Chapter 3

Projective geometry

This chapter is devoted to the theory of planar geometry and the 2D projective plane, as well as projective transformations. The set of tools we obtain through this discussion will enable a natural and very convenient way of modelling the process of image formation with pinhole cameras.

3.1 Planar geometry and the 2D projective plane

We usually represent a point in the plane by a pair of coordinates $\underline{x} = \begin{bmatrix} x_1 & x_2 \end{bmatrix}^T$, thus we often identify the plane with $\mathbb{R}^2$. In this section we introduce homogeneous coordinates as a representation of points in a plane. This unifies the concept of the intersection of two lines (even parallel lines will be shown to intersect in a well-defined point) and leads us straight into the projective plane, $\mathbb{P}^2$, with its well-defined geometric computations.

3.1.1 Homogeneous representation of lines

A line in the plane is represented by an equation of the form

$$ax + by + c = 0, \tag{3.1}$$

with different choices of a, b and c yielding different lines. For example, the lines $2x + y - 1 = 0$, with $a = 2$, $b = 1$ and $c = -1$, and $x - 1 = 0$, with $a = 1$, $b = 0$ and $c = -1$, are illustrated in Figure 3.1.

A line can be uniquely represented by the coefficient vector $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$. Using this representation, the lines in Figure 3.1 are given by $\underline{l} = \begin{bmatrix} 2 & 1 & -1 \end{bmatrix}^T$ and $\underline{l}' = \begin{bmatrix} 1 & 0 & -1 \end{bmatrix}^T$.

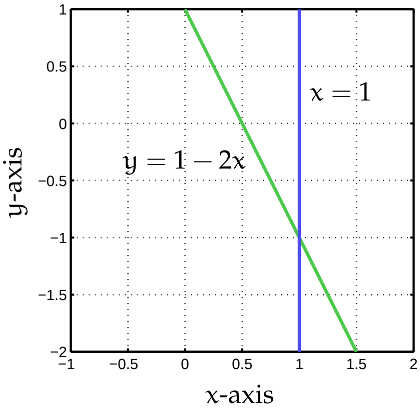

Figure 3.1: Two lines in the plane $\mathbb{R}^2$.

In this notation, the line represented by $k\underline{l} = \begin{bmatrix} ka & kb & kc \end{bmatrix}^T$, with k any nonzero constant, is the same as the line represented by $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$, since x and y satisfy (3.1) if and only if they satisfy $kax + kby + kc = 0$. For this reason we consider the vectors $\begin{bmatrix} a & b & c \end{bmatrix}^T$ and $k \begin{bmatrix} a & b & c \end{bmatrix}^T (k \neq 0)$ to be equivalent. An equivalence class of vectors under this equivalence relationship is known as a *homogeneous vector*, where any particular vector $\begin{bmatrix} a & b & c \end{bmatrix}^T$ is a representative of a whole class of equivalent vectors. The set of equivalence classes in $\mathbb{R}^3 \setminus \{\underline{0}^3\}$ (the vector space $\mathbb{R}^3$ with the zero vector $\underline{0}^3$ removed) forms the *projective space* $\mathbb{P}^2$.

3.1.2 Homogeneous representation of points

A point $\begin{bmatrix} x & y \end{bmatrix}^T$ lies on the line $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$ if and only if $ax + by + c = 0$, which can be written in terms of an inner product

$$\begin{bmatrix} x & y & 1 \end{bmatrix} \underline{l} = 0. \tag{3.2}$$

Here the point $\begin{bmatrix} x & y \end{bmatrix}^T$ in $\mathbb{R}^2$ is represented by a 3-vector by adding a third coordinate of 1. Note that, for any nonzero constant k, we have

$$\begin{bmatrix} kx & ky & k \end{bmatrix} \underline{l} = 0 \tag{3.3}$$

if and only if (3.2) holds. It is therefore natural to consider the vector $\begin{bmatrix} kx & ky & k \end{bmatrix}^T$ (with $k \neq 0$) equivalent to $\begin{bmatrix} x & y & 1 \end{bmatrix}^T$.

We therefore have, similar to the case of lines, an equivalence class where the vector $\begin{bmatrix} x & y & 1 \end{bmatrix}^T$ represents the entire class. Here, an arbitrary homogeneous vector $\begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^T$ represents the point $\begin{bmatrix} \frac{x_1}{x_3} & \frac{x_2}{x_3} \end{bmatrix}^T$ in $\mathbb{R}^2$. Points, written as homogeneous vectors, are also elements of $\mathbb{P}^2$.

Homogeneous vector representations of lines and points allow us to simply express some previously clumsy concepts, as the following theorem illustrates (the proof is easy and left as an exercise).

Theorem 3.1.1. *The intersection point of lines* $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$ *and* $\underline{l}' = \begin{bmatrix} a' & b' & c' \end{bmatrix}^T$ *is* $\underline{x} = \underline{l} \times \underline{l}'$*, where* $\times$ *represents the vector cross product. The line joining two points* $\underline{x}$ *and* $\underline{x}'$ *is given by* $\underline{l} = \underline{x} \times \underline{x}'$*.*

As an example, suppose we wish to determine the intersection point of the lines $x = 1$ and $y = x + 1$. The line $x = 1$ is written as $-x + 0y + 1 = 0$, and thus has the homogeneous representation $\underline{l} = \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}^T$, while the line $y = x + 1$ is written as $-x + y - 1 = 0$, and thus has the homogeneous representation $\underline{l}' = \begin{bmatrix} -1 & 1 & -1 \end{bmatrix}^T$. Then, from Theorem 3.1.1, we calculate the intersection point as

$$\underline{x} = \underline{l} \times \underline{l}' = \begin{vmatrix} \underline{i} & \underline{j} & \underline{k} \\ -1 & 0 & 1 \\ -1 & 1 & -1 \end{vmatrix} = \begin{bmatrix} -1 \\ -2 \\ -1 \end{bmatrix},$$

which is the homogeneous representation of the point $\begin{bmatrix} 1 & 2 \end{bmatrix}^T$, consistent with the intersection obtained in Figure 3.2(a).

