

Hole Avoidance: Experiments in Coordinated Motion on Rough Terrain

Vito Trianni[§], Stefano Nolfi[†], and Marco Dorigo[§]

[§]*IRIDIA - Université Libre de Bruxelles, Bruxelles, Belgium*

[†]*Institute of Cognitive Sciences and Technologies - CNR, Roma, Italy*

Abstract In this paper, we study coordinated motion in a swarm robotic system, called a *swarm-bot*. A *swarm-bot* is a self-assembling and self-organizing artifact, composed of a swarm of *s-bots*, mobile robots with the ability to connect to and disconnect from each other. The *swarm-bot* concept is particularly suited for tasks that require abilities of navigation on rough terrain, such as space exploration or rescue in collapsed buildings. In fact, a *swarm-bot* can exploit the cooperation of its simple components to overcome difficulties or avoid hazardous situations. As a first step toward the development of more complex control strategies, we investigate the case in which a *swarm-bot* has to explore an arena while avoiding to fall into holes. In order to synthesize the controller for the *s-bots*, we rely on artificial evolution, which proved to be a powerful tool for the production of simple and effective solutions to the hole avoidance task.

1 Introduction

The first problem to be faced when trying to control an autonomous robot is to make it move efficiently in a given environment. Depending on the robot, this task can be rather simple (i.e., the motion of a wheeled robot) or particularly complex (i.e., walking for a humanoid robot). Also the environment in which the robot is placed influences the complexity of the problem: a flat terrain is clearly less challenging than a rough terrain with holes and obstacles. An additional source of complexity is found in the *coordinated motion* task, in which the robotic system is composed of a number of independent entities that have to coordinate their actions in order to move coherently.

Coordinated motion is a well studied behavior in biology, being observed in many different animal species. For example, we can think of flocks of birds coordinately flying, or of schools of fish swimming in perfect unison. These examples are not only fascinating for the charming patterns they create, but they also represent interesting instances of self-organized behaviors. Many researchers have provided models for schooling behaviors of fish, and replicated them in artificial life simulations (see [2], chapter 11). Similarly, groups of artificial fish (called *e-boids*) have been evolved to display schooling behaviors, obtaining interesting results [13]. Finally, evolutionary computation has been used also to evolve coordinated motion behaviors in small groups of physical robots [9].

Coordinated motion is a problem of fundamental importance within the SWARM-BOTS project,¹ wherein this research is conducted. The SWARM-BOTS project aims at the develop-

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

ment of a new robotic system, called a *swarm-bot* [12]. A *swarm-bot* is defined as an artifact composed of simpler autonomous robots, called *s-bots*. An *s-bot* has limited acting, sensing and computational capabilities, and can create physical connections with other *s-bots*, thereby forming a *swarm-bot* that is able to solve problems the single individual cannot cope with. Coordinated motion is a basic ability that the *swarm-bot* should display: a *swarm-bot* should move coherently across the environment as a result of the cooperation of the *s-bots* assembled in a single structure [1].

Another basic ability for the *swarm-bot* is coping with rough terrains, holes, gaps or narrow passages. Navigating on rough terrain is an important feature for an intelligent autonomous system, that can open many possible application scenarios, like space exploration or rescue in a collapsed building. Research in this direction has focused mainly on the development of rovers provided with articulated wheels or tracks, like the *pathfinder* [7]. A different approach to rough terrain navigation is presented by reconfigurable robotics, where robots can adopt different shapes in order to cope with different environmental conditions [3, 10, 14].

In the *swarm-bot* case, navigation on rough terrain is achieved by means of the cooperation between *s-bots* which can self-assemble and build structures that can cope with hazardous situations like avoiding a hole or passing over a trough. In such cases, rigid connections serve as support for those *s-bots* that are suspended over the gap. This approach to rough terrain navigation also has a natural counterpart in ants of the species *Ecophilla longinoda* [6], which are able to build chains connecting one to the other, creating bridges that facilitate the passage of other ants.

