

ABSTRACT

Title of thesis: A COMPARISON OF OPTIC FLOW
 IN THE VISIBLE LIGHT
 AND INFRARED SPECTRUM

Michael Chinn, Master of Science, 2008

Thesis directed by: Professor J. Sean Humbert
 Department of Aerospace Engineering

Insects use a method of Wide Field Integration (WFI) to navigate efficiently through unknown environments. Using these natural paradigms, various WFI based forms of navigation can be implemented based on electro-mechanical vision devices on robotic vehicles. However, under low light and/or suspended particles in the environment, these methods become less useful. One solution to this problem is to use infrared vision sensors rather than visible light sensors. This would allow insect-like navigation for autonomous vehicles under a variety of lighting conditions, including a total lack of visible light. The results show that, using infrared sensors, it is possible to navigate under a variety of lighting conditions, even where visible light sensors become ineffective.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Comparison of Optic Flow in the Visible Light and Infrared Spectrums				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland, College Park, College Park, MD, 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Insects use a method of Wide Field Integration (WFI) to navigate efficiently through unknown environments. Using these natural paradigms, various WFI based forms of navigation can be implemented based on electro-mechanical vision devices on robotic vehicles. However, under low light and/or suspended particles in the environment, these methods become less useful. One solution to this problem is to use infrared vision sensors rather than visible light sensors. This would allow insect-like navigation for autonomous vehicles under a variety of lighting conditions including a total lack of visible light. The results show that, using infrared sensors it is possible to navigate under a variety of lighting conditions, even where visible light sensors become ineffective.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 54	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

COMPARISON OF OPTIC FLOW IN THE
VISIBLE LIGHT AND INFRARED SPECTRUMS

by

Michael Chinn

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2008

Advisory Committee:
Professor J. Sean Humbert, Chair/Advisor
Professor David Akin
Professor Robert Sanner

© Copyright by
Michael Chinn
2008

Acknowledgments

First and foremost I would like to thank my advisor Sean Humbert for making this possible. I would also like to thank my fellow graduate students, Joe Conroy and Andrew Hylsop for all their technical help. I would also like to thank my family and friends, especially Chris, Laurie, Mari, and Jess, for their support and keeping me going when times were tough. And finally I would like to thank DARPA for the funding and research opportunity under grant number N66001-07-C-2057.

Table of Contents

List of Figures	iv
1 Wide Field Integration	1
1.1 Review of Optic Flow Formulation	1
1.1.1 Summary of Planar Optic Flow	3
1.1.2 WFI Sensory Outputs	4
1.2 WFI Implementation on the Vehicle	4
2 Vehicle Design and Specifications	6
2.1 Overview	6
2.2 Vehicle Hardware	7
2.3 Dynamic Model and System ID	10
2.4 Vehicle Closed Loop Performance	13
3 Experimental Setup	19
3.1 Camera Configuration	19
3.2 Physical Layout	20
3.3 Light Measurement	20
3.4 Experiment Procedures	23
3.4.1 Pure Rotation	23
3.4.2 Pure Translation	23
3.5 Image Processing	24
3.5.1 Method of Processing Data	24
3.6 Methods of Determining Image Shift	25
3.6.1 Gradient Method	25
3.6.2 Lucas-Kanade	25
3.6.3 Data Reduction	26
4 Experimental Results	28
4.1 Low Light Conditions	28
4.1.1 Translation	28
4.1.2 Rotations	28
4.2 Obscured Line of Sight	29
4.2.1 Translation	32
4.2.2 Rotation	33
5 Conclusions	43
A Data	45
Bibliography	47

