In-Flight Orientation, Object Identification and Landing Support for an Unmanned Air Vehicle

Adam Hazeldene, Adam Sloan, Christopher Wilkin, Andrew Price
Department of Electrical and Computer Systems Engineering,
Monash University, Melbourne, Australia
Andrew.price@eng.monash.edu.au

Abstract

This paper examines three low cost, real-time passive vision systems for supporting the development of light weight, low cost Unmanned Air Vehicles as part of Monash University's Aerobotics project. The three systems add support to UAV operators at the critical times of landing a UAV, training to operate a UAV and gathering useful data about objects observed during a UAV mission. Each of the systems use computationally simple approaches and conventional vision capture hardware. This investigation was undertaken as part of an undergraduate thesis project.

Keywords: UAV, vision, landing, orientation, flight

1 Introduction

The Aerobotics group, part of the Department of Electrical and Computer Systems Engineering at Monash University has been in the process of developing low cost Unmanned Aerial Vehicles (UAV's) for a number of years. UAV's weighing up to 5kg are currently being constructed to perform missions such as search and rescue, surveillance and ground truth data acquisition. The airframes have been custom designed for specific tasks using low cost light weight materials and COTS components normally used for remote control aircraft. Vehicles have been constructed capable of travelling long distances overland (50km) or to altitudes in excess of 2000m. It's nice to have airframes that can complete such tasks, but it's even nicer to get them back safely. This paper considers the problems of a UAV in service, performing missions, and investigates what role real time vision systems can perform during a mission and at the critical moment of UAV retrieval. Three real time vision systems have been constructed: a ground based in-flight aircraft attitude determination system, a landing approach advisory system and an onboard natural/unnatural object discriminator. Each of the systems was constructed using low cost cameras and vision capture equipment in a Linux environment on conventional PC hardware [4]. Work was conducted as part of a final year undergraduate thesis project at Monash University. The vision systems that have been constructed are capable of determining the position of a UAV in 3D space as it approaches the ground to land, determining the attitude of a UAV in flight and to distinguish between objects that are likely to be of natural origin from those that are likely to be artificial. The information obtained from these systems is useful to the operators of a UAV in a number of ways including assisting the safe recovery of a UAV, monitoring in-flight characteristics and providing a useful on board instrumentation package.

2 Practical Implementation of three UAV Vision Systems

This investigation began by examining the current program of design and construction being undertaken by the Monash Aerobotics Group and determining what systems would compliment the existing programs. In determining what areas would be most beneficial to pursue, everything from the hardware to potential mission objectives, to the operators was considered. The operators of the Aerobotics group planes are all current model aircraft pilots with endorsements for electric aircraft. However it was observed that even the best pilots have a bad day. The first useful support system to be identified was a landing approach advisory system.

In considering the potential operators of UAV's further, the question of training new pilots was considered. For beginner pilots, determining the orientation of a plane in flight can be quite a challenge, in particular distinguishing whether or not a plane is flying toward an operator or away. To an operator on the ground, the controls may seem reversed, depending on the direction. This may lead an inexperienced pilot to make a critical error. A second useful system, that of in-flight attitude identification, was determined.

At present the Monash Aerobotics group is expending considerable effort in developing flight control and navigation systems. The ability to have an airframe fly a mission is only half of the problem however, if nothing useful can be achieved along the way. A third system, the ability to distinguish natural from unnatural objects in any given landscape was determined.

Each of the systems had to be passive and low cost, using simple, expendable hardware items.

2.1 Landing Approach Advisory System (LAAS)

Air Vehicles with high wing loadings are often capable of achieving very high speeds, in excess of 150km/hr. The most dangerous time in the working life of any air vehicle is most likely during take off and landing. Experienced pilots are typically used in such situations, to launch and recover the vehicle manually. Unfortunately there may not be an experienced pilot around when you need one, and even then, accidents still happen frequently. While the cost of an airframe for a small UAV may be comparable to a reasonably priced model airplane, the instrument package on board may be considerably more valuable. In any case, a damaged airframe could mean increased turnaround time in a critical emergency. Any aid that can improve the safety of landings is of value.