In this paper, we study an instance of the family of “navigation on rough terrain” tasks, that is, *hole avoidance*. A *swarm-bot* has to perform coordinated motion in an environment that presents holes too large to be traversed. Thus, holes must be recognized and avoided, so that the *swarm-bot* does not fall into them. The difficulty in this task is twofold: first, *s-bots* should coordinate their motion. Second, *s-bots* have to recognize the presence of an hole, communicate it to the whole group and re-organize to choose a safer direction of motion.

The rest of this paper is organized as follows: Section 2 describes our approach to the study of the hole avoidance problem. Section 3 and 4 are dedicated to the description of the obtained results. Finally, Section 5 concludes the paper.

2 Evolution of Hole Avoidance Behaviors

In this paper, the *s-bots* controllers are obtained using *artificial evolution*. There are multiple motivations that lay behind this choice for synthesizing controllers for a robot [8]. In particular, in a distributed multi-robot context as the one considered within the SWARM-BOTS project, handcrafting the controllers may be too complex. Here, artificial evolution can bypass this difficulty, as it directly tests the behavior displayed by the robots embedded in their environment. Furthermore, artificial evolution can exploit the richness of solutions offered by the complex dynamics resulting from robot-robot and robot-environment interactions [12].

Figure 1a shows the current *s-bot*.² In this paper, however, experiments are performed in simulation, using a software based on VortexTM, a 3D rigid body dynamics 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 we considered most important (see Figure 1b).

²Details regarding the hardware and simulation of the *swarm-bot* can also be found in the project web-site (www.swarm-bots.org).




Figure 1: (a) The *s-bot* prototype, provided of the tracks system, the body holding the rigid and the flexible grippers, and many sensors. (b) The simulated *s-bot* model. The body is transparent to show the chassis (center sphere), the motorized wheels (lighter spherical wheels) and the passive wheels (darker spherical wheels). The position of the virtual gripper is shown with an arrow painted on the *s-bot*'s body. Ground sensors are displayed as lines exiting from the *s-bot*.

The simulated *s-bot* is composed of a cylindrical turret (radius: 6 cm, height 6 cm), connected to a chassis by a motorized hinge joint. The chassis is a sphere (radius: 1.4 cm) to which 4 spherical wheels are connected (radius: 1.5 cm). The lateral wheels are connected to the chassis by a motorized joint and a suspension system and they are responsible for the motion of the *s-bot*. The front and back wheels are passive. Connections between *s-bots* are simulated creating a joint between the two bodies.

Each *s-bot* is provided with a *traction sensor* placed at the turret-chassis junction. It detects the direction and the intensity of the traction force that the turret exerts on the chassis. The traction sensor, integrating all the pulling/pushing forces created by the movement of the connected *s-bots*, provides an indication of the average direction toward which the *swarm-bot* is trying to move as a whole.³ Besides traction sensors, we also make use of 4 *ground sensors*, which are infrared proximity sensors evenly distributed around the chassis of the *s-bot* and pointed toward the ground.

Concerning the actuators, each *s-bot* can control its wheels independently. The maximum angular speed has been set to 10 rad/s, which corresponds to a maximum speed of the *s-bot* of 0.15 m/s. In addition, the movements of the *s-bot* are also influenced by the turret/chassis motor. This motor is controlled setting its desired angular speed as half of the difference between the desired angular speed of the left and right wheels. This setting helps the rotation of the chassis with respect to the turret also when one or both wheels of the *s-bot* do not touch the ground [1].