List of Figures

1.1	Planar tunnel geometry. (A) Notation and vehicle state definitions (B) Balanced planar nearness function $\mu(\gamma)$ and lateral/rotational perturbations of μ	2
1.2	Environment and Optic Flow Generation	2
2.1	Original Robot Design 1	7
2.2	Original Robot Design 2	8
2.3	Computer Flow Chart	9
2.4	X80 Base Configuration	10
2.5	Final Vehicle Configuration	11
2.6	Data from Translational Model	14
2.7	Data from Translational Model	15
2.8	Data from Rotational Model	16
2.9	Data from Rotational Model	17
2.10	Vehicle performance	18
3.1	View of the Rotation Experiment Setup	21
3.2	View of the Rotation Experiment Setup	22
4.1	Data from Translational Low Light Experiments	29
4.2	Data from Infrared Translational Low Light Experiments	30
4.3	Data from Visible Translational Low Light Experiments	31
4.4	Data from Rotational Low Light Experiments	32
4.5	Data from Infrared Rotational Low Light Experiments	33
4.6	Data from Visible Rotational Low Light Experiments	34
4.7	Calibration Images from Fog Translational Tests	35

4.8	Calibration Images from Fog Rotation Tests	36
4.9	Data from Translational Fog Experiments	37
4.10	Temporal Data from Visible Light Translational Fog Experiments . . .	38
4.11	Temporal Data from the Infrared Translational Fog Experiments . . .	39
4.12	Data from Rotational Fog Experiments	40
4.13	Temporal Data from the Visible Light Rotational Fog Experiments . .	41
4.14	Temporal Data from the Infrared Rotational Fog Experiments	42
A.1	Data from the Low Light Translational Experiments	45
A.2	Data from the Low Light Rotational Experiments	45
A.3	Data from the Fog Translational Experiments	46
A.4	Data from the Fog Rotational Experiments	46

Chapter 1

Wide Field Integration

1.1 Review of Optic Flow Formulation

Standard closed loop control design techniques involve sensing, estimation, and actuation. Typically, the sensing and control procedures are fast procedures. However, the estimation procedures generally induce the most latency into the system. This slows how quickly a closed loop system can respond to environmental factors. Biological systems do not use the massive amounts of processing and estimation for navigation and control that their human created mechanical counterparts do. Instead, insects in particular use a more efficient method for closed loop control. They eliminate the estimation process completely and directly feed the sensory input to their actuation systems. This form of Wide Field Integration (WFI) takes the visual information in tangential cells in their eyes and compares it to predetermined patterns. The differences in the patterns, being picked up by their eyes and the predetermined patterns are then used to actuate the insect in order to avoid obstacles. Using multiple patterns to compare against the perceived optic flow allows for parallel processing, which further decreases an insects response time to stimuli.

1.1.1 Summary of Planar Optic Flow

Optic flow, a relative measure of speed/depth, provides a relative proximity and speed with respect to obstacles in the environment. This can be used for navigation and collision avoidance. Figure 1.2 shows a vehicle in a general environment and the apparent velocity of the objects in its field of view. The apparent velocity is perpendicular to the ray from the center of the vehicle to the edge of the object in question. The distance $r(\gamma, \mathbf{q})$, where $\mathbf{q}$ is the position and orientation of the vehicle and γ is the viewing angle, is bounded between 0 and ∞ . If the minimum value of $r(\gamma, \mathbf{q})$ is forced to be some small, but non-zero number ϵ then the reciprocal,

$$\mu(\gamma, \mathbf{q}) = \frac{1}{r(\gamma, \mathbf{q})} \quad (1.1)$$

is bounded between 0 and some large, but finite number. This allows the formulation of an optic flow function $\dot{Q}$, which is piecewise continuous, and might have a finite number of discontinuities. For a 2-D planar case, the equation is given as

$$\dot{Q}(\gamma, \mathbf{q}, \dot{\mathbf{q}}) = -\dot{\theta} + \mu(\gamma, \mathbf{q}) (\dot{x}_b \sin \gamma - \dot{y}_b \cos \gamma), \quad (1.2)$$

where $-\dot{\theta}$ is the rotational component, and $\mu(\gamma, \mathbf{q}) (\dot{x}_b \sin \gamma - \dot{y}_b \cos \gamma)$ is the translational component, which contains information about the distances to objects in the environment.

