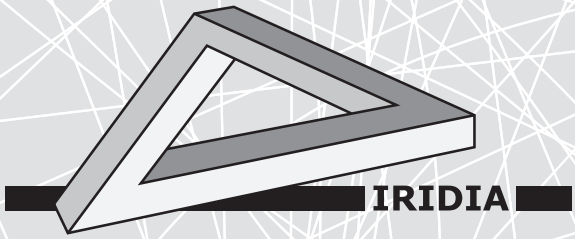

Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Evolved homogeneous neuro-controllers for robots with different sensory capabilities: coordinated motion and cooperation

Elio Tuci, Christos Ampatzis, Federico Vicentini, and Marco Dorigo

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2006-015

May 2006

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle*
Université Libre de Bruxelles
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2006-015

Revision history:

TR/IRIDIA/2006-015.001 May 2006
TR/IRIDIA/2006-015.002 July 2006

Operational Aspects of the Evolved Signalling Behaviour in a Group of Robots with Different Sensory Capabilities: Coordinated Motion and Cooperation

Elio Tuci[‡] etuci@ulb.ac.be
Christos Ampatzis[‡] campatzi@ulb.ac.be
Federico Vicentini[†] federico.vicentini@polimi.it
Marco Dorigo[‡] mdorigo@ulb.ac.be

‡ IRIDIA, CoDE, Université Libre de Bruxelles, Bruxelles, Belgium
† Robotics Lab, Mechanics Dept., Politecnico di Milano, Milano, Italy

May 2006

Abstract

This paper tackles the issue of designing homogeneous neuro-controllers with artificial evolution in order to control groups of robots that differ in terms of sensory capabilites. In order to accomplish a common goal, the agents have to complement the partial "view" they have of the environment. The results obtained prove that the agents are capable of cooperating and coordinating their actions in order to carry out a navigation task. The paper also describes operational aspects of the evolved behaviour. In particular, analyses of the sound signalling behaviour show that the robots employ the sound to remain close to each other at a safe distance with respect to the risk of collisions. Spatial discrimination of the sound sources is achieved by exploiting a rotational movement which amplifies intensity differences between the two sound sensors.

1 Introduction

In recent years, various types of agent-based simulation models have been employed to look at issues concerning communication in natural organisms and human language which can hardly be investigated with classic analytical models (see Cangelosi and Parisi, 2002; Vogt, 2005). With respect to analytical and other simulation models, agent-based models do not require the designer to make strong assumptions about the essential features on which social interactions are based—e.g, assumptions concerning what communication is and about the requirement of individual competences in the domain of categorisation and naming. This is particularly true in models in which evolutionary computation algorithms are used to design artificial neural networks as agent's controllers. These models are generally referred to as Evolutionary Robotics models (ER, see Nolfi and Floreano, 2000). Roughly speaking, ER is a methodological tool to automate the design of robots' controllers. ER is based on the use of artificial evolution to find sets of parameters for artificial neural networks that guide the robots to the accomplishment of their objective, avoiding dangers. Owing to its properties, ER can be employed to look at the effects that the physical interactions among embodied agents and their world have on the evolution of individual behaviour and social skills.

In the recent past, ER has been used in the context of social behaviour to investigate issues concerning the evolution of communication in groups of agents required to solve tasks that demanded coordination and cooperation (see Di Paolo, 2000; Quinn, 2001; Baldassarre et al., 2003;

Quinn et al., 2003; Marocco and Nolfi, 2005, 2006). The results of these studies seem to suggest that ER models are a valuable tool to study how semantics and syntax originate from the evolutionary and ontogenetic history of populations of autonomous agents (Quinn, 2001; Cangelosi, 2005). In other words, ER models can be used to look at issues concerning causal relationship between the evolution and the development of perceptual, cognitive and motor capabilities and the emergence of a communicative system and possibly language in a population of artificial and natural organisms. Following this line of investigation, we are interested in further exploring the evolution of social skills. In particular, we focus on a context in which a group of agents with different sensory capabilities are required to share their "knowledge" of the world to accomplish a common task. We consider the following experiment: three robots are placed in an arena, as shown in Figure 1. The arena is composed of walls and a light that is always turned on. The light can be situated at the bottom left corridor (*Env. L*) or at the bottom right corridor (*Env. R*). The robots are initialised with their centre anywhere on an imaginary circle of radius 12 cm centred in the middle of the top corridor, at a minimum distance of 3 cm from each other. Their initial orientation is always pointing towards the centroid of the group. The goal of the robots is (i) to navigate towards the light whose position changes according to the type of environment they are situated in, (ii) to avoid collisions.

The peculiarity of the task lies in the fact that the robots are equipped with different sets of sensors. In particular, two robots are equipped with infrared and sound sensors but they have no ambient light sensors. These robots are referred to as R_{IR} (see Figure 2a). The other robot is equipped with ambient light and sound sensors but it has no infrared sensors. We refer to this robot as R_{AL} (see Figure 2b). Robots R_{IR} can perceive the walls and other agents through infrared sensors, while the robot R_{AL} can perceive the light. Therefore, given the nature of the task, the robots are forced to cooperate in order to accomplish their goal. In principle, it would be very hard for each of them to solve the task solely based on their own perception of the world. R_{AL} can hardly avoid collisions; R_{IR} can hardly find the light source. Thus, the task requires cooperation and coordination of actions between the different types of robots. Notice that the reason why we chose the group to be composed of two R_{IR} and one R_{AL} robot is that this intuitively seems to be the smallest group capable of spatially arranging itself adaptively in order to successfully navigate the world. Although the robots differ with respect to their sensory capabilities, they are homogeneous with respect to their controllers. That is, the same controller, synthesised by artificial evolution, is cloned in each member of the group. Both types of robots are equipped with a sound signalling system (more details in Section 2). However, contrary to other studies (see Marocco and Nolfi, 2005, 2006; Baldassarre et al., 2003), we do not assume that the agents are capable of distinguishing their own sound from that of the other agents. The sound broadcasted into the environment is perceived by the agent through omnidirectional microphones. Therefore, acoustic signalling is subject to problems such as the distinction between own sound from those of others and the mutual interference due to lack of turn-taking (see Di Paolo, 2000).

The results of our study show that a quite robust and effective phototactic strategy evolves in spite of each of the agents being deprived of essential elements to accomplish the task. The successful strategies are based on cooperation and coordination of actions among the agents.

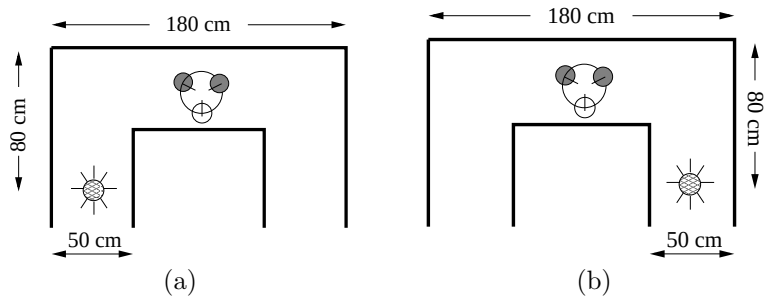

Figure 1: (a) *Env. L*; (b) *Env. R*. See text in Section 1 for details.

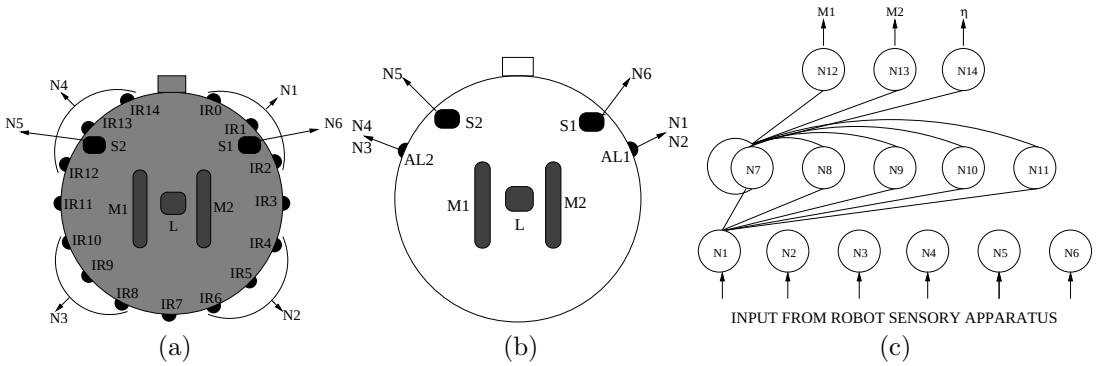

(a) (b) (c)

Figure 2: (a) The simulated robots R_{IR}; (b) The simulated robots R_{AL}; (c) the network architecture. Only the connections for one neuron of each layer are drawn. The input layer of R_{IR} takes readings as follows: neuron N_1 takes input from the infrared sensors $\frac{IR_0+IR_1+IR_2}{3}$, N_2 from $\frac{IR_4+IR_5+IR_6}{3}$, N_3 from $\frac{IR_8+IR_9+IR_{10}}{3}$, N_4 from $\frac{IR_{12}+IR_{13}+IR_{14}}{3}$, N_5 from sound sensor S_2, and N_6 from sound sensor S_1. The input layer of R_{AL} takes readings as follows: N_1 and N_2 take input from ambient light sensors AL_1, N_3 and N_4 take input from AL_2, N_5 from S_2, and N_6 from S_1. M_1 and M_2 are respectively the left and right motor. L is the loud-speaker (i.e., the sound organ).

