Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Evolving Functional Self-Assembling in a Swarm of Autonomous Robots

Vito Trianni, Elio Tuci and Marco Dorigo

Technical Report No.

TR/IRIDIA/2004-21

July 2004

Published in In S. Schaal et al., *From Animals to Animats VIII. Proceedings of the 8th International Conference on Simulation of Adaptive Behavior*, pages 405-414. MIT Press, Cambridge, MA.

Evolving Functional Self-Assembling in a Swarm of Autonomous Robots

Vito Trianni, Elio Tuci and Marco Dorigo

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

{vtrianni,etuci,mdorigo}@ulb.ac.be

Abstract

The goal of this study is the design of controllers for robots capable of physically connecting to each other, any time environmental contingencies prevent a single robot to achieve its goal. This phenomenon is referred to as *functional self-assembling*. Despite its relevance as an adaptive response, functional self-assembling has been rarely investigated within the context of collective robotics. Our task requires the robots to navigate within a rectangular corridor in order to approach light bulbs positioned on the opposite end of the corridor with respect to their starting positions. Aggregation and assembling are required in order to traverse a low temperature area, within which assembled robots navigate more effectively than a group of disconnected agents. The results of our empirical work demonstrate that controllers for a group of homogeneous robots capable of functional self-assembling can be successfully designed by using artificial neural networks shaped by evolutionary algorithms.

1 Introduction

This paper addresses the problem of synthesizing controllers for groups of autonomous robots—referred to as *s-bots*—capable of adaptively connecting to each other, forming an assembled structure—referred to as *swarm-bot*. The *swarm-bot* is a self-assembling, self-organizing artifact that exploits the cooperation among its simple components, the *s-bots*, in order to solve problems a single individual cannot cope with. The focus of this paper is on *functional self-assembling*, that is, the self-organized creation of a physically connected structure, which should be *functional* to the accomplishment of a particular task. In other words, our goal is the design of controllers for *s-bots* capable of connecting to each other (i.e., forming a *swarm-bot*) any time environmental contingencies prevent the single *s-bot* to achieve its goal.

Self-assembling occurs in a wide range of natural systems. In particular, it characterizes the behaviour of many social insects (see Anderson et al., 2002, for a review). A striking example of functional self-assembling can be observed in ants of the species *Œcophilla longinoda* (see Lioni et al., 2001). These ants are able to build chains connecting one to the other, creating bridges that facilitate the passage of other ants (see Figure 1, top).

Functional self-assembling is a problem of fundamental importance within the SWARM-BOTS project,[1] wherein this research is conducted. By drawing inspiration from the behaviour of social insects (see Bonabeau et al., 1999), functional

Technical Report IRIDIA/TR/2004-21 of IRIDIA, Université Libre de Bruxelles, Brussels, Belgium. December 2004.

[1]A project funded by the Future and Emerging Technologies Programme (IST-FET) of the European Community, under grant IST-2000-31010.

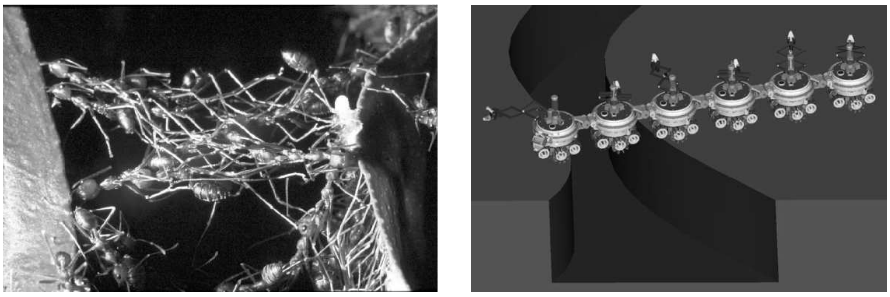

Figure 1: Examples of functional self-assembling. Left: ants of the species *Ecophilla longinoda* forming a chain (This image has been kindly provided by Dr. Arnaud Lioni). Right: a graphical representation of a *swarm-bot* structured in a chain while passing over a trough.

self-assembling can be viewed as an adaptive response of a group of autonomous agents to cope with environmental conditions which prevent them from carrying out their task. For example, the navigation of a group of robots can be hindered by the presence of a trough larger than the size of the body of a single robot, which can not be avoided and can be passed only by forming an assembled structure capable of bridging the gap (see Figure 1, bottom).

Despite its relevance as an adaptive response, functional self-assembling has been rarely investigated within the context of collective robotics. Several works in the literature have focused on the study of collective behaviour of autonomous agents, which due to the limits of their morphological structure, are not capable of self-assembling (see Cao et al., 1997; Liu and Wu, 2001). Other studies focus on the development of assembled structures of non-autonomous robotic units, that can be differently connected to each others by the experimenter and in some cases are able to self-reconfigure their shape (see Castano et al., 2000; Yim et al., 2000). With respect to this literature, the agents developed by the SWARM-BOTS project represent a step further. The *s-bots* are morphologically designed to autonomously carry out most of the tasks generally explored in collective robotics studies (e.g., foraging, clustering and sorting, navigation and surveillance), as well as to perform functional self-assembling. However, despite the facilities provided by the mechanics and electronics of our agents,[2] the design of control systems for autonomous robots capable of functional self-assembling is not a trivial task.

The complexity of functional self-assembling resides in the nature of the individual mechanisms required to bring forth the coordinated movements that lead firstly to the formation of the assembled structure, and subsequently to the collective motion of the assembled structure. Moreover, functional self-assembling is never *per se* the goal of the agents, but it generally requires to be integrated within the behavioural repertoire of agents capable of performing several different adaptive responses. In particular, within the SWARM-BOTS project, we are interested in investigating scenarios in which the *s-bots* should be capable of:

1. Independently performing a specific task. For example, if assembling is not required, *s-bots* should be capable of individually performing phototaxis.

2. Aggregating in order to allow subsequent assembling. That is, if assembling is required by particular environmental contingencies, the *s-bots* should be capable of bringing forth the conditions which facilitate functional self-assembling.