Wighting Function	Tunnel Geometry	Linearization ($\mathbf{x}_0$)
$\cos \gamma$	$\frac{4a}{3\pi(a^2-y^2)} \dot{x}_b \sin \theta \cos \theta$	$\frac{4v_0}{3\pi a} \theta$
$\cos 2\gamma$	$-\frac{y}{2(a^2-y^2)} \dot{x}_b \cos \theta$	$-\frac{v_0}{2a^2} y$

Table 1.1: Spatial Fourier Decomposition of Planar Tunnel Optic Flow for Vehicles with Sideslip Constraint (Reproduced from [4])

1.1.2 WFI Sensory Outputs

Previous work [2] and [3] shows that perturbations form a desired pattern of optic flow can be characterized by a small number of Fourier harmonics, which correlate to relative proximity and speed with respect to the environment. The information extraction model, termed Wide-Field Integration, is given in equation 1.3 where u_i is the output corresponding to a particular weighting function F_i .

$$u_i = \langle \dot{Q}, F_i(\gamma) \rangle = \frac{1}{\pi} \int_0^{2\pi} \dot{Q}(\gamma, \mathbf{q}, \dot{\mathbf{q}}) F_i(\gamma) d\gamma \quad (1.3)$$

1.2 WFI Implementation on the Vehicle

The practical application of WFI to the vehicle involves feeding back the first and second cosine harmonics for navigation. Table I shows the linearized correspondence from states of the vehicle to Fourier coefficients. Equation 1.4 shows the

general control scheme for the vehicle where u is the rotational velocity input and is used to actuate a centering response.

$$u = K_1 \langle \dot{Q}, \cos \gamma \rangle + K_2 \langle \dot{Q}, \cos 2\gamma \rangle \quad (1.4)$$

Chapter 2

Vehicle Design and Specifications

2.1 Overview

The specifications for the vehicle are given as a ground vehicle which can carry the electronics necessary to compute optic flow, have a top speed of 1 m/s, can be no larger than 1 ft in diameter, and must be able to rotate in place. A secondary specification is that the vehicle could drive both forwards and backwards. Also, the vehicle needs to be kinematic, that is, well modeled by first order differential equations rather than the second order equations derived from physics.

The first version of the vehicle was constructed from scratch. The chassis is designed to be two disks, mounted one above another. The motors, wheels, and computer are designed to be mounted on the lower disk, while the cameras, mirrors, and other various pieces of equipment are mounted on the upper disk. The motors are designed to be controlled by PIC microprocessors, and contain a feedback loop using wheel encoders for sensors. However, the encoders used were too coarse to maintain fine control in the lower speed regime. This leads to the acquisition of a prepackaged system to supply the chassis and drive system. Figures 2.1 and 2.2 show the original prototype.

The acquisition of the prepackaged chassis leads to the relaxing of the diameter of the vehicle, which is now 15 in. However, the secondary specifications are met




Figure 2.1: Original Robot Design 1

with this system as well as the slightly relaxed primary specifications. While the current vehicle does not meet the letter of the initial specifications, it meets the spirit of them.

2.2 Vehicle Hardware

The vehicle used is build off of an X80 robot chassis made by Dr. Robot. The X80s motor control board, drive motors, battery, and wheel encoders are all off the shelf items from Dr. Robot. The base X80 system is show in figure 2.4. All the Dr.




Figure 2.2: Original Robot Design 2

Robot systems are run by their proprietary software, taking in position or velocity commands for the wheels, and returning recorded velocities and positions of each wheel. The connection to the low level controller on the vehicle is a serial connection. The main onboard computer is essentially a miniaturized desktop computer built off of a Biostar motherboard with a 2.4 GHz AMD CPU. The hard drive is a 250 GB 7200 RPM drive produced by Sea Gate. A fire wire PCI expansion board by Point Grey Research provides the interface between the computer and the onboard camera. There are two main sources for power for the vehicle: an 8.4 V NiMH battery




Figure 2.3: Information and Actuation Flows

supplied by Dr. Robot which runs the X80 power systems, and a 15.8 V, 10,000 amp-hour LiPo which supplies the onboard computer. These batteries provide enough power to operate the vehicle for at least two hours of continuous use. Depending upon the load placed on the electrical system, this can be extended considerably. However, there is no good model for power consumption beyond the two hour use. Communication with the computer is done via wireless. The vehicle is typically run using a remote desktop protocol, though direct interfacing through a monitor, keyboard and mouse is also possible. Figure 2.3 shows how the different components on the vehicle interact with each other and the environment, and figure 2.5 shows the final vehicle configuration.




