CAL

User guide

Updated: March 2016 (version 2.30)

SeaGIS
www.seagis.com.au

Table of Contents

Table of Figures

General

CAL is a photogrammetric bundle adjustment program. It performs a photogrammetric bundle adjustment using optional inner constraints to provide a datum, and also implements an optional stereo constraint specifically for calibrating stereo camera systems.

CAL version 2.00 introduced a major interface change that makes the software more specifically suited to stereo camera calibrations.

Components

The CAL software package has various components that are enabled or disabled depending on your own particular installation. Some parts of this manual may describe components that are not enabled in your installation.

Currently documented optional components include:

- Conversion of movie measurement files to JPEG measurement files.

You can view the currently installed components using the menu command *About | Version* on the main CAL window.

Using this manual

This manual details all aspects of the CAL program. If you want to calibrate cameras (particularly a stereo camera system) finish reading the General section, and proceed directly to the Stereo camera calibration quick guide. This is a step-by-step guide to performing such a calibration, and it has links to other relevant parts of this manual should more information be required.

This manual is designed to be used as an electronic document as it makes extensive use of links.

This manual can be accessed through the CAL software using *About | User manual*.

Licensing

To use CAL you will require a licence. This can be in the form of a USB key provided with the software, or it can be a software-based set of keys. If you have a USB key, install the software, insert the USB key and run CAL. If you are using software-based licence keys the first time you run CAL you will need to enter the licence keys (if you already have them), otherwise contact the software distributor with the details provided by CAL to have your licence keys generated.

If you have a USB licence key take care to secure and protect the key – it is your licence for the software.

Generic editor/display window

CAL uses the same multi-column window to display data and results, and to also allow you to edit data in certain circumstances. When this window is used for the first time in each of its various applications, it will appear as a default size, with default column spacings. You will probably need to resize the window, and resize the columns to suit. Once this is done, CAL remembers the window size and placement, as well as the column widths, and will reuse these values next time the window is displayed.

This window is used to display data such as camera parameters (see Editing camera parameters) and the current project settings (see Viewing/setting current project settings). Data in this window can be sorted by any column, by clicking with the mouse on the column heading.

The data in this window can also be saved to a text file using the menu command *File* | *Save data to text file* on the window. You will be prompted to enter a filename, and the data is saved as a tab delimited text file (suitable for importing into most spreadsheet and database software packages). If you want the data exported in comma separated value (CSV) format, give the selected filename an extension of ".csv". For example saving a file with a name "my_data.csv" will produce a text file in CSV format.

If data in the window that has a green tick beside it can be edited. To edit the value, double click on it. Generally editing takes place where the value appears in the window. However, some values are edited by displaying another window. Boolean type variables are set to true using either 1, t, or T, and false using either 0, f, or F.

Closing CAL

To close the software, use the menu command *Project* | *Close application*. When CAL closes it will prompt you to save any data that has not been saved.

Version and licence information

To see the CAL version use the menu command *About* | *Version*. A dialog displays information about the version and licence type. Full text for the licence file *CAL Licence.rtf* can be found in the installed */Program files/SeaGIS/CAL/Executables* directory.

Installation details

Program files

Program files for CAL are installed in the directory */Program files/SeaGIS/CAL/Executables* or */Program files (x86)/SeaGIS/CAL/Executables*. This directory also contains the licence file for this software. CAL will also have a desktop shortcut and will appear in the start menu as *Start* | *Program files* | *SeaGIS* | *CAL* |*CAL*, or *Start* | *All programs* | *SeaGIS* | *CAL* |*CAL*.

Documentation

Documentation for the software is installed in the directory */Program files/SeaGIS/CAL/Executables* or */Program files (x86)/SeaGIS/CAL/Executables*. This directory contains the user manual and any other relevant documentation. A shortcut to the user manual is available through the start menu *Start* | *Program files* | *SeaGIS* | *CAL* | *CAL User manual*, or *Start* | *All programs* | *SeaGIS* | *CAL* | *CAL User manual*.

This manual can be accessed through the CAL software using the menu *About* | *User manual*.

Project files

This section covers all commands in the *Project* menu.

A project file is used to store filenames and other project related information to allow the project to be saved and reloaded. Each time the CAL program is closed the current project file is updated. When the CAL program is launched, the previous project file is automatically reloaded. CAL will also prompt you to save any other files that have been modified.

Reading and writing to disk

A project file can be loaded using the menu command *Project* | *Read from file*. Before the new project file is loaded, the current project file is automatically updated, and if you have any unsaved data (such as measurements, or camera parameters) you will be prompted to save these files as well.

A project file can be manually saved using the menu command *Project* | *Write to file*. This command will save the project file, and prompt you to save any other files associated with the project that need to be saved (including measurements, object points, and camera parameters). This is an effective way to save all data associated with the project. Note that CAL automatically executes this command when the program closes, or when a new project file is loaded.

Creating new projects and data organisation

To create a new project (a new project file) use the menu command *Project* | *New project*. Before the new project file is loaded, the current project file is automatically updated, and if you have any unsaved data (including measurements, object points, and camera parameters) you will be prompted to save these files as well. Beware that CAL may be asking you to save data before asking for the name of the new project file.

All data associated with the project file is cleared, and you will be prompted to save the new project file. Note that centroid parameter settings are retained when a new project is created (see Centroiding).

The CAL program uses several other files (it either creates these files, or reads data from these files). The purpose and extensions for these files are as follows:

Project file	.PrjCAL
Camera file (one for each camera)	.CamCAL
Measurement file	.ObsCAL
Object point file (for the calibration cube)	.PtsCAL

The project file stores the name of the camera files, measurement file, and object point/distance file, as well as other values that allows CAL to return to the exact project settings when a project file is loaded.

It is suggested that all files for a particular project (project file, camera files, image files, and measurement file) should all be stored in the same directory.

CAL also generates and automatic backup of the current measurements (see Saving and reading measurements).

Viewing/setting current project settings

A summary of the current project file settings can be seen using the menu command *Project* | *Display/set current project data*. A window is displayed showing the current project file name, the current picture directory, the current camera files, and the current measurement and object point files.

Renaming files

CAL provides a utility window for renaming files. This is particularly useful for renaming still image files with names that make a numeric sequence (allowing stepping through the still image sequence), and for associating files with a camera name. For example:

> Left_0001.jpg, Left_0002.jpg, Left_0003.jpg, …
> Right_0001.jpg, Right_0002.jpg, Right_0003.jpg, …

To rename files use the menu command *Project* | *Rename files*. A window is displayed. Set the directory of where the files to be renamed are located by clicking on the *Set directory* button, and selecting the directory using the standard Windows directory browser that is displayed. The selected directory is displayed underneath the *Set directory* button. At this stage some files may be displayed in the *File* list.

Change the *Display Files* value to display the type of files required. The value entered should be of the form *.ext where ext is the file extension, for example *.jpg. All files of this type in the chosen directory will now be displayed in the *File* list. In the *File* list select the files to be renamed. Use the *Shift/Ctrl* keys to make multiple selections.

Renamed files will be of the form:

```
BasenameNumber.Extension
```

Define the `Basename` by entering the required text in *Base Name*, and the `Extension` by entering text in *New Extension*. Note that you do not need to enter a '.' with the extension, and generally the new extension will be the same as the existing extension. The *Base Number* defines the `Number` of the first renamed file.

To rename a series of files to Left_0001.jpg, Left_0002.jpg, Left_0003.jpg, … the required settings are:

Base Name	Left_
New Extension	jpg
Base Number	1

Click the *Rename* button to proceed with renaming the files. After the renaming has taken place the *Base Number* is updated to represent the number of the next file that would appear in the sequence.

Take care when renaming files as the files are permanently renamed – no copies of the original files are made or kept. You can do multiple renaming operations without closing the window. When finished with the dialog click either of the *OK* or *Cancel* buttons.

Closing CAL

Use the menu command *Project* | *Close application* or the *'X'* button on the CAL window to close the program.

Cameras

Camera information and access to specific camera operations is achieved through the *Camera* menu.

Camera information is stored in a file (a separate file for each camera). The names of the current camera files are stored in the project file, and are loaded automatically when CAL starts. Note that these camera files have a .CamCAL extension. These files are different to the stereo camera files (with a .Cam extension) that CAL exports for use in stereo measurement software such as PhotoMeasure, EventMeasure or StereoLib.

Reading and writing to disk

A camera can be written to, or read from disk. To read or write a camera file, use the menu command *Camera* | *Left* | *Read from file* or *Camera* | *Left* | *Write to file*. Identical commands are available for the right camera.

If a new camera is read from disk, or saved to disk as a different name, this change is updated in the current project file.

Editing camera parameters

To view or edit a camera's parameters, use the menu command *Camera* | *Left* | *Edit parameters* (*Camera* | *Right* | *Edit parameters* for the right camera). A window is displayed with the following items:

Descriptor	Name for the camera (can be any character string)
Image rows	Number of rows in an image produced by this camera
Image columns	Number of columns in an image produced by this camera
X Pixel size	Size of a pixel in the x direction on the camera image sensor device (in mm)
Y Pixel size	Size of a pixel in the x direction on the camera image sensor device (in mm)
X PP Offset	Principal point offset in the x direction (mm)
Y PP Offset	Principal point offset in the y direction (mm)
Focal length	Camera focal length (mm)
Radial distortion (K3)	K3 Radial distortion (mm^{-2})
Radial distortion (K5)	K5 Radial distortion (mm^{-4})
Radial distortion (K7)	K7 Radial distortion (mm^{-6})
Decentring distortion (P1)	P1 Decentring distortion (mm^{-1})
Decentring distortion (P2)	P2 Decentring distortion (mm^{-1})
Orthogonality	Image sensor orthogonality (dimensionless)
Affinity	Image sensor affinity (dimensionless)

Items with a green tick beside them can be edited by double clicking on the item, and typing in a new value. Items in the *Units* column indicated the units of the parameter. Values in the *Precision* column can be either 'Fixed parameter', or a numeric value. No value at all means the camera parameter has not yet been estimated by a bundle adjustment, whereas a numeric value indicates the parameter has been estimated by the bundle adjustment, and the value is the parameter precision.

The *X Pixel size* and *Y Pixel size* are properties of the camera, and can generally be found on the manufacturer's web site, or on a specifications or data sheet for the camera. The camera manufacturer generally states these values in units of microns. Make sure you convert the values to mm when entering them as camera parameters.

Fixing camera parameters

In some circumstances it may be desirable to limit certain camera parameters. For example it is common to not use the *Radial distortion (K7)* parameter, or sometimes the *Orthogonality* or *Affinity* are not used. Parameters can be set to fixed (not estimated by the bundle adjustment) or free (estimated by the bundle adjustment) using the menu command *Camera | Left | Set fixed parameters* (*Camera | Right | Set fixed parameters* for the right camera). This displays a window where the camera parameters can be set to fixed or free. On the window, any parameter with a green tick beside it can be *fixed* by setting its value to *true*. To set a value to true double click on the parameter and enter a '1', 't', or 'T', entering any other value sets the value to false.

When a parameter is fixed, be sure to enter the actual fixed parameter value (generally 0.0). See Editing camera parameters to set the camera parameters.

Pictures

This section deals with the *Picture* menu, which allows images to be loaded and provides an interface for configuring movie sequences. General image manipulation, synchronisation and the movie player are also described.

Setting the picture directory

CAL will only load images (still images or movie files) from a specific directory that you nominate. This means all imagery for a particular project must be located in the same directory. To set the picture directory use the menu command *Picture | Set picture directory*. A standard directory chooser window is displayed, and the required directory can be chosen.

Loading images

To load an image use the menu command *Picture | Left | Open picture* for the left picture, or *Picture | Right | Open picture* for the right picture. This displays a standard Windows file dialog allowing the selection of an image file. Note that when initially loading images, you should uncheck the *Lock* check box to stop CAL trying to maintain synchronisation whilst images are loaded (see The 'Lock' function).

If you load a movie that is not part of the movie sequence definition, the movie will be automatically added to the left or right movie sequence and the movie sequence definition window will be displayed (see Movie sequences). If there are multiple movies in the movie sequence definition, check their order is correct then close the movie sequence definition window.

When CAL starts it will attempt to load the pictures that were open when the program was last shut down.

When an image is loaded the image dimensions (rows and columns) are checked against the current camera parameters (see Editing camera parameters). If the image rows and columns do not match the camera rows and columns a warning message is displayed – do not ignore this warning as the camera and image rows and columns must match. This check is also applied to movie files every time a different frame from the movie file is displayed.

Note that only images that are in the current picture directory can be opened. If you try to open valid images that are not in the current picture directory an error message will result (see Setting the picture directory).

To successfully load movie files the correct video codec for the movie file must be loaded on the computer. Note that Windows Media Player might be able to play the movie file, however CAL may not be able to open the file. If Windows

Media Player can play a movie, but CAL cannot open it, this almost certainly will be because the required codec is not installed.

For CAL to be able to use a particular movie file, the installed codec must support *frame accurate* driving (this is because CAL requires frame accurate stepping through movie files to make measurements on a specific frame).

Movie sequences

A sequence of consecutive movie files can be treated effectively as a single movie by defining a movie sequence. The movie sequence definition is used to provide a continuous time reference throughout the movie files that make up the sequence, so it is logically important that the movie sequence is made up of files that are continuous in time. There is a different movie sequence definition for the left and right camera.

