AUTOMATIC CORRECTION FOR ATMOSPHERIC DEGRADATION IN INFRARED IMAGES

D R EDMONDSON, L A WAINWRIGHT, E J PRICE, S A PARRISH, M GODFREY, N M FAULKNER, S P BROOKS
British Aerospace Systems and Equipment, Clittaford Road, Plymouth, DEVON PL6 6DE
J P OAKLEY, M J ROBINSON
University of Manchester

SUMMARY

This paper is concerned with methods to correct atmospheric degradation in thermal images in the 8-12 and 3-5 micron wavebands. Previous work [1], [2] has shown that a significant improvement in effective visibility may be achieved in the visible spectrum by using a simple model for atmospheric attenuation. The extension of this work into the infrared wavebands is reported here.

The method is based on a three-parameter model for the transmission of flux from the terrain to the imaging sensor. The measured flux is given by:

$$I = C_0 (1+C_1 \exp(-\sigma R))$$

where R is the range and C_0, C_1 and σ are constants which depend on the atmospheric parameters and the terrain radiance σ is known as the extinction coefficient. The inverse problem is solved to obtain estimates for the three parameters using the actual image data. The values of sigma obtained in this way are in good agreement with those computed using MODTRAN.

A more interesting and useful feature of this approach is that, when the model is combined with a range map corresponding to the image, the atmospheric degradation may be reversed.

Results are presented using data from the visible, 3-5 micron and 8-12 micron wavelength image sequences in a maritime environment. The data was collected under conditions of moderate-to-poor visibility with extinction coefficients in the range 0.16 to 0.62km^{-1} for the IR wavebands. The enhanced images show a significant subjective improvement in clarity. This improvement is confirmed using three metrics based on contrast, SNR and segmentability. The method has been shown to offer worthwhile improvements over the full range of visibilities - improvements capable of translation into improved detection, identification and tracking.

1. INTRODUCTION

One important limitation in using a thermal imager for surveillance and tracking is its performance in adverse weather conditions. Any method which offers the possibility of increasing the range at which targets can be detected is well worth examining. The model-based contrast-enhancement method, the subject of this paper, falls very firmly into that category.

This paper describes an image-processing method for increasing the range at which targets can be detected, which is automatic in the sense that no parameters need be entered by the user.

This method [1], [2] was originally developed as a result of a collaboration between the University of Manchester and British Aerospace Military Aircraft and Aerostructures. The initial work was directed towards improving the visibility of terrain features captured by a forward-looking airborne Charge Coupled Device (CCD) camera.

Significant enhancement of images taken in conditions of haze and low-level fog was shown [1]. The work was later extended, by British Aerospace Systems and Equipment in collaboration with the University of Manchester, to the enhancement of infrared images in adverse weather conditions. The purpose of this paper is to report on recent experiments with the algorithm using infrared images from both the 3-5 micron and the 8-12 micron bands. Visible spectrum images of some of the scenes were also collected. The algorithm used is essentially that reported in [1]. However, a full description is given here for the convenience of the reader.

2. AIMS AND OBJECTIVES OF THE STUDY

The primary aim was to quantify both the absolute and the relative performance of the algorithm in three wavebands: visible light, 3-5 micron and 8-12 micron. The experimental arrangements are described below in Section 4.

Three metrics have been used to measure the improvement in image: one for image contrast, one for Signal-to-Noise Ratio (SNR) and one which quantifies the ease with which an image can be accurately segmented. These metrics, together with the results of their application, are described in Section 5 below.

One interesting feature of the Oakley algorithm is that, as a by-product of image processing, it generates an estimate for the overall extinction coefficient (across the whole band of interest). Thus one possible check that the algorithm is working correctly is to compare the figures for extinction coefficients with those produced by current MODTRAN codes for the corresponding atmospheric conditions. The results from this comparison are also presented in Section 5.

The overall aim is to establish whether this method *truly* offers advantages if applied in a real system.

3. EXPLANATION OF THE OAKLEY METHOD

The Oakley algorithm subtracts the extra radiance introduced by the atmosphere radiating in the infrared (the "airlight") from the total radiance detected at the camera. The remaining component is then the attenuated radiance from the terrain. This is scaled by a factor which compensates for degradation by scattering and absorption. Since the grey levels of the thermal image are directly related to the irradiance at the camera lens, an enhanced image using the compensated values can be generated.