Now suppose we want to determine the line joining points $(0, 1)$ and $(1, 0)$. The homogeneous representation of $(0, 1)$ is $\underline{x} = \begin{bmatrix} 0 & 1 & 1 \end{bmatrix}^T$, and of $(1, 0)$ is $\underline{x}' = \begin{bmatrix} 1 & 0 & 1 \end{bmatrix}^T$. From Theorem 3.1.1 we calculate the homogeneous representation of the line joining the two points as

$$\underline{l} = \underline{x} \times \underline{x}' = \begin{vmatrix} \underline{i} & \underline{j} & \underline{k} \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{vmatrix} = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}.$$

The inhomogeneous representation of this line is then $x + y - 1 = 0$ or, equivalently, $y = -x + 1$. Figure 3.2(b) illustrates.

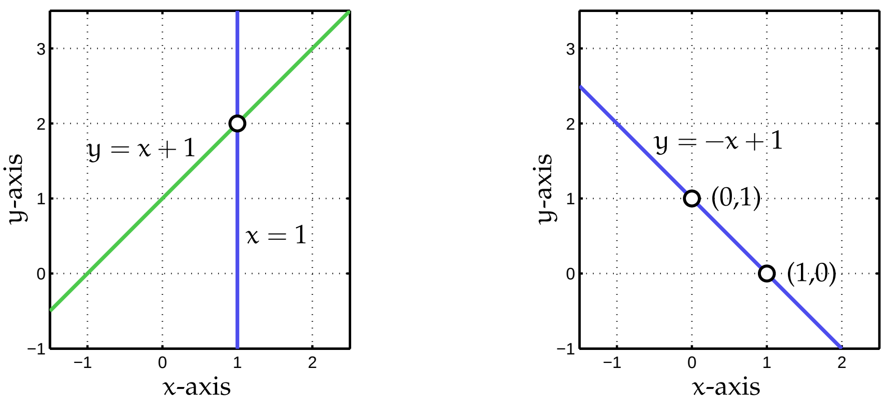

Figure 3.2: (a) intersecting lines and (b) a line joining two points in the plane $\mathbb{R}^2$.

3.1.3 Intersection of parallel lines

Consider two parallel lines $ax + by + c = 0$ and $ax + by + c' = 0$. These lines are represented by the vectors $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$ and $\underline{l}' = \begin{bmatrix} a & b & c' \end{bmatrix}^T$. Although these lines are parallel, we can compute their intersection $\underline{x}$ using Theorem 3.1.1:

$$\underline{x} = \underline{l} \times \underline{l}' = \begin{bmatrix} bc' - bc \\ ac - ac' \\ 0 \end{bmatrix} = (c' - c) \begin{bmatrix} b \\ -a \\ 0 \end{bmatrix} = \begin{bmatrix} b \\ -a \\ 0 \end{bmatrix}.$$

The last step follows from the fact that we are working in homogeneous coordinates. If we now try to find the inhomogeneous version of $\underline{x}$, we would divide by zero to get $\begin{bmatrix} \frac{b}{0} & -\frac{a}{0} \end{bmatrix}^T$, which makes no sense, unless to suggest that the point of intersection has infinitely large coordinates. This observation agrees with the usual idea, that parallel lines intersect at infinity.

Consider for example the parallel lines $x = 1$ and $x = 2$, that is $\underline{l} = \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}^T$ and $\underline{l}' = \begin{bmatrix} -1 & 0 & 2 \end{bmatrix}^T$. The intersection point of these two lines is $\underline{x} = \underline{l} \times \underline{l}' = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^T$, which is the point at infinity in the direction of the y-axis.

In general, points with homogeneous coordinates $\begin{bmatrix} x & y & 0 \end{bmatrix}^T$ do not correspond to any finite point in $\mathbb{R}^2$.

3.1.4 Ideal points and the line at infinity

The projective space $\mathbb{P}^2$ consists of all homogeneous vectors $\underline{x} = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^T$, where $\underline{x}$ represents a finite point in $\mathbb{R}^2$ if $x_3 \neq 0$. If $x_3 = 0$ we say $\underline{x}$ is an *ideal point* (or a *point at infinity*).

The set of all ideal points may be written as $\begin{bmatrix} x_1 & x_2 & 0 \end{bmatrix}^T$. Note that this set lies on a single line $\underline{l} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$, since

$$\begin{bmatrix} x_1 & x_2 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 0. \tag{3.4}$$

We call this line the *line at infinity* and denote it by $\underline{l}_\infty = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$.

