

Cooperative Transport of Objects of Different Shapes and Sizes

Roderich Groß and Marco Dorigo

IRIDIA - Université Libre de Bruxelles - Brussels, Belgium
{rgross,mdorigo}@ulb.ac.be

Abstract. This paper addresses the design of control policies for groups of up to 16 simple autonomous mobile robots (called *s-bots*) for the cooperative transport of heavy objects of different shapes and sizes. The *s-bots* are capable of establishing physical connections with each other and with the object (called *prey*). We want the *s-bots* to self-assemble into structures which pull or push the prey towards a target location. The *s-bots* are controlled by neural networks that are shaped by artificial evolution. The evolved controllers perform quite well, independently of the shape and size of the prey, and allow the group to transport the prey towards a moving target. Additionally, the controllers evolved for a relatively small group can be applied to larger groups, making possible the transport of heavier prey. Experiments are carried out using a physics simulator, which provides a realistic simulation of real robots, which are currently under construction.

1 Introduction

The transport of heavy objects by groups of robots can be motivated by low cost of manufacture, high robustness, high failure tolerance, or high flexibility – all desirable properties for a robotic system to have. Instead of using one robot powerful enough to transport the object (also called *prey* hereafter) on its own, the task is accomplished in cooperation by a group of simpler robots.

In the literature, several works can be found that consider the transport of an object by a team of robots. Deneubourg et al. (1990) proposed the use of self-organized approaches for the collection and transport of objects by robots in unpredictable environments. Each robot unit could be simple and inefficient in itself, but a group of robots could exhibit complex and efficient behaviors. Cooperation could be achieved without any direct communication among robots (Grassé, 1959; Deneubourg and Goss, 1989).

Kube and Zhang (1993a,b) studied a distributed approach to let a group of simple robots find and push a prey towards a light. Kube and Bonabeau (2000), on a follow-up of Kube and Zhang’s research, evaluated the sensitivity of their robotic system to the prey’s geometry by comparing the performance on prey of different shapes. They report that when the prey was a small cuboid, “as the number of robots increased the task took longer to complete as the robot interference was high since the limited box side space created competition




Fig. 1. (a) Example of initial placement: the prey (black) has to be transported towards a light-emitting beacon acting as a target (e.g., a robot nest). We assume that in the environments in which the robotic system will be applied, the s-bots will approach the prey from the half space on the side of the nest. To simulate this, the s-bots (gray) are placed randomly within a semi-circle 40 to 50 cm away from the prey. (b) Top view of four s-bots (gray) and of those four prey (black) having the most extreme dimensions. The mass of each prey is independent of its shape and size.

among the robots”. If the robot’s diameter was reduced while keeping the box dimensions fixed, the performance increased.

A similar observation has been reported by Groß and Dorigo (2004): they evolved neural networks to control two simple robots in order to pull or push a cylindrical prey as far as possible in an arbitrary direction. Increasing the number of robots while using heavier prey with the same cylindrical shape and dimension, the performance of the evolved strategies greatly decreased. However, if the diameter of the prey was multiplied by the same factor as the number of robots was increased, the controllers could reach again a satisfactory level of performance. In the latter case the density of robots around the prey is comparable with the one in the most successful setup used by Kube and Bonabeau (2000).

In this paper, we aim at a robotic system that is appropriate for the transport of prey of different shapes and sizes. We consider prey of cuboid as well as cylindrical shape, with footprints that may differ in size by factors up to 6.25. In addition, the prey can be of two different heights. The weight of the prey is independent of its geometry and may vary up to a factor of 2.

