BUNDLE ADJUSTMENT WITH GEOMETRIC CONSTRAINTS FOR HYPOTHESIS EVALUATION

Chris McGlone
Digital Mapping Laboratory
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891 USA
email: *jcm@cs.cmu.edu*

Commission III, Working Group 2

KEY WORDS: geometric model, reliability, bundle adjustment, statistics, image understanding

ABSTRACT

This paper describes the use of a bundle adjustment with geometric constraints to evaluate feature matches and geometric assumptions, based on the use of reliability statistics. Our evaluation procedure starts with the smallest possible redundant geometric subsets of the object, (e.g., one right-angle corner of a rectangular building), finds consistent feature matches among them, and then combines the consistent matches into larger geometric entities and continues the evaluation. We derive statistics to verify the geometric assumptions used in the solution and demonstrate their use to detect erroneous geometric constraints.

1 MOTIVATION

Our goal is the automated and semi-automated construction of accurate three-dimensional site models from multiple images. Several of our site modeling systems, such as SiteCity [Hsieh, 1996; Hsieh, 1995], MultiView [Roux and McKeown, 1994a; Roux *et al.*, 1995], and PIVOT [Shufelt, 1996a], use a hypothesize-and-test paradigm, in which a large set of object hypotheses is generated, then evaluated to extract the best descriptions of cartographic features in the scene. Evaluation methods have typically been based on image intensity, edge geometry, or shadow identification [Shufelt and McKeown, 1993] or on image geometry [McGlone and Shufelt, 1993].

Methods using image and scene geometry are especially powerful, since the geometry is independent of image intensity properties and can be established from information sources outside the image. We have been developing techniques to utilize the geometric information more fully, built around a bundle adjustment with object-space geometric constraints and which calculates a complete set of evaluation statistics. While the use of geometric constraints in a bundle adjustment is not a new technology, most previous applications have been concerned with the determination of sensor orientations [Mikhail, 1970; McGlone and Mikhail, 1982]. We intend to use the geometrically-constrained bundle adjustment to provide a rigorous evaluation procedure, both for feature matches and for object-space geometric assumptions.

The distinction between feature match verification and geometric hypothesis verification is an important one. Verification of a match means that the features identified on each image all correspond to the same physical point in object space. On the other hand, verification of a geometric hypothesis means that the features involved, whether they are lines, points, or planes, or other shapes, actually have the specified object-space configurations. For instance, four object space points may be correctly identified on all images, but may be incorrectly specified to be coplanar.

The verification of point matches is equivalent to the testing of point measurements for blunders, using standard data snooping techniques [Förstner, 1985]; the addition of geometric constraints makes such testing much more effective [McGlone, 1995]. However, the evaluation of geometric hypotheses, as expressed by geometric constraints in the bundle adjustment, requires the extension of standard statistical techniques. The derivation of such statistics is discussed in the next section. The following sections then discuss the application of reliability statistics and geometric constraints to evaluating feature matches and to evaluating the reliability of the geometric information itself.

2 CONSTRAINT RELIABILITY

A large body of work in recent years has been concerned with the derivation and interpretation of reliability statistics for point measurements in traditional photogrammetric applications [Förstner, 1985]. These techniques can be applied in a straightforward manner to the matching of point features, where corresponding images of an object must be identified across multiple images.

However, our work involves extracting and modeling complex objects, using their geometry to help us in the process. We want to use this geometric information to assist in the rigorous evaluation of our building hypotheses. As described below, the unified approach to least squares adjustment gives us this capability.

2.1 Mathematical basis

The detection of blunders is based upon the examination of residuals; therefore, we want to quantify the effect of a blunder on the residuals from an adjustment. For the classical case of least squares adjustment, the partial derivative of the residuals with respect to the input observations is [Förstner, 1987]:

$$\Delta v = -\mathbf{R}\,\Delta y = -\mathbf{Q}_{vv}\,\mathbf{W}\,\Delta y \qquad (1)$$

The **R** matrix in equation 1 is called the redundancy matrix, since the trace of the matrix is equal to the redundancy of the adjustment. The diagonal elements of the matrix, r_{ii}, whose values are between 0 and 1 since the matrix is idempotent, indicate the portion of an error in an observation visible in

the residual for that observation. The higher the r_{ii} for a particular observation, the better the probability of detecting an error in that observation. The off-diagonal terms of the $\mathbf{r}_t$ matrix, r_{ij}, show the portion of an error in observation j that appears in the residual for observation i.