In a vacuum there would be no attenuation of flux by the atmosphere and, if the solid angle corresponding to the k_{th} pixel was Ω_k and the terrain radiance was M_k, then the irradiance at the sensor from a particular patch of terrain would be given by:

$$J_0 = \Omega_k M_k \qquad (1)$$

Paper presented at the RTO SET Symposium on "E-O Propagation, Signature and System Performance Under Adverse Meteorological Conditions Considering Out-of-Area Operations", held at the Italian Air Force Academy, Naples, Italy, 16-19 March 1998, and published in RTO MP-1.

The k_{th} pixel value, P_k, is then:

$$P_k = G\Omega_k M_k \qquad (2)$$

where G is the conversion factor of the sensor.

The attenuation of the radiance from the terrain, J_r, due to scattering and absorption by the atmosphere is modelled by:

$$J_r = J_o exp(-\beta R_k) \qquad (3)$$

where J_0 is the irradiance that would theoretically come from the terrain if the intervening medium were a vacuum, β is the extinction coefficient and R_k is the distance of a terrain patch from the sensor.

The radiance introduced by the atmosphere (the "airlight"), J_b, is modelled by:

$$J_b = \frac{DP\Omega_k}{\beta} [1 - exp(-\beta R_k)] \qquad (4)$$

where D is the density of the particles in the aerosol and P is the radiant intensity of a single particle.

Therefore, the total irradiance at the camera is the sum of (3) and (4) ie:

$$J_{total} = J_b + J_r \qquad (5)$$

Figure 1 shows the solid line representing equation (5) fitted to the actual irradiance and range values for individual pixels in the image.

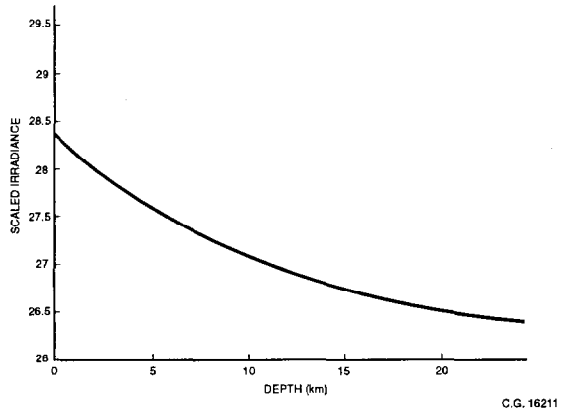

Figure 1 Solid Line Representing Equation (5)

Figure 2 shows how equation (5) comprises the two components - the attenuated terrain component decreasing with range (equation (3)) and the airlight increasing with range equation (4).

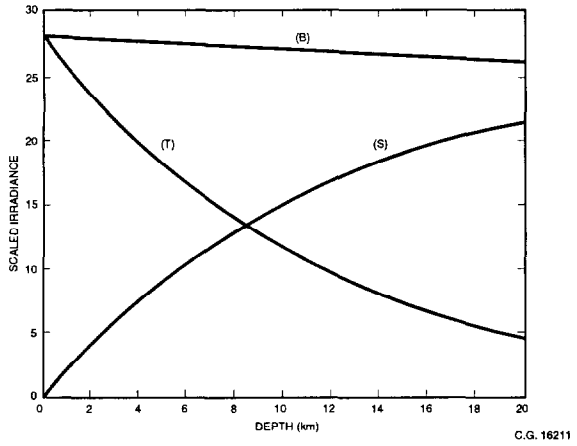

Figure 2 The Two Components of Equation (5)

The extinction coefficient β is, in general, a function of wavelength as is the radiance of the atmospheric particles.

However, for the purposes of image processing, the values of β and the airlight term $\left(\frac{DP}{\beta}\right)$ are assumed to be constant. The relationship between the irradiance at the sensor and the range may then be described using the three-parameter model:

$$J_{total} = c_0[1 + c_1 exp(-\beta R)] \qquad (6)$$

where $c_0 = DP\Omega_k$ and $c_1 = \frac{M_k \beta}{DP} - 1$

although estimates for c_0, c_1 and β are best obtained directly from the image data using a curve-fitting operation.

