

Evolutionary Conditions for the Emergence of Communication

Sara Mitri, Dario Floreano and Laurent Keller

Laboratory of Intelligent Systems, EPFL
Department of Ecology and Evolution, University of Lausanne

1 Introduction

Communication plays a central role in the biology of most species ([Maynard-Smith and Szathmàry \(1997\)](#)), particularly in social species where it allows for the transfer of vital information among group members, thereby ensuring ecological success ([Wilson \(1975\)](#)). Understanding communication and its evolution are therefore essential to our understanding of the mechanisms driving social behavior. Despite extensive efforts toward this end, the conditions conducive to the evolution of communication and the paths by which reliable systems of communication become established remain largely unknown ([Maynard-Smith and Harper \(2003\)](#)). This is a particularly challenging problem because efficient communication requires tight co-evolution between the signal emitted and the response elicited. A powerful means to study the evolution of communication would be to conduct experimental evolution (see e.g., [Griffin et al. \(2004\)](#); [Fiegna et al. \(2006\)](#)) in a species with elaborate social organization. Unfortunately, highly social species are not amenable to such experiments because they typically have long generation times and are difficult to breed in the laboratory.

In this chapter¹ we report on an experimental system using groups of foraging robots, which was designed to circumvent this problem. Robots could forage in an environment containing a food and a poison source that both emitted red light and could only be discriminated at close range (see [Fig. 1](#), right). Under such circumstances, transmitting information on food and poison location can potentially increase foraging efficiency. However, such communication also incurs direct costs to the signaler, since signaling results in higher robot density around the food. Due to spatial constraints around the food (a maximum number of robots can feed simultaneously), high robot density increases competition and interference, resulting in robots sometimes pushing each other away from the food. Thus, while beneficial to other group members, signaling of a food location effectively constitutes an altruistic act ([Hamilton \(1964, 1996\)](#)) because it decreases the food intake of signaling robots. This setting therefore mimics the natural situation where communicating almost invariably incurs costs in terms of signal production or increased competition for resources ([Zahavi and Zahavi \(1997\)](#)).

¹ This chapter is largely based on [Floreano et al. \(2007\)](#).




Fig. 1. Left: The s-bot robot used for the experiments is equipped with a panoramic vision camera and a ring of color LEDs used to emit blue light. Right: Robots emitting blue light around the food object emitting red light.

Studying why group members convey information when they also compete for limited resources requires consideration of the kin structure of groups (Hamilton (1964); Maynard-Smith (1991); Johnstone and Grafen (1992)), and the scale at which altruism and competition occur (level of selection) (West et al. (2002); Keller (1999)). We therefore conducted experimental evolution on colonies of robots with two kin structures (low and high relatedness) and two levels of selection (individual and colony level regimes). There were thus four treatments: high relatedness with colony-level selection, high relatedness with individual-level selection, low relatedness with colony-level selection, and low relatedness with individual-level selection (Fig. 2).

Artificial evolution was conducted for the four experimental treatments using a physics-based simulation of the s-bot robots (Mondada et al. (2004)) that precisely model their dynamical properties (Magenat et al. (2007)). At the end of the experiments the evolved genomes were transferred to the physical robots (Fig. 1, left) to evaluate whether the behavior of the real robots mimics that observed in simulation². Selection experiments were repeated in 20 independent selection lines (replicates of populations with newly generated genomes) for each experimental condition to determine whether different communication strategies could evolve under the different conditions (for a more detailed analysis of these results see Floreano et al. (2007)).

² The e-puck robots, which could similarly be used for these experiments are described in Appendix ???. The e-puck robot is an open-source platform that is also commercially available.

