

Different Genetic Algorithms and the Evolution of Specialization: A Study with Groups of Simulated Neural Robots

Tomassino Ferrauto^{**,†}

Consiglio Nazionale delle Ricerche
Università degli Studi dell'Aquila

Domenico Parisi^{**}

Consiglio Nazionale delle Ricerche

Gabriele Di Stefano[†]

Università degli Studi dell'Aquila

Gianluca Baldassarre^{*,**}

Consiglio Nazionale delle Ricerche

Abstract Organisms that live in groups, from microbial symbionts to social insects and schooling fish, exhibit a number of highly efficient cooperative behaviors, often based on role taking and specialization. These behaviors are relevant not only for the biologist but also for the engineer interested in decentralized collective robotics. We address these phenomena by carrying out experiments with groups of two simulated robots controlled by neural networks whose connection weights are evolved by using genetic algorithms. These algorithms and controllers are well suited to autonomously find solutions for decentralized collective robotic tasks based on principles of self-organization. The article first presents a taxonomy of role-taking and specialization mechanisms related to evolved neural network controllers. Then it introduces two cooperation tasks, which can be accomplished by either role taking or specialization, and uses these tasks to compare four different genetic algorithms to evaluate their capacity to evolve a suitable behavioral strategy, which depends on the task demands. Interestingly, only one of the four algorithms, which appears to have more biological plausibility, is capable of evolving role taking or specialization when they are needed. The results are relevant for both collective robotics and biology, as they can provide useful hints on the different processes that can lead to the emergence of specialization in robots and organisms.

Keywords

Cooperation, role taking and specialization, individual and group selection, temporary and stable groups, evolutionary regimes, navigation

1 Introduction

Cooperative behaviors are an important topic in autonomous robotics, which has received increasing attention in the last two decades ([8, 15, 22, 34, 38, 41, 49, 76]; see [20, 25, 72] for some reviews and taxonomies of multi-robot systems and the tasks that can be tackled with them). This research concerns multi-robot systems that tackle tasks that cannot be solved by single robots [12, 32, 35, 44, 70, 85] or can be solved more efficiently by multiple robots [34, 68], and it is often inspired by, or tries to

* Contact author.

** Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche, Via San Martino della Battaglia 44, 00185 Roma, Italy. E-mail: tomassino.ferrauto@istc.cnr.it (T.F.); domenico.parisi@istc.cnr.it (D.P.); gianluca.baldassarre@istc.cnr.it (G.B.)

† Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica, Università degli Studi dell'Aquila, Via Gronchi 18, 67100 L'Aquila, Italy. E-mail: gabriele.distefano@univaq.it

capture, the mechanisms underlying the highly efficient behaviors of social insects [2, 15, 39–41] and other animals acting in groups [17, 63, 67, 78]. The coordination of these organisms is based on interesting principles of self-organization [19, 26], which, if understood in depth and formalized [6, 11, 73], can often be translated into useful coordination principles [18, 41, 63] and robust collective robot controllers [9, 13, 23, 34].

This work focuses on multi-robot systems with *distributed controllers* [34, 38, 84], that is, on robot groups that do not rely upon leader robots [5, 76] or centralized control mechanisms [14, 21, 68, 85]. Rather, they base coordination on peer-to-peer interactions and self-organizing principles. Multi-robot systems with distributed control are of particular interest for autonomous robotics because, compared to systems with centralized control, they are usually more robust with respect to the failure of single robots, require little or no explicit communication [10, 15, 18, 34, 62], and allow the use of robots with simpler sensors and actuators [16, 23, 27].

The controllers presented here are *evolved* through genetic algorithms [23, 53, 62, 67, 71, 78]. As suggested by the framework of *evolutionary robotics* [52], multi-robot systems can greatly benefit from being automatically designed with evolutionary algorithms, as it is sometimes very difficult to design their controllers directly, due to the indirect and complex causal chains that link the behavior of the single robots to the behavior of the whole group, which is the ultimate target of the design (see [7] for a comparison and exploitation of synergies between evolutionary techniques and direct-design approaches). In this regard, evolutionary techniques have the potential of developing controllers that exploit the self-organizing properties of multi-robot systems in that they generate variants of controllers on the basis of random mutations and select a posteriori the best of the variants, based on the quality of the behavior exhibited by the robots in their interaction with the environment [51].

The robotic controllers used in this work are *neural networks*. The rationale of this choice is that, when one uses evolutionary techniques, such controllers have a high degree of *evolubility* [52]. Indeed, controllers based on *evolved neural networks* have been used with success in a number of works on collective robotics [13, 47, 62, 71], and this approach is now recognized as an important solution along with other approaches based on direct design, such as gradient fields [14] and behavior-based robotics [5].

The main goal of this article is to study the capacity of different genetic algorithms to evolve specialization when it is needed, and to evolve solutions based on generalist individuals when specialization is not advantageous (see [75] for a seminal work that compares different evolutionary regimes in relation to collective robotics). To this purpose, the article begins by discussing four different genetic algorithms (plus some variants), either taken from the literature of collective evolutionary robotics or inspired by biology. As we shall see, only one genetic algorithm among those considered is able to choose between specialization and solutions based on generalist individuals, depending on the demands of the task. This algorithm might correspond to an evolutionary process that actually exists and is quite diffused in nature.

In addition to demonstrating the potential for specialization of the specific algorithms considered, the research shows how various general features of these algorithms can influence the evolution of generalist or specialist solutions to collaborative problems. These features (Section 3.3 will present them in detail) are: (a) individual selection versus group selection; (b) encoding of the two individuals that make up a group in the same or in different genotypes; (c) the same or different populations for the two members of the group; (d) the temporary or permanent nature of the groups with respect to the life of each individual. The results give some indications on the influence that these features might more generally have on the emergence of specialization when one uses genetic algorithms different from those considered here.

The results might also provide useful hints to biologists concerning the conditions for the emergence of specialization, although at a very abstract level. Indeed, the similarity between the genetic algorithms presented here and some evolutionary regimes found in nature (see Section 3.3 for a first attempt to establish such a correspondence), and the analysis of their capacity to produce specialization, might give some useful suggestions concerning some issues relevant to biology (see Section 6).

The remainder of this article is organized as follows. Section 2 presents a taxonomy of the different mechanisms that might be used to implement role taking or specialization: the rest of the article will focus on two items of this taxonomy. Section 3 describes the two tasks, the fitness functions used to evolve the robots, the simulated robots used in the two tasks, and the four genetic algorithms and their possible biological counterparts. Section 4 analyzes the performance of the different genetic algorithms and their capacity to evolve role taking or specialization. Section 5 analyzes in detail the behavioral strategies evolved by the genetic algorithms. Finally, Section 6 discusses the results from a computational perspective and proposes some possible interpretations that might be suggestive for biologists in relation to the conditions for the emergence of specialization and its effects.

2 A Taxonomy of Specialization for Neural-Network Evolved Controllers

An important phenomenon studied within collective cooperative robotics is specialization [31, 72]. Specialization occurs when one or more cooperating robots focus on, and become expert in carrying out, a particular activity within the group. In adaptive multi-robot systems, specialization can take place in two conditions: (a) the accomplishment of the collective task requires that different robots engage in different activities because otherwise the task cannot be accomplished (this condition is widely studied in cooperative collective robotics); (b) the collective task does not strictly require specialization, but specialization can increase the efficiency of the group or it can simplify the robots' controllers. Specialization can decrease the complexity of controllers because each robot has to develop only a subset of the sensorimotor skills needed to accomplish the whole task.

An issue closely related to specialization is *role allocation*. This issue has received much attention within the literature of collective robotics (e.g., [28, 39, 42, 77]). It refers to the coordination mechanisms that a group of robots can exploit to allocate different roles to its members. The role assignment can be accomplished either on the basis of centralized mechanisms or in a decentralized fashion based on explicit communication or self-organizing mechanisms (as done by the robots of this study). Role taking is important for specialization because it might be a *precursor* of it, as discussed further below after the taxonomy. For this reason, we consider role taking as part of the taxonomy itself. The taxonomy includes three classes of items, each composed of some subclasses:

1. *Body specialization*: This specialization relies upon differences of the body. Body specialization can assume two forms:
 - (a) *Morphological specialization*: Robots behave differently because their bodies have permanent morphological differences in dimension, shape, or internal structure (e.g., see [31, 80]).
 - (b) *Body-state specialization*: Robots behave differently because some aspects of their morphologically similar bodies assume different states at different moments, due either to internal dynamics or to the effects of the dynamics of the external environment (e.g., a particular body configuration, or the level of energy).
2. *Controller specialization*: This specialization relies upon differences of the robots' controllers, namely differences of architecture or connection weights. Controller specialization can assume three forms:
 - (a) *Innate specialization*: Robots behave differently on the basis of controllers that have been differentiated since the beginning of the task accomplishment. This type of specialization can emerge when genetic techniques are used to evolve the robots' controllers, and different genomes, or different parts of genomes, are used to encode the architecture and/or the connection weights of the different controllers (some examples of this type of specialization will be presented here; see also [75]).

- (b) *Developmental specialization*: Robots gradually acquire different specialized behaviors and roles during a period of development in which the architecture of the controllers of the robots, and/or their connection weights, are shaped on the basis of both experience and the guidance of the genetic program (as an example of this, although not with a focus on collective robotics and specialization, see [66]).
 - (c) *Learned specialization*: The robots develop different behaviors and roles by suitably updating the connection weights of their controllers on the basis of learning algorithms and experience (see the reviews [55, 56]; these processes might also involve the dynamic update of simple but critical parameters, such as a threshold triggering a behavior; see, e.g., [69]). This specialization requires some form of long-term memory. Learned specialization differs from developmental specialization in that the former never involves changes of the architecture whereas the latter can do so, and in that the former only involves changes caused by experience whereas the latter involves changes caused by the interplay between experience and the genetic program.
3. *Role taking*: This solution relies upon differences in activation of the units of the different robots' controllers based on past and current input patterns. Role taking can assume two forms:
- (a) *Memory-based role taking*: The robots exhibit different behaviors on the basis of the different states assumed by the internal units of their controllers caused by recent percepts (see, e.g., [62]). Note that this type of role taking relies on dynamic/short-term memory mechanisms.
 - (b) *Reactive role taking*: Robots with identical controllers assume different roles, and maintain them in time, on the basis of the different input patterns perceived at each time step. Roles represent *behavioral attractors*, which are initially allocated on the basis of the different conditions in which the robots happen to be for accidental reasons (e.g., random initial conditions), and then are maintained on the basis of the dynamic interplay between the robots and the environment if this generates stable dynamic equilibria (some examples of this type of specialization will be presented here; also see [8]).

As mentioned above, role taking can be considered a precursor of specialization. Indeed, specialization requires that individuals face tasks that require them to accomplish different activities and hence to assume different roles. At the very beginning of the evolution of specialization, evolutionary processes work on individuals that are neither collaborative nor specialized. In these conditions, the exploitation of collaborative opportunities and the specialized abilities have to coevolve. In this respect, the capacity of individuals to play different roles, based on memory or behavioral attractors, can ease the discovery of collaborative activities requiring different roles, and so create the environmental pressure for the emergence of *stable* forms of role-playing based on control specialization or body specialization if this can produce further fitness advantages.

Another observation concerns the distinction between learned specialization and memory-based role taking. From a computational perspective this distinction is blurred, in the sense that there might be a continuum between the two. Indeed “connection weights” and “activations of units” of neural networks are both “memory variables” that can play similar roles (cf. [83]). The difference between the two is in terms of the time scales involved, and so it is along a continuum: in the case of learned specialization, the differences between individuals arise from prolonged experiences that change the connection weights of their controllers during relatively long time scales; in the case of memory-based role taking, they arise from recent percepts that change the dynamic activation of the units of their controllers.