In order to study the hole avoidance task, we designed a square arena (side 3 m) that contains 4 evenly distributed square holes (side 60 cm, see Figure 2b). The *swarm-bot* consists of a linear structure made of 4 *s-bots*, which are rigidly connected by means of their virtual grippers. Each *s-bot* is controlled by a simple perceptron, a neural network connecting its sensory inputs to the motor outputs. The network has 8 sensory inputs: 4 are dedicated to the readings coming from the ground sensors, and the other 4 encode the intensity and direction of traction (for more details, see [1, 11]). Moreover, the neural network is provided with one bias unit and 2 outputs that control the two wheels and the turret/chassis motor. This perceptron has in the whole 18 connections, whose weights are evolved.

³This particular kind of sensor proved to be of fundamental importance for the evolution of coordinated motion in a *swarm-bot* [1, 11].

We use a generational evolutionary algorithm. The initial population is composed of $\mu = 100$ randomly generated genotypes. Each genotype is binary encoded, and is mapped into a neural network controller for a single *s-bot*. Each weight, ranging in the interval $[-10, 10]$, was represented in the genotype by 8 bits, corresponding to a genotype length $L = 18 \times 8 = 144$ bits. This controller is cloned in each of the $n = 4$ *s-bots* involved in the experiment. The fitness F of each genotype is estimated allowing the group of *s-bots* to “live” for $M = 5$ “epochs” and then averaging the obtained value. The best $\lambda = 20$ genotypes of each generation are allowed to reproduce, each generating $\mu/\lambda = 5$ offspring. Each of their bits has a probability $2/L$ of being flipped. Parents are not copied to the offspring population (no elitism). An evolutionary experiment lasts 100 generations. This algorithm is very simple and straightforward, and we found that it is sufficient to evolve simple but efficient controllers for groups of robots [12, 1].

The fitness function is designed to favor coordinated motion, exploration of the arena and a fast reaction to the detection of an hole. The fitness estimation F_e in each epoch is given by the average of two components, F_{e_1} and F_{e_2} (see below). In order to compute the fitness components, we divide each epoch e into two sub-epochs, e_1 and e_2 . In the former, we test the genotype for its ability to perform coordinated motion in a flat environment. Here the *s-bots* start connected in a linear formation, having the orientation of their chassis randomly initialized. They are selected for the ability to move as far as possible from their initial position, which indirectly implies an ability to display coordinated movements. Therefore, the fitness estimation F_{e_1} is computed as the distance covered by the group. The sub-epoch e_1 lasts $T_{e_1} = 150$ simulation cycles, each cycle corresponding to 100 *ms* of real time.

In sub-epoch e_2 , *s-bots* are positioned at the center of the arena with holes, and start in the usual chain configuration. Their chassis are all initialized with the same random orientation. Also the chain is randomly oriented at the beginning of each sub-epoch. In this way, there is no need for a coordination phase at the beginning of the sub-epoch, the focus being on hole avoidance. The sub-epoch lasts $T_{e_2} = 200$ cycles. The fitness estimation F_{e_2} is given by the product of two sub-components: the *survival* sub-component F_s and the *exploration* sub-component F_x . The former rewards only those genotypes that reach the end of the epoch without falling into a hole. This sub-component penalizes every fall, even if it happens at the end of the sub-epoch, thus favoring more robust behaviors. The second sub-component is designed to favor those genotypes that are able to better explore the arena. In this case, the arena is virtually divided in 25 square zones of 60 *cm* side. The genotype is rewarded proportionally to the percentage of visited zones during the sub-epoch (for more details, see [11]).

3 Obtained Results

In this section, we present the results obtained evolving hole avoidance behaviors using the controller described above. We replicated the evolutionary experiment 10 times. The average fitness values, computed over all the replications, are shown in Figure 2. The average performance of the best individual and of the population are plotted against the generation number. The plot indicates that the evolutionary experiments were successful: the average fitness value of the best individuals reaches the 80% of the theoretical maximum value, which cannot be

