

On the Evolution of Homogeneous Multi-robot Teams: Clonal versus Aclonal Approach

Elio Tuci¹ and Vito Trianni²

¹ Aberystwyth University, Computer Science Department
Llandinam Building, Aberystwyth, Ceredigion, SY23 3DB, UK
elt7@aber.ac.uk, <http://users.aber.ac.uk/elt7/>

² Institute of Cognitive Science and Technology
Via San Martino della Battaglia 44, 00185 Rome, Italy
vito.trianni@istc.cnr.it, <http://laral.istc.cnr.it/trianni/>

Abstract. This study compares two different evolutionary approaches to the design of homogeneous multi-robot teams in a task that requires the agents to specialise in different roles. Our results diverge from what illustrated in a previous similar comparative study, which advocates for the superiority of the aclonal versus the clonal approach. We question this argument in view of new empirical evidence showing that the two approaches perform equally well in generating homogeneous teams.

Keywords: Swarm Robotics, Evolutionary Robotics, Dynamic Task Allocation.

1 Introduction

Homogeneous multi-robot systems are a class of autonomous multi-agent systems in which all robots of a team have identical physical structure and identical decentralised control system. Like social insects, homogeneous robots are generally required to coordinate their actions in order to maximise the efficiency of the team [9]. The synthesis of controllers for homogeneous multi-robot teams is a complex problem that has been faced with a large number of different techniques [6]. Among the various possibilities, Evolutionary Robotics represents a viable approach for the automatic synthesis of robot controllers requiring little a priori knowledge about the solution of a given problem [5]. In recent years, a number of studies have investigated the dynamics underlying the evolution of multi-robot systems in order to develop a principled understanding of how to guide self-organisation [7, 2, 10, 8]

One of the first papers to focus on this issue is the study described in [7], in which the author compares the clonal and the aclonal approach for the design of homogeneous controllers for teams of two agents required to move in a coordinated way by remaining within sensor range. The clonal approach refers to the use of a single genotype to generate a homogeneous team. The aclonal approach uses different genotypes from the same evolving population to generate teams in which the agents have different controllers (i.e., heterogeneous teams). The

study shows that, regardless of the theoretical disadvantages clearly listed and discussed in the paper, the aclonal approach can be a more efficient way than the clonal approach to generate homogeneous systems in which the team members have to autonomously specialise for the benefit of the team. To account for these counter-intuitive results, the author formulates a hypothesis according to which the aclonal approach takes advantages of specific evolutionary dynamics that are precluded to the clonal approach. In particular, in the aclonal approach behavioural roles can be developed and refined in genetically specialised agents, prior to the evolution of generalist solutions; that is, genetically identical individuals that use dynamic role allocation mechanisms to specialise in different roles. In other words, specialised solutions pave the way to the evolution of generalist ones. In the clonal approach, this gradual evolution from specialised to generalist solutions is not possible, because, given that the agents are clones, the adoption of complementary roles necessarily requires the existence of some dynamic role allocation mechanisms. Thus, behavioural roles and the mechanisms to allocate them have to (laboriously) evolve simultaneously.

This study provides further comparisons between clonal and aclonal approaches for the evolution of homogeneous multi-robot teams for tasks that require individuals to take specific roles. We are moved by the hypothesis that the results shown in [7], concerning the superiority of the aclonal versus the clonal approach, may have been affected by task-specific features, such as: a) the fact that functional differentiation between the roles may not be a prerequisite to perform the task; and b) the use of an evaluation function composed of team-based metrics (e.g., the position of the team given by the centre-point between the robots). Point (a) calls into question the causal relationship between the supposed nature of the task and the evolutionary dynamics observed. Point (b) calls into question the strong constraints that the evaluation function may have imposed to the evolutionary dynamics of the clonal approach. Contrary to what shown in [7], in our scenario, the roles are clearly different, and the evaluation function is made of robot-based (instead of team-based) factors. We show that in these conditions the clonal and aclonal approaches perform equally well in generating homogeneous teams. Our results indicate that the argument formulated in [7] concerning the superiority of the aclonal versus the clonal approach is based on an interpretation of the data that would require further empirical support to be validated. This is because we reproduced those data showing that they are open to alternative interpretations, not considered in [7].