The mutual coordination results particularly striking, so that, as already emphasised in a similar model (Di Paolo, 2000), it turns out to be very hard to speak in terms of causality. For example, (a) phototaxis is induced in the group by robot R_{AL}, but this behaviour seems to be effectively displayed by the robot R_{AL} only if it is situated in a social context—i.e., surrounded by robots R_{IR}; (b) angular movement introduces rhythm in acoustic perception, which *per se*, is not sufficient to coordinate the movements of the group. However, coordinated actions come about by the fusion of perception of sound and patterns in infrared proximity sensors. In conclusion, from these simulations, we learn something about the relationship between individual and social skills, and the potentiality of the system which can be further exploited to study the evolution of more complex forms of social interactions in similar circumstances (e.g., groups of morphologically heterogenous robots).

2 The simulated agents

The controllers are evolved in a simulation environment which models some of the hardware characteristics of the real *s-bots*. The *s-bots* are small wheeled cylindrical robots, 5.8 cm of radius, equipped with a variety of sensors, and whose mobility is ensured by a differential drive system (see Mondada et al., 2004, for details). Robot R_{IR} makes use of 12 out of 15 infrared sensors (IR_i) of an *s-bot*, while robot R_{AL} uses the ambient light sensors (AL_1) and (AL_2) positioned at $\pm 67.5°$ with respect to the orientation of the robot (see Figure 2a and 2b). The signal of the infrared sensor is a function of the distance between the robot and the obstacle. Light sensor values are simulated through a sampling technique (see Miglino et al., 1995). All robots are equipped with a loud-speaker (L) that is situated in the centre of the body of the robot, and with two omnidirectional microphones (S_1 and S_2), placed at $\pm 45°$ with respect to the robot's heading. Sound is modelled as an instantaneous, additive field of single frequency with time-varying intensity ($\eta_i \in [0.0, 1.0]$) which decreases with the square of the distance from the source, as previously modelled in (Di Paolo, 2000). Sound intensity is regulated by the firing rate of neuron $N14$ (see Section 3 for details). Robots can perceive signals emitted by themselves and by other agents. The modelling of the perception of sound is inspired by what described in (Di Paolo, 2000). There is no attenuation of intensity for self-produced signal. The perception of sound emitted by others is affected by a "self-shadowing" mechanism which is modelled as a linear attenuation without refraction, proportional to the distance (δ_{sh}) travelled by the signal within the body of

the receiver (see Di Paolo, 2000, for details). This distance is computed as follows:

$$\delta_{sh} = \delta_{sen}(1 - A), \qquad 0 \leq A < 1, \qquad A = \frac{\delta^2 - R^2}{{\delta_{sen}}^2} \tag{1}$$

where δ_{sen} is the distance between the sound source and the sensor and δ is the distance between the sound source and the centre of the body of the receiver, and R is the robot's radius (see also Figure 3). The "self" component of the sound signal is simply equal to η_i. In order to calculate the "non-self" component, we firstly scale the intensity of sound emitted by the sender (η_j) by applying the inverse square law with respect to the distance between the sound source and the microphones of the receiver. Subsequently, we multiply the scaled intensity with an attenuation factor ψ which ranges linearly from 1 when $\delta_{sh} = 0$ to 0.1 when $\delta_{sh} = 2R$. To summarise, the reading $\hat{S}_{is}$ of each sound sensor s of robot i is computed as follows:

$$\begin{aligned} &\text{self} = \eta_i \\ \hat{S}_{is} = \text{self} + \text{non-self}; \qquad &\text{non-self} = \sum_{\substack{j \in [1,3] \\ j \neq i}} \eta_j \frac{R^2}{\delta_{sen}^2} \psi \end{aligned} \tag{2}$$

The auditory receptive field of each microphone is bounded within the following interval $\hat{S}_{is} \in [0.0, 1.0]$. Therefore, the sound receptor can be saturated by the "self" emitted sound in case a robot emits at its highest intensity ($\eta_i = 1.0$). If the sound sensor is saturated by the "self" component, it is not possible for this robot to perceive sound signals emitted by others. Concerning the function that updates the position of the robots within the environment, we employed the Differential Drive Kinematics equations, as presented in (Dudek and Jenkin, 2000). 10% uniform noise was added to all sensor readings, the motor outputs and the position of the robot. The characteristics of the agent-environment model are explained in detail in (Vicentini and Tuci, 2006).

3 The controller and the evolutionary algorithm

The agent controller is composed of a network of five inter-neurons and an arrangement of six sensory neurons and three output neurons (see Fig. 2c). The sensory neurons receive input from the

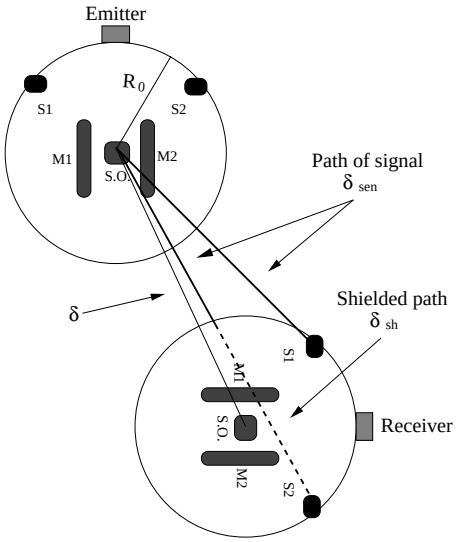

Figure 3: This picture has been adapted from (Di Paolo, 2000). It shows the working principles of the self-shadowing mechanism.

agent sensory apparatus. Thus, for robots R_{IR}, the network receives the readings from the infrared and sound sensors. For robots R_{AL}, the network receives the readings from the ambient-light and sound sensors. The inter-neuron network (from N_7 to N_{11}) is fully connected. Additionally, each inter-neuron receives one incoming synapse from each sensory neuron. Each output neuron (from N_{12} to N_{14}) receives one incoming synapse from each inter-neuron. There are no direct connections between sensory and output neurons. The network neurons are governed by the following state equation:

$$\frac{dy_i}{dt} = \begin{cases} \frac{1}{\tau_i}(-y_i + gI_i) & i \in [1,6] \\ \frac{1}{\tau_i}\left(-y_i + \sum_{j=1}^{k} \omega_{ji}\sigma(y_j + \beta_j) + gI_i\right) & i \in [7,14];\ \sigma(x) = \frac{1}{1+e^{-x}} \end{cases} \tag{3}$$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, g is a gain factor, I_i the intensity of the sensory perturbation on sensory neuron i, ω_{ji} the strength of the synaptic connection from neuron j to neuron i, β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate. The cell potentials y_i of the 12^{th} and the 13^{th} neuron, mapped into [0.0,1.0] by a sigmoid function σ and then linearly scaled into $[-6.5, 6.5]$, set the robot motors output. The cell potential y_i of the 14^{th} neuron, mapped into $[0.0, 1.0]$ by a sigmoid function σ, is used by the robot r to control the intensity of the sound emitted η_r. The following parameters are genetically encoded: (i) the strength of synaptic connections ω_{ji}; (ii) the decay constant τ_i of the inter-neurons and of neuron N_{14}; (iii) the bias term β_i of the sensory neurons, of the inter-neurons, and of the neuron N_{14}. The decay constant τ_i of the sensory neurons and of the output neurons N_{12} and N_{13} are set to 0.1. Cell potentials are set to 0 any time the network is initialised or reset, and circuits are integrated using the forward Euler method with an integration step-size of $dt = 0.1$.

A simple generational genetic algorithm is employed to set the parameters of the networks Goldberg (1989). The population contains 80 genotypes. Generations following the first one 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 individuals of the old population. Each genotype is a vector comprising 84 real values (i.e., 70 connection weights, 6 decay constants, 7 bias terms, and a gain factor). Initially, a random population of vectors is generated by initialising each component of each genotype to values chosen uniformly random from the range [0,1]. New genotypes, except "the elite", are produced by applying recombination with a probability of 0.3 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 to remain within the range [0,1]. Genotype parameters are linearly mapped to produce network parameters with the following ranges: biases $\beta_i \in [-4, -2]$ with $i \in [1, 6]$, biases $\beta_i \in [-5, 5]$ with $i \in [7, 14]$; weights $\omega_{ij} \in [-6, 6]$ with $i \in [1, 6]$ and $j \in [7, 11]$, weights $\omega_{ij} \in [-10, 10]$ with $i \in [7, 11]$ and $j \in [7, 14]$; gain factor $g \in [1, 13]$. Decay constants are firstly linearly mapped into the range $[-1.0, 1.3]$ and then exponentially mapped into $\tau_i \in [10^{-1.0}, 10^{1.3}]$. The lower bound of τ_i corresponds to the integration step-size used to update the controller; the upper bound, arbitrarily chosen, corresponds to about 1/20 of the maximum length of a trial (i.e., 400 s).

