www.photometrix.com.au

Version 2.0, May 2016

CameraCalibrator
User Manual

CAMERA CALIBRATOR

PHOTOMETRIX © 2016

How to use this *CameraCalibrator* Operational Guide

Introduction

STEP 1: Read Section 1, Overview of Automatic Camera Calibration, to familiarize yourself with the basic photogrammetric technology and requirements

STEP 2: Review the Quick Reference to Menus, Toolbar Buttons, Cursors and the Image and 3D Views provided in Appendix A, page 30

Conducting a Calibration Project – there are three options here, automated calibration using targets, automatic targetless calibration and calibration using manual image measurement

STEP 3: Follow the project set-up described in Section 2, page 9 and conduct photography

STEP 4: With acquired images loaded into the project, perform image measurement, network orientation and self-calibration via one of the following three options

Automatic Calibration using targets (coded + uncoded)

Section 3, page 12

Automatic Targetless Calibration using feature-based matching

Section 4, page 16

Calibration via manual image measurement, using targets or natural features

Appendix B, page 35

STEP 5: Output of camera calibration results, Section 5, page 20; *projectname*.cal and *projectname*.txt files generated when project is saved. Calibration report generated.

This concludes the basic camera calibration process, with further stages being optional

OPTIONAL STEP 6: Set absolute scale and XYZ axes alignment – primarily for indication of associated accuracy attained in photogrammetric 3D point determination within the self-calibrating bundle adjustment, Section 6, page 21

OPTIONAL STEP 7: Additional runs of self-calibrating bundle adjustment and output of further refined camera calibration parameters, Section 7, page 24; *projectname*.cal and *projectname*.txt files generated when project is saved. Calibration report generated.

OPTIONAL STEP 8: Camera self-calibration in the presence of object space constraints, namely camera station positions and ground control points, with a focus of calibration for UAV/drone photogrammetry, Section 8, page 27

Contents

Calibration Network: Image and Object Point Configuration

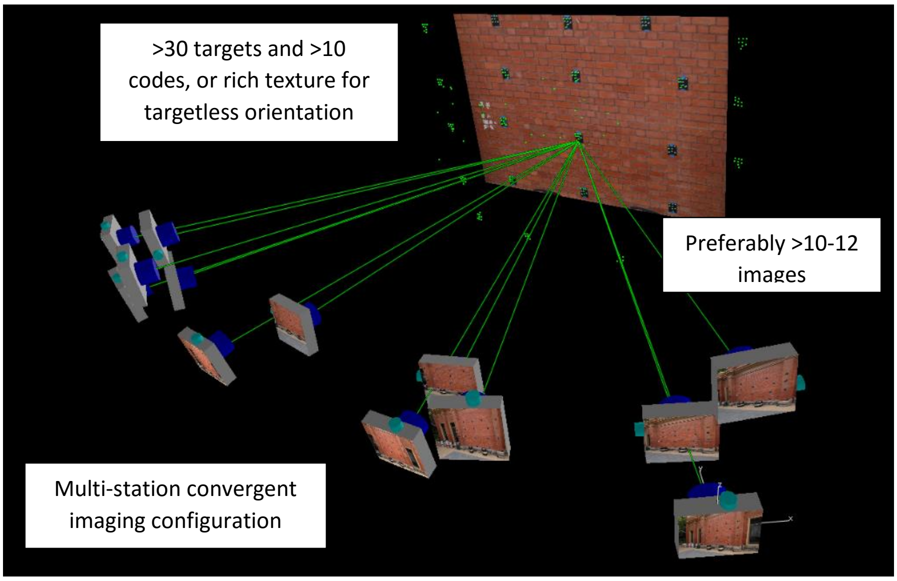

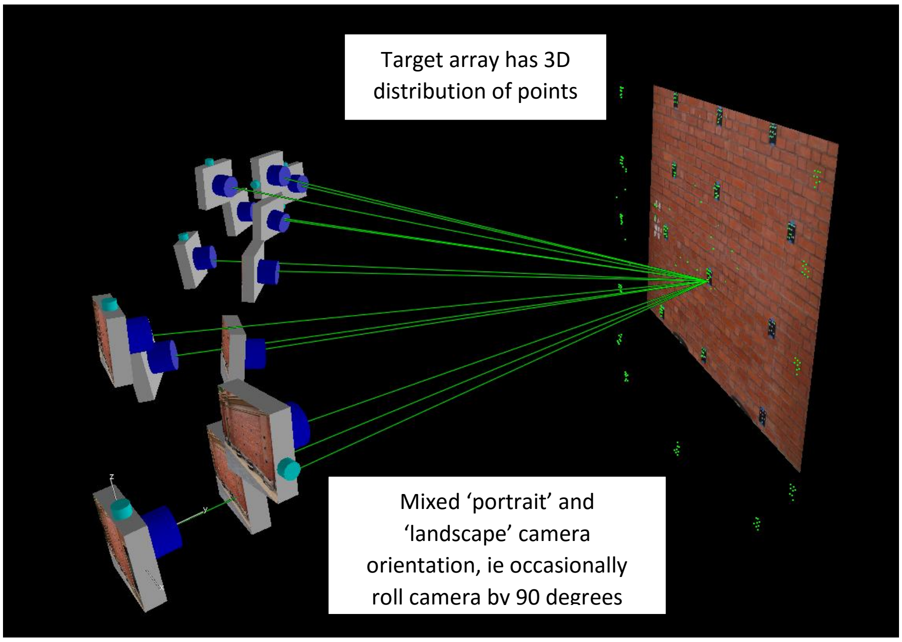

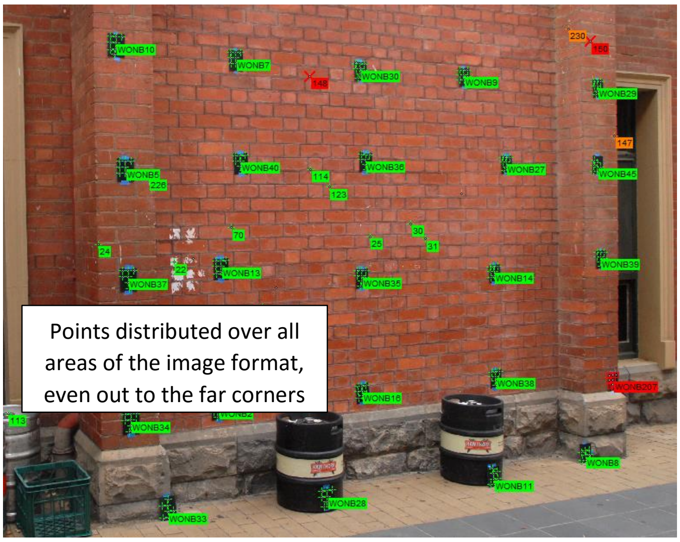

1. Overview of Automatic Camera Calibration

Two options are available within the *CameraCalibrator* for automatic camera calibration, one employs targets and the other is targetless. The principles for each are the same, since the recovery of camera calibration parameters is via a self-calibrating bundle adjustment in both cases.

1.1 Basic Rules for Camera Setting for Self-Calibration

There are four basic rules that apply to recording images for camera self-calibration:

- The camera lens should not be refocused during the photography session (Autofocus turned off).
- Similarly, the lens zoom should not be adjusted during the photography session.
- If the camera has an 'auto rotate' function, which digitally rotates the recorded image, the function should be turned off (Autorotate turned off).
- The imagery should be recorded at the highest possible resolution (i.e. image size) and highest quality. For example, a 16 megapixel camera should record 16 megapixel images at a highest possible image quality (JPEG image format is fine).

If retro-reflective targets are to be used, the camera will require a flash. If the flash is too bright and its intensity cannot be controlled, sometimes taping a small piece of white paper over the flash will cut down the intensity so as to not saturate the retrotargets.

1.2 Camera Setup and Scene Layout

If a target array is being employed, the following should be kept in mind:

- *Number of targets.* A mixed set of 30-50 different coded targets (also called 'codes') or single 'dot' targets is required for automatic calibration. Note that retroreflective codes are re-usable and by affixing codes to plastic backings, they will last a lot longer, and remain flat. Whereas a minimum number of codes for a mixed target range might be 10, it is possible to have the target array comprising only codes. The target dots should be at least 5 pixels in diameter in the images.
- *Shape of the target array.* It is preferable for the target field to have 'depth', for example having some of the targets out of the plane by 10-20% of the photographic distance. A target array with 3D distribution is much more desirable than a planar array (see Figure 1.1)
- *Camera-to-object distance*. The size of the target field will depend upon the distance from which images are to be taken. To determine this, first set the camera at the focus and zoom setting that you intend to use for subsequent photogrammetric measurement, as it is critical to calibrate the camera at the same zoom/focus that will be normally be employed. Typically, the focus will be set to infinity and the zoom at one of its 'hard stops', usually zoomed fully out (widest field of view). Once these settings are made the operator can balance the requirements of imaging distance versus target field dimensions. Note that in determining the distance, the sharpness of focus in the images is not of paramount importance if targets are employed. An accurate calibration can be performed with slightly defocused images. This can arise from infinity focus being set, but then the camera used over a distance of only 3-4m or so.
- Both target array stability and fixed focus/zoom are very important. The targets cannot be allowed to move at all. If the camera has a manual focus or focus lock, it is very useful to set this and then leave it set for the subsequent 3D measurement projects. Also, attention needs to be given to lighting, with images displaying optimal target intensities and underexposure, depending on the coded target material being used. Experiment with your camera & lens settings to determine optimal values.
- A suitable image network geometry for camera calibration is shown by the 12-station network in Figure 1.1. The convergence angle (both horizontal and vertical, if feasible) between the outer rays should be 70˚ to 100˚; it is also very desirable to have convergence in both directions.

- Note the use of orthogonal roll angles for the camera; three images are in ‘landscape’ orientation, and nine in ‘portrait’. **It is absolutely critical that this 90˚ variation of roll angle is present.** It is not important which images are rolled 90˚, but a reasonable number should be.
- A sufficient number of images is generally 10-12. Due to the process being automatic, there is essentially no cost in time in adding additional images, so even 30 images is practical.
- The target field should fill as much of the image format as possible in each image. It does not matter if some points are missed on certain images, but it is important to have image point coverage throughout the full format of the focal plane.
- Camera calibration should be performed periodically, say every 3-6 months for normal digital cameras, or more often if the camera is accidentally dropped or it is used in a relatively rough environment.

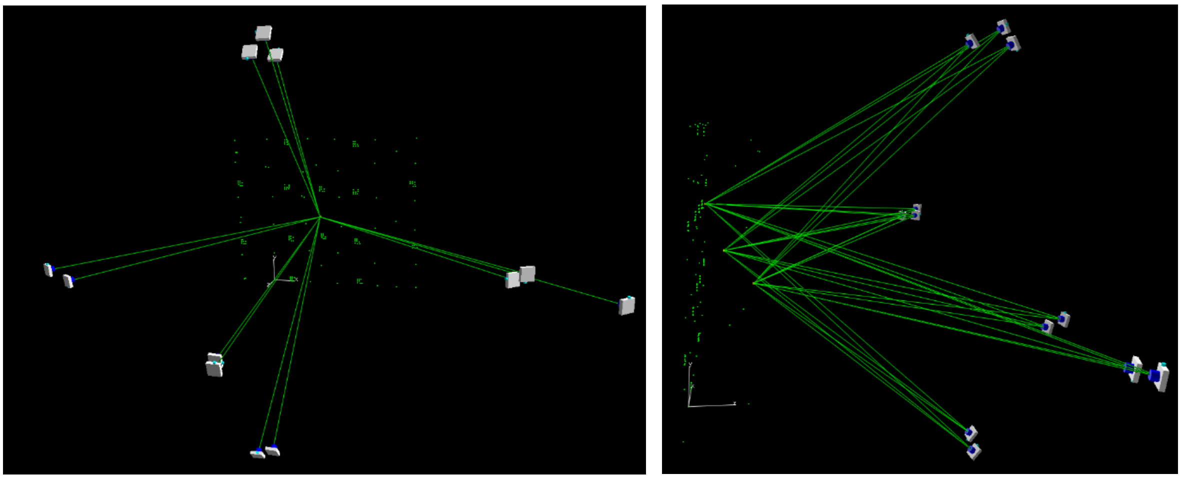

Figure 1.1

In the case of targetless orientation and self-calibration, the geometric requirements remain the same, with two special considerations: firstly, the scene must be richly textured to support feature detector-base image matching, and also the convergence angle may need to be lessened for the same reason. To compensate to some degree for the lessened convergence, additional images can be added, as indicated by the network geometry in Figure 1.2

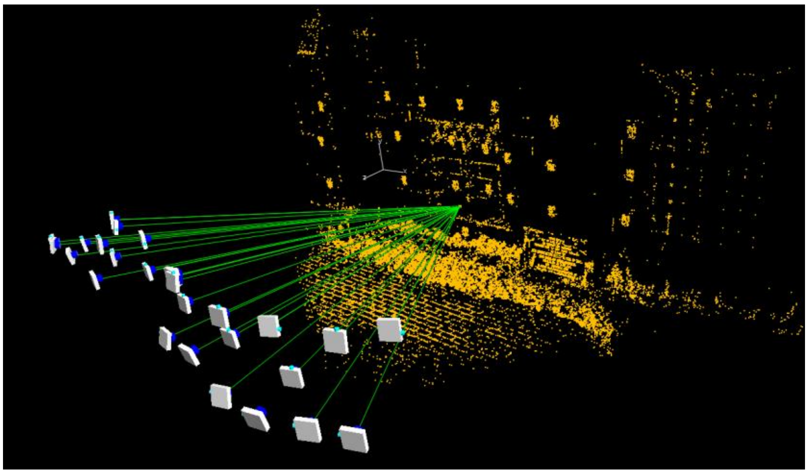

Figure 1.2

1.3 Calibration Process

The basic procedure for fully automatic camera calibration within the *CameraCalibrator* follows closely that for a normal automatic photogrammetric measurement:

i) *Establish the target array or object scene for targetless orientation*. The principle of a well-spread point array which fills the image format and is preferably non-planar should be adhered to. A suitable array is shown in Figure 1.1. (If codes-only are to be employed, as in the *AutoCal* procedure, the use of more than 15-20 codes would be warranted).

ii) *Determine the camera station geometry and record the images as JPEGS (at full resolution and highest quality).* The primary items to remember are that: a) the camera station network provides strong convergence angles; b) The images are taken such that at least one-third of them are 'rolled' (90˚ rotation); c) the codes are distinguishable with each code dot being of high contrast and larger than 5 pixels in diameter in the imagery, or the targetless scene has good texture for image matching; and d) it is preferable if all codes do not lie in the same plane.

iii) *Load the images into a new CameraCalibrator project*, as described in Chapter 2. At this point the project camera will have been identified and the image thumbnails will be displayed in the thumbnail window.

iv) *The automatic image measurement and camera self-calibration process is initiated*, the exact procedure for which will be described in Chapter 3 for automatic calibration via coded targets, and in Chapter 4 for targetless orientation. Appendix B describes self-calibration via manual image measurement.

v) The *derived camera calibration parameters* may be any or all of those listed in the dialog below, Figure 1.3, which is obtained by double-clicking on the camera icon (see Chapter 2). The corresponding image coordinate correction functions are provided in Appendix C on page 45.

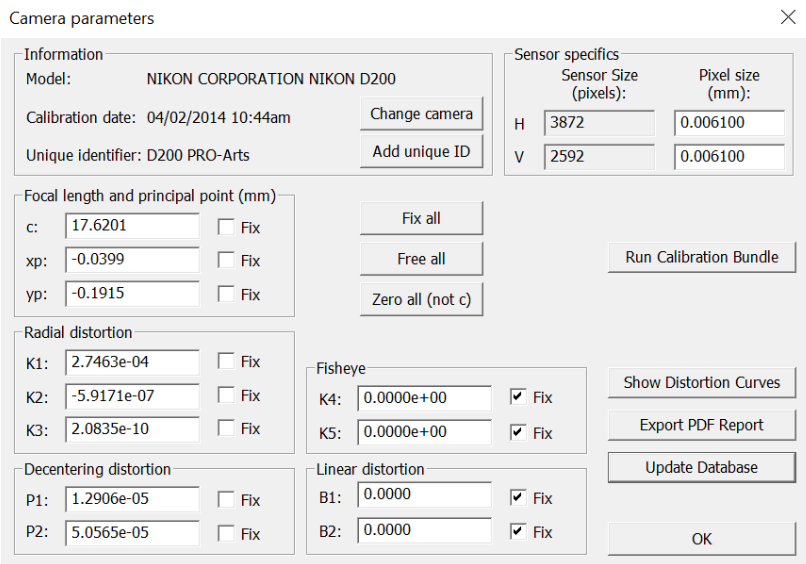

Figures 1.3

- *Focal Length and Principal Point Coordinates.* The principal point offsets indicate by how far the optical axis of the lens is displaced from the centre of the image format. Generally x_p and y_p are close to zero (e.g. 0.5mm or less). The 'focal length' actually refers to the Principal Distance, which changes with focusing. The nominal lens focal length usually relates to infinity focus.

- *Radial Distortion.* The three parameters K_1, K_2 and K_3 describe the radial lens distortion correction for the camera lens. Radial distortion can reach significant levels in digital camera lenses and it needs to be corrected when computing 3D point position to even modest accuracy levels. Radial distortion also varies with focusing which is a further reason not to alter focus or zoom within a measurement network. Note that there are also two additional radial distortion parameters, K_4 and K_5, which are only appropriate for fisheye lenses with fields of view of up to approximately 130°. Fisheye lens calibration will be discussed in a following section.
- *Decentering Distortion.* The parameters P_1 and P_2 express the effect of the decentering of optical elements within the lens assembly. This error source is generally quite small and can typically be ignored in all but high accuracy applications. Thus, P_1 and P_2 can often be constrained to zero.
- *Linear Distortion.* The parameters B_1 and B_2 effectively model any relative error in the specified pixel size. They can generally be set 'fixed' to zero and ignored.

The calibration parameters resulting from the *Self-Calibration* can be viewed in the Camera Parameters dialog (Figure 1.3) by double-clicking on the camera icon.

vi) *Save the calibration data*. Upon completion of the calibration run the user is asked whether they wish to update the local camera database with the new calibration data. The normal response would be *Yes*. The project can then be saved.

vii) *Export a full camera calibration report* in PDF format by selecting **Export PDF Report** (Figure 1.3)

1.4 Fisheye Lens Calibration

CameraCalibrator supports the self-calibration of fisheye lenses up to a maximum field of view of around 130^0, under the implicit assumption that the lens adheres to a central perspective projection. The parameters K_4 and K_5 model the 9^{th} and 11^{th} order radial distortion effects, and they should only be employed for fisheye camera networks where the geometry of the multi-image configuration is strong and there are a lot of degrees of freedom, ie many camera stations and many points well distributed in three dimensions. Fisheye camera self-calibration is a very specialized application of the *CameraCalibrator* and it requires that considerable attention be paid to network design and quality of imagery. For all other lenses (non-fisheye), the parameters K_4 and K_5 should be 'fixed' to zero.