2 Methods

2.1 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 P*, and *role F*, respectively). Roughly speaking, *role P* requires a robot to remain within the nest (i.e., an area in which the colour of the floor is

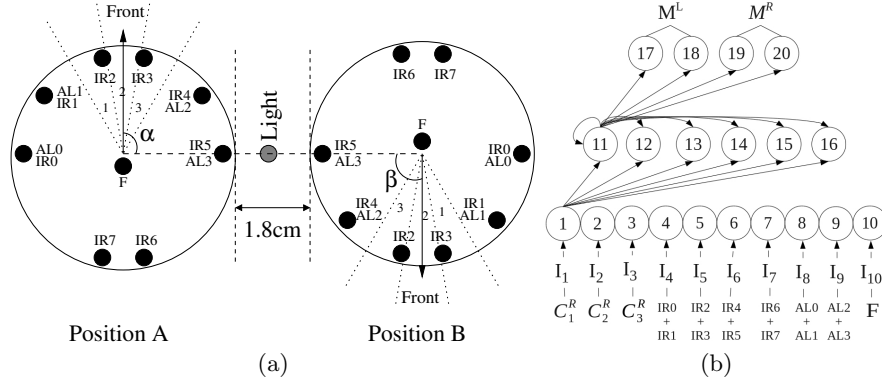


Fig. 1. a) Kheperas’ body-plan. The black circles refer to the position of infra-red (*IR*), ambient-light (*AL*), and floor sensors (*F*). The dotted lines indicated view with the camera’s sectors. α and β are the parameters defining the set of 15 different initial team positions. The grey circle is the light. b) 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, it is shown the correspondences between sensors and sensor neurons.

in shades of grey). *Role F* requires a robot to move back 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 role³.

The environment is a boundless arena with a light bulb positioned 6cm above the floor, and two red cylindrical objects (2.7cm radius, and 10cm height) positioned at 40cm 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 (15cm radius), centred around the lamp, within which the floor is in shades of grey. The inner part of the circular area (up to 5cm to the light) is black, the middle part (from 5cm to 10cm from the light) is dark grey, and the outer part (from 10cm to 15cm to the light) is light grey. The area in shades of grey represents the nest. The cylindrical objects represent the foraging sites.

The robots kinematics are simulated using a modified version of the “minimal simulation” technique described by Jakobi in [3]. Our simulation models a Khepera robot, a 2.7cm 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 (AL^i with $i = \{0, \dots, 3\}$) sensors to detect light; a linear camera; and a floor sensor (*F*) positioned facing downward on the underside of the robot (see Fig. 1a). The IR and AL sensor val-

³ See also <http://users.aber.ac.uk/elt7/suppPagn/sab2012/suppMat.html> for further methodological details, pictures, and movies of best teams.

ues are extrapolated from look-up tables provided with the Evorobot simulator (see [4]). The F sensor can be conceived of as a IR sensor capable of detecting the level of grey of the floor. It returns 0 if the robot is on white floor, 0.5 if is on light grey floor, 0.75 if is on dark grey floor, and 1 if is on black 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 in between $[0, 1]$. The camera can detect coloured objects up to a distance of 60cm. 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.

2.2 Robot controllers and the Evolutionary Algorithm

The robot controller is composed of a continuous time recurrent neural network (CTRNN) of 10 sensor neurons, 6 inter-neurons, and 4 motor neurons (see [1]). The structure of the network is shown in Fig. 1b. The states of the motor neurons are used to control the speed of the left and right wheels as explained later. The sensor neurons are simply relay units. The states of inter and motor neurons are updated using the following equations:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{10} \omega_{ji} \sigma(gI_j + \beta_j) + \sum_{j=11}^{16} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{11, \dots, 16\}; \quad (1)$$

$$\Delta y_i = -y_i + \sum_{j=11}^{16} \omega_{ji} \sigma(y_j + \beta_j); \text{ for } i = \{17, \dots, 20\}; \quad (2)$$

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 $i = \{1, \dots, 10\}$ is the activation of the i^{th} sensor neuron (see Fig. 1b 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 $i = \{11, \dots, 16\}$, 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 $M^L = f_{17} - f_{18}$, 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, and equation 1 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. 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 70 best individuals of the old population. Each genotype is a

vector comprising 135 real values (120 connections, 6 decay constants, 8 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 with a probability of 0.3 and mutation. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.05. 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 $[0,1]$.