4 The fitness function

During evolution, each genotype is translated into a robot controller, and cloned in each agent. Then, the group is evaluated twelve times, six trials in *Env. L*, and six trials in *Env. R*. The sequence order of environments within the twelve trials has no bearing on the overall performance of the group since each robot controller is reset at the beginning of each trial. Each trial (e) differs from the others in the initialisation of the random number generator, which influences the robots'

starting position and orientation, and the noise added to motors and sensors. Within a trial, the robot life-span is 400 simulated seconds (4000 simulation cycles). In each trial, the group is rewarded by an evaluation function f_e which seeks to assess the ability of the team to approach the light bulb, while avoiding collisions and staying within the range of the robots' infrared sensors.[1].

$$f_e = KP\Big(\sum_{t=i}^{T}[(d_t - D_{t-1})(tanh(S_t/R))]\Big);$$

As in Quinn et al. (2003), the simulation time steps are indexed by t and T is the index of the final time step of the trial. d_t is the Euclidean distance between the group location at time step t and its location at time step $t = 0$, and D_{t-1} is the largest value that d_t has attained prior to time step t. Therefore, the term $(d_t - D_{t-1})$ measures any gain that the team has made on its previous best distance from its initial location which is taken to be the centroid of the group.

$tanh(S_t/\rho)$ reduces any fitness increment given by $(d_t - D_{t-1})$ when one or more robots are outside of the infrared sensor range. S_t is a measure of the team's dispersal beyond the infrared sensor range R (R = 24.6 cm) at time step t. Recall that robot R_{AL} has no infrared sensors. Therefore, it does not have a direct feedback at each time-step of its distance from its group-mates. Nevertheless, the sound can be indirectly used by this robot to adjust its position within the group. If each robot is within R range of at least another, then $S_t = 0$. Otherwise, the two shortest lines that can connect all three robots are found and S_t is the distance by which the longest of these exceeds R. $tanh()$ assures that, as the robots begin to disperse, the team's score increment falls sharply.

$P = 1-(\sum_{i=1}^{3} c_i/c_{max})$ if $\sum_{i=1}^{3} c_i \leq c_{max}$ reduces the score in proportion to the number of collisions which have occurred during the trial. c_i is the number of collisions of the robot i and $c_{max} = 4$ is the maximum number of collisions allowed. $P = 0$ if $\sum_{i=1}^{3} c_i > c_{max}$. The team's accumulated score is multiplied by $K = 3.0$ if the group moved towards the light bulb, otherwise $K = 1.0$. Note that a trial was terminated early if (a) the team reached the light bulb (b) the team distance from the light bulb exceeded an arbitrary limit set to 150 cm, or (c) the team exceeded the maximum number of allowed collisions c_{max}.

5 Results

Ten evolutionary simulations, each using a different random initialisation, were run for between 2500 and 3600 generations of the evolutionary algorithm. The termination criterion for each run was set to a time equal to 86400 seconds of CPU time. Experiments were performed on a cluster of 32 nodes, each with 2 AMD Opteron244^{TM} CPU running GNU/Linux Debian 3.0 OS. 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 general, the best fitness scores take advantages of favourable conditions, which are determined by the existence of between-generation variation in starting position and orientation and other simulation parameters. In order to have a better estimate of the behavioural capabilities of the evolved controllers, we post-evaluate, for each run, the genotype with the highest fitness. The entire set of post-evaluations should establish whether a group of robots is capable of reaching the light in *Env. L* and *Env. R*. In particular, the robots of a successful group should be capable of coordinating their movement and of cooperating, in order to approach the light bulb without colliding with each other or with the walls. A trial is successfully terminated when the centroid of the group is less than 10 cm away from the light bulb. During post-evaluation, each of the best ten evolved controllers is subject to a set of 1200 trials in both environments. The number of post-evaluation trials per type of environment (i.e., 1200) is given by systematically varying the initial positions of the three robots according to the

[1]Note that, this fitness function is very similar to the one used in Quinn et al. (2003) from which it mainly differs for the parameter K. This parameter has been introduced to give a selective advantage to those groups which move towards the light bulb. In order to facilitate comparisons between our work and that detailed in Quinn et al. (2003), we provide a description of the fitness function which uses a similar mathematical notation employed in Quinn et al. (2003)

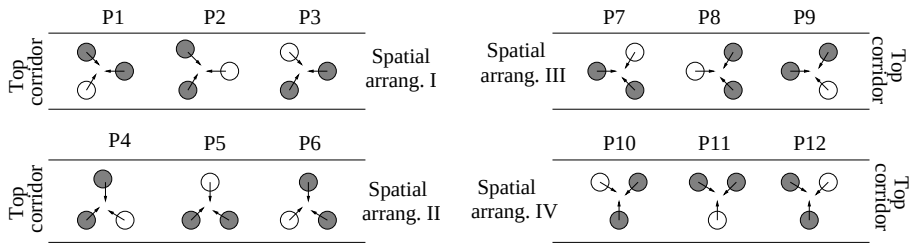

Figure 4: The robots' initial positions (from P1 to P12) during the post-evaluation phase. White circles refers to robot R_{AL}, grey circles refers to robot R_{IR}. See text in Section 5 for details.

following criteria: (i) we defined four different types of spatial arrangements in which the robots are placed at the vertices of an imaginary equilateral triangle inscribed in a circle of radius 12 cm and centred in the middle of the top corridor (see Figure 4); (ii) for each spatial arrangement, we identified three possible relative positions of the robot R_{AL} with respect to the walls of the corridor (see white circle in Figure 4); (iii) for each of these (four times three) initial positions, the post-evaluation is repeated one hundred times. The initial orientation of each robot is determined by applying an angular displacement randomly chosen in the interval $[-30°, 30°]$ with respect to a vector originating from the centre of the robot and pointing towards the centroid of the group. The four times three different arrangements take into account a set of relative positions among the robots and between the robots and the walls so that the success rate of the group is not biased by these elements. During post-evaluation, the robot life-span is more than twice longer than during evolution (i.e., 1000 s, 10000 simulation cycles). This should give the robots enough time to compensate for possible disruptive effects induced by initial positions never or very rarely experienced during evolution. At the beginning of each post-evaluation trial, the controllers are reset (see Section 3 for details). The results of the post-evaluation phase are shown in Figure 5. We notice that the best controllers are those produced by run n. 9, and 10 achieving a performance over 90% in *Env. L* and *Env. R*. Runs n. 4, and 7 display a performance over 80% in both environments. The performance of all the other genotypes is clearly unsatisfactory. Run 2, 3, 5, and 8 proved to be capable of accomplishing the task only when located in an *Env. R*, while run n. 1 is particularly effective in *Env. L*. This phenomenon can be explained by considering that the two environments require two different types of turn—a left turn in *Env. L*, and a right turn in *Env. R*. By looking at the behaviour of the groups through a simple graphical interface, we observed that, first, the successful groups employ two different navigation strategies to make the turns (see also Section 5.2). Second, those groups that systematically fail in either environment, lack the

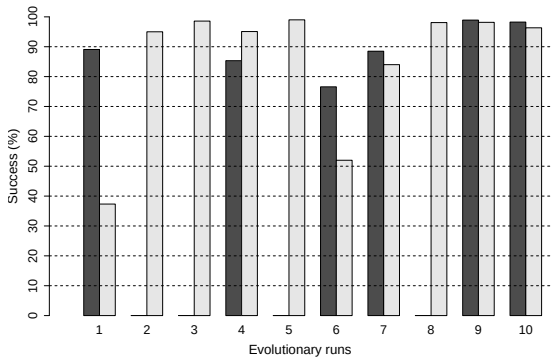

Figure 5: Results of post-evaluation showing for the best evolved controllers of each run over 1200 trials per type of environment the percentage of success (black bars refers to *Env. L*, and white bars to *Env. R*).

capability to make both turns. Note that when looking at the performances of the best evolved controllers, as shown in Figure 5, one has to take into account the arbitrary criteria we chose to determine whether or not a group of robots is successful in any given trial. We should recall that, in order to be successful, no robot has to collide with the walls or with the other robots. This is a very strict condition, which, given the nature of the task, demands each agent to be very accurate in coordinating its movement. Further post-evaluation tests proved that, if we allow the group to make a certain number of collisions (i.e., four collisions) before defining a trial as a failure, then several controllers would result almost always successful in both types of environment—the data of these post-evaluation tests are not shown. Whether or not the robots should be allowed to collide or the extent to which a single collision invalidates the performance of the group, are issues that extend beyond the interest of this paper, and shall not be discussed any further. Instead, we focus on other performance measures which tell us more about the characteristics of the best evolved controllers. For instance, by looking at the data shown in Table 1, we notice that, except for run n. 2, the majority of the failures in *Env. L*, are due to collisions. In *Env. R*, the performance of all the runs, are sensibly better than those in *Env. L* (see columns 4 and 5, Table 1). If we look at the average distances to the light (see columns 6 and 7, Table 1) and the relative standard deviations (see columns 8 and 9, Table 1), we can see that in *Env. L*, failures happen rather far away from the light. For example, for runs n. 3, 5, and 8—100% unsuccessful in *Env. L*—the final distance to the light is almost equal to the initial distance. This seems to denote a lack of coordination of movement during the initial phase, when the robots have to assume a configuration which favours the group phototaxis. In *Env. R*, the smaller final distances to the light seem to denote a problem, possibly common to several runs, in making the right turn.