2. *CameraCalibrator* Setup

This chapter describes the procedure for initiating a *CameraCalibrator* project and importing the images into the project. The steps beyond that point will depend upon whether a targeted, targetless or manual image measurement approach is adopted, and these will be described in Chapters 3 and 4, and Appendix B, respectively.

2.1 Importing Project Images

Upon running the *CameraCalibrator* and commencing a new project, the operator first selects **File|New** from the main **File** pull-down menu. The first function to perform is the importing of selected images into the project. This is carried out as follows:

- If the project is new, *CameraCalibrator* will initiate the *Import Images* function. Alternatively, this can be selected with **File|Import Images** for an existing project.
- The user then selects from the Image Browser the folder holding the images. At this point the window shown in Figure 2.1 appears. An image in the *Select Image(s)* list is transferred into the project by first highlighting the image or images and then selecting the '>' button. The '>>' button

moves all images into the project. Similarly, the ‘<’ button moves highlighted images out of the project list, and ‘<<’ removes all images.

- A single image at a time can be selected, or multiple images can be highlighted by either dragging the mouse over the images (as shown in Figure 2.1; left-mouse click and drag) or holding down the CTRL key while selecting multiple images. Holding down the SHIFT key means all images between the two selected images will be highlighted.
- *CameraCalibrator* currently supports JPEG (*.jpg) and TIFF (*.tif) image formats.

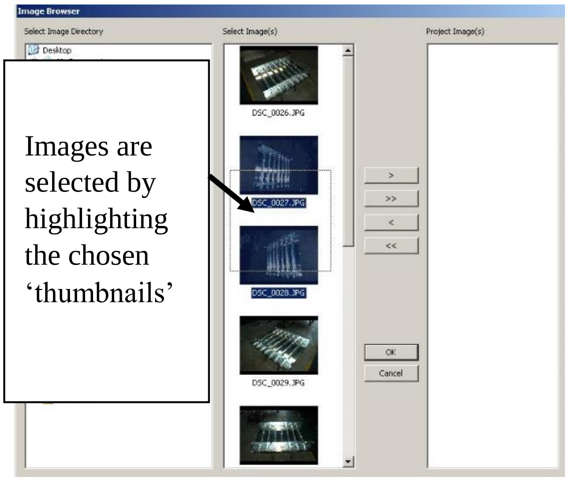

Figure 2.1

- If the image files do not contain information which identifies the camera(s), a warning message is displayed. This indicates that before the importing of images into the project, camera data must be entered either manually, or by selecting the appropriate camera from the *CameraCalibrator* camera database.
- Once the images are selected, they will appear as in Figure 2.2 and to enter them into the project the user selects **OK.**

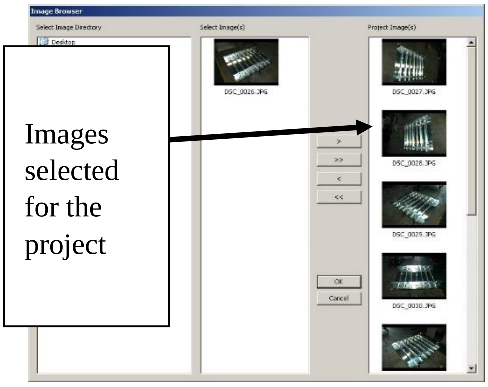

Figure 2.2

- Image thumbnails will then be shown to the left of the workspace, below the camera icon, as indicated in Figure 2.3. An image or images can be displayed in the workspace by double-clicking the image thumbnails.

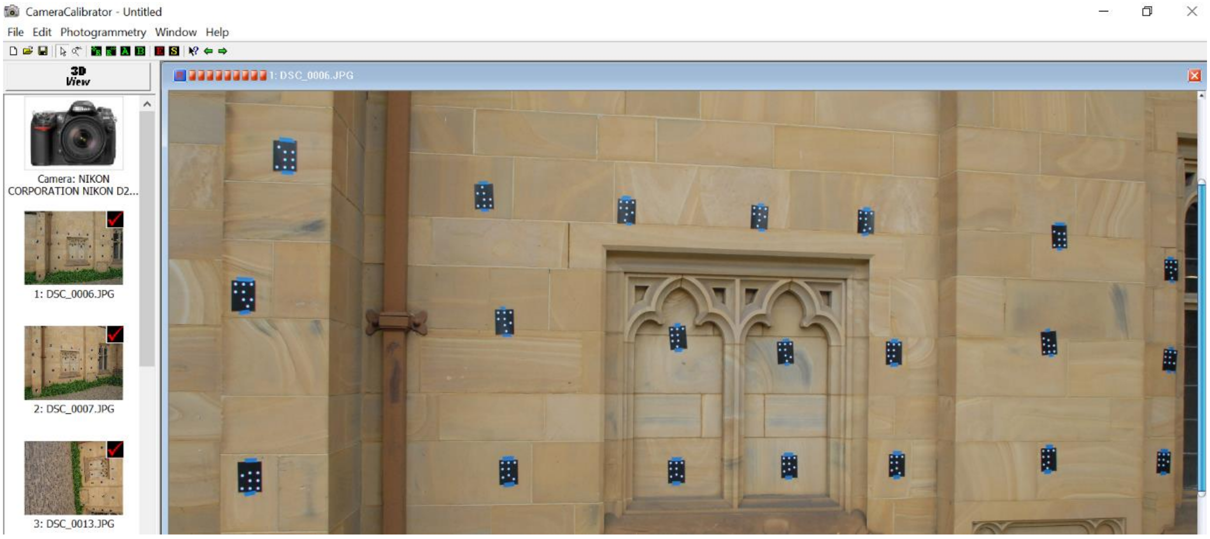

Figure 2.3

2.2 Project Settings

CameraCalibrator has options for basic project settings, which can be accessed via **Edit|Project Settings**. The *Project Settings* dialog is shown in Figure 2.4.

Note: The *Project Settings* dialog can only be accessed when either an existing project has been opened or after the images are imported for a new project.

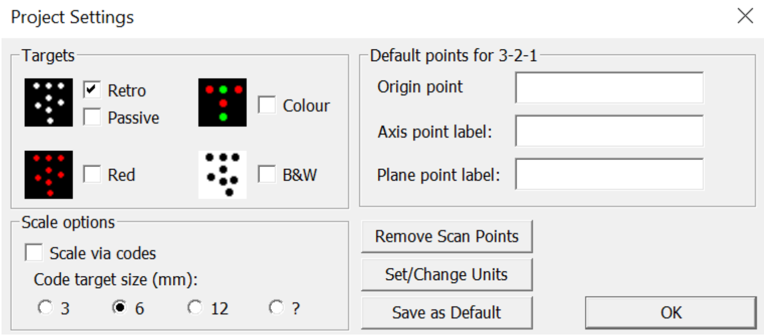

Figure 2.4

The user can select the type of targets being employed (see Quick Reference, Appendix A8), whether the network is to automatically scaled via the coded targets (different target types have different coded target size options) and whether the origin and orientation of the XYZ object point coordinate should be automatically assigned via three selected coded targets. There are also other options for removing those scanned 2D points not resulting in a 3D position and for setting project units (mm, m, inches or ft). Settings can also be left at their default values for automatic target-free orientation and calibration.

2.3 Selecting the Camera and Calibration Parameters

The *CameraCalibrator* requires a camera to be assigned to the images in the project. Generally, there is only one camera per project, but there may be more than one. An up-to-date list of cameras and their key metric design characteristics is provided with the *CameraCalibrator* software. The operator generally does not have to identify the camera or cameras used in the project; this is done automatically. Figure 2.5 shows a sample of the list of camera parameters that can be displayed in the main window by double-clicking on the camera icon. Calibration values can be changed interactively, though this is rarely required. At the outset, there should be an approximate value for the focal length/principal distance and all other parameters will likely have zero values.

Note: In situations the camera is not in the *Global Camera Database*, it is helpful – though not mandatory – to know the correct image pixel size. A default value of 0.005mm (5 microns) is otherwise assigned to pixel size and the other calibration parameters are scaled accordingly.

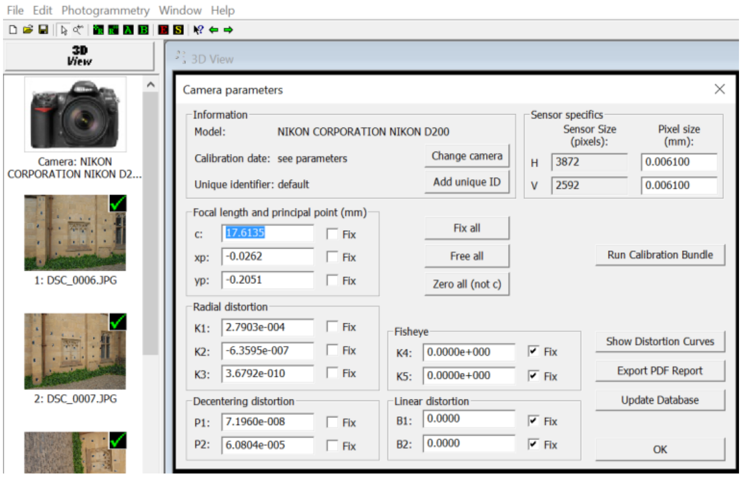

Figure 2.5

Note in Figure 2.5 that there are tick-boxes against all parameters. This allows the operator to choose which parameters are to be determined within the self-calibration. The most common parameter set is that indicated by the 'free' parameters in the figure, namely c, x_p, y_p, K_1, K_2, K_3, P_1 and P_2.

3. Automatic Calibration with Coded Targets

3.1 Procedure

The most common mode of operation of the *CameraCalibrator* is fully automatic 3D orientation, where a network of multiple images covering an array of object targets has been recorded with a specific geometric network of camera stations. An example is indicated in Figure 3.1.

The photogrammetric network of images and object points must have the following characteristics if the *CameraCalibrator* is to perform a fully automatic measurement:

i) Every targetted object point must appear in two or more images that provide good ray intersection geometry (see Figure 3.1). The targets, both codes and single points (natural or artificial), should be as highly contrasted against the background as possible, for example, as is achieved with artificial targets, either codes or 'dots', which are a light colour on a dark background.

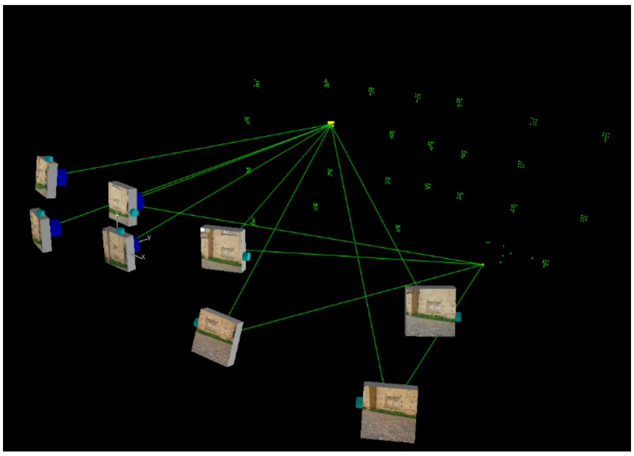

Figure 3.1

ii) There needs to be a sufficient number of coded targets, such that any five coded targets must appear on two or more images. It is not necessary that all codes and other targets are seen in all images, but codes provide the link between images so having subsets of five or more codes seen in two or more images is quite important.

iii) A strong convergent geometry is preferable, as indicated in Figure 3.1.

iv) In order to self-calibrate the camera(s), images must include both ‘portrait’ and ‘landscape’ orientations, ie the camera is rotated or ‘rolled’ 90^0 between images. It is not necessary to have exactly half with a 90^0 roll and half with no roll, but some of the images and preferably more than 30% must be orthogonally rolled.

3.2 Autoscanning Control Dialog and Parameters

After the loading of images into a new project, the first operation undertaken by the *CameraCalibrator* is the scanning of images to determine image point measurements of codes and other candidate targets. This *Autoscanning* process has a number of control parameters, which are set via the dialog shown in Figure 3.2. After opening an image, the dialog is selected using either the **Q** key within the Image View or the pull-down **Photogrammetry** menu selection **Image Scan Settings** (an Image View window must be active). The window within the dialog shows the portion of the image where the user left-clicked.

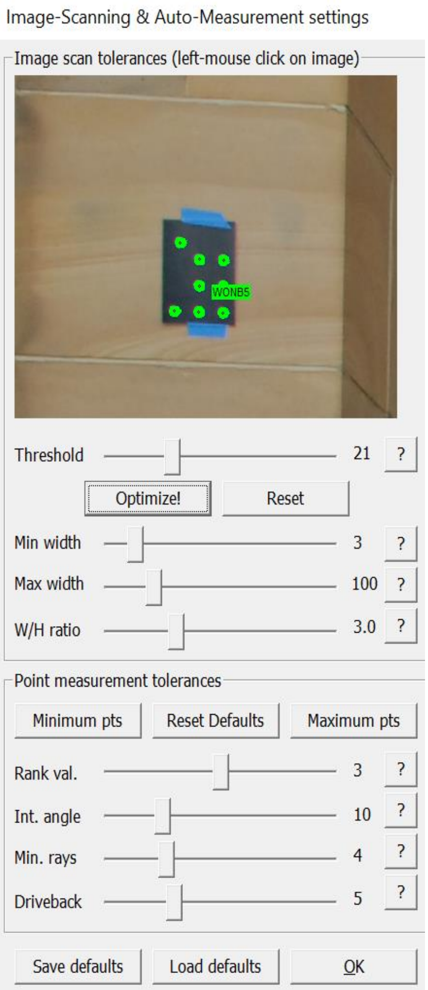

Figure 3.2

The *Autoscanning* parameters are:

1) The *Threshold* adjustment. This is used to include only those target blobs whose intensity is sufficiently greater than the surrounding background. The threshold grey value step between background and a target will be larger for a highly reflective target. The default setting is 20 grey values (the maximum for an 8-bit image is 255) and a suitable setting for retrotargets or bright white targets against a moderately dark background is anywhere from 15 to 50.

   Typically, sliding the threshold value to the right causes less blobs to be found, and more blobs (ie higher noise) are found when it is moved to the left.

2) The *Optimize!* option is selected if the user wishes to automatically set the threshold. This feature can also auto-detect the type of coded targets used, though it is best to do this via the *Project Settings*.

3) The *Maximum Width* (or diameter) simply specifies the largest size that a blob may have to be classed as a valid target. The default size is 50. The minimum size for a blob can also be set, at 3 or more pixels. The ideal sizes for target blobs are from 5 to 30 pixels.

4) The *width-to-height* or *‘W/H’ ratio* is a measure of blob shape. A value of 1 would normally indicate a near-circular blob. A value of more than 3 (the default) would indicate a very thin, elongated blob, which is likely then to be deemed invalid.

The image measurement tolerance parameters relate to the point measurement process and specifically to the determination of matching unlabelled image points. These parameters are *Rank Value*, *Intersection Angle*, *Minimum Number of Rays* and *Driveback Tolerance*. The function of these individual adjustments is as follows:

1) The *Rank Value* expresses the geometric tolerance applying to the ray-intersection geometry during the determination of corresponding non-coded targets in multiple images. There is a range of values from 1 to 5, with three being the default. Basically, the smaller the value, the larger the number of

possible point matches (for true targets and erroneous hot-spot targets). Thus, although more possible matches will be made, the automeasure process will slow down as this value is decreased towards 1. The admission of possibly many more erroneous point matches is not such a problem because these will usually be edited out in the final bundle adjustment. The recommended action here is to start with the default setting of 3. If valid target points are missed, then reduce the value to 2, and further reduce the value to 1 if true target points are still missed.

2) The *Intersection Angle Limit* is an important tolerance to apply since it is often the case in networks with many images that erroneous targets arising from such factors as sun reflections are seen in closely adjacent images, which often means that they have very small angles of intersection. This is depicted in Figure 3.3, where the point out in the right side of the figure has a maximum intersection angle of 5 degrees. The *Intersection Angle Limit* can be used to omit such points.

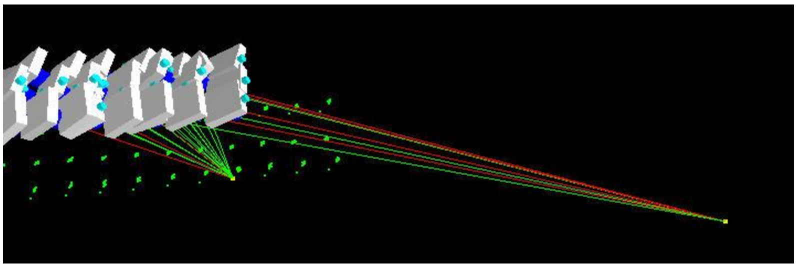

Figure 3.3

As a recommended action, the default setting of 15 (deg.) should be used. To include points with a smaller intersection angle, move the slider to the left; the minimum admissible angle is 3 degrees. To omit more points, move the slider to the right, but note that this adjustment cannot be used to reject points with an intersection angle of more than 40 degrees.

3) *Minimum Number of Rays.* This setting operates the same as the intersection angle condition, except that here points can be omitted in situations where they do not have a sufficient number of intersecting rays. The minimum number of rays for automeasure is 2 and points with more than 10 rays cannot be omitted via this adjustment. The default number, which is recommended for strong multi-image networks, is 4.

4) The *Driveback Tolerance* applies to the finding of image points corresponding to 3D object points. Based on the image orientation, the *CameraCalibrator* looks within a certain location for a candidate image point matching a given object point. The *Driveback Tolerance* expresses the size of the window in which the search will be conducted. The tolerance value goes from 1 pixel to 15 pixels, with the default being 5 pixels. In general, this tolerance rarely needs to be adjusted.

The *Camera Calibrator* has three pre-programmed *Image Scanning* and *Automeasure* Settings. These are *Minimum Pts*, which has the tightest tolerances and thus produces the fewest number of 3D points (those with the best quality); *Maximum Pts*, which has the most relaxed tolerances and thus produces the maximum number of matched 3D points; and *Default Points*, which represents the typical settings for retroreflective or high-contrast passive targets.

As to when to adjust the scanning settings, the ***Q****-selection* can be carried out before or after an *Autoscan* (eg **AutoCal** or **R++**), but if after, the autoscanning must be repeated with the new settings. The recommended time to adjust the target detection criteria is after the images are initially loaded into the project, before the first autoscan. Open a representative image, make the necessary adjustments and then run the autoscanning via **Autocal** or **R++**.

3.3 Two Options for Automatic Calibration via Targets

The *CameraCalibrator* offers two modes of automatic, target-based calibration. The first, **AutoCal**, utilises only coded targets, whereas the second, AutoMeasure or **R++**, uses both codes and any validated single dot/blob-type targets. Thus, the latter of these is generally more comprehensive and much more commonly used, but occasions my arise where it is only desired to utilise coded targets. Six to eight codes or more per image is a recommended minimum.

