

NEURAL NETS AND OPTIC FLOW FOR AUTONOMOUS MICRO-AIR-VEHICLE NAVIGATION

Paul Y. Oh* and William E. Green

Drexel Autonomous Systems Laboratory
Department of Mechanical Engineering and Mechanics
Drexel University
Philadelphia, Pennsylvania, 19104
Email: [paul.yu.oh, weg22]@drexel.edu

Geoffrey Barrows

Centeye
Washington, District of Columbia, 20009
geof@centeye.com

ABSTRACT

Reconnaissance, surveillance and target acquisition tasks in near-Earth environments like forests, caves, tunnels and buildings is a grand challenge. Micro-air-vehicles are a future line of bird-sized flying assets designed to address such a challenge. Needed are light-weight and miniature sensor suites that can provide autonomous collision avoidance in complex environments. Our demonstrations with optic flow microsensors have been promising but controller gain-tuning is often tedious. This paper describes the use of neural nets to automate gain tuning. The overall effect delivers collision avoidance over wide ranges of lighting conditions, contrast and surface textures.

INTRODUCTION

Micro-air-vehicles (MAVs) are a class of bird-sized aircraft envisioned to perform reconnaissance, surveillance and target acquisition tasks in forests, buildings, caves and tunnels. These near-Earth environments are often cluttered and dynamic where wireless communication is degraded, GPS reception is poor and illumination is varied. As such, these vehicles cannot be piloted remotely. The challenge therefore is designing sensor suites that enable autonomous flying. Oh and Green were the first to publish and successfully demonstrate both autonomous flying and landing inside buildings [6] [3] [4]. A 5-gram optic flow sen-

sor [1] was mounted on a 23-gram fixed-wing aerial testbed and interfaced to an embedded micro containing both reactive and proportional-derivative controllers. The net effect was that the testbed's flight maneuvers mimicked those of flying insects [8]. These past successes depended upon proper tuning of controller gains. This has been tedious and time-consuming because optic flow measurements are affected by lighting levels and surface texture in the area the aircraft is flown in. This paper discusses a neural net approach towards automated gain-tuning. The next section presents the theory underlying optic flow sensing. This is followed by sections describing the control system and neural net experiments. Finally, conclusions are given in the last section.

FLIGHT STRATEGEMS USING OPTIC FLOW

Insects make heavy use of vision, especially optic flow, for perceiving the environment [2]. Optic flow refers to the apparent movement of texture in the visual field relative to the insect's velocity. Insects perform a variety of tasks in complex environments by using their natural optic flow sensing capabilities. While in flight, for example, objects which are in close proximity to the insect have higher optic flow magnitudes. Thus, flying insects, such as fruit flies [9] and dragon flies, avoid imminent collisions by saccading (or turning) away from regions of high optic flow (see Figure 1). With optic flow sensors, efficient and robust navigational sensor suites for MAVs can be developed by mimicking the natural behaviors of insects.

*Address all correspondence to this author. This work was supported in part by the National Science Foundation CAREER award IIS 0347430




Figure 1. Dragon fly saccading away from regions of high optic flow in order to avoid a collision.




Figure 2. 1D optic flow during MAV flight.

Retrofitting Sensors on MAVs

Theoretically, optic flow is measured in rad/sec and is a function of the MAV's forward velocity, V , angular velocity, ω , distance D from an object, and the angle, θ , between the MAV's direction of travel and the object (see Figure 2).

$$OF = \frac{V}{D} \sin \theta - \omega \quad (1)$$

Figure 3 depicts optic flow as it might be seen by a MAV traveling a straight line above the ground. The focus of expansion (FOE) in the forward sensor view indicates the direction of travel.




Figure 3. Optic flow as seen by aerial robot flying above ground.