The line $\underline{l} = \begin{bmatrix} a & b & c \end{bmatrix}^T$ intersects the line at infinity $\underline{l}$ at $\underline{l} \times \underline{l}_\infty = \begin{bmatrix} b & -a & 0 \end{bmatrix}^T$, which is an ideal point. The line $\underline{l}' = \begin{bmatrix} a & b & c' \end{bmatrix}^T$, parallel to $\underline{l}$, intersects $\underline{l}_\infty$ at the same point. In inhomogeneous coordinates, the vector $\begin{bmatrix} b & -a \end{bmatrix}^T$ is tangent to the line $ax + by + c = 0$, and orthogonal to the line's normal, and so represents the line's direction. If we were to vary the line's direction, the ideal point $\begin{bmatrix} b & -a & 0 \end{bmatrix}^T$ would also vary over $\underline{l}_\infty$. Thus the line at infinity can be thought of as a set of directions of lines in the plane.

3.1.5 A useful way to think of $\mathbb{P}^2$

A useful way of thinking of $\mathbb{P}^2$ is as a set of rays in $\mathbb{R}^3$; see Figure 3.3. The set of vectors $k \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^T$ (with k varying) forms a ray through the origin. Such a ray can be thought of as representing a single point in $\mathbb{P}^2$. The corresponding point in $\mathbb{R}^2$ may be obtained by intersecting a particular ray with the plane $x_3 = 1$, thus effectively scaling the ray so that its third component is 1.

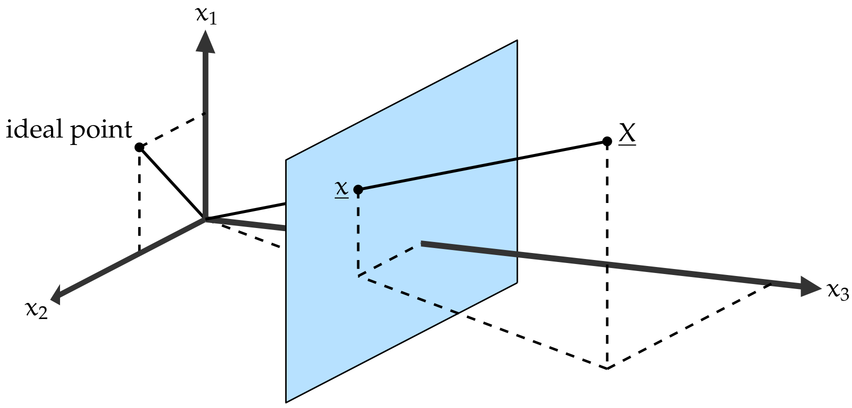

Figure 3.3: A model of the projective plane.

3.2 Projective transformations

A projective transformation is an invertible map h from $\mathbb{P}^2$ to $\mathbb{P}^2$ such that straight lines are mapped to straight lines or, equivalently, such that three points $\underline{u}$, $\underline{v}$ and $\underline{w}$ lie on a straight line if and only if the three points $h(\underline{u})$, $h(\underline{v})$ and $h(\underline{w})$ lie on a straight line. It is also called a *collineation*, a *projectivity* or a *homography*.

Theorem 3.2.1. *A mapping $h : \mathbb{P}^2 \rightarrow \mathbb{P}^2$ is a projective transformation if and only if there exists a nonsingular, homogeneous 3×3 matrix H such that for any point in $\mathbb{P}^2$ represented by a vector $\underline{x}$ it is true that $h(\underline{x}) = H\underline{x}$.*

This theorem allows us to define a projective transformation in terms of matrix multiplication. A projective transformation is a linear transformation on homogeneous 3-vectors represented by a nonsingular, homogeneous 3×3 matrix,

$$\begin{bmatrix} x'_1 \\ x'_2 \\ x'_3 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}, \tag{3.5}$$

or more briefly, $\underline{x}' = H\underline{x}$.

Note that, since the vectors $\underline{x}$ and $\underline{x}'$ are homogeneous, the projective transformation $\underline{x}' = kH\underline{x}$, for any nonzero k, is equivalent to $\underline{x}' = H\underline{x}$, i.e. an arbitrary scale factor does not change the projective transformation. Consequently, H is a homogeneous matrix since, as in the case of the homogeneous representation of a point, any nonzero multiple of H is equivalent to H.

There are 8 independent elements in the homogeneous matrix H, and so it follows that there are 8 degrees of freedom in a projective transformation.

3.2.1 Transformation of lines

By the definition of a projective transformation we know that if the n points $\underline{x}_1, \underline{x}_2, \ldots, \underline{x}_n$ lie on a line $\underline{l}$, i.e. $\underline{x}_i^T \underline{l} = 0$, $i = 1, \ldots, n$, then the transformed points $\underline{x}_1', \underline{x}_2', \ldots, \underline{x}_n'$, where $\underline{x}_i' = H\underline{x}_i$, $i = 1, \ldots, n$, lie on a line, say $\underline{l}'$. We claim that the line $\underline{l}'$ is given by $\underline{l}' = H^{-T}\underline{l}$. This claim is easily verified, since

$$(\underline{x}_i')^T(\underline{l}') = (H\underline{x}_i)^T(H^{-T}\underline{l}) = \underline{x}_i^T(H^TH^{-T})\underline{l} = \underline{x}_i^T\underline{l} = 0. \tag{3.6}$$

3.2.2 The hierarchy of transformations

The most general form of a projective transformation is given in matrix form by

$$\underline{x}' = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \underline{x}. \tag{3.7}$$

In this section, we consider special cases of the projective transformation, starting with the most restricted case, an isometry, and ending with the most general projective transformation.