2 Experimental Framework

2.1 Environment and Task

The simulated environment consists of a flat plane, a prey with characteristics that change across simulations (shape, size and mass), and a light-emitting beacon acting as a target and placed 300 cm away from the initial position of the prey. A simulation lasts 20 simulated seconds. Initially, the prey is put with random orientation on the plane and four s-bots are put at random positions and with random orientations in the neighborhood of the prey (see Figure 1 (a)). Since the s-bots are placed in the vicinity of the prey, it is not necessary for the s-bot to deploy any elaborate exploration of the environment in order to find the




Fig. 2. The s-bot: (a) front, side and top view of the s-bot model and (b) the controller network. For details see Sections 2.3 and 2.4.

prey. The s-bot controllers are supposed to let the s-bots localize and approach the prey, and self-assemble into structures physically linked to the prey, in order to pull or push it towards the beacon.

2.2 Prey

The prey is modeled as a cylinder or as a cuboid of height 12 or 20 cm. Tall prey are of special interest since they can prevent s-bots from localizing the target location. In case of a cylinder, the radius (in cm) is chosen in the range [6, 15]. If the prey is modeled as a cuboid, the length l of one horizontal side (in cm) is chosen in the range [12, 40], while the other horizontal side length is set in the range [12, $\min(l, 20)$]. The mass of the prey (in grams) is in the range [500, 1000]. Depending on its weight, the cooperative behavior of at least 2 or 3 s-bots is necessary to move the prey. Objects of height 12 cm have a uniform distribution of mass, while the taller ones are modeled as a stack of two geometrically identical objects of different masses: the mass of the object on top of the stack is 25% of the total mass. This prevents tall prey from toppling down when pushed or pulled by the s-bots. Every potential geometry has the same probability to be chosen in each simulation. Figure 1 (b) illustrates the dimensions of different prey with respect to the s-bot's size.

2.3 S-Bot Model

The s-bot model is illustrated in Figure 2 (a). It is an approximation of a real s-bot, currently under construction within the SWARM-BOTS project (Mondada et al., 2004; Dorigo et al., 2004, see also www.swarm-bots.org). The model provides a degree of freedom enabling the rotation of the s-bot's upper part (the turret) with respect to the lower part (the chassis) by means of a motorized axis. The turret is composed of a cylindrical torso, a *grripper* element fixed on the torso's edge heading outwards, and of a pillar placed on the top to support a mirror providing a camera system with an omni-directional view. The chassis is composed of a cylindrical element and of four wheels attached to it: two motorized wheels on the left and right, and two passive wheels, one in the front, and one in the back. The s-bot has a height of 16.85 cm and a mass of 660 g.

In the following, the actuators and sensors are described. Both the actuators and the sensors are affected by noise (for details see Groß (2003)).

Actuators. The s-bot is equipped with two motorized wheels on the left and right, a rotational base, a gripper element and an LED ring. Each motorized wheel can be controlled by setting the desired angular speed (in rad/s) in the range $[-8, 8]$. The rotational base actuator enables the s-bot to align its turret with respect to its chassis to any angular offset in $[-\pi, \pi]$. The gripper element is modeled as a box heading forward. If the element is in *gripping* mode, a connection will be established as soon as the gripper element touches the turret of another s-bot or the prey. If the gripper is set to the *open* mode, any gripped object is released. The gripper of the real s-bot can establish a physical link to another s-bot or to the prey by grasping the surrounding ring this object is equipped with. The LED ring can be activated with different colors. In this study it is permanently activated with a red color to facilitate the s-bots detection.

Sensors. The simulated s-bot is equipped with a gripper status sensor, a rotational base sensor, and an omni-directional camera. These sensors have their counterparts on the real s-bot. The gripper status sensor is able to detect whether the gripper element is connected to another object or not. The rotational base sensor measures the angular offset between the turret and the chassis.

The simulated camera provides data based not on recorded images but on information available in the simulator, e.g., if up to a pre-defined distance there is a green object in front of the s-bot or not. The implicit assumption is that when working with the real s-bot, such information can be obtained from the camera using feature extraction algorithms. Although the real camera is attached to the turret (that can be rotated with respect to the chassis), it is possible to simulate a camera that is aligned with the chassis, since it is possible to measure the angular mismatch between the turret and the chassis.

