

Swarm Robotic Odor Localization

Adam T. Hayes, Alcherio Martinoli, Rodney M. Goodman

Microsystems Lab 136-93, California Institute of Technology, Pasadena CA 91125
athayes,alcherio,rogo@micro.caltech.edu

Abstract

This paper presents an investigation of odor localization by groups of autonomous mobile robots using principles of Swarm Intelligence. We describe a distributed algorithm by which groups of agents can solve the full odor localization task more efficiently than a single agent. We then demonstrate that a group of real robots under fully distributed control can successfully traverse a real odor plume. Finally, we show that an embodied simulator can faithfully reproduce the real robots experiments and thus can be a useful tool for off-line study and optimization of odor localization in the real world.

1 Introduction

This paper presents an investigation of odor localization by groups of autonomous mobile robots using principles of Swarm Intelligence (SI), a computational and behavioral metaphor for solving distributed problems that takes its inspiration from biological examples provided by social insects. In most biological cases studied so far, robust and capable group behavior has been found to be mediated by nothing more than a small set of simple interactions among individuals and between individuals and the environment [1]. The application of SI principles to autonomous collective robotics aims to develop robust task solving by minimizing the complexity of the individual units and emphasizing parallelism, exploitation of direct or indirect interactions, and distributed control. The main advantages of this approach are three: first, scalability from a few to thousands of units, second, flexibility, as units can be dynamically added or removed without explicit reorganization, and third, increased system robustness, not only through unit redundancy but also through the design of minimalist units. Several examples of collective robotics tasks solved with SI principles can be found in the literature: aggregation [2] and segregation [3], exploration [4], stick pulling [5], and collaborative transportation [6].

Recently, advances have been made in understanding biological and artificial odor classification and odor localization as developed in moths [7,8] and rats [9] in air, and lobsters [10] and stomatopods [11] in water. Biology utilizes olfaction for a wide variety of tasks including finding others of the same species, communication, behavior modification, avoiding predators, and searching for food. Odors, unlike visual and auditory perceptions, are non-spatial: they possess neither spatial metric nor direction. In contrast, odorant stimuli possess both spatial

and temporal character, snaking out complex plumes that can wander over a wide area. This implies that a level of sophistication beyond gradient following is necessary for localization of an odor source.

Animals use a combination of hardware (e.g. receptor adaptation), software (temporal integration and/or spatial integration), and search strategies (both intrinsic and landmark-based) to locate odor sources. Odor localization is in essence a behavioral problem that varies from animal to animal. Some exploit fluid information at different layers (lobster), or sense several residues on the ground (ants). Others can track odors in 3-D environments (moths) or use a combination of information to locate their targets (dogs). From an engineering standpoint some tasks are facilitated by combining odor sensing with mobile robots, such as the localization of chemical leaks and the chemical mapping of hazardous waste sites. We are interested in developing small mobile robots that use odor tracking algorithms and multi sensor (e.g. odometry, anemometry, olfaction) fusion to search out and identify sources of odor.

The aim of the case study described in this paper is three-fold. Firstly, we describe a distributed algorithm by which groups of agents can solve the full odor localization task more efficiently than a single agent. Secondly, we demonstrate that a group of real robots under fully distributed control can successfully traverse a real odor plume. Thirdly, we show that an embodied simulator can faithfully reproduce the real robots experiments and thus can be a useful tool for off-line study and optimization of odor localization in the real world.

2 The Odor Localization Problem

The general odor localization problem addressed in this paper is as follows: find a single odor source in an enclosed 2D area as efficiently as possible. This can be broken down into three subtasks: plume finding - coming into contact with the odor, plume traversal - following the odor plume to its source, and source declaration - determining from odor acquisition characteristics that the source is in the immediate vicinity. Plume finding amounts to a basic search task, with the added complication, due to the stochastic nature of the plume, that a sequential search is not guaranteed to succeed. Plume traversing requires more specialized behavior, both to progress in the direction of the source and to maintain consistent contact with the plume. Source declaration does not necessarily have to be done using odor information, as typically odor sources can

be sensed via other modalities from short range, but here we propose a solution using no extra sensory apparatus.

