

follows from

$$F^T H = A^T [e']_{\times} (A - e' v^T) = A^T [e']_{\times} A$$

using $[e']_{\times} e' = 0$, since $A^T [e']_{\times} A$ is skew-symmetric.

Comparing (13.4) with the general decomposition of the fundamental matrix, as given in lemma 9.11(p255) or (9.10–p256) it is evident that they involve an identical formula (except for signs). In fact there is a one-to-one correspondence between decompositions of the fundamental matrix (up to the scale factor ambiguity k in lemma 9.11) and homographies induced by world planes, as stated here.

Corollary 13.4. *A transformation H is the homography between two images induced by some world plane if and only if the fundamental matrix F for the two images has a decomposition $F = [e']_{\times} H$.*

This choice in the decomposition simply corresponds to the choice of projective world frame. In fact, H is the transformation with respect to the plane with coordinates $(0, 0, 0, 1)^T$ in the reconstruction with $P = [I \mid 0]$ and $P' = [H \mid e']$.

Finding the plane that induces a given homography is a simple matter given a pair of camera matrices, as follows.

Result 13.5. *Given the cameras in the canonical form $P = [I \mid 0]$, $P' = [A \mid a]$, then the plane π that induces a given homography H between the views has coordinates $\pi = (v^T, 1)^T$ where v may be obtained linearly by solving the equations $\lambda H = A - av^T$, which are linear in the entries of v and λ .*

Note, these equations have an exact solution only if H satisfies the compatibility constraint (13.3) with F . For a homography computed numerically from noisy data this will not normally be true, and the linear system is over-determined.

13.2 Plane induced homographies given F and image correspondences

A plane in 3-space can be specified by three points, or by a line and a point, and so forth. In turn these 3D elements can be specified by image correspondences. In section 13.1 the homography was computed from the coordinates of the plane. In the following the homography will be computed directly from the corresponding image elements that specify the plane. This is a quite natural mechanism to use in applications.

We will consider two cases: (i) three points; (ii) a line and a point. In each case the corresponding elements are sufficient to determine a plane in 3-space uniquely. It will be seen that in each case:

- (i) The corresponding image entities have to satisfy *consistency constraints* with the epipolar geometry.
- (ii) There are *degenerate configurations* of the 3D elements and cameras for which the homography is not defined. Such degeneracies arise from collinearities and coplanarities of the 3D elements and the epipolar geometry. There may also be degeneracies of the solution method, but these can be avoided.

The three-point case is covered in more detail.




Fig. 13.3. **Degenerate geometry for an implicit computation of the homography.** The line defined by the points X_1 and X_2 lies in an epipolar plane, and thus intersects the baseline. The images of X_1 and X_2 are collinear with the epipole, and H cannot be computed uniquely from the correspondences $x_i \leftrightarrow x'_i, i \in \{1, \dots, 3\}, e \leftrightarrow e'$. This configuration is not degenerate for the explicit method.

13.2.1 Three points

We suppose that we have the images in two views of three (non-collinear) points X_i , and the fundamental matrix F . The homography H induced by the plane of the points may be computed in principle in two ways:

First, the position of the points X_i is recovered in a projective reconstruction (chapter 12). Then the plane π through the points is determined (3.3–p66), and the homography computed from the plane as in result 13.1. Second, the homography may be computed from four corresponding points, the four points in this case being the images of the three points X_i on the plane together with the epipole in each view. The epipole may be used as the fourth point since it is mapped between the views by the homography as shown in figure 13.2. Thus we have four correspondences, $x'_i = Hx_i, i \in \{1, \dots, 3\}, e' = He$, from which H may be computed.

We thus have two alternative methods to compute H from three point correspondences, the first involving an *explicit* reconstruction, the second an *implicit* one where the epipole provides a point correspondence. It is natural to ask if one has an advantage over the other, and the answer is that the implicit method should **not** be used for computation as it has significant degeneracies which are not present in the explicit method.

Consider the case when two of the image points are collinear with the epipole (we assume for the moment that the measurements are noise-free). A homography H cannot be computed from four correspondences if three of the points are collinear (see section 4.1.3(p91)), so the implicit method fails in this case. Similarly if the image points are close to collinear with the epipole then the implicit method will give a poorly conditioned estimate for H . The explicit method has no problems when two points are collinear or close to collinear with the epipole – the corresponding image points define points in 3-space (the world points are on the same epipolar plane, but this is not a degenerate situation) and the plane π and hence homography can be computed. The configuration is illustrated in figure 13.3.

We now develop the algebra of the explicit method in more detail. It is not neces-

sary to actually determine the coordinates of the points $\mathbf{x}_i$, all that is important is the constraint they place on the three-parameter family of homographies compatible with F (13.4), $H = A - \mathbf{e}'\mathbf{v}^T$, parametrized by $\mathbf{v}$. The problem is reduced to that of solving for $\mathbf{v}$ from the three point correspondences. The solution may be obtained as:

Result 13.6. *Given F and the three image point correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$, the homography induced by the plane of the 3D points is*

$$H = A - \mathbf{e}'(\mathbf{M}^{-1}\mathbf{b})^T,$$

where $A = [\mathbf{e}']_{\times} F$ and $\mathbf{b}$ is a 3-vector with components

$$b_i = (\mathbf{x}'_i \times (A\mathbf{x}_i))^T (\mathbf{x}'_i \times \mathbf{e}') / \|\mathbf{x}'_i \times \mathbf{e}'\|^2,$$

and $\mathbf{M}$ is a 3×3 matrix with rows $\mathbf{x}_i^T$.

Proof. According to result 9.14(p256), F may be decomposed as $F = [\mathbf{e}']_{\times} A$. Then (13.4) gives $H = A - \mathbf{e}'\mathbf{v}^T$, and each correspondence $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ generates a linear constraint on $\mathbf{v}$ as

$$\mathbf{x}'_i = H\mathbf{x}_i = A\mathbf{x}_i - \mathbf{e}'(\mathbf{v}^T \mathbf{x}_i), \quad i = 1, \dots, 3. \quad (13.5)$$

From (13.5) the vectors $\mathbf{x}'_i$ and $A\mathbf{x}_i - \mathbf{e}'(\mathbf{v}^T \mathbf{x}_i)$ are parallel, so their vector product is zero:

$$\mathbf{x}'_i \times (A\mathbf{x}_i - \mathbf{e}'(\mathbf{v}^T \mathbf{x}_i)) = (\mathbf{x}'_i \times A\mathbf{x}_i) - (\mathbf{x}'_i \times \mathbf{e}')(\mathbf{v}^T \mathbf{x}_i) = \mathbf{0}.$$

Forming the scalar product with the vector $\mathbf{x}'_i \times \mathbf{e}'$ gives

$$\mathbf{x}_i^T \mathbf{v} = \frac{(\mathbf{x}'_i \times (A\mathbf{x}_i))^T (\mathbf{x}'_i \times \mathbf{e}')}{(\mathbf{x}'_i \times \mathbf{e}')^T (\mathbf{x}'_i \times \mathbf{e}')} = b_i \quad (13.6)$$

which is linear in $\mathbf{v}$. Note, the equation is independent of the scale of $\mathbf{x}'$, since $\mathbf{x}'$ occurs the same number of times in the numerator and denominator. Each correspondence generates an equation $\mathbf{x}_i^T \mathbf{v} = b_i$, and collecting these together we have $\mathbf{M}\mathbf{v} = \mathbf{b}$. $\square$

Note, a solution cannot be obtained if $\mathbf{M}^T = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]$ is not of full rank. Algebraically, $\det \mathbf{M} = 0$ if the three image points $\mathbf{x}_i$ are collinear. Geometrically, three collinear image points arise from collinear world points, or coplanar world points where the plane contains the first camera centre. In either case a full rank homography is not defined.

Consistency conditions. Equation (13.5) is equivalent to six constraints since each point correspondence places two constraints on a homography. Determining $\mathbf{v}$ requires only three constraints, so there are three constraints remaining which must be satisfied for a valid solution. These constraints are obtained by taking the cross product of (13.5) with $\mathbf{e}'$, which gives

$$\mathbf{e}' \times \mathbf{x}'_i = \mathbf{e}' \times A\mathbf{x}_i = F\mathbf{x}_i.$$

Objective

Given F and three point correspondences $\mathbf{x}_i \mapsto \mathbf{x}'_i$ which are the images of 3D points $\mathbf{X}_i$, determine the homography $\mathbf{x}' = H\mathbf{x}$ induced by the plane of $\mathbf{X}_i$.

Algorithm

- (i) For each correspondence $\mathbf{x}_i \mapsto \mathbf{x}'_i$ compute the corrected correspondence $\tilde{\mathbf{x}}_i \mapsto \tilde{\mathbf{x}}'_i$ using algorithm 12.1(p318).
- (ii) Choose $A = [\mathbf{e}']_{\times} F$ and solve linearly for $\mathbf{v}$ from $M\mathbf{v} = \mathbf{b}$ as in result 13.6.
- (iii) Then $H = A - \mathbf{e}'\mathbf{v}^T$.

Algorithm 13.1. *The optimal estimate of the homography induced by a plane defined by three points.*

The equation $\mathbf{e}' \times \mathbf{x}'_i = F\mathbf{x}_i$ is a *consistency constraint* between $\mathbf{x}_i$ and $\mathbf{x}'_i$, since it is independent of $\mathbf{v}$. It is simply a (disguised) epipolar constraint on the correspondence $\mathbf{x}_i \mapsto \mathbf{x}'_i$: the LHS is the epipolar line through $\mathbf{x}'_i$, and the RHS is $F\mathbf{x}_i$ which is the epipolar line for $\mathbf{x}_i$ in the second image, i.e. the equation enforces that $\mathbf{x}'_i$ lie on the epipolar line of $\mathbf{x}_i$, and hence the correspondence is consistent with the epipolar geometry.

Estimation from noisy points. The three point correspondences which determine the plane and homography must satisfy the consistency constraint arising from the epipolar geometry. Generally *measured* correspondences $\mathbf{x}_i \mapsto \mathbf{x}'_i$ will not exactly satisfy this constraint. We therefore require a procedure for optimally correcting the measured points so that the estimated points $\tilde{\mathbf{x}}_i \mapsto \tilde{\mathbf{x}}'_i$ satisfy the epipolar constraint. Fortunately, such a procedure has already been given in the triangulation algorithm 12.1(p318), which can be adopted here directly. We then have a Maximum Likelihood estimate of H and the 3D points under Gaussian image noise assumptions. The method is summarized in algorithm 13.1.

13.2.2 A point and line

In this section an expression is derived for a plane specified by a point and line correspondence. We start by considering only the line correspondence and show that this reduces the three-parameter family of homographies compatible with F (13.4) to a 1-parameter family. It is then shown that the point correspondence uniquely determines the plane and corresponding homography.

The correspondence of two image lines determines a line in 3-space, and a line in 3-space lies on a one parameter family (a pencil) of planes, see figure 13.4. This pencil of planes induces a pencil of homographies between the two images, and any member of this family will map the corresponding lines to each other.

Result 13.7. *The homography for the pencil of planes defined by a line correspondence $l \mapsto l'$ is given by*

$$H(\mu) = [l']_{\times} F + \mu \mathbf{e}' l^T \quad (13.7)$$

provided $l^T \mathbf{e}' \neq 0$, where μ is a projective parameter.




Fig. 13.4. (a) Image lines l and l' determine planes π and π' respectively. The intersection of these planes defines the line L in 3-space. (b) The line L in 3-space is contained in a one parameter family of planes $\pi(\mu)$. This family of planes induces a one parameter family of homographies between the images.

Proof. From result 8.2(p197) the line l back-projects to a plane $P^T l$ through the first camera centre, and l' back-projects to a plane $P'^T l'$ through the second, see figure 13.4a. These two planes are the basis for a pencil of planes parametrized by μ . As in the proof of result 13.3 we may choose $P = [l \mid 0]$, $P' = [A \mid e']$, then the pencil of planes is

$$\begin{aligned}\pi(\mu) &= \mu P^T l + P'^T l' \\ &= \mu \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} A^T l' \\ e'^T l' \end{pmatrix}\end{aligned}$$

From result 13.1 the induced homography is $H(\mu) = A - e'v(\mu)^T$, with

$$v(\mu) = (\mu l + A^T l') / (e'^T l') \quad (13.8)$$

Using the decomposition $A = [e']_{\times} F$ we obtain

$$\begin{aligned}H &= ((e'^T l' I - e' l'^T) [e']_{\times} F - \mu e' l'^T) / (e'^T l') = -([l']_{\times} [e']_{\times} [e']_{\times} F + \mu e' l'^T) / (e'^T l') \\ &= -([l']_{\times} F + \mu e' l'^T) / (e'^T l')\end{aligned}$$

where the last equality follows from result A4.4(p582) that $[e']_{\times} [e']_{\times} F = F$. This is equivalent to (13.7) up to scale. $\square$

The homography for a corresponding point and line. From the line correspondence we have that $H(\mu) = [l']_{\times} F + \mu e' l'^T$, and now solve for μ using the point correspondence $x \leftrightarrow x'$.

Result 13.8. Given F and a corresponding point $x \leftrightarrow x'$ and line $l \leftrightarrow l'$, the homography induced by the plane of the 3-space point and line is

$$H = [l']_{\times} F + \frac{(x' \times e')^T (x' \times ((Fx) \times l'))}{\|x' \times e'\|^2 (l'^T x)} e' l'^T.$$

The derivation is analogous to that of result 13.6. As in the three-point case, the image point correspondence must be consistent with the epipolar geometry. This means




Fig. 13.5. The epipolar geometry induces a homography between corresponding lines $l \leftrightarrow l'$ which are the images of a line L in 3-space. The points on l are mapped to points on l' as $x' = [l']_{\times} Fx$, where x and x' are the images of the intersection of L with the epipolar plane corresponding to l_e and l'_e .

that the measured (noisy) points must be corrected using algorithm 12.1(p318) before result 13.8 is applied. There is no consistency constraint on the line, and no correction is available.

Geometric interpretation of the point map $H(\mu)$. It is worth exploring the map $H(\mu)$ further. Since $H(\mu)$ is compatible with the epipolar geometry, a point x in the first view is mapped to a point $x' = H(\mu)x$ in the second view on the epipolar line Fx corresponding to x . In general the position of the point $x' = H(\mu)x$ on the epipolar line varies with μ . However, if the point x lies on l (so that $l^T x = 0$) then

$$x' = H(\mu)x = ([l']_{\times} F + \mu e' l^T)x = [l']_{\times} Fx$$

which is independent of the value of μ , depending only on F . Thus as shown in figure 13.5 the epipolar geometry defines a point-to-point map for points on the line.

Degenerate homographies. As has already been stated, if the world plane contains one of the camera centres, then the induced homography is degenerate. The matrix representing the homography does not have full rank, and points on one plane are mapped to a line (if $\text{rank } H = 2$) or a point (if $\text{rank } H = 1$). However, an explicit expression can be obtained for a degenerate homography from (13.7). The degenerate (singular) homographies in this pencil are at $\mu = \infty$ and $\mu = 0$. These correspond to planes through the first and second camera centres respectively. Figure 13.6 shows the case where the plane contains the second camera centre, and intersects the image plane in the line l' . A point x in the first view is imaged on l' at the point x' where

$$x' = l' \times Fx = [l']_{\times} Fx.$$

The homography is thus $H = [l']_{\times} F$. This is a rank 2 matrix.

13.3 Computing F given the homography induced by a plane

Up to now it has been assumed that F is given, and the objective is to compute H when various additional information is provided. We now reverse this, and show that if H is given then F may be computed when additional information is provided. We start by introducing an important geometric idea, that of parallax relative to a plane, which will make the algebraic development straightforward.




Fig. 13.6. **A degenerate homography.** (a) The map induced by a plane through the second camera centre is a degenerate homography $H = [l']_x F$. The plane π' intersects the second image plane in the line l' . All points in the first view are mapped to points on l' in the second. (b) A point x in the first view is imaged at x' , the intersection of l' with the epipolar line Fx of x , so that $x' = l' \times Fx$.




Fig. 13.7. **Plane induced parallax.** The ray through X intersects the plane π at the point X_π . The images of X and X_π are coincident points at x in the first view. In the second view the images are the points x' and $\tilde{x}' = Hx$ respectively. These points are not coincident (unless X is on π), but both are on the epipolar line l'_x of x . The vector between the points x' and $\tilde{x}'$ is the parallax relative to the homography induced by the plane π . Note that if X is on the other side of the plane, then $\tilde{x}'$ will be on the other side of x' .

Plane induced parallax. The homography induced by a plane generates a virtual parallax (see section 8.4.5(p207)) as illustrated schematically in figure 13.7 and by example in figure 13.8. The important point here is that in the second view x' , the image of the 3D point X , and $\tilde{x}' = Hx$, the point mapped by the homography, are on the epipolar line of x ; since both are images of points on the ray through x . Consequently, the line $x' \times (Hx)$ is an epipolar line in the second view and provides a constraint on the position of the epipole. Once the epipole is determined (two such constraints suffice), then as shown in result 9.1(p243) $F = [e']_x H$ where H is the homography induced by any plane. Similarly it can be shown that $F = H^{-T} [e]_x$.

As an application of virtual parallax it is shown in algorithm 13.2 that F can be computed uniquely from the images of six points, four of which are coplanar and two are off the plane. The images of the four coplanar points define the homography, and the two points off the plane provide constraints sufficient to determine the epipole. The



Fig. 13.8. **Plane induced parallax.** (a) (b) Left and right images. (c) The left image is superimposed on the right using the homography induced by the plane of the Chinese text. The transferred and imaged planes exactly coincide. However, points off the plane (such as the mug) do not coincide. Lines joining corresponding points off the plane in the "superimposed" image intersect at the epipole.

six-point result is quite surprising since for seven points in general position there are 3 solutions for F (see section 11.1.2(p281)).

Objective

Given six point correspondences $\mathbf{x}_i \mapsto \mathbf{x}'_i$ which are the images of 3-space $\mathbf{X}_i$, with the first four 3-space points $i \in \{1, \dots, 4\}$ coplanar, determine the fundamental matrix F .

Algorithm

- (i) Compute the homography H , such that $\mathbf{x}'_i = H\mathbf{x}_i, i \in \{1, \dots, 4\}$.
- (ii) Determine the epipole $\mathbf{e}'$ as the intersection of the lines $(H\mathbf{x}_5) \times \mathbf{x}'_5$ and $(H\mathbf{x}_6) \times \mathbf{x}'_6$.
- (iii) Then $F = [\mathbf{e}']_{\times} H$.

See figure 13.9.

Algorithm 13.2. Computing F given the correspondences of six points of which four are coplanar.




Fig. 13.9. The fundamental matrix is defined uniquely by the image of six 3D points, of which four are coplanar. (a) The parallax for one point $\mathbf{X}$. (b) The epipole determined by the intersection of two parallax lines: the line joining $\tilde{\mathbf{x}}'_5 = H\mathbf{x}_5$ to $\mathbf{x}'_5$, and the join of $\tilde{\mathbf{x}}'_6 = H\mathbf{x}_6$ to $\mathbf{x}'_6$.



Fig. 13.10. **Binary space partition.** (a) (b) Left and right images. (c) Points whose correspondence is known. (d) A triplet of points selected from (c). This triplet defines a plane. The points in (c) can then be classified according to their side of the plane. (e) Points on one side. (f) Points on the other side.

Projective depth. A world point $\mathbf{X} = (\mathbf{x}^\top, \rho)^\top$ is imaged at $\mathbf{x}$ in the first view and at

$$\mathbf{x}' = \mathbf{H}\mathbf{x} + \rho\mathbf{e}' \quad (13.9)$$

in the second. Note that $\mathbf{x}'$, $\mathbf{e}'$ and $\mathbf{H}\mathbf{x}$ are collinear. The scalar ρ is the parallax *relative* to the homography $\mathbf{H}$, and may be interpreted as a “depth” relative to the plane π . If $\rho = 0$ then the 3D point $\mathbf{X}$ is on the plane, otherwise the “sign” of ρ indicates which ‘side’ of the plane π the point $\mathbf{X}$ is (see figure 13.7 and figure 13.8). These statements should be taken with care because in the absence of oriented projective geometry the sign of a homogeneous object, and the side of a plane have no meaning.

Example 13.9. Binary space partition. The sign of the virtual parallax ($\text{sign}(\rho)$) may be used to compute a partition of 3-space by the plane π . Suppose we are given $\mathbf{F}$ and three space points are specified by their corresponding image points. Then the plane defined by the three points can be used to partition all other correspondences into sets on either side of (or on) the plane. Figure 13.10 shows an example. Note, the three points need not actually correspond to images of physical points so the method can be applied to virtual planes. By combining several planes a region of 3-space can be identified. $\triangle$

Two planes. Suppose there are two planes, π_1, π_2 , in the scene which induce homographies $\mathbf{H}_1, \mathbf{H}_2$ respectively. With the idea of parallax in mind it is clear that because each plane provides off-plane information about the other, the two homographies




Fig. 13.11. The action of the map $H = H_2^{-1}H_1$ on a point x in the first image is first to transfer it to x' as though it were the image of the 3D point X_1 , and then map it back to the first image as though it were the image of the 3D point X_2 . Points in the first view which lie on the imaged line of intersection of the two planes will be mapped to themselves, so are fixed points under this action. The epipole e is also a fixed point under this map.

should be sufficient to determine F . Indeed F is over-determined by this configuration which means that the two homographies must satisfy consistency constraints.

Consider figure 13.11. The homography $H = H_2^{-1}H_1$ is a mapping from the first image onto itself. Under this mapping the epipole e is a fixed point, i.e. $He = e$, so may be determined from the (non-degenerate) eigenvector of H . The fundamental matrix may then be computed from result 9.1(p243) as $F = [e']_{\times}H_i$, where $e' = H_i e$ for $i = 1$ or 2 . The map H has further properties which may be seen from figure 13.11. The map has a line of fixed points and a fixed point not on the line (see section 2.9(p61) for fixed points and lines). This means that two of the eigenvalues of H are equal. In fact H is a planar homology (see section A7.2(p629)). In turn, these properties of $H = H_2^{-1}H_1$ define consistency constraints on H_1 and H_2 in order that their composition has these properties.

Up to this point the results of this chapter have been entirely projective. Now an affine element is introduced.

13.4 The infinite homography H_{∞}

The plane at infinity is a particularly important plane, and the homography induced by this plane is distinguished by a special name:

Definition 13.10. The infinite homography, H_{∞} , is the homography induced by the plane at infinity, π_{∞} .

The form of the homography may be derived by a limiting process starting from (13.2–p327), $H = K'(R - \mathbf{tn}^T/d)K^{-1}$, where d is the orthogonal distance of the plane from




Fig. 13.12. The infinite homography H_∞ maps vanishing points between the images.

the first camera:

$$H_\infty = \lim_{d \rightarrow \infty} H = K'RK^{-1}.$$

This means that H_∞ does not depend on the translation between views, only on the rotation and camera internal parameters. Alternatively, from (9.7-p250) corresponding image points are related as

$$x' = K'RK^{-1}x + K't/Z = H_\infty x + K't/Z \quad (13.10)$$

where Z is the depth measured from the first camera. Again it can be seen that points at infinity ($Z = \infty$) are mapped by H_∞ . Note also that H_∞ is obtained if the translation t is zero in (13.10), which corresponds to a rotation about the camera centre. Thus H_∞ is the homography that relates image points of *any* depth if the camera rotates about its centre (see section 8.4(p202)).

Since $e' = K't$, (13.10) can be written as $x' = H_\infty x + e'/Z$, and comparison with (13.9) shows that $(1/Z)$ plays the role of ρ . Thus Euclidean inverse depth can be interpreted as parallax relative to π_∞ .

Vanishing points and lines. Images of points on π_∞ are mapped by H_∞ . These images are vanishing points, and so H_∞ maps vanishing points between images, i.e. $v' = H_\infty v$, where v' and v are corresponding vanishing points. See figure 13.12. Consequently, H_∞ can be computed from the correspondence of three (non-collinear) vanishing points together with F using result 13.6. Alternatively, H_∞ can be computed from the correspondence of a vanishing line and the correspondence of a vanishing point (not on the line), together with F , as described in section 13.2.2.

Affine and metric reconstruction. As we have seen in chapter 10, specifying π_∞ enables a projective reconstruction to be upgraded to an affine reconstruction. Not surprisingly, because of its association with π_∞ , H_∞ arises naturally in the rectification. Indeed, if the camera matrices are chosen as $P = [I \mid 0]$ and $P' = [H_\infty \mid \lambda e']$ then the reconstruction is affine.

Conversely, suppose the world coordinate system is affine (i.e. π_∞ has its canonical position at $\pi_\infty = (0, 0, 0, 1)^T$); then H_∞ may be determined directly from the camera projection matrices. Suppose M, M' are the first 3×3 submatrix of P and P' respectively.




Fig. 13.13. **Reducing the search region using H_∞ .** Points in 3-space are no 'further' away than π_∞ . H_∞ captures this constraint and limits the search on the epipolar line in one direction. The baseline between the cameras partitions each epipolar plane into two. A point on one "side" of the epipolar line in the left image will be imaged on the corresponding "side" of the epipolar line in the right image (indicated by the solid line in the figure). The epipole thus bounds the search region in the other direction.

Then a point $\mathbf{X} = (\mathbf{x}_\infty^T, 0)^T$ on π_∞ is imaged at $\mathbf{x} = \mathbf{P}\mathbf{X} = \mathbf{M}\mathbf{x}_\infty$ and $\mathbf{x}' = \mathbf{P}'\mathbf{X} = \mathbf{M}'\mathbf{x}_\infty$ in the two views. Consequently $\mathbf{x}' = \mathbf{M}'\mathbf{M}^{-1}\mathbf{x}$ and so

$$H_\infty = \mathbf{M}'\mathbf{M}^{-1}. \quad (13.11)$$

The homography H_∞ may be used to propagate camera calibration from one view to another. The absolute conic Ω_∞ resides on π_∞ , and its image, ω , is mapped between images by H_∞ according to result 2.13(p37): $\omega' = H_\infty^{-T}\omega H_\infty^{-1}$. Thus if $\omega = (\mathbf{K}\mathbf{K}^T)^{-1}$ is specified in one view, then ω' , the image of Ω_∞ in a second view, can be computed via H_∞ , and the calibration for that view determined from $\omega' = (\mathbf{K}'\mathbf{K}'^T)^{-1}$. Section 19.5.2(p475) describes applications of H_∞ to camera auto-calibration.

Stereo correspondence. H_∞ limits the search region when searching for correspondences. The region is reduced from the entire epipolar line to a bounded line. See figure 13.13. However, a correct application of this constraint requires oriented projective geometry.

13.5 Closure

This chapter has illustrated a raft of projective techniques for a plane that may be applied to many other surfaces. A plane is a simple parametrized surface with 3 degrees of freedom. A very similar development can be given for other surfaces where the degrees of freedom are determined from images of points on the surface. For example in the case of a quadric the surface can be determined both from images of points on its surface, and/or (an extension not possible for planes) from its outline in each view [Cross-98, Shashua-97]. The ideas of surface induced transfer, families of surfaces when the surface is not fully determined from its images, surface induced parallax, consistency constraints, implicit computations, degenerate geometries etc. all carry over to other surfaces.

13.5.1 The literature

The compatibility of epipolar geometry and induced homographies is investigated thoroughly by Luong & Viéville [Luong-96]. The six-point solution for F appeared in Beardsley *et al.* [Beardsley-92] and [Mohr-92]. The solution for F given two planes appeared in Sinclair [Sinclair-92]. [Zeller-96] gives many examples of configurations whose properties may be determined using only epipolar geometry and their image projections. He also catalogues their degenerate cases.

13.5.2 Notes and exercises

(i) Homography induced by a plane (13.1-p326).

- (a) The inverse of the homography H is given by

$$H^{-1} = A^{-1} \left(I + \frac{\mathbf{a}\mathbf{v}^T A^{-1}}{1 - \mathbf{v}^T A^{-1} \mathbf{a}} \right)$$

provided A^{-1} exists. This is sometimes called the Sherman-Morrison formula.

- (b) Show that the homography H is degenerate if the plane contains the second camera centre. Hint, in this case $\mathbf{v}^T A^{-1} \mathbf{a} = 1$, and note that $H = A(I - A^{-1} \mathbf{a}\mathbf{v}^T)$.
- (ii) Show that if the camera undergoes planar motion, i.e. the translation is parallel to the plane and the rotation is parallel to the plane normal, then the homography induced by the plane is conjugate to a planar Euclidean transformation. Show that the fixed points of the homography are the images of the plane's circular points.
- (iii) Using (13.2-p327) show that if a camera undergoes a pure translation then the homography induced by the plane is a planar homology (as defined in section A7.2(p629)), with a line of fixed points corresponding to the vanishing line of the plane. Show further that if the translation is parallel to the plane then the homography is an elation (as defined in section A7.3(p631)).
- (iv) Show that a necessary, but not sufficient, condition for two space lines to be coplanar is $(\mathbf{l}'_1 \times \mathbf{l}'_2)^T F(\mathbf{l}_1 \times \mathbf{l}_2) = 0$. Why is it not a sufficient condition?
- (v) **Intersections of lines and planes.** Verify each of the following results by sketching the configuration assuming general position. In each case determine the degenerate configurations for which the result is not valid.
- (a) Suppose the line L in 3-space is imaged as l and l' , and the plane π induces the homography $\mathbf{x}' = H\mathbf{x}$. Then the point of intersection of L with π is imaged at $\mathbf{x} = l \times (H^T l')$ in the first image, and at $\mathbf{x}' = l' \times (H^{-T} l)$ in the second.
- (b) The infinite homography may be used to find the vanishing point of a line seen in two images. If l and l' are corresponding lines in two images, and $\mathbf{v}, \mathbf{v}'$ their vanishing points in each image, then $\mathbf{v} = l \times (H_\infty^T l')$, $\mathbf{v}' = l' \times (H_\infty^{-T} l)$.

(c) Suppose the planes π_1 and π_2 induce homographies $\mathbf{x}' = H_1\mathbf{x}$ and $\mathbf{x}' = H_2\mathbf{x}$ respectively. Then the image of the line of intersection of π_1 with π_2 in the first image obeys $H_1^T H_2^{-T} \mathbf{l} = 1$ and may be determined from the real eigenvector of the planar homology $H_1^T H_2^{-T}$ (see figure 13.11).

- (vi) **Coplanarity of four points.** Suppose F is known, and four corresponding image points $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ are supplied. How can it be determined whether their pre-images are coplanar? One possibility is to use three of the points to determine a homography via result 13.6(p331), and then measure the transfer error of the fourth point. A second possibility is to compute lines joining the image points, and determine if the line intersection obeys the epipolar constraint (see [Faugeras-92b]). A third possibility is to compute the cross-ratio of the four lines from the epipole to the image points – if the four scene points are coplanar then this cross-ratio will be the same in both images. Thus this equality is a necessary condition for co-planarity, but is it a sufficient condition also? What statistical tests should be applied when there is measurement error (noise)?
- (vii) Show that the epipolar geometry can be computed uniquely from the images of four coplanar lines and two points off the plane of the lines. If two of the lines are replaced by points can the epipolar geometry still be computed?
- (viii) Starting from the camera matrices $P = [M \mid \mathbf{m}]$, $P' = [M' \mid \mathbf{m}']$ show that the homography $\mathbf{x}' = H\mathbf{x}$ induced by a plane $\pi = (\hat{\pi}^T, \pi_4)^T$ is given by

$$H = M'(I - t\mathbf{v}^T)M^{-1} \text{ with } t = (M'^{-1}\mathbf{m}' - M^{-1}\mathbf{m}), \text{ and } \mathbf{v} = \hat{\pi}/(\pi_4 - \hat{\pi}^T M^{-1}\mathbf{m}),$$

- (ix) Show that the homography computed as in result 13.6(p331) is independent of the scale of F . Start by choosing an arbitrary fixed scale for F , so that F is no longer a homogeneous quantity, but a matrix $\hat{F}$ with fixed scale. Show that if $H = [\mathbf{e}']_{\times} \hat{F} - \mathbf{e}'(M^{-1}\hat{\mathbf{b}})^T$ with $\hat{\mathbf{b}}_i = \mathbf{c}'^T(\hat{F}\mathbf{x}_i)$, then replacing $\hat{F}$ by $\lambda\hat{F}$ simply scales H by λ .
- (x) Given two perspective images of a (plane) conic and the fundamental matrix between the views, then the plane of the conic (and consequently the homography induced by this plane) is defined up to a two-fold ambiguity. Suppose the image conics are C and C' , then the induced homography is $H(\mu) = [C'e']_{\times} F - \mu \mathbf{e}'(Ce)^T$, with the two values of μ obtained from

$$\mu^2 [(e^T Ce)C - (Ce)(Ce)^T] (e'^T C'e') = -F^T [C'e']_{\times} C' [C'e']_{\times} F.$$

Details are given in [Schmid-98].

- (a) By considering the geometry, show that to be compatible with the epipolar geometry the conics must satisfy the consistency constraint that epipolar tangents are corresponding epipolar lines (see figure 11.6-(p295)). Now derive this result algebraically starting from $H(\mu)$ above.
- (b) The algebraic expressions are not valid if the epipole lies on the conic (since then $e^T Ce = e'^T C'e' = 0$). Is this a degeneracy of the geometry or of the expression alone?

- (xi) **Fixed points of a homography induced by a plane.** A planar homography H has up to three distinct fixed points corresponding to the three eigenvectors of the 3×3 matrix (see section 2.9(p61)). The fixed points are images of points on the plane for which $\mathbf{x}' = H\mathbf{x} = \mathbf{x}$. The horopter is the locus of *all* points in 3-space for which $\mathbf{x} = \mathbf{x}'$. It is a twisted cubic curve passing through the two camera centres. A twisted cubic intersects a plane in three points, and these are the three fixed points of the homography induced by that plane.
- (xii) **Estimation.** Suppose $n > 3$ points $\mathbf{X}_i$ lie on a plane in 3-space and we wish to optimally estimate the homography induced by the plane given F and their image correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$. Then the ML estimate of the homography (assuming independent Gaussian measurement noise as usual) is obtained by estimating the plane $\hat{\pi}$ (3 dof) and the n points $\hat{\mathbf{X}}_i$ (2 dof each, since they lie on a plane) which minimizes reprojection error for the n points.

Affine Epipolar Geometry

This chapter recapitulates the developments and objectives of the previous chapters on two-view geometry, but here with affine cameras replacing projective cameras. The affine camera is an extremely usable and well conditioned approximation in many practical situations. Its great advantage is that, because of its linearity, many of the optimal algorithms can be implemented by linear algebra (matrix inverses, SVD etc.), whereas in the projective case solutions either involve high order polynomials (such as for triangulation) or are only possible by numerical minimization (such as in the Gold Standard estimation of F).

We first describe properties of the epipolar geometry of two affine cameras, and its optimal computation from point correspondences. This is followed by triangulation, and *affine* reconstruction. Finally the ambiguities in reconstruction that result from parallel projection are sketched, and the non-ambiguous motion parameters are computed from the epipolar geometry.

14.1 Affine epipolar geometry

In many respects the epipolar geometry of two affine cameras is identical to that of two perspective cameras, for example a point in one view defines an epipolar line in the other view, and the pencil of such epipolar lines intersect at the epipole. The difference is that because the cameras are affine their centres are at infinity, and there is parallel projection from scene to image. This leads to certain simplifications in the affine epipolar geometry:

Epipolar lines. Consider two points, x_1, x_2 , in the first view. These points back-project to rays which are *parallel* in 3-space, since all projection rays are parallel. In the second view an epipolar line is the image of a back-projected ray. The images of these two rays in the second view are also parallel since an affine camera maps parallel scene lines to parallel images lines. Consequently, all epipolar lines are parallel, as are the epipolar planes.

The epipoles. Since epipolar lines intersect in the epipole, and all epipolar lines are parallel, it follows that the epipole is at infinity.




Fig. 14.1. **Affine epipolar geometry.** (a) *Correspondence geometry:* Projection rays are parallel and intersect at infinity. A point x back-projects to a ray in 3-space defined by the first camera centre (at infinity) and x . This ray is imaged as a line l' in the second view. The 3-space point X which projects to x lies on this ray, so the image of X in the second view lies on l' . (b) *Epipolar lines and planes are parallel.*

These points are illustrated schematically in figure 14.1, and examples on images are shown in figure 14.2.

14.2 The affine fundamental matrix

The affine epipolar geometry is represented algebraically by a matrix termed the *affine fundamental matrix*, F_A . It will be seen in the following that:

Result 14.1. *The fundamental matrix resulting from two cameras with the affine form (i.e. the third row is $(0, 0, 0, 1)$) has the form*

$$F_A = \begin{bmatrix} 0 & 0 & * \\ 0 & 0 & * \\ * & * & * \end{bmatrix}$$

where $*$ indicates a non-zero entry.

It will be convenient to write the five non-zero entries as

$$F_A = \begin{bmatrix} 0 & 0 & a \\ 0 & 0 & b \\ c & d & e \end{bmatrix}. \quad (14.1)$$

Note that in general F_A has rank 2.

14.2.1 Derivation

Geometric derivation. This derivation is the analogue of that given in section 9.2.1- (p242) for a pair of projective cameras. The map from a point in one image to the corresponding epipolar line in the other image is decomposed into two steps, as illustrated in figure 14.5 on page 352:




Fig. 14.2. **Affine epipolar lines.** (a), (b) Two views of a hole punch acquired under affine imaging conditions. For the points marked in (c) the epipolar lines are superimposed on (d). Note that corresponding points lie on their epipolar lines, and that all epipolar lines are parallel. The epipolar geometry is computed from point correspondences using algorithm 14.1. (e) and (f) show the “flow” for selected points in the image (the lines link a point in one image to the point’s position in the other image). This demonstrates that even though the epipolar lines are parallel the movement of imaged points between the views contains both rotational and translational components.

- (i) **Point transfer via a plane π .** Since both cameras are affine, points are mapped between an image and a scene plane by parallel projection, so the map between π and the images is a planar affine transformation; the composition of the affine transformations between the first view and π , and π and the second view, is also an affine transformation, i.e. $\mathbf{x}' = \mathbf{H}_A \mathbf{x}$.
- (ii) **Constructing the epipolar line.** The epipolar line is obtained as the line through $\mathbf{x}'$ and the epipole $\mathbf{e}'$, i.e. $\mathbf{l}' = \mathbf{e}' \times \mathbf{H}_A \mathbf{x} = \mathbf{F}_A \mathbf{x}$, so that $\mathbf{F}_A = [\mathbf{e}']_{\times} \mathbf{H}_A$.

We now take note of the special forms of the affine matrix H_A , and the skew matrix $[e']_{\times}$ when e' is at infinity, and so has a zero last element:

$$F_A = [e']_{\times} H_A = \begin{bmatrix} 0 & 0 & * \\ 0 & 0 & * \\ * & * & 0 \end{bmatrix} \begin{bmatrix} * & * & * \\ * & * & * \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 & * \\ 0 & 0 & * \\ * & * & * \end{bmatrix} \quad (14.2)$$

where $*$ indicates a non-zero entry. This derives the affine form of F using only the geometric properties that the camera centres are on the plane at infinity.

Algebraic derivation. In the case that the cameras are both affine, the affine form of the fundamental matrix is obtained directly from the expression (9.1-p244) for F in terms of the pseudo-inverse, namely $F = [e']_{\times} P^+ P^+$, where $e' = P^+ C$, with C the camera centre which is the null-vector of P . Details are left as an exercise. An elegant derivation of F_A in terms of determinants formed from rows of the affine camera matrices is given in section 17.1.2(p413).

14.2.2 Properties

The affine fundamental matrix is a homogeneous matrix with five non-zero elements, it thus has 4 degrees of freedom. These are accounted as: one for each of the two epipoles (the epipoles lie on l_{∞} , so only their direction need be specified); and two for the 1D affine transformation mapping the pencil of epipolar lines from one view to the other.

The geometric entities (epipoles etc.) are encoded in F_A in the same manner as their encoding in F . However, often the expressions are far simpler and so can be given explicitly.

The epipoles. The epipole in the first view is the right null-vector of F_A , i.e. $F_A e = 0$. This determines $e = (-d, c, 0)^T$, which is a point (direction) on l_{∞} . Since all epipolar lines intersect the epipole this shows that all epipolar lines are parallel.

Epipolar lines. The epipolar line in the second view corresponding to x in the first is $l' = F_A x = (a, b, cx + dy + e)^T$. Again it is evident that all epipolar lines are parallel since the line orientation, (a, b) , is independent of (x, y) .

These properties, and others, are summarized in table 14.1.

14.3 Estimating F_A from image point correspondences

The fundamental matrix is defined by the equation $x'^T F_A x = 0$ for any pair of matching points $x \leftrightarrow x'$ in two images. Given sufficiently many point matches $x_i \leftrightarrow x'_i$, this equation can be used to compute the unknown matrix F_A . In particular, writing $x_i = (x_i, y_i, 1)^T$ and $x'_i = (x'_i, y'_i, 1)^T$ each point match gives rise to one linear equation

$$ax'_i + by'_i + cx_i + dy_i + e = 0 \quad (14.3)$$

in the unknown entries $\{a, b, c, d, e\}$ of F_A .

- F_A is a rank 2 homogeneous matrix with 4 degrees of freedom. It has the form

$$F_A = \begin{bmatrix} 0 & 0 & a \\ 0 & 0 & b \\ c & d & e \end{bmatrix}.$$

- **Point correspondence:** If $\mathbf{x}$ and $\mathbf{x}'$ are corresponding image points under an affine camera, then $\mathbf{x}'^T F_A \mathbf{x} = 0$. For finite points

$$ax' + by' + cx + dy + e = 0.$$

- **Epipolar lines:**

- ◊ $\mathbf{l}' = F_A \mathbf{x} = (a, b, cx + dy + e)^T$ is the epipolar line corresponding to $\mathbf{x}$.
- ◊ $\mathbf{l} = F_A^T \mathbf{x}' = (c, d, ax + by + e)^T$ is the epipolar line corresponding to $\mathbf{x}'$.

- **Epipoles:**

- ◊ From $F_A \mathbf{e} = \mathbf{0}$, $\mathbf{e} = (-d, c, 0)^T$.
- ◊ From $F_A^T \mathbf{e}' = \mathbf{0}$, $\mathbf{e}' = (-b, a, 0)^T$.

- **Computation from camera matrices P_A, P'_A :**

- ◊ General cameras,
 $F_A = [\mathbf{e}']_{\times} P'_A P_A^+$, where P_A^+ is the pseudo-inverse of P_A , and $\mathbf{e}'$ is the epipole defined by $\mathbf{e}' = P'_A \mathbf{C}$, where $\mathbf{C}$ is the centre of the first camera.
- ◊ Canonical cameras,

$$P_A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad P'_A = \begin{bmatrix} M_{2 \times 3} & \mathbf{t} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\begin{aligned} a &= m_{23}, & b &= -m_{13}, & c &= m_{13}m_{21} - m_{11}m_{23} \\ d &= m_{13}m_{22} - m_{12}m_{23}, & e &= m_{13}t_2 - m_{23}t_1 \end{aligned}$$

Table 14.1. Summary of affine fundamental matrix properties.

14.3.1 The linear algorithm

In the usual manner a solution for F_A may be obtained by rewriting (14.3) as

$$(x'_i, y'_i, x_i, y_i, 1) \mathbf{f} = 0$$

with $\mathbf{f} = (a, b, c, d, e)^T$. From a set of n point matches, we obtain a set of linear equations of the form $A\mathbf{f} = \mathbf{0}$, where A is a $n \times 5$ matrix:

$$\begin{bmatrix} x'_1 & y'_1 & x_1 & y_1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ x'_n & y'_n & x_n & y_n & 1 \end{bmatrix} \mathbf{f} = \mathbf{0}.$$

A minimal solution is obtained for $n = 4$ point correspondences as the right null-space of the 4×5 matrix A . Thus F_A can be computed uniquely from only 4 point

correspondences, provided the 3-space points are in general position. The conditions for general position are described in section 14.3.3 below.

If there are more than 4 correspondences, and the data is not exact, then the rank of A may be greater than 4. In this case, one may find a least-squares solution, subject to $\|f\| = 1$, in essentially the same manner as that of section 4.1.1(p90), as the singular vector corresponding to the smallest singular value of A . Refer to algorithm 4.2(p109) for details. This linear solution is the equivalent of the 8-point algorithm 11.1(p282) for the computation of a general fundamental matrix. We do not recommend this approach for estimating F_A because the Gold Standard algorithm described below may be implemented with equal computational ease, and in general will have superior performance.

The singularity constraint. The form (14.1) of F_A ensures that the matrix has rank no greater than 2. Consequently, if F_A is estimated by the linear method above it is not necessary to subsequently impose a singularity constraint. This is a considerable advantage over the estimation of a general F by the linear 8-point algorithm, where the estimated matrix is not guaranteed to have rank 2, and thus must be subsequently corrected.

Geometric interpretation. As has been seen at several points in this book, computing a two-view relation from point correspondences is equivalent to fitting a surface (variety) to points x, y, x', y' in $\mathbb{R}^4$. In the case of the equation $x'^T F_A x = 0$ the relation $ax'_i + by'_i + cx_i + dy_i + e = 0$ is *linear* in the coordinates, and the variety $\mathcal{V}_{F_A}$ defined by the affine fundamental matrix is a hyperplane.

This results in two simplifications: first, finding the best estimate of F_A may be formulated as a (familiar) plane-fitting problem; second, the Sampson error is identical to the geometric error, whereas in the case of a general (non-affine) fundamental matrix (11.9–p287) it is only a first-order approximation. As discussed in section 4.2.6(p98) this latter property arises generally with affine (linear) relations because the tangent plane of the Sampson approximation is equivalent to the surface.

14.3.2 The Gold Standard algorithm

Given a set of n corresponding image points $\{x_i \leftrightarrow x'_i\}$, we seek the Maximum Likelihood estimate of F_A under the assumption that the noise in the image measurements obeys an isotropic, homogeneous Gaussian distribution. This estimate is obtained by minimizing a cost function on geometric image distances:

$$\min_{\{F_A, \hat{x}_i, \hat{x}'_i\}} \sum_i d(x_i, \hat{x}_i)^2 + d(x'_i, \hat{x}'_i)^2 \quad (14.4)$$

where as usual $x_i \leftrightarrow x'_i$ are the measured correspondences, and $\hat{x}_i$ and $\hat{x}'_i$ are estimated “true” correspondences that satisfy $\hat{x}'_i^T F_A \hat{x}_i = 0$ exactly for the estimated affine fundamental matrix. The distances are illustrated in figure 14.3. The true correspondences are subsidiary variables that must also be estimated.

As discussed above, and in section 4.2.5(p96), minimizing the cost function (14.4)




Fig. 14.3. The MLE of F_A from a set of measured corresponding points $\{x_i \leftrightarrow x'_i\}$ involves estimating the five parameters a, b, c, d, e together with a set of correspondences $\{\hat{x}_i \leftrightarrow \hat{x}'_i\}$ which exactly satisfy $\hat{x}'_i{}^T F_A \hat{x}_i = 0$. There is a linear solution to this problem.




Fig. 14.4. In 2D a line is the analogue of the hyperplane defined by F_A , and the problem of estimating the true correspondence given a measured correspondence, is the problem of determining the closest point $(\hat{x}, \hat{y})$ on the line $ax + by + c$ to a measurement point (x, y) . The normal to the line has direction (a, b) , and the perpendicular distance of the point (x, y) from the line is $d_\perp = (ax + by + c) / \sqrt{a^2 + b^2}$, so that $(\hat{x}, \hat{y})^T = (x, y)^T - d_\perp \hat{n}$, where $\hat{n} = (a, b) / \sqrt{a^2 + b^2}$.

is equivalent to fitting a hyperplane to set of points $\mathbf{X}_i = (x'_i, y'_i, x_i, y_i)^T$ in $\mathbb{R}^4$. The estimated points $\hat{\mathbf{X}}_i = (\hat{x}'_i, \hat{y}'_i, \hat{x}_i, \hat{y}_i)^T$ satisfy the equation $\hat{x}'_i{}^T F_A \hat{x}_i = 0$ which may be written as $(\hat{\mathbf{X}}_i^T, 1)\mathbf{f} = 0$, where $\mathbf{f} = (a, b, c, d, e)^T$. This is the equation of a point in $\mathbb{R}^4$ on the plane $\mathbf{f}$. We seek the plane $\mathbf{f}$ which minimizes the squared distance between the measured and estimated points, and consequently which minimizes the sum of squared perpendicular distances to the points $\mathbf{X}_i = (x'_i, y'_i, x_i, y_i)^T$.

Geometrically the solution is very simple, and an analogue for line fitting in 2D is illustrated in figure 14.4. The perpendicular distance of a point $\mathbf{X}_i = (x_i, y_i, x'_i, y'_i)^T$ from the plane $\mathbf{f}$ is

$$d_\perp(\mathbf{X}_i, \mathbf{f}) = \frac{ax'_i + by'_i + cx_i + dy_i + e}{\sqrt{a^2 + b^2 + c^2 + d^2}}.$$

Then the matrix F_A which minimizes (14.4) is determined by minimizing the cost function

$$\mathcal{C} = \sum_i d_\perp(\mathbf{X}_i, \mathbf{f})^2 = \frac{1}{a^2 + b^2 + c^2 + d^2} \sum_i (ax'_i + by'_i + cx_i + dy_i + e)^2 \quad (14.5)$$

Objective

Given $n \geq 4$ image point correspondences $\{\mathbf{x}_i \leftrightarrow \mathbf{x}'_i\}$, $i = 1, \dots, n$, determine the Maximum Likelihood estimate F_A of the affine fundamental matrix.

Algorithm

A correspondence is represented as $\mathbf{X}_i = (x'_i, y'_i, x_i, y_i)^\top$.

- (i) Compute the centroid $\bar{\mathbf{X}} = \frac{1}{n} \sum_i \mathbf{X}_i$ and centre the vectors $\Delta \mathbf{X}_i = \mathbf{X}_i - \bar{\mathbf{X}}$.
- (ii) Compute the $n \times 4$ matrix A with rows $\Delta \mathbf{X}_i^\top$.
- (iii) Then $\mathbf{N} = (a, b, c, d)^\top$ is the singular vector corresponding to the smallest singular value of A , and $e = -\mathbf{N}^\top \bar{\mathbf{X}}$. The matrix F_A has the form (14.1).

Algorithm 14.1. *The Gold Standard algorithm for estimating F_A from image correspondences.*

over the 5 parameters $\{a, b, c, d, e\}$ of $\mathbf{f}$. Writing $\mathbf{N} = (a, b, c, d)^\top$ for the normal to the hyperplane then

$$\mathcal{C} = \frac{1}{\|\mathbf{N}\|^2} \sum_i (\mathbf{N}^\top \mathbf{X}_i + e)^2.$$

This cost function can be minimized by a very simple linear algorithm, equivalent to the classical problem of orthogonal regression to a plane. There are two steps:

The first step is to minimize $\mathcal{C}$ over the parameter e . We obtain

$$\frac{\partial \mathcal{C}}{\partial e} = \frac{1}{\|\mathbf{N}\|^2} \sum_i 2(\mathbf{N}^\top \mathbf{X}_i + e) = 0$$

and hence

$$e = -\frac{1}{n} \sum_i (\mathbf{N}^\top \mathbf{X}_i) = -\mathbf{N}^\top \bar{\mathbf{X}}$$

so the solution hyperplane passes through the data centroid $\bar{\mathbf{X}}$. Substituting for e in the cost function reduces $\mathcal{C}$ to

$$\mathcal{C} = \frac{1}{\|\mathbf{N}\|^2} \sum_i (\mathbf{N}^\top \Delta \mathbf{X}_i)^2$$

where $\Delta \mathbf{X}_i = \mathbf{X}_i - \bar{\mathbf{X}}$ is the vector $\mathbf{X}_i$ relative to the data centroid $\bar{\mathbf{X}}$.

The second step is to minimize this reduced cost function over $\mathbf{N}$. Writing A for the matrix with rows $\Delta \mathbf{X}_i^\top$, it is evident that

$$\mathcal{C} = \|\mathbf{A}\mathbf{N}\|^2 / \|\mathbf{N}\|^2.$$

Minimizing this expression is equivalent to minimizing $\|\mathbf{A}\mathbf{N}\|$ subject to $\|\mathbf{N}\| = 1$, which is our usual homogeneous minimization problem solved by the SVD. These steps are summarized in algorithm 14.1.

It is worth noting that the Gold Standard algorithm produces an identical estimate of F_A to that obtained by the factorization algorithm 18.1(p437) for an affine reconstruction from n point correspondences.




Fig. 14.5. Computing the affine epipolar line for a minimal configuration of four points. The line is computed by the virtual parallax induced by the plane π . Compare figure 13.7(p335).

Objective

Given four image point correspondences $\{x_i \leftrightarrow x'_i\}$, $i = 1, \dots, 4$, compute the affine fundamental matrix.

Algorithm

The first three 3-space points X_i , $i = 1, \dots, 3$ define a plane π . See figure 14.5.

- (i) Compute the affine transformation matrix H_A , such that $x'_i = H_A x_i$, $i = 1, \dots, 3$.
- (ii) Determine the epipolar line in the second view from $l' = (H_A x_4) \times x'_4$. The epipole $e' = (-l'_2, l'_1, 0)^T$.
- (iii) Then for any point x the epipolar line in the second view is $e' \times (H_A x) = F_A x$. Hence $F_A = [(-l'_2, l'_1, 0)^T] \times H_A$.

Algorithm 14.2. The computation of F_A for a minimal configuration of four point correspondences.

14.3.3 The minimal configuration

We return to the minimal configuration for estimating F_A , namely the corresponding images of four points in 3-space in general position. A geometric computation method of F_A for this configuration is described in algorithm 14.2. This minimal solution is useful in the case of robust estimation algorithms, such as RANSAC, and will be used here to illustrate degenerate configurations. Note that for this minimal configuration an exact solution is obtained for F_A , and the linear algorithm of section 14.3.1, the Gold Standard algorithm 14.1, and the minimal algorithm 14.2 give an identical result.

General position. The configuration of four points shown in figure 14.5 demonstrates the conditions necessary for general position of the 3-space points when computing F_A . Configurations for which F_A cannot be computed are degenerate. These fall into two classes: first, degenerate configurations depending only on the structure, for example if the four points are coplanar (so there is no parallax), or if the first three points are collinear (so that H_A can't be computed); second, those degeneracies which depend

only on the cameras, for example if the two cameras have the same viewing direction (and so have common centres on the plane at infinity).

Note once again the importance of parallax – as the point $\mathbf{X}_4$ approaches the plane defined by the other three points in figure 14.5 the parallax vector, which determines the epipolar line direction, is monotonically reduced in length. Consequently, the accuracy of the line direction is correspondingly reduced. This result for the minimal configuration is true also of the Gold Standard algorithm 14.1: as relief reduces to zero, i.e. the point set approaches coplanarity, the covariance of the estimated $\mathbf{F}_A$ will increase.

14.4 Triangulation

Suppose we have a measured correspondence $(x, y) \leftrightarrow (x', y')$ and the affine fundamental matrix $\mathbf{F}_A$. We wish to determine the Maximum Likelihood estimate of the true correspondence, $(\hat{x}, \hat{y}) \leftrightarrow (\hat{x}', \hat{y}')$, under the usual assumption that image measurement error is Gaussian. The 3D point may then be determined from the ML estimate correspondence.

As we have seen earlier in chapter 12, the MLE involves determining a “true” correspondence which exactly obeys the affine epipolar geometry, i.e. $(\hat{x}', \hat{y}', 1)\mathbf{F}_A(\hat{x}, \hat{y}, 1)^T = 0$, and also minimizes the image distance to the measured points,

$$(x - \hat{x})^2 + (y - \hat{y})^2 + (x' - \hat{x}')^2 + (y' - \hat{y}')^2.$$

Geometrically the solution is very simple, and is illustrated in 2D in figure 14.4. We seek the closest point on the hyperplane defined by $\mathbf{F}_A$ to the measured correspondence $\mathbf{X} = (x', y', x, y)^T$ in $\mathbb{R}^4$. Again, the Sampson correction (4.11–p99) is exact in this case. Algebraically, the normal to the hyperplane has direction $\mathbf{N} = (a, b, c, d)^T$ and the perpendicular distance of a point $\mathbf{X}$ to the hyperplane is given by $d_\perp = (\mathbf{N}^T \mathbf{X} + e) / \|\mathbf{N}\|$, so that

$$\hat{\mathbf{X}} = \mathbf{X} - d_\perp \frac{\mathbf{N}}{\|\mathbf{N}\|}$$

or in its full detail

$$\begin{pmatrix} \hat{x}' \\ \hat{y}' \\ \hat{x} \\ \hat{y} \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ x \\ y \end{pmatrix} - \frac{(ax' + by' + cx + dy + e)}{(a^2 + b^2 + c^2 + d^2)} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix}.$$

14.5 Affine reconstruction

Suppose we have $n \geq 4$ image point correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$, $i = 0, \dots, n-1$, which for the moment will be assumed noise-free, then we may compute a reconstruction of the 3D points and cameras. In the case of projective cameras (with $n \geq 7$ points) the reconstruction was projective. In the affine case, not surprisingly, the reconstruction is *affine*. We will now give a simple constructive derivation of this result.

An affine coordinate frame in 3-space may be specified by four finite non-coplanar basis points $\mathbf{X}_i$, $i = 0, \dots, 3$. As illustrated in figure 14.6 one point $\mathbf{X}_0$ is chosen as the




Fig. 14.6. **Affine coordinates.** (a) four non-coplanar points in 3-space (X_1 , X_2 , X_3 and origin X_0) define a set of axes in terms of which other points X can be assigned affine coordinates (X, Y, Z) . (b) Each affine coordinate is defined by a ratio of lengths in parallel directions (which is an affine invariant). For example, X may be computed by the following two operations: first, X is projected parallel to $\tilde{E}_2$ onto the plane spanned by $\tilde{E}_1$ and $\tilde{E}_3$. Second, this projected point is projected parallel to $\tilde{E}_3$ onto the $\tilde{E}_1$ axis. The value of the coordinate x is the ratio of the length from the origin of this final projected point to the length of $\tilde{E}_1$.




Fig. 14.7. **Reconstruction from two images.** The affine coordinates of the 3D point X with image x , x' in two views may be computed linearly from the projection of the basis points x_i and basis vectors $\tilde{e}_i$ of figure 14.6.

origin, and the three other points then define basis vectors $\tilde{E}_i = \tilde{X}_i - \tilde{X}_0, i = 1, \dots, 3$, where $\tilde{X}_i$ is the inhomogeneous 3-vector corresponding to X_i . The position of a point X may then be specified by simple vector addition as

$$\tilde{X} = \tilde{X}_0 + x\tilde{E}_1 + y\tilde{E}_2 + z\tilde{E}_3$$

and (x, y, z) are the affine coordinates of $\tilde{X}$ with respect to this basis. This means that the basis points X_i have the canonical coordinates $(x, y, z)^T$:

$$\tilde{X}_0 = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \quad \tilde{X}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \quad \tilde{X}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \quad \tilde{X}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \quad (14.6)$$

Given the affine projection of the four basis points in two views, the 3D affine coordinates of any other point can be directly recovered from its image, as will now be demonstrated (see figure 14.7).

Projection with an affine camera may be represented as (6.26-p172)

$$\tilde{\mathbf{x}} = \mathbf{M}_{2 \times 3} \tilde{\mathbf{X}} + \tilde{\mathbf{t}}$$

where $\tilde{\mathbf{x}} = (x, y)^T$ is the inhomogeneous 2-vector corresponding to $\mathbf{x}$. Differences of vectors eliminate $\tilde{\mathbf{t}}$. For example the basis vectors project as $\tilde{\mathbf{e}}_i = \mathbf{M}_{2 \times 3} \tilde{\mathbf{E}}_i$, $i = 1, \dots, 3$. Consequently for any point $\mathbf{X}$, its image in the first view is

$$\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_0 = x\tilde{\mathbf{e}}_1 + y\tilde{\mathbf{e}}_2 + z\tilde{\mathbf{e}}_3 \quad (14.7)$$

and similarly the image ($\tilde{\mathbf{x}}' = \mathbf{M}'_{2 \times 3} \tilde{\mathbf{X}} + \tilde{\mathbf{t}}'$) in the second view is

$$\tilde{\mathbf{x}}' - \tilde{\mathbf{x}}'_0 = x\tilde{\mathbf{e}}'_1 + y\tilde{\mathbf{e}}'_2 + z\tilde{\mathbf{e}}'_3. \quad (14.8)$$

Each equation (14.7) and (14.8) imposes two linear constraints on the unknown affine coordinates x, y, z of the space point $\mathbf{X}$. All the other terms in the equations are known from image measurements (for example the image basis vectors $\tilde{\mathbf{e}}_i, \tilde{\mathbf{e}}'_i$ are computed from the projection of the four basis points $\tilde{\mathbf{X}}_i$, $i = 0, \dots, 3$). Thus, there are four linear simultaneous equations in the three unknowns x, y, z , and the solution is straightforward. This demonstrates that the affine coordinates of a point $\mathbf{X}$ may be computed from its image in two views.

The cameras for the two views, $\mathbf{P}_A, \mathbf{P}'_A$, may be computed from the correspondences between the 3-space points $\tilde{\mathbf{X}}_i$, with coordinates given in (14.6), and their measured images. For example, $\mathbf{P}_A$ is computed from the correspondence $\tilde{\mathbf{x}}_i \leftrightarrow \tilde{\mathbf{X}}_i$, $i = 0, \dots, 3$.

The above development is not optimal, because the basis points are treated as exact, and all measurement error associated with the fifth point $\mathbf{X}$. An optimal reconstruction algorithm, where reprojection error is minimized over all points, is very straightforward in the affine case. However, its description is postponed until section 18.2(p436) because the factorization algorithm described there is applicable to any number of views.

Example 14.2. Affine reconstruction

A 3D reconstruction is computed for the hole punch images of figure 14.2 by choosing four points as the affine basis, and then computing the affine coordinates of each of the remaining points in turn by the linear method above. Two views of the resulting reconstruction are shown in figure 14.8. Note, however, that this five-point method is *not* recommended. Instead the optimal affine reconstruction algorithm 18.1(p437) should be used. $\triangle$

14.6 Necker reversal and the bas-relief ambiguity

We have seen in the previous section that in the absence of any calibration information, an affine reconstruction is obtained from point correspondences alone. In this section we show that even if the camera calibration is known there remains a family of reconstruction ambiguities which cannot be resolved in the two-view case.

This situation differs from that of perspective projection where once the internal calibration is determined the camera motion is determined up to a finite number of




Fig. 14.8. **Affine reconstruction.** (a)(b) Wireframe outline of the hole punch from the two images of figure 14.2. The circles show the points selected as the affine basis. The lines are for visualization only. (c) Two views of the 3D affine structure computed from the vertices of the wireframe.

ambiguities (from the essential matrix, see section 9.6(p257)). For parallel projection there are two important additional ambiguities: a finite reflection ambiguity (Necker reversal); and a one-parameter family rotation ambiguity (the bas-relief ambiguity).

Necker reversal ambiguity. This arises because an object rotating by ρ and its mirror image rotating by $-\rho$ generate the same image under parallel projection, see figure 14.9(a). Thus, structure is only recovered up to a reflection about the frontal plane. This ambiguity is absent in the perspective case because the points have different depths in the two interpretations and so do not project to coincident image points.

The bas-relief ambiguity. This is illustrated in figure 14.9(b). Imagine a set of parallel rays from one camera, and consider adjusting a set of parallel rays from a second camera until each ray intersects its corresponding ray. The rays lie in a family of parallel epipolar planes, and there remains the freedom to rotate one camera about the normal to these planes whilst maintaining incidence of the rays. This *bas-relief* (or *depth-turn*) ambiguity is a one-parameter family of solutions for the rotation angle and depth. The parameters of depth, ΔZ , and rotation, $\sin \rho$, are confounded and cannot be determined individually – only their product can be computed. Consequently, a shallow object experiencing a large turn (i.e. small ΔZ and large ρ) generates the same image as a deep object experiencing a small turn (i.e. large ΔZ and small ρ). The name arises from bas-relief sculptures. Fixing the depth or the angle ρ determines the structure and the motion uniquely. Extra points cannot resolve this ambiguity, but an additional view (i.e. three views) will in general resolve it.




Fig. 14.9. **Motion ambiguities under parallel projection.** (a) *Necker reversal*: a rotating object generates the same image as its mirror object rotating in the opposite sense. Under perspective projection the images are different. (b) The cameras can rotate (by ρ) and still preserve ray intersections. This cannot happen for perspective cameras. (c) The *bas-relief ambiguity*: consider a rod of length l , which rotates through an angle ρ . That is $x' - x = l \sin \rho$. This bas-relief (or depth-turn) ambiguity is so-named because a shallow object experiencing a large turn (i.e. small l and big ρ) generates the same image as a deep object experiencing a small turn (i.e. large l and small ρ).

This ambiguity casts light on the stability of reconstruction from two perspective cameras: as imaging conditions approach affine the rotation angle will be poorly estimated, but the product of the rotation angle and depth will be stable.

14.7 Computing the motion

In this section expressions for computing the camera motion from F_A will be given for the case of two weak perspective cameras (section 6.3.4(p170)). These cameras may be chosen as

$$P = \begin{bmatrix} \alpha_x & & & \\ & \alpha_y & & \\ & & 1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad P' = \begin{bmatrix} \alpha'_x & & & \\ & \alpha'_y & & \\ & & 1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} \mathbf{r}_1^T & t_1 \\ \mathbf{r}_2^T & t_2 \\ \mathbf{0}^T & 1 \end{bmatrix}$$

where $\mathbf{r}_1$ and $\mathbf{r}_2$ are the first and second rows of the rotation matrix R between the views. We will assume that the aspect ratio α_y/α_x is known in both cameras, but that the relative scaling $s = \alpha'_x/\alpha_x$ is unknown. $s > 1$ for a “looming” object and $s < 1$ for one that is “receding”. As has been seen, the complete rotation R cannot be computed from two weak perspective views, resulting in the *bas-relief* ambiguity. Nevertheless




Fig. 14.10. **The rotation representation.** (a) rotation by θ about the Z-axis; (b) subsequent rotation by ρ about a fronto-parallel axis Φ angled at ϕ to the X-axis. The Φ -axis has components $(\cos \phi, \sin \phi, 0)^T$.




Fig. 14.11. The camera rotates about the axis Φ which is parallel to the image plane. The intersection of the epipolar plane π with the image planes gives epipolar lines l and l' , and the projections of Φ in the images are orthogonal to these epipolar lines: (a) no cyclotorsion occurs ($\theta = 0^\circ$); (b) the camera counter-rotates by θ in I_1 , so the orientation of the epipolar lines changes by θ .

the remaining motion parameters can be computed from F_A , and their computation is straightforward.

To represent the motion we will use a rotation representation introduced by Koenderink and van Doorn [Koenderink-91]. As will be seen, this has the advantage that it isolates the parameter ρ of the bas-relief ambiguity, which cannot be computed from the affine epipolar geometry. In this representation the rotation R between the views is decomposed into two rotations (see figure 14.10),

$$R = R_\rho R_\theta. \quad (14.9)$$

First, there is a cyclo-rotation R_θ in the image plane through an angle θ (i.e. about the line of sight). This is followed by a rotation R_ρ through an angle ρ about an axis Φ with direction parallel to the image plane, and angled at ϕ to the positive X-axis, i.e. a pure rotation *out of* the image plane.




Fig. 14.12. The scene can be sliced into parallel epipolar planes. The magnitude of ρ has no effect on the epipolar geometry (provided $\rho \neq 0$), so it is indeterminate from two views.




Fig. 14.13. The effect of scale and rotation angles on the epipolar lines for an object moving relative to a stationary camera. This also illustrates the assumed sequence of events accounting for the transition from I_1 to I_2 : (a) I_1 ; (b) cyclotorsion (θ); (c) rotation out of the plane (ϕ and ρ); (d) scaling, giving I_2 .

Solving for s , ϕ and θ . It is now shown that the scale factor (s), the projection of the axis of rotation (ϕ) and the cyclo-rotation angle (θ) may be computed directly from the affine epipolar geometry. The solution is preceded by a geometric explanation of how the epipolar lines relate to the unknown motion parameters.

Consider a camera rotating about an axis Φ lying parallel to the image plane (figure 14.11(a)). The epipolar plane π is perpendicular to both this axis and the two images, and intersects the images in the epipolar lines I and I' . Consequently:

- The projection of the axis of rotation Φ is perpendicular to the epipolar lines.

This relation still holds if there is additionally a cyclotorsion θ in the image plane (figure 14.11(b)); the axis Φ and intersection I' remain fixed in space, and are simply observed at a new angle in the image, maintaining the orthogonality between the epipolar lines and the projected axis. The orientations of the epipolar lines in the two images therefore differ by θ . Importantly, changing the magnitude of the turn angle ρ doesn't alter the epipolar geometry in any way (figure 14.12). This angle is therefore indeterminate from two views, a consequence of the bas-relief ambiguity.



Fig. 14.14. **Computing motion from affine epipolar geometry.** (a)(b) Two views of a buggy rotating on a turntable. The computed rotation axis is superimposed on the image, drawn to pass through the image centre. The ground truth axis is, of course, perpendicular to the turntable in the world.

Figure 14.13 illustrates the effect of scale. Consider a 3D object to be sliced into parallel epipolar planes, with each plane constraining how a particular slice of the object moves. Altering the effective size of the object (e.g. by moving closer to it) simply changes the relative spacing between successive epipolar planes.

In summary, cyclotorsion simply rotates the epipolar lines, rotation out of the plane causes foreshortening along the epipolar lines (orthogonal to Φ), and a scale change uniformly alters the epipolar line spacing (figure 14.13).

It can be shown (and is left as exercise) that s , θ and ϕ can be computed directly from the affine epipolar geometry as

$$\tan \phi = \frac{b}{a}, \quad \tan(\phi - \theta) = \frac{d}{c} \quad \text{and} \quad s^2 = \frac{c^2 + d^2}{a^2 + b^2}, \quad (14.10)$$

with $s > 0$ (by definition). Note that ϕ is the angle of projection in I_2 of the axis of rotation out of the plane, while $(\phi - \theta)$ is its angle of projection in I_1 .

Example 14.3. Motion computed from the affine fundamental matrix

figure 14.14 shows two images of a buggy rotating on a turntable. The image is 256×256 pixels with an aspect ratio of 0.65. The affine fundamental matrix is computed using algorithm 14.1, and the motion parameters computed from F_A using (14.10) above. The computed rotation axis is superimposed on the image. $\triangle$

14.8 Closure

14.8.1 The literature

Koenderink and van Doorn [Koenderink-91] set the scene for affine reconstruction from two affine cameras. This paper should be read by all. The affine fundamental matrix was first defined in [Zisserman-92]. The computation of the motion parameters from F_A is described in Shapiro *et al.* [Shapiro-95], and in particular the cases where a third view does not resolve the bas-relief ambiguity. A helpful eigenvector analysis of the ambiguity is given in [Szeliski-96]. The three view affine motion case is treated quite elegantly in [Shimshoni-99].