Using the camera, the s-bot is able to scan its surrounding for teammates or for the prey. Given any horizontal direction, the s-bot can detect, up to a distance of $R = 50$ cm in the specified direction, the presence of a prey (green object) or (if not shadowed) of an s-bot (red object). In both cases, the s-bot can compute a rough estimate of the object distance. In addition, the camera can localize the light-emitting beacon (horizontal angle), as long as it is not shadowed by the prey or by a teammate.

2.4 Control Policy

The architecture used in order to control the group of s-bots is a simple recurrent neural network (Elman, 1990) with six hidden nodes and whose weights are synthesized by an evolutionary algorithm. The same network is cloned so that it is identical on each s-bot. We hypothesize that Elman networks might be more appropriate than fully reactive networks, since the proper selection of sensory patterns by motor actions might be facilitated by the use of memory.

The neural network is illustrated in Figure 2 (b). It is activated every 100 ms. The six input neurons correspond to the five, current sensor readings and to the horizontal direction in which the camera scans for the teammates and prey. The activations of the output neurons are used to control the motorized wheels, the gripper and the angular offset of the rotational base. Furthermore, the neural network may shift the horizontal direction in which the camera scans the environment of any angle (in degrees) in the range $[-18, 18]$. Given such a controllable sensor, the s-bot can select sensory patterns by its motor actions¹.

In average the s-bot's detection of the prey at the beginning of the simulation would require at least several seconds since the perception is restricted to a single horizontal direction per control step. To facilitate the detection, the scan direction of the camera is manually aligned approximately towards the prey for the first second of simulation time.

2.5 Evolutionary Algorithm

The evolutionary algorithm used is a self-adaptive $(\mu + \lambda)$ evolution strategy (Rechenberg, 1965; Schwefel, 1975). The implementation is almost identical to the one utilized by Groß and Dorigo (2004). In this paper the initial population is not biased and purely random, while in the previous study it was constructed by a Monte Carlo search.

The number of offspring is $\lambda = 60$ and the number of parents is $\mu = 20$.

2.6 Fitness

Each individual represents a common controller for a group of s-bots. The objective is to let a group of $N = 4$ s-bots approach the prey, and self-assemble into structures physically linked to the prey, in order to pull or push it towards the beacon (the faster the better). The simulation lasts $T = 20$ seconds.

To favor the evolution of solutions that make use of the gripper element, the fitness function takes the assembling performance into account. Since the s-bots start disassembled up to $R = 50$ cm away from the prey, some time is required to approach the prey and to establish a connection. Therefore, the recording of the assembling performance A starts after 10 simulation seconds have elapsed:

$$A = \frac{1}{(T - 10 + 1)N} \sum_{k=10}^T \sum_{j=1}^N A_j(t = k), \quad (1)$$

where $A_j(t = k) \in [0, 1]$ is defined by

$$A_j(t = k) = \begin{cases} 1 & \text{if } j \in M(t = k); \\ 0 & \text{if } d_j(t = k) \geq R \wedge j \notin M(t = k); \\ 0.75 & \text{if } d_j(t = k) < \frac{R}{2} \wedge j \notin M(t = k); \\ 0.65 \left(\frac{R - d_j(t = k)}{R/2} \right) + 0.1 & \text{otherwise.} \end{cases} \quad (2)$$

¹ This is called *active perception* and, within evolutionary robotics, it has already been studied, for instance, by Kato and Floreano (2001) and Marocco and Floreano (2002).




Fig. 3. Box-and-whisker plot providing information about the evolutionary progress in ten evolutionary runs (4 s-bots; prey of various dimensions and weights, see Section 2.2). Each box comprises observations ranging from the first to the third quartile. The median is indicated by a bar, dividing the box into the upper and lower part. The whiskers extend to the the farthest data points that are within 1.5 times the interquartile range. Outliers are indicated as circles. Characteristics about two types of observations for each population are displayed: the best fitness value (gray boxes) and the average fitness (white boxes) of a population.

