2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

A Sensor Fusion Approach to Odor Source Localization Inspired by the Pheromone Tracking Behavior of Moths

Adam J. Rutkowski, Roger D. Quinn., and Mark A. Willis

Abstract—**An approach to odor source localization with an aerial vehicle using the fusion of odor sensors, visual sensors, and airspeed sensors is presented. The motion of the tracking vehicle is decomposed into two components – a component normal to the wind direction and a component tangential to the wind direction. The tangential component is controlled with a strategy that moves upwind when odor is detected and moves gradually downwind when odor is lost. The normal component of velocity is controlled by two different algorithms. The first algorithm controls the rate of turning in the plane normal to the wind direction as a function of concentration. The second algorithm controls the rate of turning and the magnitude of the normal component of velocity as a function of the time derivative of concentration. Both algorithms display the potential to declare the location of an odor source in a three-dimensional space.**

I. INTRODUCTION

An aerial vehicle capable of tracking an odor to its source in a turbulent environment could be used to locate anything emitting an odor, such as lost pets and people, pipeline breaks, illegal drug labs, or improvised explosive devices. One advantage of using an aerial vehicle for these tasks is that the movement of the vehicle is not disturbed by obstacles on the ground. Also, a flying vehicle can track a plume that does not remain close to the ground. Tracking an odor to its source is a challenging problem. A strategy of simply moving in the direction of increasing odor concentration does not work because a natural odor plume consists of packets of odor created by turbulent eddies, thus no smooth gradient of chemical concentration exists [1][2].

Studies of the pheromone tracking behavior of the tobacco hornworm moth in a laboratory wind tunnel have proven beneficial for designing odor tracking algorithms. This behavior is studied by placing a pheromone source at the upwind end of the test section of the wind tunnel and releasing a male moth at the downwind end. The behavior is recorded with video cameras from above and from downwind. A digital three-dimensional reconstruction of the flight path is then obtained from the video recordings. Since this work makes reference to both biological and engineered odor tracking systems, the term 'tracker' is used to describe an odor tracking agent in general, either a vehicle or an animal, real or simulated.

The strategy that moths use to locate a pheromone source is called odor-modulated anemotaxis. Anemotaxis refers to the strategy of flying upwind, and odor modulation is a general term referring to the behavior of the moth in response to pheromone. Moths do not travel directly upwind when tracking pheromone. Instead, they counter-turn across the wind in search of pheromone while generally making upwind progress. The elapsed time between successive turns is called the inter-turn duration (ITD). For the tobacco hornworm moth, the average ITD in the horizontal and vertical directions is roughly 500 ms [3]. Although the average ITDs in the vertical and horizontal directions are similar, the temporal relationship between vertical and horizontal turns is not fixed (unpublished observations).

A behavior similar to the counter-turning behavior observed in moths can be achieved by controlling the turn rate of the component of velocity normal to the wind direction [4]. This strategy is called chemo-klino-kinesis, where the term 'kinesis' refers to a scalar response to a stimulus (as opposed to a directed orientation response) and the prefix 'klino' refers to a turning response. The motion of the tracking vehicle is decomposed into two components – a component normal to the wind direction (v_n) and a component tangential to the wind direction (v_t). Using this strategy, vertical and horizontal motion is coupled, but upwind motion can be controlled independently. This approach was tested in simulation but control of the tangential component of velocity was not explored [4]. Also, the problem was not treated as a multi-sensory problem. It was assumed that the tracker had absolute knowledge of its position in space and the wind direction.

In this work, we build on the chemo-klino-kinesis strategy for odor tracking by using a sensor fusion approach to position and wind direction estimation. Also, we present a control law for upwind motion that is designed to produce behavior that is similar to moths tracking pheromone plumes. In addition, we refine the processing of odor concentration information in an effort to make more effective use of this information. Furthermore, we formulate an odor tracking strategy that uses a variable magnitude of

Manuscript received September 15, 2006. This work was supported in part by the U.S. Air Force under grant F08630-03-1-0003 and by the Ohio Space Grant Consortium.

