CIS 580

Machine Perception

Or *Geometric Computer Vision*

Instructor: Lingjie Liu

Lec 2: Jan 22, 2025

Spring 2025 Tentative Schedule

Week	Date*	Topics
0-1	Jan 15, 22	
2	Jan 27	
3	Feb 3	
4	Feb 10	
5	Feb 17	
6	Feb 24	
7	Mar 3	
8	Mar 10 (Spring Break)	
9	Mar 17	
10	Mar 24	**Midterm (12%, on March 26)**
11	Mar 31	
12	Apr 7	
13	Apr 14	
14	Apr 21	
15	Apr 28	
16	Btw May 3-13	**Final Exam (12%)**

HW1 (Feb 3 – Feb 17)

12%

HW2 (Feb 17 – Mar 3)

12%

HW3 (March 3 – March 24)

12%

HW4 (March 31– Apr 14)

10%

two small projects(Apr 14 – May 14)

30%

HW will be released on Wed of that week and will be due two weeks later, on Wed.

HW plans might evolve

Recap: Basic Perspective Projection Equations

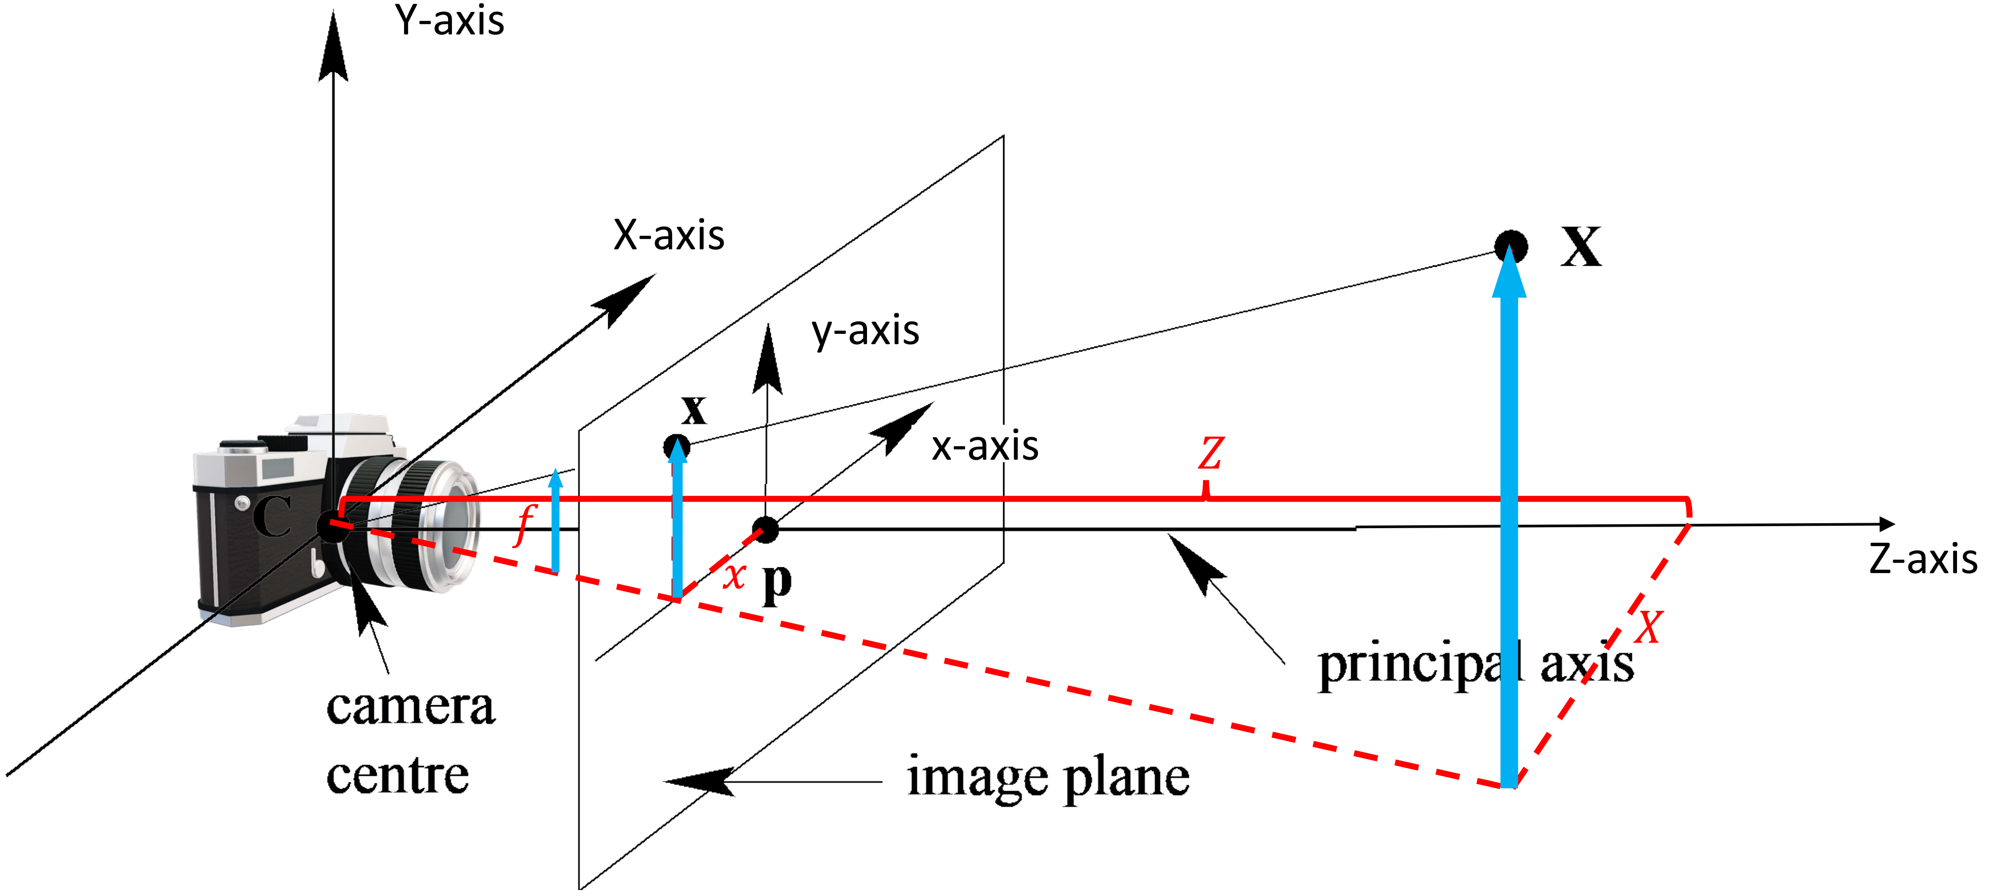

$$x = f\frac{X}{Z}, y = f\frac{Y}{Z}$$

Z&H Ch6

Recap: Perspective Effects

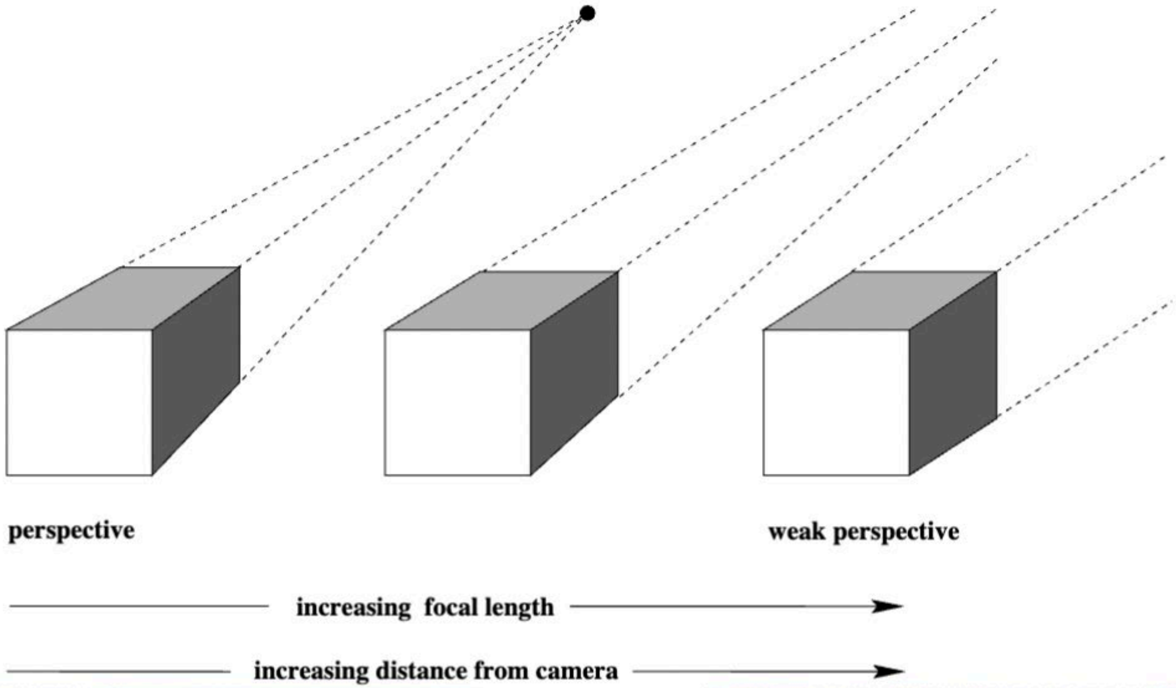

ZH Ch6

Review: (Euclidean) Geometric Concepts Through Matrix-Vector Algebra

Recap: Basic Perspective Projection Equations

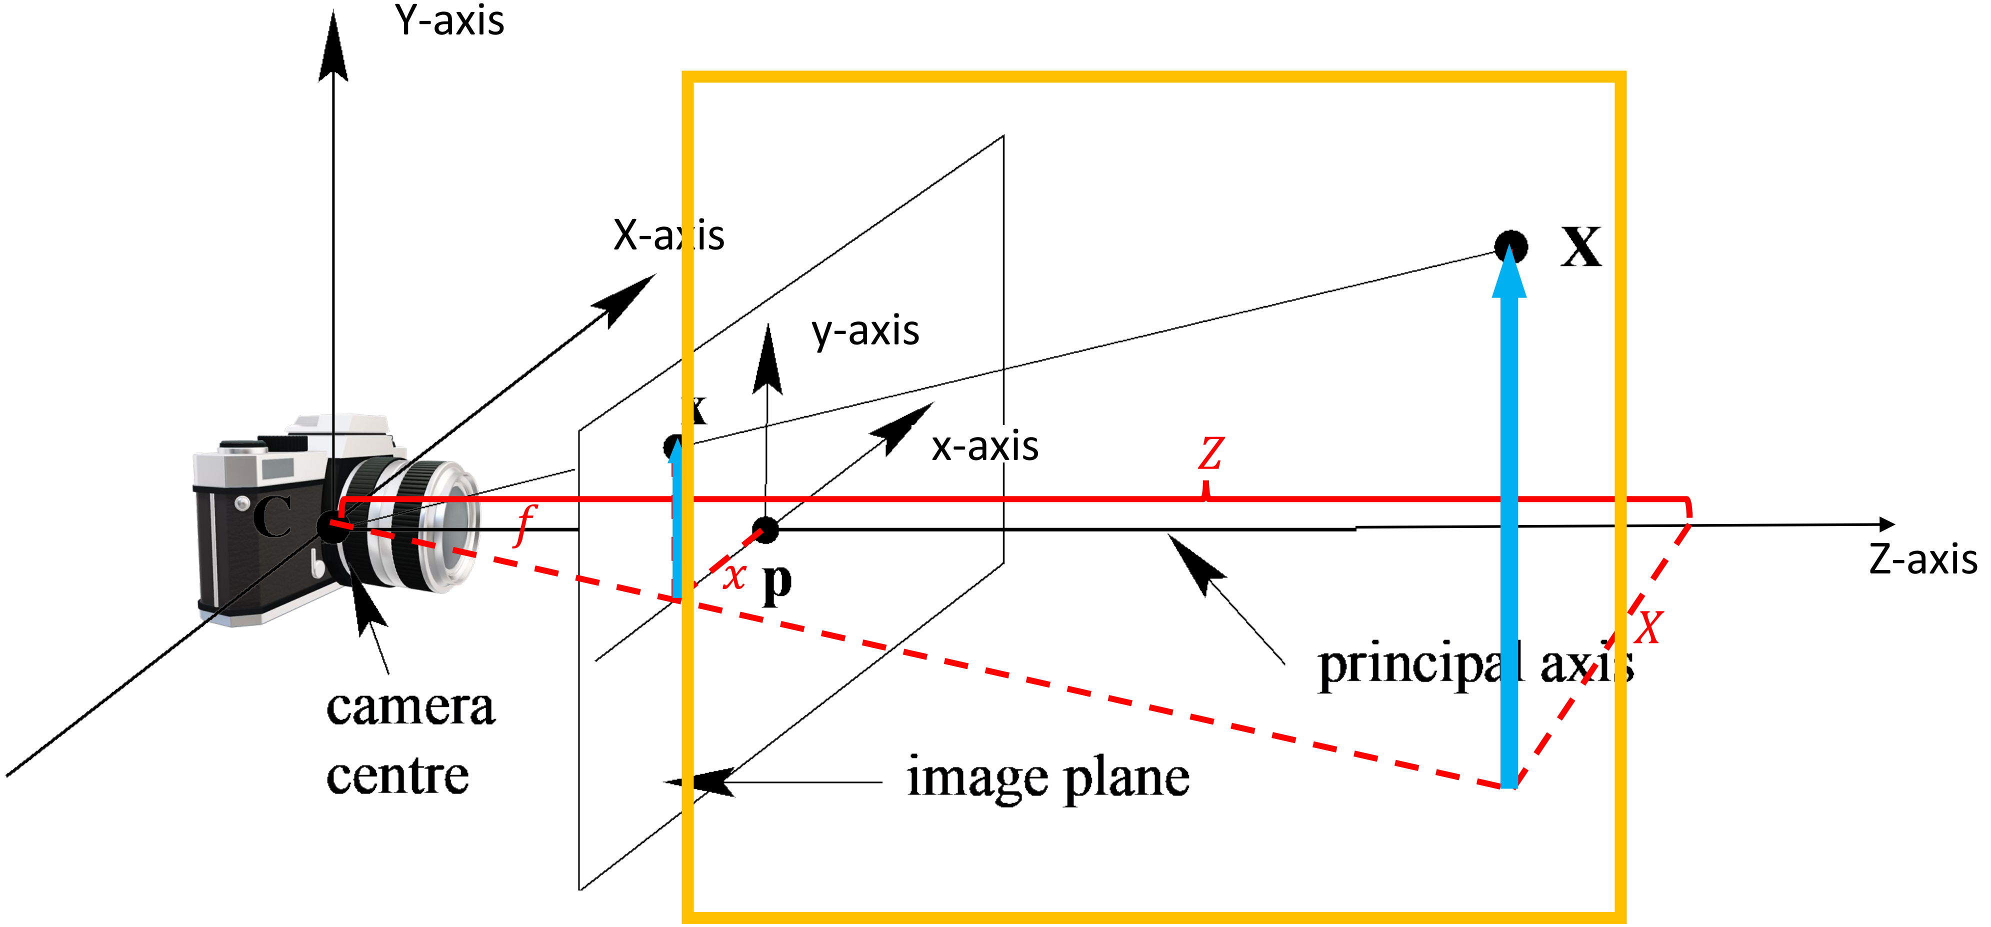

$$x = f\frac{X}{Z}, y = f\frac{Y}{Z}$$

Z&H Ch6

Points and Lines in Euclidean 2D (shortcut $\mathbb{R}^2$)

- Point in 2D space: $\boldsymbol{x} = (x, y)^T \in \mathbb{R}^2$
  - Projection onto coordinate axes starting from an origin point $(0,0)$.
  - "Cartesian coordinate system": requires origin and coordinate frame.
- Line in 2D space:
  - Collection of all points $(x, y)^T$ that satisfy an equation $ax + by + c = 0$
  - Can be indicated by the 3-D column vector: $\boldsymbol{l} = [a, b, c]^T$.
  - The equation of a line can then be written as $\boldsymbol{l}^T \begin{pmatrix} x \\ y \\ {\color{red}1} \end{pmatrix} = 0$. (dot product)

Points and Planes in Euclidean 3D (shortcut $\mathbb{R}^3$)

- Point in 3D space: $\boldsymbol{x} = (x, y, z)^T \in \mathbb{R}^3$

- Plane in 3D space:
  - Collection of all points $(x, y, z)^T$ that satisfy an equation $ax + by + cz + d = 0$
  - Can be indicated by the 4-D column vector: $\boldsymbol{\pi} = [a, b, c, d]^T$.
  - The equation of a plane can then be written as $\boldsymbol{\pi}^T \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = 0$.

Lines in $\mathbb{R}^3$

- Line in 3D space:
  - Interpolate between two points $(x_1, y_1, z_1)^T$ and $(x_2, y_2, z_2)^T$.
  - The set of all points that satisfy $(x, y, z)^T = (x_1, y_1, z_1)^T + \lambda(x_2 - x_1, y_2 - y_1, z_2 - z_1)^T$ for some value of λ
  - As a special case of interest: lines through origin: $(x, y, z)^T = \lambda(x_2, y_2, z_2)^T$

Lines in $\mathbb{R}^3$

- Alternative definition: intersection of two planes in 3D, so collection of all points that satisfy two equations:
  - $\boldsymbol{\pi}_1^T \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = 0 \text{ and } \boldsymbol{\pi}_2^T \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = 0$
  - Can rewrite in matrix form as $\begin{pmatrix} \boldsymbol{\pi}_1^T \\ \boldsymbol{\pi}_2^T \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$

(Line = 1DOF. Needs two linear equality constraints to take away 2DOF. Then 3DOF (because 3D) – 2DOF = 1DOF.)

Transformations in $\mathbb{R}^3$

- **Shift (Translation)**:
- 3 degrees of freedom (DOF)
- Effect on points: $(x,y,z)^T \rightarrow (x,y,z)^T + \boldsymbol{t} = (x,y,z)^T + (x_0, y_0, z_0)^T$
  - From this, can derive equations for how it affects planes, lines etc.

