Development of a Hemispherical Compound Eye for Egomotion Estimation

Will Maddern and Gordon Wyeth
School of Information Technology and Electrical Engineering
University of Queensland
{w.maddern, g.wyeth}@uq.edu.au,

Abstract

Biological inspiration has produced some successful solutions for estimation of self motion from visual information. In this paper we present the construction of a unique new camera, inspired by the compound eye of insects. The hemispherical nature of the compound eye has some intrinsically valuable properties in producing optical flow fields that are suitable for egomotion estimation in six degrees of freedom. The camera that we present has the added advantage of being lightweight and low cost, making it suitable for a range of mobile robot applications. We present some initial results that show the effectiveness of our egomotion estimation algorithm and the image capture capability of the hemispherical camera.

1 Introduction

The estimation of self motion (egomotion) using visual information alone has been a major area of research over the past decades. Many computational techniques have been developed that allow apparent image motion in a plane to be calculated, as well as techniques to derive the motion of the observer given such a pattern of image motion. These techniques are mostly implemented using powerful conventional computers, to allow accurate real-time estimation of egomotion using high resolution sensors. Egomotion from image motion is highly useful in the field of simultaneous localisation and mapping [1], since it allows the same hardware to be used for both place recognition and motion determination.

However, recent research in the field of miniature aerial vehicles has created a need for a simple and compact method of egomotion estimation. A new trend in the development of navigation and obstacle avoidance systems for MAVs has been to analyse how insects perform this function. Despite having low-resolution fixed focus vision, insects are capable of regulating their flight speed and position with respect to objects [2], estimating the correct time to deploy their legs before landing [3], and calculating the distance flown using visual data alone [4]. Research into the optomotor response of a fly's eye has developed a simple mathematical model that calculates the apparent image motion across two adjacent photosensors [5].

Implementations of this biologically-inspired correlation system have shown successful results. By attempting to minimise 1D image motion on a linear plane, these sensors have been used to validate the theory behind the optomotor response by comparing the response of the sensor against a real fly [6], perform altitude control and obstacle avoidance on MAVs [7], and correct for actuator mismatch in small mobile robots [6]. However, thus far these sensors have not been used to compute egomotion in more than one dimension.

The calculation of egomotion using image motion can be vastly improved by a wide-angle or omnidirectional image sensor. A larger field of view allows for more accuracy in determining total motion, since translational and rotational components are more easily separable [8]. Research into the compound eye has shown that due to its hemispherical nature, it can provide insects with an almost completely omnidirectional image [9].

This paper describes the design and construction of a wide-angle optical flow sensor based upon biologically-inspired image motion detection. The sensor is of a hemispherical shape to model a fly's eye, and is constructed using common electronic components. Phototransistors detect image intensity and are arranged in clusters to allow for modular construction and simple circuit debugging. The image can be transmitted over USB to a PC, or I^2C to a microcontroller.

The next section details the previous research done to develop methods to calculate optical flow in conventional and biologically-inspired systems, as well as methods to extract egomotion given an omnidirectional optical flow field. The investigation of the quality of visual odometry data provided by a wide-angle optical sensor based on biological principles is then demonstrated, and the details of the design and construction of the fly-eye sensor explained. The paper concludes with some preliminary results that illustrate the sensor's functionality.

2 Background

Models of egomotion estimation in insect vision are typically framed around the concept of the elementary motion detector. In this section we review the use of the elementary motion detector for sensing optical flow, and link from optical flow to egomotion estimation. In particular we illustrate the benefit of estimating self motion from a suitably arranged hemispherical sensor.

2.1 Elementary Motion Detector

Proposed in [5] and subsequently named the Hassenstein-Reichardt correlation detector (or simply Reichardt detector), this method of elementary motion detection finds the correlation between two adjacent photodetectors. Significant research into the electrophysical response of a fly's eye has shown that it can be adequately characterised using this model [10]. Figure 2 shows a simplified model of the EMD. The intensity measured by each photodetector is correlated with a phase-delayed signal from the opposing photodetector. The greatest correlation is found when the spatial intensity delay between the photodetectors is equal to the time delay constant τ. Finding the difference between the two correlations yields a direction-sensitive representation of the image motion. A planar field of EMDs can then create an image motion field similar to that calculated in Section 2.2.1, and thus can be used to calculate egomotion.

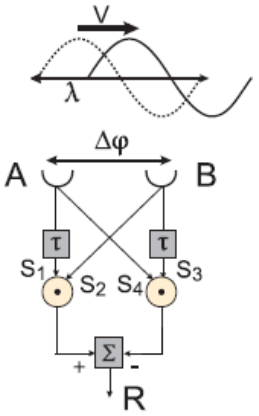

Figure 1 – Hassenstein-Reichardt correlator. A and B are photodectors, spaced a distance Δφ apart. Boxes with a τ correspond to a pure time-delay, and circes with a dot represent non-linear multiplication of two signals. The output R is equal to the difference between the two correlated signals. Reproduced from [11].

