

Understanding the Potential Impact of Multiple Robots in Odor Source Localization

Thomas Lochmatter and Alcherio Martinoli

Abstract We investigate the performance of three bio-inspired odor source localization algorithms used in non-cooperating multi-robot systems. Our performance metric is the distance overhead of the first robot to reach the source, which is a good measure for the speed of an odor source localization algorithm. Using the performance distribution of single-robot experiments, we calculate an ideal performance for multi-robot teams. We carry out simulations in a realistic robotic simulator and provide quantitative evidence of the differences between ideal and realistic performances of a given algorithm. A closer analysis of the results show that these differences are mainly due to physical interference among robots.

1 Introduction

With the advances in robotics and chemicals sensor research in the last decade, odor sniffing robots have become an active research area. Notably the localization of odor sources would allow for very interesting robotic applications, such as search and rescue operations, safety and control operations on airports or industrial plants, and humanitarian demining [21] [5] [17] [8]. Many of these applications are time-critical, i. e. odor sources should be found as fast as possible. Moreover, as the structure of plumes in the air is intermittent in both time and space [22], tracking plumes is a challenging problem.

Through real-robot [16] [15] and simulation [14] experiments, we have recently shown that the *surge-spiral* [6] [7] [2] [4] and the *surge-cast* [15] algorithms are faster and more reliable than pure *casting* [11] [10] [23] [13] [12] [1] in laminar wind flow. The experiments were run using a single robot, and the result was in-

Thomas Lochmatter, Alcherio Martinoli
Distributed Intelligent Systems and Algorithms Laboratory (DISAL), École Polytechnique Fédérale de Lausanne (EPFL), Station 2, 1015 Lausanne, Switzerland. e-mail: thomas.lochmatter@epfl.ch, alcherio.martinoli@epfl.ch

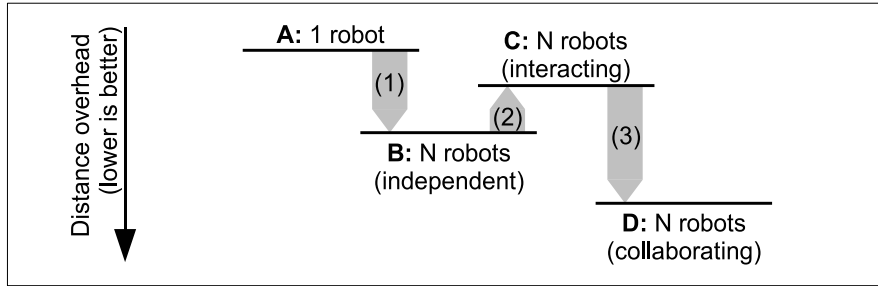


Fig. 1 Distance overhead of single-robot vs. multi-robot systems. (1) Performance gain due to randomness (mathematically derived in Section 3). (2) Performance loss due to physical interference among robots (simulated with a robotic simulator). (3) Performance gained with collaboration among robots (not discussed in this paper).

so far surprising, as the casting algorithm got much more attention by the research community up to date.

In this paper, we are studying the same algorithms with multiple non-cooperating robots. In particular, we compare the performance difference when moving from a single-robot to a homogeneous multi-robot system with 2 or 5 robots. Our performance metric is the distance overhead (traveled distance d_t divided by upwind distance d_u), which is an excellent indicator for the speed of a plume following algorithm on a holonomic robot [15]. Moreover, we only require one robot to reach the odor source, and use the distance overhead of the first robot to reach the source as the performance of the robotic team.

As sketched in Figure 1, the difference in distance overhead between single-robot (**A**) and non-cooperating multi-robot systems (**C**) consists of two components. First, randomness due to the noise in the system boosts the performance (Figure 1 (1)). This performance gain can be calculated by using the distribution of the distance overhead of single-robot experiments, and would be achieved if the robots were not interacting with each other (**B**). Second, physical interference among the robots result in a loss in performance (Figure 1 (2)), which we quantify by running simulations in a realistic robotic simulator [19]. Cooperation among robots (**D**, not discussed in this paper) would again result in a performance gain.