To do this, each pixel in the image is then assigned a range using knowledge of the scene and a range versus irradiance scatterplot generated which will exhibit a trend between received irradiance and distance from camera.

The curve fitted to the range versus irradiance data is a model which describes the theoretical relationship between the received flux P_k at the k_{th} pixel and its range R_k.

An example is shown in Figure 1. The solid line represents the model fit from (6). Note that typical sensors do not produce an output which corresponds directly to the irradiance. This is for two reasons. Firstly, they operate within a certain 'temperature' interval so that zero output does not, in general, correspond to zero flux. Secondly, the image data is Gamma-encoded for display on a video monitor. The relationship between the grey level P_k and the irradiance J_{total} may be described by the equation:

$$J_{total} = MP_k^{2.2} + c \qquad (7)$$

where M and c are constants, which are usually determined by the settings of the sensor controls. The values of these constants may be determined using a radiometric calibration procedure (see Section 4).

Once values have been obtained for c_0, c_1 and β, the enhanced image is produced using a simple pixel-by-pixel transformation:

$$J_e = (J_{total} - J_b) exp(+\beta R) \qquad (8)$$

where J_b is given by $c_0(1-exp(-\beta R))$.

The range R will, in general, be different for each pixel and so J_b must be re-evaluated for each pixel.

J_e is computed in this way for each pixel. Thus, the irradiance map is corrected for scattering and absorption by the atmosphere. The irradiance estimates thus produced are estimates for the irradiance values which would be measured if there were no degradation.

The corrected irradiance estimates may be suitably scaled and then Gamma-encoded for display on a conventional video monitor.

The different models corresponding to the various cases of c_1 being positive, negative and zero have been negotiated.

4. DATA GATHERING

4.1 Trials Requirements

The main requirement for the trial was that there must be a sufficiently large, scattered radiation component for the method to subtract.

In order to guarantee a reasonable amount of atmospheric attenuation, the visibility must be either poor or the target must be at a long range. Since poor visibility was not a certain condition during the period of the trial, a site was chosen with a large available range.

The site selected at Jennycliff, just outside Plymouth, was 38m above sea-level and allowed a maximum range of 22km out to the horizon at sea. Figure 3 below shows the area of the field of view. Other major advantages of the site were:

a) The Plymouth breakwater and the coastline from Kingsand to Penlee Point gave definite ranges.

b) Within the area, commercial and pleasure marine activity would give a number of targets at different ranges within the images.

c) The local weather station was virtually co-located with the trials site guaranteeing accurate meteorological data.

d) The site was conveniently located only a few kilometres from British Aerospace.

The cameras used in these trials were:

- Inframetrics Thermacam (3 - 5 micron)
- Steinheil Nanoquest (8 - 12 micron)
- Panasonic NV - M40B VHS (visible)

It would have been desirable to ensure that each camera was set with the same field of view and also that the three cameras were synchronised. Neither of these was possible. The magnification was set such that the field of view of each camera was as similar as possible, if conditions indicated that a reasonably large extinction coefficient was likely.

Radiometric calibration of the data was catered for by collecting the temperature of the sea, in addition to the air temperature.

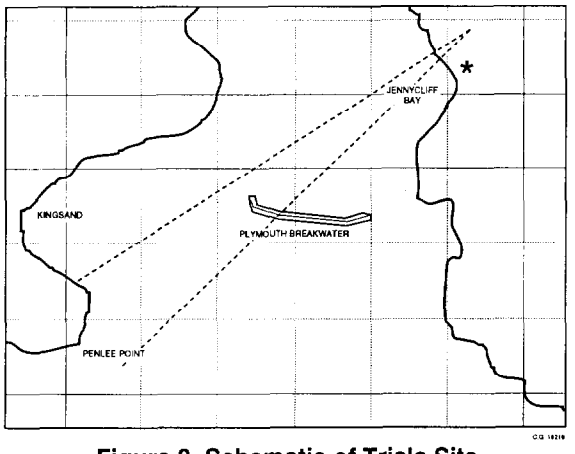

Figure 3 Schematic of Trials Site

4.2 Trials Data

Data was collected on five days between 24th July 1997 and 7th August 1997. For once the British weather did not let us down. Although an anticyclone was firmly established over the area, a large amount of moisture had been trapped. This gave conditions of high humidity (plus very misty conditions) resulting in visibilities ranging from 3km to 16km. It would have been useful if data out to the full range of 22km had emerged. Full met. data was obtained from the Meteorological Office. Visibility was assessed by the trials engineers, using the range cues of the breakwater and the shoreline. Sea temperature measurements were supplied by the Plymouth Marine laboratory. This was used as an aid to the calibration of the images.