The LAAS uses two conventional CCD cameras that are positioned either side of the potential landing site. The two cameras are angled with respect to each other so that they both observe the same area. The angle is not critical as any discrepancy is corrected during the calibration process. Since LAAS may be used at any landing sight, it is necessary to calibrate for local conditions.

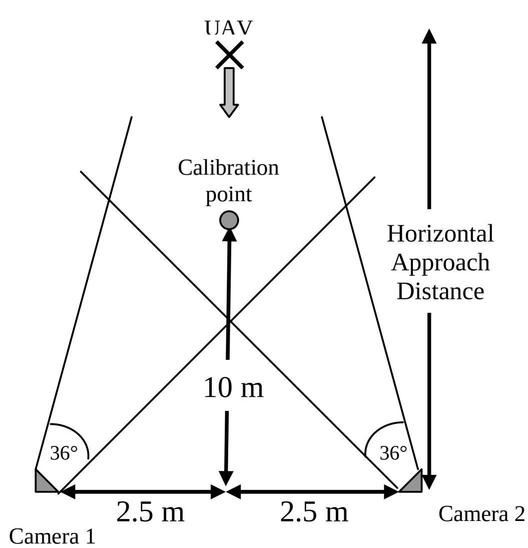

Figure 1. Layout of LAAS

The cameras are calibrated so that pixels within each image are identified with respect to a given distance, altitude and offset from a nominal centre line. At present, a landing area 20m long, 5m wide and 4 m high can be observed. It is essential that the LAAS be a real time observation. A number of computationally extensive algorithms were considered. Anything that limited the frame rate to less than 25 fps was discarded [1][2]. Figure 2 shows the Block diagram of the LAAS software.

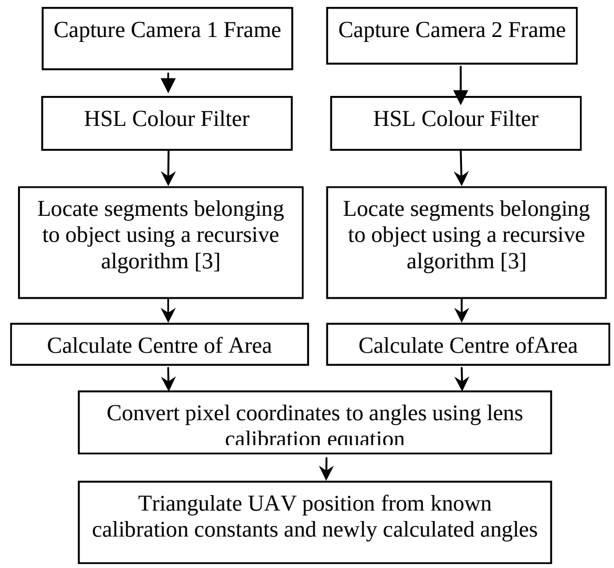

Figure 2. LAAS Software Block Diagram

The LAAS was tested by passing objects through the space. In particular, the red tail fin from one of the Aerobotics group UAV airframes was used. The following results were typical observations:

Range		**Altitude**(#)	
LAAS Measured Range (cm)	Actual Range(cm)	Actual Alt(cm)	LAAS Measured Alt (cm)
394	400	0	15
494	500	50	49
597	600	75	76
699	700	100	107
801	800	125	129
903	900	150	157
1000	1000	200	214
1098	1100	225	235
1213	1200	250	265
1306	1300	275	291
1403	1400	300	316
1506	1500	350	344
1614	1600		
1720	1700		
1813	1800		
1910	1900		
2043	2000		

Table 1.LAAS Results

The LAAS approach is sufficiently accurate to be of use as a landing assist system for small UAV airframes up to approximately 5kg flying weight. Note that the calculated reading for zero altitude takes into consideration the height of the tail fin above ground. The system is currently functional for a landing area of 20m x 5m x 4m, for larger and heavier UAVs with more unpredictable landing characteristics a larger zone would be required.

