

Negotiation of Goal Direction for Cooperative Transport

Alexandre Campo, Shervin Nouyan, Mauro Birattari,
Roderich Groß, and Marco Dorigo

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
{acampo, snouyan, mbiro, rgross, mdorigo}@ulb.ac.be

Abstract. In this paper, we study the cooperative transport of a heavy object by a group of robots towards a goal. We investigate the case in which robots have partial and noisy knowledge of the goal direction and can not perceive the goal itself. The robots have to coordinate their motion to apply enough force on the object to move it. Furthermore, the robots should share knowledge in order to collectively improve their estimate of the goal direction and transport the object as fast and as accurately as possible towards the goal.

We propose a bio-inspired mechanism of negotiation of direction that is fully distributed. Four different strategies are implemented and their performances are compared on a group of four real robots, varying the goal direction and the level of noise. We identify a strategy that enables efficient coordination of motion of the robots. Moreover, this strategy lets the robots improve their knowledge of the goal direction. Despite significant noise in the robots' communication, we achieve effective cooperative transport towards the goal and observe that the negotiation of direction entails interesting properties of robustness.

1 Introduction

There are several advantages when using a group of robots instead of a single one. Ideally, the behaviour of a group of robots is more robust, as one robot can repair or replace another one in case of failure. Furthermore, a group of robots can overcome the limitations of a single robot and solve complex tasks than can not be solved by a single robot.

Within collective robotics, swarm robotics is a relatively new approach to the coordination of a system composed of a large number of autonomous robots. The coordination among the robots is achieved in a self-organised manner: the collective behaviour of the robots is the result of *local* interactions among robots, and between the robots and the environment. The concept of locality refers to a situation in which a robot alone can not perceive the whole system. Each single robot typically has limited sensing, acting and computing abilities. The strength of swarm robotics lies in the properties of robustness, adaptivity and scalability of the group [1].

Foraging has been outlined as a canonical problem by Cao *et al.* [2] among those studied in collective robotics and is an important topic in swarm robotics

too. In foraging, a group of robots has to pick up objects that are scattered in the environment. The foraging task can be decomposed in an exploration subtask followed by a transport subtask. Foraging can be applied to a wide range of useful tasks. Examples of applications are toxic waste cleanup, search and rescue, demining and collection of terrain samples.

Central place foraging is a particular type of foraging problem in which robots must gather objects in a central place. Borrowing the terminology from biology, the central place is also called the *nest* and the objects are called *prey*. We focus on a specific case in which the transport of a prey requires the combined effort of several robots. This task is called cooperative transport. Several problems need to be solved to perform this task successfully. The coordination of the movement of the robots is one of them. This problem has been investigated by Groß *et al.* [3], in situations in which either all or some robots are able to perceive the nest.

In this paper we address the case in which all robots completely lose sight of the nest during the exploration subtask of foraging. We assume that the robots have partial knowledge of the goal direction. For instance, they may have perceived the nest earlier and kept track of its direction by means of odometry [4]. Odometry is achieved using internal, proprioceptive information [5] (*e.g.*, by measuring the rotation of the wheels of a robot). The information on the movement of a robot is integrated, thus the error made on localization increases with the distance covered. In our case, this leads to an erroneous indication about the direction of the nest. If several robots attempt to transport a heavy prey in different directions they may fail to move the prey at all. Therefore, we introduce a mechanism to let the robots negotiate the goal direction. In order to meet the general principles of swarm robotics [1], this system is fully distributed and makes use of local communication only.

The mechanism we introduce is strongly inspired by a natural mechanism that has been long studied by biologists. We rely on a particular property of models designed to explain and reproduce fish schools and bird flocks [6,7,8,9]. The models available in the literature are usually composed of three behaviours: an attraction behaviour that makes the individuals stick together, a repulsion behaviour that prevents collisions among individuals, and an orientation behaviour that coordinates the individuals' motion. It is the last of these three behaviours that we transfer and implement in our robots. Informally, the orientation behaviour lets every individual advertise locally its own orientation and update it using the mean orientation of its neighbours.

We conduct experiments with a group of four real robots that have to transport a prey moving in a direction about which they have noisy knowledge. We assess quantitatively the performance of the negotiation mechanism implemented with respect to different levels of noise and different control strategies.

In Section 2 we detail the task, the hardware, the experimental setup and the different controllers. Section 3 is devoted to the presentation of the experimental results. Section 4 concludes the paper with a discussion of the results and some ideas for future work.

