

On the evolution of self-organised role-allocation and role-switching behaviour in swarm robotics: a case study

Elio Tuci¹, Boris Mitavskiy² and Gianpiero Francesca³

^{1,2}Aberystwyth University, Llandinam Building, Aberystwyth, SY23 1RE, UK

³IRIDIA-ULB, Av. Franklin Roosevelt 50, CP 195/6, 1050 Brussels, Belgium
elt7@aber.ac.uk¹, bom4@aber.ac.uk², gfrancesca@iridia.ulb.ac.be³

Abstract

In spite of its significance for the adaptability of autonomous robotic swarms, the dynamic allocation and re-distribution of robots to tasks (i.e., role-allocation and role-switching behaviour) is still a design challenge in swarm robotics. This study investigates a simulated foraging scenario in which the variability of the environmental conditions requires that robots switch between two roles (i.e., foraging and nest-patrolling). To the best of our knowledge, this is the first simulation study that demonstrates that role-allocation and role-switching behaviour can be evolved using dynamic neural network controllers for robots with minimal communication capabilities. Initial analyses of the best evolved teams shed light on some of the characteristics and robustness of the strategies used by these teams to repeatedly face this task.

Introduction

Swarm robotics is a particular approach to the design of multi-robot systems that finds its theoretical roots in recent studies in animal societies, such as ants and bees (Dorigo and Şahin, 2004). Despite noise in the environment, errors in processing information and performing tasks, and no global information, social insects are quite successful at performing group-level tasks (Anderson et al., 2001; Camazine et al., 2001). Based on the social insect metaphor, swarm robotics emphasises aspects such as decentralisation of the control, limited communication abilities among robots, use of local information, emergence of global behaviour and robustness. These properties are meant to facilitate the design of artificial systems scalable to group size, robust to noise, and adaptive to environmental changes.

Research in swarm robotics has been focusing on mechanisms to enhance the efficiency of the group through some form of cooperation among the individual agents. Complex forms of group cooperative responses can require task-partitioning (i.e., division of a collective task into individual sub-tasks) and/or task/role-allocation (i.e., allocation of sub-tasks/roles to different individuals, see Labella et al., 2006). The latter can be a dynamic and flexible process, in that the number of individuals engaged in any given task may need to continually change, as circumstances require. Hereafter,

we use the term role-switching or task-switching behaviour to refer to the process in which one or more agents leave their current activity to join a different one for the benefit of the team. In spite of its significance for the adaptability of the swarm, the autonomous and dynamic re-distribution of robots to tasks is still a design challenge. This study aims to investigate, in a simulated scenario, the conditions for the emergence of dynamic role-allocation and role-switching behaviour in teams of autonomous agents.

We face this challenge using the Evolutionary Robotics (ER) design method. ER is based on a bottom-up *modus operandi*, where variations are introduced at the genetic level, and selection is performed on the basis of the effects that genetic variations have on the global behaviour of the swarm. With respect to other design methods, ER does not require the designer to make strong assumptions concerning what behavioural and communication mechanisms are needed by the robots. Individual behavioural strategies and rules of actions are determined by the evolutionary process that favours (through selection) those solutions which improve an agent's and the group's ability to accomplish the collective task. The operational mechanisms of the best evolved teams can be *a posteriori* analysed to gain insight into the solutions of the collective problem.

Our long term goal is to apply ER to learning, i) the evolutionary dynamics underpinning the emergence of the role-allocation and role-switching behaviour in autonomous robots; ii) the nature of the individual mechanisms underlying the group response. This study is the first step in this direction. We investigate a scenario in which teams of homogeneous robots are required to split in foragers and nest-patrollers (i.e., robots that remain in the nest). The task demands that both roles are played, but according to different rules. Under one condition, foragers have to be more numerous than patrollers, while under the other the situation is the opposite. The results show that, such a relatively complex team behaviour can be obtained using only very limited means of interaction (i.e., infra-red sensors). Initial analyses of the best evolved teams shed light on some of the characteristics and robustness of the behavioural strategies used by

these teams to repeatedly face this task.

