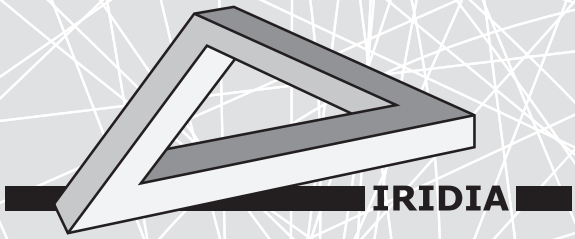

Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Evolving Autonomous Self-Assembly in Homogeneous Robots

Christos Ampatzis, Elio Tuci, Vito Trianni, Anders Lyhne Christensen and Marco Dorigo

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2008-004

January 2008

Submitted to Artificial Life

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:
IRIDIA, *Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle*
Université Libre de Bruxelles
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2008-004

Revision history:
TR/IRIDIA/2008-004.001 January 2008

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsability for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Evolving Autonomous Self-Assembly in Homogeneous Robots

Christos Ampatzis[†] campatzi@iulb.ac.be
Elio Tuci[†] etuci@ulb.ac.be
Vito Trianni[†††] vito.trianni@istc.cnr.it
Anders Lyhne Christensen[††] anders.christensen@iscte.pt
Marco Dorigo[†] mdorigo@ulb.ac.be

† CoDE-IRIDIA, Université Libre de Bruxelles (ULB), Av. F. Roosevelt 50, CP 194/6, 1050 Brussels, Belgium
†† DCTI-ISCTE, Av. das Forças Armadas, 1649-026 Lisbon, Portugal
††† ISTC-CNR, via San Martino della Battaglia 44, 00185 Roma, Italy
Corresponding author: Christos Ampatzis[†], phone: +3226502730, fax: +3226502715

January 2008

Abstract

This research work illustrates an approach to the design of controllers for self-assembling robots in which the self-assembly is initiated and regulated by perceptual cues that are brought forth by the physical robots through their dynamical interactions. More specifically, we present a homogeneous control system that can achieve assembly between the modules (fully autonomous robots) of a mobile self-reconfigurable system without a priori introduced behavioural or morphological heterogeneities. The control system is an evolved dynamical neural network that directly controls all the actuators and causes the dynamic specialisation of the robots by allocating roles between them based solely on their interaction. Our results suggest that direct access to the orientations or intentions of the other agents is not a necessary condition for robot coordination: our robots coordinate without direct or explicit communication, contrary to what assumed by most research works in collective robotics. Finally, we show that our evolved controllers prove to be successful when tested on a real hardware platform, the *swarm-bot*. The performance of our evolved neuro-controllers is similar to the one achieved by existing modular or behaviour-based approaches, owing to the effect of an emergent recovery mechanism that was not foreseen during the evolutionary simulation. However, contrary to other approaches, our system proved to be robust against changes in the experimental setup, because of the reduction in the number of user-defined preconditions for robot self-assembly.

Key Words: self-assembly, coordination, role allocation, neural network
Short Title: Evolution of Autonomous Self-Assembly

1 Introduction

According to Whitesides and Grzybowski (2002), self-assembly is defined as "the autonomous organisation of components into patterns or structures without human intervention". Nature provides many examples of animals forming collective structures by connecting themselves to one another. Individuals of various ant, bee and wasp species self-assemble and manage to build complex structures such as bivouacs, ladders etc. Self-assembly in social insects typically happens in order to accomplish some function (defence, object transport, passage formation etc., see Anderson et al., 2002). In particular, ants of the species *Œcophylla longinoda* can form chains composed of their own bodies which are used to pool leaves together to form a nest, or to bridge a passage between branches in a tree (Hölldobler and Wilson, 1978). Self-assembly is also widely observed at the molecular level (e.g., DNA molecules).

The robotics community has been largely inspired from cooperative behaviour in animal societies when designing controllers for groups of robots that have to accomplish a given task. In particular, self-assembly provides a novel way of cooperation in groups of robots. Recently, the research work carried out in the context of the SWARM-BOTS project[1] proved that it is possible to build and control a group of autonomous self-assembling robots by using swarm robotics principles. Swarm robotics represents a novel way of doing collective robotics in which autonomous cooperating agents are controlled by distributed and local rules (Bonabeau et al., 1999). That is, each agent uses only local perception to decide what action to take. Research in swarm robotics focuses on mechanisms to enhance the efficiency of the group through some form of cooperation among the individual agents. In this respect, self-assembly can enhance the efficiency of a group of autonomous cooperating robots in several different contexts. Through self-assembly a group of robots can overcome the physical limitations of each individual robot. Within the SWARM-BOTS project, it has been proved that self-assembly can offer robotic systems additional capabilities useful for the accomplishment of the following tasks: a) robots collectively and cooperatively transporting items too heavy to be moved by a single robot (Groß and Dorigo, 2008a); b) robots climbing a hill whose slope would cause a single robot to topple over (O'Grady et al., 2005); c) robots navigating on rough terrain in which a single agent might topple over (O'Grady et al., 2005). The application of such systems can potentially go beyond research in laboratories, space applications being the most obvious challenge (e.g. multi-robot planetary exploration and on-orbit self-assembly, see Izzo and Pettazzi, 2007; Izzo et al., 2005).

This research work illustrates an approach to the design of controllers for self-assembling robots in which the self-assembly is initiated and regulated by perceptual cues that are brought forth by the physical robots through their dynamical interactions. More specifically, we use Evolutionary Robotics (ER) as the methodology for the design of the control system of autonomous robots that must coordinate their motion in order to self-assemble. ER is a methodological tool to automate the design of robots' controllers (Nolfi and Floreano, 2000). It is based on the use of artificial evolution to find sets of parameters for artificial neural networks that guide the robots to the accomplishment of their task. With respect to other design methods, ER does not require the designer to make strong assumptions concerning what behavioural and communication mechanisms are needed by the robots. The experimenter defines the characteristics of a social context in which robots are required to cooperate. The agents' mechanisms for solitary and social behaviour are determined by an evolutionary process that favours (through selection) those solutions which improve an agent's or group's ability to accomplish its task (i.e., the fitness measure).

The artificial evolutionary process was exploited to synthesize dynamical neural network controllers (Continuous Time Recurrent Neural Networks—CTRNNs, see Beer and Gallagher, 1992) capable of autonomous decision-making and self-assembling in a homogeneous group of robots. Dynamical neural networks have been used in the past as a means to achieve specialisation in a robot group (see Quinn et al., 2003; Tuci et al., 2008, for example); similarly, we study self-assembly in a setup where the robots interact and eventually differentiate by allocating distinct roles. In other words, we train via artificial evolution a dynamical neural network that when downloaded on real robots allows them to coordinate their actions in order to decide who will grip whom.

It is important to notice that some characteristics of the hardware impose important constraints on the control of the modules of a self-assembling system. Some hardware platforms consist of morphologically heterogeneous modules, that can only play a predefined role in the assembly process. In others, the hardware design does not allow, for example, the assembly of more than two modules, or requires extremely precise alignment during the connection phase, that is, it requires a great accuracy. As argued by Tuci et al. (2006), the *swarm-bot* platform, thanks to its sensors and actuators and its connecting

[1] A project funded by the Future and Emerging Technologies Programme (IST-FET) of the European Commission, under grant IST-2000-31010. See also http://www.swarm-bots.org

mechanism, does not severely constrain the design of control mechanisms for self-assembly. This platform consists of identical modules, each equipped with a gripper and a large area to receive connections from other modules. In this work, we use this robotic platform to investigate autonomous self-assembly and role allocation.

The main contribution of this work lies in the design of control strategies for real assembling robots that are not constrained by either morphological or behavioural heterogeneities introduced by the hardware and control method, respectively. To the best of our knowledge, there is no system in the literature that can achieve self-assembly without a priori injected morphological or behavioural heterogeneities. Instead of a priori defining the mechanisms leading to role allocation and self-assembly, we let behavioural heterogeneity emerge from the interaction among the system's homogeneous components. We believe that by following such an approach, we can aim at obtaining more autonomous and more adaptive robotic systems, because the adaptiveness of an autonomous multi-robot system is reduced if the circumstances an agent should take into account to make a decision (concerning solitary and/or social behaviour) are defined by a set of a priori assumptions made by the experimenter. Moreover, we show that an integrated (i.e., non-modularised) neural network in direct control of all the actuators of the robots can successfully tackle real-world tasks requiring fine-grained sensory-motor coordination, such as self-assembly. Finally, we show with physical robots that coordination and cooperation in self-assembly do not require explicit signalling of internal states, as assumed, for example, by (Groß et al., 2006a). In other words, we present a setup that requires minimal cognitive and communicative capacities on behalf of the robots.

