J Intell Robot Syst (2012) 68:339–357
DOI 10.1007/s10846-012-9684-7

Horizon Profile Detection for Attitude Determination

Steven J. Dumble · Peter W. Gibbens

Received: 22 September 2011 / Accepted: 31 May 2012 / Published online: 13 June 2012

Abstract The horizon appearance is a strong visual indication of the attitude of an aircraft, so a vision based system should be able to detect the horizon and use its appearance to extract attitude measurements. Past methods have made the assumption that the horizon is straight, this neglects possible navigational and attitude information. This paper outlines a horizon detection method which allows for the actual horizon profile shape to be extracted. This horizon profile is then used for visual attitude determination. Test results for a captured flight video are presented and the proposed method is compared and evaluated against other methods.

Keywords Vision systems · Attitude determination · Image processing · Horizon detection

S. J. Dumble (✉) · P. W. Gibbens
University of Sydney, Sydney, NSW, Australia
e-mail: steven.dumble@sydney.edu.au

P. W. Gibbens
e-mail: peter.gibbens@sydney.edu.au

1 Introduction

Horizon detection in vision systems is not a new concept; there has been a large amount of past research into image processing techniques to perform horizon detection and various methods for fitting a horizon to the image. The majority of the past work on horizon detection has assumed that the horizon will be a straight line but there will be cases when this will not hold true. When a straight line is fitted to a detected horizon, non-uniformities in the horizon (such as mountains and other terrain features) can manifest themselves as biases in the fitted horizon line's slope and offset. This allows for an attitude measurement to be made, but it can be very susceptible to error when the actual horizon is not straight. Usually horizon detection will only give a bank and pitch attitude estimate [2–4, 6–10, 13, 25] due to the horizon fit being reduced to a straight line; however detecting the horizon profile allows for full attitude estimation (bank, pitch and yaw) [14] as well as estimation of navigational information [23, 24] using terrain-aided methods. A large amount of the past work has been focused on MAV (Micro Air Vehicle) or UAV (Unmanned Aerial Vehicle) attitude stabilization [2–4, 8–10, 13, 25], and only a limited number have focused on recovering the explicit attitude from the estimated horizon line [6, 7, 18]. The explicit attitude information however is of great importance in control, guidance and navigation.

The problem of horizon detection has been approached many different ways in the past. The past literature has focused on three main methods for horizon detection which are; edge based [6, 7] (horizon edge points are found then a line is fitted to them), pixel segmentation based [2, 4, 20, 25] (sky and ground are segmented first, then a horizon line is fitted) and statistical segmentation based [3, 8–10, 13] (sky and ground segmentation and horizon fitting processes occur simultaneously).

Edge based horizon detection finds points in the image with a strong image intensity gradient. It is assumed that these detected edge points in the image belong to the horizon, so a straight line is fitted to them, usually using the Hough transform [6, 7]. This method relies on the assumptions that there is a distinct straight edge between the sky and the ground and that this line is the strongest one returned by the Hough transform. When these assumptions are violated, the Hough transform can return erroneous measurements. One of the advantages of the Hough transform is its noise rejection ability when globally fitting a straight line. However this advantage quickly breaks down when the detected edge points do not lie on a straight line. A disadvantage of the Hough transform is that it can be computationally expensive. The edge detection and line fitting stages in the horizon detection process have different parameters and thresholds which need to be tuned. The tuning of these parameters affects the overall accuracy and robustness of the horizon detection process.

Segmentation based approaches first segment the image into sky and ground, then a horizon line is fitted to separate the segmented regions. The segmentation process can be based on a number of different approaches such as a color threshold [2], texture threshold or machine learning classifiers [20, 21]. A major drawback to color based sky detection methods are their inability to distinguish between sky and other sky colored and textured objects such as a body of water or haze. An advantage of the segmentation approach is the individual pixel sky/ground classification, this information allows the actual shape of the horizon to be extracted.

Statistical segmentation approaches combine the sky and ground pixel classification and horizon fitting stages together in one step which is usually accomplished by a maximization or minimization expectation search on some fit criteria. An optimization algorithm was used in [8–10] to minimize the RGB (red, green and blue) interclass variance of two segments of an image separated by a straight horizon line. Whereas [13] used the Otsu threshold method to statistically separate the background and foreground and [3] used a k-means clustering approach to separate the image into 2 regions. Statistical segmentation approaches have the advantage of being an optimization procedure, so that there are no issues with the tuning of detection parameters. However its disadvantages are its computational expense, the need for the sky and ground appearance to be significantly different and lastly the fact that the two separated regions are not guaranteed to correspond to the sky and the ground.

This paper develops a horizon detection and extraction method which differs from previous approaches [2–4, 6–10, 13, 25] by allowing the horizon to take a general shape, rather than constraining it to be a straight line. The proposed method can handle some of the situations which can cause difficulties for other methods, such as a lack of a distinct horizon edge or a lack of a significant appearance difference between the sky and the ground. The proposed method also has a higher processing rate and a higher attitude determination accuracy when compared to the two other methods investigated.

The layout of this paper is broken up into a number of sections. Section 2 investigates the physical appearance of the horizon and develops the proposed horizon detection method from these observations. Section 3 outlines the attitude determination process, which links the detected horizon profile to the attitude of the aircraft. Bank and pitch attitude equations are developed which link the fitted horizon line parameters in the image, directly to the attitude of the aircraft. Implementation and flight test results for this proposed horizon detection method and two other past approaches are presented in Section 4, while their corresponding discussions are presented in Section 5. Finally, an extension to the proposed method which can compensate for some of the issues identified in the discussion section is presented in Section 6.

2 Horizon Detection

2.1 Horizon Appearance

The physical manifestations of the appearance of the sky (specifically color, texture and spatial properties) were investigated for the purpose of the development of a horizon detection algorithm. The clear sky color is predominantly driven by the scattering of the sun's light through the earth's atmosphere. As the amount of scattering is related to its wavelength by Rayleigh's law, the blue wavelength is scattered the most, followed by green and then the red wavelengths, giving the sky an overall blue appearance [1, 17]. The shade of the color of the sky is not uniform. The atmosphere appears the thickest at the horizon which causes the sky color to saturate towards white, as the individual color component intensities increase [16]. These relationships are observable in the two vertical color component traces, shown in Figs. 1 and 2. The image traces also contain other valuable information.

The trace shown in Fig. 1 has a clearly defined horizon interface. The top of the mountain against the sky defines a very identifiable horizon edge (all three color channels have a very large gradient). This edge is easily picked in horizon edge based methods, such as [6, 7]. The color distribution in the ground region is also dramatically different compared to the sky region. In the sky region, the distribution follows Rayleigh's law whereas in the ground region, the blue component is not always the most dominant color. This clear color distribution difference allows segmentation and statistical segmentation horizon detection approaches, such as [2–4, 8–10, 13, 20] to easily identify the horizon.

There are noticeable situational differences in Fig. 2. The horizon interface no longer has a well defined edge; it only has a weak gradient or transition response. Edge based methods would be unable to detect a usable edge in these situations. The color distribution between the sky and the ground is very similar. This is due to the

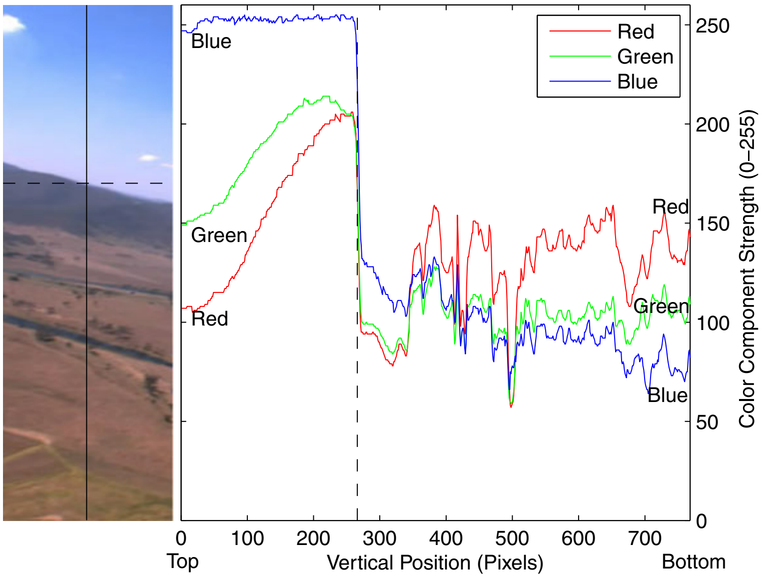

Fig. 1 Vertical color image trace (Strong Horizon)

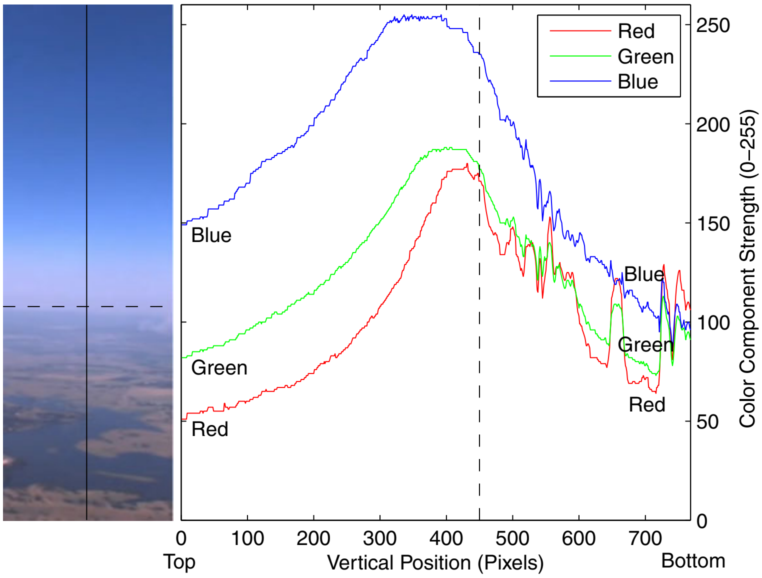

Fig. 2 Vertical color image trace (Weak Horizon)

higher altitude of the camera, which increases the distance through the atmosphere that the light has to travel from the ground. This increased distance causes the Rayleigh's law to become noticeable in ground regions, washing out the color distribution. This lack of difference between the sky and ground regions will cause statistical segmentation approaches to fail as well.

