

On the Evolution of Homogeneous Two-Robot Teams: Clonal versus Aclonal Approaches

Elio Tuci · Vito Trianni

Received: date / Accepted: date

Abstract This study compares two different evolutionary approaches (clonal and aclonal) to the design of homogeneous two-robot teams (i.e., teams of morphologically identical agents with identical controllers) in a task that requires the agents to specialise to different roles. The two approaches differ mainly in the way teams are formed during evolution. In the clonal approach, a team is formed from a single genotype within one population of genotypes. In the aclonal approach, a team is formed from multiple genotypes within one population of genotypes. In both cases, the goal is the synthesis of individual generalist controllers capable of integrating role execution and role allocation mechanisms for a team of homogeneous robots. Our results diverge from those illustrated in a similar comparative study, which supports the superiority of the aclonal versus the clonal approach. We question this result and its theoretical underpinning and we bring new empirical evidence showing that the clonal outperforms the aclonal approach in generating homogeneous teams required to dynamically specialise for the benefit of the team. The results of our study suggest that task-specific elements influence the evolutionary dynamics more than the genetic relatedness of the team members. We conclude that the appropriateness of the clonal approach for role allocation scenarios is mainly determined by

the specificity of the collective task, including the evaluation function, rather than by the way in which the solutions are evaluated during evolution.

Keywords Evolutionary Robotics · Homogeneous and Heterogeneous Teams · Role-allocation

1 Introduction

Distributed multi-robot teams are robotic systems that mimic some of the properties of natural swarm, such as the capability to develop adaptive responses without central control and with local and simple communication strategies [4]. Several research studies in robotics have been focusing on the issues related to the use of homogeneous versus heterogeneous multi-robot teams in tasks requiring the robots to take different roles [13]. In a homogeneous multi-robot team the robots share equivalent physical structure and identical control system. Each single robot has its own controller, which is an exact copy of those assigned to the other team mates. Specialisation in homogeneous multi-robot teams emerges through a dynamic or self-organising process of task/role allocation. That is, the members of a team autonomously allocate the roles among themselves. In a heterogeneous multi-robot team the team members differ in the hardware structure, in the control system, or in both of them [5]. Specialisation in heterogeneous multi-robot teams is determined by either structural, functional or both types of differences among the team members.

The idea of facing tasks requiring specialisation with the use of homogeneous multi-robot teams has some clear advantages over the alternative of using heterogeneous teams. From the point of view of robustness,

Elio Tuci
Llandinam Building, Dept. Computer Science, Aberystwyth
University, SY23 3RD, Aberystwyth, UK
Tel.: +44-(0)1970-622537
Fax: +44-(0)1970-628536
E-mail: elt7@aber.ac.uk

Vito Trianni
ISTC-CNR, Via San Martino della Battaglia 44, 00185,
Rome, Italy

heterogeneous teams are rather fragile since, due to specialisation, if an element of the team fails, its function can not be easily replaced, and the entire team is likely to fail. Similarly, the team might fail if, due to environmental changes, the collective task demands a different redistribution of agents to activities. Homogeneous teams do not suffer from these limitations since the capability of each individual to take any role makes the performance of the team less susceptible to single robot failure and to changes in the operating conditions [see 22, 1]. However, from the design point of view, it can be a difficult task to obtain complex team level responses through the definition of individual mechanisms and rules of interactions [24].

The general goal of this paper is to contribute to overcome the current limitations in the design of homogeneous multi-robot teams by shedding light on the effectiveness of different design approaches based on the use of artificial evolution. In particular, we revisit a hypothesis originally formulated by Quinn [20], and further elaborated by the author in [21], concerning the use of artificial evolution to the design of homogeneous two-robot teams for tasks requiring specialisation. From the point of view of using artificial evolution as design method, homogeneous teams make the design process less problematic than heterogeneous teams [23]. Owing to the fact that the agents of a team share identical controllers, in homogeneous teams there is no need to divvy up among the team members the reward received through their joint actions [18]. Moreover, the problem search space can be kept relatively small as fewer solutions need to be evaluated (i.e., one solution per group instead of one for each task required).

The work described in [20] indicates that the evolutionary design of controllers for homogeneous teams engaged in tasks requiring specialisation can be more effectively achieved using an alternative approach, which exploits heterogeneous teams to support the evolution of individual generalist controllers, which can be deployed also in homogeneous teams (see Section 3 for a detailed discussion). Starting from a critical analysis of this work, we first identify task-specific characteristics of the experimental scenario that we believe may have significantly contributed to the counterintuitive results shown in [20]. We design a new comparative study that tests our hypotheses. Based on the analysis of the obtained results, we formulate conclusions that revisit and revise the argument put forward in [20] to account for the results of his study.

The paper is structured as follows. In section 2, we review the relevant literature. In section 3, we briefly illustrate Quinn’s work and the hypothesis formulated by the author to account for his results. In section 4, we

discuss Quinn’s hypothesis, we propose an alternative reading to Quinn’s [20] results, and we illustrate the goal of this paper. In section 5, 6, and 7, we describe the methods of this study. In section 8, we illustrate the results. Discussion and conclusions are presented in section 9.

2 Review of relevant works

In this section, we review some of the literature focused on issues related to the use of artificial evolution for the design of multi-robot teams for tasks that require the individual to take different roles [see 13, for a more detailed review of this research area]. The great majority of these studies focuses on the analysis of the relationship between the operating conditions and the composition of the team (i.e., homogeneous versus heterogeneous).

Nitschke et al [14] advocate a particular approach called Collective Neuro-Evolution (CONE) for the evolution of collective behaviour in teams of simulated rovers. When compared to other evolutionary design methods, CONE proved to be the most effective in facilitating the emergence of behavioural specialisation in a cooperative scenario. Luke and Spector [12] show the benefits of another approach for the design of team strategies in a competitive scenario. In Bongard [3], the author illustrates a method employing genetic programming to evaluate whether a given task is more efficiently solved by behaviourally heterogeneous or homogeneous agents. The results of this study indicate that heterogeneous teams should be preferred to homogeneous teams in inherently decomposable tasks (i.e., tasks that can be functionally decomposed into different sub-tasks). In Ijspeert et al [10], homogeneous and heterogeneous robots are compared on a task in which the agents have to cooperate to pull a stick out of the ground. Results show that the performances of heterogeneous and homogeneous teams differ according to the operational circumstances. For example, heterogeneous teams are more collaborative than homogeneous teams when the robots can communicate and when there are fewer robots than sticks.

A recent series of robotic models focused on issues relevant to evolutionary biology have produced results of great interest to roboticists. In particular these studies shed light on the relationship between the genetic composition of the team and the emergence of communication, altruism, and colony efficiency in a multi-task scenario [see 26, 28, 7]. In Waibel et al [27], the authors focus on tasks that require different levels of cooperation among the agents, but no specialisation. Through a systematic investigation of all the possible

conditions arising from the interaction between the genetic composition of the team (i.e., heterogeneous or homogeneous) and the level of selection (i.e., individual or team level), the authors show that cooperative tasks are more efficiently solved by a homogeneous team of robots selected on the basis of team-performance, while non-cooperative tasks are more efficiently solved by behaviourally heterogeneous robots selected individually. The results of this latter study are rather inconsistent with the results of a previous research work illustrated by Potter et al [19], who show that there are aspects of certain cooperative tasks that make them more suitable to heterogeneous than homogeneous teams. In particular, Potter et al [19] show that the higher the number of skill sets required to solve a cooperative tasks, the more beneficial and necessary heterogeneity becomes.

The work described in Quinn [20] is one of the few exclusively dedicated to the comparison of different evolutionary methods for the design of homogeneous teams engaged in cooperative scenarios that require specialisation. The next section briefly describes this work and illustrates the hypothesis formulated by the author to account for the obtained results.