In section 2 we provide a brief review of the state of the art in the area of self-assembling robots and we discuss the limitations of these systems, justifying the methodological choices we have made. In the following sections (sections 3, 4 and 5), we describe the evolutionary machinery and the experimental scenario used to design neural network controllers. Then, in section 6 we show the results of post-evaluation tests on physical robots controlled by the best performing evolved controller and we try to shed some light on the mechanisms underpinning the behaviour of successful robots. The results presented are discussed in section 7 while conclusions are drawn in section 8.

2 Related Work

Several examples of robotic platforms in the literature consist in connecting modules. For a very comprehensive review of self-assembling robotic systems, we direct the reader to the work of Yim et al. (2002a); Groß and Dorigo (2008b); Groß et al. (2006a); Tuci et al. (2006). Following Yim et al. (2002a), it is possible to identify four different categories: chain based, lattice based, mobile and stochastic reconfigurable robots. As this work focuses on mobile self-reconfigurable robots, in the following, we provide a small overview of this category only. We go on to discuss the platform that is used in this study: the *swarm-bot*.

2.1 Mobile Self-reconfigurable Robots

The first example of a mobile self-reconfigurable robot was the CEBOT (see Fukuda and Nakagawa, 1987; Fukuda and Ueyama, 1994). CEBOT is a heterogeneous system comprised of cells with different functions (move, bend, rotate, slide). Even though there are no quantitative results to assess the performance and reliability of this system, Fukuda et al. (1988) have shown how docking can be done between a moving cell and a static object cell. Another robotic system capable of self-assembly is the Super Mechano Colony (Damoto et al., 2001; Hirose, 2001). In this system, autonomous robotic wheels, referred to as child units, can connect to and disconnect from a mother-ship. Yamakita et al. (2003) achieved docking by letting the child unit follow a predefined path. Groß et al. (2006b) recently demonstrated assembly between one and three moving child modules and a static module. Hirose et al. (1996) presented a distributed robot called Gunryu. Each robot is capable of fully autonomous locomotion and the assembled structure proved capable of navigating on rough terrain where a single unit would topple over. However, autonomous self-assembly was not studied as the units were connected from beforehand by means of a passive arm. Self-assembly is also not possible for the Millibot train (see Brown et al., 2002), composed of multiple modules that are linearly linked, since no external sensor has been implemented.

In all the above mobile self-reconfigurable systems, self-assembly is either not achieved at all or is only possible between one unit moving autonomously and a static object/unit. For the sake of consistency, we should also mention two important examples from the modular chain robot literature, CONRO and PolyBot. CONRO (Castano et al., 2000) has been used by Rubenstein et al. (2004) to demonstrate autonomous docking between two robots. It should be noted however, that the control was heterogeneous

at all levels and the generality of the approach was limited due to orientation and distance constraints. Yim et al. (2002b) demonstrated self-assembly with PolyBot: a six-modules arm connected to a spare module on a flat terrain. One end of the arm and the spare module were fixed to the walls of the arena at known positions and the motion of the arm relied on knowledge of the goal position and inverse kinematics.

2.2 Self-assembly with the *Swarm-Bot*

The *swarm-bot*, a collective and mobile reconfigurable system (see Mondada et al., 2005; Dorigo, 2005, and http://www.swarm-bots.org), consists of fully autonomous mobile robots called *s-bots*, that can physically connect to each other and to static objects (preys, also called *s-toys*). Groß et al. (2006a) presented experiments improving the state of the art in self-assembling robots concerning mainly the number of robots involved in self-assembly, the generality and reliability of the controllers and assembly speed. A significant contribution of this work is in the design of distributed control mechanisms for self-assembly relying only on local perception. In particular, self-assembly was accomplished with a modular approach in which some modules have been evolved and others hand-crafted. The approach was based upon a signalling system which makes use of colours. For example, the decision concerning which robot makes the action of gripping (the *s-bot-gripper*) and which one is gripped (the *s-bot-grippee*) is made through the emission of colour signals, according to which the *s-bots* emitting blue light are playing the role of *s-bot-gripper* and those emitting red light the role of *s-bot-grippee*. Thus, it is the heterogeneity among the robots with respect to the colour displayed, a priori introduced by the experimenter, that triggers the self-assembly process. That is, a single *s-bot* "born" red among several *s-bots* "born" blue is meant to play the role of *s-bot-grippee* while the remaining *s-bot-grippers* are progressively assembling. Once successfully assembled to another *s-bot*, each blue light emitting robot was programmed to turn off the blue LEDs and to turn on the red ones. The switch from blue to red light indicates to the yet non-assembled *s-bots* the "metamorphosis" of a robot from *s-bot-gripper* to *s-bot-grippee*. This system is therefore based on the presence of a behavioural or morphological heterogeneity. In other words, it requires either the presence of a prey lit up in red or the presence of a robot not sharing the controller of the others, which is forced to be immobile and to signal with a red colour. O'Grady et al. (2005) bypassed this requirement by handcrafting a decision-making mechanism based on a probabilistic transition between states. More specifically, the allocation of roles (which robot lights up red and triggers the process) depends solely on a stochastic process.

The research works presented above have been very successful since they also showed how assembled structures can overcome limitations of the single robots, for instance in transporting a heavy object or in navigating on rough terrain. However, this modularised architecture is based on a set of a priori assumptions concerning the specification of the environmental/behavioural conditions that trigger the self-assembling process. For example, (a) the objects that can be grasped must be red, and those that should not be grasped must be blue; (b) the action of grasping is carried out only if all the "grasping requirements" are fulfilled (among others, a combination of conditions concerning the distance and relative orientation between the robots, see Groß et al., 2006a, for details). If the experimenter could always know in advance in what type of world the agents will be located, assumptions such as those concerning the nature of the object to be grasped would not represent a limitation with respect to the domain of action of the robotic system. However, since it is desirable to have agents that can potentially adapt to variable circumstances or conditions that are partially or totally unknown to the experimenter, it follows that the efficiency of autonomous robots should be estimated also with respect to their capacity to cope with "unpredictable" events (e.g., environmental variability, partial hardware failure, etc.). For example, failure to emit or perceive red light for robots guided by the controllers presented above would significantly hinder the accomplishment of the assembly task.

We believe that a sensible step forward in this direction can be made by avoiding to constrain the system to initiate its most salient behaviours (e.g., self-assembly) in response to a priori specified agent's perceptual states. The work described in this paper represents a significant step forward in this direction. Our research work illustrates the details of an alternative methodological approach to the design of homogeneous controllers (i.e., where a controller is cloned in each robot of a group) for self-assembly in physical autonomous robots in which no assumptions are made concerning how agents allocate roles. By using dynamical neural networks shaped by artificial evolution, we managed to design mechanisms by which the allocation of the *s-bot-gripper* and the *s-bot-grippee* roles is the result of an autonomous negotiation phase between the *s-bots*, and not predetermined by the experimenter. In other words, the self-assembly process is triggered and regulated by perceptual cues that are brought forth by the agents through their dynamical interactions. Furthermore, coordination and role allocation in our system is achieved solely

through minimal sensors (distance and angle information) and without explicit communication, contrary to the above works where the agents signal their internal states to the rest of the group. Also, due to the nature of the sensory system used, the robots cannot sense the orientation of their group-mates. In this sense, our approach is similar to (and largely inspired from) the one of (Quinn, 2001; Quinn et al., 2003), where role allocation (leader-follower) or formation movement is achieved solely through infrared sensors. In addition, we also show that the evolved mechanisms are as effective as the modular and hand-coded ones described in (Groß et al., 2006a; O'Grady et al., 2005) when controlling two real *s-bots*.

3 Simulated and real *s-bot*

The controllers are evolved in a simulation environment which models some of the hardware characteristics of the real *s-bots* (see Mondada et al., 2004). An *s-bot* is a mobile autonomous robot equipped with many sensors useful for the perception of the surrounding environment and for proprioception, a differential drive system, and a gripper by which it can grasp various objects or another *s-bot* (see figure 1a). The main body is a cylindrical turret with a diameter of 11.6 cm, which can be actively rotated with respect to the chassis. The turret is equipped with a surrounding ring that receives connections from other *s-bots* through their grippers.