[2]Details regarding the hardware and simulation of the *swarm-bot* can be found in Mondada et al. (2004) and in the project web-site (`www.swarm-bots.org`).

Aggregation is one of the first steps in order to form an assembled structure—i.e., a *swarm-bot*.

3. Moving coordinately in order to assemble. That is, each *s-bot* should find the correct position with respect to another *s-bot* in order to be able to establish a connection.

4. Moving coordinately in order to improve the effectiveness of the behaviour of the assembled structure. For example, in order to pass over a trough, the assembled *s-bots* should negotiate a common direction of movement which brings the *swarm-bot* on the other side of the trough.

5. Disconnecting. That is, once the environmental contingencies do not require any longer the assembled structure, the *s-bots* should disconnect and carry out their goal independently.

This set of capabilities is commonly required by many complex tasks a *swarm-bot* should be able to carry out. In addition to the already mentioned example of passing over a trough larger than the body of a single *s-bot*, we can imagine a situation in which the *s-bots* have to climb a steep slope, or the case in which *s-bots* have to navigate through a very rough terrain. In this paper, we consider a simplification of these situations, that allows us to study the basic mechanisms that underpin functional self-assembling.

The task studied in this paper requires the *s-bots* to approach a light source located at the end of a corridor. Assembling is required to navigate in a "low temperature" area in which a *swarm-bot* can navigate more effectively than a group of disconnected *s-bots*. When located in the low temperature area, the aggregation of the *s-bots* should facilitate the subsequent assembling through their gripper element (see Section 2.1 for more details). Notice that, in our experiments, the "temperature" can be perceived by a single binary sensor which returns 1 if the *s-bots* are in a high temperature area, and 0 otherwise. This is a strong simplification with respect to other scenarios, in which the *s-bots* might be required to employ more complex sensory-motor skills in order to perceive those environmental contingencies that require assembling—e.g., the detection of a trough or of the increased roughness of the terrain. Moreover, the *s-bots* should perform individual and collective obstacle avoidance and coordinated motion. Obviously, these actions might require the use of more complex visual and proprioceptive sensors than the ones our simulated *s-bots* can rely on, in particular if the *s-bots* are located in an environment more complex than an empty corridor.

If compared to tasks that require *s-bots* to self-assemble in order to pass over a trough or to navigate on rough terrain, our scenario looks rather simple. It also relies on implementation details—such as the simulation of the gripper element, and the concept of energy, see Section 2 for further details—which might hinder the porting of the evolved controllers from simulated to real *s-bots*. However, we believe that these simplifications do not trivialise the significance of our experimental setup. They instead are required in order to allow us to begin our exploration within the domain of functional self-assembling. Another issue that induced us to take a simple or minimal approach concerns the methodology, as illustrated below.

In this study, we employed the evolutionary robotics (ER) methodology to develop integrated (i.e., single) control systems, which are capable of coordinating the actions of autonomous agents required to perform a broad repertoire of individual and collective responses. ER represents a way to automate the design of control systems for autonomous robots, using algorithms based on Darwinian evolution. The controllers are usually artificial neural networks, whose parameters are set by the evolutionary algorithm (see Harvey et al., 1992; Meyer et al., 1998; Nolfi and

Floreano, 2000, for reviews of the field). Recently, the ER approach has been successfully employed to evolve controllers for groups of homogeneous robots required to act cooperatively in order to achieve simple goals. Homogeneous groups are those in which the controllers of all the members of the group are derived from the same genotype—they are clones. For example, Quinn et al. (2003) evolved robots to perform a coordinated motion task, obtaining a situated specialization in distinct functional roles. Trianni et al. (2004) studied coordinated motion in a group of assembled robots that had to navigate in an arena with holes and open borders without falling down. Groß and Dorigo (2004) evolved cooperative transport behaviours for a group of two robots.

From the method's point of view, our task looks more demanding than those mentioned above. Moreover, at the best of our knowledge, we are not aware of any other study in which a similar methodology has been employed to design controllers for robots capable of displaying both individual and collective behaviours which require the coordination of such a rich sensory and motor apparatus, as in our task. In fact, the latter requires controllers—i.e., artificial neural networks—capable of bringing forth a variety of individual and collective responses triggered by visual stimuli (from the light sensors), acoustic stimuli (from the sound sensors), and proprioceptive stimuli (from the traction sensors), to appropriately control the state of various different actuators (loudspeakers, wheels, gripper, see Section 2.2 for more details). In consideration of this, the study of a minimal and simplified—rather than a more complex and more appealing—scenario looks more suitable to begin our exploration on functional self-assembling. Obviously, our intention is to use this study as a stepping stone toward the development of controllers for *s-bots* capable of displaying functional self-assembling in more complex environmental circumstances.

We believe that our empirical work, although preliminary, brings significant contributions: (i) at the best of our knowledge, our experiments represent one of the first works in which controllers for a group of homogeneous robots capable of functional self-assembling are described; (ii) our results prove that, despite the complexity of the task, artificial neural networks, shaped by evolutionary algorithms, can cope with complex scenarios in which a single robot's controller is required to bring forth different individual and collective responses based on the appropriate control of the state of various actuators (i.e., wheels, loudspeakers, gripper element) triggered by the local information coming from various sensors (i.e., traction sensors, microphones, ambient-light sensors).

1.1 Structure of the paper

In what follows, we first describe, in Section 2, the methodology used. In Section 2.1, we give a detailed description of the task. In Section 2.2, we describe the simulation model of the *s-bots* employed in this work. In Section 2.3, we illustrate the type of artificial neural network used to control the behaviour of the *s-bots* and the genetic algorithm used to shape the network. Section 2.4 describes in details the evaluation function. In Section 3, we show experimental results. In Section 4, some conclusions are drawn.

2 Methods

In this section, we describe the methodology defined for the evolution of controllers for *s-bots* capable of displaying individual and collective coordinated motion, phototaxis, obstacle avoidance, aggregation and self-assembling. In the following subsections, we describe in detail the task, we illustrate the morphology of the *s-bots*, the methodology employed to develop their controllers and the evaluation function used.

