

To grip, or not to grip: Evolving coordination in autonomous robots

Christos Ampatzis[†], Francisco C. Santos^{††}, Vito Trianni^{†††} and Elio Tuci^{†††}

[†] Advanced Concepts Team, European Space Agency, The Netherlands
{Christos.Ampatzis}@esa.int

^{††} MLG & CoDE-IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
^{†††} ISTC-CNR, Roma, Italy

Abstract. In evolutionary robotics, as in the animal world, performing a task which is beneficial to the entire group demands the coordination of different individuals. Whenever time-dependent dynamic allocation of roles is needed and individual roles are not pre-defined, coordination can often be hard to achieve. In this paper, we study the evolution of role allocation and self-assembling strategies in a group of two homogeneous robots. We show how robot coordination and individual choices (who will grip whom) can be successfully restated in terms of anti-coordination problems, showing how conventional game theoretical tools can be used in the interpretation and design of evolutionary outcomes in collective robotics. Moreover, we highlight and discuss striking similarities between the way our physical robots allocate roles and the way animals solve conflicts. Arguably, these similarities suggest that evolutionary robotics may offer apart from automatic controller design for autonomous robots a viable alternative for the study of biological phenomena.

Key words: anti-coordination game, evolutionary robotics, collective behavior, evolutionary game theory

1 Introduction

Robotics has largely drawn inspiration in the past decades from biology. Bio-inspired robotics refers to mimicking natural mechanisms at the hardware or the collective behaviour level. For example, social insects have often served as a source of inspiration for research on self-organized cooperative exploration in groups of robots using swarm intelligence techniques. Recently, the influence arrow that links biology to robotics has become bi-directional, as roboticists are implicitly or explicitly trying to answer questions related to biology, and in particular animal behaviour. This is because “*robots can be used as models of specific animal systems to test hypotheses regarding the control of behaviour*” [27].

Our view is that Evolutionary Robotics models (ER [17]) can be particularly suitable for testing hypotheses concerning both the **nature** and the **evolution** of the underlying mechanisms that underpin the agents’ behaviour. Such models can be complementary to other analytical modelling tools at the disposal

of biologists, such as (Evolutionary) Game Theory models (EGT, see [13], for example). While ER searches for the mechanisms to solve a problem given a fitness function, EGT defines high-level descriptions of any evolutionary process, useful for objectively identifying the conditions under which certain strategies can emerge as an evolutionary outcome. Recently, for example, ER models have contributed to the research on the evolution of communication and cooperation by proposing a low-level description and mechanisms at the neuronal level to realise signaling behaviour [12,9,2].

As a basis for our discussion we use the evolution of role allocation and self-assembling strategies in a group of two homogeneous robots (see [1,26]). The way in which we address the problem of having the agents coordinate to assume different roles (who will grip whom) is similar to anti-coordination problems. Coordination and anti-coordination problems are studied by biologists, either by direct observation of the behaviour of animals, or with the use of analytical modelling tools, as in [14], where “limited-war” type conflicts between conspecifics are studied. Game theory models allow biologists to predict the outcome of coordination/anti-coordination problems given the set of behavioural strategies available to the agent and the payoff corresponding to all the possible combinations of actions among the actors [13,10,22]. However, such analytical tools may be less suitable for testing hypotheses concerning the nature and the evolution of the underlying mechanisms that underpin the agents’ behaviour. For example, as noticed by [3], we do not know the exact mechanisms (e.g., rules, signals and cues) involved in the formation of assembled structures in natural organisms. Should we assume that agents possess means for explicitly communicating each other’s “intentions” in order to coordinate their actions? In order to address such questions, we believe that ER models can be suitable modelling tools, complementary to other analytical tools at the disposal of biologists.

Potential synergies between ER and EGT models are not limited there. We believe that the use of population dynamics and EGT provides a clear way of understanding (and designing) online adaptation and role allocation in populations of robots—central challenges in most collective robotics problems. This can provide a solid unified framework for the study of complex self-organized behavior in populations of robots.

2 Description of the task and the evolutionary process

The task considered is described in detail in [1,26], and can be roughly summarised as follows:

- Two homogeneous robots are initialised at a random distance between a lower and an upper bound with random initial orientations;
- The robots do not have means for explicit communication;
- The agents should coordinate their movement in order to allocate roles and connect with each other via the gripping mechanism.

The controllers are evolved in a simulation environment which models some of the hardware characteristics of the *s-bot* (a small mobile autonomous robot

with self-assembling capabilities [15]). Each simulated *s-bot* is provided with an omni-directional camera mounted on its turret (see figure 1), returning the distance to the other *s-bot* in each of eight 45° sectors, up to a distance of 50 cm. The actuators controlled by the neural controller are the two wheels and the gripper actuator (open/close gripper). Artificial evolution was used to train networks which are cloned on two homogeneous robots. The network used is a Continuous Time Recurrent Neural Network (CTRNN [5]), with a feed-forward architecture, and the evolutionary algorithm used to set the parameters of the networks is the (μ, λ) evolutionary strategy $((\mu, \lambda)$ -ES).

