

Evolutionary Conditions for the Emergence of Communication in Robots

Dario Floreano,¹ Sara Mitri,¹ Stéphane Magnenat,² and Laurent Keller^{3,*}

¹Laboratory of Intelligent Systems
Ecole Polytechnique Fédérale de Lausanne
Station 11

²Robotic Systems Laboratory 1
Ecole Polytechnique Fédérale de Lausanne
Station 9
CH-1015 Lausanne
Switzerland

³Department of Ecology and Evolution
University of Lausanne
Biophore
CH-1015 Lausanne
Switzerland

Summary

Information transfer plays a central role in the biology of most organisms, particularly social species [1, 2]. Although the neurophysiological processes by which signals are produced, conducted, perceived, and interpreted are well understood, the conditions conducive to the evolution of communication and the paths by which reliable systems of communication become established remain largely unknown. This is a particularly challenging problem because efficient communication requires tight coevolution between the signal emitted and the response elicited [3]. We conducted repeated trials of experimental evolution with robots that could produce visual signals to provide information on food location. We found that communication readily evolves when colonies consist of genetically similar individuals and when selection acts at the colony level. We identified several distinct communication systems that differed in their efficiency. Once a given system of communication was well established, it constrained the evolution of more efficient communication systems. Under individual selection, the ability to produce visual signals resulted in the evolution of deceptive communication strategies in colonies of unrelated robots and a concomitant decrease in colony performance. This study generates predictions about the evolutionary conditions conducive to the emergence of communication and provides guidelines for designing artificial evolutionary systems displaying spontaneous communication.

Results

In large and complex societies such as those found in social insects and humans, communication systems can be extremely sophisticated with individuals modulating their behavior in response to numerous social

signals. In addition to being a fundamental feature of the organization of highly social species, communication is also a key component ensuring their ecological success [2]. A powerful method of studying the evolution of communication would be to conduct experimental evolution [4, 5] 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. To circumvent this problem, we established an experimental system with colonies of robots that 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 Figure 1 and Experimental Procedures). Under such circumstances, foraging efficiency can potentially be increased if robots transmit information on food and poison location. However, such communication may also incur direct costs to the signaler because it can result in higher robot density and increased competition and interference nearby the food (i.e., spatial constraints around the food source allowed a maximum of eight robots out of ten to feed simultaneously and resulted in robots sometimes pushing each other away from the food). Thus, although beneficial to other colony members, signaling of a food location effectively can constitute a costly act [6, 7] because it decreases the food intake of signaling robots. This setting thus mimics the natural situation where communicating almost invariably incurs costs in terms of signal production or increased competition for resources [8].

We studied the behavior and performance of 100 colonies of 10 robots in selection experiments over 500 generations by using physics-based simulations that precisely model the dynamical properties of real robots. The specifications of the robots' neural controllers, which process sensory information and produce motor action, were encoded in artificial genomes [9, 10] (see Experimental Procedures and Figure S1 in the Supplemental Data available online). Between each generation, the genomes of the robots were subjected to mutation, sexual reproduction, and recombination (see Experimental Procedures). At the end of the experiments, we were able to successfully implement the evolved genome in real robots (Figure 1) that displayed the same behavior observed in simulation, demonstrating that the physics-based simulations allowed us to mimic the behavior of real robots (see Movie S1).

Studying why colony members convey information when it incurs costs requires consideration of the kin structure of groups [6, 11, 12] and the scale at which cooperation and competition occur (level of selection) [13, 14]. We therefore chose two kin structures (low and high relatedness) and two levels of selection (individual- and colony-level regimes) (see Experimental Procedures and Figure S2). In the individual-level selection regime, the genomes of the 20% robots with the highest individual performance ($n = 200$) were selected to form the next

*Correspondence: laurent.keller@unil.ch




Figure 1. Physical Robots

(A) The robot used for the experiments is equipped with a panoramic-vision camera and a ring of color LEDs used for emitting blue light. (B) Robots emitting blue light around the food object emitting red light.

generation, whereas in the colony-level selection regime, we randomly selected all robots ($n = 200$) from the 20% most efficient colonies. We created low-relatedness ($r = 0$) colonies by randomly grouping ten robots in the next generation of colonies and created high relatedness colonies ($r = 1$) by grouping ten genetically identical individuals. **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. **For each of the four treatments, selection experiments were repeated in 20 independent selection lines (replicates of populations with newly generated genomes) for determining whether different communication strategies could evolve.** 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). **For each treatment, we determined whether communication evolved and 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).** In all experiments, we started with 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.

