

A Case Against Epipolar Geometry

Andrew Zisserman and Stephen J. Maybank

Robotics Research Group,
Department of Engineering Science,
University of Oxford,
Parks Rd,
Oxford

Abstract. We discuss briefly a number of areas where epipolar geometry is currently central in carrying out visual tasks. In contrast we demonstrate configurations for which 3D projective invariants can be computed from perspective stereo pairs, but epipolar geometry (and full projective structure) cannot. We catalogue a number of these configurations which generally involve isotropies under the 3D projective group, and investigate the connection with camera calibration. Examples are given of the invariants recovered from real images. We also indicate other areas where a strong reliance on epipolar geometry should be avoided, in particular for image transfer.

1 Introduction

The theory of epipolar geometry is, without doubt, one of the major achievements in the applications of projective geometry to computer vision. The number of points required, ambiguities, geometric and algebraic properties are well established and understood [4, 10, 12, 16]. The essence of epipolar geometry is that it depends *only* on the relative location of the optical centers and image planes; there is no dependence on structure. The primary use of epipolar geometry is in providing a disambiguating constraint for correspondences between images planes i.e., that a point in one image constrains the corresponding point in another image to lie on a line. However, epipolar geometry is also the kernel in a number of other areas:

1. Multiple view invariants

Invariants of a 3D configuration are computed from two perspective images of the configuration, based on the epipolar geometry between the views.

2. Transfer

Typically, novel images of 3D structures are constructed from a stereo pair of images by intersecting epipolar lines [3, 7, 21] and Barrett (this volume).

3. Camera self-calibration

Constraints from the fundamental matrix are exploited to determine intrinsic camera parameters [5, 13, 17].

In this paper we investigate whether epipolar geometry is necessary for these applications, and, furthermore, whether the use of epipolar geometry as the central tool is *detrimental*.

1.1 Two View Invariants

In general invariants of 3D configurations are computed in two stages: first, the fundamental matrix is computed from point matches between the two perspective images; second, invariants of the 3D configuration are computed either from the recovered 3D structure [6, 10] or from image measurements together with the epipolar geometry [9, 20] and (Gros and Shashua, both this volume). The computation of the fundamental matrix generally uses all available scene points, not exclusively points on the target 3D configuration.

Here we investigate invariants of 3D configurations, computed from two views, where only elements of the projected configuration are used. Section 2 describes a number of configurations for which there is an isotropy group. Section 3 uses these as examples in discussing epipolar geometry and multiple view invariants. For a number of these configurations it is not possible to determine the fundamental matrix, however invariants can be computed from two images. Section 4 establishes the relation with camera calibration for particular configurations.

1.2 Image Sequences and Structure

When more than two images are available, image point correspondence ambiguity can be reduced from an area around a line (using only a single fundamental matrix) to an area around a point. Two methods are routinely used to achieve this: in the first approach, 3D structure together with 3×4 projection matrices are computed (see [27], Mohr and Hartley, this volume), the imaged point is obtained by projecting the 3D point onto the target image; in the second approach, the imaged point in the target image is obtained by intersecting epipolar lines computed between image pairs (epipolar transfer). Section 5 describes the limitations of using epipolar geometry for image transfer.

Even in the two view case, using structure might improve the accuracy of fundamental matrix computation. Assuming that outliers have been removed [28], two possible least square minimisations are: first, to minimise distances between image points and putative epipolar lines [5]; or second, to compute structure for matched points, and minimise distances between measured imaged points and putative points predicted by projecting computed 3D structure with computed 3×4 projection matrices (see Hartley, this volume). In the latter case there is a tighter constraint (point to point). The former (point to line) is only affected by the distance of the measured point to the epipolar line, but not by its distance “along” the line. Similar considerations apply in the case of parallel projection [24].

1.3 Definitions and Notation

In the following, we assume a perspective camera, with unknown intrinsic parameters, and measure only projective properties in the image.

Epipolar Geometry and Camera Projection Matrix The geometric and algebraic properties of epipolar geometry and the fundamental matrix are well documented elsewhere [5, 10] so are not repeated here. Image points in a stereo pair, corresponding to a 3D point \mathbf{X} , are \mathbf{x} and \mathbf{x}' , with $\mathbf{x} = \mathbf{P}\mathbf{X}$ and $\mathbf{x}' = \mathbf{P}'\mathbf{X}$, where \mathbf{P} and \mathbf{P}' are the 3×4 projection matrices for the ‘left’ and ‘right’ cameras. In general homogeneous vectors are used, and equality is up to a non-zero scale factor. The fundamental matrix is represented by \mathbf{F} . For corresponding point pairs, $\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$.

3D Invariants These are invariants under a projective transformation of \mathcal{P}^3 , i.e. a transformation $\mathbf{X}' = \mathbf{T}\mathbf{X}$, where \mathbf{X} and \mathbf{X}' are homogeneous four-vectors and \mathbf{T} is a 4×4 non-singular matrix. For six 3D points in general position there are three functionally independent scalar projective invariants:

$$I_1 = \frac{|\mathbf{I}_{3561}| \cdot |\mathbf{I}_{3542}|}{|\mathbf{I}_{3564}| \cdot |\mathbf{I}_{3512}|}, \quad I_2 = \frac{|\mathbf{I}_{3562}| \cdot |\mathbf{I}_{3142}|}{|\mathbf{I}_{3512}| \cdot |\mathbf{I}_{3642}|}, \quad I_3 = \frac{|\mathbf{I}_{3564}| \cdot |\mathbf{I}_{5612}|}{|\mathbf{I}_{3561}| \cdot |\mathbf{I}_{5642}|}, \quad (1)$$

where $\mathbf{I}_{abcd} = [\mathbf{X}_a, \mathbf{X}_b, \mathbf{X}_c, \mathbf{X}_d]$ and $|\cdot|$ is the determinant.

2 Isotropy Sub-Groups

An isotropy group is the set of elements of a transformation group which do not alter a configuration. For example, a line is unaffected by translation in the direction of the line, and a circle unaffected by rotation about its center. If there is an isotropy, then there may well be invariants which naive constraint-counting would not predict. A simple example is a configuration of two coplanar lines under a plane similarity transformation, where the angle between the lines is invariant. This configuration has four degrees of freedom (two for each line) and under the group of plane similarity transformations ($\dim G = 4$), naive counting predicts no scalar invariants. However, there is a one-dimensional isotropy group in this case (scaling of coordinates with origin the line intersection) which does not affect the configuration.

