

Evolution of Acoustic Communication Between Two Cooperating Robots

Elio Tuci and Christos Ampatzis

CoDE-IRIDIA, Université Libre de Bruxelles, Bruxelles, Belgium
{etuci,campatzi}@ulb.ac.be

Abstract. In this paper we describe a model in which artificial evolution is employed to design neural mechanisms that control the motion of two autonomous robots required to communicate through sound to perform a common task. The results of this work are a “proof-of-concept”: they demonstrate that evolution can exploit a very simple sound communication system, to design the mechanisms that allow the robots cooperate by employing acoustic interactions. The analysis of the evolved strategies uncover the basic properties of the communication protocol.

1 Introduction

This paper is about the evolution of acoustic communication in a two robot system, in which the agents are required to coordinate their efforts to perform a common task (see Sec. 2). The robots’ mechanisms are determined by design methods referred to as Evolutionary Robotics (see [6]). That is, an artificial evolutionary process sets the parameters of neural networks controllers. The latter are in charge of the robots’ actions by setting the states of the agents’ actuators. Although from a different perspective and with different motivations, the issue of the evolution of acoustic communication has already been investigated in several research works. Some of these works model aspects of the evolution of communication in living organisms (see [7,4,10]). Other studies aim to engineer acoustic communication systems that improve the effectiveness of the robots’ collective responses (see [9,8,1]). Either biologically or engineering inspired these studies exploit the properties of the evolutionary robotics approach in which the designer is not required 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. The results of the evolutionary process (i.e., the behaviour of the robots and the underlying mechanisms) inform the designer on the effects that the physical interactions among embodied agents and their world have on the evolution of individual behaviour and social skills.

Following this line of investigation, our work aims to demonstrate the effectiveness of a very simple sound signalling system in a context in which the robots are demanded to share individual experiences to build a common perspective of their world. The robots can communicate by using an extremely simple binary

signalling system (i.e., ON/OFF). As far as we know, this is the first study that investigates a communication scenario in which a bi-directional interaction is required by the robots to accomplish a common goal. Communication is based on the emission in time of asynchronous and mutually determined single tone signals. The results of this work should be taken as a “proof-of-concept” concerning the potentiality of the proposed approach to the design of acoustic communication mechanisms in multi-robot systems. We demonstrate that it is possible to use evolution to define the mechanism underlying a bi-directional communication protocol based on a very simple acoustic system.

2 The Task

The robot environment is a rectangular arena—120cm by 50cm—divided into two equal sides by a horizontal bar that revolves—i.e., the revolving door. There are three lights L_1 , L_2 and L_3 . When L_1 and L_2 are turned on, L_3 is turned off and vice versa. L_1 can only be seen by a robot located in the lower side of the arena while L_2 can only be seen by a robot located in the upper side of the arena. L_3 can be seen from anywhere in the environment. The arena floor is white except in the proximity of L_1 and L_2 up to a distance of 15cm from the lights, where the floor is painted in black or grey. The robots can experience four different combinations of black and grey zones (see Fig. 1). The type of environment in which the robots are located is labelled according to the combination of the colour of the floor in the two painted zones. In detail, the environments are labelled E_{xx} , where the first digit corresponds to the colour of the floor in the proximity of L_1 and the second digit to the colour of the floor near L_2 . Grey colour corresponds to 0, while black colour corresponds to 1. The four types of environment are: E_{10} , E_{01} , E_{00} , and E_{11} . The revolving door rotates from the horizontal to the vertical position if simultaneously pushed by both robots in the

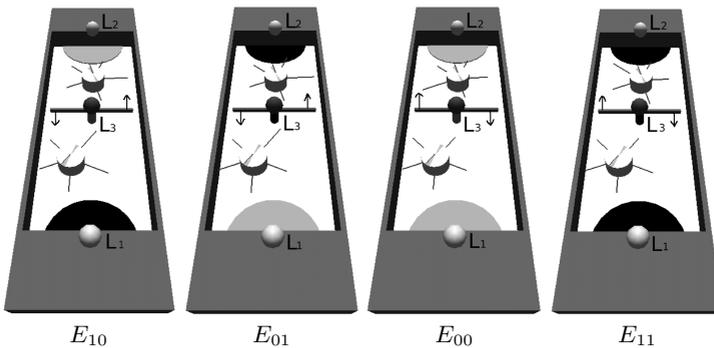


Fig. 1. The four environments E_{10} , E_{01} , E_{00} , and E_{11} . L_1 , L_2 and L_3 refer to the lights. The revolving door is indicated by the horizontal bar in the centre of the arena. In each environment, the arrows indicate the direction in which the door revolves. The cylinders with spikes on the white floor represent the robots.

proper direction. Pushing forces exerted by a single robot on the revolving door are not enough to open it. The direction of rotation changes according to the type of environment. The robots have to exert forces to make the door rotate (a) clockwise, if they are located in E_{00} or in E_{11} ; (b) anticlockwise, if located in E_{10} or in E_{10} (see the arrows in Fig. 1).

