



# Evolving Swarm Formations for Odour Source Localisation

João Macedo<sup>1,2</sup>(✉), Lino Marques<sup>1</sup>, and Ernesto Costa<sup>2</sup>

<sup>1</sup> Institute of Systems and Robotics, University of Coimbra, 3030-290 Coimbra, Portugal

{jmacedo,lino}@isr.uc.pt

<sup>2</sup> Centre for Informatics and Systems, University of Coimbra, 3030-290 Coimbra, Portugal  
ernesto@dei.uc.pt

**Abstract.** Odour source localisation is a hard problem with many applications. Over the years, researchers have drawn inspiration from Nature to devise many single-robot approaches. Swarm approaches have been growing in popularity, as they offer redundancy to the loss of agents, flexibility, scalability and enable experimenters to employ simpler robots. Many existing swarm approaches make use of robot formations. In this work, we focus on optimising the shape of a swarm formation for finding and tracking odour plumes. We do so by using a genetic algorithm, thus avoiding the cumbersome trial-and-error process that experimenters typically follow to hand-design the formations. The swarm is guided by a leader, which is controlled by a bio-inspired search strategy using the perceptions of the entire swarm. The results show that the evolved formations of three and five robots consistently outperform a single robot and that the best evolved three robot formation is more successful than the hand-designed swarms of three and five robots. As a result, one could opt by using the evolved three robot formation, minimizing the amount of robots needed. Conversely, in case there is a high risk of loss of robots, the evolved five robot formation could be preferable.

**Keywords:** Odour source localisation · Genetic algorithms · Swarm robotics

## 1 Introduction

The terms *fragrance* and *odour* are typically used to refer respectively to pleasant and unpleasant scents. Nevertheless, both refer to the animals' perception of volatile chemical compounds. Through the sense of olfaction, animals often track odours to find mates or sources of food. In the literature, odour source localisation (OSL) refers to the task of locating the source that is emitting a given odour.

OSL is a hard problem with many applications. Over the years, researchers have drawn inspiration from Nature to devise many robotic approaches, mostly

relying on a single robot. The existing approaches can be categorised as reactive, which often attempt to mimic natural behaviours such as those of the *E. coli* bacteria, the dung beetle or the male silkworm moth [18], and probabilistic approaches (e.g., Infotaxis [19]), which try to estimate the location of the chemical source. Despite the success of single-robot approaches, in Nature, animals cooperate to fulfil tasks that they would otherwise be unable to accomplish or would do so with lower performance. As an example, bees and ants form swarms to forage for food, whereas birds and fish travel in flocks and schools to protect themselves [20]. Biological swarms have been the source of inspiration for various optimisation methods and, more recently, to robotic approaches. Swarm Robotics (SR) rely on highly scalable approaches employing groups of agents with limited capabilities, which would be unable or inefficient at performing the target task. Yet, through their local interactions, complex behaviours emerge enabling them to solve the task as a group [7]. The use of multiple robots provides advantages such as robustness to the loss of agents, flexibility and scalability, which are particularly interesting when working in inhospitable environments. Moreover, swarm approaches enable experiments to employ inexpensive robots, reducing the economic impact of losing some agents. The use of multiple robots also provides the ability to sense the environment in many locations simultaneously, which is particularly helpful to locate chemical sources in environments with intermittent and meandering chemical plumes.

The existing swarm approaches can be divided based on whether the robots move in formation. The approaches that do not enforce formations often employ meta-heuristic algorithms to guide the robots. Various works have proposed Particle Swarm Optimisation (PSO) methods for guiding robotic swarms in the search for chemical sources [6, 15]. Others have proposed evolutionary approaches [13, 14]. Despite the differences inherent to the algorithms used, both types of approaches operate quite similarly. All of them iteratively evolve waypoints for the robots, which are evaluated through the chemical concentration sensed at each location. Thus, the methods attempt to find the location of maximum chemical concentration. In turn, formation-based approaches typically rely on attractive/repulsive forces, resulting in regular shapes [1, 8]. Lochmater et al. [9] proposed the crosswind formation method, which uses a minimum of two robots moving in a crosswind line formation to perform the plume tracking stage of OSL. The robots share their wind and odour perceptions, along with relative position with each other and attempt to keep the formation centred around the chemical plume while moving to the source. Their wind-tunnel experiments showed that a formation of three robots was sufficient for tracking the chemical plume. Marjovi and Marques [12], analytically studied the optimal swarm formation for finding chemical plumes, concluding that, under a set of assumptions, the maximum probability of detection is attained by using a line formation, oriented diagonally between the crosswind and upwind directions.

