

Orienting Learning By Exploiting Sociality: An Evolutionary Robotics Simulation

A. Acerbi, D. Marocco

Abstract – One of the advantages of sociality resides in the opportunity of exploiting the behavior of other individuals of the same group as a reliable source of information. In this paper we present an evolutionary simulation in which a population of 10 mobile robots has to develop a simple behavior consisting in the discrimination of two different foraging areas in the environment. We show that, given a minimal environmental pressure, a combination of learning oriented by social cues and selection at population level can lead to effective results. We further analyse the dynamic of the evolution of sociality, focusing on the fact that the global adaptiveness is a product of the combination of singularly non-adaptive processes and on the presence of a reinforcing positive feedback within populations, that is, the more 'social' a population is, the more advantageous it is to exploit social cues in that population.

I. INTRODUCTION

THE evolution of sociality among previously non-social entities has been identified as a major transition in the history of life [1]. An evolutionary explanation of the emergence of sociality requires to recognize the selective advantages that a single individual can obtain by staying in a group or, to say it better, the selective advantages that being a member of a social group give to an individual, compared to being a solitary part of a collection of non-related entities.

Here we concentrate on one of these possible advantages, namely the fact that individuals who live in a group can exploit other individuals of the same group (co-groupers) as a reliable source of information about the environment. From the point of view of a social animal, the other animals that live in the same group share one obvious characteristic: they are alive. Thus, they will behave, on average, in an adaptive manner and their behavior will be specifically tuned for the environment in which they are living [2].

From an ethological perspective, the modification of behavior during lifetime, by exploiting the interactions with conspecifics, is a useful definition of what is called social learning (for reviews, and definitions, of social learning in ethology see [3][4]).

Human beings heavily rely on social learning for the development of their behavioral repertoire [5] and they are

likely to be the only species that makes an extensive use of cognitively advanced forms of social learning, that require both the explicit copy of results and actions performed by a model [6][7] and an active role of the model in the transmission of information, such as teaching ([8]; forms of teaching have been observed also in other species, even in ants, see [9], but the issue of teaching among other species other than humans remains still controversial. For a brief review see [10]).

Nevertheless, especially in the last twenty years, researches in animal behavior have shown how social learning can be significant for the development of complex behavioral skills in primates [11] and in other vertebrates like, just to cite some influential studies, rats [12][13], birds [14], and fish [15]. In addition, even if social learning studies have been almost entirely focused on vertebrates, several pieces of evidence indicate that forms of social learning are also present in insects (for a review on social learning in insects see [16]).

Those ethological findings, taken together, suggest two considerations. Even if it is likely that the role played by social learning among human beings is greater than the role it plays among other species, although the mechanisms that permits the relative stability of cultural patterns and supports cumulative cultural evolution in human groups still remain not clear (see different answers to the last questions in [8], [5], and [17]), social learning does not seem restricted to human beings. Hence, the selective advantage of living in social group, because of social learning, can be more pervasive than generally thought.

Besides, the social transmission of behaviors in species different than humans is often realized without the need of complex cognitive machineries, but by simple processes that exploit the dynamics between learning at the individual level, the characteristics of the environment and of the population, and the genetic evolution at population level, without an easy-to-trace distinction between processes.

In relation to the last point, individual based modeling, or the development of artificial organisms (robots or simulated agents) through artificial life techniques, can be an excellent way to explore this kind of dynamics ([18][19][20]; for a general review see [21][22]). In particular, artificial life techniques are especially suitable to take into account different levels of analysis (agent – population – environment) and adaptive processes that happen at different timescales (lifetime learning – genetic evolution) which are often difficult to consider resting only on empirical observations or laboratory experiments.

A. Acerbi, is with the Department of Psychology, University of Bologna, Viale Berti Pichat, 5 40127 Bologna, Italy (e-mail: a.acerbi@unibo.it).

D. Marocco is with the School of Computing, Communications & Electronics, University of Plymouth, Drake Circus, Plymouth PL4 8AA, United Kingdom (e-mail: davide.marocco@plymouth.ac.uk).

In addition to that, using robots or simulation platforms that carefully reproduce real robots' behaviors, as in the present case, forces the experimenter to implement learning strategies in a real physical system. Hence, differently from abstract agent based modeling or analytical modeling, in which a behavior is usually represented by a single variable that can be acquired or not by an agent, robotic models are constrained to realize physically plausible learning mechanisms and to take into account the real physical dynamic of an agent that interact with other agents and with the environment [23]. Even if mathematical models have proven to be extremely important and insightful tools for understanding social learning and cultural evolution dynamics (e.g., see [5]), the additional complexity provided by robotic models can be a further useful commitment in designing new experiments.