14.8.2 Notes and exercises

- (i) A scene plane induces an *affine* transformation between two affine cameras. There is a three-parameter family of such affinities defined by the three-parameter family of planes in $\mathbb{R}^3$. Given F_A , this family of affinities may be written as (result 13.3(p328)) $H_A = [e']_x F_A + e'v^T$, where $F_A^T e' = 0$, and the 3-vector v parametrizes the family of planes. Conversely, show that given H_A , the homography induced by a scene plane, then F_A is determined up to a one-parameter ambiguity.
- (ii) Consider a perspective camera, i.e. the matrix does not have the affine form. Show that if the camera motion consists of a translation parallel to the image plane, and a rotation about the principal axis, then F has the affine form. This shows that a fundamental matrix with affine form does *not* imply that the imaging conditions are affine. Are there other camera motions which generate a fundamental matrix with the affine form?
- (iii) Two affine cameras, P_A, P'_A , uniquely define an affine fundamental matrix F_A by (9.1-p244). Show that if the cameras are transformed on the right by a common affine transformation, i.e. $P_A \mapsto P_A H_A, P'_A \mapsto P'_A H_A$, the transformed cameras define the original F_A . This shows that the affine fundamental matrix is invariant to an affine transformation of the world coordinates.
- (iv) Suppose one of the cameras is affine and the other is a perspective camera. Show that in general in this case the epipoles in both views are finite.
- (v) The 4×4 permutation homography

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

maps the canonical matrix of a finite projective camera $P = [I \mid 0]$ into the canonical matrix of parallel projection P_A :

$$P_A = [I \mid 0] H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Show, by applying this transformation to a pair of finite projective camera matrices, that the results of this chapter (such as the properties listed in table 14.1-(p348)) can be generated directly from their non-affine counterparts of the previous chapters. In particular derive an expression for a pair of affine cameras P_A, P'_A consistent with F_A .

Part III

Three-View Geometry



Lord Shiva, c. 1920-40 (print).

Shiva is depicted as having three eyes. The third eye in the centre of the forehead symbolizes spiritual knowledge and power.

Image courtesy of <http://www.healthyplanetonline.com>

Outline

This part contains two chapters on the geometry of three-views. The scene is imaged with three cameras perhaps simultaneously in a trinocular rig, or sequentially from a moving camera.

Chapter 15 introduces a new multiple view object – the trifocal tensor. This has analogous properties to the fundamental matrix of two-view geometry: it is independent of scene structure depending only on the (projective) relations between the cameras. The camera matrices may be retrieved from the trifocal tensor up to a common projective transformation of 3-space, and the fundamental matrices for view-pairs may be retrieved uniquely.

The new geometry compared with the two-view case is the ability to *transfer* from two views to a third: given a point correspondence over two views the position of the point in the third view is determined; and similarly, given a *line* correspondence over two views the position of the line in the third view is determined. This transfer property is of great benefit when establishing correspondences over multiple views.

If the essence of the epipolar constraint over two views is that rays back-projected from corresponding points are coplanar, then the essence of the trifocal constraint over three views is the geometry of a point–line–line correspondence arising from the image of a point on a line in 3-space: corresponding image lines in two views back-project to planes which intersect in a line in 3-space, and the ray back-projected from a corresponding image point in a third view must intersect this line.

Chapter 16 describes the computation of the trifocal tensor from point and *line* correspondences over three-views. Given the tensor, and thus the retrieved camera matrices, a projective reconstruction may be computed from correspondences over multiple views. The reconstruction may be upgraded to similarity or metric as additional information is provided in the same manner as in the two view case.

It is in reconstruction that there is another gain over two-view geometry. Given the cameras, in the two-view case each point correspondence provided four measurements on the three degrees of freedom (the position) of the point in 3-space. In three views there are six measurements on, again, three degrees of freedom. However, it is for lines that there is the more significant gain. In two-views the number of measurements equals the number of degrees of freedom of the line in 3-space, namely four. Consequently, there is no possibility of removing the effects of measurement errors. However, in three views there are six measurements on four degrees of freedom, so a scene line is over-determined and can be estimated by a suitable minimization over measurement errors.

The Trifocal Tensor

The trifocal tensor plays an analogous role in three views to that played by the fundamental matrix in two. It encapsulates all the (projective) geometric relations between three views that are independent of scene structure.

We begin this chapter with a simple introduction to the main geometric and algebraic properties of the trifocal tensor. A formal development of the trifocal tensor and its properties involves the use of tensor notation. To start, however, it is convenient to use standard vector and matrix notation, thus obtaining some geometric insight into the trifocal tensor without the additional burden of dealing with a (possibly) unfamiliar notation. The use of tensor notation will therefore be deferred until section 15.2.

The three principal geometric properties of the tensor are introduced in section 15.1. These are the homography between two of the views induced by a plane back-projected from a line in the other view; the relations between image correspondences for points and lines which arise from incidence relations in 3-space; and the retrieval of the fundamental and camera matrices from the tensor.

The tensor may be used to transfer points from a correspondence in two views to the corresponding point in a third view. The tensor also applies to lines, and the image of a line in one view may be computed from its corresponding images in two other views. Transfer is described in section 15.3.

The tensor only depends on the motion between views and the internal parameters of the cameras and is defined uniquely by the camera matrices of the views. However, it can be computed from image correspondences alone without requiring knowledge of the motion or calibration. This computation is described in chapter 16.

15.1 The geometric basis for the trifocal tensor

There are several ways that the trifocal tensor may be approached, but in this section the starting point is taken to be the incidence relationship of three corresponding lines.

Incidence relations for lines. Suppose a line in 3-space is imaged in three views, as in figure 15.1, what constraints are there on the corresponding image lines? The planes back-projected from the lines in each view must all meet in a single line in space, the 3D line that projects to the matched lines in the three images. Since in general three arbitrary planes in space do not meet in a single line, this geometric *incidence* condition




Fig. 15.1. A line L in 3-space is imaged as the corresponding triplet $l \leftrightarrow l' \leftrightarrow l''$ in three views indicated by their centres, C, C', C'' , and image planes. Conversely, corresponding lines back-projected from the first, second and third images all intersect in a single 3D line in space.

provides a genuine constraint on sets of corresponding lines. We will now translate this geometric constraint into an algebraic constraint on the three lines.

We denote a set of corresponding lines as $l_i \leftrightarrow l'_i \leftrightarrow l''_i$. Let the camera matrices for the three views be $P = [I \mid 0]$, as usual, and $P' = [A \mid a_1]$, $P'' = [B \mid b_1]$, where A and B are 3×3 matrices, and the vectors a_i and b_i are the i -th columns of the respective camera matrices for $i = 1, \dots, 4$.

- a_1 and b_1 are the epipoles in views two and three respectively, arising from the first camera. These epipoles will be denoted by e' and e'' throughout this chapter, with $e' = P'C$, $e'' = P''C$, where C is the first camera centre. (For the most part we will not be concerned with the epipoles between the second and third views).
- A and B are the infinite homographies from the first to the second and third cameras respectively.

As has been seen in chapter 9, any set of three cameras is equivalent to a set with $P = [I \mid 0]$ under projective transformations of space. In this chapter we will be concerned with properties (such as image coordinates and 3D incidence relations) that are invariant under 3D projective transforms, so we are free to choose the cameras in this form.

Now, each image line back-projects to a plane, as shown in figure 15.1. From result 8.2(p197) these three planes are

$$\pi = P^T l = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \pi' = P'^T l' = \begin{pmatrix} A^T l' \\ a_1^T l' \end{pmatrix} \quad \pi'' = P''^T l'' = \begin{pmatrix} B^T l'' \\ b_1^T l'' \end{pmatrix}.$$

Since the three image lines are derived from a single line in space, it follows that these three planes are not independent but must meet in this common line in 3-space. This intersection constraint can be expressed algebraically by the requirement that the 4×3 matrix $M = [\pi, \pi', \pi'']$ has rank 2. This may be seen as follows. Points on the line of intersection may be represented as $X = \alpha X_1 + \beta X_2$, with X_1 and X_2 linearly independent. Such points lie on all three planes and so $\pi^T X = \pi'^T X = \pi''^T X = 0$. It

follows that $M^T X = 0$. Consequently M has a 2-dimensional null-space since $M^T X_1 = 0$ and $M^T X_2 = 0$.

This intersection constraint induces a relation amongst the image lines l, l', l'' . Since the rank of M is 2, there is a linear dependence between its columns m_i . Denoting

$$M = [m_1, m_2, m_3] = \begin{bmatrix} 1 & A^T l' & B^T l'' \\ 0 & a_1^T l' & b_1^T l'' \end{bmatrix}$$

the linear relation may be written $m_1 = \alpha m_2 + \beta m_3$. Then noting that the bottom left hand element of M is zero, it follows that $\alpha = k(b_4^T l'')$ and $\beta = -k(a_4^T l')$ for some scalar k . Applying this to the top 3-vectors of each column shows that (up to a homogeneous scale factor)

$$l = (b_4^T l'') A^T l' - (a_4^T l') B^T l'' = (l''^T b_4) A^T l' - (l'^T a_4) B^T l''.$$

The i -th coordinate l_i of l may therefore be written as

$$l_i = l''^T (b_4 a_i^T) l' - l'^T (a_4 b_i^T) l'' = l'^T (a_i b_4^T) l'' - l'^T (a_4 b_i^T) l''$$

and introducing the notation

$$T_i = a_i b_4^T - a_4 b_i^T \quad (15.1)$$

the incidence relation can be written

$$l_i = l'^T T_i l'' \quad (15.2)$$

Definition 15.1. The set of three matrices $\{T_1, T_2, T_3\}$ constitute the *trifocal tensor* in matrix notation.

We introduce a further notation¹. Denoting the ensemble of the three matrices T_i by $[T_1, T_2, T_3]$, or more briefly $[T_i]$, this last relation may be written as

$$l^T = l'^T [T_1, T_2, T_3] l'' \quad (15.3)$$

where $l'^T [T_1, T_2, T_3] l''$ is understood to represent the vector $(l'^T T_1 l'', l'^T T_2 l'', l'^T T_3 l'')$.

Of course there is no intrinsic difference between the three views, and so by analogy with (15.3) there will exist similar relations $l'^T = l^T [T'_i] l''$ and $l''^T = l^T [T''_i] l'$. The three tensors $[T_i]$, $[T'_i]$ and $[T''_i]$ exist, but are distinct. In fact, although all three tensors may be computed from any one of them, there is no very simple relationship between them. Thus, in fact there are three trifocal tensors existing for a given triple of views. Usually one will be content to consider only one of them. However, a method of computing the other trifocal tensors $[T'_i]$ and $[T''_i]$ given $[T_i]$ is outlined in exercise (viii) on page 389.

Note that (15.3) is a relationship between image coordinates only, not involving 3D coordinates. Hence (as remarked previously), although it was derived under the assumption of a canonical camera set (that is $P = [I \mid 0]$), the value of the matrix elements $[T_i]$ is independent of the form of the cameras. The particular simple formula (15.1) for the trifocal tensor given the camera matrices holds only in the case where

¹ This notation is somewhat cumbersome, and its meaning is not quite self-evident. It is for this reason that tensor notation is introduced in section 15.2.




Fig. 15.2. **Point transfer.** A line l' in the second view back-projects to a plane π' in 3-space. A point x in the first image defines a ray in 3-space which intersects π' in the point X . This point X is then imaged as the point x'' in the third view. Thus, any line l' induces a homography between the first and third views, defined by its back-projected plane π' .

$P = [I \mid 0]$, but a general formula (17.12–p415) for the trifocal tensor corresponding to any three cameras will be derived later.

Degrees of freedom. The trifocal tensor consists of three 3×3 matrices, and thus has 27 elements. There are therefore 26 independent ratios apart from the (common) overall scaling of the matrices. However, the tensor has only 18 independent degrees of freedom. In other words once 18 parameters are specified, all 27 elements of the tensor are determined up to a common scale. The number of degrees of freedom may be computed as follows. Each of 3 camera matrices has 11 degrees of freedom, which makes 33 in total. However, 15 degrees of freedom must be subtracted to account for the projective world frame, thus leaving 18 degrees of freedom. The tensor therefore satisfies $26 - 18 = 8$ independent algebraic constraints. We return to this point in chapter 16.

15.1.1 Homographies induced by a plane

A fundamental geometric property encoded in the trifocal tensor is the homography between the first view and the third induced by a line in the second image. This is illustrated in figure 15.2 and figure 15.3. A line in the second view defines (by back-projection) a plane in 3-space, and this plane induces a homography between the first and third views.

We now derive the algebraic representation of this geometry in terms of the trifocal tensor. The homography map between the first and third images, defined by the plane π' in figure 15.2 and figure 15.3, may be written as $x'' = Hx$ and (2.6–p36) $l = H^T l'$ respectively. Notice that the three lines l, l' and l'' in figure 15.3 are a corresponding line triple, the projections of the 3D line L . Therefore, they satisfy the line incidence relationship $l_i = l'^T T_i l''$ of (15.2). Comparison of this formula and $l = H^T l''$ shows that

$$H = [h_1, h_2, h_3] \text{ with } h_i = T_i^T l'.$$




Fig. 15.3. **Line transfer.** The action on lines of the homography defined by figure 15.2 may similarly be visualized geometrically. A line, l , in the first image defines a plane in 3-space, which intersects π' in the line L . This line L is then imaged as the line l'' in the third view.

Thus, H defined by the above formula represents the (point) homography H_{13} between views one and three specified by the line l' in view two.

The second and third views play similar roles, and the homography between the first and second views defined by a line in the third can be derived in a similar manner. These ideas are formalized in the following result.

Result 15.2. The homography from the first to the third image induced by a line l' in the second image (see figure 15.2) is given by $x'' = H_{13}(l') x$, where

$$H_{13}(l') = [T_1^T, T_2^T, T_3^T] l'.$$

Similarly, a line l'' in the third image defines a homography $x' = H_{12}(l'') x$ from the first to the second views, given by

$$H_{12}(l'') = [T_1, T_2, T_3] l''.$$

Once this mapping is understood the algebraic properties of the tensor are straightforward and can easily be generated. In the following section we deduce a number of incidence relations between points and lines based on (15.3) and result 15.2.

15.1.2 Point and line incidence relations

It is easy to deduce various linear relationships between lines and points in three images involving the trifocal tensor. We have seen one such relationship already, namely (15.3). This relation holds only up to scale since it involves homogeneous quantities. We may eliminate the scale factor by taking the vector cross product of both sides, which must be zero. This leads to the formula

$$(l'^T [T_1, T_2, T_3] l'') [l]_{\times} = 0^T, \quad (15.4)$$

where we have used the matrix $[l]_{\times}$ to denote the cross product (see (A4.5–p581)), or more briefly $(l'^T [T_i] l'') [l]_{\times} = 0^T$. Note the symmetry between l' and l'' – swapping the

roles of these two lines is accounted for by transposing each T_i , resulting in a relation $(I'^T [T_i^T] I') [l]_{\times} = 0^T$.

Consider again figure 15.3. Now, a point x on the line l must satisfy $x^T l = \sum_i x^i l_i = 0$ (using upper indices for the point coordinates, foreshadowing the use of tensor notation). Since $l_i = I'^T T_i I''$, this may be written as

$$I'^T \left(\sum_i x^i T_i \right) I'' = 0 \quad (15.5)$$

(note that $(\sum_i x^i T_i)$ is simply a 3×3 matrix). This is an incidence relation in the first image: the relationship will hold for a point–line–line correspondence – that is whenever some 3D line L maps to I' and I'' in the second and third images, and to a line passing through x in the first image. An important equivalent definition of a point–line–line correspondence for which (15.5) holds results from *an incidence relation in 3-space* – there exists a 3D point X mapping to x in the first image, and to points on the lines I' and I'' in the second and third images as shown in figure 15.4(a).

From result 15.2 we may obtain relations involving points x' and x'' in the second and third images. Consider a point–line–point correspondence as in figure 15.4(b) so that

$$x'' = H_{13}(I') x = [T_1^T I', T_2^T I', T_3^T I'] x = \left(\sum_i x^i T_i^T \right) I'$$

which is valid for any line I' passing through x' in the second image. The homogeneous scale factor may be eliminated by (post-)multiplying the transpose of both sides by $[x'']_{\times}$ to give

$$x''^T [x'']_{\times} = I'^T \left(\sum_i x^i T_i \right) [x'']_{\times} = 0^T, \quad (15.6)$$

A similar analysis may be undertaken with the roles of the second and third images swapped.

Finally, for a 3-point correspondence as shown in figure 15.4(c), there is a relation

$$[x']_{\times} \left(\sum_i x^i T_i \right) [x'']_{\times} = 0_{3 \times 3}. \quad (15.7)$$

Proof. The line I' in (15.6) passes through x' and so may be written as $I' = x' \times y' = [x']_{\times} y'$ for some point y' on I' . Consequently, from (15.6) $I'^T (\sum_i x^i T_i) [x'']_{\times} = y'^T [x']_{\times} (\sum_i x^i T_i) [x'']_{\times} = 0^T$. However, the relation (15.6) is true for *all* lines I' through x' and so is independent of y' . The relation (15.7) then follows. $\square$

The various relationships between lines and points in three views are summarized in table 15.1, and their properties are investigated further in section 15.2.1, once tensor notation has been introduced. Note that there are no relations listed for point–line–line correspondence in which the point is in the second or third view. Such simple relations do not exist in terms of the trifocal tensor in which the first view is the special view.

It is also worth noting that satisfying an image incidence relation does not guarantee incidence in 3-space, as illustrated in figure 15.5.




Fig. 15.4. **Incidence relations.** (a) Consider a 3-view point correspondence $x \leftrightarrow x' \leftrightarrow x''$. If l' and l'' are any two lines through x' and x'' respectively, then $x \leftrightarrow l' \leftrightarrow l''$ forms a point-line-line correspondence, corresponding to a 3D line L . Consequently, (15.5) holds for any choice of lines l' through x' and l'' through x'' . (b) The space point X is incident with the space line L . This defines an incidence relation $x \leftrightarrow l' \leftrightarrow l''$ between their images. (c) The correspondence $x \leftrightarrow x' \leftrightarrow x''$ arising from the image of a space point X .

We now begin to extract the two-view geometry, the epipoles and fundamental matrix, from the trifocal tensor.

15.1.3 Epipolar lines

A special case of a point-line-line correspondence occurs when the plane π' back-projected from l' is an epipolar plane with respect to the first two cameras, and hence

(i) Line-line-line correspondence	$I'^T [T_1, T_2, T_3] I'' = I^T$ or $(I'^T [T_1, T_2, T_3] I'') [I]_{\infty} = 0^T$
(ii) Point-line-line correspondence	$I'^T (\sum_i x^i T_i) I'' = 0$ for a correspondence $x \leftrightarrow I' \leftrightarrow I''$
(iii) Point-line-point correspondence	$I'^T (\sum_i x^i T_i) [x'']_{\infty} = 0^T$ for a correspondence $x \leftrightarrow I' \leftrightarrow x''$
(iv) Point-point-line correspondence	$[x']_{\infty} (\sum_i x^i T_i) I'' = 0$ for a correspondence $x \leftrightarrow x' \leftrightarrow I''$
(v) Point-point-point correspondence	$[x']_{\infty} (\sum_i x^i T_i) [x'']_{\infty} = 0_{3 \times 3}$

Table 15.1. Summary of trifocal tensor incidence relations using matrix notation.




Fig. 15.5. **Non-incident configuration.** The imaged points and lines of this configuration satisfy the point-line-point incidence relation of table 15.1. However, the space point X and line L are not incident. Compare with figure 15.4.

passes through the camera centre C of the first camera. Suppose X is a point on the plane π' ; then the ray defined by X and C lies in this plane, and I' is the epipolar line corresponding to the point x , the image of X . This is shown in figure 15.6.

The plane π'' back-projected from a line I'' in the third image will intersect the plane π' in a line L . Further, since the ray corresponding to x lies entirely in the plane π' it must intersect the line L . This gives a 3-way intersection between the ray and planes back-projected from point x and lines I' and I'' , and so they constitute a point-line-line correspondence, satisfying $I'^T (\sum_i x^i T_i) I'' = 0$. The important point now is that this is true for *any* line I'' , and it follows that $I'^T (\sum_i x^i T_i) = 0^T$. The same argument holds with the roles of I' and I'' reversed. To summarize:




Fig. 15.6. If the plane π' defined by l' is an epipolar plane for the first two views, then any line l'' in the third view gives a point-line-line incidence.

Result 15.3. If x is a point and l' and l'' are the corresponding epipolar lines in the second and third images, then

$$l'^T \left(\sum_i x^i T_i \right) = 0^T \quad \text{and} \quad \left(\sum_i x^i T_i \right) l'' = 0.$$

Consequently, the epipolar lines l' and l'' corresponding to x may be computed as the left and right null-vectors of the matrix $\sum_i x^i T_i$.

As the point x varies, the corresponding epipolar lines vary, but all epipolar lines in one image pass through the epipole. Thus, one may compute this epipole by computing the intersection of the epipolar lines for varying values of x . Three convenient choices of x are the points represented by homogeneous coordinates $(1, 0, 0)^T$, $(0, 1, 0)^T$ and $(0, 0, 1)^T$, with $\sum_i x^i T_i$ equal to T_1 , T_2 and T_3 respectively for these three choices of x . From this we deduce the following important result:

Result 15.4. The epipole e' in the second image is the common intersection of the epipolar lines represented by the left null-vectors of the matrices T_i , $i = 1, \dots, 3$. Similarly the epipole e'' is the common intersection of lines represented by the right null-vectors of the T_i .

Note that the epipoles involved here are the epipoles in the second and third images corresponding to the first image centre C .

The usefulness of this result may not be apparent at present. However, it will be seen below that it is an important step in computing the camera matrices from the trifocal tensor, and in chapter 16 in the accurate computation of the trifocal tensor.

Algebraic properties of the T_i matrices. This section has established a number of algebraic properties of the T_i matrices. We summarize these here:

- Each matrix T_i has rank 2. This is evident from (15.1) since $T_i = a_i e''^T - e' b_i^T$ is the sum of two outer products.
- The right null-vector of T_i is $l''_i = e'' \times b_i$, and is the epipolar line in the third view for the point $x = (1, 0, 0)^T$, $(0, 1, 0)^T$ or $(0, 0, 1)^T$, as $i = 1, 2$ or 3 respectively.

- The epipole $\mathbf{e}''$ is the common intersection of the epipolar lines $\mathbf{l}_i''$ for $i = 1, 2, 3$.
- The left null-vector of $\mathbf{T}_i$ is $\mathbf{l}_i' = \mathbf{e}' \times \mathbf{a}_i$, and is the epipolar line in the second view for the point $\mathbf{x} = (1, 0, 0)^T$, $(0, 1, 0)^T$ or $(0, 0, 1)^T$, as $i = 1, 2$ or 3 respectively.
- The epipole $\mathbf{e}'$ is the common intersection of the epipolar lines $\mathbf{l}_i'$ for $i = 1, 2, 3$.
- The *sum* of the matrices $\mathbf{M}(\mathbf{x}) = (\sum_i x^i \mathbf{T}_i)$ also has rank 2. The right null-vector of $\mathbf{M}(\mathbf{x})$ is the epipolar line $\mathbf{l}''$ of $\mathbf{x}$ in the third view, and its left null-vector is the epipolar line $\mathbf{l}'$ of $\mathbf{x}$ in the second view.

It's worth emphasizing again that although a particular canonical form of the camera matrices $\mathbf{P}$, $\mathbf{P}'$ and $\mathbf{P}''$ is used in the derivation, the epipolar properties of the $\mathbf{T}_i$ matrices are independent of this choice.

15.1.4 Extracting the fundamental matrices

It is simple to compute the fundamental matrices $\mathbf{F}_{21}$ and $\mathbf{F}_{31}$ between the first¹ and the other views from the trifocal tensor. It was seen in section 9.2.1 (p242) that the epipolar line corresponding to some point can be derived by transferring the point to the other view via a homography and joining the transferred point to the epipole. Consider a point $\mathbf{x}$ in the first view. According to figure 15.2 and result 15.2, a line $\mathbf{l}''$ in the third view induces a homography from the first to the second view given by $\mathbf{x}' = ([\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3] \mathbf{l}'') \mathbf{x}$. The epipolar line corresponding to $\mathbf{x}$ is then found by joining $\mathbf{x}'$ to the epipole $\mathbf{e}'$. This gives $\mathbf{l}' = [\mathbf{e}']_{\times} ([\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3] \mathbf{l}'') \mathbf{x}$, from which it follows that

$$\mathbf{F}_{21} = [\mathbf{e}']_{\times} [\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3] \mathbf{l}''.$$

This formula holds for any vector $\mathbf{l}''$, but it is important to choose $\mathbf{l}''$ to avoid the degenerate condition where $\mathbf{l}''$ lies in the null-space of any of the $\mathbf{T}_i$. A good choice is $\mathbf{e}''$ since as has been seen $\mathbf{e}''$ is perpendicular to the right null-space of each $\mathbf{T}_i$. This gives the formula

$$\mathbf{F}_{21} = [\mathbf{e}']_{\times} [\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3] \mathbf{e}''.$$
 (15.8)

A similar formula holds for $\mathbf{F}_{31} = [\mathbf{e}'']_{\times} [\mathbf{T}_1^T, \mathbf{T}_2^T, \mathbf{T}_3^T] \mathbf{e}'$.

15.1.5 Retrieving the camera matrices

It was remarked that the trifocal tensor, since it expresses a relationship between image entities only, is independent of 3D projective transformations. Conversely, this implies that the camera matrices may be computed from the trifocal tensor only up to a projective ambiguity. It will now be shown how this may be done.

Just as in the case of reconstruction from two views, because of the projective ambiguity, the first camera may be chosen as $\mathbf{P} = [\mathbf{I} \mid \mathbf{0}]$. Now, since $\mathbf{F}_{21}$ is known (from (15.8)), we can make use of result 9.9 (p254) to derive the form of the second camera as

$$\mathbf{P}' = [[\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3] \mathbf{e}'' \mid \mathbf{e}']$$

and the camera pair $\{\mathbf{P}, \mathbf{P}'\}$ then has the fundamental matrix $\mathbf{F}_{21}$. It might be thought

¹ The fundamental matrix $\mathbf{F}_{21}$ satisfies $\mathbf{x}'^T \mathbf{F}_{21} \mathbf{x} = 0$ for corresponding points $\mathbf{x} \leftrightarrow \mathbf{x}'$. The subscript notation refers to figure 15.8.

Given the trifocal tensor written in matrix notation as $[T_1, T_2, T_3]$.

(i) **Retrieve the epipoles e', e''**

Let u_i and v_i be the left and right null-vectors respectively of T_i , i.e. $u_i^T T_i = 0^T$, $T_i v_i = 0$. Then the epipoles are obtained as the null-vectors to the following 3×3 matrices:

$$e'^T [u_1, u_2, u_3] = 0 \text{ and } e''^T [v_1, v_2, v_3] = 0.$$

(ii) **Retrieve the fundamental matrices F_{21}, F_{31}**

$$F_{21} = [e']_{\times} [T_1, T_2, T_3] e'' \text{ and } F_{31} = [e'']_{\times} [T_1^T, T_2^T, T_3^T] e'.$$

(iii) **Retrieve the camera matrices P', P'' (with $P = [I \mid 0]$)**

Normalize the epipoles to unit norm. Then

$$P' = [[T_1, T_2, T_3] e'' \mid e'] \text{ and } P'' = [(e'' e''^T - I) [T_1^T, T_2^T, T_3^T] e' \mid e''].$$

Algorithm 15.1. Summary of F and P retrieval from the trifocal tensor. *Note, F_{21} and F_{31} are determined uniquely. However, P' and P'' are determined only up to a common projective transformation of 3-space.*

that the third camera could be chosen in a similar manner as $P'' = [[T_1^T, T_2^T, T_3^T] e' \mid e'']$, but *this is incorrect*. This is because the two camera pairs $\{P, P'\}$ and $\{P, P''\}$ do not necessarily define the same projective world frame; although each pair is correct by itself, the triple $\{P, P', P''\}$ is inconsistent.

The third camera cannot be chosen independently of the projective frame of the first two. To see this, suppose the camera pair $\{P, P'\}$ is chosen and points X_i reconstructed from their image correspondences $x_i \leftrightarrow x'_i$. Then the coordinates of X_i are specified in the projective world frame defined by $\{P, P'\}$, and a consistent camera P'' may be computed from the correspondences $X_i \leftrightarrow x''_i$. Clearly, P'' depends on the frame defined by $\{P, P'\}$. However, it is not necessary to explicitly reconstruct 3D structure, a consistent camera triplet can be recovered from the trifocal tensor directly.

The pair of camera matrices $P = [I \mid 0]$ and $P' = [[T_1, T_2, T_3] e'' \mid e']$ are not the only ones compatible with the given fundamental matrix F_{21} . According to (9.10-p256), the most general form for P' is

$$P' = [[T_1, T_2, T_3] e'' + e' v^T \mid \lambda e']$$

for some vector v and scalar λ . A similar choice holds for P'' . To find a triple of camera matrices compatible with the trifocal tensor, we need to find the correct values of P' and P'' from these families so as to be compatible with the form (15.1) of the trifocal tensor.

Because of the projective ambiguity, we are free to choose $P' = [[T_1, T_2, T_3] e'' \mid e']$, thus $a_i = T_i e''$. This choice fixes the projective world frame so that P'' is now defined uniquely (up to scale). Then substituting into (15.1) (observing that $a_i = e'$ and $b_i = e''$)

$$T_i = T_i e'' e''^T - e' b_i^T$$

from which it follows that $e' b_i^T = T_i (e'' e''^T - I)$. Since the scale may be chosen such

that $\|e'\| = e'^T e' = 1$, we may multiply on the left by e'^T and transpose to get

$$b_i = (e'' e'^T - I) T_i^T e'$$

so $P'' = [(e'' e'^T - I) [T_1^T, T_2^T, T_3^T] e' | e'']$. A summary of the steps involved in extracting the camera matrices from the trifocal tensor is given in algorithm 15.1.

We have seen that the trifocal tensor may be computed from the three camera matrices, and that conversely the three camera matrices may be computed, up to projective equivalence, from the trifocal tensor. Thus, the trifocal tensor completely captures the three cameras up to projective equivalence.

15.2 The trifocal tensor and tensor notation

The style of notation that has been used up to now for the trifocal tensor is derived from the standard matrix–vector notation. Since a matrix has two indices only, it is possible to distinguish between the two indices using the devices of matrix transposition and right or left multiplication, and in dealing with matrices and vectors, one can do without writing the indices explicitly. Because the trifocal tensor has three indices, instead of the two indices that a matrix has, it becomes increasingly cumbersome to persevere with this style of matrix notation, and we now turn to using standard tensor notation when dealing with the trifocal tensor. For those unfamiliar with tensor notation a gentle introduction is given in appendix 1(p562). This appendix should be read before proceeding with this chapter.

Image points and lines are represented by homogeneous *column* and *row* 3-vectors, respectively, i.e. $x = (x^1, x^2, x^3)^T$ and $l = (l_1, l_2, l_3)$. The ij -th entry of a matrix A is denoted by a_{ij}^i , index i being the contravariant (row) index and j being the covariant (column) index. We observe the convention that indices repeated in the contravariant and covariant positions imply summation over the range $(1, \dots, 3)$ of the index. For example, the equation $x' = Ax$ is equivalent to $x'^i = \sum_j a_{ij}^i x^j$, which may be written $x'^i = a_{ij}^i x^j$.

We begin with the definition of the trifocal tensor given in (15.1). Using tensor notation, this becomes

$$T_i^{jk} = a_i^j b_4^k - a_4^j b_i^k. \quad (15.9)$$

The positions of the indices in T_i^{jk} (two contravariant and one covariant) are dictated by the positions of the indices on the right side of the equation. Thus, the trifocal tensor is a mixed contravariant–covariant tensor. In tensor notation, the basic incidence relation (15.3) becomes

$$l_i = l'_j l''_k T_i^{jk}. \quad (15.10)$$

Note that when multiplying tensors the order of the entries does not matter, in contrast with standard matrix notation. For instance the right side of the above expression is

$$l'_j l''_k T_i^{jk} = \sum_{j,k} l'_j l''_k T_i^{jk} = \sum_{j,k} l'_j T_i^{jk} l''_k = l'_j T_i^{jk} l''_k.$$

Definition. The trifocal tensor $\mathcal{T}$ is a valency 3 tensor $\mathcal{T}_i^{jk}$ with two contravariant and one covariant indices. It is represented by a homogeneous $3 \times 3 \times 3$ array (i.e. 27 elements). It has 18 degrees of freedom.

Computation from camera matrices. If the canonical 3×4 camera matrices are

$$P = [I \mid 0], \quad P' = [a_j^i], \quad P'' = [b_j^i]$$

then

$$\mathcal{T}_i^{jk} = a_i^j b_i^k - a_i^k b_i^j.$$

See (17.12–p415) for computation from three general camera matrices.

Line transfer from corresponding lines in the second and third views to the first.

$$l_i = l'_j l''_k \mathcal{T}_i^{jk}$$

Transfer by a homography.

(i) **Point transfer from first to third view via a plane in the second**

The contraction $l'_j \mathcal{T}_i^{jk}$ is a homography mapping between the first and third views induced by a plane defined by the back-projection of the line l' in the second view.

$$x''^k = h_i^k x^i \quad \text{where} \quad h_i^k = l'_j \mathcal{T}_i^{jk}$$

(ii) **Point transfer from first to second view via a plane in the third**

The contraction $l''_k \mathcal{T}_i^{jk}$ is a homography mapping between the first and second views induced by a plane defined by the back-projection of the line l'' in the third view.

$$x'^j = h_i^j x^i \quad \text{where} \quad h_i^j = l''_k \mathcal{T}_i^{jk}$$

Table 15.2. Definition and transfer properties of the trifocal tensor.

The homography maps of figure 15.2 and figure 15.3 may be deduced from the incidence relation (15.10). In the case of the plane defined by back-projecting the line l' ,

$$l_i = l'_j l''_k \mathcal{T}_i^{jk} = l''_k (l'_j \mathcal{T}_i^{jk}) = l''_k h_i^k \quad \text{where} \quad h_i^k = l'_j \mathcal{T}_i^{jk}$$

and h_i^k are the elements of the homography matrix H . This homography maps points between the first and third view as

$$x''^k = h_i^k x^i.$$

Note that the homography is obtained from the tensor by contraction with a line (i.e. a summation over one contravariant (upper) index of the tensor, and the covariant (lower) index of the line), i.e. l' extracts a 3×3 matrix from the tensor – think of the trifocal tensor as an operator which takes a line and produces a homography matrix. Table 15.2 summarizes the definition and transfer properties of the trifocal tensor.

A pair of particularly important tensors are ϵ_{ijk} and its contravariant counterpart ϵ^{ijk} , defined in section A1.1(p563). This tensor is used to represent the vector product. For

(i) Line-line-line correspondence	$(l_r \epsilon^{ris}) l'_j l''_k T_i^{jk} = 0^s$
(ii) Point-line-line correspondence	$x^i l'_j l''_k T_i^{jk} = 0$
(iii) Point-line-point correspondence	$x^i l'_j (x^{rk} \epsilon_{kqs}) T_i^{jq} = 0_s$
(iv) Point-point-line correspondence	$x^i (x^{rj} \epsilon_{jpr}) l''_k T_i^{pk} = 0_r$
(v) Point-point-point correspondence	$x^i (x^{rj} \epsilon_{jpr}) (x^{rk} \epsilon_{kqs}) T_i^{pq} = 0_{rs}$

Table 15.3. Summary of trifocal tensor incidence relations – the trilinearities.

instance, the line joining two points x^i and y^j is equal to the cross product $x^i y^j \epsilon_{ijk} = l_k$, and the skew-symmetric matrix $[x]_\times$ is written as $x^i \epsilon_{irs}$ in tensor notation. It is now relatively straightforward to write down the basic incidence results involving the trifocal tensor given in table 15.1. The results are summarized in table 15.3. In this table, a notation such as 0_r represents an array of zeros.

The form of the relations in table 15.3 is more easily understood if one observes that three indices i, j and k in T_i^{jk} correspond to entities in the first, second and third views respectively. Thus for instance a partial expression such as $l'_j T_i^{jk}$ cannot occur, because the index j belongs to the second view, and hence does not belong on the line l' in the third view. Repeated indices (indicating summation) must occur once as a contravariant (upper) index and once as a covariant (lower) index. Thus, we cannot write $x^{jj} T_i^{jk}$, since the index j occurs twice in the upper position. Think of the ϵ tensor as being used to raise or lower indices, for instance by replacing l'_j by $x^i \epsilon_{ijk}$. However, this may not be done arbitrarily, as pointed out in exercise (x) on page 389.

15.2.1 The trilinearities

The incidence relations in table 15.3 are *trilinear* relations or *trilinearities* in the coordinates of the image elements (points and lines). *Tri-* since every monomial in the relation involves a coordinate from each of the three image elements involved; and *linear* because the relations are linear in each of the algebraic entities (i.e. the three “arguments” of the tensor). For example in the point-point-point relation, $x^i (x^{rj} \epsilon_{jpr}) (x^{rk} \epsilon_{kqs}) T_i^{pq} = 0_{rs}$, suppose both x_1 and x_2 satisfy the relation, then so does $x = \alpha x_1 + \beta x_2$, i.e. the relation is linear in its first argument. Similarly, the relation is linear in the second and third argument. This multi-linearity is a standard property of tensors, and follows directly from the form $x^i l'_j l''_k T_i^{jk} = 0$ which is a contraction of the tensor over all three of its indices (arguments).

We will now describe the point-point-point trilinearities in more detail. There are

nine of these trilinearities arising from the three choices of r and s . Geometrically these trilinearities arise from special choices of the lines in the second and third image for the point–line–line relation (see figure 15.4(a)). Choosing $r = 1, 2$ or 3 corresponds to a line parallel to the image x -axis, parallel to the image y -axis, or through the image coordinate origin (the point $(0, 0, 1)^T$), respectively. For example, choosing $r = 1$ and expanding $x'^j \epsilon_{jpr}$ results in

$$l'_p = x'^j \epsilon_{jpr} = (0, -x'^3, x'^2)$$

which is a horizontal line in the second view through $\mathbf{x}'$ (since points of the form $\mathbf{y}' = (x'^1 + \lambda, x'^2, x'^3)^T$ satisfy $\mathbf{y}'^T \mathbf{l}' = 0$ for any λ). Similarly, choosing $s = 2$ in the third view results in the vertical line through $\mathbf{x}''$

$$l''_q = x''^k \epsilon_{kq2} = (x''^3, 0, -x''^1)$$

and the trilinear point relation expands to

$$\begin{aligned} 0 &= x^i x'^j x''^k \epsilon_{jpr} \epsilon_{kq2} T_i^{pq} \\ &= x^i [-x'^3 (x''^3 T_i^{21} - x''^1 T_i^{23}) + x'^2 (x''^3 T_i^{31} - x''^1 T_i^{33})]. \end{aligned}$$

Of these nine trilinearities, four are linearly independent. This means that from a basis of four trilinearities all nine can be generated by linear combinations. The four degrees of freedom may be traced back to those of the point–line–line relation $x^i l'_j l''_k T_i^{jk} = 0$ and are counted as follows. There is a one-parameter family of lines through $\mathbf{x}''$ in the third view. If $\mathbf{m}''$ and $\mathbf{n}''$ are two members of this family, then any other line through $\mathbf{x}''$ can be obtained from a linear combination of these:

$$\mathbf{l}'' = \alpha \mathbf{m}'' + \beta \mathbf{n}''.$$

The incidence relation is linear in $\mathbf{l}''$, so that given

$$\begin{aligned} l'_j m''_k T_i^{jk} x^i &= 0 \\ l'_j n''_k T_i^{jk} x^i &= 0 \end{aligned}$$

then the incidence relation for any other line $\mathbf{l}''$ can be generated by a linear combination of these two. Consequently, there are only two linearly independent incidence relations for $\mathbf{l}''$. Similarly there is a one-parameter family of lines through $\mathbf{x}'$, and the incidence relation is also linear in lines $\mathbf{l}'$ through $\mathbf{x}'$. Thus, there are a total of four linearly independent incidence relations between a point in the first view and lines in the second and third.

The main virtue of the trilinearities is that they are linear, otherwise their properties are often subsumed by transfer, as described in the following section.

15.3 Transfer

Given three views of a scene and a pair of matched points in two views one may wish to determine the position of the point in the third view. Given sufficient information about the placement of the cameras, it is usually possible to determine the location of




Fig. 15.7. **Epipolar transfer.** (a) The image of X in the first two views is the correspondence $x \leftrightarrow x'$. The image of X in the third view may be computed by intersecting the epipolar lines $F_{31}x$ and $F_{32}x'$. (b) The configuration of the epipoles and transferred point x'' as seen in the third image. Point x'' is computed as the intersection of epipolar lines passing through the two epipoles e_{31} and e_{32} . However, if x'' lies on the line through the two epipoles, then its position cannot be determined. Points close to the line through the epipoles will be estimated with poor precision.

the point in the third view without reference to image content. This is the point transfer problem. A similar transfer problem arises for lines.

In principle the problem can generally be solved given the cameras for the three views. Rays back-projected from corresponding points in the first and second view intersect and thus determine the 3D point. The position of the corresponding point in the third view is computed by projecting this 3D point onto the image. Similarly lines back-projected from the first and second image intersect in the 3D line, and the projection of this line in 3-space to the third image determines its image position.

15.3.1 Point transfer using fundamental matrices

The transfer problem may be solved using knowledge of the fundamental matrices only. Thus, suppose we know the three fundamental matrices F_{21} , F_{31} and F_{32} relating the three views, and let points x and x' in the first two views be a matched pair. We wish to find the corresponding point x'' in the third image.

The required point x'' matches point x in the first image, and consequently must lie on the epipolar line corresponding to x . Since we know F_{31} , this epipolar line may be computed, and is equal to $F_{31}x$. By a similar argument, x'' must lie on the epipolar line $F_{32}x'$. Taking the intersection of the epipolar lines gives

$$x'' = (F_{31}x) \times (F_{32}x')$$

See figure 15.7a.

Note that the fundamental matrix F_{21} is not used in this expression. The question naturally arises whether we can gain anything by knowledge of F_{21} , and the answer is yes. In the presence of noise, the points $x \leftrightarrow x'$ will not form an exact matched pair, meaning that they will not satisfy the equation $x'^T F_{21}x = 0$ exactly. Given F_{21} one may use optimal triangulation as in algorithm 12.1(p318) to correct x and x' , resulting in a pair $\hat{x} \leftrightarrow \hat{x}'$ that satisfies this relation. The transferred point may then be computed as $x'' = (F_{31}\hat{x}) \times (F_{32}\hat{x}')$. This method of point transfer using the fundamental matrices will be called *epipolar transfer*.




Fig. 15.8. **The trifocal plane** is defined by the three camera centres. The notation for the epipoles is $e_{ij} = P_i C_j$. Epipolar transfer fails for any point X on the trifocal plane. If the three camera centres are collinear then there is a one-parameter family of planes containing the three centres.

Though at one time used for point transfer, epipolar transfer has a serious deficiency that rules it out as a practical method. This deficiency is due to the degeneracy that can be seen from figure 15.7(b): epipolar transfer fails when the two epipolar lines in the third image are coincident (and becomes increasingly ill-conditioned as the lines become less “transverse”). The degeneracy condition that x'' , e_{31} and e_{32} are collinear in the third image means that the camera centres C and C' and the 3D point X lie in a plane through the centre C'' of the third camera; thus X lies on the *trifocal plane* defined by the three camera centres, see figure 15.8. Epipolar transfer will fail for points X lying on the trifocal plane and will be inaccurate for points lying near that plane. Note, in the special case that the three camera centres are collinear the trifocal plane is not uniquely defined, and epipolar transfer fails for all points. In this case $e_{31} = e_{32}$.

15.3.2 Point transfer using the trifocal tensor

The degeneracy of epipolar transfer is avoided by use of the trifocal tensor. Consider a correspondence $x \leftrightarrow x'$. If a line l' passing through the point x' is chosen in the second view, then the corresponding point x'' may be computed by transferring the point x from the first to the third view using $x''^k = x'^l T_i^{jk}$, from table 15.2. It is clear from figure 15.4(p371)(b) that this transfer is not degenerate for general points X lying on the trifocal plane.

However, note from result 15.3 and figure 15.6 that if l' is the epipolar line corresponding to x , then $x'^l T_i^{jk} = 0^k$, so the point x'' is undefined. Consequently, the choice of line l' is important. To avoid choosing only an epipolar line, one possibility is to use two or three different lines passing through x' , namely $l'_{jp} = x'^r \epsilon_{rjp}$ for the three choices of $p = 1, \dots, 3$. For each such line, one computes the value of x'' and retains the one that has the largest norm (i.e. is furthest from being zero). An alternative method entirely for finding x'' is as the least-squares solution of the system of linear equations $x^i (x'^j \epsilon_{jpr}) (x''^k \epsilon_{kqs}) T_i^{pq} = 0_{rs}$, but this method is probably an overkill.

The method we recommend is the following. Before attempting to compute the point x'' transferred from a pair of points $x \leftrightarrow x'$, first correct the pair of points using the fundamental matrix F_{21} , as described above in the case of epipolar transfer. If $\hat{x}$ and $\hat{x}'$




Fig. 15.9. **Degeneracy for point transfer using the trifocal tensor.** The 3D point X is defined by the intersection of the ray through x with the plane π' . A point X on the baseline B_{12} between the first and second views cannot be defined in this manner. So a 3D point on the line B_{12} cannot be transferred to the third view via a homography defined by a line in the second image. Note that a point on the line B_{12} projects to e_{12} in the first image and e_{21} in the second image. Apart from the line B_{12} any point can be transferred. In particular there is not a degeneracy problem for points on the baseline B_{23} , between views two and three, or for any other point on the trifocal plane.

are an exact match, then the transferred point $x^{nk} = \hat{x}^i l'_j T_i^{jk}$ does not depend on the line l' chosen passing through $\hat{x}'$ (as long as it is not the epipolar line). This may be verified geometrically by referring to figure 15.2(p368). A good choice is always given by the line perpendicular to $F_{21}\hat{x}$.

To summarize, a measured correspondence $x \leftrightarrow x'$ is transferred by the following steps:

- (i) Compute F_{21} from the trifocal tensor (by the method given in algorithm 15.1), and correct $x \leftrightarrow x'$ to the exact correspondence $\hat{x} \leftrightarrow \hat{x}'$ using algorithm 12.1-(p318).
- (ii) Compute the line l' through $\hat{x}'$ and perpendicular to $l'_e = F_{21}\hat{x}$. If $l'_e = (l_1, l_2, l_3)^T$ and $\hat{x}' = (\hat{x}_1, \hat{x}_2, 1)^T$, then $l' = (l_2, -l_1, -\hat{x}_1 l_2 + \hat{x}_2 l_1)^T$.
- (iii) The transferred point is $x^{nk} = \hat{x}^i l'_j T_i^{jk}$.

Degenerate configurations. Consider transfer to the third view via a plane, as shown in figure 15.9. The 3D point X is only undefined if it lies on the baseline joining the first and second camera centres. This is because rays through x and x' are collinear for such 3D points and so their intersection is not defined. In such a case, the points x and x' correspond with the epipoles in the two images. However, there is no problem transferring a point lying on the baseline between views two and three, or anywhere else on the trifocal plane. This is the key difference between epipolar transfer and transfer using the trifocal tensor. The former is undefined for *any* point on the trifocal plane.

15.3.3 Line transfer using the trifocal tensor

Using the trifocal tensor, it is possible to transfer lines from a pair of images to a third according to the line-transfer equation $l_i = l'_j l''_k T_i^{jk}$ of table 15.2. This gives an explicit

formula for the line in the first view, given lines in the other two views. Note however that if the lines l and l' are known in the first and second views then l'' may be computed by solving the set of linear equations $(l_r e^{ris}) l'_j l''_k T_i^{jk} = 0^s$, thereby transferring it into the third image. Similarly one may transfer lines into the second image. Line transfer is not possible using only the fundamental matrices.

Degeneracies. Consider the geometry of figure 12.8(p322). The line L in 3-space is defined by the intersection of the planes through l and l' , namely π and π' respectively. This line is clearly undefined when the planes π and π' are coincident, i.e. in the case of epipolar planes. Consequently, lines cannot be transferred between the first and third image if both l and l' are corresponding epipolar lines for the first and second views. Algebraically, the line-transfer equation gives $l_i = l'_j l''_k T_i^{jk} = 0$, and the equation matrix $(l_r e^{ris}) l'_j T_i^{jk}$ used to solve for l'' becomes zero. It is quite common for lines to be near epipolar, and their transfer is then inaccurate, so this condition should always be checked for. There is an equivalent degeneracy for line transfer between views one and two defined by a line in view three. Again, it occurs if the lines in views one and three are corresponding epipolar lines for these two views.

In general the epipolar geometries between views one and two, and one and three will differ, for instance the epipole e_{12} arising in the first view from view two will not coincide with the epipole e_{13} arising in the first view from view three. Thus an epipolar line in the first view for views one and two will not coincide with an epipolar line for views one and three. Consequently, when line transfer into the third view is degenerate, line transfer into the second view will not in general be degenerate. However, for lines in the trifocal plane transfer is degenerate (i.e. undefined) always.

15.4 The fundamental matrices for three views

The three fundamental matrices F_{21}, F_{31}, F_{32} are not independent, but satisfy three relations:

$$e_{23}^T F_{21} e_{13} = e_{31}^T F_{32} e_{21} = e_{32}^T F_{31} e_{12} = 0. \quad (15.11)$$

These relations are easily seen from figure 15.8. For example, $e_{32}^T F_{31} e_{12} = 0$ follows from the observation that e_{32} and e_{12} are matching points, corresponding to the centre of camera number 2.

Projectively, the three-camera configuration has 18 degrees of freedom counting 11 for each camera less 15 for an overall projective ambiguity. Alternatively, this may be accounted for as 21 for the 3×7 degrees of freedom of the fundamental matrices less 3 for the relations. The trifocal tensor also has 18 degrees of freedom and fundamental matrices computed from the trifocal tensor will automatically satisfy the three relations.

The counting argument implies that the three relations of (15.11) are sufficient to ensure consistency of three fundamental matrices. The counting argument alone is not a convincing proof of this, however, so a proof is given below.

Definition 15.5. Three fundamental matrices F_{21}, F_{31} and F_{32} are said to be *compatible* if they satisfy the conditions (15.11).

In most cases, these conditions are sufficient to ensure that the three fundamental matrices correspond to some geometric configuration of cameras.

Theorem 15.6. *Let a set of three fundamental matrices F_{21} , F_{31} and F_{32} be given satisfying the conditions (15.11). Assume also that $e_{12} \neq e_{13}$, $e_{21} \neq e_{23}$, and $e_{31} \neq e_{32}$. Then there exist three camera matrices P_1, P_2, P_3 such that F_{ij} is the fundamental matrix corresponding to the pair (P_i, P_j) .*

Note that the conditions $e_{ij} \neq e_{ik}$ in this theorem ensure that the three cameras are non-collinear. For this reason they will be referred to here as the *non-collinearity conditions*. One may show by example (left to the reader) that these conditions are necessary for the truth of the theorem.

Proof. In this proof, the indices i, j and k are intended to be distinct. We begin by choosing three points $x_i; i = 1, \dots, 3$, consistent with the three fundamental matrices. In other words, we require that $x_i^T F_{ij} x_j = 0$ for all pairs (i, j) . This is easily done by choosing first x_1 and x_2 to satisfy $x_2^T F_{21} x_1 = 0$, and then defining x_3 to be the intersection of the two epipolar lines $F_{32} x_2$ and $F_{31} x_1$.

In a similar manner, we choose a second set of points $y_i; i = 1, \dots, 3$ satisfying $y_i^T F_{ij} y_j = 0$. This is done in such a way that the four points x_i, y_i, e_{ij}, e_{ik} in each image i are in general position – that is no three are collinear. This is possible by the assumption that the two epipoles in each image are distinct.

Next we choose five world points C_1, C_2, C_3, X, Y in general position. For example, one could take the usual projective basis. We may now define the three camera matrices. Let the i -th camera matrix P_i satisfy the conditions

$$P_i C_i = 0; \quad P_i C_j = e_{ij}; \quad P_i C_k = e_{ik}; \quad P_i X = x_i; \quad P_i Y = y_i.$$

In other words, the i -th camera has centre at C_i and maps the four other world points C_j, C_k, X, Y to the four image points e_{ij}, e_{ik}, x_i, y_i . This uniquely determines the camera matrix since the points are in general position. To see this, recall that the camera matrix defines a homography between the image and the rays through the camera centre (a 2D projective space). The images of four points specify this homography completely. Let $\hat{F}_{ij}$ be the fundamental matrix defined by the pair of camera matrices P_i and P_j . The proof is completed by proving that $\hat{F}_{ij} = F_{ij}$ for all i, j .

The epipoles of $\hat{F}_{ij}$ and F_{ij} are the same, by the way that P_i and P_j are constructed. Consider the pencil of epipolar lines through e_{ij} in image i . This pencil forms a 1-dimensional projective space of lines, and the fundamental matrix F_{ij} induces a one-to-one correspondence (in fact a homography) between this pencil and the pencil of lines through e_{ji} in image j . The fundamental matrix $\hat{F}_{ij}$ also induces a homography between the same pencils. The two fundamental matrices are the same if the homographies they induce are the same.

Two 1-dimensional homographies are the same if they agree on three points (or in this case epipolar lines). The relation $x_i^T F_{ij} x_j = 0$ means that the epipolar lines through x_i in image i and x_j in image j correspond under the homography induced by F_{ij} . By construction $x_i^T \hat{F}_{ij} x_j = 0$ as well, since x_i and x_j are the projections of the point X in

the two images. Thus, both homographies agree on this pair of epipolar lines. In the same way, the homographies induced by F_{ij} and $\hat{F}_{ij}$ agree on the epipolar lines corresponding to the pairs $y_i \leftrightarrow y_j$ and $e_{ik} \leftrightarrow e_{jk}$. The two homographies therefore agree on three lines in the pencil and hence are equal; so are the corresponding fundamental matrices. (We are grateful to Frederik Schaffalitzky for this proof). $\square$

15.4.1 Uniqueness of camera matrices given three fundamental matrices

The proof just given shows that there is at least one set of cameras corresponding to three compatible fundamental matrices (provided they satisfying the non-collinearity condition). It is important to know that the three fundamental matrices determine the configuration of the three cameras uniquely, at least up to the unavoidable projective ambiguity. This will be shown next.

The first two camera matrices P and P' may be determined from the fundamental matrix F_{21} by two-view techniques (chapter 9). It remains to determine the third camera matrix P'' in the same projective frame. In principle, this may be done as follows.

- (i) Select a set of matching points $x_i \leftrightarrow x'_i$ in the first two images, satisfying $x_i'^T F_{21} x_i = 0$, and use triangulation to determine the corresponding 3D points X_i .
- (ii) Use epipolar transfer to determine the corresponding points x''_i in the third image, using the fundamental matrices F_{31} and F_{32} .
- (iii) Solve for the camera matrix P'' from the set of 3D–2D correspondences $X_i \leftrightarrow x''_i$.

The second step in this algorithm will fail in the case where the point X_i lies in the trifocal plane. Such a point X_i is easily detected and discarded, since it projects into the first image as a point x_i lying on the line joining the two epipoles e_{12} and e_{13} . Since there are infinitely many possible matched points, we can compute sufficiently many such points to compute P'' .

The only situation in which this method will fail is when all space points X_i lie in a trifocal plane. This can occur only in the degenerate situation in which the three camera centres are collinear, in which case the trifocal plane is not uniquely determined. Thus, we see that unless the three camera centres are collinear, the three camera matrices may be determined from the fundamental matrices. On the other hand, if the three cameras are collinear, then there is no way to determine the relative spacings of the cameras along the line of their centres. This is because the length of the baseline cannot be determined from the fundamental matrices, and the three baselines (distances between the camera centres) may be arbitrarily chosen and remain consistent with the fundamental matrices. Thus we have demonstrated the following fact:

Result 15.7. *Given three compatible fundamental matrices F_{21} , F_{31} and F_{32} satisfying the non-collinearity condition, the three corresponding camera matrices P , P' and P'' are unique up to the choice of a 3D projective coordinate frame.*

15.4.2 Computation of camera matrices from three fundamental matrices

Given three compatible fundamental matrices, there exists a simple method for computing a corresponding set of three camera matrices. From the fundamental matrix F_{21} , one can compute a corresponding pair of camera matrices (P, P') using result 9.14-(p256). Next, according to result 9.12(p255) the third camera matrix P'' must satisfy the condition that $P'^T F_{31} P$ and $P''^T F_{32} P'$ be skew-symmetric. Each of these matrices gives rise to 10 linear equations in the entries of P'' , a total of 20 equations in the 12 entries of P'' . From these, P'' may be computed linearly.

If the three fundamental matrices are compatible in the sense of definition 15.5 and the non-collinearity condition of theorem 15.6 holds, then there will exist a solution, and it will be unique. If however the three fundamental matrices are computed independently from point correspondences, then they will not satisfy the compatibility conditions exactly. In this case it will be necessary to compute a least-squares solution to find P'' . The error being minimized is not geometrically based. It is best to use this algorithm only when the fundamental matrices are known to be compatible.

One can think of doing three-view reconstruction by estimating the three fundamental matrices using pairwise point correspondences, then using the above algorithm to estimate the three camera matrices. This is not a very good strategy, for the following reasons.

- (i) The method for computing the three camera matrices from the fundamental matrices assumes that the fundamental matrices are compatible. Otherwise, a least-squares problem involving a non-geometrically justified cost function is involved.
- (ii) Although result 15.7 shows that three fundamental matrices may determine the camera geometry, and hence the trifocal tensor, this is only true when the cameras are not collinear. As they approach collinearity, the estimate of the relative camera placement becomes unstable.

The trifocal tensor is preferable to a triple of compatible fundamental matrices as a means of determining the geometry of three views. This is because the difficulty with the views being collinear is not an issue with the trifocal tensor. It is well defined and uniquely determines the geometry even for collinear cameras. The difference is that the fundamental matrices do not contain a direct constraint on the relative displacements between the three cameras, whereas this is built into the trifocal tensor.

Since the projective structure of the three cameras may be computed explicitly from the trifocal tensor, it follows that all three fundamental matrices for the three view pairs are determined by the trifocal tensor. In fact simple formulae, given in algorithm 15.1-(p375) exist for the two fundamental matrices F_{21} and F_{31} . The fundamental matrices determined from the trifocal tensor will satisfy the compatibility conditions (15.11).

15.4.3 Camera matrices compatible with two fundamental matrices

Suppose we are given only two fundamental matrices F_{21} and F_{31} . To what extent do these fix the geometry of the three cameras? It will be shown here that there are four

degrees of freedom in the solution for the camera matrices, beyond the usual projective ambiguity.

From F_{21} one may compute a pair of camera matrices (P, P') , and from F_{31} the pair (P, P'') . In both cases we may choose $P = [I \mid 0]$, resulting in a triple of camera matrices (P, P', P'') compatible with the pair of fundamental matrices.

However, the choice of the three camera matrices is not unique, since for any matrices H_1 and H_2 representing 3D projective transforms, the pairs $(PH_1, P'H_1)$ and $(PH_2, P'H_2)$ are also compatible with the same fundamental matrices. In order to preserve the condition that P is equal to $[I \mid 0]$ in each case, the form of H_i must be restricted to:

$$H_i = \begin{bmatrix} I & 0 \\ \mathbf{v}_i^T & k_i \end{bmatrix}.$$

We may now fix on a particular choice of the first two camera matrices (P, P') compatible with F_{21} . This is equivalent to fixing on a specific projective coordinate frame. The general solution for the camera matrices is then $(P, P', P''H_2)$, where H_2 is of the form given above and the two pairs (P, P') and (P, P'') are compatible with the two fundamental matrices.

Allowing also for the overall projective ambiguity, the most general solution is $(PH, P'H, P''H_2H)$, which gives a total of 19 degrees of freedom, 15 for the projective transformation H and 4 for the degrees of freedom of H_2 . The same number of degrees of freedom may be found using a counting argument as follows: two fundamental matrices have 7 degrees of freedom each, for a total of 14. Three arbitrary camera matrices on the other hand have $3 \times 11 = 33$ degrees of freedom. The 14 constraints imposed by the two fundamental matrices leave 19 remaining degrees of freedom for the three camera matrices.

15.5 Closure

The development of three-view geometry proceeds in an analogous manner to that of two-view geometry covered in part II of this book. The trifocal tensor may be computed from image correspondences over three views, and a projective reconstruction of the cameras and 3D scene then follows. This computation is described in chapter 16. The projective ambiguity may be reduced to affine or metric by supplying additional information on the scene or cameras in the same manner as that of chapter 10. A similar development to that of chapter 13 may be given for the relations between homographies induced by scene planes and the trifocal tensor.

15.5.1 The literature

With hindsight, the discovery of the trifocal tensor may be traced to [Spetsakis-91] and [Weng-88], where it was used for scene reconstruction from lines in the case of calibrated cameras. It was later shown in [Hartley-94d] to be equally applicable to projective scene reconstruction in the uncalibrated case. At this stage matrix notation was used, but [Vieville-93] used tensor notation for this problem.

Meanwhile in independent work, Shashua introduced trilinearity conditions relating the coordinates of corresponding points in three views with uncalibrated cameras [Shashua-94, Shashua-95a]. [Hartley-95b, Hartley-97a] then showed that Shashua's relation for points and scene reconstruction from lines both arise from a common tensor, and the trifocal tensor was explicitly identified.

In subsequent work properties of the tensor have been investigated, e.g. [Shashua-95b]. In particular [Triggs-95] described the mixed covariant–contravariant behaviour of the indices, and [Zisserman-96] described the geometry of the homographies encoded by the tensor. Faugeras and Mourrain [Faugeras-95a] gave enlightening new derivations of the trifocal tensor equations and considered the trifocal tensor in the context of general linear constraints involving multiple views. This approach will be discussed in chapter 17. Further geometric properties of the tensor were given in Faugeras & Papadopoulos [Faugeras-97].

Epipolar point transfer was described by [Barrett-92, Faugeras-94], and its deficiencies pointed out by [Zisserman-94], amongst others.

The trifocal tensor has been used for various applications including establishing correspondences in image sequences [Beardsley-96], independent motion detection [Torr-95a], and camera self-calibration [Armstrong-96a].

15.5.2 Notes and exercises

- (i) The trifocal tensor is invariant to 3D projective transforms. Verify explicitly that if $H_{3 \times 4}$ is a transform preserving the first camera matrix $P = [I \mid 0]$, then the tensor defined by (15.1–p367) is unchanged.
- (ii) In this chapter the starting point for the trifocal tensor derivation was the incidence property of three corresponding lines. Show that alternatively the starting point may be the homography induced by a plane.
Here is a sketch derivation: choose the camera matrices to be a canonical set $P = [I \mid 0]$, $P' = [A \mid a_4]$, $P'' = [B \mid b_4]$ and start from the homography H_{13} between the first and third views induced by a plane π' . From result 13.1(p326) this homography may be written as $H_{13} = B - b_4 v^T$, where $\pi'^T = (v^T, 1)$. In this case the plane is defined by a line l' in the second view as $\pi' = P'^T l'$. Show that result 15.2(p369) follows.
- (iii) Homographies involving the first view are simply expressed in terms of the trifocal tensor T_i^{jk} as given by result 15.2(p369). Investigate whether a simple formula exists for the homography H_{23} from the second to the third view, induced by a line l in the first image.
- (iv) The contraction $x^i T_i^{jk}$ is a 3×3 matrix. Show that this may be interpreted as a correlation (see definition 2.29(p59)) mapping between the second and third views induced by the line which is the back-projection of the point x in the first view.
- (v) **Plane plus parallax over three views.** There is a rich geometry associated with the plane plus two points configuration (see figure 13.9(p336)) over three views: suppose the points off the (reference) plane are X and Y . Project the

point X onto the reference plane from each of the three camera centres to form a triangle x, x', x'' , and similarly project the point Y to the triangle y, y', y'' . Then the two triangles form a Desargues's configuration and are related by a planar homology (see section A7.2(p629)). A simple sketch shows that the lines joining corresponding triangle vertices, $(x, y), (x', y'), (x'', y'')$, are concurrent, and their intersection is the point at which the line joining X and Y pierces the reference plane. Similarly the intersection points of corresponding triangle sides are collinear, and the line so formed is the intersection of the trifocal plane of the cameras with the reference plane. Further details are given in [Criminisi-98, Irani-98, Triggs-00b].

- (vi) In the case where two of the three cameras have the same camera centre, the trifocal tensor may be related to simpler entities. There are two cases.
- (a) If the second and third camera have the same centre, then $T_i^{jk} = F_{ri} H_s^k \epsilon^{rjs}$, where F_{ri} is the fundamental matrix for the first two views, and H is the homography from the second to the third view induced by the fact that they have the same centre.
 - (b) If the first and the second views have the same centre, then $T_i^{jk} = H_i^j e''^k$, where H is the homography from the first to the second view and e'' is the epipole in the third image.

Prove these relationships using the approach of chapter 17.

- (vii) Consider the case of a small baseline between the cameras and derive a differential form of the trifocal tensor, see [Astrom-98, Triggs-99b].
- (viii) There are actually three different trifocal tensors relating three views, depending on which of the three cameras corresponds to the covariant index. Given one such tensor $[T_i]$, verify that the tensor $[T'_i]$ may be computed in several steps, as follows:
- (a) Extract the three camera matrices $P = [I \mid 0]$, P' and P'' from the trifocal tensor.
 - (b) Find a 3D projective transformation H such that $P'H = [I \mid 0]$, and apply it to each of P and P'' as well.
 - (c) Compute the tensor $[T'_i]$ by applying (15.1-p367).
- (ix) Investigate the form and properties (e.g. rank of the matrices T_i) of the trifocal tensor for the special motions (pure translation, planar motion) described in section 9.3(p247) for the fundamental matrix.
- (x) Comparison of the incidence relationships of table 15.3(p378) indicates that one may replace a line l_j^i by the expression $\epsilon_{jrs} x_r^i$, and proceed similarly with l_k^i . Also, one gets a three-view line equation by replacing x^i by $\epsilon^{irs} l_i$. Can both of these operations be carried out at once to obtain an equation

$$(\epsilon^{iru} l_i) (\epsilon_{jst} x_t^j) (\epsilon_{ktw} x_w^k) T_r^{st} = 0_{vw}?$$

Why, or why not?

- (xi) **Affine trifocal tensor.** If the three cameras P , P' and P'' are all affine (definition 6.3(p166)), then the corresponding tensor $\mathcal{T}_A$ is the *affine trifocal tensor*. This affine specialization of the tensor has 12 degrees of freedom and 16 non-zero entries. The affine trifocal tensor was first defined in [Torr-95b], and has been studied in [Kahl-98a, Quan-97a, Thorhallsson-99]. It shares with the affine fundamental matrix (chapter 14) very stable numerical estimation behaviour. It has been shown to perform very well in tracking applications where the object of interest (for example a car) has small relief compared to the depth of the scene [Hayman-03, Tordoff-01].

Computation of the Trifocal Tensor $\mathcal{T}$

This chapter describes numerical methods for estimating the trifocal tensor given a set of point and line correspondences across three views. The development will be very similar to that for the fundamental matrix, using much the same techniques as those of chapter 11. In particular, five methods will be discussed:

- (i) A linear method based on direct solution of a set of linear equations (after appropriate data normalization) (section 16.2).
- (ii) An iterative method, that minimizes the algebraic error, while satisfying all appropriate constraints on the tensor (section 16.3).
- (iii) An iterative method that minimizes geometric error (the “Gold Standard” method) (section 16.4.1).
- (iv) An iterative method that minimizes the Sampson approximation to geometric error (section 16.4.3).
- (v) Robust estimation based on RANSAC (section 16.6).

16.1 Basic equations

A complete set of the (tri-)linear equations involving the trifocal tensor is given in table 16.1. All of these equations are linear in the entries of the trifocal tensor $\mathcal{T}$.

Correspondence	Relation	Number of equations
three points	$x^i x'^j x''^k \epsilon_{jqs} \epsilon_{krt} \mathcal{T}_i^{qr} = 0_{st}$	4
two points, one line	$x^i x'^j l''_r \epsilon_{jqs} \mathcal{T}_i^{qr} = 0_s$	2
one point, two lines	$x^i l''_q l''_r \mathcal{T}_i^{qr} = 0$	1
three lines	$l_p l''_q l'''_r \epsilon^{piw} \mathcal{T}_i^{qr} = 0^w$	2

Table 16.1. **Trilinear relations between point and line coordinates in three views.** The final column denotes the number of linearly independent equations. The notation 0_{st} means a 2-dimensional tensor with all zero entries. Thus, the first line in this table corresponds to a set of 9 equations, one for each choice of s and t . However, among this set of 9 equations, only 4 are linearly independent.

Given several point or line correspondences between three images, the complete set of equations generated is of the form $\mathcal{A}\mathbf{t} = \mathbf{0}$, where $\mathbf{t}$ is the 27-vector made up of the entries of the trifocal tensor. From these equations, one may solve for the entries of the tensor. Note that equations involving points may be combined with those involving lines – in general all available equations from table 16.1 may be used simultaneously. Since $\mathcal{T}$ has 27 entries, 26 equations are needed to solve for $\mathbf{t}$ up to scale. With more than 26 equations, a least-squares solution is computed. As with the fundamental matrix, one minimizes $\|\mathcal{A}\mathbf{t}\|$ subject to the constraint $\|\mathbf{t}\| = 1$ using algorithm A5.4(p593).

This gives a bare outline of a linear algorithm for computing the trifocal tensor. However, in order to build a practical algorithm out of this several issues, such as normalization, need to be addressed. In particular the tensor that is estimated must obey various constraints, and we consider these next.

16.1.1 The internal constraints

The most notable difference between the fundamental matrix and the trifocal tensor is the greater number of constraints that apply to the trifocal tensor. The fundamental matrix has a single constraint, namely $\det(\mathbf{F}) = 0$, leaving 7 degrees of freedom, discounting the arbitrary scale factor. The trifocal tensor, on the other hand, has 27 entries, but 18 parameters only are required to specify the equivalent camera configuration, up to projectivity. The elements of the tensor therefore satisfy 8 independent algebraic constraints. This condition is conveniently stated as follows.

