Biologically Inspired Direct Adaptive Guidance and Control for Autonomous Flight Systems

Final Report
September 30, 2004

Report Number GST-04-9.1

Reporting Period 1/1/02- 9/30/04

SBIR Phase II
Contract F49620-02-C-0017

Prepared for:

Lt Col Sharon Heise, USAF, PhD
Deputy Director, AFOSR/NM
Program Manager, Dynamics & Control
4015 Wilson Blvd, Rm 713
Arlington, VA 22203
Phone: 703-696-7796
Email: sharon.heise@afosr.af.mil

Prepared by:

J. Eric Corban, Cole Gilbert, Anthony J. Calise, Allen R. Tannenbaum

Guided Systems Technologies, Inc.
P.O. Box 1453
McDonough, Georgia 30253-1453
corban@mindspring.com
(770) 898-9100, x832

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-05-

0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
		01-January 2002 - 30-June 2004

4. TITLE AND SUBTITLE
Biologically Inspired Direct Adaptive Guidance and Control for Autonomous Flight Systems

5. FUNDING NUMBERS
F49620-02-C-0017

6. AUTHOR(S)
J. Eric Corban

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
DCMA Atlanta
805 Walker Street Suite 1
Marietta, GA 30060-2789

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 North Randolph Street
Suite 325, Room 3112
Arlington, VA 22203

NM

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The work at Cornell centered on developing experimental methods to characterize flesh fly pursuit evasions, and resulted in the maturation of effective means to capture the 3-D trajectory, as well as body and head orientation. The data was processed at first by hand, and later using image processing algorithms to develop 3-D visualizations at the track, including the head orientation, and ultimately to map the location of the target on the eye during the pursuit. The results provided a means to compare the guidance strategy of the fly with traditional proportional navigation, and to look for inspiration in the development of new guidance laws. Work was also completed to introduce clutter into the encounter. While a much greater understanding of the tracking and guidance strategy of the flesh fly was developed and documented, the work has not yet resulted in the discovery of a better alternative to traditional engineered guidance laws.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	46
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

ACKNOWLEDGEMENTS

The work presented herein was funded by the US Air Force Office of Scientific Research under STTR Phase II Contract F49620-02-C-0017 and coordinated with AFRL/MNG at Eglin AFB. The phase I proposal was submitted in response to the 2000 Air Force STTR solicitation topic A00T012 entitled "Implementation of Biomimetic Precision Flight in Autonomous Vehicles." The work reported herein depends greatly on previous research in many diverse areas, and the contributions of the team members and their previous research sponsors in each of these areas is gratefully acknowledged. The phase I and II experimental work also benefited greatly from the commercial products being developed by Dr. P.O. Zanen of Synceros, Inc. in association with the Cornell Laboratory. This support is greatly appreciated.

The invaluable support of graduate students Venky Madyastha and Suo Jin of Georgia Tech and John Layne of Cornell University is also gratefully acknowledged.

TABLE OF CONTENTS

A. INTRODUCTION AND BACKGROUND

A.1 PROBLEM/OPPORTUNITY

Rapid advances in many technology areas will soon make possible small or micro air vehicles that can be deployed in large numbers at low cost. In theory, military and commercial applications of such vehicles abound. However, in practice these vehicles must achieve a high level of autonomy in operations to be successfully deployed in large numbers. Furthermore, much of their utility is to be derived from operations in densely cluttered environments that include unknown obstacles. Traditional guidance and control technology is inadequate to meet the demand for autonomy in such environments, and dependence on traditional design methods will lead to expensive solutions that demand tremendous computational resources. There are, however, a multitude of biological systems that, with relatively crude sensors and limited computing resources, operate very successfully in such environments. This program seeks to emulate the success of these biological systems in an engineered guidance and control system. This is to be accomplished by developing a better understanding of the tracking, guidance and control functions of a candidate biological system (the flesh fly), and to then use this understanding to inspire, develop and demonstrate effective and affordable guidance and control technologies for autonomous operation of small or micro air vehicles.

A.2 STATED GOALS

The STTR program solicitation issued by the Air Force called for development and demonstration of biologically inspired guidance, navigation and control (GNC) sensors, components, and/or systems that will enable an autonomous small air vehicle to search for, detect, pursue, and rendezvous with an evasive target in a densely cluttered environment.

As stated in the solicitation, desirable characteristics of the GNC system include:

- Tolerance to transient sensor information distortion or obscuration;
- Capable of obstacle avoidance (e.g. buildings or trees, power lines, or overhanging branches);
- Detection, acquisition, tracking and guidance to a moving target in background clutter;
- Precision rendezvous with the target (e.g. warhead event, electronic tagging);
- Flight control robustness to wind gusts and turbulence near structures.

A.3 PROGRAM SUMMARY AND TECHNICAL OBJECTIVES

Researchers in diverse fields from Guided Systems Technologies, Inc. (GST), Cornell Universtiy, and the Georgia Institute of Technology were assembled into an STTR team. The phase I program demonstrated the feasibility and merit of combining university research in the

fields of entomology, target tracking, aerospace vehicle guidance and control, as well as image processing and computer vision, with the practical systems integration and UAV capabilities of GST.

Key objectives of the phase II program were:

- Continue the phase I experimental effort to fully characterize guidance, sensory tracking and control of oriented movements of the flesh fly in pursuit/evasion scenarios. The experiments shall include the introduction of appropriate levels of background clutter and obstacles in order to develop an understanding of the biological system response to these factors.

- Translate the developed understanding of biological system tracking, guidance and control functions into useful algorithms for performing said functions on engineered flight vehicles, and/or use the developed understanding to devise improvements for existing algorithms.

- Continue the phase I effort to develop and demonstrate an architecture for direct adaptive guidance including formulations for preplanned path following, intercept, evasion (including obstacle avoidance), and rendezvous. These formulations shall include the ability to accept new guidance strategies that are biologically inspired (i.e. employ nominal guidance strategies other than proportional navigation and its variants).

- Continue the phase I effort to develop and apply practical image processing algorithms, (inspired in whole or in part by the companion study of biological systems performing similar tasks), for separating objects of interest from a cluttered background, tracking these objects (e.g. target(s) and obstacle(s)), and generating output signals representative of these objects that are appropriate to the task of adaptive guidance and controlled active vision.

- Exploit the developed target tracking algorithms (such as the computation of optical flow) in the analysis of collected insect trajectories, and to simulate and study features of vision (such as optical flow) experienced by the fly during pursuit.

- Explore the ability to improve system performance by employing the output (e.g. optical flow calculations) of developed image processing algorithms as a direct input to the neural network used for adaptive guidance (i.e. "controlled active vision").

- Evaluate and compare performance in simulation to traditionally engineered solutions, and iterate the design. Generate feedback into the experimental effort that will characterize the solutions employed by biological systems.

- Conduct a demonstration of the developed tracking, guidance and control technologies in

simulation of the selected target engagement scenarios.

- Flight demonstrate the most promising of the developed technologies on the fixed-wing UAV testbed(s) as time and resources permit.

A.4 ORGANIZATION OF THE REPORT

This report is organized into three main topics. In Section B we discuss the experimental work done at Cornell to characterize the tracking and guidance functions employed by the male flesh fly in pursuit of the female. In Section C we discuss work done in the area of image processing, but in anticipation of ultimately employing vision as a primary sensor modality in the practical implementation of a UAV guidance and control system, and as a tool for automated 3-D registration of recorded fly pursuits. In Section D we present our efforts to develop an adaptive intercept guidance law, and present numerical simulation results. Section E completes the report with conclusions.

B. CHARACTERIZATION OF MALE FLESH FLY PURSUITS

The overall goal in this area was to understand how male flies track their targets, especially in environments cluttered with obstacles and alternative targets (other flies), and subsequently to perhaps be able to use this understanding to emulate this capability in an engineered system.

B.1 CAPTURE OF 3-D PURSUIT TRAJECTORIES

Prof. Cole Gilbert and Dr. P.O. Zanen of Cornell University developed a reliable system for eliciting aerial pursuits of females by male flesh flies (see Figure B.1). The flight arena (see Figure B.2) has a calibrated active volume of 17x17x40cm, which is a subvolume of a larger cage (50x50x80cm) in which single males and 5-8 virgin females are released several days after eclosion as adults. Flies are 15mm long or shorter with a total mass of 70mg (head: 10.4mg). As the participants fly through the active volume, we capture pursuits with a digitizing precision of $1mm^3$. To increase accuracy in digitized orientation of the sensor (head) we attach a small fiducial stick of balsa wood (8mm, 1.6mg) that indicates the seeker's line-of-sight (LOS). Pursuit sequences are captured with a Kodak MotionCorder SR500 at 250fps and directly read into the hard drive buffer of the SR500. Good sequences are archived onto Hi-8 videotape for offline digitizing and subsequent analysis by custom software [Mantid 32, Synceros Inc.] developed by Dr. P.O. Zanen in the Cornell group. This software enables digitization of any objects in the fields of view.

FLESH FLY
Neobellieria bullata (Sarcophagidae)

body length: $\leq$ 15 mm
mass: 70 mg

linear velocity: $\leq$ 1 m s^{-1}

angular velocity: $\leq$ 1000°s^{-1}

Figure B.1. Male flesh fly under study.

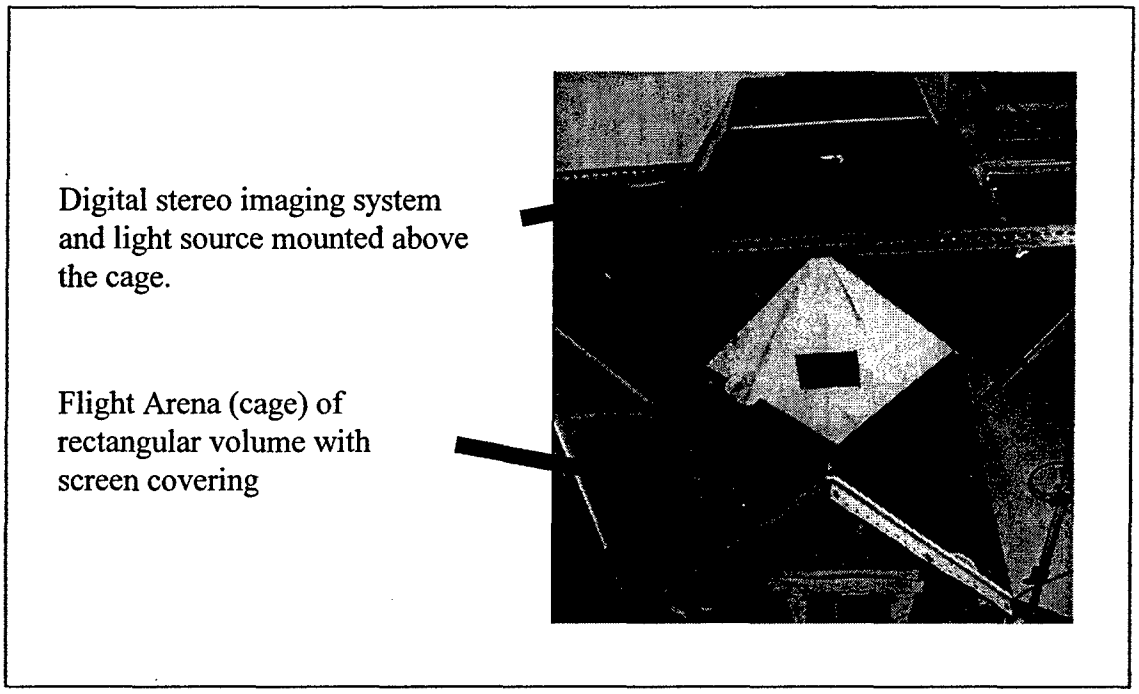

B.2 Photograph of the current flight cage with imaging equipment mounted above.

The test setup and associated procedures described above were gradually improved to ultimately obtain a number of very good pursuit/evasion trajectories that could be analyzed. Figure B.3 presents a three-dimensional plot of a "straight" trajectory and depicts the path of both the male pursuer (in black), and the female evader (in red). Points on the plot represent position samples at distinct points in time. The male has been fitted with a fiducial stick as described above and, though hard to see at this scale, the figure includes the line of sight angle by means of a vector ascribed to the male position at each sample time. Figure B.4 presents a second "curved trajectory. Figure B.5 presents a sequence in which the female successfully evades the male, but in which the male successfully reengages his target to complete the chase. Each of these figures is accompanied by a digital image sequence which is much more effective for initial description of the scenario, but which cannot be included in a printed document such as this. The video clips of these three trajectories were presented to the sponsor at the last program meeting, and delivered on a compact disk. Also included in Figure B.5 is a plot of the computed position of the female in the field of view of the male during the pursuit. Analysis of the optical flow of this type of data as it evolves during the pursuit reveals much about the sensory input to the male

target position in field of view of pursueing male on right.

B.2 View from the Cockpit of a Fly

Visual Guidance of Sexual Aerial Pursuit by Male Flesh Flies

Male flies of many species pursue females aerially at high angular and linear velocities with successful pursuit ending in physical contact with the female, yet the male fly has very limited visual information about the target and slow computational processing in a brain comprised of only about 300,000 neurons. Nevertheless, during their brief adult life many males successfully accomplish this task multiple times in cluttered environments. We have studied this 3D visually-guided pursuit task in the flesh fly, *Neobellieria bullata* [family Sarcophagidae, mass: 70mg, length: 15mm] in two project areas. 1) We have made optical measurements of the compound eye to determine the distribution of visual acuity across the sensor array. 2) In the laboratory we have made 3D high-speed digital video of females being chased by a male fitted with small balsa sticks on the head and thorax to enable more precise angular measurements of the position of the head, which bears the immobile eyes, relative to the thorax (Fig. B.1).

