ISPRS Journal of Photogrammetry & Remote Sensing 52 (1997) 149-159

Invited Review Paper

Digital camera self -calibration

Clive S. Fraser *

Department of Geomatics, The University of Melbourne, Parkville, Vic. 3052, *Australia*

Accepted 24 March 1997

Abstract

Over the 25 years since the introduction of analytical camera self-calibration there has been a revolution in close-range photogrammetric image acquisition systems. High-resolution, large-area 'digital' CCD sensors have all but replaced film cameras. Throughout the period of this transition, self-calibration models have remained essentially unchanged. This paper reviews the application of analytical self-calibration to digital cameras. Computer vision perspectives are touched upon, the quality of self-calibration is discussed, and an overview is given of each of the four main sources of departures from collinearity in CCD cameras. Practical issues are also addressed and experimental results are used to highlight important characteristics of digital camera self-calibration.

Keywords: camera calibration; digital cameras; self-calibration; close-range photogrammetry

1. Introduction

It has now been 25 years since the concept of camera system self-calibration was introduced to the wider photogrammetric community. The mathematical roots of self-calibration were laid down in the context of analytical aerial triangulation, the first experiences took place in close-range photogrammetry, and broad acceptance of the technique accompanied the development of bundle adjustment with additional parameters (APs) for precision aerial triangulation. Since the subject of this paper is digital camera self-calibration, the discussion will be confined to close-range imaging.

In the period since the early 1970s there has been a revolution in the area of image acquisition in close-range photogrammetry. Initially, large-format cameras with glass plates were overtaken by large- and medium-format film cameras with reseau arrays, focusable lenses and vacuum platens, and for medium-accuracy applications, semi-metric 70 mm film cameras became popular. Small-format CCD cameras first attracted photogrammetric attention in the mid eighties, still video cameras were introduced in the early nineties and we are now in the era of high-resolution, large-area CCD cameras with digital output. There have been numerous recent papers which report the attainment of triangulation accuracies surpassing 1:100,000 with digital cameras such as the Kodak Megaplus series, and the Kodak DCS 420 and DCS 460 (e.g. Beyer, 1995; Brown and Dold, 1995; Fraser et al., 1995; Peipe, 1995; Schneider, 1996). Digital imaging sensors are rendering film cameras obsolete for all but the niche domain of extreme high-accuracy photogrammetry (triangulation accuracies surpassing 1:250,000).

It is of interest that, whereas we have witnessed both a revolution in close-range photogrammetric camera technology, and no lack of ongoing research activity into analytical self - calibration, the familiar eight-parameter physical model reported in Kenefick et al. (1972) remains a preferred AP set, the other alternative being the same model supplemented by two further parameters. The eight-parameter model comprises the interior orientation elements of principal distance and principal point offset, and correction terms for radial and decentring distortion. It is a well established fact that perturbations such as film deformation and focal plane unflatness cannot be recovered through self-calibration, and it should therefore come as no surprise that in-plane and out-of-plane distortions in digital cameras do not readily lend themselves to correction through the use of APs - differential image axis scaling and non-orthogonality being notable exceptions to the rule.

The geometric limitations of semi-metric film cameras, especially unstable interior orientation, continue to plague today's digital still video cameras which are designed more for the mass consumer market than for photogrammetry. Attempts to overcome these limitations through analytical techniques are likely to be just as unsuccessful as earlier endeavours to metrically exploit non-metric film cameras. As will be demonstrated later in the paper, interior orientation instability and focal plane unflatness continue to be primary factors limiting the photogrammetric potential of large-area CCD cameras with onboard analog-to-digital (A/D*)* conversion.

The remainder of the paper is in the form of a review, which gives the author some license to add his own views on practical issues. Following a short reference to digital camera calibration from a computer vision perspective, the paper addresses the topic of how good a calibration should be. This is followed by a brief review of the mathematical formulation of self-calibration, and a discussion of the principal sources of departures from collinearity as they relate to digital cameras. Examples of self-calibration results are used to highlight salient characteristics.

