NPL

National Physical Laboratory

NPL REPORT MS 16

Simulation of complex, modern, computer-based measuring systems for evaluation of measurement uncertainty: case study in digital imaging and uncertainty analysis

C E Matthews, T J Esward

MARCH 2013

National Measurement System

NPL Report MS 16

Simulation of complex, modern, computer-based measuring systems for evaluation of measurement uncertainty: case study in digital imaging and uncertainty analysis

C E Matthews, T J Esward
Mathematics and Modelling

February 2013

ABSTRACT

We report the outcomes of a study of methods of achieving sub-pixel resolution in photogrammetry images to support the work of the National Physical Laboratory's Mass and Dimensional Group in large-volume metrology. We studied methods for quantifying the uncertainty associated with locating the centre of optical target reflections in images obtained using digital photography. We investigated four methods of identifying target locations in 8-bit greyscale images using both simulated and real data and describe the key results of the study and their implications for use in real-time and near real-time imaging systems. We also place the work in the wider context of the need to develop methods of associating uncertainty statements with quantitative results derived from imaging systems that are used in metrology applications.

NPL Report MS 16

ISSN 1754–2960

National Physical Laboratory,
Hampton Road, Teddington, Middlesex, United Kingdom TW11 0LW

We gratefully acknowledge the financial support of the UK Department for Business, Innovation and Skills (National Measurement Office)

Approved on behalf of the Managing Director, NPL by Mark Gee,
Knowledge Leader for the Materials Division

Contents

1 Introduction

The work reported here forms part of the project *Simulation of complex, modern, computer based measuring systems for evaluation of measurement uncertainty* that is funded by the National Measurement Office of the Department of Business, Innovation and Skills as part of its Materials and Modelling programme. This project adopts a model-based approach to evaluating uncertainties in complex, modern, computer-based dynamic measuring systems in line with the *Guide to the expression of uncertainty in measurement* [1] to ensure that all contributions to uncertainty that arise from signal acquisition, conditioning and processing are identified and quantified in a rigorous and consistent manner. The project aims to provide tools that ensure that best practice can be adopted by metrologists in an easily implementable manner, and that allow metrologists to apply their measurement results with the confidence that the uncertainties arising from their chosen measurement techniques have been accurately quantified.

We demonstrate how our approach can be applied to a quantitative image analysis problem from the discipline of photogrammetry: sub-pixel resolution to support the work of the National Physical Laboratory's Mass and Dimensional Group in large volume metrology. We have studied methods for quantifying the uncertainty associated with locating the centre of optical target reflections in images obtained using digital photography. The research was undertaken as a feasibility study to identify how we might extend the methods we have applied to dynamic one-dimensional signals (time series) to two and three-dimensional signals. This report therefore has the format of a short case study report.

2 Uncertainty analysis in digital imaging applications

From our experience in previous imaging analysis work at NPL, we have found that it is difficult to develop a high-level generalized approach to uncertainty analysis that can be applied across all imaging systems and applications. The interpretation of digital images presents a number of challenges. Typically the most significant properties of data from scientific imaging systems are the number of modalities available for analysis (these may include visible and non-visible light, ultrasound etc.), spatial and radiometric resolution, dynamic range, and signal-to-noise ratios. Each of these properties contributes sources of uncertainty when quantitative conclusions from images are required. This is true both for low-level quantification, e.g., image segmentation, and high-level decision-making, e.g., medical diagnosis. Uncertainty analysis and the association of probability statements with quantitative conclusions drawn from digital images require new methods, guidance and software tools.

Images can be regarded as possessing three types of uncertainty characteristics: static uniform (an image taken in an instance where uncertainty is uniformly spread across the image), static non-uniform (an image taken in an instance where uncertainty is non-uniformly spread, e.g. perspective scenes, images with camera aberration distortions), and dynamic non-uniform (time dependent images, e.g. video). We believe that a systematic approach that can be applied to many digital imaging systems may be developed if we concentrate on a range of lower-level features that are common to all digital imaging systems, rather than trying to assess uncertainties associated with higher-level features extraction and decision-making. Key lower-level features include the choice of imaging modalities, and signal characteristics (spatial, colour, temporal sampling and analogue to digital conversion, and signal-to-noise ratio).