Extracting the image velocity from such a motion sensor is a non-trivial task. In any given environment, the intensity at the photodetector $I(t)$ is a function of both spatial and temporal frequency. Without prior knowledge of either, it is not possible to separate these values using the above detector [12]. Also, due to the periodic nature of the correlation with a fixed time delay, temporal aliasing will occur at image motion detection above the maximum correlation [11]. This can be corrected by either using multiple parallel detectors with different τ values [13], or ensuring the maximum sensing velocity lies below the aliasing level.

The problem of separating spatial and temporal velocities can be addressed by changing the balance of the EMD [11]. By considering the response of a single EMD, with detector separation of $\Delta\varphi$, to a sinusoidal intensity with period λ moving at an angular frequency V as shown in Figure 2, the signals S_1 and S_2 are as following:

$$S_1 = \sin\left[\left(\frac{V}{\lambda} - \tau\right)t\right], \quad S_2 = \sin\left[\frac{V}{\lambda}t - \frac{\Delta\varphi}{\lambda}\right]$$

Given that S_3 and S_4 can be expressed with similar equations, the EMD output R can be calculated. The following equations show the result for the half-EMD S_1S_2 and balanced EMD $R = S_1S_2 - S_3S_4$:

$$S_1S_2 = \cos\left[\frac{V}{\lambda}\left(\frac{\Delta\varphi}{V} - \tau\right)\right], \quad R = \sin\left(\frac{\Delta\varphi}{\lambda}\right)\cdot\sin\left(\frac{V\tau}{\lambda}\right)$$

While it can be seen that R is dependent both on V and λ but provides directionally sensitive output (given $\lambda > 0$ and $-\pi < V\tau/\lambda < \pi$, the second sinusoidal term varies between -1 and 1 with the same sign as V), the result of S_1S_2 is primarily velocity sensitive (since the maximum response is at a velocity V of $\Delta\varphi/\tau$ irrespective of λ). Figure 3 shows contour plots of the response of both the half-EMD and balanced EMD.

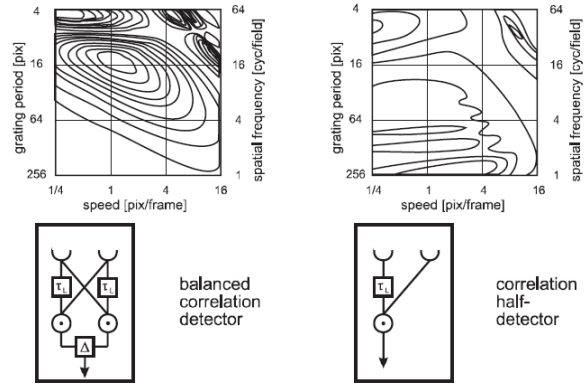

Figure 2 – Contour plots of the response to full- and half-balanced EMDs. The balanced detector exhibits a maximum response dependant on both speed and grating period, whereas the half-detector maximum is less dependant on spatial period but not direction-sensitive. Reproduced from [11].

It can be seen that in the balanced detector, the output is significantly dependant on both the speed and the grating period, and as either is increased beyond the maximum detector sensitivity, spatiotemporal aliasing is produced. Although an overall positive (direction insensitive) response is shown for the half-detector, it is not as strongly dependant on grating period, since the maximum response is independent of grating period. However, this method is still subject to spatiotemporal aliasing, and thus is insufficient for direct calculation of angular velocity. Further adjustment to the balance of the EMD can be done by introducing a balance constant, α, to more closely model an insect's motion detector system [11].

In order to unambiguously calculate velocity, multiple correlations using different time constants may be used. Figure 4 illustrates the intensity of a single balanced EMD using three different values for τ for varying spatial and temporal frequency.

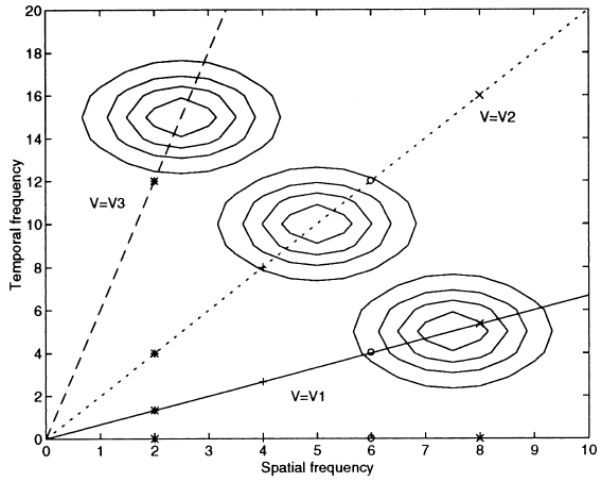

Figure 3 – Contour plot of the response of multiple full-balanced EMDs. Movement at a constant velocity will give straight lines as labelled on the plot, which can be unambiguously estimated using two or more time constants for EMD correlators. Reproduced from [14].

