

CAMERA CALIBRATION USING THE DAMPED BUNDLE ADJUSTMENT TOOLBOX

Niclas Börnin^{*,a} and Pierre Grussenmeyer^b

^a Department of Computing Science, Umeå University, Sweden,
niclas.borlin@cs.umu.se

^b ICube Laboratory UMR 7357, Photogrammetry and Geomatics Group, INSA Strasbourg, France,
pierre.grussenmeyer@insa-strasbourg.fr

WG V/1

KEY WORDS: Camera calibration; Bundle adjustment; Focal length; Convergent; Initial values; Close Range

ABSTRACT:

Camera calibration is one of the fundamental photogrammetric tasks. The standard procedure is to apply an iterative adjustment to measurements of known control points. The iterative adjustment needs initial values of internal and external parameters. In this paper we investigate a procedure where only one parameter — the focal length is given a specific initial value. The procedure is validated using the freely available Damped Bundle Adjustment Toolbox on five calibration data sets using varying narrow- and wide-angle lenses.

The results show that the Gauss-Newton-Armijo and Levenberg-Marquardt-Powell bundle adjustment methods implemented in the toolbox converge even if the initial values of the focal length are between 1/2 and 32 times the true focal length, even if the parameters are highly correlated. Standard statistical analysis methods in the toolbox enable manual selection of the lens distortion parameters to estimate, something not available in other camera calibration toolboxes.

A standardised camera calibration procedure that does not require any information about the camera sensor or focal length is suggested based on the convergence results.

The toolbox source and data sets used in this paper are available from the authors.

1. INTRODUCTION

Camera calibration is one of the fundamental photogrammetric tasks. The standard procedure is to apply an iterative adjustment to measurements of known control points. The iterative adjustment needs initial values of the internal and external orientation, and optionally any object points. The theory and formulation of radial and de-centring lens distortions from (Brown, 1971) have been adopted in close range photogrammetry as well as in e.g. the popular free Computer Vision Camera Calibration Toolbox for Matlab¹ and in the commercial Computer Vision System Toolbox for Matlab² based on the work by Zhang (2000). The aid of coded targets is now the norm for off-line automatic camera calibration (Fraser, 2013). Automatic camera calibration using self calibration is a standard in most of the image based dense matching tools available (Snively et al., 2008; Furukawa and Ponce, 2009), whatever the kind of camera used, but unfortunately detailed quality results from the adjustment are usually not available for the end users.

In this paper we focus on the progress of the bundle adjustment during the calibration process. Initial values for the internal parameters can be obtained from several sources, including the camera manufacturer, EXIF information in JPG images, or prior knowledge. We investigate a procedure where the focal length is the only parameter that is given a specific initial value; the other parameters are given standard initial values or computed from the data. Our target is to investigate the pull-in range of this procedure using the set of bundle adjustment techniques available in the free Damped Bundle Adjustment Toolbox for Matlab, described in Börnin and Grussenmeyer (2013a), and to suggest a

standardised procedure that does not require any camera knowledge.

2. CAMERA CALIBRATION

If the lens distortion (LD) parameters are included, estimating the internal camera parameters is a non-linear problem, usually performed by bundle adjustment (Luhmann et al., 2006, Ch. 7.2). After the images have been acquired and the measurements have been performed, initial values of the parameters to be estimated are given to a bundle adjustment (BA) algorithm which is allowed to iterate “until convergence”. In addition to the internal parameters (internal orientation — IO), the parameters to estimate usually include the camera positions and orientations (external orientation — EO) and possibly also object point (OP) coordinates.

2.1 Main algorithm

The following high-level algorithm was used for the camera calibration:

1. Decide which of the radial K_1 – K_3 and tangential P_1 – P_2 LD parameters that should be included in the calibration.
2. Decide the initial estimate f_0 of the focal length. This will be discussed in detail in Section 3..
3. Use standard values for the other IO parameters:
 - (a) The principal point was set to the centre of the sensor.
 - (b) All lens distortion parameters were set to zero.
 - (c) The pixel aspect ratio was assumed to be unity.
 - (d) The skew was assumed to be zero.

¹http://www.vision.caltech.edu/bouguetj/calib_doc

²<http://www.mathworks.com/products/computer-vision>

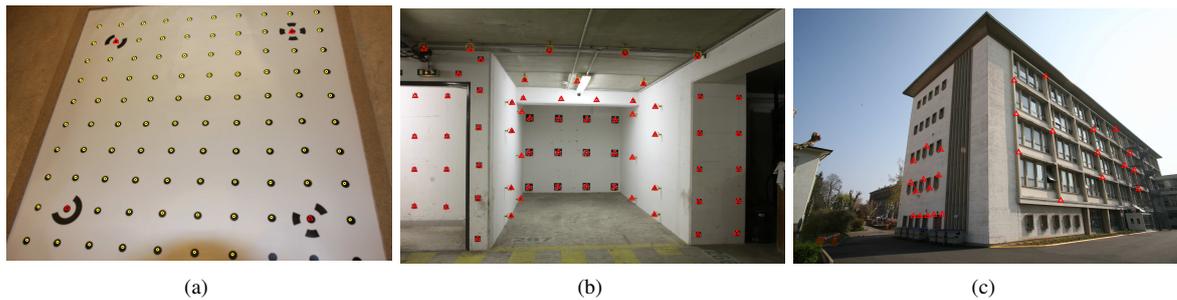


Figure 1: (a) The calibration target used in data set 2. The same type of target was used in data sets 1 and 3. (b) The calibration test field used in data set 4. (c) Part of the control point network on the INSA building used in data set 5.

