

A basis for the selection of appropriate weights for image observations in real-time 3-D measurement

M. R. Shortis¹ and T. A. Clarke²

1. Department of Geomatics, University of Melbourne, Parkville 3052, Australia.
m.shortis@unimelb.edu.au

2. Optical Metrology Centre, City University, Northampton Square, London, EC1V 0HB, UK
t.a.clarke@city.ac.uk

ABSTRACT

In the context of the quality control of manufacturing tasks it is important that errors in the measurement process can be estimated reliably. For instance to achieve assembly of aerospace structures with minimum tooling and without fixed jigs the measurement system must provide an assessment of the accuracy achieved. A typical task might be the estimation of the relative orientation of two components or the location of a machine tool with respect to a component. In either case the measurement system must be able to ensure that the manufacturing process has been performed within tolerance. Photogrammetric techniques are now beginning to be applied to such situations.

This paper will consider the use of weighted observations in photogrammetric systems with the objective of analysing how effectively errors in the observations can be detected in the object space measurements. For instance, the paper will consider:

- the difference between an equal weighting scheme compared to a scheme based on target image location precision measures. An equal weighting scheme results in a homogeneous precision throughout the 3-D co-ordinates even though the images of some targets may be very dim and poorly located while others may be bright and be located accurately.
- the rationale for a radial weighting scheme which takes into account the known increase in error sources and the typical paucity of observations at the edges of the sensor format.

The objective of this paper is to discuss and investigate at least schemes for the production of correct observation weighting which will lead to appropriate precision estimates of 3-D co-ordinates even at the cost of poorer overall results.

Keywords: Least squares, weighting, real-time, precision, quality control.

1. INTRODUCTION

This subject of weighting of photogrammetric observations has not been a prominent topic over the past few years. Users of professional bundle adjustment software such as CAP, VSTARS or BINGO would not use nor recommend weighting of image residuals for close range photogrammetric purposes. The software usually allows such weighting but it appears that the effects of weighting schemes are not well enough understood to be reliably used in practice^{1,2,3}. However, it is clear that the accuracy of observations can vary considerably within an image and sometimes from camera to camera. If all observations are equally weighted then clearly this situation is absurd as this will result in an approximately homogeneous precision in the estimates that will not reflect either the better precision of the best observations or the poor precision of the worst observations. If the observations are weighted then it might be expected that the results will reflect the actual precisions, and potentially also the accuracies, of the observations.

The question that needs to be addressed is how to compute appropriate weighting values. This paper reports on investigations of the “weighting problem” with an objective of producing correct weights for photogrammetric systems.

2. SIMPLE EXAMPLES OF LEAST SQUARES PARAMETER ESTIMATION

The analysis of the properties of errors can be conducted by the use of a simple simulation of the task of estimating the two parameters of a straight line. While the properties of the errors have been analysed in many text books on the subject^{4,5} and are well known, an experimental investigation provides a clear insight into the issues and provides a sound basis for further analysis later in this paper.

2.1 Percentage error in parameter estimation with the number of samples

A straight line was produced and noise was added to the perfect y axis values, a least squares estimation (LSE) solution for the parameters of the line was computed. The noise added had a normal distribution and a standard deviation of 0.01. The slope value was 0.1 and the intercept was 1.0. The number of observations was then chosen to vary from 2 to 2000 and the percentage error in estimate of the parameters is shown in figure 2.1 as a Log versus Log graph. The graph illustrates that the error in parameter estimation is directly related to the number of observations. If only a few observations are used and the error is a significant proportion of the signal the estimated parameters can be in considerable error.