At the beginning of the first trial and in those that follow an unsuccessful one, the robots are randomly placed in the proximity of L_3 . In trials following a successful one, the robots are not repositioned. The sequence of desired actions that each robot is demanded to carry out during a trial can be decomposed into two phases. At the beginning of the first phase, L_1 and L_2 are turned on, the revolving door is in the horizontal position and the colour of the floor in the proximity of L_1 and L_2 is set according to the type of environment that characterises the trial. During this phase, the robots are required to find the painted zone in their side of the white arena floor and remain for at least 6s on the painted zone. This exploration is facilitated by the presence of the lights that can be used as beacon (i.e., L_1 for the robot located in the lower side of the arena and L_2 for the robot located in the upper side of the arena). The first phase terminates once the 6s on the painted zones are elapsed for both robots. At this point, L_1 and L_2 are turned off, L_3 turned on, and the second phase begins. In the second phase, the two robots are required to move back towards the middle of the arena, approach the revolving door, and simultaneously push the door in order to open it and to reach the previously inaccessible opposite side of the arena. As mentioned above, the direction of rotation changes according to the type of environment. Therefore, to rotate the revolving door from the horizontal towards the vertical position the robots are required to “tell” each other the colour of the floor in the proximity of light— L_1 or L_2 —previously approached. A trial successfully terminates once both robots, by rotating the revolving door, move into the opposite side of the arena, and reach a distance of 24 cm from L_3 . At the end of a successful trial, L_3 is turned off, L_1 and L_2 are turned on, the rotating door automatically returns to the horizontal position and a new trial begins. A trial is considered unsuccessful, if a single robot exerts forces in both arms of the revolving door (i.e., west and east of L_3). This behaviour, referred to as trial-and-error strategy, is penalised by the fitness function (see Sec. 4).

Note that this task requires coordination of actions, cooperation and communication between the robots in order to successfully open the revolving door. For each robot, the perception of a grey or black floor can be associated both to a clockwise and anticlockwise rotational movement of the revolving door. Only the combination of the two coloured zones unambiguously identifies a rotational movement. Since a robot can only walk on a single zone per trial, the task can be successfully accomplished in all the environmental conditions only by a group of robots that communicate through sound. Without communication, a single robot can only exploit a trial-and-error strategy. By using a simple sound signalling system the robots should inform each other on the colour of the floor in the proximity of the light they perceive— L_1 or L_2 —and consequently push the door in the proper direction as explained above.

3 Methods

The robot and its world are simulated using simulation software based on Open Dynamic Engine (see <http://www.ode.org/>), a 3D rigid body dynamics simulator that provides primitives for the implementation of detailed and realistic physics-based simulations. Our simulation 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 [5]). Our simulated robot has a differential drive motion provided by a traction system composed of four wheels: two lateral, motorized wheels and two spherical, passive wheels placed in the front and in the back, which serve as support. The four wheels are fixed to the cylindrical body that holds the sensors. In particular, robots make use of 5 infrared sensors IR_i , two ambient light sensors AL_i , a floor sensor FS , a loudspeaker SO to emit sound and an omni-directional sound sensor SI to perceive sound (see Fig. 2a). Light levels change as a function of the robot’s distance from the lamp. FS , placed underneath the robot, detects the level of grey of the floor. It outputs the following values: 0 if the robot is positioned over white floor; 0.5 if the robot is positioned over grey floor; 1 if the robot is positioned over black floor. SO produces a binary output (on/off). SI has no directionality and intensity features. 10% uniform noise is added to IR_i and AL_i readings, the motor outputs and the position of the robot.

