Aircraft attitude estimation from horizon video

T.D. Cornall, G.K. Egan and A. Price

Measurement of aircraft attitude is a necessary step for the autonomous control of the aircraft. Autonomous aircraft are useful for a range of military and civilian purposes. A method of very low computational complexity is presented that allows the pitch and roll angles of an aircraft to be measured onboard an unmanned aircraft from real-time video of the horizon. This method can be used for maintaining level flight, and also provides a degree of automatic terrain avoidance.

Introduction: In this Letter, a simple, novel method for extracting the aircraft roll and pitch angles from a single view of the visible horizon from the front of the aircraft and also for calculating the confidence in the measurement is discussed. The following ideas are introduced:

- Use a circular mask to reduce image asymmetry and to simplify the calculations.
- The horizon position and angle can be derived simply from the average co-ordinates of the sky and ground classes.
- A validity metric has been derived that measures how well the binary image matches the model.
- The use of this method in hilly terrain leads to an obstacle avoidance behaviour that has obvious utility for aircraft.

The image processing and angle calculation is done using computing equipment onboard the aircraft. The application of a circular mask to remove asymmetries simplifies the calculations considerably. The original image can then go on to be used for other purposes, such as navigational landmark or other feature recognition. When used to maintain level flight, this method also has the advantage of providing automatic terrain avoidance because the errors that looming obstacles produce in the determination of roll and pitch angles cause the aircraft to turn away from or climb over the obstacles.

Current methods of calculating aircraft attitude using inertial measurement systems rely on integrating angular rates to obtain changes in angles. This method works well in many applications but often lacks a reliable absolute reference to refer these changes to in order to find the current angles. Thus it can fall prey to accumulated errors or fail if the rates exceed the measurable ranges. There are existing methods discussed by Taylor *et al.* [1], Chahl *et al.* [2] Brookman [3], and Hatcher and Germann [4] that use the differences in infrared or ultraviolet light levels measured by a small number of point sensors to give a measure of attitude, but these methods do not allow any other image information to be extracted and often cannot give a measure of the reliability of the measurement. There is another method by Barrows and Neely [5] that uses the image on a specialised sensor to find a measure of optical flow and use this for obstacle avoidance. Optical flow methods only pick up a change in the image and would not necessarily allow the measure of an unchanging horizon, although if the image were extractable the method discussed in this Letter could be applied. There is an image processing method discussed by Todorovic *et al.* [6] that comes close to the method discussed here in that it classifies the pixels into ground and sky, but it requires a lot more computation and relies on a video downlink to do the computation on a ground station computer and relay control signals back to the vehicle. This reduces the autonomy of the vehicle. Of course, with enough computation power onboard, the radio links would be unnecessary.

Method: The video frame has a circular mask applied to it to remove the asymmetries due to the rectangular nature of the image sensor. If these asymmetries are not removed then the accuracy of the calculations is reduced at roll angles where the sky/ground partition is most asymmetrical, and this can contribute errors of 10% or more. The red, green and blue (*RGB*) components of each pixel value are combined into a single pixel value using the formula $3B^2/(R+G+B)$. This formula is the product of the blue component of the image (B) and $3B/(R+G+B)$, which is the ratio of the B to the brightness $(R+G+B)/3$. The product of these two measures of sky colour has been shown to work better to discriminate between sky and ground than either one alone. Each resulting video image is then analysed using Otsu's histogram method [7] to determine the value of the pixel value to be used as a threshold for binarisation. The resulting two classes of pixels in the image are then considered to represent either ground or sky. With reference to Fig. 1, the average x, y co-ordinate (centroid) for each class is then calculated. If we assume that the horizon is a straight line then it is a chord of the circular view. Thus, a perpendicular bisector of the chord will pass through the centre of the circle. This divides the sky (and the ground) into two symmetrical areas, so the centroids of the sky and ground must lie on the horizon's bisector. We know the centroid positions, so we can easily calculate the angle of the line joining them. This is the bisector and we can then calculate the angle perpendicular to that to find the horizon's angle. The position of the horizon can be derived with a little more work using the relative numbers of sky and ground pixels. An example of the resulting segmentation into sky and ground can be seen in Fig. 2.