In this work, to allow robots to perceive each other, we make use of the omni-directional camera mounted on the turret. The image recorded by the camera is filtered in order to return the distance of the closest red, green, or blue blob in each of eight 45° sectors. A sector is referred to as CAM_i, where $i = 1, ..., 8$ denotes the index of the sector. Thus, an *s-bot* to be perceived by the camera must light itself up in one of the three colours using the LEDs mounted on the perimeter of its turret. An *s-bot* can be perceived in at most two adjacent sectors. Notice that the camera can clearly perceive coloured blobs up to a distance of approximately 50 cm, but the precision above approximately 30 cm is rather low. Moreover, the precision with which the distance of coloured blobs is detected varies with respect to the colour of the perceived object. We also make use of the optical barrier which is a hardware component composed of two LEDs and a light sensor mounted on the gripper (see figure 1b). By post-processing the readings of the optical barrier we extract information about the status of the gripper and about the presence of an object between the gripper claws. More specifically, the post-processing of the optical barrier readings defines the status of two virtual sensors: a) the *GS* sensor, set to 1 if the optical barrier indicates that there is an object in between the gripper claws, 0 otherwise; b) the *GG* sensor, set to 1 if a robot is currently grasping an object, 0 otherwise. We also make use of the *GA* sensor, which monitors the gripper aperture. The readings of the *GA* sensor range from 1 when the gripper is completely open to 0 when the gripper is completely closed. The *s-bot* actuators are the two wheels and the gripper.

The simulator used to evolve the required behaviour relies on a specialised 2D dynamics engine (see Christensen, 2005). In order to evolve controllers that transfer to real hardware, we overcome the limitations of the simulator by following the approach proposed in Jakobi (1997); motion is simulated with sufficient accuracy, collisions are not. Self-assembly relies on rather delicate physical interactions between robots that are integral to the task (e.g., the closing of the gripper around an object could be interpreted

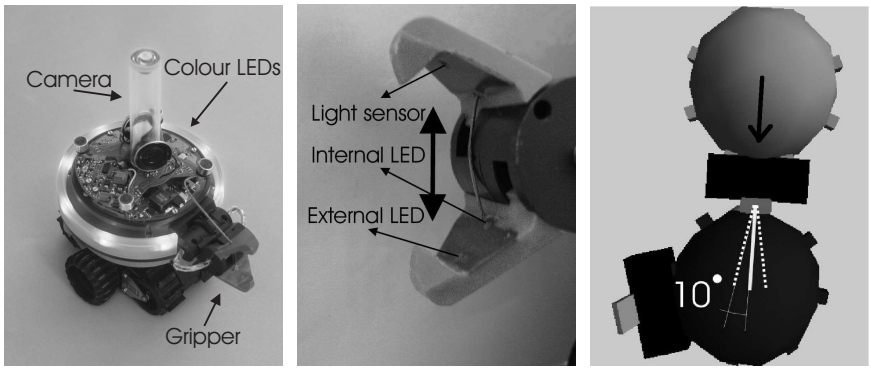

(a) (b) (c)

Figure 1: (a) The s-bot. (b) The gripper and sensors of the optical barrier. (c) Depiction of the collision manager. The arrow indicates the direction along which the *s-bot-gripper* should approach the *s-bot-grippee* without incurring into collision penalties.

as a collision). Instead of trying to accurately simulate the collisions, we force the controllers to minimise them and not to rely on their outcome. In other words, in case of a collision, the two colliding bodies are repositioned to their previous positions, and the behaviour is penalised by the fitness function if the collision can not be considered the consequence of an accepted grasping manoeuvre.

Concerning the simulation of the gripper, we modelled the two gripper claws as triangles extending from the body of the robot. As the gripper opens, these triangles are pulled in the robot body, whereas as it closes they grow out of it. Thus the size of the collision object changes with the aperture of the gripper. In order for a grip to be called successful, we require that there is an object between the claws of the (open) gripper, as close as possible to the interior of the gripper and that the claws close around it. In fact, we require that the object and the gripper socket holding the two claws collide. However, we do not penalise such a collision when the impact angle between the *s-bots* falls within the range [-10°,+10°]. Figure 1c shows how this impact angle is calculated and also depicts the simulated robots we use. In this way, we facilitate the evolution of approaching movements directed towards the turret of the robot to be gripped (see figure 1c). Robots that rely on such a strategy when attempting to self-assemble in simulation, can also be successful in reality. Other types of strategies based on rotating movements proved prone to failure when tested on real hardware. Having taken care of the collisions involved with gripping, the choice of a simple and fast simulator instead of one using a 3D physics engine significantly speeds up the evolutionary process.[2]

4 Controller and Evolutionary Algorithm

The agent controller is composed of a continuous time recurrent neural network (CTRNN) of ten hidden neurons and an arrangement of eleven input neurons and three output neurons (see figure 2a and Beer and Gallagher (1992) for a more detailed illustration of CTRNNs). Input neurons have no state. At each simulation cycle, their activation values I_i—with $i \in [1, 11]$—correspond to the sensors' readings. In particular, I_1 corresponds to the reading of the *GA* sensor, I_2 to the reading of the *GG* sensor, I_3 to I_{10} correspond to the normalised reading of the eight camera sectors CAM_i, and I_{11} corresponds to the reading of the *GS* sensor. Hidden neurons are fully connected. Additionally, each hidden neuron receives one incoming synapse from each input neuron. Each output neuron receives one incoming synapse from each hidden neuron. There are no direct connections between input and output neurons. The state of each hidden neuron y_i—with $i \in [1, 10]$—and of each output neuron o_i—with $i \in [1, 3]$—is updated as follows:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_{j=1}^{11} \omega_{ji} I_i + \sum_{k=1}^{10} \omega_{ki} Z(y_k + \beta_k); \quad o_i = \sum_{j=1}^{10} \omega_{ji} Z(y_j + \beta_j); \tag{1}$$

In these equations, τ are the decay constants, ω_{ij} the strength of the synaptic connection from neuron i to neuron j, β the bias terms, and $Z(x) = (1 + e^{-x})^{-1}$ is a sigmoid function. τ, β, and ω_{ij} are genetically

[2]Further methodological details, movies of the post-evaluation tests on real *s-bots* and data not shown in the paper can be found at `http://iridia.ulb.ac.be/supp/IridiaSupp2008-002/`

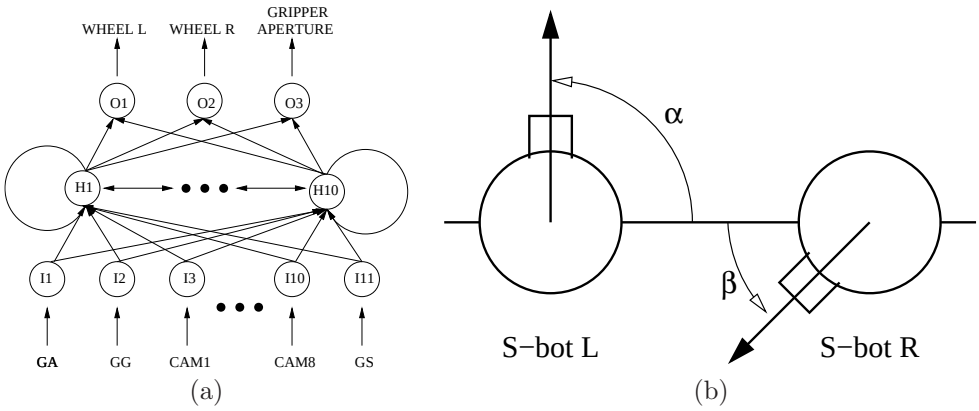

Figure 2: (a) Architecture of the neural network that controls the *s-bots*. (b) This picture shows how the *s-bots*' starting orientations are defined given the orientation duplet (α, β). *S-bot* L and *s-bot* R refer to the robots whose initial orientations in any given trial correspond to the value of α and β respectively.

specified networks' parameters. $Z(o_1)$ and $Z(o_2)$ linearly scaled into [-3.2 cm/s, 3.2 cm/s] are used to set the speed of the left and right motors. $Z(o_3)$ is used to set the gripper aperture in the following way: if $Z(o_3) > 0.75$ the gripper closes; if $Z(o_3) < 0.25$ the gripper opens. Cell potentials are set to 0 when the network is initialised or reset, and circuits are integrated using the forward Euler method with an integration step-size of 0.2.

Each genotype is a vector comprising 263 real values. Initially, a random population of vectors is generated by initialising each component of each genotype to values randomly chosen from a uniform distribution in the range $[-10, 10]$. The population contains 100 genotypes. Generations following the first one are produced by a combination of selection, mutation, and elitism. For each new generation, the five highest scoring individuals from the previous generation are chosen for breeding. The new generations are produced by making twenty copies of each highest scoring individual with mutations applied only to nineteen of them. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.25.