In our contribution we present a simple experimental scenario in which a team of 10 simulated robotic agents has to develop an adaptive behavior consisting in the discrimination of two different foraging areas present in the environment. The team is subjected to a selective pressure and an on-line genetic algorithm operates on the agents' learning behavior. In particular, as we will explain in the next section, two characteristics can be inherited: the strength of the learning and its modality (individual or social). The individual modality is a simple blind process that reinforces the 'preference' of an individual for one of the two areas, independently from the features of the specific area. In the same fashion, in the social modality the strength of the learning process is modulated by the number of individuals located nearby.

In the third section we show that, given a minimal environmental pressure, the social modality prevails over the individual one among the majority of the populations tested during the experiment: the combination of a blind, but socially oriented individual learning, together with a selection process at the population level can lead to an adaptive social transmission of behaviors. That is, it can lead to a simple form of social learning. In addition to previous works [24][25], we analyse in detail the dynamic of the evolution of sociality by showing the presence of a reinforcing positive feedback between the number of individuals that exploit the sociality and the adaptiveness of the exploitation itself. In other words, the more a population is 'social', the more advantageous it is to be social in that population.

II. EXPERIMENTAL SETUP

To investigate the adaptive role of sociality in the orientation of learning and how the further development of sociality within a population is favoured by this process, we ran a robotic simulation in which a population of robots has to survive in an unknown environment. The simulation consists of several replications of the same basic task, in which robots have been tested in different environmental conditions. In particular, each condition provides a different selective pressure under which robots have to adapt in order to survive.

A. The robotic model

The environment consists of a square arena of 200 x 200 cm surrounded by walls. The floor of the arena is grey and contains two circular areas with a diameter of 60 cm coloured in black and white, respectively (Fig. 1).

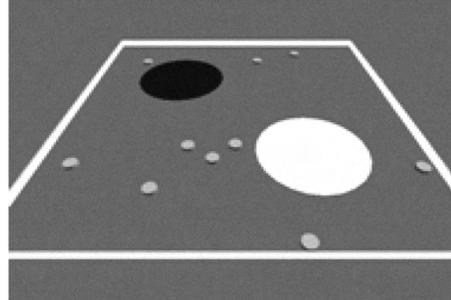


Fig. 1. The environment of the experiment. Black and white circles represent the two foraging areas. Small grey circles represent the 10 robots.

A population consists of ten simulated e-puck robots (Fig. 2(a)) which are provided with simple sensory-motor capabilities that allow them to move, gather information from the environment and to produce sound signals with a fixed intensity. The robots have a circular body with a radius of 37 mm and their sensor apparatus is provided of eight infra red sensors placed around the body that provide information about the presence of an obstacle (i.e. walls or other robots within a range of about 4 cm), one ground sensor placed at the bottom of the robot, one microphone, and one speaker.

Robots' neural controllers consist of artificial neural networks with ten sensory neurons and three motor neurons (Fig. 2(b)). Eight sensory neurons encode the activation states of the eight infra red sensors and two neurons measure the activation of the ground sensor that encodes the colour of the floor below the robot (gray [0 0], white [1 0], and black [0 1]). The first two motor neurons encode the desired speed of the two corresponding wheels normalized between $-MaxSpeed$ and $+MaxSpeed$. The third motor neuron instead encodes the maximum speed of both wheels (i.e. the $MaxSpeed$ parameter, normalized in the range [0;10]). The third motor neuron thus acts as a modulator that regulates the speed with which the robot moves or turns at each time step. The eight sensory neurons that encode the state of the infra red sensors are connected to all the three motor neurons. The two sensory neurons that encode the colour of the ground, instead, are connected to the third motor neuron only (i.e. to the motor neuron that regulates the maximum speed of the two wheels).

When robots are located in a foraging area, independently from its colour, they continually emit a sound signal of fixed intensity ((0.1 I) - See section 'Genetic Inheritance') through their speakers. Moreover, the robots can detect through their microphones the signals produced by themselves or by other robots located nearby.