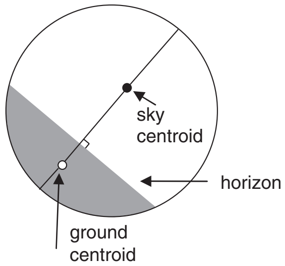

Fig. 1 *Calculation of horizon angle from class centroids*

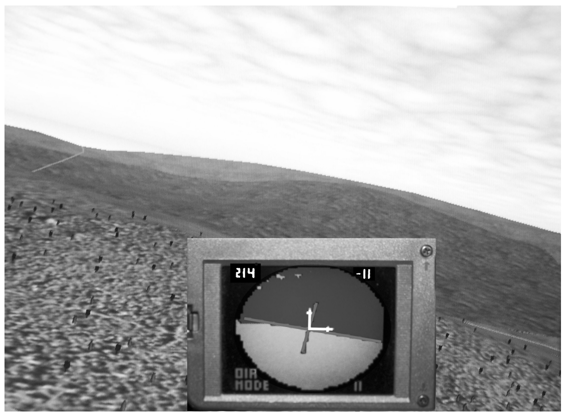

Fig. 2 *Horizon image classified into sky and ground*

The complexity of the computation needed to perform these calculations is relatively low. Each pixel value is processed once to calculate the histogram, then once to calculate the blueness value and the binary segmentation. Each pixel in the binary image is processed once to calculate the centroids. The horizon angle calculation then uses the two centroids only. If the image has N pixels, the computational complexity is of order $3N$ at worst. This can be reduced by computing the image histogram and binarisation threshold less often.

To reduce false determinations of horizons from the images, the original and the binary images are analysed to determine if they comply with the model of the horizon as an approximately straight line with distinct and separated areas of ground and sky. A number of confidence metrics are used, all readily measurable. Brightness and contrast are measured for the original image. The measure for brightness can be read directly from the camera in some cases, by reading the setting for automatic gain control (AGC) and automatic exposure control (AEC), or it can be measured by averaging the luminance part of the pixel values in the original image. An original image that is too bright might be due to a sky only view and one that is too dark might be due to a ground only view or because there is too little light available. Contrast can be derived from the variance in the global luminance, or the difference between average sky brightness and average ground brightness. Too low a contrast indicates an untrustworthy original image. Even if the original image passes on the grounds of appropriate brightness and contrast, a metric derived by the authors and called *M2P5* is measured for the binary image that represents ground and sky classes to test if the resulting classification conforms to the idea of a

Techset Composition Ltd, Salisbury Doc: {EL}Articles/Pagination/0547.3d
Image processing

straight-line horizon separating distinct sky and ground regions. *M2P5* is calculated using the formula:

$$M2P5 = \frac{(u_1 - u_2)^2 (n_1 n_2)^{1/3}}{3R^3} \quad (1)$$

where u_1 and u_2 are the x, y centroids for the two classes, n_1 and n_2 are the numbers of pixels in each class and R is the radius of the circular mask.

M2P5 is largest for images where the class centroids are furthest apart because of the $(u_1 - u_2)^2$ term, which increases as the centroids get further apart, and when the class populations are equal, due to the $(n_1 n_2)^{1/3}$ term, which is a maximum when the class populations are equal. (The $1/3$ exponent is to reduce its weight compared to the $(u_1 - u_2)^2$ term and was arrived at empirically.) The $3R^3$ term is used to normalise the metric to approximately 1. For a circularly masked image, this metric gives a maximum for an image where the sky and ground each occupy an equally sized separate hemi-circle. It discriminates against images where either of the classes contains significant pockets of the other class (because $(u_1 - u_2)^2$ decreases in this case), which would indicate a model failure. Applying a threshold to *M2P5* allows decisions to be made that can reject unsuitable processed images and the measurements from them.

Fig. 2 shows an example of original view and classified image. In the inset binary image, note the numerical value for the roll angle at top right and the value of *M2P5* (scaled to be in a range 0–255) at top left. The line joining the centroids is shown, as is the position of the calculated horizon. Arrows indicating directions of control corrections to maintain level flight indicate that the pilot should bank right and pitch down (move the pitch stick forward) to restore level flight. The arrows are not proportional to the error and are really a visual aid for development purposes. Also remember that the angle the horizon makes in the view is the opposite to the angle the aircraft makes to the horizon, so in this example the aircraft is banked too far to the left. In fact the pitch is only slightly too far up because the horizon position is only slightly lower than the centre.

Results: The horizon sensor with hardware and software to process the video and extract the measurements has been constructed using a small digital camera and a field programmable gate array (FPGA).

Fig. 3 shows the roll angle measured from the horizon video compared to the correct roll angle for a test conducted with a flight simulator. The horizon sensor's camera viewed the graphics output of the simulator and the calculations were conducted in real time over a period of 100 s in this case. Note the relatively low error of 3.5%. This test was conducted with clear skies over flat terrain. Similar tests with clouded skies and hilly terrain show similar results, though of course the differences between true roll angle and that calculated from the visible horizon were greater. It should be clear to the reader that this method will only produce angles that are strictly accurate when the visible horizon is truly horizontal. At other times, the roll and pitch angles derived by this method can nonetheless provide a useful reference.

