would have switched to antiphototactic behavior, but due to time limits it would not have been able to reach the maximum possible distance to the light (see d_f in Sect. 2.4). Consequently, the fitness of the group would have been lower.

The performance of signaling groups not only exceeds the performance of non-signaling groups in *Env B*, but also in *Env A* (pairwise Wilcoxon test with a 99% confidence interval). It seems that signaling groups are better adapted to the “danger” of discrimination mistakes in *Env A* than are non-signaling groups, and thus “early” signaling seems to be an issue that has been taken care of by evolution. Our speculation is that once signaling groups evolve, their signaling behavior is *refined*, probably by categorizing the world later than in the case of non-signaling groups. This happens in order to ensure that the chances of a potential disadvantage resulting from social behavior are minimized. In other words, the use of communication in a system can also affect aspects of the behavior not directly related to communication (i.e., the process of integration of inputs over time). This hypothesis explains the low performance recorded in the *not-other-sound* condition, compared to the *normal* condition. When robots emit signals later, the system becomes more robust because the risk of a discrimination mistake in *Env A* is minimized, at the cost of triggering antiphototaxis in *Env B* somewhat later.³ However, this is counterbalanced by the

³See <http://iridia.ulb.ac.be/supp/IridiaSupp2006-007> for data complementing the statistics of Table 3 with results in *Env A* and for data supporting our claim that signaling groups tend to initiate antiphototaxis later than non-signaling groups.

effect of the social behavior as explained above. To summarize, communication delays the moment of categorization, and at the same time anticipates the collective response: putting robustness in *Env A* and social behavior in *Env B* together, we can account for the selective advantage of communication.

4 Conclusions

In this work, we have studied the emergence of communication in a system provided with the necessary hardware (i.e., a “mouth” and “ears”) and in which the use of communication was not predetermined by the experimenter, but left to evolution to shape. It turned out that evolution produced signaling behavior tightly linked to the behavioral repertoire of the agent and that made social behavior more efficient than solitary behavior, even though the former was not explicitly rewarded by the fitness function. In fact, as we have discussed in Sect. 3.3, communication serves to increase the robustness of the categorization.

Owing to the properties of our design methodology (i.e., Evolutionary Robotics), signaling behaviors co-evolved with time-dependent categorization structures, that is, integration of perceptual information over time. In evolutionary terms, the non-reactive mechanism might have paved the way for the evolution of signaling. In fact, we can draw some hints from the evolutionary analysis we performed in Sect. 3.3 concerning the evolution of signaling, which suggest that evolution proceeds in an “incremental” way. We observed that signaling was present in the population before successful solutions started to appear, in all the evolutionary runs that produced signaling groups (data not shown). However, it seemed to have no functional meaning: signals seemed to be produced rather randomly and not with respect to the environmental contingencies. Functional signaling behaviors seem to evolve shortly after evolution produces the first groups able to solve the task without any use of signaling. In other words, communicative solutions seem to be subsequent to non-communicative ones.⁴ A possible illustration of this process is that sound production that was previously irrelevant becomes linked to the already evolved mechanisms for environmental discrimination and then, as we have shown in Sect. 3.3, the solutions making use of communication come to outperform those that do not. Another clue in support of these speculations is the comparison of the mechanisms underpinning behavior in both signaling and non-signaling groups, which was discussed in Sect. 3.2. Both solutions rely on an internal neuron integrating sensory information over time. However, for communicative solutions, the sound output also behaves similarly. What we can take from this discussion is that the evolution of signaling seems to be strongly based on already evolved cognitive structures (discrimination capabilities) of the agents (see Chap. 9 and Nolfi 2005; Mirolli and Parisi 2008).

⁴For another example that shows how the development of individual skills can subsequently favor the development of communication skills, see Chap. 11.

