

Analysing an Evolved Robotic Behaviour Using a Biological Model of Collegial Decision Making

Gianpiero Francesca¹, Manuele Brambilla¹, Vito Trianni²,
Marco Dorigo¹, and Mauro Birattari¹

¹ IRIDIA, CoDE, ULB, Brussels, Belgium

{gianpiero.francesca,mbrambill,mdorigo,mbiro}@ulb.ac.be,

² ISTC-CNR, Rome, Italy
vito.trianni@istc.cnr.it

Abstract. Evolutionary robotics can be a powerful tool in studies on the evolutionary origins of self-organising behaviours in biological systems. However, these studies are viable only when the behaviour of the evolved artificial system closely corresponds to the one observed in biology, as described by available models. In this paper, we compare the behaviour evolved in a robotic system with the collegial decision making displayed by cockroaches in selecting a resting shelter. We show that artificial evolution can synthesise a simple self-organising behaviour for a swarm of robots, which presents dynamics that are comparable with the cockroaches behaviour.

1 Introduction

In recent studies, evolutionary robotics (ER, see [1]) has been used as an instrument to investigate the evolutionary conditions for the emergence of adaptive behaviour in groups of interacting agents. The main motivation behind these studies is that the evolution of certain adaptive traits and behavioural responses is tightly linked to ecological and social conditions. These conditions are extremely difficult or impossible to be controlled and replicated with empirical studies [2], while they can be completely managed in ER studies. The use of ER to analyse adaptive behaviours has been demonstrated in several occasions. For instance, the effects of genetic relatedness on the evolution of cooperative communication strategies can be investigated by systematically varying the composition of interacting groups [3,4]. Similarly, thanks to a simple ER experiment, it has been shown that the effect of stochastic variations in the evolutionary history could be at the basis of the emergence of diverse signalling strategies [5].

At the same time, ER represents a powerful design tool for the synthesis of collective, self-organising behaviours in swarms of robots [6]. It provides an automatic design methodology to synthesise the individual mechanisms leading to an optimal group response, according to a user-defined performance metric. Additionally, ER can shed light on the evolutionary pressures leading to the emergence of observed collective behaviours. However, it is necessary to understand whether or not the target behaviour can be evolved in the artificial system, and whether it displays dynamics comparable with the natural counterpart.

In this paper, we perform this first step, that is, the validation of an ER system with respect to collegial decision making by cockroaches in selecting a resting shelter [7]. Cockroaches (*Blattella germanica*) are gregarious insects that manifest cooperative behaviour in selecting a resting site: whenever more than one site is present, the insects collectively choose to aggregate in one single place (provided that it is large enough to host them all). Experimental studies allowed to determine which are the social influences that lead to such a collegial decision-making process, and a dynamical model has been developed (see [7] and Section 3.1 for more details). The identified mechanisms have been successfully exploited for designing collective aggregation and decision-making behaviours in swarms of robots [8,9,10], allowing also mixed insect-robot experiments [11]. However, to the best of our knowledge, there has been no attempt to study the evolution of a similar decision-making behaviour in swarms of robots. In this paper, we demonstrate that similar collegial decisions can be evolved in an artificial system. Our goal is to (i) verify the evolvability of the collegial decision making in the artificial system, and (ii) determine whether the dynamics of the system correspond qualitatively and quantitatively to the ones predicted by the biological model [7]. This will allow us to determine whether or not evolutionary robotics is suitable for formulating hypotheses about the evolutionary pressures that resulted in collective decision-making in cockroaches.

The paper is organised as follows. In Section 2, we describe in detail the experimental setup for the ER experiments. In Section 3, we discuss the results obtained from the evolutionary experiments with respect to the evolvability of the decision-making behaviour in a robotic system. In Section 3.1, we present the dynamical model proposed in [7], and we discuss the methodology that leads us to fit the evolved behaviour to the model. In Section 3.2, we present a comparison of the dynamics of the evolved behaviour with the ones predicted by the model. Section 4 concludes the paper with some final remarks.