4. Based on the assumed IO parameters, determine initial EO values:
 - (a) Use the 3-point spatial resection algorithm (McGlone et al., 2004, Ch. 11.1.3.4) on the available control points (CPs).
 - (b) With more than 3 CPs available, choose the 3 triplets covering the largest measured image area.
 - (c) As the initial values, choose the resection that produces the smallest re-projection error of all measured CPs.
5. Based on the assumed IO and EO parameters, compute the position of any non-control OP with forward intersection.
6. Fine-tune the initial values with a BA algorithm.
7. Analyse the resulting parameter values for high correlations, statistical significance, etc.

Note that the algorithm does not make any additions or removals of LD parameters during the bundle iterations. Instead, the bundle is allowed to converge, and any decision of whether to add or remove any LD parameters for a subsequent run of the main algorithm is decided in step 7 after the convergence.

Throughout this paper, the pixel aspect ratio and skew were kept constant and were not estimated.

2.2 Bundle adjustment algorithms

The bundle adjustment (BA) algorithms used in this paper are from the free Damped Bundle Adjustment Toolbox³, described in Börlin and Grussenmeyer (2013a):

- GM** The classical Gauss-Markov bundle adjustment algorithm.
- GNA** The Gauss-Newton algorithm with Armijo line search.
- LM** The original Levenberg-Marquardt algorithm.
- LMP** The Levenberg-Marquardt algorithm with Powell dogleg.

The camera model defined in the toolbox used the Euler $\omega - \phi - \kappa$ angles and the Brown (1971) K_1-K_3 , P_1-P_2 lens distortion model.

³The toolbox source code and examples of data sets are available via email from niclas.borlin@cs.umu.se.

2.3 Data sets

One calibration data set for each of five different camera-lens combinations was generated. The lenses were conventional, i.e. no fish-eye lenses were used. Details about the cameras are given in Table 1. Three data sets used the 2D Photomodeler calibration sheet consisting of an approximately 1-meter-square, 10-by-10 grid of black circular targets, including 4 coded targets used as control points. One data set used a 5-by-6-by-2 meter 3D calibration coded targets indoor test field. The final data set used a 60-by-35-by-20 meter array of control points measured on the INSA Building in Strasbourg (Börlin and Grussenmeyer, 2013b). Images and the complete camera networks are shown in figures 1 and 2, respectively.

Except in data set 5, the object and control points were measured by the automatic circular target measurement technique in Photomodeler Scanner 2012. Details about each calibration data set are given in Table 2. The measured x, y coordinates and control point coordinates were exported from Photomodeler and imported into the toolbox.

For data set 3, a manufacturing flaw of the calibration sheet was discovered. Thus, the nominal coordinates of the control points could not be trusted. Instead, a minimum datum definition was used (two control points + Z coordinate of third) for the bundle. However, for the resection process in step 4 of the main algorithm, the control points were used as such with their nominal coordinates.

The toolbox does not yet include automatic blunder detection. Instead manual blunder detection was performed based on result and residual plots available in the toolbox.

3. EXPERIMENTS

For each data set, two experiments were performed to determine the pull-in range of the camera calibration with respect to the initial focal length estimate f_0 . For the FULL experiment, the full set of lens distortion (LD) parameters K_1, K_2, K_3, P_1 , and P_2 was estimated. The initial focal length $f_0 = m \cdot f_*$ value was successively set to between $m = 1/8$ and $m = 128$ times the true value f_* . Steps 1–5 of the algorithm in Section 2.1 were used to determine the other initial values. In step 6 of the algorithm, the same initial values were given to each of the four bundle algorithms listed in Section 2.2. The required number of iterations or failure to converge to the true values was recorded for each algorithm. A maximum number of 100 iterations was allowed for convergence.

For the second experiment, a STABLE set of LD parameters was determined for each data set. The posterior statistical properties