Movie sequences are used to keep synchronisation between left and right movies, particularly when using the movie player (see The movie player) and jumping to frames (see Jump to frame). To make use of these features a movie sequence should *always* be defined, *even if there is only a single movie file in the sequence*.

When a new movie file is loaded it is automatically added to the left or right movie sequence definition (see Loading images). Alternately, movies can be added to the movie sequence definition using the menu command *Picture* | *Left* | *Define movie sequence* for the left images, or *Picture* | *Right* |*Define movie sequence* for the right images. A window is displayed showing the current sequence configuration. The window can be re-sized by dragging its edge.

To add movie file(s) to the sequence click the *Add file(s)* button, and use the file dialog to select the required files. You can make multiple file selections. Note that the file dialog will default to the current picture directory. You should only add files from a single directory – and this would normally be the picture directory. If you make multiple file selections the software may take a short time before becoming responsive again. This is because each selected movie file is opened, and the number of frames in the movie and the movie frame rate are determined. The total number of frames and total time (in minutes) for the sequence are displayed above the list of files. If you make a multiple file selection, the files will be added to the movie sequence in alphabetical order.

Once the sequence files are added, they must be ordered correctly by time. The sequence order can be changed by selecting a file and using the *Move up* and/or *Move down* buttons as required. A file can be removed from the sequence by selecting it and clicking the *Remove* button. All files can be removed using the *Remove All* button.

When defining a movie sequence you will be prevented from:

- Adding the same file twice.
- Adding files that are not movies.
- Adding files that the software cannot load.

- Adding files that have different frame rates from others already in the sequence.

CAL measurement files that were created in versions prior to movie sequence support can be 'converted' by loading the measurement file and defining a movie sequence (you will need the original movie file(s) to do this).

Saving the current image as a JPEG

The currently displayed image can be saved as a JPEG image. This function is only enabled for movie files. *Right click* anywhere in the left or right image and a menu will appear. Select the *Save as JPEG* option. A standard file dialog appears allowing you to select a name for the created JPEG image file. Any current brightness and/or contrast adjustments (see Brightness and contrast) will be saved in the JPEG image file.

If you right click close to a measurement point, the point editor window will appear (rather than the pop-up menu). If this happens, click the *OK* button to close the editor, and right click away from any measurement points to access the pop-up menu.

Note that your measurements will not be overlayed in the saved image.

Jump to frame

When a movie file is loaded you can jump to a specific frame number or time in the movie. *Right click* anywhere in the left or right image and a menu will appear. Select the *Jump to frame* option. A dialog appears showing the range of valid frame numbers for the current movie file. Enter the desired frame number and click *OK*.

To jump to a specific time enter a value using the format:

H:M:S

Where H is hours, M is minutes and S is seconds. You must enter at least one colon (':') in the value as CAL uses the presence of a colon to know the value is a time. For example, to move to a position of 36.5 seconds enter a value of:

:36.5

Note that hours and minute values do not need to be entered, although the following entry would get the same result:

0:0:36.5

A jump to time request will use the movie sequence (see Movie sequences) to jump to a time in the movie sequence, and a different movie from the sequence

will be automatically loaded if required. A request for a frame number will refer to a frame in the currently loaded movie file regardless of the movie sequence definition.

If the movies are *Locked* (see The 'Lock' function), then both movies will be moved in synchronisation.

Zooming

The amount of zoom is controlled using the *Zoom* combo box. To zoom in on a particular area of the image, position the mouse in the middle of the area to be zoomed, hold down the *Ctrl* key (on the keyboard), and move the mouse slightly. The image is zoomed, centred on the mouse position.

To zoom out, position the mouse in the image, hold down the *Ctrl* key, and move the mouse slightly.

Stepping through movie and still images

The  buttons can be used to advance frames in a movie file, or to step through a still image sequence (to step through a still image sequence the image files must end in a numeric component such as pic0000.jpg, pic0001.jpg, pic0002.jpg, …). These buttons are known as 'spin controls' and can be held down continuously to repeatedly step through frames in an image sequence.

The stepping controls next to the *Zoom* combo box will step both left and right images together. The text box next to this stepping control defines the step size that is used for all stepping operations (this would normally be set to 1 for still images). The stepping controls directly above the left and right images step only the left or right images respectively. When the images are *Locked* (see The 'Lock' function), the individual image stepping controls are disabled.

Using the stepping controls with the *Ctrl* key pressed will step back or forward 5 frames. Using the stepping controls with the *Shift* key pressed will step back or forward 25 frames.

If you have defined a movie sequence (see Movie sequences), stepping outside the frame range of the currently loaded movie will cause the next or previous movie in the sequence to be loaded.

Immediately above each image is displayed the current movie frame number as well as the elapsed movie time in decimal minutes. If a movie sequence is defined (see Movie sequences), the time value will be the elapsed time through the movie sequence.

The frame stepping functionality can be accessed through the following keyboard key combinations (note the keyboard shortcuts apply stepping to both images):

- *Left* and *right arrow* keys step back or forward 1 frame.
- *Ctrl* key with the *Left* and *right arrow* keys step back or forward 5 frames.
- *Shift* key with the *Left* and *right arrow* keys step back or forward 25 frames.

The movie player

CAL has a built in movie player. If the current picture is a movie you can launch the movie player by clicking the *Play movie* button above the picture. The left image movie player can also be launched by pressing the keyboard *Space bar*.

As the movie player window is displayed, you will be prompted to disable the movie sound. This prompt will only happen if sound is available. By disabling sound you may be able to achieve faster maximum movie playback speed. The option to disable sound is only available the first time a movie is displayed in the movie player.

The movie player window will be displayed. Initially, the movie will be positioned at the same frame as the main CAL window. In summary:

- The movie player window can be dragged and re-sized. The last size and position is maintained and restored next time the player is used.
- Use  and  to start and pause movie playback.
- Use the *Rate* combo box to set the playback rate. If the selected rate cannot be used, the previous rate is automatically reset.
- Use the *Sound* button to toggle sound on and off. If the *Sound* button is a 'X', there is no sound available.
- The *Frame stepping* controls can be used to step back and forward by one frame at a time. Hold the control down to keep stepping forward or back. If the movie is being played and the *Frame stepping* controls are used, the movie will automatically pause.
- The slider control at the top of the window can be used to change the movie position. For most accurate control it is best to pause the movie when using the slider control.
- Above the slider control, the movie position is indicated in *Hours:Minutes:Seconds.tenths of second* format, with the current frame number in brackets. If you have a movie sequence defined (see Movie sequences), the displayed time will be relative to the movie sequence.
- The movie aspect can be altered between *None* (conforms to the size of the movie player window) or *Letterbox* (uses aspect ratio information embedded in the movie) by right clicking on the movie and selecting *Aspect* and then the desired setting.

When finished with the movie player window, use the *Close player* button to close the movie player. CAL will show the same movie frame as before the movie player was used. You can also close the movie player window using the *Close player and update position* button. Using this button will cause the main CAL window to change the movie to the same frame as the movie player window.

If the movies are *Locked* (see The 'Lock' function), then both movies will be moved in synchronisation.

Note that the movie player does not play through movie sequences seamlessly. You need to close the movie player and manually load the next required movie (see, for example Loading images or Stepping through movie and still images), then reload the movie player.

The movie player supports the following keyboard-based shortcuts:

- *Space bar* toggles play and pause (same as the ▶ / ⏸ button).
- *Left* and *right arrow* keys step back and forward one frame (same as the *Frame stepping* controls).
- *Ctrl* and the *left* and *right arrow* keys step back and forward by 5 frames.
- *Shift* and the *left* and *right arrow* keys step back and forward by 25 frames.
- *Ctrl* and mouse click on the left and right *Frame stepping* control steps back and forward by 5 frames.
- *Shift* and mouse click on the left and right *Frame stepping* control steps back and forward by 25 frames.
- *Up* and *down arrow* keys adjust the movie playback rate (same as the *Rate* combo box).
- *Escape* closes the movie player without updating the movie position (same as the *Close player* button).
- *Enter* closes the movie player and updates the movie position (same as the *Close player and update position* button).

The *Resect* and *Auto find* buttons

The *Resect* button can be used to force a resection at any time. At least four of the resection points (see Viewing and manipulating object space points) need to be measured before a resection can be calculated (in some special cases of camera/object geometry, a fifth resection point measurement may be required). If a resection fails an error message box will be displayed.

During the resection calculation image measurements may be automatically excluded if the measurement residual is excessive. To set the limit for excluding resection points see Values for measurement rejection. Where points are automatically excluded during a resection, their point colour will be set accordingly (see Point colours).

Once an image has been resected all other object points can be automatically measured using the *Auto find* button on the camera window. This function is little used as the auto find function is executed automatically when an image is resected. Auto finding is usually performed over all images in the current project after a successful bundle adjustment and deletion of erroneous measurements (see Find targets in all images, and Computing the adjustment). Note that the auto find function will only try to automatically measure object points that have not been measured in the current image.

To see information about resection and automatically finding targets in a measurement context, see Point measurement.

Brightness and contrast

The brightness and contrast of images can be adjusted using the menu command *Picture | Adjust image brightness and contrast*. This command will display a small window with sliders for adjusting brightness and contrast. Drag the sliders to adjust the brightness and contrast. Values of 0 for brightness and contrast represent the default picture state. The brightness and contrast adjustment is applied to both left and right images.

The brightness and contrast adjustment window can be moved to a convenient place on the screen. When finished making adjustments the window can be closed by clicking its *X*, or by using the menu command *Picture | Adjust image brightness and contrast.* If you are dealing with a series of images that require brightness and contrast adjustments, the window can be left open and the brightness and contrast settings will be applied to images as they are opened.

When the brightness and contrast adjustment window is closed, the image brightness and contrast is returned to its default state.

The 'Lock' function

The lock function is turned on and off using the *Lock* check box, located directly above the left image. The lock function is used to help maintain stereo synchronisation of the left and right images.

When the *Lock* check box is *checked* CAL does several things to try and maintain stereo synchronisation:

- The frame stepping buttons that appear above the left and right images are disabled, and images can only be stepped using the controls that apply to both the left and right images.
- If you use the movie player for either the left or right image, and click the *Close player and update position* button (see The movie player), CAL will try to adjust the position of both movies to maintain synchronisation. This only applies to movies.
- If you use the jump to frame function (see Jump to frame), CAL will try to adjust the position of both movies to maintain synchronisation. This only applies to movies.
- If you load a new image file for the left or right camera, CAL will attempt to use the movie sequence definition to maintain synchronisation. This functionality is only useful for movies; *do not use the lock function when loading still images*.
- If any of the above operations are performed and CAL is unable to maintain synchronisation, the image that cannot be synchronised will be cleared.

Compare with the *Lock* check box being *unhecked*:

- The frame stepping buttons above the left and right images are enabled, and the left and right images can be stepped independently.
- If you use the movie player for either the left or right image, and click the *Close player and update position* button (see The movie player), CAL will only adjust the movie position of the movie used in the movie player. This only applies to movies.
- If you use the jump to frame function (see Jump to frame), only the left or right movie will be adjusted. This only applies to movies.
- If you load a new file for the left or right camera, EventMeasure will not attempt to maintain synchronisation.

The suggested use of the *Lock* check box for movies is as follows:

- Uncheck the *Lock* check box.
- Define a movie sequence for the left and right images (see Movie sequences) and load the left and right movies (see Loading images).
- Establish synchronisation using combinations of the individual left and right stepping controls, movie player, or jump to frame functions.
- Once synchronised, check the *Lock* check box.
- Move through the imagery using the stepping controls and movie player.

The suggested use of the *Lock* check box for still images is as follows:

- Name the image files numerically (see Stepping through movie and still images, Renaming files).
- Uncheck the *Lock* check box.
- Load the first left and right image files using *Picture* | *Left* | *Open picture* and *Picture* | *Right* | *Open picture.*
- Check the *Lock* check box.
- Step back and forward through the still image sequence using the stepping controls (see Stepping through movie and still images).

See also the effect of the *Lock* check box with the *Images and orientations* list (see Displaying an image).

Also note that although the *Lock* function makes it difficult to accidentally upset the synchronisation, it is possible for left and right movies to drift out of synchronisation due to various factors. For this reason you must regularly check synchronisation is still valid.

Making image measurements

Refer to Point measurement to see a detailed description of how to make image measurements.

Images and orientations

This section describes the *Images and orientations* list box, located on the main window, directly below the left image.

Manipulating the list box

Before any measurements are made this list box will be empty. As image measurements are made the list box is filled with data. The list box is multi-column, and you may need to adjust the column widths to correctly see the data. Data in the list box can be sorted by any column by clicking on the column header.

As new images are measured, the image name is added to the *Image file* column of the list box. If the image comes from a movie file, the frame number of the image is displayed in the *Frame* column. If the image is a still image, the frame number will be zero. In the *Camera* column, either *Left* or *Right* will be displayed depending on which camera the image is associated with. The *No. Points* column shows the number of points that have been measured in the image. Finally, the *Oriented* column will display *Yes* when the image has been resected, and *No* when the image is not yet resected (see Point measurement). To participate in a bundle adjustment the image must be resected.

Displaying an image

An image listed in the *Images and orientations* list box can be displayed by left clicking on the item in the list box. If the image is a frame in a movie, and the images are *Locked* (see The 'Lock' function), CAL will attempt to maintain synchronisation based on the *Lock* function. If the image is a still image, the *Lock* setting has no effect.

