A comparison of some techniques for the subpixel location of discrete target images

M. R. Shortis[1], T. A. Clarke[2], & T. Short[3].

Department of Surveying and Land Information,
University of Melbourne, Parkville 3052, AUSTRALIA[1]

Department of Electrical, Electronic, and Information Engineering[2],
Department of Civil Engineering[3],
City University, Northampton Square, London EC1V 0HB. UK.

ABSTRACT

Signalising points of interest on the object to be measured is a reliable and common method of achieving optimum target location accuracy for many high precision measurement tasks. In photogrammetric metrology, images of the targets originate from photographs and CCD cameras. Regardless of whether the photographs are scanned or the digital images are captured directly, the overall accuracy of the technique is partly dependent on the precise and accurate location of the target images. However, it is often not clear which technique to choose for a particular task, or what are the significant sources of error. The research described in this paper describes aspects of target recognition, thesholding, and location. The results of a series of simulation experiments are used to analyse the performance of subpixel target location techniques such as: centroiding; Gaussian shape fitting; and ellipse fitting, under varying conditions.

Keywords: Targets, subpixel, target location, target recognition, threshold.

1. INTRODUCTION

A key area of photogrammetric measurement where reliability is essential is in target recognition and location. All the subsequent processes in high accuracy 3-D measurement rely on the results of these initial image processing stages. While many methods have been developed for this purpose it is doubtful whether all aspects of this process have been totally understood. Hence, this paper considers several differing algorithms for the recognition and location of target images. While this study is not exhaustive, several novel aspects are considered such as an alternative approach to thresholding, a scan line method of target recognition, a comparison of least squares and centroiding methods of target location, and an analysis of the effect of: threshold; size of target; quantization; DC offset; and saturation. The reason for this area requiring further work is that in the measurement environment there are many circumstances where targets have less than ideal characteristics. For instance, it is not always possible to use retro-reflective targets and so problems are encountered with uneven background illumination. Large depth ranges within objects to be measured may give rise to extreme variations in the sizes of target images. Under these conditions it can be shown that a penalty of lower accuracy of target location must be paid. The pertinent decision here is which target location algorithms to use for a range of conditions varying from favourable to adverse.

2. TARGET RECOGNITION AND LOCATION ALGORITHMS.

There are two tasks to be performed in the process of computing the subpixel locations of target images: recognition and location. The detection of the target images is required to unambiguously identify targets within a scene. The location of the target image is generally a second process which precisely and accurately determines the target image centre within the digital image frame. There have been many schemes proposed for the automated detection and measurement of targets in photogrammetric images. Many of these have concentrated on the unambiguous separation of target images from a cluttered background using a combination of patterned targets and filtering of the image (van den Heuvel et al., 1992; Wong et al., 1988). These techniques rely on sufficient contrast between the target and background, or specific patterns which are unlikely to be accidentally replicated by the background features combined with perspective distortion. Included in the patterns may be embedded target identifiers, which allow corresponding target images in different frames to be associated with a unique target. Once the targets are identified, the target images can be precisely located using edge fitting routines (Zhou, 1986) or least squares matching (Lemmens, 1988). Least squares matching can be used quite independently based on templates and searching the image for a match which satisfies predetermined criteria. The template may be artificially generated as an ideal image, or simply extracted from a typical view. However, this method becomes computationally inefficient unless there are sufficient constraints on the geometry (Gruen and Baltsavias, 1988) or the approximate locations of the target images are known from prior knowledge (Haggren and Haajanen, 1990).

An alternative strategy is to enable the detection of the target images simply by making them the brightest, or darkest, objects in the image. This strategy is adopted from that used in industrial metrology, where controlled lighting conditions and retro-reflective targets are used to good effect (Brown, 1984). In most circumstances the background clutter can be reduced dramatically by under-exposure, as the background is often of little interest other than a guide to target identification. In stable environments, the influence of the background can be virtually excluded using image subtraction of frames, first with and then without illumination of the retro-targets. High contrast, passive targets can also be used, however the maximum possible range in intensities between the targets and the background must be sought. A global threshold value, set interactively by an operator or automatically by analysis of an intensity histogram, can be used to segment the entire digital image frame and isolate the target images. The digital image may retain grey value intensities above the threshold, or be binarized for faster processing. Target detection can then be carried out using a number of schemes, such as scan-line searches on binary images (Clarke et al., 1993) or image correlation on grey scale images (Shortis et al., 1994). Because the target image can be segmented from the background with relative ease, centroiding is most commonly used to subsequently measure the location of the image (Trinder, 1989; West and Clarke, 1990).

The task of image location must take into account the two types of image that will commonly be encountered, those with a uniform background and illumination, or possibly no background illumination at all, or images where not only is there a non-uniform background, but there are also extraneous features in the image. A worst case example of this is when another feature, such as that given by a light bulb, mimics the same characteristics as a legitimate target (Figure 4). This case in not as difficult to resolve as might be imagined, as provided a good correspondence algorithm is used, it will merely result in the 3-D co-ordinates of the light bulb. Hence, if there is no solution to the problem of extraneous background illumination, then the task of recognition is somewhat more difficult. The second task, the subpixel location of the target, would appear to be more straightforward, but in practise a number of effects can cause both systematic and random errors. An example of a systematic error is that caused by poor thresholding or even in some cases the use of a threshold at all, while a random error may be caused by noise that is allowed to affect the target location. Several three dimensional graphs of targets from some typical imagery are illustrated in figures 1 - 6.