Most formation-based approaches rely on experimenters to carefully hand-design the individual behaviour of each robot, often leading to cumbersome trial-and-error processes [4]. Evolutionary Algorithms (EA) are a family of stochas-


tic search heuristics loosely inspired by the principles of evolution by natural selection and Mendel’s genetics. EAs have been successfully applied to different problems from various domains [5] which either do not have analytical solutions or where such solutions would take too much effort to be found. Their use for evolving robotic morphologies [16] and controllers [2, 17] yielded the research field of Evolutionary Robotics (ER). Genetic Algorithms (GA) are a family of evolutionary algorithms that iteratively improve a population of solutions for a given problem. GAs have already been used to evolve waypoints for a team of robots performing OSL [13]. In that work, no mutation operator was used, as the assignment of the locations to the robots and their navigation was considered to introduce enough randomness into the process. This method was able to perform well in a simulation environment without wind, being the odour dispersion dictated only by molecular diffusion. Later, [14] improved on it, introducing a directed mutation operator that biased the search towards upwind or crosswind depending on whether odour was sensed.

In this work, we focus on optimising the shape of a swarm formation for finding and tracking odour plumes. We do so by using a genetic algorithm, thus avoiding the cumbersome trial-and-error process that experimenters typically follow. The swarm is guided by a dynamically selected leader, which in turn is controlled by a bio-inspired strategy using the perceptions of the entire swarm.

## 2 Evolution of Swarm Formations


This paper proposes to locate odour sources with a robotic swarm that moves in formation to search for a chemical source. The swarm is guided by a leader, which is selected at the beginning of the trial as the one with the intended position closest to the formation centre. This approach reduces the possible behaviour instabilities caused by frequent leader changes while still being flexible, as if the leader becomes inoperative, the next closest robot may assume that role. The leader is controlled by a hand-designed reactive controller. In turn, the remaining robots are attracted to their defined positions in the formation, contributing with environmental measurements to the leader’s controller.

Contrarily to the existing works, which use regular shaped hand-designed formations, this work proposes to evolve the swarm formation with a genetic algorithm. Considering a swarm of  $N$  robots, its formation is represented by an array of  $2N$  real numbers, with each pair  $D_{r,i}, \theta_{r,i}$  being the distance and angle of robot  $i$  to the formation’s virtual centre. As an example, consider the genotype of Fig. 1, which corresponds to a formation of 5 robots.



**Fig. 1.** Example genotype for a formation of 5 robots.

The corresponding phenotype is depicted on Fig. 2, with the 5 robots (r1 to r5) shown as blue circles surrounding the virtual centre (black circle) and the large dotted circle representing the maximum admissible radius for the formation.



**Fig. 2.** Example formation of 5 robots, defined by the genotype of Fig. 1. The black circle represents the formation's virtual centre, whereas the blue circles represent the robots (r1 to r5).

The GA evolves a population of formations for a number of generations, outputting the best one found. The candidate solutions are evaluated by having the swarm perform an OSL trial, and its quality is measured by the fitness function designed in [11], which assesses how close the robots get to the source, the duration of the search and the ratio of time spent in the plume:

$$F = \sum_{i=1}^N \left( \beta \frac{d_i}{D} + \eta \frac{t}{T} + \zeta \frac{p_i}{t} \right) \quad (1)$$

where  $N$  is the number of robots in the swarm,  $d_i$  is the final distance of robot  $i$  to the odour source,  $D$  is the maximum distance to the odour source in this arena,  $t$  is the duration of the search,  $T$  is the maximum evaluation time and  $p_i$  is the time robot  $i$  spent without sensing odour.  $\beta$  and  $\eta$  and  $\zeta$  are weight coefficients and their values are all set to 1. The goal of the GA is to minimise this function, i.e., the lower its value, the better.

### 3 Experimental Setup

#### 3.1 Baseline Formations

In order to assess the performance of the evolved swarms, five baseline formations were devised: a single robot formation S1, two line formations, L3 and L5, each respectively having 3 and 5 robots spaced 2.5 m from each other; and two circular formations, C3 and C5, each with respectively 3 and 5 robots spaced equally at 5 m from the centre. Figure 3 presents these formations in the order that they were described.