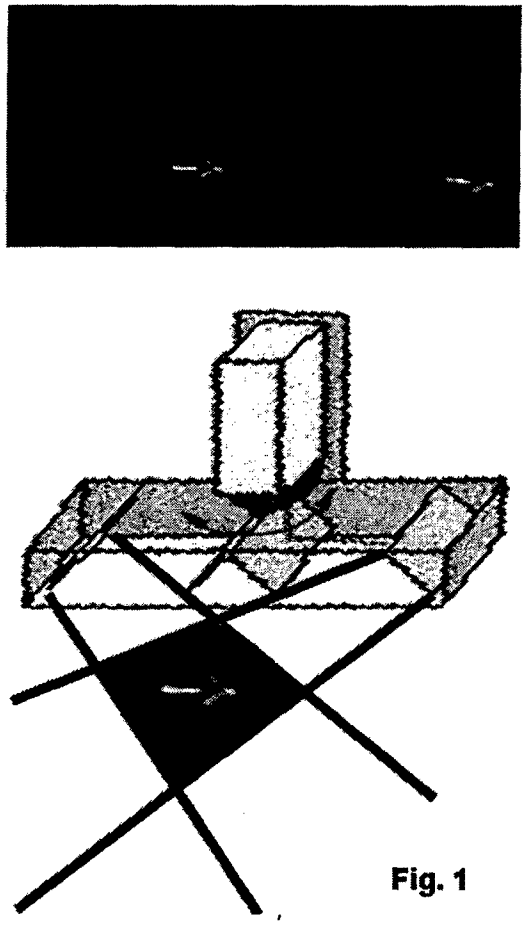

Fig. 1

We counted lens facets directly and determined the distribution of optical acuity across the sensor array by mapping regional differences across the compound eyes of males and females using an ophthalmoscopic method to quantify angular divergence of optical axes of neighboring lens facets ($\Delta\Phi$ in degrees). Male *N. bullata* have about 400 more facets than do females (Fig. 2, upper panels). Males also have a larger region of binocular overlap, as seen in the lower portion of Fig. 1, which depicts the spherical extent of the field of view of the right eye from a frontal oblique perspective. Visual acuity is not uniform across the sensor array. Both male and female flies have higher acuity in the forward viewing retina, but males have a higher maximum acuity ($\leq 1^{\circ}$) and a larger dorso-frontal zone of high acuity facets than do females (Fig. B.2, lower panels).

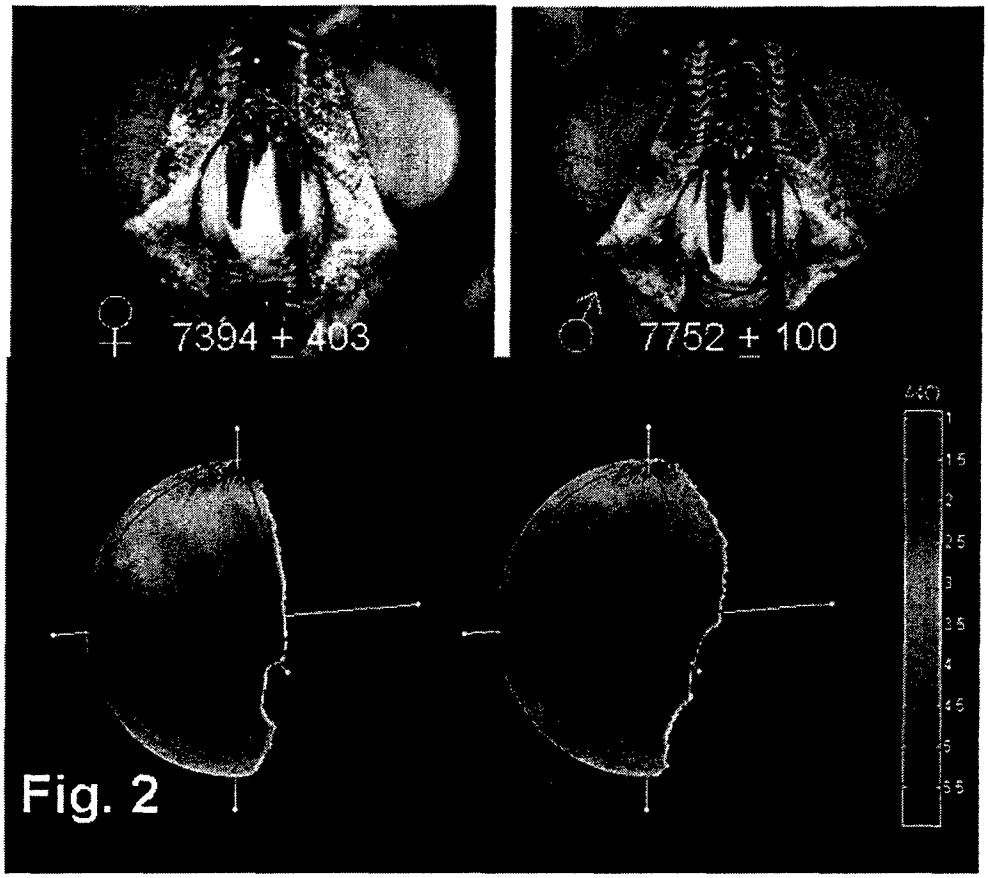

Fig. 2

To determine how regional differences in sensor acuity effect pursuit and to begin to develop simple guidance laws, we filmed pursuits in 3D at 250 frames s^{-1} with a single camera viewing

an active space of 20x20x30 cm in a larger cage through a single camera [Kodak MotionCorder] and a system of front surface mirrors (Fig. 1). The position of the female and the tips of the balsa sticks on the head and thorax of the male were digitized with $1mm^3$ precision using software by Synceros [www.synceros.com]. All data were further analysed in Matlab. Males achieve linear velocities $\geq 2\ ms^{-1}$ and angular velocities $\geq 3000\ ^{o}s^{-1}$ during pursuit and often capture females. Moreover, males often change the direction of gaze by rotating their head during pursuit as evidenced by the lack of congruence between the 3D line-of-sight (LOS) position of the target with respect to the direction of gaze (Fig. 3, left spherical projection) and with respect to the flight direction (Fig. 3, right spherical projection). Furthermore, such changes in the direction of gaze do not necessarily bring the target LOS into the high acuity region of the sensor array. Nevertheless the pursuit represented in Fig. 3 resulted in a capture beginning with the pale dot marked "start" and continuing through more saturated dots to the one marked "stop".

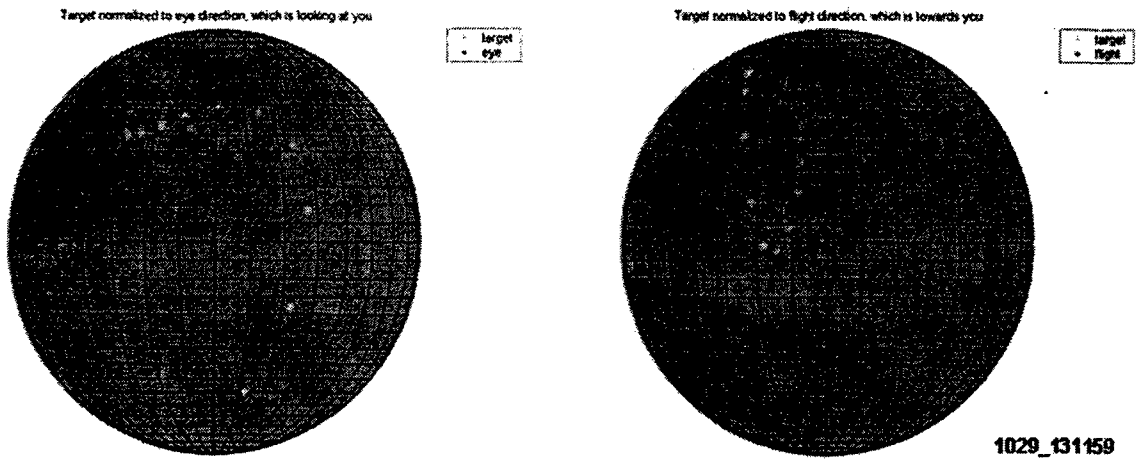

Fig. 3

In other pursuits, such as that represented in Fig. B.4, the male changed his direction of gaze and controlled the target LOS to maintain it in the region of high acuity of the sensor array for a large portion of the pursuit, in spite of extreme maneuvers by the target. Nevertheless, this pursuit was unsuccessful.

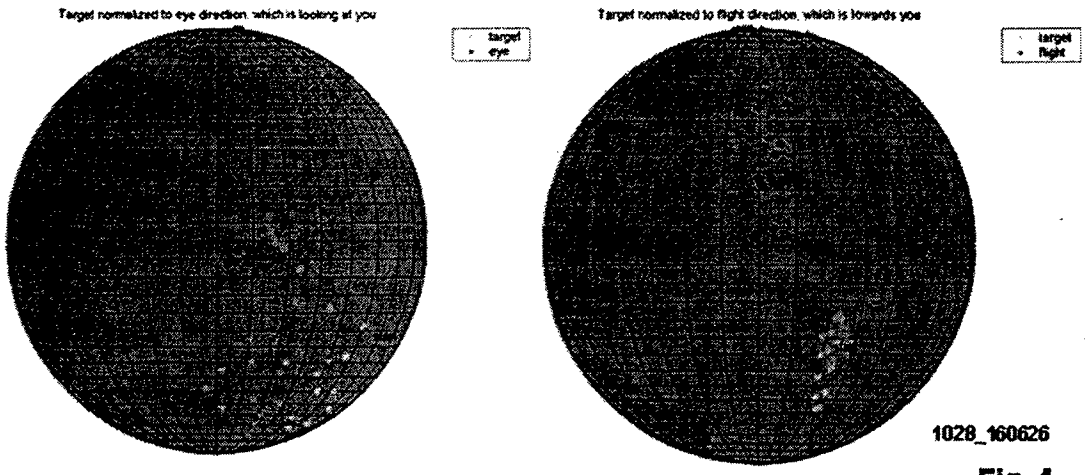

Fig. 4

Thus, the visual guidance of pursuit by the fly is more complex than a system in which a mobile seeker controls target LOS and the seeker position drives the fly's trajectory. The fly does have sensors that monitor its head position, so the signal given to the motor system may be expressed as the angular sum of LOS and head position. Our attempts at simulating a guidance algorithm using this sum are still rather naïve and capture some aspects of the pursuit trajectories, but are currently not very satisfactory. Future work will focus on an experimental paradigm that allows us more control over the stimulus, so that target trajectory and clutter can be modified under experimenter control.

This work is documented extensively in the form of 3-D visualization sequences and other multimedia files presented at the program meeting held in December of 2002 and distributed on CD.

It was confirmed that the males move their heads while chasing, resulting in the target being imaged near the forward directed retinal zone of higher acuity. The flight path, however, is not always colinear with the line of sight to the target, but in some pursuits may deviate by 30 degrees or more. Other interesting findings and new research directions are discussed in the following.

Automatic image segmentation.

Olivier Zanen worked closely with Allen Tannenbaum and one of his students, Suo Jin, to use the image segmentation software to automate the processing of captured pursuit sequences. The software will recognize the fiducial sticks on the head and thorax of the fly by their contrast. It works with both left and right 3D views to return coordinates of the ends of the sticks. The contour extraction was made to work well and a second module created that calculates the [x,y,z,] coordinates of the tips of the contours.

Chases of an experimenter controlled target by small flies.

We have gotten a species of small flies in culture and attempted to get them to chase targets that move under our control. We have had some success and a few males have chased a dummy moving in a circle. We are still adjusting the parameters, e.g., age of the male, lighting, etc. to make pursuit in this artificial situation more reliable. We will continue with this approach as it has more potential for examining the effects of visual clutter and obstacles on the pursuit by the male.

Response to Clutter

We also explored how the flies maneuver in a cluttered background. The filming/flight cage was modified to improve the fly's willingness to fly and chase. We now have them in a moving cylindrical cage with no pattern on the walls. They seem somewhat reluctant to land on the walls and are more inclined to stay active in flight, which increases the probability of chasing. We subsequently included patterns on the walls to create visual Aclutter. We have filmed, a case of distraction by multiple targets. A male chasing a female made contact, but was not able to hold on to her. After righting himself, he immediately renewed the pursuit, as is typical in such cases, but by that time another female had crossed his LOS and he pursued her instead of the original female.

Research Dissemination

Prof. Gilbert gave a seminar presentation on this project to the Department of Entomology at the University of Arizona and to a vision group in the Department of Biological Sciences at Arizona State University. He also participated in a DARPA workshop on ACognitive Arthropods@ held in Arlington, VA and hosted by Draper labs 23-24 April. Half a dozen or so project managers from DARPA and other agencies, (including Ric Wehling from AFOSR) and about 60 scientists, including five biologists, met to discuss small scale (e.g., 1 x 1 x 2 cm) autonomous, cooperative agents from the point of view of power (e.g., 1 mW for a year), bionics, and navigation/guidance. Although the task is challenging, the discussion was fruitful. DARPA is looking at a 2010-2020 time frame for funding.

C. Visual Tracking Techniques

C.1 INTRODUCTION

This section summarizes visual tracking techniques that have been improved and applied. This work draws upon team member Prof. Allen Tannenbaum's experience in image processing and computer vision employing certain geometric variational evolution equations.

Vision is a key sensor modality in both the natural and man-made domains. The prevalence of biological vision in even very simple organisms, indicates its utility in man-made machines. More practically, cameras are in general rather simple, reliable passive sensing devices which are quite inexpensive per bit of data. Furthermore, vision can capture multispectural information at a high rate at high resolution, and with a wide field of view. Finally, cameras can be used in an active manner. Namely, one can include motorized lenses mounted on mobile platforms which can actively explore the surroundings and suitably adapt their sensing capabilities. For some time now, the role of control theory in vision has been recognized. In particular, the branches of control that deal with system uncertainty, namely adaptive and robust, have been proposed as essential tools in coming to grips with the problems of both biological and machine vision. These problems all become manifest when one attempts to use a visual sensor in an uncertain environment, and to feed back in some manner the information gathered.

Specifically, Dr. Tannenbaum's research has impacted the following areas:

(a) Geometric and Statistical Methods for Deformable Contours:

One of the key techniques in active vision and tracking is that of deformable contours or snakes. These are autonomous processes which employ image coherence in order to track features of interest over time. Our general approach is based on the theory of geodesics and minimal surfaces in a conformal geometry. We have introduced some more sophisticated features into our geometric-based cost function. In particular, we incorporate non-local information into our models based on the use of adaptive filtering schemes and Bayesian statistics. This leads to a natural synthesis of our work in knowledge-based segmentation and active contours.

(b) Area-Preserving Mappings and Optical Flow:

The computation of optical flow has proven to be an important tool for problems arising in active vision, including visual tracking. The optical flow field is defined as the velocity vector field of apparent motion of brightness patterns in a sequence of images. It is assumed that the motion of the brightness patterns is the result of relative motion, large enough to register a change in the spatial distribution of intensities on the images. We have considered various constrained optimization approaches for the purpose of accurately computing optical flow. In phase I, we have weakened the optical flow constraint with one based on the theory of area-

preserving mappings. We want to also include an optical flow term in our active contour functional to better track moving images.

(c) Optimal Transport and Image Registration:

Image registration is the process of establishing a common geometric reference frame between two or more data sets from the same or different imaging modalities taken at different times. This is essential for data fusion, since registration can provide automated methods that align multiple data sets with each other and with a given object of interest. We have been using these methods for data fusion as well as a new approach to optical flow for tracking. Our measure of similarity can be based on comparing the mass densities of the images.

We now summarize some of our findings.

C.2 SEGMENTATION

One of the key problems in vision is segmentation, i.e., the partitioning of an image into homogeneous regions with ``homogeneity" defined in various contexts, e.g., texture, intensity, color, etc. Numerous approaches have been proposed for this. We have considered methods using a synthesis of partial differential equation methods based on the calculus of variations and statistical based ideas.

