



2. Image Formation



2.1 Geometric Primitives & Transformations




Outline

- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ Lens Distortions

Outline


- ❖ **Geometric primitives**
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ Lens Distortions

2D Points



2D Points


2D point (e.g., a pixel coordinate in an image):



2D Points

2D point (e.g., a pixel coordinate in an image):


$$\mathbf{x} = (x, y) \in \mathcal{R}^2.$$



2D Points

2D point (e.g., a pixel coordinate in an image):


$$\mathbf{x} = (x, y) \in \mathcal{R}^2. \quad \text{or}$$



2D Points

2D point (e.g., a pixel coordinate in an image):

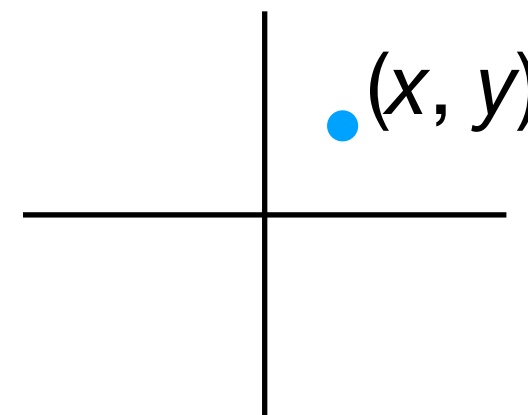
$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$



2D Points

2D point (e.g., a pixel coordinate in an image):

$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$




In **homogeneous coordinates**:

2D Points

2D point (e.g., a pixel coordinate in an image):

$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$




In **homogeneous coordinates**:

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) \in \mathcal{P}^2$$

2D Points

2D point (e.g., a pixel coordinate in an image):

$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$



In **homogeneous coordinates**:


$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) \in \mathcal{P}^2$$

$\mathcal{P}^2 = \mathcal{R}^3 - (0,0,0)$ is called the projective space.

2D Points

2D point (e.g., a pixel coordinate in an image):

$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$



In **homogeneous coordinates**:

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) \in \mathcal{P}^2$$


$\mathcal{P}^2 = \mathcal{R}^3 - (0,0,0)$ is called the projective space.

Vectors that differ only by a scale considered equivalent:

2D Points

2D point (e.g., a pixel coordinate in an image):

$$\mathbf{x} = (x, y) \in \mathcal{R}^2 \quad \text{or} \quad \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}$$



In **homogeneous coordinates**:

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) \in \mathcal{P}^2$$

$\mathcal{P}^2 = \mathcal{R}^3 - (0,0,0)$ is called the projective space.

Vectors that differ only by a scale considered equivalent:

$$s\tilde{\mathbf{x}} = \tilde{\mathbf{x}} \quad \forall s \in \mathcal{R}$$

Augmented Vectors


A homogenous vector can be converted back to an *inhomogeneous* vector by dividing by the last element:

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) = \tilde{w}(x, y, 1) = \tilde{w}\bar{\mathbf{x}}$$

Homogeneous vector


Augmented vector

2D Lines



2D Lines

2D lines can also be represented using homogeneous coordinates $\tilde{l} = (a, b, c)$.




2D Lines

2D lines can also be represented using homogeneous coordinates $\tilde{l} = (a, b, c)$.

The corresponding *line equation* is

$$\bar{x} \cdot \tilde{l} = ax + by + c = 0.$$




2D Lines

2D lines can also be represented using homogeneous coordinates $\tilde{l} = (a, b, c)$.

The corresponding *line equation* is

$$\bar{x} \cdot \tilde{l} = ax + by + c = 0.$$

We can normalize the line equation vector so that $l = (\hat{n}_x, \hat{n}_y, d) = (\hat{n}, d)$ with $\|\hat{n}\| = 1$.



2D Lines

2D lines can also be represented using homogeneous coordinates $\tilde{l} = (a, b, c)$.

The corresponding *line equation* is


$$\bar{x} \cdot \tilde{l} = ax + by + c = 0.$$

We can normalize the line equation vector so that $l = (\hat{n}_x, \hat{n}_y, d) = (\hat{n}, d)$ with $\|\hat{n}\| = 1$.

Then:

\hat{n} is the unit normal perpendicular to the line

$|d|$ is the distance of the line from the origin



2D Lines

2D lines can also be represented using homogeneous coordinates $\tilde{l} = (a, b, c)$.

The corresponding *line equation* is

$$\bar{x} \cdot \tilde{l} = ax + by + c = 0.$$


We can normalize the line equation vector so that $l = (\hat{n}_x, \hat{n}_y, d) = (\hat{n}, d)$ with $\|\hat{n}\| = 1$.

Then:

\hat{n} is the unit normal perpendicular to the line

$|d|$ is the distance of the line from the origin

$$\hat{n} = (\hat{n}_x, \hat{n}_y) = (\cos \theta, \sin \theta)$$



Intersections of 2D Lines

Intersections of 2D Lines

When using homogeneous coordinates, we can compute the intersection of two lines as

Intersections of 2D Lines


When using homogeneous coordinates, we can compute the intersection of two lines as

$$\tilde{\mathbf{x}} = \tilde{\mathbf{l}}_1 \times \tilde{\mathbf{l}}_2$$

Intersections of 2D Lines

When using homogeneous coordinates, we can compute the intersection of two lines as


$$\tilde{\mathbf{x}} = \tilde{\mathbf{l}}_1 \times \tilde{\mathbf{l}}_2$$



Intersections of 2D Lines

When using homogeneous coordinates, we can compute the intersection of two lines as

$$\tilde{\mathbf{x}} = \tilde{\mathbf{l}}_1 \times \tilde{\mathbf{l}}_2$$




Similarly, the line joining two points can be written as

Intersections of 2D Lines

When using homogeneous coordinates, we can compute the intersection of two lines as

$$\tilde{\mathbf{x}} = \tilde{\mathbf{l}}_1 \times \tilde{\mathbf{l}}_2$$




Similarly, the line joining two points can be written as

$$\tilde{\mathbf{l}} = \tilde{\mathbf{x}}_1 \times \tilde{\mathbf{x}}_2$$

Intersections of 2D Lines


When using homogeneous coordinates, we can compute the intersection of two lines as

$$\tilde{\mathbf{x}} = \tilde{\mathbf{l}}_1 \times \tilde{\mathbf{l}}_2$$



Similarly, the line joining two points can be written as

$$\tilde{\mathbf{l}} = \tilde{\mathbf{x}}_1 \times \tilde{\mathbf{x}}_2$$



3D Points


Straightforward extension from 2D:

Inhomogeneous: $\mathbf{x} = (x, y, z) \in \mathcal{R}^3$


Homogeneous: $\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{z}, \tilde{w}) \in \mathcal{P}^3$

Augmented: $\bar{\mathbf{x}} = (x, y, z, 1)$

$$\tilde{\mathbf{x}} = \tilde{w}\bar{\mathbf{x}}$$




3D Planes



3D Planes

3D planes can also be represented as homogeneous coordinates $\tilde{m} = (a, b, c, d)$




3D Planes

3D planes can also be represented as homogeneous coordinates $\tilde{\mathbf{m}} = (a, b, c, d)$

with a corresponding plane equation

$$\bar{\mathbf{x}} \cdot \tilde{\mathbf{m}} = ax + by + cz + d = 0$$




3D Planes

3D planes can also be represented as homogeneous coordinates $\tilde{\mathbf{m}} = (a, b, c, d)$

with a corresponding plane equation

$$\bar{\mathbf{x}} \cdot \tilde{\mathbf{m}} = ax + by + cz + d = 0$$

We can also normalize the plane equation as $\mathbf{m} = (\hat{n}_x, \hat{n}_y, \hat{n}_z, d) = (\hat{\mathbf{n}}, d)$ with $\|\hat{\mathbf{n}}\| = 1$.