Background and motivations

Several studies in swarm robotics are focused on issues related to the conditions that facilitate the emergence of behavioural specialisation among the robots. Many of these studies provide a comparative cost/benefit analysis of engineering specialisation by using heterogeneous teams (i.e., teams of robots with different individual controllers) versus dynamic and emergent forms of specialisation in homogeneous teams (i.e., team of robots that share the same individual controller, see Eiben et al., 2007; Luke and L.Spector, 1996; Bongard, 2000; Ijspeert et al., 2001; Quinn, 2001; Tuci and Trianni, 2012). From an evolutionary design perspective, homogeneous groups tend to imply a smaller search space than heterogeneous groups, and contrary to the latter they are not affected by the credit-assignment problem (i.e., the problem of divvying up among the team members the reward received through their joint actions, (see Panait and Luke, 2005)). Nevertheless, in those contexts in which partitioning the group task into different sub-tasks it is beneficial for the team, single controllers in homogeneous groups have to underpin behavioural skills required by the robots to undertake all the sub-tasks, as well as the decision making related to the allocation of robots to sub-tasks. The studies described in (Quinn et al., 2003; Ampatzis et al., 2009) showed that, in relatively small groups (2 to 3 robots), task-allocation and specialisation of the team members can be developed by breaking the homogeneity condition through the evolution of ritualised coupled behaviours subject to the effect of random noise inherent in sensory and motor hardware components. However, these studies looked at scenarios in which, once specialised, the agents do not need to reconsider their roles within the group. We know that in natural swarms task-allocation is quite dynamic and flexible, in that the number of individuals engaged in any given task continually changes as circumstances require (Gordon, 1996). Empirical evidence shows that internal factors, such as genetic and morphological differences among the workers, do not always account for the individual variability in task preferences. Single workers in various species of ants perform a variety of tasks, changing from one task to another by tracking contingent factors such as environmental stimuli, the number of agents currently engaged in other tasks, or the rate of interactions with other agents. For example, when a new food source suddenly become available to an harvester ant colony, which competes with other seed-eating species for food, ants previously engaged in other tasks will switch to foraging (Davidson, 1977).

Biologists are particularly interested in models focused on the evolution of emergent principles underpinning task-allocation. As stated in (Duarte et al., 2011), "... disappointingly few attempts have been made to develop realistic scenarios for how the mechanism underlying self-organised

division of labour evolve over the course of generation...". This is partially due to the limitations of classic methodologies at disposal of biologists. The authors in (Duarte et al., 2011) state that evolutionary models of division of labour tend to focus on the conditions in which specialisation is better than generalist strategies, ignoring the mechanism through which specialisation may arise. On the other hand, self-organisation models do not consider the evolutionary trajectories that may lead to task-allocation. Our long-term aim is to create models of the evolution of self-organised role-allocation and role-switching behaviour that can be used as "intuition-pump" for indicating potential evolutionary drivers and emergent behavioural rules capable of accounting for the collective behaviour of natural swarms (see Vassie and Morlino, 2012, for an epistemological account of robotic models). For example, biological evidence shows that behavioural specialisation in various insect societies evolve to minimise the costs of task-switching. Nevertheless, the fact that individuals switch tasks indicates that the evolution of task-allocation systems can not be merely the production of genetically and/or morphologically different individuals, each suited to a particular tasks. Empirically, little is still known about the selective pressures for task-switching behaviour. Robotics models can represent effective alternative methodological tools to investigate these issues. Our aim is to recreate these types of self-organised dynamics in homogeneous robotic swarms by evolving the individual mechanisms and rules of interactions/communication underpinning task-allocation and task-switching behaviour.

The Simulation Environment

In the foraging scenario studied in this paper, the environment is a boundless arena with a nest and a foraging site. The nest is a circular area indicated by a green light, in which the colour of the floor is in shades of grey. The foraging site is also a circular area indicated by a red light, in which the colour of the floor is in a different shades of grey with respect to the nest. The colour of the arena floor is white. The radius of both the nest and the foraging site is randomly defined at the beginning of each trial in the interval [20cm, 30cm]. Both lights, the green one located in the nest and the red one located in the foraging site, are positioned 6cm above the floor and, when turned on, they are visible from everywhere within the arena. In each trial, the green light is placed at the centre of the nest. The red light is randomly placed anywhere within a semicircular area of 10cm radius centred in the centre of the foraging site. The centre of the nest is 1m far from the centre of the foraging site (see Fig. 1a)

The robots kinematics are simulated using a modified version of the "minimal simulation" technique described in (Jakobi, 1997). Our simulation models a e-puck robot, a 3.55cm radius cylindrical robot. It is provided with eight




