

Evolution of self-organised path formation in a swarm of robots

Valerio Sperati, Vito Trianni, and Stefano Nolfi

Istituto di Scienze e Tecnologie della Cognizione,
Consiglio Nazionale delle Ricerche, Rome, Italy
{valerio.sperati,vito.trianni,stefano.nolfi}@istc.cnr.it

Abstract. We present a set of experiments in which a robotic swarm manages to collectively explore the environment, forming a path to navigate between two target areas, which are too distant to be perceived by an agent at the same time. Robots within the path continuously move back and forth between the two locations, exploiting visual interactions with their neighbours. The global group behaviour is obtained through an evolutionary process and presents emergent properties like robustness, path optimisation and scalability, which recall ants trail formation.

1 Introduction

Exploration and navigation in unknown environments represent basic activities for most animal species, and efficient strategies can make the difference between death and survival. For this reason, Nature presents a wide range of possibilities, each particularly adapted to the task to be accomplished and to the sensory-motor and cognitive abilities of the species under observation. In primates, as well as in other animals, navigation abilities are usually linked to mental representations of the environment, referred to as “cognitive maps”. For instance, it has been found that specific neurons of the rodents hippocampus (called “place cells”) have a high firing rate in correspondence of specific locations in the environment [15]. Neural representations seem to characterise also the behaviour of insects. A map-like organisation of spatial memory has been proposed for honeybees, which are able to retrieve the navigation path on the basis of learned landmarks around the hive [12]. A similar strategy is employed also by the desert ants of the genus *Cataglyphis*, which however couple the landmark-base strategy with their skylight (polarization) compass and path integrator (ants integrate over time the path covered through a sort of vector summation) that allow them to return to the nest following a straight line [21]. Ant species that forage in groups rely on a collective strategy for exploration and navigation, exploiting the well known mechanism of pheromone trail formation: when moving from a foraging patch to the nest, ants lay a blend of pheromones that can be exploited by other ants to reach the same patch. Thanks to this strategy, ants can efficiently navigate in the environment and optimise the path between nest and food [7, 2].

In Robotics too, much attention has been paid to the navigation and exploration problems, and several different strategies have been proposed. Map-based navigation exploits probabilistic approaches to solve the so called simultaneous localisation and mapping (SLAM) problem [18, 1], as well as biologically inspired ones [4, 9, 6]. Similarly, landmark-based navigation and path integration have been exploited, often with a close look at biology [23, 8, 10, 20]. For what concerns collective strategies inspired to the ants trails, however, research has to confront with the complex problem of finding an alternative to the pheromones, given that chemical substances are difficult to exploit in a robotic setup. Instead of chemicals, the most common approach is to rely on communication and message passing among robots, which therefore simulate pheromone attributes on a communication network [17]. Other approaches exploit the robots themselves as markers for the trail: marker robots remain static and signal a path between two locations in the environment, while explorer robots exploit this path to efficiently navigate [3, 22, 16, 14].

In this paper, we study the abilities of a robotic swarm to explore an environment and coordinately navigate between two target areas by exploiting a simple form of visual communication. We start from two basic assumptions about the usage of communication signals. On the one hand, communication can be exploited to signal the position of a target zone to the other robots, therefore facilitating the exploration task. For this purpose, robots are provided with a red LED positioned on their back, which has the same colour of the target area and therefore elicits the same effect of the visual perception of a target area [11]. On the other hand, communication can be used to coordinate the movements of the robots in the environment, therefore supporting the navigation task. For this purpose, robots are provided with a blue LED on their front, which can signal their position and heading to other robots.

By exploiting evolutionary robotics techniques [13, 5], we study the emergence of behavioural and communication strategies. Evolutionary robotics is particularly useful to synthesise self-organising collective behaviours, characterised by properties such as robustness, flexibility and scalability [19]. We analyse the system trying to identify the relevant properties of the evolved behaviours for what concerns both the individual rules followed by the robots, and the communication signals exploited. Despite not explicitly required by the fitness function, the evolved behaviour presents features that are similar to the trail formation in ants: robots form a trail between the target areas and robustly maintain it, also optimising the shape towards the shortest path (see Sec. 3). Moreover, we analyse the generalisation and scalability property of the collective behaviour, by testing it in different environmental conditions and with larger swarms (see Sec. 4). Discussions are reported in Sec. 5.

