Geometry in Computer Vision

Spring 2010

Lecture 1
Projective Geometry

A vector space

- A vector space *V* consists of a set of vectors
  - Two vectors can be added
  - A vector can be multiplied by a scalar
  - Both operations result again in a vector in *V*
  - Sets of vectors can be linearly combined into a new vector
- The dimension of *V* =
  maximal number of vectors which are linear independent
- Basis exists
- Orthogonality between two vectors defined if we have a scalar product
- Linear mappings are well-defined

A projective space

- A projective space can be defined from *V* in terms of equivalence classes:
  - Two vectors **u** and **v** are *equivalent* if there exists a non-zero scalar *s* such that **u** = *s* **v**
    $\Rightarrow$ **u** and **v** must be non-zero vectors
  - All vectors which are equivalent correspond to an element of the projective space
    (*a projective element*)
  - Projective equivalence is denoted $\mathbf{u} \sim \mathbf{v}$

A projective space

- The projective space generated from *V* consist of all such projective elements
  - Any projective element correspond to a 1D subspace of *V*
  - Any projective element has a non-unique representation of non-zero vectors in *V*
  - Any non-zero element of *V* corresponds to a unique projective element
- The projective space is (here) denoted *P(V)*

A projective space

- Dimension of $P(V) = \dim(V) - 1$
- Addition and scalar multiplications are undefined operations in *P*(V), no linear combinations
- No basis exists
- Orthogonality is well-defined!
  - Two projective elements are orthogonal iff their representative vectors are orthogonal
- A *linear* mapping $\mathbf{M} : V \rightarrow U$ produces a well-defined mapping $P(V) \rightarrow P(U)$

Projective representation

- The *n*-dimensional vector space R^n can be given a projective representation by the projective space $P(R^{n+1})$

$$\mathbf{v} = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix} \in R^n$$

The homogeneous coordinates of **y**

Is represented by the projective element corresponding to

$$\mathbf{v}_H = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \\ 1 \end{pmatrix} \in R^{n+1}$$

Example

$$\mathbf{v} = \begin{pmatrix} 1 \\ 2 \end{pmatrix} \Rightarrow \mathbf{v}_H = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix} \sim \begin{pmatrix} 2 \\ 4 \\ 2 \end{pmatrix} \sim \begin{pmatrix} \frac{1}{2} \\ 1 \\ \frac{1}{2} \end{pmatrix}$$

All these vectors in R^3 represent the same projective element

Homogeneous normalization

- Given an a vector $\mathbf{u} \in R^{n+1}$ we can scale it so that the last element = 1 $\Rightarrow$ *normalization* (can we always do this?)
- The first *n* elements in the normalized homogeneous vector are the vector in R^n that **u** represents
- This makes it possible to know which vector in R^n a specific projective element in $P(R^{n+1})$ represents

Projective representation of the Euclidean space

- The elements of vectors in R^2 and R^3 are the coordinates of points in 2D or 3D Euclidean spaces relative to some specific coordinate systems
- We use the projective representation of R^2 given by $P^2 = P(R^3)$
- We use the projective representation of R^3 given by $P^3 = P(R^4)$

Projective representation of the Euclidean space

Motivation

- A corresponding representation can be found also for lines in 2D and lines + planes in 3D.
- Operations on these geometric object are <u>much easier</u> to describe algebraically in a projective space than in standard Euclidean coordinates
  - Find the point of intersection between a 3D plane and a 3D line
- "Exceptional cases" can be included in the same representations
  - Example: All 2D lines intersect at one point, except if the lines are parallel or identical

A homogeneous representation of lines in 2D

- Let **y** = (y_1,y_2) be the Euclidean coordinates of a 2D point
- Any 2D line is characterized by an angle α and a scalar L such that

y lies on the line $\Leftrightarrow$ $y_1 \cos\alpha + y_2 \sin\alpha = L$

The defining equation of a line

A homogeneous representation of lines in 2D

L and α are well-defined also for a vertical line.

Compare to $y_2 = k\, y_1 + a$

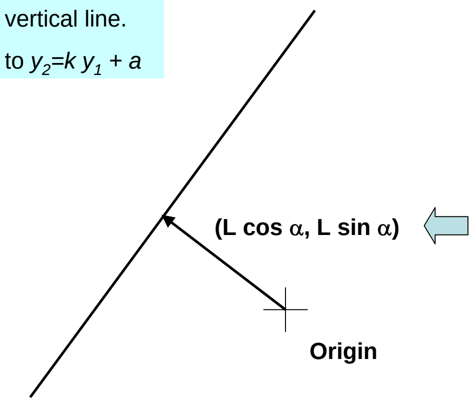

a normal vector which points from the origin to the line

L = shortest distance from the line to the origin

A homogeneous representation of lines in 2D

- **y** lies on the line ⇔ $y_1 \cos\alpha + y_2 \sin\alpha = L$

$$\mathbf{y} \text{ lies on the line} \Leftrightarrow \begin{pmatrix} y_1 \\ y_2 \\ 1 \end{pmatrix} \cdot \begin{pmatrix} \cos\alpha \\ \sin\alpha \\ -L \end{pmatrix} = 0$$

