

A Vision System for Optic-flow-based Guidance of UAVs

Dean Soccol, Saul Thurrowgood, Mandyam Srinivasan

Biorobotics Laboratory, Queensland Brain Institute The University of Queensland, Brisbane, Queensland, 4072
and ARC Centre of Excellence in Vision Science, Australia

d.soccol@uq.edu.au, s.thurrowgood@uq.edu.au, m.srinivasan@uq.edu.au

Abstract

There is considerable interest in designing guidance systems for UAVs that use passive sensing (such as vision), rather than active sensing which can be bulky, expensive and stealth-compromising. Here we describe and evaluate a sensor that uses optic flow for measurement and control of height above the ground. A video camera is used in conjunction with a specially shaped reflective surface to simplify the computation of optic flow, and extend the range of aircraft speeds over which accurate data can be obtained. The imaging system also provides a useful geometrical remapping of the environment, which facilitates obstacle avoidance and computation of 3-D terrain maps. Laboratory tests of the performance of the device, whilst undergoing complex motions, are described.

1. Introduction

There is growing evidence that flying insects use optic flow cues to regulate flight speed, to estimate and control height above ground, to guide landing, and to avoid obstacles. [Tammero and Dickinson, 2002; Srinivasan and Zhang, 2004] There is also considerable interest in incorporating this principle to guidance of aircraft [Chahl and Srinivasan, 1999; Barrows *et al.*, 2003; Srinivasan *et al.*, 2004; Ruffier and Franceschini, 2005; Zufferey and Floreano, 2006]. Here we describe a vision system that is tailored for this objective.

Vision systems that are currently used in UAVs for this purpose consist of a downward-looking video camera to measure the optic flow generated by the ground [Chahl and Srinivasan, 1999; Barrows *et al.*, 2003; Srinivasan *et al.*, 2004; Ruffier and Franceschini, 2005]. Such a configuration, while being relatively simple to implement, has practical limitations. During low-altitude terrain following or landing, the image of the ground moves very rapidly, making it difficult to obtain accurate estimates of optic flow. The aim of this study is to design and test a specially shaped mirror surface that, firstly, scales down the speed of image motion as seen by the camera, and, secondly, removes the perspective distortion (and therefore distortion in image velocity) that a camera experiences when viewing a horizontal plane that stretches out to infinity in front of the aircraft. Ideally, the

moving image that is captured by the camera through the mirror should exhibit a constant, low velocity everywhere, thus simplifying the optic flow measurements and increasing their accuracy (see Fig. 1).

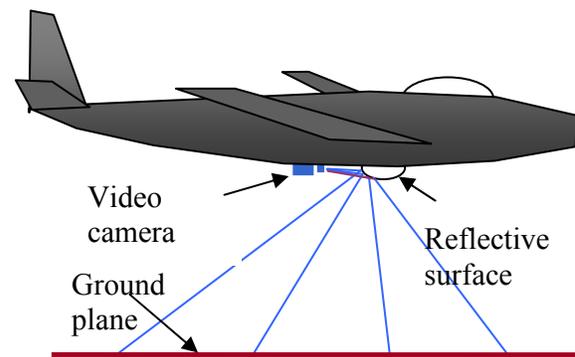


Fig. 1. Illustration of a system for visually guided terrain following and landing. The optical system is shown on an enlarged scale relative to the aircraft in order to clarify its configuration

2. Mirror profile

We seek a mirror profile which maps equal distances along the ground in the flight direction to equal displacements in the image plane of the camera. The camera-mirror configuration is shown in Fig. 2. A' is the camera image of a point A on the ground.

We want A' (the image of A) to move at a constant velocity in the image plane of the camera, independent of the position of A in the ground plane. That is, we require $\frac{d\eta}{dt} = L$, where η is the distance of the point A' from the centre of the camera's image plane, and L is the desired constant velocity. f is the focal length of the camera. The derivation of a mirror profile that meets these objectives is given in [Srinivasan *et al.*, 2006] and we shall not repeat it here. Instead, we use an example profile to illustrate the performance of the mirror.