From the large volume of data collected, sequences from three runs were selected for further analysis.

The main criteria for selection were the presence of a number of targets within the sequence and the desire to cover a good range of extinction coefficients. Contrast enhanced images were produced for each run.

5. RESULTS

5.1 Qualitative Improvements

As in previous work (1) a good preliminary test is to establish if the images appear to have better contrast after enhancement. A sample of raw and enhanced images for all wavebands are shown in pairs below in Figures 4A and 4B, Figures 4C and 4D and Figures 4E and 4F. In all the runs digitised and selected for further analysis and at all wavebands, visual inspection confirmed the contrast had improved. For a number of these runs (eg 4A and 4B), the enhanced images show targets not visible before the process.

5.2 Extinction coefficient

The meteorological data associated with each run enabled the extinction coefficient to be calculated using MODTRAN3. This value is compared with the value obtained as part of the contrast enhancement process. The results are shown in Table 1.

Table 1 Comparison of Extinction Coefficients Obtained from the Oakley Method and MODTRAN3

Date	Run No	Waveband	β_E Oakley	β_M MODTRAN
6/8/97	1	VISIBLE	-	0.359
	1	3-5 micron	0.173	0.162
	1	8-12 micron	0.213	0.221
7/8/97	1	VISIBLE	0.848	1.085
	1	3-5 micron	0.157	0.323
	1	8-12 micron	0.292	0.314
7/8/97	2	VISIBLE	0.741	0.780
	2	3-5 micron	0.154	0.291
	2	8-12 micron	0.307	0.305

For both the visible and the 8-12 micron wavebands, the extinction coefficient obtained from the Oakley method (β_E) and that obtained using MODTRAN3 (β_M) show very good agreement.

The results for 3-5 micron are poor in two of the three runs considered. However, in the MODTRAN3 results, the value of the parameter ICSTL, which controls the degree of continental aerosol influence, has been set at various values between 5 and 9.

It is possible that a better choice of ICSTL will enable the agreement of the extinction coefficients at 3-5 micron to be improved without too much degradation in the visible and 8-12 micron bands.