The number, n_I , of (functionally independent scalar) invariants under the action of a group G is given by [21]:

$$n_I = \dim S - \dim G + \dim Is \quad (2)$$

where $\dim S$ is the ‘‘dimension’’ of the structure (the number of degrees of freedom), $\dim G$ the dimension of G , in the 3D projective case 15, and $\dim Is$ the dimension of the isotropy sub-group (if any) which leaves the structure unaffected under the action of G . A simple planar example is a configuration, S , of two lines and two points (not lying on the lines), under a plane projective transformation. S has eight degrees of freedom (because each point and each line has two degrees of freedom), $\dim G = 8$ for a plane projective transformation, and there is a one-dimensional isotropy sub-group of the projectivities which leaves

the structure unchanged (given explicitly in [8]). There is one projective scalar invariant (from equation (2) $n_I = 8 - 8 + 1$) which can be expressed as:

$$I = \frac{(\mathbf{l}_1^\top \mathbf{x}_1)(\mathbf{l}_2^\top \mathbf{x}_2)}{(\mathbf{l}_1^\top \mathbf{x}_2)(\mathbf{l}_2^\top \mathbf{x}_1)} \quad (3)$$

where $\mathbf{l}_i, \mathbf{x}_i$ are homogeneous line and point coordinates respectively. The 3D analogue of this configuration is two planes and two points (not lying on the planes), which has an invariant under 3D projective transformations. Other 3D examples are given in Table 1.

| Structure (S) | dim S | dim I_S | n_{I_3} |
|-------------------------------|---------|-----------|-----------|
| 6 points general position | 18 | 0 | 3 |
| 7 points general position (*) | 21 | 0 | 6 |
| 5 points, 4 coplanar | 14 | 1 | 0 |
| 6 points, 4 coplanar (*) | 17 | 0 | 2 |
| line and 4 coplanar points | 15 | 2 | 2 |
| 2 lines and 4 coplanar points | 19 | 0 | 4 |
| four lines | 16 | 1 | 2 |

Table 1. n_{I_3} is the number of functionally independent scalar invariants for 3D configurations under the action of the projective group. In all cases general position is assumed e.g. the line is not coplanar with any two points. (*) indicates that the epipolar geometry can be determined from two views of the structure. In the case of seven points, the fundamental matrix is determined up to a finite number of possibilities.

2.1 Configurations and Coordinates

In the following we describe a number of configurations for which there is an isotropy sub-group of the 3D projective group. All the examples are variations on a configuration containing four co-planar points. The configurations are defined in Table 2.

For the six point configuration (**VI**), we may arbitrarily choose coordinates for any five of the points (not containing all four of the coplanar points):

$$\begin{aligned} \mathbf{X}_1 &= (1, 0, 0, 0)^\top \\ \mathbf{X}_2 &= (0, 1, 0, 0)^\top \\ \mathbf{X}_3 &= (0, 0, 1, 0)^\top \\ \mathbf{X}_5 &= (0, 0, 0, 1)^\top \\ \mathbf{X}_6 &= (1, 1, 1, 1)^\top \end{aligned}$$

| Type | Structure | S | $DOFS$ | $\dim I_S$ | n_{I_3} |
|------|--------------------|-----|--------|------------|-----------|
| I | single line | | 15 | 2 | 2 |
| II | single point | | 14 | 1 | 0 |
| III | point on line | | 16 | 1 | 2 |
| IV | two coplanar lines | | 18 | 1 | 4 |
| V | two skew lines | | 19 | 0 | 4 |
| VI | two points | | 17 | 0 | 2 |

Table 2. The structure S consists of four coplanar points together with other elements not on the plane. The table gives the number of degrees of freedom of S , the dimension of the isotropy group (if any), and the number of functionally independent scalar 3D projective invariants. In all cases general position is assumed. Identical results hold if the four coplanar points are replaced by four coplanar lines.

Any other coordinates of the five points can be transformed to these by a collineation of \mathcal{P}^3 . The fourth point has coordinates:

$$\mathbf{X}_4 = (\alpha, \beta, \gamma, 0)^\top$$

The points are shown in Figure 1a.

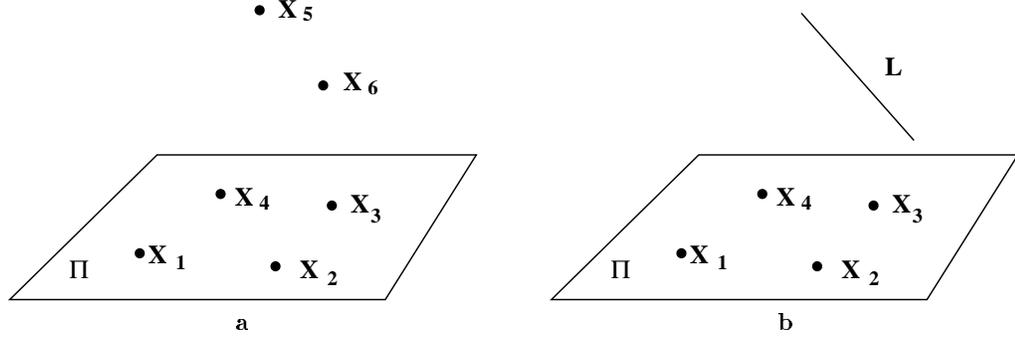


Fig. 1. The basic configurations consisting of four coplanar points and either (a) two points (configuration **VI**) or (b) a line (configuration **I**), not on the plane Π .

2.2 Four Coplanar Points and a Non-Coplanar Line

This is configuration **I** in Table 2. Coordinates are chosen in \mathcal{P}^3 such that the coplanar points $\mathbf{X}_1, \dots, \mathbf{X}_4$ have coordinates as above (configuration **VI**), and such that the line, \mathbf{L} , is represented in homogeneous parametric form as

$$\mathbf{L} = \zeta(1, 1, 1, 1)^\top + \eta(0, 0, 0, 1)^\top \quad (4)$$

See Figure 1b.

The isotropy sub-group is determined in two stages as follows. First, in order for Π to be a plane of fixed points it is necessary and sufficient that the 4×4 transformation matrix \mathbf{T} satisfies

$$\mathbf{X}_i = \lambda \mathbf{T} \mathbf{X}_i, \quad i \in \{1, \dots, 4\}$$

It is a simple matter to show that \mathbf{T} must have the form

$$\mathbf{T} = \begin{pmatrix} \mu_1 & 0 & 0 & \mu_2 \\ 0 & \mu_1 & 0 & \mu_3 \\ 0 & 0 & \mu_1 & \mu_4 \\ 0 & 0 & 0 & \mu_5 \end{pmatrix} \quad (5)$$

where $\mu_i, i \in \{1, \dots, 5\}$ parameterises the sub-group which has four DOF (only the ratio of the μ 's is significant).

