

Adaptive Behavior

<http://adb.sagepub.com>

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

Christos Ampatzis, Elio Tuci, Vito Trianni and Marco Dorigo

Adaptive Behavior 2008; 16; 5

DOI: 10.1177/1059712307087282

The online version of this article can be found at:
<http://adb.sagepub.com/cgi/content/abstract/16/1/5>

Published by:



<http://www.sagepublications.com>

On behalf of:



[International Society of Adaptive Behavior](http://www.isab-society.org)

Additional services and information for *Adaptive Behavior* can be found at:

Email Alerts: <http://adb.sagepub.com/cgi/alerts>

Subscriptions: <http://adb.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.co.uk/journalsPermissions.nav>

Citations (this article cites 18 articles hosted on the SAGE Journals Online and HighWire Press platforms):
<http://adb.sagepub.com/cgi/content/refs/16/1/5>

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

Christos Ampatzis,¹ Elio Tuci,¹ Vito Trianni,² Marco Dorigo¹

¹*IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium*

²*ISTC-CNR, Roma, Italy*

Communication is of central importance in collective robotics, as it is integral to the switch from solitary to social behavior. In this article, we study emergent communication behaviors that are not predetermined by the experimenter, but are shaped by artificial evolution, together with the rest of the behavioral repertoire of the robots. In particular, we describe a set of experiments in which artificial evolution is used as a means to engineer robot neuro-controllers capable of guiding groups of robots in a categorization task by producing appropriate actions. The categorization is a result of how robots' sensory inputs unfold in time, and, more specifically, of the integration over time of sensory input. In spite of the absence of explicit selective pressure (coded into the fitness function), which would favor signaling over non-signaling groups, communicative behavior emerges. Post-evaluation analyses illustrate the adaptive function of the evolved signals and show that these signals are tightly linked to the behavioral repertoire of the agents. Signals evolve because communication enhances group performance, revealing a "hidden" benefit for social behavior. This benefit is related to obtaining robust and fast decision-making mechanisms. More generally, we show how processes requiring the categorization of noisy dynamical information might be improved by social interactions mediated by communication. In a further series of experiments, we successfully download evolved controllers onto real *s-bots*. We discuss the challenges involved in porting neuro-controllers displaying time-based decision-making processes onto real robots. Finally, the beneficial effect of communication is shown to transfer to the case of a real robot, and the robustness of the behavior against inter-robot differences is discussed.

Keywords communication · decision-making · real robots · signaling · swarm robotics

1 Introduction

Recently, the research work carried out in the context of the SWARM-BOTS project¹ has proven that it is possible to build and control a swarm of autonomous self-assembling robots by using the principles of swarm robotics; that is, a population of simple agents that, by interacting locally with each other and with

the environment, can physically connect in order to form bigger robotic structures. Self-assembling systems can be particularly advantageous as the assembled structure can perform tasks that go beyond the capabilities of a single robot. The work presented by Tuci et al. (2006) and Groß, Tuci, Dorigo, Bonani, and Mondada (2006) shows how an assembled structure can transport an object that is too heavy to be

Correspondence to: Christos Ampatzis, IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium. E-mail: campatzi@ulb.ac.be
Tel.: +32 26502730; Fax: +32 26502715

Copyright © 2008 International Society for Adaptive Behavior (2008), Vol 16(1): 5–26.
DOI: 10.1177/1059712307087282

moved by a single robot. O'Grady, Groß, Mondada, Bonani, and Dorigo (2005) demonstrated how assembled robots can climb a hill whose slope would cause a single robot to topple over. These works highlight the mechanical and control aspects that make the implementation of self-assembly possible. However, as pointed out by Tuci et al. (2006), in these works, the conditions that trigger self-assembly are determined a priori by the experimenter. This might be a limitation, as in these cases the adaptiveness of an autonomous multi-robot system is reduced. The authors claim that when and with whom to assemble are decisions that should be governed as much as possible by robots' environmental contingencies and not determined by the experimenter. An alternative way of treating these issues is to let the switch to collective behavior be controlled by autonomous decision making. That is, the robotic group should be capable of deciding when to initiate collective responses, by identifying the environmental contingencies that demand social behavior.

The work presented in this article is about the design of robot controllers in which decision-making mechanisms to switch from solitary to social behavior (as in the case of switching from single robots to assembled structures) are integrated with the mechanisms that underpin the sensory-motor repertoire of the robot. Even though we do not go as far as integrating self-assembly in this study, we believe that our work clarifies aspects that might improve the autonomy and the adaptiveness of multi-robot systems. In particular, we prove that time-based decision-making mechanisms can be designed to allow real robots to perform a categorization task. Moreover, we look at issues directly implicated in the switch from solitary to collective behavior, such as the emergence of a communication system and its relation to the individual decision making.

This work brings the problem of decision making, together with the interest in self-organizing communicating systems, to a real world scenario. This scenario allows the empirical investigation of the switch from individual to collective behavior via an emergent communication protocol. In particular, this switch is governed by time-based decision-making structures that integrate over time sensory information available to the robot. The tool we use to implement such structures is the continuous time recurrent neural network (CTRNN; Beer & Gallagher, 1992) shaped by arti-

cial evolution. These structures should allow robots to initiate social behavior in response to the persistence of certain environmental stimuli. Because of the number of trials needed to test individuals, the design of robot controllers by means of artificial evolution is usually carried out by using simulation models. However, the digital medium might fail to take into account phenomena that impact upon the functional properties of the evolved controllers. As a consequence, controllers evolved in simulation might be less effective in managing real-world sensing and actuation (see also Matarić & Cliff, 1996). One of the main contributions of our work is to show that evolved CTRNNs successfully control real robots. This is a practice that has to be taken into account to ensure that the behaviors we want our robots to display are viable and observable in the real world and not only in a simulated environment. There are several works in the literature that deal with porting an artificial neural network (ANN) able to display memory to reality. Paine and Tani (2005), Blynel and Floreano (2003), and Jakobi (1997) all port evolved CTRNNs onto real Khepera robots, but although the networks used are non-reactive, the tasks described (variations of the T-maze) are in essence solved by switching through reactive strategies (see Ziemke & Thieme, 2002). Urzelai and Floreano (2001) downloaded a plastic neural network (PNN) on a real Khepera, but the solution to the task is also reactive. Quinn, Smith, Mayley, and Husbands (2002) report on work carried out on real hardware (on a collective task), but the network they use is based on model spiking neurons. To the best of our knowledge, there is no work reported where a CTRNN is ported to a real robot, for a task that requires the integration over time of the robot's perception. In this respect, it is worth noting that the decision-making mechanism relies on the continuum of the sensory information (i.e., how the sensory inputs unfold in time) in order to determine subsequent actions. Therefore, the main challenges in porting to reality are the possible disruptive effects on the evolved mechanisms caused by the sensor/actuator noise present in reality, as well as potential inter-robot differences.