2.1 Description of the task

The task requires navigation within a rectangular corridor in order to approach light bulbs positioned on the opposite end with respect to the *s-bots*' starting positions (see Figure 2). The corridor (4 meters long, 1 meter wide) is divided in an area of high temperature and an area of low temperature (respectively, light and dark gray in Figure 2). Aggregation and assembling are required in order to traverse a low temperature area, within which a *swarm-bot* (i.e., assembled *s-bots*) navigates more effectively than a group of disconnected *s-bots*. The effectiveness of the navigational strategies is correlated with the amount of "energy" required by the *s-bot* to explore the corridor. In the area of high temperature, each *s-bots* saves more of its energy by navigating disconnected, while in the area of low temperature, each *s-bot* saves more energy by navigating assembled—i.e., by forming a *swarm-bot*. If, while navigating, an *s-bot* exhausts its energy, it is not able to move any more. The *s-bots* do not have any information concerning their energy level. However, the *s-bots* can reach the light bulbs before running out of energy if they properly react to the characteristics of the environment. In particular, an optimal strategy requires the *s-bots* (i) to individually move toward the light bulbs as long as the temperature remains high; (ii) to aggregate by exploiting the sound signaling system they are provided with as soon as the temperature drops; (iii) to continue their phototactic behaviour in an assembled structure (i.e., by forming a *swarm-bot*) throughout the low temperature area.

It is worth mentioning that the "energy level" of an *s-bot* does not refer to any physical property of the *s-bots*. Moreover, it does not imply the use of unrealistic sensors, which cannot be instantiated on the real *s-bots*. In fact, the *s-bots* do not have access to their own energy level. The energy level has been mainly introduced to define a simple task that requires self-assembling. In fact, the energy level plays an important role in the evaluation procedure, as described in Section 2.4.

At the beginning of each trial, three *s-bots* are randomly positioned and oriented at one end of the corridor, in the area of high temperature. The light bulbs, located at the opposite end of the corridor, can be perceived by the *s-bots* from anywhere

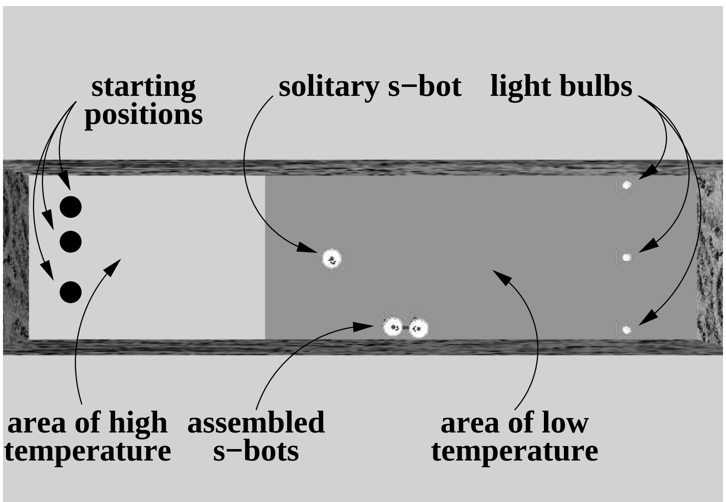

Figure 2: A graphical representation of the task. See text for details.

within the corridor. The intensity of the light which impinges upon the *s-bots* light sensors decreases quadratically with the distance from the light sources. Each *s-bot* s has an initial amount of energy $e_s = 1$, which must be higher than a certain threshold $\epsilon = 0.01$ for the *s-bot* to be able to move. The energy level of each *s-bot* can increase or decrease depending on:

1. The temperature of the area in which the *s-bot* is currently located. The temperature is 1 if the *s-bot* is in a high temperature area, 0 if it is in a low temperature area.

2. The state of the *s-bot*'s loudspeaker. An *s-bot* emits a tone to signal its position to other *s-bots*. This signaling behaviour can facilitate the aggregation of the group, which is a prerequisite for the assembling.

3. Whether the *s-bot* is assembled or disassembled.

More precisely, when *s-bot* s is assembled in a *swarm-bot* formation, it loses its energy in the area of high temperature and it increases its energy in the area of low temperature, as described by the following equation:

$$e_s(t+1) = e_s(t) + \tau \cdot ((1 - \Gamma(t)) - e_s(t)), \tag{1}$$

where $e_s(t)$ is the energy level of the s^{th} *s-bot* at cycle t, $\tau = 0.2$ is a time constant governing the speed of the energy variation and $\Gamma(t)$ is the temperature sensed by the s^{th} *s-bot* at cycle t in its current position. When an *s-bot* is not connected but it emits a sound signal, it loses energy in both areas. In the areas of low temperature its energy converges to a low but non-null value. This is described by the following equation:

$$e_s(t+1) = e_s(t) + \tau \cdot (k(1 - \Gamma(t)) - e_s(t)), \tag{2}$$

where $k = 0.1$ is a constant. In all the other situations, the *s-bot* increases its energy in areas of high temperature and loses it in areas of low temperature:

$$e_s(t+1) = e_s(t) + \tau \cdot (\Gamma(t) - e_s(t)). \tag{3}$$

The time constant τ guarantees that the *s-bots*' energy level varies smoothly according to the state of the system as described above. This smooth variation gives time to the control system of each *s-bot* to react to the new environmental situation in order to perform appropriate actions, before its energy level drops under the threshold $\epsilon = 0.01$.

The simulation is deliberately noisy, with noise added to all sensors (see Section 2.2 for more details). This is also extended to the environmental parameters: at the beginning of each trial, the point in which the temperature changes from high to low is redefined randomly within certain limits.