3.3.1 *R++*

The AutoMeasure process is initiated by selecting **R++** on the toolbar. The operator then selects **Begin**, as shown in Figure 3.4. Once the operation is complete, as indicated by the *Calibration … completed* dialog, the operator can review the brief results summary and select **Close.**

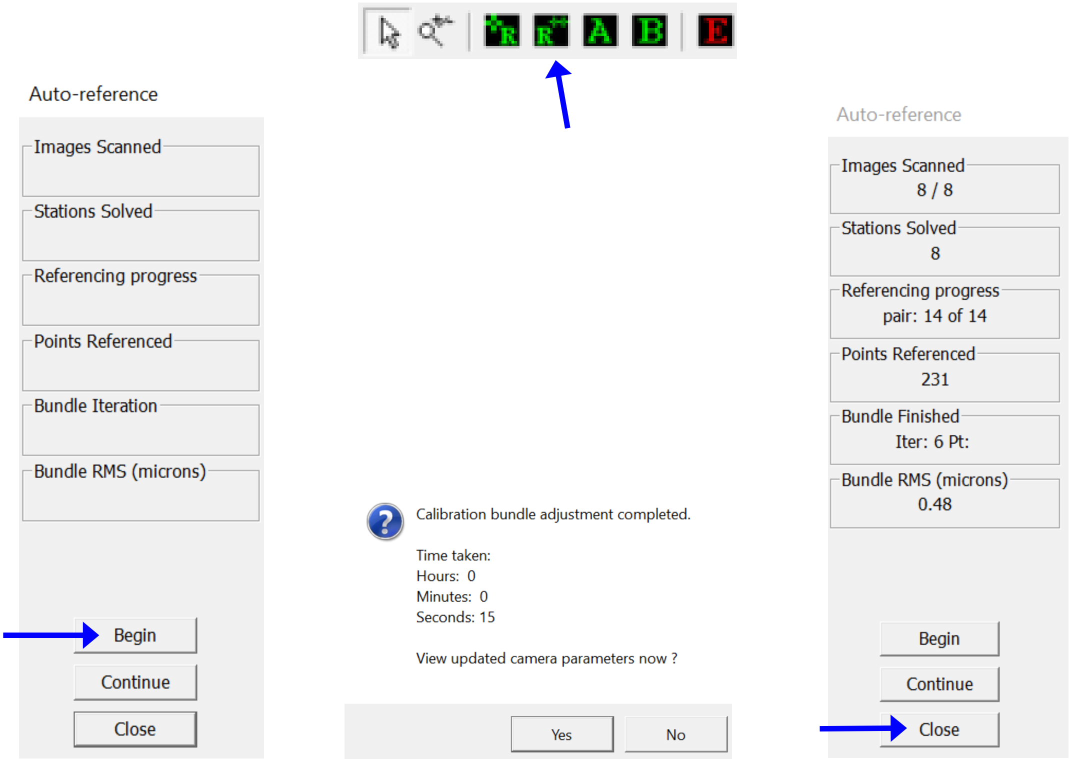

Figure 3.4

The basic 3D photogrammetric network is now measured, and displayed in the 3D View. This means that its shape is determined, but not its size or alignment with a chosen XYZ coordinate system. The operator can assess the results of the photogrammetric orientation and self-calibration by referring to the *Results Summary*, which is accessed from the **S** toolbar button.

Note: The recovery of camera calibration parameters is independent of absolute scale and coordinate alignment within the object space. However, since the *CameraCalibrator* provides an estimate of the photogrammetric triangulation accuracy attained in the XYZ object point coordinates, it may be useful to correctly scale the object point coordinate system and nominally align the axes via the 3-2-1 process (see Section 6 on page 21).

The project should be saved at this point, via the normal **File|Save** function. The project file, which can be reopened at any time is called ***projectname*.cal**.

3.3.2 **AutoCal**

Automatic orientation and calibration utilising only coded targets is initiated by choosing **AutoCal** from the pull-down **Photogrammetry** menu. This process concludes with a listing of the updated camera parameters, as shown in Figure 3.5.

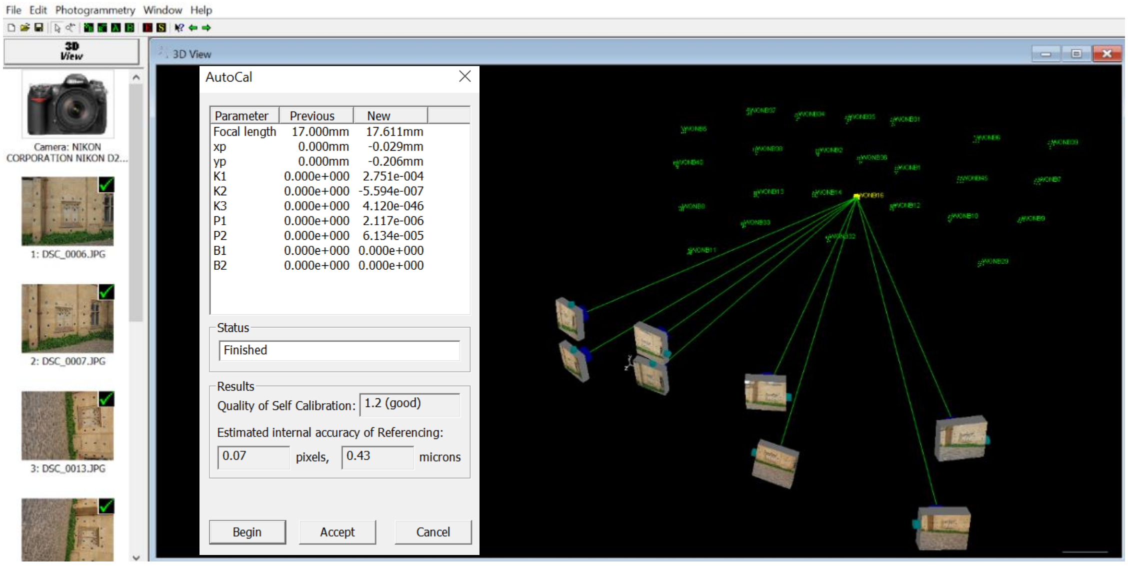

Figure 3.5

4. Targetless Automatic Calibration

A feature-based matching (FBM) approach has been implemented in the *CameraCalibrator* to facilitate fully automatic multi-image network orientation and camera calibration without the use of artificial targets. An example self-calibration network is shown in Figure 4.1. The first thing to note in the figure is the very large number of object points, generally totalling several hundred to several thousand. In this case the total is just under 9000.

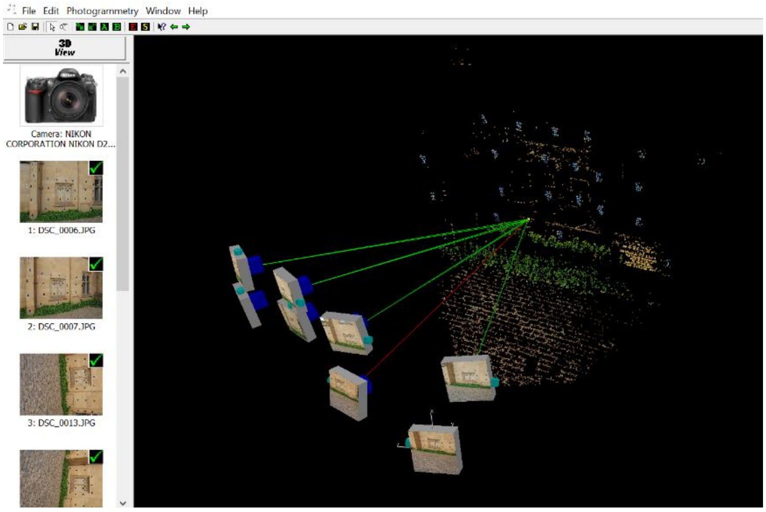

Figure 4.1

Whereas the automated target-free self-calibration functionality offers a lot of flexibility, success with the process is still contingent upon having both suitable network geometry and images with sufficient texture variation to enable unambiguous feature detector-based matching. The operational steps that follow the importing of images and the setting of desired calibration parameters (free or fixed) are as follows:

4.1 Initiating Automatic Targetless Orientation

Select the  toolbar button to initiate automatic orientation. This will bring up the *Target-Free Workflow Dialog* shown in Figure 4.2. It is not necessary to alter the Image Scan Settings (**Q**) for FBM.

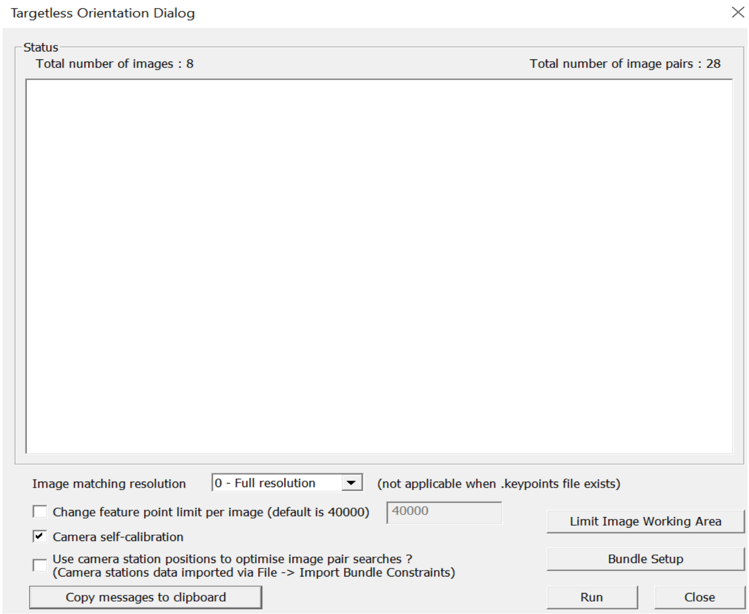

Figure 4.2

The *Targetless Orientation Dialog* indicates the number of images in the project and the maximum number of associated image-pair matchings that will be undertaken in the FBM process. A default limit of 40,000 has been set for the number of feature points per image. This can be changed by ticking the box and selecting a new limit; generally a suitable limit is between 10,000 and 80,000. The user can also apply initial bundle adjustment settings here by selecting **Bundle Setup** (Chapter 7). In general, none of the Bundle Setup parameters need to be changed at this stage, though for networks displaying very high image overlap it might be useful to change the minimum number of rays from 2 to 3 or 4.

There may be instances where it is desired to restrict the feature-based matching process to a limited working area of the image format. An example is when utilising fisheye lenses, where a practical field of view limit for the standard collinearity equation model adopted in photogrammetry is around 120^{o}. By selecting **Limit image working area**, the user can restrict the working area by radial distance or by field of view, as indicated in Figure 4.3, where the red circle indicates the boundary of the working area (the codes shown are not measured as part of the targetless orientation process).

When **Copy messages to clipboard** is selected at the end of an auto-orientation run, the screen messages are copied to the clipboard, so they can subsequently be pasted into a txt or Word document for later reference.

<u>Note on Camera station positions</u>**:** In instances where camera station positions are known, as in the case of GPS positions of images captured during a UAV/drone mission, this information can be utilised to optimise the initial pair-wise image matching. This aspect will be further discussed within the section dealing with Object Space Constraints (see Chapter 8). For regular strip/block structures of aerial images, the saving in time through utilisation of initial camera station positions **will be substantial**. To invoke pair-wise search optimisation, simply tick the box, noting that camera station positions must have first been imported, as discussed in Chapter 8.

Run is then selected.

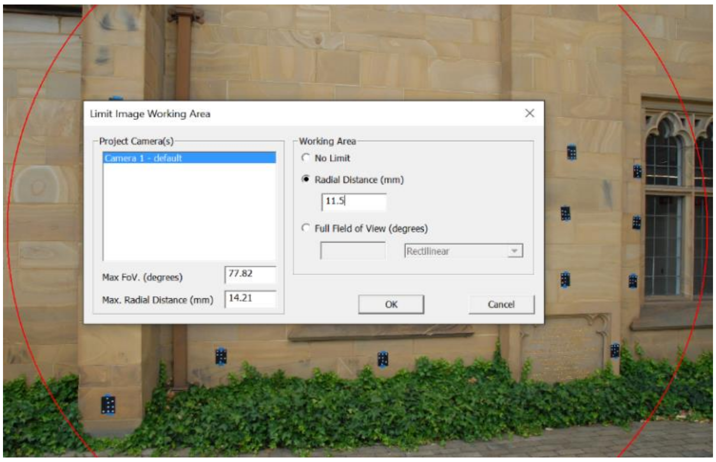

Figure 4.3

4.2 *Image Matching*

For each pair-wise matching, the number of feature points (limited to the assigned maximum) and the number of successful matches are shown, as in Figure 4.4. Feature points are only extracted once, so if a project is saved and then repeated, the FBM process is much faster.

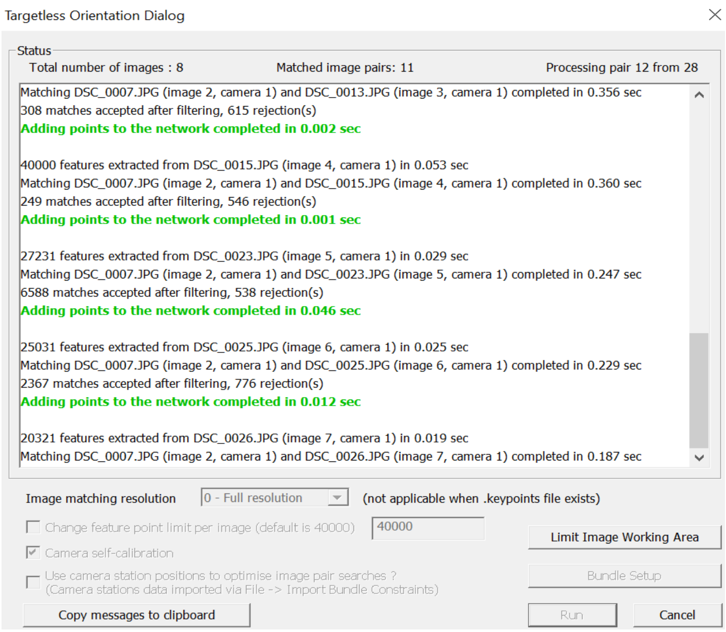

Figure 4.4

The maximum number of image-pairs to be matched is $\frac{n(n-1)}{2}$ where *n* is the number of images. Thus, this process can be time consuming when there are many images. On average, the user should expect that for a 16 megapixel camera it will take 10-15 seconds per image for the extraction of 40,000 feature points. The matching is carried at generally a few seconds or less per image pair. Thus, a network of ten x 16 megapixel images, limited to 40,000 feature points per image, is automatically oriented in some 3-4 minutes, depending on the processing speed of the computer used.

4.3 *Network Orientation and Camera Calibration*

Once the feature point matching phase is completed, two images are automatically chosen for an initial relative orientation (RO). Then, through a series of spatial resections, intersections and self-calibrating bundle adjustments, all images in the network are automatically oriented. The XYZ object point reference system will be either the coordinate system and scale of the initially provided camera station positions, or within an arbitrary datum at arbitrary scale when there is no initial camera station information. The automated calibration process is now complete, as is indicated in Figure 4.5.

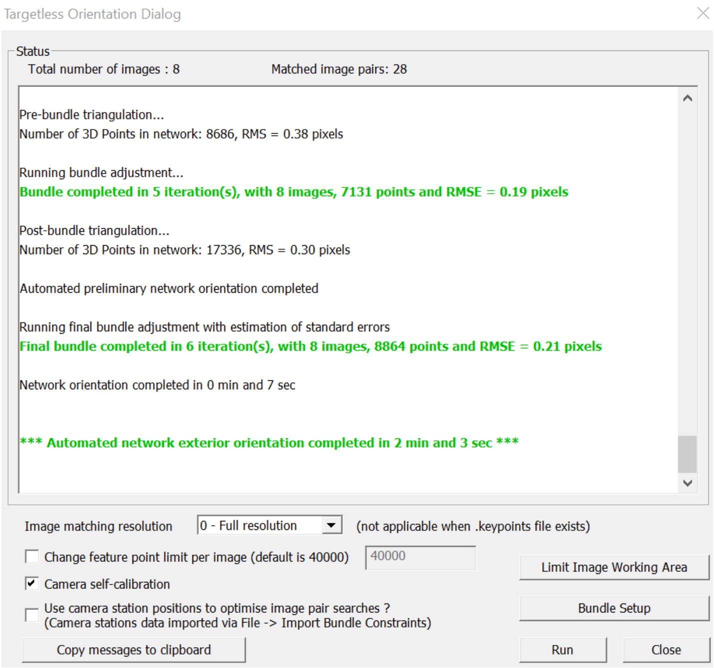

Figure 4.5

To display the oriented network, select the **3D** button, and a results summary is provided via the **S** toolbar button, an example being shown in Figure 4.6.

Note: It is possible that the network geometry and/or nature of the object or scene being imaged will not be conducive to automatic FBM-based orientation. In these instances the user can revert to manual orientation processes (starting with RO using FBM points) and even to standard manual referencing if required (see Appendix B).

4.4 The H and L Hot Keys

The **H** & **L** keys, which are described below, control different states of the Image View and 3D Graphics View when FBM points are present. FBM points are not necessarily selectable in all states, since it is helpful for certain functions to make them unselectable.

H – the H key toggles between three states in the **Image View**:

- i) Labels for manually referenced points are shown, but the labels for FBM points are not. This is the default and most useful setting because it allows the selection/ highlighting of manually measured points only. FBM points are not selectable.
- ii) Manually referenced point labels are still visible while labels for FBM points are shown on selection. All points are selectable.
- iii) No labels are shown and all points are selectable.

L – the L key toggles between four states in the **3D Graphics View**, these being

- i) Only labels for manually referenced points are shown.
- ii) Only labels for FBM points are shown.
- iii) Labels for all 3D points are shown
- iv) No 3D point labels are shown

H – the H key toggles between two states in the **3D Graphics View** (settings not appropriate for cases where all FBM points have been deleted):

i) All 3D points, FBM and referenced points are shown and thus all are selectable.

ii) Only manually referenced 3D points are shown; FBM points are not. This makes selection/highlighting of manually referenced points easier, since the scene is not cluttered with close-by FBM points.