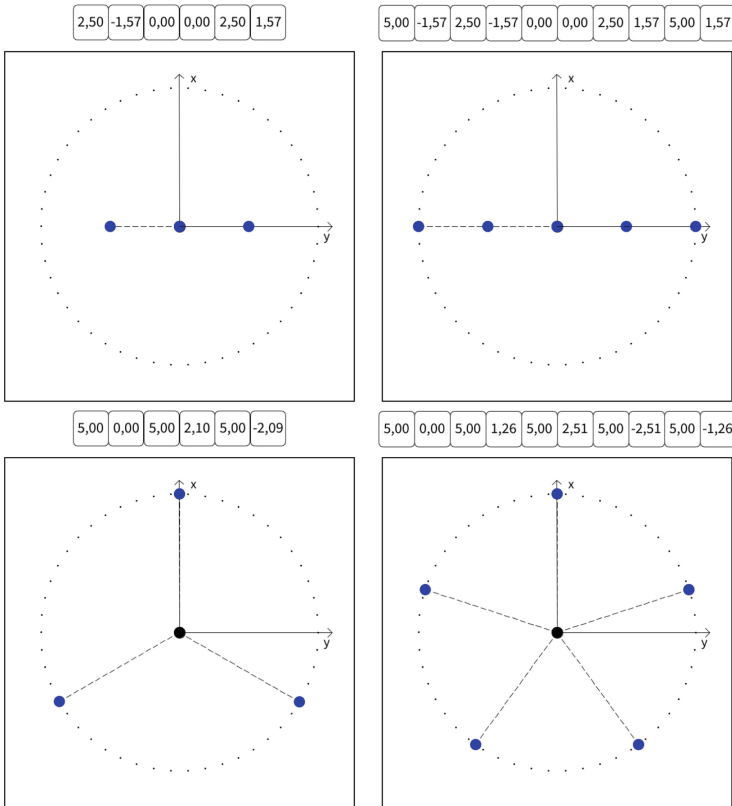


Fig. 3. Baseline formations and their genotypes.

#### 3.2 Formation Leader Controller

The formation leader uses the perceptions of the robots in the swarm to select the actions to perform. In this work, the selection is made based on a simple reactive

controller, composed by symbols that have previously been used with Genetic Programming [11]. This strategy is inspired by natural behaviours [3] moving crosswind to locate the chemical plume and upwind to track it to its source. It starts by evaluating whether odour has been sensed in the last  $t$  seconds. If so, it moves  $l_1$  meters upwind. Otherwise, it moves  $l_2$  meters crosswind, halting as soon as odour is sensed. In this work,  $t$ ,  $l_1$  and  $l_2$  are respectively set to 10s, 0.5 m and 0.5 m. Algorithm 1 presents the pseudocode for this controller.

---

**Algorithm 1.** Controller for the formation leader

---

```

if  $NPL(t)$  then
  moveDirection( $U, l_1 m, C$ )
else
  moveDirection( $X, l_2 m, SO()$ )
end if

```

---


### 3.3 Genetic Algorithm

**Table 1.** Parameters of the genetic algorithm

Parameter	Value
Number of generations	100
Size of the population	100
Size of the elite	3
Crossover rate	0.7
Mutation rate	0.003
Size of the tournament	2

The overall parameters of the GA used for evolving the formations are presented on Table 1. For a swarm of  $N$  robots, the individuals are represented by a vector of  $2N$  real numbers. The relative position of each robot in the formation is represented by a distance  $D_r$  and angle  $\theta_r$  to the virtual centre of the formation. In this work,  $D_r$  and  $\theta_r$  may respectively vary in the intervals  $[0, 5]$  m and  $[-\pi, \pi]$  rad. In order to maximise the diversity of the initial individuals, the initial population is created using Latin Hypercubes [5]. Moreover, to increase the local and global search abilities of the GA, on each generation, 10 random and 10 elitist immigrants are injected into the population. GAs work by selecting and recombining a set of individuals. In this work, tournament selection is used to choose the mates, which are recombined through arithmetic crossover, with  $\alpha = 0.5$ . Gaussian mutation may then be applied to each gene, with each  $\sigma$  being 10% of the domain's width, i.e.,  $\sigma_{D_r} = 0.5$  and  $\sigma_{\theta_r} = 0.2\pi$ . Finally, a new population is created with the individuals chosen through elitist selection.

