

Origins of communication in evolving robots

Davide Marocco, Stefano Nolfi

Institute of Cognitive Science and Technologies, CNR, Via San Martino della Battaglia 44,
Rome, 00185, Italy
{davide.marocco, stefano.nolfi}@istc.cnr.it

Abstract. In this paper we describe how a population of simulated robots evolved for the ability to solve a collective navigation problem develop individual and social/communication skills. In particular, we analyze the evolutionary origins of motor and signaling behaviors. Obtained results indicate that signals and the meaning of the signals produced by evolved robots are grounded not only on the robots sensory-motor system but also on robots' behavioral capabilities previously acquired. Moreover, the analysis of the co-evolution of robots individual and communicative abilities indicate how innovation in the former might create the adaptive basis for further innovations in the latter and vice versa.

1 Introduction

The development of embodied agents able to interact autonomously with the physical world and to communicate on the basis of a self-organizing communication system is a new exciting field of research ([13], [1], [10], [3], [9], for a review see [2], [11], [14] and [7]). The objective is that to identify methods of how a population of agents equipped with a sensory-motor system and a cognitive apparatus can develop a grounded communication system and use their communication abilities to solve a given problem. These self-organizing communication systems may have characteristics similar to that observed in animal communication [5] or human language.

In this paper we describe how a population of simulated robots evolved for the ability to solve a collective navigation problem develop individual and social/communication skills. In particular, we analyze the evolutionary origins of motor and signaling behaviors. Obtained results indicate that the signals and the meaning of signals produced by evolved robots are grounded not only on robots sensory-motor system but also on robots' behavioral capabilities previously acquired. Moreover, the analysis of the co-adaptation of robots individual and communicative abilities indicate how innovations in the former might create the adaptive basis for further innovations in the latter and vice versa.

In the next section we describe the experimental setup (for more details on the experiments and on the characteristic of the communication system at the end of the evolutionary process, see [4]). In section 3, we describe the evolutionary origin of the communication system used by evolved robots. Finally, in section 4, we summarize the main results and we briefly discuss the implications of these experiments.

2 The experimental set-up

A team of four simulated robots placed in an arena of 270x270cm (Fig.1, Left) are evolved for the ability to find and remain in the two target areas by equally dividing between the two targets. Robots communicate by producing and detecting signals up to a distance of 100cm. A signal is a real number with a value ranging between [0.0, 1.0].

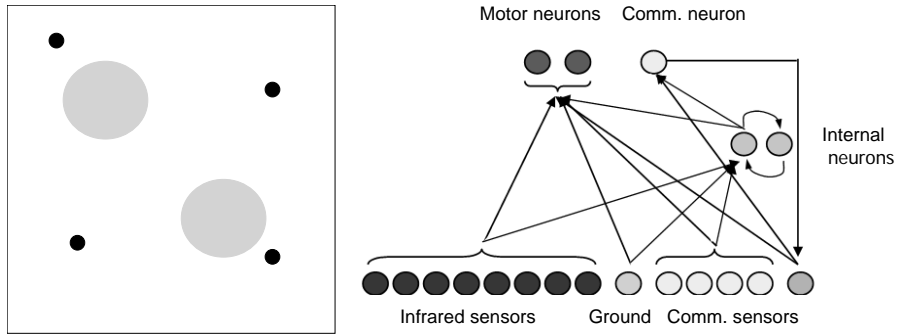


Fig. 1. Left: The environment and the robots. The square represents the arena surrounded by walls. The two gray circles represent two target areas. The four black circles represent four robots. Right: The neural controller of evolving robots.

Robots' neural controllers (Fig. 1, Right) consist of neural networks with 14 sensory neurons that encode the activation states of 8 infrared sensors, 1 ground sensor (that binarily encodes the color of the ground), 4 communicative sensors (that encode the value of the signals produced by other robots from four corresponding orthogonal directions (i.e. frontal [315°-44°], rear [135°-224°], left [225°-314°], right [45°-134°]), and the activation state of the communication neuron at times $t-1$ (i.e. each robot can hear its own emitted signal at the previous time step). These sensory neurons are directly connected to the three motor neurons that control the desired speed of the two wheels and the value of the communication signal produced by the robot. The neural controllers also include two internal neurons that receive connections from the sensory neurons and from themselves and send connections to the motor and communicating neurons [8]. The three motor neurons encode the desired speed of the two wheels of the robot and the value of the signal emitted by the robot.

The output of motor neurons is computed according to the logistic function (2), the output of sensory and internal neurons is computed according to function (3) and (4), respectively (for more details on these activation functions and on the relation with other related neural models see [6]).

$$A_j = t_j + \sum_i w_{ij} O_i \quad (1)$$

$$O_j = \frac{1}{1 + e^{-A_j}} \quad (2)$$

$$O_j = O_j^{(t-1)} \tau_j + I_j (1 - \tau_j) \quad (3)$$

$$O_j = O_j^{(t-1)} \tau_j + (1 + e^{-A_j})^{-1} (1 - \tau_j) \quad (4)$$

With A_j being the activity of the j_{th} neuron, t_j being the bias of the j_{th} neuron, w_{ij} the weight of the incoming connections from the i_{th} to the j_{th} neuron, O_i the output of the i_{th} neuron, $O_{j(t-1)}$ being the output of the j_{th} neuron at the previous time step, τ_j the time constant of the j_{th} neuron, and I_j the activity of the j_{th} sensors.