In all cases, the horizon interface is spatially correlated across the image. It is not necessary a straight line, but it should form a continuous boundary across the image.

2.2 Proposed Horizon Detection Method

The observations in Section 2.1 are used to develop the proposed horizon detection and profile extraction method described in this section. This allows the proposed method to handle some of situations that can cause difficulties for other methods, such as the lack of a distinct horizon edge or a lack of a distinct colour distribution difference between the sky and the ground. The main image feature that the proposed method uses for horizon interface detection is the image gradient. Negative gradients (light to dark) are used to define possible areas for the horizon interface. The use of gradient information allows for possible horizon detection in cases when the horizon does not return a strong edge response. The local color distribution is then used to remove gradients which are unlikely to be caused by the horizon. Spatial constraints are then used to select a continuous horizon profile from the image.

A complete outline of the proposed horizon detection method is shown in Fig. 3. The first step in the process is to smooth the individual color channels of the image to remove very high frequency content and noise. A smoothing window of size n is used in a box averaging filter to smooth the image. Let the original color image be $\boldsymbol{I}(u,v) = [\boldsymbol{I}_R(u,v), \boldsymbol{I}_G(u,v), \boldsymbol{I}_B(u,v)]$ where (u,v) are the image coordinates in pixels and $\boldsymbol{I}_R, \boldsymbol{I}_G, \boldsymbol{I}_B$ are the individual color channels corresponding to red, green and blue respectively. The different smoothed color channels are calculated using:

$$\bar{\boldsymbol{I}}_c(u,v) = \frac{1}{n^2} \sum_{i=-\frac{n}{2}}^{n/2} \sum_{j=-\frac{n}{2}}^{n/2} \boldsymbol{I}_c(u+i, v+j) \qquad (1)$$

where $\bar{\boldsymbol{I}}_c$ is the smoothed image color channel calculated for the different color channels $c \in \{R, G, B\}$.

The next step is the estimation of the vertical gradient $S(u,v)$ from the smoothed color channel images. While the horizon interface is correlated

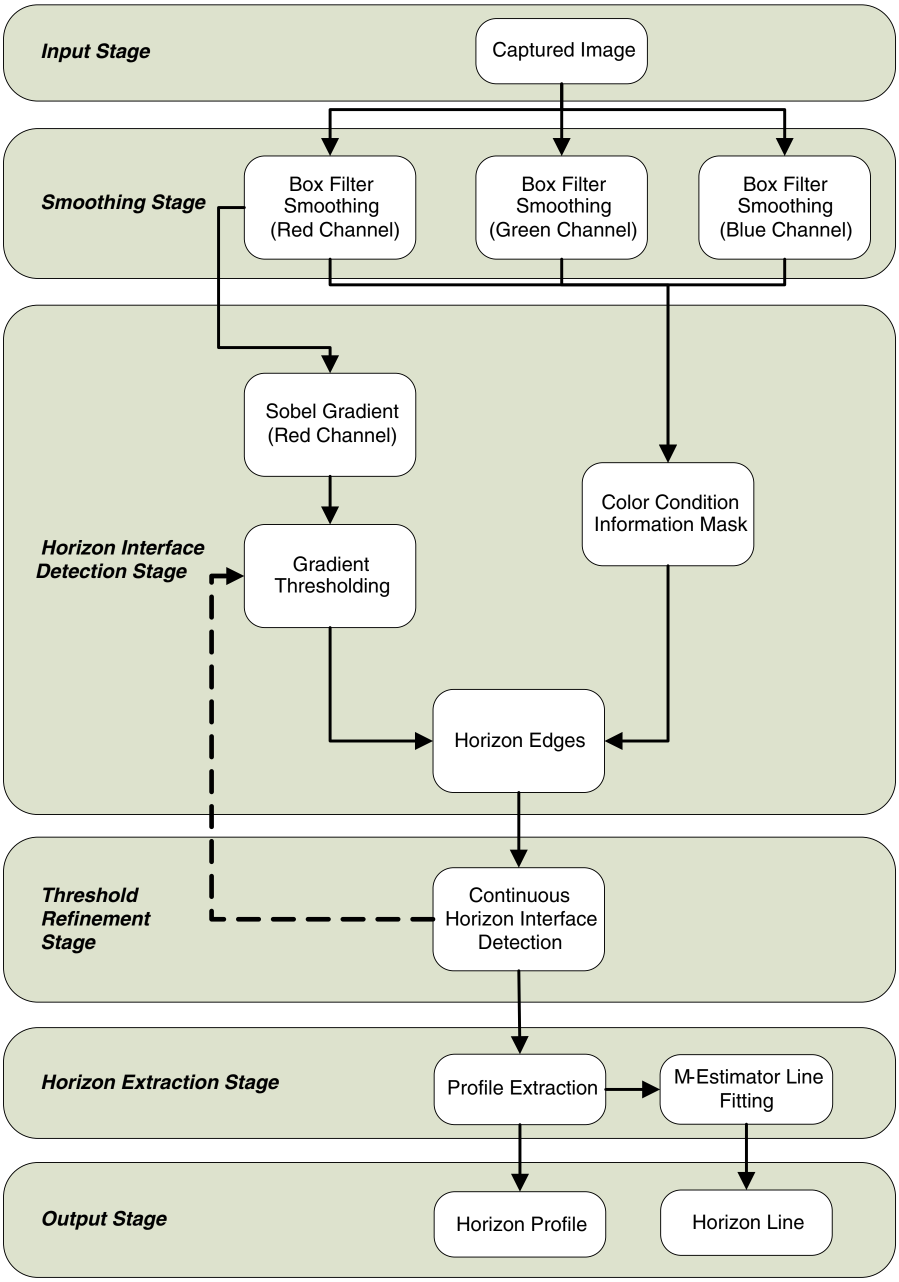

Fig. 3 Proposed horizon detection outline

with all three color channels as seen in the color traces in Figs. 1 and 2, only the red channel vertical gradient is calculated. The red channel gradient is used instead of the other channel gradients (or a combination of them) as it appears to be the most sensitive to the sky/ground transition while being less sensitive to other objects (this was also noted in [6]).

The vertical gradient $\boldsymbol{S}(u, v)$ can be calculated from the convolution of the smoothed red channel

$\bar{\boldsymbol{I}}_R$ with a 3×3 vertical[1] Sobel mask $\boldsymbol{H}$, such that $\boldsymbol{S}(u, v) = \boldsymbol{H}(u, v) * \bar{\boldsymbol{I}}_R(u, v)$ where the convolution operation is:

$$\boldsymbol{S}(u, v) = \sum_i \sum_j \boldsymbol{H}(i, j)\bar{\boldsymbol{I}}_R(u - i, v - j) \tag{2}$$

with the vertical Sobel mask [5] defined as:

$$\boldsymbol{H} = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \tag{3}$$

After the gradient calculation, possible horizon transitions or gradients are identified by thresholding the gradient image $\boldsymbol{S}(u, v)$ to form a binary image $\boldsymbol{B}_{\text{gradient}}(u, v)$ with a threshold T_s using:

$$\boldsymbol{B}_{\text{gradient}}(u, v) = \boldsymbol{S}(u, v) > T_s \tag{4}$$

Only negative gradient areas (light to dark, down an image, i.e. sky to ground) are selected. This helps to remove gradients which could not be caused by a sky/ground transition in the binary image. The color image traces shown in Figs. 1 and 2 show examples of a strong edge response horizon and weak gradient response horizon. Applying a threshold to the vertical gradient allows both cases to be detected, compared to an edge based detector would miss the horizon with a weak gradient. The selection of the threshold value T_s is an important part of the detection process and it is discussed in Section 2.3.

The gradients selected in the binary image $\boldsymbol{B}_{\text{gradient}}(u, v)$ could belong to any number of image features. Additional information is used to try to cull any gradients which are not caused by a sky/ground transition. The color distribution can provide strong information about whether a pixel belongs to the sky or the ground. Before the ground region appears washed out due to altitude effects, it can be noted that the ground region typically has a stronger green or red component when compared to the blue component. This is a strong condition that is normally only satisfied in the ground regions of the image (or in the sky due to clouds or glare, when the light intensity is very bright and partially or fully saturates the color channels). The inverse of this condition is true, the blue component is usually the strongest color in the sky due to Rayleigh's law.