3D Planes


3D planes can also be represented as homogeneous coordinates $\tilde{\mathbf{m}} = (a, b, c, d)$

with a corresponding plane equation

$$\bar{\mathbf{x}} \cdot \tilde{\mathbf{m}} = ax + by + cz + d = 0$$

We can also normalize the plane equation as $\mathbf{m} = (\hat{n}_x, \hat{n}_y, \hat{n}_z, d) = (\hat{\mathbf{n}}, d)$ with $\|\hat{\mathbf{n}}\| = 1$.

$\hat{\mathbf{n}}$ is the *normal vector* perpendicular to the plane



3D Planes

3D planes can also be represented as homogeneous coordinates $\tilde{\mathbf{m}} = (a, b, c, d)$


with a corresponding plane equation

$$\bar{\mathbf{x}} \cdot \tilde{\mathbf{m}} = ax + by + cz + d = 0$$


We can also normalize the plane equation as $\mathbf{m} = (\hat{n}_x, \hat{n}_y, \hat{n}_z, d) = (\hat{\mathbf{n}}, d)$ with $\|\hat{\mathbf{n}}\| = 1$.

$\hat{\mathbf{n}}$ is the *normal vector* perpendicular to the plane

$|d|$ is the distance of the plane from the origin




3D Lines



3D Lines


A 3D line can be represented using two points p and q that lie on the line.



3D Lines

A 3D line can be represented using two points p and q that lie on the line.

Any point r that also lies on the line can then be represented as




3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$




3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda)\mathbf{p} + \lambda\mathbf{q}$$

If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .



3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.


Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$

If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .

If we use homogeneous coordinates, we can write the line as

$$\tilde{\mathbf{r}} = \mu \tilde{\mathbf{p}} + \lambda \tilde{\mathbf{q}}.$$



3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as


$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$

If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .

If we use homogeneous coordinates, we can write the line as

$$\tilde{\mathbf{r}} = \mu \tilde{\mathbf{p}} + \lambda \tilde{\mathbf{q}}.$$

A special case of this is when the second point is at infinity, i.e., $\tilde{\mathbf{q}} = (\hat{d}_x, \hat{d}_y, \hat{d}_z, 0) = (\hat{\mathbf{d}}, 0)$.



3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$


If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .

If we use homogeneous coordinates, we can write the line as

$$\tilde{\mathbf{r}} = \mu \tilde{\mathbf{p}} + \lambda \tilde{\mathbf{q}}.$$

A special case of this is when the second point is at infinity, i.e., $\tilde{\mathbf{q}} = (\hat{d}_x, \hat{d}_y, \hat{d}_z, 0) = (\hat{\mathbf{d}}, 0)$.

where $\hat{\mathbf{d}}$ is the direction of the line.



3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$

If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .


If we use homogeneous coordinates, we can write the line as

$$\tilde{\mathbf{r}} = \mu \tilde{\mathbf{p}} + \lambda \tilde{\mathbf{q}}.$$

A special case of this is when the second point is at infinity, i.e., $\tilde{\mathbf{q}} = (\hat{d}_x, \hat{d}_y, \hat{d}_z, 0) = (\hat{\mathbf{d}}, 0)$.

where $\hat{\mathbf{d}}$ is the direction of the line.

Then:



3D Lines

A 3D line can be represented using two points \mathbf{p} and \mathbf{q} that lie on the line.

Any point \mathbf{r} that also lies on the line can then be represented as

$$\mathbf{r} = (1 - \lambda) \mathbf{p} + \lambda \mathbf{q}$$

If we restrict $0 \leq \lambda \leq 1$, we get the *line segment* joining \mathbf{p} and \mathbf{q} .

If we use homogeneous coordinates, we can write the line as


$$\tilde{\mathbf{r}} = \mu \tilde{\mathbf{p}} + \lambda \tilde{\mathbf{q}}.$$

A special case of this is when the second point is at infinity, i.e., $\tilde{\mathbf{q}} = (\hat{d}_x, \hat{d}_y, \hat{d}_z, 0) = (\hat{\mathbf{d}}, 0)$.

where $\hat{\mathbf{d}}$ is the direction of the line.

Then:


$$\mathbf{r} = \mathbf{p} + \lambda \hat{\mathbf{d}}$$



Outline

- ❖ Geometric primitives
- ❖ **2D transformations**
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ Lens Distortions


2D Translation



2D Translation

2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$




2D Translation

2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{I} is the (2×2) identity matrix




2D Translation

2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{I} is the (2×2) identity matrix

or



2D Translation


2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{I} is the (2×2) identity matrix

or

$$\bar{\mathbf{x}}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \bar{\mathbf{x}}$$



2D Translation

2D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or


$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{I} is the (2×2) identity matrix

or

$$\bar{\mathbf{x}}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \bar{\mathbf{x}}$$

Note: Whenever an augmented vector appears on both sides, it can be replaced by a full homogenous vector.



Euclidean Transformation (2D Rotation +


$$x' = Rx + t \text{ or}$$

$$x' = \begin{bmatrix} R & t \end{bmatrix} \bar{x}$$

where

$$R = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

is an orthonormal rotation matrix with $RR^T = I$ and $|R| = 1$.



Euclidean Transformation (2D Rotation +

$$x' = Rx + t \text{ or}$$


$$x' = \begin{bmatrix} R & t \end{bmatrix} \bar{x}$$

where


$$R = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

is an orthonormal rotation matrix with $RR^T = I$ and $|R| = 1$.

Preserves Euclidean distances




Similarity Transformation



Similarity Transformation


$$x' = sRx + t$$



Similarity Transformation

$$\mathbf{x}' = s\mathbf{R}\mathbf{x} + \mathbf{t}$$

$$\mathbf{x}' = \begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}} = \begin{bmatrix} a & -b & t_x \\ b & a & t_y \end{bmatrix} \bar{\mathbf{x}}$$




Similarity Transformation

$$\mathbf{x}' = s\mathbf{R}\mathbf{x} + \mathbf{t}$$

$$\mathbf{x}' = \begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}} = \begin{bmatrix} a & -b & t_x \\ b & a & t_y \end{bmatrix} \bar{\mathbf{x}}$$


Preserves angles



Affine Transformation

$\mathbf{x}' = \mathbf{A}\bar{\mathbf{x}}$, where \mathbf{A} is an arbitrary 2×3 matrix

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \end{bmatrix} \bar{\mathbf{x}}$$




Affine Transformation


$\mathbf{x}' = \mathbf{A}\bar{\mathbf{x}}$, where \mathbf{A} is an arbitrary 2×3 matrix

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \end{bmatrix} \bar{\mathbf{x}}$$

Preserves parallelism




Projective Transformation (Homography)



Projective Transformation (Homography)


$$\tilde{x}' = \tilde{H} \tilde{x}$$



Projective Transformation (Homography)

$$\tilde{x}' = \tilde{H} \tilde{x}$$

where \tilde{H} is an arbitrary 3×3 matrix.