2. The computer vision connection

Digital camera self-calibration is not unique to photogrammetry; one now finds reference to this concept in computer vision literature (e.g. Maybank and Faugeras, 1992). Although computer vision researchers could be said to have come 'late to the game' of metric calibration, they have nevertheless come with a vengeance. From their perspective in the mid eighties, camera calibration was "a somewhat neglected field in digital image processing" (Lenz, 1987). From a photogrammetric point of view, digital camera calibration had received only modest attention at the time because (a) it was early days for CCD cameras and (b) the cameras of that era showed less than encouraging photogrammetric potential (e.g. Gruen, 1996). The limitations were more a function of resolution than metric calibration. An ideal calibration technique was defined within the computer vision community as being one which is "autonomous, accurate, efficient, versatile and requires only common off-the-self cameras and lenses" (Tsai, 1986). Photogrammetric calibration processes were seen as less than optimal in pursuit of these goals.

New methods were thus developed for the recovery of intrinsic and extrinsic parameters, which is computer and machine vision jargon for interior and exterior orientation. The common feature of these approaches has tended to be the extraction of camera parameters with a minimum of geometric information. Pre-calibration to any degree seemed to be frowned upon because interior orientation parameters might change through mechanical or thermal effects, as well as through focussing. Whether these perturbations would be of a magnitude to influence the accuracy of object recognition and reconstruction did not receive as much attention as the quest to achieve 'calibration' with, seemingly, as few images and as few image points as algebraically possible. Such methods of course required control arrays comprising object points of known coordinates (Lenz and Tsai, 1988), though this requirement was dispensed with once self-calibration was introduced (Maybank and Faugeras, 1992). Although there is a reasonable degree of correspondence between the critical calibration parameters involved in the computer vision and photogrammetric approaches, there are often practical distinctions between the way these parameters are applied. For example, it is not uncommon in the former case to treat inner orientation parameters as image variant (e.g. carrying a different principal distance unknown for each image), which is generally anathema to photogrammetrists seeking a robust camera calibration.

Contradictory assessments of the importance of calibration can also be found in the computer vision literature. The following two statements are by the same author in the same year: "Camera calibration is an important task in computer vision" (Maybank and Faugeras, 1992), and "computer vision may have been slightly overdoing it in trying at all costs to obtain metric information from images" (Faugeras, 1992). In the second referenced paper the author also adds that it is not often the case that metric information is necessary for robotics applications, which is clearly fortunate given the accuracy limitations of 'self-calibrating' from two- and three-image networks comprising less than ten image point correspondences.

If photogrammetrists were to realistically ask themselves whether the present CCD camera calibration techniques developed in computer vision are beneficial in metric measurement, the answer would have to be no. Even the perceived advantages of speed and on-line processing are no longer valid. In order to automatically self-calibrate a digital camera in an on-line close-range network configuration, the photogrammetrist needs only to collect four or more images of a field of a few tens of distinct targets, with there being no requirement for object space dimensional information. Calibration to a fidelity matching the angular measurement resolution of the photogrammetric camera is then available in near real time (within a few seconds of the last image being recorded). Such fully automated self-calibration procedures are already implemented in commercially available vision metrology systems for industrial measurement (Fraser, 1997). Claims that computer vision inspired approaches such as 'object reconstruction without inner orientation' (a stereo solution for uncalibrated cameras involving only six image points) will find wide use in photogrammetry (Shan, 1996) cannot be given much credence, at least in situations requiring photogrammetric accuracies. The use of orientation techniques with their roots in computer vision is, however, clearly not precluded for preliminary orientation determination, the adoption of the closed-form resection formulation of Fischler and Bolles (1981) being a good example. The following discussion is confined to the 'photogrammetric' self -calibration approach.

3. Quality of self-calibration

At first sight the task of ascertaining the quality of digital camera calibration appears to be less than straightforward. For example, one might consider the question of how accurately interior orientation or decentring distortion needs to be determined to support a triangulation to so many parts per 100,000. The issue is further complicated by the fact that a good deal of projective compensation takes place between the terms forming the AP model, and between the self-calibration parameters and exterior orientation elements. Moreover, the issue of the 'fidelity' of the calibration model depends a good deal on what photogrammetric applications are envisaged for the camera. If one is self-calibrating a camera or cameras which is/are to be used for stereo restitution, then it is probably unwise to recover decentring distortion parameters since very few commercially available digital photogrammetric workstations accommodate such an image correction. Instead, it would generally be better to suppress these parameters and allow part of the component of the error signal to be projectively absorbed by the generally highly correlated principal point offsets (x_o, y_o). Even these parameters may be of limited practical consequence if the stereo model contains little variation in depth.