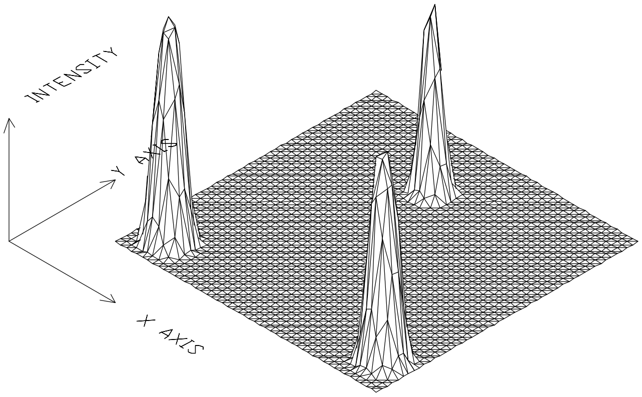

Figure 1. Ideal retro-reflective targets.

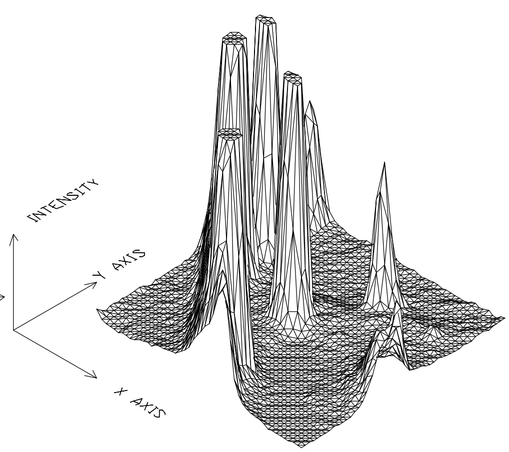

Figure 2. An example of target density and complexity.

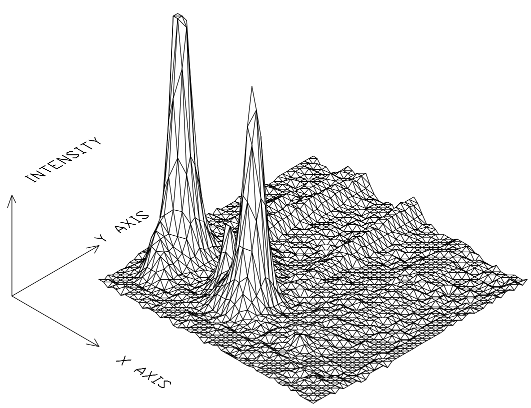

Figure 3. Targets from an industrial setting.

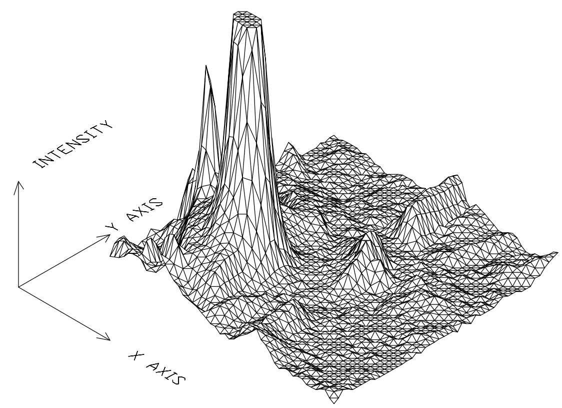

Figure 4. An image of a light bulb.

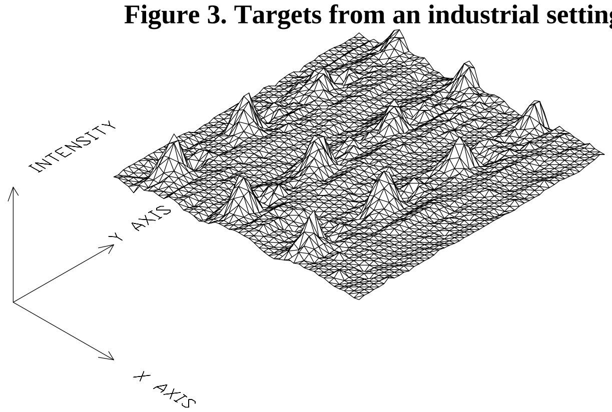

Figure 5. Low contrast targets from a research setting.

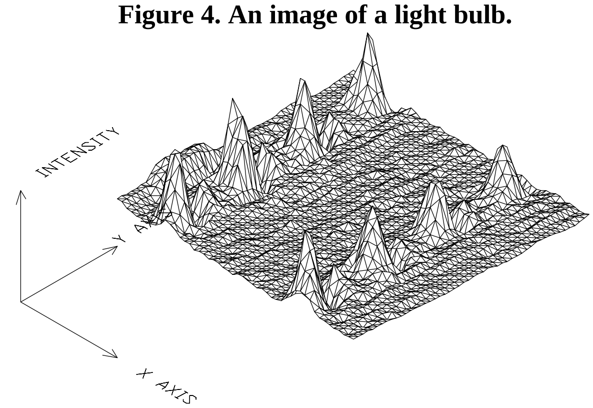

Figure 6. The problem of noise and ringing.