Second, we determine the sub-group which leaves \mathbf{L} fixed. This can be carried out using Plückerian line coordinates [23], but here we use the parametric representation (4) above. Under the action of the isotropy group the points on the line need not be fixed, but the transformed points must still lie on \mathbf{L} . The transformation of two points is sufficient to determine the transformed lines (three are required to determine the transformation of all the points on the line). By inspection \mathbf{TX}_5 and \mathbf{TX}_6 satisfy (4) iff $\mu_2 = \mu_3 = \mu_4$. Hence we arrive at

$$\mathbf{T}_2 = \begin{pmatrix} \mu_1 & 0 & 0 & \mu_2 \\ 0 & \mu_1 & 0 & \mu_2 \\ 0 & 0 & \mu_1 & \mu_2 \\ 0 & 0 & 0 & \mu_5 \end{pmatrix} \quad (6)$$

which is a two dimensional sub-group of the collineations of \mathcal{P}^3 .

It is interesting to examine the transformation of \mathcal{P}^3 under the action of \mathbf{T}_2 . The clearest way to see this is to determine the eigen-vectors of \mathbf{T}_2 . These are the fixed points of the collineation. The first three are degenerate with eigen-value μ_1 , and may be chosen as

$$\begin{aligned} \mathbf{E}_1 &= (1, 0, 0, 0)^\top = \mathbf{X}_1 \\ \mathbf{E}_2 &= (0, 1, 0, 0)^\top = \mathbf{X}_2 \\ \mathbf{E}_3 &= (0, 0, 1, 0)^\top = \mathbf{X}_3 \end{aligned} \quad (7)$$

The fourth has eigen-value μ_5 :

$$\mathbf{E}_4 = (\mu_2, \mu_2, \mu_2, \mu_5 - \mu_1)^\top \quad (8)$$

As expected any point on the plane $\mathbf{X} = \nu_1\mathbf{X}_1 + \nu_2\mathbf{X}_2 + \nu_3\mathbf{X}_3$ is unchanged by \mathbf{T}_2 (since after the transformation all the basis vectors are multiplied by μ_1). The fourth eigen-vector is a fixed point on \mathbf{L} . To see the effect of the isotropy group on points not on \mathbf{L} , consider any line \mathbf{L} containing \mathbf{E}_4 . This will intersect \mathbf{H} at some point, \mathbf{X}_H say, and any point, \mathbf{X} , on the line is given by $\mathbf{X} = \zeta\mathbf{E}_4 + \eta\mathbf{X}_H$. After the transformation the point is $\mathbf{T}_2\mathbf{X} = \mu_5\zeta\mathbf{E}_4 + \mu_1\eta\mathbf{X}_H$ which still lies on \mathbf{L} i.e. any line through \mathbf{E}_4 is a fixed line under the isotropy. Consequently, since every point in \mathcal{P}^3 lies on a line through \mathbf{E}_4 , the action of \mathbf{T}_2 on \mathcal{P}^3 is to move points towards (or away from) \mathbf{E}_4 , with only \mathbf{E}_4 and points on \mathbf{H} remaining unchanged. This eigen-vector structure is the key to understanding the isotropies in the next section.

2.3 Four Coplanar Points and a Non-Coplanar Point

This is configuration **II** in Table 2. Without loss of generality the non-coplanar point can be chosen to be $\mathbf{X}_5 = (0, 0, 0, 1)^\top$. Then for an isotropy under collineation \mathbf{T}

$$\mathbf{X}_i = \lambda\mathbf{TX}_i, \quad i \in \{1, \dots, 5\}$$

As above, a plane of fixed points restricts \mathbf{T} to the form (5). For \mathbf{X}_5 to be fixed under the isotropy

$$\mathbf{T}_1 = \begin{pmatrix} \mu_1 & 0 & 0 & 0 \\ 0 & \mu_1 & 0 & 0 \\ 0 & 0 & \mu_1 & 0 \\ 0 & 0 & 0 & \mu_5 \end{pmatrix} \quad (9)$$

which is a one dimensional sub-group of the collineations of \mathcal{P}^3 .

This form also follows immediately from the discussion of eigen-vectors in Section 2.2. The (ratio of) parameters $\{\mu_1, \mu_2, \mu_5\}$ determine the position of \mathbf{E}_4 , the fixed point not on the plane. For \mathbf{X}_5 to be fixed we require

$$\begin{aligned} \mu_2 &= 0 \\ \mu_1 - \mu_5 &\neq 0 \end{aligned}$$

which directly produces \mathbf{T}_1 as in (9).

There are a number of variations on the five point structure for which there is again a one-dimensional isotropy group. Firstly (Table 2 **III**), any line through \mathbf{X}_5 will be a fixed line (*not* a line of fixed points). This is obvious because (9) is a sub-group of (6) which preserves \mathbf{L} , but also because any line containing \mathbf{X}_5 intersects Π , and consequently there are two fixed points on \mathbf{L} , so the line is fixed. Furthermore, a star of lines through \mathbf{X}_5 will be fixed by (6) (i.e. any number of lines). In particular, there is a one-dimensional isotropy group for a configuration of two coplanar lines, not on Π , together with four points on Π (Table 2 **IV**).

3 Multiple View Invariants and Epipolar Geometry

Here we contrast configurations **I**, **II**, and **VI** in terms of whether epipolar geometry and 3D projective invariants can be determined from two views.

3.1 Six Point Configuration, Four Coplanar

This is configuration **VI** in Table 2. It has been reported by a number of authors [1, 3, 19] that it is possible to recover epipolar geometry (and hence, subsequently, projective structure) and invariants for this configuration. Here we summarise the construction.

Epipolar Geometry See Figure 2. There are six corresponding points $\mathbf{x}_i, \mathbf{x}'_i, i \in \{1, \dots, 6\}$ in two views, with the first four $i \in \{1, \dots, 4\}$ the projection of coplanar world points.

1. Compute the plane projective transformation matrix \mathbf{T} , such that $\mathbf{x}'_i = \mathbf{T}\mathbf{x}_i, i \in \{1, \dots, 4\}$.
2. Determine the epipole, \mathbf{p}' , in the \mathbf{x}' image as the intersection of the lines $(\mathbf{T}\mathbf{x}_5) \times \mathbf{x}'_5$ and $(\mathbf{T}\mathbf{x}_6) \times \mathbf{x}'_6$.