Adam J. Rutkowski is with the Department of Mechanical and Aerospace Engineering, Case Western Reserve University, Cleveland, OH 44106 USA (email: adam.rutkowski@case.edu).

Roger D. Quinn is with the Department of Mechanical and Aerospace Engineering, Case Western Reserve University, Cleveland, OH 44106 USA (email: rdq@case.edu).

Mark A. Willis is with the Department of Biology, Case Western Reserve University, Cleveland, OH 44106 USA (email: maw27@case.edu).

v_n and observes motion constraints. Finally, we discuss ideas for declaring the location of the odor source.

Tracking an odor to its source requires input from chemical sensors (to detect the odor), visual sensors, and wind sensors (to determine the direction of the wind relative to the tracker). Moths use their antenna to detect odors, their eyes to detect motion, and wind sensitive hairs along their antenna to detect the wind. Experiments have shown that moths are not capable of tracking pheromone in the absence of a visual pattern, thus vision plays an important role in odor tracking [5].

Determining the wind direction is non-trivial when suspended in the air. Sensors onboard the tracker can determine airspeed directly, but airspeed is a combination of groundspeed and the wind velocity. To estimate wind velocity, airspeed sensory information must be fused with visual information. This can be done using an aero-optical egomotion and wind velocity estimation technique [6].

The upwind control law developed in this work is inspired by the casting behavior exhibited by a moth tracking a pheromone plume. Casting is a response to loss of odor described as spiraling with a path of increasing circumference in a plane normal to the wind direction while making minimal (or even negative) upwind progress [7]. This change in behavior has a latency of about 300 ms [7]. We develop strategies that attempt to create a similar behavior.

It is common in the odor tracking literature to implement binary odor sensors [8][9]. These sensors only provide information about whether or not odor is present. In this work, we experiment with a simulated odor sensor that responds linearly to the odor concentration level.

Declaring the location of an odor source is a problem that is often ignored by odor tracking researchers. In simulation or in practice, odor tracking is terminated once the tracker reaches a certain distance from the odor source [9][10]. It is assumed that the target emitting the odor can be identified, perhaps visually, once close enough. When close to the odor source, moths continue to hover near but not necessarily directly in front of the odor source (unpublished observations). In this work, we do not assume that the odor source can be identified visually and we allow tracking to continue even if the tracker makes contact with the odor source.

II. Methods

A. Simulation Overview

The odor tracking algorithms presented in this work are tested in simulation using a vehicle with perfect motion tracking (i.e. no vehicle dynamics were modeled). The wind speed and direction are held constant and the tracker remains pointed directly upwind at all times. The tracker is capable of moving perpendicular to the wind while in this orientation and can also hover or fly backward (like a moth). Control decisions are performed at 10 Hz and collision with the odor source is ignored.

Three right-handed Cartesian coordinate systems are used as reference frames (Figure 1). The first reference frame, the odor-fixed frame, is fixed to the odor source and aligned such that the positive x-axis points upwind and the positive y-axis points opposite the direction of gravity. The primary purpose of the odor fixed frame is for plotting the path of the tracker relative to the odor source and for calculating the odor concentration of a simulated odor plume. The second reference frame, the body-fixed frame, is fixed to the body of the tracker. Since the tracker always remains pointed directly upwind, this coordinate system is always parallel to the odor-fixed frame but can translate relative to the odor-fixed frame. The third reference frame, the wind-directed frame, is aligned with the *estimated* wind direction. Ideally, since the wind direction is fixed and the tracker points directly upwind, the wind-directed frame would be the same as the body-fixed frame. However, the wind direction cannot be measured exactly from an aerial vehicle, thus the wind-directed frame will be slightly different than the body-fixed frame. The wind directed frame is constructed so that the x-axis opposes the estimated wind velocity vector and the z-axis is parallel to the ground (or perpendicular to gravity in the case of a non-level ground).

