Projective geometry for Computer Vision

Subhashis Banerjee

Department of Computer Science and Engineering
IIT Delhi

NIT, Rourkela
March 27, 2010

Overview

- Pin-hole camera
- Why projective geometry?
- Reconstruction

Computer vision geometry: main problems

Correspondence problem: Match image projections of a 3D configuration.

Reconstruction problem: Recover the structure of the 3D configuration from image projections.

Re-projection problem: Is a novel view of a 3D configuration consistent with other views? (Novel view generation)

An infinitely strange perspective

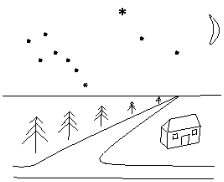

Figure 1.1: Landscape with horizon

- Parallel lines in 3D space converge in images.
- The line of the horizon is formed by 'infinitely' distant points (vanishing points).
- Any pair of parallel lines meet at a point on the horizon corresponding to their common direction.
- All 'intersections at infinity' stay constant as the observer moves.

3D reconstruction from pin-hole projections

La Flagellazione di Cristo (1460) Galleria Nazionale delle Marche by Piero della Francesca (1416-1492) (Robotics Research Group, Oxford University, 2000)

Pin-hole camera

- The effects can be modelled mathematically using the 'linear perspective' or a 'pin-hole camera' (realized first by Leonardo?)

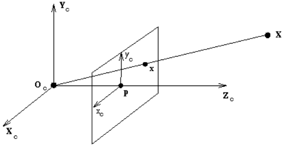

- If the world coordinates of a point are (X, Y, Z) and the image coordinates are (x, y), then

$$x = fX/Z \text{ and } y = fY/Z$$

- **The model is non-linear**.

In terms of projective coordinates

$$\lambda \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

where,

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \in \mathcal{P}^2 \text{ and } \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \in \mathcal{P}^3$$

are **homogeneous coordinates**.

Euclidean and Affine geometries

- Given a coordinate system, n-dimensional real **affine space** is the set of all points parameterized by $\mathbf{x} = (x_1, \ldots, x_n)^t \in \mathbb{R}^n$.
- An affine transformation is expressed as

$$\mathbf{x}' = \mathbf{A}\mathbf{x} + \mathbf{b}$$

where $\mathbf{A}$ is a $n \times n$ (usually) non-singular matrix and $\mathbf{b}$ is a $n \times 1$ vector representing a translation.
- By *SVD*

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = (\mathbf{U}\mathbf{V}^T)(\mathbf{V}\mathbf{\Sigma}\mathbf{V}^T) = R(\theta)R(-\phi)\mathbf{\Sigma}R(\phi)$$

where where

$$\mathbf{\Sigma} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

Euclidean and Affine geometries

- In the special case of when $\mathbf{A}$ is a rotation (i.e., $\mathbf{A}\mathbf{A}^t = \mathbf{A}^t\mathbf{A} = \mathbf{I}$, then the transformation is *Euclidean*.
- An affine transformation preserves parallelism and ratios of lengths along parallel directions.
- An Euclidean transformation, in addition to the above, also preserves lengths and angles.
- *Since an affine (or Euclidean) transformation preserves parallelism it cannot be used to describe a pinhole projection*.

Spherical geometry

- **The space $\mathcal{S}^2$:**

$$\mathcal{S}^2 = \left\{\mathbf{x} \in \mathbb{R}^3 : ||x|| = 1\right\}$$

- **lines in $\mathcal{S}^2$:** Viewed as a set in $\mathbb{R}^3$ this is the intersection of $\mathcal{S}^2$ with a plane through the origin. We will call this great circle a line in $\mathcal{S}^2$. Let ξ be a unit vector. Then, $\mathbf{l} = \left\{\mathbf{x} \in \mathcal{S}^2 : \xi^t \mathbf{x} = 0\right\}$ is the line with pole ξ.

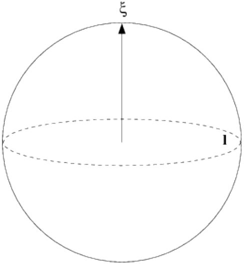

Spherical geometry

- Lines in $\mathcal{S}^2$ cannot be parallel. Any two lines intersect at a pair of antipodal points.
- A point on a line:

$$\mathbf{l} \cdot \mathbf{x} = 0 \text{ or } \mathbf{l}^T \mathbf{x} = 0 \text{ or } \mathbf{x}^T \mathbf{l} = 0$$

- Two points define a line:

$$\mathbf{l} = \mathbf{p} \times \mathbf{q}$$

- Two lines define a point:

$$\mathbf{x} = \mathbf{l} \times \mathbf{m}$$

Projective geometry

- *The projective plane $\mathcal{P}^2$ is the set of all pairs $\{\mathbf{x}, -\mathbf{x}\}$ of antipodal points in $\mathcal{S}^2$.*
- Two alternative definitions of $\mathcal{P}^2$, equivalent to the preceding one are
  1. The set of all lines through the origin in $\mathbb{R}^3$.
  2. The set of all equivalence classes of ordered triples (x_1, x_2, x_3) of numbers (i.e., vectors in $\mathbb{R}^3$) not all zero, where two vectors are equivalent if they are proportional.