In the control colonies where robots could not emit blue light, foraging efficiency greatly increased over the 500 generations of selection (Figure 2A). In each of the four experiments, robots evolved the ability to rapidly localize the food source, move in its direction, and stay nearby (more than half the robots found the food source within the first 30 s). **Both the degree of within-colony relatedness and the level of selection significantly affected the overall performance of colonies (Kruskal-Wallis test: $p < 0.001$).** **Colonies where robots were highly related and subjected to colony-level**

selection were more efficient than the three other types of colonies (Mann-Whitney test, $df = 18$, all $p < 0.001$). The two treatments with individual-level selection led to intermediate performance values (nonsignificantly different from each other $p = 0.39$ but different from the two other treatments, both $p < 0.001$). The lowest performance was achieved by robots in the low relatedness/colony-level selection treatment with performances significantly lower than in all other treatments (all $p < 0.001$). This variation of performances in the control condition where robots could not emit blue light reflects differences in selection efficiency among the four treatments (M. Waibel, L.K., and D.F., unpublished data).

In colonies where robots could produce blue light, foraging efficiency also greatly increased over the 500 generations of selection (Figure 2B). **Importantly, the ability to emit blue light resulted in a significantly greater colony efficiency compared to control experiments in three out of the four treatments (Figure 3).** 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 (see Figures 4 and 5 as well as Figure S3). 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, $r_s = 0.74$, $p < 0.01$; see Figure 4A). 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 seven out of the eight replicates where they had evolved




Figure 2. Performance
(A) Mean performance in control colonies where robots could not emit blue light (20 replicates per treatment).
(B) Mean performance of robots in colonies where robots could emit blue light (20 replicates per treatment).

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 \pm 29.5) than the alternate strategy of producing light near the poison and being repelled by blue light (197.0 \pm 16.8, Mann-Whitney test, $df = 6$, $p < 0.01$). This is probably because signaling near the food allows robots to signal in a more efficient, sustained way while they feed and because the food signal can easily be detected by other robots, even though the red light of the food is obscured by the robots feeding around it. Interestingly, 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 [15].




Figure 3. Performance Comparison
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).

The possibility to produce blue light also translated into higher performance in two other treatments: high relatedness with individual-level selection and 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 Figures 4B and 4C). 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. In the case of high relatedness and colony-level selection, robots signaled on average 82.3% of the time when detecting the poison, whereas the amount of poison signaling was only 18.3% (Mann-Whitney test, $df = 5$, $p < 0.001$) in colonies with related individuals and individual-level selection and 24.0% ($p < 0.01$) in colonies with low relatedness and 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.

The only treatment where the possibility to communicate did not translate into a higher foraging efficiency was when colonies comprised low-relatedness robots subjected to individual-level selection (Figure 4D). 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 could not 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, $p < 0.01$). A potential explanation for this surprising finding

