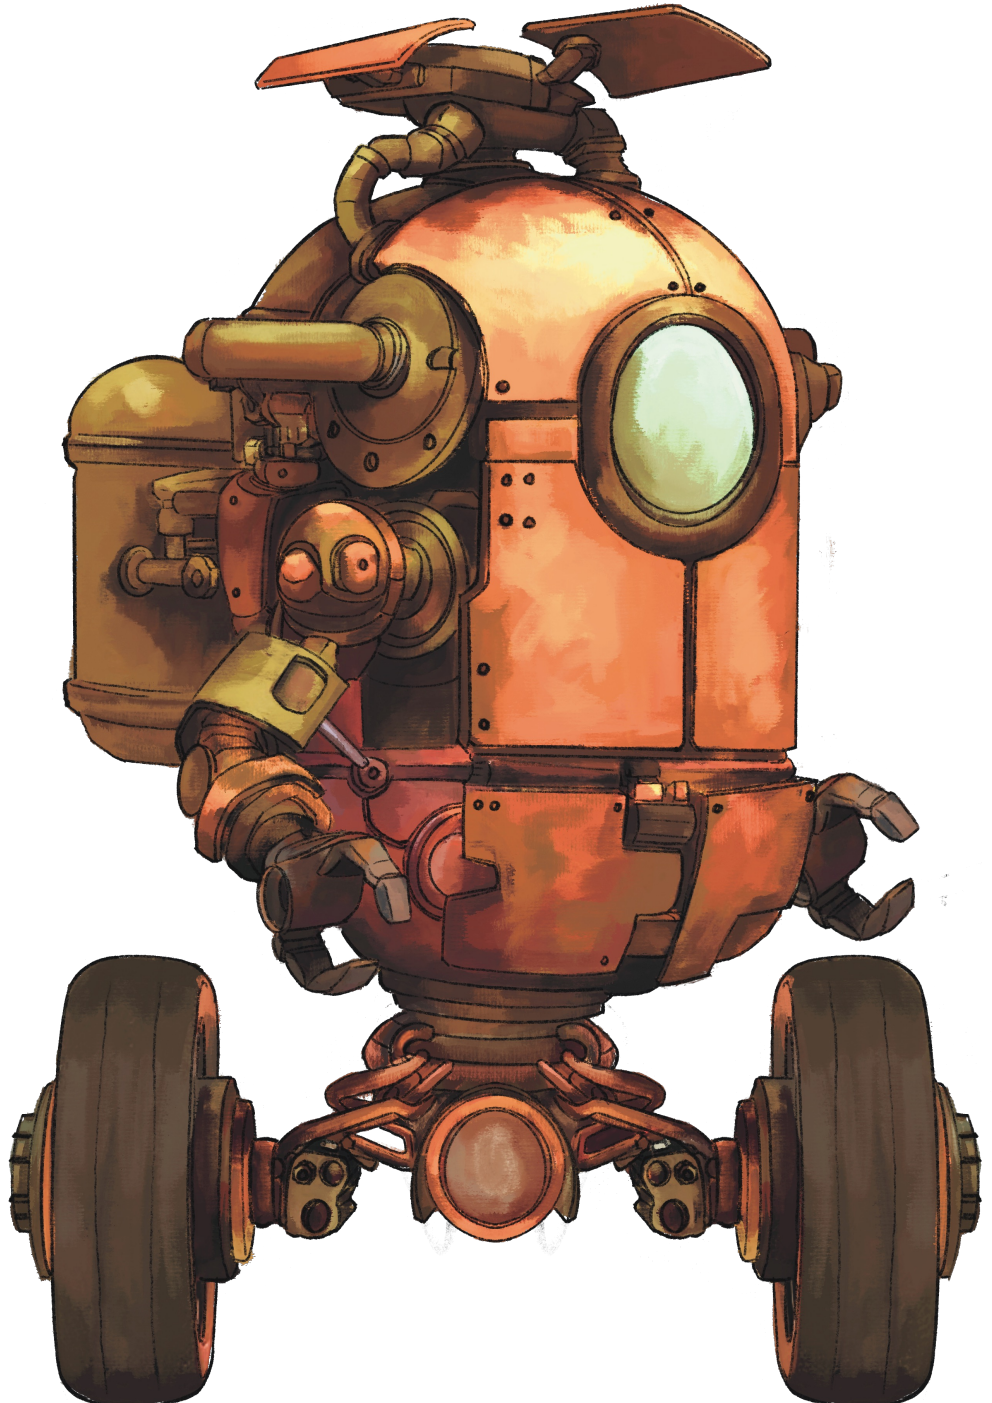
CS 3630, Fall 2026

Lecture 18:

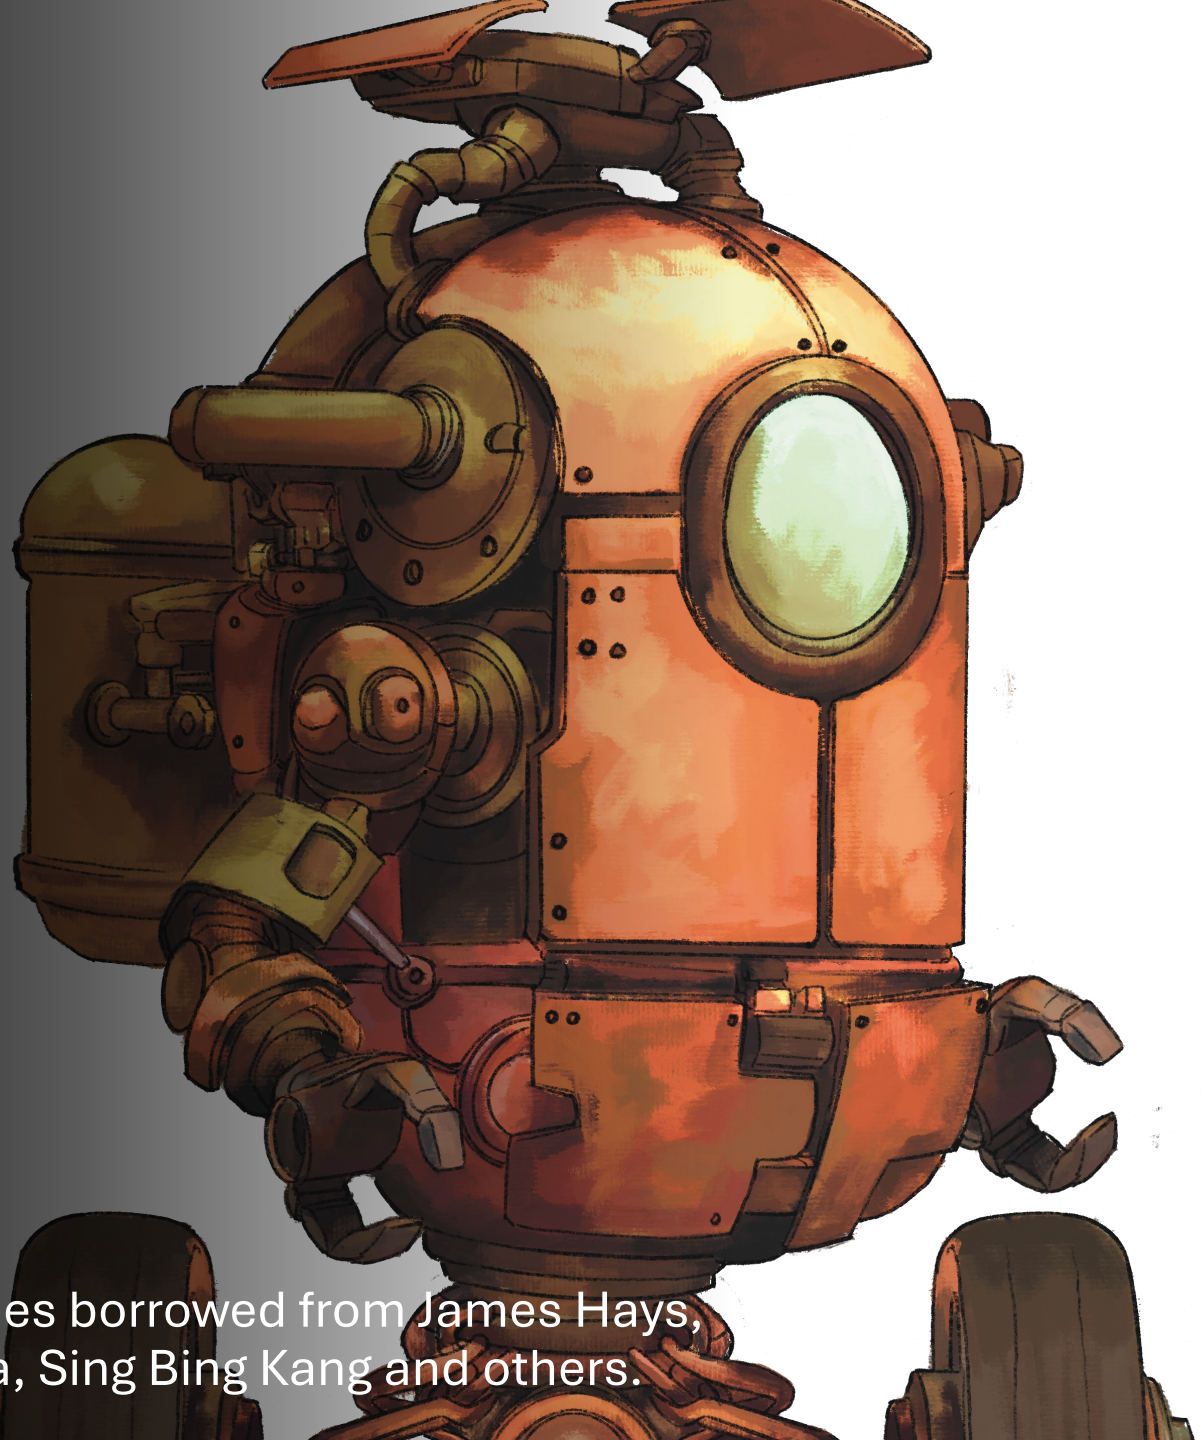
Computer Vision:

Imaging Geometry



Outline

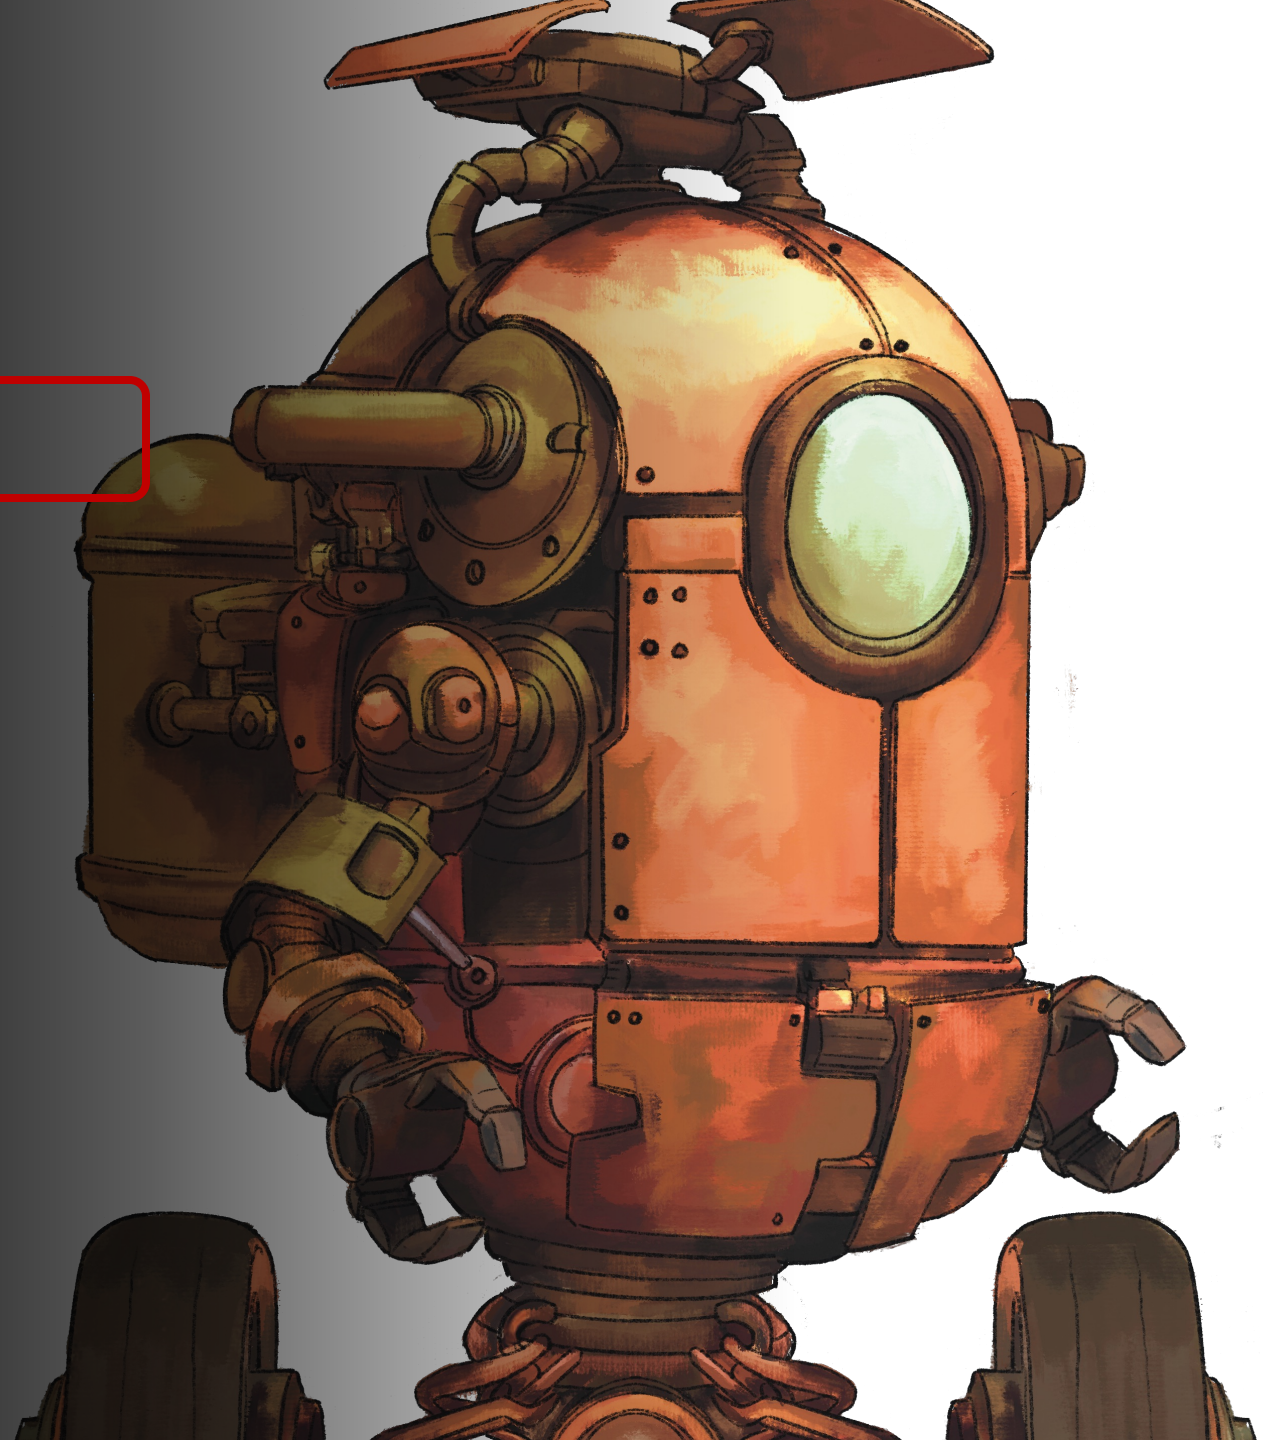
1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



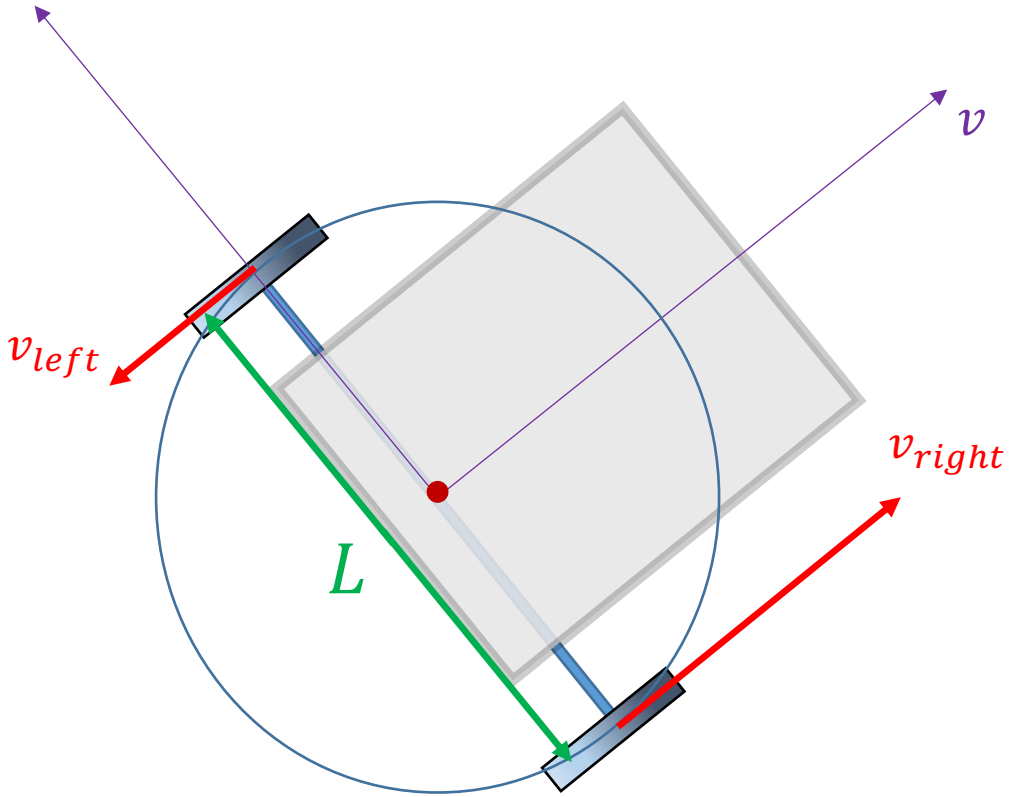
Many slides borrowed from James Hays, Irfan Essa, Sing Bing Kang and others.

Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



Forward Kinematics Jacobians



We can express these equations relative to the body-attached frame or the world frame:

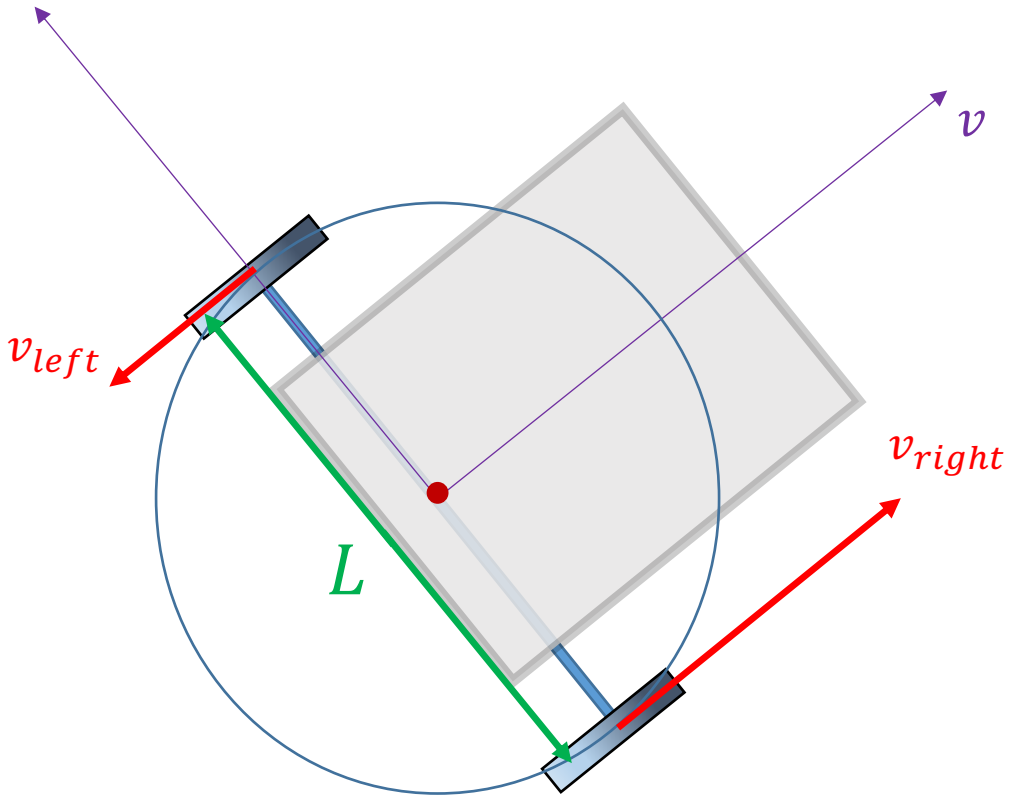
$$v^{body} = \begin{bmatrix} v_x^{body} \\ v_y^{body} \\ \omega^{body} \end{bmatrix} = \frac{r}{2} \begin{bmatrix} 1 & 1 \\ 0 & 0 \\ 2/L & -2/L \end{bmatrix} \begin{bmatrix} \dot{\phi}_R \\ \dot{\phi}_L \end{bmatrix}$$

and

These are 3x2 Jacobians

$$v^{world} = \begin{bmatrix} v_x^{world} \\ v_y^{world} \\ \omega^{world} \end{bmatrix} = \frac{r}{2} \begin{bmatrix} \cos\theta & \cos\theta \\ \sin\theta & \sin\theta \\ 2/L & -2/L \end{bmatrix} \begin{bmatrix} \dot{\phi}_R \\ \dot{\phi}_L \end{bmatrix}$$

“Inverse Kinematics” Jacobians



Forward:

$$v^{body} = \begin{bmatrix} v_x^{body} \\ v_y^{body} \\ \omega^{body} \end{bmatrix} = \frac{r}{2}\begin{bmatrix} 1 & 1 \\ 0 & 0 \\ 2/L & -2/L \end{bmatrix}\begin{bmatrix} \dot{\phi}_R \\ \dot{\phi}_L \end{bmatrix}$$

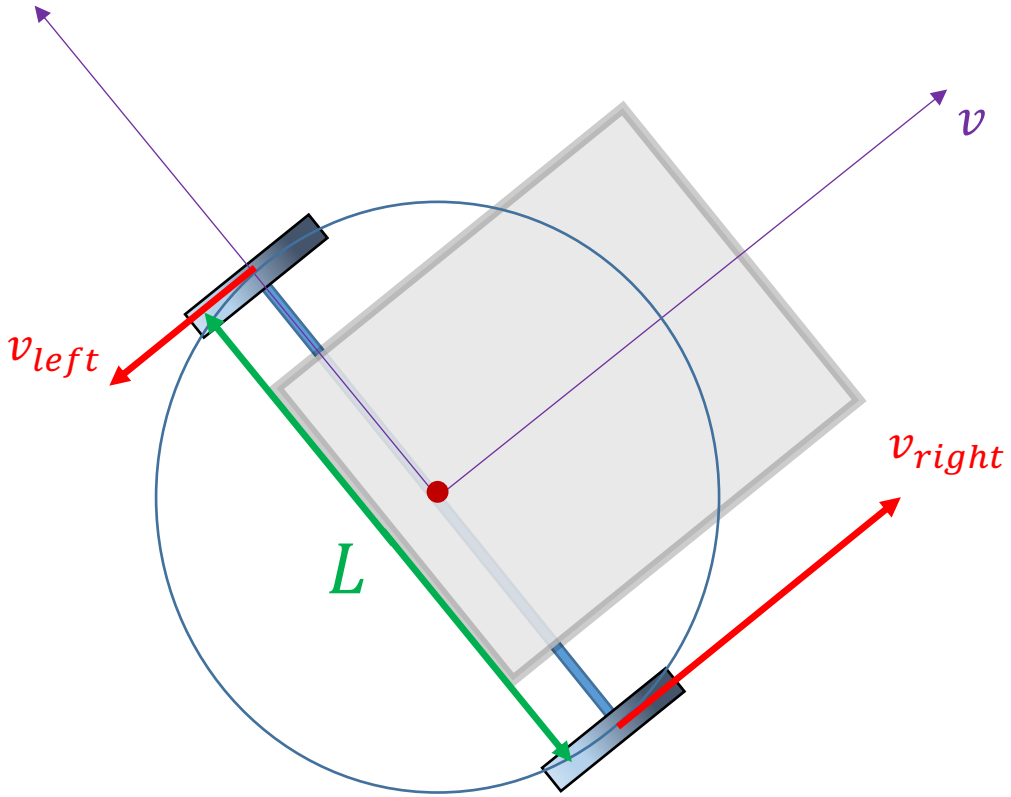
Inverse: $\dot{\phi}_R = \frac{L}{2}\frac{\omega}{r} + \frac{v_x}{r}$ and $\dot{\phi}_L = -\frac{L}{2}\frac{\omega}{r} + \frac{v_x}{r}$

$$\begin{bmatrix} \dot{\phi}_R \\ \dot{\phi}_L \end{bmatrix} = \frac{1}{r}\begin{bmatrix} 1 & L/2 \\ 1 & -L/2 \end{bmatrix}\begin{bmatrix} v_x^{body} \\ \omega^{body} \end{bmatrix}$$

Inverse is 2x2: no side-slip allowed!

Both Jacobians are *constant* in body frame !