These target images show the wide degree of variability between target images. The targets illustrated in figure 1 are of targets where the size of the target and the illumination has been optimised using a Pulnix TM6CN camera, whereas in figure 2 the same care has not been taken and overlapping targets are a possibility. Figures 3, and 4 are from a Kodak DCS 200. Figure 5 illustrates how poor targets can be in terms of contrast and background noise. These images were collected by a colour camera from within a rotating device via slip rings and recorded onto a video tape. Figure 6 is from the same source as figure 5 and clearly shows the problem of ringing which, for some target recognition processes, can produce a spurious target image close to the legitimate target.

3. TARGET RECOGNITION

Target recognition within a digital image frame can be achieved in a number of ways. Manual selection using a mouse and cursor is labour intensive and prone to human error, so the acquisition of target images is more efficiently achieved by semi- or fully automated procedures. An example of a semi-automated procedure is the resection of the frame using a few manually selected target images, followed by an automated "drive back" to all other images based on known or assumed object locations for the targets (Shortis et al, 1991). Under the circumstances that target locations are not available, techniques which do not require prior knowledge must be employed. Examples of such automated procedures, such as segmentation and filtering, have already been mentioned, however without an embedded coding system additional steps must be taken to identify corresponding target images on different frames (Chen and Clarke, 1992). Notwithstanding these restrictions, the clear advantages of fully automated procedures, and to a slightly lesser degree the semi-automated procedures, are the reliability and completeness of the acquired target image sets.

An automated method has been developed which is an extension of the technique reported by Clarke et al, (1993). In this case a recursive filling algorithm was used to both find the intensity values within a target that were above a threshold and at the same time compute the centre of gravity of the intensity values. It was noted that although this algorithm was reliable, it would be inefficient in larger images due the fact that pixels would often be visited four times during the course of the filling procedure. This method has been revised by using a scan-line filling method which is more efficient for larger targets, or for the larger sections of images that must be recognised as not being targets. The first stage in the filling process is the tracing of the outline of any object above the designated threshold (usually determined interactively by a visual segmentation or by knowledge of the imaging set up). The objects are determined by scanning each line in the image from top to bottom. When a pixel is found that is above the threshold value tracing is performed by a eight-way search to visit all points which are on the periphery of the object image and above the threshold. The image is searched from top to bottom until all of the candidate targets are found. When a candidate target is found it is temporarily removed from the image and the search continued. At this point the perimeter length and maximum **x** and maximum **y** dimensions are known. These data are stored in an structure and passed to the next part of the program. An ellipse fit algorithm can be used to provide the subpixel location of the feature. The reported statistics provide a means to accept or reject the object as a legitimate target. In the latter case while it is simple to fill a circular object, the algorithm must also be able to fill irregular shaped objects to remove them from consideration. Hence the situations illustrated in figure 7 must be considered where care must be taken to fill the section of the object to the right of the pixel marked with a cross.

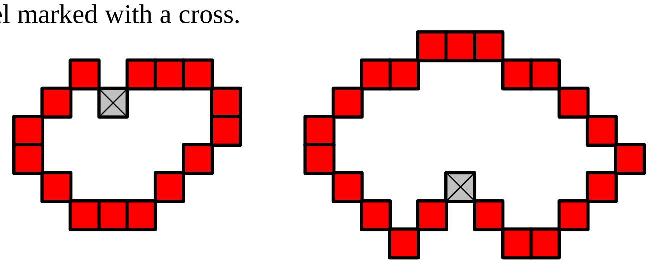

Figure 7. Cases where filling is not straight forwards.

The filling is completed by taking into account the singular points and filling between pairs of perimeter co-ordinates from the top of the feature to the bottom. At the end of this procedure a list of co-ordinates is compiled which contains all the locations of the feature. This information is finally used in the next algorithm which may be a binary centroid, normal centroid, squared centroid, or least squares template matching. The advantage of the filling method over a straight thresholding method is that there is no possibility of any extraneous intensity values which are above the threshold coming into the target location algorithm with potentially major effects.

4. TARGET LOCATION

Once the target images are identified, a second computation procedure is generally required to locate, to subpixel accuracy, the centre of the target image within the digital image frame. This second computation can be broken up into the pre-processing phase and the actual centre calculation. The pre-processing tasks are described in the next sections, followed by a description and comparison of various target image centre calculation techniques.

4.1 Pre-processing

4.1.1 Threshold removal

The first step in the target location procedure is generally the subtraction of a threshold intensity value. The threshold subtraction is based on the supposition that there will always be background noise in any digital image. The background noise may be caused by unwanted background illumination within the scene, low level electronic effects (Beyer, 1993), re-sampling effects from image recording media (Shortis et al, 1993) or, most probably, a combination of these noise contributors. To isolate the signal of the target image, this background noise may need to be removed. Template and least squares matching techniques may not require threshold value subtraction. If the template is cut from a typical digital image frame then the noise will be effectively incorporated. An ideal image template which is artificially generated may have artificial noise overlaid on the signal. The most common technique is, however, to model the noise in the least squares matching paradigm (Gruen and Baltsavias, 1988). Hence, threshold subtraction is not appropriate as the noise is statistically modelled and the effects incorporated in the precision of the match. If least squares matching is not used and the location based on signal only, the removal of noise is necessary and the critical issue is the level at which the threshold is set. A threshold which is too low will allow noise to unduly effect the locations of target images. A threshold which is too high will eliminate some of the signal from the target images. In either case the locations of the target images will be effected by either systematic errors or a falsely inflated precision of location.