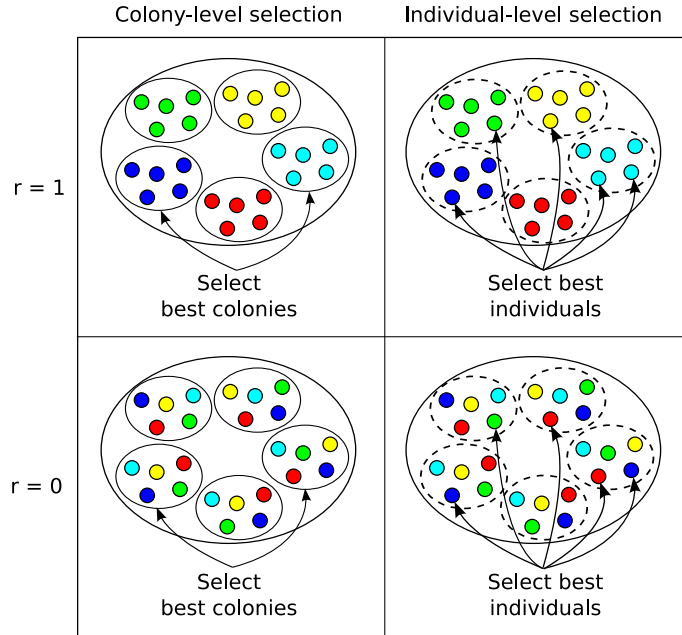


Fig. 2. Illustration of the colony composition and selection regime in the four treatments.

2 Experimental Setup

2.1 The Task

The foraging environment consisted of a $3m \times 3m$ arena that contained a food and a poison source each placed at $100cm$ from each corner. The food and poison sources constantly emitted red light that could be seen by robots in the whole foraging arena. A circular piece of gray paper was placed under the food source and a similar black paper under the poison source (see 1, right). At the beginning of each of these trials the robots were randomly placed in the foraging arena. During the trial, robots could communicate the presence of food or poison by producing blue light that could be perceived by other robots (light production was not costly).

The experiments were conducted using a physics-based simulator modeling s-bot robots (see 1, left) and later transferred to the real s-bot platform. The s-bot robots (simulated and real) were equipped with two tracks that could independently rotate in both directions, a translucent ring around the body that could emit blue light, a 360° vision system that could detect the amount and intensity of red and blue light. The paper circles laid beneath the sources

could be detected by infrared ground sensors located underneath the robot and thus allowed discrimination of food and poison (Fig. 1, left).

The simulated robots had a sensory-motor cycle of $50ms$ during which they used a neural controller to process the visual information and ground sensor input to set the direction and speed of the two wheels and control the emission of blue light accordingly during the next cycle. During each $50ms$ cycle, a robot gained one performance unit if it detected food with its ground sensors and lost one performance unit if it detected poison. The performance of each robot at the end of a $60s$ -long trial was computed as the sum of performance units obtained during that trial (1200 sensory motor cycles of $50ms$). Ten such trials were run to quantify the performance of each robot. Colony performance was equal to the average performance of all robots in the colony.

2.2 Neural Controller

The control system of each robot consisted of a feed-forward neural network with 10 input and 3 output neurons (Fig. 3). Each input neuron was connected to every output neuron with a synaptic weight representing the strength of the connection. One of the input neurons was devoted to the sensing of food and the other to the sensing of poison. Once a robot had detected the food or poison source, the corresponding neuron was set to 1. This value decayed to 0 by a factor of 0.95 at every cycle, thereby providing a short-term memory even after the robot’s sensors were no longer in contact with the gray and black paper circles placed below the food and poison. The remaining 8 neurons were used to encode the 360° visual input image, which was divided into four sections of 90° each. For each section, the average of the blue and red channels was calculated and normalized within the range of 0 and 1, such that one neural input was used for the blue and one for the red value³. The activation of each of the output neurons was computed as the sum of all inputs multiplied by the weight of the connection and passed through the continuous $\tanh(x)$ function (i.e., their output was between -1 and 1). Two of the three output neurons were used to control the two wheels, where the output value of each neuron gave the direction of rotation (forward if > 0 and backward if < 0) and velocity (the absolute value) of one of the two wheels. The third output neuron determined whether to emit blue light, which was the case if the output was greater than 0.

2.3 Artificial Evolution

The specifications of the robots’ neural controllers were encoded in artificial genomes (Fogel et al. (1990); Nolfi and Floreano (2001)). After each generation the genomes of the robots were subjected to mutation, sexual reproduction and recombination. The genotype of an individual encoded the synaptic weights of the neural network in a bit string. Each synaptic weight was encoded in 8 bits,

³ Note that no distance sensors were used on the robots. It was thus not possible for robots to directly detect the walls of the arena.