2 Experimental setup

The experimental scenario involves a swarm of wheeled robots, whose behaviours have been evolved for the ability to navigate back and forth as quickly as possible

between two target areas, located within an arena surrounded by walls (Fig. 1). Since target areas can be perceived only from a short distance, the robots should be able to find them by exploring the environment. Moreover, in order to quickly navigate from one target area to the other without relying on time-consuming exploratory behaviours, the robots should be able to preserve, in some way, some information concerning the location of previously visited areas. Since the robots are rewarded on the basis of the efficiency with which each individual is able to accomplish the task (see eq. (3)), the evolutionary process might potentially lead to the development of non-cooperative solutions in which the fact to be part of a swarm does not provide any advantage. However, as we will see, the evolutionary process rather leads to strategies in which the robots coordinate and cooperate to find the target areas. In particular, the problem of preserving a trace of the position of previously visited areas is solved by generating and maintaining a dynamic path that connects the two target areas, which allows the individuals of the swarm to efficiently navigate between them. In this section, we detail the experimental setup and the evolutionary algorithm used.

2.1 The robots and the environment

Ten simulated robots are placed in a rectangular arena (height $H = 250$ cm; width variable within the interval $W \in [250, 290]$ cm). The target areas consist in two circles painted in grey (diameter $d = 32$ cm). A target area can be perceived from distance by the robot thanks to a red LED placed over its centre. This red LED is indistinguishable from the one provided to the robots (see below).

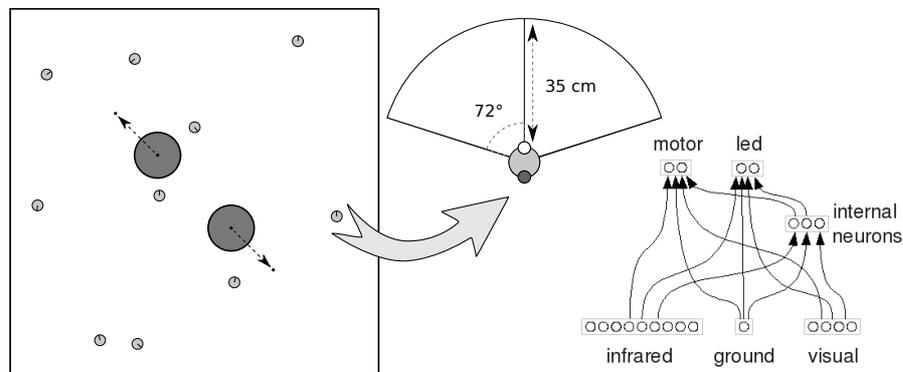


Fig. 1. The experimental setup. On the left, a snapshot of the simulated environment is shown. Ten robots randomly positioned in the environment are represented as small circles. The grey disks represent the circular target areas with a red LED in the centre. The distance between the area centres is $D = 70$ cm, while the arrows indicate which is their maximum displacement, when $D = 150$ cm. In the centre, we show a schematic representation of the robot's vision sensors, indicating the sectors and the perceptual range. The blue and red LED position is indicated as a white and a gray dot, respectively. On the right, the architecture of the robots' neural controller is shown.

The two areas are positioned symmetrically with respect to the centre of the arena at a fixed distance D (see Fig. 1). Each robot is provided with two motors that control two wheels, providing a differential drive motion (maximum speed: $v_{max} = 8.2$ cm/s). Moreover each robot is provided with a blue LED on the front and a red LED on the rear of its body. Both can be switched on and off by the robot controller. Additionally, a robot is provided with: (i) 8 infrared sensors uniformly distributed around the robot body, used to detect obstacles or other robots up to a distance of about 2.5 cm; (ii) 1 ground sensor located under the front of the robot, used to detect whether the robot is placed over a target area or not; and (iii) 4 vision sensors, used to detect the presence of red or blue LEDs (2 sensors for each colour). The vision sensors return a binary value about the presence or absence of LEDs in two 72° sectors of the image that cover the front-left and front-right area of the robot. Both red and blue LEDs can be detected up to a distance of 35 cm (see Fig. 1).

