

PROS AND CONS OF CONSTRAINED AND UNCONSTRAINED FORMULATION OF THE BUNDLE ADJUSTMENT PROBLEM

Niclas BÖRLIN¹, Pierre GRUSSENMEYER², Jerry ERIKSSON¹, Per LINDSTRÖM¹

¹Department of Computing Science, Umeå University, UMEÅ, Sweden, Niclas.Borlin@cs.umu.se

²National Institute of Applied Sciences of Strasbourg, MAP-PAGE UMR 694, STRASBOURG, France, pierre.grussenmeyer@insa-strasbourg.fr

KEY WORDS: Algorithms, Mathematics, Bundle, Generalisation, Modelling, Performance, Reliability.

ABSTRACT

Two implementations of the bundle adjustment problem were applied to a subset of the Zürich City Hall reference data set. One implementation used the standard Euler angle parameterisation of the rotation matrix. The second implementation used all nine elements of the rotation matrix as unknowns and six functional constraints. The second formulation was constructed to reduce the non-linearity of the optimisation problem. The hypothesis was that a lower degree of non-linearity would lead to faster convergence. Furthermore, each implementation could optionally use the line search damping technique known from optimisation theory.

The algorithms were used to solve the relative orientation problem for a varying number of homologous points from 33 different camera pairs.

The results show that the constrained formulation has marginally better convergence properties, with or without damping. However, damping alone halves the number of convergence failures at a minor computational cost.

The conclusion is that except to avoid the singularities associated with the Euler angles, the preferred use of the constrained formulation remains an open question. However, the results strongly suggest that the line search damping technique should be included in standard implementations of the bundle adjustment algorithm.

1 INTRODUCTION

In photogrammetry, when maximum precision is required, the bundle adjustment problem is usually solved. The bundle adjustment problem is a non-linear least squares problem, which must be solved iteratively.

Within the optimisation community, it is well-known that the convergence properties of a non-linear problem depend on at least three factors; the quality of the starting approximation, the optimisation method used, and the degree of non-linearity of the problem, e.g. (Gill et al., 1981).

Earlier papers have described modifications of the optimisation method in order to improve the convergence properties, e.g. (Börlin et al., 2003). In this paper, the discussion is extended to include different formulations of the bundle adjustment problem with and without functional constraints. The re-formulations are done in order to reduce the degree of non-linearity of the problem. Earlier results on optical photogrammetry (Hägglund, 2004) indicate that the difference in formulations have a small impact on convergence properties if the problem has strong geometry and/or high redundancy. This investigation thus focuses on the relative orientation problem which may have weak geometry and/or low redundancy.

2 BUNDLE ADJUSTMENT

Bundle adjustment is based on the collinearity equations

$$\begin{aligned} x - x_0 &= \frac{-f(m_{11}(X-X_0)+m_{12}(Y-Y_0)+m_{13}(Z-Z_0))}{m_{31}(X-X_0)+m_{32}(Y-Y_0)+m_{33}(Z-Z_0)}, \\ y - y_0 &= \frac{-f(m_{21}(X-X_0)+m_{22}(Y-Y_0)+m_{23}(Z-Z_0))}{m_{31}(X-X_0)+m_{32}(Y-Y_0)+m_{33}(Z-Z_0)}, \end{aligned}$$

which describe the projection of a world point (X, Y, Z) onto an image point (x, y) of a camera with principal distance f and principal point at (x_0, y_0) . The camera is placed at (X_0, Y_0, Z_0) and its orientation is described by the rotation matrix M .

The bundle adjustment method minimises the difference between projected object points and the corresponding measurements. For each world point i and camera j , define the residual v_i^j to be

$$v_i^j = \begin{bmatrix} x_i^j - x_0^j + f \frac{m_{11}^j(X_i - X_0^j) + m_{12}^j(Y_i - Y_0^j) + m_{13}^j(Z_i - Z_0^j)}{m_{31}^j(X_i - X_0^j) + m_{32}^j(Y_i - Y_0^j) + m_{33}^j(Z_i - Z_0^j)} \\ y_i^j - y_0^j + f \frac{m_{21}^j(X_i - X_0^j) + m_{22}^j(Y_i - Y_0^j) + m_{23}^j(Z_i - Z_0^j)}{m_{31}^j(X_i - X_0^j) + m_{32}^j(Y_i - Y_0^j) + m_{33}^j(Z_i - Z_0^j)} \end{bmatrix},$$

where (X_i, Y_i, Z_i) are the object point coordinates, and (x_i^j, y_i^j) are the measured coordinates in camera j , corrected for lens distortion.