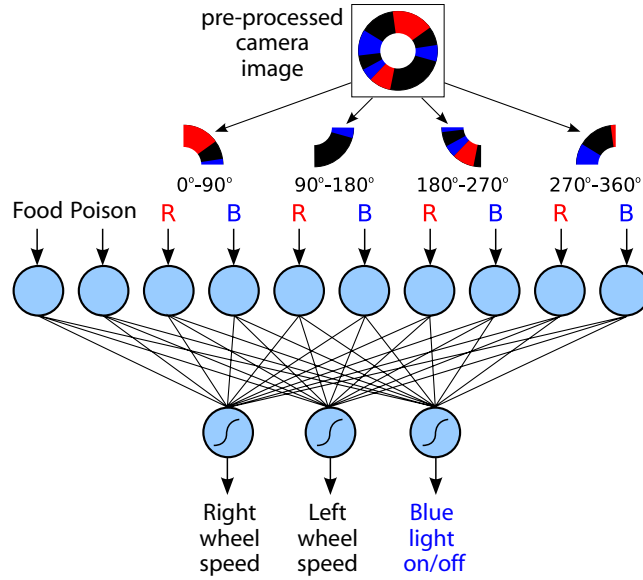


Fig. 3. Neural network architecture. The first two input neurons are activated when feeding on either food or poison. The omnidirectional camera image is pre-processed to filter out red and blue channels, divided into sections and input to the neural network as fractions of red or blue in each section (between 0 and 1). Three output neurons with *tanh*, asymptotic activation, receive weighted input from the 10 input units, and encode the speed of the wheels and whether to emit blue light.

giving 256 values that were mapped onto the interval $[-1, 1]$. The total length of the genetic string of an individual was therefore $8 \text{ bits} \times 10 \text{ input neurons} \times 3 \text{ output neurons}$ (i.e., 240 bits).

In all experiments we started with a population of completely naive robots (i.e., with randomly generated genomes that corresponded to randomly wired neural controllers) with no information about how to move and identify the food and poison sources. A population consisted of 100 colonies of 10 robots, resulting in a population of 1000 individuals. In the individual-level selection regime the genomes of robots with the 20% highest individual performance were selected to form the next generation, whereas in the colony-level selection regime, we randomly selected individuals from the 20% most efficient colonies (Fig. 2). This selected pool of robots was used to create the new generation of robots. To form colonies of related individuals $r = 1$, we randomly created (with replacement) 100 pairs of robots. A crossover operator was applied to their genomes with a probability of 0.05 at a randomly chosen point and one of the two newly

formed genomes was randomly selected and subjected to mutation (probability of mutation 0.01 for each of the 240 bits). The other genome was discarded. This procedure led to the formation of 100 new genomes that were each cloned 10 times to construct 100 new colonies of 10 identical robots. To form colonies of unrelated individuals $r = 0$, we followed the same procedure, but created 1000 pairs of robots in simulation from the selected pool of robots. These new robots were randomly distributed among the new colonies.

For each of the four treatments explored, selection experiments were repeated in 20 independent selection lines (replicates of populations with newly generated genomes) for 500 generations to determine whether different communication strategies could evolve. We quantified the benefits of communication by comparing colony performance with control colonies where robots were experimentally prevented from communicating (i.e., the blue lights were disabled).

2.4 Quantifying Behavior

To compare colony performance between treatments, we calculated the average performance of all colonies over the last 50 generations for each of the 20 replicates per treatment. The resulting mean performance values of the individual replicates were then used for comparisons.

In addition to performance measures, we were also interested in the evolved communicative behaviors, both in terms of signaling strategies (i.e., under what conditions robots lit up in blue) and response strategies (i.e., how perceiving blue light influenced the behavior of robots).

Signaling strategies were quantified by taking the difference between the average frequency of signaling near food and poison of all robots in a colony. The frequency was computed for each robot as the proportion of cycles spent near food or poison in which a robot was signaling. The signaling strategy value s can therefore vary from -1 to 1 , with a value of -1 indicating that robots signaled only when near the poison and a value of 1 that signaling occurred only when near the food. A value of 0 would indicate that robots were not more likely to signal near food or poison.

The level of response to blue light b was measured by placing each robot 35cm away from the bottom and left wall of the arena (coordinates: $x = 35, y = 35$) in a random orientation and a second stationary robot emitting blue light in the bottom left corner of the arena (coordinates: $x = 0, y = 0$). After 10 time-steps we checked the location of the moving robot relative to its original position, where a decrease in both coordinates ($x < 35$ and $y < 35$) was considered to be attraction, whereas an increase in both coordinates ($x > 35$ and $y > 35$) was counted as repulsion. All other outcomes of the test were discarded. This test was run 10 times for each robot, and the number of attractions and repulsions documented. The value of b was then calculated as the difference between the number of attractions and repulsions divided by 10. Therefore, if a robot was repelled by blue light in all tests, its score was -1 ; if it was always attracted, its score was 1 . A score of 0 indicates that there is no general tendency for the robot to be attracted or repulsed by blue light. Both s and b were calculated for

all the colonies in the population and averaged to give one value for each of the replicates of the experiment.