2 Methods

The Task. The task is the cooperative transport of a heavy prey towards a nest by a group of four real robots. The robots are physically connected to the prey using their grippers. The nest is out of sight and the robots have no means to perceive it. The initial knowledge of each individual about the goal direction is provided with a given amount of noise.

The mass of the prey is chosen such that a single robot can not transport it. At least three robots are necessary to move the prey. A high degree of coordination of the robots' motion is required to apply enough force to the prey to move it. If the robots lack coordination, that is, if they pull in different directions, they may not be able to move the prey at all.

The robots can share knowledge using visual communication in order to collectively improve their estimate of the goal direction and transport the prey as fast and as accurately as possible towards the goal.

Hardware. *The robots:* We use the *s-bot* (Figure 1(a)), a robot of 12 cm of diameter, designed and built within the context of the SWARM-BOTS project [10,11]. An *s-bot* moves using a combination of two wheels and two tracks, which we call "trees". This system notably allows the robot to efficiently turn on the spot. The robots can physically connect to a prey or to another *s-bot* using their grippers. They are supplied with a rotational base that lets them move in an arbitrary direction while maintaining the same physical connection pattern. The robots can send visual information by means of eight triplets of red, green and blue LEDs. The LEDs are positioned on a ring around the robot. An *s-bot* activating its LEDs can be perceived by another *s-bot* by means of an omnidirectional camera which provides a 360° view.

The prey: The mass of the prey is 1.5 kilograms. At least three robots are necessary to effectively pull the prey. This weight of the prey makes the transport by a group of robots very difficult if the robots are not well synchronised.

Experimental Setup. The experiments take place in an open space. Initially, four robots are connected to the prey in a regular arrangement, thus forming a cross pattern as shown in Figure 1(c). We test four levels of noise on the robots' initial estimate of the goal direction: *no noise* (0), *low noise* (L), *medium noise* (M) and *high noise* (H). In the case of no noise, the initial direction of the robots is the same and points towards the nest.

The initial imprecise knowledge of the robots about the direction of the nest is modeled by a random number drawn from a von Mises distribution, which is the equivalent of the Gaussian in circular statistics [12], and well suited for directional data. This distribution is characterised by two parameters μ and κ . The direction to the nest is indicated by μ , the mean of the distribution. The level of noise is indicated by κ . The smaller κ , the more the distribution resembles a uniform distribution in $[-\pi, \pi]$. When κ is large, the distribution resembles a Gaussian of mean μ and standard deviation σ , when $\kappa \rightarrow \infty$ the relationship $\sigma^2 = \sqrt{1/\kappa}$ holds. The three levels of noise L, M, H correspond to $\kappa = 3, 2, 1$, as displayed in Figure 2(a).

After each trial, the robots are randomly permuted, so that the possible differences among robots are averaged out and can be neglected in this study. We tested 4 possible goal directions of 0, 22.5, 45 and 67.5°. Any direction above 90° is redundant as the pattern of connected robots (a cross) is symmetrical on the two perpendicular axes, and the robots are permuted at each trial. Finally, we have tested 4 possible strategies for the robots to transport the prey towards the goal (see next section for more details). In total, we performed 256 replications: we tested 4 goal directions, 4 levels of noise and 4 distinct strategies for transport. Each combination of the aforementioned parameters was tested 4 times.

To extract the results, we used a camera placed above the initial position of the prey to record videos of each trial. The experiment is stopped either when the prey has been transported to a distance of 1 meter from its initial position or after 60 seconds (an average transport takes approximately 20 seconds). A trial can also be stopped if we judge that the robots are stuck in a situation that is potentially harmful to their hardware. Indeed, if the robots do not manage to coordinate their movements, they may pull in opposite directions and thus induce a high torque to their grippers. One gripper was broken during the experiments reported here, and we wished to avoid as much as possible any further damage. Any experiment stopped without the prey being transported for more than 1 meter of distance from the initial position is considered as a transport failure.

For each trial, we have extracted the position of the prey at each time step (5 pictures per seconds) using a simple tracking software. Using these data, we have categorized the trials in transport failure or success, and measured the duration of all trials. Furthermore, we measured the angular difference between the direction in which the prey has been moved and the goal direction, as shown in Figure 2(b). Later on, we also use the term deviation to refer to this angular difference.