One goal of our work is to provide evidence that CTRNNs are capable of displaying complex internal dynamics, such as the integration over time of the robot's perceptual flow, even when tested with physical robots. The real robot experiments are, we believe, the ultimate test-bed for the effect and performance of any

communication protocol, and also serve as a connecting link to more engineering-driven applications. We show, for example, that inter-robot differences not anticipated in simulation can lead to very different levels of performance for communicating versus non-communicating teams.

In this article, we also investigate and unveil the structure of the behavior used in a communicative context and we account for its evolution. The results of this work raise issues concerning the importance and the implementation of communication with respect to a collective robotics scenario. In previous work (Tuci, Trianni, & Dorigo, 2004; Tuci, Ampatzis, & Dorigo, 2005), a similar problem was studied, in which the actions of a simulated robot were determined by the way sensory information unfolds through time. However, these issues were not studied in a social context. The difference here is that we study the collective response to the individual decision making, based on the integration over time of sensory information. In other words, we study the group reaction to the individual categorization of the environment. Communication is the way in which the collective group response can be triggered, once one or more robots within the group take a decision (see Trianni & Dorigo, 2006; Tuci et al., 2006, for examples). 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 in both cases 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 (Nolfi, 2005). We have left the development of communicative behavior entirely to artificial evolution in this way because we believe that the co-adaptation of all the mechanisms involved can produce more effective ways to categorize sensory-motor information. Evolution can produce solutions better adapted to the problem than hand-coded signaling behavior (see Trianni & Dorigo, 2006, for an example).

1.1 Biological Background

In our study, we focus on the evolution of communication in the form of a simple signaling system. Nature abounds with examples from social species,

where simple (compared to human communication and language) signaling mechanisms are used. For example, the alarm calls of vervet monkeys given with respect to the type of predator approaching have been studied in depth (Struhsaker, 1967). Alarm calls are also observed in bird species, squirrels, and so forth (Hauser, 1997; Sherman, 1977). Food calls are another example of cooperative signaling. Animals such as chimpanzees attract conspecifics once they discover food resources. The dance of the honey bee is possibly the most elaborate and striking example (von Frisch, 1967).

Since Darwin, scientists have been trying to explain the evolution of such altruistic signals in animal societies. Ethologists justified the existence of such cooperative and honest signaling by invoking group selection theory: animals behave in such ways in order to maximize the benefit of the group or the species (see, for example, Tinbergen, 1964). However, the alternative of kin selection was presented (Hamilton, 1964) and the naïve application of group selection as an explanation was shown to be unwarranted (Williams, 1966; Dawkins, 1976). Kin selection suggests that animals can behave with apparent altruism towards conspecifics as this can be to their own long-term genetic benefit.

Game theoretical models in the 1970s and 1980s mathematically demonstrated that cheating strategies will normally invade populations of honest signalers (Maynard Smith, 1982). Thus, the interest of researchers focused on how to identify conditions that can lead to the emergence of stable cooperation (e.g., Hamilton's kin-selection theory, reputation-based models, or the effect of topology). The game theoretical models studying such issues typically consider signaling capabilities that are built into the agent's behavioral capacity. Thus, they do not allow the investigation of the origin of signaling behavior.

The experimental setup we use in this work differs in several aspects from these game theoretical models. First, we are attempting to study the origin of signaling, as signaling capabilities are evolved (sensors and effectors are available for communication, but there is no requirement that the robots use them). More specifically, we discuss the existence of possible cues that serve as precursors for the signals employed by our robots, through the process of ritualization. Second, in our work the possibility of cheating and dishonest signaling is excluded because the evaluation of the fitness of a group of individuals is

carried out at the group level and the individuals composing the group are genetically identical clones. Our aim is to understand how communication may emerge in a robotic system, in the absence of explicit selective pressures. In other words, we aim to understand the conditions under which a group of agents will switch to social behavior, and the implications of that switch for the performance of the group in a certain scenario. Our focus is more on the evolution of signaling than the evolution of cooperation. Our implementation has been influenced by an ethological perspective, but this does not mean that we are trying to perform robot ethology. Nor do we claim that our results will necessarily have any bearing on the biological literature regarding the evolution of communication.

1.2 Structure of the Article

In Section 2, we give details of the task addressed. The simulation model used is presented in Section 2.2, while the controller and the evolutionary algorithm are introduced in Section 2.3. In Section 2.4, we describe the fitness function employed to evolve the desired behavior. The results of the experiments conducted are presented in Section 3. We first report on the results of the experiments in simulation (Section 3.1), revealing the functionality of the evolved signaling behaviors (Section 3.2). We then discuss the portability of evolved controllers onto real robots (Section 3.3). In Section 4, we treat the issue of the adaptive significance of signaling. Finally, in Section 5, we draw conclusions by discussing, on the one hand, the relevance of this work for collective robotics and, on the other hand, our contribution to the understanding of the principles underlying the evolution of communication in embodied agents.

2 Methods

In this article, we exploit evolutionary robotics (ER; Nolfi & Floreano, 2000), as the methodology to design controllers capable of providing the robots with the mechanisms required to solve the task described below. Roughly speaking, ER is a methodological tool to automate the design of robot controllers. ER is based on the use of artificial evolution to find sets of parameters for ANNs that guide the robots to the accomplishment of their objective.