In the rest of this section, we concentrate on the analysis of the controller run n. 9, which proved to be the most effective at the first post-evaluation test. In section 5.1, we show the results of further post-evaluation tests that aim to verify whether the group relies on a genuine cooperation based on a mutual dependence between each member's actions and the behaviour of the others. In Section 5.2, we describe more in details the navigational strategies employed by the robots in each environment. In Sections 5.3 and 5.4, we provide an operational description of the group behaviour. These analyses help the reader to understand more about the mechanisms used by the robots to coordinate their actions and to complement the partial view that each of them has of the world. In Section 5.5 we test the robustness of the best evolved strategy with respect

Table 1: Further results of the post-evaluation test, showing for the best evolved controllers of each run: (i) the percentage of unsuccessful trials due to exceeded time limit without the group having reached the target (columns 2, and 3); (ii) the percentage of unsuccessful trials which terminated due to collisions (columns 4, and 5); (iii) the average and standard deviation of the final distance of the centroid of the group to the light during the unsuccessful trials (respectively columns 6, 8 for *Env. L*, and columns 7, 9 for *Env. R*). Note that in all trials the initial distance between the centroid of the group and the light is equal to 85.14 cm.

	(%) of failure due to time limit		**(%) of failure due to collisions**		**Distance to the light avg**		**std**	
run	***Env. L***	***Env. R***	***Env. L***	***Env. R***	***Env. L***	***Env. R***	***Env. L***	***Env. R***
n. 1	0.00	52.75	10.92	9.92	82.19	52.17	6.63	32.74
n. 2	85.33	1.83	14.67	3.17	66.30	46.17	4.36	18.22
n. 3	0.00	1.00	100.00	0.42	81.02	36.38	4.049	12.02
n. 4	0.67	0.50	14.00	4.50	57.83	69.30	13.50	15.32
n. 5	0.00	0.00	100.00	1.00	79.05	41.13	2.94	12.93
n. 6	0.00	31.00	23.42	17.08	77.50	64.27	11.89	29.92
n. 7	0.58	10.00	10.92	6.00	50.98	40.80	30.18	14.18
n. 8	0.00	0.00	100.00	1.92	80.94	53.03	2.34	11.59
n. 9	0.00	0.83	1.08	1.00	77.71	50.22	11.44	21.08
n. 10	0.00	2.17	1.75	1.50	82.28	90.37	13.19	31.81

Table 2: Results of different post-evaluation tests for the best evolved controller of run n. 9. Columns 2, and 3 show the percentage of unsuccessful trials due to exceeded time limit without the group having reached the target. Columns 4, and 5 show the percentage of unsuccessful trials which terminated due to collisions. Columns 6, 7, 8, and 9 show the average and standard deviation of the final distance of the centroid of the group to the light during the unsuccessful trials (respectively columns 6, 8 for *Env. L*, and columns 7, 9 for *Env. R*). Note that in all trials the initial distance between the centroid of the group and the light is equal to 85.14 cm. See text in Section 5.1 for details.

	(%) of failure due to time limit		**(%) of failure due to collisions**		**Distance to the light avg**		**std**	
Test	***Env. L***	***Env. R***	***Env. L***	***Env. R***	***Env. L***	***Env. R***	***Env. L***	***Env. R***
A	100.00	100.00	0.00	0.00	85.19	85.49	5.47	5.48
B	0.00	0.00	100.00	99.92	74.42	49.66	7.71	7.91
C	2.83	2.42	20.50	97.58	86.09	65.63	21.88	19.14
D	0.00	10.00	3.42	2.75	86.61	40.65	5.00	19.11

to changes in the characteristics of the environment. In particular, we show the results of tests in which the team is placed in a boundless arena and in an arena with round obstacles.

5.1 Further post-evaluation tests

In this section, we illustrate the results of further post-evaluation tests which aim to establish whether the best group navigation strategy is based on genuine cooperation among the member of the group. Alternatively, the group may employ "fixed" phototactic movement which may work as well given that the dimensions of the corridors and the positions of the lights in the two worlds do not vary. In *Test A*, the best controller of run n. 9 is cloned in a group of three robots R_{IR}. Consequently, the robots have no means to know where the light is placed. As shown in Table 2, the group is 100% unsuccessful due to time limit exceeded without having reached the target (see columns 2 and 3, Table 2). Moreover, in both environments, the average distance between the centroid of the group and the light does not differ much from the initial distance (see columns 6 and 7, Table 2). The rather small standard deviation confirms that this group of robots seems not to make any significant movement away from its initial position (see columns 8 and 9, Table 2). Indeed, it seems to be the presence of a robot R_{AL}—missing in the group in this test—that triggers the movement and guides the group towards the target. Not surprisingly, the robots are very effective in avoiding collisions (see columns 4 and 5, Table 2).

In *tests B*, *C*, and *D* the best controller run n. 9 is cloned in a group of three robots, in which, as during evolution, two are R_{IR} and one R_{AL}. Contrary to the evolution, in *test B*, the robot R_{AL} only hears its own sound; in *test C*, the robots R_{IR} do not hear the sound emitted by robot R_{AL}; in *test D*, the robots R_{IR} perceive the sound emitted by the robot R_{AL} but they can not hear each other's sound.

The kind of disruption applied to the system in *test B*, makes robot R_{AL} completely isolated from the rest of the group, since it has no means to perceive robots R_{IR}. In fact, R_{AL} has no infrared sensors, and due to the disruption, it does not perceive sound emitted by the robots R_{IR}. As shown in Table 2, the group results almost 100% unsuccessful in both environments due to collisions (see columns 4 and 5, Table 2). We also observed that, it is always the robot R_{AL} that collides against the walls. The fact that in *Env. R*, the average final distance between the centroid of the group and the light is significantly smaller than the initial distance (i.e., 85.14 cm), suggests that, the instinctive phototactic behaviour of robot R_{AL} drags the other two robots toward the light. However, due to lack of sound communication among the two types of robots, robot R_{AL} ends up systematically colliding against one of the walls. The results of *tests A* and *B* prove that, the successful navigation strategy of the group is based on a mutual dependence between each member's actions and the behaviour of the others. These dependencies are the elements which

determine the group's coordination of movement.

Data in Table 2 show that, in *Test C*, robots do not systematically fail to reach the target. Although the performance in *Env. R* is severely disrupted with almost 100% of unsuccess rate, in *Env. L* the group performance is less touched by the alterations we applied to the system (see columns 2, 3, 4 and 5, Table 2). The failure in *Env. R*, is mostly due to collisions, which seem to occur rather far away from the lights. This suggest that, the group seems to have a problem in performing phototaxis, rather than in overcoming specific manoeuvres such as the right turn (see columns 6, and 8, Table 2). In summary, the sound received by the robots R_{IR} from robot R_{AL} seems to be more important during navigation in *Env. R* than in *Env. L*. This suggests that different group behavioural strategies are employed in the two environments. We will come back to this issue in Section 5.2.

In *Test D*, we immediately notice that the rate of failure is rather low (see columns 2, 3, 4 and 5, Table 2). The success rate turns out to be quite similar to that achieved in the evolutionary conditions in which all the robots can hear the sound emitted by all the others. It seems fair to conclude that (i) communication through sound signalling among the members of the group is required in order to successfully approach the target; (ii) successful strategies of controller run n. 9 are only marginally based on communication through sound signalling between the robots R_{IR}.

5.2 The group's behaviour

In this section we provide a qualitative description of the individual and group motion of the best evolved simulated agents (i.e., controller run n. 9) as observed through a simple graphical interface. First of all, we noticed that the systematic variation of the initial positions of the robots during post-evaluation brings about contingencies in which the coordination of movements of the group toward the target requires an initial effort of the robots in re-arranging their relative positions. During this initial phase of a trial a dynamic process guided by the nature of the flow of sensation induces the specialisation of the controllers with respect to the physical characteristics of the robots, and to the relative role that they play in the group. This phase is followed by the navigation phase in which the group seems to maintain a rather regular spatial configuration; that is, the two robots R_{IR} place themselves in between the target and the robot R_{AL}. However, note that while *Env. L* requires the group to make a left turn, *Env. R* requires the group to make a right turn. This asymmetry in the environmental structures corresponds to differences in behavioural strategies employed by the group to reach the target as shown in Figure 6. While

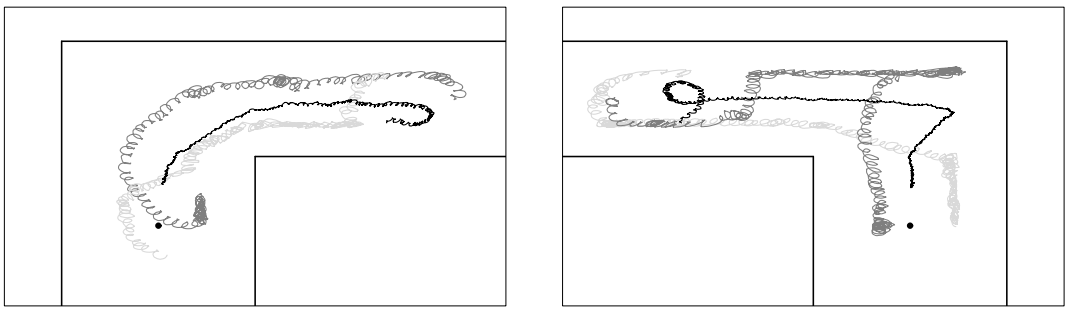

(a) (b)

Figure 6: Trajectories of the agents during a successful trial (a) in an *Env. L*, and (b) in an *Env. R*. The black lines refer to the trajectories of robot R_{AL} while the other lines refer to the trajectories of robots R_{IR}. The horizontal and vertical segments represent the walls. In each figure, we depict only the side of the corridor where the light—i.e., the small black dot—is located.

in *Env. L* the robots simply turn towards the light keeping their relative positions in the group, in *Env. R* we firstly observe an alignment of the agents along the far right wall (see Figure 6b). Subsequently, the agent close to the corner (see the dark gray line) overcomes the other two and the group starts moving towards the target once the classical configuration of the two robots R_{IR} in between the target and the robot R_{AL} is re-established.