Figure 1: (a) Experimental scenario (snapshot taken in *Env. A*, during *Phase 1*) showing the nest (dark grey circle), the foraging site (light grey circle), the five robots (white cylinders in the nest), and the lights (small circles above the nest and foraging site). (b) E-puck body-plan. The black circles refer to the position of the infra-red (IR), and floor sensor (F) positioned facing downward on the underside of the robot (see Fig. 1b). The IR sensor values are extrapolated from look-up tables provided with the Evorobot* simulator (see Nolfi and Gigliotta, 2010). The F sensor can be conceived of as an IR sensor capable of detecting the intensity of grey of the floor. It returns 0 if the robot is on white floor, 0.5 if it is on light grey floor, and 1 if it is on dark grey floor. The robots camera has a receptive field of 30° , divided in three equal sectors, each of which has three binary sensors (C_i^B for blue, C_i^G for green, and C_i^R for red, with $i = \{1, 2, 3\}$ indicating the sector). Each sensor returns a value which is 0 if no light is detected, 1 when a light is detected. The camera can detect coloured objects up to a distance of 1.5m. The robots can not see each other through the camera. The robot has left and right motors which can be independently driven forward or reverse, allowing it to turn fully in any direction. The robot maximum speed is 8cm/s.

infra-red sensors (IR^i with $i = \{0, \dots, 7\}$), which give the robot a noisy and non-linear indication of the proximity of an obstacle (in this task, an obstacle can only be another robot); a linear camera to see the lights; and a floor sensor (F) positioned facing downward on the underside of the robot (see Fig. 1b). The IR sensor values are extrapolated from look-up tables provided with the Evorobot* simulator (see Nolfi and Gigliotta, 2010). The F sensor can be conceived of as an IR sensor capable of detecting the intensity of grey of the floor. It returns 0 if the robot is on white floor, 0.5 if it is on light grey floor, and 1 if it is on dark grey floor. The robots camera has a receptive field of 30° , divided in three equal sectors, each of which has three binary sensors (C_i^B for blue, C_i^G for green, and C_i^R for red, with $i = \{1, 2, 3\}$ indicating the sector). Each sensor returns a value which is 0 if no light is detected, 1 when a light is detected. The camera can detect coloured objects up to a distance of 1.5m. The robots can not see each other through the camera. The robot has left and right motors which can be independently driven forward or reverse, allowing it to turn fully in any direction. The robot maximum speed is 8cm/s.

The Task and the Fitness Function

Teams comprising five simulated e-puck robots are evaluated in the context of a dynamic role allocation and role switching behaviour. By taking inspiration from the behaviour of social insects, the roles are nest patrolling and foraging (hereafter, we refer to them as *role P*, and *role F*, respectively). Roughly speaking, *role P* requires a robot to remain within the nest. *Role F* requires a robot to leave the nest for the foraging site, to spend a certain amount of time at the foraging site, and then to come back to the nest. A team is required to execute both roles simultaneously. Therefore, the robots have to go through a role-allocation phase in which they autonomously decide who is doing what, and then ex-

ecute their respective roles.

Moreover, the robots are required to be able to switch from one role to the other (i.e., role switching behaviour) due to the fact that they experience two different types of environment, *Env. A* and *Env. B*. In *Env. A*, *role F* is more important than *role P*. This means that in *Env. A*, a team maximises the fitness if the majority of robots (i.e., more than two robots) visits the foraging site and the minority (i.e., less than three robots) remains in the nest. In *Env. B*, *role P* is more important than *role F*. This means that a team maximises the fitness if the majority of robots (i.e., more than two robots) remains in the nest and the minority (i.e., less than three robots) visits the foraging site. Since a team, throughout its life-span, experiences twice both types of environment, not all the robots can specialise on a single role. The robots have to be able to play both roles and eventually to switch from one role to the other based on the current environmental condition and the roles allocated to the other team mates. How can a robot distinguish between *Env. A* and *Env. B*? The two types of environment can be distinguished by the intensity of grey colouring the floor in the nest site. In *Env. A*, the nest is coloured in dark grey and the foraging site in bright grey. In *Env. B*, the nest is coloured in bright grey and the foraging site in dark grey.

During evolution, each team undergoes a set of $E = 2$ evaluation sequences (hereafter, e-sequence). An e-sequence is made of $V = 4$ trials, in which the teams experience twice each type of environment in the following order: trial 1 *Env. A*, trial 2 *Env. B*, trial 3 *Env. A*, trial 4 *Env. B*. At the beginning of trial 1 of each e-sequence, the robots controllers are reset, and each robot is randomly placed within an area corresponding to a sector of the nest. The nest is divided in 6 sectors, and the robots are placed in sector 1 to 5, as illustrated in Fig. 1c. Each robot is randomly oriented in a way that the light can be within an angular distance of

$\pm 36^\circ$ from its facing direction (see Fig. 1c).