3 Quinn's work

In the work described in [20], two simulated robots equipped with only infra-red sensors, and initially placed close to each other in an empty arena, are required to move in an arbitrary direction by remaining within sensor range. To accomplish their goal, the robots differentiate their roles in robot leader (i.e., the one that, being on the front-end of the moving chain, is supposed to lead the team), and robot follower (i.e., the one that follows the leader). The author compares two different evolutionary approaches for the design of homogeneous teams engaged in this scenario requiring behavioural specialisation. In particular, the comparison is between two different ways of pairing the members of a two-robot team during evolution. In the clonal approach, a team is formed using a single genotype from the evolving population of genotypes. Thus, each genotype generates "cloned" control software for both robots. It follows that clonal teams are homogeneous by definition, because all the members of a team have a controller derived from the same genotype. In the aclonal approach instead, a team is formed from multiple genotypes (one for each team member) from the evolving population of genotypes. Each genotype generates the control software for only one robot. It follows that aclonal teams are heterogeneous because each team member has a controller derived from a different genotype (see Figure 1).

In [20], the author makes use of the aclonal approach to design artificial neural networks that, at the end of the evolution are used to control homogeneous groups. That is, while during evolution, solutions are evaluated in heterogeneous groups, after evolution, the best solutions are tested in homogeneous groups. The results of the study surprisingly indicate that, regardless of the theoretical disadvantages clearly listed and discussed in the paper, the unconventional use of the aclonal approach is a more effective way than the clonal approach to generate controllers for homogeneous systems in which the team members have to autonomously specialise for the benefit of the team. In other words, the results of the study indicate that the most efficient way to evolve homogeneous multi-robot teams for task requiring specialisation is through the aclonal approach; that is, by evolving heterogeneous multi-robot teams.

To account for these results, Quinn formulates and brings evidence in favour of a hypothesis according to which the aclonal approach takes advantages of specific evolutionary dynamics that are precluded to the clonal approach.

Analysis of the behaviour of heterogeneous teams at various stages of aclonal runs revealed that agents initially evolved to perform specialised roles within a team. The allocation of roles was thus initially genetically determined. One consequence of this is that roles could be developed and refined prior to the evolution of any dynamic allocation mechanism. This was obviously not an option in clonal runs where teams were constrained to be homogeneous. For clonal individuals, the adoption of complementary roles necessarily requires the existence of some dynamic role allocation mechanism. However, the evolution of any dynamic allocation mechanism seems unlikely before agents have the ability to perform distinct roles. Presumably then, clonal teams had to evolve behavioural roles simultaneously with the mechanisms for allocating these roles. In contrast, aclonal populations' capacity for specialisation enabled them to discover and refine behavioural roles independently of the discovery and development of dynamic role allocation mechanisms [20, p 133].

In summary, the author of [20] claims that, in the aclonal approach, behavioural roles can be developed and refined in genetically unrelated agents, owing to the emergence of specialisation. That is, agents from the same evolving population possess the mechanisms to play one role or the other but not both. Specialisation precedes and paves the way to the evolution of generalist solutions which emerge when evolution finds

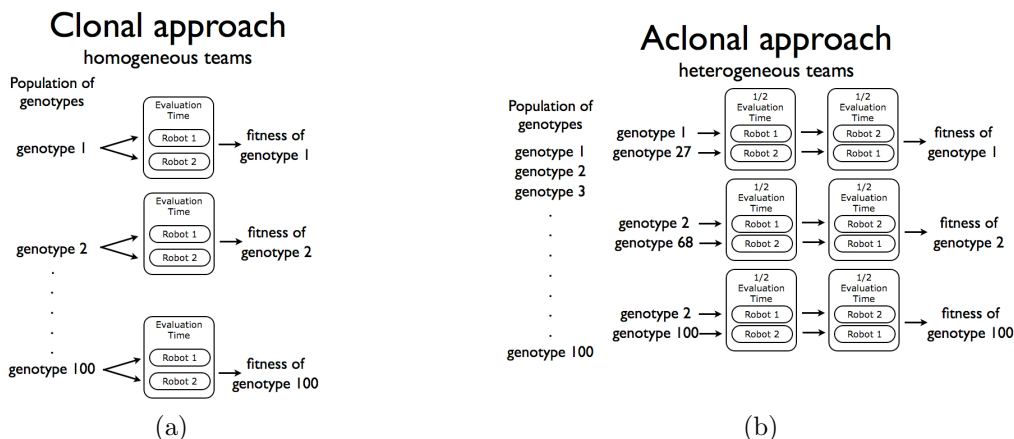


Fig. 1 (a) The clonal approach. A single genotype is used to create two identical controllers, one for robot 1 and one for robot 2. The team is evaluated and the fitness is associated to that genotype. (b) The aclonal approach. A genotype is randomly paired with another genotype. The two genotypes are used to create two different controllers. Each controller is associated to each robot for half of the evaluation time. For example, in the scenario depicted in (b), genotype 1 is randomly paired with genotype 27. The team is evaluated and the team fitness is associated to genotype 1. The same mechanism applies to genotypes 2 and 100.

the mechanisms to allow the agents to dynamically allocate the roles. The study also shows that, in the clonal approach, the gradual evolution from genetically specialised to generalist solutions is not possible, because the agents are clones, and the adoption of complementary roles necessarily requires the existence of some dynamic role allocation mechanisms. Thus, the author concludes that the clonal approach is penalised by the fact that behavioural roles and the mechanisms to allocate them have to (laboriously) evolve simultaneously.

4 Our hypothesis

This study focuses on a two-robot scenario very similar to the one illustrated in [20]. We are moved by the hypothesis that the results shown in [20] concerning the superiority of the aclonal versus the clonal approach are affected by the evaluation function, which according to us limits the potentialities of the clonal approach as a method for the design of generalist controllers for homogeneous multi-robot teams engaged in tasks requiring specialisation. We remind the reader that a generalist solution refers to a single controller capable of underpinning both role execution and role allocation processes in a team of homogeneous robots.

In [20], the author argues that in task requiring specialisation, generalist solutions are found less easily by clonal than aclonal approaches because clonal approaches, by working with homogeneous teams, are limited by the simultaneity argument, which constrains the evolutionary dynamics capable of generating successful teams. I remind the reader that the simultaneity argument refers to the idea that behavioural roles

and the mechanisms to allocate them have to evolve simultaneously. We argue that the simultaneity argument discussed in [20] is induced by the type of evaluation function that Quinn uses to design controllers for homogeneous two-robot teams. Thus, we predict that the clonal approach, if used in combination with a different type of evaluation function, can exploit alternative evolutionary paths in which the mechanisms for role allocation and for executing the roles can evolve at different evolutionary times in spite of the homogeneity of the team members.

Our hypothesis is based on the following reasoning. The evaluation function used in [20] is primarily based on a group metrics. Hereafter, we refer to this type of function as *group oriented* evaluation function. As illustrated in [20], the evaluation function rewards teams for moving the centre of mass as far as possible from its initial position, while keeping the distance between the robots below a certain threshold. In homogeneous teams, the group response targeted by this evaluation function can only be obtained by generalist solutions that possess the mechanisms to negotiate a direction of motion (e.g., to allocate the role of leader and follower), and the mechanisms to execute the movements in a coordinated way. In other words, we argue that, in the task described in [20], any fitness increase in clonal evolution is likely to be induced only by a restricted set of coordinated actions which require the existence of both the mechanisms for the allocation and execution of complementary roles. It follows that the clonal approach can generate successful generalist controllers only through a limited set of evolutionary dynamics in which the mechanisms for the allocation

and execution of complementary roles evolves simultaneously. Based on this reasoning, we hypothesise that the *group oriented* evaluation function, in combination with the homogeneity of the teams characterising the clonal approach, determines the simultaneity argument illustrated in [20], and consequently the problems of the clonal approach if compared with the aclonal approach.

