A PRIMER IN PROJECTIVE GEOMETRY

Projective geometry can be approached in many ways.

It can be conceived as an *extension* of Euclidean geometry, by adding points where *lines that never meet* (parallels) *do meet* (points at infinity). The construction of projective geometry can be pursued synthetically by altering the axioms of Euclidean geometry. In particular, projective geometry arises only from incidence axioms; no congruence – with the underlying concept of metric – or ordering are involved. The most aesthetically pleasing approaches construct projective geometry in a purely axiomatic way, and allow Euclidean geometry as a subcase. This approach is preferred by geometers, for it entails a unified view and closely relates the different geometries. We report here the 5 axioms (of existence and incidence) that constitute plane projective geometry:

1. There exists at least one line.
2. For each line there exist at least three points.
3. Not all points lie on the same line.
4. Two distinct points lie on one and only line.
5. Two distinct lines meet in one and only point.

The last axiom substitutes the parallels axioms of plane Euclieean geometry, that states that, given a line l and a point p not on l, there exists *one and only one* line passing through p and never intersecting l. In projective geometry there exists no such line, for all lines meet at a point. In order to complete the picture, one could conceive the possibility of there existing *more than one* lines passing through p and not intersecting l. This is the case in hyperbolic geometry.

Note that, in the axioms, we have not defined what a *line* is and what a *point* is. Instead, we have implicitly considered two disjoint sets, the set of point $\mathcal{P}$ and the set of lines $\mathcal{L}$, and a relation between the two sets (intended as a subset of the cartesian product $\mathcal{P} \times \mathcal{L}$). Such a relation is called *incidence*: we say that a point p *lies* on the line l if the pair (p, l) belong to the incidence relation. We also have implicitly defined the relation *meet*, as the dual of *lies*. The axioms then specify properties of the incidence relation, and can be used to construct a projective plane. For instance, we can construct a "minimal" set of points and lines that satisfies the axioms. It is easy to show that the sets $\mathcal{P} = \{p_1, \ldots, p_7\}$ and $\mathcal{L} = \{l_1, \ldots, l_7\}$ with the incidence relation defined by the table

l_1	l_2	l_3	l_4	l_5	l_6	l_7
p_1	p_2	p_3	p_4	p_5	p_6	p_7
p_2	p_3	p_4	p_5	p_6	p_7	p_1
p_4	p_5	p_6	p_7	p_1	p_2	p_3

(1)

satisfies the axioms. Here a point p_i is considered to lie on the line l_j if P_i appears in the column corresponding to l_j. It is possible to prove that a projective plane has at least 7 points, so the above construction, which is called a *Fanian plane*, is in this sense "minimal".

In the sequel we will be mostly interested in projective planes where the incidence relation can be characterized *algebraically*, as a relation between numbers of a division ring. Such an example is the so-called "real projective plane" $\mathbb{R}P^2$, often indicated with $\mathcal{P}^2$, which we describe below.

We will prefer an *analytic* approach to projective geometry. The reason is that (in these notes) we do not study it for the aesthetic pleasure, but because it is useful and we will need to perform calculations with it. We will keep as much as possible the imaging process in mind: the deformations of planar objects as seen from an ideal camera can be modeled as *projective transformations* on the plane $^{\rightarrow}$. Therefore we will reason in terms of *coordinate sets* and *allowable transformations* on that set. The group of allowable transformations acting on the set of objects of interest may be classified depending upon which properties are preserved. For instance, the distance between points and the angles between vectors in 3-D Euclidean space is invariant under the action of a rigid motion.

The combination of a set and a group of allowable transformations acting on it determines the *geometry*. For instance, the real plane, along with the action of rigid motions, is the object of plane Euclidean geometry. The classification of geometries based upon the action of a group of allowable transformations on a set was introduced by F. Klein, in order to establish a connection between geometry and group theory.