Isometric transformations

An *isometric* transformation (or isometry) of the plane $\mathbb{R}^2$ preserves Euclidean distance (*iso* = same, *metric* = measure), e.g. a rotation. The most general isometry is represented by the matrix equation

$$\begin{bmatrix} x_1' \\ x_2' \\ 1 \end{bmatrix} = \begin{bmatrix} \epsilon\cos\theta & -\sin\theta & t_x \\ \epsilon\sin\theta & \cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}. \tag{3.8}$$

where $\epsilon = \pm 1$. If $\epsilon = 1$, then the isometry is *orientation-preserving* and is a Euclidean transformation (a composition of a rotation and a translation). If $\epsilon = -1$ the isometry reverses orientation, for example, a composition of a reflection and a translation. We focus on Euclidean transformations, as they are predominant in applications. We can write a Euclidean transformation in block form as

$$\underline{x}' = H_E \underline{x} = \begin{bmatrix} R & \underline{t} \\ \underline{0}^T & 1 \end{bmatrix} \underline{x}, \tag{3.9}$$

where R is a 2×2 rotation matrix (an orthogonal matrix, satisfying $R^T R = RR^T = I$, with determinant 1), $\underline{t}$ is a translation 2-vector and $\underline{0}^T = \begin{bmatrix} 0 & 0 \end{bmatrix}$.

Degrees of freedom: a Euclidean transformation has 3 degrees of freedom: 1 for rotation and 2 for translation. Thus 3 parameters must be specified, and we can compute the transformation from two point correspondences.

Invariants: length, angles, area, parallel lines.

Similarity transformations

A *similarity* transformation (or a similarity) is an isometry together with an isotropic scaling (*isotropic* means equal in all directions). In the case of a Euclidean transformation composed with an isotropic scaling, the form of the similarity is

$$\begin{bmatrix} x_1' \\ x_2' \\ 1 \end{bmatrix} = \begin{bmatrix} s\cos\theta & -s\sin\theta & t_x \\ s\sin\theta & s\cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}, \tag{3.10}$$

where the scaling factor s is a scalar. In block form:

$$\underline{x}' = H_S \underline{x} = \begin{bmatrix} sR & \underline{t} \\ \underline{0}^T & 1 \end{bmatrix} \underline{x}. \tag{3.11}$$

Degrees of freedom: a similarity has 4 degrees of freedom (3 for the isometry plus 1 for the scaling) and we can compute a similarity from two point correspondences.

Invariants: ratio of two lengths, angles, ratio of areas, parallel lines.

Affine transformations

An *affine* transformation (or an affinity) is a non-singular linear transformation followed by a translation. In matrix form:

$$\begin{bmatrix} x_1' \\ x_2' \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}, \tag{3.12}$$

or more compactly in block form, with A a nonsingular 2 × 2 matrix:

$$\underline{x}' = H_A \underline{x} = \begin{bmatrix} A & \underline{t} \\ \underline{0}^T & 1 \end{bmatrix} \underline{x}. \tag{3.13}$$

Using the *singular value decomposition* (SVD) of the matrix A, we can develop an understanding of the geometric effects of A as the composition of a scaling and rotations. Suppose the SVD of A is $A = U\Sigma V^T$, where U and V are orthogonal matrices and $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2)$. We regroup this factorization as follows:

$$A = U\Sigma V^T = U\left(V^T V\right)\Sigma V^T = \left(UV^T\right)V\Sigma V^T, \tag{3.14}$$

since V is an orthogonal matrix. With $P = UV^T$ we note that P is also orthogonal, so that multiplication with P is equivalent to a rotation. Suppose P represents a rotation by ϕ and V a rotation by θ, then

$$A = PV\Sigma V^T = R(\phi)R(-\theta)\Sigma R(\theta), \tag{3.15}$$

where $R(\phi)$ and $R(\theta)$ represent rotations by ϕ and θ, respectively. The matrix A is therefore a composition of a rotation (by θ); a non-isotropic scaling by σ_1 and σ_2 respectively in the x- and y-directions; a rotation (by $-\theta$); and finally another rotation (by ϕ). The essence of an affinity is the scaling in orthogonal directions, oriented at a particular angle.

Degrees of freedom: the SVD argument above leads to the understanding that an affinity has two more degrees of freedom when compared to a similarity, due to the non-isotropic scaling at a particular angle, making the total degrees of freedom for an affinity 6.

Invariants:

- ratio of areas
  Area is scaled by $\det A = \sigma_1\sigma_2$, i.e. the ratio of areas is preserved.

- parallel lines
  Two parallel lines meet at $\begin{bmatrix} x_1 & x_2 & 0 \end{bmatrix}^T$. An affinity maps this ideal point to

  $$\underline{x}' = \begin{bmatrix} A & \underline{t} \\ \underline{0}^T & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ 0 \end{bmatrix} = \begin{bmatrix} A \begin{bmatrix} x \\ y \end{bmatrix} \\ 0 \end{bmatrix}.$$

  Since the last element of $\underline{x}'$ is zero, it is also an ideal point and the transformed lines are still parallel to each other.

Projective transformations

A *projective* transformation (or projectivity) as defined in (3.5) is a general non-singular linear transformation of homogeneous coordinates, and generalizes an affinity even more. It is given in block form as

$$\underline{x}' = H_P\underline{x} = \begin{bmatrix} A & \underline{t} \\ \underline{v}^T & v \end{bmatrix} \underline{x}, \tag{3.16}$$

where $\underline{v}$ is not necessarily zero and v not necessarily nonzero.

Degrees of freedom: the matrix H_P has nine elements, but since it is a homogeneous matrix, any nonzero multiple of H_P is equivalent to H_P and we therefore only have 8 degrees of freedom. We can compute the matrix H_P from a four point correspondence (with no three points being collinear).