We have used a combination of anisotropic diffusion and the Bayesian methods for the segmentation of such difficult noisy imagery such as SAR (synthetic aperture radar) and ultrasound. In our approach, we introduce a priori knowledge about the number of objects present in the image via Bayes' rule. Posterior probabilities obtained in this way are then anisotropically smoothed, and the image segmentation is obtained via MAP classifications of the smoothed data.

The model we employ begins with the assumption that the image is composed of n classes of objects. The goal of our segmentation is to determine to which class each pixel in the image belongs. We assume that the value of each pixel in a given class can be thought of as a random variable with a known normal distribution, and that these variables are independent across pixels. Given a set of intensity distributions and priors, we can apply Bayes' Rule from elementary probability theory to calculate the posterior probability that a given pixel belongs to a particular class, given its intensity. We can then calculate the posteriors P, and then to apply anisotropic smoothing to each P. Specifically, we have chosen to smooth by evolving P according to a discretized version of the affine geometric heat equation. This particular diffusion equation was chosen because of its affine invariance, because it preserves edges well, and because of its numerical stability. The final segmentation is obtained using the maximum *a posteriori* (MAP) probability estimate after anisotropic smoothing.

These methods fit very naturally into our overall research paradigm of combining curve evolution with statistical techniques. They will be carefully implemented and tested in order to

develop the optimal combination of the statistical and variational-based segmentation approaches. Furthermore, temporal information may be easily included in this scheme as is shown by the images below. These images were generated by processing computer generated imagery from a real-time simulation of pursuit/evasion. Clearly identified and tracked in the image are the horizon and the target (both indicated in red).

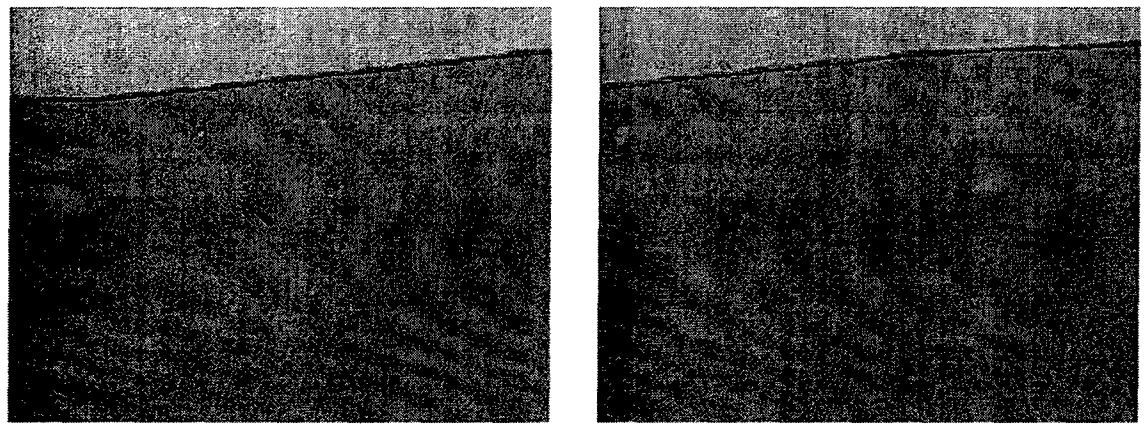

Figure C.1 Two frames of automatic tracking of a fixed-wing UAV target in flight using active contours with statistical methods.

C.3 Optical Flow

The computation of optical flow is an important tool for problems arising in active vision. The optical flow field is the velocity vector field of apparent motion of brightness patterns in a sequence of images. One assumes that the motion of the brightness patterns is the result of relative motion, large enough to register a change in the spatial distribution of intensities on the images. Thus, relative motion between an object and a camera can give rise to optical flow. Similarly, relative motion among objects in a scene being imaged by a static camera can give rise to optical flow.

A typical method is to consider a spatiotemporal differentiation method for optical flow. Even though in such an approach, the optical flow typically estimates only the isobrightness contours, it has been observed that if the motion gives rise to sufficiently large intensity gradients in the images, then the optical flow field can be used as an approximation to the real velocity field and the computed optical flow can be used reliably in the solutions of a large number of problems. Thus, optical flow computations have been used quite successfully in problems of three-dimensional object reconstruction, and in three-dimensional scene analysis for computing information such as depth and surface orientation. In object tracking and robot navigation,

optical flow has been used to track targets of interest. Discontinuities in optical flow have proved an important tool in approaching the problem of image segmentation.

One constraint which has often been used in the literature is the "optical flow constraint" (OFC). The OFC is a result of the simplifying assumption of constancy of the intensity, E = E(x,y,t), at any point in the image. It can be expressed as the following linear equation in the unknown variables u and v:

$$E_t + E_x u + E_y v = 0 .$$

Here u and v are the x and y velocity components of the apparent motion of brightness patterns in the images, respectively. It has been shown that the OFC holds provided the scene has Lambertian surfaces and is illuminated by either a uniform or an isotropic light source, the 3-D motion is translational, the optical system is calibrated and the patterns in the scene are locally rigid.

It is not difficult to see that computation of optical flow is unique only up to computation of the flow along the intensity gradient of E. This is the celebrated *aperture problem*. One way of treating the aperture problem is through the use of regularization in computation of optical flow, and consequently the choice of an appropriate constraint. A natural choice for such a constraint is the imposition of some measure of consistency on the flow vectors situated close to one another on the image.

The optical flow constraint above is of course very strong. Motivated by work on area-preserving maps, we have considered in phase I a modification that could be placed in a variational setting. Indeed, given a family of nowhere-zero 2-forms τ^t we have an explicit method to determine a family of diffeomorphisms ϕ^t such that

$$(\varphi^t)^* \tau^t = \tau^0 .$$

Differentiating this expression yields an expression very similar in form to the standard optical flow constraint with a divergence term added. Optical flow computed with respect to this new expression seems to give rather encouraging results as is illustrated by the following figures. These are data frames generated in Prof. Gilbert's lab and are images of a male flesh fly in pursuit of a female within a caged area and including another female. The red vectors represent the optical flow field associated with the motion of the three flies.

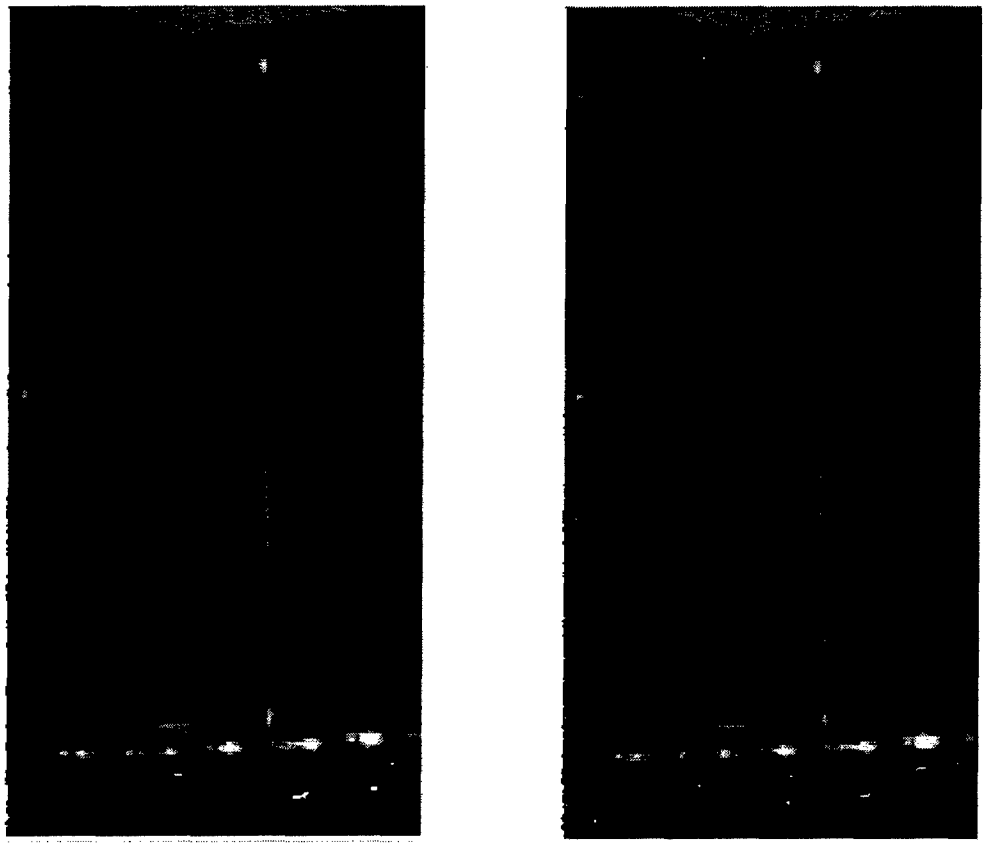

Figure C.2 Two frames of automatic tracking of flies using optical flow.

C.4 Optimal Transport, Image Registration, and Optical Flow

In this section, we discuss a method for image registration and warping based on the classical problem of optimal mass transport. The mass transport problem was first formulated by Gaspar Monge in 1781, and concerns the optimal way (in the sense of minimal transportation cost) of moving a pile of soil from one site to another. This problem was given a modern formulation in the work of Kantorovich, and so is now known as the *Monge-Kantorovich problem.*

This type of problem has appeared in econometrics, fluid dynamics, automatic control, transportation, statistical physics, expert systems, and meteorology. It also naturally fits into certain problems in computer vision. In particular, for the general tracking problem, a robust and reliable object and shape recognition system is of major importance. A key method to carry this out is via *template matching*, that is one wants to match some object to a given catalogue of objects. Typically, the match will not be exact and hence some criterion is necessary to measure the ``goodness of fit." The matching criterion can also be considered a *shape metric* for deciding the similarity of one shape to another.

We are interested in applying these techniques to the problem of registration. Indeed, many registration problems can be formulated in terms of mapping densities. The key idea is to find the optimal mapping via an equivalent problem involving certain factorizations (called ``polar") of mass-preserving mappings. We have done this via a natural gradient descent technique similar to the methods that we considered for area-preserving maps of minimal distortion.

It is important to note that there are a number of ways in which multimodal registration impacts image-guided tracking and control systems. First, it allows for fusing the different types of information of each imaging modality, providing better and more accurate information than each image viewed separately. Further, it allows quantitative comparison of images taken at different times, from which precise information about evolution over time can be inferred. Third, it allows for updating in real time a pre-existing image or model with dynamic data.

Multimodal registration is in fact not one but many problems. It greatly depends on the imaging modalities to be matched, on the objects being imaged, and the requirements of the given task. These parameters determine the assumptions and technical requirements, which in turn greatly influence the choice of models and solution methods.

Multimodal registration proceeds in several steps. First, each image or data set to be matched should be individually calibrated, corrected for imaging distortions and artifacts, and cleared of noise. Next, a measure of similarity between the data sets must be established, so that one can quantify how close one image is from another after transformations are applied. Such a measure may include the similarity between pixel intensity values, as well as the proximity of predefined image features such as fiducials, landmarks, surface contours, and ridge lines. Next, the transformation that maximizes the similarity between the transformed images is found. Often this transformation is given as the solution of an optimization problem where the transformations to be considered are constrained to be of a predetermined class. Once the transformation is obtained, it can be used to fuse the image data sets in order to plan a task, to navigate, or to track.

In our case, we are comparing densities on images via the Kantorovich-Wasserstein distance. This occurs, e.g., in functional imaging where one wants to compare mass densities of various features deforming over time, and obtain the corresponding elastic warp map. It also allows us to compare given scalar and vector fields on images.

There have been a number of algorithms considered for computing an optimal transport map. For example, methods have been proposed based on linear programming (for measures which are weighted sums of delta functions) and on Lagrangian mechanics closely related to ideas from fluid flows (Euler equation, etc.) In our work, we have used a gradient descent approach based on the notion of *polar factorization* which allows one to remove the curl from a vector field in computing the optimal transport map. These ideas can naturally be used in studying optical flow as is illustrated in the deforming heart image given below.

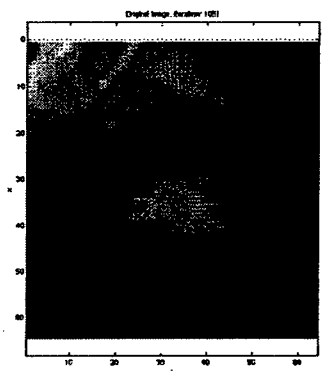

(a)

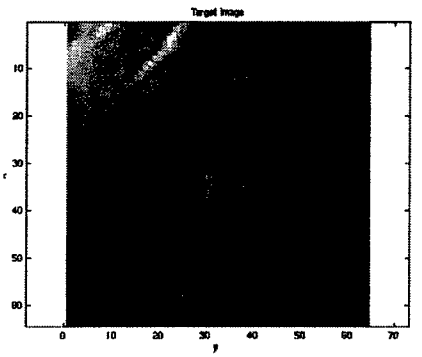

(b)

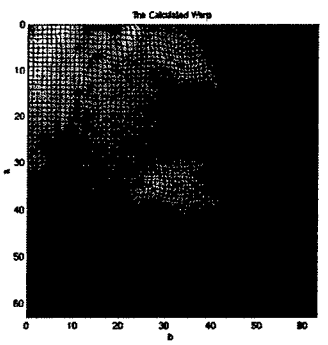

(c)

Figure C.3 Optimal transport based optical flow. (a) is original image with flow lines indicate. (b) is the target image. (c) shows the warping map.

C.5 DYNAMIC SNAKES FOR TRACKING

In this final report, we summarize our finding on the use of active contours for tracking. We should note that the active contour tracking software developed as part of this STTR project has been incorporated in the Cornell University Laboratory headed by Professor Cole Gilbert for fly tracking and has been licensed to Dr. Olivier Zanen. The software is written in C++ and runs on LINUX. A key new ingredient which we added in the present project and is included in the software is the ability to propagate velocity information from frame to frame which is essential in all dynamic tracking applications. This will be explained below in some detail.

Contour/surface trackers

We distinguish two different methods for active contour based visual tracking: the static and the dynamic approaches. In the static approach the contour is oblivious of its own state; no velocity information is propagated. Tracking can then be achieved by solving subsequent static problems (assuming that the movements of the object to be tracked are relatively slow and the contour does not leave the object's capture range from one frame to the next), or by incorporating temporal information (e.g. optical flow) into an energy functional to be minimized. The dynamic approach is based on a dynamical systems perspective, where points on the contour possess an inherent kinetic energy. They are typically associated with a mass, a velocity vector, are related to their neighboring points through elasticity and rigidity constraints, and move based on an underlying potential field.