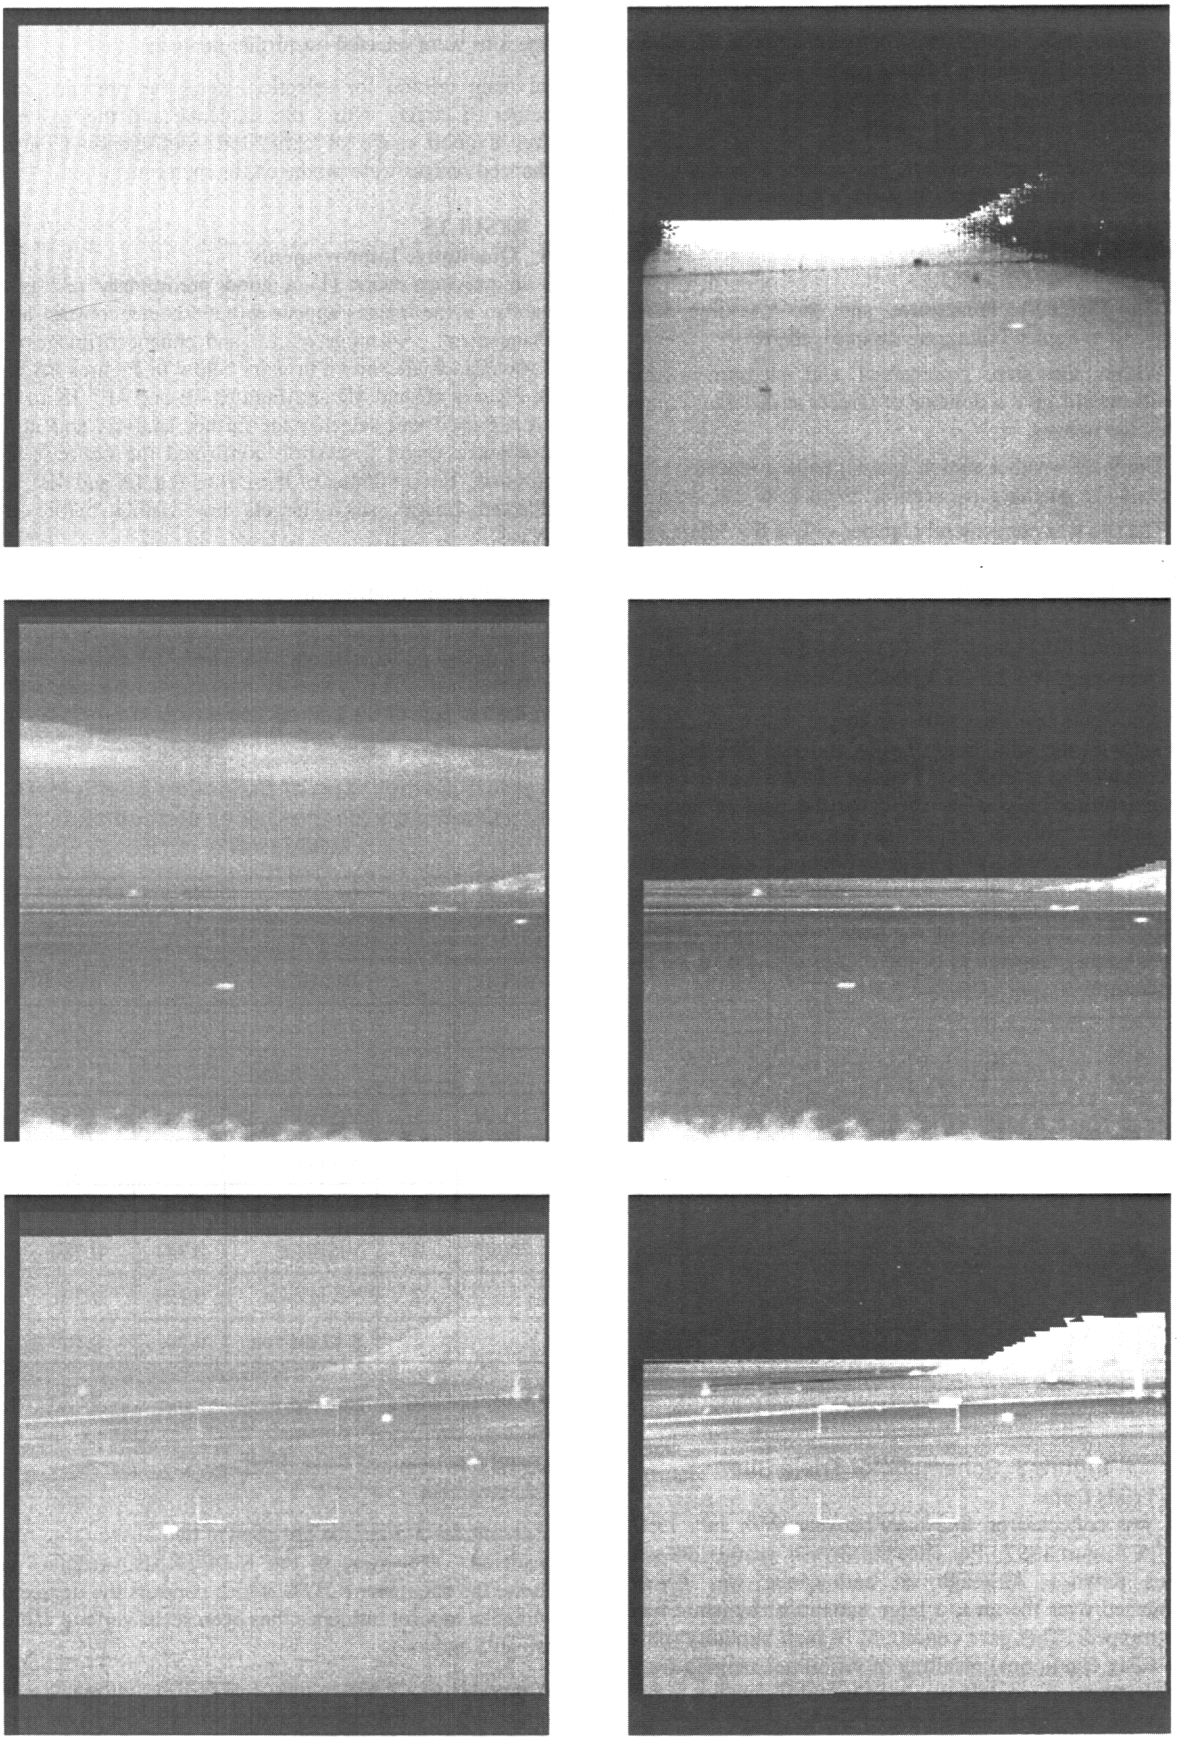

Figure 4E Original - 8-12 micron

Figure 4F Enhanced - 8-12 micron

Before and After Images for the Same Run Covering the Three Wavebands

5.3 Contrast Enhancement Metric

Since the aim of the Oakley algorithm is to improve contrast, it is an obvious choice to look at this metric first. The original form of the metric considered was very simple. Within the selected region, the maximum and minimum grey level values were determined for both the original image and the enhanced image. The contrast improvement (percent) C is given by:

$$C = \left(\left(\frac{E_{max} - E_{min}}{O_{max} - O_{min}} \right) - 1 \right) x\ 100 \qquad (9)$$

where E_{max} and E_{min} are the maximum and minimum grey levels respectively for the enhanced image and O_{max} and O_{min} the corresponding grey levels for the original image.

Though this metric should give a reasonable idea of the contrast improvement, it is rather crude. A better method is to sum the percentage improvement along each row in the defined region and then to average this over all the rows.

Thus, the average contrast improvement C_{av} is:

$$C_{av} = \frac{\sum_{r=1}^{n} C_r}{n} \qquad (10)$$

where C_r is the percentage improvement for row r (found by applying equation [9] to row r) and n is the total number of rows within the region. Results for C and C_{av} for a number of runs are given in Table 2.

The results obtained in Table 2 show the contrast improvement obtained over the whole image. The results show that, in general, the method gives significant improvements in contrast. For most cases, the contrast improvement measured over the whole picture shows a lower level of improvement. Two cases in the averaged case show an actual reduction in contrast. It is significant that they are both for the same run in the IR wavebands.

