Adaptive Behavior

Tracking of Fluid-Advected Odor Plumes: Strategies Inspired by Insect Orientation to Pheromone

Wei Li[1,2], Jay A. Farrell[1], Ring T. Cardé[3]
[1]*Department of Electrical Engineering, University of California, Riverside*
[2]*Department of Computer Science, California State University, Bakersfield*
[3]*Department of Entomology, University of California, Riverside*

Autonomous vehicles with plume-tracing capabilities would be valuable for finding chemical sources in fluid flows. This article considers strategies allowing autonomous vehicles to find and trace an odor plume to its source. These strategies are inspired by the maneuvers of moths flying upwind along a pheromone plume. Although moth maneuvers are well documented, the mechanisms underlying sensory perception and navigation are not fully understood; therefore, a key objective was to define sensor, signal-processing, and actuation algorithms for autonomous vehicles. The strategies presented do not precisely mimic insect orientation to odors. Optimizing performance, however, suggests orientation strategies that may have biological counterparts. The results demonstrate the importance of cross-plume counterturning strategies for maintaining intermittent contact with the chemical plume, given noisy sensory information. It is important for the searcher to maintain intermittent contact with the plume because flow direction while detecting odor is the main indicator of the instantaneous desired direction of motion.

Keywords robotics plume tracking · insect behavior · biomimetic behavior · odor plume · attraction

1 Introduction

Olfactory-based mechanisms have been hypothesized for a variety of biological behaviors (Dusenbery, 1992; Vickers, 2000; Zimmer & Butman, 2000): homing by Pacific salmon (Hasler & Scholz, 1983), foraging by Antarctic procellariiform seabirds (Nevitt, 2000), foraging by lobsters (Atema, 1995; Basil & Atema, 1994; Devine and Atema, 1982), foraging by blue crabs (Weissburg & Zimmer-Faust, 1994), and mate seeking and foraging by insects (Cardé, 1996). Typically, olfactory-based mechanisms proposed for biological entities combine a large-scale orientation behavior based in part on olfaction with a multisensor local search in the vicinity of the source. The long- range olfactory-based search is documented in moths at ranges of 100–1,000 m (Rau and Rau, 1929; Elkinton, Schal, Ono, & Cardé, 1987) and in Antarctic procellariiform seabirds over thousands of kilometers (Nevitt, 2000).

Recently there has been interest in the development of algorithms to replicate these feats on autonomous vehicles (e.g., Larcombe & Helsall, 1984; Kuwana, Nagasawa, Shimoyama, & Kanzaki, 1999; Ishida, Nakamoto, Moriizumi, Kikas, & Janata, 2001). The goal of the autonomous vehicle would be to locate the source of a chemical that is transported in a turbulent fluid flow. Such autonomous vehicle capabilities have applicability in searching for environmentally interesting phenomena such as unexploded ordinance, undersea wreckage, and sources of hazardous chemicals or pollutants. In this article, we develop strategies, inspired by moth plume-following behaviors, that allow an autonomous vehicle to trace a chemical plume from a possible large initial distance to the odor source.

Correspondence to: J. Farrell, Department of Electrical Engineering, University of California at Riverside, Riverside, CA 92521, USA.
E-mail: j.a.farrell@ieee.org
Tel.: +1-909-7872159, *Fax*: +1-909-7872425

(2001), Vol 9(3–4): 143–170
[1059–7123 (200110) 9:3-4; 143–170; 028765]

The location of pheromone-emitting females by flying male moths is considered a remarkable case of odor-guided navigation. Upwind flight along pheromone plumes is guided by two complementary mechanisms. First, direction with respect to wind flow is gauged by the optomotor reaction: Feedback from movement in the visual surround provides information about the wind's direction and the organism's progress upwind. This maneuver is termed odor-induced optomotor anemotaxis. Second, the precise upwind heading and velocity seem to be governed mainly by moment-to-moment contact with small-scale filaments of pheromone. Filament encounter rates above 5 Hz promote rapid flight with a mainly due-upwind heading (Mafra-Neto & Cardé, 1994; Vickers & Baker, 1994a). When the filament encounter rate falls below 5 Hz, males tend to fly a slower zigzag course, counterturning across the wind with a reduced rate of upwind displacement.

Male moths that lose contact with the plume "cast," that is, they cease upwind movement and their crosswind excursions progressively widen (Kuenen & Cardé, 1994). Casting can continue for several seconds until contact with the plume is reestablished. If moths fail to reenter the plume, they may return to "ranging behaviors" used for initial location of a plume (Elkinton & Cardé, 1983). There are two principal reasons that upwind flight within the plume can lead to loss of contact with odorant (Murlis, Elkinton, & Cardé, 1992). First, the internal structure of plumes is patchy, that is, in certain locations the signal is not present or it is below the threshold of detection, two effects that are magnified as the plume is carried downwind. Second, the upwind direction may lead out of the plume due to the meander of the plume centerline. In this article, the three types of behavior described above are referred to as plume finding (ranging), plume maintaining (upwind orientation), and plume reacquisition (casting).

Strategies of walking and flying organisms following a fluid-advected plume differ in one important respect. In walking organisms, the direction of fluid flow is sensed by mechanoreceptors (e.g., comparative deflection of the two antennae), with reference to the organism's position on the walking substrate (Bell, 1984). The straightforward mechanism for detection of fluid directionality used by walking insects is not available to flying insects but could be readily employed by a ground-based autonomous vehicle.

Because the first male moth to locate a female typically secures a mating, there is considerable selective pressure for males to be efficient at locating a plume of pheromone and then quickly to navigate a course along the plume to the female (Cardé, 1986). Therefore, a biological system's plume-tracing strategy should be optimized, given the sensor and mobility constraints of the biological system. It should also be robust to typical variations in environmental conditions. The presumed optimality and robustness of biologically based plume-tracing strategies are valuable traits to mimic in an autonomous vehicle designed to locate a source of an odor (e.g., a source of environment contamination) via plume tracing.

Much of what we know about plume-following maneuvers of flying insects is based on studies of moths flying along odor plumes generated in wind tunnels. In the natural environment, where the fluid flow is more complex, systematic behavioral studies are difficult to perform. Translation of the observed insect behavior into an algorithm for plume tracing suitable for autonomous vehicle applications is still in its infancy.

Robotic tracking of chemical trails is discussed in for example, the work of Deveza and colleagues (Deveza, Thiel, Russell, & Mackay-Sim, 1994). Tracing of chemical plumes is studied in several articles (e.g., Belanger & Willis, 1996a, 1996b, 1998; Belanger & Arbas, 1998; Grasso, Consi, Mountain, & Atema, 2000; Ishida, H., Kawaga, Y., Nakamoto, T., & Moriizumi, 1996; Ishida, H., Nakamoto, T., & Moriizuma, T., 1998; Ishida, H., Kobayashi, A., Nakamoto, T., & Moriizumi, T., 1999; Kuwana et al., 1999; Rozas, Morales, & Vega, 1991; Russell, Thiel, Devesa, & Mackay-Sim, 1995; Russell, 2001). Belanger and Willis (1996a, 1996b, 1998) and Belanger and Arbas (1998) follow an approach similar to that used herein, implementing proposed strategies as software algorithms (see the appendix of Belanger & Arbas, 1998) that can be evaluated in computer-based simulations. The main advantages of this approach are that alternative strategies can be evaluated under identical experimental conditions and the causes of the searcher maneuvers (i.e., central nervous system processing of living entities) are known. The main disadvantage, relative to experimental evaluation, is that the validity of the results is a function of the extent to which the characteristics of the environment are accurately modeled. The simulated environment of

Belanger and Willis (1996a, 1996b, 1998) and Belanger and Arbas (1998) is similar to a small wind tunnel. The centerline of the plume aligns with the mean wind vector and points directly back to the source. Three strategies are proposed in Belanger and Arbas (1998). These will be further discussed in Section 6.

Grasso et al. (2000) present and analyze chemotaxis-based (concentration information only) algorithms using two concentration sensors mounted on a robot lobster. The article includes results of experiments in a flume where their robo-lobster initiated its search 1.0 m or less from the source in the downstream center of a well-developed plume. Grasso et al. (2000) discuss some of the experimental difficulties that result from the intermittency due to a turbulent flow. Because the experiments occur in a flume they are not able to analyze the challenges imposed by plumes with significant meander.

Ishida et al. (1996) present on-vehicle experimental results that use an array of sensors calculating a cross-plume concentration gradient to track the plume centerline. Ishida et al. (1998) propose a method to estimate the source location remotely by using a Kalman filter and nonlinear optimization methods to estimate parameters of a Gaussian plume-distribution model. It is interesting, in comparison with the counterturning algorithms suggested herein, that accurate estimation of the desired parameters by this method requires the use of cross-plume Z-shaped paths. Although not directly comparable to the results herein, the algorithms of Ishida et al. (1998) required on the order of 300 s to trace a 2-m plume. This slow rate of travel was required due to slow sensor characteristics and the need to calculate statistical properties and gradient information. The approaches presented herein can respond more quickly because their decisions are based on instantaneous concentration measurements. The remote estimation approach in Ishida et al. (1998) is similar to that suggested in Stacey, Cowen, Powell, Dobbins, Monismith, & Koseff (2000). Ishida et al. (1999) present a multisensor methodology for estimating the three-dimensional direction to the odor source. The method assumes that the local time-averaged concentration gradient at the vehicle location points toward the odor source. This assumption is only true in the vicinity (a few meters) of the odor source; at greater distances slow changes in the direction of the wind will cause the plume centerline to meander (see, for example, Figure 2). Also, at large distances from the source detection of the plume gradient might require a large intersensor separation.

Kuwana et al. (1999) present experimental results for a miniature robot instrumented with silkworm antennae and using a chemotaxis strategy tracing a 10-cm pheromone plume in a wind tunnel. The focus of this article was the use of antennae as pheromone sensors. Rozas et al. (1991) present an electronic nose design, sensor-processing algorithms, and gradient-following strategies. Russell et al. (1995) include robotic implementation of algorithms that estimate statistics of the plume such as the plume centroid. The strategy then maneuvers relative to the plume centroid in an upwind direction. Such strategies are not expected to work well in situations where the flow can change significantly, resulting in plume meander, because the centroid would then be time varying. Russell (2001) also includes robotic implementation experiments where the chemical is constrained to a multiple-duct tunnel system. The approach focuses on selecting the correct tunnels to follow to locate the odor source.

In this article, we investigate the problem of maintaining contact with a plume once it has been detected. In contrast to the articles cited above that studied plume tracing near the source where the wind direction and plume axis are aligned, we focus on the development of effective algorithms for tracking the plume at distances from the source where it is possible that neither the plume centerline nor the wind direction point to the source. The article is organized as follows. Section 2 reviews key issues from the topic of chemical dispersion in turbulent media and summarizes the plume model used in the Monte Carlo simulation studies of this article. Section 3 proposes a strategy to maintain contact with the plume, presents the variables that will be used to quantify performance, and evaluates how changes in the key strategy parameters affect performance. Section 4 proposes and optimizes a strategy with improved ability to maintain contact with the intermittent plume. The optimized strategies that result from Section 4 yield generally zigzag-shaped ground tracks similar in form to those exhibited by a range of biological entities during plume-tracking. Section 5 evaluates plume-tracking performance using fluid flow and concentration data obtained from flume experiments. Section 6 compares virtual plume-tracking strategies with strategies used by insects and

other organisms. The Appendix A includes a table defining various notations used in the body of the article.

2 Plumes and Plume-Maintaining Behavior

Moth plume-tracing behavior can be used as a basis for designing an autonomous system to search for the unknown location of an odor source; however, the sensitivity of search performance to various environmental factors and strategy parameters must be investigated before accepting such biologically based search strategies as appropriate for autonomous systems. For example, we should examine how we can alter strategy parameters to decrease the time to find the plume, increase the time in contact with the plume, and reduce the time to reacquire the plume when contact is lost. This article focuses on the problem of maximizing the time in contact with the plume while making progress toward the source and being robust to changes in environmental conditions. Monte Carlo simulation methods are used for performance evaluation. Both simulated air plumes and experimental water flume data were used for the performance evaluations. The simulated plumes are more difficult to track than the plumes in the flume data sets due to the significant meander of the plume centerline that can be included in the simulations by temporal and spatial variation of the wind direction and speed. Therefore, the simulated plumes are used throughout the development portions of this study. Important aspects of diffusion in a turbulent media are summarized in Section 2.1. Section 2.2 summarizes important aspects of the plume simulation implemented for these studies. Experimental flume data are used for the performance evaluations presented in Section 5 to ensure that the strategy performance has not inadvertently been optimized to the characteristics of the simulation.

2.1 Diffusion in a Turbulent Media

Pheromone is transported downwind by advection. The pheromone is dispersed mainly by the forces of turbulence while it is transported downwind in turbulent eddies, which range upward in scale from centimeters. Individual filaments of pheromone may be transported relatively undiluted for many meters

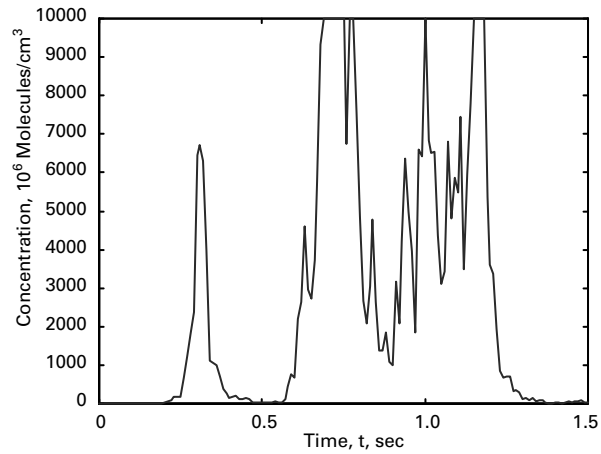

Figure 1 Concentration time series with threshold $\tau = 2{,}000$.

before being torn apart and mixed with clean air by small-scale turbulence (Murlis & Jones, 1981; Murlis et al., 1992; Murlis, Willis, & Cardé, 2000).

A typical time series of concentration at a fixed downwind location is shown in Figure 1. The turbulent nature of the transport mechanism causes the sensed concentration to be only intermittently above a positive sensor threshold τ. Such intermittency has several consequences. First, algorithms based on instantaneous gradient-following in an intermittent plume would be extremely inefficient. The instantaneous gradient varies too rapidly to be calculated. The gradient of a *mean* concentration theoretically could be traced; however, plume meander, the required spatial sampling, and the requirement of a long averaging time prohibit on-vehicle calculation of the mean concentration function and, therefore, gradient-search methods. Second, intermittency causes the concepts of "in the plume" and "not in the plume" to be difficult to define, as discussed in the following paragraphs.