Homogeneous coordinates of **y**

A homogeneous representation of lines in 2D

- Suggests a homogeneous representation of the line:

$$\mathbf{l}^H = \begin{pmatrix} \cos\alpha \\ \sin\alpha \\ -L \end{pmatrix}$$

- **y** lies on the line ⇔ $\mathbf{y}_H \cdot \mathbf{l}^H = 0$
- $\mathbf{l}^H$ is the (dual) homogeneous coordinates of the line

Dual homogeneous normalization

- Given a non-zero vector in R^3 we can determine which line it represents by scaling it such that
  - The norm of elements 1 and 2 equals 1
  - Third element is non-positive (≤ 0)
- The elements of the normalized vector directly gives α and L

The cross product

- The cross product $\mathbf{a} \times \mathbf{b}$ for $\mathbf{a}, \mathbf{b} \in R^3$

  $\mathbf{a} \times \mathbf{b}$ is orthogonal to **a** and **b**

  $\mathbf{a} \times \mathbf{b} = \mathbf{0}$ if $\mathbf{a} = \mathbf{b}$

  $\mathbf{a} \times \mathbf{b} = -(\mathbf{b} \times \mathbf{a}) \sim (\mathbf{b} \times \mathbf{a})$

The cross product operator

- For a fix vector **a**, the cross product with **b** is a linear mapping on **b**
- The "**a** $\times$" mapping can be represented by an anti-symmetric 3 $\times$ 3 matrix $[\mathbf{a}]_\times$:

$$[\mathrm{a}]_\times = \begin{pmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{pmatrix} \text{ with } \mathbf{a} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}$$

such that $\mathbf{a} \times \mathbf{b} = [\mathbf{a}]_\times \mathbf{b}$

The intersection of two lines

- Let $\mathbf{l}_1^H$ and $\mathbf{l}_2^H$ be the dual homogeneous representation of two lines in 2D
- Wanted: the intersection point **x** between the lines
- Its homogeneous representation $\mathbf{y}_H$ must satisfy

$$\mathbf{y}_H \cdot \mathbf{l}_1^H = \mathbf{y}_H \cdot \mathbf{l}_2^H = 0$$

$\Rightarrow$ $\mathbf{y}_H$ is orthogonal to both $\mathbf{l}_1^H$ and $\mathbf{l}_2^H$

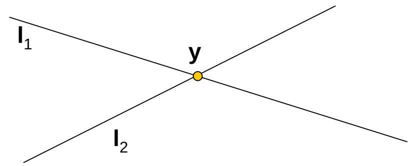

$$\mathbf{y}_H = \mathbf{l}_1^H \times \mathbf{l}_2^H$$

Special case 1

- If the two lines are identical
  - **y** is not unique
  - any point on the line is an intersection point
- In this case: $\mathbf{l}_1^H \times \mathbf{l}_2^H = \mathbf{0}$
- We can use the result = **0** to flag that the lines are identical, i.e., the result is not a specific point

Result = 0 $\Rightarrow$ multiple solutions exist

Special case 2

- If the two lines are distinct but parallel, **y** is "undefined", but …

$$\mathrm{l}_1^H = \begin{pmatrix} \cos\alpha \\ \sin\alpha \\ -L_1 \end{pmatrix} \qquad \mathrm{l}_2^H = \begin{pmatrix} \cos\alpha \\ \sin\alpha \\ -L_2 \end{pmatrix}$$

Cannot be normalized to represent a 2D point

$$\mathrm{l}_1^H \times \mathrm{l}_2^H = (L_1 - L_2) \begin{pmatrix} \sin\alpha \\ -\cos\alpha \\ 0 \end{pmatrix} \sim \begin{pmatrix} \sin\alpha \\ -\cos\alpha \\ 0 \end{pmatrix}$$

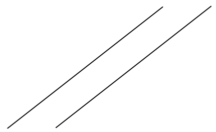

Assuming the lines are distinct

Points at infinity

- The result in this case can be used to represent a "point at infinity"
  - The normalization suggests that the corresponding 2D point lies at infinite distance from the origin
- This is a single point even though there are two directions to look for this point
  - An abstraction of an orientation of a line in 2D
- Given "for free" as elements in P^2
- The result of an operation which maps onto P^2 is either a proper 2D point or a point at infinity

Line intersecting two points

- Let $\mathbf{y}_{1H}$ and $\mathbf{y}_{2H}$ be the homogeneous coordinates of two points in 2D
- We want to find the line which intersects both points
- Its dual homogeneous representation $\mathbf{l}^H$ must satisfy

$$\mathbf{y}_{H1} \cdot \mathbf{l}^H = \mathbf{y}_{H2} \cdot \mathbf{l}^H = 0$$

$\Rightarrow$ $\mathbf{l}^H$ is orthogonal to $\mathbf{y}_{H1}$ and $\mathbf{y}_{H2}$

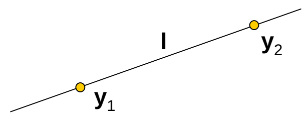

$$\mathbf{l}^H = \mathbf{y}_{H1} \times \mathbf{y}_{H2}$$

Special case 1

- If the two points are identical, the line is not unique: any line going through one point goes through the other
- In this case: $\mathbf{y}_{H1} \times \mathbf{y}_{H2} = \mathbf{0}$
- We can (again) use the result = **0** to flag that the points are identical, i.e., the result is not a specific line but rather a set of lines