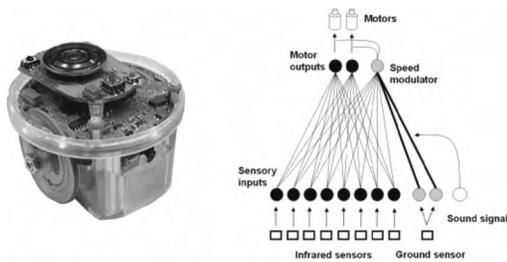


Fig. 2. (left) The e-puck robot and (right) the neural controller. In (b) thick lines represent the weights that are modified in the simulation. As it is shown, the speed modulator does not operate on the outputs but it acts directly on the motors' activations. The intensity of the sound signal is used to modify the learning rate in case of social individuals. See the text for details.

As in the case of [24] and [25], in order to provide robots with basic sensory-motor capabilities that allow them to move around the environment and avoid obstacles, the robots are pre-trained through an evolutionary method [26] for developing an obstacles avoidance behavior and for exploring the environment in an arena that does not contain coloured areas. This pre-training process only concerns the strength of the connections between the eight infra red sensors and the three motor neurons (see [24] for further details).

Because of that procedure, the robots' behavior outside the foraging areas is already determined and cannot be modified during the simulation, therefore robots can only modify their behavior inside the areas.

At the beginning of the simulation the population is placed in the environment, each robot at a different random position, and the task of the robots is to survive as long as possible in an environment that contains the white and the black target areas.

Each robot starts its life with a fixed amount of energy (i.e. 2000 units) and the environment provides two different sources of selective pressure. In all experimental conditions, every time step spent outside the areas causes the loss of 1 unit of energy, while every time step spent in the black area causes a variable loss of energy units, depending on the experimental condition. The white target area does not affect the energy level of the robots. In the simulation we test the robots in eleven different conditions in which the loss of energy in the black area varies from 0 to 5 energy units per time step with an increment of 0.5 units at every condition. In this way, the manipulation of the amount of energy lost in the black area determines the selective pressure under which robots have to survive.

Therefore, with the exception of the first three experimental conditions (i.e. energy loss from 0 to 1.0), in which even staying in the black area is more advantageous or equal than exploring the environment, the simulation task consists in locating and spending time in the white target area, while avoiding the black one. In practice, given the pre-evolved exploratory behavior, the free parameters that encode the strength of the connections between the two ground sensor neurons and the third motor neuron should be set so to allow the robot to slow down and, eventually, to

stop in the white target area and to quickly move out from the black target area.

It is worthy to note, however, that robots do not have any direct feedback about the quality of a target area or about their own energy. In fact, they only have information about the colour of an area when they are in it, but they do not know how remaining in a particular area affects their energy state. This means that the experience of a particular area does not provide any explicit information about the correct behavior that should be performed in that area.

In all experimental conditions, the population is evaluated for one trial lasting 20000 time steps (of 100ms each) and every robot is allowed to live for about 2000 time steps. When a robot finishes up its energy it dies and when a robot lives for 2000 time steps it has a small probability of dying for every further time step of its life (1‰). When a robot dies, it is replaced with a new robot placed in the environment in a randomly chosen position and orientation, with full initial energy.

The dynamic of the adaptation to the environment consists of two interplaying processes: one that acts at the phylogenetic level (genetic inheritance) and another that acts at the ontogenetic level (learning).

B. Genetic inheritance

Each robot of the population has a genetic code consisting of two parameters (i.e. two genes) that determine two different aspects of its learning behavior. The first parameter ω , which takes values 0 or 1, determines if a robot learns only on the basis of its own experience (individual modality) or if it also relies on social cues provided by the other robots (social modality). The second parameter ϕ (which varies continuously in the range [0;1]) determines the strength of learning. More details about the role of ω and ϕ are provided in the next section.

At the beginning of the simulation, 50% of the population presents an individual modality and 50% a social modality. That is, the parameter ω of 50% of the population (5 robots) is set to 0, while in the other 5 robots is set to 1. In this way, at the beginning of the simulation, any bias of the population towards one of the two modalities is prevented. The value of the parameter ϕ is randomly set for each robot (notice the difference with respect to Marocco and Acerbi 2007, where the parameters were both randomly set at the beginning of the simulation). When a robot dies during the simulation, an on-line genetic algorithm acts on the population, by replacing the dead robot with a newborn that inherits the two genes from a randomly individual selected among the robots that live from at least 2000 time steps, i.e. the oldest robots of the population. Notice that the duration of a robot's life is an indirect indication of the adaptiveness of its behavior. Indeed, the fact that a robot lives more than 2000 time steps is a good clue of the fact that it has acquired the capacity to correctly discriminate between the two areas during its ontogenetic development or, at least, that it did not develop a 'preference' for the black area (given that there is a selective pressure for staying in the white one).