Using the notation in (Mikhail et al., 2001, Appendix B), the value to be minimised is

$$\phi = v^T W v,$$

where v is a vector containing all v_i^j , and W is a weight matrix. The minimisation is performed over a set of unknown parameters w . To emphasise that ϕ and v are functions of the unknown parameters, this will be written

$$\phi(w) = v(w)^T W v(w).$$

In optimisation, the function $\phi(w)$ is known as the *objective function*.

At each iteration k , the problem is linearised and the normal equations (Mikhail et al., 2001, Equation B-19)

$$N \Delta = t \quad (1)$$

or system equations (Mikhail et al., 2001, Equation B-27)

$$\begin{bmatrix} -N & C^T \\ C & 0 \end{bmatrix} \begin{bmatrix} \Delta \\ k_c \end{bmatrix} = \begin{bmatrix} -t \\ g \end{bmatrix} \quad (2)$$

are solved and the unknown parameter vector w is updated

$$w_{k+1} = w_k + \Delta. \quad (3)$$

The process is repeated until the update Δ is small enough, or the maximum number of iterations has been exceeded.

3 PROBLEM FORMULATIONS

3.1 Euler angles

The standard parameterisation of the rotation matrix M is by the Euler angles for some specified sequence of rotation axes. This results in a formulation with three unknowns and no functional constraint. Thus, in each iteration, the normal equations (1) would be solved.

3.2 Rotation matrix

An alternative formulation is to treat all rotation matrix elements $m_{11}, m_{12}, \dots, m_{33}$ as unknowns and use the six functional constraints $g_i(w) = 0, i = 1, \dots, 6$, where

$$\begin{aligned} g_i(w) &= m_i^T m_i - 1, \quad i = 1, 2, 3, & g_4(w) &= m_1^T m_2, \\ g_5(w) &= m_1^T m_3, & g_6(w) &= m_2^T m_3, \end{aligned} \quad (4)$$

where $M = [m_1, m_2, m_3]$. This formulation has nine parameters and six functional constraints and replaces the highly non-linear trigonometric functions with less non-linear constraints. Furthermore, the rotation matrix is always well-defined for a given rotation, and is thus singularity-free. However, the constraints are satisfied both for a rotation matrix ($|M|=+1$) and a reflection matrix ($|M|=-1$). For this parameterisation, the system equations (2) would be solved in each iteration.

4 ALGORITHMS

Two versions of the bundle adjustment algorithm were implemented in Matlab (The Mathworks, Inc.) with the parameterisations of sections 3.1 and 3.2.

4.1 Line search

Furthermore, the algorithms were implemented to work with or without *line search*, a damping method common within the optimisation community (Gill et al., 1981, Nash and Sofer, 1996, Björck, 1996). With line search, the update formula (3) is replaced by

$$w_{k+1} = w_k + \alpha_k \Delta, \quad (5)$$

where the scalar α is taken as the first value of the sequence $1, \frac{1}{2}, \frac{1}{4}, \dots$ such that the updated estimate w_{k+1} of the parameters are *sufficiently better* than the current estimate w_k (for details, see (Börlin et al., 2003)). E.g. for




Figure 1: The Olympus subset of the Zürich data set with camera stations indicated.

the unconstrained formulation, the line search method will ensure that the next objective function value $\phi(w_{k+1})$ is smaller than the current $\phi(w_k)$, i.e. that the residual sum will decrease for each iteration. Since only objective function evaluations are used, the method is computationally cheap.

Thus, the following bundle adjustment algorithms were available:

Alg. 1U Euler angles, undamped (no line search). Corresponds to classical bundle adjustment.

Alg. 1D Euler angles, damped (with line search). This algorithm corresponds to the Gauss-Newton algorithm in e.g. (Börlin et al., 2003).

Alg. 2U All nine rotation matrix elements and the six orthogonality constraints (4). Undamped. Corresponds to bundle adjustment with functional constraints.

Alg. 2D All nine rotation matrix elements and the six orthogonality constraints (4). Damped. Corresponds to the GNC (Gauss-Newton, Constrained) algorithm in (Börlin et al., 2003).