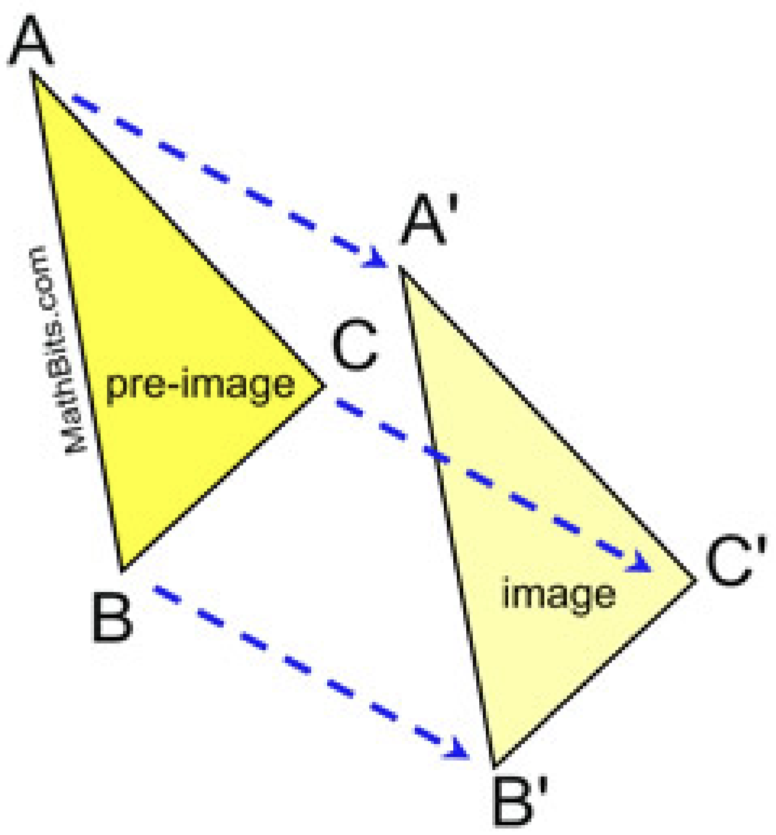

$$(x,y,z)^T = (x_1, y_1, z_1)^T + \lambda(x_2 - x_1, y_2 - y_1, z_2 - z_1)^T$$

Transformations in $\mathbb{R}^3$

The determinant of a 2×2 matrix is

$$\begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc,$$

and the determinant of a 3×3 matrix is

$$\begin{vmatrix} a & b & c \\ d & e & f \\ g & h & i \end{vmatrix} = aei + bfg + cdh - ceg - bdi - afh.$$

- **Rotation**:
- 3DOF
- Effect on points: $(x, y, z)^T \rightarrow R_{3\times3}(x, y, z)^T$
  - Where $R_{3\times3} = [\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3]$ is an *orthonormal* matrix, satisfying:
    - $R^T R = I$
    - Or equivalently, $\boldsymbol{r}_i^T \boldsymbol{r}_i = ||\boldsymbol{r}_i||_2 = 1$ and $\boldsymbol{r}_i^T \boldsymbol{r}_j = 0$ for all $i \neq j$
    - Determinant: Det($R_{3\times3}$) = 1
    - 9 entries in R matrix, but only 3 actual DOF.
  - Can be expressed in terms of geometric angular rotations in various formalisms. (See: https://en.wikipedia.org/wiki/Rotation_formalisms_in_three_dimensions)

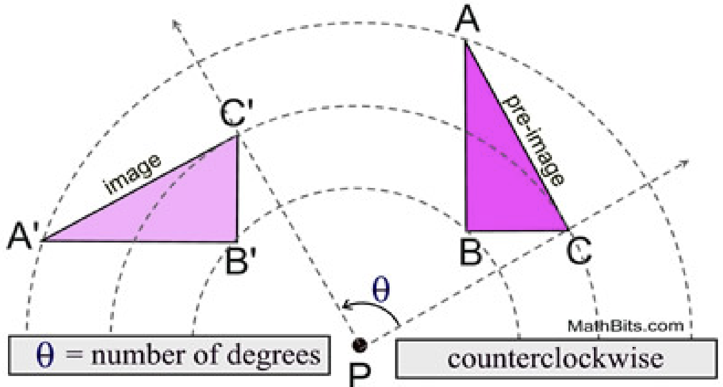

Transformations in $\mathbb{R}^3$

- **Reflection**
- Determinant: Det($R_{3\times 3}$) = -1

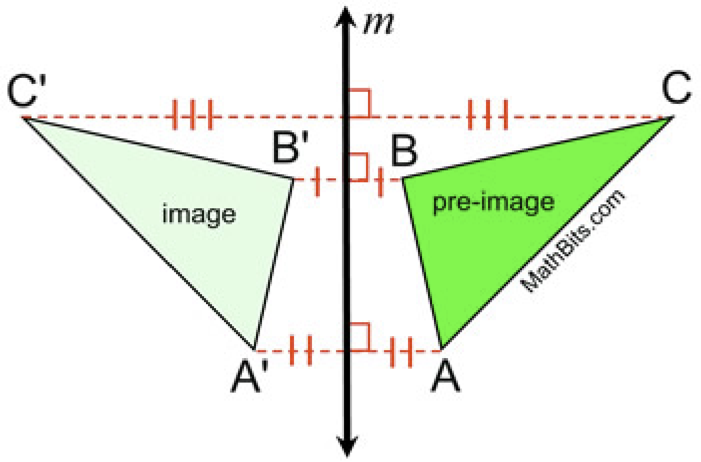

How many DOF does a reflection transformation in $\mathbb{R}^3$ have?

Transformations in $\mathbb{R}^3$

- “Euclidean Transformations” – **Rigid Transformation**
- The rigid transformations include rotations, translations, reflections, or any sequence of these. **Reflections are sometimes excluded from the definition of a rigid transformation** by requiring that the transformation also preserve the handedness of objects in the Euclidean space. (A reflection would not preserve handedness; for instance, it would transform a left hand into a right hand.) To avoid ambiguity, a transformation that preserves handedness is known as a **rigid motion**, a **Euclidean motion**, or a **proper rigid transformation**.
- arbitrary compositions of rotations and shifts. 6 DOF
  - Special because they preserve all lengths (“Isometry”), angles, perimeter, and area.

Transformations in $\mathbb{R}^3$

- **Scaling**- 1DOF
  - Effect on points: $(x, y, z)^T \rightarrow \lambda(x, y, z)^T$

- **"Similarity Transformations"** = combinations of rotations, shifts, scaling.
- 3+3+1=7DOF. Preserves lengths up to a scaling factor

- **"Affine Transformations"**:
  - Remove the orthonormality constraint on R. It is no longer meaningfully treatable as a "rotation". So we have a new, arbitrary 3x3 matrix A
    - $(x, y, z)^T \rightarrow A_{3\times3}(x, y, z)^T + \boldsymbol{t}$
    - 12 DOF = 9 DOF for A + 3 DOF for translation
    - Note: does it make sense to do "scaling" any more?
      - Hint: it is equivalently just absorbed into $A_{3\times3}$.
  - Still preserves parallel lines

Transformations in $\mathbb{R}^3$

- **"Affine Transformations"**: $(x, y, z)^T \rightarrow A_{3\times3}(x, y, z)^T + \boldsymbol{t}$
  - preserves parallel lines

Proof?

Hierarchy of Transformations

Affine

Minus Orthonormality Constraint

Similarity

Plus Scaling

Euclidean

Shifts

Rotations

Soon an even more general transformation (“projective”)!