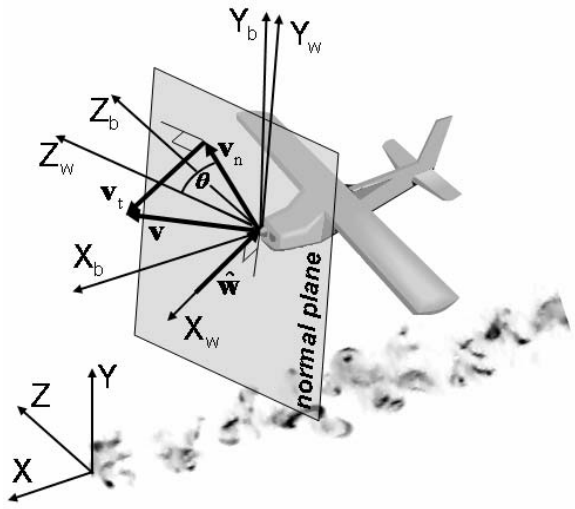

Figure 1. Three reference frames are used in the odor tracking simulation. The first reference frame is the odor-fixed frame (X,Y,Z), the second is the body-fixed frame (X_b,Y_b,Z_b), and the third is the wind-directed reference frame (X_w,Y_w,Z_w)

B. Odor Plume Model

Since our algorithms are being tested in simulation, it is necessary to construct a simulated odor plume that is similar to a real odor plume. Real turbulent odor plumes have the following characteristics.

1. The "width" of the odor plume increases with increasing downwind distance from the odor source

2. The average odor concentration decreases gradually with increasing distance from the plume centerline
3. The odor plume is discontinuous in nature at any instant in time, but is continuous when averaged over time.

In a previous study, the *probability of odor detection* in the normal plane followed a bivariate normal distribution [4]. In this study, since we wish to utilize odor sensors with a linear response to concentration instead of a binary response, the bivariate normal distribution is used to simulate the *average odor concentration*, $\bar{c}$. At the odor source, the plume has an average concentration of $\bar{c}_0$ and a standard deviation of σ_0. The standard deviation of the average odor concentration, σ, increases linearly by a factor of a with increasing downwind distance from the odor source. To satisfy conservation of mass, the average odor concentration along the centerline decreases by a factor of $(\sigma / \sigma_0)^2$. To model the discontinuous nature of the plume, the odor concentration at any point in the normal plane follows a Weibull distribution that has a mean given by $\bar{c}$. A Weibull distribution was chosen because this distribution produces only positive odor concentration values. This behavior is qualitatively desirable since a negative odor concentration is nonsense. The plume is formulated using the set of equations given in (1), where the function $U(0,1)$ generates a uniformly distributed pseudo-random number between 0 and 1 on each time step, and r is the normalized distance from the odor plume centerline.

$$\begin{aligned} \sigma &= \sigma_0 - ax \\ r^2 &= (y^2 + z^2)/\sigma^2 \\ \bar{c}(x,y,z) &= (\bar{c}_0 \sigma^2 / \sigma_0^2) \exp\left(-r^2/2\right) \\ c(x,y,z) &= \begin{cases} 2\bar{c}\sqrt{-\log(U(0,1))/\pi} & \text{if x<0} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{1}$$

The plot on the left side of Figure 2 is representative of the normalized odor concentration level at any given point in time. The plot on the right side of Figure 2 shows the concentration levels averaged over 1000 samples at each location. This model satisfies the characteristics of a natural plume that are important for testing odor tracking strategies.

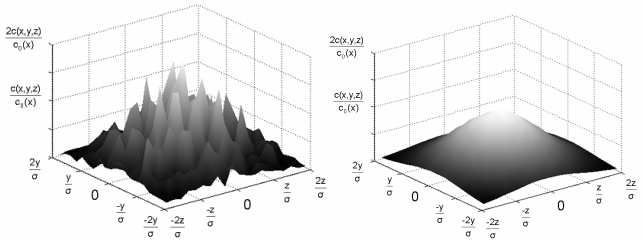

Figure 2. Representative plots of simulated normalized concentration levels at an instant in time (left) and averaged over 1000 samples (right).

The numerical values of the odor plume concentration constants are arbitrary for the simulation since the response of a real odor sensor to a natural odor plume can be rescaled. The physical dimensions were chosen to be similar to those of the plume generated in moth experiments (unpublished). The values of $\bar{c}_0 = 15$, $\sigma_0 = 0.02$ meters, and $a = 0.05$ were used.

