



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

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Evolved homogeneous neuro-controllers for 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 capabilities. 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. A preliminary analysis of the mechanisms underlying the group behaviour is provided.

1 Introduction

Embodied autonomous systems are relatively recent methodological tools which can be used to investigate various aspects of social interactions and behavioural coordinations in artificial and natural organisms (see [9, 2]). In this type of systems, social behaviour is investigated by firstly determining the characteristics of the agents’ embodiment (e.g., sensory and motor capabilities of the agent) and the world that they inhabit, and by subsequently looking at how the latter features influence social skills.

This approach is particularly prominent in a subset of embodied autonomous systems, generally referred to as Evolutionary Robotics models (ER, see [7]). 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 [8, 10, 1, 11, 5]). 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 Fig. 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




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

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 Fig. 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 Fig. 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 Sec. 2). However, contrary to other studies (see [5, 1]), 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 the own sound from those of others and the mutual interference due to lack of turn-taking (see [8]).

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. The mutual coordination results particularly striking so that, as already emphasised in a similar model [8], 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, 12 cm of diameter, equipped with a variety of sensors, and whose mobility is ensured by a differential drive system (see [6] 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_6) positioned at $\pm 67.5^\circ$ with respect to the orientation of the robot (see Fig. 2). 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.

All robots are equipped with a sound output (S.O.) that is situated in the centre of the body of the robot, and with two omnidirectional microphones (S_1 and S_2), placed at $\pm 45^\circ$ with respect to the robot's heading. Sound is modelled as an instantaneous, additive field of single frequency with time-varying intensity ($\eta \in [0.0, 1.0]$) which decreases with the square of the distance from the source, as previously modelled in [8]. Sound intensity is regulated by the firing rate of neuron N_{14} (see Sec. 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 [8]. There is no attenuation of intensity for self-produced signals which can in principle be loud enough ($\eta = 1.0$) to make it impossible for a robot to perceive sound signals emitted by others. 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 travelled by the signal within the body of the receiver (see [8] for details).

Concerning the function that updates the position of the robots within the environment, we employed the Differential Drive Kinematics equations, as presented in [3]. 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 [12].

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 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




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 $\frac{Ir_0+Ir_1+Ir_2}{3}$, 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 S_1 , and N_6 from S_2 . The input layer of R_{AL} takes readings as follows: N_1 and N_2 take input from AL_1 , N_3 and N_4 take input from AL_6 , N_5 from S_1 , and N_6 from S_2 . M_1 and M_2 are respectively the left and right motor.

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:

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

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^{th} to neuron i^{th} , β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate. The cell potentials y_i of the 12th and the 13th neuron, mapped into $[0,1]$ 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 14th neuron, mapped into $[0,1]$ by a sigmoid function σ , is used by the robot to control the intensity of the sound emitted η . 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_j 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 [4]. 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 six times, three trials in *Env. L*, and three trials in *Env. R*. The sequence order of environments within the six 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. By taking inspiration from the work of Quinn et al. [11], the fitness score is computed as follows:

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

As in [11], 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 . 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.

$P = 1 - \left(\sum_{i=1}^3 c_i / c_{max} \right)$ 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 1000 and 1500 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 Opteron244TM CPU running GNU/Linux Debian 3.0 OS. 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 south of 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 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 Fig. 3b); (ii) for each spatial arrangement, we identified three possible relative positions of the robot R_{AL} with respect to the walls' corridor (see white circle in Fig. 3b); (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 $[-30^\circ, 30^\circ]$ 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




Figure 3: (a) Results of post-evaluation showing the percentage of success of the best evolved controllers of each run over 1200 trials per type of environment. White bars refers to *Env. L*, and black bars to *Env. R*. (b) The robots’ initial positions during the post-evaluation phase. White circles refers to R_{AL} , grey circles refers to R_{IR} .

reset (see Sec. 3 for details).

The results of the post-evaluation phase are shown in Figure 3a. We notice that the best controller is the one produced by run n. 2, achieving a performance over 90% in *Env. L* and *Env. R*. Runs n. 4, 9, and 10 display a performance over 80%, run n. 1, and 7 displays a performance around 75% in both environments. Note that when looking at the performances of the best evolved controllers, as shown in Figure 3a, 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, for all the runs, the majority of the failures are due to collisions. An exception is run n. 1 which seems to be only minimally affected by this factor (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 the robots guided by this controller seem to be capable of covering much of their initial distance to the light. Therefore, the small percentage of collisions is indeed the result of an effective coordination of actions among the agents invalidated by the lack of time to complete the task due to a slow phototactic movement, rather than, for instance, a consequence of the lack of movement of the group towards the light.