In this chapter we are going to describe the basic principles of plane projective geometry and euclidean 3-D geometry in a fairly unformal manner. We aim at giving a general intuitive grasp of the subject, and refer the reader to more orthodox literature for a rigorous treatement of the subject.

We will also make wide use of the notion of *embedding* $\rightarrow$ by viewing what happens on a particular space from a bigger and simpler one. For instance, we will work on the projective plane by *immerging* it into the euclidean 3-D space. The wide use of *homogeneous coordinates* follows this intent; the aim is both to help the reader in developing intuition and to use the wide body of intruments in linear algebra that are typical heritage of an engineering background.

1 Coordinates on the projective plane

[figure here] Consider four coplanar points in 3-D space, disposed as to form a quare of side one. When we view such a square with a camera, it is very unlikely that its image will resemble a square. In particular, opposite sides are no longer parallel, nor have equal size. The class of plane projective transformations describes exactly all the possible deformations a planar set of 3-D points can undergo when viewed from an ideal camera (a camera that performs an ideal perspective projection). Since the notion of parallelism has no significance under perspective projection (for it is not conserved), we may uniform parallel lines to incident lines by saying that the former intersect at their "point at infinity".

> We need to find a way of formally "adding" points at infinity to the Euclidean plane. We may think such an operation as "inserting a flag" to the coordinates of each point of the plane. Intuitively, consider the the points of the Euclidean plane, represented using $[x, y]^T \in \mathbb{R}^2$ that belong to two lines with coefficients $[\alpha, \beta, \gamma]$, and $[\alpha', \beta', \delta]$. Such points satisfy the following equations
>
> $$\begin{aligned} \alpha x + \beta y + \gamma &= 0 \\ \alpha' x + \beta' y + \delta &= 0 \end{aligned}$$
>
> In order to find the coordinates of the intersecting point, we just solve the above system of equations. If the two lines are parallel, then we have $\alpha = \alpha', \beta = \beta'$ and
>
> $$\gamma \neq \delta. \tag{2}$$
>
> Therefore, by solving the above system, we would have an inconsistent result
>
> $$\gamma = \delta. \tag{3}$$
>
> This would not happen if we added a "flag" z, which could be either 0 or 1, that multiplies γ and δ:
>
> $$\begin{aligned} \alpha x + \beta y + \gamma z &= 0 \\ \alpha' x + \beta' y + \delta z &= 0. \end{aligned} \tag{4}$$
>
> If the two lines are incident, we just set $z = 1$ and we fall in the usual Euclidean case. If the lines are parallel, then we have the equation
>
> $$(\gamma - \delta) z = 0 \tag{5}$$
>
> which i solved for $z = 0$, since $\gamma \neq \delta$ unless the two lines are coincident.

From this naïf argument, it looks like, to specify the coordinates of a point p on the projective plane, we actually need three numbers, $[x, y, z]^T$. Such coordinates are the so-called *"homogeneous coordinates"*. It is also evident that, if z is allowed to be any non-zero real number $k \in \mathbb{R}$ rather than just 0 or 1, we could represent the point p using the coordinates $[kx, ky, k]^T$, since the equations in (4) are *homogeneous*, and therefore

$$\alpha kx + \beta ky + \gamma k = k(\alpha x + \beta y + \gamma) = 0. \tag{6}$$

Therefore − intuitively − in order to represent a point in the projective plane, we use three coordinates, but we agree that scaled versions of these coordinates also represent the same point. We must also exclude the origin $[0, 0, 0]^T$ from the projective plane, for all lines would pass through it, violating the axioms of incidence.

The above discussion can be summarized in the simple statement

$$\mathcal{P}^2 \doteq \mathbb{R}^3 - \{0\} \backslash \mathbb{R}. \tag{7}$$

This statement is indeed a definition, and the discussion above is only ment as an intuitive reason why such a definition is appropriate. Now that we have this statement, we can investigate alternative ways of interpreting the real projective plane.