2.2 The simulation model

We developed a simulation software based on Vortex$^{\text{TM}}$, a 3D rigid body dynamics simulator that provides primitives for the implementation of detailed and realistic physics-based simulations (see also Mondada et al., 2004, for more details about the simulator). We have defined a simple *s-bot* model that at the same time allows fast simulations and preserves those features of the real *s-bot* that were important for the experiments (see Figure 3). The *s-bot* has a differential drive motion provided by a traction system composed of four wheels: two lateral, motorized wheels and two spherical, passive wheels placed in the front and in the back, which serve as support. The four wheels are fixed to the chassis, which also holds the cylindrical

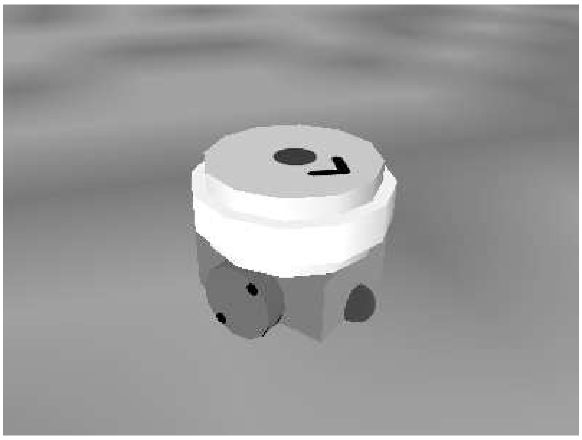

Figure 3: The simulation model of the *s-bot*.

rotating turret. The turret can rotate around its axis, and it holds many sensory systems. Connections among *s-bots* can be made using a virtual gripper, which is modeled by dynamically creating a joint between two *s-bots* when needed. The position of the virtual gripper is represented by an arrow painted on the turret. Finally, the turret also holds a loudspeaker that can be controlled to produce a tone perceived by the other *s-bots*.

Each *s-bot* is provided with a *traction sensor*, which detects the forces that are applied to the junction between the chassis and the rotating turret. This particular sensor proved to be important to evolve coordinated motion strategies in a *swarm-bot* and can also work as a bumper for individual and collective obstacle avoidance (Baldassarre et al., 2003; Trianni et al., 2004). Four variables encode the traction force information from four different preferential orientations with respect to the chassis (front, right, back and left, see Baldassarre et al., 2003, for more details). Each *s-bot* is also equipped with three directional microphones, used to detect the tone emitted by other *s-bots*, and with a temperature sensor. Finally, the ambient light is sensed using eight light sensors mounted on the rotating turret. We have defined two virtual light sensors positioned on the front and on the back of the chassis, whose readings are computed using the two closest light sensors of the turret. The information provided by the virtual light sensor does not vary significantly with the rotation of the turret with respect to the chassis, and it facilitates the evolution of phototactic behaviours. Noise is simulated for all sensors, adding a random value uniformly distributed within the 5% of the sensors saturation value.

Concerning the actuators, *s-bots* can control the two wheels, independently setting their speed in the range $[-6.5, 6.5]$ rad/s. The loudspeaker can be switched on, simulating the emission of a continuous tone, or it can be turned off. The virtual gripper can be closed or open, in order to connect to another *s-bot*. However, we force the gripper to stay open if a connection fails. Finally, the motor controlling the rotation of the turret is used, even though it is not directly controlled by the evolved neural network. When *s-bots* are not connected, this motor ensures the alignment between the turret and the chassis. On the contrary, when an *s-bot* is connected to other *s-bots* to form a *swarm-bot*, the turret can rotate freely.

2.3 The controller and the genetic algorithm

Homogeneous groups of *s-bots* are controlled by artificial neural networks, whose parameters are set by an evolutionary algorithm. A single genotype is used to create

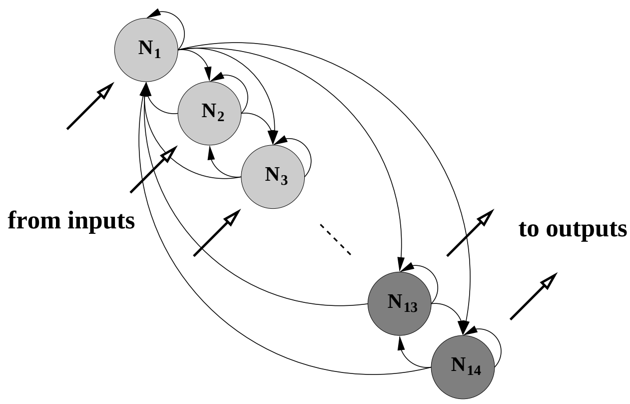

Figure 4: A graphical representation of the artificial neural network employed to control the *s-bots*. The nodes in gray represent those which receive input from the *s-bots* sensors. The nodes in black represent those whose activation values are used to set the *s-bots* actuators.

a group of individuals with an identical control structure. Each *s-bot* is controlled by a fully connected, 14 neuron continuous time recurrent neural network (CTRNN, see also Beer, 1995, for details). We chose to employ a 14 neurons network because it corresponds to the smallest network which has at least one neuron associated with each input signal and one neuron for each output signal (see Figure 4). Each neuron, indexed by i, is governed by the following state equation:

$$\frac{dy_i}{dt} = \frac{1}{\tau_i}\left(-y_i + \sum_{j=1}^{14} \omega_{ji} z_j + gI_i\right) \qquad \text{where} \qquad z_j = \frac{1}{1+e^{-(y_j+\beta_j)}} \tag{4}$$

Here, by analogy with real neurons, y_i is the cell potential, τ_i the decay constant, β_j the bias term, z_j the firing rate, ω_{ji} is the strength of synaptic connections from the j^{th} neuron to the i^{th} neuron, I_i the intensity of the sensory perturbation on sensory neuron i. The neurons either receive direct sensor input or are used to set the state of an *s-bot*'s actuators (i.e., the wheels speed, the state of the gripper and the loudspeaker). There are no internal neurons. All but four of the neurons receive direct input from the robot sensors (for the remaining four, $gI_i = 0$). Each input neuron is associated with a single sensor, receiving a real value in the range $[0.0, 1.0]$, which is a simple linear scaling of the reading taken from its associated sensor.[3]

The four remaining neurons are used to control the *s-bot*'s actuators, after mapping their cell potential y_i onto the range $[0.0, 1.0]$ by a sigmoid function. Two of them are used to set the *s-bot*'s wheels speed, linearly scaling the output into $[-6.5, 6.5]$. The third motor neuron is used to set the state of the loudspeaker, which is turned on if the neuron output is higher than 0.5, and off otherwise. The last motor neuron controls the gripper actuator, trying to set up a connection if the neuron output is higher than 0.5, and keeping the gripper open otherwise.

