ALGORITHM DEVELOPMENTS FOR AUTOMATED OFF-LINE VISION METROLOGY

J. O. Otepka [a], H.B. Hanley [b], C.S.Fraser [b,*]

[a] Institute of Photogrammetry & Remote Sensing, Vienna University of Technology, 1040 Vienna, Austria - johannes@avt.at
[b] Department of Geomatics, University of Melbourne, Victoria 3010, Australia - c.fraser@unimelb.edu.au, hanley@sunrise.sli.unimelb.edu.au

Commission V, WG V/1

KEY WORDS: Vision metrology, automation, sensor orientation, close-range photogrammetry, algorithms

ABSTRACT:

The goal of developing automated off-line vision metrology systems has largely been accomplished over recent years. The emphasis of research into measurement automation has thus moved on to implementation of more flexible and robust algorithms and processes in order to both improve performance and better accommodate data errors. This paper describes key features of the automatic measurement operation adopted for the *Australis* off-line digital close-range photogrammetric system. Particular aspects covered include target detection and validation within image scanning; the design of an exterior orientation device to support initial image orientation; image point correspondence determination based on the geometry of epipolar planes; and the integration of these processing phases with both a preliminary and final bundle adjustment. As well as discussing the fundamental models and computational schemes adopted, the paper will address error handling strategies. It will also consider practical aspects of these developments and briefly report on results of experimental performance tests.

1. INTRODUCTION

In general terms, the process of image measurement and photogrammetric triangulation in automated off-line vision metrology (VM) systems follows that of traditional multi-station close-range photogrammetry: image coordinates are measured and labelled, preliminary exterior orientation (EO) and object point triangulation are performed, and bundle adjustment with sensor self-calibration is employed for the final object point determination. With the automated approach, however, the traditional sequential point-wise process for image mensuration has given way to image scanning, where in a single operation all potential targets are detected and validated. Moreover, through the incorporation of specific target patterns, the concepts of automatic resection via an EO device and network building through coded targets have been realised, as has robust feature point correspondence determination.

The data processing scheme for off-line VM systems now generally follows a sequence of image scanning with target detection and validation; identification of target groupings constituting EO devices and coded targets; initial EO; correspondence determination, object point labelling and initial triangulation; and finally bundle adjustment. With the implementation of such an automated computational sequence, there is a need to take account of the impact of data errors. When compared to photogrammetric surveys with manual image measurement, two features of the automated process stand out as far as data errors are concerned: there are likely to be more observational blunders, and it is probable that these will arise through the addition of invalid targets rather than the omission of valid observations. As a consequence, the automatic measurement process must prove itself to be robust and reliable in the presence of both different image qualities, and many more data errors than might otherwise be expected.

In this paper the integrated image measurement and photogrammetric triangulation process of the *Australis* software system for off-line VM (Fraser & Edmundson, 2000) is described, with particular reference to the goal of achieving robust performance in the treatment of inherent data errors. The coverage includes the operations of image scanning and initial EO determination, image point correspondence determination, a 3-stage approach to photogrammetric triangulation, as well as a brief account of the results of performance tests. A more comprehensive account is provided in Otepka (2001).

2. IMAGE SCANNING

2.1 Target Quality

The first data processing step within an automated VM process is the detection and measurement of artificial targets in each image. Whereas image recognition remains a significant challenge in image processing and artificial intelligence, the use of high-contrast targets has considerably simplified the development of straightforward target detection algorithms for VM, though there is a necessity for the image scanning algorithms to take account of differing image quality. Under ideal conditions, the use of retro-reflective targets will yield near binary images and 'high-quality' targets, as indicated in Figure 1a.

In many practical instances, however, the contrast difference between the targets and background may not be as distinct, for example when utilising non retro-reflective targets or when operating in daylight conditions. It is thus necessary to account for such 'low-quality' images, exemplified by Figure 1b, as far as is practicable within the automatic image mensuration process. A more complex detection and validation process is warranted for low-quality images, which by their nature can lead in many instances to lower photogrammetric triangulation accuracy.

In this section the procedure developed for image recognition and detection, target validation and centroid measurement within the *Australis* system is described, with the discussion being confined to the use of greyscale images only. The strategy adopted takes into account the anticipated target quality variations, as indicated in Figure 1c, by including distinct processes for high- and low-quality images.

An analysis of the imaged targets in Figure 1c shows that irrespective of the contrast difference between the targets and

their background, all have a common feature, namely distinct edges. Figure 2 shows the respective intensity profiles through these targets. Although the distinct jump in grey values at the edge of the targets varies in magnitude, the edge gradient in each case is more or less the same (recall that we design targets to yield such edge gradients). Thus, the gradient is generally a better criterion for the detection of targets than the commonly adopted simple pixel-value thresholding (e.g. Atkinson, 1996; Luhmann, 2000).