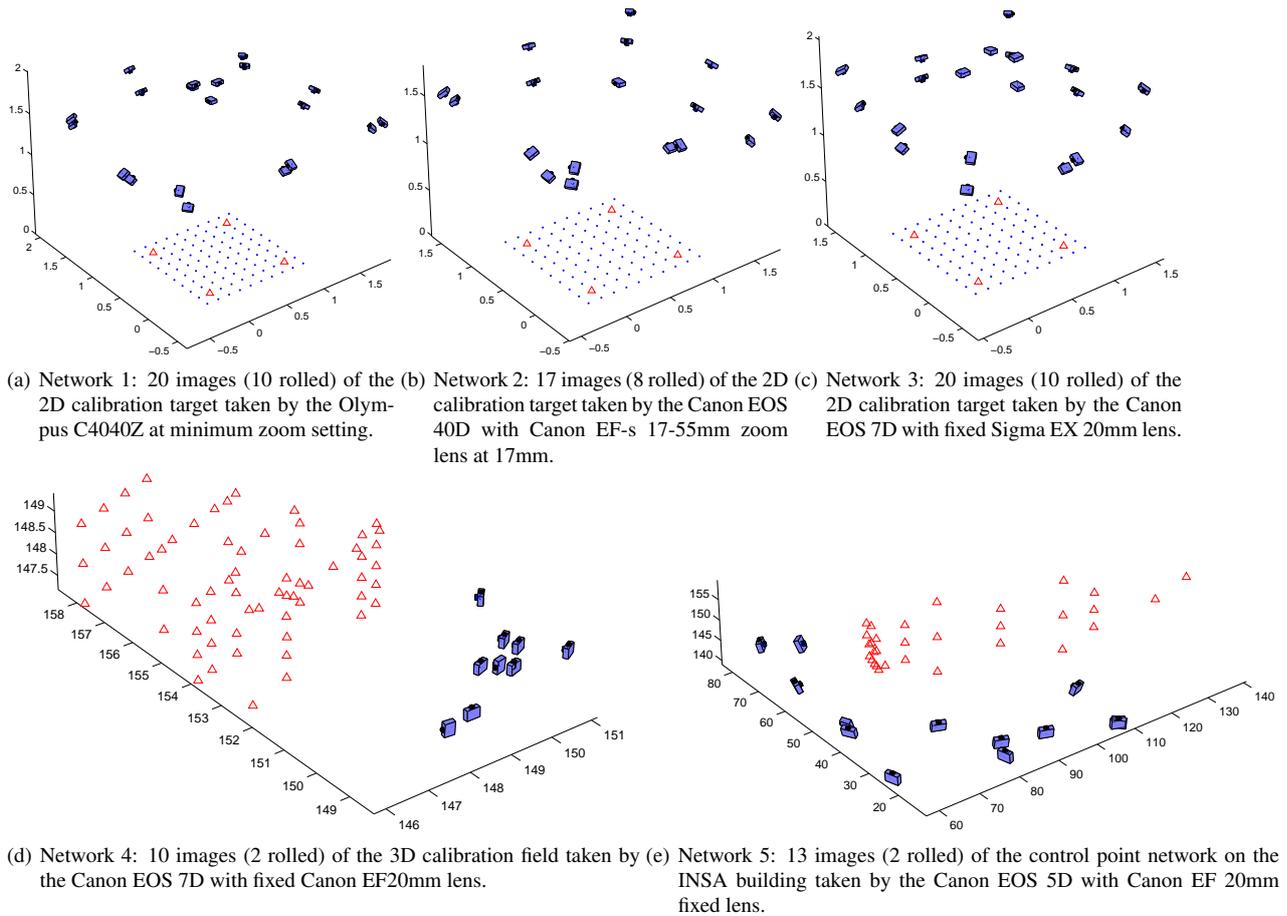


Figure 2: The camera networks used for the calibrations. Control points are plotted as red triangles, object points as blue dots. The object space unit is meters.

Table 1: The cameras and lenses used in the experiments. The listed crop factor is the ratio of the diagonal of the 35mm film format (36-by-24 mm) to the sensor diagonal. The angle of view value is computed across the sensor diagonal. All cameras were focused at infinity. Zoom lenses were set at their smallest setting. The sensor sizes have been retrieved from public sources. The EXIF info shows what focal length and/or sensor width information was available in the images.

Data set	Camera	Lens	Sensor size w x h (mm)	Crop factor	Angle of view (deg)	Image size (pixels)	EXIF info (mm)	
							Focal	Sensor
1	Olympus C4040Z	built-in (zoom)	7.18 x 5.32	4.8	62	2272 x 1704	7.3	-
2	Canon EOS 40D	Canon EF-s 17–55	22.2 x 14.8	1.6	76	3888 x 2592	17.0	22.25
3	Canon EOS 7D	Sigma EX 20	22.3 x 14.9	1.6	68	5184 x 3456	20.0	23.04
4	Canon EOS 7D	Canon EF 20	22.3 x 14.9	1.6	68	5184 x 3456	20.0	23.04
5	Canon EOS 5D	Canon EF 20	35.8 x 23.9	1.0	94	4368 x 2912	20.0	35.94

Table 2: Statistics about the measurements of each data set. The target depth-to-width and distance values are computed with respect to a stationary camera, i.e. as if the calibration object was moving. The reported redundancy value is the number of observation minus the number of unknowns. The average max angle is the average of the maximum intersection angle of the observation rays for each target. * In data set 3, the control point coordinates were used by the resection process only. For details, see the text.

Data set	Number of images		Computed target		Number of points		Measurement type		Radial image coverage	Avg. ray count	Redundancy	Avg. max angle (deg)
	total	rolled	depth-to-width (m)	distance (m)	ctrl	object	control	image				
1	20	10	1.5-to-2.0	1.4–2.9	4	96	assumed	auto	90%	19.7	3533	86
2	17	8	1.7-to-2.0	1.1–2.8	4	96	assumed	auto	88%	16.7	2939	85
3	20	10	1.5-to-2.0	1.2–2.8	4*	96	assumed	auto	94%	19.2	3416	81
4	10	2	7.9-to-5.3	2.4–10	66	0	total station	auto	92%	8.1	1005	33
5	13	2	66-to-66	15–81	31	0	total station	manual	81%	5.6	265	74