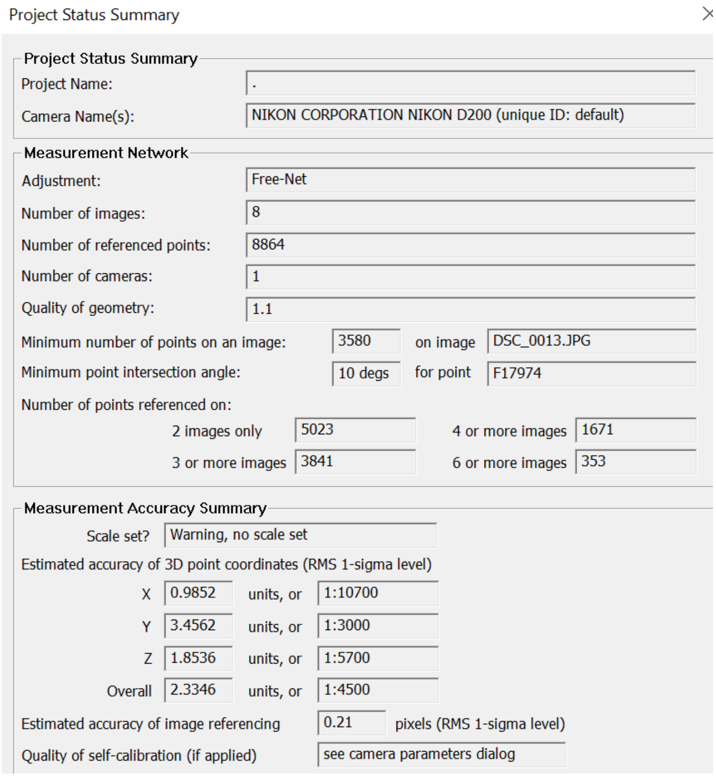

Figure 4.6

5. Output of Calibration Results

5.1 Computed Calibration Parameters

At the conclusion of the automatic calibration process, the dialog in Figure 5.1 will appear. This indicates that the camera calibration parameters have been successfully determined. Responding with a *Yes* will open the camera dialog, which lists the updated parameter values, as shown in Figure 5.2.

The operator can view the radial and decentring distortion profiles by selecting *Show Distortion Curves*, which results in the plots shown in Figure 5.3. Values can be read from these two plots simply by moving the cursor along the plots.

The red vertical line in Figure 5.3 indicates the maximum radial distance encountered within any image in the camera calibration. It is highly desirable that this line, which is also an indicator of how much of the image format was covered by measured image points, be as far to the right as possible (even more so than in this particular case), since the dashed sections of the distortion curves represent extrapolated values and thus should be minimized.

Note: Upon the saving of a *CameraCalibrator* project (***projectname*.cal**), an ASCII file (***projectname*.txt**) listing the camera name, image dimensions and pixel sizes, along with the computed calibration values will be generated and written within the Working Folder.

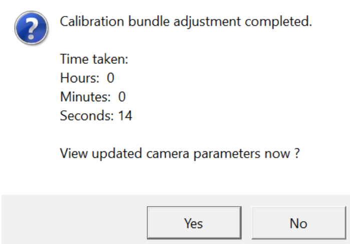

Figure 5.1

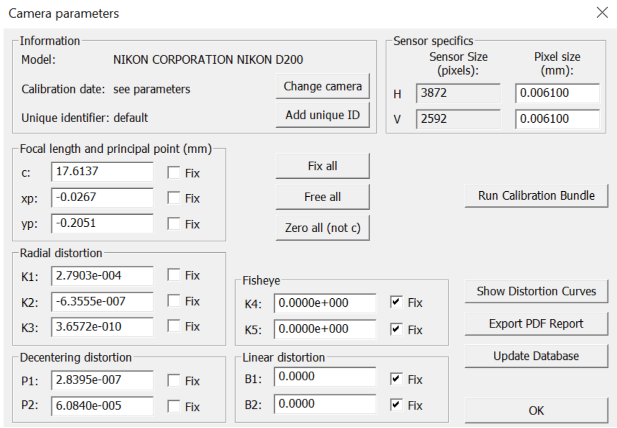

Figure 5.2

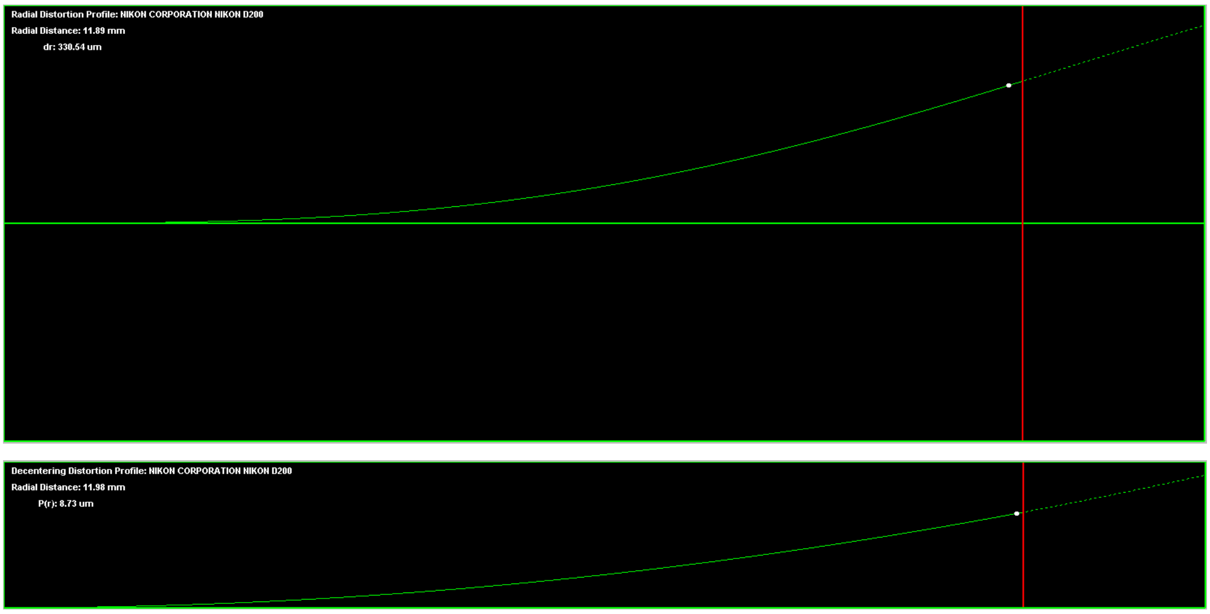

Figure 5.3

5.2 Calibration Report

Once the user is satisfied with the self-calibration outcome, including the choice of calibration parameters, a Camera calibration report can be generated by selecting **Export PDF Report** from the Camera Dialog, Figure 5.2. A sample calibration report is shown in Appendix D.

5.3 Further CameraCalibrator Processes

For the majority of users, the generation of the calibration report will signify completion of the camera calibration process. However, users may choose to conduct additional bundle adjustment runs with different calibration parameter sets, or they may choose to assign a correct absolute scale and a particular XYZ axes orientation to the network to gain an estimate of the quality of the photogrammetric triangulation expressed in terms of the resulting accuracy of object point determination.

6. Assigning Scale and Coordinate Axes

Shown in Figure 6.1 is an 8-station *CameraCalibrator* network comprising 8,800 points. From the moderately convergent camera station configuration and the dense 3D distribution of object points, coupled with the Accuracy of Image Referencing of 0.21 pixels, one can infer that this has been a reasonably strong camera calibration. It is also beneficial to gain an indication of the corresponding accuracy of the XYZ object point coordinates, as a measure of overall photogrammetric measurement accuracy.

An overall accuracy of 1:4500 of the principal diameter of the field is indicated, and this measure of relative accuracy at the RMS 1-sigma level, is invariant with network scale. However, in order to gain an estimate of accuracy in absolute units, it is necessary to apply a true scale to the network, since the accuracy figure of 2.33 only relates to an arbitrary scale. As well as assigning true scale, it is often useful to assign a particular alignment for the XYZ coordinate axes, for example so the accuracy in 'depth' (usually Z) can be distinguished from XY.

This section describes the assigning of absolute scale and the setting of a particular XYZ coordinate system. As mentioned, neither of these two steps is necessary for camera calibration, but some users may find the generation of absolute accuracy indicators for the photogrammetric triangulation to be useful.

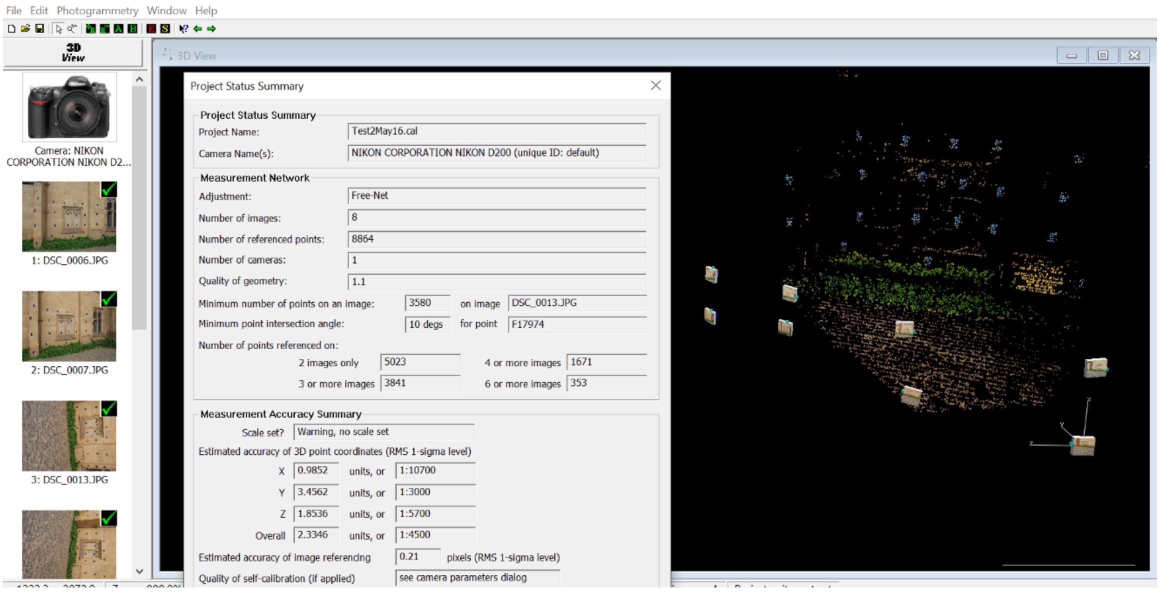

Figure 6.1

6.1 Scaling via Point-to-Point Distances

Photogrammetric orientation in the *CameraCalibrator* results in 3D object point coordinates at an arbitrary scale. In order to introduce true scale to the XYZ coordinates it is necessary to specify one or more point-to-point distances. The most straightforward scaling process proceeds as follows:

i) Right-click in an *Image View* or *3D View* window and select **Set Scale** from the menu.

ii) When **Set Scale** is first selected (and only on the first occasion), the user is asked to specify the project units. All coordinate and distance information, such as distances input to define scale are assumed to be in these units. Note that the choice of units does not relate to camera parameters; the chosen units refer only to object XYZ coordinates. The units are defined by choosing one of the four options in Figure 6.2.

iii) Next, the dialog box shown in Figure 6.3 appears. The operator selects the two points A and B (from the point number list) forming the known distance, and this distance is entered.

iv) To register the scaled distance, click **Apply**. Additional scaling distances can be added at any time and previously entered distances can be altered or deleted. Once all data is entered, press **Close**. Red lines will then be drawn in the 3D view between the pairs of nominated 'scale points'. The word 'pairs' is used because multiple scaled distances should be entered to enhance scaling accuracy and reliability. The final scale will be a weighted average of the nominated scaled distances.

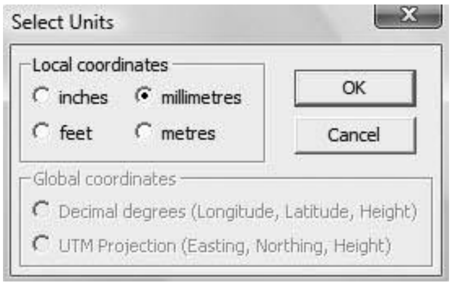

Figure 6.2

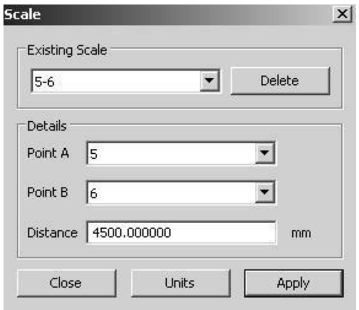

Figure 6.3

v) As an alternative to entering the numbers for points A and B, as in (iv), first select/highlight two points. Left-click in Select mode in either the 3D View or on an Image View and then CTRL/left-click for the second point. Then, select **Set Scale**. The two selected points will be shown as Points A and B and the operator need only enter the distance and press **Apply** followed by **Close**. This is shown in Figure 6.4, for the case of an image. Note that scale can also be set using image-to-image distances, if these distances are known.

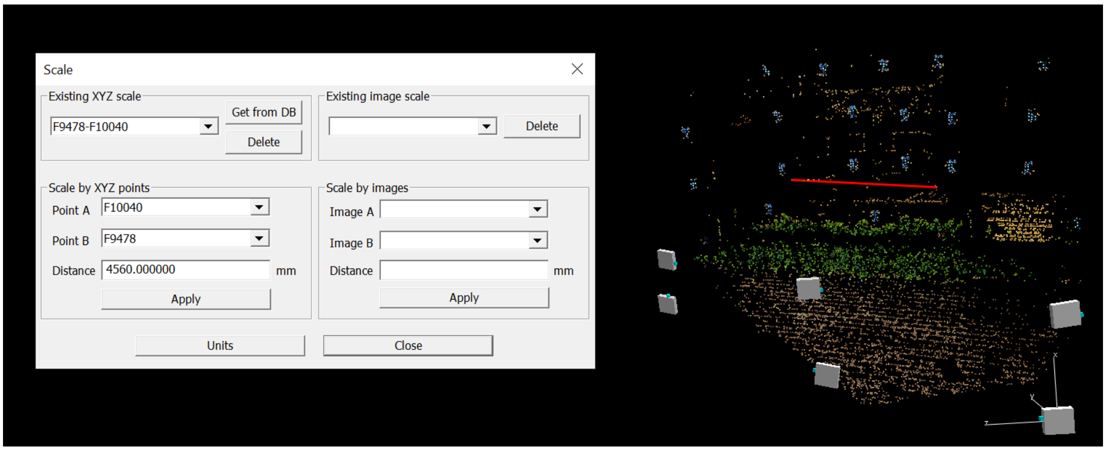

Figure 6.4

The project units can be changed at any time by using the pull-down **Edit|Project Settings** menu. Clicking on **Set/Change Units** brings up the standard Units dialog allowing the preferred units to be set. The **S** keyboard key can be used to toggle the red scale distances shown in the 3D View on and off. The **S** key also has the same function in the Image View.

<u>Cautionary Note</u>: In the event of a drastic rescaling of the network, the 3D view may look too small or large and may require rescaling via either the mouse wheel or the **Increase** / **Decrease view scale** options. The camera size may also need rescaling ('1'/'2' keys). Pressing the space-bar may also help.

6.2 Assigning a Coordinate System via the 3-2-1 Process

Photogrammetric orientation can be carried out within an arbitrary right-handed coordinate system, which is notated as X, Y, Z. In the *CameraCalibrator*, this coordinate system has its default origin at one of the first two referenced images, and its XY plane is aligned with the focal plane of the camera at that station, as shown in Figure 6.4. The assignment of a specific coordinate system origin and orientation is carried out via the so-called 3-2-1 process. First, a point is selected to define the origin (X,Y,Z values of zero). Next, a point through which the X axis will pass is defined, and finally a third point is selected to define the XY plane and therefore the direction of the Z coordinate axis (the Z axis is traditionally vertical).

In order to set the XYZ coordinate system, the following steps are followed:

i) Select (highlight) three points in either an Image View or the 3D View, in the order: origin point, X-axis point and XY-plane point, then right-click and select **3-2-1** from the menu.

ii) A dialog box then appears, as shown in Figure 6.5. It shows the three points and the default axial directions. The newly assigned axes are also shown, along with the old, in the 3D View.

 <u>Note</u>: The 3 points for the 3-2-1 process can also be selected (highlighted) in the Image View, but the 3D View is necessary for initially viewing the newly assigned axes. The 3-2-1 process can also be selected without highlighting points, in which case point numbers are entered via the drop-down options in the 3-2-1 dialog box.

iii) By interactively checking the selected +X, -X, ... , -Z boxes, the axes can be swapped, keeping the right-handed Cartesian nature of the coordinate system. Once the desired axial directions are established, as in Figure 6.5, press the **OK** button. At this point the 3D coordinates of all points and camera stations are transformed to the new system.

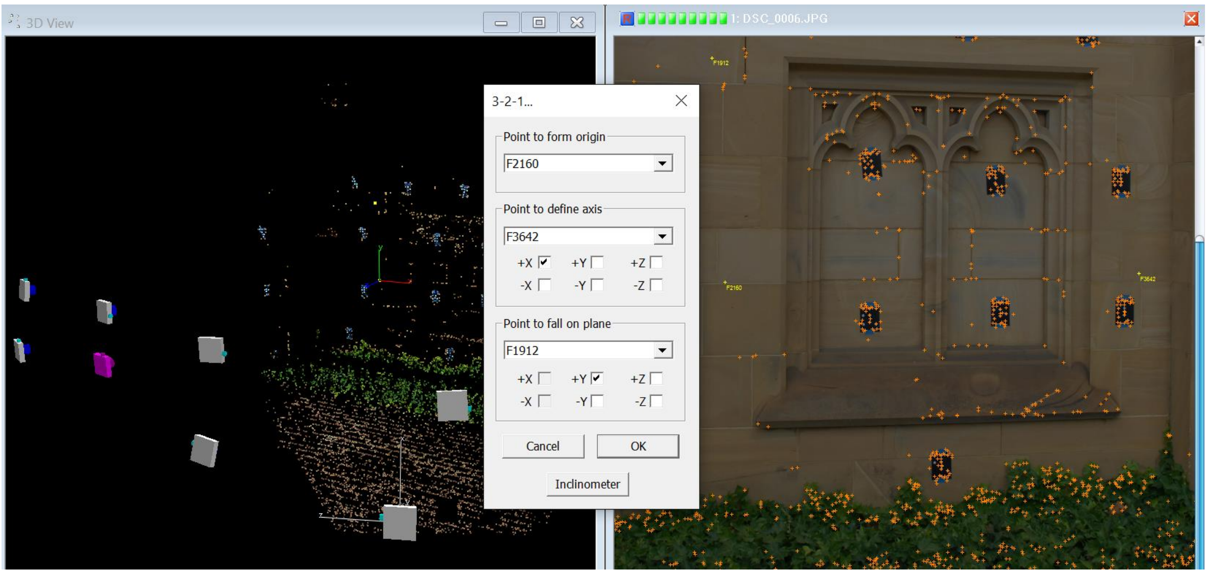

Figure 6.5

<u>Note</u>: It is possible to show each of the axial planes in the form of a grid within the 3D View. In order to do this, simply select the **X**, **Y** or **Z** key while the 3D View window is active. To change the size of the grid, hold down either the '-' or '+' key and then hit either X,Y or Z. To change the size of the grid cells, hold down '+' and hit '-' to make the grid cells smaller or vice versa to make them larger.

Once the true scale and desired XYZ coordinate orientation have been set, the mean standard error values for the X,Y and Z in the *Summary* will be correct indicators of absolute accuracy, as indicated in Figure 6.6. In the figure it can be seen that the 1:4500 relative accuracy has a corresponding absolute RMS (1-sigma) value of 1.9mm.