“Inverse Kinematics” Jacobians in world frame



Forward:

Nonlinear in θ

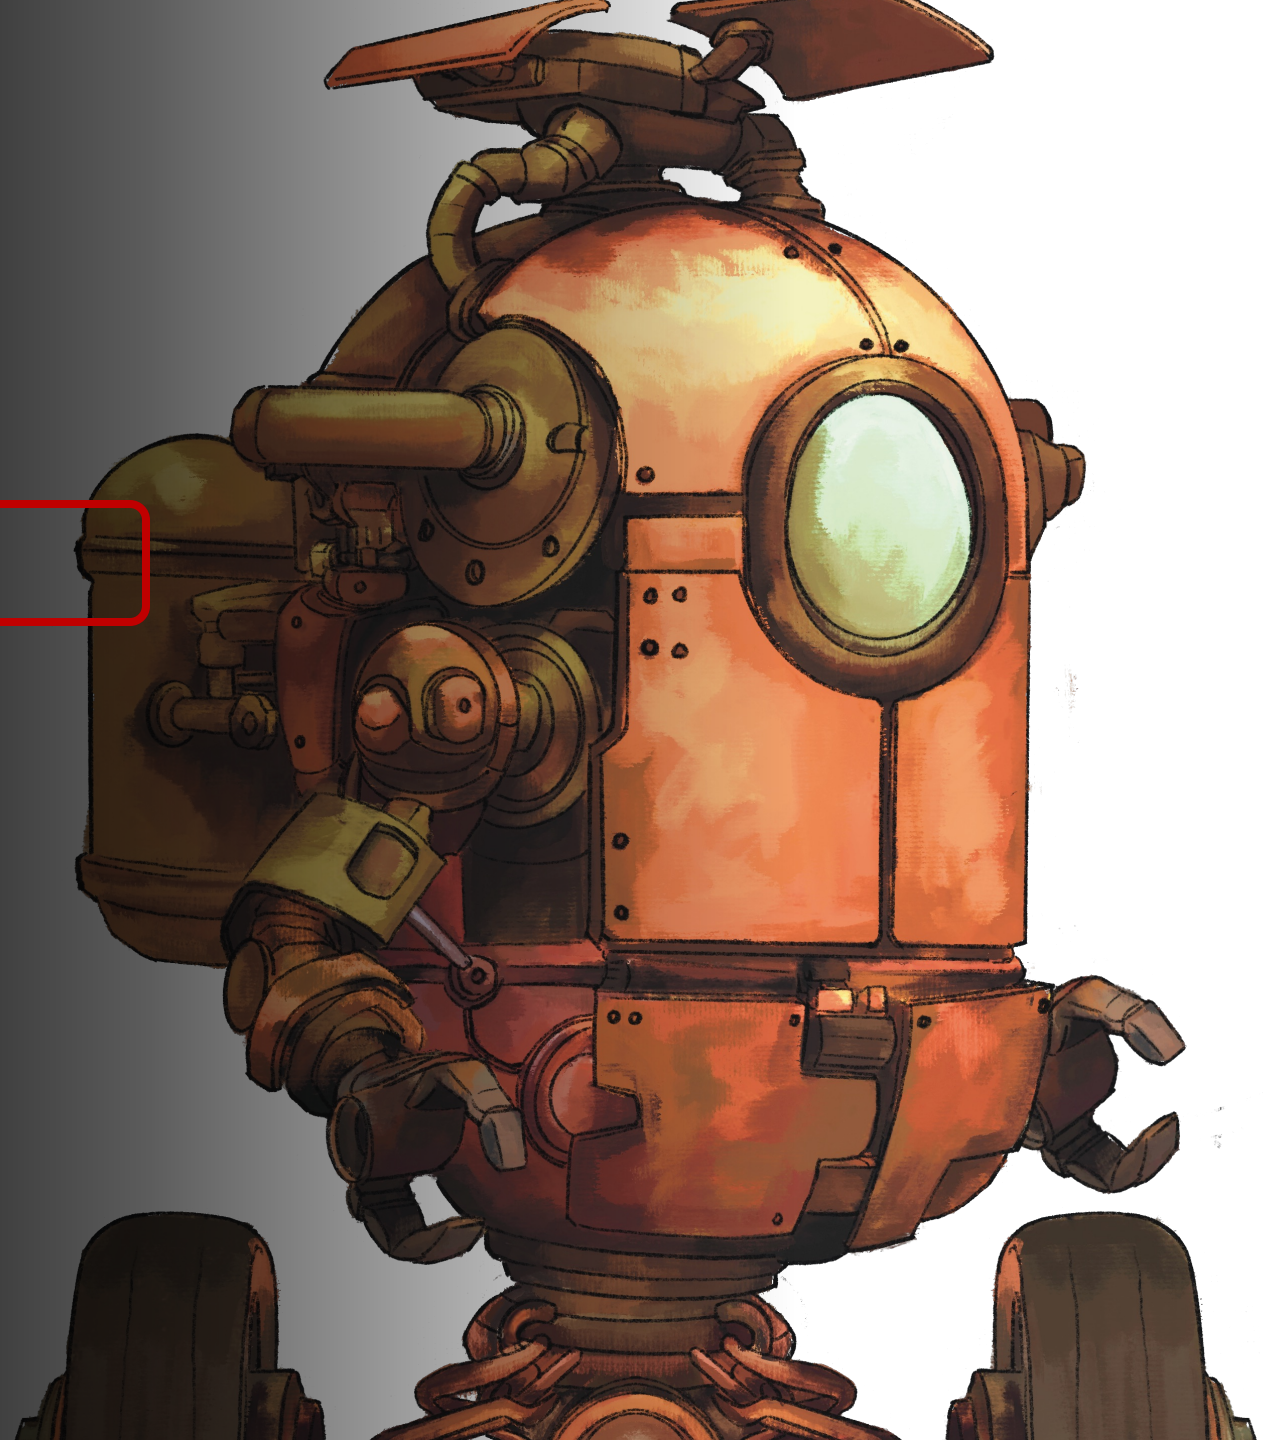
$$v^{world} = \begin{bmatrix} v_x^{world} \\ v_y^{world} \\ \omega^{world} \end{bmatrix} = \frac{r}{2} \begin{bmatrix} \cos\theta & \cos\theta \\ \sin\theta & \sin\theta \\ 2/L & -2/L \end{bmatrix} \begin{bmatrix} \dot{\phi}_R \\ \dot{\phi}_L \end{bmatrix}$$

Inverse:

Does not exist!
Not all world velocities
can be realized using a DDR!

Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm

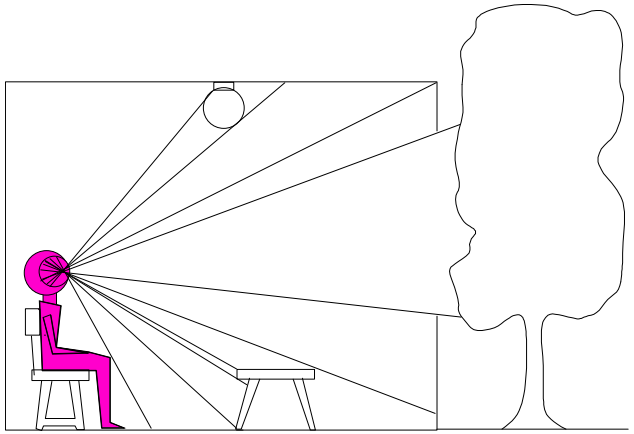


Motivation

- We need to model the image formation process
- The camera can act as an (angular) measurement device
- Need a mathematical model for a simple camera
- Two cameras are better than one: metric measurements

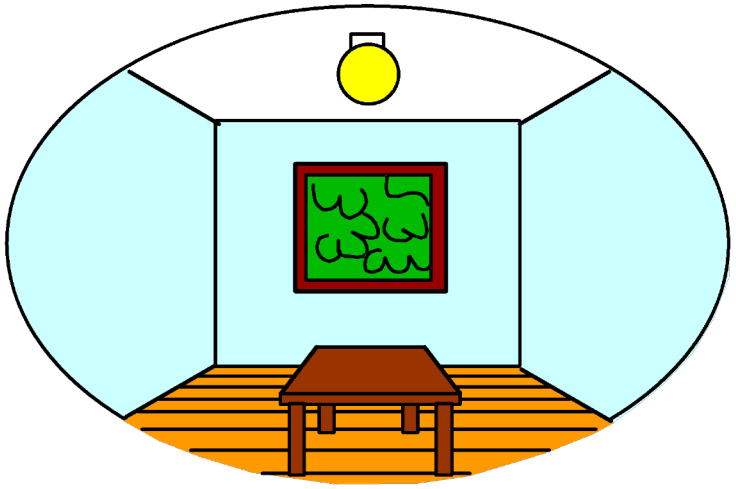
Perspective Cameras

3D world



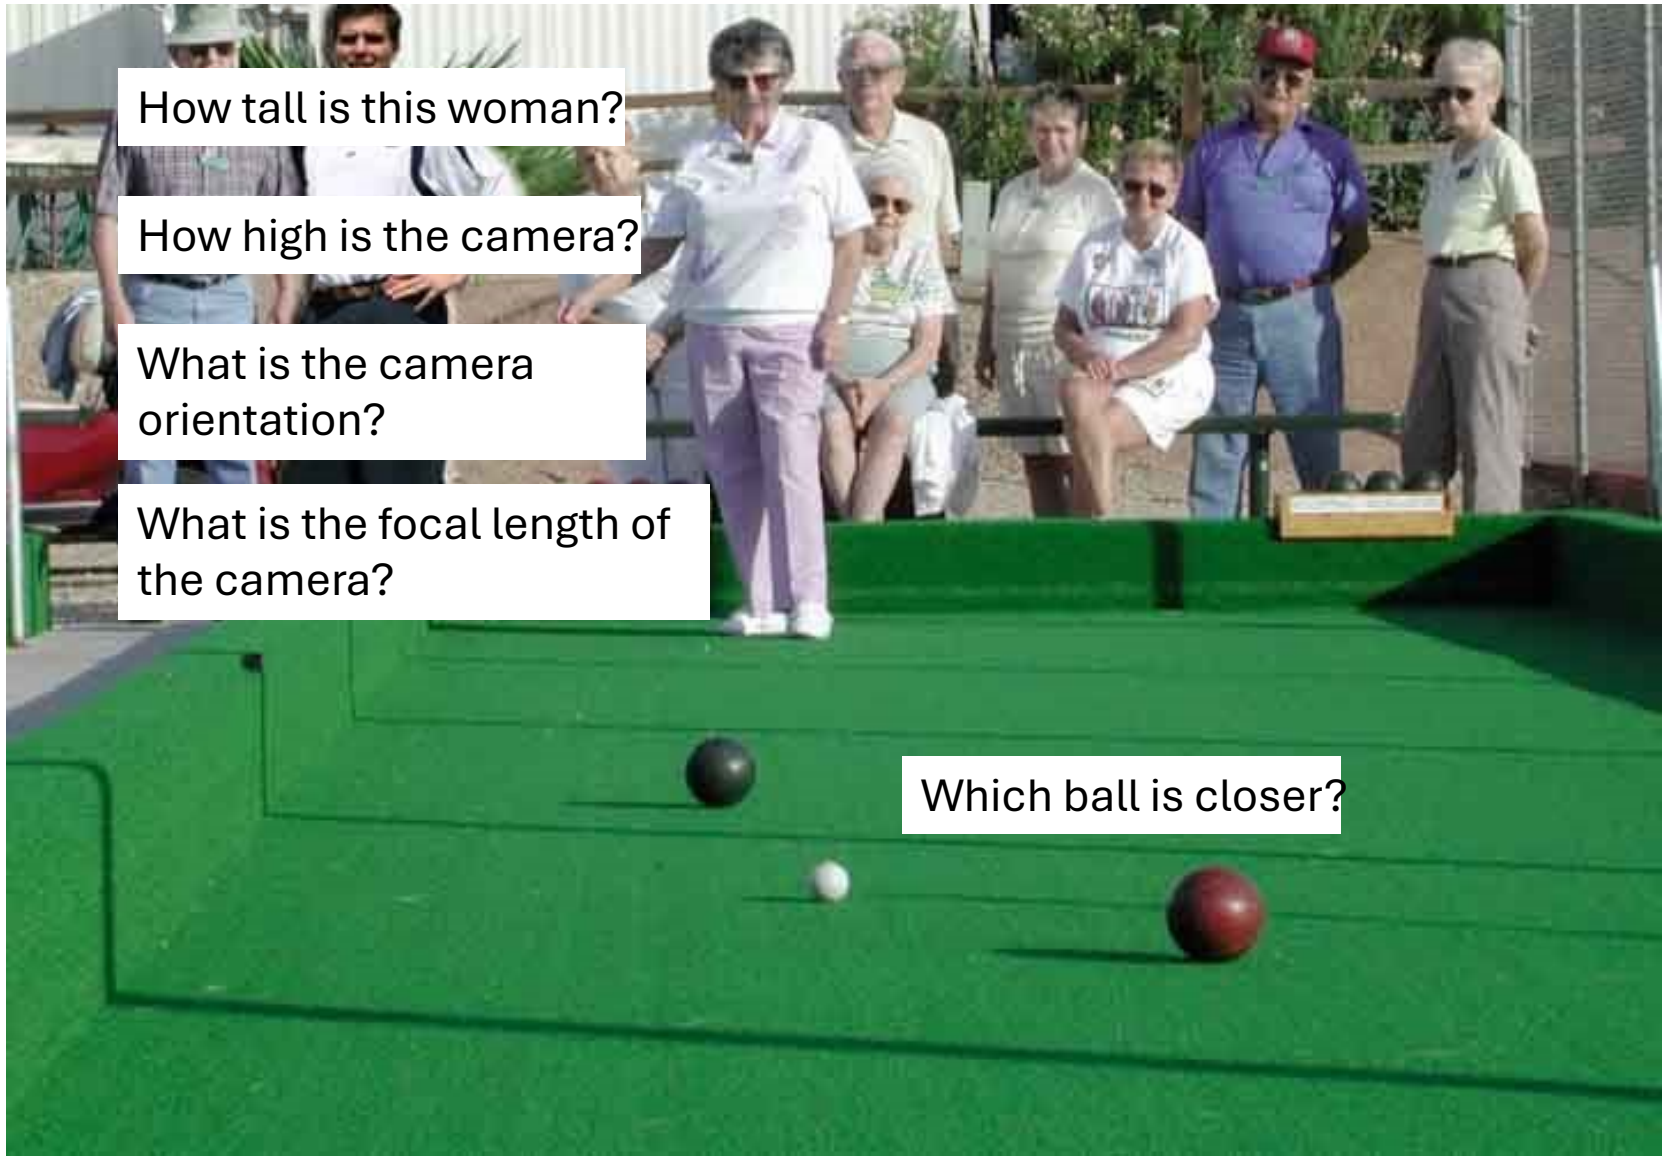
Point of observation

2D image



- Computer Vision: Images to Models
- To do this, we first need to understand the image formation process.
- We concentrate here on *geometry* (not photometry)

Camera and World Geometry



Slide source: Seitz

Projection can be tricky...



Slide source: Seitz

Projection can be tricky...

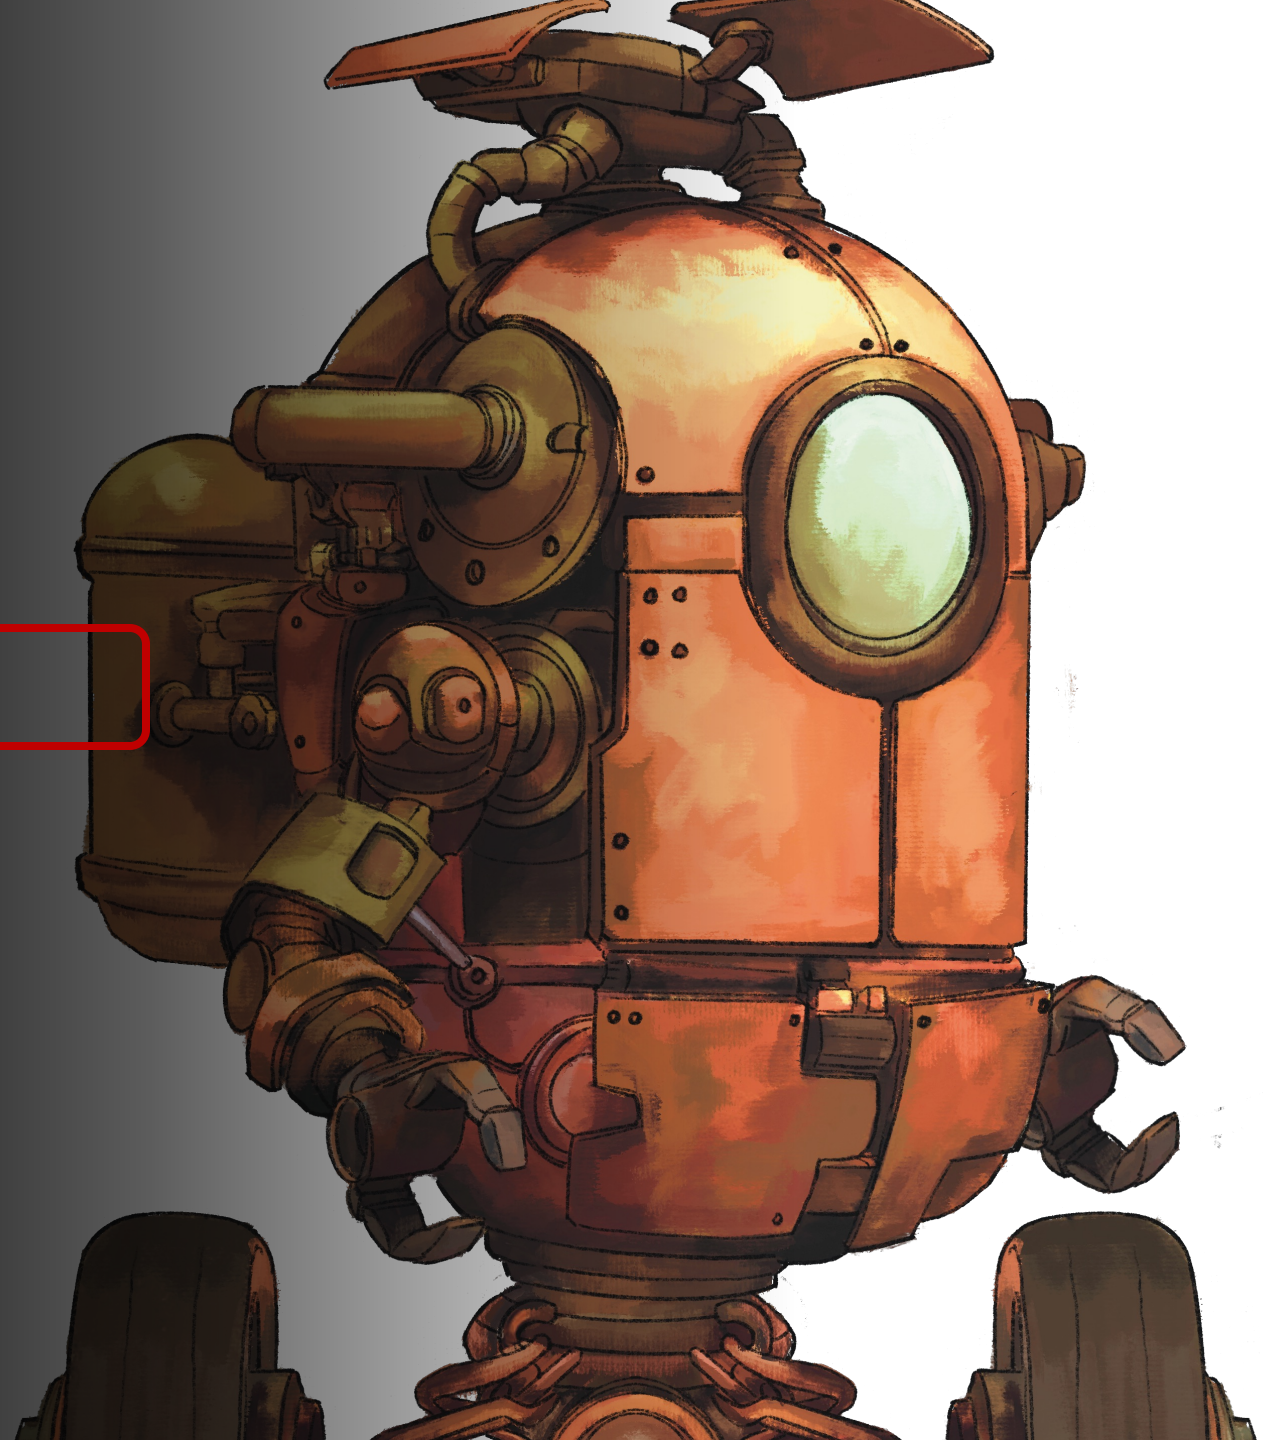




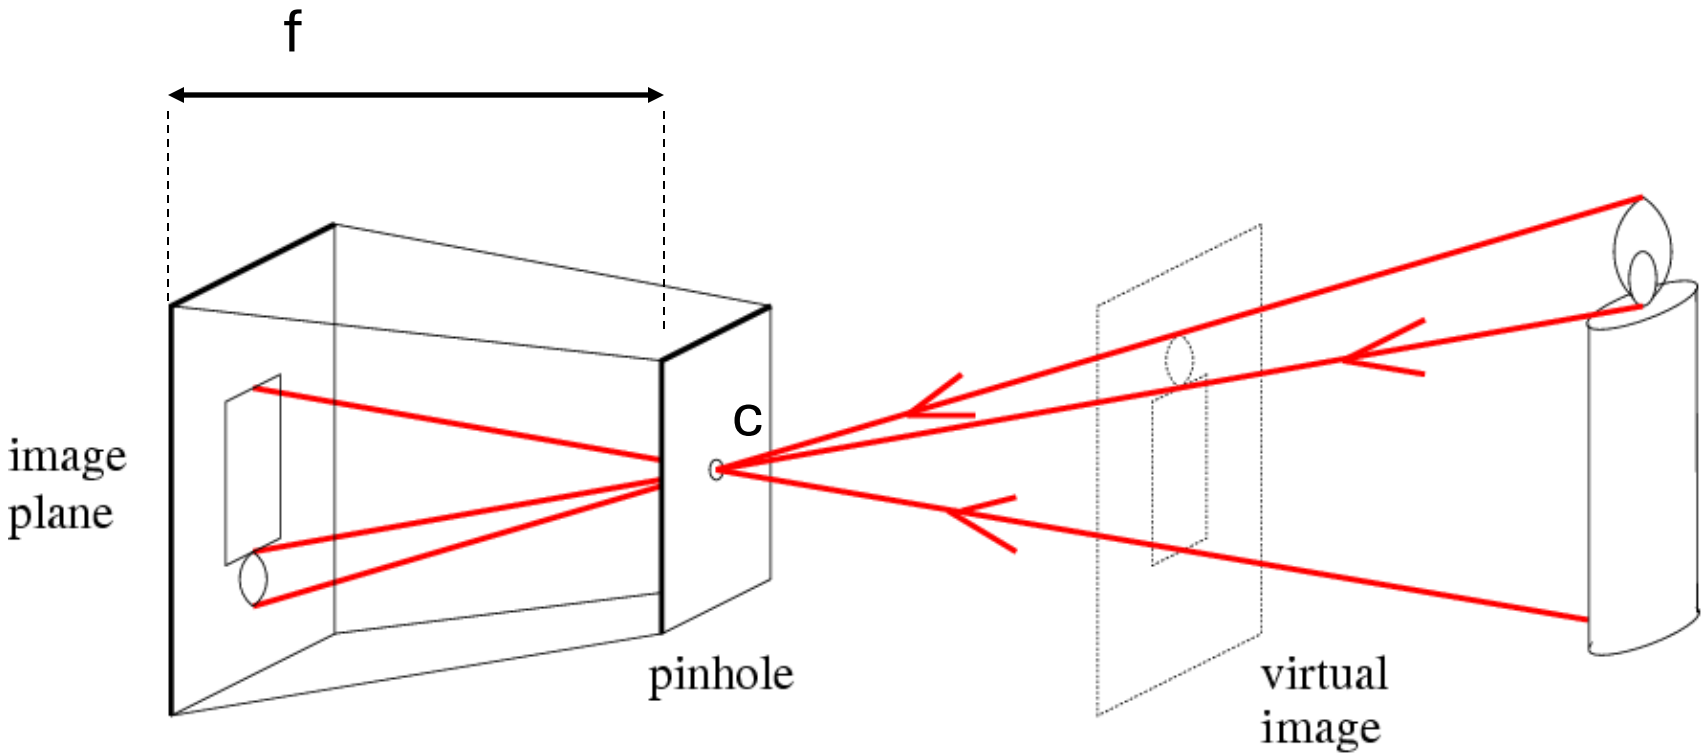


Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



Pinhole camera model



f = focal length

c = center of the camera

Figure from Forsyth

Camera obscura: the pre-camera

- Known during classical period in China and Greece (e.g. Mo-Ti, China, 470BC to 390BC)