A last observation concerns role taking. In [8] role taking (in particular, reactive role taking) was called “situated specialization” to emphasize the continuum existing between role taking and controller specialization. Here, however, we use the expression “role taking” to better reflect the use of the term “specialization” in the literature, where it is usually used to refer to body or controller specialization.

3 The Experimental Setup

3.1 The Tasks Used to Test the Algorithms

This article uses two tasks to investigate the capacity of the genetic algorithms considered here to produce solutions based on generalist individuals or specialization. The first task is a *light-approaching task*, initially proposed in [8], where a target light periodically switches its position. This task requires that the robots repeatedly switch roles during its accomplishment, and so it can be best performed with role taking (namely, reactive role taking, as the neural networks used here do not have a dynamic memory; see Section 3.2). The second is a *coordinated motion task*, initially proposed in [61], where robots do not perceive any reference landmark in the environment: this task requires a difficult initial allocation of roles between the robots, and so it can be best performed with specialization (namely, innate specialization, as the neural networks used here do not learn or develop during life, but acquire their connection weights based on genetic algorithms; see Section 3.3).

The two tasks were chosen because they have some features that make them suitable for investigating the issues targeted here. In particular, the coordinated motion task can be solved best with specialization, although it can also be solved with role taking at the cost of lower fitness. This generates a certain evolutionary pressure for the emergence of specialization, and so the task allows checking whether a genetic algorithm is capable of producing specialization when necessary. On the contrary, the light-approaching task can be solved best with role taking, although it can also be solved with specialization, obtaining a smaller fitness. For this reason, the task can be used to determine if a genetic algorithm is capable of developing a solution based on generalist individuals. This was particularly important because, as we shall see, some algorithms tend to *force* specialization even when it is not advantageous.

3.1.1 Light-Approaching Task

In the light-approaching task the robots had to approach a light target while staying close to each other. The environment used for this purpose was a simulated rectangular arena of 1×2 m surrounded by walls (Figure 1). The light target was formed by two halogen light bulbs of 230 W located in the middle of the west and east (shorter) walls of the arena at a height of 1.5 cm from the ground. The lights turned on in an alternating fashion (see below). The robots could perceive the lights through infrared sensors used in passive mode, and they could perceive each other through those sensors used in active mode and through directional microphones (see Section 3.2 for the explanation of how these sensors were simulated).




Figure 1. The arena with two robots (two circles) and two light bulbs (light and gray semicircles: only one light was on at each time). The radial segment on each robot marks its front (the side with six sensors; see text) and indicates the direction of movement of the wheels in response to positive motion commands.

Each group of robots was tested for four epochs, which lasted 1,500 time steps each (each time step corresponded to 100 ms of real time). At each time step, the robots' sensors sent signals to the controllers, the controllers calculated the activation of the output units, this activation was used to issue the motion commands to the wheel motors, and finally the robots moved. At the beginning of each epoch the robots of a group were randomly placed in the arena with random orientations, and only the light bulb on the east wall of the arena was turned on. Then robots were left free to move, and when the barycenter of the group reached a distance smaller than 300 mm from the light currently on, this was turned off and the other light was turned on.

To reward robots able to stay close to each other and at the same time move toward the light turned on, a fitness function formed by two components was used: a *group compactness component* (GCC) and a *group speed component* (GSC) (cf. [8]). GCC was computed at each step t as follows:

$$GCC_t = \max \left[1 - \frac{D}{600 \text{ mm}}, 0 \right]$$

where D was the distance, measured in millimeters, between the centers of the two robots. According to this formula, GCC is close to one for very compact groups, decreases linearly for more dispersed groups, and is equal to zero for robots that are more than 600 mm apart.

GSC was computed at each time step t as follows:

$$GSC_t = -\frac{1}{2} \left(\frac{\Delta GD_t - 7 \text{ mm}}{7 \text{ mm}} \right)$$

where ΔGD_t was the variation of the distance between the group's barycenter and the light target, and 7 mm corresponded to the maximum displacement in space that a robot could accomplish in one time step. Given this formula, GSC ranged in $[0, 1]$ and was less than 0.5 if the group moved away from the light, 0.5 if it was still, and greater than 0.5 if it moved toward the light.

The total fitness F of a each group was computed as the average of the two components over the duration of all the epochs:

$$F = \frac{1}{E T} \sum_{e=1}^E \sum_{t=1}^T \left(\frac{1}{2} GCC_{et} + \frac{1}{2} GSC_{et} \right)$$

where E was the total number of epochs, and T was the total number of time steps for each epoch. As explained in Section 3.3, in the case of the algorithms using temporary grouping the fitness of single robots was computed based on epochs involving different companions.

3.1.2 Coordinated Motion Task

The coordinated motion task used an empty open arena (without light targets or walls). The robots had only eight infrared sensors to perceive each other, and no light sensors (see Section 3.2). The robots had to move as far as possible in any direction from the initial position while staying close to each other. At the beginning of each epoch, the robots were assigned random orientations and were set at a distance of 15 mm from each other. This distance ensured that each robot could detect its companion through infrared sensors with a limited sensing range (see Section 3.2). Each group was tested for 40 epochs, which lasted 150 steps each. The number of epochs had to be high, in comparison with the light-approaching task (four epochs), because the initial random orientation of the robots could cause very different levels of performance.

Similarly to the light-approaching task, the fitness was formed by two components: a group compactness component (GCC) and a group speed component (GSC). The group compactness component was computed at each time step t as follows:

$$GCC_t = \max \left[1 - \frac{D}{300 \text{ mm}}, 0 \right]$$

Note that this formula is more demanding, with respect to group compactness, than the analogous formula used for the light-approaching task (300 vs. 600 mm): this was necessary to evolve robots capable of staying within the small infrared perceptual range (about 45 mm; see Section 3.2). In this respect, note that the parameter 300 mm is used to set the slope of the decrease of the fitness function with respect to the distance between the robots, but the sensory contact between the robots is lost well before 300 mm.

The group speed component was computed at each time step as follows:

$$GSC_t = \frac{\Delta GD_t}{7 \text{ mm}}$$

where ΔGD_t was the variation of the distance between the group's barycenter and the starting point at time step t , and 7 mm was a constant equal to the maximum distance a robot could cover in one time step. As in the light-approaching task, the total fitness F of a group was computed as the average of the two components over the duration of all epochs.

3.1.3 The Challenges Posed by the Two Tasks

As mentioned in Section 3.1, the two tasks were chosen because they require either role taking or specialization. In particular, the light-approaching task requires role taking, as the switching of roles can aid the groups' behavior in response to the light switches. On the other side, the coordinated motion task requires specialization to perform the difficult initial allocation of roles.

Before closing this section, it is important to note that the two tasks, although they might appear simple at first sight, are rather difficult to perform in view of the very limited local sensors of the robots (see Section 3.2). So, for example, the light-approaching task requires each robot to perform a fine dynamical regulation of the distance to the other robot and at the same time to move toward the light, which sometimes is shadowed by the companion (on the difficulty of the task, see also [8], where the task was proposed). Similarly, the coordinated motion task requires each robot to keep at a precise distance from the companion, to avoid losing it and at the same time to avoid crashing into it, on the basis of infrared sensors that have a very short sensing range (see Section 3.2), and while moving in a common direction in coordination with the companion without the possibility of relying upon external landmarks (on the difficulty of this task, see also [61] and [62], where the task was proposed). The thorough analyses of the behaviors evolved by the genetic algorithms, presented in detail in Sections 5.1 and 5.2, will show the difficulties produced by these conditions.

3.2 The Robots and the Neural Network Controller

Two simulated Khepera robots were used in all experiments (Figure 2; [48]). The robots were equipped with: (a) eight infrared sensors used to detect the presence of walls and other robots, up to a distance of 45 mm, in both tasks; (b) four light sensors used to detect the light target, up to 4 m, in the light-approaching task (these "sensors" were actually obtained using infrared sensors in passive mode, and by averaging the activations of pairs of them); (c) four directional microphones used to detect the position of other robots in the light-approaching task (these sensors, not present on the real robots, were simulated as described below). Each robot had a diameter of 55 mm and two independently controllable




Figure 2. (a) The Khepera robot. (b) Scheme of the robot's sensors and actuators. A dotted line connecting a light sensor with two infrared sensors indicates that the output of that light sensor was obtained as the average of the activations of the two infrared sensors used in passive mode.

motors connected to the two wheels and able to regulate the speed and direction of motion. In the light-approaching task the robots also had a simulated loudspeaker, which continuously emitted a sound having a fixed amplitude and frequency, used to signal the own position to the other robot (this actuator, not present in the real robots, was simulated as described below). Note that the microphones and loudspeakers allowed robots to detect each other at a greater distance than did the infrared sensors.

To achieve a greater level of realism of the simulations, the data obtained from a sampling procedure carried out with the real robot were used to compute the activation of the infrared sensors and of the light sensors (see [47, 52]). This sampling procedure set a physical robot at 180 different angles and 20 different distances with respect to walls, or another robot, or a halogen light bulb, and measured the activation of the infrared sensors of the robot (used in active mode to detect the wall or the other robot, and in passive mode to detect the light). A geometrical simulation of shadows was also implemented to have a more realistic activation of light sensors in the light-approaching task. In order to simulate the effects of different activations of the motors, the change of orientation and displacement in space of a physical robot was sampled in response to different commands issued to the motors (the sampled data relative to the infrared sensors activated by walls and other robots, and the sampled data relative to motors, were those used in the simulator Evorobot [52]; the sampled data relative to the infrared sensors activated by the halogen light were those used in [12]).

As the physical robots were not endowed with directional microphones and loudspeakers, these were simulated as in [12]. In particular, the sound amplitude A , decreasing with increasing distance D from the source, was computed as follows:

$$A = \frac{1}{1 + \left(\frac{D}{1000 \text{ mm}}\right)^2} AF$$

where AF is an attenuation factor that simulates the effects of the microphone orientation with respect to the sound source:

$$AF = 1 - 0.9 \frac{\alpha}{180^\circ}$$

where α is the convex angle, measured in degrees, between the direction pointed to by the microphone and the direction of the sound source. The sound amplitude actually perceived by the microphone, denoted by PA , was obtained by passing A through a logistic function to simulate the saturation of sensors:

$$PA = \frac{2}{1 + e^{-A}} - 1$$

In the light-approaching task, each robot was controlled by a neural network (Figure 3) having 16 input units, each corresponding to a particular sensor of the robot (eight for the active infrared sensors, four for the light ones, and four for the microphones), plus a bias unit. These units were directly connected to the two output units, for a total of 34 connection weights ($= [8 + 4 + 4 + 1] \times 2$). The activation of these two output units was mapped onto the desired speed of the two wheels, which ranged over ± 7 mm in linear displacement. In the coordinated motion task, the robots used only the infrared sensors used in active mode, and so the controllers had only 18 connection weights ($= [8 + 1] \times 2$).

3.3 The Four Genetic Algorithms and Their Biological Counterparts

This section explains the four genetic algorithms tested here. Two algorithms had two variants each. The algorithms had various features in common but differed in a few important respects. The most important feature they had in common was the equal sharing of the payoff between the robots of each group. This is a fundamental assumption that eliminates the possibility of *cheating*, that is, the possibility that some members of the group perform actions directed to gain higher benefits or pay lower costs than those of other members. This is a strong simplification with respect to biology, where the problem of cheating is one of the major obstacles to the emergence of cooperative behavior [37, 46]. The justification for the introduction of this assumption is that this research aimed to focus on the issue of specialization and so avoided tackling the issue of cheating, which would have complicated the interpretation of the results on specialization.

