Extension of Plume Tracking Behavior to Robot Swarms

S. Kazadi
Jisan Research Institute
28 North Oak Avenue
Pasadena, CA 91107
USA
sanza@jisan.org

Abstract

Generating an effective algorithm for plume tracking is a relatively straightforward task in plumes of large density. However, in the low density regime, plume tracking can be problematic, as a single packet odorant may not be sufficient to locate the source of the plume. We present work on the application of swarm engineering to the plume tracking problem. A swarm of mobile robots is designed which can track a virtual plume to it's source. We illustrate that the sensitivity of the swarm is significantly greater than that of the single robots, though the basic algorithm is identical on both sets of robots. We provide some motivation for the use of swarms of robots in specific regimes not well served by single robots.

keywords: swarm engineering, plume tracking

1 Introduction

Odor plumes are dynamic structures composed of moving packets of odorant which emanate from a single source and are carried by a sustained wind blowing in a well defined direction. Recently, a great deal of work has been done on the development of methods of tracking the plume to its source. While many animals are capable of carrying out this seemingly difficult task, few robotic systems have been developed that demonstrate a capability to carry out this task. Most research programs designed to tackle the plume tracking problem have centered around the use of single robots which are capable of tracking the plume to a single destination. All algorithms will fail, however, in the face of weak plumes of small concentration. This low concentration limit does not seem to be a limitation of the sensor, but rather a limitation of the search algorithm. Extending this limit, therefore, is an important part of allowing robots to locate the sources of plumes, even in the most challenging cases.

The most successful strategies in plume tracking are those that utilize wind direction information. This piece of information is so important, that the lack of it can be the main reason for failure. Despite the great need for this piece of information, the way in which it is obtained can vary widely. Moreover, the accuracy of this information is less of a critical issue than one might expect, and four directions can normally suffice for good performance, though only two need be used.

Russell et al. [12] report the design of a robot which is capable of tracking a camphor plume up-wind to its source. In these experiments, the wind direction is used in conjunction with gradient information. Assuming that the robot can reliably determine the center of the plume by following a gradient to the center of the plume, the robot's two behaviors include acquisition of an approximate center of the plume by determining the differential concentration at two crystal-based sensors spaced apart, and the following of the plume up-wind, as determined by an on-board wind vane. An obstacle avoidance behavior was included by adding a sensitive bumper to the system and forcing the robot to make on-the-spot turns when it came into contact with an obstacle, effectively avoiding it. The robot was capable of reliably tracking the plume and avoiding obstacles from a meter away.

Ishida et al. [6] describe a robotic system in which an 'active probe' is constructed. This 'active probe' uses a fan to draw air to the sensor, effectively handling the boundary layer problem[1], and the difficulty associated with a limited amount of available odorant. Wind direction information is obtained by rotating this sensor by 360° and recording the direction that has the largest reading. Later versions included stationary directional sensors fabricated from four sensors arranged in a cross, providing quicker but less accurate direction information. The new design could also be used to find three-dimensional plume information, though this information was unused.

The simplest algorithm for determining the position of the plume source is one employed by Kuwana et al. [9] in which the robot remains completely within the plume as it is carrying out its tracking behavior. The robot is inactive when outside of the plume, and active when inside the plume. Its behavior is simple, forcing the robot to execute turns when it approaches the edge of a plume, and to walk straight when completely within the plume. This strategy has two main weaknesses. Firstly, the robot cannot find the plume, but rather must be initially placed within the plume in order for it to function at all. Secondly, half the time the robot will tend to go downstream, and half upstream. A bad turn produced by noisy or unexpected sensor data might cause the robot to move in the wrong direction down the plume.

A similar algorithm may be used to track saline plumes in water. Grasso et al. [4] report the construction of a robotic 'lobster' designed to track saline plumes in water. The robot is endowed with sensors which can determine the local concentration of saline. Using this information, the robot climbs a concentration gradient from the plume source. The speed of each motor is controlled by the sensed concentration of the salt in the water. Turns can then be affected by changes in the concentration. The robot, which initially moves from within the plume, will also employ a backing behavior if the concentration becomes too low, indicating that the plume has been exited.A possible subgroup is that which uses dynamic properties of the odorant packets sensed by the robot. However, only one group [3] has reported the possible use of this. Airborne plumes typically move too quickly to make use of this data.