This figure illustrates that for different spatiotemporal frequencies, a given velocity *V* (represented by the straight lines on the graph) produces a unique response in terms of all three EMDs [14], even though the response of each individual EMD is periodic and subject to aliasing (not shown on the contour plot). This method yields a directional velocity-tuned response, but doubles or triples the number of correlations required for each EMD [13].

2.2 Motion Estimation from Optical Flow

Given an arbitrary 2D optical flow vector field, the average image motion in a translational direction can be found simply by summing every vector in the field, and dividing the magnitude of the resultant vector by the number of vectors in the field [1]. Likewise, the average rotation can be found by the sum of multiplying each vector by its perpendicular distance to the centre of rotation [1]. However, given a hemispherical image surface, the image motion vector field is more difficult to calculate, and may be better represented by polar or spherical co-ordinates, rather than Cartesian.

Due to recent research in the areas of wide angle and omnidirectional computer vision, a number of strategies to extract egomotion from a non-planar velocity field have been developed. Figure 5 illustrates the advantages of using a wide angle spherical optical flow field to estimate egomotion in comparison to a 2D planar image sensor.

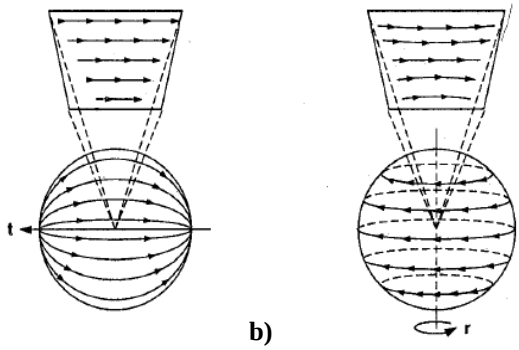

Figure 4 – Translational and rotational motion viewed from a planar and spherical projection. a) corresponds to uniform translation t and b) to uniform rotation r. It can be seen that the 2D planar projections are very similar, however the spherical projections are markedly different. Reproduced from [8].

It can be seen that while translational and rotational egomotion appear very similar using a 2D image projection, it is easy to differentiate between the two with an omnidirectional viewing angle [8]. For a spherical view, given a translational vector *T* and a rotational vector Ω, the motion of any arbitrary point *P* projected on a unit circle (yielding $\hat{P}$) is given as follows:

$$U(\hat{P}) = \frac{1}{\|P\|}\left((T \cdot P)\hat{P} - T\right) - \Omega \times \hat{P}$$

$U(\hat{P})$ gives the instantaneous velocity at any point $\hat{P}$ on the spherical image surface, which would be provided by an optical flow algorithm such as the system described in Section 2.1. Therefore, knowing $U(\hat{P})$, a method to determine *T* and Ω can be found by deriving a depth independent constraint (the epipolar constraint [15]), then setting the derivative to 0 to find the minimum. This yields the following equation:

$$T \cdot \left(\hat{P} \times \left(U(\hat{P}) + \left(\Omega \times \hat{P}\right)\right)\right) = 0$$

From this, a least-square estimate of Ω as a function of *T* can be found, which can then be re-substituted into the above equation to solve for *T*, yielding the translational and rotational egomotion [15].

3 Design of the Fly Eye Sensor

The preceding sections illustrate the benefit of using a hemispherical array of elementary motion detectors as the basis for estimating egomotion, and offer some explanation for the formation of the insect eye. This section describes our implementation of a hemispherical camera suitable for use as an array of EMDs.

3.1 Requirements

To model the arrangement of elementary motion detectors in an insect's compound eye three main properties are required; a hemispherical low-resolution field of view, a high bandwidth and a wide spectral sensitivity. Conventional wide-angle imaging sensors make use of a planar array of photodetectors (using CCD or CMOS technology) coupled to a lens or mirror. While this allows for a large pixel density, it introduces significant spatial

distortion and requires high bandwidth processing of the image. Additionally, most low-cost image sensors are limited to a maximum frame rate of 30 fps, whereas for high velocity image motion it is desired to sample at rates above 200fps. Therefore, it was decided to construct a hemispherical sensor using narrow-angle, wide spectral sensitivity photodetectors to provide the required capabilities.