In order to be of genuine and practical use to a pilot attempting to land a UAV manually, information must be relayed without delay. Once the pilot is committed to retrieving his UAV, the landing process lasts a matter of seconds. For this reason an audible tone was considered more useful than a voice speaking the altitude. The tone is frequency modulated reflecting the proximity of the ground once the aircraft is in the zone. The LAAS approach requires two cameras to be accurately placed and calibrated prior to use. Accuracy of altitude and range measurements depends on this, however even an un-calibrated tone can be of benefit in the landing process, providing feedback to the operator regarding control movements. The system is portable and can be applied to any landing site.

2.2 In Flight Attitude Determination

Learning to fly a small, high-powered air vehicle can be a difficult task, particularly for non-model flyers. In dark or cloudy sky, airframes tend to become black silhouettes. A common effect is the appearance of an airframe flying towards an observer, when it is in fact flying away. It may not be possible for an observer on the ground to correct the attitude of an airframe once it gets out of shape. The In Flight Attitude Determination system has benefits not only for the Monash UAV program but also for pilots training to fly model aeroplanes. Even the so-called ‘Buddy-box’ system of two radio transmitters linked together, with an experienced pilot on one and a novice on the other can break down if the experienced pilot is distracted or in any way becomes complacent.

The In Flight Attitude Determination System [1] can be used to determine in real time the track of a plane flying at distances up to 70m in daylight. This is adequate for practice flights and testing. It is a real time system, and can alert a pilot when an undesirable attitude is achieved. At present it has been designed as an open loop system however there is significant potential to be able to override the pilot controls and initiate a recovery sequence to protect the hardware. The system shown in Figure 3 uses a conventional Tri-Pod mounted video camera with zoom function and can operate on any Linux [4] based computer system capable of capturing live video.

Figure 3 Attitude Determination Hardware Setup

Figure 4 Shows the Block Diagram of the software used to determine aircraft attitude.

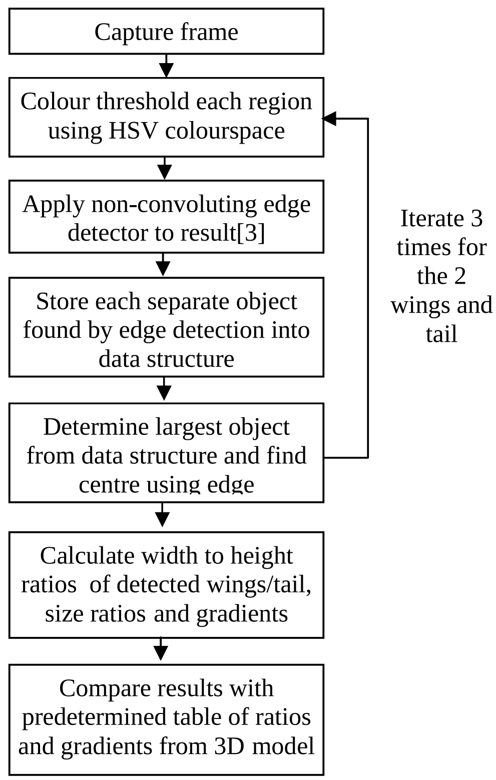

Figure 4 Attitude Determination Software Model

The Attitude Determination System uses existing markings on the airframe as a basis for determining the attitude. The system is adaptable, however the further the UAV is from the camera, the less information that is available. Shadows, silhouettes and reflection can cause misreads, particularly at long distances.

Figure 5 shows the Attitude Determination System during testing of the algorithm using a model of an airframe under real lighting conditions. The wings and tail of the airframe are encoded with red, green and purple markings. Provided no other object enters the captured image that is larger in pixel content than the actual UAV the attitude determination system will give a reasonable response. Underneath the captured image is the computer's own interpretation of the attitude. While it is still possible to 'trick' the attitude determination system in individual frames due to the problems mentioned previously, time averaging reduces the effects of outliers in the calculated data.