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 (i.e., by integrating their perceptual inputs over time). At the start of each trial, two simulated robots are placed in a circular arena with a radius of 120 cm (see Figure 1), 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, centered around the lamp covering an area between 40 and 60 cm from it. The band is divided into three subzones 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 Figure 1a). In the other type, referred to as *Env B*, the band completely surrounds the light (see Figure 1b). 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. Once they find it, they should move 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




Figure 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.

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*. Communication is not required to solve the task considered. In particular, the fitness function we use does not explicitly reward the use of signaling, in contrast with Tuci et al. (2004). However, 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. This issue is discussed further in Section 3.

2.2 Simulation Model

The controllers are evolved in a simulation environment, which models some of the hardware characteristics of the *s-bots* (see Figure 2a). 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 (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, 15 infra-

red proximity sensors placed around the turret (P_1-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 Figure 2b). 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 possible signaling. Light and proximity sensor values are simulated through a sampling technique (Miglino, Lund, & Nolfi, 1995). The robot floor sensors assume the following values: 0 if the sensor is positioned over the white floor; $1/3$ if the sensor is positioned over the light gray floor; $2/3$ if the sensor is positioned over the dark gray floor; 1 if the sensor is positioned over the black floor. The loudspeaker produces a binary output (on/off); the sound sensor has no directionality or intensity features. During evolution, 10% random noise was 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. The reason for this last choice is the fact that the sound sensor proved to be 100% reliable in reality. Of course, adding noise to the sound sensor would force the simulation to address the issue of the reliability of

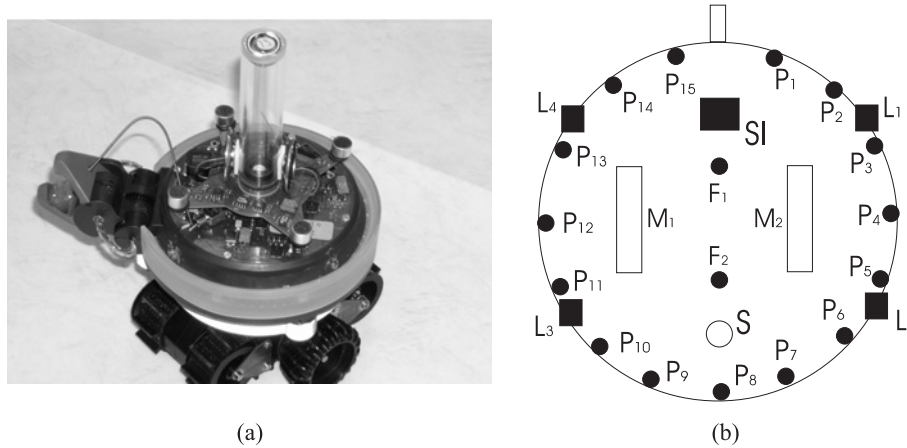


Figure 2 (a) A picture of an *s-bot*. (b) Sensors and motors of the simulated robot. The robot is equipped with four ambient light sensors (L_1 – L_4), two floor sensors (F_1 and F_2), 15 proximity sensors (P_1 – P_{15}) and a binary sound sensor, called *SI* (see text for details). The wheel motors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




Figure 3 The fully connected CTRNN architecture. Neurons are represented as circles. Circles with a light gray outline represent the input neurons, while circles with a heavy gray outline represent the output neurons. Only the efferent connections for N_i 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.

the evolved signals and thus produce neural mechanisms able to cope with noisy communication. This issue, while an interesting one, is beyond the scope of the current article.

2.3 Controller and Evolutionary Algorithm

We use fully connected, 13 neuron CTRNNs (Beer & Gallagher, 1992, see Figure 3 for a depiction of the network). All neurons are governed by the following state equation:

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

$$\sigma(x) = \frac{1}{1 + e^{-x}}. \quad (1)$$

Here, using terms derived from an analogy with real neurons, τ_i is the decay constant, y_i represents the cell potential, ω_{ji} is the strength of the synaptic connection from neuron j to neuron i , $\sigma(y_j + \beta_j)$ is the firing rate, β_j is the bias term, g is the gain, and I_i is the intensity of the sensory perturbation on sensory neuron i . The connections of all neurons to sensors and actuators is shown in Figure 3. Neurons N_1 – N_8 receive as input a real value in the range $[0, 1]$. Neuron N_1 takes as input

$$\frac{L_1 + L_2}{2}$$

and

$$N_2 \leftarrow \left(\frac{L_3 + L_4}{2} \right),$$

$$N_3 \leftarrow F_1,$$

$$N_4 \leftarrow F_2,$$

$$N_5 \leftarrow \left(\frac{P_1 + P_2 + P_3 + P_4}{4} \right),$$

$$N_6 \leftarrow \left(\frac{P_5 + P_6 + P_7 + P_8}{4} \right),$$

$$N_7 \leftarrow \left(\frac{P_9 + P_{10} + P_{11} + P_{12}}{4} \right),$$

$$N_8 \leftarrow \left(\frac{P_{13} + P_{14} + P_{15}}{3} \right).$$

Neuron N_9 receives a binary input (i.e., 1 if a tone is emitted by either agent, and 0 otherwise) from the microphone SI , while neurons N_{10} – N_{13} do not receive input from any sensor. 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. It is important to mention that the

speed that these values translate to is not the maximum possible speed of the robot, but only half of it. This is because, after some initial experimentation, we have found that if we use a faster robot, we have a higher chance of obtaining a false reading from the floor sensors and, in general, a worse sensory-motor coordination. 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 individuals of the old population. New genotypes, except “the elite,” are produced by applying recombination with a probability of 0.1 and mutation. Mutation entails that a random Gaussian offset 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 larger than the average time required for a robot to reach and perform a complete loop of the band in shades of gray.

2.4 Fitness Function

During evolution, each genotype is coded into a robot controller, and is evaluated for 10 trials, five in each environment. Both robots in the 10 trials have the same controller (homogeneous system). The sequence order of environments within the 10 trials does not influence the overall performance of the group as 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° and 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 10 trials. In each trial, the fitness function E is given by the following formula:

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

Here, n_c is the number of (virtual) collisions in a trial, which is the number of times the robots come 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.