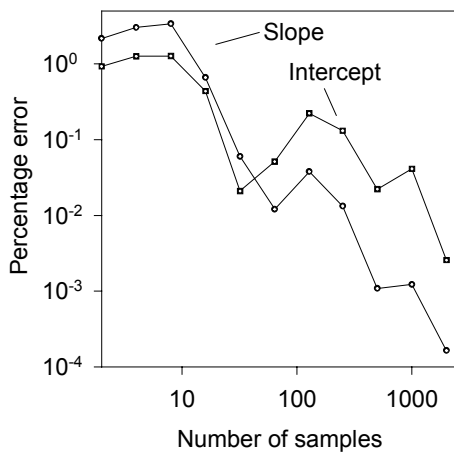


Figure 2.1 Log percentage error versus log number of samples




Figure 2.2 Percentage error plotted against standard deviation of the noise

2.2 Percentage error in parameter estimation with the level of noise

As the level of noise is changed for a fixed number of observations, the accuracy of parameter estimation should change. To test this a further experiment was conducted where the noise level was changed for 2000 samples. Figure 2.2 illustrates the fact that the parameters can be extracted from the signal but the error in extracting the parameters is directly related to the noise level.

2.3 The effect of mixing measurements with differing error levels

Having established the dependence of parameter estimation on the number of points and the level of error, the effect of a mixed error signal is considered. To achieve this the same experiments are performed again, this time with the errors mixed together. Table 2.1 illustrates the results.

Parameter	1000 x 0.001 error	1000 x 0.1 error	2000 x 0.1 and 0.001 error	2000 x 0.1 error
Slope error	0.000000049	0.0000077	0.0000036	0.0000067
Intercept error	0.0000272	0.00837	0.0042	0.006751

Table 2.1 Errors in parameter estimation for differing observation errors

When the two sets of error are combined the parameters are only marginally better than for the case of the large errors alone. This result is as expected from the theory of error propagation and is illustrated very clearly in this simulation experiment. The dominance of the size of the errors on the results is a potential problem for mixed accuracy target location situations. Attempting to rectify the situation with blunder detection is a difficult task.

2.4 The effect of weighting

To test the effect of weighting in the LSE of a straight line, the error in slope and intercept were computed for two sets of 2000 data points with normal distribution errors added first of 0.1 and then of 0.001. These provided the basic performance of the estimation under both conditions (Table 2.2).

Error	2000x $\sigma = 0.1$	2000x $\sigma = 0.001$	1000x $\sigma = 0.001$	1000x $\sigma = 0.001$
	No weights	No weights	1000x $\sigma = 0.1$ No weights	1000x $\sigma = 0.1$ Weighted
Slope error (%)	0.00671	-0.000056	0.003128	-0.000057
Intercept error (%)	-0.6751	0.00498	-0.2527	0.00651

Table 2.2 Results of the experiments with differing weights

When half of the data samples were given the large errors and half the small errors, the results were slightly improved with respect to the large errors only and very poor with respect to the small errors. In other words the results are swamped by the larger errors (as expected by the theory). When the LSE was performed with appropriate weights the results were only marginally worse than the best results. This result is of course expected. It would not be expected that the results would be better than the case of 2000 good measurements, as there are only 1000 good measurements. Figure 2.3 illustrates these results graphically. Although this situation is exaggerated for clarity it does represent a realistic situation that will be demonstrated later in this paper.




Figure 2.3 Data used in the experiment (200 samples out of the 2000 used). The approximately straight line at zero is the small error data.

3. PHOTOGRAMMETRIC SIMULATION

The simple example of the linear regression in the previous section can be extended to a close range photogrammetric self-calibration. A typical scenario for a pre-calibration of a standard CCD camera such as the Pulnix TM series is a three dimensional calibration field and multiple convergent images exposed from a range at which the calibration field almost fills the shorter dimension of the rectangular sensor. Typically, 100 or more retro-reflective targets would be employed on the calibration field, which would have an overall size suitable for the intended metrology application. In this instance an eleven by eleven array of target was used, with ten targets substantially out of the plane of the remaining 111. Several

exposure positions would be chosen, and in this case eight locations have been adopted, seven in a ring which give a 90 degree intersection at the centre of the calibration field, and one in a central position. The camera is given a 90 degree roll at each exposure station to minimise parameter correlations.

