# Self-Assembly in Physical Autonomous Robots:
## the Evolutionary Robotics Approach

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

[†]IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
[††]ISTC-CNR, Rome, Italy
[†††]DCTI-ISCTE, Lisbon, Portugal
{etuci,campatzi,mdorigo}@ulb.ac.be, vito.trianni@istc.cnr.it, anders.christensen@iscte.pt

## Abstract

This research work illustrates the details of a methodological approach to the design of homogeneous neuro-controllers for self-assembly in physical autonomous robots in which no assumptions are made concerning how agents allocate roles. Artificial evolution is used to set the parameters of a dynamical neural network that when ported on two physical robots allows them to coordinate their actions in order to decide who will grip whom. The neural network directly controls the state of all the actuators. To the best of our knowledge, this work is the first example in which physical robots manage to self-assemble without relying on a priori injected morphological or behavioural heterogeneities. The results shed a light on the minimal requirements necessary to achieve self-assembly in autonomous robots.

## 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 (e.g., defence, object transport, passage formation, etc.; see Anderson et al., 2002). Ants of the species *Œcophylla longinoda* can form chains composed of their own bodies which are used to pull 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). Although ubiquitous in nature, self-assembly remains in general a phenomenon whose operational principles are not easy to grasp, both in non-living and living organisms, at any scale. This is because "it is impractical to change many of the parameters that determine the behaviour of the system components" (see Whitesides and Grzybowski, 2002). However, self-assembly is particularly appealing to various scientific disciplines. For example, understanding the mechanisms of self-assembly in the cell may provide further insights into the emergence of life starting from chemical reactions. From an engineering point of view, understanding self-assembly may inspire the design of artificial self-assembling components. 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).

Building artificial models that capture the main properties of natural phenomena can provide the means to formulate and test hypotheses concerning the underlying mechanisms of the observed phenomena (see Webb, 2000). Several examples of robotic platforms in the literature consist of connecting modules[1]. Among the various autonomous self-assembling systems that have been proposed in the literature, the work done by Groß et al. (2006) using the robots called *s-bot* is particularly relevant to the subject of our study. Groß et al. (2006) 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 the assembly speed. A significant contribution of Groß et al. work is in the design of distributed control mechanisms for self-assembly relying only on local perception. In particular, self-assembly is accomplished with a modular approach in which some modules are evolved and others handcrafted. The approach is based upon a signalling system which makes use of colours. For example, the decision concerning which robot makes the action of gripping (i.e., the *s-bot-gripper*) and which one is gripped (i.e., the *s-bot-grippee*) is made through the emission of colour signals, according to which the robots emitting blue light are playing the role of *s-bot-grippers* and those emitting red light the role of *s-bot-grippees*. 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 robot "born" red among several

---

[1]The reader can find comprehensive reviews of the work on autonomous self-assembling systems in (Yim et al., 2002; Groß and Dorigo, 2008b; Tuci et al., 2006).

robots "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 robot, each blue light emitting robot is 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 robots 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 an object 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 in (Groß et al., 2006) and in (O'Grady et al., 2005) 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, the modularised control architecture used in these works to allow the robots to self-assemble 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., 2006, 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.

In this work we aim at designing control structures by which the self-assembly mechanisms do not rely on a priori designer-specified morphological or behavioural differences between the robots, and the individual behaviours are not triggered by a priori designer-specified agents' perceptual cues. To accomplish our objective we exploit the properties of a particular type of design method referred to as Evolutionary Robotics (ER). ER is a methodological tool to automate the design of robots' controllers 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 (Nolfi and Floreano, 2000). With respect to other design methods, ER provides the methodological tools to generate control structures for artificial agents such as autonomous robots, in a relatively prejudice-free fashion. For example, 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 the fitness (i.e., a measure of an agent's or group's ability to accomplish its task).

In this work, we study self-assembly in a setup where the robots interact and eventually differentiate by allocating distinct roles (i.e., *s-bot-gripper* versus *s-bot-grippee*). In particular, two identical robots, placed in a boundless arena at 25/30 cm from each other with random orientations, are required to approach each other and to self-assemble; that is, a robot should physically connect to the other one via a gripper. 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 show that an integrated (i.e., non-modularised) dynamical 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. 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., 2006). Coordination and role allocation in our system is achieved solely through minimal sensory information and without explicit communication. Groß and Dorigo (2008a) have reached a similar conclusion in a cooperative transport task with simulated robots. 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) 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., 2006; O'Grady et al., 2005) when controlling two physical robots.