The fitness assigned to each genotype after evaluation of the robots behaviour is the average of the fitness achieved in 40 trials with predefined robot initialisations in order to ensure that successful controllers can cope with a large and representative sample of all possible initialisations. Notice that this set comprises both symmetrical and asymmetrical initial conditions, that is, when robots share the same initial perceptions, or not, respectively. In each trial, a group is rewarded by the following evaluation function, which assesses the ability of two robots to approach each other and assemble through the gripper, without dictating the role allocation strategy the robot should use:

$$\mathbf{F} = \mathbf{A} \cdot \mathbf{C} \cdot \mathbf{S} \quad (1)$$

A is the aggregation component used for bootstrapping purposes, computed as follows (with d the robot-robot distance at the end of the trial):

$$A = \begin{cases} \frac{1.0}{1.0 + \text{atan}(\frac{d-16}{16})} & \text{if } d > 16 \text{ cm;} \\ 1.0 & \text{otherwise;} \end{cases} \quad (2)$$

C is the collision component aiming to gradually and smoothly punish collisions (with n the number of robot-robot collisions), computed as follows:

$$C = \begin{cases} 1.0 & \text{if } n = 0; \\ 0.0 & \text{if } n > 20; \\ \frac{1.0}{0.5 + \sqrt{n}} & \text{otherwise;} \end{cases} \quad (3)$$

S is the self-assembly component, computed at the end of a trial as follows (with $K(t)$ a bootstrapping component):

$$S = \begin{cases} 100.0 & \text{if robots are assembled;} \\ 1.0 + K(t) & \text{otherwise;} \end{cases} \quad (4)$$

3 Successful strategies tested on real robots

One of the evolved neurocontrollers was first extensively tested in simulation under varying noise conditions to ensure its robustness and its generalisation

capabilities with respect to various initial conditions. Subsequently, this controller was downloaded to two *s-bots*, and tested against different environmental conditions, with a very high success rate (see [1,26]).¹




Fig. 1. Snapshots from a successful trial. (a) Initial configuration. (b) Starting phase. (c) Role allocation phase. (d) Gripping phase. (e) Success (grip).

For all trials, histories of interactions can be described by transitions between a few phases which exhaustively “portray” the observed phenomena. The robots leave their respective starting positions (see figure 1a) and during the starting phase (see figure 1b) they approach each other. The robots then move from the starting phase to what we call the role allocation phase (RA-phase, see figure 1c), in which each *s-bot* tends to remain on the right side of the other. They slowly move by following a trajectory corresponding to an imaginary circle centred in between the *s-bots*. Moreover, each robot rhythmically changes its heading by turning left and right. The RA-phase ends once a robot (*s-bot-gripper*) stops oscillating and heads towards the other (*s-bot-grippee*), which instead orients itself to facilitate the gripping (gripping phase, see figure 1d). The *s-bot-gripper* approaches the *s-bot-grippee*’s turret and grips it. A successful trial terminates when the robots are connected (see figure 1e).

4 Self-assembly as an anti-coordination game

The results of post-evaluation analyses illustrate that the role allocation is the result of an autonomous negotiation phase between the robots [1,26]. The outcome of any action an agent chooses depends on the action of the other and none of the two agents can predict its final role from its initial perception. According to [11], in coordination and anti-coordination problems, “*two (or more) agents must choose one of several alternative actions*”. The author continues by stressing that “*the outcome of any action an agent might choose depends on the action of the other agents*”, as in every frequency dependent process [18]. In [1,26] it is shown that the system is characterised by two basic operational principles:

- For almost all asymmetrical initial configurations, the experimenter can predict the result of the role allocation;

¹ Videos available at <http://iridia.ulb.ac.be/supp/IridiaSupp2008-002>

- For all symmetrical initial configurations, we cannot predict who will grip whom and the agents assume both roles with 50% probability, as random noise in sensors/actuators is the element that breaks the symmetry.

This kind of scenario can be described as an anti-coordination problem, in which the agents have to estimate the strategy of their opponent. If the agents had the means to directly and precisely access the other’s strategies, the role allocation would be immediate. However, our agents are deprived of any such information and must interact to assess and estimate each other’s strategy.