The free parameters of the robots' neural controllers have been evolved through a genetic algorithm. Each team of four robots was allowed to "live" for 20 trials (each trial lasting 100 seconds, i.e. 1000 lifecycles of 100 ms each). At the beginning of each trial the position and the orientation of the robots was randomly assigned outside the target areas. The fitness of the team of robots consists of the sum of 0.25 scores for each robot located in a target area and a score of -1.00 for each extra robot (i.e. each robot exceeding the maximum number of two) located in a target area. The total fitness of a team is computed by summing the fitness gathered by the four robots in each time step.

The initial population consisted of 100 randomly generated genotypes that encoded the connection weights, the biases, and the time constants of 100 corresponding neural controllers. Each parameter was encoded with 8 bits and normalized in the range $[-5.0, +5.0]$, in the case of connection weights and biases, and in the range $[0.0, 1.0]$, in the case of time constants. Each genotype was translated into 4 identical neural controllers that were embodied in the four corresponding robots, i.e. teams were homogeneous and consisted of four identical robots. For a discussion about this point and alternative selection schemas see [7]. The 20 best genotypes of each generation were allowed to reproduce by generating five copies each, with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasted 2000 generations (i.e. the process of testing, selecting and reproducing robots is iterated 2000 times). The experiment was replicated 10 times starting by 10 different initial populations.

By analyzing the fitness thorough out generations we observed that evolving robots are able to accomplish their task to a good extent in all replications from generation 500 on (evolving robots are able to find and remain in the two target areas by equally dividing between the two areas in 58.3% of the trials). Further increases of performance observed from generation 500 on, are due to slight improvements with respect to the ability to solve the task faster (the average time required by the four robots of all replications to reach the two target areas goes from 74s to 67s in generation 500 and 2000, respectively) and better (the percentage of trials in which the task is solved correctly increase from 58.3% to 67.5%, in generation 500 and 2000 respectively).

By comparing these results with the results obtained in a control condition in which robots were not allowed to detect signals (i.e. in which the state of the communication sensors was always set to 0.0) we observed that, in all replications, the fitness reach a stable state after 150 generations, which is significantly lower than the case in which robots are allowed to communicate (i.e. robots are able to solve the problem only in 36.7% of the trials after 2000 generations).

The comparison between the results obtained in the normal and in the control condition in which robots are not allowed to detect other robots' signals indicates how the possibility to produce and detect other robots' signals is necessary to achieve optimal or close to optimal performance.

In the next subsection we will analyze the evolutionary origins of robots ability to solve their task and of the communication system displayed by evolved individuals.

3 Origins and evolution of a self-organized communication system

To understand the evolutionary origins of robots' communication system we analyzed the motor and signaling behavior of evolving robots through out generations. To reconstruct the chain of variations that led to the final evolved behavior we analyzed, for each replication, the lineage of the best individual of the last generation (i.e. the 1999 individuals, one for each generation, that constitute the ancestors of the best individual of generation 2000). Below we report the results of this analysis by focusing in particular on the best replication of the experiment. The analysis of the other replications of the experiment (not shown) produced qualitatively similar results (although the values of the signals serving a given function and the length of different evolutionary phases vary significantly).

As shown in Fig. 2 and Fig. 3, in the case of the best replication of the experiment, the fitness quickly increases by reaching high level performance during the first 50 generations (the team of robots of generation 50 is able to solve the problem in 64% of the trials and the average time required to reach the two target areas is 65.4s). From generation 50 to generation 1700, the fitness remains rather stable, beside small and unstable increases in performance. From about generation 1700 on, performance stabilizes again on a slightly higher value with respect to previous generations (the team of robots of generation 2000 are able to solve the problem in 79% of the trials and the average time required to reach the two target areas is 57.6s).

By analyzing the motor and signaling behavior through out generations in the case of the best replication of the experiment (the same replication shown in Fig 2 and 3) we observed the following phases:

Generation 1. At this stage robots move in the environment by producing curvilinear trajectories and by avoiding obstacles (in most of the cases). Robots produce two stable signals with a value of 0.53 and 0.33 when they are located inside or outside a target area, respectively, and far from other robots. Moreover, robots produce highly variable signals when they interact with other robots located nearby.