2 Models of the real projective plane

It is easy to think at adding the points at infinity by wrapping the Euclidean plane around a sphere. Lines are then represented as great circles, and they always intersect. [figure here] As a matter of fact, they always intersect *twice*. Therefore, we must agree that – in this model – opposite points coincide. We can view this process as a simple representation of the projective plane using its definition (7). Since points in the projective plane are represented as points along a *ray* (a line) in $\mathbb{R}^3$, and all points on such a line equivalently represent a point in the projective plane, we may as well agree on chosing a particular one. For instance, we can choose the intersection of a ray in $\mathbb{R}^3$ with [figure here] the unit-sphere $\mathbf{S}^2 = \{x \in \mathbb{R}^3 \mid \|x\| = 1\}$. Of course such an intersection consists of two points, so we must either agree on one, or consider the two equivalent. Points at infinity correspond to the border of the half-sphere.

[figure here] Alternatively, we can consider as representative of a point in the projective plane the intersection of a ray in $\mathbb{R}^3$ with a plane. For instance, the plane $\{[x_1, x_2, x_3]^T \in \mathbb{R}^3 | x_3 = 1\}$. But this plane is nothing else than the Euclidean plane. Indeed we must also consider rays that do not intersect such a plane, for they have coordinates $k[x_1, x_2, 0]^T$. These are exactly the points at infinity. Such a model is sometimes referred to in the literature as the "ray-space model".

The ray-space is the model we are going to use for the image formation process. Note that, although an image is naturally bounded and therefore points at infinity cannot be observed, points that are visible (and therefore finite) at some time-instant can be tranformed into points at infinity under the motion of the camera.

3 Representations of general projective spaces

A simple way of generalizing the intuitive description of the projective plane in terms of "ray-space" (embedding space) consists in defining the general projective space associated with a given vector space V as the set of its one-dimensional subspaces (rays):

Definition 3.1 *Let V be a finite-dimensional vector space over the field K. $\mathcal{P}(V)$, the projective space associated with V is the set of all vector subspace of V of dimension one.*

We will deal mostly with the real projective plane: $V = \mathbb{R}^3$, $\mathcal{P}(\mathbb{R}^3) \doteq \mathcal{P}^2$. If V is an n-dimensional vector space, $\mathcal{P}(V)$ is $n-1$-dimensional.

Note that each element $P \in V$ generates a subspace of dimension one via

$$< P >= \{\lambda P \mid \lambda \in \mathbb{R} - \{0\}\} \tag{8}$$

and the coordinates of the point $P = \mathbf{X}$ determine the equivalence class

$$[\mathbf{X}] = \{\lambda \mathbf{X} \mid \lambda \in \mathbb{R} - \{0\}\} \tag{9}$$

which identifies nothing else than the *homogeneous (or projective) coordinates* of the point P interpreted as an element of the projective space $\mathcal{P}(V)$.

4 Projective basis

We have seen that a point in a projective space can be represented by an equivalence class of coordinates in the *embedding* space. For instance, a point p in the real projective plane $\mathcal{P}^2$ can be represented by the coordinates of a point P of $\mathbb{R}^3$, where it is agreed that all scalar multiples of such coordinates also identify the same point. In general, we may represented a point on an n-dimensional projective space using homogeneous coordinates, which is, cartesian coordinates on the $n+1$-dimensional Euclidean space, defined up to a scaling factor.

Therefore, it looks like, in order to define a basis for $\mathcal{P}^n$, we need $n+1$ *Euclidean points*[1] in E^{n+1}. If we consider the basis of E^{n+1}:

$$E^{n+1} = < \mathbf{e}_1, \dots, \mathbf{e}_n \ , \ \mathbf{e}_{n+1} > \tag{10}$$

where e_i has canonical coordinates $[0 \dots 1_i \dots 0]^T$ then a point $p \in \mathcal{P}^n$ can be written as