2.1 Biological Inspiration

As an odor source dissolves into a fluid medium, an odor plume is formed. The turbulent nature of fluid flow typically breaks the plume into isolated packets, areas of relative high concentration surrounded by fluid that contains no odor. The task of odor localization thus becomes one of plume traversal, or following the trail of odor packets upstream to the source.

Although the approach of moving slowly and continually sampling odor and flow data to reduce environmental noise is used in nature (starfish) and has been applied to robotic systems [12,13], environmental and behavioral constraints (e.g. significant plume sparseness or meander, time critical performance) can render these systems ineffective. In that case, upon sensing an odor signal, a good policy is to move directly upwind, because a good immediate local indication of source direction under such circumstances is the instantaneous direction of flow [14]. When the odor is no longer present, a good strategy is to perform a local search until it is reacquired, as the location of the previous packet encounter provides the best immediate estimate of where the next will occur. This type of behavior has been observed in moths [15], and its performance has been studied in simulation [8].

The previous work on this algorithm was aimed at studying biology, which limited the sensory and behavioral time scales investigated. When applying these ideas to robots, however, the separation between algorithm and underlying hardware is much more clear, and it no longer makes sense to constrain behavior strictly by sensory response characteristics. Therefore, in this work, key aspects of the search behavior, such as surge duration and casting locality, are treated as algorithm parameters.

2.2 The Spiral Surge Algorithm

The basic odor localization algorithm used in this study, Spiral Surge (SS), is shown in Figure 1. It consists of different behaviors related to the three different subtasks.




Fig. 1. Spiral Surge odor localization behavior.

Plume finding is performed by an initial outward spiral search pattern (SpiralGap1). This allows for thorough

TABLE I
Spiral Surge Algorithm Parameters

SpiralGap1	Initial spiral gap width
SpiralGap2	Plume reacquisition spiral gap width
StepSize	Surge distance post odor hit
CastTime	Length of time before reverting from reacquisition to initial search spiral
SrcDecThresh	Significance threshold between separate odor hits
SrcDecCount	Number of significant differences before source declaration

coverage of the local space if the total search area is large and initial information can be provided by the deployment point (an external 'best guess' as to source location). Alternatively, if no a priori knowledge is available, a spiral with a gap much greater than the arena size (producing essentially straight line search paths) provides an effective, although not optimal [16], search procedure. Future work will address search efficiency in greater detail.

Plume traversal is performed using a type of surge algorithm. When an odor is encountered during spiraling, the robot samples the wind direction and moves upwind for a set distance (StepSize). If during the surge another odor packet is encountered, the robot resets the surge distance but does not resample the wind direction. After the surge distance has been reached, the robot begins a spiral casting behavior, looking for another plume hit. The casting spiral can be tighter than the plume finding spiral (SpiralGap2), as post surge the robot has information about packet density and a thorough local search is a good strategy. If the robot subsequently re-encounters the plume, it will repeat the surging behavior, but if there is no additional plume information for a set amount of time (CastTime), the robot will declare the plume lost and return to the plume finding behavior (with a wider, less local, spiral gap parameter).

Source declaration can be accomplished using the fact that a robot performing the plume traversal behavior at the head of a plume will tend to surge into an area where there is no plume information, and then spiral back to the origin of the surge before receiving another odor hit. If the robot keeps track internally of the post spiral inter-hit distances (using odometry, for example, which is sufficient because information must be accurate only locally), a series of small differences can indicate that the robot has ceased progress up the plume, and must therefore be at the source. However, because small inter-hit distances can occur in all parts of the plume, this method is not foolproof, and tuning the significance threshold (SrcDecThresh), as well as the number of observed occurrences before source declaration (SrcDecCount), is required to obtain a particular performance within a given plume. See Table I for a summary of individual SS parameters.

SS uses only binary odor information generated from a single plume sensor. This is motivated partially because this is the most simple and reliable type of information that can be obtained from real hardware. However, due to the highly stochastic nature of turbulent fluid flow and the odor-packet nature of the plume, it is unclear that more complex sensing -- via graded intensity information or larger sensor arrays -- would benefit an agent when flow information is available through other means.