During the genetic transmission process the gene ω can switch its original value with a probability of 10% and the

gene ϕ is mutated by adding a random value in the range $[-0.1;0.1]$.

C. Learning

This process affects the way in which robots modify their behavior during the lifetime. In particular, the two synaptic weights that connect the two ground sensor neurons and the third motor neuron are trained by means of a non-directional learning algorithm. In practice, at the beginning of the life the weights are initialized with a null value, that is, the robots tend to go away from both the areas. When a robot happens to stay inside a target area, a learning algorithm acts to reinforce with a positive feedback the individual preference for that area, independently from the colour of the area and, consequently, from the effect on the robot's energy. This positive feedback mechanism is realized by modifying the two weights every time an individual is inside one area through the addition of a randomly chosen value in the interval $[-lrate;+lrate]$ to their current value, and by retaining only the variations that lead to an increase of the time spent by the individual in the corresponding target area.

The learning rate is modulated depending on the value of ω :

$$lrate = \begin{cases} \phi & \text{if } \omega = 0 \\ \phi \cdot Q_I & \text{if } \omega = 1 \end{cases}$$

Where Q is the intensity of the signal perceived by the robot inside the area expressed in I units and it varies linearly in the range $[0.1;1.0]$ according to the number of individuals located in the same foraging area: $0.1I$ corresponds to one robot in the area (the learner robot itself), and $1.0I$ to all the 10 robots in the same area.

In practice, if ω is equal to 1 (social modality) the learning rate is a function both of ϕ and of the intensity of the signal detected.

III. RESULTS

The simulation presented in this work has been run considering eleven different values of energy loss in the black foraging area, i.e. from 0 to 5. In particular, for each replication, 100 different populations of ten robots have been randomly generated and tested in the same environmental conditions.

A. The adaptiveness of social exploitation

The results of the simulation show an interesting relationship between the exploitation of social cues by the robots and the amount of selective pressure provided by the environment.

To measure the consistency of this relationship with respect to the different environmental conditions, for every replication we calculated, at the end of the simulation, the average value of the parameter ω among a given population. On the basis of this measure, the graph in figure 3 shows that the average value of ω within populations moves towards values above 0.5 according to the quantity of selective pressure provided by the environment. This means

that, in relation to the increase of the energy loss in the black target area, the populations tend to be composed by a majority of individuals with ω equal to 1, thus, individuals that employ the social modality to orient the learning.

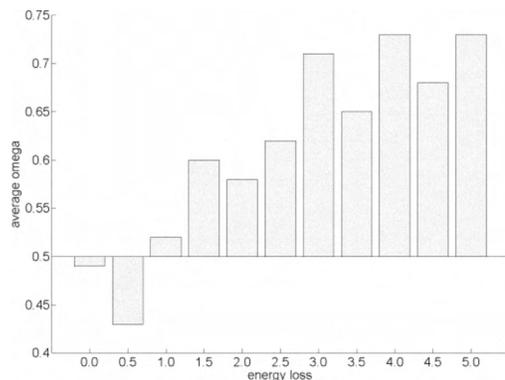


Fig. 3. Average value of parameter ω , at the end of simulations for every experimental condition. Value 0.5 indicates the starting condition of the populations, in which 50% of the population rely on individual modality and 50% rely on social modality.

We can also notice that also a minimal adaptive pressure leads the population towards the exploitation of sociality. In fact, only in the case of no adaptive pressure in differentiating between the two areas (energy loss lower than 1.0) we observe values lower than 0.5 for the average of ω . In those cases, in fact, there is no difference between the two areas and the punishment provided outside the areas leads the robots to quickly locate and remain in one of the two target areas without distinctions. This fact, in turn, makes the individual modality, i.e. the fact that robots do not rely on social cues, more advantageous over the social modality.

However, as soon as the adaptive pressure provided by the increased energy loss in the black area starts to push the population towards the need of locating the white area and avoid the black one, the average of the population moves toward values of ω greater than 0.5.