In the rest of this section, we concentrate on the analysis of the controller of run n. 2, which

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.

run	(% of failure due to time limit)		(% of failure due to collisions)		Distance to the light			
	<i>Env. L</i>	<i>Env. R</i>	<i>Env. L</i>	<i>Env. R</i>	avg		std	
n. 1	15.75	20.08	3.17	3.17	22.22	24.29	14.35	22.20
n. 2	1.42	0.00	7.50	0.42	71.63	87.96	20.47	0.46
n. 3	19.17	4.67	69.00	15.50	45.11	36.67	23.06	13.73
n. 4	0.00	4.92	11.25	12.25	62.72	52.80	22.24	23.98
n. 5	20.75	11.33	15.58	21.83	48.83	47.31	38.70	25.72
n. 6	43.83	61.67	22.58	5.75	35.11	30.76	16.12	12.80
n. 7	0.00	10.17	12.00	12.33	67.60	42.93	15.98	28.05
n. 8	36.33	3.58	12.33	23.92	31.22	58.94	22.05	23.46
n. 9	0.67	7.50	5.00	6.25	55.49	29.10	22.73	17.29
n. 10	0.00	6.42	3.67	6.83	54.72	50.59	18.08	32.26

proved to be the most effective at the first post-evaluation test. In particular, we try 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.

5.1 Further analysis of the best evolved navigation strategy

In an effort to understand how the robots manage to cooperate and coordinate their actions in order to solve the task, we repeated the post-evaluation test described at the beginning of Sec. 5 for groups of robots controlled by controller run n. 2. However, in these series of tests, the robots are deprived in various ways of sensory information which may or may not turn out to be crucial for the achievement of the task.¹ Recall that, only by listening to the sound signals emitted by robots R_{IR} , robot R_{AL} can avoid collisions. Sound signalling *and/or* coordination through the infrared sensors might play a significant role in guiding the group towards the target.

First, we run two tests, referred to as *Test A* and *Test B*, which should reveal to us whether the robots employ effective navigational strategies based on cooperation and coordination of actions or rather 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. 2 is cloned on three robots R_{IR} . Consequently, the robots have no means to know where the light is placed. As shown in Table 2, the group was 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

¹In the post-evaluation tests in which alterations concern the agents' received sound signal, or the nature of the group (i.e., what types of robot are part of the group), and/or the characteristics of the environment, the changes are applied after 10 s (i.e., 100 simulation cycles) from the beginning of each trial. This give time to the controllers to reach a functional state different from the initial one, arbitrarily chosen by the experimenter, in which the cell potential of the neurons is set to 0 (see Sec. 3).

Table 2: Results of different post-evaluation tests for the best evolved controller of run n. 1. See text in Sec. 5.1 for details. Note that in all trials the initial distance between the centroid of the group and the light is equal to 85.14 cm.

Test	(% of failure due to time limit)		(% of failure due to collisions)		Distance to the light			
	<i>Env. L</i>	<i>Env. R</i>	<i>Env. L</i>	<i>Env. R</i>	avg		std	
	<i>Env. L</i>	<i>Env. R</i>	<i>Env. L</i>	<i>Env. R</i>	<i>Env. L</i>	<i>Env. R</i>	<i>Env. L</i>	<i>Env. R</i>
A	100	100	0.00	0.00	85.32	85.50	8.24	8.31
B	0.00	0.00	100.00	100.00	90.83	104.90	3.64	5.36
C	100	100	0.00	0.00	122.57	127.94	5.19	4.73
D	2.42	0.00	46.17	0.58	75.06	84.83	25.32	4.90
E	0.25	0.00	12.08	6.75	70.03	37.03	22.33	23.06

columns 4 and 5, Table 2). In *Test B* a single robot R_{AL} is controlled by the best controller of run n. 2. The results tell us that R_{AL} , if left without robots R_{IR} , systematically collides with walls (see columns 4 and 5, Table 2). *Test A* and *B* suggest that the successful strategies of run n. 2 are based on effective coordination of actions and cooperation among the different types of agents of the group. In brief, there is neither phototaxis nor any other movement along the corridors if robot R_{AL} is missing in the group. There is neither obstacle avoidance nor successful phototaxis if a single robot R_{AL} is left alone in this simple world. Surprisingly, while robots R_{IR} retain their capability to avoid obstacles if situated in an “odd” group of all R_{IR} robots, a single robot R_{AL} can hardly perform phototaxis if left alone. This is shown by the results of *Test C* in which a single robot R_{AL} is placed in a boundless arena (no walls) with only a light at around 85 cm away from it. As proved by the final distance to the light (see column 6 and 7, Table 2), a single robot R_{AL} is not capable of approaching the light source. Oddly enough, it displays an anti-phototactic movement. In summary, the different types of robot complement each other not only to accomplish the task, but also to carry out those functions for which they are more apt (e.g., phototaxis in R_{AL}).