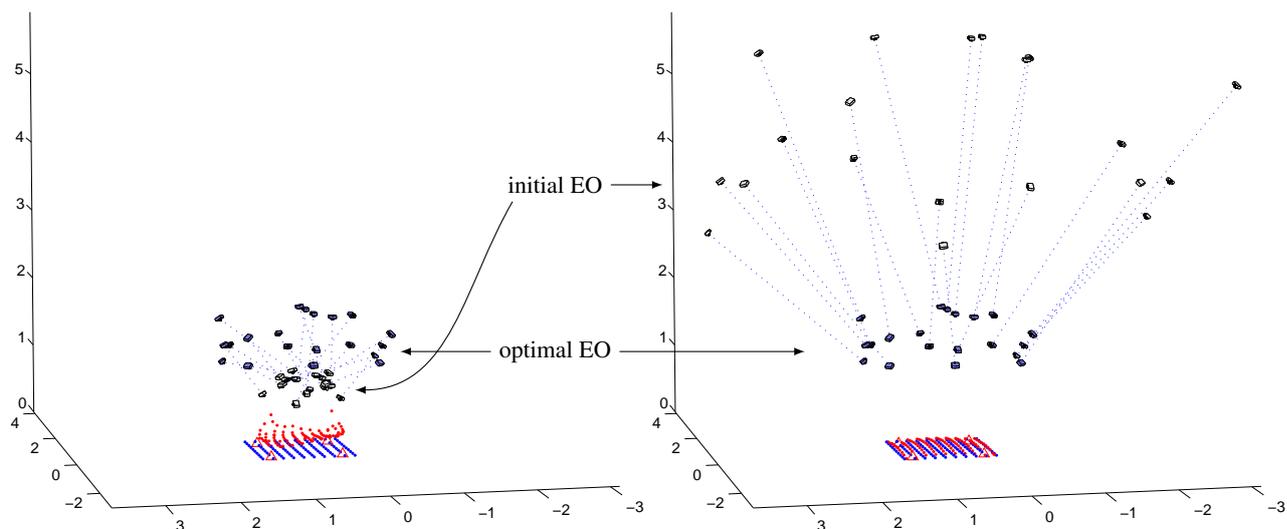


Figure 3: Initial EO and OP values for $f_0 = 0.3f_*$ (left) and $f_0 = 3f_*$ (right) for data set 3 ($f_* = 20.942$ mm). Initial (white) and optimal (blue) camera positions are connected by dotted lines. Initial OP coordinates are red, optimal blue. Given the smaller f_0 value, the distance to the target is underestimated and the cameras cluster above the centre of the target. Their incorrect position affects the initial OP values as well. Given the larger f_0 value, the distance to the target is overestimated and the cameras fan out radially away from the target. In both cases, the orientations of the cameras are approximately correct.

of the LD parameters from the FULL experiment were analysed to determine which parameters were stable and should be part of the second experiment. Parameters with correlation values above 95% were considered unstable and were excluded, as were parameters whose estimated values were not statistically significant. The parameters were tested in the following order: K_3 , (P_1, P_2) , K_2 , and finally K_1 , where (P_1, P_2) were tested together. If any parameter was removed, the remaining parameters were re-estimated and the process was repeated until only stable parameters remained. The parameter testing procedure was applied once for each data set after which the pull-in investigation described above was executed.

For each data set and experiment, the “true” calibration parameters were estimated by one of the bundle algorithms based on the best available initial values after which consensus on the optimal values was verified by the other algorithms.

The effect of too small or too large initial focal length on the other parameters is illustrated in Figure 3 for data set 3. Given a smaller f_0 than the true, i.e. $m < 1$, the distance to the target is underestimated and the initial EO position are grouped together closer to the centre of the target. The opposite is true for a larger than true f_0 with $m > 1$: The distance to the target is overestimated by approximately a factor of m and the initial camera positions are distributed radially away from the target. In both cases, the orientations of the cameras are approximately correct.

4. RESULTS

The results of the calibration and parameter selection process are given in Table 3 together with σ_0 and computed distortion of each camera. The σ_0 for data set 5 is significantly higher than for the other projects, consistent with the use of manual measurements. The other σ_0 values based on coded target measurements are similar to those reported by Fraser (2013). Furthermore, the distortion for the compact Olympus camera is more than twice that of the DSLR cameras.

Table 3: The calibration results of the FULL experiment (upper half) and STABLE experiment (lower half). The maximum distortion is computed at the sensor corners and is given in percent of the sensor half-diagonal.

Data set	σ_0 (px)	f_* (mm)	Estimated parameters	Max distortion
Experiment 1 — FULL				
1	0.17	7.460 ± 0.001	K_1, K_2, K_3, P_1, P_2	8.2%
2	0.36	17.919 ± 0.003	K_1, K_2, K_3, P_1, P_2	4.3%
3	0.11	20.942 ± 0.001	K_1, K_2, K_3, P_1, P_2	4.0%
4	0.43	20.695 ± 0.002	K_1, K_2, K_3, P_1, P_2	3.6%
5	1.3	20.62 ± 0.03	K_1, K_2, K_3, P_1, P_2	3.9%
Experiment 2 — STABLE				
1	0.18	7.468 ± 0.001	K_1, K_2, P_1, P_2	8.9%
2	0.36	17.918 ± 0.003	K_1, K_2, P_1, P_2	4.2%
3	0.11	20.942 ± 0.001	K_1, K_2, P_1, P_2	4.0%
4	0.45	20.696 ± 0.002	K_1, K_2, P_1, P_2	3.6%
5	1.3	20.61 ± 0.02	K_1, K_2, P_1, P_2	2.5%

Regarding the convergence, the same pattern is present in the results of all data sets and experiments. As a typical example, Figure 5 shows the number of required iterations for the four BA methods on data set 3. In a region near the true focal length, all methods converge within a few iterations. The width of the region vary by method. Near the lower end of the region, around $f_0 = 1/4f_*$, the resection process fails, and all methods signal failure.