7. Self-Calibration with Different Parameter Sets

7.1 Additional Bundle Adjustment Runs – for advanced users

The main reason for conducting additional bundle adjustments would be to refine the camera calibration parameter set, since the full 'physical' model c, x_p, y_p, K_1, K_2, K_3, P_1 and P_2 (and, rarely, B_1 and B_2) may not always be either applicable or optimal for the camera/lens combination being calibrated. While the following section covers some typical parameter set selections, here we consider how to conduct additional self-calibration runs.

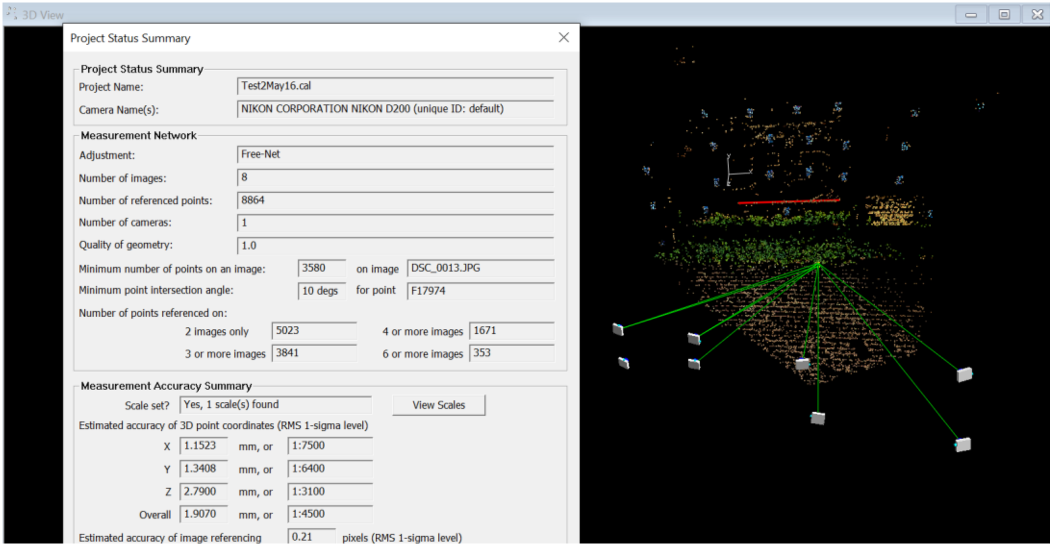

Figure 6.6

Firstly, the camera dialog is opened so that individual camera parameters can be either ‘fixed’ to pre-set values (generally zero for parameters other than focal length) or treated as ‘free’ unknowns to be computed. An example is shown in Figure 7.1, where the decentring distortion parameters have been suppressed from the self-calibration.

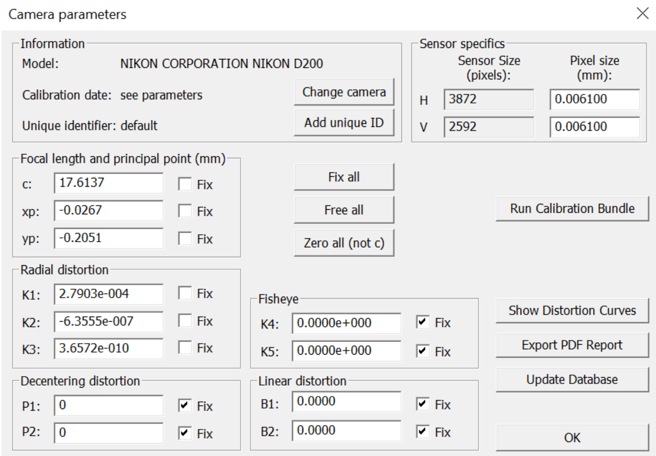

Figure 7.1

There are two options for running an additional bundle adjustment, the first being the **Run Calibration Bundle** button in the Camera Dialog above in Figure 7.1. The second, which provides further options for bundle adjustment control parameters, is to select the **B** toolbar button, which brings up the dialog shown in Figure 7.2, from which the set-up dialog shown in Figure 7.3 can be opened.

Following the self-calibration run, initiated from the **Run** button in Figure 7.2, the values in this dialog will be updated. The most important 1-number indicator of the quality of the calibration is the *RMS* (microns). This value is important in an absolute sense, with the expected magnitude being equivalent to 0.05 to 0.3 pixels (depending upon the camera and whether targets or a targetless approach have been adopted). In a relative sense, the magnitude of the *RMS* error will vary with the selection of calibration parameters. So, for example, if the suppression of the decentring parameters leads to either no change or a minimal change (eg <0.1 micron) in the *RMS* value, then it can be concluded for practical purposes that these parameters do not need to be modelled. The fact that the systematic error signal introduced by ignoring P_1 and P_2 is insignificant may either be because there is no measurable decentring distortion, or this error signal is projectively absorbed by the principal point coordinates x_p and y_p, and so P_1 and P_2 are unnecessary.

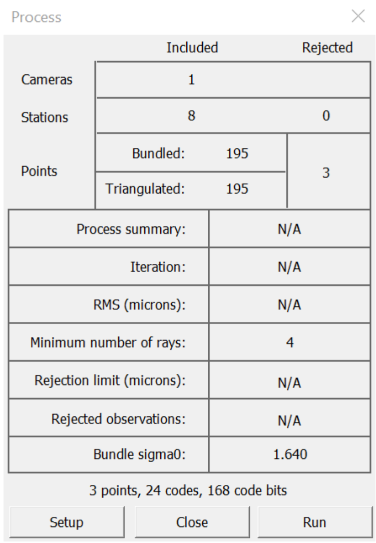

Figure 7.2

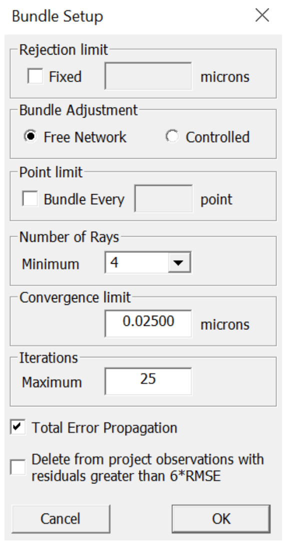

Figure 7.3

It is often desired to re-run the bundle adjustment with altered tolerance values, the two most common being the *Minimum number of rays* per point and the *Rejection limit*. The user can also decide whether to include every point in the bundle adjustment, or every n^{th} point instead. Running the bundle with, say, every 4^{th} point might be useful and much faster when there are thousands of points in the network.

There is also the option of either running a free-network bundle adjustment (which uses no control points or explicit exterior orientation constraints) or an adjustment with Control. The free-net is the default solution. Implementation of Control is discussed in Chapter 8.

A 'fixed' rejection limit (in micrometres) for image coordinate residuals can be set by ticking *Fixed* (see top of Figure 7.3) and then entering the desired value, which will typically be somewhat looser than the automatically set value. The Minimum Number of Rays often warrants a change in value. For strong, multi-ray networks a setting of >4 for the minimum number of rays can enhance the reliability and accuracy of the calibration.

7.2 Typical Calibration Parameter Sets

The following typical camera calibration sets apply to often adopted camera/lens combinations:

- c, x_p, y_p, K_1, K_2, K_3, P_1 and P_2 - short focal length cameras with wide- to medium-angle lenses (especially zoom lenses), and in cases where maximum metric accuracy is required.
- c, x_p, y_p, K_1, K_2, K_3 - medium to long focal length cameras; generally yields a calibration of sufficient metric quality for medium accuracy application.
- c, x_p, y_p, K_1 - medium to long focal length cameras; can yield a calibration of sufficient metric quality for medium accuracy application, especially when the distortion profile is near-cubic, which is typical of modest barrel distortion.
- c, K_1 – very long focal length lenses; recovery of other than radial distortion can be very problematic for self-calibration techniques because of projective absorption and there can be cases where the near linear nature of radial distortion in the central regions of a very long lens means that focal length may not be recoverable with high accuracy.

A notable benefit of the *CameraCalibrator* is that it is possible to try different combinations of parameters, since the tools to assess the integrity of each calibration run are readily at hand.

8. Object Space Constraints

This section describes the process of camera self-calibration in the presence of object space constraints. These most often take the form of GPS-recorded camera station positions in aerial imaging networks (eg with drones/UAVs) and/or ground control points (so-called GCPs). It is important to note that such object space constraint are NOT necessary for close-range camera self-calibrations that employ the recommended imaging and object point configurations (eg 'strong', convergent, multi-image networks). However, for 'normal' imaging geometry (stereo or near parallel optical axes between images) and object fields that have limited depth variation in relation to the camera-to-object distance, the recovery of interior orientation parameters requires the implementation of constraints, and generally both camera station positions and at least a small number of GCPs.

Since object space constraints generally only apply in stereo networks, the description here will be confined to the case applying camera station and object point constraints within networks of near-vertical images from UAVs/drones. However, the application of constraints can be extended to any desired calibration situation, noting of course that the presence of such constraints may not necessarily enhance the quality of the camera calibration process.

8.1 Importing Initial Camera Station Positions

To import camera station control, select **File/Import Bundle Constraints** and then the **Station Constraints** tab. There are three ways by which camera station positions are imported into a *CameraCalibrator* project. The first of these is an import of the geographic coordinates (Latitude, Longitude, Height) from the EXIF header of each of the images. The geographic coordinates are converted on input into either UTM coordinates (Easting, Northing, Height) in the appropriate UTM Zone or a Local Space Rectangular (LSR) coordinate system (X, Y, Z). An obvious prerequisite here is that the image EXIF header information includes camera position (GPS coordinates). This is becoming increasingly more commonplace, especially for networks of UAV images.

The second means is to import camera station positions from a file. Importing geographic (decimal degrees), UTM or Cartesian-style coordinates is supported, and geographic coordinates are converted to either of the latter two systems upon import.

As a final input option for camera station constraints, the coordinates can be entered manually in a spreadsheet-type interface, see Figure 8.2.

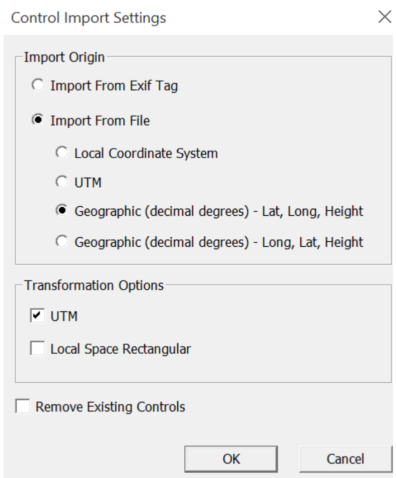

Figure 8.1

In order to enter camera station positions, select **File|Import Bundle Constraints** and the **Station Constraints** tab. The user then chooses a coordinate system via the dialog shown in Figure 8.1. Note that when **Import from File** is selected, the user needs to specify whether the coordinate system is Cartesian, UTM or geographic. If one of the latter two is selected, a further choice must be made as to whether to process the coordinates in either UTM (default) or LSR.

Note: This same procedure applies to the import of object space control points, eg ground control points, in a network of UAV images.

The ASCII file to be imported needs to have the following structure:

```
IMG_7897 619710.8816 5760687.1240 175.2992 0.10000 0.10000 0.10000
IMG_7898 619716.5438 5760738.7347 173.2267 0.10000 0.10000 0.10000
IMG_7899 619714.0350 5760790.8645 173.3083 0.10000 0.10000 0.10000
IMG_7900 619716.7431 5760842.3101 172.0480 0.10000 0.10000 0.10000
IMG_7901 619713.3989 5760894.2176 171.6542 0.10000 0.10000 0.10000
IMG_7902 619720.1593 5760945.3374 172.2147 0.10000 0.10000 0.10000
    ...
```

For geographic coordinates, Latitude and longitude values must be in decimal degrees.

Note that an image name must match that of the corresponding image, without the '.jpg' extension. The coordinate data is extracted from the camera stations file only for the images in the project (ie the camera stations file can contain more or indeed less stations than the number of images in the project). The last three columns, which may be left blank, are positional standard errors. This information is used within a 'controlled' bundle adjustment and the standard error values can be set either individually or as a block as per Figure 8.2.

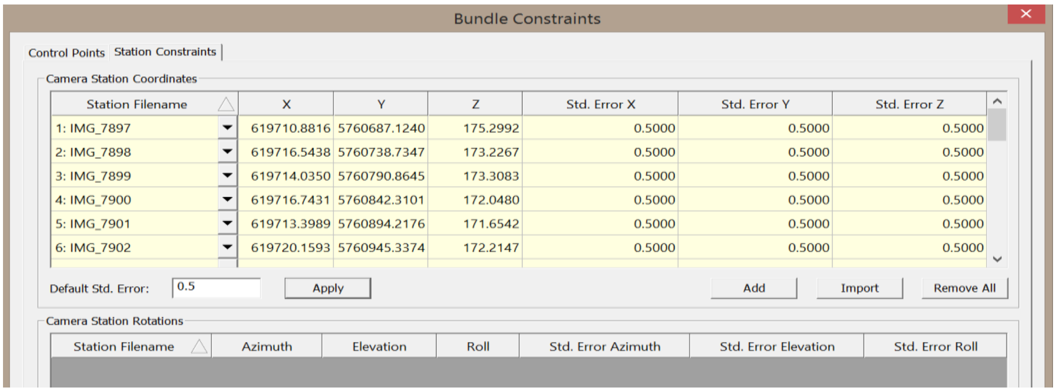

Figure 8.2

Once the camera station coordinates are imported, the network configuration can be displayed by selecting the 3D View button 

, the result being shown in Figure 8.3. Simply toggle with **SHIFT-C** to move from the display of Figure 8.3a, to the camera station numbers of Figure 8.3b, through to a blank display.

8.2 Importing Object Point Control (GCPs)

To import GCPs, select **File/Import Bundle Constraints** and then the **Control Points** tab. The procedure then follows that described above for input from file or via manual means of the camera station positions.

Note: It will be necessary to measure the 3D positions of the control points within the photogrammetric network. Targetted points may lend themselves to automatic measurement, but in general GCPs need to be measured manually, as described in Appendix B.

8.3 Controlled Bundle Adjustment

To run a bundle adjustment with camera station constraints and GCPs, simply select the **Controlled** radio button in the *Bundle Setup* dialog, followed by **Run**.

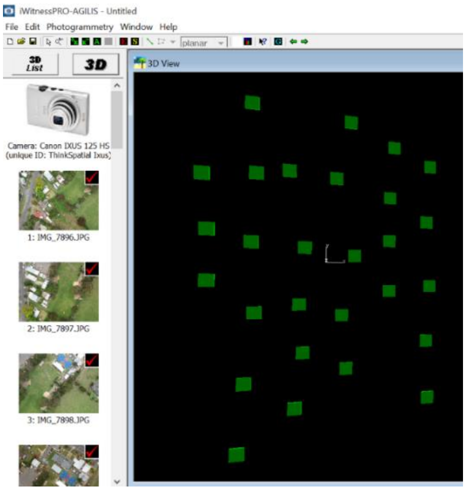
iWitnessPRO-AGILIS - Untitled
File Edit Photogrammetry Window Help
3D View
Camera: Canon IXUS 125 HS
(unique ID: ThinkSpatial Ixus)
1: IMG_7896.JPG
2: IMG_7897.JPG
3: IMG_7898.JPG

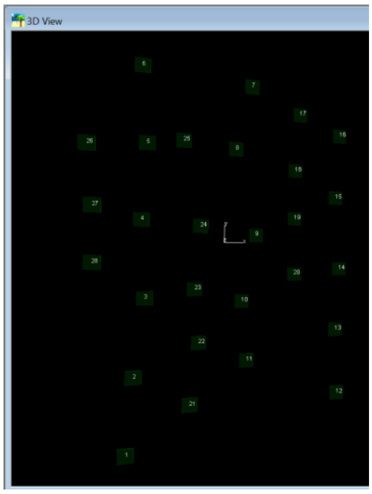
3D View

(a) (b)

Figure 8.3

Appendix A

Quick Reference to Menus, Toolbar, Cursors and Image and 3D Views

A1. Main Menus and Toolbar Buttons

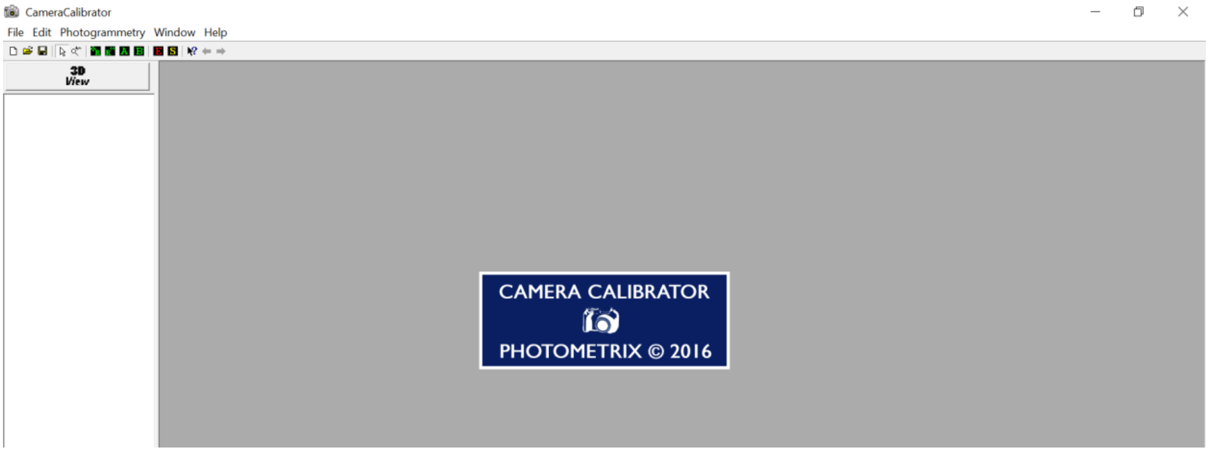

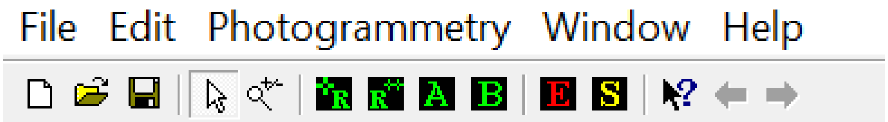

A1.1 Pull-down Menus

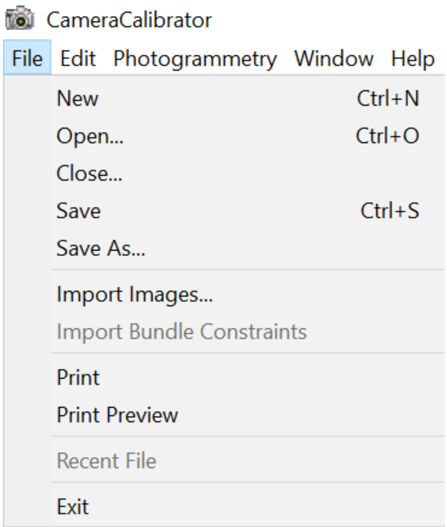

- New project
- Open project
- Close project
- Save project
- Save project under another name
- Import images into new or current project
- Import camera station or object point control
- Print project results summary
- Preview print of summary
- Most recent projects
- Exit software