According to the average value of ω , that reflects the number of robots which rely on social cues in each population, we consider a population as 'social' if its average ω is greater than 0.5, meaning that the majority of the robots have a ω value equal to 1. In particular, in the following sections we will call *social population* a population in which more than five robots use social cues to orient their individual learning and, similarly, *social individual* a robot that expresses that particular behavior.

The adaptiveness of social exploitation is also confirmed by the differences of the mortality (i.e. the number of robots that 'die' during each trial) between social populations and the non-social ones. Indeed, figure 4 shows the difference between the total number of died individuals between social and non-social populations for different energy losses. By looking at the graph we can see that, with the exception of energy loss 0, the values are positive, that is the mortality rate of the former is always lower than the mortality of the

latter, and the difference continues to increase roughly following the progressive increment of the selective pressure.

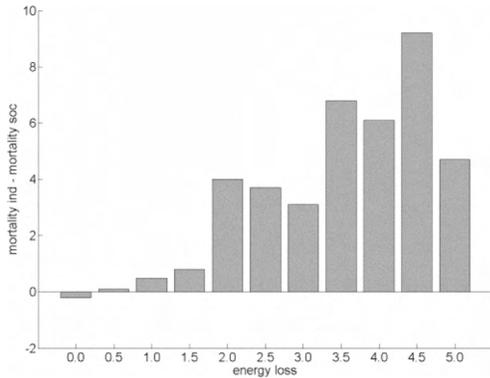


Fig. 4. Differences of the average mortality rate (i.e. the number of robots died during a trial) between non-social population and social population for every experimental condition.

Moreover, by looking more in detail at the final compositions of populations for all the replications in different experimental conditions (figure 5), we can notice, as we said above, that the number of populations with a large majority of social individuals increases, following the increment of the selective pressure.

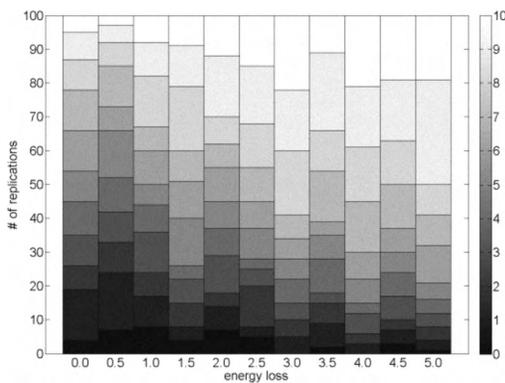


Fig. 5. Final composition of the populations for all the replications in the different experimental conditions. Depending on the number of social individuals populations go to black (no social individuals) to white (10 social individuals). See color bar on the right.

However, in addition to that, is interesting to note that the increasing of social populations is mainly due to the increasing of groups in which all or most individuals are social (groups made by 8, 9, or 10 social individuals account for approximately 50-60% of the replications after energy loss equal to 3.0).

B. The dynamic of social exploitation

In order to analyse with more details the emergence of sociality we focus on the experimental condition in which energy loss is equal to 5.0, that is, the condition that presents the highest selective pressure.

Figure 6 shows the average value of ϕ , i.e. the parameter that determines the strength of learning (see above), through time for social and non-social populations.

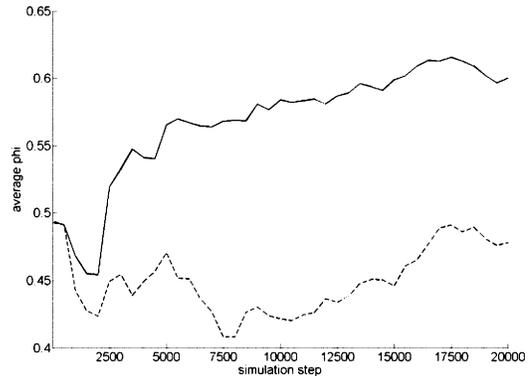


Fig. 6. Average value of the parameter ϕ through time for condition of energy loss 5.0. Dashed line: non-social populations. Continuous line: social populations.

In the former case, the initial value of ϕ , which is randomly initialized at the beginning of the simulation, is around 0.5, but it suddenly drops and stays lower for about 2500 steps. After 2500 steps it starts to grow up again until it reaches a value around 0.6 at the end of the simulation. A possible interpretation of this dynamic can illustrate how the two inherited parameters (ϕ and ω) interact. In fact, during the first steps of the simulation, the ratio between social and non-social individuals in the population is roughly 50%: non-social individuals can be present with equal probability in both areas and social individuals can hardly, for that reason, make use of social information to correctly orient their learning. At this time of the simulation learning is in general a risky job, therefore the selection process favours those individuals that learn slowly, as they tend to continue in exploring the environment without stopping in the target areas. This fact reduces the average value of ϕ in the population.