Invariants: cross ratio of four points (ratio of ratios of lengths on a line).

Note that if a projective transformation preserves parallel lines, it is an affinity. To prove this, suppose a projective transformation is such that parallel lines are mapped to parallel lines, i.e.

$$\begin{bmatrix} A & \underline{t} \\ \underline{v}^T & v \end{bmatrix} \begin{bmatrix} \underline{x} \\ 0 \end{bmatrix} = \begin{bmatrix} A\underline{x} \\ \underline{v}^T\underline{x} \end{bmatrix} = \begin{bmatrix} x_1' \\ x_2' \\ 0 \end{bmatrix}. \tag{3.17}$$

Then it must hold that $\underline{v}^T\underline{x} = 0$ for all $\underline{x}$, so that $\underline{v}$ must necessarily be equal to the zero vector. Since the matrix must be non-singular, it must then also hold that $v \neq 0$, so that the transformation matrix is of the form

$$\begin{bmatrix} A & \underline{t} \\ \underline{0}^T & c \end{bmatrix} \tag{3.18}$$

for a nonzero constant c. The transformation is therefore an affinity.

3.2.3 Applying projective transformations to images

Since projective transformations map points in the 2D plane to points in the 2D plane, they can be used to geometrically transform images in much the same way as the mappings we've seen in section 1.7 of these notes. We would essentially map every pixel's 2D location in a given image to a new 2D location, by multiplying a homogeneous representation of the input (x, y)-coordinate with the given projective transformation matrix, to create a new image.

Of course, in practice we would first initialize the output image, loop through its pixels, apply the inverse transformation

$$\underline{x} = H^{-1}\underline{x}', \tag{3.19}$$

convert $\underline{x}$ to (x, y) coordinates in the input image through division of the third element of $\underline{x}$, and perform interpolation over the pixels surrounding (x, y).

A few examples of an image transformed in this way are shown in Figure 3.4. See if you can convince yourself of the invariants of the various transformations, as listed in the previous section. For example, do straight lines remain straight under all of these transformations; are angles and lengths preserved in the case of the isometry; do parallel lines remain parallel under the affinity; etc.

Figure 3.4: Different types of projective transformations on an input image.

3.3 Removing perspective distortion from a plane

Consider the *central projection* in Figure 3.5. Projection along rays through a common point (the centre of the projection) defines a mapping from one plane to another, say the world plane to the image plane. It maps lines to lines, and so is a projective transformation. If a coordinate system is defined in each of the planes, say world coordinates and image coordinates, and points are represented in homogeneous coordinates then the mapping may be expressed as

$$\underline{x}' = \mathsf{H}\underline{x}, \tag{3.20}$$

where H is a non-singular 3×3 matrix. This transformation is called a *perspective transformation* (or plane-to-plane homography).

Shape is distorted under such a perspective transformation but, since the image plane is related to the world plane via a projective transformation, we can undo this distortion. This is done by computing the inverse projective transformation and applying it to the image, resulting in an image where the objects have their correct geometric shape.

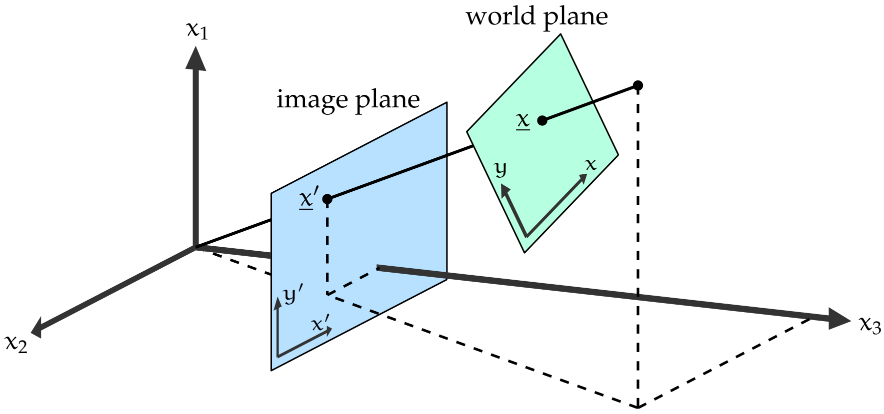

Figure 3.5: Transforming from a world plane to the image plane.

In order to compute this inverse transformation, we choose local inhomogeneous coordinates $\underline{x} = (x, y)$ and $\underline{x}' = (x', y')$ known to correspond, both measured directly from the world and the image plane respectively. Then the projective transformation

$$\begin{bmatrix} x_1' \\ x_2' \\ x_3' \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \tag{3.21}$$

can be written in inhomogeneous form as

$$x' = \frac{x_1'}{x_3'} = \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}}, \quad y' = \frac{x_2'}{x_3'} = \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}}. \tag{3.22}$$

Each point correspondence leads to two linear equations (linear in the unknown coefficients h_{ij}) of the form

$$\begin{aligned} x'(h_{31}x + h_{32}y + h_{33}) &= h_{11}x + h_{12}y + h_{13} \\ y'(h_{31}x + h_{32}y + h_{33}) &= h_{21}x + h_{22}y + h_{23} \end{aligned}$$

or equivalently,

$$\begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -x'x & -x'y & -x' \\ 0 & 0 & 0 & x & y & 1 & -y'x & -y'y & -y' \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ \vdots \\ h_{33} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \tag{3.23}$$

Since H has eight degrees of freedom and each point correspondence results in two linear equations, we need four point correspondences of the form $\underline{x} \leftrightarrow \underline{x}'$ to solve for H. Of course, if one has more than four point correspondences it is even better. The resulting overdetermined system for the null space can then be solved in a least-squares sense. Since measurement errors corrupt the system, the rank of the coefficient matrix will now be 9, but with a small 9th singular value. The 9th singular vector is therefore used as an approximation for the null space.