Most active contours are static. However, variational formulations many times appear to be dynamic, because the Euler-Lagrange equations are solved by gradient descent, introducing an *artificial* time parameter. To use static active contours for visual tracking one usually uses a two step approach. First, the curve evolves on a static frame until convergence (or for a fixed number of evolution steps). Second, the location of the curve in the next frame is predicted. In the simplest case this prediction is the current location. Better prediction results can be achieved by using optical flow information for example. Here, the curve is not moving intrinsically, but instead is placed in the solution's vicinity by an external observer (the prediction algorithm). The curve is completely unaware of its state. In the static case, level set methods are known to handle sharp corners, topological changes (for tracking multiple objects), and to be numerically robust. In this program, we now have a fully automatic dynamic active contour algorithm ideal for tracking multiple objects (such as several flying insects in the presence of clutter and occlusions).

Dynamic Active Contours

The dynamic active contour approach for the first time allows one to track objects without the need for an external prediction step. Moreover, since it is based on curvature driven flows, it is naturally multi-scale, and so much less sensitive to noise as conventional active contours. It is the natural dynamic generalization of geodesic active contours. We consider the evolution of closed curves of the form $C(p,t): S^1 \times [0,\tau] \to R^2$ in the plane. Here $p \in [0,1]$ parameterizes the curve and t parameterizes the family. Since the curves are closed, we have $C(0,t) = C(1,t)$. The

The classical formulation for dynamic curve evolution is derived by means of minimization of the action integral

$$\mathbf{L} = \int_{t=t_0}^{t=t_1} L(t, C, C_t)dt \quad (1)$$

where the subscripts denote partial derivatives. The Lagrangian $L = T - U$ is the difference between the kinetic and the potential energy. In previous works, the choice of the Lagrangian was not intrinsic to the geometry of the evolving contour since it was dependent upon its parametrization p, and was sensitive to noise and local minima.

Accordingly, we have considered the following Lagarangian:

$$L := \int_0^1 (\frac{1}{2}\mu \| C_t \|^2 - g) \| C_p \| dp.$$

Taking the first variation of the action integral (1) results in the Euler-Lagrange equation

$$\mu C_{tt} = -\mu T_{ts} C_t - \mu C_t \cdot C_{ts} T - (\frac{1}{2}\mu \| C_t \|^2 - g)\kappa N - (\nabla g \cdot N)N \qquad (2)$$

which is intrinsic and a natural extension of the geodesic active contour approach. Here N is the unit inward normal, and $T = C_s$, the unit tangent vector to the curve. $\kappa = C_{ss} \cdot N$ and $ds = \| C_p \| dp$ is arc-length.

A common choice for the potential function is

$$g(x,y) = \frac{1}{1 + \| G_\sigma * \nabla I(x,y) \|^r}$$

where I is the image, r is a positive integer, and G_σ is a Gaussian of variance σ^2.

Equation (2) describes a curve evolution that is only influenced by inertia terms and information on the curve itself. To increase robustness, we added region-based terms in the formulation. Automatic generation of the fly's coordinates was also generated as part of our package provided to the Cornell team.

D. ADAPTIVE INTERCEPT GUIDANCE

Summary

This report presents summaries the research performed in the area of biologically inspired adaptive intercept guidance. The objective of this task is to develop a Neural Network (NN) based adaptive control approach to the intercept guidance problem in which vision is the only sensing modality. The primary role of adaptation is to provide or improve the means by which target acceleration is estimated. In one approach we examine the possibility of employing a direct adaptive control approach that compensates the guidance law for target acceleration, without the traditional use of a Kalman filter to estimate target acceleration. In a second approach, we use adaptation to compensate an extended Kalman filter rather than the guidance law. The point of both of these approaches is that estimator based methods rely on a model for target behavior, whereas a direct adaptive approaches that are NN based might be free of target modeling assumptions, relying instead on the universal approximation properties of NNs to cancel the effect of the target evasive strategy. Such an approach, if successful, should be less vulnerable to intelligent targets that can take advantage of the modeling assumptions imbedded in estimator based design approaches, or in approaches that do not attempt to estimate target acceleration. For the later case, we consider augmenting a simple proportional navigation (PN) guidance law. For the former case we adaptively augment an extended Kalman filter that estimates target acceleration, and apply this estimate to an augmented PN (APN) guidance law. In this approach the premise that the nominal estimator is model based, and the target may perform an evasive maneuver that is not representative of the modeled behavior. Therefore, the role of the NN in this setting is to correct for modeling error in the target state estimator.

The report first presents a review of some of the basic concepts of PN guidance. This is followed by a brief description of the theoretical issues involved in attempting to develop a NN based adaptive guidance law as an extension of our earlier work in flight control. This approach does not employ a Kalman filter in the design. Then we discuss the approach for augmenting an extended Kalman filter instead of the APN guidance law. Finally, the two approaches to adaptive guidance (adaptation in the guidance law versus adaptation in the target acceleration estimator) are compared in a nonlinear simulation representative of pursuit of a maneuvering low speed target. Conclusions and recommendations for future research are given at the end.

D.1 PROPORTIONAL NAVIGATION

Equations of Motion

The engagement scenario is depicted in Figure 1.

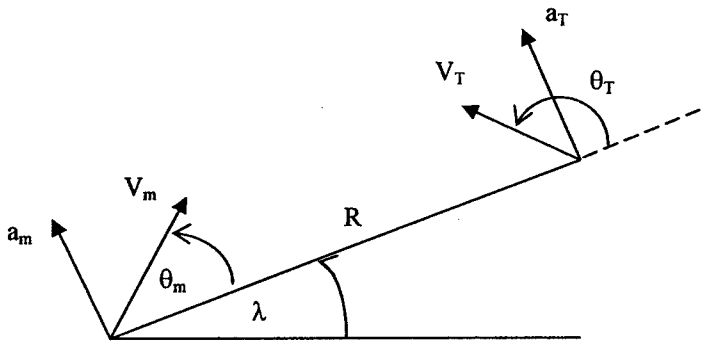

Figure D.1. Planar engagement model.

The line-of-sight (LOS) rate and range rate are given by:

$$\omega = \dot{\lambda} = \frac{V_T \sin\theta_T - V_M \sin\theta_m}{R}$$
$$\dot{R} = V_T \cos(\theta_T) - V_m \cos(\theta_m) \qquad \text{(D.1)}$$

Defining a_m as the purser acceleration normal to the LOS, and assuming that it is the result of a net acceleration normal to the velocity, and similarly defining a_T, then

$$V_m(\dot{\theta}_m + \dot{\lambda}) = a_m / \cos\theta_m$$
$$V_T(\dot{\theta}_T + \dot{\lambda}) = a_T / \cos\theta_T \qquad \text{(D.2)}$$

To simplify the analysis we introduce the following two assumptions:

Assumption 1: The velocities are constant
Assumption 2: The range rate is constant

Differentiating the first of (D.1), and using $R=-\tau\dot{R}$, where τ denotes time-to-go, and invoking the above assumptions we obtain

$$\dot{\omega} = 2\omega/\tau + (a_T - a_m)/R \qquad \text{(D.3)}$$

Required Acceleration

Substituting the PN guidance law

$$a_m = -N\dot{R}\omega = NV_c\omega \tag{D.4}$$

into (Δ.3), where V_c is the closing velocity, gives

$$\dot{\omega} = (2-N)\omega/\tau + a_T/R \tag{D.5}$$

It is well known that PN guidance with navigation gain N=3 is optimal in the sense of minimizing the integral square of the control effort against non-maneuvering targets ($a_T = 0$).[1] With Assumptions 1 and 2, and with an initial heading error of zero, the pursuer's acceleration against a constant accelerating target is given by

$$a_m/a_T = \frac{N}{N-2}\left[1-\left(1-t/t_f\right)^{N-2}\right] \tag{D.6}$$

Figure D.2 illustrates the normalized acceleration resulting from (6) as it depends on N and τ.

APN guidance is optimal in the same sense as PN guidance when flown against a constant maneuvering target. This form of guidance is given by

$$a_m = -N\dot{R}\omega + a_T/2 \tag{D.7}$$

which assumes that an estimate of a_T is available. The corresponding solution for the pursuer's acceleration with zero initial heading error, against a constant accelerating target is

$$a_m/a_T = 0.5N(1-t/t_f)^{N-2} \tag{D.8}$$

Figure D.3 illustrates the normalized acceleration for this guidance law. Note that the acceleration profiles are monotonically increasing functions of t/t_f for PN, and the reverse is true for APN. Bear in mind that for N=3 and a_T = constant, the APN law is optimal, whereas the PN law is not. The obvious advantage of APN is that the acceleration requirement near the final time is minimal, whereas for PN it reaches a maximum. This is crucial in situations where the acceleration limit is encountered during an engagement. It is much better to encounter a limit earlier in the engagement rather than later, since encountering a limit near the point of closest approach will result in a large miss distance.

There are several important factors to bear in mind when implementing a PN or APN law. Both laws require an estimate of closing velocity. The manner in which this estimate is obtained

may vary significantly depending on the method used for sensing. When employing passive sensing, such as might the case in an optical or visual system, only angle and possibly subtended angle may be measured. If the target size is known, then subtended angle may be viewed as equivalent to a range measurement. In either case, both LOS rate and closing velocity have to be inferred somehow from these measurements. Therefore some form of estimation is required even if no attempt is made to compensate for target acceleration. If an active seeker is employed, then the closing velocity is directly measured from the Doppler effect, but the other basic measurement is still LOS and not LOS rate. Therefore even in the case of an active seeker, it is still necessary to provide some means of estimating the LOS rate. In the simplest mechanization of PN, LOS rate is commonly estimated by employing a signal in the high gain seeker tracking loop that is nearly proportional to the LOS rate. A simple analog representation of a seeker tracking loop showing the point where the LOS rate signal is taken is shown in Figure D.4[3], where ε represents the bore sight error. It should be noted that for passive systems, where range is inferred from subtended angle (image size), estimation of closing velocity is difficult, even when employing a nonlinear estimator. Since the viewpoint taken in this study is that a vision based method of guidance is employed, when an extended Kalman filter (EKF) is employed, the estimates of $\dot{\lambda}$, a_T and V_c are provided as a byproduct of the state estimates. When PN guidance is implemented, the LOS rate is obtained using the approach in Figure 4, and a constant value is used for the estimate of V_c.

In a modern guidance law, the dynamic model used to design the estimator is linear, and time-invariant, and is based on Assumptions similar to Assumptions 1 and 2 presented earlier. To obtain a linear time-invariant form R(t) is assumed known, or directly measured. Only the steady state Kalman filter solution is normally implemented. As noted previously, the basis for this effort is that a vision based sensor is employed, in which the basic measurements are LOS angle and R. However, we employ a nonlinear dynamic model for the intercept dynamics to design the estimator, and the resulting estimator is an EKF. The reason for this is that we are interested in examining engagements having highly nonlinear dynamics (large rotations in the LOS), and that may involve multiple attempts to intercept the target. We then augment the EKF with a NN. This required the development of a new approach for adaptively augmenting an EKF[4], whereas previous results were limited to augmenting a linear, time invariant estimator[5].

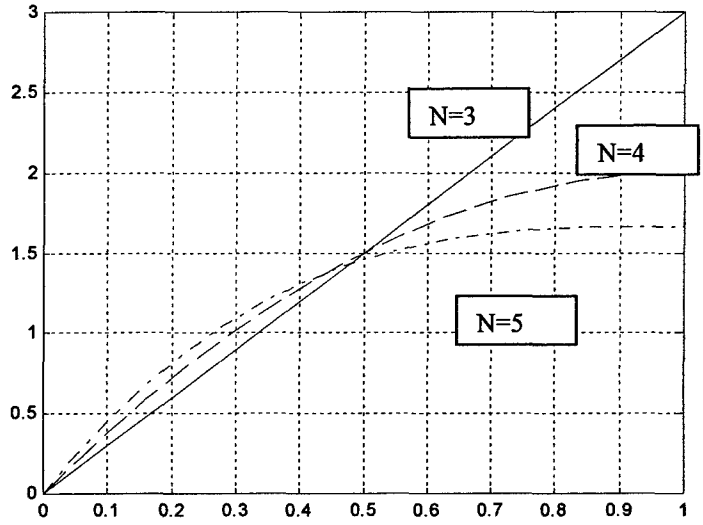

Figure D.2. Normalized acceleration due to target maneuver for PN guidance.

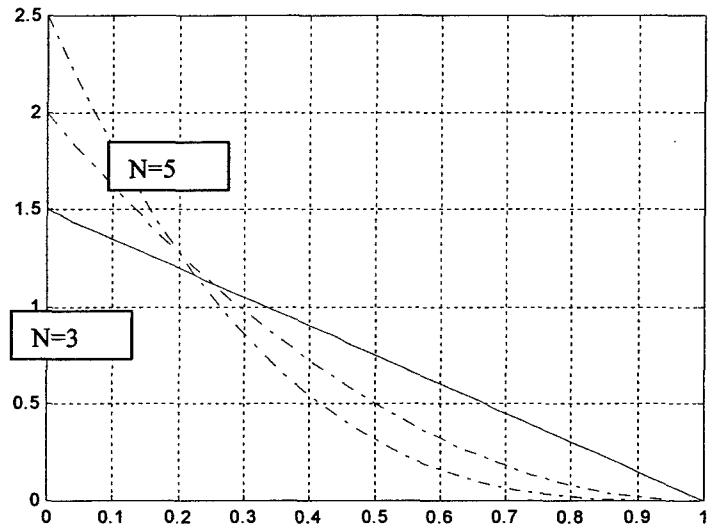

Figure D.3. Normalized acceleration due target maneuver for APN guidance.

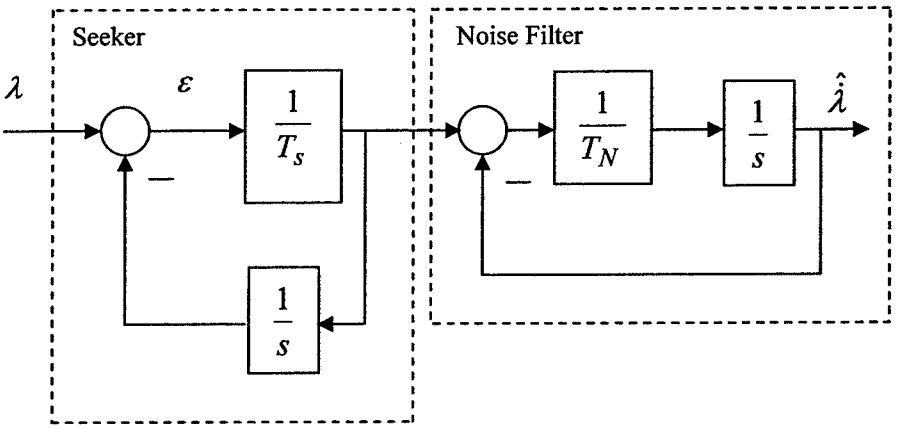

Figure D.4. Classical estimation of LOS rate.