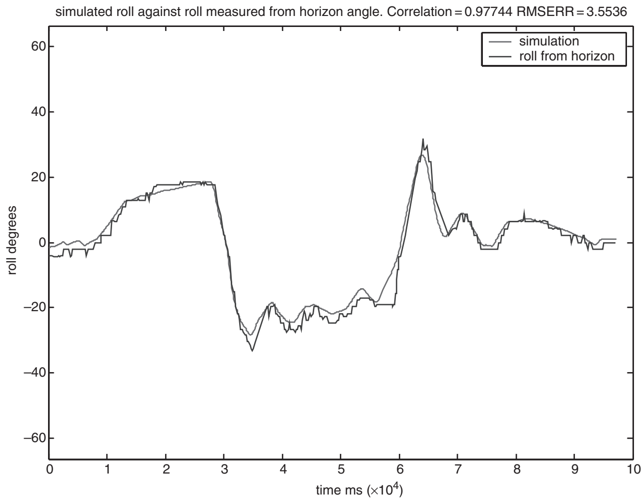

Fig. 3 *True roll angle compared to roll angle calculated from horizon*

When closed-loop tests were conducted with the flight simulator and the horizon derived angles were used as inputs to the control system to correct deviations from level flight, it was possible for the simulation to continue to fly without human intervention, sometimes for hours The flightpaths automatically follow terrain features such as valleys and saddles and avoid obstacles such as hills and ridges because the obstacles have the effect of making the horizon appear higher or more angled than it really is. In correcting the perceived angle error, an aircraft will change its pitch and/or roll in such a way that the aircraft climbs or rolls away from the obstacle. Many tests were conducted in simulation of this feature and the majority showed that the aircraft successfully avoided the terrain. There were some impacts when the aircraft did not have the speed to enable it to climb over the obstacles.

Despite the success of most aspects of this system, it has been observed that there remain some failure mechanisms. The main danger is from a situation where the true horizon has been lost and the system has no current attitude information against which to judge new measurements. If in those circumstances there appear images that sufficiently resemble horizons but which are not, then the system will be fooled. The most successful closed-loop trials that have been conducted to date use a combination of horizon sensor measurements as absolute references and angular rate measurements to give the expected change against which to judge the next horizon measurement.

Conclusions: A method for extracting aircraft attitude information from video has been derived and tested. The results show that it is a promising technique for flight stabilisation and terrain avoidance. A metric that measures the confidence value for the measurements has also been proposed.

00 00 xxxx
Electronics Letters online no: 2006xxxx
doi: 10.1049/el:2006xxxx

T.D. Cornall, G.K. Egan and A. Price (*Department of Electrical and Computer Systems Engineering, Monash University, Melbourne, Australia*)

E-mail: terry.cornall@eng.monash.edu.au

References

1 Taylor, B., Bil, C., Watkins, S., and Egan, G.: 'Horizon sensing attitude stabilisation: a VMC autopilot'. 18th Int. UAV Systems Conf., Bristol, UK, 2003
2 Chahl, J.S., Stange, G., Thakoor, S., Le Bouffant, N., Sriivasan, M.V., Hine, B., and Zornetzer, S.: 'Bioinspired engineering of exploration systems: a horizon sensor/attitude reference system based on the Dragonfly Ocelli for Mars Exploration applications', *J. Rob. Sys.*, 2003, **20**, (1), pp. 32–42
3 Brookman, A.K.: 'An albedo horizon sensor using hybrid circuitry'. RAE TR 71153 July 1971
4 Hatcher, N.M., and Germann, E.F.: 'Study of a proposed IR horizon scanner for use in space orientation control systems', NASA TN D-1005 Langley January 1962
5 Barrows, G., and Neely, C.: 'Mixed-mode VLSI optic flow sensors for in-flight control of a micro air vehicle', *Proc. SPIE*, 2000, **4109**, pp. 52–63
6 Todorovic, S., Nechyba, M.C., and Ifju, P.G.: 'Sky/ground modeling for autonomous MAV flight'. Proc. IEEE Int. Conf. on Robotics and Automation, May 2003, Tiawan
7 Otsu, N.: 'A threshold selection method from gray level histograms', *IEEE Trans. Syst. Man Cybern.*, 1979, **9**, pp. 62–66