3.1 Uniform weighting of observations

It is well known that if uniform weighting of observables is used in least squares estimation, the internal (image space coordinates) precision and the external (object space coordinates and derived camera calibration parameters) precision are linearly related. This process can be demonstrated using a simulation in a similar fashion to the linear regression. The parameters and observables are generated from the ideal case and then normally distributed random deviates are added to the observables to simulate noise. The data is then processed through a conventional photogrammetric network solution using self-calibration and a least squares estimation solution.

The linearity of the relationship between the internal and external precisions is shown in table 3.1 using three cases of internal precision. The three cases are noise levels of 10 microns, 1 micron and 0.1 microns, corresponding to approximately 1, 0.1 and 0.01 pixels respectively. These values cover the extremes of the range of likely target image measurement precision, from visual estimation to binary centroiding to leading edge techniques such as weighted centroiding, least squares matching or edge matching techniques.

Number of exposures at each station	Network Redundancy	Object Space Precision (mm)			Precision of the Calibrated PD		
		Image Precision			Image Precision		
		10	1	0.1	10	1	0.1
1	1400	16.86	1.69	0.17	18.9	1.9	0.2
2	3000	12.02	1.20	0.12	13.1	1.3	0.1
4	7000	8.43	0.84	0.08	9.2	0.9	0.1
8	14000	5.95	0.59	0.06	6.4	0.6	0.1
16	30000	4.19	0.42	0.04	4.5	0.5	0.0

Table 3.1 Results of simulations of a self-calibrating photogrammetric network solution with uniform weighting of observables (all units are micrometres unless stated).

In table 3.1, the object space precision and the precision of the calibrated principal distance increase and decrease in proportion to the image space precision. The relationship is close to perfectly linear within the limits of rounding errors.

3.2 Effect of the number of observations

Table 3.1 also shows the relationship between the number of observations and the external precision measures. Simulations were also carried out for multiple exposures at each station, doubling the number of available camera images four times. Once more it is well known from statistical theory and error propagation that increasing the number of observables improves the external precision measures by a factor which is the square root of the increase in the number of observables.

In the table it is evident that a fourfold increase in the number of observables improves the external precision measures by the expected factor of two. In practice this relationship is not as well behaved as the linear relationship between internal and external precision variations. The complication is that additional exposures would also have 90 degree rolls applied, and would be inadvertently or deliberately moved slightly between exposures. These strategies once more have the effect of minimising correlations and also randomising the locations of target images to avoid systematic placement of targets on exposures, respectively. This practice was also simulated here and the slight variations from the square root relationship are evidence that there is an effect from these strategies.

An alternative mechanism for changes in the number of observables is variation in the number of targets, rather than the number of exposed images. Shown in table 3.2 and figure 3.1 are the results of varying the number of targets for the self-calibrating network with four exposures per station and a one micrometre noise level. Target numbers were reduced by

eliminating rows and columns of targets within the calibration field, whilst always retaining corner targets (to maintain the relative size of the field) and at least one out of plane target (to ensure that the field retained a three dimensional nature throughout).

Number of targets	Network Redundancy	Image RMS error	Object Space Precision (mm)	Precision of the Calibrated PD
6	80	0.51	1.54	6.9
9	230	0.68	1.16	3.3
18	750	0.85	1.00	2.2
36	1800	0.94	0.90	1.6
66	3500	0.98	0.87	1.2
121	7000	0.98	0.84	0.9

Table 3.2 Results of simulations of a self-calibrating photogrammetric network solution with variation in the number of targets (all units are micrometres unless stated).

The relationship between the number of targets and the external precision measures is clearly non-linear when a range of targets from six to more than one hundred is considered. However within this range, there is a transition at approximately twenty or fifty targets, depending on the intent of the network. If the intent is a measurement application and the target coordinate precisions are of prime importance, there is a small but significant degradation of precision for networks of less than approximately twenty targets. If the primary aim of the photogrammetric network is the calibration parameters then there is an accelerating degradation of precision for networks of less than approximately fifty targets.