The short case study reported here attempts to show how the location of the centre of target images and the uncertainty associated with the determination of that location are dependent on quantisation and image resolution effects and the choice of analysis algorithm.

3 The CATMESS system

3.1 Image analysis requirements

Metrology Space (CATMESS) is a photogrammetry-based measuring system [3]. A diverging laser beam, of wavelength $633\,\mathrm{nm}$, passes through a polarized beam splitter and is retroreflected from $10\,\mathrm{mm}$ diameter spheres, with high refractive index. It has been shown that glass spheres of refractive index very close to $n = 2$ can be used as omnidirectional retroflective targets [4]. The reflected beam is focused on to a CCD (charge-coupled device) image plane. The task in the analysis described here is to determine the coordinates (in pixel space) on the CCD that correspond to the centre of the target sphere. The recorded images vary greatly in quality and are affected by pixelation, camera distortion, and noise. Sharp focus settings on the camera produce clear target reflections with high contrast with respect to the background,but these target images are smaller and often saturated, making it difficult to identify targets far from the camera. The relative performance of location methods on targets of different size and quality is therefore of interest because the uncertainty value of the centroid algorithm in image space directly impacts a photogrammetry system's uncertainty in locating the spherical targets in real object space [6].

The location method must work in near real-time, so that a fully automated process is required, with no scope for user input. If an image is very poor quality, it is preferable for the location method to fail in real-time. For this analysis, we are concerned only with locating the centre of a given target. It can be assumed that some previous process has already identified the approximate target location. There will be background noise present in the pixels surrounding the sphere, but there will be one and only one target sphere in each image analysed.

Throughout this section of the report, unless stated otherwise, figure axes show pixel number.

3.2 Methods for centre location

The files analysed are 8-bit greyscale images (note that for clarity of presentation, images are shown in full colour scale in this report). Four different methods are investigated for locating the target centre.

- Threshold mean. A threshold is applied to all pixels, producing a binary representation of the image. The means of the x-coordinates and the y-coordinates of the pixels with intensity exceeding the threshold gives, respectively, the x- and y-coordinates of the estimated centre location.
- Weighted mean. An edge detection algorithm is applied to the image. A weighted mean of the x- and y-coordinates of the pixels inside the defined edge gives, respectively, the x- and y-coordinates of the estimated centre location. The weighting of each pixel is defined by the pixel intensity.
- Circle-fitting. An edge detection algorithm is applied to the image. The centre of the least-squares best-fit circle to the defined edge gives the estimated centre location.
- Reflection. Vertical and horizontal axes of reflection are tested by calculating the sum of squares of differences between the pixel intensities in the actual image and the image reflected along the specified axes. The intersection of the horizontal and vertical axes providing the minimum sum of squares difference gives the estimated central location. The axes tested run along the edge or through the centre of pixels. This method therefore produces a discrete estimate, in the sense that the resulting centre point will always lie at the centre, corner or mid-edge of a pixel.

3.3 Simulated images

Simulation allows testing of the four location methods on images of varying degrees of quality. The effects of both noise and pixelation are investigated through simulated images. Pixelation is simulated by using two grids of different resolution. A *continuous* image is generated on a high-resolution grid (e.g. 800×800 pixels), as a circle of given radius, with intensity decreasing linearly with squared distance from the centre. Gaussian noise is added to the image at the high-resolution stage. Pixelation is then simulated by reducing the grid down to a lower resolution (e.g. 50×50 pixels), with each new pixel covering a uniform area of high-resolution sub-pixels (figure 1). The subdivision of the high-resolution grid is applied as a systematic tiling, beginning at a randomly selected pixel within the top left 8×8 pixels. Therefore the known, true centre of the circle, when translated on to the low-resolution grid, is not necessarily located at the centre of a pixel. The intensity of each low-resolution pixel is proportional to the sum of intensities of its sub-pixels.