In the unified approach to least squares adjustment all quantities, including parameters and constraint values, are treated as observations. This enables the use of *a priori* weights to more closely model the information input to the adjustment—parameters are seldom completely unknown and in fact may be known very well, while constraints are seldom absolutely true. This also allows all equations to be treated consistently as observation equations, which we will exploit in our testing of constraint and parameter residuals.

The basic condition equations for the case of unified adjustment are [Mikhail, 1980]:

$$\begin{bmatrix} \mathbf{A} & 0 & 0 \\ 0 & \mathbf{A}_c & 0 \\ 0 & 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} v \\ v_c \\ v_x \end{bmatrix} + \begin{bmatrix} \mathbf{B} \\ \mathbf{C} \\ -\mathbf{I} \end{bmatrix} \Delta = \begin{bmatrix} f^0 \\ f_c^0 \\ f_x \end{bmatrix} \quad (2)$$

$$\mathbf{A}_t v_t + \mathbf{B}_t \Delta = f_t \quad (3)$$

The first row of equation 2 refers to the condition equations in the traditional sense; the second row, with the "c" subscripts, is the constraint equations, while the third row is the observation equations on the parameters. v, v_c, v_x are the residual vectors for each type of observation, $\mathbf{A}$ and $\mathbf{A}_c$ are the partial derivatives of the condition equations with respect to the observations, Δ is the vector of corrections to the parameters, and $\mathbf{B}$ and $\mathbf{C}$ are the partials of the equations with respect to the parameters. The vectors l^0, l_c^0, x^0 are approximations for the image observations, the constraint observations, and the parameters respectively, while the discrepancy vectors are f^0, f_c^0, and f_x. Each class of observation has its own covariance matrix, Q, Q_{cc}, and Q_{xx}.

The normal equations are:

$$\left[\mathbf{B}^t\mathbf{W}_e\mathbf{B} + \mathbf{C}^t\mathbf{W}_{ec}\mathbf{C} + \mathbf{W}_{xx}\right]\Delta = \left[\mathbf{B}^t\mathbf{W}_e f^0 + \mathbf{C}^t\mathbf{W}_{ec} f_c^0 - \mathbf{W}_{xx} f_x\right] \quad (4)$$

or,

$$\mathbf{N}\Delta = t \quad (5)$$

with the equivalent weight matrices $\mathbf{W}_e = \left(\mathbf{A}\mathbf{Q}\mathbf{A}^t\right)^{-1}$ for the image measurements and $\mathbf{W}_{ec} = \left(\mathbf{A}_c\mathbf{Q}_{cc}\mathbf{A}_c^t\right)^{-1}$ for the constraint observations, and weight matrix W_{xx} for the parameter observations.

Each class of "observations", the original measurements, the parameters, and the constraints, has a residual and a residual cofactor associated with it.

The residuals are:

$$\begin{aligned} v &= \mathbf{Q}\mathbf{A}^t\mathbf{W}_e\left(f^0 - \mathbf{B}\Delta\right) \\ v_c &= \mathbf{Q}_c\mathbf{A}_c^t\mathbf{W}_{ec}\left(f_c^0 - \mathbf{C}\Delta\right) \\ v_x &= f_x + \Delta \end{aligned} \quad (6)$$

and $\mathbf{R}_t$, the total redundancy matrix, is:

$$\mathbf{R}_t = -\mathbf{Q}_{vv_t}\mathbf{W}_t \quad (7)$$

$$= \begin{bmatrix} \mathrm{I} - \mathrm{B}\mathrm{N}^{-1}\mathrm{B}^t\mathrm{W} & -\mathrm{B}\mathrm{N}^{-1}\mathrm{C}^t\mathrm{W}_c & \mathrm{B}\mathrm{N}^{-1}\mathrm{W}_{xx} \\ -\mathrm{C}\mathrm{N}^{-1}\mathrm{B}^t\mathrm{W} & \mathrm{I} - \mathrm{C}\mathrm{N}^{-1}\mathrm{C}^t\mathrm{W}_c & \mathrm{C}\mathrm{N}^{-1}\mathrm{W}_{xx} \\ \mathrm{N}^{-1}\mathrm{B}^t\mathrm{W} & \mathrm{N}^{-1}\mathrm{C}^t\mathrm{W}_c & \mathrm{I} - \mathrm{N}^{-1}\mathrm{W}_{xx} \end{bmatrix}$$

Using the calculated residuals and cofactors for the total set of observations, we can perform the same statistical testing on constraint equations and parameter values that we do on the image measurements. We can examine constraint observation residuals and decide whether the constraint equation is valid, and we can test *a priori* parameter values to ensure that the input covariance information was not too optimistic.

We can also examine the terms of the $\mathbf{R}_t$ matrix to determine how much of an error will appear in any particular observation. In the classical approach this matrix is used to quantify the effects of errors in the image measurements on the image measurement residuals. In our extended application, we can also study the effects of errors in constraint values on the constraint residuals (i.e., constraining the distance between two points to the wrong value) or of errors in parameter approximations (inputing a bad approximation with a too-small variance) on the final parameter values. By examining the off-diagonal terms in $\mathbf{R}_t$ we can look at the interactions between types of observations—for instance, how much of an error in an image measurement appears in the value of a constraint residual.

Of course, the same considerations in testing residuals apply as in the classical case [Förstner, 1994]. The ability to isolate bad observations is completely dependent upon the redundancy and the geometric strength of the adjustment. While adding constraints increases both the redundancy and the geometric strength, it may still be the case that bad observations cannot be unambiguously identified by examination of individual residuals. Constraint equations typically introduce more correlation between the residuals; if the determining geometry of the other observations is not strong enough, a blunder will be distributed among all the constrained points.

2.2 Effects of constraint formulation and redundant constraints

Previous work [McGlone, 1995] has shown differences in the reliability characteristics of constrained solutions, depending on the form in which the constraint equations were expressed. Specifically, a rectangular building was modeled as a scaled rectangular prism and also by a set of coplanarity and right angle constraints. The solution using the rectangular prism model was much more resistant to blunders in the corner points than was the solution utilizing the combination of constraints. Further investigation has shown that this is a function of the constraint equation weights; when the solution emulates the classical technique and uses very high constraint weights (small standard deviations), the two solutions are exactly equivalent. As the constraint standard deviations are relaxed to values which are more physically realistic, the two types of solutions begin to exhibit different behaviors.

This behavior also sheds light on the question of using redundant constraints. In the classical case redundant constraints lead to a singular solution, since the equations are dependent. For instance, constraining all four angles of a rectangle to be right angles would be singular, since the fourth angle can be determined from the values of the other three. In the unified

approach, however, the constraint equations are actually observation equations; adding redundant observation equations actually improves the solution.

3 EVALUATION OF FEATURE MATCHING USING GEOMETRIC INFORMATION

One of the most common problems in computer vision is the matching of point features between multiple images. Corner detectors or interest operators generate a large number of features, often similar in appearance; this forces a choice between accepting only the most obvious matches, and thereby discarding the majority of the features, or dealing with matches containing a large number of errors.

A simple strategy for evaluating feature matches would be to perform a bundle adjustment for each plausible combination of potential feature matches, with assumed object space geometry modeled by geometric constraints. By examining the statistics for each adjustment, we could then decide which feature matches are blunders and which are valid.

The obvious problem with this strategy is combinatorics. For the evaluation to be effective we need redundant solutions, i.e., at least three image rays for each object point, at least four points to determine a general plane, etc. However, to obtain sufficient redundancy we must process the many possible feature matching combinations. For instance, if a given object point has three possible matches on three images, 27 solutions would have to be run. If we add another image with another three possible matches to obtain better redundancy, 81 solutions would be required. Adding geometric constraints between features only makes the combinatorics worse, since we must test all possible matches of all features involved.

