

# Lecture 16: Image formation

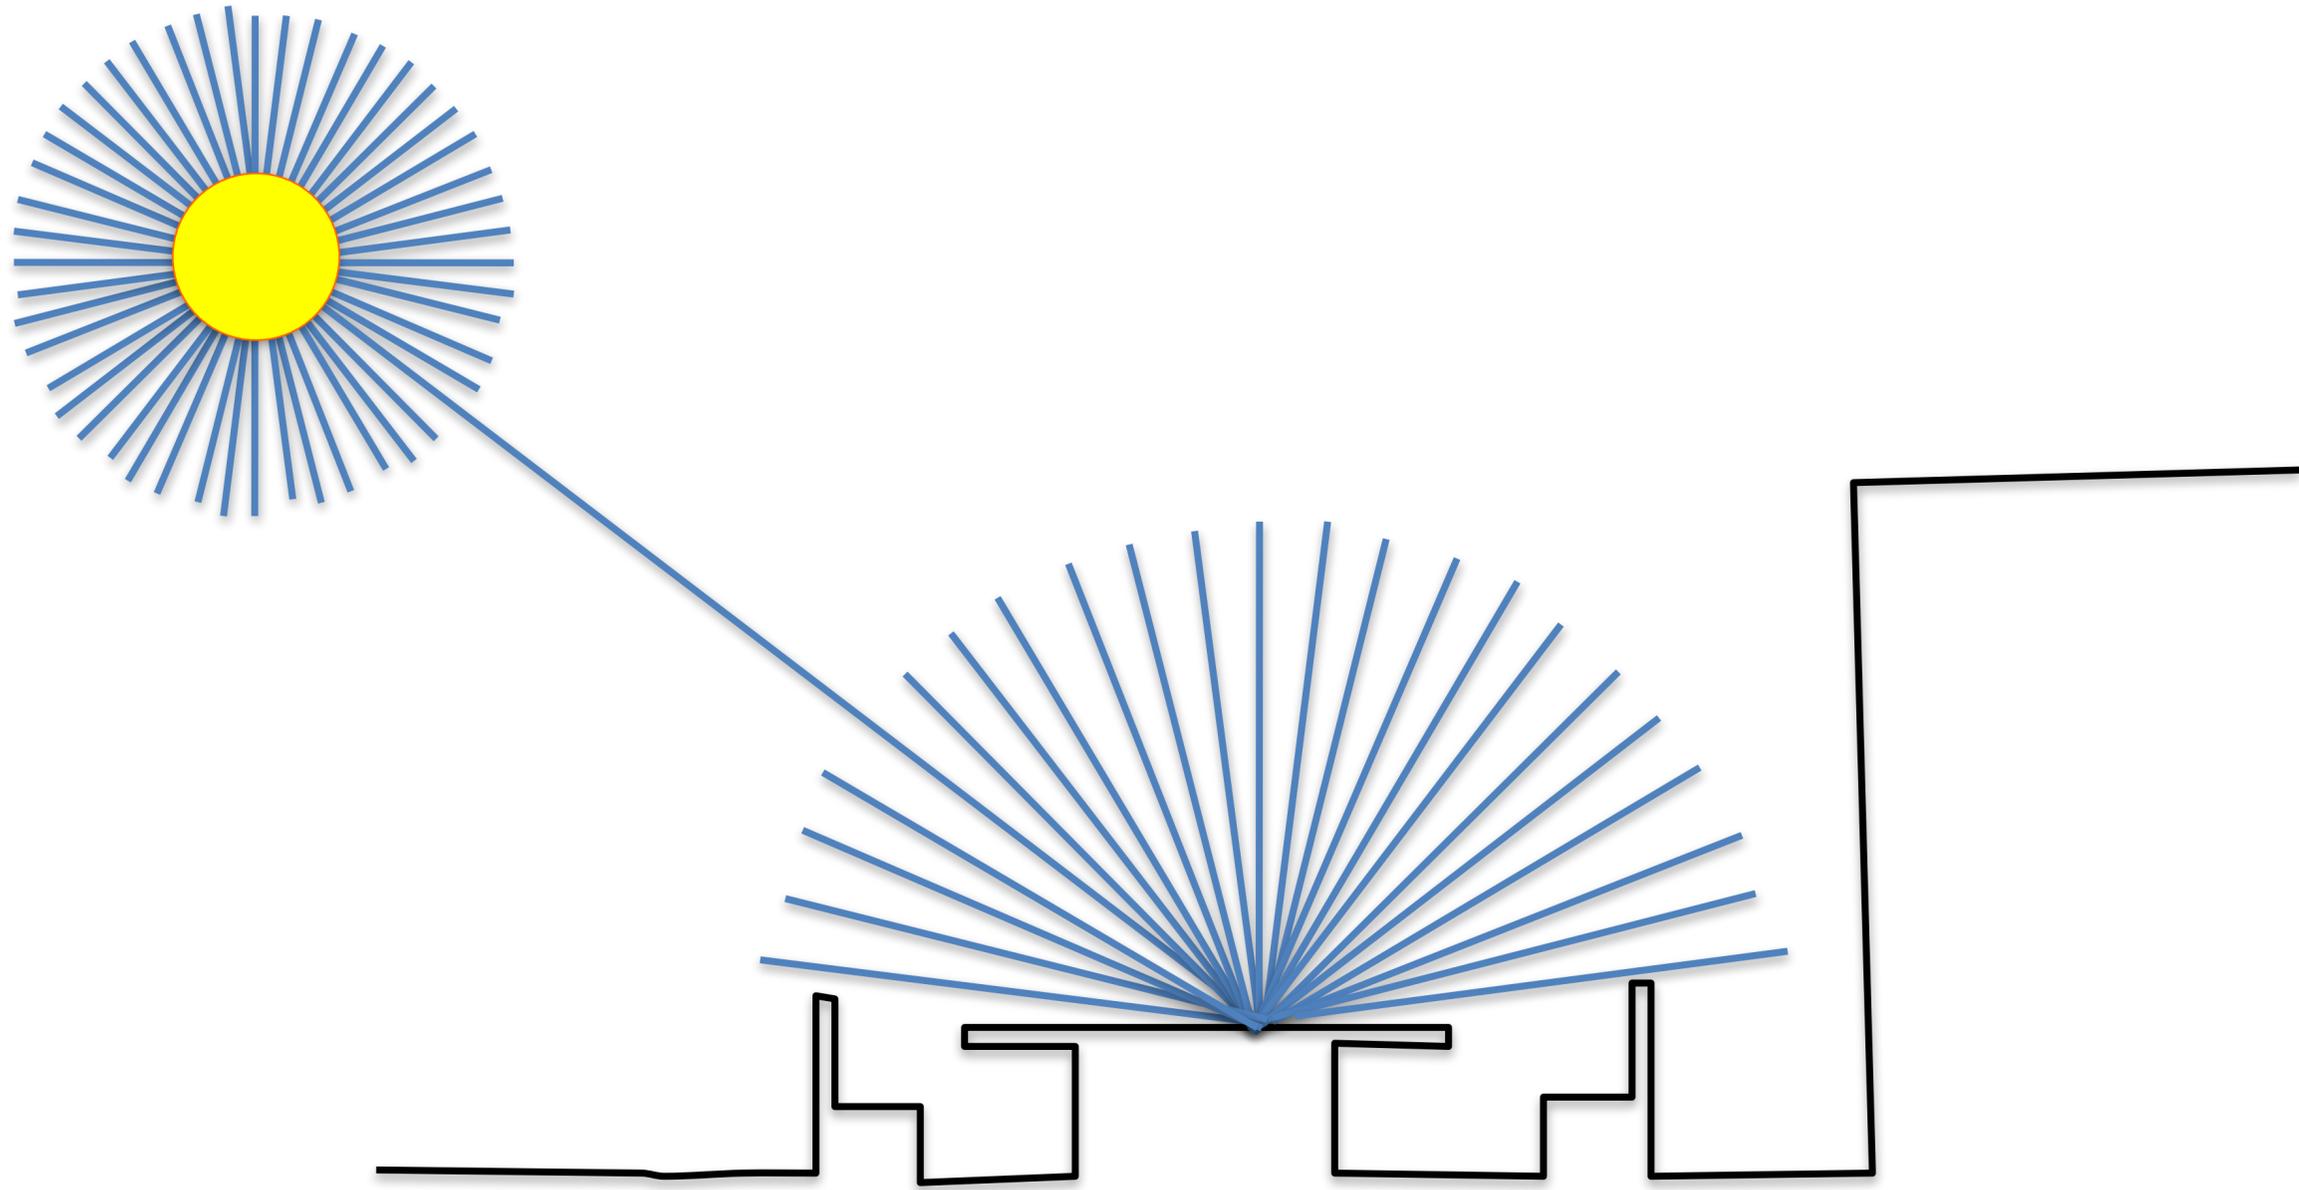
# Announcements

- Reminder: PS2 grades out
- Next problem sets:
  - PS7: representation learning
  - PS8: panorama stitching