Another important qualitative element is that each of the members of the group is characterised by a movement with a strong angular component (anti-clockwise). In other words, the robots proceed toward the light by rotating on the spot. Within a trial, pure linear movement replaces the rotational behaviour only sporadically and for a very short interval (see Figure 6a and b). This can happen to avoid an imminent danger of collision or if required by the navigational strategy of the group. The evolution of the rotational movement is not particularly surprising if we think about its effect on the perception of sound. In particular, the rotational movement may introduce rhythm in perception. The oscillations of perceived sound, produced by the rotational movement and/or by the oscillations manifested in signalling behaviour, may provide the robots the cues to adjust their positions relative to each other, since intensity differences between the two sound receptors—known as interaural level difference (ILD, Kandel et al., 2000)—can be a valuable cue for spatial discrimination of sound sources. Further and deeper investigations on the nature of sound signals and its relationship with the robots' motion will be carried out in the next section.

The effect of the starting position and the rotational movement are phenomena that have a strong effect on the time it takes to the group to reach the target. Indeed, as resulted from the post-evaluation test, most of the successful trials of the best evolved group of robots last longer than the 400 s given to the groups to complete the task during the evolutionary phase (data not shown).

5.3 A description of the signalling behaviour

Each robot of the group is required to coordinate its actions in order (i) to remain close to the other two agents without incurring into collisions, and (ii) to make actions which bring the group closer to the target. What is the role of signalling behaviour for the achievement of these goals? Is signalling behaviour used by robot R_{AL} to communicate to robots R_{IR} information concerning the relative position of the target? Similarly, is signalling behaviour used by robots R_{IR} to inform robot R_{AL} on the position of obstacles against which it may collide? In order to provide an answer to this type of questions, we carried out a series of tests that look at the properties of the sound signals perceived by each robot during a successful trial in each environment. Our goal is to identify oscillatory phenomena or other distinctive features in sound production/perception whose properties can be exploited by the robots to coordinate their actions. Despite the fact that our analysis is limited to two successful trials, one for each type of environment, we hope that the results help us to formulate "general" hypotheses concerning the role of signalling behaviour for the coordination of actions of the group.

Before proceeding further, we should remind the reader that the intensity of sound perceived at each microphone results from the summation of two components—the "self" and the "non-self"—and the noise. The "self" component (i.e., the agent's own signal) is only determined by the intensity of the sound emitted by the robot itself. The "non-self" component is determined by the intensity at which the sound is emitted from the loud-speaker of a sender as well as by the relative distance and orientation of the loud-speaker with respect to the receiver's microphones (see Section 2). Although the agents have no means to distinguish between the "self" and the "non-self" components of the perceived sound, they can act in a way to determine patterns in the flow of sensations which are informative on their spatial relationships.

We proceeded by separately recording the "self" and the "non-self" components of the sound perceived by each robot at each microphone, during a successful trial in each environment. Moreover, we recorded the "orientation signal" of each robot. The latter refers to how the orientation of each robot changes in time with respect to a global framework of reference. As shown in Section 5.2, each robot combines phototaxis with a rotational movement. The latter, due to the simulated physics, may introduce rhythms in the perception of sound through its effects on the

Table 3: Frequencies (Hz) of the main harmonic components fn_i in the orientation signal (columns 3, 4 and 5), and in the "non-self" component of the perceived sound signal at sensor S1 (columns 6) and at sensor S2 (columns 7) for each robot.

		orientation signal			**non-self**	
					S1	**S2**
Env.	**robot**	fn_1 **(Hz)**	fn_2 **(Hz)**	fn_3 **(Hz)**	fn_1 **(Hz)**	fn_1 **(Hz)**
Env. L	R^1_{IR}	0.187	0.375	0.566	0.185	0.185
	R^2_{IR}	0.187	0.377	0.556	0.181	0.181
	R_{AL}	0.205	0.411	0.617	0.205	0.205
Env. R	R^1_{IR}	0.186	0.378	0.563	0.184	0.184
	R^2_{IR}	0.184	0.358	0.559	0.181	0.181
	R_{AL}	0.212	0.426	0.633	0.201	0.201

"non-self" component as mentioned above. With a Fast Fourier Transform analysis (FFT), we transform the orientation signal and the "self" and "non-self" components of the sound signal perceived by each robot at each microphone from the time domain into the frequency domain. By looking at the power spectral density (PSD) we observe that: (i) the "self" component of each robot does not have any harmonic at frequency different from 0Hz; (ii) for all the robots, there are three main harmonic components in the spectrum of the orientation signal (see Table 3 columns 3, 4, 5); (iii) the "non-self" component of each robot has only one main harmonic (see Table 3 columns 6, and 7); (iv) fn_1 of robot R_{AL} differs from fn_1 of both robots R_{IR}.

These observations allow us to conclude that for each robot there are no oscillatory phenomena in sound production. Oscillations are instead observed in the perceived sound. These periodic phenomena are produced by the rotational movement of each robot through the effect that the movement has on the characteristic of the "non-self" components. A further evidence of the causal relationship between the rotational movement and the oscillation of the "non-self" components is given by the fact that the main harmonic of the "non-self" components has the same frequency of the first harmonic of each robot's orientation signal (see Table 3, columns 3 and 6). Moreover, the similarities of the first harmonic of the "non-self" component between robots R_{IR} and the differences between robots R_{IR} and robot R_{AL} confirm that there is a dynamic speciation of the characteristics of the homogeneous controllers with respect to the physical properties of the robots. In particular, robot R_{AL} seems to rotate slightly faster than the other two robots.

So far, we have identified periodic phenomena and their relative frequencies in sound signals and in the rotational movement of the robots. The next step of our analysis focuses on the characteristics of the "non-self" components. We use the frequencies of the main harmonic fn_1 obtained from the PSD analysis to filter the sound signals. In particular, we applied a narrow bandpass filter at frequencies 0 Hz and fn_1. In this way, we transform the "non-self" components into sinusoidal signals $ns_i = \alpha + p * \sin(2\pi fn_1 t)$, where α is the DC offset of the signal, p is the peak amplitude, and fn_1 is the frequency of oscillation.

From our analysis, it results that the average amplitude and standard deviation of the "self" components and the DC offset and peak of ns_i with $i \in [1, 3]$ do not substantially differ (i) among the robots; (ii) between the two sensors (S1 and S2); and (iii) between the two environments (i.e., *Env. L* and *Env. R*). In particular, the mean value of the "self" components contributes to more than 90% of the perceived sound (see Table 4 columns 2, and 8).[2] Given the high intensity of the "self" component, the "non-self" component can only induce changes in the perception of sound that are less than 10% of the sensors' receptive-field. By looking at the DC offset of the ns_i components (see Table 4 columns 4, 6, 10 and 12) we can infer that the "non-self" components are already very "weak", possibly due to the relatively "long" robot-robot distances. Despite this, if we sum, for each robot, for each sensor and for each environment, the average intensity of the "self"component", the DC offset and the peak of the ns_i, we notice that the reading of

[2]The "self" components are described by referring to their average and standard deviation since they do not present any periodic oscillations.

the sound sensory neuron saturates ($S1 \wedge \vee\ S2 > 1$). From this we infer that, if not attenuated by the shadowing effect, the "non-self" plus the "self" component may be sufficient to saturate the sensors' receptive field of the receiver. If we combine this data with the fact that the "non-self" components oscillate due to the rotational behaviour of the robots, we may formulate the following hypothesis: during navigation, the readings of the sound sensors of each robot may go through oscillations constrained between an upper and a lower bound. The upper bound is reached when the sum of the "self" and the "non-self" component corresponds to a value equal or bigger than the saturation value of the sound sensors (1.0). The lower bound is close to the intensity of the "self" component that is reached when the "non-self" components are strongly attenuated by the shadowing effect. These oscillations are very small since they concern less than 10% of the auditory receptive field, and certainly not very regular since the random noise applied to the sensors reading may disrupt the regularity of the oscillations determined by the contingencies (i.e., rotational movements and robots' relative distances). However, in spite of being small and noisy, these oscillations seem to be the only phenomenon related to the perception of sound that may play a significant role in the coordination of action of the group. In fact, given a controller sufficiently sensitive to capture them, they may represent a valuable perceptive cue for the receiver to spatially discriminate sound sources and consequently relative position and orientation of the emitter/s. Our hypothesis is that this phenomenon is exploited by the robots to remain close to each other while avoiding collisions and moving towards the target. Given the lack of complexity in robots' sound production, we exclude that signalling behaviour carries information concerning the classification of the environment. In the following Section, we show the results of further post-evaluation tests that are meant to assess the significance of our hypothesis on the role of signalling behaviour for spatial discrimination and coordination of actions.