1. 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 $E_i = (d_i - d_f)/d_i$.
2. Otherwise (i.e., if the band is reached in *Env B*), $E_i = 1 + [(d_f - 40)/(d_{\max} - 40)]$.

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. In cases where robot i ends up in the *target area* in *Env A*, we set $E_i = 2$. From the above equations, we can see that this is also the maximum value of E_i that a robot can obtain in *Env B*, which corresponds to the robot ending up 120 cm from the light ($d_f = 120$). So, if both robots are successful, the trial obtains 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 performing the correct action in each environment, regardless of any use of sound signaling. That is, a genotype that controls a group which solves the task without any signaling or communication has the same fitness as one that makes use of communication.

3 Results: From Simulated Agents to Real Robots (The *S-Bots*)

In this section, we present a series of post-evaluation tests concerning both simulated and real robots. In particular, in Section 3.1, we select and re-evaluate the best evolved strategies of a series of 20 evolutionary simulations. In Section 3.2, we show that sound signaling is a functional element of the behavioral strategies in the majority of successful groups of robots. In Section 3.3, we report on the results of experiments in which we test the capability of one of the best neural networks evolved in simulation when controlling the behavior of real robots engaged in the task illustrated in Section 2.1. In Section 4, we run further post-evaluation tests aimed at unveiling the adaptive significance of sound signaling behavior.

3.1 Simulated Agents: A First Series of Post-Evaluation Tests

We ran 20 evolutionary simulation runs, each using a different random initialization, for 12,000 generations. Of these, 13 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*, and (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 those 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.

The post-evaluation tests are meant to provide a better estimate of the behavioral capabilities of these groups. In fact, the fitness of the best evolved controllers during evolution may have been an overestimation of their ability to guide the robots in the task. In

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* (see columns 2 and 3) and in *Env B* (see columns 6 and 7); the average and standard deviation of the percentage of time-steps sound was emitted by either robot over 500 trials in *Env A* (see columns 4 and 5) and in *Env B* (see columns 8 and 9).

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

general, the best fitness scores take advantage of favorable conditions, which are determined by the existence of inter-generational variation in starting position and orientation and other simulation parameters. The entire set of post-evaluations should establish whether the groups chosen from the 13 successful runs can effectively solve the task and at the same time ascertain whether signaling behavior characterized the successful strategies. 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 see that all these groups achieve an average fitness score in each environment higher than 1.8 (see Table 1, columns 2, 3, 6, and 7). 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, columns 4, 5, 8, and 9, which refer to the average percentage and standard deviation of the time either robot emits a signal during a trial). In groups $g_{10}, g_{14}, g_{16},$ and g_{18} , the robots do not emit sound during post-evaluation in either environment.

3.2 Sound Signaling and Communication

The results of post-evaluation analyses carried out so far have shown that in nine of the best evolved groups, the robots emit sound during the accomplishment of the task in *Env B*. Note that the emission of sound is not demanded in order to navigate towards the target

and discriminate *Env A* from *Env B*. Indeed, the task and the fitness function do not require the robots to display signaling behavior (see Section 2.4). Mechanisms for phototaxis, antiphototaxis, and memory are sufficient for a robot to accomplish the task. Therefore, in this section we show the results of further post-evaluation tests on those 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 cannot 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. In fact, we have already shown that in *Env A* the robots of signaling groups either do not emit sound at all, or they do it in such a way that it is clear that the sound is not functional within that particular environment (see Table 1, columns 4, and 5, groups g_1 , g_2 , g_5 , g_6 , g_7 , g_8 , g_9 , g_{13} , and g_{19}). We show only the results of one signaling group (i.e., g_2) as it turned out that the groups that emit sound in *Env B* share the same behavioral strategies. Therefore, everything that is said for group g_2 with respect to sound signaling, applies to groups g_1 , g_5 , g_6 , g_7 , g_8 , g_9 , g_{13} , and g_{19} .

In Figures 4a and 4b, solid and dashed lines refer to the robot–light distances in the *normal* and *not-other-sound* conditions, respectively. 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 Figure 4a) only slightly differs from what is observed in the *not-other-sound* condition (see Figure 4b). 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.

Figures 4c and 4d 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 light on 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 solid line in Figure 4c around time-step 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 Figure 4d around time-step 830).

3.2.2 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 mecha-




Figure 4 The graphs show some features of the behavior of the group of robots g_2 at each time-step of a successful trial in *Env B*. (a) and (b) show the robots’ distance to the light. The areas in shades of gray represent the circular band. (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. (a) and (c) refer to the *normal* condition. (b) and (d) refer to the *not-other-sound* condition (i.e., the robots do not hear each other’s sound). Robot 1 (see solid lines) is always initialized closer to the light than robot 2 (see dashed lines).

nisms 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 the *normal* condition), it changes the behavior of the other robot (i.e., robot 2 in the *normal* condition), which stops circuiting around the light and initiates antiphototaxis (see Figures 4a and 4b).

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 (i.e., 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 that concerning g_2 , and illus-

Table 2 *Deaf* setup (robots' sound inputs set to 0). The results of the post-evaluation test showing for group g_2 the average and standard deviation of the fitness over 500 trials in *Env A* (see columns 1 and 2) and in *Env B* (see columns 5 and 6); the average and standard deviation of the percentage of time-steps the sound was on by either robot over 500 trials in *Env A* (see columns 3 and 4) and in *Env B* (see columns 7 and 8); the average and standard deviation of the final distance (d_f) of each robot to the light in *Env B* (see columns 9–12). The row in bold 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

trated in Table 2, have been observed for all the other signaling groups. As far as it concerns *Env A*, the average fitness of the group does not differ much from the average fitness obtained in the *normal* setup (see Table 2, columns 1 and 2). Concerning *Env B*, the average fitness of the group is lower than the average fitness recorded in the *normal* setup (see Table 2, columns 5 and 6). 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, columns 9–12). 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 cannot 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.