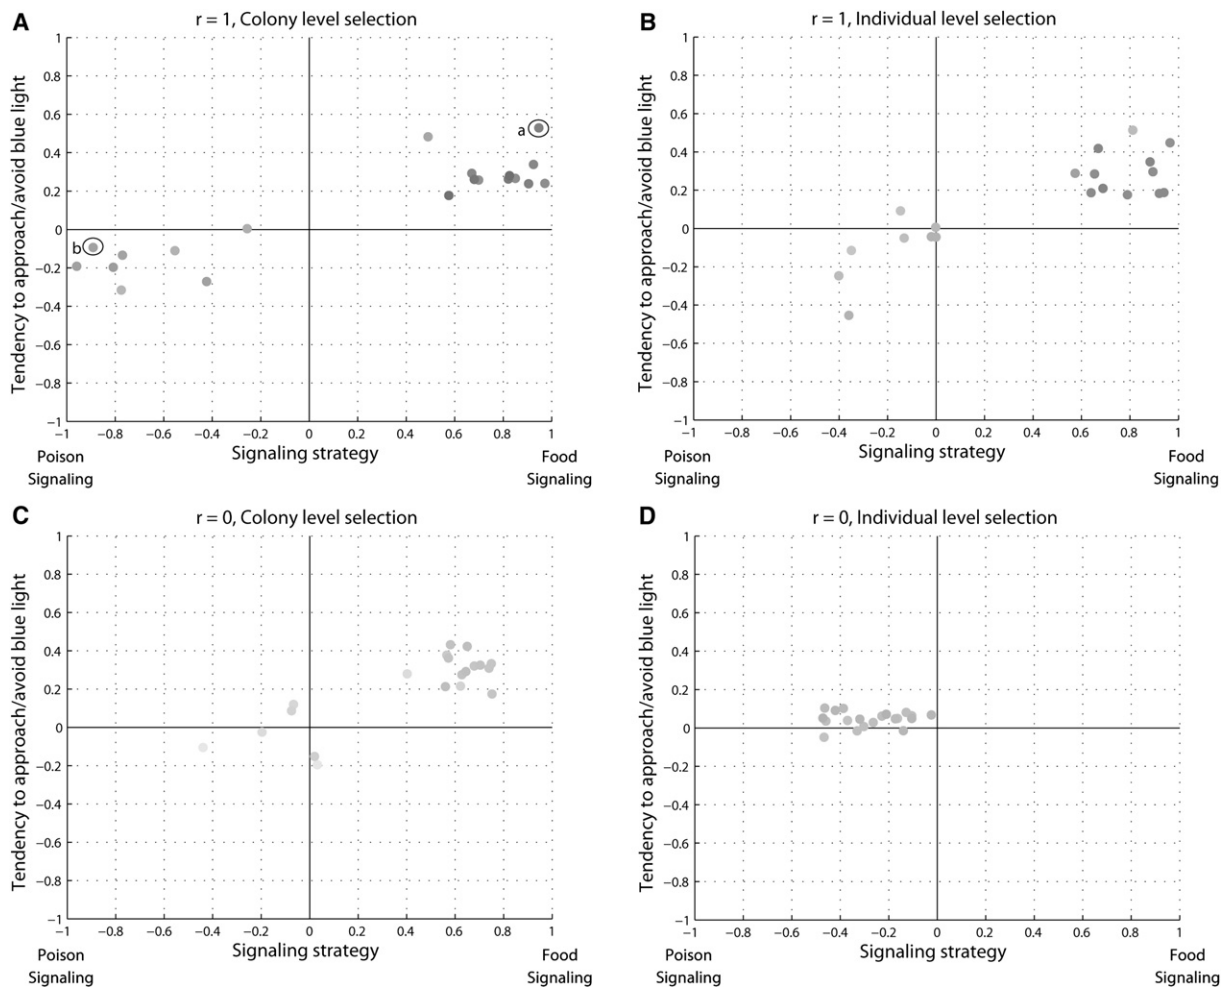


Figure 4. Relationship between Signaling Strategies and Behavioral Responses

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 indicate a tendency to signal close to the poison. Positive values for the tendency to approach or avoid blue light indicate an attraction to blue light, and negative values indicate an aversion (see Supplemental Data for definitions). The darkness of the points is proportional to the mean performance. The different signaling strategies of robots are shown in Figures 5A and 5B.

is that in an early stage of selection, robots randomly produced blue light, and this resulted in robots being selected to be attracted by blue light because blue light emission was greater near food where robots aggregated. 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 greater level of blue light emission associated with the greater 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 for decreasing competition near the food. Consistent with this view, the tendency of robots to be attracted by blue light significantly decreased during the last 200 generations (Mann-Whitney test, $df = 18$, $p < 0.05$).