The value $d_j(t = k)$ denotes the distance between s-bot j and the prey at time k . $M(t = k)$ is the set of s-bots that are physically linked to the prey at time t . This comprises s-bots that directly grip the prey, but also those that are connected to a chain of s-bots of which the foremost is gripped to the prey. The assembling performance A rewards s-bots for being close to the prey and even more for being physically linked to the prey.

The distance gain D is given by $D = D(t = 0) - D(t = T)$, where $D(t = k)$ denotes the horizontal distance (in cm) between the prey and the beacon after k seconds of simulation time.

Finally, the quality measure q is defined as

$$q = \begin{cases} A & \text{if } D \leq 0; \\ 1 + (1 + \sqrt{D})A^2 & \text{otherwise.} \end{cases} \quad (3)$$

The fitness value (to be maximized) is given by the average quality observed in five tests with different initial configurations and different prey. The sample of tests is changed at each generation.

3 Results

The experimental setup described above has been used in ten independent evolutionary runs of 850 generations each. Every run lasted about three weeks on a machine equipped with an AMD Athlon XPTM 2400+ processor. Figure 3 shows the development of the best fitness (gray boxes) and the average fitness (white boxes) for all ten runs. Altogether there is a continuous increase of the best and average fitness for the entire range of 850 generations. The fitness values are noisy and obtained running five tests per individual so that they are only a very rough estimate of the attained quality level.




Fig. 4. Connection structures and frequencies by which the structures have been observed at the end of the simulation (in 2,500 tests of the best individual of run 10).

For each evolutionary run, we measure the transport quality of the $\mu = 20$ parent individuals of the last generation on a sample of 500 different tests. This set of μ parent individuals comprises all the genetic material that would be exploited in subsequent generations in case the evolution were continued. For every evolutionary run, the individual exhibiting the highest average of the observed transport quality values is selected. These individuals are post-evaluated for a second time, on a new sample of 2,500 tests.

The median assembling performance $A \in [0, 1]$ for the best individuals for all ten runs is in the range $[0.92, 1.0]$. The corresponding average values range from 0.79 to 0.92. In the following, we focus on the individuals with the highest average distance gain of the two most successful evolutions (here labeled as run 9 and run 10). Figure 4 illustrates all the possible connection structures in which s-bots that are connected to the prey can be organized. Each of these structures has appeared during the post-evaluation of the best individual of run 10 (frequencies indicated in the figure).

Figure 5 presents a box-and-whisker plot of the distance the prey has gained with respect to the beacon as observed in the 2,500 tests for the two best individuals. To evaluate the quality of these results, we proceed as follows.

First we experimentally measure the maximum distance that can be covered in 20 seconds by a single s-bot connected from simulation start (call it $s\text{-bot}_{con}$) to a prey of mass 125 g (250 g): this distance turns out to be 222 cm (141 cm). Then, we assume that within the same time interval (20 s) a group of n $s\text{-bot}_{con}$ connected from start to a prey of weight $n \cdot 125$ g ($n \cdot 250$ g), will transport it for the same maximum distance (i.e., 222 cm (141 cm)). Last, we experimentally observe that the average time for a group of 4 s-bots to self-assemble is 9.9 seconds (6.2 seconds) for the best controller of run 9 (run 10).

Therefore, we can use the distance covered by a group of n $s\text{-bot}_{con}$ within the average time available to the n s-bots for transport as a reference to the distance achievable on average by our group of n s-bots. For the best evolved controller of run 9 (run 10), the time available to a group of 4 s-bots for transporting the prey is 10.1 seconds (13.8 seconds) and the reference distance is 112 cm (153 cm) for the $n \cdot 125$ g prey, and 71 cm (97 cm) for the $n \cdot 250$ g prey (on average).