5.4 The role of signalling behaviour

In this Section, we present the results of post-evaluation tests which aim to assess whether ILD is a cue used by the robots to coordinate their actions, as hypothesised in the previous Section. We run two series of post-evaluation tests. In the first series, we interfere with the propagation of sound in the environment by disrupting the orientation of the robot emitter with respect to the heading of the receiver. We refer to this as the *orientation test* (see Figure 7a). In the second series, we interfere with the propagation of sound in the environment by disrupting the sender-receiver distance. We refer to this as the *distance test* (see Figure 7b). In each of these tests, the robots undergo a set of 1200 trials in each type of environment. For all the simulation cycles following the first 10 seconds of each trial, the sound sensor readings of each agent are computed with respect to a hypothetical state of the system in which the senders are supposed to be:

orientation test: re-oriented by a fixed angular displacement, ranging from a minimum of 18° to a maximum of 180°, with a randomly chosen direction (clockwise or anti-clockwise) with respect to the heading of the receiver.

Table 4: This table shows, for each robot and for each microphone—$S1$ and $S2$—the average and standard deviation of the intensity of the "self" component (columns 2, 3, 8, and 9), and the DC offset (columns 4, 6, 10, and 12) and peak (columns 5, 7, 11, and 13) of the filtered "non-self" component ns_i, of sound intensity values recorded during a successful trial in each type of environment. See text in Section 5.3 for details on the filtering.

	Env. L						***Env. R***					
	self		**ns_i**				**self**		**ns_i**			
			S1		**S2**				**S1**		**S2**	
	avg	**std**	**DC**	**peak**	**DC**	**peak**	**avg**	**std**	**DC**	**peak**	**DC**	**peak**
R^1_{IR}	0.934	0.03	0.115	0.076	0.106	0.056	0.937	0.027	0.103	0.070	0.093	0.051
R^2_{IR}	0.932	0.035	0.133	0.018	0.120	0.014	0.933	0.035	0.107	0.054	0.098	0.041
R_{AL}	0.926	0.018	0.115	0.064	0.115	0.060	0.921	0.0196	0.135	0.014	0.136	0.012

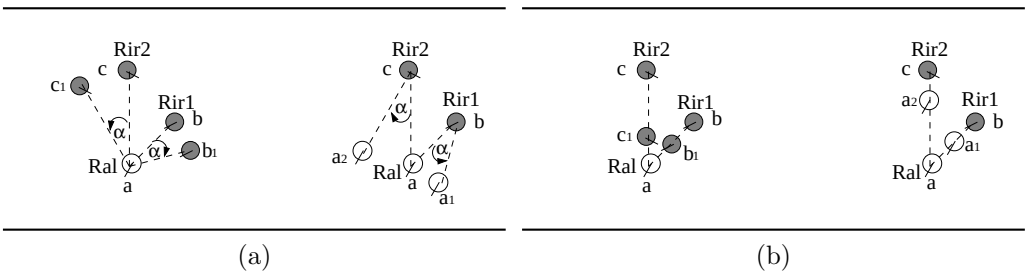

(a) (b)

Figure 7: (a) *orientation test*: on the left, robots R_{IR} (the grey circles) are displaced of an angle α with respect to robot R_{AL} (empty circle). On the right, robot R_{AL} is displaced of an angle α with respect to robots R_{IR}. (b) *distance test*: on the left, robots R_{IR} are moved to a fixed distance to robot R_{AL}. On the right, robot R_{AL} is moved to a fixed distance to robots R_{IR}. The left side of Figure a and b represents an hypothetical state in which the readings of the sound sensors of robot R_{AL} are computed considering R^1_{IR} located in position b_1 instead of b, and R^2_{IR} located in position c_1 instead of c. The right side of Figure a and b represents an hypothetical state in which the readings of the sound sensors of robot R^1_{IR} are computed considering R_{AL} located in position a_1 instead of a.

distance test: at a fixed distance to the receiver, ranging from a minimum of 2 cm to a maximum of 32 cm.

Note that, the updating of the infrared sensors of robots R_{IR} and of the ambient light sensors of R_{AL} does not undergo any disruption during these tests. The hypothetical states are taken into account only as far as it concerns the updating of the sound sensors' reading of one type of robot at the time. That is, during a set of trials, the sound perceived by robot R_{AL} is computed with reference to a hypothetical state in which the orientation/position of both robots R_{IR} is changed in order to meet the angular displacement/distance requirements. In this type of tests no disruptions are applied to update the sound perceived by robots R_{IR}. For the *orientation test* the results are shown in Figure 8a and 8b. For the *distance test*, the results are shown in Figure 9a and 9b. In a different set of tests, the sound perceived by the robots R_{IR} is computed with reference to a hypothetical state in which the orientation/position of robot R_{AL} is changed in order to meet the angular displacement/distance requirements. In this type of tests no disruptions are applied to update the sound perceived by robot R_{AL}. For the *orientation tests* the results are shown in Figure 8c and 8d. For the *distance test*, the results are shown in Figure 9c and 9d.

Generally speaking, by varying the sender-receiver orientation/distance, we indirectly increase/decrease the magnitude of the "non-self" component. In particular, those hypothetical states in which the sender-receiver distances tend to be decreased with respect to normal conditions, produced an increase of the magnitude of the "non-self" component and consequently an increase of the proportion of time in a trial the sound sensors are saturated. The same effect is obtained by applying angular displacements which increase the attenuation factor ψ. On the contrary, those hypothetical states in which the sender-receiver distances tend to be increased, produce a decrease of the magnitude of the "non-self" component and consequently a decrease of the proportion of time in a trial the sound sensors are saturated. The same effect is obtained by applying angular displacements which decrease the attenuation factor ψ. However, while the *distance test* preserves the intensity differences between the sound perceived in each ear due to the relative orientation of the sender with respect to the receiver, the *orientation test* disrupts any kind of regularities in the perception of sound which are linked to sender-receiver relative orientation. Therefore, a drop in performance at the *orientation test* is a sign that ILD is a cue exploited by the robots to successfully carry out their task. Contrary to the *orientation test*, the *distance test* informs us on the robustness of the mechanisms that exploit ILD with respect to variation concerning the magnitude of the oscillations in sound perception mentioned in Section 5.3.

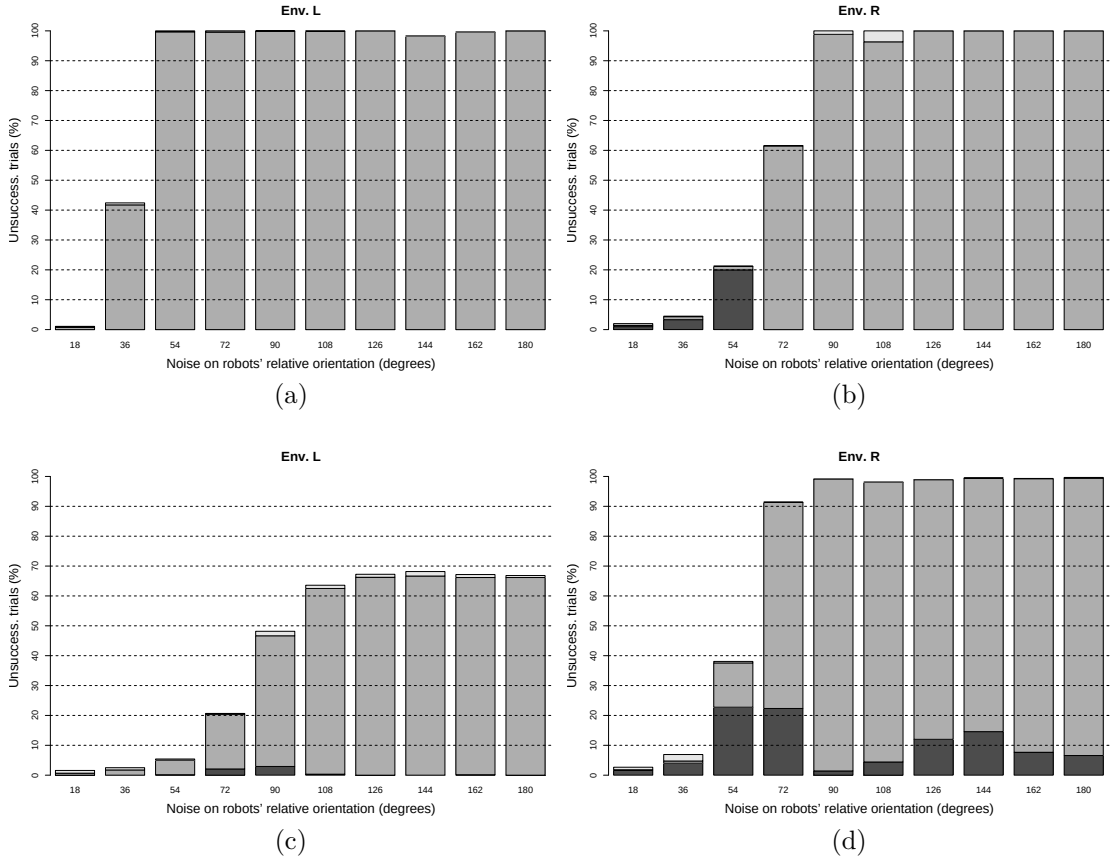

Figure 8: Percentage of failure during 1200 trials in each type of environment in post-evaluation tests with alterations applied to the relative orientation of the robots during the computation of the perceived sound. In (a) and (b) the robots R_{IR}, during all the simulation cycles following the first 10 seconds of any trial, are considered to be re-oriented with respect to the heading of robot R_{AL} by applying the angular displacement indicated on the horizontal axis and randomly choosing the direction of displacement (i.e., clockwise or anti-clockwise). In (c) and (d) the robot R_{AL} is re-oriented with respect to the heading of each robot R_{IR} as explained above. (a) and (c) refer to tests in *Env. L*; (b) and (d) refer to tests in *Env. R*. The black area of the bars refers to the percentage of trials terminated without collisions and with the group not having reached the target. The light grey area of the bars refers to the percentage of trials terminated due to robot-robot collisions. The dark grey area of the bars refers to the percentage of trials terminated due to robot-wall collisions.