The Effect of Target Maneuvers

The PN and APN laws when implemented in their ideal forms result in zero miss against an evading target. Real implementations result in non-zero miss distances due to a combination of factors. The primary sources of miss are: saturation, nonlinearities in the engagement dynamics that represent a significant departure from Assumptions 1 and 2, lags within the guidance loop (mainly due to autopilot lags), sensor noise, and other sources of tracking error such as radome slope error. Sensor errors also place a lower limit on the effective time constant of the guidance loop.[2] An illustration of the miss that results when a first order lag with a time constant of T = 1.0 sec is modeled in the guidance loop, is given in Figure D.5. This shows the miss due to a constant 3G target acceleration as a function of time-to-go at the start of the maneuver.

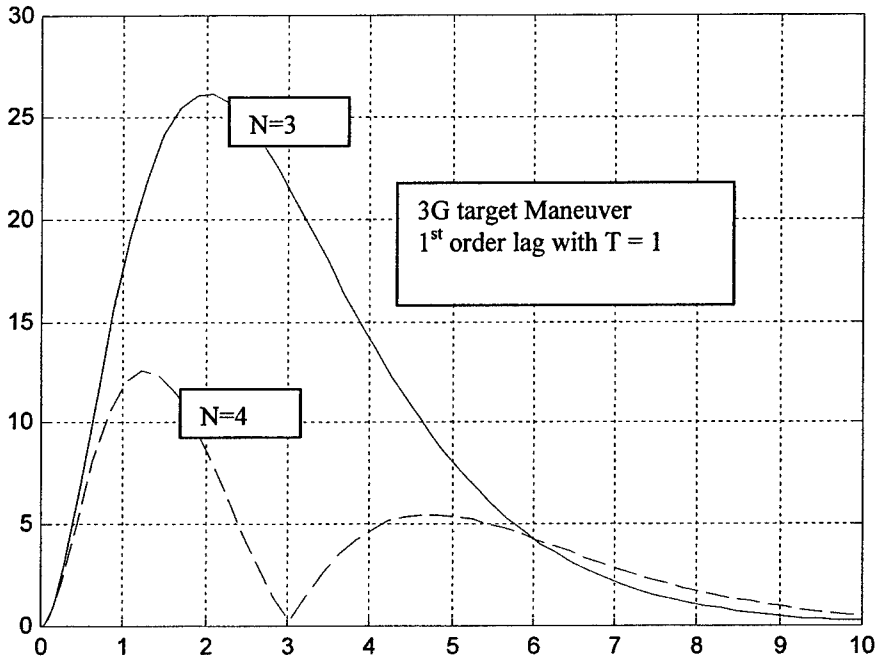

Figure D.5. Miss distance (feet) versus time-to-go at launch due to target maneuver.

While Figure D.5 illustrates the effect of a constant target maneuver, it can be shown that the optimal evasive maneuver is bang-bang, with reversals occurring at times-to-go corresponding to the peaks in the above curves.[2] So for N = 3, the optimal evasive strategy is to reverse the maneuver at approximately $\tau = 2.5$ seconds. This would have the net effect of doubling the peak miss distance, or achieving a total miss of approximately 54 feet. Kalman filters have traditionally been used in the guidance loop to estimate target acceleration in order to implement APN. While this is very effective against target maneuvers that belong to the class for which the filter is designed, it also makes the overall design more vulnerable to evasive strategies that take advantage of these modeling assumptions.

D.2 NN Based Adaptive Guidance

There are 3 major challenges in considering the application of adaptive control to guidance problems. The first concerns the fact that guidance problems are fundamentally finite time problems, whereas adaptive control is formulated as a problem in asymptotic stability. That is, Lyapunov stability analysis methods are employed to derive or justify the adaptive laws, and therefore the focus is on the behavior about an equilibrium point with respect to the error dynamics. One approach to overcome this problem is to consider the following time transformation

$$t_1 = \ln(1/\tau) \tag{D.9}$$

where $\tau = t_f - t$. This transforms the LOS rate dynamics in (D.3) to the following linear, time invariant form

$$d\omega / dt_1 = 2\omega + (a_m - a_T)/\dot{R} \tag{D.10}$$

This suggests that the adaptive control problem should be formulated in the t_1 time scale. The adaptive laws would be integrated in the real time scale, but with an integration step of $dt_1 = dt/\tau$, where dt is the increment in real time. *Thus an estimate of time-to-go is required to implement the adaptive law when employing the transformation in (D.9).*

The second difficulty surrounds the fact that the guidance problem is fundamentally an output feedback problem. Fortunately, under AFOSR sponsorship, have recently developed a direct approach to adaptive output feedback design that avoids the introduction of a model-based observer.[6] This formulation is also applicable to problems with unmodeled dynamics.

A drawback to the method in Ref. 6 is that it is based on feedback inversion of the plant dynamics, which constitutes the third major barrier to be surmounted. Feedback inversion is not appropriate for intercept guidance applications for numerous reasons. Ideally, we seek an adaptive approach that augments a nominal intercept guidance solution (such as PN) to correct for evasive target maneuvering. One such architecture that appears to be suitable for surmounting this difficulty is shown in Figure D.6.[7] This architecture assumes that PN guidance will be augmented by an adaptive process in such a way that the plant behaves like the model response, where the model response is based on (D.10) with $a_T = 0$. In the ideal case, the main difference between these responses is due to the target acceleration. Therefore, the role of the adaptive process is limited to augmenting the PN law to cancel the effect of the target maneuver. In effect, its role is to estimate the target acceleration. We have shown in Ref 7 that the error dynamics for this architecture fits the form addressed in Ref 6, and therefore the NN based adaptive control laws developed there are applicable.

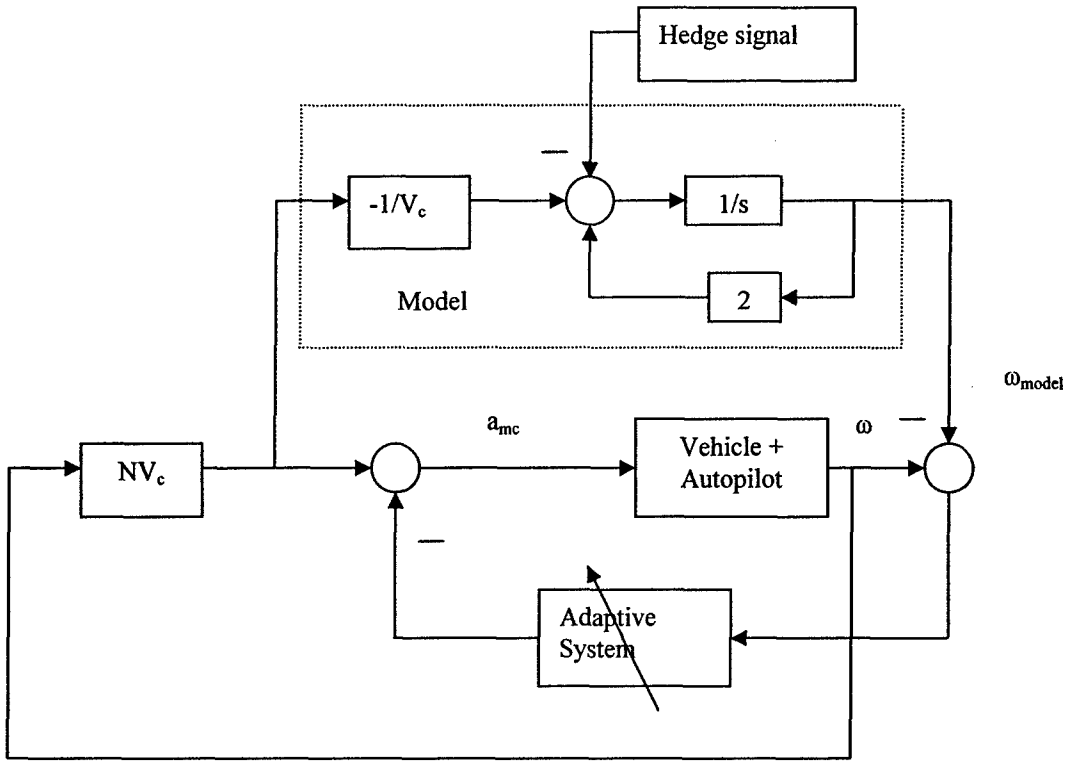

Figure D.6. Adaptive guidance architecture.

We conducted a preliminary test using the architecture in Figure D.6 and the adaptive control approach of Ref. 6. To keep matters simple, it was assumed that LOS rate and τ are known exactly. A second order model was used to represent the autopilot lag in the guidance loop, with $\omega_n = 2/T$ and $\zeta = 2/T/\omega_n$. This choice results in T being the effective guidance loop time constant, in the sense that the second order model approximates the rise time of a first order system with this time constant. We chose T = 0.5 for the study. For this model, the optimal target maneuver against a PN guidance law with N = 3 is a reversal at time-to-go = 1.2 seconds, close to the value predicted by Figure D.4 for a first order lag if the time axis is viewed as normalized by 1/T. The PN guidance law alone resulted in a miss distance of 44.8 feet. With adaptive guidance, the miss distance was reduced to 14.9 feet (approximately a 3 fold improvement).

Figure D.7 compares the resulting acceleration profiles in Gs. Note that the adaptive controller appears to be effective in estimating target acceleration. Following an initial transient, its acceleration profile decreases with time, similar in form to the ideal guidance law result that has perfect knowledge of target acceleration (see Figure D.4). Note that the initial response of the non-adaptive PN law is similar to that depicted in Figure D.2.

A second essential feature is that adaptive hedging[8] was also implemented to account for the autopilot lag, which changes the relative degree of the sensed LOS rate output. The entry

point of this signal is illustrated in Figure D.6. Its purpose is to subtract from the model response, the deficit in the achieved acceleration due to lag in the guidance loop.

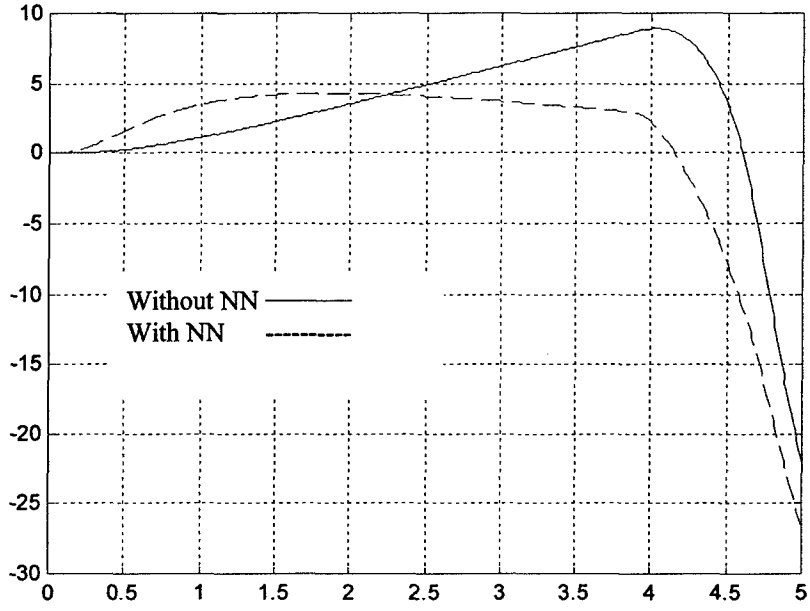

Figure D.7. Comparison of acceleration profiles.

As a next step we investigated the use of hedging to protect the adaptive loop from control saturation. Figure 8 examines the effect of an acceleration command limit of 5 g's, for the case of PN guidance (N=3) with an autopilot lag of 0.5 seconds, and a constant 3g target maneuver. Because the PN law results in an acceleration profile that increases linearly with time, the limit ultimately results in a miss distance of approximately 209 feet. The Adaptive law results in negligible miss distance. Note that the similarity of this result and the ideal behavior in Figure D.4 is even greater than what was observed in Figure D.7.

These results demonstrate the ability to correct for target acceleration without the use of a Kalman filter to estimate target acceleration. The approach also corrects the guidance law for the effect of nonlinear engagement dynamics that result in large rotations of the LOS angle. However, this is not apparent from the results in this example. This will be better illustrated in Section D.4.

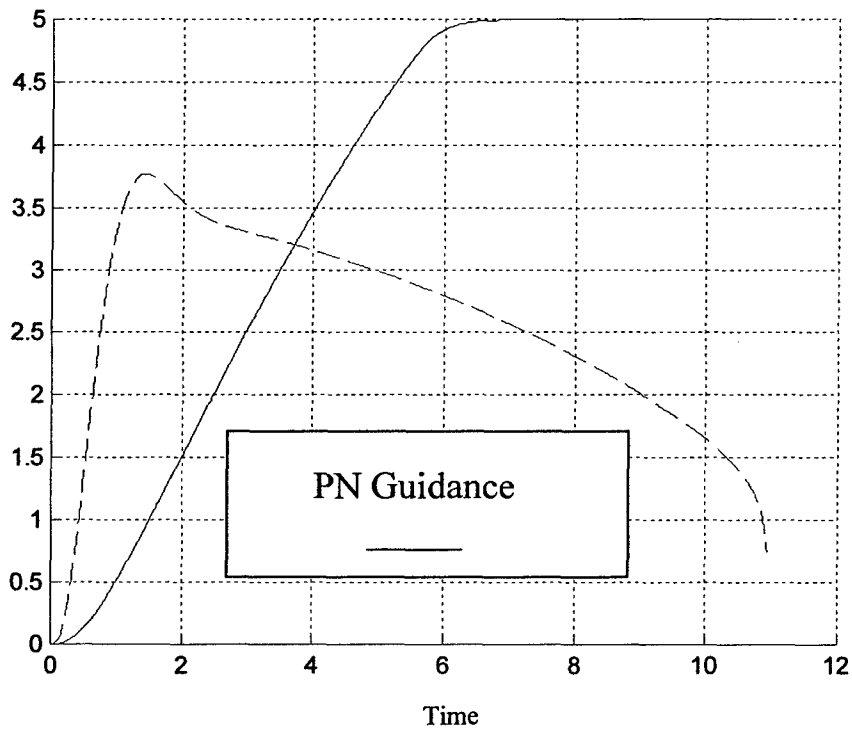

Figure D.8. Effect of an acceleration limit.

D.3 Adaptive Estimation

In this section we describe an alternative approach in which an EKF is employed to estimate LOS rate, closing velocity and target acceleration. First we describe the EKF problem formation, and then discuss how this formulation is augmented with an adaptive element to correct for modeling error.