3. Example of mirror profile

Fig. 3 shows one example of a profile of the reflective surface, and includes the computed ray paths. In this example the camera faces forward, in the direction of flight. The nodal point of the camera is at (0.0). The

image plane of the camera is to the left of this point and is not included in the figure, but its reflection about the nodal point (an equivalent representation) is depicted by

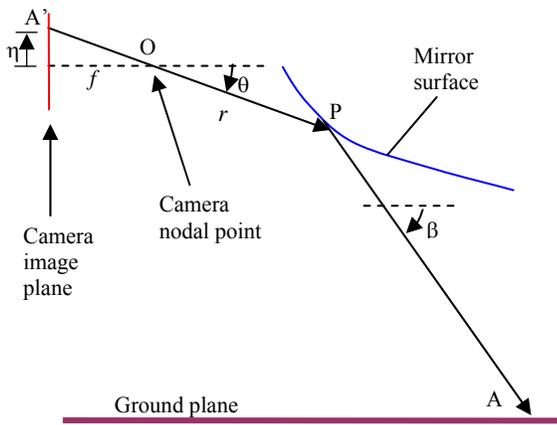


Fig. 2. Geometry of camera/mirror configuration for design of mirror profile

the vertical blue line to the right of the nodal point. Parameters used in this design are: $V=1000.0$ cm/sec; $h=100.0$ cm; $r_0 = 10.0$ cm; $f = 3.5$ cm; $L=2.0$ cm/sec, where V is the speed of the aircraft, h is the height above the ground, L is the image velocity, f is the focal length of the camera, and r_0 is the distance from the nodal point of the camera to the tip of the reflective surface. If the camera was looking directly downward at the ground, the image velocity of the ground would have been 35 cm/sec; with the mirror, the image velocity is reduced to 2.5 cm/sec. Thus, the mirror scales down the image velocity by a factor of 14.

The curvature of the mirror is highest in the region that images the ground directly beneath the aircraft, because this is the region of the ground that moves at the highest angular velocity with respect to the camera, and which therefore requires the greatest reduction of motion. We see from Fig. 3 that equally spaced points on the ground along the line of flight map to equally spaced points in the camera image. This confirms the correct operation of the surface.

Fig. 4 illustrates the imaging properties of the mirror, positioned with its axis parallel to and above a plane carrying a checkerboard pattern. Note that the mirror has removed the perspective distortion (foreshortening) of the image of the plane. The scale of the mapping depends upon the radial direction. Compression is lowest in the vertical radial direction and highest in the horizontal radial direction.

Fig. 5 shows a digitally remapped view of the image in Fig. 4, in which the polar co-ordinates of each pixel in Fig. 4 are plotted as Cartesian co-ordinates. Here the vertical axis represents radial distance from the centre of the image of Fig 4, and the horizontal axis represents the angle of rotation about the optic axis. The circle in Fig. 4 maps to the horizontal line in Fig. 5. Regions below the line represent areas in front of the aircraft, and regions

above it areas behind. Thus, the mirror endows the camera with a large field of view that covers regions in front of, below and behind the aircraft.

A consequence of the geometrical mapping produced by the mirror is that, for straight and level flight parallel to the ground plane, the optic flow vectors will have constant magnitude along each radius but will be largest along the vertical radius and smallest (zero) along the horizontal radius.