Our solution to this dilemma is to work with the smallest possible redundant subsets—the smallest geometric configuration which can be redundantly specified with the available features. In this example, the redundant geometric subsets are the right angles at each building corner. The 3D coordinates of each point defining the right angle are determined by the intersection of the image rays; constraining the 3D points lie at a right angle within a horizontal plane means that their positions are redundantly determined. This redundancy, or extra information, allows us to evaluate the point matches.

Each subset is solved for every combination of feature matches, and feature matches which do not form any consistent subsets are eliminated. Feature matches which are part of consistent subsets are used to form larger subsets. The process is repeated until the final solution is obtained, and, ideally, only one consistent set of match hypotheses is left.

The computational savings due to this decomposition depends upon the number of possible feature matches compared to the number of subsets used. For example, if each point has 4 possible matches on each image, each subset solution will require 4^3 solutions. Since there are four subsets of the total solution, a total of 4^4 subset solutions will be required. Doing the complete geometric solution would require 4^4 separate solutions. The subset solutions are smaller (fewer parameters and points) and therefore less expensive than the complete solution, but we must still do some number of complete solutions after editing points using the subset solutions. We would still prefer to do the subset solutions over doing the complete solution, simply for the reason that the subsets are simpler to edit and understand since fewer features are involved. Also, we do not necessarily have to do every possible geometric subset, as long as all features are included in some subset.

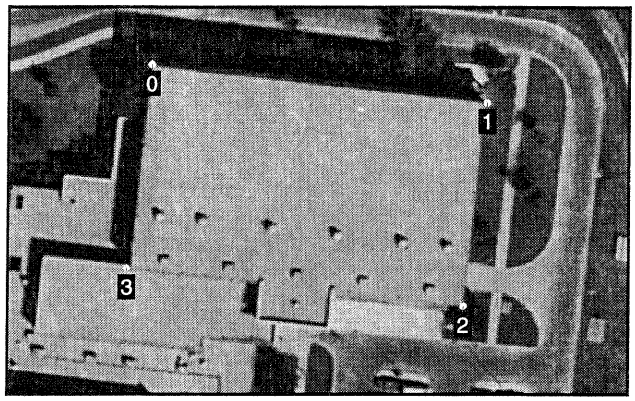

Figure 1: Image fhn715 showing building corners to be matched.

This section describes the application of this technique to point matching, although it can be applied to matching edges or even combinations of edges. In the following examples we use three images, two vertical (fhn715, fhn717) and one oblique (fhov1627), taken over Ft. Hood, Texas, as part of the RADIUS program [Gee and Newman, 1993]. We start with four building corners in fhn715, shown in Figure 1. Figures 2 and 3 show the initial set of corners extracted by the BUILD [Shufelt and McKeown, 1993] corner finder for images fhn717 and fhov1627, along with the epipolar lines corresponding to the points of interest on fhn715.

We first filter the corners using criteria such as epipolar search bounds (propagated from the image orientation covariance), point elevation ranges, corner direction, [Roux and McKeown, 1994b; Roux and McKeown, 1994a], or corner type determined by vanishing point analysis [Shufelt, 1996b; McGlone and Shufelt, 1993]. Next, potentially matching image points are intersected and their standardized residuals are examined to eliminate bad matches. The results of the these two steps are shown in Figures 4 and 5. The smaller dots are the points remaining after preliminary filtering using the epipolar bounds, an expected elevation range, and the corner direction, while the larger dots represent points which have apparently valid matches. Note that no valid corners survived the filtering for point 3 on image fhov1627.

At this point, we introduce the object space geometry into the evaluation. Since we are looking for the horizontal roof of a rectangular building, we can apply constraints forcing the four hypothesized roof corners to lie in a horizontal plane and to form right angles. However, applying the constrained solution to each possible combination of potentially matching points would require an impractical number of solutions.