The more useful approach to examining the quality of calibration involves essentially three simple factors: the distribution of points within the images, the photogrammetric network configuration, and the variance factor or standard error of unit weight of the self-calibrating bundle adjustment. The first item relates specifically to lens distortion which is modelled in terms of polynomial functions that are notoriously poor extrapolators. If a representative distortion modelling (radial and decentring) is sought over the full image format, then the image point distribution must encompass the full image area, albeit not in all images.

Photogrammetric network design considerations for self-calibration are well known, among these being the need for a highly convergent imaging configuration, incorporation of orthogonal camera roll angles, and the use of four or more images (motivated in large part by blunder detection considerations). In addition, there is the preference for an object point field which is well distributed in three dimensions, though a two-dimensional distribution will suffice. In digital camera networks there is usually no reason why high levels of geometric strength and redundancy cannot be employed if sensor calibration is the aim; the image mensuration task is essentially instantaneous and thus no significant increase in workload accompanies the use of, say, twelve images in a single-camera network rather than four. The issue of determinability of self-calibration parameters will not be further discussed in detail in this paper and the reader is referred to, for example, Gruen and Beyer (1992) and Fraser (1992) for reviews of this important issue.

Since the introduction of self-calibration, the ability to fully evaluate the fidelity of the camera calibration parameters has been limited by the accuracy of image coordinate mensuration. The high degree of consistency between the a-priori and a-posteriori estimates of image coordinate precision have been maintained as image mensuration accuracies for distinct target features have increased from around 2 μm for manual measurements on film to 0.02-0.04 pixels for a CCD sensor. If we assume a pixel size of 6 μm, this translates to 0.1-0.2 μm, a better than ten-fold increase in angular measurement resolution. Yet the 'standard' self-calibration model does not yield image coordinate residuals of a magnitude inconsistent with a-priori estimates of measurement precision. The quality of the 'standard' AP model, while satisfying the most stringent practical accuracy demands, has long frustrated attempts to seek incremental improvements in the mathematical models employed for self-calibration. Whereas graphical analysis of image coordinate residuals from bundle adjustment typically indicates the presence of systematic trends, the magnitude of the residuals is all too often 'in the noise'. Having spoken in general terms about the self-calibration model and bundle adjustment with APs, these aspects will now be looked at in more detail.

4. The self-calibration model

The mathematical basis of the self-calibrating bundle adjustment is the well-known extended collinearity model:

$$x - x_0 + \Delta x = -c\frac{R_1}{R_3}$$
$$y - y_0 + \Delta y = -c\frac{R_2}{R_3} \tag{1}$$

where

$$\begin{pmatrix} R_1 \\ R_2 \\ R_3 \end{pmatrix} = R\begin{pmatrix} X - X^0 \\ Y - Y^0 \\ Z - Z^0 \end{pmatrix}$$

These equations describe the perspective transformation between the object space (object point X, *Y*, Z and perspective centre X^0, Y^0, Z^O with rota- tion matrix R*)* and image space (image point *x*, y). The calibration terms in Eq. 1 comprise the principal point offsets x_0, y_0 and the principal distance c (the interior orientation parameters), and the image coordinate perturbation terms Δx and Δy which account for the departures from collinearity due to lens distortion and focal plane distortions. Upon linearisation, the observation equations for the least-squares bundle adjustment are formed. This matrix equation system is given below, supplemented by a constraint function which can be used to impose certain geometric relationships between the parameters of the bundle adjustment:

$$\begin{aligned} A_1x_1 + A_2x_2 + A_3x_3 + w &= 0 \\ H_ix_i + w_h &= 0 \end{aligned} \tag{2}$$