Multi-robot odor source localization experiments with an algorithm called *spiral surge* (which is close to the *surge-spiral* algorithm used here) have previously been carried out by Hayes et al. [6] [7]. Hayes ran experiments with up to 6 real robots, and up to 10 robots in simulation. Results showed that increasing the number of robots is beneficial in terms of time to find the source. To our knowledge, casting strategies have never been tested with multiple robots.

In two other projects [9] [18], multi-robot odor source localization algorithms based on Particle Swarm Optimization (PSO) were tested in simulation. In both papers, the robots were communicating with each other.

The remainder of this paper is structured as follows. In Section 2, we present the three algorithms used in this paper. In Section 3, we derive the ideal performance

for non-cooperative multi-robot systems. The simulator and the odor propagation model are introduced in Section 4. Finally, we discuss the results in Section 5 and conclude in Section 6.

2 Algorithms

All three algorithms used in this paper are bio-inspired [20] and use binary odor information. That is, they either *perceive the odor* or *do not perceive any odor*, but ignore different concentrations levels. Finally, all three algorithms need a wind sensor to measure the wind direction.

With the *casting* algorithm, the robots move in a zig-zag pattern under an upwind angle β towards the source. With *surge-spiral*, the robots move upwind as long as they are in the plume, and spiral with a gap length d_{gap} to reacquire the plume whenever they lose it. The *surge-cast* algorithm, finally, works in a very similar fashion, except that the robot casts in crosswind direction for a distance d_{cast} to reacquire the plume. A detailed description of these algorithms can be found in [15].

To avoid collisions, all robots are running a Braitenberg obstacle avoidance algorithm using 9 on-board infrared proximity sensors. Both algorithms run (independently from each other) in parallel at all times. While obstacle avoidance has virtually no influence on the behavior of the robot in open space, it overrides the plume tracking algorithm when the robot is close to an obstacle (i. e., another robot in our case).

Finally, a robot gives up and stops after having lost the plume for too long, or reached the arena boundary.

In this paper, we only consider plume traversal and intentionally omit plume finding (i. e. randomized or systematic search until the plume is found) and source declaration (i. e. declaring that the source is in close vicinity), to prevent those two parts from interfering in the results. Hence, the robots start *in* the plume, and source declaration is done by a supervisor (ideal source declaration). Experiments are considered successful as soon as the first robot has come in physical vicinity of the source, and unsuccessful if all robots gave up.

3 Expected Performance of Multi-Robot Experiments

Assume a performance value q that can be associated with each experimental run. In this paper, this metric is the distance overhead (traveled distance d_t divided by upwind distance d_u) of the first robot that reaches the source. Hence, a small q value stands for a good performance, with $q = 1$ being the optimum.

The distance overhead of an experiment with a single-robot algorithm (with a fixed set of parameters) can be expressed as a distribution Q_1 , which can be ap-




Fig. 2 Q_1 : Experimentally measured distribution of the distance overhead with a single robot. Q_2 and Q_5 : The expected distributions for 2 resp. 5 robots, based on the assumption that the robots do not physically interfere with each other. The triangles on top of the diagram indicate the mean values of the respective distributions.

proximated by performing a large number of runs. Examples of such distributions estimated with 200 runs are shown in Figure 2.

If two independent robots are going for the same source, their performances q_a and q_b are random samples drawn from Q_1 . Clearly, the smaller of these two numbers (corresponding to the faster robot) will set the overall performance of the team,

$$q_{ab} = \min(q_a, q_b) \quad (1)$$

Hence, the performance distribution of a system with two independent robots is the distribution of q_{ab} , and can be expressed as

$$Q_2(q) = \frac{1}{c_Q} \iint [\min(q_a, q_b) = q] Q_1(q_a)Q_1(q_b) dq_a dq_b \quad (2)$$

$$\text{with } c_Q = \iint Q_1(q_a)Q_1(q_b) dq_a dq_b \quad (3)$$

where $[\cdot]$ stands for the Iverson bracket. Generalizing this for N robots is straightforward.