In all genetic algorithms, the genotype directly encoded the connection weights of the neural network controllers of the robots in the ways explained below. For each algorithm, the evolution of controllers was replicated 10 times, starting with different initial genotype populations. In each replication the evolution lasted 600 generations. Each robot group corresponded to a certain genotype (or to two genotypes; see Section 3.3) and was tested four times (*epochs*; each epoch lasted 1,500 time steps) in the

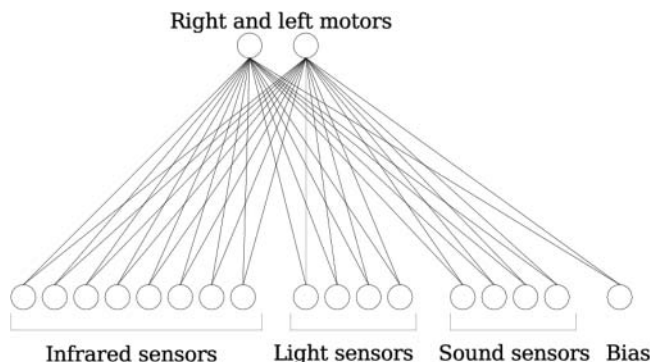


Figure 3. The neural-network controller used in the light-approaching task. The controller used in the coordinated motion task differed from this controller in that it did not have the input units corresponding to the light sensors and the sound sensors.

study of the light-approaching task, and 40 times (each lasting 150 time steps) in the study related to the coordinated motion task. Each group was formed by two robots. The best 20 genotypes of each generation of the population (40 in the single-population genetic algorithms; see below) were reproduced by generating five copies each. During reproduction, each connection weight encoded in the genotype was mutated with a probability of 5% by adding it to a random real number uniformly drawn from the interval $[-10, +10]$.

Table 1 summarizes the features of the genetic algorithms (note that the third and fourth genetic algorithms have two variants each). The names assigned to the algorithms are supposed to capture their most important features. As shown in the table, the algorithms differ along four major dimensions:

1. *Unit of selection.* This indicates the element copied in the duplication phase of the genetic algorithm. The unit of selection could be either the single controller of robots (standard selection) or the controllers of both robots of a group (group selection).
2. *Genotype information.* The genotype that was reproduced could encode either the connection weights of the controller of a single robot or the weights of the controllers of all robots of a group.

Table 1. List of the four genetic algorithms and their variants considered in this article (first column), and summary of their properties (other columns), in particular: the unit of selection (individual vs. group); the information encoded in the genotype (information related to only one robot controller vs. information related to two different controllers); the number of populations (one vs. two); the companion used in the tests (always the same companion vs. a varying companion). Note that the third and fourth genetic algorithms had two variants each. The last column indicates examples of biological systems with evolutionary regimes that share some features with the genetic algorithms (see text for details).

The four genetic algorithms	Unit of selection		Genotype information		Populations		Companion in tests		Biological counterparts
	Indiv.	Group	Indiv.	Group	One	Two	Same	Diff.	
	1. Homogeneous group selection		✓	✓		✓		✓	
2. Heterogeneous group selection		✓		✓	✓		✓		Cells of multicellular organisms (e.g., epidermal cells, red skeletal muscle cells, and neurons)
3.1 Multiple population, stable grouping	✓		✓			✓	✓		Endosymbionts (e.g., single-celled algae inside reef-building corals)
3.2 Multiple population, temporary grouping	✓		✓			✓		✓	Ectosymbionts (e.g., cleaner fish)
4.1 Single population, stable grouping	✓		✓		✓		✓		Animal demographic societies (e.g., cooperatively breeding birds, social spider colonies). Slime molds (plasmodial and cellular).
4.2 Single population, temporary grouping	✓		✓		✓			✓	Animal casual societies (e.g., monkey feeding troops, bird flocks, fish schools)

3. *Population.* The genetic algorithms could use either one population of genotypes, in which case both controllers of the robots forming a group were drawn from a single population of genotypes, or as many distinct populations of genotypes as the number of members of each group (two in the experiments presented here), in which case the controllers of robots forming a group were formed with genotypes drawn from different populations, which were never mixed.
4. *Companion in tests.* Each robot could be tested either with always the same (randomly selected) companion robot in all the epochs of the test, or with a (randomly selected) different companion in each epoch.

The four genetic algorithms and their variants considered here do not cover all the possible combinations of these features. Rather, a subset of the possible algorithms was chosen on the basis of presence in the literature on collective artificial intelligence or correspondence to existing known biological evolutionary regimes.

The details of the four genetic algorithms are now presented, together with some examples of their possible biological counterparts. Regarding the latter, note that real biological systems are often based on mixtures of regimes, and in this respect the genetic algorithms proposed here should be considered as abstract stylizations of reality rather than detailed models of it. This was done because the consideration of “pure cases” could ease the isolation of the principles that underlie the emergence of specialization in real and artificial systems. The first two genetic algorithms are characterized by group selection, the other two by individual selection. The algorithms have the following features:

1. *Homogeneous group selection.* This genetic algorithm is based on a population of 100 genotypes. Each genotype encodes the connection weights of a neural network controller that is identically copied into the two robots forming a group (as these robots have the same body and controller, they are sort of *clones*). The unit of selection is hence the group, and the fitness obtained by the group is shared equally between the two robots. Each robot is tested with the same companion in all epochs. We now illustrate the reasons why this algorithm was chosen, from the point of view of computational models and biology.

Computational models. This algorithm was chosen because it had been used with success in various previous works on evolutionary collective robotics [8, 9, 12, 59, 62, 70, 71] and so it was interesting to investigate its capabilities to produce turn taking or specialization. Indeed, given that in this algorithm the robots are clones, the only type of solution it could produce to perform the tasks was role taking. For this reason we expected that it would have led to high-performing groups in the task requiring this solution, but not in the task requiring specialization.

Biology. This genetic algorithm is also interesting from a biological perspective. The algorithm might be considered a stylization of the evolutionary regime operating in social insects, as ants and bees, where individuals descending from the same queen have very similar genotypes and compete against other groups descending from different queens ([81]; note that here we ignore the possible contextual presence, rather common in social insects, of morphological, developmental, and learned specialization). Animal communities formed by highly cooperating individuals, in particular social-insect societies, are often associated with this evolutionary regime [36], probably because it contributes to solving the cheating problem mentioned above. The reason is that it can be considered an extreme form of *kin-selected altruism*. This refers to situations where the organisms of a group are close relatives and so have very similar genotypes. This is a strong incentive to cooperate, as “altruistic” actions help to propagate the own genome [46]. From the evolutionary point of view, this regime implies that a group of such organisms tends to form a *superorganism* [65] acting as a single unit of selection.

2. *Heterogeneous group selection.* This genetic algorithm is based on a population of 100 genotypes. Each genotype encodes the two different sets of connection weights of the two neural network controllers of the robots forming a group, so the group is the unit of selection of the algorithm. Each robot is tested with the same companion in all epochs.

Computational models. From a computational perspective, this algorithm was chosen for two reasons. The first is that we wanted to verify two expectations related to it. The first expectation was that the algorithm had a strong tendency to produce specialization and so to be highly performing in tasks requiring differentiated behaviors. This expectation was based on the fact that the algorithm encodes the controllers of the robots forming a group with different parts of the genotype, so it was expected to be able to produce robots with different, complementary behaviors. We also had a second expectation, namely that the algorithm would perform poorly on tasks requiring generalists, as it would have had difficulty in developing robots with similar or identical behaviors. The second reason was that this algorithm had been used with some success in some collective-robotics works (e.g., [59]), so we wanted to compare its performance with that of the other algorithms.

Biology. This genetic algorithm might be thought of as a stylization of genetic mechanisms operating in multicellular organisms. The genotype contained in the cells of these organisms is formed by different parts that are expressed in cells having different specialization, for example, forming different parts of the body. For instance, the human body has about 210 different types of cells (e.g., epidermal cells, red skeletal muscle cells, and neurons; see the Web site http://en.wikipedia.org/wiki/Human_cell for a whole list): these are generated by the same genome but differentiate on the basis of their position within the body, the interaction with the external environment, and so on [82].

3. *Multiple populations.* This genetic algorithm is based on two populations, each corresponding to one of the two robots forming a group. Each population is formed by 100 genotypes. Each genotype encodes the set of connection weights of the neural network controller of one of the two robots of a group, and the unit of selection is the single robot. The groups are formed by randomly choosing the controller of the first robot from the first population and randomly choosing the controller of the second robot from the second population (the selection is conducted in such a way that each robot is tested the same number of times). Two variants of the algorithm were tested:

- (a) *Stable grouping.* The formation of groups is accomplished only once for all the epochs of the fitness test, at the beginning of the test itself (but of course groups are generated again for each generation).
- (b) *Temporary grouping.* This formation of groups is accomplished at the beginning of each epoch, so that each robot is tested with several different companions. In this second variant of the algorithm, the robots of a group share the group's fitness obtained in each epoch in two equal parts, so each robot accumulates the fitness gathered with different companions during life.

Computational models. From a computational perspective, the temporary grouping version of this algorithm is relevant as some authors have claimed that it has a high potential to lead to the emergence of specialization [29]. The reason is that in this algorithm the members of groups are selected from different populations and the different populations can develop individuals with complementary specialized abilities while avoiding the risk that individuals with the same capabilities will be grouped together. We also tested the stable grouping version of this algorithm to evaluate whether this feature affected the results.

Biology. From a biological perspective, the algorithm might correspond to *symbiosis* [1, 50]. In this case, organisms belonging to two different species coevolve to form tightly linked systems within which the survival of one species is highly dependent on the survival of the other species—for example, because they specialize in producing different nutritional substances that are then suitably exchanged (note that the forms of symbiosis relevant here are those where the organisms receive mutual advantage from the interaction). The stable grouping variant of the algorithm might be related to *endosymbiosis*, where one organism lives inside the body of another organism and so can create with it a stable collaboration forming a whole system (to the extent that the endosymbiont or the host might not survive without the other, giving rise in this case to *obligate symbiosis* such as that arising between some bacteria and algae [24]). The temporary grouping variant of the algorithm might be related to *ectosymbiosis*, where one organism lives on the surface of the body of another organism. This relationship can give rise to a collaboration that varies during the life of the two organisms (e.g., as in some cases of the clownfish–sea-anemone symbiosis). Symbiosis, and especially endosymbiosis, is increasingly recognized as a main drive of evolution ([64, 79]; cf. Section 6).

4. *Single population.* This genetic algorithm is based on a population of 200 genotypes. Each genotype encodes the connection weights of the neural network controller of one robot, and the unit of selection is the single robot. Groups are formed by randomly choosing pairs of robots from the population (the selection is accomplished in such a way that each robot is tested the same number of times). Two variants of the algorithm were tested:

- (a) *Stable grouping.* For each generation, the formation of groups is accomplished only once for all the epochs composing the fitness test.
- (b) *Temporary grouping.* The formation of groups was accomplished at the beginning of each epoch. As for the multiple-population algorithm, in this second variant the robots of a group share the groups' fitness of each epoch in two equal parts.

Computational models. The questions investigated in relation to this genetic algorithm were whether and how it had the capacity to lead to the emergence of specialization. We expected that this algorithm would have difficulties in developing specialization because, being individuals chosen randomly, specialized individuals would have had a high risk of getting grouped with individuals with the same type of specialization so having a fitness lower than non-specialized individuals. We shall see that in this respect the results were instead quite surprising.