Deleting an image and its measurements

An image can be removed from the project, and all its associated measurements can be deleted by clicking on the appropriate line in the list box (thereby selecting it), and pressing the *Delete* key (on the keyboard). A message box asks for confirmation before deleting the data. This list box only allows single selections. If you want to delete multiple images and their data from the project see View orientation data.

Viewing the orientation data for an image

The camera position and orientation data for a particular image can be viewed by double right clicking on an item in the list box and selecting *View camera position and orientation* from the pop-up menu. A window is displayed, and the image is

displayed in the appropriate camera window. In the window the *Filename*, *Frame*, and *Camera* details are displayed (these are the same as displayed in the main list box, note that camera 1 is the left camera, and camera 2 is the right). If the image has been resected the camera orientation (*Omega, Phi, Kappa*) and position (*Camera X, Camera Y, Camera Z*) information is displayed. If the image has been resected, and involved in a bundle adjustment, the camera position and orientation parameters will also have a precision estimate associated with them. If the image has not been resected, the position and orientation parameters will be *Not estimated*.

Exporting a stereo camera file for an image

A stereo camera file can be generated for any particular image in the *Images and orientations* list box. The stereo camera file is generated by combining:

- The current camera parameters for the image, and
- The current camera position and orientation for the image.

To export a stereo camera file, double right click on the required item (image) in the list box and select *Export this orientation as a stereo camera file* from the pop-up menu. If the image has not been resected CAL will not allow you to export the camera file. If the image has been resected, but the position and orientation have not been refined in a bundle adjustment CAL will warn you before exporting the camera file. Similarly, if the camera is not adequately defined CAL will not allow the file to be exported. CAL displays the parameters that will be exported in a window. Inspect the parameters then close the window. You are then prompted to save the file.

Note that this process *is not* the same as generating a relative orientation for calibrating a pair of stereo cameras. To fully calibrate a pair of stereo cameras see Stereo constraints.

Viewing/deleting the image measurements

The image measurements for a particular image can be viewed by double left clicking on an item in the list box. A window is displayed, and the image is displayed in the appropriate camera window (see also Displaying an image).

In the window, the first column *Point ID (Picture, frame, camera)* displays the point ID, along with the picture name, frame and camera number (camera 1 is the left camera, and camera 2 is the right). The second column *(x, y) image point, pixels* displays the actual image coordinates of the point measurement in pixel coordinates.

The third column *Residual magnitude (x, y) pixels, source* displays residual information for the measurement. There will only be data in this column after the image has been resected, intersected, or involved in a bundle adjustment, with the source component identifying either *resection, intersection,* or *bundle* as the

source of the residual information. The residual is stated as a magnitude $(\sqrt{x^2 + y^2})$ and as the individual x and y components in pixels. Large residuals indicate poor quality measurements that are often erroneous.

The fourth column *Rejected* indicates whether a component of the point measurement has been rejected. This column will only contain data when an image has been resected, or involved in an intersection or bundle adjustment. Either or both of the x and y components of the measurement can be excluded. There are several adjustment settings that control the exclusion of image measurements based on their residual, see Values for measurement rejection.

Measurements can be deleted by selecting them with the mouse and pressing the *Delete* key (on the keyboard). Multiple selections can be achieved using the *Shift* and *Ctrl* keys. Before any deletions occur a window is displayed asking for confirmation. This delete function is often used after a bundle adjustment to delete points that have been excluded by the adjustment. Identify any rejected points by clicking twice on the *Rejected* column header. This will sort all rejected measurements to the top of the list. They can then easily be selected and deleted.

Note that the window will be empty if no points have been measured.

Object points and distances

This section of the manual describes operations that are generally associated with the *Object points* menu. These menu items control the loading, saving and manipulation of object point coordinates and distances measured in the object space. In the case of most stereo camera calibrations the object points and distances refer to points and distances on a calibration cube.

Reading and writing the object point file

The object point file stores object point coordinates and distance measurements. In camera calibration applications the object point file holds coordinates and distances for the calibration cube.

To read an object point file use the menu command *Object points* | *Read from file*. To save an object point file use the *Object points* | *Write to file*.

If you are using a calibration cube to calibrate stereo camera systems, there is no need to save an object point file.

Viewing and manipulating object space points

To view a list of the current object space points, use the menu option *Object points* | *View/edit/delete points*. This displays a window showing for each point:

- *Point ID.* The point identifier (a number).
- *X, Y, Z coordinates.* Coordinates in object space units.
- *Coordinate precisions*. Precision of the coordinates. These fields will be blank until the coordinates are estimated with a bundle adjustment, or intersection.
- *Source*. How the coordinate and precision values have been derived. This will be *bundle* when the coordinates have been estimated using a bundle adjustment, *intersection* when the estimates come from an intersection, or blank if the coordinates have not been estimated yet.
- *No. Obs.* The number of measurements used to estimate the coordinates and precisions.
- *Constraints*. This field lists whether the object point is constrained in x, y, or z, or if the point is to be used when resecting camera positions. In order to resect camera positions there must be at least 4 points that are resection points (usually there will be at least 5 points nominated as resection points).

To delete a point, highlight the point using the mouse and press the delete key (on the keyboard). You can select multiple points using combinations of the *Shift* and *Ctrl* keys with the mouse.

To edit the parameters for a particular point, double click on the corresponding line in the window. Another window is displayed, allowing point details to be edited:

- *Point ID*. The point identifier cannot be edited in this dialog. To alter the point ID the point must deleted, and a new point added by making an image measurement.
- *X, Y, Z*. The coordinates of the point.
- *X Precision, Y Precision, Z Precision*. The precision of the coordinate estimates. Note that precision values will only be present if the point has been involved in an intersection or bundle adjustment. The source of the precision estimate is also given. The precisions cannot be edited.
- *Resection point*. Denotes whether this point is used for resecting camera positions or not. There must be at least 4 resection points in order to perform resections.
- *Coordinate constraints*. This advanced feature is not covered in detail in this manual. Use of coordinate constraints is not required for standard stereo camera system calibrations.

View object point summary

This menu option (*Object points* | *Object point summary*) displays a summary for the current object points. This information is only valid after a bundle adjustment or intersection has been performed. Data displayed is from the last bundle adjustment or intersection calculation.

Item	**Meaning**
Number of points	The number of object points.
Number of resection points	The number of object points to be used as resection points (must be at least 4 to perform resections).
Number of points not estimated	The number of points that could not have a position estimated. This is usually due to not enough image measurements to define the point's location.
Number of point constraints	Number of object point constraints (counted as one for every X, Y, Z ordinate that has a constraint).
Average valid image observations per point	The average number of image observations made to each object point. The average does not include image observations excluded due to large residuals
Minimum point standard deviation	The best object point standard deviation, with units.
Maximum point standard deviation	The worst object point standard deviation, with units.
Average point standard deviation	The average object point standard deviation, taken over all object points, with units.
Overall average point standard deviation	The average standard deviation taken over X, Y, Z ordinates for all points, with units.
Span of object space points	The diagonal distance between corners of the smallest 3D rectangle that encloses all object points.

Relative precision	Relative precision of the network.

NOTE: When this report is generated from coordinates produced in an intersection, the standard deviations and relative precision are affected by the *Image measurement standard deviation* set in the adjustment settings. If this setting is optimistic, the standard deviations will be optimistic as well. If this report is generated from coordinates produced in a bundle adjustment, results will be affected in the same way by the *Image measurement standard deviation* if the *Re-weight observations using sigma zero* is false (generally this option is set to true). See Adjustment settings.

Viewing and manipulating distances

Distance measurements are generally used to provide an object space scale. In the case of calibration cubes, there are usually several measured distances between points on the calibration cube surface. To view a list of current distance measurements use the menu command *Object points | View/edit/delete distances.* This a list of the current distances showing:

- *From/to points.* The object point identifiers for the end points of the distance.
- *Distance.* The measured distance.
- *SD.* The standard deviation of the distance measurement.
- *Residual.* The field is filled with the distance residual from the last bundle adjustment. It will be empty if no such adjustment has been performed.
- *Excluded.* Whether the distance measurement was excluded from the bundle adjustment based on its residual (see Values for measurement rejection).

A distance can be deleted by selecting it and pressing the delete key (on the keyboard). Multiple distances can be selected using a combination of the *Shift* and *Ctrl* keys.

To edit a distance, double click on it. A window is displayed allowing various distance details to be edited. The from point ID, to point ID, distance and precision can all be edited by double clicking on the value to change. If a bundle adjustment has been performed the residual and exclude flag state are also displayed.

Measurement

This section covers the process of making point measurements, as well as describing the use of items in the *Measurement* menu.

Point measurement

Measuring points is the fundamental measurement operation for performing a calibration. In general, you will need to manually measure four points in an image. These points are the resection points (see Viewing and manipulating object space points). Once the four resection points have been identified, the image can be resected, and the remaining object points are measured automatically.

Points can be measured in two ways, either manually or using centroiding. To measure a point manually, position the mouse over the centre of the point and press the left mouse button. To measure a point using centroiding, position the mouse near the centre of the point to be measured, hold down the *Shift* key, and press the left mouse button. While the *Shift* key is depressed a highlight square is drawn in the image showing the area to be centroided. More information on centroiding and centroid parameters is found in Centroiding. Whilst performing point measurements, the zoom functionality can also be used (see Zooming). Centroiding is the preferred method of point measurement as it offers much better automation and accuracy.

When a point is measured it is automatically assigned a point number. The number assigned to a point will be the next available object point number, that has not already been measured in the current image (see Viewing and manipulating object space points to view the object point numbers). Resection points are always set to be the first (numerically) object points. In the case of a calibration cube, points 100, 101, 102, 103, 104 are set as resection points, and other points on the calibration cube would start at 105 and continue upwards.

If a point is measured and the automatically assigned point number is incorrect, the number can be manually adjusted by right clicking on the point. This displays a window that allows the point number to be adjusted. On the window, retype the point number as required and click the *OK* button to accept the change. A point can be deleted by right clicking on it and clicking the *Delete* button on the window that pops up, or simply pressing the *esc* (escape) key on the keyboard.

The procedure for measuring an image is as follows. Firstly, measure (using centroiding in preference to manual measurement) the resection points, starting with the numerically smallest point ID and working upward. Generally, after the first four resection points have been measured the image will automatically be resected. Once the image resects the software attempts to automatically identify all other object points. If the cameras have a previous calibration loaded nearly all other object points will be measured automatically. If the resection is not computed automatically after measuring four resection points, measure a further

resection point and the resection will be computed and all other object points will be automatically measured.

At any time a resection computation can be forced by clicking the *Resect* button above the left or right image. Similarly, once an image has been resected all targets in the image can be automatically found by clicking on the *Auto find* button above the left or right image. See The Resect and Auto find targets buttons.

Using automatic stereo resection and measurement

This feature makes use of the fixed camera geometry of a stereo camera system to automatically make image measurements. To use this feature:

- The imagery must be correctly synchronised stereo imagery.
- The imagery must be suitable for centroided measurement, and the centroiding parameters must be correctly configured (see Centroiding).
- You must set *Use stereo for automatic resection/measurement* to *True* in the Adjustment settings (the default setting is *True*, see Stereo constraints). Note that if your imagery is not suitable for centroiding, you may need to set this parameter to *False* to stop CAL trying to make automatic measurements.

When you begin measuring, you will need to measure the resection points in both the left and right images. After you have measured one stereo pair (possibly two or three if the current camera parameters are not particularly good), CAL will try to automatically measure the other image: so if you measure the left image of a stereo pair, CAL will try to automatically measure the right and vice versa.

Saving and reading measurements

All image measurements are saved in a file. This file also contains the camera positions and orientations (but not the actual camera parameters), the picture directory setting, and the movie sequence definitions. Measurement files have the extension .ObsCAL.

To load a measurement file use the menu command *Measurement* | *Read measurement file*. A standard Windows file dialog is displayed, and the appropriate file can be selected.

To save the current measurements use the menu command *Measurement* | *Write measurement file.* A standard Windows file dialog is displayed allowing the file to be selected.

CAL automatically saves a backup of the current measurement data. The backup is saved every 5 minutes (approximately). The file is saved using the current filename with the text "_AUTO" appended. For example, if the current measurement filename is *Meas.ObsCAL*, CAL will automatically save

measurements to the file *Meas.ObsCAL_AUTO*. If there is currently no filename for the measurement data, CAL will generate a file called *_AUTO*. If the computer crashes, or you forget to save your measurement data, you can use the backup by renaming the file and removing "_AUTO" from the name, then loading the measurement file with *Measurement | Read measurement file*.

Convert movie to JPEG

This function allows measurement files associated with movie imagery to be converted to measurement files using JPEG imagery. Once measurements in a movie file (or several movie files) have been made, this function searches the movie file(s) and extracts all frames that have measurements associated with them. The extracted frames are saved as JPEG images, and a new measurement file is created, the new measurement file making use of the extracted JPEG images.

This functionality provides an effective form of compression. High quality uncompressed movie files can be used to make measurements. Once measurements are complete, this function extracts JPEG images of frames in the movie files where measurements have been made, and creates a new measurement file that relates all measurements to the JPEG images. The large movie file(s) can then be archived or deleted, with all measurements expressed in the newly created series of JPEG images.

Note that this functionality is an optional component and will only appear if you have the component installed.