2.3 Collaborative Spiral Surge

One way to increase the performance of a robot swarm is collaboration. In particular, if collaboration is obtained with simple explicit communication schemes such as binary signaling, the team performance can be enhanced without losing autonomy or significantly increasing complexity at the individual level. Several simple types of communication can be integrated into the basic SS. Though this issue is not explored in this paper, the effects of communication strategies can change depending on the environment, so communication type should be a tunable system parameter.

2.4 Plume Traversal

This paper will focus on the plume traversal subtask because it contains most of the plume related complexity present in the full odor localization task, and due to experimental limitations it is not feasible to study all phases with real robots at this time. To study plume traversal, we place groups of agents within a starting area at the distal end of an odor plume in an enclosed arena. Over repeated trials we measure the time and distance traveled by the whole group until the first agent comes within a given radius of the plume source (T_{sf} , D_{sf}).

To justify the high density of agents in the plume (which would be unlikely given that in the general problem the plume area is a small percentage of the total search area), we allow communication between the agents that causes all downwind agents (locally determined from previous individual measurement and odometry) to surge toward an agent that has received an odor hit and is initiating its own surge behavior. This provides an attractive force that holds the group together as it traverses the plume.

Efficiency for the plume traversal task cannot be defined in the general case. Instead, there are two basic measures of task performance: time and group energy (which can be considered proportional to the sum of the individual distances traveled). Since these measures are physically independent, a composite metric incorporating a particular weighting of these two basic factors can be considered.

$$P = \frac{2}{\left[\frac{E(T_{sf})}{T_{min}} \right]^\alpha + \left[\frac{E(D_{sf})}{D_{min}} \right]^\beta}$$

This metric is an arbitrary weighting of time and distance, which are normalized by the optimum values for the given task (T_{min} , D_{min}). The form ensures that for any exponent α and β greater than 0, the optimal system will achieve a performance of 1, and any that require more time or distance will have a performance less than 1. By choosing specific values for α and β , the appropriate relationship can be generated for evaluating any particular application.

3 Materials and Methods

3.1 Real Robots

We use Moorebots, as shown in Figure 2. The plume traversal arena is 6.7 by 6.7 m, and the robots are 24 cm in diameter. In addition to the basic setup, as described in [17], each robot is equipped with four infra-red range sensors for collision avoidance, a single odor sensor tuned to sense water vapor, and a hot wire anemometer.



Fig. 2. Moorebots in plume traversal arena.

The odor sensor detects the presence of an airborne substance through a change in the electrical resistance of a chemically sensitive carbon-doped polymer resistor [18]. We generate a water plume using a pan of hot water and an array of fans. Mapping the plume using a random walk behavior (see Figure 3a) indicates that the plume is stable.




Fig. 3. (a) Plume hits received by 6 real robots over 1 hour while performing a random walk behavior. (b) Plume hits received by 6 simulated robots over 1 hour.

The anemometer is enclosed in a tube which gives it unidirectional sensitivity, which, combined with a scanning behavior, allows the robot to measure wind direction. A wind map of 2102 individual samples averaged spatially is shown in Figure 4a.




Fig. 4. (a) Average wind direction in plume traversal arena as measured by the real robots. Plume source at upper right. Arrow lengths are proportional to the mean flow magnitude at the tail of each arrow. (b) Webots plume traversal arena with average plume intensity map.

An overhead camera tracking system, combined with a radio LAN among the robots and an external workstation, is used to log position data during the trials, reposition the robots between trials, and emulate the binary communication signals. Trials of different group size are interleaved and inactive robots are automatically positioned at recharging stations.

3.2 Inherent Odor Localization Task Complexity

When studying the performance of distributed robotic systems, it can be useful to model the system using different levels of abstraction. Probabilistic analytic models are ideal, but it can be difficult to formalize all relevant local interactions at the macroscopic level. Less abstract model types include probabilistic numerical models (microscopic-level), non-embodied point simulations, and finally embodied simulations. Successful modeling provides a way of understanding the essential aspects of the system, as well as a significantly decreased evaluation time, which allows a more complete investigation of the system parameter space.