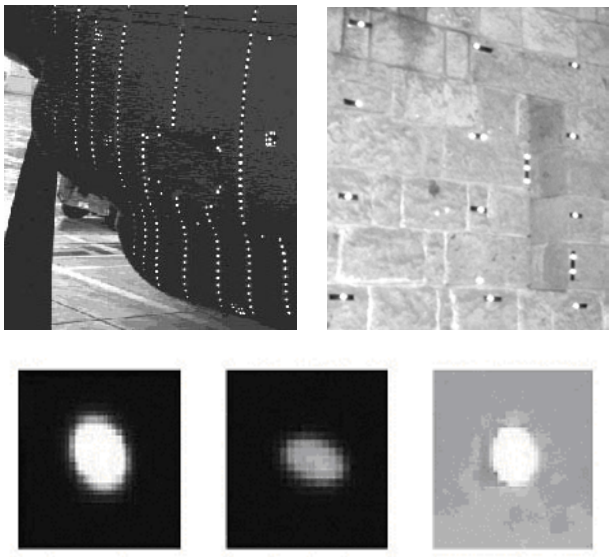

Figure 1. Target images in VM: high-quality, near-binary (a), low-quality (b) and a range of images (c).

The grey value gradient is expressed via the first derivative of the intensities. With the image being interpreted as a two-dimensional discrete function, the first derivatives in row (x) and column (y) direction can be given as

$$g'^{x}_{i,j} = g_{i,j+1} - g_{i,j}$$
$$g'^{y}_{i,j} = g_{i+1,j} - g_{i,j} \qquad (1)$$

In the literature, alternative operators are also described (e.g. Gonzalez & Woods, 1992), with some using more than two neighbouring pixels to compute the gradient at a certain position. However, derivatives using only two adjacent pixel values are well suited to near-circular targets, and they have the advantage of simple and fast computation. Performance considerations are very important for VM scanning algorithms, with higher contrast imagery generally leading to optimal computational speed.

2.2 Image Segmentation

Given the edge characteristics discussed, a line-orientated image segmentation algorithm that uses the gradient as a target recognition criterion would appear as a good candidate for VM applications. The gradient criterion alone, however, can deliver unsatisfactory image segmentation results in conjunction with low-quality images, since brighter areas are always recognised as targets if they are surrounded by darker regions. To alleviate this problem an extended segmentation algorithm has been developed for *Australis*, which uses both the gradient criterion and a minimum grey value threshold to detect target regions.

In comparison to algorithms that only use minimum grey value thresholding, the value of the adopted minimum threshold in the approach described is quite low. The threshold is not used so much to differentiate between isolated bright areas and targets. Instead, it only guarantees that pixels below a certain grey value, which are definitely non-target pixels, are rejected from further target measurement processes. The two thresholds used in the scanning process, the gradient threshold and the grey value threshold, need to be carefully chosen to ensure satisfactory results in the segmentation process. A fixed gradient threshold is generally appropriate for quasi-binary images as well as for low-quality images. However, the grey value threshold has to be separately selected for each image using histogram analysis.

With this development, the line-orientated scanning algorithm of *Australis* is implemented as follows: The process starts at the first pixel line of the image and goes through every pixel in that line. If the current pixel position fulfills both the grey value and gradient threshold criteria, the beginning of a target blob is indicated. The end of the target blob is detected either if the pixel grey value is below the minimum threshold or the gradient threshold is exceeded in the negative direction. If the 'negative' gradient criterion is detected at multiple positions, the last position is taken as the end of the target. It may happen that the criterion is met, without the detection of a target beginning. In this case, if the actual pixel is above the minimum threshold, the start of the target region has obviously been missed. To find the beginning of the missed blob, the algorithm searches backwards until the pixel grey value drops below the minimum threshold.

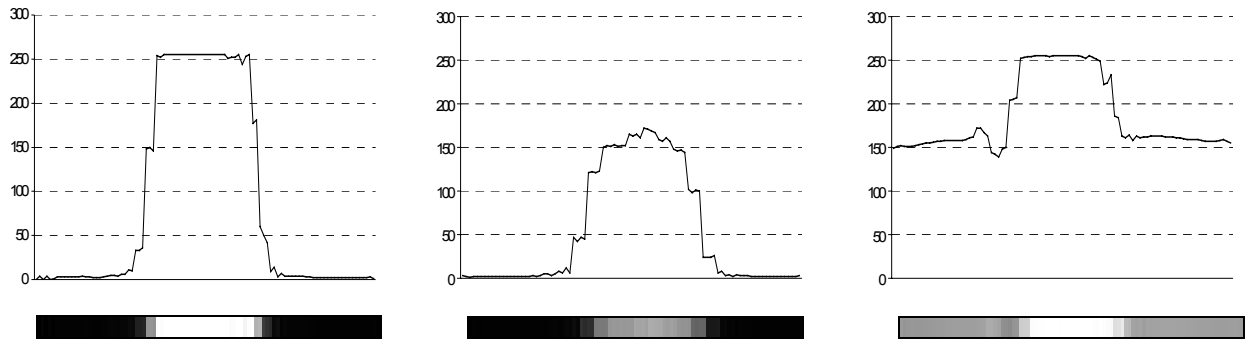

Figure 2. Intensity profiles through the targets of Figure 1c.