### 3.4 Evaluation Environment



**Fig. 4.** Time average of air-flow (grey) and gas dispersion (green) of an instance of the evaluation environment. The green circles represent a screenshot of the odour plume at the last simulation step. The rectangle surrounding the chemical source (black circle) denotes the region from which its location is sampled, whereas the rectangle at the lower right corner represents the start region for the robots. (Color figure online)

A simulator specific for ER and odour source localisation experiments is used [10], as it showed to be much faster than real-time while adequately modelling the real world phenomena of odour dispersion and air-flow. The simulator uses a filament-based model of chemical dispersion, where each odour filament contains a given amount of molecules. An environment is devised based on a  $40 \times 40 \text{ m}^2$  arena containing no obstacles and a single odour source, being its parameters presented on Table 2. The stability of the wind and chemical emission rate were set to create an intermittent meandering plume. Each robot is modelled as a circular two-wheeled differential unit, measuring 16 cm in diameter. The robots are able to measure the wind velocity and chemical concentration at their location, as well as the distance to nearby obstacles through simulated laser range finders. Each laser range finder has 10 equally spaced beams, with a field of view of  $\pi$  radians and a maximum range of 0.5 m. The robots move with a maximum linear speed of 0.5 m/s and a maximum angular speed of 0.3 rad. For further details, we direct the reader to [10]. Figure 4 presents one instance of the evaluation environment. Thirty instances are created with the same parameters, one for each independent trial. The position of the chemical source and the initial pose of the robot are sampled from the respective regions for each environment instance and kept fixed for all evaluations in the same trial. The robot successfully finds the source if it reaches a position closer than 0.5 m from it, within 1500 s.

**Table 2.** Environmental parameters


Parameter	Value
Filament emission rate ( $F_r$ )	0.1 Hz
Initial filament radius ( $f_r(0)$ )	0.0316 m
Filament growth rate ( $\gamma$ )	0.01 m <sup>2</sup> /s
Chemical emission rate ( $\bar{Q}$ )	9.86 mg/s
Detection threshold ( $D_t$ )	1 ng/cm <sup>3</sup>
Initial wind speed ( $W_s$ )	0.7 m/s
Initial wind direction ( $W_d$ )	0 rad
Std. dev. of the Gaussian noise emulating turbulence ( $W_v$ )	0.075
Size of each grid cell for computing the wind velocity ( $W_c$ )	6 × 6 m <sup>2</sup>
Diffusivity constant ( $K_x$ )	6 m <sup>2</sup> /s
Dimensions of the arena	40 × 40 m <sup>2</sup>
Robot's initial pose	$x \sim U(36, 39)$ m, $y \sim U(36, 39)$ m, $\theta \sim N(-\pi/2, \pi/4)$ rad
Source position	$x \sim U(8, 14)$ m, $y \sim U(19.2, 21.2)$ m
Simulation step	0.5 s
Simulation time	1500 s

## 4 Experimental Results


Thirty independent trials of the evolutionary approach were conducted for evolving formations of three and five robots. Afterwards, a validation step took place, consisting on taking the best formation from each trial and re-evaluating it on the thirty environment instances. The performance of the swarms is measured through the success rates and duration of successful trials achieved in validation. As plotted in Fig. 5, using five robots seems to lead to higher success rates and lower search times. In order to draw more robust conclusions, statistical hypothesis tests are applied. The Kolmogorov-Smirnov test was applied, showing that, at a 95% confidence interval, none of the data could be considered to follow normal distributions. As a result, the Wilcoxon test was selected to perform pairwise comparisons of the success rates and duration of the searches. Its results showed that there are no statistically significant differences between the success rates of swarms of 3 and 5 robots ( $p = 9.7848e-01$ ) but that there are significant differences between the duration of the successful searches ( $p = 4.6796e-03$ ), implying that using 5 robots leads to locating the sources significantly faster.

To compare the performance of the evolutionary approach with the hand-designed formations, the best formation for each amount of robots must be selected. This choice is made by selecting the formation that attains the highest success rates in validation. In case of ties, the formation that takes the least





**Fig. 5.** Boxplots of the success rate (left) and duration of successful runs (right) of the evolved strategies in validation.



**Fig. 6.** Best formations evolved for three (left) and five (right) robots.

amount of time to locate the chemical source is chosen. Figure 6 presents the best formations found for three (E3) and five (E5) robots.