The strengths of the synaptic connections, the decay constants, bias terms and the gain factor are all genetically encoded parameters. Cell potentials are set to 0

[3]Specifically, neurons N_1 to N_4 take input from the 4 variables encoding the traction force, neurons N_5 to N_7 take input from the sound sensors (i.e., the directional microphones), N_8 and N_9 from the virtual light sensors, and N_{10} from the temperature sensor.

each time a network is initialised or reset. State equations are integrated using the forward Euler method with an integration step-size of 0.1.

In order to set the parameters of the *s-bots*' controllers, a simple generational genetic algorithm (GA) is employed (Goldberg, 1989). Initially, a random population of 100 genotypes is generated. Each genotype is a vector of 1800 binary values—8 bits for each of the 225 parameters, that is 196 connections, 14 decay constants, 14 bias terms, and 1 gain factor. Subsequent generations were produced by a combination of selection with elitism and mutation. Recombination is not used. At every generation, the best 20 genotypes are selected for reproduction, and each generates 4 offspring. The genotype of the selected parents is copied in the subsequent generation; the genotype of their 4 offspring is mutated with a 5% probability of flipping each bit. One evolutionary run lasts 1000 generations.

The binary values of a genotype were mapped to produce CTRNN parameters with the following ranges:

- connection weights: $\omega_{ji} \in [-6, 6]$.
- biases: $\beta_j \in [-2, 2]$.
- gain factor: $g \in [1, 13]$.

Concerning the decay constants, the genetically encoded parameters were firstly mapped onto the range $[-1, 1]$ and then exponentially mapped onto $\tau_i \in [10^{-1}, 10]$.

2.4 The evaluation function

During the evolution, a genotype is mapped into a control structure that is cloned and downloaded in all the *s-bots* taking part to the experiment (i.e., homogeneous group of *s-bots*). Groups of $n = 3$ *s-bots* are evaluated 5 times—i.e., 5 trials. Each trial differs from the others in the initialisation of the random number generator, which influences mainly the *s-bots* starting positions and the point beyond which the temperature drops from 1 to 0. In each trial θ, the lifetime of an *s-bot* is limited to 600 simulation cycles, corresponding to 60 seconds of real time. The behaviour of the *s-bots* is evaluated according to an evaluation function F_θ that takes into account the individual contribution of each *s-bot* s:

$$F_\theta = \frac{1}{n^3} \cdot \left(\sum_{s=1}^{n} d_s \cdot \sum_{s=1}^{n} e_s \cdot c \right), \tag{5}$$

where the factors d_s, e_s and c are explained below.

- d_s rewards *s-bots* that perform phototaxis; this fitness component is computed as follows:

$$d_s = \begin{cases} 0.1 \cdot \dfrac{x_{f,s} - x_{i,s}}{x_\Gamma - x_{i,s}} & \text{if } x_{f,s} \leq x_\Gamma \\ \\ 0.1 + 0.9 \cdot \dfrac{x_{f,s} - x_\Gamma}{x_M - x_\Gamma} & \text{if } x_\Gamma < x_{f,s} \leq x_M \\ \\ 1 & \text{otherwise} \end{cases} \tag{6}$$

  where $x_{i,s}$ and $x_{f,s}$ are respectively the initial and final x coordinate of the s^{th} *s-bot* position, x_Γ is the x coordinate in which the temperature drops from 1 to 0, and x_M is the x coordinate of the light bulbs position.[4]

[4]The coordinate system used has the x and y axes parallel respectively to the long and short wall of the corridor. The origin of the axes is positioned at the bottom left corner of the corridor.

- e_s is the final energy possessed by the s^{th} *s-bot*, at cycle $t = 600$. The variation of the energy $e_s(t)$ of the s^{th} *s-bot* at cycle t is regulated by (1), (2) and (3) as discussed in Section 2.1.

  This fitness component rewards *s-bots* that end their lifetime with a high amount of energy. For example, if we compare groups of *s-bots* that managed to reach the end of the corridor close to the light bulbs, those which proved to be capable of assembling early in response to the decrease in the environmental temperature will get a higher fitness score than those which did not perform such collective response.

- c is the maximum size of a *swarm-bot* observed at the end of the trial, ranging from 1 (no connections among *s-bots*) to n (all *s-bots* connected in a single *swarm-bot*). This fitness component rewards *s-bots* that reach the end of the corridor assembled in a *swarm-bot* formation. Recall that, due to the characteristics of the environment—an initial area of high temperature is followed by an area of low temperature at the end of which the light bulbs are located—successful *s-bots* should terminate the trial in *swarm-bot* formation close to the opposite end of the corridor with respect to their starting position.

The average performance of the group F is computed averaging the evaluations F_θ performed in each trial θ. This value corresponds to the fitness of the genotype: it is used to select which genotypes will reproduce in the current generation, but is not in any sense a reinforcement directly available to the *s-bots*.

3 Results

Ten evolutionary runs, each using different random seeds, were run for 1000 generations each. Two runs out of ten ended up successfully by producing controllers capable of displaying functional self-assembling. Figure 5 shows the fitness of the best group of *s-bots* and the average fitness of the population for each generation. Two prototypic runs are shown: a particularly successful one (top) and an unsuccessful one (bottom). Note that, the fitness curve does not increase monotonically because each genotype is evaluated every generation. Therefore, even in the absence of mutations, genotypes with relatively high fitness—e.g., those corresponding to the elite— might not obtain the same fitness in the following generation.