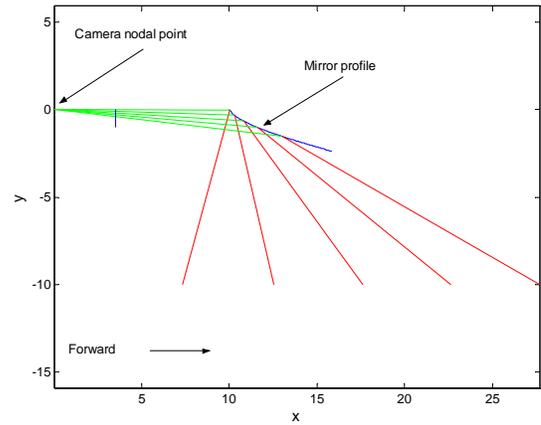


Fig. 3. Example of computed mirror profile

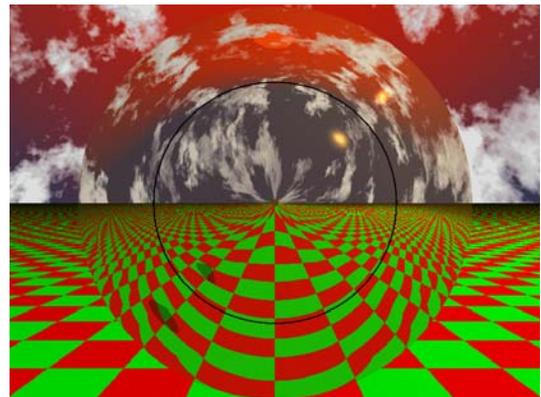


Fig. 4. Illustration of imaging properties of mirror

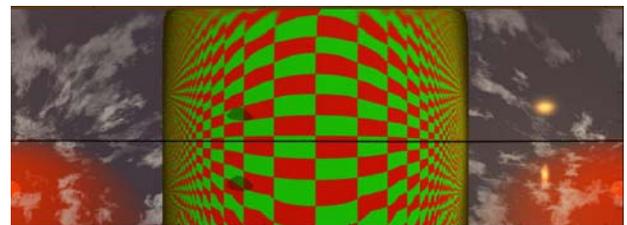
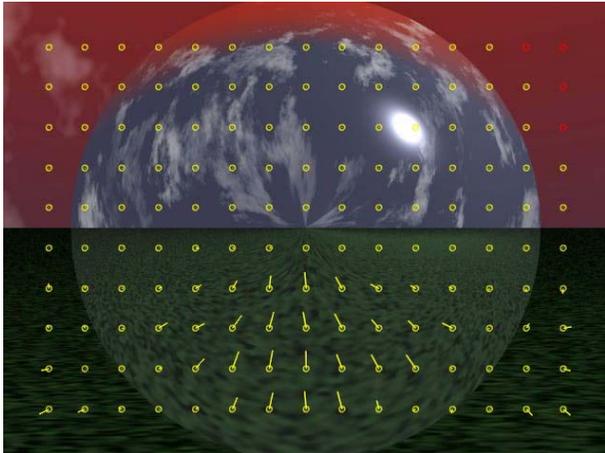


Fig. 5. Remapped version of image in Fig 4.

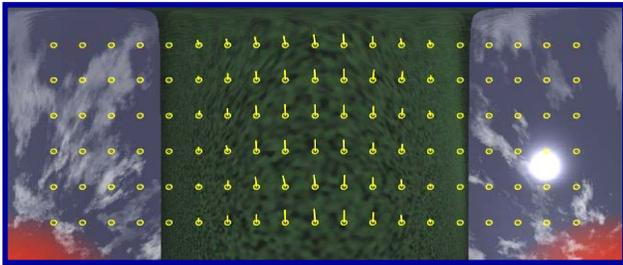
This is illustrated in Fig. 6, which shows the optic flow vectors that are generated by a simulated level flight over an ocean. Fig 6a shows the flow field in the raw image, and Fig. 6b the flow field in the remapped image. In the remapped image, all flow vectors are oriented vertically. The magnitudes of the vectors are constant within each column, and decrease progressively with increasingly lateral directions of view. These vectors provide

information on the height above the ground, the topography of the ground, and the ranges to objects in the field of view (which is quite large).

If the system undergoes pure translation along its optic axis, as in Fig. 6, the magnitude of the optic flow vector at each point on the unwarped image will be inversely proportional to the distance of the viewed point from the optic axis of the system. Thus, the flow vectors provide



a



b

Fig. 6. Optic flow field generated by simulated level flight over an ocean. (a) shows the flow in the raw image, and (b) the flow in the remapped image.

information on the profile of the terrain in a cylindrical co-ordinate system relative to the aircraft. The mapping that is provided by the mirror should be particularly useful for aircraft guidance. If a cylinder of “clear space” is desired for obstacle-free flight along a given trajectory, the maximum permissible flow magnitude is determined by the speed of the aircraft and the radius R of this cylinder (see Fig. 7). This simplifies the problem of determining in advance whether an intended flight trajectory through the environment will be collision-free, and of making any necessary adjustments to the trajectory to ensure safe flight. The system will provide information on the height above the ground, as well as the distance of potential obstacles, as measured from the optical axis.