Figure 2.4: X80 Vehicle

2.3 Dynamic Model and System ID

In order to talk meaningfully about control a dynamic model of the vehicle must be determined. It is assumed that vehicle motion can be described with the following kinematic model




Figure 2.5: Final Vehicle Configuration

$$\begin{aligned}\dot{x} &= \dot{x}_b \cos \theta \\ \dot{y} &= \dot{x}_b \sin \theta\end{aligned}\tag{2.1}$$

$$\dot{\theta} = u_i\tag{2.2}$$

where it is assumed that $\dot{x}_b = V_0 = \text{constant}$. To test the validity of these assumptions, a system identification test was performed on the vehicle. The following frequency domain model was assumed:

$$V(s) = G_t(s)\dot{X}_b(s) \quad (2.3)$$

$$U(s) = G_r(s)\dot{\theta}(s)$$

where $G_t(s) = \frac{K_t}{1+T_t s}$ and $G_r(s) = \frac{K_r}{1+T_r s}$ and V, U were the translational and rotational inputs respectively. If the kinematic assumption is valid, $K_t \simeq K_r \simeq 1$ and $T_t \simeq T_r \simeq 0$.

Data on the vehicle's position and velocity is collected using the Vicon System. This system is capable of retro reflective tracking markers on a vehicle at 350 Hz. The system makes a model of the vehicle and returns positions and orientations. The number of markers and placement dictates the error in the data. However, most models give an accuracy of less than 2 mm, and an orientation of less than 0.1 radians. Several markers are placed asymmetrically on the vehicle for tracking. A model is created in the Vicon software. The vehicle stores all commands sent to the wheels, in addition to external sensors and time stamps of the data. To control the vehicle, a Labview program was created that takes in commands from a joystick and translates them to wheel velocity commands. Two different modes are modeled during the tests: translation and rotation. Using the joystick, a frequency rich series of commands is sent to the vehicle. To synchronize the data between the Vicon system and the vehicle data, an integrated rate gyro is placed on the vehicle. This is used with the heading information provided by Vicon to match the vehicle data with the positions and velocities. At the start of each test, a series of pseudo-random rotations are made to make this process easier [1]. The data on both

Vicon and on the onboard computer is taken at 500 Hz. Due to issues with memory allocation during the start of a program, the onboard computer data is linearly interpolated to match the Vicon data. This, with the initial rotations, allows the data from the two systems to be correlated.

The results show that a first order model predicts the behavior of the vehicle well, both in rotation and in translation.

Results for the models and the recorded data can be seen on Figure 2.6 for translation, and Figure 2.8 for rotation. Figures 2.7 and 2.9 show that the model tracks the general behavior, but is unable to account for the higher order oscillations in the velocity.

K_t is determined to be 1.140, and T_t is 0.0376. Likewise, K_r is 1.066 and T_r is 0.0209. A second order mode is also present with a natural frequency of about 14 Hz, but is much smaller in magnitude than the first order mode. Since the time constants T_t and T_r are very small, the kinematic approximation is a good model for the vehicle.

2.4 Vehicle Closed Loop Performance

In order to verify the vehicle performance, several different tests are run. An important test is the vehicle's closed loop performance. One variant of the vehicle uses a single Firefly camera looking up at a parabolic mirror to obtain a full field of view of the environment around the vehicle. Using the optic flow algorithms developed previously and in [4], the vehicle is capable of traversing a number corridors




Figure 2.6: Measured and modeled outputs for the Translational system

and cluttered environments without colliding with obstacles. Figure 2.10 shows the performance of the vehicle in a corridor with a 90° bend, as well as the Fourier coefficients and the optic flow at various points. In this experiment, the vehicle was run closed loop down the corridor with just the sensory information from the onboard camera. The vehicle was started near and pointed towards the right wall. Figure 2.10C shows the optic flow for the vehicle turning to avoid the wall, as well as the large mismatch in the magnitude of the left and right sides optic flow. Once the vehicle has centered itself in the corridor, it then turns right to avoid crashing into the walls and attempts to stay centered on its new heading. Figure 2.10D shows




Figure 2.7: Measured and modeled outputs for the Translational system

a typical optic flow pattern for driving down the middle of a straight corridor or hallway.