As the true centre of the circle in an image is known, the four location methods can be tested for accuracy for different levels of noise and pixelation. Figure 2 shows the estimate from

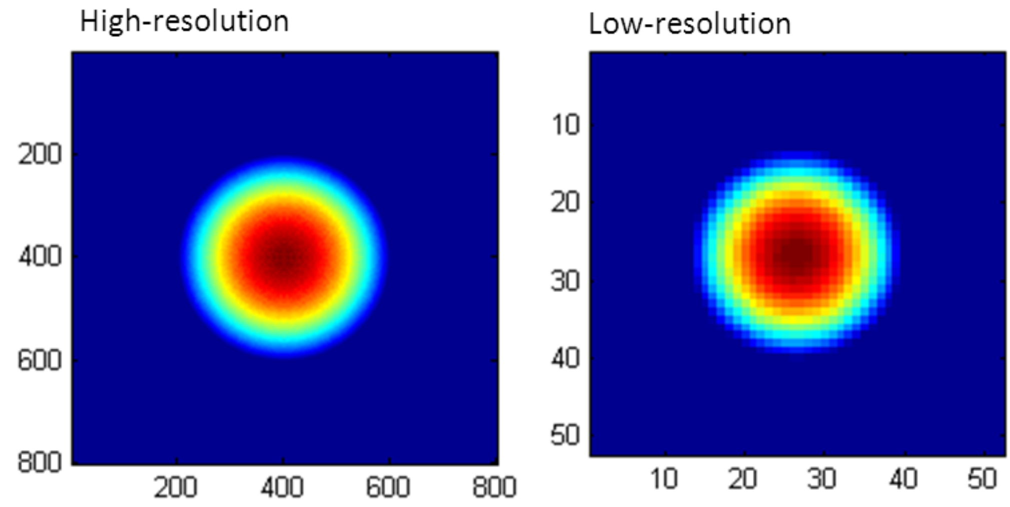

Figure 1: Simulated high-resolution and low-resolution images.

each method for a single simulation. The reflection method performs appreciably worse than the other three methods and its restriction to discrete pixel edges or centres can be seen in the right-hand figure.

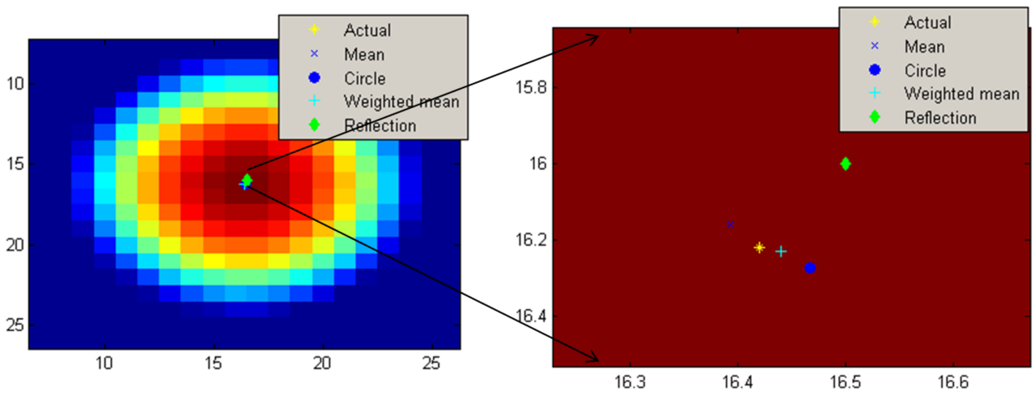

Figure 2: Centre location estimates from four different methods.

A Monte Carlo method approach is used to estimate the uncertainties associated with pixelation and image noise [2]. Figure 3 shows the mean and standard deviation of the errors for each method over 50 simulations. The amount of pixelation varies for each image, with the low-resolution scaling ranging from 400×400 pixels to 25×25 pixels. Random noise is also added to each image. For each simulation, the threshold mean method was applied both with a relatively high and relatively low threshold value. The error is calculated as the Euclidean distance between the estimated and true centres. The weighted mean method produces the smallest mean and standard deviation of errors. As seen in the single simulation, the discrete nature of the reflection method leads to much larger errors than for the other methods. The two different threshold values used in the simple mean method do not appear to have an effect on the performance of the method.