The controller of each agent is composed of two modules referred to as M_C and M_M (see Fig. 2b). The modularisation is hand-coded to facilitate the evolution of successful behavioural strategies. M_C is a non-reactive module, that is a six neurons fully connected continuous time recurrent neural network (CTRNN,

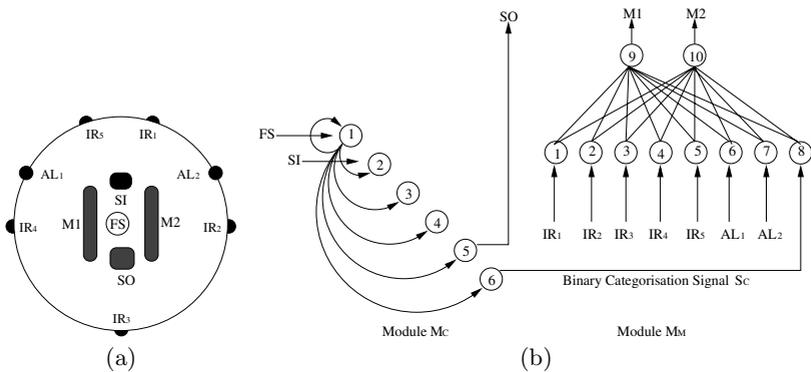


Fig. 2. (a) The simulated robot. $IR_i, i \in [1, 5]$ are the infrared sensors; $AL_i, i = [1, 2]$ are the ambient light sensors; FS is the floor sensor; SI is the sound sensor (i.e., the microphone); SO is the sound actuator (i.e., the loudspeaker); M_1 and M_2 are respectively the left and right motor. (b) The network architecture: module M_C and module M_M . For M_C only the efferent connections for one neuron are drawn. S_C is the binary categorisation signal sent, at each updating cycle, by M_C to M_M .

see also [2]). M_C is required to detect in which type of environment the robot is currently located. The categorisation has to be based on the FS 's readings of both robots. Thus, it demands communication between the agents. For this reason, M_C takes input from FS and SI and it outputs the state of the SO and SC (i.e., the binary categorisation signal). In other words, at every updating cycle, M_C is in charge of (a) managing sound by producing the signal the robot emits and by receiving the signal of either robot, and (b) "informing" M_M on the type of environment in which the robot is currently located by setting the value of the binary categorisation signal SC either to 0 or 1. M_M is a reactive module, that is a feed-forward artificial neural network made of eight sensory neurons and two output neurons. M_M is demanded to (a) guide the robot avoiding collisions with the arena walls, and (b) "parse" the value of SC to determine in which side to push the revolving door (i.e., anticlockwise if current trial in E_{10} or E_{01} , clockwise if current trial in E_{00} or E_{11} , see also Fig. 1). M_M takes input from IR_i , $i \in [1, 5]$, from AL_i , $i = [1, 2]$, and SC , and it outputs the speed of the robot's wheels. The following associations (a) $SC = 1$, robots located in E_{10} or E_{01} , anticlockwise rotational direction of the revolving door, and (b) $SC = 0$, robots located in E_{00} or E_{11} , clockwise rotational direction of the revolving door, are determined *a priori* by the experimenter (see Sec. 4). The neural mechanisms and the communication protocol required by the robots to build these relationships from the sensors' readings are set by evolution.

The states of the neurons of M_C and M_M are governed by the equations (1) and (2) respectively:

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

$$dy_i = \begin{cases} -y_i + gI_i & i \in [1, 8] \\ \left(-y_i + \sum_{j=1}^8 \omega_{ji} \sigma(y_j + \beta) \right) & i \in [9, 10]; \end{cases} \quad (2)$$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i is 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 the bias term, $\sigma(y_j + \beta)$ the firing rate. The parameters ω_{ji} , τ , β and g are genetically encoded. 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 $dt = 0.1$. Note that the cell potentials of M_M 's neurons do not depend on time (see equation (2)). That is, the neurons' decay constant τ is set to 0.1, as the integration step-size dt . In M_C , the cell potentials y_i of the 5th and the 6th neuron, mapped into $[0, 1]$ by a sigmoid function σ , set the state of the robot's sound actuator SO and of the binary categorisation signal SC . The robot emits a sound if $SO \geq 0.5$. $SC = 1$ if $\sigma(y_6 + \beta_6) \geq 0.5$ otherwise $SC = 0$. In M_M , the cell potentials y_i of the 9th and the 10th neuron, mapped into $[0, 1]$ by a sigmoid function σ and then linearly scaled into $[-6.5, 6.5]$, set the robot motors output.