2.2 The controller and the evolutionary algorithm

Each robot is controlled by a feed-forward neural network (see Fig. 1) with 13 sensory neurons (8 infrared, 1 ground, 4 vision sensors), 3 internal leaky integrators neurons and 4 motor neurons (2 wheels, 1 blue LED and 1 red LED). Connections weights, biases and time constants of the leaky integrators are genetically encoded parameters subject to artificial evolution [13, 5]. The free parameters of the robot’s neural controller are encoded in a binary genotype, using 8 bits for each real number. The connection weights and biases can vary in the range $[-5, 5]$, while the time constants vary in $[0, 1]$. Evolution works on a population of 100 randomly generated genotypes. After evaluation of the fitness, the 20 best genotypes survive in the next generation (elitism), and reproduce by generating four copies of their genes with a 3% mutation probability of flipping each bit. The evolutionary process lasts 500 generations.

In order to evaluate the fitness, a genotype is translated into N identical neural controllers which are downloaded onto N identical robots (i.e., the group is homogeneous). Each group of robots is tested for 15 trials, each lasting 6000 time-steps (one time-step corresponds to 100 ms). Regarding the fitness computation, the trial is split into two periods, T_{add} and T_{eff} : the latter lasts 5400 time-steps and is the actual period during which the fitness is estimated. The former lasts 600 time-steps, and is the time dedicated to the exploration of the environment and to the initial coordination of the robots. At the beginning of each trial the positions and the orientations of the robots are randomly initialised, while the target areas are positioned systematically choosing a value in $D \in D_{set} = \{70, 90, 110, 130, 150\}$ cm (see Fig. 1). The performance F of the group during a trial is obtained evaluating how often and how quickly each robot in the group moves from one area to the other. For this purpose, each robot i cumulates a reward f every time it enters in a target area different from the one previously visited. This reward is computed according to the energy e_i saved in

moving from one target area to the other.

$$f_i(t) = f_i(t-1) + \begin{cases} e_i(t) & \text{if robot enters a new target area} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $f_i(t)$ is the reward cumulated at time t , and $e_i(t)$ is the energy saved. Equation (1) states that when a robot enters in a target area different from the one previously visited, it stores the current energy load $e_i(t)$ as reward. At the same time, the energy level is reset to a quantity proportional to the distance between the target areas. Otherwise, the robot consumes its energy while moving, proportionally to its speed:

$$e_i(t) = \begin{cases} 1 + E_{ab} & \text{if robot enters a new target area} \\ e_i(t-1) - \delta_i(t) & \text{otherwise} \end{cases} \quad (2)$$

where E_{ab} is the energy that a robot would consume to move in a straight line between the two target areas, and $\delta_i(t)$ is the energy consumed in a single time-step, proportional to the wheels speed¹. With an optimal behaviour (moving straight between the two areas), a robot would store a quantity $e_i(t) = 1$ each time it enters a new target area, independently from the distance between the two. The performance of the robot is computed at the end T of the trial, and is normalised according to the maximum performance that can be achieved by a robot behaving optimally:

$$\bar{f}_i = f_i(T)/f_{max}, \quad f_{max} = v_{max} \cdot T_{eff}/D_{ab} \quad (3)$$

where $D_{ab} = D - d$ is the minimum distance that must be covered between two target areas. Finally, the fitness of the group F in a trial is computed as the average across the group:

$$F = \sum_{i=1}^N \bar{f}_i \quad (4)$$

The final fitness of the genotype is the average of F over 15 different trials.