To perform a conversion use the menu command *Measurement | Convert movie to JPEG*. If the current measurement file is not saved, you will be prompted to save it. If you intend to convert the current measurement file it *must* be saved otherwise the conversion will not be complete (unsaved measurements will not be converted). A dialog is displayed. On the dialog select:

- *Input measurement file*. This is the measurement file to be converted from movie imagery to JPEG imagery. Double click on the item to show a standard file dialog and select the required file.
- *Output JPEG directory*. The directory where the generated JPEG images will be saved. Double click on the item to display a directory chooser dialog (note that a new directory can be created using this dialog). If the chosen directory already contains files with the same naming convention as those generated by this conversion process, the existing files will be overwritten (see the note about naming convention below).
- *Output JPEG measurement file*. The name of the new measurement file associated with the generated JPEG images. Double click on the item to show a standard file dialog.

It is recommended that you make a new directory for the JPEG images, and to save the *Output JPEG measurement file* in this directory with the JPEG images.

Once these parameters are set click the *Process* button (or *Cancel* to quit). A dialog with a *Cancel* button is displayed. The dialog reports progress as the JPEG images are extracted from the movie file(s) and saved. If the *Cancel* button is pressed the JPEG image generation is cancelled and the new measurement file is not written (you will need to manually delete any JPEG files that were created before the process was cancelled). To use the new measurement file (with JPEG images) you need to load the file (see Saving and reading measurements). In the newly generated measurement file the picture directory is set to the chosen *Output JPEG directory*.

The conversion process will only generate JPEG images for movie frames where there are measurements.

JPEG files generated by this process are named as follows:

CamXX_YYYYYY.jpg

Where:

XX is the camera number (1 is left, 2 is right).

YYYYYY is the image sequence number.

The sequence number is calculated by sorting the required movie frames by movie file name(s) and movie frame number. The sorted frames are then assigned a sequence number from 1 through to the number of frames for the camera. The relationship between left and right camera frames used in stereo constraints (for stereo camera calibrations) is maintained throughout this process.

The movie to JPEG conversion utility will deal with requests to convert measurement files that use a mixture of still and movie imagery or only still imagery. Any still images involved in the conversion will have JPEG images generated for them.

Note if you are using movie (usually AVI) imagery compressed with the DivX codec you may need to alter a DivX setting. This problem was noted with DivX version 6.4, and causes the generated JPEG images to be one frame away from the required frame (so measurements do not match to the extracted image). The problem may be present with other versions of the codec. To resolve this issue in DivX version 6.4, in the *DivX playback configuration* (accessed in the *Decoder configuration utility)* make sure the setting *Reduce CPU spikes using one frame delay* is NOT checked.

View orientation data

Use the menu command *Measurement |View orientation data* to view a list of the current camera position and orientation data (also know as camera exterior orientations, or EOs). The list appears in a window with an entry for each image. The image is identified with a *Filename (Frame No., camera)*. In stereo

calibrations camera 1 is the left camera, camera 2 is the right. If the image has not been resected the *X* coordinate will say *NOT ESTIMATED*. If the image has been resected or involved in a bundle adjustment the *X, Y,* and *Z* coordinates for the camera position will be provided as will the *Omega, Phi,* and *Kappa* camera orientations. If the image has been involved in a successful bundle adjustment then the camera position will have precision estimates (*sX, sY, sZ*), as will the camera orientation parameters (*sOmega, sPhi, sKappa*). The final column in the window lists the number of points measured in the image.

You can remove images from the project by selecting them and pressing the *delete* key (on the keyboard). Multiple selections can be achieved by holding down the *Shift* and/or *Ctrl* keys while making the selection.

Double click on any of the orientations to view the image measurements associated with the image. A window with measurements is displayed. Information about each image measurement and its residual (if applicable) is provided. This window is described in Viewing/deleting the image measurements.

View/delete image measurements

Use the menu command *Measurement* | *View/delete image measurements* to display a window that lists all image measurements for all images in the project. This window is identical to that described in Viewing/deleting the image measurements, *except* it lists *all* image measurements, rather than just image measurements associated with a particular image.

This command is particularly useful for deleting points that have been excluded from a bundle adjustment. To group together all the excluded measurements simply click twice on the *Rejected* column header. This sorts the data by this column, and brings the rejected measurements to the top of the list. The rejected measurements can be selected (using the *Shift/Ctrl* keys for multiple selections) and deleted using the *Delete* key (on the keyboard).

This command is also useful for deleting all measurements to a particular object point. Make sure the data is sorted by point ID by clicking on the *Point ID (Picture, frame, camera)* column. Scroll through the list to find the group of measurements to a particular point ID and select them all (using the *Shift/Ctrl* keys), then delete the measurements by pressing the *Delete* key (on the keyboard).

Image measurement summary

The menu command *Measurement* | *Image measurement summary* displays a window summarising measurement residuals for each image. Each image is identified with the *Filename (frame, camera)*. In stereo calibrations camera 1 is the left camera, camera 2 is the right. The second column provides the average image residual as a magnitude and in x and y components in both pixels and units x 1000 (generally microns). This is the average residual over all included image

measurements. The third column provides the maximum image residual as a magnitude and in x and y components in both pixels and units x 1000 (generally microns). This represents the maximum residual over all included image measurements. Values in this column are limited by the thresholds set for measurement exclusion in Values for measurement rejection. The final column gives the count of included image measurements (note that it is possible for the x *or* y component of an image measurement to be excluded). The *average* entry gives the average value over all images for average residual column, and the maximum over all images for the maximum residual column.

Data in this window is helpful in identifying poorly measured images. If an image has average residuals that are noticeably worse than other images it indicates a problem with the image. If such an image is noticed, it should be examined for correct identification of all object points (its often easier to just delete all measurements, and re-measure the image). Above average image residuals may be caused by the image quality itself. If the image has blurred targets due to excessive camera or object movement during exposure the image residuals will be high. If this is the case it is desirable to remove the image (or re-capture it if possible).

Centroiding

For more details about centroiding and setting the centroid parameters, see Centroiding.

Stereo constraints

This section covers menu options under *Measurement* | *Stereo constraints*. These operations are used to configure stereo constraints and compute a relative orientation for a pair of stereo cameras (the relative orientation describes the position and orientation of one camera with respect to the other, and is required to make measurements with a stereo camera system). These options are chiefly provided to generate stereo camera files (.cam extension) for use with PhotoMeasure, EventMeasure and StereoLib.

In stereo camera calibrations, camera 1 refers to the left camera, and camera 2 refers to the right camera. CAL retains a numbering system rather than using the text ‘Left’ and ‘Right’ to facilitate generic calibrations using systems with more than two cameras.

Configuring constraints

Use the menu command *Configure stereo constraints* to set up the constraints. A window is displayed. The main purpose of the window is to allow definition of stereo pairs of left and right images from which the relative orientation of the cameras can be derived.

Select the appropriate *Left camera* and *Right camera* using the combo boxes. In most instances the rest of the configuration can be achieved using the *Automatic* button. Click the *Automatic* button and *pairings* list box is populated with combinations of *Left filename (frame no, camera no)* and *Right filename (frame no, camera no)*. Ensure that the pairings represent stereo pairs – if they do not you will need to manually edit the list. Note that in standard stereo calibrations, camera 1 is the left camera and camera 2 is the right.

To delete a pairing, select it and click the *Delete* button. To manually add a pairing, select a left camera *Filename (frame no, camera no)* and a right *Filename (frame no, camera no)* and click the *Add* button. If the any of the newly added data is present in the current *Pairings* list, it is automatically deleted and replaced with the new stereo pairing.

Click *OK* to accept the new configuration, or *Cancel* to ignored the changes.

Estimating constraints

Once the constraints are configured, the relative orientation can be estimated. Use the menu command *Estimate constraints*. A window displays the computed relative orientation. Details of the window are explained in Viewing the relative orientation.

Viewing the constraints

Once the stereo constraints have been configured (see Configuring constraints) and estimated (see Estimating constraints, or estimated in a bundle adjustment), the quality of the estimation can be assessed using the menu command *View stereo constraints*. This command shows a window. The window displays each stereo pairing on a line, beginning with the left and right camera *Filename (frame number, camera number)*. The following columns give residuals for the *Base, Delta omega, Left phi, Left kappa, Right phi,* and *Right kappa*. The residuals are the difference between the relative orientation of each individual left/right stereo pair and average relative orientation computed over all stereo pairs.

The final column lists any components that have been automatically excluded from the relative orientation computation based on their residual (see Stereo constraints for setting thresholds for these exclusions, and when these thresholds may need to be adjusted). The magnitude of residuals for the stereo constraints largely depends on the quality of stereo camera synchronisation. If one of the constraints has many components excluded it should be suspected that the stereo synchronisation of this left/right pairing is incorrect (possibly an error in the selection of the stereo pair while configuring the constraints, see Configuring constraints). If many of the stereo pairs are being excluded, and the configuration of stereo pairs is correct, the thresholds for exclusion may need to be increased (you will need to increase the values for either/or the *Critical value for base rejection* and *Critical value for orientation rejection* in Stereo constraints, then re-estimate the constraints as per Estimating constraints). Always investigate

synchronisation and image pairing as the primary reason for rejections before adjusting these thresholds.

Viewing the relative orientation

Use the menu command *View relative orientation* to view the current relative orientation parameters. A window is displayed. On the window the *Parameter source* will be *Bundle adjustment* if the parameters are estimated in the bundle adjustment (see Stereo constraints), otherwise it will be *Camera positions and orientation*s (meaning the relative orientation has been derived from the current camera positions and orientations). The number of constraints will be equal to the number of constraints configured in Configuring constraints. The *Base separation (X)* represents the physical separation between the cameras. In typical stereo camera configurations *Delta omega, Left kappa* and *right kappa* will be close to zero, and *Left phi* and *Right phi* represent the amount of inward convergence for each camera (although these values will be around zero as well if the cameras do not converge). Each of the parameters has a precision associated with its calculation.

Exporting stereo camera files

In order to export stereo camera files the stereo constraints must be configured (see Configuring constraints), estimated (see Estimating constraints, or estimated in a bundle adjustment), and the results verified (see Viewing the constraints). The stereo camera files can then be created using the menu command *Export stereo camera files*.

If the constraints have not been configured or estimated an error message will be given. The left camera parameters will be displayed in a window. Click the *Close dialog* button and a message box prompts you to save the left camera parameters. Click *Yes* and a standard Windows file dialog is displayed, allowing you to select the name of the stereo camera file for the left camera. After the left camera file is saved a window will appear with the right camera parameters. Click the *Close dialog* button and a message box prompts you to save the right camera parameters. Click *Yes* to display a Windows file dialog where the right camera filename can be selected.

Resect all images

The menu command *Measurement* | *Resect all images* on the main CAL window computes a resection in all images associated with the current project. At the conclusion of the resection process a window with a summary report is displayed. The window lists all images involved in the resection operation, whether the resection was successful, and the RMS error associated with the resection. See The Resect and Auto find targets buttons for more information about resections. Note that at least four resection points must be measured in an image before the image can be resected.

This option is generally not used as images are usually resected individually as they are measured.

Find targets in all images

The menu command *Measurement* | *Find targets in all images* runs an automatic process where all images in the current project are loaded, and an attempt is made to measure any unmeasured object points in each image. While this process runs a progress window is displayed. The process can be cancelled at any time by clicking the *Cancel* button on the progress window.

This option is generally used after a successful bundle adjustment. Following a bundle adjustment the camera parameters are of a better quality, and it is likely that more object points can be automatically measured. The general procedure is to delete measurements excluded by the bundle adjustment then run this process to automatically measure any un-measured object points. The bundle adjustment can then be re-computed (see more details in Computing the adjustment).

Note that this process will only try to find (in each image) object points that currently have no measurement. If a point is already measured, the current point measurement is maintained without change.

Intersect all points

The menu command *Measurement* | *Intersect all points* uses the current camera positions and orientations for all images, and intersects all image measurements to compute object coordinates. This functionality is used to estimate the object space coordinates for newly measured points.

When performing bundle adjustments with calibration cubes this option is not used.

Delete rejected point measurements

The menu command *Measurement* | *Delete rejected point measurements* will delete all point measurements that have been marked as rejected (point measurements can be rejected by resections, intersections, or the bundle adjustment, see Point colours, Values for measurement rejection). This function is typically used following a bundle adjustment to delete all rejected point measurements.

Adjustment

This section deals with the menu commands under *Adjustment* menu. These commands are used to configure and execute a bundle adjustment. Before computing a bundle adjustment you will need to resect all images (see Point measurement), and measure most of the targets in most of the images (preferably all visible targets in all images).

Convergence criteria

During the bundle adjustment parameters (object point coordinates, camera positions and orientations, and camera calibration parameters) converge towards a final value. The degree of convergence can be fine tuned if required. See also Computing the adjustment.

Object points

To adjust the object point convergence criteria use the menu command *Adjustment | Object point convergence criteria*. A window is displayed allowing the convergence criteria to be adjusted. This value can generally be set to 1/1000th of the object space units, for example 1 micron if the object space units are millimetres.

Camera positions and orientations

To adjust the camera position and orientation convergence criteria use the menu command *Adjustment | Camera position and orientation convergence criteria*. A window is displayed allowing the convergence criteria to be adjusted. There are two values, one for camera orientations in units of seconds of arc, the other for camera position. The units for camera position will be in object space units x 1000. The default values are 1 second of arc for orientation, and 1/1000th of the object space units for position (for example 1 micron if the object space units are millimetres).

Camera parameters

The convergence criteria for camera parameters can be set using the menu command *Adjustment | Camera convergence criteria*. Each of the camera parameters (see Editing camera parameters) has its own convergence criteria. The following table provides default values for camera convergence criteria.