Each trial differs from the others in the initialisation of the random number generator, which influences the robots initial position and orientation, all the randomly defined features of the environment, and the noise added to motors and sensors (see Jakobi, 1997, for further details on sensors and motor noise). Within a trial, the team life-span is $T=900$ simulation cycles (with 1 simulation cycle lasting 0.1s). Robots are frozen (i.e., don't move and do not contribute to the team fitness) if they exceed the arena limits (i.e., a circle of 120cm radius, centred in the middle point between the nest and the foraging site). Trials are terminated earlier if all the robots are frozen, or the team exceeds the maximum number of collisions (i.e., 10). In trials following the first one of each e-sequence (trial 2,3, and 4), the robots are repositioned only if the previous trial has been terminated earlier, or with one or more robot frozen.

Each trial is divided into three phases. During *Phase 1*, which lasts 12s, the green light is on and the red light is off. The robots are required to stay within the nest. During *Phase 2*, which can last from a minimum of 47,5s to a maximum of 52.5s, the red light is on and the green light is off. During *Phase 2*, a team is required to behave according to the rules of the task. That is, in *Env. A*, the majority of robots (i.e., more than two robots) has to visit the foraging site and the a minority (i.e., less than three robots) has to remain for the entire length of this phase in the nest. In *Env. B*, the majority of robots has to remain for the entire length of *Phase 2* in the nest and the minority has to visit the foraging site. A robot is considered having visited the foraging site if, during *Phase 2*, it spends more then 100 consecutive time steps within the foraging site. During *Phase 3*, which starts at the end of *Phase 2* and terminates at the end of the trial, the green light is on again and the red light is off. The robots that were foraging during *Phase 2* are required to return in the nest to rejoin their team mates.

The fitness of a genotype is its average team evaluation score after it has been assessed for two e-sequences (i.e., for a total of 8 trials). In each trial (v) of each e-sequence (e), the team is rewarded by an evaluation function F_{ev} which corresponds to: $F_{ev} = ((PH^1 \times PH^3) + PH^2) \times PEN$. $PH^1 \in [0, 1]$ is computed during *Phase 1*, and it corresponds to the robot average proportion of time steps “inside” the nest. PH^2 is computed during *Phase 2*. $PH^2 = 5$ if the robots follow the rules of the task (see above). If the team does not behave according to the rules of the task, then PH^2 corresponds to i) the proportion of foraging robots multiplied by two, in *Env. A* ii) the proportion of robots remained in the nest multiplied by two in *Env. B*. PH^3 is computed during *Phase 3*. $PH^3 = 2$ if all the robots terminate the phase within the nest. PH^3 corresponds to the proportion of robots that terminated the trail within the nest, otherwise. PH^3 is set to 1 if the trial is terminated before a team reaches *Phase 3*. The team collision penalty




Figure 2: The neural network.

PEN is inversely proportional to the number of collisions, with $PEN = 1$ with no collisions, and $PEN = 0.4$ with 10 collisions in a trial; The average team evaluation score is $F = \frac{1}{E} \frac{1}{V} \sum_{e=1}^E \sum_{v=1}^V F_{ev}$.

Controller and the Evolutionary Algorithm

The robot controller is composed of a continuous time recurrent neural network (CTRNN) of 11 sensor neurons, 4 inter-neurons, and 4 motor neurons (see Beer and Gallagher, 1992). The structure of the network is shown in Fig. 2. The states of the motor neurons are used to control the speed of the left and right wheels as explained later. The values of sensory, internal, and motor neurons are updated using equations 1, 2, and 3.

$$y_i = gI_i; \text{ for } i \in \{1, \dots, 11\}; \quad (1)$$

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{15} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{12, \dots, 15\}; \quad (2)$$

$$y_i = \sum_{j=12}^{15} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{16, \dots, 19\}; \quad (3)$$

with $\sigma(x) = (1 + e^{-x})^{-1}$. In these equations, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, g is a gain factor, I_i with $i = \{1, \dots, 11\}$ is the activation of the i^{th} sensor neuron (see Fig. 2 for the correspondence between robot's sensors and sensor neuron), ω_{ji} the strength of the synaptic connection from neuron j to neuron i , β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate (hereafter, f_i). All sensory neurons share the same bias (β^I), and the same holds for all motor neurons (β^O). τ_i and β_i with $i = \{12, \dots, 15\}$, β^I, β^O , all the network connection weights ω_{ij} , and g are genetically specified networks' parameters. At each time step, the output of the left motor is $M^L = f_{16} - f_{17}$, and the right motor is $M^R = f_{18} - f_{19}$, with $M_L, M_R \in [-1, 1]$. Cell potentials are set to 0 when

the network is initialised or reset, and equation 2 is integrated using the forward Euler method with an integration time step $\Delta T = 0.1$.