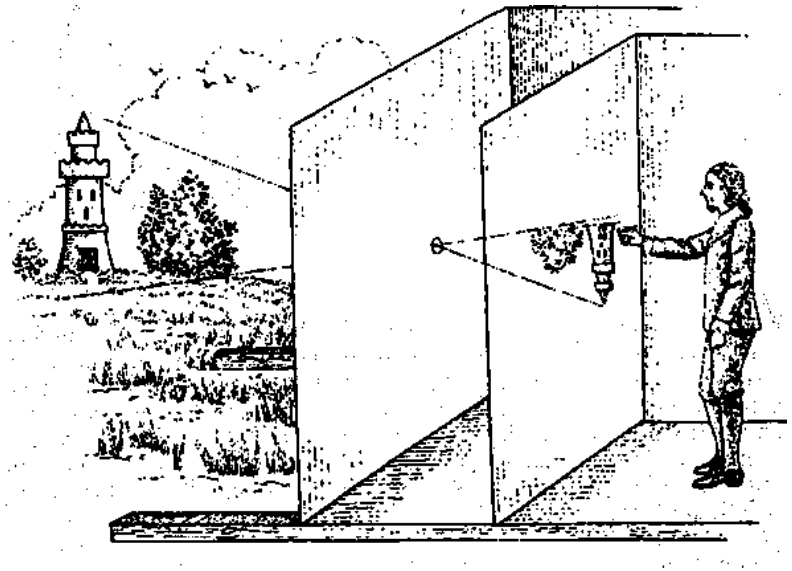


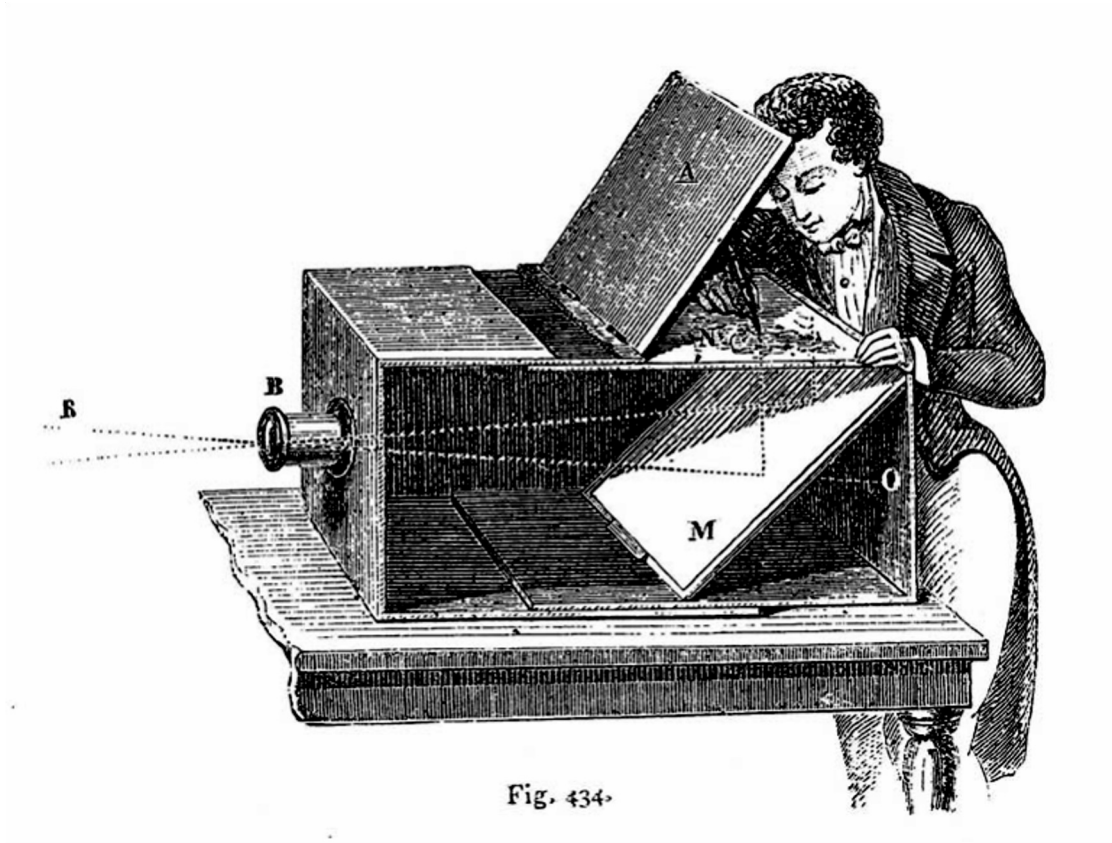
Illustration of Camera Obscura



Freestanding camera obscura at UNC Chapel Hill

Photo by Seth Ilys

Camera Obscura used for Tracing



Lens Based Camera Obscura, 1568

First Photograph

Oldest surviving photograph

- Took 8 hours on pewter plate



Joseph Niepce, 1826

Photograph of the first photograph



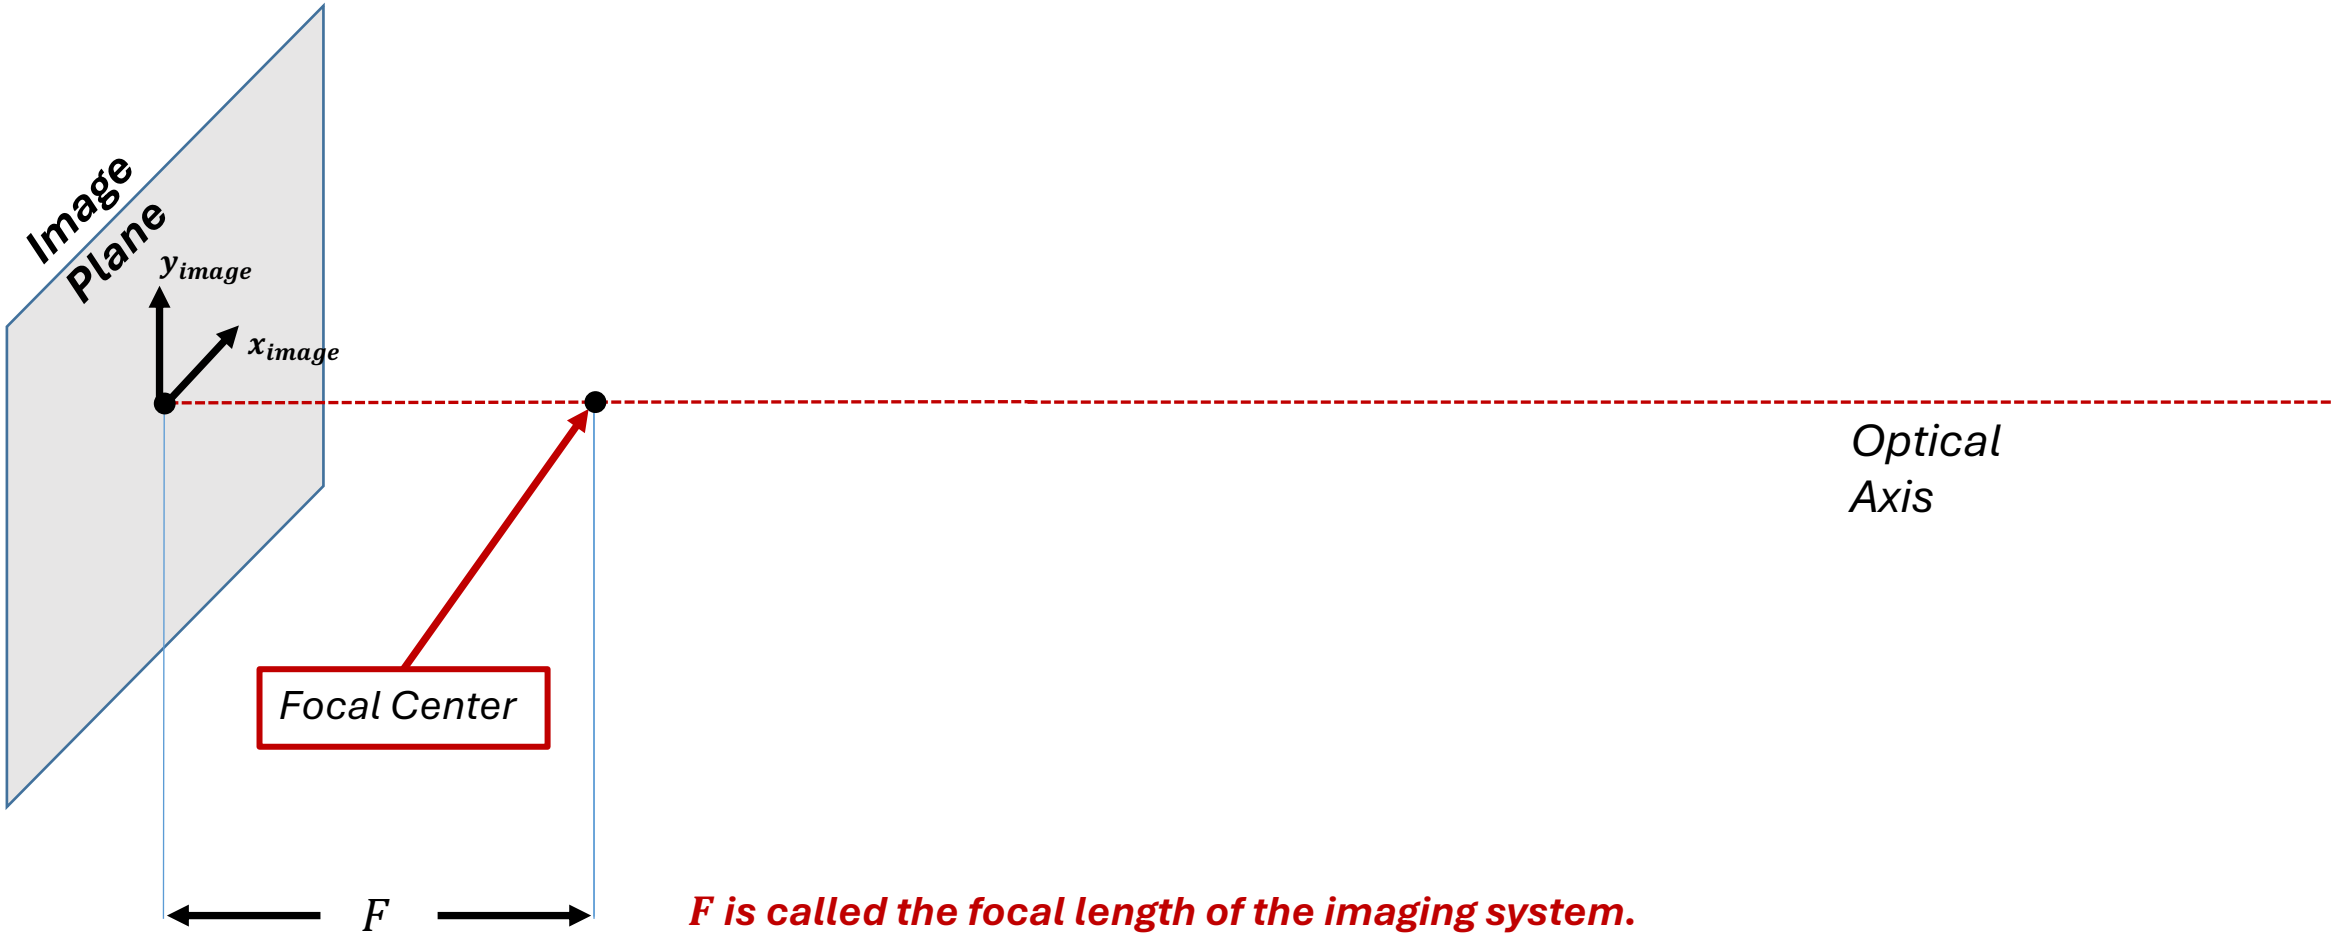
Stored at UT Austin

Niepce later teamed up with Daguerre, who eventually created Daguerrotypes

Pinhole Camera Geometry

The imaging geometry for the pinhole camera has several important properties:

- The image plane is located at distance F behind the focal center.
- The optical axis passes through the focal center, perpendicular to the image plane.

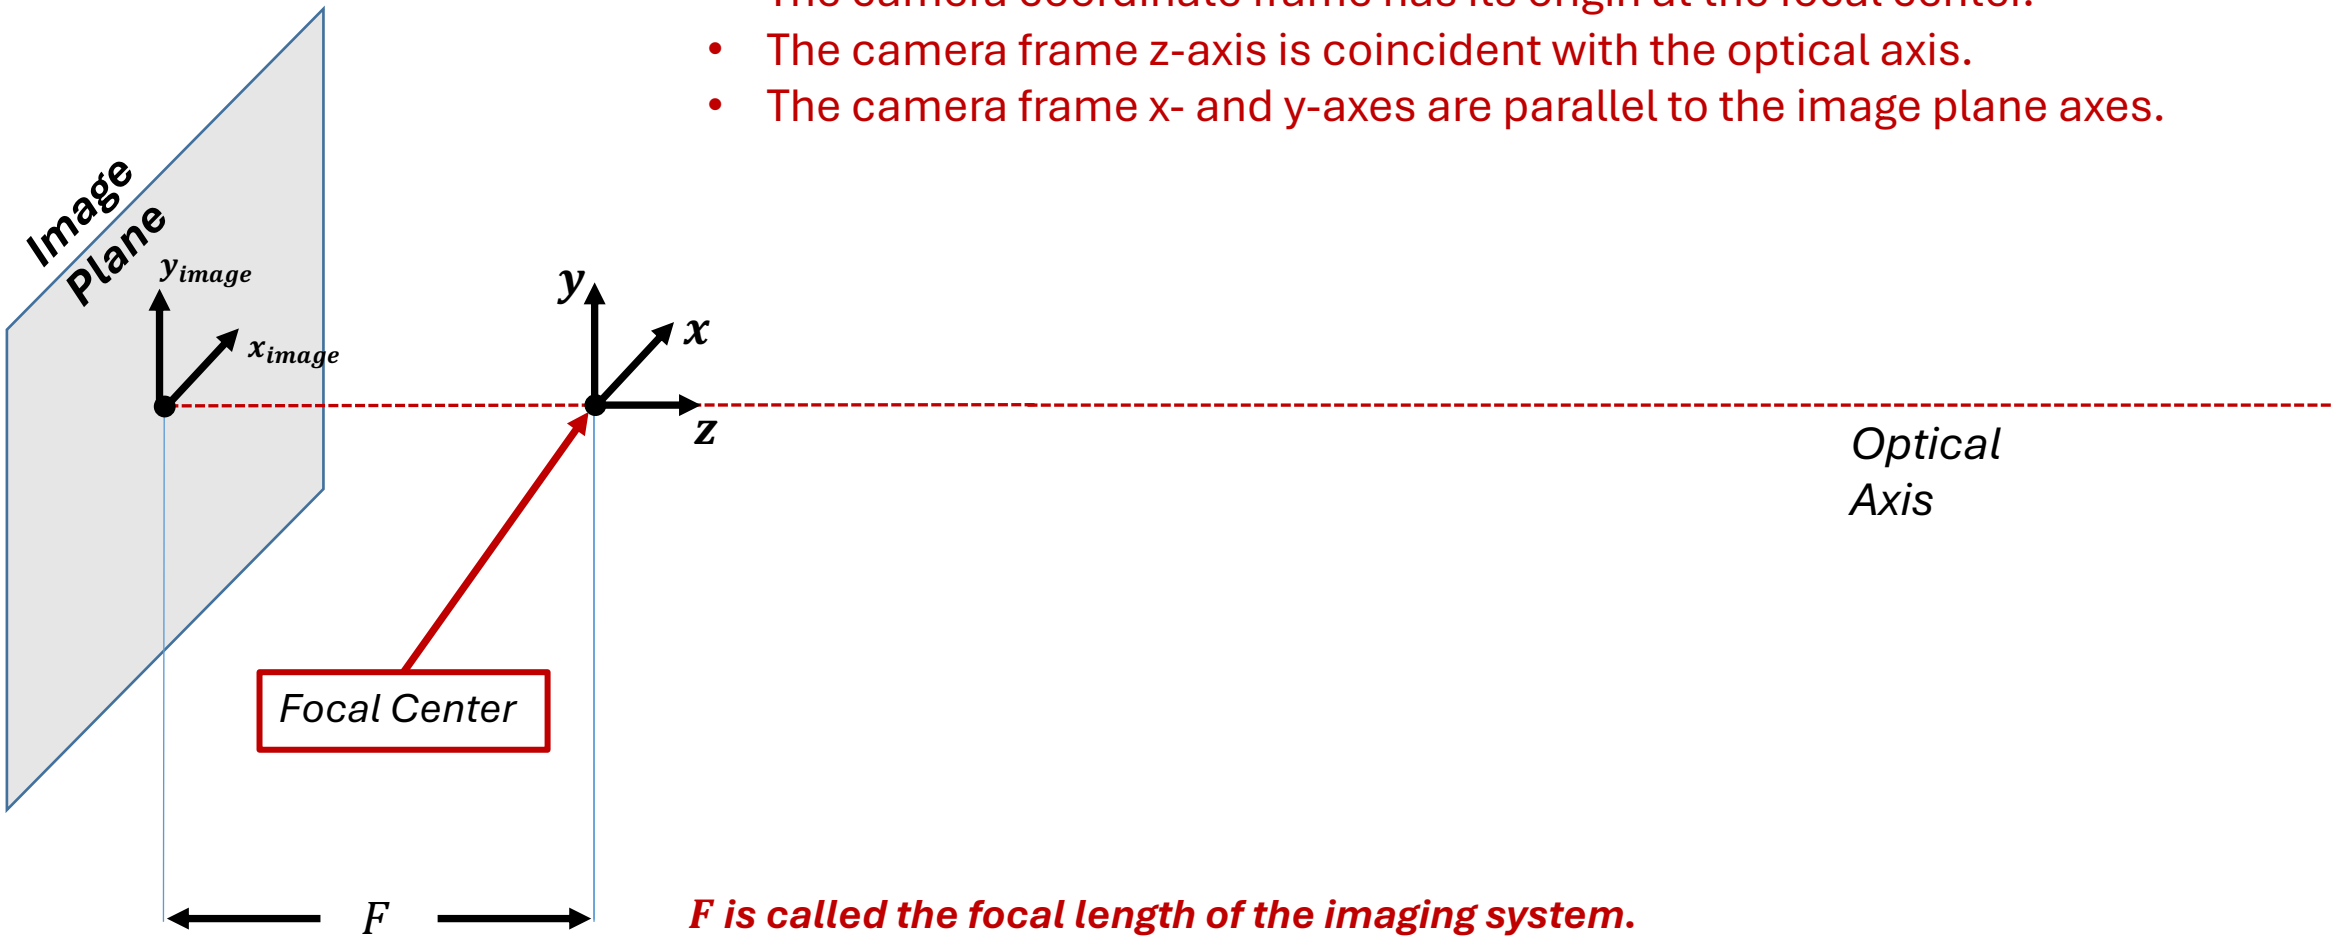


F is called the focal length of the imaging system.

Pinhole Camera Geometry

The imaging geometry for the pinhole camera has several important properties:

- The image plane is located at distance F behind the focal center.
- The optical axis passes through the focal center, perpendicular to the image plane.
- The camera coordinate frame has its origin at the focal center.
- The camera frame z-axis is coincident with the optical axis.
- The camera frame x- and y-axes are parallel to the image plane axes.

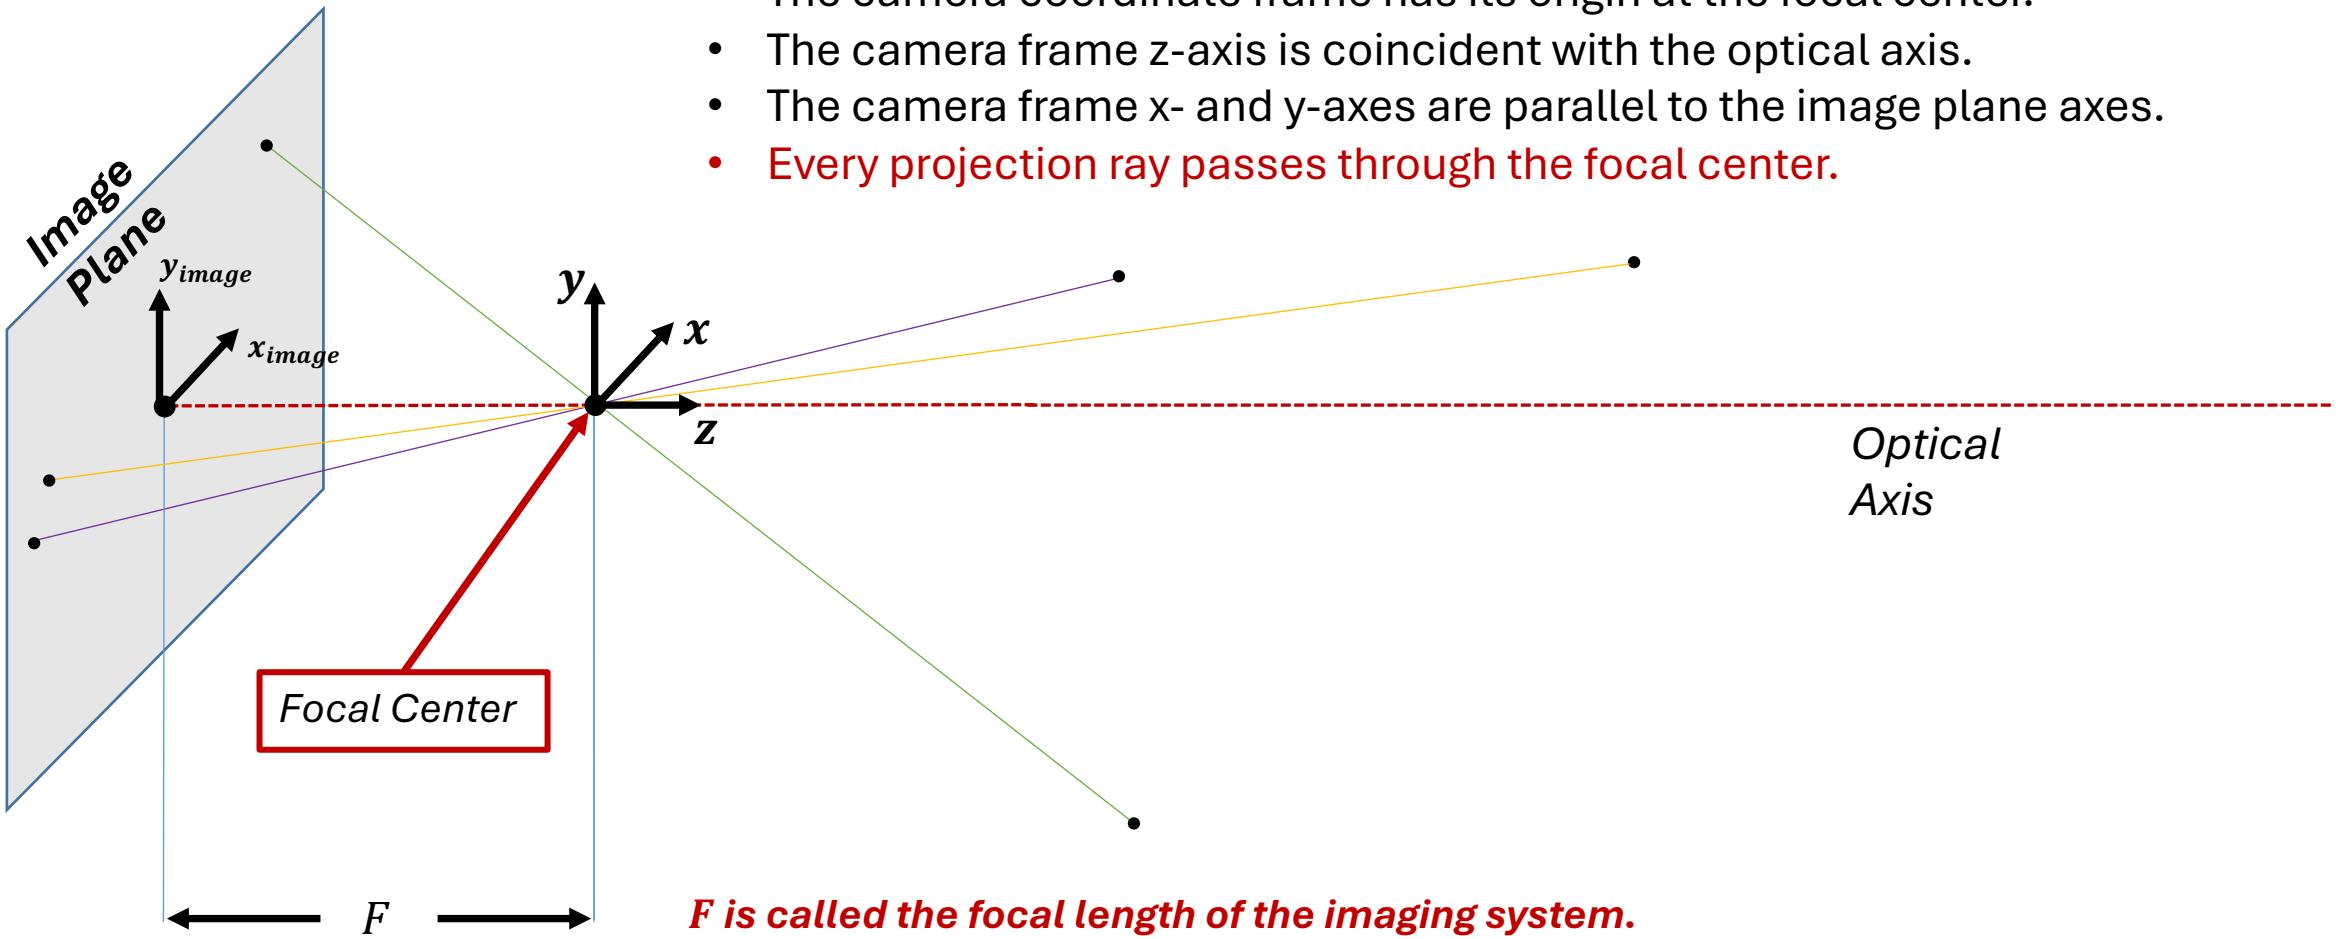


F is called the focal length of the imaging system.

Pinhole Camera Geometry

The imaging geometry for the pinhole camera has several important properties:

- The image plane is located at distance F behind the focal center.
- The optical axis passes through the focal center, perpendicular to the image plane.
- The camera coordinate frame has its origin at the focal center.
- The camera frame z-axis is coincident with the optical axis.
- The camera frame x- and y-axes are parallel to the image plane axes.
- Every projection ray passes through the focal center.

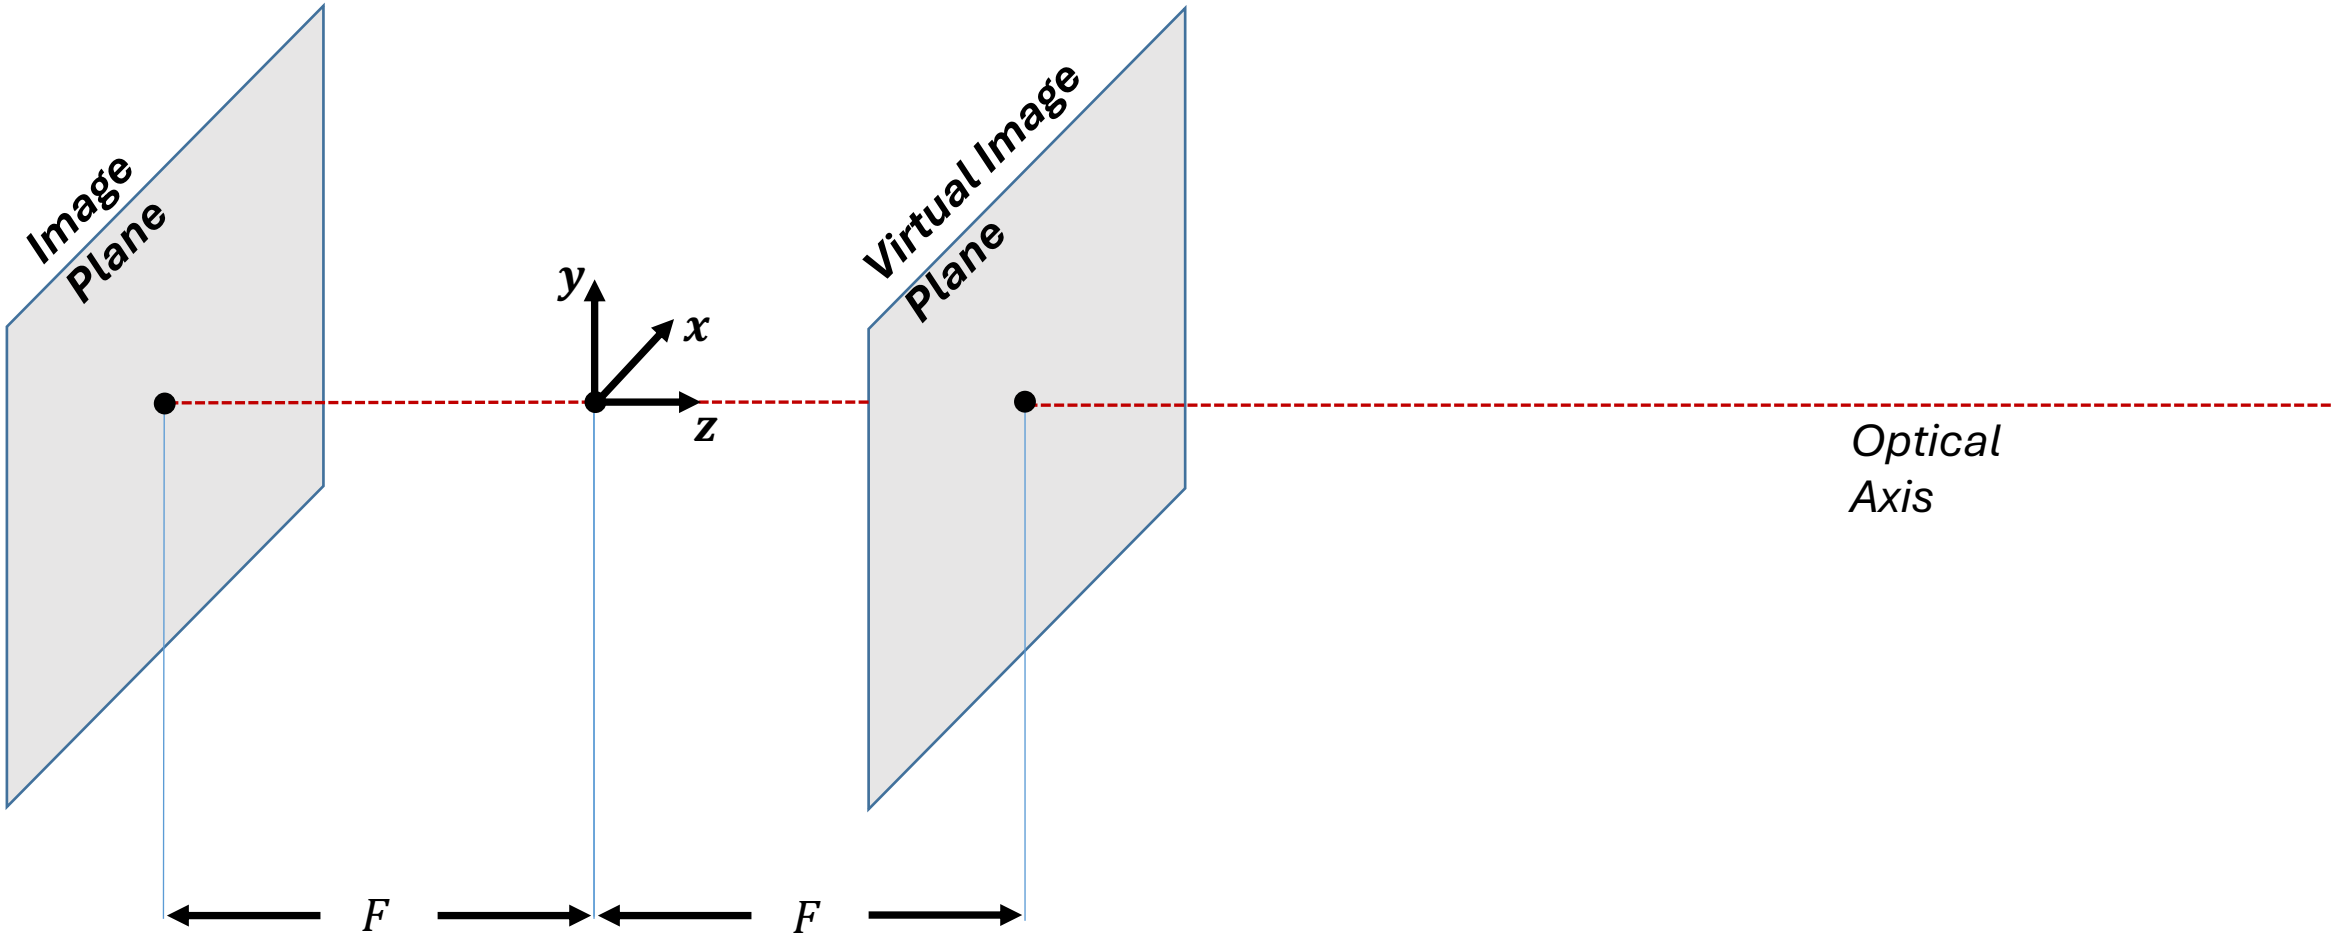


F is called the focal length of the imaging system.