Definition 16.1. A trifocal tensor $\mathcal{T}_i^{jkl}$ is said to be “geometrically valid” or “satisfy all internal constraints” if there exist three camera matrices $\mathbf{P} = [\mathbf{I} \mid \mathbf{0}]$, $\mathbf{P}'$ and $\mathbf{P}''$ such that $\mathcal{T}_i^{jkl}$ corresponds to the three camera matrices according to (15.9-p376).

Just as with the fundamental matrix it is important to enforce these constraints in some way so as to arrive at a geometrically valid trifocal tensor. If the tensor does not satisfy the constraints, there are consequences similar to a fundamental matrix which is not of rank 2 – where epipolar lines, computed as $\mathbf{F}\mathbf{x}$ for varying $\mathbf{x}$, do not intersect in a single point (see figure 11.1(p280)). For example, if the tensor does not satisfy the internal constraints and is used to transfer a point to a third view, given a correspondence over two views as described in section 15.3, then the position of the transferred point will vary depending on which set of equations from table 16.1 is used. In the following the objective is always to estimate a geometrically valid tensor.

The constraints satisfied by the trifocal tensor elements are not so simply expressed (as $\det = 0$), and some have thought this an impediment to accurate computation of the trifocal tensor. However, in reality, in order to work with or compute the trifocal tensor it is not necessary to express these constraints explicitly – rather they are implicitly enforced by an appropriate parametrization of the trifocal tensor, and rarely cause any trouble. We will return to the issue of parametrization in section 16.3 and section 16.4.2.

16.1.2 The minimum case – 6 point correspondences

A geometrically valid trifocal tensor may be computed from images of a 6 point configuration, provided the scene points are in general position. There are one or three real solutions. The tensor is computed from the three camera matrices which are obtained using algorithm 20.1(p511), as described in section 20.2.4(p510). This minimal six point solution is used in the robust algorithm of section 16.6.

16.2 The normalized linear algorithm

In forming the matrix equation $At = 0$ from the equations on $\mathcal{T}$ in table 16.1 it is not necessary to use the complete set of equations derived from each correspondence, since not all of these equations are linearly independent. For instance in the case of a point–point–point correspondence (first row of table 16.1) all choices of s and t lead to a set of 9 equations, but only 4 of these equations are linearly independent, and these may be obtained by choosing two values for each of s and t , for instance 1 and 2. This point is discussed in more detail in section 17.7(p431).

The reader may verify that the three points equation obtained from table 16.1 for a given choice of s and t may be expanded as

$$x^k(x^i x^{tm} T_k^{jl} - x^{lj} x^{tm} T_k^{il} - x^{ti} x^{ml} T_k^{jm} + x^{lj} x^{ml} T_k^{im}) = 0^{ijlm} \quad (16.1)$$

when $i, j \neq s$ and $l, m \neq t$. Equation (16.1) collapses for $i = j$ or $l = m$, and swapping i and j (or l and m) simply changes the sign of the equation. One choice of the four independent equations is obtained by setting $j = m = 3$, and letting i and l range freely. The coordinates x^3, x'^3 and x''^3 may be set to 1 to obtain a relationship between the observed image coordinates. Equation (16.1) then becomes

$$x^k(x^i x^{tl} T_k^{33} - x^{tl} T_k^{i3} - x^{ti} T_k^{3l} + T_k^{il}) = 0. \quad (16.2)$$

The four different choices of $i, l = 1, 2$ give four different equations in terms of the observed image coordinates.

How to represent lines

The three lines correspondence equation of table 16.1 may be written in the form

$$l_i = l'_i l''_i T_i^{jk},$$

where, as usual with homogeneous entities, the equality is up to scale. In the presence of noise, this relationship will only be approximately satisfied by the *measured* lines l, l' and l'' , but will be satisfied exactly for three lines $\hat{l}, \hat{l}'$ and $\hat{l}''$ that are close to the measured lines.

The question is whether two sets of homogeneous coordinates that differ by a small amount represent lines that are close to each other in some geometric sense. Consider the two vectors $l_1 = (0.01, 0, 1)^T$ and $l_2 = (0, 0.01, 1)^T$. Clearly as vectors they are not very different, and in fact $\|l_1 - l_2\|$ is small. On the other hand, l_1 represents the line $x = 100$, and l_2 represents the line $y = 100$. Thus in a geometric sense, these lines are totally different. Note that this problem is alleviated by scaling. If coordinates are

Objective

Given $n \geq 7$ image point correspondences across 3 images, or at least 13 line correspondences, or a mixture of point and line correspondences, compute the trifocal tensor.

Algorithm

- (i) Find transformation matrices H, H' and H'' to apply to the three images.
- (ii) Transform points according to $x^i \mapsto \hat{x}^i = H_j^i x^j$, and lines according to $l_i \mapsto \hat{l}_i = (H^{-1})_i^j l_j$. Points and lines in the second and third image transform in the same way.
- (iii) Compute the trifocal tensor $\hat{\mathcal{T}}$ linearly in terms of the transformed points and lines using the equations in table 16.1 by solving a set of equation of the form $At = 0$, using algorithm A5.4(p593).
- (iv) Compute the trifocal tensor corresponding to the original data according to $\mathcal{T}_i^{jk} = H_i' (H'^{-1})_s^j (H''^{-1})_t^k \hat{\mathcal{T}}_r^{st}$.

Algorithm 16.1. *The normalized linear algorithm for computation of $\mathcal{T}$.*

scaled by a factor of 0.01, then the coordinates for the lines become $\mathbf{l}_1 = (1, 0, 1)^T$ and $\mathbf{l}_2 = (0, 1, 1)^T$, which are quite different.

Nevertheless, this observation indicates that care is needed when representing lines. Suppose one is given a correspondence between three lines l, l' and l'' . Two points $\mathbf{x}_1$ and $\mathbf{x}_2$ lying on l are selected. Each of these points provides a correspondence $\mathbf{x}_s \leftrightarrow l' \leftrightarrow l''$, for $s = 1, 2$, between the three views, in the sense that there exists a 3D line that maps to l' and l'' in the second and third images and to a line (namely l) passing through $\mathbf{x}_s$ in the first image. Two equations of the form $x_s^i l_j' l_k'' \mathcal{T}_i^{jk} = 0_s$ for $s = 1, 2$ result from these correspondences. In this way one avoids the use of lines in the first image, though not the other images. Often lines in images are defined naturally by a pair of points, possibly the two endpoints of the lines. Even lines that are defined as the best fit to a set of edge points in an image may be treated as if they were defined by just two points, as will be described in section 16.7.2.

Normalization

As in all algorithms of this type, it is necessary to carry out prenormalization of the input data before forming and solving the linear equation system. Subsequently, it is necessary to correct for this normalization to find the trifocal tensor for the original data. The recommended normalization is much the same as that given for the computation of the fundamental matrix. A translation is applied to each image such that the centroid of the points is at the origin, and then a scaling is applied so that the average (RMS) distance of the points from the origin is $\sqrt{2}$. In the case of lines, the transformation should be defined by considering each line's two endpoints (or some representative line points visible in the image). The transformation rule for the trifocal tensor under these normalizing transformations is given in section A1.2(p563). The normalized linear algorithm for computing $\mathcal{T}$ is summarized in algorithm 16.1.

This algorithm does not consider the constraints discussed in section 16.1.1 that should be applied to $\mathcal{T}$. These constraints ought to be enforced before the denormal-

ization step (final step) in the above algorithm. Methods of enforcing these constraints will be considered next.

16.3 The algebraic minimization algorithm

The linear algorithm 16.1 will give a tensor not necessarily corresponding to any geometric configuration, as discussed in section 16.1.1. The next task is to correct the tensor to satisfy all required constraints.

Our task will be to compute a geometrically valid trifocal tensor $\mathcal{T}_i^{jkl}$ from a set of image correspondences. The tensor computed will minimize the algebraic error associated with the input data. That is, we minimize $\|\mathbf{A}\mathbf{t}\|$ subject to $\|\mathbf{t}\| = 1$, where $\mathbf{t}$ is the vector of entries of a geometrically valid trifocal tensor. The algorithm is quite similar to the algebraic algorithm (section 11.3(p282)) for computation of the fundamental matrix. Just as with the fundamental matrix, the first step is the computation of the epipoles.

Retrieving the epipoles

Let $\mathbf{e}'$ and $\mathbf{e}''$ be the epipoles in the second and third images corresponding to (that is being images of) the first camera centre. Recall from result 15.4(p373) that the two epipoles $\mathbf{e}'$ and $\mathbf{e}''$ are the common perpendicular to the left (respectively right) null-vectors of the three $\mathbf{T}_i$. In principle then, the epipoles may be computed from the trifocal tensor using the algorithm outlined in algorithm 15.1(p375). However, in the presence of noise, this translates easily into an algorithm for computing the epipoles based on four applications of algorithm A5.4(p593).

- (i) For each $i = 1, \dots, 3$ find the unit vector $\mathbf{v}_i$ that minimizes $\|\mathbf{T}_i \mathbf{v}_i\|$, where $\mathbf{T}_i = \mathcal{T}_i^{::}$. Form the matrix $\mathbf{V}$, the i -th row of which is $\mathbf{v}_i^T$.
- (ii) Compute the epipole $\mathbf{e}''$ as the unit vector that minimizes $\|\mathbf{V}\mathbf{e}''\|$.

The epipole $\mathbf{e}'$ is computed similarly, using $\mathbf{T}_i^T$ instead of $\mathbf{T}_i$.

Algebraic minimization

Having computed the epipoles the next step is to determine the remaining elements of the camera matrices $\mathbf{P}', \mathbf{P}''$ from which the trifocal tensor can be calculated. This step is linear.

From the form (15.9-p376) of the trifocal tensor, it may be seen that once the epipoles $\mathbf{e}'^i = \mathbf{a}_4^i$ and $\mathbf{e}''^k = \mathbf{b}_4^k$ are known, the trifocal tensor may be expressed linearly in terms of the remaining entries of the matrices $\mathbf{a}_j^i$ and $\mathbf{b}_j^k$. This relationship may be written linearly as $\mathbf{t} = \mathbf{E}\mathbf{a}$ where $\mathbf{a}$ is the vector of the remaining entries a_j^i and b_j^k , $\mathbf{t}$ is the vector of entries of the trifocal tensor, and $\mathbf{E}$ is the linear relationship expressed by (15.9-p376). We wish to minimize the algebraic error $\|\mathbf{A}\mathbf{t}\| = \|\mathbf{A}\mathbf{E}\mathbf{a}\|$ over all choices of $\mathbf{a}$ constrained such that $\|\mathbf{t}\| = 1$, that is $\|\mathbf{E}\mathbf{a}\| = 1$. This minimization problem is solved by algorithm A5.6(p595). The solution $\mathbf{t} = \mathbf{E}\mathbf{a}$ represents a trifocal tensor satisfying all constraints, and minimizing the algebraic error, subject to the given choice of epipoles.

Objective

Given a set of point and line correspondences in three views, compute the trifocal tensor.

Algorithm

- (i) From the set of point and line correspondences compute the set of equations of the form $At = 0$, from the relations given in table 16.1.
- (ii) Solve these equations using algorithm A5.4(p593) to find an initial estimate of the trifocal tensor T_i^{jk} .
- (iii) Find the two epipoles e' and e'' from T_i^{jk} as the common perpendicular to the left (respectively right) null-vectors of the three T_i .
- (iv) Construct the 27×18 matrix E such that $t = Ea$ where t is the vector of entries of T_i^{jk} , a is the vector representing entries of a_i^j and b_i^k , and where E expresses the linear relationship $T_i^{jk} = a_i^j e''^k - e'^j b_i^k$.
- (v) Solve the minimization problem: minimize $\|AEa\|$ subject to $\|Ea\| = 1$, using algorithm A5.6(p595). Compute the error vector $\epsilon = AEa$.
- (vi) **Iteration:** The mapping $(e', e'') \mapsto \epsilon$ is a mapping from $\mathbb{R}^6$ to $\mathbb{R}^{27}$. Iterate on the last two steps with varying e' and e'' using the Levenberg–Marquardt algorithm to find the optimal e', e'' . Hence find the optimal $t = Ea$ containing the entries of T_i^{jk} .

Algorithm 16.2. Computing the trifocal tensor minimizing algebraic error. *The computation should be carried out on data normalized in the manner of algorithm 16.1. Normalization and denormalization steps are omitted here for simplicity. This algorithm finds the geometrically valid trifocal tensor that minimizes algebraic error. At the cost of a slightly inferior solution, the last iteration step may be omitted, providing a fast non-iterative algorithm.*

Iterative method

The two epipoles used to compute a geometrically valid tensor T_i^{jk} are determined using the estimate of T_i^{jk} obtained from the linear algorithm. Analogous to the case of the fundamental matrix, the mapping $(e', e'') \mapsto AEa$ is a mapping $\mathbb{R}^6 \rightarrow \mathbb{R}^{27}$. An application of the Levenberg–Marquardt algorithm to optimize the choice of the epipoles will result in an optimal (in terms of algebraic error) estimate of the trifocal tensor. Note that the iteration problem is of modest size, since only 6 parameters, the homogeneous coordinates of the epipoles, are involved in the iteration problem.

This contrasts with an iterative estimation of the optimal trifocal tensor in terms of geometric error, considered later. This latter problem requires estimating the parameters of the three cameras, plus the coordinates of all the points, a large estimation problem.

The complete algebraic method for estimating the trifocal tensor is summarized in algorithm 16.2.

16.4 Geometric distance**16.4.1 The Gold Standard method for the trifocal tensor**

As with the computation of the fundamental matrix, best results may be expected from the maximum likelihood (or “Gold Standard”) solution. Since this has been adequately described for the case of the fundamental matrix computation, little needs to be added for the three-view case.

Objective

Given $n \geq 7$ image point correspondences $\{\mathbf{x}_i \leftrightarrow \mathbf{x}'_i \leftrightarrow \mathbf{x}''_i\}$, determine the Maximum Likelihood Estimate of the trifocal tensor.

The MLE involves also solving for a set of subsidiary point correspondences $\{\hat{\mathbf{x}}_i \leftrightarrow \hat{\mathbf{x}}'_i \leftrightarrow \hat{\mathbf{x}}''_i\}$, which exactly satisfy the trilinear relations of the estimated tensor and which minimize

$$\sum_i d(\mathbf{x}_i, \hat{\mathbf{x}}_i)^2 + d(\mathbf{x}'_i, \hat{\mathbf{x}}'_i)^2 + d(\mathbf{x}''_i, \hat{\mathbf{x}}''_i)^2$$

Algorithm

- (i) Compute an initial geometrically valid estimate of $\mathcal{T}$ using a linear algorithm such as algorithm 16.2.
- (ii) Compute an initial estimate of the subsidiary variables $\{\hat{\mathbf{x}}_i, \hat{\mathbf{x}}'_i, \hat{\mathbf{x}}''_i\}$ as follows:
 - (a) Retrieve the camera matrices P' and P'' from $\mathcal{T}$.
 - (b) From the correspondence $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i \leftrightarrow \mathbf{x}''_i$ and $P = [I \mid 0]$, P' , P'' determine an estimate of $\hat{\mathbf{X}}_i$ using the triangulation method of chapter 12.
 - (c) The correspondence consistent with $\mathcal{T}$ is obtained as
 $\hat{\mathbf{x}}_i = P\hat{\mathbf{X}}_i$, $\hat{\mathbf{x}}'_i = P'\hat{\mathbf{X}}_i$, $\hat{\mathbf{x}}''_i = P''\hat{\mathbf{X}}_i$.
- (iii) Minimize the cost

$$\sum_i d(\mathbf{x}_i, \hat{\mathbf{x}}_i)^2 + d(\mathbf{x}'_i, \hat{\mathbf{x}}'_i)^2 + d(\mathbf{x}''_i, \hat{\mathbf{x}}''_i)^2$$

over $\mathcal{T}$ and $\hat{\mathbf{X}}_i, i = 1, \dots, n$. The cost is minimized using the Levenberg–Marquardt algorithm over $3n + 24$ variables: $3n$ for the n 3D points $\hat{\mathbf{X}}_i$, and 24 for the elements of the camera matrices P', P'' .

Algorithm 16.3. *The Gold Standard algorithm for estimating $\mathcal{T}$ from image correspondences.*

Given a set of point correspondences $\{\mathbf{x}_i \leftrightarrow \mathbf{x}'_i \leftrightarrow \mathbf{x}''_i\}$ in three views, the cost function to be minimized is

$$\sum_i d(\mathbf{x}_i, \hat{\mathbf{x}}_i)^2 + d(\mathbf{x}'_i, \hat{\mathbf{x}}'_i)^2 + d(\mathbf{x}''_i, \hat{\mathbf{x}}''_i)^2 \quad (16.3)$$

where the points $\hat{\mathbf{x}}_i, \hat{\mathbf{x}}'_i, \hat{\mathbf{x}}''_i$ satisfy a trifocal constraint (as in table 16.1) exactly for the estimated trifocal tensor. As in the case of the fundamental matrix one needs to introduce further variables corresponding to 3D points $\mathbf{X}_i$ and parametrize the trifocal tensor by the entries of the matrices P' and P'' (see below). The cost function is then minimized over the position of the 3D points $\mathbf{X}_i$ and the two camera matrices P' and P'' with $\hat{\mathbf{x}}_i = [I \mid 0]\mathbf{X}_i$, $\hat{\mathbf{x}}'_i = P'\mathbf{X}_i$, and $\hat{\mathbf{x}}''_i = P''\mathbf{X}_i$. Essentially one is carrying out bundle adjustment over three views. The sparse matrix techniques of section A6.3-(p602) should be used.

A good way to find an initial estimate is the algebraic algorithm 16.2, though the final iterative step can be omitted. This algorithm gives a direct estimate of the entries of P' and P'' . The initial estimate of the 3D points $\mathbf{X}_i$ may be obtained using the linear triangulation method of section 12.2(p312). The steps of the algorithm are summarized in algorithm 16.3.

The technique can be extended to include line correspondences. To do this, one needs to find a representation of a 3D line convenient for computation. Given a 3-view line correspondence $l \leftrightarrow l' \leftrightarrow l''$, the lines being perhaps defined by their endpoints in each image, a very convenient way to represent the 3D line during the LM parameter minimization is by its projections $\hat{l}'$ and $\hat{l}''$ in the second and third views. Given a candidate trifocal tensor, one can easily compute the projection of the 3D line into the first view using the line transfer equation $\tilde{l}_i = \hat{l}_j' \hat{l}_k'' T_i^{jk}$. Then one minimizes the sum-of-squares line distance

$$\sum_i d(l_i, \hat{l}_i)^2 + d(l'_i, \hat{l}'_i)^2 + d(l''_i, \hat{l}''_i)^2$$

for some appropriate interpretation of the distance $d(l'_i, \hat{l}'_i)^2$ between the measured and estimated line. If the measured line is specified by its endpoints, then the obvious distance metric to use is the distance of the estimated line from the measured endpoints. In general a Mahalanobis distance may be used.

16.4.2 Parametrization of the trifocal tensor

If the tensor is parametrized simply by its 27 entries, then the estimated tensor will not satisfy the internal constraints. A parametrization which ensures that the tensor does satisfy its constraints, and so is geometrically valid, is termed *consistent*.

Since, from definition 16.1, a tensor is geometrically valid if it is generated from three camera matrices $P = [I \mid 0]$, P' , P'' by (15.9-p376), it follows that the three camera matrices give a consistent parametrization. Note that this is an *overparametrization* since it requires 24 parameters to be specified, namely the 12 entries each of the matrices $P' = [A|a_4]$ and $P'' = [B|b_4]$. There is no need to attempt to define a minimal set of (18) parameters, which is a difficult task. Any choice of cameras is a consistent parametrization, the particular projective reconstruction has no effect on the tensor.

Another consistent parametrization is obtained by computing the tensor from six point correspondences across three views as in section 20.2(p508). Then the position of the points in each image is the parametrization – a total of 6 (points) $\times$ 2 (for x, y) $\times$ 3 (images) = 36 parameters. However, only a subset of the points need be varied during the minimization, or the movement of the points can be restricted to be perpendicular to the variety of trifocal tensors.

16.4.3 First-order geometric error (Sampson distance)

The trifocal tensor may be computed using a geometric cost function based on the Sampson approximation in a manner entirely analogous to the Sampson method used to compute the fundamental matrix (section 11.4.3(p287)). Again the advantage is that it is not necessary to introduce a set of subsidiary variables, as this first-order geometric error requires a minimization only over the parametrization of the tensor (e.g. only 24 parameters if P', P'' is used as above). The minimization can be carried out with a simple iterative Levenberg–Marquardt algorithm, and the method initialized by the iterative algebraic algorithm 16.2.

The Sampson cost function is a little more complex computationally than the corresponding cost function for the fundamental matrix (11.9–p287), because each point correspondence gives four equations, instead of just one for the fundamental matrix. The more general case was discussed in section 4.2.6(p98). The error function (4.13–p100) in the present case is

$$\sum_i \epsilon_i^T (J_i J_i^T)^{-1} \epsilon_i \quad (16.4)$$

where ϵ_i is the algebraic error vector $A_i t$ corresponding to a single 3-view correspondence (a 4-vector in the case of 4 equations per point), and J is the 4×6 matrix of partial derivatives of ϵ with respect to the coordinates of each of the corresponding points $x_i \leftrightarrow x'_i \leftrightarrow x''_i$. As in the programming hint given in exercise (vii) on page 129, the computation of the partial derivative matrix J may be simplified by observing that the cost function is multilinear in the coordinates of the points x_i, x'_i, x''_i .

The Sampson error method has various advantages:

- It gives a good approximation to actual geometric error (the optimum), using a relatively simple iterative algorithm.
- As in the case of actual geometric error, non-isotropic and unequal error distributions may be specified for each of the points without significantly complicating the algorithm. See exercises in chapter 4.

16.5 Experimental evaluation of the algorithms

A brief comparison is now given of the results of the (iterative) algebraic algorithm 16.2 along with the Gold Standard algorithm 16.3 for computing the trifocal tensor. The algorithms are run on synthetic data with controlled levels of noise. This allows a comparison with the theoretically optimal ML results, and a determination of how well these algorithms are able to approximate the theoretical lower bound on residual error, achieved by an optimal ML algorithm.

Computer-generated data sets of 10, 15 and 20 points were used to test the algorithm, and the cameras were placed at random angles around the cloud of points. The camera parameters were chosen to approximate a standard 35mm camera, and the scale was chosen so that the size of the image was 600×600 pixels.

For a given level of added Gaussian noise in the image measurement, one may compute the expected residual achieved by an ML algorithm, according to result 5.2(p136). In this case, if n is the number of points, then the number of measurements is $N = 6n$, and the number of degrees of freedom in the fitting is $d = 18 + 3n$, where 18 represents the number of degrees of freedom of the three cameras (3×11 less 15 to account for projective ambiguity) and $3n$ represents the number of degrees of freedom of n points in space. Hence the ML residual is

$$\epsilon_{\text{res}} = \sigma(1 - d/N)^{1/2} = \sigma \left(\frac{n-6}{2n} \right)^{1/2}.$$




Fig. 16.1. **Comparison of trifocal tensor estimation algorithms.** The residual error RMS-averaged over 100 runs is plotted against the noise level, for computation of the trifocal tensor using 10, 15 and 20 points. Each graph contains three curves. The top curve is the result of the algebraic error minimization, whereas the lower two curves, actually indistinguishable in the graphs, represent the theoretical minimum error, and the error obtained by the Gold Standard algorithm using the algebraic minimization as a starting point. Note that the residual errors are almost exactly proportional to added noise, as they should be.

16.5.1 Results and recommendations

The results are shown in figure 16.1. We learn two things from these results. Minimization of the algebraic error achieves residual errors within about 15% of the optimal and using this estimate as a starting point for minimizing geometric error achieves a virtually optimal estimate.

All the algorithms developed above, except the linear method of section 16.1, enforce the internal constraints on the tensor. The linear method is not recommended for use on its own, but is necessary for initialization in most of the other methods. As in the case of estimating the fundamental matrix our recommendations are to use the iterative algebraic algorithm 16.2 or the Sampson geometric approximation of section 16.4.3. Both give excellent results. Again to be certain of getting the best results, if Gaussian noise is a viable assumption, implement the Gold Standard algorithm 16.3.

16.6 Automatic computation of $\mathcal{T}$

This section describes an algorithm to compute the trifocal geometry between three images automatically. The input to the algorithm is simply the triplet of images, with no other *a priori* information required; and the output is the estimated trifocal tensor together with a set of interest points in correspondence across the three images.

The fact that the trifocal tensor may be used to determine the exact image position of a point in a third view, given its image position in the other two views, means that there are fewer mismatches over three views than there are over two. In the two view case there is only the weaker geometric constraint of an epipolar line against which to verify a possible match.

The three-view algorithm uses RANSAC as a search engine in a similar manner to its use in the automatic computation of a homography described in section 4.8- (p123). The ideas and details of the algorithm are given there, and are not repeated here. The method is summarized in algorithm 16.4, with an example of its use shown

Objective Compute the trifocal tensor between three images.

Algorithm

- (i) **Interest points:** Compute interest points in each image.
 - (ii) **Two-view correspondences:** Compute interest point correspondences (and F) between views 1 & 2, and 2 & 3 using algorithm 11.4(p291).
 - (iii) **Putative three-view correspondences:** Compute a set of interest point correspondences over three views by joining the two-view match sets.
 - (iv) **RANSAC robust estimation:** Repeat for N samples, where N is determined adaptively as in algorithm 4.5(p121):
 - (a) Select a random sample of 6 correspondences and compute the trifocal tensor using algorithm 20.1(p511). There will be one or three real solutions.
 - (b) Calculate the distance $d_{\perp}$ in $\mathbb{R}^6$ from each putative correspondence to the variety described by $\mathcal{T}$, as in section 16.6.
 - (c) Compute the number of inliers consistent with $\mathcal{T}$ by the number of correspondences for which $d_{\perp} < t$.
 - (d) If there are three real solutions for $\mathcal{T}$ the number of inliers is computed for each solution, and the solution with most inliers retained.
- Choose the $\mathcal{T}$ with the largest number of inliers. In the case of ties choose the solution that has the lowest standard deviation of inliers.
- (v) **Optimal estimation:** Re-estimate $\mathcal{T}$ from all correspondences classified as inliers using the Gold Standard algorithm 16.3 or the Sampson approximation to this.
 - (vi) **Guided matching:** Further interest point correspondences are now determined using the estimated $\mathcal{T}$ as described in the text.

The last two steps can be iterated until the number of correspondences is stable.

Algorithm 16.4. Algorithm to automatically estimate the trifocal tensor over three images using RANSAC.

in figure 16.2, and additional explanation of the steps given below. Figure 16.3 shows a second example which includes automatically computed line matches.

The distance measure – reprojection error. Given the match $\mathbf{x} \leftrightarrow \mathbf{x}' \leftrightarrow \mathbf{x}''$ and the current estimate of $\mathcal{T}$ we need to determine the minimum of the reprojection error – $d_{\perp}^2 = d^2(\mathbf{x}, \hat{\mathbf{x}}) + d^2(\mathbf{x}', \hat{\mathbf{x}}') + d^2(\mathbf{x}'', \hat{\mathbf{x}}'')$, where the image points $\hat{\mathbf{x}}, \hat{\mathbf{x}}', \hat{\mathbf{x}}''$ are consistent with $\mathcal{T}$. As usual the consistent images points may be obtained from the projection of a 3-space point $\hat{\mathbf{X}}$

$$\hat{\mathbf{x}} = [\mathbf{I} \mid \mathbf{0}] \hat{\mathbf{X}}, \quad \hat{\mathbf{x}}' = \mathbf{P}' \hat{\mathbf{X}}, \quad \hat{\mathbf{x}}'' = \mathbf{P}'' \hat{\mathbf{X}}$$

where the camera matrices $\mathbf{P}', \mathbf{P}''$ are extracted from $\mathcal{T}$. The distance $d_{\perp}^2$ is then obtained by determining the point $\hat{\mathbf{X}}$ which minimizes the image distance between the measured points $\mathbf{x}, \mathbf{x}', \mathbf{x}''$ and the projected points.

Another way of obtaining this distance is to use the Sampson error (16.4), which is a first-order approximation to the geometric error. However, in practice it is quicker to estimate the error directly by non-linear least-squares iteration (a small Levenberg–Marquardt problem). Starting from an initial estimate of $\hat{\mathbf{X}}$, one iterates varying the coordinates of $\hat{\mathbf{X}}$ to minimize the reprojection error.




Fig. 16.2. Automatic computation of the trifocal tensor between three images using RANSAC. (a - c) raw images of Keble College, Oxford. The motion between views consists of a translation and rotation. The images are 640×480 pixels. (d - f) detected corners superimposed on the images. There are approximately 500 corners on each image. The following results are superimposed on the (a) image: (g) 106 putative matches shown by the line linking corners, note the clear mismatches; (h) outliers - 18 of the putative matches. (i) inliers - 88 correspondences consistent with the estimated $\mathcal{T}$; (j) final set of 95 correspondences after guided matching and MLE. There are no mismatches.

Guided matching. We have an initial estimate of $\mathcal{T}$ and wish to use this to generate and assess additional point correspondences across the three-views. The first step is to extract the fundamental matrix F_{12} between views 1 & 2 from $\mathcal{T}$. Then two-view



Fig. 16.3. **Image triplet matching.** The trifocal tensor is computed automatically from interest points using algorithm 16.4, and subsequently used to match line features across views. (a) Three images of a corridor sequence. (b) Automatically matched line segments. The matching algorithm is described in [Schmid-97].

guided matches are computed using loose thresholds on matching. Each two-view match is corrected using F_{12} to give points $\hat{x}, \hat{x}'$ which are consistent with F_{12} . These corrected two-view matches (together with T) define a small search window in the third view in which the corresponding point is sought. Any three-view point correspondence is assessed by computing $d_{\perp}$, as described above. The match is accepted if $d_{\perp}$ is less than the threshold t . Note, the same threshold is used for inlier detection within the RANSAC stage and guided matching.

In practice it is found that the stage of guided matching is more significant here, in that it generates additional correspondences, than in the case of homography estimation.

Implementation and run details. For the example of figure 16.2, the search window was ± 300 pixels. The inlier threshold was $t = 1.25$ pixels. A total of 26 samples were required. The RMS pixel error after RANSAC was 0.43 (for 88 correspondences), after MLE it was 0.23 (for 88 correspondences), and after MLE and guided matching it was 0.19 (for 95 correspondences). The MLE required 10 iterations of the Levenberg–Marquardt algorithm.

Note, RANSAC has to do far less work than in algorithm 11.4(p291) to estimate F and correspondences, because the two-view algorithm has already removed many outliers before the putative correspondences over three views are generated.

16.7 Special cases of $\mathcal{T}$ -computation

16.7.1 Computing $\mathcal{T}_i^{jk}$ from a plane plus parallax

We describe here the computation of $\mathcal{T}_i^{jk}$ from the image of a special configuration consisting of a world plane (from which a homography between views can be computed) and two points off the plane. Of course, it is not necessary for the plane to actually be present. It may be virtual, or the homography may simply be specified by the images of four coplanar points or four coplanar lines. The method is the analogue of algorithm 13.2(p336) for the fundamental matrix.

The solution is obtained by constructing the three camera matrices (up to a common projective transformation of 3-space) and then computing the trifocal tensor from these matrices according to (15.9-p376). The homography induced by the world (reference) plane between the first and second view is H_{12} , and between the first and third views is H_{13} . As shown in section 13.3(p334) the epipole e' may be computed directly from the two point correspondences off the plane for the first and second views, and the camera matrices chosen as $P = [I \mid 0]$, $P' = [H_{12} \mid \mu e']$, where μ is a scalar. Note the scale of both H_{12} and e' is considered fixed here, so they are no longer homogeneous quantities. Similarly, e'' may be determined from the two point correspondences for views one and three and the camera matrices chosen as $P = [I \mid 0]$, $P'' = [H_{13} \mid \lambda e'']$, where λ is a scalar.

It is then easily verified that a consistent set of cameras for the three views (see the discussion on consistent camera triplets on page 375) is given by

$$P = [I \mid 0], \quad P' = [H_{12} \mid e'], \quad P'' = [H_{13} \mid \lambda e''] \quad (16.5)$$

where μ has been set to unity. The value of λ is determined from one of the point correspondences over three views, and this is left as an exercise. For more on plane-parallax reconstruction, see section 18.5.2(p450).

Note that the estimation of the trifocal tensor for this configuration is over-determined. In the case of the fundamental matrix over two views the homographies determine all but 2 degrees of freedom (the epipole), and each of the point correspondence provides one constraint, so that the number of constraints equals the number of degrees of freedom of the matrix. In the case of the trifocal tensor the homography determines all but 5 degrees of freedom (the two epipoles and their relative scaling). However, each point correspondence provides three constraints (there are six coordinate measurements less three for the point's position in 3-space), so that there are six constraints on 5 degrees of freedom. Since there are more measurements than degrees of freedom in this case, the tensor should be estimated by minimizing a cost function based on geometric error.

16.7.2 Lines specified by several points

In describing the reconstruction algorithm from lines, we have considered the case where lines are specified by their two endpoints. Another common way that lines may

be specified in an image is as the best line fit to several points. It will be shown now how that case may easily be reduced to the case of a line defined by two endpoints. Consider a set of points $\mathbf{x}_i$ in an image, normalized to have third component equal to 1. Let $\mathbf{l} = (l_1, l_2, l_3)^T$ be a line, which we suppose is normalized such that $l_1^2 + l_2^2 = 1$. In this case, the distance from a point $\mathbf{x}_i$ to the line $\mathbf{l}$ is equal to $\mathbf{x}_i^T \mathbf{l}$. The squared distance may be written as $d^2 = \mathbf{l}^T \mathbf{x}_i \mathbf{x}_i^T \mathbf{l}$, and the sum-of-squares of all distances is

$$\sum_i \mathbf{l}^T \mathbf{x}_i \mathbf{x}_i^T \mathbf{l} = \mathbf{l}^T \left(\sum_i \mathbf{x}_i \mathbf{x}_i^T \right) \mathbf{l}.$$

The matrix $\mathbf{E} = (\sum_i \mathbf{x}_i \mathbf{x}_i^T)$ is positive-definite and symmetric.

Lemma 16.2. *Matrix $(\mathbf{E} - \epsilon_0 \mathbf{J})$ is positive-semidefinite, where $\mathbf{J}$ is the matrix $\text{diag}(1, 1, 0)$ and ϵ_0 is the smallest solution to the equation $\det(\mathbf{E} - \epsilon \mathbf{J}) = 0$.*

Proof. We start by computing the vector $\mathbf{x} = (x_1, x_2, x_3)^T$ that minimizes $\mathbf{x}^T \mathbf{E} \mathbf{x}$ subject to the condition $x_1^2 + x_2^2 = 1$. Using the method of Lagrange multipliers, this comes down to finding the extrema of $\mathbf{x}^T \mathbf{E} \mathbf{x} - \xi(x_1^2 + x_2^2)$, where ξ denotes the Lagrange coefficient. Taking the derivative with respect to $\mathbf{x}$ and setting it to zero, we find that $2\mathbf{E}\mathbf{x} - \xi(2x_1, 2x_2, 0)^T = \mathbf{0}$. This may be written as $(\mathbf{E} - \xi \mathbf{J})\mathbf{x} = \mathbf{0}$. It follows that ξ is a root of the equation $\det(\mathbf{E} - \xi \mathbf{J}) = 0$ and $\mathbf{x}$ is the generator of the null-space of $\mathbf{E} - \xi \mathbf{J}$. Since $\mathbf{x}^T \mathbf{E} \mathbf{x} = \xi \mathbf{x}^T \mathbf{J} \mathbf{x} = \xi(x_1^2 + x_2^2) = \xi$, it follows that to minimize $\mathbf{x}^T \mathbf{E} \mathbf{x}$ one must choose ξ to be the minimum root ξ_0 of the equation $\det(\mathbf{E} - \xi \mathbf{J}) = 0$. In this case one has $\mathbf{x}_0^T \mathbf{E} \mathbf{x}_0 - \xi_0 = 0$ for the minimizing vector $\mathbf{x}_0$. For any other vector $\mathbf{x}$, not necessarily the minimizing vector, one has $\mathbf{x}^T \mathbf{E} \mathbf{x} - \xi_0 \geq 0$. Then, $\mathbf{x}^T (\mathbf{E} - \xi_0 \mathbf{J}) \mathbf{x} = \mathbf{x}^T \mathbf{E} \mathbf{x} - \xi_0 \geq 0$, and so $\mathbf{E} - \xi_0 \mathbf{J}$ is positive-semidefinite. $\square$

Since the matrix $\mathbf{E} - \xi_0 \mathbf{J}$ is symmetric it may be written in the form $\mathbf{E} - \xi_0 \mathbf{J} = \mathbf{V} \text{diag}(r, s, 0) \mathbf{V}^T$ where $\mathbf{V}$ is an orthogonal matrix and r and s are positive. It follows that

$$\begin{aligned} \mathbf{E} - \xi_0 \mathbf{J} &= \mathbf{V} \text{diag}(r, 0, 0) \mathbf{V}^T + \mathbf{V} \text{diag}(0, s, 0) \mathbf{V}^T \\ &= r \mathbf{v}_1 \mathbf{v}_1^T + s \mathbf{v}_2 \mathbf{v}_2^T \end{aligned}$$

where $\mathbf{v}_i$ is the i -th column of $\mathbf{V}$. Therefore $\mathbf{E} = \xi_0 \mathbf{J} + r \mathbf{v}_1 \mathbf{v}_1^T + s \mathbf{v}_2 \mathbf{v}_2^T$. Then for any line $\mathbf{l}$ satisfying $l_1^2 + l_2^2 = 1$ we have

$$\begin{aligned} \sum_i (\mathbf{x}_i^T \mathbf{l})^2 &= \mathbf{l}^T \mathbf{E} \mathbf{l} \\ &= \xi_0 + r(\mathbf{v}_1^T \mathbf{l})^2 + s(\mathbf{v}_2^T \mathbf{l})^2. \end{aligned}$$

Thus, we have replaced the sum-of-squares of several points by a constant value ξ_0 , which is not capable of being minimized, plus the weighted sum-of-squares of the distances to two points $\mathbf{v}_1$ and $\mathbf{v}_2$. To summarize: when forming the trifocal tensor equations involving a line defined by points $\mathbf{x}_i$, formulate two point equations expressed in terms of the points $\mathbf{v}_1$ and $\mathbf{v}_2$ with weights $\sqrt{r}$ and $\sqrt{s}$ respectively.

Orthogonal regression. In the proof of lemma 16.2 above, it was shown that the line l that minimizes the sum of squared distances to the set of all points $\mathbf{x}_i = (x_i, y_i, 1)^T$ is obtained as follows.

- (i) Define matrices $E = \sum_i \mathbf{x}_i \mathbf{x}_i^T$ and $J = \text{diag}(1, 1, 0)$.
- (ii) Let ξ_0 be the minimum root of the equation $\det(E - \xi J) = 0$.
- (iii) The required line l is the right null-vector of the matrix $E - \xi_0 J$.

This gives a least-squares best fit of a line to a set of points. This process is known as *orthogonal regression* and it extends in an obvious way to higher-dimensional fitting of a hyperplane to a set of points in a way that minimizes the sum of squared distances to the points.

16.8 Closure

16.8.1 The literature

A linear method for computing the trifocal tensor was first given in [Hartley-97a], where further experimental results of estimation using both point and line correspondences on real data are reported. An iterative algebraic method for estimating a consistent tensor was given in [Hartley-98d].

Torr and Zisserman [Torr-97] developed an automatic algorithm for estimating a consistent tensor $\mathcal{T}$ from three images. This paper also compared several parametrizations of the iterative minimization. Several methods of representing and imposing the constraints on the tensor are given by Faugeras and Papadopoulos [Faugeras-97].

[Oskarsson-02] gives minimal solutions for reconstruction for the two cases of “four points and three lines in three views”, and “two points and six lines in three views”.

16.8.2 Notes and exercises

- (i) Consider the problem of estimating the 3-space point $\mathbf{X}$ which minimizes re-projection error from measured image points $\mathbf{x}, \mathbf{x}', \mathbf{x}''$, given the trifocal tensor. This is the analogue of the triangulation problem of chapter 12. Show that for general motion the one parameter family parametrization of epipolar lines developed in chapter 12 does not extend from two views to three. However, in the case that the three camera centres are collinear the two-view parametrization can be extended to three and a minimum determined by solving a polynomial in one variable. What is the degree of this polynomial?
- (ii) An affine trifocal tensor may be computed from a minimal configuration of 4 points in general position. The computation is similar to that of algorithm 14.2-(p352), and the resulting tensor satisfies the internal constraints for an affine trifocal tensor. How many constraints are there in the affine case?
If more than 4 point correspondences are used in the estimation then a geometrically valid tensor is estimated using the factorization algorithm of section 18.2(p436).
- (iii) The transformation rule for tensors is $\mathcal{T}_i^{jkc} = A_i^t (B^{-1})_s^j (C^{-1})_l^k \tilde{\mathcal{T}}_t^{slc}$. This may be computed easily as

```

Binv = B.inverse();
Cinv = C.inverse();

for (i=1; i<=3; i++) for (j=1; j<=3; j++) for (k=1; k<=3; k++)
{
    T[i][j][k] = 0.0;

    for (r=1; r<=3; r++) for (s=1; s<=3; s++) for (t=1; t<=3; t++)
        T[i][j][k] +=
            A[r][i] * Binv[j][s] * Cinv[k][t] * T_hat[r][s][t];
}

```

How many multiplications and loop iterations does this involve? Find a better way of computing this transformation.

- (iv) In the computation of the trifocal tensor using plane plus parallax (section 16.7.1), show that if ρ is the projective depth of one of the points off the plane (i.e. $\mathbf{x}' = H_{12}\mathbf{x} + \rho\mathbf{e}'$ see (13.9-p337)), then the scalar λ in (16.5) may be computed from the equation $\mathbf{x}'' = H_{13}\mathbf{x} + \rho\lambda\mathbf{e}''$.

Part IV

N-View Geometry



Untitled 1947, (Oil on sackcloth) by Asger Jorn (1914-1973)
© 2003 Artists Rights Society (ARS), New York / COPY-DAN, Copenhagen

Outline

This part is partly a recapitulation and partly new material.

Chapter 17 is the recapitulation. We return to two- and three-view geometry but now within a more general framework which naturally extends to four- and n -views. The fundamental projective relations over multiple views arise from the intersection of lines (back-projected from points) and planes (back-projected from lines). These intersection properties are represented by the vanishing of determinants formed from the camera matrices of the views. The fundamental matrix, the trifocal tensor, and a new tensor for four views – the *quadrifocal tensor* – arise naturally from these determinants as the multiple view tensors for two, three, and four views respectively. The tensors are what remains when the 3D structure and non-essential part of the camera matrices are eliminated. The tensors stop at four views.

These tensors are unique for each set of views, and generate relationships which are multi-linear in the coordinates of the image measurements. The tensors can be computed from sets of image correspondences, and subsequently a camera matrix for each view can be computed from the tensor. Finally, the 3D structure can be computed from the retrieved cameras and image correspondences.

Chapter 18 covers the computation of a reconstruction from multiple views. In particular the important factorization algorithm is given for reconstruction from affine views. It is important because the algorithm is optimal, but is also non-iterative.

Chapter 19 describes the auto-calibration of a camera. These are a set of methods for computing the internal parameters of a camera based on constraints over multiple images. In contrast to the traditional approach to calibration described in chapter 7, no explicit scene calibration object is used, but simply constraints such as that the internal parameters are common across the images, or that the camera rotates about its centre and does not change aspect ratio.

Chapter 20 emphasises the duality between points and cameras, and how this links various configurations and algorithms that have been given throughout this book. This chapter contains an algorithm for computing a reconstruction of six points imaged in 3-views.

Chapter 21 investigates the issue of whether points are in front of or behind one or more cameras. This is an issue that goes beyond the homogeneous representation used throughout the book which does not distinguish the direction of a ray.

Chapter 22 covers the important topic of those configurations for which the estimation algorithms described in this book will fail. An example is for resectioning, where the camera matrix cannot be computed if all the 3D points and the camera centre lie on a twisted cubic.

N-Linearities and Multiple View Tensors

This chapter introduces the quadrifocal tensor Q^{ijkl} between four views, which is the analogue of the fundamental matrix for two and the trifocal tensor for three views. The quadrifocal tensor encapsulates the relationships between imaged points and lines seen in four views.

It is shown that multiple view relations may be derived directly and uniformly from the intersection properties of back-projected lines and points. From this analysis the fundamental matrix F , trifocal tensor T_i^{jk} , and quadrifocal tensor Q^{ijkl} appear in a common framework involving matrix determinants. Specific formulae are given for each of these tensors in terms of the camera matrices.

We also develop general counting arguments for the degrees of freedom of the tensors and the number of point and line correspondences required for tensor computation. These are given for configurations in general position and for the important special case where four or more of the elements are coplanar.

17.1 Bilinear relations

We consider first the relationship that holds between the coordinates of a point seen in two separate views. Thus, let $\mathbf{x} \leftrightarrow \mathbf{x}'$ be a pair of corresponding points which are the images of the same point $\mathbf{X}$ in space as seen in the two separate views. It will be convenient, for clarity of notation, to represent the two camera matrices by A and B , instead of the usual notation, P and P' . The projection from space to image can now be expressed as $k\mathbf{x} = A\mathbf{X}$ and $k'\mathbf{x}' = B\mathbf{X}$ where k and k' are two undetermined constants. This pair of equations may be written down as one equation

$$\begin{bmatrix} A & \mathbf{x} & 0 \\ B & 0 & \mathbf{x}' \end{bmatrix} \begin{pmatrix} \mathbf{X} \\ -k \\ -k' \end{pmatrix} = 0$$

and it may easily be verified that this is equivalent to the two equations. This can be written in a more detailed form by denoting the i -th row of the matrix A by $\mathbf{a}^i$, and similarly the i -th row of the matrix B by $\mathbf{b}^i$. We also write $\mathbf{x} = (x^1, x^2, x^3)^T$ and

$\mathbf{x}' = (x'^1, x'^2, x'^3)^\top$. The set of equations is now

$$\left[\begin{array}{cc|cc} \mathbf{a}^1 & x^1 & & \\ \mathbf{a}^2 & x^2 & & \\ \mathbf{a}^3 & x^3 & & \\ \hline \mathbf{b}^1 & & x'^1 & \\ \mathbf{b}^2 & & x'^2 & \\ \mathbf{b}^3 & & x'^3 & \end{array} \right] \begin{pmatrix} \mathbf{X} \\ -k \\ -k' \end{pmatrix} = 0. \quad (17.1)$$

Now, this is a 6×6 set of equations which by hypothesis has a non-zero solution, the vector $(\mathbf{X}^\top, -k, -k')^\top$. It follows that the matrix of coefficients in (17.1) must have zero determinant. It will be seen that this condition leads to a bilinear relationship between the entries of the vectors $\mathbf{x}$ and $\mathbf{x}'$ expressed by the fundamental matrix F . We will now look specifically at the form of this relationship.

Consider the matrix appearing in (17.1). Denote it by X . The determinant of X may be written as an expression in terms of the quantities x^i and x'^i . Notice that the entries x^i and x'^i appear in only two columns of X . This implies that the determinant of X may be expressed as a quadratic expression in terms of the x^i and x'^i . In fact, since all the entries x^i appear in the same column, there can be no terms of the form $x^i x^j$ or $x'^i x'^j$. Briefly, in terms of the x^i and x'^i , the determinant of X is a bilinear expression. The fact that the determinant is zero may be written as an equation

$$(x'^1, x'^2, x'^3) F (x^1, x^2, x^3)^\top = x'^i x^j F_{ji} = 0 \quad (17.2)$$

where F is a 3×3 matrix, the fundamental matrix.

We may compute a specific formula for the entries of the matrix F as follows. The entry F_{ij} of F is the coefficient of the term $x'^i x^j$ in the expansion of the determinant of X . In order to find this coefficient, we must eliminate the rows and columns of the matrix containing x'^i and x^j , take the determinant of the resulting matrix and multiply by ± 1 as appropriate. For instance, the coefficient of $x'^1 x^1$ is obtained by eliminating two rows and the last two columns of the matrix X . The remaining matrix is

$$\begin{bmatrix} \mathbf{a}^2 \\ \mathbf{a}^3 \\ \mathbf{b}^2 \\ \mathbf{b}^3 \end{bmatrix}$$

and the coefficient of $x'^1 x^1$ is equal to the determinant of this 4×4 matrix. In general, we may write

$$F_{ji} = (-1)^{i+j} \det \begin{bmatrix} \sim \mathbf{a}^i \\ \sim \mathbf{b}^j \end{bmatrix}. \quad (17.3)$$

In this expression, the notation $\sim \mathbf{a}^i$ has been used to denote the matrix obtained from A by *omitting* the row $\mathbf{a}^i$. Thus the symbol $\sim$ may be read as *omit*, and $\sim \mathbf{a}^i$ represents two rows of A . The determinant appearing on the right side of (17.3) is therefore a 4×4 determinant.

A different way of writing the expression for F_{ji} makes use of the tensor ϵ_{rst} (defined in section A1.1(p563)) as follows:¹

$$F_{ji} = \left(\frac{1}{4}\right) \epsilon_{ipq} \epsilon_{jrs} \det \begin{bmatrix} \mathbf{a}^p \\ \mathbf{a}^q \\ \mathbf{b}^r \\ \mathbf{b}^s \end{bmatrix}. \quad (17.4)$$

To see this, note that F_{ji} is defined in (17.4) in terms of a sum of determinants over all values of p, q, r and s . However, for a given value of i , the tensor ϵ_{ipq} is zero unless p and q are different from i and from each other. This leaves only two remaining choices of p and q (for example if $i = 1$, then we may choose $p = 2, q = 3$ or $p = 3, q = 2$). Similarly, there are only two different choices of r and s giving rise to non-zero terms. Thus the sum consists of four non-zero terms only. Furthermore, the determinants appearing in these four terms consist of the same four rows of the matrices A and B and hence have equal values, except for sign. However, the value of $\epsilon_{ipq} \epsilon_{jrs}$ is such that the four terms all have the same sign and are equal. Thus, the sum (17.4) is equal to the single term appearing in (17.3).

17.1.1 Epipoles as tensors

The expression (17.3) for the fundamental matrix involves determinants of matrices containing two rows from each of A and B. If we consider instead the determinants of matrices containing all three rows from one matrix and one row from the other matrix, the resulting determinants represent the epipoles. Specifically we have

$$\mathbf{e}^i = \det \begin{bmatrix} \mathbf{a}^i \\ \mathbf{B} \end{bmatrix} \quad \mathbf{e}^j = \det \begin{bmatrix} \mathbf{A} \\ \mathbf{b}^j \end{bmatrix} \quad (17.5)$$

where $\mathbf{e}$ and $\mathbf{e}'$ are the epipoles in the two images. To see this, note that the epipole is defined by $\mathbf{e}^i = \mathbf{a}^i \mathbf{C}'$, where $\mathbf{C}'$ is the centre of the second camera, defined by $\mathbf{B}\mathbf{C}' = \mathbf{0}$. The formula (17.5) is now obtained by expanding the determinant by cofactors along the first row (in a similar manner to the derivation of (3.4-p67)).

17.1.2 Affine specialization

In the case where both the cameras are affine cameras, the fundamental matrix has a particularly simple form. Recall that an affine camera matrix is one for which the final row is $(0, 0, 0, 1)$. Now, note from (17.3) that if neither i nor j is equal to 3, then the third rows of both A and B are present in this expression for F_{ij} . The determinant has two equal rows, and hence equals zero. Thus, F is of the form

$$\mathbf{F}_A = \begin{bmatrix} & a \\ & b \\ c & d & e \end{bmatrix}$$

¹ Of course the factor 1/4 is inessential since F is defined only up to scale. It is included here just to show the relationship to (17.3).

with all other entries being zero. Thus the affine fundamental matrix has just 5 non-zero entries, and hence 4 degrees of freedom. Its properties are described in section 14.2(p345).

Note that this argument relies solely on the fact that both cameras have the same third row. Since the third row of a camera matrix represents the principal plane of the camera (see section 6.2.1(p158)), it follows that the fundamental matrix for two cameras sharing the same principal plane is of the above form.

17.2 Trilinear relations

The determinant method of deriving the fundamental matrix can be used to derive relationships between the coordinates of points seen in three views. This analysis results in a formula for the trifocal tensor. Unlike the fundamental matrix, the trifocal tensor relates both lines and points in the three images. We begin by describing the relationships for corresponding points.

17.2.1 Trifocal point relations

Consider a point correspondence across three views: $\mathbf{x} \leftrightarrow \mathbf{x}' \leftrightarrow \mathbf{x}''$. Let the third camera matrix be $\mathbf{C}$ and let $\mathbf{c}^i$ be its i -th row. Analogous to (17.1) we can write an equation describing the projection of a point $\mathbf{X}$ into the three images as

$$\begin{bmatrix} \mathbf{A} & \mathbf{x} & & \\ \mathbf{B} & & \mathbf{x}' & \\ \mathbf{C} & & & \mathbf{x}'' \end{bmatrix} \begin{pmatrix} \mathbf{X} \\ -k \\ -k' \\ -k'' \end{pmatrix} = 0. \quad (17.6)$$

This matrix, which as before we will call $\mathbf{X}$, has 9 rows and 7 columns. From the existence of a solution to this set of equations, we deduce that its rank must be at most 6. Hence any 7×7 minor has zero determinant. This fact gives rise to the trilinear relationships that hold between the coordinates of the points $\mathbf{x}$, $\mathbf{x}'$ and $\mathbf{x}''$.

There are essentially two different types of 7×7 minors of $\mathbf{X}$. In choosing 7 rows of $\mathbf{X}$, we may choose either

- (i) Three rows from each of two camera matrices and one row from the third, or
- (ii) Three rows from one camera matrix and two rows from each of the two others.

Let us consider the first type. A typical such 7×7 minor of $\mathbf{X}$ is of the form

$$\begin{bmatrix} \mathbf{A} & \mathbf{x} & & \\ \mathbf{B} & & \mathbf{x}' & \\ \mathbf{c}^i & & & x''' \end{bmatrix}. \quad (17.7)$$

Note that this matrix contains only one entry in the last column, namely x''' . Expanding the determinant by cofactors down this last column reveals that the determinant is equal to

$$x''' \det \begin{bmatrix} \mathbf{A} & \mathbf{x} \\ \mathbf{B} & \mathbf{x}' \end{bmatrix}.$$

Apart from the factor x''^i , this just leads to the bilinear relationship expressed by the fundamental matrix, as discussed in section 17.1.

The other sort of 7×7 minor is of more interest. An example of such a determinant is of the form

$$\det \begin{bmatrix} \mathbf{A} & \mathbf{x} \\ \mathbf{b}^j & x'^j \\ \mathbf{b}^l & x'^l \\ \mathbf{c}^k & x''^k \\ \mathbf{c}^m & x''^m \end{bmatrix}. \quad (17.8)$$

By the same sort of argument as with the bilinear relations it is seen that setting the determinant to zero leads to a trilinear relation of the form $f(\mathbf{x}, \mathbf{x}', \mathbf{x}'') = 0$. By expanding this determinant down the column containing x^i , we can find a specific formula as follows.

$$\det \mathbf{X}_{uv} = -\frac{1}{2} x^i x'^j x''^k \epsilon_{ilm} \epsilon_{jqu} \epsilon_{krv} \det \begin{bmatrix} \mathbf{a}^l \\ \mathbf{a}^m \\ \mathbf{b}^q \\ \mathbf{c}^r \end{bmatrix} = 0_{uv} \quad (17.9)$$

where u and v are free indices corresponding to the rows omitted from the matrices $\mathbf{B}$ and $\mathbf{C}$ to produce (17.8). We introduce the tensor

$$\mathcal{T}_i^{qr} = \frac{1}{2} \epsilon_{ilm} \det \begin{bmatrix} \mathbf{a}^l \\ \mathbf{a}^m \\ \mathbf{b}^q \\ \mathbf{c}^r \end{bmatrix}. \quad (17.10)$$

The trilinear relationship (17.9) may then be written

$$x^i x'^j x''^k \epsilon_{jqu} \epsilon_{krv} \mathcal{T}_i^{qr} = 0_{uv}. \quad (17.11)$$

The tensor $\mathcal{T}_i^{qr}$ is the trifocal tensor, and (17.11) is a trilinear relation such as those discussed in section 15.2.1 (p378). The indices u and v are free indices, and each choice of u and v leads to a different trilinear relation.

Just as in the case of the fundamental matrix, one may write the formula for the tensor $\mathcal{T}_i^{qr}$ in a slightly different way

$$\mathcal{T}_i^{qr} = (-1)^{i+1} \det \begin{bmatrix} \sim \mathbf{a}^i \\ \mathbf{b}^q \\ \mathbf{c}^r \end{bmatrix}. \quad (17.12)$$

As in section 17.1, the expression $\sim \mathbf{a}^i$ means the matrix $\mathbf{A}$ with row i omitted. Note that we omit row i from the first camera matrix, but *include* rows q and r from the other two camera matrices.

In the often-considered case where the first camera matrix $\mathbf{A}$ has the canonical form $[\mathbf{I} \mid \mathbf{0}]$, the expression (17.12) for the trifocal tensor may be written simply as

$$\mathcal{T}_i^{qr} = b_i^q c_4^r - b_4^q c_i^r. \quad (17.13)$$

Note that there are in fact 27 possible trilinear relations that may be formed in this way (refer to (17.8)). Specifically, note that each relation arises from taking all three rows from one camera matrix along with two rows from each of the other two matrices. This gives the following computation.

- 3 ways to choose the first camera matrix from which to take all three rows.
- 3 ways to choose the row to omit from the second camera matrix.
- 3 ways to choose the row to omit from the third camera matrix.

This gives a total of 27 trilinear relations. However, among the 9 ways of choosing two rows from the second and third camera matrices, only 4 are linearly independent (we return to this in section 17.6). This means that there are a total of 12 linearly independent trilinear relations.

It is important to distinguish between the number of trilinear relations, however, and the number of different trifocal tensors. As is shown by (17.11), several different trilinear relations may be expressed in terms of just one trifocal tensor. In (17.11) each distinct choice of the free indices u and v gives rise to a different trilinear relation, all of which are expressible in terms of the same trifocal tensor $\mathcal{T}_i^{qr}$. On the other hand, in the definition of the trifocal tensor given in (17.10), the camera matrix A is treated differently from the other two, in that A contributes two rows (after omitting row i) to the determinant defining any given entry of $\mathcal{T}_i^{qr}$, whereas the other two camera matrices contribute just one row. This means that there are in fact three different trifocal tensors corresponding to the choice of which of the three camera matrices contributes two rows.

17.2.2 Trifocal line relations

A line in an image is represented by a covariant vector l_i , and the condition for a point x to lie on the line is that $l_i x^i = 0$. Let X^j represent a point X in space, and a_j^i represent a camera matrix A . The 3D point X^j is mapped to the image point as $x^i = a_j^i X^j$. It follows that the condition for the point X^j to project to a point on the line l_i is that $l_i a_j^i X^j = 0$. Another way of looking at this is that $l_i a_j^i$ represents a plane consisting of all points that project onto the line l_i .

Consider the situation where a point X^j maps to a point x^i in one image and to some point on lines l'_q and l''_r in two other images. This may be expressed by equations

$$x^i = k a_j^i X^j \quad l'_q b_j^q X^j = 0 \quad l''_r c_j^r X^j = 0.$$

These may be written as a single matrix equation of the form

$$\begin{bmatrix} A & x \\ l'_q b^q & 0 \\ l''_r c^r & 0 \end{bmatrix} \begin{pmatrix} X \\ -k \end{pmatrix} = 0. \quad (17.14)$$

Since this set of equations has a solution, it follows that $\det X = 0$, where X is the matrix on the left of the equation. Expanding this determinant down the last column

gives

$$\begin{aligned}
 0 = -\det X &= \frac{1}{2} x^j \epsilon_{ilm} \det \begin{bmatrix} \mathbf{a}^l \\ \mathbf{a}^m \\ l'_q \mathbf{b}^q \\ l''_r \mathbf{c}^r \end{bmatrix} = \frac{1}{2} x^j l'_q l''_r \epsilon_{ilm} \det \begin{bmatrix} \mathbf{a}^l \\ \mathbf{a}^m \\ \mathbf{b}^q \\ \mathbf{c}^r \end{bmatrix} \\
 &= x^j l'_q l''_r \mathcal{T}_i^{qr}.
 \end{aligned} \tag{17.15}$$

This shows the connection of the trifocal tensor with sets of lines. The two lines l'_q and l''_r back-project to planes meeting in a line in space. The image of this line in the first image is a line, which may be represented by l_i . For any point x^i on that line the relation (17.15) holds. It follows that $l'_q l''_r \mathcal{T}_i^{qr}$ is the representation of the line l_i . Thus, we see that for three corresponding lines in the three images

$$l_p = l'_q l''_r \mathcal{T}_p^{qr} \tag{17.16}$$

where, of course, the two sides are equal only up to a scale factor. Since the two sides of the relation (17.16) are vectors, this may be interpreted as meaning that the vector product of the two sides vanishes. Expressing this vector product using the tensor ϵ^{ijk} , we arrive at an equation

$$l_p l'_q l''_r \epsilon^{ipw} \mathcal{T}_i^{qr} = 0^w. \tag{17.17}$$

In an analogous manner to the derivation of (17.11) and (17.15) we may derive a relationship between corresponding points in two images and a line in a third image. In particular, if a point x^j in space maps to points x^i and x'^j in the first two images, and to some point on a line l''_r in the third image, then the relation is

$$x^i x'^j l''_r \epsilon_{jqu} \mathcal{T}_i^{qr} = 0_u. \tag{17.18}$$

In this relation, the index u is free, and there is one such relation for each choice of $u = 1, \dots, 3$, of which two are linearly independent.

We summarize the results of this section in table 17.1, where the final column denotes the number of linearly independent equations.

Correspondence	Relation	Number of equations
three points	$x^i x'^j x''^k \epsilon_{jqk} \epsilon_{krv} \mathcal{T}_i^{qr} = 0_{uv}$	4
two points, one line	$x^i x'^j l''_r \epsilon_{jqk} \mathcal{T}_i^{qr} = 0_u$	2
one point, two lines	$x^i l'_q l''_r \mathcal{T}_i^{qr} = 0$	1
three lines	$l_p l'_q l''_r \epsilon^{ipw} \mathcal{T}_i^{qr} = 0^w$	2

Table 17.1. *Trilinear relations (see also table 16.1(p391)).*

Note how the different equation sets are related to each other. For instance, the second line of the table is derived from the first by replacing $x''^k \epsilon_{krv}$ by the line l''_r and deleting the free index v .

17.2.3 Relations between two views and the trifocal tensor

To this point we have considered correspondences across three views and the trifocal tensor. Here we describe the constraints that arise if the correspondence is only across two views. From the two view case, where point correspondences constrain the fundamental matrix, we would expect some constraint on $\mathcal{T}$.

Consider the case of corresponding points x^{lj} and x^{mk} in the second and third images. This means that there is a point X in space mapping to the points x^{lj} and x^{mk} . The point X also maps to some point x^i in the first image, but x^i is not known. Nevertheless, there exists a relationship $x^i x^{lj} x^{mk} \epsilon_{jqu} \epsilon_{kqv} \mathcal{T}_i^{qr} = 0_{uv}$ between these points. For each choice of u and v , denote $A_{i,uv} = x^{lj} x^{mk} \epsilon_{jqu} \epsilon_{kqv} \mathcal{T}_i^{qr}$. The entries of $A_{i,uv}$ are linear expressions in the entries of $\mathcal{T}_i^{qr}$ that may be determined explicitly in terms of the known points x^{lj} and x^{mk} . There exists a point x such that $x^i A_{i,uv} = 0$. For each choice of u, v we may consider $A_{i,uv}$ as being a 3-vector indexed by i , and for the different choices of u and v , there are 4 linearly independent such expressions. Thus, A may be considered as a 3×4 matrix, and the condition that $x^i A_{i,uv} = 0$ means that $A_{i,uv}$ has rank 2. This means that every 3×3 subdeterminant of A is zero, which leads to cubic constraints on the elements of the trifocal tensor. For geometric reasons, it appears that the equations $x^i A_{i,uv}$ are not algebraically independent for the four choices of u and v . Consequently we obtain a single cubic constraint on $\mathcal{T}_i^{jk}$ from a two-view point correspondence. Details are left to the reader.

In the case where the point correspondence is between the first and second (or third) views the analysis is slightly different. However, the result in each case is that although a point correspondence across two views leads to a constraint on the trifocal tensor, this constraint is not a linear constraint as it is in the case of correspondences across three views.

17.2.4 Affine trifocal tensor

In the case where all three cameras are affine, the trifocal tensor satisfies certain constraints. A camera matrix is affine if the last row is $(0, 0, 0, 1)$. It follows that if two of the rows in the matrix in (17.12) are of this form, then the corresponding element of $\mathcal{T}_i^{jk}$ is zero. This is the case for elements $\mathcal{T}_1^{j3}, \mathcal{T}_2^{j3}, \mathcal{T}_1^{3k}, \mathcal{T}_2^{3k}$ and $\mathcal{T}_3^{33}$ – a total of 11 elements. Thus the trifocal tensor contains 16 non-zero entries, defined up to scale. As in the case of the affine fundamental matrix, this analysis is equally valid for the case of cameras sharing the same principal plane.

17.3 Quadrilinear relations

Similar arguments work in the case of four views. Once more, consider a point correspondence across 4 views: $x \leftrightarrow x' \leftrightarrow x'' \leftrightarrow x'''$. With camera matrices A, B, C and D ,

the projection equations may be written as

$$\begin{bmatrix} A & x & & \\ B & & x' & \\ C & & & x'' \\ D & & & & x''' \end{bmatrix} \begin{pmatrix} X \\ -k \\ -k' \\ -k'' \\ -k''' \end{pmatrix} = 0. \quad (17.19)$$

Since this equation has a solution, the matrix X on the left has rank at most 7, and so all 8×8 determinants are zero. As in the trilinear case, any determinant containing only one row from one of the camera matrices gives rise to a trilinear or bilinear relation between the remaining views. A different case occurs when we consider 8×8 determinants containing two rows from each of the camera matrices. Such a determinant leads to a new quadrilinear relationship of the form

$$x^i x^j x^k x^l \epsilon_{ipm} \epsilon_{jqx} \epsilon_{kry} \epsilon_{lsz} Q^{pqrs} = 0_{wxyz} \quad (17.20)$$

where each choice of the free variables w, x, y and z gives a different equation, and the 4-dimensional *quadrifocal tensor* Q^{pqrs} is defined by

$$Q^{pqrs} = \det \begin{bmatrix} a^p \\ b^q \\ c^r \\ d^s \end{bmatrix}. \quad (17.21)$$

Note that the four indices of the four-view tensor are contravariant, and there is no distinguished view as there is in the case of the trifocal tensor. There is only one four-view tensor corresponding to four given views, and this one tensor gives rise to 81 different quadrilinear relationships, of which 16 are linearly independent (see section 17.6).

As in the case of the trifocal tensor, there are also relations between lines and points in the case of the four-view tensor. Equations relating points are really just special cases of the relationship for lines. In the case of a 4-line correspondence, however, something different happens, as will now be explained. The relationship between a set of four lines and the quadrifocal tensor is given by the formula

$$l_p l'_q l''_r l'''_s Q^{pqrs} = 0 \quad (17.22)$$

for any set of corresponding lines l_p, l'_q, l''_r and l'''_s . However, the derivation shows that this condition will hold as long as there is a single point in space that projects onto the four image lines. It is not necessary that the four image lines correspond (in the sense that they are the image of a common line in space). This configuration is illustrated in figure 17.1a.

Now, consider the case where three of the lines (for instance l'_q, l''_r and l'''_s) correspond by deriving from a single 3D line (figure 17.1b). Now let l_p be any *arbitrary* line in the first image. The back-projection of this line is a plane, which will meet the 3D line in a single point, X , and the conditions are present for (17.22) to hold. Since this is true for any arbitrary line l_p , it must follow that $l'_q l''_r l'''_s Q^{pqrs} = 0^p$. This gives three linearly independent equations involving l'_q, l''_r and l'''_s . However, given a set of corresponding




Fig. 17.1. **Four line "correspondence"**. The four lines l, l', l'', l''' satisfy the quadrilinear relation (17.22) since their back-projections intersect in a common point X . (a) No three planes intersect in a common line. (b) The three lines $l' \leftrightarrow l'' \leftrightarrow l'''$ are the image of the same line L in 3-space.

lines in four images, as above, we may choose a subset of three lines, and for each line-triplet obtain three equations in this way. Since there are four choices of line-triplets, the total number of equations is thus 12.

However only 9 of these equations are independent, and this may be seen as follows. Suppose that $l = l' = l'' = l''' = (1, 0, 0)^T$. This is equivalent to the general situation since by applying projective transformations to each of the four images we can transform an arbitrary line correspondence to this case. Now the equation $l'_q l''_r l'''_s Q^{pqrs} = 0^p$ means that any element $Q^{p111} = 0$. Applying this argument to all four triplets of three views, we find that $Q^{pqrs} = 0$ whenever at least three of the indices are 1. There are a total of 9 such elements. Since the set of equations generated by the line correspondence is equivalent to setting each of these elements to zero, among the total of 12 equations there are just 9 independent ones.

The four-view relations are summarized in table 17.2. No equation is given here for the case of three lines and one point, since this gives no more restrictions on the tensor than just the three-line correspondence.

Correspondence	Relation	Number of equations
four points	$x^i x^j x^{(k} x^{l)} \epsilon_{ipw} \epsilon_{jqx} \epsilon_{kry} \epsilon_{lsz} Q^{pqrs} = 0_{wxyz}$	16
three points, one line	$x^i x^j x^{(k} l^{l')} \epsilon_{ipw} \epsilon_{jqx} \epsilon_{kry} Q^{pqrs} = 0_{wxyz}$	8
two points, two lines	$x^i x^j l^{l'} l^{l''} \epsilon_{ipw} \epsilon_{jqx} Q^{pqrs} = 0_{wx}$	4
three lines	$l_p l'_q l''_r Q^{pqrs} = 0^s$	3
four lines	$l_p l'_q l''_r Q^{pqrs} = 0^s, \quad l_p l'_q l'''_s Q^{pqrs} = 0^r, \quad \dots$	9

Table 17.2. *Quadrilinear relations.*

17.4 Intersections of four planes

The multi-view tensors may be given a different derivation, which sheds a little more light on their meaning. In this interpretation, the basic geometric property is the intersection of four planes. Four planes in space will generally not meet in a common point. A necessary and sufficient condition for them to do so is that the determinant of the 4×4 matrix formed from the vectors representing the planes should vanish.

Notation. In this section only we shall represent the determinant of a 4×4 matrix with rows $\mathbf{a}$, $\mathbf{b}$, $\mathbf{c}$ and $\mathbf{d}$ by $\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c} \wedge \mathbf{d}$. In a more general context, the symbol $\wedge$ represents the meet (or intersection) operator in the double algebra (see literature section of this chapter). However, for the present purposes the reader need only consider it as a shorthand for the determinant.