The results of the tests shown in Figure 8 and 9 are very informative. First, the performance of the group is significantly disrupted by alterations which concern the orientation of the sender with respect to the heading of the receiver. Figure 8 shows that the bigger the magnitude of the disruption the higher the percentage of failure of the system. This proves that intensity differences between the sound perceived in each ear have a bearing on the development of effective navigational strategies as hypothesised above. In particular, regularities in the oscillation of the sound sensors' reading linked to the environmental contingencies and to the "variation" of the "non-self" component, are important perceptual cues exploited by the agents to coordinate their movements. The majority of failure are due to robot-wall collision. In particular, by looking at the behaviour of the group in these conditions, we noticed that, under the effects induced by the disruptions, the robots are not capable of remaining close to each other—i.e., within the infrared

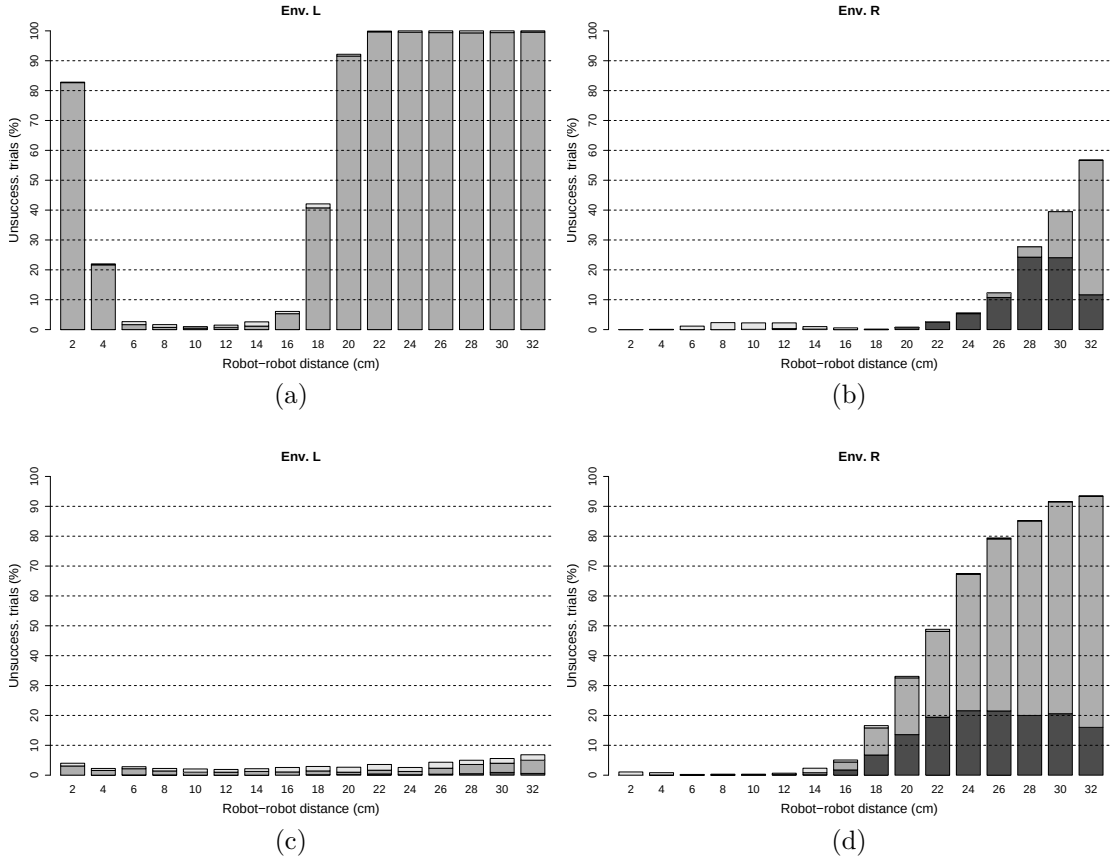

Figure 9: Percentage of failure during 1200 trials in each type of environment in post-evaluation tests with alterations applied to the robot-robot distance during the computation of the sound perceived by the receiver. In (a) and (b) the robots R_{IR}, during all the simulation cycles following the first 10 seconds of any trial, are both considered to be at the distance to robot R_{AL} indicated on the horizontal axis. In (c) and (d) the robot R_{AL} is considered to be at the distance to robots R_{IR} indicated on the horizontal axis as explained above. (a) and (c) refer to tests in *Env. L*; (b) and (d) refer to tests in *Env. R*. The black area of the bars refers to the percentage of trials terminated without collisions and with the group not having reached the target. The light grey area of the bars refers to the percentage of trials terminated due to robot-robot collisions. The dark grey area of the bars refers to the percentage of trials terminated due to robot-wall collisions.

sensors' range. When the distances becomes too high, the robots start wandering around the arena, and the trial terminates due to a collision of the robot R_{AL} with the arena walls. Only in few circumstances the robots do not lose contact to each other but they are not capable of reaching the target within the time-limits (see Figure 8 black area of the bars).

The results shown in Figure 9 tell us that the mechanisms which exploit ILD for spatial discrimination and behavioural coordination are quite robust to a general increment of the sound intensity. An exception is the case in which the robots R_{IR} are placed very close to robot R_{AL} in an environment *Env. L* (see Figure 9a). By looking at the behaviour of the group in these conditions, we noticed that contrary to what observed in the *orientation test*, in the unsuccessful trials the robots manage to remain close to each other—i.e., within the infrared sensors' range. However, the robot R_{AL} is not capable of making the left turn hitting the walls close to the corner. This is a quite general problem in these type of tests. That is, the robots manage to approach the turn (left or right) relatively close to each other but they fail to complete the turn due to

Table 5: Results of two post-evaluation tests for the best evolved controller of run n. 9. See caption Table 2 for details.

	(%) of failure due to time limit		(%) of failure due to collisions		Distance to the light avg		Distance to the light std	
Test	**left turn**	**right turn**	**left turn**	**right turn**	**left turn**	**right turn**	**left turn**	**right turn**
E	76.50	81.75	0.00	0.00	20.73	20.34	2.54	2.27
F	15.33	10.00	16.25	1.67	29.49	28.76	7.09	12.01

the lack of behavioural coordination of robot R_{AL} during the turn. Another significant result is that the robustness with respect to this type of disruptions is not the same for both types of robots. In general, the most disruptive effects are recorded in those tests in which discrepancies are artificially induced between the current state of the system and the perception of sound of robot R_{AL}. Disruptions on the perception of sound of robots R_{IR} when the group is located in *Env. L* do not alter the performance of the system with respect to the normal conditions (see Figure 9c). This suggests that, in *Env. L* robots R_{IR} "favour" infrared over sound sensors to coordinate their actions.

5.5 Robustness to environmental changes

In this section, we show the results of further post-evaluation analyses that aim to test the robustness of the group navigation strategy in environments that differ from those experienced by the robots during the evolutionary phase. In particular, in *tests E* and *F* the best controller run n. 9 is cloned in a group of three robots, in which, as during evolution, two are R_{IR} and one R_{AL}. Contrary to the evolution, in *test E*, we simply remove the walls. The absence of walls causes a sensible increase in the percentage of failure in both environments (see columns 2, 3, Table 5). The magnitude of the disruption seems to suggests that the walls are environmental structures systematically exploited by the robots to approach the target. For example, the group navigation may be based on a wall-following strategy that, given the structure of the world, guarantees the group to reach the target. However, the fact that the average distance of the team's centroid to the light in the unsuccessful trials is around 20 cm seems to indicate that the absence of walls does not hinder the capability of the group to perform phototaxis. Rather, the walls seem to play an important role during the last phase of the group navigation, when the centroid of the team is required to get into the proximity of the target.

In *test F*, the team is located in two different types of environment as depicted in Figure 10. These environments differ from those experienced by the team during the evolutionary phase because the walls are replaced by round obstacles of the diameter of a robot (see Figure 10). Note that the distance between any two round obstacles allows a robot to pass in between them. This

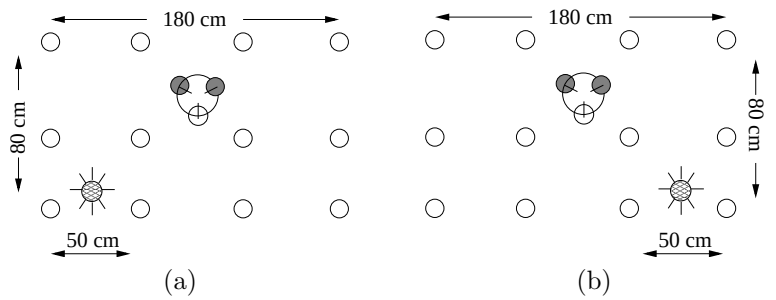

Figure 10: Environment that requires (a) a left turn; (b) a right turn. These environments differ from *Env. L* and *Env. R* because the walls are replaced by round obstacles of the diameter of a robot.

means that the probability of robot R_{AL} to collide against one of those round objects depends very much on the spatial configuration of the group during navigation. The results of this test tell us more about the obstacle avoidance capability of the group. The fact that the percentage of unsuccessful trials is around 31% when the light is on the left and only around 11% when the light is on the right indicates that the navigation strategy of the group is quite robust to cope with the kind of environmental changes we made.