Special case 2

- This operation still works also when only one of the two points is a point at infinity:
  - The resulting line goes through the first point
  - In the orientation given by the second point (the point at infinity)

Special case 3

- The operation even works when both points are points at infinity:

$$\mathbf{x}_{1H} = \begin{pmatrix} \cos\alpha \\ \sin\alpha \\ 0 \end{pmatrix} \quad \mathbf{x}_{2H} = \begin{pmatrix} \cos\beta \\ \sin\beta \\ 0 \end{pmatrix}$$

Cannot be normalized to represent a 2D line

$$\mathbf{x}_{1H} \times \mathbf{x}_{2H} = \sin(\alpha-\beta) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \sim \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Assuming the points at inifity are distinct

The line at infinity

- The result in this case can be used to represent a “line at infinity”
  - The normalization suggests that the corresponding 2D line lies at infinite distance from the origin
- There is only one single line at infinity
  - Represents the line which intersects with any distinct pair of points at infinity
  - An abstraction of a circle at infinite distance from the origin
  - Given “for free” as an element of P^2
  - The result of an operation which maps onto P^2 can be either a proper 2D line or the line at infinity

Notation

- In the following, most vectors are homogeneous representations of points or lines
  - Drop the H
  - Use **y** to denote homogeneous coordinates of a 2D point. **y** is then an element of P^2
  - The corresponding 2D point is also called **y** !
  - Use **l** to denote dual homogeneous coordinates of a 2D line. **l** is then an element of P^2
  - The corresponding line is also called **l** !

Affine coordinate transformations

- A 2D point **y** is transformed to **y’** such that the corresponding Euclidean 2D coordinates are related as

$$\begin{pmatrix} y_1' \\ y_2' \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

- This transformation is called *affine*

Affine coordinate transformations

- Translation:

$$\begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

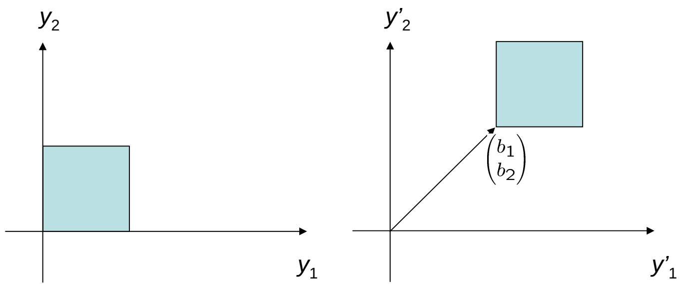

Affine coordinate transformations

- Rotation:

$$\begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

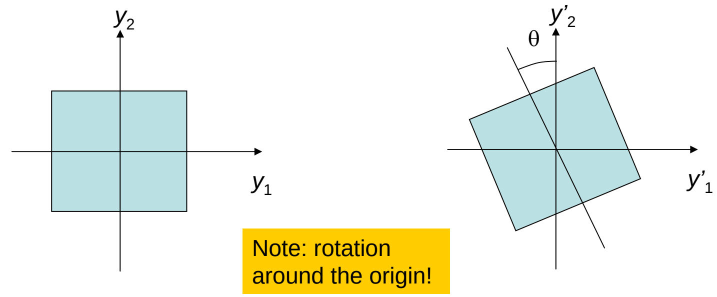

Coordinate transformations

- Scaling

$$\begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix} = \begin{pmatrix} s & 0 \\ 0 & s \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

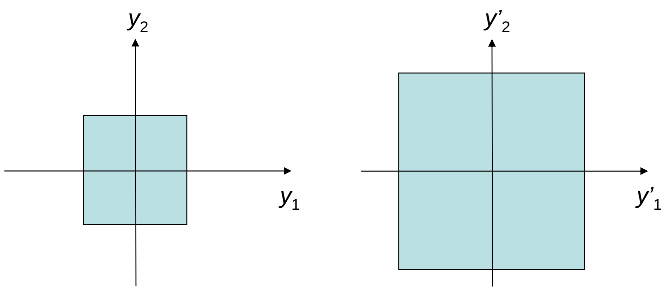

Coordinate transformations

- Mirroring

$$\begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix} = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

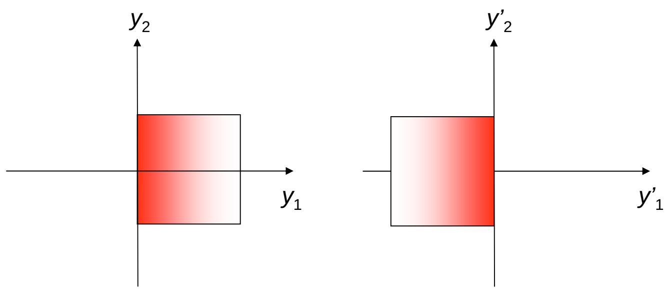

Affine coordinate transformations

- Skewing

$$\begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

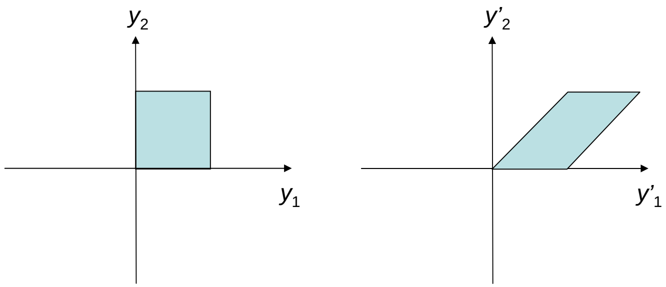

Affine coordinate transformations

- In homogeneous coordinates:

$$\mathbf{y}' = \begin{pmatrix} y'_1 \\ y'_2 \\ 1 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & b_1 \\ a_{21} & a_{22} & b_2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ 1 \end{pmatrix}$$

$$\mathbf{y}' = \begin{pmatrix} y'_1 \\ y'_2 \\ 1 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & b_1 \\ a_{21} & a_{22} & b_2 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{y} = \mathbf{T}\,\mathbf{y}$$

All these transformations are represented as linear mappings **T** onto the homogeneous coordinates

Coordinate transformations

- The composition of two such matrices is again a matrix of this type: a matrix group
  - Rotation around an arbitrary point can be represented as a composition of translation-rotation-translation
- The 3 $\times$ 3 transformation matrix is itself an element of a projective space
  - A multiplication by a non-zero scalar onto the matrix can be moved to either of the two homogeneous vectors **y** or **y**' which gives equivalent homogeneous vectors
- The transformation matrix can be more general than described here
  - More on this after the 3D case has been described

2D coordinate transformations

- Let **y** $\in P^2$ homogeneous coordinates of a 2D point
- Let **T** be a 3 $\times$ 3 matrix which represents some coordinate transformation: **y' = T y**
  - Note that y' represent the same point as y but in a different coordinate system!!
- Let **l** be a line that includes **y**: **l**$\cdot$**y** = 0 (why?)
- If then follows that **l** transforms to **l**'=($\mathbf{T}^T)^{-1}$ **l**
- ($\mathbf{T}^T)^{-1}$ is called the *dual* transformation of **T**

A homogeneous representation of planes in 3D

- Let (x_1,x_2,x_3) be the Euclidean coordinates of a 3D point **x**
- Any 3D plane is characterized by a unit vector $\mathbf{n}=(n_1,n_2,n_3)$ and a scalar L such that

x lies on the plane $\Leftrightarrow$ $x_1\, n_1+x_2\, n_2+x_3\, n_3 = L$

The defining equation of a plane

A homogeneous representation of planes in 3D

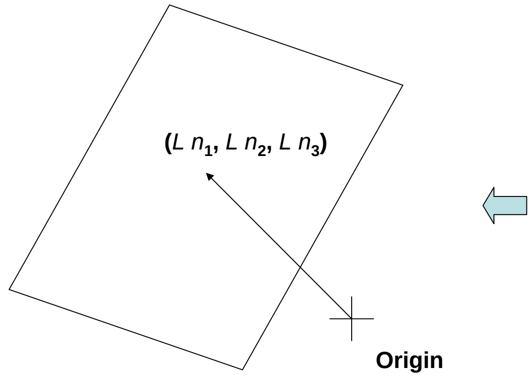

a normal vector which points from the origin to the plane

***L* = shortest distance from the plane to the origin**

A homogeneous representation of planes in 3D

- **x** lies on the plane $\Leftrightarrow$ $x_1\, n_1+x_2\, n_2+x_3\, n_3 = L$

$$\mathbf{x} \text{ lies on the plane} \Leftrightarrow \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix} \cdot \begin{pmatrix} n_1 \\ n_2 \\ n_3 \\ -L \end{pmatrix} = 0$$

Homogeneous coordinates of **x**

A homogeneous representation of planes in 3D

- Suggests a homogeneous representation of the plane:

$$\mathbf{p} = \begin{pmatrix} n_1 \\ n_2 \\ n_3 \\ -L \end{pmatrix}$$

- **x** lies on the plane $\Leftrightarrow$ $\mathbf{x} \cdot \mathbf{p} = 0$
- **p** are the (dual) homogeneous coordinates of the plane

Dual homogeneous normalization

- Given a vector in R^4 we can determine which plane it represents by scaling it such that
  - The norm of elements 1 to 3 equals 1
  - Fourth element is non-positive (≤ 0)
- The elements of the normalized vector directly gives **n** and *L*
- Similar to the 2D case:

  $\mathbf{x'} = \mathbf{T}\,\mathbf{x} \;\Leftrightarrow\; \mathbf{p'} = (\mathbf{T}^T)^{-1}\,\mathbf{p}$

Points and planes at infinity

Similar to the 2D case:

- In 3D there are points at infinity
  - Have last homogeneous coordinate = 0
- There is a single 3D plane at infinity
  - Intersects all 3D points at infinity

Affine transformations in 3D

- Simple extension from the 2D case!!

BREAK!