Taking the potential matches for three out of the four corner points at a time, we perform solutions constraining the three points to form a right angle in a horizontal plane. Points in consistent subsets are flagged for incorporation in the overall solution, which includes all combinations of potential matching points for all four corners of the building. Figures 6 and 7 show the points remaining after the right-angle constraints solution as small dots, and the final points as larger dots. Note that point 2 on image fhov1627 has two possibilities,

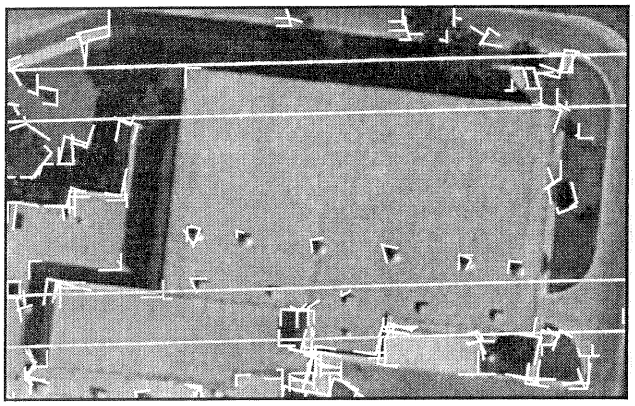

Figure 2: Image fhn717, showing all detected corners and epipolar lines with image fhn715

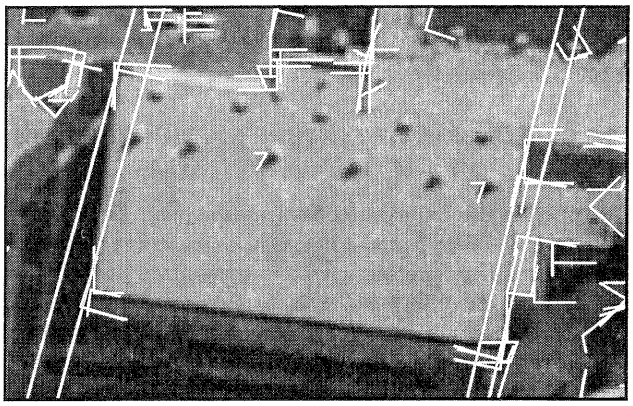

Figure 3: Image fhov1627, showing all detected corners and epipolar lines with image fhn715

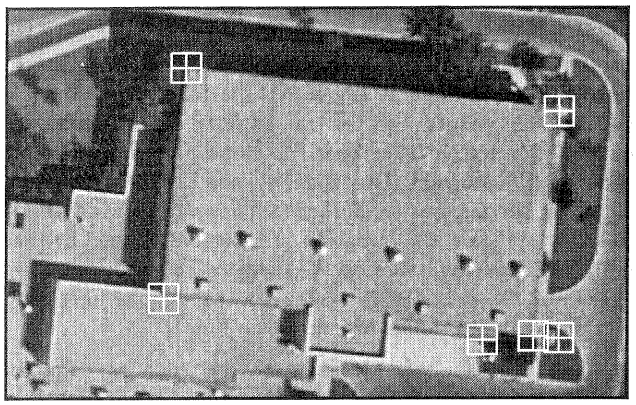

Figure 4: Image fhn717, showing corners after preliminary filtering (crosses) and the point match step (squares). For this image, they were the same points.

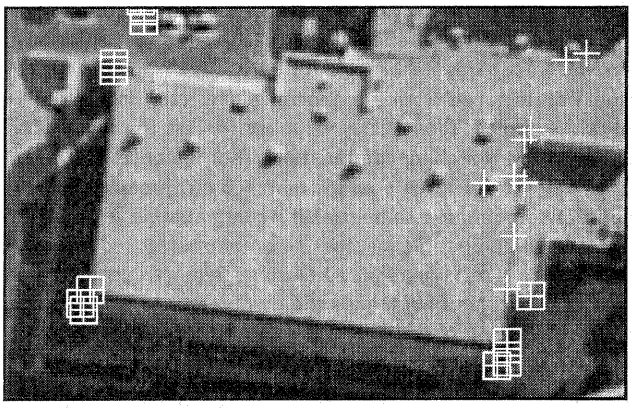

Figure 5: Image fhov1627, showing corners after preliminary filtering (crosses) and points after the point match step (squares).