3 Honest Communication

Both in the control colonies where robots could and could not emit blue light, foraging efficiency greatly increased over the 500 generations of selection (Fig. 4, top). In all treatments robots evolved the ability to rapidly localize the food source, move in its direction and stay nearby. Both the degree of within-group relatedness and the level of selection significantly affected the overall performance of colonies, as can be seen from Fig. 4 (top left). This variation of performances in the control condition where robots could not emit blue light reflects differences in selection efficiency between the four treatments. These differences are due to a number of factors, such as the accuracy of evaluating the performance of a given genome, the strength of the correlation between a robot's performance and its likelihood of being selected, and the number of different genomes tested in the environment (see [Waibel et al.](#) for details).

Compared to control experiments, the ability to emit blue light resulted in a significantly greater colony efficiency in three out of the four treatments (Fig. 4). An analysis of the robot behavior revealed that this performance increment was associated with the evolution of effective systems of communication.

In colonies of related robots with colony-level selection, two distinct communication strategies evolved. In 12 of the 20 evolutionary replicates, robots preferentially produced light in the vicinity of the food, whereas in the other eight, robots tended to emit light near the poison (Fig. 5, 6). The response of robots to light production was tightly associated with these two signaling strategies, as shown by the strong positive association between the tendency of robots to be attracted to blue light and the tendency to produce light near the food rather than the poison source across the 20 replicates (Spearman's rank correlation test, $r_S = 0.74$, $p < 0.01$, Fig. 5, top left). Overall, robots were positively attracted to blue light in all the 12 replicates where they signaled in the vicinity of the food and repelled by blue light in 7 out of the 8 replicates where they had evolved a strategy of signaling near the poison. The communication strategy where robots signaled near the food and were attracted by blue light resulted in higher performance (mean \pm SD, 259.6 ± 29.5) than the alternate strategy of producing light near the poison and being repelled by blue light (197.0 ± 16.8 , Mann-Whitney test, $df = 6$, $p < 0.01$). However, once one type of communication was well-established, we observed no transitions to the alternate strategy over the last 200 generations. This is because a change in either the signaling or response strategy would completely destroy the communication system and result in a performance decrease. Thus, each communication strategy effectively constitutes an adaptive peak separated by a valley with lower performance values ([Wright \(1932\)](#)).

The possibility to produce blue light also translated into higher performance in two other treatments: high relatedness with individual-level selection and