3.2 One Dimensional Prototype

A planar arc of sensors was used a one dimensional prototype in order to perform some initial investigation of suitable devices and parameters for the hemispherical array. The one dimensional prototype consists of 8 phototransistors separated by 15° along an arc of radius 30 mm. The prototype is illustrated below (see Figure 6) and forms the basis for the preliminary results illustrated in the results section.

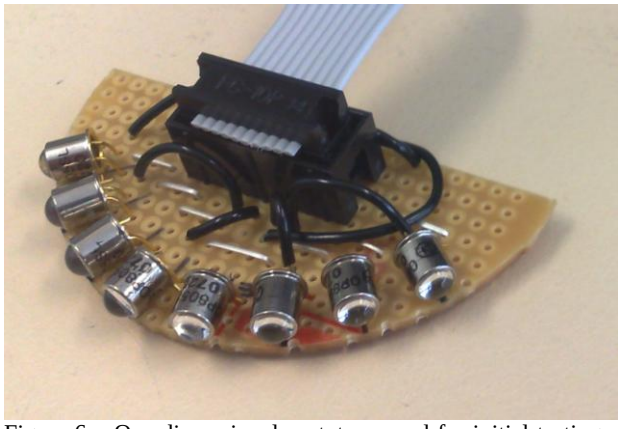

Figure 6 – One dimensional prototype used for initial testing of egomotion estimation and component selection.

3.3 Physical Construction of the Hemisphere

To provide a hemispherical structure upon which the photodetectors could be uniformly positioned, a section of a truncated icosohedron was chosen (see Figure 7). By packing the faces of the polyhedron with photodetectors, a maximum angular separation of 15° could be achieved, equal to that used in the pilot implementation. A total of 116 sensors were used in this configuration.

Figure 7 – The completed hemispherical sensor. The phototransistors are arranged on hexagonal and pentagonal PCBs and aligned by precision machined Perspex covers.

3.3.1 Construction

Given that the sensor is configured as a number of pentagons and hexagons connected to form a three-dimensional polyhedron, it was decided to base the circuit board construction on these shapes. A total of 6 pentagonal, 10 hexagonal and 5 quadrilateral (at the base of the hemisphere) printed circuit boards were constructed (see Figure 8). Each board has a side length of 12mm and every edge contains a 5 pin connector. These connectors allow the boards to be mechanically and electrically connected to each adjacent board, forming a robust structure.

In order to align the sensors radially from the centre, a number of Perspex polygons with sensor holes were constructed using a CNC machine (see Figure 7). Alignment to the centre of the hemisphere was achieved by bending the phototransistor pins to conform to the precision milled holes in the Perspex polygons. The polygons are held in place using an interference fit over the body of the phototransistors, which allows them to be removed if required.

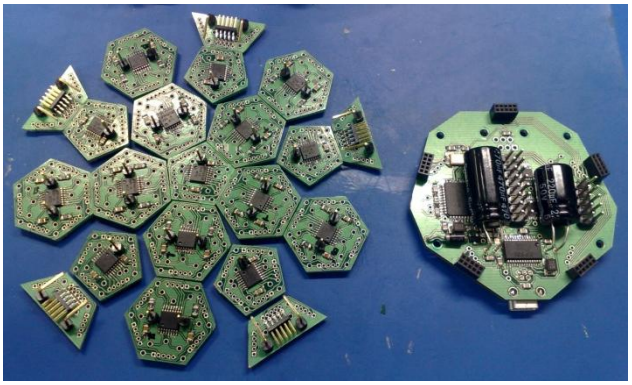

Figure 8 – An exploded view of the fly-eye sensor. This view shows the underside of the PCBs with the surface mount analog multiplexers. Boards are interconnected by the holes in the edges, and connect by jumper connections to microcontroller on the motherboard (shown on the right).

3.4 Electronic Implementation

The electronics of the eye were designed to provide rapid, synchronous acquisition of illumination values from the sensor array. Figure 9 shows a block diagram of the electronic implementation of the sensor.

3.4.1 Sensor Selection

The photodetector chosen for this application was the Optek OP802SL phototransistor, available for less than $2 in quantity. This sensor provides a wide spectral response of 400 to 1100nm (and is therefore sensitive to both fluorescent and incandescent light sources) as well as a 12 degree half-angle of sensitivity, giving a very narrow field of view. Given a 15 degree angular separation, this provides a small amount of overlap between sensors, which is desirable for image interpolation. The OP802SL in particular has a saturation current of less than 3mA, allowing for low power consumption when used in large numbers.