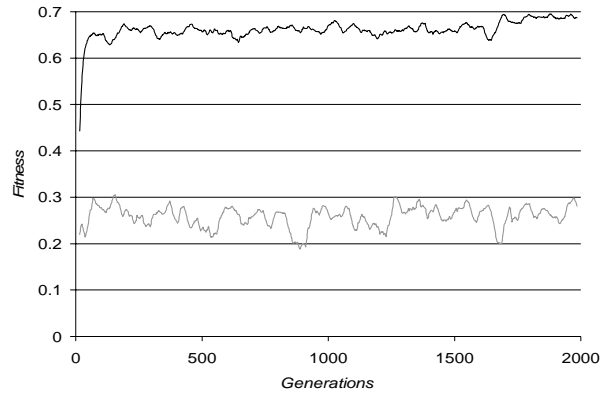


Fig. 2. Fitness of the lineage of the best individual of generation 2000 through out generations in the case of the most successful replication of the experiment. The black and gray lines represent the performance in a normal and no-signal condition (in which robots are not allowed to detect other robots' signals). Lines indicate the moving average over 30 generations. Each individual have been tested for 100 trials.

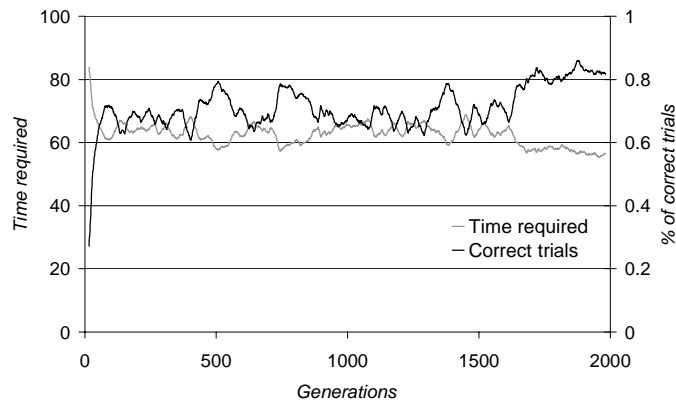


Fig. 3. Percentage of trials in which robots accomplish the task successfully (within 100s) and average time required by the robots to reach the target areas by equally dividing between the two areas throughout generations. Average performance obtained by testing each team for 100 trials.

In particular, when a robot located outside a target area starts to detect the signal emitted by another robot, it modifies the signal produced by a stable signal with a value of about 0.33 (a signal that we will call **A** that is produced by robots that do not detect signals produced by other robots) to an highly variable signal with an average value of 0.28 (a signal **B** that is produced by robots detecting the signal **A** or **B** produced by another robot). Signal **B** increases robots' exploratory abilities (i.e. the probability to reach target areas). Indeed, by testing the robots in a normal condition and in a control condition in which they are not allowed to produce the signal **B**, we

observed that the average time spent by the robots to reach a target area for the first time is 58.8s and 70.4s, in the normal and control condition respectively. Therefore, the functionality of signal **A** is that to trigger the production of signal **B**. The functionality of signal **B** is that to increase robots navigation ability, as described above.

We do not assign an identification letter to the other signals produced by robots at this stage since these signals does not seem to have any clear adaptive function. Some of these signals however, for example an highly varying signal (average value 0.45) and a stable signal (with a value of about 0.53) produced by robots located in a target area interacting or not interacting with other robots located nearby will acquire a functional role in successive generations.

On the basis of these individual and social behaviors (i.e. an individual obstacle avoidance behavior, individual exploration behavior, and a social behavior based on signals that alters other robots' trajectories in a way that enhances their chance to reach target areas) robots are able to spend almost half of their lifetime on target areas (Fig. 4). A typical behavior observed at this stage is shown in Fig.5 (Gen. 1).

Generations 2-7. During this phase robots progressively evolve an individual ability to remain in target areas. Indeed, at generation 7, robots located on target areas rotate on the spot so to remain there for the rest of the trial. Moreover, robots produce several differentiated signals. However, as in the previous phase, only two of these signals have an adaptive function.

As in the case of generation 1, robots located outside target areas produce a signal **A** (a signal with a value of about 0.34 produced by robots located far from other robots), and a signal **B** (a varying signal with an average value of 0.24 produced by robots interacting with other robots located outside target area). We keep the same labels introduced above since, although the value and the effect of the signals slightly varied, the functionality of the signals is very similar to that of the signals described in the previous section. As for generation 1, the functionality of signal **A** is that to trigger the production of signal **B** and the functionality of signal **B** is that to increase robots' ability to reach target areas. Moreover, as in the case of generation 1, robots located on target areas produce two non-adaptive signals: (1) an highly varying signal with an average value of about 0.73 (produced by robots that interact with other robots located nearby outside target areas), and (2) a stable signal with a value of about 0.82 (produced by robots that do not detect signals produced by other robots). These two signals do not have any adaptive function, but rather produce a decrease in robots' performance. Indeed, the production of these two signals reduce the chances that robots located outside target area join target areas that already contain a single robot.

As a result of the newly developed individual behavior that allows robots to remain on target areas, however, the percentage of lifecycles in which one or two robots are located on a target area increases from 35% to 45% and from 10% to 22%, respectively (see Fig. 4). A typical behavior observed at this stage is shown in Fig.5 (Gen. 7).