4. Height estimation and obstacle detection during complex motions

If the aircraft undergoes rotation as well as translation,

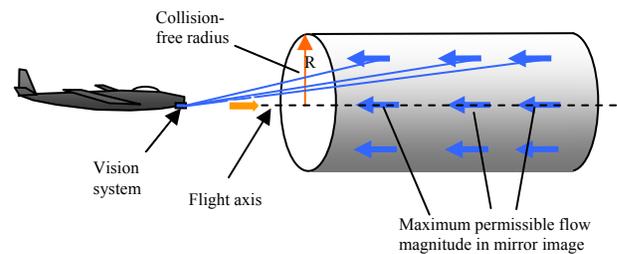


Fig. 7. Illustration of collision-free cylinder mapping achieved by the system

the optic flow vectors that are induced by the rotational components of aircraft motion will contaminate the measurement. They must be subtracted from the total optic flow field, to obtain a residual flow field that represents the optic flow that is created only by the translational component of aircraft motion. The rotations of the aircraft can be estimated through gyroscopic signals. Since each component of rotation (yaw, roll, pitch) produces a known, characteristic pattern of optic flow which depends on the magnitude of the rotation but is independent of the ranges of objects or surfaces in the environment, the optic flow vector fields that are generated by the rotations can be predicted computationally as a sum of the characteristic optic flow fields for yaw, pitch and roll, each weighted by the magnitude of the measured rotation about the corresponding axis. This composite rotational flow field must then be subtracted from the total flow field to obtain a residual flow field that represents the flow due to just the translational component of motion. The distances to objects and surfaces from the optical axis can then be computed from this residual field.

5. Hardware realization and system tests

The mirror profile shown in Fig. 3 was machined in aluminium on a numerically controlled lathe. It was mounted on a bracket which also carried analog video CCD camera (320 x 240 pixels) with its optical axis aligned with the axis of the mirror. The nodal point of the camera’s lens was positioned 10cm from the tip of the mirror, as per the design specifications (see above).

In order to extract the range of objects during complex motions, it is necessary to first determine the flow signatures, or “templates”, that characterise the patterns of optic flow during pure yaw, pure roll and pure pitch. This was done by using a robotic gantry to move the vision system in a richly textured visual environment. The environment consisted of a rectangular arena 3.05 m long, 2.2 m wide and 1.13 m tall (Fig. 8). The walls and floor of the arena were lined with a texture composed of black circles of five different diameters (150mm, 105mm, 90mm, 75mm and 65mm) on a white background. The rich visual texture permitted dense and accurate measurements of the optic flow in the lower hemisphere of the visual field. A raw image of the arena, as acquired by the system, is shown in Fig. 9a. An unwarped version of this image is shown in Fig. 9b.

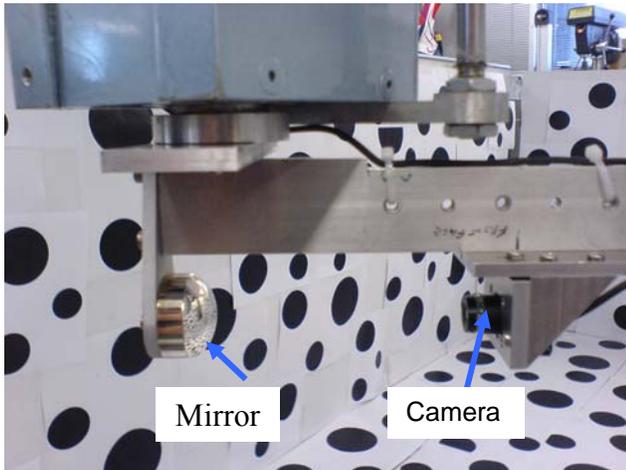
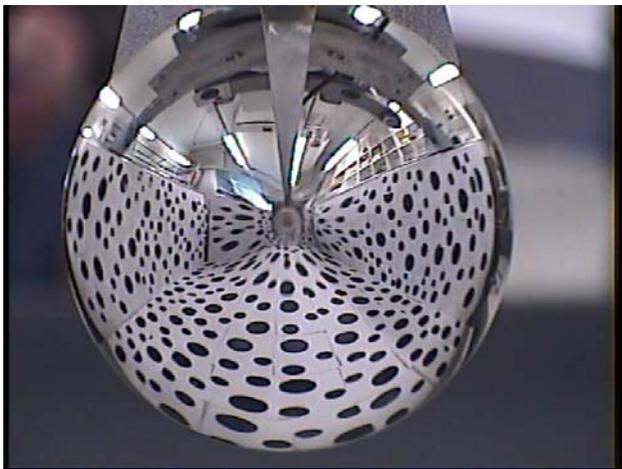