5 The fitness function

During evolution, each genotype is translated into a robot controller, and cloned onto each agent. At the beginning of each trial, two *s-bots* are positioned in a boundless arena at a distance randomly generated in the interval [25 cm, 30 cm], and with predefined initial orientations α and β (see figure 2b). Our initialisation is inspired from the initialisation used in (Quinn, 2001). In particular, we defined a set of orientation duplets (α, β) as all the combinations with repetitions from a set:

$$\Theta_n = \left\{ \frac{2\pi}{n} \cdot i \mid i = 0, \ldots, n-1 \right\}, \tag{2}$$

where n is the cardinality of the set. In other words, we systematically choose the initial orientation of both *s-bots* drawing from the set Θ_n. The cardinality of the set of all the different duplets—where we consider $(\alpha, \beta) \equiv (\beta, \alpha)$—corresponds to the total number of combinations with repetitions, and can be obtained by the following equation:

$$\frac{(n+k-1)!}{k!(n-1)!}, \tag{3}$$

where $k = 2$ indicates that combinations are duplets, and $n = 4$ lets us define the set of possible initial orientations $\Theta_4 = \{0°, 90°, 180°, 270°\}$. From this, we generate 10 different (α, β) duplets. Each group is evaluated 4 times at each of the 10 starting orientation duplets for a total of 40 trials. Each trial (e) differs from the others in the initialisation of the random number generator, which influences the robots initial distance and their orientation by determining the amount of noise added to the orientation duplets (α, β). During a trial, noise affects motors and sensors as well. In particular, uniform noise is added in the range ± 1.25 cm for the distance, and in the range $\pm 1.5°$ for the angle of the coloured blob perceived by the camera. 10% uniform noise is added to the motor outputs $Z(o_i)$. Uniform noise randomly chosen in the range $\pm 5°$ is also added to the initial orientation of each *s-bot*. Within a trial, the robots life-span is 50 simulated seconds (250 simulation cycles), but a trial is also terminated if the robots incur in 20 collisions. In each trial e, each group is rewarded by an evaluation function $F_e = A_e \cdot C_e \cdot S_e$ which seeks to assess the ability of the two robots to get closer to each other and to physically assemble through the gripper.

A_e is the aggregation component, computed as follows:

$$A_e = \begin{cases} \frac{1.0}{1.0 + atan(\frac{d_{rr}-16}{16})} & \text{if } d_{rr} > 16 \text{ cm;} \\ 1.0 & \text{otherwise;} \end{cases} \tag{4}$$

with d_{rr} corresponding to the distance between the two *s-bots* at the end of the trial e;

C_e is the collision component, computed as follows:

$$C_e = \begin{cases} 1.0 & \text{if } n_c = 0; \\ 0.0 & \text{if } n_c > 20; \\ \frac{1.0}{0.5 + \sqrt{n_c}} & \text{otherwise;} \end{cases} \tag{5}$$

with n_c corresponding to the number of robot-robot collisions recorded during trial e;

S_e is the self-assembly component, computed at the end of a trial ($t = T$ with $T \in (0, 250]$), as follows:

$$S_e = \begin{cases} 100.0 & \text{if } GG(T) = 1, \text{ for any robot;} \\ 1.0 + \frac{29.0 \sum_{t=0}^{T} K(t)}{T} & \text{otherwise;} \end{cases} \quad (6)$$

$K(t)$ is set to 1 for each simulation cycle t in which the sensor GS of any *s-bot* is active, otherwise $K(t) = 0$.

Notice that, given the way in which F_e is computed, no assumptions are made concerning which *s-bot* plays the role of *s-bot-gripper* and which one the role of *s-bot-grippee*. The way in which collisions are modelled in simulation and handled by the fitness function is an element that favours the evolution of assembly strategies in which the *s-bot-gripper* moves straight while approaching the *s-bot-grippee* (see section 3). This has been done to ease transferability to real hardware. The fitness assigned to each genotype after evaluation of the robots behaviour is given by

$$FF = \frac{1}{E} \sum_{e=1}^{E} F_e \text{ with } E = 40; \quad (7)$$

6 Results

As stated in section 1, the goal of this research work is to design through evolutionary computation techniques dynamical neural networks to allow a group of two homogeneous *s-bots* to physically connect to each other. To pursue our objective, we run for 10.000 generations twenty randomly seeded evolutionary simulations. Although several evolutionary runs produced genotypes that obtained the highest fitness score (i.e., $FF = 100$, see section 5), the ranking based on the evolutionary performances has not been used to select a suitable controller for the experiments with real robots. The reason for this is that during evolution, the best groups may have taken advantage of favourable conditions, determined by the existence of between-generation variation in the starting positions and relative orientation of the robots and other simulation parameters. Thus, the best evolved genotype from generation 5.000 to generation 10.000 of each evolutionary run has been evaluated again on a series of 136.000 trials, obtained by systematically varying the *s-bots*' starting orientations. In particular, we evaluated the evolved genotypes using a wider set of 16 initial orientations Θ_{16}, defined by equation 2. This set covers all the possible perceptual configurations for the starting condition of one *s-bot*, which may perceive the other *s-bot* through one or two camera sectors (see figure 4 for more details). From this set, equation 3 tells us that we can derive 136 different duplets (α, β). Each starting condition (i.e., orientation duplet) was tested in 1.000 trials, each time randomly choosing the robots' distance from a uniform distribution of values in the range [25 cm, 30 cm]. Noise is added to initial orientations, sensors readings and motor outputs as described in section 5.

The best performing genotype resulting from the set of post-evaluations described above was decoded into an artificial neural network which was then cloned and ported onto two real *s-bots*. In what follows, first we provide the results of post-evaluation tests aimed at evaluating the success rate of the real *s-bots* at the self-assembly task as well as the robustness of the self-assembly strategies in different setups (see section 6.1). Subsequently, we illustrate the results of analyses carried out with simulated *s-bots*, aimed at unveiling operational aspects underlying the best evolved self-assembling strategy (see section 6.2).

6.1 Post-evaluation tests on real *s-bots*

The *s-bots*' controllers are evaluated four times on each of 36 different orientation duplets (α, β), obtained drawing α and β from Θ_8. The cardinality of this set of duplets is given by equation 3, with $n = 8$, $k = 2$. In each post-evaluation experiment, successful trials are considered those by which the robots manage to self-assemble, that is, when one robot manages to grasp the other one. Note that, for real *s-bots*, the trial's termination criteria was changed with respect to those employed with the simulated *s-bots*. We set no limit on the maximum duration of a trial, and no limit on the number of collisions allowed. In each trial, we let the *s-bots* interact until physically connected. In a single case we terminated the trial before the robots self-assembled because the *s-bots* moved so far away from each other that they ended up outside the perceptual range of their respective camera. This trial has been terminated after one minute of robot-robot distance higher than 50 cm and the trial has been considered unsuccessful. As illustrated later in the section, these new criteria allowed us to observe interesting and unexpected behavioural

sequences. In fact, the *s-bots* sporadically committed inaccuracies during their self-assembly manoeuvres. Unexpectedly, the robots demonstrated to possess the required capabilities to autonomously recover from these inaccuracies. In what follows, we provide the reader a detailed description of the performance of the real *s-bots* in these post-evaluation trials.[2]

The first two tests with physical robots are referred to as test G25 and test G30. These are tests in which the *s-bots* light themselves up in green and are initialised at a distance from each other of 25 cm and 30 cm, respectively. The *s-bots* proved to be 100% successful in both tests. That is, they managed to self-assemble in all trials. Table 1 gives more details about the *s-bots*' performances in these trials. In particular, we notice that the number of successful trials at the first gripping attempt is 28 and 29 trials out of 36 respectively for G25 and G30 (see Table 1, 2^{nd} column). In a few trials, the *s-bots* managed to assemble after two/three grasping attempts (see Table 1, 3^{rd} and 7^{th} column). The failed attempts were mostly caused by inaccurate manoeuvres—referred to as inaccuracies of type I_1—, in which a series of maladroit actions by both robots makes impossible for the *s-bot-gripper* to successfully grasp the *s-bot-grippee*'s cylindrical turret. In a few other cases, the group committed a different inaccuracy—referred to as I_2—, in which both robots assume the role of *s-bot-gripper*. In such circumstances, the *s-bots* head towards each other until a collision between their respective grippers occurs. Note that, in both G25 and G30, the *s-bots* always managed to recover from the inaccuracies and end up successful.

As mentioned in section 3, the *s-bots* have to turn on their coloured LEDs in order to perceive each other through the camera. However, as discussed in section 2.1, a significant advantage of our control design approach is that the specific colour displayed has no functional role within the neural machinery that brings forth the *s-bots*' actions. In order to empirically demonstrate that the mechanisms underpinning the *s-bots* self-assembling strategies do not depend on the specific colour displayed by the LEDs, we repeated a third and a fourth time the 36 post-evaluation trials, both times by deliberately changing the colour of the *s-bots*' LEDs. The *s-bots* are placed at an initial distance of 30 cm from each other, and they are evaluated with the LEDs displaying blue light—this test is referred to as B30—and with the LEDs displaying red light—this test is referred to as R30.