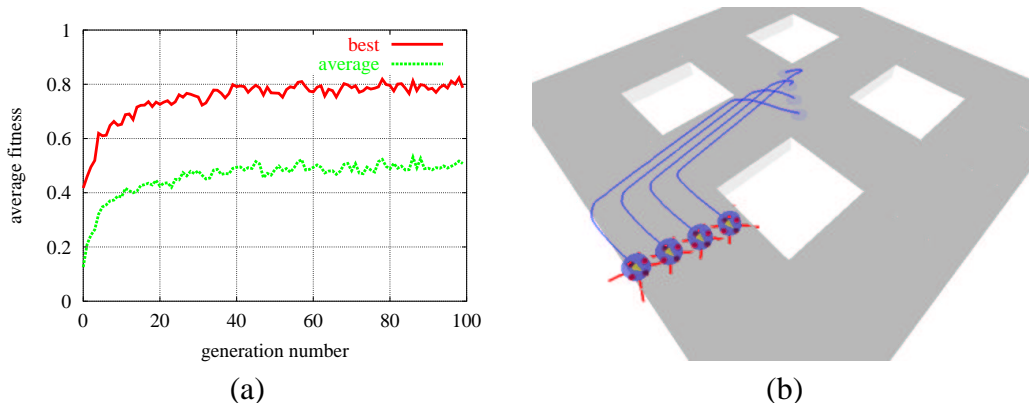


Figure 2: Hole avoidance results: (a) Average fitness over 10 replications of the experiment. (b) Trajectories displayed by a *swarm-bot* performing a hole avoidance task.

achieved due to the particular experimental setup.⁴

In order to test the performance of the evolved controllers, we evaluated the best individuals of the last generation of each replication of the experiment. The corresponding results are shown in Table 1. All individuals perform reasonably well, even if it can be noted that the average performance of every controller is lower than the average value achieved at the last generation of the evolutionary runs, showed in Figure 2a. This is due to the small sampling size used for the estimation of the fitness during the evolution (5 epochs). In fact, a small sampling size usually leads to an over-estimation of the fitness of the best individual. Thus, the post-evaluation analysis with a larger sampling size (100 epochs, in this case) gives a better approximation of the real performance of the evolved controller.

Table 1: Mean performance of the best individuals for each replication of the experiment, averaged over 100 epochs. The best evolved individual is highlighted in bold.

Replication	1	2	3	4	5
Performance	0.6640	0.6541	0.6502	0.6079	0.5835
Replication	6	7	8	9	10
Performance	0.6376	0.6866	0.6397	0.6640	0.6458

Direct observation of the behaviors evolved showed that all solutions rely on similar strategies. We observed the evolved behaviors placing the *swarm-bot* in the arena with holes, and starting with different orientation of the chassis of the *s-bots*.⁵ At the beginning, the *s-bots* start to move in the direction they were positioned, resulting in a rather disordered overall motion. Within few simulation cycles, the physical connections transform this disordered motion into traction forces, that are exploited to coordinate the group. When an *s-bot* feels a traction force, it rotates its chassis in order to cancel this force. Once the chassis of all the *s-bots* are oriented in the same direction, the traction forces disappear and the coordinated motion of the *swarm-bot* starts (see Figure 2b). Then, when one *s-bot* detects an edge, it rotates the chassis and changes the direction of motion in order to avoid falling. This change in direction creates

⁴The theoretical maximum value could be reached only if in the first sub-epoch *s-bots* started with their chassis perfectly aligned, so that no coordination phase is required, allowing the *swarm-bot* to cover the maximum distance.

⁵See <http://www.swarm-bots.org/hole-avoidance.html> for some movies of these behaviors.

a traction force for the other *s-bots*, which they perceive by means of their traction sensors. At this point, a new coordination phase is triggered, which ends up in a new direction of motion that leads the *swarm-bot* away from the edge. A key role in the functioning of this strategy is played by the motor controlling the rotation of the chassis with respect to the turret of an *s-bot*. In fact, this motor has a stabilizing effect on the rotation of the chassis even if one of the wheels is suspended out of the edge. This gives to the *s-bot* the chance of changing its direction of motion, even when partially suspended. Consequently, the *s-bot* can exert a traction force that can be felt by the other *s-bots*.