A binary image is formed where $\boldsymbol{B}_{\text{color}}(u, v)$ represents whether or not that particular pixel satisfies the binary condition of:

$$\begin{aligned} \boldsymbol{B}_{\text{color}}(u, v) = &\left(\bar{\boldsymbol{I}}_G(u, v) > \bar{\boldsymbol{I}}_B(u, v)\right) \\ &\vee \left(\bar{\boldsymbol{I}}_R(u, v) > \bar{\boldsymbol{I}}_B(u, v)\right) \end{aligned} \tag{5}$$

where $\vee$ is the logical disjunction (OR) operation. Using the distribution rather than a threshold for this condition, helps the condition to be invariant to various lighting conditions.

This color condition is used to remove any of the possible gradients identified by $\boldsymbol{B}_{\text{gradient}}(u, v)$ which could not be caused by a sky/ground transition, such as ground/ground transitions (i.e. roads) or sky/sky transitions (i.e. clouds). This forms the binary image $\boldsymbol{B}_{\text{edges}}$ using the logical operation:

$$\boldsymbol{B}_{\text{edges}}(u, v) = \boldsymbol{B}_{\text{gradient}}(u, v) \wedge \neg \boldsymbol{B}_{\text{color}}(u, v) \tag{6}$$

where $\wedge$ is the logical conjunction (AND) and $\neg$ is the negation (NOT) operation.

The detection steps so far have used color and gradient information. The last step in the process is to use spatial information to select the horizon. This step is usually where other methods [2–4, 6–10, 13, 25] make the assumption (in some form) that the horizon is straight. Instead, the proposed method assumes that the horizon should be spatially correlated with each column of the image by implying that it should be continuous from one border of the image to another. In normal attitude operating conditions, this means that the horizon should pass through the left and right borders of the image.[2] This condition is used to

---

[1] Only the vertical mask is used as it was assumed that it should be roughly normal to the horizon. This assumption is justified in most flight conditions for a civilian fixed-wing aircraft. If extreme attitude ranges where expected, then the Sobel mask direction could be updated in 45° increments based upon the current estimated attitude. This was not necessary for any of the results presented in this paper.

[2] The assumption that the horizon passes through the left and right borders of the image is modified in Section 6 for improved detection at non-nominal attitudes when this condition would not be true.

filter the possible gradients to select one that fulfills this criterion. The binary image $\boldsymbol{B}_{\text{edges}}$ is morphologically reconstructed with dilation [22] to filter out any gradients which do not span the image. This forms the horizon binary edge image $\boldsymbol{B}_{\text{horizon}}(u, v)$. The horizon profile is then extracted from this remaining horizon edge.

To preform the morphological reconstruction, two binary images, a left binary image $\boldsymbol{B}_{\text{left}}$ and right binary image $\boldsymbol{B}_{\text{right}}$ are required. The goal of the process is to select the gradient which connects both binary images together. The left and right binary images are simply generated with:

$$\boldsymbol{B}_{\text{left}}(u, v) = \begin{cases} \text{true} & \text{if } u = 0 \\ \text{false} & \text{otherwise} \end{cases} \tag{7}$$

$$\boldsymbol{B}_{\text{right}}(u, v) = \begin{cases} \text{true} & \text{if } u = \text{width} \\ \text{false} & \text{otherwise} \end{cases} \tag{8}$$

The left binary image $\boldsymbol{B}_{\text{left}}$ is morphologically reconstructed with the edge binary image $\boldsymbol{B}_{\text{edges}}$ to select all gradients which are connected to the left binary image. This forms the left connected horizon edge $\boldsymbol{B}_{\text{left,horizon}}$ binary image. The left morphological process iteratively dilates the left connected horizon edge $\boldsymbol{B}_{\text{left,horizon}}$ binary image until $\boldsymbol{B}_{\text{k-1,left,horizon}} = \boldsymbol{B}_{\text{k,left,horizon}}$ using:

$$\boldsymbol{B}_{\text{1,left,horizon}} = \boldsymbol{B}_{\text{left}} \tag{9}$$

$$\boldsymbol{B}_{\text{k+1,left,horizon}} = (\boldsymbol{B}_{\text{k,left,horizon}} \oplus 1_{3\times3}) \cap \boldsymbol{B}_{\text{edges}} \tag{10}$$

where $\oplus$ is the morphological dilation operation and $\cap$ is the intersection. The same process is done with the right side to form the right connected horizon edge binary image $\boldsymbol{B}_{\text{right,horizon}}$. The right connected horizon edge binary image is iteratively dilated until $\boldsymbol{B}_{\text{k-1,right,horizon}} = \boldsymbol{B}_{\text{k,right,horizon}}$ using:

$$\boldsymbol{B}_{\text{1,right,horizon}} = \boldsymbol{B}_{\text{right}} \tag{11}$$

$$\boldsymbol{B}_{\text{k+1,right,horizon}} = (\boldsymbol{B}_{\text{k,right,horizon}} \oplus 1_{3\times3}) \cap \boldsymbol{B}_{\text{edges}} \tag{12}$$

The left and right reconstructed images are combined to only select the edge which is connected to both borders to form the horizon edge binary image $\boldsymbol{B}_{\text{horizon}}(u, v)$ by the operation:

$$\boldsymbol{B}_{\text{horizon}}(u, v) = \boldsymbol{B}_{\text{left,horizon}} \wedge \boldsymbol{B}_{\text{right,horizon}} \tag{13}$$

If there is no gradient which fulfills this condition (i.e. the threshold is too strong) then the horizon edge binary image $\boldsymbol{B}_{\text{horizon}}$ would be empty (all pixels are false). However, if any gradients are returned, then it is probable that it belongs to the horizon.

The horizon profile is extracted from the binary horizon image $\boldsymbol{B}_{\text{horizon}}(u, v)$, by tracing along the image left to right $u = (0 \rightarrow \text{width})$, extracting the top most edge i for each column, such that $(\boldsymbol{B}_{\text{horizon}}(u, i) \neq \boldsymbol{B}_{\text{horizon}}(u, i + 1))$. This forms the one dimensional horizon profile[3] $v = H_p(u)$.

2.3 Threshold Selection

The threshold value T_s plays a large part in the detection process. Correct selection of T_s is an important process which affects the accuracy and robustness of the extracted horizon profile.

There are two methods for the selection of the threshold T_s value. The first method is to set T_s to a fixed value, one which is low (or sensitive) enough to ensure horizon detection. An advantage of this fixed threshold method is its ability to select a horizon with a weak response over a false horizon with a stronger response. This situation might occur when there is a strong edge on the ground which spans the image. A fixed threshold however has a large disadvantage, the low value required can make the whole detection process very sensitive to noise (a noisy profile can still be used for attitude determination when the horizon profile is approximated by a straight line, however it makes robust profile extraction difficult).

The second method for the selection of T_s is to use an adaptive threshold. The proposed method has a clear indication of whether a threshold value is too high (no horizon is found after the

---

[3] Again this assumes that the horizon profile is normally horizontal across the image otherwise shape information is lost.

morphological reconstruction stage). This can be used as an adaptive criterion; the adaptive threshold should select the largest threshold value in which a horizon is detected. The assumption here is that a more sensitive threshold (lower) will only allow a greater sensitivity to noise. If no prior information is known, then the threshold value can be set high and iteratively reduced until a horizon is detected. If a previous threshold value is known, then this value can become the starting point for the iterative search. If the previous value is sensitive enough to detect a horizon in the first iteration, then the threshold is iteratively increased until the largest threshold value that results in a detected horizon is found. Otherwise the threshold is iteratively reduced until a horizon is found. Due to computational performance constraints, a limit N on the maximum number of threshold iterations performed each frame can be put in place. An outline of the adaptive algorithm is shown in Algorithm 1.