3 Obtained results and behavioural analysis

The evolutionary process has been replicated 10 times—hereafter, evolutionary runs—starting from randomly generated populations. At the end of each evolutionary run, we selected a single genotype to be analysed thoroughly. To do so, we evaluated the performance of the best genotype of the last 100 generations, and we selected the one showing the highest average over 500 trials as the representative of each evolutionary run. The results are shown in Fig. 2. A qualitative analysis of the evolved behaviours reveals that 6 evolutionary runs out of 10 result in a good collective exploration and navigation behaviour (runs number

¹ This is in $[0.0, 0.0025]$, which means a robot wastes at most 1 unit of energy in 400 time-steps, if moving at maximum speed.

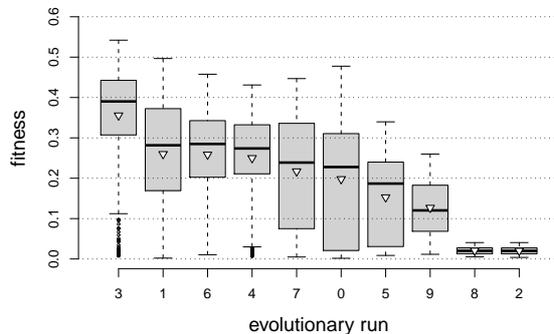


Fig. 2. Performance of the best evolved individual for each evolutionary run. Each boxplot corresponds to the performance obtained in 500 trials. Boxes represent the inter-quartile range of the data, while the horizontal lines inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. Circles mark the outliers. The symbol ∇ indicates the average performance.

3, 1, 6, 4, 7 and 0). Two runs (number 5 and 9) produced sub-optimal strategies, and two others (number 8 and 2) resulted in unsatisfactory behaviours both at the individual and collective level. Given that the successful runs produced qualitatively similar behaviours, in the following we describe the one of the best genotype, corresponding to the evolutionary run number 3 (see Fig. 2).²

The sequence displayed in Fig. 3(a) shows how a typical trial unfolds in time³. Initially, robots move independently and explore the environment. In doing so, they signal their position and heading to other robots keeping the front blue LED switched on, while the red LED is used only in certain conditions (see below). The visual interactions mediated by these signals allow the group to converge to a coherent motion between the two target areas. Eventually, the robots form two rows moving from one target to the other in opposite directions (see the last frame in Fig. 3(a)). We refer to this structured spatio-temporal pattern formed by the robots as *dynamic chain*. The term *dynamic* well illustrates two interesting features of this structure. Firstly, each robot in the chain is not static, but moves continuously along it, swinging between the target areas as requested by the fitness function. Secondly, the chain connecting the two targets adapts its shape according to the current distance D between areas: it adapts the chain direction by choosing the shortest path between the two areas, and adapts the inter-robot distance to fit all robots in the chain (see Fig. 3(b)).

This collective behaviour is the result of simple rules followed by each individual robot and encoded in the neural controller. When a robot has no objects

² In this case the best genotype belongs to the 475th generation.

³ See <http://lalar.istc.cnr.it/esm/sperati-et-al-ANTS2010/> for videos

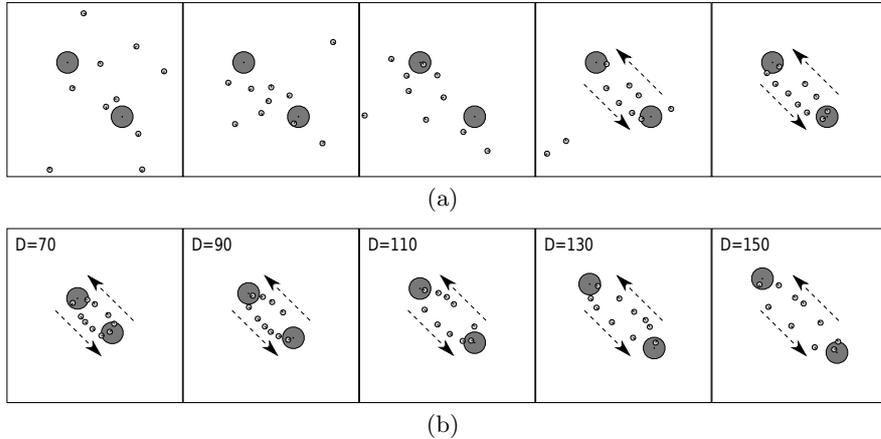


Fig. 3. (a) Temporal sequence recorded in a generic successful trial ($D = 110$ cm), showing the formation of the *dynamic chain*. (b) From left to right, each snapshot displays the final configuration achieved by the swarm, at the end of 5 different standard trials where the target areas are positioned according to the distance values in D_{set} .