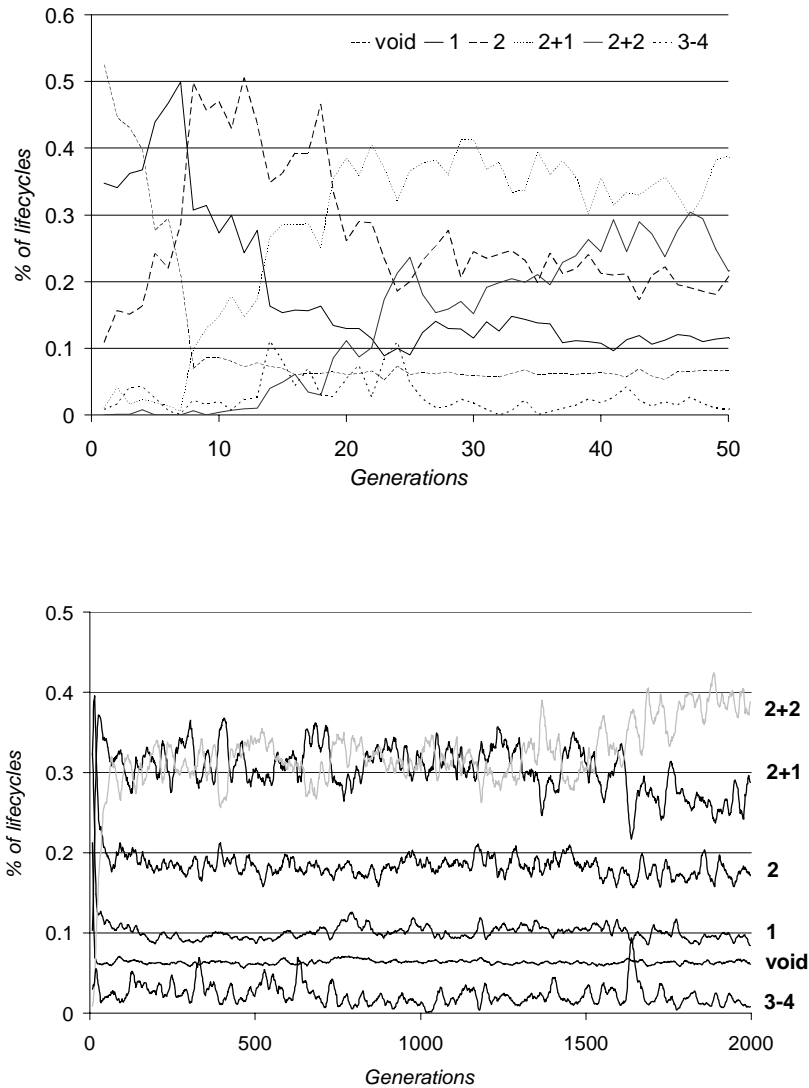


Fig. 4. Percentage of lifecycles spent by a team of robots in 6 possible states: void = all robots are outside target areas, 1 = only one robot is located in a target area, 2 = two robots are located in target areas (either in the same or in two different areas), 2+1 = 3 robots are located inside two different target areas, 2+2 = all robots are located inside the target area equally divided between the two areas, 3-4 = 3 or 4 robots are located in the same target area. The data refer to the lineage of the best individual of the last generation for the best replication of the experiment. Each robot have been tested for 100 trials lasting 1000 lifecycles. Top graph: data up to generation 50. Bottom graph: data from generation 0 to 2000.

