

Evolutionary Conditions for the Emergence of Communication in Robots

Dario Floreano¹, Sara Mitri¹, Stéphane Magnenat² and Laurent Keller^{3*}

¹ Laboratory of Intelligent Systems, Ecole Polytechnique Fédérale de Lausanne, EPFL, Station 11, CH-1015 Lausanne, Switzerland.

² Robotic Systems Laboratory 1, Polytechnique Fédérale de Lausanne, EPFL, Station 9, CH-1015 Lausanne, Switzerland.

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

* Correspondence:

Laurent Keller

Department of Ecology and Evolution

University of Lausanne

Biophore

CH-1015 Lausanne

Switzerland

Laurent.keller@unil.ch. Tel: +41 21 692 41 73

Summary

Living organisms acquire, store, process and transmit information and, as such, information transfer plays a central role in the biology of most species [1]. This is particularly true in social species where communication plays a pivotal organising role, allowing the transfer of vital information among colony members, for example to detect predators and find food sources [2]. While much is known on the neurophysiological processes by which signals are produced, conducted, perceived and interpreted, the conditions conducive to the evolution of communication and the paths by which reliable systems of communication become established remain largely unknown because communication does not leave a fossil record. This is a particularly challenging problem because efficient communication requires tight co-evolution between the signal emitted and the response elicited [3]. Here 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 are composed of genetically similar individuals and/or when selection acts at the colony level. We identified several distinct systems of communication 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 on the evolutionary conditions that may lead to the emergence of sophisticated forms of communication, including cooperative communication and deceptive signalling, and it provides guidelines for designing artificial evolutionary systems that may display 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 behaviour in response to numerous social signals. In addition to being a fundamental feature of the organisation of highly social species, communication is also a key component ensuring their ecological success [2]. A powerful mean to study the evolution of communication would be to conduct experimental evolution [4, 5] in a species with elaborate social organisation. 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 section). 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 signaller 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 8 robots out of 10 to feed simultaneously and resulted in robots sometimes pushing each other away from the food). Thus, while beneficial to other colony members, signalling of a food location effectively can constitute a costly act [6, 7] because it decreases the food intake of signalling 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 behaviour 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 section, Figure S1). Between each generation the genomes of the robots were

subjected to mutation, sexual reproduction and recombination (see Experimental Procedures section). At the end of the experiments we were able to successfully implement the evolved genome in real robots (Figure 1) that displayed the same behaviour observed in simulation, demonstrating that the physics-based simulations allowed us to mimic the behaviour of real robots (see video in Supplemental data).

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) [14, 15]. We therefore chose two kin structures (low and high relatedness) and two levels of selection (individual and colony level regimes) (see Experimental Procedures section, 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 generation, whereas in the colony-level selection regime, we randomly selected all robots (N=200) from the 20% most efficient colonies. Low relatedness ($r = 0$) colonies were created by randomly grouping 10 robots in the next generation of colonies while high relatedness colonies ($r = 1$) were formed by grouping 10 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) to determine 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 localise the food source, move in its direction and stay nearby (more than half the robots found the food source within the first 30 seconds). 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 (non-significantly 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 between the four treatments (M. Waibel, L. Keller and D. Floreano, 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 behaviour 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 Figure 4, 5 and S3 in supplemental data). The response of robots to light production was tightly associated with these two signalling 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 signalled 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 signalling near the poison. The communication strategy where robots signalled near the food and were attracted by blue light resulted in higher performance (mean \pm s.d., 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, $df = 6$, $P < 0.01$). This is probably because signalling near the food is more efficient as robots can signal in a 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 signalling 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].

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, signalling strategies evolved that were similar to those observed in the selection experiments with high relatedness and colony-level selection (see Figure 4B, C). 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 signalling close to food yielded higher performance than the alternative poison signalling strategy (both $P < 0.01$). However, when robots signalled 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 signalled on average 82.3% of the time when detecting the poison, whereas the amount of poison-signalling 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 signalling strategy permitted a switch to a food signalling 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 were comprised of low relatedness robots subjected to individual-level selection (Figure 4D). In this case, the ability to signal resulted in a deceptive signalling strategy associated with a significant decrease in colony performance compared to the situation where robots could not emit blue light. An analysis of individual behaviours 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 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, the tendency of robots to be attracted by blue light significant 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 and/or increased competition for resources resulting from information transfer about food location³. Thus, cooperative communication is expected to occur principally among kin or when selection takes place at colony rather than 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 favour cooperation [19] and costly communication.