In order to test our hypothesis, we propose to compare the clonal and the aclonal approaches in a task-allocation scenario for two-robot teams in which the teams are rewarded by an *individual oriented* evaluation function. This type of function rewards the groups with reference to how each single individual contributes to each activity rather than to how the group collectively performs the task. We believe that the *individual oriented* evaluation function gives the clonal approach the possibility to find progressively better solutions for a single role independently of the solutions in place for the other role. With an *individual oriented* evaluation function, the mechanisms for playing the roles can evolve independently from the mechanisms to allocate them, in spite of the homogeneity condition. Unfortunately, the roles in the Quinn’s study (i.e., leader versus follower) are not necessarily based on different behavioural competencies¹. For this reason, we found it difficult to design an *individual oriented* evaluation function that was general enough to avoid dictating specific solutions to the task described in [20], and at the same time coherent to the principles illustrated above. Consequently, we had no choice than to change the task, while preserving those elements indicated by Quinn as responsible for the phenomena illustrated in his study. We wish to emphasise that the hypothesis formulated by Quinn refers only to the criteria (i.e., clonal and aclonal approaches) for the evolution of homogeneous controllers for two-robot teams engaged in a task requiring the individuals to take different roles. Thus, we believe that any task that complies with the above mentioned characteristics is suitable to test this hypothesis.

We have designed a two-robot task in which, as in [20], the robots interact only through the activation of their proximity sensors, and controllers are designed using exactly the same clonal and aclonal evolutionary approaches. Contrary to [20], in our task, the two roles are based on different behavioural responses, and we use an *individual oriented* instead of a *group oriented* evaluation function (see Table 1 for a summary of similarities and differences between our and Quinn’s task). The re-

¹ In [20], the roles are *a posteriori* identified based on the characteristics of the best evolved strategies. They are not part of the definition of the task, and their evolution is not imposed by the design of the evaluation function.




Fig. 2 Experimental scenario with the two robots placed within the nest (floor in shades of grey), and the foraging sites referred to as *L-site*, and *R-site*.

sults of this study show that: i) the clonal approach, in combination with an *individual oriented* evaluation function, can generate evolutionary trajectories where the mechanism for role-allocation and role-execution evolve at different evolutionary times; ii) without the limitations imposed by the *group oriented* evaluation function, the clonal approach outperforms the aclonal approach; iii) the aclonal approach exploits evolutionary paths that proceed from specialist to generalist solutions as illustrated in [20]. However, in our scenario, this transition is not as frequent as detailed in [20]. In view of the results of this work, we revise the hypothesis formulated in [20], we comment on the implications of the results of our study, and we indicate directions for future work.

5 The Task and the Simulation Environment

Teams comprising two simulated Khepera mini-robots are evaluated in the context of a dynamic role-allocation task. By taking inspiration from the behaviour of social insects, the roles are nest patrolling and foraging (hereafter, we refer to them as *role-patrolling*, and *role-foraging*, respectively). Roughly speaking, *role-patrolling* requires a robot to remain within the nest (i.e., an area in which the colour of the floor is in shades of grey). *Role-foraging* requires a robot to move back

Table 1 Summary of similarities and differences between our and Quinn [20]’s task.

| Similarities | Differences |
|---|-----------------------------|
| The number of robot | The task |
| The robot model | The evaluation function |
| The number of roles | The neuro-controller |
| Means of interaction/communication between the robots | The robot sensory apparatus |
| The robots’ initial relative orientations | |
| The clonal and aclonal approaches | |




Fig. 3 Kheperas' body-plan. The black circles refer to the position of infra-red (*IR*), ambient-light (*AL*), and floor sensors (*FS*). The dotted lines indicated view with the three camera's sectors.

and forth between the nest and any of the two foraging sites located in the environment. The robots are required to execute both roles simultaneously. Therefore, they should go through a role-allocation phase in which they autonomously decide who is doing what, and then execute their role.²

The environment is a boundless arena with a light bulb positioned 6 cm above the floor, and two red cylindrical objects (2.7 cm radius, and 10 cm height) positioned at 40 cm on the left and on the right of the light, respectively, and referred to as *L-site*, and *R-site*. The colour of the arena floor is white except for a circular area (15 cm radius), centred around the lamp, within which the floor is in shades of grey. The inner part of the circular area (up to 5 cm from the light) is black, the middle part (from 5 cm to 10 cm from the light) is dark grey, and the outer part (from 10 cm to 15 cm from the light) is light grey. The area in shades of grey represents the nest and the cylindrical objects represent the foraging sites (Figure 2).

Our simulation models a Khepera robot, a 2.7 cm radius cylindrical robot. It is provided with eight 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 be another robot or a foraging site); four ambient light sensors (AL^i with $i = 0, \dots, 3$) to detect light; a simple camera; and a floor sensor (*FS*) positioned facing downward on the underside of the robot (Figure 3). The IR and AL sensor values are extrapolated from look-up tables provided with the Evorobot simulator [15]. The IR sensors' range is approximately 4 cm. AL sensors have an angle of acceptance of 120°. Light levels change as a function of the robot's distance from the lamp. The floor sensor can be conceived of as a IR sensor capable of detecting

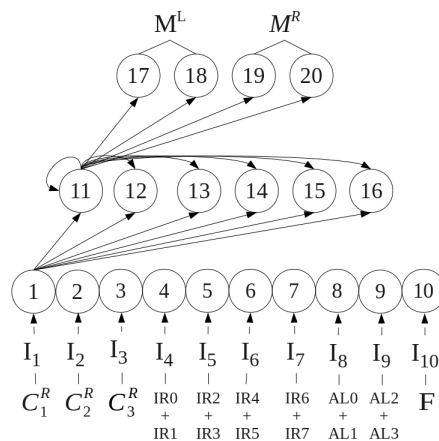


Fig. 4 (a) The neural network. Continuous line arrows indicate the efferent connections of the first neuron of each layer. Neurons on the same layer share the same type of efferent connections. Underneath the input layer, the correspondences between sensors, the notation used in equation 2 to refer to them, and the input neurons are shown.

the level of grey of the floor. It returns 0 if the robot is on white floor, 0.5 on light grey floor, 0.75 on dark grey floor, and 1 on black floor. The robots camera has a receptive field of 30°, divided in three equal sectors, each of which has three 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 in between [0,1]. The camera can detect coloured objects up to a distance of 60 cm. The robots can not see each other through the camera. The robots kinematics are simulated using the Differential Drive Kinematics equations, as illustrated in [6]. The robot has left and right motors which can be independently driven forward or in reverse, allowing it to turn fully in any direction. The robot is assumed to have negligible mass, so that the motor output can be taken as the tangential velocity of the robot to the motor mount point. The robot maximum speed is 8 cm/s. High levels of noise are applied to motor outputs, to guarantee that the simulated controller will transfer to a physically realised robot with no loss of performance.

6 Robot controllers and the Evolutionary Algorithm

The robot controller is a continuous time recurrent neural network (CTRNN) with 10 input neurons, 6 hidden neurons, and 4 output neurons [2]. Each hidden neuron is connected to all the other hidden neurons including itself. Additionally, each hidden neuron receives one incoming synapse from each input neuron. Each output neuron receives one incoming synapse from each hidden neuron. There are no direct connections between input and output neurons (Figure 4). The states of the

² See also <http://users.aber.ac.uk/elt7/suppPagn/TA2013/suppMat.html> for further methodological details, pictures, and videos.

output neurons are used to control the speed of the left and right wheels. The states of input, hidden and output neurons are updated using equations (1), (2), and (3).

$$y_i = gI_i; \text{ for } i \in \{1, \dots, 10\}; \quad (1)$$

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{16} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{11, \dots, 16\}; \quad (2)$$

$$y_i = \sum_{j=11}^{16} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{17, \dots, 20\}; \quad (3)$$

where, $\sigma(x) = (1 + e^{-x})^{-1}$. 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 $i = \{1, \dots, 10\}$ is the activation of the i^{th} input 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 input neurons share the same bias (β^I), and the same holds for all output neurons (β^O). τ_i and β_i with $i = \{11, \dots, 16\}$, β^I , β^O , all the network connection weights ω_{ij} , and g are genetically specified network parameters. At each time step, the output of the left motor is $M^L = f_{17} - f_{18}$, and the right motor is $M^R = f_{19} - f_{20}$, with $M_L, M_R \in [-1, 1]$. Cell potentials are set to 0 when the network is initialised or reset. Equation 2 is integrated using the forward Euler method with an integration time step $\Delta T = 0.1$.