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 that shown in Figure 4c. 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 comes 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 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, columns 4 and 5, groups g_2 and g_{13}) were probably a result of 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 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} , and g_{18}), if this increase is not interrupted by the encounter of the *way in* zone, it eventually induces antiphototaxis.² For groups that emit sound (i.e., g_1 , g_2 , g_5 , g_6 , g_7 , g_8 , g_9 , g_{13} , and g_{19}), this mechanism is linked to the behavior of neuron N_{13} , as shown in Figure 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 Transfer to Real Robots

The task described in this article is characterized by the fact that not only the change but also the persist-

ence of particular sensorial states is directly linked to the effectiveness of the evolved strategies (see previous section). These strategies are generated by robot controllers developed in a simulated world that is responsible for modeling the sensory states of *s-bots* acting in *Env A* or *Env B*. Our simulated world (see Section 2.2) models only a small subset of the *s-bot* world physics, as it has been designed to speed up a particularly long evaluation process (i.e., 12,000 generations, 100 genotypes, 10 evaluation trials for each genotype, 1000 simulated time cycles for each trial). As mentioned in Section 2.2, we compensate for the effect of those physical phenomena not modeled (e.g., acceleration, friction, etc.), by adding random noise to the light and proximity sensor readings, the motor outputs, the position of the robot, and the reading of the two floor sensors. However, there is always the risk that the physics of our simulated world are insufficiently or incorrectly defined, and that the evolved behavioral strategies exploit loopholes that limit their effectiveness to an unrealistic scenario. Porting the controllers evolved in simulation onto a real robot is the best way to rule out the above-mentioned problem (Brooks, 1992). As pointed out in Section 1, this step has not been taken in previous research work in which CTRNNs have been evolved to deal with tasks that required the integration over time of sensory states. In this article, we provide evidence of the “portability” of the evolved controllers by showing the results of tests in which real robots are repeatedly evaluated in *Env A* and *Env B*. We chose to re-evaluate the controller of the successful group g_2 because this group during post-evaluation achieved a very high performance. Also, in preliminary tests, among other equally suc-

cessful controllers, this group seemed to achieve the best sensory-motor coordination when downloaded on a group of real robots. Experiments are performed with groups of two and four *s-bots*.

Jakobi (1997) claims that the robot does not have to move identically in simulation and reality in order for the porting to be called successful. In fact, it is enough that its behavior satisfies some criteria defined by the experimenter. Following this principle, real robots are considered successful if they carry out the main requirements of our task. That is, the robots have to reach the band in shades of gray regardless of the type of environment and subsequently (i) end up in the *target area* in *Env A*, without crossing the inner black edge of the circular band or (ii) end up as far as possible from the light in *Env B*. The robots should also avoid collisions.

3.3.1 Experiments with Two S-Bots In our real-world experimental setup, two *s-bots* ($s-bot_1$ and $s-bot_2$) are randomly positioned at a distance of 85 cm from the light. We performed 40 trials, 20 in each environment. Each trial differs from the others for the randomly defined initial position and orientation of the robots, and for the position of the *way in* zone in *Env A*. The initial position of the robots is randomly chosen among one of the 16 possible starting positions that surround the light (see Figure 5a). The width of the *way in* zone is fixed to 45° , which is the smallest value encountered during evolution and the most difficult case for a possible misidentification of an *Env A* for an *Env B*. The *s-bots* proved to be 100% successful in both environments: there were no mistakes in




Figure 5 The experimental setup. (a) A picture of the arena, with the points around the band showing the locations where the robots were randomly positioned. (b) A snapshot of a trial in which two robots find the *way in* zone in *Env A*. (c) A snapshot of a trial in *Env B*. The robot with the lighter turret color is the one that has signaled the absence of a *way in* zone. Both robots have left the band and are performing antiphototaxis.

discrimination, no collisions, and no crossing of the black edge of the band.³ As was the case for the simulated robots of group g_2 , the *s-bots* accomplished the task by using sound in a communicative context. That is, the sound emitted by one *s-bot* triggers antiphototaxis in both robots. The following paragraphs provide further quantitative descriptions of the behavior of simulated and real robots. These data will help to quantify the extent to which the behavior of simulated robots diverges from the behavior of real robots and to evaluate the reliability of our simulated world as a tool for developing controllers for real robots.

Given the nature of the successful strategy of group g_2 , the start of the emission of a tone can be used as a sign that precisely indicates when an *s-bot* has reached the conclusion that it is located in *Env B* rather than *Env A*. We compute the offset between the entrance position in the circular band of the robot that first emits a signal and the position at which this robot starts to signal. This measure, called offset Δ , takes value 0° if the robot signals exactly after covering a complete loop around the circular band. Negative values of offset Δ suggest that the robot signals before having performed a complete loop, while positive values correspond to the situation in which the robot emits a tone after having performed a loop around the light (see Tuci et al., 2004, for details on how to calculate Δ). The offset Δ is used to compare the behavior of simulated and real robots.

During the tests on real robots, we observed that in *Env B* it is always *s-bot*₁ that emits a signal. As shown in Table 3, we see that the *s-bot* that first emits a signal does so on average before completing a loop. However, given that the magnitude of the offset Δ is smaller than the width of the *way in* zone, the group does not run the risk of misinterpreting an *Env A* as an

Env B. Further tests have proved that, if left to act alone in an *Env B*, *s-bot*₂ always signals after completing a loop (i.e., positive offset Δ ; data not shown). This result can be accounted for by noting the existence of various arbitrary mechanical and sensor differences between the two *s-bots*; inter-robot differences that are impractical to include in the simulated world. Contrary to the *s-bots*, the simulated robots of group g_2 signal on average after completing the loop (see Table 3). The mismatch between the behavior of simulated and real robots controlled by the same neural network is an estimate of the magnitude of the divergence between the simulated and real worlds. However, given that our real robots were 100% successful in both environments, we conclude that the noise injected into the simulated world was sufficient to cross the “reality gap” (Jakobi, 1997) and to capture the variability of the behavior of sensors and actuators of real hardware, which can easily disrupt the effectiveness of the evolved neural mechanisms. Note that the successful porting of the controller of a group (i.e., g_2) does not necessarily imply that controllers of other groups that were successful in simulation would be equally successful in guiding real robots. For example, if the effects of inter-robot differences on the mechanisms used for environmental discrimination induce robots to anticipate (with respect to what the group does in simulation) the emission of a signal, then simulated groups with $\Delta \in [-45^\circ, 0^\circ)$ could be more likely to fail. In fact, these groups might fall into the error of signaling and consequently performing antiphototaxis even if placed in an *Env A*.