Kazadi [8] produced a four-legged robot capable of tracking a plume to its source by traversing the plume and executing turns upon partial loss of the plume. The robot utilized re-

---

[1] The *boundary layer problem* refers to the tendency of an odorant to pass over a sensor without directly interacting with it. This is a serious problem when dealing with odorants whose presence is intermittent or sporadic.

sistive polymer sensors [10] and simple analog circuitry. This robot, like those of Kuwana et. al. and Grasso et. al., is active only when inside the plume. The robot's legs' motions are tied to the sensors; when the sensor on a given side of the robot is active (or responding to odorant), the legs on the other side of the robot start moving. The resulting zig-zag behavior through the plume takes the robot upstream. Moreover, the placement of the sensors and their detailed properties seem to bias to robot toward upstream turns, which makes the robot track upstream despite initially facing downstream.

In this paper, we explore one method of extending the apparent sensory capability of plume tracking systems. The method, based in swarm engineering methodology, utilizes large groups of simple plume tracking robots. These robots, designed as simply as possible, and utilizing a minimal repertoire of capabilities, can behaviorally extend the functional sensitivity of the individual robot, making target identification much more reliable than single robots. The remainder of the paper is organized as follows. In the Section 2, we briefly review the odor plume model and our spiral-based algorithm upon which we base this work. Section 3 discusses the groups of robots that make up our study, including individual non-interacting agents, self-avoiding agents, and communicating swarms. Section 4 describes the performance of the agents in the plume tracking task. Finally, Section 5 summarizes the work, offering some concluding remarks and future avenues of research.

2 Odor Plumes and the Wind-Based Spiral Algorithm (WBSA)

The problem with plume tracking algorithms primarily comes from the loss of the plumes in turbulent fluids, and the reaquisition of the plume. This is a result of the plume's overall structure. If the plume is strong, its parts are connected to one another. However, in many cases, this is not the case, and individual packets of odorant travel downstream. The number of packets can be small or large, depending on the strength of the plume. The plume tracker, then, must be able to deal with these structural discontinuities and still succeed in approaching and identifying the source.

In this study, we are interested in exploring the regions of sensitivity of single simulated robots, and their swarm counterparts. We utilize a plume simulator created by Jay Farrell of the University of California at Riverside in all of these simulation experiments [2]. This simulator delivers odorant to the search space in the form of packets. Each packet is spherical in shape, has a well-defined size, and has a density which drops off as the distance squared $\left(r^{-2}\right)$ from the center of the plume. Each packet is carried by a simulated wind downstream. As the packet ages, it widens, lowering its peak concentration, and increasing the area over which it has some influence. The wind is designed to emulate natural wind, and to produce plumes with meander. This allows one to investigate the effectiveness of plume strategies in variable wind conditions, though the meander does not closely model real wind.

Many parameters are important in specifying a plume. We summarize these in Table 3.1.

Quantity	Value
Packet Initial Size	0.1 cm
Packet Growth Rate	10.0%/s
Wind Speed	1 m/s
Plume Length	10 m
Wind Meander Parameter	2.5

Table 3.1: This table gives several of the parameters associated with the plume. The values given here reflect the condition of the plume in use. The wind meander parameter is a characteristic of the plume's inherent spectral noise [15]. It is beyond the scope of this paper to thoroughly discuss this parameter.

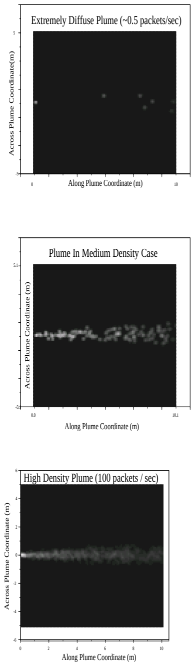

Figure 1: Simulated plumes with 1, 10, and 100 packets/second released into the plume.

A simple variation of the original simulator allows the concentration of the plume to be greatly varied, making study of the behavior of different algorithms under widely variable conditions possible. In Figures 2-4 we illustrate the various concentration levels, spanning three orders of magnitude. Notably, in the most sparse plume, the probability of encounter is exceedingly small.

The main question, then, is how the performance under these extreme conditions may be improved by use of a swarm. In the next section, we illustrate the sensitivity limitations of single agents. The intention is to use this as a baseline for measuring the performance of minimally communicating swarms, and of explicitly communicating swarms.