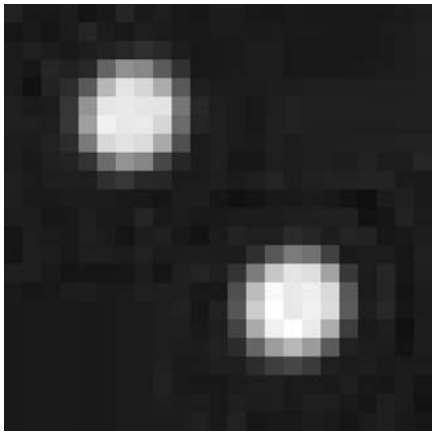

```
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 2 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 2 2 2 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 2 2 2 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 2 2 2 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 2 2 2 2 2 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

Figure 3. Test image and its corresponding segmentation matrix after scanning.

An image segmentation matrix and an array of region objects is built up during the scanning. A region object contains the connected pixels of a target. This allows a fast validation later on in the image measurement process. In detail, when the starting and finishing pixels of a target are detected in an image line (row or column), a new pixel region is created. This will contain the two edge pixels and all pixels between. Then, neighbouring pixel regions from subsequent lines are merged. If a pixel has been added to a region, the region label is set at the corresponding position in the image segmentation matrix. This allows fast detection of new pixels within the neighbouring rows. Figure 3 shows an example of a small test image and its segmentation matrix.

2.3 Region Validation

Once the image has been scanned and initially segmented, the important step of region validation is carried out so that regions can be classified as targets or non-targets. Various validation strategies can be employed to eliminate non-target regions prior to sub-pixel target centroiding. However, strategies which aim to differentiate target and non-target blobs with 100% correctness are both very difficult to design and require excessive computational effort and time. Considering that in VM we design for high redundancy and optimal geometric strength of the image station network, gross errors in target validation and centroiding can be tolerated to a considerable extent because such observation errors can be detected and eliminated within the blunder detection process of the bundle triangulation. Consequently, it is generally sufficient to design fast validation methods which reject most (say 90-95%) of the non-target regions, while accepting all legitimate target regions.

A second component of the target validation process is shape testing. Considering that the perspective image of a circular target will be an ellipse, an obvious shape verification mechanism is the best-fitting ellipse. However, simple tests such as blob size can also be applied to eliminate large bright areas (e.g. sky) and small reflectance hot-spots. In the procedure described, attention has been given to evaluation of the image scanning algorithm in the presence of differing image qualities. Practical experience has suggested that slightly different target validation strategies are warranted for high- and low-quality images (e.g. Figures 1a & 1b, respectively), even though the same segmentation process can be employed for both.

In the case of near-binary, high-quality images, it has been found that image scanning with only a simple and therefore fast validation process based on size testing is generally sufficient, with the small percentage of wrongly validated blobs being rejected as valid observations in the subsequent bundle adjustment processing. For low-quality images on the other hand, many more non-target regions are typically identified (incorrectly) in the image segmentation process. A shape test via ellipse fitting combined with a normalisation process for the detected target blobs has turned out to be most suitable for low-quality images. Within *Australis*, therefore, the following three validation processes are performed for low-quality images: normalisation, a size test and shape validation via a best-fitting ellipse computation. High-quality images generally require only the size testing.

The normalisation process removes 'dark blunder' pixels from the target region, these being falsely added because of inherent weaknesses within the segmentation algorithm. The procedure simply computes the mean and standard deviation of intensities within the region and classifies as dark blunders those pixels whose grey values are below the threshold mean minus a given multiple of the standard deviation. The process is needed to determine the correct boundary of the region for the ellipse fit test. The normalisation process also helps to remove regions that are small bright spots, e.g. reflectance hot-spots. Whereas VM targets can be expected to have very homogenous reflection intensity, as exemplified in Figure 2, the intensity distribution of small, non-saturated light-spots is generally quite inhomogeneous. The normalisation process classifies the pixels of such regions with strong variations in intensity as dark blunder pixels. This can then reduce the 'valid' region to the point where it is rejected within the subsequent size testing.