A generational genetic algorithm is employed to set the network parameters [8]. At generation 0, a random population of 100 vectors is generated by initialising each component of each vector to a value chosen uniformly random in the range $[0, 1]$. Each vector comprises 135 real values (120 connections ω_{ji} , 6 decay constants τ_i , 8 bias terms β , and a gain factor g shared by all the input neurons). Hereafter, a vector is referred to as genotype and its components as genes.

Generations following the first one are produced by a combination of selection with elitism and mutation. At each new generation, the three highest scoring genotypes (“the elite”) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness-proportional selection from the 70 best genotypes of the old population. New genotypes, except “the elite”, are produced by applying mutation. Mutation is a random Gaussian offset applied to each gene, with a probability of 0.07. The mean of the Gaussian is 0, and its standard deviation is 0.1. During evolution, all genes are constrained to remain within the range $[0, 1]$. That is, if mutations cause a gene value to fall below zero, its value is fixed to 0; if it rises above 1, its value is fixed to 1.

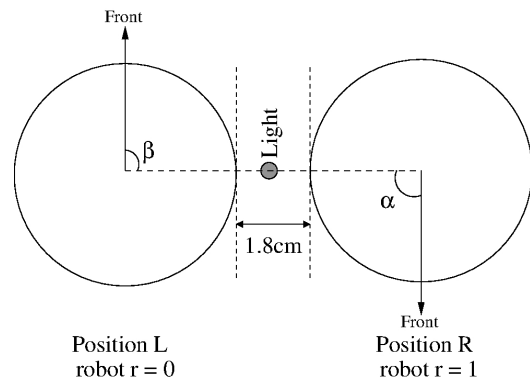


Fig. 5 The robots initial relative orientations and position with respect to the light. α and β are the parameters defining the set of 15 different initial team positions.

Genes are linearly mapped to produce network parameters with the following ranges: $\beta^I \in [-4, 4]$, $\beta^O \in [-5, 5]$, $\beta_i \in [-5, 5]$, with $i \in \{11, \dots, 16\}$, $\omega_{ji} \in [-8, 8]$, with $j \in \{1, \dots, 10\}$, and $i \in \{11, \dots, 16\}$, $\omega_{ji} \in [-10, 10]$, with $j \in \{11, \dots, 16\}$, and $i \in \{11, \dots, 20\}$, gain factor $g \in [1, 13]$. Decay constants τ_i with $i \in \{11, \dots, 16\}$, are firstly linearly mapped into the range $[-1.0, 1.7]$ and then exponentially mapped into $\tau_i \in [10^{-1.0}, 10^{1.7}]$. The lower bound of τ_i corresponds to the integration step-size used to update the controller; the upper bound, arbitrarily chosen, corresponds to about 8% of the maximum length of a trial. A trial is an evaluation sequence of up to 400 simulation cycles. All mapping ranges were chosen on the basis of having proved useful in other CTRNN experiments [25].

7 Evaluation Regime and Fitness Function

At the beginning of each trial, the robots are placed in the nest, located symmetrically on the left and on the right of the light, at 1.8 cm away from each other (Figure 5). Their controllers are reset. The initial relative orientation of the two robots is sufficiently described by a vector of two variables (α , β , see Figure 5). A sample set of starting configuration is chosen such that $\alpha, \beta \in (0, \frac{2\pi}{5}, \frac{4\pi}{5}, \frac{6\pi}{5}, \frac{8\pi}{5})$, leading to 25 combinations. From these combinations, 10 have been removed because they are rotational duplicates. This leaves the set of 15 relative starting orientations that have been used. For each orientation pair, uniform noise randomly chosen in the range ± 5 cm is added to the robots’ initial distance, and uniform noise randomly chosen in the range $\pm 5^\circ$ is added to α and β . The 15 orientation pairs include 5 symmetrical conditions in which $\alpha = \beta$. In symmetrical orientation pairs, in spite of the noise, the robots share the same perception at the beginning of the trial. Asymmetrical orientation pairs are those in which $\alpha \neq \beta$.

Each trial differs from the others in the initialisation of the random number generator, which influences the robots' initial distance and orientation, and the noise added to motors and sensors. See [11] for further details on sensors and motor noise. Within a trial, the team life-span is 40s ($T=400$ simulation cycles). Trials are terminated earlier if either one of the robots exceeds the arena limits (i.e., a circle of 120 cm radius, centred on the light), or the team exceeds the maximum number of collisions (i.e., 10), or a robot completes two foraging trips (i.e., two trips between any of the food sites and the nest).

The parameters of the evolutionary algorithms (i.e., clonal and aclonal approach) are identical to those illustrated in [20]. In clonal runs, the fitness of a genotype is its average team evaluation score after it has been assessed twice for each of the 15 starting configurations, for a total of $E = 30$ trials. The fitness of a genotype in an aclonal run is the average evaluation score of the team in which it participates. In aclonal runs, a genotype is evaluated four times for each starting configuration, twice from each of the robots positions (i.e., position L and position R, see Figure 5) comprising each configuration, for a total of $E = 60$ trials. Each one of an aclonal individual 60 trials is undertaken with a different, randomly chosen, partner. Note that, hereafter, we refer to the agent initialised in position L and R as robot $r = 0$ and $r = 1$, respectively.

Contrary to the approach of Quinn [20], we designed an evaluation function which rewards groups based on how the single individuals contribute to each sub-role. We first compute for each robot the performance in both roles—*role-foraging* and *role-patrolling*—independently from each other. Then, we compute a team performance considering the two possible role allocations, that is, robot 0 as patroller and robot 1 as forager and the other way round. We then consider the maximum team score between the two role allocations. Additionally, two group penalties are considered accounting for collisions and for exceeding the arena limits.

More precisely, the average team evaluation score is given by the following function:

$$F = \frac{1}{E} \sum_{e=1}^E \arg \max_{r \in \{0,1\}} (CP_e^r \times CF_e^{1-r}) \times Z_e^a \times Z_e^b;$$

where, for each trial e , $CP_e^r \in [0, 1]$ rewards robot $r = \{0, 1\}$ for staying in the nest; $CF_e^r \in [0, 4]$ rewards robot r for travelling twice the distance from the nest to any of the two food sites; the team collision penalty Z_e^a is inversely proportional to the number of collisions, with $Z_e^a = 1$ if no collisions are recorded, and $Z_e^a = 0$ if 10 collisions are recorded; Z_e^b is the team penalty for

exceeding the arena's limits, with $Z_e^b = 1$ if none of the robots exceeds the limits, $Z_e^b = 0.3$ otherwise. The maximum value of the average team evaluation score F is 4.

For each robot r and for each trial e , the fitness components for playing *role-patrolling* and *role-foraging* are computed using the following:

$$CP_e^r = 1.0 - \sqrt{\frac{S_r}{T}}; \quad (4)$$

$$CF_e^r = AT_FOOD + AT_NEST; \quad (5)$$

$$AT_FOOD = \min[V^{food} + (1 - \bar{D}^{food}); 2] \quad (6)$$

$$AT_NEST = \min[V^{nest} + (1 - \bar{D}^{nest}); 2]; \quad (7)$$

where, S_r is the number of time steps a robot spends outside the nest. $\min[V^{food} + (1 - \bar{D}^{food}); 2]$ rewards foraging behaviour. V^{food} is the number of visit to food sites, \bar{D}^{food} the normalised distance to the nearest food site. A food site is considered visited by a robot when $D^{food} < 4.6$ cm. \bar{D}^{food} is set to zero if the robot is inside the nest or is looking for the nest after a visit to a food site. $\min[V^{nest} + (1 - \bar{D}^{nest}); 2]$ rewards homing behaviour. V^{nest} is the number of visits to the nest following a visit to a food source; \bar{D}^{nest} is the normalised distance to the nest. \bar{D}^{nest} is set to zero if the robot is inside the nest or is currently foraging.