## Simulated and Real *S-bot*

An *s-bot* is a mobile autonomous robot equipped with many sensors useful for the perception of the surrounding environment or for proprioception, a differential drive system, and a gripper by which it can grasp various objects or another *s-bot* (see Figure 1a, and Mondada et al., 2004, for further

details on the robot). The main body is a cylindrical turret with a diameter of 11.6 cm, which can be actively rotated with respect to the chassis.

In this work, to allow robots to perceive each other, we make use of the omni-directional camera. 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 the eight $45°$ sectors. Each sector is referred to as $C_i$ with $i \in [1, 8]$. 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 its turret. Notice that the camera can clearly perceive coloured blobs up to a distance of approximately 50 cm, but the precision above 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 valuable information concerning 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 anytime a robot has gripped 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 0 when the gripper is completely closed to 1 when the gripper is completely open.

The controllers are evolved in a simulation environment which models some of the hardware characteristics of the real *s-bots*. The simulator used is based on a specialized 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 inte-
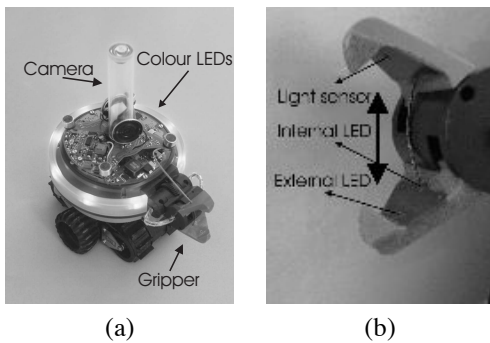


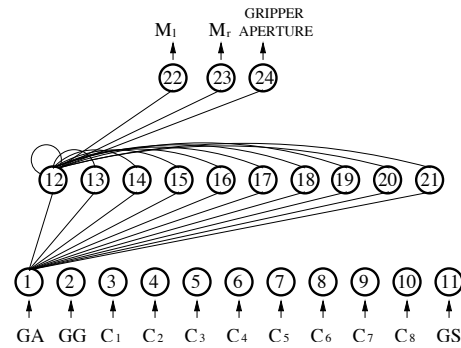Figure 1: (a) The s-bot. (b) The gripper and sensors of the optical barrier.



Figure 2: Neural network architecture: only the efferent connections of the first node of each layer are drawn. See text for the meaning of the labels.

gral to the task (e.g., the closing of the gripper around an object can be seen 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. 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.

## 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 2 and also Beer and Gallagher, 1992). At each simulation cycle, the activation values $y_i$ of input neurons correspond to: the reading of the *GA* sensor for $i = 1$; the reading of the *GG* sensor for $i = 2$; the normalised reading of the eight camera sectors $C_j$ with $j \in [1, 8]$ for $i \in [3, 10]$; the reading of the *GS* sensor for $i = 11$. 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 and output neuron is updated as follows:

$$\tau_i \dot{y}_i = \begin{cases} \sum_{j=1}^{11} \omega_{ji} y_j + \sum_{k=12}^{21} \omega_{ki} \sigma(y_k + \beta_k) - y_i; & i \in [12, 21] \\ \sum_{j=12}^{21} \omega_{ji} \sigma(y_j + \beta_j) - y_i; & i \in [22, 24]; \end{cases}$$

with $\sigma(x) = \dfrac{1}{1 + e^{-x}}$

$$(1)$$

In these equations, $\tau_i$ is the decay constant, $\omega_{ij}$ the strength of the synaptic connection from neuron $j$ to neuron $i$, $\beta_i$ the bias term. $\tau_i$ with $i \in [12, 24]$, $\beta_i$ with $i \in [12, 24]$, and all the network connection weights $\omega_{ij}$ are genetically specified networks' parameters. $\sigma(y_{22})$ and $\sigma(y_{23})$ linearly scaled into $[-3.2s^{-1}, 3.2s^{-1}]$ are used to set the speed of the left and right motors ($M_l$, and $M_r$). $\sigma(y_{24})$ is used to set the gripper aperture in the following way: if $\sigma(y_{24}) > 0.75$ the gripper closes; if $\sigma(y_{24}) < 0.25$ 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 $\Delta T = 0.2$.