A simple evolutionary algorithm using linear ranking is employed to set the parameters of the networks (Goldberg, 1989). The population contains 100 genotypes. Generations following the first one are produced by a combination of selection with elitism, recombination, and mutation. For each new generation, the three highest scoring individuals (“the elite”) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness-proportional selection from the 60 best individuals of the old population. Each genotype is a vector comprising 87 real values (76 connections, 4 decay constants, 6 bias terms, and a gain factor). Initially, a random population of vectors is generated by initialising each component of each genotype to values chosen uniformly random from the range $[0,1]$. New genotypes, except “the elite”, are produced by applying recombination and mutation. Each new genotype has a 0.3 probability of being created by combining the genetic material of two parents. During recombination, one crossover point is selected. Genes from the beginning of the genotype to the crossover point is copied from one parent, the other genes are copied from the second parent. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.04. The mean of the Gaussian is 0, and its standard deviation is 0.1. During evolution, all vector component values are constrained to remain within the range $[0,1]$.

Results

Our objective is to design neuro-controllers for homogeneous teams of robots required to exhibit the following skills: dynamic role-allocation, and role-switching behaviour. This means that teams have to be capable of dynamically allocating roles to robots and simultaneously executing both roles in each trial. Moreover, not all the robots can specialise in a single role (i.e., playing only a single role throughout an e-sequence). This is because there are two different types of environment: *Env. A*, in which the majority of the robots has to play *role F*; and *Env. B*, in which the majority of the robots has to play *role P*. The robots (at least one) have to switch role between consecutive trials for the majority to be distributed as required by the task.

10 evolutionary runs, each using a different random initialisation were carried out for 2500 generations. Seven evolutionary runs managed to generate teams with the highest fitness score (data not shown¹). In order to have a better estimate of the behavioural capabilities of the evolved teams, we post-evaluated, for each run, the fittest team of each generation for the last 500 generations. The post-evaluation test

¹See <http://users.aber.ac.uk/elt7/suppPagn/ECAL2013/suppMat.html> for further methodological details, graphs and movies of the best teams.

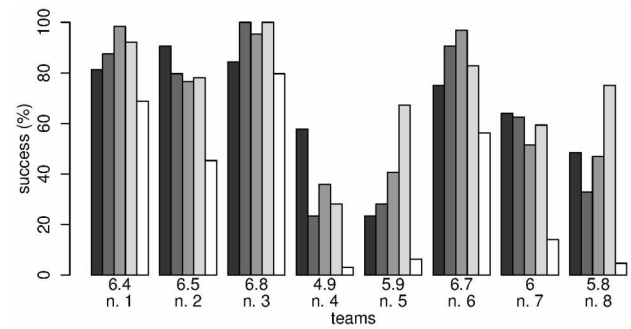


Figure 3: Graph showing the performances of the best teams from 8 different evolutionary runs, re-evaluated for 64 sequences. For each team, the number below the bars refers to the average score F . The four bars in shades of grey (from black to light grey) indicate the % of success in each of the four temporal sections (or trials) of an e-sequence, respectively. The fifth bar (the white one) indicates the % of successful e-sequences out of 64.

consists of 64 e-sequences per team (for a total of 256 trials, 4 trials times 64 e-sequences). The performance of each team is measured using the metrics F illustrated above, with $E = 64$. For each run, the team with the highest average re-evaluation score F is assumed to be an adequate measure of the success of the run.

The graph in Fig. 3 shows for each best team, the fitness (F), the success rate (%) in each temporal section (i.e., trial) of the e-sequence (i.e., black bars for the first trials, dark grey bars for the second trials, medium grey bars for the third trials, and light grey bars for the fourth trials), and the percentage of successful e-sequence out of 64 (see white bars). An e-sequence (e) is considered successful if a team manages to get the highest score ($F_{ev} = 7$) in each of the four trials (v). Note that, the results of the worst two runs have been omitted from Fig. 3 because the percentage of successful e-sequence of the best teams was 0.

The numbers just below the bars in Fig. 3 refers to the average fitness score of each team. Four teams (team n. 1, 2, 3, and 6) managed to get an average fitness score quite close to the optimum $F = 7$, with the team generated by run n. 3 (hereafter, team n. 3) being the most successful. The bars in the graph explain the meaning of these fitness scores in term of teams’ performance and robustness. For example, team n. 1, 2, 3, and 6, have a relatively high success rate in each of the four temporal sections of the e-sequence (see black to light grey bars in Fig. 3, for team n. 1, 2, 3, 6). Nevertheless, these scores do not necessarily correspond to a high percentage of successful e-sequences. Team n. 1, 2, 6 have a percentage of successful e-sequence below 70% (see white bars in Fig. 3, for team 1, 2, and 6). Team n. 3, instead, manages to successfully complete about 80% of the

64 post-evaluation e-sequences, proving to be quite robust and effective. As shown in Table 1, team 3 is also the only team for which unsuccessful e-sequences are all caused by only a single unsuccessful trial.