Algorithm 1 Adaptive threshold procedure

```
Data: Previous Adaptive Threshold Value and Iteration Limit
Result: Current Adaptive Threshold Value
begin
    for $i = 1 \to N$ do
        Horizon Detection Using Current Threshold
        if Horizon Found then
            if Threshold Increased from Last Iteration then
                return Threshold From Current Iteration
            else
                Decrease Threshold
            end
        else
            if Threshold Decreased from Last Iteration then
                return Threshold From Last Iteration
            else
                Increase Threshold
            end
        end
    end
    return No Threshold Found
end
```

2.4 Examples

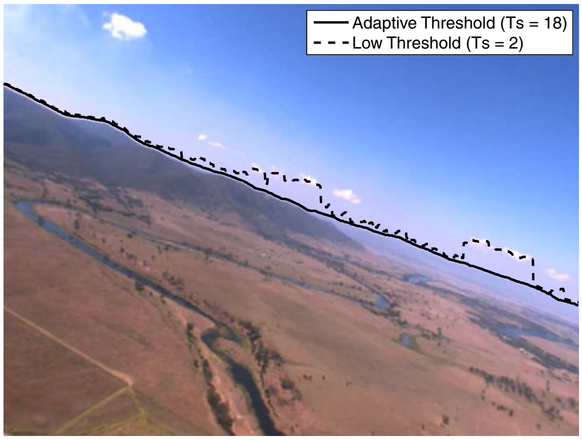

(a) Detected Horizon

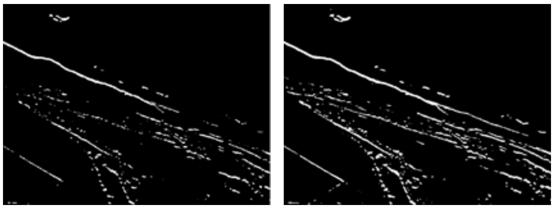

(b) $T_s = 40$ Threshold Image (c) $T_s = 30$ Threshold Image

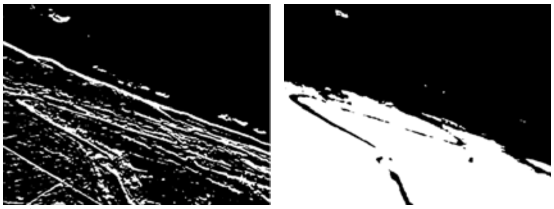

(d) $T_s = 10$ Threshold Image (e) Color Condition Image

Fig. 4 Example image 1

Two examples of the horizon detection process are shown in Figs. 4 and 5. Each example shows the detected horizon using an adaptive threshold and a fixed threshold as explained in Section 2.3. The adaptive threshold produces the best extracted horizon profile, while the fixed low threshold is very sensitive to noise and obstructions.

A number of different $\boldsymbol{B}_{\text{gradient}}$ binary images with a reducing threshold are shown in Figs. 4b–d and 5b–d. As the threshold lowers, the amount of noise in the binary image increases. The adaptive threshold process tries to find the largest threshold value which causes a detected gradient to span the image. In the first example, Fig. 4e shows the $\boldsymbol{B}_{\text{color}}$ binary image, which is used to remove unwanted gradients from $\boldsymbol{B}_{\text{gradient}}$. It can be seen in this example that the color condition would remove a lot of the unwanted gradients

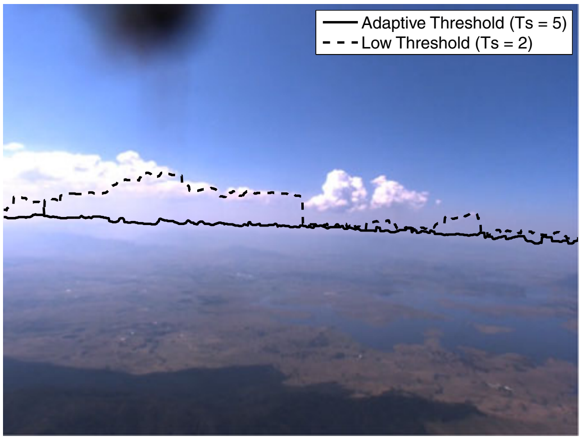

(a) Detected Horizon

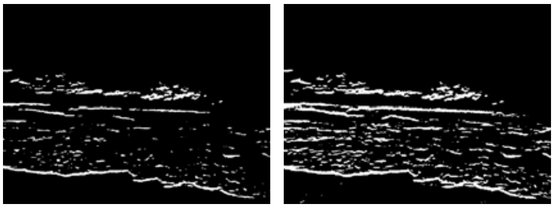

(b) $T_s = 15$ Threshold Image (c) $T_s = 10$ Threshold Image

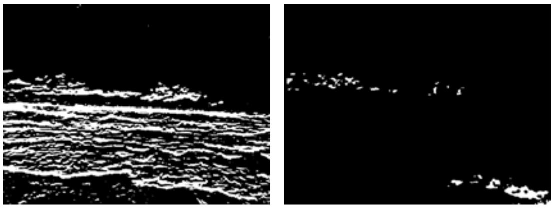

(d) $T_s = 5$ Threshold Image (e) Color Condition Image

Fig. 5 Example image 2

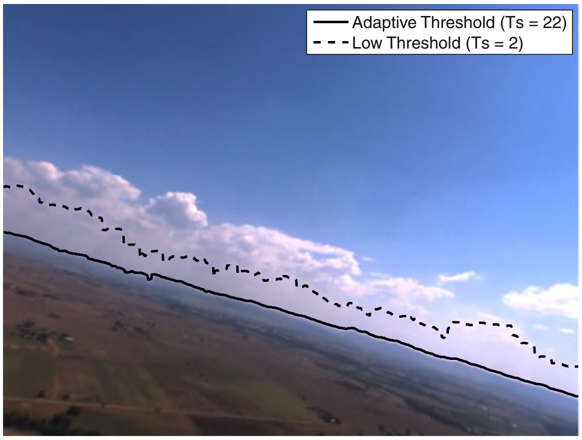

(a) Example Image 3

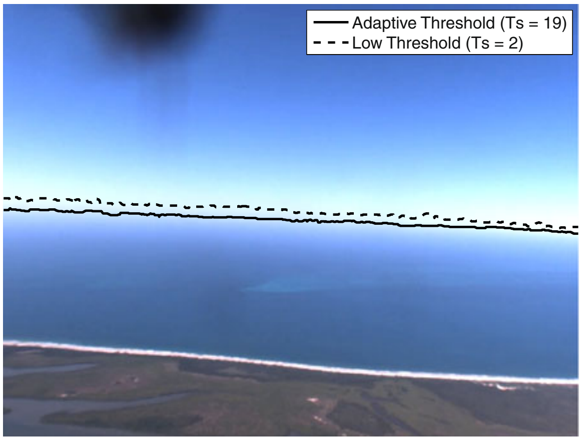

(b) Example Image 4

Fig. 6 Horizon detection examples

and noise associated with the ground. However, the $\boldsymbol{B}_{\text{color}}$ binary image for the second example shown in Fig. 5e would provide little help. The second example image in Fig. 5 has been taken at a higher altitude (>1000 ft) and the ground color distribution has become washed out due to the atmosphere. The washed out color distribution and lack of distinct horizon edge makes this example difficult for other horizon detection methods, however the proposed method is able to find a good representation of the true horizon in this situation. Two other horizon detection examples are shown in Fig. 6a and b which can provide difficulties for other methods. The example shown in Fig. 6b can be particularly difficult for statistical segmentation based approaches, as the sky and the ocean have very similar color and texture distributions.

There are situations when the proposed method can return erroneous results. Two examples of such situations are shown in Fig. 7a and b. The adaptive threshold process finds a gradient spanning the image at a higher threshold value than the value required for correct horizon detection. This can happen in cases when the horizon only has a weak gradient response and there exists a stronger edge feature spanning the image. Another difficult situation for the proposed algorithm is when part of the horizon is discontinuous. The discontinuity can cause the adaptive algorithm to saturate towards the lower threshold limit; searching for any continuous gradient spanning the image. Figure 8a shows the situation