in its perceptual field, it moves counterclockwise in large circles, the front blue LED always switched on. When a target area is in sight, a robot approaches it in a straight line and makes a u-turn when it reaches the grey circle. When two robots encounter, they avoid each other by always dodging to the right, exploiting the blue visual signal emitted by the robots. This constitutes the basic mechanism for the formation of the dynamic chain: in fact, a robot does not necessarily follow the robot in front moving in the same direction, but rather keeps on its left the robots coming in the opposite direction. It is clear that a minimal number of robots is necessary to support this behaviour, because a dynamic chain is stable as long as there are robots moving in opposite directions. In order to aggregate all robots in the chain formation, the red signal is exploited. In fact, the red signal mimics the colour of the target area, and in general induces an approaching behaviour. The red LED is switched on in two conditions. On the one hand, a robot flashes while moving towards a target area. In this case, the signal allows nearby robots to react by approaching the target area themselves, even though they do not directly perceive it. On the other hand, robots signal while avoiding each other. In this case, the function of the signal appears linked also to the enhancement of the stability of the dynamic chain. Although the identification of the exact roles played by the red and blue signals needs further analysis that we plan to carry on in future research, these observations clearly indicate that communication plays a crucial role for the formation of the dynamic chain and more generally for the ability of the robots to coordinate and cooperate.

This brief qualitative analysis suggests the following considerations. First of all, the formation of the chain is the outcome of a self-organising process that

results solely from the robot-robot interactions. We observed that the dynamic chain forms rather abruptly out of a disordered group motion. We believe that this may correspond to a phase transition that depends mainly on the density of robots between the two target areas. In other words, when enough robots are attracted in this area thanks to the visual communication, the chain forms. A second important remark concerns the function of the dynamic chain, which is exploited by the robots to maintain the right heading toward the target areas when they are not in sight. From this point of view, the group behaviour compensates for the limited sensory range of the robots, which collectively discover and preserve information concerning the direction of the two areas, thanks to the exploitation of the communication signals.

To better evaluate the performance of the group, we tested the collective behaviour systematically varying the distance between the target areas $D \in D_{set}$. The obtained results are presented in Fig. 4(a). Here the performance of the group is compared to the performance of a single individual evolved in a control experiment.⁴ First of all, we notice that the group always outperforms the single individual. This confirms that a coordinated behaviour has been evolved, which goes beyond the capability of the individual robot. Looking at the performance of the group, it is possible to notice that the behaviour seems adapted mostly for an intermediate distance, in which it scores the highest average performance. With larger distances, the performance across different trials is more changing. This suggests that the group may be able to coordinate in some cases, and in others is not. We ascribe this variability to the limited duration of the trial, hypothesising that in some cases robots do not have enough time to coordinate and form a dynamic chain. To test this hypothesis, we performed an identical test, but now increasing the duration of the initial coordination period ($T_{add} = 18600, T_{eff} = 5400$ time-steps). The results plotted in Fig. 4(b) confirm that for all distances the group attains a good score, which is also very stable across different trials. Moreover, the results indicate that the group behaves better for large distances. In fact, with short distances, the dynamic chain is overcrowded and robots interfere with each other, therefore obtaining a lower performance. In this conditions, a smaller group behaves better (data not shown). Finally we note how, when $D = 150$, the performance of the individual continue to be very low, despite the extended time. This means that a robot alone is not capable to face all the D values, while the swarm is.

