

# An Autonomous Vehicle Navigation System using Panoramic Machine Vision Techniques

Kevin Rushant,  
Department of Computer Science,  
University of Sheffield, GB.  
email: krusha@dcs.shef.ac.uk

Libor Spacek,  
Department of Computer Science,  
University of Essex, Wivenhoe Park,  
Colchester, CO4 3SQ, GB.  
email: spacl@essex.ac.uk

*Keywords:* Horizon, panoramic images, autonomous vehicle navigation.

## Abstract

This paper describes a navigation system for an autonomous farm vehicle using machine vision techniques applied to panoramic images of the horizon. The navigation system finds the position using triangulation of the selected horizon features. Simple machine vision techniques are sufficient for the horizon analysis. Restraints considered throughout this work include the processing cost of the procedures used and the ability to achieve results in a variety of outdoors farm environments, along with the financial cost of the system.

## 1 Introduction

The goal is to produce a system capable of visually triangulating the position of the moving vehicle in an unfamiliar outdoors environment. The approach examined in this paper is based on the empirical observation that the horizon line often provides the strongest contrast with the ‘cleanest’ features. The method consists of identifying and tracking ‘interesting’ horizon features such as those produced by tops of trees, bushes and buildings.

The initial motivation for this work came from a seminar given at the University of Essex by Ulrich Nehmzow of the University of Manchester [4]. A small desktop robot was used as part of the presentation to demonstrate navigation using a four way light sensor (‘a light compass’). The robot was set on the table and pushed away in a variety of directions. It then made use of dead reckoning information calculated from the number of wheel turns and the light sensor information to return to the original position. An error of approximately 10% of the distance traveled was observed. This was explained as being due to errors in the dead reckoning caused by wheel slippage. Clearly, there is a need for improved dead-reckoning methods, particularly over uneven terrain.

Further motivation for the research came from a need to develop robust navigation methods for an autonomous farm vehicle being designed by the University of Essex in cooperation with the Writtle Agricultural College.

## 2 Horizon Imaging

Initial experiments used multiple images of the horizon to assemble one panoramic strip. The image data was obtained by mounting a video camera on a tripod and capturing 16 images at regular intervals while turning the camera a full 360°. The images were then assembled into a single strip. The panoramic strip was used for preliminary tests of the visual navigation methods.

This process suffered from time delay problems. Specifically, the camera’s automatic aperture made it difficult to make sensible comparisons of the gray-level values across the strip (illumination problems). For example, a panoramic strip obtained inside a laboratory exhibited the lowest intensity values at the brightest point, which was a sunlit glass doorway. Due to the automatic aperture this otherwise clear landmark was reduced to an unidentifiable dark area.

Perhaps the most serious problems are caused by the independent motions of the vehicle and the objects while the horizon is being scanned, and by the vehicle and camera shake over the rough terrain.

For all these reasons an alternative solution is required, namely to capture the whole horizon in a single instant. By using a spherical mirror mounted above the upward facing camera it is possible to capture the entire horizon in a single image, as shown in Figure 1. The mirror is preferred to a ‘fish-eye’ lens in order to avoid mounting the camera high above the vehicle.