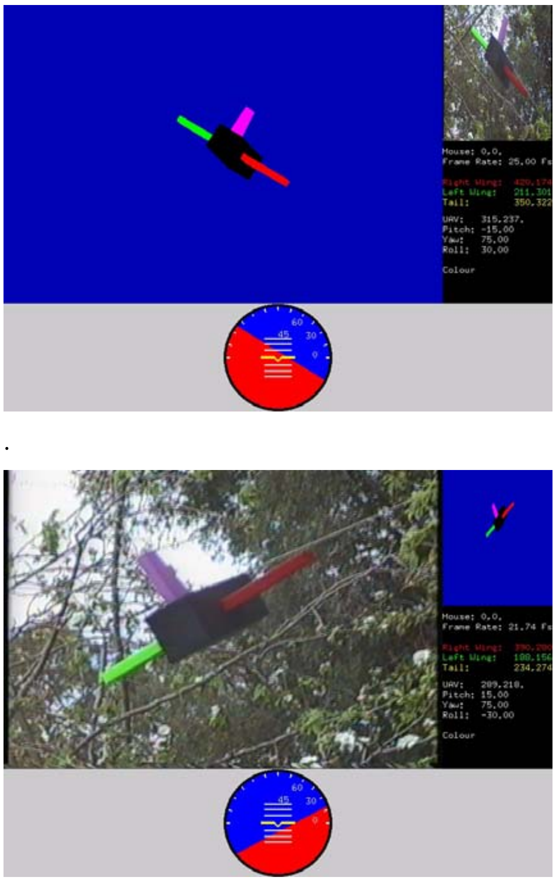

Figure 5 Testing the In-Flight Attitude Determination System with model airframe in real lighting conditions.

Present field tests have accurately determined the attitude of a small UAV, with a 20cm wingspan to within 10 degrees in each of the three planes (roll, pitch, yaw) Accuracy diminishes with increased separation. The use of a Zoom lens can extend the effective range. No special hardware is needed, as the determination algorithm begins with pixel data in the HSV space. Objects within the captured image can be the same color as the predefined markers, but not larger in terms of pixels than the identifying markers. The software is robust enough to reject objects that are the same color as the markers but are smaller.

Outside testing is not as robust as inside testing in more ideal conditions. Specular reflections and the use of color as a determining factor are known limitations. The use of color also dramatically impinges the performance in low light conditions such as evening or night. For this problem it is proposed to replace the colors with high intensity LEDS encoded so as to make it possible to identify each of the critical surfaces individually (left and right wings, tail).

2.3 Natural and Unnatural Object Determination

Many mission scenarios that a UAV may undertake involve the determination of an unnatural object in a background of natural objects. These objects may be of similar color, or objects of interest may simply be swamped by the detail of the surrounding landscape. Ships and rocks at sea as well as road vehicles in bush land are examples of the types of information needed to be obtained by a UAV. In such situations color alone is insufficient to determine the existence of an unnatural object. Dark rocks and dark ships all look the same, in color terms. It would be fairly pointless having a UAV successfully seeking out a large ship-like slab of granite.

Our approach to Natural and Unnatural Object Determination involves a recursive analysis and labelling of adjacent pixels to define interesting objects in the frame. A method of determining straight, possibly parallel line segments from the interesting objects is used to identify the unnatural objects in the image. Figure 6 shows actual results from a UAV experiment highlighting unnatural objects such as cars and buildings.

Figure 6. Actual footage from a UAV with the possible unnatural objects detected

While natural objects do have regular edges, straight lines and sometimes parallel lines, man made objects can have a significantly higher proportion of both straight lines and lines that are parallel, as well as other features such as symmetry of lines. By using these methods, a reasonable assumption can be made about the origin of an object of interest being observed from a UAV.