EKF Problem formulation

The dynamics are written in a slightly different form, wherein the angles θ_m and θ_T are defined relative to an inertial frame, as depicted in Figure D.9

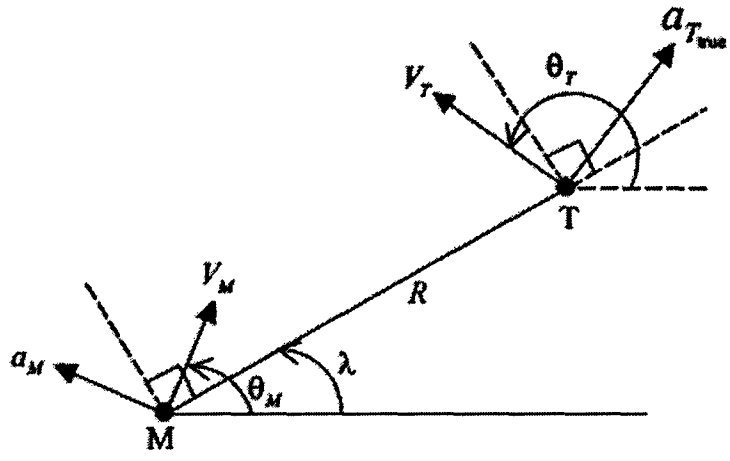

Figure D.9. State variables for the EKF formulation

With the angles defined in this manner the process dynamics model for the EKF design are:

$$\begin{aligned}
\dot{\lambda} &= \frac{V_T \sin(\theta_T - \lambda) - V_M \sin(\theta_M - \lambda)}{R} \\
\ddot{\lambda} &= -2\dot{\lambda}\frac{\dot{R}}{R} + \frac{a_T - a_M}{R} \\
\dot{R} &= V_T \cos(\theta_T - \lambda) - V_M \cos(\theta_T - \lambda) \qquad \text{(D.11)} \\
\ddot{R} &= \dot{\lambda}^2 R \\
\dot{a}_T &= (-a_T + w_T)/\tau_T
\end{aligned}$$

where w_T represents process noise, and a_T is the component of target acceleration normal to the LOS. The second and fourth equations in (D.11) are approximately correct so long as the lead

angle $(\theta_m - \lambda)$ is small and the component of target acceleration along the LOS is small. These approximations are one source of modeling error in this formulation. With this model, target acceleration is modeled as a first order Markov process, having a time constant τ_T. This represents a second potential source of modeling error, depending on the true target behavior. Note that the estimates for LOS rate, the closing velocity and the target acceleration come directly from the estimates of the first and fourth and fifth state variables in the model.

The measurement equations are:

$$\begin{aligned} y_R &= R + v_R \\ y_\lambda &= \lambda + v_\lambda \end{aligned} \qquad \text{(D.12)}$$

In addition, we treat a_M(t) as a know quantity when implementing the EKF. The standard discrete EKF filter equations that apply to the process dynamics in (D.11) and the measurement equation in (D.12) are obtained by linearizing and discretizing the dynamics and measurement equation about the last state estimate, to obtain

$$\begin{aligned} \delta x_{n+1} &= \phi(\hat{x}_n)\delta x_n + Bw_k \\ y_n &= Hx_n + v_n \end{aligned} \qquad \text{(D.13)}$$

where x is the state vector, $\delta x_n = x_n - \hat{x}_n^-$, $B^T = [0 \quad 0 \quad 0 \quad 0 \quad 1]$, $H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$, $\phi = I + \partial f / \partial x * dt$, and $f(x)$ represents the right hand side of (D.11) excluding the process noise. The filter update and prediction equations are:

Update

$$\begin{aligned} \mathbf{K}_k &= \mathbf{P}_k \mathbf{H}^T (\mathbf{H}\mathbf{P}_k^- \mathbf{H}^T + \mathbf{R})^{-1} \\ \mathbf{P}_k &= (\mathbf{I} - \mathbf{K}_k \mathbf{H})\mathbf{P}_k^- \\ \hat{x}_k &= \hat{x}_k^- + \mathbf{K}_k (y_k - \hat{y}_k^-) \end{aligned} \qquad \text{(D.15)}$$

where $\mathbf{K}_k$ is the Kalman Gain, $\mathbf{P}_k$ is the covariance associated with the estimation error, $y^T = [y_\lambda \quad y_R]$ is the measurement vector and $\mathbf{R}$ is the covariance matrix associated with the measurement error .

Prediction

$$\begin{aligned}\hat{x}_k^- &= \hat{x}_{k-1} + f(\hat{x}_{k-1}) * dt \\ \mathbf{P}_{k+1}^- &= \phi_{k-1}\mathbf{P}_k\phi_{k-1}^T + \mathbf{Q} * dt \\ \hat{y}_k^- &= \mathbf{H}\hat{\mathbf{x}}_k^-\end{aligned} \tag{D.16}$$

where dt is the sampling time interval, $\phi_{k-1} = \phi(\hat{x}_{k-1})$, $\mathbf{Q} = \sigma_w BB^T$, and σ_w^2 is process variance assumed for the target model.

Adaptive Augmentation of the EKF

The adaptive estimation problem is formulated in continuous time, but implemented in discrete time, much the same as is commonly done for the adaptive control problem. Thus the description given here is in the continuous time setting of (D.11), (D.12). Let the actual dynamics be represented in the following form

$$\begin{aligned}\dot{x} &= f(x) + B_1 z_1 \\ \dot{z}_1 &= h(z_1) \\ y &= Hx\end{aligned} \tag{D.17}$$

where $B_1^T = [0 \quad 1 \quad 0 \quad 0 \quad 0]$, and $z_1 = \tilde{a}_T / R$, with $\tilde{a}_T$ representing the unmodeled portion of the target acceleration (see the second equation in (D.11)). It is assumed that z_1 is an unknown but bounded process. We consider a NN approximation of z_1 have the form

$$z_1 = M^T\sigma(\mu) + \varepsilon(\mu) \tag{D.18}$$

where M is a vector of network weights, $\sigma(*)$ represents a vector of NN radial basis functions, and μ is a vector of network inputs comprised of difference quotients of the measurements. It has been shown that given any $\varepsilon^* > 0$ there exists a set of network weights $M = M^*$ and a delay $d > 0$ (used to form the difference quotients of the measurements), such at $h(z_1)$ can be approximated on a bounded set such that $|\varepsilon(\mu)|_F < \varepsilon^*$.[9]

We propose the following adaptive estimator for the dynamics in (D.11)

$$\begin{aligned}\dot{\hat{x}} &= f(\hat{x}) + K(t)(y - \hat{y}) + B_1\hat{W}^T\sigma(\mu) \\ \hat{y} &= H\hat{x}\end{aligned} \tag{D.19}$$

The NN weights are adapted according to the following adaptation law

$$\dot{\hat{M}} = -\gamma(\sigma(\mu)\tilde{y}^T + k_\sigma|\tilde{y}|\hat{M})$$
$$\tilde{y} = y - \hat{y} \qquad \text{(D.20)}$$

where $\gamma > 0$ is the adaptation gain and $k_\sigma > 0$ is the sigma modification gain. In Ref. 4 we prove that with this adaptation law, the estimation error and the NN weight errors remain bounded. Note that the form of the adaptive estimator in (D.19) is that of an EKF in which the dynamics of the estimator are augmented with an adaptive element. Ideally, the adaptive element is intended to estimate and cancel the effect of $z_1(t)$, which in this application amounts to the unmodeled portion of the target acceleration normal to the LOS. In applying this result to the discrete estimation equations in (D.16) we simply add in the adaptive term from (D.19) as a part of the prediction equation for the state estimate. We also replace the weight update law in (D.20) with its discrete equivalent. Finally, the estimate for target acceleration is constructed as

$$\hat{a}_T = \hat{x}_5 + \hat{M}^T\sigma(\mu) * \hat{R} \qquad \text{(D.21)}$$

D.4 NONLINEAR ENGAGEMENT RESULTS

In this section we examine two settings in which adaptation may be employed for a highly nonlinear, low-speed engagement scenario involving pursuit of a maneuvering target. The engagement is nonlinear because it involves very large angle rotations in the LOS and time periods during which the pursuing vehicle flies at maximum load factor. Some of the important engagement parameters are:

$T_s = 0.05$	guidance update rate (sec)
$\tau = 0.5$	effective guidance time constant (sec)
$\tau_t = 0.5$	target maneuver time constant (sec)
$a_{max} = 3.5$	pursuer acceleration limit (G's)
$T_s = 0.02$	seeker time constant (sec)
$T_n = 0.2$	noise filter time constant (sec)

Sensor errors include the following:

σ_{LOS} =0.01	standard deviation of the range independent error in LOS (rad)
σ_{LOS_g} =0.05/R	standard deviation of the glint error in LOS (rad)
σ_R =0.05*R	standard deviation of error in range (ft)

This model is somewhat pessimistic since it is doubtful that a vision based sensor will exhibit glint error. Also, the range error will depend on the density of the pixels in the image plane, and will not increase proportionately with range. However, beyond a certain range is will be

impossible to discern the image size, and range information will not available. So the standard deviation of the range error is more likely one that is range independent out to a certain range, and then suddenly increases to a large value beyond that range.

The target performs a random maneuver in which the maneuver direction is decided by the sign of a random number generator in 1 second intervals. Therefore the target acceleration command appears like a random square wave of amplitude $\pm 2g's$. This command is passed through a first order filter, whose time constant is τ_t. The resulting acceleration profile is shown in Figure D.10. We regard this as a highly maneuvering target in the sense that reversals in the acceleration profile occur frequently, in relation to the effective guidance time constant of the pursuer.

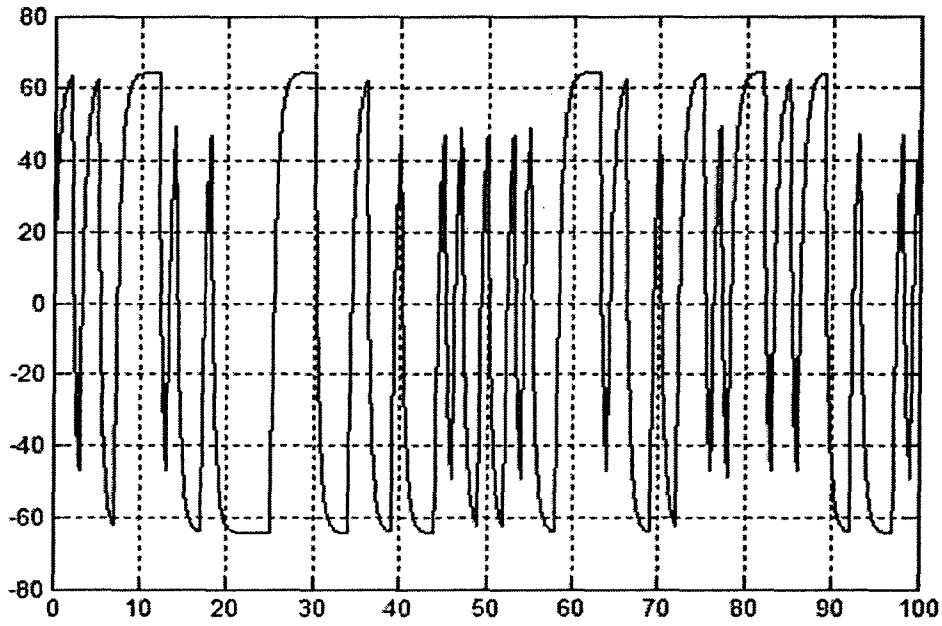

Figure D.10. Target acceleration profile (ft/s/s).

Adaptive Guidance Results

We first show the results obtain when augmenting the PN guidance law as depicted previously in Figure D.6. These results were obtained using the classical method for estimating LOS rate shown in Figure D.4, with the LOS measurement error and seeker and filter time constant settings given above. Closing velocity was estimated using a constant value of 20 ft/s, and time-to-go was estimated using $\tau = \hat{R}/20$, where $\hat{R}$ is the simulated range measurement, with the range dependent measurement error proportional to R as specified above.

Figure D.11 shows the resulting trajectories for the pursuer and the target, and Figure

D.12 shows the range profile. These figures correspond to the case with adaptive augmentation of the PN guidance law. Note that the pursuer makes several passing attacks on the target. In three of these attacks the range is reduced to a relatively small value, approximately between 5 and 40 feet. The pursuer's acceleration profile is shown in Figure D.13, where is it evident that there are several periods during which the control is saturated.

To see the effect of target evasive maneuver, the sample rate in the random number generator used in the target model was increased from 1 to 2 second intervals. This has the effect of simulating a less aggressive target having the same maximum target acceleration. The target acceleration profile is shown in Figure D.14, and the resulting effect on the trajectory and range profiles is shown in Figures D.15 and D.16. Note that in this case the target was intercepted (range < 1 ft).

The trajectory and range profiles without adaptation (using the original target model) are depicted in Figures D.14 and D.15. Note that once the point of closest approach is reached, the pursuer never turns back toward the target. This is characteristic of PN guidance.

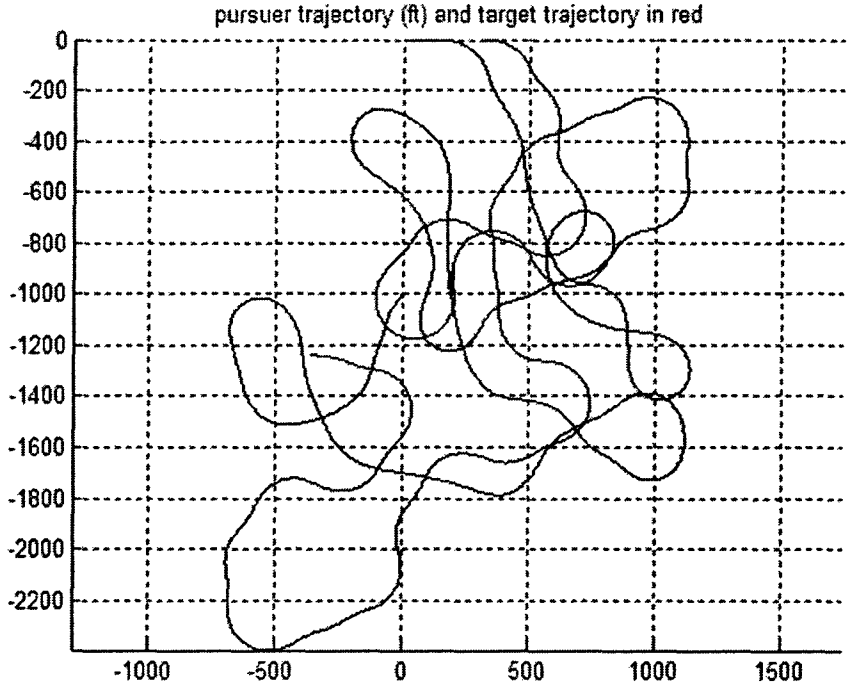

Figure D.11. Pursuer and target trajectory with adaptation.