Conventional time-averaged concentration models (e.g., Sutton, 1953) are time independent. Such models assume that the plume was fully developed prior to the initiation of data collection. Concentration distributions (assuming stationarity and ergodicity) of such models allow time-independent predictions, *over an ensemble* of plumes, of the average concentration as a function of position. However, measured *instantaneous* concentrations of specific exemplar plumes deviate significantly from the ensemble average (Aylor, 1976; Aylor, Parlange, & Granett, 1976; Jones,

1983; Mylne, 1992). This deviation is the result of the intermittency and of the plume meander. Because flying moths are making behavioral decisions using odor concentrations sensed over fractions of a second, whereas time-averaged models calculate concentrations over a span of several minutes, time-averaged concentration models are not useful for behavioral studies (Elkinton, Cardé, & Mason, 1984). Time-averaged concentrations also are not useful as a basis for maneuver design. A fleet of searchers would be required to implement the spatial sampling, and the time averaging would prohibit rapid plume tracing. We instead pursue an approach where an individual searcher tracks a plume by reacting to the sensed instantaneous concentration.

For the strategies developed herein, a single search vehicle with a wind direction and a binary concentration sensor is assumed. The sensed concentration at time t is denoted $c(t)$. The searcher will consider itself to be instantaneously "in the plume" at time t if it has detected above-threshold odor within the last T_w seconds (i.e., the searcher is in the plume at time t if $c(t - T) > \tau$, for some T such that $0 \le T < T_w$).[1]

The studies to follow use a plume simulation model that generates individual realizations of odor plumes.[2] The simulated plume concentration can be evaluated at *any* location with a time resolution of 0.01 s. Such simulated plume models enable certain types of analysis that would be precluded by the loss of time-signature information in mean concentration plume models and allows performance analysis and optimization under repeatable conditions that are not possible with plumes in natural environments [the tradeoffs between simulation and flight testing are discussed by Belanger and Arbas (1998)].

2.2 Plume Simulation Model

To generate realistic simulated plumes we considered the following three characteristic structural features to be of primary concern. First, the sensed plume at a fixed location should have an intermittent internal structure that closely duplicates experimental fine time-scale observations. Second, the plume shape that results from a semicontinuous release of odor should be sinuous and time varying. Third, the plume shape and fluid flow history should be coherent.

The first characteristic feature is achieved by representing the odor plume as a sequence of puffs. Puffs are released sequentially at the source location. Each puff is composed of n filaments. After release, each filament center location is determined as a function of time by integration of the advective fluid flow velocity. To achieve the second characteristic feature, the simulation incorporates a dynamic model of the advective fluid flow that is the solution of a K-closure-based partial differential equation (Pielke, 1984) driven by boundary conditions that are implemented as colored noise processes (p. 123 in Jazwinski, 1970). As the boundary conditions change, the partial differential equation is solved numerically over the region of interest. This model allows the fluid flow, which advects the filaments, to be a continuous but varying function of position and time. Changes in fluid flow direction cause the plume to meander (form a snake-like path when viewed from above). Because the fluid flow direction and magnitude change, spatially and temporally, the instantaneous fluid flow direction within the plume often will not point toward the odor's source nor be coincident with the plume's centerline. The meander and lack of correlation between the wind direction and the plume long axis in the simulated plumes reproduce similar features in odor plumes observed in the field (Elkinton et al., 1987; Brady, Gibson, & Packer, 1989) that significantly complicate the plume-tracking process.

A detailed description of this filament-based odor dispersion model is contained in Farrell et al.'s work (Farrell, Murlis, Long, Li, & Cardé, 2002). The resulting plumes duplicate both the time-averaged and instantaneous features of plume structures measured in the field (Jones, 1983) without sacrificing computational feasibility.

Related plume models based on the advection of puffs have been suggested in previous work. In Gifford's (1959) work, puffs are represented as two-dimensional (i.e., planar) disks, but the advection model was not specified. The model of Belanger and Arbas (1998) represents the puffs as homogeneous patches and appears to assume a constant unidirectional fluid flow model. The fluid flow model used in this study represents a significant extension to such puff-based plume models.

The filament-based model, presented in detail in Farrell et al. (2002), is general and is designed so that it can be tuned for a range of applications. For the simulation results presented here, the simulation parameters were set to correspond to odor plumes in

wind as perceived by male moths. The receptor system of some male moths is capable of resolving 10 odor pulses per second (Rumbo and Kaissling, 1989). A flying male can react to odor filaments within 200 ms (Mafra-Neto and Cardé, 1994, 1996). Therefore, we designed the present model of plume dispersion with an integration time step of 10 ms. Similarly, the units of the sensed concentration were specified relative to the measured threshold of the male gypsy moth (*Lymantria dispar*) (Elkinton et al., 1984) and the rate of emission at the mean rate of pheromone release by the female (Charlton and Cardé, 1982).

A graphic "snapshot" illustrating the nature of the plume that results from this simulation model is contained in Figure 2. This figure shows a vehicle, with heading ψ, searching for the plume. The search area is a rectangle with corners: $(\underline{x},\underline{y})$, $(\underline{x},\overline{y})$, $(\overline{x},y)$, and $(\overline{x},\overline{y})$. For Figure 2, $\underline{x} = 0$, $\underline{y} = -50$, $\overline{x} = 100$, and $\overline{y} = 50$ m. The grayscale indicates the local odor concentration as a function of position. Meander of the plume is evident clearly as is dispersion relative to the meandering plume centerline. The array of small arrows on the figure indicates the direction of the local fluid flow vector at the tail of each arrow. Note that the fluid flow vector varies spatially and that the correlation of the fluid flow direction and plume axis direction decreases with the distance from the source. The source location is a simulation variable. For the simulation studies in this article, the source will be located at (20,0; Figure 2) to allow analysis of strategy performance when the searcher starts upwind from the source.

The sensed concentration is modeled as a low-pass-filtered version of the instantaneous concentration at the sensor location, followed by a threshold setting operation:

$$\frac{\mathrm{d}\rho(t)}{\mathrm{d}t} = -\alpha\rho(t) + \alpha C(x_s(t)) \tag{1}$$

$$c(t) = \begin{cases} \rho(t) & \text{if} \quad \rho(t) > \tau \\ 0 \end{cases} \tag{2}$$

where α is the filter bandwidth, $\rho(t)$ is an internal state of the filter, $x_s(t)$ is the time-varying sensor position. The filter input is the instantaneous concentration at the sensor location $C(x_s(t))$. The filter (i.e., sensor) output is $c(t)$. For the simulation results of this article, the sensor threshold was 0.01; the filament release rate was 3 filaments s^{-1}; the simulation time step was 0.01 s; and the mean fluid velocity was 1 m s^{-1}.

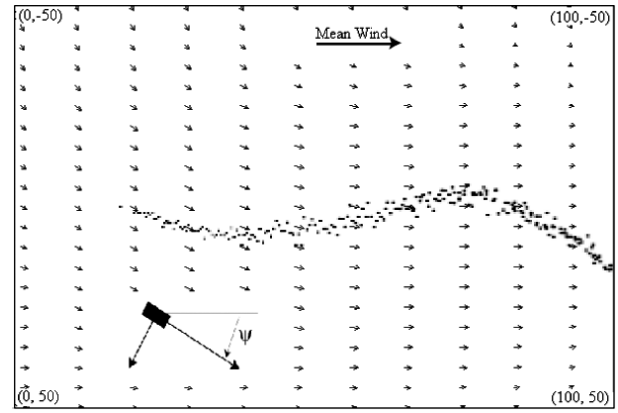

Figure 2 Vechicle (not to scale) and plume within the simulated search area. The *x* coordinate is between $\underline{x} = 0$ and $\overline{x} = 100$ m and is positive in the mean downwind (i.e., horizontal) direction. The *y* coordinate is between $\underline{y} = -50$ and $\overline{y} = 50$ m and is positive down. Heading is measured positively with respect to the *z*-axis in a right-handed coordinate system. The grayscale indicates above threshold concentration. The arrows indicate the magnitude and direction of the local wind vector at the tail of the arrow.

2.3 Fluid Flow Variation

An important issue in Monte Carlo analysis is performance sensitivity to variations in the parameters of the fluid flow model, particularly those that affect meander. Note that the simulation model produces a fluid flow field that varies continuously as a function of time and space. The model parameters determine the rate and magnitude of these changes. Therefore, the analyses to follow use two settings of the fluid flow model parameters that will be referred to as "large meander" and "small meander." The small meander parameter settings yield plumes with meander shapes typified as in Figure 2. The typical plume has instantaneous centerline lateral deviation from the mean centerline (i.e., the *x*-axis) at $y = 100$ that are normally less than 10 m. The large meander parameter settings yield plumes with instantaneous centerline lateral deviation from the mean centerline at $y = 100$ that are normally less that 25 m. Both sets of parameter settings generate plumes of realistic extent in the *y* direction. The plumes that result with the small meander parameter settings are significantly easier to track.

3 Analysis of Maintaining Behavior

The plume-tracing task is decomposed into the four components: finding the plume, maintaining (intermittent) contact with the plume, reacquiring contact with the plume, and "declaring" that the source has been found. The focus of this article is on the strategies for maintaining and reacquiring contact with the plume.

From an initial starting position, the first objective is to find the plume. The plume is found once the first above-threshold odor concentration $c(t)$ is detected. The total time spent finding the plume is denoted by T_F. The searcher has no prior information about the likely odor source location within the search area; therefore, the plume-finding task is an uninformed search problem. There are many known solutions for uninformed search problems (Dusenbery, 1992). The plume-finding strategy used throughout this article is described in the Appendix. When the area to be searched is large, finding the plume can consume a large amount of resources (i.e., time and energy).

Once the plume has been found the objective is to maintain contact with it while making progress toward the source. The total time in the plume, referred to as the maintain time, will be denoted by T_M. If the searcher loses contact with the plume, then the behavior switches to one that is likely to reacquire contact with the plume so that the searcher does not need to revert to the (resource-consuming) plume-finding behavior. The total time spent trying to reacquire the plume will be denoted by T_R. Finally, the searcher should be able to declare that the source has been found. In biological entities, the source-found declaration is often influenced by multiple sensory modalities, including visual and tactile cues. Design of algorithms to declare that the source has been found using only sensed concentration is a significant problem that, due to its complexity, will be considered elsewhere. The methodology used in this article for declaration that the source has been found is described in Appendix B.

Due to the turbulent nature of the fluid flow, the sensed concentration is an intermittent signal, as shown in Figure 1. Figure 3 depicts the situation and the parameters that will be used to define practically the concept of "in the plume." The top portion of Figure 3 is a schematic rendition of a searcher trajectory that finds the plume and subsequently makes

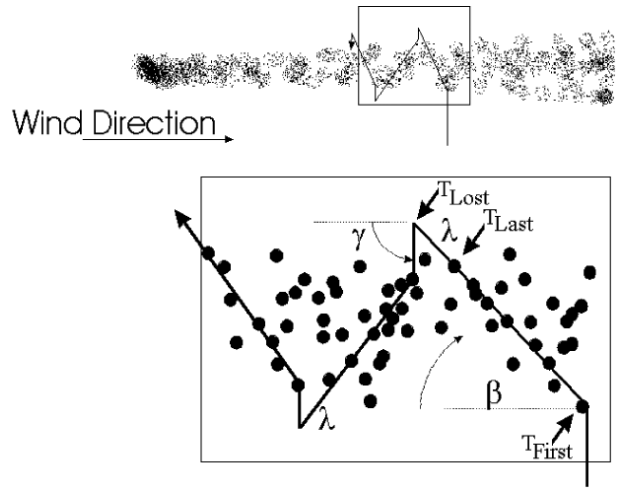

Figure 3 Illustration of the designed searcher behavior "in the plume." Black dots depict above-threshold paches of odor. After first odor detection, the searcher travels at an angle β relative to upwind. The angle β is maintained until above-threshold odor has not been detected for λ seconds, at which time the searcher counterturns to travel at an angle γ with respect to the wind. T_{First} denotes the time at which the sensed concentration is first above threshold. T_{Last} denotes the most recent time above-threshold odor was detected. T_{Last} denotes the time at which the searcher declares that contact with the plume is lost.

along-plume progress. The lower portion of the figure presents a magnified depiction of such an in-plume odor encounter scenario. The time at which the sensed concentration is first above threshold is denoted T_{First}. The searcher will be defined to be maintaining contact with the plume at time t if the sensed concentration at any time in the last λ s has been above threshold (i.e., if $c(v) > \tau$ for any $v \in [t - \lambda, t]$). Note that even while the searcher is within the time-averaged boundary of the plume, there will be periods where the sensed concentration is below threshold (see also Figure 1). With reference to Figure 3, let T_{Last} denote the most recent time that above-threshold odor was detected. There may be many above-threshold detections between T_{First} and T_{Last}. Let T_{Lost} denote $T_{Last} + \lambda$. Note that T_{Lost} is the time at which the searcher declares that contact with the plume is lost. The maintain time for this cross-plume transit will be calculated as

$$T_M = T_{Lost} - T_{First} - \lambda. \tag{3}$$

Finding the plume expends both time and energy. Therefore, once the plume is found, the searcher does

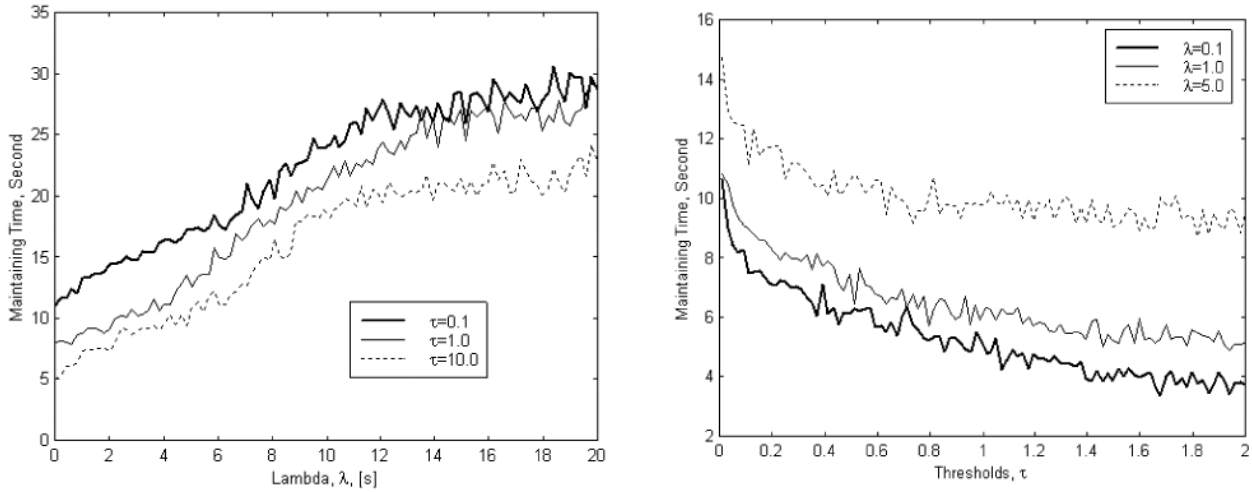

Figure 4 Sensitivity of maintaining time (T_M) to waiting time λ and threshold τ.

not want to lose (intermittent) contact with the plume. In addition, significant progress toward the source can be made once the plume has been found. Therefore, a critical question is how to design algorithms that are effective at maintaining plume contact while advancing toward the source. The nominal strategy (motivated by studies of moth behavior) for maintaining contact with the plume is to define the searcher heading[3] by

$$\psi_v(t) = \psi_u(t) + \beta(t) \quad \text{if } c(v) > \tau \text{ for some } v \in [t-\lambda, t] \tag{4}$$

$$\psi_v(t) = \psi_u(t) - \gamma(t)\ \text{sign}(\beta) \text{ if } c(v) > \tau \text{ for some } v \in [t - T_W, t-\lambda] \tag{5}$$

where T_W is a strategy parameter, $\psi_u(t) = \psi_w(t) + 180$ is the heading that instantaneously points upwind at the searcher location, $\beta(t)$ has constant magnitude with a sign that changes each time that behavior of Equation 5 occurs, and γ is a constant to be defined later. This nominal behavior is illustrated in Figure 3. If $c(v) < \tau$ for all $v \in [t - T_W, t]$), then the searcher reverts to the plume-finding behavior defined in the Appendix.