Figure 3.1 Results of simulations of a self-calibrating photogrammetric network solution with variation in the number of targets (all units are micrometres unless stated).

A more unexpected result of reducing the number of targets is the apparent improvement in the internal precision, as indicated by the decreasing image RMS (root mean square) error. As this is simulated data with introduced normally distributed noise, the reduction must be spurious. Whilst there is an effect from the mathematical relationship between the estimate of unit weight and the RMS value, for networks with small numbers of targets there are also effects from parameter correlations. With so few targets, systematic effects such as lens distortion are no longer modelled effectively and the "fit" of each exposure to the target image data is falsely improved. The ineffective modelling of the systematic errors in the network solution will inevitably lead to the degradation of accuracy.

The transitions shown here are specific to the network used and will be influenced by target field layout, however the general conclusions that can be drawn are nevertheless valid. For large numbers of targets in a photogrammetric network the effect of eliminating individual targets is very small. On the other hand, when small numbers of targets are involved the

elimination of a single target may cause a significant degradation in the precision of the target coordinates or the calibration parameters or both. Further, for low numbers of targets, the accuracy of the results may also be called into question due to the weakening of the geometry of the network and the increased influence of correlations. More detail on these effects can be found in a number of publications on photogrammetric network design⁶.

In terms of the weighting of observations within the least squares estimation solution, the relevant conclusion is that, for relatively large numbers of targets, the external precision measures of coordinates and calibration parameters are virtually independent of the number of targets. This is an important consideration when groups of targets must be eliminated from or separately analysed within a single photogrammetric network. If there is a significant change in the external precisions then it is clear that there is another factor present, such as inappropriate weights for the observed image coordinates.

3.3 Non-uniform weighting of observations

The introduction of observables with differing error levels can be simulated for a photogrammetric network. As noted earlier, this non-uniformity could be caused by considering that some target images within a calibration field or object to be measured are of different size or brightness. To illustrate the effect in a straightforward analysis, the self-calibration network was re-processed with half the exposures with one micrometre simulated noise and half the exposures with ten micrometre simulated noise.

The networks with varying numbers of exposures were then processed in two modes. The first used non-uniform weighting in the least squares estimation solution, such that the weighting of the observables corresponded to the applied noise level. As can be seen in table 3.3, the results for the non-uniformly weighted solution, shown in the columns headed "non-uniform", exhibit the expected trends. Once more the square root relationship between numbers of observables and external precision is clearly evident. A comparison with table 3.1 also shows that the linear relationship between internal and external precisions is also in force. Like the linear regression (see table 2.2), the networks in table 3.3 which were processed with non-uniform weights behave virtually as if the observations with ten micrometre noise are not present, that is, as if the networks were composed of only observations with one micrometre noise. The results correspond with the data for observations with one micrometre noise in table 3.1 with half the number of exposures per station. This result is expected, as the data with one micrometre noise dominates the least squares solution and the data with ten micrometre noise has no significant impact because of the different weighting.

Number of exposures at each station	Image RMS Error						Object Space Precisions (mm)		Precision of the Calibrated PD	
	All exposures		1 um exposures		10 um exposures		Non-uniform	Uniform	Non-uniform	Uniform
	Non-uniform	Uniform	Non-uniform	Uniform	Non-uniform	Uniform				
2	7.32	6.80	0.91	2.57	10.14	9.24	1.68	8.63	1.9	9.5
4	7.24	6.94	0.96	1.93	10.04	9.62	1.20	6.10	1.3	6.5
8	7.21	7.03	0.98	1.58	10.01	9.80	0.84	4.30	0.9	4.3
16	7.16	7.03	0.99	1.35	9.94	9.83	0.59	3.01	0.6	2.9

Table 3.3 Results of simulations of a self-calibrating photogrammetric network solution with non-uniform and uniform weighting of observables with differing error levels (all units are micrometres unless stated).