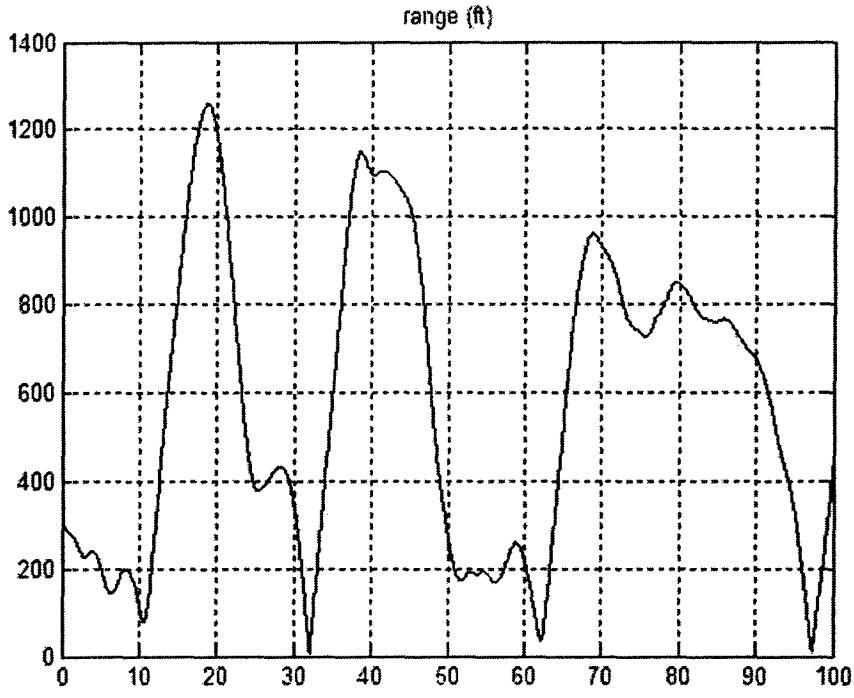

Figure D.12. Range profile with adaptation.

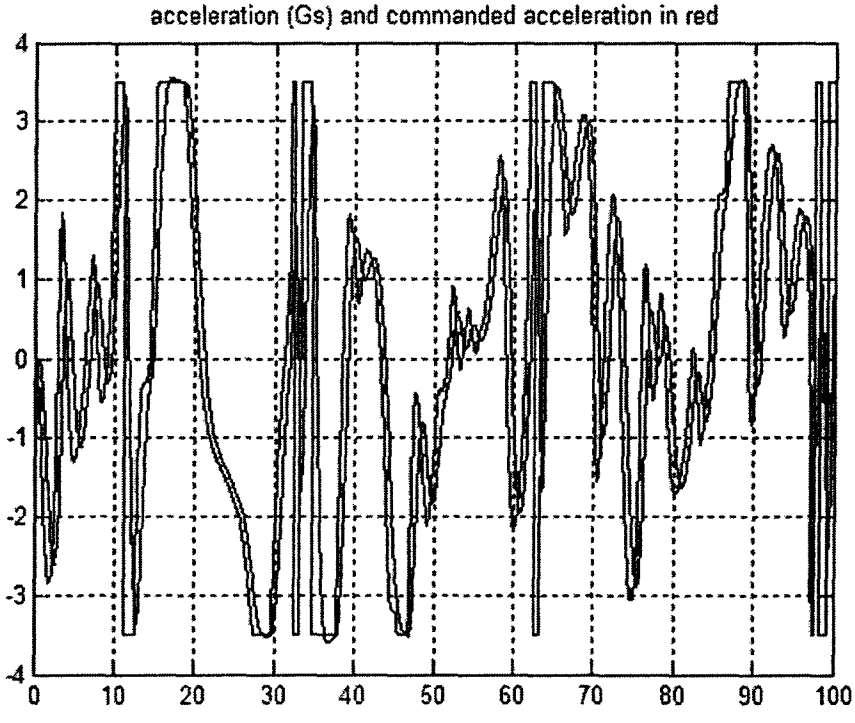

Figure D.13. Pursuer's acceleration profile with adaptation.

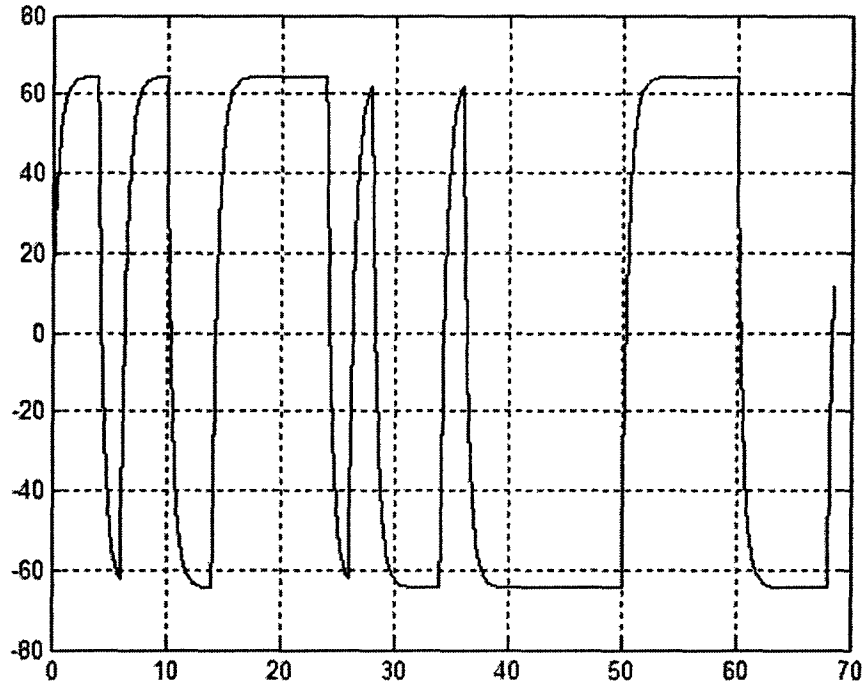

Figure D.14. The less aggressive target acceleration profile, (ft/s/s).

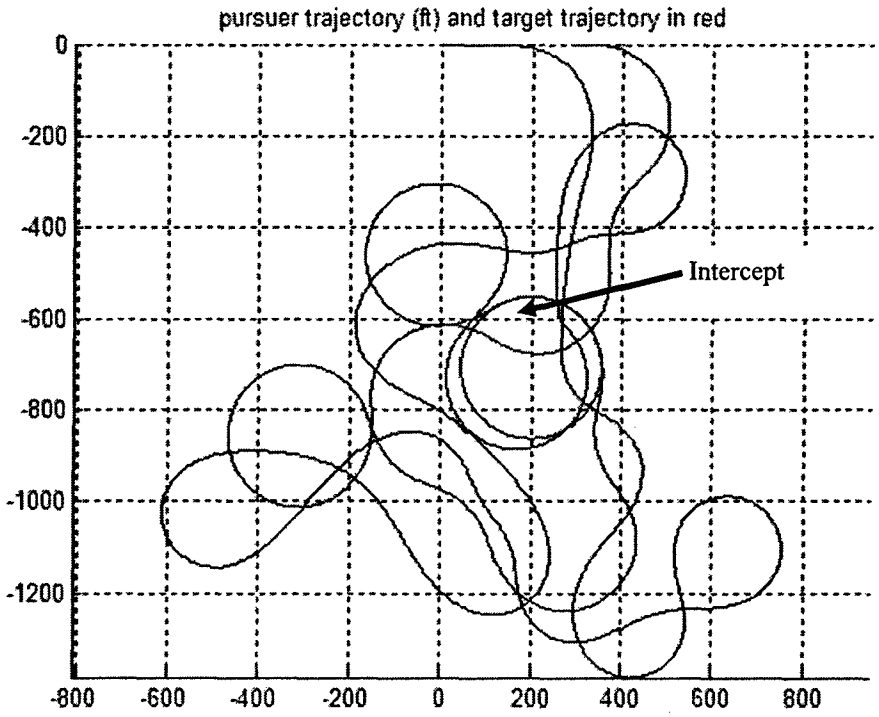

Figure D.15. Trajectories with adaptation for the less aggressive target.

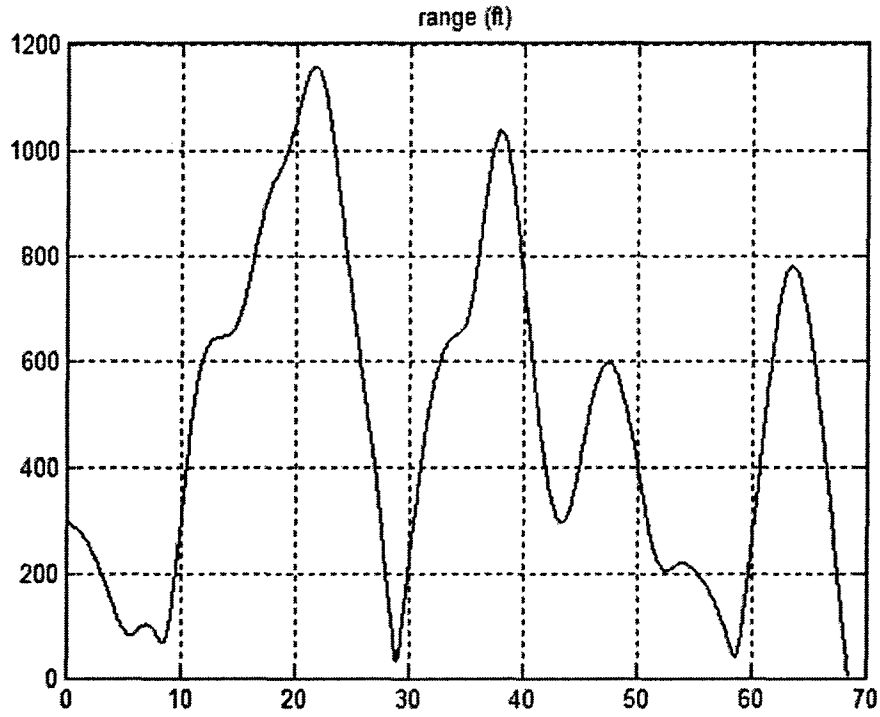

Figure D.16. Range profile with adaptation for the less aggressive target.

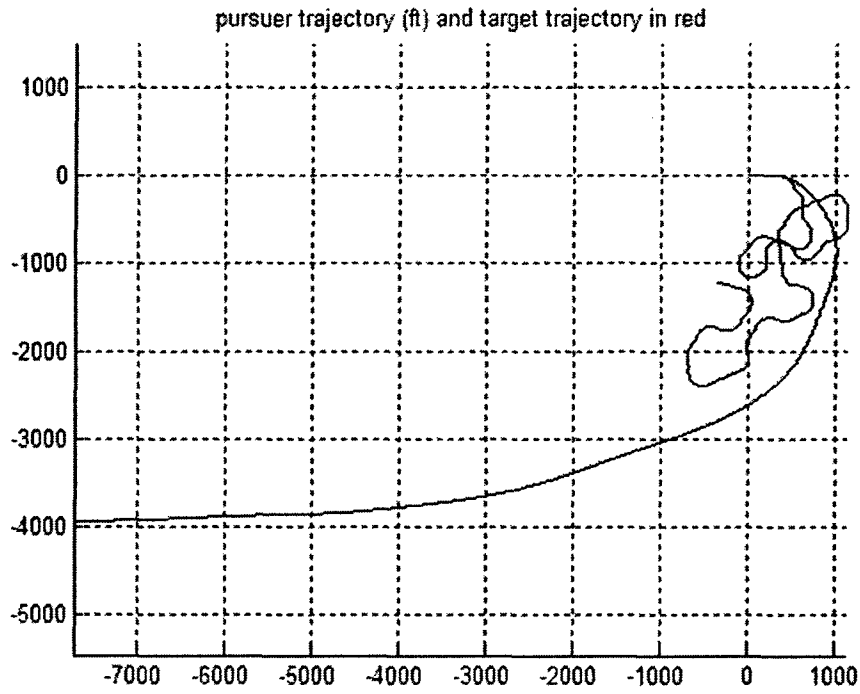

Figure D.17. Pursuer and target trajectory without adaptation.

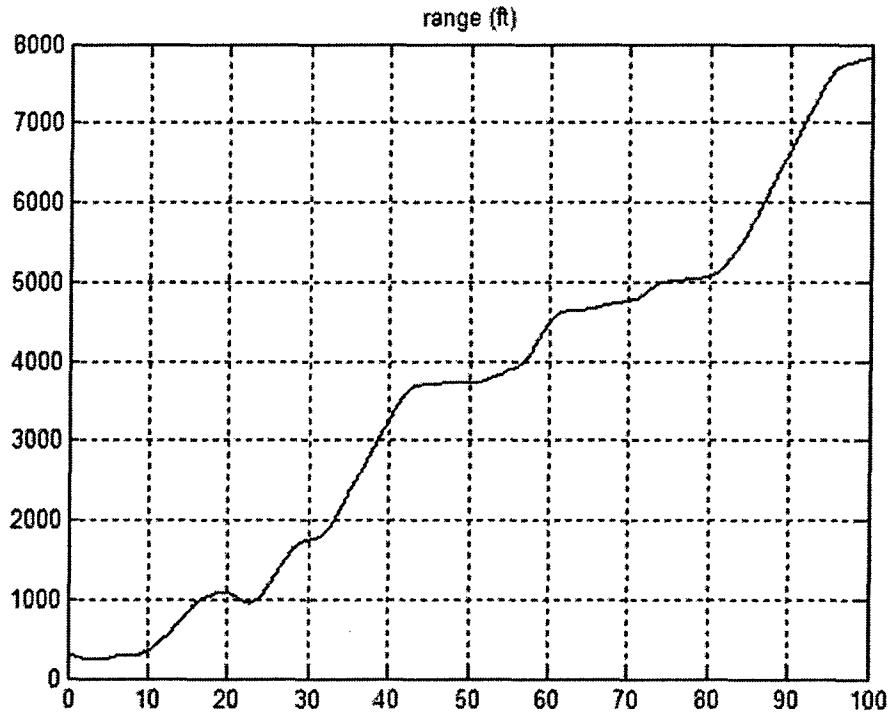

Figure D.18. Range profile without adaptation.

Adaptive Estimation Results

We now examine the performance of the APN guidance law that employs an EKF to estimate LOS rate, closing velocity and target acceleration. Figure 19 shows the trajectories for the pursuer and the target, and Figure D.20 shows the range profile for the EKF with adaptation. Comparing Figure D.20 with Figure D.12 it is evident that the performance is not as good as that obtained when the adaptive PN law. Figures D.21-23 show the performance of the adaptive EKF as a state estimator. Note that the estimation performance is not that good with respect the LOS rate, range rate and target acceleration. Also, we were not able to intercept the less aggressive maneuvering target. In fact, there was no discernable improvement in this case.