$$p \sim x_1\mathbf{e}_1 + \dots + x_{n+1}\mathbf{e}_{n+1} \tag{11}$$

and represented using a vector of homogeneous coordinates $[\mathbf{x}] = [x_1, \dots, x_{n+1}]^T$. A change of coordinates in $\mathcal{P}^n$ can be accomplished simply as a change of homogeneous coordinates in E^{n+1}, by a map represented as a generic $n+1 \times n+1$ invertible matrix M:

$$[\mathbf{x}'] \sim M[\mathbf{x}] \qquad\qquad M \in \mathcal{GL}(n+1) \tag{12}$$

where the equality is intended, as usual, up to a scaling factor. $\mathcal{GL}(n)$ is the general linear group of $n \times n$ non-singular matrices.

We have just stated that in order to define a basis of the $n-$dimensional *projective space* we need $n+1$ *Euclidean points*. But since we usually want to define the basis of a space using elements of *that same* space, the question arises naturally of *how many projective points we need in order to define a projective basis*. The question resorts to addressing how many projective points we need in order to define a linear nonsingular map M up to a scale factor. Suppose we are given pairs of homogeneous coordinates $[\mathbf{x}^i]$ and $[\mathbf{x}^{i'}]$ as in (12) and, from a number $i = 1 \dots m$ of pairs, we want to determine the elements of M. For each pair, a relation holds of the form

$$\lambda_i' \mathbf{x}^{i'} = M \lambda_i \mathbf{x}^i \tag{13}$$

where we have made explicit the arbitrary scaling factors. The matrix M is itself defined up to a scaling: αM. We may collect the effects of these three arbitrary scales into just one parameter $\mu_i = \frac{\lambda_i'}{\alpha \lambda_i}$. Given m pairs of coordinates, we can then write a system of linear equations in the form

$$\mu_i \mathbf{x}^{i'} = M \mathbf{x}^i \qquad\qquad \forall\ i = 1 \dots m \tag{14}$$

which has $m + (n+1)^2 - 1$ unknowns (the scaling factors μ_i and the elements of M up to a scale) and $m(n+1)$ equations. In order for these equations to admit a solution for the elements of the matrix M, we need at least $m = n+2$. In the case of the projective plane, therefore, we need at least four projective points. It is easily shown that, provided that no $n+1$ of these points are *collinear* (i.e. their homogeneous coordinates are independent in $\mathbb{R}^{n+1}$), $n+2$ points are also sufficient in order to determine a basis.

Therefore a basis of the projective plane consists of four points, not three of them collinear. The canonical choice consists in adding to the canonical coordinates of the Euclidean embedding space E^{n+1} the so-called *unity point* $\mathbf{u}$, with coordinates $[1 \dots 1]^T$:

$$\mathcal{P}^n = < \mathbf{e}_1, \dots, \mathbf{e}_n \ , \ \mathbf{e}_{n+1} \ , \ \mathbf{u} > . \tag{15}$$

From now on we will omit the brackets in writing homogeneous coordinates $[\mathbf{X}]$ when it is clear from the context that they represent a projective point.

[1] In this paragraph we indicate with E^n the Euclidean space of dimension n, modeled by $\mathbb{R}^n$, when we want to emphasize the Euclidean structure.

5 Projective transformations

A projective transformation is a linear nonsingular transformation of the projective coordinates. The names "projective transformation", "homography", "collineation" and "projectivity" are all equivalent. A projective transformation of an $n-$dimensional projective space can be represented by a non-singular $n+1 \times n+1$ matrix M, that transforms a point with projective coordinates $[\mathbf{x}]$ into $[\mathbf{x}']$ via

$$\mathbf{x}' \sim M\mathbf{x} \tag{16}$$

where $\sim$ indicates equality up to a scale factor.

5.1 Euclidean, affine and projective transformations of the plane

A Euclidean transformation of the plane is a transformation that does not change the distance between any two points. Such transformations can always be decomposed into a translation and a rotation, so that a point of coordinates (x, y) is transformed into

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = R \begin{bmatrix} x \\ y \end{bmatrix} + \mathbf{T} \tag{17}$$

where $R = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$ is a rotation matrix and $\mathbf{T} = [T_x \; T_y]$ is a translation vector.