2 Experimental Setup

We study the evolution of collegial decision making in a swarm of robots that have to aggregate in one of two areas within the experimental arena. Our experimental setup is based on the one used in Amé et al. [7]. The robots operate in a dodecagonal arena (Figure 1 left) of area 4.91 m^2 surrounded by walls. The floor of the arena is white with two black circular areas having the same radius ($r_a = r_b = 35 \text{ cm}$) and centred at 67 cm from the walls. In the following, we refer to the two black areas as area a and b , and the remaining white area as c .

The experiments are carried out in simulation using ARGoS [12], a multi-engine simulator of swarm robotics systems. The robots and the environment are modelled using a 2D dynamic physics engine. We use a simulated model of the e-puck robot (Figure 1 right), a small wheeled robot designed for research and education [13]. In our experimental setup, each robot can perceive walls and other robots through eight infrared proximity sensors placed all around its chassis. It can sense the colour of the floor using three ground sensors placed under its front.




Fig. 1. Left: the simulated experimental arena used for the experiments. Right: the e-puck robot, used for the simulated evolutionary experiments presented in this paper.

Additionally, each robot features another sensor called *range and bearing* [14]. This sensor allows the robot to communicate locally with other robots by sending and receiving messages. In our experiments, the robot uses such a sensor only to perceive the number of other robots within a 70 cm range. To normalise the output of the sensor, we use the preprocessing function $z(n) = 1 - (\frac{2}{1+e^n})$, where n is the number of robots perceived at any given moment. Since the real e-puck can perceive no more than 5 robots at a given time, $z(n)$ saturates to 1 for $n > 5$.

The controller that governs each robot is an artificial neural network. We assume that robots can achieve aggregation using a memoryless behaviour, that is, the behaviour of each robot depends only on the present values of sensors without any kind of internal state. For this reason, we use as controller a fully connected, feed-forward neural network. This neural network has 12 inputs, one for each sensor (8 infrared proximity, 3 ground sensors, 1 from the range and bearing), 2 outputs, one for each wheel, and no hidden units. The input values are linearly scaled in $[0,1]$ when necessary. The activation of the output neurons is computed as the weighted sum of all input units plus a bias term, filtered through a standard logistic function. The two output neurons control the speed of the two wheels, by scaling their activation in the range $[-v_m, v_m]$, with $v_m = 16$ cm/s.

We use a simple evolutionary algorithm to set the parameters of the neural network. Each parameter is represented in the genotype by a real number in the range $[-5,5]$. The evolutionary algorithm works on a population of 100 genotypes, evolved for 200 generations. The population of the first generation is randomly generated. Subsequent generations are created using a selection and reproduction process that involves elitism and mutation. The 20 best genotypes—i.e., the elite—are included unchanged in the next generation. The remaining genotypes of the population are generated by mutation of the genotypes of the elite. The mutation is done by adding a random value to each element of the genotype. The random value is drawn from a normal distribution with mean 0 and variance 1.

The genotype is mapped into a controller that is instantiated in all the robots of the group ($N = 10$). To evaluate the performance, 10 trials of $T = 250$ seconds are run. The evaluation of the performance of the genotype is based on the function $f(t)$:

$$f(t) = \frac{|x_a(t) - x_b(t)|}{N} \in [0, 1] \quad (1)$$

where $x_i(t)$ is the number of robots in area $i \in \{a, b\}$ at time t and N is the total number of robots. The function $f(t)$ is equal to zero when a and b contain the same number of robots. On the contrary, $f(t)$ is equal to 1 when all the robots aggregate on the same area. Fluctuations of $f(t)$ are smoothed through an exponential moving average with time constant $\alpha = 0.9$:

$$G(t) = \alpha G(t - 1) + (1 - \alpha)f(t) \in [0, 1] \quad (2)$$

where $G(0) = 0$. Finally, the fitness F of the genotype is the average of $G(T)$ over all the 10 trials.