The robots manage to coordinate their actions, “hovering” around the conditions that lead to assuming the *s-bot-gripper* role (oscillatory movement). The way the robots solve the anti-coordination problem bears striking similarities to how animals solve conflicts. In fact, the conflict between two strategies as attack and flee (or surrender), approach or avoid, is very common in nature. For example, gulls during fights adopt in turns agonistic (aggressive) postures which are abandoned as the birds turn broadside to the antagonist. Eventually, one bird will abandon the offensive and will adopt an appeasement posture or run away [25]. The similarity of this behaviour with the one of our robots is striking: it has even been observed that fighting birds walk parallel or around each other, as our robots circle around each other (see figure 1c). Similar coordination rituals are observed in gulls mating, in the “dance-fighting” observed in the male starling [7], in the fighting behaviour but also the mating “zig-zag dance” of the stickleback [24], and in the parallel walks engaged in by red deer stags [6], which allow for each animal to assess the other’s size and strength and to investigate possible asymmetries [16]. Parallels can also be drawn between our system and simultaneous hermaphrodites, as snails, slugs and fish species, where individuals take single mating decisions (“one-shot” games) that require anti-coordination, since assuming the same role will end up costly for both [4].

Obviously, the way in which our robots solve the self-assembly task is determined by the way in which we set up our evolutionary process. By isolating the particular part of the fitness function related with the gripping decision, one immediately obtains a payoff matrix of a simple anti-coordination game, i.e., a simple version of the “Hawk-Dove” game² [13,14]. The fitness function selectively rewards the robot group to achieve self-assembly: the robots must coordinate to decide who will grip whom. A failure to take a decision will result in low fitness scores. The same goes if both robots decide to assume the *s-bot-gripper* role; the robots will collide and will be therefore punished by the fitness function. If we make a simplification and we assume that i) correct role allocation yields the maximum fitness (100), ii) a failure in the decision-making (with both robots playing *s-bot-gripper*) leads to a fitness score of 0, and iii) in case of failure to allocate the roles during the length of a trial, robots receive just the aggregation fitness component ($A \ll 1$), the payoff matrix would be the one in table 1.

Let’s assume that each controller defines an abstract behaviour p , which in turn defines the probability of choosing to grip or not to grip. From table 1, it

² This game is based on the principle that the outcome where neither player yields to the other is the worst possible one for both players.

Table 1. The payoff matrix of the game our robots are evolved to play. One robot chooses a strategy from the columns, and the other from the rows. The payoff refers to the fitness score assigned to the group after the end of the trial. A is the aggregation component of the fitness function (see equation 2).

	<i>s-bot-gripper</i>	<i>s-bot-grippee</i>
<i>s-bot-gripper</i>	0	100
<i>s-bot-grippee</i>	100	A

is trivial to see that $p = 1/2$ is the only evolutionary outcome of this system, whenever $A \rightarrow 0$. In other words, this problem has only one solution (the action of every agent is optimal according to what the other agent does): that the agents should do the opposite of what the other is doing and thus allocate roles. More specifically, the circular movement with oscillations can be seen as the sum of two components: assuming the *s-bot-gripper* role and abandoning it. A premature decision on behalf of one robot to assume the *s-bot-gripper* role might lead to a decision-making error and the robots would end up receiving a fitness score of 0. Thus, a robot has to assume this role while the other assumes the *s-bot-grippee* role. This solution is optimal regardless of the selection type (group or individual); even if we were using heterogeneous pairs of robots which were not evaluated collectively but individually, this solution would still be optimal.

Describing the fitness function in game theory terms offers us the opportunity to view the experiment from a more high-level point of view and to realise that our experimental setup is very close to an anti-coordination game. This clarifies the outcome selected by evolution and shows that the solution found is not just a random one in a possible universe; instead, the principles characterising this solution could only be the ones they are. Also, this way of looking at our fitness function after understanding the basic mechanisms underpinning behaviour in our two robot system helps by providing a more high-level description of our system, that sometimes may be obscured by the complexity involved in the effort to break down its behaviour to a set of transparent rules or states.

Furthermore, even if it provides an excellent starting point, it is important to note that the complexity offered by the above one-shot analysis fails to embrace the complexity of the online decision process. The reason is twofold: first, robots do not play mixed strategies, but, instead, reactions to environmental inputs are preassigned by evolution; second, the online decision process and continuous integration of inputs from the environment, implies not a one-shot game, but a repeated anti-coordination game, in which more complex strategies can emerge. This is particularly relevant in the analysis of the case of symmetric initial configurations. In these cases, robots negotiate on-the-fly their strategies, as they have the possibility to wait in order to assess the other’s decision and, at the same time, reshape their own strategy. In order to ease the differentiation and facilitate the convergence to complementary roles, evolution produced a new strategy characterized by the “dancing” movement (RA-phase, see figure 1c) which amplifies the effect of noise as a symmetry breaking factor. As

expected, whenever the same decision process is studied in the absence of noise (in simulations, not shown) the RA-phase continues indefinitely.