Figure 2.8: Measured and modeled outputs for the Rotational system




Figure 2.9: Measured and modeled outputs for the Rotational system




Figure 2.10: Vehicle performance, Fourier coefficients, and optic flow. (A) shows the closed loop performance down a corridor with a 90° bend, (B) shows the Fourier coefficients during the experiment, and (C) and (D) show the optic flow at two different points during the experiment.

Chapter 3

Experimental Setup

3.1 Camera Configuration

Two different cameras are used for these experiments. These are a Thermal-Eye 4550AS camera from L-3 communications and a Firefly MV camera from Point Grey Research. The field of view on the Thermal-Eye camera is 50° and the field of view of the Firefly is 30° . Both cameras are mounted facing the same direction, perpendicular to the forward direction of the vehicle. The field of view for the Thermal-Eye is given in the data sheet shipped with the camera. The field of view for the Firefly is measured, since only part of the Firefly's maximum image capture size is used.

A Matrox Morphis Quad frame grabber card is used to collect the images from the Thermal-Eye camera. Due to hardware conflicts between the chipset on the robot and the frame grabber card, the Thermal-eye images are taken and stored on an off-board computer. While the Thermal-eye is mounted on the vehicle, the power and communications with it are run on a tether to a desktop computer. All the data is stored on the desktop.

The Firefly camera is plugged directly into the PCI fire wire card on the vehicle. All data and power for this camera is supplied by the vehicle itself and all data is stored on the onboard computer. The frame rate of the Firefly in closed loop

is 30 fps +/- .1 fps.

3.2 Physical Layout

Two lights, each with three 40 W light bulbs, are placed approximately 2 ft away from the vehicle. These lights are controlled by a dimmer switch which allows the full range of light settings from darkness to the maximum light output for the lamps. These lights are placed parallel to the wall, either in line with the vehicle for the rotation experiments, or 1 ft behind the vehicles travel path for the translation experiments. The texture used for the experiments is a thermal blanket draped over the walls of the environment. The coils in the blanket are set to increase the heat by 5°. The actual temperature difference is irrelevant, however, as long as there is enough thermal contrast between the coils in the blanket and the environment. The visual texture of the thermal blanket is sufficient to generate image shifts. Figures 3.1 and 3.2 show the physical layout for the rotational experiment.

3.3 Light Measurement

A Sekonic L-359 Light Master light meter is used to take all light value measurements. This light meter is typically used by photographers to establish how fast the shutter speed on a camera should be set at and/or how long and how bright a flash should be set. While the light meter has many different types of data, only the Exposure Value (EV) readings are used. The other output options are more camera specific. The conversion between EV and lumens or candle light per square foot is



Figure 3.1: Rotation Experiment Setup



Figure 3.2: Rotation Experiment Setup

given in equation 3.1. The low EV values for the various experiments (0-0.5 EV) are at about the light level of a dark parking lot at night. The high EV values (4-5 EV) corresponds to lighting high enough to see clearly, but not overly bright.

$$EV = \log_2\left(\frac{\text{lumens}}{10}\right) + 10 \quad (3.1)$$

3.4 Experiment Procedures

3.4.1 Pure Rotation

The pure rotation experiments have the vehicle turning at a constant 0.25 rad/sec turn rate for a set duration of time. The average distance from the robot to the textures being viewed is 6 ft +/- 8 in. The light measurements are taken at three different points along the arc of the circle, one at either end of the arc, and one at the center. The arc that the cameras on the vehicle view is 8 ft +/- 1ft, or about 1.3 radians.

3.4.2 Pure Translation

The translation experiments have the vehicle moving in a straight line parallel to a wall. The distance from the wall to the cameras is 3.5 ft +/- 3 in. The speed of the vehicle was 8 in/s +/- 1 in/s. Light measurements are taken at three different points along the wall: at the beginning, middle, and end.