2.3 Evaluation and Fitness Function

At the beginning of each evaluation trial, the robots are placed in the nest, located symmetrically on the left and on the right of the light, at 1.8cm away from each other. Their controllers are reset. The initial relative orientation of the two robots is sufficiently described by a vector of two variables (α, β) , see Fig. 1a). A sample set of starting configuration is chosen such that $\alpha, \beta \in (0, \frac{2\pi}{5}, \frac{4\pi}{5}, \frac{6\pi}{5}, \frac{8\pi}{5})$. From these combinations, 10 have been removed because they are rotational duplicates. This leaves the set of 15 team starting relative orientations that have been used. 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 [3] 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 robot exceeds the arena limits (i.e., a circle of 120cm 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., twice the journey to any of the food sites and back to the nest).

The evaluation procedure for what concerns the clonal and aclonal runs substantially matches the one described in [7]. 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 A and position B, see Fig. 1a) comprising each configuration, for a total of $E = 60$ trials. Each one of an aclonal individual 60’s trails is undertaken with a different, randomly chosen, partner.

In each trial e , the team is rewarded by an evaluation function F_e which corresponds to the product of the following components: $F_e = \max(C_1^{role P} \times C_2^{role F}; C_1^{role F} \times C_2^{role P}) \times P' \times P''$. $C_r^{role P} \in [0, 1]$ rewards a robot $r = \{1, 2\}$ for staying in the nest; $C_r^{role F} \in [0, 4]$ rewards a robot r for travelling twice the distance from the nest and to any of the two food sites; the team collision penalty P' is inversely proportional to the number of collisions, with $P' = 1$ with no collisions, and $P' = 0$ with 10 collisions in a team; P'' is the team penalty for exceeding the arena’s limits, with $P'' = 1$ if none of the robots exceeds the limits, $P'' = 0.3$ otherwise. The average team evaluation score is $F = \frac{1}{E} \sum_{e=1}^E F_e$.

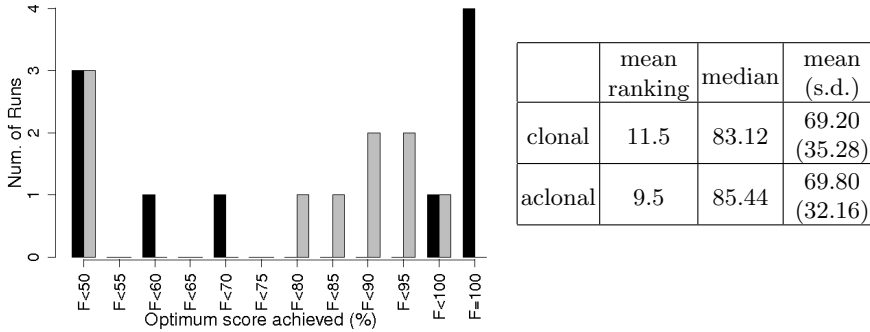


Fig. 2. The histogram shows the distributions of the highest average re-evaluation scores achieved by each run of the clonal (black bars) and aclonal approach (grey bars). Values represent percentage of the theoretical optimum average evaluation score F . The Table on the right shows, for clonal and aclonal approach, mean ranking, median, mean and standard deviation of the scores mentioned above.

3 Results

10 evolutionary runs, each using a different random initialisation, were carried out for each of the two approaches (i.e., clonal and aclonal). Each run lasted 2500 generations. Recall that our objective is to compare the performances of the clonal and aclonal approach for the evolution of homogeneous team of two robots capable of dynamically allocating and simultaneously executing *role P* (i.e., nest patrolling) and *role F* (i.e., foraging). Following the procedure illustrated in [7], 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 Sec. 2.3). In these tests, the 10 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 Sec. 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.