Fig. 8. View of vision system, carried by a robotic gantry in a visually textured arena.



a



b

Fig 9: Raw (a) and unwarped (b) images of the arena as viewed by the system.

The gantry was used to position the optical axis of the system at a height of 650 cm above the floor. Optic-flow templates for yaw, pitch and roll were obtained for the remapped images by using the gantry to rotate the vision system by small, known angles (ranging from 0.25deg to 2.5 deg, in steps of 0.25 deg) about each of the three axes, in turn. Measurements were repeated with the vision system positioned at several different locations in the

arena and the results were pooled and normalized to obtain reliable and dense estimates of the optic-flow templates for a 1-deg rotation (In theory, the rotational optic-flow templates should be independent of the position or attitude of the vision system within the arena.) The optic flow was computed using a correlation algorithm [Fua, 1993]. The resulting rotational template for yaw is shown in Fig. 10, for a rotation of 1 deg in the CCW direction.

For small angular rotations, the magnitude of the flow vector at any given point in the visual field should increase approximately linearly with the magnitude of the rotation, and the direction of the flow vector should remain approximately constant. This is verified in Fig. 11, which compares the magnitudes and directions of the flow vectors for yaw rotations ranging from 0.25 deg to 2.5 deg, at two different locations indicated by the circles in Fig. 10. This result validates the use of scalable optic-flow templates for rotation.

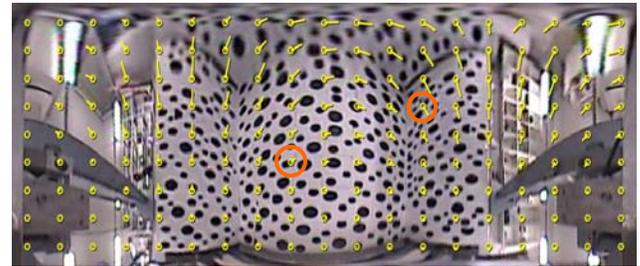


Fig. 10: Rotational flow template for yaw.

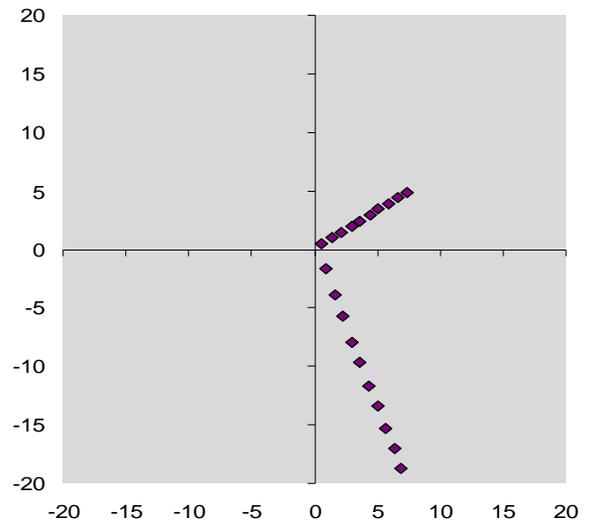


Fig. 11. Variation of magnitude and direction of optic flow vectors for yaw amplitudes ranging from 0.25 deg to 2.5 deg at two different locations in the visual field, indicated by the orange circles in Fig. 10.

The next step was to test performance when the system executed compound motions that combined translation and rotation. The aim was to investigate whether the system could determine the range to objects in the environment (and specifically, the height above the ground) whilst executing complex motions. This investigation was performed by using the gantry to move

the system along a curved trajectory, as shown in Fig. 12.