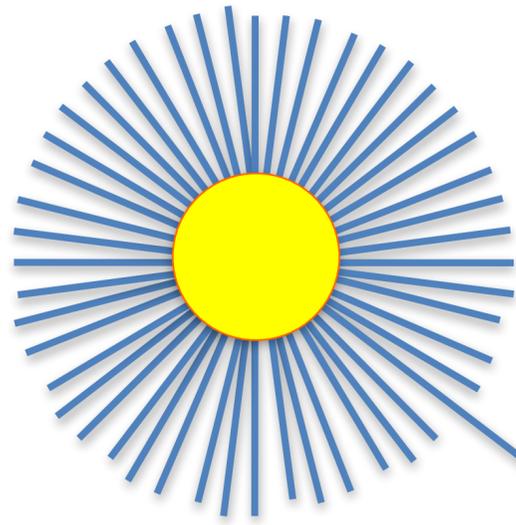
# Today

- Camera models
- Projection equations

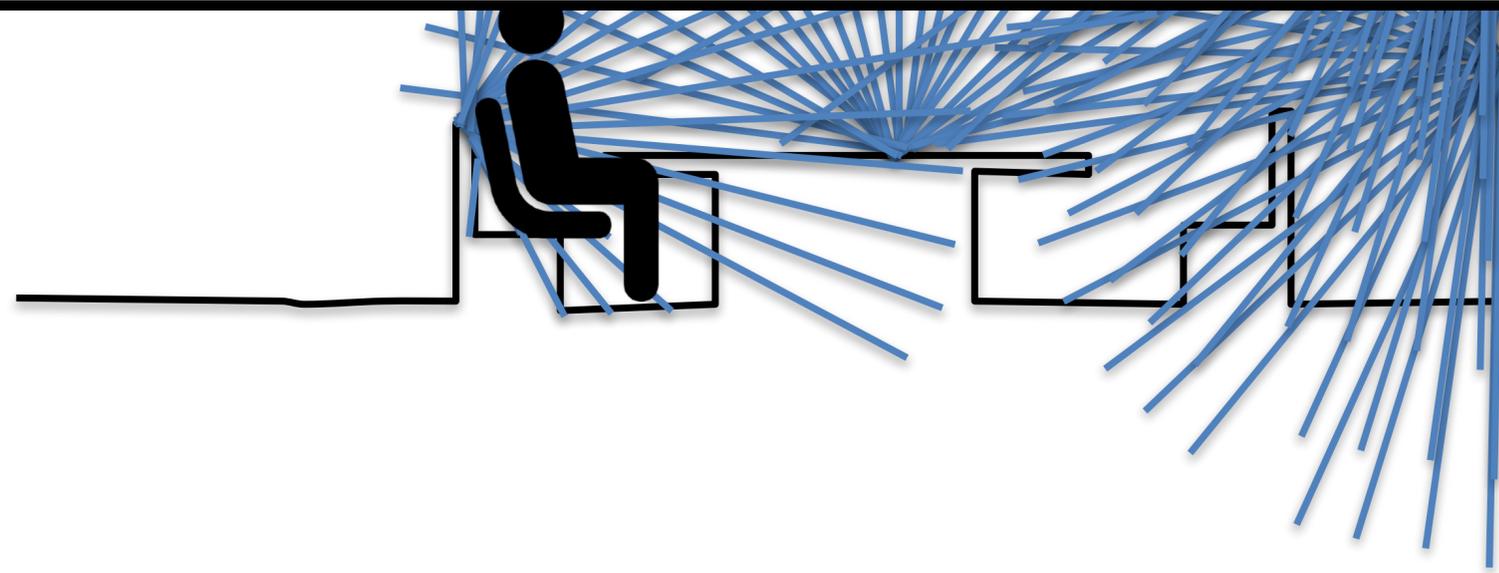
# The structure of ambient light



# The structure of ambient light

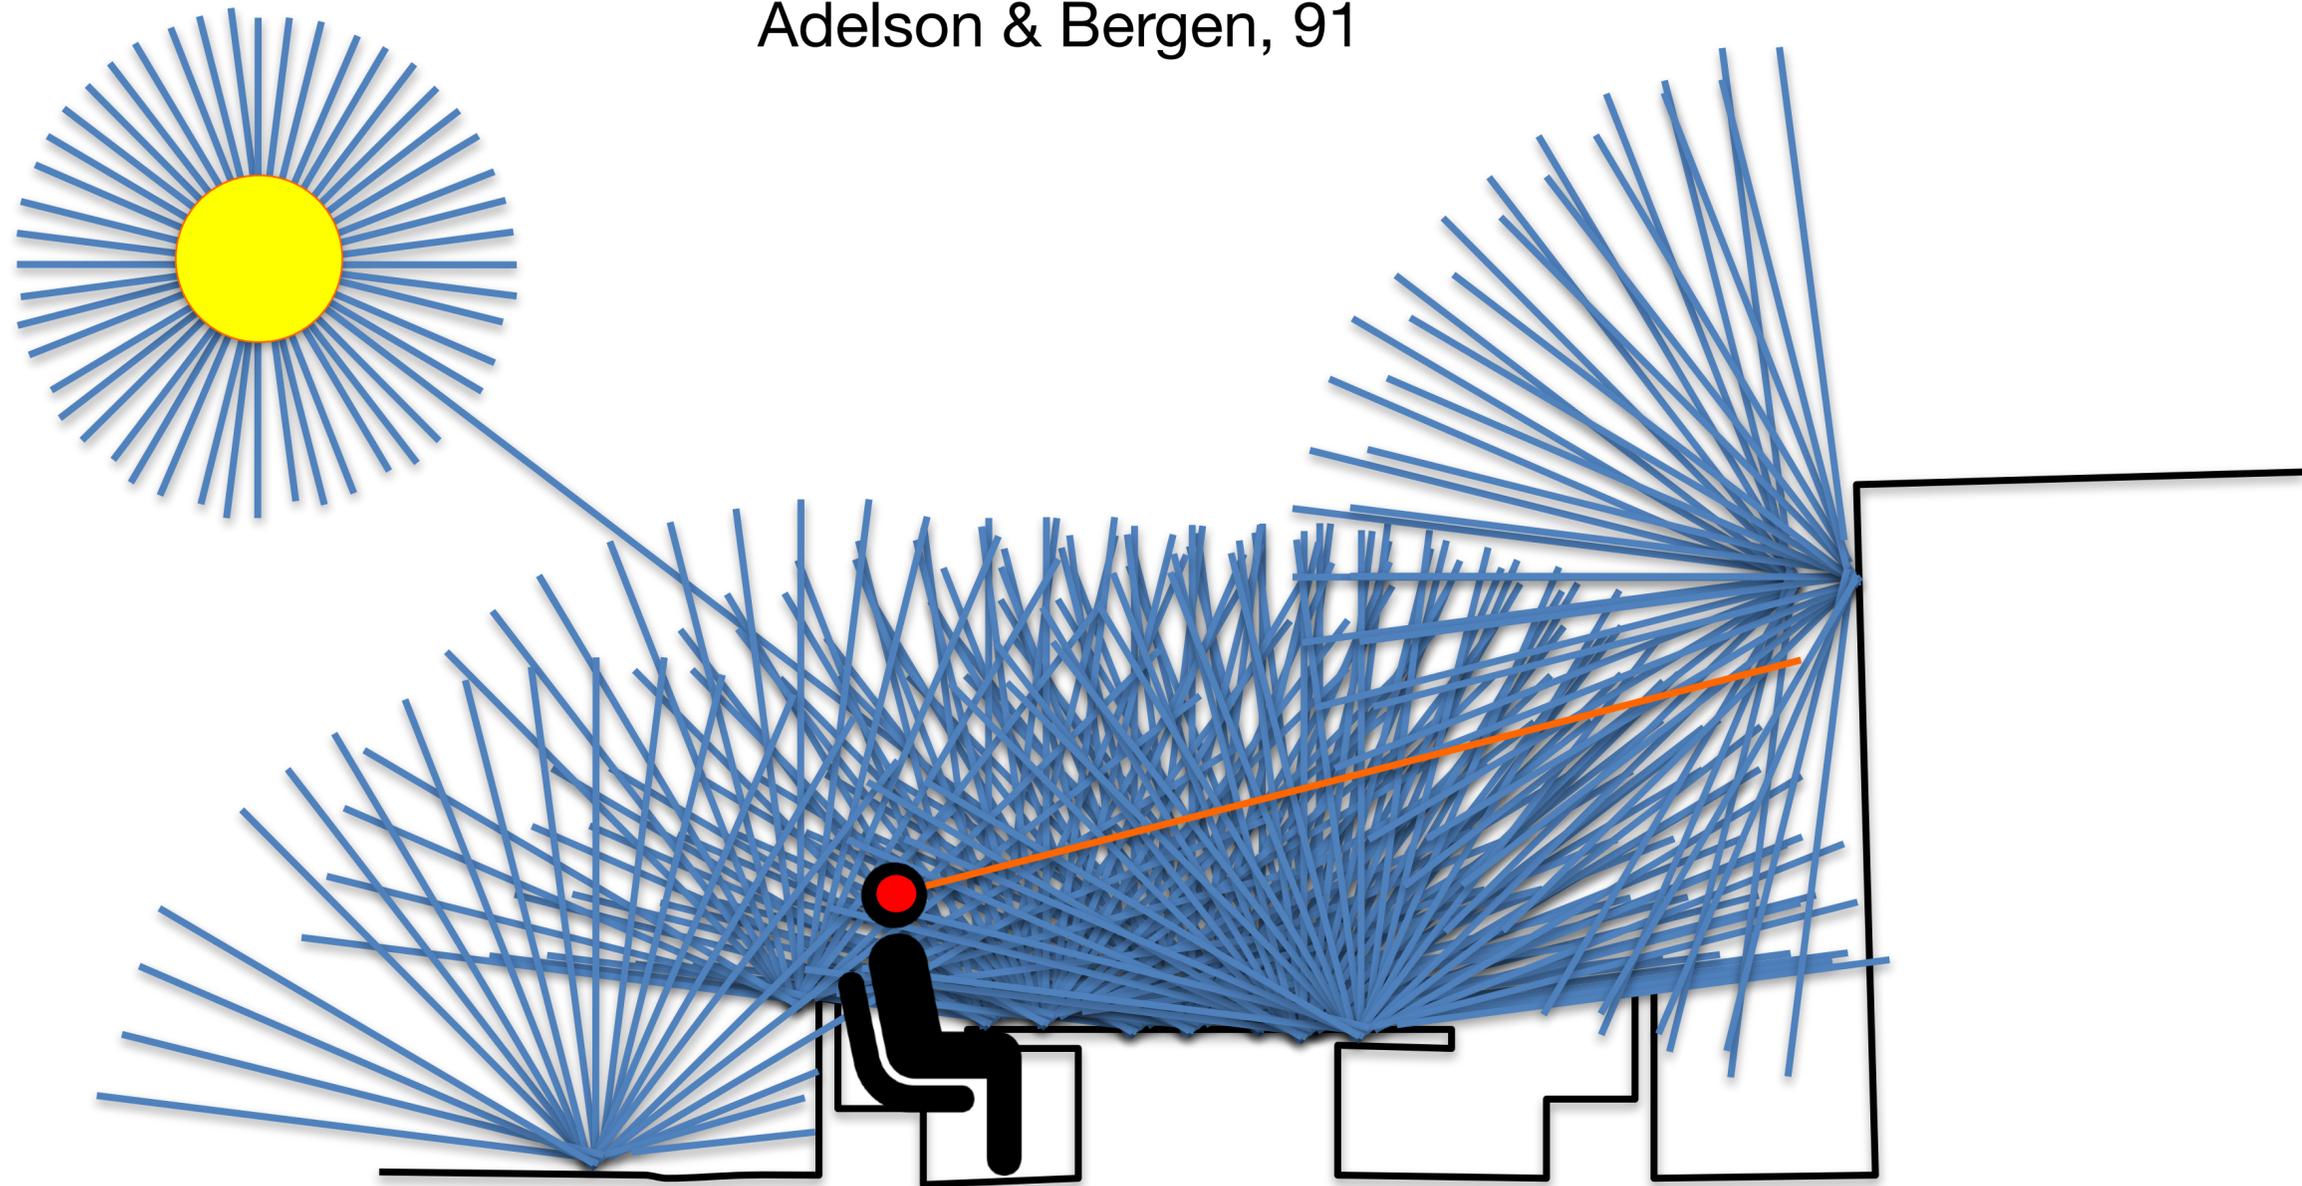


What information does this light provide?



# The Plenoptic Function

Adelson & Bergen, 91

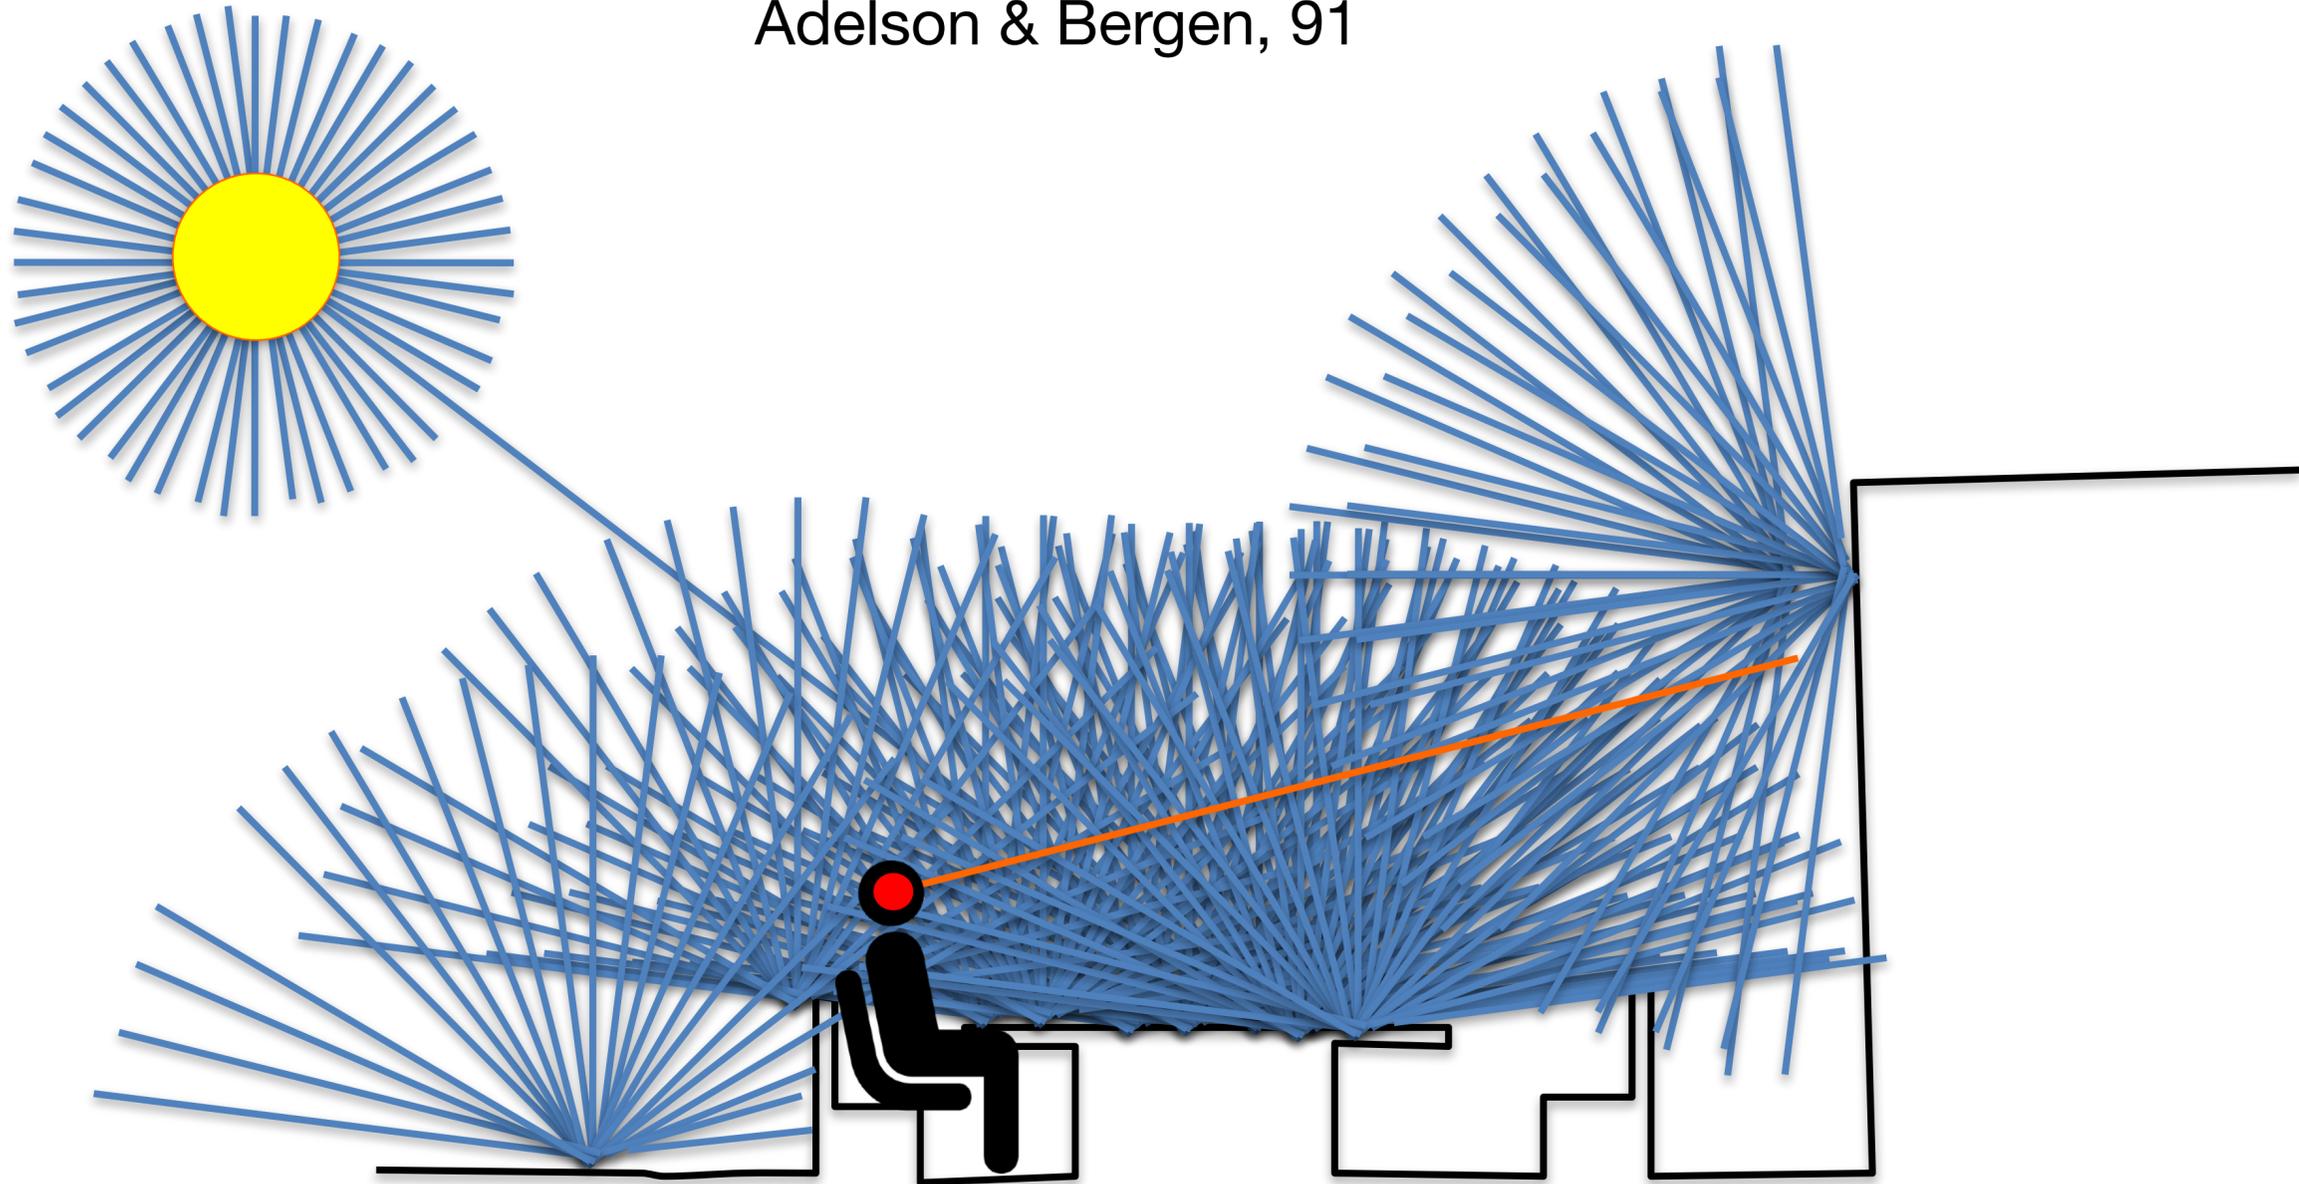


The intensity  $P$  can be parameterized as:

$$P(\text{Eye position}, X, Y, Z)$$

# The Plenoptic Function

Adelson & Bergen, 91



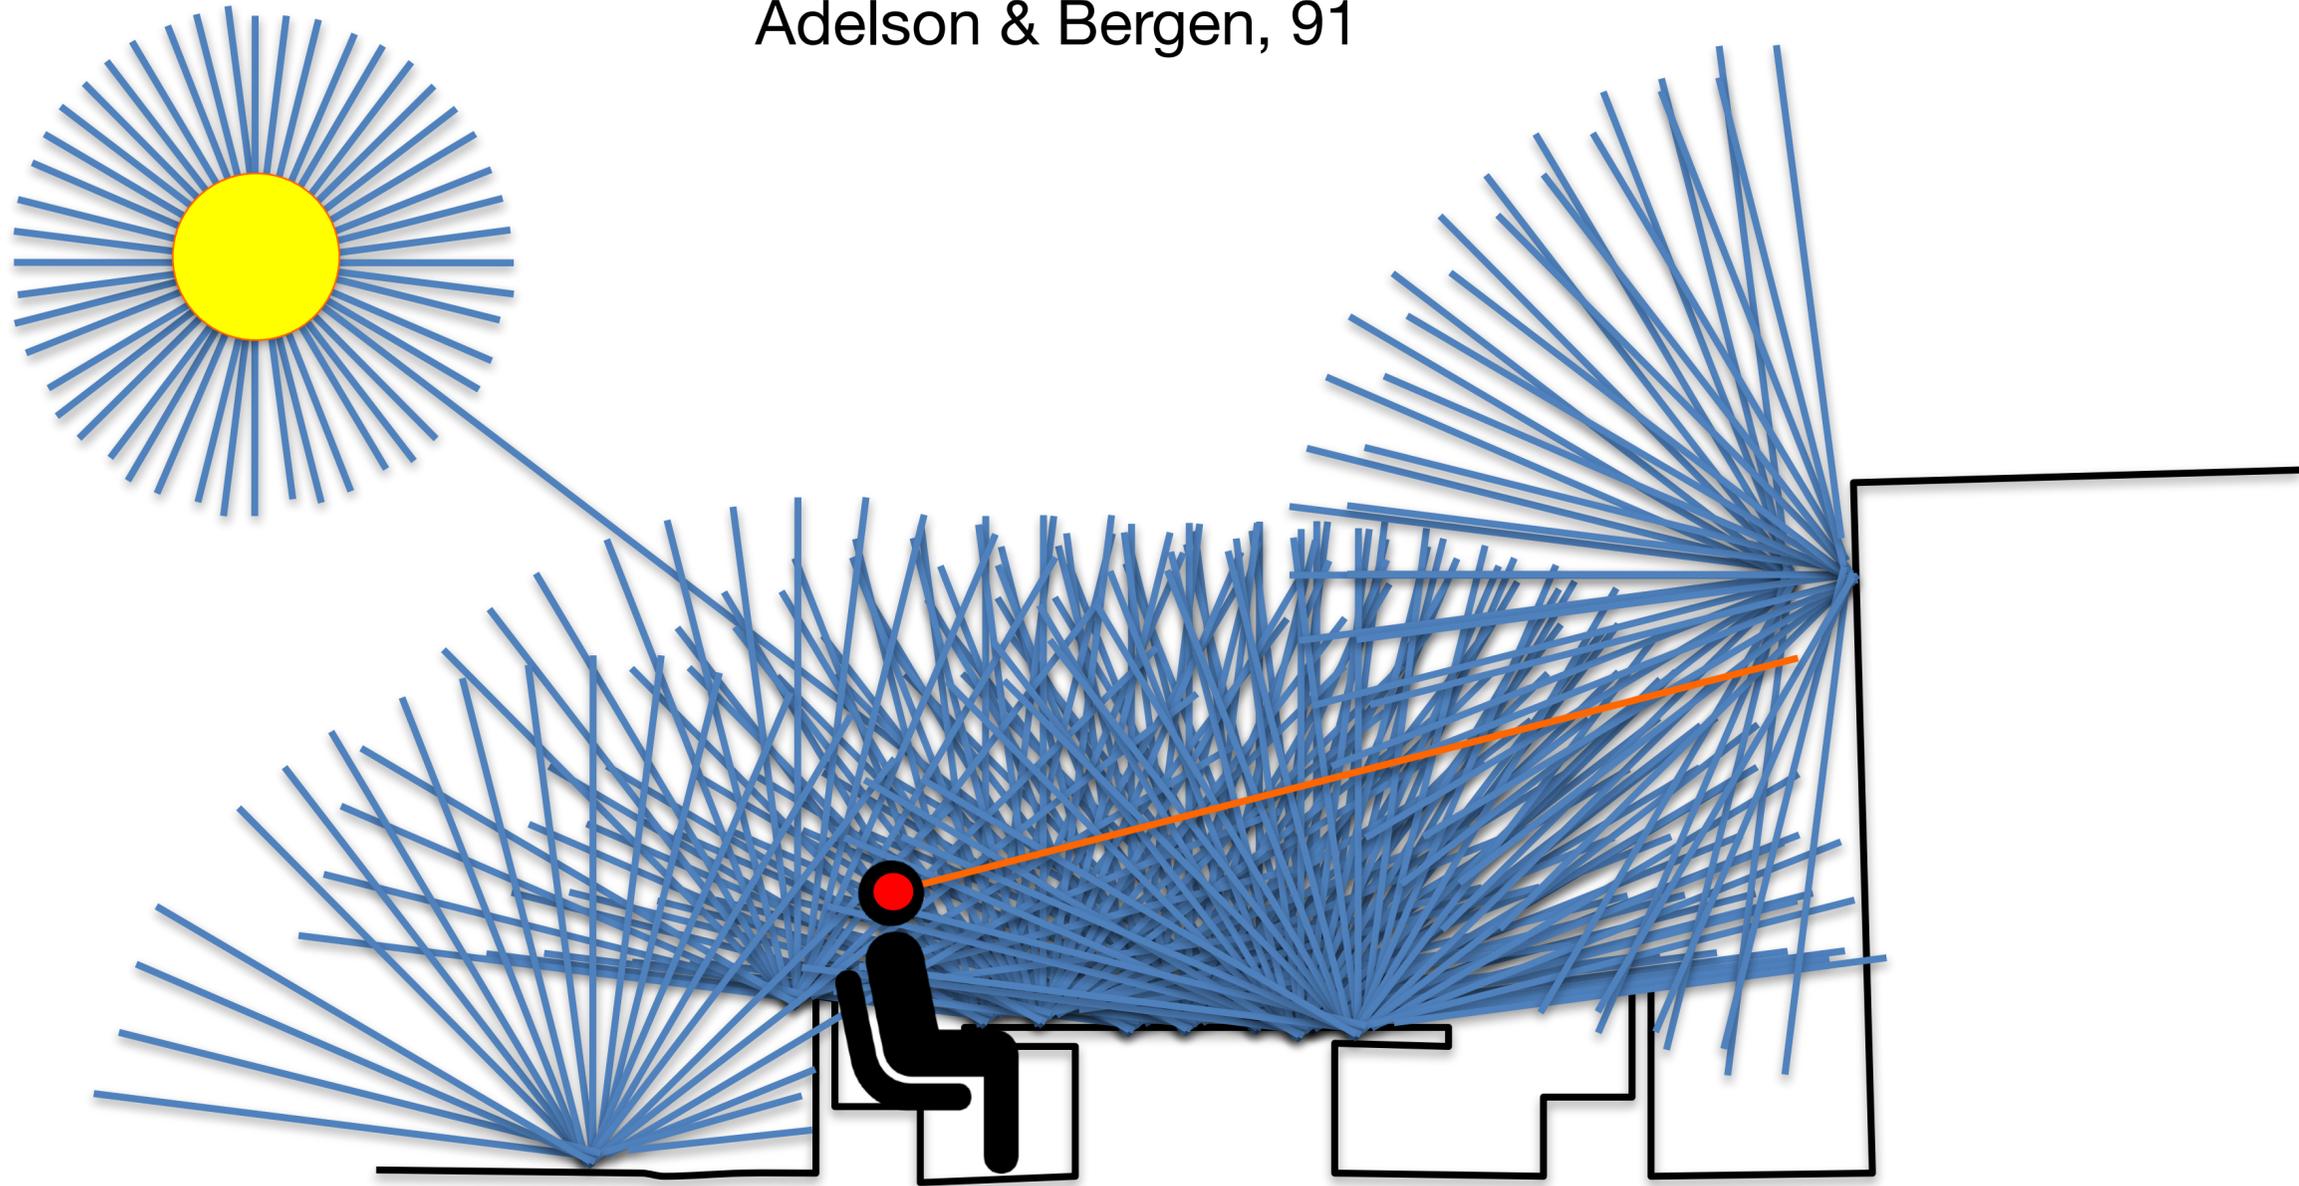
The intensity  $P$  can be parameterized as:

$$P(\theta, \phi, X, Y, Z)$$

Angle

# The Plenoptic Function

Adelson & Bergen, 91



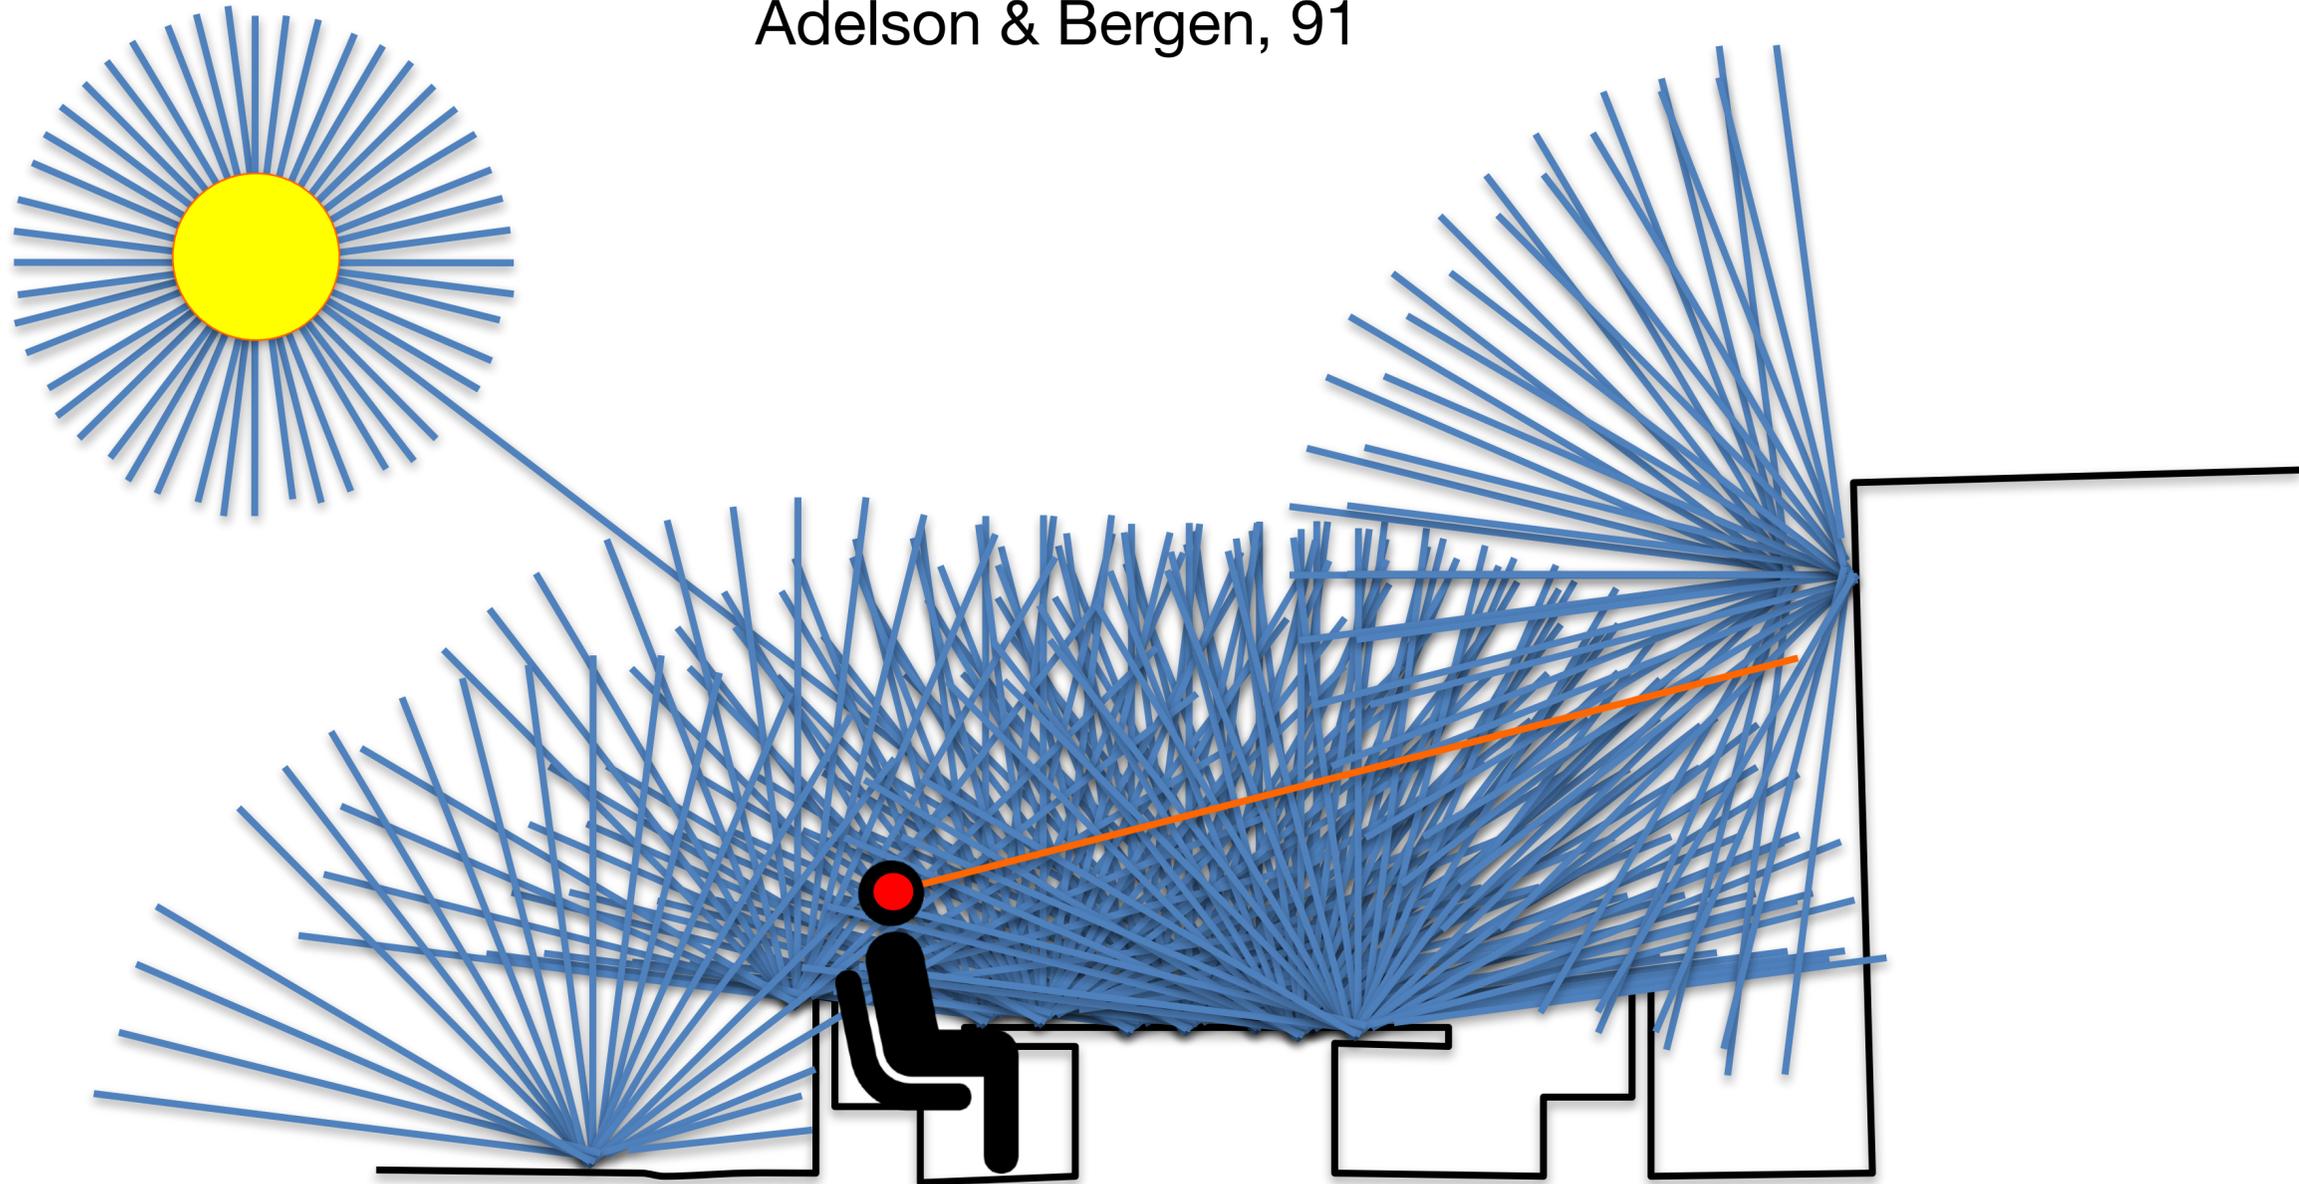
The intensity  $P$  can be parameterized as:

$$P(\theta, \phi, \lambda, t, X, Y, Z)$$

Wavelength, time

# The Plenoptic Function

Adelson & Bergen, 91

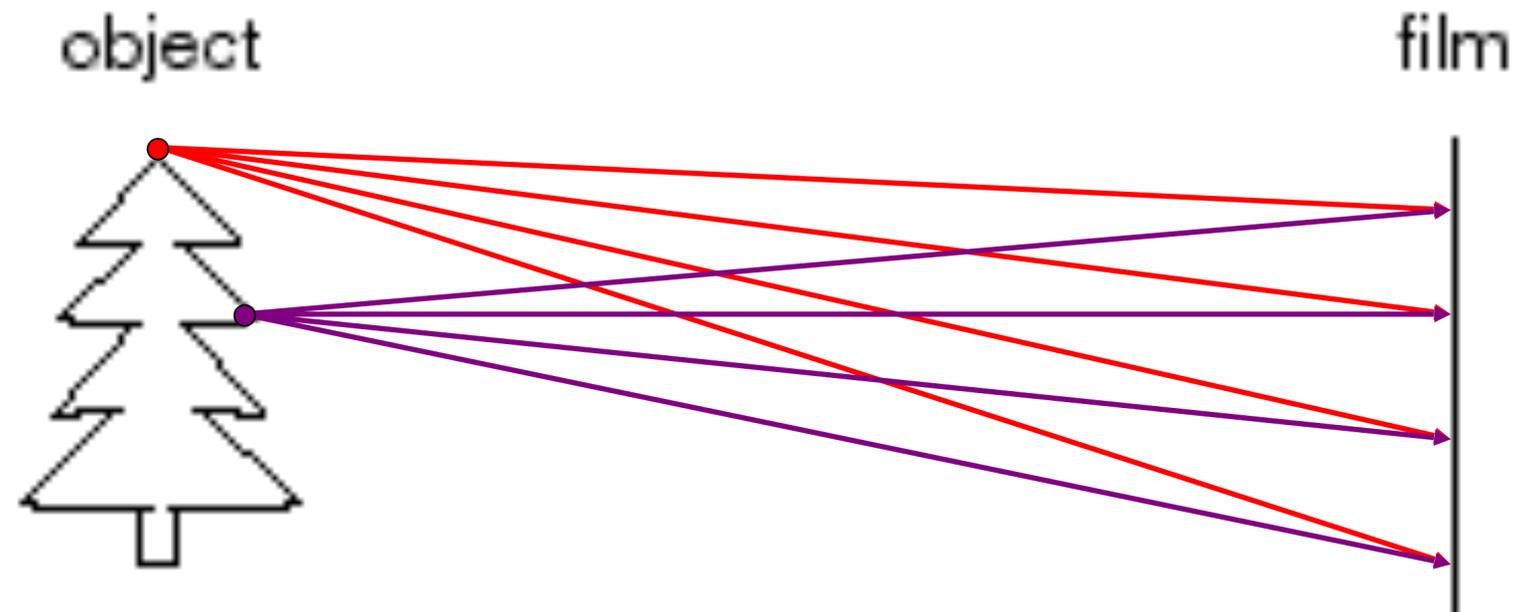


The intensity  $P$  can be parameterized as:

$$P(\theta, \phi, \lambda, t, X, Y, Z)$$

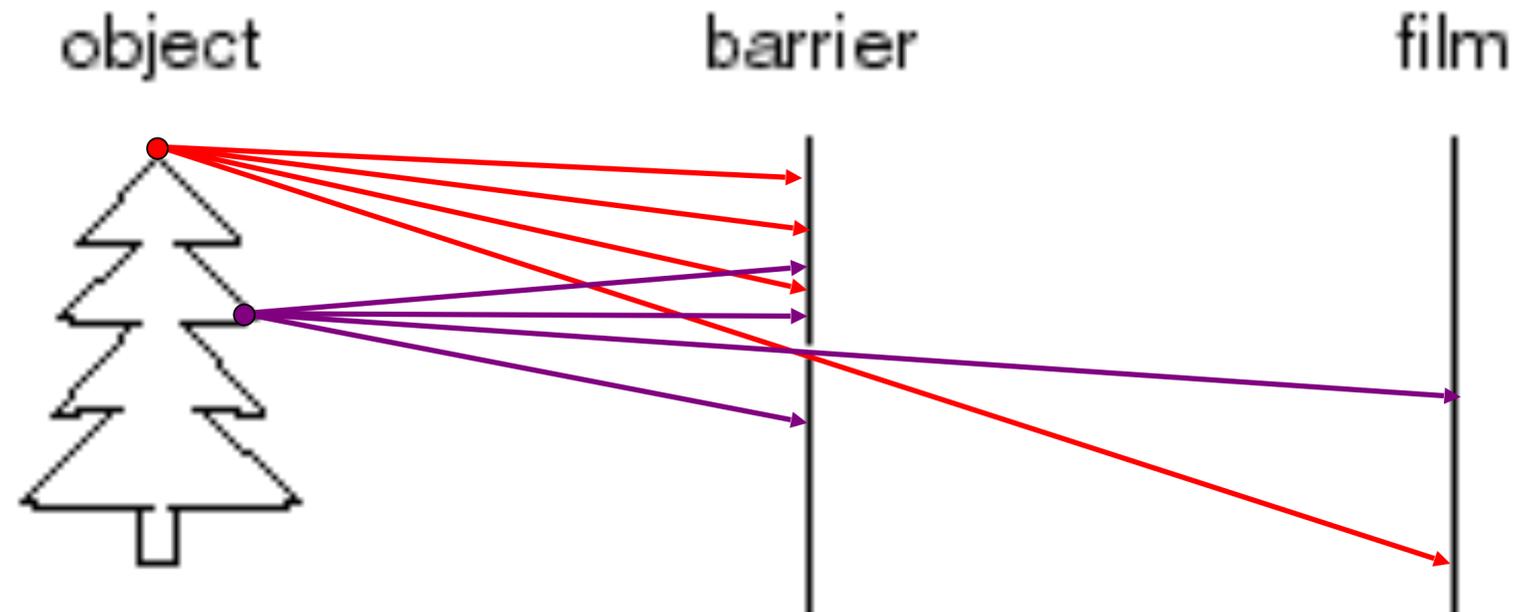
Full plenoptic function

# Making a camera



Idea #1: put a piece of film in front of an object.