Here, x_1 represents the sensor exterior orientation parameters, x_2 the object point coordinates and x_3 the self-calibration parameters. The A_i matrices are the corresponding configuration matrices and *w* is the image coordinate discrepancy vector. There is nothing to restrict the constraint function from embracing more than one parameter set, but where it is employed it is usually confined to one. Examples might be the use of redundant coordinate, distance or angle constraints in the object space, the employment of geometric constraints on the exterior orientation (as with a theodolite mounted digital camera), the enforcement of prior determined lens distortion values (Shortis et al., 1995), and the imposition of lens distortion variation conditions in multi-sensor self-calibrations of the same camera at different focal settings (Fraser, 1980). Although these and other constraints have been applied over the years, they have not yielded any significant improvements in the recovery of calibration parameters in networks which were of sufficient geometric strength to allow determination of the APs, x_3, in the absence of the constraints. With the greatly enhanced opportunities provided by digital cameras to adopt highly redundant, geometrically strong network configurations, there seems little justification in resorting to such constraints from the point of view of optimising the self-calibration.

In seeking appropriate parameters for the image coordinate correction functions Δx and Δy it is necessary to consider the four principal sources of departures from collinearity which are 'physical' in nature. These are symmetric radial distortion, decentring distortion, image plane unflatness and in-plane image distortion. The net image displacement at any point will amount to the cumulative influence of each of these perturbations. Thus,

$$\Delta x = \Delta x_r + \Delta x_d + \Delta x_u + \Delta x_f$$
$$\Delta y = \Delta y_r + \Delta y_d + \Delta y_u + \Delta y_f \qquad (3)$$

where the subscript r is for radial distortion, d for decentring distortion effects, u for out-of-plane unflatness influences and f for in-plane image distortion. The relative magnitude of each of the four image coordinate perturbations depends very much on the nature of the camera system being employed. The presence of radial lens distortion is usually seen in the form of barrel distortion, decentring distortion is typically small in magnitude, unflatness effects in CCD cameras would arise through chip bowing or the 'crinkling' of thin wafers, and in-plane distortion can be introduced through electronic influences such as clock synchronisation and rate errors, and long-period effects from line jitter (Beyer, 1992).

5. Radial lens distortion

Symmetric radial distortion in analytical photogrammetry is universally represented as an odd ordered polynomial series, as a consequence of the nature of Seidel aberrations:

$$\Delta r = K_1 r^3 + K_2 r^5 + K_3 r^7 \qquad (4)$$

where the *Ki* terms are the coefficients of radial distortion and *r* is the radial distance from the principal point:

$$r^2 = \bar{x}^2 + \bar{y}^2 = (x - x_0)^2 + (y - y_0)^2 \qquad (5)$$

The necessary corrections to the *x*,y image coordinates follow as $\Delta x_r = \bar{x}\Delta r/r$ and $\Delta y_r = \bar{y}\Delta r/r$. The K_1 term alone will usually suffice in medium-accuracy applications of digital cameras with c- or f-mount lenses to account for the commonly encountered third-order barrel distortion. Inclusion of the K_2 and K_3 terms may be warranted for higher-accuracy applications and wide-angle lenses. The decision as to whether to incorporate one, two or three radial distortion terms can be based on statistical tests of significance, though from a practical point of view this is hardly necessary in the presence of a 'strong' self-calibration network. Fortuitously, although the K_i terms are typically highly correlated, their coupling with other exterior orientation and APs in the eight-parameter physical model is low. Thus, overparameterization in this regard rarely leads to numerical difficulties, and still yields a valid radial distortion profile. In such circumstances, multi-dimensional statistical analysis must be employed should an estimate of the precision of radial distortion be sought, but this would have to be a rare occurrence in practise.

There is a projective coupling between the linear component of radial lens distortion and the principal distance which gives rise to an interesting and beneficial feature in the self-calibration of selected CCD cameras. It is not uncommon for these cameras to utilise only a modest portion of the available field of view of the lens, as exemplified by the Kodak DCS series of still video cameras. Hence, any variation within the essentially paraxial, linear section of the distortion curve will be largely compensated by the projective coupling, with the result that the lens will appear to have very little radial distortion.

The radial distortion profile Δr associated with a particular principal distance value c is termed Gaussian distortion. It is well known that radial lens distortion varies both with focussing and within the field of view. Variation with focus due to changing principal distance is of limited consequence in self-calibration so long as the camera is used at a fixed focus. The amount of variation of distortion with object distance is a function of the distortion gradient and this effect can be of metric significance for lenses exhibiting large radial distortion (Fraser and Shortis, 1992). Variation of distortion is most pronounced at high magnifications, within object distances of less than about 15 times the focal length. For a lens of 240 mm focal length on a large-format film camera, this is clearly of concern for many close-range applications call for object distances of less than 3.6 m.