Camera parameter	Convergence criteria	Unit
X PP Offset	0.00001000	mm
Y PP Offset	0.00001000	mm
Focal length	0.00001000	mm

Radial distortion (K3)	0.00000001	mm^{-2}
Radial distortion (K5)	0.00000010	mm^{-4}
Radial distortion (K7)	0.00000010	mm^{-6}
Decentring distortion (P1)	0.00000010	mm^{-1}
Decentring distortion (P2)	0.00000010	mm^{-1}
Orthogonality	0.00000010	Dimensionless
Affinity	0.00000010	Dimensionless

Adjustment settings

The bundle adjustment settings are controlled using the menu command *Adjustment* | *Adjustment settings*. A window is displayed allowing the adjustment settings to be configured. The following describes the settings and gives suggested default values.

Values for measurement rejection

These values control the automatic exclusion of measurements from certain calculations when it is determined the measurement is likely to be in error.

Name: *Critical value for resection rejection*.
Default: 3.00 pixels.
Description: When a measured point is used in a resection computation, and its resulting residual is larger than this value, it is excluded from the calculation. When this is the case the point colour is changed to *Point colour, residual failed* (see Point colours).

Name: *Critical value for intersection rejection*.
Default: 3.00 pixels.
Description: When a measured point is used in an intersection computation, and its resulting residual is larger than this value, it is excluded from the calculation. When this is the case the point colour is changed to *Point colour, residual failed* (see Point colours).

Name: *Critical value for bundle rejection*.
Default: 2.50 pixels.
Description: When a measured point is used in a bundle adjustment, and its resulting residual is larger than this value, it is excluded from the adjustment. When this is the case the point colour is changed to *Point colour, residual failed* (see Point colours).

Name: *Rejection factor for object space distances*.
Default: 5.00.
Description: This factor is used to set a critical value (critical value = distance measurement standard deviation x *Rejection factor for object space distances*) for removing distances from a bundle adjustment. When a distance residual from a bundle

adjustment exceeds the critical value, it is excluded from the bundle adjustment.

Name: *Rejection factor for coordinate constraints*.
Default: 5.00.
Description: This factor is used to set a critical value (critical value = coordinate constraint standard deviation x *Rejection factor for coordinate constraints*) for removing coordinate constraints from a bundle adjustment. When a coordinate constraint residual from a bundle adjustment exceeds the critical value, it is excluded from the bundle adjustment.

Name: *% of critical value for warning point colour*.
Default: 80.0%.
Description: See Point colours, this value can be set on the point colours dialog, and is described there.

Image measurement precision

Name: *Image measurement standard deviation*.
Default: 0.10 Pixels.
Description: Sets the standard deviation for image measurements. As a general rule use 0.10 pixels for measurements made with centroiding, and 0.50 to 1.00 pixels for measurements made manually (by mouse pointing only). This value influences the precision of parameters estimated with the bundle adjustment *ONLY* if the following parameter is set to *FALSE*. Setting this value correctly is important for simulations.

Name: *Re-weight image observations using sigma zero*.
Default: True.
Description: When set to true and computing a bundle adjustment, the image measurement standard deviation (described above) will be automatically adjusted to give a sigma zero value of 1 (± 0.05) . This automatically provides the 'best' estimate of parameter precisions. This value should normally be set to true unless you are computing a simulation. Note that only the image measurement standard deviations are altered (the standard deviation of distance measurements or coordinate constraints is not altered). The new estimate of the image measurement standard deviation is shown in the adjustment results summary (see Computing the adjustment).

Stereo constraints

The following parameters are only of concern when the purpose of the bundle adjustment is to estimate the parameters of a stereo camera configuration.

Name: *Use stereo constraint*.
Default: False.
Description: When set to true, the relative orientation parameters are estimated (or rather constrained) during the bundle adjustment. Effectively, each left and right image pairing are forced to result in the same relative orientation within the limits set by *SD for stereo base constraint* and *SD for stereo orientation constraint* (see directly below).

Name: *Critical value for base rejection*.
Default: 5000.0 microns.
Description: During the computation of a relative orientation (either as part of a bundle adjust or using Estimating constraints) any base component that exceeds this value for a left/right pairing of images will be ignored. The units used for this parameter will depend on the current project. The setting of this value may need to be increased if the stereo synchronisation of the cameras is not perfect (often with video camera systems).

Name: *Critical value for orientation rejection*.
Default: 1800 seconds.
Description: During the computation of a relative orientation (either as part of a bundle adjust or using Estimating constraints) any orientation component that exceeds this value for a left/right pairing of images will be ignored. The setting of this value may need to be increased if the stereo synchronisation of the cameras is not perfect (often with video camera systems).

Name: *SD for stereo base constraint*.
Default: 100.0 microns.
Description: When *Use stereo constraint* (see above) is true, this is the constraint placed on all left/right image pairs of the relative orientation base component. The units used for this parameter will depend on the current project units. This constraint may need to be relaxed if the stereo synchronisation of the cameras is not perfect (often with video camera systems).

Name: *SD for stereo orientation constraint*.
Default: 100.0 seconds.
Description: When *Use stereo constraint* (see above) is true, this is the constraint placed on all left/right image pairs of the relative orientation components. This constraint may need to be relaxed if the stereo synchronisation of the cameras is not perfect (often with video camera systems).

Name: *Use stereo for automatic resection/measurement*.
Default: True.
Description: When set to true CAL will automatically try to 'measure' the second image of a stereo pair. Once the first image of a stereo pair (either the left or right) has been resected by manually measuring the resection points, CAL uses the stereo geometry of all currently measured images to try and automatically find the resection points, resect and measure all targets in the second image. You will need to manually measure at least one stereo pair for the automatic function to start working. The automatic function relies on targets that can be centroided – if you do not have suitable targets, turn of this feature (by setting the value to *False*).

Adjustment

Name: *Maximum iterations*.
Default: 200.
Description: Defines the maximum number of iterations for the bundle adjustment. Where the calibration network is strong, and all object points have been measured, there should be no need to increase this value. If an adjustment will not converge, this value can be increased to allow convergence (generally after the solution converges the first time, erroneous measurements will be automatically excluded and the adjustment will converge in far fewer iterations).

Name: *Apply scale constraint*.
Default: False.
Description: Flag that determines whether the scale is constrained through the bordering of normal equations during the bundle adjustment. This should be set to false when the scale is supplied though object space distance measurement(s), or object point constraints.

Point colours

The colours used to display point measurements can be altered using the menu command *Adjustment* | *Point colours*. A window is displayed allowing the four different point colours to be set. The different point colours are used to indicate the status of the point, and the quality of the point measurement. To adjust any of the point colours, click the button associated with the colour and a standard Windows colour dialog is displayed. Select the required colour from this dialog and click OK.

When points are initially measured they are displayed using the *Point colour, no residual* colour. After a point has been involved in either a resection, an intersection or a bundle adjustment computation, the point will have a residual

associated with it, and the point colour will change accordingly. If the point residual is acceptable, the point is displayed using the *Point colour, residual OK* colour (default colour is green).

If the point's residual fails the acceptable limit (for either resection, intersection, or bundle adjustment, as appropriate) the *Point colour, residual failed* colour is used (default colour is red). When a point is this colour it indicates the point has been automatically removed from the calculation. If a point is this colour *after a successful bundle adjustment has been* computed, the point should be investigated as an erroneous measurement, and either be re-measured or deleted.

Where the point's residual is above the *Percentage of critical value for warning colour* it is displayed using the *Point colour, residual warning colour*. The default for this colour is orange, and is used as a warning that the point's residual is acceptable, but is close to failing. The *Percentage of critical value for warning colour* can also be set in the Adjustment settings window.

Computing the adjustment

After all images are resected, and targets have been measured in all images, the bundle adjustment can be computed. To run the bundle adjustment use the menu command *Adjustment* | *Compute bundle adjustment* on the main CAL window. A progress window is displayed with a *Cancel* button. At any time the adjustment can be cancelled by clicking the *Cancel* button (the process will not cancel until the current iteration is computed). The window displays information about the adjustment progress (the iteration number), and sigma zero (which indicates the quality of the adjustment and should be ideally close to 1). The adjustment first computes an *Initial solution*, followed by a series of *Re-weighting solutions* where poor quality measurements are automatically removed from the adjustment. Finally the *Parameter precisions and correlations* are computed.

At this stage there will either be a converged successful adjustment, or an adjustment that has failed to converge. A new window is displayed. That window has two list boxes, the upper list box gives a text *Adjustment summary*, and the lower list box gives *Details* about the adjustment.

Considering the *Adjustment summary*, the first line will report either a successful converged adjustment or a failure. Other information provided will be:

- Whether image observations were re-weighted to give a sigma zero of 1 (see Image measurement precision), and whether the re-weighting was successful. The sigma zero value is given, as is the new estimate of the image measurement precision.
- Whether any aspect of the adjustment process failed, and if so give an indication of what went wrong. This can include failure to converge within the maximum number of iterations (see Adjustment), or a failure to solve the normal equations (this usually indicates that not enough object points have been measured in enough images, or can be caused by nonsensical object point constraints).

- The average image residual (should generally always be less than 0.2 pixels).
- The total number of adjustment model parameters, constrained parameters, and redundancies. The number of redundancies for a stereo camera calibration will generally be in the order of 5,000.
- The number of eligible and excluded image and distance measurements, and object point constraints. Excluded measurements have failed the parameters set in Values for measurement rejection. Ideally there should be few rejected measurements.
- Warnings if there are object space points that do not have coordinates estimated for them. This will happen if new object points are added without coordinate estimates (this should never happen with calibration cube projects). Use intersection to estimate the coordinates (see Intersect all points).
- Warnings if there are distances measured to object points that do not exist, or to object points that do not have estimated coordinates.
- Warnings if there are un-resected camera positions (see Point measurement).
- Warnings if there are image measurements made to object points that do not exist, or to object points that do not have an estimated coordinate.
- Warnings if there are measurements associated with a camera that is not part of a project (can happen if a camera is deleted from an existing project, but there are still measurements associated with the deleted camera).
- Warnings, if using stereo constraints, and the constraints have not been estimated, or the constraints involve invalid or un-estimated camera positions (see Stereo constraints).

If the adjustment does not converge, the parameters that are failing to converge will be reported. Non-convergence is generally associated with poor network geometry, or not enough measurements in enough different images (this should never be the case with calibration networks involving calibration cubes as in Appendix A Calibration cube orientations). In order to force the adjustment to converge, the convergence criteria for the parameters that are failing can be adjusted to force the network to converge (see Convergence criteria), or the *Maximum iterations* for the bundle adjustment can be temporarily increased (see Adjustment).

The *Details* list box provides more detailed information on the adjustment results. All items listed can be viewed in further detail by double clicking on the item. The following information is provided:

- *Camera parameters*: For each camera, a summary of the camera parameters. This is the same information as displayed in Editing camera parameters. Ensure that the *Focal length* is as expected. All other parameters should have a small magnitude – parameters with a large magnitude indicate a problem.
- *Object point summary*. This is the same information as displayed in View object point summary. Check the relative precision is as expected. The *Number of points not estimated* should be 0. The range between

maximum and minimum point standard deviations should be small. If the *Maximum point standard deviation* is very large, it indicates that some object points have not been measured enough to accurately estimate the coordinates. These points should be identified and more image measurements made to these points.

- *Object point coordinates/precisions*. This is the same information as displayed in Viewing and manipulating object space points. In this window, identify points that have larger than average standard deviations (*sX, sY, sZ*). They will generally have a lower number of measurements (*No. Obs.*) compared to other points. More image measurements need to be made to these points. Make a quick check that the object point coordinates are all in the range expected.
- *Camera orientation and positions (and image observations)*. This is the same information as displayed in View orientation data. In this window examine all camera positions and orientations, and make sure the values are reasonable (look particularly for camera positions that do not fit). Look at the precisions of the camera positions (*sX, sY, sZ*) and orientations (*sOmega, sPhi, sKappa*), and also the number of measurements in each image (*No. obs.)*. These should be of similar magnitude over all images. If one image is particularly worse it may be because it does not have enough object point measurements, or the image quality may be sub-standard. You can also examine the measurements and residuals for a particular image by double clicking on it, although the *Image residual summary* below can be more useful.
- *Image residual summary*. This is the same information as described in Image measurement summary. Study the *average* image residuals and identify any images that have a noticeably higher average than the others. Investigate the measurements and/or image quality associated with these images. Often, an image that has a large average may have misidentified resection points, or measurements made to misidentified targets. In such cases it is prudent to delete all measurements for the image with a large average re-measure it. A large average can also be caused by poor image quality (often motion blur).
- *All image measurements*. This is the same information as displayed in View/delete image measurements. Sort data in this window by clicking twice on the *Rejected* column to bring all rejected measurements to the top. Look to see if the rejected measurements tend to all be associated with one or more object points. If this is the case, investigate all measurements to the particular object point(s) as measurement may have been made to the wrong point.
- *Distance measurements*. This will only be displayed if there are distance measurements. The information is the same as displayed in Viewing and manipulating distances. Check for any distances being excluded (this should not happen with a calibration cube, and indicates the cube or targets thereon may have been physically damaged).
- *View correlations*. Double clicking on this item displays another window allowing parameter correlations to be viewed. Enter the desired *Correlation coefficient*, any parameters above the specified correlation coefficient will be reported. You can choose to include or exclude the camera position and orientation (exterior orientation) parameters as

required. Once configured, click the *Close dialog button*, and a further window is displayed showing any correlated parameters. If the window is empty, there were no parameters above the prescribed limits.