This study demonstrates that sophisticated forms of communication, including cooperative communication and deceptive signalling 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 signallers 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 honey bee dance language [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, showing that transfer of knowledge from evolutionary biology can be useful to design efficient groups of cooperative robots.

Experimental Procedures

Experimental setup

For each colony of 10 robots, we conducted 10 foraging trials. At the beginning of each of these trials the robots were randomly placed in a 300 cm × 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 using a physics-based simulator which 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 grey 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 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 10 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 10 input and 3 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, thereby providing a short-term memory even after the robot's sensors were no longer in contact with the grey 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 normalised 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 to control 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, which 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 \times 10 input neurons \times 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 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) [22]. 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 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 including detailed methods, two figures and one movie are available at www.current-biology.com/.

Acknowledgements 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 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. The authors declare no competing financial interests.

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* (Belknap Press).

3. Maynard-Smith, J. and Harper, D. (2003). *Animal Signals* (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. *Journal of Theoretical Biology* 7.
7. Lehmann, L. and Keller, L. (2006). The evolution of cooperation and altruism – a general framework and a classification of models. *Journal of Evolutionary Biology* 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. *Biological Cybernetics* 63, 487-493.
10. Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics* (Cambridge, MA: MIT Press).
11. Maynard-Smith, J. (1991). Honest signaling - The Philip Sidney game. *Animal Behavior* 42, 1034-1035.
12. Johnstone, R. A. and Grafen, A. (1992). The continuous Sir Philip Sidney game: a simple model of biological signaling. *Journal of Theoretical Biology* 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 University Press, Princeton).
15. Wright, S. (1932). In *Proceedings of the VI International Congress of Genetics*, (ed. Jones, D. F.), pp. 356–366.
16. Wilson, E. O. (1971). *The Insect Societies* (Belknap Press, Cambridge, MA).
17. Bourke, A. F. G. and Franks, N. R. (1995). *Social Evolution in Ants* (Princeton University Press, Princeton, New Jersey).
18. Keller, L. and Surette, M. G. (2006). Communication in bacteria. *Nature Reviews Microbiology* 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* (Harvard Univ. Press, Cambridge, Massachusetts).
22. Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems* (University of Michigan Press, Ann Arbor).

Figure 1. Physical robots. **A**, The robot used for the experiments is equipped with a panoramic vision camera and a ring of colour LEDs used to emit blue light. **B**, Robots emitting blue light around the food object emitting red light.