Nevertheless, selection does favour also those individuals that, for random reasons, stop in the white area, so that, after a while, more robots are located in the white area than in the black one. At this point, learning starts to be reliable for a social individual, so that ϕ increases among social populations and remains above the starting value for the rest of the simulation.

This analysis also demonstrates that the trend of ϕ is strictly related to the trend of ω (see figure7). In fact, higher values of ϕ are not selected until most of the population is composed by social individuals, because only in that condition the information provided by the behavior of co-groupers is reliable. This interpretation is further confirmed by the fact that the average value of ϕ for non-social populations, in the same experimental condition, (energy loss 5.0), never goes over 0.5 (see figure 6).

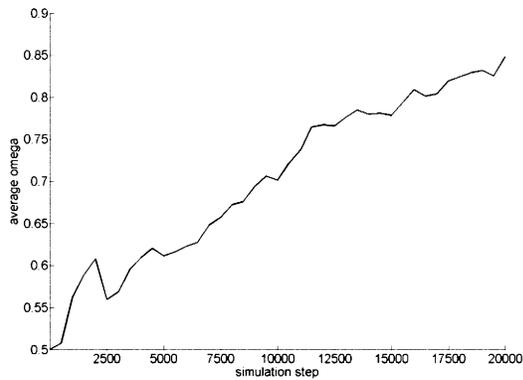


Fig. 7. Average value of parameter ω (for social populations) through time for condition of energy loss 5.0.

Moreover, ϕ is also related to the strength of the selection at population level, not only for the trivial fact that it is subjected to selection in certain conditions, but even for the fact that these conditions are realized just because the selection has already acted in order to remove from the population those individuals with non adaptive behavior.

Not only a high value of ϕ is not per se adaptive, but the same is true for the social modality of learning, i.e. ω equal to 1. Figure 8 shows the difference, through time, between the average life span of social individuals that belong to a social population (SocSoc) and the average life span of social individual that belong to a non-social population (SocNSoc) in the same experimental condition (energy loss 5.0).

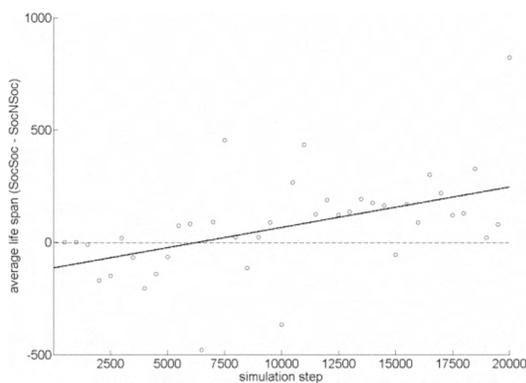


Fig. 8. Dots represent the difference between the average life span of social individuals that belong to a social population (SocSoc) and the average life span of social individuals that belong to a non-social population (SocNSoc). The continuous line represents a linear fitting of the dataset.

A first thing to notice is that most of the points of the plot lay over the dashed line, that is, in general that difference is positive. This result indicates that SocSoc tend to live longer than SocNSoc. Therefore, being a social individual is more adaptive in a social population than in a non-social one.

Even more interesting is the dynamic through time of this dataset. The amount of the difference, in fact, increases significantly through time (Spearman's $\rho = 0.57$, $p <$

0.001) meaning that, at the beginning of the simulation, being a SocSoc is not better than being a SocNSoc, if not even worse (for a discussion related to this point see Acerbi et al. 2007). Thus, we can conclude that being a social individual is not adaptive per se, but it is adaptive only with respect to the overall condition of the population.

V. DISCUSSION AND CONCLUSIONS

The findings presented in the previous section suggest that, in the cases in which we observe the evolution of sociality, we are in presence of a reinforcing positive feedback towards sociality, that guides the adaptive search at both phylogenetic and ontogenetic level.