A simple generational genetic algorithm is employed to set the parameters of the networks [3]. 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 (also known as roulette wheel selection) from the 64 best individuals of the old population. Each genotype is a vector comprising 67 real values, chosen uniformly random from the range $[0, 1]$. The first 18 genes are used to set the parameters of M_M (i.e., 16 connection weights, 1 bias term and 1 gain factor both shared by all the input neurons). The other 49 genes are used to set the parameters of M_c (i.e., 36 connection weights, 6 decay constants, 6 bias terms, and 1 gain factor). More details on the genetic algorithm and on the genotype-networks mapping can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2007-005>.

4 The Fitness Function

During evolution, each genotype is translated into a robot controller (i.e., modules M_C and M_M see Sec. 3), and cloned in each agent. Then, the two robot group is evaluated two times in each environment type E_{11} , E_{00} , E_{01} , and E_{10} , for a total of eight trials. Note that the sequence order of the environment type experienced by the robots—randomly chosen at the beginning of each generation—has a bearing on the overall performance of the group since the robots’ controllers are reset only at the beginning of the first trial. Each trial differs from the others in the initialisation of the random number generator, which influences the robots’ starting position and orientation anytime the robots are positioned, and the noise added to motors and sensors. The robots are randomly placed in the arena at the beginning of the first trial and repositioned in subsequent trials following an unsuccessful one. Within a trial, the robots life-span is 90 simulated seconds (900 simulation cycles). A trial is terminated earlier in case a robot crashes with the arena walls, or if the group successfully accomplishes its task. For each trial $e \in [1, 8]$, the group is rewarded by an evaluation function which seeks to assess the ability of the robots to open the revolving door located at the centre of the arena (see Sec. 2). This requires the robots to be able to determine the nature of the environment (i.e., E_{11} , E_{00} , E_{01} , or E_{10}) by using acoustic communication. The final fitness F attributed to a group controlled by a specific genotype is the average group score over a set of eight trials. A detailed illustration of the fitness function can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2007-005>. Note that F doesn’t refer anyhow to signalling behaviour. F rewards the robots for accomplishing the task as detailed in Sec. 2. However, due to the nature of the task, the robots can be successful only if they coordinate their actions using the sound signalling system. By leaving signalling behaviour out of the fitness function, we clean our model from preconceptions concerning what (i.e., semantics) and how (i.e., syntax) successful group communicates, and we let evolution determine the characteristics of the communication protocol.

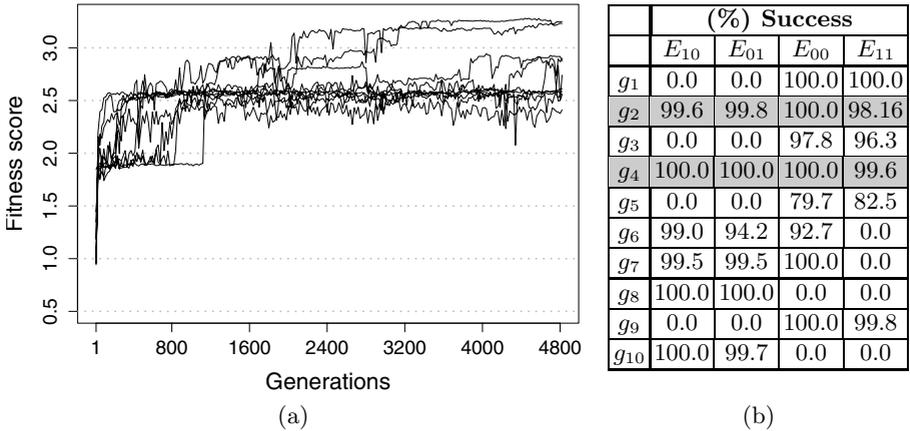


Fig. 3. (a) Fitness F of the best groups at each generation of ten evolutionary runs. (b) Results of post-evaluation tests, showing for the best evolved groups of each run the (%) of successful trials in each type of environment. In grey the successful groups.