8 Results

Each experiment (i.e., clonal and aclonal) consisted of 20 evolutionary runs, each using a different random initialisation. Each run lasted 3000 generations. Recall that our objective is to compare the performances of the clonal and aclonal approach for the evolution of homogeneous two-robot teams capable of dynamically allocating and simultaneously executing *role-patrolling* and *role-foraging*. Following the procedure illustrated in [20], at the end of the evolutionary phase, we run a first set of re-evaluations consisting of 60 trials per team (i.e., 4 times for each of the 15 starting orientation mentioned in section 2.3). As in [20], the 7 fittest genotypes of each generation of both clonal and aclonal runs are re-evaluated in a homogeneous setup. The average re-evaluation score of each genotype is measured using the metrics F illustrated in section 2.3. The highest average re-evaluation score recorded during a given run is assumed to be an adequate measure of the success of that run.

The results of our re-evaluations are summarised in Table 2, which shows a comparisons of mean, median scores, and the mean ranking of both approaches. Each measure shows the clonal out-performing the aclonal, and the difference between the two set of results is statistically significant (Mann-Whitney U test, $p < 0.01$).

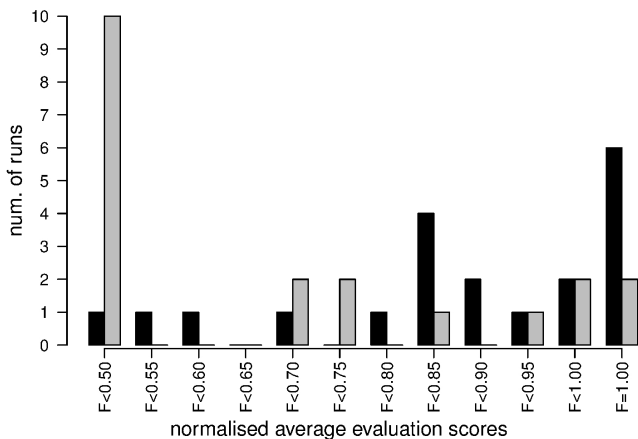


Fig. 6 The histogram shows the distributions of the highest average re-evaluation scores achieved by each run of the clonal (black bars) and aclonal approach (grey bars). Values represent normalised average evaluation scores F .

In particular, eight runs of the clonal approach produced high-scoring teams exceeding 95% of the optimal score, with six of them 100% successful (see Figure 6, black bars). In contrast, only four out of twenty of the aclonal runs generated a homogeneous team that exceeds 95% of the optimal score, with two of them capable of completing the 60 re-evaluation trials with the highest score. From a statistical point of view, there is enough evidence to prefer one approach over the other for the evolution of homogeneous multi-robot teams engaged in this dynamic role allocation scenario. Moreover, from the point of view of generating optimal controllers, the clonal approach does better than the aclonal one. That is, 6 teams generated with the clonal approach yielded a 100% success rate (see Figure 6, black bars). Only 2 aclonally generated teams yielded the same result (see Figure 6, grey bars).

Our results clearly show that the clonal approach outperforms the aclonal approach in generating homogeneous teams. This evidence not only diverges from what was shown in [20], but also questions the hypoth-

Table 2 Table showing, for clonal and aclonal approaches, mean ranking, median, mean and standard deviation of the trials' scores obtained at the re-evaluation test. Recall that the maximum score in a trial is 4.

| | mean ranking | median | mean (s.d.) |
|---------|--------------|--------|-------------|
| clonal | 25.8 | 3.51 | 3.35 (0.79) |
| aclonal | 15.2 | 2.26 | 1.99 (1.52) |

esis formulated by Quinn to account for the superiority of the aclonal over the clonal approach in dynamic task-allocation scenarios. Recall that our original hypothesis, motivating this study, was that the results shown by [20] were determined by task-specific elements. The analysis we illustrate in the next sections indicates to what extent the type of evaluation function affects the evolutionary dynamics in both approaches. In view of the results of our tests, we revise the argument formulated by [20].

8.1 The simultaneity argument

In this section, we investigate the relationships between the characteristics of the evolutionary scenario and the evolutionary dynamics observed in successful groups generated clonally. Recall that, according to what illustrated in [20], the clonal approach is penalised by the fact that the mechanisms to play the roles and the mechanisms to allocate them have to evolve simultaneously. Since in our scenario, clonal runs outperform aclonal runs, we first question the simultaneity argument to account for the difference between our results and those shown in [20]. The tests illustrated in the next paragraphs aim to verify to what extent the simultaneity argument concerns successful groups generated clonally.

We analyse the evolutionary trajectory of the best six evolved groups generated clonally from different evolutionary runs and we re-evaluated all their ancestors four times for each of the 15 orientation pairs. This phylogenetic analysis is made possible by the fact that we did not use recombination during evolution, meaning that every genotype has only one parent. Figure 7 shows the results of this analysis only for two best clonally generated groups, indicated as run 1 and run 2. This is because the other four successful runs produced evolutionary dynamics very similar to either the one produced by run 1 or by run 2. Moreover, in Figure 7, the performances in symmetrical (i.e., $\alpha = \beta$) and asymmetrical (i.e., $\alpha \neq \beta$) trials are plotted separately. This is due to the fact that, from previous tests, we noticed that all successful groups, generated both clonally and aclonally, employ different behavioural strategies for symmetrical and asymmetrical trials. In asymmetrical conditions, the robots use the differences in the initial perceptual states to break the system homogeneity (i.e., the controllers are in the same initial state) and to allocate roles. In symmetrical conditions, random fluctuations are integrated over time and lead to breaking the initial perceptual symmetries and the homogeneity condition, and thereby resulting in the allocation of roles (data shown in the supplementary materials,

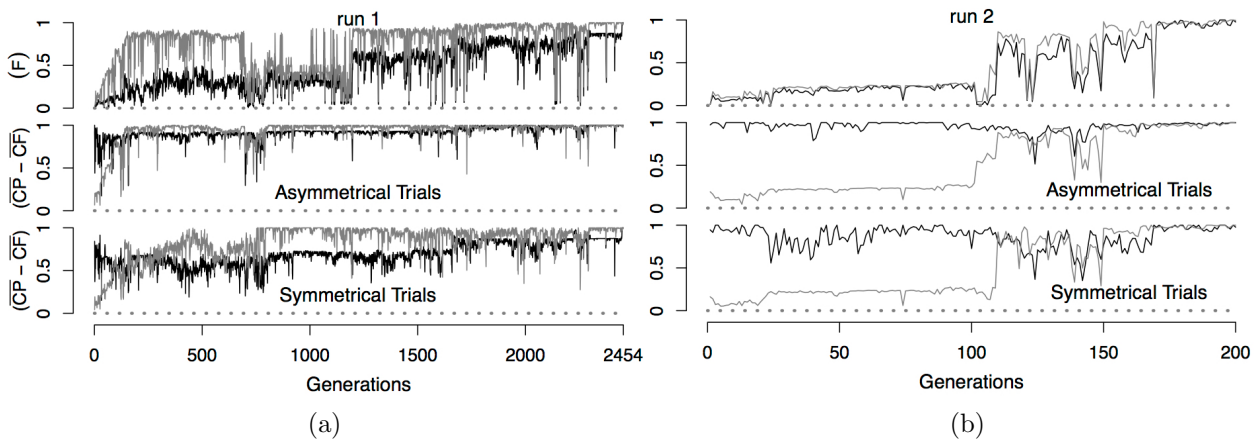


Fig. 7 Graph (a) refers to a successful homogeneous group generated by clonal run n. 1; graph (b) to a different successful homogeneous group generated by clonal run n. 2. Both graphs show various aspects of the evolutionary history of these two different groups. In both (a) and (b), the top graph shows, for each generation, the ancestors' fitness measured on symmetrical trials (black line) and on asymmetrical trials (grey line). The middle and the bottom graph show, for each ancestor, the normalised average values of two fitness components in asymmetrical and symmetrical trials, respectively. The first component rewarding *role-patrolling* (\bar{CP}), is indicated with the black line, and computed using $\bar{CP} = \frac{1}{E} \sum_{e=1}^E CP_e^r$. The second component rewarding *role-foraging* (\bar{CF}), is indicated with the grey line, and computed with $\bar{CF} = \frac{1}{E} \sum_{e=1}^E CF_e^r$. The role that an ancestor plays in a trial is determined by how it contributes to the team fitness in that trial. For example, robot 0 plays *role-patrolling* and robot 1 plays *role-foraging* if, in a trial e , $CP_e^r \times CF_e^{1-r}$ is bigger than $CF_e^r \times CP_e^{1-r}$, or vice versa. $E = 30$ in the symmetrical condition, and $E = 60$ in the asymmetrical condition.

see footnote 2). In view of the the above mentioned behavioural differences, and of the causal relationship that links the conditions during evaluation and the evolutionary dynamics observed, we decided to consider the symmetrical and asymmetrical trials as two distinctive cases in the evolutionary analysis of successful groups generated clonally.