We examined the behaviour of the best evolved group of *s-bots* of the last generation for each run, in order to establish whether it evolved functional self-assembling. The fitness of each of these groups was re-evaluated 250 times by using the equations illustrated in Section 2.4. The corresponding results are shown in Table 1. The first two runs are the most successful ones, showing an average fitness higher than all the others, which did not result in a satisfying performance.

Table 1: Average fitness and standard deviation of the best evolved controller of the last generation of each run.

seed	1	2	3	4	5
avg.	0.505	0.572	0.052	0.031	0.028
std.	0.218	0.289	0.028	0.003	0.002
seed	6	7	8	9	10
avg.	0.006	0.030	0.292	0.031	0.058
std.	0.007	0.002	0.325	0.002	0.029

An analysis of the controllers produced by the unsuccessful runs revealed that these groups of *s-bots* were able to solve the task only in part. We observed that,

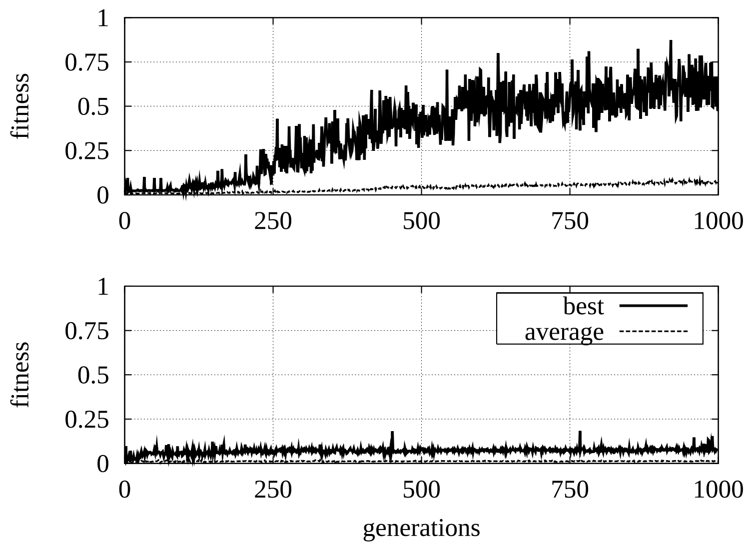

Figure 5: The graphs show the fitness of the best group of *s-bots* (thick line) and the normalized average fitness of the population (thin line), for each generation, for a successful run (top graph) and an unsuccessful one (bottom graph).

while in these runs the *s-bots* were capable of phototaxis and obstacle avoidance, only in few runs they were able to properly react to the decrease in temperature, and only in the run n. 8 the *s-bots* were occasionally capable of functional self-assembling. On the contrary, in the two successful runs, the groups of *s-bots* showed the complete repertoire of behaviours required by the task.

3.1 Behavioural analysis: self-assembling

In this section, we illustrate the results of post-evaluation tests performed on one of the successful evolutionary run. Similar results can be observed in the post-evaluation tests of the other successful run, with a small difference as far as it concerns the time required for the formation of the self-assembled structure (i.e., the *swarm-bot*).

Figure 6 shows how the covered distance[5] and the energy level of each *s-bot* vary over time. Looking at these graphs, it is possible to distinguish four different phases: individual phototaxis, aggregation, self-assembling and collective phototaxis. At the beginning of the trial—from cycle 0 to the time indicated by the empty circle—the three *s-bots*, located in the high temperature area and with full energy, perform individual phototaxis, as shown by the continuous line.

The second phase starts when the *s-bots* enter the low temperature area. Three phenomena can be observed in Figure 6: aggregation, decrease in the energy level and signaling behaviour. Aggregation is indicated by the covered distances of the three *s-bots* (see continuous lines), which reach similar values before the end of the phase. The decrease in the energy level, according to equations (2) and (3), indicates that the *s-bots* move independently. Since the energy level converges, for each *s-bot*, to the value $e_s = k$ (see equation (2)), we can deduce that the *s-bots* react to the temperature decrease by switching on their loudspeaker, signaling their position to the other *s-bots*. This should in principle facilitate the aggregation. However, we observed that the *s-bots* tend also to exploit other "affordances", such

[5]The covered distance refers to the distance between the current position of the *s-bot* and the starting position, along the x axis.

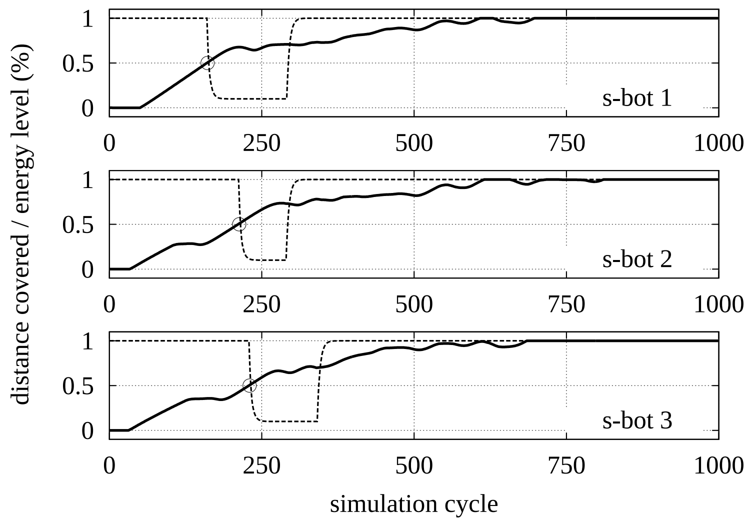

Figure 6: The graphs refer to a post-evaluation of the best evolved group of three *s-bots*. In particular, each graph shows how the covered distance along the corridor (continuous line) and the energy level (dashed line) of an *s-bot* vary during a post-evaluation which lasts 1000 simulation cycle. The empty circles indicate the time when an *s-bot* enters the low temperature area.

as the walls of the corridor, in order to get close to each other (Gibson, 1977). A complete understanding of this process would require further analysis, that will be carried out in the future.

