EMERGENCE OF COMMUNICATION IN EMBODIED AGENTS: CO-ADAPTING COMMUNICATIVE AND NON-COMMUNICATIVE BEHAVIOURS

DAVIDE MAROCCO STEFANO NOLFI

Institute of Cognitive Science and Technologies, CNR, Viale Marx 15
Rome, 00137, Italy
[davide.marocco; s.nolfi]@istc.cnr.it

We show how a population of simulated robots developed their communication capabilities in order to solve a collective navigation problem. The self-organized emergent vocabulary includes four different signals that influence both the motor and signalling behaviour of other robots. The analysis of the evolved behaviours also indicates: (a) the emergence of a simple form of communication protocol that allows individuals to switch signalling on and off, (b) the emergence of tightly co-adapted communicative and non-communicative behaviours, and (c) the exploitation of properties resulting from the dynamical interactions between motor and signalling behaviours produced by interacting robots.

1. Introduction

Existing models of emergence of communication in embodied agents often focus on how a shared communication system and/or communication ability might emerge in a population of interacting individuals (Cangelosi, 2001; Cangelosi & Parisi, 2002; Kirby, 2001; Steels, 1999; Steels and Kaplan, 2001). In contrast, this paper focuses on the more general question of how a population of agents that have to solve a given adaptive problem might develop their ability to display complex forms of interactions and communications that enhance their adaptive capability (Quinn, 2001; Quinn et al. 2003; Baldassarre et al., 2003).

The reason for this choice is twofold. The theoretical motivation is that communication and communication systems are adaptive capabilities shaped by their function. What, when and how individuals communicate (and whether individuals do or do not communicate) depends on the adaptive function of communication. Similarly the type of communication system that might self-organize in a population of interacting agents will strongly depend on the type of behaviour individuals display in isolation and on the complementary functions that interactions and communications might have. The underlying assumption is that communication and language can be properly understood by considering their relation with other important behavioural, social and cognitive processes. The practical motivation is that, from an application point of view, the possibility of developing embodied agents able to solve real life problems by exploiting complex forms of interaction and communication might have huge application potentials.

In this perspective, three additional aspects play a crucial role. In fact, we are interested in models that (a) allow the emergence of a communication ability and a shared communication system and (b) that allow the discovery of categories (or coupled internal/external dynamical processes) that are useful from the communication and cognitive point of view but not explicitly or implicitly identified in the experimental set. This claim is based on the theoretical assumption that communication abilities might enhance the cognitive/adaptive abilities of interacting agents for two main reasons: (1) the fact that the indirect adaptive advantages of communication might force the development of useful and compact ways to categorize the continuous flow of sensory-motor information. (2) The fact that the discoveries of useful categories might be easier through social interactions as opposed to being in isolation.

We are interested in models in which individuals, in addition to having signalling and interaction capabilities, also have a reach sensory and motor non-communicative repertoire that might allow them to improve their ability to solve their cognitive/adaptive problems by improving both their individual and their social/communication capabilities. This claim is based on the assumption that only by co-adapting their behavioural non-communicative and communicative abilities, individuals might develop a really useful communication system grounded in the physical and behavioural characteristics of communicating individuals and furthermore, be able to exploit active perceptual capabilities. Moreover, this claim is based on the assumption that one of the key aspects of communication is the possible reliance of the individual on implicit information that does not need to be communicated.

Finally we are interested in models in which forms of communication of different complexity might be used. By forms of communication we refer to the protocol with which individuals interact during communication and to the way with which communication signals are organized. Forms of communication might range from simple continuous broadcasted signalling to complex regulated communication protocols. Furthermore, communication acts might be episodic and asynchronous or, for example, communication protocols might be negotiated on the fly between the two communicating agents and communication acts might consist of a sequence of signals organized according to a grammar. This claim is based on the assumption that more complex forms of communication are not effective in general terms. Therefore, agents should have the liberty, where possible, to select the communication forms that are most useful to them, given their current behavioural/cognitive capabilities.

In the following section we briefly report the preliminary results obtained in a set of experiments in which a team of robots has been selected to solve a collective navigation problem.

2. Evolving a team of robots for the ability to solve a collective navigation problem

A group of four simulated robots live in a square arena of 270x270cm surrounded by walls that contains two circular target areas (see Figure 1, left). The robots have a circular body and are provided with: two motors controlling the two corresponding wheels, one communication actuator capable of sending signals with varying intensities, eight infrared sensors, one ground sensor (i.e. a binary sensor that detects whether the robot is located on a target area), and four communication sensors that detect signals produced by other robots up to a distance of 100cm from four corresponding directions (i.e. frontal [315°-45°], rear [135°-225°], left [225°-315°], right [45°-135°]).