We start with the quadrifocal tensor for which the derivation is easiest. Consider four lines l , l' , l'' and l''' in images formed from four cameras with camera matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ and $\mathbf{D}$. The back projection of a line l through camera $\mathbf{A}$ is written as the plane $l_i \mathbf{a}^i$, with notation as in (17.14). The condition that these four planes are coincident may be written as

$$(l_p \mathbf{a}^p) \wedge (l'_q \mathbf{b}^q) \wedge (l''_r \mathbf{c}^r) \wedge (l'''_s \mathbf{d}^s) = 0.$$

However, since the determinant is linear in each row, this may be written as

$$0 = l_p l'_q l''_r l'''_s (\mathbf{a}^p \wedge \mathbf{b}^q \wedge \mathbf{c}^r \wedge \mathbf{d}^s) \stackrel{\text{def}}{=} l_p l'_q l''_r l'''_s Q^{pqrs}. \quad (17.23)$$

This corresponds to the definition (17.21) and line relation (17.22) for the quadrifocal tensor. The basic geometric property is the intersection of the four planes in space.

Trifocal tensor derivation. Consider now a point–line–line relationship $x^i \leftrightarrow l'_j \leftrightarrow l''_k$ for three views and let l^1_p and l^2_q be two lines in the first image that pass through the image point $\mathbf{x}$. The planes back-projected from the four lines meet in a point (see figure 17.2). So we can write:

$$l^1_l l^2_m l'_q l''_r (\mathbf{a}^l \wedge \mathbf{a}^m \wedge \mathbf{b}^q \wedge \mathbf{c}^r) = 0.$$




Fig. 17.2. A point-line-line correspondence $\mathbf{x} \mapsto l' \mapsto l''$ involving three images may be interpreted as follows. Two arbitrary lines are chosen to pass through the point $\mathbf{x}$ in the first image. The four lines then back-project to planes that meet in a point in space.

The next step is an algebraic trick – to multiply this equation by the $\epsilon^{ilm} \epsilon_{ilm}$. This is a scalar value (in fact equal to 6, the number of permutations of (ilm)). The result after regrouping is

$$(l_l^1 l_m^2 \epsilon^{ilm}) l_q' l_r'' \epsilon_{ilm} (\mathbf{a}^l \wedge \mathbf{a}^m \wedge \mathbf{b}^q \wedge \mathbf{c}^r) = 0.$$

Now the expression $l_l^1 l_m^2 \epsilon^{ilm}$ is simply the cross-product of the two lines l_l and l_m , in other words their intersection point, x^i . Thus finally we can write

$$0 = x^i l_q' l_r'' (\epsilon_{ilm} (\mathbf{a}^l \wedge \mathbf{a}^m \wedge \mathbf{b}^q \wedge \mathbf{c}^r)) \stackrel{\text{def}}{=} x^i l_q' l_r'' \mathcal{T}_i^{qr} \quad (17.24)$$

which are the definition (17.10) and basic incidence relation (17.15) for the trifocal tensor.

Fundamental matrix. We can derive the fundamental matrix in the same manner. Given a correspondence $\mathbf{x} \leftrightarrow \mathbf{x}'$, select pairs of lines l_p^1 and l_q^2 passing through $\mathbf{x}$, and l_r^1 and l_s^2 passing through $\mathbf{x}'$. The back-projected planes all meet in a point, so we write

$$l_p^1 l_q^2 l_r^1 l_s^2 (\mathbf{a}^p \wedge \mathbf{a}^q \wedge \mathbf{b}^r \wedge \mathbf{b}^s) = 0.$$

Multiplying by $(\epsilon_{ipq} \epsilon^{ipq})(\epsilon_{jrs} \epsilon^{jrs})$ and proceeding as before leads to the coplanarity constraint

$$0 = x^i x'^j (\epsilon_{ipq} \epsilon_{jrs} (\mathbf{a}^p \wedge \mathbf{a}^q \wedge \mathbf{b}^r \wedge \mathbf{b}^s)) \stackrel{\text{def}}{=} x^i x'^j F_{ji} \quad (17.25)$$

which can be compared with (17.4).

17.5 Counting arguments

In this section we specify the number of points or lines required to carry out reconstruction from several views. This analysis is related to counting the number of degrees of freedom of the associated tensors. However, in doing this it is necessary to distinguish between the number of degrees of freedom of the tensor viewed as an unconstrained

algebraic object, and the number of degrees of freedom that arises from a configuration of cameras and their camera matrices.

For instance, consider the fundamental matrix. At one level, the fundamental matrix may be considered as a homogeneous 3×3 matrix, and hence has 8 degrees of freedom (9 minus 1 for the indeterminate scale). On the other hand, the fundamental matrix arising from a pair of camera matrices according to (17.3–p412) must satisfy the additional constraint $\det F = 0$. Hence, such a fundamental matrix has just 7 degrees of freedom. Since the camera matrices may be determined up to a 3D projectivity from the fundamental matrix (and vice versa), we may count the number of degrees of freedom of a fundamental matrix by counting the degrees of freedom of the camera matrices. Two camera matrices have a total of 22 degrees of freedom (two homogeneous 3×4 matrices). A 3D homography is represented by a 4×4 homogeneous matrix, and so has 15 degrees of freedom. This gives a total of $7 = 22 - 15$ degrees of freedom for the configuration of the two cameras, modulo a projective transformation. This agrees with the 7 degrees of freedom of the fundamental matrix, providing a check on the previous calculation.

Similarly, the trifocal tensor encodes the projective structure of three camera matrices, and hence has $18 = 3 \cdot 11 - 15$ degrees of freedom. In the same way the quadrifocal tensor has $29 = 4 \cdot 11 - 15$ degrees of freedom. Generally for m cameras, we have

$$\# \text{ dof} = 11m - 15.$$

Since the trifocal and quadrifocal tensors, considered just as homogeneous algebraic arrays have 26 and 80 degrees of freedom respectively, they must satisfy an additional set of constraints, dictated by the camera geometry – 8 constraints for the trifocal tensor and 51 constraints for the quadrifocal tensor.

In computing geometric structure, we may use linear algebraic methods based on estimating the multifocal tensor by solving the constraints from the multi-linearities. The number of correspondences required is determined from the number of equations generated by each point or line correspondence. This *linear method* takes no account of the constraints imposed on the tensors by their geometry.

On the other hand, the number of correspondences required may be determined by counting the number of geometric constraints given by each correspondence and comparing with the total number of degrees of freedom of the system. Consider a configuration of n points in m views. The total number of degrees of freedom of this system is $11m - 15 + 3n$, since each of the n 3D points has 3 degrees of freedom. To estimate the projective structure the available data are the images of the n points in m images, a total of $2mn$ measurements (each 2D point having two coordinates). Thus for reconstruction to be possible, we require $2mn \geq 11m - 15 + 3n$, or $(2m - 3)n \geq 11m - 15$. Thus the required number of points is

$$n \geq \frac{11m - 15}{2m - 3} = 5 + \frac{m}{2m - 3}.$$

One may also think of this by observing that each point correspondence contributes

# views	tensor	# elems	# dof	linear		non-linear	
				# points	# lines	# points	# lines
2	F	9	7	8	—	7*	—
3	$\mathcal{T}$	27	18	7	13	6*	9*?
4	Q	81	29	6	9	6	8?

Table 17.3. **Projective degrees of freedom and constraints.** *The linear column indicates the minimum number of correspondences across all views required to solve linearly for the tensor (up to scale). The non-linear is the minimum number of correspondences required. A star indicates multiple solutions, and a question-mark indicates that no practical reconstruction method is known.*

$2m - 3$ constraints on the cameras, that is $2m$ for the coordinates of the points in each view, less 3 for the added degree of freedom of the 3D point.

An analogous argument applies for line correspondences. A line has four degrees of freedom in 3-space, and its image line is described by 2 degrees of freedom, so that each line correspondence provides $2m - 4$ constraints and the number of lines l required is

$$l \geq \frac{11m - 15}{2m - 4}.$$

In either of these cases, if the number of constraints (equations) is equal to the number of degrees of freedom (unknowns) of the cameras and points or lines configuration, then we can in general expect multiple solutions, except in the linear case. However, if there are more equations than unknowns, the system is over-determined, and in the generic case there will exist a single solution.

Table 17.3 summarizes the number of correspondences required for reconstruction. A star (*) indicates that multiple solutions are expected. For the case of non-linear solutions, actual methods of solution are not always known, other than by brute-force generation and solution of a set of simultaneous polynomial equations. Situations where no simple method is known are marked with a question-mark. Note that not much is known about non-linear solutions for lines. Specific non-linear algorithms are:

- (i) 7 points in 2 views: see section 11.1.2(p281). Three solutions are possible.
- (ii) 6 points in 3 views: see section 20.2.4(p510). Three solutions are possible. This is the dual (see chapter 20) of the previous problem.
- (iii) 6 points in 4 views. Solve using 6 points in 3 views, then use the camera resection DLT algorithm (section 7.1(p178)) to find the fourth view and eliminate all but one solution. However, unlike the previous two cases, the case of 6 points in 4 views is overdetermined, and this solution is only valid for perfect data. The estimation for noisy data is discussed in chapter 20.

# views	tensor	# non-zero elements	# dof	linear		non-linear	
				# points	# lines	# points	# lines
2	F_A	5	4	4	—	4	—
3	T_A	16	12	4	8	4	6*?
4	Q_A	48	20	4	6	4	5*?

Table 17.4. **Affine degrees of freedom and constraints.** *The camera for each view is affine. See caption of table 17.3 for details.*

17.5.1 Affine cameras

For affine cameras, the reconstruction problem requires fewer correspondences. The plane at infinity may be identified in a reconstruction, being the principal plane of all of the cameras, and the reconstruction ambiguity is affine, not projective. The number of degrees of freedom for m affine cameras is

$$\# \text{ dof} = 8m - 12.$$

Each view adds 8 degrees of freedom for the new 3×4 affine camera matrix, and we subtract 12 for the degrees of freedom of a 3D affine transformation.

As in the projective case a point correspondence imposes $2m - 3$ constraints, and a line correspondence $2m - 4$ constraints. As before the number of required points may be computed as $n(2m - 3) \geq 8m - 12$, i.e.

$$n \geq \frac{8m - 12}{2m - 3} = 4.$$

For lines the result is

$$l \geq \frac{8m - 12}{2m - 4} = 4 + \frac{2}{m - 2}.$$

As for linear methods, the number of elements in the tensor is 3^m . In the affine case the tensor is, as always, only defined up to scale, but additionally many of the elements are zero, as seen in section 17.1.2 and section 17.2.4. This cuts down on the required number of correspondences. Counting results are given in table 17.4. Note that for point correspondences, the linear algorithms apply with the minimum number of correspondences given by the above equation. Hence, the non-linear algorithms are the same as the linear ones.

17.5.2 Knowing four coplanar points – plane plus parallax

The previous counting arguments were for the case of points and lines in general position. Here we consider the important case that four or more of the 3D points are coplanar. It will be seen that the computation of the tensors, and consequently projective structure, is significantly simplified. This discussion is given in terms of knowing

# views	tensor	# dof	points		lines	
			# points	# constraints	# lines	# constraints
2	F	2	2	2	—	
3	$\mathcal{T}$	5	2	5	5	5
4	Q	8	2	8	4	8

Table 17.5. Number of additional correspondences required to compute projective structure given 2D homographies between the views, induced by a plane. The homographies may be computed from four or more point matches derived from coplanar points, or by any other means. Points used to compute the homographies are not counted in this table.

four coplanar points in the images. However, all that is really important here is that inter-image homographies induced by a plane should be known.

Computation of the fundamental matrix knowing four coplanar points was considered in section 13.3(p334) (see algorithm 13.2(p336)), and for the trifocal tensor in section 16.7.1. Now, we consider the four-view case. From the correspondence of 4 (or more) points derived from 3D points in a plane, we can compute the homographies H' , H'' and H''' from the first view to the second, third and fourth views respectively, induced by the plane. In a projective reconstruction, we choose the plane containing the points as the plane at infinity, in which case H' , H'' and H''' are the infinite homographies. Assuming further that the first image is at the origin, the four camera matrices may now be written as

$$A = [I \mid 0] \quad B = [H' | t'] \quad C = [H'' | t''] \quad D = [H''' | t''']$$

The vectors t' , t'' , t''' may be determined only up to a common scale.

Since the left hand 3×3 blocks of the camera matrices are now known, it is easily seen from (17.21) that the entries of Q^{pqrs} are linear in the remaining entries t' , t'' and t''' of the camera matrices. Using (17.21) we can write down an expression for the entries of Q^{pqrs} in terms of the entries of t' , t'' and t''' . In fact, we may write this explicitly as $q = Mt$, where M is an 81×9 matrix and q and t are vectors representing the entries of Q and $t', \dots, t'''$. Thus, the quadrifocal tensor may be linearly parametrized in terms of 9 homogeneous coordinates and hence has 8 degrees of freedom.

Now, given a set of equations $Eq = 0$ derived from correspondences as in table 17.2, they may be rewritten in terms of the minimal parameter set t by substituting $q = Mt$ arriving at $EMt = 0$. This allows a linear solution for t and hence the camera matrices, and (if desired) the tensor, $q = Mt$. Note the important advantage here that the tensor so obtained automatically corresponds to a set of camera matrices, and so satisfies all geometric constraints – the troublesome 51 constraints that must be satisfied by the quadrifocal tensor disappear into thin air.

The above analysis was for the computation of the quadrifocal tensor, but it applies equally well to the fundamental matrix and trifocal tensor as well.

Counting argument. We return to the general case of m views and consider the situation from a geometric point of view. The number of degrees of freedom parametrizing the tensor is equal to the number of geometric degrees of freedom remaining in the camera matrices, namely $3m - 4$ where m is the number of views. This arises from $3(m - 1)$ for the last column of each camera matrix apart from the first, less one for the common scaling.

Counting the number of constraints imposed by a point or a line correspondence is a little tricky, however, and shows that one must always be careful with counting arguments not to neglect hidden dependencies. First, consider a line correspondence $l \leftrightarrow l' \leftrightarrow l''$ across three views – the argument will hold in general for $m \geq 3$ views. The question to be resolved is just how much information can be derived from the measurement of the image lines. Surprisingly, knowledge of the plane homographies between the images reduces the amount of information that a line correspondence provides. For simplicity of argument, we may assume that planar homographies have been applied to the images so that the four coplanar reference points map to the same point in each image. As a result any further point in the reference plane will map to the same point in each image. Now the 3D line L from which the line correspondence is derived must meet the reference plane in a point X . Since it lies on the reference plane, X projects to the same point x in all images, and x must lie on all three image lines l, l', l'' . Thus, corresponding lines in the three views cannot be arbitrary – they must all pass through a common image point. In the general case the number of degrees of freedom of the measurements of the lines in m views is $m + 2$. To specify the point x requires 2 degrees of freedom, and each line through the point then has one remaining degree of freedom (its orientation). Subtracting 4 for the four added degrees of freedom of the line in space, we see

- *Each line correspondence across m views generates $m - 2$ constraints on the remaining degrees of freedom of the cameras.*

Note how the condition that the image lines must meet in a point restricts the number of equations generated by the line correspondence. Without observing this condition, one would expect $2m - 4$ constraints from each line. However, with perfect data the set of $2m - 4$ equations will have rank only $m - 2$. With noisy data the image lines will not be exactly coincident, and the system may have full rank, but this will be entirely due to noise, and the smallest singular values of the system will be essentially random. One should, however, include all available equations in solving the system, since this will diminish the noise effects.

For point correspondences the argument is similar. The line through any two image points is the image of a line in space. The projections of this 3D line are constrained as discussed in the preceding discussion, and this imposes constraints on the matching points. The measurements have $3m + 2$ degrees of freedom: 2 for the intersection point of the line with the plane and, in each view, one for the line orientation and two corresponding to the position of each of the two points on the line. Subtracting 6 for the degrees of freedom of the two points in 3-space, the result is

- *Two point correspondences across m views generate $3m - 4$ constraints on the remaining degrees of freedom of the cameras.*

Since this is the same as the number of degrees of freedom of the cameras, 2 points are sufficient to compute structure.

Since the number of geometric constraints on the cameras is the same as the number of degrees of freedom of the tensor, there is no distinction between constraints on the tensor and constraints on the geometry. Thus, point and line correspondences may be used to generate linear constraints on the tensor. There is no need for non-linear methods. The number of required correspondences is summarized in table 17.5.

17.6 Number of independent equations

It was asserted in table 17.2 that each four-view point correspondence gives rise to 16 linearly independent equations in the entries of the quadrifocal tensor. We now examine this point more closely.

Given sufficiently many point matches across four views, one may solve for the tensor Q^{pqrs} . Once Q is known, it is possible to solve for the camera matrices and hence compute projective structure. This was shown in [Heyden-95b, Heyden-97c, Hartley-98c] but is not discussed further in this book. A curious phenomenon occurs however when one counts the number of point matches necessary to do this. As indicated above, it appears to be the case that each point match gives 16 linearly independent equations in the entries of the tensor Q^{pqrs} , and it seems unlikely that the equations derived from two totally unrelated sets of point correspondences could have any dependencies. It would therefore appear that from five point correspondences one obtains 80 equations, which is enough to solve for the 81 entries of Q^{pqrs} up to scale. From this argument it would seem that it is possible to solve for the tensor from only five point matches across four views, and thence one may solve for the camera matrices, up to the usual projective ambiguity. This conclusion however is contradicted by the following remark.

- *It is not possible to determine the positions of four (or any number of) cameras from the images of five points.*

This follows from the counting arguments of section 17.5. Obviously there is some error in our counting of equations. The truth is contained in the following.

Result 17.1. *The full set of 81 linear equations (17.20) derived from a single point correspondence $\mathbf{x} \leftrightarrow \mathbf{x}' \leftrightarrow \mathbf{x}'' \leftrightarrow \mathbf{x}'''$ across four views contains 16 independent constraints on Q^{pqrs} . Furthermore, let the equations be written as $\mathbf{A}\mathbf{q} = \mathbf{0}$ where $\mathbf{A}$ is an 81×81 matrix and $\mathbf{q}$ is a vector containing the entries of Q^{pqrs} . Then the 16 non-zero singular values of $\mathbf{A}$ are all equal.*

What this result is saying is that indeed as expected one obtains 16 linearly independent equations from one point correspondence, and in fact it is possible to reduce this set of equations by an orthogonal transformation (multiplication of the equation matrix $\mathbf{A}$ on the left by an orthogonal matrix $\mathbf{U}$) to a set of 16 orthogonal equations. The proof

is postponed until the end of this section. The surprising fact however is that the equation sets corresponding to two unrelated point correspondences have a dependency, as stated in the following result.

Result 17.2. *The set of equations (17.20–p419) derived from a set of n general point correspondences across four views has rank $16n - \binom{n}{2}$, for $n \leq 5$.*

The notation $\binom{n}{2}$ means the number of choices of 2 among n , specifically, $\binom{n}{2} = n(n-1)/2$. Thus for 5 points there are only 70 independent equations, not enough to solve for Q^{pqrs} . For $n = 6$ points, $16n - \binom{n}{2} = 81$, and we have enough equations to solve for the 81 entries of Q^{pqrs} .

We now prove the two results above. The key point in the proof of result 17.1 concerns the singular values of a skew-symmetric matrix (see result A4.1(p581)).

Result 17.3. *A 3×3 skew-symmetric matrix has two equal non-zero singular values.*

The rest of the proof of result 17.1 is quite straightforward as long as one does not get lost in notation.

Proof. (Result 17.1) The full set of 81 equations derived from a single point correspondence is of the form $x^i \epsilon_{ipw} x^j \epsilon_{jqz} x^m \epsilon_{krq} x^n \epsilon_{lrs} Q^{pqrs} = 0_{wxyz}$. A total of 81 equations are generated by varying w, x, y, z over the range 1, ..., 3. Thus, the equation matrix A may be written as

$$A_{(wxyz)(pqrs)} = x^i \epsilon_{ipw} x^j \epsilon_{jqz} x^m \epsilon_{krq} x^n \epsilon_{lrs} \quad (17.26)$$

where the indices $(wxyz)$ index the row and $(pqrs)$ index the column of A . We will consider a set of indices, such as $(wxyz)$ in this case, as a single index for the row or column of a matrix. This situation will be indicated by enclosing the indices in parentheses as here, and referring to them as a *combined index*.

We consider now the expression $x^i \epsilon_{ipw}$. This may be considered as a matrix indexed by the free indices p and w . Furthermore, since $x^i \epsilon_{ipw} = -x^i \epsilon_{ipw}$, we see that it is a skew-symmetric matrix, and hence has equal singular values. We denote this matrix by S_{wp} . Writing result 17.3 using tensor notation, we have

$$U_a^w S_{wp} V_c^p = k D_{ac} \quad (17.27)$$

where the diagonal matrix D is as in result 17.3. The matrix A in (17.26) may be written as $A_{(wxyz)(pqrs)} = S_{wp} S_{xq} S_{yr} S_{zs}$. Consequently, applying (17.27) we may write

$$U_a^w U_b^x U_c^y U_d^z A_{(wxyz)(pqrs)} V_e^p V_f^q V_g^r V_h^s = k k' k'' k''' D_{ac} D_{bf} D_{eg} D_{dh} \quad (17.28)$$

Now, writing

$$\hat{U}_{(abcd)}^{(wxyz)} = U_a^w U_b^x U_c^y U_d^z \quad \hat{V}_{(efgh)}^{(pqrs)} = V_e^p V_f^q V_g^r V_h^s \quad \hat{D}_{(abcd)(efgh)} = D_{ac} D_{bf} D_{eg} D_{dh}$$

and $\tilde{k} = k k' k'' k'''$ we see that (17.28) may be written as

$$\hat{U}_{(abcd)}^{(wxyz)} A_{(wxyz)(pqrs)} \hat{V}_{(efgh)}^{(pqrs)} = \tilde{k} \hat{D}_{(abcd)(efgh)} \quad (17.29)$$

As a matrix, $D_{(abcd)(efgh)}$ is diagonal with 16 non-zero diagonal entries, all equal to

unity. To show that (17.29) is the SVD of the matrix $A_{(wxyz)(pqrs)}$, and hence to complete the proof, it remains only to show that $U_{(abcd)}^{(wxyz)}$ and $V_{(efgh)}^{(pqrs)}$ are orthogonal matrices. It is necessary only to show that the columns are orthonormal. This is straightforward and is left as an exercise. $\square$

Proof. (Result 17.2) We consider two point correspondences across four views, namely $x^i \leftrightarrow x'^i \leftrightarrow x''^i \leftrightarrow x'''^i$ and $y^j \leftrightarrow y'^j \leftrightarrow y''^j \leftrightarrow y'''^j$. These give rise to two sets of equations $A^x q = 0$ and $A^y q = 0$ of the form (17.26). Each of these matrices has rank 16, and if the rows of A^x are independent of the rows of A^y , then the rank of the combined set of equations is 32. However, if there is a linear dependence between the rows of A^x and those of A^y then their combined rank is at most 31.

We define a vector s_x with combined index $(pqrs)$ by $s_x^{(pqrs)} = x^p x'^q x''^r x'''^s$. A vector s_y is similarly defined. We will demonstrate that $s_y^T A^x = s_x^T A^y$, which means that the rows of A^x and A^y are linearly dependent.

Expanding $s_y^T A^x$ gives

$$\begin{aligned} s_y^T A^x &= s_y^{(wxyz)} A_{(wxyz)(pqrs)}^x \\ &= (y^w y'^x y''^y y'''^z) x^i \epsilon_{ipw} x'^j \epsilon_{jqx} x''^k \epsilon_{kr y} x'''^l \epsilon_{lsz} \\ &= (x^i x'^j x''^k x'''^l) y^w \epsilon_{wpi} y'^x \epsilon_{xqj} y''^y \epsilon_{yrk} y'''^z \epsilon_{zsl} \\ &= s_x^T A^y. \end{aligned}$$

This demonstrates that the rows of A^x and A^y are dependent, and their combined rank is at most 31. We now consider the possibility that the combined rank is less than 31. In such a case, all 31×31 subdeterminants of the matrix $[A^x{}^T, A^y{}^T]^T$ must vanish. These subdeterminants may be expressed as polynomial expressions in the coefficients of the points $x, x', x'', x''', y, y', y'', y'''$. These 24 coefficients together generate a 24-dimensional space. Thus, there is a function $f: \mathbb{R}^{24} \rightarrow \mathbb{R}^N$ for some N (equal to the number of such 31×31 subdeterminants), such that the equation matrix has rank less than 31 only on the set of zeros of the function f . Any arbitrarily chosen example (omitted) may be used to show that the function f is not identically zero. It follows that the set of point correspondences for which the set of equations has rank less than 31 is a variety in $\mathbb{R}^{24}$, and hence is nowhere dense. Thus, the set of equations generated by a general pair of point correspondences across four views has rank 31.

We now turn to the general case of n point correspondences across all four views. Note that the linear relationship that holds for two point correspondences is non-generic, but depends on the pair of correspondences. In general, therefore, given n point correspondences, there will be $\binom{n}{2}$ such relationships. This reduces the dimension of the space spanned by the set of equations to $16n - \binom{n}{2}$ as required. $\square$

The three-view case. Similar arguments hold in the three-view case. It can be proved in the same manner that the nine equations arising from a point match contain four independent equations. This is left to an exercise (see page 432).

17.7 Choosing equations

In section 17.6 a proof was given that the singular values of the full set of equations derived from four point equations are all equal. The proof may easily be adapted to the three-view case as well (see exercise, page 432). The key point in the argument is that the two non-zero singular values of a 3×3 skew-symmetric matrix are equal. This proof may clearly be extended to apply to any of the other sets of equations derived from line or point correspondences given in sections 17.2 and 17.3.

We consider still the four-view case. The results on singular values show that it is in general advisable to include all 81 equations derived from this correspondence, rather than selecting just 16 independent equations. This will avoid difficulties with near singular situations. This conclusion is supported by experimental observation. Indeed, numerical examples show that the condition of a set of equations derived from several point correspondences is substantially better when all equations are included for each point correspondence. In this context, the condition of the equation set is given by the ratio of the first (largest) to the n -th singular value, where n is the number of linearly independent equations.

Including all 81 equations rather than just 16 means that the set of equations is larger, leading to increased complexity of solution. This can be remedied as follows. The basis for the equality of the singular values is that $S_{wp} = x^i \epsilon_{ipw}$ and the other similarly defined terms are skew symmetric matrices, and hence have equal singular values. The same effect can be achieved by any other matrix S with equal singular values. We require only that the columns of S should represent lines passing through the point x (otherwise stated $S_{wp} x^w = 0_p$). Matrix S will have equal singular values if its columns are orthonormal. These conditions may be achieved with S being a 3×2 matrix. If this is done for the point in each view the total number of equations will be reduced from $3^4 = 81$ to $2^4 = 16$, and the 16 equations will be orthogonal. A convenient way of choosing S is to use a Householder matrix (see section A4.1.2(p580)) as shown below.

This discussion also applies to the trifocal tensor, allowing us to reduce the number of equations from 9 to 4, while retaining equal singular values. Summarizing this discussion for the trifocal tensor case:

- Given a point correspondence $x \leftrightarrow x' \leftrightarrow x''$ across three views, generate equations of the form

$$x^i \tilde{l}_{qx}^i \tilde{l}_{ry}^j T_{ij}^{qr} = 0_{xy} \text{ for } x, y = 1, 2$$

where $\tilde{l}_{q1}^i$ and $\tilde{l}_{q2}^i$ are two lines through x' represented by orthonormal vectors (and similarly $\tilde{l}_{ry}^j$). A convenient way to find $\tilde{l}_{qx}^i$ and $\tilde{l}_{ry}^j$ is as follows.

- Find Householder matrices h_{qx}^i and h_{ry}^j such that $x^{iq} h_{qx}^i = \delta_{3x}$ and $x^{jr} h_{ry}^j = \delta_{3y}$.
- For $x, y = 1, 2$ set $\tilde{l}_{qx}^i = h_{qx}^i$ and $\tilde{l}_{ry}^j = h_{ry}^j$.

It is evident that essentially this method will work for all the types of equations summarized in table 17.1(p417) and table 17.2(p421).

This chapter suggests that the most basic relations are the point–line–line correspondence equation $x^i l_q^j l_r^k T_i^{qr} = 0$ in the three-view case, and the line–correspondence equation $l_p^i l_q^j l_r^k Q_i^{pqr} = 0$ for four views. Indeed numerical robustness may be enhanced by reducing other correspondences to this type of correspondence, for carefully selected lines.

17.8 Closure

17.8.1 The literature

Although using a slightly different approach, this chapter summarizes previous results of [Triggs-95] and Faugeras and Mourrain [Faugeras-95a] on the derivation of multilinear relationships between corresponding image coordinates. The formulae for relations between mixed point and line correspondences are extensions of the results of [Hartley-95b, Hartley-97a].

The enumeration of the complete set of multilinear relations given in table 17.1-(p417) and table 17.2-(p421), formulae for the multifocal tensors, and the analysis of the number of independent equations derived from point correspondences, are adapted from [Hartley-95a]. A similar analysis of the multifocal tensors has also appeared in [Heyden-98].

The quadrifocal tensor was probably first discovered by [Triggs-95]. The quadrilinear constraints and their associated tensor have been described in several papers [Triggs-95, Faugeras-95a, Shashua-95b, Heyden-95b, Heyden-97c].

The double (or Grassmann–Cayley) algebra was introduced into the computer vision literature in [Carlsson-93], see also [Faugeras-95a, Faugeras-97] for further applications.

An algorithm for computing the quadrifocal tensor and an algorithm for reconstructing it based on the reduced tensor was given in ([Heyden-95b, Heyden-97c]). A later paper [Hartley-98c] refined this algorithm.

17.8.2 Notes and exercises

- (i) Determine the properties of the *affine quadrifocal tensor*, i.e. the quadrifocal tensor computed from affine camera matrices. In particular using the determinant definition of the tensor (17.21–p419), verify the number of non-zero elements given in table 17.4.
- (ii) Show that the 9 linear equations (17.11–p415) derived from a single point correspondence $x \leftrightarrow x' \leftrightarrow x''$ across three views contains 4 linearly independent equations. Furthermore, let the equations be written as $At = 0$ where A is a 9×27 matrix. Then the 4 non-zero singular values of A are equal. Unlike the four-view case, the equations resulting from different point matches are linearly independent, so n point matches produce $4n$ independent equations.
- (iii) If a canonical affine basis is chosen for the image coordinates such that three corresponding points have coordinates $(0, 0)$, $(1, 0)$, $(0, 1)$ in each view, then the resulting tensors have a simpler form. These “reduced tensors” have a greater number of zero elements than the general form of the tensors. For example, in

the case of the reduced fundamental matrix the diagonal elements are zero, and the reduced trifocal tensor has only 15 nonzero entries. Also the tensors are specified by fewer parameters, e.g. four in the case of the reduced fundamental matrix, as effectively the basis points specify the other parameters. Further details are given in [Heyden-95b, Heyden-95a].

- (iv) Show that if the four camera centres are coplanar then the quadrifocal tensor has 28 geometric degrees of freedom.

N-View Computational Methods

This chapter describes computational methods for estimating a projective or affine reconstruction from a set of images – in particular where the number of views is large.

We start with the most general case which is that of bundle adjustment for a projective reconstruction. This is then specialized to affine cameras and the important factorization algorithm introduced. A generalization of this algorithm to non-rigid scenes is given. A second specialization of bundle adjustment is then described for the case of scenes containing planes. Finally we discuss methods for obtaining point correspondences throughout an image sequence and a projective reconstruction from these correspondences.

18.1 Projective reconstruction – bundle adjustment

Consider a situation in which a set of 3D points $\mathbf{X}_j$ is viewed by a set of cameras with matrices $\mathbf{P}^i$. Denote by $\mathbf{x}_j^i$ the coordinates of the j -th point as seen by the i -th camera. We wish to solve the following reconstruction problem: given the set of image coordinates $\mathbf{x}_j^i$ find the set of camera matrices, $\mathbf{P}^i$, and the points $\mathbf{X}_j$ such that $\mathbf{P}^i \mathbf{X}_j = \mathbf{x}_j^i$. Without further restriction on the $\mathbf{P}^i$ or $\mathbf{X}_j$, such a reconstruction is a projective reconstruction, because the points $\mathbf{X}_j$ may differ by an arbitrary 3D projective transformation from the true reconstruction.

Bundle adjustment. If the image measurements are noisy then the equations $\mathbf{x}_j^i = \mathbf{P}^i \mathbf{X}_j$ will not be satisfied exactly. In this case we seek the Maximum Likelihood (ML) solution assuming that the measurement noise is Gaussian: we wish to estimate projection matrices $\tilde{\mathbf{P}}^i$ and 3D points $\tilde{\mathbf{X}}_j$ which project exactly to image points $\hat{\mathbf{x}}_j^i$ as $\hat{\mathbf{x}}_j^i = \tilde{\mathbf{P}}^i \tilde{\mathbf{X}}_j$, and also minimize the image distance between the reprojected point and detected (measured) image points $\mathbf{x}_j^i$ for every view in which the 3D point appears, i.e.

$$\min_{\tilde{\mathbf{P}}^i, \tilde{\mathbf{X}}_j} \sum_{i,j} d(\tilde{\mathbf{P}}^i \tilde{\mathbf{X}}_j, \mathbf{x}_j^i)^2 \quad (18.1)$$

where $d(\mathbf{x}, \mathbf{y})$ is the geometric image distance between the homogeneous points $\mathbf{x}$ and $\mathbf{y}$. This estimation involving minimizing the reprojection error is known as *bundle adjustment* – it involves adjusting the bundle of rays between each camera centre and

the set of 3D points (and equivalently between each 3D point and the set of camera centres).

Bundle adjustment should generally be used as a final step of any reconstruction algorithm. This method has the advantages of being tolerant of missing data while providing a true ML estimate. At the same time it allows assignment of individual covariances (or more general PDFs) to each measurement and may also be extended to include estimates of priors and constraints on camera parameters or point positions. In short, it would seem to be an ideal algorithm, except for the fact that: (i) it requires a good initialization to be provided, and (ii) it can become an extremely large minimization problem because of the number of parameters involved. We will discuss briefly these two points.

Iterative minimization. Since each camera has 11 degrees of freedom and each 3-space point 3 degrees of freedom, a reconstruction involving n points over m views requires minimization over $3n + 11m$ parameters. In fact, since entities are often over-parametrized (e.g. using 12 parameters for the homogeneous P matrix) this may be a lower bound. If the Levenberg–Marquardt algorithm is used to minimize (18.1) then matrices of dimension $(3n + 11m) \times (3n + 11m)$ must be factored (or sometimes inverted). As m and n increase this becomes extremely costly, and eventually impossible. There are several solutions to this problem:

- (i) **Reduce n and/or m .** Do not include all the views or all the points, and fill these in later by resectioning or triangulation respectively; or, partition the data into several sets, bundle adjust each set separately and then merge. Such strategies are discussed further in section 18.6.
- (ii) **Interleave.** Alternate minimizing reprojection error by varying the cameras with minimizing reprojection error by varying the points. Since each point is estimated independently given fixed cameras, and similarly each camera is estimated independently from fixed points, the largest matrix that must be inverted is the 11×11 matrix used to estimate one camera. Interleaving minimizes the same cost function as bundle adjustment, so the same solution should be obtained (provided there is a unique minimum), but it may take longer to converge. Interleaving is compared with bundle adjustment in [Triggs-00a].
- (iii) **Sparse methods.** These are described in appendix 6(p597).

Initial solution. Several methods for initialization are described in the following sections. If the problem is restricted to affine cameras then factorization (section 18.2) gives a closed form optimal solution provided points are imaged in every view. Even with projective cameras there is an (iterative) factorization method (section 18.4) available provided points are imaged in every view. If there is more information available on the data, for example that it is partly coplanar, then again a closed form solution is possible (section 18.5). Finally, hierarchical methods can be used as described in section 18.6 for the case where points are not visible in every view.

18.2 Affine reconstruction – the factorization algorithm

In this section we describe reconstruction from a set of image point correspondences for images acquired by affine cameras. As described in section 17.5.1 the reconstruction in this case is affine.

The factorization algorithm of Tomasi and Kanade [Tomasi-92] to be presented below and summarized in algorithm 18.1 has the following property:

- *Under an assumption of isotropic mean-zero Gaussian noise independent and equal for each measured point, factorization achieves a Maximum Likelihood affine reconstruction.*

This fact was first pointed out by Reid and Murray [Reid-96]. However, the method requires a measurement of each point in all views. This is a limitation in practice, since matched points may be absent in some views.

An affine camera may be characterized by having its last row equal to $(0, 0, 0, 1)$. In this section, however, we will denote it somewhat differently, separating out the translation and the pure linear transformation part of the camera map. Thus we write

$$\begin{pmatrix} x \\ y \end{pmatrix} = M \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} + \mathbf{t}$$

where M is a 2×3 matrix and $\mathbf{t}$ a 2-vector. From here on for ease of readability $\mathbf{x}$ represents an *inhomogeneous* image point $\mathbf{x} = (x, y)^T$, and $\mathbf{X}$ an *inhomogeneous* world point $\mathbf{X} = (X, Y, Z)^T$.

Our goal is to find a reconstruction to minimize geometric error in image coordinate measurements. That is, we wish to estimate cameras $\{M^i, \mathbf{t}^i\}$ and 3D points $\{\mathbf{X}_j\}$ such that the distance between the estimated image points $\hat{\mathbf{x}}_j^i = M^i \mathbf{X}_j + \mathbf{t}^i$ and measured image points $\mathbf{x}_j^i$ is minimized

$$\min_{M^i, \mathbf{t}^i, \mathbf{X}_j} \sum_{ij} \|\mathbf{x}_j^i - \hat{\mathbf{x}}_j^i\|^2 = \min_{M^i, \mathbf{t}^i, \mathbf{X}_j} \sum_{ij} \|\mathbf{x}_j^i - (M^i \mathbf{X}_j + \mathbf{t}^i)\|^2. \quad (18.2)$$

As is common in such minimization problems the translation vector $\mathbf{t}^i$ can be eliminated in advance by choosing the centroid of the points as the origin of the coordinate system. This is a consequence of the geometric fact that an affine camera maps the centroid of a set of 3D points to the centroid of their projections. Thus, if the coordinate origin is chosen as the centroid of the 3D points and of each set of image points then it follows that $\mathbf{t}^i = \mathbf{0}$. This step requires that the same n points be imaged in all views, i.e. that there are no views in which the image coordinates of any point are unknown. An analytical derivation of this result goes like this. The minimization with respect to $\mathbf{t}^i$ requires that

$$\frac{\partial}{\partial \mathbf{t}^i} \sum_{kj} \|\mathbf{x}_j^k - (M^k \mathbf{X}_j + \mathbf{t}^k)\|^2 = \mathbf{0}$$

which after a brief calculation reduces to $\mathbf{t}^i = \langle \mathbf{x}^i \rangle - M^i \langle \mathbf{X} \rangle$, where the centroids are

Objective

Given $n \geq 4$ image point correspondences over m views $\mathbf{x}_j^i$, $j = 1, \dots, n$; $i = 1, \dots, m$, determine affine camera matrices $\{\mathbf{M}^i, \mathbf{t}^i\}$ and 3D points $\{\mathbf{X}_j\}$ such that the reprojection error

$$\sum_{ij} \|\mathbf{x}_j^i - (\mathbf{M}^i \mathbf{X}_j + \mathbf{t}^i)\|^2$$

is minimized over $\{\mathbf{M}^i, \mathbf{t}^i, \mathbf{X}_j\}$, with $\mathbf{M}^i$ a 2×3 matrix, $\mathbf{X}_j$ a 3-vector, and $\mathbf{x}_j^i = (x_j^i, y_j^i)^\top$ and $\mathbf{t}^i$ are 2-vectors.

Algorithm

- (i) **Computation of translations.** Each translation $\mathbf{t}^i$ is computed as the centroid of points in image i , namely

$$\mathbf{t}^i = \langle \mathbf{x}^i \rangle = \frac{1}{n} \sum_j \mathbf{x}_j^i.$$

- (ii) **Centre the data.** Centre the points in each image by expressing their coordinates with respect to the centroid;

$$\mathbf{x}_j^i \leftarrow \mathbf{x}_j^i - \langle \mathbf{x}^i \rangle.$$

Henceforth work with these centred coordinates.

- (iii) **Construct the $2m \times n$ measurement matrix $\mathbf{W}$** from the centred data, as defined in (18.5), and compute its SVD $\mathbf{W} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$.

- (iv) Then the matrices $\mathbf{M}^i$ are obtained from the first three columns of $\mathbf{U}$ multiplied by the singular values:

$$\begin{bmatrix} \mathbf{M}^1 \\ \mathbf{M}^2 \\ \vdots \\ \mathbf{M}^m \end{bmatrix} = \begin{bmatrix} \sigma_1 \mathbf{u}_1 & \sigma_2 \mathbf{u}_2 & \sigma_3 \mathbf{u}_3 \end{bmatrix}.$$

The vectors $\mathbf{t}^i$ are as computed in step (i) and the 3D structure is read from the first three columns of $\mathbf{V}$

$$\begin{bmatrix} \mathbf{X}_1 & \mathbf{X}_2 & \dots & \mathbf{X}_n \end{bmatrix} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \end{bmatrix}^\top.$$

Algorithm 18.1. *The factorization algorithm to determine the MLE for an affine reconstruction from n image correspondences over m views (under Gaussian image noise).*

$\langle \mathbf{x}^i \rangle = \frac{1}{n} \sum_j \mathbf{x}_j^i$ and $\langle \mathbf{X} \rangle = \frac{1}{n} \sum_j \mathbf{X}_j$. The origin of the 3D frame is arbitrary, so may be chosen to coincide with the centroid $\langle \mathbf{X} \rangle$, in which case $\langle \mathbf{X} \rangle = \mathbf{0}$ and

$$\mathbf{t}^i = \langle \mathbf{x}^i \rangle. \quad (18.3)$$

It follows that if we measure the image coordinates with respect to a coordinate origin based at the centroid of the projected points, then $\mathbf{t}^i = \mathbf{0}$. Thus, we replace each $\mathbf{x}_j^i$ by $\tilde{\mathbf{x}}_j^i = \mathbf{x}_j^i - \langle \mathbf{x}^i \rangle$. Henceforth we will assume that this has been done, and work with the centred coordinates. With respect to these new coordinates $\mathbf{t}^i = \mathbf{0}$, and so (18.2) reduces to

$$\min_{\mathbf{M}^i, \mathbf{X}_j} \sum_{ij} \|\mathbf{x}_j^i - \tilde{\mathbf{x}}_j^i\|^2 = \min_{\mathbf{M}^i, \mathbf{X}_j} \sum_{ij} \|\mathbf{x}_j^i - \mathbf{M}^i \mathbf{X}_j\|^2. \quad (18.4)$$

The minimization problem now has a very simple form when written as a matrix. The *measurement* matrix W is the $2m \times n$ matrix composed of the centred coordinates of the measured image points

$$W = \begin{bmatrix} x_1^1 & x_2^1 & \dots & x_n^1 \\ x_1^2 & x_2^2 & \dots & x_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ x_1^m & x_2^m & \dots & x_n^m \end{bmatrix}, \quad (18.5)$$

Since each $x_j^i = M^i X_j$, the complete set of equations may be written as

$$W = \begin{bmatrix} M^1 \\ M^2 \\ \vdots \\ M^m \end{bmatrix} \begin{bmatrix} X_1 & X_2 & \dots & X_n \end{bmatrix},$$

In the presence of noise this equation will not be satisfied exactly, so instead we seek a matrix $\hat{W}$ as close as possible to W in Frobenius norm, such that $\hat{W}$ may be decomposed as

$$\hat{W} = \begin{bmatrix} \hat{x}_1^1 & \hat{x}_2^1 & \dots & \hat{x}_n^1 \\ \hat{x}_1^2 & \hat{x}_2^2 & \dots & \hat{x}_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ \hat{x}_1^m & \hat{x}_2^m & \dots & \hat{x}_n^m \end{bmatrix} = \begin{bmatrix} M^1 \\ M^2 \\ \vdots \\ M^m \end{bmatrix} \begin{bmatrix} X_1 & X_2 & \dots & X_n \end{bmatrix}, \quad (18.6)$$

In this case it may be verified that

$$\|W - \hat{W}\|_F^2 = \sum_{ij} (w_{ij} - \hat{w}_{ij})^2 = \sum_{ij} \|x_j^i - \hat{x}_j^i\|^2 = \sum_{ij} \|x_j^i - M^i X_j\|^2$$

Comparing this with (18.4) we find that minimizing the required geometric error is equivalent to finding such a $\hat{W}$ as close as possible to W in Frobenius norm.

Note that a matrix $\hat{W}$ satisfying (18.6) is the product of a $2m \times 3$ *motion matrix* $\hat{M}$, and a $3 \times n$ *structure matrix* $\hat{X}$; consequently $\hat{W} = \hat{M}\hat{X}$ has rank 3. In other words we seek a rank 3 matrix which is closest to W in Frobenius norm. Such a matrix may be determined by the SVD of W truncated to rank 3. In more detail, if the SVD of $W = UDV^T$ then $\hat{W} = U_{2m \times 3} D_{3 \times 3} V_{3 \times n}^T$ is the rank 3 matrix which is closest to W in the Frobenius norm, where $U_{2m \times 3}$ consists of the first 3 columns of U , $V_{3 \times n}^T$ consists of the first 3 rows of V^T , and $D_{3 \times 3}$ is the diagonal matrix containing the first 3 singular values, $D_{3 \times 3} = \text{diag}(\sigma_1, \sigma_2, \sigma_3)$.

Note that the choice of $\hat{M}$ and $\hat{X}$ is not unique. For example $\hat{M}$ may be chosen as $\hat{M} = U_{2m \times 3} D_{3 \times 3}$ and $\hat{X} = V_{3 \times n}^T$, or as $\hat{M} = U_{2m \times 3}$, $\hat{X} = D_{3 \times 3} V_{3 \times n}^T$ since in either case $\hat{W} = \hat{M}\hat{X} = U_{2m \times 3} D_{3 \times 3} V_{3 \times n}^T$.

Affine ambiguity. In fact for any such choice there is an additional ambiguity since an arbitrary 3×3 rank 3 matrix A may be inserted in the decomposition as $\hat{W} = \hat{M} A A^{-1} \hat{X} = (\hat{M}A)(A^{-1}\hat{X})$. This means that the camera matrices M^i , which are obtained from $\hat{M}$, and

the 3D points X_j , which are obtained from $\hat{X}$, are determined up to multiplication by a common matrix A . In other words the MLE reconstruction is affine.

This affine reconstruction may be upgraded to a metric reconstruction by supplying metric information on the scene as described in section 10.4.2(p272), or by using auto-calibration methods as described in chapter 19, or a combination of the two. Note that in the case of affine cameras only three internal parameters need be specified (compared to five for projective cameras) and the auto-calibration task is correspondingly simpler.

18.2.1 Affine multiple view tensors

The factorization algorithm provides an optimal method for computing the affine multiview tensors from image point correspondences. These tensors are the affine fundamental matrix, affine trifocal tensor, and affine quadrifocal tensor. In each case the algorithm determines the camera matrices up to an overall affine ambiguity. The tensors may then be computed directly from the camera matrices (as for instance in chapter 17). The affine ambiguity of 3-space is irrelevant when computing the tensors since they are unaffected by affine transformations of 3-space. In fact it is not necessary to compute the full SVD of W since only the U part of the decomposition is required. If the number of points n is large compared with the number of views then very great savings can be made in the computation of the SVD by not determining V (see table A4.1(p587)).

An alternative to using the SVD is to use the eigenvalue decomposition of WW^T , since $WW^T = (UDV^T)(UDV^T)^T = UD^2U^T$. In the case of three views (computation of the trifocal tensor) the matrix WW^T has dimension only 9×9 . Thus this approach can mean significant savings. However, it is numerically inferior, since forming WW^T causes the condition number of the matrix to be squared (see the discussion of SVD in [Golub-89]). Since we need just the three largest eigenvectors, that may not be such a problem in this case. However, the savings of this approach will not be so great given an implementation of the SVD that avoids computing V .

The factorization method may be used to compute any of the multiple-view affine tensors. For the affine fundamental matrix, however, algorithm 14.1(p351) described in chapter 14 is more direct. The results of both the methods are identical.

18.2.2 Triangulation and reprojection using subspaces

The factorization algorithm also provides an optimal method for computing the images of new points or of points not observed in all views. Again the affine ambiguity of 3-space is irrelevant.

A column of W is the set of all corresponding image points for the point X_j and is referred to as a point's *trajectory*. The rank 3 decomposition (18.6) of W as $W = M\hat{X}$ shows that all trajectories lie in a 3 dimensional subspace. In particular the trajectory (i.e. all image projections) of a new point X may be obtained as $M\hat{X}$. This is simply a linear weighting of the three columns of M .

Suppose we have observed a new point X in some (not all) views, and wish to predict its projection in the other views. This is carried out in two steps: first triangulation to find the pre-image $\hat{X}$, and then reprojection as $M\hat{X}$ to generate its image in

all views. Note that the projected points will not coincide exactly with the measured (noisy) points. In the triangulation step we wish to find the point $\mathbf{X}$ that minimizes re-projection error, and this corresponds to finding the point in the linear subspace spanned by the columns of $\tilde{\mathbf{M}}$ closest to the trajectory. This closest point is found by projecting the trajectory onto the subspace (in a similar manner to algorithm 4.7(p130)).

In more detail suppose we have computed a set of affine cameras $\{\mathbf{M}^i, \mathbf{t}^i\}$ then the triangulation problem may be solved linearly for any number of views. The image points $\mathbf{x}^i = \mathbf{M}^i \mathbf{X} + \mathbf{t}^i$ give a pair of linear equations $\mathbf{M}^i \mathbf{X} = \mathbf{x}^i - \mathbf{t}^i$ in the entries of $\mathbf{X}$. Given sufficiently many such equations (arising from known values of $\mathbf{x}^i$) one can find the linear least-squares solution for $\mathbf{X}$, using algorithm A5.1(p589), the pseudo-inverse (see result A5.1(p590)) or algorithm A5.3(p591). Note that if the data $\mathbf{x}^i$ is centred using the same transformation applied in step (ii) of algorithm 18.1, then the translation vectors $\mathbf{t}^i$ in the affine triangulation method may be taken to be zero.

In practice triangulation and reprojection provides a method of 'filling in' points that are missed during tracking or multiple view matching.

18.2.3 Affine Reconstruction by Alternation

Suppose a set of image coordinates $\mathbf{x}_j^i$ are given as in algorithm 18.1, and we wish to perform affine reconstruction. We have seen that affine triangulation may be carried out linearly. Thus, if the affine camera matrices represented by $\{\mathbf{M}^i, \mathbf{t}^i\}$ are known, then the optimal point positions $\mathbf{X}_j$ may be computed by a linear least-squares method such as the normal equations method of algorithm A5.3(p591).

Conversely, if the points $\mathbf{X}_j$ are known, then the equations $\mathbf{x}_j^i = \mathbf{M}^i \mathbf{X}_j + \mathbf{t}^i$ are linear in $\mathbf{M}^i$ and $\mathbf{t}^i$. So it is once again possible simply to solve for $\{\mathbf{M}^i, \mathbf{t}^i\}$ by linear least-squares.

This suggests a method of affine reconstruction in which linear least-squares methods are used to solve alternately for the points $\mathbf{X}_j$ and the cameras $\{\mathbf{M}^i, \mathbf{t}^i\}$. This method of alternation is not to be recommended as a general method for reconstruction, or for solving optimization problems in general. In the case of affine reconstruction, however, it can be proven to converge rapidly to the optimal solution, starting from a random starting point. This method of affine reconstruction has the advantage of working with missing data, or with covariance-weighted data which algorithm 18.1 will not, though in the missing data or covariance-weighted case, global optimal convergence is not guaranteed in all cases.

18.3 Non-rigid factorization

Throughout the book it has been assumed that we are viewing a rigid scene and that only the relative motion between the camera and scene is to be modelled. In this section we relax this assumption and consider the problem of recovering a reconstruction for a deforming object. It will be shown that if the deformation is modelled as a linear combination over basis shapes then the reconstruction *and the basis shapes* may be recovered with a simple modification of the factorization algorithm of section 18.2.

An example where this type of situation arises is in a sequence of images of a per-




Fig. 18.1. **Shape basis.** A face template is represented by N equally spaced 2D points (here $N = 140$). The central face of the seven is the mean shape and the faces to the left or right are generated by adding or subtracting, respectively, the basis shape that accounts for the maximum variation in the training set. In this case the basis spans expressions from surprised to disgruntled. Facial expressions are learnt by tracking the face of an actor with the template whilst he changes expression but does not vary his head pose. Each frame of the training sequence generates a set of N 2D points, and the coordinates of these are rewritten as a $2N$ -vector. A $2N \times f$ matrix is then composed from these vectors, where f is the number of training frames, and the basis shapes are computed from the singular vectors of this matrix. Figure courtesy of Richard Bowden.

son's head which moves and also changes expression. The motion of the head may be modelled as a rigid rotation, and the change of expression relative to the fixed head may be modelled as a linear combination over basis sets. For example the mouth outline may be represented by a set of points.

Suppose the set of n scene points $\mathbf{X}_j$ may be represented as a linear combination of l basis shapes $\mathbf{B}_k$ so that at a particular time i :

$$\begin{bmatrix} \mathbf{X}_1^i & \mathbf{X}_2^i & \dots & \mathbf{X}_n^i \end{bmatrix} = \sum_{k=1}^l \alpha_k^i \begin{bmatrix} \mathbf{B}_{1k} & \mathbf{B}_{2k} & \dots & \mathbf{B}_{nk} \end{bmatrix} = \sum_k \alpha_k^i \mathbf{B}_k$$

where here both the scene points $\mathbf{X}_j^i$ and the basis points $\mathbf{B}_{jk}$ are inhomogeneous points represented by 3-vectors, and $\mathbf{B}_k$ is a $3 \times n$ matrix. Typically the number of basis shapes, l , is much smaller than the number of points, n . The coefficients α_k^i may be different at each time i , and the resulting differing combination of basis shapes generates the deformation. An example is shown in figure 18.1.

In the forward model of image generation each view i is acquired by an affine camera and gives the image points

$$\mathbf{x}_j^i = \mathbf{M}^i \sum_k \alpha_k^i \mathbf{B}_{jk} + \mathbf{t}^i.$$

It will again be assumed that image point matches are available for all views. Our goal is to estimate cameras $\{\mathbf{M}^i, \mathbf{t}^i\}$ and 3D structure $\{\alpha_k^i, \mathbf{B}_{jk}\}$ from the measured image points $\{\mathbf{x}_j^i\}$, such that the distance between the estimated image points $\hat{\mathbf{x}}_j^i = \mathbf{M}^i \sum_k \alpha_k^i \mathbf{B}_{jk} + \mathbf{t}^i$ and measured image points is minimized

$$\min_{\mathbf{M}^i, \mathbf{t}^i, \alpha_k^i, \mathbf{B}_{jk}} \sum_{ij} \left\| \mathbf{x}_j^i - \hat{\mathbf{x}}_j^i \right\|^2 = \min_{\mathbf{M}^i, \mathbf{t}^i, \alpha_k^i, \mathbf{B}_{jk}} \sum_{ij} \left\| \mathbf{x}_j^i - \left(\mathbf{M}^i \sum_k \alpha_k^i \mathbf{B}_{jk} + \mathbf{t}^i \right) \right\|^2.$$

As in affine factorization the translation may be eliminated by centring the measured image points, and it will be assumed from here on that this has been done. Then the

problem reduces to

$$\min_{M^i, \alpha_k^i, B_{jk}} \sum_{ij} \left\| \mathbf{x}_j^i - \bar{\mathbf{x}}_j^i \right\|^2 = \min_{M^i, \alpha_k^i, B_{jk}} \sum_{ij} \left\| \mathbf{x}_j^i - M^i \sum_k \alpha_k^i B_{jk} \right\|^2. \quad (18.7)$$

The complete set of equations $\dot{\mathbf{x}}_j^i = M^i \sum_k \alpha_k^i B_{jk}$ may be written

$$\tilde{W} = \begin{bmatrix} M^1 (\alpha_1^1 B_1 + \alpha_2^1 B_2 + \dots + \alpha_l^1 B_l) \\ M^2 (\alpha_1^2 B_1 + \alpha_2^2 B_2 + \dots + \alpha_l^2 B_l) \\ \vdots \\ M^m (\alpha_1^m B_1 + \alpha_2^m B_2 + \dots + \alpha_l^m B_l) \end{bmatrix} = \begin{bmatrix} \alpha_1^1 M^1 & \alpha_2^1 M^1 & \dots & \alpha_l^1 M^1 \\ \alpha_1^2 M^2 & \alpha_2^2 M^2 & \dots & \alpha_l^2 M^2 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^m M^m & \alpha_2^m M^m & \dots & \alpha_l^m M^m \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ \vdots \\ B_l \end{bmatrix} \quad (18.8)$$

This rearrangement shows that the $2m \times n$ matrix $\tilde{W}$ may be decomposed as a product of a $2m \times 3l$ motion matrix $\tilde{M}$ and $3l \times n$ structure matrix $\tilde{B}$, and consequently $\tilde{W}$ has maximum rank $3l$.

As in rigid factorization a rank $3l$ decomposition may be obtained from the measurement matrix $\tilde{W}$ by truncating the SVD of $\tilde{W}$ to rank $3l$. Also, as in rigid factorization, the decomposition $\tilde{W} = \tilde{M}\tilde{B}$ is not uniquely defined since an arbitrary $3l \times 3l$ rank $3l$ matrix A may be inserted in the decomposition as $\tilde{W} = \tilde{M}A A^{-1} \tilde{B} = (\tilde{M}A)(A^{-1} \tilde{B})$. In the rigid case this resulted in a straightforward affine ambiguity in the reconstruction. However, in the non-rigid case there is the additional requirement that the motion matrix has the replicated block structure of (18.8), and we return to this below. This block structure is not required for determining a point's image motion, as will now be discussed.

18.3.1 Subspaces and tensors

In the case of rigid factorization (18.6), as discussed in section 18.2.2, the trajectories lie in a 3 dimensional subspace, and any trajectory may be generated as a linear combination of the columns of $\tilde{M}$ (the $2m \times 3$ motion matrix). Similarly in the case of non-rigid factorization, (18.8), the trajectories lie in a $3l$ dimensional subspace, and any trajectory may be generated as a linear combination of the columns of $\tilde{M}$ (the $2m \times 3l$ motion matrix).

Suppose we observe a new point in a subset of the views, how many images are required before its position can be predicted in all the other views? This is simply a question of triangulation: in rigid factorization a 3-space point has 3 degrees of freedom, and must be observed in two views to obtain the necessary 3 measurements. In non-rigid factorization we need to specify $3l$ degrees of freedom (the number of rows in the $\tilde{B}$ matrix), and this requires $3l/2$ images. For example, if $l = 2$ the subspace is six dimensional (the columns of the $\tilde{B}$ matrix are 6-vectors), and given the image position in three views, the image position in all views is then determined by an analogue of affine triangulation (section 18.2.2) even though the object is deforming.

Independently moving objects. Low-rank factorization methods also arise when there are independently moving objects in the scene. For instance suppose the scene is divided into two objects, each moving independently of the other, and viewed by




Fig. 18.2. **Non-rigid motion sequence.** Top row: alternate frames from a sequence in which a giraffe gracefully walks and flexes its neck, whilst the camera pans to match its speed. Bottom row: point tracks showing the motion over (a) the 10 previous, and (b) the 10 forthcoming frames. These tracks are computed using non-rigid factorization and lie in a six dimensional subspace. Note the very different trajectories of the (rigid) background from the (deforming) foreground. Yet these are all spanned by the six basis vectors of the motion matrix. The rank can be accounted for as follows: the sequence motion is effectively that of two planes of points moving independently relative to the camera. The background is a rigid object represented by a plane and contributes 2 to the rank. The giraffe in the foreground is represented as a non-rigid object by a set of $l = 2$ planar basis shapes and contributes 4 to the rank. Figures courtesy of Andrew Fitzgibbon and Aeron Morgan.

a moving affine camera. In this case, the columns of the measurement matrix corresponding to points on one object will have rank 3, and those corresponding to the other object will also have rank 3. The total rank of the measurement matrix will be 6. In degenerate configurations in which one object's points all lie in a plane, its contribution to the rank will be only 2. This multibody factorization problem has been studied in some depth in [Costeira-98].

An example of point tracks residing in a low dimensional subspace is shown in figure 18.2.

The existence of the analogue of the fundamental matrix and trifocal tensor depends on the dimension of the subspace. For example suppose the subspace has odd dimension (e.g. $l = 3$ so it is 9 dimensional) then given point measurements in $\lceil 3l/2 \rceil$ views (e.g. 4 views) the corresponding point in any other view is constrained to a line, the analogue of an epipolar line, since there is one fewer measurement than degrees of freedom of the subspace. However, if the dimension is even (e.g. $l = 2$ so it is 6 dimensional) then given point measurements in $l/2$ views (e.g. 3 views) the corresponding point in any other view is completely determined. Multi-view tensors can be built for

the non-rigid $l > 1$ subspaces using methods similar to those developed in chapter 17 for $l = 1$ 3-space points.

18.3.2 Recovering the camera motion

In rigid-factorization the camera matrices are obtained relatively easily from the motion matrix $\hat{M}$ – all that is required is to remove a global affine ambiguity specified by a 3×3 matrix A as described on page 438.

In the non-rigid case, the analogous problem is not so straightforward. It is a simple matter to obtain the motion matrix:

- (i) Construct the $2m \times n$ measurement matrix W from the centred data, as defined in (18.5), and compute its SVD $W = UDV^T$.
- (ii) Then the motion matrix $\hat{M}$ is obtained from the first $3l$ columns of U multiplied by the singular values as $\hat{M} = \begin{bmatrix} \sigma_1 \mathbf{u}_1 & \sigma_2 \mathbf{u}_2 & \dots & \sigma_{3l} \mathbf{u}_{3l} \end{bmatrix}$,

but the matrix obtained by this route will *not* in general have the block structure of (18.8). As in the case of rigid-factorization the motion matrix is determined up to post-multiplication by a matrix A , which here is $3l \times 3l$. The task is then to determine A such that $\hat{M}A$ has the required block structure of (18.8) *and* also to remove the usual affine ambiguity such that each block conforms to any available constraints on the camera calibration (for example identical internal parameters over all views).

Various methods for determining A have been investigated (see [Brand-01, Torresani-01]), but these do not impose the full block structure, and currently there is not a satisfactory solution to this problem. Once an initial solution has been obtained by some means, then the correct form can be imposed by bundle adjustment of (18.7).

18.4 Projective factorization

The affine factorization method does not apply directly to projective reconstruction. It was observed in [Sturm-96], however, that if one knows the “projective depth” of each of the points then the structure and camera parameters may be estimated by a simple factorization algorithm similar in style to the affine factorization algorithm.

Consider a set of image points $\mathbf{x}_j^i = P^i \mathbf{X}_j$. This equation representing the projective mapping is to be interpreted as true only up to a constant factor. Writing these constant factors explicitly, we have $\lambda_j^i \mathbf{x}_j^i = P^i \mathbf{X}_j$. In this equation, and henceforth in the description of the projective factorization algorithm, the notation $\mathbf{x}_j^i$ means the 3-vector $(x_j^i, y_j^i, 1)^T$ representing an image point. Thus the third coordinate is equal to unity, and x_j^i and y_j^i are the actual measured image coordinates. Provided that each point is visible in every view, so that $\mathbf{x}_j^i$ is known for all i, j , the complete set of equations may be written as a single matrix equation as follows:

$$\begin{bmatrix} \lambda_1^1 \mathbf{x}_1^1 & \lambda_2^1 \mathbf{x}_2^1 & \dots & \lambda_n^1 \mathbf{x}_n^1 \\ \lambda_1^2 \mathbf{x}_1^2 & \lambda_2^2 \mathbf{x}_2^2 & \dots & \lambda_n^2 \mathbf{x}_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_1^m \mathbf{x}_1^m & \lambda_2^m \mathbf{x}_2^m & \dots & \lambda_n^m \mathbf{x}_n^m \end{bmatrix} = \begin{bmatrix} P^1 \\ P^2 \\ \vdots \\ P^m \end{bmatrix} [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n] \quad (18.9)$$

Objective

Given a set of n image points seen in m views:

$$\mathbf{x}_j^i : i = 1, \dots, m, j = 1, \dots, n$$

compute a projective reconstruction.

Algorithm

- (i) Normalize the image data using isotropic scaling as in section 4.4.4(p107).
- (ii) Start with an initial estimate of the projective depths λ_j^i . This may be obtained by techniques such as an initial projective reconstruction, or else by setting all $\lambda_j^i = 1$.
- (iii) Normalize the depths λ_j^i by multiplying rows and columns by constant factors. One method is to do a pass setting the norms of all rows to 1, then a similar pass on columns.
- (iv) Form the $3m \times n$ measurement matrix on the left of (18.9), find its nearest rank-4 approximation using the SVD and decompose to find the camera matrices and 3D points.
- (v) **Optional iteration.** Reproject the points into each image to obtain new estimates of the depths and repeat from step (ii).

Algorithm 18.2. *Projective reconstruction through factorization.*

This equation is true only if the correct weighting factors λ_j^i are applied to each of the measured points $\mathbf{x}_j^i$. For the present, let us assume that these depths are known. As with the affine factorization algorithm, we would like the matrix on the left – denote it by W – to have rank 4, since it is the product of two matrices with 4 columns and rows respectively. The actual measurement matrix can be corrected to have rank 4 by using the SVD. Thus, if $W = UDV^T$, all but the first four diagonal entries of D are set to zero resulting in $\hat{D}$. The corrected measurement matrix is $\hat{W} = U\hat{D}V^T$. The camera matrices are retrieved from $[P_1^T, P_2^T, \dots, P_m^T]^T = U\hat{D}$ and the points from $[X_1, X_2, \dots, X_n] = V^T$. Note that this factorization is not unique, and in fact we may interpose an arbitrary 4×4 projective transformation H and its inverse between the two matrices on the right of (18.9), reflecting the fact that reconstruction has a projective ambiguity.

The steps of the projective factorization method are summarized in algorithm 18.2.

18.4.1 Choosing the depths

The weighting factors λ_j^i are called the *projective depths* of the points. The justification of this terminology is the relation of these λ_j^i to the actual depths if camera matrices are known in a Euclidean frame. Refer to section 6.2.3(p162) and in particular figure 6.6-(p162). The main difficulty with this projective factorization algorithm is that we need to know these projective depths up front, but we do not have this knowledge. There are various techniques to estimate the depths.

- (i) Start with an initial projective reconstruction obtained by other means, such as those discussed in section 18.6 below. Then compute λ_j^i by reprojecting the 3D points.
- (ii) Start with initial depths all equal to 1, compute the reconstruction and reproject to obtain a new estimate of the depths. This step may be repeated to obtain

improved estimates. However, there is no guarantee that the procedure will converge to a global minimum.

The original paper [Sturm-96] gives a method of computing the depths by stringing together pairwise estimates of the depth obtained from the fundamental matrix, or the trifocal tensor. This method is quite similar to obtaining an initial projective reconstruction by stringing together triples of images (see section 18.6), whilst ensuring that the scale factors are consistent for a common projective reconstruction.

18.4.2 What is being minimized?

In the case of noise, or incorrect values for λ_j^i , the equations (18.9) are not satisfied exactly. We determine a corrected measurement matrix $\tilde{W}$ that is closest to W in Frobenius norm, subject to having rank 4. Denoting the entries of this matrix as $\tilde{\lambda}_j^i \tilde{x}_j^i$, then the computed solution minimizes the expression

$$\|W - \tilde{W}\|^2 = \sum_{ij} \|\lambda_j^i x_j^i - \tilde{\lambda}_j^i \tilde{x}_j^i\|^2 = \sum_{ij} (\lambda_j^i x_j^i - \tilde{\lambda}_j^i \tilde{x}_j^i)^2 + (\lambda_j^i y_j^i - \tilde{\lambda}_j^i \tilde{y}_j^i)^2 + (\lambda_j^i - \tilde{\lambda}_j^i)^2 \quad (18.10)$$

Because of the last term, at a minimum $\tilde{\lambda}_j^i$ must be close to λ_j^i . Assuming they are equal, (18.10) reduces to $\sum_{ij} (\lambda_j^i)^2 \|x_j^i - \tilde{x}_j^i\|^2$. Noting that $\|x_j^i - \tilde{x}_j^i\|$ is the geometric distance between the measured and estimated points, what is being minimized is a weighted sum-of-squares geometric distance, where each point is being weighted by λ_j^i . If all the geometric depths λ_j^i are close to equal, then the factorization method minimizes an approximation to geometric distance scaled by the common value of λ_j^i .

18.4.3 Normalizing the depths

Projective depths as defined here are not unique. Indeed suppose that $\lambda_j^i x_j^i = P^i X_j$. If we replace P^i by $\alpha^i P^i$ and X_j by $\beta_j X_j$, then we find that

$$(\alpha^i \beta_j \lambda_j^i) x_j^i = (\alpha^i P^i)(\beta_j X_j) .$$

In other words, the projective depths λ_j^i may be replaced by multiplying the i -th row of (18.9) by a factor α^i and the j -th column by a factor β_j . In the light of the previous paragraph, the closer all λ_j^i are to unity, the more exactly the error expression represents geometric distance. Therefore, it is advantageous to renormalize the values of the λ_j^i so that they are as close to unity as possible, by multiplying rows and columns of the measurement matrix by constant values α^i and β_j . A simple heuristic manner of doing this is to multiply each row by a factor α^i so that it has unit norm, followed by a similar pass normalizing the columns. The row and column passes may be iterated.

18.4.4 Normalizing the image coordinates

As with most numerical algorithms involving homogeneous representations of image coordinates described in this book, it is important to normalize the image coordinates. A reasonable scheme is the isotropic normalization method described in section 4.4.4 (p107). One can see the necessity of normalization quite clearly in this case. Consider two image points $x = (200, 300, 1)^T$ and $\tilde{x} = (250, 375, 1)^T$. Obviously these points

are very far apart in a geometric sense. However, the error expression (18.10) measures not geometric error, but the distance between homogeneous vectors $\|\lambda \mathbf{x} - \hat{\lambda} \hat{\mathbf{x}}\|$. Choosing $\lambda = 1.25$ and $\hat{\lambda} = 1.0$, the error is $\|(250, 375, 1.25)^\top - (250, 375, 1)^\top\|$ which is proportionally quite small. On the other hand, the distance between points $\mathbf{x} = (200, 300, 1)^\top$ and $\hat{\mathbf{x}} = (199, 301, 1)^\top$, which are much closer geometrically can not be made so small by choice of λ and $\hat{\lambda}$ (except for small values). The reader may observe that if the points are scaled down by a factor of 200, then this anomalous situation no longer occurs. In short, with normalized coordinates, the error is a closer approximation to geometric error.