3.3.2 Experiments with Four S-Bots We also performed a further experiment by porting the controllers of group g_2 to a group of four different *s-bots*: *s-bot*₃, *s-bot*₄, *s-bot*₅, and *s-bot*₆. The aim of these tests is twofold. First, we test the ability of the evolved controllers to accomplish the task, despite the cardinality of the group being higher than that experienced during the evolutionary phase. Second, we evaluate the effectiveness of our controllers with respect to individual differences among the robots. The experiment consists of evaluating for 10 trials the four-robot group in each environment. The results are once again almost perfect.⁴ In all trials in *Env A*, the robots found the *way in* zone without erroneously emitting a tone or crossing the inner black edge of the band in shades of gray. In

Table 3 Average and standard deviation of the offset Δ recorded for different group types.

Groups	Offset Δ	
	Avg (degree)	SD (degree)
Two <i>s-bots</i>	- 30.6	11.75
Four <i>s-bots</i>	+ 18.22	12.97
Simulated robots	+ 31.6	16.05

Env B, we noticed that *s-bot*₃ was always signaling first, but never too early. In all trials, *s-bot*₅ never emitted a tone, and in a separate test we discovered that it was signaling much too late, even after more than two full loops around the band.⁵ Nevertheless, all robots reacted properly to the signal emitted by *s-bot*₃, left the band, and reached the appropriate distance from the light. In Table 3 we see that, with respect to offset Δ , the behavior of the four-robot group is closer to simulation than the behavior of the two-robot group.

4 On the Adaptive Significance of Signaling

The results illustrated in Section 3.1 have shown that the majority of successful strategies employ signaling behavior and communication among the members of the groups. This suggests that our decision to equip the robots with “ears” and a “mouth” turned out to be helpful. However, by simply looking at the characteristics of our model, we cannot necessarily see why evolution exploited these robot’s structures to develop a simple form of communication. In principle, groups in which the use of sound is functionally relevant for the success of the group, and groups in which it is not, can be equally successful. Yet, the majority of the evolutionary runs that ended successfully (i.e., 9 out of 13 best evolved groups) are characterized by group strategies that make use of sound signaling and communication among the robots (see Table 1). How can we account for this result?

It might be that there is in fact no selective advantage for groups in which the use of sound is functionally relevant to their success with respect to alternative types of group. The evolution of signaling might simply be a result of the effect of statistical drift of genetic material over time in populations of simulated agents (i.e., genetic drift). However, we have collected evidence that inclines us to rule out the genetic drift hypothesis, and that supports the idea that there are selective pressures which favor signaling over non-signaling groups. The rest of this section is dedicated to this issue.

4.1 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 Section 3.2. To do this, we post-evaluated (500 times in each type of environment) all the best groups at each generation (1–12,000) 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,⁶ (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 into 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 Figures 4a and 4c). 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 that makes a group of signaling robots better adapted than other group types.

4.2 Social Function of Sound Signaling as a Means to Obtain Robustness

The statistics shown in Table 4 refer to a series of tests in which we post-evaluated (500 times in each envi-

Table 4 This table shows the statistics of post-evaluation tests for 100 different groups of robots of five different evolutionary runs (runs 2, 10, 14, 16, and 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

ronment) 100 different groups of robots of five different evolutionary runs (runs 2, 10, 14, 16, and 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-evaluated: (i) 100 groups that use sound signaling in the *normal* setup (see Table 4, second row “sig”) and in the *not-other-sound* setup (see Table 4, fourth row “not-other”); (ii) 100 groups that do not use sound signaling (see Table 4, 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 Section 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 4 we notice the following. (i) 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). (ii) 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). (iii) 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 (i) and (ii) as empirical evidence that 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 quartiles 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 (iii) as evidence suggesting that the beneficial effect of signaling is not linked to the “solitary” function, because 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, because of 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, as signaler and hearer react at the same time. Moreover, the effectiveness of the mechanisms that 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 that by itself is not capable or not yet ready 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. The experiments performed on real hardware provide perfect examples of the benefits of communication and social behavior, given the presence of severe disruptions as a result of inter-robot differences. For example, in the experiments with two real robots, we have seen that *s-bot*₂ signals always later than *s-bot*₁ (see Section 3.3.1). In the four-robot experiments (see Section 3.3.2), we noticed that *s-bot*₄, *s-bot*₅, and *s-bot*₆ repeatedly benefit from the sound signal emitted by *s-bot*₃, which is the fastest (as well as proving extremely accurate) robot to signal the absence of the way in zone in *Env B*. 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 because of time limits it would not have been able to reach the maximum possible distance to the light (see d_f in Section 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 non-signaling groups. 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 (e.g., see Table 3: simulated robots of group g_2 signal on average after completing a loop—rather late). 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 (high offset Δ), 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* later.⁷ However, this is counterbalanced by the effect of the social behavior, as explained above. To summarize, communication delays the moment of categorization (larger offset Δ), 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.

5 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 agents and this made social behavior more efficient than solitary behavior, even though the former was not explicitly rewarded by the fitness function. In fact, as discussed in Section 4, communication serves to increase the robustness of the categorization. This study contributes to the understanding of issues concerning the evolution of communication, and, more specifically, the identification of conditions that might facilitate the emergence of communication in populations of embodied agents.

It should definitely be acknowledged that there are elements in our experimental setup that facilitate the evolution of cooperative behavior: (i) the fact that our robot group is composed of genetically identical clones; (ii) the fact that it is the behavior of the group that we evaluate after each trial. The above factors leave no room for conflicts of interest or cheating.