In these computational experiments, we utilize a simple algorithm that may be implemented with a minimum of computational overhead on our robots. This is called the *wind-based spiral algorithm (WBSA)* is studied. This particular algorithm assumes the presence of a wind direction sensor, and is designed to deal with acquisition and re-acquisition problems of the plume. The algorithm is specified as

1. Initiate a spiral from the current direction using a given chosen direction.
2. If an odorant is detected above threshold, go to 4. Otherwise go to 3.
3. Continue current spiral. Return to 2.
4. Turn up-wind and move directly up-wind.
5. If no odorant is detected, go to 6. Otherwise go to 5.
6. Reverse spiral direction. Go to 1.

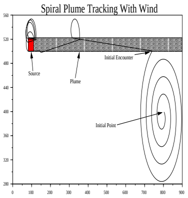

Figure 2: This is a case of the WBSA locating the source of a plume.

This algorithm has been shown to be extremely effective in finding the source in near optimal distance and time. Moreover, the algorithm has been shown to be noise resistant, and rather fault tolerant. In the presence of a plume of suitable strength, this is an extremely effective algorithm.

3 Sensitivity Measurements of Plume Tracking Robots and Swarms

In this section, we report on the performance of differing groups of robots in the location of the source of the plume. These groups are groups of non-interacting robots which act without knowledge of the other members of the group, groups of robots which try to avoid one another, so as to cover a particular area, and a swarm of communication robots. We report on the sensitivity measurements for all three groups locating sources of plumes of differing intensities.

3.1 Measurements

Our measures [1] of performance are scaled performance measures, giving the relative performance of the robot using the given algorithm to one with perfect knowledge which may turn directly toward the target and move right to it. These are defined by

$$tdf = \frac{t}{t_p} \tag{1}$$

and

$$ddf = \frac{d}{d_p} \tag{2}$$

where t and d represent the time and distance expended and traveled, respectively, and t_p and d_p represent their perfect counterparts.

Moreover, we report on the arrival times and distances of the first agent to arrive given a particular plume concentration and number of agents. As we increase the number of solitary searchers, we expect this performance to improve, on average. The amount that it improves as well as the numerical values of these average runs provides us with a view of how different methods of communication may be exploited to more quickly locate an odor source.

Finally, we report on the performance of the group of robots. For each parameter setting, we report on how many robots make it to the source in a span of ten seconds. All robots are initiated in a random initial position centered around twelve meters downstream. As the plume extends only ten meters, the robots must first locate the plume, and then track it upstream. The robot is considered to have arrived if it is located to the left of the origin and is between -2 m and 2 m from the origin. It must also have detected some odorant to be counted as having arrived. A sample single robot track is given in Figure 3.

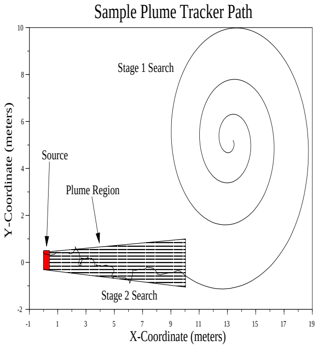

Figure 3: This is a typical robot track, illustrating the two regions of search. In the first region of search, the robot is locating the plume. In the second region of search, the robot is using the plume to find the plume source.

Figure 3.1: This is a typical robot track, illustrating the two regions of search. In the first region of search, the robot is locating the plume. In the second region of search, the robot is using the plume to find the plume source.

The agents have a maximum speed of 17.0 meters per second, an angular resolution on the wind sensor of 90 ^{o}, noise on both each of the two wheel motors causing a variation in the speed equaling up to 10% of the commanded speed, and noise on the wind direction sensor with a Gaussian distribution and standard deviation 45 ^{o}. This robot configuration is used throughout this paper. In this study, we use virtual data and sensors. The sensors utilized can have the sensitivity arbitrarily set, and we set it to perfect detection; the sensor can detect any odorant of arbitrarily low concentration.

Each of the behaviors in Section 3.2 is applied to the search at each of the concentrations given in Figures 4.1 and 4.2 in ten independent runs. The results are averaged, and their scores are presented in Section 4.

3.2 Behaviors

While all agents are started in a group scattered nearby a common starting point, the different groups have different behaviors. The first group of agents consists of independent searchers. These agents act simultaneously, but without noticing each other during the search. In this respect, each agent is point-like, and neither affects or is affected by any other agent. These agents may be used to baseline the performance of the search algorithm.