Projective Transformation (Homography)

$$\tilde{x}' = \tilde{H} \tilde{x}$$

where \tilde{H} is an arbitrary 3×3 matrix.

\tilde{H} is homogenous:

Two \tilde{H} matrices that differ only by a scale factor are equivalent.



Projective Transformation (Homography)


$$\tilde{x}' = \tilde{H} \tilde{x}$$

where \tilde{H} is an arbitrary 3×3 matrix.

\tilde{H} is homogenous:

Two \tilde{H} matrices that differ only by a scale factor are equivalent.

$$x' = \frac{h_{00}x + h_{01}y + h_{02}}{h_{20}x + h_{21}y + h_{22}} \quad \text{and} \quad y' = \frac{h_{10}x + h_{11}y + h_{12}}{h_{20}x + h_{21}y + h_{22}}$$



Projective Transformation (Homography)

$$\tilde{x}' = \tilde{H} \tilde{x}$$


where \tilde{H} is an arbitrary 3×3 matrix.

\tilde{H} is homogenous:

Two \tilde{H} matrices that differ only by a scale factor are equivalent.

$$x' = \frac{h_{00}x + h_{01}y + h_{02}}{h_{20}x + h_{21}y + h_{22}} \quad \text{and} \quad y' = \frac{h_{10}x + h_{11}y + h_{12}}{h_{20}x + h_{21}y + h_{22}}$$


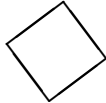
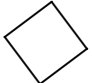

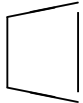
Preserves straight lines




Summary of 2D Transformations

Nested set of groups

- ❖ Closed under composition
- ❖ Each transformation has an inverse that is a member of the same group

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} \mathbf{I} & & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths	
similarity	$\begin{bmatrix} s\mathbf{R} & & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	



Co-vectors

Co-vectors

We now know how to transform points.
How do we transform lines?

Co-vectors

We now know how to transform points.
How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

Co-vectors

We now know how to transform points.
How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

$$\tilde{x}' = \tilde{H} \tilde{x}$$

Co-vectors

We now know how to transform points.

How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

$$\tilde{x}' = \tilde{H} \tilde{x}$$

$$\tilde{l}' \cdot \tilde{x}' = \tilde{l}'^T \tilde{H} \tilde{x} = (\tilde{H}^T \tilde{l}')^T \tilde{x} = \tilde{l} \cdot \tilde{x} = 0$$

Co-vectors

We now know how to transform points.

How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

$$\tilde{x}' = \tilde{H} \tilde{x}$$

$$\tilde{l}' \cdot \tilde{x}' = \tilde{l}'^T \tilde{H} \tilde{x} = (\tilde{H}^T \tilde{l}')^T \tilde{x} = \tilde{l} \cdot \tilde{x} = 0$$

Thus

Co-vectors

We now know how to transform points.
How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

$$\tilde{x}' = \tilde{H} \tilde{x}$$

$$\tilde{l}' \cdot \tilde{x}' = \tilde{l}'^T \tilde{H} \tilde{x} = (\tilde{H}^T \tilde{l}')^T \tilde{x} = \tilde{l} \cdot \tilde{x} = 0$$

Thus

$$\tilde{l}' = \tilde{H}^{-T} \tilde{l}.$$

Co-vectors

We now know how to transform points.

How do we transform lines?

$$\tilde{l} \cdot \tilde{x} = 0$$

$$\tilde{x}' = \tilde{H} \tilde{x}$$

$$\tilde{l}' \cdot \tilde{x}' = \tilde{l}'^T \tilde{H} \tilde{x} = (\tilde{H}^T \tilde{l}')^T \tilde{x} = \tilde{l} \cdot \tilde{x} = 0$$

Thus

$$\tilde{l}' = \tilde{H}^{-T} \tilde{l}.$$

i.e., the action of a projective transformation on a *co-vector* such as a 2D line can be represented by the transposed inverse of the matrix.

Outline

- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ **3D transformations**
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ Lens Distortions

3D Translation

3D translations can be written as $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ or

$$\mathbf{x}' = \begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{I} is the (3×3) identity matrix

Euclidean Transformation (3D Rotation + Translation)

$$\mathbf{x}' = \mathbf{R}\mathbf{x} + \mathbf{t}$$

$$\mathbf{x}' = \begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix} \bar{\mathbf{x}}$$

where \mathbf{R} is a 3×3 orthonormal rotation matrix with $\mathbf{R}\mathbf{R}^T = \mathbf{I}$ and $|\mathbf{R}| = 1$

Preserves Euclidean distances

Similarity Transformation

$$x' = sRx + t$$

$$x' = \begin{bmatrix} sR & t \end{bmatrix} \bar{x}$$

Preserves angles

Affine Transformation

$\mathbf{x}' = \mathbf{A}\bar{\mathbf{x}}$, where \mathbf{A} is an arbitrary 3×4 matrix

$$\mathbf{x}' = \begin{bmatrix} a_{00} & a_{01} & a_{02} & a_{03} \\ a_{10} & a_{11} & a_{12} & a_{13} \\ a_{20} & a_{21} & a_{22} & a_{23} \end{bmatrix} \bar{\mathbf{x}}$$

Preserves parallelism


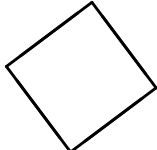
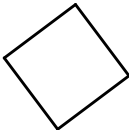
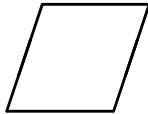
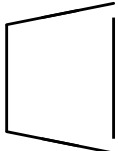
Projective Transformation (Homography)

$$\tilde{x}' = \tilde{H} \tilde{x}$$

where \tilde{H} is an arbitrary 4×4 homogeneous matrix

Preserves straight lines


Summary of 3D Transformations

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} \mathbf{I} & & \mathbf{t} \end{bmatrix}_{3 \times 4}$	3	orientation	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & & \mathbf{t} \end{bmatrix}_{3 \times 4}$	6	lengths	
similarity	$\begin{bmatrix} s\mathbf{R} & & \mathbf{t} \end{bmatrix}_{3 \times 4}$	7	angles	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{3 \times 4}$	12	parallelism	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{4 \times 4}$	15	straight lines	

Outline


- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ **3D rotations**
- ❖ 3D to 2D projections
- ❖ Lens Distortions

3D Rotations: Axis/Angle Representation

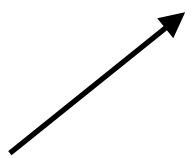



3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$



3D Rotations: Axis/Angle Representation


$$\omega = \theta \hat{n}$$




3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$


Amount of rotation



3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$


Amount of rotation



3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$

Amount of rotation Axis of rotation




3D Rotations: Axis/Angle Representation


$$\omega = \theta \hat{n}$$

Amount of rotation Axis of rotation

Let \mathbf{u} be the result of rotating vector \mathbf{v} about axis $\hat{\mathbf{n}}$ by the angle θ .