18.4.5 When is the assumption $\lambda_j^i = 1$ reasonable?

According to result 6.1(p162), if camera matrices are normalized such that $p_{31}^2 + p_{32}^2 + p_{33}^2 = 1$, and 3D points are normalized to have last coordinate $\tau = 1$, then λ_{ij} defined by $\lambda_{ij}^i(x_j^i, y_j^i, 1) = P^i \mathbf{X}_j$ are the true depths of the points from the camera in a Euclidean frame. If all points are equidistant from the cameras throughout a sequence then we may reasonably assume that each $\lambda_j^i = 1$, for (18.9) will have at least the solution where P^i and $\mathbf{X}_j$ are the true cameras and points, normalized in the manner just stated. More generally, suppose that points are located at different depths, but each point $\mathbf{X}_j$ remains at approximately the same depth d_j from the cameras through the whole sequence. In this case a solution will exist with all $\lambda_j^i = 1$ in which the computed $\mathbf{X}_j = d_j^{-1}(X_j, Y_j, Z_j, 1)^\top$. Similarly, by allowing multiplication of the camera matrices by a factor, we find

- *If the ratios of true depths of the different 3D points $\mathbf{X}_j$ remain approximately constant during a sequence, then the assumption $\lambda_j^i = 1$ is a good first approximation to projective depth.*

This is for instance the case of an aerial image camera pointing straight down from constant altitude.

18.5 Projective reconstruction using planes

It was seen in section 17.5.2 that if four points visible in each view are known to be coplanar then the computation of the multifocal tensors relating the image points becomes significantly more simple. A major advantage is that a tensor satisfying all its constraints may be computed using a linear algorithm. We now continue with that particular line of investigation, and show that the use of linear techniques extends to estimation of motion and structure for any number of views.

The condition that four of the image correspondences are derived from coplanar points is equivalent to knowing the homographies between the images induced by a plane in space, since a homography may be computed from the four points. It is only the homographies that are important in the following approach. These homographies may be computed from four or more point correspondences, or line correspondences, or estimated directly from the images by direct correlation methods.

What do the plane-plane homographies tell us? The key to projective reconstruction using planes is the observation that knowledge of homographies between the images means we know the first 3×3 part of the camera matrices:

$$P = \left[M \begin{vmatrix} t_1 \\ t_2 \\ t_3 \end{vmatrix} \right]$$

Hence, it remains only to compute their last columns, namely the vectors t .

Since we are interested at this point only in obtaining a *projective* reconstruction of the scene, we may suppose that the plane inducing the homographies is the plane at infinity, with points $X_j = (X_j, Y_j, Z_j, 0)^T$. Camera matrices may be written in the form $P^i = [M^i | t^i]$, where M is a 3×3 matrix and t^i is a column vector. A reasonable assumption is that the camera centres do not lie on the plane inducing the homographies (for otherwise the homographies will be degenerate). This means that the matrix M is non-singular. For simplicity, the first camera may be assumed to have the form $P^1 = [I | 0]$, where I is the identity matrix.

Now, if x_j^i is the point in image i corresponding to the 3D point $X_j = (X_j, Y_j, Z_j, 0)^T$ lying on the homography-induced plane, then

$$x_j^1 = P^1(X_j, Y_j, Z_j, 0)^T = (X_j, Y_j, Z_j)^T$$

whereas

$$x_j^i = M^i(X_j, Y_j, Z_j)^T = M^i x_j^1.$$

Thus M^i represents the homography from the first image to the i -th image induced by the plane. Conversely, if M^i is the known plane-induced homography that maps a point in the first image to its matching point in the i -th image, then the set of camera matrices can be assumed to have the form $P^i = [M^i | t^i]$, where the M^i are known and their scale is fixed, but the final columns t^i are not.

Known camera orientation. We have just shown that knowledge of homographies implies the knowledge of the left-hand 3×3 submatrix of each camera matrix. The same will hold if we know the orientation (and calibration) of all the cameras. For instance, a reasonable approach to reconstruction, knowing the calibration of each camera, is to estimate the orientation of each camera separately from the translation (for example from two or more scene vanishing points). Once the orientation (R^i) and calibration (K^i) of each camera is known, the left-hand block of each camera matrix is $K^i R^i$.

18.5.1 Direct solution for structure and translation

We describe two separate methods for computation of the projective structure given plane-induced homographies between images. The first method solves for the 3D points and the camera motion simultaneously by solving a single linear system. Suppose point $X = (x, y, z, 1)^T$ is not on the plane at infinity, that is, the plane inducing the homographies.

The equation for point projection is

$$\lambda \mathbf{x} = \mathbf{P}\mathbf{X} = [\mathbf{M}|\mathbf{t}]\mathbf{X} = [\mathbf{M}|\mathbf{t}] \begin{pmatrix} \tilde{\mathbf{X}} \\ 1 \end{pmatrix}$$

where the (unknown) scale factor λ has been explicitly written. More precisely, we may write

$$\lambda \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{bmatrix} \mathbf{m}_1^T & t_1 \\ \mathbf{m}_2^T & t_2 \\ \mathbf{m}_3^T & t_3 \end{bmatrix} \begin{pmatrix} \tilde{\mathbf{X}} \\ 1 \end{pmatrix} = \begin{pmatrix} \mathbf{m}_1^T \tilde{\mathbf{X}} + t_1 \\ \mathbf{m}_2^T \tilde{\mathbf{X}} + t_2 \\ \mathbf{m}_3^T \tilde{\mathbf{X}} + t_3 \end{pmatrix}$$

where $\mathbf{m}_i^T$ is the i -th row of the matrix $\mathbf{M}$.

The unknown scale factor λ may be eliminated by taking the vector product of the two sides of this equation, resulting in

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \times \begin{pmatrix} \mathbf{m}_1^T \tilde{\mathbf{X}} + t_1 \\ \mathbf{m}_2^T \tilde{\mathbf{X}} + t_2 \\ \mathbf{m}_3^T \tilde{\mathbf{X}} + t_3 \end{pmatrix} = \mathbf{0}.$$

This provides two independent equations

$$\begin{aligned} x(\mathbf{m}_3^T \tilde{\mathbf{X}} + t_3) - (\mathbf{m}_1^T \tilde{\mathbf{X}} + t_1) &= 0 \\ y(\mathbf{m}_3^T \tilde{\mathbf{X}} + t_3) - (\mathbf{m}_2^T \tilde{\mathbf{X}} + t_2) &= 0 \end{aligned}$$

which are linear in the unknowns $\tilde{\mathbf{X}} = (x, y, z)^T$ and $\mathbf{t} = (t_1, t_2, t_3)^T$. The equations may be written as

$$\begin{bmatrix} x\mathbf{m}_3^T - \mathbf{m}_1^T & -1 & 0 & x \\ y\mathbf{m}_3^T - \mathbf{m}_2^T & 0 & -1 & y \end{bmatrix} \begin{pmatrix} \tilde{\mathbf{X}} \\ t_1 \\ t_2 \\ t_3 \end{pmatrix} = \mathbf{0}.$$

Thus, each measured point $\mathbf{x}_j^i = \mathbf{P}^i \mathbf{X}_j$ generates a pair of equations, and with m views involving n points a $2nm$ set of equations in $3n + 3m$ unknowns is generated in this way. These equations may be solved by linear or linear least-squares techniques to obtain the structure and motion.

A few remarks about this method are offered.

- (i) In contrast to factorization methods (section 18.2) we do not need all points to be visible in all views. Only equations corresponding to the measured points are used.
- (ii) Since it is assumed that points have final coordinate equal to one, it is necessary to exclude points that lie on the plane at infinity (the plane inducing the homography) which have final coordinate equal to zero. A test to detect points lying on or close to the plane is necessary.
- (iii) Both points and cameras are computed at once. For a large number of points and cameras this may be a very large estimation problem. However, if the point tracks have a banded form, then sparse solution techniques may be used to solve the equation set efficiently, as in section A6.7(p613).


This method and its implementation are discussed in depth in [Rother-01, Rother-03]. The details given here are different from those given in [Rother-01], where the structure and motion computation is carried out in a specific projective frame related to the matched points on the plane, involving a coordinate change in the images.

18.5.2 Direct motion estimation

The second method for planar reconstruction knowing homographies solves for the camera matrices first and subsequently computes the point positions.

We start from the set of camera matrices which again can be assumed to have the form $P^i = [H^i | t^i]$, where the H^i are known and their scale is fixed, but the final columns t^i are not. We may assume that $P^1 = [I | 0]$, so that $t^1 = 0$. The set of all remaining t^i have $3m - 4$ degrees of freedom, since the t^i are defined only up to a common scale. Now assume that several point or line correspondences across two or more views are known (three views are required for lines). These correspondences must derive from 3D points or lines that do not lie in the reference plane (used to compute the H^i). Each point correspondence across two views leads to a linear equation in the entries of the fundamental matrix. Similarly, correspondences of points or lines across three or four views lead to linear equations in the entries of the trifocal or quadrifocal tensor.

The key point (as explained in section 17.5.2) is that we may express the entries of the fundamental matrix (or trifocal or quadrifocal tensor) linearly in the entries of the vectors t^i . Therefore each linear relation induced by a point or line correspondence may be related back to a linear relationship in terms of the entries of the t^i . Thus, for example, a correspondence across views i , j and k gives rise to a set of linear equations in the entries of the three vectors t^i , t^j and t^k . A set of correspondences across many views can be broken down into correspondences across sets of consecutive views. Thus, for example, a single point correspondence across $m > 4$ views will give a set of equations of the form



where each row represents a set of equations derived from a quadrifocal tensor relationship. Each black square represents a block with 3 columns corresponding to one of the vectors t^i . In the diagram above, we choose to wrap the equations around from the last to the first view to add greater rigidity. Otherwise, the values of the t^i can drift from the first to the last view. Other schemes for selecting groups of views are possible, and it is not necessary to restrict to consecutive views.

Linear relations may be generated between any subset of sufficiently many images (2, 3 or 4 depending on which tensor is used to generate the equations). One must trade off the added stability of the solution against the added computational cost of adding more equations. A mixture of bifocal, trifocal and quadrifocal constraints may be used

in generating the set of all equations, and it is not necessary that all points be visible in all views.

Numbers of equations generated. Let the total number of views be m . Consider a subset of s views ($s = 2, 3$ or 4) and let n point correspondences be given between these views. We briefly consider the problem of reconstruction from this subset of s views in isolation. From these point correspondences we can generate a set of equations At' between the entries t' of the s views, and thence estimate the values of the $3s$ entries of t' . In doing this, we can assume that the first view has $t = 0$, and the vectors t from the remaining $s - 1$ views are only determined up to a common scale. Thus, the A occurring in the equation set At' has a right null-space of dimension at least 4, corresponding to the 4 degrees of freedom of the solution. In general, then:

Result 18.1. *Ignoring the effects of noisy data, the total rank of the set of equations generated from $n \geq 2$ point correspondences in s views is $3s - 4$. This is independent of the number of point (or line) correspondences used to generate them.*


To be exact, the argument above showed that the rank was *at most* $3s - 4$. For 2-view, 3-view and 4-view correspondences this is equal to 2, 5 or 8 respectively. However, as long as there are two correspondences the maximum rank is achieved. This is because two points are sufficient for reconstruction from $s = 2, 3$ or 4 views as shown by the counting arguments of section 17.5.2.

Now consider the total set of m views. The total number of retrievable parameters of all the t^i is $3m - 4$. Therefore, for a solution to be possible, the number of equations must exceed $3m - 4$, which gives the following result.


Result 18.2. *If S subsets of s_k views are chosen from among m views, then in order to solve for all the vectors t^i representing final columns of the camera matrices, it is necessary that*

$$\sum_{k=1}^S (3s_k - 4) \geq 3m - 4.$$

One can verify that if 2-view correspondences are to be used, involving equations derived from the fundamental matrix constraints, then it is not sufficient to use just pairs of consecutive views in a configuration such as



for in this case, the total number of equations generated is $m(3s - 4) = m(3 \cdot 2 - 4) = 2m$, whereas the total number of equations required is $3m - 4$. Thus for $m > 4$ there are not enough equations. This is related to the fact that the fundamental matrices between consecutive views are not sufficient to define the structure of the sequence of views. It is necessary to add additional constraints from non-consecutive views, such as



Note though that the discussion of section 15.4(p383) suggests that it is preferable to use trifocal or quadrifocal constraints over triplets or quadruplets of views. Implementation details for this method are given in [Kaucic-01].

18.6 Reconstruction from sequences

In this final section we bring together several ideas from earlier in the book. The objective here is to compute a reconstruction from a sequence of frames provided by a video. There are three stages to this problem: (i) compute corresponding features throughout the sequence; (ii) compute an initial reconstruction which may be used as a starting point for (iii) bundle adjustment (as described in section 18.1).

Here the features we will consider are interest points, though others such as lines could equally well be used. The correspondence problem is exacerbated because an interest point feature will generally not appear in all of the images, and often will be missing from consecutive images. Bundle adjustment, however, is not hindered by missing correspondences.

There are several advantages of a video sequence over an arbitrary set of images: (i) there is an ordering on the images; (ii) the distance between camera centres (the baseline) for successive frames is small. The small baseline is important because it enables possible feature matches between successive images to be obtained and assessed more easily. Matches are more easily obtained because the image points do not move “far” between views so a proximity search region can be employed; matches are more easily assessed (as to whether they arise from the same point in 3-space) because nearby images are similar in appearance. The disadvantage of a small baseline is that the 3D structure is estimated poorly. However, this disadvantage is mitigated by tracking over many views in the sequence so that the effective baseline is large.

An overview of the method is given in algorithm 18.3. There are several strategies that may be used to obtain the initial reconstruction, though this area is still to some extent a black art. Three possibilities are:

1. Extending the baseline. Suppose a reasonable number of scene points are visible throughout the sequence. Correspondences may be carried through from the first to the last frame using the pairwise matches (from F), or the triplet matched points (from T). Indeed if the baseline between consecutive frames is small (compared to the structure depth), then pairwise matches may be obtained using homography computation (algorithm 4.6(p123)) – this provides a stronger matching constraint (point to point) than F (point to line).

A trifocal tensor can then be estimated from corresponding points in the first, middle (say), and end frames of the sequence. This tensor determines a projective reconstruc-

Objective

Given a sequence of frames in a video, compute correspondences and a reconstruction of the scene structure and the camera for each frame.

Algorithm

- (i) **Interest points:** Compute interest points in each image.
- (ii) **2 view correspondences:** Compute interest point correspondences and F between consecutive frames using algorithm 11.4(p291) (frames may be omitted if the baseline motion is too small).
- (iii) **3 view correspondences:** Compute interest point correspondences and T between all consecutive image triplets using algorithm 16.4(p401).
- (iv) **Initial reconstruction:** See text.
- (v) **Bundle adjust** the cameras and 3D structure for the complete sequence.
- (vi) **Auto-calibration:** see chapter 19 (optional).

Algorithm 18.3. Overview of reconstruction from a sequence of images.

tion for those points and frames. The cameras for the intermediate frames may then be estimated by resectioning, and the scene points not visible throughout the sequence estimated by triangulation.

2. Hierarchical merging of sub-sequences. The idea here is to partition the sequence into manageable sub-sequences (there can be several hierarchical layers of partitioning). A projective reconstruction is then computed for each sub-sequence and these reconstructions are “zipped” (merged) together.

Consider the problem of merging two triplets which overlap by two views. It is a simple matter to extend the correspondences over the views: a correspondence which exists across the triplet 1-2-3 and also across the triplet 2-3-4 may be extended to the frames 1-2-3-4, since the pair 2-3 overlaps for the triplets. The camera matrices and 3D structure are then computed for the frames 1-2-3-4, for example by first resectioning and then bundle adjustment. This process is extended by merging neighbouring groups of frames until camera matrices and correspondences are established throughout the sequence. In this manner error can be distributed evenly over the sequence.

3. Incremental bundle adjustment. A fresh bundle adjustment is carried out as the correspondences from each new frame are added. The disadvantage of this method is the computational expense and also the possibility that error systematically accumulates.

Of course these three methods may be combined together. For example, the sequence can be partitioned into a sub-sequence where common points are visible, and a reconstruction built for the sub-sequence using the extended baseline method. These sub-sequences may then be combined hierarchically.

In this manner structure and cameras may be computed automatically for sequences consisting of hundreds of frames. These reconstructions may form the basis for such tasks as navigation (determining the camera/ego-position) and virtual model genera-




Fig. 18.3. **Corridor sequence.** (a) A three dimensional reconstruction of points and lines in the scene, and (b) cameras (represented by their image planes) computed automatically from the images. A texture mapped triangulated graphical model is then automatically constructed as described in [Baillard-99]. (c) A rendering of the scene from a novel viewpoint, different from any in the sequence. (d) VRML model of the scene with the cameras represented by their image planes (texture mapped with the original images from the sequence).

tion. Often it is necessary first to compute a metric reconstruction from the projective one, using the methods described in chapter 10 and chapter 19. Metric reconstruction and virtual model generation is illustrated in the following examples.

Example 18.3. Corridor sequence

A camera is mounted on a mobile vehicle for this sequence. The vehicle moves along the floor turning to the left. The forward translation in this sequence makes structure recovery difficult, due to the small baseline for triangulation. In this situation, the benefit of using all frames in the sequence is significant. Figure 18.3 shows the recovered structure. △




Fig. 18.4. **Wilshire:** 3D points and cameras for 350 frames of a helicopter shot. Cameras are shown for just the start and end frames for clarity, with the camera path plotted between.

Example 18.4. “Wilshire” sequence

This is a helicopter shot of Wilshire Boulevard, Los Angeles. In this case reconstruction is hampered by the repeated structure in the scene – many of the feature points (for example those on the skyscraper windows) have very similar intensity neighbourhoods, so correlation-based tracking produces many false candidates. However, the robust geometry-guided matching successfully rejects the incorrect correspondences. Figure 18.4 shows the structure. △

18.7 Closure

It is probably fair to say that no fully satisfactory technique for reconstruction from a sequence of projective images is known, and many ad-hoc techniques have been used, with reasonable success. Four views is the limit for the closed-form solutions based on multiview tensors. For larger numbers of views there is no such neat mathematical formulation of the problem. One exception to this is the m -view technique based on duality (see chapter 20), but this technique is limited to six to eight points, depending on which dual tensor (fundamental matrix, trifocal tensor or quadrifocal tensor) is used. Most sequences will contain many more matched points than this.

18.7.1 Literature

The Tomasi-Kanade algorithm was first proposed for orthographic projection, [Tomasi-92], but was later extended to paraperspective [Poelman-94]. It has been extended to lines and conics, e.g. [Kahl-98a], but the MLE property no longer applies, and it is unclear what is being minimized in the affine reconstruction. Others have investigated subspace methods for multiple affine views in the case of planes [Irani-99], and the case of multiple objects moving independently [Boult-91, Gear-98]. Non-rigid factorization was formulated by [Brand-01, Torresani-01], though the elements of the idea are present in [Bascle-98]. Affine reconstruction with uncertainty (covariance-weighted data) has been discussed by Irani and Anandan [Irani-00, Anandan-02]. The method of affine reconstruction by alternating triangulation and camera estimation is mentioned in [Huynh-03], under the name "PowerFactorization."

The extension of factorization to projective cameras is due to Sturm and Triggs [Sturm-96]. Methods of iteration using this approach have been proposed by [Heyden-97a, Triggs-96].

A method of computing multiple cameras based on a plane homography was employed in [Cross-99], initializing the t^i vectors using planar auto-calibration.

Methods for obtaining an initial projective reconstruction from a sequence are described in [Avidan-98, Beardsley-94, Beardsley-96, Fitzgibbon-98a, Laveau-96a, Nister-00, Sturm-97b]. [Torr-99], and more recently [Pollefeys-02], discuss the important problem of scene and motion degeneracies that may be encountered in sequence reconstruction.

18.7.2 Notes and exercises

- (i) The affine factorization algorithm can be applied to obtain a reconstruction in situations where a set of cameras $\{P^i\}$ have a common third row, even though the cameras are not affine. The third row is the principal plane of the camera (see section 6.2(p158)) and the condition of a common third row is equivalent to coplanar principal planes. For example if a camera translates in a direction perpendicular to its principal axis, then all the camera centres will lie on a plane, and the principal planes are coplanar. The affine factorization algorithm can be applied in this case because the set of cameras can be transformed as $P^i H_{4 \times 4}$ to

the affine form by a 4×4 homography H satisfying $P^3 H_{4 \times 4} = (0, 0, 0, 1)$, where P^3 is the last row of P^i .

More generally, if the camera centres are restricted to a plane then the images may be synthetically rotated such that the cameras effectively have coplanar principal planes. For example in the case of planar motion (section 19.8) or single axis rotation (section 19.9(p490)) if all the images are rotated such that the principal axis is parallel to the rotation axis (by applying a homography to each image which maps the horizon to infinity in the case of a vertical rotation axis), then the principal planes of all the cameras are parallel. However, if the cameras are not actually affine, then the algorithm will not give the ML estimate of the reconstruction.

Auto-Calibration

Auto-calibration is the process of determining internal camera parameters directly from multiple uncalibrated images. Once this is done, it is possible to compute a metric reconstruction from the images. Auto-calibration avoids the onerous task of calibrating cameras using special calibration objects. This gives great flexibility since, for example, a camera can be calibrated directly from an image sequence despite unknown motion and changes in some of the internal parameters.

The root of the method is that a camera moves rigidly, so the absolute conic is fixed under the motion. Conversely, then, if a unique fixed conic in 3-space can be determined in some way from the images, this identifies Ω_∞ . As we have seen in earlier chapters, once Ω_∞ is identified, the metric geometry can be computed. An array of auto-calibration methods are available for this task of identifying Ω_∞ .

This chapter has four main parts. First we lay out the algebraic structure of the auto-calibration problem, and show how the auto-calibration equations are generated from constraints on the internal or external parameters. Second, we describe several *direct* methods for auto-calibration which involve computing the absolute conic or its image. These include estimating the absolute dual quadric over many views, or the Kruppa equations from view pairs. Third, are *stratified* methods for auto-calibration which involve two steps – first solving for the plane at infinity, then using this to solve for the absolute conic. The fourth part covers a number of special configurations including a camera rotating about its centre, a camera undergoing planar motion, and the motion of a stereo rig.

19.1 Introduction

Auto- (or self-) calibration is the computation of metric properties of the cameras and/or the scene from a set of uncalibrated images. This differs from conventional calibration where the camera calibration matrix K is determined from the image of a known calibration grid (chapter 7) or properties of the scene, such as vanishing points of orthogonal directions (chapter 8). Instead, in auto-calibration the metric properties are determined directly from constraints on the internal and/or external parameters.

For example, suppose we have a set of images acquired by a camera with fixed internal parameters, and that a projective reconstruction is computed from point cor-

respondences across the image set. The reconstruction computes a projective camera matrix P^i for each view. Our constraint is that for the actual cameras the internal parameter matrix K is the same (but unknown) for each view. Now, each camera P^i of the projective reconstruction may be decomposed as $P^i = K^i[R^i | t^i]$ but in general the calibration matrix K^i will differ for each view. Thus the constraint will *not* be satisfied by the projective reconstruction.

However, we have the freedom to vary our projective reconstruction by transforming the camera matrices by a homography H . Since the actual cameras have fixed internal parameters, there will exist a homography (or a family of homographies) such that the transformed cameras P^iH do decompose as $P^iH = KR^i[I | t^i]$, with the same calibration matrix for each camera, so the reconstruction is consistent with the constraint. Provided there are sufficiently many views and the motion between the views is general (see later), then this consistency constrains H to the extent that the reconstruction transformed by H is within a similarity transformation of the actual cameras and scene, i.e. we achieve a metric reconstruction.

Although the particular constraints used to achieve a metric reconstruction may differ, this example illustrates the general approach:

- (i) Obtain a projective reconstruction $\{P^i, X_j\}$.
- (ii) Determine a rectifying homography H from auto-calibration constraints, and transform to a metric reconstruction $\{P^iH, H^{-1}X_j\}$.

Various flavours of auto-calibration will be covered in the following sections. They differ in the constraints used, and the methods whereby the homography H is determined. The methods may be divided into two classes: those that directly determine H ; and those that are stratified, determining first the projective and then the affine components of H . The advantage of the latter approach is that once an affine reconstruction is achieved, i.e. π_∞ is known, the solution for a metric reconstruction is linear.

If camera calibration rather than metric scene reconstruction is the goal, then it is not always necessary to compute an explicit projective reconstruction, and sometimes the camera calibration may be computed more directly than via a rectifying transformation. This is the case, for instance, if a camera rotates about its centre without translation, as is discussed in section 19.6.

19.2 Algebraic framework and problem statement

Suppose we have a projective reconstruction $\{P^i, X_j\}$; then based on constraints on the cameras' internal parameters or motion we wish to determine a rectifying homography H such that $\{P^iH, H^{-1}X_j\}$ is a metric reconstruction.

We start from the true metric situation with calibrated cameras, and structure represented in a Euclidean world frame. Thus in actuality there are m cameras P_M^i which project a 3D point X_M to an image point $x^i = P_M^i X_M$ in each view, where the subscript M indicates that the cameras are calibrated and the world frame is Euclidean. The cameras may be written as $P_M^i = K^i[R^i | t^i]$ for $i = 1, \dots, m$.

In a projective reconstruction we obtain cameras P^i which are related to P_M^i by

$$P_M^i = P^i H \quad i = 1, \dots, m \quad (19.1)$$

where H is an unknown 4×4 homography of 3-space. Our goal is to determine H .

To be precise we are not concerned with the absolute rotation, translation and scale of the reconstruction, and we will now factor out this similarity component. We choose the world frame to coincide with the first camera, so that $R^1 = I$ and $t^1 = 0$. Then R^i, t^i specifies the Euclidean transformation between the i -th camera and the first, and $P_M^1 = K^1[I \mid 0]$. Similarly, in the projective reconstruction we choose the usual canonical camera for the first view, so that $P^1 = [I \mid 0]$. Then writing H as

$$H = \begin{bmatrix} A & t \\ v^T & k \end{bmatrix}$$

the condition $P_M^1 = P^1 H$ from (19.1) becomes $[K^1 \mid 0] = [I \mid 0]H$, which implies that $A = K^1$ and $t = 0$. In addition, since H is non-singular, k must be non-zero, so we may assume $k = 1$ (this fixes the scale of the reconstruction). This shows that H is of the form

$$H = \begin{bmatrix} K^1 & 0 \\ v^T & 1 \end{bmatrix}.$$

This has factored out the similarity component.

The vector v , together with K^1 , specifies the plane at infinity in the projective reconstruction since the coordinates of π_∞ are

$$\pi_\infty = H^{-T} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \begin{bmatrix} (K^1)^{-T} & -(K^1)^{-T}v \\ 0 & 1 \end{bmatrix} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -(K^1)^{-T}v \\ 1 \end{pmatrix}.$$

We will write $\pi_\infty = (p^T, 1)^T$ so that $p = -(K^1)^{-T}v$. To summarize so far we have shown:

Result 19.1. A projective reconstruction $\{P^i, X_j\}$ in which $P^1 = [I \mid 0]$ can be transformed to a metric reconstruction $\{P^i H, H^{-1} X_j\}$ by a matrix H of the form

$$H = \begin{bmatrix} K & 0 \\ -p^T K & 1 \end{bmatrix} \quad (19.2)$$

where K is an upper triangular matrix. Furthermore,

- (i) $K = K^1$ is the calibration matrix of the first camera.
- (ii) The coordinates of the plane at infinity in the projective reconstruction are given by $\pi_\infty = (p^T, 1)^T$.

Conversely, if the plane at infinity in the projective frame and the calibration matrix of the first camera are known, then the transformation H that converts the projective to a metric reconstruction is given by (19.2).

From this result it follows that to transform a projective to a metric reconstruction it is sufficient to specify 8 parameters – the three entries of $\mathbf{p}$ and five entries of $\mathbf{K}^1$. This agrees with a geometric counting argument. Finding metric structure is equivalent to specifying the plane at infinity and the absolute conic (three and five degrees of freedom respectively). In a metric reconstruction the calibration $\mathbf{K}^i$ of each camera, its rotation $\mathbf{R}^i$ relative to the first camera, and its translation $\mathbf{t}^i$ relative to the first camera up to a single common scaling, i.e. $\mathbf{t}^i \mapsto s\mathbf{t}^i$, are all determined.

We now develop the basic auto-calibration equations. We denote the cameras of the projective reconstruction as $\mathbf{P}^i = [\mathbf{A}^i \mid \mathbf{a}^i]$. Multiplying out (19.1) using (19.2) gives

$$\mathbf{K}^i \mathbf{R}^i = (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T) \mathbf{K}^1 \quad \text{for } i = 2, \dots, m \quad (19.3)$$

which may be rearranged as $\mathbf{R}^i = (\mathbf{K}^i)^{-1} (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T) \mathbf{K}^1$ $i = 2, \dots, m$. Finally, the rotation $\mathbf{R}^i$ may be eliminated using $\mathbf{R} \mathbf{R}^T = \mathbf{I}$, leaving

$$\mathbf{K}^i \mathbf{K}^{iT} = (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T) \mathbf{K}^1 \mathbf{K}^{1T} (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T)^T.$$

Note now that $\mathbf{K}^i \mathbf{K}^{iT} = \omega^*$, the dual image of the absolute conic (or DIAC) – see (8.11–p210). Making this substitution gives the basic equations for auto-calibration:

$$\begin{aligned} \omega^{*i} &= (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T) \omega^{*1} (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T)^T \\ \omega^i &= (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T)^{-T} \omega^1 (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T)^{-1} \end{aligned} \quad (19.4)$$

the second equation being simply the inverse of the first, with ω the image of the absolute conic (or IAC). These equations relate the unknown entries of ω^{*i} or ω^i $i = 1, \dots, m$ and unknown parameters $\mathbf{p}$ with the *known* entries of the projective cameras $\mathbf{A}^i, \mathbf{a}^i$.

The art of auto-calibration is to use constraints on the $\mathbf{K}^i$, such as that one of the elements of $\mathbf{K}^i$ is zero, to generate equations on the eight parameters of $\mathbf{p}$ and $\mathbf{K}^1$ from (19.4). All auto-calibration methods are variations on solving these equations, and in the following sections we describe several of these methods. Generally methods proceed by computing ω^i or ω^{*i} first and extracting the values of the calibration matrices $\mathbf{K}^i$ from these – though iterative methods (such as bundle adjustment) may be parametrized by $\mathbf{K}^i$ directly. The equations (19.4) have a geometric interpretation as mappings on the absolute conic, and we will return to this in section 19.3 and section 19.5.2.

We start with a simple example to illustrate how equations on the eight parameters are generated from (19.4).

Example 19.2. Auto-calibration equations for identical $\mathbf{K}^i$

Suppose that all the cameras have the same internal parameters, so $\mathbf{K}^i = \mathbf{K}$, then (19.4) becomes

$$\mathbf{K} \mathbf{K}^T = (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T) \mathbf{K} \mathbf{K}^T (\mathbf{A}^i - \mathbf{a}^i \mathbf{p}^T)^T \quad i = 2, \dots, m. \quad (19.5)$$

Each view $i = 2, \dots, m$ provides an equation, and we can develop a counting argument

for the number of views required (in principle) in order to be able to determine the 8 unknowns. Each view other than the first imposes 5 constraints since each side is a 3×3 symmetric matrix (i.e. 6 independent elements) and the equation is homogeneous. Assuming these constraints are independent for each view, a solution is determined provided $5(m - 1) \geq 8$. Consequently, provided $m \geq 3$ a solution is obtained, at least in principle. Clearly, if m is much larger than 3 the unknowns K and p are very over-determined. $\triangle$

One could imagine using (19.5) as a basis for a direct estimation of the rectifying transformation H . This may be framed as a parametrized minimization problem in which the eight parameters of (19.2) are allowed to vary with the purpose of minimizing a cost function on how well the equations (19.4) are satisfied or measuring closeness to metric structure. Of course, a method of obtaining an initial solution would also be required. In essence these two steps, initial solution and iterative minimization, are what will be investigated in the following sections – though under constraints less restrictive than identical internal parameters.

19.3 Calibration using the absolute dual quadric

The absolute dual quadric, Q_{∞}^* is a degenerate dual (i.e. plane) quadric represented by a 4×4 homogeneous matrix of rank 3. Its importance here is that Q_{∞}^* encodes both π_{∞} and Ω_{∞} in a very concise fashion, for instance π_{∞} is the null-vector of Q_{∞}^* , and it has an algebraically simple image projection:

$$\omega^* = PQ_{\infty}^*P^T \quad (19.6)$$

which is simply the projection (8.5-p201) of a (dual) quadric. In words, Q_{∞}^* projects to the dual image of the absolute conic $\omega^* = KK^T$.

The idea of auto-calibration based on Q_{∞}^* is to use (19.6) to transfer a constraint on ω^* to a constraint on Q_{∞}^* via the (known) camera matrix P^i . In this manner the matrix representing Q_{∞}^* may be determined in the projective reconstruction from constraints on K^i , as will be seen below. Indeed, Q_{∞}^* was introduced as a convenient representation for auto-calibration in [Triggs-97].

Once Q_{∞}^* has been determined, then the rectifying homography (19.2) H that we seek is also determined as shown below. Thus we have a general framework for auto-calibration based on specifying constraints on K^i to determine Q_{∞}^* , and then from Q_{∞}^* determining H . This general approach is summarized in algorithm 19.1. In section 19.3.1 we will concentrate on the second step of this algorithm, estimation of Q_{∞}^* . We first fill in some details.

Simple properties of the absolute dual quadric. Section 3.7(p83) gives a full description of Q_{∞}^* . For the purposes of auto-calibration particularly important properties are summarized here. In a Euclidean frame Q_{∞}^* has the canonical form

$$\tilde{I} = \begin{bmatrix} I_{3 \times 3} & 0 \\ 0^T & 0 \end{bmatrix}, \quad (19.7)$$

Objective

Given a set of matched points across several views and constraints on the calibration matrices K^i , compute a metric reconstruction of the points and cameras.

Algorithm

- (i) Compute a projective reconstruction from a set of views, resulting in camera matrices P^i and points X_j .
- (ii) Use (19.6) along with constraints on the form of the ω^{*i} arising from K^i to estimate Q_∞^* .
- (iii) Decompose Q_∞^* as $H\tilde{I}H^T$, where $\tilde{I}$ is the matrix $\text{diag}(1, 1, 1, 0)$.
- (iv) Apply H^{-1} to the points and H to the cameras to get a metric reconstruction.
- (v) Use iterative least-squares minimization to improve the solution (see section 19.3.3).

Alternatively, the calibration matrix of each camera may be computed directly:

- (i) Compute ω^{*i} for all i using (19.6).
- (ii) Compute the calibration matrix K^i from the equation $\omega^* = KK^T$ by Cholesky factorization.

Algorithm 19.1. Auto-calibration based on Q_∞^* .

In a projective coordinate frame Q_∞^* has the form $Q_\infty^* = H\tilde{I}H^T$, where $\tilde{I}$ is the matrix in (19.7). This follows from the projective transformation rule (3.17–p74) for dual quadrics, $Q_\infty^* \mapsto HQ_\infty^*H^T$. Consequently:

Result 19.3. *In an arbitrary projective frame, the dual absolute quadric is represented by a symmetric 4×4 matrix with the following properties.*

- (i) *It is singular of rank 3, since Q_∞^* is a degenerate conic.*
- (ii) *Its null space is the vector representing the plane at infinity, since $Q_\infty^* \pi_\infty = 0$.*
- (iii) *It is positive semi-definite (or negative – depending on the homogeneous scale).*

These properties are immediate for Q_∞^* in its canonical form in a Euclidean frame, and easily extend to an arbitrary frame.

Extracting the rectifying homography from Q_∞^* . Given an estimated Q_∞^* in a projective coordinate frame we wish to determine the homography H . Extracting H is a simple matter of decomposing the expression as follows.

Result 19.4. *If Q_∞^* is decomposed as $Q_\infty^* = H\tilde{I}H^T$ (see notation above), then H^{-1} is a 3D (point) homography that takes the projective coordinate frame to a Euclidean frame.*

Note that a camera is transformed by the inverse of the transformation applied to points, so H is the correct matrix to apply to a camera to give $P_M = PH$. Thus H is the rectifying transformation to apply to cameras. A decomposition of Q_∞^* as $H\tilde{I}H^T$ may be easily computed from its eigenvalue decomposition (see section A4.2(p580) for Jacobi's algorithm for this).

Equivalence to auto-calibration equations. Equations (19.6) which describe the image projection of Q_∞^* are simply a geometric representation of the auto-calibration equations (19.4) as will now be demonstrated.

The forms of $\omega = (KK^T)^{-1}$ and $\omega^* = \omega^{-1} = KK^T$ for a camera with calibration matrix K as in (6.10–p157) are

$$\omega^* = \begin{bmatrix} \alpha_x^2 + s^2 + x_0^2 & s\alpha_y + x_0y_0 & x_0 \\ s\alpha_y + x_0y_0 & \alpha_y^2 + y_0^2 & y_0 \\ x_0 & y_0 & 1 \end{bmatrix} \quad (19.9)$$

and

$$\omega = \frac{1}{\alpha_x^2 \alpha_y^2} \begin{bmatrix} \alpha_y^2 & -s\alpha_y & -x_0\alpha_y^2 + y_0s\alpha_y \\ -s\alpha_y & \alpha_x^2 + s^2 & \alpha_y s x_0 - \alpha_x^2 y_0 - s^2 y_0 \\ -x_0\alpha_y^2 + y_0s\alpha_y & \alpha_y s x_0 - \alpha_x^2 y_0 - s^2 y_0 & \alpha_x^2 \alpha_y^2 + \alpha_x^2 y_0^2 + (\alpha_y x_0 - s y_0)^2 \end{bmatrix} \quad (19.10)$$

If the skew is zero, i.e. $s = 0$, then the expressions simplify to

$$\omega^* = \begin{bmatrix} \alpha_x^2 + x_0^2 & x_0y_0 & x_0 \\ x_0y_0 & \alpha_y^2 + y_0^2 & y_0 \\ x_0 & y_0 & 1 \end{bmatrix} \quad (19.11)$$

and

$$\omega = \frac{1}{\alpha_x^2 \alpha_y^2} \begin{bmatrix} \alpha_y^2 & 0 & -\alpha_y^2 x_0 \\ 0 & \alpha_x^2 & -\alpha_x^2 y_0 \\ -\alpha_y^2 x_0 & -\alpha_x^2 y_0 & \alpha_x^2 \alpha_y^2 + \alpha_y^2 x_0^2 + \alpha_x^2 y_0^2 \end{bmatrix} \quad (19.12)$$

Table 19.1. The image of the absolute conic, ω , and dual image of the absolute conic, ω^* , written in terms of the camera internal parameters.

We have seen that in a projective frame Q_∞^* has the form $H\tilde{I}H^T$. The projective reconstruction is related to the metric reconstruction by (19.2), so in detail

$$Q_\infty^* = H\tilde{I}H^T = \begin{bmatrix} K^1 K^{1T} & -K^1 K^{1T} \mathbf{p} \\ -\mathbf{p}^T K^1 K^{1T} & \mathbf{p}^T K^1 K^{1T} \mathbf{p} \end{bmatrix} = \begin{bmatrix} \omega^{*1} & -\omega^{*1} \mathbf{p} \\ -\mathbf{p}^T \omega^{*1} & \mathbf{p}^T \omega^{*1} \mathbf{p} \end{bmatrix}. \quad (19.8)$$

On applying (19.6) with $P^i = [A^i \mid \mathbf{a}^i]$ we obtain once again the auto-calibration equations (19.4)

$$\omega^{*i} = P^i Q_\infty^* P^{iT} = (A^i - \mathbf{a}^i \mathbf{p}^T) \omega^{*1} (A^i - \mathbf{a}^i \mathbf{p}^T)^T.$$

This is a geometric interpretation of (19.4) – Q_∞^* is a fixed quadric under the Euclidean motion of the camera, and the DIAC ω^{*i} is the image of Q_∞^* in each view.

19.3.1 Linear solutions for Q_∞^* from a set of images

The objective here is to estimate Q_∞^* in a projective reconstruction directly from constraints on the internal parameters. We will start by describing three cases for which a linear solution may be obtained. It is convenient at this point to summarize the forms of the DIAC and also the IAC. Refer to table 19.1.

Specifying linear constraints on Q_∞^* . Linear constraints on Q_∞^* may be obtained if the principal point is known. Assume that this point is known, then we may change

the image coordinate system so that the origin coincides with the principal point. Then $x_0 = 0$, $y_0 = 0$, and from table 19.1 the DIAC becomes

$$\omega^* = \begin{bmatrix} \alpha_x^2 + s^2 & s\alpha_y & 0 \\ s\alpha_y & \alpha_y^2 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (19.13)$$

The linear equations on Q_∞^* are then generated from the zero entries in (19.13) by applying the projection equation (19.6) $\omega^* = PQ_\infty^*P^T$ to each view i . For example the two equations

$$(P^i Q_\infty^* P^{iT})_{13} = 0 \text{ and } (P^i Q_\infty^* P^{iT})_{23} = 0 \quad (19.14)$$

follow immediately from $\omega_{13}^{*i} = \omega_{23}^{*i} = 0$.

If there are additional constraints on K^i which result in further relationship between the entries of ω^* , then these may provide additional linear equations. For instance, an assumption that skew is zero means that the (1,2)-entries of (19.13) vanish, which provides one more linear equation on the entries of Q_∞^* similar to (19.14). Known aspect ratio provides a further constraint. Table 19.2 summarizes the possible constraints that may be used.

Condition	constraint	type	# constraints
zero skew	$\omega_{12}^* \omega_{33}^* = \omega_{13}^* \omega_{23}^*$	quadratic	m
principal point (p.p.) at origin	$\omega_{13}^* = \omega_{23}^* = 0$	linear	$2m$
zero skew (p.p. at origin)	$\omega_{12}^* = 0$	linear	m
fixed (unknown) aspect ratio (zero skew and p.p. at origin)	$\frac{\omega_{11}^{*i}}{\omega_{22}^{*i}} = \frac{\omega_{11}^{*j}}{\omega_{22}^{*j}}$	quadratic	$m - 1$
known aspect ratio $r = \alpha_y / \alpha_x$ (zero skew and p.p. at origin)	$r^2 \omega_{11}^* = \omega_{22}^*$	linear	m

Table 19.2. **Auto-calibration constraints derived from the DIAC.** The number of constraints column gives the total number of constraints over m views, assuming the constraint is true for each view. Each additional item of information generates additional equations. For example, if the principal point is known and skew is zero then there are 3 constraints per view.

Linear solution. Since it is symmetric, Q_∞^* may be parametrized linearly by 10 homogeneous parameters, namely the 10 diagonal and above-diagonal entries. These 10 entries may be represented by a 10-vector $\mathbf{x}$. In the usual manner the linear equations on Q_∞^* may be assembled into a matrix equation of the form $A\mathbf{x} = 0$, and a least-squares solution for $\mathbf{x}$ obtained via the SVD. For example, the two equations (19.14) provide two rows of the matrix from each view. From five images a total of 10 equations are obtained (assuming only that principal point is known), and a linear solution is possible. From four images eight equations are generated. In the same way as with the

computation of the fundamental matrix from seven points there is a 2-parameter family of solutions. The condition $\det Q_{\infty}^* = 0$ gives a fourth-degree equation and so up to four solutions for Q_{∞}^* .

Example 19.5. Linear solution for variable focal length

Suppose the camera is calibrated apart from the focal length – the principal point is known, the aspect ratio is unity (if it is not the equations can be transformed so that it is unity from the known value), and skew is zero – the focal length is unknown and may vary between views. In this case from table 19.2 there are four linear constraints on Q_{∞}^* available from each view. In the case of two views there are eight linear constraints and up to four solutions are obtained using the condition $\det Q_{\infty}^* = 0$. If $m \geq 3$ a unique linear solution exists. $\triangle$

More will be said about determination of focal lengths in this minimal case in example 19.8.

19.3.2 Non-linear solutions for Q_{∞}^*

We now describe various non-linear equations that can be obtained from the form of (19.6). It has been seen that each element of $\omega^{*i} = P^i Q_{\infty}^* P^{iT}$ is expressible as a linear expression in terms of the parameters of Q_{∞}^* . It follows that any relationship between the entries of the various ω^{*i} translates into an equation involving the entries of Q_{∞}^* . In particular, linear or quadratic relationships between entries of ω^* generate respectively linear or quadratic relationships between entries of Q_{∞}^* . Given sufficiently many such equations, we may solve for Q_{∞}^* .

Constant internal parameters. If the internal parameters of all cameras are the same, then $\omega^{*i} = \omega^{*j}$ for all i and j , which expands to $P^i Q_{\infty}^* P^{iT} = P^j Q_{\infty}^* P^{jT}$. However, since these are homogeneous quantities, the equality holds only up to an unknown scale. A set of five equations are generated:

$$\omega_{11}^{*i}/\omega_{11}^{*j} = \omega_{12}^{*i}/\omega_{12}^{*j} = \omega_{13}^{*i}/\omega_{13}^{*j} = \omega_{22}^{*i}/\omega_{22}^{*j} = \omega_{23}^{*i}/\omega_{23}^{*j} = \omega_{33}^{*i}/\omega_{33}^{*j}$$

This gives a set of quadratic equations in the entries of Q_{∞}^* . Given three views, a total of 10 equations result, which may be solved to find Q_{∞}^* .

Calibration assuming zero skew. Under the assumption of zero skew in each of the cameras, the form of the DIAC is simplified, as given in (19.11). In particular, we obtain the following constraints between the entries of ω^* in the zero-skew case

$$\omega_{12}^* \omega_{33}^* = \omega_{13}^* \omega_{23}^* \quad (19.15)$$

This gives a single quadratic equation in the entries of Q_{∞}^* . From a set of m views we obtain m quadratics. However there is also one extra equation $\det Q_{\infty}^* = 0$ derived from the fact that the absolute dual quadric is degenerate. Since Q_{∞}^* has 10 homogeneous linear parameters, it may be computed (at least in principle) from 8 views.

These different calibration constraints are also summarized in table 19.2.

19.3.3 Iterative methods

As we have seen on many occasions throughout this book there is a choice between minimizing an algebraic or geometric error. In this case a suitable algebraic error is provided by (19.4). In previous cases, such as (4.1–p89), the unknown scale factor has been eliminated by forming a cross product. Here the scale factor can be eliminated by using a matrix norm. The cost function is

$$\sum_i \|K^i K^{iT} - P^i Q_\infty^* P^{iT}\|_F^2 \quad (19.16)$$

where $\|M\|_F$ is the Frobenius norm of M , and $K^i K^{iT}$ and $P^i Q_\infty^* P^{iT}$ are both normalized to have unit Frobenius norm. The cost function is parametrized by the (at most eight) unknown elements of Q_∞^* , and the unknown elements of each $\omega^{*i} = K^i K^{iT}$. It is possible to use the expansion (19.8) to parametrize the absolute dual quadric. For example, in the case of example 19.5 where the focal length is the only unknown per view, (19.16) would be minimized over $m+3$ parameters. These are the focal length f^i of each view, and the three components of $\mathbf{p}$. Note that this parametrization ensures that Q_∞^* has rank 3 throughout the minimization.

Since the above cost function has no particular geometric meaning, it is advisable to follow this up with a complete bundle adjustment. In fact, given a good initial linear estimate one can proceed directly to bundle adjustment. There is no difficulty in incorporating assumptions on calibration parameters into a full bundle adjustment as described in section 18.1(p434).

Example 19.6. Metric reconstruction for general motion

Figure 19.1(a–c) shows views of an Indian temple acquired by a hand held camera. A projective reconstruction is computed from image point correspondences as described in section 18.6, and a metric reconstruction obtained using algorithm 19.1 under the constraint of constant camera parameters with known principal point. The computed cameras and 3D point cloud are shown in figure 19.1(d) and (e). $\triangle$

19.3.4 A counting argument

The constraints we have seen have been of two types: a parameter has a known value; or a parameter is fixed across views but its value is unknown. The actual constraints that apply depend on the physical circumstances of the image production from acquisition by the camera, through digitization and cropping, to the final image. For example, for an image sequence in which the lens is zoomed it might be the case that the skew and aspect ratio are fixed (but unknown), but that the focal length and principal point vary through the sequence. Often it is the case that the pixels are square or have a known aspect ratio, so that both the skew (which is zero) and aspect ratio are known.

We will now consider the number of constraints that are required to determine a metric reconstruction fully.

The number of parameters that must be computed to perform calibration is 8. This is equal to the number of essential parameters of the absolute dual quadric, including the scale ambiguity and rank-3 constraint. Consider m views and suppose that k of the




Fig. 19.1. **Metric reconstruction for general motion.** (a)-(c) 3 views (of 5) acquired by a hand held camera. (d) and (e) Two views of a metric reconstruction computed from interest points matches over the five views. The cameras are represented by pyramids with apex at the computed camera centre. (f) and (g) two views of a texture mapped 3D model computed from the original images and reconstructed cameras using an area based stereo algorithm. Figures courtesy of Marc Pollefeys, Reinhard Koch, and Luc Van Gool.

internal parameters are known in all views, and f are fixed over the views but unknown (where $k + f \leq 5$). A fixed and known calibration parameter provides one constraint per view via the condition $\omega^{*i} = P^i Q^* P^{iT}$, for a total of mk constraints. A fixed but unknown calibration parameter provides one fewer constraint, since just the value of the unknown parameter is missing. Thus f fixed parameters provide a total of $f(m-1)$ constraints. The requirement for calibration is then that

$$mk + (m-1)f \geq 8.$$

Table 19.3 gives values for m for several combinations of constraints. It is important to remember that degenerate configurations are of course possible in which some of the constraints are dependent. This will increase the number of required views.

19.3.5 Limitations of the absolute quadric approach to calibration

The following considerations apply to calibration using this method.

Limitations of least-squares algebraic solution. Since the least-squares solution (e.g. of $Ax = 0$ in the linear solution for the ω^*) minimizes but does not enforce constraints, the solution obtained will not precisely satisfy the required conditions. This observation holds in over-constrained cases. For instance in the case of estimating focal lengths in example 19.5, the entries ω_{11}^{*i} and ω_{22}^{*i} will not be in the required

Condition	fixed f	known k	views m
Constant internal parameters	5	0	3
Aspect ratio and skew known, focal length and principal point vary	0	2	4*
Aspect ratio and skew constant, focal length and principal point vary	2	0	5*
Skew zero, all other parameters vary	0	1	8*
p.p. known all other parameters vary	0	2	4*, 5(linear)
p.p. known skew zero	0	3	3(linear)
p.p., skew and aspect ratio known	0	4	2, 3(linear)

Table 19.3. *The number of views m required under various conditions in order for there to be enough constraints for auto-calibration. For those cases marked with an asterisk there may be multiple solutions, even for general motion between views.*

ratio, nor will the off-diagonal entries be precisely zero. This means that the K^i will not be precisely of the desired form. The absolute dual quadric computed by linear means will not have rank 3 in general, since this is not enforced by the linear equations. A rank 3 matrix for Q_∞^* can be obtained by setting the smallest eigenvalue to zero in its eigenvalue decomposition (in a similar manner to using the SVD to obtain a rank 2 matrix for F in the 8-point algorithm in section 11.1.1(p280)). This rank 3 matrix may then be directly decomposed to obtain the rectifying homography (19.2) using result 19.4. Alternatively, the rank 3 matrix can provide the starting point for an iterative minimization as described in section 19.3.3.

The positive-definiteness condition. The most troublesome failing of this method is the difficulty in enforcing the condition that Q_∞^* is positive semi-definite, (or negative semi-definite if the sign is reversed). This is related to the condition that $\omega^* = PQ_\infty^*P^T$ should be positive-definite. If ω^* is not positive-definite, then it can not be decomposed using Cholesky factorization to compute the calibration matrix. This is a recurring problem with auto-calibration methods based on estimation of the IAC or DIAC. If the data is noisy, then this problem may occur indicating that the data are not consistent with metric reconstruction. It is not appropriate if this occurs to seek the closest positive-definite solution, since this will generally be a boundary case leading to a spurious calibration.

19.4 The Kruppa equations

A different method of auto-calibration involves the use of the Kruppa equations, which were originally introduced into computer vision by Faugeras, Luong and Maybank [Faugeras-92a] and historically are seen as the first auto-calibration method. They




Fig. 19.2. **Epipolar tangency for a conic.** A conic C_w on a world plane π is imaged as corresponding conics $C \leftrightarrow C'$ in two views. The imaged conics are consistent with the epipolar geometry of the view pair. Upper: An epipolar plane tangent to the world conic C_w defines corresponding epipolar lines which are tangent to the imaged conics. Lower: epipolar lines, l_1, l_2 tangent to the imaged conic in the first view correspond to the epipolar lines l'_1, l'_2 , respectively, tangent to the imaged conic in the second.

are two-view constraints that require only F to be known, and consist of two independent quadratic equations in the elements of ω^* .

The Kruppa equations are an algebraic representation of the correspondence of epipolar lines tangent to a conic. The geometry of this correspondence is illustrated in figure 19.2. Suppose the conics C and C' are the images of a conic C_w on a world plane in the first and second views respectively, and that C^* and C'^* are their duals. In the first view the two epipolar tangent lines l_1 and l_2 may be combined into a single degenerate point conic (see example 2.8(p32)) as $C_t = [e]_{\times} C^* [e]_{\times}$. (It may be verified that any point x on the lines l_1 and l_2 satisfies $x^T C_t x = 0$). Similarly, in the second view the corresponding epipolar lines l'_1 and l'_2 may be written as $C'_t = [e']_{\times} C'^* [e']_{\times}$. The epipolar tangent lines correspond under the homography H induced by any world plane π . Since C_t is a point conic it transforms according to result 2.13(p37) $C'_t = H^{-T} C_t H^{-1}$, and the correspondence of the lines requires that

$$\begin{aligned} [e']_{\times} C'^* [e']_{\times} &= H^{-T} [e]_{\times} C^* [e]_{\times} H^{-1} \\ &= F C^* F^T \end{aligned} \quad (19.17)$$

the last equality following from $F = H^{-T} [e]_{\times}$ (see page 335). Note, this equation does not enforce the condition that the epipolar tangent lines map individually to their corresponding lines, only that their symmetric product maps to their symmetric product.

The development to this point applies to any conic. However, in the case of interest

here the world conic is the absolute conic on the plane at infinity, so that $C^* = \omega^*$, $C^{*'} = \omega^{*'}$, (and $H = H_\infty$), and (19.17) specializes to

$$[e']_\times \omega^{*'} [e']_\times = F \omega^* F^T \quad (19.18)$$

If the internal parameters are constant over the views then $\omega^{*'} = \omega^*$ and so $[e']_\times \omega^* [e']_\times = F \omega^* F^T$, which are the Kruppa equations in a form originally given by Viéville [Vieville-95]. On eliminating the homogeneous scale factor, one obtains equations quadratic in the elements of ω^* .

Although (19.18) concisely expresses the Kruppa equations, it is not in a form that can be easily applied. A succinct and easily usable form of the Kruppa equations is now given. We show that the null-space of $[e']_\times$, which is common to both sides of (19.18), can be eliminated leaving an equation between two 3-vectors.

Result 19.7. *The Kruppa equations (19.18) are equivalent to*

$$\begin{pmatrix} \mathbf{u}_0^T \omega^{*'} \mathbf{u}_2 \\ -\mathbf{u}_1^T \omega^{*'} \mathbf{u}_2 \\ \mathbf{u}_1^T \omega^{*'} \mathbf{u}_1 \end{pmatrix} \times \begin{pmatrix} \sigma_1^2 \mathbf{v}_1^T \omega^* \mathbf{v}_1 \\ \sigma_1 \sigma_2 \mathbf{v}_1^T \omega^* \mathbf{v}_2 \\ \sigma_2^2 \mathbf{v}_2^T \omega^* \mathbf{v}_2 \end{pmatrix} = \mathbf{0} \quad (19.19)$$

where $\mathbf{u}_i$, $\mathbf{v}_i$ and σ_i are the columns and singular values of the SVD of F . This provides three quadratic equations in the elements ω_{ij}^* of ω^* , of which two are independent.

Proof. The fundamental matrix has rank 2, and thus has an SVD expansion

$$F = UDV^T = U \begin{bmatrix} \sigma_1 & & \\ & \sigma_2 & \\ & & 0 \end{bmatrix} V^T$$

where the null-vectors are $F^T \mathbf{u}_3 = \mathbf{0}$ and $F \mathbf{v}_3 = \mathbf{0}$. This means that the epipoles are $\mathbf{e} = \mathbf{v}_3$ and $\mathbf{e}' = \mathbf{u}_3$. Substituting this expansion into (19.18) we obtain

$$[\mathbf{u}_3]_\times \omega^{*'} [\mathbf{u}_3]_\times = UDV^T \omega^* VDU^T. \quad (19.20)$$

We now use the property that U is an orthogonal matrix. On pre-multiplying (19.20) by U^T , and post-multiplying by U , the LHS becomes

$$\begin{aligned} U^T [\mathbf{u}_3]_\times \omega^{*'} [\mathbf{u}_3]_\times U &= \begin{bmatrix} \mathbf{u}_2 & -\mathbf{u}_1 & \mathbf{0} \end{bmatrix}^T \omega^{*'} \begin{bmatrix} \mathbf{u}_2 & -\mathbf{u}_1 & \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{u}_2^T \omega^{*'} \mathbf{u}_2 & -\mathbf{u}_2^T \omega^{*'} \mathbf{u}_1 & 0 \\ -\mathbf{u}_1^T \omega^{*'} \mathbf{u}_2 & \mathbf{u}_1^T \omega^{*'} \mathbf{u}_1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \end{aligned}$$

and the RHS of (19.20) becomes

$$\begin{aligned} DV^T \omega^* VD &= \begin{bmatrix} \sigma_1 & & \\ & \sigma_2 & \\ & & 0 \end{bmatrix} V^T \omega^* V \begin{bmatrix} \sigma_1 & & \\ & \sigma_2 & \\ & & 0 \end{bmatrix} \\ &= \begin{bmatrix} \sigma_1^2 \mathbf{v}_1^T \omega^* \mathbf{v}_1 & \sigma_1 \sigma_2 \mathbf{v}_1^T \omega^* \mathbf{v}_2 & 0 \\ \sigma_1 \sigma_2 \mathbf{v}_1^T \omega^* \mathbf{v}_2 & \sigma_2^2 \mathbf{v}_2^T \omega^* \mathbf{v}_2 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \end{aligned}$$

It is evident that both sides have reduced to symmetric matrices each with three independent elements. These three elements may be represented by a homogeneous 3-vector on each side:

$$\begin{aligned} \mathbf{x}_{\text{LHS}}^T &= \left(\mathbf{u}_2^T \boldsymbol{\omega}^{*'} \mathbf{u}_2, -\mathbf{u}_1^T \boldsymbol{\omega}^{*'} \mathbf{u}_2, \mathbf{u}_1^T \boldsymbol{\omega}^{*'} \mathbf{u}_1 \right) \\ \mathbf{x}_{\text{RHS}}^T &= \left(\sigma_1^2 \mathbf{v}_1^T \boldsymbol{\omega}^* \mathbf{v}_1, \sigma_1 \sigma_2 \mathbf{v}_1^T \boldsymbol{\omega}^* \mathbf{v}_2, \sigma_2^2 \mathbf{v}_2^T \boldsymbol{\omega}^* \mathbf{v}_2 \right). \end{aligned}$$

The two sides are only equal up to a scale factor. However, equalities are obtained in the usual way using a vector cross-product, $\mathbf{x}_{\text{LHS}} \times \mathbf{x}_{\text{RHS}} = \mathbf{0}$. An alternative derivation is given in [Hartley-97d]. $\square$

Note, the Kruppa equations involve the DIAC, rather than the IAC, since they arise from tangent line constraints, and line constraints are more simply expressed using a dual (line) conic.

We now discuss the solution of the Kruppa equations, beginning with a simple example where all the internal parameters are known apart from the focal length. An alternative method for solving this problem using the absolute dual quadric was given in example 19.5.

Example 19.8. Focal lengths for a view pair

Suppose two cameras have zero skew and known principal point and aspect ratio, but unknown and different focal lengths (as in example 19.5). Then from (19.11) by a suitable change of coordinates their DIACs may be written as

$$\boldsymbol{\omega}^* = \text{diag}(\alpha^2, \alpha^2, 1), \quad \boldsymbol{\omega}^{*'} = \text{diag}(\alpha'^2, \alpha'^2, 1)$$

where α, α' are the unknown focal lengths of the first and second view, respectively. Writing the Kruppa equations (19.19) as

$$\frac{\mathbf{u}_2^T \boldsymbol{\omega}^{*'} \mathbf{u}_2}{\sigma_1^2 \mathbf{v}_1^T \boldsymbol{\omega}^* \mathbf{v}_1} = -\frac{\mathbf{u}_1^T \boldsymbol{\omega}^{*'} \mathbf{u}_2}{\sigma_1 \sigma_2 \mathbf{v}_1^T \boldsymbol{\omega}^* \mathbf{v}_2} = \frac{\mathbf{u}_1^T \boldsymbol{\omega}^{*'} \mathbf{u}_1}{\sigma_2^2 \mathbf{v}_2^T \boldsymbol{\omega}^* \mathbf{v}_2}$$

it is evident that each numerator is linear in terms of α'^2 and each denominator linear in α^2 . Cross-multiplying provides two simple quadratic equations in α^2 and α'^2 which are easily solved. The values of α and α' are found by taking the square roots. Note, if the internal parameters are the same for the two views (that is $\alpha = \alpha'$) then each equation of result 19.7 provides a quadratic in the single unknown α^2 . $\triangle$

Extending the Kruppa equations to multiple views. In the absence of knowledge of the internal parameters, other than that the parameters are constant across views, the Kruppa equations provide two independent constraints on the five unknown parameters. Thus given three views, with F known between each pair, there are in principle six quadratic constraints, which is sufficient to determine $\boldsymbol{\omega}^*$. Using any five of these equations results in five quadratics in five unknowns, a total of 2^5 possible solutions. Solving this set of equations is not a particularly promising approach, although solutions have been obtained by homotopy continuation [Luong-92] and by minimizing algebraic residuals using every view pair for a sequence of images [Zeller-96].

Ambiguities. If there is no rotation between views then the Kruppa equations provide no constraint on ω^* . This may be seen from (19.18) which, in the case of pure translation, reduces to $[e']_{\times} \omega^* [e']_{\times} = [e']_{\times} \omega^* [e']_{\times}$, since $F = [e']_{\times}$.

The Kruppa equations are closely related to the calibration constraint provided by the transfer of the IAC under the infinite homography, as discussed later in section 19.5.2. It will follow from that discussion that the constraint placed on ω^* by the Kruppa equations for a pair of views is weaker than that placed by the infinite homography constraint (19.25). Consequently ambiguities of ω^* imposed by the Kruppa equations are a superset of those imposed by (19.25).

The application of the Kruppa equations to three or more views provides weaker constraints than those obtained by other methods such as the modulus constraint (section 19.5.1) or the absolute dual quadric (section 19.3.1). This is because the Kruppa constraints are a view-pair constraint for *conics* obtained as a projection of a 3D (dual) *quadric*. They do not enforce that the (dual) quadric is degenerate, or equivalently do not enforce a common support plane for Ω_{∞} over the multiple views. Consequently, there are additional ambiguous solutions as described in [Sturm-97b].

Although the application of the Kruppa equations to auto-calibration is chronologically the first example in the literature, the difficulty of their solution and the problem with ambiguities has seen them losing favour in the face of more tractable methods such as the dual quadric formulation. However, if only two views are given then the Kruppa equations are *the* constraint available on ω^* .

19.5 A stratified solution

Determining a metric reconstruction involves simultaneously obtaining both the camera calibration K and the plane at infinity, π_{∞} . An alternative approach is first to obtain by some means π_{∞} , or equivalently an affine reconstruction. The subsequent determination of K is then relatively simple because there exists a linear solution. This approach will now be described starting with methods of determining π_{∞} , and hence H_{∞} , and followed by methods of computing K given H_{∞} .

19.5.1 Affine reconstruction – determining π_{∞}

For general motion and constant internal parameters, (19.3–p461) can be rearranged into providing a constraint only on π_{∞} , known as the *modulus constraint*. This allows the coordinates p of π_{∞} to be solved for directly, and is described below.

The modulus constraint

The modulus constraint is a polynomial equation in the coordinates of π_{∞} . Assume the internal parameters are constant; then from (19.3–p461) with $K^i = K$

$$A - ap^T = \mu K R K^{-1} \quad (19.21)$$

where the scale factor μ is explicitly included, and for clarity the superscripts are omitted. Since $K R K^{-1}$ is conjugate to a rotation, it has eigenvalues $\{1, e^{i\theta}, e^{-i\theta}\}$. Consequently, the eigenvalues of $(A - ap^T)$ are $\{\mu, \mu e^{i\theta}, \mu e^{-i\theta}\}$, and thus have equal moduli. This is the modulus constraint on the coordinates of the plane at infinity, p .

To develop this constraint further consider the characteristic polynomial of $A - \mathbf{ap}^T$ which is

$$\begin{aligned}\det(\lambda I - A + \mathbf{ap}^T) &= (\lambda - \lambda_1)(\lambda - \lambda_2)(\lambda - \lambda_3) \\ &= \lambda^3 - f_1\lambda^2 + f_2\lambda - f_3\end{aligned}$$

where λ_i are the three eigenvalues, and

$$\begin{aligned}f_1 &= \lambda_1 + \lambda_2 + \lambda_3 = \mu(1 + 2\cos\theta) \\ f_2 &= \lambda_1\lambda_2 + \lambda_1\lambda_3 + \lambda_2\lambda_3 = \mu^2(1 + 2\cos\theta) \\ f_3 &= \lambda_1\lambda_2\lambda_3 = \mu^3.\end{aligned}$$

Eliminating the scalar μ and angle θ we obtain

$$f_3 f_1^3 = f_2^3.$$

Looking more closely at the characteristic polynomial we observe that $\mathbf{p}$ appears only as part of the rank-1 term $\mathbf{ap}^T$. This means that the elements of $\mathbf{p}$ appear only linearly in the determinant $\det(\lambda I - A + \mathbf{ap}^T)$, and hence linearly in each of f_1, f_2, f_3 . Hence the modulus constraint may be written as a *quartic* polynomial in the three elements p_i of $\mathbf{p}$. This polynomial equation is only a necessary condition for the eigenvalues to have equal moduli, not sufficient.

Each view pair generates a quartic equation in the coordinates of π_∞ . Thus, in principle, three views determine π_∞ , but only as the intersection of three quartics in three variables – a possible 4^3 solutions. However, for three views an additional cubic equation is available from the modulus constraint, and this equation can be used to eliminate many of the spurious solutions. This cubic equation is developed in [Schaffalitzky-00a]. The modulus constraint may also be combined with scene information. For example, if a corresponding vanishing line is available in two views, then π_∞ is determined up to a one-parameter ambiguity. Applying the modulus constraint resolves this ambiguity, and results in a quartic equation in one variable.

The modulus constraint may be considered the cousin of the Kruppa equations: the Kruppa equations are equations on ω^* which do not involve π_∞ ; conversely, the modulus constraint is an equation on π_∞ which does not involve ω^* . Once one of ω^* or π_∞ is known the other can subsequently be determined.

Other methods of finding π_∞

Because of the problem with solving sets of simultaneous quartic equations, the modulus constraint is not very satisfactory as a practical means of finding the plane at infinity. In fact, finding the plane at infinity is the hardest part of auto-calibration and the place where one is most likely to run into difficulties.

The plane at infinity may be identified by various other methods. Several of these are described in chapter 10. One straightforward method (which is outside the province of pure auto-calibration) is to use properties of the scene geometry. For example, the correspondence of a vanishing point between two views determines a point on π_∞ , and three such correspondences determine π_∞ in a projective reconstruction. Indeed,

a good approximation for π_∞ may well be obtained from the correspondence of three points that are distant in the scene. A second method is to employ a pure translation between two views, i.e. the camera translates but does not rotate or change internal parameters; π_∞ is determined uniquely by such a motion.

As seen in result 19.3(p463) the plane at infinity may also be computed from the absolute dual quadric and this method is quite attractive if the principal point is known. Methods of bounding the position of π_∞ using cheirality inequalities will be described in chapter 21. These use information about which points are in front of the cameras to get a *quasi-affine* reconstruction, which is close to an affine reconstruction. A method for finding the plane at infinity by iterative search from this initial quasi-affine reconstruction is described in [Hartley-94b]. More recently, the bounds imposed by cheirality have been used [Hartley-99] to define a rectangular region of 3D parameter space inside which the vector $\mathbf{p}$ representing the plane at infinity must lie. Then an exhaustive search is undertaken to find the elusive plane at infinity inside this region.

19.5.2 Affine to metric conversion – determining $\mathbf{K}$ given π_∞

Once the plane at infinity has been determined, an affine reconstruction is effectively known. The remaining step is to transform from affine to metric. It turns out that this is a far easier step than the step from projective to affine. In fact, a linear algorithm is available based on the transformation of the IAC or its dual.

The infinite homography. The infinite homography H_∞ is the plane projective transformation between two images induced by the plane at infinity π_∞ (see chapter 13). If the plane at infinity $\pi_\infty = (\mathbf{p}^T, 1)^T$ and camera matrices $[A^i | \mathbf{a}^i]$ are known in any projective coordinate frame, an explicit formula for the infinite homography can be derived from result 13.1(p326)

$$H_\infty^i = A^i - \mathbf{a}^i \mathbf{p}^T \quad (19.22)$$

where H_∞^i represents the homography from a camera $[I | 0]$ to the camera $[A^i | \mathbf{a}^i]$. So H_∞^i may be computed from a projective reconstruction once the plane at infinity is known.

If the first camera is not in the canonical form $[I | 0]$, then one can still compute the homography H_∞^i from the first image to the i -th by writing

$$H_\infty^i = (A^i - \mathbf{a}^i \mathbf{p}^T) (A^1 - \mathbf{a}^1 \mathbf{p}^T)^{-1}. \quad (19.23)$$

This is not strictly necessary, however, since one can *invent* a new view that *does* have camera matrix in the canonical form $[I | 0]$, and express the infinite homographies with respect to this view. In the following discussion, we write $\mathbf{K}$ and ω (without superscripts) to refer either to this reference view, or to the first view, if it is in canonical form.