The *s-bots* proved to be very successful both in B30 and R30 (see Table 1). In the large majority of the trials the *s-bots* managed to self-assemble at the first grasping attempt. In a few trials, two or three grasping manoeuvres were required (see Table 1, 3^{rd} and 7^{th} column). A new type of inaccuracy emerged in test R30. That is, in three trials, after grasping, the connected structure got slightly elevated at the connection point. We refer to this type of inaccuracy as I_3. Notice also that in a single trial, in test B30, the *s-bots* failed to self-assemble (see Table 1, last column). In this case, the *s-bots* moved so far away from each other that they ended up outside the perceptual range of their respective camera. This trial in which the *s-bots* spent more than 1 minute without perceiving each other has been terminated, and it was considered unsuccessful.

Table 1: Results of post-evaluation tests on real *s-bots*. G25 and G30 refer to the tests in which the *s-bots* light themselves up in green and are initialised at a distance from each other of 25 cm and 30 cm, respectively. B30 and R30 refer to the tests in which the *s-bots* light themselves up in blue and red respectively, and are initialised at a distance of 30 cm from each other. Trials in which the physical connection between the *s-bots* requires more than one gripping attempt, due to inaccurate manoeuvres I_i, are still considered successful. I_1 refers to a series of maladroit actions by both robots which makes impossible for the *s-bot-gripper* to successfully grasp the *s-bot-grippee*'s cylindrical turret. I_2 refers to those circumstances in which both robots assume the role of *s-bot-gripper* and collide at the level of their grippers. I_3 refers to those circumstances in which, after grasping, the connected structure gets slightly elevated at the connection point. Failures correspond to trials in which the robots do not manage to return to a distance from each other smaller than their visual field.

Test	Number of successful trials per gripping attempt and types of inaccuracy									$N.^\circ$ failure
	1^{st}	2^{nd}				3^{rd}				
	$N.^\circ$	$N.^\circ$	I_1	I_2	I_3	$N.^\circ$	I_1	I_2	I_3	
G25	28	7	6	1	0	1	2	0	0	0
G30	29	6	3	3	0	1	1	1	0	0
B30	26	5	3	2	0	4	8	0	0	1
R30	21	12	8	0	2	4	7	0	1	0

For each single test (i.e., G25, G30, B30, and R30), the sequences of *s-bots*' actions are rather different from one trial to the other. However, these different histories of interactions can be succinctly described by a combination of few distinctive phases and transitions between phases which exhaustively "portray" the observed phenomena. Figure 3 shows some snapshots from a successful trial which represent these phases. The robots leave their respective starting positions (see figure 3a) and during the starting phase (see figure 3b) they tend to get closer to each other. In the great majority of the trials, the robots move from the starting phase to what we call the role allocation phase (RA-phase, see figure 3c). In this phase, each *s-bot* tends to remain on the right side of the other. They slowly move by following a circular trajectory corresponding to an imaginary circle centred in between the *s-bots*. Moreover, each robot rhythmically changes its heading by turning left and right. The RA-phase ends once one of the two *s-bots*—that is, the one assuming the role of the *s-bot-gripper*—stops oscillating and heads towards the other *s-bot*—that is, the one assuming the role of the *s-bot-grippee*—which instead orients itself in order to facilitate the gripping (gripping phase, see figure 3d). The *s-bot-gripper* approaches the *s-bot-grippee*'s turret and, as soon as its *GS* sensor is active, it closes its gripper. A successful trial terminates as soon as the two *s-bots* are connected (see figure 3e).

As mentioned above, in a few trials the *s-bots* failed to connect at the first gripping attempt by committing what we called inaccuracies I_1 and I_3. These inaccuracies seem to denote problems in the sensory-motor coordination during grasping. Recovering from I_1 can only be accomplished by returning to a new RA-phase, in which the *s-bots* negotiate again their respective roles, and eventually self-assemble. Recovering from I_3 is accomplished by a slight backward movement of both *s-bots* which restores a stable gripping configuration. Given that I_3 has been observed only in R30, it seems plausible to attribute the origin of this inaccuracy to the effects of the red light on the perceptual apparatus of the *s-bots*. In particular, it could be that, due to the red light, the *s-bot-gripper* perceives through its camera the *s-bot-grippee* at a farther distance than the actual one. Alternatively, it could be that the red light perturbs the regular functioning of the optical barrier and consequently the readings of the *GS* and *GG* sensors. Both phenomena may induce the *s-bot-gripper* to keep on moving towards the *s-bot-grippee* up to the occurrence of I_3, even though the distance between the robots and the status of the gripper of the *s-bot-gripper* would require a different response. I_2 seems to be caused by the effects of the *s-bots*' starting positions on their behaviour. In those trials in which I_2 occurs, after a short starting phase, the *s-bots* head towards each other until they collide with their grippers without going through the RA-phase. The way in which the robots perceive each other at starting positions seems to be the reason why they skip the RA-phase. Without a proper RA-phase, the robots fail to autonomously allocate between themselves the roles required by the self-assembly task (i.e., *s-bot-gripper* and *s-bot-grippee*), and consequently they incur in I_2. In order to recover from I_2, the *s-bots* move away from each other and start a new RA-phase in which roles are eventually allocated. In the future we will further investigate the exact cause of the inaccuracies.

As shown in Table 1, except for a single trial in test B30 in which the *s-bots* failed to self-assemble, the robots proved capable of recovering from all types of inaccuracies. This is an interesting result because it is evidence of the robustness of our controllers with respect to contingencies never encountered during evolution. Indeed, as mentioned in section 3, in order to speed up the evolutionary process, the simulation in which controllers have been designed does not handle collisions with sufficient accuracy. In those cases in which, after a collision, the simulated robots had another chance to assemble, the agents were simply re-positioned at a given distance to each other. In spite of this, *s-bots* guided by the best evolved controllers proved capable of engaging in successful recovering manoeuvres which allowed them to eventually assemble.[2]

(a) (b) (c) (d) (e)

Figure 3: Snapshots from a successful trial. (a) Initial configuration (b) Starting phase (c) Role allocation phase (d) Gripping phase (e) Success (grip)

6.2 An operational description

Our research work illustrates the details of an alternative methodological approach to the design of controllers for self-assembly in autonomous robots in which no assumptions are made concerning how agents allocate roles in the self-assembly task. The evolved mechanisms are as effective as those described in (Groß et al., 2006a; O'Grady et al., 2005). Contrary to modular or hand-coded approaches, the evolutionary one proved to be robust with respect to changes of the colour of the light displayed by the LED. The controllers described in (Groß et al., 2006a; O'Grady et al., 2005) require to be re-structured by the experimenter in order to cope with the same type of changes.

In view of the results shown in section 6.1, we believe that evolved neuro-controllers are a promising approach to the design of mechanisms for autonomous self-assembly. However, it is important to remark that the operational principles of self-assembly used by the *s-bots*, controlled by this type of neural structures, are less "transparent" than the modular or hand-coded control described in (Groß et al., 2006a; O'Grady et al., 2005). Further research work and experimental analysis are required to unveil the operational principles of the evolved neural controllers. What are the strategies that the *s-bots* use to carry out the self-assembly task? How do they decide who is the *s-bot-gripper*, and who is the *s-bot-grippee*? Although extremely interesting, providing an answer to this type of questions is not always a simple task. The methodologies we have in order to look for the operational mechanisms of evolved neural networks are limited to distributed systems with a small number of neurons, or to cases in which the neural networks control simple agents that can only move in a one-dimension world, or by discrete steps (see Beer, 2003, 2006; Keinan et al., 2006, for details). Due to the nature of our system, most of these methods cannot be directly employed to investigate which mechanisms control the process by which two homogeneous *s-bots* differentiate into *s-bot-gripper* and *s-bot-grippee*. In spite of these difficulties, in this section we describe the results of an initial series of studies focused on the relationship between the *s-bots*' starting orientations and the role allocation process.