Table 2 Contrast Metric Results - Overall Picture

Date	Run No	Waveband	C	C_{av}
7/8/97	1	VISIBLE	392	377
	2	VISIBLE	49	70
6/8/97	1	3-5 micron	132	-12
7/8/97	1	3-5 micron	67	16
	2	3-5 micron	131	64
6/8/97	1	8-12 micron	44	-16
7/8/97	1	8-12 micron	95	30
	2	8-12 micron	49	70

In practice, the contrast improvement in the region of a target is of greater interest. To investigate this, one run was selected (7/8/97 Run 1) and the change in contrast improvement measured as the region of interest surrounding the target was increased.

During this process the box was kept centred on the target. It was also necessary to select targets which were well separated from other objects and also away from the border of the image. These results are shown in Table 3.

Table 3 Change in Contrast Metric as the Region of Interest Surrounding a Target is Increased

Waveband	Small Box	Medium Box	Large Box
VISIBLE	196	99	58
VISIBLE	115	105	94
3 - 5 micron	- 4	- 9	- 12
3 - 5 micron	137	105	104
8 - 12 micron	26	6	- 2
8 - 12 micron	203	165	119

These results show that, for individual targets, a similar pattern of contrast improvement results. This improvement is maximised when the bounding box around the target is small, ie the target dominates the area.

5.4 Signal-to-Noise Ratio

The signal to noise ratio (SNR) metric is the most commonly used measure of the quality of an image. The reason for this is that it is simple to implement and gives a result that is easy to interpret.