In *tests D* and *E*, the best controller run n. 2 is cloned on a group of three robots, in which, as during evolution, two are R_{IR} and one R_{AL} . Contrary to the evolution, in *test D*, the robots R_{IR} only hear their own sound; in *test E*, the robots R_{IR} can potentially perceive the sound emitted by the robot R_{AL} but they can not hear each other’s sound. These tests should help us to understand more about the significance of sound signalling. Data in Table 2 show that, in *Test D*, robots do not systematically fail to reach the target. Although the performance in *Env. L* is severely disrupted with almost 50% of unsuccess rate, in *Env. R* the group performance is not touched by the alterations we applied to the system (see columns 2, 3, 4 and 5, Table 2). The failure in *Env. L*, is mostly due to collisions, which seem to occur rather far away from the lights (see columns 6, and 8, Table 2). In summary, the sound received by the robots R_{IR} from robot R_{AL} seems to play a significant role in carrying out obstacle avoidance in *Env. L*.

In *Test E*, 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. 2 are only marginally based on communication through sound signalling between the robots R_{IR} . Initially, we thought that this latter phenomenon was a side effect of the spatial arrangement of the group during navigation. For instance, if the robots form a chain in which the robot R_{AL} is in the middle position and the other two robots R_{IR} are at the two ends of the chain, then the latter robots may not hear each other because of the distance between them. Consequently, preventing the robots R_{IR} from hearing each other can not affect in any way a navigational strategy of a group that does not rely on this element. However, by looking at the spatial arrangement of the robots during navigation, we saw that, within a trial, they tend to dispose themselves in various spatial configurations in which the two robots R_{IR} do perceive each other’s signals. This implies that

robots R_{IR} might be capable of discriminating among agents of different type (i.e., R_{IR} , R_{AL}). However, this and other issues related to management of the coordination and cooperation of the group can not be inferred from this preliminary analysis, and they need to be further investigated.

6 Conclusions

In a context in which robots differ in their sensory capabilities, cooperation and coordination of actions evolved 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. The agents (i) emit sound signals that are not too loud to hinder the perception of the sound emitted by the others, but loud enough to be captured by the other robots if relatively close to the emitter; (ii) negotiate a common direction of movement; and (iii) navigate safely (i.e, without collisions) towards the target. The “dynamic speciation” of the homogenous controller, whose mechanisms underpin sensory-motor coordination and social interactions in structurally different agents, is particularly significant. From an engineering point of view, these results suggest that homogeneous controllers can be efficiently exploited to control morphologically identical as well as morphologically different groups of robots. This element 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. Moreover, a better coordination of actions might be achieved by varying the characteristics of the sound and/or morphological features of the sound signalling systems—e.g., the number and/or the position of the loudspeakers and microphone. Finally, further investigations need to be carried out to provide a deeper operational explanation of the properties of the system. Does the variability in the emission of sound reflect a simple “vocabulary” grounded on sensor-motor activity of the agent? This issue is an interesting subject for future investigations.

7 Acknowledgements

Elio Tuci is supported by the *ECAgents* project, an european project funded by the Future and Emerging Technologies programme of the European Commission, under grant IST-1940. 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. Marco 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

- [1] G. Baldassarre, S. Nolfi, and D. Parisi. Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9:255–267, 2003.
- [2] A. Cangelosi and D. Parisi, editors. *Simulating the Evolution of Language*. Springer Verlag, London, UK, 2002.
- [3] G. Dudek and M. Jenkin. *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, UK, 2000.
- [4] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.

- [5] D. Marocco and S. Nolfi. Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. In A. Cangelosi, G. Bugmann, and R. Borisjuk, editors, *Modeling language, cognition and action: 9th Neural Computation and Psychology Workshop*. World Scientific, London, UK, 2005.
- [6] F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221, 2004.
- [7] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA, 2000.
- [8] E. Di Paolo. Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):27–48, 2000.
- [9] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, Cambridge, MA, 2001.
- [10] M. Quinn. Evolving communication without dedicated communication channels. In J. Kelemen and P. Sosik, editors, *Advances in Artificial Life: 6th European Conf. on Artificial Life*, volume 2159 of *LNCS*, pages 357–366. Springer Verlag, Berlin, Germany, 2001.
- [11] M. Quinn, L. Smith, G. Mayley, and P. Husbands. 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, 2003.
- [12] F. Vicentini and E. Tuci. Swarmod: a 2d s-bot’s simulator. Technical Report TR/IRIDIA/2006-005, IRIDIA, Université Libre de Bruxelles, 2006. This paper is available at <http://iridia.ulb.ac.be/IridiaTrSeries>.