The performance level of this strategy depends on environmental factors and the parameters defining the plume-maintaining behavior. To account for environmental factors the Monte Carlo approach analyzes statistical performance based on numerous simulations. Each simulation in a Monte Carlo batch uses a distinct instantiation of the plume. Our approach to designing a plume-maintaining behavior includes evaluation of the effect of changing the parameters of the plume-maintaining strategy. Strategies will be evaluated by the duration of time that they maintain contact with the plume. Possible alternative metrics are discussed in Section 7. Increased maintain times are considered desirable, but in addition to increasing maintain time, it is necessary to make progress toward the odor source. The results in Sections 4.2 and 4.3 show that it is possible to achieve long maintain times without making significant progress toward the source. Section 4.3 presents an alternative strategy that is successful at maintaining contact with the plume while making significant progress toward the source. This improved behavioral strategy may have biological counterparts.

3.1 Maintain Time Dependency on Waiting Time (T_W) and Threshold (τ)

For the maintain strategy of Equations 4 and 5 with $\beta = 0$, the sensitivity of T_M to the parameters λ and τ is plotted in Figure 4. Each iteration of the evaluation method consisted of 100 searchers starting at $(x,y) = (85, 0)$ m. All searchers followed exactly the same path during the plume-finding phase. Once the plume was detected, each vehicle used a distinct set of λ and τ values for its maintain strategy until the plume was lost. After all searchers lost the plume, the resulting set of T_M values was saved. Then, the process was repeated, with all 100 searchers restarting at the same time instant at $(x,y) = (85,0)$ m. This process was repeated for 100 iterations. The time-varying nature of the plume causes each of the iterations to be distinct. Beginning all 100 searchers of each

iteration at the same time and location with the same plume-finding behavior causes the evaluation of their performance to occur (for that iteration) under identical conditions.

Decreasing the detection threshold results in increased T_M as shown in Figure 4. This effect is expected, because decreasing the threshold increases the number of times that odor is detected. The plots showing that T_M increases with λ are more interesting; however, the result is not unexpected, because increasing λ gives the searcher additional time to encounter odor. What Figure 4 does not reveal is that increasing λ can result in the searcher being significantly further from the plume when it decides that the plume has been lost; this increased separation complicates plume reacquisition. Note also that T_M saturates as λ increases. The results of Figure 4 were found via Monte Carlo simulation for a given set of experimental conditions. Changing the simulation parameters (experimental conditions) changes the specific values along the plots, but not the general trends.

4 Design and Optimization of Plume Maintaining

Designing a behavior to keep the searcher in contact with the plume while making upwind progress requires the design of two interacting behaviors: the "in the plume" (maintaining) behavior and the reacquisition behavior. The two behaviors are difficult to consider or optimize independently. Plume reacquisition is considered first, because those strategies are required for the subsequent development of plume-maintaining behavior. Our explicit goal is to define a strategy of plume tracing for autonomous vehicles. The process of this design has suggested principles relevant to plume following in wind by flying and walking organisms.

4.1 Plume Reacquisition

The plume-reacquisition problem is that λ seconds after the last above-threshold detection of odor, the vehicle performs a maneuver intended to recontact the odor plume. If the plume is not recontacted within a time T_W, then the vehicle reverts to plume-finding behavior. Because plume finding is expensive in terms of both time and energy, maximizing the probability that plume reacquisition succeeds is critical.

The plume-maintaining behavior involves travel at an angle β offset to upwind as indicated in Figure 3. For reacquisition, a key question is "when the plume-maintaining strategy results in the searcher losing contact with the plume, in which direction should the searcher turn to maximize the likelihood of recontacting the plume?" Although the obvious answer is to turn toward the plume, this approach cannot be directly implemented, as the plume location is unknown to the searcher. The information available to the searcher is the vehicle heading $\psi_v(t)$, wind heading $\psi_w(t)$, sensed odor concentration $c(t)$, and vehicle position $(x(t), y(t))$. We study plume reacquisition in two steps. First, we analyze recovery time as a function of the recovery angle γ. Second, we analyze the statistical relationship between measured information (the vehicle motion direction minus the wind direction) and the optimal recovery angles. The goal is to determine a means of using ψ_v and ψ_w to determine the optimal turn direction γ.

To analyze the first issue a set of simulations was run in which each iteration involved 36 vehicles, indexed by the variable i in [1,36], started at (x,y) = (80,10) m and remaining at that location until odor was detected. After detection, each searcher traveled with the heading $\psi_u(t)$ defined after Equation 5. Each searcher maintains this direction of travel until contact with the plume is lost. When the plume is lost, the heading of the ith vehicle becomes

$$\psi_v(t) = \psi_u(t) - \gamma_i + \phi \tag{6}$$

where $\gamma_i = 10i$ and ϕ is a normal random variable with a standard deviation, σ. The random component ϕ is inserted to allow analysis of the effect of error in measuring wind heading. Each vehicle continues until either it reacquires the plume or the simulation times out (T = 200 s). At the end of each iteration, the reacquisition time for each vehicle is saved. Therefore, each iteration of the above simulation results in a record of T_R versus γ. Each iteration was classified into one of two categories based on this record. Category 1 contains those records closest in shape to $T_R = 200$ for γ in [0, 180] and $T_R = 0$ for γ in [180,360], which corresponds to those iterations where negative γ yielded better performance. Category 2 contains those records closest in shape to $T_R = 0$ for γ in [0,180] and $T_R = 200$

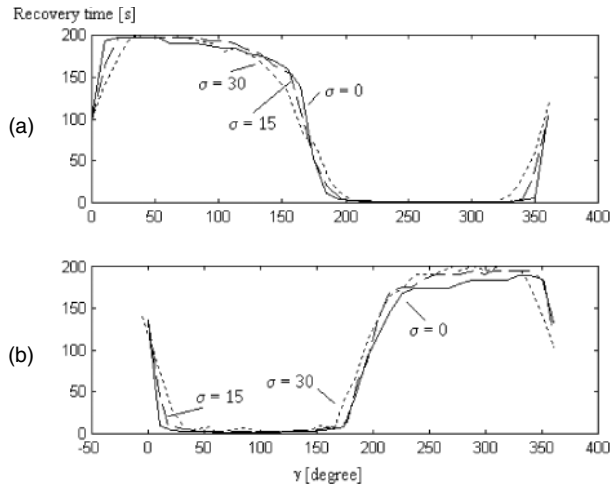

Figure 5 Mean reacquisition time versus the angle γ. The parameter σ indicates the standard deviation of the normal random variable ϕ. (a) Mean of all records for which $\gamma < 0$ (right of upwind) yielded better performance. (b) Mean of all records for which $\gamma > 0$ (left of upwind) yielded better performance.

for γ in [180,360], which corresponds to those iterations where positive γ yielded better performance. The average of all records in Category 1 is shown in Figure 5a. The average of all records in Category 2 is shown in Figure 5b. This analysis was repeated for three different values of the parameter σ, with similar results.

Note that the objective in this section is to minimize recovery time. Figure 5 shows that crosswind (i.e., 270° for Figure 5a and 90° for Figure 5b) provides both the minimal reacquisition time and the greatest robustness to error (± 70°) in the vehicle heading. Note that at this point in the discussion the optimal magnitude of γ is known and the records have been sorted a posteriori according to the sign of γ that yielded optimal performance. The remaining goal is to determine a means for reliably predicting the sign of γ, during the search, based on the data available to the searcher (i.e., $\psi_w(t)$, $\psi_v(t)$, vehicle position, and $c(t)$).

To solve this remaining problem, we further analyze the data of the previous simulation for the cases where $\gamma = -90$ or $\gamma = +90°$. In this further analysis, we consider $\psi_w(t)$ and $\psi_v(t)$ at the time $t = T_{\text{Lost}}$. For each of the 100 iterations, we determine which sign of γ yielded the minimum value of T_R. For this value of γ we further analyze $\psi_w(t)$ and $\psi_v(t)$. For the 58 iterations classified as Category 1 ($\gamma = -90°$ is optimal) 90% had

$$\psi_r(t) = \psi_v(t) - \psi_w(t) < -180°.$$

Table 1 Summary of Strategy 1

Action	*Condition*
$\psi_v(t) = \psi_u(t)$	If $c(\nu) > \tau$ for some $\nu \in [t-\lambda, t]$
$\psi_v(t) = \psi_u(t) - \gamma$	If $c(\nu) < \tau$ for all $\nu \in [t-\lambda, t]$ and $c(\nu) > \tau$ for some $\nu \in [t-T_w, t-\lambda]$
$\gamma = +90°$	If $\psi_v(T_{\text{Lost}}) - \psi_w(T_{\text{Lost}}) > -180°$
$\gamma = -90°$	If $\psi_v(T_{\text{Lost}}) - \psi_w(T_{\text{Lost}}) < -180°$

For the 42 iterations classified as Category 2 ($\gamma = +90$ is optimal) 90% had

$$\psi_r(t) = \psi_v(t) - \psi_w(t) > -180°.$$

These results lead to the conclusion that by defining

$$\gamma = +90° \text{ if } \psi_v(t) - \psi_w(t) > -180° \text{ and} \quad (7)$$

$$\gamma = +90° \text{ if } \psi_v(t) - \psi_w(t) < -180° \quad (8)$$

at $t = T_{\text{Lost}}$ the strategy should result in successful plume reacquisition with approximately 90% success.

The above results motivate Strategy 1 shown in Table 1, for maintaining plume contact. Two comments relative to the physical interpretation of the results of this section and the practicality of Strategy 1 are important. First, because Strategy 1 commands the vehicle to travel upwind (relative to the recent wind direction), the quantity $\psi_r(T_{\text{Lost}}) = \psi_v(T_{\text{Lost}}) - \psi_w(T_{\text{Lost}})$ is related to the derivative of the wind direction. Therefore, the above conclusions have the physical interpretation that to reacquire contact with the plume, the vehicle should turn in the direction opposite to the change in the wind direction. Second, because Strategy 1 uses $\beta = 0°$ (i.e., upwind) and the wind direction cannot be measured perfectly, Strategy 1 may not be robust in the sense that γ may have the incorrect sign due to noise in measurement of the wind direction.

Figure 6 shows a contour plot of the source finding time of Strategy 1, with a vehicle speed of 1 m s^{-1}, λ = 1.0 s, and T_W = 3.0 s when combined with the plume-finding and source-declaration methods of the Appendix. The data for the contour plot is generated, with the large meander plume-simulation parameters (see Section 2.3), by defining a regular 10 × 10 grid of starting locations over the 100 × 100 m search area. The searcher starts at each location 100 times (i.e., 10,000 total searches). The time to find the source from each starting location is averaged to produce a matrix of average times to find the source as a

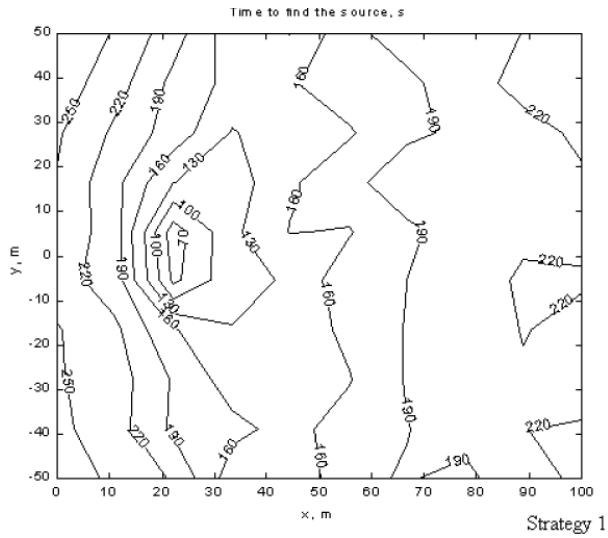

Figure 6 Contour plot of the average time to find the odor source as a function of the initial searcher position for a searcher using Strategy 1. The average finding time for each initial position is evaluated over 100 runs from the starting position. The large meander set of plume parameters was used. The nominal speed at which the searcher travels is 1 m s^{-1}.

function of position. In the simulations corresponding to Figure 6 (and the contour plots in later sections), if a searcher failed to find the source within the 300 s simulation duration, then the value 300 s was used as its source-finding time. The odor source is located at (20, 0) m. This form of contour plot will be used throughout the remainder of this article to illustrate overall search performance.