3 Results

We performed 20 evolutionary runs starting from different randomly generated populations. For each run, we selected the best controller within the final population: we evaluated the performance of every controller of the last generation for $K = 200$ trials, and we selected the one with the highest average fitness. All the evolutionary runs were able to produce controllers with high performance (data available as supplementary material in [15]).

A qualitative analysis of the obtained controllers reveals that the evolved behaviours are quite similar one to the other. In general, the robots act differently according to their position in the arena. When a robot is in the white area c , it explores the arena following a wide curved trajectory. If the robot reaches the external wall of the arena, it starts to follow it. The robot motion is influenced by the presence of other robots: curves become sharper when other robots are nearby. Such a perturbation makes the robot leave the border of the arena and eventually enter in one of the two black areas. When the robot is in one black area it follows a circular trajectory. The radius of the trajectory decreases as the number of robots in the area increases. In this way, if the area is empty the robot follows a wide trajectory and eventually leaves. On the contrary if the area is crowded the robot almost rotates on its axis. If the robot goes out of the black area it starts again to explore the arena. Example videos of the obtained controller are available as supplementary material [15].

There are qualitative similarities between the evolved behaviour just described and the self-organizing aggregation behaviour observed in groups of cockroaches. In particular, we observed that the probability that a robot leaves an area is inversely proportional to the number of the robots located in the area itself. To determine whether or not the evolved behaviour presents dynamics quantitatively similar to the biological system, we check the adherence of the evolved robotic behaviour¹ with the model introduced in [7]. In Section 3.1, we introduce the model and the methodology we used to estimate its parameters. In Section 3.2, we compare the dynamics of the evolved behaviour with the predictions of the mathematical model.

¹ To this aim, we select the best obtained controller among all evolutionary runs.

3.1 Model

In Amé et al.'s model [7], the behaviour of each individual insect is characterised by J_i , its probability to join area i , and L_i , its probability to leave area i . Both probabilities depend on x_i , the number of insects located in area i , and on S , the *carrying capacity*, that is, the maximum number of insects that can be hosted in a single area.

The joining probability J_i decreases slightly with the number of insects in area i because of crowding effects. This accounts for the observation that it is less probable to join an area that is already densely populated. Amé et al. define J_i as:

$$J_i = \mu \left(1 - \frac{x_i}{S}\right), \quad i = [a, b]; \quad (3)$$

where μ represents the area quality, that is, the probability that an individual joins the area without social influences, x_i is the number of insects already in area i , and S is the carrying capacity.

Similarly, the leaving probability L_i is inversely proportional to the number of individuals in area i . This accounts for social influences among individuals, which tend to stay close together. L_i is low when the area is densely populated and high when it is sparsely populated. Amé et al. define L_i as:

$$L_i = \frac{\theta}{1 + \rho \left(\frac{x_i}{S}\right)^2}, \quad i = [a, b]; \quad (4)$$

where θ depends on the quality of the area, and ρ is a reference surface ratio related to the area carrying capacity. Using J_i and L_i it is possible to describe the time evolution of the number of individuals in the different areas through a system of differential equations:

$$\frac{dx_i}{dt} = J_i x_c - L_i x_i = \mu x_c \left(1 - \frac{x_i}{S}\right) - \frac{\theta x_i}{1 + \rho \left(\frac{x_i}{S}\right)^2}, \quad i = [a, b] \quad (5)$$

$$N = x_c + x_a + x_b \quad (6)$$

where N is the total number of individuals and x_c is the number of individuals in c , that is, the individuals outside the black areas. This model therefore describes the dynamics of the aggregation behaviour in terms of the number of individuals present in the different areas of the arena (see [7] for details).

To evaluate the correspondence of the evolved behaviour with the model, we estimated the model parameters from the results of simulated experiments.