In contrast, if the observables are uniformly weighted with an expected precision of one micrometre, the results of the network processing are quite different. This case of uniform weighting corresponds to the practical case in which poorly defined target images are included, but not appropriately weighted in the least squares estimation solution.

Results for this case are shown also in table 3.3, under the columns headed "uniform". Whilst the square root relationship between numbers of observables and external precision is still evident, the linear relationship between internal and external precisions is no longer applicable. Not only have the external precisions have been inflated substantially by the *a posteriori* estimate of unit weight, but the relative effect of the different error levels in the data changes with the number of exposures. As the number of exposures, and therefore the total number of observables, increases then the separation of the two

populations of observables is more distinct. Evidence for this is the change in the RMS image errors with the number of exposures, which approach the actual noise levels for larger numbers of observables.

These simulations demonstrate the important concept that appropriate weighting of observables is crucial for least square estimation solutions for any application, including photogrammetric networks. Although it may be feasible to separate out different groups of observables using a variance component analysis, this technique is not widely known and becomes less effective with low numbers of observables. A more likely outcome of inappropriate weighting is an excessive rejection of observables and an unwarranted inflation of external precision measures.

4. SYSTEMATIC ERROR SOURCES IN CLOSE RANGE PHOTOGRAMMETRY

4.1 Errors due to retro-reflective target angular characteristics

Retro-reflective material properties are created by a complex interaction with incident light due to diffraction, occlusion, and variations in transmittance and reflectance⁷. As a result materials have a range of characteristics depending on the design. The parameter of interest here is the fall off in returned light with angle and distance from a light source. The latter effect is primarily due to the inverse square law variation in light intensity with distance together with the ability of the lens to collect the returned light, which diminishes with distance from the target. The fall off with angle effect is due to physical occlusion of the incident light return path, the variation in area with angle, and the ability of the manufacturer to coat the rear of the spherical balls with reflective material. As a result of these physical constraints retro-reflective target users find that there is usually dramatic fall-off in intensity generated by each target beyond around 40-55 degrees. Methods exist to avoid this effect but usually at the expense of another desirable features such as the efficiency of retro-reflection. A typical photogrammetric network will place cameras in positions where the angle between the main object surface normal and the camera are between 35-50 degrees. The point at which the light return from the target has fallen by 90 percent from its maximum is illustrated in figure 4.1.

Type (3M)	3991 series	3970 series	5870 series	6710 series	2290 series
10% point (degrees)	40	50	55	55	40

Table 4.1 Angle at which the return light is 10 percent of the maximum for various materials

It is not often possible to constrain the targets to all be oriented in the same plane or at the same distance from the cameras. An inevitable result is that target images will be both bright and dim, and often the same target viewed from a different direction will have both bright and dim images. A direct consequence of this will be a reduction in the accuracy with which targets can be located as a result of quantization and image noise. The requirement for a weighting scheme dependent upon intensity is the logical conclusion of target intensity variations^{8,9}.

4.2 Errors due to the physical shape of the sensor

Sensors are typically rectangular with either a 4/3 aspect ratio or, more recently, a 16/9 aspect ratio. The fact that the sensor is this shape has implications for camera calibration in that the observations at the corners are not representative of the characteristics of the lens. In the extreme there may only be one observations at a corner of the sensor. It is not unreasonable to suggest that there will be a penalty to be paid for this in terms of the accuracy that can be achieved with observations at the corners of the field of view. Figure 4.1 shows an estimate of the percentage loss of targets per unit area.

In addition to this effect the operator will often point the camera to the centre of the object to be measured. As a result the centre of the image format will have more targets per unit area than around the edges. This combined effect of both of these target distributions could be compensated for in the weights attributed to these observations to ensure that the number of distribution of observations does not have an adverse effect of the measurement process.