5 Results

Ten evolutionary simulations, each using a different random initialisation, were run for 4800 generations. Given the nature of the fitness function, the highest fitness score that a group can reach is 3.4. This score corresponds to the behaviour of a group in which each robot (i) finds the coloured zone on the white arena floor; (ii) communicates to the robot at the opposite side of the arena the colour encountered in its side; (iii) uses the combination of colours to properly set the binary categorisation signal S_C ; and (iv) pushes the revolving door in the proper direction until it reaches the opposite side of the arena. Fig. 3a shows the fitness of the best groups at each generation for each evolutionary run. Notice that only two evolutionary runs managed to produce groups whose average fitness F is close to the maximum score. However, fitness scores lower than 3.4 might be associated to equally successful alternative strategies.¹ Thus, in order to have a better estimate of the behavioural capabilities of the best evolved controllers, we post-evaluate, for each run, the genotype with the highest fitness. These groups are referred to as $g_i, i \in [1, 10]$. The entire set of post-evaluations (i.e., 2400 trials, 100 evaluations for each permutation, $100 \cdot N!$ with $N=4$) should establish whether a group of robots is capable of accomplishing the task as described in Sec. 2 in all the four types of environment. The results of the post-evaluation tests are shown in Fig. 3b. The data confirm that only two groups g_2 and g_4 have a success rate higher than 98% in all four types of environment (see Fig. 3b, grey rows); g_1, g_3, g_5, g_8 and g_9 are capable of carrying out the task only when the door revolves clockwise, and g_{10} only when the door revolves anticlockwise; g_6

¹ Data not shown, movies of successful strategies, and further methodological details can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2007-005>.

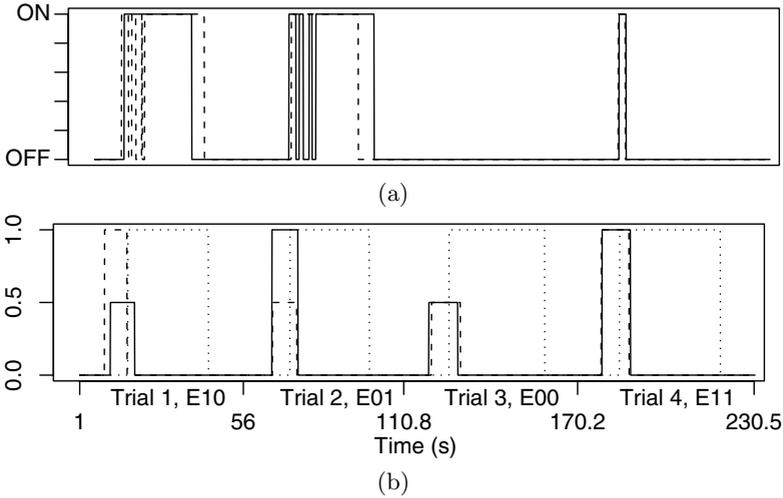


Fig. 4. Post-evaluations of group g_4 . Dashed lines refers to the robot placed at the beginning of trial 1, in the upper side of the arena; continuous lines refer to the robot placed in the lower side of the arena. (a) Sound signals. (b) Floor sensors readings. Dotted line indicates the state of L_3 , 1 = ON, 0 = OFF. On the x axis is indicated the time of start and end of each trial.

and g_7 fail in only one type of environment. From a behavioural point of view, the failure are due to trial-and-error strategy (data not shown, see footnote 1). That is, during the second phase of the task, both robots push the revolving door both west and east of L_3 instead of exerting forces directly on the proper side of the bar. Failure due to collisions are very rare. The lower success rate of g_{10} in E_{00} and E_{11} is mainly due to the fact that the robots of this group are not able to exert enough forces in order to rotate the revolving door (data not shown, see footnote 1). From a mechanism point of view, the failure of each single robot can be caused by either (a) M_C not capable of correctly categorising the environment by properly setting S_C as made explicit in Sec. 3 or (b) M_M not capable of “interpreting” the value of S_C as produced by M_C . Post-evaluation tests show that for almost all the unsuccessful groups it is M_C that by setting incorrectly the value of S_C , does not allow M_M to choose the correct direction of rotation of the revolving door (data not shown, see footnote 1). It seems that robots of unsuccessful groups are not capable of informing each other about the colour of the painted zone in the proximity of L_1 and L_2 . Consequently, in the absence of an effective communication protocol, it turns out to be impossible for M_C to properly set S_C . In the following paragraphs, we analyse the communication protocol used by a successful group.