Robot’s Controller. *Vision software:* We employ a specific vision software that allows a robot to perceive the direction pointed by a neighbour in his visual range. The perception algorithm implemented in the software is probabilistic and approximates the directions communicated by the local neighbours using a triangular pattern shown by the LEDs (see Figure 1(b)). In order to assess the quality of the vision software, we have performed a series of basic tests. We have run in total 8000 times the vision software on 8 different pictures to obtain a distribution of direction estimates. Figure 3 summarises the pooled results of the tests for a communicated direction pointing towards direction 0. The tests show that it is possible to achieve a reliable estimate of the direction pointed by neighbouring robots. As the mechanism of negotiation of direction should be robust to noise, there is no need to improve the output of the vision software with any kind of signal filter. We directly feed the negotiation mechanism with a single estimate.

Negotiation mechanism: The negotiation mechanism is bio-inspired and implemented in a straightforward manner, following closely the rules used to model the orientation behaviour of fish schools or bird flocks [9]. Let n be the total number of robots. For each robot $i \in [1, n]$, let $\mathcal{N}_i(t)$ be the set of robots in the




Fig. 1. (a) The *s-bot*. (b) An *s-bot* displaying a direction using a triangular LED pattern. (c) Star-like formation of four *s-bots* around the prey as used in the experiment.

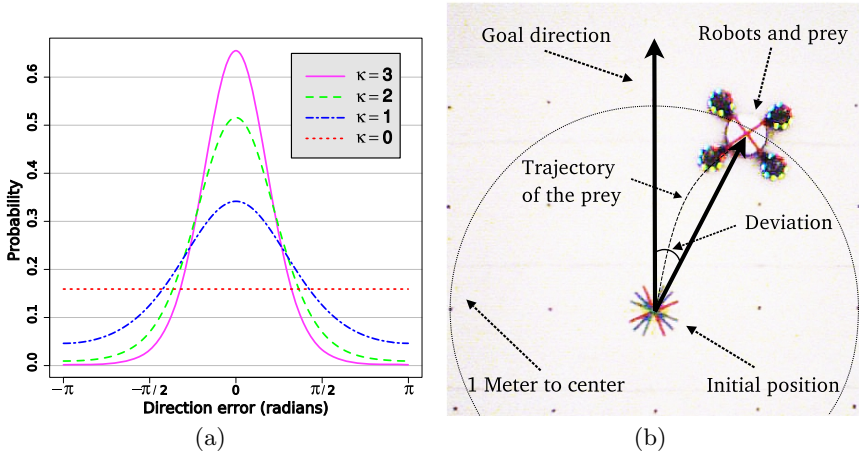


Fig. 2. (a) The effect of parameter κ on a von Mises distribution. (b) A snapshot describing the final situation of a successful transport. Note how the deviation of the transport direction from goal direction is measured.

visual range of robot i at time t . This defines the topology of the communication network. Let $d_i(t) \in [-\pi, \pi]$ be the goal direction estimated by robot i at time t . Let $D_j^i(t) = d_j(t) + \epsilon_j^i(t)$ with $j \in \mathcal{N}_i(t)$ the direction of robot j perceived by robot i assuming noise $\epsilon_j^i(t)$.

If robot i communicates and exchanges information with his neighbours, it will calculate what we call a desired direction \bar{d}_i by using Equation 1 that basically computes a mean direction. To do so, we use the sum of unit vectors, which is a classical method in circular statistics [12]:

$$\bar{d}_i(t) = \arctan \left(\frac{\sin(d_i(t)) + \sum_{j \in \mathcal{N}(i)} (\sin(D_j^i(t)))}{\cos(d_i(t)) + \sum_{j \in \mathcal{N}(i)} (\cos(D_j^i(t)))} \right). \quad (1)$$




Fig. 3. We used four robots arranged in a circular pattern to display a common direction and a central robot to take pictures and estimate the direction pointed by the surrounding robots. For eight distinct directions, the central robot produced 1000 estimates each. The resulting pooled distribution of errors shows that the vast majority of the estimates matches the direction pointed by the surrounding robots.