Observing Figure 5, it is clear that, although in some tests the performance achieved is very good, the average is still rather low and there is surely room for future improvements.




Fig. 5. Distance gain (in cm) within 20 seconds of simulation time by the two best individuals for different prey weights (4 s-bots; 500 observations per box). Average values are indicated by diamonds.

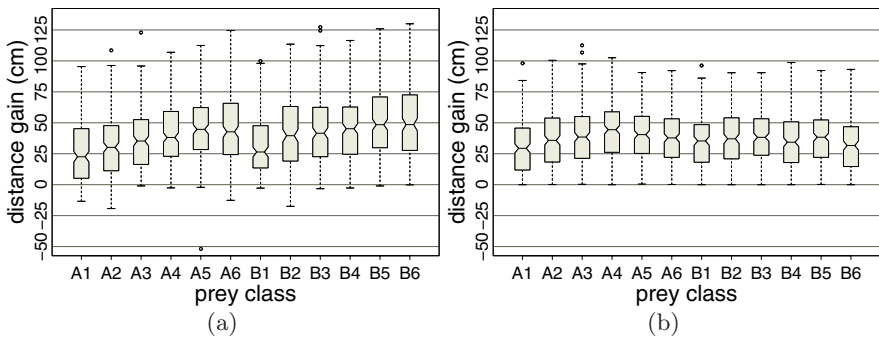


Fig. 6. Sensitivity of performance to different shapes and sizes of prey (4 s-bots; mass of the prey: 500 to 1,000 g): distance gain within 20 seconds by the best controller of (a) run 9 (188-215 observations per box) and (b) run 10 (205-229 observations per box).

Sensitivity to Different Shapes and Sizes of Prey

To assess the flexibility of the best two evolved controllers with respect to the shape and the size of the prey, the data concerning the 2,500 tests performed is partitioned according to different classes of prey (all measures in cm):

- The sets A1, A2, ..., A6 comprise the data concerning cylindrical prey. A_i corresponds to tests in which the cylinders' radii reside in the range $[6 + 1.5(i - 1), 6 + 1.5i)$.
- The sets B1, B2, ..., B6 comprise the data concerning prey of cuboid shape. B1, B2, ..., B6 refer to tests in which the length of the longer side is in the ranges $[12, 16)$, $[16, 20)$, $[20, 25)$, $[25, 30)$, $[30, 35)$ and $[35, 40)$. The length of the shorter side is always in the range $[12, 20)$.

Figure 6 shows the distance gain values observed for different classes of prey. The best individual of run 10 exhibits a decrease in performance if the prey is very small (i.e. for A1, A2 and B1), while the other performs robustly with respect to all combinations of size and shape.




Fig. 7. Ability of the group to adapt the direction of transport according to a new target location for the best individual of (a) run 9 and (b) run 10. The distance gained with respect to the new target is illustrated. Eight different angular displacements of the target with respect to the prey are considered (4 s-bots; mass of the prey: 500 to 1,000 g; 1,250 observations per box).

Moving Target Location

We examine to what extent a group of s-bots engaged in the transport of a prey is able to dynamically adapt the direction of transport according to a new target location. The simulation period of each test is doubled from 20 to 40 seconds. In the first half of the simulation period the experimental setup has been retained unchanged. As soon as the first half of the simulation period is elapsed, the beacon (i.e., the target) is moved to a new location. The (horizontal) angular displacement (radian measure) of the beacon with respect to the prey is chosen randomly in $\{\frac{i}{8}2\pi \mid i \in \{0, 1, 2, \dots, 7\}\}$. For each angular displacement the distance gained during the second part of the simulation is recorded for both individuals on 1,250 tests (see Figure 7). If the target direction does not change, the median distance gain of both individuals is 56 cm. Depending on the angular displacement, the median distance gain of the best individual of run 9 (run 10) may decrease to 17 cm (32 cm).