Rotation matrices are a very special subset of all possible 2×2. If instead of forcing R to have the structure of a rotation matrix, i.e. that $RR^T = R^T R = I$, we allow R to be generic, we obtain an *affine* transformation:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = A \begin{bmatrix} x \\ y \end{bmatrix} + \mathbf{T} \tag{18}$$

where A is a nonsingular 2×2 matrix. Therefore, the set of Euclidean transformations of the plane are a subset of affine transformations. In order to see that affine (and therefore also Euclidean) transformations of the plane are then a subset of projective transformations, we embed the affine plane into the projective plane by considering the projective coordinates $[\mathbf{x}] = [x \; y \; 1]^T$, so that the affine (or Euclidean) transformation can be written as

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} A & \mathbf{T} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \doteq M \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \tag{19}$$

where the matrix M has the special structure above with A a non-singular matrix for an affine transformation, or a rotation matrix for a Euclidean transformation.

5.2 Geometric interpretation of plane-projective transformations

In section 2 we have described the "ray-space" as a model of the projective plane: points in the projective plane are represented as one-dimensional subspace of the three-dimensional Euclidean space. Since each subspace (ray) identifies an equivalence class (the homogeneous coordinate), we can choose one element as representative of the class, for instance the intersection of the rays with a plane. In doing so, we have to be careful to add the points at infinity, represented by subspaces parallel to the chosen plane.

Of course, if we choose a different plane, we get different points. More in general, by moving the plane rigidly in three-dimensional space, we induce a transformation on the intersection of the rays with the plane. We now show that such a transformation is a projective transformation. If we choose a reference frame in the three-dimensional Euclidean space, with the origin in the center of the rays, we can describe a plane not passing through the origin as

$$\Pi = \{\mathbf{X} \in \mathbb{R}^3 \mid \nu^T \mathbf{X} = 1\} \tag{20}$$

where ν is the normal vector to the plane, and $\|\nu\|$ is the distance of the plane from the origin. Any other plane is obtained by rigidly moving Π:

$$\Pi' = \{R\mathbf{X} + \mathbf{T} \in \mathbb{R}^3 \mid \mathbf{X} \in \Pi\} \tag{21}$$

where R is a 3×3 rotation matrix, such that $RR^T = R^T R = I$, and $\mathbf{T}$ is a three-dimensional translation vector. We now consider a point $\mathbf{X}' \in \Pi'$, and write the tautology

$$\mathbf{X}' = R\mathbf{X} + \mathbf{T}\ 1 = R\mathbf{X} + \mathbf{T}\nu^T\mathbf{X} = (R + \mathbf{T}\nu^T)\mathbf{X} \tag{22}$$

which we can write, after having defined $M \doteq R + \mathbf{T}\nu^T$, as

$$\mathbf{X}' = M\mathbf{X}. \tag{23}$$

Here M is non-singular, as a consequence of the fact that $\det(R) = 1$. If we now interpret the above equation up to a scale factor, we get that $\mathbf{X} \doteq [\mathbf{x}]$ represent the homogeneous coordinates of a point in the projective plane, and $\mathbf{X}' \doteq [\mathbf{x}']$ represent its transformed via

$$[\mathbf{x}'] \sim M[\mathbf{x}]. \tag{24}$$

Therefore, every rigid motion of a plane induces a projective transformation of the ray-space model.

What is less obvious to see is that every projective transformation can be interpreted as a rigid motion of a plane in the ray-space model. This comes as a consequence of a result that can be state as follows:

Claim

For any non-singular 3×3 matrix M there exists an orthonormal matrix R, a scalar α and a rank-one matrix $\mathbf{T}\nu^T$, $\mathbf{T}, \nu \in \mathbb{R}^3$ such that

$$M = \alpha R + \mathbf{T}\nu^T \tag{25}$$