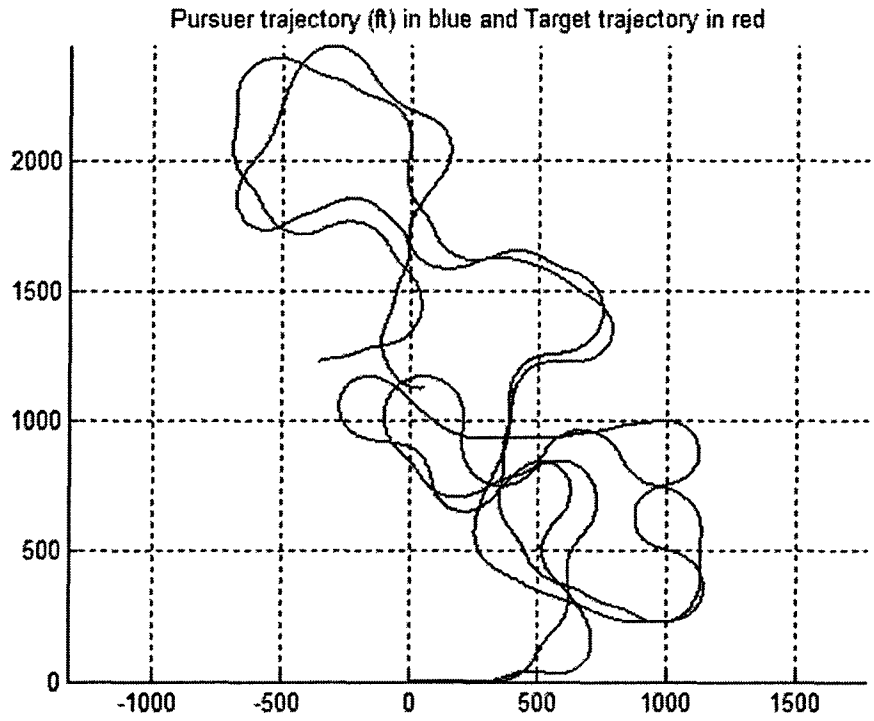

Figure D.19. Pursuer and target trajectory with adaptation in the EKF.

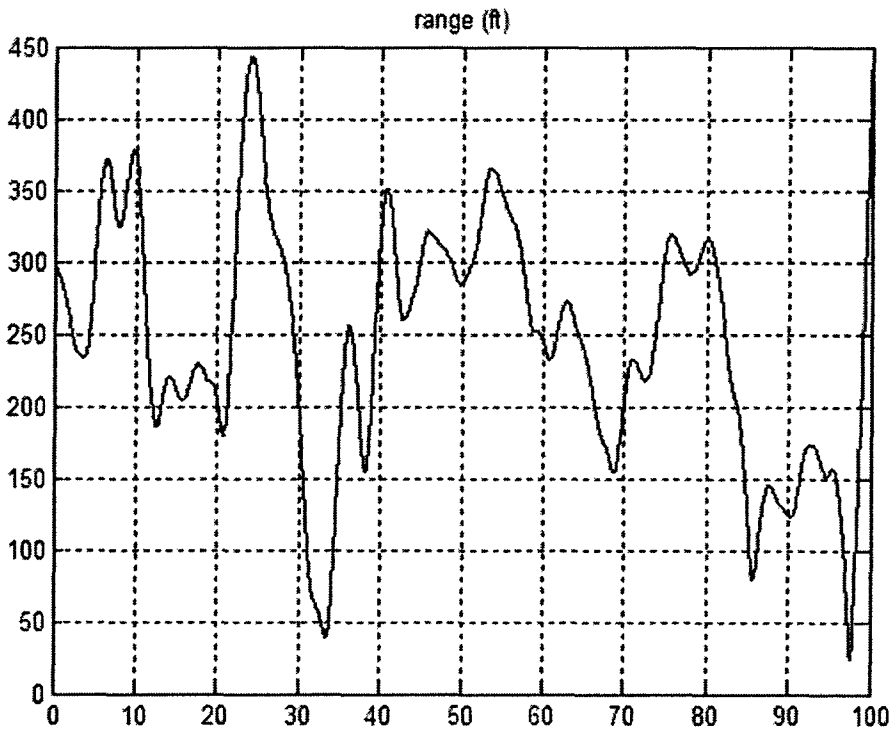

Figure D.20. Range profile with adaptation in the EKF.

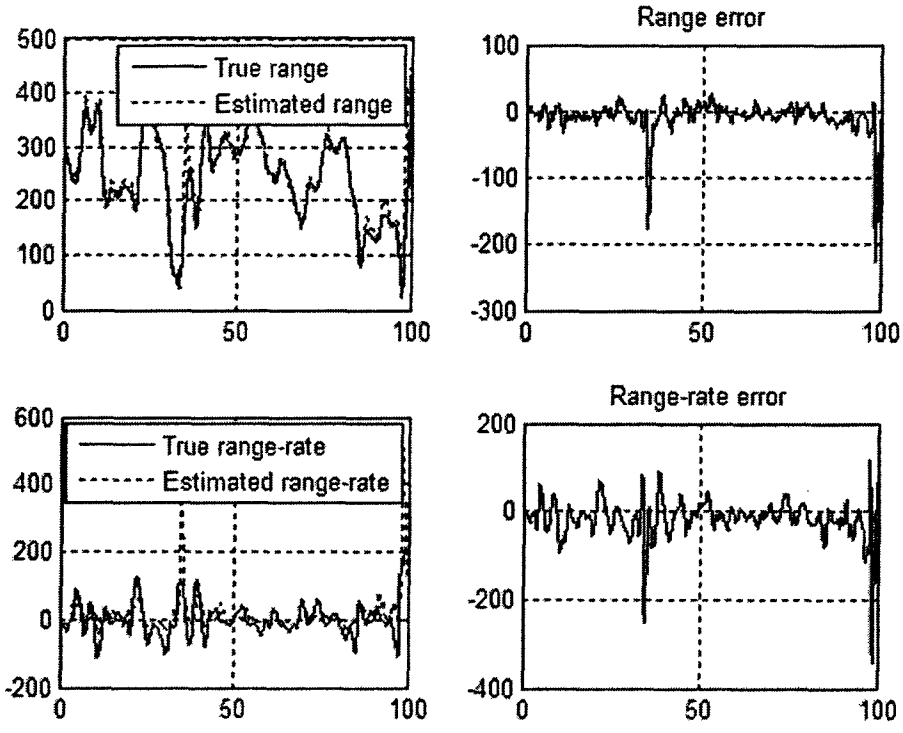

Figure D.21. Estimation performance for range and range rate, with adaptation.

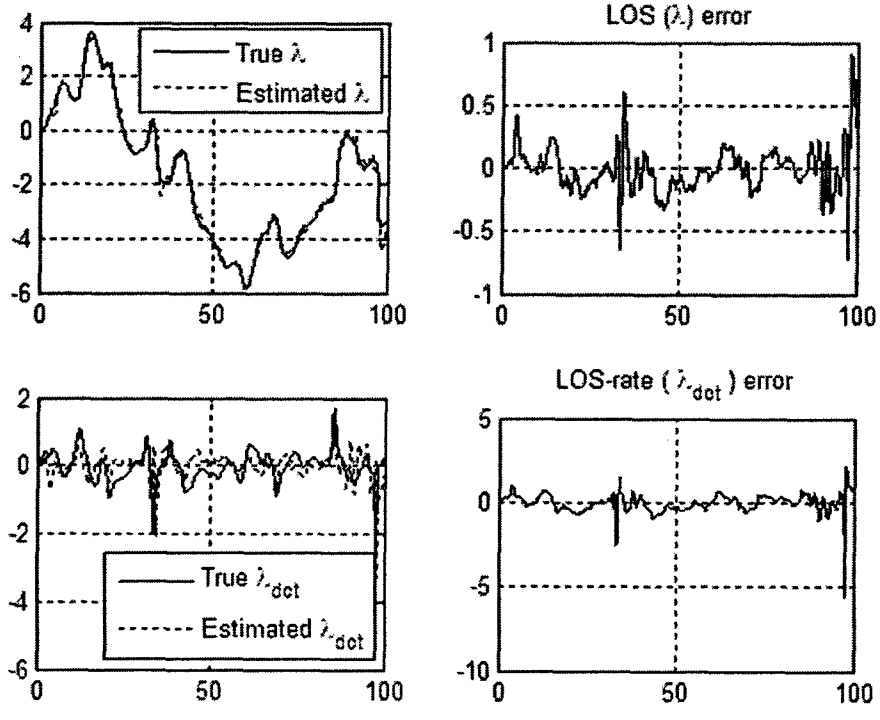

Figure D.22. Estimation performance for LOS and LOS rate, with adaptation.

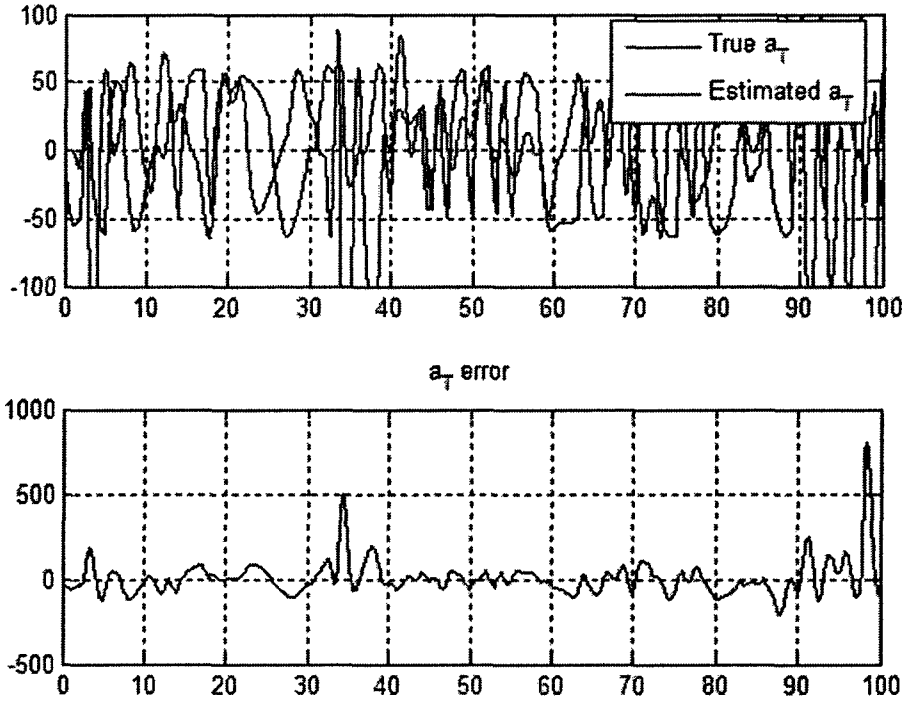

Figure D.23. Estimation performance for target acceleration, with adaptation.

D.5 Summary and Future Research

Two methods for adaptive guidance have been explored. The first applies adaptation to a PN guidance law and attempts to compensate for target acceleration, without employing a Kalman filter. The second approach adaptively augments an EKF designed to estimate target acceleration, and uses this estimate in an APN guidance law. The former approach seems to be superior to the latter. However, it should be noted that the status of the research in adaptive estimation as it applies to an EKF is not complete at this point, and further research in this direction is warranted. Another direction that can be taken is to combine adaptive guidance with adaptive estimation.

E. SUMMARY AND CONCLUSIONS

The program produced results in three main areas: understanding of the tracking and guidance functions associated with male flesh fly pursuits, improvements to vision based tracking algorithms; and the development of adaptive methods for pursuit guidance.

The work at Cornell centered on developing experimental methods to characterize flesh fly pursuit evasions, and resulted in the maturation of effective means to capture the 3-D trajectory, as well as body and head orientation. The data was processed at first by hand, and later using image processing algorithms to develop 3-D visualizations of the track, including the head orientation, and ultimately to map the location of the target on the eye during the pursuit. The results provided a means to compare the guidance strategy of the fly with traditional proportional navigation, and to look for inspiration in the development of new guidance laws. Work was also completed to introduce clutter into the encounter. While a much greater understanding of the tracking and guidance strategy of the flesh fly was developed and documented, the work has not yet resulted in the discovery of a better alternative to traditional engineered guidance laws.

Work by Prof. Allen Tannenbaum and his students at Georgia Tech was focused on specific improvements to the set of image processing tools needed to employ vision as the primary sensor modality in the development and implementation of an advanced UAV guidance and control system. Specific emphasis was placed on the use of dynamic snakes for tracking. Unexpectedly, the collaboration of scientists in entomology with researchers in the field of image processing resulted in the opportunity to employ image processing tools in the experimental work to great advantage.

The work in guidance law development focused on the introduction of adaptation to target maneuvers. Two approaches were pursued, in the first, a neural network was used to augment a proportional navigation law directly, without the use of estimation. In the second, a neural network was used to augment an extended Kalman filter. Both approaches were successful, and represent an important new contribution to the field of pursuit guidance, and to the field of estimation in general.

Numerical simulation of characteristic pursuits was employed to illustrate the final results. While program time and resources did not permit transition to an actual pair of UAVs for flight demonstration, the effort created a very strong foundation for the development of an AFOSR MURI program in controlled active vision that is now carrying collaborative research in fields of image processing and guidance and control forward to planned flight demonstration on a pair of fixed-wing UAVs.

REFERENCES

1. Bryson, A.E., and Ho, Y.C., *Applied Optimal Control,* Blaisdell, Waltham, MA, 1969.

1. Zarchan, P., *Tactical and Strategic Missile Guidance,* Progress in Astronautics and Aeronautics, A. Richard Seebass, Volume 124, AIAA, Washington, DC, 1990.

3. Nesline, F.W., Zarchan, P., "A New Look at Classical vs Modern Homing Missile Guidance," AIAA J. of Guidance and Control, Vol. 4, No. 1, Jan.-Feb. 1981.

4. Madyastha, V.K., Calise, A.J, "An Adaptive Filtering Approach to Target Tracking," American Control Conference, June 2005.

5. Hovakimyan, N., Calise, A.J., and Madyastha, V.K., "An Adaptive Observer Design Methodology for Bounded Nonlinear Processes," Conference on Decision and Control, December 2002.

6. Calise, A. J., N. Hovakimyan, and M. Idan, "Adaptive Output Feedback Control of Nonlinear Systems Using Neural Networks." Automatica, Vol. 37, No. 8, August 2001.

7. Calise, A. J., Yang, B.-J., and Craig, J. I., "Augmenting Adaptive Approach to Control of Flexible Systems," Journal of Guidance, Control and Dynamics, Vol.27, No.3, 2004.

8. Johnson, E.N., Calise, A., "Limited Authority Adaptive Flight Control for Reusable Launch Vehicles," Journal of Guidance, Control and Dynamics, Vol. 26, No. 6, Nov-Dec. 2003.

9. Lavretsky, E., Hovakimyan, N., Calise, A.J., Upper Bounds for Approximation of Continuous-Time Dynamics Using Delayed Outputs and Feedforward Neural Networks, IEEE Transactions on Automatic Control, vol. 48 (9), 2003.