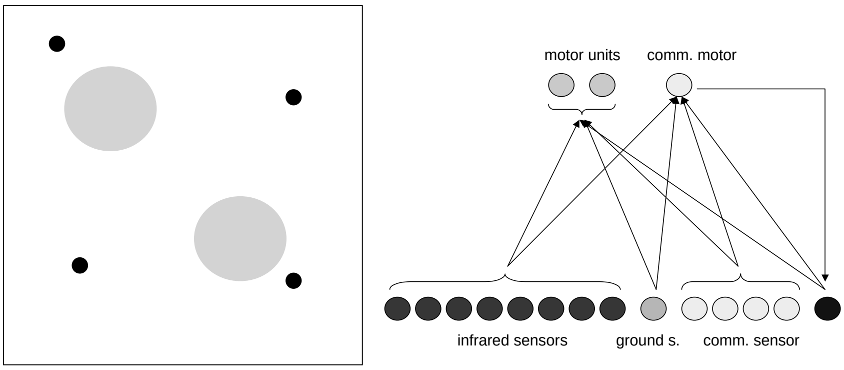

Figure 1. **Left**: The environment and the robots. The square represents the arena surrounded by walls. The two grey circles represent two target areas. The four black circles represent four robots. **Right**: The neural controller of evolving robots.

The robots' neural controllers consist of feed-forward neural networks with 14 sensory neurons (that encode the activation states of the corresponding 13 sensors and the activation state of the communication actuators at times t-1) directly connected to the three motor neurons that control the desired speed of the two wheels and the intensity of the communication signal produced by the robot.

Robots were evolved with the ability to find and remain in the two target areas by subdividing themselves equally between the two areas. Each team of four robots was allowed to "live" for 20 epochs, lasting 1000 time steps of 100 ms each. At the beginning of each epoch the position and the orientation of the four 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 -1.25 for each additional robot located in a target area in which there are two others robots already. The total fitness of a team is computed by summing the fitness gathered by the four robots in each time step.

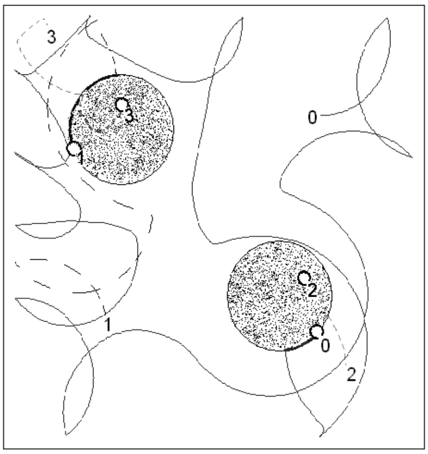

Figure 2. The behaviour of a team of evolved robots of one of the best replications of the experiment. The square and the circles indicate the arena and the target area respectively. Lines inside the arena indicate the trajectory of the four robots during a trial. The numbers indicate the starting and ending position of the corresponding robot (the ending position is marked with a white circle).

The initial population consisted of 100 randomly generated genotypes that encoded the connection weights of 100 corresponding neural controllers (connection weights are encoded with 8 bits each and normalized in the range [–5.0, +5.0]). Each genotype is translated into 4 identical neural controllers that are embodied in the four corresponding robots (i.e. teams consist of four identical robots). 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 100 generations. The experiment was replicated 10 times.

By analysing the behaviour of evolved robots of one of the best replications of the experiment one can see how they are able to coordinate and to solve the collective navigation problem. In the example shown in Figure 2, robots 2 and 3 quickly reach an empty target area. Later on, robot 1 joins robot 2 in the top-left target area. Finally, robot 0 approaches and then avoids the same target area (that already contains two robots) and later joins the bottom-left target area.

3. Evolved Ontology: grounding and behavioural effects of signals

The analysis of the signals produced by these robots shows how, in order to cooperate, robots developed a non-trivial communication systems. More specifically, the robots described in Figure 2 produce at least four different classes of signals:

(a) a signal *A* with an intensity of about 0.07 produced by robots located outside the target areas not interacting with other robots located inside the target areas;

(b) a signal *B* with an intensity of about 0.45 produced by robots located alone inside a target area;

(c) a signal *C* with an intensity of 0.25 produced by robots located inside a target area that also contains another robot;

(d) a signal *D* with an intensity of about 0.01 produced by robots outside the target areas that are approaching a target area and are interacting with another robot located inside the same target area.

