Extraction of Focal Lengths from the Fundamental Matrix *

Richard I. Hartley

G.E. CRD, Schenectady, NY, 12301.
Email : hartley@crd.ge.com

Abstract

The 8-point algorithm is a well known method for solving the relative orientation and scene reconstruction problem for calibrated cameras. In that algorithm, a matrix, called the fundamental matrix of the two calibrated cameras is computed using 8 or more point correspondences between the pair of images. The relative orientation of the two cameras may be computed from the fundamental matrix. In the present paper it is shown that the assumption of completely known calibration is unnecessary. In fact, the fundamental matrix may be used to solve the relative orientation problem for a pair of pinhole cameras with known principal points, but unknown (possibly different) focal lengths. A simple non-iterative algorithm is given for finding the two focal lengths. Once this is done, known techniques may be used to find the relative orientation of the cameras, and to reconstruct the scene.

1 Introduction

The essential matrix was introduced by Longuet-Higgins to solve the relative placement problem for calibrated cameras. Given a set of correspondences between points in a pair of images, he gave a non-iterative algorithm for computing the orientation and location of the second camera with respect to the first, as well as a method for computing the 3D locations of the points in space corresponding to the matched image points. This algorithm is known as the 8-point algorithm, though it may also be applied when more than eight point matches are known. The matrix may also be defined for uncalibrated cameras, in which case it has come to be known as the fundamental matrix. It is a compact way of representing the epipolar geometry of the two images.

The fundamental matrix is a 3×3 matrix with 7 degrees of freedom, of which 5 are required to compute the relative placement of the two cameras. It was shown in [1] that the two extra degrees of freedom of the fundamental matrix may be used to compute the focal lengths of the two cameras, provided all other internal parameters of the two cameras are known. The algorithm given in [1] was quite complicated, non-intuitive and hard to understand. It relied on complicated algebraic manipulations. Furthermore, no evaluation of the performance of the algorithm was given. In the present paper, a very simple algorithm is given for the computation of the two focal lengths, based on projective geometry. This approach allows one to gain an intuitive understanding of the method,

*The research described in this paper has been supported by DARPA Contract #MDA972-91-C-0053

as well as allowing the identification of critical configurations where the algorithm will fail.

Experiments are carried out to evaluate the performance. It is shown that the algorithm performs relatively well, and is quite robust in the presence of noisy data, provided the camera geometry is sufficiently far from one of the critical configurations.

The results of this paper may be compared with those of Maybank and Faugeras ([5]) on self-calibration. In [5] the problem considered is that of calibration of cameras from a set of images taken all with the **same** camera. It results ([5, 4] that three view are sufficient to calibrate a camera, but no very tractable algorithm for actually carrying out this calibration is given. In this paper we assume two views taken with **different** cameras, and give a simple algorithm for computing a partial calibration (specifically, the focal lengths) of the two cameras.

2 The Camera Model

The imaging process carried out by a pinhole camera may be broken down into two steps, corresponding to so-called external and internal camera orientation. The first step, corresponding to external camera orientation, consists in transforming 3D spatial coordinates into a camera-centred coordinate frame. A point $\mathbf{x} = (x, y, z)^\top$ in space is mapped to the point $\hat{\mathbf{x}} = R(\mathbf{x} - \mathbf{t})$, where R is a rotation matrix, and $\mathbf{t}$ is the camera centre. This may be represented in matrix notation as

$$\hat{\mathbf{x}} = (R \mid -R\mathbf{t})(x, y, z, 1)^\top \ .$$

where $(R \mid -R\mathbf{t})$ is a 3×4 matrix decomposed into a 3×3 block and a 3×1 column.

The second step corresponds to central projection. The point $\hat{\mathbf{x}} = (\hat{x}, \hat{y}, \hat{z})^\top$ is mapped to the point $(\hat{u}, \hat{v}) = (f\hat{x}/\hat{z} + p_u, f\hat{y}/\hat{z} + p_v)$, where f is the focal length of the camera, and (p_u, p_v) are the coordinates of the principal point. This may be written in the form $\hat{u} = u/w$ and $\hat{v} = v/w$, where

$$\begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} f & 0 & p_u \\ 0 & f & p_v \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \hat{x} \\ \hat{y} \\ \hat{z} \end{pmatrix} \ .$$