Discussion

Our results provide a clear experimental demonstration of how the kin structure and the level of selection jointly

influence the evolution of cooperative communication. Under natural conditions, most communication systems are also costly because of the energy required for signal production or increased competition for resources resulting from information transfer about food location [3]. Thus, cooperative communication is expected to occur principally among kin or when selection takes place at a colony rather than an individual level. Consistent with this view, most sophisticated systems of communication indeed occur in animals forming kin groups as exemplified by pheromone communication in social insects [16, 17] and quorum sensing in clonal colonies of bacteria [18]. Humans are a notable exception, but other selective forces such as direct and reputation-based reciprocity may operate to favor cooperation [19] and costly communication.

This study demonstrates that sophisticated forms of communication including cooperative 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




Figure 5. Spatial Signaling Frequency

Measured in each area of the arena for robots from two colonies at generation 500.

(A) The colony was one where robots signal the presence of food (colony “a” in Figure 4A).

(B) In this colony, robots signal the presence of poison (colony “b” in Figure 4A). The darkness of each square is proportional to the amount of signaling in that area of the arena.

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 [20], which are formed by the randomness of the evolutionary selection process, leading, for example, to different dialects in the language of the honey-bee dance [21]. 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, indicating that transfer of knowledge from evolutionary biology can be useful for designing efficient groups of cooperative robots.

Experimental Procedures

Experimental Setup

For each colony of ten robots, we conducted ten foraging trials. At the beginning of each of these trials, the robots were randomly placed in a 300 × 300 cm foraging arena that contained a food and a poison source each placed at 100 cm from one of two opposite corners. The 10-cm-radius food and poison sources constantly emitted red light that could be seen by robots in the whole foraging arena.

All experiments were conducted with a physics-based simulator that accurately models the dynamical properties of real robots (Figure 1A). The robots were equipped with two tracks that could independently rotate in both directions, a translucent ring around the body that could emit blue light, and a 360° vision system that could detect the amount and intensity of red and blue light. A circular piece of gray paper with a radius of 25 cm was placed under the food source and a similar black paper under the poison source. These paper circles could be detected by infrared ground sensors located between the tracks underneath the robot and thus allowed discrimination of food and poison when robots were very close (Figure 1B).

The robots had a sensory-motor cycle of 50 ms during which they used a neural controller to process the visual information and used ground-sensor input to set the direction and speed of the two tracks and control the emission of blue light accordingly during the next

50 ms cycle. During each 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 trial was computed as the sum of performance units obtained during that trial (1200 sensory motor cycles of 50 ms), and the robot performance was quantified as the sum of performance units over all ten trials. Colony performance was equal to the average performance of all robots in the colony.

Neural Controller

The control system of each robot consisted of a feed-forward neural network with ten input and three output neurons. Each input neuron was connected to every output neuron with a synaptic weight representing the strength of the connection (Figure S1). 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 every 50 ms and thereby provided 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 eight neurons were used for encoding 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 value. 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 for controlling the two tracks, 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 tracks. The third output neuron determined whether to emit blue light; such was the case if the output was greater than 0.

The 30 genes of an individual each controlled the synaptic weights of one of the 30 neural connections. Each synaptic weight was encoded in 8 bits, giving 256 values that were mapped onto the interval $[-1, 1]$. The total length of the genetic string of an individual was therefore 8 bits × 10 input neurons × 3 output neurons (i.e., 240 bits).

Selection and Recombination

For each of the four treatments, selection experiments were repeated in 20 independent selection lines (replicates), each consisting of 100 colonies of 10 robots. In the individual-level selection treatment, we selected the best 20% of individuals from the

population of 1000 robots (Figure S2). This selected pool of 200 robots was used for creating 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) [22]. The other genome was discarded. This procedure led to the formation of 100 new genomes that were each cloned ten 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 from the selected pool of 200 robots. The 1000 new robots were randomly distributed among the 100 new colonies.