However, we should stress again that the evolution of signaling and thus cooperation is neither trivial or obvious, in particular because it is not explicitly favored by the fitness function. The use of a functional and meaningful signaling system is not a question with a binary answer based on chance. In order for a signaling system to evolve, evolution must produce appropriate signals, appropriate reactions to signals and a reorganization of the decision-making mechanisms to ensure the robustness of the system, as seen in the previous section. Still, we can certainly ask what might happen if the individuals in the robotic group were not genetically identical, and each robot had a different controller and was evaluated only on the basis of its own performance. Floreano, Mitri, Magnenat, and Keller (2007) report on a series of experiments aimed at studying the evolutionary conditions for the emergence of visual communication. They note that “under individual selection, the ability to produce visual signals resulted in the evolution of deceptive communication strategies in colonies of unrelated robots and a concomitant decrease in colony performance.” Clearly, the aspects of our experimental setup that prevented conflicts of interest and deceptive communication between the robots have had some bearing on our observed results. In future work, we intend to withdraw these assumptions and address the broader issue of conditions under which communication may evolve despite the absence of explicitly group-level selection.

Owing to the properties of our design methodology (i.e., ER), signaling behaviors co-evolved with time-dependent categorization structures, that is, integration over time. In evolutionary terms, these non-reactive mechanisms might have paved the way for the evolution of signaling. In fact, we can draw some hints from the evolutionary analysis we performed in Section 4 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. 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 solutions. A possible illustration of this process is that sound production that was previously irrelevant becomes linked to the already evolved mechanisms for environmental discrimination. Then, as shown in Section 4, 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, as discussed in Section 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 also Nolfi, 2005).

The selective advantage of signaling over non-signaling groups, as detailed in Section 4, is the reason why we observe the evolution of signaling groups. Moreover, it is the social function of signaling (i.e., 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. Ethologists (see, for example, Tinbergen, 1964) considered the existence of cues (or derived activities) as precursors of signals and their subsequent ritualization into signals crucial notions in an effort to explain the evolution of communication. 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 indi-

vidual 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.

In Section 4.2, 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, Smith, and Kacelnik (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 reshaped (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.

In this article, we have demonstrated the portability of time-dependent decision-making mechanisms onto real robots. Even though the controllers were evolved in a simulated world and the simulation did not go as far as implementing possible inter-robot differences (see Section 3.3), the system was always successful and we did not observe any mistakes in categorizing the environment. Our results also show that the use of communication was particularly beneficial in the real world, as the inter-robot differences did in fact severely disrupt the individual decision-making mechanisms of certain agents (see, for example, *s-bot₅* in the four-robot generalization test, which would need, if left alone, more than two loops around the light to initiate antiphototaxis). However, it is easy to imagine a case where a robot takes the wrong decision about the state of the world, and initiates antiphototaxis emitting a tone in *Env A*. This would cause the collapse of the whole system, as all robots would perform the wrong action in that environment, even if their individual discrimination mechanism would have produced a correct categorization. This event was never observed in reality (see Section 3.3), sug-

gesting that the evolved behavior is also very robust against inter-robot differences.

In parallel with studying the effectiveness of the evolved signaling mechanisms when tested on real robots, we found that signaling evolved even in the absence of explicit selected pressure coded in the fitness function. In fact, our analysis revealed a “hidden” benefit for communication. This raises the following issue: should we always equip our robots with “ears” and a “mouth” to make possible the switch from solitary to social behavior, even in cases when the benefit of communication in such a system is obscure to us, the experimenter? The work presented in this article is an example in which communication proves to be beneficial and the ER methodology manages to discover ways to use it by linking it to the rest of the robot behavior and in the process enhancing the robustness of the system. That is, evolution found an efficient way to use these hardware tools (i.e., “ears” and a “mouth”), and it is our belief that this might also be the case with other more complex tasks. We cannot go so far as to claim that incorporating long-range communication devices should be standard, but we can say that evolution seems to discover efficient ways to use them. In other words, by using ER within the context of collective or swarm robotics, such hardware tools for long-range signaling might end up being beneficial for the group’s performance on a certain task, even though communication might appear pointless to the experimenter at the time of defining the building blocks of the behavior.

Acknowledgments

The authors would like to thank Dr. Stefano Nolfi, Dr. Francisco Santos, and Dr. Mauro Birattari for comments on the work at an earlier stage, the three anonymous reviewers for their constructive comments on the initial version of this article, Rehan O’ Grady for assistance in proofreading, and Dr. Jason Noble for assistance in producing the final version.

E. Tuci and M. Dorigo acknowledge European Commission support via the ECAgents project, funded by the Future and Emerging Technologies program (grant IST-1940). M. Dorigo acknowledges support from the Belgian FNRS, of which he is a Research Director. M. Dorigo and C. Ampatzis acknowledge support from the “ANTS” project, an “Action de Recherche Concertée” funded by the Scientific Research Directorate of the French Community of Belgium. The information provided is the sole responsibility of the authors and

does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

Notes

- 1 A project funded by the Future and Emerging Technologies Program (IST-FET) of the European Commission, under grant IST-2000-31010. See also <http://www.swarm-bots.org>.
- 2 See <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 The movies that correspond to these experiments can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2006-007>.
- 4 The movies that correspond to these experiments and a more detailed description can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2006-007>.
- 5 The light sensors of *s-bot₅* are the reason for this behavior. In fact, their readings proved to be significantly different from those of the other robots. By comparing the behavior of this robot with *s-bot₁*, we can get an idea of the magnitude of the inter-robot differences.
- 6 We do not further analyze the cases in which the robots signal throughout the entire duration of a trial as we consider it obvious that in these cases the sound-emitting behavior does not serve any specific function.
- 7 See <http://iridia.ulb.ac.be/supp/IridiaSupp2006-007> for data complementing the statistics of Table 4 with results in *Env A* and for data supporting our claim that signaling groups tend to initiate antiphototaxis later than non-signaling groups.