The upper-triangular 3×3 matrix in this equation is called the calibration matrix of the camera. One sees that $\mathbf{u} = (u, v, w)^\top$ are the homogeneous coordinates of the image point. It will be convenient, similarly, to represent 3D points in homogeneous coordinates as $\mathbf{x} = (x, y, z, 1)^\top$. Putting both the exterior and interior mappings together, one sees that $\mathbf{u} = K(R \mid -R\mathbf{t})\mathbf{x}$, where K is the calibration matrix. The matrix $M = (KR \mid -KR\mathbf{t})$ is called the camera matrix, which depends on both exterior and interior orientation of the camera. Given a matrix M it is a very simple matter to obtain K and R, using the QR-decomposition of the left-hand block of M.

It is sometimes customary to let K be an arbitrary upper-triangular matrix by allowing different magnifications f_u and f_v in different axial directions, and by introducing a *skew* parameter in the $(1, 2)$ position of the matrix. This will not be done here. For original images produced by a true pinhole camera there is no skew, and scaling (denoted by f) is isotropic. Enlarging of images (as is done in printing from negatives) or digitization using a digitizing camera may introduce skew and non-isotropic scale, if the image plane is not perpendicular to the principal axis of the enlarger or digitizing camera. Usually, this effect will be minimal. In CCD cameras non-square sensors will introduce unequal scaling in the two directions, but this will be a parameter of the sensor plane, and will not be effected by changing lenses or zooming. The effect of non-square pixels can be compensated for by adjusting the image coordinates. It is therefore valid to assume, as we shall, that skew is zero and scaling is isotropic. The scaling factor f present in the calibration matrix will represent the combined effect of focal length, enlargement and pixel size. Despite this, we shall still call it the focal length.

Given two images, it is not possible to estimate the focal lengths of the two cameras as well as the principal point offset. In this paper, it is assumed that the principal point of the camera is known. It then proves possible to compute the focal lengths of the two cameras. The principal point of a camera is rarely known precisely, but for practical purposes, it is sufficient to assume that the principal point is at the centre of the image, unless the image has been cropped. For instance, in [6] good scene reconstruction is achieved under this assumption. The most evident application of this work is for computing the focal length of a camera fitted with a zoom or changeable lens. It is reasonable to assume that the principal point will not move during zooming, though this is dependent on the quality of the zoom lens.

3 Properties of the Fundamental Matrix

The fundamental matrix is a 3×3 matrix expressing the epipolar correspondence between two images of a scene. It was introduced by Longuet-Higgins ([3]) for calibrated cameras, but applies to uncalibrated cameras as well. Several properties of the fundamental matrix are listed below. For proofs of these properties, see [3, 1].

Proposition 3.1. *Let F be the fundamental matrix for an ordered pair of images (J, J') of a common scene.*

1. *$F^\top$ is the fundamental matrix for the ordered pair of images (J', J). In other words, swapping the roles of the two images corresponds to transposing the fundamental matrix.*
2. *If $\mathbf{u} \leftrightarrow \mathbf{u}'$ is a pair of matching points in the two images, then $\mathbf{u}'^\top F \mathbf{u} = 0$.*
3. *If $\mathbf{u}$ is a point in the image J, then the corresponding epipolar line in the second image is the line $F\mathbf{u}$.*
4. *The rank of F is equal to two.*
5. *If $\mathbf{p}$ and $\mathbf{p}'$ are the two epipoles, then $F\mathbf{p} = \mathbf{p}'^\top F = 0$.*
6. *If the calibration matrix $K = I$ for both cameras, then F factors as $F = R[\mathbf{t}]_\times$, where R is a rotation matrix representing the orientation of the second camera relative to the first and $[\mathbf{t}]_\times$ is a skew-symmetric matrix given by*

$$[\mathbf{t}]_\times = \begin{pmatrix} 0 & -t_z & t_y \\ t_z & 0 & -t_x \\ -t_y & t_x & 0 \end{pmatrix}$$

where $\mathbf{t} = (t_x, t_y, t_z)^\top$ *represents the displacement of the second camera from the first.*

4 Computation of the Scale Factors

As observed by Maybank and Faugeras ([5]), the key to camera calibration is the absolute conic. We begin with some general observations about the absolute conic. For a more thorough treatment the reader is referred to [5]. The absolute conic is a conic on the plane at infinity consisting of points $(x, y, z, t)^\top$ such that $t = 0$ and $x^2 + y^2 + z^2 = 0$. The absolute conic contains no real points. Writing $\mathbf{x} = (x, y, z)^\top$, the defining equation for the absolute conic is $\mathbf{x}^\top \mathbf{x} = 0$.