In order to demonstrate SS as an odor localizing strategy, we attempted to apply the numerical probabilistic modeling methodology described in [4]. However, we were unsuccessful because that framework is not able to capture the influence of agent trajectory across different functional states. In the previously studied exploration task, agent trajectories were randomized via wall avoidance between state transitions, so the assumptions of the model (that position and heading within each state are random) were approximately correct. In the odor localization task, transitions between areas where plume information is available to areas where there is none do not require an intermediate avoidance procedure. Thus the random position and heading assumptions of the modeling methodology do not hold, and it cannot be successfully applied. Note that it may yet be possible to develop a more sophisticated model that properly incorporates all aspects of the algorithm and dynamics of the environment.

The next lower level of investigation is non-embodied point simulation. Again, we attempted to evaluate SS at

this level, but we found that the source declaration aspect of the algorithm, a sub-task in which agent density can be elevated around the source, is very sensitive to the inter-agent repulsion parameters. Since these are intended only to approximate the behavior of the real robots, we could not hope to obtain accurate performance information using non-embodied simulation.

3.3 Embodied Simulation

In absence of a functional higher level alternative, we used Webots [19], a 3D sensor-based, kinematic simulator, originally developed for Khepera robots [20], to systematically investigate the performance of SS in simulation. This embodied simulator has previously been shown to generate data that closely matches real Khepera [5,2] and Moorebot [4] experiments, so we were confident that real robot behavior was accurately captured.

The physical arena was captured in Webots, as shown in Figure 4b. To properly capture the plume stimulus, we incorporated a series of leaky source 2D plume images generated in a water flume by Philip Roberts and Donald Webster at Georgia Tech. Such 'plume movies', even though they do not capture the influence of the agents on plume dynamics, offer a good approximation to the discretized (packet-like) nature of odor stimulus received in real environments. We scaled the recorded plume data to imitate the average speed and envelope of the real plume data (see Figure 3a and Figure 3b), and tuned the odor sensitivity threshold (higher threshold leads to less odor information) based on performance observed in our real arena. Odor hit frequency differences between the real and simulated maps are due to different polling rates of the respective measurement systems and differences in response bandwidth of the real and simulated sensors. Flow information was taken directly from the real robot data (as shown in Figure 4a) and introduced into the embodied simulations.

4 Results and Discussion

4.1 Real Robots

We tested real robot plume traversal performance using two sets of SS parameters and two control experiments. We varied only `SpiralGap2` and `StepSize` because we considered only the plume traversal aspect of the task. Parameter set SS1 represents a non-local search in that its search paths are straight and its surges extend to the boundaries of the arena. SS2 uses a smaller spiral gap and surge length to perform a more local exploration of the arena. Random Odor uses SS2 parameters, and receives odor hits that are generated from the time sequence of SS2 odor hits but are not correlated with robot position in the arena. This control experiment investigates whether an algorithm incorporating precise odor packet location information is more efficient than a blind upwind surging behavior. Random Walk takes straight line paths and random avoidance turns at boundaries (using no odor or

TABLE II
Plume Traversal Parameter Values

Agent Speed	.325 m/s
Plume Length	9 m
Plume Speed	1 m/s
Source Dec. Radius	.88 m
Plume:Search Area	1:2.3
Goal:Search Perimeter	1:18.0
α, β	1
T_{min}	19.0 s
D_{min}	6.2 m
SS1: SpiralGap2	1785 km
SS1: StepSize	9.1 m
SS2: SpiralGap2	.357 m
SS2: StepSize	.91 m




Fig. 5. Normalized time to find the source across group size for real robot trials. Lower values are better.

flow information) to provide a traversal performance baseline. Specific parameters relating to the real robot tests are listed in Table II. 15 trials of each group size were run for SS1, SS2 and Random Odor, and 30 trials were run for Random Walk due to the high performance variance.