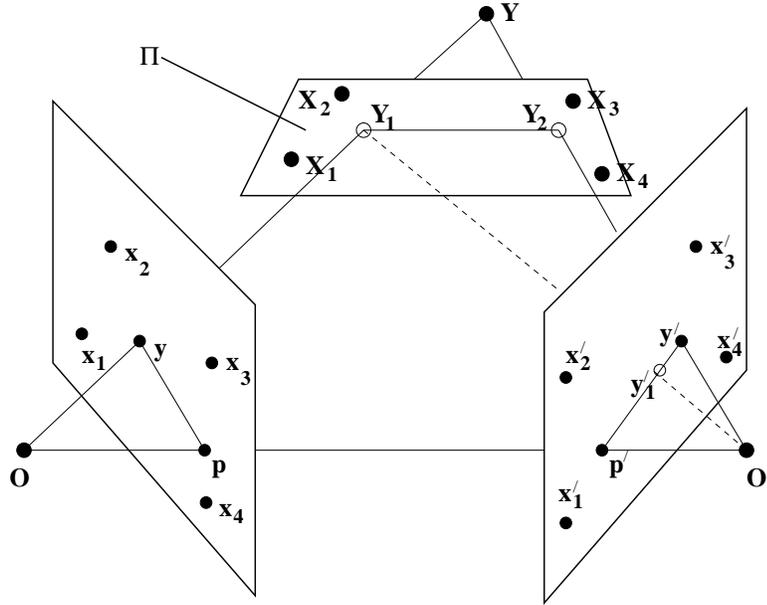


Fig. 2. Epipolar geometry. The points X_1, \dots, X_4 are coplanar, with images x_i and x'_i in the first and second images respectively. The epipolar plane defined by the point Y and optical centers O and O' intersects the plane Π in the line $L(Y) = \langle Y_1, Y_2 \rangle$, where Y_1 and Y_2 are the intersections of Π with the lines $\langle Y, O \rangle$ and $\langle Y, O' \rangle$ respectively. The epipolar line may be constructed in the second image as follows: Determine the plane projective transformation such that $x'_i = Tx_i, i \in \{1, \dots, 4\}$. Use this transformation to transfer the point y to $y'_1 = Ty$. This determines two points in the second image, y' and Ty , which are projections of points $(Y$ and $Y_1)$ on the line $\langle O, Y \rangle$. This defines the epipolar line of y in the second image. A second point, not on Π , will define its corresponding epipolar lines. The epipole lies on both lines, so is determined by their intersection. A similar construction gives epipolar lines and hence the epipole in the first image.

3. The epipolar line in the x' image of any other point x is given by $(Tx) \times p'$.
4. Hence $F = [p']_{\times} T$, where $[p']_{\times}$ is the skew matrix formed from the components of p' .

Projective Invariants This configuration has seventeen degrees of freedom (three for each point less one for the planarity constraint), and there is no isotropy group. From the counting argument in Section 2, the configuration has two projective invariants. The invariants can be computed using (1) for the 3D invariants of six points (one of the three invariants will be zero due to planarity of four points). Equivalently the invariants can be computed from the planar invariants of the four coplanar points together with the intersection of the line through X_5 and X_6 and the plane Π . See Figure 3.

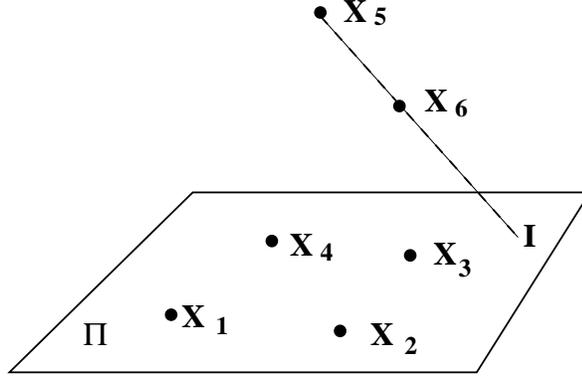


Fig. 3. The projective invariant of 6 points, 4 coplanar (points 1-4), can be computed by intersecting the line, \mathbf{L} , through the non-planar points (5 and 6) with Π . There are then 5 coplanar points, for which two invariants to the plane projective group can be calculated.

1. The \mathbf{x}' image of the point of intersection, \mathbf{X}_I , of the line $\langle \mathbf{X}_5, \mathbf{X}_6 \rangle$ and the plane, is given by the intersection [22] of the lines $(\mathbf{T}\mathbf{x}_5) \times (\mathbf{T}\mathbf{x}_6)$ and $\mathbf{x}'_5 \times \mathbf{x}'_6$.
2. Two plane projective invariants can be calculated from five points (in this case the four coplanar points and \mathbf{x}'_I) by

$$I_1 = \frac{|m_{431}| |m_{I21}|}{|m_{421}| |m_{I31}|} \quad I_2 = \frac{|m_{421}| |m_{I32}|}{|m_{432}| |m_{I21}|} \quad (10)$$

where m_{jkl} is the matrix $[\mathbf{x}'_j, \mathbf{x}'_k, \mathbf{x}'_l]$ and $|m_{jkl}|$ its determinant.

3.2 Five Point Configuration, Four Coplanar

As shown in Section 2.3 there is a one dimensional isotropy group for configuration **II**. There are only fourteen degrees of freedom, and consequently (equation 2) no projective invariants. Furthermore, the epipolar geometry cannot be determined.

It is instructive to consider how the epipolar geometry is constrained by images of this configuration. It is clear from Figure 4 that in each image the epipole is constrained to lie on a line. To see this algebraically we again use the notation in Section 2 for the 3D points $\mathbf{X}_i, i \in \{1, \dots, 5\}$. In this case we can, without loss of generality, choose \mathbf{X}_4 to have coordinates $\mathbf{X}_4 = (1, 1, 1, 0)^T$. We choose the projective coordinates of their images in both views to be:

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{x}'_1 = (1, 0, 0)^T \\ \mathbf{x}_2 &= \mathbf{x}'_2 = (0, 1, 0)^T \\ \mathbf{x}_3 &= \mathbf{x}'_3 = (0, 0, 1)^T \\ \mathbf{x}_4 &= \mathbf{x}'_4 = (1, 1, 1)^T \end{aligned}$$

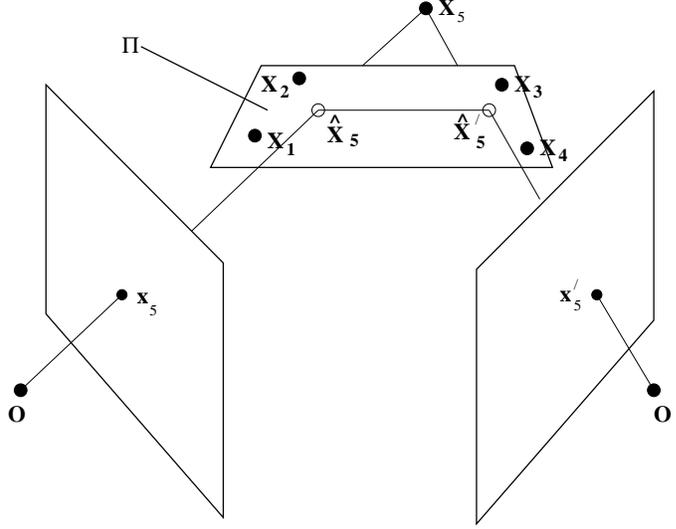


Fig. 4. As in Figure 2 the line $\langle \hat{\mathbf{X}}_5, \hat{\mathbf{X}}'_5 \rangle$ is coplanar with the line joining the optical centers $\langle \mathbf{O}, \mathbf{O}' \rangle$. Thus one point off the plane constrains the epipole to lie on a line in each image, namely the projection of $\langle \hat{\mathbf{X}}_5, \hat{\mathbf{X}}'_5 \rangle$.