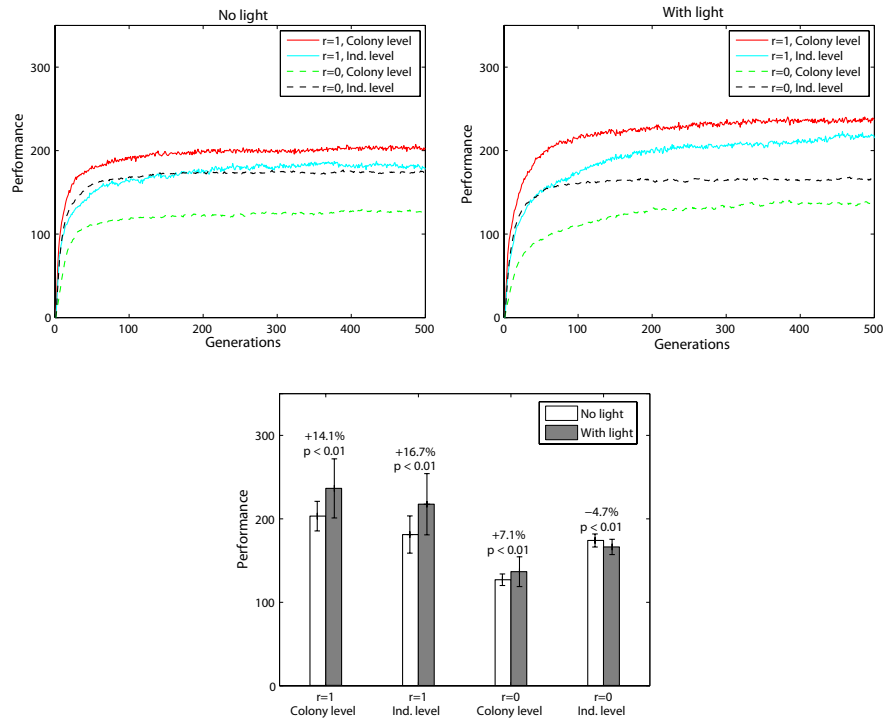


Fig. 4. Top left: Mean performance of robots in control colonies where robots could not emit blue light (20 replicates per treatment), where $r = 0$ stands for unrelated colonies and $r = 1$ for colonies composed of clones. Top right: Mean performance in colonies where robots could emit blue light (20 replicates per treatment). Bottom: Mean (\pm SD) performance of robots during the last 50 generations for each treatment when robots could versus could not emit blue light (20 replicates per treatment; percentages show differences in mean performance).

low relatedness with colony-level selection. In both cases, signaling strategies evolved that were similar to those observed in the selection experiments with high relatedness and colony-level selection (see Fig. 5, top right and bottom left). There was also a strong positive correlation between the tendency to signal close to food and being attracted to blue light, (high relatedness/individual-level selection: $r_S = 0.81$, $p < 0.01$; low relatedness/colony-level selection: $r_S = 0.60$, $p < 0.01$). Moreover, in both treatments the strategy of signaling close to food yielded higher performance than the alternative poison signaling strategy (both $p < 0.01$). However, when robots signaled near the poison, they were less efficient than in the treatments with high relatedness and colony-level selection (both $p < 0.01$). In the latter case robots signaled on average 82.3% of the time when detecting the poison, whereas the amount of poison-signaling was only 18.3%

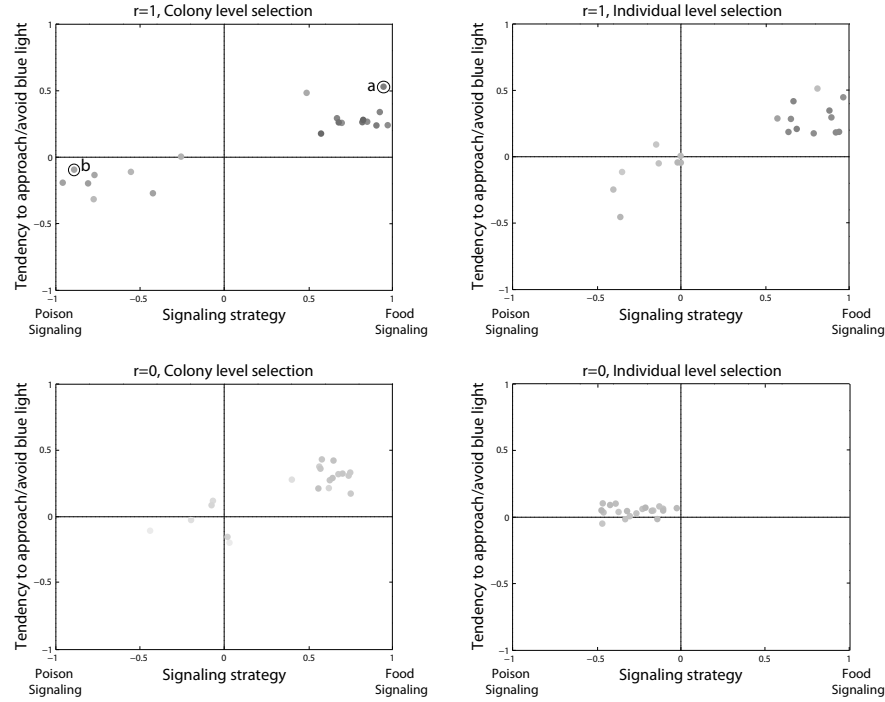


Fig. 5. Relationship between the signaling strategy and response to blue light. Each dot is the average for the 100 colonies in one replicate after 500 generations of selection. Positive values for the signaling strategy indicate a tendency to signal close to the food and negative values a tendency to signal close to the poison. Positive values for the tendency to approach/avoid blue light indicate an attraction to blue light and negative values an aversion. The darkness of the points is proportional to the mean performance. The signaling strategies of robots in replicates *a* and *b* are illustrated in Fig. 6.