3.4.2 Sensor Circuitry

As the photodetector in use is a phototransistor, there is no need for any external drive circuitry or amplification. The phototransistors are wired in a common-collector configuration, with a 150 kΩ resistor between the emitter and ground. This value was determined experimentally using typical light conditions found in indoor environments, and allows a maximum current of 30μA to pass through the sensor. A 1nF ceramic capacitor is connected in parallel with each resistor to form a first-order lowpass filter with a cutoff frequency of 100Hz, to provide limited attenuation of flicker from fluorescent light sources without reducing temporal response.

Each sensor is connected via a series of analog multiplexers to the ADC port on an Atmel AVR microcontroller. The analog signals are sampled at 10-bit resolution, and the microcontroller allows for a maximum frame sample rate of 5 kHz. The microcontroller chosen was the ATMega164P clocked at 16MHz using an external crystal. At full load, the entire sensor uses under 20mA at 5V.

Communication to the PC is performed using a USB to UART converter, allowing for data rates up to 3Mbps. The connection appears as a virtual serial port on the PC and can be accessed using most terminals or programming languages. Additionally, the I^2C port of the microcontroller is brought out to a RJ-11 connector, to allow the sensor to communicate with other microcontrollers. The wiring of the connector allows a direct connection to the Lego NXT brick sensor inputs, providing a ready to use test bed for experimenting with a mobile platform.

4 Testing and Results

Testing of self motion estimation has been performed using the one dimensional prototype described in the previous section. The results of those tests are described below. The full hemispherical sensor has been implemented and is shown in this section to be able to capture wide angle images of the environment, which are compared to images captured using a conventional camera and a panoramic mirror.

4.1 One Dimensional Sensor

The one dimensional sensor was used to perform some initial experiments to gauge the effectiveness of the phototransistor array in accurately capturing self motion information. These preliminary results provide valuable insight into the performance of the hemispherical sensor.

4.1.1 Experimental Setup

In order to test the performance of the one-dimensional sensor, well-controlled optical conditions are required. It was decided to use grating patterns to provide an optical flow field, and evaluate the estimated angular velocity against a measured velocity provided by an encoder.

The sensor was placed at the centre of a circular table with a diameter of 75 cm, ringed by a patterned wall. The testing pattern was printed onto 2.1 × 0.3 m sheets, to fill a 300° field on the surface of the table. The field of view of the sensors from the centre of the table to the edge is only 15 cm high, well short of the height of the pattern, and therefore the entire field of view of the one-dimensional sensor is filled by the pattern. The intensity pattern illustrated here is a sinusoid with wavelength that is double the spatial wavelength of the sensors. The table was uniformly lit using standard fluorescent lamps in a laboratory environment.

For this experiment the sensor was driven at a known rate using a Lego NXT motor. The readings from the encoder attached to the motor were compared to the motion estimate from the sensor. Sensor readings were sampled at 25 Hz using a 10 bit ADC on an AVR microcontroller. The encoder readings are based a 1080 counts per revolution encoder. The motor was actuated with a commanded velocity of 71, 143 and 263 deg/s.

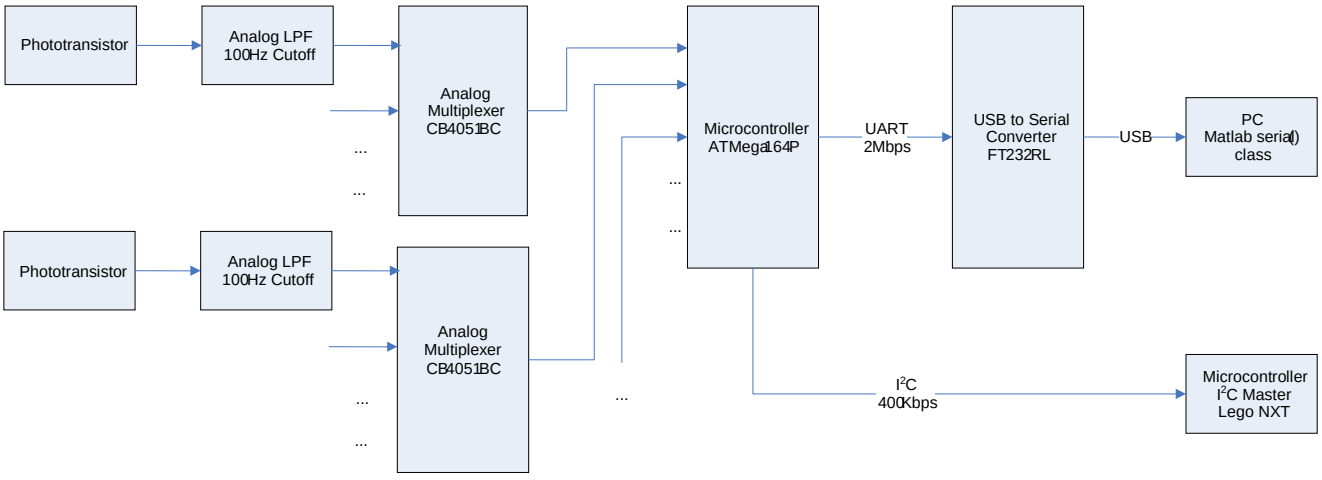

Figure 9 – Block diagram showing the electronic implementation of the hemispherical sensor. The electronic design is chiefly constrained by the challenge of streaming 116 sensor values into a single serial port.