Projective geometry

The space $\mathcal{P}^2$ can be thought of as the infinite plane tangent to the space $\mathcal{S}^2$ and passing through the point $(0, 0, 1)^t$.

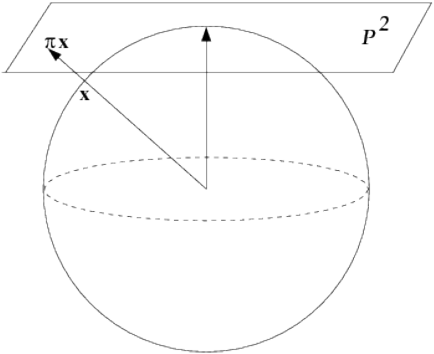

Projective geometry

- Let $\pi : \mathcal{S}^2 \rightarrow \mathcal{P}^2$ be the mapping that sends $\mathbf{x}$ to $\{\mathbf{x}, -\mathbf{x}\}$. The π is a two-to-one map of $\mathcal{S}^2$ onto $\mathcal{P}^2$.
- A line of $\mathcal{P}^2$ is a set of the form $\pi\mathbf{l}$, where $\mathbf{l}$ is a line of $\mathcal{S}^2$. Clearly, $\pi\mathbf{x}$ lies on $\pi\mathbf{l}$ if and only if $\xi^t\mathbf{x} = 0$.
- **Homogeneous coordinates:** In general, points of real n-dimensional **projective space**, $\mathcal{P}^n$, are represented by $n+1$ component column vectors $(x_1, \ldots, x_n, x_{n+1}) \in \mathbb{R}^{n+1}$ such that at least one x_i is non-zero and $(x_1, \ldots, x_n, x_{n+1})$ and $(\lambda x_1, \ldots, \lambda x_n, \lambda x_{n+1})$ represent the same point of $\mathcal{P}^n$ for all $\lambda \neq 0$.
- $(x_1, \ldots, x_n, x_{n+1})$ is the homogeneous representation of a projective point.

Canonical injection of $\mathbb{R}^n$ into $\mathcal{P}^n$

- Affine space $\mathbb{R}^n$ can be embedded in $\mathcal{P}^n$ by

$$(x_1, \ldots, x_n) \to (x_1, \ldots, x_n, 1)$$

- Affine points can be recovered from projective points with $x_{n+1} \neq 0$ by

$$(x_1, \ldots, x_n) \sim (\frac{x_1}{x_{n+1}}, \ldots, \frac{x_n}{x_{n+1}}, 1) \to (\frac{x_1}{x_{n+1}}, \ldots, \frac{x_n}{x_{n+1}})$$

- A projective point with $x_{n+1} = 0$ corresponds to a **point at infinity**.
- The ray $(x_1, \ldots, x_n, 0)$ can be viewed as an additional **ideal point** as $(x_1, \ldots, x_n)$ recedes to infinity in a certain direction. For example, in $\mathcal{P}^2$,

$$\lim_{T \to 0} (X/T, Y/T, 1) = \lim_{T \to 0} (X, Y, T) = (X, Y, 0)$$

Lines in $\mathcal{P}^2$

- A line equation in $\mathbb{R}^2$ is

$$a_1 x_1 + a_2 x_2 + a_3 = 0$$

- Substituting by homogeneous coordinates $x_i = X_i / X_3$ we get a homogeneous linear equation

$$(a_1, a_2, a_3) \cdot (X_1, X_2, X_3) = \sum_{i=1}^{3} a_i X_i = 0, \ \mathbf{X} \in \mathcal{P}^2$$

- A line in $\mathcal{P}^2$ is represented by a homogeneous 3-vector (a_1, a_2, a_3).
- A point on a line: $\mathbf{a} \cdot \mathbf{X} = 0$ or $\mathbf{a}^T \mathbf{X} = 0$ or $\mathbf{X}^T \mathbf{a} = 0$
- Two points define a line: $\mathbf{l} = \mathbf{p} \times \mathbf{q}$
- Two lines define a point: $\mathbf{x} = \mathbf{l} \times \mathbf{m}$

The line at infinity

- The **line at infinity** ($\mathbf{l}_\infty$): is the line of equation $X_3 = 0$. Thus, the homogeneous representation of $\mathbf{l}_\infty$ is $(0, 0, 1)$.
- The line (u_1, u_2, u_3) intersects $\mathbf{l}_\infty$ at the point $(-u_2, u_1, 0)$.
- Points on $\mathbf{l}_\infty$ are directions of affine lines in the embedded affine space (can be extended to higher dimensions).