The estimate of the goal direction of a particular robot is not updated directly. Indeed, the noise present in perception might induce oscillations if the update of the robots' estimates is done too fast. Therefore, we use a damping factor δ to stabilise the system (we chose $\delta = 0.05$ for our experiments). The update of the estimate of the goal direction for robot i is described by Equation 2:

$$d_i(t + \Delta t) = (1 - \delta) \cdot d_i(t) + \delta \cdot \bar{d}_i(t). \quad (2)$$

The motion control of each robot is implemented by a simple algorithm [3] that sets the speed and orientation of the robot's *treels* to pull the prey in the estimated direction d of the nest.

Control strategies: We have defined and implemented four distinct strategies. To refer to the strategies, we employ a notation in which **T** means transport, **N** means negotiation, and **:** marks the end of an optional and preliminary negotiation phase. If this preliminary phase takes place, it lasts 30 seconds. The second phase always involves transport and lasts 60 seconds.

- **Transport directly (T):** a naive strategy that we use as a yardstick to show the improvement brought by the negotiation mechanism. The robots move along their initial direction. No communication and no update of the estimated direction is done.
- **Negotiate then transport (N:T):** robots first negotiate their estimate of the direction of the goal for 30 seconds without moving. Afterwards, they all start moving without either communicating or updating their estimates.
- **Negotiate then transport and negotiate (N:NT):** robots start by negotiating the direction of the goal for 30 seconds without moving. After this preliminary negotiation, they all start moving and at the same time they keep on negotiating together.

- **Negotiate and transport (NT)**: from the very beginning of the experiment, the robots start both moving and negotiating.

At the beginning of the experiments, robots have each a rough estimate of the direction of the goal, but they never perceive directly the goal. The three last strategies may appear identical to the reader, but in fact two important aspects, namely time and noise in communication should be considered. On one hand, the duration of the negotiation process affects the degree of synchronisation of the robots. On the other hand, visual communication is imperfect. When robots don't move, errors in visual communication are persistent and may have a strong impact on the outcome of the negotiation process. When robots move, they modify slightly their relative locations and this results in a reduction of the errors in visual communication.

3 Results

We report here the experimental results of the task of cooperative transport for all the strategies and levels of noise tested. We examine three different aspects of the system: the ability of the system to succeed in transporting the prey for a certain distance, the duration of transport and the accuracy in direction of transport. Data analysis was performed with the R software and the package *circular* [13].

Success in Transporting. We first study the ability of the robots to transport the prey. If the robots are not able to move the prey over a distance of at least 1 meter from the initial position within 60 seconds, we consider the trial as a transport failure. Figure 4 presents the performances in transport of the four strategies for the different levels of noise.

First, we observe that in absence of noise (level θ), the robots manage very well to transport the prey without negotiating the direction. Therefore, negotiation is not necessary and it is desirable that strategies employing the negotiation mechanism do not perform worse. The strategy $N:T$ yields only 75% of successful transports when there is no noise in the initial direction of the nest. When this strategy is employed, it is possible that negotiation is stopped while robots are not perfectly coordinated and no further correction can be done on the direction of the robots. The two other strategies $N:NT$ and NT do not decrease the capability of the group of robots to transport the prey with respect to strategy T . We have observed that during motion, the formation of robots can alter slightly, mainly due to slippage of the grippers on the prey. Strategies using negotiation during transport allowed robots to quickly correct their direction and remain coordinated. Conversely, the strategies T and $N:T$ were very sensitive to small errors.

When noise is present, the performance of the group of robots using strategy T decreases. For medium and high noise it is close to 10%. This result was expected as robots are not able to coordinate their motion at all and are initialised with different initial directions. We also notice that, although noise




Fig. 4. The percentage of successful and failed transports grouped by strategy and by level of noise

has a non neglectable impact on the transport capability, the performances stay quite similar for different levels L , M and H of noise considering the strategies $N:T$, $N:NT$ and NT . All strategies relying on the negotiation mechanism achieve better performances, and especially strategy NT is much less sensitive to noise than the others.

Duration of Transport. We focus now on the duration of the transport. For all the trials, we consider whether or not transport is successful. Figure 5 shows for all strategies and all levels of noise boxplots of the duration of the 16 transport tasks. Note that we do not take into account the preliminary negotiation period that lasts 30 seconds when strategies $N:T$ or $N:NT$ are employed.

Once again, the performance of strategy T in absence of noise is the best with respect to any other pair of strategy and level of noise. Only strategy NT reaches a comparable performance.