Pinhole Camera

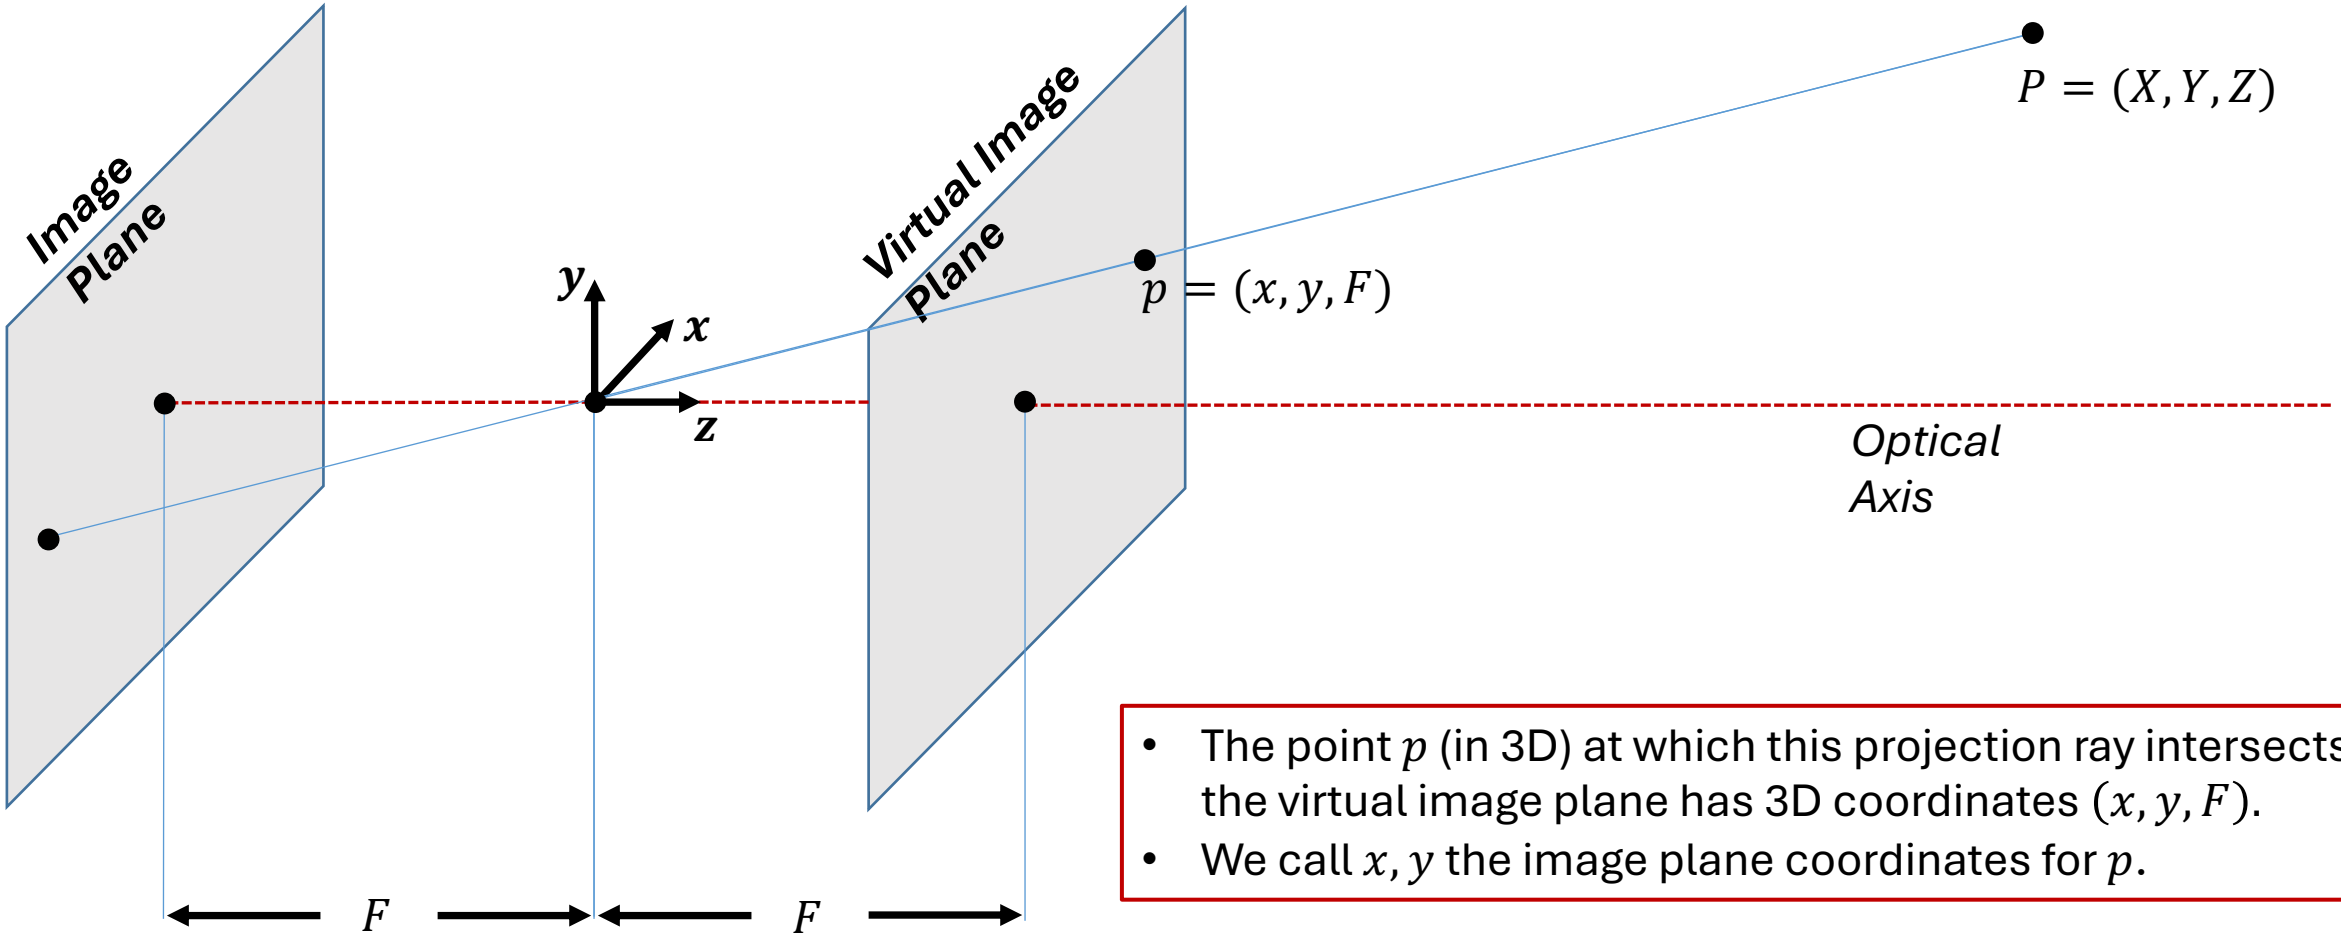
Life is so much easier if we insert a ***virtual image plane*** in front of the focal center.

No more need for upside-down image geometry!



Pinhole Camera

The point $P = (X, Y, Z)$ lies on a projection ray that passes through P and the focal center, and that intersects both the image plane and the virtual image plane.



- The point p (in 3D) at which this projection ray intersects the virtual image plane has 3D coordinates (x, y, F).
- We call x, y the image plane coordinates for p.

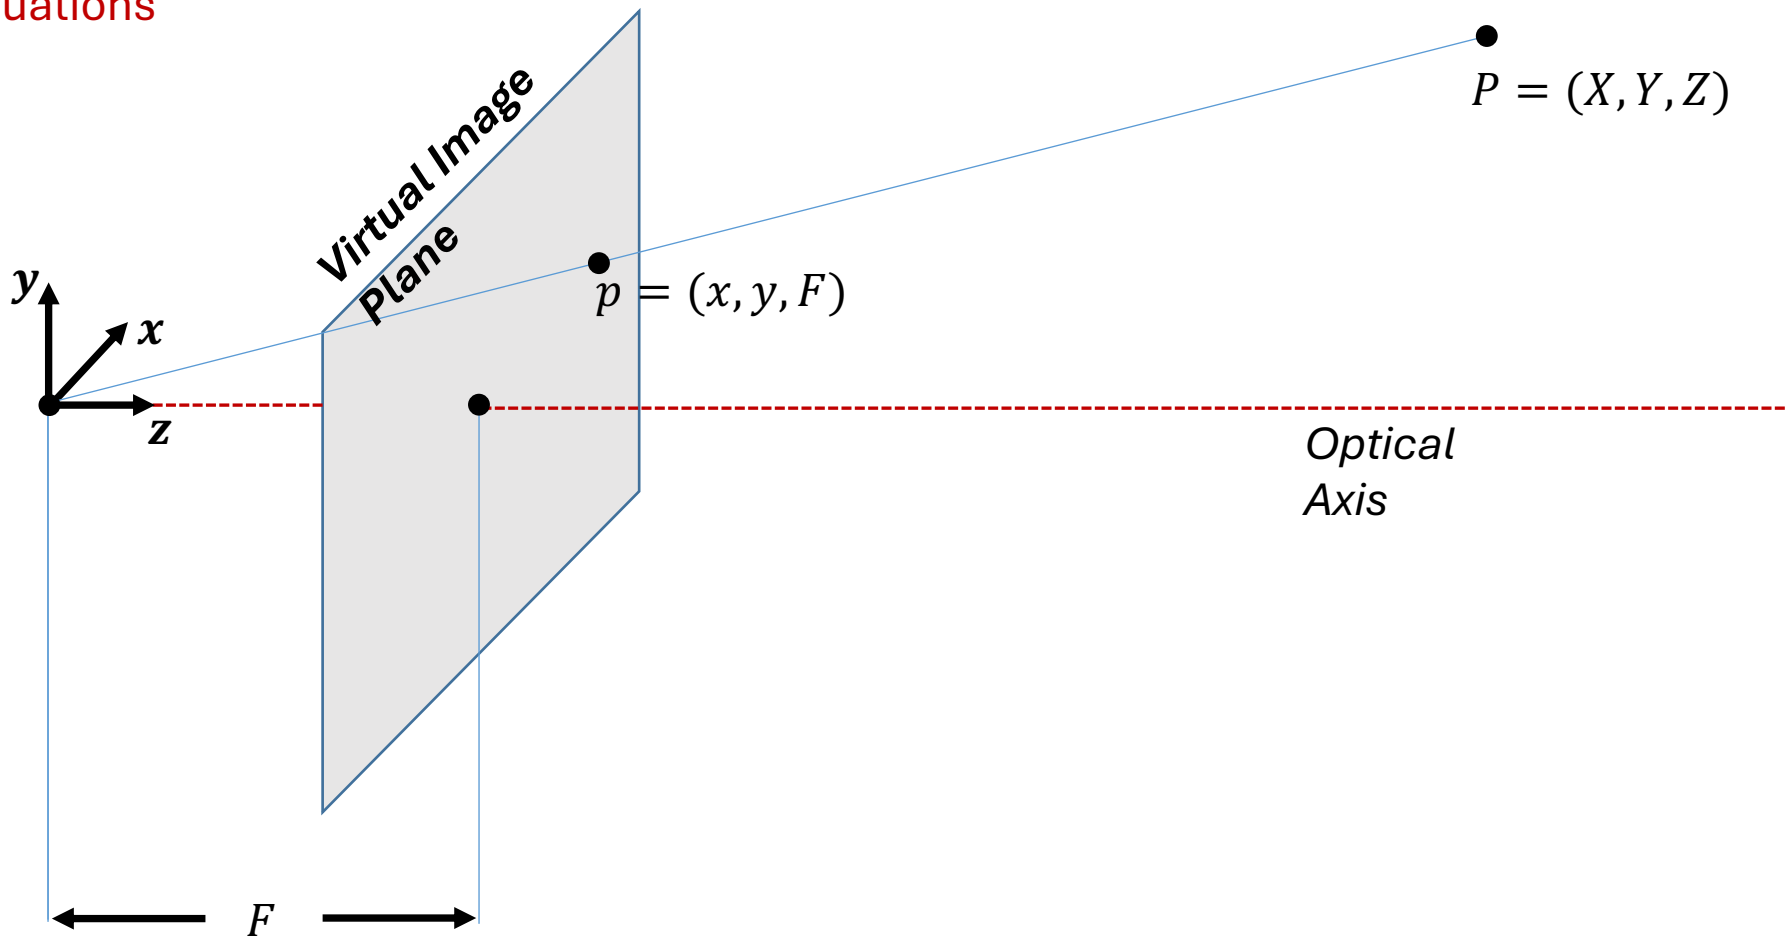
Pinhole Camera

Because p and P lie on the same projection ray through the origin, we have

$$\lambda p = P$$

We can write this as three equations

$$\lambda x = X$$
$$\lambda y = Y$$
$$\lambda F = Z$$



Pinhole Camera

Because p and P lie on the same projection ray through the origin, we have

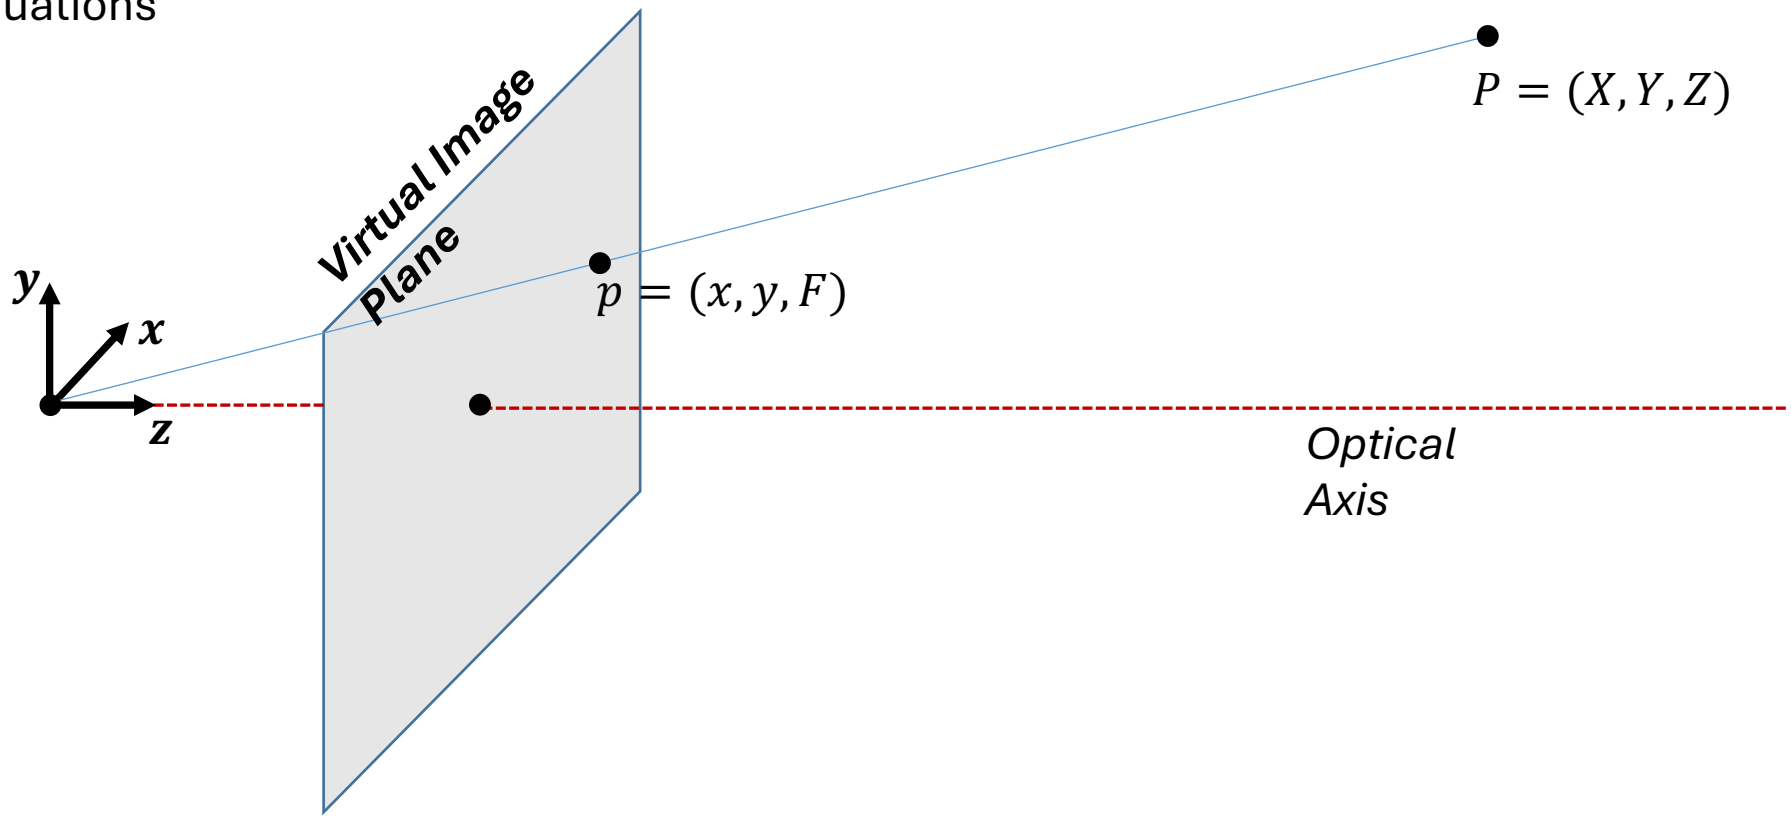
$$\lambda p = P$$

We can write this as three equations

$$\lambda x = X$$
$$\lambda y = Y$$
$$\lambda F = Z$$

Solving for λ yields

$$\lambda = \frac{Z}{F}$$



Pinhole Camera

Because p and P lie on the same projection ray through the origin, we have

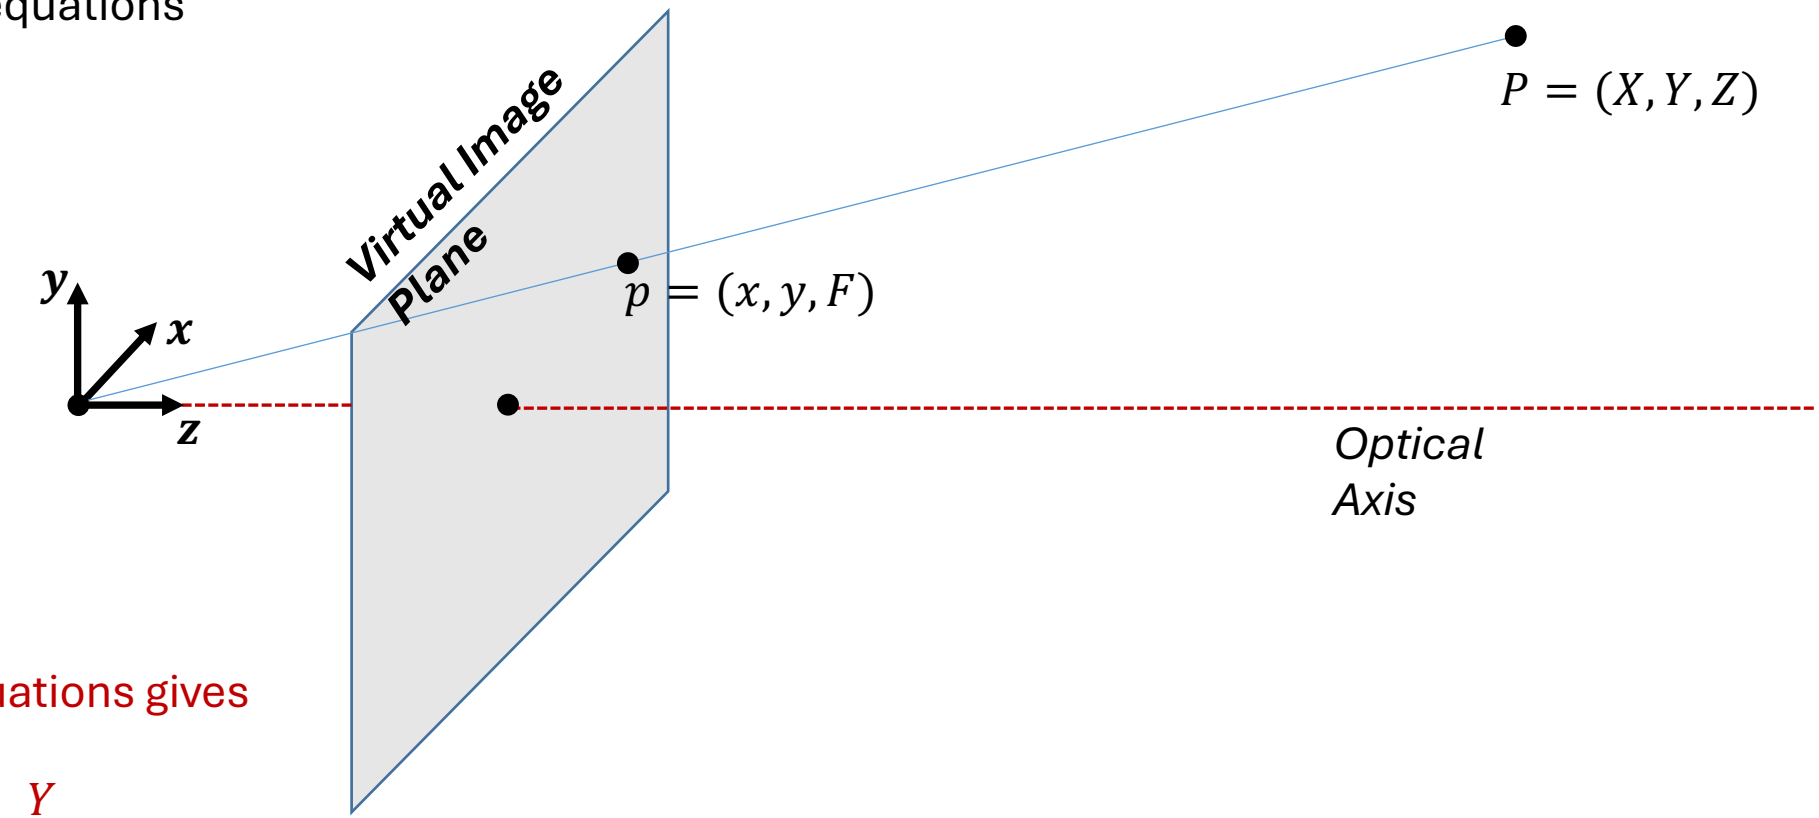
$$\lambda p = P$$

We can write this as three equations

$$\lambda x = X$$
$$\lambda y = Y$$
$$\lambda F = Z$$

Solving for λ yields

$$\lambda = \frac{Z}{F}$$



Substituting into the first equations gives

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

Pinhole Camera

Because p and P lie on the same projection ray through the origin, we have

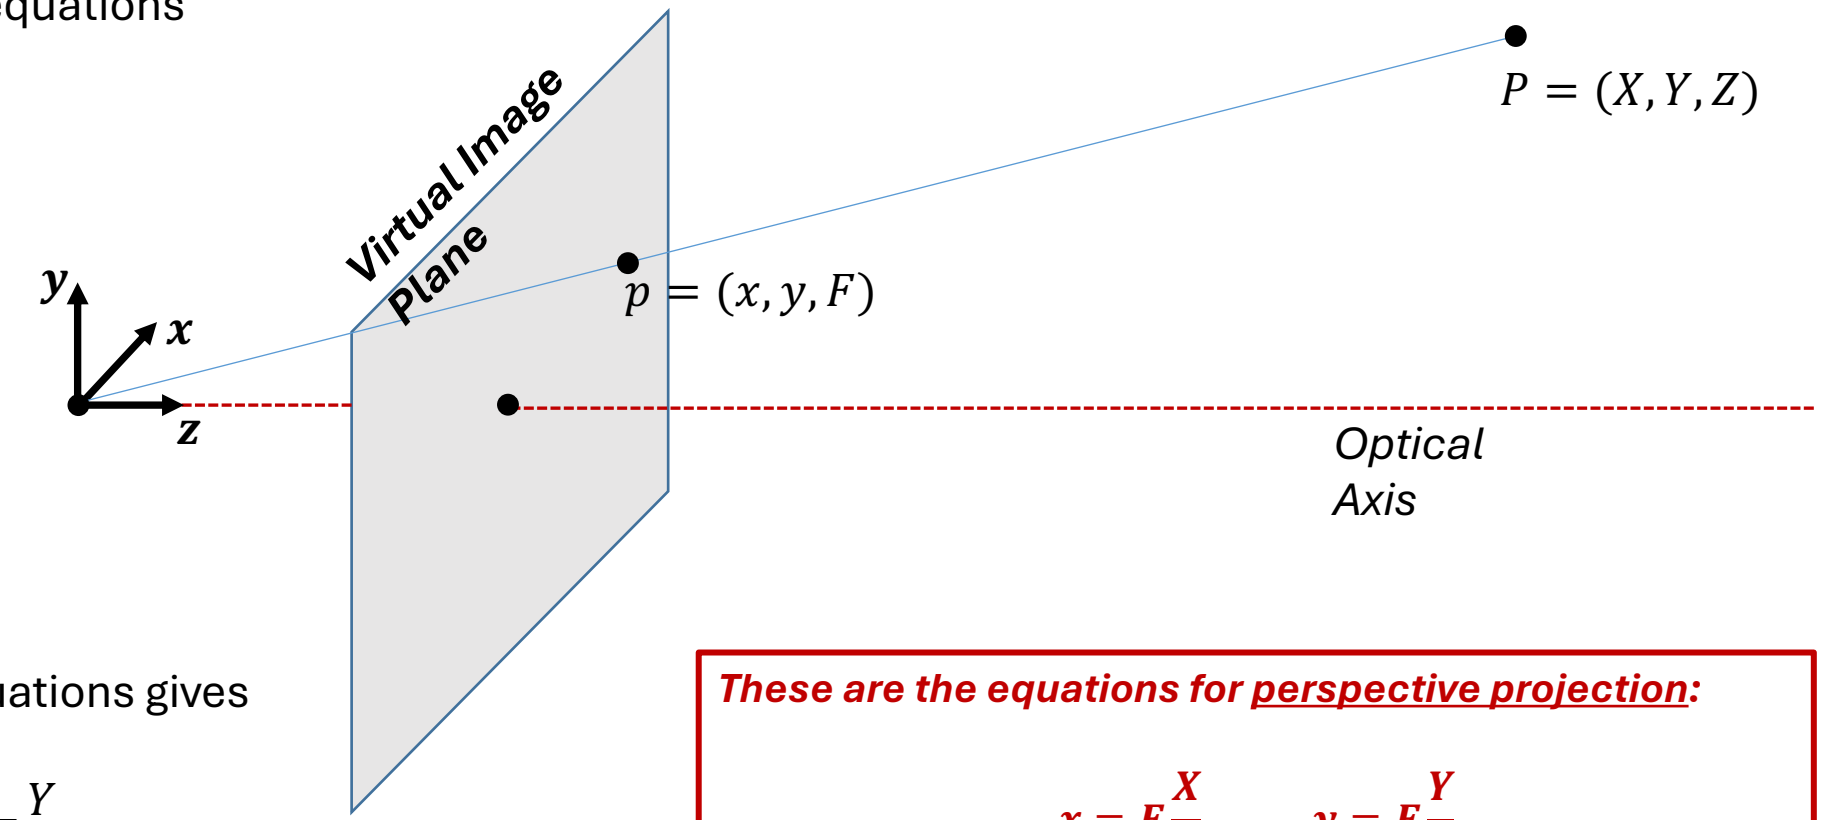
$$\lambda p = P$$

We can write this as three equations

$$\lambda x = X$$
$$\lambda y = Y$$
$$\lambda F = Z$$

Solving for λ yields

$$\lambda = \frac{Z}{F}$$



Substituting into the first equations gives

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

These are the equations for perspective projection:

$$\boldsymbol{x = F\frac{X}{Z}, \qquad y = F\frac{Y}{Z}}$$

Computer Vision Convention

- Fundamental equation:

$$(X, Y, Z) \rightarrow (x, y) = \left(F\frac{X}{Z}, F\frac{Y}{Z}\right)$$

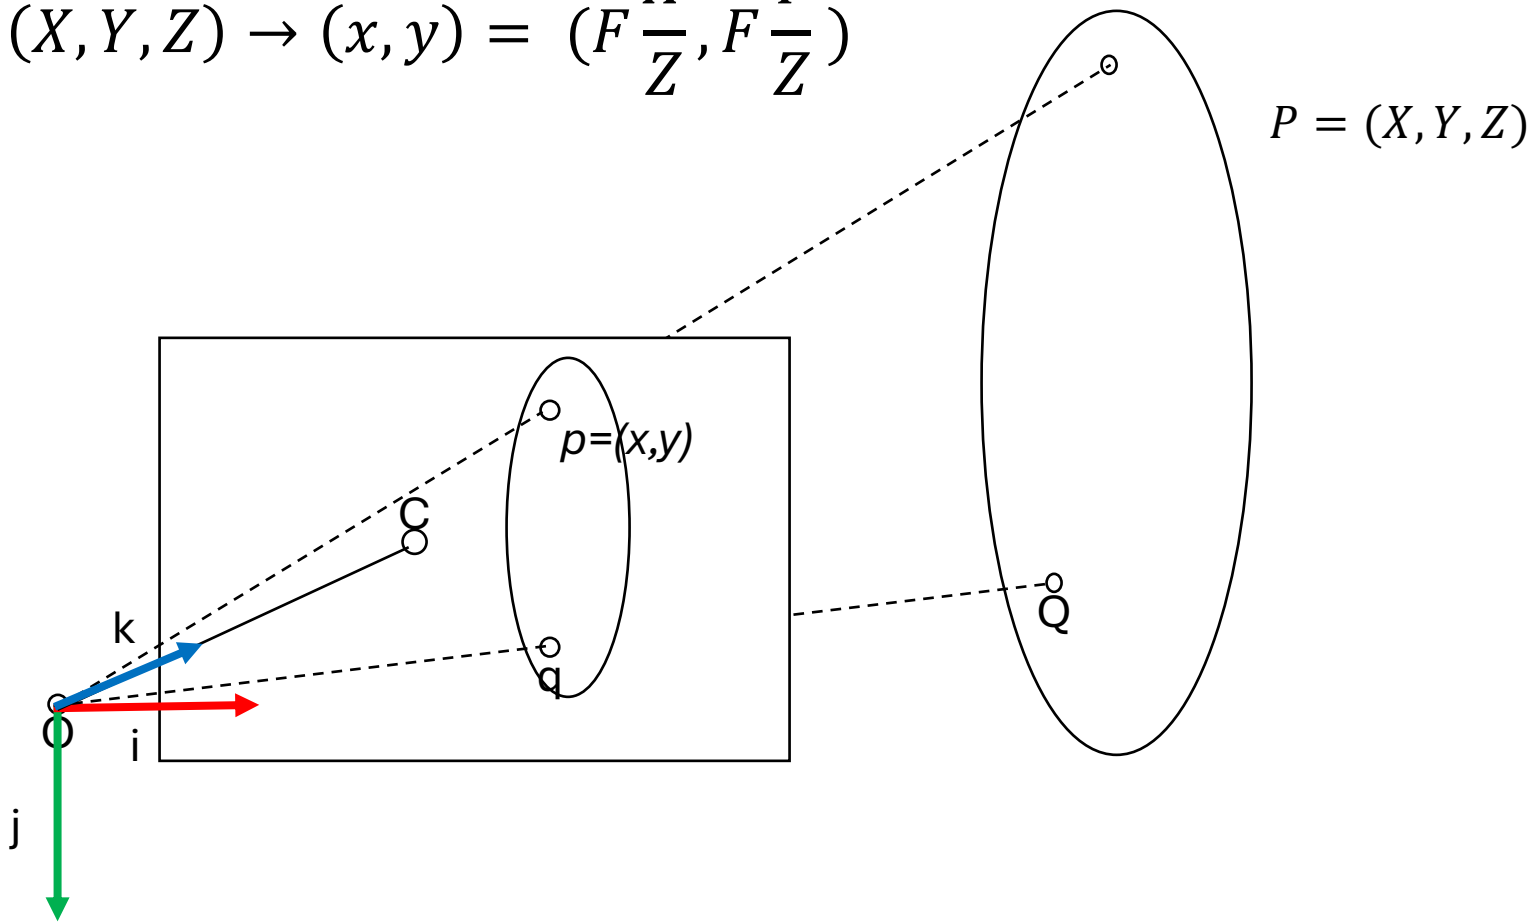
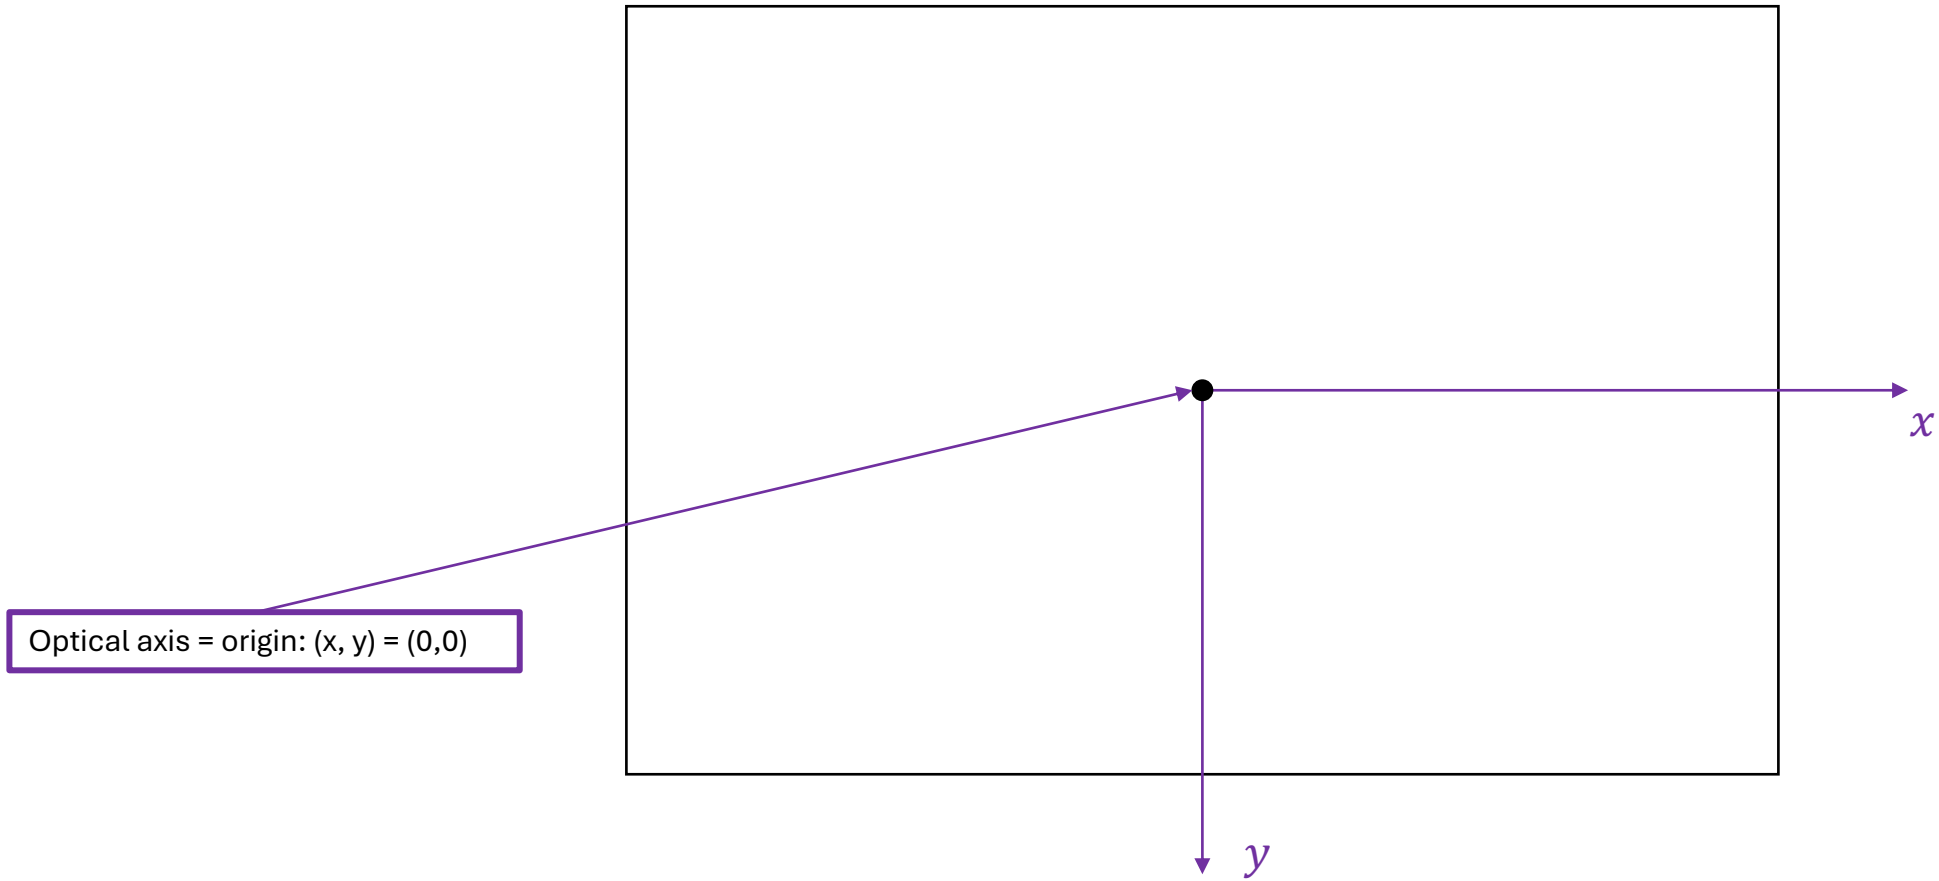


Image plane coordinates (2D) convention

Image plane coordinate frame (x, y):

$$(x, y) = (F\frac{X}{Z}, F\frac{Y}{Z})$$

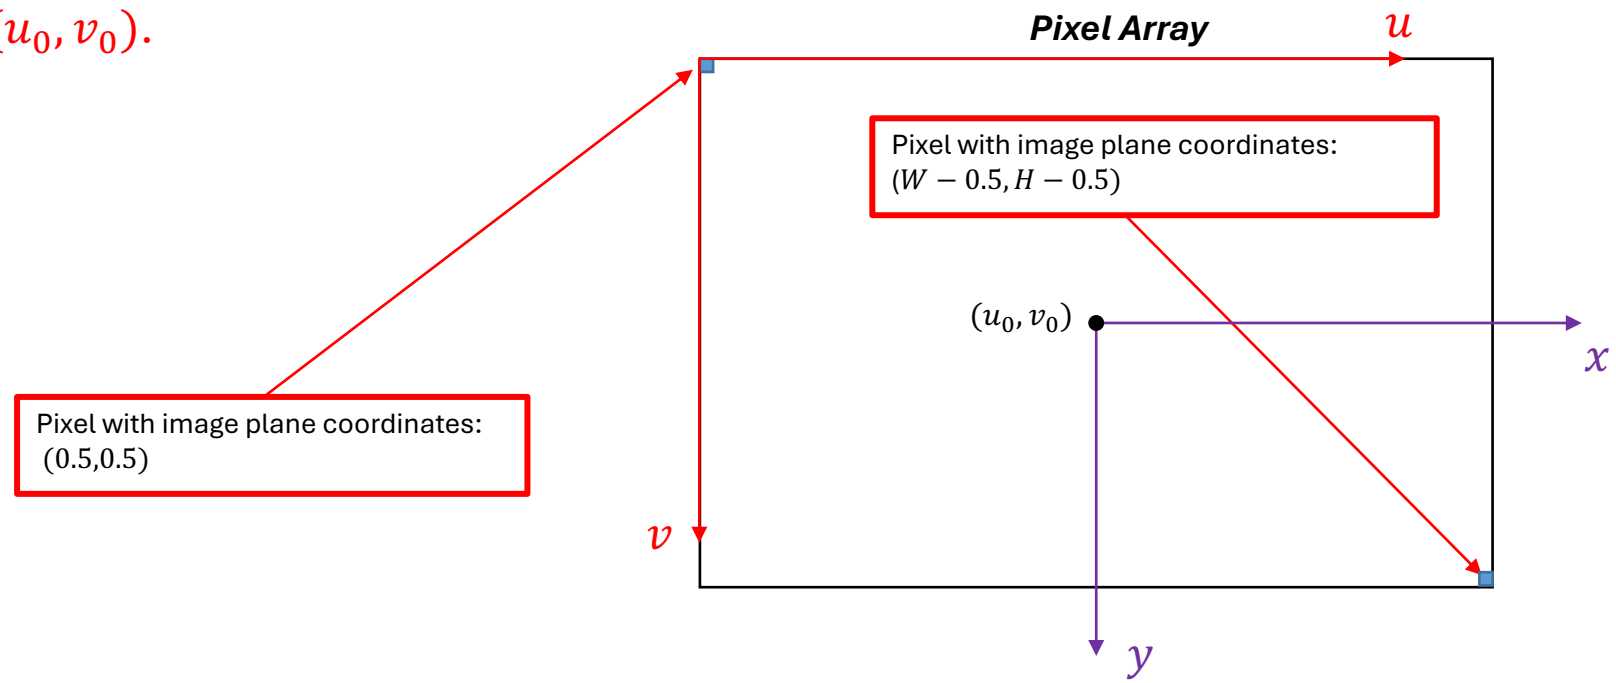


Sensor Coordinates

- Instead of a continuous image plane, real cameras have a 2D array of sensors that correspond to pixels in the image.
- When we make measurements in an image, we measure ***sensor coordinates (u,v)***, not image plane coordinates (x, y).

Sensor coordinate frame:

- The top, left pixel is location $(\mathrm{r}, \mathrm{c}) = 0, 0$ in the sensor array.
- The bottom, right pixel has location $(\mathrm{r}, \mathrm{c}) = H - 1, W - 1$ in the sensor array.
- The sensor coordinates or a pixel, u, v correspond to the center of the corresponding pixel.
  - Top, left pixel is $(0.5,0.5)$
  - Bottom right pixel is $(W - 0.5, H - 0.5)$
- Note that the v-axis points ***down***.
- The optical axis has coordinates (u_0, v_0).



Sensor Coordinates

From image-plane coordinates to sensor coordinates

To convert from image-plane coordinates to sensor coordinates u, v

- Scale x by pixel width
- Scale y by pixel height
- Shift coordinates by u_0, v_0:

$$u = u_0 + \alpha x, \qquad v = v_0 + \beta y$$

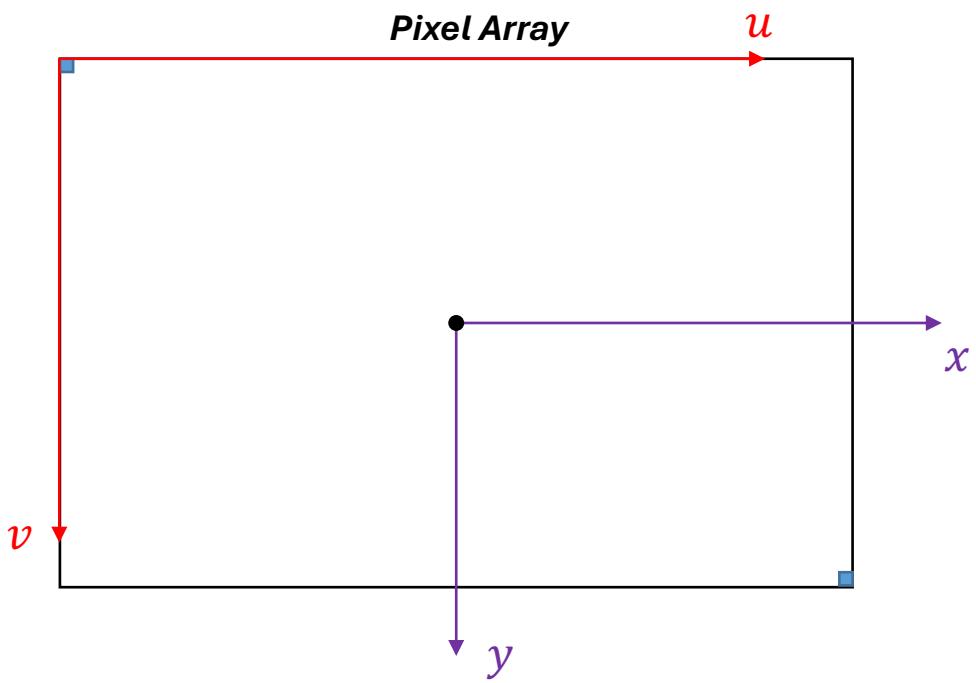
If we now substitute the perspective projection equations for x and y we obtain

$$u = u_0 + \alpha F \frac{X}{Z}, \qquad v = v_0 + \beta F \frac{Y}{Z}$$

If the camera happens to have square pixels, then $\alpha = \beta$ and we can simplify this to

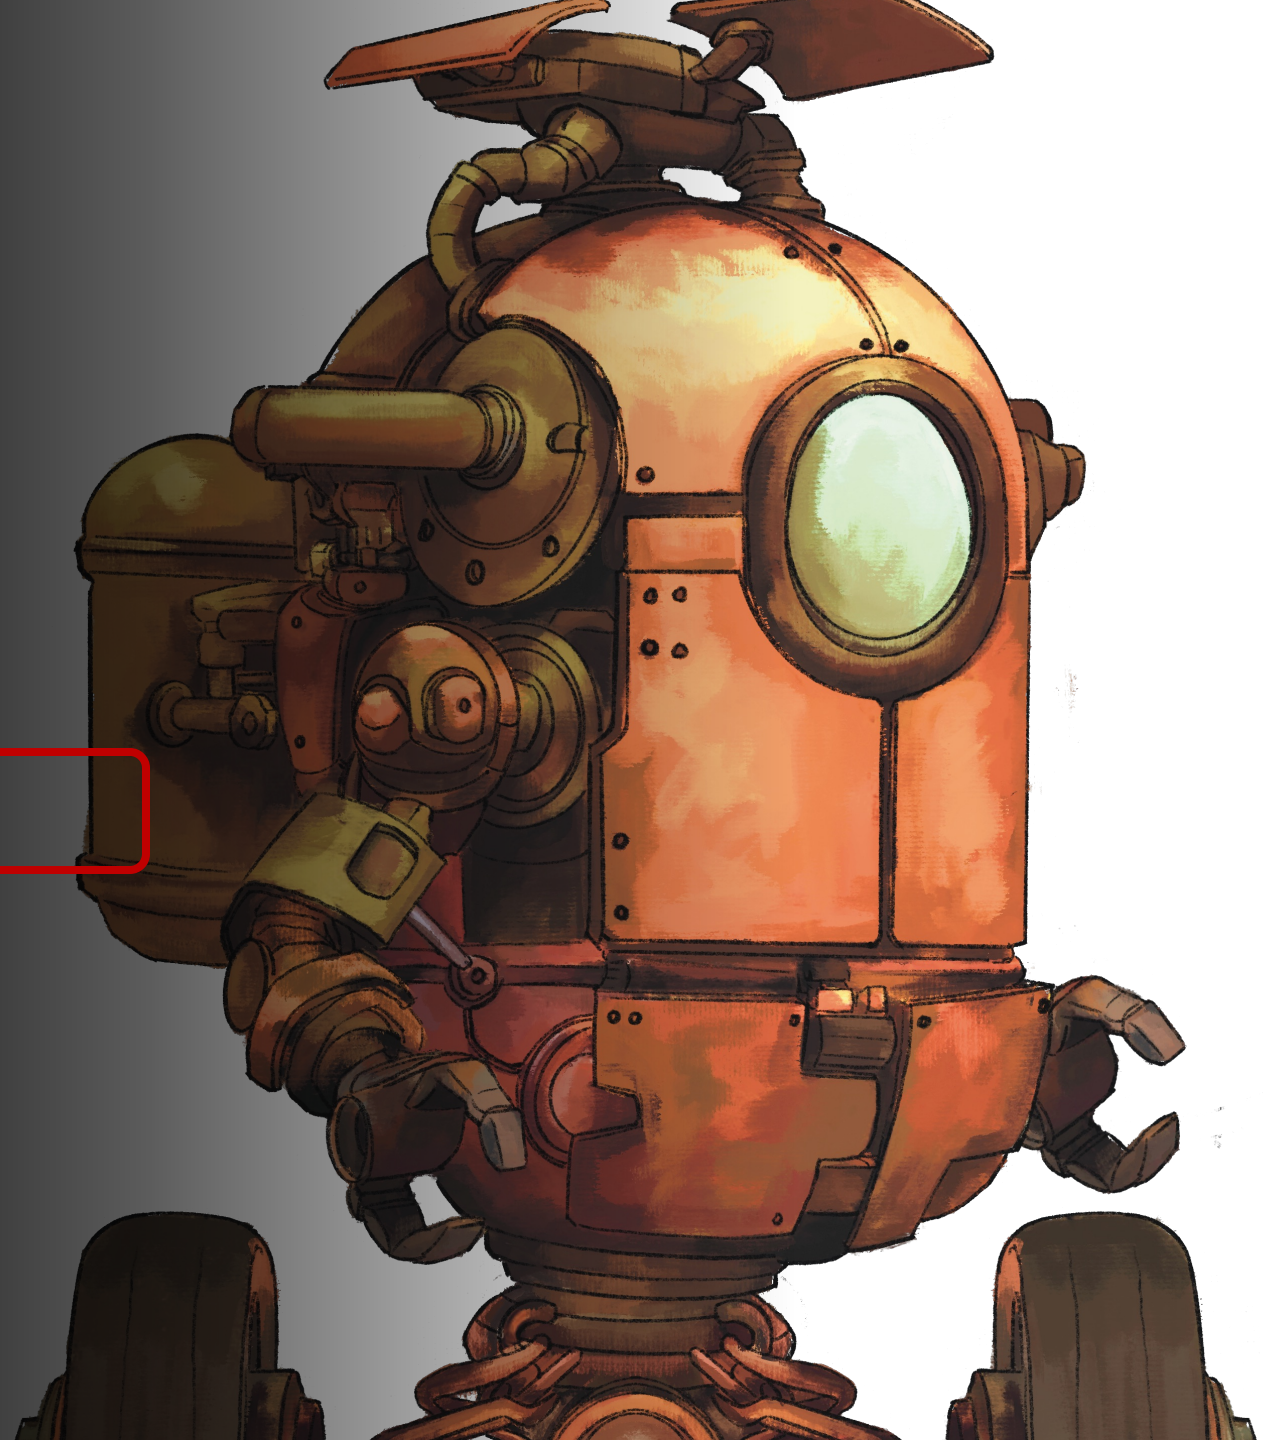
$$u = u_0 + f \frac{X}{Z}, \qquad v = v_0 + f \frac{Y}{Z}$$

Camera calibration is used to determine the values of u_0, v_0 and f.