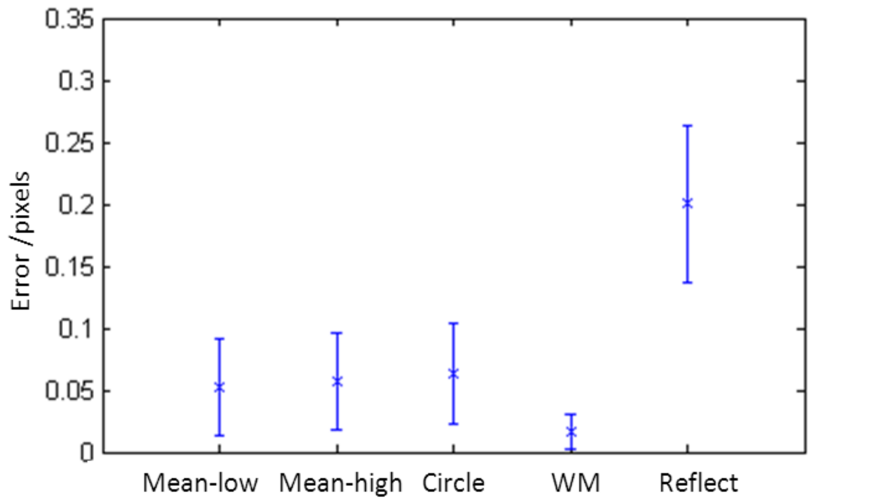

Figure 3: Mean error for different location methods over 50 simulations with varying degrees of noise and pixelation. Location methods are Threshold Mean, with two different threshold values (Mean-low, Mean-high), Circle-fitting (Circle), Weighted Mean (WM) and Reflection (Reflect). Error bars show the standard deviation for each method.

In practice, as seen in section 3.4 of this report, often the intensity of pixels making up the target image becomes saturated and for these cases a simulation comprising a more uniform intensity, with variation only at the sphere borders may be more appropriate than the graded images in figure 1.

3.4 Real images

The quality of an image and the resulting performance of the centre location methods will be affected by camera settings, such as bit-depth and exposure time, as well as the distance of the target from the camera. It is useful to assess the relative performance of the methods for the different types of images likely to be produced. Figure 4 shows the images of spheres located at different distances from the camera. Identifying the centre of the furthest sphere would be very difficult without using a longer exposure time. However, a longer exposure time may increase noise in the closer images from non-target objects and lead to saturation of more pixels, therefore losing detail. The images also appear to be affected by the relative orientation of the camera and a sphere, suggesting a lack of uniformity of the spheres themselves. This effect will be seen in more detail in section 3.5 below.