Figure 4.1 Variation in the number of targets per unit area for a typical CCIR camera with a 4/3 aspect ratio

4.3 Errors due to lens aberrations

Lens aberrations would normally be regarded as a second order effect in photogrammetry and have largely been ignored in the digital camera era. However, these effects do exist and should form part of either the photogrammetric model or the weighting model. Lens aberrations, especially chromatic aberrations, have been considered in camera calibration since the early 1950s. For instance, Hothmer¹⁰ considered the effect of wavelength between calibration at one wavelength and use of the lens at another, a difference of 15 μm was cited from a study conducted at the time. Carman and Brown¹¹ found consistent differences between visual methods and photographic methods of camera calibration using laboratory methods. This was attributed to chromatic effects and resulted in calibration methods that closely simulated the conditions in which the cameras were used. Wight¹² concludes that “second order measuring errors due to lateral colour aberrations may be significant... Chromatic variations in astigmatism and field curvature contribute further to the displacement of the “centre of energy” of off-axis image points. Some residual coma is not unusual in wide-field lenses, and the chromatic flare varies as a function of aperture; this will cause an apparent shift in the centroid”. Wight also cites variations in intensity as another source of error as the effect of a comatic tail is only apparent in higher intensity images and lost in lower intensity images due to the thresholding effect of the film response. The influence of the radiometric qualities of target images has been investigated for digital still cameras¹³ and the optical quality of target images has been shown to be an influence on the precision and accuracy of photogrammetric networks¹⁴. Given that chromatic aberrations and coma are dependent upon radial distance from the principal point, a model for the systematic errors or a radial weighting scheme should be investigated.

5. 3-D OBSERVATIONS

5.1 Empirical assessment of the effect of target quality on target location accuracy

To see whether weighting can be successfully applied to photogrammetric measurements a series of experiments were performed with a metrology surface plate that was in good condition but of unknown flatness. First, 84 identical targets were placed on the plate in a regular pattern. The results from the photogrammetric adjustment process using four different Pulnix TM6CN cameras and an Epix frame-grabber were obtained as follows. The cameras were pre-calibrated using a 3-D test-field with appropriate rotations about the camera optical axis at each station. The lighting was varied between four epochs and the image residuals were noted. Most of the targets were of roughly equal intensity. The surface plate is approximately 460 x 600 mm in dimension.

Average target intensity	Image residuals (microns)
230	0.187
150	0.207
80	0.246
35	0.456

Table 5.1 Results from four adjustments with varying lighting

The accuracy (residuals from a plane fitted to the data using LSE) in the best case was 1 part in 22,000 of the object maximum dimension. The reported precision from the adjustment equated to 1 part in 33,000. This level of precision is around the best that could be expected from a CCIR format camera especially considering each of the cameras had been calibrated some weeks previously and only one image per camera station was used. It should be noted that in the case of the surface fitting, the surface itself is not perfectly flat and the random surface error is estimated to be of the order of 5-10 micrometres. What is of importance here is the variation in the image residuals with target intensity. The relationship between image residuals and target intensity is clearly indicated. However, the effect is not as dramatic as would be expected from theoretical image location algorithm capabilities, suggesting that there are other random or systematic errors that must be accounted for in a model of the image precision estimates for any weighting scheme.

To put the information gathered in the experiments to use, a number of additional targets were added to the surface plate. This time the targets were made from a retro-reflective material that was less efficient compared to the material used for the original targets. In this way target intensities that varied considerably were obtained. The area and peak values for these targets are illustrated in the figures 5.1 and 5.2.