5 DATA SET

The Zürich City Hall data set (Streilein et al., 1999) was selected to test the algorithms. The data set consists of 31 images taken with an Olympus C1400L and a Fuji DS300 camera. For this investigation, the 14 Olympus images were selected. The difference in focal settings described in (Rottensteiner et al., 2001) was ignored. From the Olympus images, 33 camera pairs were identified with a suitable number of homologous points, see Figure 1 and Table 1. The smallest number of homologous points were 42.

Table 1: Camera pair listing. Stations numbers are from the Zürich data set and Figure 1. Distance, object depth, and image coverage are average for both cameras.

Pair #	Stations	Homologous points	Base [m]	Intersection angle [°]	Distance [m]	Object depth [m]	Image coverage [%]
1	5, 14	674	133	96	77	34	15
2	5, 13	667	53	51	51	24	29
3	13, 14	659	85	46	81	22	19
4	4, 25	294	38	74	28	13	42
5	14, 15	283	29	8	91	29	11
6	28, 31	263	26	24	39	14	32
7	26, 27	260	9	2	24	14	43
8	15, 28	259	89	85	54	17	24
9	15, 31	259	81	65	64	13	11
10	4, 26	257	57	95	33	21	39
11	4, 5	249	50	71	34	16	25
12	14, 28	245	103	77	65	16	23
13	14, 31	245	88	57	74	12	11
14	4, 30	211	21	32	27	10	44
15	4, 29	209	23	27	28	8	44
16	5, 29	205	27	44	28	10	41
17	5, 30	200	29	39	30	10	39
18	24, 25	176	12	22	21	6	53
19	4, 24	172	26	53	26	10	42
20	25, 26	167	20	24	24	9	29
21	4, 27	166	65	94	38	21	31
22	29, 30	165	2	6	22	8	44
23	27, 28	157	14	14	26	21	33
24	5, 12	121	58	57	51	24	26
25	12, 13	120	6	6	57	13	29
26	12, 14	111	79	40	84	22	17
27	26, 28	85	22	13	27	21	26
28	13, 15	65	56	37	68	20	21
29	24, 26	52	31	42	25	7	15
30	5, 15	51	105	88	67	31	16
31	4, 28	49	79	108	43	34	18
32	12, 15	43	51	32	72	20	19
33	25, 27	42	28	23	27	5	8

The measured points and internal camera parameters were exported from Photomodeler (EOS Systems Inc.) and loaded into Matlab for processing.

6 ADJUSTMENTS

For each camera pair, an initial guess of the relative orientation between the cameras was calculated with the normalised 8-point algorithm and the essential matrix, as described in (Hartley and Zisserman, 2000). All homologous points were used. Given the camera relative orientation, initial object point estimates were calculated by forward intersection.

The four algorithms were given the same starting approximations and used to solve the two-camera bundle adjustment problem. To handle the datum problem, the first camera position and orientation and one coordinate of the second camera were kept fixed, i.e. a dependent relative orientation (Mikhail et al., 2001). The success or failure to converge and the used number of iterations were recorded




Figure 2: Algorithm behaviour for Case 2, all points. Large initial steps for the undamped algorithms lead to failure after a few iterations. Transparent cameras indicate position estimates during the iterations, opaque cameras indicate final camera position. Light points indicate initial object point approximations, dark point indicate final object points (converged case only).

for each algorithm. The maximum number of iterations was set to 20. The termination criteria was

$$\|J_k p_k\| \leq \varepsilon_r \|r_k\| \text{ and } \max |g_i(w_k)| \leq \varepsilon_g$$

from (Börlin et al., 2003, Equation (16a)) with constants $\varepsilon_r = 10^{-5}$, $\varepsilon_g = 10^{-8}$, roughly corresponding to termination when the largest update was around 1 mm. Other algorithmic constants were set to $\mu = 0.1$, $\alpha_{min} = 10^{-3}$.

Furthermore, the process was repeated on 500 random subsets of random size from each camera pair. The number of points in each subset were taken from a uniform distribution between 8 and n_{max} , where n_{max} is the number of homologous points for the specific camera pair. The same points and starting approximations were given to all algorithms.

7 RESULTS

The optimisation results using all homologous points for each camera pair are shown in Table 2. Based on the convergence results, the cases may be classified as follows:

Table 2: Number of iterations used by the four algorithms for each camera pair. Failure to converge in 20 iterations is indicated by a dash. σ_0 was below 11 microns for all converged cases.