4 Generalisation abilities

In the previous section, we have described the features of the evolved behaviour, and observed how the system always converges to a dynamic path formation if enough time is granted for coordination. In this section, we test the ability of the

⁴ We explicitly evolved a neural controller for a single robot confronted with the same task, with identical evolutionary conditions. We performed 30 evolutionary runs, each one lasting 1000 generations, and selected the best genotype with the same procedure described above.

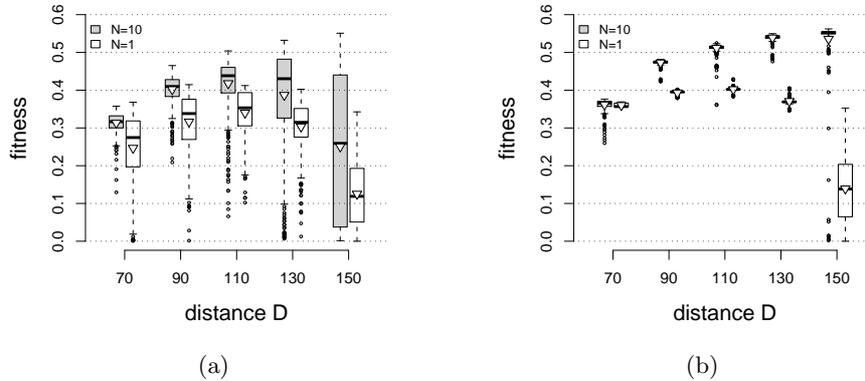


Fig. 4. (a) Performance test with varying distance $D \in D_{set}$. Performance is computed for groups of $N = 10$ robots, and for a single robot evolved for the same task in a control experiment ($N = 1$). (b) The same test performed with a longer initial period $T_{add} = 18600$, during which performance is not computed.

system to generalise to different conditions never met during the evolutionary optimisation. In particular, we want to understand whether robots are able to form a path with larger distances and with larger groups. We test the performance of the group in 16 new conditions, obtained coupling 4 groups ($N \in \{20, 30, 40, 50\}$) with 4 distances ($D \in \{200, 250, 300, 350\}$ cm). These tests have been performed in a larger arena (fixed height $H = 350$ cm, variable width $W \in [350, 390]$ cm), and in longer trials ($T_{add} = 36600$). The quantitative results are shown in Fig. 5. We can immediately notice that, when the number of robots is sufficiently large, the swarm is successful also when the distance between the two target areas is much wider compared to the conditions experienced during the evolutionary process. With distance $D = 200$ cm, groups with 20 robots perform best, while larger groups are less efficient. Groups of 30 robots have a fairly good performance, which however presents a large variability. With distance $D = 250$ cm it is possible to notice a similar pattern. However, this time $N = 30$ is the optimal group size. Finally, for $D = 300$ cm and $D = 350$ cm, the size $N = 40$ performs best, with a larger variability in the latter case, in which also $N = 50$ presents a fairly good performance in many trials.

This analysis confirms our expectations: the larger the distance between the target areas, the larger the number of robots required to form a stable chain. In fact, as mentioned above, the dynamic chain is maintained as long as there are constantly robots moving in opposite directions uniformly distributed along the path, which implies larger groups for larger distances. The analysis also confirms that a minimum number of robots is necessary to form a path over a certain distance. Similarly, large groups suffer overcrowding when the distance D is too short, as there is no space available to distribute all the robots along

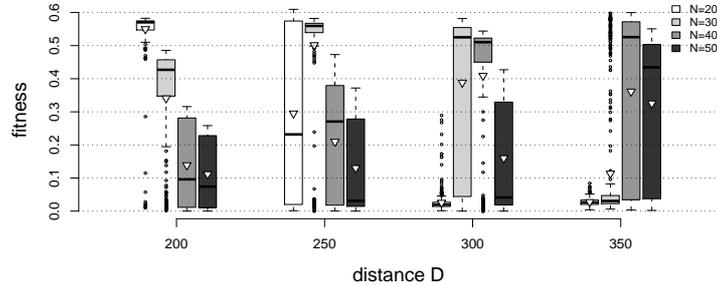


Fig. 5. Generalisation ability for groups of increasing size N and for increasing distance D . Each boxplot corresponds to the performance obtained in 500 trials ($T_{eff} = 5400$, $T_{add} = 36600$).

the path. However, the dynamic chain can adapt to a wide range of distances. For instance, groups of 30 robots present good performance up to $D = 300$ cm, and only with larger distances the performance systematically drops.