As previously discussed, global threshold values can be set as part of the process to recognise target images, particularly images of retro-reflective or high contrast targets. However, threshold values for target location are typically set locally, within a window into the digital image frame which surrounds the target image. Rather than a pre-set value for all images, the threshold is set dynamically, based on local conditions within the window, which allows the threshold to adapt to different illumination conditions within the scene (Zhou 1990). Local threshold values can be set by a number of techniques. One of the earliest threshold values tested was the average of the minimum and mean intensities in the window, which was successfully used with passive targets in ambient lighting (Wong and Wei-Hsin, 1986). This threshold value considers all pixels within the window, and is an estimator of the minimum of the expected bi-modal distribution of pixel intensities. Computation and analysis of the actual histogram of pixel intensities within the window would reveal a more accurate threshold, however the computational overhead would be substantially increased. The alternative to the consideration of all pixels in the window is to evaluate only pixels at the edge of the window. Here, the assumption is made that the target is centred in the window and therefore the edge pixels are representative of the background noise. This has the advantage that the uncertainty of the transition zone between the noise and signal at the periphery of the target image has no effect on the threshold computation. The disadvantage of the technique is that additional steps must be taken when the image is not centred in the window and intrudes into the window edge.

Threshold algorithms using window edge pixels vary in complexity. The simplest method is to add an arbitrary constant, usually a few grey levels, to the maximum intensity at the window edge (Snow et al, 1993). This technique ensures that all background noise will be eliminated, but may remove some of the target signal. Alternatively, an amount proportional to the peak intensity in the window can be added, under the assumption that this is representative of the overall brightness of the image. A more complex mechanism is a statistical analysis of the distribution of the intensities of the edge pixels. Given the assumption of random background noise, the mean and standard deviation of the intensities can be computed. A threshold value can be set by then calculating a critical value for the maximum expected intensity value, based on a suitable confidence level. The statistical method requires a global limit to be set, or iteration of the threshold and centring of the window on the image. Without these precautions, falsely high threshold values may be computed under the influence of a target image intrusion into the window edge.

4.1.2 Blob Testing

All algorithms for the determination of threshold values for high contrast targets run the risk of outlier pixels. Such pixels have intensity values above the threshold, but are not part of the main "blob" of the target image. Outlier pixels are caused by unexpectedly high noise, intrusions into the window by other target images or extraneous background detail. The removal of outlier pixels is dependent on the assumption that the window is centred on the target image and that the target centre is known at least approximately. The initial location may be obtained from a simple centroid calculation for example. The outlier removal process is governed by the detection of the edge of the target blob. The edge detection criterion is the first value of pixel intensity below the threshold (or a zero value if the threshold has already been subtracted). Scanning the pixels outward from the target image centre, the target image edge will be detected, assuming the window size is sufficiently large. Once the image edge is detected, all subsequent pixels encountered outward to the window edge are assumed to be non-blob pixels and are set to zero intensity. Scanning for outliers is most effective if it is carried out in two perpendicular directions, typically along and across the image scan lines. Removal of the outlier pixels is essential to eliminate bias in the image location. The scanning technique will successfully remove the outliers in the vast majority of cases. The technique will fail if there are major intrusions into the window, commonly caused by target images in close proximity, or if there are discontinuities in the image blob, such as a target which has been partially occluded. In the former case an iterative solution with a reducing window size may eliminate the intrusion, in the latter case the target image should be rejected as a valid location.

4.1.3 Geometry Testing

The image blob can subsequently be tested for the correct shape using a series of straightforward geometry tests. Once more, these tests are prefaced on knowledge of the location and extent of the target image within the window, as well as prior knowledge of the expected size and shape of target images.

The simplest test is a size range criterion, which will reject targets below a minimum size and above a maximum size. The ratio of the extents in two perpendicular directions, using either the scan line axis or the line of maximum extent as one of the directions, can also be tested. Knowledge of the extents of the actual targets, the image scale and the expected perspective distortion are required to set realistic criteria against which to test the size range and extent ratio. These tests will generally eliminate spurious targets such as background light sources (see figure 4), specular reflections and target labels or identification tags. The ratio of the number of non-zero intensity pixels to the area of the smallest possible enclosing window of the target image can also be tested against an expected range. Often called the black-white ratio test, for circular targets the expected value is one quarter of pi. However due to the effects of sampling by discrete pixels and perspective distortion, a range of acceptable values must be pre-determined from knowledge of the imagery and prior experience. Similar tests, such as the ratio of the perimeter to the area of the target image blob, are all designed to eliminate target images which are not of the correct shape, but nevertheless are acceptable to the range and extent criteria. The blob and geometry tests are not applicable to least squares matching if high contrast targets are not being used. Templates which are extracted from a typical image, to be matched to other images, include the local background and any imperfections will either degrade the precision of the match or reject the location as a potential target image.