6 Conclusions

In a context in which robots differ in their sensory capabilities, cooperation and coordination of actions evolved for the group to achieve a common goal. Behavioural capabilities of the single agents become effective in a social context in which mutual dependencies at various operational levels characterise the system more than causal explanations. These dependencies have been investigated in Section 5.1, in which a series of post-evaluation tests proved that the group navigation strategy is based on genuine cooperation among the members of the group. No single agent can accomplish the task in the absence of one of the original elements of the group. Due to the characteristics of our robots, cooperation and coordination of actions is mainly achieved through sound signalling behaviour. Post-evaluations tests unveiled that (i) all the robots emit sound at a very high intensity and (ii) signalling behaviour is not characterised by oscillatory phenomena (see Section 5.3). However, periodic phenomena characterised the perception of sound, since they are generated by the receiver through a rotational movement associated to the phototaxis. The absence of complex structures in signalling behaviour and its homogeneity among the agents and between the two types of environment, let us exclude the possibility that signalling behaviour is linked to the characteristics of the environment and that it is used to generate alternative group navigation strategies. The results of investigations shown in Section 5.4, seem to suggest that sound signalling behaviour is mainly used in the context of spatial discrimination of sound sources. That is, the robots' rotational movement generates oscillation in sound perception which causes intensity differences between the sound perceived in each ear (ILD). These differences, which are potentially informative of the relative position of sound sources, are exploited by a receiver to generate adaptive actions to safely navigate (i.e, without collisions) towards the target.

From our point of view, the simulation model presented in this paper is characterised by distinctive features, such as (i) the model of the sound, (ii) the way in which the controllers are wired-up with the sensory apparatus of the robots, and (iii) the "dynamic speciation" of the homogenous controller, whose mechanisms underpin sensory-motor coordination and social interactions in structurally different agents. These features make the results particularly relevant for the ER practitioners.

As far as it concerns the model of sound, although inspired by the work of Di Paolo (2000), it presents peculiarities of particular interest. As in Di Paolo (2000), and contrary to other experimental works in ER, we did not make use of directional microphones, or any other form of hand-coded mechanisms to discriminate between different sound sources or between "self" and "non-self" produced sound. For example, in the work of Marocco and Nolfi (2006), four directional microphones capture the sound of the nearest robot located within $\pm 45.0°$ left or right of each microphone. As well as hardly portable on a physical system, these type of models, might preclude the possibility to investigate the principles underlying behavioural coordination through sound signalling in a team of autonomous agents, since the problem of synchronisation or turn-taking to avoid mutual interference, and of spatial discrimination of sound sources are eluded thanks to the implementation details. With respect to what described in Di Paolo (2000), we strongly simplified the characteristics of the robot's controller. In particular, we did not implement the neural structures which provide the agents in Di Paolo'work the means to further regulate the intensity of emission of sound (i.e., regulation for sound effector, see Di Paolo, 2000) and the receptiveness of the sound sensory neurons (i.e., sensory gain regulation, see Di Paolo, 2000). We simplified the mechanisms to constraint the production of sound by fixing a limit to the intensity of the signal which also correspond to the saturation level of the sound sensors. That is, the "self"

produced sound can completely saturate the sound sensors of the emitter. Although arbitrarily implemented by the authors, these simplifications were introduced to compensate for an increase in structural complexity of the controller due to the nature of the agents' sensory apparatus. In particular, while in the work described in (Di Paolo, 2000) the agents are equipped only with sound sensors, in this work the agents are equipped with sound receptors as well as light or infrared sensors. Moreover, we investigated teams of three robots instead of two robots. Possibly due to these differences, the evolved solutions in Di Paolo's work and in ours diverge significantly. While in Di Paolo's 2000 model, oscillations and synchronisation in sound production underpin behavioural coordination, in our model, there is no oscillation in sound production.

From an engineering point of view, it is worth to mention that, although extremely effective, the best evolved navigation strategy does not display a very efficient phototaxis. On the one hand, the strong rotational movement allows for behavioural coordination through sound signalling, as explained above. On the other hand, it slows down the movement towards the light. We believe that alternative navigation strategies, can potentially be achieved by reintroducing some of the mechanisms originally proposed in (Di Paolo, 2000). These mechanisms facilitate the evolution of oscillatory behaviour in sound production and the distinction between "self" and "non-self" component, without having to model phenomena such as time varying frequencies, Doppler effect, etc. A group of robots in which each agent is capable of differentiating between "self" and "non-self" and of associating the intensity of the sound perceived in each ear with the distance to the sound source may favour linear over rotational movements. Other hardware specifications, such as the position of the microphones on the robot body might certainly facilitate the evolution of more efficient phototaxis. These issues will be the subject for future investigations.

As far as it concerns the way in which the controllers are wired-up with the sensory apparatus of the robots, we would like to provide further justifications for our implementation choices. Our goal was to generate through artificial evolution a controller capable of guiding both types of robots. For this reason, we chose to keep the group homogenous with respect to the controllers. That is, at time 0 of each trial, each robot is equipped with exactly the same control structure. However, the properties of the controllers allow for a "dynamic speciation": that is, a differentiation of the functionalities of each controller induced by the singularities in the flow of perception experienced by each robot. Moreover, we wanted to reduce at a minimum the number of parameters which define the search space of the evolutionary algorithm. For this reason, we decided to use neural structures in which the same input neurons in different networks are linked to different type of sensors (see Section 3 for details). Our results suggest that implementation details make possible to generate through artificial evolution homogeneous controllers that can efficiently guide morphologically identical as well as morphologically different groups of robots. In our case, the differences in the flow of sensation coming from different sensory channels (i.e., infrared sensors, ambient light and sound sensor) contribute to induce the specialisation of the controllers with respect to the physical characteristics of the robots, and to the relative role that they play in the group (i.e., the "dynamic speciation"). This latter mechanism can be also exploited in case of hardware failure, in which an on-line re-assignment of association between agent's sensors and network's input neurons might provide a robust mechanism to preserve the functionality of multi-robot systems. However, in order to efficiently exploit our methodological choices in the context such as hardware failure, further investigations are required to determine the plasticity of controllers in those circumstances in which they have already undertaken a process of "dynamic speciation". That is, it is an open question whether a neuro-controller already specialised to receive as input the reading of a paricular set of sensors is capable of "redefining" its functionalities to guide a robot with a different set of sensors.

7 Acknowledgements

E. Tuci and M. Dorigo acknowledge European Commission support via the*ECAgents* project, funded by the Future and Emerging Technologies programme (grant IST-1940), and by COMP2SYS, a Marie Curie Early Stage Training Site (grant MEST-CT-2004505079). The authors

thank their colleagues at IRIDIA for stimulating discussions and feedback during the preparation of this paper, and the two anonymous reviewers for their helpful comments. M. Dorigo acknowledges support from the Belgian FNRS, of which he is a Research Director, and 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.

References

Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9:255–267.

Cangelosi, A. (2005). Approaches to grounding symbols in perceptual and sensorimotor categories. In Cohen, H. and Lefebvre, C., editors, *Handbook of Categorization in Cognitive Science*, pages 719–737. Elsevier.

Cangelosi, A. and Parisi, D., editors (2002). *Simulating the Evolution of Language*. Springer Verlag, London, UK.

Di Paolo, E. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):27–48.

Dudek, G. and Jenkin, M. (2000). *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, UK.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA.

Kandel, E., Schwartz, J., and Jessell, T., editors (2000). *Principles of Neural Science*. McGraw-Hill/Appleton and Lange, fourth edition.

Marocco, D. and Nolfi, S. (2005). Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. In Cangelosi, A., Bugmann, G., and Borisyuk, R., editors, *Modeling language, cognition and action: 9th Neural Computation and Psychology Workshop*. World Scientific.

Marocco, D. and Nolfi, S. (2006). Self-organization of communication in evolving robots. In Rocha, L., Yaeger, L., Bedau, M., Floreano, D., Goldstone, R., and Vespignani, A., editors, *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems (Artificial Life X*, pages 178–184. MIT Press, Cambridge, MA.

Miglino, O., Lund, H., and Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4):417–434.

Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221.

Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA.

Quinn, M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life: 6th European Conf. on Artificial Life*, pages 357–366. Springer Verlag, Berlin, Germany.

Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Phil. Trans. of the Royal Soc. of London, Series A*, 361:2321–2344.

Vicentini, F. and Tuci, E. (2006). *Swarmod*: a 2d *s-bot*'s simulator. Technical Report TR/IRIDIA/2006-005, IRIDIA, CoDE, Université Libre de Bruxelles. This paper is available at `http://iridia.ulb.ac.be/IridiaTrSeries.html`.

Vogt, P. (2005). Language acquisition and evolution. *Adaptive Behavior*, 13(4):265–268.