Biology. This algorithm was chosen because it reflects an evolutionary regime that is very common in nature. In particular it corresponds to cases where individuals of the same species have different genotypes but have some advantages in performing some cooperating tasks. The permanent and temporary versions of this algorithm, both based on single populations, have a parallel in respectively the *demographic societies* and the *casual societies* formed by social animals and defined and studied in [81]. According to that author, a demographic society is “a society that is stable enough through time, usually owing to its being relatively closed to newcomers, for the demographic processes of birth and death to play a significant role in its composition.” Examples of demographic societies are groups of cooperatively breeding birds (e.g., [33]), social spider colonies [3], female-based coati bands [30], and some instances of cooperative hunting [54]. On the contrary, a *casual society*, or *casual group*, is “a temporary group formed by individuals within a society. The casual society is unstable, being open to new members and losing old ones at a high rate. Examples include feeding groups of monkeys within a troop and groups

of playing children” (see [17], on monkeys feeding in a group). Other examples are bird flocks [57, 60] and fish schools [58]. The stable version of the regime might also have a counterpart within unicellular organisms. In particular, *cellular slime molds* [43, 74], such as *Dictyostelium discoideum*, are formed by many individual amoebae capable of aggregating so as to form whole systems that move and assume complex structures within which the single individuals might strongly specialize (e.g., front and rear individuals to perform the movement). Similarly, *plasmodial slime molds*, such as *Physarum*, are amoebae capable of merging their cytoplasm to form a whole *syncytium* having a multitude of nuclei and capable of moving [45]. The wide dispersion of this evolutionary regime in nature increases the relevance of the unexpected results obtained by using it.

4 Different Genetic Algorithms and the Emergence of Specialization

This section compares the different levels of performance in the two tasks achieved by all the genetic algorithms illustrated in Section 3.3 and analyzes their capacity to lead to the emergence of role taking or specialization depending on the opportunities offered by the tasks. For this purpose, Table 2 summarizes the type of strategy developed by the different genetic algorithms when they are used to tackle the two tasks. In particular, the table reports the strategy used by the best evolutionary run of each algorithm (i.e., either role taking or specialization), detected on the basis of direct inspection and the techniques illustrated in detail in Section 5. The table also reports the levels of performance achieved by the algorithms. The information reported in the table is complemented by Figures 4 and 5, which allow visual comparison of the performance of all the four genetic algorithms and their variants. The performance of the four genetic algorithms and the strategy they developed are now analyzed in detail.

Table 2. Type of strategy (role taking vs. specialization) and performance (average and SD) of the four genetic algorithms and their variants in the light-approaching task and the coordinated motion task. Data are related to the average performance in a test of 100 epochs of the best group that emerged in the best evolutionary run of each algorithm.

Type of genetic algorithm	Light-approaching task				Coordinated motion task			
	Specialization		Fitness		Specialization		Fitness	
	Role taking	Spec.	Average	SD	Role taking	Spec.	Average	SD
1. Homogeneous group selection	✓		0.875176	0.003089	✓		0.891230	0.020010
2. Heterogeneous group selection		✓	0.844204	0.002297		✓	0.948215	0.002420
3.1 Multiple population, stable grouping		✓	0.854000	0.005853		✓	0.952055	0.003636
3.2 Multiple population, temporary grouping		✓	0.852918	0.004057		✓	0.947073	0.002650
4.1 Single population, stable grouping	✓		0.865800	0.001821		✓	0.946535	0.003686
4.2 Single population, temporary grouping	✓		0.862780	0.006053	✓		0.892165	0.008137




Figure 4. Performance and SD in the light-approaching task (y axis) of the four genetic algorithms and their variants (indicated by the numbers on the x axis, for which see Table 1). “Average over 10 runs”: performance of the best individuals of the last generation of 10 evolutionary runs, averaged over 100 epochs of test per individual and over the 10 evolutionary runs per algorithm. “Best run”: performance of the best individuals of the last generation of the best evolutionary run (out of 10) of each genetic algorithm, averaged over 100 epochs of test per individual. Light-gray and dark-gray bars indicate the algorithms or variants that developed role taking and specialization, respectively.

4.1 Genetic Algorithms Based on Homogeneous and Heterogeneous Group Selection

Homogeneous group selection can produce only role taking, as all the robots have the same controller. In contrast, heterogeneous group selection has a strong tendency to produce specialization. The reason is that in this case role taking does not tend to emerge as it would require different parts of the genotype encoding the controllers of the different group members to develop functionally equivalent controllers, a condition that evolution finds difficult to produce. For these reasons, as




Figure 5. Performance and SD in the coordinated motion task (y axis) of the four genetic algorithms and their variants (numbers on the x axis). Data collected, processed, and plotted as for Figure 4.

shown in Figures 4 and 5, homogeneous group selection always developed role taking, even in the coordinated motion task where this was detrimental. On the other hand, heterogeneous group selection always developed specialization, even in the light-approaching task where this led to a lower fitness. Section 5 illustrates in detail the behaviors found by these strategies.

These features of the two algorithms explain the differences in performance exhibited in the two tasks and reported in Table 2. In the light-approaching task, where role taking is more advantageous because it allows quick switching of roles (Section 5), homogeneous group selection outperforms heterogeneous group selection, as indicated by the fact that the performance of the best individual of the best evolutionary runs of the two algorithms is statistically different in a test of 100 epochs (t -test, $p < 0.001$). But in the coordinated motion task, where specialization facilitates the initial role allocation (Section 5), heterogeneous group selection outperforms homogeneous group selection, as indicated by the fact that the performance of the best individual of the best evolutionary runs of the two algorithms is statistically different in a test of 100 epochs (t -test, $p < 0.001$).

4.2 Genetic Algorithm Based on Multiple Populations with Stable and Temporary Grouping Variants: The Differentiation-Bias Problem

The outcome of the tests run with the groups evolved with the multiple-population genetic algorithm proposed in [29] showed that this type of algorithm is affected by what might be called the *differentiation-bias problem*, consisting of a tendency of the algorithm to enforce the emergence of specialization in any circumstance, even when it is not needed by the task, both with stable and with temporary grouping. Indeed, direct inspection of the strategies developed by the genetic algorithm under the two grouping conditions in the two tasks revealed that it always led to the emergence of behavioral strategies typically based on specialization (Section 5).

In order to quantitatively detect the presence of specialization in both tasks, and with both variants of the genetic algorithm, we measured the correlation between the connection weights of the genotypes within each of the two populations and between the two populations. In all the four simulations related to the two tasks and the two variants of the genetic algorithm, this measure revealed the presence of a strong correlation between all the possible pairs of genotypes within each population and a close-to-null correlation between all the possible pairs of genotypes belonging to the two populations. In particular, in the four simulations the within-population and between-population measures were as follows: (a) light-approaching task, temporary grouping: 0.9728 and 0.0691; (b) light-approaching task, stable grouping: 0.9877 and 0.2751; (c) coordinated motion task, temporary grouping: 0.9808 and 0.1717; (d) coordinated motion task, stable grouping: 0.9587 and -0.1449 .

This strong tendency to lead to the emergence of specialization impaired performance in the light-approaching task, where role taking was needed to switch roles on the fly. In particular (see Table 2 and Figures 4 and 5), in the light-approaching task both variants of the multiple-population algorithm were outperformed by the homogeneous group selection algorithm, as shown by the statistical comparison of the performance of their best runs (t -test, $p < 0.001$ for both comparisons). On the contrary, in the coordinated motion task the performance of the best run of two variants of the multiple-population algorithm slightly outperformed (stable grouping, 0.952 on average vs. 0.948) or was very similar to (temporary grouping, 0.947 vs. 0.948) that of the best run of the heterogeneous group selection algorithm, and both algorithms correctly produced specialization.

4.3 Genetic Algorithm Based on Single Populations and Stable and Temporary Grouping Variants: The Same-Specialization Problem

The *temporary grouping variant* of the single-population genetic algorithm was capable of developing role taking, as shown by direct observation of behavior (see Section 5) and the fact that its performance is statistically higher than that of the heterogeneous algorithm and the multiple-population algorithms (t -test, $p < 0.001$ for all comparisons). However, the algorithm was not capable of developing specialization when needed. In this respect, in the coordinated motion task the performance of the algorithm

was statistically lower than that of the heterogeneous algorithm, the multiple-population algorithms, and also the stable grouping variant of the algorithm (t -test, $p < 0.001$ for all comparisons).

A possible explanation of this is that this evolutionary regime incurs what might be called the *same-specialization problem*, which arises if (too many) individuals with the same specialization are grouped together and so the whole group cannot accomplish some functions. In the case considered here, individuals that tend to develop a specialized behavior have a 50% probability of being grouped with a companion with the same type of specialization, in which case they get a very low fitness. This risk increases exponentially with the number of epochs that form the fitness test, as at each epoch each individual is grouped with a new companion (the chance that at least one bad grouping takes place is $1 - (1/2)^n$ with n epochs). As a consequence, non-specialized generalist robots tend to diffuse in the population, as they do not incur such costs.

To show that this latter event actually took place in the simulations, Figure 6 shows the correlation existing between all possible pairs of genotypes of the population. To visually highlight the similarities between genotypes, we created a 200×200 matrix, storing in each cell the correlation between a pair of genotypes: The x axis and the y axis report the different genotypes, whereas the z axis reports the correlation between each pair of genotypes. In the matrix, we sorted the rows in decreasing order of the values in the first column. We considered all the rows one by one, starting from the first one, and swapped the current row with the row having the next higher first-column entry among the remaining rows. For each pair of rows swapped, we also swapped the corresponding columns so as to maintain the same order of the genotypes along both the rows and columns. This arrangement implies that the genotypes are ranked in decreasing order of correlation with the first genotype: if the population splits into two subpopulations, this arrangement should show it.

The figure shows that all the pairs of genotypes have a correlation close to 1, indicating that a non-specialized genome had diffused through the whole population at the end of the evolutionary process. Measurement of the correlations was also carried out during the evolutionary process and revealed that weakly correlated subpopulations tended to emerge during the first generations of the process, but they were soon wiped out in favor of a whole population of generalist robots (data not reported).

The *stable grouping variant* of the single population genetic algorithm led to various unexpected and interesting results. Notably, this algorithm was the only one among those analyzed here to lead to the emergence of a type of strategy more suitable for the task, as indicated by direct inspection of the type of developed strategy and performance (Table 2). In particular, in the light-approaching task,

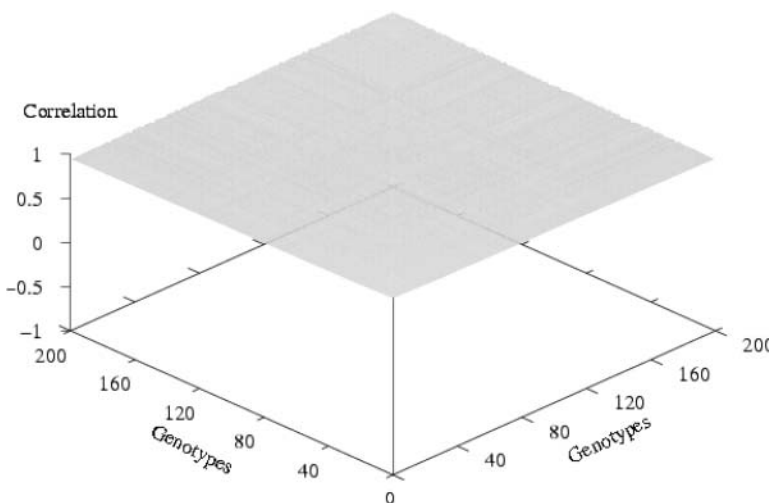


Figure 6. Correlations between the genotypes of the last generation of the best evolutionary run of the single-population genetic algorithm, temporary grouping variant, applied to the coordinated motion task (a graph very similar to this was obtained with the light-approaching task, data not reported).




Figure 7. Correlations between the genotypes of the last generation of the best evolutionary run of the single-population genetic algorithm, stable grouping variant, in the coordinated motion test.

where the ability to switch roles on the fly was advantageous, half of the evolutionary runs led to the emergence of role taking, as indicated by the types of behavior developed by the robots. This allowed the best run of the algorithm to outperform all other algorithms with the exception of the homogeneous one (t -test, $p < 0.001$ for all comparisons).