From four point correspondences we get the system of equations $A\underline{h} = \underline{0}$, where A is an 8×9 matrix of the form

$$A = \begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1'x_1 & -x_1'y_1 & -x_1' \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -y_1'x_1 & -y_1'y_1 & -y_1' \\ x_2 & y_2 & 1 & 0 & 0 & 0 & -x_2'x_2 & -x_2'y_2 & -x_2' \\ 0 & 0 & 0 & x_2 & y_2 & 1 & -y_2'x_2 & -y_2'y_2 & -y_2' \\ x_3 & y_3 & 1 & 0 & 0 & 0 & -x_3'x_3 & -x_3'y_3 & -x_3' \\ 0 & 0 & 0 & x_3 & y_3 & 1 & -y_3'x_3 & -y_3'y_3 & -y_3' \\ x_4 & y_4 & 1 & 0 & 0 & 0 & -x_4'x_4 & -x_4'y_4 & -x_4' \\ 0 & 0 & 0 & x_4 & y_4 & 1 & -y_4'x_4 & -y_4'y_4 & -y_4' \end{bmatrix}, \tag{3.24}$$

and $\underline{h} = \begin{bmatrix} h_{11} & h_{12} & \cdots & h_{33} \end{bmatrix}^T$. Choosing points so that A has full rank, i.e A is of rank 8, implies that no three points can be collinear. The vector $\underline{h}$ is a vector in the null space of the matrix A, or any nonzero multiple of it. One way of obtaining $\underline{h}$ is to compute the SVD of the matrix A, i.e. $A = U\Sigma V^T$. Then the last column of V will be a basis for the null space of A, so that we can choose $\underline{h}$ as the last column of V.

Applying this to a perspective distorted image, we can correct the distortion. For example in Figure 3.6(a) we recognise that the four corners of the window are the corners of a rectangle in world coordinates. We then choose the points $\underline{x}_i = (x_i, y_i)$, $i = 1, \ldots, 4$, as the four corners of the window and map them to the four corners of a rectangle to get the projective transformation H. Once we have H, we apply its inverse to the entire distorted image (following the procedure in section 3.2.3) to obtain the rectified image shown in Figure 3.6(b). A second example is shown in Figure 3.6(c) and (d).

(a) distorted image

(b) corrected image

(c) distorted image

(d) corrected image

Figure 3.6: Removing perspective distortion with plane-to-plane homographies.

Note that the computation of the matrix H does not require any knowledge of the camera's parameters (such as the focal length), or the orientation of the plane.

3.4 Recovering affine properties of images

In the previous section we corrected a perspectively distorted image up to a similarity. Suppose now that we are given less information and asked to correct the image up to an affinity (such that parallel lines are restored).

We may accomplish this task by identifying the image of the line at infinity (where parallel lines meet), and mapping it to infinity.

3.4.1 Affinities and the line at infinity

Under a projective transformation ideal points may be mapped to finite points, since

$$\begin{bmatrix} A & \underline{t} \\ \underline{v}^T & v \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ 0 \end{bmatrix} = \begin{bmatrix} A\underline{y} \\ \underline{v}^T\underline{y} \end{bmatrix}, \tag{3.25}$$

where $\underline{y} = \begin{bmatrix} x_1 & x_2 \end{bmatrix}^T$. If $\underline{v}^T\underline{y} \neq 0$, this point is finite. Consequently, a projective transformation may map the line at infinity $\underline{l}_\infty = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$ to a finite line. We have already witnessed this phenomenon: parallel lines in the distorted images in Figure 3.6 intersect at finite points.

However, if the transformation is an affinity, say

$$H_A = \begin{bmatrix} A & \underline{t} \\ \underline{0}^T & 1 \end{bmatrix}, \tag{3.26}$$

then $\underline{l}_\infty$ is mapped to $\underline{l}_\infty$, and not a finite line. To prove this, we use the fact that if the line $\underline{l}$ is mapped to $\underline{l}'$ under the transformation H, then the line $\underline{l}'$ is given by $H^{-T}\underline{l}$ (refer to section 3.2.1). Then

$$\underline{l}' = H_A^{-T}\underline{l}_\infty = \begin{bmatrix} A^{-T} & \underline{0} \\ -\underline{t}^T A^{-T} & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \underline{l}_\infty. \tag{3.27}$$

The converse is also true, i.e. an affine transformation is the most general linear transformation that maps $\underline{l}_\infty$ to $\underline{l}_\infty$. This can be justified as follows. If an ideal point, say $\underline{x} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T$, is mapped to another ideal point, then the first element of $\underline{v}$ must be zero. Similarly, the second element of $\underline{v}$ must be zero. Therefore the transformation is an affinity.

Note that the line $\underline{l}_\infty$ is not fixed point-wise under an affinity. For example, the affinity

$$\mathsf{H}_A = \begin{bmatrix} 1 & 2 & -\frac{1}{2} \\ -1 & \frac{1}{2} & 1 \\ 0 & 0 & 1 \end{bmatrix},$$

maps the ideal point $\begin{bmatrix} 1 & 1 & 0 \end{bmatrix}^\mathsf{T}$ to another ideal point, namely

$$\begin{bmatrix} 1 & 2 & -\frac{1}{2} \\ -1 & \frac{1}{2} & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ \frac{1}{2} \\ 0 \end{bmatrix}.$$

We will use this property of affinities to recover affine properties by identifying the line at infinity $\underline{l}_\infty$.