To summarise, several best evolved teams look quite effective in repeatedly solving individual temporal sections of the task (i.e., single trials), even if none of the best teams proved to be 100% successful. One evolutionary run (n. 3) managed to generate several successful and robust teams, with the best of them capable of executing about 80% of successful e-sequences at the post-evaluation test. The performances of best evolved teams from other evolutionary runs are significantly reduced if evaluated with respect to the percentage of successful e-sequences. This data indicates that, for most of the teams, the evolutionary scores have been an overestimation of the effective behavioural capabilities of the teams. We will further discuss the reasons of the post-evaluation performance drop of some of the best evolved teams in the next section.

Analysis of the best evolved team

In this Section, we show data collected during the post-evaluation test aimed to illustrate some of the features of the best evolved team (i.e., team n. 3). In particular, we show some qualitative data referring to role-switching behaviour. Recall that, an e-sequence is made of 4 trials, during which the robots controller is not reset. In an e-sequence, a team experiences both types of environment in the following order: trial 1 *role F*; trial 2 *role P*; trial 3 *role F*; trial 4 *role P*. Role-switching is a robot's behaviour that happens between two consecutive trials. Within a single trial, a robot can only play either *role P* or *role F*. A role-switching event refers to any change of role from *role P* to *role F* or vice-versa.

How many robots of team n. 3 switch role during an e-sequence? Due to the nature of the task, a team can employ different strategies with respect to role-switching. For example, a strategy with relatively small amount of role-switching behaviour is one in which one robot systemati-

Table 1: Table showing, for each team, num. of unsuccessful e-sequence with failure in: 1 trial (col. 2), 2 trials (col. 3), 3 trials (col. 4), 4 trials (col. 5). Best team in grey.

Team n.	Num. trial failed			
	1	2	3	4
1	14	6	0	0
2	23	11	1	0
3	13	0	0	0
4	8	16	29	9
5	7	20	25	8
6	23	3	2	0
7	22	19	12	2
8	15	29	15	2

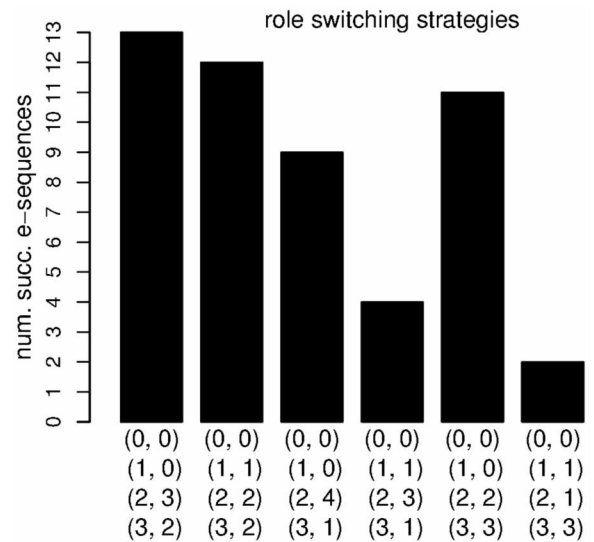


Figure 4: Graphs showing, for team n. 3, the different behavioural strategies (indicated by the four 2-tuples on the x-axis) and their frequencies recorded in the 51 successful e-sequences of the re-evaluation test. On the x-axis, the first element of each 2-tuple refers to the number of role-switching event and the second element refers to the number of robots.

cally switches role at the beginning of a trial (for a total of 3 switches), two robots systematically play *role F*, and two robots systematically play *role P*. A strategy with a significant amount of role-switching behaviour is one in which all the robots change roles in response to any type of environmental change. Fig. 4 shows all the behavioural strategies of team n. 3 observed among all the 51 successful e-sequences recorded at the post-evaluation test. A behavioural strategy is described by 4 2-tuples, in which the first element of each 2-tuple refers to the number of role-switching event and the second element refers to the number of robots. For example, the first bar in Fig. 4 refers to the strategy in which during an e-sequence, 3 robots switch role two times, and 2 robots switch role three times. This is the strategy employed more frequently by team n. 3 during successful post-evaluation e-sequences. However, we can see that this team employs up to 6 different strategies, all of them characterised by the fact that all the robots switch role at least once during a successful e-sequence. The team relies on a variety of strategies which are highly dynamics with respect to role-switching behaviour. Thus, the evolutionary conditions characterised by our experimental scenario are sufficient to generate controllers that are plastic enough to avoid behavioural specialisation (i.e., robots that play only a single role throughout an entire e-sequence).