At the upper end of the region, around $f_0 = 2f_*$, the GM and LM algorithms switch abruptly from convergence to failure. In contrast, the GNA and LMP algorithms have a gradual increase in the number of required iterations, roughly following a parabola, and their tolerance to too high initial f_0 values is significantly higher than for the GM and LM algorithms.

Figure 4 shows the difference in behaviour of the algorithms for data set 3 at $f_0 = 3f_*$. The iteration trace shows that the GM

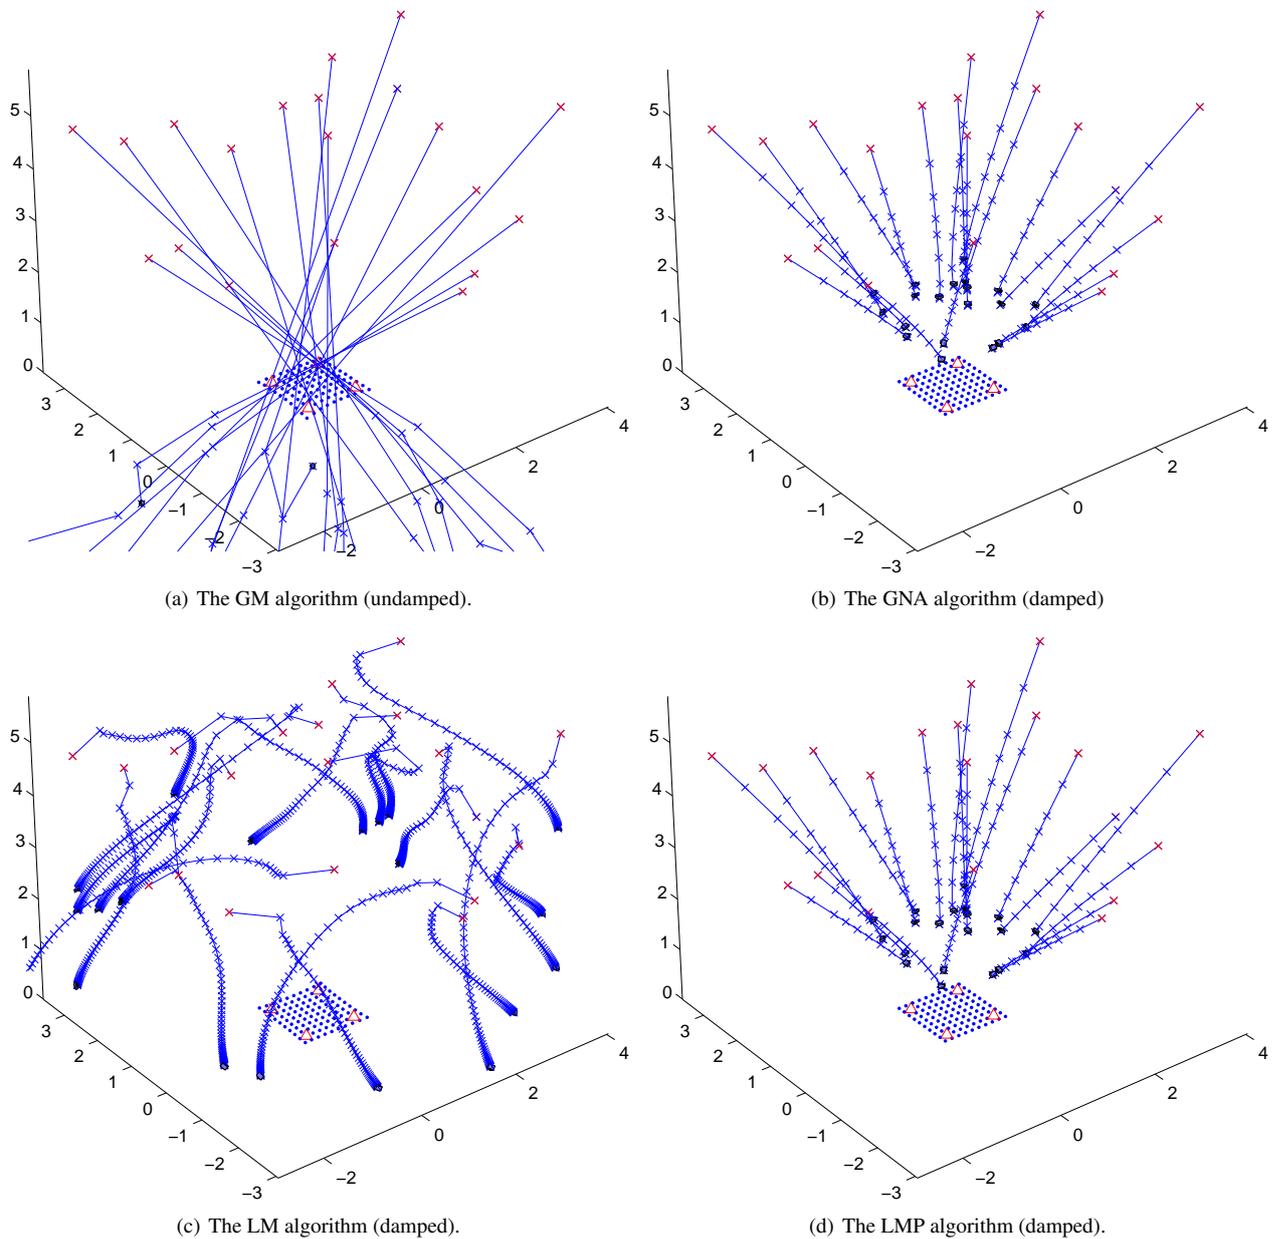


Figure 4: Iteration trace of the EO parameters for the four bundle algorithms on data set 3 with $f_0 = 3f_*$. The initial EO positions, same for all algorithms, are indicated by red crosses. The EO positions at subsequent iterations are indicated by blue crosses. The GM method (a) makes an early over-correction where the cameras end up on the wrong side of the target. After that the GM method never recovers. For the the GNA (b) and LMP (d) algorithms, damping is active during the first iterations and stops the updates from becoming too large. Near the solution, damping is relaxed and no damping is required during the final iterations. The LM algorithm (c) oscillates between damped and undamped and makes an early over-correction of the rotation angles after which it does not recover in the allowed number of iterations. The shown results are typical for $m > 2$ for all data sets and experiments.

algorithm makes an initial over-correction that results in the cameras appearing on the wrong side of the target. After that, the GM algorithm never recovers. In contrast, the GNA and LMP algorithms follow approximately the same pattern. Initially, their damping schemes are active and stop the updates from becoming too large. Later, the damping is relaxed as they approach the minimum. For cases where the GNA and LMP algorithms did not converge, the trace plots show that they were making progress but did not reach the optimal solution in the allowed number of iterations. The behaviour of the LM algorithm is more complex. It oscillates between damped and undamped steps and makes an early over-correction of the orientation angles from which it does not recover within the allowed number of iterations.

The results of all experiments and data sets show the same pattern. At the lower end, all BA algorithms fail below a certain m due to resection failure. At the upper end, the GM and LM algorithms fail abruptly at roughly the same m , whereas GNA and LMP converges for higher m values, albeit in an increasing number of iterations. Figure 6 shows the pull-in region for each method and experiment, defined as the interval where convergence was achieved in less than 100 iterations. The undamped GM method converged when the initial f_0 was within a factor of 2 of the true. The damped GNA and LMP algorithms converged from a wider range of initial