To estimate the parameters of J_i and L_i , we gathered the empirical probabilities by performing 200 simulated experiments in the same conditions as presented in Section 2 ($N = 10$, $r = 0.35$ cm, $S = 29$). We separately fitted the parameters for J_i and L_i to our data using the non-linear least squares method. The obtained parameters are: $\mu = 0.008$, $\theta = 0.008$ and $\rho = 138.574$. We measured the quality of the fitting by computing the coefficient of determination R^2 .




Fig. 2. The bifurcation diagram of our model for different values of $d = \frac{S}{N}$. The percentage of robots in area a and area c .

While the fitting on L_i is excellent ($R^2 = 0.979, p\text{-val} < 0.001$), the fitting on J_i is not as good ($R^2 = 0.560, p\text{-val} = 0.148$). This is due to the fact that in our robotic system, J_i appears to be non-linear, differently from Amé et al.’s model. Even though the fitting is not good, we decided to be consistent with Amé et al.’s model and not change J_i . A discussion of the possible effects of this decision is presented in Section 4.

Following the analysis presented in [7], we studied the system behaviour described by eq. (5) and (6) for different values of $d = \frac{S}{N}$. In Fig. 2, it is possible to see the bifurcation diagram of the model. Four different situations can be observed: (i) For d lower than 0.5, the areas are too small to host all the robots; the robots fill completely the areas and some remain in c . (ii) For $0.5 \leq d < 1$, a single area is too small to host all the robots, so the areas are filled equally. However, in this second case, since there is enough space on the areas for all the robots, only few robots are on c . (iii) For $1 < d \leq 4.2$ the areas are big enough for aggregation to happen. Two stable solutions are found, corresponding to area a or area b hosting the majority of the robots. Additionally an unstable solution is found, corresponding to both areas filled equally. (iv) For d greater than 4.2 the areas are too big and the robots are less likely to perceive the presence of other robots in the same area. Thus, the stable solution corresponds to both areas filled equally.

The number of robots present in c described by eq. (6) also varies with d . Two different situations can be observed: For d lower than 0.5, x_c/N decreases sharply: as areas a and b get bigger, more space is available and the areas can

host more and more robots; For d greater than 0.5, the population fraction on c increases steadily. This is due to the fact that, as S becomes bigger, the probabilities J_i and L_i increase, resulting in a system less likely to converge on a state in which all robots are in areas a or b .

We consider that a collective decision has occurred when $x_a/N > 0.8$. In the bifurcation diagram in Fig. 2 this happens only for d between 1 and 2.8. For d between 2.8 and 4.2 the model predicts a more variable condition with a still unbalanced distribution of robots among the two areas, and an increasing number of robots that move from one area to the other. In the following, we verify these model predictions with respect to the experimental data, presenting a comparison between the results obtained in simulation and those obtained with the model.

3.2 Dynamics of Robotics and Model Simulations

We compared the results obtained from simulated robotics experiments and Monte Carlo experiments for different values of $d = S/N$. We carried out two different analyses. In the first one, the different values of d are obtained by keeping the number of robots fixed to $N = 10$ and varying the carrying capacity S , by changing r_i . In the second, we keep $r_i = 35$ cm (which corresponds to $S = 29$) and we vary the number of robots. For each value of d , we run 1000 trials of $T = 500$ seconds, both for the robotic and the Monte Carlo simulations. For each trial, we collected the final group distribution x_i over the different areas.

Figures 3 and 4 show the obtained results. For each value of d , one bar for each area of the arena is reported. The colours in the stacked bars show the frequency of individual distributions. We divided the distributions in five classes (0-20%, 20-40%, 40-60%, 60-80%, 80-100%), giving each class a different colour. The size of each class in the figures is proportional to its frequency. If the robots are able to perform a collegial decision and aggregate in one single area most frequently, the bars of the areas a and b are mostly dark blue, corresponding to a bimodal distribution with peaks in 0-20% and 80-100%. On the contrary, if the group splits by aggregating in both areas, the bars of the areas a and b are mostly white, corresponding to a unimodal distribution centred in 40-60%. Area c is depicted in dark red when empty (0-20%) and white when full (80-100%).