3.4.2 Using the line at infinity to recover affine properties

Once the image of the line at infinity has been identified, say $\underline{l} = \begin{bmatrix} \ell_1 & \ell_2 & \ell_3 \end{bmatrix}^\mathsf{T}$, it is possible to construct a projectivity that will project it back to $\underline{l}_\infty$. Since an affinity maps $\underline{l}_\infty$ to itself, this transformation is only defined up to an affinity.

It is straightforward to verify that, if $\ell_3 \neq 0$, the projective map

$$\mathsf{H} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \ell_1 & \ell_2 & \ell_3 \end{bmatrix} \tag{3.28}$$

maps $\underline{l}$ to $\underline{l}_\infty$, since

$$\mathsf{H}^{-\mathsf{T}}\underline{l} = \begin{bmatrix} 1 & 0 & -\frac{\ell_1}{\ell_3} \\ 0 & 1 & -\frac{\ell_2}{\ell_3} \\ 0 & 0 & \frac{1}{\ell_3} \end{bmatrix} \begin{bmatrix} \ell_1 \\ \ell_2 \\ \ell_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \tag{3.29}$$

This prescription does not work if $\ell_3 = 0$. We leave it as an exercise to verify that if $\ell_3 = 0$ and $\ell_1 \neq 0$ we may choose

$$\mathsf{H} = \begin{bmatrix} 0 & 1 & -1 \\ 0 & 0 & 1 \\ \ell_1 & \ell_2 & 0 \end{bmatrix}, \tag{3.30}$$

or if $\ell_3 = 0$ and $\ell_1 = 0$ (implying that $\ell_2 \neq 0$),

$$\mathsf{H} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & -1 \\ 0 & \ell_2 & 0 \end{bmatrix}. \tag{3.31}$$

Once a mapping H is specified, we may transform the image accordingly (following the procedure discussed in section 3.2.3).

It remains to identify the image of the line at infinity. One way to do this is to identify lines that were parallel on the original (world) plane. Their point of intersection provides an image of a point on $\underline{l}_\infty$. If we can find a second such point, then the image of $\underline{l}_\infty$ is the line through these two points. Let's illustrate with an example.

For the image in Figure 3.7(a), we notice that the frame of the window should be rectangular, so that the top and bottom of the window frame should be parallel and the two sides should be parallel. We identify the four corners of the window frame, and then use them to calculate the lines that run with the sides of the frame, labelled $\underline{l}$ and $\underline{l}'$ in Figure 3.7(b). We know that these lines are the images of parallel lines, so that their intersection point, labelled $\underline{x}$ in Figure 3.7(c), is the image of an ideal point. Similarly, we calculate the intersection point of the two lines running with the top and bottom of the window frame, to get $\underline{x}'$. Now that we have two points, that are each the image of an ideal point, we can calculate the imaged line at infinity as $\underline{l}'_\infty = \underline{x} \times \underline{x}'$. We construct the matrix H that maps this imaged line at infinity to the canonical form of the line at infinity, i.e. $\underline{l}_\infty = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$, and apply the transformation H to the entire image to get the image in Figure 3.7(d), where the sides of the window frame are parallel.

(a) distorted image

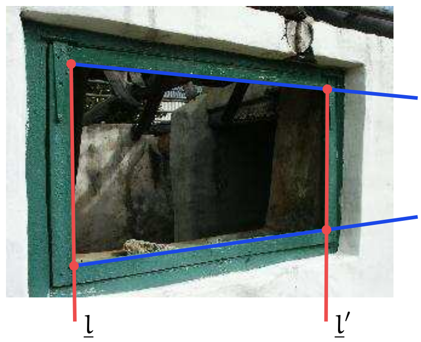

(b) imaged parallel lines identified

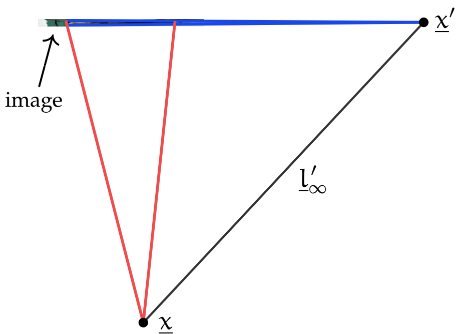

(c) intersection points of parallel lines and imaged $\underline{l}'_\infty$ identified

(d) image corrected up to an affinity: parallel lines are parallel

Figure 3.7: Recovering affine properties from a distorted image using $\underline{l}_\infty$.

3.5 Cross ratios and vanishing points

Let us now consider the projective geometry of a line, $\mathbb{P}^1$, which proceeds much the same as that of a plane. A point $\underline{x}$ on a line is represented by homogeneous coordinates $\begin{bmatrix} x_1 & x_2 \end{bmatrix}^T$, and a point for which $x_2 = 0$ is an ideal point of the line. A projective transformation of a line is represented by a non-singular 2×2 homogeneous matrix H_2, so that $\underline{x}' = H_2\underline{x}$, and has three degrees of freedom.