Recap: Basic Perspective Projection Equations

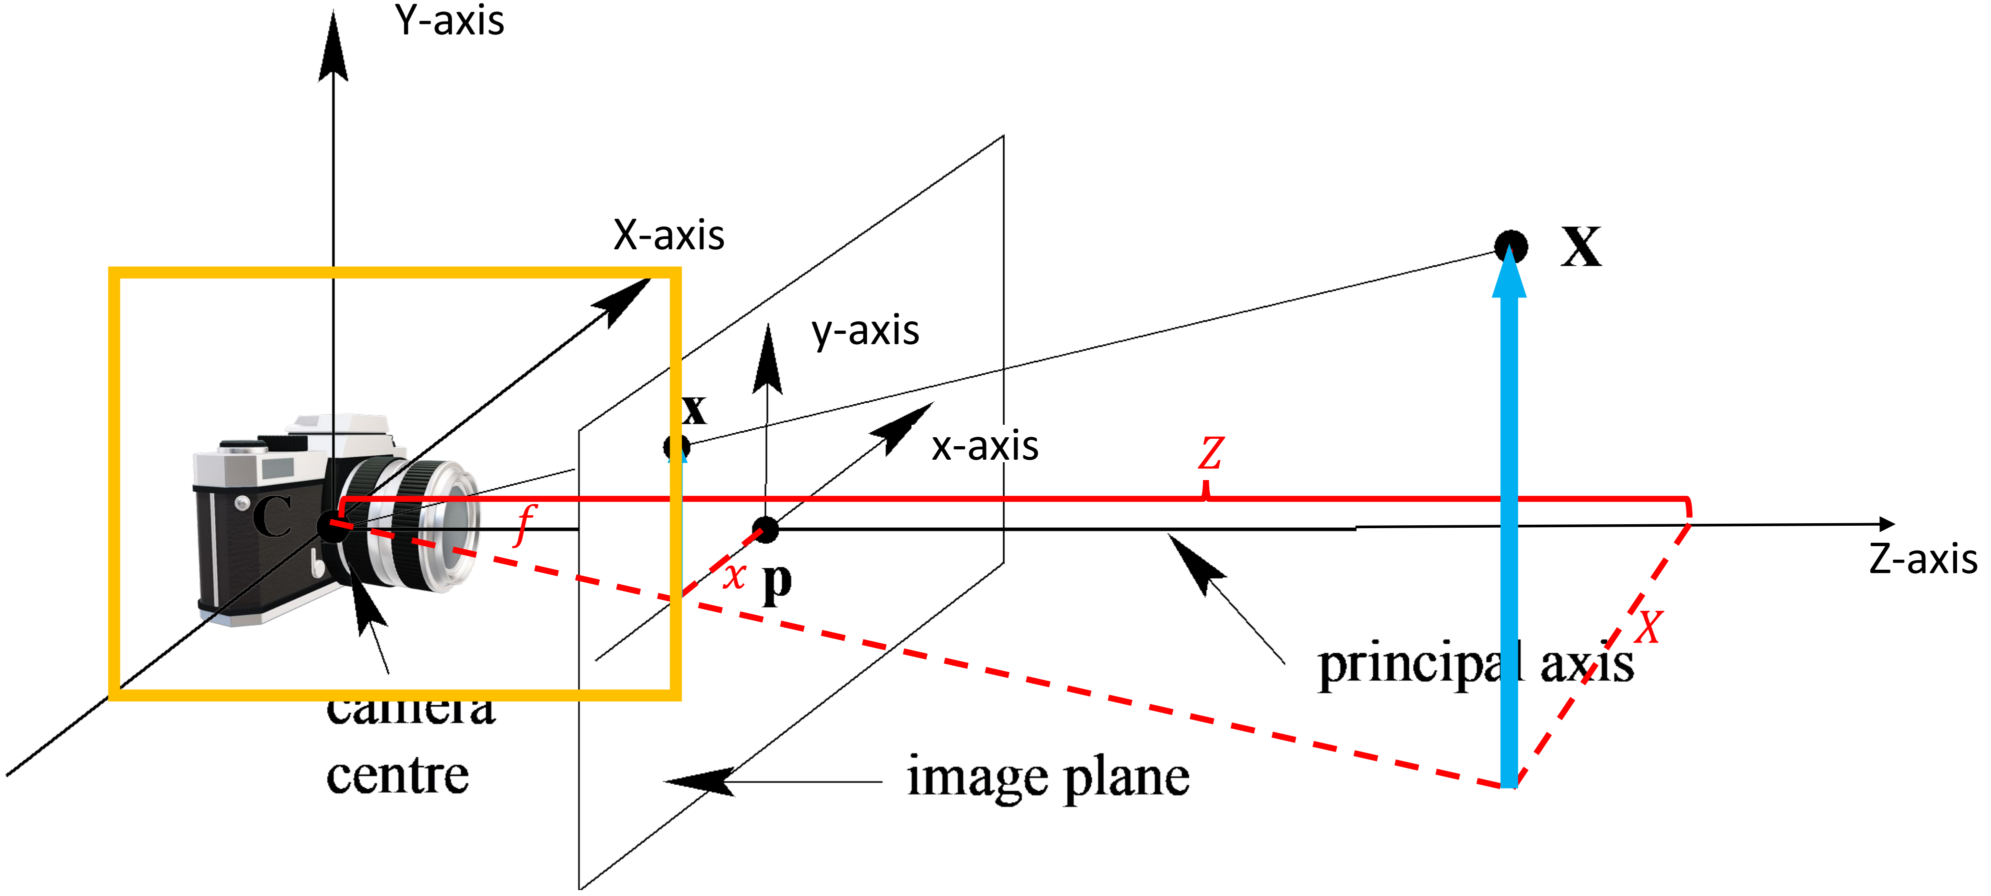

$$x = f\frac{X}{Z}, y = f\frac{Y}{Z}$$

Projective Geometry

Based on slides by Jianbo Shi, Hyun Soo Park, Kostas Daniilidis

Vanishing points

A vanishing point is a point on the image plane of a perspective rendering where the two-dimensional perspective projections of mutually parallel lines in three-dimensional space appear to converge.