and \mathbf{X}_5 to project to $\mathbf{x}_5 = (a, b, c)^\top$ and $\mathbf{x}'_5 = (a', b', c')^\top$ in the left and right images respectively. With this notation coordinates for points on the plane, Π , are identical for both images (and given by the first three homogeneous 3D coordinates). The projection matrix can be shown to be:

$$\mathbf{P} = \begin{pmatrix} \rho & 0 & 0 & a \\ 0 & \rho & 0 & b \\ 0 & 0 & \rho & c \end{pmatrix}. \quad (11)$$

Up to an overall scaling, this has one degree of freedom (since only the ratio of $\{a, b, c\}$ is significant). Solving for the optical center $\mathbf{C} = (C_1, C_2, C_3, C_4)^\top$ via $\mathbf{P}\mathbf{C} = 0$ gives

$$\mathbf{C} = C_4 \begin{pmatrix} \frac{a}{\rho} \\ \frac{b}{\rho} \\ \frac{c}{\rho} \\ 1 \end{pmatrix} = \alpha \begin{pmatrix} a \\ b \\ c \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

which, as expected, constrains \mathbf{C} to lie on the line through \mathbf{X}_5 intersecting Π in the point $\hat{\mathbf{X}}_5 = (a, b, c, 0)^\top$. Corresponding results for the primed image are obtained by priming a, b, c etc. Consequently, both optical centers are constrained to lie on lines containing \mathbf{X}_5 (i.e. the lines are coplanar), and so the image of one optical center from the other (the epipole) is constrained to a line.

A complementary picture is obtained by considering the action of the isotropy group: Under the (one dimensional) isotropy group, T_1 , the configuration $\mathbf{X}_i, i \in$

$\{1, \dots, 5\}$ is unchanged, but the optical centers *are* moved - they travel along the line $\langle \hat{\mathbf{X}}_5, \mathbf{X}_5 \rangle$ and $\langle \hat{\mathbf{X}}'_5, \mathbf{X}_5 \rangle$ respectively (i.e. the orbit is along the fixed lines as described in Section 2.2). Clearly, a projective transformation cannot alter projective invariants, and indeed the action of the isotropy group maintains the image projective coordinates.

3.3 Four Points and a Line Configuration, Four Points Coplanar

As shown in Section 2.2 there is a two dimensional isotropy group for configuration **I**. There are fifteen degrees of freedom, and consequently (equation 2) two projective invariants. These are determined in the same manner as the six point configuration **VI**.

Projective Invariants

1. Compute the plane projective transformation matrix \mathbf{T} , such that $\mathbf{x}'_i = \mathbf{T}\mathbf{x}_i, i \in \{1, \dots, 4\}$.
2. The \mathbf{x}' image of the point of intersection, \mathbf{x}'_I , of the plane Π and line \mathbf{L} , is given by the line intersection $(\mathbf{T}^{-1}\mathbf{l}) \times \mathbf{l}'$, where \mathbf{l} and \mathbf{l}' are the images of \mathbf{L} [22].
3. Two plane projective invariants can be calculated for the five points (in this case the four coplanar points and \mathbf{x}'_I) as in (10).

The epipolar geometry cannot be determined in this case. The epipole is not even constrained to a line. This is clear from the action of the isotropy group on the optical center, which moves \mathbf{C} on a *surface* (as the isotropy group is two dimensional) whilst the structure, S , is unchanged. As this is a projective transformation both the image projective invariants, and the 3D projective invariants of S are unaltered. So there is a surface, the orbit of \mathbf{C} under Is , over which \mathbf{C} can move which has no effect on projective image measurements. Conversely, the image projective invariants are *only* affected by motions of \mathbf{C} transverse to this surface. Consequently, image projective invariants can only place constraints on this transverse motion, so \mathbf{C} is restricted to a one-parameter family of surfaces. From Section 2.2 it can be shown that this orbit (the surface) is a plane.

Counting constraints and unknowns indicates why 3D projective invariants can be recovered from two images in this case: There are four “unknowns” (two 3D invariants and the one degree of freedom of each optical center which affects the image i.e. can be measured) and four measurements (the two projective coordinates of the line in each image). Consequently, it is possible to solve for the 3D invariants in this case.

Experimental Results Invariants using (10) are measured from the images shown in Figure 6. Points are extracted as follows: A local implementation of Canny’s edge detector is used to find edges to sub-pixel accuracy. These edge chains are linked, extrapolating over any small gaps. A piecewise linear graph

is obtained by incremental straight line fitting. Edgels in the vicinity of tangent discontinuities (“corners”) are excised before fitting as the edge operator localisation degrades with curvature. Vertices are obtained by extrapolating and intersecting the fitted lines. The results are given in Table 3 with the labelling indicated in Figure 5.

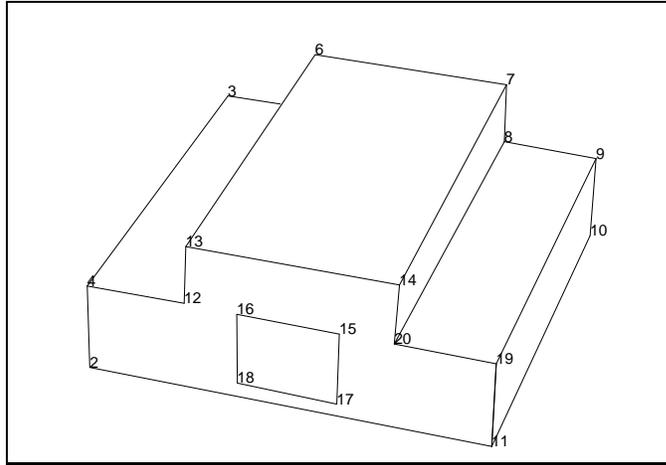


Fig. 5. Line drawing of the hole punch extracted from image A in Figure 6. Points 1 and 5 are occluded in this view.

| Images | I_1 | I_2 |
|--------|-------|--------|
| D,B | 0.378 | -1.117 |
| B,A | 0.371 | -1.170 |
| C,E | 0.370 | -1.150 |
| F,A | 0.333 | -1.314 |

Table 3. Line and four coplanar point two view invariants extracted from several combinations of views using points 2,4,14,17 and the line between points 6 and 13.