The selective advantage of signaling over non-signaling groups, as detailed in Sect. 3.3, is the reason why we observe the evolution of signaling groups. Moreover, it is the social function of signaling, that is, the communication resulting from it, that makes these groups more fit than others. In other words, we can attribute the evolution of signaling to its social function, and thus to the effect of emitted signals on other members of the group. This observation justifies the use of the word “signal” in order to describe the emission of sound. In fact, according to Maynard Smith and Harper (2003), a signal evolves because of its effect on others. A *signal* is defined as “an act or structure that alters the behavior of another organism, which evolved because the receiver’s response has also evolved.” In contrast, a *cue* is defined as in Hasson (1994): “a feature of the world, animate or inanimate, that can be used by an animal as a guide to future action.” Obviously our robots do emit a sound “as a guide to future action” (to trigger the action of antiphototaxis), but this is not the reason why signaling behaviors emerged in the first place, even if they also display the latter functionality. In an effort to explain the evolution of communication, ethologists considered the existence of cues as precursors of signals and their subsequent *ritualization* into signals (see Tinbergen 1964) as crucial notions. They saw ritualization as the process of the conversion of a movement or action initially not containing any communicative value into one that does. In our case this description is absolutely relevant and we could summarize by saying that the individual categorization seems to be the cue that later on is ritualized into the (acoustic categorization) signal. Indeed, as we said above, social solutions to the problem seem to be subsequent to solitary ones (see also Chap. 9).

In Sect. 3.3.2 we have seen that signaling groups become more robust as they tend to categorize the environment by initiating antiphototaxis later than non-signaling groups. In other words, we observe that the social context has a bearing and effectively alters the behavior of the robots with respect to their decision-making. This observation brings to mind examples from zoology and in particular social foraging. It has been reported that the foraging behavior of animals changes if the animals are situated in a social context. For example, Elgar (1987) shows that social companionship in house sparrows leads to higher feeding rates, as each individual eventually spends less time scanning for predators. Similarly, Fernandez-Juricic et al. (2005) show that while foraging, starlings spend more time scanning for predators once social information is reduced. Overall, we can say that the behavior of our robots is re-shaped (through evolution) as a consequence of the social context in which they are located and the availability at some point in evolution of social information—categorization signals.

To conclude, this work provides another example of evolving communicative skills, even if communication is not directly rewarded by the fitness function. The evolved signals proved to be encoding sensory information integrated over time and to be contributing to increasing the reliability of the categorization process. In other words, the possibility to exploit communication allows the group of robots to perform a more robust discrimination between the two types of environment. Even if the task could also be solved by agents relying only on individual time-dependent strategies, evolution produced mainly social solutions to the problem.

References

- Ampatzis, C., Tuci, E., Trianni, V., & Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, *16*(1), 5–26.
- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, *1*, 91–122.
- Dudek, G., & Jenkin, M. (2000). *Computational principles of mobile robotics*. Cambridge: Cambridge University Press.
- Elgar, M. A. (1987). Food intake rate and resource availability: flocking decisions in house sparrows. *Animal Behaviour*, *35*, 1168–1176.
- Fernandez-Juricic, E., Smith, R., & Kacelnik, A. (2005). Increasing the costs of conspecific scanning in socially foraging starlings affects vigilance and foraging behaviour. *Animal Behaviour*, *69*, 73–81.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Reading: Addison-Wesley.
- Hasson, O. (1994). Cheating signals. *Journal of Theoretical Biology*, *167*, 223–238.
- Maynard Smith, J., & Harper, D. (2003). *Animal signals*. Oxford: Oxford University Press.
- Miglino, O., Lund, H. H., & Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, *2*(4), 417–434.
- Mirolli, M., & Parisi, D. (2008). How producer biases can favour the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, *16*(1), 27–52.
- Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., & Dorigo, M. (2004). Swarm-bot: a new distributed robotic concept. *Autonomous Robots*, *17*(2–3), 193–221.
- 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., & Floreano, D. (2000). *Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines*. Cambridge: MIT Press.
- Tinbergen, N. (1964). The evolution of signalling devices. In W. Etkin (Ed.), *Social behavior and organization among vertebrates* (pp. 206–230). Chicago: University of Chicago Press.
- Tuci, E., Trianni, V., & Dorigo, M. (2004). ‘Feeling’ the flow of time through sensory-motor coordination. *Connection Science*, *16*(4), 301–324.