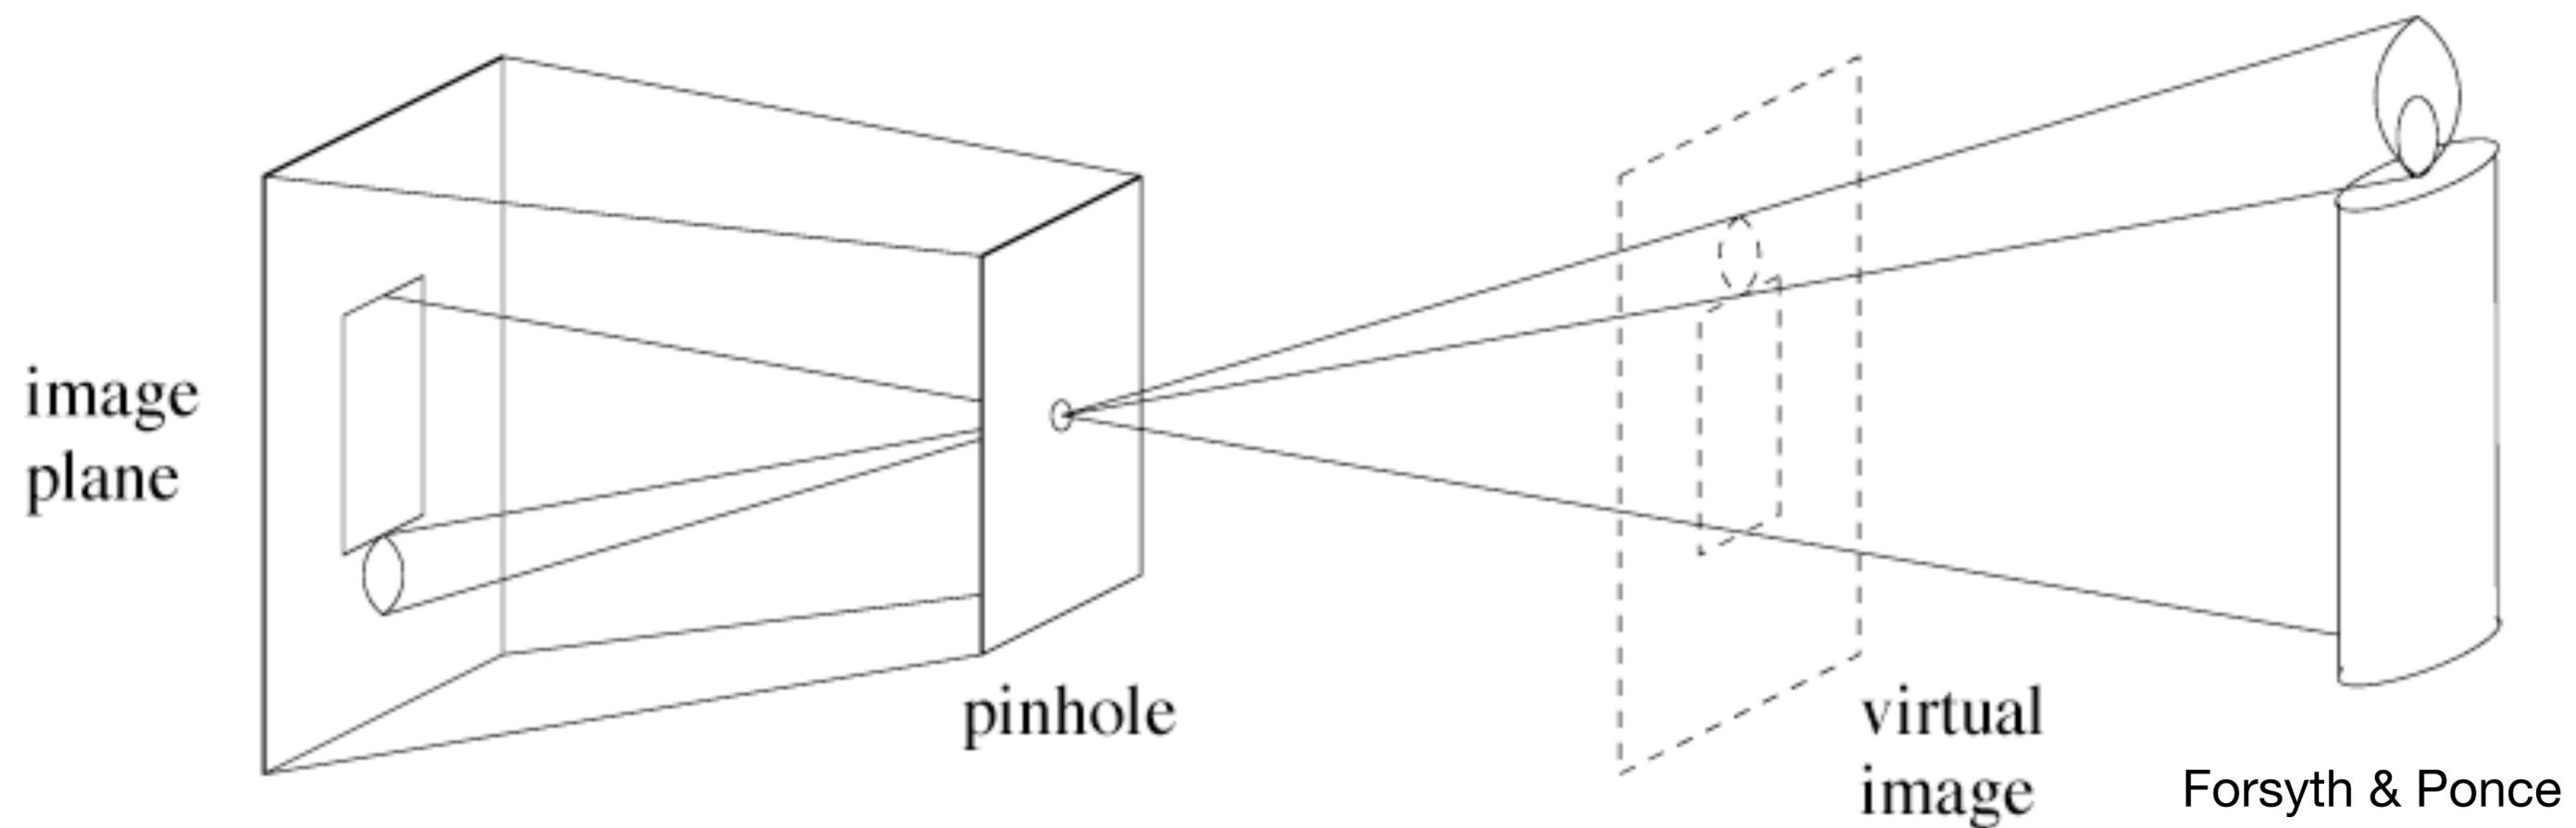
# Pinhole camera



Add a barrier to block off most of the rays

- This reduces blurring
- The opening known as the **aperture**

# Upside down images



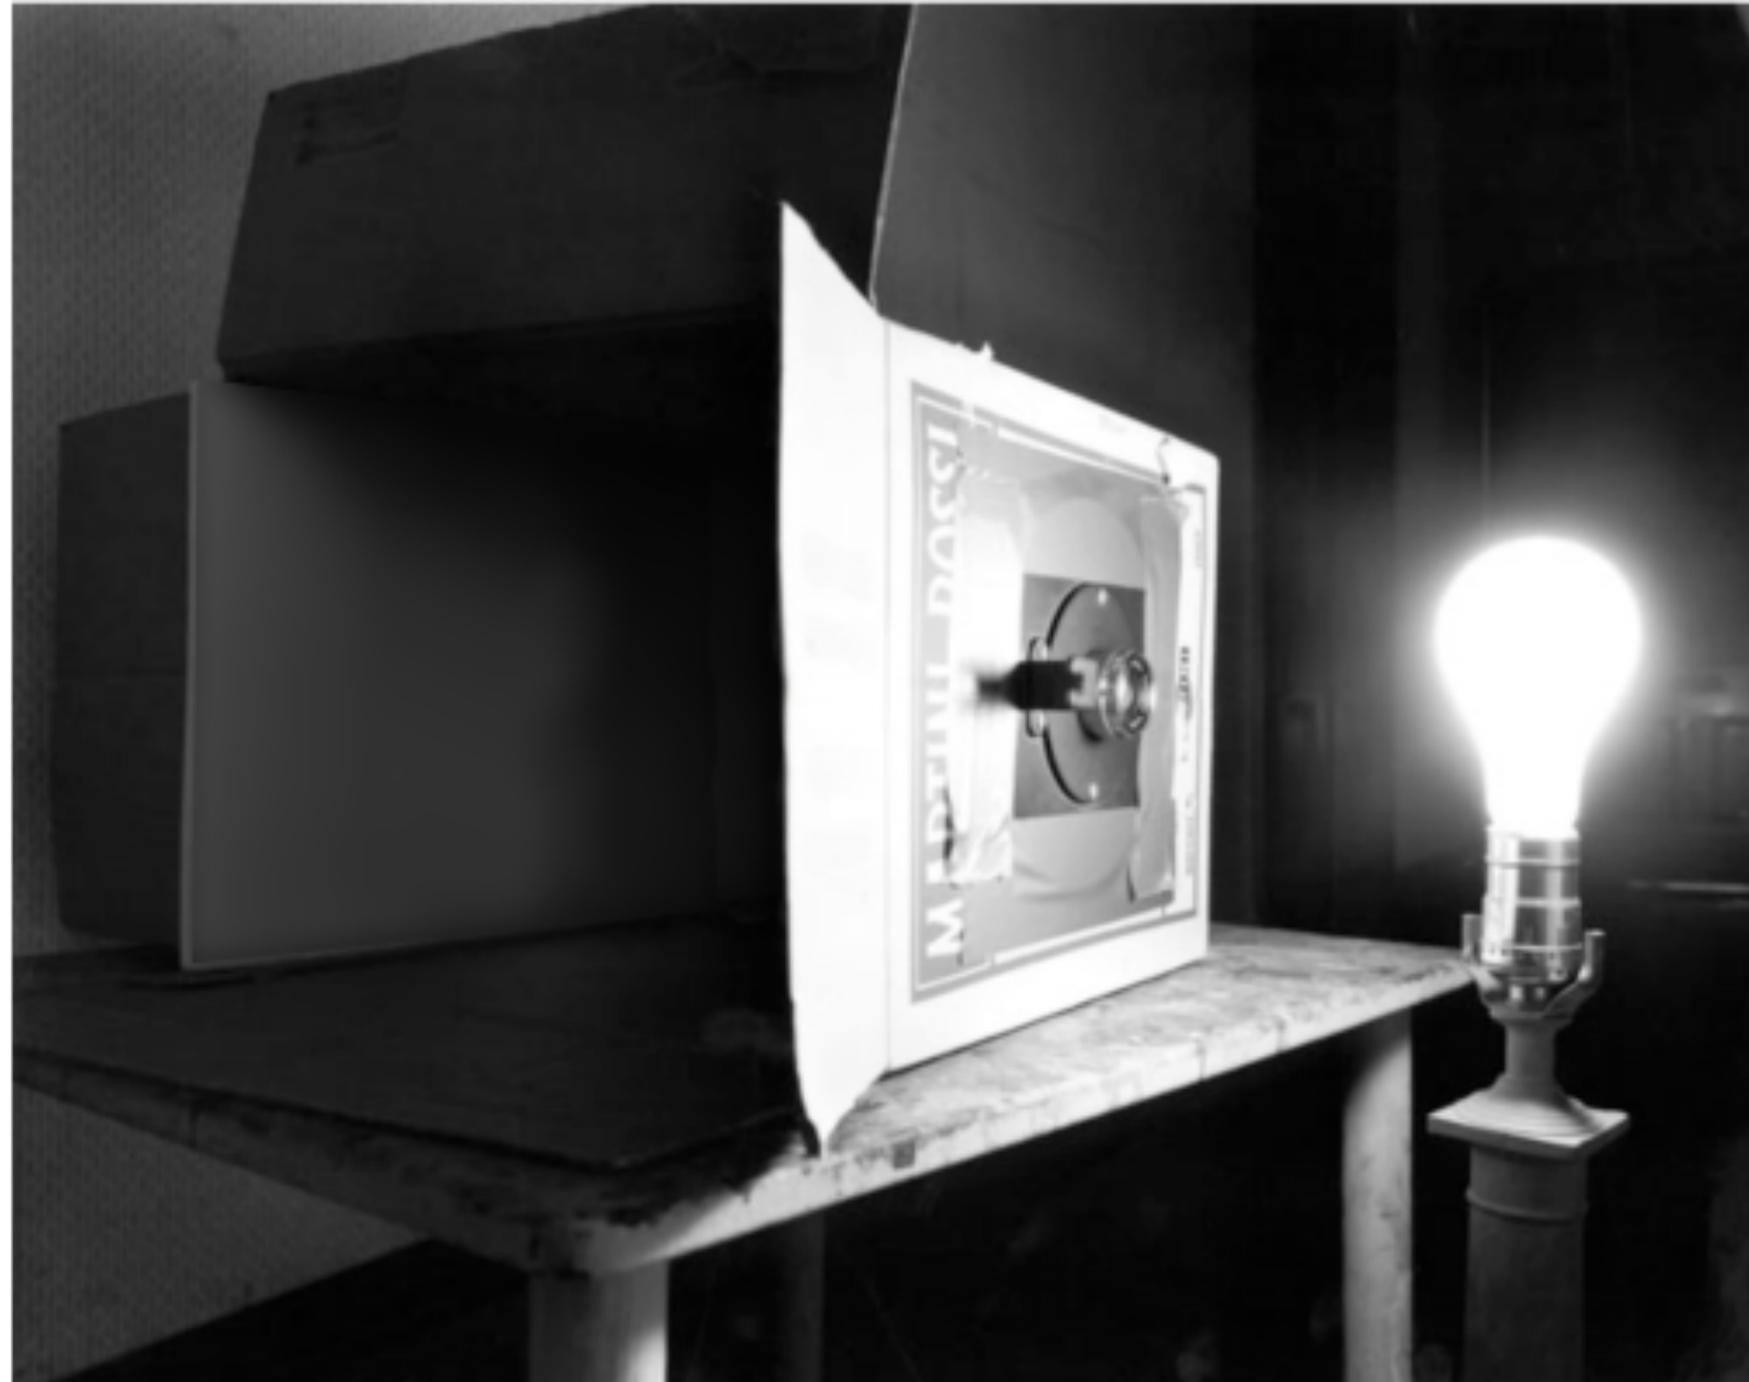
Useful concept: virtual image