3D Rotations: Axis/Angle Representation



Let \mathbf{u} be the result of rotating vector \mathbf{v} about axis $\hat{\mathbf{n}}$ by the angle θ .

First, project the vector \mathbf{v} onto the axis $\hat{\mathbf{n}}$:



3D Rotations: Axis/Angle Representation


$$\omega = \theta \hat{n}$$

Amount of rotation Axis of rotation

Let \mathbf{u} be the result of rotating vector \mathbf{v} about axis $\hat{\mathbf{n}}$ by the angle θ .

First, project the vector \mathbf{v} onto the axis $\hat{\mathbf{n}}$:

$$\mathbf{v}_{\parallel} = \hat{\mathbf{n}}(\hat{\mathbf{n}} \cdot \mathbf{v}) = (\hat{\mathbf{n}}\hat{\mathbf{n}}^T)\mathbf{v}$$



3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$


Amount of rotation Axis of rotation

Let \mathbf{u} be the result of rotating vector \mathbf{v} about axis $\hat{\mathbf{n}}$ by the angle θ .

First, project the vector \mathbf{v} onto the axis $\hat{\mathbf{n}}$:

$$\mathbf{v}_{\parallel} = \hat{\mathbf{n}}(\hat{\mathbf{n}} \cdot \mathbf{v}) = (\hat{\mathbf{n}}\hat{\mathbf{n}}^T)\mathbf{v}$$

Next, compute the perpendicular residual:



3D Rotations: Axis/Angle Representation

$$\omega = \theta \hat{n}$$

Amount of rotation Axis of rotation


Let \mathbf{u} be the result of rotating vector \mathbf{v} about axis $\hat{\mathbf{n}}$ by the angle θ .

First, project the vector \mathbf{v} onto the axis $\hat{\mathbf{n}}$:


$$\mathbf{v}_{\parallel} = \hat{\mathbf{n}}(\hat{\mathbf{n}} \cdot \mathbf{v}) = (\hat{\mathbf{n}}\hat{\mathbf{n}}^T)\mathbf{v}$$

Next, compute the perpendicular residual:

$$\mathbf{v}_{\perp} = \mathbf{v} - \mathbf{v}_{\parallel} = (\mathbf{I} - \hat{\mathbf{n}}\hat{\mathbf{n}}^T)\mathbf{v}$$




3D Rotations: Axis/Angle Representation



3D Rotations: Axis/Angle Representation

We can rotate this vector by 90° using the cross product,

$$\mathbf{v}_\times = \hat{\mathbf{n}} \times \mathbf{v} = [\hat{\mathbf{n}}]_\times \mathbf{v},$$




3D Rotations: Axis/Angle Representation

We can rotate this vector by 90° using the cross product,

$$\mathbf{v}_\times = \hat{\mathbf{n}} \times \mathbf{v} = [\hat{\mathbf{n}}]_\times \mathbf{v},$$

where $[\hat{\mathbf{n}}]_\times$ is the matrix form of the cross product operator with the vector $\hat{\mathbf{n}} = (\hat{n}_x, \hat{n}_y, \hat{n}_z)$,




3D Rotations: Axis/Angle Representation

We can rotate this vector by 90° using the cross product,

$$\mathbf{v}_\times = \hat{\mathbf{n}} \times \mathbf{v} = [\hat{\mathbf{n}}]_\times \mathbf{v},$$

where $[\hat{\mathbf{n}}]_\times$ is the matrix form of the cross product operator with the vector $\hat{\mathbf{n}} = (\hat{n}_x, \hat{n}_y, \hat{n}_z)$,

$$[\hat{\mathbf{n}}]_\times = \begin{bmatrix} 0 & -\hat{n}_z & \hat{n}_y \\ \hat{n}_z & 0 & -\hat{n}_x \\ -\hat{n}_y & \hat{n}_x & 0 \end{bmatrix}$$



3D Rotations: Axis/Angle Representation

We can rotate this vector by 90° using the cross product,

$$\mathbf{v}_\times = \hat{\mathbf{n}} \times \mathbf{v} = [\hat{\mathbf{n}}]_\times \mathbf{v},$$

where $[\hat{\mathbf{n}}]_\times$ is the matrix form of the cross product operator with the vector $\hat{\mathbf{n}} = (\hat{n}_x, \hat{n}_y, \hat{n}_z)$,


$$[\hat{\mathbf{n}}]_\times = \begin{bmatrix} 0 & -\hat{n}_z & \hat{n}_y \\ \hat{n}_z & 0 & -\hat{n}_x \\ -\hat{n}_y & \hat{n}_x & 0 \end{bmatrix}$$

Note that rotating this vector by another 90° is equivalent to taking the cross product again,


$$\mathbf{v}_{\times\times} = \hat{\mathbf{n}} \times \mathbf{v}_\times = [\hat{\mathbf{n}}]_\times^2 \mathbf{v} = -\mathbf{v}_\perp,$$

and hence

$$\mathbf{v}_\parallel = \mathbf{v} - \mathbf{v}_\perp = \mathbf{v} + \mathbf{v}_{\times\times} = (\mathbf{I} + [\hat{\mathbf{n}}]_\times^2) \mathbf{v}.$$



3D Rotations: Axis/Angle Representation




3D Rotations: Axis/Angle Representation

We can now compute the in-plane component of the rotated vector \mathbf{u} as

$$\mathbf{u}_{\perp} = \cos \theta \mathbf{v}_{\perp} + \sin \theta \mathbf{v}_{\times} = (\sin \theta [\hat{\mathbf{n}}]_{\times} - \cos \theta [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

Putting all these terms together, we obtain the final rotated vector as

$$\mathbf{u} = \mathbf{u}_{\perp} + \mathbf{v}_{\parallel} = (\mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$



3D Rotations: Axis/Angle Representation


We can now compute the in-plane component of the rotated vector \mathbf{u} as

$$\mathbf{u}_{\perp} = \cos \theta \mathbf{v}_{\perp} + \sin \theta \mathbf{v}_{\times} = (\sin \theta [\hat{\mathbf{n}}]_{\times} - \cos \theta [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

Putting all these terms together, we obtain the final rotated vector as

$$\mathbf{u} = \mathbf{u}_{\perp} + \mathbf{v}_{\parallel} = (\mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

We can therefore write the rotation matrix corresponding to a rotation by θ around an axis $\hat{\mathbf{n}}$ as