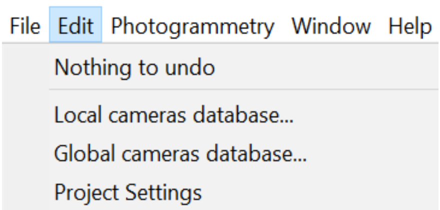

- Undo last point referencing
- List cameras in local database (cameras can be deleted)
- List cameras in global database (read-only)
- Project Settings dialog, described in Section A8

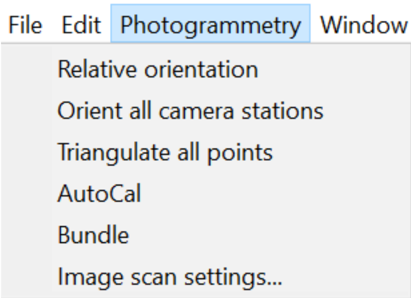

- Perform relative orientation between two images
- Orient all camera stations (image resection)
- Perform spatial intersection of object points
- Auto-calibration using only codes
- Run a bundle adjustment
- Adjust settings for image scanning

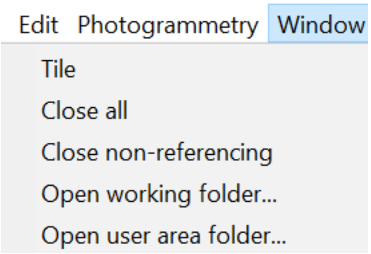

- Tile open windows in the workspace.
- Close all windows; image views and 3D view.
- Close all Image Views except for two being referenced.
- Opens working folder in Windows Explorer mode
- Opens folder of internal data files

A1.2 Main Toolbar Buttons

 - Commence a **new** *CameraCalibrator* project

 - **Open** an existing *CameraCalibrator* project

 - **Save** a project

 - **Select Cursor**: used for 'select' functions such as highlighting, marquee dragging, etc

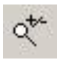 - **Navigate Cursor**: used for moving around within the image windows

 - Enter **Referencing Mode**: green 'pencil'  used to mark corresponding pairs of feature points in the two images being 'referenced'

 - Initiate fully automatic orientation & calibration using coded & uncoded target array

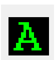 - Automatic orientation and calibration from targetless network (FBM)

 - Initiate bundle adjustment

 - Enter **Edit/Review Mode**: Selects the Review mode, where the quality of each feature point referencing can be verified and adjusted

 - List **Results Summary**: Provides a printable summary of the project status

 - Lists *CameraCalibrator* hot keys

 - Step forwards or backwards through the project images in full image view

A1.3 The 3D View Button

- Opens 3D Graphics View to display camera/object point network

A2. Right-Click Menus

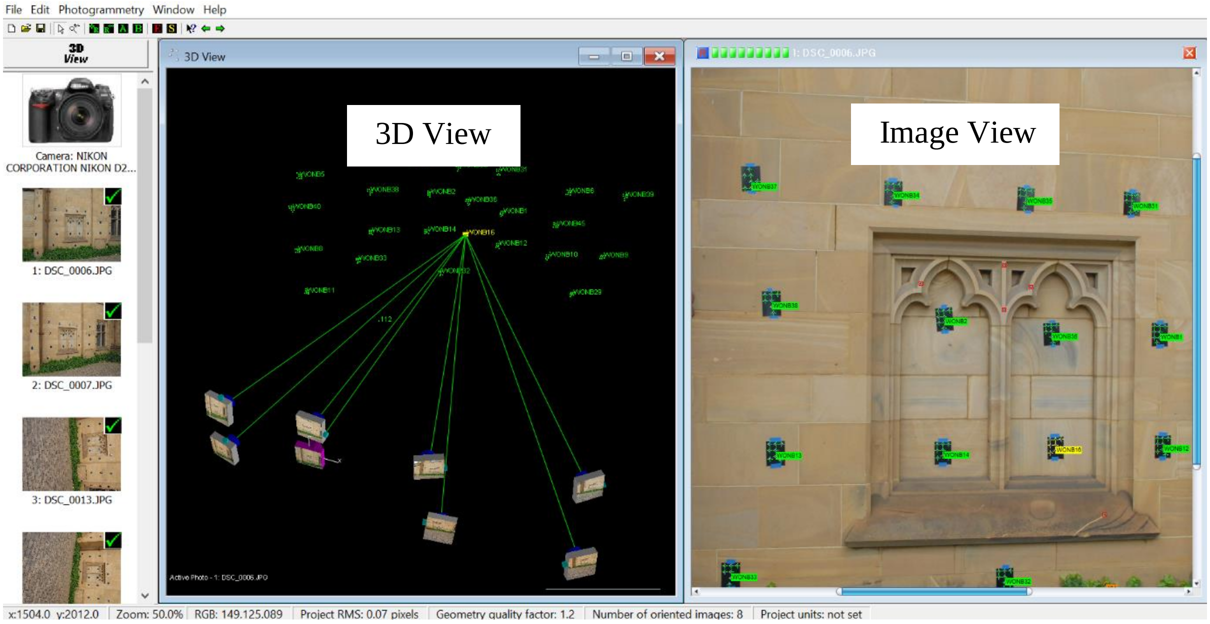

A2.1 Image View

Set scale...	- Set the scale for the 3D coordinate system via point to point distance(s)
3-2-1...	- Set a desired origin & orientation for the XYZ coordinate system
Unreference	- Unreference a feature point marking (label reverts to red)
Delete	- Delete a referenced point; all markings unreferenced
Close non-referencing images	- Close Image Views of non-referencing images
Photogrammetry >	- See below
Rotate >	- Rotate image view (left or right 90 deg. or 180 deg.)

Orient camera station	- Orient image in network
Driveback	- Image measurement completed by driveback

A2.2 3D View

Show / Hide Cameras	C	- Show/make transparent/hide icons for 3D camera stations (hotkey **C**)
Show / Hide Rays	R	- Show/hide plots of imaging rays (hotkey **R**)
Show / Hide Axes	A	- Show/hide XYZ axes (hotkey **A**)
Show / Hide Labels	L	- Show/hide point labels (hotkey **L**)
Increase View Scale >		- Vary plot scaling – see hotkey list
Decrease View Scale >		
Set scale...		- Set the scale for the 3D coordinate system
3-2-1...		- Set a desired origin & orientation for the XYZ coordinate system
Show / Hide FOV	V	- Show/hide camera field of view for highlighted station
Show / Hide Images		- Show/hide images mapped to camera station icons
Show / Hide Codes		- Show/hide coded targets

A3. Right-Click Menu from Image Thumbnail (some are supported by multi-select also)

Set Image Path...	- Set new path to images when project moved from original folder/directory
Change Camera...	- Change camera associated with this image
Remove Image	- Removes image from project
Set Un-oriented	- Set oriented image to un-oriented, hence not used in bundle triangulation

A4. **Hotkeys** (select via Help menu or toolbar button)

CameraCalibrator Hotkeys

Image - General | Image - Measurement | 3D View - General

View

+ / -- keys	Zoom image in /
ALT +	Drag for zoom-to-selection
CTRL + SHIFT +	Rotate image 90 degrees left
CTRL + SHIFT +	Rotate image 90 degrees right
SHIFT + F	Fit all images to view windows
F	Fit image to view window
SHIFT + T	Tile all
SPACE (hold down)	Image navigation mode
A	Show / Hide axes
H	Show / Hide point labels
V	Show / Hide residual vectors
W	Show / Hide reference guiding line
S	Show / Hide scale
P	Show 3D point information for selection
B	Blue (predicted) points on/off
N	Blue (predicted) polylines on/off
R	Referencing mode on/off
M	Single image point
G	Open 3D Graphics view
[/] keys	Decrease / Increase coordinate axes
Arrow keys (up / down)	Increase / Decrease image brightness
DEL	Delete selected objects

Image - General | Image - Measurement | 3D View - General

3D graphics view hotkey list

Arrow keys (left / right)	Rotate view left/right
Arrow keys (up / down)	Zoom view in / out
ESCAPE	Unselect all points/lines
SPACE	Centre view on selection
A	Show / Hide axes
C	Show / Hide cameras
L	Show / Hide point labels
N	Show / Hide Control Points
P	Show 3D information for selected point
R	Show / Hide camera to point rays
S	Show / Hide scale
V	Show / Hide camera field of view
[/]	Decrease / Increase coordinate axes
, / .	Decrease / Increase point size
9 / 0	Decrease / Increase text size
1 / 2	Decrease / Increase camera size
Shift + Right Mouse	Pans the view
F	Fit to view window

Image - General | Image - Measurement | 3D View - General

Measurement

Arrow keys (left / right)	Move to next / previous point in guided referencing mode
1 / 2 keys	Move to next / previous point in guided referencing mode
Q	Edit color target scanning parameters
C (hold down)	Brings up the zoom window for dark color target centroiding
X (hold down)	Brings up the zoom window for light color target centroiding
Z/Y (hold down)	Brings up the zoom window for referencing/marking

A5. Panning in the Image View

To ‘pan’ or roam within an enlarged (zoomed-in) image, hold down the wheel of the mouse (or central button) and move the mouse. The slider bars can also be used for this function, as can dragging the mouse with left-button held down in Navigate mode. Also, the space-bar can be held down in any mode and the mouse dragged to pan.

A6. **3D View Functions** (the cursor must be in the 3D View area)

- To zoom in or out within the 3D view (make sure the cursor is over the 3D window), the wheel/roller ball on the mouse can be used.
- To rotate the 3D View about the axis coming out of the screen, hold down the **CRTL** key and the right mouse button and move the mouse.

- To pan in the 3D View, use the **SHIFT** key and the right mouse button and move the mouse.
- To bring a point to the centre of the 3D View display window, first highlight the point (left mouse button and drag) and then hit the **Spacebar**.
- Moving the mouse with the right button pressed, and the point mentioned in the previous dot-point still highlighted, rotates the display about the highlighted point. Also, the rotations described in the second dot-point will be about the highlighted point.
- To centre the network in the 3D View, use the **F** key; **Alt+Right Click** does a fit-to-screen.
- An Image View can be opened from the 3D View by double clicking the corresponding, highlighted camera icon.
- The **0** and **9** keys can be used to increase and decrease the point label size.

A7. Zooming within the Image View

The four zoom options available in the active image window are:

- The **Z** (or **Y**) key. When the pencil cursor is over the point of interest, hold down the Z key to generate a zoom window. The cursor can then be accurately placed on the feature point of interest. The window will remain open as long as the Z (or Y) key is held down.
- The **ALT** zoom. Press the ALT key and use the left mouse button to draw a marquee box that will enclose the enlarged portion of the image. The image will be zoomed upon release of the mouse button.
- The '**+**' keyboard key to enlarge the image and the '**-**' key to reduce it.
- The mouse scroll wheel (if present) can be used to zoom in and out. This is the easiest method.

A8. Project Settings Dialog

Initial settings for the *CameraCalibrator* can be made using the *Project Settings Dialog*. Further details are provided in Section 2.2.

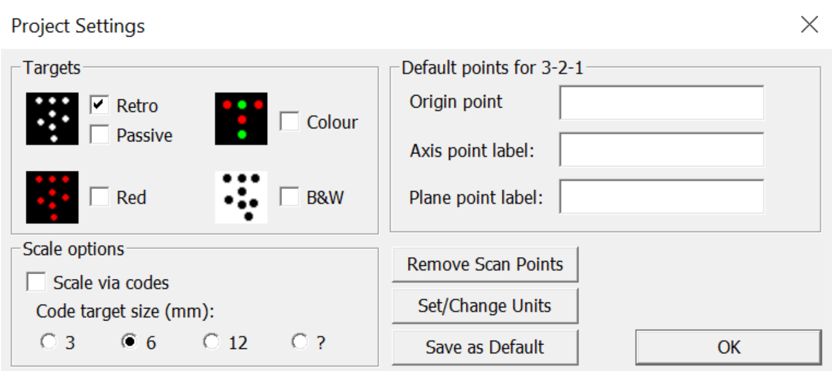

Targets Options
- *Retro* for white-on-black retroreflective codes
- *Passive* for white-on-black printable (non-retroreflective) codes
- *Red* for red retro-reflective coded targets
- *Colour* for red/green calibration cards
- *Black & White* for black-on-white printable coded targets

Appendix B

Manual Image Measurement by Point Referencing

Although the *CameraCalibrator* is primarily designed for automatic calibration, which includes automatic image measurement and point correspondence determination for both targeted and untargeted scenes, the image measurement process, the so-called marking and referencing of feature points, can also be carried out manually. The process is slower and typically a little less accurate than automatic measurement, but in some circumstances it may be necessary.

Moreover, the manual measurement process can be used to add additional points to self-calibration networks that have been oriented automatically.

This chapter briefly describes the manual image measurement and sequential, interactive network orientation associated with manual point referencing. The procedure starts once the images are loaded into the project, which may be a new project or an existing project.

B1 Referencing the First Two Images for Relative Orientation

The determination of 3D point position requires the spatial intersection of two or more imaging rays, and hence two or more images are needed. The first task is thus to identify and measure corresponding feature points in an initial image pair. This is termed *referencing*.

Shown in Figure B1 are the two images to be initially referenced. The procedure is as follows:

- Open the two images in the workspace by double-clicking on the image thumbnails.
- The brightness of the image display can be changed via the up/down keyboard arrows.
- Select two images that have a good 'base'. This means that the camera stations should be separated by a distance that is 15% or more of the distance from the camera positions to the object of interest. For example, if the imaging distance is 20m, a desirable camera station separation is 3m or more.
- To enter referencing mode, select either the green $^{+}$**R** button on the toolbar or the **R** key on the keyboard. Alternatively, left-click on the two red **R** buttons at the top left of each image view (they will then turn green). The cursor will change to a green 'pencil'. Referencing simply entails the successive *marking* of the same point, sequentially in the two images, with the order being either right to left or left to right. In FigureB1, after Point 1 was marked in the left image, a guiding line connects that point to the referencing cursor for ease of marking in the right image. (Ignore the fact that the images have targets; these are not being used here.)

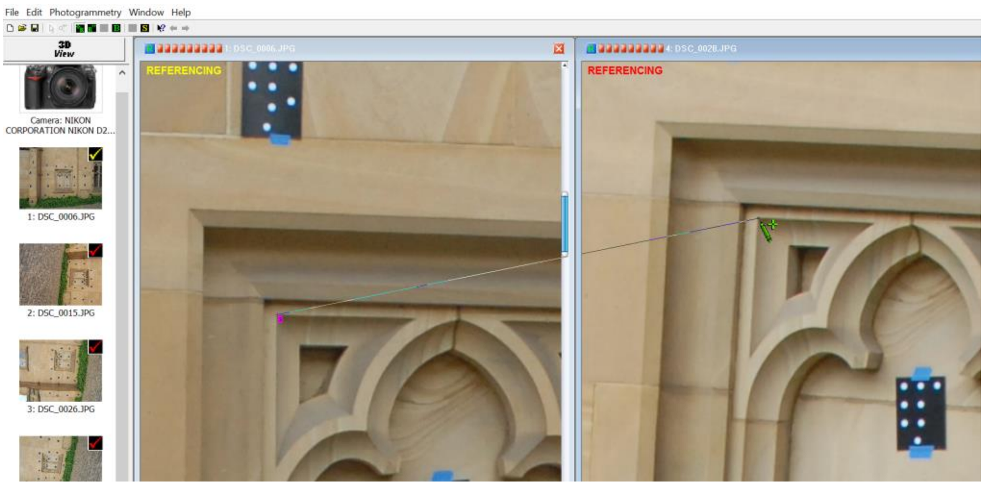

Figure B1

- It is generally desirable to enlarge the image of the point to be marked. This is achieved through one of the zoom options just described, typically by scrolling the mouse wheel in and out.

Panning and Zooming to Aid Referencing

Recall that you can **pan** (roam) within the image by holding down the scroll wheel of the mouse (or central button) and moving the mouse, or dragging the mouse with left-button held down in Navigate mode (hold down space bar).

The **zoom** options available in Referencing Mode within the active image window are:

i) The **Z** (or **Y**) key. When the pencil cursor is over the point of interest, hold down the Z or Y key to generate a zoom window, as shown here. The window will remain as long as the key is pressed.

ii) The **ALT** zoom. Press the ALT key and use the mouse to draw a marquee box that will enclose the enlarged portion of the image.

iii) The '**+**' keyboard key to enlarge the image; the '**-**' key to reduce it.

iv) The mouse wheel (if present) can be used to zoom.

v) As per the discussion of image centroiding, the Z-window will also automatically appear when either the **X** or **C** centroiding selections are made.

Note regarding automated marking. The automatic and semi-automatic marking and centroiding of distinct, targeted points will be described in Sections B3 and B4.

- With the pencil cursor positioned on the point of interest, click the left mouse button and a purple cross will mark the desired point. A sequence number, which has no importance at this stage, will also appear. Now, go to the same point in the second image and mark the corresponding position. You will note that the two label markers now turn either orange, white or green and a number is assigned to the point. The point is now referenced.
- In order to 'orient' the images, six or more corresponding points need to be referenced, and depending upon the quality of the point marking and the imaging geometry, *CameraCalibrator* may require more referenced points (usually 7-10) before it decides that an acceptable and sufficiently accurate Relative Orientation (RO) has been achieved. Referenced points should be well distributed in two dimensions (they cannot be co-linear) and preferably have a three dimensional separation.
- The thumbnails of the two images being referenced will have yellow tick marks, which indicates that they are being referenced, but have not yet been oriented. Upon successful orientation the tick marks will turn green.

Undoing the Last Point Referencing

To undo the most recent point referencing operation, simply select **CTRL+Z** or choose the pull-down menu selection **Edit|Undo last referenced point: #** (where # is the point number).

- Keeping in mind the requirement to have a well distributed set of referenced points, proceed with the referencing until *CameraCalibrator* automatically computes the relative orientation (RO). This is indicated by the referenced points turning green or white, and green squares appearing in the orientation progress indicators, as shown in Figure B2. (the points are yellow in the figure because they have been 'selected')
- With the RO established, the user can click on the **3D** button to show the graphics view of the oriented cameras and the 3D object point positions, as shown in Figure B2. At this stage the scale of the network is arbitrary, as is the reference coordinate system (XYZ).
- The **H** key can be used to toggle point labels of marked (red) and referenced points (green) on and off in the image display.