4.2 Design of Maintaining Behavior with Passive Strategy

The previous section determined γ based on $\psi_r(T_{\text{Lost}})$ for the case where $\beta = 0$. This section extends the analysis to consider the more general case where β is nonzero. Throughout this section, the plume finding behavior is that described in Appendix B. The strategy for maintaining plume contact is as follows:

$$\psi_v(t) = \psi_u(t) + \phi \qquad \text{if } c(t) > \tau \tag{9}$$

$$\psi_v(t) = \psi_u(t) + \beta + \phi \quad \text{if } c(t) < \tau, \text{ but } c(v) > \tau \text{ for some } v \in [t - \lambda, t] \tag{10}$$

where for the ith vehicle $\beta = -60 + 120i/36$, i is the vehicle index in [1, 36], and ϕ is a uniform random variable with maximum amplitude A. With this strategy, the searcher turns toward upwind while sensing above-threshold concentration and turns to an angle β relative to upwind when not instantaneously detecting above-threshold odor. If the vehicle has not detected above-threshold odor within the last λ seconds, then that iteration of that vehicle's simulation ends.

For each iteration of the simulation 36 vehicles started at $(x,y) = (85,0)$ with velocity 1.0 m s^{-1}. All 36 vehicles implemented the same plume-finding behavior until odor was detected. After odor detection all vehicles implemented the strategy described above with a different value of β for each vehicle. Each vehicle continued until either the simulation timed out, odor was not detected for λ s, or the source was found. For each searcher and each iteration, the total amount of time in contact with the plume (i.e., the time duration that Equations 9 and 10 were used) and $\psi_r(T_{\text{Lost}}) = \psi_w(T_{\text{Lost}}) - \psi_v(T_{\text{Lost}})$ was recorded. This process was repeated for 100 iterations.

Each iteration of the above simulation results in a record of T_{M} versus β. Each iteration was classified into one of two categories by a process similar to that described in the previous section. Category 1 corresponds to those iterations where a positive β yielded better performance. Category 2 corresponds to those iterations where a negative β yielded better performance. The average of all records in Category 1 is shown in Figure 7a. The average of all records in Category 2 is shown in Figure 7b. The batch of 100 iterations was repeated twice, once with the random component ϕ equal to zero and once with the random component uniformly distributed with $\phi \in [-10°,10°]$.

For the Category 1 iterations, $\psi_r(T_{\text{Lost}}) > -180°$ degrees for 81% of iterations when $\phi = 0°$ and 78% of iterations when ϕ is uniform over $\pm 10°$. For the Category 2 iterations, $\psi_r(T_{\text{Lost}}) < -180°$ for 91% of iterations when $\phi = 0°$ and 86% of iterations when ϕ is uniform over $\pm 10°$. Combining these statistical relations with the conclusion of the previous section, we conclude that we can use sign(β) to determine the sign of γ instead of using $\psi_r(T_{\text{Lost}})$. This is useful since β is an internal variable whereas $\psi_r(T_{\text{Lost}})$ is a possibly noisy measurement. The resulting strategy is referred to as Strategy 2 and is summarized in Table 2. Note that Strategy 2 is a counterturning strategy. The

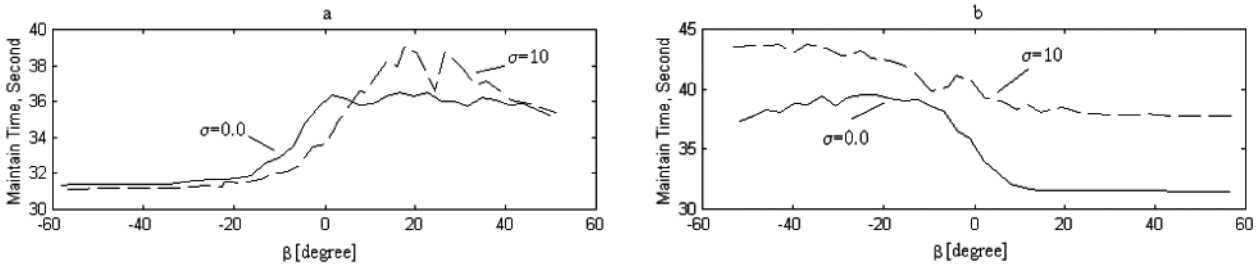

Figure 7 Mean of records from each reacquistion iteration after classification. (a) Mean of all records classified as Category 1 (i.e., $\beta < 0$ is better). (b) Mean of all records classifed as Category 2 (i.e., $\beta > 0$ is better).

Table 2 Summary of Strategy 2

Action	*Condition*
$\psi_v(t) = \psi_u(t)$	If $c(t) > \tau$
$\psi_v(t) = \psi_u(t) + \beta$	If $c(t) < \tau$, but $c(v) > \tau$ for some $v \in [t - \lambda, t]$
$\psi_v(t) = \psi_u(t) - \gamma$	If $c(v) < \tau$ for all $v \in [t - \lambda, t]$ and $c(v) > \tau$ for some $v \in [t - T_w, t - \lambda]$ where $\gamma = 90° \text{ sign}(\beta(T_{\text{Lost}}))$

reacquisition behavior requires the searcher to travel in the opposite crosswind direction to the maintain behavior. Figure 7 shows that any $\beta > 10°$ yields good performance. Improving performance by allowing β to be a time-varying function will be considered in the next section.

It is interesting to compare, for a biological entity, the two strategies of (1) traveling with $\beta = 0$ and attempting to reacquire based on $\psi_r(T_{\text{Lost}})$, or (2) traveling with $\beta \neq 0$ and counterturning based on sign(β). Strategy 1 requires a low noise sensor with a bandwidth high enough to detect $\psi_r(T_{\text{Lost}})$ rapidly and accurately. Note that with $\beta = 0$ accurate decisions concerning whether $\psi_r(T_{\text{Lost}})$ is greater or less than 180º are not robust to wind sensor noise. For biological flying entities, a combination of airspeed and visual cues could yield a low bandwidth wind direction estimate (see discussion in Cardé and Knols, 2000), but this system of measurement is unlikely to produce an accurate high bandwidth estimate of $\psi_r(T_{\text{Lost}})$. Strategy 2 requires only a low bandwidth (relative to the vehicle dynamics) wind sensor to maintain the wind relative direction of travel. The counterturning of Strategy 2 is inherently more robust to wind sensor noise, since the counterturning direction is not directly dependent on instantaneously sensed wind direction. A disadvantage of Strategy 2 is that its mean "in the plume" path to the source will be longer than that of Strategy 1 since for Strategy 2 $\beta \neq 0$. Strategy 1, which has similarities with aim-then-shoot strategies, described by Dindonis and Miller (1980) and named by Kennedy (1986), is advantageous either when the flow field is nearly uniform and time invariant or when the searcher is very near the source where the wind direction is reasonably well correlated with the direction of the plume's long axis; however, when Strategy 1 causes the searcher to exit the plume the searcher may not be able to turn reliably in the correct direction to rapidly reacquire the plume. Strategy 2 has a plume reacquisition advantage, which is beneficial when the searcher is following long, meandering plumes.

A physical interpretation for the utility of $\beta \neq 0$ for predicting the correct direction to turn to reacquire the plume is that if the vehicle travels due upwind within the plume (i.e., $\beta = 0$), then at the moment when the vehicle exits the plume the vehicle is equally likely to be on either side of the plume. Therefore, without additional information, either direction the vehicle turns has an equal chance of immediately recontacting the plume. Alternatively, if the vehicle travels to the left of upwind ($\beta < 0$) while in the plume, the vehicle is more likely to exit the plume via its left edge. Therefore, a right turn (counterturn) improves the likelihood of quickly recontacting the plume.

Note that it may be possible to increase the performance by using both the information ψ_r and crosswind travel ($\beta \neq 0$). We have not pursued this line of analysis since the increased distance of travel is not large if $\beta < 45°$ and we do not expect ψ_r to be measurable with high resolution or bandwidth on autonomous vehicles. Also, our experience in this project showed that failure of the plume reacquisition process resulted in a large time penalty, whereas crosswind travel

resulted in only a small penalty (for $\beta < 45°$). The active strategy of the next section tries to balance these two issues by making upwind progress (β near 0) when entering the plume and increasing the likelihood of counterturning success by crosswind travel (β approaching 90°) as the time in the plume increases.

4.3 Design of Maintaining Behavior with Active Strategy

The previous section demonstrated the advantage of the counterturning behavior with $\beta \neq 0$. In this section, we analyze the possible benefits of using a time-varying upwind angle $\beta(t)$ to improve the trade-off between up-plume progress and maintaining contact with the plume. In this "active" strategy, the vehicle heading, within the plume, will be

$$\psi_v(t) = \psi_u(t) + \phi + \beta(t), \tag{11}$$

where ϕ is an error component simulating sensor inaccuarcy and

$$\beta(t) = \begin{cases} \beta_1 & (t - T_{\text{First}}) \leq \tau_1 \\ 2\beta_1 & \tau_1 < (t - T_{\text{First}}) \leq \tau_2 \\ 3\beta_1 & (t - T_{\text{First}}) > \tau_2 \end{cases}$$

is a time-varying offset to upwind with T_{First} redefined after each plume reacquisition.

The time variation is illustrated in Figure 8. Upon first entering the plume, the searcher turns toward $\psi_v = \psi_u(t)$ (i.e., β small). As the time in the plume increases, $\beta(t)$ increases in magnitude (toward either $\pm 90°$). The sign of β is selected based on the direction the searcher was moving upon entering the plume, so that the vehicle continues progress into the plume. Therefore, on first entering the plume [i.e., $\beta(t)$ near 0] significant up-plume progress is achieved. As $|\beta(t)|$ increases, the searcher traveling across the plume will, by design, leave the plume. Since the direction of cross-plume travel is known at T_{Lost}, the searcher knows the direction in which to turn (from the results of the previous section) to maximize the likelihood of rapidly reacquiring the plume. Upon reacquisition, the time varying $\beta(t)$ process is repeated with the opposite signs.

Optimization of this strategy requires analysis of the effects of various parameters. To analyze the effects of the strategy parameters, 100 iterations of the following set of simulations were performed. For each iteration, 100 searchers were initiated at $(x, y) = (60, 5)$ m. All 100 searchers performed exactly the same plume-finding maneuvers until odor was detected. At the time of detection, each searcher initiated the plume-maintaining behavior, with a distinct set of strategy parameters. For example, for the left column of plots in Figure 9, in the approaches using constant β (i.e., the bottom curve), the ith vehicle used $\beta = 0.85i$. Each iteration results in a record of T_{M} versus β. The 100 records from the 100 iterations are averaged to produce the lower curve of the left column of Figure 9. For this study, only one parameter was changed at a time. For all the simulations, the reacquisition strategy parameters of the previous section were $\gamma = \pm 90°$ and $\lambda = 1.0$ s. If the searcher reacquired the plume within λ seconds, then T_{M} kept accumulating. Otherwise, the simulation of that searcher and that iteration terminated.

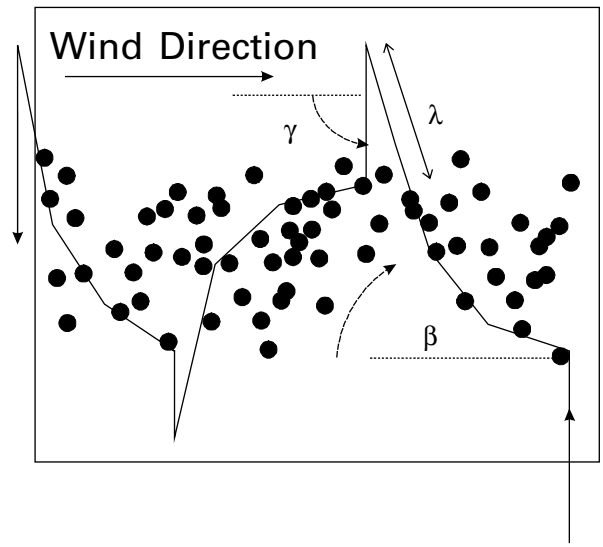

Figure 8 Depiction of the plume-maintaining behavior with a time-varying offset relative to upwind $\beta(t)$. Black dots depict above-threshold patches of odor.

The left column of graphs in Figure 9 evaluates the effect of the search angle parameter β. The top left graphs shows the results of simulations using the set of simulation parameters that results in large-scale plume meander. The bottom left graphs shows the results of simulations using the set of simulation parameters that results in small-scale plume meander. Note that the maintain time is significantly longer when the

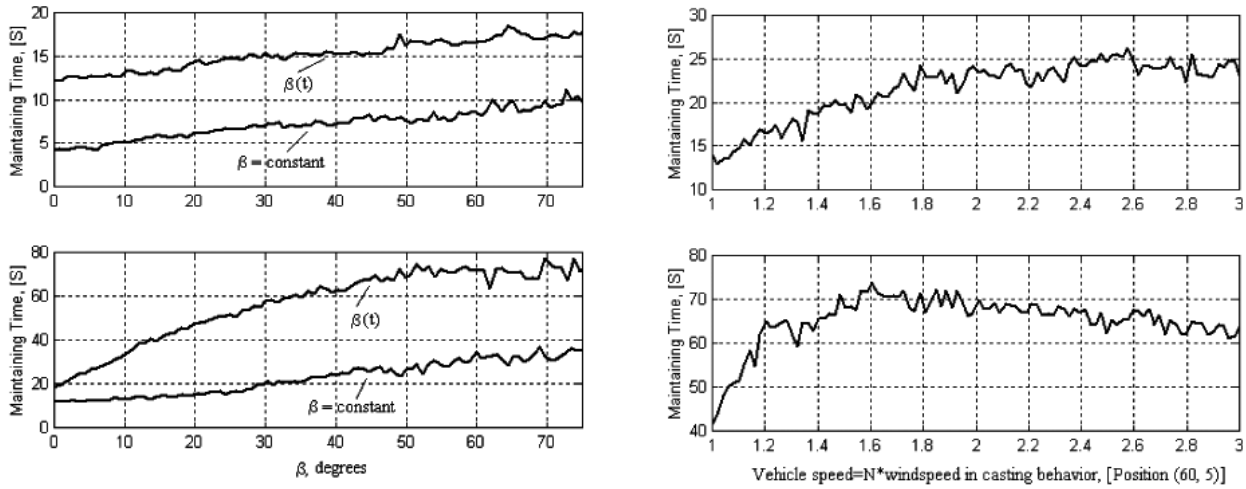

Figure 9 Maintain time dependence on strategy and environment parameters. Each plotted point is the average over 100 repetitions. Note the different scales. Left: Maintain time versus the angle β. Graphs in the top row use simulation parameters resulting in large amplitude, low frequency plume meander. Graphs in the bottom row use simulation parameters resulting in small amplitude, high frequency, plume meander. The bottom curve on each of the left-column graphs uses a constant β. The upper curve on each of the left-column graphs uses the time varying $\beta(t)$. For the $\beta(t)$ curves, the horizontal axis should be read as the maximum of $\beta(t)$ (i.e., $3\beta_1$ in the definition of $\beta(t)$ after Equation 11). Right top: Maintain time versus vehicle speed during reacquisition phase. Right bottom: Maintain time versus vehicle speed during maintain phase.