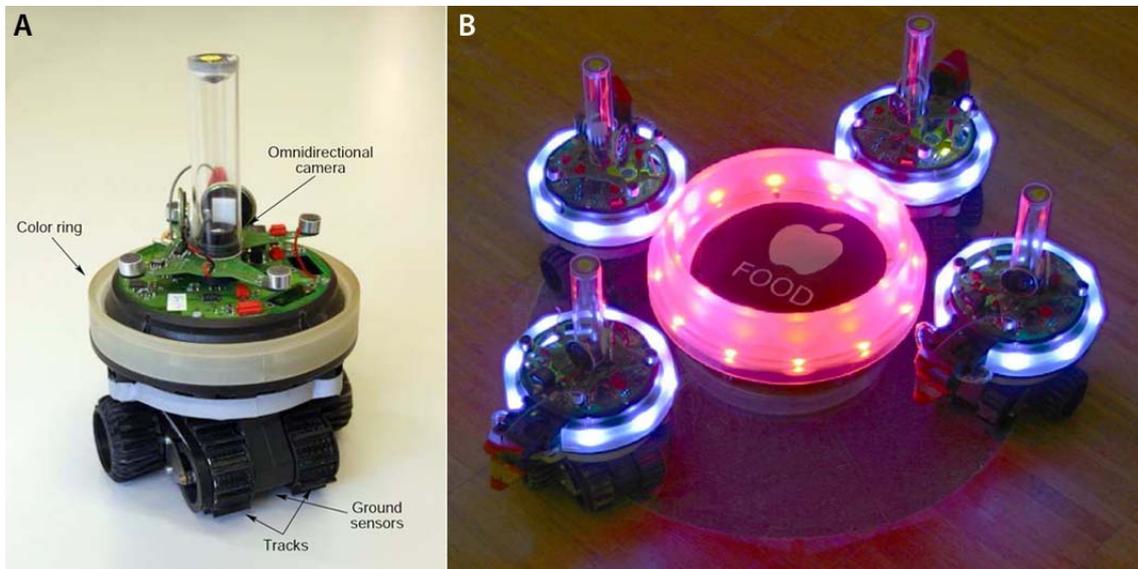


Figure 2. Performance. **A.** Mean performance in colonies where robots could not emit blue light (20 replicates per treatment). **B.** Mean performance of robots in control colonies where robots could emit blue light (20 replicates per treatment).

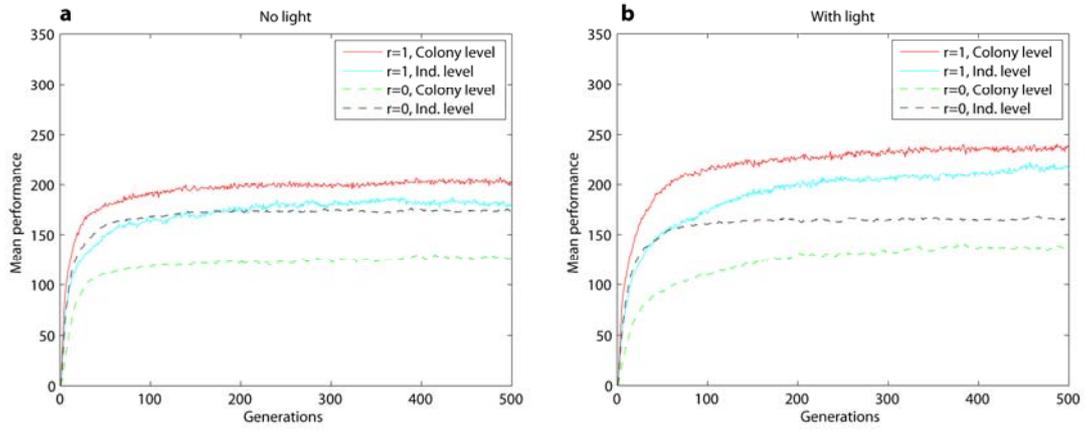


Figure 3. Performance comparison. Mean (\pm s.d.) performance of robots during the last 50 generations for each treatment when robots could vs. could not emit blue light (20 replicates per treatment).