Contrary to what illustrated in [7], the results of our re-evaluations show no significant difference between the performances of homogeneous teams generated clonally and aclonally. 5 runs of the clonal approach produced high-scoring teams exceeding 95% of the optimal score, with 4 of them 100% successful (see Fig. 2, black bars). In contrast, only 1 out of 10 of the aclonal runs generated a homogeneous team that exceeds 95% of the optimal score, with no teams capable of completing the 60 re-evaluation trials with the highest score. The Table on the right of Fig. 2 shows a comparisons of mean and median scores and the mean ranking of both approaches. The difference between the two set of results is not significant (Mann-Whitney U test, $p > 0.1$);

From a statistical point of view, there is not enough evidence to prefer one approach over the other for the evolution of homogeneous multi-robot teams engaged in a dynamic task allocation scenario. From the point of view of generating

optimal controllers, the clonal approach does better than the aclonal one. The 4 out of 10 teams generated clonally, with a 100% success rate at the re-evaluation test (see Fig. 2, black bars), proved to be robust enough to deal with the variability concerning the robot relative starting positions and noise injected into the simulation. None of the teams generated aclonally showed a similar robustness. This evidence not only diverges from what shown in [7], but also questions the argument, put forward in Quinn’s paper, concerning the superiority of the aclonal approach. Recall that, according to Quinn, the aclonal approach takes advantage of the fact that roles can be evolved and refined prior to the evolution of any dynamical allocation mechanism. This is assumed not possible in the clonal approach for which the roles and the allocation mechanisms have to evolve simultaneously. Our results do not fit into the explanatory framework proposed by Quinn. To account for this divergence, we question the simultaneity argument. Should we always assume that in the clonal approach, behavioural roles and mechanisms to allocate them have to evolve simultaneously?

To answer this question, we analyse the evolutionary trajectory of the best clonal runs looking for the emergence of roles. In our scenario, the clonal approach generated optimal (in term of fitness) solutions by avoiding the simultaneous appearance of behavioural roles and role-allocation mechanisms. All the best clonal runs follow a very similar evolutionary trajectory. Fig. 3a shows the evolutionary history of the fitness components of the best genotypes of one of the best clonal run. The graph shows that early generation teams are very good in *role P* (see $\hat{C}^{role P}$ in Fig 3a), and very bad in *role F* (see $\hat{C}^{role F}$ in Fig 3a). After some generations, in which no progress can be observed, the fitness starts to increase owe to the evolution of the behavioural skills required to perform *role F*. By visual inspection of the team behaviour, we noticed that foraging behaviour appears first in a limited number of trials. These trials are those in which the robots, due to their initial relative orientations, experience different perceptual states which break the team symmetry and facilitate the role-allocation process. In following generations, we observe that robots perform foraging behaviour in a larger set of trials, and the role-allocation process becomes less dependent on the robots’ initial perceptual states. Eventually, the robots acquire the skills to negotiate their role regardless of their initial perceptual states and become capable of performing the task in all the experienced initial conditions.

In Sec. 4, we will come back to the issue of simultaneity with further speculations. In what remains of this section, we focus on another important aspect of the Quinn’s argument, the specialisation in heterogeneous teams generated aclonally. We run a second set of re-evaluation tests identical to the one used in [7] to investigate whether in the aclonal approach specialisation appears before the role allocation ability. In these tests, the 10 fittest genotypes from every generation of the best aclonal run are taken and retested in a heterogeneous setup. That is, each genotype is re-evaluated in combination with every other fittest genotypes of its generation, for 60 trials per combination. As in [7], for each generation, we selected the highest average score F among those produced by each combination, and we compared them with the highest average scores

