

Calculating Attitude from Horizon Vision

Terry Cornall¹, Greg Egan²

¹ *CTIE, Electrical and Computer Systems Eng., Monash University, Wellington Rd, Clayton, Victoria, 3168, Australia*

² *CTIE, Electrical and Computer Systems Eng., Monash University, Wellington Rd, Clayton, Victoria, 3168, Australia*

Summary: The horizon angle is calculated as a function of the average coordinates for the ground and sky classes and a measure of the aircraft pitch is determined by the displacement of the horizon from the centre of the view.

A prototype camera and image processing system has been built and used to test and validate the procedures. Trials of the system in simulation and real flights in a remote controlled glider have been carried out. The results are given and discussed.

Keywords: UAV, unmanned aircraft, sky, ground, image processing, computer vision, horizon detection and tracking.

Introduction

The displacement and angle of the horizon in a video frame from a video camera onboard an aircraft can inform us about the attitude of the camera and hence of the aircraft. This is ongoing work being done by the authors as well as other groups. [1][2][3][4][5][6][7][8][9][10]

Given an algorithm that can segment a video frame into ground and sky parts and a measure of how reliable the segmentation is it is possible to measure the approximate angle and position of the horizon. Methods for the image processing to produce a binary representation of the horizon and measure of the reliability of the representation are discussed in [2]. It should be noted that although for the purposes of this paper the difficulties of obtaining this segmentation are mostly ignored, this is no mean task and in reality the entire success or otherwise of the technique relies on this image processing and the reliability measures that it produces.

For the authors' purposes, a good segmentation has clearly defined sky and ground classes with little or no overlap and a well defined interface, the horizon. A circularly shaped view is required for our work because it makes the measurement of the horizon angle simpler, given the average coordinates (centroids) of the classes. Essentially it is because the area under and above the horizon line in a rectangular view is not invariant with the angle that the horizon makes with the horizontal (X axis) of the image. With a circular view this asymmetry disappears and this makes it easier to calculate the angle and position of the horizon from the spatial relationships of the ground and sky centroids in the binary image.

Algorithm

As discussed in [1], if the segmented horizon image makes a perfect chord of the circular view, then the gradient m and angle f that the horizon makes to the camera horizontal are given by Eqn. 1.

$$m = \frac{X_S - X_G}{Y_S - Y_G}$$
$$f = \arctan(m)$$

Eqn. 1

Where X_S, Y_S and X_G, Y_G are, respectively, the Cartesian coordinates of the sky and ground centroids in the binary image of the horizon, that is, the class centroids. In order to measure the centroids a threshold in the blue component of the RGB video images is applied to create a binary image which then has a circular mask applied to it in software as the pixel data are received. The coordinates of pixels inside the mask that arrive with a value of 1, indicating that they had been classified as sky, are added to an accumulated total for the X_S and Y_S , and similarly treated are the coordinates of pixels classified as ground. The number of pixels in each class are also counted on the fly. At the end of the frame the X_S, Y_S and X_G, Y_G totals are divided by the number of pixels in each class to arrive at the average coordinates and this information is used with Eqn. 1 to arrive at the horizon angle. All of these procedures are done on a pixel by pixel basis so there is no need for a frame memory, which simplifies the hardware and software design considerably and reduced memory transfer overheads. The arctan function is calculated using a Taylor's series approximation. This calculation can be carried out easily by a low-powered processor at the rate of arrival of the binary thresholded pixel data. *Fig. 1* shows an example of this process where, after adjusting for the Y coordinate

increasing downwards instead of upwards we have: $\arctan\left(\frac{90-45}{-(55-100)}\right) \times \frac{180}{p} = -45$

Unlike the horizon angle, in the trial discussed in this paper the horizon displacement was not calculated onboard the aircraft but rather was calculated in the post processing stage. It is a relatively simple task that will be added to the airborne processor for future work. The actual pitch angle of the aircraft is not calculated because to do so would require knowledge of the altitude, the horizon displacement being a function of both altitude and pitch angle (and camera mount angle, which is being ignored herein).

The horizon displacement is measured as an angle by measuring the distance in pixels that the line representing the horizon is from the centre of the view. Given that the field of view of the camera is known it is possible to work out how many radians is subtended per pixel and thus the angle of the pitch displacement can be calculated, but for our purposes, it is probably just as easy to leave the horizon displacement expressed in terms of pixels rather than convert to radians.