5 Discussion and conclusions

In this paper we reported a series of experiments in which the behaviour of a swarm of robots has been evolved for the ability to navigate back and forth between two target areas which can only be perceived locally. The analysis of the obtained results indicates that the robots solve the problem by exploring the environment and by forming a dynamic chain constituted by two rows of robots moving from one target to the other in opposite directions. The dynamic chain emerges rather abruptly from the robot-robot interactions mediated by light signals, and afterwards adapts by converging toward a configuration that corresponds to the shortest path and that is characterised by a rather uniform distribution of distances between the robots. The analysis of the evolved behaviour demonstrates how it generalises to situations in which the distance between the two target areas is much wider compared to the conditions experienced during the evolutionary process provided that the number of robots forming the swarm is sufficiently large. Similarly to pheromone trails in ants, dynamic chains allow the swarm to efficiently navigate between the two target areas. Indeed, in both cases the stability of the structure is a result of a sustained flux of individuals that support it, in one case by pheromone laying, in the other by coloured signals. Another common feature is the ability to identify the shortest path between two locations. Even though we only performed tests in an obstacle-free arena, we observed that dynamic chains may initially form in curved paths (especially with large groups), which slowly straighten until the shortest route is taken. In future work, we plan to analyse in more detail the evolved behaviour, in order to

better understand the properties of the dynamic chain and its relationship with similar behaviours observed in Nature.

In future work, we plan to test the evolved behaviour in hardware exploiting the foot-bot robotic platform developed within the European project *Swarmanoid* (grant IST-022888, see <http://www.swarmanoid.org>). Within this project, we also aim at testing coordinated behaviours among groups of heterogeneous robots. In particular, we plan to exploit the eye-bot robotic platform to work as 'smart' target areas. The eye-bot is in fact an aerial robot with the ability to attach to the ceiling, from which it can monitor the environment and detect relevant areas. Foot-bots can perceive an eye-bot only when they are approximately underneath it, thanks to a camera pointing upward. As a consequence, eye-bots could be exploited by the foot-bots as target areas, and dynamic chains can be formed between them, creating for instance a delivery line from a target to a goal location, which is known by the eye-bots thanks to their privileged viewpoint. Moreover, eye-bots can move in the environment and exploit the robustness of the dynamic chain to purposely modify its length or its shape. Finally, we will investigate how to form more complex dynamic chains that can trace the shortest path between more than two eye-bots or target areas.

Acknowledgements: The authors wish to thank Onofrio Gigliotta, Tomassino Ferrauto and Gianluca Massera for the fruitful discussions and their help. This work was supported by the Swarmanoid project, funded by Future and Emerging Technologies programme (IST-FET), of the European Commission, under grant IST-022888.

References

1. Bailey, T., Durrant-Whyte, H.: Simultaneous localization and mapping: part II. *IEEE Robotics & Automation Magazine* 13(3), 108–117 (2006)
2. Detrain, C., Denebourg, J.: Collective decision and foraging patterns in ants and honeybees. *Advances in Insect Physiology* 35, 123–173 (2009)
3. Drogoul, A., Ferber, J.: From Tom Thumb to the Dockers: some experiments with foraging robots. In: *From Animals to Animats 2, Second International Conference on Simulation of Adaptive Behavior (SAB-92)*. pp. 451–459. MIT Press Cambridge, MA (1993)
4. Filliat, D., Meyer, J.: Map-based navigation in mobile robots - I. A review of localization strategies. *Journal of Cognitive Systems Research* 4, 243–282 (2003)
5. Floreano, D., Husband, P., Nolfi, S.: Evolutionary robotics. In: Siciliano, B., Oussama, K. (eds.) *Handbook of Robotics*, pp. 1423–51. Berlin: Springer Verlag (2008)
6. Gigliotta, O., Nolfi, S.: On the coupling between agent internal and agent/environmental dynamics: Development of spatial representations in evolving autonomous robots. *Adaptive Behavior* 16, 148–165 (2008)
7. Goss, A., Aron, S., Denebourg, J., Pasteels, J.: Self-organized shortcuts in the argentine ant. *Naturwissenschaften* 76(12), 579–581 (1989)