The third phase corresponds to self-assembling. In Figure 6, this phase is indicated by an increase in the energy level (dashed line), caused by the *s-bots* connecting to each other when located in the low temperature area (see equation (1)). In this particular case, *s-bots* 1 and 2 self-assemble first, while *s-bot* 3 joins the *swarm-bot* later. Collective phototaxis is performed during the last phase. Here, *s-bots* move assembled in a *swarm-bot* that approaches the light bulbs, as indicated in Figure 6 by the synchronous increase of the covered distance (see continuous lines).

3.2 Disassembling of the *swarm-bot*

In an additional series of post-evaluation tests, we looked at the capability of the *swarm-bot* to disassemble as a reaction to an increase of the environmental temperature. Recall that this circumstance has never been encountered by the *s-bots* during the evolutionary phase. Therefore, disassembling should be considered an additional capability of the evolved controllers, which confers robustness to the system. We placed the *s-bots* in a corridor with four temperature areas: two high temperature and two low temperature areas (see Figure 7).

The graphs in Figure 8 show how the covered distance and the energy level of each *s-bot* vary in time while the *s-bots* move down the corridor toward the light bulbs. In this case, we focus our attention on how the *s-bots* react to the transition from low to high temperature areas. In fact, the transitions from high to low temperature areas result in a variation of the covered distance and of the energy levels similar to what was observed and discussed for Figure 6.

The transition of the *s-bots* from low to high temperature areas is indicated in the graphs by a filled circle. Roughly speaking, this transition is characterized by

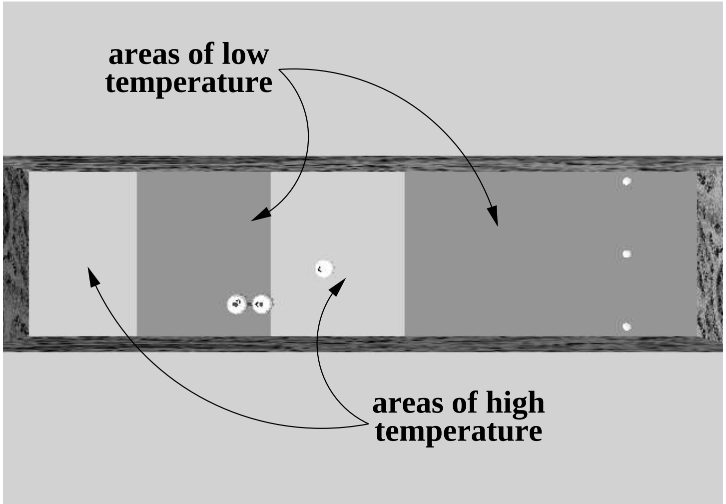

Figure 7: A graphical representation of the environment with two high temperature and two low temperature areas. This environment has been used for post-evaluation to check whether the *s-bots* capable of assembling were also capable of disassembling in response to an increase in the environmental temperature.

two different phases. Initially, a decrease in the energy level is observed, according to equation (1), when an *s-bot*, still assembled in a *swarm-bot* formation, perceives the new environmental condition (high temperature). Subsequently, the *s-bots* progressively disconnect from each other, which results in a gain in the energy level. In the particular case illustrated in Figure 8, *s-bot* 1 is the first to perceive the high temperature area and consequently to disassemble from the *swarm-bot*. It is possible to notice that *s-bot* 1, after disconnecting, moves back and forth, experiencing twice the low-to-high temperature transition. Similarly, *s-bot* 2 disconnects from *s-bot* 3 as soon as it ends up in the high temperature area. Consequently, *s-bot* 3 finds itself alone in the area of low temperature. It is possible to notice that its energy level drops, according to equation (2), due to the fact that the *s-bot* has the loudspeaker turned on. Nevertheless, the *s-bot* still has enough energy to perform individual phototaxis and to approach the high temperature area. Once in the high temperature area, its energy level increases again, indicating that the *s-bot* has switched off the loudspeaker. Its covered distance indicates that the *s-bot* approaches faster the light bulbs, reaching and finally connecting to the other 2 *s-bots*.

4 Conclusions

In this paper, we have described a set of experiments aimed at the design of controllers for a homogeneous group of robots—the *s-bots*. Our goal was the evolution of a single neural network capable of integrating a repertoire of individual and collective behaviours. In particular, we focused on functional self-assembling, that is, on the ability of the *s-bots* to react to particular environmental contingencies forming a physically connected structure—the *swarm-bot*.

Our results show that the evolutionary robotic methodology is promising: the evolved controllers are capable of displaying individual and collective obstacle avoidance, individual and collective phototaxis, aggregation and self-assembling. To the best of our knowledge, our experiments represent one of the first works in which

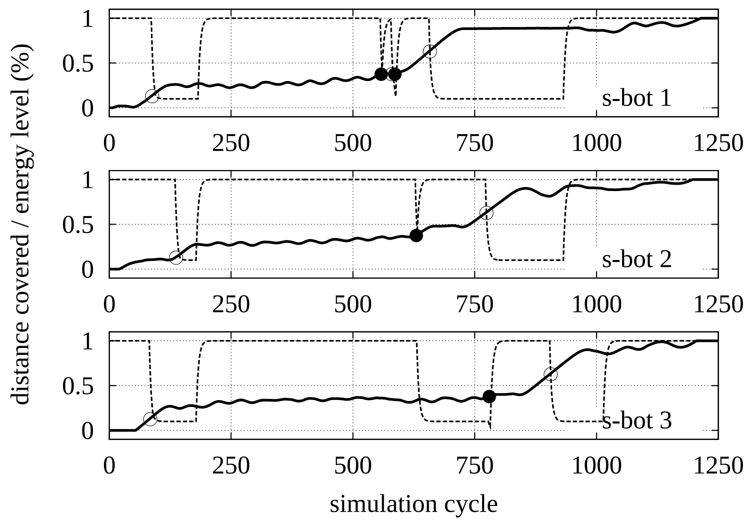

Figure 8: Each graph shows how the distance to the light bulbs (continuous line) and the energy level (dashed line) of an *s-bot* vary during a post-evaluation which lasts 1250 simulation cycle. The empty circles indicate the time when an *s-bot* enters a low temperature area. The filled circles indicate the time when an *s-bot* enters a high temperature area.