Figures 5 and 6 show that for all conditions studied, traversal time decreases with group size while group distance traveled increases. Time and distance are normalized to the minimum values possible for this task.

Figure 7 shows that while single robots are generally most efficient in this arena, SS1 gives the best results for each group size, demonstrating successful plume tracing. Random Odor performs worse than SS2 for all group sizes, indicating that location of odor information is an important aspect of the search algorithm. Also, SS2 performs worse than SS1, suggesting that local search is not a good strategy in this small arena where the goal-to-search perimeter ratio is high (i.e., it is likely to find the goal by chance). Note that as α and β change, giving more weight to time or energy in the performance function, the values in Figure 7 will tend toward the inverse of the data shown in Figure 5 or Figure 6. In other words, as time becomes more important than energy consumed, larger group sizes




Fig. 6. Normalized distance to find the source across group size for real robot trials. Lower values are better.




Fig. 7. Performance of real robot (RR) and Webots trials across group size. Higher values indicate better performance.

will become more efficient, and vice versa. All error bars in the plots represent standard error.

4.2 Webots

We successfully reproduced the real robot performance data in Webots, as shown in Figure 7. Data represents 1000 trials per group size. All parameters in Table II apply to the Webots data as well.

Because our Webots data closely matches our available real robot data, it is reasonable that further simulated experiments will accurately reflect real world behavior. The main limitations to our real robot experiments thus far are arena size and restriction to the plume traversal subtask, thus in simulation we ran a set of trials involving both the plume finding and plume traversal subtasks in a 25x (area) larger arena. The simulated plume remained the same, the start area remained the same size but was moved out of the plume to a corner of the arena, and both SS algorithms used a SpiralGap1 of 1785 km (producing straight line plume finding search paths) and a CastTime of 96 s. Figure 8 shows that in the larger arena the local search of SS2 is the best strategy. Single robots are no longer the most efficient because the penalty for




Fig. 8. Performance on the plume finding and traversal tasks of Webots trials across group size in larger arena. Higher values indicate better performance.

losing contact with the plume is high. While larger group sizes ensure that the plume is never lost, they also bring higher interference and search overlap as well. Optimal balance for this environment and parameter set is at a group size of 4 for SS2. SS1 performs worse because its non-local search has a higher likelihood of losing the plume across all group sizes. Random Walk performance decreases most drastically, as the probability of encountering the goal by chance is highly dependent on the goal-to-search perimeter ratio. Note that the SS1 and the Random Walk performance curves have optimal values like SS2, but they occur at group sizes above 10.

5 Conclusion

In this paper we have described a distributed algorithm for solving the full odor localization task, and shown that group performance can exceed that of a single robot. We have demonstrated that one subtask, plume traversal, can be successfully accomplished by real robots. Furthermore, we have established that an embodied simulator can accurately replicate the real robots results, and shown that it can be a useful tool for exploring system performance.

Achievement of near optimal performance on the full odor localization task in the real world will require efficient search of a large parameter space, which will call for the combination of accurate simulation and machine-learning techniques.

Acknowledgements

We would like to thank O. Holland, S. Kazadi, J. Pugh, R. Enright, L. Van Tol, and A. Lundsten for contributions on various hardware and software aspects of the project. We would also like to thank our collaborators in the DARPA-ONR Chemical Plume Tracing Program for their valuable input. This work is supported in part by the Center for Neuromorphic Systems Engineering under grant EEC-9402726, by DARPA under grant DAAK60-97-K-9503, by the ONR under grant N00014-98-1-0821, and by the ARO under MURI grant DAAG55-98-1-0266. A. Hayes is supported by a NSF Graduate Research Fellowship.