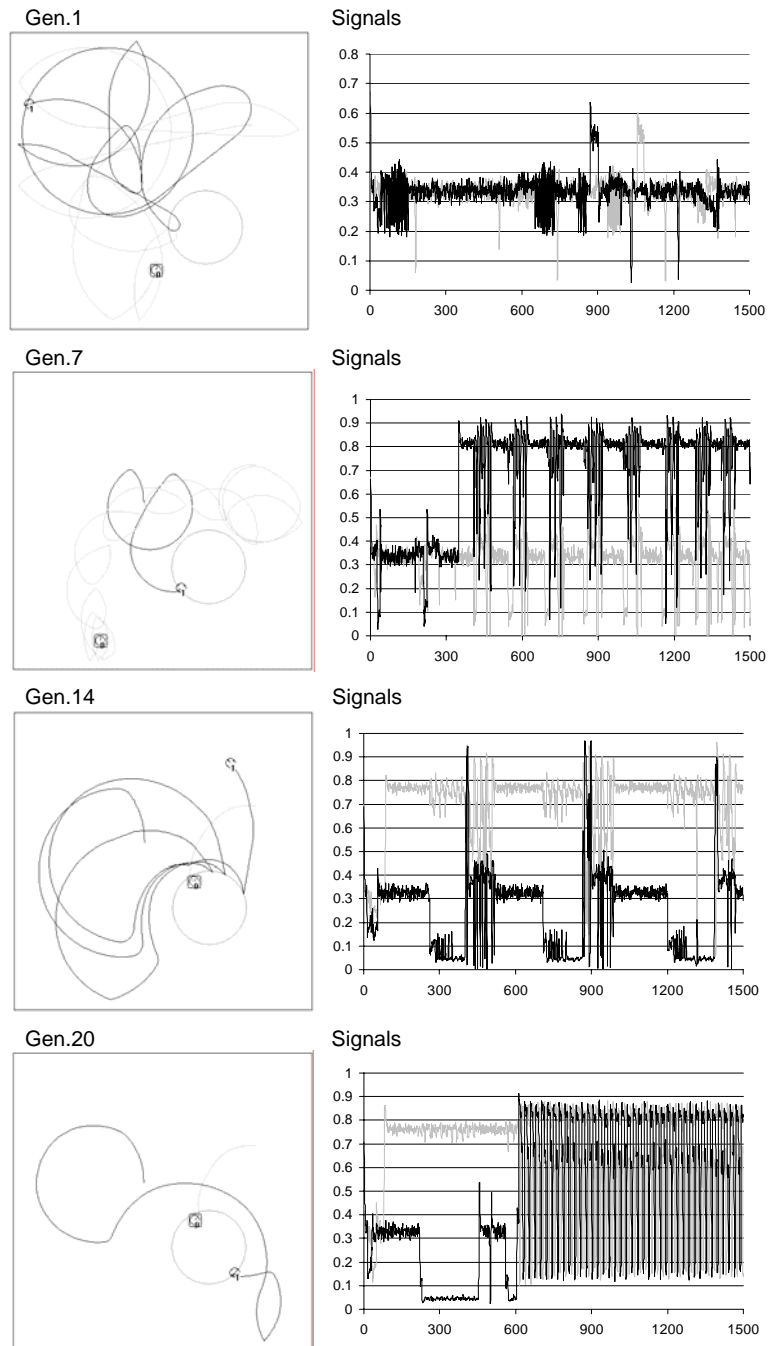


Fig. 5. Motor and signaling behavior observed at different generations. Left: the trajectory produced by two robots tested in an environment including a single target area. Right: the

signals produced by the two robots during the test shown in the left part of the figure. The motor trajectory and the signal of the first and of the second robot are shown with black and gray lines, respectively.

Generations 8-14. The development of an individual ability to remain on target areas developed in previous generations posed the adaptive basis for the development of a cooperative behavior that allows robots located on a target area alone to attract other robots toward the same target area. As we said in the previous section, the highly varying signal produced at generation 7 by robots located inside a target area interacting with other robots located outside the area reduced the chances that the latter robots join the area. At generation 14, however, this highly varying signal is not produced anymore. This innovation results from the fact that robots located outside target areas interacting with robots located inside target areas now produce signal **D**, (i.e. a signal with a value of 0.04). Since producing a signal with an almost null value is equivalent to stop signaling, the production of signal **D** implies that robots located inside a target area alone now produce a signal **C** (a stable signal with a value of about 0.78) independently of whether they interact or not with another robot located nearby outside the target area. Since the signal **C**, produced by a robot located inside a target area, increases the chances that other robots will enter in the same target area, the innovation that allows robots located outside the target area to switch their signaling behavior off (as soon as they detect signal **C**) produces a significant adaptive advantage.

To summarize, during this phase robots develop an ability to produce a new signal (signal **D**) whose functionality is that to allow robot located inside target area to keep producing signal **C** even when other robots are located nearby. This in turn allows robots to exploit the effect of signal **C**, that consists in attracting other robots toward the source of the signal (i.e. toward the corresponding target area). This effect of signal **C** on other robots motor behavior already existed in previous generations. However it could not be exploited since robots located in target area were able to produce signal **C** only when no other robots were located in the communicative range.

The acquisition of an ability to switch signaling behavior off leads to a specialization of the role of the two interacting robots since, in these situations, the robot located in the target area and producing the signal **C** acts as a speaker and the robot located outside the target area producing signal **D**, acts as a hearer. The social interaction between the two robots in this circumstance, therefore, can be described as a form of information exchange (in which a speaker robot located inside a target area informs the hearer robot on the location of the target area and in which the hearer robot reacts to the signal by moving toward the direction of the area) or as a form of manipulation (in which the speaker robot drives the hearer robot toward the target area by exploiting the tendency of the hearer robot to alter its motor trajectory as a result of a detected signal).

At this stage, robots are not still able to remain in a target area in couple (see Fig. 5, Gen. 14). In fact, as soon as a second robot reaches a target area, the two robots start to produce two different signals (i.e. two highly varying signals with an average

value of 0.63 for the former and 0.38 for the latter robot) that are maladaptive since they increase the chances that one of the two robots abandons the area.

As a result of the innovations occurring during this phase (that mainly consist in the variations that leads to the production of signal **D**) the percentage of lifecycles in which two and three robots are located on a target area increases from 22% to 50% and from 0% to 18% (Fig. 4).