The graphs at the top in Figure 7 shows, for run 1 and 2, their evolutionary history through the normalised average fitness in symmetrical (black line) and asymmetrical trials (grey line). The middle and the bottom graphs in Figure 7 show for asymmetrical and symmetrical trials, respectively, the evolutionary trend of the two main fitness components that contributed to generate the evolutionary trajectories shown in the top graphs of Figure 7. For middle and bottom graphs, the black lines refers to the fitness component that rewards agents for playing *role-patrolling*. The grey lines refers to the fitness component that rewards agents for playing *role-foraging* (see the caption of Figure 7 for a detailed description of these graphs).

If we look at the top graphs, we notice that while run 2 reaches the maximum fitness in less then 200 generations, run 1 takes about 2500 generations to reach a similar performance. In spite of this marked difference, both runs are characterised by similar evolutionary trends for what concerns asymmetrical trials, where the neural machinery required to solve the task appears, for both runs, relatively early during evolution (see Figure 7, middle graphs). For asymmetrical trials, in both

runs, the mechanisms required to play *role-patrolling* appear before the mechanisms required to play *role-foraging* (see Figure 7, middle graphs, black and grey lines). In other words, the oldest ancestors in both runs, are robots that, in asymmetrical trials, tend to remain in the proximity of their initial position. This means that the group fitness, in the very early stages of the evolution, is largely determined by the component that rewards nest patrolling behaviour (see Figure 7, middle graphs, black line). The mechanisms required to play *role-foraging* are progressively acquired in subsequent evolutionary times (see Figure 7, middle graphs, grey line).

For what concerns symmetrical trials, in run 1 successful groups emerge after a quite long evolutionary process, while in run 2 evolution finds quite quickly the way to successful strategies (see Figure 7, bottom graphs). In spite of this difference, we notice that in both runs the mechanisms required to play each role appear at different evolutionary times. Particularly interesting is the trend observed in run 1, where, contrary to what happens in run 2, solutions in symmetrical and asymmetrical trials evolve in a completely different way. In asymmetrical trials, the robots are able to get very quickly to the maximum score in *role-patrolling* (see Figure 7a, middle graph, black line). In symmetrical trials, *role-patrolling* appears to be a slightly more complex response that takes longer time to be optimally developed (see Figure 7a, bottom graphs, black line). The mechanisms to play *role-foraging* also take longer

to develop in symmetrical then in asymmetrical trials (see Figure 7a, middle and bottom graphs, grey line). Other two successful clonal runs produced evolutionary dynamics similar to those observed in run 1, where solutions for symmetrical trials take longer to appear than for asymmetrical trials. Seemingly, the combination of the initial perceptual symmetry and the homogeneity condition can make the task particularly challenging for the evolutionary process.

In summary, we have shown that, for what concerns the six successful groups generated clonally, the mechanisms required to play *role-patrolling* and *role-foraging* evolve at different evolutionary times, both for symmetrical and asymmetrical conditions. Although we do not have any direct evidence of the evolutionary history of the mechanisms to allocate the roles, we can indirectly infer that these mechanisms progressively emerged thanks to the capabilities of the ancestors to improve their performances in one role without losing fitness on the other. We claim that such dynamics are enormously facilitated by the nature of our fitness function which, by rewarding the group based on the best combination of individual-per-task generates alternative evolutionary paths to those described in [20]. In other words, the results of these tests show that under particular evolutionary conditions, the clonal approach can generate successful homogeneous groups engaged in task allocation scenario by going through alternative path to the simultaneous emergence of the mechanisms to play the role and those to allocate them. The use of an *individual oriented* rather than a *group oriented* fitness function creates the conditions for the gradual and diachronic evolution of mechanisms for playing the roles. The independent and progressive improvement of the groups in both roles is automatically linked to the emergence of the mechanisms for allocating them.

8.2 The specialisation argument

In this section, we focus on the role of specialisation in a clonal runs. According to what was claimed in [20], in the a clonal approach the evolution of homogeneous solutions is preceded by an intermediate evolutionary step in which populations are composed of specialised agents. These are agents capable of playing one role or the other, but not both. Whenever two complementary specialised agents meet in a group, they successfully complete the collective task. The author in [20] claims that specialisation paves the way to the following evolution of generalist or homogeneous solutions, and this is what makes the a clonal approach more effective than the clonal approach in a dynamic role-allocation scenario. Since in our scenario a clonal runs did not per-

form as well as the clonal runs, we question the significance of specialisation. In particular, we investigate to what extent a clonally generated successful homogeneous solutions exploit the evolutionary path described by Quinn [20].

To carry out this analysis, we re-evaluate for 90 trials (i.e., six times in each of the 15 orientation pairs) all the ancestors of five homogeneous solutions, generated from five different a clonal runs, whose scores were higher than 85% of the optimum. The ancestors are re-evaluated in three different conditions: 1) *Test Homo*, homogeneous groups; 2) *Test Hete-A*, heterogeneous groups, with a partner specialised in *role-patrolling*; 3) *Test Hete-B*, heterogeneous groups, with a partner specialised in *role-foraging*. Each specialised partner is a robot whose controller is generated by a genotype chosen from all the genotypes produced by the 20 a clonal evolutionary runs. Specialised partners have been selected because, when repeatedly evaluated both in homogeneous and heterogeneous conditions, they showed a strong preference for one or the other role. That is, the robot controlled by the neural network generated by a specialised genotype systematically plays one of the two roles. These tests should help us to clarify whether, and during which evolutionary time, the ancestors of successful a clonally generated homogeneous groups were specialists. The rationale is that, if ancestors are generalists, they will play, both in *Test Hete-A* and in *Test Hete-B*, the complementary role to the one played by the specialised partner. If ancestors are specialist, then in either *Test Hete-A* or *Test Hete-B*, they will play the same role of the specialised partner.

Figure 8 shows the results of these tests for a particular successful a clonally generated group, and only for a particularly interesting evolutionary time window when we observed the appearance of the mechanisms for the dynamic allocation of roles. Figure 8a refers to the results of *Test Homo*, in which the ancestors are evaluated in homogeneous groups. The graph shows the fitness of the a clonal ancestors, re-evaluated in a homogeneous setup. We clearly notice that around generation 200, there is a marked increment in the group fitness. That is, around generation 200, ancestors of a successful a clonally generated solution become capable of successfully accomplishing the task in homogeneous conditions. The fitness increment can only be determined by the evolution of the mechanisms for the dynamic allocation of the roles. The graphs in Figure 8b and 8c tell us more about the behavioural capabilities of the ancestors preceding generation 200.