In the case of CCD cameras, however, we face a different situation. For a lens of 10 or 20 mm focal length, the influences of variation of distortion could be expected to be confined to object distances of 15 cm to 30 cm, which should not pose a significant metric problem. To test this assumption the author employed a plumbline calibration technique (Fryer and Brown, 1986) to ascertain the variation of distortion for a 20 mm Nikkor lens mounted on a Kodak DCS420 camera at object distances of 1.5 m, 2.3 m, 3 m and 4.6 m, the focus being set at 1.5 m. The Gaussian distortion profile of the lens is cubic and reaches a magnitude of 85 µm at a radius of 7 mm. At this same radial distance the difference in the profiles for the 1.5 and 4.6 m distances amounted to only 0.4 µm. There was certainly a consistent change in radial distortion between each distance setting, but it reached only 0.2 µm which is of limited metric consequence, even in high-precision measurement applications. For the major portion of the working format of the 20 mm lens, the variation in distortion over a 3 m object distance range was 0.1 µm or less.

The question of stability and recoverability of radial lens distortion through self-calibration is also of interest. In the author's experience, radial lens distortion is the most stable of all calibration parameters (assuming the lens is not refocussed!). To illustrate this point the same 20 mm Nikkor lens referred to above was self-calibrated in a multi-station bundle adjustment some four months after the plumbline calibrations. Although the lens has no 'stops' and was simply re-focussed to 1.5 m, the variation in radial distortion from the plumbline determination was again 0.2 µm or less for all but the outer extremities of the image format. There has been a recently reported investigation which indicated that a 20 mm Nikkor lens might display a metrically significant variation of distortion with object distance, but the evidence was by no means conclusive (Shortis et al., 1996).

Another approach to ascertaining both the repeatability of radial lens distortion and its independence from other parameters in a self-calibration adjustment is via the technique of multi-sensor system self-calibration whereby a number of digital cameras are calibrated simultaneously. One recent application of this technique involved three lenses and two DCS camera bodies, or six different camera/lens combinations in all. Highly repeatable (sub-micrometre) Gaussian radial distortion profiles were obtained for 20 mm and 28 mm Nikkor lenses, demonstrating the very low degree of camera-body specific projective coupling between distortion and interior orientation parameters. For a wider angle lens of 15 mm focal length, the variation between the two radial distortion profile determinations was higher, averaging 1 µm and reaching 4 µm at the extremity of the field of view. One explanation for this, so far unverified, is that an element of focal plane unflatness was being projectively absorbed by the Gaussian distortion profile of the wider angle lens.

6. Decentring distortion

A lack of centring of lens elements along the optical axis gives rise to a second category of lens distortion which has metric consequences in anaIytical restitution, namely decentring distortion. The misalignment of lens components causes both radial and tangential image displacements which can be modelled by correction equations due to Brown (1966):

$$\begin{aligned} \Delta x_d &= P_1(r^2 + 2\bar{x}^2) + 2P_2\overline{xy} \\ \Delta y_d &= 2P_1\overline{xy} + P_2(r^2 + 2\bar{y}^2) \end{aligned} \qquad (6)$$

A useful means of representing the magnitude of decentring distortion is via the profile function *P(r)* which is obtained from the parameters P_1 and P_2 as follows:

$$P(r) = (P_1^2 + P_2^2)^{1/2} r^2 \qquad (7)$$

The maximum magnitudes for the radial and tangential components of decentring distortion are then obtained as 3P(r) and P(r), respectively. For lenses employed with digital cameras, the magnitude of decentring distortion, as determined in a self-calibration, rarely exceeds 10 µm at the extremity of the image format. Decentring distortion also varies with focussing, but the resulting image coordinate perturbations are typically very small and the distortion variation is universally ignored in analytical photogrammetry. In the case of the 20 mm Nikkor I lens examined in the previous section, the decentring profile function reached a value of P(r) = 1.6 µm at a radial distance of 7.5 mm. No measurable variation with object distance was present. There is a strong projective coupling between the decentring distortion parameters P_1 and P_2 and the principal point offsets x_0 and y_0. Correlation coefficient values of up to 0.98 are frequently encountered. This correlation has practical consequences in self-calibration for it means that to a significant extent decentring distortion effects can be compensated for by a shift in the principal point (and an effective tilting of the optical axis). The projective compensation can usually be anticipated with CCD cameras and hence a self–calibration may indicate that the lens be treated as if it were largely free of decentring distortion.