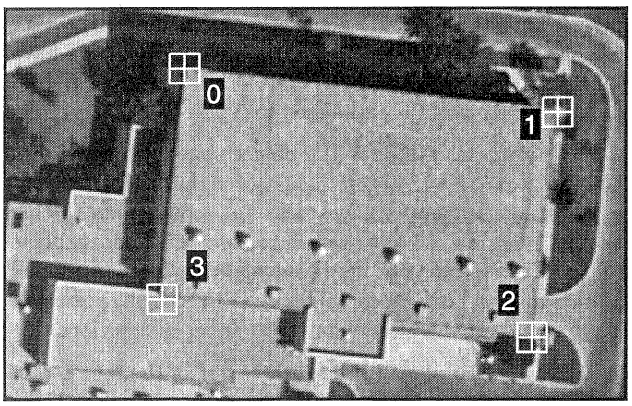

Figure 6: Image fhn717, showing points after the right angle tests (crosses) and final point matches (squares). For this image they were the same points.

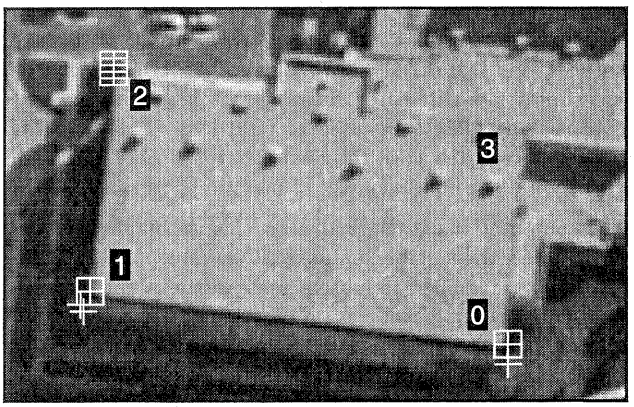

Figure 7: Image fhov1627, showing points after the right angle tests (crosses), and final point matches (squares).

about two pixels apart, remaining after the final step.

It is important to remember that hypothesis evaluation based on image or scene geometry is strictly limited by the available geometry. In this case, the oblique image was taken in a direction perpendicular to the baseline of the two vertical images, giving nearly optimal intersection geometry and therefore eliminating many ambiguous matches. A corresponding test case, run with a vertical image (fhn713) from the same flight line as fhn715 and fhn717, resulted in a large number of hypotheses remaining after the geometric verification procedure. Since the epipolar lines between the three images from the same strip were nearly parallel, many matches appear to be geometrically consistent. This problem can be alleviated somewhat by using strips with 60% sidelap.

The experimental results given above were run with the image parameter covariances set to very small values and therefore not allowed to adjust, with the intent being to prevent the image parameters from absorbing any of the matching errors. The same set of experiments was run with the image parameter covariances as determined from the original block adjustment to see what effect it might have. This made no significant difference in the results for the example described above. The runs with realistic image covariances produced one additional point match which passed the individual right angle evaluation, but the final point match selections were the same for both cases. This is significant in terms of computational expense since it implies that instead of a full simultaneous solution incorporating weighted image parameters, only the points and constraints would need to be included. This possibility should be verified for more test cases.

4 EVALUATION OF GEOMETRIC HYPOTHESES

The previous section has shown how the use of geometric information improves the detection of bad feature matches. However, we must also be able to evaluate these geometric assumptions. It is entirely possible that, although we have matched a feature correctly across several images, the corresponding object-space point may not be the building corner that we have assumed it is.

By using the statistics derived in Section 2, we can evaluate the constraint equation residuals and determine whether the geometric conditions are being met. A problem is separating the effects of bad feature matches from bad geometric conditions; a bad constraint typically results in high image residuals for all of the features involved.

Our approach of starting from minimal subsets and building up to final complete solution provides an answer for this. As described above, point or feature matches are first evaluated using only the image information, without added geometry. As individual features are combined into small geometric units and constraints added to the solution, a bad geometric assumption will affect only some of the solutions.