- *Stereo constraints and residuals*. This data will only appear if using stereo constraints in the bundle adjustment (see Stereo constraints). The information and interpretation thereof is the same as in Viewing the constraints. If many of the stereo constraint components are being excluded the constraint weighting may need to be relaxed (see Stereo constraints).
- *Stereo constraints – relative orientation*. This data will only appear if using stereo constraints in the bundle adjustment (see Stereo constraints). The information and interpretation thereof is the same as in Viewing the relative orientation.

If the adjustment was successful, the *Accept* button at the bottom of the window will be enabled. After assessing the adjustment results, and being satisfied the result is valid, click the *Accept* button to adopt the bundle adjustment result. When the *Accept* button is clicked the camera parameters, object point coordinates, and camera positions and orientations are all updated with results computed in the bundle adjustment. If the adjustment has failed, or the adjustment result is not satisfactory, click the *Cancel* button to ignore the adjustment results. NEVER *Accept* the results of a questionable bundle adjustment.

After completing a successful bundle adjustment it is a good idea to delete image measurements that were excluded by the adjustment (see View/delete image measurements).

If there were many object points that were not successfully measured (as will often be the case when there is no previous calibration information for a camera), these points should now be successfully measured automatically. Once the bundle adjustment is successfully computed, there will be better camera calibration parameters, allowing the unmeasured targets to be automatically measured. The procedure for automatic measurement of all images in the project is given in Find targets in all images. After further targets have been measured the bundle adjustment can be recomputed. Note that it will never be possible to measure all targets in all images, as some targets will inevitably be occluded.

Reasons a bundle adjustment might fail

Sometimes the bundle adjustment will fail, and it can be difficult to identify exactly why this is happening. Here are some common reasons and cures:

- The adjustment may not converge because of one or two really bad measurements. If you can force that adjustment to converge once, these bad measurements will be identified and removed.

  Adjust the maximum iterations using Adjustment | Adjustment settings*, and set the* Maximum iterations *to 400*.

- Not enough measurements to enough points. There is sometimes a tendency to compute a bundle adjustment before making enough measurements to enough object points in enough images. When this is the case the bundle adjustment may fail completely, or may fail to converge. You simply need to make more measurements. The bundle adjustment relies on redundancy (many measurements) to reliably detect and eliminate erroneous measurements. If there are not enough measurements the correct detection of erroneous measurements is difficult and can lead to failure of the bundle adjustment.

  Don't bother computing the bundle adjustment until you have measured the entire calibration sequence (see Appendix A Calibration cube orientations).

- An adjustment can fail if measured points have been assigned the incorrect point ID. If a point is consistently measured and given the incorrect point ID in many images this will cause confusion in the bundle adjustment (that is, the bundle adjustment will not be able to accurately identify that the incorrect point has been measured). This can be a real problem if the initial resection points are misidentified, as, in this case, all other automatically measured object points in the image are likely to be incorrectly identified as well.

  Work through all images, and carefully ensure that the resection points are correctly identified in each image. If you find an image with incorrectly measured resection points, it is easiest to delete the image and re-measure it.

- The initial camera parameters are incorrect. The *Image rows* and *Image columns* and the *X Pixel size* and *Y Pixel size* must be correct. The *Focal length* for the camera must also be a good estimate of the true value (remembering that in water the focal length is 4/3 the focal length in air). If there is any doubt about the quality of the camera parameters, set all parameters except the *Image rows, Image columns, X Pixel size, Y Pixel size,* and *Focal length* to 0 (see Editing camera parameters).

- If you are having difficulty with a bundle adjustment, especially if you have accepted the results of a poor adjustment, it is a good idea to re-load the original object point file (a clean version of the file, not one that has been created by the adjustment causing problems, see Reading and writing the object point file).

  After re-loading a 'clean' version of the object point file, use the menu command Measurement | Resect all images*, then compute the bundle adjustment.*

  If this is not successful, re-load a clean version of the object point file using Object points | Read from file*. Delete all measurements, apart from the resection points:* Measurement | View/delete image measurements*,*

sort by the Point ID *column, and delete all point measurements apart from the resection points (on SeaGIS cubes these are points 100, 101, 102, 103 and 104). Resect all images with* Measurement | Resect all images*, then find the targets with* Measurement | Find targets in all images. *Recompute the bundle adjustment.*

- When an adjustment fails to converge, look at the *Image residual summary* report on the *Adjustment results* window. Note the *Valid points* and *Average magnitude*.

  Investigate any image that has a higher than usual Average magnitude *(anything more than three times the average), and any image that has a lower than usual* Valid points *(anything with less than half the targets measured). Generally it is easiest to delete and re-measure these images.*

See also the following section.

Failing adjustments with a SeaGIS calibration cube

Often the reason for a failing adjustment is poor estimates of the camera parameters. When you have good estimates for camera parameters (from previous calibrations) there are generally no problems with computing an adjustment. Where an adjustment is failing with a two-faced calibration cube a useful strategy is to compute the calibration using only the front face of the cube. This will give good estimates of the camera parameters, and a calibration using both cube faces can then be computed.

To get to a point where the bundle adjustment is failing you will have already made a set of measurements on the calibration footage. The instructions below assume you have created a project, and have a complete set of measurements:

- Save a copy of your measurements to date using *Measurement | Write measurement file*. Call the file 'Complete', or similar, to indicate this was your original complete set of measurements.
- Load an original 'clean' copy of the calibration cube (.PtsCAL) file using *Object points | Read from file*.
- Modify the object points to create a cube with just a front face:
  - *Object points | View/edit/delete points*, sort by column 1 (*Point ID*), and delete all points in the 200s (use the *Ctrl* and *Shift* keys to make multiple selections with the mouse, and the *Delete* keyboard key to delete).
  - *Object points | View/edit/delete distances*, delete all distances that involve point IDs in the 200s (use the *Ctrl* and *Shift* keys to make multiple selections with the mouse, and the *Delete* keyboard key to delete). Note that your cube may not have distances on the back face.
- Delete all your images measurements, apart from measurements to the resection points. Use the menu command *Measurement | View/delete image measurements,* sort by the *Point ID* column, and delete all point

measurements apart from the resection points 100, 101, 102, 103 and 104 (use the *Ctrl* and *Shift* keys to make multiple selections with the mouse, and the *Delete* keyboard key to delete).

- Resect all images using *Measurement* | *Resect all images.*
- Find the targets on the front face of the cube with *Measurement* | *Find targets in all images*.
- Compute the bundle adjustment.
- When the adjustment succeeds re-load your original object point file using *Object points* | *Read from file* – don't bother saving the modified front face only object point file.
- Find targets on the back face of the cube with the menu command *Measurement* | *Find targets in all images*.
- Perform another bundle adjustment (now using points measured on the front and back cube faces) and analyse the results per the usual steps.

Centroiding

This section describes the use and configuration of centroiding parameters. Note that the centroid parameters are stored in the registry and automatically remembered by the software (in versions before CAL 2.0, CAL stored centroid parameters in a file). When a new project is created, the current centroid parameter settings are retained (see Creating new projects and data organisation).

Making measurements using centroiding

To measure a point using centroiding, locate the mouse at approximately the centre of the dot to measure, and hold down the *Shift* key. Once the *Shift* key is held an outline square is drawn around the mouse location. This square indicates the region of the image that will be involved in the centroiding process. The dot to be centroided should fit within the outline square. Press the left mouse button to effect the centroid measurement. If the process is not successful an error message box is displayed. You may need to use the error message to fine-tune the centroid parameters (see next section).

Note that centroiding can be used in conjunction with image zooming. If centroiding fails for a point, it can be manually measured by simply using a left mouse click (without holding the *Shift* key).

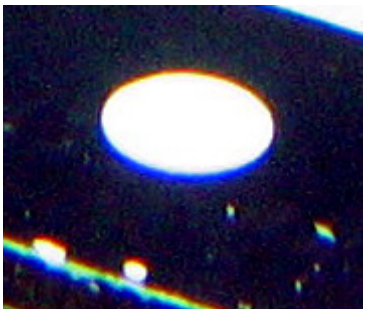

A suitable target for centroiding

Editing and fine tuning settings

To edit the current centroid parameters use *Measurement* | *Edit centroid parameters*. This displays a window allowing centroid settings to be adjusted.

Default settings for centroiding circular targets of high intensity on a dark background are:

Window size	20
Minimum Pixel Intensity Range	30
Minimum Target Size	2
X to Y Ratio	0.15
Invert Pixels	Unchecked
Apply shape test	Checked
Display dynamic dialog	Checked
Threshold Algorithm	Random
Target Shape	Circle
Data channel	Greyscale

With these default values you should only need to alter the *Window size* to suit your particular image resolution and target size.

Window size. If the target to be centroided does not fit within the centroiding window (the outline square drawn when the *Shift* key is held down), the *Window size* may need to be increased. When this is the case you will get a centroiding error message 'Image larger than window'. This parameter is the most likely to need adjustment.

Minimum Pixel Intensity Range. This value sets the minimum difference required between the target intensity and the background intensity. If this setting is too small you will get a centroiding error message 'Intensity range too low'. If you make this setting too small spurious targets will be detected.

Minimum target size. This is the minimum required size of a target to be centroided.

X to Y Ratio. This setting is only applied if the *Apply shape test* option is checked. The *X to Y Ratio* places a limit on the 'roundness' of the target, and is useful to exclude targets that are too oblique, or in reducing the detection of targets that are not proper circular targets. When a target fails this setting the centroiding error message 'Image shape error' is given. If the centroiding window is open, the window will display the calculated ratio of the target.

Invert Pixels. Check this option if the targets for centroiding are dark targets on a light background.

Apply shape test. Check this option to apply a geometric shape test to centroided targets. If a target fails this test the centroiding error message 'Image shape error' is given.

Display dynamic dialog. Check this option to display a window that shows the details of target centroiding (uncheck to hide the window). The displayed window can be resized and repositioned to suit, and the application will remember the last size and position of this window and re-create in the same place. The window shows the chosen target area from the source image (the left or upper image), the target after thresholding and noise removal (the right or lower image), and any error messages about the centroid. If the centroid is successfully computed, the centroid coordinates are displayed. This window is a useful diagnostic tool when configuring centroiding settings. Note that the window will not be displayed or hidden until the configuration window is closed.

Threshold Algorithm. Sets the threshold algorithm used when centroiding. The *Random* algorithm is recommended.

Target Shape. Use *Circle*.

Data channel. This setting specifies the image colour channel to use. In general is should be set to greyscale for white/black targets. If you are using targets of a specific colour (particularly red or green or blue), set the *Data channel* appropriately.

Stereo camera calibration quick guide

Use this guide to step through the process of calibrating stereo cameras. Once you have captured and downloaded imagery of a calibration cube, follow these steps:

Create a new project

1. Using Windows File Explorer, create a new directory for the project.
2. Put all image files associated with the calibration into this new directory.
3. Copy the object point file (.PtsCAL) for the calibration cube into this directory.
4. If you have previously calibrated the camera system, also copy the camera files (.CamCAL) from a previous calibration into this new directory. Do not use a previous in air calibration result for an in water calibration or vice versa, as the parameters are significantly different. There should be one .CamCAL file for the left camera, and one for the right.
5. In CAL create a new project using *Project* | *New project*. Enter a name for the project, and save the project file in the new directory you have just created for this project (see Creating new projects and data organisation).

Configure the cameras

1. If you have camera files from a previous calibration, load these files for the left and right cameras using the menu commands *Camera* | *Left* | *Read from file* and *Camera* | *Right* | *Read from file* (see Reading and writing to disk). Once you have loaded a file for each camera, you have finished with this section.
2. If you do not have existing camera files, you will need to set some camera parameters manually. For each camera you must manually enter:

   Image rows, Image columns, X Pixel size, Y Pixel size, and the *Focal length*. All other parameters are set to 0. You should also enter a *Descriptor* for each camera.

   See Editing camera parameters for notes on how to manually enter parameters, and a description of what these parameters are. If the cameras are being calibrated in water, their focal length will be approximately 4/3 of the focal length in air (a 4.8mm focal length lens will be approximately 6.4mm in water). Use this adjustment to give a good estimate of the focal length.

   You will also need to fix certain camera parameters (see Fixing camera parameters). The parameters *Radial distortion (K7)* and *Orthogonality* should be fixed, all other camera parameters should be free.

   Save the camera files once you have finished this basic configuration (see Reading and writing to disk).

Load other files

1. Load the calibration cube's object point file. Use the menu command *Object points* | *Read from file* (see Reading and writing the object point file). If you are using a calibration cube, always start the project with a copy of the original cube .PtsCAL file.
2. Set the directory where the images for this project are located. To do this, use the menu command *Picture* | *Set picture directory* (see Setting the picture directory).
3. Create a new measurement file. Although you have not yet made any measurements, creating this file will save the picture directory you have just set, and finish preparation of the project. Use the menu command *Measurement* | *Write measurement file.* This displays a Windows file dialog allowing you to create the file. Name the file, and save it in the directory you have created for this project. See Saving and reading measurements.
4. Before proceeding, check the current project status using *Project* | *Display/set current project data.* All items listed should have a filename, and generally all files should come from the same directory. If any other data is missing check all steps previous to this (see also Viewing/setting current project settings).