Generations 15-20. The development of an ability to attract nearby robots toward target areas that contain a single robot described in the previous section leads to an increase in performance but also poses new adaptive opportunities, namely the need to develop an ability to remain into target areas that contain a single robot and the need to produce a signal that keep other robots away from a target area that contains two robots. These two problems are solved in this phase through variations that allow robots to not exit from target areas when they detect the signal produced by another robot located in the same target area. This is achieved through the development of a new signal **E** (an highly varying signal with an average value of 0.61 produced by robots located in a target area that contains two robots). Signal **E** plays two adaptive functions: (1) it does not push the other robot located in the target area out (unlike the signals previously produced in this circumstance), and (2) it reduces the chances that other robots located outside the target area will join the area itself. Interestingly, signal **E** (i.e. the signal produced by two interacting robots located in the same target area) allows the two robots to generate an information (i.e. that encode the fact that the area contains two robots) that is not directly available to none of the two robots.

Generations 21-1700. During this long evolutionary phase the performances of the robots, the number of signals, and the functionalities of signals remain rather stable. Evolving robots display close to optimal performance, few simple but crucial individual behaviors (that allow the robots to explore the environment, avoid obstacles, and remain into target areas) and an effective communication system that now includes 5 signals (i.e. signals **A**, **B**, **C**, **D**, and **E** described in previous sections) that modulate the robots' behavior by producing an enhanced exploratory behavior, a target approaching behavior, and a target avoidance behavior. Since each of these individual and communication abilities provides a clear adaptive advantage, all of them are preserved during the rest of the evolutionary process.

Despite of that, some characteristics of the individual behavior exhibited by the robots, the value of the signals serving a given function, and the impact of signals on other robots' behavior vary significantly.

Variations of individual behaviors mainly concern how robots explore the environment while they do not detect signals produced by other robots. This fact can be explained by considering that robots' ability to find target areas on their own plays a limited adaptive value at this stage in which individuals possess a reliable ability to find target areas by exploiting the signals produced by other robots. Variation on individual behavioral abilities, however, can be tolerated only within limits. To illustrate this point let us consider how robots' individual exploratory behavior varies during this phase. As we reported above, robots located outside target areas tend to produce a curvilinear trajectory and to avoid obstacles. The combination of these two behaviors allow the robots to explore different parts of the environment and to encounter target areas relatively quickly. The turning angle with which robots move

forward, however, should be sufficiently large so to avoid turning on the same position indefinitely. The turning angle of the robots in this circumstance is indeed a character that is subjected to significant variations until a certain threshold is reached. Variations that overcome the threshold tend to be maladaptive since they lead to robots that are unable to explore the environment without the help of other robots (as shown in Fig. 6, Gen. 225). However, their negative effects only manifest in robots that do not receive the necessary social help during their lifetime. As a consequence, these variations might be retained and might cause a drop in performance in successive generations until characters similar to those previously lost are restored (for an example, see Fig. 6, Gen. 226-230). This analysis illustrates how individual behavior, such as individual abilities to explore the environment, does not only poses the evolutionary basis for the emergence of the communication system, but still plays a fundamental role when the communication system is established. This individual behavior, in fact, also constitutes a pre-requisites for the ability of the robots to collect information to be communicated or to create the conditions for receiving useful signals.

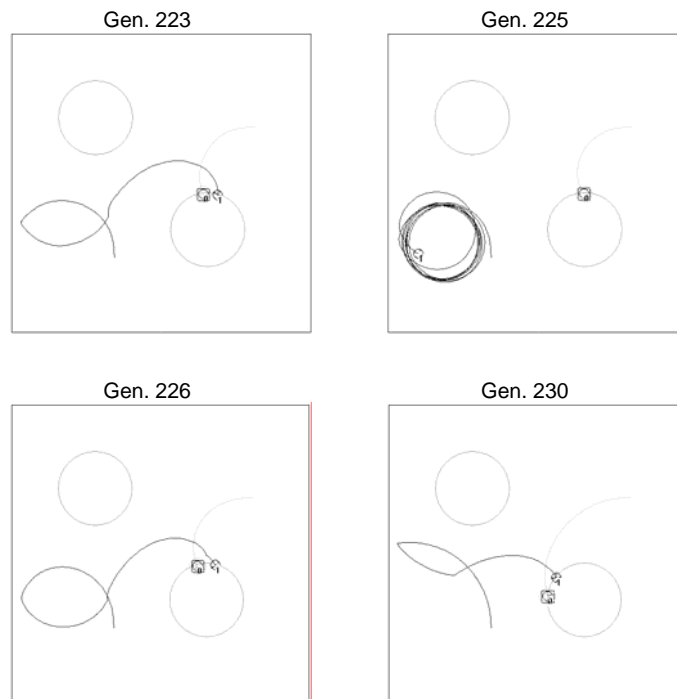


Fig. 6. Behavior exhibited by robots of different generations. For space reasons only the behavior produced by two robots of different generation is displayed. As can be seen, at generation 225 robots lose their ability to explore the environment and keep circling in the same area. The exploration ability is recovered in successive generations.