Each genotype is a vector comprising 263 real values (i.e., 240 genes for the weights, 13 genes for the time constants, 10 genes for the biases). 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 [0,1]. 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. Genotype parameters are linearly mapped to produce CTRNN parameters with the following ranges: biases $\beta_i \in [-10, 10]$ and weights $\omega_{ji} \in [-10, 10]$. Decay constants are firstly linearly mapped onto the range $[-1.0, 1.5]$ and then exponentially mapped into $\tau_i \in [10^{-1.0}\ 10^{1.5}]$.
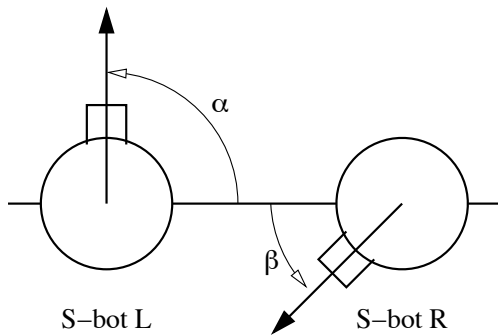


Figure 3: This picture shows how the *s-bots*' starting orientations are defined given the orientation duplet $(\alpha, \beta)$. *S-bot* L and *s-bot* R refer to the robots whose initial orientations in any given trial correspond to the value of $\alpha$ and $\beta$, respectively.

## The Task and 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 $\alpha$ and $\beta$ (see Figure 3). Our initialisation is inspired from the initialisation used in (Quinn, 2001). In particular, we define a set of orientation duplets $(\alpha, \beta)$ as all the combinations with repetitions from the set:

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

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 $\Theta_n$. The cardinality of the set of all the different duplets—where $(\alpha, \beta) \equiv (\beta, \alpha)$—corresponds to the total number of combinations with repetitions, and can be obtained by the following formula:

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

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 $(\alpha, \beta)$ 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 $(\alpha, \beta)$. During a trial, noise affects motors and sensors as well. In particular, uniform noise is added in the range $\pm1.25$ cm for the distance, and in the range $\pm1.5°$ for the angle of the object perceived by the camera. Note that, in simulation, colours are not considered. The camera returns distances and angles of the closest object in each sector. 10% uniform noise is added to the motor outputs $\sigma(y_{22})$, $\sigma(y_{23})$. Uniform noise randomly chosen in the range $\pm5°$ 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). A trial can be terminated earlier if the robots successfully self-assemble in less than 50 simulated seconds, or if they 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:

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

where $d_{rr}$ is the distance between the two *s-bots* at the end of the trial $e$.

$C_e$ is the collision component, computed as follows:

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

where $n_c$ is 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:

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

$(6)$

where $K(t)$ 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. This has been done to ease transferability to real hardware. The fitness assigned to each genotype after the evaluation of the robots is given by $FF = \frac{1}{E} \sum_{e=1}^{E} F_e$, with $E = 40$.

## Results

As stated in the Introduction, in this work we aim at designing 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., $FF = 100$), 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 $\Theta_{16}$, defined by equation 2. From this set, equation 3 tells us that we can derive 136 different duplets $(\alpha, \beta)$. 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.

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, 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.

### Post-evaluation Tests on Real *S-bots*

The $s$-bots' controllers are evaluated four times on each of 36 different orientation duplets $(\alpha, \beta)$, obtained drawing $\alpha$ and $\beta$ from $\Theta_8$. 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 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 this 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 show 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.

The *s-bots* have to turn on their coloured LEDs in order to perceive each other through the camera. However, 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

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 assembly 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 hinder the *s-bot-gripper* from successfully grasping the *s-bot-grippee*'s 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.