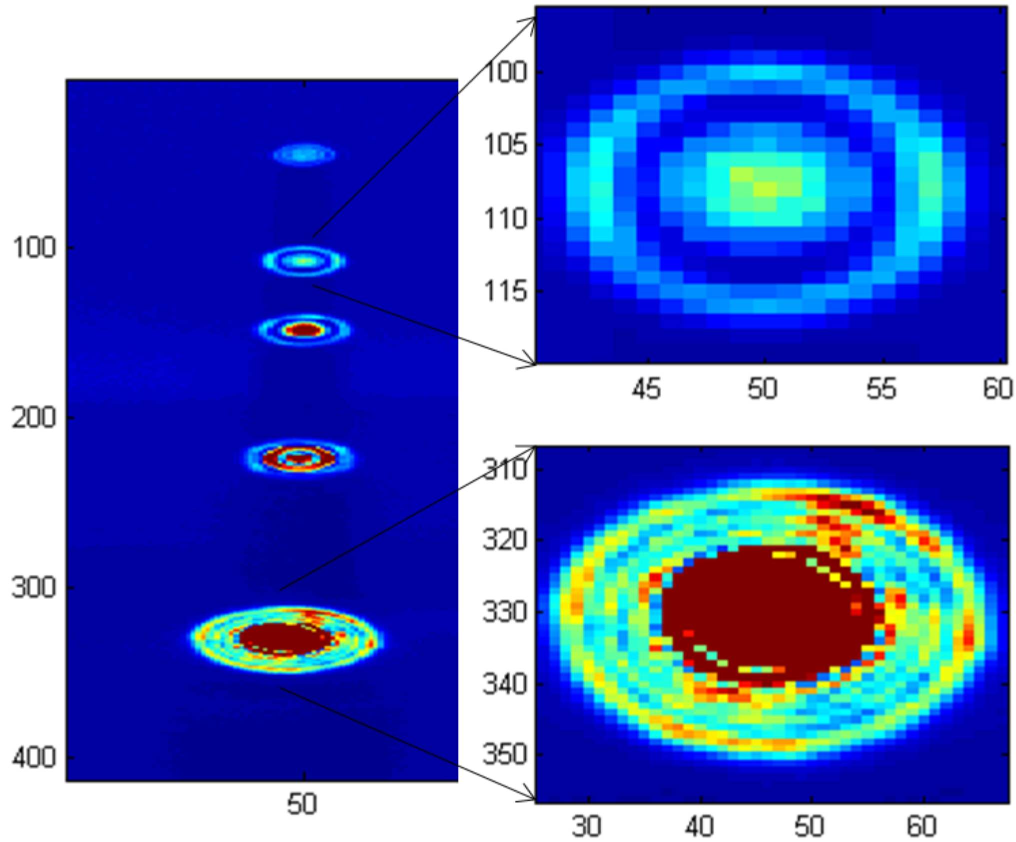

Figure 4: Image containing five spheres at different distances from the camera. The camera settings were 8-bit mode and 5 ms exposure time.

The sets of five sphere images, as seen in figure 4, are useful for testing the location methods on images of different quality and with different camera settings. However, with no knowledge of the true centre for each sphere, and no repeat images for each experimental set-up, it is only possible to obtain a qualitative assessment of the performance of the location methods. A new set of experiments was performed, described below, to enable a better assessment of the uncertainty associated with the estimates from each method.

3.5 Talyrond images

A series of images were obtained while rotating the sphere on a Talyrond roundness machine. The set-up of the roundness machine comprises a turntable, with the sphere mounted at the centre. The Talyrond would typically be used to obtain a roundness profile of the sphere. However, it can be used in this instance to rotate the sphere incrementally while centering the sphere to within $1.5\,\mu m$ horizontally. The distance between the camera and sphere was approximately $2.5\,m$, and the camera focal length is $23\,mm$. The $1.5\,\mu m$ centering in object space therefore corresponds to $1.5 \times (23/2500) = 0.014\,\mu m$ in image space, or 1/400th of a pixel. It is assumed that the sphere has no movement vertically. The set of 2D sphere images will then nominally all have the same centre location (1/30th of a pixel is regarded as excellent centering in photogrammetry terms [5]), and the reproducibility of the location methods may be tested.

In total, nine sets of test images were produced with the Talyrond machine. Images were taken at three different focuses:

- focused - small, sharp spot with little gradient at boundary;
- one turn out of focus (focal length increased by $0.5\,mm$) - larger, less sharp spot;
- two turns out of focus (focal length increased by $1\,mm$) - larger and less sharp again.

Each different focus setting was taken at three different exposures: $1\,ms$, $10\,ms$ and $30\,ms$. There are 18 images in each of the nine test sets, with the sphere rotated by $20°$ between images.

3.5.1 Reference set

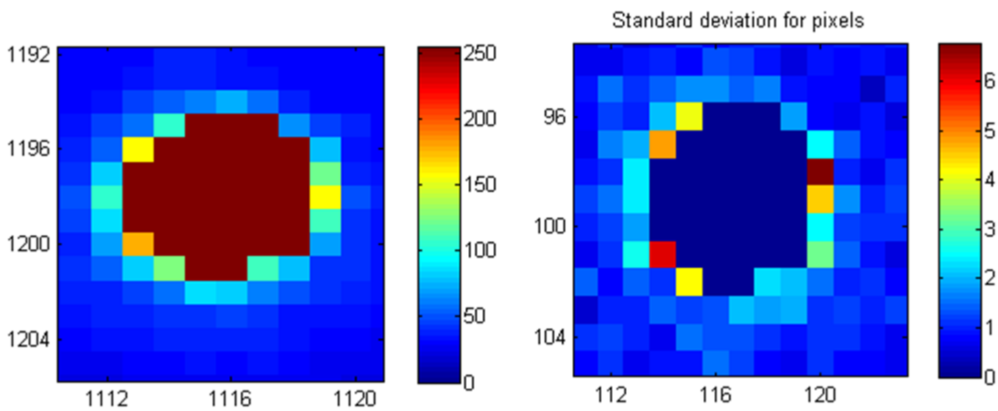

Figure 5: Left: Typical image from the reference set of spheres rotated through $360°$. Right: Standard deviation of pixel values over all ten images in the set.