4.2 Description of the location methods.

A number of methods have can be used to find the subpixel location of a target image. However, they have rarely been compared. Each algorithm has characteristics which it is necessary to know about if they are to be used in practise. A simulation is a good way of discovering the relative merits of each of the methods. In this paper the methods are analysed by varying the size, shape and quantization levels of a perfect Gaussian shape target. While it would be advantageous to include in this simulation various sources of noise other than quantization, nevertheless the optimum characteristics of each method can be analysed. The simulation provides the raw image data and when necessary the perimeter of the target image which is above a pre-selected threshold.

4.2.1 Average of perimeter.

This simple method averages the co-ordinates of the perimeter of the target image chosen with reference to a pre-selected threshold. The equation is given by:

$$\bar{x} = 1/n \cdot \sum_{i=1}^{n} x_i \qquad (1)$$

where: $\boldsymbol{x_i}$ are the $\boldsymbol{x}$ co-ordinates; $\boldsymbol{n}$ is the number of co-ordinates; and $\bar{x}$ is the subpixel location of the target.

4.2.2 Binary centroid.

The equation for a binary centroid is given by:

$$\bar{x} = \sum_{j=1}^{m} \sum_{i=1}^{n} i \cdot I_{i,j} \,/\, \sum_{j=1}^{m} \sum_{i=1}^{n} I_{i,j} \qquad (2)$$

where $I_{i,j}$ is one or zero depending on the threshold and the intensity at the ***i,j*** th. pixel location; and **m** and **n** are the dimensions of the window in which the centroid is being computed.

4.2.3 Grey scale centroid.

The equation for a grey scale centroid is the same as for (2) with no threshold applied to the intensity values.

4.2.4 Squared grey scale centroid.

The equation for a squared centroid is the same as for 4.2.3 but with the intensity values squared.

4.2.5 Ellipse Fitting

In this method an attempt is made to find the least-squares fit of an arbitrary ellipse to a set of points considered to lie on the perimeter of a target. These perimeter points are found by the recognition algorithm described in section 3. The five parameters determined by the least-squares estimation routine are the ellipse centre, its semi-major and semi-minor axes, and its rotation, theta, in the **x,y**-plane. This final parameter can be difficult to determine when the data is nearer to circular than elliptical in shape, which produces a small element at the corresponding position in the normal coefficients matrix. When necessary theta is set to zero and removed from the solution. Starting parameters for the solution are computed firstly from a simple centroid to determine the ellipse centre ($\mathbf{x_0}$, $\mathbf{y_0}$), and then from a second moment to find the axes and the rotation.

4.2.6 Gaussian Distribution Fitting

In a similar approach to the ellipse fitting routine, an array of grey-scale values are taken to approximate a 2-D Gaussian (normal) distribution of the form:

$$I = f(x,y) = \frac{K}{2ps_x s_y \sqrt{(1-r^2)}} \exp\{\frac{-1}{2(1-r^2)}[(x/s_x)^2 - 2r(x/s_x)(y/s_y) + (y/s_y)^2]\} \qquad (3)$$

The six parameters are :
- x, y – the deviations from the mean
- s_x, s_y - the standard deviations
- r - the correlation coefficient
- K - the scaling parameter

The scaling factor, K, is used to modify the basic surface equation, whose volume is unity, to the figure determined by the magnitudes of the grey scales. Simple intensity weighted means and standard deviations of the population provide good starting values for the parameters in the iterative adjustment. The correlation is normally set to zero, and the scaling factor is determined by the sum of the grey-values. The correlation coefficient plays a comparable role to the rotation in the ellipse routine and is removed from the solution if required.

5. SIMULATION TRIALS

5.1 Simulation trials.

The simulation program creates a perfect target image of known, but variable, size and shape. The location of this target image is then moved by known subpixel amounts within a defined limit. At each location the target is quantized. Each algorithm is then applied to the resulting target image and the difference between the known target location and the reported location recorded. It should be noted that although it is sufficient to analyse the results on one orthogonal direction because of the symmetry of the image, it is necessary to displace the target in the other direction to obtain all of the possible permutations of intensity values. To compare results the root mean square (Rms.) error is computed and the errors are plotted to assess whether any systematic effects are present.

The algorithms that were tested in this manner were the: average perimeter co-ordinate mean; centroid of binary intensity; centroid of intensity; centroid of squared intensity; least squares ellipse fitting; and least squares Gaussian shape fitting. It should be noted that it is not possible to compare the performance of pattern matching techniques, where patches are taken from more than one image, in the same manner. The relationships between target size and location precision as well as between target peak intensity and location precision have been studied elsewhere (Trinder 1989; Clarke et al. 1993) revealing that for the centroid method there is an optimum target size beyond which little improvement in location precision is achieved. Furthermore, the location precision is approximately doubled for a doubling of peak intensity. For the squared centroid method, the optimum size is larger than for the centroid method at which point an improvement in the precision is observed. The corresponding features of the other methods have not been tested in comparison with these two methods and so are considered in this paper.