If the FOE is located inside a rapidly diverging region, then a collision is imminent. A rapidly expanding region to the right of the FOE (like the one seen in the Figure 3) corresponds to an obstacle approaching on the right side of the MAV. Thus, the MAV should turn left, or away from the region of high optic flow, to avoid the collision. Similarly, the MAV can estimate its height from the optic flow in the downward direction; faster optic flow indicates a low flight altitude. By equipping a MAV with sensors capable of measuring the optic flow in front of and below the aircraft, the above flight patterns can be embedded in a sensor suite for autonomous navigation.

Optic Flow Microsensors

Mixed-mode and mixed-signal VLSI techniques are often used to create compact circuits. Centeye has developed the one-dimensional *Ladybug* optic flow microsensors based on such techniques and is shown in Figure 4. These sensors are inspired by the general flow model of animal visual systems. A lens focuses an image of the environment onto a focal plane chip, which contains photoreceptor circuits and other circuits necessary to compute optic flow. Low level feature detectors respond to different spatial or temporal entities in the environment, such as edges, spots, or corners. The elementary motion detector (EMD) is the most basic structure or entity that senses visual motion, though its output may not be in a form easily used. Fusion circuitry fuses information from the EMDs to reduce errors, increase robustness, and produces a meaningful representation of the optic flow for specific applications.

Figure 5 depicts a simple realization of the feature tracker EMD algorithm used [1]. On the left is the basic EMD architecture, on the upper right is an edge detection kernel implemented by a differential amplifier, and on the lower right are sample traces of feature signals and feature location signals. This EMD measures one-dimensional optic flow in one part of the visual field, thus the complete sensor would have many such EMDs




Figure 4. mixed-mode VLSI optic flow microsensor is slightly bigger than a US quarter.

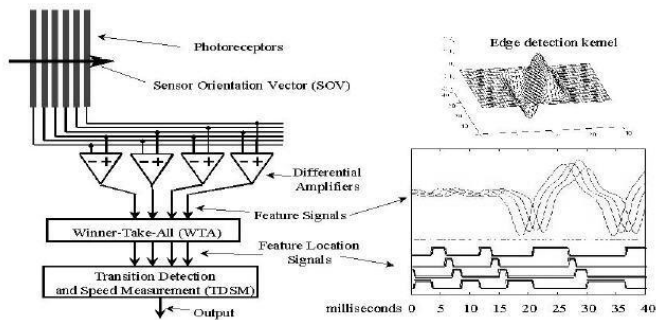


Figure 5. The feature tracker elementary motion detector (EMD).

replicated throughout the visual field. Functionally there are four sections, as shown in Figure 5: photoreceptors, feature detectors (shown here as differential amplifiers), a winner-take-all (WTA), and a transition detection and speed measurement section (TDSM).

A section of the focal plane is sampled with an array of elongated rectangular photoreceptors laid out so that the array is positioned along the sensor orientation vector (SOV). The photoreceptor rectangles are arranged so that their long axes are perpendicular to the SOV. This layout filters out visual information perpendicular to the SOV direction while retaining information in the parallel direction. One effect of these rectangular photoreceptors is that the measurement of the sensor is actually a measurement of the projection of the two dimensional optic flow vector onto the SOV vector [1].

The outputs from the photoreceptors are sent to an array of four feature detectors that output four analog feature signals. A feature detector circuit attains its highest output value when the feature to which it is tuned appears on its input photoreceptors. For example, suppose the feature detectors are differential ampli-

fiers. Then their effective response function is the edge detection kernel, shown in the upper right part of Figure 5. A feature signal will have a high value when an edge is located between the input photoreceptors with the brighter side on the positively connected photoreceptor.

The four analog feature signals are then sent to a winner-take-all (WTA). The WTA has four analog inputs and four digital outputs. The WTA determines which input has the highest value and sets the corresponding output a digital high (or 1), and all the other outputs low (or 0). The location of the high value indicates where on the photoreceptor array the image is most like the feature defined by the configuration vector. The WTA outputs are thus also called feature location signals. As an edge moves across the photoreceptors shown in Figure 5, the high value will move sequentially across the WTA outputs. This is easily visualized with the aid of the trace in the lower right part of Figure 5. Shown are four feature signals and their corresponding feature location signals when the photoreceptors are exposed to a moving black and white bar pattern.