3.1 Calculating Q_N

Closed-form expressions for Q_N only exist for a few well-known distributions. If Q_1 is exponentially distributed with mean $\frac{1}{\lambda}$, for instance, then Q_N is exponentially distributed with mean $\frac{1}{N\lambda}$.

The algorithms used in this paper yield complicated distributions, however, and an approximation by an exponential distribution would be very rough for the *surge-spiral* and *surge-cast* algorithms, and not justifiable for the *casting* algorithm. We therefore calculated the distributions for multiple robots numerically, by randomly sampling from the distribution Q_1 (Monte-Carlo simulation). Formally, we estimated the distribution Q_N with 100000 samples of the form

$$\min(q_1, q_2, \dots, q_N) \quad (4)$$

where q_1, q_2, \dots, q_N are randomly selected performance samples of the single-robot runs. Distributions obtained in this way for 2 and 5 robots executing the *casting* algorithm are shown in Figure 2. On that figure, it can be observed how the distribution and its mean value shift towards the left (lower distance overhead) as the number of robots increases.

4 Simulation Experiments

We are using Webots [19] for the experiments. Webots is a commercial realistic robotic simulator, which ships with a calibrated model of the Khepera III robot that we used for the real-robot experiments [15]. The simulation environment (Figure 4) was enhanced with a wind and odor propagation model, and the robot model was extended with the corresponding sensors to measure the odor concentration and a wind direction (Figure 3).

The simulation setup was kept close to the setup in the wind tunnel that we used in previous work for the experiments with real robots [16] [15]. The simulation time step, Δt , was set to 32 ms.

4.1 Experimental Arena

The experimental arena was a rectangular area of 18 m length and 4 m width, which corresponds roughly to the dimensions of the wind tunnel. At 1 m from one end of the arena, a circular odor source of radius 12 cm was placed. The robots were placed at roughly 14.5 m downwind from the source.

4.2 Advection Model

A constant wind field of 1 m/s was used, which corresponds to a constant laminar flow comparable to the one in the wind tunnel. In the coordinate system indicated in Figure 4, the wind vector at position u , $a(u)$, can be written as

$$a(u) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \quad (5)$$

4.3 Odor Propagation Model

The odor propagation model closely resembles the filament-based model proposed by Farrell et al. [3]. This model is easy to implement and requires only a very limited amount of CPU power. Yet, it generates an intermittent plume which is similar to the real plume in the wind tunnel.

Odor is thereby simulated as a set of filaments ($i = 0, \dots, N$), each containing a constant amount $s = 8.3 \cdot 10^9$ of molecules or particles. Each filament is defined by its position, $p_{i,t}$, and its width, $w_{i,t}$.

In each time step, the position of a filament is updated according to the wind flow and a stochastic process:

$$p_{i,t+\Delta t} = p_{i,t} + a(p_{i,t})\Delta t + v_p \quad (6)$$

The stochastic component v_p is a vector of three independent Gaussian random variables, $N(0, \sigma_p^2)$, with standard deviation $\sigma_p = 0.1$ m.

To model molecular dispersion, filaments become wider with time while their peak concentration decreases. The width of a filament evolves as

$$w_{i,t+\Delta t} = w_{i,t} + \frac{\gamma}{2w_{i,t}} \quad \text{with } \gamma = 4 \cdot 10^{-7} \quad (7)$$

Our virtual odor source released 100 such filaments per second with an initial width of $w_{i,0} = 10$ cm and an initial position which was uniformly distributed over the circular area of the source. This resulted in a plume comparable to the real plume in the wind tunnel.

4.4 Odor Sensor Model

The odor concentration at time t and position u was calculated as the sum over the concentration contribution of all filaments,

$$C_t(u) = \sum_{i=0}^N c_{i,t}(u) \quad (8)$$

and each filament i contributed

$$c_{i,t}(u) = \frac{s}{w_{i,t}^3} \exp\left(-\frac{|u - p_{i,t}|}{w_{i,t}^2}\right) \quad (9)$$

to the concentration. Hence, the concentration decayed exponentially with increasing distance from the center of a filament.