Using some simple trigonometry, the relationship that gives the distance h (also known as the apothem) of the horizon chord line from the centre of the view as a function of the number of pixels n of the image underneath the chord can be found as expressed in Eqn. 2, where R is the radius of the circular view mask.

$$n = R^2 \arccos\left(\frac{h}{R}\right) - h\sqrt{R^2 - h^2}$$

Eqn. 2

Unfortunately, this non-algebraic function cannot be re-written to simply express h in terms of n . However, the function can be approximated for a known value of R . For example, given $R=72$, we can use the linear equation given by Eqn. 3.

$$h = 65 - 0.00837n$$

Eqn. 3

If more accuracy is required, a non-linear polynomial approximation can be used.

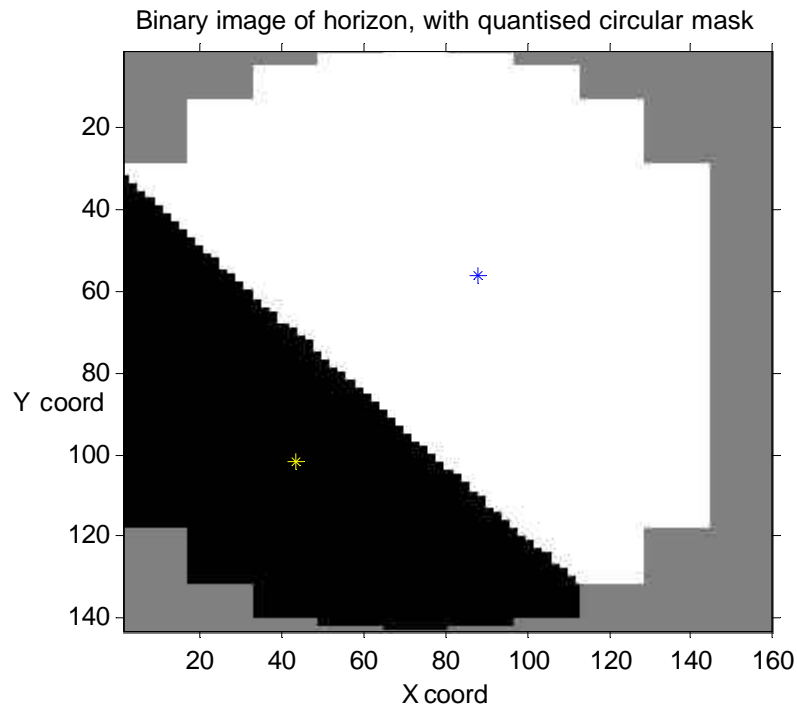


Fig. 1 Masked, thresholded horizon with class centroids

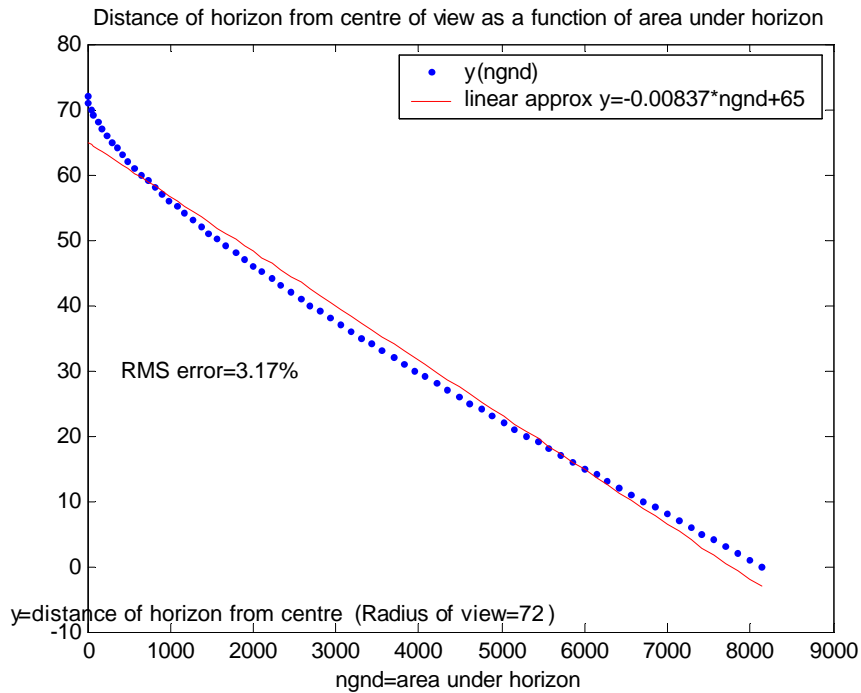


Fig. 2 Linear approximation for horizon displacement

Method and Equipment

Although other groups are carrying out the investigation of this type of technique using UAVs and MAVs fitted with video telemetry systems and doing the image processing on the ground and then transmitting control signals back to the aircraft, [3][4][8], the authors of this paper are determined that the image capture, processing, angle measurements and control responses should all happen onboard the aircraft. This gives the great advantage of not having to rely on the quality of the video transmission to the ground station and the control signals from the ground station. However, it does severely limit the size and weight of the image processing system because it is intended to be flown in an aircraft with a payload capacity measured in hundreds of grams, and with a limited electrical power budget.

For the trials and results presented in this paper the airborne equipment consisted of a digital video camera with some simple image processing capabilities, a microcontroller programmed to control the camera and carry out the analysis and angle measurements and relay the results to the ground station via a digital radio link. A second video camera and video telemetry radio were used to transmit the horizon imagery to the ground station for confirmation of the onboard measurements. See *Fig. 3*. Power for the equipment was supplied by a three cell lithium-polymer battery. The radio control receiver, servos and battery were kept quite separate from the other equipment for reliability and safety reasons. The ground station consisted of the relevant radio receivers and antennas for the two radio links, a laptop computer for recording the serial data stream and a digital handy-cam for recording the video telemetry. On the ground were also the pilot and radio control transmitter.

The digital video camera used was the OV6620 from OmniVision Technologies [11] with an image processing board designed and programmed by students [12] of Carnegie-Melon University. This camera system, which we will call the CMUCam, has, amongst other attributes, the ability to apply a threshold to the frame and send a binary frame as a result, thus