C. Egomotion and Wind Velocity Estimation Algorithm

The velocity of the tracker relative to the ground and the wind direction are simultaneously measured using an aero-optical egomotion estimation algorithm [6]. This algorithm uses the fusion of airspeed and optical flow to estimate the height of the tracker above the ground. Once the height is estimated, the Cartesian components of the groundspeed vector $(\hat{v}_x, \hat{v}_y, \hat{v}_z)$ are determined directly from the optical flow of the ground beneath the tracker. The wind velocity vector, with components $(\hat{w}_x, \hat{w}_y, \hat{w}_z)$, is the difference between the groundspeed vector and the airspeed vector.

D. Control of Upwind Motion

A simple algorithm controls the upwind component (v_t) of the velocity of the tracker. The tracker is commanded to progress upwind if odor has been detected recently (within 300 ms – the behavioral latency of a moth to the loss of odor) at a rate of v_t=300 mm/s (the average upwind velocity of a moth). Otherwise, the tracker moves slowly downwind (v_t=-70 mm/s). A minimum odor concentration threshold ($c_{threshold}$=0.2) is used to determine the presence or absence of odor.

E. Control of Motion in Normal Plane

The rate of turning in the normal plane, ψ, is controlled as a function of odor concentration, a strategy known as chemo-klino-kinesis (CKK) [11]. The parameter ψ is the time derivative of the angle θ from Figure 1. We adapted the original chemo-klino-kinesis algorithm presented in [4] such that the turn rate in the normal plane is a continuous function of concentration. The turn rate is set to 1 rad/s if the odor concentration is at or below $c_{threshold}$, 6 rad/s if the odor concentration is at or above a saturation level ($c_{saturation}$=1.0), and a linear function of concentration if the concentration is between $c_{threshold}$ and $c_{saturation}$. In CKK the magnitude of the normal component of velocity is held constant at 300 mm/s.

A strategy that controls the magnitude of v_n was also tested. This strategy is called chemo-ortho-klino-kinesis (COKK), where the prefix 'ortho' refers to a functional relationship between speed and stimulus. This strategy controls the turn rate as a function of the *time derivative* of the concentration with the hope that the algorithm will be more adaptive to varying concentration levels. This is important since the tracker will be moving upwind and the measured concentration levels will in general be increasing. Four design criteria for a function to control the magnitude of the turn rate were identified.

1. Decrease the turn rate if the odor concentration is increasing ($\dot{c}$ >0)
2. Increase the turn rate if the odor concentration is decreasing ($\dot{c}$ <0)
3. Decrease the turn rate gradually if the odor concentration is constant ($\dot{c}$ = 0).
4. The magnitude of the turn rate is always positive.

Criteria 1 and 2 seem counter-intuitive at first, but are logical if one thinks that a decrease in odor concentration causes an increase in the turn rate, thereby directing the tracker back to an area of higher concentration. Also, an increase in the concentration causes the tracker to decrease its turn rate, thereby allowing the tracker to traverse the normal plane in search of areas of higher concentration. Criterion 3 is based on the casting response observed in moths tracking pheromone plumes. If the moth loses contact with the odor plume, it will spiral in the normal plane with an increasing radius of curvature, or a decreasing turn rate. This allows the tracker to reacquire the odor plume if it has been lost. Criterion 4 allows the magnitude and sign of the turn rate to be calculated separately. The magnitude is controlled by some function of $\dot{c}$ and the directional sense (positive or negative) is controlled by an estimate of the odor source location (this will be discussed later).