Figure 8b refers to the results of *Test Hete-A*, in which the ancestors of successful a clonally generated homogeneous solutions are evaluated in hetero-

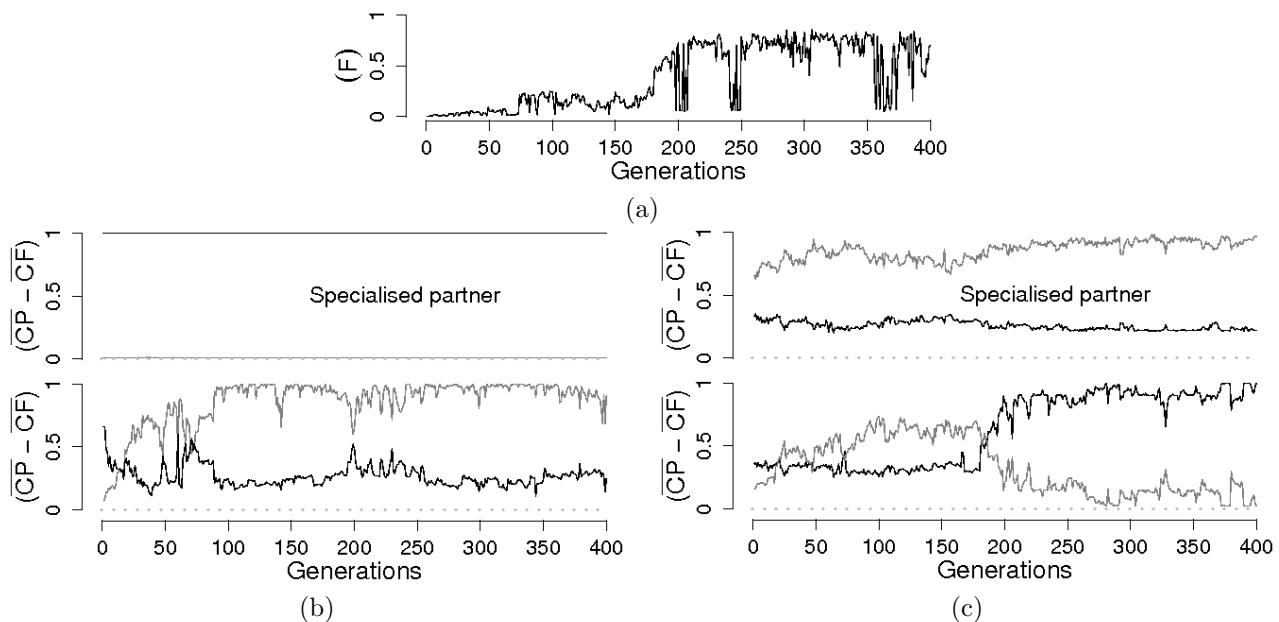


Fig. 8 Graphs showing the results of three different tests on the ancestors of a successful homogeneous group generated aclonally. The graph in (a) shows the fitness of the aclonal ancestors, re-evaluated in a homogeneous setup. In (b) and (c), the graphs show the normalised average values of the fitness components \overline{CP} rewarding *role-patrolling* (see black lines), and \overline{CF} rewarding *role-foraging* (see grey lines), for the aclonal ancestors (bottom graph), and for the specialised partner (top graph). In (b), the partner is specialised in *role-patrolling*. In (c), the partner is specialised in *role-foraging*.

geneous conditions, with a partner that systematically plays *role-patrolling*. As expected the specialised robot is optimal in playing *role-patrolling* (see Figure 8b, top graph, black lines). The ancestors, except from the initial 100 generations, proved to be capable of adjusting to the preference of the partner by playing *role-foraging*. This can be inferred by the fact that the average values of the fitness component \overline{CF} , rewarding *role-foraging*, are higher than the average values of the fitness component \overline{CP} , rewarding *role-patrolling* (see Figure 8b, bottom graph, grey line for \overline{CF} , and black line \overline{CP}). We conclude that, for a period of about 100 generations preceding the evolution of the mechanism for the dynamic allocation of roles, the ancestors are capable of successfully playing *role-foraging* when re-evaluated with a partner specialised in *role-patrolling*.

Figure 8c refers to the results of *Test Hete-B*, in which the ancestors of successful aclonally generated homogeneous solutions are evaluated in heterogeneous conditions, with a partner that systematically play *role-foraging*. As expected the specialised partner is rather effective in playing *role-foraging* (see Figure 8c, top graph, grey line). The ancestors are not as good as the partner in playing the complementary *role-patrolling* (see Figure 8c, bottom graph, black line). For large part of the evolutionary time preceding the appearance of the mechanism for the dynamic allocation of role (i.e., from generation 50 to generation 200), the aver-

age values of the fitness component \overline{CF} , rewarding *role-foraging*, are higher than the average values of the fitness component \overline{CP} , rewarding *role-patrolling* (see Figure 8c, bottom graph, grey line for \overline{CF} , and black line \overline{CP}). We also notice that, during the first 200 generation, the ancestors are not as good in playing *role-patrolling* when re-evaluated with a partner that is specialised in *role-foraging* (see Figure 8c, bottom graph, black lines) as they are in playing *role-foraging*, when re-evaluated with a partner that is specialised in *role-patrolling* (see Figure 8b, bottom graph, grey line). This evidence indicates that the ancestors preceding the evolution of the mechanisms for the dynamic allocation of roles, show a strong preference for playing *role-foraging*, regardless of the role played by the partner.

In other words, Figure 8 indicates that specialisation characterises the early part of the evolutionary history of solutions generated aclonally (i.e., from about generation 50 to about generation 200). Generalist solutions chronologically follow the evolution of specialisation. Since qualitatively similar evolutionary trends have been observed in all the evolutionary history of the five successful homogeneous solutions generated aclonally (data shown in the supplementary materials, see footnote 2), we conclude that, as stated in [20], specialisation precedes the emergence of generalist or homogeneous solutions in aclonal evolutionary runs. We will further speculate on this in section 9.

9 Discussion and conclusions

More than a decade ago, Quinn [20] discussed the differences between the aclonal and the clonal approach for the evolution of homogeneous multi-robot teams for tasks which require robots to take specific roles, showing that the former outperform the latter. Our study described a similar robotic task. We designed a task-allocation scenario where two robots, structurally identical to those used in [20], have to carry out a collective task made of two distinctive activities: that is, foraging and nest-patrolling. The robots have to autonomously decide who is doing what and then carry out the sub-tasks. As in [20], the robots can interact only through the activation of the proximity sensors. Contrary to what is illustrated in [20], the activities require different behavioural skills (see section 5 for details). Moreover, we used a *individual oriented* rather than a *group oriented* evaluation function (see section 7 for details).

Contrary to the results of Quinn [20], we found that the clonal approach outperforms the aclonal approach in generating successful homogeneous teams capable of solving the task. The analysis of the evolutionary trajectories of clonal and aclonal runs produced evidence that only partially support the reasoning put forward by Quinn [20] to account for his results. As illustrated in [20], we also found that aclonal runs exploit evolutionary dynamics based on the early appearance of specialised solutions, which, in our scenario, only in a limited number of runs, have been followed by generalist solutions. Additionally, in our study, the clonal approach has found successful collective strategies by travelling on alternative evolutionary paths than those suggested by [20]. Clonal evolutions have searched the space of possible solutions by capitalising on gradual improvements on the execution of single behavioural roles, and on the appearance of allocation mechanisms (initially) bounded to specific ecological conditions. That is, in our scenario, solutions tend to appear earlier in asymmetrical than symmetrical trials (see section 8.1 for details).

The results of our study induce us to review the simultaneity argument formulated in [20]. We showed that, in a scenario in which the differences between the roles is captured by an evaluation function that multiply robot-based (instead of team-based) factors, the simultaneous evolution of the mechanisms for role-allocation and role-execution is not the only way to the emergence of successful collective strategies. Our results clarify that the clonal approach should be considered in view of the task-allocation scenario the robots are required to solve rather than in view of the simultaneity argument formulated in [20]. Mechanisms for

role-allocation and role-execution can evolve at different evolutionary time also in clonal runs. The simultaneity argument does not necessarily limit the effectiveness of this evolutionary approach for the design of two-robot teams engaged in task-allocation scenarios.