To better assess the performance of the evolved formations, their performance is compared to those of the baseline formations, evaluated on the thirty environmental instances (Table 3). As can be seen, using one robot leads to the overall worst success rate and slowest searches. Also, the evolved formations consistently present higher success rates than the hand designed ones. In particular, E3 presents a higher success rate than all hand-designed formations, even those using 5 robots. Conversely, E5 is slower than L5 and C5. The statistical analysis is once again performed, with the results of the Kolmogorov-Smirnov test showing that, at a 95% confidence interval, none of the data can be considered to follow normal distributions. As a result, the Wilcoxon test is selected for the remaining analysis.

**Table 3.** Success rates and duration of successful searches of the baseline and best evolved formations

Formation	Success rate	Mean search duration
S1	70%	685.24 ± 284.21 s
L3	76.67%	507.67 ± 287.67 s
C3	80%	521.69 ± 177.25 s
L5	93.33%	348.29 ± 127.23 s
C5	76.67%	371.91 ± 140.43 s
E3	100%	409.57 ± 130.70 s
E5	96.67%	382.97 ± 185.84 s

The first analysis consists on comparing the success and duration of the searches of S1, E3 and E5, to assess whether the amount of robots used produces substantial performance gains. The Bonferroni correction is applied, adjusting the significance value to 0.0167. The Wilcoxon test shows that there are statistically significant differences between the success of S1 and E3 ( $p = 2.6998e-03$ ) and of S1 and E5 ( $p = 1.1412e-02$ ), but that there are no significant differences when comparing E3 and E5 ( $p = 3.1731e-01$ ). Regarding the duration of the searches, the Wilcoxon test shows that there are statistically significant differences between S1 and E3 ( $p = 3.1108e-05$ ) and S1 and E5 ( $p = 8.1842e-05$ ), but that there are no significant differences between E3 and E5 ( $p = 6.4351e-01$ ). These results show that using only one robot leads to significantly lower success rates and longer searches than using 3 or 5 robots in the formations specified by E3 and E5. Also, using E3 could be preferred to E5, as less robots are required to attain equivalent performances. However, using E5 should be more robust when there is a high risk of loosing robots.

The next step in the analysis consists on comparing the performance of hand-designed and evolved formations for the same amount of robots. Once again, the Wilcoxon test is applied with the adjusted significance value of 0.0167, showing that there are significant differences between the success of E3 and C3 ( $p = 1.4306e-02$ ) and E3 and L3 ( $p = 8.1510e-03$ ), but there are no significant differences between C3 and L3 ( $p = 6.5472e-01$ ). Regarding the duration of the searches, the Wilcoxon test shows no significant differences between the three robot formations. Analysing the 5 robot formations, the Wilcoxon test shows no significant differences between their success rates or duration of the searches. As a result, one may conclude that the shape of the formation is more important when using less robots.

## 5 Conclusions and Future Work

This paper proposed an evolutionary approach for evolving the shape of a formation of a robotic swarm tasked with locating odour sources. The formation is

guided by a leader, which is controlled by a simple reactive controller. The other robots are then attracted to their intended positions in the formation, which are defined as distances and angles to the formation's virtual centre. Formations of three and five robots were evolved and the best solutions were compared to each other, as well as to hand-designed formations of one, three and five robots. The results showed that there are no significant differences in the success rates of the evolved formations of three and five robots, but that using five robots leads to significantly faster searches. When comparing the overall best evolved swarms with hand-designed formations, a set of conclusions could be drawn: (1) using evolved formations of three and five robots consistently outperforms using a single robot; (2) the best evolved formation of 3 robots is significantly more successful than the hand-designed swarms, but their search times are equivalent; and (3) there are no statistically significant differences between the performances of all five robots formations. As a result, one could opt by using the evolved formation of three robots, minimizing the amount of robots needed. Conversely, in case there is high risk of loss of robots, the evolved five robots formation could be preferred, as it presents slightly better performance indicators than the hand-designed ones.

In the future, studies shall be conducted to analyse the influence of the loss of agents. Furthermore, co-evolution techniques shall be investigated for jointly evolving the shape of the formation and the controller for its leader.