The *cross ratio* is the basic projective invariant of $\mathbb{P}^1$. Given four points on a line, say $\underline{a}$, $\underline{b}$, $\underline{c}$ and $\underline{d}$, then the cross ratio is defined as

$$\text{Cross}(\underline{a}, \underline{b}, \underline{c}, \underline{d}) = \frac{|\underline{a}\,\underline{b}|\ |\underline{c}\,\underline{d}|}{|\underline{a}\,\underline{c}|\ |\underline{b}\,\underline{d}|}, \tag{3.32}$$

where

$$|\underline{a}\,\underline{b}| = \det \begin{bmatrix} a_1 & b_1 \\ a_2 & b_2 \end{bmatrix}, \tag{3.33}$$

and similarly for the other points.

Cross ratios have the following important properties:

- The value of the cross ratio does not depend on which homogeneous representative of a point is used, since the scale cancels between the numerator and denominator.
- The definition of the cross ratio is still valid if one of the points is an ideal point.
- The value of the cross ratio is invariant under any projective transformation of the line (see Figure 3.8) since if $\underline{x}' = H\underline{x}$, then

$$\begin{aligned}
\text{Cross}(\underline{a}', \underline{b}', \underline{c}', \underline{d}') &= \frac{|\underline{a}'\,\underline{b}'|\ |\underline{c}'\,\underline{d}'|}{|\underline{a}'\,\underline{c}'|\ |\underline{b}'\,\underline{d}'|} \\
&= \frac{|(\lambda_1 H\underline{a})\ (\lambda_2 H\underline{b})|\ |(\lambda_3 H\underline{c})\ (\lambda_4 H\underline{d})|}{|(\lambda_1 H\underline{a})\ (\lambda_3 H\underline{c})|\ |(\lambda_2 H\underline{b})\ (\lambda_4 H\underline{d})|} \\
&= \frac{(\lambda_1\lambda_2\,|H|\,|\underline{a}\,\underline{b}|)\ (\lambda_3\lambda_4\,|H|\,|\underline{c}\,\underline{d}|)}{(\lambda_1\lambda_3\,|H|\,|\underline{a}\,\underline{c}|)\ (\lambda_2\lambda_4\,|H|\,|\underline{b}\,\underline{d}|)} \\
&= \frac{|\underline{a}\,\underline{b}|\ |\underline{c}\,\underline{d}|}{|\underline{a}\,\underline{c}|\ |\underline{b}\,\underline{d}|} \\
&= \text{Cross}(\underline{a}, \underline{b}, \underline{c}, \underline{d}).
\end{aligned}$$

Suppose $\underline{a}$, $\underline{b}$ and $\underline{c}$ are collinear points in the world reference frame, i.e. they fall on a line in $\mathbb{R}^3$, and suppose further that the ratio of the two intervals defined by these three points is known. If the images of these points, $\underline{a}'$, $\underline{b}'$ and $\underline{c}'$ can be identified, the projective map between the two lines can be calculated. Using this map, the point at infinity on the line is mapped to its image. More specifically, if the length ratio is given by $\frac{d(\underline{a},\underline{b})}{d(\underline{b},\underline{c})} = \frac{m}{n}$ (where $d(\underline{a}, \underline{b})$ is the Euclidean

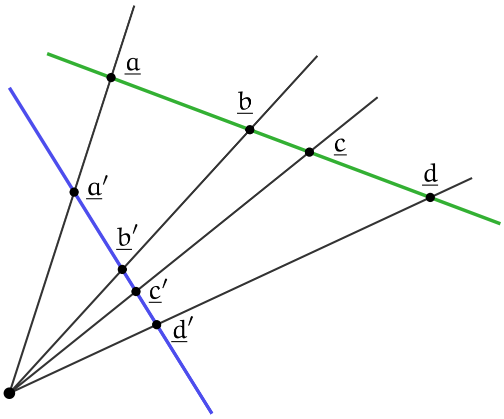

Figure 3.8: Invariance of the cross ratio under a projective transformation.

distance between $\underline{a}$ and $\underline{b}$), one calculates the image of the point at infinity using the following procedure.

1. Measure the distance ratio in the image, $\frac{d(\underline{a}',\underline{b}')}{d(\underline{b}',\underline{c}')} = \frac{m'}{n'}$.

2. Represent the points $\underline{a}$, $\underline{b}$ and $\underline{c}$ as 0, m and $m+n$ in a coordinate frame on the line going through $\underline{a}$, $\underline{b}$ and $\underline{c}$. These points may be represented by homogeneous 2-vectors $\begin{bmatrix} 0 & 1 \end{bmatrix}^T$, $\begin{bmatrix} m & 1 \end{bmatrix}^T$ and $\begin{bmatrix} m+n & 1 \end{bmatrix}^T$. Similarly, $\underline{a}'$, $\underline{b}'$ and $\underline{c}'$ can be represented by the homogeneous 2-vectors $\begin{bmatrix} 0 & 1 \end{bmatrix}^T$, $\begin{bmatrix} m' & 1 \end{bmatrix}^T$ and $\begin{bmatrix} m'+n' & 1 \end{bmatrix}^T$.

3. Relative to these coordinate frames, compute the 1D projective transformation H_2 that maps $\underline{a}$ to $\underline{a}'$, $\underline{b}$ to $\underline{b}'$, and $\underline{c}$ to $\underline{c}'$.

4. Calculate the image of the ideal point $\begin{bmatrix} 1 & 0 \end{bmatrix}^T$ under H_2. It is the image of the vanishing point on the line through $\underline{a}'$, $\underline{b}'$ and $\underline{c}'$.

How would you find the image of a point at infinity using the cross ratio, assuming you know the world coordinates of three collinear points, as well as their images, as above?