The image point corresponding to such an object point as mapped by a camera with matrix $K(R \mid -R\mathbf{t})$ is given by $\mathbf{u} = KR\mathbf{x}$, since $t = 0$, from which we obtain $\mathbf{x} = R^{-1}K^{-1}\mathbf{u}$. Then from $\mathbf{x}^\top\mathbf{x} = 0$ follows $\mathbf{u}^\top K^{-\top} R R^{-1} K^{-1} \mathbf{u} = \mathbf{u}^\top (KK^\top)^{-1}\mathbf{u} = 0$. In other words, $\mathbf{u}$ is on the image of the absolute conic if and only if $\mathbf{u}^\top (KK^\top)^{-1}\mathbf{u} = 0$. Thus, the image of the absolute conic is a plane conic represented by the matrix $(KK^\top)^{-1}$.

We now specialize to the case where the calibration matrices are diagonal, of the form $\mathrm{diag}(k, k, 1)$ and $\mathrm{diag}(k', k', 1)$. In other words, it is assumed that the principal point is known, and is equal to the origin $(0, 0)$ of image coordinates. Under these assumptions, the image of the absolute conic has matrix $\mathrm{diag}(k^2, k^2, 1)^{-1} \approx diag(1, 1, k^2)$ for one camera and $\mathrm{diag}(k'^2, k'^2, 1)^{-1} \approx \mathrm{diag}(1, 1, k'^2)$ for the other[2]. These are circles of imaginary radius centred at the origin. Their equations are $u^2 + v^2 + k^2 = 0$ and $u'^2 + v'^2 + k'^2 = 0$.

Consider now two images J and J'. Let π be a plane in space passing through the two camera centres and tangent to the absolute conic. This plane is mapped into the image J as a line $\boldsymbol{\lambda}$ through the epipole $\mathbf{p}$ tangent to the image of the absolute conic. Similarly it maps to a line $\boldsymbol{\lambda}'$ in the second image. The two lines $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$ in the two images are a matching pair of epipolar lines. There are in fact two tangent planes through the two camera centres tangent to the absolute conic, which results in two pairs of epipolar lines $\boldsymbol{\lambda}_1 \leftrightarrow \boldsymbol{\lambda}'_1$ and $\boldsymbol{\lambda}_2 \leftrightarrow \boldsymbol{\lambda}'_2$ all tangent to the image of the absolute conic.

Let the two epipoles be $\mathbf{e}$ and $\mathbf{e}'$. We assume that neither $\mathbf{e}$ nor $\mathbf{e}'$ lies at the origin, which means that neither camera lies on the principal ray of the other. Simply by rotating each of the images, it may be ensured that the two epipoles lie on the positive u-axis. Rotating the image about the origin is equivalent to rotating the camera about the principal ray. Note that this observation is true only because of the simple form of the calibration matrix.

Suppose that the two epipoles are $\mathbf{e} = (e_1, 0, e_3)$ and $\mathbf{e}' = (e'_1, 0, e'_3)$. Under these conditions, the fundamental matrix has a particularly simple form. From the conditions $F\mathbf{e} = F(e_1, 0, e_3)^\top = 0$ and $\mathbf{e}'^\top F = (e'_1, 0, e'_3)F = 0$, we derive

$$F \approx \begin{bmatrix} e'_3 & & \\ & 1 & \\ & & -e'_1 \end{bmatrix} \begin{bmatrix} a & b & a \\ c & d & c \\ a & b & a \end{bmatrix} \begin{bmatrix} e_3 & & \\ & 1 & \\ & & -e_1 \end{bmatrix} \tag{1}$$

for real numbers a, b, c and d.