The virtual odor sensor reported this concentration $C_t(u)$ without adding any additional noise, as the perceptual noise related to the chemical-to-electrical transduc-

Fig. 3 Simulated Khepera III robot equipped with an odor sensor (small cylinder on top of the robot) and a wind sensor (big cylinder). The hexagons in the air represent odor filaments.





Fig. 4 Simulated environment (18 m by 4 m arena) in Webots.



tion is negligible even on the real platform [15]. Furthermore, since the concentration is anyway thresholded and filtered through d_{lost} by the algorithms in use here, a precise calibration of the odor propagation and odor sensor model was not required.

4.5 Wind Direction Sensor Model

The wind sensor reported a noisy wind measurement,

$$a_s(u) = a(u) + v_a \quad (10)$$

where v_a was a vector with samples of a zero-mean normal distribution ($N(0, \sigma_a^2)$). Since the wind field was constant in all our simulations, the reported value in world coordinates was simply

$$a_s(u) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} N(0, \sigma_a^2) \\ N(0, \sigma_a^2) \\ N(0, \sigma_a^2) \end{pmatrix} \quad \text{with } \sigma_a = 0.1 \text{ m} \quad (11)$$

This vector was rotated into the local reference system of the robot to account for the robot's pose.

4.6 Experiments

For all three algorithms, we run experiments with 9 different parameters (upwind angle β for *casting*, spiral gap d_{gap} for *surge-spiral*, and cast distance d_{cast} for *surge-*

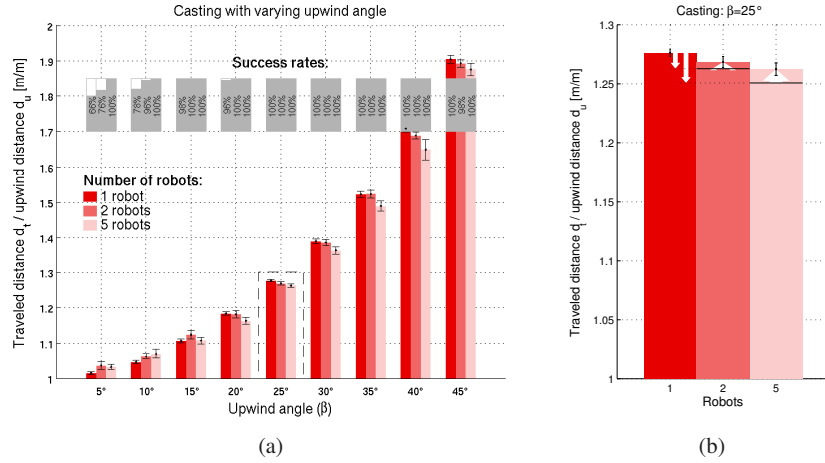


Fig. 5 (a) Results obtained with the *casting* algorithm. The error bars indicate the 95 % confidence interval for the mean. (b) Close-up for $\beta = 25^\circ$. The thin arrows indicate the intrinsic performance gain by passing from a single-robot to a multi-robot system (Figure 1 (1)), while the thick arrows indicate the performance loss due to physical interaction between the robots (Figure 1 (2)).

cast), each with 1, 2, or 5 robots. For $\beta = 25^\circ$, $d_{\text{gap}} = 22$ cm and $d_{\text{cast}} = 34$ cm, we performed 200 independent runs and calculated the ideal performance as described in Section 3. For all other configurations, 50 runs were carried out.

In each run, the robots were released in the odor at fixed positions (evenly spaced) between 14.5 m and 16 m downwind from the source. If one robot reached the odor source, the run was stopped and considered successful. During the run, the trajectory, the measured odor concentration and the measured wind direction were recorded for each simulation step. Distance and upwind distance were derived from the trajectory.

The forward speed of the robot (on straight lines) was 10.6 cm/s and therefore same as with the real-robot experiments in the wind tunnel. The plume threshold was set to $c = 100$.