Decentring distortion appears least pronounced in narrow-ang1e lenses, though this may be due more to projective compensation than to small distortion per se. Burner (1995) has discussed the ability of decentring distortion parameters to absorb the error signal, and specifically the perturbation of the photogrammetric principal point position, which arises from the misalignment of zoom lenses. The important practical consequence of this property is that for all but the most stringent accuracy requirements (0.01 pixel level or better), there is little to be gained in precisely aligning lenses such that the optical axis passes through the foot of the perpendicular from the perspective centre to the image plane (i.e. the photogrammetric principal point).

Practical experience (e.g. Fraser et al., 1995) has shown that given a stable interior orientation, and notwithstanding the high level of projective coupling between P_1, P_2 and x_0, y_0, very repeatable results can be obtained for decentring distortion through self-calibration. Although the profile function P(r) may only reach a maximum value approaching half to one pixel, decentring distortion cannot be ignored in high-accuracy close-range digital photogrammetric measurement. For self-calibrations aimed at calibrating a digital camera for stereo restitution, the parameters P_1 and P_2 can usually be suppressed, for the reasons already alluded to in Section 3.

7. Interior Orientation

Interior orientation (IO) instability tends to be the bane of the photogrammetrist seeking to carry out precision measurement with CCD cameras. Whereas cameras such as the Kodak Megaplus series can be rendered 'fully metric' with little effort, still video cameras such as the DCS series make only modest concessions in their design to the photogrammetric requirements for a stable IO. Error sources include movement of the CCD sensor with respect to the camera body, movement of the c-mount lens also with respect to the body, and differential movement of lens elements (which also affects decentring distortion). Instabilities are often exacerbated by the requirement to employ orthogonal roll angles in the self-calibration network. As with decentring distortion, one can expect to find the most pronounced effects of IO instability in convergent, multi-station network configurations; projective compensation operates to a much greater degree in stereo configurations. Unfortunately, there is no analytical solution to unstable IO. The self-calibration process can certainly identify its presence but it cannot rectify the problem.

In regard to IO stability, mixed results have been obtained with DCS still video cameras. Whereas some commercial concerns 'pin' chips to enhance stability, and also stabilize lens elements, the average user is constrained by the knowledge that mechanically tampering with the camera may void the warranty. In the multi-sensor self-calibration reported by Fraser et al. (1995) a high level of IO stability and parameter determination repeatability was obtained for two unmodified still video cameras, a DCS420 and a DCS200. on the other hand, a DCS460 was recently encountered, for which the action of turning the camera upside down caused the chip to move by an estimated 0.3 mm, the mechanical support being extremely unstable. The self-calibration process will indicate, but not adequately quantify IO instability in a strong multi-station network. A practical recommendation in this regard is to employ an object target array which is well distributed in three dimensions. Planar target fields tend to offer a better scope for projective compensation of IO instability; with 3D target fields the resulting degradation of the photogrammetric triangulation is more clearly pronounced.

8. Out-of-plane distortion

Systematic image coordinate errors due to focal plane unflatness constitute a major factor limiting the accuracy of the photogrammetric triangulation process. The induced radial image displacement Δr_u is a function of the incidence angle of the imaging ray. Thus, narrow-angle lenses of long focal length are much less influenced by out-of-plane image deformation than short focal length, wide-angle lenses. It is unfortunate that, due to practical necessities, many vision metrology systems employ wide-angle lenses to achieve a workable field of view in cameras with CCD arrays of small format.