greatly reducing the amount of data that needs to be handled by higher level processing systems.

The microcontroller that is the heart of the horizon angle measuring system was a low-powered device by modern standards, being a 20 MHz PIC 16F876 manufactured by Microchip [13]. The software to control the camera and receive and process the intermediate image results was written by the authors in the C programming language. The frames were processed at a rate of 5 frames per second and this limit was imposed mainly by the time taken to receive the intermediate image via the serial connection between the camera and the microcontroller and process it. Despite the simplicity of the algorithm, the software loop that processed the binary results from the camera and applied the circular mask and accumulated the average coordinates of the classes could not have handled a faster frame rate. An assembler language inner loop was written and tested but eventually discarded in favour of the slightly slower but more easily maintained C code.




Fig. 3 Airborne equipment, left to right, microcontroller with data telemetry radio, OV6620 camera with CMUCam image processor, analog video camera and video telemetry radio

Trial and Results

One trial and its outcome will be discussed in this section. The conditions that the trial was carried out under were considered to be very good and there was a known weakness in the thresholding method used for the trials. The results are considered as a proof of concept rather than proof that the system will work under all conditions. The trial, known as Grampians2, was conducted on a clear day with little or no cloud and in the early afternoon. The air was moderate, with a light breeze and little turbulence. The ground visual appearance was lightly treed grassland and bush, with low mountains in the distance, and a lake just on the horizon. This trial was over farmland acquired by one of the professors at Monash University for the testing of robotic vehicles.

Launch of the glider was by ‘bungee’ consisting of a stretched rubber tube and nylon line. The flight was on the order of a few minutes as there was not a lot of lift. No feedback from the visual system was applied and the radio control pilot was in control at all times. Video and data telemetry was recorded on the ground for later processing and graphing. The horizon angle measurements were calculated by the airborne equipment. Graphing and image processing for the purposes of presentation were conducted by the ground equipment.

In this trial, a fixed threshold method was used for segmenting the image into binary ground and sky classes. This far from optimal method was adopted for reasons of simplicity as the trial was really to prove the concept and the equipment rather than be the final tests. Later

tests will be carried out using an adaptive thresholding method, or a k-means clustering method as discussed in [2].