Figure 5.1 Area of targets for the four images




Figure 5.2 Peak intensity of targets for the four images



Figure 5.3 Image of the surface plate

5.2 Assessment of the use of a weighting scheme on 3-D measurements

Given such a dramatic variation in size of the images of the targets a clear difference between the two sets was expected to be distinguishable in the final 3-D data. Given that the poor target images were sometimes of just one or two pixels in size, the weighting for the group of targets was chosen to be 3 micrometres and a weighting of 0.3 micrometres was used for the good targets. It may be that these values would not be extreme enough given that an image residual of around 0.2 micrometres is typically achieved for good quality target images. An image of the surface plate is shown in figure 5.3.

A photogrammetric network solution using standard LSE was used to compute the 3-D co-ordinates using the weighting scheme indicated. The results were then sorted according to the reported errors produced by the software. The errors fell into two distinct groups. The in-plane locations of the 3-D data are plotted in figure 5.4 where the triangles represent targets with the smallest errors and the squares represent targets with the largest errors. The in-plane locations of the targets in figure 5.4 are exactly correlated with the good and poor targets.




Figure 5.4 Selection of target grouping from error estimates




Figure 5.5 Graph of residuals from surface fitting to 3-D data - unequal weighting

Least squares fits of a plane were again performed to analyse these data sets. Given that the surface was expected to be flat it is possible to provide a statistical analysis of the Z coordinate accuracy. The results for the best data were an average standard deviation of 28 micrometres and a RMS error from the plane fitting software of 20 micrometres. For the poor data the results were an average standard deviation of 125 micrometres and a RMS error of 207 micrometres. An accuracy of approximately 1 part in 27,000 is a good result considering that both poor and good measurements were being used at the same time. To illustrate this result graphically the residual errors from the plane fit are plotted in figure 5.5 where the square symbols relate to the best data and the round symbols relate to the poorer data.

This result is analogous to the example of the two sets of data shown in figure 2.3. The method used could be refined further by establishing the correct weights for each target area or peak and applying a continuum of values. For the purposes of this preliminary investigation this was thought unnecessary as it would cause some difficulty in easily analysing the data.

5.3 Comparison of the unequal weighted scheme with the equal weighted scheme

As a comparison, the typical equally weighted scheme was tried with the same data set. Figure 5.6 shows the coordinate precisions as derived from the photogrammetric network solution, which shows a generally curved pattern for the X and Y coordinate precisions and a distinct layering for the Z coordinate precisions. The curved trend for the X and Y coordinate precisions is caused by more favourable intersection angles for targets near the centre of the surface plate. The layering for the Z coordinate precisions is caused by targets being imaged on either three or four exposures. This effect is also present

for the X and Y coordinate precisions, but is less obvious. However the precisions shown here are much more uniform than the results for the unequally weighted photogrammetric network solution.

Once more a plane was fitted to the data and figure 5.7 shows the residuals from the plane of best fit. There are no significant differences between the extents and randomness of the patterns of the residuals between figures 5.5 and 5.7, indicating that there is no bias in the equally weighted data set. If all data points from the unequally weighted data set are used to fit to a plane then the RMS error for the surface is 87 micrometres. The RMS error of the surface for the uniformly weighted data set is only marginally worse at 90 micrometres. The conclusion that can be drawn from this result is that, as could be expected, variation in the precision of the observables does not imply a change in accuracy. A LSE solution minimises the squares of the residual errors taking into account their weights, however in the absence of a bias the most probably true solution is relatively unaffected. However it is also true that error analysis is much more problematic for the equally weighted case, as compared to the unequally weighted case.




Figure 5.6 Graph of equal weighted error estimates




Figure 5.7 Graph of residuals from surface fitting to 3-D data - equal weighting

To assess whether a more accurate surface could be recovered from the equally weighted data, sixty targets with the largest residuals were removed from the data and the standard deviation computed again. This time a RMS error for the surface fit of 20 micrometres was obtained, which is an improvement over using good targets alone. The result obtained here is from good targets as well as poor targets that may have randomly been included in the set and is effectively the result of rejecting the worst of the better targets and retaining the best of the poorer targets. As could be expected, the target selection is dominated by the better targets. Figure 5.8 shows the in-plane locations of the selected targets and there does not appear to be any significant systematic effects as the locations are reasonably well distributed over the whole plate.