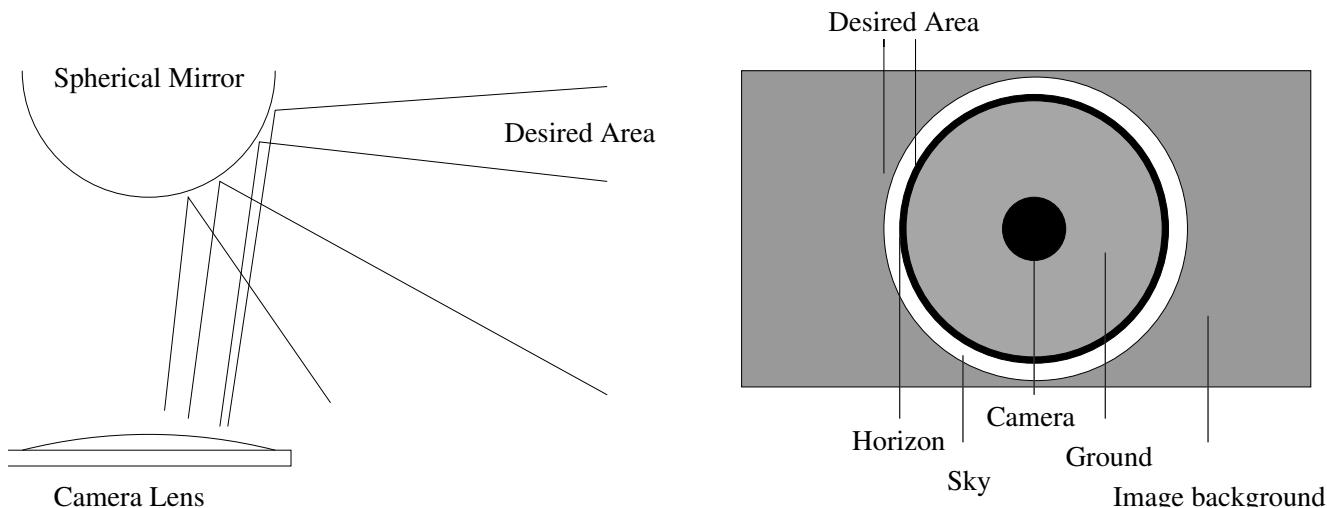


Figure 1: *The device for capturing an image of the full horizon*

The image area depicting the horizon line is relatively small. It is possible, and desirable, to improve the resolution of the horizon region by using specially made conical mirror resulting in more appropriate image projection, followed by a projection correcting image transformation. Due to cost and availability this study was undertaken using a spherical mirror and the crude supporting bracket, which is clearly visible in the images. However, the bracket, being in a fixed position, does not interfere with the tracking of the horizon features. It can even be used to give a rough indication of the vehicle orientation relative to the horizon features.

### 3 Image Transformation

To allow a time efficient solution to be produced the original captured image is transformed to a panoramic strip image. This step reduces the projection distortion. The time consuming trigonometry functions and inter-pixel interpolations need to be calculated only once and all subsequent computations are carried out on the new strip image.

The original image and the transformed image are shown in Figure 2 and Figure 3 respectively. It can be seen that very little detail is lost during the transformation process. A factor influencing the strip image size is the resolution of the camera. The camera used to capture the image shown has the spatial resolution of 768 x 576 pixels. The resulting strip has a length of 720 pixels which means 0.5 degrees per pixel for calculation purposes. Thus this angular accuracy is within the range of the cheapest cameras.




Figure 2: The original image, showing the mounting bracket, the camera, and the horizon




Figure 3: The transformed image showing the extracted horizon line, plotted clockwise. It is inverted vertically for programming convenience only. The two gaps correspond to the wire bracket locations.

### 3.1 Locating the Horizon Extremities in the Image

By detecting the horizon within the original image, the position of the center of the spherical mirror can be identified and thus the transformation can be performed correctly. It would not be necessary to compute this 'registration' of the center if the mirror was always guaranteed to be in the same fixed position with relation to the camera.

A simple differencing edge detector operator is used. The template has dimensions 12x24 pixels with the left six columns containing weights of +1's and the right six columns containing -1's. By convolving this template over the left side of the image it is possible to identify the left position of the horizon. The positive weight values correlate with the relatively bright sky and the negative weight values correlate with the darker tree line, hedgerows, or ground.

This procedure is repeated with different orientations of the grid to locate the top and the right side extremities of the depicted horizon. These are then used to calculate the center and the radius of the horizon in the original image for the transformation process. See Figure 4.




Figure 4: Locating the horizon line in the image

### 3.2 Image Smoothing

An unfiltered transformation would copy the pixel value from the original image directly to the new panoramic strip image. Unfortunately, the image noise produced from camera shake during capture would not be removed. To reduce this noise a filter was used to average the value of the pixel depending upon the values of its nearest neighbours [1, 2, 5]. Experiments were conducted with both a median and a mean filter with varying filter sizes of between 9 and 49 pixels. The median filters which work well with salt and pepper noise failed to reduce the noise sufficiently. However the larger mean filters produced good results. These filters also smoothed the horizon line and thus allowed later horizon comparisons to be made more reliable and stable. The 25 pixel grid was chosen as the more economical version capable of producing adequate results.

### 3.3 Locating the Bracket

Once the panoramic strip image has been created it is necessary to remove the areas containing the wire mounting bracket. This is done by locating two vertical lines with the strongest contrast and checking that they are directly opposite each other. They are then masked to remove the possibility of identifying the bracket as a horizon feature.

The panoramic strip is then realigned to the first masked mounting bracket support to reduce the horizon rotation between frames. For navigation purposes the number of pixels the image was shifted by is stored to allow calibration of navigation data.

## 4 Detecting the Horizon Contour Line

A simple edge detector was used to locate and extract the continuous horizon line. This works well, as the horizon line usually has the strongest contrast out of all the edges in the image. The edge detector grid is a single pixel wide by 8 pixels high. This is used to locate the horizon at each of the 720 positions along the horizon strip image.

The most useful property of the horizon line is its varying height. The detected height values (y coordinates) are stored in an array of unsigned chars to reduce the storage capacity needed to characterise each field location.

### 4.1 Horizon Line Smoothing

The panoramic strip image and thus the resulting plot array still contain undesired noise. This can be seen as small random changes in height values. To reduce this noise another mean smoothing filter was used, this time using the surrounding 6 height values to calculate an average height.



Figure 5: The edge detected horizon used for the feature tracking process

## 5 Horizon Features Identification

We define a horizon feature is an area of the horizon which is easily identifiable. This most commonly consists of a tall tree in the middle of a low hedgerow but could just as likely be a small group of bushes or a building.

A variety of ways are available to identify the best features. These range from examining high peaks and low troughs to selecting the area of highest standard deviation of the horizon heights. Also the width of the features make a great difference to system accuracy.