**Acknowledgements.** This work was partially supported by the Portuguese Foundation for Science and Technology (FCT), under projects UID/EEA/00048/2020, UID/CEC/00326/2020 and Ph.D. studentship COVID/BD/152379/2022, co-funded by the European Social Fund, through the Regional Operational Program Centro 2020. It was also partially supported by the Ultrabot project, funded by the Portuguese National Innovation Agency (ANI), under reference CENTRO-01-0247-FEDER-072644.

## References

1. Barnes, L., Fields, M.A., Valavanis, K.: Unmanned ground vehicle swarm formation control using potential fields. In: 2007 Mediterranean Conference on Control & Automation, pp. 1–8. IEEE (2007)
2. Bredeche, N., Fontbonne, N.: Social learning in swarm robotics. *Philos. Trans. R. Soc. B* **377**(1843), 20200309 (2022)
3. Chen, X., Huang, J.: Odor source localization algorithms on mobile robots: a review and future outlook. *Robot. Auton. Syst.* **112**, 123–136 (2019)
4. Coppola, M., McGuire, K.N., De Wagter, C., De Croon, G.C.: A survey on swarming with micro air vehicles: Fundamental challenges and constraints. *Front. Robot. AI* **7**, 18 (2020)
5. Eiben, A.E., Smith, J.E.: Introduction to Evolutionary Computing. NCS, vol. 53. Springer, Heidelberg (2003). <https://doi.org/10.1007/978-3-662-05094-1>
6. Feng, Q., et al.: An experimental and numerical study on a multi-robot source localization method independent of airflow information in dynamic indoor environments. *Sustain. Urban Areas* **53**, 101897 (2020)

7. Hamann, H.: *Swarm Robotics: A Formal Approach*, vol. 221. Springer, Heidelberg (2018). <https://doi.org/10.1007/978-3-319-74528-2>
8. Hettiarachchi, S., Spears, W.M.: Moving swarm formations through obstacle fields. In: IC-AI, pp. 97–103 (2005)
9. Lochmatter, T., Aydın Göl, E., Navarro, I., Martinoli, A.: A Plume Tracking Algorithm Based on Crosswind Formations. In: Martinoli, A., et al. (eds.) *Distributed Autonomous Robotic Systems*. Springer Tracts in Advanced Robotics, vol. 83, pp. 91–102. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-32723-0\\_7](https://doi.org/10.1007/978-3-642-32723-0_7)
10. Macedo, J., Marques, L., Costa, E.: A comparative study of bio-inspired odour source localisation strategies from the state-action perspective. *Sensors* **19**(10), 2231 (2019)
11. Macedo, J., Marques, L., Costa, E.: Designing fitness functions for odour source localisation. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 103–104 (2021)
12. Marjovi, A., Marques, L.: Optimal swarm formation for odor plume finding. *IEEE Trans. Cybern.* **44**(12), 2302–2315 (2014)
13. Marques, L., Nunes, U., de Almeida, A.T.: Cooperative odour field exploration with genetic algorithms. In: *Proceedings of 5th Portuguese Conference on Automatic Control (CONTROLO 2002)*, pp. 138–143. Citeseer (2002)
14. Marques, L., Nunes, U., De Almeida, A.T.: Odour searching with autonomous mobile robots: an evolutionary-based approach. In: *Proceedings of the IEEE International Conference on Advanced Robotics*, pp. 494–500 (2003)
15. Marques, L., Nunes, U., de Almeida, A.T.: Particle swarm-based olfactory guided search. *Autonom. Robot.* **20**(3), 277–287 (2006)
16. Murata, S., Kurokawa, H.: Self-reconfigurable robots. *IEEE Robot. Automa. Mag.* **14**(1), 71–78 (2007)
17. Nolfi, S., Bongard, J.C., Husbands, P., Floreano, D.: Evolutionary robotics. *Commun. ACM.* **56**, 74–83 (2016)
18. Andrew Russell, R., Bab-Hadiashar, A., Shepherd, R.L., Wallace, G.G.: A comparison of reactive robot chemotaxis algorithms. *Robot. Autonom. Syst.* **45**(2), 83–97 (2003)
19. Vergassola, M., Villermaux, E., Shraiman, B.I.: ‘infotaxis’ as a strategy for searching without gradients. *Nature* **445**(7126), 406–409 (2007)
20. El Zoghby, N., Loscri, V., Natalizio, E., Cherfaoui, V.: Chapter 8: Robot cooperation and swarm intelligence. In *Wireless Sensor and Robot Networks: From Topology Control to Communication Aspects*, pp. 163–201. World Scientific (2014)