Outline

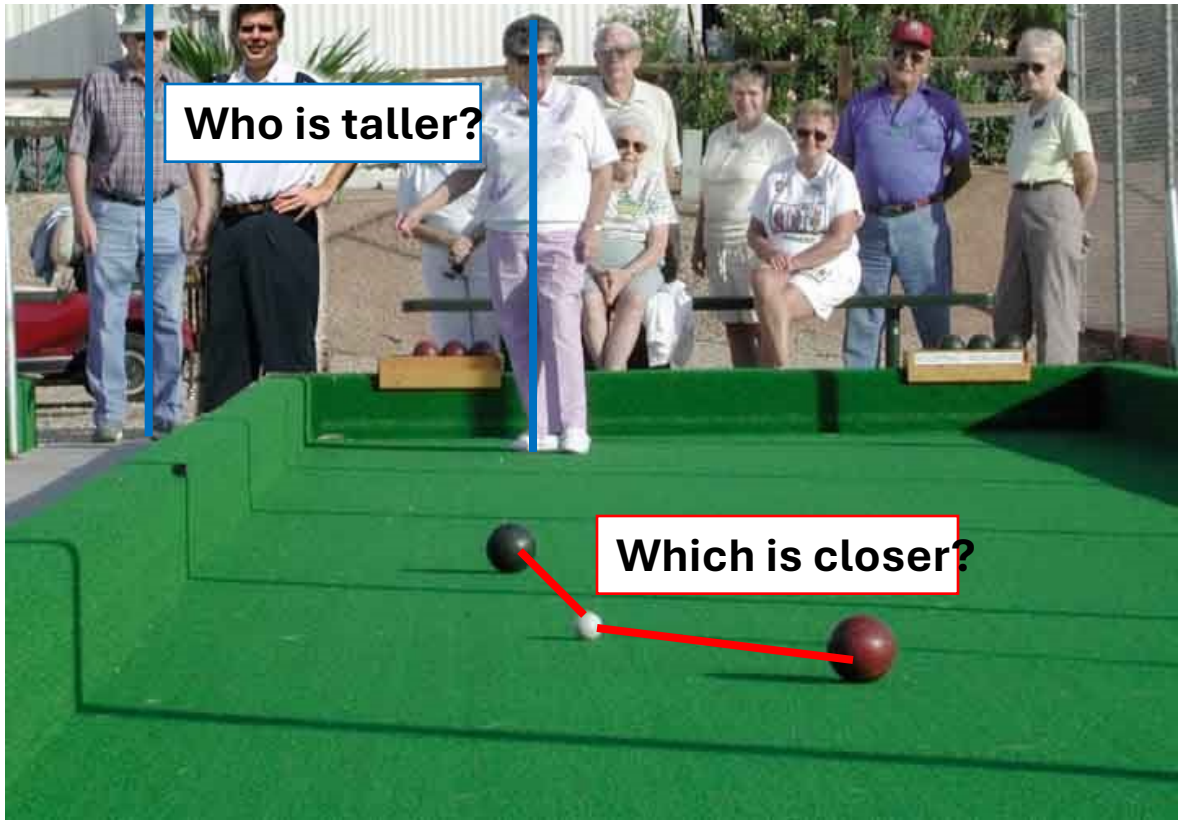
1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



Properties of projective Geometry

What is lost?

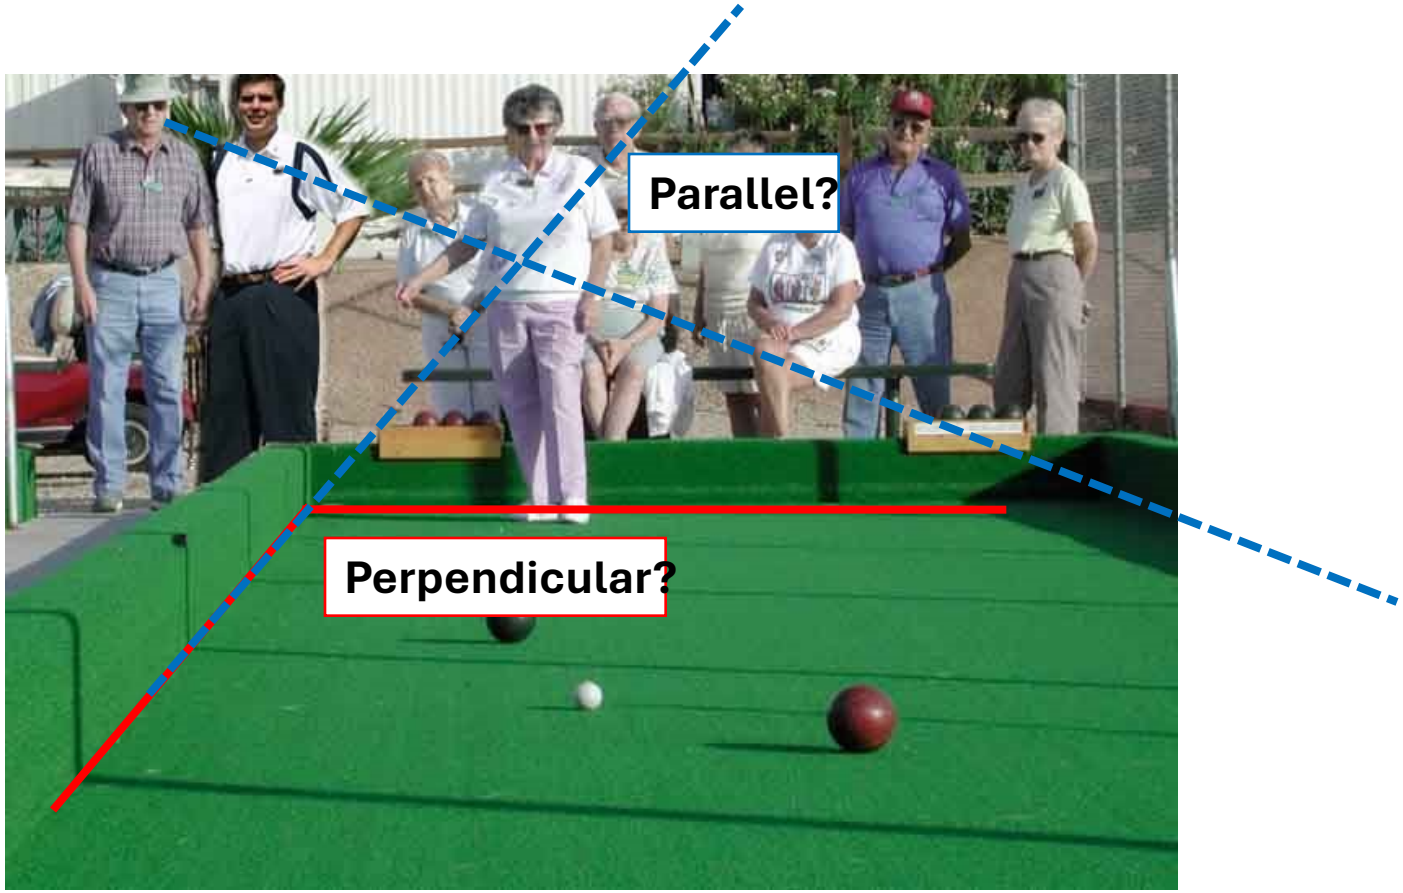
- Length



Properties of projective Geometry

What is lost?

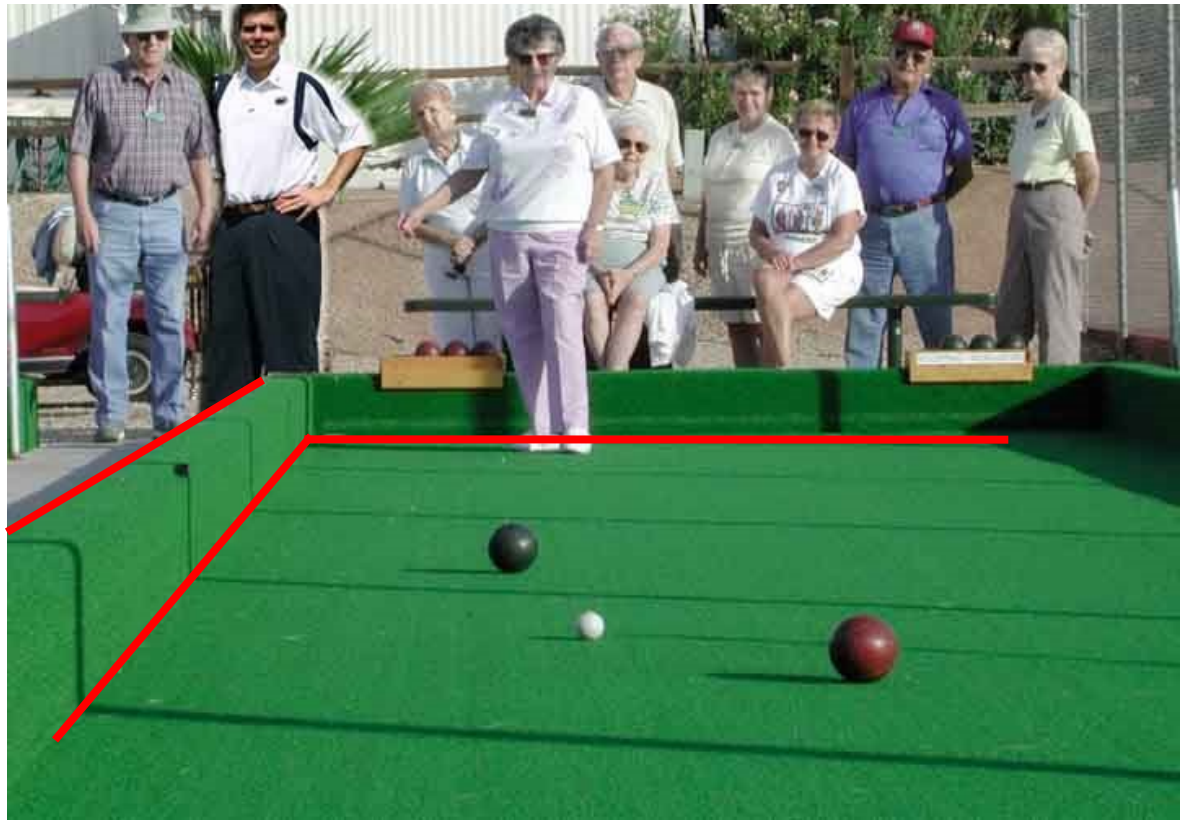
- Length
- Angles



Properties of projective Geometry

What is preserved?

- Straight lines are still straight



We can see infinity !

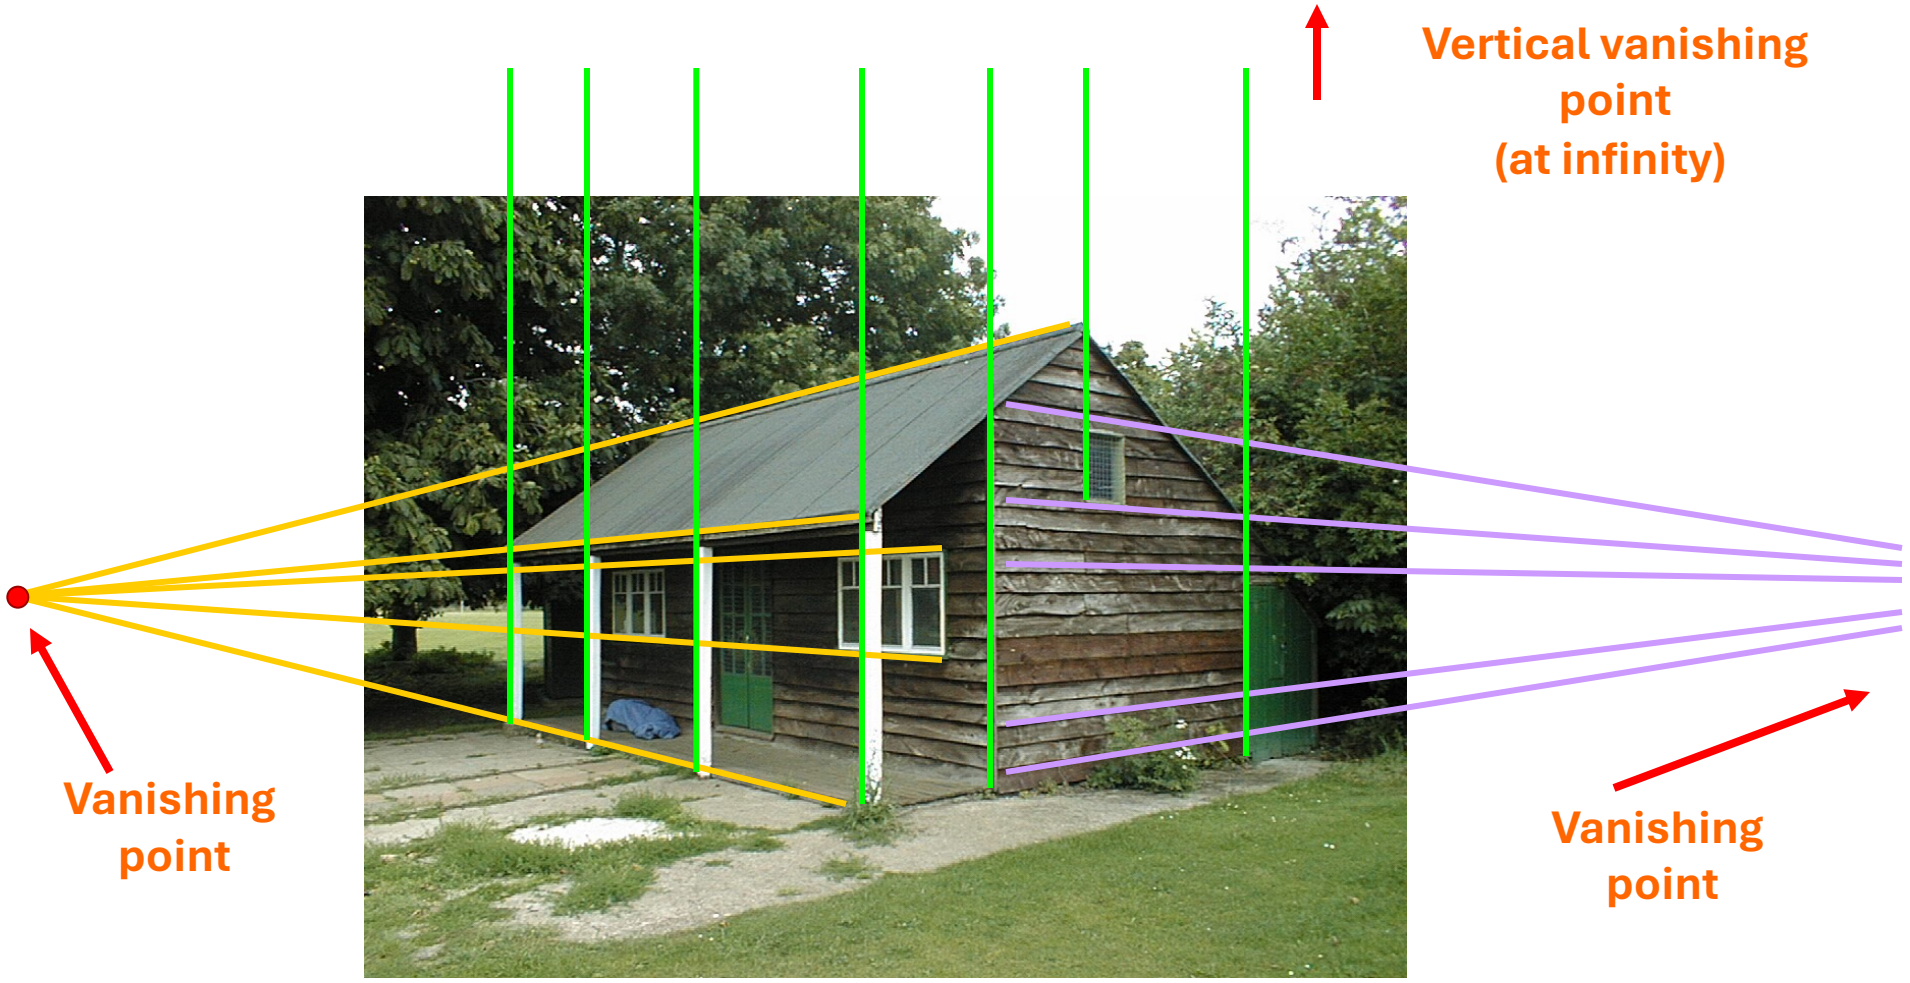
Where do parallel lines meet?

At infinity.



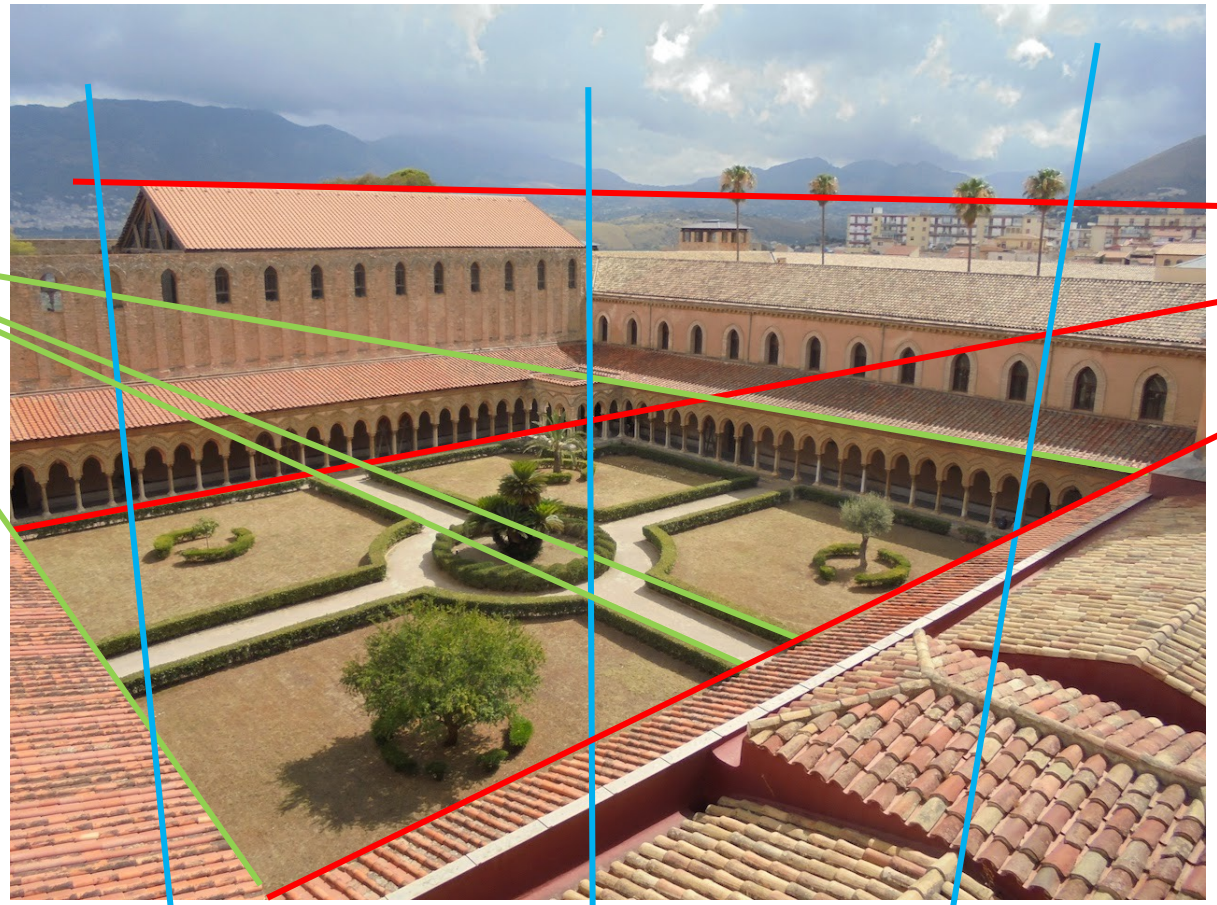
Railroad: parallel lines

Vanishing points and lines



Slide from Efros, Photo from Criminisi

Vanishing points and lines



Y-axis vanishing point

X- axis vanishing point

Z-axis vanishing point

Cattedrale di Monreale, Palermo 2014, by Frank Dellaert

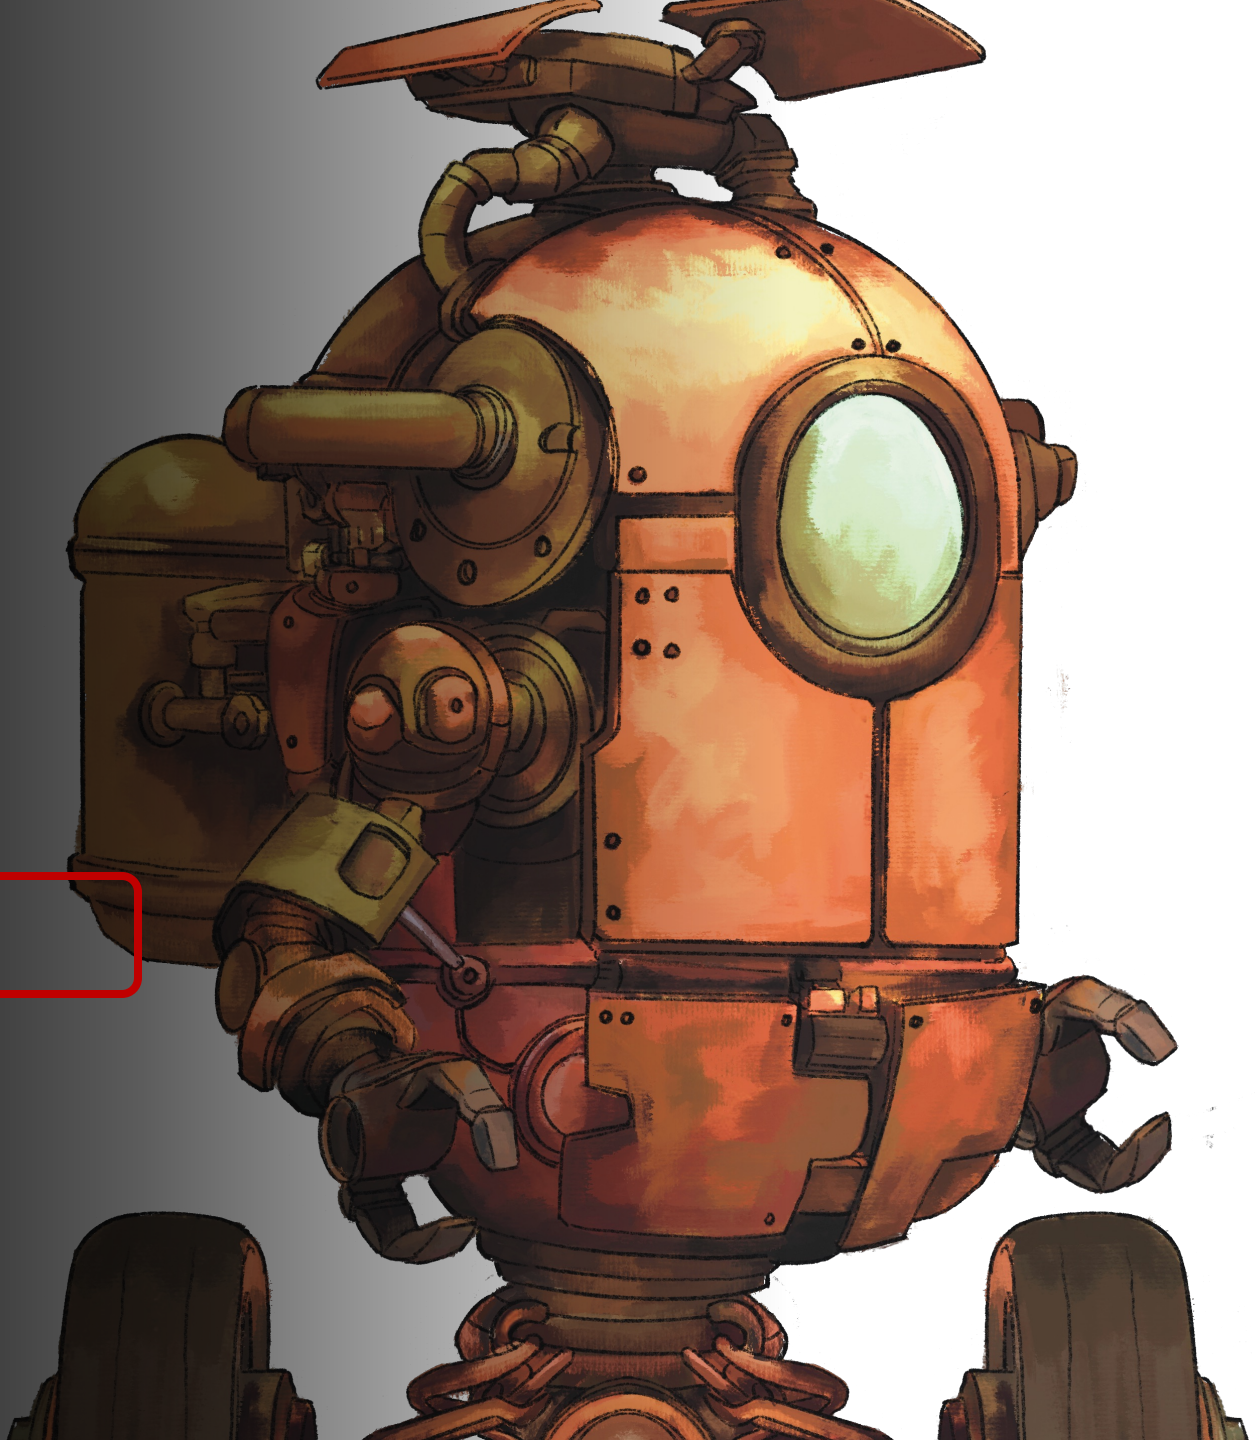
Vertical vanishing point



Cathedral in Antwerp

Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



Stereo Vision

- Stereo is used in the Human Visual System
- Very useful in computer vision as well
- Eliminates scale ambiguity