($p < 0.001$) in groups with related individuals and individual-level selection and 24.0% ($p < 0.01$) in groups of low relatedness individuals subjected to colony-level selection. Interestingly, the less efficient poison signaling strategy permitted a switch to a food signaling strategy in the last 200 generations of selection in three replicates for related robots selected at the individual level and in one replicate for low relatedness robots selected at the colony level.

4 Deceptive Communication

The only treatment where the possibility to communicate did not translate into a higher foraging efficiency was when colonies were comprised of low relatedness

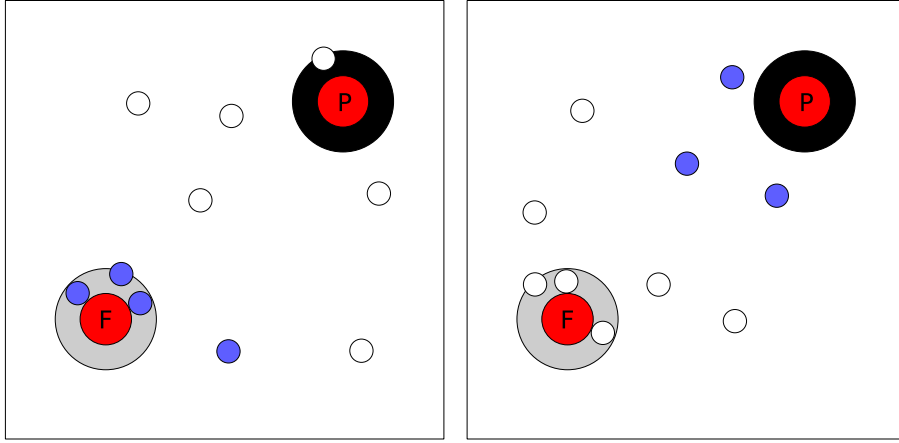


Fig. 6. Illustration of the two honest signaling strategies evolved in simulation. Left: robots (small circles) signal the presence of food F (illustrating strategy of replicate *a* in Fig. 5 top left). Right: robots signal the presence of poison P (illustrating strategy of replicate *b* in Fig. 5 top left).

robots subjected to individual-level selection (Fig. 5, bottom right). In this case, the ability to signal resulted in a deceptive signaling strategy associated with a significant decrease in colony performance compared to the situation where robots did not have the opportunity to emit blue light. An analysis of individual behaviors revealed that in all replicates robots tended to emit blue light when far away from the food. However, contrary to what one would expect, the robots still tended to be attracted rather than repelled by blue light (17 out of 20 replicates, binomial test z-score: 3.13, p-value < 0.01). A potential explanation for this surprising finding is that in an early stage of selection, blue light provided a useful cue about food location, hence selecting for a positive response by robots to blue light. Indeed, in another set of experiments (data not shown) we found that, when constrained to produce light randomly, robots were attracted by blue light because the higher level of blue light emission associated with the higher density of robots near food provided a useful cue about food location. Emission of light far from the food would then have evolved as a deceptive strategy to decrease competition near the food. Consistent with this view, there was a significant decrease during the last 200 generations in the tendency of robots to be attracted by blue light (Mann-Whitney test, $df = 18$, $p < 0.05$). This co-evolution between signalers and receivers with conflicting interests is similar to the processes described in chapter ? (see also [Mirolli and Parisi \(2008\)](#)).

5 Conclusion

With this chapter, we have provided a clear experimental demonstration of the role of kin structure of groups and the level of selection on the evolution of communication, a long-standing question in sociobiology (Maynard-Smith and Harper (2003); Searcy and Nowicki (2005)). Under natural conditions, most communication systems are also costly because of the energy required for signal production and/or increased competition for resources resulting from information transfer about food location (Maynard-Smith and Harper (2003)). Thus, altruistic communication is expected to occur principally in groups composed of kin or when selection takes place at the level of the group rather than the individual. Consistent with this view, most sophisticated systems of communication indeed occur in social animals forming family-based groups as exemplified by pheromone communication in social insects (Wilson (1971); Bourke and Franks (1995)) and quorum sensing in clonal groups of bacteria (Trivers (1971)). Humans are a notable exception but other selective forces such as reciprocal altruism and reputation-based systems of reciprocity may operate to favor altruism (Nowak and Sigmund (2005)) and costly communication (?).