Figure 10 – The testing arena for the one-dimensional sensor array. The sinusoidal grating provides an exemplary environment for assessing the egomotion measurement characteristics of the sensor.

4.1.2 Experimental Results

The experimental results show the measured angular velocity from both sensors, and the total angle moved calculated from the integral of the velocity. The results are illustrated in Figures 11, 12 and 13, which show the curves for commanded speeds of 71, 143 and 263 deg/s respectively.

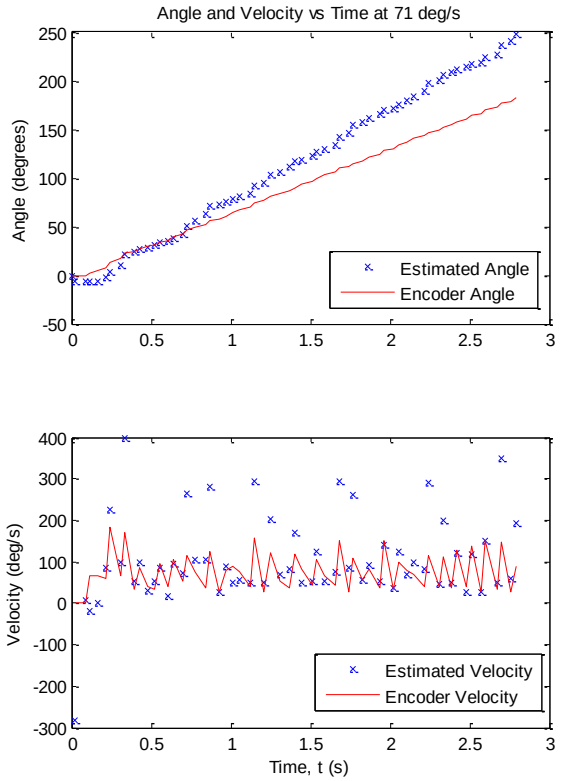

Figure 11 – Comparison of estimated egomotion at 71 deg/s.

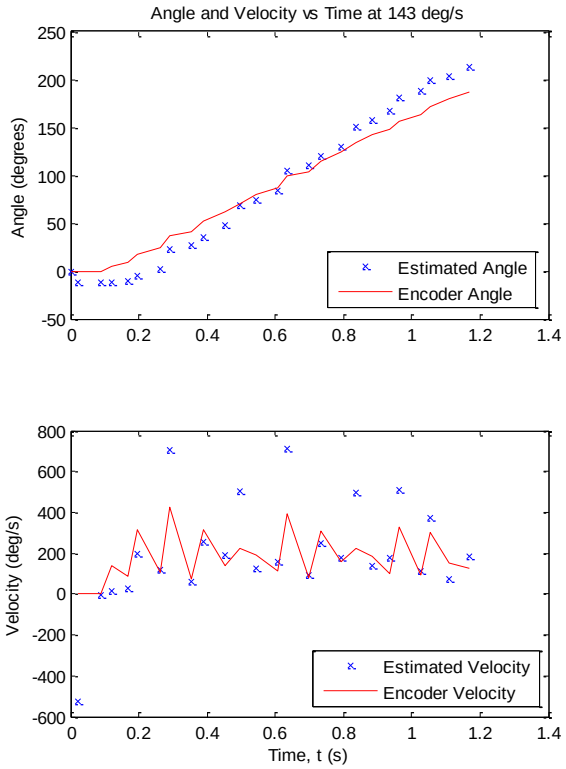

Figure 12 – Comparison of estimated egomotion at 143 deg/s.