meander is small. The bottom curve of each of the left-column graphs indicates average maintain time as a function of the constant angle $\beta \in [0°,75°]$, as shown in Figure 3. This curve shows that maintain time is increased by traveling at larger crosswind angles; however, larger crosswind angles result in less up-plume progress per transect. The top curve of each of the left column graphs indicates average maintain time as a function of the time-varying angle $\beta(\tau)$, as shown in Figure 8 and defined following Equation 11. To produce the $\beta(t)$ curves, β_1 (=0.25i for the ith vehicle) varied (in different batch simulations) from 0° to 25°. Therefore, $3\beta_1$, which is the maximum of $\beta(t)$, was varied from 0° to 75°. The left set of graphs shows that there is a significant advantage in terms of time in the plume when using the time-varying β approach. In addition, the approach makes substantial up-plume progress each time the searcher reenters the plume as $\beta(t)$ reinitializes to near 0°. The right column of graphs in Figure 9 shows the effect of searcher speed on T_M. The top graph shows T_M versus speed of travel during reacquisition. The bottom graph shows T_M versus speed of travel within the plume. For both graphs, the horizontal axis is the multiple of wind speed N (i.e., vehicle speed equals N times the average wind speed.). The function

$$\beta(t) = \begin{cases} 10 & (t - T_{\text{First}}) < 0.2s \\ 65 & 0.2s < (t - T_{\text{First}}) < 0.5s \\ 85 & (t - T_{\text{First}}) > 0.5s \end{cases} \quad (12)$$

with T_{First} redefined after each plume reacquisition, was used during the simulations that generated the right column of Figure 9 and will be used in all subsequent simulations.

4.4 Strategy Summary

Sections 4.1—4.3 have developed alternative strategies for maintaining contact with an odor plume. These strategies are summarized in Table 3. The results of Section 4.1 showed that counterturning behavior was beneficial to the process of reacquiring the plume but used a mechanism (difference between the wind and vehicle directions) that was not robust to wind sensor noise. Section 4.2 showed that crosswind travel while in the plume simplifies the definition of

Table 3 Summary of the Strategies Specified in Sections 3 and 4

Strategy	*Maintain behavior*	*Reacquisition behavior*	*Parameters*
1. Upwind with wind-relative counterturns	If the searcher has detected odor within the last λ seconds, then the searcher travels straight upwind.	If the searcher has detected odor within the last T_w seconds, but not the last λ seconds, then the searcher travels at an angle $\gamma = \pm 90°$ to upwind. The sign of γ is defined based on the vehicle heading and wind direction to cause the vehicle to reverse its crosswind direction (i.e., counterturn)	$\lambda = 1$, $T_w = 4$, and $\gamma = \pm 90°$
2. Constant crosswind with counterturn	If the searcher is currently sensing odor, it turns toward straight upwind. If the searcher is not currently sensing odor, but has sensed odor within the last λ seconds, then the searcher travels at an angle β with respect to upwind.	If the searcher has detected odor within the last T_w seconds, but not the last λ seconds, then the searcher travels at an angle $\gamma = \pm 90°$ to upwind. The sign of γ is defined as the opposite of the sign of β. This generates counterturning travel and the sign of γ is not sensitive to the measured wind direction or vehicle heading	$\lambda = 1$, $T_w = 4$, $\beta = 35$, and $\gamma = \pm 90°$
3. Progressive crosswind with counterturn	If the searcher has detected odor within the last λ seconds, then it travels at an angle $\beta(t)$ relative to upwind, where t is the time since odor was first detected on the present crosswind traverse of the plume.	If the searcher has detected odor within the last T_w seconds, but not the last λ seconds, then the searcher travels at an angle $\gamma = \pm 90°$ to upwind. The sign of γ is defined as the opposite of the sign of β. This generates counterturning travel and the sign of γ is not sensitive to the measured wind direction or vehicle heading	$\lambda = 1$, $T_w = 4$, $\beta(t)$ described by Eq. 12, and $\gamma = \pm 90°$

the counterturn and makes that definition relatively immune to wind sensor noise. Section 4.3 improved the rate of travel up the plume by varying the angle of crosswind travel while within the plume.

The relative performance of the three strategies under identical conditions is shown in Figure 10. This figure was generated by starting searchers using each strategy at positions along the x-axis [i.e., $(i,0)$, $i = 0, \ldots, 100$]. Each strategy was started 500 times from each starting location. The average time to find the source versus starting location for each strategy is shown in Figure 10. For initial locations far from the odor source, Strategy 3 yields 30–40 s improvement over Strategy 1 and 10 s improvement over Strategy 2. Near the odor source, where the wind and plume axis are nearly in the same direction, all three strategies have nearly identical performance.

The comparison between the strategies of Table 3 and the behavior of flying moths and other biological entities is not necessarily straightforward for the following reasons. For the strategies developed herein, the behaviors are explicitly known, because they are programmed. For moths, the behaviors are known, but the sensory inputs are inferred, based on observations of moths flying in wind tunnels where only the approximate instantaneous position of the plume is known (e.g., Justus, Schofield, Murlis, & Cardé, 2002). For the simulation analyses herein, the method allows explicit knowledge of when the searcher is detecting odor. Finally, the simulations involve a large search area and plumes with significant meander. Much of the literature on biological plume tracing involves biological entities confined within a flume or wind tunnel. Such experiments necessarily are restricted to a relatively small search area within which plume meander is constrained by the vertical and lateral expanse of the wind tunnel or flume. Such wind tunnel experiments correspond to the portion of Figure 10 in which all three strategies yield nearly identical levels of performance.

5 Performance Evaluations Using Simulated-Odor Plumes and Experimental Water Flumes

This section presents performance evaluation for the complete Strategy 3 (including plume finding and source declaration as described in Appendix B) with β defined in Equation 12, relative to the simulated plume

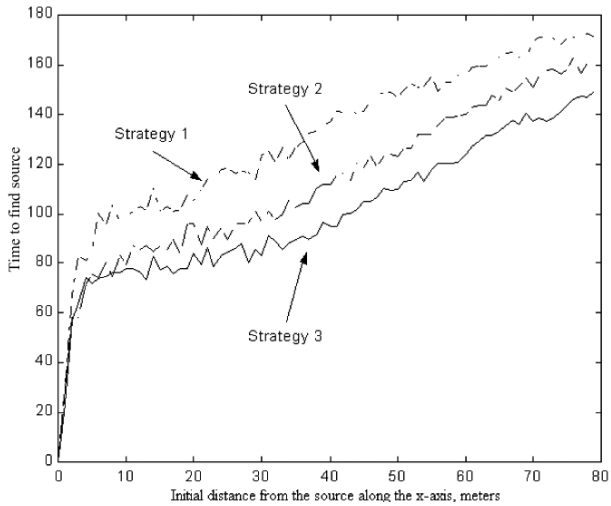

Figure 10 Time to find the source versus initial *x*-axis starting position for the three strategies summarized in Table 3. The starting positions were $(i,0)$ for $i = 0, \ldots, 100$. Performance from each starting location was evaluated 500 times.

and three sets of experimental flume data. During the simulation evaluation, the measured fluid flow was corrupted by additive noise ϕ that is a white normal random process (Jazwinski, 1970) that changes value at a 1.0-Hz rate. If $c(\nu) < \tau$ for all $\nu \in [t - T_W, t - \lambda]$, then the vehicle reverts to the plume-finding behavior.

Four sets of batch simulations were evaluated. For each batch evaluation, the strategy and vehicle model were identical, but the source of the concentration and flow data were unique. Four sources of concentration and flow data were used: the simulation model, flume data from Stanford University (Crimaldi & Koseff, 2001), Cornell University (Cowen, Chang, & Liao, 2001), and Georgia Institute of Technology (Webster, Roberts, & Ra'ad, 2000; Webster & Weissburg, 2001). Each of the flume data sets contained both concentration and flow information obtained by exciting a fluorescent dye with a laser sheet within a portion of a flume. Images of this area were obtained by video camera and stored. The image intensity is interpreted as concentration. Correlation of concentration data between sequential images was used to determine the local flow. Interpolation between images based on the fluid velocity was used to generate plumes changing continuously in time from the discrete images. The image, flow vector, and time were scaled so that the image fit a 20 × 60-m portion of the 100 × 100-m search area. The 20 × 60-m flume image was placed within the search area so that the source appeared to be at (20,0) m. Outside the flume image area, the concentration was set equal to zero and the flow was equal to the mean flow over the flume image.

Figure 11 shows contour plots of the source-finding time of Strategy 3 with $\lambda = 1.0$ s and $T_W = 4.0$ s. Figure 11 contains data from four different batches of simulations. The data for each contour plot is generated by defining a 10 × 10 grid of starting locations over the 100 × 100-m search area. The searcher starts at each location 100 times (i.e., 10,000 total searches). The time to find the source from each starting location is averaged to produce a 10 × 10 matrix of average time to find the source as a function of position. The odor source is located at (20,0) m. For the upper left contour plot, the large meander plume simulation was used. Proceeding clockwise, the contour plots were generated using experimental data sets from Stanford, Georgia Tech, and Cornell. Figure 11 should be compared with Figure 6 to observe the improved performance achieved by optimization of the plume-maintaining and reacquisition behaviors. Note that, by comparison of the four plots of Figure 11, performance is least proficient for the simulated plume. This is caused by the flumes lacking significant lateral meander due to the width constraint of the flume. Therefore, the simulated plume, with the large lateral meander typical of plumes in natural environments, is significantly more challenging to trace than (currently available) wind tunnel or flume-generated experimental data sets. Also, note that 97% of all searchers succeeded in locating the source within a 1-m radius within the 300 s time limit.

6 Plume-following strategies of Animats and Biological exemplars

The selective pressures for optimizing strategies of plume following can be expected to vary with the adaptive value of rapid location of the odor's source, and with the penalty for particular plume-following strategies that result in loss of the scent. It may be that maneuvers such as high upwind velocity that seemingly promote rapid location are also more prone to cause the searcher to exit and possibly lose the plume. Plume following by flying male moths seeking a pheromone-emitting female can be expected to be highly constrained to behaviors that promote rapid

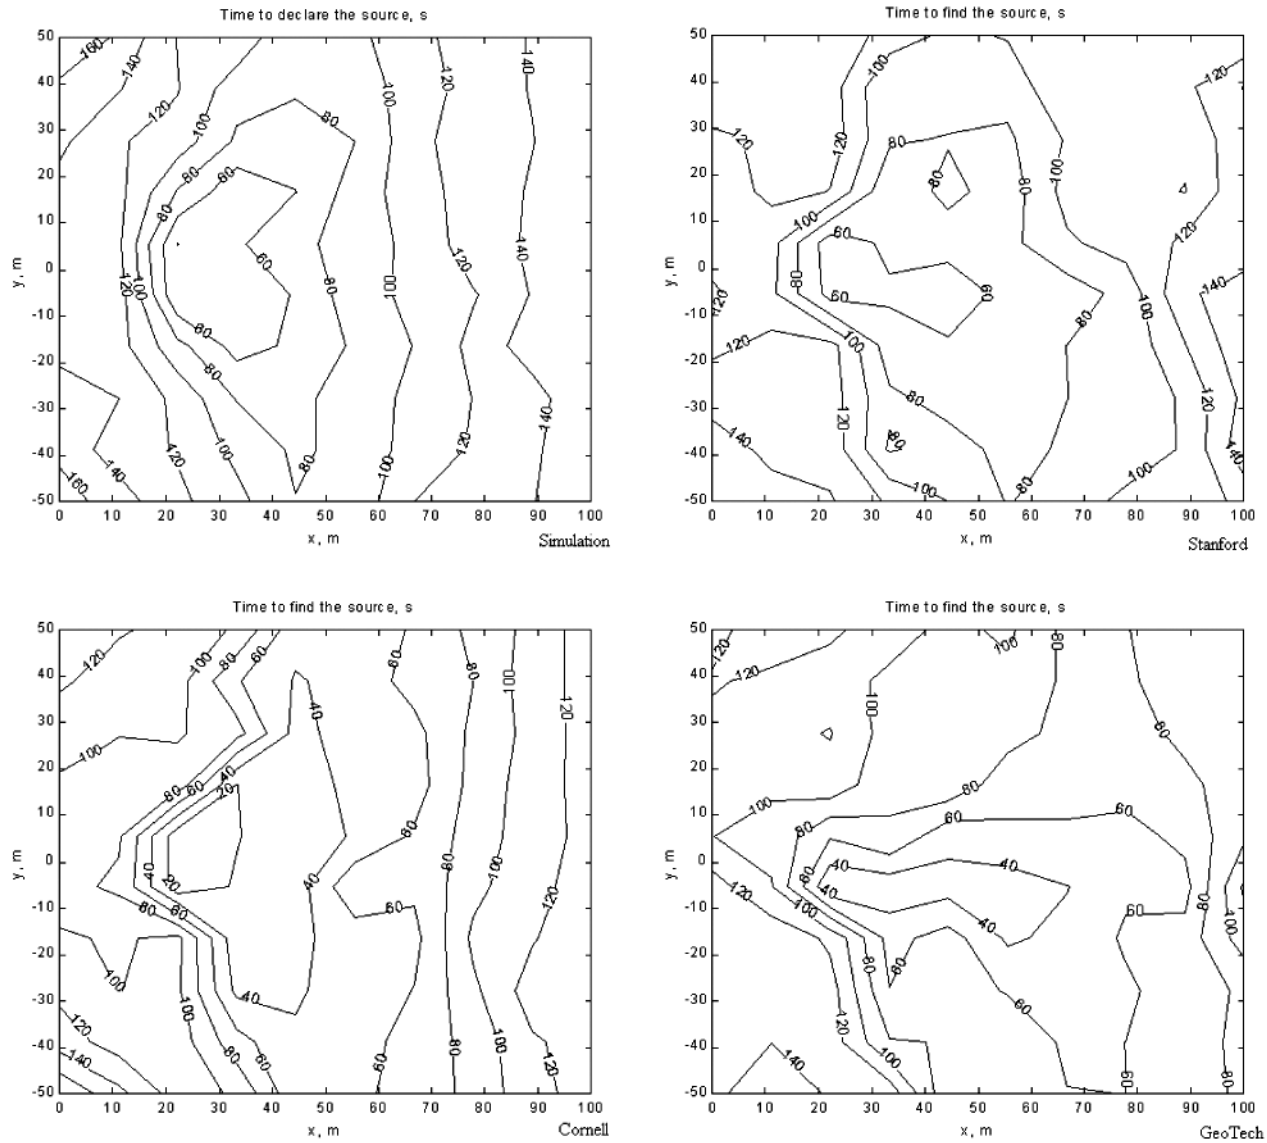

Figure 11 Contour plots of Strategy 3 performance as measured by the total time to maneuver to within 1 m of the source location. The four plots correspond to results using the plume simulation and three experimental flume data sets.

location of the female, given that the first-arriving male is generally most apt to mate with the female (Cardé, 1986). In contrast, location of a host plant by tracing its odor plume typically may not be so time constrained in that the plant remains as an available resource until it is located. Finding a prospective plant host after several minutes of searching is not as likely to have as strong a selective disadvantage as failing to find a mate.