Pair #	Stations	Number of iterations			
		Alg. 1U	Alg. 1D	Alg. 2U	Alg. 2D
1	5, 14	—	12	8	7
2	5, 13	—	9	—	9
3	13, 14	4	5	4	5
4	4, 25	3	3	3	3
5	14, 15	3	3	3	3
6	28, 31	5	7	5	7
7	26, 27	4	4	4	4
8	15, 28	6	6	6	6
9	15, 31	4	4	4	4
10	4, 26	3	3	3	3
11	4, 5	4	4	4	4
12	14, 28	4	4	4	5
13	14, 31	4	4	4	4
14	4, 30	4	4	4	4
15	4, 29	4	4	4	4
16	5, 29	4	4	4	4
17	5, 30	3	3	3	3
18	24, 25	11	6	10	6
19	4, 24	4	4	4	4
20	25, 26	5	5	5	5
21	4, 27	3	3	3	3
22	29, 30	—	—	—	—
23	27, 28	4	4	4	4
24	5, 12	3	3	3	3
25	12, 13	—	—	—	—
26	12, 14	4	4	4	4
27	26, 28	4	4	4	4
28	13, 15	—	8	—	8
29	24, 26	4	4	4	4
30	5, 15	4	4	3	3
31	4, 28	4	4	4	4
32	12, 15	—	8	—	9
33	25, 27	4	4	4	4
Total # failures		6	2	5	2

1. Damped algorithms converge, undamped fails. Case 2, 28, 32.
2. Damped algorithms use fewer iterations. Case 18.
3. Undamped algorithms use fewer iterations. Case 6.
4. All algorithms differ. Case 1.
5. All algorithms fail. Case 22, 25.
6. Small differences. All other cases.

Examples from classes 1–4 are illustrated in figures 2–5. In Class 1 (Figure 2), the undamped algorithms overshoot the target in the first iteration and never recovers. The damped algorithms take smaller steps (initially down to $\alpha = 1/16$) but converge. In Class 2 (Figure 3), the undamped algorithms overshoot the target, but converge, although slower than the damped algorithms. In Class 3 (Figure 4), the undamped algorithms are faster than the damped algorithms. Finally, in Class 4 (Figure 5), the constrained algorithms are better than their unconstrained versions, with faster




Figure 3: Algorithm behaviour for Case 18, all points. The undamped algorithms overshoots the target, but converges. The damped algorithms take shorter initial steps and converge faster.

convergence for Algorithm 2U than 1U and convergence for Algorithm 1U despite failure for Algorithm 1D.

Further analysis of the Class 5 cases reveal that in case 22, the starting approximation algorithm suggest that the two cameras were 45 cm apart and that one point was *behind* both cameras. Modifying the starting approximation by mirroring that point through the midpoint of the baseline produced convergence in 9 iterations for the damped algorithms, but still failure for the undamped algorithms. Increasing the maximum number of iterations revealed that in case 25, the damped algorithms converge in 22 iterations.

The optimisation results for the 500×33 subset optimisations are shown in Table 3. The results for each case are similar to their “all point” cases, but with larger differences between individual cases. On average, the constrained algorithms are marginally better than the unconstrained algorithms. The damped algorithms have equal or lower failure percentage for all cases, but the difference varies substantially. In case 1, 90% of the failures are eliminated by the




Figure 4: Algorithm behaviour for Case 6, all points. The damped algorithms use more iterations than the undamped.

damped algorithms; in case 25 only 5%.

In the 12865 cases where all algorithms converged, the average number of iterations used by each algorithm were 4.41 (Alg. 1U), 4.52 (Alg. 1D), 4.38 (Alg. 2U), and 4.52 (Alg. 2D).

No case of convergence towards a matrix M with determinant -1 was detected.

8 CONCLUSIONS

When a non-linear problem is solved as a sequence of linearised problems, the convergence properties will be affected by how good each linear approximation is of the original problem. Any linearisation of a continuous function is “good” in a neighbourhood of the approximation point. However, the size of the neighbourhood varies. If the problem is highly non-linear, there is a risk that the updates calculated by solving equation (3) or (5) are outside this neighbourhood, which may lead to slow convergence or failure.