3.4 Summary

We have demonstrated that multiple view invariants of 3D configurations can be recovered without epipolar calibration being necessary, and argued that if a configuration has a non trivial isotropy group under projective transformations of 3D, then it is not possible to recover epipolar geometry using projective measurements.

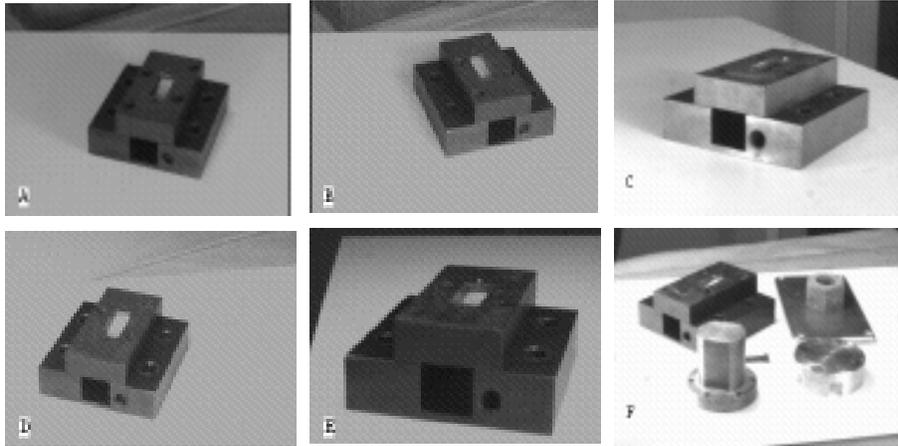


Fig. 6. Images of a hole punch captured with different lenses and viewpoints.

The discussion applies as well to analogues of these configuration, for example: four coplanar lines and two non-coplanar points, and five lines (four coplanar). If the configurations are enlarged by adding lines or points to the plane Π , there are no additional constraints on the optical center position or epipolar geometry, since Π is a plane of fixed points under the isotropy group. However, additional planar projective invariants will then be available.

4 Isotropy Sub-Groups and Camera Calibration

Suppose a configuration S for which a non-trivial isotropy T_s exists, is used to determine epipolar geometry. Under a projective transformation T of 3D all projective image measurements are unaffected. In particular under T_s the configuration S is unchanged. However, the optical centers are moved by T_s (e.g. for configuration **I** they move on a surface). Consequently, it is not possible to determine the optical centers, C , since they can move on an orbit without affecting image measurements. To summarise:

Theorem *Suppose a configuration S is perspectively imaged onto one or more images. Then, if there exists an isotropy sub-group, T_s , for S under the group of 3D projective transformations, the optical center(s) cannot be uniquely determined from projective image measurements.*

If $\mathbf{x} = P\mathbf{X}$ and $\mathbf{X}' = T\mathbf{X}$, then for image points to be preserved, $\mathbf{x} = PT^{-1}T\mathbf{X} = P'\mathbf{X}'$, where $P' = PT^{-1}$. For the isotropy group and points on S ,

$\mathbf{X}' = \mathbf{T}_s \mathbf{X} = \mathbf{X}$. Consequently, $\mathbf{x} = \mathbf{P} \mathbf{T}_s^{-1} \mathbf{X} = \mathbf{P} \mathbf{X}$, so \mathbf{P} is not uniquely defined. If \mathbf{C} is a solution for the optical center, then any point on the orbit $\mathbf{T}_s^{-1} \mathbf{C}$ is also a solution, since $\mathbf{P}(\mathbf{T}_s^{-1} \mathbf{C}) = \mathbf{P} \mathbf{C} = 0$.

4.1 Twisted Cubics

In certain cases the orbit $\mathbf{T}_s^{-1} \mathbf{C}$ can be described easily, using the following theorem stated by Buchanan [2].

Theorem *Let \mathcal{R} be a set of reference points in real 3-space. Let \mathbf{P} be a projection of maximal rank from 3-space onto 2-space with $K = \text{Ker}(\mathbf{P})$. Then there exists a projection \mathbf{P}' , distinct from \mathbf{P} , with $\mathbf{P}(R) = \mathbf{P}'(R)$ for all $R \in \mathcal{R}$ if and only if*

- (a) $\mathcal{R} \cup \{K\}$ is a subset of a possibly composite twisted cubic or
- (b) \mathcal{R} is a subset of the set theoretic union of a line and a plane, and K lies on the line.

Note that the kernel $\text{Ker}(\mathbf{P})$ is the optical center of the camera. Buchanan lists the curves in case (a) as follows:

1. a proper (i.e. non singular) twisted cubic;
2. the union of a conic and a line which cuts the conic but which is not coplanar with the conic.
3. the union of two skew lines and a common transversal.

Properties of the Twisted Cubic Let c be a non-singular twisted cubic [23]. Then c is not contained in any plane of \mathcal{P}^3 ; it intersects a general plane at three distinct points. There is a unique c through six points in general position. The curve c is rational, and has a parameterisation of the form $t \mapsto \mathbf{T}(1, t, t^2, t^3)^\top$ where \mathbf{T} is an invertible 4×4 matrix. It follows that there is a collineation of \mathcal{P}^3 that maps c to a standard twisted cubic $t \mapsto (1, t, t^2, t^3)^\top$. Thus any two proper twisted cubics are projectively equivalent.

4.2 The Twisted Cubic and Camera Calibration

The twisted cubic plays a key rôle in (extrinsic) camera calibration because of the following property. Let \mathbf{A} , \mathbf{B} be two distinct points on the non-singular twisted cubic c , and \mathbf{X} a variable point on c . The chords $\langle \mathbf{X}, \mathbf{A} \rangle$, as \mathbf{X} moves on c , form a subset of the star of lines $\text{st}(\mathbf{A})$, and similarly for \mathbf{B} . Then there is a unique collineation $\mathbf{U} : \text{st}(\mathbf{A}) \rightarrow \text{st}(\mathbf{B})$ such that $\mathbf{U}(\langle \mathbf{X}, \mathbf{A} \rangle) = \langle \mathbf{X}, \mathbf{B} \rangle$ for all \mathbf{X} in c .

It is the existence of the collineation \mathbf{U} that is the basis of the above theorem quoted from [2]. If the true optical center \mathbf{C} and data points \mathbf{X}_i used in the camera calibration all lie on a twisted cubic c then all points on c are candidates for the optical center of the camera. Each putative optical center yields a camera calibration compatible with the \mathbf{X}_i and their projections \mathbf{x}_i to the image. There

is always a twisted cubic though six points in general position in space. Thus a unique camera calibration cannot be obtained from five or fewer data points.