Measurement

When performing a stereo camera calibration, it is important that the images you measure for the calibration are a stereo pair. This is particularly important where the imagery is in a movie file. Make sure the frames selected for measurement from the movie files are stereo pairs. If the movie frames used in the camera calibration are not stereo pairs, the second part of the calibration that determines the relative orientation of the cameras will fail. With movie files, use the movie player to help to quickly synchronise the frames, then *Lock* the movies (see The 'Lock' function).

Repeat this procedure for each of the calibration images. There should generally be 20 images for each camera (see Appendix A Calibration cube orientations):

1. Load a new image that has not been measured yet:
    i. To load an image use *Picture* | *Left* | *Open picture* and *Picture* | *Right* | *Open picture*.
    ii. If you load a movie that is not part of the movie sequence definition, it is automatically added to the left or right movie sequence definition, and the movie sequence definition window is displayed. If there are multiple movies in the sequence make sure they are correctly ordered and click *OK*. See Movie sequences.
    iii. For still images, or use the  buttons to advance through still pictures that have filenames in a numeric sequence.
    iv. For movies, once the movie is loaded, use the  buttons or keyboard shortcuts to advance frames in a movie file or use the *Play movie* button to play the movie to an appropriate frame.

Ensure that the images are correctly synchronised, and make use of the *Lock* check box (see The 'Lock' function, Stepping through movie and still images, The movie player).

2. Measure the image using centroided measurement (with the *Shift* key held down) to identify the resection points. The image should be automatically resected, and other points measured automatically. After one or more stereo pairs have been measured, the corresponding stereo image will be measured automatically (see Point measurement, Using automatic stereo resection and measurement). Note that when calibrating cameras that have no previous calibration data, not all of the image points will be automatically found. If the estimated camera focal length is good there will usually be enough points identified to compute a bundle adjustment. After the initial bundle adjustment the object points that were initially not found can be able to be automatically measured.

Adjustment

1. Check the adjustment parameters are correctly set (see Adjustment settings).
2. Compute the bundle adjustment using *Adjustment* | *Compute bundle adjustment*. Check the bundle adjustment results (see Computing the adjustment, it is important to thoroughly review all items listed, and follow the procedure given). If the bundle adjustment is failing, read Reasons a bundle adjustment might fail and Failing adjustments with a SeaGIS calibration cube. Once an adjustment completes successfully, follow this procedure:
    i. Accept the adjustment results (using the *Accept results* button).
    ii. Delete the rejected measurements using *Measurement* | *Delete rejected point measurements.*
    iii. Search for missing targets using *Measurement* | *Find targets in all images.*
    iv. Compute another bundle adjustment and *Accept results*.
    v. Delete the rejected measurements using *Measurement* | *Delete rejected point measurements.*
    vi. Compute a final bundle adjustment and *Accept results.*
3. Note the bundle adjustment needs to be of acceptable quality before proceeding.

Generating stereo camera files

Note, all menu operations mentioned in this section are under *Measurement* | *Stereo constraints.*

1. Configure the stereo constraints using *Configure stereo constraints.* On this dialog it is usually adequate to make sure the *Left camera* and *Right camera* are properly selected, and then click the *Automatic* button to generate the left/right pairings. Check the left/right image pairings are correct (see Configuring constraints).
2. Estimate the constraints using *Estimate constraints* (see Estimating constraints).

3. Check the relative orientation using *View stereo constraints*. Check to see constraint residuals are acceptable, and not too many constraints have been excluded from the calculation (*Exclusions* column). If more than around 4 constraints are rejected re-check the stereo synchronisation. For more details, see Viewing the constraints.
4. Generate the stereo camera files using *Export stereo camera files* (see Exporting stereo camera files).

Saving results

The simplest way to save the project and all the associated data is to close CAL, and the program will prompt you to save all unsaved data. If you are using a SeaGIS calibration cube, you don't need to save the object point (.PtsCAL) file.

Shortcut keys

Movie player

Key(s)	**Action**
Space bar	Play/pause playback
Left arrow	Step back 1 frame, puts movie in pause mode if playing
Ctrl + Left arrow	Step back 5 frames, puts movie in pause mode if playing
Shift + Left arrow	Step back 25 frames, puts movie in pause mode if playing
Right arrow	Step forward 1 frame, puts movie in pause mode if playing
Ctrl + Right arrow	Step forward 5 frames, puts movie in pause mode if playing
Shift + Right arrow	Step forward 25 frames, puts movie in pause mode if playing
Ctrl + mouse click on frame stepping controls	Step back/forward 5 frames, puts movie in pause mode if playing
Shift + mouse click on frame stepping controls	Step back/forward 25 frames, puts movie in pause mode if playing
Up arrow	Decrease playback rate
Down arrow	Increase playback rate
Escape	Close player, don't update movie position in main window
Enter	Close player, update movie position in main window

Main window

Key(s)	Action
Space bar	Launch movie player (for the left movie)
Left arrow	Step back 1 frame (both left and right movies)
Ctrl + Left arrow	Step back 5 frames (both left and right movies)
Shift + Left arrow	Step back 25 frames (both left and right movies)
Right arrow	Step forward 1 frame (both left and right movies)
Ctrl + Right arrow	Step forward 5 frames (both left and right movies)
Shift + Right arrow	Step forward 25 frames (both left and right movies)
Ctrl + mouse click on frame stepping controls	Step back/forward 5 frames
Shift + mouse click on frame stepping controls	Step back/forward 25 frames
Ctrl + d	Deletes all rejected point measurements in the currently displayed images

Release information

Version 1.11

New features:

- Installable component (see Components) to convert movie measurement files to measurement files based on JPEG images. Once measurement of movie files is finished, frames from the movie files that were used in measurements are automatically extracted and saved as JPEG images, and a new measurement file referencing the JPEG images is created. See Convert movie to JPEG.

- Export a stereo camera file for any orientation (exterior orientation for a particular image). See Exporting a stereo camera file for an image.

- Ability to use coordinates as constraints in a bundle adjustment.

- An option to apply brightness and contrast settings to all images. See Viewing/setting current project settings and Brightness and contrast.

Improved functionality:

- The image rows and columns as defined in the camera settings are used to check the currently loaded image dimensions, with a warning issued if there is a mismatch. Where movie files are used, this check is performed every time a different frame from the movie is displayed (see Loading images).

- Before performing any calculations involving the cameras, the camera parameters are checked. If there is not a valid x and y pixel size, focal length, and image rows and columns definition, a warning is issued (see Editing camera parameters).

- Improved licence checking to avoid the need to manually order computer adapters.

Version 1.12

Improved functionality:

- Software licence details are now stored on a machine basis rather than a user basis. Allows all machine users to access the software without entering licence details.

- New function that allows the current image in a camera window to be saved as a JPEG image.

Version 1.20

New features:

- Added HASP USB licence keys.
- Visual Studio 2005 build.

Version 1.21

New features:

- Software uses the camera x and y pixel size (see Editing camera parameters) to correctly draw imagery with 16:9 aspect ratio (or any other aspect ratio).
- Added support for mpeg movie formats. To use mpeg formats an appropriate frame accurate mpeg codec must be installed (see Loading images).

Improved functionality:

- Increased allowed range for stereo constraint rejections (see Stereo constraints).
- Fixed problem to allow correct loading of AVIs using the Cineform codec.
- Limited the range of centroiding X to Y ratio to be a range of 0 to 1 (see Editing and fine tuning settings).
- Changed bitmap drawing mode to halftone to improve drawing of interlaced sources.

Version 1.22

- Added support for exporting comma separated value (CSV) data files (see Generic editor/display window).
- Fixed a bug that caused a crash if 'View correlations' was selected and the adjustment had not completed successfully (see Computing the adjustment).
- Added automatic backup of observation (measurement) data (see Saving and reading measurements).
- Built with Unicode libraries.

- Removed the ‘standard’ brightness and contrast mode (see Brightness and contrast).

Version 1.30

- Found and fixed a bug where measurements could be recorded without being associated with a valid image filename.
- Upgraded USB licence library to HASP 3.5.
- Visual Studio 2008 build.
- Added debug window for movie file loading (activated with registry flag).
- Altered default centroiding parameters (see Editing and fine tuning settings). Added logic for automatically setting the centroid window size.
- Changed default maximum iterations for bundle adjustment.

Version 1.31

- Fixed bug introduced in version 1.30. Some movie files will not load because of slight variations in reported frame rate.

Version 1.32

- Changed default base rejection for stereo constraints to 5mm (see Stereo constraints).
- Added registry setting to force movie handler to DirectShow or Video for Windows.

Version 1.33

- Changed HASP runtime installer to use command line utility.
- Small interface changes to allow the user to cancel adding a new point or a new distance, and to cancel movie to JPEG generation.
- Improved the speed of writing measurement files over network connections.
- When loading images in the camera windows, the file dialog defaults to the current picture directory.

Version 2.00

- Major interface change, CAL now runs in a single window, is specifically set up to process stereo calibrations, and looks and feels more like EventMeasure. The easiest way to learn the new interface is to process a calibration using the updated quick guide (see Stereo camera calibration quick guide).

- The software is fixed for stereo calibration, there is no longer any need to set the number of cameras.

- Added a movie player and *Lock* function to enable easier movie manipulation and help maintain synchronisation (see The movie player, The 'Lock' function). Because of the addition of a movie player, the frame stepping acceleration has been removed, and control of frame stepping is done with the text box located next to *Lock* check box (see Stepping through movie and still images).

- Setting the picture directory is now done through the *Picture* menu (*Picture* | *Set picture directory*). This used to be done through the *Project* menu. See Setting the picture directory.

- Functions for loading pictures have moved into the *Picture* menu (see Loading images). This menu also has commands for configuring movie sequence definitions (see Movie sequences). Movie sequence should be defined whenever you use movie files.

- Brightness and contrast adjustment is now located under the *Picture* menu. There is one window that simultaneously controls the brightness and contrast settings of the left and right image (see Brightness and contrast).

- Some menu items that were not used in stereo camera calibrations have been hidden for simplicity.

- The 'observation and orientation' file is now just referred to as the 'measurement' file, for example *Measurement* | *Read measurement file*, *Measurement* | *Write measurement file*.

- The *Show measurements* check box, used to toggle the display of measurements, has been removed.

- Centroid parameters are now stored by the application (in the registry), rather than in a file (see Centroiding). As such, no centroid filename is reported in the project settings (see Viewing/setting current project settings).

- Functions to jump to a specific frame and to save a jpeg image of the current movie frame are now accessed via a pop-up menu. These functions are only available with movie imagery. To access the menu, right click in either the left or right image (see Saving the current image as a JPEG, Jump to frame). Note that the jump to frame function will use the *Lock* function (see The

'Lock' function), if checked, to jump in stereo. The jump to frame function now supports jumping to a time as well as jumping to a frame number.

- Updated the tips and tricks to deal with a failing bundle adjustment (see Reasons a bundle adjustment might fail, Failing adjustments with a SeaGIS calibration cube).

- Made installation of the HASP runtime environment an installer option.

Version 2.01

- Added keyboard shortcuts for frame stepping (see Stepping through movie and still images).

- The movie player can now be launched by pressing the keyboard space bar, also there are several keyboard shortcuts for the movie player window (see The movie player).

- Corrected units displayed against camera distortion model parameters.

- Expanded information and descriptions for capturing calibration footage (see Appendix A Calibration cube orientations).

Version 2.02

- Added automated measurement. After one or more stereo pairs have been measured, CAL will try to automatically complete the measurement of the second image of a stereo pair once the first image (left or right) has been measured (see Using automatic stereo resection and measurement, Stereo constraints).

- Improved the response time when deleting data from the data editor window – for example, deleting rejected image measurements.

- Removed the warning message displayed when exporting the left and right stereo camera files (see Exporting stereo camera files).

Version 2.03

- Changed the label 'Principal distance' to 'Focal length' – physically the same thing, just a more user-friendly name.

Version 2.04

- Added table of shortcut keys (see Shortcut keys).

Version 2.10

- Installer now uses HASP version 6.51 (Windows 8 compatible).
- The *Ctrl* and *Shift* keyboard keys work with the mouse activated frame stepping controls (see Stepping through movie and still images).
- The *Ctrl* and *Shift* keyboard keys work with the mouse activated frame stepping controls in the movie player (see The movie player).

Version 2.11

- Added a link to the user manual from within the software: *About* | *User manual*. See Using this manual, Documentation.
- Added an option to uninstall the software to the 'Start' menu.
- Added *Ctrl* + *d* shortcut to delete rejected point measurements in the current images (see Main window shortcuts).
- Added menu function *Measurement* | *Delete rejected point measurements*. This is a quick method to remove rejected points after a bundle adjustment (rather than the previous method of selecting and deleting the rejected points manually). Updated the calibration quick guide accordingly. See Delete rejected point measurements, Adjustment.

Version 2.20

- Updated Appendix A Calibration cube orientations to include a note about adding extra calibration images for 1920x1080 (16:9) format images.
- Upgraded HASP licence key driver to version 7.40.
- Movie files are now added automatically to the movie sequence definition when they are loaded. See Loading images.
- Changed default bundle adjustment maximum iterations to 200 (was previously 80).
- Increased the maximum playback rate in the movie player.

Version 2.30

- Moved to Windows SDK version 7.1, target platform Windows 7.
- Removed zoom '1' setting (caused confusion when zoom appeared not to work if '1' was selected). Added additional zoom levels. See Zooming.
- Added Media Foundation implementation for movie handling and the movie player window (ability to handle MP4 files).