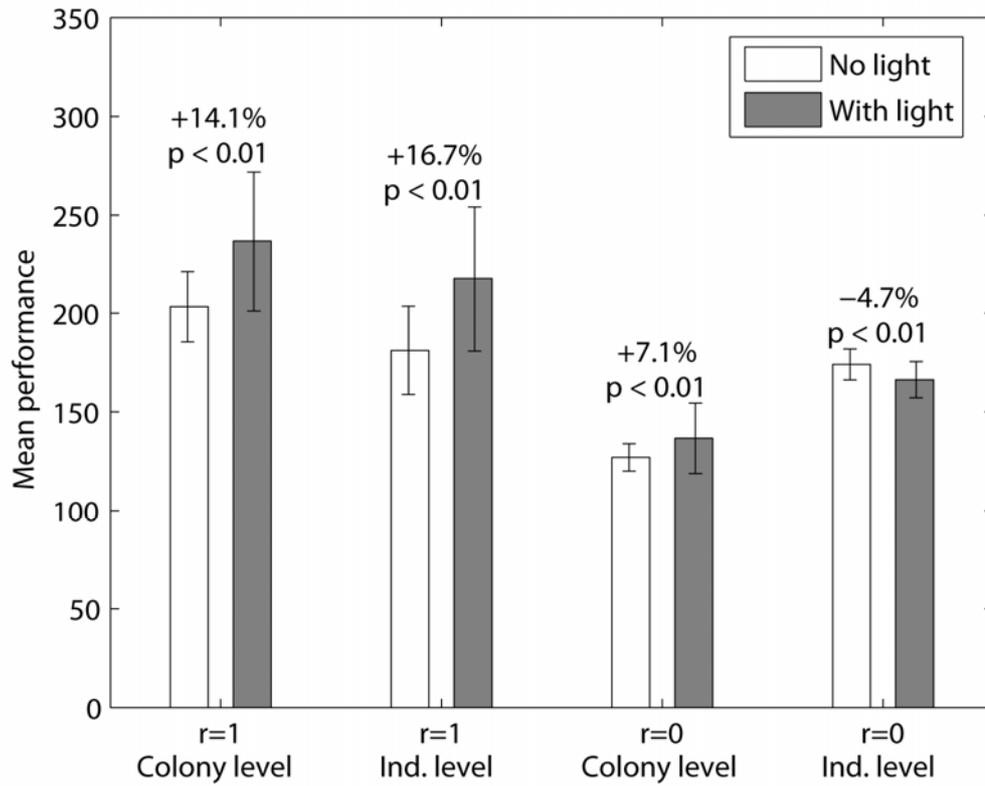


Figure 4. Relationship between signalling strategies and behavioural responses.

Each dot is the average for the 100 colonies in one replicate after 500 generations of selection. Positive values for the signalling 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 (see supplemental data for definitions). The darkness of the points is proportional to the mean performance. The different signalling strategies of robots are shown in Figure 5A and B.