In metric film cameras, focal plane topography can be measured directly, and the induced image coordinate perturbations can be modelled through third- or fourth-order polynomials of the form:

$$\begin{Bmatrix} \Delta x_u \\ \Delta y_u \end{Bmatrix} = \begin{Bmatrix} \bar{x}/r \\ \bar{y}/r \end{Bmatrix} \sum_{i=0}^{n} \sum_{j=0}^{i} a_{ij} \bar{x}^{(i-j)} \bar{y}^{(j)} \qquad (8)$$

The applicability of this approach to CCD matrix arrays is uncertain, however, for after installation within the camera, the CCD chip surface does not often lend itself to direct surface contour measurement. Moreover, information regarding chip topography seems very difficult to come by, especially from manufacturers for whom ‘flat’ implies a much freer tolerance than the micrometre level sought by photogrammetrists. It is even conceivable that the CCD sensor surface may exhibit a degree of planarity that does not warrant any unflatness correction. However, focal plane unflatness should not be ignored; at an incidence angle of 45^o a departure from planarity of 10μm will give rise to an image displacement of the same magnitude. The ‘10μm’ figure happens to coincide with the unflatness tolerance of one commercially available 1K X 1K CCD sensor. Moreover, the influence of unflatness is most insidious in that is invariably leads to significant accuracy degradation in the object space, without aggravating the magnitude of triangulation misclosures.

In seeking to compensate for focal plane unflat ness there is really no satisfactory alternative to measuring the surface topography directly and applying corrections as per Eq. 8. Results of the investigation referred to earlier involving multi-sensor system self-calibration of DCS cameras (Fraser et al., 1995) only reinforced the notion that focal plane unflatness cannot be adequately modelled, nor fully compensated for through the AP approach.

In an attempt to ascertain the typical flatness of a DCS420 CCD sensor, the author obtained two 14 mm x 9 mm KAF-1600 chips from the Kodak Company. No manufacturer's specification regarding flatness was available, although Kodak's anticipation was that the 'bow' of the chip would be ±1 μm or less. The surface topography of these two chips was then measured at the Melbourne Branch of Australia's National Measurement Laboratory using both a flatness interferometer and a Wyko phase shifting interferometer. It was found by the Laboratory that the CCD chip functioned optically as a diffraction reflection grating giving not only specular reflections but also higher-order reflections at angles other than the angle of incidence. This property was used to distinguish between specular reflections from the protective glass window and higher-order reflections from the CCD sensor surface.

Flatness measurements of the physical chip surface (and not the unknown 'electronic' surface) were then made using the higher-order reflections. The resulting measured surface topography is shown in Fig. 1 for both CCD chips. Flatness, as expressed by the maximum peak-to-valley height difference, was 1.7 μm for each chip, with the measuring accuracy being ±0.1 μm. The RMS departure from a best-fitting plane in each case was 0.3 μm, which is encouraging given that this value is also reasonably representative of the image coordinate mensuration precision anticipated in high-accuracy vision metrology applications, namely 0.02 to 0.04 pixels. It is unlikely that either sensor would give rise to a clearly measurable deformation in object space triangulation if the camera employed had a medium-angle field of view of, say, 50°. We can only hope that as CCD matrix arrays increase in size they keep their level of 'bowing' down to that of the two chips examined. The author was informed by Kodak that measurements of the 24.6 x 24.6 mm KAF-1000 chip had shown a maximum peak-to-valley height difference of 5 *.urn,* which is clearly of more photogrammetric concern.

(a)

(b)

Fig. 1. Measured surface topography of two Kodak KAF-I600 14 mm x 9 mm CCD chips; total 'height' range for both is 1.7 μm and the RMS departure from planarity is 0.3 μm.

9. In-plane distortion

The problems of physical in-plane distortion that adversely influenced film-based photogrammetry are fortunately absent from digital systems employing high-resolution, large-area CCD cameras. The geometric integrity of the layout of the pixel array is typically precise to the 0.1 μm level (Shortis , and Beyer, 1996). Nevertheless, electronic effects can give rise to apparent in-plane distortion of the image. The first of two prominent effects which have received photogrammetric attention is the introduction of a differential scaling between *x* and *y* image coordinates due to differences between the frequency of the 'analog' CCD pixel shift clock and the sampling frequency of the A/D converter within the framegrabber. This error source can be readily modelled, and can be alleviated through the use of either pixel-synchronous framegrabbing or 'digital' CCD cameras such as the DCS still video cameras which employ on-board A/D conversion (Beyer, 1992).