focal lengths and have a large tolerance to too high initial values with a pull-in range of at least 32 times f_* . The behaviour of the LM algorithm is again more complex. On some experiments, it shows a pull-in range comparable to the undamped GM algorithm. On others, the pull-in range was significantly smaller. On the STABLE experiment on data set 5, the LM algorithm failed to converge even for $f_0 = f_*$. An analysis of the iteration trace revealed that the reason was a combination of poor initial EO values for two cameras combined with the damped-undamped oscillations described above.

5. DISCUSSION

5.1 Unknown focal length and/or sensor size

In this paper we used the actual sensor size and focal lengths, available either from the EXIF information or from knowledge about the used camera. Our data suggests that an initial value of $f_0 \leq 32f_*$ will give convergence, at least for geometrically strong networks with good image coverage.

If the sensor size is unknown, we observe that the choice of unit for the sensor size and focal length is arbitrary, since only their ratios appears in the collinearity equations. Furthermore, said ratios are functions of the angle of view of the camera. Since it is reasonable to assume that the angle of view is restricted to a reasonable range, we suggest that the sensor height is fixed arbitrarily and the other internal parameters are scaled accordingly. Barring other knowledge, we suggest fixing the sensor height to 24 mm to match the 36-by-24 mm “35 mm” standard film format, in which case the calibrated focal length will be the “35 mm equivalent” focal length of the camera.

For our cameras, the angle of view is 60–90 degrees and the focal length is approximately equal to the sensor height. Based on that and the conclusions from the previous paragraphs, we suggest an f_0 of approximately 25 times the sensor height, or 600 mm in “35 mm equivalent” focal length, corresponding to an angle of view of about 4 degrees. Unless the real focal length is very long (angle of view is very small), this should be on the high side of the actual focal length and within the pull-in range of the best damped algorithms. If the angle of view is significantly higher than 90 degrees, a lower f_0 should be used.

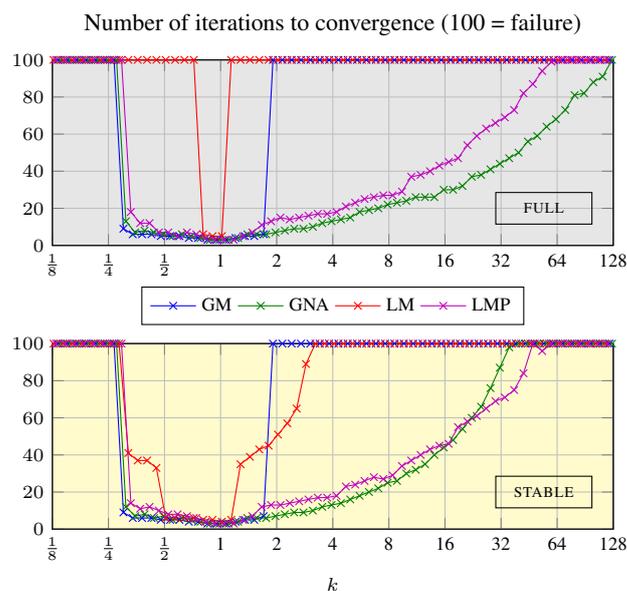


Figure 5: Convergence results for data set 3: The Canon EOS 7D with Sigma lens for the FULL (upper panel) and STABLE (lower panel) experiments. The panels show the number of iterations required for each BA algorithm for varying values of m . A value of 100 indicate failure. Near the true focal length ($m = 1$), all algorithms converge quickly. Below $m = 1/4$ the resection fails, so neither bundle works. Above approximately $m = 2$, the GM (blue) and LM (red) algorithms fail whereas the GNA (green) and LMP (purple) algorithms degrade gracefully and converge up to at least $m = 32$.