Scalability

In this section, we study the applicability of evolved controllers to larger groups of s-bots in order to transport prey of bigger weight. We consider groups of 4, 8, 12 and 16 robots. Along with the group size, the mass of the prey is also increased proportionally from 500 to 1,000, 1,500 and 2,000 g respectively.

The two best individuals have been evaluated on 250 tests per group size. Since for larger groups more time is necessary for assembling, the simulation period has been extended to 30 seconds. To ensure a non-overlapping placement of up to 16 s-bots in the immediate vicinity of the prey, the s-bots are initially put randomly within a full circle 25 to 50 cm away from the prey. Given the locality of sensing, it is not possible to position many more s-bots in the environment

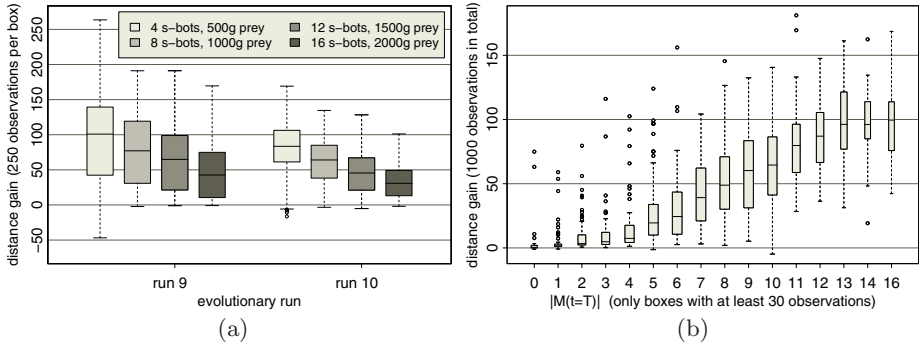


Fig. 8. Distance gain (cm) within 30 seconds of simulation time (a) with 4, 8, 12 and 16 robots transporting prey of mass 500, 1,000, 1,500 and 2,000 g and (b) with respect to $|M(t = T)|$ (i.e. the number of robots physically linked to the prey at time T) for 16 robots transporting prey of mass 2,000 g. The boxes are drawn with widths proportional to the square-roots of the number of observations in the groups.

(without having contacts among them) while ensuring at the same time that each one can locate the prey from its starting point.

For all group sizes the evolved individuals are able to cooperatively move the object. Overall, the performance decreases with group size (see Figure 8 (a)). We observed that a too high density of robots makes it difficult for the robots to self-assemble into structures connected to the prey. However, once the majority of s-bots is assembled, they exhibit a high performance even for group size 16 (see Figure 8 (b)).

4 Conclusions

In this paper we addressed the design of control policies for a group of simple autonomous mobile robots (the s-bots) for the cooperative transport of heavy objects called prey. We used artificial evolution to synthesize neural network controllers in order to let the s-bots localize and approach a prey, self-assemble into structures and pull or push the prey towards a target location. We consider this to be a challenging task; the evolved controllers are responsible for selecting sensory patterns by motor actions, and for controlling the wheels, a gripper device and the orientation of the s-bot’s turret with respect to the chassis. The group has to cope with prey of different shapes and sizes. Moreover, for an s-bot the target can be shadowed by teammates or by the prey.

We have considered prey of cuboid and cylindrical shape with footprints that may differ in size by factors up to 6.25. One of the best individuals has shown to perform robustly with respect to all combinations of size and shape. Furthermore, the evolved controllers can be applied to let a group of robots pull or push a prey towards a target whose location changes in time.

As a final point, we wish to discuss to what extent our study can be considered a swarm robotics work. Recently, four criteria have been identified by Dorigo and

Şahin (2004) in order to distinguish swarm robotics research from other multi-robot studies:

1. The study should be relevant for the coordination of large numbers of robots.
2. The system being studied should consist of relatively few homogeneous groups or robots, and the number of robots in each group should be large.
3. The robots being used in the study should be relatively simple and incapable, so that the tasks they tackle require the cooperation of the individual robots.
4. The robots being used in the study should only have local and limited sensing and communication abilities.