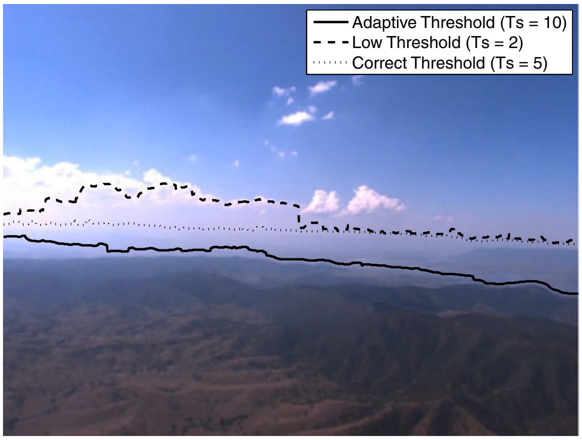

(a) Incorrect Horizon Detection 1

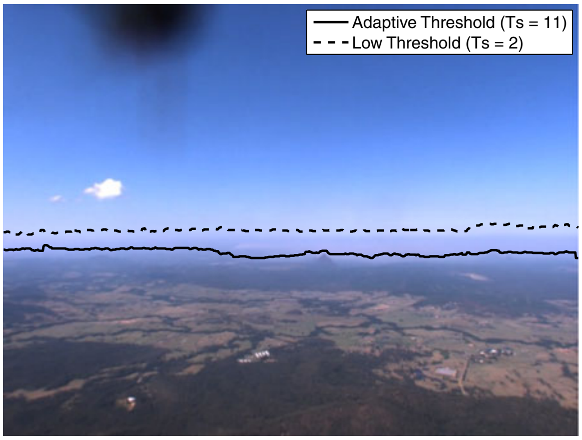

(b) Incorrect Horizon Detection 2

Fig. 7 Incorrect horizon detection examples

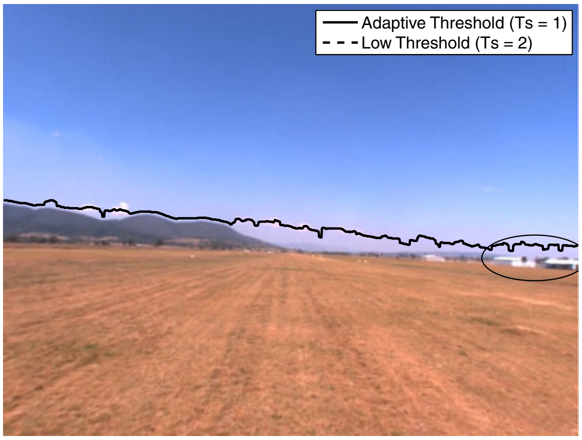

(a) Incorrect Horizon Detection Due to Obstruction

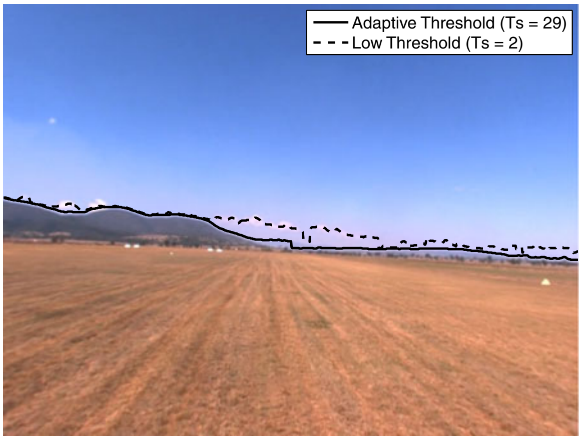

(b) Correct Horizon Detection Without Obstruction

Fig. 8 Incorrect horizon detection due to obstruction

when part of a strong horizon is obstructed with white buildings. Due to the color and shape of the buildings, they have no vertical gradient response. Once the buildings leave the image frame, as shown in Fig. 8b, then the adaptive algorithm can correctly find the strong horizon.

3 Attitude Determination

The appearance of the detected horizon needs to be linked to the attitude of the aircraft so that attitude determination can be carried out. For the purpose of this paper, a straight horizon line is fitted to the detected horizon profile, so that the proposed method can be compared to other horizon detection methods which use a straight horizon line approximation [2–4, 6–10, 13, 25]. The detection of the horizon profile would also allow the additional shape information (which should ideally be the local terrain) to be used in different methods. These methods could be a terrain-aided horizon method for attitude determination [14] (which would be more accurate than the horizon line approximation) or a terrain-aided horizon method for position estimation [23, 24]. Section 3.1 explains the line fitting method used in the proposed method to approximate the horizon profile as a straight line. A mathematical model which links the detected line parameters to the infinite horizon line for attitude determination is presented in Section 3.2. The attitude measurement equations in Section 3.2 assume a forward facing camera, so Section 3.3 explains a transformation process which can correct for camera mounting rotational offsets.

3.1 Horizon Line Fitting

For attitude determination using an infinite horizon line model, a line needs to be fitted to the detected horizon profile. The line fitting is done after the horizon extraction stage (rather than during). The line fitting is done using an m-estimator to help reduce the impact of the terrain profile and noise in the extracted profile on the horizon line estimate. Ordinary least squares could have be used, however since the horizon profile has the freedom to take on a general shape, a robust estimation method is used to minimize the biases that the general shape could have on the overall line fit. An m-estimator [15] is used to fit a line of the form $v = au + b$ to the extracted profile $v = H_p(u)$ by using iterated reweighed least squares to minimize:

$$\min \sum_i w\left(r_i^{k-1}\right) r_i^2 \tag{14}$$

where r_i is the fit residual calculated by $r_i = (v_i - au_i - b)$ and the influence function used is the "fair" function[4] [19], which has the weighting function $w(x)$ of:

$$w(x) = \frac{1}{1 + |x|/1.3998} \tag{15}$$

3.2 Infinite Horizon Line Model

The papers [2–4, 8–10, 13, 25] which investigate horizon detection for MAV or UAV stabilization did not explicitly link the appearance of the horizon to the attitude of the platform. General measures, such as a pitch percentage (ratio of the sky to ground area in the image) were used to regulate the attitude. In [6] a mathematical model was presented which linked the infinite horizon line (vanishing line of the ground plane) to the attitude of the aircraft based upon a normalized focal length of the camera. The model assumed a forward facing camera. The equations presented in [6] are:

$$\phi = \arctan\left(\frac{-m_y}{m_x}\right) \tag{16}$$

$$\theta = \arctan\left(\pm\frac{u\sin(\phi) + v\cos(\phi)}{f}\right) \tag{17}$$

where m_x and m_y are the gradient components of the infinite horizon line, (u, v) are a point along the horizon line, (ϕ, θ) are the bank and pitch of the aircraft and f is the camera's focal length.

The mathematical model in [6] can be extended for a pin-hole camera model [11], with (α, β) being the scaled focal lengths in the (u, v) directions and (u_0, v_0) being the coordinates of the principal point. The focal parameters α and β are a scaled version of the camera's focal length expressed in pixels rather than as a distance. The scale focal lengths can be different which allows for non-square pixels, which is typical for low cost CCD cameras. The horizon line can also be expressed in the line slope intercept form of $v = au + b$ which results in the following horizon line equations:

$$a = -\frac{\beta}{\alpha}\tan(\phi) \tag{18}$$

$$b = v_0 + \frac{\beta}{\alpha}u_0\tan(\phi) + \beta\tan(\theta)\sec(\phi) \tag{19}$$

The horizon line parameter equations (18) and (19) can be inverted to find the attitude of the aircraft as a function of the infinite horizon line parameters:

$$\phi = -\arctan\left(a\frac{\alpha}{\beta}\right) \tag{20}$$

$$\theta = \arctan\left(\frac{(b - v_0)}{\beta}\cos(\phi) - \frac{u_0}{\alpha}\sin(\phi)\right) \tag{21}$$

Equations 20 and 21 are used during the results to calculate the measured attitude of the aircraft from the fitted horizon line parameters. The attitude equations in this form allow for a direct

---

[4] The selection of the influence function was not investigated. Conceivably other influence functions could produce similar or better fitting results.

calculation of the attitude from the fitted line parameters in the image.

3.3 Camera Rotation Offset Correction

The attitude measurements calculated from the infinite horizon line model described in Section 3.2 assumed a forward facing camera. Any rotation offsets in the camera mounting alignment can bias the attitude measurements. These camera rotation offsets can be accounted for by using DCM (direction cosine matrix) transformation equations. The world to camera frame rotation matrix R_{cw} can be evaluated using the attitude measurements:

$$R_{cw} = R_x(\phi) R_y(\theta) \tag{22}$$

where ϕ and θ are the uncorrected attitude measurements from the infinite horizon line model and the rotation matrix $R_{x,y,z}(\cdot)$ represents the standard orthogonal rotation about a single axis. The body to camera frame rotation matrix R_{cb} is calculated using:

$$R_{cb} = R_x(\phi_c) R_y(\theta_c) R_z(\psi_c) \tag{23}$$

where the body to camera frame rotation offset Euler angles are ϕ_c, θ_c, ψ_c. The true world to body frame rotation matrix R_{bw} can be recovered from:

$$R_{bw} = R_{cb}^{-1} R_{cw} \tag{24}$$

where the corrected $\hat{\phi}$ and $\hat{\theta}$ measurement angles are extracted using:

$$\hat{\phi} = \arctan\left(\frac{R_{bw(1,3)}}{R_{bw(3,3)}}\right) \tag{25}$$

$$\hat{\theta} = \arcsin\left(-R_{bw(1,3)}\right) \tag{26}$$

4 Results

For comparison with the proposed method, two other horizon detection methods were also implemented. The first method for comparison is developed in [8–10] by Ettinger et al. and is hereafter referred to as the Ettinger method. The second method for comparison is developed in [6, 7] by Dusha et al. and is hereafter referred to as the Dusha method. The computational time for each method is compared in Section 4.1 and the attitude determination accuracy of each method is evaluated for a test flight video data set in Section 4.2.

4.1 Implementation Results

The three methods (Dusha, Ettinger and the proposed) were implemented in C++ using OpenCV (Open Source Computer Vision processing libraries) with IPP (Intel Performance Primitives) optimizations. The methods were tested on an Intel i5 3.30 Ghz Processor and the frame rates shown in Table 1 were achieved.[5] The parameters shown in Table 2 for the Dusha method, Table 3 for the Ettinger method and in Table 4 for the proposed method were used during the timing and horizon detection tests. Refer to [6–10] for the method details and parameter descriptions. The values of these parameters can greatly affect the computational efficiency and accuracy of the horizon detection methods.

At a resolution of 1024 × 768, the Dusha method performed at 9 Hz while the proposed method was able to run at 42 Hz (using a single iteration or fixed threshold value). The Ettinger method was able to run at 20 Hz however it was on a significantly down-sampled image of size 245 × 192. The other methods can also achieve faster speeds if their images are down-sampled as well. A maximum rate of 48 Hz for the Dusha method and 100 Hz for the proposed method were achieved when down-sampled to a resolution of 640 × 480.

---

[5] The Dusha method was stated in [6] to be able to run at 15 Hz at a resolution of 352 × 288 on a Pentium 4 3 Ghz processor with the method implemented in C using the OpenCV libraries and IPP optimizations. While the Ettinger method was stated in [10] to be able to run at 30 Hz at a resolution of 320 × 240 (down-sampled to 80 × 60) on a 900 Mhz ×86 processor with high and low level optimizations as outlined in [8].