When the level of noise increases, the duration of transport of the strategy T increases too, in a quasi linear manner. Strategies that rely on the negotiation mechanism are much less sensitive to noise. The duration of transport using those strategies is very similar for the different levels of noise L , M and H , but strategy $N:T$ has produced more failures. Because robots can not correct their coordination with this strategy, they easily rotate while transporting the prey. This constant error produces round or even circular trajectories and prevents the robots to quickly move the prey away from its initial position. Strategies can be clearly ranked: the slowest ($N:T$), the average ($N:NT$) and the fastest (NT).

Deviation From the Direction of the Nest. The last measure we study is the deviation of the direction of transport with respect to the direction of the nest. Again, we take into account all trials. The study of deviation from the direction of the nest confirms all previous observations (see Figure 6).

In absence of noise, the naive strategy T performs very well, and the only other strategy with a comparable result is strategy NT . When noise is introduced, the performance of strategy T decreases. The strategies that make use of negotiation perform better, and show only small differences for the different levels of noise tested. Among these strategies, the best is NT .

We have fitted von Mises distributions with the distributions of deviations in order to study strategy NT in further detail. The fit with a von Mises distribution




Fig. 5. Box-and-whisker plot [14] showing the duration of transport of the prey (in seconds), taking into account successful and failed transports. The distributions are grouped by strategy and by level of noise.

yields an estimate of κ , which corresponds to the error of the transport direction. The robots start with an initial knowledge affected by a noise that corresponds to an individual error of respectively 33.1° , 40.5° and 57.3° . After the application of the strategy *NT*, the final values of the error of transport direction for the levels of noise *L*, *M*, *H* are respectively $42.8^\circ \pm 76.2^\circ$, $42.3^\circ \pm 75.5^\circ$ and $42.5^\circ \pm 75.7^\circ$ (degrees \pm standard error). These values are not significantly different. Moreover, it is observed that for the level of noise *H*, the strategy *NT* improves the robots' estimate of the direction of the nest.

4 Discussion

Achievements. We have compared different strategies to achieve efficiently the cooperative transport of a prey with partial knowledge of the direction of the nest. We performed systematic experiments to evaluate the characteristics of the different strategies under study for four distinct levels of noise. The comparison of the strategies has shown that negotiation during transport of a prey improves the coordination of motion. It has also been shown that negotiation without moving prior to transport (*N:NT*) performs worse than the straightforward strategy *NT* consisting in negotiating and transporting the prey at the same time.

It has been observed that the strategy *NT* is neutral: if negotiation is not mandatory to achieve efficient transport, making use of this strategy does not alter the transport performances with respect to the naive strategy *T*. Hence, it is not necessary to choose which strategy to employ depending on the level of noise affecting robots' knowledge of the direction of the nest. The strategy *NT* can be used at any time.

Besides the coordination of motion, our experimental results have also shown that the group of robots could improve their knowledge of the direction of the nest by means of visual negotiation. Strategy *NT* improves the robots' estimate of the direction of the nest and shows no discernible difference of the errors for the levels of noise *L*, *M*, *H*. The improvement of the accuracy of direction of transport with respect to the *s-bots* initial knowledge is most striking when the level of noise is high.




Fig. 6. Box-and-whisker plot [14] showing the average deviation (in radians) between direction of transport and direction of the nest. Both successful and failed transports are taken into account. The distributions are grouped by strategy and by level of noise.

Improvements. With respect to mechanisms of coordination of motion that use a traction sensor [3,15], our system appears to be more flexible, as visual communication is also available when the prey is not in motion, and it is not distracted if the prey moves in irregular steps. Additionally, visual communication leaves the door open to collective motion with or without transport or physical connections. The topology of the network of communications is also likely to be more flexible, allowing the robots to school in very diverse patterns.

The negotiation mechanism we have introduced is not only able to supply a group of robots with collective motion, but also to let each individual improve its own estimate of the goal direction by sharing knowledge with its neighbours. This mechanism may also be used to correct measures of odometry in multi-robot experiments, in a fully distributed fashion. This self-organised negotiation is likely to display properties of scalability besides the robustness shown in this paper.

Perspectives and Future Work. The difference in performance between strategies $N:NT$ and NT is counter-intuitive, as the negotiation in the first strategy lasts in total longer and thus the robots are expected to achieve a better performance because they are granted more time to negotiate. However, the preliminary phase of negotiation without movement negatively affects the performance of the robots. This might be due to persistent errors in visual communication in absence of movement. It is also possible that transport is more efficient when robots align their tracks gradually, and not immediately as it happens with strategy $N:NT$. We plan to further investigate the exact reasons for this phenomenon.

