

# A Plume Tracking Algorithm based on Crosswind Formations

Thomas Lochmatter, Ebru Aydın Göl, Iñaki Navarro and Alcherio Martinoli

**Abstract** We introduce a plume tracking algorithm based on robot formations. The algorithm is inherently designed for multi-robot systems, and requires at least two robots to collaborate. The robots try to keep themselves centered around the plume while moving upwind towards the source, and share their odor concentration and wind direction measurements with each other. In addition, robots know the relative poses of other team members. Systematic experiments with up to 5 real robots in a wind tunnel show that the robots achieve close-to-optimal performance in our scenario, and by far outperform previous approaches. The performance gain is attributed to the fact that robots continuously share information about the plume (odor concentration, wind direction) without spatially competing for acquiring it.

## 1 Introduction

With the advances in robotics and chemical 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 [15] [12] [3]. 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 [17], tracking plumes is a challenging problem.

---

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

Iñaki Navarro  
Universidad Politécnica de Madrid, ETSI Industriales, c/ José Gutiérrez Abascal 2, E-28006 Madrid, Spain.

In previous work [9], we analyzed three bio-inspired plume tracking algorithms based on upwind surge, casting, and spiraling, and carried out experiments with real robots and in simulation. Experiments with the multi-robot versions of these algorithms [8] thereby revealed that robots are competing for space even when they are communicating and collaborating, a limitation which is especially important in narrow plumes. When robots avoid each other (to prevent collisions), they often lose the plume and switch to plume reacquisition, which obviously results in performance degradation. Reasons for this are the state-machine nature of the bio-inspired algorithms on one hand, but also the fact that robots do not plan their path with respect to the positions of the other robots. The low-level controller (obstacle avoidance) interferes with the high-level controller (bio-inspired odor source localization algorithm).

In this paper, we present a novel algorithm to tackle exactly this problem. The algorithm is based on a loose robot formation which collaboratively moves through the plume towards the source. Robots thereby communicate with each other, and share all observations (wind and odor concentration) as well as their relative positions. A simple reactive control scheme keeps the robot formation centered (in crosswind direction) around the plume while the robots are moving upwind towards the source.

This algorithm is *not* an extension of some existing single-robot algorithm, but inherently designed for multi-robot system. (At least 2 robots are necessary for the algorithm to work.) *By design*, robots do not compete for space, and no low-level avoidance algorithm is necessary to prevent robots from bumping into each other<sup>1</sup>. In addition, robots truly and continuously collaborate: every single observation is shared, and the control algorithm uses the observations of all robots in the formation as input.

The remainder of this paper is structured as follows. In Section 2, we present existing approaches for multi-robot odor source localization. Section 3 formally introduces the crosswind formation algorithm. After depicting the experimental setup in Section 4, we show real-robot experiments with a static source (Section 5) and a moving source (Section 6). Finally, we conclude in Section 7.

## 2 Related Work

Most odor source localization algorithms found in the scientific literature are intended for single-robot systems. To our knowledge, there only exist four multi-robot approaches to date.

Hayes et al. applied a bio-inspired algorithm based on spiraling and upwind surge to multiple robots [2] [1], and studied the effect of a primitive broadcasting communication scheme. In particular, they studied a communication scheme called KILL

---

<sup>1</sup> An obstacle avoidance algorithm may still be necessary to prevent robots from bumping into obstacles, of course.

in which all robots stop as soon as one perceives an above-threshold odor concentration, and a scheme called ATTRACT whereby robots that do not perceive any odor join others that have found plume information. Experiments were carried out with real robots and in an embodied simulation, and the performance metric was a linear combination of time and group energy, the latter being proportional to the sum of the distances the robots traveled. The KILL strategy was found to save significant power, whereas the ATTRACT strategy did not reveal any advantage in their setup. The multi-robot experiments we carried out in simulation [8] are conceptually similar.

Second, all algorithms based on PSO [13] [5] [6] [4] are intended for use in multi-robot systems. PSO requires robots to communicate at least locally, and robots must be aware of each other's position. In the standard PSO algorithm, collaboration is however limited. The only variable they share is the (locally or globally) highest concentration, and the position where it was measured. As a result, robots will have a tendency to move towards the same local optimum if communication is global, and may bump into each other. Jatmiko et al. therefore introduced an extension called CPSO [6] in which robots share their positions and use a repulsive force to avoid collisions and make sure robots remain spread over some area. Robots however do not directly take concentration measurements of other robots into account.