Given the limited time frame of a current small scale UAV mission, any time wasted is undesirable. By the same token, payload capacity dramatically limits the onboard computational ability. Any approach must be economical in terms of time and power consumption.

Figure 7. The Straight Lines Detected

Figure 7 shows a captured frame from a UAV mission showing the highest concentrations of straight lines. Figure 8 shows the approach we are developing to classify the dataset. Several algorithms have been employed, including a computationally inexpensive edge detector, a recursive segmentation algorithm that defines possible objected and a trigonometric line detector.. Objects that may be of interest due to their grayscale representative color are extracted using low computationally intensive algorithms [2]. A recursive algorithm [3] is then used to segment edge pixels into lines. By means of thresholding the segmented data, the relative density of straight lines and parallel lines is determined.

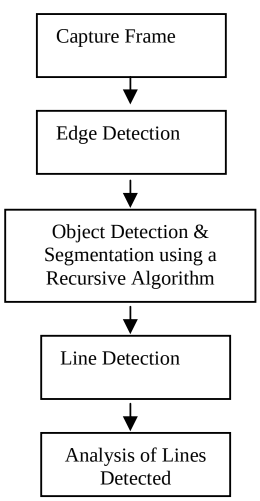

Figure 8. Natural/Unnatural object discrimination.

The hypothesis that natural objects have fewer parallel straight lines, regular shapes and unnatural colors is the basis for this method. Figure 9 shows the block diagram of the process.

After the image is passed through the edge detection algorithm, the output will resemble possible object edges plus a certain degree of noise. This noise will be dependant on the threshold value chosen for the

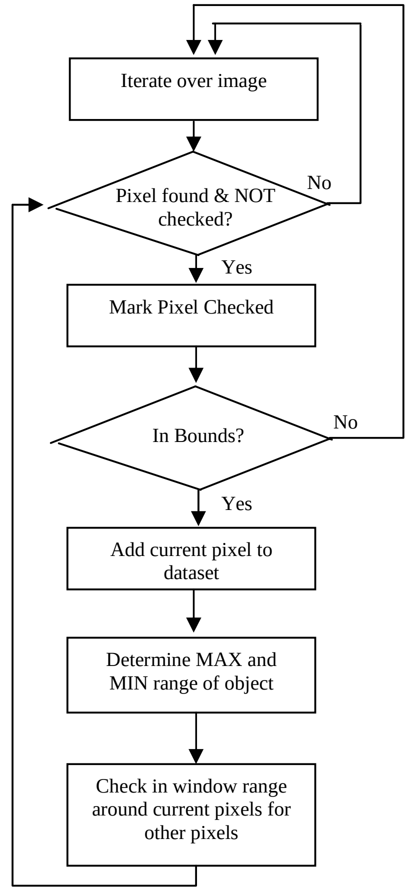

Figure 9. Recursive search algorithm for natural/unnatural object detection

edge detector. This output is then passed into the object detection algorithm, which segments regions of the frame where it believes an object is present.

It uses a recursive algorithm to traverse around the adjacent connected pixels within a certain specified window of search. A data structure was setup for the storage of interest pixels that may be considered objects. Within this data structure were variables such as the coordinates used by the recursive algorithm, the size of the object and the maximum and minimum x and y values of the object.

Before the recursive algorithm is called, the capture frame is iterated upon, setting all pixels that it comes across as *'checked'* until it finds a pixel that is of interest. The pixel location is passed into the recursive algorithm, where a specified size window is passed over the surrounding pixels in the frame to determine whether adjacent interest pixels are present. If the adjacent pixels are also of interest, and they haven't already been processed or checked in this current frame, then the recursive algorithm is called again on this new interest pixel and it is set to '*checked*' also.