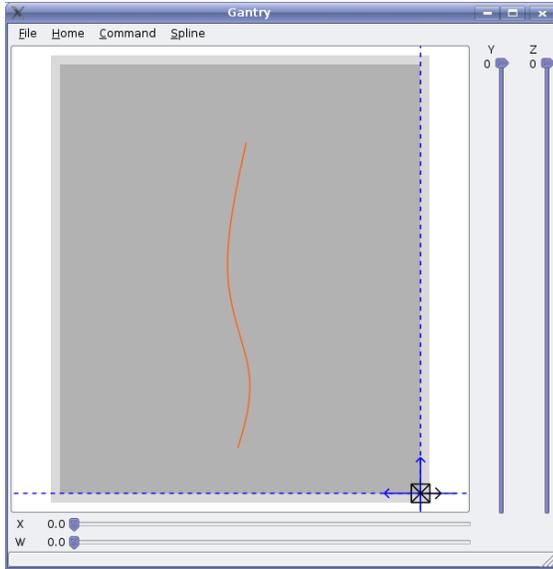


Fig 12: Plan view of a curved trajectory used to test the system

The height of the system above the floor was held constant at 650 mm throughout the trajectory. The trajectory consisted of a sequence of stepwise motions in the horizontal plane. The optical axis of the system was always aligned along the instantaneous direction of translation, i.e it was parallel to the local tangent to the trajectory. Each step, in general, consisted of an elementary translation (of 50 mm) along the optical axis of the vision system, followed by an elementary yaw rotation of known, but variable magnitude (ranging from +2.9 deg to -1.5 deg). A visual frame was acquired from the camera at the end of each elementary step (translation or rotation).

Fig. 13 shows the optic flow generated by a compound step, consisting of an elementary translation followed by an elementary rotation. This is the flow that would be generated between two successive frames if the system had moved along a smooth curve.

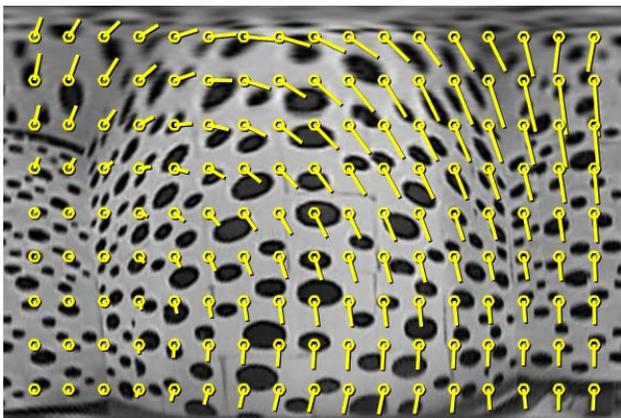


Fig. 13. Flow measured after a compound step consisting of translation followed by yaw

Fig. 14 shows the flow that would have been induced by the rotational component of motion that occurred during this compound step. This flow pattern was computed by weighting the template for yaw rotation by the (known) magnitude and polarity of the yaw that occurred during the compound step. (During flight, this yaw information would be obtained from a rate gyro.)

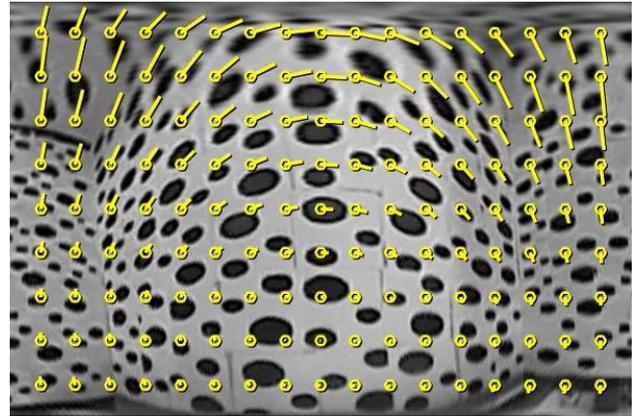


Fig. 14. Flow induced by the rotational (yaw) component, calculated from the template for yaw.