The shape testing process by best-fitting an ellipse to the target first determines the boundary of the region. The centre coordinates of boundary pixels are then used for the ellipse fit computation. Recent ellipse fit tests have been described, for example, by Fraser & Shao (1997) and Luhmann (2000). Experience has shown that shape testing via a best-fitting ellipse is a reliable criterion for region validation for large targets (say greater than 5x5 pixels). However, small regions invariably pass this test, since the redundancy in the boundary point distribution is often too little for reliable least-squares estimation of ellipse parameters, and for subsequent computation of departures of the region boundary from the best-fitting ellipse.

2.4 Sub-Pixel Centroiding

Once a target region has passed the validation process its exact centre needs to be determined. Though various centroiding and template matching methods can be found in the literature (e.g. Atkinson, 1996; Luhmann, 2000; Shortis et al. 1994), the intensity-weighted centroiding approach is most often adopted in automated VM because it is simple and very fast to compute:

$$\begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = \frac{\sum_{i=1}^{n}\sum_{j=1}^{m} g_{ij} \begin{pmatrix} x_{ij} \\ y_{ij} \end{pmatrix}}{\sum_{i=1}^{n}\sum_{j=1}^{m} g_{ij}} \qquad (2)$$

Here, x_{ij} and y_{ij} are row and column coordinates of pixels within the target blob, g_{ij} is the corresponding grey value, and (x_0, y_0) are the final centroid coordinates. The main issue of concern in any high-precision centroiding method is background thresholding. The threshold has to be chosen as low as possible, so it removes all background noise without deleting valid edge pixels of the target region. The target edge pixels preserve the important geometric information for sub-pixel centroiding, which in VM can reach accuracies to 2-3% of the pixel size. For today's large area CCD cameras, this can translate to an image coordinate accuracy of close to 0.2 μm.

3. EO DEVICE DETECTION

Effective automatic detection of the EO device is necessary for initial spatial resection of the images forming the VM network, which is in turn critical to a successful solution of the image point correspondence problem using epipolar geometry. Multiple strategies for automatic recognition of target groupings within an image can be designed. Theoretically, recognition strategies for coded targets could also be employed for the EO device, but the main difference between coded targets and an EO device is that the points on the EO device have known XYZ coordinates within an arbitrary datum, which initially defines the reference coordinate system of the full VM network. One reason that EO devices are designed to be noticeably larger than coded targets is to guarantee sufficient geometric strength for the calculation of EO parameters for all images via either closed-form resection or, less commonly, relative orientation

Whereas EO devices in commercial VM systems such as *V-STARS* (GSI, 2002) use recognisable fixed patterns of targets, a more flexible recognition and initial EO concept has been developed for *Australis,* where the user has the flexibility to define their own EO device under minimal design constraints. Figure 4 shows an example EO device suited to *Australis.* This design incorporates one basic constraint, namely that all points must be enclosed within a bounding shape, a circle in this case (a rectangle, square or triangle would also suffice). To achieve better approximations of orientation parameters, one of the five targets is usually non-coplanar with the others. Use of the EO device requires no more than the user specifying the *XYZ* coordinates of the EO target points. No labelling is necessary so long as the targets are asymmetrically distributed, since their relative positions are decoded in the recognition process.

The EO detection procedure can be subdivided into three stages: device detection, target recognition and pattern decoding, and closed-form spatial resection. For the adopted procedure, the pattern decoding and resection stages are closely interrelated. At the initial image scanning stage, the primary recognition criterion for the EO device is the closed boundary, which need not be of a specific shape. The segmentation process treats the boundary in a similar way to other targets when forming an appropriate segmentation matrix, an extract from which is shown in Figure 5 for an EO-device boundary section. During the EO validation process it is checked whether the segmented region has a 'hole', ie an enclosed region within the boundary, the test being performed only if the enclosed area contains a specified minimum number of pixels.

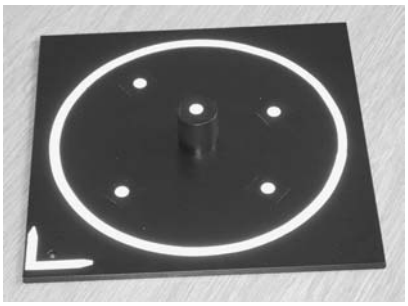

Figure 4. An example EO device for *Australis.*