5 Results and Discussion

5.1 Casting

The results for the *casting* algorithm are presented in Figure 5. The differences between the single-robot and the multi-robot experiments are very small and statistically not significant for most configurations. However, as a general trend, multiple robots seem to yield slightly better performance for upwind angles $\beta > 20^\circ$, and worse performance otherwise.

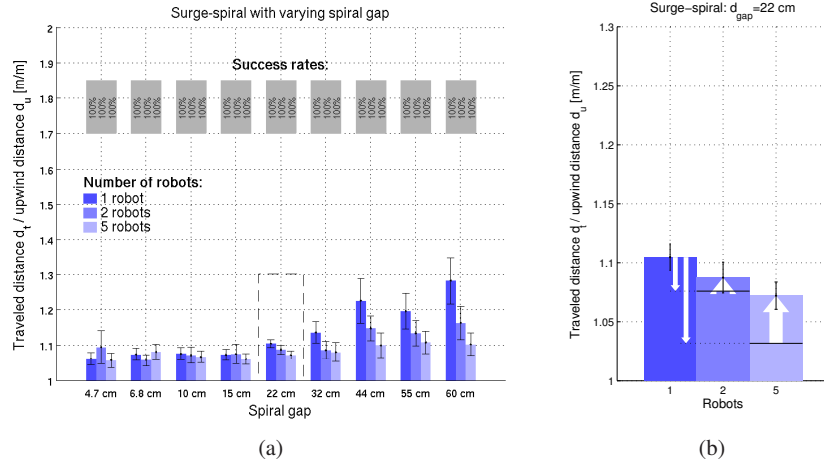


Fig. 6 (a) Results obtained with the *surge-spiral* algorithm. (b) Close-up for $d_{gap} = 22$ cm.

As Figure 5 (b) reveals, even the ideal performance (for robots that are not physically interfering) is not much better than the single-robot performance. Indeed, single-robot experiments with the *casting* algorithm yield a compact — almost normal — performance distribution with a small variance, and the resulting “left shift” of the distribution for multiple robots is small.

A noticeable gain can be observed for the success rate, however. For small angles, where the success rate with a single robot is small, a team of robots can achieve very high success rates. This robustness is an advantage often cited in the context of multi-robot systems. Surprisingly, physical interference seems not to have a big influence here. As the following table shows, the actual success rates obtained in the multi-robot experiments are close to the expected success rates calculated based on the success rate of the single-robot runs.

	1 robot	2 robots		5 robots	
	actual	actual	expected	actual	expected
<i>Casting</i> , $\beta = 5^\circ$	0.66	0.76	0.884	1.0	0.995
<i>Casting</i> , $\beta = 10^\circ$	0.78	0.96	0.952	1.0	0.999

5.2 Surge-Spiral

The picture for the *surge-spiral* algorithm looks pretty different. As the performance distribution of single-robot runs resembles an exponential distribution, its mean value decreases as $\frac{1}{N}$ with increasing numbers N of robots. Hence, large performance gains are expected.

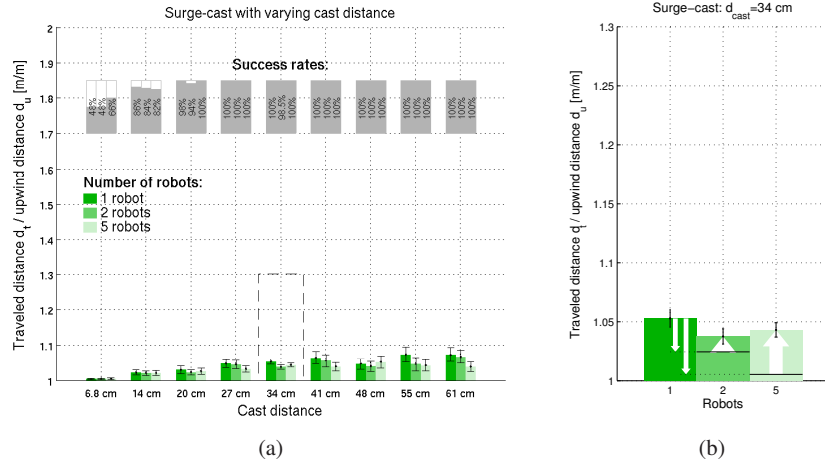


Fig. 7 (a) Results obtained with the *surge-cast* algorithm. (b) Close-up for $d_{cast} = 34$ cm.