3D Rotations: Axis/Angle Representation

We can now compute the in-plane component of the rotated vector \mathbf{u} as


$$\mathbf{u}_{\perp} = \cos \theta \mathbf{v}_{\perp} + \sin \theta \mathbf{v}_{\times} = (\sin \theta [\hat{\mathbf{n}}]_{\times} - \cos \theta [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

Putting all these terms together, we obtain the final rotated vector as

$$\mathbf{u} = \mathbf{u}_{\perp} + \mathbf{v}_{\parallel} = (\mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

We can therefore write the rotation matrix corresponding to a rotation by θ around an axis $\hat{\mathbf{n}}$ as

$$\mathbf{R}(\hat{\mathbf{n}}, \theta) = \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2$$



3D Rotations: Axis/Angle Representation

We can now compute the in-plane component of the rotated vector \mathbf{u} as


$$\mathbf{u}_{\perp} = \cos \theta \mathbf{v}_{\perp} + \sin \theta \mathbf{v}_{\times} = (\sin \theta [\hat{\mathbf{n}}]_{\times} - \cos \theta [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

Putting all these terms together, we obtain the final rotated vector as

$$\mathbf{u} = \mathbf{u}_{\perp} + \mathbf{v}_{\parallel} = (\mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2) \mathbf{v}.$$

We can therefore write the rotation matrix corresponding to a rotation by θ around an axis $\hat{\mathbf{n}}$ as

$$\mathbf{R}(\hat{\mathbf{n}}, \theta) = \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2 \quad \text{(Rodriguez' formula)}$$




3D Rotations: Axis/Angle Representation

$$\mathbf{R}(\hat{\mathbf{n}}, \theta) = \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2 \quad \text{(Rodriguez' formula)}$$

For small rotations:

$$\mathbf{R}(\boldsymbol{\omega}) \approx \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} \approx \mathbf{I} + [\theta \hat{\mathbf{n}}]_{\times} = \begin{bmatrix} 1 & -\omega_z & \omega_y \\ \omega_z & 1 & -\omega_x \\ -\omega_y & \omega_x & 1 \end{bmatrix}$$



Unit Quaternions

$$\mathbf{q} = (x, y, z, w) \text{ with } \|\mathbf{q}\| = 1$$

\mathbf{q} and $-\mathbf{q}$ represent the same rotation.


Quaternions can be derived from the axis/angle representation through the formula

$$\mathbf{q} = (\mathbf{v}, w) = \left(\sin \frac{\theta}{2} \hat{\mathbf{n}}, \cos \frac{\theta}{2} \right),$$

where $\hat{\mathbf{n}}$ and θ are the rotation axis and angle.

Rodriguez' formula now becomes (see textbook):

$$\begin{aligned} \mathbf{R}(\hat{\mathbf{n}}, \theta) &= \mathbf{I} + \sin \theta [\hat{\mathbf{n}}]_{\times} + (1 - \cos \theta) [\hat{\mathbf{n}}]_{\times}^2 \\ &= \mathbf{I} + 2w [\mathbf{v}]_{\times} + 2[\mathbf{v}]_{\times}^2. \end{aligned}$$



Quaternion Algebra


$$\mathbf{q} = (\mathbf{v}, w) = \left(\sin \frac{\theta}{2} \hat{\mathbf{n}}, \cos \frac{\theta}{2} \right)$$

Composition (multiplication): $\mathbf{q}_2 = \mathbf{q}_0 \mathbf{q}_1 = (\mathbf{v}_0 \times \mathbf{v}_1 + w_0 \mathbf{v}_1 + w_1 \mathbf{v}_0, w_0 w_1 - \mathbf{v}_0 \cdot \mathbf{v}_1)$

$$\mathbf{R}(\mathbf{q}_2) = \mathbf{R}(\mathbf{q}_0) \mathbf{R}(\mathbf{q}_1).$$

Inverse: flip the sign of \mathbf{v} or w (but not both).


i.e., if $\mathbf{q} = (\mathbf{v}, w)$, then $\mathbf{q}^{-1} = (-\mathbf{v}, w) = (\mathbf{v}, -w)$.



Outline


- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ **3D to 2D projections**
- ❖ Lens Distortions

Orthographic (Parallel) Projection




Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.




Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)




Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!



Orthographic (Parallel) Projection


- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!
- ❖ In inhomogeneous (Euclidean) coordinates:



Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!
- ❖ In inhomogeneous (Euclidean) coordinates:

$$\mathbf{x} = [\mathbf{I}_{2 \times 2} | \mathbf{0}] \mathbf{p}$$




Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!
- ❖ In inhomogeneous (Euclidean) coordinates:


$$\mathbf{x} = [\mathbf{I}_{2 \times 2} | \mathbf{0}] \mathbf{p}$$

where \mathbf{p} is the 3D point and \mathbf{x} is the projected 2D image point




Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!
- ❖ In inhomogeneous (Euclidean) coordinates:
$$\mathbf{x} = [\mathbf{I}_{2 \times 2} | \mathbf{0}] \mathbf{p}$$
where \mathbf{p} is the 3D point and \mathbf{x} is the projected 2D image point
- ❖ In practice, we also need to scale the x and y coordinates from metres to pixels:



Orthographic (Parallel) Projection

- ❖ Reasonable approximation to perspective projection when % depth variation within field of view is small.
- ❖ This is often the case for telephoto lenses (long viewing distances, small field of view)
- ❖ Given camera-aligned world coordinate frame, simply drop the z component!



- ❖ In inhomogeneous (Euclidean) coordinates:


$$\mathbf{x} = [\mathbf{I}_{2 \times 2} | \mathbf{0}] \mathbf{p}$$

where \mathbf{p} is the 3D point and \mathbf{x} is the projected 2D image point

- ❖ In practice, we also need to scale the x and y coordinates from metres to pixels:


$$\mathbf{x} = [s\mathbf{I}_{2 \times 2} | \mathbf{0}] \mathbf{p}.$$

Perspective Projection



Perspective Projection


- ❖ Points projected onto image plane by dividing them by their z component.