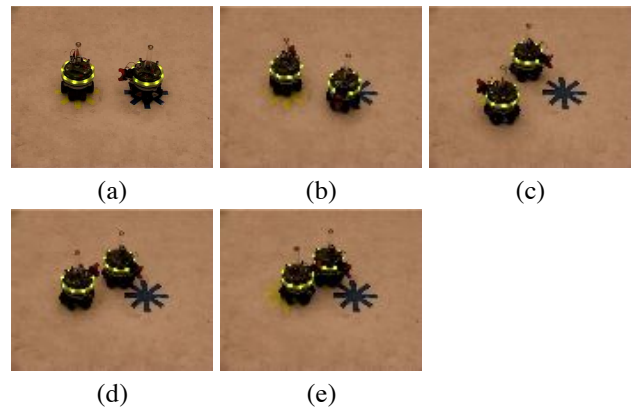| Test | Number of successful trials per gripping attempt and types of inaccuracy | | | | | | | | |
|------|------|------|-------|-------|-------|------|-------|-------|-------|
| | $1^{st}$ | $2^{nd}$ | | | | $3^{rd}$ | | | |
| | N.° | N.° | $I_1$ | $I_2$ | $I_3$ | N.° | $I_1$ | $I_2$ | $I_3$ |
| G25 | 28 | 7 | 6 | 1 | 0 | 1 | 2 | 0 | 0 |
| G30 | 29 | 6 | 3 | 3 | 0 | 1 | 1 | 1 | 0 |
| B30 | 26 | 5 | 3 | 2 | 0 | 4 | 8 | 0 | 0 |
| R30 | 21 | 12 | 10 | 0 | 2 | 4 | 7 | 0 | 1 |



Figure 4: Snapshots from a successful trial. (a) Initial configuration; (b) Starting phase; (c) Role allocation phase; (d) Gripping phase; (e) Success (grip).

connection point. We refer to this type of inaccuracy as $I_3$. In a single trial in test B30, the *s-bots* failed to self-assemble. In this case, the *s-bots* ended up outside the perceptual range of their respective cameras. This trial in which the *s-bots* spent more than 1 minute without perceiving each other has been terminated, and it was considered unsuccessful.

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 interaction can be succinctly described by a combination of few distinctive phases and transitions between phases which exhaustively "portray" the observed phenomena. Figure 4 shows some snapshots from a successful trial which represent these phases. The robots leave their respective starting positions (see Figure 4a) and during the starting phase (see Figure 4b) 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 4c). 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 two *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 4d). 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 4e). 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. Recov-

ering from $I_1$ can only be accomplished by returning to a new RA-phase, in which the *s-bots* re-establish 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, 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 from 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]

## Conclusion

In this article, we have presented the results of an evolutionary methodology for the design of control strategies for self-assembling robots. 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 decision-making process between two ho-

mogeneous 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. (2006) and in 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, we have designed a self-assembling system that exhibits recovery capabilities that have not been selected during the evolutionary design phase 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. (2006) a recovery mechanism had to be designed as a specific behavioural module to be activated every time the robots failed to achieve assembly.

Our system is not as "transparent" as a hand-coded or modular rule-based one, as we can not 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, preliminary results not shown in the paper suggest that there is an effect of the starting configuration on the final outcome of a trial (how roles are allocated). In short, our analysis revealed that, in those trials in which the two robots have different initial perceptions ($\alpha \neq \beta$), the role that each *s-bot* assumes can be predicted knowing the combination of $\alpha$ and $\beta$. However, it is important to notice that perceiving the other robot at a specific distance and through a given camera sector does not inform a robot about the role it will assume during the trial. In other words, it is this combination of $\alpha$ and $\beta$ which determines the roles. In those cases in which the robots start with an identical perception ($\alpha = \beta$), 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's 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. We would also like 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. We will also 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

## References

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

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

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.

Groß, R., Bonani, M., Mondada, F., and Dorigo, M. (2006). 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*. In press.

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.

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

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.

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, *Proc. of the $8^{th}$ European Conference on Artificial Life (ECAL05)*, volume 3630 of *LNCS*, pages 272–281. Berlin, Germany:Springer-Verlag.

Quinn, M. (2001). Evolving communication without dedicated communication channels. In *Advances in Artificial Life: Sixth European Conference on Artificial Life (ECAL01)*, volume 2159 of *LNCS*. Berlin, Germany:Springer-Verlag.

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.

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.

Webb, B. (2000). What does robotics offer animal behaviour? *Animal Behaviour*, 60:545–558.

Whitesides, G. M. and Grzybowski, B. (2002). Self-assembly at all scales. *Science*, 295:2418–2421.

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