Other characteristics that significantly vary during this phase are the value of the signals and the way in which signals affect robots behavior. Although the functionality of the five signals described above remains rather constant during this phase, the value associated to each signal significantly vary (Fig. 7). This fact can be explained by considering that the functionality of a signal depend both on the value of the signal and the effect that the signal produces on robots. The possibility to co-adapt the value of signals and the impact of a signal on robots' motor and signaling behavior, ensures that the functionality can be preserved while the signals and their effects co-vary.

In principle, these neutral variations could lead to new organizations of the communication system, that might represent a pre-requisite for further innovations of individual and communicative abilities. Some preliminary evidences suggest that this is indeed one of the reasons that explain the evolutionary transition that leads to better behaviors in the next phase (see below). This evidences however are only preliminary and should be integrated with further analysis that we plan to conduct in the future.

Generations 1700-2000. After a long phase in which performances remain rather stable, a small but stable increase in performance is observed from generation 1700 on (Fig. 2). Indeed, as shown in Fig. 3, the percentage of times in which the robots are able to accomplish their task correctly increase from 62% to 79% and the average time required for solving the problem decreases from 67s to 57s during this evolutionary phase.

The evolutionary transition that leads to this improvement involves a significant re-organization of the values of the five signals (see Fig. 7). Although the number and the general functionality of the signals remains the same, from generation 1700 on, the values of the five different signals are distributed on a wider range and the value of each signal is distributed on a smaller range with respect to previous generations.

Other variations occurring in this phase might affect the way in which signals are exploited. In particular, from generation 1700 on robots often display: (a) an enhanced ability to avoid target areas that already contains two robots without remaining plugged into unfruitful conditions, (b) an ability to reach the target areas faster by taking the risk to end up in a target area that already contains two robots but by also being able to exit from these areas, (c) an enhanced ability to negotiate situations in which robots concurrently receive signals from several robots. However, further analysis should be conducted to clarify the nature and the adaptive role of the innovations occurring during this phase.

4 Conclusion

In this paper we described how a population of simulated robots, evolved for the ability to solve a collective navigation problem, develop an effective communication system. By analyzing the evolutionary origins of motor and signaling behaviors we observed that the co-adaptation of robots' motor and communicative abilities plays a crucial role on the evolutionary dynamic.

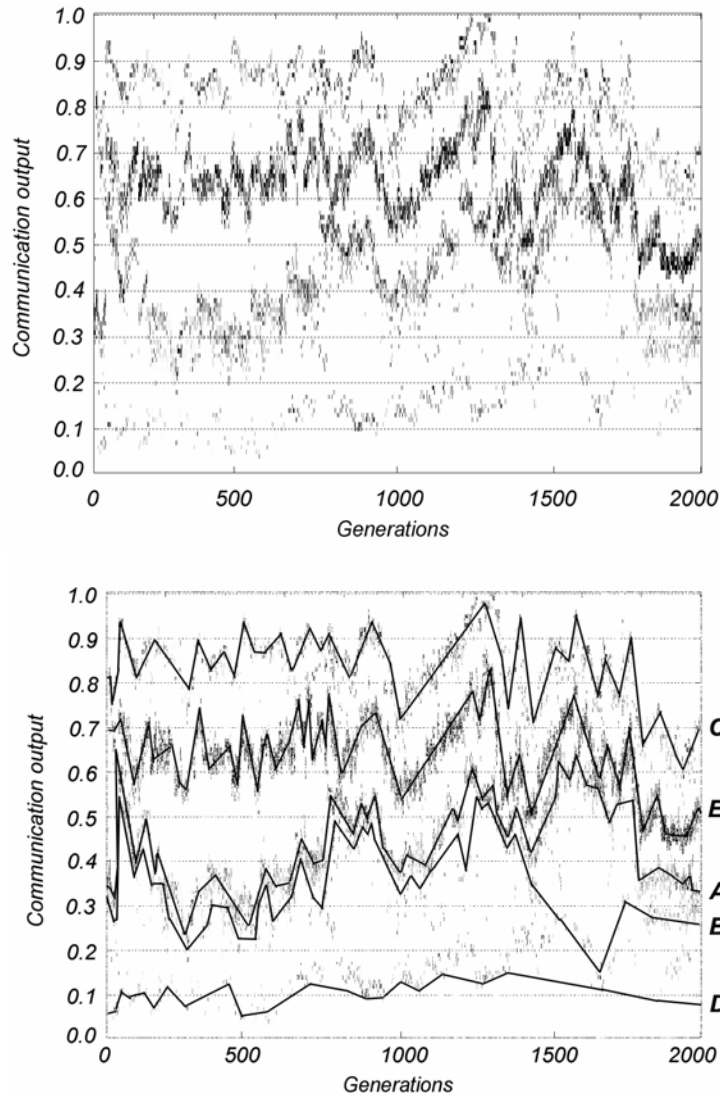


Fig. 7. The value of the signals produced by robots throughout generations. Each point represents the mean value of each signal. The gray scale of each point indicates the variance of the signal with respect of the mean value (i.e. the darkness of each point is proportional to the variability of the corresponding signal). The data displayed on the graph have been obtained by filtering out signals that are produced only occasionally. Oscillatory signals have been identified through a wavelet analysis. The bottom figure displays the same data of the top figure with a superimposed schematization of signals average value through out generations.