5.2 The effect of quantization.

Simulation trials were conducted for an ideal Gaussian shaped target with a standard deviation, or sigma, in both x and y of 1.25 and a threshold of one. The location of the target was placed into 2,500 known positions and the Rms. error computed for each method using the same shifted target image in each case. The results of these trials are illustrated in figure 8. The results indicate that the performance of the Gaussian fit, the centroid, and the squared centroid are broadly similar with a clear trend toward better location accuracy with more quantization levels. The ellipse fit, the average co-ordinates, and the binary centroid are also similar to each other, but in this case there is no clear trend. The lack of an improvement with greater quantization levels is to be expected as there will not be a great change in the size and shape of the target periphery upon which all of the targets rely for their location accuracy. The Rms. error of the intensity based methods is between twice as good as the methods based on edge threshold and ten times as good.

5.3 The effect of threshold variations.

A simulated target of sigma of 1.25, a peak intensity of 128 was selected and 2,500 locations were computed for each method. Again the characteristics of the methods fall into two clear groups (figure 9), those which use the intensity of the target image and those that do not. In the case of the binary, average, and ellipse methods a trend towards poorer target location exists which may be explained by the gradual decrease in the size of the target which, in the case of the ellipse, appears to have a significant effect. For the intensity based methods the centroid method decreases in location accuracy as the threshold is increasing. This is consistent with the fact that imbalances in the image geometry will exist as pixels are progressively removed from the periphery of the target image. The squared centroid method shows a lesser decrease in accuracy in line with its lack of emphasis on the lower intensity values. The Gaussian fit method does not appear to be affected by the change in threshold level making it more suitable for situations where a sub-optimal thresholding is likely to take place for each target image.

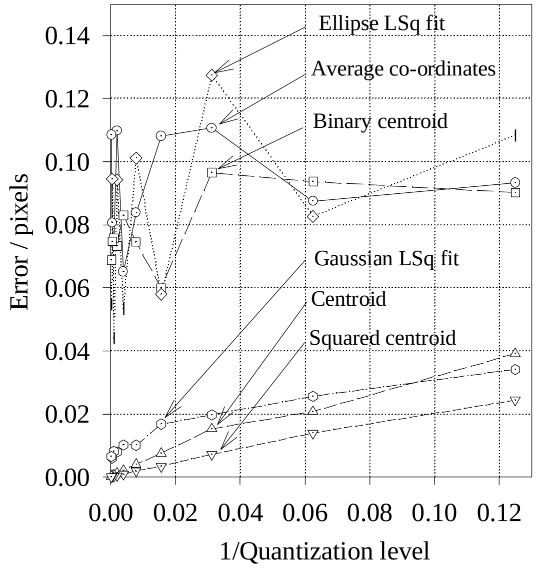

Figure 8. The effect of quantization on target image location.

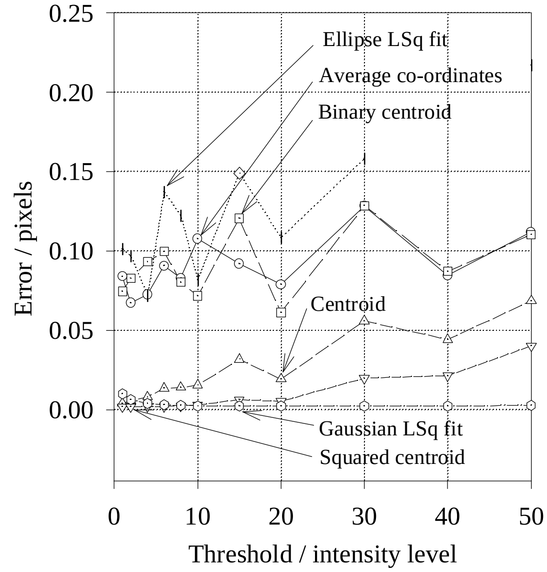

Figure 9. The effect of varying the threshold of the target images.

5.4 The effect of target size on target location.

A simulated target with a sigma of 1.25 and a peak intensity of 128 was used to analyse the effect of change in size of the target image. The results are illustrated in figure 10. These trials indicate that for two of the methods which do not use the intensity information, the ellipse fit and the binary centroid, a decrease in the error of location occurs for an increase in the target size. This is consistent with the fact that for each of these methods there is more information on which to base the location. However, the average co-ordinates method does not appear to have the same improvement in location precision. For the intensity based methods there appears to be little improvement in location accuracy with increase in size of target with the best results being obtained by the squared centroid method followed by the centroid and then the Gaussian fitting method. This confirms the results obtained by Clarke et al. (1993) that with respect of quantization alone there is no reason why targets above a sigma of about 1.25 need be used.

5.5 The effect of saturation on target location accuracy.

In many cases target images will have central areas which are at the saturation level of intensity. This may be caused by the target being too bright for the sensor because of incorrect exposure, variable lighting conditions or a large range of depth in the object being imaged. The effect of saturation was analysed by varying the peak intensity from 128 to 290 whilst applying a maximum intensity limit of 128 to the simulated target image, effectively cutting off the top of the intensity profile (figure 11). For the non-intensity binary centroid and ellipse fitting methods there is an initial improvement in location accuracy followed by a worsening in location accuracy. While the initial improvement may be attributed to a marginal increase in the size of the target with peak intensity, the subsequent worsening is not so easily explained. The location accuracy of the intensity based methods shows a general worsening in location accuracy with increased saturation of the image, however, in this case the centroid method performs better than the squared centroid method and the Gaussian fit method. The squared centroid method would be expected to give this result as the bias of the location computation is with the higher intensity values which are in this case all equal. The edge pixels would not be highly weighted in the computation but, in this case, contain the significant location information.

