

Appendix A: Detailed Field Procedures

Camera Calibration Considerations

Over the course of generating camera-lens calibration files for this project and other research, it was found that the Canon 7D (crop sensor) returned more accurate calibrations with the same lens, compared to the Canon 5D Mk II. This is interpreted to be related to sensor dimensions. Because the lenses used were designed for full frame sensors, they are expected to cast the same size circle of light. However, the crop sensor, being smaller is intercepting only the fraction of the circle which is closest to the centre, where distortion is smallest.

This is corroborated by calibrations conducted using the same camera and lens, but different apertures. In these cases, wider apertures (using more of the outsides of the lens) resulted in decreased accuracy. It should be noted that this change was often only on the order of 0.01 of a pixel, which is largely insignificant.

The Adam Technology “Object Distance” spreadsheet provides relations between the performance of crop sensor cameras and full frame sensors. One such relation is the equivalent full frame aperture, which is higher than the set aperture. This increase is taken to be a result of the decreased pixel dimensions, with allow for diffraction (a problem at tight apertures where light beams fall over multiple pixels) to occur at wider apertures. This is partly why the camera was set to f7.1 at Doe Run.

Use of Teleconverters

Teleconverters, which were not used directly in this study, but were employed in related surface work, were found to reduce calibration accuracy. This is partly related to the effect of the teleconverter on effective aperture. The teleconverter employed was found to halve the light transmitted to the sensor, such that an aperture setting of f6 on the camera is interpreted to correspond to an aperture setting of f12 on the actual lens. This was investigated by calibrating a 200 mm prime lens with the teleconverter at multiple apertures. In this case an aperture of f6 or f8, which was found to provide accurate calibrations without the teleconverter, returned much less accurate calibrations with the teleconverter attached. However, when the aperture was increased to f11 with the teleconverter, the accuracy was increased, although not to the level achieved

without. When aperture was further increased to f13, the accuracy was found to decrease again.

Calibration Project Accuracy and Evidence of Errors

When examining the results of a camera-lens calibration, an accuracy of between 0.2 and 0.3 pixels should be achievable. Several calibrations with the Canon 7D achieved accuracies on the order of 0.1 pixels. In addition to calibration project accuracy, projects which use the calibration file can be checked for errors by examining the direction of the point residuals. 3DM Calibcam produces scaled vectors for each point representing the x-y error associated with calculated and observed co-ordinates of the point. For a well calibrated project, these vectors should be randomly distributed. Projects with calibration file issues will often exhibit swirling or other patterns in the point residuals.

Sub-Hyperfocal Distance Calibration

As discussed, object distance specific calibration files are required for photos taken at less than hyperfocal distance of the camera and lens. When setting up calibration projects at sub-hyperfocal distance, the same techniques apply with minor adjustments. A pyramid camera setup is still employed, but the maximum and minimum distances from the face are set based on the depth of focus of the camera-lens setup. The following steps outline the process:

1. Review the Object Distance spreadsheet to determine the depth of focus for the camera-lens setup at the chosen aperture.
2. Confirm that the depth of focus covers the entire area of interest and adjust aperture as needed.
3. Note the extents of the depth of focus.
4. Layout the calibration pyramid with the middle row at the planned object distance for the actual project.
5. Set the top and bottom distances as close to the maximum and minimum extents of the depth of focus as possible, while allowing for the depth of the calibration object.
6. Set the pyramid width to the same as the distance between the front and back of the pyramid.

Figure A-1 summarizes the recommended pyramid setup.