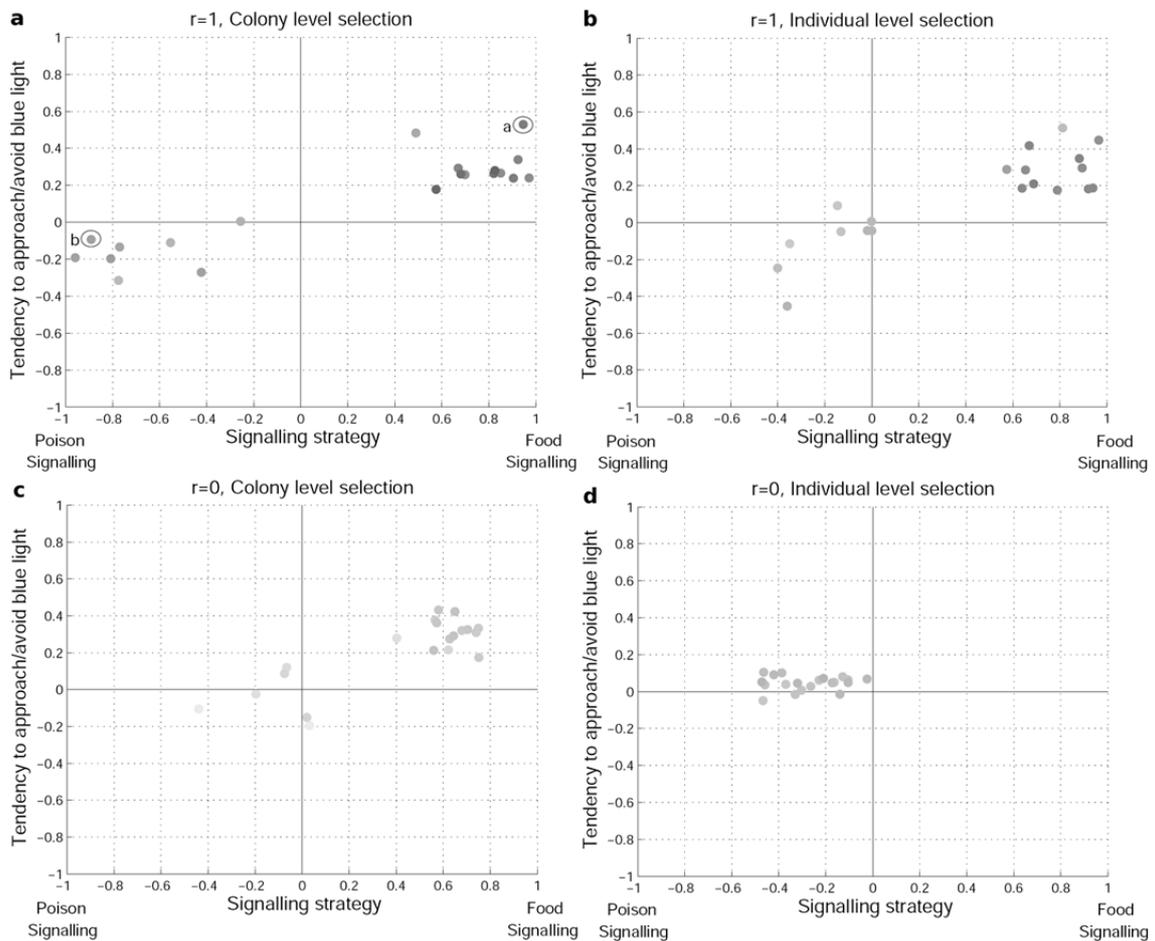
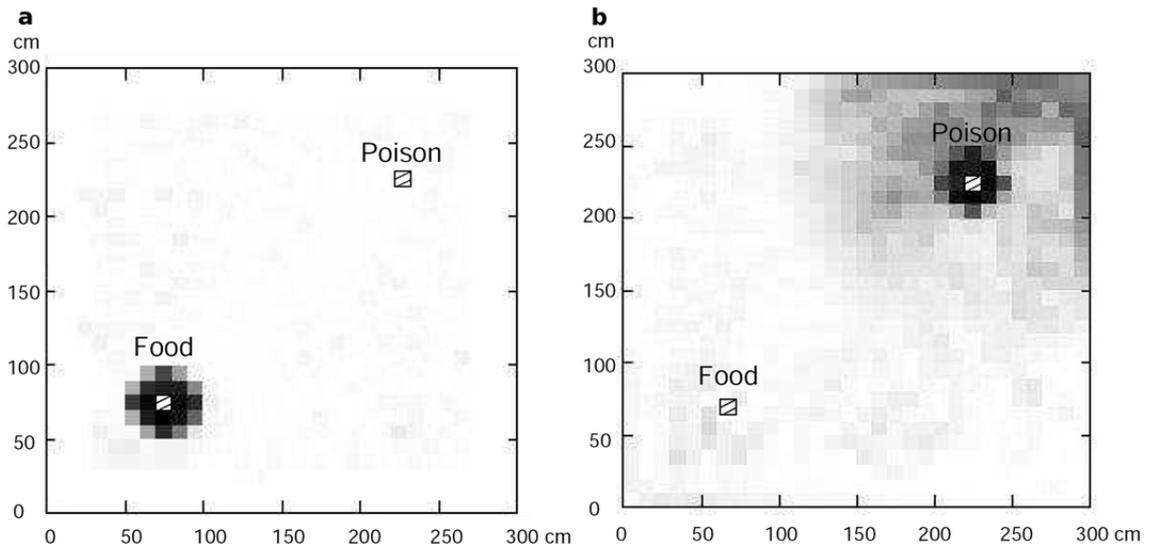


Figure 5. Spatial signalling frequency. Measured in each area of the arena for robots from two colonies at generation 500. **A**, The colony was one where robots signal the presence of food (colony *a* in Figure 4A). **B**, In this colony robots signal the presence of poison (colony *b* in Figure 4A). The darkness of each square is proportional to the amount of signalling in that area of the arena.



Supplemental Data

Supplemental Text. Data Analysis

To compare colony performance between treatments in the simulations, we calculated the average performance of the 100 colonies over the last 50 generations for each of the 20 replicates per treatment. The resulting 20 values per treatment were compared with non-parametric (Kruskal-Wallis and Mann-Whitney) tests as some of the data did not follow a normal distribution (Kolmogorov-Smirnov test).