Robots receiving these four types of signals modify their motor and/or signalling behaviour on the basis of the signal received and on other available sensory information in order to coordinate and solve the collective navigation problem. More specifically:

(1) robots located outside the target areas receiving signal *A* tend to modify their motor behaviour to better explore the environment (i.e. they modify their motor trajectory so to explore different areas of the environment);

(2) robots located outside target areas receiving signal *B* tend to modify their motor behaviour (by approaching the robot emitting the signal in order to enter into the corresponding target area) and their signalling behaviour (i.e. by producing signal *D*);

(3) robots located outside the target areas that receive the signal *C* (i.e. the signal produced by two robots located inside the target area) modify their motor behaviour so **as** to move away from the signal source.

4. Coupling and self-organization of primitive forms of communication protocols

The analysis of the evolved individuals also shows that robots are capable of modulating their signalling behaviour on the basis of the signals their receive from other robots. This type of analysis was conducted by testing evolved robots in new environmental conditions that might facilitate the analysis of the interaction between robots and the analysis of the relations between signalling and non-signalling behaviours.

For example, consider a test in which two evolved robots are placed in an arena that includes a single target area (Figure 3). At the beginning the two robots are both outside the target area and produce a signal with an intensity of about 0.07 (signal *A*) that push them apart so that they will have more chance to explore different areas of the environment. Later on individual 0 reaches the target area and starts to produce a signal with an intensity of about 0.45 (signal *B*). Once robot 1 gets close enough to robot 0 (i.e. within a range of 100cm) it modifies its trajectory so as to follow the signal emitted by the other robot and turn off its signalling behaviour (i.e. it starts to produce signal *D*). Later on, when robot 1 joins the area, the two robots start to produce a noisy signal with an intensity of about 0.25 (signal *C*) that will keep other coming robots out the target area.

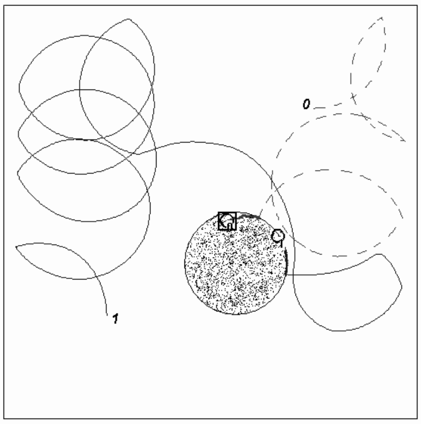

Figure 3. The behaviour of two robots tested in an arena including a single target area. The dashed and full lines represent the trajectory of robot 0 and 1, respectively. The numbers indicate both the starting and ending positions of the corresponding robots.

By analyzing the signals produced by the two robots during the behaviuor shown in Figure 3, we can clearly see two interesting examples in which the signals produced by one robot alter the signals produced by the other robots (see Figure 4).

More specifically, data shown in Figure 4 indicates the emergence of a primitive communication protocol in which signalling is switched on and off according to the needs of the robots. In this experimental setting, the signalling behaviour can be turned off by producing signals with null or close to null intensity. Indeed during the second phase (in which robot 0 is in and robot 1 is out the target area), robot 1 turns its signalling behaviour off as soon as it detect the signal *B* produced by robot 0. The reason that the signalling behaviour is switched off in these circumstances is probably due to the need to not alter the other robot's signalling behaviour. Indeed, by preventing robots from the possibility to switch their signalling behaviour off in these circumstances, a significant decrease in performance can be observed.