# Pinhole camera



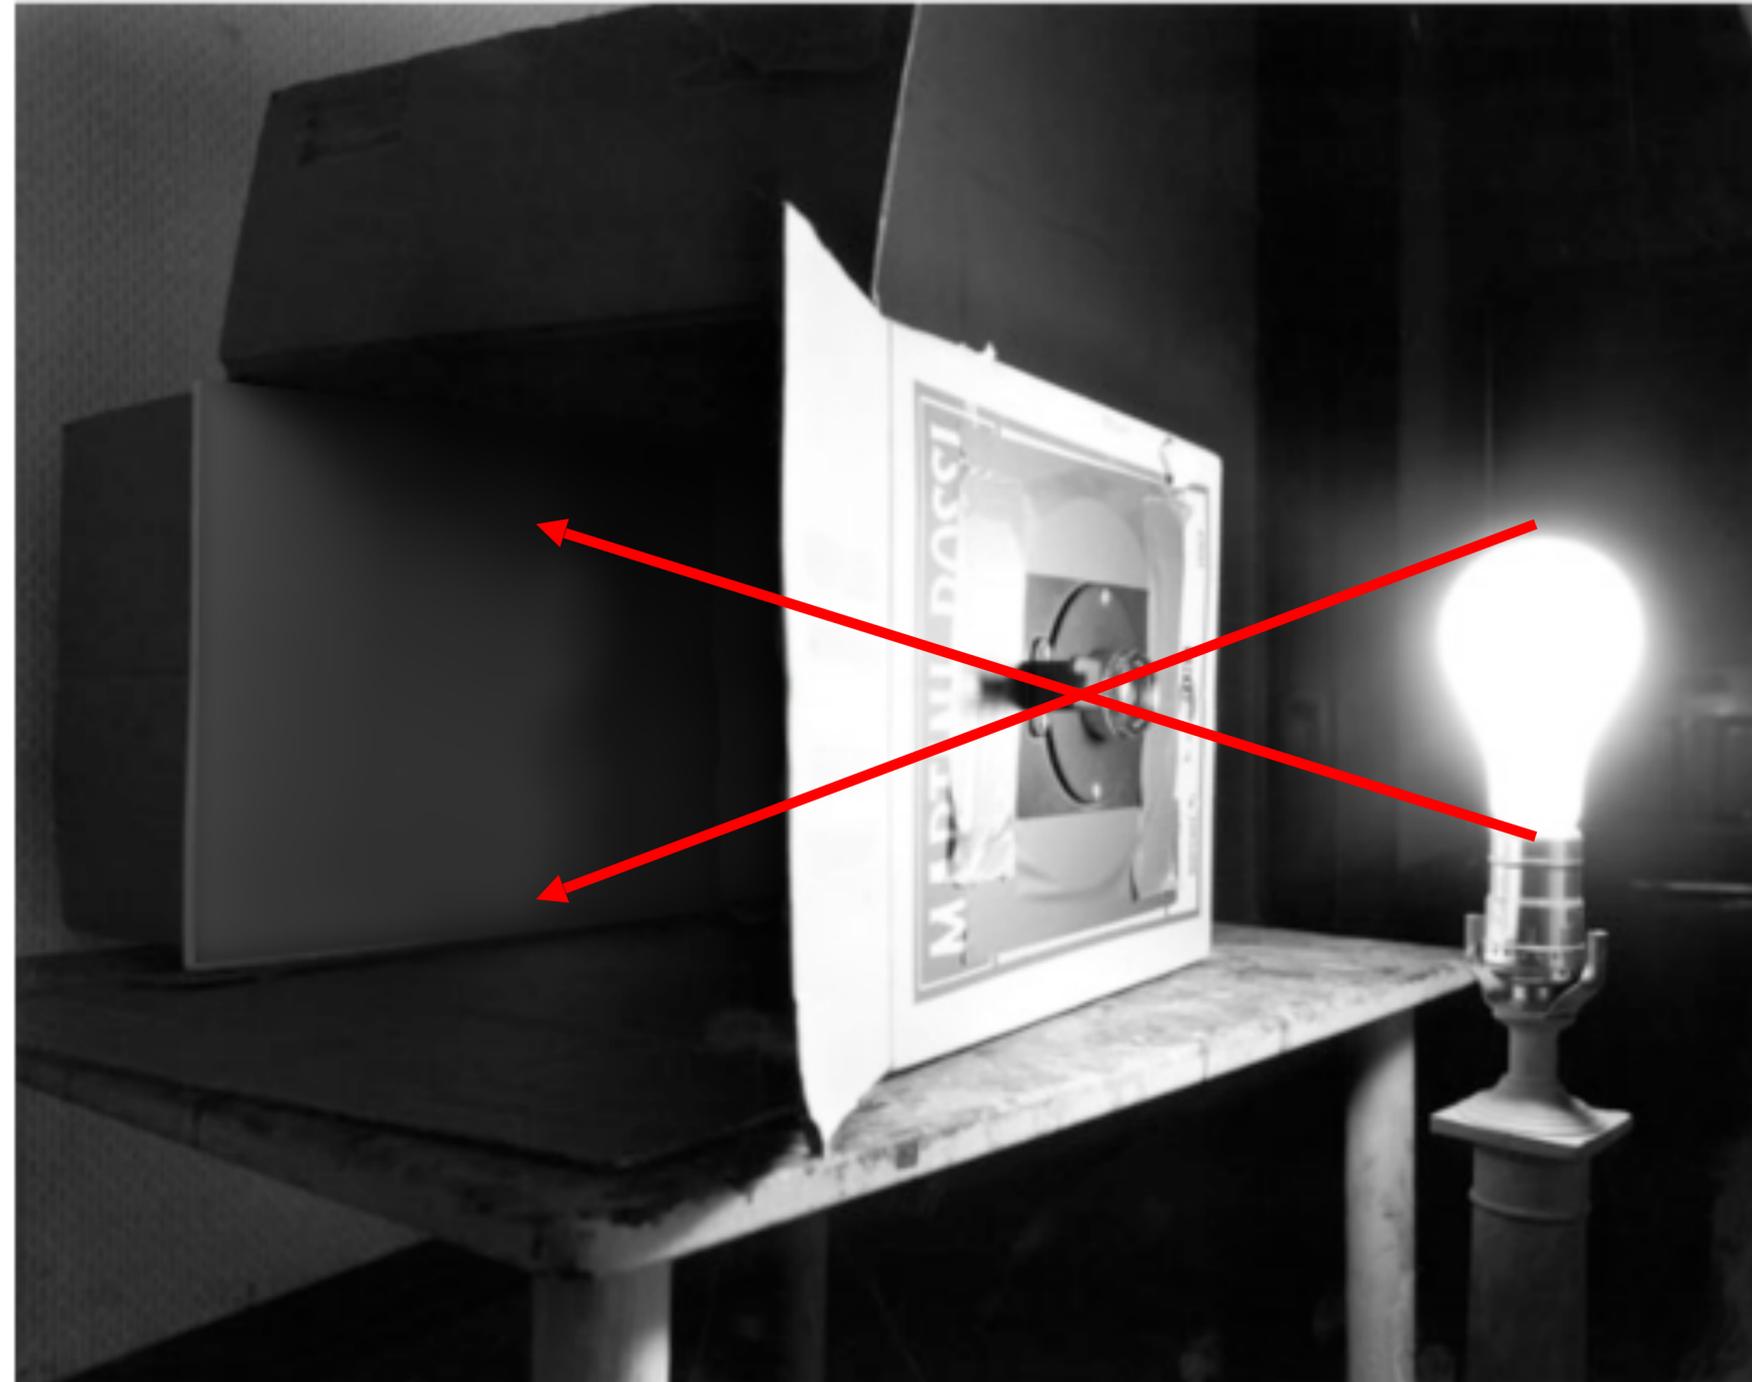
Photograph by Abelardo Morell, 1991

# Pinhole camera



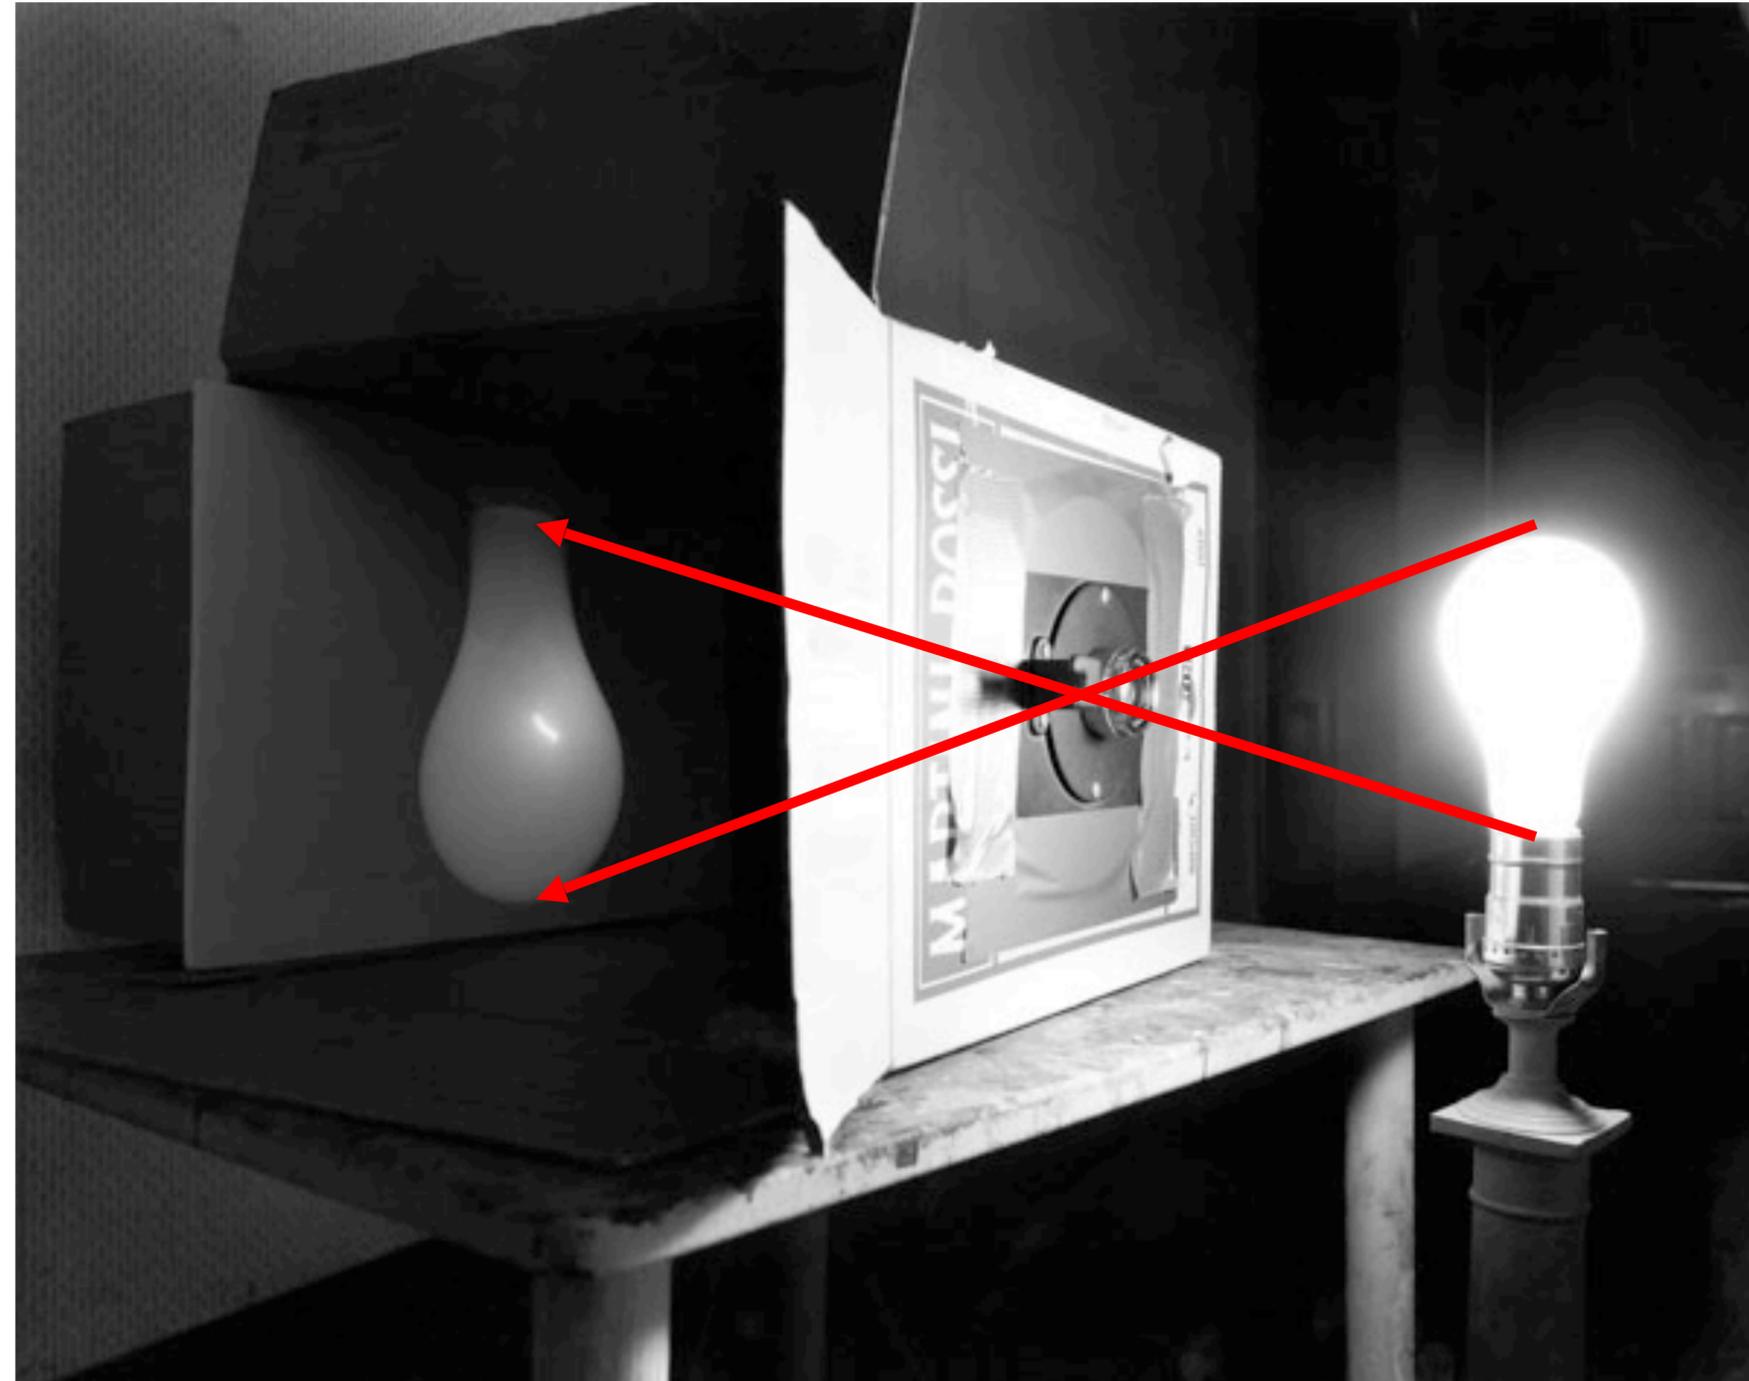
Photograph by Abelardo Morell, 1991

# Pinhole camera

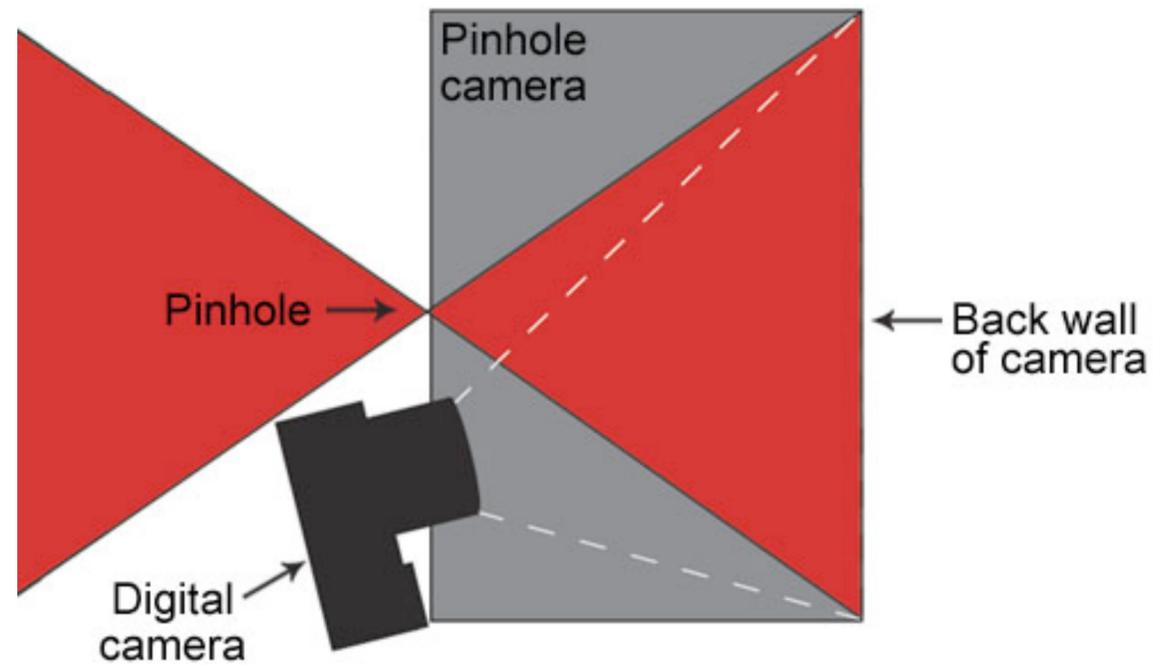


Photograph by Abelardo Morell, 1991