In contrast, in the coordinated-motion task, where breaking the initial sensorial symmetry could result in a considerable fitness advantage, the algorithm developed specialization in all evolutionary runs. This allowed the best run of the algorithm to have a performance very similar to that of the heterogeneous algorithm (0.947 vs. 0.948; note that the average performance of the different evolutionary runs with this algorithm is equally high in this case).

The specialization was also confirmed by the results of a quantitative correlation test similar to the one illustrated above with regard to the temporary-grouping version of the algorithm. In this respect, whereas in the case of the light-approaching task the correlation graph was similar to the one reported in Figure 6 (data not reported), in the coordinated motion task the correlation graph, reported in Figure 7, clearly showed that two different subpopulations emerged within the whole population. This indicates that in this case specialization did emerge. The measurement carried out during evolution showed that the emergence of the two subpopulations took place during the first generations and then was maintained until the end of evolution (data not reported).

The reason why stable grouping, as opposed to temporary grouping, led to the emergence of specialization is that in this evolutionary regime groups were formed only once for all the epochs composing the fitness test, so a robot with a good specialized behavior had a 0.5 probability of being paired with a companion having complementary skills. In this lucky case the robot had a higher fitness than non-specialist robots and so a higher chance of being selected and of producing several offspring. This implied a high chance that at least one of its many offspring was again paired with a companion with complementary functions, and so on through the generations, so that specialized robots had the chance to emerge and wipe out the population of non-specialized robots.

5 Analysis of the Behavioral Strategies Based on Role Taking and Specialization in the Two Tasks

This section analyzes the solutions based on role taking and specialization developed by the genetic algorithms to solve the light-approaching and the coordinated motion tasks. This is useful in many

respects: (a) it furnishes specific examples of robot behaviors based on role taking or specialization; (b) it gives a clear idea of the behavior exhibited by the robots when they use either one of the two strategies (e.g., it can be used to immediately recognize if a genetic algorithm has evolved a solution to the two tasks based either on role taking or on specialization, as done in Section 4). The analysis focuses on the outcome of the tests run with the homogeneous and the heterogeneous group selection algorithms in the two tasks. The reason for doing this is that, as mentioned in Section 3.3, these two algorithms have a strong tendency, and are best suited, to produce solutions based respectively on role taking and on specialization, so they best allow studying the behaviors, based on either strategy, that could be evolved to solve the two tasks.

5.1 The Strengths of Role Taking in the Light-Approaching Task

Table 2 indicated that homogeneous group selection outperforms heterogeneous group selection in the light-approaching task. As mentioned in Section 4, this is due to the fact that only homogeneous group selection is capable of developing the role taking needed to best solve this task.

In relation to this, a specific test was run that indicated that the different level of performance of the two algorithms was due to the different behaviors that the corresponding groups of robots exhibited when the light target changed position. In particular, the test measured the speed (*GSC*) of homogeneous and heterogeneous groups during two subphases of the test run for 1,500 steps: (1) when the barycenter of the group was within a distance between 30 and 58 cm from the off light (this phase took place right after the light target changed position and the groups were rearranging their formation); (2) when the barycenter of the groups was within a distance between 58 and 142 cm from the on light (in this phase the groups were usually moving in a stable formation). Figure 8 reports the average performance of the best groups that emerged in the 10 replications of the evolutionary runs of the homogeneous genetic algorithms and the heterogeneous algorithm. The results show that heterogeneous groups are less efficient than homogeneous groups in rearranging the formation after the light target changes position. In fact, although the performance of the algorithms does not statistically differ in the phase where they move in stable formation (*t*-test, $p = 0.21$), it does differ in the phase where they rearrange the formation (*t*-test, $p < 0.01$).

Direct inspection of the behaviors of the robots evolved with the two algorithms suggests that they have different performance in the two phases of the test because they rely upon different types of strategies. During the stable-formation phase, in which their performance does not differ, the

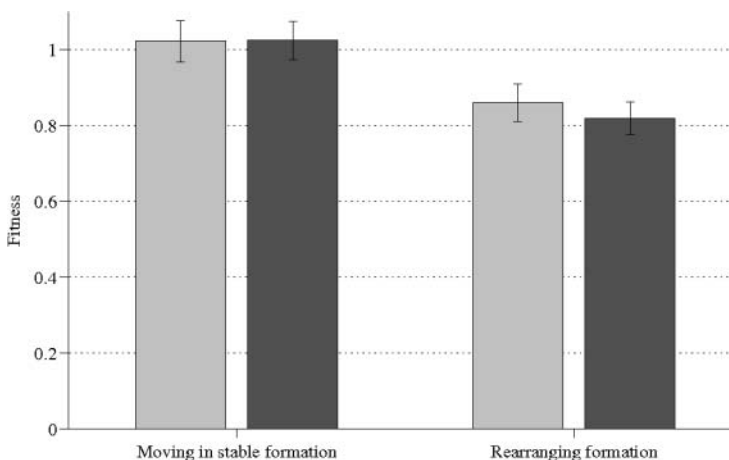


Figure 8. Average performance and SD of homogeneous (light gray bars) and heterogeneous groups (black bars) when moving toward the light target in stable formation in the central part of the arena (“Moving in stable formation”) and when rearranging the formation after the light target changes position (“Rearranging formation”).




Figure 9. Traces left by the robots of the homogeneous and heterogeneous groups in the light-approaching task. The graphs on the left side show the groups while moving in stable formation (the light target is on the west wall), while the graphs on the right side show the groups engaged in changing direction of motion (from going toward the west light to going toward the east light). (a, b) The behavior of a typical group produced by the homogeneous algorithm in two different tests where the robots, indicated by the gray and the white circles, play different initial roles within the group (e.g., the right front and the left rear positions); note how the robots rearrange the formation (in particular by switching their roles) after the light target changes position in both tests, on the basis of the current percepts. (c) The behavior of a typical group produced by the heterogeneous algorithm: Note how the gray robot always occupies the left front position as established by its genotype, even after the light switches position.

robots of both groups move at maximum speed toward the light target while maintaining stable reciprocal positions. In contrast, during the formation-rearrangement phase the robots of homogeneous groups quickly invert their orientations by about 180° by rotating on the spot (Figure 9a, b). By doing so, the robots rapidly and efficiently switch roles—in particular, their positions with respect to the front of motion of the group. For example, the robot that is at the left front of the group goes to its right rear after the light switches, whereas the robot that is at the right rear goes to the left front. This indicates that the homogeneous groups are tackling the task by relying upon role taking; that is, they play a role depending on the initial perception, stably maintain it on the basis of their different situated input patterns, and switch roles if their perception appreciably changes. On the contrary, the robots of heterogeneous groups rearrange the formation so as to always occupy the same positions within the group when the light target changes position, and this decreases their efficiency (Figure 9c). This behavior indicates that the robots are tackling the task by relying on specialization, that is, they play a certain fixed role in the group depending on their genotype.

Figure 10 reports the results of a test directed to give a quantitative measure of the different strategies that homogeneous and heterogeneous groups use in rearranging formations when the light target changes. The test measured the difference between the distances of the two robots forming a group from the light target during 1,500 cycles (several switches of the light target took place during this time). The measure was carried out in the central part of the arena, located 55 cm away from both light targets, to eliminate the noise related to the rearrangement of the formations due to the light switch.

The results of the test show that the robots of homogeneous groups have similar average distances from the light target, while the robots of heterogeneous groups have different average distances from it (the variety of differences between the robots' distances from the light was caused by the fact that in some groups robots sometimes moved side by side). This indicates that the robots of all evolved homogeneous groups play all roles within the group, that is, they are formed by generalist individuals, while the robots of all heterogeneous groups maintain the same roles during the whole duration of the test, that is, they use specialization.

In order to furnish some examples of the qualitative strategy used by the robots of the homogeneous and heterogeneous groups to switch roles or to maintain fixed roles during the light target changes, we ran some tests with evolved robots set in the arena at 40 cm from each other and both facing the target (cf. the analyses reported in [12]; recall that the robots could move either forward or backward, so they happened to develop different fronts of motion in different runs of the evolutionary process).

The results of the tests run with homogeneous groups of different evolutionary runs, exemplified in Figure 11a, b, show that (note that the length of the traces left by the robots indirectly indicates their speed): (a) the robot at the group rear moves quickly toward the other robot to improve the group compactness; (b) the robot at the front either moves backward toward the other robot and then toward the light when the group is compact, or moves slowly toward the light to allow the other robot to reach it, or approaches the light irrespective of the behavior of the second robot (in this case, group compactness is accomplished when the robots reach the first light target). As the robots are identical, if the tests are repeated by switching the initial positions of the two robots, identical results are obtained (data not shown). This confirms that behavioral strategies of homogeneous groups are based on role taking, and each robot assumes and maintains a particular role within the group, depending on the activation of its sensors. For example, the robot in front perceives the sound emitted by the second robot as coming from the opposite direction to that of the light target, whereas the second robot perceives sound and light (when not shadowed by the first robot) as coming from the same direction: These perceptual differences allow the robots to assume and maintain different roles (see [12] for further analyses of this strategy).

The results of the test run with heterogeneous groups, reported in Figure 11c–f, show that, initially, the same robot always moves to occupy the front position within the group, whereas the other robot always moves to occupy the rear position (in one out of 10 seeds the robots move side by side; however, one of them is always at the right side of the formation, while the other one is always at the left side of it). The specific strategies emerged in the different evolutionary runs are different: In some cases one of the robots, usually the one playing the staying-at-the-front role, tends to approach the light while ignoring the other robot, whereas in other cases it waits for the other robot to come closer before starting to move toward the light target. The robot that plays the staying-at-the-rear role always waits for the other robot to arrive at the front of the group and then starts to follow it (in some cases it moves backward to reach the other robot and so allows the group to

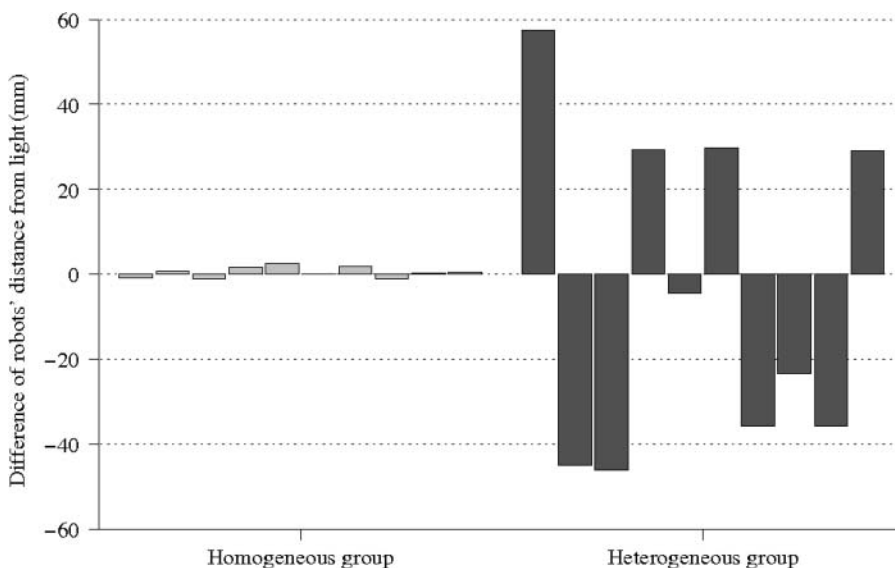


Figure 10. Difference between the distances from the light target of the two robots forming the groups produced by the homogeneous and heterogeneous algorithms, averaged over 1,500 time steps. Different bars refer to the different evolutionary runs relative to the homogeneous groups (light gray bars) and heterogeneous groups (dark gray bars). The values reported are the averages over 100 replications of the tests for each evolutionary run. Bars with positive values indicate that the robot arbitrarily designated as first was at the front of the group, whereas bars with negative values indicate that it was at the rear of the group. The bar of the fifth heterogeneous group is short because the robots forming it moved side by side.