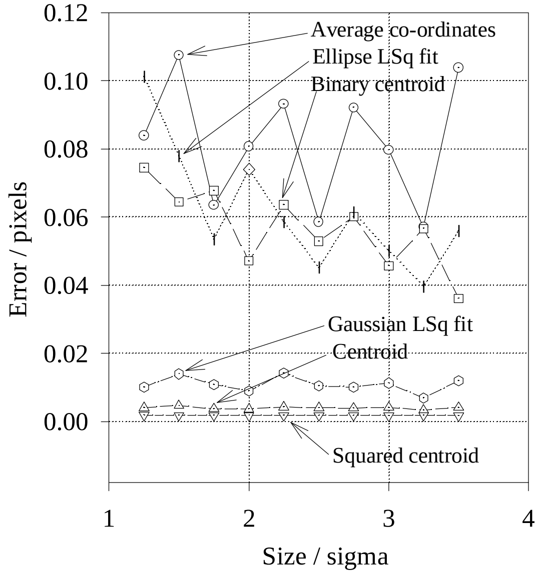

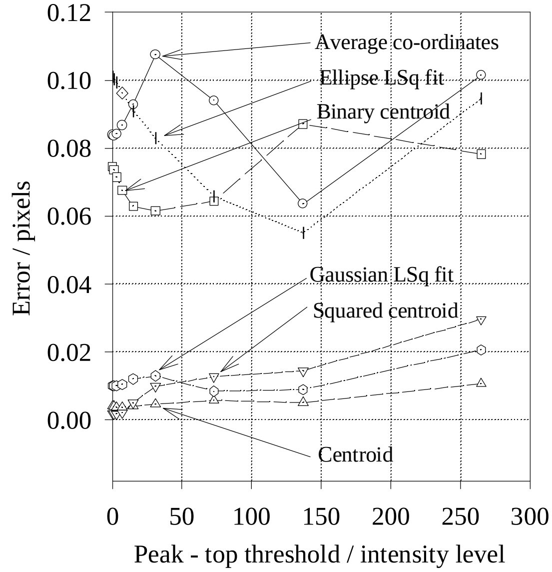

Figure 10. The effect of target size on location accuracy. **Figure 11. The effect of saturation on target location accuracy.**

5.6 The effect of DC offset on target location accuracy.

It has been shown by Clarke (1994) that a DC offset to the input to the A-D converter on a frame grabber will adversely affect the accuracy of target location. In this simulation experiment varying levels of DC offset were subtracted from the target image which had a peak height of 128 grey levels and a sigma of 1.25. The results are illustrated in figure 12.

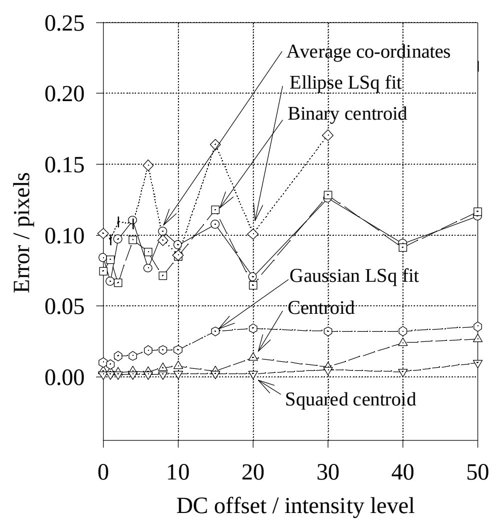

Figure 12. The effect of DC offset on target location.

The degrading effect of the DC offset does not significantly affect the thresholded methods beyond the effect caused by the decrease in target size as the threshold gets larger. The squared centroid method performs most accurately, followed by the centroid and the Gaussian least squares fit methods.

5.7 Conclusions.

It is clear from these simulation trials that each algorithm has its own characteristics. However, it may be generalised that intensity based algorithms are better than the threshold based algorithms. Furthermore, for the most accurate results the following rules would appear to apply.

> (i) If small target images are used there is no advantage in using a target image with a sigma of greater than 1.25, however care must be taken to avoid saturated or flat top target images, as in this case the squared centroid or Gaussian fitting method will not perform as well as a centroid. For the binary centroid and ellipse fitting methods the accuracy improves with the size of the target image.
>
> (ii) With respect to quantization alone each extra quantization bit results in an improvement in target location accuracy for all intensity based methods. This is not the case for the threshold methods.
>
> (iii)For target images which require background noise removal via a threshold subtraction, there is a degradation in location performance for the centroid method that is significant. This effect is smaller for the squared centroid method and non-detectable for the Gaussian fit method.
>
> (iv) For large targets the binary centroid and ellipse methods improve their location accuracy with increased size.
>
> (v) A DC offset caused by an incorrect setting of a frame grabber will result in a degradation of target location performance which is most noticeable in the centroid and Gaussian fitting method.

As a result of these trials certain characteristics of location algorithms have been investigated. Further work is required to add into the simulation the effect of noise and more extreme conditions such as very large target images. It is recognised that the simulation trial used here is not a complete indicator of what happens in real imaging situations, however, it does help to isolate the characteristics of each method of target location and allow rules to be constructed for each.