# Pinhole camera

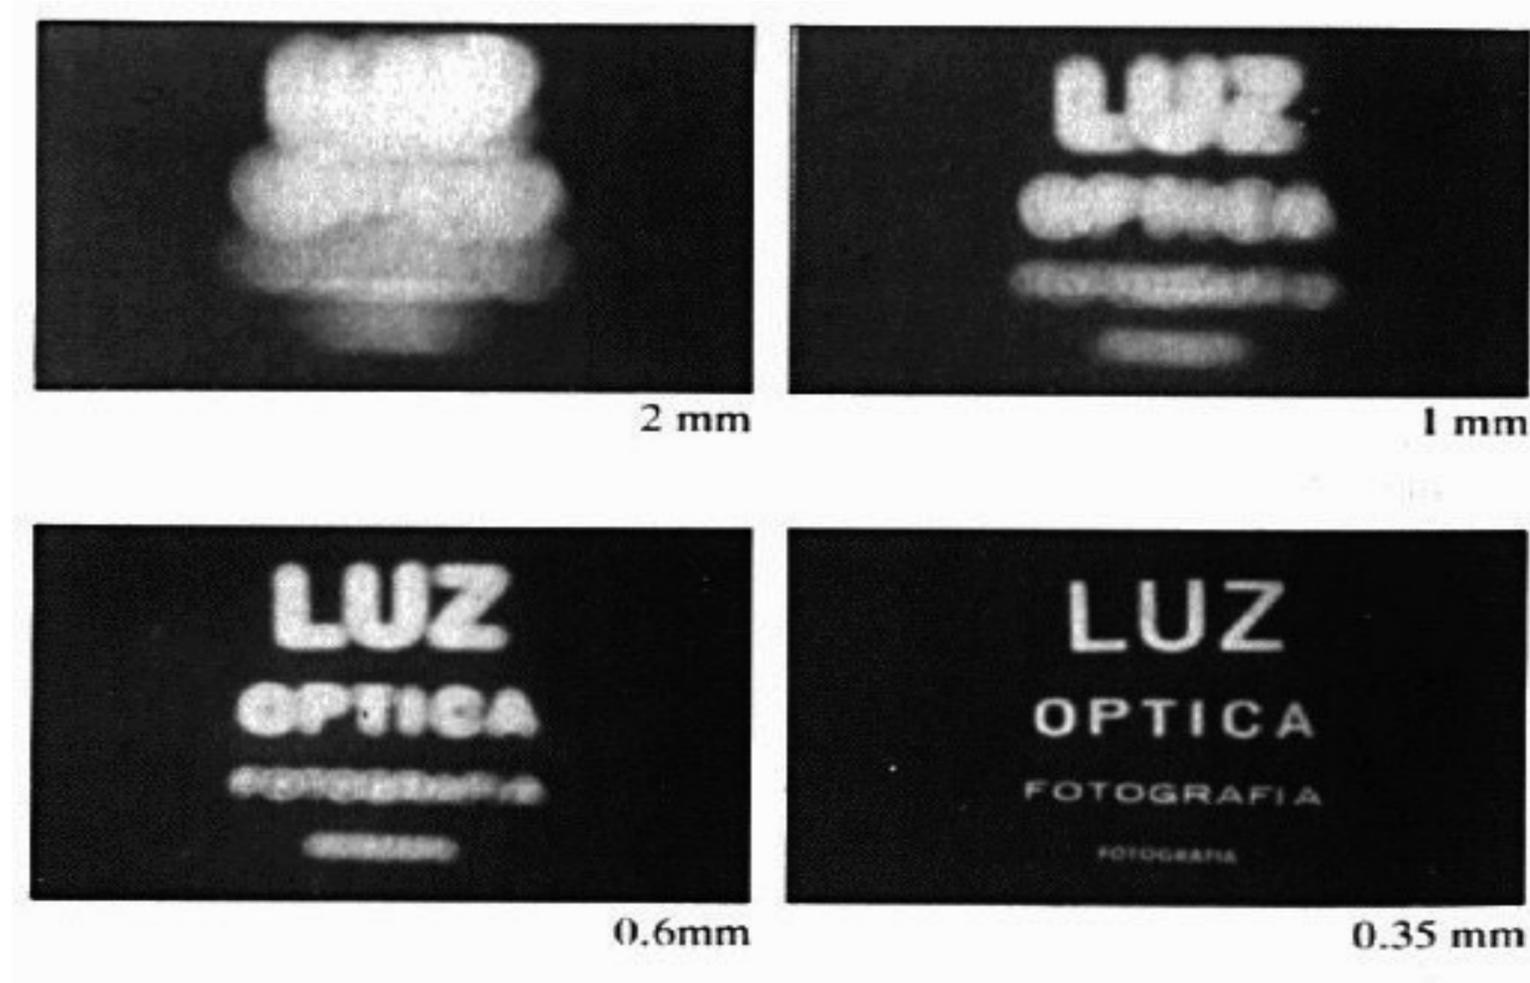


Photograph by Abelardo Morell, 1991



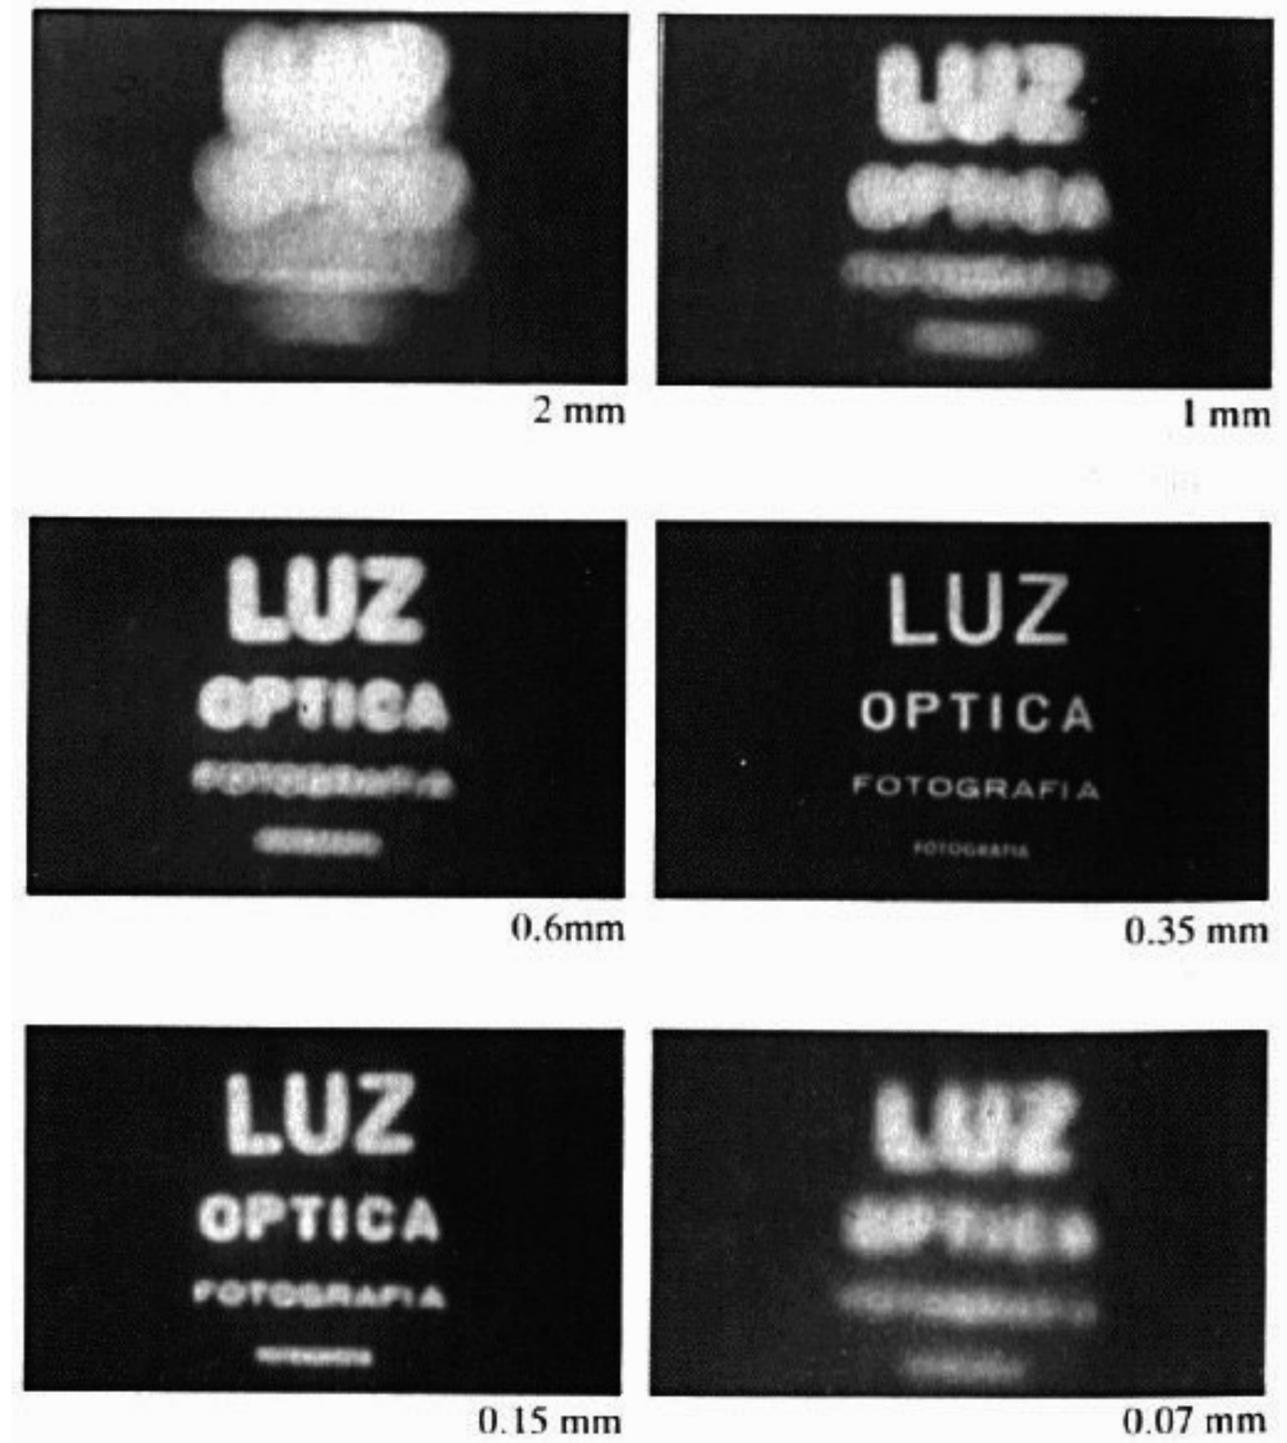


# Shrinking the aperture

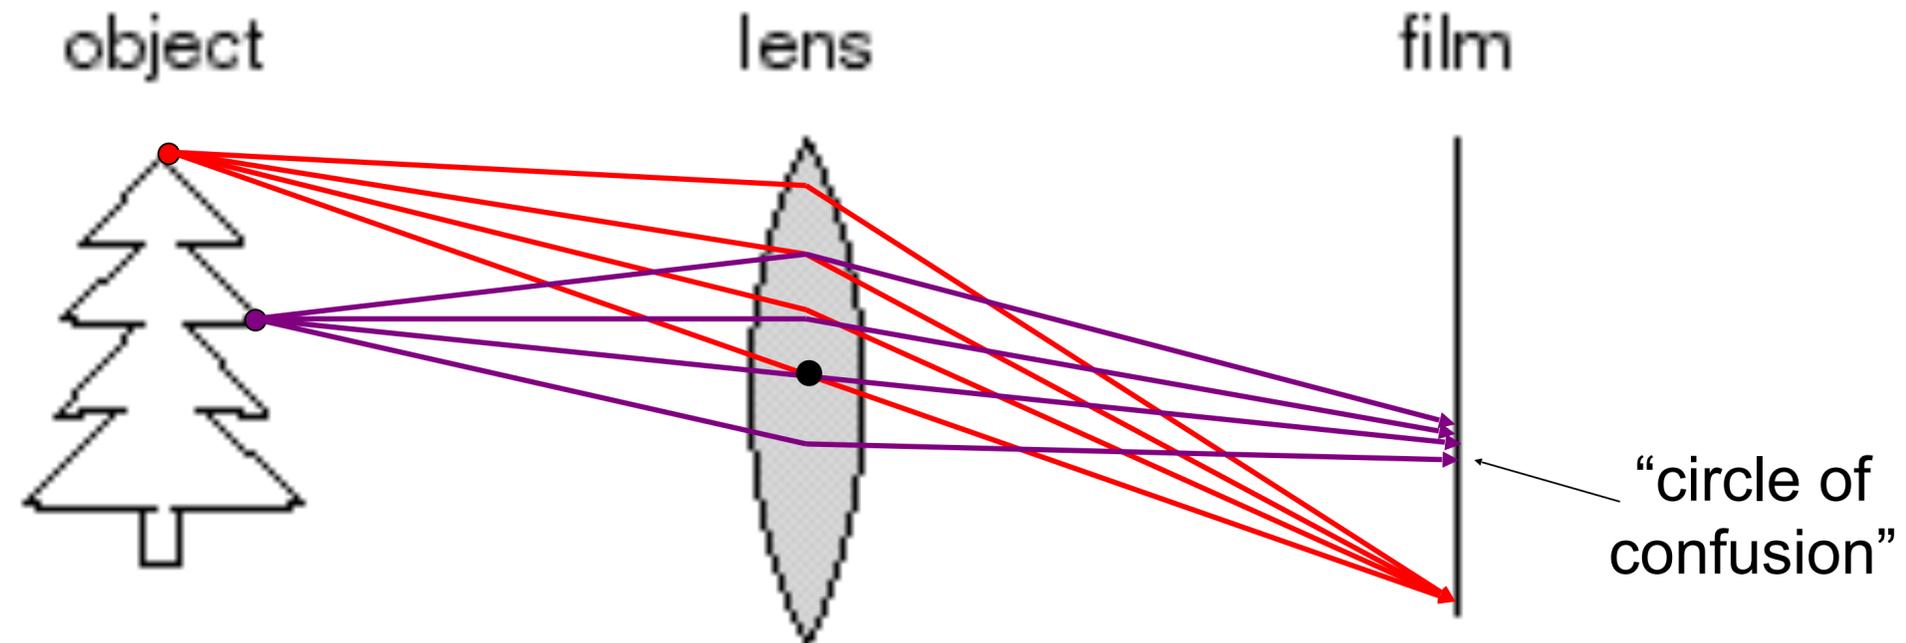


- Why not make the aperture as small as possible?
  - Less light gets through
  - *Diffraction* effects...