4 Generalization

The evolved strategy for hole avoidance is very robust, being able to work in a number of different situations. This is a result of the physical connections among *s-bots* and, above all, of the use of the traction sensors.

As a first experiment, we tested the scalability of the evolved controllers varying the size and the shape of the *swarm-bot*. We observed that the evolved controllers perform well in many different conditions. For example, Figure 3a shows the case of a *swarm-bot* comprising 8 *s-bots* connected in a “star” shape. The *swarm-bot* is placed in a square arena without holes, but with open borders. The *swarm-bot* is still able to avoid to fall out of the arena, notwithstanding the higher inertia of the star formation.

Another interesting feature of the evolved controllers is that they are able to perform collective obstacle avoidance. In fact, when an *s-bot* hits an obstacle, its turret exerts a force on the chassis in a direction opposite to the obstacle. This force is felt as a traction pulling the *s-bot* away from the obstacle. In response to this traction, the *s-bot* rotates its chassis to cancel it, as explained before. Moreover, the rigid connections between *s-bots* transmit the force resulting from the collision to the whole group, triggering a fast change in the direction of movement of the *swarm-bot*. As shown in Figure 3b, the *swarm-bot* is able to avoid both holes and obstacles, represented here by walls surrounding the arena. It is worth noting that the traction sensor works as an omni-directional bumper distributed on the whole body of the *swarm-bot*, allowing collective obstacle avoidance.

Finally, we tested the evolved controllers when the *s-bots* are linked using flexible, rather than rigid, connections. Flexible connections allow the relative motion of connected *s-bots*, and, therefore, the use of this type of connections allows the shape of the *swarm-bot* to change during motion. Because of the flexibility of the connections, traction can be only partially transmitted. Nevertheless, the evolved strategies still work. We performed tests with both a star and a chain formation composed of 8 *s-bots* each. The flexible star formation case is shown in Figure 3c, where the *swarm-bot* was placed in a square arena with four big cylindrical obstacles and no walls on the perimeter. Figure 3c shows that the flexible formation was able to perform coordinated motion, obstacle and hole avoidance, changing shape when it had to go through a narrow passage having an obstacle on the left and the arena border on the right. The flexible formation adapts more easily to the environment, and in some situations can avoid holes more efficiently than a rigid structure. In fact, the *s-bots* do not completely feel the inertia of the *swarm-bot*, because they can move deforming the structure and adapting to the edge of the hole. This fact is even more evident in Figure 3d, where a chain formation was placed in the arena with holes. Here, when the chain reached the edge, it completely deformed without having a single *s-bot* being completely pushed out of the arena.




Figure 3: Generalization properties: (a) Size and shape change. (b) Obstacle avoidance. (c) Obstacle and hole avoidance of a big star formation with flexible connections. (d) Hole avoidance of a big linear formation with flexible connections.

5 Conclusions

We presented a set of experiments for the evolution of hole avoidance behaviors in a group of simulated *s-bots* that are physically connected to form a *swarm-bot*. The solutions found by evolution are simple and in many cases they generalize to different environmental situations. This demonstrates that evolution is able to produce a self-organizing system that relies on simple and general rules, a system that is consequently robust to environmental changes and to the number of *s-bots* involved in the experiment. The evolved strategies strongly rely on the traction forces produced by those *s-bots* that feel the presence of a hazard. Using the information given by the traction sensors, the whole group can change the direction of motion when heading toward a hole.