The signalling strategy s was quantified by estimating the average frequency of signalling near food F and poison P of the 10 robots present in a colony. This was quantified by:

$$s = \frac{1}{10} \sum_{r=1}^{10} \left(\frac{1}{N^F} \sum_{n=1}^{N^F} b_m^F - \frac{1}{N^P} \sum_{n=1}^{N^P} b_m^P \right) \quad (1)$$

where N^F represents the total number of cycles spent near the food (i.e., over the paper circle) and N^P near the poison, and b_m^F and b_m^P whether robot r was emitting blue light (a value of 1 if it was and 0 if it was not) at cycle n near the food or poison, respectively. The signalling strategy value s can therefore vary from -1 to 1, with a value of -1 indicating that robots signalled only when near the poison and a value of 1 that signalling 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 tendency b of robots to be attracted by blue light was quantified by placing robots at 50 cm from a stationary robot emitting blue light. After 10 time-steps we checked the location of the moving robot r relative to its original position, where a decrease in the distance from the signalling robot was counted as attraction (increment a_r by 1) and an

increase in distance as avoidance (increment v_r by 1). The test was run 4 times for each of the 10 robots of a colony and b was calculated as follows:

$$b = \frac{1}{4 \times 10} \sum_{r=1}^{10} (a_r - v_r), \quad (2)$$

representing an average over the 4 tests for the 10 robots. Therefore, if all robots were repelled by blue light in all tests, the score was -1; if they were all attracted, the score was 1. A score of 0 would indicate that there is no general tendency for the robots to be attracted or repulsed by blue light. Both s and b were calculated for the 100 colonies in the population and averaged to give one value for each of the 20 replicates of the experiment.

Figure S1. 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 sigmoid, asymptotic activation receive weighted input from the 10 input units, which encode the speed of the tracks and whether to emit blue light.