Table 1 Horizon detection method timings

Method	Resolution	Timing (ms)	Frame rate (Hz)
Ettinger	245 × 192	50	20
Dusha	1024 × 768	112	9
Dusha	800 × 600	78	13
Dusha	640 × 480	48	21
Proposed[a]	1024 × 768	24	42
Proposed[a]	800 × 600	17	59
Proposed[a]	640 × 480	10	100
Proposed ($N = 2$)	1024 × 768	31	32
Proposed ($N = 5$)	1024 × 768	61	16
Proposed ($N = 10$)	1024 × 768	90	11

[a] A fixed threshold was used

Table 2 Dusha method parameters

Parameter	Value
Smoothing window size	15 × 15 pixels
Edge threshold	$\sqrt{4\text{mean}}$
Hough transform θ bin size	1°
Hough transform ρ bin size	1 pixel

Table 3 Ettinger method parameters

Parameter	Value
Coarse resolution	64 × 48 pixels
Coarse search size	36 × 36 (θ, σ) parameter space
Fine resolution	245 × 192 pixels
Fine refinement limit	10 iterations

Table 4 Proposed method parameters

Parameter	Value
Smoothing window size	6 × 6 pixels
Fixed threshold	$T_s = 4$
Adaptive threshold range	$T_s \in \{40, 35, \ldots, 10, 7.5, 5, 2.5\}$
Adaptive iteration limit	$N = 10$

Table 5 Proposed method (1024 × 768) timing breakdown

Stage	Timing (ms)
Smoothing stage	9.88
Color binary calculation	2.05
Gradient calculation	1.45
Gradient culling	0.37
Spatial filtering stage (per iteration)	7.26
Horizon extraction stage	3.51

The timings for the proposed method are broken down into the different stages and are shown in Table 5 for a image resolution of 1024 × 768. The majority of the calculation time is spent in the smoothing stage and spatial filtering stage. If an adaptive threshold method is used, then the spatial filtering stage can be repeated a number of times. A frame iteration limit N can be set to keep the algorithm running at an acceptable frame rate for a real-time system as described in Section 2.3.

4.2 Attitude Determination Results

The data set used for comparison was an on-line available subset of the ASL (Airborne Systems Laboratory) data set presented in [7]. The data set is of a test flight video of the Australian Centre for Aerospace Automation (ARCAA) Cessna 172. The details of the system can be found in [12]. The test data contains position and attitude truth information given by a highly accurate NovAtel SPAN system running at 50 Hz. Video footage for a forward facing camera is acquired. The camera is a Point Grey Flea IIDC camera with a resolution of 1024 × 768 with a 60° FOV (Field of View) running at 20 Hz. The test flight was conducted around southeast Queensland, Australia. The available data set contains 38 min of flight data from flight time 1769 s to 4049 s. The test flight attitude sequence is shown in Fig. 9 while

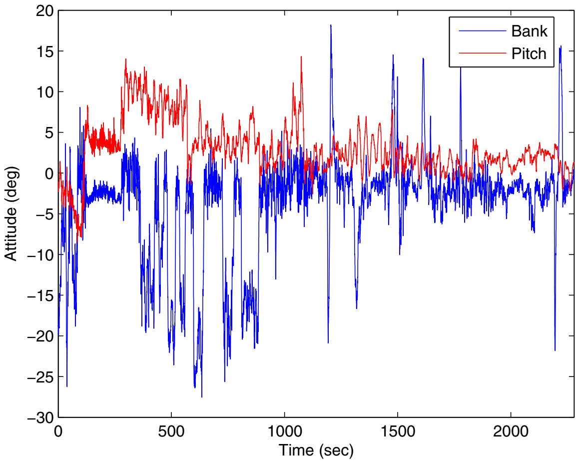

Fig. 9 Test flight attitude sequence

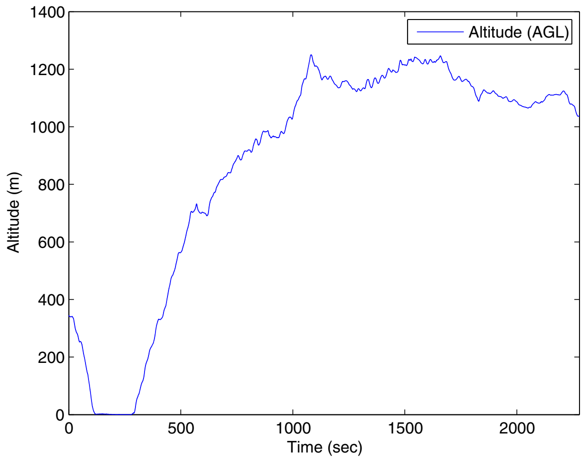

Fig. 10 Test flight altitude sequence

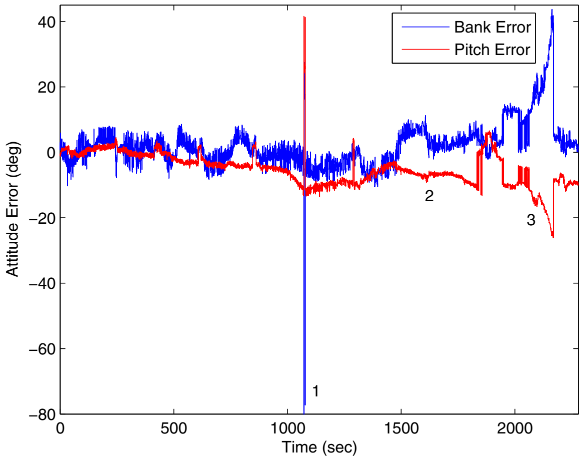

Fig. 11 Ettinger method attitude error

the altitude sequence is shown in Fig. 10. The flight sequence contains significant variations in the bank and pitch of the aircraft, while reaching a maximum altitude of approximately 4000ft.

The camera calibration properties for the data set were not given in [7], however they were estimated to be the values shown in Table 6, with (ϕ_c, θ_c, ψ_c) being the rotation offset mounting Euler angles between the aircraft body and camera frame. Unfortunately, theses offset angles have a noticeable shift when the aircraft is taxiing along the ground which introduces slight biases in the attitude error results during this time period.

Results for the proposed method are presented for the two different threshold selection techniques. The adaptive threshold technique is labeled as (AT) and the fixed threshold is labeled as (FT). For real-time purposes, the adaptive method described in Section 2.3 would limit the iterations per frame. However, since all the processing was done off-line for comparison purposes, the maximum frame iteration limit was removed. The processing rate for the proposed method would then vary between 42 Hz (1 iteration) and 11 Hz (10 iterations) depending upon the image content.

Table 6 Camera calibration properties

Parameter	Value
Focal length α	788.41 pixels
Focal length β	832.96 pixels
Principal point u_0	512 pixels
Principal point v_0	384 pixels
Camera offset ϕ_c	−4.0 deg
Camera offset θ_c	4.5 deg
Camera offset ψ_c	1.0 deg

The tuning of the detection parameters for the proposed method shown in Table 4 were done by hand. The video was of reasonable quality, so a small smoothing window size could be used. The fixed threshold value was selected in conjunction with the smoothing window size, so that a very weak horizon could be detected. The adaptive

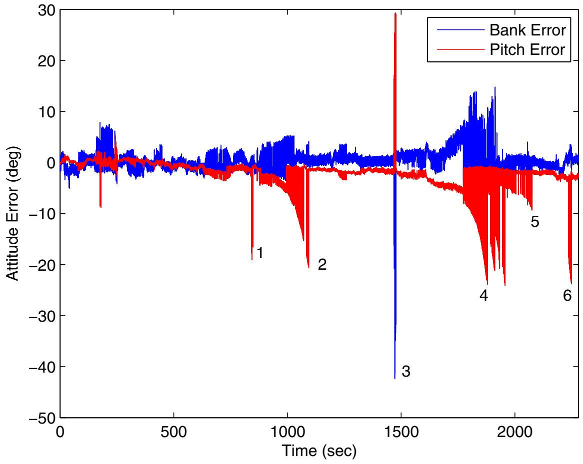

Fig. 12 Dusha method attitude error

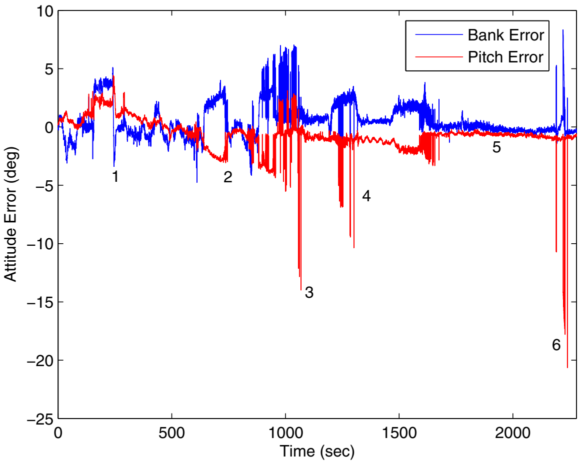

Fig. 13 Proposed method (AT) attitude error

Table 7 Attitude error comparisons

Parameter	Ettinger	Dusha	Proposed (AT)	Proposed (FT)
Max ϕ error	43.56°	20.74°	8.34°	6.68°
Max θ error	41.51°	29.36°	4.84°	5.00°
Min ϕ error	−80.34°	−42.36°	−7.45°	−6.91°
Min θ error	−26.15°	−24.02°	−20.72°	−11.79°
Mean ϕ error	2.08°	0.38°	0.48°	0.28°
Mean θ error	−5.12°	−2.07°	−0.68°	−0.0078°
STD ϕ error	6.68°	2.20°	1.39°	1.00°
STD θ error	5.16°	3.00°	1.40°	0.87°

threshold range was selected based on the typical horizon strength for a strong edge smoothed by a filter using the selected window size. The adaptive threshold step size was decreased with the decreasing threshold to minimise the excess noise content in the sensitive threshold range. Thresholds below zero are not considered as the gradient direction sensitivity would be reversed.