How much is the role that a robot play determined by the characteristics of the environment? The graph in Fig. 5 shows, for each robot of the best evolved team (n. 3), the




Figure 5: Graph showing, for each robot of team n. 3, the percentage of successful post-evaluation e-sequences, in which the robot plays a different role in different trials presenting the same type of environment. Black bars refers to *Env. A*, white bars to *Env. B*.

percentage of successful e-sequences, at the post-evaluation test, in which the robot plays a different role in trials associated to the same type of environment. This graphs indicates that the type of environment experienced by a robot is not the only cue that determines the role a robot plays within a trial. This is because, for both *Env. A* and *Env. B*, the robots do not necessarily play the same role in different trials associated to the same type of environment. The simple rule “one role for each type of environment” does not systematically apply to any robot of the team n. 3. There is a certain amount of variability among the robots, with robot 1 (i.e., the robot initialised in sector 1, see Fig. 1c) being the most reluctant to play different roles in different trials of an e-sequence presenting the same type of environment (see Fig. 5). Nevertheless, this evidence suggests that the roles are genuinely determined by a combination of factors. In particular, the physical interactions among the robots bring forth contingent phenomena that, biased by the current environmental condition, guide the role-allocation process.

Without further evidence, it is premature to speculate on the operational mechanisms that guides the decision process in this scenario. In similar studies, in which homogeneous teams have been evolved to solve dynamic role-allocation tasks using infra-read sensors as means for communication, the decision making process turned out to heavily rely on the noise injected into the simulation (e.g., Ampatzis et al., 2009). It is likely that, even in our scenario, the robots exploit the noise to break the symmetries due to the homogeneity of the system, and to “diverge” on different roles. However, this scenario differs from those described in the studies above mentioned for two main features: first, the robots are required to repeatedly go through the decision making process, without the robot controllers being reset. Thus, from a functional point of view, the symmetry condition of the robots controller applies only to trial 1. Second,

variable environmental structures and the robots interactions have to equally contribute to the dynamic allocation of roles to robots. This means that, the noise may not be so important in guiding the decision making as it turned out to be in the above mentioned studies.

The relative positions of the robots at the beginning of each trial can potentially be another element that influences the role-allocation process. By visually inspecting the team behaviour, we noticed that most of the best evolved strategies are particularly sensitive to the variability in the initial relative positions of the robots. Too much variability is highly disruptive. Another evidence in favour of the significance of the “initial relative positions” hypothesis for the role allocation process can be found in Fig. 3. The graphs shows that the best three teams (n. 1, 3, 6) did worst in the first temporal section (trial) than in the following three sections at the post-evaluation test (compare black with all the other bars in Fig. 3, for team n. 1, 3, 6). Recall that, in trials following the first one, the robots can position themselves, through their movement, in a way to facilitate the execution of the task. In the first trial of each e-sequence the robots are pseudo-randomly positioned. The robots’ positioning algorithm has been intentionally designed to introduce some variability into the system. However, it seems that, even the best evolved strategy can only cope with a limited portion of this variability. Future work and analysis on the operational mechanisms used by the robots to allocate role may shed light on whether and the extent to which the robots relative position at the beginning of the trial has an impact on the performance of the team, and on the role that each robot play within the team. However, it seems plausible to think that the sensitivity of the best evolved teams to the variability in the robots’ initial relative position can account for the fitness drop observed between evolution and post-evaluation.

Conclusion

In this study, we have investigated the evolutionary conditions that facilitate the emergence of role-allocation and role-switching behaviour in teams of homogeneous robots. To the best of our knowledge, this is the first simulated study in which: i) role-allocation and role-switching behaviour are both required for the benefit of the team. ii) role-allocation and role-switching behaviour are evolved in team of more than 3 agents; iii) a relatively complex team’s behaviour, based on a different distribution of robots-to-roles, is obtained using only infra-red sensors as means for interaction.

We consider this study the first step towards the development of swarm robotics models that could shed light on the evolution of self-organised role-allocation and role-switching behaviour. Similarly to other swarm robotics studies, we are motivated by engineering and biological objectives. From an engineering perspective, our objective is to look at task/role-allocation and task/role-switching behaviour in order to generate design principles that preserve

the adaptability and flexibility of both the system components and of the resulting processes. With respect to this, the results of this study demonstrate that a task requiring complex group response in simulated agents with limited sensory and communication capabilities can be solved by teams of homogeneous robots controlled by dynamic neural network synthesised by artificial evolution.