In the colony-level selection treatment, we followed exactly the same procedure as in the individual-level selection treatment, but the selected pool of 200 robots was formed with all of the robots from the best 20% of the 100 colonies (Figure S2).

Supplemental Data

Supplemental Data include additional Experimental Procedures, three figures, and one movie and are available with this article online at <http://www.current-biology.com/cgi/content/full/17/6/514/DC1/>.

Acknowledgments

We thank Michel Chapuisat, Philippe Christe, Andy Gardner, Rob Hammond, Christoph Hauert, Sara Helms Cahan, Karen Parker, Rick Riolo, Ian Sanders, Claus Wedekind, and two anonymous reviewers for useful comments on the paper. This research has been supported by the ECAgents project funded by the Future and Emerging Technologies (IST-FET) program of the European Community under European Union's Framework Programme for Research and Technological Development contract IST-2003-1940 and by the Swiss National Science Foundation.

Received: November 22, 2006

Revised: January 11, 2007

Accepted: January 17, 2007

Published online: February 22, 2007

References

1. Maynard-Smith, J., and Szathmary, E. (1997). *The Major Transitions in Evolution* (New York: Oxford University Press).
2. Wilson, E.O. (1975). *Sociobiology: The New Synthesis* (Cambridge, MA: Belknap Press).
3. Maynard-Smith, J., and Harper, D. (2003). *Animal Signals* (Oxford: Oxford University Press).
4. Griffin, A.S., West, S.A., and Buckling, A. (2004). Cooperation and competition in pathogenic bacteria. *Nature* 430, 1024–1027.
5. Fiegna, F., Yuen-Tsu, N.Y., Kadam, S.V., and Velicer, G.J. (2006). Evolution of an obligate social cheater to a superior cooperator. *Nature* 441, 310–314.
6. Hamilton, W.D. (1964). The genetical evolution of social behaviour, part I. *J. Theor. Biol.* 7, 1–16.
7. Lehmann, L., and Keller, L. (2006). The evolution of cooperation and altruism – a general framework and a classification of models. *J. Evol. Biol.* 19, 1365–1376.
8. Zahavi, A., and Zahavi, A. (1997). *The Handicap Principle. A Missing Piece of Darwin's Puzzle* (New York: Oxford University Press).
9. Fogel, D., Fogel, L., and Porto, V. (1990). Evolving Neural Networks. *Biol. Cybern.* 63, 487–493.
10. Nolfi, S., and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (Cambridge, MA: MIT Press).
11. Maynard-Smith, J. (1991). Honest signaling - The Philip Sidney game. *Anim. Behav.* 42, 1034–1035.
12. Johnstone, R.A., and Grafen, A. (1992). The continuous Sir Philip Sidney game: A simple model of biological signaling. *J. Theor. Biol.* 156, 215–234.
13. West, S.A., Pen, I., and Griffin, A.S. (2002). Cooperation and competition between relatives. *Science* 296, 72–75.
14. Keller, L., ed. (1999). *Levels of Selection in Evolution* (Princeton, NJ: Princeton University Press).
15. Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the VI International Congress of Genetics*, D.F. Jones, ed., pp. 356–366.
16. Wilson, E.O. (1971). *The Insect Societies* (Cambridge, MA: Belknap Press).
17. Bourke, A.F.G., and Franks, N.R. (1995). *Social Evolution in Ants* (Princeton, NJ: Princeton University Press).
18. Keller, L., and Surette, M.G. (2006). Communication in bacteria. *Nat. Rev. Microbiol.* 4, 249–258.
19. Nowak, M.A., and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature* 437, 1291–1298.
20. Jacob, F. (1981). *Le Jeu des Possibles* (Paris: Librairie Arthème Fayard).
21. von Frisch, K. (1967). *The Dance Language and Orientation of Bees* (Cambridge, MA: Harvard University Press).
22. Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems* (Ann Arbor, MI: University of Michigan Press).