In this paper, two parameterisations of the rotation matrix were compared with respect to their effect on the conver-

Table 3: Percentage failure for each algorithm on 500 point subsets of each camera pair.

Pair #	Stations	Failure [%]			
		Alg. 1U	Alg. 1D	Alg. 2U	Alg. 2D
1	5, 14	68	6	60	4
2	5, 13	60	22	58	23
3	13, 14	4	3	4	3
4	4, 25	10	3	9	3
5	14, 15	5	5	5	5
6	28, 31	7	4	7	3
7	26, 27	3	2	3	2
8	15, 28	14	4	14	4
9	15, 31	12	3	12	3
10	4, 26	2	1	2	1
11	4, 5	5	3	5	3
12	14, 28	8	3	9	3
13	14, 31	9	4	9	4
14	4, 30	3	1	3	1
15	4, 29	6	2	6	2
16	5, 29	4	3	5	3
17	5, 30	3	1	3	1
18	24, 25	16	3	18	3
19	4, 24	6	4	7	4
20	25, 26	3	1	3	1
21	4, 27	3	2	4	2
22	29, 30	77	53	77	53
23	27, 28	22	6	22	6
24	5, 12	15	8	13	9
25	12, 13	94	90	95	90
26	12, 14	23	11	23	11
27	26, 28	14	5	14	6
28	13, 15	87	16	87	16
29	24, 26	14	11	14	11
30	5, 15	15	11	15	11
31	4, 28	11	7	11	7
32	12, 15	67	18	66	19
33	25, 27	18	7	17	7
Average		21.5	9.9	21.2	9.8

gence properties of the bundle adjustment algorithm. The hypothesis was that since the rotation matrix parameterisation is less non-linear than the Euler angle parameterisation, the linear approximations used by e.g. bundle adjustment would be better, and the convergence properties would improve.

In addition to the two parameterisations, the *line search* technique was optionally applied. The line search technique is a different approach to reducing the risk of adding updates that are too large. Different line search techniques exist, but the common denominator is that a *fraction* $\alpha \leq 1$ of the update is added. The fraction is determined such that the new parameter approximations are better than the current.

The results indicate that the re-parameterisation hypothesis holds, but that the difference is probably too small to justify using it, except to avoid singularities. However, the results show that using line search, the number of convergence failures on average is reduced by 54% at the insignificant extra cost of 0.15 extra iterations. Thus, we conclude that a technique such as line search should be part of any standard implementation of bundle adjustment.







Figure 5: Algorithm behaviour for case 1, all points. Both damped algorithms converge, as does the undamped constrained algorithm.

REFERENCES

Björck, A., 1996. Numerical methods for least squares problems. SIAM.

Börlin, N., Lindström, P. and Eriksson, J., 2003. A globally convergent Gauss-Newton algorithm for the bundle adjustment problem with functional constraints. In: A. Gruen and H. Kahmen (eds), Optical 3-D Measurement Techniques VI, Vol. 2, Wichmann-Verlag, pp. 269–276.

Gill, P. E., Murray, W. and Wright, M. H., 1981. Practical Optimization. Academic Press.

Hägglund, H., 2004. Photogrammetric camera calibration and constrained optimization. Master's thesis, Department of Physics, Umeå University.

Hartley, R. I. and Zisserman, A., 2000. Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521623049.

Mikhail, E. M., Bethel, J. S. and McGlone, J. C., 2001. Introduction to Modern Photogrammetry. Wiley.

Nash, S. G. and Sofer, A., 1996. Linear and Nonlinear Programming. McGraw-Hill.

Rottensteiner, F., Grussenmeyer, P. and Geneva, M., 2001. Experiences with the digital photogrammetric program package ORPHEUS based on CIPA's "Zurich city hall" dataset for architectural photogrammetry. In: XVIII CIPA International Symposium, Potsdam, Germany. and Int. Archives of Photogrammetry and Remote Sensing, ISSN: 1682-1750, Vol. XXXIV, Part 5/C7, pp. 639-646.

Streilein, A., Grussenmeyer, P. and Hanke, K., 1999. Zurich city hall — A reference data set for digital close range photogrammetry. In: XVII CIPA International Symposium, Olinda, Brazil. and Int. Archives of Photogrammetry and Remote Sensing, ISSN 0256-1840 Vol. XXXII Part 5C2 pp. 104-107.