The absolute conic lies on π_∞ so its image is mapped between views by H_∞ . Under the point transformation $\mathbf{x}^i = H_\infty^i \mathbf{x}$, where H_∞^i is the infinite homography between the reference view and view i , the transformation rules for a dual conic (result 2.14(p37))

and a point conic (result 2.13(p37)) lead to the relations

$$\omega^{*i} = H_{\infty}^i \omega^* H_{\infty}^{iT} \quad \text{and} \quad \omega^i = (H_{\infty}^i)^{-T} \omega (H_{\infty}^i)^{-1}. \quad (19.24)$$

where ω^i is the IAC in the i -th view. It may be verified that these equations are precisely the auto-calibration equations (19.4–p461), and this is another geometric interpretation of those equations.

These are among the most important relationships for auto-calibration. They are the basis for obtaining metric reconstruction from affine reconstruction, and also for calibrating a non-translating camera, as will be seen later in section 19.6. The significance of this relation for auto-calibration is that if H_{∞}^i is known, then these are *linear* relations between ω^i and ω (and similarly for ω^*). This means that constraints placed on ω^i in one view can easily be transferred to another and in this manner sufficient constraints may be assembled to determine ω by linear means. Once ω is determined, then K follows by the Cholesky decomposition. We will illustrate this approach for the example of fixed internal parameters.

Sketch solution for identical internal parameters. If the internal parameters are constant over m views then $K^i = K$ and $\omega^{*i} = \omega^*$ for $i = 1, \dots, m$, and the ω^* equation of (19.24) becomes

$$\omega^* = H_{\infty}^i \omega^* H_{\infty}^{iT}. \quad (19.25)$$

A very important point here relates to the scale factors in the equation (19.25).

- Although (19.25) is a relationship between homogeneous quantities, the scale factor in the homogeneous equation can be chosen as unity provided H_{∞}^i is normalized as $\det H_{\infty}^i = 1$.

This results in six equations for the independent elements of the symmetric matrix ω^* . Then (19.25) can be written in a homogeneous *linear* form

$$Ac = 0, \quad (19.26)$$

where A is a 6×6 matrix composed from the elements of H_{∞}^i , and c is the conic ω^* written as a 6-vector. As discussed below, c is not uniquely determined by one such equation since A has at most rank 4. However, if linear equations (19.26) from $m \geq 2$ view pairs are combined, so that A is now a $6m \times 6$ matrix, and provided the rotations between the views are about different axes, then in general c is determined uniquely.

Related to the issue of uniqueness is the issue of numerical stability. Under a single motion the linear computation of K from H_{∞} is extremely sensitive to the accuracy of H_{∞} . If H_{∞} is inaccurate it is not always possible to obtain a positive-definite matrix ω (or ω^*) and thus to apply the Cholesky decomposition to obtain K . This sensitivity is reduced if further motions are made, and ω obtained from the combined constraints of a number of H_{∞} 's.

Advantage of using the IAC. In an analogous manner to the linear solution for ω^* , a linear solution may be obtained for ω starting from (19.24). In fact using the equations involving the IAC is attractive for the following reason. In the zero-skew case the form of (19.12–p464) for the IAC is simpler and more clearly reflects the role of each calibration parameter than does the corresponding formula (19.11–p464) for the DIAC. In order to obtain linear equations using the DIAC equations (19.11) it is necessary to assume that the principal point is known. This is not necessary for equations deriving from the IAC. An assumption of zero skew is quite natural and is a safe assumption for most imaging conditions. However, an assumption of known principal point is much less tenable. For this reason it is usually preferable to use the IAC constraints of (19.24) for auto-calibration rather than using the DIAC constraints.

Other calibration constraints. The algorithm just described was for constant but arbitrary internal parameters. If more is known about K , such as the value of the aspect ratio or that the skew is zero, the corresponding constraints may be simply added to the set of equations on ω (or ω^*), and imposed as soft constraints. The possible constraints are shown in table 19.2(p465) for the DIAC (the same constraints that are used to compute Q_∞^* in section 19.3) and table 19.4 for the IAC.

As mentioned above, the constraints derived from the IAC are generally linear, whereas the constraints derived from the DIAC, are linear only under the assumption that the principal point is known (and at the origin).

Just as with the absolute dual quadric method, it is possible to allow varying internal parameters for the cameras, as long as sufficiently many constraints are imposed. The constraint of constant internal parameters for the cameras imposed a total of $5(m-1)$ constraints on the calibration parameters of the m views. We can make do with fewer constraints, letting certain parameters vary. The method of calibration with varying internal parameters is quite analogous to that used in the case of the absolute dual quadric. Each entry of ω^i is expressible as a linear expression in the entries of ω according to (19.24). A linear constraint on some entry of ω^i therefore maps back to a linear constraint on the entries of ω .

Note: To avoid treating the first image differently, any constraints imposed on the first camera, such as $\omega_{12}^1 = 0$ (camera 1 has zero skew) should be treated by adding this equation to the complete equation set, rather than by decreasing the number of parameters used to describe ω^1 . The latter method would cause the zero-skew constraint to be enforced exactly in the first image (a hard constraint), but it would be a soft constraint in the other images.

Algorithm 19.2 summarizes the stratified method for both constant and varying parameters. One could imagine implementing this algorithm directly for a camera mounted on a robot: the camera is first moved by a pure translation in order to determine π_∞ ; and in subsequent motions the camera may both translate and rotate until sufficient rotations have accumulated to determine K uniquely.

Condition	constraint	type	# constraints
zero skew	$\omega_{12} = 0$	linear	m
principal point at origin	$\omega_{13} = \omega_{23} = 0$	linear	$2m$
known aspect ratio $r = \alpha_y / \alpha_x$ (assuming zero skew)	$\omega_{11} = r^2 \omega_{22}$	linear	m
fixed (unknown) aspect ratio (assuming zero skew)	$\omega_{11}^i / \omega_{22}^i = \omega_{11}^j / \omega_{22}^j$	quadratic	$m - 1$

Table 19.4. **Auto-calibration constraints derived from the IAC.** These constraints are derived directly from the form of (19.10–p464) and (19.12–p464). The number of constraints column gives the total number of constraints over m views, assuming the constraint is true for each view.

Using hard constraints. Algorithm 19.2 comes down to solving a homogeneous set of equations of the form $\mathbf{A}\mathbf{c} = \mathbf{0}$, where $\mathbf{c}$ represents $\boldsymbol{\omega}$ arranged as a 6-vector. Generally the supplied information on $\mathbf{K}$, such as that the skew is zero, will not be satisfied exactly. As discussed in section 8.8(p223) known information can be imposed as a hard constraint by parametrizing $\boldsymbol{\omega}^i$ in each view to satisfy this constraint. For instance if the camera is known to have square pixels then the remaining parameters for the IAC of each view can be represented by a homogeneous 4-vector. Linear equations for the unknown parameters in each view may again be obtained from (19.24). A homogeneous set of equations of the form $\mathbf{A}\mathbf{c} = \mathbf{0}$ may then be assembled, where $\mathbf{c}$ now represents the unknown parameters of $\boldsymbol{\omega}^i$ over all views, and a solution which minimizes $\|\mathbf{A}\mathbf{c}\|$ obtained in the usual manner via the SVD. An alternative is to include all the parameters in each view and use algorithm A5.5(p594) to minimize $\|\mathbf{A}\mathbf{c}\|$, while satisfying constraints $\mathbf{C}\mathbf{c} = \mathbf{0}$ exactly.

19.5.3 The ambiguities in using the infinite homography relation

In this section we describe the ambiguities in determining the internal parameters from (19.25) that occur if only a single rotation axis is used. It will be assumed that the internal parameters are unknown but fixed.

A rotation matrix $\mathbf{R}$ has an eigenvector $\mathbf{d}_r$ with unit eigenvalue, $\mathbf{R}\mathbf{d}_r = 1\mathbf{d}_r$, where $\mathbf{d}_r$ is the direction of the axis of the rotation. Consequently, the matrix $\mathbf{H}'_\infty = \mathbf{K}\mathbf{R}'\mathbf{K}^{-1}$ also has an eigenvector with unit eigenvalue (provided $\mathbf{H}'_\infty$ is normalized as $\det \mathbf{H}'_\infty = 1$). This eigenvector is $\mathbf{v}_r = \mathbf{K}\mathbf{d}_r$, and the image point $\mathbf{v}_r$ corresponds to the vanishing point of the rotation axis direction. Suppose $\boldsymbol{\omega}_{\text{true}}^*$ is the true $\boldsymbol{\omega}^*$; then, it may be verified that if $\boldsymbol{\omega}_{\text{true}}^*$ satisfies (19.25) with $\mathbf{H}'_\infty = \mathbf{K}\mathbf{R}'\mathbf{K}^{-1}$, then so does the one-parameter family of (dual) conics

$$\boldsymbol{\omega}^*(\mu) = \boldsymbol{\omega}_{\text{true}}^* + \mu \mathbf{v}_r \mathbf{v}_r^T \quad (19.27)$$

where μ parametrizes the family. In a similar manner there is a one parameter family (pencil) of solutions for the IAC equation of (19.24). This argument indicates that

Objective

Given a projective reconstruction $\{P^i, X_j\}$, where $P^i = [A^i \mid a^i]$, determine a metric reconstruction via an intermediate affine reconstruction.

Algorithm

- (i) **Affine rectification:** Determine the vector p that defines π_∞ , using one of the methods described in section 19.5.1. At this point an affine reconstruction may be obtained as $\{P^i H_p, H_p^{-1} X_j\}$ with

$$H_p = \begin{bmatrix} I & 0 \\ -p^T & 1 \end{bmatrix}.$$

- (ii) **Infinite homography:** Compute the infinite homography between the reference view and the others as

$$H_\infty^i = (A^i - a^i p^T).$$

Normalize the matrix so that $\det H_\infty^i = 1$.

- (iii) **Compute ω :**

- In the case of constant calibration: rewrite the equations $\omega = (H_\infty^i)^{-T} \omega (H_\infty^i)^{-1}$, $i = 1, \dots, m$ as $A c = 0$ with A a $6m \times 6$ matrix, and c the elements of the conic ω arranged as a 6-vector, **or**
- For variable calibration parameters, use the equation $\omega^i = (H_\infty^i)^{-T} \omega (H_\infty^i)^{-1}$ to express linear constraints on entries of ω^i (e.g. zero skew) as linear equations in the entries of ω .

- (iv) Obtain a least-squares solution to $A c = 0$ via SVD.

- (v) **Metric rectification:** Determine the camera matrix K from the Cholesky decomposition $\omega = (K K^T)^{-1}$. Then a metric reconstruction is obtained as $\{P^i H_p H_\lambda, (H_p H_\lambda)^{-1} X_j\}$ with

$$H_\lambda = \begin{bmatrix} K & 0 \\ 0^T & 1 \end{bmatrix}.$$

- (vi) Use iterative least-squares minimization to improve the solution (see section 19.3.3).

Algorithm 19.2. *Stratified auto-calibration algorithm using IAC constraints.*

although the infinite homography constraint seemingly provides six constraints on the five degrees of freedom of ω^* , only four of these constraints are linearly independent.

Removing the ambiguity. The one-parameter ambiguity may be resolved in several ways. First, if there is another view available related by a rotation around an axis with a direction different to d_r , then the combination of both sets of constraints will not have this ambiguity. A linear solution is easily obtained in the manner of (19.26). Thus with a minimum of three views (i.e. more than one rotation) a unique solution can generally be obtained. A second method of resolving the ambiguity is to make assumptions on the internal parameters of the cameras: for instance an assumption of zero skew (see table 19.4). The equations enforcing zero skew may be added as hard constraints to the set of equations being solved.

An alternative (but equivalent) method enforces the constraints *a posteriori* in the following manner. An ambiguity in solving for c , from the linear equation system

$Ae = 0$, occurs when A has a 2-dimensional (or greater) right null-space. In this case in solving for ω there would be a family of solutions of the form

$$\omega(\alpha) = \omega_1 + \alpha\omega_2.$$

Here ω_1 and ω_2 are known from the null-space generators, and α must be determined. It remains simply to find the value of α that leads to a solution satisfying the chosen constraint condition in table 19.4. This is solved linearly. One could do the same thing solving for the DIAC, but then the constraint condition would be quadratic (see table 19.2(p465)), and one of the solutions would be spurious.

In certain cases, these additional constraints do not resolve the ambiguity. For example, skew-zero does not resolve the ambiguity if the rotation is about the image x - or y -axes. Such exceptions are described in more detail in [Zisserman-98], and we give a few commonly occurring examples now.

Typical ambiguities. The one-parameter family of solutions given in (19.27) for $\omega^*(\mu)$ corresponds to a one-parameter family of calibration matrices obtained from $\omega^*(\mu)$ as $\omega^*(\mu) = K(\mu)K(\mu)^T$. For simplicity we will assume that the true camera K (which is a member of this family) has skew zero, so K has four unknown parameters.

If the rotation axis is parallel to the camera x -axis, then $d_r = (1, 0, 0)^T$ and $v_r = Kd_r = \alpha_x(1, 0, 0)^T$. From the form (19.11-p464) of ω^* with no skew, the family (19.27) is

$$\omega^*(\mu) = \omega_{\text{true}}^* + \mu v_r v_r^T = \begin{bmatrix} \alpha_x^2(1 + \mu) + x_0^2 & x_0 y_0 & x_0 \\ x_0 y_0 & \alpha_y^2 + y_0^2 & y_0 \\ x_0 & y_0 & 1 \end{bmatrix}. \quad (19.28)$$

Note that the entire family has skew-zero, and in this case only the element ω_{11}^* is varying. This means that the principal point and α_y are unambiguously determined – since they may be read-off from elements which are unaffected by the ambiguity. However, it is apparent that α_x cannot be determined because it only appears in the varying element $\omega_{11}^*(\mu)$. To summarize this, and two other canonical cases:

- If K is computed from the infinite homography relation (19.25) assuming a zero-skew camera, then for some motions, there remains one undetermined calibration parameter. For rotation about various axes this ambiguity is as follows.

- (i) **X-axis:** α_x is undetermined;
- (ii) **Y-axis:** α_y is undetermined;
- (iii) **Z-axis (principal axis):** α_x and α_y are undetermined, but their ratio α_y/α_x is determined.

Geometric note. These ambiguities are not limited to calibration from a pair of views, but apply to complete sequences. For instance if the set of rotations in a camera motion are all about the x -axis of the camera, then there will be a reconstruction ambiguity, and

the same is true for Y-axis rotations. One can see this geometrically as follows. Consider a metric reconstruction of a scene from a sequence of images with only Y-axis rotations of the camera. One can define a coordinate system in which the world Y-axis is aligned with the direction of the camera's y -axis. Now, consider "squashing" the whole reconstruction (points and camera positions) so that their Y coordinate is multiplied by some factor k . From the imaging geometry, it is easy to see that this will have the effect of multiplying the y image coordinate of any imaged point by the same factor k , but not affecting the x -coordinate. However, this effect can be undone by multiplying the scale factor α_y of coordinates in the image by the inverse factor k^{-1} , thereby leaving image coordinates unchanged. This shows that α_y is not unambiguously determined, in fact it is unconstrained. In summary there is a one-parameter ambiguity parallel to the rotation axis in the metric reconstruction and a corresponding one-parameter ambiguity in the internal parameters. This argument shows that the problem of ambiguity is intrinsic to the motion, and not to any particular auto-calibration algorithm.

Relationship to the Kruppa equations. Writing (19.24) for two views as $\omega^{*'} = H_\infty \omega^* H_\infty^T$ and multiplying before and after by the matrix $[e']_\times$ leads to

$$[e']_\times \omega^{*'} [e']_\times = [e']_\times H_\infty \omega^* H_\infty^T [e']_\times = F \omega^* F^T$$

since $F = [e']_\times H_\infty$. This is simply the Kruppa equations (19.18–p471), which shows that they follow immediately from the infinite homography constraint. Since $[e']_\times$ is not invertible, one can not go the other direction and derive the infinite homography constraint from the Kruppa equations. Thus, the Kruppa equations are a weaker constraint.

However the difference is that to apply (19.24) one needs to know the plane at infinity (and hence affine structure), since it is true only for the infinite homography, and not for an arbitrary H . The Kruppa equations, on the other hand, do not involve any knowledge of affine structure of the scene. Nevertheless, this relationship shows that for a sequence of images, any calibration ambiguity under the infinite homography relation is also an ambiguity of the Kruppa equations.

19.6 Calibration from rotating cameras

In this section, we begin consideration of calibration under special imaging conditions. The situation considered here is the one in which the camera rotates about its centre but does not translate. We will consider both the case of fixed internal parameters, and the case of some parameters known and fixed whilst others are unknown and varying.

This situation is one that occurs frequently. Examples include: pan-tilt and zoom surveillance cameras; cameras used for broadcasts of sporting events which are almost invariably fixed in location but free to rotate and zoom; and hand-held camcorders which are very often panned from a single viewpoint. Even though the rotation is not exactly about the centre, in practice the translation is generally negligible compared to the distance of scene points, and a fixed centre is an excellent approximation.

The calibration problem from rotating cameras is mathematically identical with the affine-to-metric calibration step in stratified reconstruction, as given in section 19.5.2.

Objective

Given $m \geq 2$ views acquired by a camera rotating about its centre with fixed or varying internal parameters, compute the parameters of each camera. It is assumed that the rotations are not all about the same axis.

Algorithm

- (i) **Inter-image homographies:** Compute the homography H^i between each view i and a reference view such that $\mathbf{x}^i = H^i \mathbf{x}$ using, for example, algorithm 4.6(p123). Normalize the matrices such that $\det H^i = 1$.
- (ii) **Compute ω :**
 - In the case of constant calibration: rewrite the equations $\omega = (H^i)^{-T} \omega (H^i)^{-1}$, $i = 1, \dots, m$ as $A\mathbf{c} = \mathbf{0}$ where A is a $6m \times 6$ matrix, and $\mathbf{c}$ the elements of the conic ω arranged as a 6-vector, **or**
 - For variable calibration parameters, use the equation $\omega^i = (H^i)^{-T} \omega (H^i)^{-1}$ to express linear constraints on entries of ω^i in table 19.4 (e.g. unit aspect ratio) as linear equations in the entries of ω .
- (iii) **Compute K :** Determine the Cholesky decomposition of ω as $\omega = UU^T$, and thence $K = U^{-T}$.
- (iv) **Iterative improvement:** (Optional) Refine the linear estimate of K by minimizing

$$\sum_{i=2, m; j=1, n} d(\mathbf{x}_j^i, K R^i K^{-1} \mathbf{x}_j)^2$$

over K and R^i , where $\mathbf{x}_j, \mathbf{x}_j^i$ are the position of the j -th point measured in the first and i -th images respectively. Initial estimates for the minimization are obtained from K and $R^i = K^{-1} H^i K$.

Algorithm 19.3. Calibration for a camera rotating about its centre.

From a non-translating camera it is impossible to achieve an affine (or any) reconstruction, because there is no way to resolve depth. Nevertheless, we may compute the infinite homography between the images, which is all that is needed to determine the camera calibration.

As has been shown earlier (section 8.4(p202)) the images of two cameras with a common centre are related by a plane projective transformation. Indeed, if $\mathbf{x}^i$ and $\mathbf{x}$ are corresponding image points then they are related by $\mathbf{x}^i = H^i \mathbf{x}$, where $H^i = K^i R^i (K)^{-1}$, and R^i is the rotation between view i and the reference view. Furthermore, since this map is independent of the depth of the points imaged at $\mathbf{x}$, it applies also to points at infinity, so as shown in section 13.4(p338),

$$H^i = H_{\infty}^i = K^i R^i (K)^{-1}.$$

Thus we have a convenient means of measuring H_{∞} directly from images.

Given H_{∞} a solution for the calibration matrices K^i of all the images in the set acquired by the rotating camera may be obtained as described in section 19.5.2. The method may be applied to either fixed or variable internal parameters and is summarized in algorithm 19.3. We will illustrate this by a number of examples.



Fig. 19.3. **Calibrating a camera rotating about its centre.** (upper) Five images of the US Capitol acquired by approximately rotating the camera about its centre. (lower) A mosaic image constructed from the five images (see example 8.14(p207)). The mosaic image shows very clearly the distortion effect of the infinite homography between the images. Analysis of this distortion provides the basis for the auto-calibration algorithm. The calibration is computed as described in algorithm 19.3.

Example 19.9. Rotation about centre with fixed internal parameters

The images in figure 19.3 were obtained using a 35mm camera with ordinary black and white film to produce negatives. The camera was hand-held, and no particular care was taken to ensure that the camera centre remained stationary.

Prints enlarged from these negatives were digitized using a flat-bed scanner. The enlargement process can lead to a non-zero value of s and unequal values of α_x and α_y if the negative and print paper are not precisely parallel. The resulting image size was 776×536 pixels.

The constraint applied here is that the internal parameters are constant. The camera




Fig. 19.4. **Rotation for varying internal parameters assuming square pixels.** These are for the panned sequence of figure 8.9(p206). (a) Focal length. (b) Principal point. (c) Pan angle. (d) Tilt angle. Figures courtesy of Lourdes de Agapito Vicente.

matrix computed as described in algorithm 19.3 is

$$K_{\text{linear}} = \begin{bmatrix} 964.4 & -4.9 & 392.8 \\ & 966.4 & 282.0 \\ & & 1 \end{bmatrix} \quad K_{\text{iterative}} = \begin{bmatrix} 956.8 & -6.4 & 392.0 \\ & 959.3 & 281.4 \\ & & 1 \end{bmatrix}.$$

There is little difference between the linear and iterative estimates, and the computed aspect ratio (virtually unity) and principal point are very reasonable. $\triangle$

Example 19.10. Rotation about centre with varying internal parameters

The images in figure 8.9(p206) were acquired by panning a camcorder approximately about its centre. The camera was not zoomed, but due to auto-focus there might be slight variations in the focal length and principal point.

In this example the constraint used is that the pixels are square: i.e. that the skew is zero and the aspect ratio unity; but the focal length and principal point are unknown and not fixed. Then from table 19.4 we have two linear constraints on ω from each view. From zero skew $(H_{\infty}^i)^{-T} \omega (H_{\infty}^i)^{-1})_{12} = 0$, and from unit aspect ratio $(H_{\infty}^i)^{-T} \omega (H_{\infty}^i)^{-1})_{11} = (H_{\infty}^i)^{-T} \omega (H_{\infty}^i)^{-1})_{22}$. These constraints are assembled to give a linear system of equations on ω as described in algorithm 19.3.

The internal parameters of the computed camera matrix for each view are shown in figure 19.4. It is evident that the recovered focal length and principal point are quite constant (even though this was not imposed), and the pan and tilt angles are very reasonable for this hand-held sequence. $\triangle$

19.7 Auto-calibration from planes

For a set of images of a planar scene, the two-step approach of estimating a projective reconstruction followed by computation of a rectifying transformation to take it to metric does not work. This is because it is not possible to determine the cameras without depth relief. As seen in section 17.5.2(p425), a minimum of two points not lying on the plane are required. Nevertheless, auto-calibration from scene planes *is* possible. This was shown by [Triggs-98] who gave a solution in the case of constant internal parameters. The method is especially interesting from the point of view of potential applications. Scenes consisting of planes are extremely common in man-made environments, e.g. the ground plane. Furthermore, in aerial images acquired by a high-flying aircraft or a satellite, the depth-relief of the scene is small compared with the extent of the image, and the scene may be accurately approximated as a plane and the auto-calibration method will apply.

The starting point of the algorithm is a set of image-to-image homographies induced by the world plane. These can all be related back to the first image, providing a set of homographies H^i . Geometrically auto-calibration from planes is then a marriage of two ideas. First, the circular points on the plane, which are the intersection of the absolute conic with the plane, are mapped from image to image via the homographies. Second, as we have seen in example 8.18(p211), the calibration matrix K may be determined from the imaged circular points of a plane (with two constraints provided by each image).

Thus suppose the images of the circular points (4 dof) are determined in the first image (by some means), then they may be transferred to the other views by the known H^i . In each view then there are two constraints on ω , since the two imaged circular points lie on ω . In detail, if we denote the (at this stage unknown) imaged circular points in the first view by $c_j, j = 1, 2$. Then the auto-calibration equations are

$$(H^i c_j)^T \omega^i (H^i c_j) = 0, \quad i = 1, \dots, m \quad j = 1, 2 \quad (19.29)$$

where $H^1 = I$. In solving these equations the unknowns are the coordinates of the circular points in the first image, and some number of unknown calibration parameters. Although the circular points are complex points, they are complex conjugates of each other, so 4 parameters suffice to describe them. If in addition there are a total of v unknowns in the internal parameters K^i of all m views, then a solution is possible provided $2m \geq v + 4$, since each view provides two equations.

Restrictions on the internal parameters of the cameras lead to further algebraic constraints according to table 19.4 and these are used to supplement the constraints imposed by (19.29). Various cases are considered in table 19.5. In most cases, calibration from a plane is a non-linear problem and nothing more will be said here about computational aspects of how to find the solution. Iterative methods are necessary, minimizing some cost function. See [Triggs-98] for more details of minimization methods.

Implementation. A considerable implementational advantage of this method is that it only requires the homography between planes, and not the point correspondences

Condition	dof(v)	views
Unknown but constant internal parameters	5	5
Constant known skew and aspect ratio. Constant unknown principal point and focal length	3	4
All internal parameters known except varying focal length	m	4
Varying focal length, all other internal parameters fixed but unknown	$m + 4$	8

Table 19.5. **The number of views required for calibration from a plane under various conditions.** Calibration is (in principle) possible if $2m \geq v + 4$.

arising from 3-space points that are generally required to estimate multiple view tensors, such as the fundamental matrix. The matching transformation between planes is a much simpler, stabler and accurate computation because of the constrained nature of the inter-image transformation which is point-to-point. The method of algorithm 4.6-(p123) may be used to estimate this transformation between two images. Alternatively, correlation-based methods that estimate the parametrized homography directly from image intensity may be used.

Including additional information. If additional information is available on the plane or the motion then the complexity of auto-calibration using scene planes may be reduced. For example, if the vanishing line of the imaged plane can be determined then only two parameters are required to specify the circular points since the imaged circular points lie on the vanishing line. Indeed if the plane provides sufficient information to estimate its imaged circular points directly (such as a square grid) then the problem reduces to that of calibration from scene constraints discussed in section 8.8(p223).

Similarly constraints on the motion may be used to simplify the problem. One particular example is where the rotation axis describing the camera motion is parallel to the scene plane normal. In this case the imaged circular points may be computed directly from the fixed points of the homography and two linear constraints placed on ω . This situation is discussed further in Note (vii) at the end of the chapter. An example of such a motion is planar motion which is discussed in detail in the following section.

19.8 Planar motion

A case of some practical importance is that of a camera moving in a plane and rotating about an axis perpendicular to that plane. This is the case for a roaming vehicle moving on a ground plane, with a camera fixed with respect to the vehicle body. In this case, the camera must move in a plane parallel to the (horizontal) ground, and as the vehicle turns, the camera will rotate about a vertical axis. It is *not* assumed that the camera is pointing horizontally, or is in any other particular orientation with respect to the vehicle. However, we assume constant internal calibration for the camera. The constrained nature of the motion makes the calibration task significantly simpler.

It will be shown that given three or more images from a sequence of planar motions an affine reconstruction may be computed. To do this, we need to determine the plane at infinity. This will be done by identifying three points on the plane at infinity, thereby defining the plane. These points will be identified as being fixed points in the sequence of images.

Fixed image points under planar motion. According to section 3.4.1(p77), any rigid motion (for instance the camera motion) can be interpreted as a rotation about screw axis along with a translation along the direction of the axis. For a planar motion, the rotation axis is perpendicular to the plane of motion, and the translational part of the screw motion is zero. Think of the vehicle as being swung horizontally around the screw axis. The position of the screw axis with respect to the camera remains fixed, and so it will constitute a line of fixed points in the image. If a second motion takes place about a different axis, then the image of the second axis will be fixed in the second pair of images. The images of the two axes will in general be different, but will intersect in the image of the point where the two screw axes meet. Since the two axes are vertical, they are perforce parallel, and so will meet at their common direction on the plane at infinity. This direction projects into the images at the intersection of the images of the screw axes, which is the vanishing point of the screw axes direction. This image point is called the apex. We now have one fixed point over the views, and this will be used to determine one point on the plane at infinity in a projective reconstruction.

As we have seen the image of the screw axis is a line of fixed points for an image pair. There is also a fixed line which is identifiable from a pair of images as follows. Because the plane of motion of the camera (which will be called the ground plane) is fixed, the set of points in the plane is mapped to the same line in all the images. This line is called the horizon line and is the vanishing line of the ground plane. Since each of the cameras lie on the ground plane their epipoles must all lie on the horizon line. Unlike the image of the screw axis, the horizon line is a fixed line, but not a line of fixed points.

Although the horizon line is not fixed pointwise, it contains two points that are fixed in the image pair, namely the image of the two circular points on the ground plane. These circular points are the intersection of the absolute conic with the ground plane. Since the image of the absolute conic is fixed under rigid motion, and the image of the ground plane is fixed, the images of two circular points must be fixed. In fact, they will be fixed in all images from the planar motion sequence. This is illustrated in figure 19.5.

So far we have described the fixed points of the motion sequence. Computing these fixed points is equivalent to affine reconstruction, since we can back-project from the fixed image points to find the corresponding 3D points on the plane at infinity. Although the apex may be computed from two views; it requires three views to compute the imaged circular points.

Computing the fixed points. The set of points in space that map to the same image point in two images is called the horopter. In general, the horopter is a twisted cubic,




Fig. 19.5. **Fixed image entities for planar motion.** (a) For two views the imaged screw axis is a line of fixed points in the image under the motion. The horizon is a fixed line under the motion. The epipoles e, e' and imaged circular points of the ground plane c_i, c_j lie on the horizon. (b) The relation between the fixed lines obtained pairwise for three images under planar motion. The image horizon lines for each pair are coincident, and the imaged screw axes for each pair intersect in the apex v . All the epipoles lie on the horizon line.

but in the case of planar motion it degenerates to a line (the screw axis) and a conic on the ground plane. The image of the horopter is the conic defined by $F + F^T$, where F is the fundamental matrix (see section 9.4(p250)). In the case of planar motion this will be a degenerate conic consisting of two lines, the image of the screw axis and the horizon line (see figure 9.11(p253)). By decomposing the conic, these two lines are determined. From three images we can compute the image of the horopter for each of three pairs, and thereby obtain three sets of horizon lines and imaged axes. The horizon line will be a common component of these sets, and the other components (the images of the screw axes) will intersect in the image at the apex.

Now we turn to computing the circular points. It is useful to understand the geometry of a pair of horopters corresponding to two pairs of images. Let C^{12} be a conic representing the portion of the horopter for images 1 and 2, lying in the ground plane. This conic will pass through the two circular points, the two camera centres and the intersection of the screw axis with the ground plane. Since the conic contains the circular points it is a circle. Let C^{23} be the corresponding circle defined from images 2 and 3. The two circular points and the centre of the second camera will lie on both circles. Since two conics intersect in general in four points, there must be a further (real) intersection point. However, this can be discarded, since the points of interest are the two complex intersection points, namely the circular points on the ground plane.

In implementations, authors have chosen different ways of finding the two circular points. In [Armstrong-96b] the method is based on finding fixed lines in three views through the apex using the trifocal tensor. In [Faugeras-98, Sturm-97b] the method involves computing the trifocal tensor of a 1D camera, applied to imaging the 2D ground plane to a 1D image. In both cases the positions of the circular points on the horizon are obtained as the solution of a cubic in one variable.

The main steps of this affine calibration method are summarized in algorithm 19.4.

Objective

Given three (or more) images acquired by a camera with constant internal parameters undergoing planar motion compute an affine reconstruction.

Algorithm

- (i) **Compute a projective reconstruction.** from the trifocal tensor $\mathcal{T}$ for the three views. The trifocal tensor may be computed by, e.g. algorithm 16.4(p401).
- (ii) **Compute pairwise fundamental matrices from $\mathcal{T}$.** See algorithm 15.1(p375). Decompose the symmetric part of each fundamental matrix into two lines, a horizon and the image of the screw axes. See section 9.4.1(p252).
- (iii) **Compute the apex.** Intersect the three imaged screw axes to determine the apex $\mathbf{v}$.
- (iv) **Compute the horizon for the triplet.** Obtain the six epipoles from the three fundamental matrices, and determine the horizon by an orthogonal regression fit to these epipoles.
- (v) **Compute the imaged circular points.** Compute the position of the imaged circular points on the horizon (see text) $\mathbf{c}_i, \mathbf{c}_j$.
- (vi) **Compute the plane at infinity.** Triangulate points on the plane at infinity from the corresponding image points $\mathbf{x} \mapsto \mathbf{x}'$, with $\mathbf{x} = \mathbf{x}'$ for each of $\mathbf{v}, \mathbf{c}_i, \mathbf{c}_j$. This determines three points on π_∞ , and hence determines the plane.
- (vii) **Compute an affine reconstruction.** Rectify the projective reconstruction using the computed π_∞ as in algorithm 19.2.

Algorithm 19.4. *Affine calibration for planar motion.*

Metric reconstruction. Once the apex and two circular points are found, the plane at infinity is computed, and the infinite homographies between the images can be computed. Calibration and metric reconstruction now proceeds in the usual way. However, it must be noted that the constrained nature of the motion means that there is a one-parameter family of solutions for the calibration because all the camera rotations are about the same axis. We have the sort of calibration ambiguity considered in section 19.5.3. It is necessary to make an assumption about internal calibration in order to find a unique result. If the y -axis of the camera is parallel with the rotation axis (which may be true in a practical situation), then we have seen that the zero-skew constraint is not sufficient. The best plan is to enforce a zero-skew and known aspect ratio constraint (e.g. if the pixels are square).

Example 19.11. Metric reconstruction for planar motion.

Figure 19.6(a) shows four of seven images of a planar motion sequence. For this sequence the elevation angle of the camera is approximately 20° . The computed imaged screw axes and horizon lines using all possible pairs are shown in figure 19.6(b), with the resulting estimated apex and horizon line. The positions of the imaged circular points were estimated as $x = 104 \pm 362i$, $y = -86 \mp 2i$. Assuming an aspect ratio of 1.1, the internal parameters of the calibration matrix $\mathbf{K}$ were computed as $\alpha_x = 330$, $\alpha_y = 363$, $x_0 = 123$, $y_0 = 50$. The accuracy of the metric reconstruction was assessed by measuring metric invariants. Typical results are shown in figure 19.6(c). $\triangle$




Fig. 19.6. **Planar motion sequence.** (a) Four images (of seven used) acquired from a camera mounted on a vehicle moving on a plane. (b) The computed epipoles ($\times$), horizon lines (grey solid), and imaged screw axes (grey dashed) for all image pairs. (c) Euclidean invariants measured in the metric reconstruction the right angle is measured as 89° , the ratio of the non-parallel lengths is measured as 0.61 (compared to the approximate veridical value of 0.65).

19.9 Single axis rotation – turntable motion

In this section we discuss auto-calibration of single axis motions where the relative motion between scene and cameras is equivalent to a rotation about a single fixed axis. This is a specialization of the planar motion case of section 19.8 where here the screw axes for each motion are coincident. This situation occurs, for example, in the case of a static camera viewing an object rotating on a turntable. A second example is that of a camera rotating about a fixed axis (offset from its centre). A third example is that of a camera viewing a rotating mirror.

We will consider here turntable motion and for ease of discussion assume that the axis is vertical so that the motion occurs in a horizontal plane. Again it is *not* assumed that the camera is pointing horizontally, or is in any other particular orientation with respect to the axis. It is assumed that the internal calibration of the camera is constant.

The fixed image points under a sequence of single axis rotation are those of planar motion described in section 19.8, and shown in figure 19.5(a), with the addition that the imaged screw axis is a line of fixed points. The constraint of figure 19.5(b) is not available here as all the imaged screw axes are coincident. The consequence is that it is not possible to determine the apex v directly, and only the two circular points on π_∞ may be recovered from imaged fixed points. This means that the reconstruction, in the absence of constraints on the internal parameters, is not affine, but a particular

parametrized projective transformation. Metric structure is known in the horizontal planes (since the circular points of these planes is known) but there is a 1D projective transformation in the vertical (Z) direction. The resulting ambiguity is the transformation $\mathbf{X}_p = \mathbf{H}\mathbf{X}_e$ where

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \gamma & 0 \\ 0 & 0 & \delta & 1 \end{bmatrix} \quad (19.30)$$

and γ and δ are scalars which determine the intersection of the Z axis with π_∞ and the relative scaling between the horizontal and vertical directions. An example of the projective transformations represented by family of mappings is shown in figure 1.4-(p16).

Computing the fixed points. One way to proceed is to determine the imaged circular points directly. As is evident from figure 19.7(b) the point tracks are ellipses which are images of circles. In 3-space these circles lie in parallel horizontal planes and intersect in the circular points on π_∞ . In the image, ellipses may be fitted to these tracks, and the common intersection points of the image conics are the (complex conjugate) imaged circular points. The 3D circular points may then be determined by triangulation from two or more views. This is the approach taken by [Jiang-02].

An alternative more algebraic way to proceed is to model the camera matrices as $\mathbf{P}^i = \mathbf{H}_{3 \times 3}[\mathbf{R}_Z(\theta^i) \mid \mathbf{t}]$ where

$$\mathbf{P}^i = \begin{bmatrix} \mathbf{h}_1 & \mathbf{h}_2 & \mathbf{h}_3 \end{bmatrix} \begin{bmatrix} \cos \theta^i & -\sin \theta^i & 0 & t \\ \sin \theta^i & \cos \theta^i & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (19.31)$$

with $\mathbf{h}_k$ the columns of $\mathbf{H}_{3 \times 3}$. This division of the internal and external parameters means that $\mathbf{H}_{3 \times 3}$ and $\mathbf{t}$ are fixed over the sequence, and only the angle of rotation, θ^i , about the Z axis varies for each camera $\mathbf{P}^i$.

Given this parametrization, the estimation problem can then be precisely stated as determining the common matrix $\mathbf{H}_{3 \times 3}$ and the angles θ_i in order to estimate the set of cameras $\mathbf{P}^i$ for the sequence. Thus a total of $8 + m$ parameters must be estimated for m views, where 8 is the number of degrees of freedom of the homography $\mathbf{H}_{3 \times 3}$. This is a considerable saving over the $11m$ that would be required for a projective reconstruction of a general motion sequence.

For single axis motion the fundamental matrix has the special form described in section 9.4.1(p252), and writing the matrix in terms of the camera matrices (19.31) gives

$$\mathbf{F} = \alpha[\mathbf{h}_2]_\times + \beta \left((\mathbf{h}_1 \times \mathbf{h}_3)(\mathbf{h}_1 \times \mathbf{h}_2)^T + (\mathbf{h}_1 \times \mathbf{h}_2)(\mathbf{h}_1 \times \mathbf{h}_3)^T \right)$$

which means that the columns of $\mathbf{H}_{3 \times 3}$ are partially determined once the fundamental matrix is computed. This is the approach taken by [Fitzgibbon-98b] where a fundamental matrix of the special form is fitted to point correspondences (see section 11.7.2-(p293)). It may be shown that from 3 or more views the first two columns $\mathbf{h}_1, \mathbf{h}_2$ of




Fig. 19.7. **Dinosaur sequence and tracking:** (a) six frames from a sequence of 36 of a dinosaur rotating on a turntable (Image sequence courtesy of the University of Hannover [Niem-94]). (b) A subset of the point tracks—only the 200 tracks which survived for longer than 7 successive views are shown. (c) **Track lifetimes:** Each vertical bar corresponds to a single point track, extending from the first to last frame in which that point was seen. The horizontal ordering is according to where the point first appeared in the sequence. The measurement matrix is relatively sparse, and few points survive longer than 15 frames.

$H_{3 \times 3}$ and the angle θ^i can be fully determined, but h_3 is only determined up to the two-parameter family corresponding to the ambiguity of (19.30).

Metric reconstruction. The two parameter ambiguity in the reconstruction given by (19.30) can be resolved by providing additional information on the internal parameters, for example that the pixels are square. However, if the camera is horizontal with image y axis parallel to the rotation axis, then square pixels only provides one additional constraint (from the aspect ratio, as the skew zero constraint does not provide an additional constraint). In this case further information on the camera is required (e.g. the y coordinate of the principal point) or the aspect ratio of the scene may be used.

Example 19.12. Reconstruction from a turntable sequence

Figure 19.7 shows frames from a sequence of a model dinosaur rotating on a turntable, and the resulting image point tracks. Feature extraction is performed on the luminance component of the colour signal. The projective geometry of the turntable motion is de-




Fig. 19.8. **Dinosaur:** (a) 3D scene points (about 5000) and camera positions for the *Dinosaur* sequence. (b) The automatic computation of a 3D graphical model for this sequence is described in [Fitzgibbon-98b].

terminated from these tracks (and no other information) and the resulting reconstruction of cameras and 3D points is shown in figure 19.8. Effectively the camera circumnavigates the object. A 3D texture mapped graphical model may then be computed, in principle, by back-projecting cones defined by the dinosaur silhouette in each frame, and intersecting the set of cones to determine the visual hull of the 3D object. $\triangle$

19.10 Auto-calibration of a stereo rig

In this section we describe a stratified method for calibrating a “fixed” two-camera stereo rig. Fixed here means that the relative orientation of the cameras on the rig is unchanged during the motion, and the internal parameters of each camera are also

unchanged. It will be shown that from a single motion of the rig the plane at infinity is determined uniquely.

Suppose a fixed stereo rig undergoes a general motion. The projective structure of the scene can be obtained before ($\mathbf{X}$) and after ($\mathbf{X}'$) the motion. Since $\mathbf{X}$ and $\mathbf{X}'$ are two projective reconstructions of the same scene they are related by a 4×4 projective transformation H_P , as

$$\mathbf{X}' = H_P \mathbf{X}.$$

However, the actual motion of the rig is Euclidean, and it follows (see below) that the homography H_P is conjugate to the Euclidean transformation representing the motion. Conjugacy is the key result because under a conjugacy relation fixed entities are mapped to fixed entities. Consequently the fixed entities of the Euclidean motion (in particular the plane at infinity) can be *accessed* from the fixed entities of the projective motion represented by H_P .

Conjugacy relation. Suppose $\mathbf{X}_E$ represents a point in 3-space in a Euclidean coordinate frame attached to the rig, and $\mathbf{X}'_E$ represents the same point after the motion of the rig. Then the points are related as

$$\mathbf{X}'_E = H_E \mathbf{X}_E \quad (19.32)$$

where H_E is a non-singular 4×4 Euclidean transformation matrix that represents the rotation and translation of the rig. Suppose also that the point is represented in a projective coordinate frame attached to the rig (and which is obtained by a projective reconstruction); then

$$\mathbf{X}_E = H_{EP} \mathbf{X} \quad \mathbf{X}'_E = H_{EP} \mathbf{X}' \quad (19.33)$$

where H_{EP} is a non-singular 4×4 matrix that relates projective to metric structure. An essential point to note here is that the two projective reconstructions, before and after the camera must be in the same projective coordinate frame, in other words, the same pair of cameras matrices must be used before and after.

From (19.32) and (19.33) it follows that

$$H_P = H_{EP}^{-1} H_E H_{EP} \quad (19.34)$$

so that H_P is conjugate to a Euclidean transformation. There are two important properties of this conjugacy relation:

- (i) H_P and H_E have the same eigenvalues.
- (ii) If $\mathbf{E}$ is an eigenvector of H_E then the corresponding eigenvector of H_P , with the same eigenvalue, is $(H_{EP}^{-1} \mathbf{E})$, i.e. the eigenvectors of H_E are mapped to the eigenvectors of H_P by the point transformation (19.33). This follows from (19.34), for if $H_E \mathbf{E} = \lambda \mathbf{E}$ then $H_{EP} H_P H_{EP}^{-1} \mathbf{E} = \lambda \mathbf{E}$, and pre-multiplying by H_{EP}^{-1} gives the desired result.

Fixed points of a Euclidean transformation. Consider the Euclidean transformation represented by the matrix

$$H_E = \begin{bmatrix} R & t \\ 0^T & 1 \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 & 0 \\ \sin \theta & \cos \theta & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

This is a rotation by θ about the Z-axis together with a unit translation along the Z-axis (it is a general screw motion). The eigenvectors of H_E are the fixed points under the transformation (refer to section 2.9(p61)). In this case the eigenvalues are $\{e^{i\theta}, e^{-i\theta}, 1, 1\}$ and the corresponding eigenvectors of H_E are

$$E_1 = \begin{pmatrix} 1 \\ i \\ 0 \\ 0 \end{pmatrix} \quad E_2 = \begin{pmatrix} 1 \\ -i \\ 0 \\ 0 \end{pmatrix} \quad E_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \quad E_4 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}.$$

All the eigenvectors lie on π_∞ . This means that π_∞ is fixed as a set, but is *not* a plane of fixed points. The eigenvectors E_1 and E_2 are the circular points for planes perpendicular to the Z (rotation) axis. The other two (identical) eigenvectors E_3 and E_4 are the direction of the rotation axis.

Computing π_∞ . If the point transformation matrix is H_E then from (3.6-p68) the plane transformation matrix is H_E^{-T} . The eigenvectors of H_E^{-T} are the fixed planes under the motion. The matrix H_E^{-T} also has two equal, unit eigenvalues and a single eigenvector corresponding to these which is the plane π_∞ as may easily be verified. The eigenvectors of H_E^{-T} are mapped to those of H_p^{-T} , in the same manner as the mapping of eigenvectors of H_E to those of H_p described above. Consequently, π_∞ in the projective reconstruction is the eigenvector corresponding to the (double) real eigenvalue of H_p^{-T} . Thus,

- π_∞ may be computed uniquely as the real eigenvector of H_p^{-T} , or equivalently, and more simply, as the real eigenvector of H_p^T .

We observe here that although the real eigenvalue has algebraic multiplicity of two, its geometric multiplicity (in the case of non-planar motions) is one. This is what enables us to find the plane at infinity.

The procedure for affine calibration is summarized in algorithm 19.5.

Metric calibration and ambiguities. Once π_∞ is identified, the metric calibration may proceed as described in the stratified algorithm 19.2(p479). Since the rig is fixed, the parameters of the left camera are unchanged during the motion (and similarly for the right camera). From a single motion the internal parameters of each camera are determined up to the one-parameter family resulting from a single rotation as described in section 19.5.3.

As usual the ambiguity from a single motion is removed by additional motions or

Objective

Given two (or more) stereo pairs of images acquired by a fixed stereo rig undergoing general motions (i.e. both R and t are non-zero, and t not perpendicular to the axis of R), compute an affine reconstruction.

Algorithm

- (i) **Compute an initial projective reconstruction X :** Using the first stereo pair compute a projective reconstruction ($P^L, P^R, \{X_j\}$) as described in chapter 10. This involves computing the fundamental matrix F and point correspondences between the images of the first pair $x_j^L \leftrightarrow x_j^R$, e.g. use algorithm 11.4(p291).
- (ii) **Compute a projective reconstruction X' after the motion:** Compute correspondences between the images of the second stereo pair $x_j'^L \leftrightarrow x_j'^R$. Since both the internal and relative external parameters of the cameras are fixed, the second stereo pair has the same fundamental matrix F as the first. The same camera matrices P^L, P^R are used for triangulating points X'_j in 3-space from the computed correspondences $x_j'^L \leftrightarrow x_j'^R$ in the second stereo pair.
- (iii) **Compute the 4×4 matrix H_p which relates X to X' :** Compute correspondences between the left images of the two stereo pairs $x_j^L \leftrightarrow x_j'^L$ (e.g. again use algorithm 11.4(p291)). This establishes correspondences between the space points $X_j \leftrightarrow X'_j$. The homography H_p may be estimated linearly from five or more of these 3-space point correspondences, and then the estimate refined by minimizing a suitable cost function over H_p . For example, minimizing $\sum_j (d(x_j^L, P^L H X'_j)^2 + d(x_j^R, P^R H X'_j)^2)$ minimizes the distance between the measured and reprojected image points.
- (iv) **Affine reconstruction:** Compute π_∞ from the real eigenvector of H_p^T and thence an affine reconstruction.

Algorithm 19.5. *Affine calibration of a fixed stereo rig.*

by supplying additional constraints, such as that the pixels are square. If there are additional motions then an improved estimate of π_∞ may also be computed. The outcome of metric calibration is the complete calibration of the rig (i.e. the relative external orientation of the cameras and their internal parameters).

Planar motion. In the special case of orthogonal (planar) motion, where the translation is orthogonal to the rotation axis direction, the space of eigenvectors corresponding to the repeated (real) eigenvalue is two-dimensional. Consequently, π_∞ is determined only up to a one-parameter family. We are therefore unable to find the plane at infinity uniquely (this is examined in detail in example 3.8(p81)). The ambiguity may be resolved by a second orthogonal motion about an axis with a different direction from the first.

Example 19.13. Auto-calibration from two stereo pairs

Figure 19.9(a)(b) shows the two stereo pairs used for the affine calibration of the stereo rig following the procedure of algorithm 19.5. The accuracy of the calibration is assessed by computing a vanishing point in the right image in two ways: first, as the intersection of imaged parallel lines; second, by determining the corresponding vanishing point in the left image (from images of the same parallel lines), and mapping this vanishing point to the right image using the infinite homography computed from




Fig. 19.9. **Auto-calibration of a stereo rig.** The input stereo pairs before (a) and after (b) the motion of the rig. The stereo rig moves left by about 20 cm, pans by about 10° and changes elevation by about 10° . The accuracy of the computed H_∞ is assessed on another stereo pair acquired by the same rig as follows: In (c), the left image (of a stereo pair), a vanishing point is computed by intersecting the imaged sides of the table (which are parallel in the scene). In (d), the right image (of a stereo pair), the corresponding vanishing point is computed. The white square (near the line intersection) is the vanishing point from the left image mapped to the right using the computed H_∞ . In the absence of error the points should be identical. (e) Following metric calibration the computed angle between the desk sides (shown in white) from the 3D reconstruction is 90.7° , in very good agreement with the veridical value.

the eigenvector of H_p^T . The discrepancy between the vanishing points is a measure of the accuracy of the computed H_∞ . The metric calibration uses the zero skew constraint to resolve the one-parameter ambiguity. Angles in the resulting metric reconstruction are accurate to within 1° . $\triangle$

19.11 Closure

Research in the area of auto-calibration is still quite active, and better methods than those described in this chapter may yet be developed. There is still a lack of closed form solutions from multiview tensors, and of algorithms to automatically detect critical motion sequences (see below).

Critical motion sequences. It has been seen in this chapter that for certain classes of motion it is not possible to completely determine the rectifying homography H . The resulting reconstruction is then at some level between metric and projective. For example, for constant internal parameters in the case of planar motion there is a one-parameter scaling ambiguity parallel to the rotation axis; and for pure translation under constant internal parameters the reconstruction is affine. Sequences of camera motions for which such ambiguities arise are termed “Critical motion sequences” and have been systematically classified by Sturm [Sturm-97a, Sturm-97b] in the case of constant internal parameters. This classification has been extended to more general calibration constraints, such as varying focal lengths [Pollefeys-99b, Sturm-99b]. For recent work see [Kahl-99, Kahl-01b, Ma-99, Sturm-01].

Recommendations. It may seem that auto-calibration offers a complete solution to metric reconstruction. Calibrated cameras are not necessary and we can do with constraints as weak as the zero-skew constraint on the camera. Unfortunately one must be wary of putting complete trust in auto-calibration. Auto-calibration can work well in the right circumstances, but used recklessly it will fail. Several specific recommendations can be made.

- (i) Take care to avoid ambiguous motion sequences. It has been seen that calibration degeneracies occur if the motion is too restricted, such as being about a single axis. The motion should not be too small, or cover too small a field of view. Auto-calibration often comes down to estimating the infinite homography, the effects of which are not apparent on small fields of view.
- (ii) Use as much information as you have. Although it is possible to calibrate from minimal information such as zero skew, this should be avoided if other information is available. For instance the known aspect ratio constraint should be used if it is valid, as should the knowledge of the principal point. Even if the values are known only imprecisely, this information can be incorporated into a linear auto-calibration method by including an equation, but with low weight.
- (iii) This recommendation relates to bundle adjustment as well. Generally it is best to finish off with a bundle adjustment. In doing this, it is recommended that the internal parameters of the camera not be left to float unbounded. For instance, even if the principal point is not known exactly, it is usually known within some reasonable bounds (it is not close to infinity for instance). Similarly, aspect ratio normally lies between 0.5 and 3. This knowledge should be incorporated in a bundle adjustment by adding further constraints to the cost function, with small weights (standard deviations) if necessary. This can give an enormous improvement in results where auto-calibration is poorly conditioned (and hence unstable) by preventing the solution from wandering off into remote regions of parameter space in quest of a minor and insignificant improvement in the cost function.
- (iv) Methods that use restricted motions usually are more reliable than those that allow general motion. For instance the methods that involve a rotating but non-translating camera are generally much more reliable than general motion methods. The same is true of affine reconstruction from a translational motion.

19.11.1 The literature

The idea of auto calibrating a camera originated in Faugeras *et al.* [Faugeras-92a] where the Kruppa equations were used. The early papers considered the case of constant internal parameters. [Hartley-94b] and Mohr *et al.* [Mohr-93] investigated bundle-adjustment like methods for more than two views.

The affine reconstruction solution for the case of pure translation was given by Moons *et al.* [Moons-94], and was extended to a combination of pure translation followed by rotation for a full metric reconstruction by Armstrong *et al.* [Armstrong-94]. The case of auto-calibration for a camera rotating about its centre

was given by [Hartley-94a]. The modulus constraint was first published by Pollefeys *et al.* [Pollefeys-96].

The original method for auto-calibration of a stereo rig was given by Zisserman *et al.* [Zisserman-95b], with alternative parametrizations given in Devernay and Faugeras [Devernay-96], and Horaud and Csürka [Horaud-98]. The special case of planar motion of a stereo rig is covered in [Beardsley-95b, Csürka-98]. For planar motion of a monocular camera the original method was published by Armstrong *et al.* [Armstrong-96b], and an alternative numerical solution was given in Faugeras *et al.* [Faugeras-98].

In more recent papers less restrictive constraints than constant internal parameters have been investigated. A number of “existence proofs” have been given: Heyden and Åström [Heyden-97b] showed that metric reconstruction is possible knowing only skew and aspect ratio, and [Pollefeys-98, Heyden-98] showed that skew-zero alone was sufficient.

Triggs [Triggs-97] introduced the absolute (dual) quadric as a numerical device for formulating auto-calibration problems, and applied both linear methods and sequential-quadratic programming to solve for Q_{∞}^* . Pollefeys *et al.* [Pollefeys-98] showed that computations based on Q_{∞}^* could be used to compute metric reconstructions for varying focal length under general motion for real image sequences.

For the case of a rotating camera de Agapito *et al.* [DeAgapito-98] gave a non-linear solution for varying internal parameters based on the use of the DIAC. This was modified in [DeAgapito-99] to an IAC-based linear method.

19.11.2 Notes and exercises

- (i) [Hartley-92a] first gave a solution for the extraction of focal lengths from the fundamental matrix, but the algorithm given there is unwieldy. A simple elegant formula is given in [Bougnoux-98]:

$$\alpha^2 = -\frac{\mathbf{p}'^T [\mathbf{e}']_{\times} \tilde{\mathbf{F}} \mathbf{p} \mathbf{p}^T \mathbf{F}^T \mathbf{p}'}{\mathbf{p}'^T [\mathbf{e}']_{\times} \tilde{\mathbf{F}} \mathbf{F}^T \mathbf{p}'} \quad (19.35)$$

where $\tilde{\mathbf{F}}$ is the matrix $\text{diag}(1, 1, 0)$ and $\mathbf{p}$ and $\mathbf{p}'$ are the principal points in the two images. Unit aspect ratio and zero skew are assumed. The formula for α'^2 is given by reversing the roles of the two images (and transposing $\mathbf{F}$).

Note that the final step of the algorithm is to take a square root. It is assured that α^2 and α'^2 as computed by (19.35) are positive, *given good data and a good guess of the principal point*. However in practice this does not always pertain, and negative values can result. This is the same problem as mentioned previously in section 19.3.5(p468). In addition, as flagged in [Newsam-96], the method has an intrinsic degeneracy when the principal rays of the two cameras meet in space, in which case it is impossible to compute the focal lengths independently. This occurs when both cameras are trained on the same point, quite a common occurrence.

A further degeneracy occurs when the plane defined by the baseline and the

principal ray of one camera is perpendicular to the plane defined by the baseline and principal ray of the other camera. Generally speaking, our opinion is that this method is of doubtful value as a means of computing focal lengths.

- (ii) Show that if the internal parameters are constant then the constraints on $\mathcal{Q}_\infty^*$ obtained from two views (19.6–p462) are equivalent to the Kruppa equations (19.18–p471). Hint, from (9.10–p256) the cameras may be chosen as $P^1 = [I \mid \mathbf{o}]$, $P^2 = [[\mathbf{e}'] \times F \mid \mathbf{e}']$.
- (iii) Show from (19.21–p473) that in the case of a camera translating with constant internal parameters and without rotating, then the plane at infinity may be computed directly from a projective reconstruction.
- (iv) The infinite homography relation (19.24–p476) may be derived in two lines simply from the definition $H_\infty^{ij} = K^i R^{ij} (K^j)^{-1}$, in section 13.4(p338). This may be rearranged as $H_\infty^{ij} K^j = K^i R^{ij}$. Eliminating the rotation using orthogonality as $R^{ij} R^{ijT} = I$ gives $H_\infty^{ij} (K^j K^{jT}) H_\infty^{ijT} = (K^i K^{iT})$.
- (v) Under H_∞ , points on π_∞ (i.e. with $x_4 = 0$) are mapped to points on π_∞ by the 3×3 homography $\mathbf{x}_\infty \mapsto R\mathbf{x}_\infty$. Under this point transformation a conic on π_∞ maps according to result 2.13(p37) $C \mapsto R^{-T} C R^{-1} = R C R^T$. The absolute conic Ω_∞ is fixed since $R I R^T = I$. Now, denote by $\mathbf{a}$ the direction of the rotation axis, so that $R\mathbf{a} = \mathbf{a}$. The (degenerate) point conic $\mathbf{a}\mathbf{a}^T$ is also fixed. It follows that there is a pencil of fixed conics $\mathcal{C}_\infty(\mu) = I + \mu\mathbf{a}\mathbf{a}^T$ under the mapping since

$$\begin{aligned} R(\Omega_\infty + \mu\mathbf{a}\mathbf{a}^T)R^T &= R I R^T + \mu R\mathbf{a}\mathbf{a}^T R^T \\ &= \Omega_\infty + \mu\mathbf{a}\mathbf{a}^T. \end{aligned}$$

The scalar μ parametrizes the pencil. This shows that under a particular similarity there is a one-parameter family of fixed conics on π_∞ . However, it is the case that Ω_∞ is the only conic fixed under any similarity.

- (vi) A further calibration ambiguity exists for a common type of robot head, namely a pan-tilt (or alt-azimuth) head. In [DeAgapito-99] it was shown that since the camera may rotate about its X or Y axes its set of orientations form only a 2-parameter family, rather than a 3-parameter family of general rotations. This limitation causes an ambiguity in the aspect ratio α_x of the camera, and consequently of x_0 as well.
- (vii) The method of auto-calibration from planes is generally non-linear. However, for special motions linear constraints on ω can be obtained. Suppose we have two images of a plane that induces a planar homography H between views, and imagine that the motion of the camera relative to the plane is a general screw motion but with the screw axis parallel to the plane normal.

Consider the action of this screw motion on the plane. Since this action (a rotation about the plane's normal and a translation) does not change the plane's orientation, its line of intersection with the plane at infinity is unchanged (as a set). The absolute conic is fixed (as a set) under Euclidean motion. Consequently, the plane's intersection with the absolute conic, which defines the circular points for the plane, is also unchanged under the motion.

Consider now applying the action to the camera viewing the plane. Since the two circular points are fixed (as 3-space points) they have the same image before and after the motion. As the homography H maps points on the plane between images, the imaged circular points must correspond to two of its fixed points (see section 2.9(p61)) and can thus be determined directly from the homography. Each circular point places a linear constraint on ω . Further details of this method are given in [Knight-03].

Duality

It has been known since the work of Carlsson [Carlsson-95] and Weinshall *et al.* [Weinshall-95] that there is a dualization principle that allows one to interchange the role of points being viewed by several cameras and the camera centres themselves. In principle this implies the possibility of dualizing projective reconstruction algorithms to obtain new algorithms. In this chapter, this theme is developed to outline an explicit method for dualizing any projective reconstruction algorithm. At the practical implementation level, however, it is shown that there are difficulties which need to be overcome in order to allow application of this dualization method to produce working algorithms.

20.1 Carlsson–Weinshall duality

Let $E_1 = (1, 0, 0, 0)^T$, $E_2 = (0, 1, 0, 0)^T$, $E_3 = (0, 0, 1, 0)^T$ and $E_4 = (0, 0, 0, 1)^T$ form part of a projective basis for $\mathbb{P}^3$. Similarly, let $e_1 = (1, 0, 0)^T$, $e_2 = (0, 1, 0)^T$, $e_3 = (0, 0, 1)^T$ and $e_4 = (1, 1, 1)^T$ be a projective basis for the projective image plane $\mathbb{P}^2$.

Now, consider a camera with matrix P . We assume that the camera centre C does not sit on any of the axial planes, that is none of the four coordinates of C is zero. In this case, no three of the points PE_i for $i = 1, \dots, 4$ are collinear in the image. Consequently, one may apply a projective transformation H to the image so that $e_i = HPE_i$. We assume that this has been done, and henceforth denote HP simply by P . Since $PE_i = e_i$, the form of the matrix P is

$$P = \begin{bmatrix} a & & d \\ & b & d \\ & & c & d \end{bmatrix}.$$

Definition 20.1. A camera matrix P is called a *reduced* camera matrix if it maps E_i to e_i for each $i = 1, \dots, 4$. In other words $e_i = PE_i$.

Now, for any point $\mathbf{X} = (X, Y, Z, T)^T$ one verifies that

$$\mathbf{P} = \begin{bmatrix} a & & d \\ & b & \\ & & c & d \\ & & & d \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ T \end{pmatrix} = \begin{pmatrix} aX + dT \\ bY + dT \\ cZ + dT \end{pmatrix}. \quad (20.1)$$

Notice the symmetry in this equation between the entries of the camera matrix and the coordinates of the point. They may be interchanged as follows

$$\begin{bmatrix} a & & d \\ & b & \\ & & c & d \\ & & & d \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ T \end{pmatrix} = \begin{bmatrix} X & & T \\ & Y & \\ & & Z & T \end{bmatrix} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix}. \quad (20.2)$$

The interchange of the roles of cameras and points may be interpreted as a form of duality, which will be referred to as Carlsson–Weinshall duality, or more briefly Carlsson duality. The consequences of this duality will be explored in the rest of this chapter.

20.1.1 Dual algorithms

First of all, we will use it for deriving a dual algorithm from a given projective reconstruction algorithm. Specifically, it will be shown that if one has an algorithm for doing projective reconstruction from n views of $m + 4$ points, then there is an algorithm for doing projective reconstruction from m views of $n + 4$ points. This result, observed by Carlsson [Carlsson-95], will be made specific by explicitly describing the steps of the dual algorithm.

We consider a projective reconstruction problem, which will be referred to as $\mathcal{P}(m, n)$. It is the problem of doing reconstruction from m views of n points. We denote image points by $\mathbf{x}_j^i$, which represents the image of the j -th object space point in the i -th view. Thus, the upper index indicates the view number, and the lower index represents the point number. Such a set of points $\{\mathbf{x}_j^i\}$ is called *realizable* if there are a set of camera matrices $\mathbf{P}^i$ and a set of 3D points $\mathbf{X}_j$ such that $\mathbf{x}_j^i = \mathbf{P}^i \mathbf{X}_j$. The projective reconstruction problem $\mathcal{P}(m, n)$ is that of finding such camera matrices $\mathbf{P}^i$ and points $\mathbf{X}_j$ given a realizable set $\{\mathbf{x}_j^i\}$ for m views of n points. The set of camera matrices and 3D points together are called a *realization* (or *projective realization*) of the set of point correspondences.

Let $\mathcal{A}(n, m + 4)$ represent an algorithm for solving the projective reconstruction problem $\mathcal{P}(n, m + 4)$. An algorithm will now be exhibited for solving the projective reconstruction $\mathcal{P}(m, n + 4)$. This algorithm will be denoted $\mathcal{A}^*(m, n + 4)$, the dual of the algorithm $\mathcal{A}(n, m + 4)$.

Initially, the steps of the algorithm will be given without proof. In addition, difficulties will be glossed over so as to give the general idea without getting bogged down in details. In the description of this algorithm it is important to keep track of the range of the indices, and whether they index the cameras or the points. Thus, the following may help to keep track.




Fig. 20.1. **Left:** Input to algorithm $\mathcal{A}^*(m, n+4)$. **Right:** Input data after transformation.

- Upper indices represent the view number.
- Lower indices represent the point number.
- i ranges from 1 to m .
- j ranges from 1 to n .
- k ranges from 1 to 4.

The dual algorithm

Given an algorithm $\mathcal{A}(n, m+4)$ the goal is to exhibit a dual algorithm $\mathcal{A}^*(m, n+4)$.

Input:

The input to the algorithm $\mathcal{A}^*(m, n+4)$ consists of a realizable set of $n+4$ points seen in m views. This set of points can be arranged in a table as in figure 20.1(left).

In this table, the points x_{n+k}^i are separated from the other points x_j^i , since they will receive special treatment.

Step 1: Transform

The first step is to compute, for each i , a transformation T^i that maps the points x_{n+k}^i , $k = 1, \dots, 4$ in the i -th view to the points e_k of a canonical basis for projective 2-space $\mathbb{P}^2$. The transformation T^i is applied also to each of the points x_j^i to produce transformed points $x_j'^i = T^i x_j^i$. The result is the transformed point array shown in figure 20.1(right). A different transformation T^i is computed and applied to each column of the array, as indicated.




Fig. 20.2. **Left:** Transposed data. **Right:** Transposed data extended by addition of extra points.

Step 2: Transpose

The last four rows of the array are dropped, and the remaining block of the array is transposed. One defines $\hat{x}_i^j = x_j^i$. At the same time, one does a mental switch of points and views. Thus the point $\hat{x}_i^j$ is now conceived as being the image of the i -th point in the j -th view, whereas the point x_j^i was the image of the j -th point in the i -th view. What is happening here effectively is that the roles of points and cameras are being swapped – the basic concept behind Carlsson duality expressed by (20.2). The resulting transposed array is shown in figure 20.2(left).

Step 3: Extend

The array of points is now extended by the addition of four extra rows containing points e_k in all positions of the $(m+k)$ -th row of the array, as shown in figure 20.2(right). The purpose of this extension will be explained in section 20.1.2.

Step 4: Solve

The array of points resulting from the last step has $m+4$ rows and n columns, and may be regarded as the positions of $m+4$ points seen in n views. As such, it is a candidate for solution by the algorithm $\mathcal{A}(n, m+4)$, which we have assumed is given. Essential here is that the points in the array form a realizable set of point correspondences. Justification of this is deferred for now. The result of the algorithm $\mathcal{A}(n, m+4)$ is a set of cameras $\hat{P}^j$ and points $\hat{X}_i$ such that $\hat{x}_i^j = \hat{P}^j \hat{X}_i$. In addition, corresponding to the last four rows of the array, there are points $\hat{X}_{m+k}$ such that $e_k = \hat{P}^j \hat{X}_{m+k}$ for all j .

Step 5: 3D transformation

Since the reconstruction obtained in the last step is a projective reconstruction, one may transform it (equivalently, choose a projective coordinate frame) such that the points

$\widehat{\mathbf{X}}_{m+k}$ are the four points $\mathbf{E}_k$ of a partial canonical basis for $\mathbb{P}^3$. The only requirement is that the points $\widehat{\mathbf{X}}_{m+k}$ obtained in the projective reconstruction are not coplanar. This assumption is validated later.

At this point, one sees that $\mathbf{e}_k = \tilde{\mathbf{P}}^j \widehat{\mathbf{X}}_{m+k} = \tilde{\mathbf{P}}^j \mathbf{E}_k$. From this it follows that $\tilde{\mathbf{P}}^j$ has the special form

$$\tilde{\mathbf{P}}^j = \begin{bmatrix} a^j & & d^j \\ & b^j & d^j \\ & & c^j & d^j \end{bmatrix} . \quad (20.3)$$

Step 6: Dualize

Let $\widehat{\mathbf{X}}_i = (X_i, Y_i, Z_i, T_i)^T$, and $\tilde{\mathbf{P}}^j$ be as given in (20.3). Now define points $\mathbf{X}_j = (a^j, b^j, c^j, d^j)^T$ and cameras

$$\mathbf{P}^i = \begin{bmatrix} X_i & & T_i \\ & Y_i & T_i \\ & & Z_i & T_i \end{bmatrix} .$$

Then one verifies that

$$\begin{aligned} \mathbf{P}^i \mathbf{X}_j &= (X_i a^j + T_i d^j, Y_i b^j + T_i d^j, Z_i c^j + T_i d^j)^T \\ &= \tilde{\mathbf{P}}^j \widehat{\mathbf{X}}_i \\ &= \widehat{\mathbf{X}}_i^j \\ &= \mathbf{X}_j^i . \end{aligned}$$

If, in addition, one defines $\mathbf{X}_{n+k} = \mathbf{E}_k$ for $k = 1, \dots, 4$, then $\mathbf{P}^i \mathbf{X}_{n+k} = \mathbf{e}_k$. It is then evident that the cameras $\mathbf{P}^i$ and points $\mathbf{X}_j$ and $\mathbf{X}_{n+k}$ form a projective realization of the transformed data array obtained in step 1 of this algorithm.

Step 7: Reverse transformation

Finally, defining $\mathbf{P}^i = (\mathbf{T}^i)^{-1} \mathbf{P}'^i$, and with the points $\mathbf{X}_j$ and $\mathbf{X}_{n+k}$ obtained in the previous step, one has a projective realization of the original data. Indeed, one verifies

$$\mathbf{P}^i \mathbf{X}_j = (\mathbf{T}^i)^{-1} \mathbf{P}'^i \mathbf{X}_j = (\mathbf{T}^i)^{-1} \mathbf{X}_j^i = \mathbf{X}_j^i .$$

This completes the description of the algorithm. One can see that it takes place in various stages.

- (i) In step 1, the data is transformed into canonical image reference frames based on the selection of four distinguished points.
- (ii) In steps 2 and 3 the problem is mapped into the dual domain, resulting in a dual problem $\mathcal{P}(n, m+4)$.
- (iii) The dual problem is solved in steps 4 and 5.
- (iv) Step 6 maps the solution back into the original domain.
- (v) Step 7 undoes the effects of the initial transformation.

20.1.2 Justification of the algorithm

To justify this algorithm, one needs to be sure that at step 4 there indeed exists a solution to the transformed problem. Before considering this, it is necessary to explain the purpose of step 3, which extends the data by the addition of rows of image points $\mathbf{e}_k$, and step 5, which transforms the arbitrary projective solution to one in which four points are equal to the 3D basis points $\mathbf{E}_k$.