Fig. 15 shows the optic flow that is obtained when the rotational component of the flow (from Fig. 14) is subtracted from the total flow that is generated by the compound step (Fig. 13). This residual flow should represent the flow that is induced solely by the translational component of the system's motion. It is

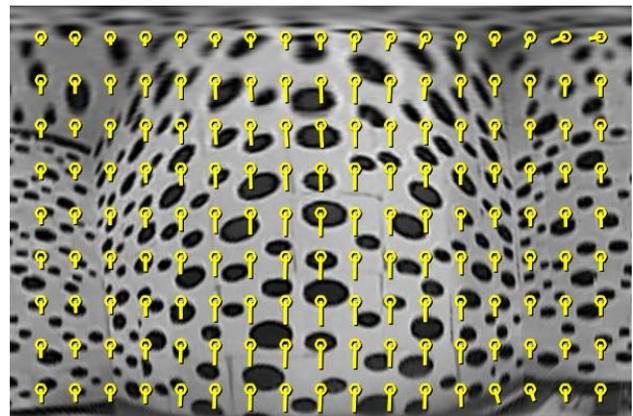


Fig. 15 Residual flow, obtained by subtracting the yaw-induced flow (Fig. 14) from the measured flow (Fig. 13).

evident that all of the vectors in the residual flow field are parallel, as would be expected during pure translation. Furthermore, this pattern of optic flow is in excellent agreement with the pattern of flow that is actually generated by pure translation at this particular location in the arena. This latter flow pattern, shown in Fig. 16, is the flow measured between the two frames that bracketed the translatory segment of the compound step. A comparison of Figs. 15 and 16 reveals that these flow patterns are virtually identical. This result, which was obtained

consistently for all of the compound steps in the trajectory, reveals that the pattern of optic flow that is

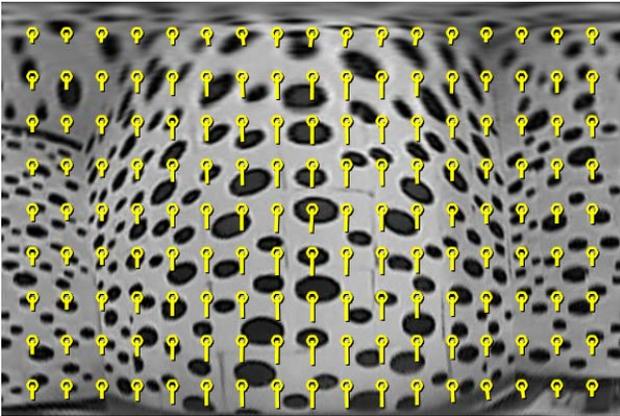


Fig. 16. Flow induced by the translatory component of motion in the compound step

generated during a small, but complex motion can be successfully ‘de-rotated’ to extract the optic flow that is created purely by the translatory component of motion that occurred during the period.

6. Extracting information on range and topography

The final step is to examine whether the residual component of the optic flow (the translation-induced component), can be used to obtain accurate information on the range and profile of the terrain over which flight occurs. With reference to Fig. 15, for horizontal flight over level ground, the magnitudes of the translation-induced optic flow vectors should be a maximum in the central column (corresponding to the ground directly beneath the aircraft) and should fall off as a cosine function of the lateral angle of view (in the columns to the left and right; [Srinivasan et al., 2006]). This should be true for any row of vectors.

This prediction is tested in Fig. 17, which shows the variation of translatory flow magnitude with lateral angle for flow vectors in any given row, in data sets such as that shown in Fig. 14. The results show the magnitude profiles for the vectors in the second row from the bottom, computed for each step of the trajectory. The second row represents a view that is oriented at approximately 90 deg to the axial direction (a lateral view). A least-squares analysis reveals that each of the profiles approximates a cosine function quite well. The red curve shows the mean of the cosine functions fitted to each of the profiles obtained along the trajectory. If the speed of the aircraft is known, the amplitude of this curve provides information on the height above the ground (the amplitude is inversely proportional to height). The mean amplitude of the curves is 9.076 pixels and the standard deviation is 0.062 pixels, indicating that the estimate of flight height is consistent and reliable throughout the trajectory of Fig. 12.