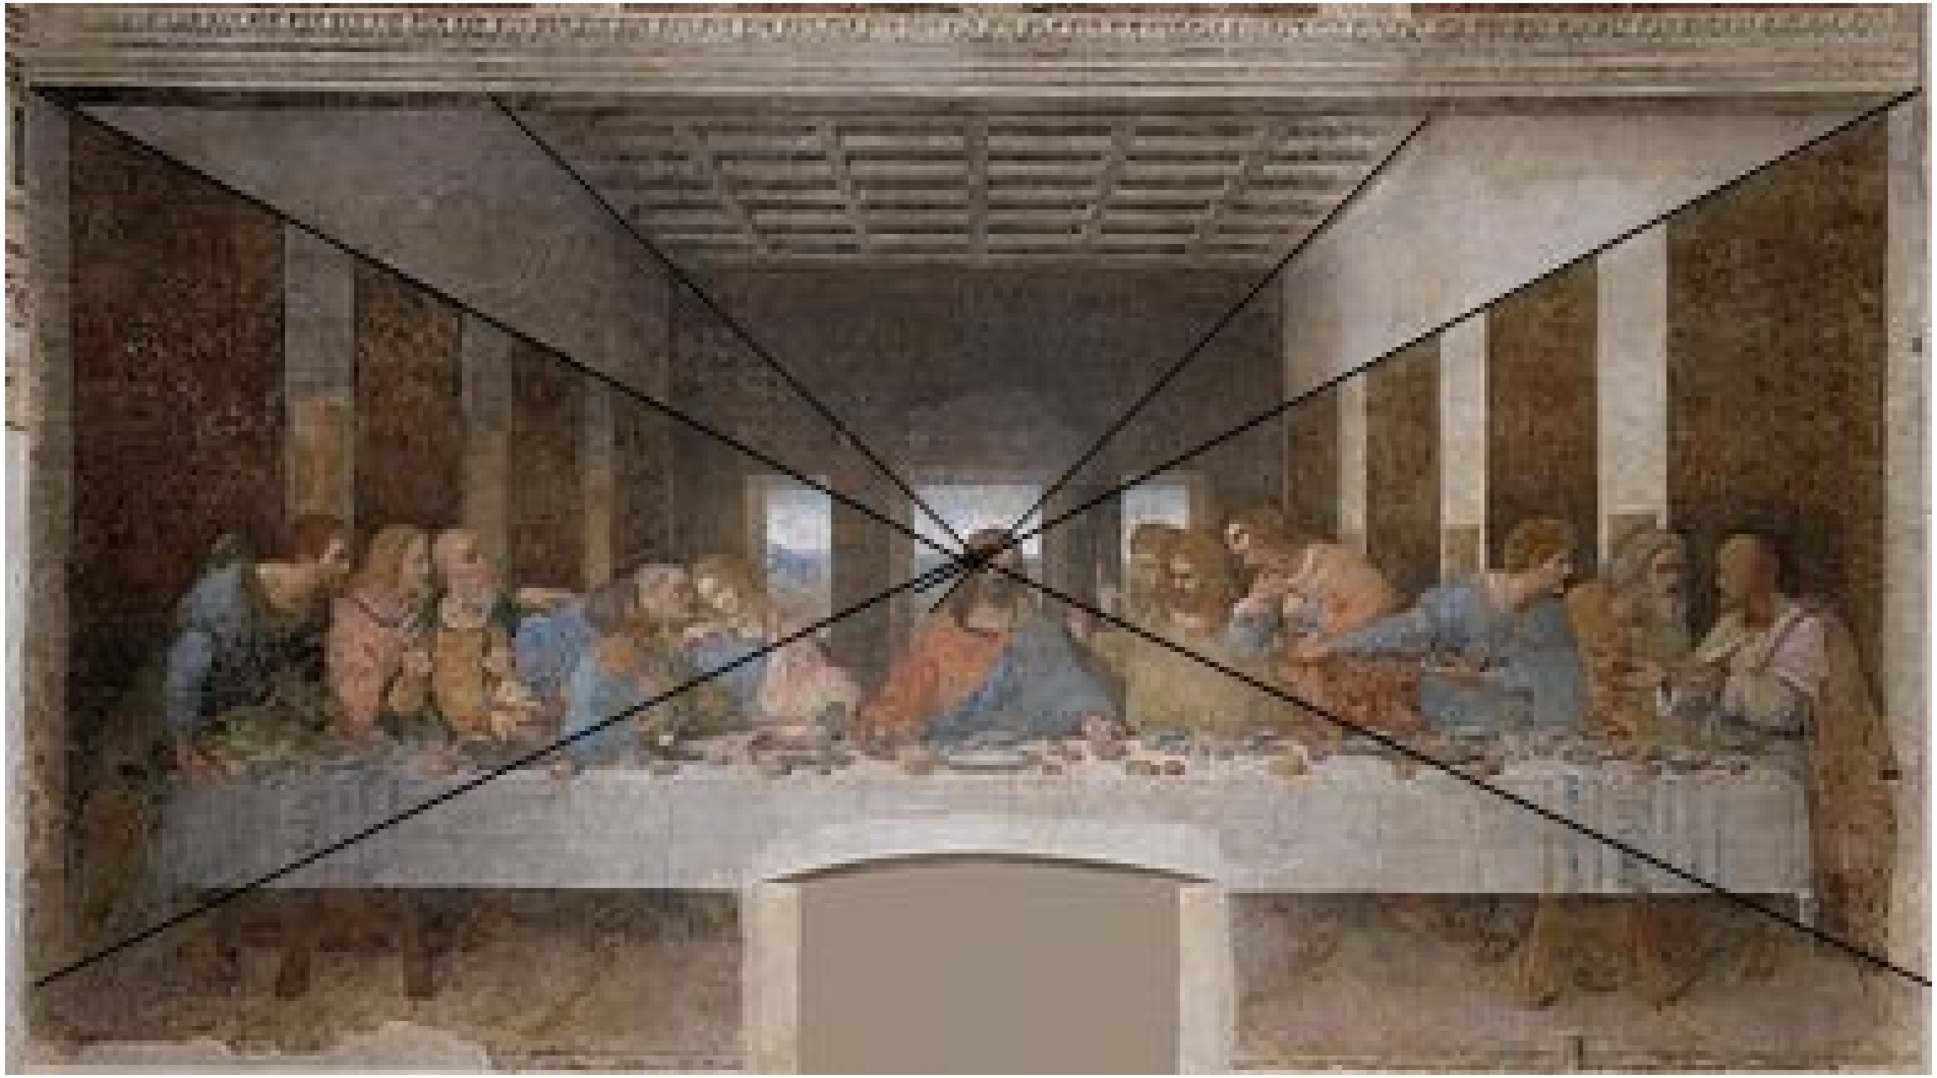

Da Vinci's "The Last Supper" c. 1495-98. http://pennpaint.blogspot.com/

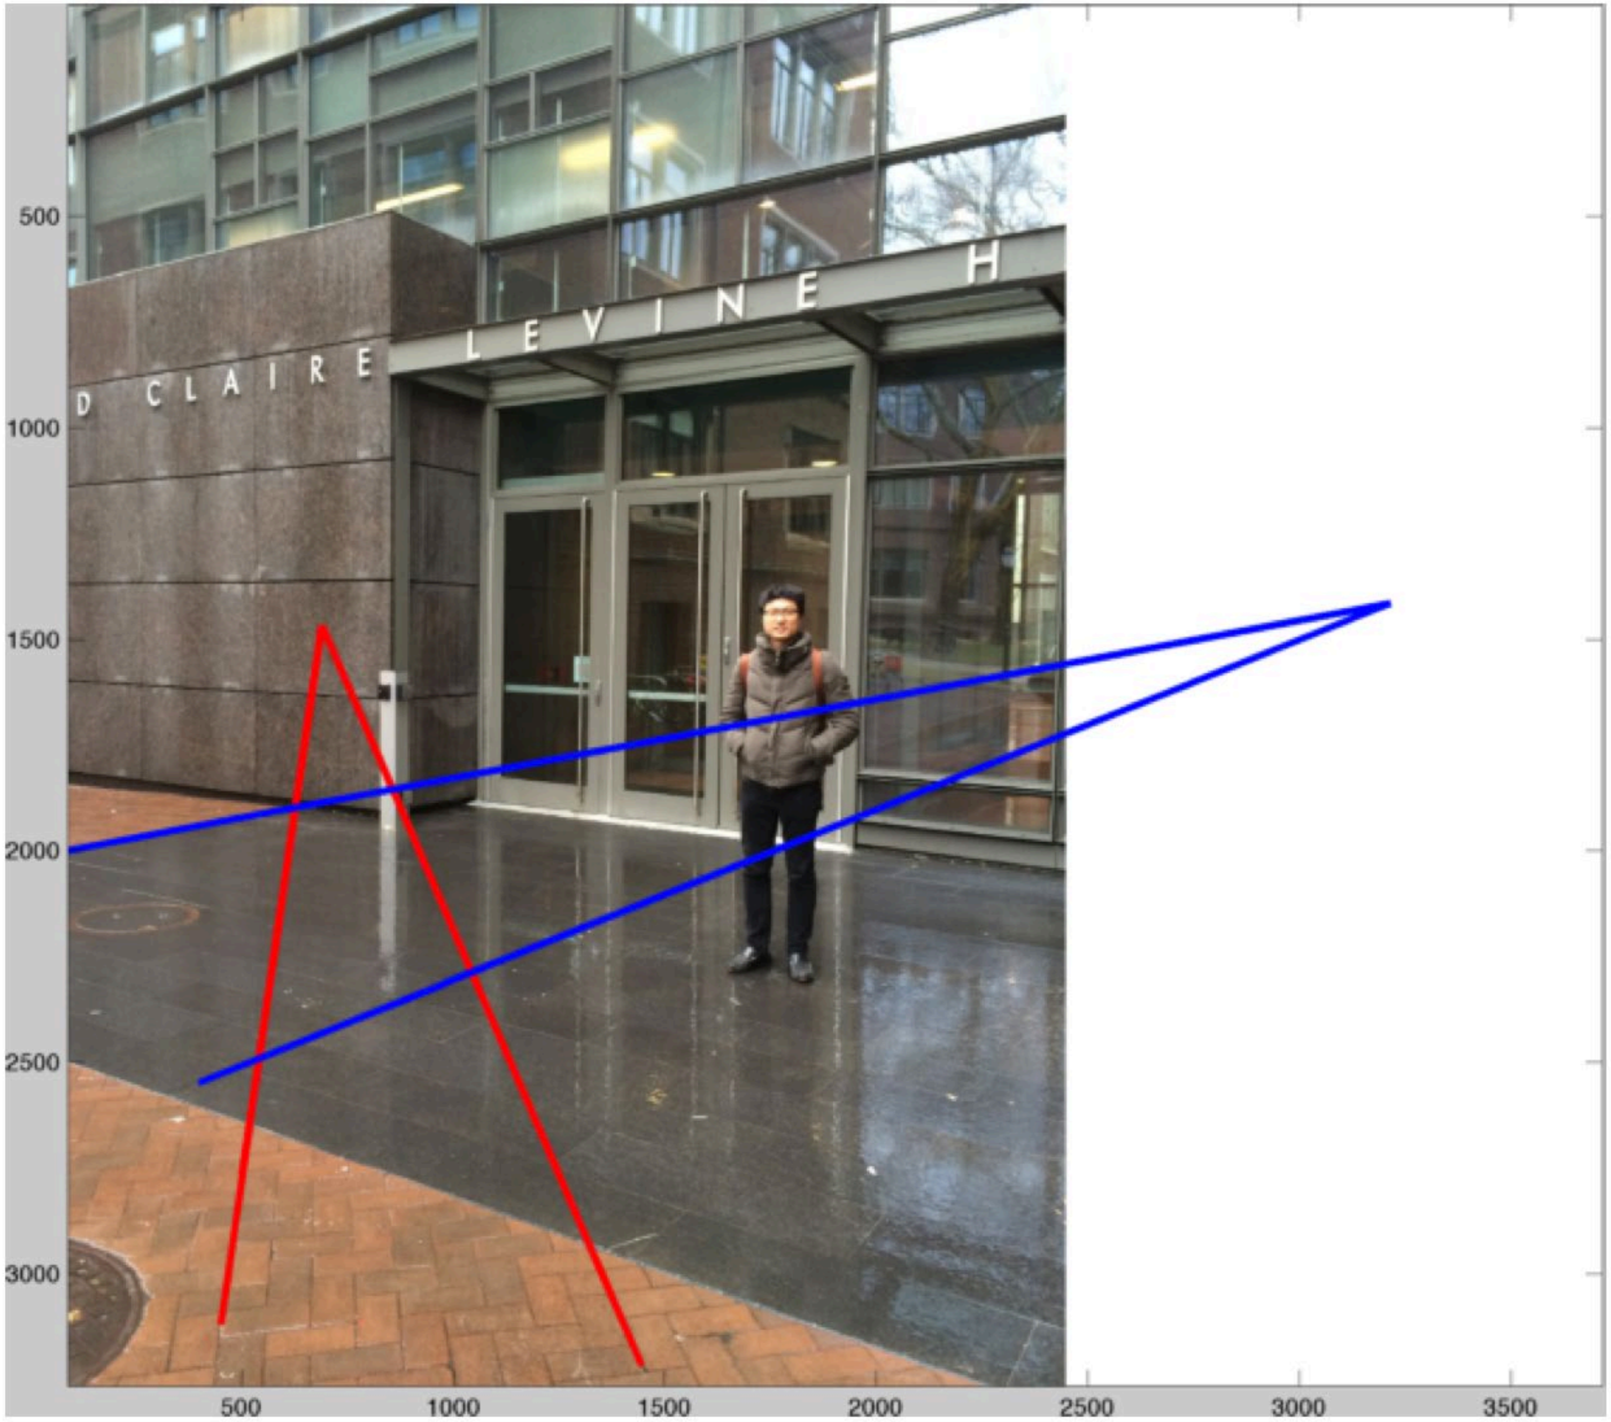
D CLAIRE
LEVINE H
500
1000