We have seen that when there is an isotropy (for example as in configuration **II**) it is not possible to solve for the optical centers. Here we show that this ambiguity is a result of the points lying on a twisted cubic. Let Π be the plane containing the four points and let the optical center be a point \mathbf{C} not contained in Π . Then there is a unique conic through the five points $\mathbf{X}_{\Pi} = (\langle \mathbf{C}, \mathbf{X}_5 \rangle \cap \Pi)$ and $\mathbf{X}_i, i \in \{1, \dots, 4\}$. Consequently, the configuration **II** together with \mathbf{C} is contained in the degenerate case (2) of a twisted cubic, and there is not a unique solution for \mathbf{C} . Any point on the line is a candidate for the optical center. For configuration **VI**, there is no longer freedom to choose the line. In general, the configuration and the true optical center \mathbf{C} do not lie on a twisted cubic. Consequently, it is possible to obtain a unique solution for \mathbf{C} .

4.3 Critical Surfaces

The twisted cubic of Section 4.2 also appears in the theory of critical surfaces for reconstruction from two views of points [18]. Let ψ be a critical surface and let \mathbf{O}, \mathbf{C} be the optical centres from which the images of ψ are obtained. It is well known that ψ is in general a hyperboloid of one sheet and that it contains \mathbf{O}, \mathbf{C} . The reconstruction is ambiguous thus there is a second critical surface ϕ defining the same set of image correspondences as ψ . Let the optical centres for ϕ be \mathbf{O}, \mathbf{C}' . It can be shown that ψ and ϕ contain all three points $\mathbf{O}, \mathbf{C}, \mathbf{C}'$.

Let \mathbf{X} be a general point on ψ or ϕ . If the coordinates of \mathbf{X} are known then it is possible to distinguish the true reconstruction from the false. For example, if \mathbf{X} is on ψ but not on ϕ , then ψ is the true reconstruction and ϕ is false. There are certain exceptional points \mathbf{X} which cannot be used to distinguish between the two reconstructions even when their coordinates are known. These points are contained in $\psi \cap \phi$.

The intersection $\psi \cap \phi$ is a space curve of degree four which splits into a line \mathbf{L} and a twisted cubic c . The line \mathbf{L} contains \mathbf{O} whilst c contains \mathbf{C} and \mathbf{C}' . Let $\mathbf{x} \leftrightarrow \mathbf{x}'$ be a pair of corresponding image points such that \mathbf{x} is the image of \mathbf{L} . If the optical centre of the camera is at \mathbf{C} then there is a point \mathbf{X}_1 on \mathbf{L} projecting to \mathbf{x}' . Similarly, if the optical centre of the camera is at \mathbf{C}' then there is a point \mathbf{X}_2 on \mathbf{L} projecting to \mathbf{x}' . The points $\mathbf{X}_1, \mathbf{X}_2$ are in general different; this is possible because the entire line \mathbf{L} projects to a single point \mathbf{x} in the first image. The true reconstruction can be found once it is known which of the points $\mathbf{X}_1, \mathbf{X}_2$ gives rise to the correspondence $\mathbf{x} \leftrightarrow \mathbf{x}'$.

The true reconstruction cannot be distinguished from the false by giving the coordinates of the points on c . To see this, let \mathbf{X}_1 be a general point of c projecting to \mathbf{x} when the camera is at \mathbf{O} and to \mathbf{x}' when the camera is at \mathbf{C} . Let \mathbf{X}_2 be the point of ϕ projecting to \mathbf{x} at \mathbf{O} and to \mathbf{x}' at \mathbf{C}' . The line $\langle \mathbf{O}, \mathbf{X}_1 \rangle$ cuts ϕ at a unique point, \mathbf{X}_1 , distinct from \mathbf{O} . It follows that $\mathbf{X}_1 = \mathbf{X}_2$. The images of c taken from \mathbf{C} and \mathbf{C}' are indistinguishable. In the case of calibrated cameras this means indistinguishable up to a rotation of the camera; for uncalibrated cameras this means indistinguishable up to a collineation of

the image. Both cases are examples of the rôle of the twisted cubic in camera calibration, as described in Section 4.2.

5 Transfer

The process of rendering new images given only $image(s)$ of the original structure is known as *transfer*. Typically, two images are provided, and a new view of a 3D structure is rendered based on a number of “reference” points. Epipolar geometry is often used to carry out this process (see Barrett [21] and [3, 7]). For example, in the affine case (parallel projection) only four points are required to determine epipolar geometry. Suppose four points are in correspondence between two acquisition images, and a third target image. This is sufficient to determine the epipolar geometry between each acquisition image and the target image. Any other point imaged in both acquisition images can then be transferred by intersecting epipolar lines, as shown in Figure 7.

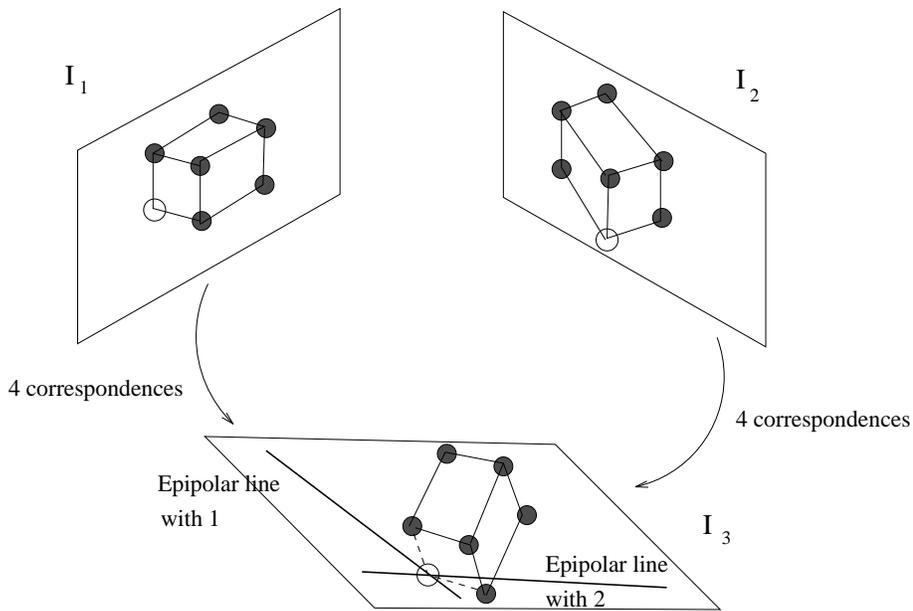


Fig. 7. Transfer by intersecting epipolar lines in the affine case.

Clearly, this process fails when epipolar lines are coincident, and becomes increasingly ill-conditioned as the lines become less “transverse”. This occurs in the affine case when a camera rotates about a fixed object axis. In this case the epipolar lines in the target image are parallel.