A high SNR is an intuitive indication of a clear image. The SNR formula used for this report is as follows:

$$SNR = \frac{Average\ of\ Target - Average\ of\ Local\ Background}{\sigma\ of\ Local\ Background} \qquad (11)$$

where 'average' and 'σ' refer to the mean and standard deviation of the pixel grey levels respectively.

This formula is implemented by defining a target within an image and then defining a local background area. The target is defined by hand picking pixels within an image which the user perceives as belonging to a target. The hand-picked target pixels are then assigned a bounding box which is defined as the rectangle which clips the extremes of the target in the horizontal and vertical directions. The local background area is then defined as the area between the target bounding box and another box which is a number of times larger than the bounding box of the target. For this project, the background box is scaled up by a factor of three on the target bounding box.

The results of the SNR calculation using the calculations described above are shown below in Table 4. The objects are taken from a single run but for the three wavebands. As expected, the process of contrast enhancement makes little difference to the SNR as defined in this paper. In the visible waveband, some improvement can be seen but, for the IR wavebands, little or no change is seen.

Table 4 7th August 1997 Run 1 - SNR Results for Objects in the Three Wavebands

Waveband	Original	Averaged	Enhanced
VISIBLE	6.29	8.07	11.51
VISIBLE	1.91	4.32	5.60
3-5 micron	12.60	13.44	12.41
3-5 micron	3.51	3.43	3.68
8-12 micron	8.58	8.64	8.51
8-12 micron	2.66	2.59	2.32

5.5 Automatic Segmentation

The most significant image processing application is that of Automatic Tracking and Recognition (ATR). Clearly, if the tracking of a particular target in an enhanced image is better than in the original, this suggests that the enhanced image is an improvement on the original. The automatic segmentation metric used here is a simple method of quantifying the trackability of an image.

The tracking system consists of two stages. The first stage is to segment the input image and the second stage (the tracking stage) attempts to follow targets in the segmented image. The latter process does not change, in the sense that the algorithms are identical, no matter what input is used. Therefore, the tracking quality is controlled entirely by the properties of the segmented image. The trick is to measure how useful the segmented image is to a tracking application.

The object of the segmentation process is to form an image in which possible targets are separated from the background. Usually the output is an image in which anything not considered as a target is black, ie zero brightness. Therefore, possible targets can be picked out easily by the tracker because they are simply the non-zero pixels. The quality of the segmented image can be measured by checking to see how well the automatically designated targets match up with real targets.

The problem of how to define the real targets remains. Fortunately, the human brain is extremely good at picking up targets within an image. On the whole, this is due to the huge target database that the brain contains which enables it to recognise objects within the image as targets. Further to recognising the objects, the brain can extrapolate to decide which pixels actually make up the target even though noise in the image may make it virtually impossible to see the edges of the object. This means that a target can be hand designated by eye with a great deal of accuracy.

The method used for this investigation is as follows. Firstly, a target in an image is designated by selecting individual pixels. The output from this process is a simple binary image in which the designated target pixels appear white and the rest of the image is black. The next step is to pass the image through an automatic segmentation process which will produce an image in which many targets can be identified. Finally, the hand picked or 'real' targets and the automatically selected targets can be compared.

The segmentation technique used in this project was the same as that used as part of many BASE tracking products.

A quality assessment is done by simply finding any automatically selected targets that coincide with the real target and calculating three values for each one in turn.

The three values calculated are as follows:

- The number of pixels designated as target pixels by both processes.
- The number of pixels designated as the real target but not selected by the automatic segmentation process.
- The number of pixels not designated as the real target but were nonetheless selected by the automatic process. This is known as the number of false alarms.

These three values are illustrated in Figure 6 below.

Table 5 below shows the results obtained for various targets in three of the runs.