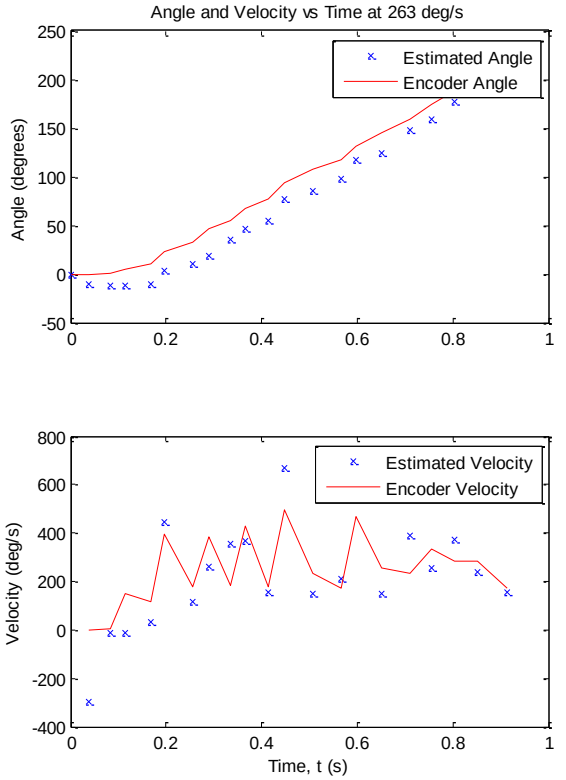

Figure 13 – Comparison of estimated egomotion at 263 deg/s.

Each of the results show that the optical sensor

approximates the encoder sensor well, if with additional noise, in the velocity measurement. The integrated angle of each speed shows a strong correspondence between the angle measure from the encoder and the angle measured from the sensor. At all speeds, there is a notable offset in the early movement, and an evident velocity bias as the movement progresses. It is interesting that the offset and the bias are consistent over time rather than angle, with the intersection of the trajectories always occurring at about 0.6 seconds. The bias and offset are caused by angular misalignment in the construction of the prototype, leading to the decision to have strong alignment on the complete sensor by using the Perspex covers.

4.2 Hemispherical Sensor

At the time of writing the hemispherical sensor had only recently been completed, and the egomotion algorithms were yet to be implemented for the full hemispherical arrangement. The egomotion results will be compiled in the coming weeks using a robot arm to create a known motion that can be used as ground truth for evaluation of the sensor performance. In this section, we present some results illustrating the functionality of the sensor in detecting light levels, and projecting those light levels on to a model.

4.2.1 Experimental Method

The hemispherical sensor is interfaced over the serial communications line to MATLAB which is being used for visualisation of the data from the sensor array. Each sensor value is decoded from the serial stream and indexed to its location on a rendering of the truncated icosohedron. By mapping the sensor value as the grey level in the MATLAB rendering, it is possible to visualise the data being produced by the sensor.

The sensor was placed in a disused office with two fluorescent lights on the ceiling and a window on one wall. The window had adjustable blinds. After capturing images with different combinations of light and window blind settings, another set of images was taken with a conventional camera using a panoramic mirror.

4.2.2 Experimental Results

The results show a side by side comparison of the images from the hemispherical sensor with the images from the conventional camera. Figures 14, 15 and 16 show the output of the MATLAB rendering compared to images taken with a conventional camera using a panoramic mirror. Note that the images taken with conventional camera have an obstructed area in the centre where the image is blocked by the camera lens and mirror support.

The results show that there is some noticeable inconsistency in the response of the individual sensors, most likely due to tolerances of the biasing resistors or variations in the phototransistor gains. We will implement a calibration procedure to create a uniform response across all sensors, with any gross adjustments made by altering resistor values, and fine adjustments implemented in software.

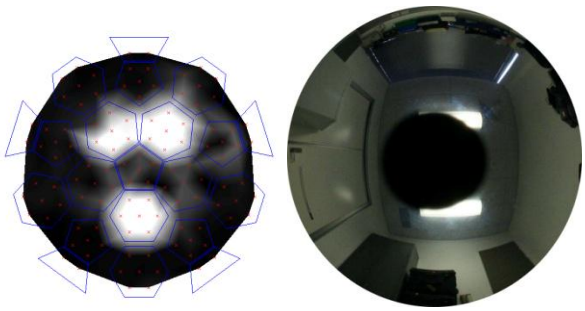

Figure 14 – Comparison of hemispherical sensor with panoramic image in office with both lights on and blinds closed.

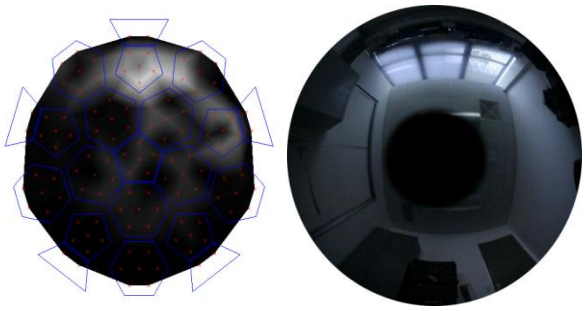

Figure 15 – Comparison of hemispherical sensor with panoramic image in office with both lights off and blinds open.