Conics in $\mathcal{P}^2$

A **conic** in affine space (inhomogeneous coordinates) is

$$ax^2 + bxy + cy^2 + dx + ey + f = 0$$

Homogenizing this by replacements $x = X_1/X_3$ and $y = Y_1/Y_3$, we obtain

$$aX_1^2 + bX_1X_2 + cX_2^2 + dX_1X_3 + eX_2X_3 + fX_3^2 = 0$$

which can be written in matrix notation as $\mathbf{X}^T\mathbf{C}\mathbf{X} = 0$ where C is symmetric and is the *homogeneous representation* of a **conic**.

Conics in $\mathcal{P}^2$

- The line $\mathbf{l}$ tangent to a conic $\mathbf{C}$ at any point $\mathbf{x}$ is given by $\mathbf{l} = \mathbf{C}\mathbf{x}$.
- $\mathbf{x}^t\mathbf{C}\mathbf{x} = 0 \implies (\mathbf{C}^{-1}\mathbf{l})^t\mathbf{C}((\mathbf{C}^{-1}\mathbf{l}) = \mathbf{l}^t\mathbf{C}^{-1}\mathbf{l} = 0$ (because $\mathbf{C}^{-t} = \mathbf{C}^{-1}$). This is the equation of the *dual conic*.

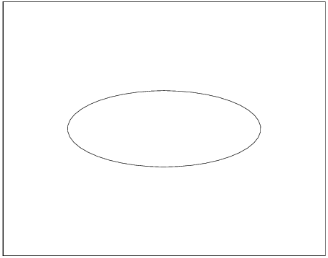

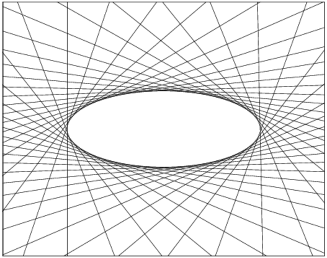

Conics in $\mathcal{P}^2$

- The *degenerate conic* of rank 2 is defined by two line $\mathbf{l}$ and $\mathbf{m}$ as

$$\mathbf{C} = \mathbf{l}\mathbf{m}^t + \mathbf{m}\mathbf{l}^t$$

Points on line $\mathbf{l}$ satisfy $\mathbf{l}^t\mathbf{x} = 0$ and are hence on the conic because $(\mathbf{x}^t\mathbf{l})(\mathbf{m}^t\mathbf{x}) + (\mathbf{x}^t\mathbf{m})(\mathbf{l}^t\mathbf{x}) = 0$. (Similarly for $\mathbf{m}$). The dual conic $\mathbf{x}\mathbf{y}^t + \mathbf{y}\mathbf{x}^t$ represents lines passing through $\mathbf{x}$ and $\mathbf{y}$.

Projective basis

Projective basis: A **projective basis** for $\mathcal{P}^n$ is any set of $n+2$ points no $n+1$ of which are linearly dependent.

Canonical basis:

$$\underbrace{\begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \ldots}_{\text{points at infinity along each axis}} \quad \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}}_{\text{origin}}, \underbrace{\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}}_{\text{unit point}}$$

Projective basis

Change of basis: Let $\mathbf{e_1}, \mathbf{e_2}, \ldots, \mathbf{e_{n+1}}, \mathbf{e_{n+2}}$ be the standard basis and $\mathbf{a_1}, \mathbf{a_2}, \ldots, \mathbf{a_{n+1}}, \mathbf{a_{n+2}}$ be any other basis. There exists a non-singular transformation $[\mathbf{T}]_{(n+1)\times(n+1)}$ such that:

$$\mathbf{Te_i} = \lambda_i \mathbf{a_i}, \forall i = 1, 2. \ldots, n+2$$

$\mathbf{T}$ is unique up to a scale.