[2]The notation $\approx$ denotes equality up to a non-zero factor

Now, we consider the tangent lines from the epipole $(e_1, 0, e_3)^\top$ to the co nic with matrix $\mathrm{diag}(1, 1, k^2)$. The polar ([7]) of the point $(e_1, 0, e_3)^\top$ with respect to the conic $\mathrm{diag}(1, 1, k^2)$ is the lin e $(e_1, 0, k^2 e_3)^\top$, in other words the line $e_1 u + k^2 e_3 = 0$. The point s of tangency from $(e_1, 0, e_3)^\top$ to the conic are therefore the intersections of the line $e_1 u + k^2 e_3 = 0$ with the conic $u^2 + v^2 + k^2 = 0$. These are the points $\hat{\mathbf{u}} = (-k^2 e_3, ik(e_1^2 + k^2 e_3^2)^{1/2}, e_1)^\top$ and the complex conjugate point. Transforming this point by F, we wil l obtain the corresponding epipolar line in the second image, namely the tangent line from $\mathbf{e}' = (e'_1, 0, e'_3)^\top$ to the conic $\mathrm{diag}(1, 1, k'^2)$. Writing $\Delta = (e_1^2 + k^2 e_3^2)^{1/2}$, we compute

$$
\begin{aligned}
F\hat{\mathbf{u}} &\approx \begin{pmatrix} (-a\Delta^2 + bik\Delta)e'_3 \\ -c\Delta^2 + dik\Delta \\ (a\Delta^2 - bik\Delta)e'_1 \end{pmatrix} \approx \begin{pmatrix} (-a\Delta + bik)e'_3 \\ -c\Delta + dik \\ (a\Delta - bik)e'_1 \end{pmatrix} \\
&\approx \begin{pmatrix} -e'_3 \\ \frac{-c\Delta + dik}{a\Delta - bik} \\ e'_1 \end{pmatrix} = \begin{pmatrix} -e'_3 \\ \frac{(-c\Delta + dik)(a\Delta + bik)}{a^2\Delta^2 + b^2k^2} \\ e'_1 \end{pmatrix} \\
&= \begin{pmatrix} -e'_3 \\ \frac{-ac\Delta^2 - bdk^2 + ik(ad - bc)\Delta}{a^2\Delta^2 + b^2k^2} \\ e'_1 \end{pmatrix} \stackrel{\mathrm{def}}{=} \begin{pmatrix} -e'_3 \\ \mu + i\nu \\ e'_1 \end{pmatrix}
\end{aligned} \tag{2}
$$

The other tangent line is obtained as the complex conjugate of this line, namely $(-e'_3, \mu - i\nu, e'_1)$. On the other hand, these two tangent lines are tangents from the epipole $\mathbf{e}'$ to the conic $\mathrm{diag}(1, 1, k'^2)$, which is symmetric about the u-axis. In other words, if $(-e'_3, \mu + i\nu, e'_1)$ is one tangent line, t hen the other is $(-e'_3, -(\mu + i\nu), e'_1)$. Consequently, $-(\mu + i\nu) = \mu - i\nu$, and so $\mu = 0$. This gives

$$
\mu = \frac{-ac\Delta^2 - bdk^2}{a^2\Delta^2 + b^2k^2} = 0 \tag{3}
$$

whence

$$
-ac\Delta^2 - bdk^2 = -ac(e_1^2 + k^2 e_3^2) - bdk^2 = 0 \ .
$$

Finally, this leads to

$$
k^2 = \frac{-ace_1^2}{ace_3^2 + bd} \ . \tag{4}
$$

A formula for k' may be computed by reversing the role of the two images. This corresponds to taking the transpose of the fundamental matrix. Consequently, we may write

$$
k'^2 = \frac{-abe_1'^2}{abe_3'^2 + cd} \ . \tag{5}
$$

Since the focal lengths must be positive, we select the positive values of k and k' satisfying (4) and (5).

5 Algorithm Outline

Now, we are able to describe the complete algorithm for computation of the magnifications of the two cameras. It is assumed that the cameras have camera matrices diag$(k, k, 1)$ and diag$(k', k', 1)$ and that the fundamental matrix F corresponding to the (ordered) pair of cameras is known. Matrix F has rank 2. It is required to determine k and k'.