Perspective Projection

- ❖ Points projected onto image plane by dividing them by their z component.

$$\bar{\mathbf{x}} = \mathcal{P}_z(\mathbf{p}) = \begin{bmatrix} x/z \\ y/z \\ 1 \end{bmatrix}$$




Perspective Projection

- ❖ Points projected onto image plane by dividing them by their z component.

$$\bar{x} = \mathcal{P}_z(\mathbf{p}) = \begin{bmatrix} x/z \\ y/z \\ 1 \end{bmatrix}$$

- ❖ In homogeneous coordinates:




Perspective Projection

- ❖ Points projected onto image plane by dividing them by their z component.

$$\bar{x} = \mathcal{P}_z(\mathbf{p}) = \begin{bmatrix} x/z \\ y/z \\ 1 \end{bmatrix}$$

- ❖ In homogeneous coordinates:

$$\tilde{x} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tilde{p}$$



Camera Intrinsics

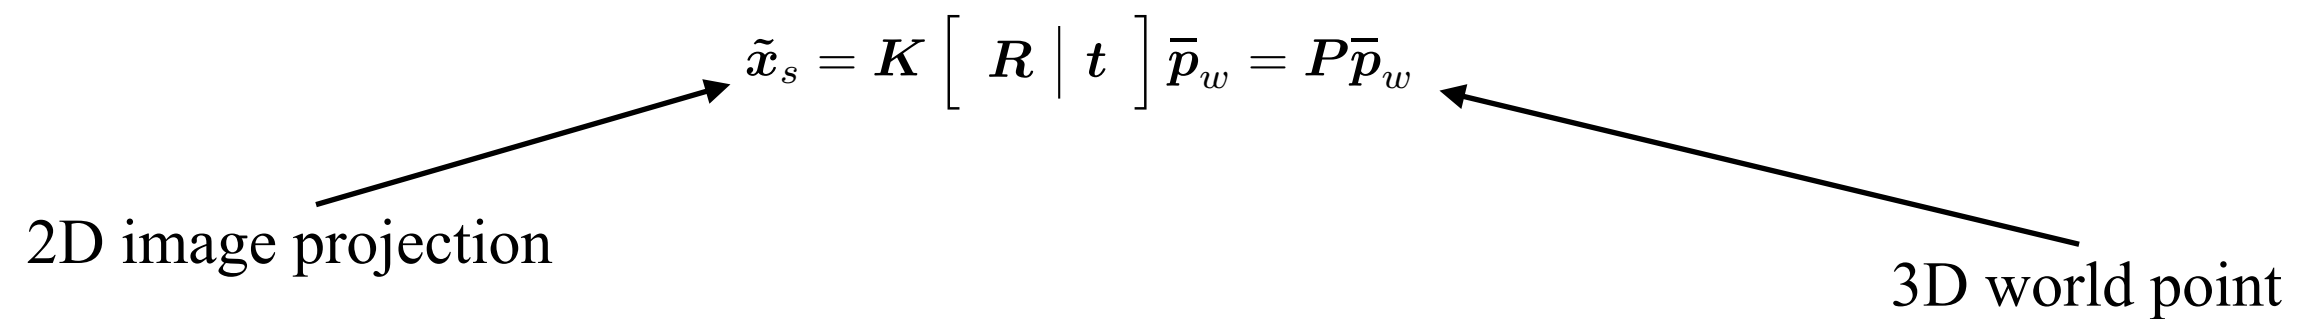
$$\tilde{\mathbf{x}}_s = \mathbf{K} \left[\mathbf{R} \mid \mathbf{t} \right] \bar{\mathbf{p}}_w = \mathbf{P} \bar{\mathbf{p}}_w$$

Camera Intrinsics


$$\tilde{\mathbf{x}}_s = \mathbf{K} \left[\mathbf{R} \mid \mathbf{t} \right] \bar{\mathbf{p}}_w = \mathbf{P} \bar{\mathbf{p}}_w$$