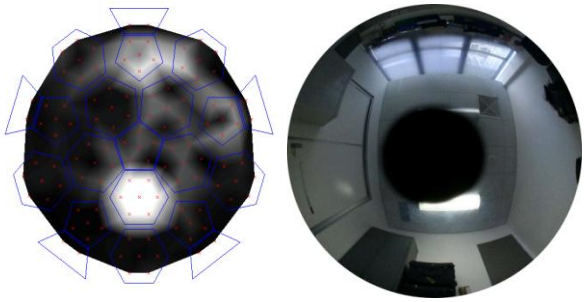

Figure 16 – Comparison of hemispherical sensor with panoramic image in office with one lights on and blinds open.

5 Conclusions

This paper has presented a unique hemispherical sensor designed to compute egomotion in six degrees of freedom. The sensor's unique spherical shape facilitates the decoupling of translational and rotational motion, which is a fundamental problem with egomotion estimation when using planar cameras. The camera produces low resolution images suitable for real time computation on low end processors.

The construction of the sensor was shown to be feasible with low cost off-the-shelf components, illustrating the viability of producing the sensor for a range of applications. The camera can produce frame rates of 600 fps, while only consuming 100 mW of power. Initial studies have shown the acceptable performance of the egomotion estimation algorithm around a single axis of rotation, and illustrated the sensing capabilities of the hemispherical sensor.

5.1 Future Work

Current work is centred on the calibration of the sensor so

that all photosites produce an equal response to light. The next step will be to implement the six degree of freedom version of the egomotion estimation algorithm and test it by moving the sensor through a known trajectory with a robot arm. These results will be available for presentation at the conference. With these results in place, we will interface the sensor with the mobile robots in our laboratory for further experiments in mapping and localistaion.

References

[1] D. Nister, O. Naroditsky and J. Bergen, "Visual odometry for ground vehicle applications," *Journal of Field Robotics*, vol 23, issue 1, pp. 3-20, Jan 2006.

[2] M. V. Srinivasan, M. Lehrer, W. H. Kirchner and S. W. Zhang, "Range perception through apparent image speed in freely-flying honeybees," *Visual Neuroscience*, vol 6, pp. 519-535, 1991.

[3] H. Wagner, "Flow-field variables trigger landing in flies," *Nature* (*London*), vol 297, pp. 147-148, 1982.

[4] M. V. Srinivasan, S. W. Zhang, M. Lehrer and T. S. Collett, "Honeybee navigation en route to the goal: visual flight control and odometry," *Journal of Experimental Biology*, vol 199, pp. 237-244, 1996.

[5] W. Reichardt, "Movement perception in insects," in *Processing of optical data by organisms and machines*, W. Reichardt, New York: Academic, 1969, pp. 465-493.

[6] R. R. Harrison and C. Koch, "A Robust Analog VLSI Motion Sensor Based on the Visual System of the Fly," *Autonomous Robots*, vol 7, issue 3, pp. 211-224, Oct 2004.

[7] J. C. Zufferey and D. Floreano, "Fly-Inspired Visual Steering of an Ultralight Indoor Aircraft," *IEEE Transactions on Robotics*, vol 22, issue 1, Feb 2006.

[8] J. Gluckman and S. K. Nayar, "Ego-motion and omnidirectional cameras," in *Sixth International Conference on Computer Vision*, 1998, pp. 999-1005.

[9] S. Thakoor, J. Chahl, M. V. Srinivasan, L. Young, F. Werblin and B. Hine, "Bioinspired engineering of exploration systems for NASA and DoD," *Artificial Life*, vol 8, issue 4, pp. 357-369, 2002.

[10] M. Egelhaaf, A. Borst and W. Reichardt, "Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system," Journal of the Optical Society of America A, vol 6, pp. 1070-1087, 1989.

[11] J. M. Zanker, M. V. Srinivasan and M. Egelhaaf, "Speed tuning in elementary motion detectors of the correlation type," Biological Cybernetics, vol 80, pp. 109-116, 1999.

[12] M. Egalhaaf and A. Borst, "Transient and steady-state response properties of movement detectors," Journal of the Optical Society of America, vol 6, pp. 116-127, 1989.

[13] H. Glunder, "Correlative velocity estimation: visual motion analysis, independent of object form, in arrays of velocity-tuned bilocal detectors," Journal of the Optical Society of America, vol 7, issue 2, pp. 255-263, Feb 1990.

[14] M. V. Srinivasan, M. Poteser and K. Kral, "Motion detection in insect orientation and navigation," Vision Research, vol 39, pp. 2749-2766, 1999.

[15] A. Bruss and B. K. P. Horn, "Passive Navigation," Computer Vision, Graphics and Image Processing, vol 21, pp. 3-20, 1983.