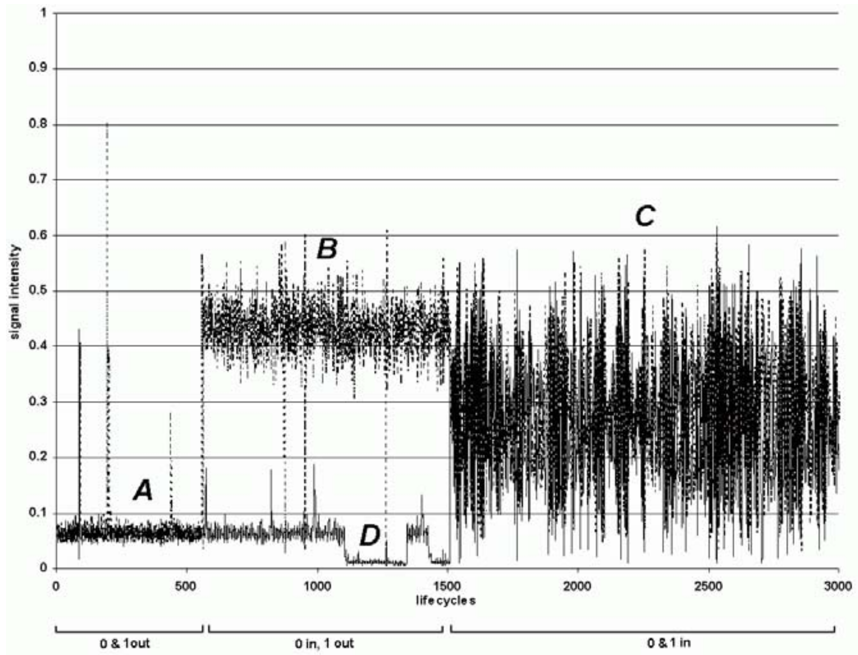

Figure 4. Intensity of the signals produced by the two robots during the behaviour shown in Figure 3. Dashed and full lines indicate the intensity of the signals produced by robot 0 and 1, respectively. Letters (*A*,*B*,*C*, and *D*) indicate the 4 classes of signals produced by the robots. The bottom lines indicate the three phases in which: (1) both robots are out the target area, (2) robot 0 is in and robot 1 is out, and (3) both robots are inside the target area.

More generally, the emergence of the ability to switch signalling behaviours on and off can be explained by considering that communicative actions might have counter-adaptive effects on other robots and on the adaptive ability of the population as a whole. For instance, communication acts might negatively affect the behaviour of other robots thus preventing or delaying the ability of these robots to accomplish their tasks. Signalling might even interfere with the behaviour of the robot producing the signal when, as in the case of the experiments described in this paper, robots perceive their own signals.

A second interesting phase is that in which the two robots located in the target area produce signal *C* that keeps other approaching robots away from the target area. During this phase the bi-directional interaction between the two robots leads to a coupled dynamical process that seems crucial both from the point of view of allowing the two robots to detect the presence of the other robot located inside the area and from the point of view of allowing the two robots to produce signals that push other approaching robots away.

When a single robot is located inside the area, it starts to rotate and to produce signal *B* that attracts other approaching robots. When a second robot joins the same area, it also starts to rotate and to signal. However, as a result of the interaction between the two motor and signalling behaviours, the two robots start to produce signal *C*, a highly varying signal with an average intensity of about 0.25, that rejects other approaching robots.

The fact that the interaction between signalling and not signalling behaviours are crucial for determining the ability of the robots to solve their problem can be demonstrated by running tests in which these forms of behaviours are selectively perturbed. For instance, by preventing one of the two robots to freely rotate at the desired speed when both robots are located inside the target area, we observed that the two robots are no longer able to remain inside the area. Similarly, by preventing one of the two robots from being able to emit signals (while preserving the signalling behaviour of the robot intact), we observed that other approaching robots would erroneously join the area. More generally, we observed that perturbations of the motor behaviour tend to significantly affect signalling behaviour and vice versa and produce a significant degradation of performance.

Overall, the tests performed indicate that signals *B* and *C* are the emergent results of the dynamical interaction between the motor and signalling behaviours of the two robots.

5. Conclusion

In this paper we showed how a population of simulated robots evolved for the ability to solve a collective navigation problem can develop non-trivial communication abilities despite the fact that communication is not explicitly rewarded. The self-organized emergent vocabulary includes four different signals that influence both the motor and signalling behaviours of other robots.

The analysis of the evolved behaviours also indicates: (a) the emergence of a simple form of communication protocol that allows individuals to switch communication on and off according to the current needs of the robots, (b) the emergence of tightly co-adapted communicative and non-communicative behaviours, and (c) the exploitation of properties resulting from the dynamical interactions between communicative and non communicative behaviours produced by interacting robots.

In future research we plan to: (1) study more complex tasks, (2) provide the neural controllers of evolving robots with internal neurons and recurrent connections that might allow them to integrate sensory-motor information over time, and (3) allow evolving individuals to switch the impact of perceived signals on and off.

Acknowledgments

The research has been supported by the ECAGENTS project founded by the Future and Emerging Technologies programme (IST-FET) of the European Community under EU R&D contract IST-2003-1940.

References

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

Cangelosi A. (2001). Evolution of communication and language using signals, symbols and words. *IEEE Transactions in Evolutionary Computation*, 5(2): 93-101.

Cangelosi A, Parisi D. (Eds.) (2002). *Simulating the Evolution of Language*. London: Springer-Verlag.

Kirby S (2001). Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation and Cognitive Science*, 5(2): 102-110.

Quinn M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosik, P. (Eds.) *Advances in Artificial Life: Sixth European Conference on Artificial Life* (ECAL 2001). Springer Verlag.

Quinn M., Smith L., Mayley G., Husbands P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London*, Series A: Mathematical, Physical and Engineering Sciences 361, pp. 2321-2344.

Steels L. (1999). *The Talking Heads Experiment*, Antwerpen, Laboratorium. Limited Pre-edition.

Steels L. and Kaplan F. (2001). AIBO's first words: The social learning of language and meaning. *Evolution of Communication*, 4:3-32.