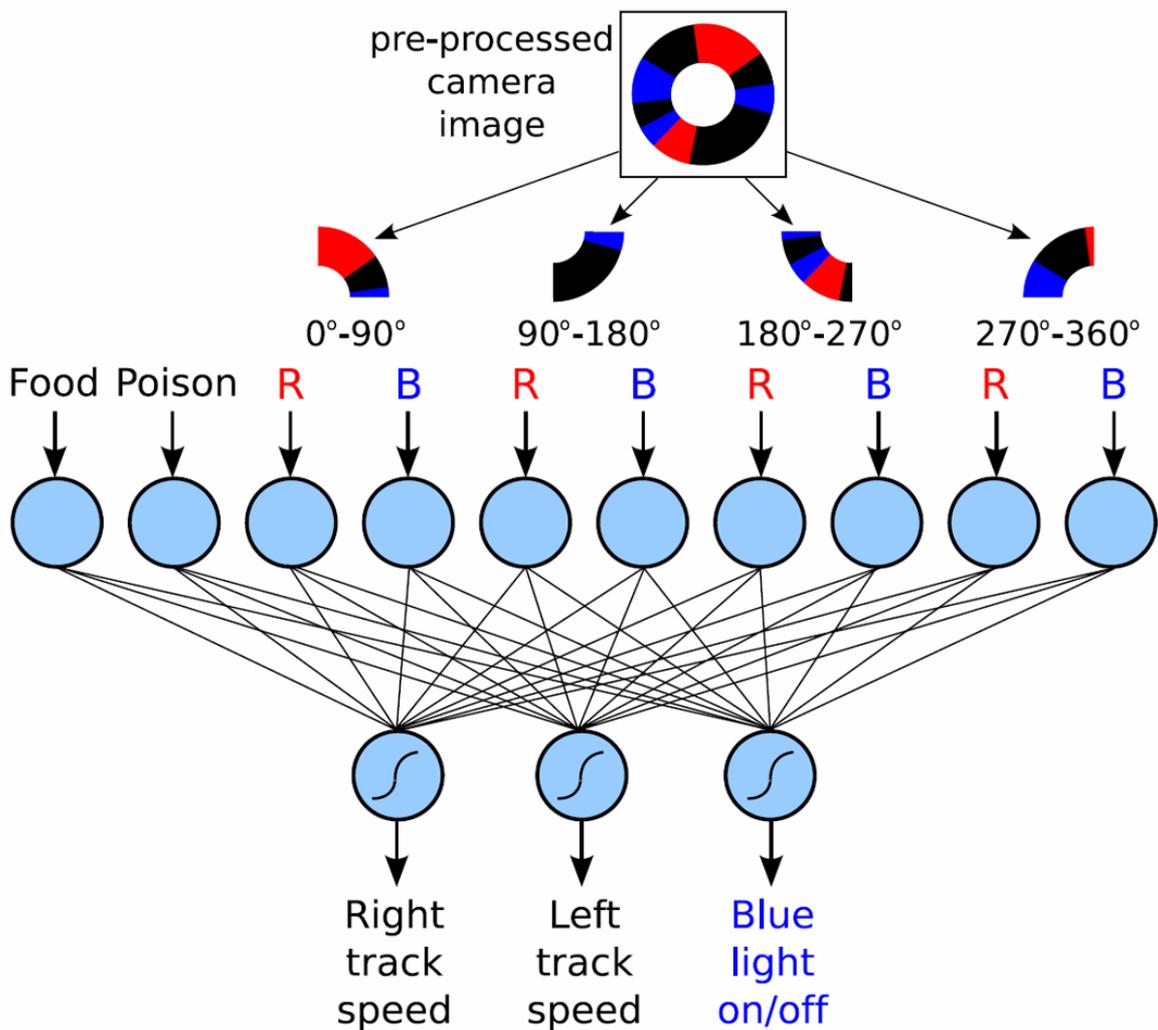


Figure S2. Four treatments. Illustration of the colony composition and selection regime in the four treatments.