References

- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1, 91–122.
- Blynel, J., & Floreano, D. (2003). Exploring the t-maze: Evolving learning-like robot behaviors using CTRNNs. In G. R. Raidl, J.-A. Meyer, M. Middendorf, S. Cagnoni, J. J. R. Cardalda, D. Corne, J. Gottlieb, A. Guillot, E. Hart, C. G. Johnson, & E. Marchiori (Eds.), *Applications of Evolutionary Computing, EvoWorkshops2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, EvoSTIM, Lecture Notes in Computer Science*, Vol. 2611 (pp. 598–609). Berlin: Springer.
- Brooks, R. A. (1992). Artificial life and real robots. In F. J. Varela, & P. Bourguine (Eds.), *Towards a practice of autonomous systems: Proceedings of the 1st European Conference on Artificial Life* (pp. 3–10). Cambridge, MA: MIT Press.
- Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.
- Dudek, G., & Jenkin, M. (2000). *Computational principles of mobile robotics*. Cambridge, UK: 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.
- Floreano, D., Mitri, S., Magnenat, S., & Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17, 514–519.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley.
- Groß, R., Tuci, E., Dorigo, M., Bonani, M., & Mondada, F. (2006). Object transport by modular robots that self-assemble. In *Proceedings of the 2006 IEEE International Conference on Robotics and Automation* (pp. 2558–2564). Los Alamitos, CA: IEEE Computer Society Press.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7, 1–52.
- Hasson, O. (1994). Cheating signals. *Journal of Theoretical Biology*, 167, 223–238.
- Hauser, M. D. (1997). *The evolution of communication*. Cambridge, MA: MIT Press/BradfordBooks.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6, 325–368.
- Matarić, M. J., & Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19, 67–83.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge: Cambridge University Press.
- 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, 417–434.
- 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, 193–221.
- 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. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.

- O'Grady, R., Groß, R., Mondada, F., Bonani, M., & Dorigo, M. (2005). Self-assembly on demand in a group of physical autonomous mobile robots navigating rough terrain. In M. S. Capcarrere, A. A. Freitas, P. J. Bentley, C. G. Johnson, & J. Timmis (Eds.), *Advances in artificial life: Proceedings of the 8th European Conference (ECAL 2005), Lecture Notes in Computer Science*, Vol. 3630 (pp. 272–281). Berlin: Springer.
- Paine, R. W., & Tani, J. (2005). How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior*, *13*, 211–225.
- Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2002). Evolving teamwork and role allocation with real robots. In K. Russell, M. A. B. Standish, & H. A. Abbass (Eds.), *Artificial life VII: Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior (SAB'02)* (pp. 302–311). Cambridge, MA: MIT Press.
- Sherman, P. W. (1977). Nepotism and the evolution of alarm calls. *Science*, *197*, 1246–1253.
- Struhsaker, T. T. (1967). Auditory communication among vervet monkeys (*Cercopithecus aethiops*). In S. A. Altmann (Ed.), *Social communication among primates*. Chicago: University of Chicago 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.
- Trianni, V., & Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, *95*, 213–231.
- Tuci, E., Ampatzis, C., & Dorigo, M. (2005). Evolving neural mechanisms for an iterated discrimination task: A robot based model. In M. S. Capcarrere, A. A. Freitas, P. J. Bentley, C. G. Johnson, & J. Timmis (Eds.), *Advances in artificial life: Proceedings of the 8th European Conference (ECAL 2005), Lecture Notes in Computer Science*, Vol. 3630 (pp. 231–240). Berlin: Springer.
- Tuci, E., Groß, R., Trianni, V., Mondada, F., Bonani, M., & Dorigo, M. (2006). Cooperation through self-assembly in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, *1*, 115–150.
- Tuci, E., Trianni, V., & Dorigo, M. (2004). “Feeling” the flow of time through sensory-motor coordination. *Connection Science*, *16*, 301–324.
- Urzelai, J., & Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evolutionary Computation*, *9*, 495–524.
- von Frisch, K. (1967). *The dance language and orientation of bees*. Cambridge, MA: Harvard University Press.
- Williams, G. C. (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.
- Ziemke, T., & Thieme, M. (2002). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, *10*, 185–199.

About the Authors



Christos Ampatzis received a degree from the Department of Electrical Engineering and Computer Engineering from the Aristotle University of Thessaloniki, Greece in 2001, a degree of Advanced Studies – Master in Artificial Intelligence from the Katholieke Universiteit Leuven, Belgium in 2003, and a Diplôme d'Études Approfondies in applied sciences from the Université Libre de Bruxelles, Brussels, Belgium in 2004. Currently, he is a Ph.D. student at IRIDIA, Université Libre de Bruxelles, Brussels, Belgium. His research interests are in the field of evolutionary robotics, artificial life and swarm intelligence.



Elio Tuci has been a researcher at IRIDIA since November 2003. In February 2004, he received a D.Phil. in computer science and artificial intelligence from the University of Sussex (UK). His research interests are within the field of artificial life and evolutionary robotics. *E-mail:* etuci@ulb.ac.be



Vito Trianni received a Laurea (Master of Technology) degree in computer science engineering from the Politecnico di Milano, Milan, Italy, in 2000, a Master's in information technology from the ICT Center of Excellence for Research, Innovation, Education and Industrial Labs partnership (CEFRIEL), Milan, in 2001, and a Diplôme d'Études Approfondies and Ph.D. degree in applied sciences from the Université Libre de Bruxelles, Brussels, Belgium, in 2003 and 2006, respectively. Currently, he is a Research Fellow with the Laboratory of Autonomous Robotics and Artificial Life, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche (LARAL-ISTC-CNR), Rome, Italy. His research interests are in the fields of evolutionary robotics, swarm intelligence, and self-organization. *E-mail:* vito.trianni@istc.cnr.it



Marco Dorigo is a research director of the FNRS, the Belgian National Funds for Scientific Research, and of IRIDIA, the artificial intelligence laboratory of the Université Libre de Bruxelles. He is the inventor of the ant colony optimization metaheuristic. His current research interests include swarm intelligence, swarm robotics, and metaheuristics for discrete optimization. Dr Dorigo is the Editor-in-Chief of the journal *Swarm Intelligence*. He is an Associate Editor for *IEEE Transactions on Evolutionary Computation*, *IEEE Transactions on Systems, Man, and Cybernetics*, and *ACM Transactions on Autonomous and Adaptive Systems*, and he is a member of the editorial boards of numerous international journals. In 1996 he was awarded the Italian Prize for Artificial Intelligence, in 2003 the Marie Curie Excellence Award, and in 2005 the Dr A. De Leeuw-Damry-Bourlart award in applied sciences. *E-mail:* mdorigo@ulb.ac.be