The second group of agents consists of self-avoiding agents. These agents search for odorant, as in the case of the first set of agents. However, when encountering another agent, these agents first move as directly away from nearby agents as possible, and then re-initiate their spiraling behavior. This serves to spread the agents over a large distance quickly, making the search more evenly spread out. The hope of this algorithm is to make it easier to cover a space. A sample trajectory of an agent engaged in this search mode is given in Figure 3.2.

The trajectory consists primarily of long lines, indicating a primary focus on avoidance, at least initially.

In designing the swarm of agents, we apply a swarm criterion, or design criterion for the swarm. This criterion is given by:

> **Swarm Criterion**: The swarm must, at the individual agent level, use the spiral search algorithm. It must expand in areas of little or no density of odorant, and contract in areas of higher density, while making progress up the plume in measurable odorant density areas.

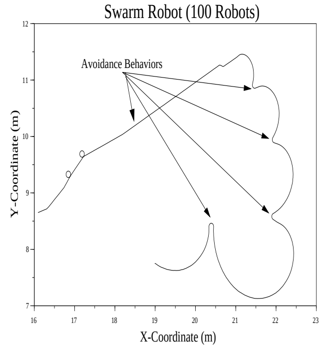

Figure 4: This figure illustrates the trajectory of a robot in the avoidance behavior.

This criterion does a number of things at once. First, it requires that the swarm be capable of implementing a standard WBSA at the individual level. In the absence of any other agents, a single-agent swarm will be virtually indistinguishable from any other single agent. Next, it requires that the swarm be a space-filling entity in areas of low odorant density. This behavior would seem to allow the swarm to efficiently fill a space and discover any source of odorant or any plumes. The swarm must contract upon a plume or odorant source in order to concentrate resources in areas where the odorant may be found, but also be sparse. Finally, the swarm must progress upstream. This precludes any kind of communication that brings the robots back downstream.

We implement these conditions by making a few addendums to the WBSA.

> **Swarm Design**: Agents located in high density areas will seek low density areas. Agents in low density areas will execute WBSA. If encountering an odorant, a robot will turn on a beacon to call other agents, and continue to execute WBSA. Those agents located downstream of the transmitting agent will move directly toward the transmitting agent, avoiding collisions only. Convergence behavior will take precedence over expansion behavior. The beacon will stay on only for a predetermined period after the loss of the signal.

This behavior serves to fulfill the swarm requirements, while utilizing a minimal strategy.

4 Plume tracking with groups of agents

We report four different comparative performance measures of the agent groups. As described above, we report the TDF and the DDF. We also report the minimal search time for each of the algorithms, and the success rate of the robots in the group.

Figure 4.1 illustrates the TDF and DDF performances of the groups in the plume tracking problem. These are calculated by averaging the scores of all agents in the groups. Most notably, the non-interacting robots actually perform better than the self-avoiding robots in performance at low plume concentrations. Moreover, the non-interacting robots perform comparably to

Figure 4.1: These figures give the TDF and DDF values for the three groups of interacting robots. Not only does the swarm perform much better, on average, at high density, but the turning point between minimal and maximal performance is significantly lower in plume concentration.

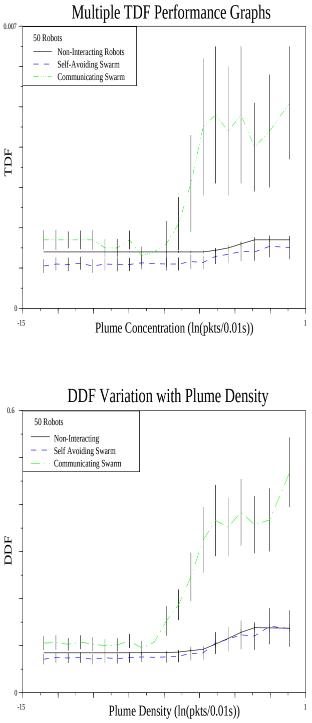

Figure 5: These figures give the TDF and DDF values for the three groups of interacting robots. Not only does the swarm perform much better, on average, at high density, but the turning point between minimal and maximal performance is significantly lower in plume concentration.