Homography

The invertible transformation $\mathbf{T} : \mathcal{P}^n \rightarrow \mathcal{P}^n$ is called a **projective transformation** or **collineation** or **homography** or **perspectivity** and is completely determined by $n + 2$ point correspondences.

- Preserves straight lines and cross ratios
- Given four collinear points $\mathbf{A}_1$, $\mathbf{A}_2$,$\mathbf{A}_3$ and $\mathbf{A}_4$, their **cross ratio** is defined as

$$\frac{\overline{\mathbf{A}_1\mathbf{A}_3}\ \overline{\mathbf{A}_2\mathbf{A}_4}}{\overline{\mathbf{A}_1\mathbf{A}_4}\ \overline{\mathbf{A}_2\mathbf{A}_3}}$$

- If $\mathbf{A}_4$ is a point at infinity then the cross ratio is given as

$$\frac{\overline{\mathbf{A}_1\mathbf{A}_3}}{\overline{\mathbf{A}_2\mathbf{A}_3}}$$

- The cross ratio is independent of the choice of the projective coordinate system.

Homography

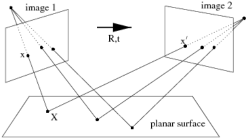

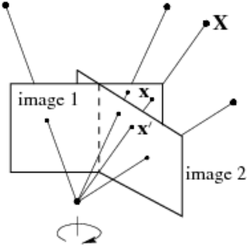

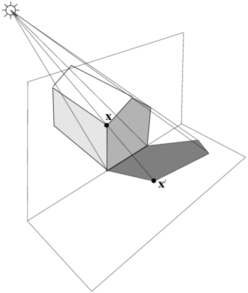

Homography

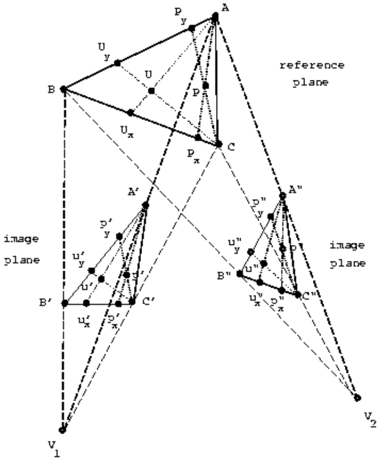

A
P y
U y
U
B
P
reference
plane
U x
C
P x
A'
A"
p" y
p' y
image
plane
u" y
p"
image
plane
u' y
u'
p'
B"
u"
C"
B'
C'
u" x
p" x
u' x
p' x
V 1
V 2

Projective mappings of lines

- If the points $\mathbf{x}_i$ lie on the line $\mathbf{l}$, we have $\mathbf{l}^T\mathbf{x}_i = 0$.
- Since, $\mathbf{l}^T\mathbf{H}^{-1}\mathbf{H}\mathbf{x}_i = 0$ the points $\mathbf{H}\mathbf{x}_i$ all lie on the line $\mathbf{H}^{-T}\mathbf{l}$.
- Hence, if points are transformed as $\mathbf{x}'_i = \mathbf{H}\mathbf{x}_i$, lines are transformed as $\mathbf{l}' = \mathbf{H}^{-T}\mathbf{l}$.

Projective mappings of conics

- Note that a conic is represented (homogeneously) as

$$\mathbf{x}^T \mathbf{C} \mathbf{x} = 0$$

- Under a point transformation $\mathbf{x}' = \mathbf{H}\mathbf{x}$ the conic becomes

$$\mathbf{x}^T \mathbf{C} \mathbf{x} = \mathbf{x}'^T [\mathbf{H}^{-1}]^T \mathbf{C} \mathbf{H}^{-1} \mathbf{x}' = \mathbf{x}'^T \mathbf{H}^{-T} \mathbf{C} \mathbf{H}^{-1} \mathbf{x}' = 0$$

- This is the quadratic form of $\mathbf{x}'^T \mathbf{C}' \mathbf{x}'$ with $\mathbf{C}' = \mathbf{H}^{-T} \mathbf{C} \mathbf{H}^{-1}$. This gives the transformation rule for a conic.

The affine subgroup

In an affine space $\mathcal{A}^n$ an **affine transformation** defines a correspondence $\mathbf{X} \leftrightarrow \mathbf{X}'$ given by:

$$\mathbf{X}' = \mathbf{A}\mathbf{X} + \mathbf{b}$$

where $\mathbf{X}$, $\mathbf{X}'$ and $\mathbf{b}$ are n-vectors, and $\mathbf{A}$ is an $n \times n$ matrix. Clearly this is a subgroup of the projective group. Its projective representation is

$$\mathbf{T} = \begin{bmatrix} \mathbf{C} & \mathbf{c} \\ \mathbf{0}_n^T & t_{33} \end{bmatrix}$$

where $\mathbf{A} = \frac{1}{t_{33}}\mathbf{C}$ and $\mathbf{b} = \frac{1}{t_{33}}\mathbf{c}$.