Do the robots' orientations at the beginning of a trial influence the way in which roles (i.e., *s-bot-gripper* versus *s-bot-grippee*) are allocated? We start our analysis by looking at the results of the post-evaluation tests mentioned at the beginning of section 6. In particular, we look at those data concerning the behaviour of the *s-bots* controlled by the best performing genotype; that is, the genotype used to build the networks ported on the real robots. Recall that, in these tests, the simulated *s-bots* have been evaluated on a series of 136 starting orientation duplets (α, β) obtained from Θ_{16}. For each orientation duplet the *s-bots* underwent 1.000 evaluation trials, each time randomly choosing the agents' distance from a uniform distribution of values in the range [25 cm, 30 cm]. Each duplet (α, β) defines a perceptual scenario at the beginning of a trial characterised by the sector/s through which the robots perceive each other. We defined two subsets of Θ_{16}: Θ_A and Θ_B. These sets encompass those initial orientations that correspond to a perception of the other robot through respectively two camera sectors and one camera

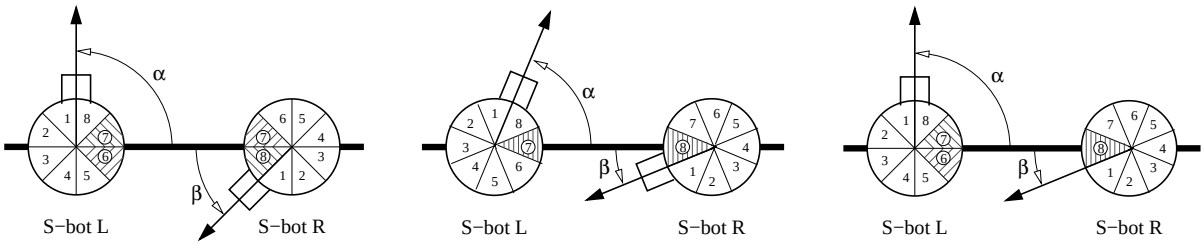

$\alpha, \beta \in \Theta_A$,
$\Theta_A = \{45° \cdot i \mid i = 0, \ldots, 7\}$
(a)

$\alpha, \beta \in \Theta_B$,
$\Theta_B = \{45° \cdot i + 22.5°$
$\mid i = 0, \ldots, 7\}$
(b)

$\alpha \in \Theta_A, \beta \in \Theta_B$,
OR
$\alpha \in \Theta_B, \beta \in \Theta_A$
(c)

Figure 4: Depiction of three different *s-bots*' starting conditions. In each picture, circles represent the *s-bots*, thick arrows indicate the robots' headings and the thin arrows their orientation. The numbers within the circles refer to the camera sectors CAM_i with $i \in [1, 8]$. Filled sectors are those through which the *s-bots* perceive each other. (a) In those trials in which α and β are drawn from the set of starting orientations Θ_A the robots perceive each other in two camera sectors. (b) In those trials in which α and β are drawn from the set of starting orientations Θ_B the robots perceive each other in one camera sector. (c) In those trials in which $\alpha \in \Theta_i$ and $\beta \in \Theta_j$, with $i \neq j$, the *s-bots* perceive each other in one and two camera sectors.

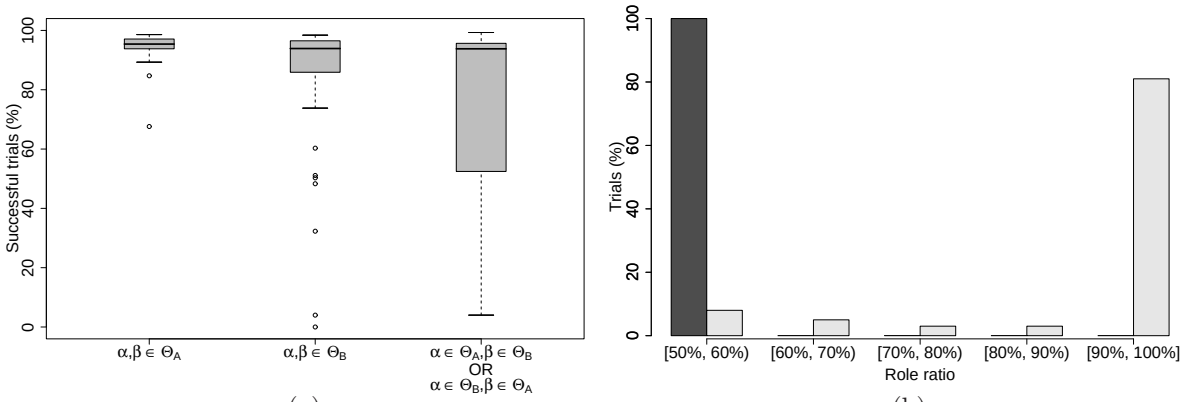

(a) (b)

Figure 5: (a) The graph shows the percentage of successful trials of simulated *s-bots* controlled by the best evolved genotype for each of the three types of starting conditions: (1) α, $\beta \in \Theta_A$, see figure 4a; (2) α, $\beta \in \Theta_B$, see figure 4b; (3) $\alpha \in \Theta_i$, $\beta \in \Theta_j$, and $i \neq j$, see figure 4c. Boxes represent the inter-quartile range of the data, while the horizontal bars inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. Empty circles mark the outliers. (b) The graph shows the percentage of trials per level of *role ratio*. Black bars refer to orientation duplets (α, β) with $\alpha = \beta$, referred to as symmetrical orientation duplets. Grey bars refer to orientation duplets (α, β) with $\alpha \neq \beta$, referred to as asymmetrical orientation duplets.

sector. As a consequence, each duplet (α, β) may identify one out of three different conditions: $\alpha, \beta \in \Theta_A$ (see figure 4a), $\alpha, \beta \in \Theta_B$ (see figure 4b), or $\alpha \in \Theta_i, \beta \in \Theta_j$ and $i \neq j$ (see figure 4c). Note that, given the distance up to which a coloured blob can be perceived by the *s-bots*' camera (approximately 50 cm), and the dimension of each camera sector (45°), the three categories illustrated in figure 4 take into account all the possible perceptual scenarios that the two *s-bots* system can experience at the beginning of a trial, assuming that the robots are at less than 50 cm from each other. The 136 orientation duplets include 16 symmetrical conditions in which $\alpha = \beta$. In symmetrical orientation duplets, the robots share the same perception at the beginning of the trial. That is, they perceive each other through the same sector/s of their corresponding camera. Asymmetrical orientation duplets are those in which $\alpha \neq \beta$. These tests have been repeated twice: once without adding any noise to the robots' orientations (α and β) and once by applying a random offset to the robots' orientations. These offsets are chosen in order not to disrupt the perceptual scenario—i.e., the sector/s through which the robots perceive each other—determined by the corresponding orientation duplet (α, β). The results are qualitatively similar, therefore in the following we discuss only the results of the tests with noise. Recall that, contrary to the real *s-bots*, the simulated robots, due to the way our simulator handles collisions, are not allowed to use any recovery manoeuvres. That is, in these post-evaluation tests, the simulated *s-bots* are scored according to a binary criterion: a trial can be either successful or unsuccessful. Unsuccessful trials are considered those in which the robots did not manage to self-assemble within the time-limit, as well as those that terminated due to the occurrence of collisions that are not considered the result of an accepted grasping manoeuvre (see section 3 for details).

Figure 5a refers to the success rate of the simulated *s-bots*, controlled by the best evolved genotype. The three boxes correspond to the three initial perceptual scenarios depicted in figure 4. We notice that the medians of the three distributions are all over 90%, which indicates that the robots are quite successful in all scenarios. The 2^{nd} and 3^{rd} scenario include some trials where the performance is rather low. However, we cannot conclude that a particular initial perception scenario (e.g., perceiving a robot through one or two camera sectors) has a direct impact on the success rate. Before proceeding with our analysis of the mechanisms for role allocation we need to clarify a few issues. First, the reader should bear in mind that, in our tests, the *s-bot-gripper* is operationally defined as the robot that successfully grips the other one. This definition has been chosen for being the most reliable in discriminating the roles. Second, given the way in which α and β are varied in our tests, the operational definition of the roles, and the *s-bots* success rate (see figure 5a), we can already exclude that the system works by following simple rules by which the role is determined by the initial individual perception. This is because the robots proved to be successful even in symmetrical trials. In other words, having the same initial

perception does not hinder the robots from allocating different roles. Therefore, either the system has to be governed by more complex principles based on the combination of α and β, or the initial orientations have no bearing on the role allocation process. In the remainder of this section, we carry out an analysis that helps us further clarify this issue.

For our role allocation analysis we chose only the orientation duplets for which the robots self-assembled in 85% or more of the trials (104 out of 136 duplets). For the remaining 32 orientation duplets—1 symmetrical and 31 asymmetrical— for which the robots self-assembled in less than 85% of the trials, we considered the sample size to be too small for a meaningful role allocation analysis[3]. By looking at the frequency (i.e., how many trials out of 1.000) with which each *s-bot* (i.e., *s-bot* L and *s-bot* R) plays the role of *s-bot-gripper* for any given value of α and β, our analysis is intended to unveil any relationship between the robots initial orientations and the role they assume during the trial. In particular, we looked at the *role ratio*.