A feature common to moths maneuvering along a pheromone plume is a zigzagging[4] path. In our simulation such counterturning ($\beta \neq 0$) would seem to have two advantages: increased ability to maintain intermittent contact with the plume; and decreased sensitivity to the sensed wind direction. In moths, counterturning usually is assumed to reflect an underlying motor program that is suppressed when the trajectory is aimed more directly upwind (Baker, 1990). The importance of counterturning (producing a generally zigzag path) to the success of upwind orientation has not been fully resolved. Preiss and Kramer (1986) proposed that a zigzag path was simply a consequence of moths being unable to gauge upwind direction accurately; therefore, reiterative corrections toward upwind yielded the

counterturning. The fact that some moths can fly upwind without evident counterturns (Cardé & Mafra-Neto, 1996), however, argues against this explanation. Zigzagging also can be viewed as a strategy to enhance the information available to gauge upwind direction (Cardé, 1984): moths heading off the windline experience more transverse image flow, the visual cue used to set their heading. Zigzagging as well enables the moth to sample a wider expanse of the plume, which might also increase the male's ability to trace the plume (Kennedy, 1983). Although a counterturn generally occurs when a moth "exits the plume," and therefore brings the moth back within the plume, counterturns also can occur within the plume's boundaries (Kennedy Ludlow, & Sanders, 1980; Willis & Baker, 1984; Justus & Cardé, 2002), and therefore their occurrence and temporal regularity seem dictated by a putative counterturn generator.

The sensory inputs and causal mechanisms governing counterturning remain unverified. The observer of the experiment specifies when the moth is "in the plume." In some cases, the intermittent nature of the plume could be causing the moth's sensed concentration to be on occasion below threshold even when the experimental observer declares that the moth is "in the plume." Counterturning need not be mandated by an internal motor program but rather, it could be dictated by loss of the plume followed by a turn back toward the direction where it was last sensed (as proposed explicitly by Kennedy, 1977, cf. Kennedy, 1983). Even though there is debate over the neurobiological source of counterturning (i.e., a counterturn generator) and the ethological mechanisms (induced by a loss of scent or a corrective navigational maneuver), a main conclusion of these simulations is that there is an advantage to counterturning for robustly maintaining contact with a plume.

The strategies used by male moths flying upwind along a pheromone plume also include linear surges aligned with the instantaneous wind flow and a zigzag trajectory with counterturns typically expressed several times a second. In the almond moth, *Cadra cautella*, a continuum between both forms of these flight tracks is expressed, contingent upon the temporal pattern of the moth's contact with filaments of pheromone (Mafra-Neto & Cardé, 1994). Counterturning with little or no progress upwind is caused by sporadic encounters with odor filaments, usually at a rate of < 5 Hz, and a zigzag path upwind is evoked by rates of > 5 Hz. In some moths, rates above 5 Hz produce an accelerated velocity and a straightened-out track (Justus et al., 2002). In such wind tunnel assays, fast, relatively straight tracks are more apt to result in location of the odor sources, and to decease the time spent in orientation (Mafra-Neto & Cardé, 1994).

The simulation Strategy 1—heading due or nearly due upwind when odor has been recently sensed—matches the behavior of some moths encountering filaments of pheromone at rates of at least 5 Hz (e.g., Mafra-Neto & Cardé, 1994; Vickers & Baker, 1994a). In the context of a wind tunnel assay over a distance of less than 2 m with an unvarying wind velocity and direction, this type of strategy enhanced the likelihood and rapidity of finding a pheromone source. In Strategies 2 and 3 the track is angled upwind while contact with the plume is maintained; turns occur when the search exits the plume. The direction of the turn is back toward the plume's putative centerline and toward the windline. These two strategies superficially resemble the zigzag path assumed by many moths flying along a pheromone plume (Cardé, 1996) and other insects walking upwind in the presence of pheromone (Bell, 1984).

There are two important distinctions. In many wind tunnel assays in which the boundaries of the plume are visually defined by the observer, a moth will perform some counterturns while within the plume's boundaries (e.g., Kennedy et al., 1980; Willis & Baker, 1984), a response not dictated by Strategies 2 and 3; however, this behavior also could result if the counterturn was triggered by plume intermittency even while the searcher was supposed to be "in the plume."

Another issue is the spatial and temporal scale of wind tunnel assays of moth trajectory. Our model's searcher velocity is set at 1 m s^{-1} with a simulation step size of 0.01 s; therefore, the functional step size of the simulation is 0.01 m every 0.01 s. Wind tunnel assays are highly constrained by the wind tunnel's dimensions (which rarely exceed 3 m in length and 1 m in width), and movements are monitored many times a second (e.g., 0.03 s^{-1}). However, in nearly all wind tunnel assays the plume is filamentous, so that the odor signal flickers many times a second (Justus et al., 2002). In such plumes, males may counterturn several times per second. In experimental manipulations of the plume structure in a wind tunnel, the spatial

scale over which the plume was generated could have contributed to the finding that the moths seeming to use Strategy 1 had a clear advantage over those seeming to use Strategies 2 and 3, in contrast to our simulations. In any case, the size of the wind tunnel did not allow observations of more than 4 s of flight over more than 1–2 m and so the applicability of these observations to plume following in the field cannot yet be assessed.

The spatial and temporal scale at which our simulation is defined also may explain why the highly successful Strategy 3 in our simulations does not appear to have an exact counterpart observed to date in flying moths, or in the reactions of other flying or walking insects navigating along an odor plume in wind tunnels. Not unexpectedly, wind tunnel studies in general seem to simulate plume structure conditions found near the source: a narrow plume and an alignment of the plume's long axis and the windline. Strategies 2 and 3 implement a counterturning behavior that requires only a low bandwidth measurement of the wind direction and a binary concentration sensor. Currently, it is not clear whether similar behaviors in biological entities are the result of the entity's limited sensing ability (e.g., sensing wind direction to 20º, see Preiss & Kramer, 1986) or whether these behaviors are the result of the evolutionary optimization process, because they give the searcher some advantage in the task of tracing the plume. Some speculations based on the analysis herein have been included. Although the methods of this article provide no evidence of how these behaviors might be implemented biologically, the studies of this article demonstrate that the counterturning behavior with at least 15º deflection from upwind (Strategy 3) provides a significant advantage to behaviors that maintain a constant, nearly due upwind, direction within the plume. Strategy 3 may be most valuable at moderate to long distances from the odor source where the local wind vector neither aligns with the plume's centerline nor points toward the odor source. This is exemplified in Figure 10: Strategies 1, 2 and 3 yield equivalent performance within about 5 m of the source.

Other efforts to model plume-following tactics have not addressed the range of issues considered in our simulations. Williams (1994) modeled tsetse fly seeking of traps baited with host odor, comparing the rapidity and probability of source location by upwind, anemotactic flight while in the plume. Location of the odor source was essentially contingent on how accurately flies were able to detect wind direction, in effect to bias their two-dimensional "random-walk" toward upwind. The plume's fine-scale structure and plume meander were not explicit components of the Williams model.

Kramer (1996) has provided another approach, modeling the upwind walking behavior of the commercial silkworm moth (*Bombyx mori*) and a hypothetical flying moth. Within the plume's overall boundaries, the plume was assumed to have an internal filamentous structure (with a mean pulse interval of 0.33 s) that induced upwind displacement. The precision of upwind heading was set to improve when the odor-pulse rate increased up to 3–5 s^{-1}, whereas at higher rates upwind displacement ceased. Counterturning was governed by a small internal tendency to turn, the sign of which "flip-flops," producing an upwind zigzag and casting. The task set forth in these models was to demonstrate that these design features would yield tracks that resembled the actual walking and flight tracks and to suggest how a modulation of "internal turning tendency" might alter the success of orientation. Actual odor plumes in wind have much higher rates of filament production than 5 s^{-1} (Murlis et al., 2000).

Belanger and colleagues (Belanger & Willis 1996a, 1996b, 1998; Belanger & Arbas, 1998), who also perform simulation studies motivated by moth behavior, suggest strategies with similar underlying ideas to those proposed herein. A short description of their Surger strategy [see pp. 355–356 in Belanger & Arbas (1998) for a detailed discussion] is that recent odor detection causes the searcher to move upwind at a commanded wind-relative course generated once per detection (after accounting for latencies and olfactory refraction) with mean 0º and standard deviation 40º. If odor detection does not occur before an internal timer achieves a threshold, then casting flight begins. Casting flight continues until either odor is detected or the simulation times out. Their Surger strategy has some similarities with Strategies 1 and 2 proposed herein. Note however that the Surger strategy does not force the surges to alter their crosswind relative direction. Strategies 1 and 2 proposed herein both attempt to force counterturning after contact with the plume is lost. The Belanger Counterturner strategy [see pp. 353–354 in Belanger & Arbas (1998) for a detailed discussion] generates counterturns when the time

since the last turn exceeds a time-varying threshold. If odor is detected within the previous oscillator cycle, then the commanded course angle is generated once per oscillator cycle as a uniform random variable with magnitude in [0°,70°] and sign defined to enforce a counterturn. If odor is not detected during the previous oscillator cycle, then the commanded course angle is generated once per oscillator cycle as a uniform random variable with magnitude in [85°,105°] and sign defined to enforce a counterturn. This Counterturner strategy has similarities with Strategy 3 proposed herein. The main difference is the counterturns of Strategy 3 are triggered by lack of recent odor detection whereas those of Belanger are triggered by an internal clock. Also, after detection, the course commands of Belanger are defined to be upwind plus a random component. In Strategy 3, the after-detection course commands initially point toward upwind but vary with time toward crosswind, purposefully to increase the likelihood that the searcher knows from which plume edge it has departed.[5]

This article and the studies of Belanger et al. have different objectives, which result in different methods of analysis. The goal of Belanger et al. was "to determine whether any of the models … in the literature to account for moth orientation … could reproduce the behavior … in *M. sexta*" (Belanger & Arbas, 1998, p. 357). Therefore, they faithfully attempted to implement such models in software and were only interested in biologically feasible model parameters. Consequently, model parameters were based on those in the literature and numerous simulations were performed both to analyze performance sensitivity to parameter settings and to evaluate model performance relative to moth wind tunnel experiments. Alternatively, our goal was to start with models based on moth behaviors and develop them into good algorithms for autonomous vehicles. Our simulations were focused on the design process of analyzing good parameter settings and of identifying structural improvements in the strategies that could significantly improve performance. The fact that at the end of this process Strategy 3 and the parameters defined herein are biologically relevant is noteworthy.

The different foci of these articles also generated different experimental conditions. Because Belanger et al. were attempting to replicate wind tunnel experiments, they "released" their searchers near the center downwind edge of a 1×2-m search area with the odor source at the center of the upwind edge. Note that the release is 2 m directly downwind from the odor source. The flow is essentially laminar. Therefore, although their intermittency characteristics are tunable, the plume has no meander. The best source-finding performance Belanger et al. presented (Belanger & Willis 1996a, 1996b, 1998; Belanger & Arbas, 1998) was 56% success for a 300-s simulation. Herein, we have studied plume tracking over a 100×100-m area. The flow changes with space and time, generating significant plume meander and decorrelation of the wind direction, plume centerline, and direction to the source as the distance from the source increases. Also, searchers are released on a 10×10 grid of starting locations, which include searchers starting up to 20 m upstream of the source and 45 m offset from the mean plume location ($y = 0$) in the crosswind direction. The success rate of Strategy 3 in finding the source averaged over all starting locations was 97% for a 300-s simulation. Therefore, it is interesting to consider possible reasons for this improvement in performance for a significantly more difficult problem. First there is the difference in objectives. An objective of our simulations was to start from biological models and to improve their performance, whereas the objective of Belanger et al. was to evaluate strategies relative to models in the literature. Optimization of their models is further considered (in Belanger & Willis, 1998). Second, we decompose the problem into four parts: plume finding, plume maintaining, plume reacquisition, and source declaration. Casting is a portion of the plume reacquisition behavior. If plume reacquisition does not succeed within a few seconds, then we revert to plume-finding behaviors. In contrast, the strategies of Belanger et al. continue casting until they detect odor. As they state, many of their simulations failed because the moth moved so far from the plume that subsequent casting maneuvers never again brought the moth into contact with the plume. Third, the timing of the Belanger counterturns is set by an internal clock and the post-detection behavior may command a course that is due upwind. Such a course may result in the subsequent counterturn taking the searcher well outside the plume. The reacquisition counterturns and the in-plume maneuvers of Strategy 3 are defined to increase the likelihood that the first turn upon leaving the plume is in the correct direction. The importance of the sign of this first turn is also recognized by Belanger et al. (see p. 358 in Belanger & Arbas, 1998).

7 Caveats and Conclusions

The strategies that would enable robotic vehicles to follow a meandering, filamentous odor plume upwind to its point of origin must be robust to a range of environmental conditions. Optimization of an overall strategy for localization of an upwind and distant chemical source is a difficult problem. Therefore, the overall problem was decomposed into plume finding, plume maintaining, plume reacquisition, and source-found declaration. The motivation for decomposition was inspired by the known maneuvers of plume tracing by biological entities, particularly flying moths. The decomposition has been fruitful in the sense that each portion of the decomposition has proved tractable and successful strategies have been defined. Implementation of an overall strategy requires fusion of these behaviors into an overall approach, with the appropriate behavior being selected at each time instant. The literature on behavior-based robotics is useful for this research direction (Brooks, 1986; Li, 1994).