1500
2000
2500
3000
500
1000
1500
2000
2500
3000
3500

Where vanishing points come from

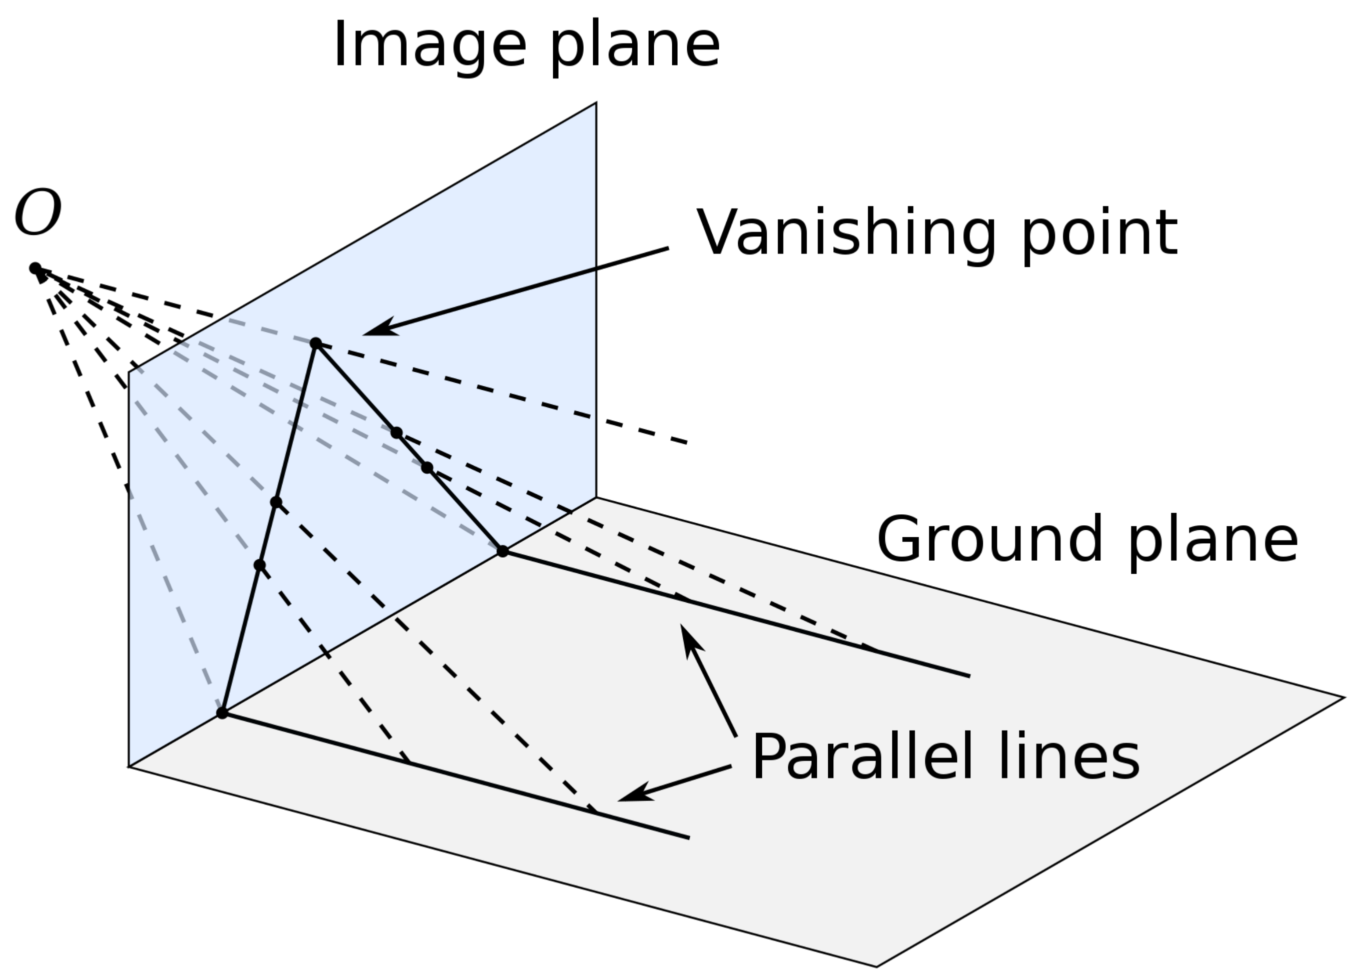

ZHCh8

Images of points “at infinity” are often finite!