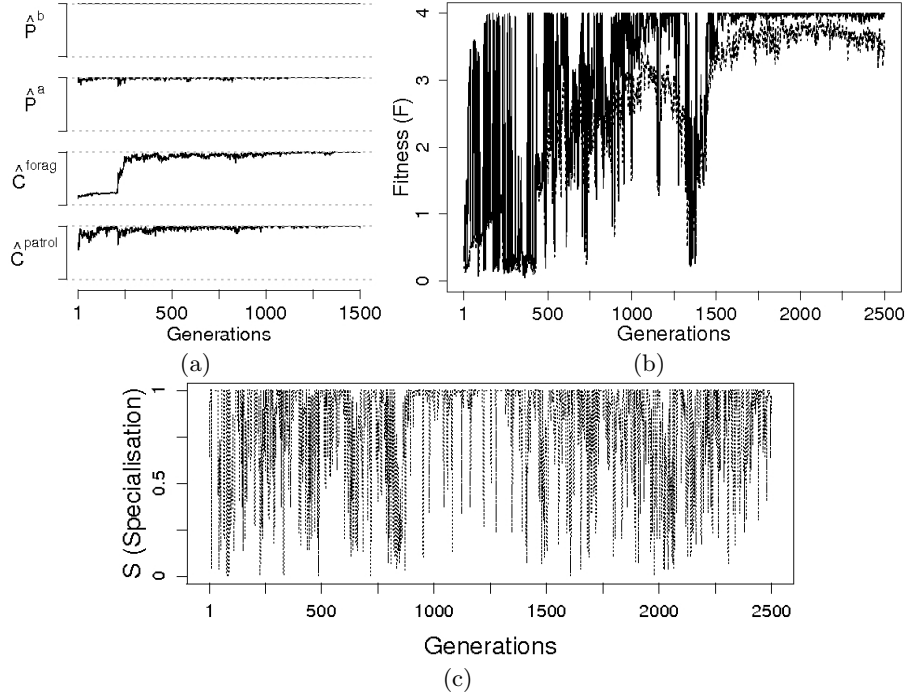


Fig. 3. a) The evolutionary history of the fitness components of the best genotypes of one of the best clonal run. At each generation, $\hat{C}^{\text{role } P}$, $\hat{C}^{\text{role } F}$, \hat{P}' , and \hat{P}'' are the mean values of $C_e^{\text{role } P}$, $C_e^{\text{role } F}$, P'_e , and P''_e , respectively, over 60 trials. b) Average evaluation score (F) of best heterogeneous combination of genotypes (continuous line), and best homogeneous team (dashed line), computed during a set of re-evaluation tests on the 10 fittest genotypes of the best aclonal run. c) Level of specialisation (S) in the best heterogeneous combination of genotypes that contributed to the fitness curve (continuous line) shown in (b).

F obtained by these genotypes when re-evaluated in a homogeneous setup (i.e., during the first set of re-evaluation tests). The continuous line in Fig 3b refers to the average re-evaluation score (F) of the best heterogeneous combinations; the dashed line refers to the average score of the best homogeneous team.

As in [7], we also have, for significant period of this aclonal run, a certain disparity between the fitness of heterogeneous and homogeneous teams. This disparity appears to be more prominent right after generation 1000, when the fitness of heterogeneous teams oscillates around the optimum, while the fitness of homogeneous teams is clearly decreasing. In [7], this disparity is accounted to by the existence of genetic specialisation in aclonally generated heterogeneous teams. The logic is rather simple. If robots are genetically predisposed to play either one or the other role, they get high fitness (F) when they are evaluated in combination with a partner that is specialised to play the complementary role, and low or lower fitness when they are evaluated with their clones. In [7], evo-

lutionary dynamics driven by the “specialisation factor” represent the basis of the superiority of the aclonal over the clonal approach. Contrary to clonal runs, aclonal runs can partition the task allocation scenario in two successive phases: a first one, in which evolution generates genetically specialised behavioural roles; a second one, in which evolution finds the way to move from specialised to generalist solutions (i.e., a single genotype that generates agents capable of dynamic allocation and execution of the roles of the task).