Figure 11. Behaviors of the robots of homogeneous and heterogeneous groups that first gather and then move toward the west-wall light target. (a) Behavior exhibited by the best group that emerged in an evolutionary process run with the homogeneous group selection genetic algorithm (the behavior of the group is identical if the initial positions of the two robots forming it are switched, because the robots are clones; data not reported). Note how the robot closer to the light moves toward the other robot before approaching the target. (b) Behavior of the best group that emerged with another run of the same genetic algorithm. Note how the robot closer to the light moves toward it while ignoring the other robot. (c) Behavior exhibited by the best group that emerged with a run of the heterogeneous group selection genetic algorithm when the robot that plays the staying-at-the-front role is set close to the light. Note how such a robot does not wait for the companion before moving toward the light. (d) Behavior exhibited by the same group when the robot that plays the staying-at-the-front role is set farther from the light. Note how the other robot waits for it before moving toward the light. (e, f) Behavior exhibited by the best group of another run of the same genetic algorithm. Note how each robot waits for its companion when they are initially set close to the light.

become compact). This analysis confirms that the behavioral strategies of the heterogeneous groups are based on specialization.

5.2 The Strengths of Specialization in the Coordinated Motion Task

As shown in Table 2, heterogeneous group selection outperforms homogeneous group selection in the coordinated motion task. Indeed, in this task the initial allocation of roles is quite challenging, as the robots are in perfectly symmetric conditions (same sensors, same body, and no landmarks as external reference points), so specialized robots can rapidly solve the problem by assuming a unique role.

To support this explanation, some tests were run where the performance of homogeneous and heterogeneous groups was measured while systematically varying the initial orientations of the two robots forming the groups (the orientations were varied in intervals of 2°). The results of the tests run in the heterogeneous groups, reported in Figure 12, show that they have high performance independently of the initial robots' orientation. In contrast, homogeneous groups have performance that, as shown in Figure 13, is strongly dependent on the robots' initial orientations. In particular, when the two robots have similar absolute orientations, and hence a strong asymmetry in orientation with respect to the center of the group (this implies that robots have different perceptions; this condition corresponds to the surface of the graph along the main diagonal), the performance is higher than when they have different absolute orientations, that is, when they have more symmetrical




Figure 12. Measure of the fitness of heterogeneous groups for combinations of 180 different initial orientations of each of the two robots, averaged over the 10 evolutionary runs, in the coordinated motion task.

orientations with respect to the group's center (and hence similar perceptions). This demonstrates that homogeneous robots have difficulty in breaking the initial perceptual symmetry, a facility needed for role taking to assign different roles to the robots.

Why was the problem of the initial role allocation not encountered with the light-approaching task, as indicated by the high performance of homogeneous groups in that task? The reason is that the light target furnished the robots an external landmark that allowed them to immediately allocate roles on the basis of their asymmetric perception. Indeed, the robots had a different perceptual input pattern in each possible condition of the task; for example, when one robot perceived the light from the front and the sound from the back, the other robot perceived both the light and the sound from the front.




Figure 13. Measure of the fitness of homogeneous groups for 180 different initial orientations of the two robots, averaged over the 10 evolutionary runs, in the coordinated motion task.

The qualitative analysis of the behavioral strategies of the heterogeneous and homogeneous groups at the beginning of the coordinated motion test shows in detail why the former were advantaged with respect to the latter by specialization in the initial role allocation. In particular, Figure 14 shows that homogeneous robots exhibit good behavior only when the initial orientations are quite asymmetric with respect to the center of the group—for example, when the robots have opposite orientations with respect to that center (Figure 14a, b). In fact, in this condition the different perceptual pattern allows them to easily assume different roles, for example, the being-at-the-front or the being-at-the-rear role. On the contrary, in cases in which their initial orientation is symmetric, they have difficulty in assuming different roles (Figure 14c, d). In contrast, heterogeneous robots solve the problem of the initial allocation of roles very efficiently in any condition, as they always assume the different roles written in their genotype, independent of their initial orientation (Figure 14e–h).

Figure 15 shows the results of a test run for the coordinated motion task similar to the one run for the light-approaching task and directed to furnish a quantitative measure of the differences between the strategies used by the homogeneous and the heterogeneous groups in that task. In particular, the test measured the difference between the distances of the two robots from the starting position, averaged over the whole duration of 100 epochs lasting 150 time steps each (this was done for the final best groups of all evolutionary runs). The results of this test for the two algorithms indicate that in the homogeneous groups either one of the robots can happen to be at the front or at the rear of the group, depending on the initial orientations: This is a typical strategy based on role taking. On the contrary, in the heterogeneous groups robots always occupy a fixed position within the group with respect to the direction of motion, indicating that they use a behavioral strategy based on specialization (note that the differences of the outcome of the test run with the heterogeneous groups are due to the fact that in some cases robots move side by side).

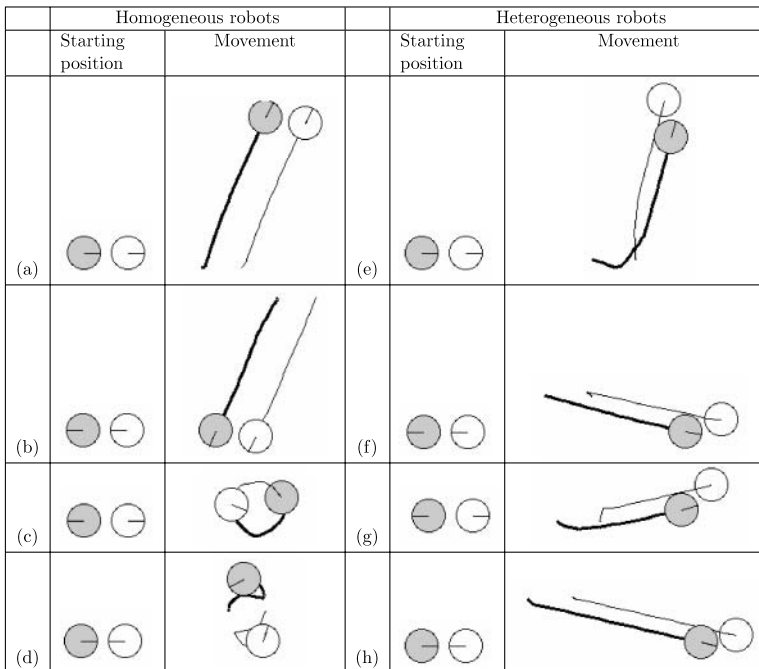


Figure 14. Behavior of homogeneous and heterogeneous groups of robots at the beginning of the coordinated motion test. (a, b) Behavior of a homogeneous group with robots having two different asymmetric initial orientations. Note how, depending on the initial orientations, each robot (e.g., the gray one) assumes different roles (e.g., staying-at-the-left or staying-at-the-right roles). (c, d) Behavior of the same homogeneous group when the robots have symmetric initial orientations, different for the two cases. Note how the robots have difficulty in assuming different roles. (e–h) Behavior of a heterogeneous group of robots with different initial orientations. Note how each robot always assumes the same role.




Figure 15. Difference of the distances from the initial positions of the two robots forming the groups produced by the homogeneous and heterogeneous algorithms when tested in the coordinated motion test. The light gray bars report the difference, averaged over 100 runs lasting 150 cycles each, of the best groups of the 10 different evolutionary runs carried out with the homogeneous algorithm, whereas the dark gray bars report the same measure for the best groups of the 10 evolutionary runs carried out with the heterogeneous algorithm. Bars with positive values indicate that the robot arbitrarily designated as first was at the front of the group, whereas bars with negative values indicate that it was at the rear of the group. Bars relative to the sixth evolutionary run of heterogeneous groups are close to zero because the robots moved side by side.

In order to qualitatively understand which difficulties were encountered by the homogeneous groups that used role taking and which specific strategies were exploited by heterogeneous groups that relied on specialization, we ran some tests in which one robot of the group could move and the other robot kept still (its wheel motors were always set to a null value). To this purpose, the robots were set in the arena within 1.5 cm of each other (this allowed them to perceive each other through the infrared sensors) and with specific orientations.

In these conditions, homogeneous groups exhibited one of two different strategies (see Figure 16; the fact that only two different strategies emerged in the coordinated motion task—contrary to what happened in the light-approaching task, where many different strategies emerged—is likely due to the limited number of sensors involved in the present task). In the evolutionary runs that develop the first strategy (*side-by-side*), the robot of the group that is free to move moves around the still robot clockwise (*clockwise-side-by-side* role) or counterclockwise (*counterclockwise-side-by-side* role), depending on its starting position. In the evolutionary runs that develop the second strategy (*front-rear*), if the freely moving robot faces the second robot with its front of motion, then it tends to go toward the second robot and hit it (*pusher* role), whereas if it has its rear toward the second robot, then it tends to keep at a particular distance from it (*distance-keeper* role). As the robots have identical controllers, the tests produce similar results if the front and rear initial positions of the two robots are switched.

In general, this analysis indicates that when both robots can freely move, they succeed in moving in space only if their initial orientations allow them to assume different roles. In particular, the first strategy is successful only if one robot tends to turn left while the other tends to turn right. The second strategy is successful only if one robot plays the role of pusher while the other plays the role of distance keeper. This confirms that strategies based on role taking might have difficulty in tasks that furnish little information on how to accomplish the initial allocations of roles.

All evolved heterogeneous groups exhibit the same two strategies exhibited by homogeneous groups (Figure 16). However, the roles played by heterogeneous robots are fixed. In the case of

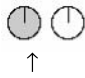

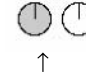

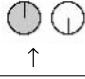

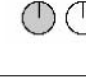

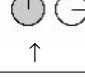
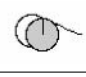
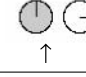
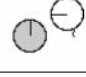
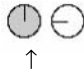
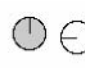


	Homogeneous robots, side-by-side strategy			Heterogeneous robots, side-by-side strategy	
	Initial position	Final position		Initial position	Final position
(a)			(e)		
(b)			(f)		
	Homogeneous robots, front-rear strategy			Heterogeneous robots, front-rear strategy	
	Initial position	Final position		Initial position	Final position
(c)			(g)		
(d)			(h)		

Figure 16. Behaviors of homogeneous and heterogeneous groups of robots in the coordinated motion task. The robots marked by an arrow in the initial position are fixed, whereas the other ones can move. (a, b) A homogeneous group where the free robot moves around the fixed robot in different directions, depending on its starting orientation with respect to the group. (c, d) A homogeneous group where the free robot keeps its distance or goes toward the other robot, depending on its initial orientation. (e, f) A heterogeneous group where one robot moves around the second one clockwise whereas the second robot moves around the first one counterclockwise. (g, h) A heterogeneous group where one robot keeps a certain distance from the second one while the second one goes toward the first one.

the first strategy, the robots move in space so as to always occupy the same position and play either the counterclockwise-side-by-side or the clockwise-side-by-side role, independent of their initial orientation. In the case of the second strategy, the robots move in space so as to always occupy the same position and play the pusher or distance-keeper role. This implies that heterogeneous robots can exploit specialization to solve the role allocation problem in tasks where the information available to the robots is highly symmetric.