A selection of frames from the resulting video with the horizon angle and horizon displacement superimposed is shown in *Fig. 4*. This selection shows a good correspondence between the video and the measured angles, although there are certainly other sections in the video that are less satisfactory. The horizon angle shown by the red line in the blue circle was reported from the airborne sensor via telemetry, although the horizon displacement was calculated during post-processing. The blue line along the image of the horizon was calculated by the post-processing method to get a measurement to compare the telemetry angle to.

Fig. 5 shows the angle reported by the airborne processor compared to the horizon angle derived by post-processing the video using a more sophisticated method, which was Otsu's histogram analysis thresholding [14], followed by filling in of all objects less than ten percent of half of the view size, then edge pixel finding followed by fitting a line to the edge pixels. Observation of the video produced during this post-processing indicates that the post-processing angle derived from the video telemetry usually corresponds well to the angle that a human observer would decide. There are some exceptions such as that due poor telemetry, sun-glare or when the horizon goes completely out of view and these have not been compensated for. It was found that the angles reported by telemetry had a non-zero offset compared to the angles derived during post-processing and that was subtracted before comparison. This offset was probably due to the camera mounting. Angles in data packets that were detected as corrupt or missing in the telemetry were replaced by prior good angles. Contemplation of *Fig. 5* shows that there are areas where the two angles substantial in agreement, but others where there is substantial disagreement, such as in the region of frame 150. Some of this disagreement is due to sun glare and different responses when the horizon goes out of view, but once again a lot is thought to be due to the simple fixed threshold used for segmentation in the airborne equipment.



Fig. 4 Horizon angles superimposed on video. Grampians2

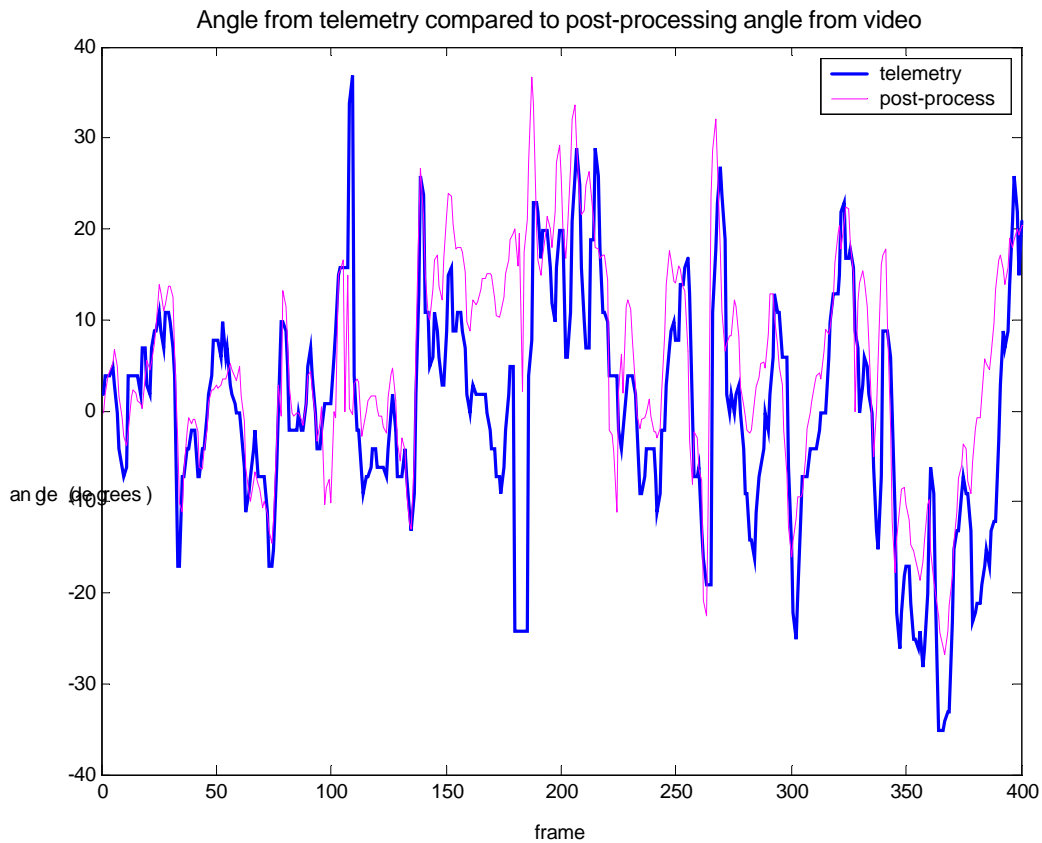


Fig. 5 Angle measured by airborne equipment compared to that measured after post-processing the video telemetry