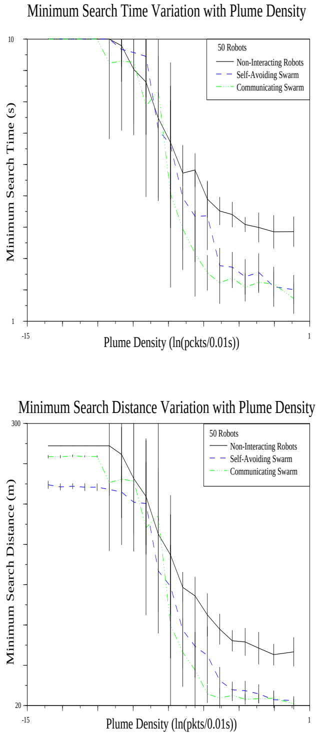

Figure 6: This figure gives the performance of the three algorithms as measured by their minimal time or minimal distance.

the swarm of robots at low density. This makes a great deal of sense, as the difficulty for these systems at low concentrations is in locating the plume in the first place. However, the lower performance of the self-avoiding robots may be attributed to interference in a group size of fifty robots.

At higher density, the performance of all three groups increases, though the performance of the communicating swarm sharply increases disproportionately when compared to the others. Moreover, the increase occurs at a significantly smaller concentration than for the other algorithms. This indicates that the swarm of agents is not only able to find the plume reliably at lower concentrations, but also capable of bringing a majority of its agents to the source in the process. While the performance notably smaller than the perfect score of 1.0, this reflects the inability of the swarm of robots to directly move to the source, even in cases of high plume density. The much lower scores of the self-avoiding or singleton groups indicate that while they reliably find the source at concentrations of e^{-2} packets/0.01 sec, they are incapable of influencing other agents to come to the source, and so the average performance score is quite low.

We may also ask the question of how much time a search will take before the *first* robot arrives at the target, or how *far* the robot will have to move before arriving at the target. This is important if we are planning to make robots with time-critical missions or with limited power. Figure 4.2 illustrates the minimal time and distance of any member of each of these groups on the search task as a function of the packet density.

As with the previous measurements, the initial performance is comparable for all groups at low density when considering the search time. However the distance for the self-avoiding swarm is much smaller than either of the other two algorithms. This is a result of the expansive behavior of the robot group. The swarm of agents also benefits from this, though it is a significantly lesser degree due to the contraction behavior that occurs once a plume is detected. After the turning point, however, both the expansive and swarm groups have a significantly smaller minimal search time and distance than the non-interacting robots. This indicates that the expansion of the group, which allows the agents to quickly find the plume, is a factor in helping the agents reach the source.

At medium concentrations, the swarm does significantly better than the self-avoiding group. This indicates a capability of the swarm to explore efficiently which is not part of the self-avoiding group. However, at high concentrations, this effect disappears, indicating that much of the difference is swamped by the ability of the group of agents to locate the plume.

Thus, in both of these studies, we see that the interaction protocols improve the performance of the group of agents, though the communication protocols of the swarm have an overall, but not individual performance increase.

5 Discussion and Concluding Remarks

This is, of course, only a first look at the practical design of robot systems of this type. Much more work needs be done. However, the algorithms given in this paper may easily be ported to a real robot system, as the WBSA has been [7]. All communication protocols are binary, and may be altered easily to become fuzzy. Information a given robot must have is local, and need not include concentrations as measured by other robots. This means that the relative calibrations of sensors on different robots is not a critical issue in the generation of swarm algorithms, relieving the pressure that exists in the community of producers of olfactory technology. Indeed, if it were, it would be difficult to have been realized up to now.

The second purpose of this work is to develop a method of handling plumes of very low concentration, where the probability of interacting with the plume is very small. What we have illustrated above is that a swarm of interacting robots is capable of tracking plumes of very much smaller density than those that can be tracked by either independent or semi-independent swarms. The distinction is important. It is not sufficient to cover an area with robots. Rather, the robots must share information about the area, and use this to focus the search. This means that the swarm, viewed as an entity, must have the ability to react to local conditions, and to rearrange itself based on these conditions. In this way, the swarm may accurately locate the source of the plume. Despite these considerations, we have only been able to demonstrate a second order effect. The main effect seems to have been tied up in the dispersion effect in terms of delivering a first agent. This would seem to be a problem only in cases, as in those presented, that contain swarms that are initially in close proximity to the plume. Swarms starting further away from the plume would seem to be much more likely to exhibit a greater dominance of this second order effect.