In addition to the incrementally rotated sets (rotated $20°$ between each image), a reference set of ten images was taken in which the sphere was rotated a full $360°$ between each image. Figure 5 shows a typical image of the sphere from the reference set, and the standard deviation of pixel values across all images in the set. The centre of the sphere is identical, but unknown, in each image. When the four location methods are applied to each image, the circle-fitting and reflection methods return identical estimates in each case. The slight variation in the border pixels leads to a slight variation in the centre estimates for the weighted mean location method. A large variation is seen in the estimates from the mean pixel method if a low threshold value is used (figure 6). This variation is significantly reduced if the threshold is increased.

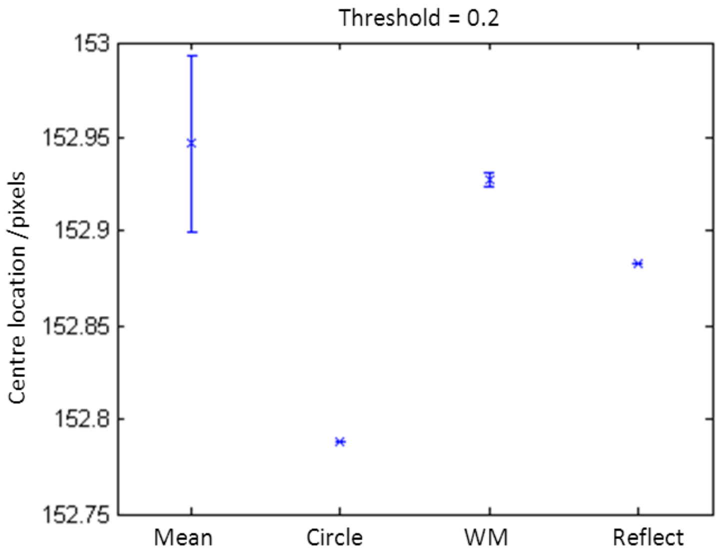

Figure 6: Mean estimated centre location from the four location methods over the set of reference images. Error bars show the standard deviation of estimates. A threshold value of 0.2 is used in the mean pixel method.

3.5.2 Test sets

All four location methods were applied to the 18 images in each of the nine test sets. The mean pixel method was applied with a low (0.4) and high (0.8) threshold value.

For each of the nine sets, there is a significant change in the 2D image of the sphere as it is rotated (e.g. figure 7). On analysis of all test sets, it was found that the irregular images (such as those in the centre and right of the top row in figure 7) were found at the same rotation values in all sets. This suggests a systematic defect in the sphere itself. Figure 8 shows mean results of the location methods for a single test set: one turn out of focus, 1 ms exposure. Note that in the figure, the pixel intensities as well as the centre locations are the means over the 18 rotated images in the set. In addition to the mean estimate for a given set of images, given the change observed for a single sphere image during rotation (figure 7), it is also of interest to consider the variation of location estimates provided by each method, over a given set.

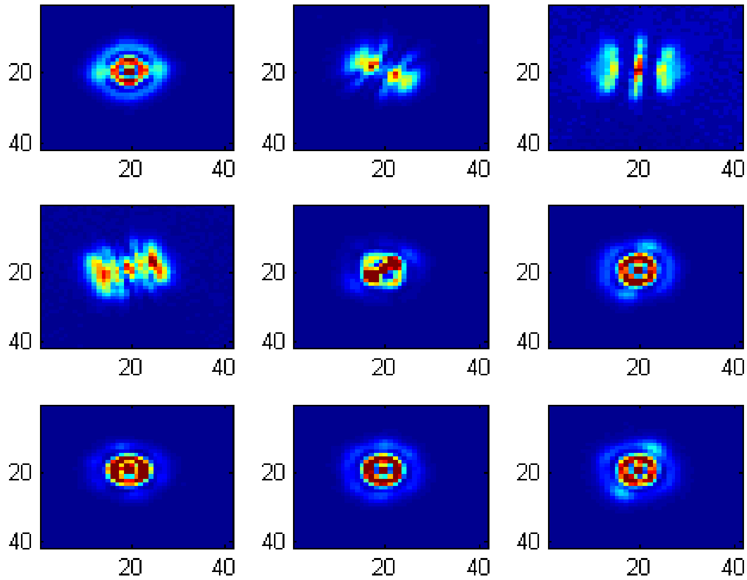

Figure 7: Images of the same sphere, under the same focus and exposure conditions, at different rotations.