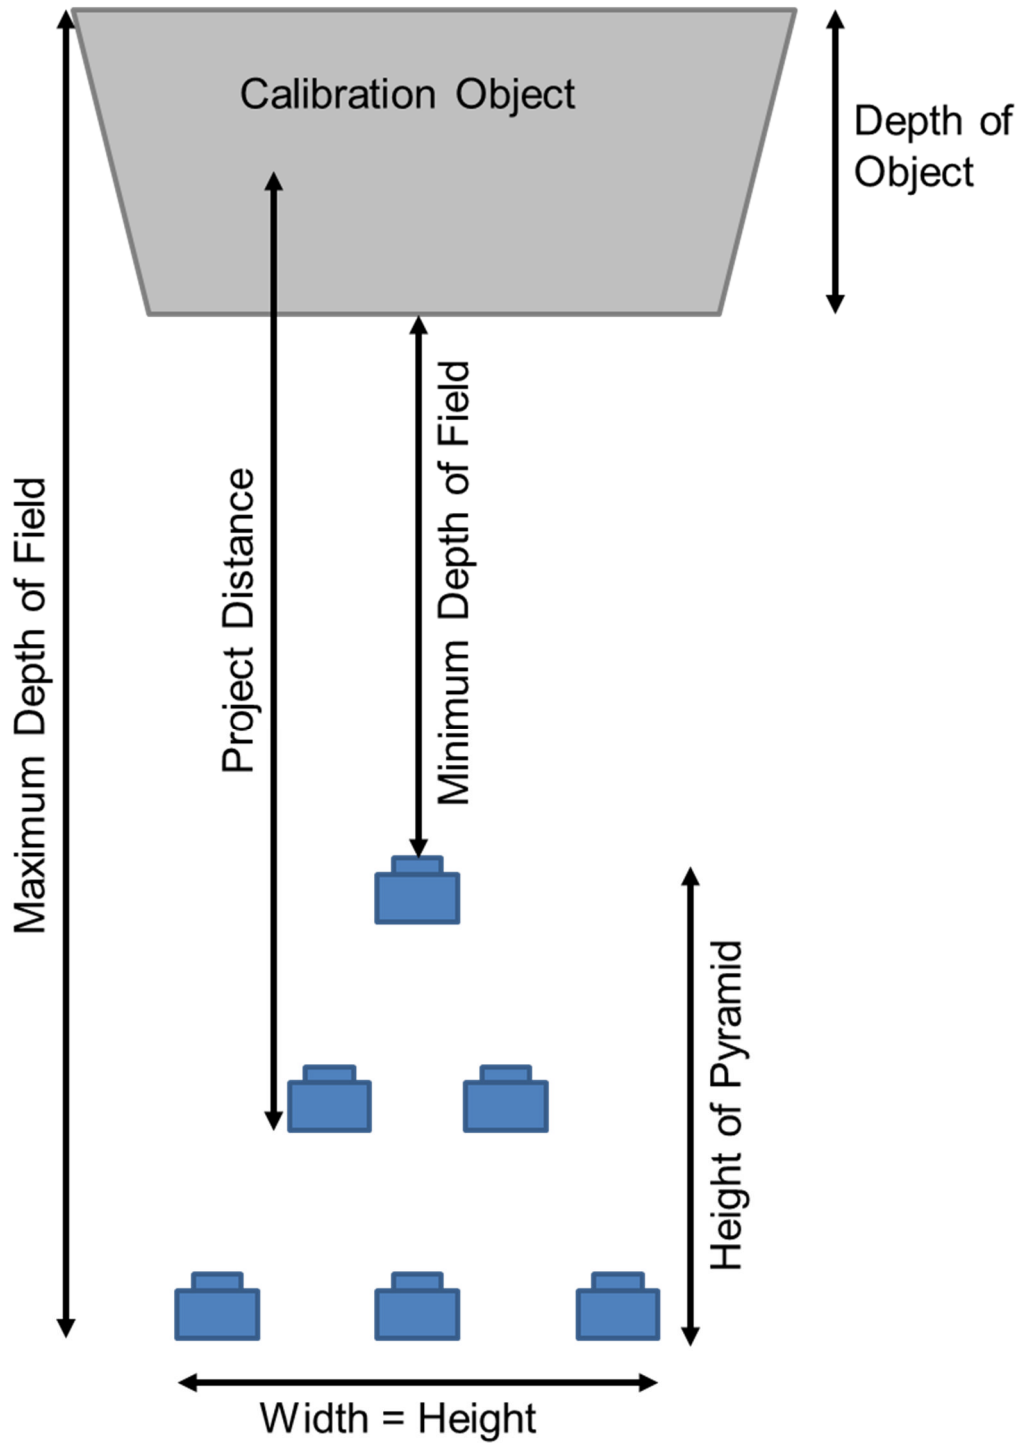


Figure A-1: Recommended Pyramid for Sub-Hyperfocal Calibration

This process can be assisted in the field by downloading one of the available applications for smart phones which calculate depth of field and hyperfocal distance.

Model Generation Considerations

Early experiments were conducted to determine the DTM generator settings which would return the highest possible density of point cloud while maintaining quality. The recommended settings are listed in Table A-1.

Table A-1: DTM Generation Settings used for Maximum Point Density

Setting	Value
Match similarity	0.85
Operator WinSize	19
Grid Size (Pixel)	10
Seed Size (Pixel)	64
Density (Pixel)	4
Feature Rate Left	10
Feature Rate Right	10

The Fit-Existing bundle adjustment function was not necessary for this project due to the presence of primary and secondary control points. However, on large surface project, the Fit Existing function is invaluable. Photos taken from the same or similar viewpoints can be automatically matched between observations. Once matching is complete, the second observation photos are disabled and bundle adjustments are conducted on the first observation photos. Then the second observation photos are enabled and the Fit Existing function is used. This function performs another bundle adjustment, but does not alter any of the points from images which the user has fixed (the first observation in this case). This tool can reduce the need to repeated control surveys as well as reduce cumulative error from minor errors in the camera-lens calibration (discussed below). When using the tool, the stereopair which is checked remains stationary and the unchecked photos are aligned. It appears that a control network adjustment should be conducted on the stable images directly before running the fit existing adjustment. Subsequent fit existing adjustments appear to affect both stereopairs regardless of checked status.

Birch (2006) observed that minor errors in the scale of a photomodel project can result from slightly inaccurate camera calibration. These errors become evident upon review of the control point bundle adjustment residuals. Plotting the magnitude and vector of these residuals will highlight any scaling issues. These errors can be corrected for by allowing the software to adjust the camera co-ordinates. The software will then calculate new co-ordinates for the cameras, either further or closer to the object, which produce the correct object size. This is only an issue when camera co-ordinates are used as part of the control network. These errors have been observed on the order of 0.1% which is insignificant for local work, but can result in significant issues at an open pit scale, especially when monitoring minor displacements is the goal. Projects should therefore be setup such that the control network on, or around the object is sufficiently robust that the camera co-ordinates can be used as a check on the calibration accuracy. Use of the fit existing function can partly correct for these errors between observations reducing the overall error to that of the reference photoset.