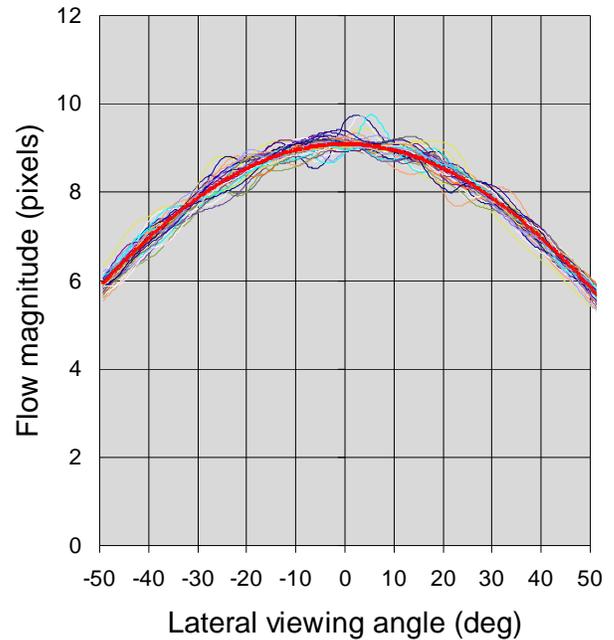


Fig. 17. Variation of flow magnitude with lateral viewing angle in the residual flow field at each of the steps along the trajectory of Fig. 12. The thick red curve represents a cosine function fitted to the profiles.

7. Conclusions

This study has described the design of a vision sensor, based partly on principles of insect vision and optic flow analysis, for measurement and control of flight height and for obstacle avoidance. A video camera is used in conjunction with a specially shaped reflective surface to simplify the computation of optic flow, and extend the range of aircraft speeds over which accurate data can be obtained. The imaging system also provides a useful geometrical remapping of the environment, which facilitates obstacle avoidance and computation of 3-D terrain maps. By using calibrated optic flow templates for yaw, roll and pitch, accurate range information can be obtained even when the aircraft executes complex motions. Future work will involve open-loop and closed-loop tests on flying vehicles, to examine the utility of this approach for vehicle guidance.

8. Acknowledgements

This work was supported partly by the US Army Research Office MURI ARMY-W911NF041076, Technical Monitor Dr Tom Doligalski, US ONR Award N00014-04-1-0334 and the ARC Centre of Excellence Grant CE0561903.

9. References

[Barrows *et al.*, 2003] G.L. Barrows, J.S. Chahl and M.V. Srinivasan. Biologically inspired visual sensing and flight control. *The Aeronautical Journal, London: The Royal Aeronautical Society* 107:1069, 159-168, 2003.

[Chahl and Srinivasan, 1999] J.S. Chahl, and M.V. Srinivasan. Panoramic vision system for imaging, ranging and navigation in three dimensions. *Proceedings, Field and Service Robotics Conference, Pittsburgh*, August 29-31, 127-132, 1999.

[Fua, 1993] P. Fua. A parallel stereo algorithm that produces dense depth maps and preserves image features. *Machine Vision and Applications* 6 (1): 35-49, December 1993.

[Ruffier and Franceschini, 2005] F. Ruffier and N. Franceschini. Optic flow regulation: the key to aircraft automatic guidance. *Robotics and Autonomous Systems* 50:177-194, 2005.

[Srinivasan and Zhang, 2004] M.V. Srinivasan and S.W. Zhang. Visual motor computations in insects. *Annual Review of Neuroscience* 27: 679-696, 2004.

[Srinivasan *et al.*, 2004] M.V. Srinivasan., S.W. Zhang, J S Chahl, G Stange and M Garratt. An overview of insect inspired guidance for application in ground and airborne platforms. *Proc Inst Mech Engrs Part G*, 218: 375-388, 2004.

[Srinivasan *et al.*, 2006] M.V. Srinivasan, S. Thurrowgood and D. Soccol. An optical system for guidance of terrain following in UAVs. *Proceedings, IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS '06)*, Sydney, 51-56, 2006.

[Tammero and Dickinson, 2002] L.F. Tammero and M.F. Dickinson. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *Journal of Experimental Biology* 205: 327-343, 2002.

[Zufferey and Floreano, 2006] J.C. Zufferey and D. Floreano. Fly-inspired visual steering of an ultralight indoor aircraft. *IEEE Transactions on Robotics* 22: 137-146, 2006.