As far as it concerns the aclonal approach, we also believe that there may be other circumstances to those already shown by [20], in which the aclonal approach may be more effective than the clonal one. For example, we have observed that perceptual symmetries are better handled in aclonal than in clonal evolution. We believe that future comparative work is certainly required to estimate how useful the aclonal approach can be for the evolution of homogeneous multi-robot teams engaged in task-allocation scenarios. A step further on this direction may be represented by studies aimed to clarify how aclonal populations make the transition from specialist to generalist solutions. We found that the aclonal approach found specialist solutions in 8 out 20 runs. Only 5 of them produced sufficiently successful generalist solutions (i.e., homogeneous groups with a performance higher than 85% of the fitness optimum). If all the 8 aclonal runs had moved from specialist to generalist solutions, the aclonal approach would have done better than the clonal one. We speculate that the number of genotypes against which each solution is evaluated can be a crucial parameter to induce the specialist to generalist transition in aclonal runs. Recall that during the phase in which agents are specialised, the populations are characterised by agents with different preferences. If a specialised agent is lucky enough to encounter only agents that prefer complementary roles, it gets the highest fitness. In such case, its fitness would not be different from the one of a generalist agent. Thus, it would not be possible for evolution to favour the latter to the former. However, the higher the number of different agents against which a single solution is evaluated, the higher the probability to encounter agents with similar preferences, the greater the difference between specialised and generalist solutions. Ideally, evaluating every solution against all the others in the population would maximise the selective advantage of generalist over specialised solutions. However, such an approach may be too computationally expensive, because it depends on the population size, and its effects remain subject to the influence of all the other evolutionary parameters.

In natural swarms, tasks are allocated to workers not only on the basis of morphological structure (in polymorphic species) and/or age of the agents, but also on the basis of short-term or long-term emergencies [see 17]. It seems that genetic diversity in various eusocial species is associated to the capability of “genetically related” workers to carry out multiple concomitant tasks

in an efficient way [see 16]. In particular, various studies show that division of labour in ants and bees may be, in part, a consequence of genetic variation for the agents' tendency to switch among different tasks. Workers in a colony respond in different way to stimuli associated to various tasks. These differences are the causal factor that regulates the distribution of agents to tasks. The process is guided by a combination of positive and negative feedback mechanisms. If the need for a particular activity increases due to change in colony conditions, the workers with a high response threshold for a task, which in normal condition would not perform it, are progressively attracted to the task in response to an increase level of the stimulus associated to it. For example, in Gordon [9], the author shows that in a species of harvester ants (*Pogonomyrmex barbatus*), foragers can be recruited from workers originally performing other tasks (e.g., nest patrolling, nest maintenance) when the quantity of food close to the nest is experimentally manipulated. These mechanisms can also underpin the inverse process. For example, if the increase in the number of agents performing a task decreases the stimulus associated to the task, then those workers with a higher response threshold are likely to abandon the task, reducing the number of agents performing it.

The role of genetic variability in coordinating the colony response to changing conditions is still under scrutiny by entomologists. However, these biological evidences suggest to roboticists that genetic variability can be an effective mechanisms to generate dynamic task-allocation processes in group of cooperating robots. This suggests that further research is needed to provide effective solutions to the issues mentioned in the previous paragraph, in order to make, within the context of dynamic task-allocation, a more effective use of heterogeneous groups with evolutionary design methods.

Acknowledgements V. Trianni acknowledges funding from the H²SWARM research project (European Science Foundation).

References

1. Ampatzis C, Tuci E, Trianni V, Christensen A, Dorigo M (2009) Evolving self-assembly in autonomous homogeneous robots: Experiments with two physical robots. *Artificial Life* 15(4):465–484
2. Beer RD, Gallagher JC (1992) Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior* 1(1):91–122
3. Bongard J (2000) The legion system: A novel approach to evolving heterogeneity for collective problem solving. In: Poli R (ed) *Genetic Programming: Proceedings of EuroGP-2000*, Springer-Verlag, LNCS, pp 16–28
4. Camazine S, Deneubourg JL, Franks N, Sneyd J, Theraulaz G, Bonabeau E (2001) *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ
5. Dorigo M, Floreano D, Gambardella L, Mondada F, Nolfi S, Baaboura T, Birattari M, Bonani M, Brambilla M, Brutschy A, Burnier D, Campo A, Christensen A, Decugniere A, Di Caro G, Ducatelle F, Ferrante E, Forster A, Gonzales J, Guzzi J, Longchamp V, Magnenat S, Mathews N, Montes de Oca M, OGrady R, Pinciroli C, Pini G, Retornaz P, Roberts J, Sperati V, Stirling T, Stranieri A, Stutzle T, Trianni V, Tuci E, Turgut A, Vausard F (2013) Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *Robotics Automation Magazine, IEEE* 20(4):60–71
6. Dudek G, Jenkin M (2000) *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, UK
7. Floreano D, Mitri S, S M, Keller L (2007) Evolutionary conditions for the emergence of communication in robots. *Current Biology* 17:514–519
8. Goldberg DE (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA
9. Gordon D (1996) The organisation of work in social insects. *Nature* 380:121–124
10. Ijspeert A, Martinoli A, Billard A, Gambardella L (2001) Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots* 11(2):149–171
11. Jakobi N (1997) Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior* 6:325–368
12. Luke S, Spector L (1996) Evolving teamwork and coordination with genetic programming. In: *Genetic Programming 1996. Proceedings of the 1st Annual Conference*, MIT Press, Cambridge, MA, pp 150–156
13. Nitschke G, Schut M, Eiben A (2007) Emergent specialization in biologically inspired collective behavior systems. In: *Intelligent complex adaptive systems*, IGI, New York, pp 100–140
14. Nitschke G, Schut M, Eiben A (2009) Collective neuro-evolution for evolving specialised sensor resolutions in multi-rover task. *Evolutionary Intelligence* 3(1):13–29
15. Nolfi S (2000) *EvoRob 1.1 User Manual*. Institute of Psychology, National

- Research Council (CNR), available at <http://gral.ip.rm.cnr.it/evorobot/simulator.html>
16. Oldroyd B, Fewell J (2007) Genetic diversity promotes homeostasis in insect colonies. *Trends in Ecology and Evolution* 22(8):408–413
 17. Page R (1997) The evolution of insects society. *Endeavour* 21(3):114–120
 18. Panait L, Luke S (2005) Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems* 11:387–434
 19. Potter M, LAMeeden, schultz A (2001) Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists. In: *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, pp 1337–1343
 20. Quinn M (2001) A comparison of approaches to the evolution of homogeneous multi/robot teams. In: *Proceedings of them International Conference on Evolutionary Computation*, vol 1, pp 128–135
 21. Quinn M (2004) The evolutionary design of controllers for minimally-equipped homogeneous multi-robot systems. PhD thesis, University of Sussex, School of Informatics
 22. Quinn M, Smith L, Mayley G, Husbands P (2003) Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of The Royal Society A* 361:2321–2343
 23. Trianni V, Nolfi S (2011) Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial Life* 17(3):183–202
 24. Trianni V, Nolfi S, Dorigo M (2008) Evolution, self-organisation, and swarm robotics. In: Blum C, Merkle D (eds) *Swarm Intelligence. Introductions and Applications*, Natural Computing Series, Springer Verlag, Berlin, Germany, pp 163–192
 25. Tuci E (2009) An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. *Biological Cybernetics* 101(3):183–199
 26. Waibel M, Floreano D, Magnenat S, Keller L (2006) Division of labour and colony efficiency in social insects: effect of interactions between genetic architecture, colony kin structure and rate of perturbation. *Proceedings of the Royal Society, B* 273:1815–1823
 27. Waibel M, Keller L, Floreano D (2009) Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transaction of Evolutionary Computation* 13(3):648–660
 28. Waibel M, Floreano D, Keller L (2011) A quantitative test of hamilton’s rule for the evolution of altruism. *PLoS Biology* 9(5):1–7