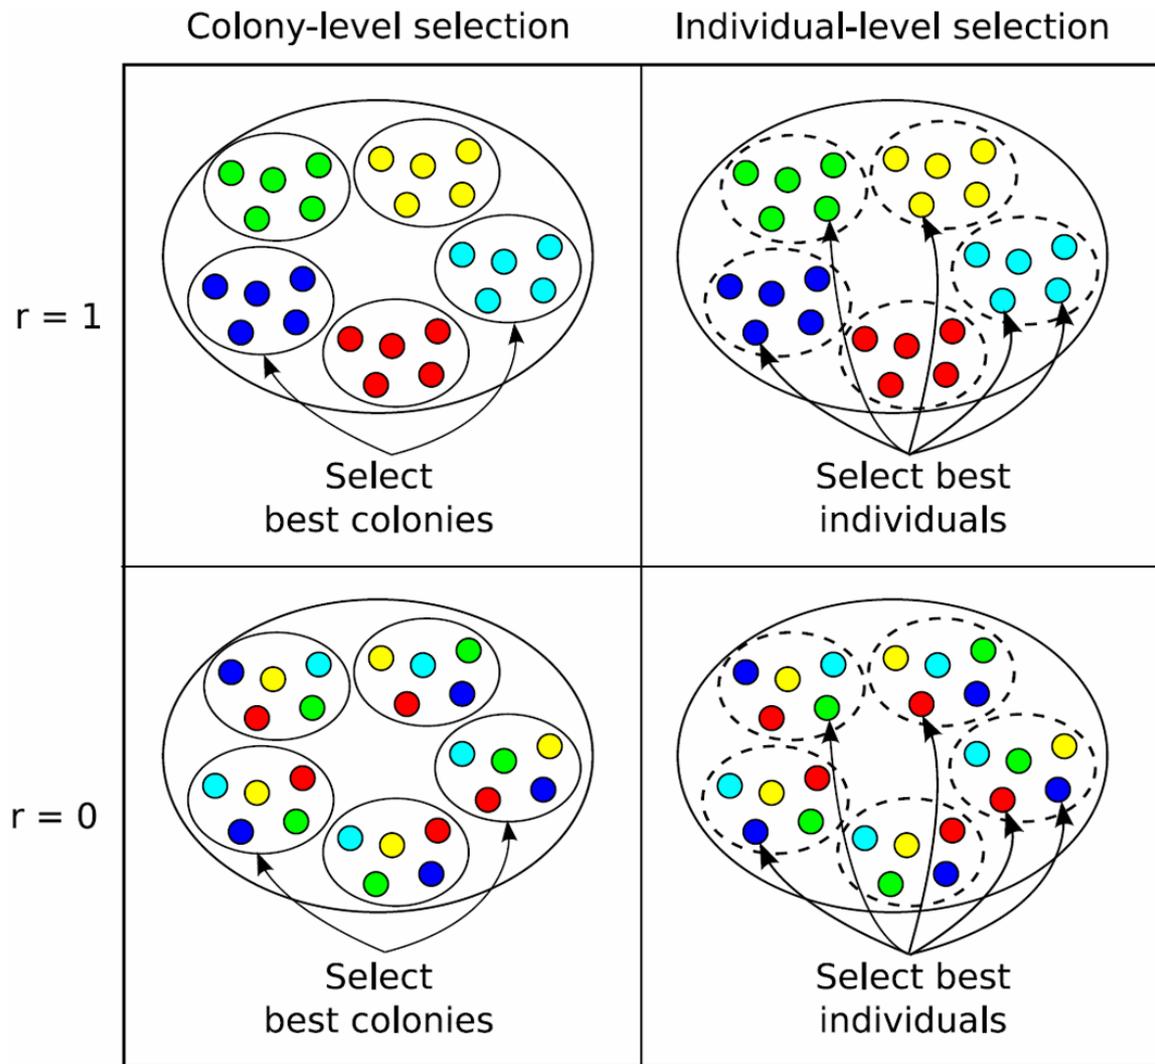
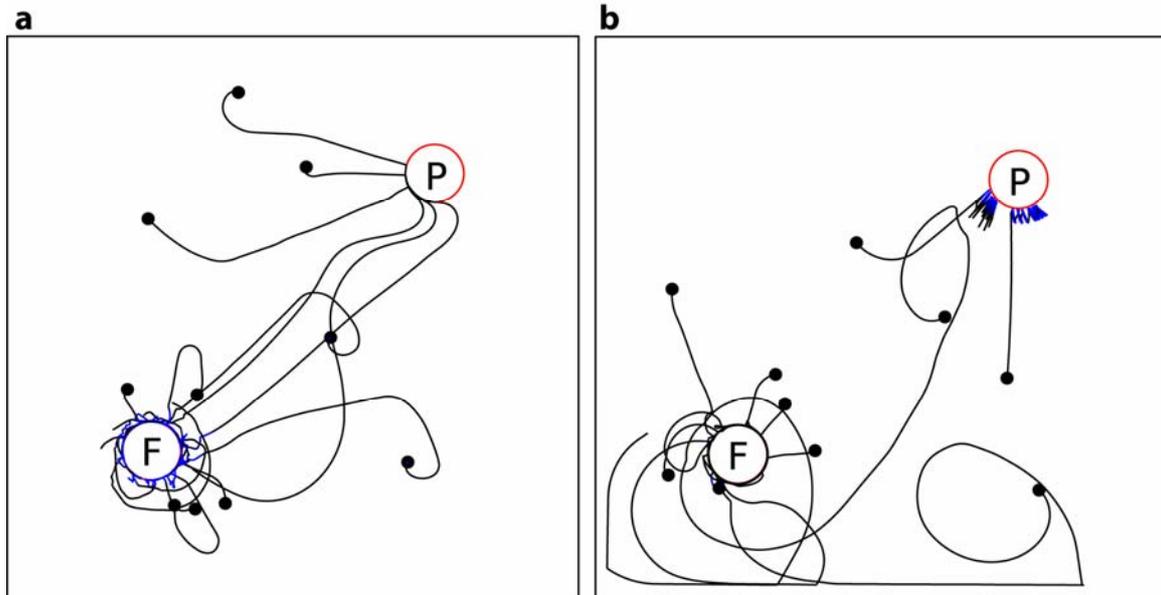


Figure S3. Trajectories and signalling strategies. The trajectory of each of the 10 robots is shown, where the dots represent the starting point of each robot and the blue sections of the lines show when the robots were signalling. **a**, food signalling strategy; **b**, poison signalling strategy.



Supplemental Movie M1. Evolved robot signalling behaviour. Example of the collective and signalling behaviours in a group of 6 robots (For technical reasons, only 6 physical robots were available, instead of the 10 we discuss in the paper. The simulation was therefore rerun with 6 robots prior to the transfer). The robot controller was taken from a colony selected with colony-level selection and high relatedness at generation 500. In this trial, the signalling strategy is to emit blue light near the food (left object marked with an apple not visible to the robots).