The purpose of these steps is to ensure that one obtains a solution to the dual reconstruction problem in which $\tilde{\mathbf{P}}^j$ has the special form given by (20.3) in which the camera matrix is parametrized by only 4 values. The dual algorithm is described in this manner so that it will work with any algorithm $\mathcal{A}(n, m+4)$ whatever. However, both steps 3 and 5 may be eliminated if the known algorithm $\mathcal{A}(n, m+4)$ has the capability of enforcing this constraint on the camera matrices directly. Algorithms based on the fundamental matrix, trifocal or quadrifocal tensors may easily be modified in this way, as will be seen.

In the mean time, since $\tilde{\mathbf{P}}^j$ of the form (20.3) is called a *reduced camera matrix*, we call any reconstruction of a set of image correspondences in which each camera matrix is of this form a *reduced reconstruction*. Not all sets of realizable point correspondences allow a reduced realization, however. The following result characterizes sets of point correspondences that do have this property.

Result 20.2. *A set of image points $\{\mathbf{x}_j^i : i = 1, \dots, m; j = 1, \dots, n\}$ admits a reduced realization if and only if it may be augmented with supplementary correspondences $\mathbf{x}_{n+k}^i = \mathbf{e}_k$ for $k = 1, \dots, 4$ such that*

- (i) *The total set of image correspondences is realizable, and*
- (ii) *The reconstructed points $\mathbf{X}_{n+k}$ corresponding to the supplementary image correspondences are non-coplanar.*

Proof. The proof is straightforward enough. Suppose the set permits a reduced realization, and let $\mathbf{P}^i$ be the set of reduced camera matrices. Let points $\mathbf{X}_{n+k} = \mathbf{E}_k$ for $k = 1, \dots, 4$ be projected into the m images. The projections are $\mathbf{x}_{n+k}^i = \mathbf{P}^i \mathbf{X}_{n+k} = \mathbf{P}^i \mathbf{E}_k = \mathbf{e}_k$ for all i .

Conversely, suppose the augmented set of points is realizable and the points $\mathbf{X}_{n+k}$ are non-coplanar. In this case, a projective basis may be chosen such that $\mathbf{X}_{n+k} = \mathbf{E}_k$. Then for each view, one has $\mathbf{e}_k = \mathbf{P}^i \mathbf{E}_k$ for all k . From this it follows that each $\mathbf{P}^i$ has the desired form (20.3). $\square$

One other remark must be made before proving the correctness of the algorithm.

Result 20.3. *If a set of image points $\{\mathbf{x}_j^i : i = 1, \dots, m; j = 1, \dots, n\}$ permits a reduced realization then so does the transposed set $\{\hat{\mathbf{x}}_i^j : j = 1, \dots, n; i = 1, \dots, m\}$ where $\hat{\mathbf{x}}_i^j = \mathbf{x}_j^i$ for all i and j .*

This is the basic duality property, effectively proved by the construction given in step 6 of the algorithm above. Now it is possible to prove the correctness of the algorithm.

Result 20.4. Let $\mathbf{x}_j^i$ and $\mathbf{x}_{n+k}^i$ as in figure 20.1(left) be a set of realizable image point correspondences, and suppose

- (i) for each i , the four points $\mathbf{x}_{n+k}^i$ do not include three collinear points.
- (ii) the four points $\mathbf{x}_{n+k}$ in a projective reconstruction are non-coplanar.

Then the algorithm of section 20.1.1 will succeed.

Proof. Because of the first condition, transformations T^i exist for each i , transforming the input data to the form shown in figure 20.1(right). This transformed data is also realizable, since the transformed data differs only by a projective transformation of the image.

Now, according to result 20.2 applied to figure 20.1(right), the correspondences $\mathbf{x}_j^i$ admit a reduced realization. By result 20.3 the transposed data figure 20.2(left) also admits a reduced realization. Applying result 20.2 once more shows that the extended data figure 20.2(right) is realizable. Furthermore, the points $\hat{\mathbf{X}}_{m+k}$ are non-coplanar, and so step 5 is valid. The subsequent steps 6 and 7 go forward without problems. $\square$

The first condition may be checked from the image correspondences $\mathbf{x}_j^i$. It may be thought that to check the second condition requires reconstruction to be carried out. It is, however, possible to check whether the reconstructed points will be coplanar without carrying out the reconstruction. This is left as an exercise for the reader (page 342).

20.2 Reduced reconstruction

In this section, we concentrate on and reevaluate steps 3 – 5 of the algorithm described in the preceding section. To recapitulate, the purpose of these steps is to obtain a reduced reconstruction from a set of image correspondences. Thus, the input is a set of image correspondences $\hat{\mathbf{x}}_i^j$ admitting a reduced realization (see figure 20.2(left)). The output is a set of *reduced* camera matrices $\hat{\mathbf{P}}^j$ and points $\hat{\mathbf{X}}_i$ such that $\hat{\mathbf{P}}^j \hat{\mathbf{X}}_i = \hat{\mathbf{x}}_i^j$ for all i, j .

As we have seen, one way to do this (as in steps 3 – 5) of the given algorithm is to augment the points by the addition of four extra synthetic point correspondences $\hat{\mathbf{x}}_{m+k}^j$, carrying out projective reconstruction, and then applying a 3D homography so that the 3D points $\hat{\mathbf{X}}_{m+k}$ are mapped to the points $\mathbf{E}_k$ of a projective basis for $\mathbb{P}^3$. The problem with this is that in the presence of noise, the projective reconstruction is not exact. Thus, the camera matrices obtained by this method will map points $\mathbf{E}_k$ to points close to, but not identical with $\mathbf{e}_k$. This means that the camera matrices are not exactly in reduced form. Therefore, we now consider methods of computing a realization of the point correspondences in which the cameras are exactly reduced.

20.2.1 The reduced fundamental matrix

The most evident applications of these dual methods are to dualize the reconstruction algorithms involving the fundamental matrix and trifocal tensor. These will lead to reconstruction algorithms for 6 or 7 points (respectively) across N views. In this section, we consider reconstruction from 6 points. The dual of a reconstruction problem

$\mathcal{P}(N, 6)$ is a problem $\mathcal{P}(2, N + 4)$, namely reconstruction from $N + 4$ points in 2 views. The method of chapter 10 involving the fundamental matrix is a standard way of solving such a problem.

To this end we define a reduced fundamental matrix:

Definition 20.5. A fundamental matrix $\tilde{F}$ is called a *reduced* fundamental matrix if it satisfies the condition $\mathbf{e}_i^T \tilde{F} \mathbf{e}_i = 0$ for $i = 1, \dots, 4$.

It is evident that since a reduced fundamental matrix already satisfies constraints derived from four point correspondences, it may be computed from a small number of additional points, in fact linearly from four points, or non-linearly from three points.

20.2.2 Computation of the reduced fundamental matrix

For a reduced fundamental matrix, the condition $\mathbf{e}_i^T \tilde{F} \mathbf{e}_i = 0$ for $i = 1, \dots, 3$ implies that the diagonal entries of $\tilde{F}$ are zero. The requirement that $(1, 1, 1)^T \tilde{F} (1, 1, 1)^T = 0$ gives the additional condition that the sum of entries of $\tilde{F}$ is zero. Thus one may write $\tilde{F}$ in the form

$$\tilde{F} = \begin{bmatrix} 0 & p & q \\ r & 0 & s \\ t & -(p+q+r+s+t) & 0 \end{bmatrix} \quad (20.4)$$

thereby parametrizing a fundamental matrix satisfying all the linear constraints (though not the condition $\det \tilde{F} = 0$). Now, a further point correspondence $\mathbf{x} \leftrightarrow \mathbf{x}'$ satisfying $\mathbf{x}'^T \tilde{F} \mathbf{x} = 0$ is easily seen to provide a linear equation in the parameters $p, \dots, t$ of $\tilde{F}$. Given at least four such correspondences, one may solve for these parameters, up to an inconsequential scale. Given only three such correspondences, the extra constraint $\det \tilde{F} = 0$ may be used to provide the extra constraint necessary to determine $\tilde{F}$. There may be one or three solutions. This computation is analogous to the method used in section 11.1.2 (p281) to compute the fundamental matrix from seven point correspondences. With four correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ or more, one finds a least-squares solution.

20.2.3 Retrieving reduced camera matrices

Computing a pair of reduced camera matrices that correspond to a reduced fundamental matrix is surprisingly tricky. One can not assume that the first camera is $[I \mid 0]$ as in the usual projective camera case, since this is non-generic, the camera centre corresponding with the basis point $\mathbf{E}_4 = (0, 0, 0, 1)^T$. However, we may instead assume that the pair of cameras are of the form

$$P = \begin{bmatrix} 1 & & 1 \\ & 1 & \\ & & 1 & 1 \end{bmatrix} \quad \text{and} \quad P' = \begin{bmatrix} a & & d \\ & b & \\ & & c & d \end{bmatrix} \quad (20.5)$$

since the centre of the first camera is $(1, 1, 1, 1)^T$ which is a generic point with respect to the basis points $\mathbf{E}_1, \dots, \mathbf{E}_4$. Further, if $d \neq 0$, then we may assume that $d = 1$, but we prefer not to do this.

The reduced fundamental matrix corresponding to this pair of cameras is

$$\tilde{F} = \begin{bmatrix} 0 & b(d-c) & -c(d-b) \\ -a(d-c) & 0 & c(d-a) \\ a(d-b) & -b(d-a) & 0 \end{bmatrix} \quad (20.6)$$

which the reader may verify satisfies the four linear constraints as well as the zero-determinant condition. The task at hand is to retrieve the values of (a, b, c, d) given the fundamental matrix. This seemingly requires a solution of simultaneous quadratic equations, but there is a linear method as follows.

- (i) The ratio $a : b : c$ may be found by solving the set of homogeneous linear equations

$$\begin{bmatrix} f_{12} & f_{21} & 0 \\ f_{13} & 0 & f_{31} \\ 0 & f_{23} & f_{32} \end{bmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \mathbf{0} \quad (20.7)$$

where f_{ij} is the ij -th entry of $\tilde{F}$. The matrix appearing here clearly has the same rank as $\tilde{F}$ (namely 2), so there is a unique solution (up to scale) to this equation set. The solution $a : b : c = A : B : C$ provides a set of homogeneous equations in a, b and c , namely $Ba = Ab$, $Ca = Ac$ and $Cb = Bc$, of which two are linearly independent.

- (ii) Similarly, one may find the ratio $d - a : d - b : d - c$ by solving $(d - a, d - b, d - c)\tilde{F} = \mathbf{0}$. Once again, the solution is unique. This provides two more linear equations in the values of a, b, c, d .
- (iii) From the set of equations one may solve for (a, b, c, d) up to scale, and hence reconstruct the second camera matrix according to (20.5).

20.2.4 Solving for six points in three views

The minimal case involving the computation of the reduced fundamental matrix involves three point, in which case there may be three solutions. By dualization, this leads to a solution to the reconstruction problem for six points in three views. Its use in outlier detection using the trifocal tensor has been described in algorithm 16.4(p401). Because of its intrinsic interest as a minimum-case solution, and because of its practical use, the algorithm is given explicitly as algorithm 20.1. The algorithm consists essentially of putting together what has been described previously. However there are a few minor variations.

In algorithm 20.1, the final estimation of the camera matrices takes place in the domain of the original point measurements. As an alternative, one could apply a DLT algorithm in the dual domain, as in the basic algorithm of section 20.1.1. However, the present method seems simpler. It has the advantage of avoiding the need for applying the inverse transformations T, T', T'' . More significantly the final computation of the camera matrices is carried out using the original data. This is important, because the transformations can severely skew the noise characteristics of the data.

The basis for this algorithm and the n -view case to follow is the dual fundamental

Objective Given a set of six point correspondences $\mathbf{x}_j \leftrightarrow \mathbf{x}'_j \leftrightarrow \mathbf{x}''_j$ across three views, compute a reconstruction from these points, consisting of the three camera matrices $\mathbf{P}$, $\mathbf{P}'$ and $\mathbf{P}''$ and the six 3D points $\mathbf{X}_1, \dots, \mathbf{X}_6$.

Algorithm

- (i) Select a set of four points no three of which are collinear in any of the three views. Let these be the correspondences $\mathbf{x}_{2+k} \leftrightarrow \mathbf{x}'_{2+k} \leftrightarrow \mathbf{x}''_{2+k}$ for $k = 1, \dots, 4$.
- (ii) For the first view, find a projective transformation $\mathbf{T}$ that maps each $\mathbf{x}_{2+k}$ to $\mathbf{e}_k$. Apply $\mathbf{T}$ to the other two points $\mathbf{x}_1$ and $\mathbf{x}_2$ resulting in points $\hat{\mathbf{x}}_1 = \mathbf{T}\mathbf{x}_1$ and $\hat{\mathbf{x}}_2 = \mathbf{T}\mathbf{x}_2$.
- (iii) From the *dual correspondence* $\hat{\mathbf{x}}_1 \leftrightarrow \hat{\mathbf{x}}_2$ derive an equation in the entries p, q, r, s, t of the reduced fundamental matrix $\hat{\mathbf{F}}$ as in (20.4). The equation is derived from the relationship $\hat{\mathbf{x}}_2^T \hat{\mathbf{F}} \hat{\mathbf{x}}_1 = 0$.
- (iv) In the same way as in the last two steps, form two more equations from the points in the other two views. This results in a set of three homogeneous equations in five unknowns. There will be a two-parameter family of solutions for the reduced fundamental matrix, generated by two independent solutions $\hat{\mathbf{F}}_1$ and $\hat{\mathbf{F}}_2$.
- (v) The general solution is $\hat{\mathbf{F}} = \lambda \hat{\mathbf{F}}_1 + \mu \hat{\mathbf{F}}_2$. From the requirement $\det \hat{\mathbf{F}} = 0$ one derives a homogeneous cubic equation in (λ, μ) , which one may solve to find (λ, μ) , and hence $\hat{\mathbf{F}}$. There will be 1 or 3 real solutions. Further steps are applied to each solution in turn.
- (vi) Use the method of section 20.2.3 to extract the parameters $(a, b, c, d)^T$ of the second reduced camera matrix $\hat{\mathbf{P}}'$ defined in (20.5).
- (vii) We complete the reconstruction in the original measurement domain. The dual of camera $\hat{\mathbf{P}}'$ defined by (a, b, c, d) is the point $\mathbf{X}_2 = (a, b, c, d)^T$. Thus the six 3D points are $\mathbf{X}_1 = (1, 1, 1, 1)^T$, $\mathbf{X}_2 = (a, b, c, d)^T$ and $\mathbf{X}_{2+k} = \mathbf{E}_k$ for $k = 1, \dots, 4$. This gives the structure of the reconstructed scene. One may then compute the three camera matrices using the DLT algorithm for camera calibration described in section 7.1(p178). Since we require the camera matrices defined with respect to the original camera coordinates here, we use the original coordinates to solve for $\mathbf{P}$, $\mathbf{P}'$ and $\mathbf{P}''$ such that $\mathbf{P}\mathbf{X}_j = \mathbf{x}_j$, etc. Since the solution will be exact, the DLT solution will be sufficient.

Algorithm 20.1. Algorithm for computing a projective reconstruction from six points in three views.

matrix, $\hat{\mathbf{F}}$. Note how the dual fundamental matrix expresses a relationship between points in the same image. Indeed the equations used to solve for $\hat{\mathbf{F}}$ are constructed from points in the same image. This contrasts with the standard fundamental matrix where the relationships being encoded are between points seen in separate images.

20.2.5 Six points in n views

The method for six points in three views can be applied with little modification to the case of six points in many views. The main difference is that the reduced fundamental matrix $\hat{\mathbf{F}}$ will be uniquely determined by the data. Specifically at step 4 of the algorithm of section 20.2.4, each view will contribute one equation. With four views or more, this will be sufficient to determine $\hat{\mathbf{F}}$.

In this redundant-data case, one must be careful with the effects of noise. For this reason, it appears preferable to carry out the last step of the algorithm, as shown, with original untransformed points. This mitigates the effect of noise distortion that would result from working with transformed points.

20.2.6 Seven points in n views

The problem of seven points in n views is dual to the case of three views of $n+4$ points and is solved by computing the reduced trifocal tensor from n point correspondences.

Definition 20.6. A trifocal tensor $\mathcal{T}$ is called a *reduced trifocal tensor* if it satisfies the linear constraints imposed by synthetic point correspondences $\mathbf{e}_k \leftrightarrow \mathbf{e}'_k \leftrightarrow \mathbf{e}''_k$ for $i = 1, \dots, A$.

The general method for reconstruction from 7 points is similar to the method for six points in n views, except that the reduced trifocal tensor is used instead of the reduced fundamental matrix. There are, however some minor differences.

In computing the reduced trifocal tensor, the constraints corresponding to the synthetic correspondences $\mathbf{e}_j \leftrightarrow \mathbf{e}_j \leftrightarrow \mathbf{e}_j$ should be satisfied exactly, whereas the other correspondences used to compute the tensor are subject to noise, and will only be satisfied approximately. Otherwise the computed tensor will not be exactly in reduced form. In the case of the reduced fundamental matrix this was handled by giving a specific parametrization of the reduced fundamental matrix. That is, it was parametrized in such a way that the constraints generated by the synthetic correspondences were automatically satisfied (see (20.4)). In the case of the trifocal tensor, it is not obvious that such a convenient parametrization is possible. The synthetic constraints are of the form

$$e^i e^p e^q \epsilon_{jpr} \epsilon_{kqs} \mathcal{T}_i^{jrk} = 0_{rs}$$

which is rather more complicated than the linear constraints on the reduced fundamental matrix. Instead, one may proceed in the following manner.

In the usual linear method of computing the trifocal tensor, one must solve a set of linear equations of the form $\mathbf{A}\mathbf{t} = \mathbf{0}$, or more precisely, find the vector $\mathbf{t}$ that minimizes $\|\mathbf{A}\mathbf{t}\|$ subject to $\|\mathbf{t}\| = 1$. In solving for the reduced trifocal tensor, the matrix $\mathbf{A}$ may be divided into two parts, corresponding to those constraints from the synthetic correspondences, which should be satisfied exactly, and the constraints from real correspondences, which must be satisfied in a least-squares sense. The first set of constraints are of the form $\mathbf{C}\mathbf{t} = \mathbf{0}$, and the second set may be written as $\hat{\mathbf{A}}\mathbf{t} = \mathbf{0}$. The problem becomes: find $\mathbf{t}$ that minimizes $\|\hat{\mathbf{A}}\mathbf{t}\|$ subject to $\|\mathbf{t}\| = 1$ and $\mathbf{C}\mathbf{t} = \mathbf{0}$. An algorithm for solving this problem is given as algorithm A5.5(p594).

For the problem of extracting the three camera matrices from the reduced tensor no simple method seems to be available, similar to that described in section 20.2.3 for the fundamental matrix. Instead, one may use the method described in steps 5 and 6 of the general dual algorithm (see page 505).

The minimal configuration of this type is seven points in two views. In this case, the problem is best solved directly by the method of section 11.1.2(p281), rather than by dualization to the case of three points in six views.

20.2.7 Performance issues

Dual reconstruction algorithms based on the reduced fundamental matrix and trifocal tensor have been implemented and tested. The results of these tests were contained

in a student report presented in August 1996 by Gilles Debunne. Since this report is effectively unavailable, the results are summarized here.

The most serious difficulty is the distortion of the noise distribution by the application of projective transformations T^i to the images. Application of projective transformations to the image data has the effect of distorting any noise distribution that may apply to the data. The problem is related to the need to choose four points that are non-collinear in any of the images. If the points are close to collinear in any of the images, then the projective transformation applied to the image in step 1 of the algorithm may entail extreme distortion of the image. This sort of distortion can degrade performance of the algorithm severely.

Without special attention being paid to the noise-distortion, performance of the algorithms was generally unsatisfactory. Despite great care being taken to minimize errors due to noise in steps 4–6 of the algorithm (page 505ff), when the inverse projective transformations are applied in step 7, the average error became very large. Some points retained quite small error, whereas in those images where distortion was significant, quite large errors resulted.

Normalization in the sense of section 4.4.4(p107) is also a problem. It has been shown to be essential for performance of the linear reconstruction algorithms to apply data normalization. However, what sort of normalization should be applied to the transformed data of figure 20.1(right), which is geometrically unrelated to actual image measurements, is a mystery.

To get good results, it would seem that one would need to propagate assumed error distributions forward in step 1 of the algorithm to get assumed error distributions for the transformed data figure 20.1(right), and then during reconstruction to minimize residual error relative to this propagated error distribution. Equivalently, cost functions being minimized during reconstruction must be related back to measurement error in the original image points. Recent work reported in [Hartley-00a, Schaffalitzky-00c] has shown that this indeed gives significantly improved results.

20.3 Closure

20.3.1 The literature

The idea behind Carlsson–Weinshall duality was first described in a pair of papers appearing simultaneously: [Carlsson-95] and Weinshall [Weinshall-95] and subsequently in a joint paper [Carlsson-98]. The treatment given here for the general method of dualizing an algorithm was given in [Hartley-98b], derived from these earlier papers. Details of methods for handling with noise propagation were given in [Hartley-00a], building on earlier implementations contained in an unavailable report by Gilles Debunne (August 1996).

The problem of reconstruction from six points in three views was perhaps first treated in a technical report [Hartley-92b] (later published as [Hartley-94c]) where the existence of up to eight solutions was shown. The problem was given a complete solution in [Quan-94] where it was shown that only three solutions are possible. This was also pointed out in [Ponce-94]. The paper [Carlsson-95] established that this problem is

dual to the two-view, seven point problem for which a solution was known. This enabled the formulation of the method given in this chapter. The minimum six point configuration was used in [Torr-97] for robust estimation of the trifocal tensor. An alternative method for computing a reconstruction from six points in $n \geq 3$ views is given in [Schaffalitzky-00c]. This method does not require that images are first projectively transformed to a canonical basis.

20.3.2 Notes and Exercises

- (i) In the dual algorithm of section 20.1.1 it was noted that the method will only work if the four points used to define the image transformations are non-coplanar. However, note that in this case, the algorithm of section 18.5.2(p450) will compute a projective reconstruction linearly.
- (ii) If the four chosen points are coplanar, then the homographies T^i will map the plane to a common coordinate system. The transformed points $\mathbf{x}_j^i$ will then satisfy the condition of section 17.5.2(p425), namely the lines joining any pair of points $\mathbf{x}_j^i$ and $\mathbf{x}_k^i$ (j and k fixed and a different line for each i) will meet in a common point. The duality in the case is described in [Criminisi-98, Irani-98].
- (iii) Still in the case of the four chosen points being coplanar: after applying the T^i , any point on the plane of the four points will map to the same point in all images. Thus, the fundamental matrix consistent with the point correspondences will be skew-symmetric.

Cheirality

When a projective reconstruction of a scene is carried out from a set of point correspondences, an important piece of information is typically ignored – if the points are visible in the images, then they must have been in front of the camera. In general, a projective reconstruction of a scene will not bear a close resemblance to the real scene when interpreted as if the coordinate frame were Euclidean. The scene is often split across the plane at infinity, as is illustrated in two dimensions by figure 21.1. It is possible to come much closer to at least an affine reconstruction of the scene by taking this simple constraint into account. The resulting reconstruction is called “quasi-affine” in that it lies part way between a projective and affine reconstruction. Scene objects are no longer split across the plane at infinity, though they may still suffer projective distortion.

Converting a projective reconstruction to quasi-affine is extremely simple if one neglects the cameras and requires only that the scene be of the correct quasi-affine form – in fact it can be accomplished in about two lines of programming (see corollary 21.9). To handle the cameras as well requires the solution of a linear programming problem.

21.1 Quasi-affine transformations

A subset B of $\mathbb{R}^n$ is called convex if the line segment joining any two points in B also lies entirely within B . The convex hull of B , denoted $\overline{B}$, is the smallest convex set containing B . Our main concern will be with 3-dimensional point sets, so $n = 3$. We view $\mathbb{R}^3$ as being a subset of $\mathbb{P}^3$, consisting of all non-infinite points. The infinite points constitute the plane at infinity, denoted π_∞ . Thus, $\mathbb{P}^3 = \mathbb{R}^3 \cup \pi_\infty$. A subset of $\mathbb{P}^3$ will be called convex if and only if it is contained in $\mathbb{R}^3$ and is convex in $\mathbb{R}^3$. Hence, according to this definition, a convex set does not contain any infinite points.

Definition 21.1. Consider a point set $\{X_i\} \subset \mathbb{R}^3 \subset \mathbb{P}^3$. A projective mapping $h : \mathbb{P}^3 \rightarrow \mathbb{P}^3$ is said to preserve the convex hull of the points $\{X_i\}$ if

- (i) $h(X_i)$ is a finite point for all i , and
- (ii) h maps the convex hull of points $\{X_i\}$ bijectively onto the convex hull of the points $\{h(X_i)\}$.




Fig. 21.1. (a) an image of a comb, and (b) the result of applying a projective transformation to the image. The projective transformation does not however preserve the convex hull of the set of points constituting the comb. In the original image, the convex hull of the comb is a finite set contained within the extent of the visible image. However, some of the points in this convex hull are mapped to infinity by the transformation.

An example, shown in figure 21.1, may help in understanding this definition. The example deals with 2D point sets, but the principle is the same. The figure shows an image of a comb and the image resampled according to a projective mapping. The projective mapping does *not* however preserve the convex hull of the comb. Most people will agree that the resampled image is unlike any view of a comb seen by camera or human eye.

The property of preserving the convex hull of a set of points may be characterized in various different ways, as is shown by the theorem to be given shortly. In order to state this theorem, we introduce a new notation.

Notation. The symbol $\hat{\mathbf{X}}$ denotes a homogeneous representation of a point $\mathbf{X}$ in which the last coordinate is equal to 1.

Commonly in this chapter we will be interested in the exact equalities (not equalities up to scale) between vectors representing homogeneous quantities, such as points in 3D. Thus for instance if $\mathbf{H}$ is a projective transformation, we may write $\mathbf{H}\hat{\mathbf{X}} = \mathbf{T}'\hat{\mathbf{X}}'$ to mean that the transformation takes point $\hat{\mathbf{X}}$ to point $\hat{\mathbf{X}}'$, but that the scale factor $\mathbf{T}'$ is required to make this equality exact.

Now for the theorem.

Theorem 21.2. Consider a projective transformation $h : \mathbb{P}^3 \rightarrow \mathbb{P}^3$ and a set of points $\{\mathbf{X}_i\}$. Let π_∞ be the plane mapped to infinity by h . The following statements are equivalent.

- (i) h preserves the convex hull of the points $\{\mathbf{X}_i\}$.
- (ii) $\tilde{\mathbf{X}} \cap \pi_\infty = \emptyset$ for any point $\tilde{\mathbf{X}}$ in the convex hull of the points $\{\mathbf{X}_i\}$.
- (iii) Let $\mathbf{H}$ be a matrix representing the transformation h , and suppose that $\mathbf{H}\hat{\mathbf{X}}_i = \mathbf{T}'_i\hat{\mathbf{X}}'_i$. Then the constants $\mathbf{T}'_i$ all have the same sign.

Proof. This will be proved by showing that $(i) \Rightarrow (ii) \Rightarrow (iii) \Rightarrow (i)$.

$(i) \Rightarrow (ii)$. If h preserves the convex hull of the points, then $h(\bar{X})$ is a finite point for any point $\bar{X}$ in the convex hull of the X_i . So $\bar{X} \cap \pi_\infty = \emptyset$.

$(ii) \Rightarrow (iii)$. Consider the chord joining two points X_i and X_j and suppose that T'_i and T'_j (as in part (iii) of the theorem) have opposite signs. Since T'_i is a continuous function of the coordinates of X_i , there must be a point $\bar{X}$ lying on the chord from X_i to X_j for which T' is equal to zero. This means that $H\bar{X} = (\bar{X}', \bar{Y}', \bar{Z}', 0)^T$. Since $\bar{X}$ is in the convex hull of the points $\{X_i\}$, this contradicts (ii) .

$(iii) \Rightarrow (i)$. We assume that there exist constants T'_i all of the same sign such that $H\bar{X}_i = T'_i\bar{X}'_i$. Let S be the subset of $\mathbb{R}^n$ consisting of all points X satisfying the condition $H\bar{X} = T'X'$ such that T' has the same sign as the T'_i . The set S contains $\{X_i\}$. It will be shown that S is convex. If X_i and X_j are points in S with corresponding constants T'_i and T'_j , then any intermediate point X between X_i and X_j must have T' value intermediate between T'_i and T'_j . To see this, consider a point $\bar{X} = \alpha\bar{X}_i + (1-\alpha)\bar{X}_j$ where $0 \leq \alpha \leq 1$. This point lies between X_i and X_j . Denote by h_4^T the last row of H . Then,

$$\begin{aligned} T' &= h_4^T \bar{X} \\ &= h_4^T (\alpha \bar{X}_i + (1-\alpha) \bar{X}_j) \\ &= \alpha h_4^T \bar{X}_i + (1-\alpha) h_4^T \bar{X}_j \\ &= \alpha T'_i + (1-\alpha) T'_j \end{aligned}$$

which lies between T'_i and T'_j as claimed. Consequently, the value of T' must have the same sign as T'_i and T'_j , and so X lies in S also. This shows that S is convex.

Now, let $\tilde{S}$ be a convex subset of S . It will be shown that $h(\tilde{S})$ is also convex. This is easily seen to be true, since h maps a line segment in $\tilde{S}$ to a line segment that does not cross the plane at infinity. Thus, h maps any convex set $\tilde{S}$ such that $S \supset \tilde{S} \supset \{X_i\}$ to a convex set $\tilde{S}'$ such that $S' \supset \tilde{S}' \supset \{X'_i\}$. However, if H satisfies condition (iii) , then it is easily seen that H^{-1} does also. From this it follows that the above correspondence between convex sets $\tilde{S}$ and $\tilde{S}'$ is bijective. Since the convex hull of $\{X_i\}$ (or $\{X'_i\}$ respectively) is the intersection of all such convex sets, it follows that h preserves the convex hull of the points. $\square$

The projective transformations that preserve the convex hull of a given set of points form an important class, and will be called *quasi-affine* transformations.

Definition 21.3. Let B be a subset of $\mathbb{R}^n$ and let h be a projectivity of $\mathbb{P}^n$. The projectivity h is said to be “quasi-affine” with respect to the set B if h preserves the convex hull of the set B .

It may be verified that if h is quasi-affine with respect to B , then h^{-1} is quasi-affine with respect to $h(B)$. Furthermore, if h is quasi-affine with respect to B and g is quasi-affine with respect to $h(B)$, then $g \circ h$ is quasi-affine with respect to B . Thus, quasi-affine projectivities may be composed in this fashion. Strictly speaking, however, quasi-affine projectivities with respect to a given fixed set of points do not form a group.

We will be considering sets of points $\{X_i\}$ and $\{X'_i\}$ that correspond via a projectivity. When we speak of the projectivity being *quasi-affine*, we will mean with respect to the set $\{X_i\}$.

Two-dimensional quasi-affine mappings

Two-dimensional quasi-affine mappings arise as transformations between planar point sets in 3D and their image under a projective camera mapping, as stated formally below.

Theorem 21.4. *If B is a point set in a plane (the “object plane”) in $\mathbb{R}^3$ and B lies entirely in front of a projective camera, then the mapping from the object plane to the image plane defined by the camera is quasi-affine with respect to B .*

Proof. That there is a projectivity h mapping the object plane to the image plane is well known. What is to be proved is that the projectivity is quasi-affine with respect to B . Let L be the line in which the principal plane of the camera meets the object plane. Since B lies entirely in front of the camera, L does not meet the convex hull of B . However, by definition of the principal plane $h(L) = L_\infty$, where L_∞ is the line at infinity in the image plane. Thus, one deduces that $h(\overline{B}) \cap L_\infty = \emptyset$, and hence by theorem 21.2 the transformation is quasi-affine with respect to B . $\square$

Note that if points x_i are visible in an image, then the corresponding object points must lie in front of the camera. Applying theorem 21.4 to a sequence of imaging operations (for instance, a picture of a picture of a picture, etc.), it follows that the original and final images in the sequence are connected by a planar projectivity which is quasi-affine with respect to any point set in the object plane visible in the final image.

Similarly, if two images are taken of a set of points $\{X_i\}$ in a plane, $\{x_i\}$ and $\{x'_i\}$ being corresponding points in the two images, then there is a quasi-affine mapping (with respect to the x_i) mapping each x_i to x'_i , and so theorem 21.2 applies, yielding the following:

Result 21.5. *If $\{x_i\}$ and $\{x'_i\}$ are corresponding points in two views of a set of object points $\{X_i\}$ lying in a plane, then there is a matrix H representing a planar projectivity such that $Hx_i = w_i x'_i$ and all w_i have the same sign.*

21.2 Front and back of a camera

The depth of a point $X = (X, Y, Z, T)^T$ with respect to a camera was shown in (6.15–p162) to be given by

$$\text{depth}(X; P) = \frac{\text{sign}(\det M)w}{T\|m^3\|} \quad (21.1)$$

where M is the left hand 3×3 block of P , m^3 is the third row of M , and $PX = w\hat{x}$. This expression is not dependent on the particular homogeneous representation of X or M , that is it is unchanged by multiplication by non-zero scale factors. This definition of depth is used to determine whether a point is in front of a camera or not.

Result 21.6. *The point X lies in front of the camera P if and only if $\text{depth}(X; P) > 0$.*

In fact, depth is positive for points in front of the camera, negative for points behind the camera, infinite on the plane at infinity and zero on the principal plane of the camera. If the camera centre or the point X is at infinity, then depth is not defined.

Usually, in this section we will only be concerned with the sign of depth and not its magnitude. We may then write

$$\text{depth}(X; P) \doteq w^T \det M \quad (21.2)$$

where the symbol $\doteq$ indicates equality of sign. The quantity $\text{sign}(\text{depth}(X; P))$ will be referred to as the *cheirality* of the point X with respect to the camera P . The cheirality of a point is said to be reversed by a transformation if it is swapped from 1 to -1 or vice versa.

21.3 Three-dimensional point sets

In this section the connection between cheirality of points with respect to a camera and convex hulls of point sets will be explained. The main result is stated now.

Theorem 21.7. *Let P^E and P'^E be two cameras, X_i^E a set of points lying in front of both cameras, and x_i and x'_i the corresponding image points. (The superscript E stands for Euclidean.)*

- (i) *Let $(P, P', \{X_i\})$ be any projective reconstruction from the image correspondences $x_i \leftrightarrow x'_i$, and let $PX_i = w_i \hat{x}_i$ and $P'X_i = w'_i \hat{x}'_i$. Then $w_i w'_i$ has the same sign for all i .*
- (ii) *If each X_i is a finite point, and $P\hat{X}_i = w_i \hat{x}_i$ with w_i having the same sign for all i , then there exists a quasi-affine transformation H taking each X_i to X_i^E .*

Of course, the existence of a projective transformation taking each X_i to X_i^E is guaranteed by theorem 10.1(p266). The current theorem gives the extra information that the transformation is quasi-affine, and hence one has a quasi-affine reconstruction.

Note that the condition that each $w_i w'_i$ have the same sign is unaffected by multiplying P, P' or any of the points X_i by a scale factor, and hence is invariant of the choice of homogeneous representative for any of these quantities. In particular, if P is multiplied by a negative constant, then so is w_i for all i . Thus the sign of $w_i w'_i$ is inverted for each i , but they still all have the same sign. Similarly, if one point X_i is multiplied by a negative constant, then both w_i and w'_i change signs, and so the sign of $w_i w'_i$ is unchanged. In the same way, the condition that each w_i (in part (ii) of the theorem) have the same sign for all i is unaffected if the camera matrix is multiplied by a negative (or of course positive) constant.

Proof. The points X_i^E lie in front of the cameras P^E and P'^E , and hence have positive depth with respect to these cameras. According to (21.2)

$$\text{depth}(X_i^E, P^E) \doteq \det(M^E) w_i T_i^E.$$

Hence, $\det(M^E) w_i T_i^E > 0$ for all i . Similarly for the second camera, $\det(M'^E) w'_i T_i^E > 0$. Multiplying these expressions together, and cancelling T_i^E because it appears twice,

gives $w_i w'_i \det M^E \det M'^E > 0$. Since $\det M^E \det M'^E$ is constant, this shows that $w_i w'_i$ has constant sign.

This was shown in terms of the true configuration. Note however that for any H , one has $w_i \hat{x}_i = P^E X_i^E = (P^E H^{-1})(H X_i^E)$, and hence $w_i w'_i$ has the same sign for the projective reconstruction $(P^E H^{-1}, P'^E H^{-1}, \{H X_i^E\})$. Since any projective reconstruction is of this form (except for homogeneous scale factors), and the condition that $w_i w'_i$ has the same sign is independent of choice of homogeneous representatives of X_i^E , P^E and P'^E , it follows that in any projective reconstruction, $w_i w'_i$ has the same sign for all i . This proves the first part of the theorem.

To show the second part, suppose in the projective reconstruction $w_i \hat{x}_i = P \hat{X}$ all the w_i have the same sign. Since this is a projective reconstruction, there exists a transformation represented by H such that $H \hat{X}_i = \eta_i \hat{X}_i^E$ and $PH^{-1} = \epsilon P^E$ for some constants η_i and ϵ . Then,

$$w_i \hat{x}_i = P \hat{X}_i = (PH^{-1})(H \hat{X}_i) = (\epsilon P^E)(\eta_i \hat{X}_i^E) \quad .$$

and so

$$P^E \hat{X}_i^E = (w_i / \epsilon \eta_i) \hat{x}_i$$

for all i . However, since $\text{depth}(X_i^E, P^E) > 0$, one has $\det(M^E) w_i / \epsilon \eta_i > 0$ for all i . Since $\det(M^E) / \epsilon$ is constant, and by hypothesis w_i has the same sign for all i , it follows that η_i has the same sign for all i . Thus the mapping H such that $H \hat{X}_i = \eta_i \hat{X}_i^E$ is a quasi-affine map with respect to the points $\hat{X}_i$, according to theorem 21.2.

Note that the condition that w_i have the same sign for all i needs to be checked for one of the cameras only. However, defining $P' \hat{X}_i = w'_i \hat{x}_i$, according to part (i) of the theorem, $w_i w'_i$ has the same sign for all i . Thus, if all w_i have the same sign, then so do all w'_i . $\square$

21.4 Obtaining a quasi-affine reconstruction

According to theorem 21.7, any projective reconstruction in which $P \hat{X}_i = w_i \hat{x}_i$ and w_i has the same sign for all i is a quasi-affine reconstruction. The advantage of a quasi-affine reconstruction is that it gives a closer approximation to the true shape of the object than does an arbitrary projective transformation. It may be used as a stepping stone on the way to a metric reconstruction of the scene, as in [Hartley-94b]. In addition, one may retrieve the convex hull of the object or determine such questions as whether two points lie on the same side of a plane.

It turns out that quasi-affine reconstruction is extremely simple, given a projective reconstruction, as shown in the following theorem.

Theorem 21.8. *Any projective reconstruction in which one of the cameras is an affine camera is a quasi-affine reconstruction.*

Proof. Recall that an affine camera is one for which the last row is of the form $(0, 0, 0, 1)$. In this case, writing $w_i \hat{x}_i = P \hat{X}_i$, one immediately verifies that $w_i = 1$ for all i , and in particular they all have the same sign. According to theorem 21.2 this

means that the reconstruction differs by a quasi-affine transformation from the truth. $\square$

The following result follows immediately.

Corollary 21.9. *Let $(P, P', \{X_i\})$ be a projective reconstruction of a scene in which $P = [I \mid 0]$. Then by swapping the last two columns of P and of P' , as well as the last two coordinates of each X_i , one obtains a quasi-affine reconstruction of the scene.*

This is similar to result 10.4(p271), in which it was shown that if the camera P is known in reality to be an affine camera, then the above procedure provides an affine reconstruction.

21.5 Effect of transformations on cheirality

At this point, it is desirable to derive a slightly different form of the formula for depth defined in (21.2). Let P be a camera matrix. The centre of P is the unique point C such that $PC = 0$. One can write an explicit formula for C as follows.

Definition 21.10. Given a camera matrix P , we define C_P^T to be the vector (c_1, c_2, c_3, c_4) , where

$$c_i = (-1)^i \det \bar{P}^{(i)}$$

and $\bar{P}^{(i)}$ is the matrix obtained by removing the i -th column of P .

We denote by $[P/V^T]$ the 4×4 matrix made up of a 3×4 camera matrix P augmented with an final row V^T . Definition 21.10 leads to a simple formula for $\det[P/V^T]$. Cofactor expansion of the determinant along the last row gives $\det[P/V^T] = V^T C_P$ for any row vector V^T . As a special case, if p_i^T is the i -th row of P , then

$$p_i^T C_P = \det[P/p_i^T] = 0$$

where the last equality is true because the matrix has a repeated row. Since this is true for all i , it follows that $PC_P = 0$, and so C_P is the camera centre, as claimed.

Note that submatrix $\bar{P}^{(4)}$ is the same as matrix M in the decomposition $P = [M \mid v]$, and so $\det M = c_4$. This allows us to reformulate (21.2), as follows.

$$\text{depth}(X; P) \doteq w(E_4^T X)(E_4^T C_P) \quad (21.3)$$

where E_4^T is the vector $(0, 0, 0, 1)$. It is significant to note here that E_4 is the vector representing the plane at infinity – a point X lies on the plane at infinity if and only if $E_4^T X = 0$.

We now consider a projective transformation represented by matrix H . If $P' = PH^{-1}$ and $X' = HX$ then the image correspondences are preserved by this transformation. When speaking of a projective transformation being applied to a set of points and to a camera, it is meant that a point X is transformed to HX and the camera matrix is transformed to PH^{-1} .

In this section we will consider such projective transformations and their effect on the cheirality of points with respect to a camera. First, we wish to determine what

happens to C_P when P is transformed to PH^{-1} . To answer that question, consider an arbitrary 4-vector V . We see that

$$V^T H^{-1} C_{PH^{-1}} = \det(PH^{-1}/V^T H^{-1}) = \det(P/V^T) \det H^{-1} = V^T C_P \det H^{-1}.$$

Since this is true for all vectors V , it follows that $H^{-1} C_{PH^{-1}} = C_P \det H^{-1}$, or

$$C_{PH^{-1}} = H C_P \det H^{-1}. \quad (21.4)$$

At one level, this formula is saying that the transformation H takes the camera centre $C = C_P$ to the new location $C_{PH^{-1}} \approx HC$. However, we are interested in the exact coordinates of $C_{PH^{-1}}$, especially the sign of the last coordinate c_4 which appears in (21.3). Thus, the factor $\det H^{-1}$ is significant.

Now, applying (21.4) to (21.3) gives

$$\begin{aligned} \text{depth}(HX; PH^{-1}) &\doteq w(E_4^T HX)(E_4^T C_{PH^{-1}}) \\ &\doteq w(E_4^T HX)(E_4^T H C_P) \det H^{-1}. \end{aligned}$$

One may interpret the expression $E_4^T H$ as being the plane π_∞ mapped to infinity by H . This is because a point X lies on π_∞ if and only if the last coordinate of HX is zero – that is $E_4 HX = 0$. On the other hand, X lies on π_∞ if and only if $\pi_\infty^T X = 0$. Finally, denoting the fourth row of the transformation matrix H by h_4^T , and $\text{sign}(\det H)$ by δ , we obtain

Result 21.11. *If π_∞ is the plane mapped to infinity by a projective transformation H and $\delta = \text{sign}(\det H)$, then*

$$\text{depth}(HX; PH^{-1}) \doteq w(\pi_\infty^T X)(\pi_\infty^T C_P) \delta.$$

This equation will be used extensively. It may be considered to be a generalization of (21.3). It will be seen in the next section that $\delta = \text{sign}(\det H)$ is an indicator of whether H is an orientation-reversing or orientation-preserving transformation. Thus, the effect on cheirality of applying a transformation H is determined only by the position of the plane mapped to infinity π_∞ , and whether H preserves or reverses orientation.

We now consider the effect of different transformations on the cheirality of points with respect to a camera. The effect of an affine transformation is considered first.

Result 21.12. *An affine transformation with positive determinant preserves the cheirality of any point with respect to a camera. An affine transformation with negative determinant reverses cheirality.*

Proof. An affine transformation preserves the plane at infinity, hence $\pi_\infty = E_4$. The result then follows by comparing (21.3) and result 21.11. $\square$

We now determine how an arbitrary projective transformation affects cheirality.

Result 21.13. *Let H represent a projective transformation with positive determinant, and let π_∞ be the plane in space mapped to infinity by H . The cheirality of a point X is preserved by H if and only if X lies on the same side of the plane π_∞ as the camera centre.*

Proof. Since $\det H > 0$, we see from (21.3) and result 21.11 that $\text{depth}(X; P) \doteq \text{depth}(HX; PH^{-1})$ if and only if $(\pi_\infty^T X)(\pi_\infty^T C) \doteq (E_4^T x)(E_4^T C)$. Suppose the point X and the camera P are located at finite points so that the cheirality is well defined, and let them be scaled so that both X and C have final coordinate equal to 1. In this case, $(E_4^T X)(E_4^T C) = 1$ and we see that cheirality is preserved if and only if $(\pi_\infty^T X)(\pi_\infty^T C) \doteq 1$, or otherwise expressed $\pi_\infty^T X \doteq \pi_\infty^T C$. This condition may be interpreted as meaning that the points C and X both lie on the same side of the plane π_∞ . Hence, the cheirality of a point X is preserved by a transformation H , if and only if it lies on the same side of the plane π_∞ as the camera centre. $\square$

21.6 Orientation

We now consider the question of image orientation. A mapping h from $\mathbb{R}^n$ to itself is called orientation-preserving at points X where the Jacobian of h (the determinant of the matrix of partial derivatives) is positive, and orientation-reversing at points where the Jacobian is negative. Reflection of points of $\mathbb{R}^n$ with respect to a hyperplane (that is mirror imaging) is an example of an orientation-reversing mapping. A projectivity h from P^n to itself restricts to a mapping from $\mathbb{R}^n - \pi_\infty$ to $\mathbb{R}^n$, where π_∞ is the hyperplane (line, plane) mapped to infinity by H . Consider the case $n = 3$ and let H be a 4×4 matrix representing the projectivity h . We wish to determine at which points X in $\mathbb{R}^n - \pi_\infty$ the map h is orientation-preserving. It may be verified (quite easily using Mathematica [Mathematica-92]) that if $H\hat{X} = w\hat{X}'$ and J is the Jacobian matrix of h evaluated at X , then $\det(J) = \det(H)/w^4$. This gives the following result.

Result 21.14. *A projectivity h of $\mathbb{P}^3$ represented by a matrix H is orientation-preserving at any point in $\mathbb{R}^3 - \pi_\infty$ if and only if $\det(H) > 0$.*

Of course, the concept of orientability may be extended to the whole of $\mathbb{P}^3$, and it may be shown that h is orientation-preserving on the whole of $\mathbb{P}^3$ if and only if $\det(H) > 0$. The essential feature here is that as a topological manifold, $\mathbb{P}^3$ is orientable.

Two sets of points $\{X_i\}$ and $\{\bar{X}_i\}$ that correspond via a quasi-affine transformation are said to be *oppositely oriented* if the transformation is orientation-reversing. As an example, consider the transformation given by a diagonal matrix $H = \text{diag}(1, 1, -1, 1)$. This transformation has negative determinant, and hence is orientation-reversing. On the other hand, it is affine, and hence quasi-affine. Therefore, it is possible always to construct oppositely oriented quasi-affine reconstructions of a scene. It may appear therefore that the orientation of a scene may not be determined from a pair of images. Although this is sometimes true, it is sometimes possible to rule out one of the oppositely oriented quasi-affine reconstructions of a scene, and hence determine the true orientation of the scene.

Common experience provides some clues here. In particular a stereo pair may be viewed by presenting one image to one eye and the other image to the other eye. If this is done correctly, then the brain perceives a 3D reconstruction of the scene. If, however, the two images are swapped and presented to the opposite eyes, then the perspective will be reversed – hills become valleys and vice versa. In effect, the brain is able to compute two oppositely oriented reconstructions of the image pair. It seems, therefore,




Fig. 21.2. Stereo pairs of images that may be viewed by cross-fusing (the eyes are crossed so that the left eye sees the right image and vice versa). The two bottom images are the same as the top pair, except that they are swapped. In the top pair of images one sees an L-shaped region raised above a planar background. In the bottom pair of images the L-shaped region appears as an indentation. The two "reconstructions" differ by reflection in the background plane. This demonstrates that the same pair of images may give rise to two differently oriented projective reconstructions.

that in certain circumstances, two oppositely oriented realizations of an image pair exist. This is illustrated in figure 21.2.

It may be surprising to discover that this is not always the case, as is shown in the following theorem. As used in this theorem and elsewhere in this chapter, a projective realization of a set of point correspondences is known as a *strong realization* if the reconstructed 3D points X_i are in front of all the cameras.

Theorem 21.15. *Let $(P, P', \{X_i\})$ be a strong realization of a uniquely realizable set of point correspondences. There exists a different oppositely oriented strong realization $(\bar{P}, \bar{P}', \{\bar{X}_i\})$ if and only if there exists a plane in $\mathbb{R}^3$ such that the perspective centres of both cameras P and P' lie on one side of the plane, and the points X_i lie on the other side.*

Proof. Consider one strong realization of the configuration. By definition, all the points lie in front of both cameras. Suppose that there exists a plane separating the points from the two camera centres. Let G be a projective transformation mapping the given plane to infinity and let A be an affine transformation. Suppose further that

$\det G > 0$ and $\det A < 0$. Let H be the composition $H = AG$. According to result 21.13 the transformation G is cheirality-reversing for the points, since the points are on the opposite side of the plane from the camera centres. According to result 21.12 A is also cheirality-reversing, since $\det A < 0$. The composition H must therefore be cheirality-preserving, and it transforms the strong configuration to a different strong configuration. Since H has negative determinant, however, it is orientation-reversing, so the two strong realizations have opposite orientations.

Conversely, suppose that two strong oppositely oriented realizations exist and let H be the transformation taking one to the other. Since H is orientation-reversing, $\det H < 0$. The mapping H is by definition cheirality-preserving on all points, with respect to both cameras. If π_∞ is the plane mapped to infinity by H , then according to result 21.13 the points X must lie on the opposite side of the plane π_∞ from both camera centres. $\square$

21.7 The cheiral inequalities

In section 21.4 a very simple method was given for obtaining a quasi-affine reconstruction of a scene directly from a projective reconstruction. However, the reconstruction obtained there did not respect the condition that the points must lie in front of all cameras. In fact, the first camera in this construction is an affine camera, for which front and back are not well defined. By taking full advantage of the fact that visible points must lie in front of the camera, it is possible to constrain the reconstruction more tightly, leading to a closer approximation to a true affine reconstruction of the scene.

The method will be given for the case of a reconstruction derived from several images. One is given a set of image points $\{x_i^j\}$, where x_i^j is the projection of the i -th point in the j -th image. Not all points are visible in each image, so for some (i, j) the point x_i^j is not given, in which case it is not known whether the i -th point lies in front of the j -th camera or not. On the other hand, the existence of an image point x_i^j implies that the point lies in front of the camera.

We start from an assumed projective reconstruction of the scene, consisting of a set of 3D points X_i and cameras P^j such that $x_i^j \approx P^j X_i$. Writing the implied scalar constant explicitly in this equation gives $w_i^j \tilde{x}_i^j = P^j X_i$. In this equation, P^j and X_i are arbitrarily chosen homogeneous representatives of the respective matrix or vector. Related to theorem 21.7 one may state for several views:

Result 21.16. *Consider a set of points X_i^E and cameras P^{jE} , and let $x_i^j = P^{jE} X_i^E$ be defined for some indices (i, j) such that point X_i^E lies in front of camera P^{jE} . Let $(P^j; X_i)$ be a projective reconstruction from x_i^j . Then there are camera matrices $\tilde{P}^j = \pm P^j$ and $\tilde{X}_i = \pm X_i$ such that for each (i, j) for which x_i^j is defined, one has*

$$\tilde{P}^j \tilde{X}_i = w_i^j \tilde{x}_i^j \text{ with } w_i^j > 0.$$

Briefly stated, one can always adjust a projective reconstruction, multiplying camera matrices and points by -1 if necessary, so that w_i^j is positive whenever image point x_i^j exists. The simple proof is omitted. To find the matrices $\tilde{P}^j$ and points $\tilde{X}_i$, one may assume that one of the cameras $\tilde{P}_1 = P_1$, for otherwise all points and cameras may be multiplied by -1 . The condition $P_1 \tilde{X}_i = w_i^1 x_i^1$ with $w_i^1 > 0$ determines whether to

choose $\tilde{X}_i = X_i$ or $-X_i$ for all i such that x_i^j is defined. Each known $\tilde{X}_i$ determines $\tilde{P}^j$ for all j such that x_i^j is defined. Continuing in this way, one easily finds the factor ± 1 to apply to each P^j and X_j to find $\tilde{P}^j$ and $\tilde{X}_j$. We assume that this has been done, and replace each P^j by $\tilde{P}^j$ and X_i by $\tilde{X}_i$. In future we drop the tildes and continue to work with the corrected P^j and X_i . We now know that $w_i^j > 0$ whenever image point x_i^j is given.

Now, we seek a transformation H that will transform the projective reconstruction to a quasi-affine reconstruction for which all points lie in front of the cameras as appropriate. Denoting by 4-vector $\mathbf{v}$ the plane π_∞ mapped to infinity by H , this condition may be written as (see result 21.11):

$$\text{depth}(X_i; P^j) \doteq (\mathbf{v}^T X_i)(\mathbf{v}^T C) \delta > 0$$

where $\delta = \text{sign}(\det H)$. This condition holds for all pairs (i, j) where x_i^j is given.

Since we are free to multiply $\mathbf{v}$ by -1 if necessary, we may assume that $(\mathbf{v}^T C^1) \delta > 0$ for the centre of the camera P^1 . The following inequalities now follow easily:

$$\begin{aligned} X_i^T \mathbf{v} &> 0 \quad \text{for all } i \\ \delta C'^T \mathbf{v} &> 0 \quad \text{for all } j. \end{aligned} \tag{21.5}$$

These equations (21.5) may be called the *cheiral inequalities*. Since the values of each X_i , C and C' are known, they form a set of inequalities in the entries of $\mathbf{v}$. The value of δ is not known *a priori*, and so it is necessary to seek a solution for each of the two cases $\delta = 1$ and $\delta = -1$.

To find the required transformation H , first of all we solve the cheiral inequalities to find a value of $\mathbf{v}$, either for $\delta = 1$ or $\delta = -1$. The required matrix H is any matrix having $\mathbf{v}^T$ as its last row and satisfying the condition $\det H \doteq \delta$. If the last component of $\mathbf{v}$ is non-zero, then H can be chosen to have the simple form in which the first three rows are of the form $\pm [I \mid 0]$.

If a Euclidean reconstruction (or more specifically a quasi-affine reconstruction) is possible, then there must be a solution either for $\delta = 1$ or $\delta = -1$. In some cases there will exist solutions of the cheiral inequalities for both $\delta = 1$ and $\delta = -1$. This will mean that two oppositely oriented strong realizations exist. The conditions under which this may occur were discussed in section 21.6.

Solving the cheiral inequalities

Naturally, the cheiral inequalities may be solved using techniques of linear programming. As they stand however, if $\mathbf{v}$ is a solution then so is $\alpha \mathbf{v}$ for any positive factor α . In order to restrict the solution domain to be bounded, one may add additional inequalities. For instance, if $\mathbf{v} = (v_1, v_2, v_3, v_4)^T$, then the inequalities $-1 < v_i < 1$ serve to restrict the solution domain to be a bounded polyhedron.

To achieve a unique solution we need to specify some goal function to be linearized. An appropriate strategy is to seek to maximize the extent to which each of the inequalities is satisfied. To do this, we introduce one further variable, d . Each of the inequalities $\mathbf{a}^T \mathbf{v}$ of the form (21.5) for appropriate $\mathbf{a}$ is replaced by an inequality $\mathbf{a}^T \mathbf{v} > d$. We seek

to maximize d while satisfying all the inequalities. This is a standard linear programming problem, for which many methods of solution exist, such as the simplex method ([Press-88]).¹ If a solution is found for which $d > 0$ then this will be a desired solution.

Summary of the algorithm

Now, we give the complete algorithm for computing a quasi-affine reconstruction of a scene using the cheiral inequalities. The algorithm as outlined above was discussed for the case of two views. In the present case it will be presented for an arbitrary number of views. The extension to more views is straightforward.

Objective

Given a set of 3D points $\mathbf{X}_i$ and camera matrices $\mathbf{P}^j$ constituting a projective reconstruction from a set of image points, compute a projective transformation $\mathbf{H}$ transforming the projective to a quasi-affine reconstruction.

Algorithm

- (i) For each pair (i, j) such that point $\mathbf{x}_i^j$ is given let $\mathbf{P}^j \mathbf{X}_i = w_i^j \tilde{\mathbf{x}}_i^j$.
- (ii) Replace some cameras $\mathbf{P}^j$ by $-\mathbf{P}^j$ and some points $\mathbf{X}_i$ by $-\mathbf{X}_i$ as required to ensure that each $w_i^j > 0$.
- (iii) Form the cheiral inequalities (21.5) where $\mathbf{C}^j = \mathbf{C}_{\mathbf{P}^j}$ is defined by definition 21.10.
- (iv) For each of the values $\delta = \pm 1$, choose a solution (if it exists) to the set of cheiral inequalities. Let the solution be $\mathbf{v}_\delta$. A solution must exist for at least one value of δ , sometimes for both values of δ .
- (v) Define a matrix $\mathbf{H}_\delta$ having last row equal to $\mathbf{v}_\delta$ and such that $\det(\mathbf{H}) \doteq \delta$. The matrix $\mathbf{H}_\delta$ is the required transformation matrix. If both $\mathbf{H}_+$ and $\mathbf{H}_-$ exist, then they lead to two oppositely oriented quasi-affine reconstructions.

Algorithm 21.1. *Computing a quasi-affine reconstruction.*

Bounding the plane at infinity

A quasi-affine reconstruction is of course not unique, being defined only up to a quasi-affine transformation with respect to the points and camera centres. However, once one has been found, it is possible to set bounds on the coordinates of the plane at infinity. Thus, let $\mathbf{P}^j$ and $\mathbf{X}_i$ constitute a quasi-affine reconstruction of a scene. One may choose the sign of $\mathbf{P}^j$ and $\mathbf{X}_i$ such that the last coordinates of $\mathbf{X}_i$ and the determinants of each $\mathbf{M}^j$ are positive. One may apply a translation to the points and cameras so that the coordinate origin lies inside the convex hull of the points and camera centres. For simplicity, the centroid of these points may be placed at the origin.

It is possible to apply a further quasi-affine transformation $\mathbf{H}$ to obtain an alternative reconstruction. Let π_∞ be the plane mapped to infinity by $\mathbf{H}$. We confine our interest to orientation-preserving transforms, and wish to find constraints on the coordinates of π_∞ such that $\mathbf{H}$ is quasi-affine. A plane π_∞ has this property if and only if it lies entirely outside the convex hull of the points and camera centres. Since the plane π_∞

¹ The Simplex algorithm given in [Press-88] is not suitable for use as stands, since it makes the unnecessary assumption that all variables are non-negative. It needs to be modified to be used for this problem.

cannot cross the convex hull, it cannot pass through the origin. Representing π_∞ by the vector $\mathbf{v}$, this says that the last coordinate of $\mathbf{v}$ is non-zero. One may then write $\mathbf{v} = (v_1, v_2, v_3, 1)^T$. Since the origin lies on the same side of the plane as all the points, the cheiral inequalities become

$$\begin{aligned} \mathbf{X}_i^T \mathbf{v} &> 0 \text{ for all } i \\ \mathbf{C}^j{}^T \mathbf{v} &> 0 \text{ for all } j \end{aligned} \quad (21.6)$$

for $\mathbf{v} = (v_1, v_2, v_3, 1)^T$. One may find upper and lower bounds for each v_i by solving the linear programming problem to maximize v_i or $-v_i$ subject to these constraints. None of the v_i can be unbounded, since otherwise the plane π_∞ represented by the vector $\mathbf{v}$ could lie arbitrarily close to the origin.

Before solving for this system, good practice is to apply an affine transformation to normalize the set of points and camera centres so that their centroid is at the origin and their principal moments are all equal to 1.

The complete algorithm for computing the bounds on the position of the plane at infinity is given in algorithm 21.2.

Objective

Given a quasi-affine reconstruction of a scene, establish bounds on the coordinates of the plane at infinity.

Algorithm

- (i) Normalize the points $\mathbf{X}_i = (X_i, Y_i, Z_i, T_i)^T$ so that $T_i = 1$, and cameras $\mathbf{P}^j = [\mathbf{M}_j \mid \mathbf{t}^j]$ so that $\det \mathbf{M}^j = 1$.
- (ii) Further normalize by replacing $\mathbf{X}_i$ by $\mathbf{H}^{-1} \mathbf{X}_i$ and $\mathbf{P}^j$ by $\mathbf{P}^j \mathbf{H}$, where $\mathbf{H}$ is an affine transformation moving the centroid to the origin and scaling in the principal axis directions so that it has equal principal axes.
- (iii) Letting $\mathbf{v} = (v_1, v_2, v_3, 1)^T$, form cheiral inequalities (21.6). Any orientation-preserving transformation $\mathbf{H}$ mapping the reconstruction to an affine reconstruction of the image must have the form

$$\mathbf{H} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ v_1 \ v_2 \ v_3 & 1 \end{bmatrix}$$

for a vector $\mathbf{v}$ satisfying these inequalities.

- (iv) Upper and lower bounds for each v_i may be found by running a linear programming problem six times. The coordinates of a desired transformation $\mathbf{H}$ must lie inside the box defined by these bounds.

Algorithm 21.2. *Establishing bounds on the plane at infinity.*

21.8 Which points are visible in a third view

Consider a scene reconstructed from two views. We consider now the question of determining which points are visible in a third view. Such a question arises when one is given two uncalibrated views of a scene and one seeks to synthesize a third view. This can be done by carrying out a projective reconstruction of the scene from the




Fig. 21.3. The point set $\mathbf{X}$ is in front of all three cameras as shown. However, if an orientation-preserving projective transformation $\mathbf{H}$ is applied taking the plane π_∞ to infinity, then the point set will subsequently lie in front of the cameras $\mathbf{P}^1$ and $\mathbf{P}^2$ but behind the camera $\mathbf{P}^3$. Thus, suppose the point set $\mathbf{X}$ is reconstructed from images captured by cameras $\mathbf{P}^1$ and $\mathbf{P}^2$, and let $\mathbf{P}^3$ be any other camera matrix. If a plane exists separating the centre of camera $\mathbf{P}^3$ from the other camera centres, and not meeting the convex hull of the point set $\mathbf{X}$, then it cannot be determined whether the points lie in front of $\mathbf{P}^3$.

first two views and then projecting into the third view. In this case, it is important to determine if a point lies in front of the third camera and is hence visible, or not.

If the third view is given simply by specifying the camera matrix with respect to the frame of reference of some given reconstruction, then it may be impossible to determine whether points are in front of the third camera or behind it in the true scene. The basic ambiguity is illustrated in figure 21.3. Knowledge of a single point known to be visible in the third view serves to break the ambiguity, however, as the following result shows. By applying theorem 21.7(p519) to the first and third views, one obtains the following criterion.

Result 21.17. Let points $(\mathbf{P}^1, \mathbf{P}^2, \{\mathbf{X}_i\})$ be a realization of a set of correspondences $\mathbf{x}_i^1 \leftrightarrow \mathbf{x}_i^2$. Let $\mathbf{P}^3$ be the camera matrix of a third view and suppose that $w_j^3 \hat{\mathbf{x}}_i = \mathbf{P}^t \mathbf{X}_j$ for $i = 1, \dots, 3$. Then $w_j^1 w_j^3$ has the same sign for all points $\mathbf{X}_j$ visible in the third view.

In practice, it will usually be the case that one knows at least one point $\mathbf{X}_0$ visible in the third view. This serves to define the sign $w_0^1 w_0^3$, and any other point $\mathbf{X}_j$ will be in front of the camera $\mathbf{P}^3$ if and only if $w_j^1 w_j^3 \doteq w_0^1 w_0^3$.

As an example, once a projective reconstruction has been carried out using two views, the camera matrix of the third camera may be determined from the images of six or more points known to be in front of it by solving directly for the matrix $\mathbf{P}^3$ given the correspondences $\mathbf{x}_i^3 = \mathbf{P}^3 \mathbf{X}_i$ where points $\mathbf{X}_i$ are the reconstructed points. Then one can determine unambiguously which other points are in front of $\mathbf{P}^3$.




Fig. 21.4. As shown, the point x_1 is closer to the camera than x_2 . However, if an orientation-reversing projectivity is applied, taking π_∞ to infinity, then both x_1 and x_2 remain in front of the camera, but x_2 is closer to the camera than x_1 .


21.9 Which points are in front of which

When we are attempting to synthesize a new view of a scene that has been reconstructed from two or more uncalibrated views it is sometimes necessary to consider the possibility of points being obscured by other points. This leads to the question: given two points that project to the same point in the new view, which one is closer to the camera, and hence obscures the other? In the case where the possibility exists of oppositely oriented quasi-affine reconstructions it may once again be impossible to determine which of a pair of points is closer to the new camera. This is illustrated in figure 21.4. If a plane exists, separating the camera centres from the point set, then two oppositely oriented reconstructions exist, and one cannot determine which points are in front of which. The sort of ambiguity shown in figure 21.4 can only occur in the case where there exists a plane π_∞ that separates the camera centres from the set of all visible points. If this is not the case, then one can compute a quasi-affine reconstruction and the problem is easily solved. To avoid the effort of computing a quasi-affine reconstruction, however, we would like to solve this problem using only a projective reconstruction of the scene. How this may be done is explained next.

One may invert (21.1–p518) to get an expression for $\text{depth}^{-1}(X; P) = 1/\text{depth}(X; P)$. This inverse depth function is infinite on the principal plane of the camera, zero on the plane at infinity, positive for points in front of the camera and negative for points behind the camera. For notational simplicity, we write $\chi(X; P)$ instead of $\text{depth}^{-1}(X; P)$.

For points X lying on a ray through the camera centre, the value of $\chi(X; P)$ decreases monotonically along this ray, from zero at the camera centre, decreasing through positive values to zero at the plane at infinity, thence continuing to decrease through negative values to $-\infty$ at the camera centre. A point X_1 lies closer to the front of the camera than X_2 if and only if $\chi(X_1) > \chi(X_2)$. This is illustrated in figure 21.5.

Now, if the configuration undergoes an orientation-preserving transformation H taking the plane π_∞ to infinity, then the parameter χ will be replaced by a new parameter χ' defined by $\chi'(X) = \chi(HX; PH^{-1})$. The value of χ' must also vary monotonically along the ray. Since H is orientation-preserving, points just in front of the camera centre will remain in front of the camera after the transformation, because of result 21.13.

Fig. 21.5. The parameter χ .

Thus both χ and χ' decrease monotonically in the same direction along the ray. If X_1 and X_2 are two points on the line, then $\chi'(X_1) > \chi'(X_2)$ if and only if $\chi(X_1) > \chi(X_2)$.

In the case where the projective transformation has negative determinant, then the front and back of the camera are reversed locally. In this case the direction of increase of the parameter χ' will be reversed. Consequently, $\chi'(X_1) > \chi'(X_2)$ if and only if $\chi(X_1) < \chi(X_2)$.

In the case where the projective transformation transforms the scene to the “true” scene, of two points that project to the same point in the image, the one with the higher value of χ' is closer to the camera. This leads to the following result that allows us to determine from an arbitrary projective reconstruction which of two points is closer to the front of the camera.

Result 21.18. *Suppose that X_1 and X_2 are two points that map to the same point in an image. Consider a projective reconstruction of the scene and let the parameter χ be defined (by formula (21.3–p521)) in the frame of the projective reconstruction. If the projective reconstruction has the same orientation as the true scene, then the point that lies closer to the front of the camera in the true scene is the one that has the greater value of χ . On the other hand, if the projective transformation has the opposite orientation, then the point with smaller value of χ will lie closer to the front of the camera in the true scene.*

As remarked previously, unless there exists a plane separating the point set from the cameras used for the reconstruction, the orientation of the scene is uniquely determined, and one can determine whether the projective transformation of result 21.18 has positive or negative determinant. However, to do this may require one to compute a strong realization of the configuration by the linear programming method as described in section 21.7. If differently oriented strong realizations exist, then as illustrated by figure 21.4, there is an essential ambiguity. However, this ambiguity may be resolved by knowledge of the relative distance from the camera of a single pair of points.

21.10 Closure

21.10.1 The literature

The topic of this chapter belongs to *Oriented projective geometry*, which is treated in a standard and readable text [Stolfi-91]. Laveau and Faugeras apply the ideas of

oriented projective geometry in [Laveau-96b]. The concepts of front and back of the camera were also used in [Robert-93] to compute convex hulls in projective reconstructions. This chapter derives from the paper [Hartley-98a] which also treats such topics of invariants for quasi-affine mappings and conditions under which arbitrary correspondences allow quasi-affine reconstructions.

Cheirality, and specifically the cheirality inequalities have been useful in determining quasi-affine reconstructions as an intermediate step towards affine and metric reconstruction in [Hartley-94b, Hartley-99]. More recently Werner and Pajdla in [Werner-01] have used oriented projective geometry to eliminate spurious line correspondences, and to constrain correspondences of five points [Werner-03], over two views.

Degenerate Configurations

In past chapters we have given algorithms for the estimation of various quantities associated with multiple images – the projection matrix, the fundamental matrix and the trifocal tensor. In each of these cases, linear and iterative algorithms were given, but little consideration was given to the possibility that these algorithms could fail. We now consider under what conditions this might happen.

Typically, if sufficiently many point correspondences are given in some sort of “general position” then the quantities in question will be uniquely determined, and the algorithms we have given will succeed. However, if there are too few point correspondences given, or else all the points lie in certain critical configurations, then there will not be a unique solution. Sometimes there will be a finite number of different solutions, and sometimes a complete family of solutions.

This chapter will concentrate on three of the main estimation problems that we have encountered in this book, camera resectioning, reconstruction from two views and reconstruction from three views. Some of the results given here are classical, particularly the camera resectioning and two-view critical surface problems. Others are more recent results. We consider the different estimation problems in turn.

22.1 Camera resectioning

We begin by considering the problem of computing the camera projection matrix, given a set of points in space and the corresponding set of points in the image. Thus, one is given a set of points X_i in space that are mapped to points x_i in the image by a camera with projection matrix P . The coordinates of the points in space and the image are known, and the matrix P is to be computed. This problem was considered in chapter 7. Before considering the critical configurations for this problem, we will digress to look at an abstraction of the camera projection.