In some cases the development of new motor skills poses the basis for the successive development of new social abilities. For instance, the development of an ability to remain in target areas constitutes a pre-requisite for the development of an ability to communicate the location of the area to other robots so to increase the chances that other robots will join the same target area. In other cases, the development of social/communication abilities pose the basis for the development of new motor skills. For instance the development of an ability to detect the number of robots located in a target area through bi-directional signaling interactions creates the basis for the development of an effective avoidance behavior that allow robots to avoid entering in crowded target areas and to look for another target areas.

Interestingly the co-adaptation process of motor and social/communicative abilities may potentially lead to open-ended evolutionary dynamics in which innovations create the adaptive basis for further innovations thus leading to a progressive increase in performance and to a progressive complexification of agents abilities. Indeed, while during the first phase of the evolutionary experiment robots can only rely on few environmental cues (that provide information on whether they are located on a target area or not and whether they are close to obstacles), in later generation they can exploit a much larger number of cues (that, for example, provide information also on the location of target areas and on the number of robots located in target areas).

Finally, we observed how the complexification of robots' motor and social skills involve different aspects, and can be characterized along several dimensions: (a) an increase in the number of elementary behaviors exhibited by the robots, (b) an increase in the number of signals produced by robots, (c) an increase in the number of ways in which the same signal affect robots' behaviors in different contexts, (d) a differentiation of the modalities with which communication is regulated (e.g. the transformation of symmetrical interaction forms in which communicating robots act concurrently as speakers and hearers to specialized asymmetrical interaction forms in which one robot acts as a speaker and one robot acts as an hearer).

From a scientific point of view, these types of experiments and results can allow us to understand better how 'meanings' originate and how signals are grounded in agents sensory-motor and behavioral abilities. From an application point of view, these methods can allow us to develop a new generation of artifacts able to solve practical problems by cooperating and communicating on the basis of a self-organized communication system.

Acknowledgments

The research has been supported by the ECAGENTS project funded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract IST-1940.

References

1. Cangelosi A., Parisi D.: The emergence of a 'language' in an evolving population of neural networks. *Connection Science*. Taylor & Francis (1998) 10: 83-97.
2. Kirby S.: Natural Language from Artificial Life. *Artificial Life*. The MIT Press, Cambridge (2002) 8(2):185-215.
3. Marocco D., Cangelosi A., Nolfi S.: The emergence of communication in evolutionary robots. *Philosophical Transactions of the Royal Society A*, London (2003) 361: 2397-2421.
4. Marocco D., Nolfi S.: Emergence of Communication in Embodied Agents Evolved for the Ability to Solve a Collective Navigation Problem. *Connection Science*. Taylor & Francis (in press).
5. Marocco D., Nolfi S.: Communication in Natural and Artificial Organisms. Experiments in evolutionary robotics. In: Lyon C., Nehaniv C., Cangelosi A. (eds.): *Emergence of Communication and Language*, Springer-Verlag (in press).
6. Nolfi S.: Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science*. Taylor & Francis (2002) (14) 3:231-244.
7. Nolfi S.: Emergence of Communication in Embodied Agents: Co-Adapting Communicative and Non-Communicative Behaviours. *Connection Science*. Taylor & Francis (2005) 17 (3-4): 231-248.
8. Nolfi S., Marocco D.: Evolving Robots Able To Integrate Sensory-Motor Information Over Time. *Theory in Biosciences*. Urban & Fischer Verlag, Berlin, (2002) 120: 287-310.
9. Quinn M., Smith L., Mayley G., Husband P.: Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society A*, London (2003) 361: 2321-2344.
10. Steels L. The Talking Heads Experiment. Antwerpen, Laboratorium (1999). Limited Pre-edition.
11. Steels L. Evolving grounded communication for robots. *Trends in Cognitive Science*, Elsevier (2003) 7(7): 308-312.
12. Steels L., Kaplan F.: AIBO's first words: The social learning of language and meaning. *Evolution of Communication*. John Benjamins (2001) 4: 3-32.
13. Steels L., Vogt P.: Grounding adaptive language games in robotic agents. In: Husband P., Harvey I. (eds.): *Proceedings of the 4th European Conference on Artificial Life*. The MIT Press, Cambridge (1997).
14. Wagner K., Reggia J.A., Uriagereka J., Wilkinson G.S.: Progress in the simulation of emergent communication and language. *Adaptive Behavior* Sage Publications (2003), 11(1): 37-69.