The adopted solution consisted of simply marking the high peaks within partitioned zones (see below for the partitioning method). The three features with the highest standard deviation values are then to be tracked into the next image.

### 5.1 Horizon Features Separation

The horizon was split into six zones which were positioned away from the masked mounting bracket areas to prevent an identified feature becoming occluded by the mounting bracket in the following image. Six zones were used to allow redundancy of horizon features within the system. Above all, this solution ensures that the selected features have a significant angular separation and therefore increases the accuracy of the triangulation calculations.




Figure 6: The horizon zones

## 6 Horizon Features Matching

Each selected feature is stored in a template to be matched in the next image. The templates used are 45 pixels wide with zone widths of 90 pixels. This was determined experimentally.

The next image is processed and the first template is compared to each position within a 30 degrees range of its original position. This was considered beyond the maximum distance of travel for a feature between images.

The first part of the matching process is to evaluate the average difference in height between the template and the new image. This is done by summing the differences between each pair of pixels and taking the mean average as the height normalising offset.

The next step is to evaluate the similarity of the templates using the sum of absolute values of the differences of normalised heights.

## 6.1 Feature Matching Experiment




Figure 7: The first image in the test sequence: features selection




Figure 8: The second image in the test sequence: matching and re-selection




Figure 9: The third image in the test sequence

Figure 7 is used as the origin for the test image sequence, therefore no features are yet matched within it.

The black lines represent the boundaries of the three features selected. As can be seen, the selected features are from zones 1, 2 and 4. These features are encoded into the templates for matching in Figure 8.

The grey lines represent the areas that the templates have been matched to. In zone 2, the matched feature has been selected again in the correct position of the match and therefore the grey lines have been overwritten with black.

The newly selected features from Figure 8 are successfully matched in Figure 9. Note that although the matching process scans beyond the zone boundaries, the feature selection process does not.

## 6.2 Tracking

Experiments show that a particular feature cannot continue to be matched well (or tracked) through a long sequence of several images. It is therefore necessary to re-select new features in each new image and only match each feature between a pair of two adjacent frames (subsequent images). This provides robustness and redundancy, with the navigation continuing to function even if significant numbers of the previously selected horizon features become occluded.

## 7 Navigation

Once the system is placed in the field, the first stage is to calibrate it. This is done by moving a measured distance and capturing two separate images. It is then possible to calculate the distances to the selected features. With this information it is possible to calculate the distance moved between subsequent images (dead reckoning) and therefore build up a map of the surroundings, a movement path, and subsequently navigation information.

### 7.1 Bearing error

The test images used were captured by positioning a camera, with its attached horizon imaging device, in various positions in a field and reading a bearing taken from a hand held compass. This process however is not sufficiently accurate for producing navigation information, as 0.5 degree of error can cause significant position errors for distant objects.

Features within the images can be seen to converge or separate which gives approximate motion direction and could be the subject to detailed optical flow analysis. However, without separating the rotational and translational flow fields, it is difficult to calculate the actual movement accurately.

Further work has been conducted which demonstrates that with the addition of accurate electronic compass measurements taken on board, good navigation results are possible. This was proved by manufacturing a sequence of images with feature positions changing in relation to a constant movement of 5 metres per image. The resulting images were processed by the previous routines and a map created. The map reproduces accurately the calculated vehicle movement of 5 metres per image.

## 8 Conclusion

The image capture and image transformation processes used are sufficient to produce accurate dead-reckoning results from the feature tracking process (with the addition of an electronic compass mounted in a fixed position relative to the camera and the lens, and thus subtracting out the bearing error).

This paper has shown that a relatively inexpensive navigation system is available for an autonomous outdoors vehicle, making use of a variety of basic machine vision techniques [3, 2].

This navigation method is particularly suited for autonomous agricultural vehicles in situations where there may not be any nearby objects for navigation by stereopsis or other traditional methods.

The central observation which makes this approach work is that the horizon line has the strongest contrast, and thus can be detected reliably and coherently even with simple techniques. The horizon line also typically contains several well spaced fixed features, such as trees, poles, and buildings, which can be successfully matched for navigation purposes.

The visual methods used in this work are fully general, ie. they do not depend on recognition and classification of the horizon objects, for example as trees and buildings. It is not necessary to use any special markers, or rely on detailed knowledge of any particular fixed environment.

## Acknowledgements

Many thanks to Steffen Schlachter who provided much assistance and ample advice, as did Sue Sharples.

## References

- [1] Michael C. Fairhurst. *Computer Vision for Robotic Systems*. Prentice Hall, 1988.
- [2] Jain, Kasturi, and Schunck. *Machine Vision*. McGraw Hill International Books, 1995.
- [3] Eric Paul Krotkov. *Active Computer Vision by Cooperative Focus and Stereo*. Springer Verlag, 1989.
- [4] Ulrich Nehmzow and Brendan McGonigle. Robot navigation by light. *European Conference on Artificial Life, ECAL*, May 1993.
- [5] Libor Spacek. Edge detection and motion detection. *Image and Vision Computing*, 4(1), pp.43-56, February 1986.