In fact, they will turn out to be surprisingly important.

Need ways to deal with them conveniently and well!

Enter Projective Geometry

We will return to deal with vanishing points and much more once we know some projective geometry!

Projective Geometry

- Extension of Euclidean geometry that deals with points at infinity
  - Other than that, Projective Geometry keeps many of the same features as Euclidean geometry.

- A key property is that two lines always meet in a point
  - (sometimes a point at infinity)

"Homogeneous coordinates": Euclidean $\mathbb{R}^n \rightarrow$ Projective $\mathbb{P}^n$

- A point in n-D Euclidean space $\mathbb{R}^n$ can be injected into n-D projective space $\mathbb{P}^n$ through the "homogeneous coordinates" notation: simply add one more vector element, and set it to 1.

Homogeneous coordinates

– represent coordinates in $\mathbb{P}^2$ with a 3-vector

$$\begin{bmatrix} x \\ y \end{bmatrix} \xrightarrow{\text{homogeneous coords}} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

We will normally use w to denote this homogeneous vector element

Projective $\mathbb{P}^n \rightarrow$ Euclidean $\mathbb{R}^n$

- Just drop the 1? Not so fast!
  - Points in $\mathbb{P}^2$ do not all have $w = 1$ e.g. $(x, y, w = 2)^T$
- Key property: $(x, y, 1) = (2x, 2y, 2) = (wx, wy, w)$
  - So $(x, y, 2) \; = \; (x/2, y/2, 1)$ in $\mathbb{P}^2$ = $(x/2, y/2) \in \mathbb{R}^2$
- Geometric intuition matching to perspective projection:
  - a point in the image is a ray in projective space

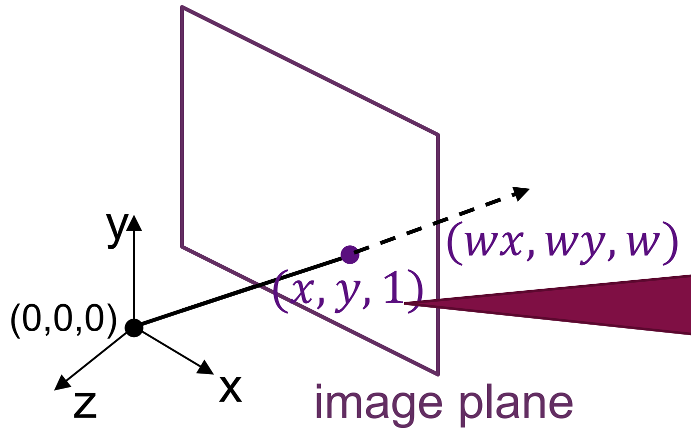

Reduce any $(wx, wy, w) \in \; \mathbb{P}^2$ to $w = 1$ to get image plane Euclidean $\mathbb{R}^2$ coordinates (x, y)

- Each *point* (x, y) on the plane is represented by a *ray through origin* (wx, wy, w)
  - – all points on the ray are equivalent: $(x \; y, 1) \approx (wx, wy, w)$

What about when $w = 0$?

- If the homogeneous coordinate is 0, $(x, y, w = 0)$ in projective coordinates corresponds to what points, what rays?
  - Points at infinity! (Think $x/0, y/0$)
  - Equivalently, rays through origin (camera center), *parallel to the projective image plane*

- This property of projective geometry allows us to handle points at infinity through simple vector operations!

Note: $\mathbb{P}^2$ does not contain $(0,0,0)$.

Projective geometry is ubiquitous in geometric vision, which models image formation as a map from 3-D projective space into 2-D projective space.

Perspective Projections are Linear in $\mathbb{P}$

Recall: $\boxed{x = f\frac{X}{Z}, y = f\frac{Y}{Z}}$

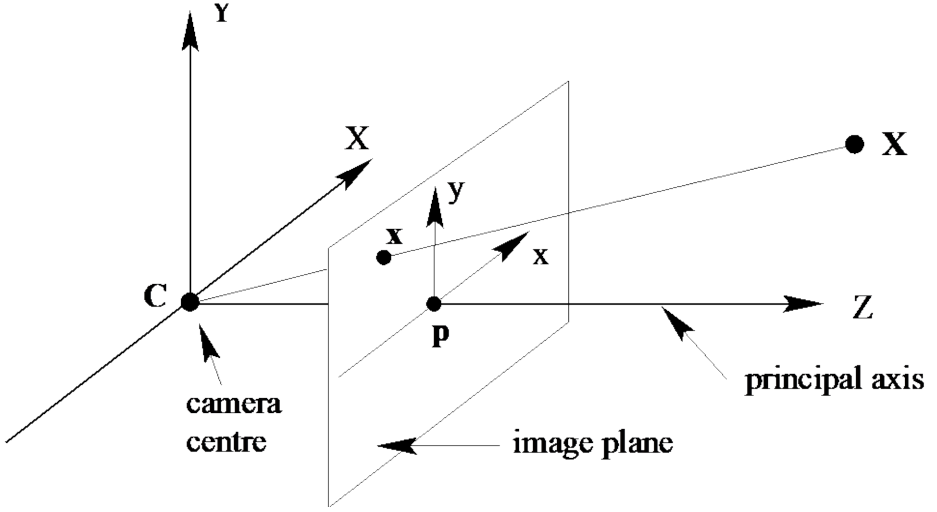

$$\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \rightarrow \begin{bmatrix} fX \\ fY \\ Z \end{bmatrix}$$

This encounter with the camera projection equation will not be our last.
More general versions soon!

The projective plane in computer vision

The projective image plane $\mathbb{P}^2$ consists of all rays through the origin $\lambda \begin{bmatrix} u \\ v \\ w \end{bmatrix}$,

which correspond either to points in the image plane $\begin{bmatrix} x = u/w \\ y = v/w \\ 1 \end{bmatrix}$,

or to points at infinity $\begin{bmatrix} u \\ v \\ 0 \end{bmatrix}$.

Projective lines

- What does a line in the image correspond to in projective space?

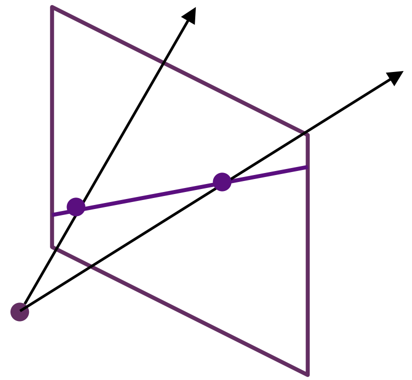

- A line is a *plane* of rays through origin
  - all rays (x, y, w) satisfying: $ax + by + cw = 0$

$$\begin{bmatrix} a & b & c \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix} = 0$$

$$\boldsymbol{l}^T \boldsymbol{x} = 0$$

Projective Lines

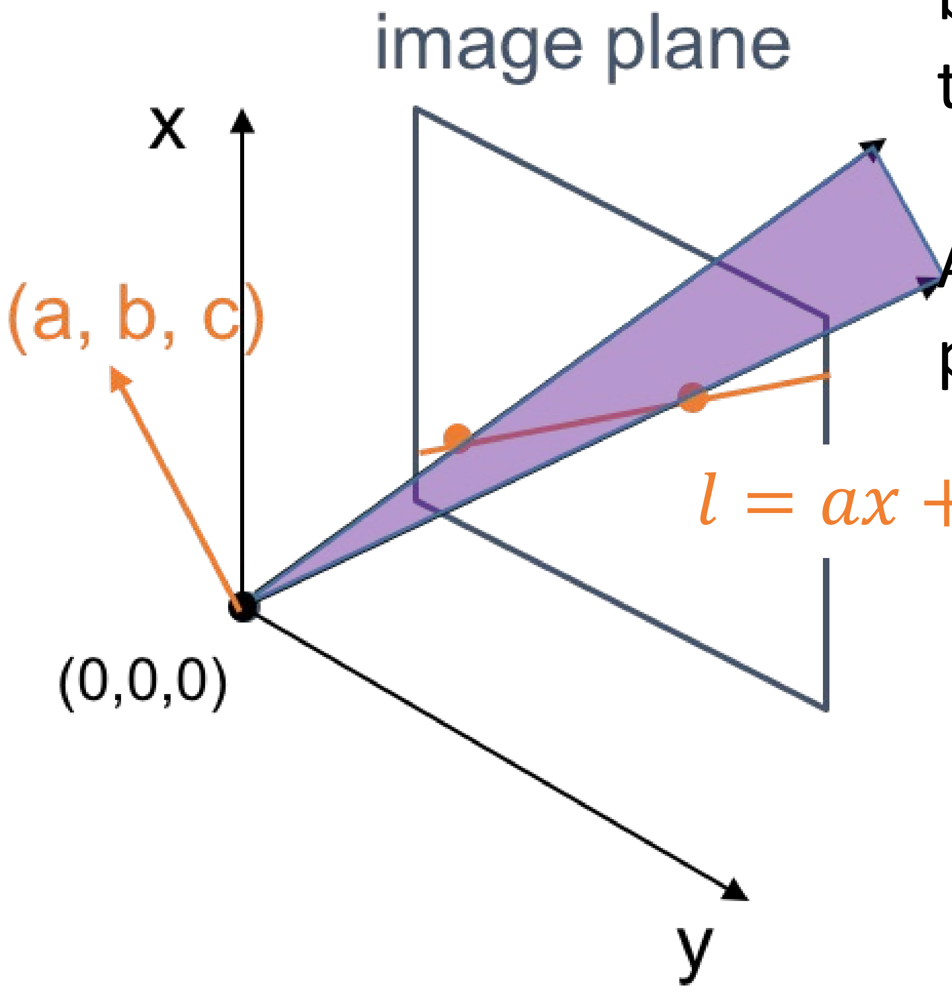

$$l = ax + by + cw = 0$$

Q: what is the relationship between $l = (a, b, c)^T$ and the plane?