In the projective case at least seven points are needed to determine epipolar geometry. Transfer fails for 3D points on the tri-focal plane (the plane containing

all three optical centers). The epipolar lines for such points are coincident in the target image. The tri-focal plane intersects the target image in the line containing the epipoles of the two acquisition views. For points “close to” the trifocal plane, transfer will be very poorly conditioned. Examples of transfer using this method are shown in Barrett’s paper in this volume. The worst situation is when the three optical centers are collinear, in which case transfer is not possible for any point by epipolar line intersection.

This problem is avoided if structure is used. For example, consider the affine case with four points again in correspondence between the three images. From the two acquisition images the 3D affine coordinates of all corresponding points can be determined. The projection matrix P from this 3D affine structure to the target image is also determined from four points, and is used to project the 3D structure onto the target image. In fact, the actual reconstruction of 3D structure is not required [25]. A similar argument applies in the projective case, where seven points are required in general.

Shashua [26] derives a trilinear function of image coordinates in three views, which can be used for transfer. Unlike epipolar transfer, this method does not fail when the optical centers are collinear.

6 Discussion

We have demonstrated that for a number of configurations, although it is possible to recover invariants in a stable fashion from two views, it is not possible to recover structure or epipolar geometry. This is important in terms of stability: As is well known for configurations “close to” critical surfaces [13, 18] numerical procedures for recovering structure are ill-conditioned. In the same way, one would expect ill-conditioning in configurations close to an isotropy (e.g. if a significant part of the point set used to determine the epipolar geometry is near planar).

It has also been shown that transfer from two views to a third can be ill-conditioned, and may fail, if epipolar geometry alone is used.

Epipolar geometry is also not necessary for self-calibration. Hartley [11] describes a self-calibration method for the case of a camera rotating about its optical center, but not translating. In this case there is no epipolar geometry. Luong and Viéville [14] and Hartley (this volume), show that intrinsic parameters can be determined from projection matrices computed from two or more views of structure known up to an affine ambiguity.

It remains to be seen whether other procedures, such as structure recovery, are more tractable and better conditioned if epipolar geometry is displaced from its central role.

Acknowledgements

A number of people contributed to the ideas in this paper, particularly: David Forsyth, Richard Hartley, Joe Mundy, Larry Shapiro, Phil Torr and Charlie

Rothwell. Thanks to Amnon Shashua for questioning the relation of isotropies to critical surfaces, which lead to the clarification of the relationship with twisted cubics. Sabine Demey [3] carried out the numerical experiments.

References

1. Beardsley, P., Sinclair, D., Zisserman, A., Ego-motion from six points, *Insight meeting, Catholic University Leuven*, Feb. 1992.
2. Buchanan, T., The twisted cubic and camera calibration. *Computer Vision, Graphics, and Image Processing*, 42, 130-132, 1988.
3. Demey, S., Zisserman, A., and Beardsley, P., Affine and projective structure from motion, *Proc. BMVC92, Springer-Verlag*, 1992
4. Faugeras, O.D. and Maybank, S.J., Motion from point matches: multiplicity of solutions, *IJCV-4*, 225-246, 1990.
5. Faugeras, O.D., Luong, Q.T., and Maybank, S.J., Camera self-calibration: theory and experiments, *Proc. ECCV-92, Springer-Verlag*, 321 – 334, 1992.
6. Faugeras, O.D., What can be seen in three dimensions with an uncalibrated stereo rig?, *Proc. ECCV-92, Springer-Verlag*, 563 – 578, 1992.
7. Faugeras, O.D. and Robert, L., What can two images tell us about a third one? *Proc. ECCV-94, Springer-Verlag*, 485-492, 1994.
8. Gros, P. and Quan, L., Présentation de la théorie des invariants sous une forme utilisable en vision par ordinateur, *TR, LIFIA*, 1991.
9. Hartley, R.I., Invariants of points seen in multiple images, *G.E. Internal Report 1992*, also to appear *PAMI*.
10. Hartley, R.I., Gupta, R. and Chang, T., Stereo from uncalibrated cameras, *Proc. of CVPR-92*, 1992.
11. Hartley, R.I., Self-calibration of stationary cameras. *Proc. ECCV-94*, 1994.
12. Longuet-Higgins, H.C., A computer algorithm for reconstructing a scene from two projections, *Nature, Vol.*, 293, 10, Sept. 1981.
13. Luong, Q.T., *Matrice Fondamentale et Autocalibration en Vision par Ordinateur*. PhD thesis, Université de Paris-Sud, France, 1992.
14. Luong, Q.T. and T. Viéville. Canonic representations for the geometries of multiple projective views. Tech. Report UCB/CSD-93-772, University of California, Berkeley, USA, 1993.
15. Maybank, S.J., The projective geometry of ambiguous surfaces, *Phil. Trans. Roy. Soc. Lond. Series A Vol. 332*, 1-47, 1990.
16. Maybank, S.J., Properties of essential matrices, *Int. J. Imaging Systems and Technology*, 2, 380–384, 1990.
17. Maybank, S.J. and Faugeras, O.D., A theory of self-calibration of a moving camera. *International Journal of Computer Vision*, 8:123–151, 1992.
18. Maybank, S.J., *Theory of Reconstruction from Image Motion*. Springer Series in Information Science vol 28. Springer-Verlag: Berlin, 1993.
19. Mohr, R., Projective geometry and computer vision, *Handbook of Pattern Recognition and Computer Vision*, Chen, Pau and Wang editors, 1992.
20. Mohr, R., Sparr, R. and Faugeras, O., Multiple image invariants, *Proc. ESPRIT Workshop on Invariants in Vision, Santa Margherita Ligure*, 1992.
21. Mundy, J. L. and Zisserman, A. (editors), *Geometric Invariance in Computer Vision*, MIT Press, Cambridge Ma, 1992.

22. Quan, L. and Mohr, R., Towards structure from motion for linear features through reference points, *Proc. IEEE Workshop on Visual Motion*, 1991.
23. Semple, J.G. & Kneebone, G.T., *Algebraic Projective Geometry*. Oxford: Clarendon Press, 1953 (reprinted 1979).
24. Shapiro, L.S., Zisserman, A. and Brady, J.M., Motion from point matches using affine epipolar geometry, *Proc. ECCV-94*, 1994.
25. Shapiro, L.S., *Affine analysis of image sequences*, PhD thesis, Oxford University, 1993.
26. Shashua, A., Trilinearity in visual recognition by alignment, *Proc. ECCV-94*, 1994.
27. Szeliski, R. and Kang, S.B., Recovering 3D shape and motion from image streams using non-linear least squares, *DEC Tech. Rep.*, 93/3, 1993.
28. Torr, P.H.S. and Murray, D.W., Stochastic motion clustering, *Proc. ECCV-94*, Springer-Verlag, 328-337, 1994.