Cameras as points

Suppose the existence of a set of correspondences $X_i \leftrightarrow x_i$. Let us suppose that there is a unique camera matrix P such that $x_i = PX_i$. Now, let H be a matrix representing a projective transformation of the image, and let $x'_i = Hx_i$ be the set of transformed image coordinates. Then it is clear that there is a unique camera matrix P' such that

$\mathbf{x}'_i = P'\mathbf{x}_i$, namely the camera matrix $P' = HP$. Conversely, if there exists more than one camera matrix P mapping $\mathbf{X}_i$ to $\mathbf{x}_i$, then there exists more than one camera matrix P' mapping $\mathbf{X}_i$ to $\mathbf{x}'_i$. Thus, the existence or not of a unique solution to the problem of determining the projection matrix P is dependent on the image points $\mathbf{x}_i$ *only up to a projective transformation* H of the image.

Next, observe that applying a projective transformation H to a camera matrix P does not change the camera centre. Specifically, the point C is the camera centre if and only if $PC = 0$. However $PC = 0$ if and only if $HPC = 0$. Thus, the camera centre is preserved by a projective transformation of the image. Next, we show that this is essentially the only property of the camera that is preserved.

Result 22.1. *Let P and P' be two camera matrices with the same centre. Then there exists a projective image transformation represented by a non-singular matrix H such that $P' = HP$.*

Proof. If the centre C is not at infinity, then the camera matrices are of the form $P = [M \mid -Mc]$ and $P' = [M' \mid -M'c]$, where c is the inhomogeneous 3-vector representing the camera centre. Then clearly, $P' = M'M^{-1}P$. If C is a point at infinity, then one chooses a 3D projective transformation G such that GC is a finite point, say $\bar{C}$. In this case, the two camera matrices PG^{-1} and $P'G^{-1}$ both have the same centre, namely $\bar{C}$. It follows that $P'G = HPG$ for some H . Cancelling G gives $P' = HP$. $\square$

This result may be interpreted as saying that an image is determined up to projectivity by the camera centre alone. Thus, we see that in considering the problem of uniqueness of the camera matrix, one may ignore all the parameters of the camera, except the camera centre, since this alone determines the projectivity type of the image, and hence the uniqueness or not of a solution.

Images as equivalence classes of rays

To gain insight into the critical configurations of camera resectioning, we turn first to consider 2-dimensional cameras, mapping $\mathbb{P}^2$ to $\mathbb{P}^1$. Consider a camera centre c and a set of points $\mathbf{x}_i$ in space. The rays $c\mathbf{x}_i$ intersect an image line l at a set of points $\bar{\mathbf{x}}_i$; thus points $\bar{\mathbf{x}}_i$ are the images of the points $\mathbf{x}_i$. The projection of the points $\mathbf{x}_i$ to points $\bar{\mathbf{x}}_i$ in the 1D image may be described by a 2×3 projection matrix as described in section 6.4.2(p175).

As shown in chapter 2, the projective equivalence type of the set of rays $c\mathbf{x}_i$ is the same as that of the image points $\bar{\mathbf{x}}_i$. This is illustrated in figure 22.1. Thus, instead of considering an image as being the set of points on the image line, the image may be thought of as the projective equivalence class of the set of rays from the camera centre through each of the image points. In the case of just 4 image points, the cross ratio of the points $\bar{\mathbf{x}}_i$ (or equivalently, the rays) characterizes their projective equivalence class. To give a specific notation, we denote by $\langle c; \mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n \rangle$ the projective equivalence class of the set of rays $c\mathbf{x}_i$.

The same remarks are valid for projections of 3D points into a 2-dimensional image. One may also extend the above notation by writing $\langle C; \mathbf{X}_1, \dots, \mathbf{X}_n \rangle$ to represent the




Fig. 22.1. Projection of points in the plane using a 2D camera. A 1D image is formed by the intersection of the rays $l_i = \overline{cX_i}$ with the image line l . The set of image points $\{\bar{x}_i\}$ is projectively equivalent to the set of rays $\{l_i\}$. For four points, the projective equivalence class of the image is determined by the cross ratio of the points.

projective equivalence class of all the rays $\overline{cX_i}$. As in the 2-dimensional case, this is an abstraction of the projection of the points X_i relative to a camera with centre at C .

We will be considering configurations of camera centres and 3D points, which will be denoted by $\{C_1, \dots, C_m; X_1, \dots, X_n\}$ or variations thereof. Implicit is that the symbols appearing before the semicolon are camera centres, and those that come after are 3D points. In order to make the statements of derived results simple, the concept of *image equivalence* is defined.

Definition 22.2. Two configurations

$$\{C_1, \dots, C_m; X_1, \dots, X_n\} \text{ and } \{C'_1, \dots, C'_m; X'_1, \dots, X'_n\}$$

are called *image-equivalent* if $\langle C_i; X_1, \dots, X_n \rangle = \langle C'_i; X'_1, \dots, X'_n \rangle$ for all $i = 1, \dots, m$.

The concept of image equivalence is distinct from projective equivalence of the sets of points and camera centres involved. Indeed, the relevance of this to reconstruction ambiguity is that if a configuration $\{C_1, \dots, C_m; X_1, \dots, X_n\}$ allows another image-equivalent set which is not projective equivalent, then this amounts to an ambiguity of the projective reconstruction problem, since the projective structure of the points and cameras is not uniquely defined by the set of images. In this case, we say that the configuration $\{C_1, \dots, C_m; X_1, \dots, X_n\}$ *allows an alternative reconstruction*.

22.1.1 Ambiguity in 2D – Chasles' theorem

Before considering the usual 3D cameras, we discuss the simpler case of 2D cameras. The analysis of the uniqueness of 2D camera projections involves planar conics. Ambiguity in determining the camera centre from the projection of a set of known points x_i means that the projection of the points is the same from two different camera centres c and c' . The question is for what configurations of the points this may occur. The answer to this question is given by Chasles' Theorem.




Fig. 22.2. The points $x_1, \dots, x_n$ project in equivalent ways to the two camera centres c and c' .

Theorem 22.3 Chasles' Theorem. Let x_i be a set of n points and c and c' two camera centres, all lying in a plane. Then

$$\langle c; x_1, \dots, x_n \rangle = \langle c'; x_1, \dots, x_n \rangle$$

if and only if one of the following conditions holds

- (i) The points c , c' and all x_i all lie on a non-degenerate conic, or
- (ii) All points lie on the union of two lines (a degenerate conic), both camera centres lying on the same line.

These two configurations are shown in figure 22.2.

Note that as a simple corollary of this theorem, if c'' is any other camera centre lying on the same component of a conic (degenerate or non-degenerate) as c , c' and the x_i , then the projection of the points to the centre c'' is equivalent to their projection to the original camera centres. Furthermore, any number of extra points x_j may be added without breaking the ambiguity, as long as they lie on the same conic.

22.1.2 Ambiguity for 3D cameras

We now address the problem of ambiguity of camera resection in the case of full 3D cameras. Twisted cubics (described in section 3.3(p75)) play an analogous role in the 3D case to that played by conics in the case of 2D cameras. The degenerate form of a twisted cubic consisting of a conic plus a line intersecting the conic arises in this context; so does the degenerate cubic consisting of three lines. As in the 2D case, when ambiguity arises from points lying on a degenerate cubic, the camera centres must both lie on the same component.

A complete classification of the point and camera configurations leading to ambiguous camera resectioning is given in figure 22.3 and also in the following definition. For the present we describe the geometric configurations. The exact relevance to ambiguous camera configurations will be given afterwards.

Definition 22.4.

A *critical set* for camera resectioning consists of two parts:

- (i) An algebraic curve $\mathcal{C}$ containing the camera centres, plus any number of the 3D points. This curve may be
 - (a) a non-degenerate twisted cubic (degree 3),
 - (b) a planar conic (degree 2),
 - (c) a straight line (degree 1), or
 - (d) a single point (degree 0).
- (ii) A union of linear subspaces L (lines or planes) containing any number of 3D points.

The curve $\mathcal{C}$ and the linear subspaces satisfy three conditions:

- (i) Each of the linear subspaces must meet the curve $\mathcal{C}$.
- (ii) The sum of the degree of the curve $\mathcal{C}$ and the dimensions of the linear subspaces is at most 3.
- (iii) Except in the case where $\mathcal{C}$ is a single point, the cameras do not lie at the intersection point of $\mathcal{C}$ and the linear subspaces.

The various possibilities are shown in figure 22.3, and it is easily verified that these completely enumerate all configurations in accordance with definition 22.4.

Result 22.5. *The different possible critical sets for camera resectioning are:*

- (i) A non-degenerate twisted cubic $\mathcal{C}$ (degree 3).
- (ii) A plane conic $\mathcal{C}$ (degree 2), plus a line (dimension 1) that meets the conic.
- (iii) A line $\mathcal{C}$ (degree 1) plus up to two other lines (total dimension 2) that meet the first line.
- (iv) A line $\mathcal{C}$ (degree 1), plus a plane (dimension 2).
- (v) A point $\mathcal{C}$ (degree 0) and up to three lines passing through the point (total dimension 3).
- (vi) A point $\mathcal{C}$ (degree 0) and a line and a plane both passing through this point.

The exact relationship of these critical sets to ambiguous camera resectioning is given by the following result.

Result 22.6. *Let P and P' be two different camera matrices with camera centres C_0 and C_1 . Then the two camera centres and the points X_i that satisfy $PX_i = P'X_i$ all lie on a critical set as given by definition 22.4.*

Furthermore, if $P_\theta = P + \theta P'$ has rank 3¹, then the camera centre of the camera defined by P_θ lies on the component $\mathcal{C}$ of the critical set containing the two original camera centres C_0 and C_1 , and $P_\theta X = PX = P'X$ for any point X in the critical set.

Conversely, let P be a camera matrix with centre C_0 . Let C_1 and a set of 3D points X_i lie in a critical set for camera resectioning. Let C_1 be any other point lying on the component $\mathcal{C}$ of the critical set, with $C_1 \neq C_0$ unless $\mathcal{C}$ consists of a single point. Then there exists a camera matrix P' (different from P) with camera centre C_1 such that $PX_i = P'X_i$ for all points X_i .

¹ We include the case $P_\infty = P'$.




Fig. 22.3. The different critical configurations for camera resectioning from a single view. The open circles represent centres of projection and the filled circles represent points. Each case consists of an algebraic curve (or single point) C containing the camera centres, plus a set of linear subspaces (lines or planes).

Proof. An outline of the proof will be given, with details left for the reader to fill in. We temporarily need to distinguish between equality of homogeneous quantities up to a scale factor (which will be denoted by $\approx$) and absolute equality (which is denoted by $=$). Suppose that X maps to the same image point with respect to cameras P and P' . One can write $PX \approx P'X$. Taking account of the scale factor, this can be written as $PX = -\theta P'X$ for some constant θ . From this it follows that $(P + \theta P')X = 0$. Conversely, suppose that $(P + \theta P')X = 0$, for some θ . It follows that $PX = -\theta P'X$, and so $PX \approx P'X$. Thus, the critical set is the set of points X in the right null-space of $P + \theta P'$ for some θ .

Define $P_\theta = P + \theta P'$. The rest of the proof involves finding the set of all points X satisfying $P_\theta X = 0$ as θ runs over all values. If P_θ is a camera matrix (has rank 3) then such an X is the centre of the camera P_θ . If however P_θ is rank-deficient for some value of θ_i , then the set of points X such that $P_{\theta_i} X = 0$ is a linear space. The total critical set therefore consists of two parts:

- (i) The locus of the camera centre of P_θ over all values θ for which P_θ has full rank (that is rank 3). This is a curve $\mathcal{C}$ containing the two camera centres C_0 and C_1 .
- (ii) A linear space (line or plane) corresponding to each value of θ for which P_θ has rank 2 or less. If P_θ has rank 2, the points such that $P_\theta \mathbf{X} = \mathbf{0}$ form a line, and if it has rank 1, then they form a plane.

Let the 4-vector C_θ be defined by $C_\theta = (c_1, c_2, c_3, c_4)^T$, where $c_i = (-1)^i \det P_\theta^{(i)}$ and $P_\theta^{(i)}$ is the matrix P_θ with the i -th column removed. Since each $P_\theta^{(i)}$ is a 3×3 matrix, and the entries of P_θ are linear in θ , we see that each $c_i = (-1)^i \det P_\theta^{(i)}$ is a cubic polynomial in θ . According to the discussion following definition 21.10(p521), $P_\theta C_\theta = \mathbf{0}$, hence if $C_\theta \neq \mathbf{0}$ then it is the camera centre of P_θ , and as θ varies the point C_θ traces out the curve $\mathcal{C}$. Since the coordinates of C_θ are cubic polynomials, this is in general a twisted cubic. If however the four components of C_θ have a simultaneous root θ_i , then the degree of the curve $\mathcal{C}$ is diminished. In this case P_{θ_i} is rank-deficient, and there exists a linear space of points $\mathbf{X}$ such that $P_{\theta_i} \mathbf{X} = \mathbf{0}$. Thus the linear subspace components of the critical set correspond to values of θ where C_θ vanishes. Clearly there can be at most three such values. Further details are left to the reader.

The last part of the theorem provides a converse result – namely that if the points and one camera centre lie in a critical configuration, then there exist alternative resection solutions with the camera placed at any location in $\mathcal{C}$. The exact form of the camera matrix P is not important here, only its camera centre, as has been observed above. For most of the configurations in figure 22.3 it is clear enough geometrically that the image of the critical set is unchanged (up to projectivity) by moving the camera along the locus $\mathcal{C}$. In the case where $\mathcal{C}$ is a planar conic, this follows fairly easily from the 1D camera case (theorem 22.3). The exception is the twisted cubic case. It is illustrated graphically in figure 22.4. We leave this proof for now, and come back to it later (result 22.25(p551)). $\square$

22.2 Degeneracies in two views

Notation. For the rest of this chapter, the camera matrices are represented by P and Q , 3D points by $\mathbf{P}$ and $\mathbf{Q}$. Thus cameras and 3D points are distinguished only by their type-face. This may appear to be a little confusing, but the alternative of using subscripts or primes proved to be much more confusing. In the context of ambiguous reconstructions from image coordinates we distinguish the two reconstructions by using P and $\mathbf{P}$ for one, and Q and $\mathbf{Q}$ for the other.

Now we turn to the case of two views of an object. Given sufficiently many points placed in “general position” one may determine the placement of the two cameras, and reconstruct the point set up to a projective transformation. This may be done using one of the projective reconstruction algorithms discussed in chapter 10. We now wish to determine under what conditions this technique may fail.

Thus, we consider a set of image correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$. A realization of this set of correspondences consists of a pair of camera matrices P and P' and a set of 3D points $\mathbf{P}_i$ such that $\mathbf{x}_i = P\mathbf{P}_i$ and $\mathbf{x}'_i = P'\mathbf{P}_i$ for all i . Two realizations $\{P, P', \mathbf{P}_i\}$




Fig. 22.4. Separate views of a set of points on the twisted cubic $(t^3, t^2, t, 1)^T$ as viewed from a centre of projection. The visible points are viewed from the points with $t = 3, 4, 5, 10, 50, 1000$ with the image suitably magnified to prevent the point set from becoming too small. As is plausible from the image, the sets of points differ by a projective transformation. From a viewpoint on a twisted cubic, the twisted cubic has the appearance of a conic, in this particular case a parabola.

and $\{Q, Q', Q_i\}$ are projectively equivalent if there is a 3D projective transformation, represented by a matrix H such that $Q = PH^{-1}$, $Q' = P'H^{-1}$, and $Q_i = HP_i$ for all i .

Because of a technicality, this definition of equivalence is not quite appropriate to the present discussion. Recall from the projective reconstruction theorem, theorem 10.1-(p266), that one cannot determine the position of a point lying on the line joining the two camera centres. Hence, non-projectively-equivalent reconstructions will always exist if some points lie on the line of camera centres, the two reconstructions differing only by the position of the points P_i and Q_i on this line. This type of reconstruction ambiguity is not of great interest, and so we will modify the notion of equivalence by defining two reconstructions to be equivalent if H exists such that $Q = PH^{-1}$ and $Q' = P'H^{-1}$. As in the proof of the projective reconstruction theorem, such an H will also map Q_i to P_i , except possibly for reconstructed points Q_i lying on the line of the camera centres. According to theorem 9.10(p254), this condition is also equivalent to the condition that $F_{P'P} = F_{Q'Q}$, where $F_{P'P}$ and $F_{Q'Q}$ are the fundamental matrices corresponding to the camera pairs (P, P') and (Q, Q') . Accordingly, we make the following definition.

Definition 22.7. Two configurations of cameras and points $\{P, P', P_i\}$ and $\{Q, Q', Q_i\}$ are said to be *conjugate configurations* provided

- (i) $PP_i = QQ_i$ and $P'P_i = Q'Q_i$ for all i ,
- (ii) The fundamental matrices $F_{Q'Q}$ and $F_{P'P}$ corresponding to the two camera matrix pairs (P, P') and (Q, Q') are different.

A configuration $\{P, P', P_i\}$ that allows a conjugate configuration is called *critical*.

An important remark is that being a critical configuration depends only on the camera centres and the points, and not on the particular cameras.

Result 22.8. *If $\{P, P', P_i\}$ is a critical configuration and $\hat{P}$ and $\hat{P}'$ are two cameras with the same centres as P and P' respectively, then $\{\hat{P}, \hat{P}', P_i\}$ is a critical configuration as well.*

Proof. This is easily seen as follows. Since $\{P, P', P_i\}$ is a critical configuration there exists an alternative configuration $\{Q, Q', Q_i\}$ such that $PP_i = QQ_i$ and $P'P_i = Q'Q_i$ for all i . However, since P and $\hat{P}$ have the same camera centre, $\hat{P} = HP$ according to result 22.1 and similarly $\hat{P}' = H'P'$. Therefore

$$\hat{P}P_i = HPP_i = HQQ_i \text{ and}$$

$$\hat{P}'P_i = H'P'P_i = H'Q'Q_i.$$

It follows that $\{HQ, H'Q', Q_i\}$ is an alternative configuration to $\{\hat{P}, \hat{P}', P_i\}$, which is therefore critical. $\square$

The goal of this section is to determine under what conditions non-equivalent realizations of a set of point correspondences may occur. This question is completely resolved by the following theorem, which will be proved incrementally.

Theorem 22.9. (i) *Conjugate configurations of cameras and points generically come in triples. Thus a critical configuration $\{P, P', P_i\}$ has two conjugates $\{Q, Q', Q_i\}$ and $\{R, R', R_i\}$.*
 (ii) *If $\{P, P', P_i\}$ is a critical configuration, then all the points P_i and the two camera centres C_P and $C_{P'}$ lie on a ruled quadric surface S_P .*
 (iii) *Conversely, suppose that the camera centres of P, P' and a set of 3D points P_i lie on a ruled quadric (excluding the quadrics (v) and (viii) in result 22.11), then $\{P, P', P_i\}$ is a critical configuration.*

By the words “ruled quadric” in this context is meant any quadric surface that contains a straight line. This includes various degenerate quadrics, as will be seen. A general discussion and classification of quadric surfaces has been given in section 3.2.4 (p74). A *quadric* is usually defined to be the set of points X such that $X^T S X = 0$, where S is a symmetric 4×4 matrix. However, if S is any 4×4 matrix, not necessarily symmetric, then one sees that for any point X , one has $X^T S X = (X^T S_{\text{sym}} X)$, where $S_{\text{sym}} = (S + S^T)/2$ is the symmetric part of S . Thus, $X^T S X = 0$ if and only if $X^T S_{\text{sym}} X = 0$, and so S defines the same quadric surface as S_{sym} . In investigating reconstruction ambiguity, it will often be convenient to use non-symmetric matrices S to represent quadrics.

Proof. We begin by proving the first part of the theorem. Let F and F' be two distinct fundamental matrices satisfying the relation $x_i'^T F x_i = x_i'^T F' x_i = 0$ for all correspondences $x_i' \leftrightarrow x_i$. Define $F_\theta = F + \theta F'$. One easily verifies that $x_i'^T F_\theta x_i = 0$. However, F_θ is a fundamental matrix only if $\det F_\theta = 0$. Now, $\det F(\theta)$ is generally a polynomial of degree 3 in θ . This polynomial has roots $\theta = 0$ and $\theta = 1$ corresponding to F and

F' respectively. The third root corresponds to a third fundamental matrix, and hence a third non-equivalent reconstruction. In special cases, $\det F(\theta)$ has degree 2 in θ and there are only two conjugate configurations. $\square$

The second part of the theorem is concluded by proving the following lemma.

Lemma 22.10. *Consider two pairs of cameras (P, P') and (Q, Q') , with corresponding different fundamental matrices $F_{P'P}$ and $F_{Q'Q}$. Define quadrics $S_P = P'^T F_{Q'Q} P$, and $S_Q = Q'^T F_{P'P} Q$.*

- (i) *The quadric S_P contains the camera centres of P and P' . Similarly, S_Q contains the camera centres of Q and Q' .*
- (ii) *S_P is a ruled quadric.*
- (iii) *If P and Q are 3D points such that $PP = QQ$ and $P'P = Q'Q$, then P lies on the quadric S_P , and Q lies on S_Q .*
- (iv) *Conversely, if P is a point lying on the quadric S_P , then there exists a point Q lying on S_Q such that $PP = QQ$ and $P'P = Q'Q$.*

Proof. Recall that according to result 9.12(p255) the matrix F is the fundamental matrix corresponding to a pair of cameras (P, P') if and only if $P'^T F P$ is skew-symmetric. Since $F_{P'P} \neq F_{Q'Q}$, however, the matrices S_P and S_Q defined here are not skew-symmetric, and hence represent non-trivial quadrics.

We adopt a convention that a camera represented by a matrix such as P or P' has camera centre denoted by C_P or $C_{P'}$.

- (i) The camera centre of P satisfies $PC_P = 0$. Then

$$C_P^T S_P C_P = C_P^T (P'^T F_{Q'Q} P) C_P = C_P^T (P'^T F_{Q'Q}) PC_P = 0$$

since $PC_P = 0$. So, C_P lies on the quadric S_P . In a similar manner, $C_{P'}$ lies on S_P .

- (ii) Let e_Q be the epipole defined by $F_{Q'Q} e_Q = 0$, and consider the ray passing through C_P consisting of all points such that $e_Q = PX$. Then for any such point, one verifies that $S_P X = P'^T F_{Q'Q} P X = P'^T F_{Q'Q} e_Q = 0$. Thus the ray lies on S_P and so S_P is a ruled quadric.
- (iii) Under the given conditions one sees that

$$P^T S_P P = P^T P'^T F_{Q'Q} P P = Q'^T (Q'^T F_{Q'Q} Q) Q = 0$$

since $Q'^T F_{Q'Q} Q$ is skew-symmetric. Thus, P lies on the quadric S_P . By a similar argument, Q lies on S_Q .

- (iv) Let P lie on S_P and define $x = PP$ and $x' = P'P$. Then, from $P^T S_P P = 0$ we deduce $0 = P^T P'^T F_{Q'Q} P P = x'^T F_{Q'Q} x$, and so $x \leftrightarrow x'$ are a corresponding pair of points with respect to $F_{Q'Q}$. Therefore, there exists a point Q such that $QQ = x = PP$, and $Q'Q = x' = P'P$. From part (iii) of this lemma, Q must lie on S_Q .

$\square$

This lemma completely describes the sets of 3D points giving rise to ambiguous image correspondences. Note that any two arbitrarily chosen camera pairs can give rise to ambiguous image correspondences, provided that the world points lie on the given quadrics.

22.2.1 Examples of ambiguity

At this point it remains to prove the converse of theorem 22.9. This needs to be done for all types of ruled quadrics and any placement of the two camera centres on the quadric. The different types of ruled quadrics, including degenerate cases, are (see section 3.2.4(p74)): a hyperboloid of one sheet; a cone; two (intersecting) planes; a single plane; a single line. A complete enumeration of the types of placement of two camera centres on a ruled quadric is given in the following result.

Result 22.11. *The possible configurations of two distinct points (particularly camera centres) on a ruled quadric surface are as enumerated:*

- (i) *hyperboloid of one sheet, two points not on the same generator*
- (ii) *hyperboloid of one sheet, two points on the same generator*
- (iii) *cone, one point at the vertex and one not*
- (iv) *cone, two points on different generators, neither one at the vertex*
- (v) *cone, two points on the same generator, neither one at the vertex*
- (vi) *pair of planes, two points on the intersection of the two planes*
- (vii) *pair of planes, one point at the intersection and one not*
- (viii) *pair of planes, neither point at the intersection, but points on different planes*
- (ix) *pair of planes, neither point on the intersection, both points on the same plane*
- (x) *single plane with two points*
- (xi) *single line with two points*

Any two quadric/point-pairs in the same class are projectively equivalent.

It is clear by enumeration of cases that this list is complete. The only case in which it is not immediately obvious that any two configurations in the same category are projectively equivalent is in the non-degenerate case of the hyperboloid of one sheet. The proof of this fact is left as exercise (i) at the end of the chapter (page 559).

Now, consider an example of a critical configuration $\{P, P', \hat{P}_i\}$ in which all P_i lie on a quadric S_P also containing the two camera centres. The quadric and two camera centres belong to one of the categories given in result 22.11.

Let the centres of two new cameras $(\hat{P}, \hat{P}')$ and a set of points $\hat{P}_i$ lie on a quadric $\hat{S}_P$. Suppose that $\hat{S}_P$ and the two camera centres lie in the same category as the given example. Since two configurations in the same category are projectively equivalent, we may assume that $\hat{S}_P = S_P$ and the centres of P and $\hat{P}$ are the same, as are the centres of P' and $\hat{P}'$. Since the points $\hat{P}_i$ lie on S_P , it follows that $\{P, P', \hat{P}_i\}$ is a critical configuration, and hence by result 22.8 so is $\{\hat{P}, \hat{P}', \hat{P}_i\}$.

This shows that to demonstrate the converse section of theorem 22.9 it is sufficient merely to demonstrate an example of a critical configuration belonging to each of the categories given in result 22.11 (except for cases (v) and (viii)). Examples will be given

for several of the categories, though not all. The fact that cases (v) and (viii) are not critical is not shown here. The remaining cases are left to the motivated reader.

Examples of critical configurations

Consider the case $P = Q = [I \mid 0]$ and $P' = [I \mid t_0]$. In this case, one sees that

$$S_P = \begin{bmatrix} I \\ t_0^T \end{bmatrix} F_{Q'Q} [I \mid 0] = \begin{bmatrix} I \\ t_0^T \end{bmatrix} [F_{Q'Q} \mid 0] .$$

Consequently,

$$S_{P_{\text{sym}}} = \frac{1}{2} \begin{bmatrix} F_{Q'Q} + F_{Q'Q}^T & F_{Q'Q}^T t_0 \\ t_0^T F_{Q'Q} & 0 \end{bmatrix} .$$

Given the fundamental matrix $F_{Q'Q}$ of rank 2, and the camera matrix $Q = [I \mid 0]$, one may easily find the other camera matrix Q' . This is done by using the formula of result 9.13(p256). We now consider several examples of critical configurations belonging to different categories of result 22.11.

Example 22.12. Hyperboloid of one sheet – two centres not on the same generator. We

choose $F_{Q'Q} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & -1 \end{bmatrix}$ and $t_0 = (0, 2, 0)^T$. Then $t_0^T F_{Q'Q} = (0, 2, 0)$ and we find that

$$S_{\text{sym}} = \begin{bmatrix} 1 & & & \\ & 1 & & 1 \\ & & -1 & \\ & 1 & & 0 \end{bmatrix} .$$

This is the quadric with equation $x^2 + y^2 + 2y - z^2 = 0$, or $x^2 + (y+1)^2 - z^2 = 1$. This is a hyperboloid of one sheet. Note that in this case, the camera centres are $C_P = C_Q = (0, 0, 0)^T$ (in inhomogeneous coordinates). The camera centre $C_{P'} = (0, 2, 0)$. We note that the line from C_P to $C_{P'}$ does not lie on the quadric surface. The camera centre $C_{Q'}$ is not uniquely determined by the information already given, since the fundamental matrix does not uniquely determine the two cameras. However, the epipole e such that $F_{Q'Q}e = 0$ is $e = (1, 0, -1)$. Since $e = QC_{Q'}$, we see that $C_{Q'} = (1, 0, -1, k^{-1})$. In inhomogeneous coordinates $C_{Q'} = k(1, 0, -1)$ for any k . We verify that the complete line from C_P to $C_{Q'}$ lies on the quadric, but $C_{Q'}$ may lie anywhere along this line. $\triangle$

Example 22.13. Hyperboloid of one sheet, both centres on the same generator. We choose the same $F_{Q'Q}$ as in the previous example, and $t_0 = (3, 4, 5)^T$. Then $t_0^T F_{Q'Q} = (-2, 4, -2)$ and we find that

$$S_{\text{sym}} = \begin{bmatrix} 1 & & & -1 \\ & 1 & & 2 \\ & & -1 & -1 \\ -1 & 2 & -1 & 0 \end{bmatrix} .$$

This is the quadric with equation $(x-1)^2 + (y+2)^2 - (z+1)^2 = 4$, which is once again

a hyperboloid of one sheet. It may be verified that the line $(x, y, z) = (3t, 4t, 5t)$ lies entirely on the quadric and contains the two camera centres, $(0, 0, 0)^T$ and $(3, 4, 5)^T$. $\triangle$

Example 22.14. Cone – one centre at the vertex of the cone. We choose $F_{Q/Q}$ to be the same as in the previous example, but $t_0 = (1, 0, 1)^T$. In this case we see that $t_0^T F_{Q/Q} = 0^T$, and so $S_{\text{sym}} = \text{diag}(1, 1, -1, 0)$. This is an example of a cone. The camera centre of both P and Q is at the vertex of the cone, namely the point $C_P = (0, 0, 0, 1)^T$. $\triangle$

Example 22.15. Cone – neither centre at the vertex of the cone. We choose $F_{Q/Q} = \text{diag}(1, 1, 0)$ and $t_0 = (0, 2, 0)^T$. In this case, we see that

$$S_{\text{sym}} = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 0 & \\ & & & 1 \end{bmatrix}$$

This is the quadric with equation $x^2 + y^2 + 2y = 0$, or $x^2 + (y + 1)^2 = 1$, which is a cylinder parallel with the Z-axis, thus projectively equivalent to the cone with vertex at the infinite point $(0, 0, 1, 0)^T$. Neither of the camera centres lies at the vertex. $\triangle$

Example 22.16. Two planes. We choose $F_{Q/Q} = \text{diag}(1, -1, 0)$ and $t_0 = (0, 0, 1)^T$, so that $t_0^T F_{Q/Q} = 0^T$. In this case, we see that $S_P = S_{\text{sym}} = \text{diag}(1, -1, 0, 0)$, which represents a pair of planes. In this case the camera centres are on the intersection of the two planes $\triangle$

Example 22.17. Single plane. We choose $F_{Q/Q} = \begin{bmatrix} 1 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $t_0 = (0, 0, 1)^T$. Then $S_{\text{sym}} = \text{diag}(1, 0, 0, 0)$, which represents a single plane, $x = 0$. $\triangle$

Example 22.18. Single line. We choose $F_{Q/Q} = \text{diag}(1, 1, 0)$ and $t_0 = (0, 0, 1)^T$. In this case, we see that $S_{\text{sym}} = \text{diag}(1, 1, 0, 0)$, which represents a single line – the Z-axis. All the points and the two camera centres lie on this single line. $\triangle$

Apart for the impossible cases (v) and (viii) this gives examples of all possible degenerate configurations except for cases (vii) and (ix) in result 22.11. These remaining cases are left for the reader.

Minimal case – seven points As seen in section 11.1.2(p281), seven is the minimum number of points in general position from which one can do projective reconstruction. The method comes down to solving a cubic equation, for which either one or three real solutions exist. Looked at from the point of view of critical surfaces, the seven points and two camera centres in one configuration must lie on a quadric surface (since 9 points lie on a quadric). If this quadric is ruled, then there will be three conjugate solutions. On the other hand, if it is not a ruled quadric (for instance an ellipsoid) then there will be only one solution. This shows that the distinction between the cases where the cubic equation has one or three solutions arises from the difference between the

points and camera centres lying on a ruled or unruled quadric – a pleasing connection between the algebra and geometry.

22.3 Carlsson–Weinshall duality

The duality explored in chapter 20 between cameras and points may be exploited so as to dualize degeneracy results, as will be explained in this section. We give a more formal treatment of Carlsson–Weinshall duality here.

The basis of Carlsson–Weinshall duality is the equation

$$\begin{bmatrix} a & & -d \\ & b & -d \\ & & c & -d \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ T \end{pmatrix} = \begin{bmatrix} X & & -T \\ & Y & -T \\ & & Z & -T \end{bmatrix} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix}.$$

The camera matrix on the left corresponds to a camera with centre $(a^{-1}, b^{-1}, c^{-1}, d^{-1})^T$. We are interested in describing cameras in terms of their camera centres. As this shows, in swapping camera centres with 3D points, one must invert the coordinates of the camera centre and point. For instance, the point $(X, Y, Z, T)^T$ is dual to a camera with camera centre $(X^{-1}, Y^{-1}, Z^{-1}, T^{-1})^T$.

We denote the reduced camera matrix

$$P = \begin{bmatrix} a^{-1} & & -d^{-1} \\ & b^{-1} & -d^{-1} \\ & & c^{-1} & -d^{-1} \end{bmatrix} \tag{22.1}$$

with centre $C = (a, b, c, d)^T$ by P_C . Now, defining the points $\bar{X} = (X^{-1}, Y^{-1}, Z^{-1}, T^{-1})^T$ and $\bar{C} = (a^{-1}, b^{-1}, c^{-1}, d^{-1})^T$, one immediately verifies that

$$P_C X = P_{\bar{X}} \bar{C} \tag{22.2}$$

Thus, this transformation interchanges the results of 3D points and camera centres. Thus, a camera with centre C acting on X gives the same result as camera with centre $\bar{X}$ acting on $\bar{C}$.

This observation leads to the following definition

Definition 22.19. The mapping of $\mathbb{P}^3$ to itself given by

$$(X, Y, Z, T)^T \mapsto (YZT, ZTX, TXY, XYZ)^T$$

will be called the Carlsson–Weinshall map, and will be denoted by Γ . We denote the image of a point X under Γ by $\bar{X}$. The image of an object under Γ is sometimes referred to as the *dual* object.

The Carlsson–Weinshall map is an example of a *Cremona* transformation. For more information on Cremona transformations, the reader is referred to Semple and Kneebone ([Semple-79]).

Note. If none of the coordinates of X is zero then we may divide $\bar{X}$ by $XYZT$. Then Γ is equivalent to $(X, Y, Z, T)^T \mapsto (X^{-1}, Y^{-1}, Z^{-1}, T^{-1})^T$. This is the form of the mapping

that we will usually use. In the case where one of the coordinates of $\mathbf{X}$ is zero, then the mapping will be interpreted as in the definition. Note that any point $(0, Y, Z, T)^T$ is mapped to the point $(1, 0, 0, 0)^T$ by Γ , provided none of the other coordinates is zero. Thus, the mapping is not one-to-one.

If two of the coordinates of $\mathbf{X}$ are zero, then $\bar{\mathbf{X}} = (0, 0, 0, 0)^T$, which is an undefined point. Thus, Γ is not defined at all points. In fact, there is no way to extend Γ continuously to such points.

Define the *reference tetrahedron* to be the tetrahedron with vertices $\mathbf{E}_1 = (1, 0, 0, 0)^T$, $\mathbf{E}_2 = (0, 1, 0, 0)^T$, $\mathbf{E}_3 = (0, 0, 1, 0)^T$ and $\mathbf{E}_4 = (0, 0, 0, 1)^T$. As we have just seen, Γ is one-to-one other than on the faces of the reference tetrahedron. It maps a face of the reference tetrahedron to the opposite vertex, and is undefined on the edges of the reference tetrahedron. Next, we investigate the way in which Γ acts on other geometric objects.

Theorem 22.20. *The Carlsson–Weinshall map, Γ acts in the following manner:*

- (i) *It maps a line passing through two general points $\mathbf{X}_0$ and $\mathbf{X}_1$ to the twisted cubic passing through $\bar{\mathbf{X}}_0, \bar{\mathbf{X}}_1$ and the four reference vertices $\mathbf{E}_1, \dots, \mathbf{E}_4$.*
- (ii) *It maps a line passing through any of the points $\mathbf{E}_i$ to a line passing through the same $\mathbf{E}_i$. We exclude the lines lying on the face of the reference tetrahedron, since such lines will be mapped to a single point.*
- (iii) *It maps a quadric $\mathbf{S}$ passing through the four points $\mathbf{E}_i$, $i = 1, \dots, 4$ to a quadric (denoted $\bar{\mathbf{S}}$) passing through the same four points. If $\mathbf{S}$ is a ruled quadric, then so is $\bar{\mathbf{S}}$. If $\mathbf{S}$ is degenerate then so is $\bar{\mathbf{S}}$.*

Proof.

(i) A line has parametric equation $(\mathbf{X}_0 + a\theta, \mathbf{Y}_0 + b\theta, \mathbf{Z}_0 + c\theta, \mathbf{T}_0 + d\theta)^T$, and a point on this line is taken by the Carlsson–Weinshall map to the point

$$((Y_0 + b\theta)(Z_0 + c\theta)(T_0 + d\theta), \dots, (X_0 + a\theta)(Y_0 + b\theta)(Z_0 + c\theta))^T.$$

Thus, the entries of the vector are cubic functions of θ , and the curve is a twisted cubic. Now, setting $\theta = -X_0/a$, the term $(X_0 + a\theta)$ vanishes, and the corresponding dual point is $((Y_0 + b\theta)(Z_0 + c\theta)(T_0 + d\theta), 0, 0, 0)^T \approx (1, 0, 0, 0)^T$. The first entry is the only one that does not contain $(X_0 + a\theta)$, and hence the only one that does not vanish. This shows that the reference vertex $\mathbf{E}_1 = (1, 0, 0, 0)^T$ is on the twisted cubic. By similar arguments, the other points $\mathbf{E}_2, \dots, \mathbf{E}_4$ lie on the twisted cubic also. Note that a twisted cubic is defined by six points, and this twisted cubic is defined by the given six points $\mathbf{E}_i, \bar{\mathbf{X}}_0, \bar{\mathbf{X}}_1$ that lie on it, where $\mathbf{X}_0$ and $\mathbf{X}_1$ are any two points defining the line.

(ii) We prove this for lines passing through the point $\mathbf{E}_0 = (1, 0, 0, 0)^T$. An analogous proof holds for the other points $\mathbf{E}_i$. Choose another point $\mathbf{X} = (x, y, z, t)^T$ on the line, such that $\mathbf{X}$ does not lie on any face of the reference tetrahedron. Thus $\mathbf{X}$ has no zero coordinate. Points on a line passing through $(1, 0, 0, 0)^T$ and $\mathbf{X} = (x, y, z, t)^T$ are all of the form $(\alpha, y, z, t)^T$ for varying values of α . These points are mapped by the transformation to $(\alpha^{-1}, y^{-1}, z^{-1}, t^{-1})^T$. This represents a line passing through the two points $(1, 0, 0, 0)^T$ and $\bar{\mathbf{X}} = (x^{-1}, y^{-1}, z^{-1}, t^{-1})^T$.

(iii) Since the quadric S passes through all the points E_i , the diagonal entries of S must all be zero. This means that there are no terms involving a squared coordinate (such as x^2) in the equation for the quadric. Hence the equation for the quadric contains only mixed terms (such as XY , YZ or XT). Therefore, the quadric S may be defined by an equation $aXY + bXZ + cXT + dYZ + eYT + fZT = 0$. Dividing this equation by $XYZT$, we obtain $aZ^{-1}T^{-1} + bY^{-1}T^{-1} + cY^{-1}Z^{-1} + dX^{-1}T^{-1} + eX^{-1}Z^{-1} + fX^{-1}Y^{-1} = 0$. Since $\bar{X} = (X^{-1}, Y^{-1}, Z^{-1}, T^{-1})^T$, this is a quadratic equation in the entries of $\bar{X}$. Thus Γ maps quadric to quadric. Specifically, suppose S is represented by the matrix

$$S = \begin{bmatrix} 0 & a & b & c \\ a & 0 & d & e \\ b & d & 0 & f \\ c & e & f & 0 \end{bmatrix} \quad \text{then } \bar{S} = \begin{bmatrix} 0 & f & e & d \\ f & 0 & c & b \\ e & c & 0 & a \\ d & b & a & 0 \end{bmatrix}$$

and $X^T S X = 0$ implies $\bar{X}^T \bar{S} \bar{X} = 0$. If S is ruled, then it contains two generators passing through any point, and in particular through each E_i . By part (ii), these are mapped to straight lines, which must lie on $\bar{S}$. Thus $\bar{S}$ is a ruled quadric. One may further verify that $\det S = \det \bar{S}$, which implies that if S is a non-degenerate quadric (that is $\det S \neq 0$), then so is $\bar{S}$. In this non-degenerate case, if S is a hyperboloid of one sheet, then $\det S > 0$, from which it follows that $\det \bar{S} > 0$. Thus $\bar{S}$ is also a hyperboloid of one sheet. □

The action of Γ on other geometric entities is investigated in exercises (page 559); We wish to interpret duality equation (22.2) in a coordinate-free manner. The matrix P_C has by definition the form given in (22.1), and maps E_i to e_i for $i = 1, \dots, 4$. The image $P_C X$ may be thought of as a representation of the projection of X relative to the projective basis e_i in the image. Alternatively, $P_C X$ represents the projective equivalence class of the set of the five rays $\overline{CE_1}, \dots, \overline{CE_4}, \overline{CX}$. Thus $P_C X = P_{C'} X'$ if and only if the set of rays from C to X and the four vertices of the reference tetrahedron is projectively equivalent to the set of rays from C' to X' and the four reference vertices. In terms of the notation introduced earlier, we may write (22.2) in a different form as

$$\langle C; E_1, \dots, E_4, X \rangle = \langle \bar{X}; E_1, \dots, E_4, \bar{C} \rangle. \tag{22.3}$$

The duality principle

The basis of duality is (22.2) which states that $P_C X = P_{\bar{X}} \bar{C}$, with notation as in (22.2) The notation $P_C X$ represents the coordinates of the projection of point X with respect to the canonical image coordinate frame defined by the projections of the corners of the reference tetrahedron. Equivalently, P_C may be considered as representing the projective equivalence class of the five projected points $P_C E_i$ and $P_C X$. In the notation of this chapter, this is $\langle C; E_1, \dots, E_4, X \rangle$. Thus the duality relation may be written as

$$\langle C; E_1, \dots, E_4, X \rangle = \langle \bar{X}; E_1, \dots, E_4, \bar{C} \rangle \tag{22.4}$$

where the bar represents the Carlsson–Weinshall map.

Although P_C was defined in terms of the canonical projective basis, there is nothing special about the four points $E_1, \dots, E_4$ used as vertices of the reference tetrahedron,

other than the fact that they are non-coplanar. Given any four non-coplanar points, one may define a projective coordinate system in which these four points are the points E_i forming part of a projective basis. The Carlsson–Weinshall map may then be defined with respect to this coordinate frame. The resulting map is called the Carlsson–Weinshall map with respect to the given reference tetrahedron.

To be more precise, it should be observed that five points (not four) define a projective coordinate frame in $\mathbb{P}^3$. In fact, there is more than one projective frame (in fact a 3-parameter family) for which four non-coplanar points have coordinates E_i . Thus the Carlsson–Weinshall map with respect to a given reference tetrahedron is not unique. However, the mapping given by definition 22.19 with respect to any such coordinate frame may be used.

Given a statement or theorem concerning projections of sets of points with respect to one or more projection centres one may derive a dual statement. One requires that among the four points being projected, there are four non-coplanar points that may form a reference tetrahedron. Under a general duality mapping with respect to the reference tetrahedron

- (i) Points (other than those belonging to the reference tetrahedron) are mapped to centres of projection.
- (ii) Centres of projection are mapped to points.
- (iii) Straight lines are mapped to twisted cubics.
- (iv) Ruled quadrics containing the reference tetrahedron are mapped to ruled quadrics containing the reference tetrahedron. If the original quadric is non-degenerate, so is its image under the duality mapping.

Points lying on an edge of the reference tetrahedron should be avoided, since the Carlsson–Weinshall mapping is undefined for such points. Using this as a sort of translation table, one may dualize existing theorems about point projection, giving new theorems for which a separate proof is not needed.

Note. It is important to observe that only those points not belonging to the reference tetrahedron are mapped to camera centres by duality. The vertices of the reference tetrahedron remain points. In practice, in applying the duality principle, one may select any four points to form the reference tetrahedron, as long as they are non-coplanar. In general, in the results stated in the next section there will be an assumption (not always stated explicitly) that *point sets considered contain four non-coplanar points*, which may be taken as the reference tetrahedron.

22.3.1 Single view ambiguity

It will be shown in this section how various ambiguous reconstruction results may be derived simply from known or obvious geometric statements by applying duality.

Camera resectioning from five points

Five 3D–2D point correspondences are not sufficient for camera resectioning for projective cameras. It is interesting to know what one can determine from five point correspondences, however.




Fig. 22.5. **Left:** Any point on the line passing through C and X is projected to the same point from projection centre C . **Right:** The dual statement – from any centre of projection C lying on a twisted cubic passing through X and the vertices of the reference tetrahedron, the five points are projected in the same way (up to projective equivalence). Thus a camera is constrained to lie on a twisted cubic by its image of five known points.

As a simple example of what can be deduced using Carlsson duality, consider the following simple question: when do two points project to the same point in an image? The answer is obviously when the two points lie on the same ray (straight line) through the camera centre. Dualizing this simple observation, figure 22.5 shows that the centre of the camera constrained by five point 3D–2D correspondences must lie on a twisted cubic passing through the five 3D points.

The horopter

In a similar manner one may compute the form of the horopter determined by two cameras. The horopter is the set of space points that map to the same point in two images. The argument is illustrated in figure 22.6 and begins with a simple observation concerning straight lines.

Result 22.21. *Given points X and X' , the locus of camera centres C such that*

$$\langle C; E_1, \dots, E_4, X \rangle = \langle C; E_1, \dots, E_4, X' \rangle$$

is the straight line passing through X and X' .

This is illustrated in figure 22.6(left) The dual of this statement determines the horopter for a pair of cameras (see figure 22.6(right)).

Result 22.22. *Given projection centres C and C' , non-collinear with the four points E_i of a reference tetrahedron, the set of points X such that $\langle C; E_1, \dots, E_4, X \rangle = \langle C'; E_1, \dots, E_4, X \rangle$ is a twisted cubic passing through $E_1, \dots, E_4$ and the two projection centres C and C' .*

Ambiguity of camera resection

Finally, we consider ambiguity of resection. This is very closely related to the horopter. To visualize this, the reader may refer again to figure 22.6, though it is not exactly pertinent in this situation.

Result 22.23. *Consider a set of camera centres $C_1, \dots, C_m$ and a point X_0 all lying on*




Fig. 22.6. **Left:** From any centre of projection $C, C', \dots$ lying on the line passing through X and X' , the points X and X' are projected to the same ray. That is, $\langle C; E_i, X \rangle = \langle C; E_i, X' \rangle$ for all C on the line. **Right:** The dual statement – all points on the twisted cubic passing through C and C' and the vertices of the reference tetrahedron are projected in the same way relative to the two projection centres. That is, $\langle C; E_i, X \rangle = \langle C'; E_i, X \rangle$ for all X on the twisted cubic. This curve is called the horopter for the two centres of projection.

a single straight line and let $E_i, i = 1, \dots, 4$ be the vertices of a reference tetrahedron. Let X be another point. Then the two configurations

$$\{C_1, \dots, C_m; E_1, \dots, E_4, X\} \text{ and } \{C_1, \dots, C_m; E_1, \dots, E_4, X_0\}$$

are image-equivalent configurations if and only if X lies on the same straight line as X_0 and the cameras.

In passing to the dual statement, according to theorem 22.20 the straight line becomes a twisted cubic through the four vertices of the reference tetrahedron. Thus the dual statement to result 22.23 is:

Result 22.24. Consider a set of points X_i and a camera centre C_0 all lying on a single twisted cubic also passing through four reference vertices E_i . Let C be any other camera centre. Then the configurations

$$\{C; E_1, \dots, E_4, X_1, \dots, X_m\} \text{ and } \{C_0; E_1, \dots, E_4, X_1, \dots, X_m\}$$

are image-equivalent if and only if C lies on the same twisted cubic.

Since the points E_i may be any four non-coplanar points, and a twisted cubic cannot contain 4 coplanar points, one may state this last result in the following form:

Result 22.25. Let $X_1, \dots, X_m$ be a set of points and C_0 a camera centre all lying on a twisted cubic. Then for any other camera centre C the configurations

$$\{C; X_1, \dots, X_m\} \text{ and } \{C_0; X_1, \dots, X_m\}$$

are image-equivalent if and only if C lies on the same twisted cubic.

This is illustrated in figure 22.6 (right). It shows that camera pose cannot be uniquely determined whenever all points and a camera centre lie on a twisted cubic. This gives an independent proof of result 22.6(p537) covering the case that was left unfinished previously.

Using similar methods one can show that this is one of only two possible ambiguous situations. The other case in which ambiguity occurs is when all points and the two

camera centres lie in the union of a plane and a line. This arises as the dual of the case when the straight line through the camera centres meets one of the vertices of the reference tetrahedron. In this case, the dual of this line is also a straight line through the same reference vertex (see theorem 22.20), and all points must lie on this line or the opposite face of the reference tetrahedron.

Note in both these examples how the use of duality has taken intuitively obvious statements concerning projections of collinear points and derived a result somewhat less obvious about points lying on a twisted cubic.

22.3.2 Two-view ambiguity

The basic result theorem 22.9(p541) about critical surfaces from two views may be stated as follows.

Theorem 22.26. *A configuration $\{C_1, C_2; X_1, \dots, X_n\}$ of two camera centres and n points allows an alternative reconstruction if and only if both camera centres C_1, C_2 and all the points X_j lie on a ruled quadric surface. If the quadric is non-degenerate (a hyperboloid of one sheet), then there will always exist a third distinct reconstruction.*

One may write down the dual statement straight away as follows.

Theorem 22.27. *A configuration $\{C_1, \dots, C_n; X_1, \dots, X_6\}$ of any number of cameras and six points allows an alternative reconstruction if and only if all camera centres $C_1, \dots, C_n$ and all the points $X_1, \dots, X_6$ lie on a ruled quadric surface.¹ If the quadric is non-degenerate (a hyperboloid of one sheet) then there will always exist a third distinct reconstruction.*

This result was originally proved in [Maybank-98]. Just to emphasize how a duality proof works, a proof for theorem 22.27 will be given.

Proof. Consider the configuration $\{C_1, \dots, C_n; X_1, \dots, X_6\}$. One renumbers the points so that the configuration is denoted by $\{C_1, \dots, C_n; E_1, \dots, E_4, X_1, X_2\}$ where $E_1, \dots, E_4$ are four non-collinear points, taken to be the vertices of a reference tetrahedron. If this configuration has an alternative reconstruction, then there exists another configuration $\{C'_1, \dots, C'_n; E_1, \dots, E_4, X'_1, X'_2\}$ such that for all $i = 1, \dots, n$ and $j = 1, 2$, one has $\langle C_i; E_1, \dots, E_4, X_j \rangle = \langle C'_i; E_1, \dots, E_4, X'_j \rangle$. Dualizing this using (22.3) yields

$$\langle \bar{X}_j; E_1, \dots, E_4, \bar{C}_i \rangle = \langle \bar{X}'_j; E_1, \dots, E_4, \bar{C}'_i \rangle \text{ for all } j = 1, 2 \text{ and } i = 1, \dots, n.$$

Now, theorem 22.26 applies, and one deduces that the two camera centres $\bar{X}_j$, the reference vertices $E_1, \dots, E_4$ and the points $\bar{C}_i$ all lie on a ruled quadric surface $\bar{S}$. Applying the reverse dualization, using theorem 22.20(iii), one sees that the points X_1, X_2 and the camera centres C_i all lie on the quadric surface S . This proves the forward implication of the theorem. The reverse implication is handled in a similar manner.

¹ In this statement, it is assumed that the set of points $X_1, \dots, X_6$ includes four non-coplanar points forming a reference tetrahedron and that none of the other two X_j nor any of the camera centres C_i lies on a face of this tetrahedron. Whether or not this condition is essential is not resolved.

The existence of a third distinct solution follows from the fact that the dual of a non-degenerate quadric is non-degenerate. $\square$

The minimum interesting case of theorem 22.27 is when $n = 3$, as studied in section 20.2.4(p510). In this case there are nine points in total (three cameras and six points). One can construct a quadric surface passing through these nine points (a quadric is defined by nine points). If the quadric is a ruled quadric (a hyperboloid of one sheet in the non-degenerate case), then there are three possible distinct reconstructions. Otherwise the reconstruction is unique. As with reconstruction from seven points in two views, algorithm 20.1(p511) for six points in three views requires the solution of a cubic equation. As with seven points, the distinction between cases where the cubic has one or three real solutions is explained as the difference between whether the quadric is ruled or not.

22.4 Three-view critical configurations

We now turn to consider the ambiguous configurations that may arise in the three-view case.

In this section, to distinguish the three cameras, we use superscripts instead of primes. Thus, let P^0, P^1, P^2 be three cameras and $\{P_i\}$ be a set of points. We ask under what circumstances there exists another configuration consisting of three other camera matrices Q^0, Q^1 and Q^2 and points $\{Q_i\}$ such that $P^j P_i = Q^j Q_i$ for all i and j . We require that the two configurations be projectively inequivalent.

Various special ambiguous configurations exist.

Points in a plane

If all the points lie in a plane, and $P_i = Q_i$ for all i , then any of the cameras may be moved without changing the projective equivalence class of the projected points. It is possible to choose P^j and Q^j with centres at any two preassigned locations in such a way that $P^j P_i = Q^j Q_i$.

Points on a twisted cubic

A similar ambiguous situation arises when all the points plus one of the cameras, say P^2 , lie on a twisted cubic. In this case, we may choose $Q^0 = P^0$ and $Q^1 = P^1$ and the points $Q_i = P_i$ for all i . Then according to the ambiguity of camera resectioning for points on a twisted cubic, (section 22.1.2), for any point C_0^2 on the twisted cubic a camera matrix Q^2 may be chosen with centre at C_0^2 such that $P^2 P_i = Q^2 Q_i$ for all i .

These examples of ambiguity are not very interesting, since they are no more than extensions of the one-view camera resectioning ambiguity. In the above examples, the points P_i and Q_i are the same in each case, and the ambiguity lies only in the placement of the cameras with respect to the points. More interesting ambiguities may also occur, as we consider next.

General three-view ambiguity

Suppose that the camera matrices (P^0, P^1, P^2) and (Q^0, Q^1, Q^2) are fixed, and we wish to find the set of all points such that $P^i P = Q^i Q$ for $i = 0, 1, 2$. Note that we are trying here to copy the two-view case in which both sets of camera matrices are chosen in advance. Later, we will turn to the less restricted case in which just one set of cameras is chosen in advance.

A simple observation is that a critical configuration for three views is also a critical set for each of the pairs of views. Thus one is led naturally to assume that the set of points for which $\{P^0, P^1, P^2, P_i\}$ is a critical configuration is simply the intersection of the point sets for which each of $\{P^0, P^1, P_i\}$, $\{P^1, P^2, P_i\}$ and $\{P^0, P^2, P_i\}$ is a critical configuration. Since by lemma 22.10(p542) each of these point sets is a ruled quadric, one is led to assume that the critical point set in the three-view case is simply an intersection of three quadrics. Although this is not far from the truth, the reasoning is somewhat fuzzy. The crucial point missing in this argument is that the corresponding conjugate points may not be the same for each of the three pairs.

More precisely, corresponding to the critical configuration $\{P^0, P^1, P_i\}$, there exists a conjugate configuration $\{Q^0, Q^1, Q_i^{01}\}$ for which $P^j P_i = Q^j Q_i^{01}$ for $j = 0, 1$. Similarly, for the critical configuration $\{P^0, P^2, P_i\}$, there exists a conjugate configuration $\{Q^0, Q^2, Q_i^{02}\}$ for which $P^j P_i = Q^j Q_i^{02}$ for $j = 0, 2$. However, the points Q_i^{02} are not necessarily the same as Q_i^{01} , so we cannot conclude that there exist points Q_i such that $P^j P_i = Q^j Q_i$ for all i and $j = 0, 1, 2$ – at least not immediately.

We now consider this a little more closely. Considering just the first pairs of cameras (P^0, P^1) and (Q^0, Q^1) , lemma 22.10(p542) tells us that if P and Q are points such that $P^j P = Q^j Q$, then P must lie on a quadric surface S_p^{01} determined by these camera matrices. Similarly, point Q lies on a quadric S_q^{01} . Likewise considering the camera pairs (P^0, P^2) and (Q^0, Q^2) one finds that the point P must lie on a second quadric S_p^{02} defined by these two camera pairs. Similarly, there exists a further quadric defined by the camera pairs (P^1, P^2) and (Q^1, Q^2) on which the point P must lie. Thus for points P and Q to exist such that $P^j P = Q^j Q$ for $j = 0, 1, 2$ it is necessary that P lie on the intersection of the three quadrics: $P \in S_p^{01} \cap S_p^{02} \cap S_p^{12}$. It will now be seen that this is almost a necessary and sufficient condition.

Result 22.28. *Let (P^0, P^1, P^2) and (Q^0, Q^1, Q^2) be two triplets of camera matrices and assume $P^0 = Q^0$. For each of the pairs $(i, j) = (0, 1), (0, 2)$ and $(1, 2)$, let S_p^{ij} and S_q^{ij} be the ruled quadric critical surfaces defined for camera matrix pairs (P^i, P^j) and (Q^i, Q^j) as in lemma 22.10(p542).*

- (i) *The centre of camera P^0 lies on $S_p^{01} \cap S_p^{02}$, P^1 lies on $S_p^{01} \cap S_p^{12}$, and P^2 lies on $S_p^{12} \cap S_p^{02}$.*
- (ii) *If there exist points P and Q such that $P^i P = Q^i Q$ for all $i = 0, 1, 2$, then P must lie on the intersection $S_p^{01} \cap S_p^{02} \cap S_p^{12}$ and Q must lie on $S_q^{01} \cap S_q^{02} \cap S_q^{12}$.*
- (iii) *Conversely, if P is a point lying on the intersection of quadrics $S_p^{01} \cap S_p^{02} \cap S_p^{12}$, but not on a plane containing the three camera centres C_q^0, C_q^1 and C_q^2 , then there exists a point Q lying on $S_q^{01} \cap S_q^{02} \cap S_q^{12}$ such that $P^i P = Q^i Q$ for all $i = 0, 1, 2$.*

Note that the condition that $P^0 = Q^0$ is not any restriction of generality, since the projective frames for the two configurations (P^0, P^1, P^2) and (Q^0, Q^1, Q^2) are independent. One may easily choose a projective frame for the second configuration in which this condition is true. This assumption is made simply so that one may consider the point P in relation to the projective frame of the second set of cameras.

The extra condition that the point P does not lie on the plane of camera centres C_Q^i is necessary, however the reader is referred to [Hartley-00b] for justification of this claim. Note that this case will usually not arise, however, since the intersection point of the three quadrics with the trifocal plane will be empty, or in special cases consist of a finite number of points. Where it does arise is through the possibility that the three camera centres C_Q^0, C_Q^1 and C_Q^2 are collinear, in which case any other point is coplanar with these three camera centres.

Proof. The first statement follows directly from lemma 22.10(p542). For the second part, the fact that the points P and Q lie on the intersections of the three quadrics follows (as pointed out before the statement of the theorem) from lemma 22.10(p542) applied to each pair of cameras in turn.

To prove the final assertion, suppose that P lies on the intersection of the three quadrics. Then from lemma 22.10(p542), applied to each of the three quadrics S_P^{ij} , there exist points Q^{ij} such that the following conditions hold:

$$\begin{aligned} P^0 P &= Q^0 Q^{01} & P^1 P &= Q^1 Q^{01} \\ P^0 P &= Q^0 Q^{02} & P^2 P &= Q^2 Q^{02} \\ P^1 P &= Q^1 Q^{12} & P^2 P &= Q^2 Q^{12}. \end{aligned}$$

It is easy to be confused by the superscripts here, but the main point is that each line is precisely the result of lemma 22.10(p542) applied to one of the three pairs of camera matrices at a time. These equations may be rearranged as

$$\begin{aligned} P^0 P &= Q^0 Q^{01} = Q^0 Q^{02} \\ P^1 P &= Q^1 Q^{01} = Q^1 Q^{12} \\ P^2 P &= Q^2 Q^{02} = Q^2 Q^{12}. \end{aligned}$$

Now, the condition that $Q^1 Q^{01} = Q^1 Q^{12}$ means that the points Q^{01} and Q^{12} are collinear with the camera centre C_Q^1 of Q^1 . Thus, assuming that the points Q^{ij} are distinct, they must lie in a configuration as shown in figure 22.7. One sees from the diagram that if two of the points are the same, then the third one is the same as the other two. If the three points are distinct, then the three points Q^{ij} and the three camera centres C_Q^i are coplanar, since they all lie in the plane defined by Q^{01} and the line joining Q^{02} to Q^{12} . Thus the three points all lie in the plane of the camera centres C_Q^i . However, since $P^0 P = Q^0 Q^{01} = Q^0 Q^{02}$ it follows that P must lie along the same line as Q^{01} and Q^{02} , and hence must lie in the same plane as the camera centres C_Q^i . $\square$

Thus, this result shows that the points in a 3-view critical configuration lie on the intersection of three quadrics, whereas the camera centres lie on the intersections of




Fig. 22.7. Configuration of the three camera centres and the three ambiguous points. If the three points Q^{ij} are distinct, then they all lie in the plane of the camera centres Q_q^i .

pairs of the quadrics. In general, the intersection of three quadrics will consist of eight points. In this case, the critical set with respect to the two triplets of camera matrices will consist of these eight points, less any such points that lie on the plane of the three cameras Q^i . In fact, it has been shown (in a longer unpublished version of [Maybank-98]) that of the eight points of intersection of three quadrics, only seven are critical, since the eighth point lies on the plane of the three cameras.

In some cases, however, the camera matrices may be chosen such that the three quadric surfaces meet in a curve. This will occur if the three quadrics S_p^{ij} are linearly dependent. For instance if $S_p^{12} = \alpha S_p^{01} + \beta S_p^{02}$, then any point P that satisfies $P^T S_p^{01} P = 0$ and $P^T S_p^{02} P = 0$ will also satisfy $P^T S_p^{12} P = 0$. Thus the intersection of the three quadrics is the same as the intersection of two of them. In this case, the three cameras must also lie on the same intersection curve. We define a non-degenerate *elliptic quartic* to be the intersection curve of two non-degenerate ruled quadrics, consisting of a single curve. This is a fourth-degree space curve. Other ways that two quadrics can intersect include a twisted-cubic plus a line, or two conics. Examples of elliptic quartics are shown in figure 22.8.

Example 22.29. Three-view critical configuration – the elliptic quartic.

We consider the quadrics represented by matrices A and $B = \tilde{B} + \tilde{B}^T$, where

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 \end{bmatrix} \quad \text{and} \quad \tilde{B} = \begin{bmatrix} p & q & s-t & -s-u \\ 0 & r & s+t & -s+u \\ 0 & 0 & -p-q-r & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (22.5)$$

Thus, B is a member of a 5-parameter family of quadrics generated by $\{p, q, r, s, t, u\}$, (remembering that scale is irrelevant). The camera matrices

$$P^0 = [I | 0], \quad P^1 = [I | (-1, -1, -1)^T] \quad \text{and} \quad P^2 = [I | (1, 1, -1)^T]$$

have centres lying on the intersection of these three quadrics.

We show that a configuration consisting of these three cameras, and any number of points on the intersection of the two quadrics is critical. This is demonstrated by explicitly exhibiting two alternative configurations consisting of three cameras Q_i and for each point P lying on the intersection of the two quadrics, a conjugate point Q such




Fig. 22.8. Examples of elliptic quartic curves generated as the intersection of two ruled quadrics.

that $P^i P = Q^i Q$. In fact, two different conjugate configurations are given in table 22.1 and table 22.2.

It may be verified directly that $P^i P = Q^i Q$ for all points $P = (x, y, xy, 1)^T$ and corresponding points Q , provided that P lies on the quadric B . (It always lies on quadric A). The easiest way to see this is to verify that $(P^i P) \times (Q^i Q) = 0$ for all such points. In fact for $i = 0, 1$, the cross-product is always zero, whereas for $i = 2$ it may be verified by direct computation that

$$(P^2 P) \times (Q^2 Q) = (P^T B P) (4, -4x, 4)^T$$

for the first solution, and

$$(P^2 P) \times (Q^2 Q) = (P^T B P) (-4y, 4, 4)^T$$

for the second solution. Thus $P^2 P = Q^2 Q$ if and only if P lies on B .

Note that in this example, A is the matrix representing the quadric S_p^{01} . $\triangle$

This example is quite general, since if A' and B' are two non-degenerate ruled quadrics containing the centres of three cameras, then by a projective transformation, we may map A' to A and the three camera centres to those of the given P^i , provided only that no two of the camera centres lie on the same generator of A' . In addition, B' will map to $A + \lambda B$ for some λ . Thus, the pencil generated by A' and B' , and hence also their intersection curve, are projectively equivalent to those generated by A and B .

The possibility that two of the camera centres lie on the same generator of A' is not a difficulty, since if the line of the camera centres lies on all quadrics in the pencil, then the quadric intersection can not be an non-degenerate elliptic quartic. Otherwise,

The camera matrices are

$$Q^0 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad Q^1 = \begin{bmatrix} -4 & 0 & 0 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & -2 & 1 \end{bmatrix}$$

and $Q^2 =$

$$\begin{bmatrix} -4(2p+q-t+u) & 8r & 4(p+q+2r+s+t) & -2(p+q-s-t) \\ 0 & 8(r+s-u) & -2(q-t+u) & -q+t-u \\ 8p & -8r & -2(2p+q-2s+3t+3u) & 2p+q-2s-t-u \end{bmatrix}$$

The conjugate point to $P = (x, y, xy, 1)^T$ is

$$Q = ((x-1)x, (x-1)y, (x-1)xy, -2x(-2+y+xy))^T.$$

Table 22.1. First conjugate solution to reconstruction problem for cameras P^i and points on the intersection of quadrics A and B given in (22.5).

The camera matrices are

$$Q^0 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad Q^1 = \begin{bmatrix} 0 & 0 & -2 & 1 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 2 & 1 \end{bmatrix}$$

and $Q^2 =$

$$\begin{bmatrix} -8(p+s+u) & 0 & 2(q+t-u) & -q-t+u \\ -8p & 4(q+2r+t-u) & -4(2p+q+r+s-t) & -2(q+r-s+t) \\ 8p & -8r & 2(q+2r-2s-3t-3u) & q+2r-2s+t+u \end{bmatrix}$$

The conjugate point to $P = (x, y, xy, 1)^T$ is

$$Q = ((y-1)x, (y-1)y, (y-1)xy, 2y(-2+x+xy))^T.$$

Table 22.2. Second conjugate solution to reconstruction problem.

we can choose A' to be one of the quadrics not containing the line of the two camera centres. This demonstrates

Result 22.30. Any configuration consisting of three cameras and any number of points lying on a non-degenerate elliptic quartic is critical.

22.5 Closure

22.5.1 The literature

The twisted cubic as the critical curve for camera resectioning was brought to the attention of the computer-vision community by [Buchanan-88]. For more about critical sets of two views, the reader is referred to [Maybank-90] and the book [Maybank-93]. In both these cases, the critical point sets were known much earlier. In fact [Buchanan-88] refers the reader to the German photogrammetric literature [Krames-42, Rinner-72]. For two views, the result that (v) and (viii) are *not* critical in theorem 22.9(p541) is due to Fredrik Kahl (unpublished).

The discussion of critical configurations for three views given in this chapter is only a part of what is known about this topic. More can be found in [Hartley-00b, Hartley-02a, Kahl-01a]. In particular, the elliptic-quartic configuration is extended to any number of cameras in [Kahl-01a]. A critical configuration for any number of cameras, consisting of points on a twisted cubic and cameras on a straight line is considered in [Hartley-03]. Earlier work on this area includes an investigation of the critical camera positions for sets of six points in [Maybank-98], and an unpublished report [Shashua-96] deals with critical configurations in three views.

Nothing has been said here about critical configurations of lines in three or more views, but this topic has been treated in [Buchanan-92]. In addition, critical configuration for linear reconstruction from lines (the linear line complex) have been identified in [Stein-99].

22.5.2 Notes and exercises

- (i) Fill out the details of the following sketch to prove that any two configurations consisting of a hyperboloid of one sheet and two points on the hyperboloid are projectively equivalent (via a projectivity of $\mathbb{P}^3$) provided that the points in both pairs either do or do not lie on the same generator.

Since any hyperboloid of one sheet is projectively equivalent to $x^2 + y^2 - z^2 = 1$, any two hyperboloids of one sheet are projectively equivalent to each other, and also to the hyperboloid given by $z = xy$. Define a 1D projective transformation $h_X(x) = x' = (ax + b)/(cx + d)$. One computes that

$$\begin{bmatrix} a & d & c & b \\ c & & & d \end{bmatrix} \begin{pmatrix} x \\ y \\ xy \\ 1 \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ x'y' \\ 1 \end{pmatrix}.$$

This is a 3D projective transformation taking the surface $z = xy$ to itself. Composing this with a similar transformation of Y one finds a projective transformation that takes $(x, y, xy, 1)^T$ to $(x', y', x'y', 1)^T$ while fixing the quadric $z = xy$. Since h_X and h_Y are arbitrary 1D projective transformations, this gives enough freedom to map any two points to two others.

- (ii) Show that Γ maps a line that meets one of the edges of the reference tetrahedron to a conic.
- (iii) Show that Γ maps a straight line meeting two opposite edges of the reference tetrahedron to a straight line meeting the same two edges.
- (iv) How are these configurations related to degenerate configurations for camera resectioning, as shown in figure 22.3(p538).
- (v) Two-view degeneracy occurs when all points and the two cameras lie on a ruled quadric. Given eight points on the corners of a Euclidean cube and two camera centres, show that these 10 points always lie on a quadric. If this is a ruled quadric, then the configuration is degenerate, and reconstruction is not possible from 8 points. Investigate under what conditions the quadric is ruled. *Hint:*

there is a two-parameter family of quadrics passing through the cube vertices. What does this two-parameter family look like?

- (vi) Extend result 22.30 by showing that a configuration of any number of cameras and points lying on a non-degenerate elliptic quartic is critical. This does not require complex computations. If stuck, refer to [Kahl-01a].