3D world point 


Camera Intrinsics




Camera Intrinsics



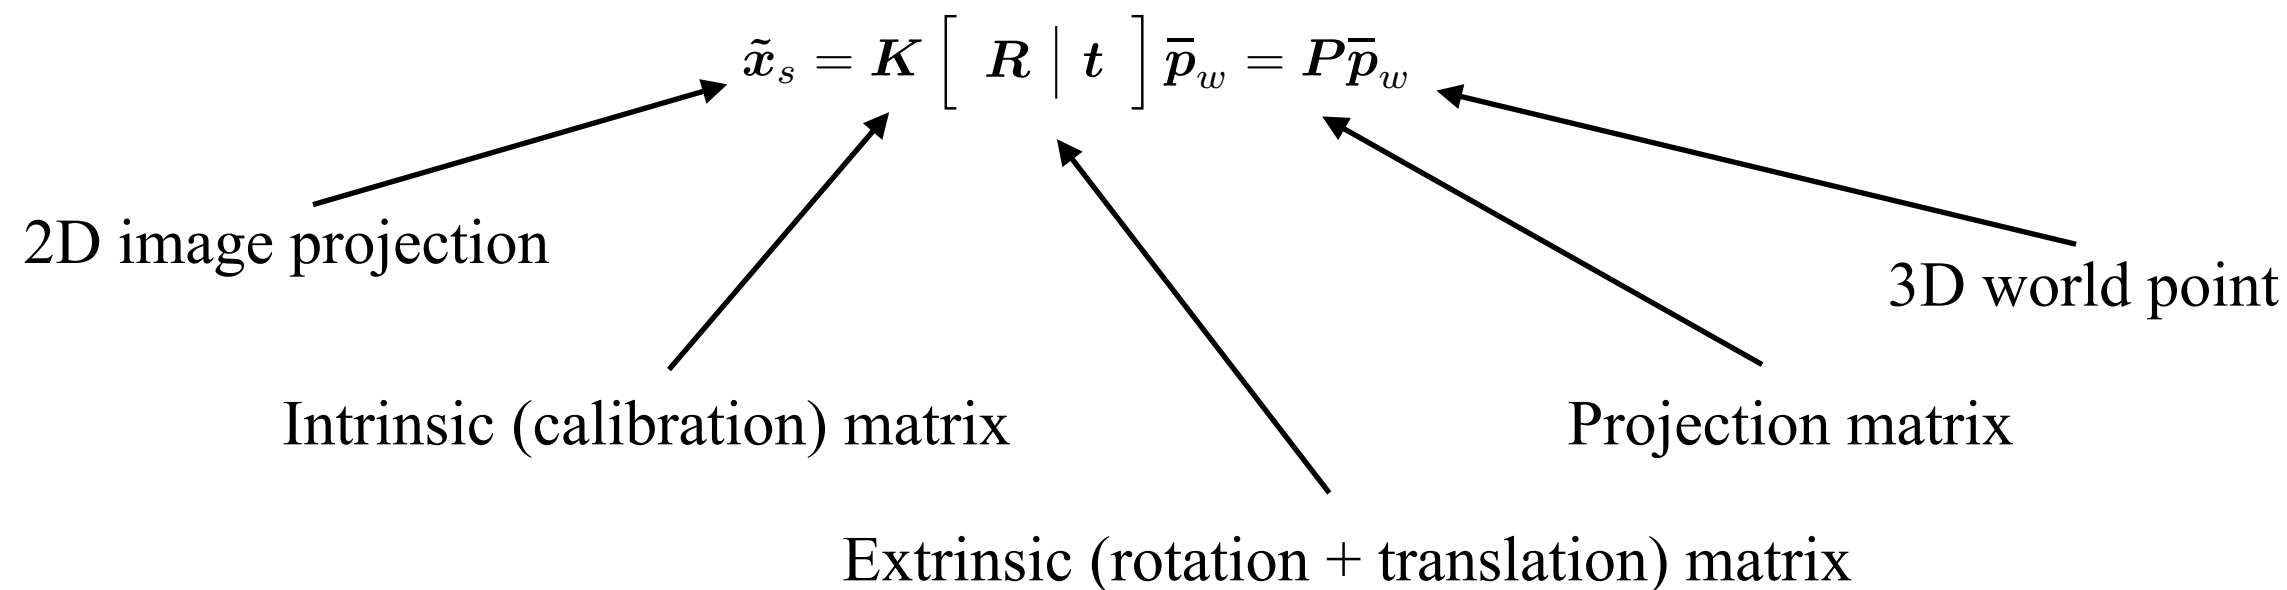
Camera Intrinsics



Camera Intrinsics

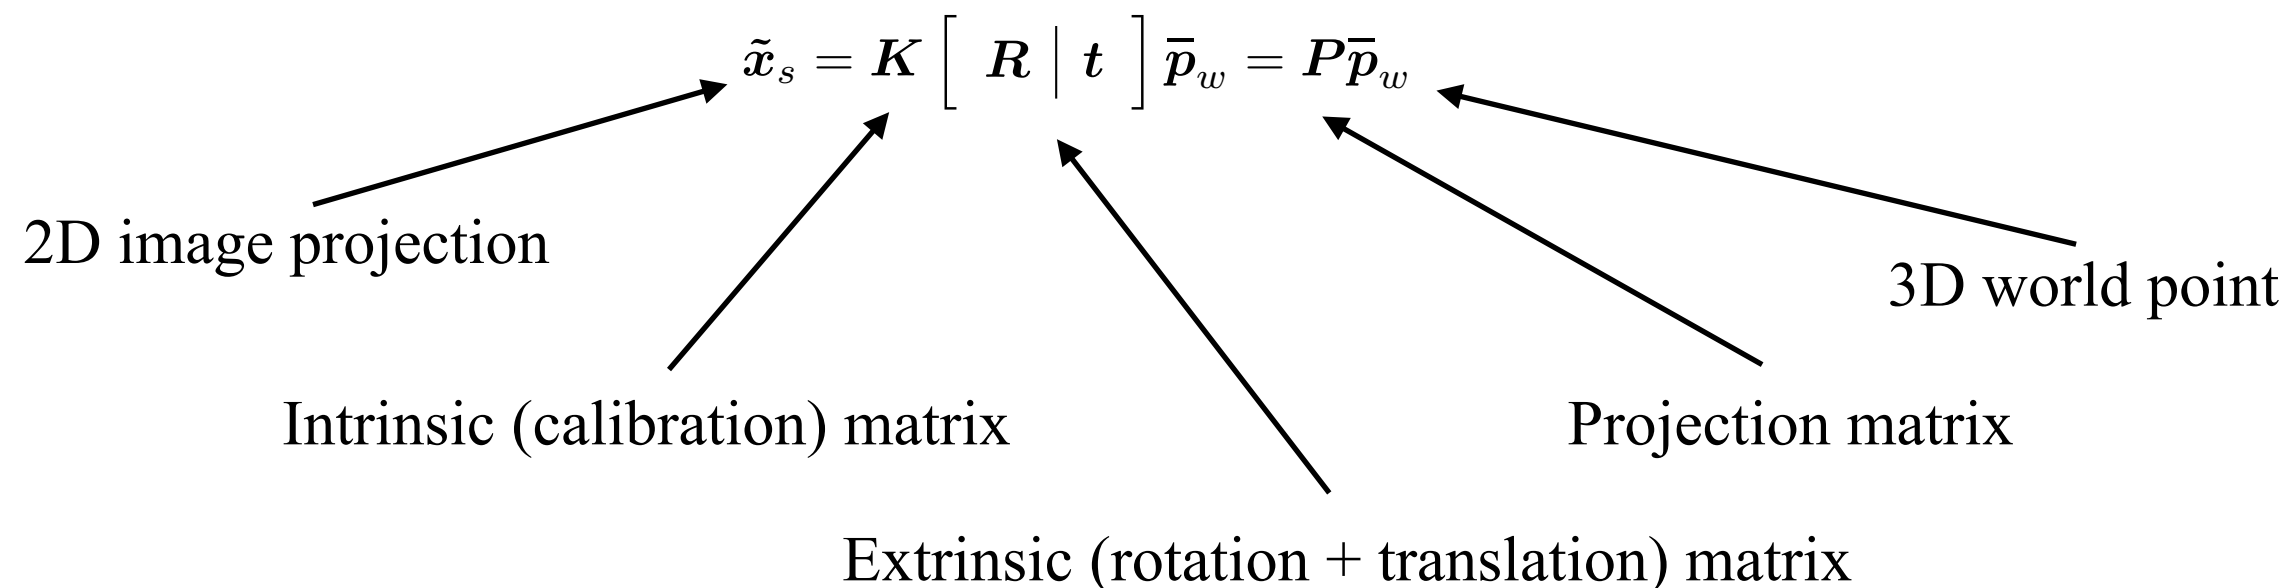


Camera Intrinsics



$$K = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

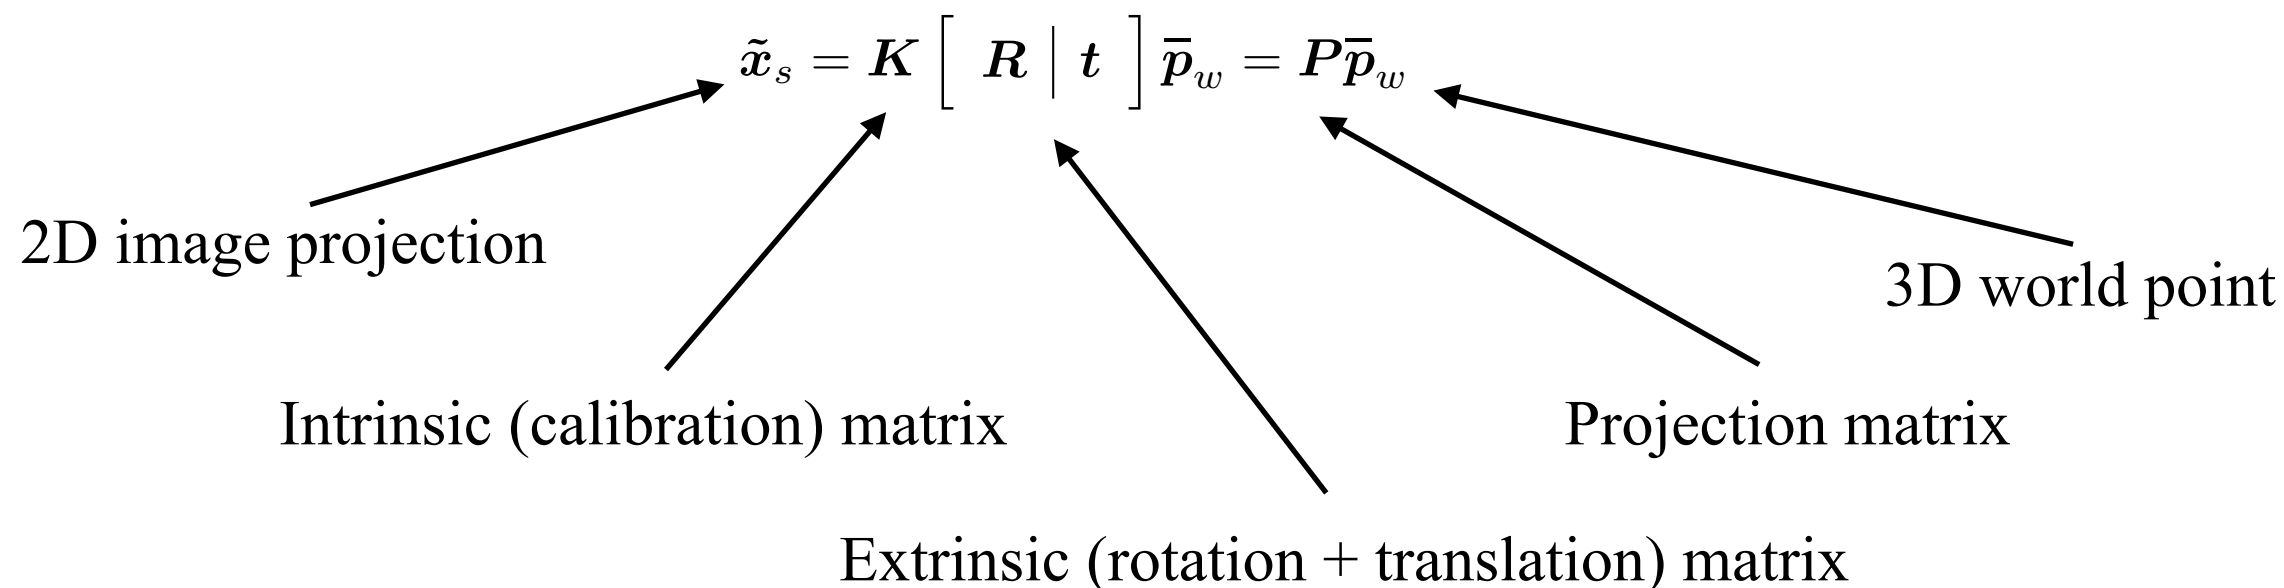
Camera Intrinsics



$$\mathbf{K} = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

f_x and f_y : encode focal length and pixel spacing, which may be slightly different in x and y dimensions.

Camera Intrinsics




$$K = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

f_x and f_y : encode focal length and pixel spacing, which may be slightly different in x and y dimensions.

c_x and c_y : encode principal point (intersection of optic axis with sensor plane) - usually very close to centre of image

Camera Intrinsics




$$K = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

f_x and f_y : encode focal length and pixel spacing, which may be slightly different in x and y dimensions.

c_x and c_y : encode principal point (intersection of optic axis with sensor plane) - usually very close to centre of image

s : encodes possible skew between sensor axes (usually close to 0).

Camera Intrinsics




$$K = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

f_x and f_y : encode focal length and pixel spacing, which may be slightly different in x and y dimensions.

c_x and c_y : encode principal point (intersection of optic axis with sensor plane) - usually very close to centre of image


s : encodes possible skew between sensor axes (usually close to 0).



Focal Lengths

- ❖ Focal length can be measured either in pixels or in mm.


$$\tan \frac{\theta}{2} = \frac{W}{2f} \quad \text{or} \quad f = \frac{W}{2} \left[\tan \frac{\theta}{2} \right]^{-1}$$



Focal Lengths

- ❖ Focal length can be measured either in pixels or in mm.

$$\tan \frac{\theta}{2} = \frac{W}{2f} \quad \text{or} \quad f = \frac{W}{2} \left[\tan \frac{\theta}{2} \right]^{-1}$$



- ❖ Example: Consider the FLIR BlackFly S BFS-PGE-122S6C-C paired with a 10mm lens:

- ❖ Resolution: 4096 x 3000 pixels
- ❖ Sensor width: 1.1" = 27.94mm



Outline

- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ **Lens Distortions**

Lens Distortions

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.
- ❖ This is not true in real cameras, due to lens distortions.

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.
- ❖ This is not true in real cameras, due to lens distortions.
- ❖ Wide-angle lenses produce noticeable radial distortion

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.
- ❖ This is not true in real cameras, due to lens distortions.
- ❖ Wide-angle lenses produce noticeable radial distortion
- ❖ Let (x_c, y_c) be image coordinates after perspective projection but before scaling by focal length and shifting to the optical centre.

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.