The *role ratio* can be considered a property of each orientation duplet. It indicates how often a given robot (i.e., *s-bot* L or *s-bot* R) played the role of *s-bot-gripper* when repeatedly evaluated on a given orientation duplet. In particular, the *role ratio* corresponds to the highest frequency of playing the *s-bot-gripper* role between the one recorded by *s-bot* L and by *s-bot* R. Thus, the *role ratio* can vary between 50%, when both robots played the *s-bot-gripper* role with the same frequency, to 100% when only one robot plays the *s-bot-gripper* role in all the trials that start with the same perceptual scenario (i.e., evaluation trials repeated for a given orientation duplet). If the *role ratio* is around 50% for both symmetrical and asymmetrical trials, then this is an evidence that the initial individual orientations have no bearing on the role allocation process. If instead, for certain orientation duplets (α, β) the *role ratio* diverges from the 50% value, then we conclude that the system is governed by principles based on the combination of the robots initial individual orientations.

In figure 5b, the *role ratio* is divided in five categories, represented on the x-axis, and the orientation duplets are divided in two categories represented by the two types of bars. The black bars refer to the percentage of symmetrical orientation duplets for each category of *role ratio*. The grey bars refer to the percentage of asymmetrical orientation duplets for each category of *role ratio*. By looking at figure 5b, we clearly see that while the totality of the orientation duplets corresponding to symmetrical starting positions falls into the category of *role ratio* [50%, 60%), the large majority of the orientation duplets corresponding to asymmetrical starting positions falls into the category of *role ratio* [90%, 100%]. This means that, while in many of the asymmetrical trials the role of *s-bot-gripper* is for the large majority (if not the totality) of the successful trials played by the same robot (i.e., either *s-bot* L or *s-bot* R), in all the symmetrical trials both robots play the role of *s-bot-gripper* with more or less the same frequency. We also found out that in asymmetrical trials with $\alpha = 0°$, it is the *s-bot* L that systematically plays the role of *s-bot-gripper* (data not shown[2]). For all the other asymmetrical starting conditions with $\alpha \neq 0°$, the role of *s-bot* L depends on the value of β. That is, except for the orientation 0°, for all the other duplets, no associations can be made between an *s-bot*'s initial orientation and its role. In these circumstances (i.e., $\alpha, \beta \neq 0°$), it seems to be the combination of values of α and β which determines whether it is *s-bot* L or *s-bot* R playing the *s-bot-gripper* role.

To summarise, our analysis revealed that, who is the *s-bot-gripper* and who is the *s-bot-grippee* is purely the result of an autonomous negotiation phase between the two *s-bots*. The role allocation unfolds in time during the entire duration of a trial. However, in the case in which the two robots have different initial perceptions, the role that each *s-bot* assumes can be predicted knowing the combination of α and β. This means that it is this combination which determines the roles. In other words, perceiving the other robot at a specific distance and through a given camera sector does not inform a robot about the role it assumes during the trial. In those cases in which the robots start with an identical perception, this symmetry does not seem to hinder the robots from autonomously allocating different roles to successfully accomplish their goal. At the moment, it is unclear how the initial symmetry is broken. Perhaps, the driving forces have to be searched in the way in which the robots mutually affect each other behaviour. Perhaps, the random noise injected into the system is the causal factor that drives the system through sequences of actions that turn out to be successful. Stochastic phenomena may take over any causal relationship between environmental structures (i.e., how the robots perceive each other at the beginning of a trial) and the role allocation process. Future analyses are certainly required to see whether any invariants can be found among the history of interactions between the robots and what significance can be attributed to them.

[3]This is mainly due to the operational definition of *s-bot-gripper*. Since there is no *s-bot-gripper* in unsuccessful trials, the lower the *s-bots* success rate the smaller the sample (i.e., the number of trials) in which the relationship between roles and starting perceptual scenario can be observed.

7 Discussion

In a context free of assumptions concerning the nature of the mechanisms underlying the agents behavioural repertoire, our evolutionary robotics model exploits an automatic design process which mimics the mechanisms of natural evolution to define the control structures which allow the robots to autonomously self-assemble by playing complementary roles (i.e., *s-bot-gripper* and *s-bot-grippee*). The results of post-evaluation analyses shown in section 6.2 illustrate that the allocation of roles is the result of an autonomous negotiation phase between the two robots. The outcome of any action an agent chooses depends on the action of the other agent. In other words, none of the two agents can know the role it will assume at the end of the trial, judging only from its initial perception.

We have shown on real hardware that explicit communication to directly access the "intention" of the other agent (through explicit signals, as the ones used in Groß et al., 2006a, for example) is not a necessary condition for coordination. Our robots coordinate without direct and explicit communication. Noble (1998) reached a similar conclusion with an evolutionary simulation model involving two simulated animals contesting the possession of a resource. Groß and Dorigo (2008a) have also concluded that cooperative behaviour can be achieved without explicit means of communication. More specifically, in a cooperative transport task, simulated robots could find effective transport strategies exploiting indirect communication, that is, by interacting with each other through the object to be transported. Finally, our results are very similar to the results obtained in Quinn (2001); Quinn et al. (2003), where role allocation (leader-follower) and formation movement is achieved solely through infrared sensors and the control structure is once again an evolved dynamic neural network. In particular, the work presented in Quinn (2001) reports on role allocation between two robots for symmetrical and non-symmetrical cases. Whilst the author qualitatively explains how the difference in the initial perceptions influences the role allocation for non-symmetrical cases, an analysis of the evolved behaviour in case of "insufficient differences" is not performed. In the analysis performed in section 6.2, we have tried to explain quantitatively and qualitatively the effect of the starting configuration on the final outcome of a trial (how roles are allocated); the great majority of non-symmetrical configurations severely bias the role allocation process, while symmetrical configurations may be governed by stochastic phenomena that take over causal relationships between initial conditions and the final role allocation.

An important parameter for the evolutionary processes is the robots initialisation. Our choice aimed to evolve a system that can cope with all possible orientation duplets. Altering the proportion of symmetrical and asymmetrical orientation duplets experienced throughout evolution might have an impact on the evolved role allocation strategies. For example, it is possible that presenting evolution with more symmetrical examples can lead to the prevalence of strategies where the role ratio is around 50%. We believe that this is a very important and interesting issue to be considered in future work.

8 Conclusion

In this article, we have presented the results of an evolutionary methodology for the design of control strategies for self-assembling robots. More specifically, to the best of our knowledge, the control method we have proposed for the physical connection of two robots is the only existing in the literature where the role allocation between gripper and grippee is the result of an autonomous negotiation phase between the homogeneous robots; there is no a priori injected behavioural or morphological heterogeneity in the system. Instead, the behavioural heterogeneity emerges through the interaction of the robots. Moreover, the communication requirements of our approach are reduced to the minimum; simple coordination by means of the dynamical interaction between the robots—as opposed to explicit communication of internal states—is enough to bring forth differentiation within the group. We believe that reducing the assumptions on necessary conditions for assembly is an important step to obtain more adaptive and more general controllers for autonomous self-assembly.

The results of this work are a proof-of-concept: they proved that dynamical neural networks shaped by evolutionary computation techniques directly controlling the robots' actuators can provide physical robots all the required mechanisms to autonomously perform self-assembly. Contrary to the modular or hand-coded controllers described in Groß et al. (2006a); O'Grady et al. (2005), the evolutionary robotics approach did not require the experimenter to make any a priori assumptions concerning the roles of the robots during self-assembly (i.e., either *s-bot-gripper* or *s-bot-grippee*) or about their status (e.g., either capable of moving or required not to move). The evolved mechanisms proved to be robust with respect to changes in the colour of the light displayed by the LEDs. Furthermore, in section 6.1 we have presented a system that exhibits recovery capabilities that could not be observed during the artificial evolution and

that were not coded or foreseen by the experimenter. Such a feature in our case comes for free, while in the case of Groß et al. (2006a) a recovery mechanism had to be designed as a specific behavioural module to be activated every time the robots failed to achieve assembly.

As mentioned in previous sections, our system is not as "transparent" as a hand-coded control system is, as we cannot break its behaviour down to a set of rules or states. Such an endeavour seems to be very challenging and particularly difficult, especially when the network sizes are large and/or the movement of the robots takes place in a continuous and noisy world, such as the real world. However, we would like to stress that we do not consider this step a necessary precondition for the success of research work using Evolutionary Robotics as a design methodology. Our view is that it is more important to identify those choices that made the implementation and experimentation successful. In other words, we put the stress on better understanding which principles make the evolutionary machinery able to produce efficient rules to guide groups of robots, than on identifying each and every one of these rules.