A homogeneous representation of lines in 3D

- 3D lines can be represented in several slightly different ways
- Here we will use
  - so called *Plücker coordinates* in the form of an anti-symmetric matrix
  - Parametric representation:

    $\mathbf{x} = \mathbf{x}_0 + t\,\mathbf{n}$ eller $\mathbf{x} = t\,\mathbf{x}_1 + (1 - t)\,\mathbf{x}_2$

Parametric representation of lines in 3D

- Let $\mathbf{x_1}$ and $\mathbf{x_2}$ be two distinct 3D points with $\mathbf{x_1}, \mathbf{x_2} \in P^3$
- Any point $\mathbf{x}$ on the line can be written

$$\mathbf{x} = t\,\mathbf{x}_1 + (1 - t)\,\mathbf{x}_2 \quad \text{for some } t \in R$$

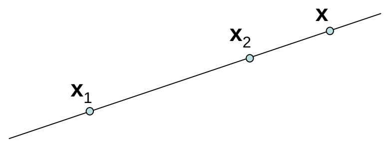

A homogeneous representation of lines in 3D

Define $\mathbf{L} = \mathbf{x_1}\,\mathbf{x_2}^T - \mathbf{x_2}\,\mathbf{x_1}^T$

- $\mathbf{L}$ is a homogenous representation of the line which intersects $\mathbf{x_1}$ and $\mathbf{x_2}$
- $\mathbf{L}$ is a 4×4 anti-symmetric matrix: $\mathbf{L}^T = -\mathbf{L}$
- $\mathbf{L}$ can be seen as a projective element (why?)
- Referred to as *Plücker coordinates* of the line
- As a projective element $\mathbf{L}$ is independent of which two distinct points on the line are used (why?)

Intersection between a line and a plane in 3D

- Let $\mathbf{L}$ be the Plücker coordinates of a 3D line
- Let $\mathbf{p}$ the dual homogeneous coordinates of a plane
- Which is the intersection point $\mathbf{x}_0$?
  - Characterized by $\mathbf{x}_0 \cdot \mathbf{p} = 0$

$$\mathbf{x}_0 \sim \mathbf{L}\,\mathbf{p}$$

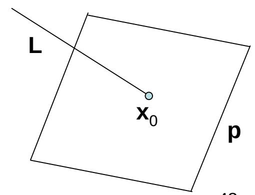

Dual Plücker coordinates

- Alternatively, let $\mathbf{p}_1$ and $\mathbf{p}_2$ be two planes that intersect the 3D line
- $\mathbf{L}' = \mathbf{p}_1\mathbf{p}_2^T - \mathbf{p}_2\mathbf{p}_1^T$ is the *dual Plücker* coordinates of the line
- Independent of which 2 planes we use (as long as they are distinct and intersect the line)
- $\mathbf{L}'\mathbf{x}$ gives the plane that includes the line and point $\mathbf{x}$
- Relation between **L** and **L**' ?

2D coordinates on a 3D plane

- The Euclidean coordinates of a 3D point in a plane can be described as

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} o_1 \\ o_2 \\ o_3 \end{pmatrix} + \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} s + \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} t$$

Some point on the plane

Tangent vectors of the plane

All three vectors define a 2D coordinate system in the plane

2D coordinates on a 3D plane

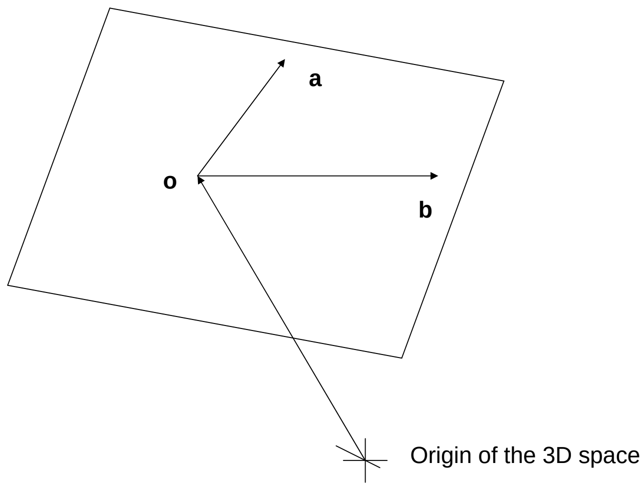

2D coordinates on a 3D plane

- In homogeneous coordinates:

Homogeneous 2D coordinates: y

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix} = \begin{pmatrix} a_1 & b_1 & o_1 \\ a_2 & b_2 & o_2 \\ a_3 & b_3 & o_3 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} s \\ t \\ 1 \end{pmatrix}$$

Mapping from 2D to 3D homogeneous coordinates: P

or $\mathbf{x} = \mathbf{P}\,\mathbf{y}$

P uniquely defines the 2D coordinate system on the plane

2D to 2D projective mappings

- Given
  - two 3D planes, each with its own 2D coordinate system, $\mathbf{P}_1$ and $\mathbf{P}_2$
  - a 3D point $\mathbf{n}$

  there is a unique mapping from one plane to the other:

Project a point $\mathbf{x}_1$ on the first plane through $\mathbf{n}$ onto the second plane which gives $\mathbf{x}_2$

2D to 2D projective mappings

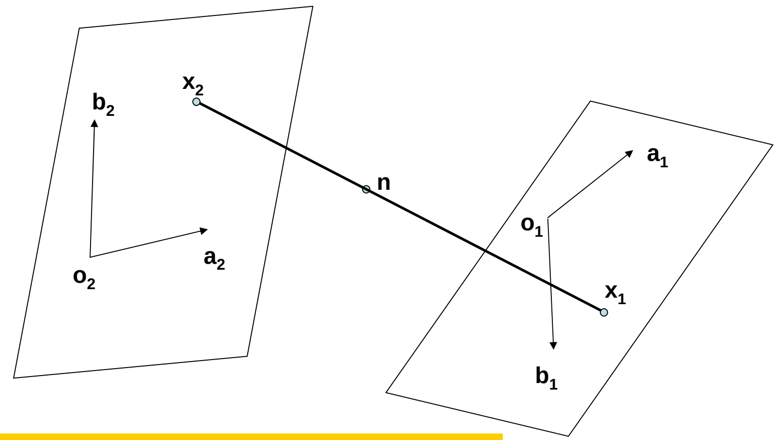

In this example: **n** is located in between the planes
This is not required!