3.5 Image Processing

3.5.1 Method of Processing Data

Data from the experiments is received as a series of images. Additional information, namely frame rates and vehicle ideal optic flow, are also supplied. To process the information, an appropriate sequence of images is chosen. This is done by throwing out the images that correspond either to when the vehicle not moving, or during the initial jerk of starting and stopping. The remaining images are then run through a processing program to determine the image shift between frames. Each pair of frames $(n, n + 1)$ is taken in by the program, then 40 lines of pixel shift are computed along horizontal lines. These lines are averaged vertically, giving a single line of average pixel shifts across the entire image. Next, the lines of optic flow are averaged to give a single average pixel shift for the entire image. Finally, the average pixel shift per image is averaged across the entire data set, giving the final global average of pixel shift for the experiment. Using the frame rate and the physical parameters from the experiment, that is the field of view of the camera, the pixel shift is converted to units of rad/s and compared to the ideal optic flow. The conversion from pixels in the image to field of view is made by dividing the number of horizontal pixels by the field of view measured from a static image. Equation 3.2 gives the conversion from the pixels per frame to radians per second.

$$Q_{dot} = \frac{\text{pixel}}{\text{frame}} \frac{\text{frame}}{\text{sec}} \frac{\text{FOV rad}}{\text{pixel}} \frac{\text{rad}}{\circ} \quad (3.2)$$

3.6 Methods of Determining Image Shift

3.6.1 Gradient Method

The gradient method for determining image shifts is merely solving equation 3.3 along a set line on the image. $\frac{dI}{dt}$ is evaluated as the estimated image shift between two frames in time. $\frac{dx}{dI}$ is the shift in intensities along the prescribed line. One problem with this method is that when only using data along the prescribed lines, the ability to track pixels, which get shifted off of the given line is lost. This gives a tendency to give false readings, particularly if there is any vibration in the image.

$$\frac{dx}{dt} = \frac{dx}{dI} \frac{dI}{dt} \quad (3.3)$$

3.6.2 Lucas-Kanade

The Lucas-Kanade method is built off of the gradient method [5]. One of the major improvements from the gradient method is the fact that this works in two dimensions, effectively using the entire image rather than just a prescribed line. This method uses a form of non-linear optimization to converge on an image shift. In order to do this, a low resolution version of the image is processed for an initial estimate on image shift. Once an estimate is found, a slightly higher resolution is used from the original image, with the estimated initial image shift used as the starting point for the optimization routine. This continues for a prescribed number of iterations. The result is a two dimensional vector field for the shift of each pixel

in the image. The benefits of this method over the gradient method are that all of the data from the image can be used, rather than just data along a line. However, the improved estimates come at the expense of increased computation time.

3.6.3 Data Reduction

Once the image shift per pixel on a line is been determined, the average pixel shift per line is computed. These lines are then averaged again, resulting in a global average across the picture. This is valid, since the experiments consist of either constant translation or rotation. There should be no local image changes such as expansion or contraction. Finally, the average shifts between images are averaged over the duration of the experiment, giving a global average shift across the experiment. Given the physical dimensions of the experiment and the velocity of the vehicle, the ideal optic flow is calculated and compared against the computed optic flow.

In addition to the average optic flow over time, the data is given temporally as well. At each data point the data is show as a histogram. These histograms are generated by taking the optic flow between each pair of images and binning it for each test. The Firefly camera generates a data point every 0.033 sec of test time (with images being taken at 30 Hz), and the Thermal-Eye camera every 0.091 sec (a frame rate of 11 Hz). The final value in the average plots is just the average of all the data shown in each histogram. Typically, a few of the data points in the high end of the histogram come from the initial start of the experiment, and a few of low

end points are influenced by the inertia of the system even though the motors have been told to stop.

Chapter 4

Experimental Results

4.1 Low Light Conditions

The low light experiments measure the amount of optical flow detected by the infrared and visible light cameras under various lighting conditions. These range from near darkness to standard indoor lighting. The results show a near constant estimate of optic flow from the Thermal-Eye camera over the entire range of testing. However, the Firefly loses the ability to track objects as the light levels drop.

4.1.1 Translation

The translational data (Figures 4.1, 4.5 and 4.3) shows that the optic flow detected by the Thermal-Eye camera is constant over the entire light range, while the Firefly camera's ability to detect shifts in images drops off. This is expected since the Firefly camera's ability to detect image shifts disappears as the images approach a uniform black image.