The attitude determination error results for the different methods are shown in Figs. 11, 12, 13 and 14. Each figure plots the attitude error for each frame between the NovAtel SPAN system and the attitude measurement from the horizon detection method. These results are for the raw measurements returned by the different methods without any fault detection or filtering methods applied. Fault detection or filtering methods could be applied to all of the methods, improving the results, however a baseline method comparison was sought. The Ettinger and Dusha method results are shown in Figs. 11 and 12 respectively. The proposed method attitude error results for the different threshold selection techniques are shown in Figs. 13 and 14. A summary of all the results from the different methods is shown in Table 7, which contains the maximum and minimum error, mean error and error standard deviation for both bank and pitch.

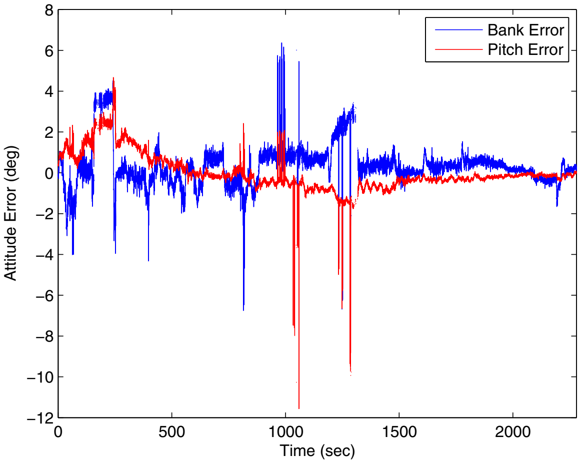

Fig. 14 Proposed method (FT) attitude error

The proposed method performed the best for the test sequence. It returned the smallest error range, mean and standard deviation. The Dusha method also returned small error results for the mean and standard deviation, however they were slightly larger than the proposed method. The maximum and minimum error for the Dusha method were not as well contained. The Ettinger method performed the worst with a large error mean and standard deviation, which was multiple times larger than the other methods.

5 Discussion

During the beginning of the test sequence (< 500sec), all the methods produced similar results. During this time, the altitude of the aircraft was low enough, that the appearances of the sky and ground were dissimilar and the regions were separated by a strong horizon edge. These features allowed all the algorithms to perform fairly well. At these low altitudes, the main source of attitude error was due to the terrain introducing biases in measurements, since a straight horizon was assumed. During the middle and end of the test sequence, the altitude was higher. The higher altitude effected the accuracy of the different methods in different ways. The individual method results are discussed separately in Sections 5.1–5.4 below.

5.1 Ettinger Method

At higher altitudes, the Ettinger method had trouble identifying the difference between the sky and the ground. The primary cause of this, is the fact that the sky and ground regions have similar color distributions at high altitudes due to the terrain color becoming washed out with atmospheric effects. This region similarity makes the algorithm very sensitive to the image content. This is the main source of large attitude errors returned by the method. A number of attitude error features can be seen in Fig. 11. The situation and causes of the significant errors are described below with reference to the numbered features in Fig. 11:

1. The method sensitivity has caused the algorithm to separate the image into a dark sky region and a light sky combined with ground region.
2. The algorithm has locked onto a region of the ground close to the aircraft which is not washed out due to the altitude.
6. The algorithm has locked onto the bottom edge of a region of smoke on the ground. As the aircraft flies towards the region, the translation of the aircraft quickly causes the attitude error to diverge.

The Ettinger method performed well when there was a significant difference between the sky and the ground. However when this difference diminishes, the algorithm becomes very sensitive to disturbances (such as clouds, rivers and smoke) and the accuracy of the algorithm reduces to an unacceptable level for control and navigational purposes. The Ettinger method was designed for stabilization of MAVs [8–10] and as such it was only designed and tested at low altitudes due to the limited performance envelope of MAVs.

5.2 Dusha Method

The primary source of error in the Dusha method, was the need for a distinct straight edge in the image for the algorithm to detect. At high altitudes, the horizon does not produce a strong edge response. The identification and selection of other strong straight edges in the image becomes the main source of erroneous measurements. A number of labeled attitude error features can be seen in Fig. 12, the situation and causes are described below:

1. The algorithm has selected a strong edge of a body of water.
2. The algorithm has locked onto and tracked a strong edge in the image caused by a straight mountain ridge.
3. The algorithm selected the top edge of a series of clouds in the sky.
4. The algorithm has locked onto a bottom edge of haze covering the ground.
5. The bottom edge of a cloud of smoke from a bushfire has been selected.
6. The strong edge of the coast line has been selected.

The Dusha method performed significantly better than the Ettinger method on the data set. The main problem with the algorithm was its inability to detect the horizon when the horizon has a gradual transition between sky and ground. Other strong edges in the image are also easily returned. The paper which developed the Dusha method [6] used a Kalman filter and optical flow to track candidate horizon lines over time and cull any lines which had an optical flow component caused by the translation of the aircraft. This filtering method could not be applied in real time (<2 Hz). Paper [7] reported an attitude determination error standard deviation of 1.79° for bank and 1.75° for pitch, for the complete data set with the optical flow candidate horizon line filtering method implemented. The proposed method worked in real-time, resulting in smaller errors even without any fault detection or filtering methods applied.

5.3 Proposed Method (AT)

As with the other methods, the terrain introduced biases into the attitude measurements during the beginning of the test sequence. However at higher altitudes, the proposed method was able to select the gradient transition caused by the horizon. This transition region is a very strong and stable indication of the attitude of the aircraft. The transition region is usually the most distant part of the image, so the infinite horizon line approximation for the attitude determination performs extremely well. When this transition region is correctly identified, it can produce an unbiased attitude measurement with a small variance. Attitude error features in Fig. 13 are described below:

1. Attitude errors are introduced by the terrain profile when the aircraft is close to the ground. Change in camera mounting angles (due to the aircraft being on the ground) introduced a noticeable step bias.
2. A stronger continuous edge of a mountain ridge which spans the image was selected over the weaker horizon due to the adaptive threshold.
3. A very low threshold was needed. This caused the detection noise in the image to increase. It selected a weak group of noisy edges caused by mountains, rather than the slightly weaker horizon.
4. Fog and haze obscured a large part of the horizon, such that there was no gradient response across it (even with the lowest threshold). At this point the algorithm is very sensitive, so it selected the bottom of the haze where it transitions to the clear ground.
5. This segment highlights one of the stronger advantages of this method. It cleanly identifies the horizon transition region producing excellent attitude information. This is a situation when the other methods have trouble due to lack of distinguishing features.
6. A strong coastline edge across the image is picked up before the threshold can be adapted to the weak horizon.

5.4 Proposed Method (FT)

The fixed threshold allows the horizon gradient transition region to be consistently identified. This resulted in fairly stable and unbiased attitude measurements during the high altitude phases. In these phases, the measurements were not biased due to the terrain or the view distance. Since the threshold was fixed at a low value, it did not select stronger edges on the ground (ie. the coastline) over a weak horizon. This was the main reason the fixed threshold produced better attitude results compared to the adaptive threshold. However the low threshold value makes the extracted horizon profile very sensitive to noise and disturbances such as clouds or other weak edges; this can clearly be seen in the example Fig. 4. At low altitudes, the returned profile was very noisy and not as precise due to the sensitive threshold value. At high altitudes however this was less of a problem as accurate profile extraction is less relevant. For accurate horizon profile extraction (such as at low altitude, when there is a terrain profile to detect), the fixed threshold is inappropriate as the profile is too sensitive to noise. An extension to the adaptive method is developed in Section 6, which allows the proposed method with an adaptive threshold to perform as well as the fixed threshold for attitude determination while maintaining its horizon profile extraction performance.

6 Attitude-Aided Detection Extension

The two incorrect horizon detection examples shown in Fig. 7 and most of the erroneous detections from the test sequence (using an adaptive threshold) are caused by a stronger gradient spanning the image which is detected before the weaker horizon. These situations can be corrected by using extra information in the detection process.