Another algorithm inherently designed for multi-robot systems is the *fluxotaxis* algorithm [16]. As its name suggests, it is based on the chemical mass flux, which is a product of the chemical density and the flow velocity. Fluxotaxis is therefore a hybrid between chemotaxis and anemotaxis. As positive divergence of mass flux indicates a source, this algorithm basically climbs up the mass flux gradient. This is done with a flock of robots that share their observations (concentration and air-flow) with a central controller. Fluxotaxis has been shown to perform well even in complex scenarios with obstacles.

Finally, the infotaxis algorithm has been extended to multiple robots [14]. Robots thereby share all their observations (concentrations and positions) to collaboratively infer the location of the source. All information from all robots is integrated into a single model, which is the maximum amount of information robots can share. The authors reported that in some scenarios, super-linear performance increase can be achieved by using multiple robots.

### 3 The Crosswind Formation Algorithm

The underlying idea of the algorithm presented in this paper is to keep some robots on the left side of the plume, and some robots on the right side of the plume. While they are going upwind, they try to stay centered in the plume, i. e. keep the (average) concentrations on the left and on the right approximately equal. The formation chosen here is a line formation in crosswind direction.

Each robot can measure the odor concentration and the wind direction (relative to its pose) at its current location, and shares this information with all members of the



**Fig. 1** Sketch of the formation algorithm from the perspective of robot 3. Calculation of the forces  $f_u$  and  $f_c$  is carried out in a reference system defined by the wind direction. The resulting force vector is then rotated into the robot's reference system. All robots carry out the exact same calculation, but from their own perspective.

formation. In addition, robots can measure the relative poses of each other. A global reference system is not required, but the wind measurement actually provides the robots with an estimate of a global direction.

In a reference system defined by the wind direction, each robot calculates a (virtual) crosswind and a (virtual) upwind force, as depicted in Figure 1. The upwind force,  $f_u$  is

$$f_u = u + \frac{1}{N} \sum_i y_i \quad (1)$$

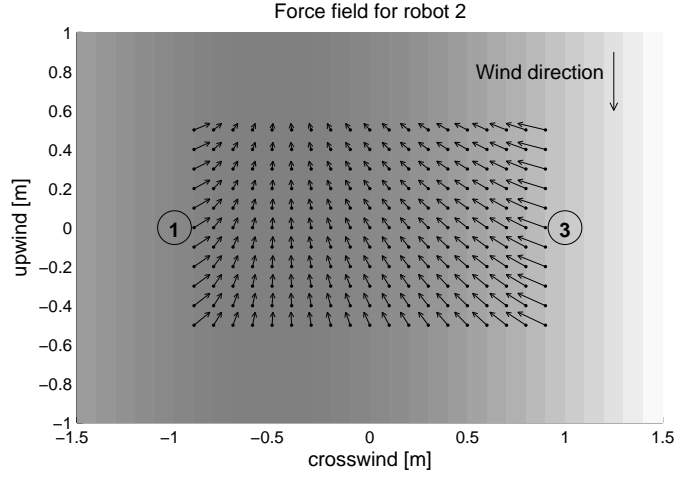
where  $N$  denotes the number of robots and  $u$  the constant upwind drag, a parameter of the algorithm.  $f_u$  keeps the robot aligned with the other robots, such that they all have approximately the same downwind distance from the odor source. If, for instance, one robot is behind, the  $y_i$  tend to be more positive in the coordinate system of this robot and the resulting force is stronger.

The crosswind force,  $f_c$ , is a weighted difference (with weights  $a$  and  $r$ ) between an attractive and a repulsive force. Formally,

$$f_c = af_a - rf_r \quad (2)$$

$$f_a = \frac{\sum_i x_i c_i}{\sum_i c_i} \quad (3)$$

$$f_r = \frac{1}{N} \sum_{i, i \neq \text{me}} \frac{1}{x_i} \quad (4)$$



**Fig. 2** Crosswind and upwind forces (scaled) for different positions of robot 2 in a formation with 3 robots. Robot 1 is currently measuring a 4 times higher concentration as compared to robot 3. The gray shading stands for the odor concentration, which attains its maximum at  $x = -0.5$  m.

The attractive force,  $f_a$ , takes into account the odor concentrations,  $c_i$ , measured by all other robots and is responsible for keeping the formation centered in the plume. Robots measuring a high concentration contribute more weight, and therefore pull the other robots towards them. The repulsive force,  $f_r$ , keeps the robot at a certain distance from all other robots.