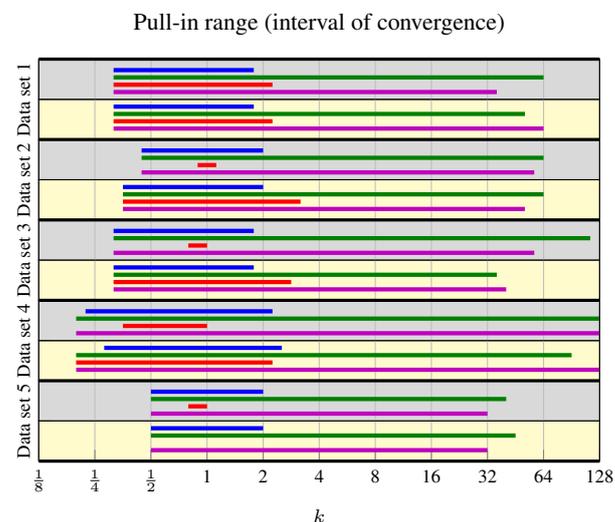


Figure 6: The pull-in range for varying values of m for all data sets and experiments. Gray background corresponds to the FULL experiment, yellow to the STABLE experiment. In general, the undamped GM method (blue) converges when the initial focal length is roughly within a factor of 2 of the true for all experiments and data sets. The damped GNA (green) and LMP (purple) methods converge in less than 100 iterations for $m = 1/2$ up to at least $m = 32$. The damped LM (red) behaviour is more complex and is discussed in the text. Data set 4 has the widest pull-in range, data set 5 the narrowest.

If the EXIF focal length is available it could be used. However, using the standard f_0 value suggest above instead does handle the situation when the EXIF information is incorrect, something that has happened to the first author using another copy of the zoom lens used in data set 2.

5.2 Proposed algorithm

In summary, we propose the following rules of thumb for using the algorithm in Section 2.1 for camera calibration:

- If the sensor size is known, use it. Otherwise, assume a sensor height of 24 mm and scale the sensor width to match the assumed (unit or otherwise) pixel aspect ratio.
- If the angle of view is roughly 60 degrees, use 25 times the sensor height as f_0 . If the angle of view is above 90 degrees, use 10 times the sensor height as f_0 .
- Use the GNA or LMP algorithms for the bundle.
- Initially, estimate all LD parameter with the bundle. After convergence, exclude any unstable parameters and re-estimate until only stable parameters remain.

5.3 Discussion

Even though photogrammetric state of the art often include self-calibration or on-the-job-calibration, off-line camera calibration is still important to investigate the stability of the camera, investigate the metric quality of new cameras, or as initial values for self-calibration (Shortis et al., 1998; Luhmann et al., 2006; Fraser, 2013).

The theory and formulation of radial and de-centring lens distortions from Brown (1971) have been adopted in close range photogrammetry as well as in the field of computer vision. Whereas the photogrammetric literature focus on quality, the importance of strong networks, high redundancy, good initial values and the use of the Gauss-Markov estimation model (McGlone et al., 2004; Luhmann et al., 2006; Fraser, 2013), the computer vision literature focus on simplicity, high level of automation, and the use of the Levenberg-Marquardt method (Zhang, 2000; Hartley and Zisserman, 2003; Snavely et al., 2008; Furukawa and Ponce, 2009).

While several computer vision-based camera calibration software are publicly available, including their source, the computation of statistical quality parameters is generally lacking, especially with respect to correlations between the estimated parameters (Remondino and Fraser, 2006). In contrast, photogrammetric software often provide a detailed statistical analysis, including automatic exclusion of unstable parameters. However, if the bundle fails to converge, the available statistical analysis may be anything between limited and useless. Good photogrammetric software include tools and suggest guidelines, but due to the closed-source nature of photogrammetric software, the detailed information needed to determine the reason for the failure, be it poor camera network configuration, measurement blunders, or poor initial values due to erroneous EXIF information, is often hidden from the end user.

In this paper we investigated a camera calibration algorithm based on the freely available Damped Bundle Adjustment Toolbox for Matlab that required initial values of only one parameter — the focal length. The results show that the damped GNA and LMP bundle algorithms have a wide pull-in range, especially compared to the GM and LM algorithms, even on highly correlated camera

parameter sets. This is consistent with the results reported by Börlin and Grussenmeyer (2013a). The ability to converge with correlated parameters, enables a posteriori statistical analysis and exclusion of unstable parameters starting with the full parameter set rather than adding parameter incrementally.

Other techniques for estimating the initial focal length, not investigated in this paper, include the DLT method for 3D targets by Abdel-Aziz and Karara (1971) and the relative orientation techniques by Nistér (2004) and Snavely et al. (2008). These techniques could of course be used as a complement, should our standardised approach fail.