The algorithms discussed in this paper are not generated by any desire to emulate their biological counterparts. However, we have noted in discussions and readings with some of the people working in the biological plume tracking domain that the algorithms created here are very similar to those found in the natural world [1], [11], [13]. Belanger and Willis report on the flight patterns of the moth *Manducta sexta* while tracking a pheromone plume up-wind. The behaviors found are strangely reminiscent of our WBSA algorithm while in the plume. This would seem to indicate that the algorithm is exploiting the information in the plume in a way that behaviorally is similar to that the moth produces. Holldobler and Wilson [5] report on the recruitment of leaf-cutter ants while searching for food. The mechanisms of recruitment, are similar to those we use, essentially a "sound beacon" which may be received by other leaf-cutter ants. The investigation of similarities and differences, both in the information used and methods of using the information might prove interesting and rewarding, both in the engineering and biological domains.

6 Acknowledgments

This work is supported in part by the Center for Neuromorphic Systems Engineering, in part as part of the National Science Foundation Engineering Research Center Program under grant EEC-9402726; and by the California Trade and Commerce Agency, Office of Strategic Technology under grant C94-0165. This work is also supported in part by DARPA/ONR under grant N00014-93-1-0990, and by AFOSR University Research Initiative grant F49620-93-1-0332.This research is also supported by the ONR under grant F49620-97-1-0514, and by DARPA under grant DAAK60-97-K-9503. ONR also provided funding under grant N00014-98-1-0821 Chemical Plume Tracing Via Biological Inspiration and Electronic Olfaction. We also acknowledge support from the DARPA program in Distributed Robotics.

References

[1] J. Belanger and M. Willis. 1996. *Adaptive control of odor-guided locomotion: behavioral flexibility as an antidote to environmental unpredictability.* **Adaptive Behavior**. 4:(3/4), 217-253.

[2] J. Farrell , J. Murlis, X. Long, W. Li, and R. Carde . *Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes.* **Environmental Fluid Mechanics**. 2: 143-169,2002.

[3] F. Grasso, T. Consi, D. Mountain, and J. Atema. 1996. *Locating Odor Sources in Turbulence with a Lobster Inspired Robot*, **Proc. 4th International Conference on Simulation of Adaptive Behavior,** Maes, Mataric, Meyer, Pollack, and Wilson (Eds), Cambridge, Massachusetts : MIT Press.

[4] F. Grasso, T. Consi, D. Mountain, and J. Atema. 1996. *Behavior of Purely Chemotactic Robot Lobster Reveals Different Odor Dispersal Patterns in the Jet Region and the Patch Field of a Turbulent Plume*, Biol. Bull.**,** 191: 312-313.

[5] B. Holldobler. and E. Wilson. 1990. **The Ants**. Cambridge, Massachusetts : Harvard University Press.

[6] H. Ishida et. al. 1996. *Odour Source Localization System Mimicking Behaviour of Silkworm Moth.* **Sensors and Actuators**, 99: 225-230.

[7] S. Kazadi. **Swarm Engineering**. Ph.D. Thesis, California Institute of Technology, 2000.

[8] S. Kazadi, R. Goodman, D. Tsikata, D. Green, H. Lin. *An Autonomous Water Vapor Plume Tracking Robot Using Passive Resistive Polymer Sensors.* **Autonomous Robots**, 9(2): 175-188, 2000.

[9] Y. Kuwana, I. Shimoyama, Y. Sayama, and H. Miura. 1996. *Synthesis of Pheromone-Oriented Emergent Behavior of a Silkworm Moth. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 1722-1729.*

[10] N. Lewis et. al. 1996. *Array based vapor sensing using chemically sensitive, carbon black-polymer resistors.* **Chem. Mater.** 8 (9):2298-2312.

[11] A. Mafra-Neto and R. Carde. 1994. *Fine scale structure of pheromone plumes modulates up-wind orientation of flying moths.* **Nature**, 369, 142-144.

[12] R. Russell, D. Thiel, R. Deveza , and A. Mackay-Sim. 1994. *A Robotic System to Locate Hazardous Chemical Leaks.* **IEEE International Conference on Robotics and Automation**. 13 (3): 232-239.

[13] P. Zanen and R. Carde. 1996. *Effects of host-odour plume altitude and changing wind velocity on up-wind flight manoeuvres of a specialist braconid parasitoid.* **Physiological Entomology**, 21, 329-338.