1. **Determination of the epipoles :** The two epipoles $\mathbf{e}$ and $\mathbf{e}'$ are determined by solving the equations $F\mathbf{e} = 0$ and $F^\top \mathbf{e}' = 0$.
2. **Normalizing the position of the epipoles :** The images are rotated so that the two epipoles both lie on the x-axis, namely at points $(e_1, 0, e_3)^\top$ and $(e'_1, 0, e'_3)^\top$ respectively. Subsequently, correct the fundamental matrix to reflect this change. More precisely, if T and T' are the two image rotations, then F is replaced by the new F equal to $T'^{-\top} F T^{-1}$. This new F will be of the form (1).
3. **Computation of k and k' :** Compute k and k' according to the formulae (4) and (5).

Once the scales k and k' are known, it is an easy matter to compute the relative placement of the two cameras. One simply computes the essential matrix $E = \mathrm{diag}(k', k', 1)^\top F \, \mathrm{diag}(k, k, 1)$ for cameras with identity calibration matrices and then computes the camera placement using the algorithm given in [3] or [1].

6 Failure

No solution possible : There may be no solution possible, if the right sides of (4) or (5) are negative. This indicates that the data is faulty. The fundamental matrix F does not correspond to a pair of cameras with the assumed simple calibration matrices.

The Principal Rays Meet : There are other cases, however, in which the algorithm may fail. The most interesting situation is when the principal rays of the cameras intersect. Any point along the principal ray is projected to the principal point in the image, which is assumed to be the origin $(0, 0, 1)^\top$. Therefore, the point of intersection of the principal rays maps to the origin in both images. This is an image correspondence $(0, 0, 1)^\top \leftrightarrow (0, 0, 1)^\top$. It follows that $(0, 0, 1)F(0, 0, 1)^\top = 0$, and hence that the $(3, 3)$-entry of F is zero. With F as in (1), this means $a = 0$. Now, from (3), we have $\mu = -bdk^2/b^2k^2 = -d/b$. It follows that $\mu = 0$ if and only if $d = 0$, and this condition is independent of k. This means that k and k' can not be determined from the fundamental matrix. This condition is somewhat troublesome in practice, since if two images of the same object are taken, then there is a good chance that they may both be taken with the camera aimed at the same point in the scene.

In this case, however, the two magnification factors k and k' may not be varied independently. In fact, if $a = d = 0$, then from (2) we see that $F\hat{\mathbf{u}} = (e'_3, ic\Delta/bk, -e'_1)^\top$. These are the two tangent lines to the image of the absolute conic in the second image. On the other hand, the points of tangency from point $\mathbf{e}' = (e'_1, 0, e'_3)^\top$ to the conic $u'^2 + v'^2 + k'^2 = 0$ are the two points $(-k'^2 e'_3, \pm ik'\Delta', e'_1)^\top$. However $F\hat{\mathbf{u}}$ is a tangent line, and hence passes through one of these points. Multiplying out, we get

$$\Delta'^2 \pm \frac{c\Delta k'\Delta'}{bk} = 0$$

from which it follows that $b\Delta'/k' = \pm c\Delta/k$. Squaring, substituting for Δ and Δ' and simplifying gives

$$b^2(e_3'^2 + e_1'^2 k'^{-2}) = c^2(e_3^2 + e_1^2 k^{-2}) \ .$$

The two magnification factors may vary freely as long as they satisfy this relationship.

Other Configurations : A similar situation occurs when $b = 0$. From (3) one deduces as before that $c = 0$, and the condition that $\mu = 0$ is independent of k.

Uniqueness It is evident that the geometric configurations that lead to failure of the algorithm through ambiguity are somewhat special cases. In fact, in the space of all possible configurations of cameras, the set of configurations that give rise to ambiguous solutions in the determination of the focal lengths constitute a critical set of lower dimension. This allows us to state the following uniqueness result.

Theorem 6.2. *For almost all configurations of a pair of pinhole cameras with known principal points, the focal lengths and relative orientation of the two cameras are determined uniquely by the correspondence of points seen in the two images.*

The words “almost all configurations” are to be interpreted as meaning all configurations, except those lying in a lower dimensional critical set. Similarly, the term pinhole camera is intended to denote a camera carrying out central projection from object space onto an image plane. As explained previously, such a camera has isotropic coordinates in the image plane, and no skew.