The transition detection and speed measurement (TDSM) circuit converts the movement of the high WTA output into a velocity measurement. Essentially this circuit is a state machine that responds to the WTA outputs and interprets the 1-2-3-4 motions of the high feature location signal as visual motion. Whenever the high feature location signal moves in a manner that indicates visual motion, the EMD generates a measurement of the optic flow. The direction of the visual motion is determined by the direction of travel of the high feature location signal. Likewise the speed is obtained with the lag-time from one feature location to the next. This lag-time is also referred to as the transition interval. Then the actual optic flow can be determined from the physical geometry of the photoreceptor array and the sensor optics.

The resulting sensor, including optics, imaging, processing, and I/O weighs 4.8 grams. This sensor grabs frames up to 1.4 kHz, measures optic flow up to 20 rad/s (4 bit output), and functions even when texture contrast is just several percent (see Figure 6).

AUTONOMOUS FLIGHT MANEUVERS

Optic flow microsensors can be oriented to perceive information about oncoming collisions and altitude. For example, positioning sensors such that the optical axis faces in the forward direction will allow the measurement of the optic flow field in front of the aircraft. Likewise, measuring the optic flow on the ground requires placing a sensor on the belly of the MAV. Such information can be used to mimic insect flight patterns to perform autonomous collision avoidance and landings for MAVs.




Figure 6. Contrast variation prompts full rudder deflection as a result of bang-bang control.




Figure 7. Optic flow control system block diagram.

Autonomous Landing

Oh and Green were the first researchers to demonstrate autonomous landing of a fixed-wing aerial robot inside a building [4]. The approach involves keeping the optic flow on the landing surface constant. When measuring the optic flow on the landing surface, the obstacle is now the ground and thus $\theta = 90$. To further simplify this task, the rotational component of optic flow arising from changes in aircraft pitch are assumed smaller than the translational component. Thus, Equation (1) reduces to

$$OF = \frac{V}{D} \quad (2)$$

Keeping Equation (2) constant (where D is the altitude) demands the aircraft's control system decrease forward speed in proportion to altitude.

The control system block diagram and flow chart are shown in Figures 7 and 8 respectively. When approaching a landing, an embedded microcontroller (see Figure 9), will implement a function to gradually throttle down the motor while continuing to take readings throughout the landing process. The error, $e(t)$, is computed between the desired optic flow, o_i which was estimated beforehand, and the actual optic flow value, $o_f(t)$. When the optic flow on the landing surface becomes larger than the desired optic flow, the error is negative and two conditions are possible. One,

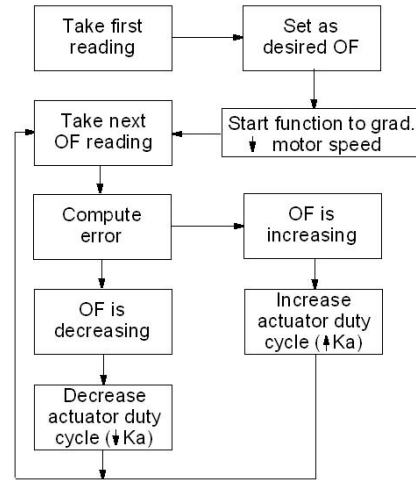


Figure 8. Flow chart of landing control system.

the forward velocity, V , could be significantly increasing which is not possible based on the motor function. Two, the altitude, D , can be decreasing at a faster rate than V . Here, the controller will send a signal to the elevator to decrease the vehicle's descent rate based on the error magnitude and proportional constant, K_a . The other possibility is that the optic flow could start to dip below the desired level causing the error to be positive. The two possible cases that arise here are one, D is increasing but again this is not practical while in landing mode and two, V is decreasing faster than D . In this case, the controller will need to command the elevator to increase the descent rate. After a control sequence has been implemented to force the optic flow back to the desired value, the elevator resets to its neutral position. By implementing this control scheme, we were able to successfully demonstrate an autonomous landing (see Figure 10).