6 Discussion and Conclusions

6.1 Computational Issues

The goal of this article was to analyze the potentialities of different genetic algorithms to lead to the emergence of specialization when needed, and at the same time be capable of not forcing specialization if it is not advantageous. To this purpose, the article first proposed a taxonomy of different types of mechanisms that can be used to implement specialization in evolved robots controlled by neural networks. This taxonomy includes specialization based on body differences or controller differences, the latter being due either to evolved behavior or to learned behavior. The taxonomy also includes turn taking (i.e., the capacity of dynamically assuming a role within a group) as an important precursor of stable specialized behaviors. The collective-robotics literature has proposed other taxonomies of multi-agent systems (e.g., [25]), also specifically focused on specialization [28], but the one proposed here is new in that it is formulated with a focus on controllers based on evolved neural networks: These have a specific importance for collective robotics (cf. Section 1 and [7]) and are especially relevant for biology (see Section 6.2).

Two tasks were used to pursue that goal: a light-approaching task [8], which could be performed best with turn taking, and a coordinated motion task [61], which could be performed best with specialization. The behaviors that could be evolved to perform these two tasks based on either role taking or specialization, were investigated in depth. This investigation showed that the light-approaching task could be performed with four strategies, and the collective motion task with two strategies. The knowledge of these strategies allows an immediate recognition of the fact that a genetic algorithm has evolved a strategy for the two tasks, based on either role taking or specialization. Given this knowledge and the features of the two tasks illustrated in this article, the tasks themselves represent a useful tool to study the propensity of genetic algorithms to evolve specialization or generalist solutions.

The test of the genetic algorithms with these two tasks allowed us to understand that when role taking or specialization are needed, the algorithms succeed in producing them if they can solve two evolutionary problems that we called the differentiation-bias problem and the same-specialization problem. The differentiation-bias problem refers to the difficulty of evolutionary algorithms in independently evolving individuals playing identical or quasi-identical behavioral roles. The same-specialization problem refers to the fitness costs that genetic algorithms pay if they try to evolve individuals with complementary specializations and these individuals are grouped with other individuals having the same specialization.

At a more specific level, the four genetic algorithms considered here performed as follows with respect to the two problems and hence with respect to the evolution of specialization. A first algorithm, which used as the units of selection groups of robots with identical controllers, was not capable of producing specialization, but led to the development of very powerful role-taking behaviors. In this case the difficulty of the algorithm in producing specialization is not due to the same-specialization problem, but rather to the fact that, for its functioning, the algorithm cannot produce differentiated individuals. A second algorithm, based on the selection of groups of robots with heterogeneous controllers, showed itself to be very powerful in producing specialization, but failed to produce role taking when needed. This failure can be ascribed to the differentiation-bias problem, which prevents the algorithm from developing group members with the same capabilities. Similarly, a third algorithm, which uses a different population of controllers for each robot of the group, had a strong bias to produce specialization even when not needed, and this prevented it from producing role taking.

A last genetic algorithm, based on only one population of controllers from which the controllers of the groups were randomly drawn, was the only algorithm, among those analyzed here, that exhibited the capacity to produce either role taking or specialization, depending on the task demand. This allowed the algorithm to display performance in the two tasks that was comparable to the performance of the algorithms that were most suited for those tasks, namely, the homogeneous or the heterogeneous group selection algorithms. However, this happened only with the stable grouping condition, as the temporary grouping condition was affected by the same-specialization problem and so failed to produce specialization when needed in one of the two tasks.

Interestingly, the analysis of the single-population algorithm with stable grouping revealed the reasons why it was not affected by either the differentiation-bias problem or the same-specialization problem. The results of the tests indicated that: (a) when role taking was needed, the algorithm succeeded in overcoming the differentiation-bias problem by exploiting the genetic mechanism of *reproduction* so as to produce multiple copies of virtually identical controllers; (b) when specialization was needed, the algorithm succeeded in overcoming the same-specialization problem by exploiting the fact that, although a specialized controller had indeed some chance of being grouped with a companion playing the same role, nevertheless it also had a high chance of having *at least some of its siblings* grouped with a companion playing a complementary role, so overall its genome could diffuse in the population.

Although this remains to be proven in future experiments, the single-population algorithm with stable grouping is expected to scale to groups with a high number of individuals if the number of roles requested by the task is small relative to the number of group members. The reason for this is that, statistically, if the individuals of a group are randomly drawn from the entire population, there is a high chance that the group is formed by individuals that cover all the needed roles (the larger the group and the smaller the number of different needed roles, the higher this chance). Moreover, it is also possible that the evolutionary pressure related to the fact that groups with suitable proportions of

individuals playing different roles get a higher fitness than do unbalanced groups tends to produce a suitable proportion of individuals within the whole population. Indeed, if the individuals with a certain specialization are too abundant in the population, they will tend (a) to be abundant in some groups, (b) to lower the fitness of these unbalanced groups, (c) to have relatively few offspring, and so (d) to decrease in number within the whole population. In contrast, if the individuals with a certain specialization are too scarce in the population, they will tend to obtain a higher fitness in groups where they are abundant than in groups where they are not, and so they will tend to reproduce a lot and to achieve a suitable number within the whole population. If confirmed in future work, these features might render the algorithm particularly interesting.

6.2 Biological Issues

We think that the results and the type of computational modeling presented here might give useful hints to biologists in relation to the effects of different features of evolutionary regimes on the emergence of cooperation and specialization. In this respect, the evolutionary regimes investigated here could be roughly placed along a line of increasing degree of interdependence between the individuals, measurable in terms of the degrees of cooperation and specialization between them. This might be relevant for using computational models to study the stages that lead to the emergence of more complex organisms and more complex units of selection [46]. Note that in what follows we assume that the evolutionary regimes themselves evolve and follow the success or failure destiny of the species that rely upon them.

At a first stage, single-population temporary grouping might represent the standard evolutionary condition, where cooperation emerges only rarely on account of the free-riding problem [4] and the same-specialization problem illustrated here. In this case, cooperative tasks requiring roles can be performed only on the basis of role taking, as the organisms need to function fully as individuals. The stable grouping variant of this single-population regime might be an important first step that allows the evolution of tighter collaboration within animal groups, and hence stable forms of specialization, either physical or behavioral, as it ameliorates the free-riding and same-specialization problems. This allows, for example, the formation of demographic societies of animals [81] that carry out collaborative activities and manifest some specialization, often based on individual differences (e.g., as in cooperatively breeding birds [33], or some instances of cooperative hunting [54], or cellular slime molds [43]). The borderline nature of this regime is indicated by one of the main results of this research: It is the only regime that is capable of developing either role taking or specialization, depending on the ecological challenges and opportunities.

The multiple-population regime with temporary grouping, and even more with stable grouping, represents an evolutionary regime with an even tighter interdependence between the individuals of the group and has some resemblance to cooperative ectosymbioses and endosymbioses [1, 50]. This regime tends to favor close collaboration and specialization, and so eventually to lead to a full dependence between the group members (obligate symbiosis), as the same-specialization problem is fully solved by having different populations for different specializations. Interestingly, in the biological literature endosymbiosis has been proposed to be a main drive of evolution as a means to leverage on “networking” instead of “competition” [64, 79], and for example it might have led to the development of certain organelles of the eukaryotic cells (e.g., mitochondria and chloroplasts).

Heterogeneous group selection (e.g., corresponding to multicellular systems) generates even tighter group dependences where the unit of selection is now the group: An individual in the group cannot survive by itself, and the degree of specialization within the group reaches its maximum level. Note that this regime is not strongly related to symbiosis, as it is difficult to see how the different DNA of symbiotic individuals might have become the unique DNA of multicellular organisms (indeed, as an example consider how the genetic material of the mitochondria has remained rather distinct from that of the cell nucleus). So this regime might have emerged as an extreme case of kin selection, physical connection, and specialization (slime molds might represent a good “transition example” of this [43]).

The regime of homogeneous group selection, which might correspond to social insects, represents an interesting variant with respect to the endosymbionts and multicellular systems, where the individuals remain physically distinct and relatively autonomous (and this explains the fact that they often exploit role taking) but nevertheless are highly dependent and form *superorganisms* with common or very similar genomes selected as a whole [65].

Although these considerations are tentative and preliminary, we think they indicate how the research started here might turn out to be useful for biological investigation of the evolution of specialization in real organisms and the role it might play in the emergence of more complex organisms, or more complex societies of organisms, from simpler ones.

Acknowledgments

This research was in part supported by the Swarm-bots project funded by the Future and Emerging Technologies Programme (IST-FET) of the European Commission under Grant IST-2000-31010. We thank Stefano Nolfi for valuable interactions during the project, which prompted this research and led us to develop some of the ideas presented in the article.

References

- Ahmadjian, V., & Paracer, S. (2000). *Symbiosis: An introduction to biological associations*. Oxford, UK: Oxford University Press.
- Anderson, C., Theraulaz, G., & Deneubourg, J. L. (2002). Self-assemblages in insect societies. *Insectes Sociaux*, 49, 1–12.
- Avilés, L., & Tufiño, P. (1998). Colony size and individual fitness in the social spider *Anelosimus eximius*. *American Naturalist*, 152, 403–418.
- Avilés, P. (2002). Solving the freeloaders paradox: Genetic associations and frequency-dependent selection in the evolution of cooperation among non-relatives. *Evolutionary Ecology Research*, 1, 459–477.
- Balch, T., & Arkin, R. C. (1998). Behavior-based formation control for multirobot groups. *IEEE Transactions on Robotics and Automation*, 14(6), 926–939.
- Baldassarre, G. (2008). Self-organization as phase transition in decentralized groups of robots: A study based on Boltzmann entropy. In M. Prokopenko (Ed.), *Advances in applied self-organizing systems* (pp. 127–146). Berlin: Springer-Verlag.
- Baldassarre, G., & Nolfi, S. (2009). Strengths and synergies of evolved and designed controllers: A study within collective robotics. *Journal of Artificial Intelligence*, 173, 857–875.
- Baldassarre, G., Nolfi, S., & Parisi, D. (2003). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9, 255–267.
- Baldassarre, G., Nolfi, S., & Parisi, D. (2003). Evolution of collective behaviour in a group of physically linked robots. In G. Raidl, A. Guillot, & J. A. Meyer (Eds.), *Applications of Evolutionary Computing—Proceedings of the Second European Workshop on Evolutionary Robotics* (pp. 581–592). Berlin: Springer Verlag.
- Baldassarre, G., Parisi, D., & Nolfi, S. (2004). Coordination and behavior integration in cooperating simulated robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, & J. A. Meyer (Eds.), *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior* (pp. 385–394). Cambridge, MA: MIT Press.
- Baldassarre, G., Parisi, D., & Nolfi, S. (2004). Measuring coordination as entropy decrease in groups of linked simulated robots. In A. Minai & Y. Bar-Yam (Eds.), *Proceedings of the Fifth International Conference on Complex Systems (ICCS2004)*. Cambridge, MA: The New England Complex Systems Institute. Available at <http://www.necsi.org/events/iccs/2004proceedings.html>.
- Baldassarre, G., Parisi, D., & Nolfi, S. (2006). Distributed coordination of simulated robots based on self-organization. *Artificial Life*, 12(3), 289–311.
- Baldassarre, G., Trianni, V., Bonani, M., Mondada, F., Dorigo, M., & Nolfi, S. (2007). Self-organized coordinated motion in groups of physically connected robots. *IEEE Transactions in Systems, Man and Cybernetics—Part B Cybernetics*, 37(1), 224–239.