8. Gutiérrez, Á., Campo, A., Santos, F.C., Pinciroli, C., Dorigo, M.: Social odometry in populations of autonomous robots. In: Proceedings of the Sixth International Conference on Ant Colony Optimization and Swarm Intelligence (ANTS-2008). Lecture Notes in Computer Science, vol. 5217, pp. 371–378. Springer Verlag, Berlin, Germany (2008)
9. Hafner, V.V.: Cognitive maps in rats and robots. *Adaptive Behavior* 13, 87–96 (2005)
10. Lambrinos, D., Kobayashi, H., Pfeifer, R., Maris, M., Labhart, T., Wehner, R.: An autonomous agent navigating with a polarized light compass. *Adaptive Behavior* 6(1), 131–161 (1997)
11. Maynard-Smith, J., Harper, D.G.: *Animal Signals*. Oxford University Press, Oxford, UK (2003)
12. Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., Bundrock, G., Hülse, S., Plümpe, T., Schaupp, F., Schüttler, E., Stach, S., Stindt, J., Stollhoff, N., Watzl, S.: Honey bees navigate according to a map-like spatial memory. *Proceedings of the National Academy of Sciences of the United States of America* 102(8), 3040–3045 (2005)
13. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Mit Press/Bradford Books, Cambridge, MA (2000)
14. Nouyan, S., Campo, A., Dorigo, M.: Path formation in a robot swarm. *Swarm Intelligence* 2(1), 1–23 (2007)
15. O’Keefe, J., Nadel, L.: *The Hippocampus as a Cognitive Map*. Originally published by Oxford University Press, Oxford, UK (1978)
16. Ostergaard, E., Sukhatme, G., Matarić, M.: Emergent bucket brigading: a simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In: Proceedings of the Fifth International Conference on Autonomous Agents. pp. 2219–2223 (2001)
17. Payton, D., Daily, M., Estkowski, R., Howard, M., Lee, C.: Pheromone robotics. *Autonomous Robots* 11(3), 319–324 (2001)
18. Thrun, S.: Robotic mapping: a survey. In: Gerhard Lakemeyer, G., Nebel, B. (eds.) *Exploring artificial intelligence in the new millennium*, pp. 1–35. Morgan Kaufmann Publishers Inc. San Francisco, CA (2003)
19. Trianni, V., Nolfi, S., Dorigo, M.: Evolution, self-organisation and swarm robotics. In: Blum, C., Merkle, D. (eds.) *Swarm Intelligence. Introduction and Applications*, pp. 163–192. Natural Computing Series, Springer Verlag, Berlin, Germany (2008)
20. Vickerstaff, R.J., Di Paolo, E.A.: Evolving neural models of path integration. *Journal of Experimental Biology* 208, 3349–3366 (2005)
21. Wehner, R.: Desert ant navigation: how miniature brains solve complex tasks. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 189(8), 579–588 (2003)
22. Wergler, B., Matarić, M.: Robotic food chains: Externalization of state and program for minimal-agent foraging. In: *From Animals to Animats 4, Fourth International Conference on Simulation of Adaptive Behavior (SAB-96)*. pp. 625–634. MIT Press Cambridge, MA (1996)
23. Zeil, J., Boeddeker, N., Stürzl, W.: Visual homing in insects and robots. In: Floreano, D., Zufferey, J.C., Srinivasan, M., Ellington, C. (eds.) *Flying Insects and Robots*, pp. 87–100. Springer Verlag, Berlin, Germany (2009)