Autonomous Collision Avoidance

Autonomous collision avoidance while flying fixed-wing aircraft inside buildings was first successfully demonstrated by Oh and Green [6]. The general approach is to command the MAV to turn away from regions of high optic flow. Optic flow must be detected in front of the vehicle in order to avoid collisions and thus, the sensor must be positioned at some angle forward. Unlike with autonomous landing, where the sensor was oriented at 90 degrees to the direction of travel, the angle θ to the obstacle will be a factor. Assuming the MAV is traveling in a straight path with a relatively constant translational velocity, V , we have from Equation 1

$$OF = \frac{V}{D} \sin \theta \quad (3)$$



Figure 10. The optic flow on the basketball gym floor is kept constant by the control system. That is, the aircraft (encircled) forward velocity is decreased in proportion with its altitude to land smoothly. Left: Aircraft just after hand launch. Middle: Aircraft midway through landing sequence at proportionally lower altitude and velocity. Right: Aircraft comes to a smooth landing within 25 meters from starting point.



Figure 11. Optic flow is used to sense when an obstacle is within two turning radii of the aircraft. The aircraft avoids the collision by fully deflecting the rudder.

An optic flow threshold was set to correspond to an obstacle being within two turning radii of the aircraft. Thus, when the threshold is exceeded, the sensor suite will implement proportional rudder control to safely avoid the obstacle. By implementing this method, autonomous collision avoidance was successfully demonstrated (see Figure 11).

NEURAL NETWORKS

Artificial neural networks are modeled after biological nervous systems, such as the brain, and represent a methodology to process and interpret raw information. Like neurons in the human brain, neural networks consist of many interconnected nodes which function collectively to communicate and disseminate information to solve specific problems. The ability to perform distributed computation, tolerate noisy inputs, and learn and adapt to unseen conditions is what makes neural networks so attractive. Such attributes make applying neural networks a promising approach to the characterization of optic flow microsensors.

There are several parameters, if varied, which could affect

the overall performance of optic flow sensors. The three most significant are light intensity, contrast and texture. The difference in terms of light intensity (measured in lux) for natural sunlight and artificial lighting can be as high as two orders of magnitude. The sensor output for an object in identical motion in both environments could yield extremely different results. Furthermore, optic flow readings are almost non-existent in poor lighting conditions (e.g. at dusk or in a shadow). Similarly, objects which are dull or low in contrast (e.g. a white wall) will yield very low optic flow magnitudes even when within close proximity of the sensor. The net result is that these realistic conditions which yield contradicting sensor outputs could be fatal to a MAV if not accommodated for.

Neural networks can be taught to deal with this type of data. Networks, like people, learn by example. A network must be trained for it to successfully implement a desired task. To be able to adapt to different lighting conditions as well as objects with different textures and contrast, the network must be presented with actual data that represent a specific state of the MAV's world. For example, one state for an approaching object (e.g. boulder) might include: *high* light intensity, *rich* object tex-




Figure 9. A microcontroller is used to read the digital output of the optic flow sensor and implement the control algorithms. The control signal is then sent through an H-bridge to deflect the aircraft's control surfaces.

ture, and *high* contrast, while another state could consist of: *low* light intensity, *rich* object texture, and *high* contrast. While the possibilities of different scenarios seem endless, a network need only be trained so that it remains generalized; overtraining (i.e. presenting the network with every possible scenario) can lead to performance degradation.

We conducted an experiment to see how effective a neural network would be in interpreting data from the optic flow sensor at different light intensities. Our network and experimental setup are shown in Figures 12 and 13, respectively. It is a multilayer feed-forward network consisting of two hidden layers. The network was trained and validated with two inputs: (1) a one-dimensional optic flow sensor acquired readings of a model railcar (in rad/sec) as it passed by the sensor at a constant reproducible linear velocity and (2) a digital light sensor measured the intensity of the ambient lighting conditions (in lux). The network output was trained using the actual distance from the sensor to the railcar (in inches). Different scenarios were achieved by varying the florescent lighting conditions from 0-500 lux (a bright office is apx. 400 lux) as well as the actual distance from 0 to 63 inches.