The function to control the magnitude of ψ is inspired by a modified spring-mass-damper system with a forcing term. In the spring-mass-damper system, the position of the mass is influenced by an external force, elastic and viscous forces, and inertia. The forcing term, *F*, is analogous to $-\dot{c}$, and the position, *x*, is analogous to ψ. Requirement 3 is satisfied by a spring with stiffness *k*. If the mass, with a mass of m, is released from rest at a certain distance above the floor, the spring will act to pull the mass toward the floor (decrease ψ). Requirement 4 is satisfied by adding a hard stop at a height x_{min} above the floor. This keeps the position of the mass positive. In our system, we can set some lower bound on the value of ψ. A linear viscous term, η_{lin}, reduces oscillation of the system. An additional nonlinear viscous term, η_{inv}, is added to gradually increase damping as the turn rate approaches the lower bound on ψ. The hard stop keeps the mass from hitting the floor where the damping is infinite. The control law for ψ is described by the differential equation in (2) and the turn rate in response to a test function for the differential odor signal is shown in Figure 3. It can be seen that this function meets the design criteria for a function to control the turn rate.

$$m\ddot{\psi} + \left(\eta_{lin} + (\eta_{inv} / \psi)\right)\dot{\psi} + k\psi = -\dot{c} \tag{2}$$

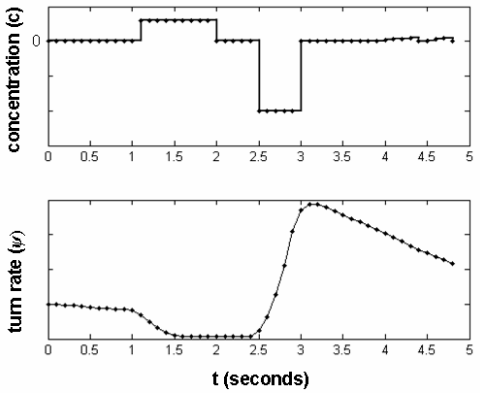

Figure 3. Turn rate in response to a test function for the differential concentration.

The tracker is directed to turn toward the estimated source location using (3), where $(\hat{y}_s, \hat{z}_s)$ is the estimated source location. The desired magnitude of the normal component of velocity is computed to achieve maximum acceleration given a maximum attainable acceleration ($a_{n,max}$) and a maximum attainable velocity ($v_{n,max}$) using (4). The turn rate is integrated to calculate a desired angle, θ, for the velocity vector relative to the z-axis of the wind-directed frame. Finally, the desired velocity vector relative to the body frame is computed using (5), where $\mathbf{C}_{bw}$ is a coordinate transformation matrix from the wind directed frame to the body fixed frame.

$$\text{sign}(\psi(t_i)) = \begin{cases} -1 \text{ if } \left(\hat{v}_y(t_i)\hat{z}_s(t_i) - \hat{v}_z(t_i)\hat{y}_s(t_i)\right) < 0 \\ 1 \text{ otherwise} \end{cases} \tag{3}$$

$$v_n = \min(a_{n,\max} / \psi, v_{n,\max}) \tag{4}$$

$$\mathbf{v} = \mathbf{C}_{bw}\begin{bmatrix} v_t \\ -v_n \sin(\theta) \\ v_n \cos(\theta) \end{bmatrix} \qquad \theta = \int_0^{t_i} \psi dt \tag{5}$$

F. Source Location Estimation

The location of the odor source projected into the normal plane can be estimated using the odor concentration information. One method of estimating the projected odor source location is to mark the position in the normal plane where the highest odor concentration was measured. Another method is to use the average position of previously scanned locations in the normal plane weighted by the odor concentration measured at those positions. The latter method was used in [4], however, the formulation was performed in the odor-fixed reference system.

The influence of previous odor samples is decreased gradually over time using a decay factor, a_{decay}, with a value slightly less than one (a_{decay}=0.95 was used). This minimizes the effect of propagated errors in the estimated motion. The odor is sampled with two odor detectors separated by a

distance d_s=30 mm (the approximate antennal separation of the tobacco hornworm moth). The odor concentration measured by the left antenna is c_{left} and the odor concentration measured by the right antenna is c_{right}. The left antenna has a horizontal position relative to the body of $z_{left}=-d_s/2$ and the right antenna has a horizontal position of $z_{right}=d_s/2$. A transformation that magnifies the effect of higher concentration values is constructed using an exponent p=2 on the concentration values. The final form of the odor source location estimate is given in (6).