Figure 3 shows the comparison between robotics and model simulations when the number of robots is fixed to $N = 10$. Apart from low values of S , there is a good correspondence between the model and the evolved behaviour. Moreover, the evolved behaviour looks more stable for $d > 2.9$, indicating that the robots have a better tendency to perform collegial decision than predicted by the model. For $S = \{5, 11\}$ —corresponding to $r_i = \{15, 20\}$ cm—the robots find the areas with difficulty due to the small radius and aggregates are less stable.

Figure 4 shows the results when the carrying capacity S is fixed to 29. The evolved behaviour presents a smoother transition from equally occupying the areas at low d to collegial decisions at high d . For $8 < N < 12$ there is a good correspondence, as the robotic system is close to the evolutionary conditions. Differently, for $N \geq 13$ robots split more frequently than aggregating, while the

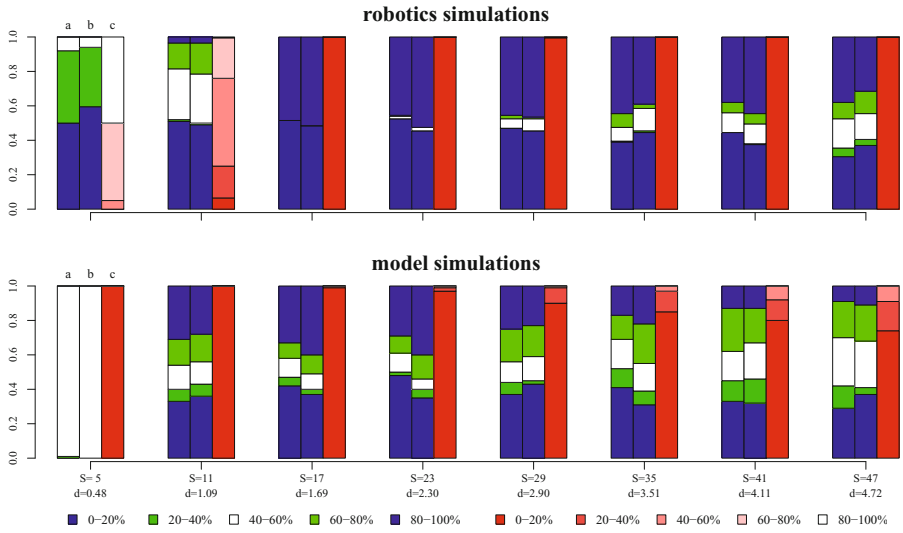


Fig. 3. Comparison between the behaviour of the simulated experiments and the Monte Carlo experiments keeping the number of robots fixed to $N = 10$, and varying the size of the black areas from $r_i = 15$ cm to $r_i = 50$ cm




Fig. 4. Comparison between the behaviour of the simulated experiments and the Monte Carlo experiments keeping a constant radius $r_i = 35$ cm ($S = 29$) and varying the number of robots from $N = 5$ to $N = 40$

model predicts splitting only when one area is saturated. Overall, we observe a good qualitative correspondence between robotics simulations and the model, but mostly within the range of parameters used to evolve the robotics behaviour. A more detailed discussion about these discrepancies follows in the next section.

4 Conclusions

In this paper, we demonstrated that evolutionary robotics techniques can be used to synthesise a collegial decision making behaviour similar to the one observed in cockroaches. This is an important result, especially considering that the robotic controllers are simple feed-forward neural networks without internal states. That is, also in a robotic system collegial decisions can emerge solely from simple individual behaviours modulated by social interactions.