The standard deviation of x- and y-coordinates of the centre location estimations from each method are shown in figure 9, for all nine test sets. The sets are ordered horizontally with focus decreasing from left to right, and three increasing exposure times within each focus group. Where there is no data point, the method failed to find a centre for at least one of the 18 rotated images. The weighted mean and circle-fitting methods both failed for one set. These two methods both rely on identifying a boundary in the image. The algorithm performing this edge detection requires a parameter to control the minimum allowable size of items identified in the image. When multiple, isolated shapes are detected within an image, such as seen in the top, right cell in figure 7, we want the location methods to recognise this as representing one sphere only. However, particularly for sharp, long exposure images, surrounding noise in the image may also be identified as being part of the sphere, with the result that the algorithm will attach these isolated patches of noise to the pixels covering the sphere. The parameter for the detection algorithm therefore needs to be set such that the algorithm is sensitive enough to recognise different parts of the sphere in poorer quality images, but not so sensitive that it begins to include noise in the image. For a given image, with some pre-analysis, it is straightforward to find a value of this parameter that will allow the method to run successfully. However, as the quality and exposure time of the test images varies overs the different sets, it is not possible to find a single value that is appropriate for every image. Given the requirement for near real-time analysis it is not

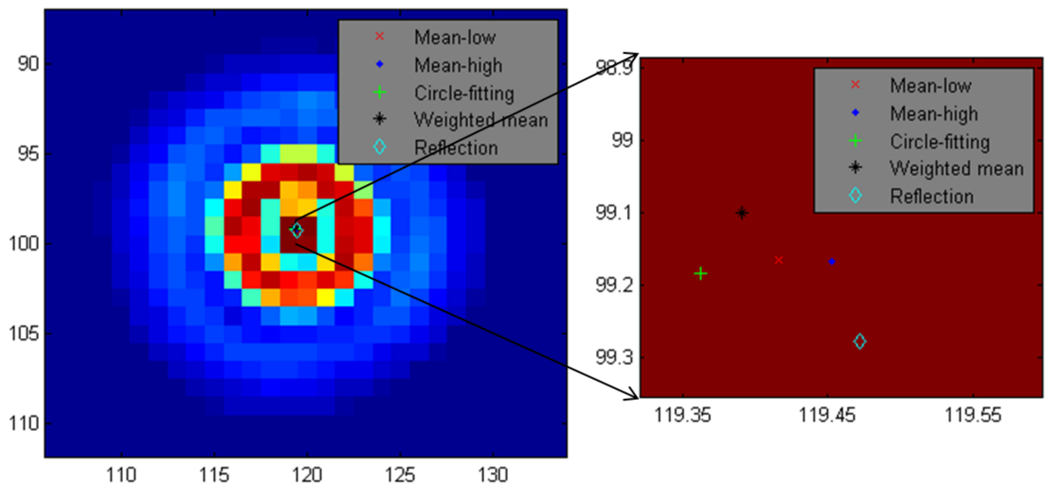

Figure 8: Mean estimated centre location from the four location methods over the set of test images taken with one turn out of focus and 1 ms exposure. Pixel intensity is also shown as the mean over the set of images.

practical to assign the parameter by user input and therefore a single value was chosen prior to running the analysis and this same value was used on every image. The result of this approach is that for the sharp images with the longest exposure time, background noise (in fact, the stand that the sphere was mounted on) was identified as being part of the required image. To attempt to classify this as a failed location attempt, rather than still returning a value, a roundness test is applied to the selected pixels.

For all four location methods, there is a clear general trend of increased standard deviation as the focus becomes less sharp and as the exposure time increases. Both of these camera settings are likely to increase the noise in the images. Despite the weighted mean method failing to return an estimate for every image in the sharp focus, 30 ms test set, across the remaining test sets this method appears to perform best, on average, in terms of minimum variation of centre locations.