$$\hat{y}_s(t_i)=\frac{c_{sum}(t_{i-1})\left(\hat{y}_s(t_{i-1})-\hat{v}_y(t_i)\Delta t\right)}{c_{sum}(t_i)}$$

$$\hat{z}_s(t_i)=\frac{c_{sum}(t_{i-1})\left(\hat{z}_s(t_{i-1})-\hat{v}_z(t_i)\Delta t\right)+c_{right}(t_i)^p z_{right}(t_i)+c_{left}(t_i)^p z_{left}(t_i)}{c_{sum}(t_i)} \quad (6)$$

$$c_{sum}(t_i)=a_{decay}c_{sum}(t_{i-1})+c_{right}(t_i)^p+c_{left}(t_i)^p$$

III. Results

Flight tracks of the CKK and the COKK algorithms were plotted with respect to the odor-fixed reference frame and inspected for qualitative characteristics. Figure 4 shows a typical flight track resulting from the CKK algorithm, and Figure 5 shows the flight track as viewed from downwind. The estimated source location is also shown. It was found that the estimated source location tends to drift more when the aero-optical algorithm is used than when perfect knowledge of position and velocity is available [4]. Figure 6 shows the horizontal position of the tracker and height above the floor with respect to time. The estimated height above the floor, as calculated using the aero-optical egomotion estimator, is also shown. In general, the estimated height is in good agreement with the true height of the tracker. The tracker reaches the upwind position of the odor in about 7 seconds, and then tends to oscillate upwind and downwind near the odor source while continuing to survey the plane normal to the wind direction.

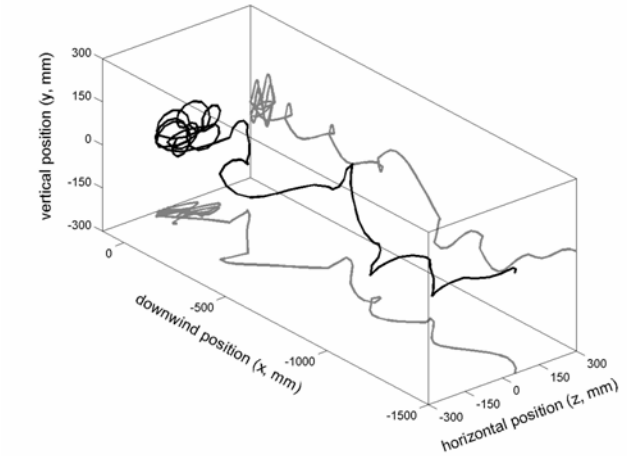

Figure 4. Three dimensional flight track representative of the COKK algorithm

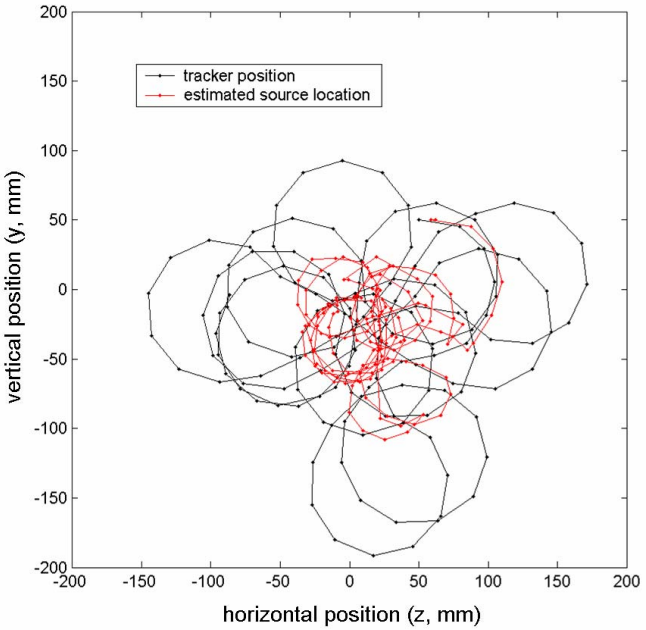

Figure 5. Downwind view of a typical flight track obtained using the CKK algorithm

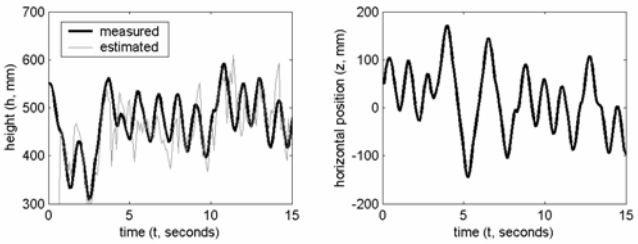

Figure 6. Height of the tracker above the ground (left) and horizontal position in the odor-fixed reference frame (right) representative of the CKK algorithm.