4.1.2 Rotations

The results from the rotation experiments are similar to the translational experiments. Notably, the quality of the visible light optic flow falls off as the environment gets darker, while the infrared optic flow quality is unaffected. In figures ??




Figure 4.1: Measured Optic Flow from the Low Light Translational Experiments

and 4.6 it can be seen that the data spread is less in some of the data points compared to others. This is because the data came from two different sets of tests. This variation comes from differences in the charge of the battery which runs the wheels and the low level control.

4.2 Obscured Line of Sight

To obscure the line of sight between the vehicle and its environment, a liquid based fog is pumped into the air. The lighting is the same in the clear environment




Figure 4.2: Temporal Data from the Low Light Infrared Translational Experiments

as in the bright end of the low light condition tests. To quantify the level of fog, a black and white image, figures 4.7 and 4.8, is posted within the image field for the Firefly camera. For each test, the intensity of the black in the image is compared to the intensity of the white in the image as show in equation 4.1. In order to get consistent results, the light and dark areas are the same. In the case where the black and white image is completely obscured, an estimate is made using the same




Figure 4.3: Temporal Data from the Low Light Visible Translational Experiments

method and location of where the image should be based upon the initial conditions of the test. To get the different levels of fogginess, fog is pumped into the test area. The tests are then run, as the fog slowly dissipates out of the test area. In order to keep the fog from dissipating too quickly, plastic sheets are used to enclose the test section.




Figure 4.4: Measured Optic Flow from the Low Light Rotational Experiments

$$Fog = \int (\text{Dark Area})dA - \int (\text{Light Area})dA \quad (4.1)$$

4.2.1 Translation

The data from the translational tests (figure 4.9) shows that significant optic flow is detected at all ranges of fog. This is due to the fact that the fog is not perfectly uniform and thus the shift of the image can be detected. This however raises a problem, since the Firefly camera is no longer giving information about




Figure 4.5: Temporal Data from the Low Light Infrared Rotational Experiments

the environment being traversed, but instead giving information about the local air velocity in the local vicinity. Figures 4.10 and 4.11 show the data arrayed temporally.

4.2.2 Rotation

Like the fog translation experiments, the rotation case (figure 4.12) also shows the same issues with optic flow being detected, but not relevant to the task of navigation and control. Like the translation experiment, the rotational data is also shown temporally in figures 4.13 and 4.14.




Figure 4.6: Temporal Data from the Low Light Visible Rotational Experiments





Figure 4.7: Calibration Images from Fog Translational Tests

0.74%




10.90%



12.89%




34.40%




37.49%



49.23%



81.60%



100.00%




Figure 4.8: Calibration Images from Fog Rotation Tests




Figure 4.9: Measured Optic Flow from the Fog Translational Experiments




Figure 4.10: Temporal Data from the Visible Light Translational Fog Experiments




Figure 4.11: Temporal Data from the Infrared Translational Fog Experiments




Figure 4.12: Measured Optic Flow from the Fog Rotational Experiments




Figure 4.13: Temporal Data from the Visible Light Rotational Fog Experiments




Figure 4.14: Temporal Data from the Infrared Rotational Fog Experiments

Chapter 5

Conclusions

The vehicle presented in this thesis meets a modified set of the initial parameters. This vehicle is capable of testing, among other things, the viability of infrared optic flow for navigation and control. Two sets of experiments document this; a set of low light experiments, and a set of reduced visibility experiments. The low light experiments show that as the light level drops the ability for a visible light sensor to compute meaningful optic flow decreases. This trend is the same in both the translational and rotational experiments. This is expected, since in total darkness, there would be no change in the image to generate optic flow. Unlike the visible light spectrum, infrared light is unaffected by a loss in visible light directly. The only way that this could change is if the light source produces enough infrared radiation, and there is little thermal contrast naturally in the environment.

The obscured vision tests show that low levels of fog does not seem to affect optic flow. However once the visibility drops to between 10 to 20% the optic flow drops off. The actual usefulness of this optic flow can be somewhat questionable, since the relative motion of the fog to the vehicle can change the measured optic flow. Like the low light experiments, the rotation and translation experiments show the same trends. Also, the infrared optic flow is unaffected but the fog, which again is expected.