Figure 5.8 Selection of targets based on small surface fit residuals

6. CONCLUSIONS

This paper has discussed a topic of weighting of observed image coordinates that has received little attention in the close range photogrammetric literature in spite of its vital importance to rigorous results. The paper has also provided a simple review of the requirement for weighting and has applied it to a single case to justify the use of a variable weighting scheme. The results have been analysed by the use of a surface plate to provide an independent method of assessment of accuracy.

To use equal weighting is not an optimal solution in many cases. Use of equal weighting of observables where it is not appropriate has clear consequences for the ability of the LSE solution to attribute realistic errors to measurements. Use of variable weighting complicates the results of any given measurement projects as a single value will not describe the whole measurement process. However, being able to make a realistic assessment of any measurement is a valuable component in the error propagation process for most applications. It should be noted that the motivation for use of variable weighting is not to procure higher accuracy results, although identification of measurements may lead to better results in some cases and parameter estimates may be less biased by poor results. Further work is necessary to meet the objective of providing end users with error estimates that correctly reflect the precision of image measurements for close range photogrammetric applications.

ACKNOWLEDGEMENTS

The authors would like to thank Rudiger Kotowski, Erwin Kruck, and Giuseppe Ganci for their assistance with information on weighting in photogrammetric network computations.

REFERENCES

1. Personal communication, October 1998. Dr. R. Kotowski, K² Photogrammetry, Münsterstraße 19, D-48565, Steinfurt, Germany.
2. Personal communication, October 1998. Dr. -Ing. Erwin Kruck, Leibnizstr. 28, D-73461 Aalen, Germany.
3. Personal communication, October 1998. Mr. G. Ganci, Geodetic Services Inc., 1511 South Riverview Drive, Melbourne, Florida 32901, USA.
4. Topping, J. 1963. *Errors of observation and their treatment*. Chapman and Hall Ltd. London. 119 pages.
5. Mikhail, E. M., 1981. *Analysis and Adjustment of Survey Measurements*. Van Nostrand Reinhold.
6. Fraser, C. S., 1996. Network Design. Chapter 9, *Close Range Photogrammetry and Machine Vision* (K.B. Atkinson, Ed.), Whittles, Scotland, 371 pages, pp. 256-282.
7. Clarke, T. A. 1997. Spherical ball retro-reflective materials, Technical Report for Peter Cook PLC, 72 pages.
8. Clarke, T. A. Cooper, M. A. R. and Fryer, J. G., 1993. An estimator for the random error in subpixel target location and its use in the bundle adjustment. *Optical 3-D measurements techniques II*, Pub. Wichmann, Karlsruhe:161-168.
9. Shortis, M. R., Clarke, T. A., and Robson, S. 1995. Practical testing of the precision and accuracy of target centring algorithms. *Videometrics IV*, SPIE Vol. 2598, pp. 65-76.
10. Hothmer, J. 1958. Possibilities and limitations for elimination of distortion in aerial photographs. *Photogrammetric Record*, Vol. II, No 12. pp. 426-445.
11. Carman, P. D. and H. Brown, 1956. Differences between visual and photographic calibration of air survey cameras. *Photogrammetric Engineering*. 22(4): 623.
12. Slama, C. C. (Ed), 1980. *Manual of Photogrammetry*. (Fourth Edition) American Society for Photogrammetry, Falls Church, Virginia, 1056 pages.
13. Robson, S. and Shortis, M. R., 1998. Practical influences of geometric and radiometric image quality provided by different digital camera systems. *Photogrammetric Record*, 16(92) : 225-248.
14. Robson, S., Shortis, M. R. and Ray, S. - Vision metrology with super wide angle and fish-eye optics. *Videometrics VI*, SPIE Vol. 3641, pp 199-206 (1999)