The traction sensor was found to be a very powerful mean of achieving coordination in the *swarm-bot*. In fact, it allows the exploitation of the complex dynamics arising from the interactions among *s-bots* and between the *s-bots* and the environment. It provides robustness and adaptivity features with respect to environmental or structural changes of the *swarm-bot*. Besides, traction forces are used as a sort of communication of the presence of a hazard. This communication among *s-bots* is neither direct nor explicit, but can be considered as an implicit stigmergic communication, as it takes place through the environment, that is, through the bodies and the physical connections among *s-bots* [4, 5]. Finally, the traction sensor can work also as a distributed bumper for the *swarm-bot*, allowing collective obstacle avoidance.

The hole avoidance task represents the first step toward the solution of more difficult problems. We plan to continue studying problems that belong to the “navigation on rough terrain” family, like passing over a trough or coping with an uneven terrain. Finally, we will face the challenge given by functional self-assembling for all-terrain navigation, that is, we will study the problem of forming or disbanding *swarm-bots* with a shape functional to the environmental conditions and to the task to be performed, in order to maximize the efficiency in the navigation.

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. 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. 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.

References

- [1] G. Baldassarre, S. Nolfi, and D. Parisi. Evolution of collective behavior in a team of physically linked robots. In R. Gunther, A. Guillot, and J.-A. Meyer, editors, *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics (EvoWorkshops2003: EvoROB)*, pages 581–592. Springer-Verlag, Berlin, Germany, 2003.
- [2] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- [3] A. Castano, W. Shen, and P. Will. CONRO: Towards deployable robots with inter-robot metamorphic capabilities. *Autonomous Robots*, 8:309–324, 2000.
- [4] M. Dorigo, E. Bonabeau, and G. Theraulaz. Ant algorithms and stigmergy. *Future Generation Computer Systems*, 16(8):851–871, 2000.
- [5] P. P. Grassé. La reconstruction du nid et les coordinations interindividuelles 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.
- [6] A. Lioni, C. Sauwens, G. Theraulaz, and J.-L. Deneubourg. Chain formation in *Æcophilla longinoda*. *Journal of Insect Behaviour*, 15:679–696, 2001.
- [7] L. Matthies, E. Gat, R. Harrison, B. Wilcox, R. Volpe, and T. Litwin. Mars microrover navigation: Performance evaluation and enhancement. *Autonomous Robots, Special Issue on Autonomous Vehicles for Planetary Exploration*, 2(4):291–312, 1995.
- [8] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
- [9] M. Quinn, L. Smith, G. Mayley, and P. Husband. Evolving teamwork and role allocation with real robots. In R. K. Standish, M. A. Bedau, and H. A. Abbass, editors, *Proceedings of the 8th International Conference on Artificial Life*, pages 302–311. MIT Press, Cambridge, MA, 2002.
- [10] K. Støy, W.-M. Shen, and P. Will. Global locomotion from local interaction in self-reconfigurable robots. In W. M. Shen, C. Torras, and H. Yuasa, editors, *Proceedings of the 7th International Conference on Intelligent Autonomous Systems (IAS-7)*, pages 309–316. IOS Press, Amsterdam, The Netherlands, 2002.
- [11] V. Trianni. Evolution of coordinated motion behaviors in a group of self-assembled robots. Technical Report TR/IRIDIA/2003-25, IRIDIA, Université Libre de Bruxelles, Belgium, May 2003. DEA Thesis.
- [12] V. Trianni, R. Gross, T. H. Labella, E. Şahin, and M. Dorigo. Evolving aggregation behaviors in a swarm of robots. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, and J. Ziegler, editors, *Advances in Artificial Life - Proceedings of the 7th European Conference on Artificial Life (ECAL)*, volume 2801 of *Lecture Notes in Artificial Intelligence*, pages 865–874. Springer-Verlag, Berlin, Germany, 2003.
- [13] C. R. Ward, F. Gobet, and G. Kendall. Evolving collective behavior in an artificial ecology. *Artificial Life*, 7(2):191–209, 2001.
- [14] M. Yim, D. G. Duff, and K. D. Roufas. 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, 2000.