6. CONCLUSIONS

In this paper a comparison has been made between several different methods of target location. In addition a description of the problems of target recognition and threshold selection have also been covered. Various factors which affect the accuracy of target location have been analysed in a simulation experiment revealing many facets of the problems encountered by any algorithm used in practise. It is certain that further work is required to fully model all of the effects found in industrial measurement tasks but the approach appears to offer an excellent means of testing algorithms and, where necessary, tuning them for a particular application.

7. ACKNOWLEDGEMENTS

The research conducted by Dr. M.R. Shortis is supported by an Australian Research Council Grant entitled "Vision Metrology System Development". Dr. T.A. Clarke is currently supported by the Engineering and Physical Sciences Research Council of the UK. The contribution of Mr. X. Wang is acknowledged in the development of the scan-line filling algorithm.

8. REFERENCES

Beyer, H. A., 1993. Determination of radiometric and geometric characteristics of frame grabbers. Proceedings, SPIE Vol. 2067 Videometrics II, Boston, Massachusetts, September 1993, pp 93-103.

Brown, D. C., 1984. A large format, microprocessor controlled film camera optimized for industrial photogrammetry. Presented paper, Commission V, 15th Congress of the International Society for Photogrammetry, Rio de Janeiro, Brazil, 29 pages.

Chen, J. and Clarke, T. A., 1992. The automatic location and identification of targets in digital photogrammetric engineering measurement. Int. Archives of Photogrammetry and Remote Sensing, XVII(5), pp. 686-693.

Clarke, T.A. Cooper, M.A.R. & Fryer, J.F., 1993. An estimator for the random error in subpixel target location and its use in the bundle adjustment. Optical 3-D measurement techniques II. Pub. Wichmann, Karlsruhe, pp. 161-168.

Clarke, T.A., 1994. The effect of frame grabber adjustment on subpixel target location. To be published.

Gruen, A. and Baltsavias, E., 1988. Geometrically constrained multiphoto matching. Photogrammetric Engineering and Remote Sensing, 54(5) : 633-641.

Haggren, H. and Haajanen, L., 1990. Target search using template images. Proceedings, SPIE Vol. 1395 Close Range Photogrammetry Meets Machine Vision, Zurich, Switzerland, pp 572-578.

Lemmens, M., 1988. A survey on stereo matching techniques. Proceedings, Commission V, 16th Congress of the International Society for Photogrammetry and Remote Sensing, Kyoto, Japan, pp 11-23.

Shortis, M. R., Burner, A. W., Snow, W. L., and Goad, W. K., 1991. Calibration tests of industrial and scientific CCD cameras. Invited Paper (Paper 6, Volume 1), First Australian Photogrammetric Conference, Sydney, Australia, 12 pages.

Shortis, M. R., Snow, W. L., Childers, B. A. and Goad, W. K., 1993. The influence of storage media on the accuracy and repeatability of photogrammetric measurements using CCD cameras. Proceedings, SPIE Vol. 2067 Videometrics II, Boston, Massachusetts, September 1993, pp 80-92.

Shortis, M. R., Thompson, M. W. and Hall, C. J., 1994. An application of digital close range photogrammetry to the monitoring of large scale engineering structures. International Archives of Photogrammetry and Remote Sensing, Vol. 30 (5), pp 365-373.

Snow, W. L., Childers, B. A. and Shortis, M. R., 1993. The calibration of video cameras for quantitative measurements. Presented paper, 39th International Instrumentation Symposium, Albuquerque, New Mexico, May 1993, 28 pages.

Trinder, J. C., 1989. Precision of digital target location. Photogrammetric Engineering and Remote Sensing, 55 (6) : 883-886.

van den Heuvel, F. A., Kroon, R. J. G. A. and Le Poole, R. S., 1992. Digital close-range photogrammetry using artificial targets. Proceedings, Commission V, 17th Congress of the International Society for Photogrammetry and Remote Sensing, Washington D.C., U.S.A., pp 222-229.

West, G.A.W. and Clarke, T.A., 1990. A survey and examination of subpixel measurement techniques. Proceedings, SPIE Vol. 1395 Close Range Photogrammetry Meets Machine Vision, Zurich, Switzerland, pp 456-463.

Wong, K. W. and Wei-Hsin, H., 1986. Close range mapping with a solid state camera. Photogrammetric Engineering and Remote Sensing, 52 (1) : 67-74.

Wong, K. W., Lew, M. and Wiley, A. G., 1988. 3D metric vision for engineering construction. Proceedings, Commission 5, 16th Congress of the International Society for Photogrammetry and Remote Sensing, Kyoto, Japan, pp 647-656.

Zhou, G., 1986. Accurate determination of ellipse centres in digital imagery. Proceedings Volume 4, ACSM-ASPRS Annual Convention, Washington D.C., U.S.A., pp 256-264.

Zhou, H. B., 1990. Object points detection in a photogrammetric test field. Proceedings, SPIE Vol. 1395 Close Range Photogrammetry Meets Machine Vision, Zurich, Switzerland, pp 1227-1233.

PAPER REFERENCE

Shortis, M.R. Clarke, T.A., Short, T. 1994., A comparison of some techniques for the subpixel location of discrete target images, Videometrics III. SPIE Vol. 2350. Boston. pp. 239-250.