The vector  $(f_u, f_c)$  is then rotated into the reference system defined by the robots heading,

$$f_v = f_u \cos(-\alpha) - f_c \sin(-\alpha) \quad (5)$$

$$f_h = f_u \sin(-\alpha) + f_c \cos(-\alpha) \quad (6)$$


where  $\alpha$  denotes the wind angle relative to the robots heading. The resulting vector  $(f_v, f_h)$  is finally transformed into differential drive wheel speeds as follows:

$$s_l = s(k_v f_v + k_h f_h) \quad (7)$$

$$s_r = s(k_v f_v - k_h f_h) \quad (8)$$

$k_v$  and  $k_h$  are thereby factors to scale the forward and differential speed appropriately, and  $s$  denotes the mean forward speed.

All robots keep executing these steps continuously in a loop. The loop speed is approximately the same on all robots, but robots are not tightly synchronized. In each iteration of the loop, a robot takes one measurement with the wind direction sensor and one with the odor sensor, and broadcasts the latter to all other robots. To



**Fig. 3** Schematic drawing of the arena (not to scale) with the approximate location of the plume.

calculate the forces, it uses the last received odor concentration and relative position values of each robot.


## 4 Experimental Setup

The experiments were carried out in a 16 m long and 4 m wide wind tunnel. The setup was the same as described in [10], except that the arena inside the tunnel was enlarged to approximately 15 m by 3.5 m. The same setup was also used to analyze three bio-inspired algorithms [7]. In the following paragraphs, we briefly repeat the most important figures.

The wind field in the wind tunnel was laminar at roughly 1 m/s speed. The ethanol odor plume was therefore a straight line (see Figure 3), and the concentration peaks were slightly decreasing as the plume moved downwind. A constant amount of ethanol vapor was released by means of a pump. To reduce the turbulence created by the odor source, the pump was placed outside of the arena and connected with a tube to the source outlet. Nevertheless, the outlet created some turbulence right downwind the source, which sometimes disturbed the laminar wind flow in that area. The starting area was 14 meters downwind from the outlet, as depicted in Figure 3.

### 4.1 Robotic Platform

The robot used in the experiments was a Khepera III robot (K-Team SA, Switzerland) equipped with an odor sensor and a wind sensor board, as depicted in Figure 4 (a).



**Fig. 4** (a) The Khepera III robot with the wind sensor and the odor sensor board. (b) Picture of the wind sensor board. The two LEDs on top were used for tracking the robots with overhead cameras.

The odor sensor was a MiCS-5521 volatile organic compound (VOC) sensor, which has a very fast response time ( $\approx 0.1$  s). This sensor reacts to a wide range of organic compounds in the air, with an sensitivity to ethanol comparable to that of a human nose ( $\approx 10$  ppm). To take advantage of the sensor's low response time, air was taken in and released with a small pump.

A new version of the wind sensor board with 6 thermistors was used on the robots. The board reports the wind direction at a rate of  $\approx 10$  Hz with a standard error of  $4^\circ$  in our flow.

## 4.2 Relative Positioning

Relative positions were emulated using a camera system and sent to the robots via wireless LAN at a 10 Hz update rate. The camera system consisted of 6 overhead cameras, which recorded the whole arena. Each robot was equipped with two LEDs, and SwisTrack [11] was used to detect these markers on the recorded images. The global accuracy after calibration was about 8 cm, while the precision was around 4 cm.

## 5 Experiments with a Static Source

We tested the crosswind formation algorithm with the following settings:

Algorithm	Robots	Start position	Runs
A Crosswind Formation OSL	3	left	10
B Crosswind Formation OSL	3	middle	10
C Crosswind Formation OSL	3	right	10
D Crosswind Formation OSL	5	middle	5

With the start position *left* (resp. *right*), the robots started slightly at the left (resp. right) of the plume, and only the rightmost (resp. leftmost) robot was measuring an above-baseline odor concentration. With the start position *middle*, one robot was placed in the plume center at the beginning of the experiment, while an equal number of robots started on the left and on the right of the plume. The mean forward speed of the robots was  $s = 7.1$  cm/s, and the parameters of the force model were set as follows:

	$u$	$a$	$r$	$k_v$	$k_h$
Experiments with 3 robots (A, B, C)	1	1	0.1225	1	1
Experiments with 5 robots (D)	1	1	0.4225	1	1

Note that no attempt was made to systematically optimize these parameters, as the main objective is to demonstrate odor source localization using formations, and not formation control itself.