The advantage of sociality derives primarily by the complex interaction among three forces that, at different levels, channel the adaptation process of the populations: the individual learning algorithm, the social interaction, and the genetic transmission. In fact, by considering that the learning algorithm provided to the robots is totally 'blind' with respect to the adaptiveness of the behavior, at the beginning the robots tend to distribute roughly equally in the two areas (also see Acerbi et al. 2007 for a detailed description of this process). However, the selective pressure at the population level favours those individuals that, for random reasons, have developed a 'preference' for the white area because they tend to survive more than others. At this point, at the ontogenetic level, those individuals that can exploit social information are favoured, as well as individuals that learn faster. This fact, thanks to the transmission mechanism that favours individuals with adaptive behavior in propagating their genes, produces an increment in the number of social individuals within a given population. In turn, this makes more advantageous with respect to the early stages of the evolution, for a newborn, to be social, because in a population mainly composed by social individuals the behavior of co-groupers tends to be a more reliable source of information than in a population mainly composed by non-social one.

We claim that the described adaptive mechanism, able to generate a form of social transmission - or diffusion - of behaviors between robots, can be considered as a genuine form of social learning. In fact, as we pointed out at the beginning of our contribution, social learning should not be considered only as a form of explicit copy of behaviors that requires high-level cognitive abilities, but also as an emergent property of particular systems that possess certain specific characteristics. In this view, therefore, social learning is not a mechanism that can be directly implemented in populations of artificial agents, or discovered as a cognitive machinery in the heads of observed animals. Rather, it can be viewed as a complex emergent outcome of a multilevel interaction between different and non necessarily - if taken singularly - adaptive processes, such as, in the case we analysed, individual learning, social influence and selective pressure at population level.

In future work, we plan to investigate the scalability of the mechanism in different environmental conditions and selectively manipulate specific properties of the system.

Many aspects of the simulation, in fact, can be modified in order to add complexity and degrees of freedom to the parameters we take into account. Major transitions in the obtained results are expected, e.g., by linking the adaptive pressure provided by the environment, i.e. the energy loss, to the other characteristics of the behavior of robots. If, for example, the amount of adaptive pressure directly depends on the number of robots in the same area, how will this characteristic affect the ratio between social and non-social individuals in the populations? In that situation, is it always the best solution to learn faster?

Other possible improvements are related to the different modalities that regulate the social interaction between robots. In this work, for example, robots located in a target area always emit a continuous signal. What does it happen if the robots can modulate on their own the emission of the signal? Under what condition will we observe, if any, the emergence of selfish robots, i.e. robots that exploit social information without providing social information to others, within a population?

Moreover, the idea that robots might self-regulate their communication activities, and in turn social interactions, open new perspectives in the development of language-like communication acts that might refer both to environmental and social features. Such additional level of complexity will require additional skills for the agents, that would be capable of flexible and scalable categorisation abilities, in order to interact fruitfully with the physical and social environment.

We think that an approach based upon the theory of Neural Modeling Fields [27] could endow the robots of the necessary flexibility and scalability required to deal with the difficult task of merging many different sensory information and use the outcome of such a fusion to improve robots abilities to categorise the world and to communicate about it.

Cangelosi and colleagues [28] have already shown the applicability of the Neural Modeling Fields in the language and cognition integration domain. On the same perspective, the presented set up allows to test and extend the range of applicability of the theory to a different level of complexity, less focused on the scaling up of individual capabilities, such as actions and language, rather more focused on the social dimension.

ACKNOWLEDGMENTS

This work was partially supported by the ITALK project, a project funded by the Cognitive Systems and Robotics Integrating Project, and by the CULTAPTATION project funded by the New and Emerging Science and Technology of the European Commission. 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] J. Maynard Smith and E. Szathmáry, *The Major Transitions in Evolution*. Oxford: Oxford University Press, 1995.

[2] B. G. Galef and K. N. Laland, "Social learning in animals: Empirical studies and theoretical models", in *BioScience*, 55(6), 2005, pp. 489-499.

[3] T. R. Zentall and B.G. Galef, *Social learning: psychological and biological perspectives*. Hillsdale: Lawrence Erlbaum, 1988.

[4] C.M. Heyes and B.G. Galef, *Social learning in animals: the roots of culture*. San Diego: Academic Press, 1996.

[5] P. J. Richerson and R. Boyd, *Not by genes alone: How culture transformed human evolution*. Chicago: Chicago University Press, 2004.

[6] J. Call and M. Carpenter, "Three sources of information in social learning", in *Imitation in Animals and Artifacts*, Dautenhahn, K., Nehaniv, C. L., Eds., Cambridge: MIT Press, 2002, pp. 211-227.