There are several avenues of research based off of these experiments. One branch is to quantify the how obscured visibility effects the closed loop performance of the vehicle. Open loop tests show that there is significant optic flow under high levels of fog. It is, however, undetermined if this optic flow is useful for navigation and control. Closed loop control experiments are one of the ways to determine what level of vision obstruction causes control problems. Another research direction is to develop methods for using infrared and normal cameras in unison.

Appendix A

Data

EV	IR (rad/s)	VL (rad/s)	Ideal (rad/s)
0.17	0.118355	0.01636246	0.14
1	0.119446	0.00681769	0.14
1.5	0.139626	0.00954477	0.14
1.97	0.149444	0.03681554	0.14
2.63	0.119446	0.04772385	0.14
2.97	0.136899	0.08590292	0.14
3.73	0.148353	0.11181016	0.14
3.87	0.1309	0.11453723	0.14
4.33	0.112356	0.13089969	0.14

Figure A.1: Data from the Low Light Translational Experiments

EV	IR (rad/s)	VL (rad/s)	Ideal (rad/s)
0.4	0.127491	0.06217735	0.14
0.7	0.13874	0.05563237	0.14
1	0.110992	0.01636246	0.14
1.3	0.13874	0.02519819	0.14
2.1	0.110242	0.0212712	0.14
2.3	0.14249	0.05857761	0.14
3.1	0.13349	0.08672105	0.14
3.5	0.14549	0.07592182	0.14
4	0.104243	0.10471976	0.14
4.3	0.126741	0.05563237	0.14

Figure A.2: Data from the Low Light Rotational Experiments

Raw Measure	% Fog	IR (rad/s)	VL (rad/s)	Ideal (rad/s)
122	1.383848	0.026098	0.00981748	0.025
3666	41.58348	0.018749	0.02138028	0.025
1909	21.65381	0.017099	0.01963495	0.025
4477	50.78267	0.021148	0.02214386	0.025
2341	26.55399	0.026098	0.02312561	0.025
5203	59.0177	0.022348	0.02378011	0.025
5555	63.01044	0.024898	0.02487094	0.025
6427	72.90154	0.022198	0.02454369	0.025
8353	94.74819	0.025798	0.0263981	0.025
8816	100	0.030598	0.0263981	0.025

Figure A.3: Data from the Fog Translational Experiments

Fog Measure	% Fog	IR (rad/s)	VL (rad/s)	Ideal (rad/s)
122	0.746543	0.027298	0.004254	0.0275
1781	10.8983	0.032548	0.01849	0.0275
2107	12.89316	0.027148	0.022417	0.0275
5621	34.39603	0.025498	0.023071	0.0275
6146	37.60862	0.034048	0.026507	0.0275
8045	49.22898	0.022198	0.024707	0.0275
13335	81.59956	0.034648	0.028962	0.0275
16342	100	0.032848	0.029616	0.0275

Figure A.4: Data from the Fog Rotational Experiments

Bibliography

- [1] Conroy, J., Humbert, S., and Pines, D., *System Identification of a Rotary Wing Micro Air Vehicle* (Submitted to the Journal of the American Helicopter Society)
- [2] J. S. Humbert and R. M. Murray and M. H. Dickinson, *Sensorimotor Convergence in Visual Navigation and Flight* (Proceedings of the 16th IFAC World Congress, Praha, Czech Republic)
- [3] J. S. Humbert and R. M. Murray and M. H. Dickinson, *Pitch-altitude control and terrain following based on bio-inspired visuomotor convergence* (Proceedings of the AIAA Guidance, Navigation and Control Conference, San Francisco, CA)
- [4] J. S. Humbert, Andrew Hyslop, Michael Chinn, *Experimental Validation of Wide-Field Integration Methods for Autonomous Navigation* (International Conference on Intelligent Robots and Systems, San Diego, 2007)
- [5] B. Lucas and T. Kanade, *An iterative image registration technique with an application to stereo vision* (Proc. of the 7th International Conference of Artificial Intelligence)