From a biological perspective, we aim at generating new insights into the general properties of large scale distributed natural systems. In particular, this project aims at providing a principled understanding of the temporal development of task/role-allocation and task/role-switching behaviour by designing models that look at the effect of evolution on operating and design principles. Further work and analysis of the working principles of the evolved solutions is certainly required to be able to understand how our simulated teams solve the task, and how this can help us to understand natural swarms. However, at first glance, we can already mention a couple of phenomena that seem to play a significant role for the evolution and development of self-organised role-allocation and role-switching behaviour. The first element is the way in which the robots access and leave the nest. In our scenario, the robots have been left free to move in and out the nest from any directions. However, the best evolved strategies are based on movements which induce all the foragers to exit and re-enter the nest from a specific limited area. This suggests that structural properties of the nest may interfere and maybe facilitate the adaptive re-distribution of agents-to-roles. Another element is the amount of environmental variability the system can cope with. Our robots, which communicated only through infra-red sensors, bring forth strategies that seem quite fragile with respect to various sources of variability, such as the cardinality of the team, the distance between nest and foraging site, the length of a trial, the order in which the different environmental conditions are experienced, etc. Future work will concentrate on the investigation of alternative means of communication that could strengthen the effectiveness of the self-organised re-distribution of agents-to-roles.

References

- Ampatzis, C., Tuci, E., Trianni, V., Christensen, A., and Dorigo, M. (2009). Evolving self-assembly in autonomous homogeneous robots: Experiments with two physical robots. *Artificial Life*, 15(4):465–484.
- Anderson, C., Franks, N., and McShea, D. (2001). The complexity and hierarchical structure of tasks in insect societies. *Animal Behaviour*, 62:643–651.
- Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122.
- Bongard, J. (2000). The legion system: A novel approach to evolving heterogeneity for collective problem solving. In Poli, R., editor, *Genetic Programming: Proc. of EuroGP-2000*, LNCS, pages 16–28. Springer-Verlag.
- Camazine, S., Deneubourg, J.-L., Franks, N., Sneyd, J., Theraulaz, G., and Bonabeau, E. (2001). *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ.
- Davidson, D. (1977). Foraging ecology and community organization in desert seed-eating ants. *Ecology*, 58(4):725–737.
- Dorigo, M. and Şahin, E. (2004). Guest editorial. Special issue: Swarm robotics. *Aut. Rob.*, 17(2–3):111–113.
- Duarte, A., Weissing, F., Penn, I., and Keller, L. (2011). An evolutionary perspective on self-organised division of labour in social insects. *Annual Review of Ecology, Evolution, and Systematics*, 42:91–110.
- Eiben, A., Nitschke, G., and Schut, M. (2007). Collective specialization for evolutionary design of a multi-robot system. In Sahin, E., Spears, W., and Winfield, A., editors, *Proc. of the 2nd Int. Workshop on Swarm Robotics*, volume 4433, pages 189–205. Springer.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA.
- Gordon, D. (1996). The organisation of work in social insects. *Nature*, 380:121–124.
- Ijspeert, A., Martinoli, A., Billard, A., and Gambardella, L. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Aut. Rob.*, 11(2):149–171.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6:325–368.
- Labella, T., Dorigo, M., and Deneubourg, J.-L. (2006). Division of labour in a group of robots inspired by ants’ foraging behavior. *ACM TAAS*, 1(1):4–25.
- Luke, S. and L.Spector (1996). Evolving teamwork and coordination with genetic programming. In *Genetic Programming 1996. Proc. of the 1st Annual Conf.*, pages 150–156. MIT Press, Cambridge, MA.
- Nolfi, S. and Gigliotta, O. (2010). Evorobot*. In Nolfi, S. and Mirolli, M., editors, *Evolution of Communication and Language in Embodied Agents*, pages 327–332. Berlin-Heidelberg: Springer-Verlag.
- Panait, L. and Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11:387–434.
- Quinn, M. (2001). A comparison of approaches to the evolution of homogeneous multi/robot teams. In *CEC*, volume 1, pages 128–135.
- Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Phil. Trans. of the Royal Society of London, Series A*, 361:2321–2344.
- Tuci, E. and Trianni, V. (2012). On the evolution of homogeneous multi-robot teams: clonal versus aclonal approach. In Ziemke, T., Balkenius, C., and Hallam, J., editors, *SAB XII*, volume 7426 of *LNAI*, pages 391–400. Springer.
- Vassie, K. and Morlino, G. (2012). Natural and artificial systems: Compare, model or engineer? In Ziemke, T., Balkenius, C., and Hallam, J., editors, *SAB XII*, volume 7426 of *LNAI*, pages 1–11. Springer.