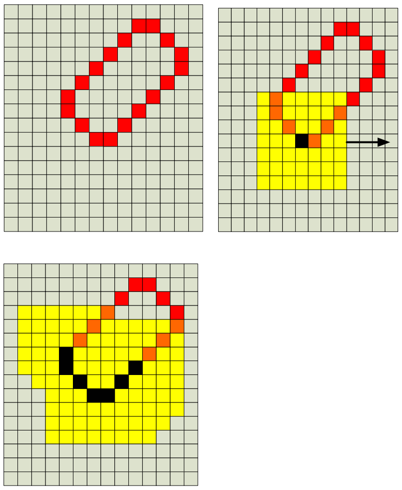

Figure 10. The Algorithm running in consecutive recursions

Figure 10 demonstrates the object detection algorithm running in consecutive recursions. The first image shows the single pixel object outline, while the second shows the window of search that takes place once an initial interest pixel is found and the third shows the next consecutive search that has taken place.
The black pixels represent the already checked pixels, the orange pixels represent the pixels that are both part of the object and within the current search window/s. Therefore, these pixels will be passed into the recursion function next.
Finally, the yellow pixels represent both the current search and already searched pixel windows.

After the object extraction takes place, and the minium and maximum values for both x and y are known, each possible object is dynamically assigned to a smaller window frame. The size of the smaller frame is proportional to the object size, which is known from the minimum and maximum values of x and y. This ensures that any further processing that is done, occurs significantly faster, unless a considerable sized object is found.
This optimisation assumes the fact that smaller sized object windows are a lot faster to search when compared to the full sized capture window.

Each of the individual segmented objects were stored into a data structure, which held an array of each object's properties. As well as this, each segmented object was also stored into another data structure which held the total number of interest pixels in the whole frame. The maximum number of objects that could be found was assumed to be 50.

Figure 11 shows the results of a frame containing 3 cars and vegetation, processed to reveal the positions of only the 3 cars.

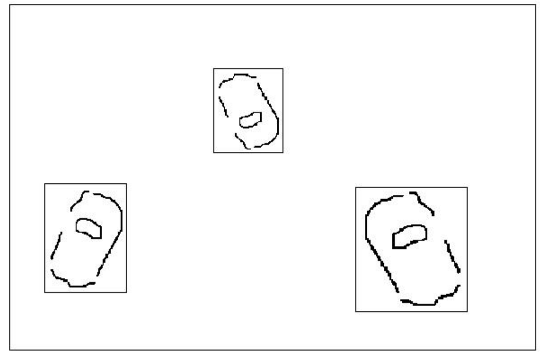

Figure 11. The Processed Frame with the detected unnatural objects highlighted in boxes

3 Conclusions

Small scale, light weight and expendable UAV platforms have a very wide potential in surveillance, search and rescue and data gathering applications. This paper presents three vision based systems, using low computation approaches and low cost image capture hardware to determine the attitude of an airplane in flight, the nature of an object of interest on the ground and most significantly, the position of a UAV as it comes in to land. Of particular consideration in this investigation was the use of economical vision capture hardware. This has lead to considerable problems in terms of accuracy due to lens distortion, density of pixel information, image clarity and speed. The nature of the research platform however, that of small low cost and expendable UAV systems that may be employed in harsh or hostile environments, suggests that an economical approach is advisable. UAV airframe attitude has been determined within 10 degrees at the time of writing. Position of a UAV on approach has been determined within 10 centimetres over a landing field 20m x 5m x 4m. The process of determining natural and unnatural objects has been successfully tested on video images obtained from real UAV's in the field. Buildings and vehicles have been successfully detected in clusters of trees.

4 References

[1] Forsyth, D. Ponce, J.., Computer Vision, A Modern Approach, *Pearson Educational* (2003).

[2] Du Quan Huynh, Feature based stereo vision on a mobile platform, PhD Thesis, University of Western Australia, (1994).

[3] Price, A.R. Mobility and Vision for Mobile Robots in Non-Deterministic, Competitive Environments, PhD Thesis, Deakin University, (1999).

[4] Video4Linux Web Resources and Reference http://linux.bytesex.org/v4l2