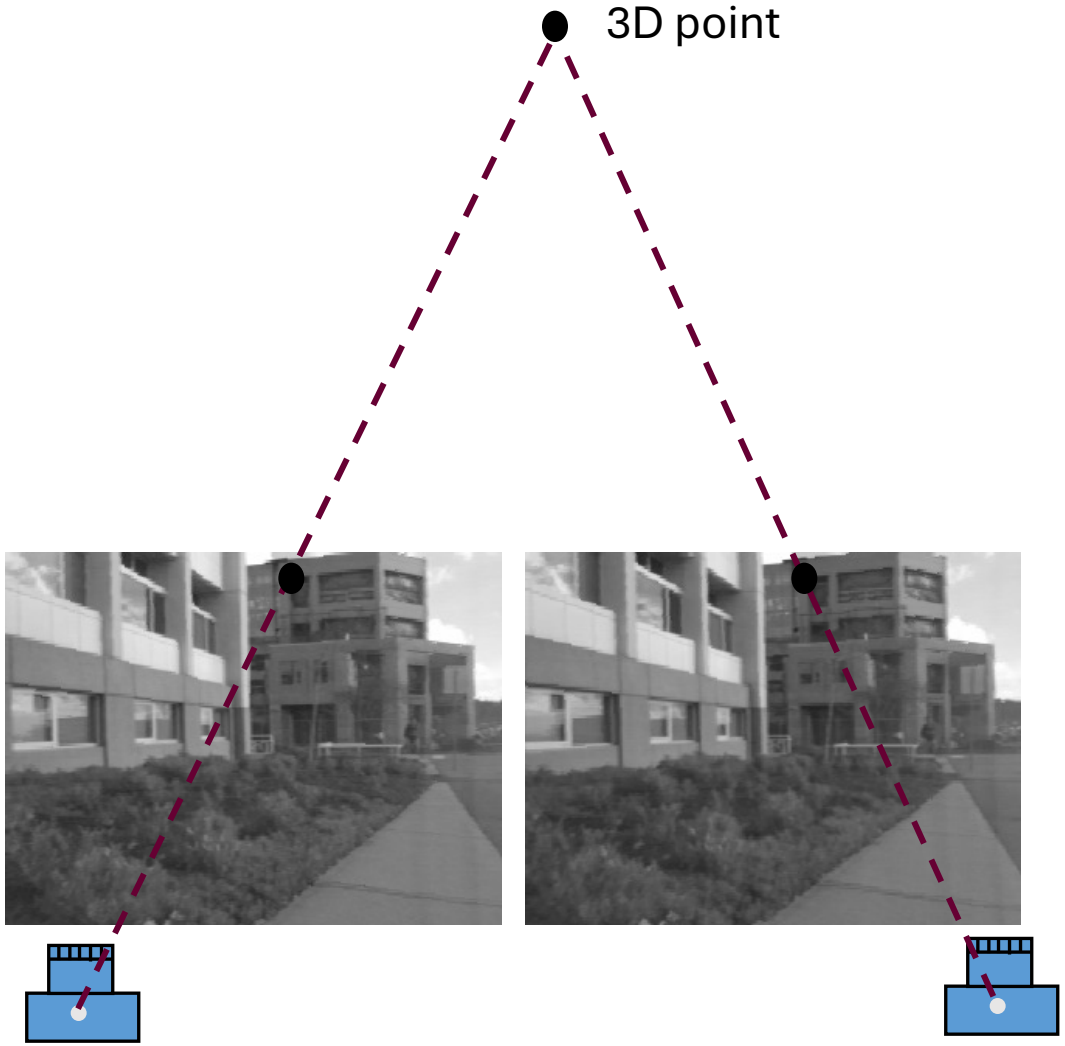
- Many slides adapted from F&P and Sing Bing Kang guest lecture

Etymology

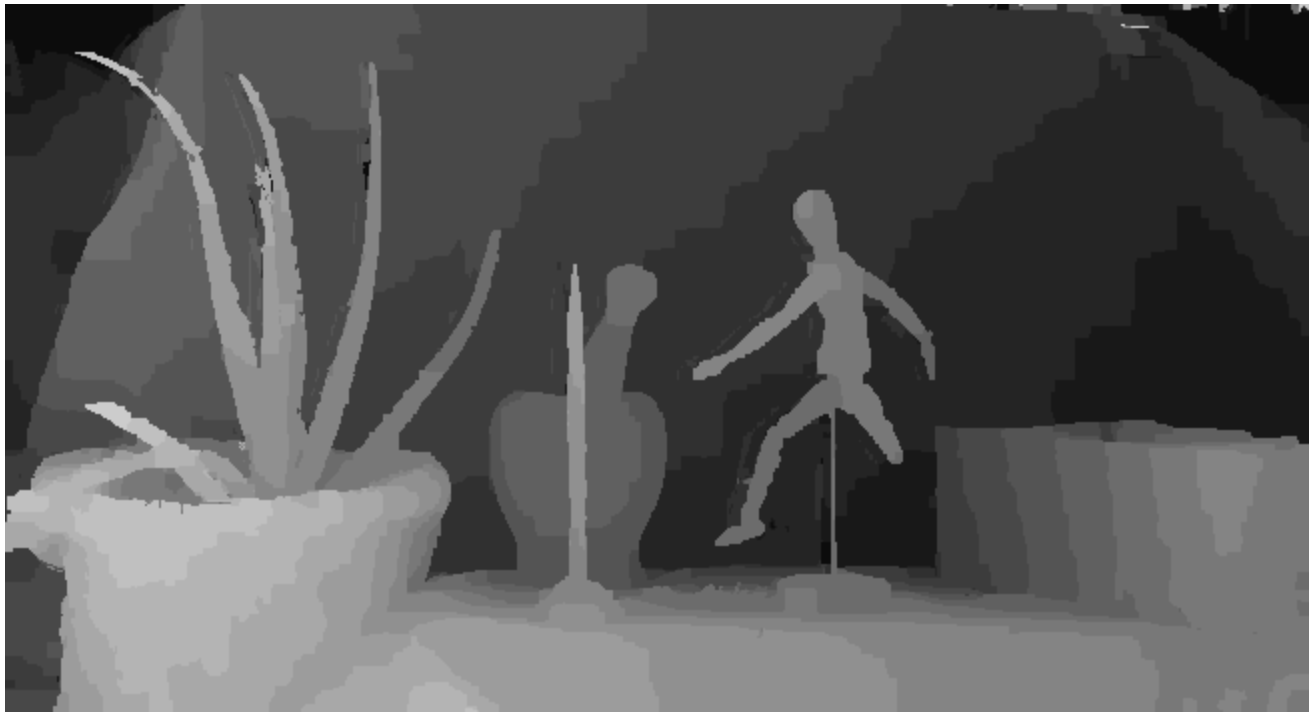
Stereo comes from the Greek word for *solid* (στερεο), and the term can be applied to any system using more than one channel

Effect of Moving Camera

- As camera is shifted (viewpoint changed):
  - 3D points are projected to different 2D locations
  - Amount of shift in projected 2D location depends on depth
- 2D shifts= stereo disparity



Example

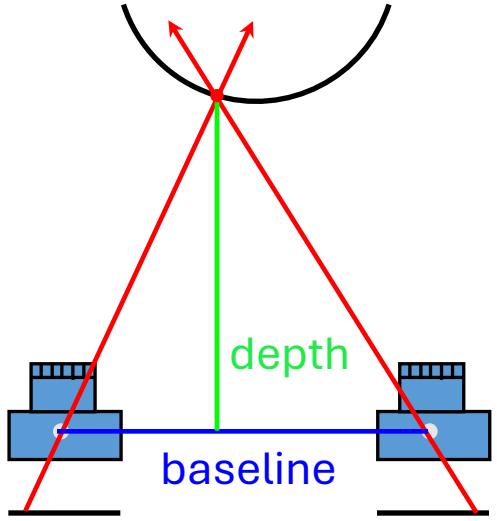


Right image

View Interpolation



Basic Idea of Stereo



Triangulate on two images of the same point to recover depth.

- Feature matching across views
- Calibrated cameras

Left



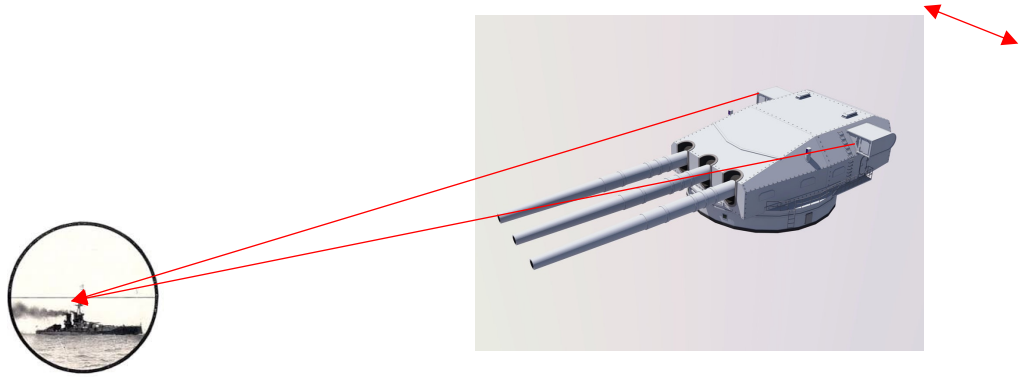
Right



Matching correlation windows across scan lines

Why is Stereo Useful?

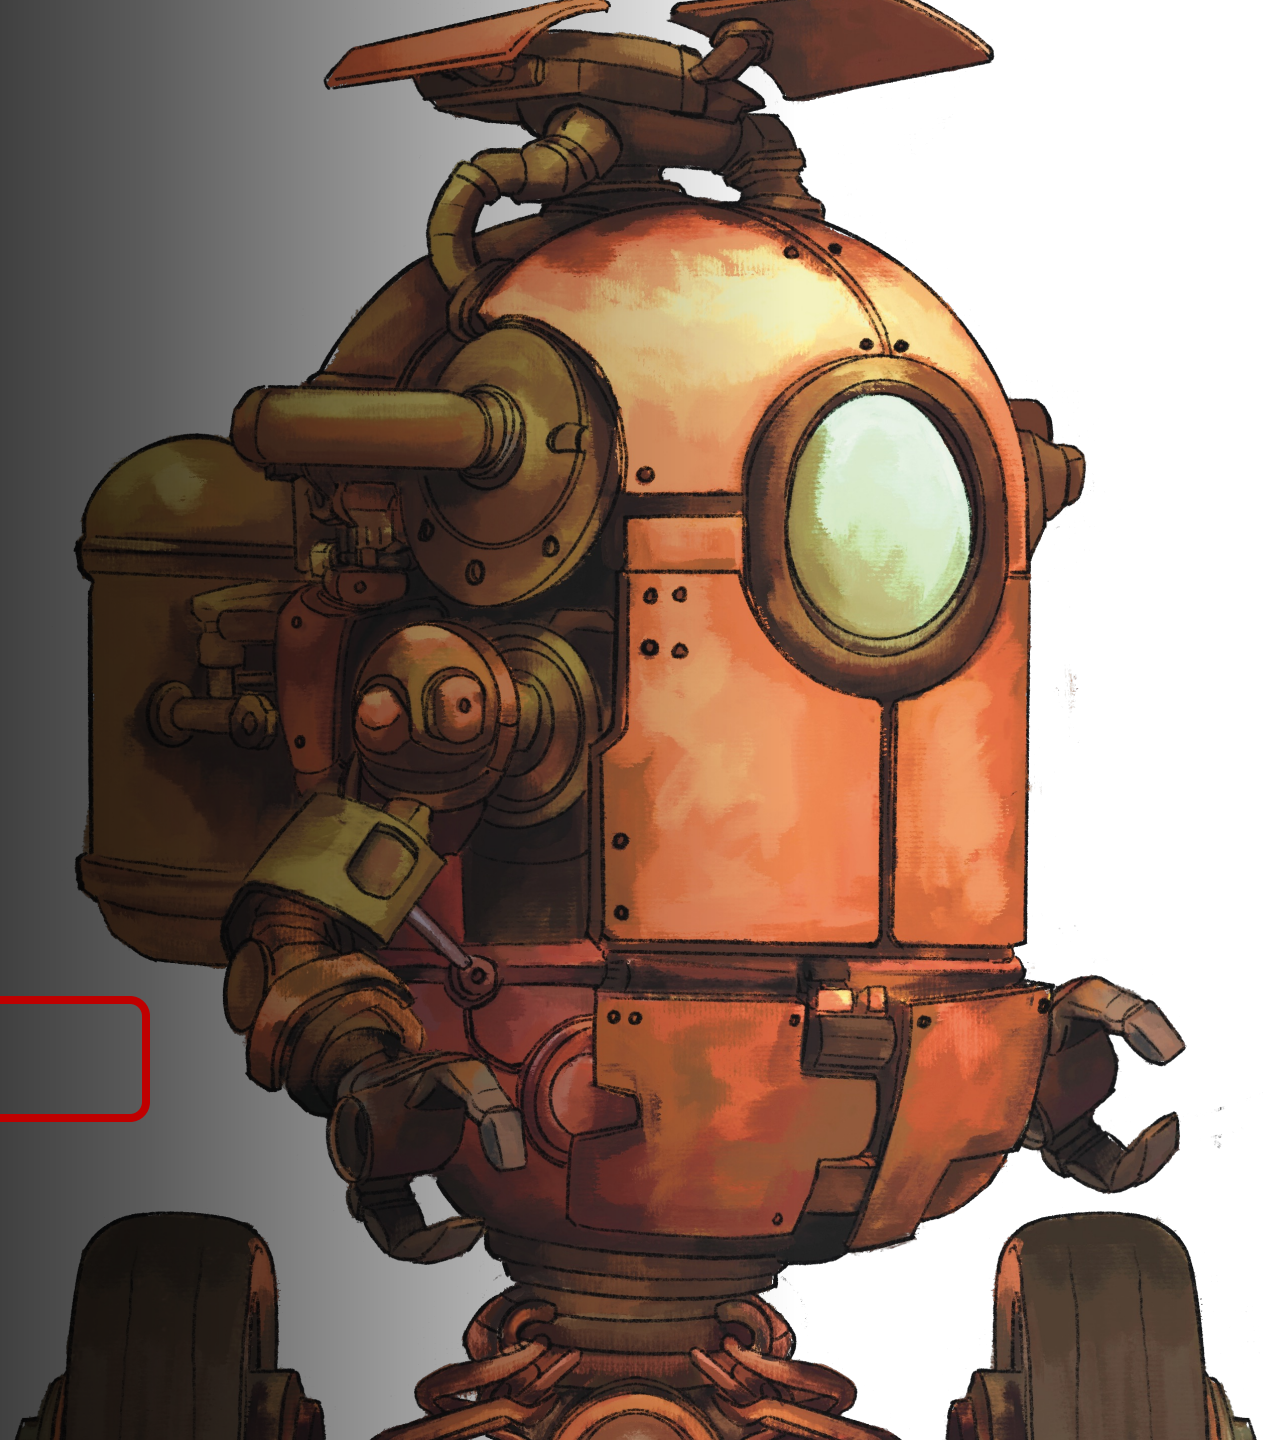
- Passive and non-invasive
- Robot navigation (path planning, obstacle detection)
- 3D modeling (shape analysis, reverse engineering, visualization)
- Photorealistic rendering





Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm

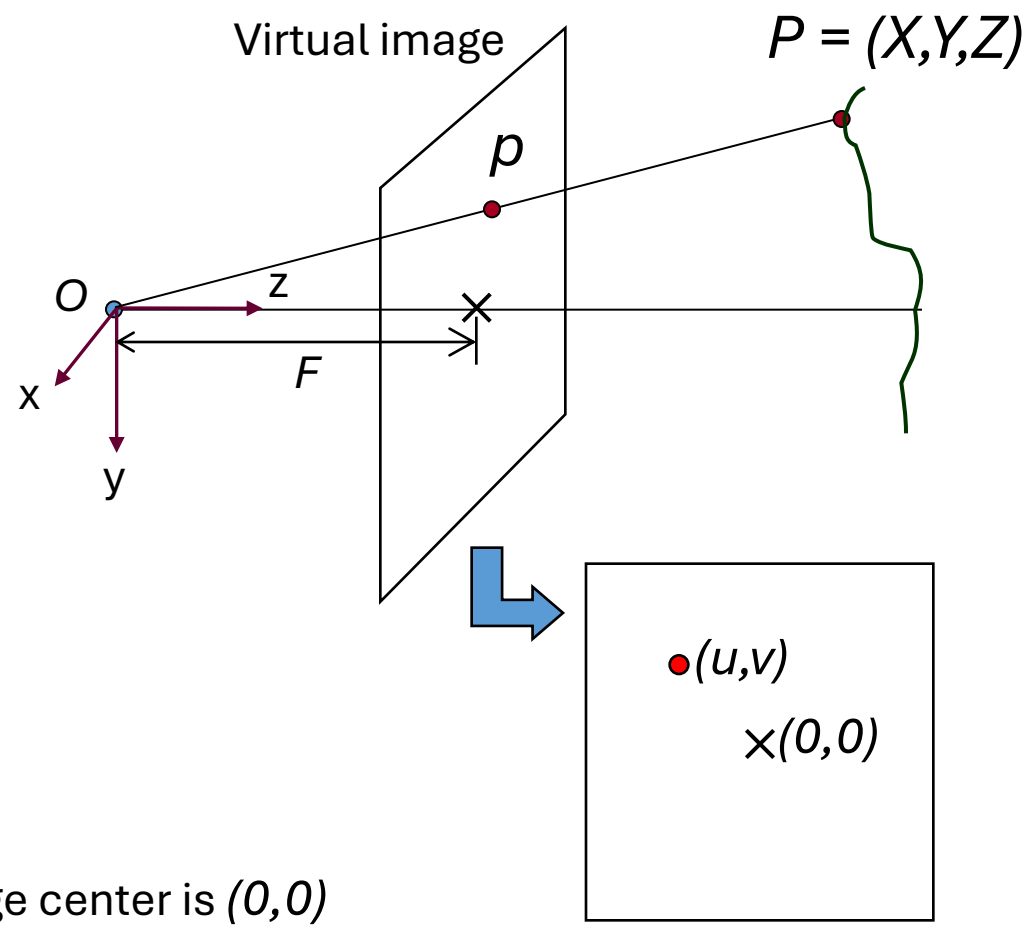


Stereo Geometry

- Recall: Pinhole model
- Now we have two !
- How to recover depth from two measurements?

Review: Pinhole Camera Model

3D scene point *P* is projected to a 2D point *Q* in the virtual image plane

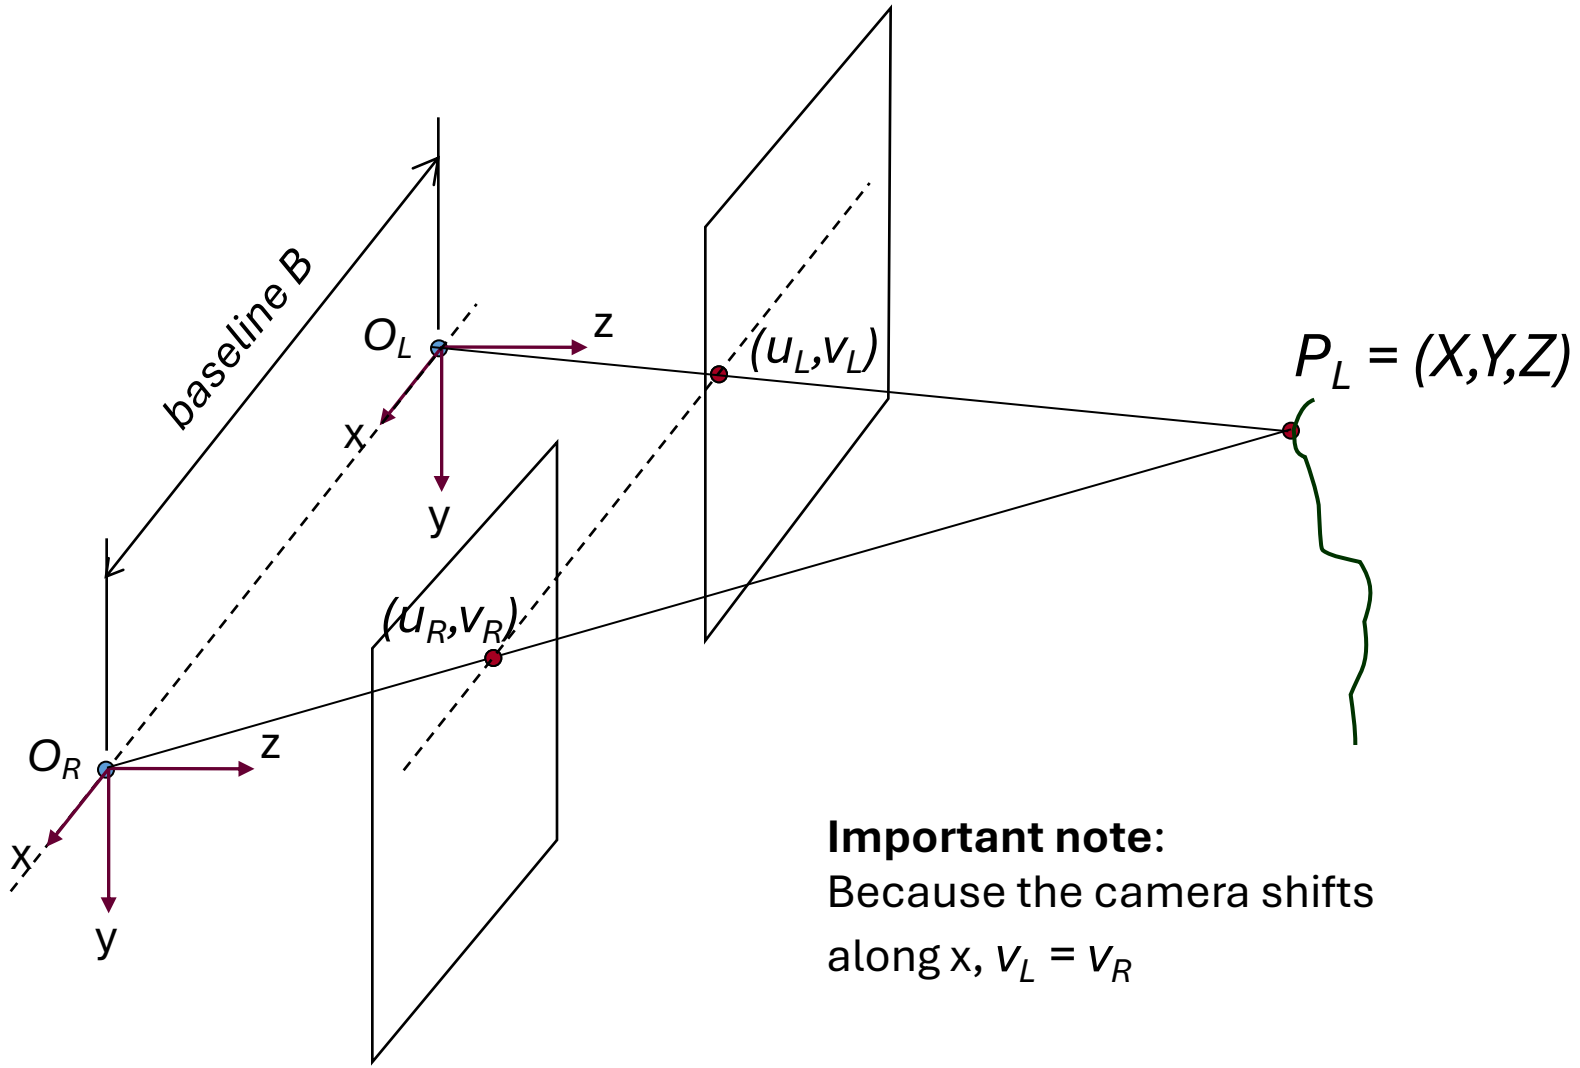


The 2D coordinates in the image are given by

$$(u, v) = \left(f\frac{X}{Z}, f\frac{Y}{Z} \right)$$

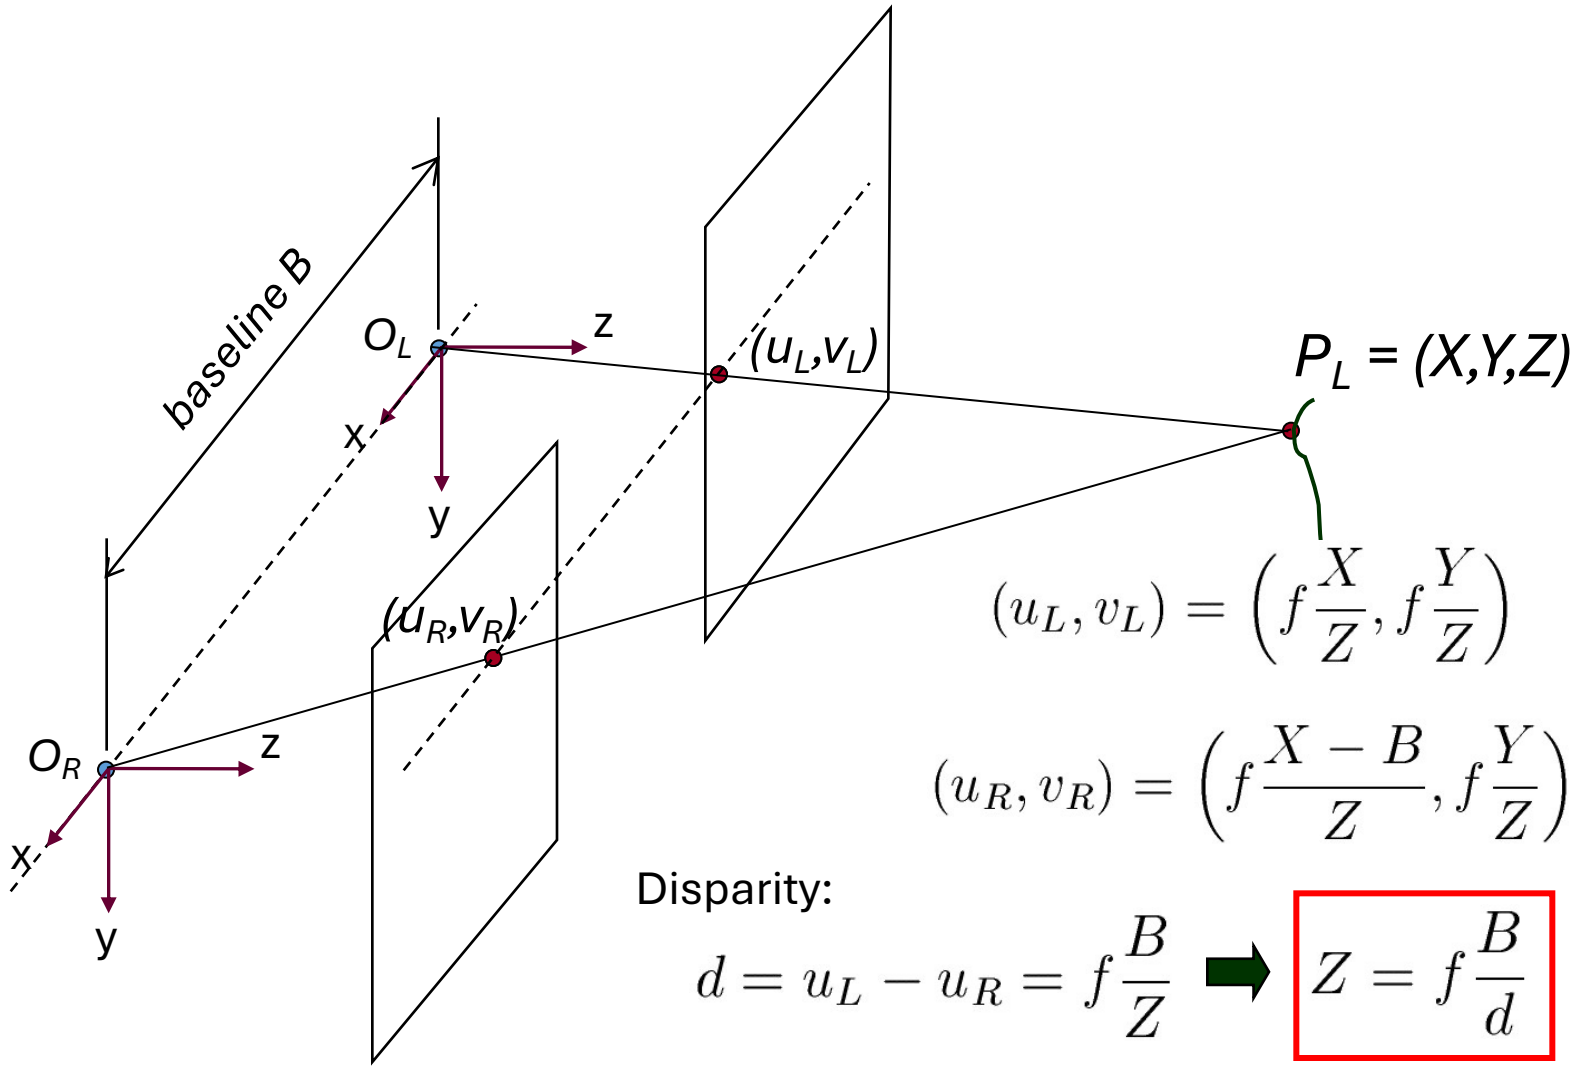
Let's assume for now image center is *(0,0)*