For instance, suppose that in the point matching example discussed above, point 2 had been incorrectly identified at the roof vent nearest the actual building corner and that corresponding points had been matched on the other images (Figure 8). While we would have a perfectly valid 3D point it is not a roof corner, it is not coplanar with the other corners, nor will it form right angles with them.

The evaluation of the point match shows that this is indeed a valid match. However, during the next stage, forming right-angle corners with subsets of three points, the three subsets involving point 2 show up as bad. The statistics are summarized in Table 9, where the image residuals given are the root-mean-square of each point's vector standardized residuals on each image. Subsets 0-1-2 and 1-2-3 both have bad point standardized image residuals and bad constraint residuals; subset 2-3-1 has good image residuals but a bad constraint residual on the angle constraint. This can be explained by noticing that, for this solution, the bad point is on the long side of the building so that its angular offset from the correct location is less than for the other cases, where it was on the short side of the angle or was the central point.

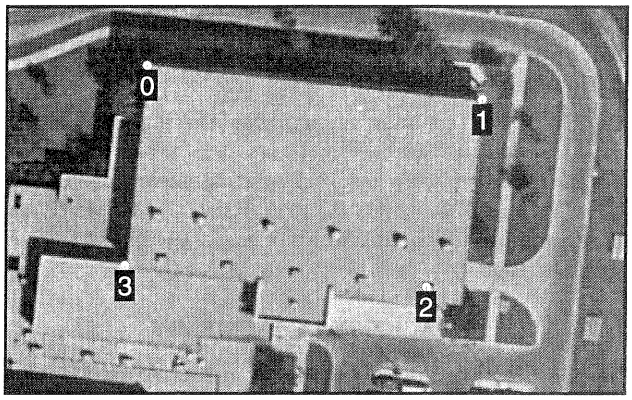

Figure 8: Image fhn715 with bad point 2.

corner points	0-1-2	1-2-3	2-3-0	3-0-1
pt 0	6.2	–	2.0	0.5
pt 1	11.4	11.4	–	0.7
pt 2 (bad)	9.9	13.8	1.6	–
pt 3	–	7.8	2.6	1.0
right angle constraint	21.6	25.2	4.3	0.3

Table 9: Standardized residuals from individual corner solutions with bad point.

Systematic detection of bad geometry requires searching the bad subsets for the common elements; in this case, point 2 is the common element among the three bad subsets. A bad subset must be identified by examining the constraint residuals also; the table shows that the image residuals for angle 2-3-0 were satisfactory, with only the constraint residual bad, due to the location of the bad point.

5 CONCLUSIONS

Using geometric constraints in a bundle adjustment to verify feature matches or geometric hypotheses can be part of an effective hypothesis testing strategy, when supported by the image geometry and when knowledge or inferences of the scene geometry are available. The combinatoric problems inherent in the method can be avoided by early editing of the hypotheses and by utilizing sequential evaluations of minimal redundant geometric subsets.

Use of these methods in full-sized systems will require that the computations be optimized as much as possible. While the tests in this paper were done using a general-purpose simultaneous orientation solution, a real application should use special purpose routines, each optimized to verify a particular geometric configuration such as right angles, coplanarity, etc.

There appears to be no need to include the image parameters in the calculation, although this should be verified with further tests.

Our work will continue on applying these techniques to more complicated geometric objects such as buildings, and on incorporating them into our existing site modeling systems, such as SiteCity [Hsieh, 1996; Hsieh, 1995], MultiView [Roux and McKeown, 1994a; Roux *et al.*, 1995], and PIVOT [Shufelt, 1996a].

6 ACKNOWLEDGEMENTS

I would like to thank my colleagues in the Digital Mapping Laboratory for their support and interaction.

Digital Mapping Laboratory publications are available at http://www.cs.cmu.edu/~MAPSLab

The research reported in this paper was supported by the Advanced Research Projects Agency (ARPA/ISO) and the U.S. Army Topographic Engineering Center under contracts DACA76-91-C-0014, DACA76-92-C-0036, and DACA76-95-C-0009. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Topographic Engineering Center, the Advanced Research Projects Agency, or of the United States government.