- ❖ This is not true in real cameras, due to lens distortions.
- ❖ Wide-angle lenses produce noticeable radial distortion
- ❖ Let (x_c, y_c) be image coordinates after perspective projection but before scaling by focal length and shifting but the optical centre.
- ❖ Then without distortion, we should have

Lens Distortions

- ❖ In perspective projection, straight lines project to straight lines.
- ❖ This is not true in real cameras, due to lens distortions.
- ❖ Wide-angle lenses produce noticeable radial distortion
- ❖ Let (x_c, y_c) be image coordinates after perspective projection but before scaling by focal length and shifting but he optical centre.
- ❖ Then without distortion, we should have

$$x_c = \frac{\mathbf{r}_x \cdot \mathbf{p} + t_x}{\mathbf{r}_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{\mathbf{r}_y \cdot \mathbf{p} + t_y}{\mathbf{r}_z \cdot \mathbf{p} + t_z}$$

Radial Distortion



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$

- ❖ In radial distortion, points are displaced radially by an amount that increases with their distance from the image centre



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$

- ❖ In radial distortion, points are displaced radially by an amount that increases with their distance from the image centre
 - ⦿ Barrel distortion: points are displaced away from the image centre



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$

- ❖ In radial distortion, points are displaced radially by an amount that increases with their distance from the image centre
 - Barrel distortion: points are displaced away from the image centre
 - Pincushion distortion: points are displaced towards the image centre



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$

- ❖ In radial distortion, points are displaced radially by an amount that increases with their distance from the image centre
 - Barrel distortion: points are displaced away from the image centre
 - Pincushion distortion: points are displaced towards the image centre
 - Radial distortion can be modelled by a 4th-order perturbation on these coordinates:



Radial Distortion

$$x_c = \frac{r_x \cdot \mathbf{p} + t_x}{r_z \cdot \mathbf{p} + t_z}$$
$$y_c = \frac{r_y \cdot \mathbf{p} + t_y}{r_z \cdot \mathbf{p} + t_z}$$

- ❖ In radial distortion, points are displaced radially by an amount that increases with their distance from the image centre
 - Barrel distortion: points are displaced away from the image centre
 - Pincushion distortion: points are displaced towards the image centre
 - Radial distortion can be modelled by a 4th-order perturbation on these coordinates:

$$\hat{x}_c = x_c(1 + \kappa_1 r_c^2 + \kappa_2 r_c^4)$$
$$\hat{y}_c = y_c(1 + \kappa_1 r_c^2 + \kappa_2 r_c^4),$$



Outline

- ❖ Geometric primitives
- ❖ 2D transformations
- ❖ 3D transformations
- ❖ 3D rotations
- ❖ 3D to 2D projections
- ❖ Lens Distortions