Following the image scanning, the segmentation matrix for a detected enclosed region is tested to establish whether it contains others targets. If the number of such detected targets matches that for an EO device, the region is flagged as an EO-device candidate. The design for *Australis* can account for the detection of multiple EO devices in an image, though this feature has not been applied in practical applications so far. Once the correct enclosed region and all EO-device targets are detected and measured, the pattern decoding and closed-form resection are performed.

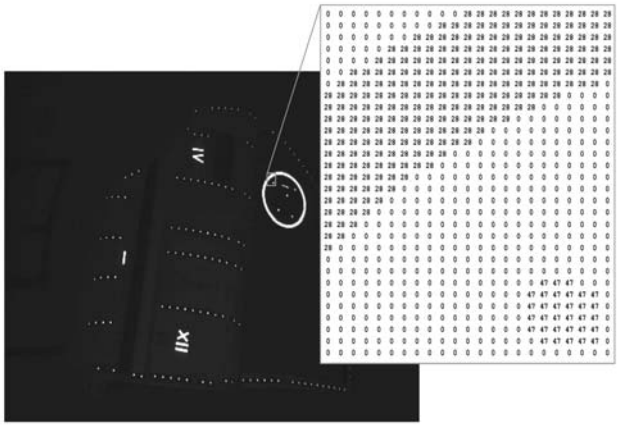

Figure 5. Image of EO device with part of segmentation matrix.

To determine which image point corresponds to the correct object point, a simple but rigorous strategy is employed. An algorithm tries all permutations of possible resections to find the most likely solution. If the geometric arrangement of EO points is favourable (an asymmetric pattern), the sought after permutation is the one with the lowest RMS value of image coordinate residuals in the resection. This algorithm has proven to be very robust in practical applications, even in cases where sensor calibration parameters (primarily lens distortion and focal length) are only approximately known, or where the lens distortion is assumed to be zero and a very course initial principal distance is adopted.

As can be imagined, the resolving of point correspondences for EO takes quite some computation time if all point permutations on the EO device are considered. A total of 120 (or 5!) resections need to be performed for a 5-point array before the correct solution is known. On a Pentium II/266 PC, the EO device detection takes two seconds whereas the scanning of a single high-quality image (1536 x 1024 resolution) is performed in only one second. Though computation time is usually not critical in off-line VM, an optimised detection algorithm has been implemented in *Australis* to further enhance processing speed. If the targets on the EO device are labelled in a specific order, the correct EO can usually be achieved with only four resections. In the adopted labelling scheme, the point closest to the centre of gravity of all EO targets (the raised point in Figure 4) is initially labelled, and then the remaining points are incrementally labelled in a clockwise fashion. EO devices which are labelled as such are detected within one second.

4. POINT CORRESPONDENCE DETERMINATION USING EPIPLOAR GEOMETRY

4.1 Epipolar Plane Angle Geometry

Epipolar geometry can be employed to locate homologous image points in a multi-image VM network. Therefore, reasonable approximations of the EO parameters for all camera stations are required, this being provided via the EO device. In the literature, various techniques having different characteristics in relation to speed and robustness can be found for image point correspondence determination (e.g. Furnee et al. 1997, Chen et al. 1993). All techniques, however, must deal with correspondence ambiguities that can easily arise in the case of only two images or in multi-image networks where there is a dense array of targets (real and false). The technique adopted for *Australis* is based on epipolar plane angles and is similar to that proposed by Sabel (1999). It is fast and capable of accounting for lens distortion effects in correspondence determination, which is often critical when using modern CCD cameras.

The basic principle of the epipolar plane angle method is that rays of corresponding image points subtend identical angles to any epipolar plane, as illustrated in Figure 6. Epipolar planes always contain the projection centres of both camera stations. In order to calculate the required angles for image points of both images, an arbitrary epipolar plane is first defined. The normalised vectors n_1 and n_2 are then computed with the applied constraints that both the vectors and the baseline are normal to each other, and that n_1 is contained by the defined epipolar plane. The corresponding epipolar plane angle for each image point in both images can then be calculated:

$$\vec{r}_i = R\begin{pmatrix} x_i \\ y_i \\ -c \end{pmatrix} \qquad (3)$$

$$\theta = \arctan(\vec{n}_2 \cdot \vec{r}_i \, , \vec{n}_1 \cdot \vec{r}_i) \qquad (4)$$

Here, x_i and y_i are the coordinates of an image point corrected for lens distortion, R is the rotation matrix for the current camera station and c is the corresponding focal length.

4.2 Resolving Ambiguities