2D to 2D projective mappings

- The geometric relation between $\mathbf{x}_1$, $\mathbf{x}_2$, and $\mathbf{n}$ together with $\mathbf{x}_1 = \mathbf{P}_1\, \mathbf{y}_1$ and $\mathbf{x}_2 = \mathbf{P}_2\, \mathbf{y}_2$ leads to (why?)

$$\mathbf{y}_2 = \mathbf{H}\, \mathbf{y}_1$$

- **H** is a 3×3 general non-singular matrix
- Depends on the two planes and on **n**

Homography

- This mapping on the 2D coordinates in the two planes is more general than the affine transformations described earlier!
- Called *homography* or *projective transformation*
- Any 3×3 non-singular **H** is a homography
- Describes e.g. how a pinhole-camera maps points on a plane to the image plane

Homography

- Includes the affine transformations!
  - In the special case that the planes are parallel
  - In other cases: there are points at infinity that are mapped to normal points and vice versa
- We assume that **n** is not included in any of the two planes ⇒ **H** is always invertible
  - We can uniquely go from image coordinates to coordinates in the plane
- H always maps a line to a line (why?)

Homographies

- Any homography is determined by how it maps 4 distinct points

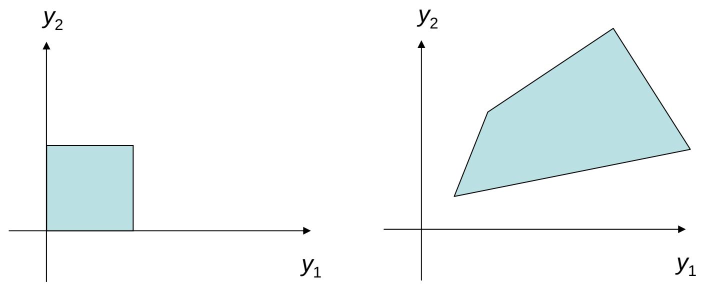

Homographies

- One or two of the 4 points may be at infinity

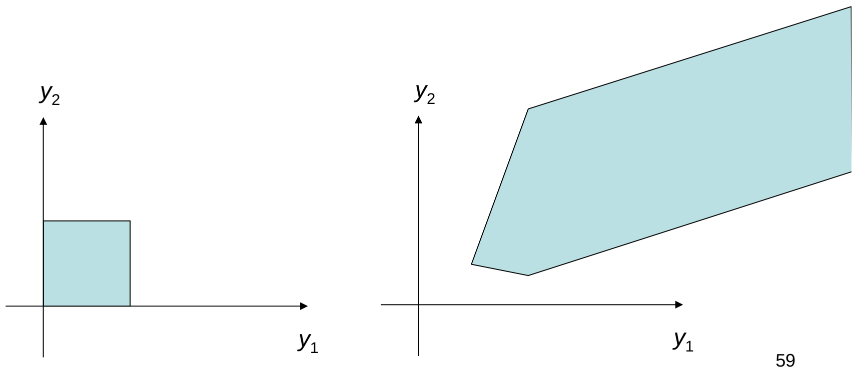

3D homography transformations

- The group of 4 × 4 non-singular matrices define the group of 3D homography transformations
- Analogue to the 2D case, but cannot be characterized in terms of projective mappings in a simple way

The pinhole camera model

The image plane and the camera center define a camera-centered coordinate system (x_1,x_2,x_3):

x_1,x_2 are parallel to the image plane, x_3 is perpendicular to the plane and defines the viewing direction of the camera

Principal or optical axis

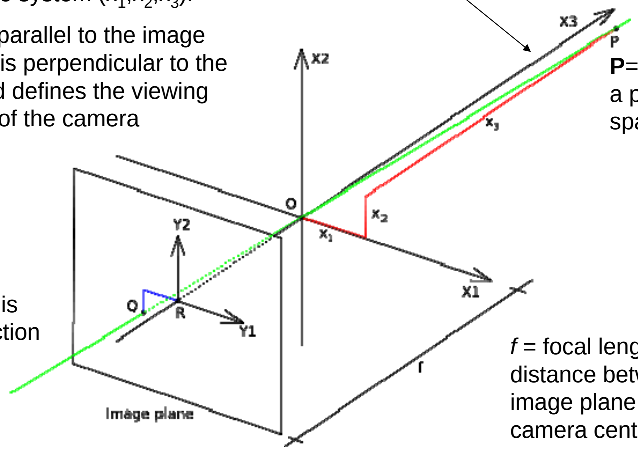

P=(x_1,x_2,x_3) is a point in 3D space

Q=(y_1,y_2) is the projection of **P**

f = focal length, the distance between the image plane and the camera center

The pinhole camera model

- If we look at the camera coordinate system along the x_2 axis:

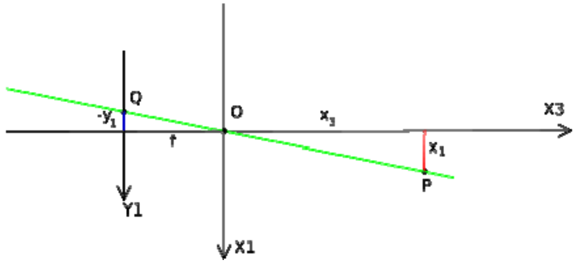

Two similar triangles give:

$$\frac{-y_1}{f} = \frac{x_1}{x_3} \quad \text{or} \quad y_1 = -\frac{f\,x_1}{x_3}$$

The pinhole camera model

- Looking along the x_1 axis gives a similar expression for y_2
- This can be summarized as:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = -\frac{f}{x_3} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

The virtual image plane

- The projected image is rotated 180^o relative to how we “see” the 3D world
  - Reflection in both y_1 and y_2 coordinates = rotation
- Must be de-rotated before we can view it
- Mathematically this is equivalent to placing the image plane *in front* of the focal point
- Called a *virtual image plane*

The pinhole-camera

- We now have:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \frac{f}{x_3} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

The mapping of 3D coordinates to 2D image coordinates defined by the pinhole-camera in camera centered coordinates (virtual image plane!)

- In homogeneous image coordinates

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ 1 \end{pmatrix} = \frac{f}{x_3} \begin{pmatrix} x_1 \\ x_2 \\ x_3/f \end{pmatrix} \sim \begin{pmatrix} f\, x_1 \\ f\, x_2 \\ x_3 \end{pmatrix}$$

The pinhole-camera

- Using also homogeneous 3D coordinates:

$$\mathbf{y} \sim \begin{pmatrix} f\, x_1 \\ f\, x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix}$$

Defines a 3×4 matrix C

$$\mathbf{y} \sim \mathbf{C}\,\mathbf{x}$$

C is the camera (projection) matrix

The normalized camera

- In the case of a *normalized camera*: f = 1

$$\mathbf{C}_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Notation for the normalized camera matrix

The camera center

- In the camera centered coordinate system, the camera center (focal point) has 3D coordinates (0,0,0)
- The camera matrix maps this point to:

$$\begin{pmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \mathbf{0}$$

- The homogeneous representation of the camera center lies in the *null space* of the camera matrix

The general camera matrix

- The camera matrix defined so far assumes that both 2D and 3D coordinates are given in a camera centered coordinate system
- We want to be able to use
  - 3D coordinates in any coordinate system of our choice, *world coordinates*
  - 2D image coordinates in a pixel based coordinate system, often with the origin at the top left corner and first coordinate down

The general camera matrix

- Assuming that the world coordinate system we use is Euclidean, there is always a rotation and translation of the 3D coordinate system that align it with the camera centered system

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{pmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} + \begin{pmatrix} d_1 \\ d_2 \\ d_3 \end{pmatrix}$$

Camera centered 3D coordinates

Rotation transformation

3D world coordinates

World origin in camera centered coordinates

The general camera matrix

- In homogeneous coordinates:

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \\ 1 \end{pmatrix} = \begin{pmatrix} r_{11} & r_{12} & r_{13} & d_1 \\ r_{21} & r_{22} & r_{23} & d_2 \\ r_{31} & r_{32} & r_{33} & d_3 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix}$$

Camera centered 3D coordinates

Defines the transformation matrix T_e

World coordinates

$$\mathbf{x}' = \mathbf{T}_e \, \mathbf{x}$$

The general camera matrix

- The normalized image coordinates are then given as

$$\mathbf{y}_0 = \mathbf{C}_0 \, \mathbf{x}' = \mathbf{C}_0 \, \mathbf{T}_e \, \mathbf{x}$$

$$\mathrm{C}_0 \, \mathrm{T}_e = \begin{pmatrix} r_{11} & r_{12} & r_{13} & d_1 \\ r_{21} & r_{22} & r_{23} & d_2 \\ r_{31} & r_{32} & r_{33} & d_3 \end{pmatrix}$$

$$\mathbf{C}_0 \, \mathbf{T}_e = (\, \mathbf{R} \mid \mathbf{d} \,)$$

Image coordinates

- Normalized image coordinates
  - f = 1
  - Origin at the image center
  - First coordinate right, second up
  - Same length unit as in 3D space

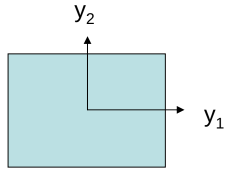

- Standard image coordinates
  - Arbitrary f > 0
  - Origin at the image top left
  - First coordinate down, second right
  - Pixel based length unit

or vice versa

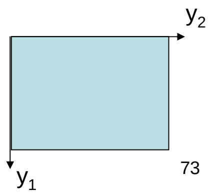

Image coordinates

- To transform from $\mathbf{y}_0$ to standard image coordinates $\mathbf{y}$

$$\mathbf{y} = \begin{pmatrix} 0 & -r\,f & c_1 \\ r\,f & 0 & c_2 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{y}_0$$