Before each run, the odor concentration baseline was determined individually for each robot by taking a few measurement samples in fresh air. The sensitivities of the odor sensors were not systematically calibrated, but believed to be approximately the same. Slight sensitivity differences would result in a little drift of the formation in crosswind direction, but hardly affect the performance.

Two metrics were calculated for each run:

- ▷ The *success rate* of the team is the number of successful runs (where robots keep the formation and reach the source) divided by the total number of runs.
- ▷ The *distance overhead* of a robot is its traveled distance divided by the distance of the shortest path to the source (straight line). The distance overhead of the team is the average over the distance overheads of all team members.

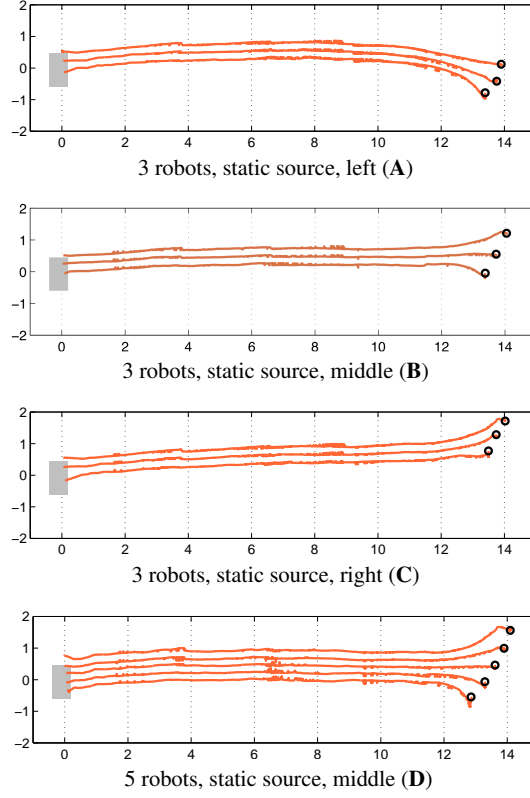
Figure 5 shows one run of each setting. No matter where the robots started, the robots found the center of the plume within the first 2 m upwind distance and then continued going straight upwind. The distance overheads are therefore extremely low, as shown in Figure 6. The success rate was 100 % in all settings.

In our setup, this algorithm clearly outperforms all bio-inspired algorithms [7] [9] in terms of distance overhead and success rate. The distance overhead for most runs was below 2 %, and for some runs even below 1 %. Included in this overhead is the initial phase in which robots get into the predefined formation shape. Without this, the results would be even closer to the optimal performance.

This is not surprising: with sensors on the left and on the right of the plume, the formation obtains direct feedback about its position with respect to the plume, and can correct for that long before leaving the plume completely. A plume reacquisition phase, as it is used with the bio-inspired algorithms is not necessary any more. The robots keep going upwind without ever losing the plume.

Using 5 instead of 3 robots did not improve the performance in our experiments. There is no a priori reason for which 5 robots should yield worse results. After all, 5 robots collect even more information about the plume than 3 robots, and should therefore do at least an equally good job. However, the information gain when switching from 3 to 5 robots might be tiny, and therefore irrelevant in our






**Fig. 5** Real-robot trajectories produced by crosswind formation algorithm. The gray rectangle represents the target area with the odor source while the black circles denote the starting positions of the robots. The robots go almost straight upwind towards the source, yielding very low distance overheads and high success rates.


setup. Main reason for the performance drop here are presumable the outermost robots, which started at a suboptimal position quite far away from the plume center, and first had to move closer to the center. We believe, however, that increasing the number of robots would be advantageous in settings with a sparser plume.

## 6 Experiments with a Moving Source

Since this algorithm measures the odor concentration at several points at the same time, it is particularly well suited for scenarios with moving sources. With a single sensor, an algorithm is unable to tell with a single measurement in which direction the source moved. Such information can only be deduced from multiple (sequential or parallel) measurements at different locations.



**Fig. 6** Distance overheads (by starting position) of the experiments with three robots. The dots stand for the distance overhead of the robots in individual runs, and are classified by the robot's position within the formation. The bars indicate the mean distance overhead over all runs and all robots.



**Fig. 7** Real-robot trajectories produced by the crosswind formation algorithm tracking a source moving in crosswind direction. The gray rectangle represents the target area with the odor source while the black circles denote the starting positions of the robots. The robots nicely follow the movement of the plume.

Multi-robot algorithms provide just this: taking several measurements at the same time. This allows the formation to know immediately whether the source moved towards the left or towards the right. The force model takes advantage of that information in that it tries to keep the robots centered around the plume.