If an estimate of the current attitude is known, then this information can be used to improve the detection process. Attitude information allows the estimation of two sets of image border pixels which must be connected together by the horizon profile. This decreases the set size used in the horizon spatial filtering stage. Stronger continuous

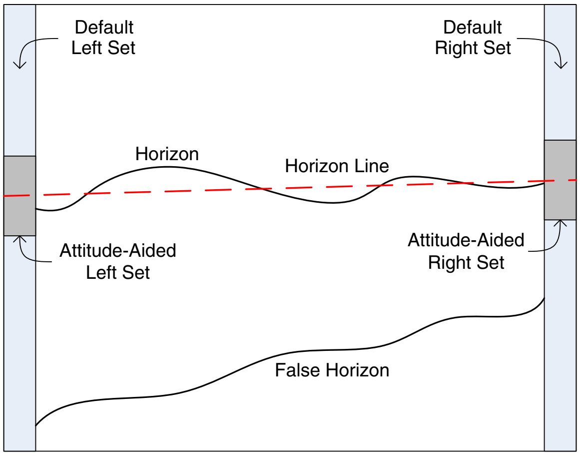

Fig. 15 Attitude-aided spatial filtering set generation

gradients which are not the horizon are removed by this process, as they would not fall within the revised attitude-aided spatial constraints. The diagram shown in Fig. 15 explains how the estimated attitude can be linked to the horizon line in the image; this can be used to generate attitude-aided border sets for constraints. Gradients which do not continuously span these two border sets are filtered out.

Using the horizon line equations (18) and (19) which express the horizon line in the form of $v = au + b$ with $(a, b) = f(\phi, \theta)$, it is possible to relate the current estimated attitude to the image border pixel sets. Chi-Squared tests then can be applied to the image border coordinates to extract two sets of border pixels which must be connected by the horizon. These two pixel border sets are used in the left and right morphological reconstruction process $(B_{\text{left}}, B_{\text{right}})$ for the spatial filtering stage of the proposed algorithm.

Taking partial derivatives of Eqs. 18 and 19 with respect to the estimated attitude gives:

$$\frac{\partial a}{\partial \phi} = -\frac{\beta}{\alpha}(1 + \tan^2 \phi) \tag{27}$$

$$\frac{\partial a}{\partial \theta} = 0 \tag{28}$$

$$\frac{\partial b}{\partial \phi} = \frac{\beta}{\alpha}\frac{(u_0 + \alpha \sin \phi \tan \theta)}{\cos^2 \phi} \tag{29}$$

$$\frac{\partial b}{\partial \theta} = \frac{\beta}{\cos \phi \cos^2 \theta} \tag{30}$$

These partial derivatives allow the estimated attitude variances to be transformed to the horizon line parameter variances using:

$$\sigma_a^2 = \left(\frac{\partial a}{\partial \phi}\right)^2 \sigma_\phi^2 \tag{31}$$

$$\sigma_b^2 = \left(\frac{\partial b}{\partial \phi}\right)^2 \sigma_\phi^2 + \left(\frac{\partial b}{\partial \theta}\right)^2 \sigma_\theta^2 + 2\frac{\partial b}{\partial \phi}\frac{\partial b}{\partial \theta}\sigma_{\phi\theta} \tag{32}$$

$$\sigma_{ab} = \frac{\partial a}{\partial \phi}\frac{\partial b}{\partial \phi}\sigma_\phi^2 + \frac{\partial a}{\partial \phi}\frac{\partial b}{\partial \theta}\sigma_{\phi\theta} \tag{33}$$

where $\sigma_a^2, \sigma_b^2, \sigma_{ab}$ are the horizon line variances and $\sigma_\phi^2, \sigma_\theta^2, \sigma_{\phi\theta}$ are the attitude variances associated with the current estimated attitude. Image border pixels can be tested to see if they should be contained in one of the connection border sets using a Chi-Squared test with a single degree of freedom and a probability threshold of P:

$$\frac{(au_i + b - v_i)^2}{u_i^2\sigma_a^2 + \sigma_b^2 + 2u_i\sigma_{ab}} < \chi^2(P, 1) \tag{34}$$

with the test being performed on all the image border pixels belonging to $(u, v)_i \in \{(u, v)_{u=\{0,\text{width}\}}, (u, v)_{v=\{0,\text{height}\}}\}$.

This method can only be used when the test identifies two unconnected border pixel sets. If the sets are connected, then a horizon cannot be found which joins them. Since the sets need to be unconnected, this method cannot be applied at extreme attitudes when the horizon does not span two distinct border regions.

7 Conclusion

The development of robust horizon detection methods for attitude determination is a key step towards the development of a passive attitude sensor scheme for navigation and control of small aeronautical platforms. The proposed horizon detection method removes the need for the straight horizon line assumption in the detection process and is able to extract the horizon profile shape from the image. It is able to perform more accurately than past approaches (when using an horizon line assumption), returning a mean error of 0.28° and −0.0078° for bank and pitch with a standard deviation of 1.00° and 0.87° for the test sequence. The ability to extract the horizon profile also means that the shape information could be used in additional attitude and navigational processes which would result in greater accuracy. The proposed method is able to run with a higher processing rate compared to other past methods, making the proposed method suitable for real-time implementation.

References

1. Bohren, C.F., Huffman, D.R.: Absorption and Scattering of Light by Small Particles. Wiley, New York (1983)
2. Cornall, T.D., Egan, G.K.: Measuring horizon angle from video on a small unmanned air vehicle. In: 2nd International Conference on Autonomous Robots and Agents (2004)
3. Cornall, T.D., Egan, G.K.: Heaven and Earth: how to tell the difference. Technical Report MECSE-21-2005, Department of Electrical and Computer Systems Engineering, Monash University (2005)
4. Cornall, T.D., Egan, G.K., Price, A.: Aircraft attitude estimation from horizon video. Electron. Lett. **42**(13), 744–745 (2006)
5. Davies, E.R.: Machine Vision: Theory, Algorithms, Practicalities. Elsevier, Amsterdam (2005)
6. Dusha, D., Boles, W., Walker, R.: Fixed-wing attitude estimation using computer vision based horizon detection. In: 12th Australian International Aerospace Congress, pp. 1–19 (2007)
7. Dusha, D., Mejias, L., Walker, R.: Fixed-wing attitude estimation using temporal tracking of the horizon and optical flow. J. Field Robot. **28**(3), 355–372 (2011). doi:10.1002/rob.20387
8. Ettinger, S.M.: Design and Implementation of Autonomous Vision-Guided Micro Air Vehicles. University of Florida (2001)
9. Ettinger, S.M., Nechyba, M.C., Ifju, P.G., Waszak, M.: Towards flight autonomy: Vision-based horizon detection for micro air vehicles. In: Conference on Recent Advances in Robotics, Florida (2002)
10. Ettinger, S.M., Nechyba, M.C., Ifju, P.G., Waszak, M.: Vision-guided flight stability and control for micro air vehicles. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 2134–2140 (2002)
11. Forsyth, D., Ponce, J.: Computer Vision: A Modern Approach. Prentice Hall, Englewood Cliffs (2002)
12. Greer, D.G., Mudford, R., Dusha, D., Walker, R.A.: Airborne systems laboratory for automation research. In: Proceedings of 27th International Congress of the Aeronautical Sciences, Acropolis Conference Centre, Nice (2010)
13. Gui-Qiu, B., Shen-Shu, X., Zhao-Ying, Z.: Vision-based horizon extraction for micro air vehicle flight control. IEEE Trans. Instrum. Meas. **54**(3), 1067–1072 (2005)
14. Gupta, V., Brennan, S.: Terrain based vehicle orientation estimation combining vision and inertial measurements. J. Field Robot. **25**(3), 181–202 (2008)
15. Holland, P.W., Welsch, R.E.: Robust regression using iteratively reweighted least-squares. Commun. Stat., Theory Methods **6**(9), 813–827 (1977)
16. Luo, J., Etz, S.: A physical model-based approach to detecting sky in photographic images. IEEE Trans. Image Process. **11**(3), 201–212 (2002)
17. Minnaert, M.G.J.: The nature of light and colour in the open air. In: Dover Books on Earth Sciences. Dover, New York (1954)
18. Mondragón, I.F., Campoy, P., Martinez, C., Olivares, M.: Omnidirectional vision applied to Unmanned Aerial Vehicles (UAVs) attitude and heading estimation. Robot. Auton. Syst. **58**(6), 809–819 (2010)
19. Rey, W.J.J.: Introduction to Robust and Quasi-Robust Statistical Methods. Springer, New York (1983)
20. Sergiy, F., Volha, S., Lawrence, O.H., Dmitry, B.G.: Horizon detection using machine learning techniques. In: 5th International Conference on Machine Learning and Applications, pp. 17–21 (2006)
21. Todorovic, S., Nechyba, M.C., Ifju, P.G.: Sky/ground modeling for autonomous MAV flight. In: IEEE International Conference on Robotics and Automation, vol. 1421, pp. 1422–1427 (2003)
22. Vincent, L.: Morphological grayscale reconstruction in image analysis: applications and efficient algorithms. IEEE Trans. Image Process. **2**(2), 176–201 (1993)
23. Woo, J., Kweon, I., Kim, G., Kim, I., Yuseong, P., Daejeon, K.: Robust horizon and peak extraction for vision-based navigation. In: IAPR Conference on Machine Vision Applications, Tokyo, Japan (2005)
24. Woo, J., Son, K., Li, T., Kim, G., Kweon, I.: Vision-based UAV navigation in mountain area. In: IAPR Conference on Machine Vision Applications, Tokyo, Japan (2005)
25. Yuan, H.Z., Zhang, X.Q., Feng, Z.L.: Horizon detection in foggy aerial image. In: International Conference on Image Analysis and Signal Processing, pp. 191–194 (2010)