REFERENCES

[Förstner, 1985] W. Förstner. The reliability of block triangulation. *Photogrammetric Engineering and Remote Sensing*, 51(8):1137–1149, 1985.

[Förstner, 1987] W. Förstner. Reliability analysis of parameter estimation in linear models with applications to mensuration problems in computer vision. *Computer Vision, Graphics, and Image Processing*, 40:273–310, 1987.

[Förstner, 1994] W. Förstner. Diagnostics and performance evaluation in computer vision. In *Proceedings, NSF/ARPA Workshop on Performance versus Methodology in Computer Vision*, pages 11–25, 1994.

[Gee and Newman, 1993] S. Gee and A. Newman. RADIUS: Automating image analysis through model-supported exploitation. In *Proceedings of the DARPA Image Understanding Workshop*, pages 185–196, Washington, D. C., April 19–21 1993. Morgan Kaufmann Publishers, Inc.

[Hsieh, 1995] Y. Hsieh. Design and evaluation of a semi-automated site modeling system. Technical Report CMU–CS–95–195, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, November 1995.

[Hsieh, 1996] Y. Hsieh. Design and evaluation of a semi-automated site modeling system. In *Proceedings of the ARPA Image Understanding Workshop*, Palm Springs, California, February 1996. Advanced Research Projects Agency, Morgan Kaufmann Publishers, Inc.

[McGlone and Mikhail, 1982] Chris McGlone and Edward M. Mikhail. Geometric constraints in multispectral scanner data. In *Proceedings of the 48th Annual Meeting, American Society of Photogrammetry*, pages 563–572. American Society of Photogrammetry, March 1982.

[McGlone and Shufelt, 1993] J. C. McGlone and J. A. Shufelt. Incorporating vanishing point geometry into a building extraction system. In *Proceedings of the DARPA Image Understanding Workshop*, pages 437–448, Washington, D. C., April 19–21 1993. Defense Advanced Research Projects Agency, Morgan Kaufmann Publishers, Inc.

[McGlone, 1995] Chris McGlone. Bundle adjustment with object space constraints for site modeling. In *Proceedings of the SPIE: Integrating Photogrammetric Techniques with Scene Analysis and Machine Vision*, volume 2486, pages 25–36, April 1995.

[Mikhail, 1970] Edward M. Mikhail. Relative control for extraterrestrial work. *Photogrammetric Engineering*, 36(4), April 1970.

[Mikhail, 1980] Edward M. Mikhail. *Observations and Least Squares.* Harper and Row, New York, 1980.

[Roux and McKeown, 1994a] M. Roux and D. M. McKeown, Jr. Feature matching for building extraction from multiple views. In *Proceedings of the ARPA Image Understanding Workshop*, pages 331–349, Monterey, California, November 13–16 1994. Advanced Research Projects Agency, Morgan Kaufmann Publishers, Inc.

[Roux and McKeown, 1994b] M. Roux and D. M. McKeown, Jr. Feature matching for building extraction from multiple views. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 46–53, Seattle, Washington, June 19–23 1994.

[Roux *et al.*, 1995] M. Roux, Y. C. Hsieh, and D. M. McKeown, Jr. Performance analysis of object space matching for building extraction using several images. In *Proceedings of the SPIE: Integrating Photogrammetric Techniques with Scene Analysis and Machine Vision II*, volume 2486, pages 277–297, 1995.

[Shufelt and McKeown, 1993] J. A. Shufelt and D. M. McKeown. Fusion of monocular cues to detect man-made structures in aerial imagery. *Computer Vision, Graphics, and Image Processing*, 57(3):307–330, May 1993.

[Shufelt, 1996a] J. Shufelt. Exploiting photogrammetric methods for building extraction in aerial images. In *International Archives of Photogrammetry and Remote Sensing*, volume XXXI, B3, 1996.

[Shufelt, 1996b] J. Shufelt. Performance evaluation and analysis of vanishing point detection techniques. In *Proceedings of the ARPA Image Understanding Workshop*, Palm Springs, California, February 1996. Advanced Research Projects Agency, Morgan Kaufmann Publishers, Inc.