(i) functional self-assembling in a homogeneous group of robots has been achieved and (ii) evolved neural controllers successfully cope with such a complex scenario, producing different individual and collective responses based on the appropriate control of the state of various actuators triggered by the local information coming from various sensors.

We approached the study of functional self-assembling by identifying possible scenarios in which the *s-bots* could require the formation of a *swarm-bot*. In Section 1, we provided some examples of these scenarios, such as passing over a trough or navigation on a rough terrain. Instead of directly facing these challenges, we have chosen to simplify the task, in order to focus our attention on the basic mechanisms that underpin functional self-assembling. The task studied in this paper is the result of a careful analysis of the common features observed in those scenarios that require self-assembling. If compared to tasks that require *s-bots* to self-assemble in order to pass over a trough or to navigate on rough terrain, our scenario looks simpler. However, we believe that this is a important initial step to be taken before proceeding further.

The obtained results are satisfying, though only two evolutionary runs out of ten were successful. We believe that, on the one hand, it is important to improve our methodology in order to consistently obtain successful behaviours. On the other hand, these results are very promising, as they suggest that solutions to complex tasks can be obtained using the evolutionary robotics methodology.

In the future, we aim at the improvement of the results discussed in this paper. First of all, we would like to consistently evolve successful controllers, by improving the evaluation function and by exploring the parameter space of the neural controllers and of the genetic algorithm we used. Furthermore, we would like to synthesize more efficient behaviours, in particular concerning the use of the sound signaling system for aggregation and self-assembling.

Another interesting issue is given by the disassembling capabilities of the *swarm-bot*. Disassembling can be improved in many different ways. In particular, we would like the *s-bots* to be capable of using the sound signalling system both to favour aggregation and self-assembling (switching on the loudspeaker) and to trigger the disassembling of connected *s-bots* (turning it off). In this way, disassembling would not be exclusively under the control of the *s-bot* that makes a connection, but it could also be influenced by the *s-bot* that "receives" a connection.

Finally, we intend to study functional self-assembling triggered by a temporal integration of the information available to an *s-bot*. That is, self-assembling should arise when no other possible solution has been found while the *s-bots* are experiencing the environment. This implies that an *s-bot* should be capable of integrating some information obtained through its sensors during time, before trying to self-assemble with other *s-bots*. For example, when trying to reach a location that is barred by the presence of a trough, an *s-bot* should first try to search the surroundings for a feasible passage. When such a passage does not exist, the iterated experience of the presence of the trough should trigger the self-assembling phase.

Acknowledgments

This work was supported by the "SWARM-BOTS" project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-2000-31010. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate, through the grant "Virtual Swarm-bots", contract no. 9.4515.03, and from the "ANTS" project, an "Action de Recherche Concertée" funded by 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.-L. (2002). Self-assemblage in insects societies. *Insectes Sociaux*, 49:99–110.

Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolution of collective behavior in a team of physically linked robots. In Gunther, R., Guillot, A., and Meyer, J.-A., (Eds.), *Proceedings of the Second European Workshop on Evolutionary Robotics*, pages 581–592. Springer Verlag, Berlin, Germany.

Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215.

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

Cao, Y. U., Fukunaga, A. S., and Kahng, A. B. (1997). Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:1–23.

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

Gibson, J. J. (1977). The theory of affordances. In Shaw, R. and Bransford, J., (Eds.), *Perceiving, Acting and Knowing. Toward an Ecological Psychology*, chapter 3, pages 67–82. Lawrencw Erlbaum Associates, Hillsdale, NJ.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, Reading, MA.

Groß, R. and Dorigo, M. (2004). Evolving a cooperative transport behavior for two simple robots. In Liardet, P., Collet, P., Fonlupt, C., Lutton, E., and Schoenauer, M., (Eds.), *Artificial Evolution – 6^{th} International Conference, Evolution Artificielle, EA 2003, Marseille, France, October 2003*, volume 2936 of *Lecture Notes in Computer Science*, pages 305–317. Springer Verlag, Berlin, Germany.

Harvey, I., Husbands, P., and Cliff, D. (1992). Issues in evolutionary robotics. In Meyer, J.-A., Roitblat, H., and Wilson, S., (Eds.), *From Animals to Animats II: Proceedings of the 2^{nd} International Conference on Simulation of Adaptive Behavior*, pages 364–373. MIT Press/Bradford Books, Cambridge, MA.

Lioni, A., Sauwens, C., Theraulaz, G., and Deneubourg, J.-L. (2001). Chain formation in *Œcophilla longinoda*. *Journal of Insect Behaviour*, 15:679–696.

Liu, J. and Wu, J. (2001). *Multiagent Robotic Systems*, volume 21 of *International Series on Computational Intelligence*. CRC Press, Boca Raton, FL.

Meyer, J.-A., Husbands, P., and Harvey, I. (1998). Evolutionary robotics: A survey of applications and problems. In Husbands, P. and Meyer, J.-A., (Eds.), *Proceedings of the First European Workshop on Evolutionary Robotics*. Springer Verlag, Berlin, Germany.

Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots, Special Issue on Swarm Robotics*. To appear.

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

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.

Trianni, V., Nolfi, S., and Dorigo, M. (2004). Hole avoidance: Experiments in coordinated motion on rough terrain. In *Proceedings of the 8^{th} International Conference on Intelligent Autonomous Systems (IAS-8)*, Amsterdam, The Netherlands. To appear.

Yim, M., Duff, D., and Roufas, K. (2000). PolyBot: A modular reconfigurable robot. In *Proceedings of the 2000 IEEE/RAS International Conference on Robotics and Automation*, volume 1, pages 514–520. IEEE Robotics and Automation Society, Piscataway, NJ.