Future work will focus on mainly two issues. Firstly, we want to test the scalability of our system. Can the controllers still manage to achieve assembly if there are more than two robots involved? Some initial experimentation[2] looks very promising. However, we plan to introduce coordinated motion capabilities to the robots behavioural repertoire before we systematically address this issue. In other words, the assembled structure of two or more robots must be able to move coordinately, in order to actively participate in the assembly process. For example, it could interact with other assembled structures or individual robots by either receiving connections from them or grasping them. Secondly, we will study more complex scenarios in which self-assembly is functional to the achievement of particular objectives that are beyond the capabilities of a single robot.

Acknowledgements

The authors would like to thank: Dr. Francisco Santos, Dr. Roderich Groß and Marco Montes de Oca for comments, recommendations and fruitful discussion during the production of the manuscript.

E. Tuci and M. Dorigo acknowledge European Commission support via the *ECAgents* project, funded by the Future and Emerging Technologies programme (grant IST-1940). M. Dorigo acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Research Director. M. Dorigo and C. Ampatzis acknowledge support from the "ANTS" project, an "Action de Recherche Concertée" funded by the Scientific Research Directorate of the French Community of Belgium. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

References

Anderson, C., Theraulaz, G., and Deneubourg., J. (2002). Self-assemblages in insect societies. *Insectes Sociaux*, 49(2):99–110.

Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243.

Beer, R. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18:3009–3051.

Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1:91–122.

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY.

Brown, H., Weghe, J. V., Bererton, C., and Khosla, P. (2002). Millibot trains for enhanced mobility. *IEEE/ASME Trans. Mechatron.*, 7:452–461.

Castano, A., Shen, W., and Will, P. (2000). CONRO: Towards deployable robots with inter-robot metamorphic capabilities. *Auton. Robots*, 8:309–324.

Christensen, A. (2005). Efficient neuro-evolution of hole-avoidance and phototaxis for a swarm-bot. DEA thesis TR/IRIDIA/2005-14, Université Libre de Bruxelles, Bruxelles, Belgium.

Damoto, R., Kawakami, A., and Hirose, S. (2001). Study of super-mechano colony: concept and basic experimental set-up. *Advanced Robotics*, 15(4):391–408.

Dorigo, M. (2005). Swarm-bot: A novel type of self-assembling robot. In Murase, K., Sekiyama, K., Kubota, N., Naniwa, T., and Sitte, J., editors, *Proceedings of the 3rd International Symposium on Autonomous Minirobots for Research and Edutainment (AMiRE 2005)*, pages 3–4. Springer-Verlag, Berlin, Germany.

Fukuda, T. and Nakagawa, S. (1987). A dynamically reconfigurable robotic system (concept of a system and optimal configurations). In *Proc. of the 1987 IEEE Int. Conf. on Industrial Electronics, Control and Instrumentation*, pages 588–595. IEEE Computer Society Press, Los Alamitos, CA.

Fukuda, T., Nakagawa, S., Kawauchi, Y., and Buss, M. (1988). Self organizing robots based on cell structures - CEBOT. In *Proc. of the 1988 IEEE Int. Workshop on Intelligent Robots*, pages 145–150. IEEE Computer Society Press, Los Alamitos, CA.

Fukuda, T. and Ueyama, T. (1994). *Cellular Robotics and Micro Robotic Systems*. World Scientific Publishing, London, UK.

Groß, R., Bonani, M., Mondada, F., and Dorigo, M. (2006a). Autonomous self-assembly in swarm-bots. *IEEE Transactions on Robotics*, 22(6):1115–1130.

Groß, R. and Dorigo, M. (2008a). Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling. *Adaptive Behavior*. in press.

Groß, R. and Dorigo, M. (2008b). Self-assembly at the macroscopic scale. *Proceedings of the IEEE*. Accepted for publication.

Groß, R., Dorigo, M., and Yamakita, M. (2006b). Self-assembly of mobile robots—from swarm-bot to super-mechano colony. In *Proc. of the* 9^{th} *Int. Conf. on Intelligent Autonomous Systems*, pages 487–496. IOS Press, Amsterdam, The Netherlands.

Hirose, S. (2001). Super mechano-system: New perspectives for versatile robotic systems. In Rus, D. and Singh, S., editors, *Proc. of the 7th Int. Symposium on Experimental Robotics, (ISER)*, volume 271 of *Lecture Notes in Control and Information Sciences*, pages 249–258. Springer, Berlin, Germany.

Hirose, S., Shirasu, T., and Fukushima, E. (1996). Proposal for cooperative robot "Gunryu" composed of autonomous segments. *Robotics and Autonomous Systems*, 17:107–118.

Hölldobler, B. and Wilson, E. O. (1978). The multiple recruitment systems of the african weaver ant, oecophylla longinoda (latreille) (hymenoptera: Formicidae). *Behavioural Ecology and Sociobiology*, 3:19–60.

Izzo, D. and Pettazzi, L. (2007). Autonomous and distributed motion planning for satellite swarm. *Journal of Guidance Control and Dynamics*, 30(2):449–459.

Izzo, D., Pettazzi, L., and Ayre, M. (2005). Mission concept for autonomous on orbit assembly of a large reflector in space. In *56th International Astronautical Congress*. Paper IAC-05-D1.4.03.

Jakobi, N. (1997). Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6:325–368.

Keinan, A., Sandbank, B., Hilgetag, C., Meilijson, I., and Ruppin, E. (2006). Axiomatic scalable neurocontroller analysis via the shapley value. *Artificial Life*, 12:333–352.

Mondada, F., Gambardella, L. M., Floreano, D., Nolfi, S., Deneubourg, J.-L., and Dorigo, M. (2005). The cooperation of swarm-bots: Physical interactions in collective robotics. *IEEE Robotics & Automation Magazine*, 12(2):21–28.

Mondada, F., Pettinaro, G., Guignard, A., Kwee, I., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L., and Dorigo, M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221.

Noble, J. (1998). Tough guys don't dance: intention movements and the evolution of signalling in animal contests. In *Proceedings of the fifth international conference on simulation of adaptive behavior on From animals to animats 5*, pages 471–476, Cambridge, MA. MIT Press.

Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA.

O'Grady, R., Groß, R., Mondada, F., Bonani, M., and Dorigo, M. (2005). Self-assembly on demand in a group of physical autonomous mobile robots navigating rough terrain. In Capcarrere, M., Freitas, A., Bentley, P., Johnson, C., and Timmis, J., editors, *Proceedings of the* 8^{th} *European Conference on Artifcial Life (ECAL05)*, number 3630 in LNCS, pages 272–281. Springer Verlag, Berlin, Germany.

Quinn, M. (2001). Evolving communication without dedicated communication channels. In *Advances in Artificial Life: Sixth European Conference on Artificial Life (ECAL01)*, Prague, Czech Republic.

Quinn, M., Smith, L., Mayley, G., and 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, Series A: Mathematical, Physical and Engineering Sciences*, 361:2321–2344.

Rubenstein, M., Payne, K., Will, P., and Shen, W.-M. (2004). Docking among indepenpent and autonomous CONRO self-reconfigurable robots. In *Proc. of 2004 IEEE Int. Conf. on Robotics and Automation (ICRA'04)*, volume 3, pages 2877–2882. IEEE Computer Society Press, Los Alamitos, CA.

Tuci, E., Ampatzis, C., Vicentini, F., and Dorigo, M. (2008). Evolving homogeneous neuro-controllers for a group of heterogeneous robots: coordinated motion, cooperation, and communication. *Artificial Life*. In press.

Tuci, E., Groß, R., Trianni, V., Bonani, M., Mondada, F., and Dorigo, M. (2006). Cooperation through self-assembling in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):115–150.

Whitesides, G. and Grzybowski, B. (2002). Self-Assembly at All Scales. *Science*, 295:2418 – 2421.

Yamakita, M., Taniguchi, Y., and Shukuya, Y. (2003). Analysis of formation control of cooperative transportation of mother ship by smc. In *Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA'03)*, volume 1, pages 951–956.

Yim, M., Zhang, Y., and Duff, D. (2002a). Modular robots. *IEEE Spectrum*, 39(2):30–34.

Yim, M., Zhang, Y., Roufas, K., Duff, D., and Eldershaw, C. (2002b). Connecting and disconnecting for chain self-reconfiguration with polybot. *IEEE/ASME Transactions on Mechatronics*, 7(4):442–451.