7 Experiments

Simulations were carried out to evaluate the algorithm in the presence of varying degrees of noise, and with varying proximity to a critical configuration. The set of experiments were carried out with a simulated 35mm camera with a focal length of 28mm. This focal length is equal to the lower limit of focal lengths for one commonly available zoom lens. Such a camera has a field of view approximately 64°. This is a fairly wide field of view, but also within typical ranges for aerial cameras. For this experiment, both simulated cameras had the same focal length. Images were assumed to be digitized with pixels of size 70μ, which means 500 pixels across the 35mm width of the image. The f parameter of the images was 400 (the focal length in pixels).

Two images were assumed taken with one camera rotated 30 degrees from the other. The two principal axes were skew lines in space chosen to be varying distances apart. A set of 30 points lying in a sphere of diameter 1.5 units were viewed by the two cameras. The cameras were placed so that the set of points approximately filled the field of view of each of the cameras. The images of the points were computed in each of the two cameras.

As seen in Section 6, if the principal axes meet, then the algorithm will fail and the focal lengths can not be determined. Suppose the principal axis of each camera could actually be seen in the image taken with the other camera. If the two principal axes meet in space, then the image of the principal axis of one camera, as seen by the other camera, will pass through the principal point of the image. The distance (in pixels) from the principal point in an image to the image of the principal axis of the other camera is a measure of how far the configuration differs from the critical configuration. This parameter will be denoted by α. For the configurations used in the experiment, the value of α was the same for both images (though this need not be so).

To determine how closeness to the critical configuration affects the stability of the computation a set of tests were carried out with values of α equal to 20, 39, 58 and 75 pixels. For each α, varying degrees of gaussian noise was added to each image pixel coordinate. The standard deviation of the added noise varied from 0.25 pixels to 5 pixels in each axial direction, for each point. For each noise level and value of α the computation was repeated 100 times with different added noise.

To measure the performance of the algorithm, four criteria were used.

1. Successful completion of the algorithm. In some cases, because of noise, the values on the right of (4) or (5) will be negative, and no solution is possible.

2. Accuracy of the computation of the focal length.

3. Accuracy of the ratio of the focal lengths of the two cameras. It turns out that errors in the computed focal lengths of the two cameras are strongly correlated. The ratio of the focal lengths is more accurately computed than the individual focal lengths.

4. Reconstruction error. Once the focal lengths were computed, the scene was reconstructed and compared with the actual values of the object points. To do this the reconstructed scene was superimposed on the actual scene using the algorithm of Horn ([2]) and the RMS construction error was measured.

The results of these experiments are reported in Figs 1 – 4.

8 Conclusions

The algorithm described in this paper provides a very simple method of computing the focal lengths and relative position and orientation of a pair of pinhole cameras given a set of image correspondences. As long as the principal axes of the two cameras are sufficiently distant from each other the algorithm performs quite well in the presence of reasonable amounts of noise. Realistic noise levels should be on the order of at most 2 pixels in each axial direction. The deviation of the estimated from the correct values of focal length are due to the inherent instability of the problem, rather than to an inadequacy of the algorithm. This has been demonstrated by using least-squares minimization methods to find the optimal solution iteratively. Images taken with a narrow field of view and long focal length cameras differ only slightly from orthographic views, and what difference there is can easily be swamped by noise.

References

[1] R. I. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Computer Vision - ECCV '92, LNCS-Series Vol. 588, Springer-Verlag*, pages 579 – 587, 1992.

[2] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America, A*, Vol. 4:629 – 642, 1987.

[3] H.C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, Sept 1981.

[4] Q.-T Luong. *Matrice Fondamentale et Calibration Visuelle sur l'Environnement.* PhD thesis, Universite de Paris-Sud, Centre D'Orsay, 1992.

[5] S. J. Maybank and O. D. Faugeras. A theory of self-calibration of a moving camera. *International Journal of Computer Vision*, 8:2:123 – 151, 1992.

[6] Jean Ponce, Todd Cass, and David Marimont. Relative stereo and motion reconstruction. Report UIUC-BI-AI-RCV-93-07, Beckman Institute, University of Illinois, 1993.

[7] J.G. Semple and G. T. Kneebone. *Algebraic Projective Geometry.* Oxford University Press, Oxford, 1952.