The second influence for which a self-calibration correction function has been formulated is the introduction of image axis non-orthogonality. Systematic effects from line jitter were originally seen as the likely source of this image coordinate perturbation (e.g. Beyer, 1992). Once again, this distortion can be rendered metrically insignificant through the use of , ‘digital' CCD cameras, or through pixel synchronous A/D conversion with due attention being paid to aspects such as camera warm-up and power supply fluctuations.

Given the potential problems of overparameterization and the lack of any known models for higher-order distortions of CCD sensors, the AP correction model for in-plane distortion of the image reduces to two terms in the x-coordinate only in the present formulation. One is to account for differential scaling between the horizontal and vertical pixel spacings, and the other is to model non-orthogonality between the *x* and y axes:

$$\Delta x_f = b_1\bar{x} + b_2\bar{y} \qquad (9)$$

In instances of self-calibration where CCD sensors with digital output are employed, b_1 and b_2 are best regarded as empirical correction terms. In the author's experience with self-calibration of Kodak DCS cameras, the 'shear' term, b_2, is invariably insignificant. One very interesting result of the simultaneous multi-sensor self-calibration of DCS200 and DCS420 cameras referred to earlier, was that the affinity term, b_1, was found to be modelling an error signal from the lens and not from the image plane of the CCD sensor (Fraser et al., 1995). It was concluded that, firstly, neither of the two CCD sensors employed displayed any significant differential scaling between horizontal and vertical pixel spacing and, secondly, that the standard lens distortion model fell short of providing complete functional model fidelity.

10. Concluding Remarks

Substitution of the individual image coordinate correction functions comprising Eq. 3 yields the following lO-AP model for digital camera self-calibration, which is simply the 'standard' eight-parameter model supplemented with two terms for first-order in-plane image distortion:

$$\Delta x = -x_0 - \frac{\bar{x}}{c}\Delta c + \bar{x}r^2K_1 + \bar{x}r^4K_2 + \bar{x}r^6K_3 + (2\bar{x}^2 + r^2)P_1 + 2P_2\overline{xy} + b_1\bar{x} + b_2\bar{y}$$
$$\Delta y = -y_0 - \frac{\bar{y}}{c}\Delta c + \bar{y}r^2K_1 + \bar{y}r^4K_2 + \bar{y}r^6K_3 + 2P_1\overline{xy} + (2\bar{y}^2 + r^2)P_2 \qquad (10)$$

Eq. 10 could be said to constitute the current preferred model for digital close-range camera calibration. This image coordinate correction function is appropriate to any number of cameras employed in the bundle adjustment, and it is most interesting to recall that the eight-parameter model formed by omission of the 'empirical' correction terms b_1 and b_2 is the same as that proposed for film camera self-calibration 25 years ago (Kenefick et al., 1972). The implicit message arising from this is that solutions via AP models are as inapplicable for in-plane and out-of-plane distortions in modern digital cameras as they were for film cameras. Image coordinate perturbations such as focal plane unflatness can, if present, only be rectified adequately through physically measuring the CCD chip surface, and there is no solution short of mechanical modification to solve for an unstable interior orientation. There is nothing to preclude the supplementing of the self-calibration model of Eq. 10 with further, higher-order correction terms, though experience suggests that this is generally not a wise course of action.

As a final word, the 10-parameter model employed for close-range digital camera self-calibration has yielded object space triangulation accuracies of well beyond 1:100,000, and these accuracies have been verified through external measures in controlled tests. Medium-accuracy photogrammetric applications, employing still video cameras for stereo restitution for example, may well require only four calibration parameters, these being c, x_0, y_0 and K_1 .One of the benefits of the self-calibration approach is that it is a simple matter to assess the metric impact of changing the AP model. It will be interesting to see if the 'standard' self-calibration model stands the test of time in the future, or whether improved correction functions are developed to model image coordinate perturbations that, while being currently 'in the noise', nevertheless exhibit clear systematic trends.

Acknowledgements

This paper was prepared while the author was on study leave at the Institute of Photogrammetry, University of Stuttgart. The author gratefully acknowledges both Professor Dieter Fritsch and the University of Melbourne for making this study leave possible.