(c_1,c_2) = coordinates of the image center in the standard coodinates

- **f = focal length**
- **r = pixel resolution [pixel / length unit]**

Defines the transformation matrix $\mathbf{T}_i$

$$\mathbf{y} = \mathbf{T}_i\,\mathbf{y}_0$$

The general camera matrix

- We can now summarize all this as

$$\mathbf{y} = \mathbf{T}_i\,\mathbf{y}_0 = \mathbf{T}_i\,\mathbf{C}_0\,\mathbf{T}_e\,\mathbf{x} = \mathbf{C}\,\mathbf{x}$$

- The general camera matrix **C** is given by

Internal (intrinsic) camera parameters

The normalized camera matrix

$$\mathbf{C} = \mathbf{T}_i\,\mathbf{C}_0\,\mathbf{T}_e$$

External (extrinsic) camera parameters

The general camera matrix

- $\mathbf{T}_e$ depends on where the camera (camera center!) is positioned in 3D space and how it is oriented. May be variable or fixed depending on application
- $\mathbf{T}_i$ depends on the type of camera, and its setting such as zoom, resolution, etc. Typically fixed.
- Since **C** is the product of three matrices of rank 3, 3, and 4 $\Rightarrow$ **C** has rank 3
- To determine **C** is referred to as *camera calibration* (separate lecture)

Equivalent cameras

- Let $\mathbf{C}_1$ and $\mathbf{C}_2$ be the camera matrices of two pinhole cameras with the *same camera center* $\mathbf{n}$

  $\mathbf{y}_1 = \mathbf{C}_1\,\mathbf{x} \qquad \mathbf{C}_1\,\mathbf{n} = \mathbf{0}$

  $\mathbf{y}_2 = \mathbf{C}_2\,\mathbf{x} \qquad \mathbf{C}_2\,\mathbf{n} = \mathbf{0}$
- In this case: there is a homography mapping $\mathbf{H}$ from $\mathbf{y}_1$ to $\mathbf{y}_2$ defined by $\mathbf{C}_1$ and $\mathbf{C}_2$ such that

  $\mathbf{y}_1 = \mathbf{H}\,\mathbf{y}_2 \qquad \mathbf{y}_2 = \mathbf{H}^{-1}\,\mathbf{y}_1$ (why?)
- The images in the two cameras are identical except for a geometric transformation
  - In practice the images crop different parts!

Affine camera

- In certain applications the 3D points have a distance *d* to the camera that does not vary much relative to the distance

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \frac{f}{x_3}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \approx \frac{f}{d}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

- In homogeneous coordinates:

$$\mathrm{y} = \begin{pmatrix} y_1 \\ y_2 \\ 1 \end{pmatrix} \approx\sim \begin{pmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 0 & d \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix}$$

The *affine camera* matrix: it always has bottom row (0 0 0 d)

The orthographic camera

- An identical case appears when the 3D points are at a large distance from the camera
- Referred to as an *orthographic camera*
- Note: the affine/orthographic property is derived from propoerties of the 3D points, not of the camera

Conics (in 2D)

- (y_1, y_2) lies on a conic curve centered on the origin if

$$\begin{pmatrix} y_1 & y_2 \end{pmatrix} \mathbf{A} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = 1$$

A is 2 × 2 symmetric and determines the character of the curve

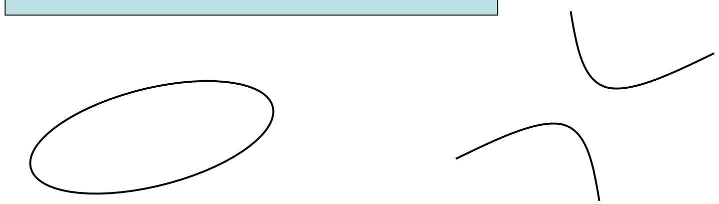

Conics (in 2D)

- In homogeneous coordinates the defining equation becomes

$$\mathbf{y}^T \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0}^T & -1 \end{pmatrix} \mathbf{y} = \mathbf{y}^T \mathbf{Q}\, \mathbf{y} = 0$$

Q is 3 × 3 symmetric

- Generalizes to conics at arbitrary positions by appropriate translations

Conics (in 2D)

Assuming that **y** lies on the conic

- We can interpret **Q y** as a line that must pass through **y** (why?)
- This line is in fact the tangent $\mathbf{l}_t$ of the conic at point **y**

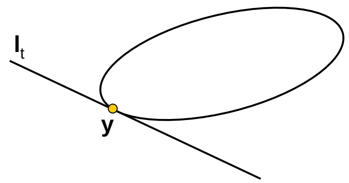

Dual conics

- $\mathbf{y}^T\mathbf{Q}\,\mathbf{y} = 0$ defines the points **y** that lie on a conic
- Follows: $\mathbf{l}^T\, \mathbf{Q}^{-1}\, \mathbf{l} = 0$ defines the lines that are tangent to the same conic (why?)
- $\mathbf{Q}^{-1}$ is the *dual conic* relative to **Q**
- $\mathbf{Q}^{-1}\mathbf{l}$ gives the tangent point of tanget line **l**