In order to deepen our understanding on the causal relationship between specialisation and the fitness disparity observed in Fig 3b, we measured the level of specialisation (S) of the best heterogeneous combinations that contributed to the generation of the fitness curve in Fig 3b, continuous line. In particular, we looked at the number of times, in each combination, the robots play each role during the 60 re-evaluation trials. The role that a robot plays in a trial is determined by how it contributes to the team fitness in that trial. For example, robot 1 plays *role P* and robot 2 plays *role F* if $C_0^{role P} \times C_1^{role F}$ is bigger than $C_0^{role F} \times C_1^{role P}$, and vice versa. For each combination, $S = \frac{|N^{role P} - N^{role F}|}{60}$, with $N^{role P}$ and $N^{role F}$ being the number of times in 60 trials in which one of the robot plays *role P* and *role F*, respectively. $S = 1$ means that robots are highly specialised; $S = 0$ no specialisation at all. Results are show in Fig 3c.

By comparing the graph in Fig 3b, continuous line, with the graph in Fig 3c, in particular focusing on the performances after generation 1500, we clearly see that not all the best combinations are highly specialised. While specialisation rises and falls, the best average score of these heterogeneous combinations only slightly fluctuates around the optimal score. This indicates that both specialisation and a certain level of generalist solutions are simultaneously present in the aclonal population. As far as it concerns the argument formulated in [7], it seems that the disparity between the scores of heterogeneous and homogeneous teams is not necessarily a sign of specialisation among the individuals of an aclonal population, as mentioned in [7]. To conclude, in our dynamic task allocation scenario, the aclonal population dynamics appear to be more articulated than what described in [7]. Specialisation did not turn out to be completely alternative and antecedent to generalist solutions. We believe that this evidence points at interesting evolutionary dynamics, not considered in [7], which suggests that the argument concerning the superiority of the aclonal versus the clonal approach in task-allocation scenario needs to be revisited.

4 Conclusions

This study aimed at testing the hypothesis that the results shown in [7], concerning the superiority of the aclonal versus the clonal approach, have been affected by task-specific features. We designed a similar task-allocation scenario which mainly differs from what shown in [7] for the nature of the behavioural roles. Basically, in our task the roles are clearly distinct, whereas in [7] the roles are rather similar. Contrary to what illustrated in [7], we found no significant difference between the aclonal and clonal approach. The analysis of the evolutionary

trajectories of clonal and aclonal runs produced evidence that did not support the reasoning put forward in [7], to account for his results. First, 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 clonal approach can rely on a broader range of evolutionary trajectories not necessarily limited by the simultaneity argument formulated in [7]. In the clonal approach, the fitness landscape can be successfully explored 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. Second, we showed that the author in [7] drew conclusions on the presence of role specialisation among the individuals generated aclonally from empirical evidence (i.e., the fitness disparity shown in Fig. 3b) that do not seem to be uniquely produced by genetically specialised individuals. We showed that in the aclonal approach, specialised and generalist solutions can live together in the same populations, with role allocation mechanisms evolving together with, and not necessarily after, genetic specialisation. In conclusion, the evolutionary dynamics of both clonal and aclonal approaches seem to be richer than what assumed in [7]. Although further investigation is required, in view of this, we speculate that the argument formulated in [7] applies to a rather restricted set of task-allocation scenarios, in which the selective pressures limit the possible successful evolutionary trajectories in the clonal approach.

References

1. R. D. Beer and J. C. Gallagher. Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122, 1992.
2. D. Floreano, S. Mitri, Magnenat S, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519, 2007.
3. N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6:325–368, 1997.
4. S. Nolfi. *EvoRob 1.1 User Manual*. Institute of Psychology, National Research Council (CNR), 2000. Available at <http://gral.ip.rm.cnr.it/evorobot/simulator.html>.
5. S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-organising Machine*. MIT Press, Cambridge, MA, 2001.
6. L. Parker. *Springer Handbook of Robotics*, chapter 40. Springer Verlag, Berlin, Germany, 2008.
7. M. Quinn. A comparison of approaches to the evolution of homogeneous multi/robot teams. In *Proceedings of the International Conference on Evolutionary Computation (CEC)*, volume 1, pages 128–135, 2001.
8. V. Trianni and S. Nolfi. Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial Life*, 17(3):183–202, 2011.
9. M. Waibel, D. Floreano, S. Magnenat, and L. Keller. 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, 2006.
10. M. Waibel, L. Keller, and D. Floreano. Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transaction of Evolutionary Computation*, 13(3):648–660, 2009.