For small spiral gaps, no performance gain is noticed in the simulation results. For large gaps, however, the difference between the single-robot and the multi-robot runs are significant. Even though the theoretical model would expect an even larger difference, the multi-robot runs were clearly faster than the single-robot runs and achieved similarly good results as the single-robot runs with small spiral gaps. The reason for this is that a spiraling robot spends enough time aside the plume, such that other robots can safely overtake. This could be interpreted as an indirect and implicit coordination scheme (without communication), whereby robots losing the plume try to make room for other robots in the plume.

5.3 Surge-Cast

The results for the *surge-cast* are similar: the longer the cast distance, the more performance is gained by using multiple robot. This has to be taken with a grain of salt, though, since longer cast distances yield worse performance in the single-robot case and are therefore not desired anyway. Hence, in well-configured systems with near-optimal cast distances (here 27 cm - 34 cm), no performance gains are visible.

Contrary to the *casting* experiments, using multiple robots does not increase the robustness of the algorithm here. In some cases, the success rate even got worse. The *surge-cast* algorithm in its present form is clearly not robust with respect to physical collisions. Especially during plume reacquisition, a robot blocking the way at the plume boundary can cause another robot to lose the plume completely. This could certainly be improved by adaptively increasing the cast distance until the plume is found.

6 Conclusion

We carried out single-robot and non-cooperative multi-robot odor source localization experiments in simulation with three different bio-inspired algorithms, and compared their results in terms of success rate and distance overhead. The setup was similar to the single-robot experiments carried out with real Khepera III robots in the wind tunnel [15].

While the theoretically ideal performance of the multi-robot teams was expected to be significantly better than that of the single-robot runs, the actual performance was found to be comparable for most configurations. Statistically significant differences were mainly found for the *surge-spiral* algorithm with large spiral gaps. With such large gaps, the robots leave the plume for a significant time and distance, which allows other robots to overtake without interference. This underlines the negative effect of physical interference among the robots on the team performance. In particular, uncoordinated teams of robots have troubles overtaking each other during plume following. A simple, local coordination scheme dealing with this problem would presumably offer a significant performance gain.

In future work, we will test the algorithms in turbulent flow and/or meandering plume conditions, and test different multi-robot coordination schemes in simulation and using real robots.

Acknowledgements This work was supported by the National Competence Center in Research on Mobile Information and Communication Systems NCCR-MICS, a center supported by the Swiss NSF under grant number 5005-67322.

References

1. Eugene Balkovsky and Boris I. Shraiman. Olfactory search at high Reynolds number. *PNAS*, 99(20):12589–12593, October 2002.
2. Jim H. Berlinger and Mark A. Willis. Adaptive control of odor-guided locomotion: behavioral flexibility as an antidote to environmental unpredictability. *Adaptive Behavior*, 4(3-4):217–253, August 1996.
3. Jay A. Farrell, John Murlis, Xuezhong Long, Wei Li, and Ring T. Cardé. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environmental Fluid Mechanics*, 2:143–169, 2002.
4. Gabriele Ferri, Emanuele Caselli, Virgilio Mattoli, Alessio Mondini, Barbara Mazzolai, and Paolo Dario. A biologically-inspired algorithm implemented on a new highly flexible multi-agent platform for gas source localization. In *Proceedings of the First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics (BIOROB 2006)*, February 2006.
5. Douglas W. Gage. Many-robot MCM search systems. In *Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium*, pages 9.56–9.64, April 1995.
6. Adam T. Hayes, Alcherio Martinoli, and Rodney M. Goodman. Distributed odor source localization. *IEEE Sensors Journal*, 2(3):260–271, June 2002.
7. Adam T. Hayes, Alcherio Martinoli, and Rodney M. Goodman. Swarm robotic odor localization: Off-line optimization and validation with real robots. *Robotica*, 21:427–441, 2003.