[7] C. Tennie, Call C. and M. Tomasello, "Push or pull: imitation versus emulation in great apes and children", *Ethology* 112, pp. 1159-1169, 2006.

[8] M. Tomasello, S. Kruger and H. Ratner, "Cultural learning", *Behavioral and Brain Sciences*, 16, 495-552, 1993.

[9] N.L. Franks and T. Richardson, "Teaching in tandem-running ants", *Nature* 439, 153, 2006.

[10] G. Csibra, "Teachers in the wild", *Trends in Cognitive Science* 11(3), 95-96, 2007.

[11] M. Tomasello, "Do apes ape?", in *Social learning in animals: The roots of culture*, C. M. Hayes and B. G. Galef, Eds., San Diego: Academic Press, 1996, pp. 319-346.

[12] K.N. Laland and H. Plotkin, "Social learning and social transmission of foraging information in Norway rats (*rattus norvegicus*)", *Animal Learning and Behavior* 18, pp. 246-251, 1990.

[13] B. G. Galef, "Introduction: social learning and imitation", in *Social learning and imitation: The roots of culture*, C. M Heyes and B. G. Galef, Eds., New York: Academic Press, 1996, pp. 3-16.

[14] D.F. Sherry and B.G. Galef, "Cultural transmission without imitation: Milk bottle opening by birds", *Animal Behavior* 32, pp. 937-938, 1984.

[15] L.A. Dugtakin, "Copying and mate choice", in *Social learning in animals: the roots of culture*, C. M Heyes and B. G. Galef, Eds., New York: Academic Press, 1996, pp. 85-105.

[16] E. Leadbeater and L. Chittka. "Social Learning in Insects – From Miniature Brains to consensus Building", *Current Biology* 17, pp. 703-713, 2007.

[17] D. Sperber and L. Hirschfeld, "The cognitive foundations of cultural stability and diversity", *Trends in Cognitive Sciences* 8(1), pp. 40-46, 2004.

[18] J. Noble, E. Tuci and P.M. Todd, "Social learning and information sharing: an evolutionary simulation model of foraging in Norway rats", in *Proceedings of the Fifth European Conference on Artificial Life*, Floreano, D., Mondada, F., Nicoud, J. D., Eds, Berlin: Springer, 1999, pp. 514-523.

[19] J. Noble and P. M Todd, "Imitation or something simpler? Modelling simple mechanisms for social information processing", in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds., Cambridge: MIT Press, 2002, pp. 423-440.

[20] D.J. van der Post and P. Hogeweg, "Diet traditions and cumulative cultural processes as side-effects of grouping", *Animal Behavior*, 75, pp. 133-144, 2008.

[21] K. Dautenhahn and C.L. Nehaniv, *Imitation in animals and Artifacts*. Cambridge: MIT Press, 2002.

[22] C.L. Nehaniv and K. Dautenhahn, *Imitation and Social Learning in Robots, Human and Animals: Behavioural, Social and Communicative Dimensions*, Cambridge: Cambridge University Press, 2007.

[23] T. Ziemke, "On the Role of Robot Simulations in Embodied Cognitive Science", *AISB Journal* 1(4), pp. 389-399, 2003.

[24] A. Acerbi, D. Marocco and S. Nolfi, "Social facilitation on the development of foraging behaviors in a population of autonomous robots", in *Advances in Artificial Life. Proceedings of ECAL 2007*, Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., Coutinho, A., Eds, Berlin: Springer, 2007, pp. 625-634.

[25] D. Marocco, A. Acerbi, "Adaptation and Social Facilitation in a Population of Autonomous Robots", in *Proceedings of the Seventh International Conference on Epigenetic Robotics*, Berthouze, L., Prince, C. G., Littman, M., Kozima, H., Balkenius, C., Eds, Lund: LUCS, 2007.

- [26] S. Nolfi and D. Floreano, "Synthesis of autonomous robots through evolution", *Trends in Cognitive Sciences*, 6(1), pp. 31-37.
- [27] L.I. Perlovsky, *Neural Networks and Intellect: using model-based concepts*, Oxford University Press: New York, NY, 2001.
- [28] A. Cangelosi, V. Tikhanoff, J.F. Fontanari and E. Hourdakis, "Integrating language and cognition: A cognitive robotics approach", *IEEE Computational Intelligence Magazine*, 2(3), 2007, pp. 65-70.