Fig. 4a illustrates the structures of signalling behaviour of the successful group g_4 . In this post-evaluation test, the group undergoes 4 trials with the environment presented in the following sequence: E_{10} , E_{01} , E_{00} , and E_{11} . In each trial the robots don’t emit sound before reaching the coloured zones. The

perception of grey doesn't induce the emission of sound. Therefore, in E_{00} no robots emit sound (see Fig. 4 trial 3). The absence of sound in the environment lets M_C set S_C to 0 in both robots. $S_C = 0$ is correctly "interpreted" by M_M modules so that both robots push the revolving door clockwise. The perception of a black zone induces the robots to emit intermittent bursts of sound (see Fig. 4a trials 1, 2 and 4). In trials E_{10} and E_{01} , the perception of these intermittent bursts induces the robot that is on grey to emit a continuous tone. The perception of a continuous tone induces the robot on black to imitate its fellow, so that at the time when L_3 turns on (see Fig. 4b, dotted line) both robots emit a continuous tone. The presence of sound in the environment lets M_C set S_C to 1 in both robots. $S_C = 1$ is correctly "interpreted" by M_M modules so that both robots push the revolving door anticlockwise. Both robots autonomously stop emitting sound before the end of a trial in E_{10} or E_{01} , few seconds after the aperture of the revolving door. Thus, at the beginning of the following trial both robots are in the state of not emitting sound. In trials E_{11} , the asynchronous emission of intermittent bursts of sound by both robots determines moments of silence which inhibit signalling behaviour. At the time when L_3 turns on, none of the robots is signalling. The absence of sound in the environment lets M_C set S_C to 0 in both robots. $S_C = 0$ is correctly "interpreted" by M_M modules so that both robots push the revolving door clockwise as in E_{00} .

6 Conclusions

In this paper, we described a model in which artificial evolution is employed to design neural mechanisms that control the motion of autonomous robots required to communicate through sound to perform a common task. The results of this work are a "proof-of-concept": they demonstrate that evolution can exploit a simple sound system, detailed in Sec. 3, to design the mechanisms that allow two robots cooperate by using bi-directional acoustic interactions. Post-evaluation tests illustrate the nature of the robots' communication protocol based on entirely evolved asynchronous and mutually determined single tone signals. Concerning future work, we believe that priority should be given to investigations aimed to limit the amount of *a priori* assumptions that we have been forced to make in this first study. In particular, we are referring to the modularisation of the control structures and the arbitrary associations detailed in Sec. 3. In spite of this, we believe that the results are particularly encouraging. A complex syntax may emerge in scenarios in which semantic categories are linked to more articulated sensory-motor structures (e.g., neural structures that underpin object recognition processes rather than the perception of coloured zones).

Acknowledgements. This research work was supported by the *ECAgents* project (grant IST-1940), the *SWARMANOID* project (grant IST-022888), and 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. The authors thank Carlo Pinciroli and their colleagues at IRIDIA for stimulating discussions and feedback during the preparation of this paper.

References

1. Ampatzis, C., Tuci, E., Trianni, V., Dorigo, M.: Evolution of signalling in a group of robots controlled by dynamic neural networks. In: Sahin, E., Spears, W.M., Winfield, A.F.T. (eds.) Proc. 2nd Int. Workshop on Swarm robotics, Berlin, Germany, vol. 4433, pp. 173–188. Springer, Heidelberg (2006)
2. Beer, R.D., Gallagher, J.C.: Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior* 1(1), 91–122 (1992)
3. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading (1989)
4. Marocco, D., Nolfi, S.: Origins of communication in evolving robots. In: Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.-A., Miglino, O., Parisi, D. (eds.) SAB 2006. LNCS (LNAI), vol. 4095, pp. 789–803. Springer, Heidelberg (2006)
5. Mondada, F., Pettinaro, G.C., Guignard, A., Kwee, I.V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L.M., Dorigo, M.: SWARM-BOT: A new distributed robotic concept. *Autonomous Robots* 17(2–3), 193–221 (2004)
6. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge (2000)
7. Di Paolo, E.: Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior* 8(1), 27–48 (2000)
8. Trianni, V., Dorigo, M.: Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics* 95, 213–231 (2006)
9. Tuci, E., Ampatzis, C., Vicentini, F., Dorigo, M.: Evolved homogeneous neuro-controllers for robots with different sensory capabilities: coordinated motion and cooperation. In: Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.-A., Miglino, O., Parisi, D. (eds.) SAB 2006. LNCS (LNAI), vol. 4095, pp. 679–690. Springer, Heidelberg (2006)
10. Wischmann, S., Pasemann, F.: The emergence of communication by evolving dynamical systems. In: Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.-A., Miglino, O., Parisi, D. (eds.) SAB 2006. LNCS (LNAI), vol. 4095, pp. 777–788. Springer, Heidelberg (2006)