An image point can have an infinite number of homologous points in the second image, the applicable constraint being that they all have the same epipolar plane angles, or in other words, they are positioned on epipolar lines. When lens distortion is taken into account, the epipolar line transforms to a curve. Such correspondence ambiguities can be resolved or at least greatly reduced by using a third image. Consider the case of Figure 7, where the point I_1 in image 1 has two possible homologous points in image 2 (J_1 and J_2), because they are positioned on the corresponding epipolar line in the image. The two object points P_1 and P_2 can be computed and are projected into the third image. If one of these projected points is coincident with an existing image point (K_1), within a given tolerance, the ambiguity is resolved and P_1 is then the correct 3D object point.

In the resection-driveback image-point measurement process, where the object point coordinates and EO are utilised to predict image point location, lens distortion has to be considered for the back projection to deliver correct image point locations, and therefore correct correspondence determination. An inverse function for lens distortion is therefore needed. Because no closed-form of this function exists, a numerical inversion that typically converges within a few iterations is performed.

4.3 Correspondence Determination Algorithm

The entire correspondence determination process within *Australis*, which was designed for highly-convergent VM networks where most images 'see' a large portion of the target range, is briefly outlined below. The process takes as input several images with known EO and their corresponding image point measurements.

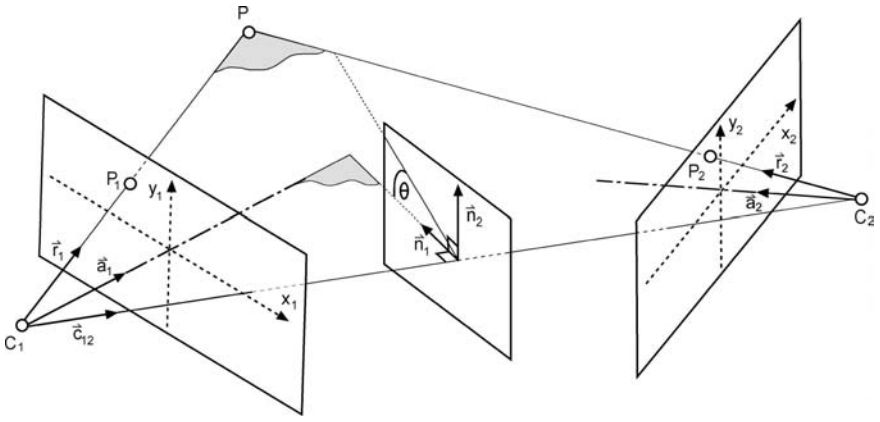

Figure 6. Epipolar plane angle geometry.

The main loop iterates over all possible image pair combinations. Within this loop, the algorithm computes all epipolar plane angles of the two selected images and stores the values in an angle array. To find homologous image points, the angle array is sorted. At this stage only object point candidates are determined. To guarantee a certain accuracy of the candidate point coordinates, only rays with a specified minimum intersection angle are accepted. An object point candidate is classified as accepted if, using back projection, it is found in at least one other image. Consequently, the complete triangulation information of an object point is built up once it is found to be an object point candidate. In the case of ambiguities (Figure 7), only one of all possible object point candidates may be accepted as the true object point. After rejecting or accepting all object point candidates of the image pair, the next two images are selected. In the new image pair only epipolar plane angles for unassigned image points are computed, which makes the computation faster for every new image pair.

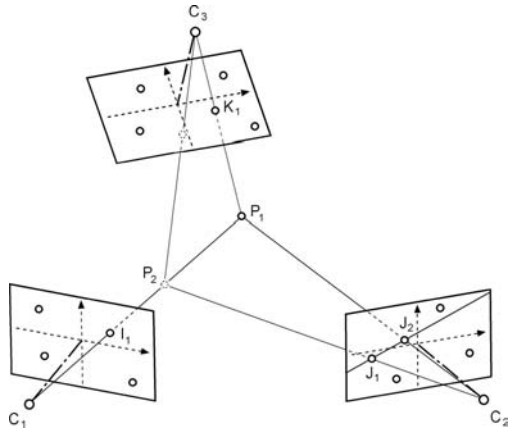

Figure 7. Resolving correspondence ambiguity using a third image.

5. BUNDLE ADJUSTMENT AND QUALITY MONITORING

The integrated bundle adjustment forms a logical final step within an automated off-line VM computation process. Beside the determination of object point coordinates and EO parameters, the final bundle adjustment typically includes sensor self-calibration. Sensor calibration is a basic requirement for the achievement of high triangulation accuracy, especially in the case of off-the-shelf digital cameras. Bundle adjustment strategies include methods for blunder detection, which leads to the idea of employing an initial bundle adjustment for the detection and correction of gross errors that arise in the automated measurement process. Least-squares adjustment methods are well suited to error detection because all decisions can be statistically based.

Considering possible observational blunders in both image scanning and point correspondence determination, a correction algorithm needs to be developed. There are three types of errors that can occur:

- *Recognition errors in image scanning*: It is assumed that all targets are found correctly and that only a few non-target regions (e.g. bright areas or spots) are falsely accepted. If a non-target feature is found in at least three images, the correspondence process will then find a corresponding object point. Because the located object point is a real, physical object feature the final triangulation network is rarely disturbed. Additionally, these object points are often rejected in early iterations of the bundle adjustment because of their anticipated larger image coordinate residuals. Artificial targets (e.g. retro-reflective targets) have a very high centroiding precision which is rarely matched by 'natural' target features. Consequently, the interference of these errors in the automatic measurement process is rarely of any practical consequence.

- *Correspondence determination errors*: This type of error, which arises from ambiguities in the epipolar-based point matching, can always be expected to be present to some extent, although it rarely causes serious concern in strong VM networks. In rare situations, often with very dense target arrays, unresolved false detections occur. Whereas with recognition errors the effect is to generate additional object points, here incorrect points with completely wrong coordinates might result. Though these observational blunders interfere with initial triangulation, they are usually rejected in the bundle adjustment if the number of images in which they appear is low.