4 Conclusions and next steps

4.1 Conclusions from the CATMESS work

The methods of simply calculating a pixel mean or pixel intensity weighted mean appear to perform better, on average, than the reflection or circle-fitting methods both in terms of accuracy (in the simulated case) and variation of results. Further simulations, with a greater

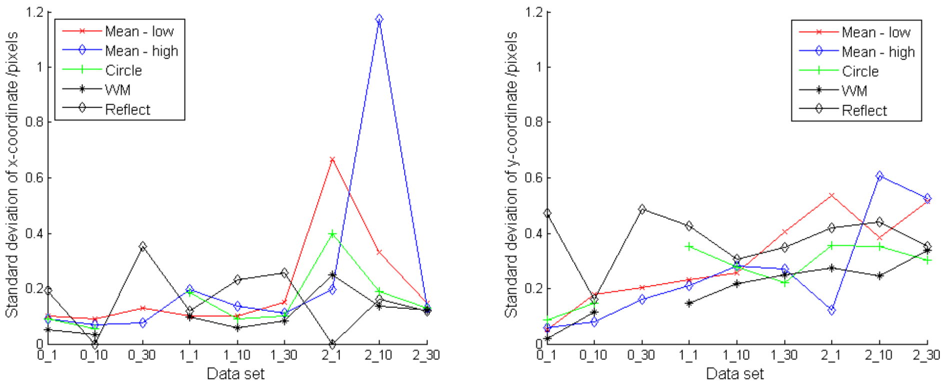

Figure 9: Standard deviation of pixel x- and y-coordinates for centre estimates over 18 rotation images in nine test sets. Horizontal axes denote the test conditions: (turns out of focus)_(exposure time /ms). For each image the centre location estimate is found using the four different methods, and using two different threshold values for the mean pixel method. Where no data point is shown, the method failed to find a centre point for at least one of the 18 images.

variety in the quality of images would help to further understand the relative accuracy of the methods. A better understanding of the features of the sphere that are causing irregularities would help to generate more realistic simulations. The reflection method, while the most robust in terms of dealing with any image quality or shape, is not suitable without further development, due to the discrete nature of its estimates.

The possibility of developing and using multiple methods may also be considered. In this case a merit function could be used to quantify the performance of each method and the results used to select the best estimate. Taking into account the uncertainties associated with each of the estimates, a weighted mean or largest consistent subset approach could be taken. Alternatively an accuracy/robustness trade-off could be implemented by using the weighted mean location method where possible, but applying the more robust reflection method if the weighted mean failed.

The need for a near-real time algorithm makes it difficult to apply any method requiring a user-defined parameter. In the cases where some form of decision must be made (e.g. threshold value for greyscale images, minimum allowable size for identification of target etc.), further work would be required to determine the optimal values for parameters, taking in to account relative performance of the method with different parameters on different quality images. This analysis should also consider that it is preferable for a method to fail completely than to return a large error estimate.

In the analysis described here we have looked only at effects of sphere orientation on the estimates from each algorithm. As a next step we may consider the effects of image size and digitisation. For example, how do the results differ if the estimate is based on a circle of 3-pixel diameter, or 50-pixel diameter?

4.2 Beyond CATMESS

We have shown that the approach to identifying and quantifying uncertainty sources in images that we set out in the introduction to this report can be applied successfully to a photogrammetry application where the task is to locate the centre of a target image of an apparently simple object such as a sphere. We intend to seek opportunities to apply our methods in other metrology applications areas where quantitative conclusions drawn from images are of importance. Examples of such application areas include nanostructure and microstructure analysis of materials, biological cell analysis including single molecule detection, earth observation, video surveillance and information fusion for biometric identification and tracking systems.

References

[1] JCGM 100:2008 *Evaluation of measurement data - Guide to the expression of uncertainty in measurement* Joint Committee for Guides in Metrology

[2] JCGM 101:2008 *Evaluation of measurement data - Supplement 1 to the "Guide to the expression of uncertsinty in measurement" - Propagation of distributions using a Monte Carlo method* Joint Committee for Guides in Metrology

[3] Hughes E., Warden M.S., Veal D. 2012 *Coordinate System and Method* UK patent application GB1205563.8

[4] Yang B, Friedsam H. 1999 *Ray-tracing studies for a whole-viewing-angle retroreflector* Argonne National Laboratory, Argonne, IL USA

[5] Robson S., Shortis M.R. *Photogrammetry Workshop* National Physical Laboratory, Teddington UK, 22 June 2011

[6] Luhmann T., Robson S., Kyle S., Harley I. 2006 *Close Range Photogrammetry: Principles, Techniques and Applications* Whittles Publishing, Scotland