Table 5 Automatic Segmentation Before and After Contrast Enhancement

Raw Image			Enhanced Image		
Missed	Both	False Alarms	Missed	Both	False Alarms
2	170	54	0	172	25
0	103	102	0	103	14
0	87	68	0	87	34

The results shown in Table 5 are most encouraging. After enhancement, the 'false alarm' pixels for individual targets within an image are considerably reduced. This will enable better detection, identification and tracking to take place

6. DISCUSSION AND CONCLUSIONS

The aim of this work was to collect data and identify the improvement obtained by the Oakley method in a quantifiable manner. The main findings are these:

- In most cases the overall contrast of an image is improved in all wavebands. These improvements are large in a number of cases.
- SNR is relatively unaffected by contrast enhancement. This is to be expected as changes in contrast across the whole image will have little effect.
- Contrast enhancement tends to improve the performance of image segmentation. This is shown in the considerably lower number of 'false alarm' regions produced after enhancement.
- Contrast Enhancement is a valuable process and should aid the overall tracking process. It is particularly useful in adverse weather conditions. The estimation of the extinction coefficient is an interesting by-product of this algorithm.

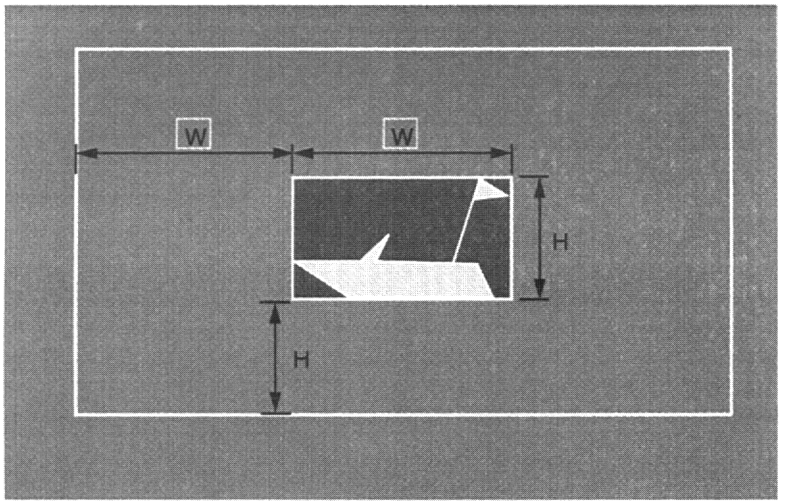

C.G. 16218

Figure 5 The Target Bounding Box and Local Background Area used in the SNR Formula (11)

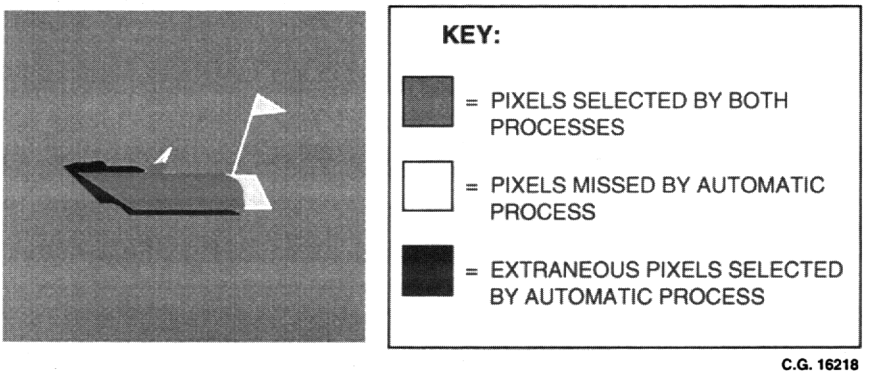

C.G. 16218

Figure 6 Parameters Used to Evaluate the Quality of the Segmented Image

7. FUTURE WORK

Four main activities are planned for the future:

a) The current simulation is to be linked into a digital terrain database. This will enable the range of each pixel within the image to be calculated in a simple fashion.

b) Work will be carried out to establish if this method aids in the classification of objects. The main features used in our current classification schemes will be examined and a suitable metric devised.

c) The effect of frame averaging on the various metrics will be examined.

d) The design and construction of a hardware demonstrator.

8. REFERENCES

1. J R Oakley and B L Satherly IEEE Transactions on Image Processing - Vol 7 No 2 February 1988 pp 167-179.
2. Patent Application PCT/GB96/01657 Image Enhancement.
International Filing Date 12th July 1996.

9. ACKNOWLEDGEMENTS

The authors would like to thank Mike Everett and Dave Parker of British Aerospace Military Aircraft and Aerostructures for their help and encouragement.