The main objective of our simulations was to design sensor, signal-processing, and actuation algorithms to enable autonomous vehicles to track chemical plumes to their source. Secondary objectives were to improve the understanding of the nature of this problem and the methods that biological entities use for its solution. Our approach has involved "optimization" of strategy parameters and redesign of algorithms based on successive iterations of the optimization process. Because the parameters of each strategy are interrelated, multiple iterations of the optimization could further improve the performance (i.e., multiple iterations of a line-search approach); therefore, the parameters of the strategies are not strictly "optimal." We did not pursue full optimization for two reasons. First, such an optimization would be tuned to the particular simulation environment of the optimization process. An additional important criterion is robustness of performance to environmental factors. Second, we were most interested in the structural features of the strategies that were required to yield good and robust performance. A single iteration of the optimization process was used to understand the trade-offs involved in parameter selection and to select reasonably good choices for the strategy parameters.

The main structural features that we identified related to strategy design are as follows. First, maintaining intermittent contact with the plume is critical to success. Flow direction while detecting the odor is critical information indicative of direction to the source. If contact with the plume is lost for more than a few seconds, then the searcher could spend large amounts of time and energy searching to reacquire contact with the plume without making progress toward the source. Therefore, the strategy design approaches described in this article selected maintain time while making progress along the plume as a key criterion to be minimized. Second, the ability of a searcher to maintain intermittent contact with a plume is determined by its ability to turn in the correct direction (back toward the plume) if it leaves the plume. Our analysis has shown that cross-plume travel with counterturning significantly increases the likelihood and robustness (to sensor noise) of turning in the correct direction. Because the counterturn is performed after leaving the plume and the counterturn involves turning through the upwind direction, this results in the searcher making some (possibly most) of its upwind progress while it is out of the plume. Similar up-plume progress while out of the plume has been observed in the field with released gypsy moths (David et al., 1983). In those observations of moth flight within 15 m of a point source of pheromone in an open field, the instantaneous wind direction usually pointed toward the odor source (cf. Elkinton et al., 1987; Brady et al., 1989).

The overall objective is to maneuver the vehicle to the odor source location in a resource-efficient (e.g., time, energy) manner. Different strategies will result from different weightings of the resources. For example, time optimality (e.g., an advantage in rapid mate finding) will result in higher energy usage than energy optimality. Also, the penalty or reward functions that are defined have a strong effect on the definition of the optimal strategy. For the autonomous vehicle plume-tracing application, the total time and energy expended in the search do not need to be minimized, but they do need to satisfy the mission time and vehicle energy capacity constraints. Also, the mission success is mainly evaluated by whether the vehicle ultimately and accurately determines the odor source location. Therefore, the vehicle plume-tracing application is closer to an organism finding a long-lasting resource, such as a plant, rather than being the first to locate a prospective mate. Strategy 3 developed herein is remarkably successful in finding an odor source (i.e., 97%) with a high accuracy (i.e., within 1 m)

within 300 s for a 100-m square search area. This is well within reasonable vehicle energy and mission time constraints. This performance was also robust across simulated air plumes and experimental water flume data. Verification of this performance awaits testing of these plume-following strategies with robots in natural environments.

Section 3 includes a heuristic definition of "in the plume." Various alternative definitions could be suggested. A main criterion for this definition is that the searcher itself must be able to evaluate this definition based on its recent sensory information. Therefore, time-average contours or out-of-plane images are not relevant. Also, a decision based on the instantaneous sensed concentration is overly sensitive, because turbulent diffusion is known to result in patches of above- and below-threshold chemical within the plume body with the pulse duration and pulse return distributions determined by the flow conditions. The definition of Section 3 was selected since it addressed the above issues and was efficiently computable by both engineered and biological entities.

We also have assumed a planar world in which the plume's (unspecified) height above ground is essentially unvarying. In nature, plumes undulate to some extent and, under some non-adiabatic (unstable) atmospheric conditions, may rise vertically (e.g., Fares, Sharpe & Maguson, 1980; Schal, 1986) or fall. Certainly flying insects must navigate plumes in three dimensions, whereas organisms such as lobsters and blue crabs sample the odor relatively close to the substrate on which they walk. Moths stabilize their flight altitude via an optomotor reaction, most likely with feedback from the equatorial region of the eyes (Vickers & Baker, 1994b). An example of the importance of the two-dimensional application, as described in Stacey et al. (2000), occurs when stratification of the fluid flow can result in an effectively two-dimensional search problem even in a three-dimensional environment.

Many features of the output of our three navigation strategies superficially resemble the general zigzag tracks of moths and other organisms orienting upwind along an odor plume. Even though we have explicitly designed these strategies to be biomimetic, in the sense that they use many described features of sensory perception and a navigational strategy of counterturning across the windline while heading upwind, we cannot be certain that we have employed identical causal mechanisms (Webb, 2000). This difficulty is illustrated by the "robolobster," a lobster biomimic that has been designed to track an odor plume in a water flow using artificial odor sensors that are positioned to match a lobster's chemosensory array. These robots successfully track an odor-scented plume upstream over several meters, but the precise form of the tracks and success rates of the robolobster and lobsters differ (Grasso et al., 2000). Such a divergence in behavioral outputs indicates that the robot and the lobsters may not be extracting information about odor flux in precisely the same way, that the real time control of movement by the lobster has not been replicated in the robot, or that other sensory cues, such as the direction of fluid movement, need to be incorporated into the animat's response.

Notes

1 Theoretically, an instantaneous definition of the plume extent could be defined as an area with a high probability of detecting an above-threshold concentration of odor. This area is time varying. Derivation of this area is extremely complex and would require the search vehicle to have knowledge of the odor source location and the fluid flow field over the entire search region for a significant range of past time. Therefore, this theoretically appealing definition of the plume is impractical for on-vehicle maneuvering.

2 A single simulation image is contained in this article. The time-varying simulation output is more interesting. Therefore, an executable version (for Windows) of the simulation described herein is available for download at http://www.ee.ucr.edu/~farrell. Source code is available by contacting the second author.

3 The vehicle heading dynamics are defined by $d\psi(t)/dt = -a\psi(t) + a\psi_v(t)$ for $a > 0$. In this article, $a = 5.0$. Therefore, ψ converges toward ψ_v as t increases. Throughout this article, ψ_v should be interpreted as a commanded reference heading. For a constant commanded heading, convergence is approximately complete after $(3/a)$ s.

4 Zigzagging is counterturning with significant upwind or downwind displacement (Kennedy, 1986).

5 Note that when the searcher leaves the plume its instantaneous heading will typically have a significant crosswind component. The reacquisition course command will be due crosswind in the opposite direction. If the searcher reverses its direction with a turn through upwind, then the searcher will have made upwind progress while out of the plume. Depending on the parameter settings and the plume

characteristics, it is possible that the searcher may make the majority of its upwind progress while it is outside the plume. In such a situation, the searcher would only be crossing the plume between upwind "sprints." The plume crossings may ensure that the upwind sprints are centered on the plume and have not passed the odor source. The issue of moths making significant progress toward the odor source while outside the plume has been discussed in the moth literature (David, Kennedy, & Ludlow, 1983; Ludlow, 1983).

Acknowledgments

This research was supported by a grant from the Office of Naval Research (ONR N00014-98-1-0820) under the ONR/DARPA Chemical Plume Tracing Program. We thank Todd Cowen, John Crimaldi, Jeffery Koseff, Phil Roberts, Don Webster, and Marc Weissburg for providing data from their flume experiments. The experimental flume data sets also were generated through support from the ONR/DARPA Chemical Plume Tracing Program. John Murlis provided valuable insight into the processes of plume dispersion.

References

Atema, J. (1995). Chemical signals in the marine environment: dispersal, detection, and temporal signal analysis. *Proceedings of the National Academy of Sciences of the United States of America*, *92*, 62–66.

Aylor, D. E. (1976). Estimating peak concentrations of pheromones in the forest. In J. F. Anderson & M. K. Kaya (Eds.), *Perspectives in forest entomology* (pp. 117–188). New York: Academic Press.

Aylor, D. E., Parlange, J.-Y., & Granett, J. (1976). Turbulent dispersion of disparlure in the forest and male gypsy moth response. *Environmental Entomology*, *5*, 1026–1032.

Baker, T. C. (1990). Upwind flight and casting flight: Complimentary phasic and tonic systems used for location of sex pheromone sources by male moths. In K. B. Døving (Ed.), *International Symposium on Olfaction and Taste X* (pp. 18–25). Oslo: Graphic Communications Systems.

Basil, J., & Atema, J. (1994). Lobster orientation in turbulent odor plumes: Simultaneous measurements of tracking behavior and temporal odor patterns. *Biological Bulletin*, *187*, 272–273.

Belanger, J. H., & Arbas, E. (1998). Behavioral strategies underlying pheromone-modulate flight in moths: Lessons from simulation studies. *Journal of Comparative Physiology A*, *183*, 345–360.

Belanger, J. H., & Willis, M. A. (1996a). Centrally-generated and reflexive control strategies in the adaptive behavior of real and simulated animals. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 155–162). Cambridge, MA: MIT Press.

Belanger, J. H., & Willis, M. A. (1996b). Adaptive control of odor-guided location: Behavioral flexibility as an antidote to environmental unpredictability. *Adaptive Behavior*, *4*, 217–253.

Belanger, J. H., & Willis, M. A. (1998). Biologically-inspired search algorithms for locating unseen odor sources. In *Proceedings of the IEEE International Symposium on Intelligent Control* (pp. 265–270). New Jersey: IEEE.

Bell, W. J. (1984). Chemo-orientation in walking insects. In W. J. Bell & R. T. Cardé (Eds.), *Chemical Ecology of Insects* (pp. 93–109). London: Chapman and Hall.

Brady, J. H., Gibson, G., & Packer, M. J. (1989). Odour movement, wind direction, and the problem of host finding by tsetse flies. *Physiological Entomology*, *14*, 369–380.

Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation, 2*, 14–23.

Cardé, R. T. (1984). Chemo-orientation: Flying insects. In W. J. Bell & R. T. Cardé (Eds.), *Chemical Ecology of Insects* (pp. 111–124). London: Chapman and Hall.

Cardé, R. T. (l986). Epilogue: Behavioural mechanisms. In T. Payne, M. Birch, & K. Kennedy (Eds.), *Mechanisms of insect olfaction* (pp. 175–186). Oxford: Claredon Press.

Cardé, R. T. (1996). Odour plumes and odour-mediated flight in insects. In *Olfaction in mosquito-host interactions* (pp. 54–70). CIBA Foundation Symposium. 200. Chichester, UK: Wiley.

Cardé, R. T., & Knols, B. G. J. (2000). Effects of light levels and plume structure on the orientation manoeuvres of male gypsy moths flying along pheromone plumes. *Physiological Entomology*, *25*, 141–150.

Cardé, R. T., & Mafra-Neto, A. (1996). Mechanisms of flight of male moths to pheromone. In R. T. Cardé & A. K. Minks (Eds.), *Insect pheromone research. New directions* (pp. 275–290). New York: Chapman and Hall.

Charlton, R. E., & Cardé, R. T. (1982). Rate and diel periodicity of pheromone emission from female gypsy moths (*Lymantria dispar*) determined with a glass-adsorption collection system. *Journal of Insect Physiology, 28*, 423–430.

Cowen, E. A., Chang, K.-A., and Liao, Q. (2001). A single-camera coupled PTV-LIF technique. *Experiments in Fluids, 31*, 63–73.

Crimaldi, J. P. & Koseff, J. R. (2001). High-resolution measurements of the spatial and temporal scalar structure of a turbulent plume. *Experiments in Fluids*, *31*, 90–102.

David, C. T., Kennedy, J. S., and Ludlow, A. R. (1983). Finding of sex pheromone source by gypsy moths released in the field. *Nature*, *303*, 804–806.

Deveza, R., Thiel, D., Russell, A., & Mackay-Sim, A. (1994). Odor sensing for robot guidance. *International Journal of Robotics Research, 13*, 232–239.

Devine, D. V., & Atema, J. (1982). Function of chemoreceptor organs in spatial orientation of the lobster *Homarus americanus*: Differences and overlap. *Biological Bulletin, 163*, 144–153.

Dindonis, L. L., & Miller, J. R. (1980). Host-finding behavior in onion flies, *Hylemia antiqua. Environmental Entomology, 9*, 769–772.

Dusenbery, D. B. (1992). *Sensory ecology: How organisms acquire and respond to information*. New York: Freeman.

Elkinton, J. S., & Cardé, R. T. (1983). Appetitive flight behavior of male gypsy moths (Lepidoptera: Lymantriidae). *Environmental Entomology, 12*, 1702–1707.

Elkinton J. S., Cardé, R. T., & Mason C. J. (1984). Evaluation of time-average dispersion models for estimating pheromone concentration in a deciduous forest. *Journal of Chemical Ecology, 10*, 1081–1108.

Elkinton, J. S., Schal, C., Ono, T., & Cardé, R. T. (l987). Pheromone puff trajectory and upwind flight of male gypsy moths in a forest. *Physiological Entomology, 12*, 399–406.

Fares, Y., Sharpe, P. J. H., & Magnuson, C. E. (1980). Pheromone dispersion in forests. *Journal of Theoretical Biology, 84*, 335–359.

Farrell, J. A., Murlis, J., Long, X., Li, W., & Cardé, R. T. (2002). Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environmental Fluid Mechanics 12*, 143–169.

Gifford, F. A. (1959). Statistical properties of a fluctuating plume dispersion model. In F. N. Frankiel & R. A. Shepard (Eds.), *Advances in geophysics: Vol. 6. Atmospheric diffusion and air pollution*. (p. 117). New York: Academic Press.

Grasso, F. W., Consi, T. R., Mountain, D. C., & Atema, J. (2000). Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges. *Robotics and Autonomous Systems, 30*, 115–131.

Hasler, A. D., & Scholz, A. T. (1983). *Olfactory imprinting and homing in salmon*. New York: Springer.

Ishida, H., Kagawa, Y., Nakamoto, T., Moriizumi, T. (1996). Odor-source localization in the clean room by an autonomous mobile sensing system. *Sensors and Actuators B, 33*, 115–121.