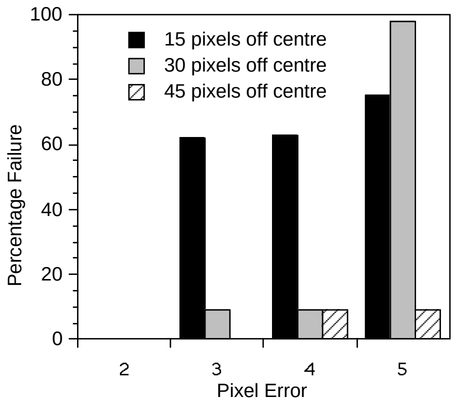

Figure 1: **Failure.** *The algorithm fails if the value of* k^2 *or* k'^2 *computed from (4) and (5) is negative. This plot shows the percentage failure of the algorithm. For RMS pixel errors of less than 3 pixels, the algorithm succeeded in all cases. Similarly for a value of* α *of 75 pixels, the algorithm always succeeded for all noise levels. The graphs shows the percentage failure for values of* α *of 20, 39 and 58 pixels at RMS noise levels of 3, 4 and 5 pixels. For a value of* α *of 20 pixels, the algorithm is not reliable for RMS pixel errors exceeding 2 pixels.*

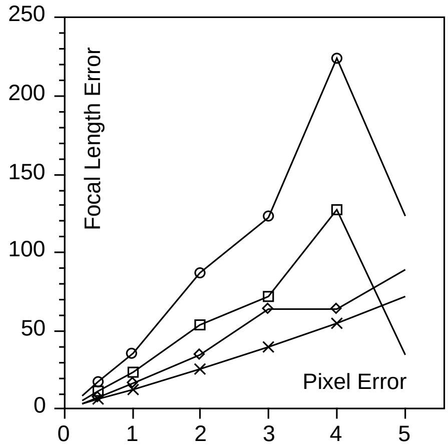

Figure 2: **Error in estimation of focal length.** *The plot shows the standard deviation in the estimate of focal length at different noise levels. In all cases, the mean estimated focal length was close to 400 (the correct value). The four plots show the standard deviation for values of* α *equal to (from the top) 20, 39, 58 and 75 pixels. Thus, in most cases, for noise levels less than 2 pixels, the estimated focal length lies in the range between 350 and 450. The apparent improvement in performance for pixel errors of 5 pixels in the top two plots is due to the fact that the algorithm failed for a substantial percentage of such tests, which were therefore not included in this statistic.*

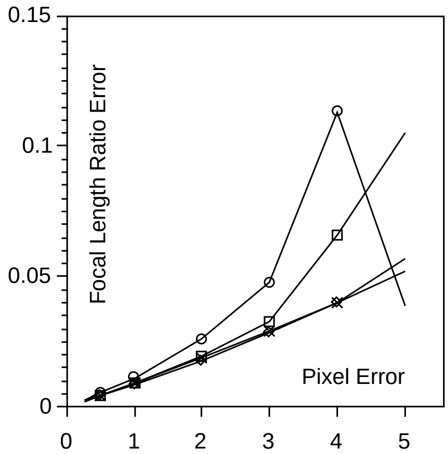

Figure 3: **Error in focal length ratio.** *This plot shows standard deviation in in the estimated ratio of the focal lengths of the two cameras for differing noise levels. Each curve represents a different value of* α *– from the top 20, 39, 58 and 75 pixels. The graph shows that except for high noise levels, the estimated ratio of the two focal lengths lies between 0.95 and 1.05. Once more the apparent improvement in performance for* $\alpha = 20$ *pixels at a noise level of 5 pixels is due to the exclusion of those cases where the algorithm fails.*

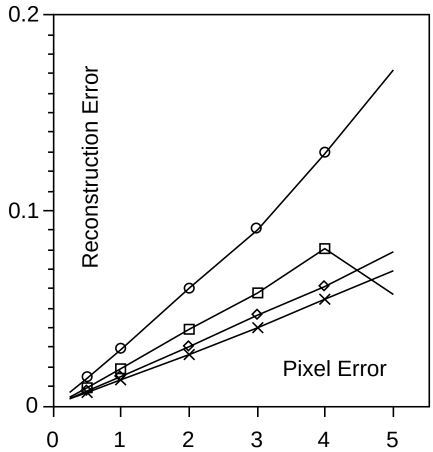

Figure 4: **Reconstruction Error.** *This graph shows the RMS reconstruction error for different noise levels and the same values of* α *as before. The vertical axis gives the absolute RMS reconstruction error in length units where the reconstructed set of points has radius about 1.5 units.*