A: It is perpendicular to the plane.

Line passing through two points in $\mathbb{P}^2$

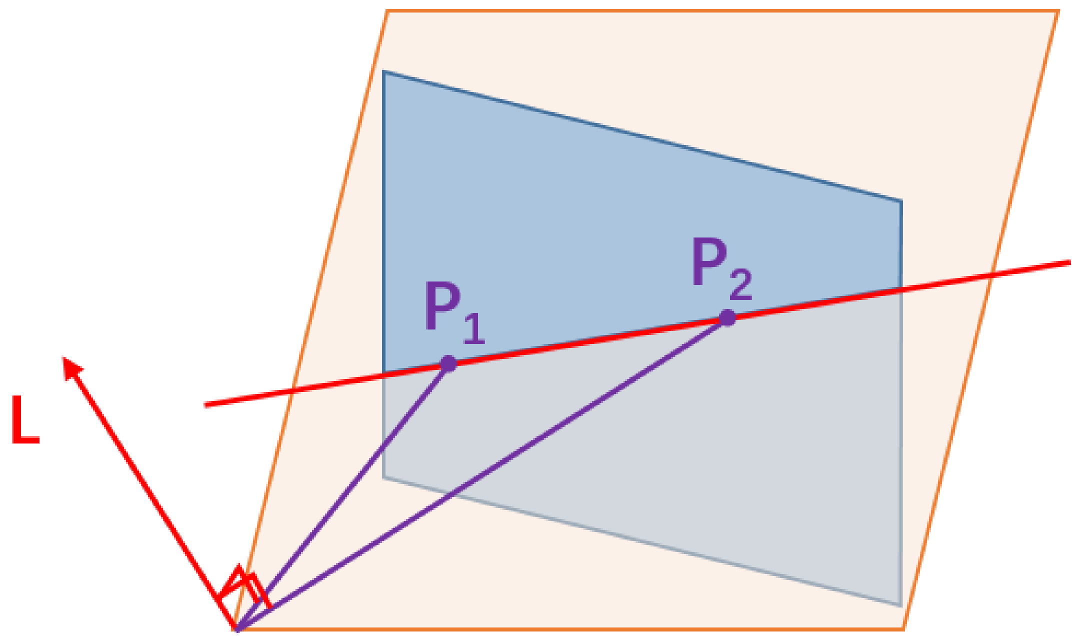

Line passing through two points in $\mathbb{P}^2$

Two points: $\boldsymbol{p_1}$ and $\boldsymbol{p_2}$

Cross product

$\boxed{\text{Define a line}}$ $\boldsymbol{l} = \boldsymbol{p_2} \times \boldsymbol{p_1}$

Claim: $\boldsymbol{l}$ is the line passing through the two points

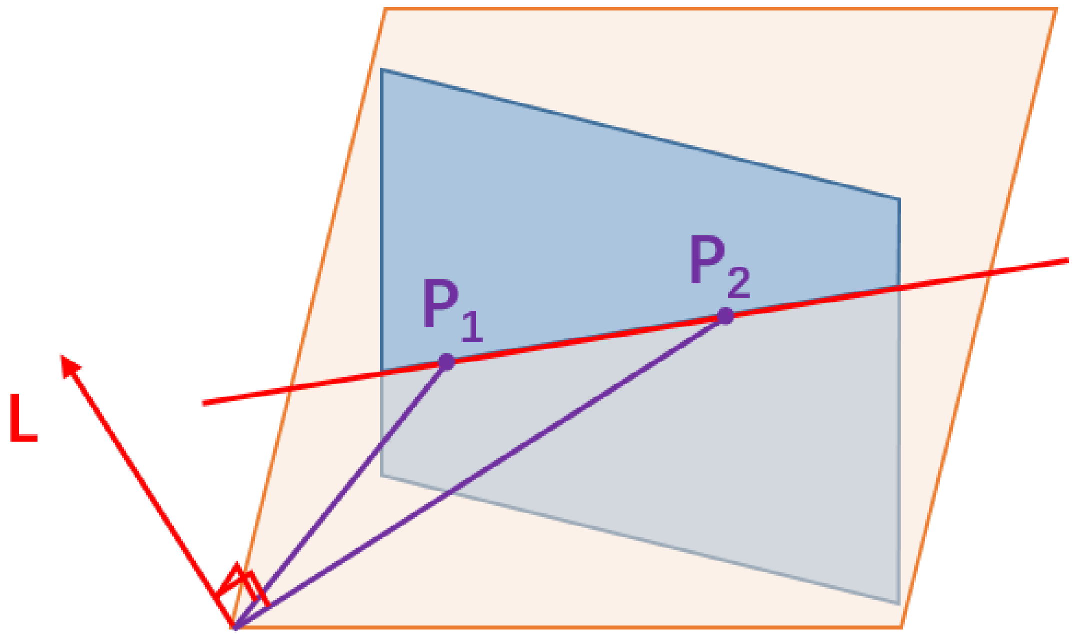

Vector Algebra Recap: Cross Products

- $\boldsymbol{a} \times \boldsymbol{b}$
  - Output is a vector with magnitude = $||\boldsymbol{a}||||\boldsymbol{b}|| \sin \boldsymbol{\theta}_{ab}$ and direction perpendicular to both vectors.
  - Right-hand thumb rule to compute direction.
  - Area of the parallelogram with $\boldsymbol{a}$ and $\boldsymbol{b}$ as sides.

- Computed as the determinant $\begin{vmatrix} \boldsymbol{i} & \boldsymbol{j} & \boldsymbol{k} \\ a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \end{vmatrix}$

- $\boldsymbol{a} \times \boldsymbol{b} = -(\boldsymbol{b} \times \boldsymbol{a})$

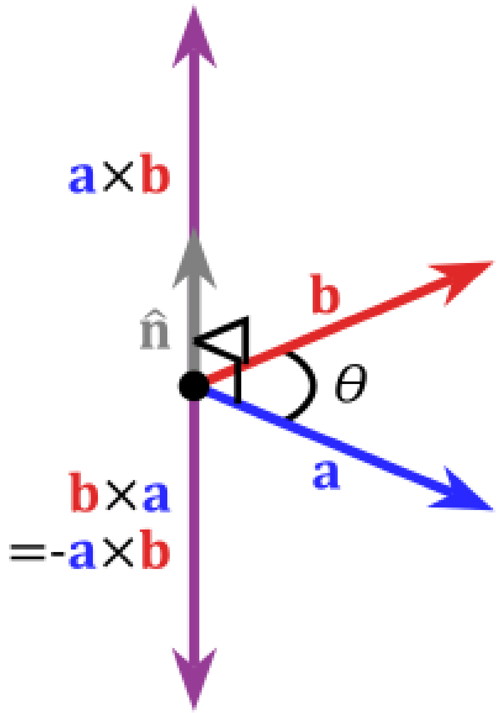

Line passing through two points in $\mathbb{P}^2$

Two points: p_1 and p_2

$\boxed{\text{Define a line}}$ $\quad l = p_1 \times p_2$

Claim: l is the line passing through the two points

Proof:

The scalar triple product of the vectors **a**, **b**, and **c**: $\vec{a} \cdot (\vec{b} \times \vec{c}) = (\vec{a} \times \vec{b}) \cdot \vec{c}$

$$p_1 . l = p_1 . (p_1 \times p_2) = 0 \text{ (property of cross-product)}$$

$$p_2 . l = p_2 . (p_1 \times p_2) = 0$$

The other way: Intersection point of two lines in $\mathbb{P}^2$

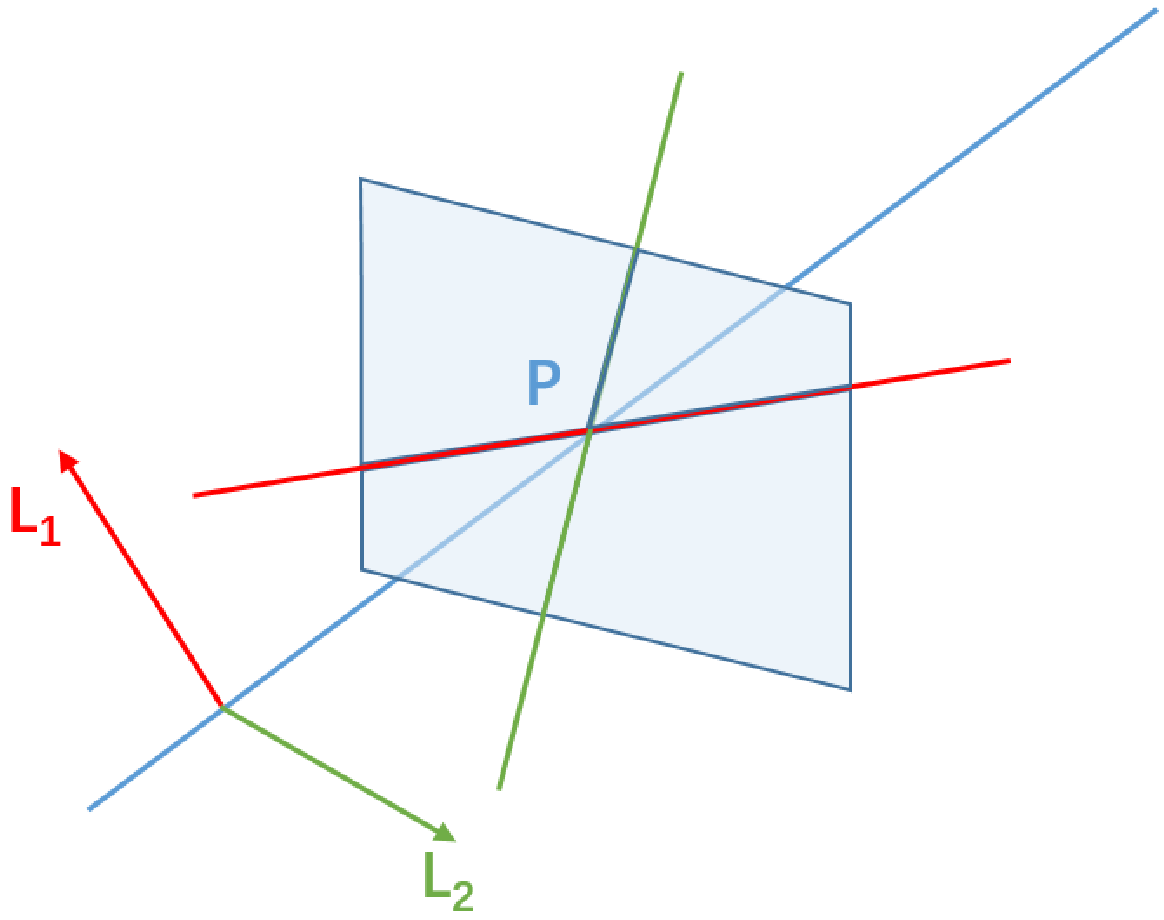

Intersection point of two lines in $\mathbb{P}^2$

Q: Why? Hint: solve for the intersection by finding points that satisfy both line equations.

Given two lines: l_1 and l_2

Define a point $p = l_1 \times l_2$

p is the intersection of the two lines

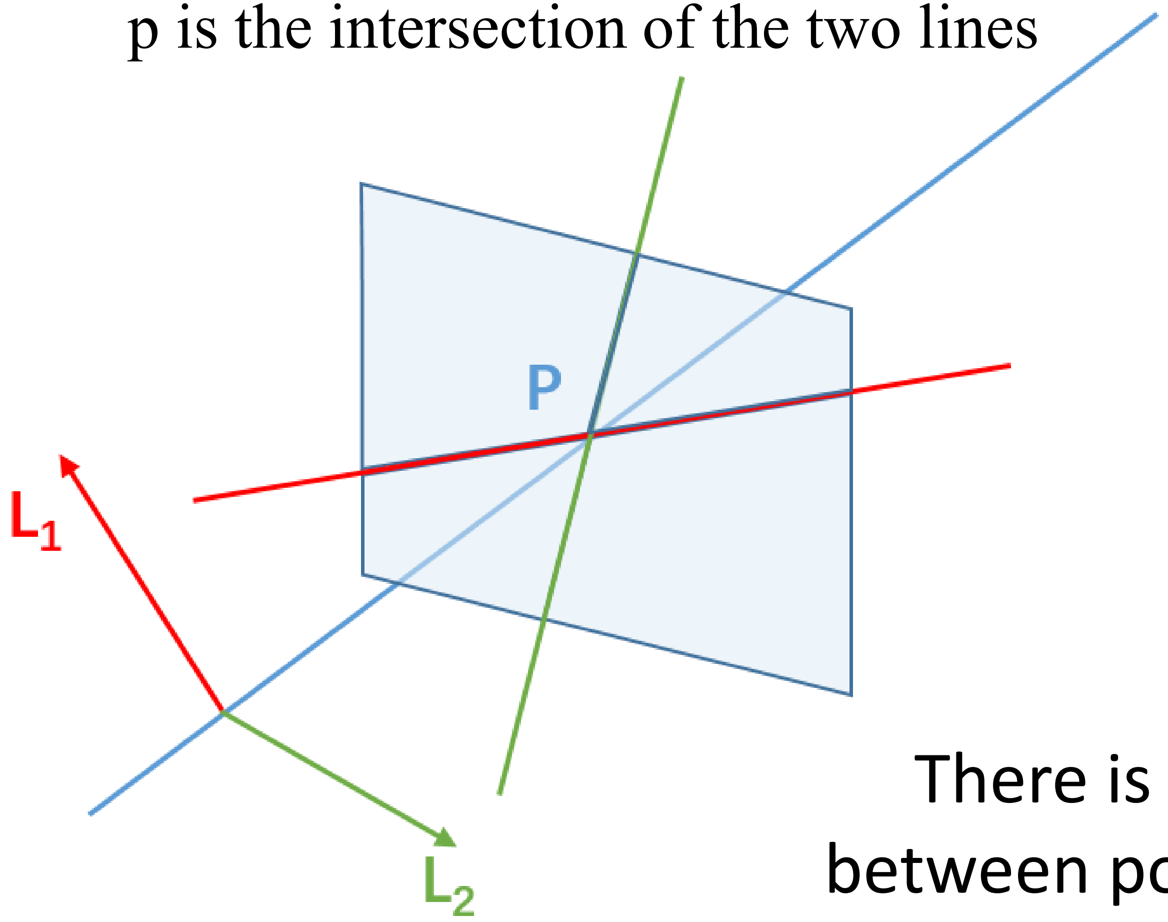

There is a close symmetry between points and lines in $\mathbb{P}^2$!

Exercise

- Where do the lines $x = 1$ and $x = 2$ on image plane intersect?

Intersection point of two lines in $\mathbb{P}^2$

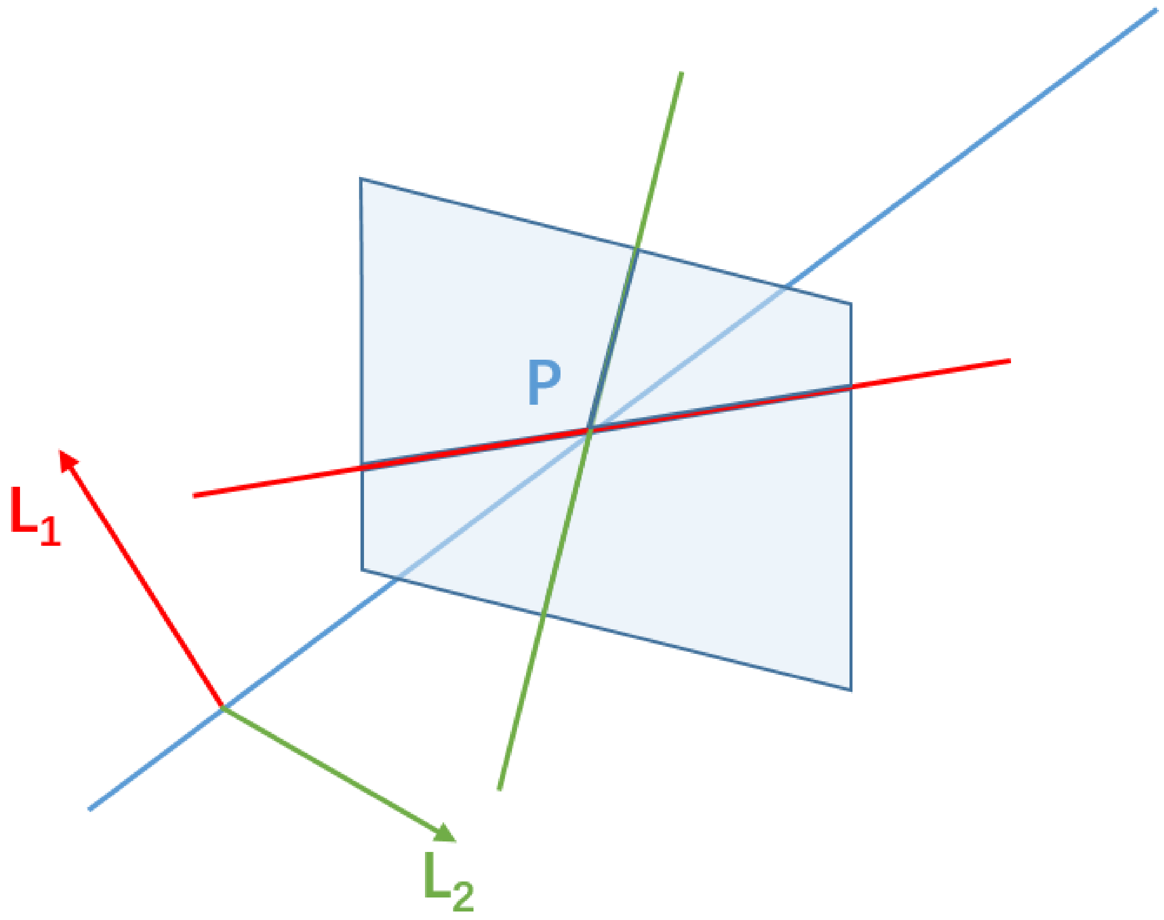

When does P have the form (x,y,0)?