Backpropagation updating was used during the training and validation phase. The activation function of each non-input node in the network is a sigmoid [7]

$$\sigma = \frac{1}{1 + e^{-x}} \quad (4)$$

where x is the nodal input. Once the network was trained and validated, it was presented with an unseen state (i.e. 300 lux and a distance of 18 inches) to test its efficiency. In 900 data




Figure 12. A neural network with two hidden layers was created using JavaNNS v1.1. The network was used to characterize the output of an optic flow sensor in terms of distance to the obstacle.




Figure 13. Experimental setup used to collect training and validation data for neural network.

points, the average error of the network output was 3.5 inches. There were two outliers which accounted for the maximum error of 16 inches; the remaining errors fell in the error range of 0-5 inches. A summary of the results can be found in Figure 14. Two hidden layers are sufficient in this experiment, but more nodes and layers will be required when trying to characterize sensor outputs for more than two varying parameters.

Training and Validation Data	
Light Intensity (lux)	Actual Distance (in.)
15	7
15	63
250	7
250	31
500	7
500	63
Testing Data	
Light Intensity (lux)	Actual Distance (in.)
300	18
Network Output Error (900 trials)	
Average Error (in.)	Maximum Error (in.)
3.5175	16.0000

Figure 14. Results of applying a neural network to optic flow sensors in different lighting levels.

CONCLUSIONS

Near-Earth environments which occlude conventional navigational methods, such as GPS satellites and the horizon, are time consuming and labor intensive to patrol and safekeep. Lightweight optic flow microsensors, based on the vision system of flying insects, are suitable for micro-air-vehicle payload capacities. This paper presented details on leveraging sensors for navigation. The underlying control laws for collision avoidance and automated landings were also detailed. A neural net to automate controller gain-tuning was formulated. The results were promising, suggesting a viable method to bypass tedious and time-consuming calibration procedures.

REFERENCES

- [1] Barrows, G., "Mixed-Mode VLSI Optic Flow Sensors for Micro Air Vehicles", *Ph.D. Dissertation*, University of Maryland, College Park, MD, Dec. 1999.
- [2] Gibson, J.J., *The Ecological Approach to Visual Perception*, Houghton Mifflin, 1950.
- [3] Green, W.E., Oh, P.Y., "An Aerial Robot Prototype for Situational Awareness in Closed Quarters", *IEEE/RSJ Int Conf on Intelligent Robots and Systems*, pp. 61-66, Las Vegas, NV, Oct. 2003.
- [4] Green, W.E., Oh, P.Y., Barrows, G., Sevcik, K., "Autonomous Landing for Indoor Flying Robots Using Optic

- Flow", *ASME Int. Mech. Eng. Congress and Expo.*, v2, pp. 1341-1346, Wash., D.C., Nov. 2003.
- [5] Netter, T., Francheschini, N., "A Robotic Aircraft that Follows Terrain Using a Neuromorphic Eye", *IEEE/RSJ Int Conf on Intelligent Robots and Systems*, V1, pp. 129-134, Lausanne, Switz., Sept. 2002.
- [6] Oh, P.Y., Green, W.E., "Closed Quarter Aerial Robot Prototype to Fly In and Around Buildings", *Int. Conference on Computer, Communication and Control Technologies*, Vol. 5, pp. 302-307, Orlando, FL, July 2003.
- [7] Russell, A., Norwig, R., *Advanced Artificial Intelligence*, McGraw-Hill, 1999.
- [8] Srinivasan, M.V., Chahl, J.S., Weber, K., Venkatesh, S., Nagle, M.G., Zhang, S.W., *Robot Navigation Inspired By Principles of Insect Vision in Field and Service Robotics*, A. Zelinsky (ed), Springer Verlag Berlin, NY 12-16.
- [9] Tammero, L.F., Dickinson, M.H., "The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*", *Journal of Experimental Biology*, v205, pp. 327-343, 2002.