Appendix A Calibration cube orientations

The following details the steps required for capturing imagery suitable for performing a camera calibration. The instructions cover several different hardware variations that result in slightly different approaches to the calibration. However, regardless of the hardware configuration used, the aim is to collect a sequence of 20 images of the calibration hardware (generally a calibration cube). Stereo camera calibrations should conclude with recording of an independent length or scale bar in order to verify the calibration results.

Preliminary

- Check that camera settings are correct (you will generally have a SeaGIS document that summarises how to configure the cameras).
- Ensure that cameras are mounted securely in the housings, and all other housing fasteners are secured (check other SeaGIS camera and housing documentation).
- Make sure the camera batteries are charged and there is sufficient recording memory available.

Setup

There are several different 'styles' of calibration, using slightly different hardware configurations. These are described below.

Using a l000x1000x500mm cube

- These calibrations are best done with the camera base bar resting on a plastic bin (or similar) that is approximately 500mm tall, with the calibration cube being rotated.
- Rest the cube on rubber matting to help define the location of the cube, and protect the cube and pool.

Using a 500x500x300mm cube

- These calibrations are generally performed with the camera system being held still, and the calibration cube being rotated. The camera system can be placed on the bottom of the pool, and the cube rotated in front of the camera system.

For 'whale shark' style calibrations

- These calibrations are performed using a portable target array that is configured on the pool bottom for each calibration. During the calibration the target array stays still on the pool bottom, and the camera system is moved around and rotated to gather the required imagery.

Recording imagery

- Position camera(s) relative to the cube so that the cube fills a large portion of the field of view. However, leave enough margin so that you *do not cut off parts of the cube*. If you are performing a stereo calibration, check the positioning of both cameras.
- Consider the position of the sun and/or other lighting sources to avoid excessive glare.
- Make sure camera(s) are recording.
- If you are performing a stereo calibration, make sure the cameras are synchronised and/or turn on the synchronisation diode.
- Move the cube (or camera(s), depending on the style of calibration) to achieve the required sequence of 20 different calibration images. The following provide a verbal description, diagram, and sequence of 'real' calibration images.
- When the cube is turned away from parallel to the camera system, the angle should be approximately 20 degrees.
- With wide format footage (like 1920 x 1080 or 16:9 video) the cube will not fill the left and right edges of the images. To improve the calibration modelling at the image edges you can collect some extra images of the cube where the cube is positioned towards the left and right edges of the images. In CAL, just measure these extra images in addition to the standard twenty.

NOTE

- The order in which the sequence of images is collected does not matter.
- At each different cube position hold the cube and camera system as still as possible whilst capturing the imagery (this is true for either still or video systems) – for video systems capture about 1 second worth of imagery in this still state. Capturing stable imagery will produce the best quality imagery to measure from, as well as reducing the influence of any camera synchronisation error.
- If the camera system is used to make measurements in water, the calibration should be done in water.

Recording imagery – description

- Take 4 images with the calibration cube square on to the camera(s), rotating the cube through 90 degrees for each image.

- Turn the left face of the cube towards the camera(s) by about 20 degrees and take a series of 4 images, rotating the cube through 90 degrees each time.

- Turn the top face of the cube towards the camera(s) by about 20 degrees and take a series of 4 images, rotating the cube through 90 degrees each time.

- Turn the right face of the cube towards the camera(s) by about 20 degrees and take a series of 4 images, rotating the cube through 90 degrees each time.

- Turn the bottom face of the cube towards the camera(s) by about 20 degrees and take a series of 4 images, rotating the cube through 90 degrees each time.

Recording imagery – diagram

The ‘bolded’ side is provided for rotational reference.

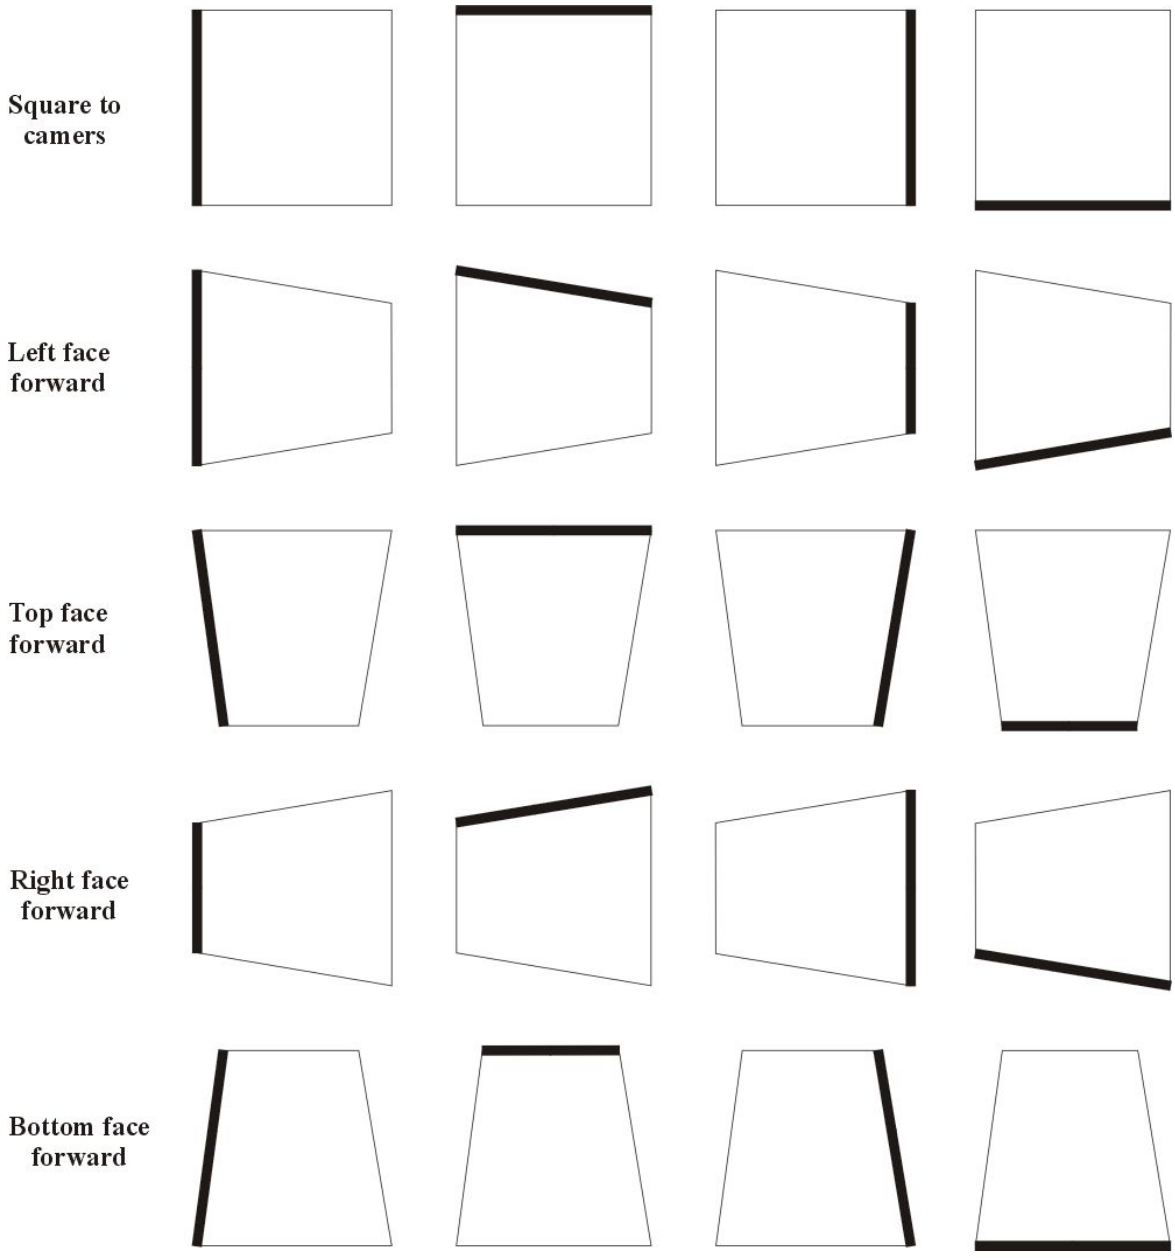

Schematic calibration sequence

Recording imagery – picture sequence

Note the position of the reflective ‘bar’ on the cube face for rotational reference.

Square to cameras

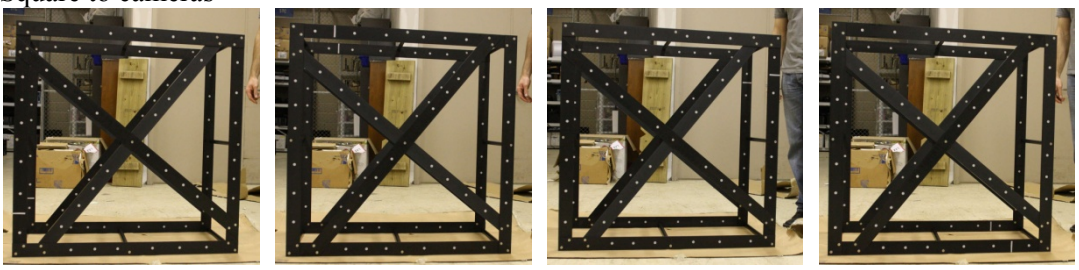

Left face forward

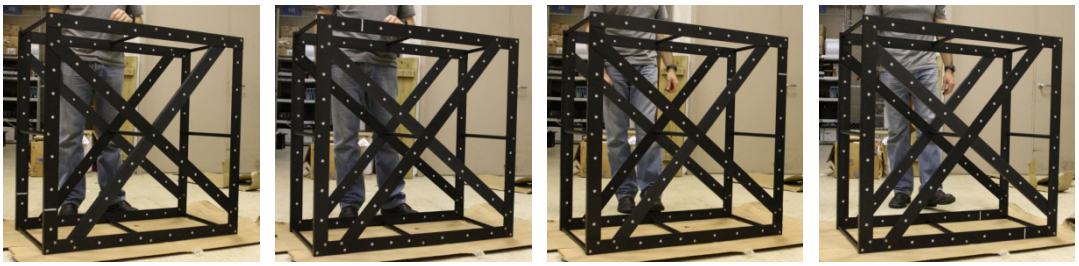

Top face forward

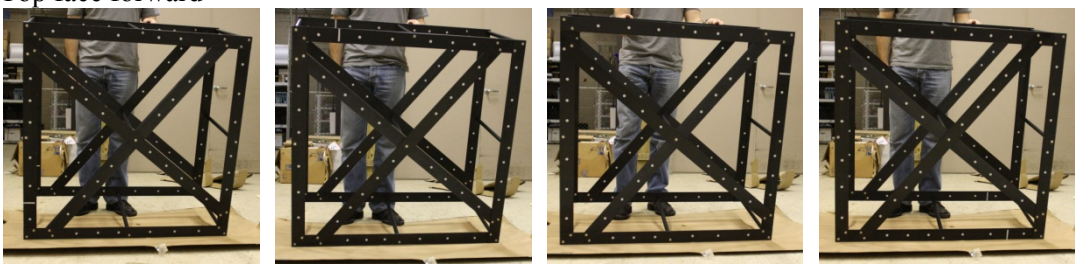

Right face forward

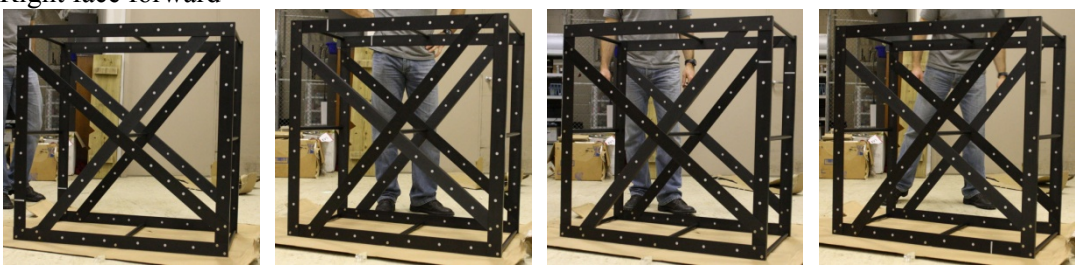

Bottom face forward

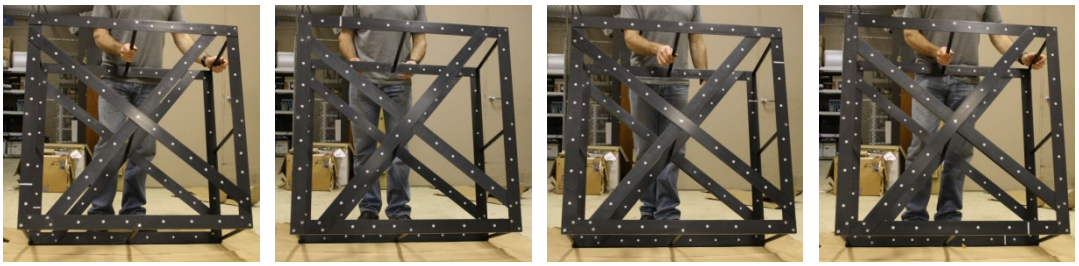

Pictorial calibration sequence

Calibration check

For stereo camera system calibrations, after capturing the calibration sequence, capture imagery of a scale bar/test object at a variety of ranges and angles relative to the camera system. Once the calibration has been processed in CAL, load the newly generated stereo camera files (.Cam files) into your measurement software (usually either PhotoMeasure or EventMeasure), and verify the calibration by making measurements of the scale bar/test object.