Figure 7 shows a plot of the behavior that is typical of the COKK algorithm, and Figure 8 shows the flight track as viewed from downwind. The behavior of the tracker appears to be more natural than the CKK algorithm until the tracker passes the source. At this point, the sudden drop in odor concentration triggers a dramatic increase in the turn rate, yet the directed turning response toward the estimated source location creates a very jagged back and forth turning motion. Figure 9 shows the horizontal position and height above the floor with respect to time. The counter-turning behavior is generally much less regular with the COKK algorithm than with the CKK algorithm. The same near-source oscillatory behavior in the downwind direction that was observed with the CKK algorithm was also observed with the COKK algorithm. In this case, the estimated height was also in good agreement with the true height of the tracker, also a general trend of the algorithm. This is typical of most of the simulation results, but there was one case where the tracker became unstable and lost the odor source after 30 seconds.

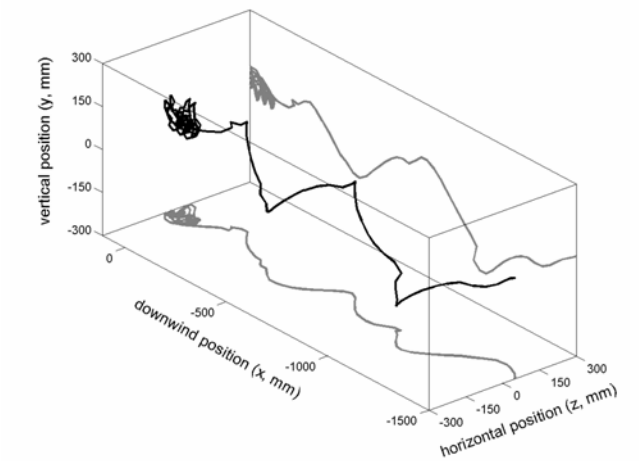

Figure 7. Three dimensional flight track representative of the COKK algorithm

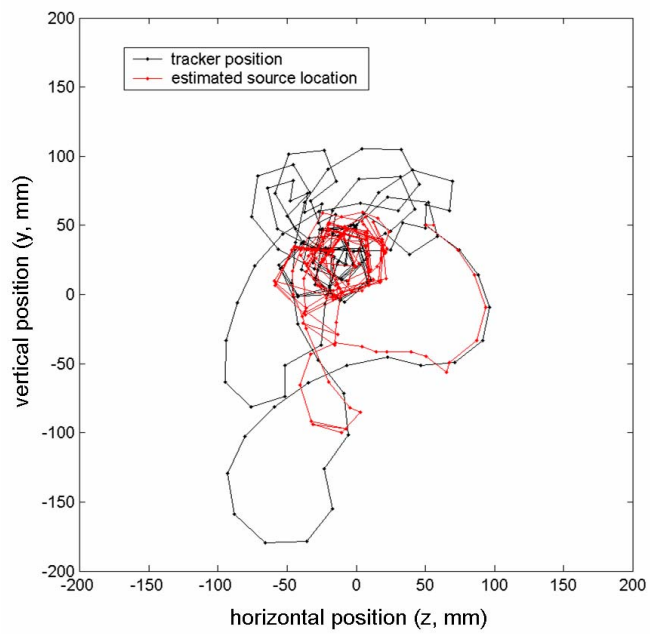

Figure 8. Downwind view of a typical flight track obtained using the COKK algorithm

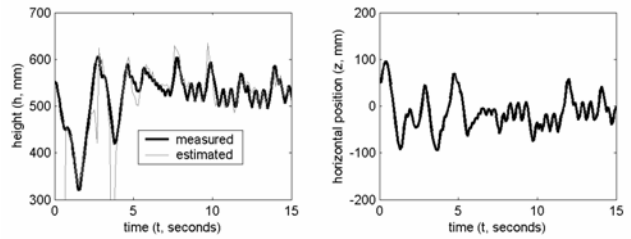

Figure 9. Height of the tracker above the ground (left) and horizontal position in the odor-fixed reference frame (right) representative of the COKK algorithm.