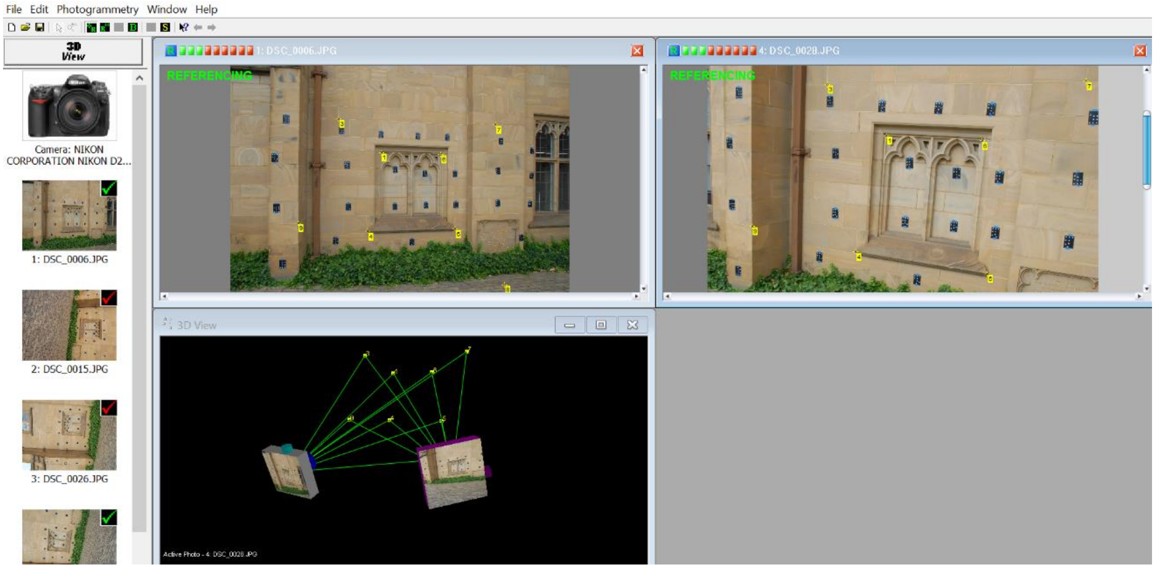

Figure B2

Deleting a Point During Referencing of the First Two Images

Points may be deleted during referencing as follows:

i) If the point selected in the first image (coloured purple) is deemed to be in error, hitting the Delete key will cancel the marking.

ii) For a referenced pair of points (in reference mode) hold down the CTRL key and the cursor will change to the Select cursor (arrow). A marquee box can be drawn over the points(s) to be either unreferenced or deleted. The point label in all images containing that point will turn yellow. The DELETE key will then unreference the points and they will revert to red 'marked' status only (they are no longer 3D points). Clicking on these points in reference mode (after the CTRL key is released) will turn them purple and the DELETE key will remove them completely.

iii) When in reference mode, a right-click on the mouse will present the user with the option to either unreference or delete highlighted (yellow) points. For points already referenced in other image pairs, unreferencing will occur only in the current image.

iv) If the 3D view is open, the left-mouse button can be pressed and a box drawn over points to highlight them. They will again turn yellow, in both the 3D view and the images. Hitting the DELETE key will unreference these points and remove them from the 3D points list.

B2 Referencing Points after Two-Image Relative Orientation

Once the first two images are oriented, the next step is to continue referencing points of interest. Although it is always possible to leave and return to the referencing mode at any time, it is highly recommended that further referenced points be added before moving on to the third image. Two reasons for this are:

i) Additional points make the orientation 'stronger' and can therefore enhance accuracy.

ii) Referenced points in the first two images serve as orientation points for subsequent images, so it is good to have a reasonable number to choose from.

With these aspects in mind, we now return to the referencing sequence.

- As shown in Figure B3, immediately a point to be referenced is marked in one image, a blue line is projected in the second image. The point to be marked in the second image should lie along this line. The line is called the epipolar line and it serves as a useful guide for point location. Note that the line may not always pass exactly through the target point, for its position is a function of both the quality of the initial relative orientation and the calibration of the camera. The referencing of the point along the epipolar line proceeds as described previously. The projected epipolar line spans only that part of the image where the resulting feature point will produce an intersection angle of greater than 3 degrees.

- Once all desired point pairs are referenced within the first pair of images, the operator can leave the referencing mode by either clicking the **R** button on the Toolbar, or by pressing the **R** key. Figure B4 shows the fully referenced first two images, which contain 20 marked and measured points.
- To **delete 3D referenced points** from the project using the 3D Points List simply highlight the point labels in the List, either as a single or multiple select (via the CTRL or SHIFT key), and choose **DELETE**. The deleted 3D points will revert to unreferenced point markings (red colour).

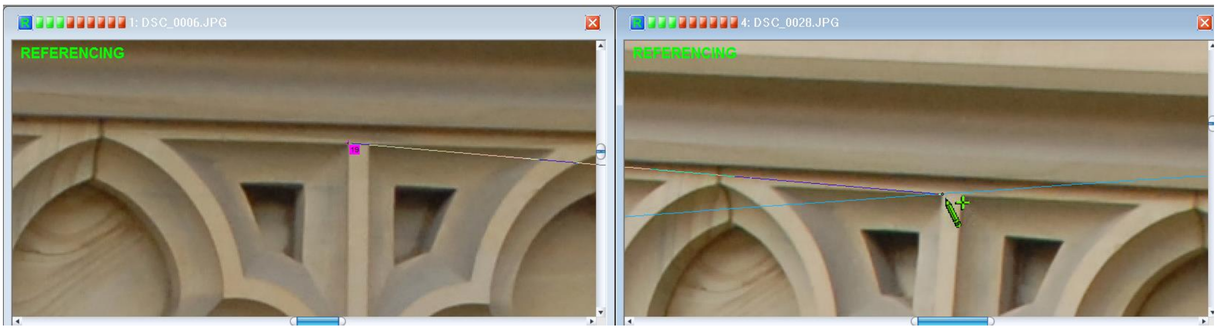

Figure B3

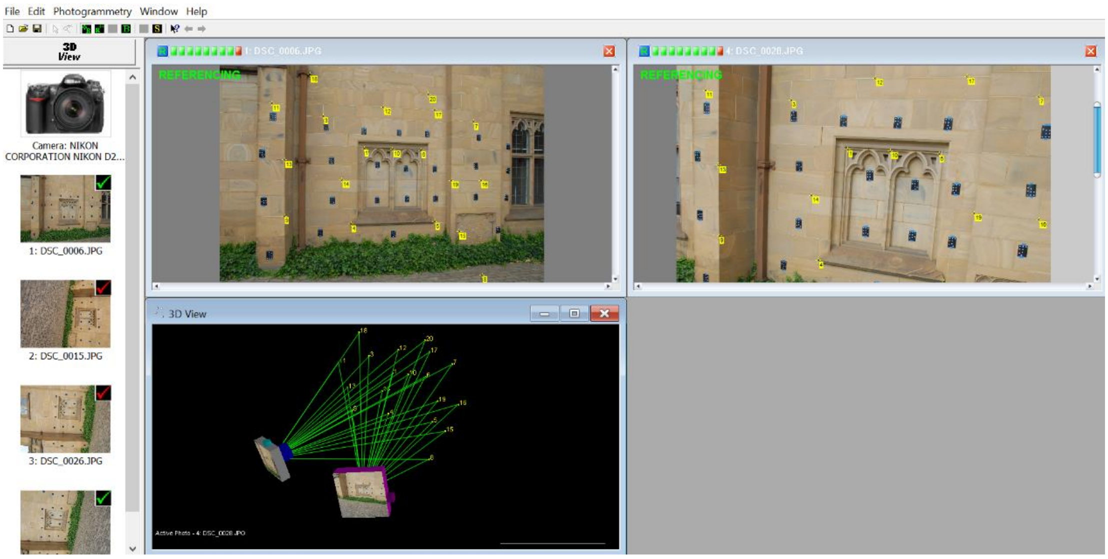

Figure B4

B3 Target Centroiding in Referencing Mode

Target Centroiding. The optimum targets for precise marking and referencing are likely to be high contrast dots. Where such targets are utilised, *CameraCalibrator* can provide a precise centre-of-target (centroid) determination using a fine-measure auto-assist function. In order to perform an automatic precise marking in either Referencing or Marking mode, the Centroiding function can be employed by selecting the **X** key for white or light blobs on a dark background, or the **C** key for dark blobs on a light background. (Recall: select Referencing mode first, followed by the centroiding function). In this case the centre of target G in Figure B5 is required:

Figure B5

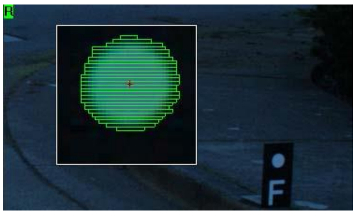

Figure B6

The centroiding operation proceeds as follows:

i) Select Referencing or Marking mode.

ii) The cursor is placed over the target and the **X** key (or **C** key for dark blobs) is selected. The zoom window appears as shown in Figure B6.

iii) Moving the mouse within the window adjusts the intensity profile of the target, with the red cross always indicating the centroid (centrepoint). To see the image window clearly, simply move the mouse to the top right-hand corner of the zoom window. The aim here is to isolate an image blob which is distinct from its background, i.e. it has a closed border as indicated in Figure B7, and its red cross is quite stable with small movements of the mouse.

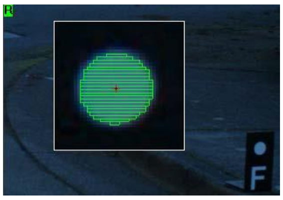

Figure B7

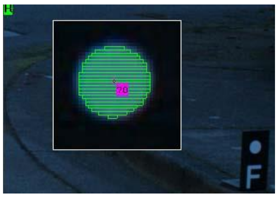

Figure B8

iv) A left-click of the mouse records the centroid position, as shown in Figure B8. A final centroid is thus determined to an accuracy that may well be five times better than with manual marking for suitably exposed high contrast targets (such as the feature point marker shown).

Users of the centroiding tool should ensure that they employ high-contrast 'blob' (e.g. circular) targets if they expect optimal centroiding performance. White blobs are preferable to black, though the centroiding works equally well for both.

B4 Referencing the Third and Subsequent Images

To extend the referencing to a third image, firstly open the desired image by double-clicking on the thumbnail. This image should share a sufficient number of feature points with an already referenced image. Figure B9 shows the addition of the third image, which contains most of the points already referenced in the first two images. The procedure for referencing and subsequently orienting the third image is then as follows:

- Select reference mode, e.g. by clicking the red **R** button on the two images to be referenced.
- Reference a minimum of 5-7 well-distributed referenced points in the already oriented image to the corresponding points in the third (new) image. Continue to reference to previously referenced points until the third image is oriented. You will know this has occurred when:

  a) the tick mark at the top right of the thumbnail of the third image turns to green, and orange newly referenced points will also turn white or green,

  b) the third camera will appear in the 3D view, and

  c) the predicted positions in the third image of already referenced points in the first two images (but not yet in the third) will be displayed in blue.

This is shown in Figure B9. The predicted point location concept is similar to the epipolar line prediction previously discussed; the ‘true’ image position of the point to be referenced should lie near (but not necessarily at) the position indicated by the blue cross.

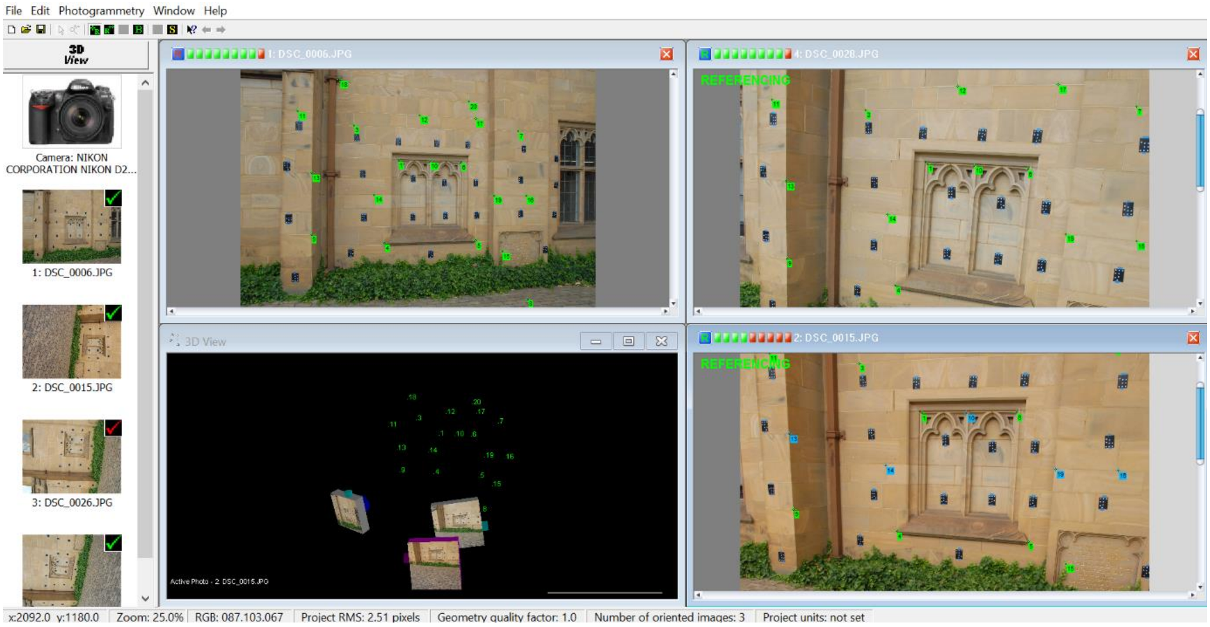

Figure B9

- Referencing of the ‘blue points’ proceeds by first clicking on the already referenced point, followed by a precise marking of the target in the third image. At this point the green reference label will appear and the blue marker for that point will disappear. If the blue labels are in the way of a precise marking, simply use the **B** key to toggle the blue labels on and off.

Warning Note: In referencing the points indicated by the blue labels, do not simply click over the blue marker, since this is the predicted position, which may be in error. The actual point should be precisely marked in the referencing process. A click near to a point marked in the *Image Scanning* process is OK.

- The referencing process just described requires two clicks per 3D point, one in each image. There is also a process that allows 1-click referencing of the predicted (blue) feature points:

  i) Zoom-in in both images.

  ii) With the green reference pencil positioned in the new image (the image with the blue points), press either the **right arrow key** or the number **2** on the keyboard. The next point to be referenced is then displayed in purple in the middle of one viewing window, whereas the already referenced corresponding point is shown in the window of the image being referenced to, with the point label also being coloured purple. To fully reference the point, simply precisely position on the feature point and left-click. The referenced pair will then both change to green. To move to the next point, press the right arrow key or the **2** key again. To move backwards, use the left arrow key or **1** key. The **1** and **2** keys stop functioning when all points are measured.

Once the third image is fully referenced, a fourth image is opened for referencing to any or all three of the already referenced images. The same process continues until the desired number of images and points are added to the network.

Note on the number of open images: The choice of the number of images tiled in the workspace is left to the operator. The minimum for referencing (green pencil) is two. Already referenced images can be ‘closed’ (click the “X” on the image titlebar) to free up work space. To close all images not currently being referenced (all but two), select **Window|Close NonReferencing** or Right-Click and select **Close Non-Referencing Images**. To close all windows, select **Window|Close All** from the main menu.

***Feature*: Display of Orientation, Reference Accuracy Indicators and Geometric Quality**

A 9-point scale shows orientation quality for an image. It starts as fully red when referencing begins, and stays red until a successful orientation has been achieved. A number of squares will then turn green. As the orientation quality improves, as more points are referenced in more images, the number of green squares will increase. The aim is to ensure that as many as practicable are green. As a general rule, all squares will be green once about 20 points have been referenced.

Two measures of overall network orientation quality are displayed at the base of the workspace:

The **Total Project RMS** indicates the **accuracy of referencing;** desirable values are in the range of 0.1 to 2. The **Geometry quality factor** indicates the strength of the network geometry, with 1 to 1.2 being good, 1.3 to 2 being average, and above 2 being poor. Geometry should be improved by additional images and/or points if a quality indicator >1.5 is displayed and as strong camera calibration is sought.

B5. 3D Point Information Dialog and Point Descriptions

Information on 3D points, and blue 'predicted points', are obtained in two ways. The first is via tooltips, which will appear in the Image View when the cursor is hovered over a green or blue point. A second approach is via the ***3D Point Information*** dialog. To obtain this dialog, either select **P** when a point is highlighted, or double-click on the Point Label in the **3D List**. The dialog box lists the point standard errors (SX, SY & SZ), the point quality factor, the names of the images upon which the point has been referenced, and the residual triangulation error vector length for each imaging ray to the point. It is also possible using this dialog to change the label of the 3D point. The button **Open all images that see:** *i* (point label) in the dialog box open the Image Views for all images that 'see' the particular point.

B6. Performing the Self-Calibrating Bundle Adjustment

Once the manually-built network comprises the desired number of images (4 or more) and points (say 30 or more), the self-calibration process can be initiated. Shown in Figure B10 is an example network of four images and 20 points (in practise more images and points would be warranted). Note in the figure how the Project RMS value is 2.5 pixels, which indicates that the camera has likely not been calibrated, as otherwise this figure would be expected to be less than 1 pixel.

The first step is to double-click on the camera icon to bring up the camera parameters dialog, as shown in Figure B11. The desired camera parameters are then unticked so they are no longer 'fixed' at pre-assigned values (generally zero except for focal length) and are free to be computed in the self-calibrating bundle adjustment. Here, the chosen parameters are the quite common set of focal length, principal point offset and radial distortion.

The self-calibrating bundle adjustment is now initiated via the B toolbar button. Shown in Figures B12, B13 and B14 are bundle adjustment setup dialogs and the summary dialog from the adjustment, where is can be seen that the RMS value now falls to 5 micrometres or 0.8 pixel. The final camera calibration parameters are shown in Figure B15.

From this point, all other functions of the *CameraCalibrator* for post orientation can be applied, just as in the automated calibration. Such functions include the application of a network scale and XYZ coordinate alignment (Section 6) (not necessary for camera calibration, but useful for analysing network measurement accuracy), the running of different calibration parameter sets (Section 7), and the outputting of the final Calibration Report (Section 5).