Conclusion

The fundamental idea of using low complexity image processing and computation onboard a UAV for detecting and measuring horizon angle and displacement does work, at least under good visual conditions when there is good contrast between ground and sky. Improvement needs to be made in the image segmentation and a reliability measure needs to be calculated and applied to the results before action is taken by the control systems using that information. There will be occasions when the technique cannot be used because the video image is not suitable or because the terrain of the horizon does not match the model used, and these situations also must be detected and appropriate actions taken. The use of other sensors such as GPS and inertial measurement systems should be integrated to work in concert with this visual sensing method.

The CMUCam has served its purpose well to allow us to test the above ideas, but it will probably be replaced in future work. There is a new version, the CMUCam2, available that has the ability to measure a histogram of a color component, which fits in very well with the methods discussed in [2] and a frame buffer built in that decouples raw frame rate (and hence exposure control) and data-transfer rates. Work is also being carried out at Monash to finish designing and implementing a vision processing system based on a FPGA (field programmable gate-array) device coupled with the OV6620 digital front end. This should allow us much greater flexibility in the image processing, allowing fast parallel measurements of statistics using dedicated hardware described in HDL (hardware description language) in conjunction with more conventional procedural software routines.

References

- [1] Terry Cornall and Greg Egan. "Measuring Horizon Angle from Video on a Small Unmanned Air Vehicle" ICARA 2004, Palmeston Nth N.Z. Dec 13-15.
- [2] Terry Cornall and Greg Egan, "Heaven and Earth: How to tell the difference." Paper no WC0055 Australian International Aerospace Congress, Melbourne Australia, March 2005.
- [3] S. Ettinger, P. Ifju and M. C. Nechyba, "Vision-Guided Flight for Micro Air Vehicles", <http://mil.ufl.edu/~nechyba/mav/index.html#vision1>, 2002, September.
- [4] S. Ettinger, M. Nechyba, P. Ifju and M. Waszak, "Vision-guided flight stability and control for micro air vehicles", International Conference on Intelligent Robots and Systems (IEEE/RSJ). Sept, 2002, volume 3, number 30, pages 2134 to 2140,
- [5] Luc Fety, Michel Terre and Xavier Noreve, "Image Processing for the Detection of the Horizon and Device for the Implementation Thereof", Thompson TRT Defense, 1991, United States Patent, number = 5,214,720,
- [6] G. Stange, S. Stowe, J.S. Chahl and A. Massaro, "Anisotropic imaging in the dragonfly median ocellus: a matched filter for horizon detection.", Journal of Comparative Physiology A., 2002, VOLUME 188, PAGES 455 to 467
- [7] S. Todorovic, M. C. Nechyba and P. G. Ifju, "Sky/Ground Modeling for Autonomous MAV Flight", Proc. IEEE Int. Conf. on Robotics and Automation., May, 2003, Taiwan.
- [8] Centre for MAV Research, University of Florida, "Horizon Detection and Tracking", <http://www.mil.ufl.edu/mav/research/vision/horizontracking>, 2002, September.
- [9] G. Barrows, J. S. Chahl and M. V. Srinivasan, "Biomimetic Visual Sensing and Flight Control", Bristol UAV Conference, April, 2002, Bristol, UK.
- [10] T. Netter and N. Franceschini, "A Robotic Aircraft that Follows Terrain using a Neuromorphic Eye", (IEEE/RSJ) International Conference on Robots and Systems", October, 2002 Lausanne.
- [11] OmniVision Technologies Inc., "Advanced Information Preliminary regarding OV6620 Single-chip CMOS CIF Color Digital Camera", 2000, 930 Thompson Place Sunnyvale, California 94086 U.S.A., <http://www.ovt.com>
- [12] Anthony Rowe, Chuck Rosenberg and Illah Nourbakhsh, "A Low Cost Embedded Color Vision System", Proceedings of IROS 2002
- [13] Microchip, "PIC16F87XA Data Sheet", 2003, <http://www.microchip.com>
- [14] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," IEEE Transactions on Systems, Man, and Cybernetics, vol. 9, no. 1, pp. 62-66, 1979.