We compared the dynamics of the evolved robotic behaviour with the predictions of the model proposed in [7], finding some qualitative correspondence. However, quantitative comparisons revealed similar dynamics mostly for a small parameter range around the evolutionary conditions ($N = 10$, $S = 29$). We identify two reasons for these discrepancies: (i) the evolved system exploits geometric regularities, such as the arena dimension and the positioning of the areas; (ii) the sensing radius for the robots (75 cm) is quite large with respect to the arena dimensions. Both these issues have a bearing on the probability of joining an area, which also explains the non perfect fit of the model parameters observed in Section 3.1. In practice, we observe that the evolved behaviour depends on both S and N , and not only on their ratio d , as predicted by the model. In future work, by removing geometric regularities from the evolutionary setup, we hope to obtain a better quantitative matching with the model predictions. If successful, we plan to exploit this artificial experimental setup to investigate the optimality of the evolved behaviour with respect to different selective pressures, genetic relatedness among individuals in the group, and variable ecological conditions. We believe this can be useful to better understand the evolutionary path leading to collegial decision making.

Acknowledgements. The research leading to the results presented in this paper has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n 246939. M. Brambilla, M. Birattari and M. Dorigo acknowledge support from the Belgian F.R.S.-FNRS.

References

1. Nolfi, S., Floreano, D.: Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines. MIT Press, Cambridge (2000)
2. Adami, C.: Digital genetics: unravelling the genetic basis of evolution. *Nature Reviews Genetics* 7(2), 109–118 (2006)
3. Waibel, M., Keller, L., Floreano, D.: Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation* 13(3), 648–660 (2009)

4. Mitri, S., Floreano, D., Keller, L.: The evolution of information suppression in communicating robots with conflicting interests. *Proceedings of the National Academy of Sciences* 106(37), 15786–15790 (2009)
5. Wischmann, S., Floreano, D., Keller, L.: Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots. *Proceedings of the National Academy of Sciences* 109(3), 864–868 (2012)
6. Trianni, V., Nolfi, S.: Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial Life* 17(3), 183–202 (2011)
7. Amé, J., Halloy, J., Rivault, C., Detrain, C., Deneubourg, J.-L.: Collegial decision making based on social amplification leads to optimal group formation. *Proceedings of the National Academy of Sciences* 103(15), 5835–5840 (2006)
8. Garnier, S., Gautrais, J., Asadpour, M., Jost, C., Theraulaz, G.: Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior* 17(2), 109–133 (2009)
9. Campo, A., Garnier, S., Dédriche, O., Zekkri, M., Dorigo, M.: Self-organized discrimination of resources. *PLoS ONE* 6(5), e19888 (2011)
10. Brambilla, M., Pinciroli, C., Birattari, M., Dorigo, M.: Property-driven design for swarm robotics. In: *Proceedings of 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Richland, SC, pp. 139–146. IFAA-MAS (2012)
11. Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., Said, I., Durier, V., Canonge, S., Amé, J.M., Detrain, C., Correll, N., Martinoli, A., Mondada, F., Siegart, R., Deneubourg, J.-L.: Social integration of robots into groups of cockroaches to control self-organized choices. *Science* 318(5853), 1155–1158 (2007)
12. Pinciroli, C., Trianni, V., O’Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Caro, G.D., Ducatelle, F., Stirling, T., Gutiérrez, A., Gambardella, L.M., Dorigo, M.: ARGoS: a modular, multi-engine simulator for heterogeneous swarm robotics. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*, pp. 5027–5034. IEEE Computer Society Press, Los Alamitos (2011)
13. Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J., Floreano, D., Martinoli, A.: The e-puck, a robot designed for education in engineering, pp. 59–65. *IPCB: Instituto Politécnico de Castelo Branco* (2009)
14. Gutiérrez, Á., Campo, A., Dorigo, M., Donate, J., Monasterio-Huelin, F., Magdalena, L.: Open E-puck range & bearing miniaturized board for local communication in swarm robotics. In: *ICRA 2009: Proceedings of the 2009 IEEE International Conference on Robotics and Automation*, pp. 3111–3116. IEEE Press, Piscataway (2009)
15. Francesca, G., Brambilla, M., Trianni, V., Dorigo, M., Birattari, M.: Analysing an evolved robotic behaviour using a biological model of collegial decision making: Complete data (2012), Supplementary information page at <http://iridia.ulb.ac.be/supp/IridiaSupp2012-009/>