Basic Stereo Derivations



Important note:
Because the camera shifts along x, $v_L = v_R$

Basic Stereo Formula



$$(u_L, v_L) = \left(f\frac{X}{Z}, f\frac{Y}{Z}\right)$$

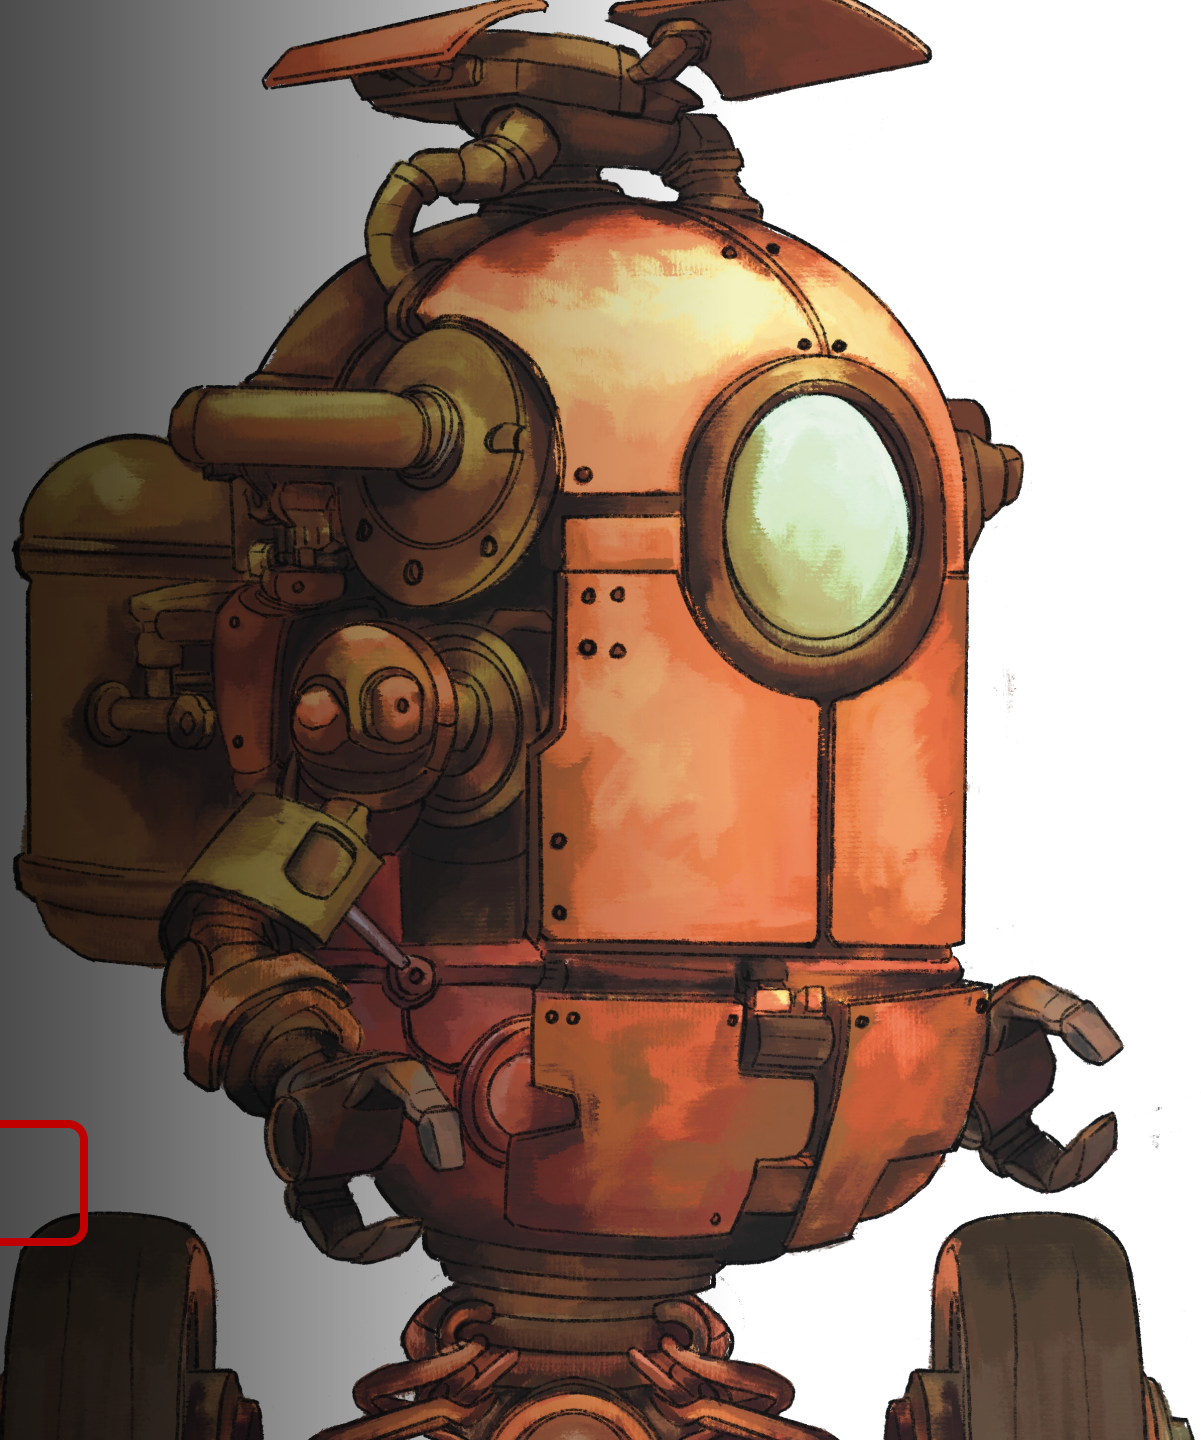
$$(u_R, v_R) = \left(f\frac{X-B}{Z}, f\frac{Y}{Z}\right)$$

Disparity:

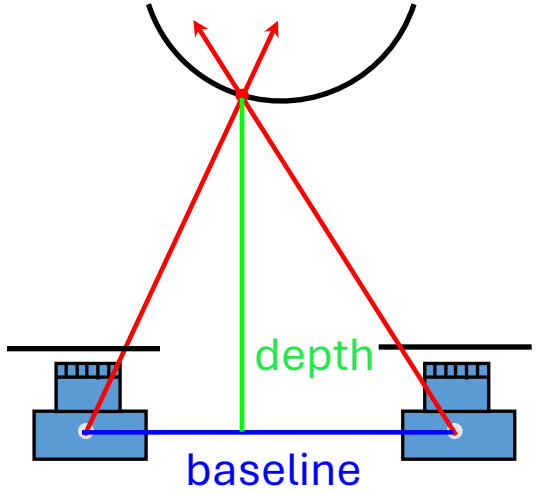
$$d = u_L - u_R = f\frac{B}{Z} \Rightarrow \boxed{Z = f\frac{B}{d}}$$

Outline

1. Differential Drive Recap
2. Perspective Cameras
3. Pinhole Camera Model
4. Properties of Projective Geometry
5. Stereo Vision
6. Stereo Geometry
7. Classical Stereo Algorithm



Classical Stereo Algorithm



$$Z(x,y) = \frac{f\,B}{d(x,y)}$$

$Z(x, y)$ is depth at pixel (x, y)
$d(x, y)$ is disparity

Left



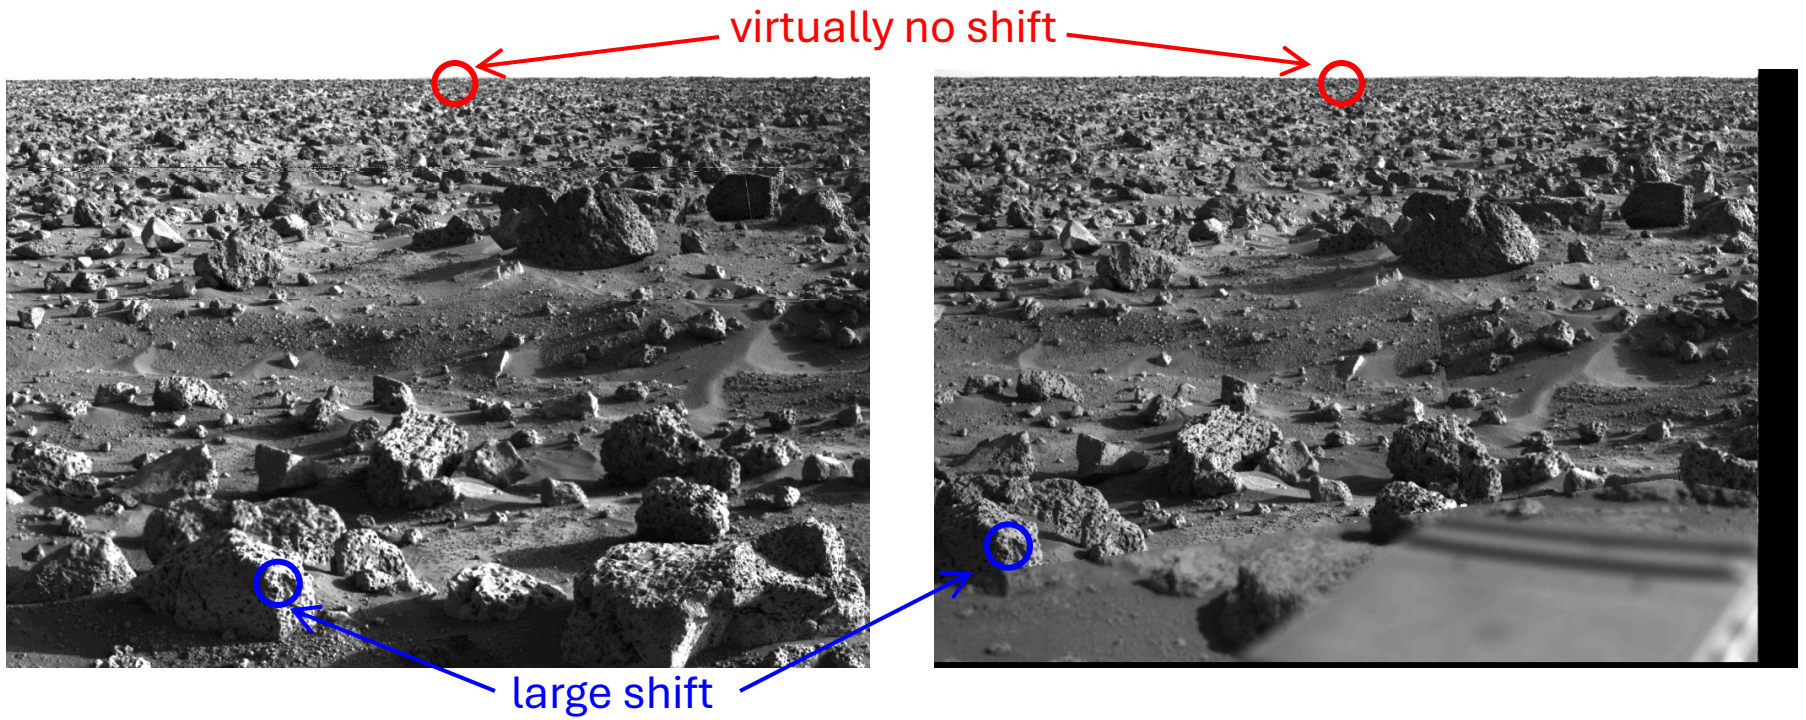
Right



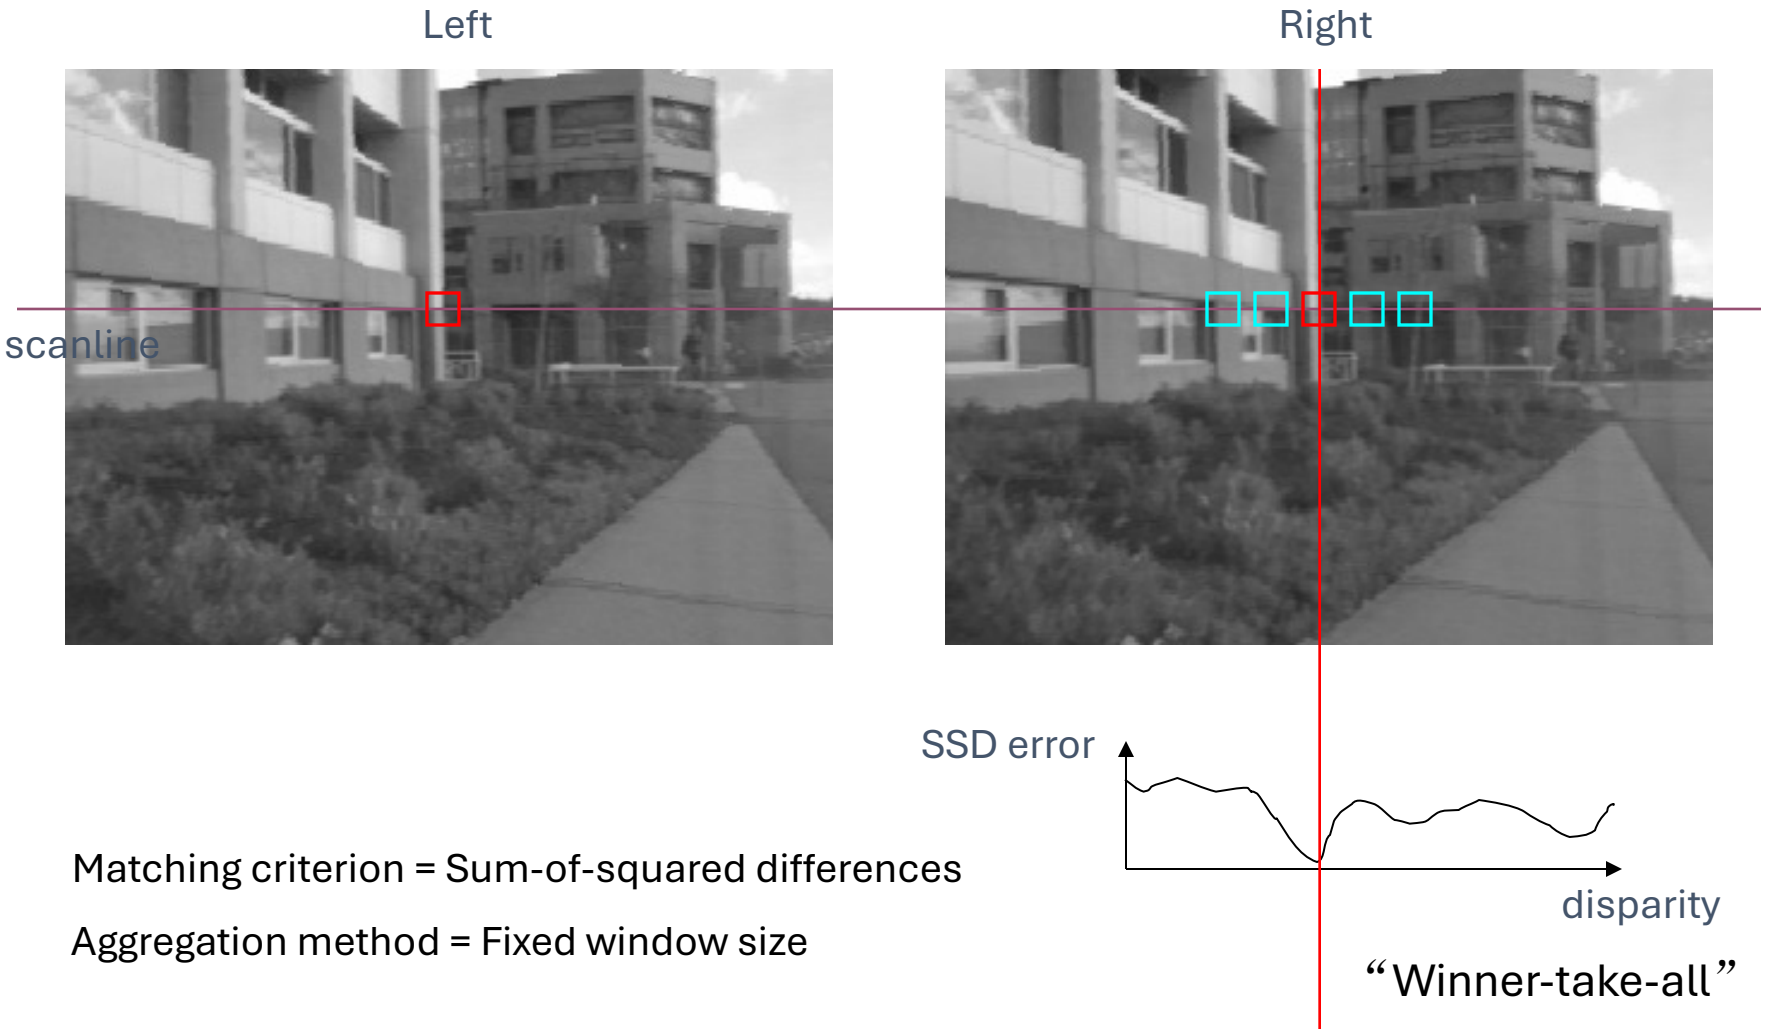
Matching correlation windows across scan lines

Stereo Correspondence

- Search over disparity to find correspondences
- Range of disparities can be large



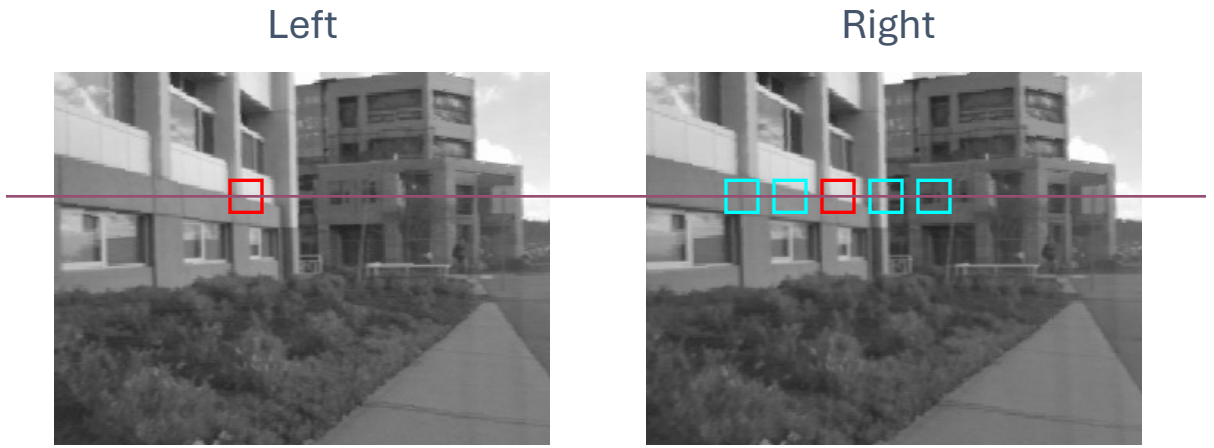
Correspondence Using Window-based Correlation



Matching criterion = Sum-of-squared differences

Aggregation method = Fixed window size

Sum of Squared (Intensity) Differences



w_L and w_R are corresponding m by m windows of pixels.

We define the window function:

$$W_m(x,y) = \{u,v \mid x - \tfrac{m}{2} \le u \le x + \tfrac{m}{2}, y - \tfrac{m}{2} \le v \le y + \tfrac{m}{2}\}$$

The SSD cost measures the intensity difference as a function of disparity:

$$C_r(x,y,d) = \sum_{(u,v)\in W_m(x,y)} [I_L(u,v) - I_R(u-d,v)]^2$$

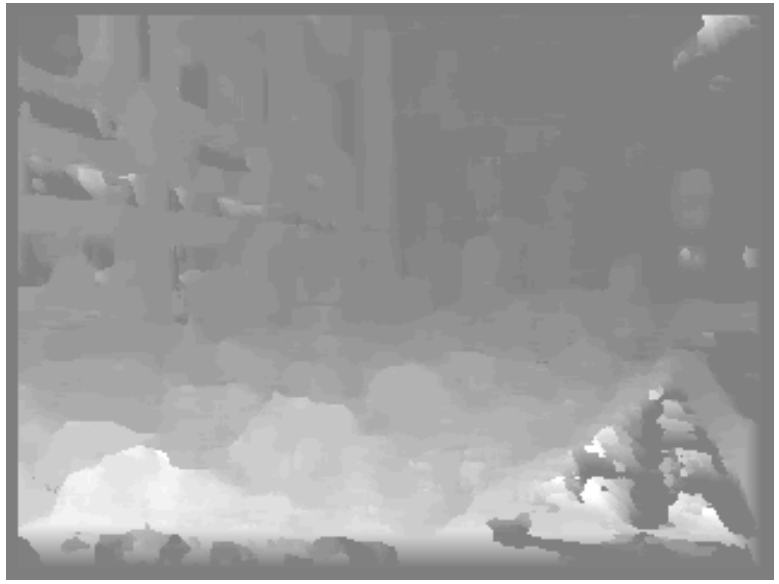
Correspondence Using Correlation



Left



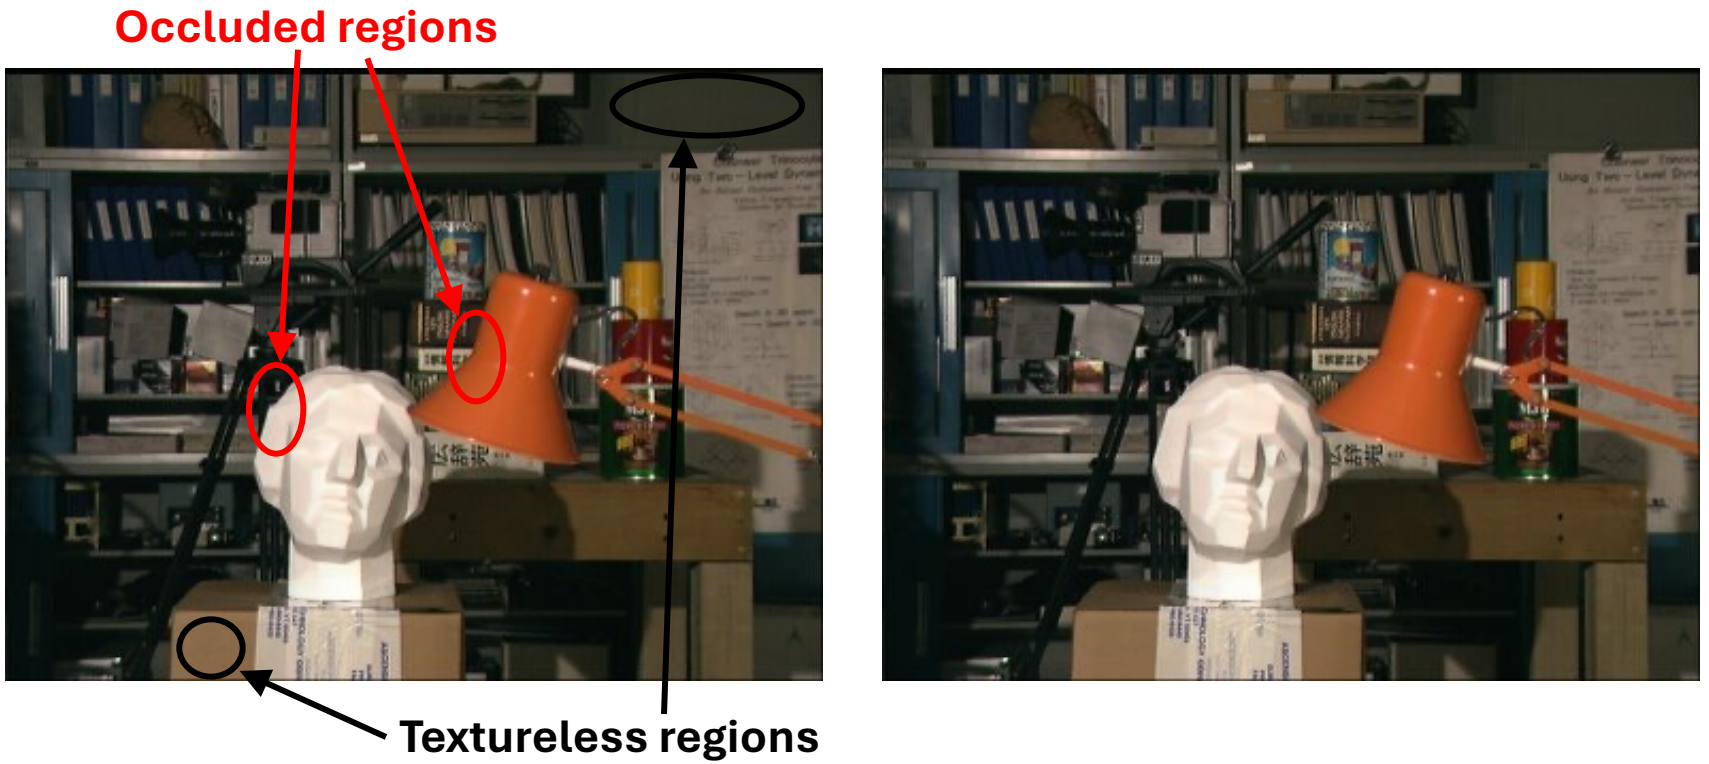
Disparity Map



Images courtesy of Point Grey Research

Two major roadblocks

- Textureless regions create ambiguities
- Occlusions result in missing data



Stereo is Hard! Learning solves many of those issues

- Depth discontinuities
- Lack of texture (depth ambiguity)
- Non-rigid effects (highlights, reflection, translucency)





Summary

- Perspective Cameras Intro
- Pinhole Camera Model defined
- Properties of Projective Geometry
- Stereo Vision can recover metric structure
- Stereo Geometry is simply Z = f B/d
- Amazing Stereo Algorithms were elusive, BUT, deep learning!!!!