Ishida, H., Kobayashi, A., Nakamoto, T., & Moriizumi, T. (1999). Three-dimensional odor compass. *IEEE Transactions on Robotics and Automation, 15*, 251–257.

Ishida, H., Nakamoto, T., & Moriizumi, T. (1998). Remote sensing of gas/odor source location and concentration distribution using mobile system. *Sensors and Actuators B, 49*, 52–57.

Ishida, H., Nakamoto, T., Moriizumi, T., Kikas, T., & Janata, J. (2001). Plume-tracking robots: A new application of chemical sensors. *Biological Bulletin, 200*, 222–226.

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. San Diego: Academic Press.

Jones, C. D. (1983). On the structure of instantaneous plumes in the atmosphere. *Journal of Hazardous Materials, 7*, 87–112.

Justus, K. A. & Cardé, R. T. (2002). Flight behaviour of males of two moths, *Cadra cautella* and *Pectinophora gossypiella*, in homogeneous clouds of pheromone. *Physiological Entomology, 27*, 67–75.

Justus, K. A., Schofield, S. W., Murlis, J., & Cardé, R. T. (2002). Flight behaviour of *Cadra cautella* males in rapidly pulsed pheromone plumes. *Physiological Entomology, 27*, 58–66.

Kennedy, J. S. (1977). Olfactory responses to distant plants and other odor sources. In H. H. Shorey & J. J. McKelvey, Jr. (Eds.), *Chemical control of insect behavior* (pp. 67–91). New York: Wiley Interscience.

Kennedy, J. S. (1983). Zigzagging and casting as a preprogrammed response to wind-borne odour: A review. *Physiological Entomology, 8*, 109–120.

Kennedy, J. S. (1986). In T. Payne, M. Birch, & K. Kennedy (Eds.), *Mechanisms in insect olfaction* (pp. 175–186). Oxford: Clarendon Press.

Kennedy, J. S., Ludlow, A. R., & Sanders, C. J. (1980). Guidance of flying moths by wind-borne pheromone. *Nature, 288*, 475–477.

Kramer, E. (1996). A tentative intercausal nexus and its computer model on insect orientation in windborne pheromone plumes. In R. T. Cardé & A. K. Minks (Eds.), *Insect pheromone research. New directions* (pp. 232–247). New York: Chapman and Hall.

Kuenen, L. P. S., & Cardé, R. T. (1994). Strategies for recontacting a lost pheromone plume: Casting and upwind flight in the male gypsy moth. *Physiological Entomology, 19*, 15–29.

Kuwana, Y., Nagasawa, S., Shimoyama, I., & Kanzaki, R. (1999). Synthesis of the pheromone-oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors. *Biosensors and Bioelectronics, 14*, 195–202.

Larcombe, M. H. E., & Helsall, J. R. (1984). Robotics in nuclear engineering: Computer-assisted teleoperation in hazardous environments with particular reference to radiation fields. Graham and Trotman, for Commission of the European Communities.

Li, W. (1994). Fuzzy-logic-based reactive behavior of an autonomous mobile system in unknown environments.

Engineering Application of Artificial Intelligence, 7, 521–531.

Ludlow, A. R. (1983). *Applications of computer modelling to behavioural coordination*. Unpublished doctoral dissertation, University of London.

Mafra-Neto, A., & Cardé R. T. (1994). Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature, 369*, 142–144.

Mafra-Neto, A., & Cardé R. T. (1996). Dissection of the pheromone-modulate flight of moths using the single-pulse response as a template. *Experientia, 52*, 373–379.

Murlis, J., & Jones, C. D. (1981). Fine scales structure of odor plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiological Entomology, 6*, 71–86.

Murlis, J., Elkinton, J. S., & Cardé, R. T. (1992). Odor plumes and how insects use them. *Annual Review of Entomology, 37*, 505–532.

Murlis, J., Willis, M. A., & Cardé, R. T. (2000). Spatial and temporal structure of pheromone plumes in fields and forests. *Physiological Entomology, 25*, 211–222.

Mylne, K. R. (1992). Concentration fluctuation measurements in a plume dispersing in a stable surface layer. *Boundary-Layer Meteorology, 60*, 15–48.

Nevitt, G. A. (2000). Olfactory foraging by antarctic procellariiform seabirds: Life at high Reynolds numbers. *Biological Bulletin, 198*, 245–253.

Pielke, R. A. (1984). *Mesoscale meteorological modeling*. San Diego: Academic Press.

Preiss, R., & Kramer, E. (1986). Mechanism of pheromone orientation in flying moths. *Naturwissenschaften, 73*, 555–557.

Rau, P., & Rau, N. L. (1929). The sex attraction and rhythmic periodicity in the giant saturniid moths. *Transactions of the Academy of Science of Saint Louis, 26*, 82–221.

Rozas, R., Morales, J., & Vega, D. (1991). Artificial smell detection for robotic navigation. *Fifth International Conference on Advanced Robotics. Robots in Unstructured Environments, 2*, 1730–1733.

Rumbo, E. R., & Kaissling, K.-E. (1989). Temporal resolution of odour pluses by three types of pheromone receptor cells in *Antheraea polyphemus. Journal of Comparative Physiology A, 165*, 281–295.

Russell, R. A. (2001). Tracking chemical plumes in constrained environments. *Robotica, 19*, 451–458.

Russell, R. A., Thiel, D., Deveza, R., & Mackay-Sim, A. (1995). A robotic system to locate hazardous chemical leaks. *Proceedings of 1995 IEEE International Conference on Robotics and Automation, 1*, 556–561.

Schal, C. (1986). Interspecific vertical stratification as a mate-finding mechanism in tropical cockroaches. *Science, 215*, 1405–1407.

Stacey, M. T., Cowen, E. A., Powell, T. M., Dobbins, E., Monismith, S. G., & Koseff, J. R. (2000). Plume dispersion in a stratified, near-coastal flow: Measurements and modeling. *Continental Shelf Research, 20*, 637–663.

Sutton, O. G. (1953). *Micrometeorology*. New York: McGraw-Hill.

Vickers, N. J. (2000). Mechanisms of animal navigation in odor plumes. *Biological Bulletin, 198*, 203–212.

Vickers, N. J., & Baker, T. C. (1994a). Reiterative responses to single strands of odor promote sustained upwind flight and odor source location by moths. *Proceedings of the National Academy of Sciences of the United States of America, 91*, 5756–5760.

Vickers, N. J., & Baker, T. C. (1994b). Visual feedback in the control of pheromone-mediated flight of *Heliothis virescens* males (Lepidoptera: Noctuidae). *Journal of Insect Behavior, 7*, 605–632.

Webb, B. (2000). What does robotics offer animal behaviour? *Animal Behaviour, 60*, 545–558.

Webster, D. R., Roberts, P. J. W., & Ra'ad, L. (2000). Simultaneous DPTV/PLIF measurements of a turbulent jet. *Experiments in Fluids, 30*, 65–72.

Webster, D. R., & Weissburg, M. J. (2001). Chemosensory guidance cues in a turbulent chemical odor plume. *Limnology and Oceanography, 46*, 1034–1047.

Weissburg, M. J., & Zimmer-Faust, R. K. (1994). Odor plumes and how blue crabs use them in finding prey. *Journal of Experimental Biology, 197*, 349–375.

Willams, B. (1994). Models of trap seeking by tsetse flies: Anemotaxis, klinotaxis and edge detection. *Journal of Theoretical Biology, 168*, 105–115.

Willis, M. A., & Baker, T. C. (1984). Effects of intermittent and continuous pheromone stimulation on the flight behaviour of the oriental fruit moth, *Grapholita molesta. Physiological Entomology, 9*, 341–358.

Zimmer R. K., & Butman, C. A. (2000). Chemical signaling processes in the marine environment. *Biological Bulletin,* 198, 168–187.

Appendix A

Table A1 Notation used in the text

Parameter	*Units*	*Interpretation*
τ		Sensor detection threshold
(x,y)	m	Coordinates of searcher
λ	s	Time from last above threshold detection until the searcher begins the plume reacquisition behavior
T_{W}	s	Time from last above threshold detection until the searcher reverts to the plume finding behavior
T_{F},	s	Time to first detect the odor, depends on starting location and time
T_{M}	s	Time in contact with the plume
T_{R}	s	Total time to recontact the plume after its contact is lost
$c(t)$		Sensed concentration at time t
$\psi_u(t)$	Deg	Instantaneous direction of the wind plus 180º (i.e., upflow direction)
β	Deg	Angle of travel relative to $\psi_u(t)$ during the maintain plume behavior
γ	Deg	Angle of travel relative to $\psi_u(t)$ during plume reacquisition behavior
ϕ	Deg	Random variable
σ	Deg	Standard deviation
$[\underline{x},\overline{x}]$	m	Minimum and maximum x position
$[\underline{y},\overline{y}]$	m	Minimum and maximum y position
$\psi_r(t)$	Deg	$\psi_w(t) - \psi_v(t)$
$\psi_v(t)$	Deg	Commanded vehicle heading at time t
$\psi_w(t)$	Deg	Wind heading at time t at the vehicle location

Appendix B

The main body of this article focuses on the problem of designing a strategy to maintain contact with an intermittent plume. Some of the evaluations also require a strategy for the searcher to find the plume and a method to declare the successful completion of a search. Both of these issues are discussed in this appendix.

Plume Finding

The article considers plume tracing within a rectangular search area with corners defined by $(\underline{x},\underline{y})$, $(\underline{x},\overline{y})$, $(\overline{x},\underline{y})$, and $(\overline{x},\overline{y})$ The initial searcher location (x_0,y_0) is assumed to be located in the search area. The plume-finding behavior is defined by crosswind search within the rectangular search area and reflection off the boundaries of the search area when encountered. Initially, the vehicle heading is defined as

$$\psi_v(t) = \psi_w(t) + \theta_0 \,\mathrm{sign}((\underline{y} + \overline{y})/2 - y_0)$$

with θ_0 initialized to 70° where $\psi_w(t)$ is the wind direction. This results in initial crosswind flight toward the furthest y boundary. This heading is maintained until either odor is detected, in which case the strategy switches to the maintain behavior, or the edge of the

search region is encountered. If a y boundary is encountered (i.e., $x \in [\underline{x}, \overline{x}]$ and $y \notin [\underline{y}, \overline{y}]$), then the vehicle heading is defined by

$$\psi_v(t) = \psi_w(t) + \theta_{i+1}$$

(i.e., reflection reverses crosswind direction) where $\theta_{i+1} = -\theta_i$ and i is incremented by 1 at the time T_y at which the y boundary was encountered. If an x boundary is encountered (i.e., $x \notin [\underline{x}, \overline{x}]$ and $y \in [\underline{y}, \overline{y}]$), then the vehicle heading is defined by

$$\psi_v(t) = \psi_w(t) + \theta_{i+1}$$

where

$$\theta_{i+1} = 70 \text{ sign}(\theta_i) \quad \text{if } x < \underline{x}$$

(i.e., reflection to slightly downwind search) and

$$\theta_{i+1} = 110 \text{ sign}(\theta_i) \text{ if } x > \overline{x}$$

(i.e., reflection to slightly upwind search) and i is incremented by 1 at the time T_x at which the x boundary was encountered. This plume-finding behavior continues until the plume is found or the maximum simulation time is exceeded.

Source-Finding Declaration

The evaluations herein did not consider the problem of the searcher declaring that the source had been found. Instead, each search was declared to have successfully found the source when the distance from searcher to source,

$$\rho = \sqrt{(x - x_{\text{src}})^2 + (y - y_{\text{src}})^2}$$

was smaller than $R = 1$. When this occurs, the vehicle is considered to have reached the odor source and the simulation is stopped.

Because the searcher could wander randomly and meet this stopping criteria, the performance of such a random search was evaluated for comparison. The results of the random search are not included, because the random search takes so long that it is completely ineffective. (Except when the searcher starts very near the source, the simulation typically timed out unsuccessfully.)

About the Authors

Wei Li received B.S. (1982) and M.S. (1984) degrees in electrical engineering from the Northern Jiaotong University, Beijing, P.R. China, and Ph.D. (1991) in electrical and computer engineering from the University of Saarland, Germany. He was a faculty member in computer science and technology at Tsinghua University (1993–2001). In 2001, he became associate professor of computer science at California State University, Bakersfield. Dr. Li received the 1995 National Award for Outstanding Postdoctoral Researcher in China and the 1996 Award for Outstanding Young Researcher at Tsinghua University. He was a Croucher Foundation Research Fellow (1996) at City University of Hong Kong and an Alexander von Humboldt Foundation Research Fellow at the Technical University of Braunschweig, Germany (1997–1998). His research interests are intelligent systems, robotics, fuzzy logic control and neural networks, multisensor fusion and integration, and graphical simulation. *Address*: Department of Electrical Engineering, University of California at Riverside, Riverside, CA 92521, USA. E-mail: wli@cs.csubak.edu

Jay A. Farrell received B.S. degrees (1986) in physics and electrical engineering from Iowa State University, and M.S. (1988) and Ph.D. (1989) degrees in electrical engineering from the University of Notre Dame. At Charles Stark Draper Lab (1989–1994), he was principal investigator on projects involving intelligent and learning control systems for autonomous vehicles. Dr. Farrell received the Engineering Vice President's Best Technical Publication Award in 1990, and Recognition Awards for Outstanding Performance and Achievement in 1991 and 1993. He is a professor and former chair of the department of electrical engineering at the University of California, Riverside. His research interests include identification and on-line control for nonlinear systems, integrated GPS/INS navigation, and artificial intelligence techniques for autonomous vehicles. He is author of the book *The Global Positioning System and Inertial Navigation* (McGraw-Hill, 1998) and over 85 additional technical publications.

Ring Cardé holds a B.S. degree (1966) in biology from Tufts University and M.S. (1968) and Ph.D. (1971) degrees in entomology from Cornell University. He is Distinguished Professor of Entomology and holds the Boyce Chair in Entomology at the University of California, Riverside. His research group focuses on chemical communication in insects, with particular attention to the orientation maneuvers that flying insects use to locate odor sources. These studies include mate finding by male moths using pheromones and host location by female mosquitoes. His work covers manipulation of conditions in laboratory wind tunnels to determine how flight maneuvers are modulated by visual and wind cues and the odor plume's spatial distribution. Other studies underway examine the patterns of mating behaviors, the genetic architecture of variation in pheromone production and response, and the turbulent diffusion of odors. *Address*: Department of Entomology, University of California at Riverside, Riverside, CA 92521, USA. E-mail: ring.carde@ucr.edu