Given its popularity in the computer vision community, the performance of the LM algorithm was surprisingly poor. The reason was largely attributed to an implementation detail of the algorithm in Börlin and Grussenmeyer (2013a); the choice of λ_0 , the initial and minimum damping value. The λ_0 was introduced to ensure bias-free co-variance estimates at the minimum, should the algorithm converge. Advocates for the LM method may argue that the algorithm was unfairly treated and it may certainly be possible to fine-tune the LM parameters to allow the LM method to have a wider pull-in range. However, since neither the GNA nor LMP algorithms had the same problems, this rather reinforces the argument in Börlin and Grussenmeyer (2013a) that unless a biased convergence is acceptable, the λ damping scheme used by LM introduces an implementation detail that may — and in this case, did — cause problems.

The STABLE experiment may be seen as multi-stage process to estimate both which parameters are stable and the actual parameter values, similar to the established multi-stage practice used by many photogrammetric software. The difference is twofold; whereas the established “forward” practice usually starts with few parameters and add new parameters incrementally until no more significant parameters can be added, our “backward” procedure starts with all parameters and removes any statistically insignificant and/or highly correlated parameters until all remaining parameters are stable. A further difference is that in the “forward” procedure, the estimated parameter values at one stage is used as the initial values for the next. In our “backward” strategy, all remaining parameters to be estimated are reset to their original initial values before adjustment. The “backward” strategy was chosen for simplicity, but our main algorithm could also be modified to use the “forward” strategy.

The plotting features of the toolbox, especially the camera network plots illustrated in figures 2–4, where useful to understand the reason for bundle failure. Other plots were helpful for blunder detection. The toolbox did not do any automatic parameter exclusion. However, the computed statistical analysis provided the information necessary to determine which parameters were stable. Furthermore, the availability of the source code simplified the automation of several key tasks.

The fixed unit aspect ratio and manual outlier detection are limitations to this study, as is the relatively small number of cameras and lenses. However, the wide pull-in range reported does indicate that the suggested algorithm should work for a large majority of conventional camera-lens combinations.

5.4 CONCLUSIONS

The Gauss-Newton-Armijo (GNA) and Levenberg-Marquardt-Powell (LMP) algorithms of the Damped Bundle Adjustment Toolbox applied to the camera calibration problem have superior pull-in range compared to the classical Gauss-Markov algorithm

and the Levenberg-Marquardt method. The GNA and LMP algorithms converged even if the initial focal length estimate was 32 times too large. This was the case even for highly correlated parameter sets, enabling a statistical analysis to remove unstable parameters.

The source of the Damped Bundle Adjustment Toolbox is freely available and includes implementations of the above mentioned algorithms. The toolbox estimate the standard deviations and correlations of the estimated parameters, something only partially available in other free camera calibration toolboxes. The plotting features of the toolbox (see figures 2–4) were especially useful to understand the behaviour of the bundle algorithms during the iterations.

We propose a standardised procedure that does not require any camera knowledge. The procedure has been demonstrated to work on both narrow- and wide-angle lenses.

Future work includes extending the camera model to include the skew and aspect ratio in the estimation process and to extend the algorithms to handle calibration of special lenses, such as the fish-eye.

References

- Abdel-Aziz, Y. I. and Karara, H. M., 1971. Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. In: Proceedings of ASP Symposium on Close-range Photogrammetry, University Illinois at Urbana-Champaign, Urbana, IL, pp. 1–18.
- Börlin, N. and Grussenmeyer, P., 2013a. Bundle adjustment with and without damping. *Photogramm Rec* 28(144), pp. 396–415.
- Börlin, N. and Grussenmeyer, P., 2013b. Experiments with metadata-derived initial values and linesearch bundle adjustment in architectural photogrammetry. *ISPRS Annals of the Photogrammetry, Remote Sensing, and Spatial Information Sciences II-5/W1*, pp. 43–48.
- Brown, D. C., 1971. Close-range camera calibration. *Photogrammetric Engineering* 37(8), pp. 855–866.
- Fraser, C. S., 2013. Automatic camera calibration in close range photogrammetry. *Photogramm Eng Rem S* 79(4), pp. 381–388.
- Furukawa, Y. and Ponce, J., 2009. Accurate camera calibration from multi-view stereo and bundle adjustment. *Int J Comp Vis* 84(3), pp. 257–268.
- Hartley, R. I. and Zisserman, A., 2003. *Multiple View Geometry in Computer Vision*. 2nd edn, Cambridge University Press, ISBN: 0521540518.
- Luhmann, T., Robson, S., Kyle, S. and Harley, I., 2006. *Close Range Photogrammetry*. Whittles, Scotland, UK.
- McGlone, C., Mikhail, E. and Bethel, J. (eds), 2004. *Manual of Photogrammetry*. 5th edn, ASPRS.
- Nistér, D., 2004. An efficient solution to the five-point relative pose problem. *IEEE T Pattern Anal* 26(6), pp. 756–770.
- Remondino, F. and Fraser, C., 2006. Digital camera calibration methods: Considerations and comparisons. *International Archives of Photogrammetry, Remote Sensing, and Spatial Information Sciences XXXVI(5)*, pp. 266–272.
- Shortis, M. R., Robson, S. and Beyer, H. A., 1998. Principal point behaviour and calibration parameter models for kodak DCS cameras. *Photogramm Rec* 16(92), pp. 165–186.
- Snavely, N., Seitz, S. M. and Szeliski, R., 2008. Modeling the world from internet photo collections. *Int J Comp Vis* 80(2), pp. 189–210.
- Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE T Pattern Anal* 22(11), pp. 1330–1334.