IV. Conclusions

A sensor fusion approach to odor tracking was developed and two odor tracking algorithms were described and compared – the chemo-klino-kinesis (CKK) algorithm and the chemo-ortho-klino-kinesis (COKK) algorithm. Both algorithms used the aero-optical egomotion estimation technique to estimate the wind direction. The agreement between the true height and the estimated height was good for both algorithms. This is an encouraging result since the quality of the wind direction estimate and the source location estimate depend on the accuracy of the height estimate.

It was found that the COKK algorithm was more difficult to work with than the CKK algorithm because of the number of parameters involved. However, the COKK algorithm tended to estimate the odor source location more accurately. A full analysis using quantitative performance measures will be performed in the future to compare these two algorithms in more detail.

Results from both algorithms suggest a technique for declaring the location of the odor source. Both the CKK and the COKK algorithms tended to oscillate in the direction of the wind when near the source, but this behavior was not observed farther downwind of the source. A method of detecting this oscillation in motion in the direction of the wind, coupled with the odor source location estimate in the plane normal to the wind direction, could be used to declare the location of the source. This idea will be explored in the future.

References

[1] J. S. Elkinton and R. T. Cardé, "Odor Dispersion," in Chemical Ecology of Insects, W. J. Bell and R. T. Carde, Eds. Sunderland, MA: Sinauer, 1984.

[2] J. Atema, "Distribution of chemical stimuli," in Sensory Biology of Aquatic Animals, J. Atema, R. R. Fay, A. M. Popper, and W. N. Tavolga, Eds. New York: Springer-Verlag, 1987.

[3] A. E. Arbas, M. A. Willis, and R. Kanzaki, "Organization of goal-oriented locomotion: Pheromone-modulated flight behavior of moths," in Biological neural networks in invertebrate neuroethology and robotics, R. D. Beer, R. E. Ritzmann, and T. McKenna, Eds. San Diego: Academic Press, 1993.

[4] A. J. Rutkowski, M. A. Willis, and R. D. Quinn, "Simulated Odor Tracking in a Plane Normal to the Wind Direction," IEEE ICRA conference proceeding, 2006

[5] N. J. Vickers, T. C. Baker, "Visual Feedback in the Control of Pheromone-Mediated Flight of Heliothis virescens Males (Lepidoptera: Noctuidae)", Journal of Insect Behavior, Vol. 7, no. 5 (1994) 605-632

[6] A. J. Rutkowski, R. D. Quinn, and M. A. Willis, "Biologically Inspired Self Motion Estimation using the Fusion of Airspeed and Optical Flow", American Control Conference Proceedings, 2006.

[7] N. J. Vickers, T. C. Baker, "Latencies of behavioral response to interception of filaments of sex pheromone and clean air influence flight track shape in Heliothis virescens (F.) males", J Comp Physiol A (1996) 178: 831-847

[8] R. A. Russell, Odour Detection by Mobile Robots, World Scientific Publishing, 1999.

[9] J. Farrell, W. Li, S. Pang, and R. Arrieta, "Chemical Plume Tracing Experimental Results with a REMUS AUV", MTS/IEEE Oceans 2003, San Diego, CA.

[10] J. H. Belanger, E. A. Arbas, "Behavioral Strategies underlying pheromone-modulated flight in moths: lessons from simulation studies", J Comp Physiol A (1998) 345-360.

[11] G. Fraenkel and D. Gunn, The Orienation of Animals: Kineses, Taxes, and Compass Reactions, Dover Publications Inc., New York, New York, 1961.