14. Barfoot, T. D., & Clark, C. M. (2004). Motion planning for formations of mobile robots. *Journal of Robotics and Autonomous Systems*, 46, 65–78.
15. Beckers, R., Holland, O. E., & Deneubourg, J. L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. A. Brooks & P. Maes (Eds.), *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems (Artificial Life IV)* (pp. 181–189). Cambridge, MA: MIT Press.
16. Böhringer, K., Brown, R., Donald, B., Jennings, J., & Rus, D. (1997). Distributed robotic manipulation: Experiments in minimalism. In O. Khatib et al. (Eds.), *Experimental robotics IV* (pp. 11–25). Berlin: Springer Verlag.
17. Boinski, S., & Garber, P. A. (Eds.). (2000). *On the move: How and why animals travel in groups*. Chicago: University of Chicago Press.
18. Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford, UK: Oxford University Press.
19. Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-organization in biological systems*. Princeton, NJ: Princeton University Press.
20. Cao, Y. U., Fukunaga, A. S., & Kahng, A. B. (1997). Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4, 1–23.
21. Desai, J. P., Ostrowski, J. P., & Kumar, V. (2001). Modeling and control of formations of nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 17(6), 905–908.
22. Dorigo, M., & Şahin, E. (2004). Swarm robotics—Special issue editorial. *Autonomous Robots*, 17(2–3), 111–113.
23. Dorigo, M., Trianni, V., Şahin, E., Groß, R., Labella, T. H., Baldassarre, G., Nolfi, S., Deneubourg, J. L., Floreano, D., & Gambardella, L. M. (2004). Evolving self-organizing behavior for a swarm-bot. *Autonomous Robots*, 17(2–3), 223–245.
24. Düblier, N., Mülders, C., Ferdelman, T., de Beer, D., Pernthaler, A., Klein, M., Wagner, M., Erséus, C., Thiermann, F., Krieger, J., Giere, O., & Amann, R. (2001). Endosymbiotic sulphate-reducing and sulphide-oxidizing bacteria in an oligochaete worm. *Nature*, 411(6835), 298–302.
25. Dudek, G., Jenkin, M., Milius, E., & Wilkes, D. (1996). A taxonomy for multi-agent robotics. *Autonomous Robots*, 3, 375–397.
26. Dussutour, A., Fourcassié, V., Helbing, D., & Deneubourg, J. (2004). Optimal traffic organization in ants under crowded conditions. *Nature*, 428(6978), 70–73.
27. Fredslund, J., & Mataric, M. J. (2002). A general algorithm for robot formations using local sensing and minimal communication. *IEEE Transactions on Robotics and Automation*, 18(5), 837–846.
28. Gerkey, B. P., & Mataric, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9), 939–954.
29. Gomez, F., & Miikkulainen, R. (2002). *Robust non-linear control through neuroevolution* (TR AI-2002-292). Austin: The University of Texas at Austin.
30. Gompper, M. E. (1997). Population ecology of white-nosed coati (*Nasua narica*) on Barro Colorado Island, Panama. *Journal of Zoology*, 241, 441–445.
31. Grabowski, R., Navarro-Serment, L., Paredis, C., & Khosla, P. (2000). Heterogeneous teams of modular robots for mapping and exploration. *Autonomous Robots*, 8(3), 293–308.
32. Groß, R., Bonani, M., Mondada, F., & Dorigo, M. (2006). Autonomous self-assembly in mobile robotics. In K. Murase, K. Sekiyama, N. Kubota, T. Naniwa, & J. Sitte (Eds.), *Proceedings of the Third International Symposium on Autonomous Minirobots for Research and Edutainment* (pp. 314–322). Berlin: Springer Verlag.
33. Heinsohn, R. G. (1992). Cooperative enhancement of reproductive success in white-winged coughts. *Evolutionary Ecology*, 6, 247–257.
34. Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5, 173–202.
35. Ijspeert, A. J., Martinoli, A., Billard, A., & Gambardella, L. M. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11, 149–171.

36. Keller, L., & Chapuisat, M. (1999). Cooperation among selfish individuals in insect colonies. *BioScience*, *49*, 899–909.
37. Keller, L., & Ross, K. (1998). Selfish genes: A green beard in the red fire ant. *Nature*, *394*, 573–575.
38. Khatib, O., Yokoi, K., Chang, K., Ruspini, D., Holmberg, R., & Casal, A. (1996). Coordination and decentralized cooperation of multiple mobile manipulators. *Journal of Robotic Systems*, *13*(11), 755–764.
39. Krieger, M. J. B., Billeter, J. B., & Keller, L. (2000). Ant-like task allocation and recruitment in cooperative robots. *Nature*, *406*, 992–995.
40. Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, *30*, 85–101.
41. Kube, C. R., & Zhang, H. (1993). Collective robotics: From social insects to robots. *Adaptive Behavior*, *2*(2), 189–219.
42. Lerman, K., Jones, C., Galstyan, A., & Mataric, M. (2006). Analysis of dynamic task allocation in multi-robot systems. *The International Journal of Robotics Research*, *25*(3), 225.
43. Marée, A. F., & Hogeweg, P. (2001). How amoeboids self-organize into a fruiting body: Multicellular coordination in *Dictyostelium discoideum*. *Proceedings of the National Academy of Science*, *98*(7), 3879–3883.
44. Martinoli, A. (1999). *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies*. PhD thesis, Computer Science Department, Ecole Polytechnique Fédérale de Lausanne.
45. Matsumoto, K., Takagi, S., & Nakagaki, T. (2008). Locomotive mechanism of *Physarum plasmodia* based on spatiotemporal analysis of cytoplasmic streaming. *Biophysical Journal*, *94*(7), 2492–2504.
46. Maynard Smith, J., & Szathmáry, E. (1997). *The major transitions in evolution*. Oxford, UK: Oxford University Press.
47. Miglino, O., Lund, H. H., & Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, *4*, 417–434.
48. Mondada, F., Franzi, E., & Jenne, P. (1993). Mobile robot miniaturisation: A tool for investigation in control algorithms. In T. Yoshikawa & F. Miyazaki (Eds.), *Proceedings of the Third International Symposium on Experimental Robotics*. Berlin: Springer Verlag.
49. Mondada, F., Pettinaro, G., Guignard, A., Kwee, I., Floreano, D., Denebourg, J. L., Nolfi, S., Gambardella, L. M., & Dorigo, M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, *17*(2–3), 193–221.
50. Moran, N. A. (2006). Symbiosis. *Current Biology*, *16*(20), 866–871.
51. Nolfi, S. (2006). Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour. *Complexus*, *2*(3–4), 195–203.
52. Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
53. Nolfi, S., Floreano, D., Miglino, O., & Mondada, F. (1994). How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks & P. Maes (Eds.), *Proceedings of the International Conference on Artificial Life IV* (pp. 190–197). Cambridge, MA: MIT Press.
54. Packer, C., & Ruttan, L. (1988). The evolution of cooperative hunting. *American Naturalist*, *132*(2), 159–198.
55. Panait, L., & Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, *11*, 387–434.
56. Parker, L., Touzet, C., & Fernandez, F. (2001). Techniques for learning in multi-robot teams. In B. Tucker & L. E. Paker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A K Peters.
57. Parrish, J. K., & Hamner, W. M. (Eds.). (1997). *Animal groups in three dimensions: How species aggregate*. Cambridge, UK: Cambridge University Press.
58. Partridge, B. (1982). The structure and function of fish schools. *Scientific American*, *June*, 114–123.

59. Pérez-Uribe, A., Floreano, D., & Keller, L. (2003). Effects of group composition and level of selection in the evolution of cooperation in artificial ants. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Advances in Artificial Life: Proceedings of the Seventh European Conference on Artificial Life (ECAL2003)* (pp. 128–137). Berlin: Springer Verlag.
60. Potts, W. (1984). The chorus line hypothesis of manoeuvre coordination in avian flocks. *Nature*, *309*, 344–345.
61. Quinn, M. (2001). A comparison of approaches to the evolution of homogeneous multi-robot teams. *Proceedings of the 2001 Congress on Evolutionary Computation (CEC2001)* (pp. 128–135). New York: IEEE Press.
62. 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 of London—A: Mathematical, Physical and Engineering Sciences*, *361*, 2321–2344.
63. Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, *21*(4), 25–34.
64. Sagan, D., & Margulis, L. (1997). *Microcosmos: Four billion years of evolution from our microbial ancestors*. Berkeley: University of California Press.
65. Seeley, T. D. (1989). The honey bee colony as a superorganism. *American Scientist*, *77*, 546–553.
66. Soltoggio, A., Bullinaria, J. A., Mattiussi, C., Dürr, P., & Floreano, D. (2008). Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios. In *Proceedings of the 11th International Conference on the Simulation and Synthesis of Living Systems (Alife XI)* (pp. 569–576). Cambridge, MA: MIT Press.
67. Spector, L., Klein, J., Perry, C., & Feinstein, M. (2005). Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, *8*(1), 111–125.
68. Sugar, T. G., & Kumar, V. (2002). Control of cooperating mobile manipulators. *IEEE Transactions on Robotics and Automation*, *18*(1), 94–103.
69. Theraulaz, G., Bonabeau, E., & Deneubourg, J. (1998). Response threshold reinforcements and division of labour in insect societies. *Proceedings of the Royal Society of London—Series B: Biological Sciences*, *265*, 327–332.
70. Trianni, V., Ampatzis, C., Christensen, A. L., Tuci, E., Dorigo, M., & Nolfi, S. (2007). From solitary to collective behaviours: Decision making and cooperation. In F. Almeida e Costa, M. Rocha Luis, E. Costa, I. Harvey, & A. Coutinho (Eds.), *Proceedings of the 9th European Conference on Artificial Life (ECAL 2007)* (pp. 575–584). Berlin: Springer Verlag.
71. Trianni, V., Nolfi, S., & Dorigo, M. (2006). Cooperative hole-avoidance in a swarm-bot. *Robotics and Autonomous Systems*, *54*(2), 97–103.
72. Tucker, B., & Paker, L. E. (Eds.). (2002). *Robot teams: From diversity to polymorphism*. Natick, MA: A K Peters.
73. Turgut, A. E., Huepe, C., Çelikkanat, H., Gökçe, F., & Şahin, E. (2008). Modeling phase transition in self-organised flocks. In M. Dorigo, M. Birattari, C. Blum, M. Clerc, T. Stützle, & A. F. T. Winfield (Eds.), *Proceedings of the 6th International Conference on Ant Colony Optimization and Swarm Intelligence (ANTS 2008)* (pp. 108–119). Berlin: Springer-Verlag.
74. Tyler, M. S. (2000). *Developmental biology: A guide for experimental study*. Sunderland, MA: Sinauer.
75. Waibel, M., Keller, L., & Floreano, D. (2009). Genetic team composition and level of selection in the evolution of multi-agent systems. *IEEE Transactions on Evolutionary Computation*, *13*(3), 648–660.
76. Wang, P. K. C. (1991). Navigation strategies for multiple autonomous mobile robots moving in formation. *Journal of Robotic Systems*, *8*(2), 177–195.
77. Wang, Z. D., Nakano, E., & Takahashi, T. (2003). Solving function distribution and behavior design problem for cooperative object handling by multiple mobile robots. *IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans*, *33*(5), 537–549.
78. Ward, C. R., Gibet, F., & Kendall, G. (2001). Evolving collective behaviour in artificial ecologies. *Artificial Life*, *7*, 191–209.
79. Wernegreen, J. J. (2004). Endosymbiosis: Lessons in conflict resolution. *PLoS Biology*, *2*(3), E68.

80. Wilson, E. O. (1971). *The insect societies*. Cambridge, MA: Belknap Press of Harvard University Press.
81. Wilson, E. O. (1975). *Sociobiology: The new synthesis*. Cambridge, MA: Belknap Press of Harvard University Press.
82. Wolpert, L., Beddington, R., Jessell, T., Lawrence, P., Meyerowitz, E., & Smith, J. (2002). *Principles of development*. Oxford, UK: Oxford University Press.
83. Yamauchi, B., & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3), 219–246.
84. Yang, X., Watanabe, K., Izumi, K., & Kiguchi, K. (2004). A decentralized control system for cooperative transportation by multiple non-holonomic mobile robots. *International Journal of Control*, 77(10), 949–963.
85. Zhu, W. H., & De Schutter, J. (1999). Control of two industrial manipulators rigidly holding an egg. *IEEE Control Systems Magazine*, 19(2), 24–30.