8. Hiroshi Ishida, Takamichi Nakamoto, Toyosaka Moriizumi, Timo Kikas, and Jiri Janata. Plume-tracking robots: A new application of chemical sensors. *Biological Bulletin*, (200):222–226, April 2001.
9. Wisnu Jatmiko, Kosuke Sekiyama, and Toshio Fukuda. A pso-based mobile robot for odor source localization in dynamic advection-diffusion with obstacles environment. *IEEE Computational Intelligence Magazine*, pages 37–51, May 2007.
10. L. P. S. Kuenen and H. C. Rowe. Cowpea weevil flights to a point source of female sex pheromone: analyses of flight tracks at three wind speeds. *Physiological Entomology*, 31(2):103, June 2006.
11. Wei Li, Jay A. Farrell, and Ring T. Cardé. Tracking of fluid-advected odor plumes: Strategies inspired by insect orientation to pheromone. *Adaptive Behavior*, 9(3-4):143–170, 2001.
12. Wei Li, Jay A. Farrell, Shuo Pang, and Richard M. Arrieta. Moth-inspired chemical plume tracing on an autonomous underwater vehicle. *IEEE Transactions on Robotics*, 22(2):292–307, April 2006.
13. Achim J. Lilienthal, Denis Reiman, and Andreas Zell. Gas source tracing with a mobile robot using an adapted moth strategy. In *Autonome Mobile Systeme (AMS), 18. Fachgespräch*, pages 150–160. GDI, December 2003.
14. Thomas Lochmatter and Alcherio Martinoli. Simulation experiments with bio-inspired algorithms for odor source localization in laminar wind flow. In *Proceedings of the The Seventh International Conference on Machine Learning and Applications (ICMLA 2008)*. IEEE, December 2008. To appear.
15. Thomas Lochmatter and Alcherio Martinoli. Tracking odor plumes in a laminar wind field with bio-inspired algorithms. In *Proceedings of the 11th International Symposium on Experimental Robotics 2008 (ISER 2008)*, Springer Tracts in Advanced Robotics (2010), Athens, Greece, July 2008. To appear.
16. Thomas Lochmatter, Xavier Raemy, Loïc Matthey, Saurabh Indra, and Alcherio Martinoli. A comparison of casting and spiraling algorithms for odor source localization in laminar flow. In *Proceedings of the 2008 IEEE International Conference on Robotics and Automation (ICRA 2008)*, pages 1138–1143, May 2008.
17. Matt Long, Aaron Gage, Robin Murphy, and Kimon Valavanis. Application of the distributed field robot architecture to a simulated demining task. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005)*, pages 3193–3200, April 2005.
18. Lino Marques, Urbano Nunes, and A. T. de Almeida. Particle swarm-based olfactory guided search. *Autonomous Robots*, 20(3):277–287, June 2006.
19. Olivier Michel. Webots: Professional mobile robot simulation. *International Journal of Advanced Robotic Systems*, 1(1):39–42, 2004.
20. R. Andrew Russell. *Odour Detection by Mobile Robots*, volume 22 of *World Scientific Series in Robotics and Intelligent Systems*. World Scientific Publishing Company, 1999.
21. Gary S. Settles. Sniffers: Fluid-dynamic sampling for olfactory trace detection in nature and homeland security—the 2004 freeman scholar lecture. In *Journal of Fluids Engineering*, volume 127 of *Transactions of the ASME*, pages 189–218, 2005.
22. Massimo Vergassola, Emmanuel Villermaux, and Boris I. Shraiman. ‘Infotaxis’ as a strategy for searching without gradients. *Nature*, 445:406–409, January 2007.
23. Barbara Webb, Reid R. Harrison, and Mark A. Willis. Sensorimotor control of navigation in arthropod and artificial systems. *Arthropod Structure and Development*, 33:301–329, May 2004.