Our system fully meets criteria 2) and 3), and in part² also criteria 4). For testing the fulfillment of criterion 1) further experiments with a greater number of s-bots need to be carried out. In this paper we have shown that for groups of size 4, 8, 12 and 16 robots, the best evolved controllers are able to cooperatively transport prey of mass 500, 1,000, 1,500 and 2,000 g respectively towards a target location. Once the majority of s-bots is physically linked to the prey, also for group size 16 a high performance level is reached. However, the assembling performance seems to decrease as the number of robots is increased. This is likely to be caused by the very high density of robots that are positioned in the immediate vicinity of the prey.

In the future, we want to extend this work by the integration of exploration and recruitment strategies aiming at a complete and scalable prey retrieval solution. In addition, we plan to transfer the solutions from simulation to the real s-bots.

Acknowledgments

This work was supported by the Belgian FNRS, of which Marco Dorigo is a Senior Research Associate, via the grant “Virtual Swarm-bots” (contract no. 9.4515.03), by the Scientific Research Directorate of the French Community of Belgium via an “Action de Recherche Concertée” (the “ANTS” project), and by the Future and Emerging Technologies programme (IST-FET) of the European Commission, via the “Swarm-bots” project (grant IST-2000-31010). 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. The authors also wish to thank all the members of the SWARM-BOTS project for their support, suggestions and comments.

References

J.-L. Deneubourg and S. Goss. Collective patterns and decision-making. *Ethology, Ecology and Evolution*, 1:295–311, 1989.

² The target of transport is visible for several meters (conflict with locality of sensing).

- J.-L. Deneubourg, S. Goss, G. Sandini, F. Ferrari, and P. Dario. Self-organizing collection and transport of objects in unpredictable environments. In *Proceedings of Japan – U.S.A Symposium on Flexible Automation*, pages 1093–1098. ISCIÉ, Kyoto, Japan, 1990.
- M. Dorigo and E. Şahin. Swarm robotics – special issue editorial. *Autonomous Robots*, 17(2–3):111–113, 2004.
- M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2–3):223–245, 2004.
- J. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
- P. Grassé. La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis et Cubitermes sp.* La théorie de la stigmergie : essai d’interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6: 41–81, 1959.
- R. Groß. Swarm-intelligent robotics in prey retrieval tasks. Technical Report TR/IRIDIA/2003-27, IRIDIA, Université Libre de Bruxelles, Belgium, 2003. DEA thesis.
- R. Groß and M. Dorigo. Evolving a cooperative transport behavior for two simple robots. In *Artificial Evolution – 6th International Conference, Evolution Artificielle (EA 2003)*, volume 2936 of *Lecture Notes in Computer Science*, pages 305–317. Springer Verlag, Berlin, Germany, 2004.
- T. Kato and D. Floreano. An evolutionary active-vision system. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC’01)*, pages 107–114. IEEE Computer Society Press, Los Alamitos, CA, 2001.
- C. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1-2):85–101, 2000.
- C. Kube and H. Zhang. Collective robotic intelligence. In *From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 460–468. MIT Press, Cambridge, MA, 1993a.
- C. Kube and H. Zhang. Collective robotics: from social insects to robots. *Adaptive Behaviour*, 2(2):189–218, 1993b.
- D. Marocco and D. Floreano. Active vision and feature selection in evolutionary behavioral systems. In *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*, pages 247–255. MIT Press, Cambridge, MA, 2002.
- F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221, 2004.
- I. Rechenberg. *Cybernetic Solution Path of an Experimental Problem*. PhD thesis, Royal Aircraft Establishment, Hants Farnborough, UK, 1965.
- H.-P. Schwefel. *Evolutionsstrategie und numerische Optimierung*. PhD thesis, Technische Universität Berlin, Germany, 1975.