- *Duplicate point errors*: Situations can arise such that a single target is identified as two (or more) discreet object points. This occurs when two or more subsets of images exist where the target passes the correspondence test in each subset, but not the set as a whole. The algorithm, not knowing that the target is unique, assigns it as if it were two discreet object points with slightly different XYZ coordinates. The correction for this type of error can occur later, after the final bundle adjustment reveals that the coordinates for these discreet points are coincident. This situation occurs when the EO within each image subset is relatively strong but not strong between each subset.

Analysis of the bundle adjustment progress within various *Australis* projects has shown that the first two error types are accommodated (bad data rejected) in early iterations of a bundle adjustment. Even the third error type can be identified early on in the processing because later iterations effect only very small shifts of object points and these multi-object points are close to coincident in space. These considerations have led to the development of a final three-stage triangulation process for *Australis*:

- *A 'fast' bundle adjustment:* This is an initial adjustment to 'near' convergence, where a larger than usual RMS threshold for image coordinate residuals is used as a convergent criterion; 0.5-3μm turns out to be an appropriate threshold range for highly convergent VM networks. These networks typically achieve a final RMS value of 0.15 – 0.3 μm. The camera is not fully self-calibrated in this initial adjustment, though the focal length and the principal point coordinates are usually carried as additional parameters. This results in a performance gain, and it is generally a more robust calculation strategy if only an inaccurate camera calibration (or no calibration) is at hand.

- *Data cleansing stage:* Next, a correction process is performed where the results of the initial bundle adjustment are analysed. Object points with less than three active imaging rays are removed and all the corresponding point measurements are omitted. This process also corrects blunders arising from the first two error sources listed above. To correct for duplicate point errors, a simple merging process is performed where close neighbouring

object points are merged into a single point, with appropriate re-labelling of image point observations.

- *Final bundle adjustment*: A final bundle adjustment with sensor self-calibration is carried out for all validated observational data.

The final computation in three stages ensures robust and accurate results, without the influence of observational blunders or gross errors made in automatic feature matching and labelling. Throughout the 'fast' bundle adjustment, the performance of the process is successively improved with each iteration, especially within VM networks comprising large numbers of images and points.

6. PERFORMANCE ANALYSIS USING PRACTICAL MEASUREMENT PROJECTS

The performance and reliability of the presented algorithms, as implemented in *Australis*, have been proven in numerous practical applications covering networks of various size and objects of various scale. Here we use two practical measurement projects to illustrate the performance of the automated VM procedure developed. The projects where conducted with two different cameras, and the two objects were of approximately the same size, the first, Project 1, being a car door with 130 retro-reflective targets (Figure 8), and the second, Project 2, a plastic mannequin with more than 300 targets (Figure 9). Whereas the car door was recorded with six images using a *GSI INCA* camera, nine images were recorded with a *Kodak DCS 420* camera in the second test project.

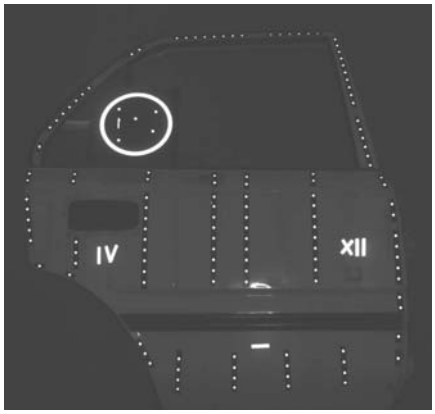

Figure 8. Image from car door project, Project 1.

For quality assessment of the automated measurement process, the correctness and accuracy of the results were analysed. It was also investigated if all possible target information was used for the final bundle adjustments. Additionally the performance of the process on two different computer systems was assessed. An initial quality statement can be made by examining the RMS value of the image coordinate residuals from the final bundle adjustment.