This study demonstrates that sophisticated forms of communication, including altruistic communication and deceptive signaling can evolve in groups of robots with simple neural networks. Importantly, our results show that once a given system of communication has evolved, it may constrain the evolution of more efficient communication systems because it would require going through a stage where communication between signalers and receivers is perturbed. This finding supports the idea of the possible arbitrariness and imperfection of communication systems, which can be maintained despite their suboptimal nature. Similar observations have been made about evolved biological systems (Jacob (1981)), which are formed by the randomness of the evolutionary selection process, such as the existence of different dialects in the honey bee dance language (Dyer (2002)). Finally, our experiments demonstrate that the evolutionary principles governing the evolution of social life also operate in groups of artificial agents subjected to artificial selection, showing that transfer of knowledge from evolutionary biology can be useful to design efficient groups of cooperative robots⁴.

⁴ This research has been supported by the ECAgents project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract IST-2003-1940 and by the Swiss National Science Foundation grant nr. K-23K0-117914/1 on the Evolution of Altruistic Communication.

Bibliography

- A. F. G. Bourke and N. R. Franks. *Social Evolution in Ants*. Princeton University Press, Princeton, New Jersey, 1995.
- F. C. Dyer. Biology of the dance language. *Annual Review of Entomology*, 47: 917–949, 2002.
- F. Fiegna, N. Y. Yuen-Tsu, S. V. Kadam, and G. J. Velicer. Evolution of an obligate social cheater to a superior cooperator. *Nature*, 441:310–314, 2006.
- D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519, 2007.
- D. Fogel, L. Fogel, and V. Porto. Evolving neural networks. *Biological Cybernetics*, 63:487–493, 1990.
- A. S. Griffin, S. A. West, and A. Buckling. Cooperation and competition in pathogenic bacteria. *Nature*, 430:1024–7, August 2004.
- W. D. Hamilton. The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7, 1964.
- W. D. Hamilton. *Narrow Roads of Gene Land Vol. 1: Evolution of Social Behaviour*, 1996.
- F. Jacob. *Le Jeu des Possibles*. Paris: Librairie Arthème Fayard, 1981.
- R. A. Johnstone and A. Grafen. The continuous Sir Philip Sidney game: a simple model of biological signalling. *Journal of Theoretical Biology*, 156: 215–234, 1992.
- L. Keller, editor. *Levels of Selection in Evolution*. Princeton University Press, Princeton, 1999.
- S. Magnenat, M. Waibel, and A. Beyeler. Enki: The fast 2D robot simulator, 2007. <http://lis.epfl.ch/resources/enki>.
- J. Maynard-Smith. Honest signaling - the philip sidney game. *Animal Behavior*, 42:1034–1035, 1991.
- J. Maynard-Smith and D. Harper. *Animal Signals*. Oxford University Press, 2003.
- J. Maynard-Smith and E. Szathmàry. *The Major Transitions in Evolution*. New York: Oxford University Press, 1997.
- M. Mirolli and D. Parisi. How producer biases can favor the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16(1): 27–52, 2008.

- F. Mondada, G. C. Pettinaro, A. Guignard, I. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. Swarm-Bot: a new distributed robotic concept. *Autonomous Robots*, 17(2-3):193–221, 2004.
- S. Nolfi and D. Floreano. *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines*. MIT Press, Cambridge, MA, 2001. 2001 (2nd print), 2000 (1st print).
- M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437:1291–1298, October 2005.
- W. A. Searcy and S. Nowicki. *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems*. Princeton University Press, Princeton and Oxford, 2005.
- R. L. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57, 1971.
- M. Waibel, L. Keller, and D. Floreano. Genetic Team Composition and Level of Selection in the Evolution of Cooperation. *IEEE Transactions on Evolutionary Computation*. In Press.
- S. A. West, I. Pen, and A. S. Griffin. Cooperation and competition between relatives. *Science*, 296:72–75, april 2002.
- E. O. Wilson. *The Insect Societies*. Belknap Press, Cambridge, MA, 1971.
- E. O. Wilson. *Sociobiology: The New Synthesis*. Belknap Press, 1975.
- S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones, editor, *Proceedings of the VI International Congress of Genetics*, pages 356 – 366, 1932.
- A. Zahavi and A. Zahavi. *The Handicap Principle. A Missing Piece of Darwin's Puzzle*. New York: Oxford University Press, 1997.