5 Conclusions

We have introduced the idea of the feasibility of ER models as models for biological behaviour, complementary to game theory models. Using the game theoretical framework to design the rules of the game, the properties of the interaction space between agents and the costs and benefits they may share, we can practically build a fitness function within the ER framework incorporating these rules. This can be particularly beneficial when identifying the relative importance that should be given to different fitness components in order to obtain a systematic evolution of the desired behaviour. Subsequently, artificial evolution can propose bias-free low-level behavioural mechanisms instead of high-level strategies to achieve the solutions to the “game” specified.

In this manuscript, only pairwise interactions were considered. However, both engineering and biological applications are often characterised by more than two individuals engaging in coordination and co-existence dilemmas [19]. In this case, allocation of roles and online social dynamics becomes even more intricate. Factors such as selection pressure, interaction structure, population size and average number partners [20,19,8] are known to play decisive roles, making a game theoretical analysis necessary to understand and design efficient populations of controllers. Similarly, to achieve coordination of actions, information needs to be transmitted and properly interpreted. Recent advances in dynamics of signaling and information processing [21,23] are able to capture the most essential aspects by means of simplified models based in game theory. Hence, also here, a high level population-based perspective for ER may provide a unified framework for the study of the evolution of communication, filtering, integration and resultant action. This is an open field of research that can complement the already present momentum that ER research on the evolution of communication has generated [12,9,2], offering a novel way of doing robotics, inspired by biology and population dynamics and less biased by the experimenter.

References

1. C. Ampatzis, E. Tuci, V. Trianni, A. L. Christensen, and M. Dorigo. Evolving self-assembly in autonomous homogeneous robots: experiments with two physical robots. *Artificial Life*, 15(4):1–20, 2009.
2. C. Ampatzis, E. Tuci, V. Trianni, and M. Dorigo. Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26, 2008.
3. C. Anderson, G. Theraulaz, and J.-L. Deneubourg. Self-assemblages in insect societies. *Insectes Sociaux*, 49(2):99–110, 2002.
4. N. Anthes, A. Putz, and N. Michiels. Gender conflicts, sex role preferences and sperm trading in hermaphrodites: a new framework. *Animal Behaviour*, 72:1–12, 2006.

5. R. Beer and J. Gallagher. Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.
6. T. Clutton-Brock, S. D Albon, R. M. Gibson, and F. E. Guinness. The logical stag: adaptive aspects of fighting in red deer (*cervus elaphus l.*). *Animal Behaviour*, 27:211–225, 1979.
7. C-JR. Ellis. Agonistic behaviour in the male starling. *The Wilson Bulletin*, 78(2):208–224, 1966.
8. J. M. Pacheco F. C. Santos and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *PNAS*, 103:3490–3494, 2006.
9. D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519, 2007.
10. J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.
11. D. Lewis. *Convention: A philosophical study*. Harvard University Press, Cambridge, MA, 1969.
12. H. Lipson. Evolutionary robotics: Emergence of communication. *Current Biology*, 17:R330–R332, 2007.
13. J. Maynard-Smith. *Evolution and the theory of games*. Cambridge University Press, Cambridge, UK, 1982.
14. J. Maynard-Smith and G. Price. The logic of animal conflict. *Nature*, 246(5427):15–18, November 1973.
15. F. Mondada, G.C. Pettinaro, A. Guignard, I.V. 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.
16. J. Noble. Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests. In *Proceedings of SAB2000*, pages 481–490. MIT Press, Cambridge, MA, 2000.
17. S. Nolfi and D. Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA, 2000.
18. M. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, 2006.
19. J. M. Pacheco, F. C. Santos, M. Souza, and B. Skyrms. Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proceedings Royal Society B*, 276(1655):315–321, 2009.
20. F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes cooperation in public goods games. *Nature*, 454:213–216, 2008.
21. B. Skyrms. *Evolution of the social contract*. Cambridge University Press, 1996.
22. B. Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, 2003.
23. B. Skyrms. Evolution of signalling systems with multiple senders and receivers. *Phil. Trans. R. Soc. B*, 364(1518):771–779, 2009.
24. N. Tinbergen. The curious behavior of the stickleback. *Scientific American*, 187:22–26, 1952.
25. N. Tinbergen. *The herring gull's world*. Collins, London, UK, 1953.
26. E. Tuci, C. Ampatzis, A. L. Christensen, V. Trianni, and M. Dorigo. Self-assembly in physical autonomous robots: the evolutionary robotics approach. In *Proceedings of ALIFE XI*, pages 616–623. MIT press, 2008.
27. B. Webb. What does robotics offer animal behaviour? *Animal Behaviour*, 60:545–558, 2000.