Affine transformations preserve the plane/line at infinity

$$\begin{bmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{0}^t & 1 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ 0 \end{pmatrix} = \begin{pmatrix} \mathbf{A} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \\ 0 \end{pmatrix}$$

A general projective transformation moves points at infinity to finite points.

The Euclidean subgroup

- The absolute conic: The conic $\mathbf{\Omega}_\infty$ is intersection of the quadric of equation:

$$\sum_{i=1}^{n+1} x_i^2 = x_{n+1} = 0 \text{ with } \pi_\infty$$

- In a metric frame $\pi_\infty = (0,0,0,1)^T$, and points on $\mathbf{\Omega}_\infty$ satisfy

$$\left.\begin{array}{r} X_1^2 + X_2^2 + X_3^3 \\ X_4 \end{array}\right\} = 0$$

- For directions on π_∞ (with $X_4 = 0$), the absolute conic $\mathbf{\Omega}_\infty$ can be expressed as

$$(X_1, X_2, X_3)\mathbf{I}(X_1, X_2, X_3)^T = 0$$

- **The absolute conic, $\mathbf{\Omega}_\infty$, is fixed under a projective transformation H if and only if H is an Euclidean transformation.**

Affine calibration of a plane

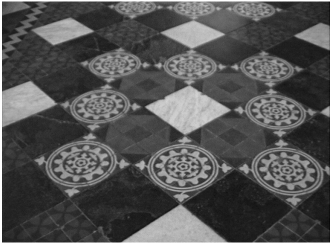

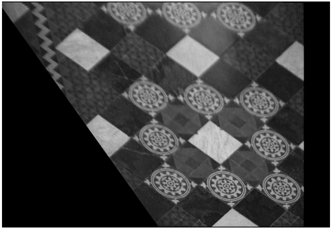

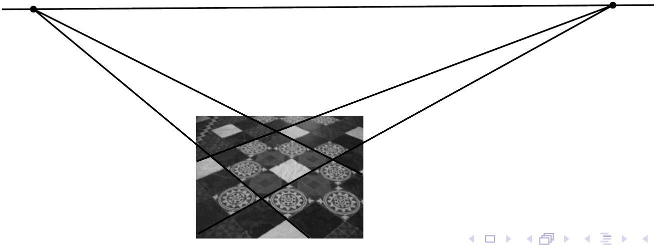

Affine calibration of a plane

If the imaged line at infinity is $\mathbf{l} = (l_1, l_2, l_3)^t$, then provided $l_3 \neq 0$ a suitable projective transformation that maps $\mathbf{l}$ back to $\mathbf{l}_\infty$ is

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ l_1 & l_2 & l_3 \end{bmatrix} \mathbf{H_A}$$

Reconstruction

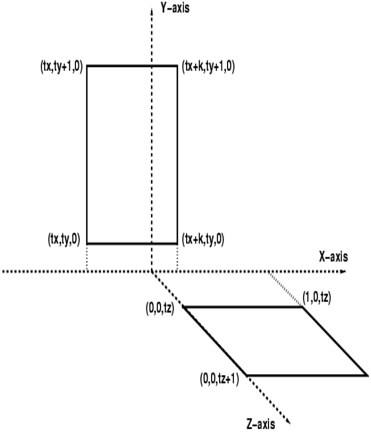

Camera recovery

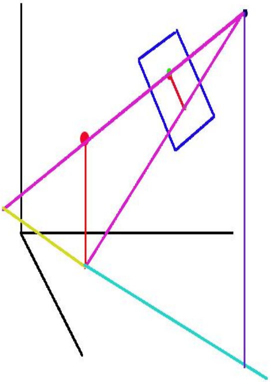

Metrology

Surfaces of revolution

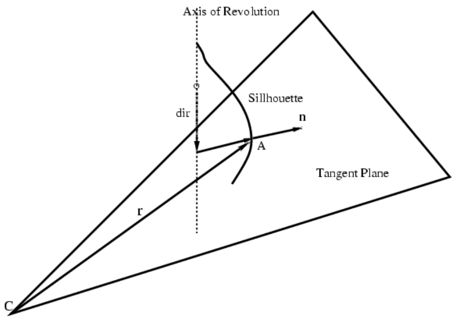

Modeling of structured scenes

A walkthrough