In the car door network, an RMS value of 0.15 μm (better than 0.02 pixel) was reached, which is an appropriate accuracy for an *INCA* camera project. Though the plastic body was imaged in a 9-station network, the final RMS of the residuals reached only 0.24 μm (0.03 pixel). The *DCS 420* camera did not have a stabilised interior orientation, and consequently the achieved RMS is also assessed as an acceptable value under the circumstances (the same camera often yields residuals at the 0.25μm level in *Australis* projects with retro-reflective targets). The focus of attention in these test measurements was upon internal rather than external accuracy, though it is noteworthy that in each project the mean standard error of *XYZ* object point coordinates surpassed a relative accuracy of 1:100,000 of the object size.

The investigation of the two final VM networks showed that all targets were both found and computed correctly in each project, and no additional non-targets were triangulated. Analysis of the imaging-ray statistics for object points showed that most targets were measured in all images, though occlusions or highly oblique incidence angles for certain targets prevented this in some images. The results of the automated process were very satisfactory and it should be mentioned that the data cleansing procedure within the 3-stage bundle triangulation process improved the results of the overall measurement, especially for the second, 300-point project.

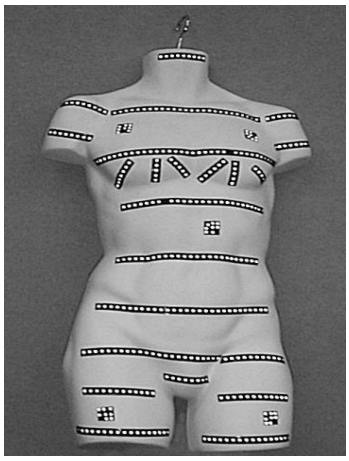

Figure 9. Targetted plastic mannequin, Project 2.

Regarding performance, the projects were computed both on a *Pentium II Celeron* 433 Mhz computer with 128 MB RAM and on an *AMD Athlon XP* 1600 computer with 512MB RAM. The *Pentium II* system took 12 and 30 seconds, respectively, for the complete measurement process for Projects 1 and 2. The corresponding computation times with the *AMD* system were substantially shorter, being 7 and 14 seconds. Details can be found in Table 1. These computation times are certainly acceptable for practical automated off-line VM, and it is doubtless the case that improvements in PC technology will bring even faster processing times.

Table 1: Computation times for the automated measurement processes in *Australis* (all values are in seconds).

Processor	Project	Scanning and EO detection	Correspondence Process	Fast Bundle Adjustment	Data Cleaning	Final Bundle Adjustment	**Overall Process**
AMD Athlon XP 1600	1	3.9	0.3	0.5	0.2	0.7	**7.3**
	2	5.4	1.4	2.1	0.4	2.9	**14.3**
Intel Pentium II Celeron 433 Mhz	1	9.0	0.5	1.0	0.2	1.5	**12.1**
	2	12.4	4.5	5.3	0.4	7.2	**29.9**

7. CONCLUDING REMARKS

The automated image measurement and multi-stage photogrammetric triangulation process developed for *Australis* has proved both robust and computationally efficient. Not surprisingly, optimal reliability, speed and accuracy are typically attained with 'high-quality' images of retro-reflective targets. The results obtained with 'low-quality' images are, nevertheless, also reasonably robust, showing that automated processes have potential in situations where lower contrast, non retro-reflective targets are employed, and where the image scanning process identifies a higher than desirable number of candidate targets which turn out to be invalid data.

8. REFERENCES

Atkinson, K.B. (Ed.), 1996. *Close Range Photogrammetry and Machine Vision*, Whittles, Scotland, 371 pages.

Chen, J., Clarke, T.A. & Robson, S., 1993. An alternative to the epipolar line method for automatic target matching in multiple image 3-D measurement, In: *Optical 3-D Measurement Techniques II*, (Gruen, A., Kahmen, H., eds.), Wichmann, Karlsruhe, pp. 197-204

Furnee, E.H., Jobbagy, A., Sabel, J.C., van Veenendaal, H.L.J, Martin, F. & Andriessen, D.C.W.G., 1997. Marker-referred movement measurement with grey-scale coordinate extraction for high-resolution real-time 3D at 100 Hz, *SPIE Proceedings,* Vol. 3173, pp 357-369.

Fraser, C.S. & Shao, J., 1997. An Image Mensuration Strategy for Automated Vision Metrology. In: *Optical 3-D Measurement Techniques IV*, (Gruen, A., Kahmen, H., eds.), Wichmann, Heidelberg, pp. 187-197.

Fraser, C.S. & Edmundson, K.L., 2000. Design and Implementation of a Computational Processing System for Off-Line Digital Close-Range Photogrammetry. *ISPRS Journal of Photogrammetry & Remote Sensing,* 55(2): 94-104.

Geodetic Services, Inc., 2002. Company website, accessed May: www.geodetic.com.

Gonzalez, R.C. & Woods, R.E., 1992. *Digital Image Processing*, Addison-Wesley, pp 418-423.

Luhmann, T., 2000. *Nahbereichsphotogrammetrie, Grundlagen, Methoden und Anwendungen.* Wichmann, Heidelberg, 571 pages.

Otepka, J., 2001. *Algorithms and their Implementation in an Automated Close-Range Photogrammetric System*, Diploma thesis, Vienna University of Technology.

Sabel, J.C, 1999. *Calibration and 3D Reconstruction for Multi Camera Marker Based Motion Measurement*, PhD thesis, Faculty of Applied Physics, Technical University of Delft, Netherlands.

Shortis, M.R., Clarke, T.A. & Short, T., 1994, Comparison of some techniques for the subpixel location of discrete target images, *SPIE Proceedings*, Vol. 2350, Paper #25.