We plan to integrate the cooperative transport in a more complex and challenging scenario of foraging, such as for instance the one used by Nouyan *et al.* [16]. This scenario would include an exploration phase preliminary to transport, in which robots lose sight of the nest before finding the prey. In this context, robots have a rough estimate of the direction of the nest by means of odometry. Improvement of this knowledge by means of negotiation is a critical feature of the scenario, necessary to let the robots transport the prey efficiently to the nest, even in presence of noisy communications and failed robots.

Acknowledgements. This research has been supported by COMP²SYS, a Marie Curie Early Stage Research Training Site funded by the European Community's Sixth Framework Programme under contract number MEST-CT-2004-505079, and by the ANTS project, an *Action de Recherche Concertée* funded by the Scientific Research Directorate of the French Community of Belgium. Marco Dorigo acknowledges support from the Belgian FNRS of which he is a Research Director. The information provided is the sole responsibility of the authors and does not reflect the opinion of the sponsors. The European Community is not responsible for any use that might be made of data appearing in this publication.

The authors would like to thank Jean-Louis Deneubourg for fruitful discussions.

References

1. Dorigo, M., Şahin, E.: Swarm robotics – special issue editorial. *Autonomous Robots* **17**(2–3) (2004) 111–113
2. Cao, Y.U., Fukunaga, A.S., Kahng, Andrew, B.: Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots* **4**(1) (1997) 7–27
3. Groß, R., Mondada, F., Dorigo, M.: Transport of an object by six pre-attached robots interacting via physical links. In: Proc. of the 2006 IEEE Int. Conf. on Robotics and Automation, IEEE Computer Society Press, Los Alamitos, CA (2006) 1317–1323
4. Dieter, F., Wolfram, B., Hannes, K., Sebastian, T.: A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots* **8**(3) (2000) 325–344
5. Borenstein, J., Feng, L.: Measurement and correction of systematic odometry errors in mobile robots. *IEEE Trans. on Robotics and Automation* **12**(5) (1996) 869–880
6. Aoki, I.: A simulation study on the schooling mechanism in fish. *Bulletin of the Japanese Society of Scientific Fisheries* **48**(8) (1982) 1081–1088
7. Reynolds, C.W.: Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics* **21**(4) (1987) 25–34
8. Huth, A., Wissel, C.: The simulation of the movement of fish schools. *Journal of Theoretical Biology* **156** (1992) 365–385
9. Couzin, I.D., Krause, J., James, R., Ruxton, G.D., Franks, N.R.: Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology* **218**(1) (2002) 1–11
10. Dorigo, M., Tuci, E., Groß, R., Trianni, V., Labella, T., Nouyan, S., Ampatzis, C., Deneubourg, J.L., Baldassarre, G., Nolfi, S., Mondada, F., Floreano, D., Gambardella, L.: The SWARM-BOTS project. In E. Şahin, Spears, W., eds.: Proceedings of the First International Workshop on Swarm Robotics at SAB 2004. Volume 3342 of LNCS., Springer Verlag, Berlin, Germany (2004) 31–44
11. Dorigo, M.: Swarm-bot: An experiment in swarm robotics. In Arabshahi, P., Martinioli, A., eds.: Proceedings of SIS 2005 – 2005 IEEE Swarm Intelligence Symposium, IEEE Press, Piscataway, NJ (2005) 192–200
12. Jammalamadaka, S.R., SenGupta, A.: Topics in Circular Statistics. World Scientific Press, Singapore (2001)
13. R Development Core Team: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. (2005)

14. J. M. Chambers, W. S. Cleveland, B. Kleiner, P. A. Tukey: Graphical Methods for Data Analysis. The Wadsworth statistics / probability series. Wadsworth and Brooks/Cole, Pacific Grove, CA (1983)
15. Trianni, V., Nolfi, S., Dorigo, M.: Cooperative hole avoidance in a *swarm-bot*. *Robotics and Autonomous Systems* **54**(2) (2006) 97–103
16. Nouyan, S., Groß, R., Bonani, M., Mondada, F., Dorigo, M.: Group transport along a robot chain in a self-organised robot colony. In: Proc. of the 9th Int. Conf. on Intelligent Autonomous Systems, IOS Press, Amsterdam, The Netherlands (2006) 433–442