References

- [1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, New York, US, 1999.
- [2] A. Martinoli, A. J. Ijspeert, and F. Mondada, "Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots," *Robotic and Autonomous Systems*, vol. 29, pp. 51-63, 1999.
- [3] O. E. Holland and C. Melhuish, "Stigmergy, self-organization, and sorting in collective robotics," *Artif. Life*, vol. 5, pp. 173-202, 1999.
- [4] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Comparing distributed exploration strategies with simulated and real autonomous robots," in *Proc. of the fifth Int. Symp. on Distributed Autonomous Robotic Systems DARS-2000*, L. E. Parker, G. Bekey, and J. Barhen, Eds., Knoxville, Tennessee, October 2000, pp. 261-270, Springer Verlag.
- [5] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella, "Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment," *Autonomous Robots*, vol. 11, no. 3, pp. 149-171, 2001.
- [6] C. R. Kube and E. Bonabeau, "Cooperative transport by ants and robots," *Robotics and Autonomous Systems*, vol. 30, pp. 85-101, 2000.
- [7] R. T. Carde and A. Mafra-Neto, "Effect of pheromone plume structure on moth orientation to pheromone," in *Perspectives on Insect Pheromones. New Frontiers*, R. T. Carde and A. K. Minks, Eds., pp. 275-290. Chapman and Hall, N.Y., 1996.
- [8] J. H. Belanger and M. A. Willis, "Adaptive control of odor guided locomotion: Behavioral flexibility as an antidote to environmental unpredictability," *Adaptive Behavior*, vol. 4, pp. 217-253, 1996.
- [9] U. Bhalla and J. M. Bower, "Multi-day recording from olfactory bulb neurons in awake freely moving rats: Spatial and temporally organized variability in odorant response properties," *J. of Computational Neuroscience*, vol. 4, pp. 221-256, 1997.
- [10] J. Atema, "Eddy chemotaxis and odor landscapes: Exploration of nature with animal sensors," *Bio. Bull.*, vol. 191, pp. 129-138, 1996.
- [11] M. J. Weissburg, "From odor trails to vortex streets: Chemo and mechanosensory orientation in turbulent and laminar flows," in *Orientation and Communication in Arthropods*, M. Lehrer, Ed. Birkhauser, Basel, 1997.
- [12] T. Nakamoto, H. Ishida, and T. Moriizumi, "A sensing system for odor plumes," *Analytical Chemistry*, vol. 71, no. 15, pp. 531A-537A, 1999.
- [13] S. Kazadi, R. Goodman, D. Tsikata, and H. Lin, "An autonomous water vapor plume tracking robot using passive resistive polymer sensors," *Autonomous Robots*, vol. 9, no. 2, pp. 175-188, 2000.
- [14] C. T. David, J. S. Kennedy, J. S. Ludlow, and J. N. Perry, "A reappraisal of insect flight towards a point source of wind-borne odor," *J. of Chemical Ecology*, vol. 8, pp. 1207-1215, 1982.
- [15] N. J. Vickers and T. C. Baker, "Reiterative responses to single strands of odor promote sustained upwind flight and odor source location by moths," *Proc. of the Nat. Acad. of Sciences USA*, vol. 91, pp. 5756-5760, 1994.
- [16] D.W. Gage, "Randomized search strategies with imperfect sensors," in *Proc. of SPIE Mobile Robots VIII*, Boston, September 1993, vol. 2058, pp. 270-279.
- [17] A. F. T. Winfield and O. E. Holland, "The application of wireless local area network technology to the control of mobile robots," *Microprocessors and Microsystems*, vol. 23, pp. 597-607, 2000.
- [18] M. S. Freund and N. S. Lewis, "A chemically diverse conducting polymer-based electronic nose," *Proc. of the Nat. Acad. of Sciences USA*, vol. 92, pp. 2652, 1995.
- [19] O. Michel, "Webots: Symbiosis between virtual and real mobile robots," in *Proc. of the First Int. Conf. on Virtual Worlds, VW'98*, Paris, France, July 1998, pp. 254-263, Springer Verlag.
- [20] F. Mondada, E. Franzi, and P. lenne, "Mobile robot miniaturization: A tool for investigation in control algorithms," in *Proc. of the Third Int. Symp. on experimental Robotics ISER-93*, T. Yoshikawa and F. Miyazaki, Eds., Kyoto, Japan, 1993, pp. 501-513, Springer Verlag.