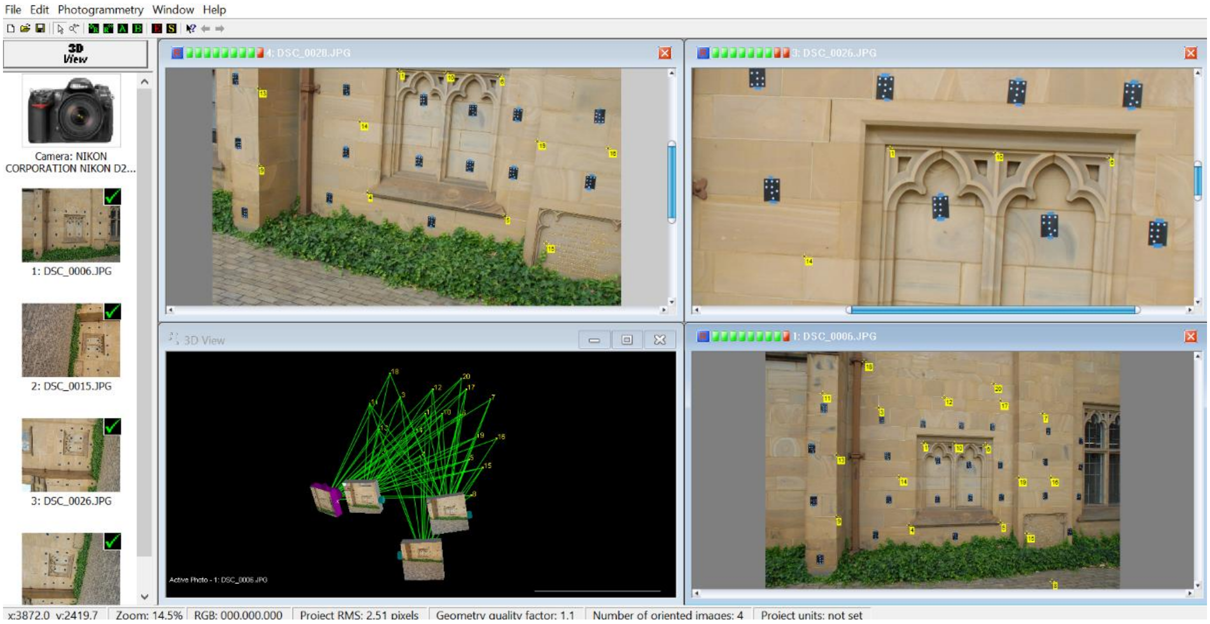

Figure B10

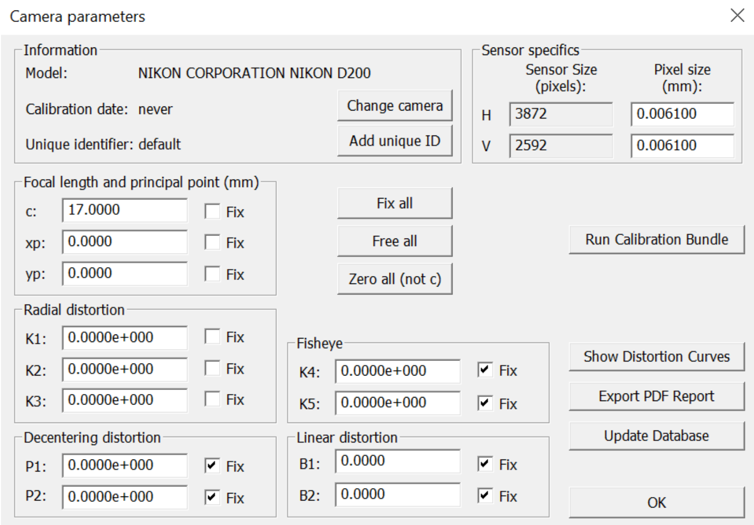

Figure B11

B7. Review Mode

Review Mode, which is initiated by selecting the **Edit/Review** button on the Toolbar (Indicated with a capital **E**), allows a point-by-point review of all image point markings (ie image point referencing operations). As the name implies, Review Mode represents a final quality control procedure which is commenced once the network has been formed. It is really only appropriate for manually referenced points, so it will only be briefly described here. By presenting a visual display of every marking for a given point, as indicated by point 22 in Figure B16, the review process allows the operator to verify that the same physical feature point has been precisely marked in every image. If this is not the case, the operator can either move the marked point to the correct position or delete it if desired.

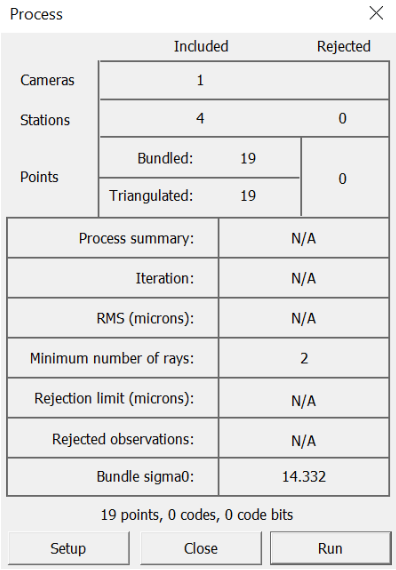

Figure B12

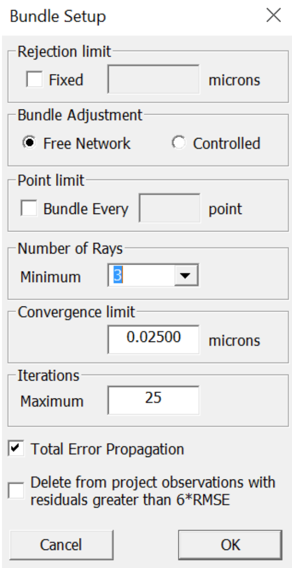

Figure B13

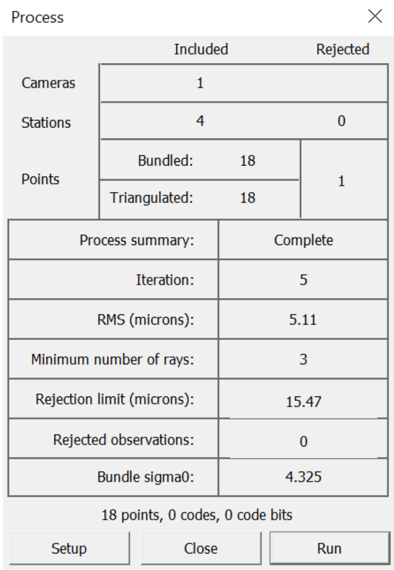

Figure B14

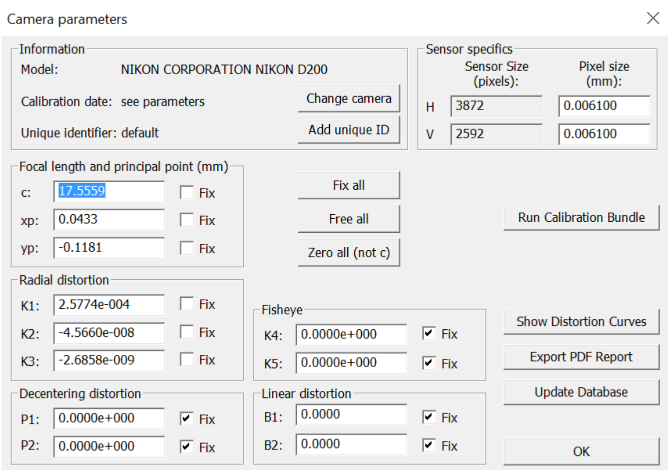

Figure B15

Figure B16

The review process is carried out as follows:

i) At the desired stage, typically after all images have been oriented and points referenced, select the **Review Mode** button on the toolbar (the **E** button). The dialog in Figure B17 will appear, allowing the user to choose the number of images to be displayed, the options being 1, 4, 6 or 9.

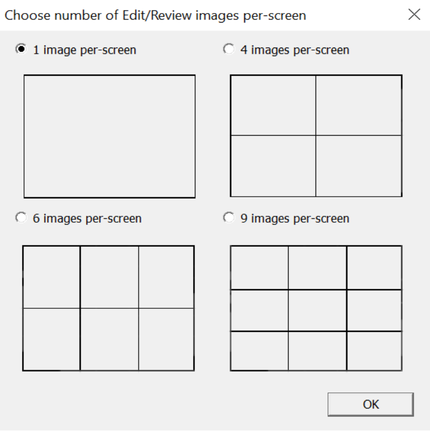

Figure B17

ii) The workspace will then display the images. In Figure B16, the manually referenced point 22 is shown. Note the small dialog box which allows the user to move forward to the next point, move backwards to the previous point or move to any selected point (select the number from the list).

iii) Also, the number of imaging rays to the point is shown, the maximum intersection angle, and a Quality Value which is the RMS value of image measurement errors (in pixels) for the individual triangulated point. Values in the range of 0.05 to 2 would be expected for this. Finally, there are zoom buttons to further enlarge or reduce the images (all together).

<u>Note on number of images displayed</u>: The operator can use the '>>' or '<<' buttons to move forward to the next subset of 1, 4, 6 or 9 image chips of the point currently being reviewed, or back to the previous subset. Thus for a network with, say, 20 images, the user must proceed, if desired, through a number of image sets to review all possible images.

iv) The marked, referenced point in any image can be dragged to a new position using the blue pencil cursor and by holding down the left mouse button (if the blue predicted points are not being displayed, hit the **B** key).

v) Every time a point is 'moved' in Review Mode, the full network is immediately recomputed, so small changes will be seen in the Point Table list, especially for the point being reviewed.

vi) In instances where multiple points appear in an enlarged image view, only the current point can be adjusted.

vii) Where a point does not lie within the format of an image, a large red cross will be shown through the centre of the image.

viii) On the other hand, if the ray from an object point to a camera station does fall within the image format then the predicted point location will be indicated in blue, as in Figure B12. It is now possible to measure (reference) this feature point, if visible, by simply positioning the cursor at the correct position and clicking the left mouse button. The point will then show in green.

ix) Within Review Mode, it is possible to drive to the next/previous 3D point that has predicted positions (ie unreferenced blue points) by using the #1 or #2 keys, or the left and right arrow keys.

<u>Note</u>: The operator can perform all the normal functions within the Review Mode images such as zooming, un-referencing points and deleting referenced points.

Upon completion of Review Mode (select the Finish button) *CameraCalibrator* returns to the status it was at before Review Mode was selected.

Appendix C

Image Coordinate Correction Function in the *CameraCalibrator*

C1. Calibration Parameters and Correction Equations

The image coordinate correction function in the *CameraCalibrator* is the commonly used 10-parameter 'physical' model employed in digital close-range photogrammetry. The calibration parameters can be grouped as follows:

- Camera interior orientation: c, x_p, y_p (principal distance & principal point offset)
- Radial distortion parameters: K_1, K_2, K_3
- Decentring distortion parameters: P_1, P_2
- Affinity, non-orthogonality parameters: b_1, b_2 (rarely used; typically inappropriate for CCDs)

The corrected image coordinates (x_{corr}, y_{corr}) can be calculated from the measured coordinates (x_{meas}, y_{meas}) by using the formulas

$$x = x_{meas} - x_p$$

$$y = y_{meas} - y_p$$

where x and y are now with respect to the principal point, and

$$r^2 = x^2 + y^2$$

$$dr = K_1 \cdot r^3 + K_2 \cdot r^5 + K_3 \cdot r^7$$

$$x_{corr} = x_{meas} - x_p + x \cdot dr / r + P_1 \cdot \left(r^2 + 2x^2\right) + 2 \cdot P_2 \cdot x \cdot y + b_1 \cdot x + b_2 \cdot y$$

$$y_{corr} = y_{meas} - y_p + y \cdot dr / r + P_2 \cdot \left(r^2 + 2y^2\right) + 2 \cdot P_1 \cdot x \cdot y$$

It is noteworthy that b_1 and b_2 are invariably set to zero (ie suppressed).

The additional parameters (calibration values) extracted from the *CameraCalibrator* should be applied as per these correction equations, without change of sign. **Thus, calibration terms in the *CameraCalibrator* can be thought of as <u>corrections</u> and not calibration 'errors'.**

C2. Image Coordinate Reference System

The pixel and image coordinate systems used in the *CameraCalibrator* are shown in Figure C1.

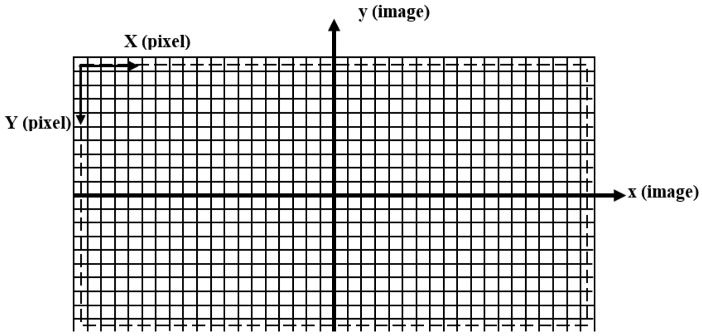

Figure C1: Image coordinate system is centred on the image while the pixel coordinate system origin is placed at the centre of the top-left pixel, with positive Y-axis directed downwards.

The pixel coordinate system origin is at the centre of the pixel at the top left of the sensor with the positive X-axis directed towards the right and the positive Y-axis directed downwards. The image coordinate system is centred on the sensor with the positive x-axis again directed towards the right and the positive y-axis directed up. In Figure C1, the solid outer border represents the physical sensor extents, while the dashed interior border is meant as an aid in visualising the position of the pixel coordinate system.

C3. Converting from Pixel to Image Coordinate System

The location of an artificial target or natural feature is recorded to sub-pixel accuracy in the pixel coordinate system as shown in Figure C1. A half pixel is subtracted from both X and Y to account for the fact that the pixel coordinate system is centred at the top-left pixel. Before use in any adjustment, the measured coordinates are then converted to the image coordinate system as follows.

i) The centre of the sensor in the pixel coordinate system is determined by

$$CentreX = \left(\frac{nXPixels}{2}\right) - 0.5$$

$$CentreY = \left(\frac{nYPixels}{2}\right) - 0.5$$

ii) The measured image coordinates are then

$$\mathrm{x}_{\mathrm{image}} = (XPixel - CentreX) \times XPixelSize = \mathrm{x}_{\mathrm{meas}}$$

$$\mathrm{y}_{\mathrm{image}} = (CentreY - YPixel) \times YPixelSize = \mathrm{y}_{\mathrm{meas}}$$

Appendix D

Sample Camera Calibration Report (first four pages)

CAMERA CALIBRATION REPORT

PROJECT DETAILS

Camera: NIKON CORPORATION NIKON D200

Filename: ...a - Projects\Close-Range Projects\Calibrator\Nikon D200-18mm Cal.cal

Calibration Date: 05/05/2016 12:49pm

METRIC CALIBRATION PARAMETERS

Resolution = 3872 x 2592 pixels

Pixel width = 0.0061mm, Pixel height = 0.0061mm

	VALUE	STANDARD ERROR
Principal distance	c = 17.6137mm	0.002mm
Principal point offset in x-image coordinate	xp = -0.0267mm	0.001mm
Principal point offset in y-image coordinate	yp = -0.2051mm	0.001mm
3rd-order term of radial distortion correction	K1 = 2.79029e-004	8.9424e-007
5th-order term of radial distortion correction	K2 = -6.35554e-007	1.3535e-008
7th-order term of radial distortion correction	K3 = 3.65725e-010	6.1678e-011
Coefficient of decentering distortion	P1 = 2.8395e-007	1.105e-006
Coefficient of decentering distortion	P2 = 6.0840e-005	9.818e-007
No significant differential scaling present	B1 = 0.0000e+000	1.640e-020
No significant non-orthogonality present	B2 = 0.0000e+000	1.640e-020
9th-order term of radial distortion correction	K4 = 0.00000e+000	1.6399e-036
11th-order term of radial distortion correction	K5 = 0.00000e+000	1.6399e-040

STANDARD CORRECTION EQUATION

The corrected image coordinates x(corr) & y(corr) can be calculated from the measured coordinates x(meas) & y(meas) by using the formulas:

x = x(meas) - xp

y = y(meas) - yp

x and y are now with respect to the principal point,

r^2 = x^2 + y^2

dr = K1•r^3 + K2•r^5 + K3•r^7 + K4•r^9 + K5•r^11

x(corr) = x(meas) - xp + x•dr/r + P1•(r^2 + 2x^2) + 2•P2•x•y

y(corr) = y(meas) - yp + y•dr/r + P2•(r^2 + 2y^2) + 2•P1•x•y

Camera self-calibration determined in a network of 8 images and 195 points, to an image measurement accuracy (RMS 1-sigma) of 0.07 pixels or 0.44 um, and qf of 1.2.

Produced by CameraCalibrator from Photometrix - http://www.photometrix.com.au PAGE 1 of 5

CAMERA CALIBRATION REPORT

GAUSSIAN RADIAL DISTORTION CORRECTION PROFILE (dr)

For principal distance c, Gaussian radial distortion correction dr (microns) is given for any radial distance r (mm) as:

dr = K1•r^3 + K2•r^5 + K3•r^7 + K4•r^9 + K5•r^11

correction dx = x•dr/r

correction dy = y•dr/r

	VALUE	STANDARD ERROR
c =	17.614mm	0.0019mm
K1 =	2.79029e-004	8.9424e-007
K2 =	-6.35554e-007	1.3535e-008
K3 =	3.65725e-010	6.1678e-011
K4 =	-1.90452e-056	1.6399e-036
K5 =	-2.76127e-062	1.6399e-040

r(mm)	dr(microns)
0.00	0.0
1.00	0.3
2.00	2.2
3.00	7.4
4.00	17.2
5.00	32.9
6.00	55.4
7.00	85.3
8.00	122.8
9.00	167.6
10.00	219.1
11.00	276.2
12.00	337.1
13.00	400.0
14.00	462.4

Produced by CameraCalibrator from Photometrix - http://www.photometrix.com.au PAGE 2 of 5

CAMERA CALIBRATION REPORT

BALANCED RADIAL DISTORTION CORRECTION PROFILE(dr)

For 'balanced' principal distance cb, radial distortion correction dr (microns) is given for any radial distance r (mm) as:

dr = K0•r + K1•r^3 + K2•r^5 + K3•r^7 + K4•r^9 + K5•r^11

cb = 17.2415mm

K0 = -2.11307e-002

K1 = 2.73133e-004

K2 = -6.22124e-007

K3 = 3.57997e-010

K4 = -1.86427e-056

K5 = -2.70292e-062

r(mm)	dr(microns)
0.00	0.0
1.00	-20.9
2.00	-40.1
3.00	-56.2
4.00	-67.7
5.00	-73.4
6.00	-72.5
7.00	-64.4
8.00	-48.8
9.00	-26.1
10.00	3.2
11.00	37.9
12.00	76.4
13.00	116.8
14.00	156.8

Distortion profile is 'balanced' (dr = 0.0) about a radial distance of r = 9.9mm

Produced by CameraCalibrator from Photometrix - http://www.photometrix.com.au PAGE 3 of 5

CAMERA CALIBRATION REPORT

GAUSSIAN RADIAL DISTORTION PLOT [dr shown in micrometres]

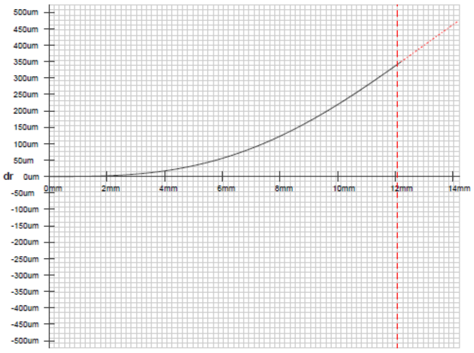

DECENTRING DISTORTION PLOT [P(r) shown in micrometres]

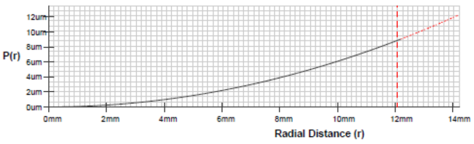

(If present, – – – – indicates the maximum radial distance encountered in the self-calibration.)

Produced by CameraCalibrator from Photometrix - http://www.photometrix.com.au PAGE 4 of 5