We carried out 5 runs with the same algorithm tracking a moving source. The source was thereby moved back and forth by 92 cm in crosswind direction at constant speed. All other parameters of the setup and the algorithm were kept the same. Two of these runs are drawn in Figure 7. While all runs were successful, it is not possible to calculate the distance overhead (as defined above) in this case. The trajectories however reveal that the algorithm works fine.

## 7 Conclusion

In this paper, we presented the *crosswind formation* algorithm, an odor source localization algorithm which is inherently designed for multi-robot systems. In its present form, the algorithm requires at least 2 robots to collaborate. Experiments were carried out with 3 and 5 real robots in a scenario with laminar flow and no obstacles, and showed that this algorithm achieves close-to-optimal performance in terms of distance overhead and success rate. We also demonstrated the algorithm’s robustness in a scenario with a moving source.

The algorithm performs substantially better than previous attempts with bio-inspired algorithms that were extended to multi-robot algorithms. With the formation-based approach used here, a single controller takes care of avoiding collisions among robots and tracking a plume at the same time. Hence, robots do not compete for space, and do not block each other’s way. Instead, robots take advantage of their relative positions, as well as their wind direction and concentration measurements, and strive for a common goal: finding the source. Nevertheless, the algorithm is simple and has — except for relative localization — low requirements.

Future work could address different formations or failing robots. In addition, the positioning requirements may be reduced to direct-neighbor-only information. The algorithm should also be tested in more complex scenarios, such as scenarios with obstacles, turbulence, more complex source motions models, or multiple sources. Also of interest is a comparison with single-robot algorithms based on multiple odor sensors, which could be regarded as multi-robot algorithms with perfect formations.

**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. Adam T. Hayes, Alcherio Martinoli, and Rodney M. Goodman. Distributed odor source localization. *IEEE Sensors Journal*, 2(3):260–271, June 2002.
2. 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.

3. 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.
4. W. Jatmiko, K. Sekiyama, and T. Fukuda. A pso-based mobile sensor network for odor source localization in dynamic environment: Theory, simulation and measurement. In *IEEE Congress on Evolutionary Computation, 2006. CEC 2006*, pages 1036–1043, 2006.
5. Wisnu Jatmiko, Yusuke Ikemoto, Takayuki Matsuno, and Toshio Fukuda. Distributed odor source localization in dynamic environment. In *Proceedings of the 4th IEEE Conference on Sensors (Sensors 2005)*, pages 254–257. IEEE, 2005.
6. 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.
7. 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)*, volume 54 of *Springer Tracts in Advanced Robotics* (2009), pages 473–482, Athens, Greece, 2008. Springer.
8. Thomas Lochmatter and Alcherio Martinoli. Understanding the potential impact of multiple robots in odor source localization. In *Proceedings of the 9th International Symposium on Distributed Autonomous Robotic Systems (DARS 2008)*, Distributed Autonomous Robotic Systems 8, pages 239–250, Tsukuba, Ibaraki, Japan, November 2008.
9. Thomas Lochmatter and Alcherio Martinoli. Theoretical analysis of three bio-inspired plume tracking algorithms. In *ICRA'09: Proceedings of the 2009 IEEE international conference on Robotics and Automation*, pages 3195–3202, Piscataway, NJ, USA, 2009. IEEE Press.
10. 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.
11. Thomas Lochmatter, Pierre Roduit, Christopher Cianci, Nikolaus Correll, Jacques Jacot, and Alcherio Martinoli. SwisTrack: A flexible open source tracking software for multi-agent systems. In *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, pages 4004–4010. IEEE/RSJ, 2008.
12. 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.
13. Lino Marques, Urbano Nunes, and A. T. de Almeida. Particle swarm-based olfactory guided search. *Autonomous Robots*, 20(3):277–287, June 2006.
14. J.-B. Masson, M. Bailly Bechet, and Massimo Vergassola. Chasing information to search in random environments. *Journal of Physics A: Mathematical and Theoretical*, 42, 2009.
15. 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.
16. Diana F. Spears, David R. Thayer, and Dimitri V. Zarzhitsky. Foundations of swarm robotic chemical plume tracing from a fluid dynamics perspective. *International Journal of Intelligent Computing and Cybernetics*, 2(4):745–785, 2009.
17. Massimo Vergassola, Emmanuel Villermaux, and Boris I. Shraiman. ‘Infotaxis’ as a strategy for searching without gradients. *Nature*, 445:406–409, January 2007.