# Shrinking the aperture



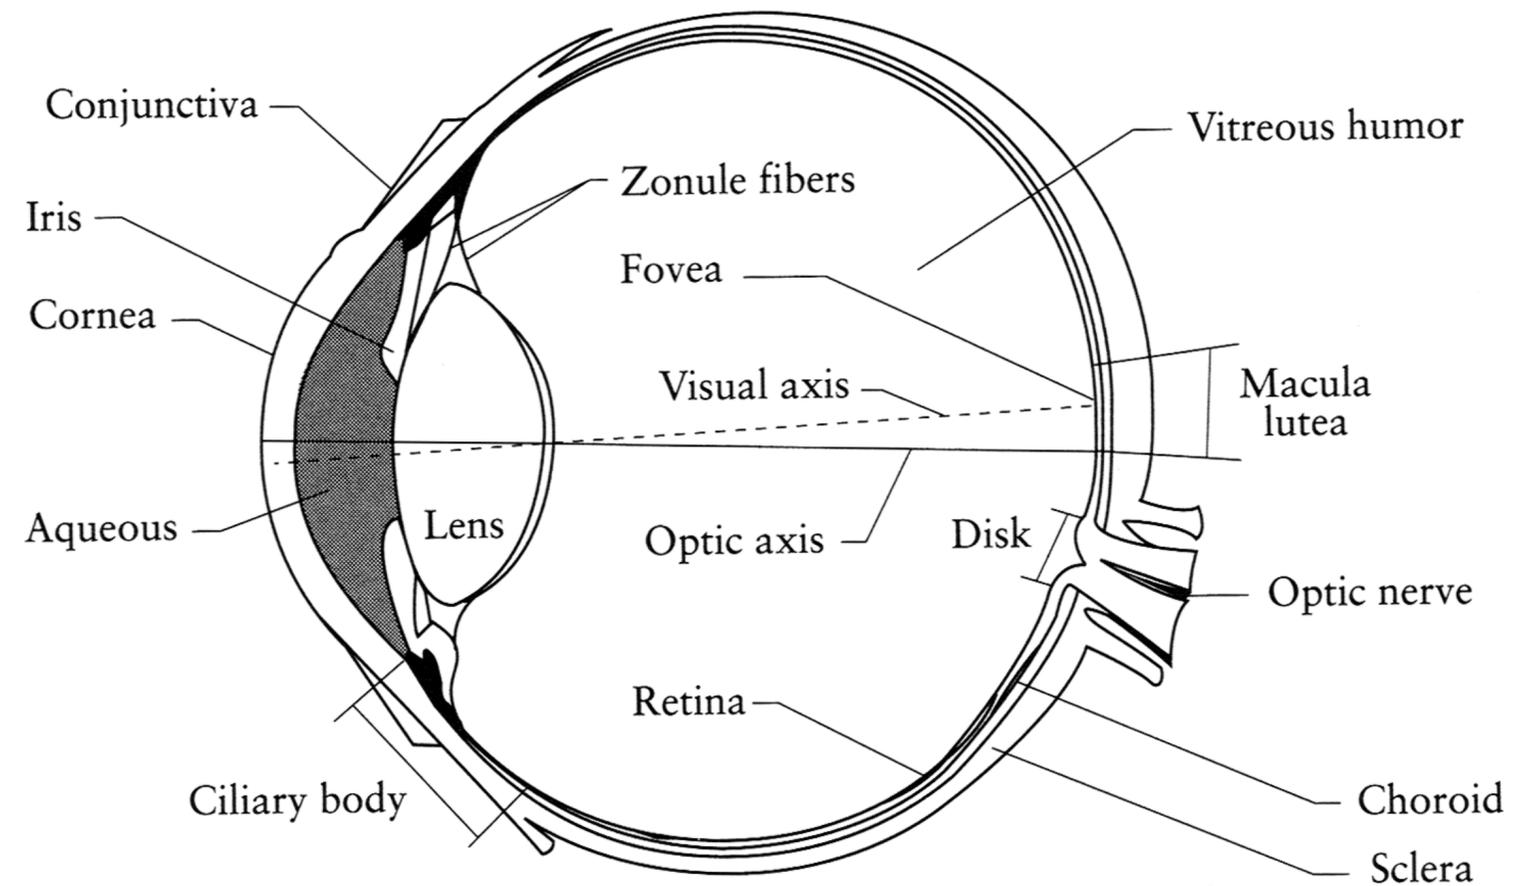
# Adding a lens



A lens focuses light onto the film

- There is a specific distance at which objects are “in focus”
  - other points project to a “circle of confusion” in the image
- Changing the shape of the lens changes this distance

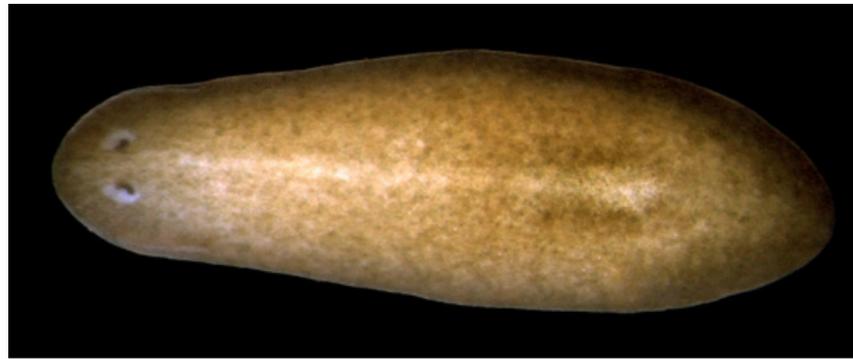
# The eye



## The human eye is a camera

- **Iris** - colored annulus with radial muscles
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- What's the “film”?
- Photoreceptor cells (rods and cones) in the **retina**

# Eyes in nature: eyespots to pinhole camera



[http://upload.wikimedia.org/wikipedia/commons/6/6d/Mantis\\_shrimp.jpg](http://upload.wikimedia.org/wikipedia/commons/6/6d/Mantis_shrimp.jpg)

# Pinhole cameras in unexpected places



Tree shadow during a solar eclipse

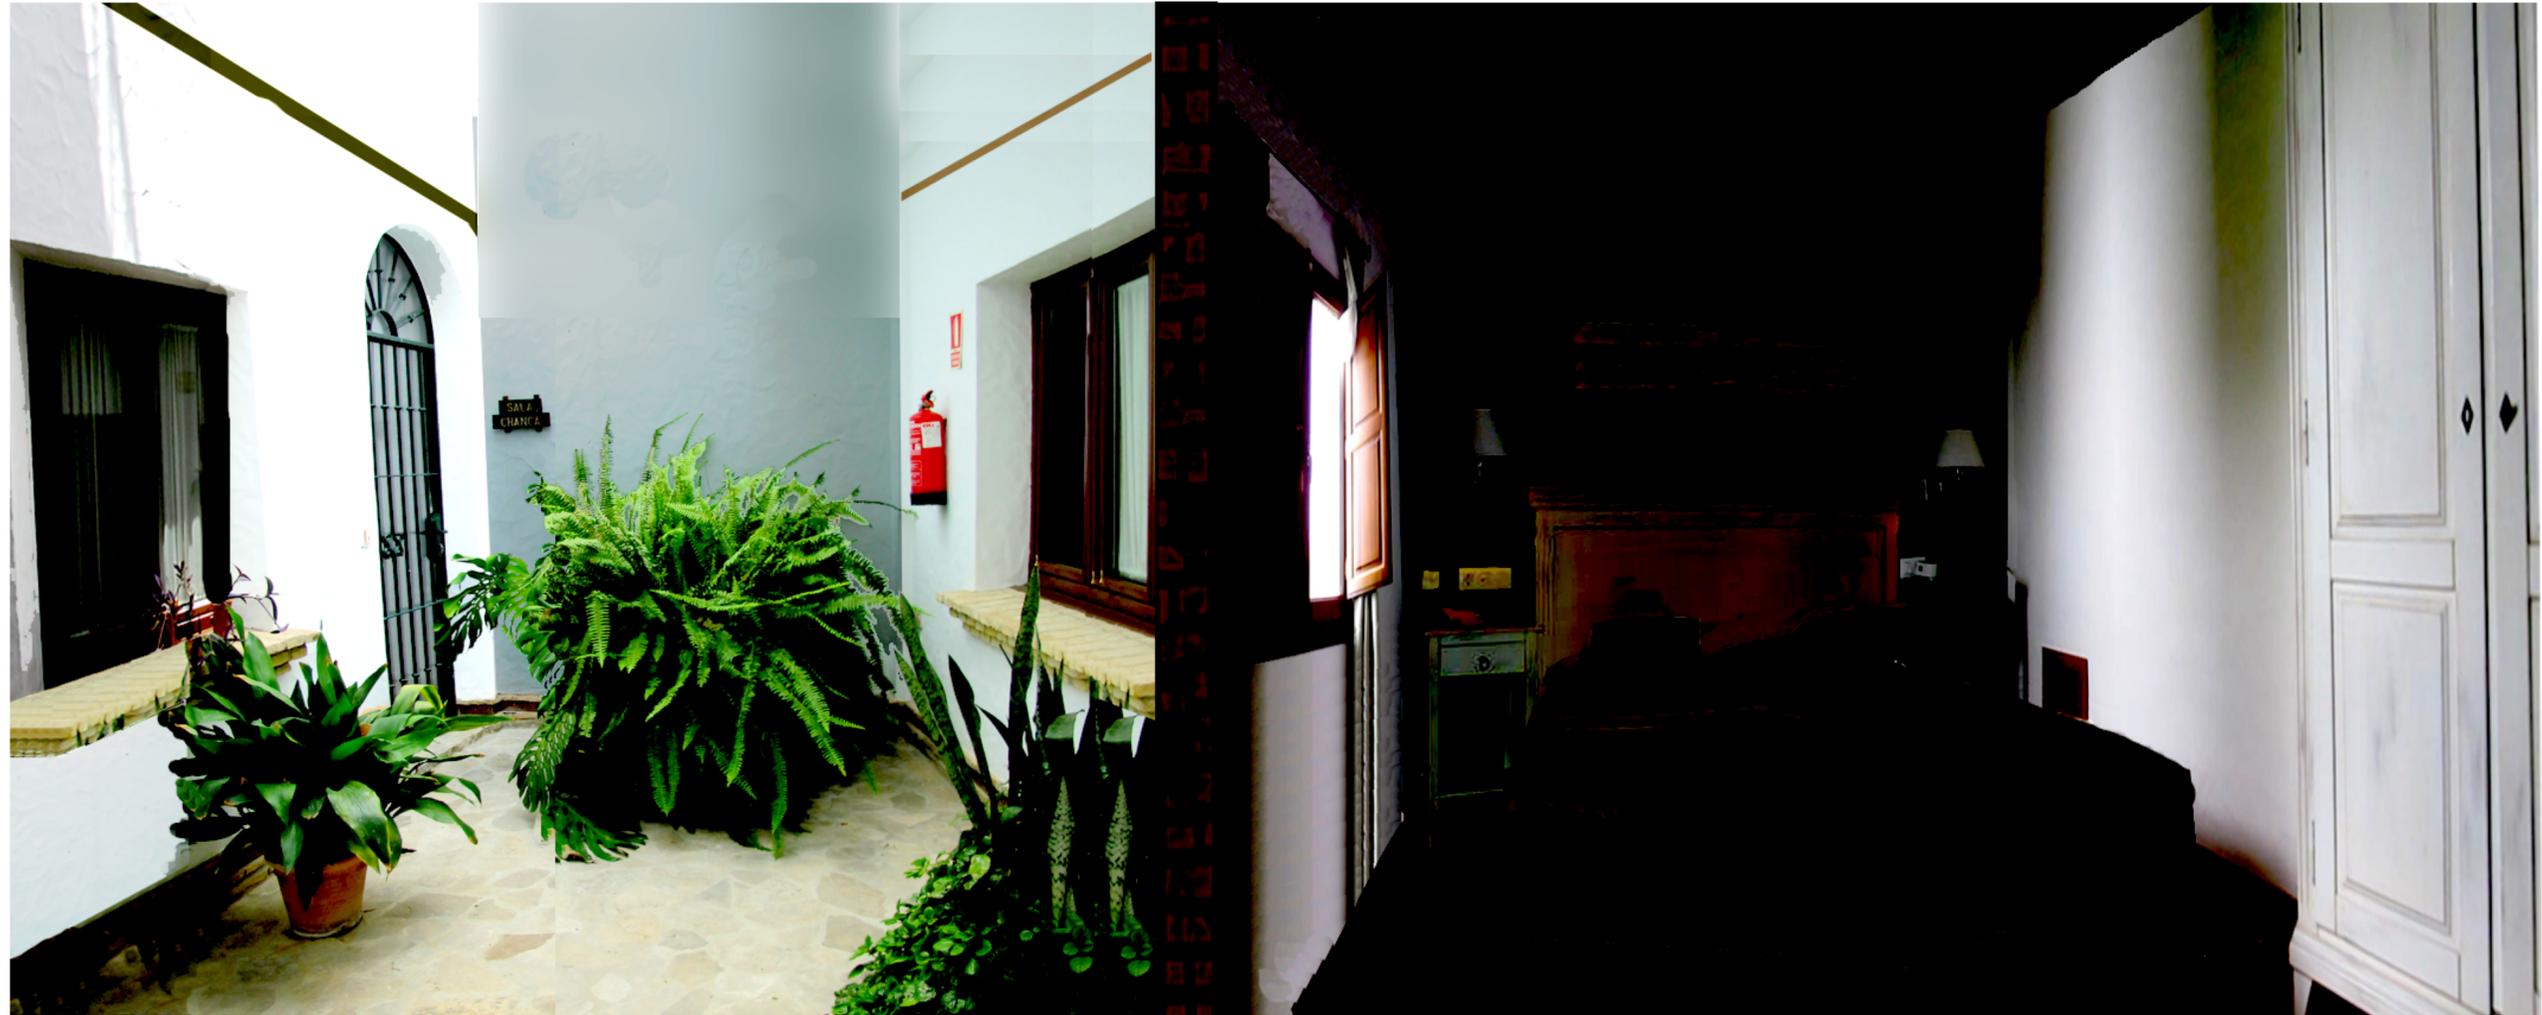
photo credit: Nils van der Burg

<http://www.physicstogo.org/index.cfm>



Shadows?





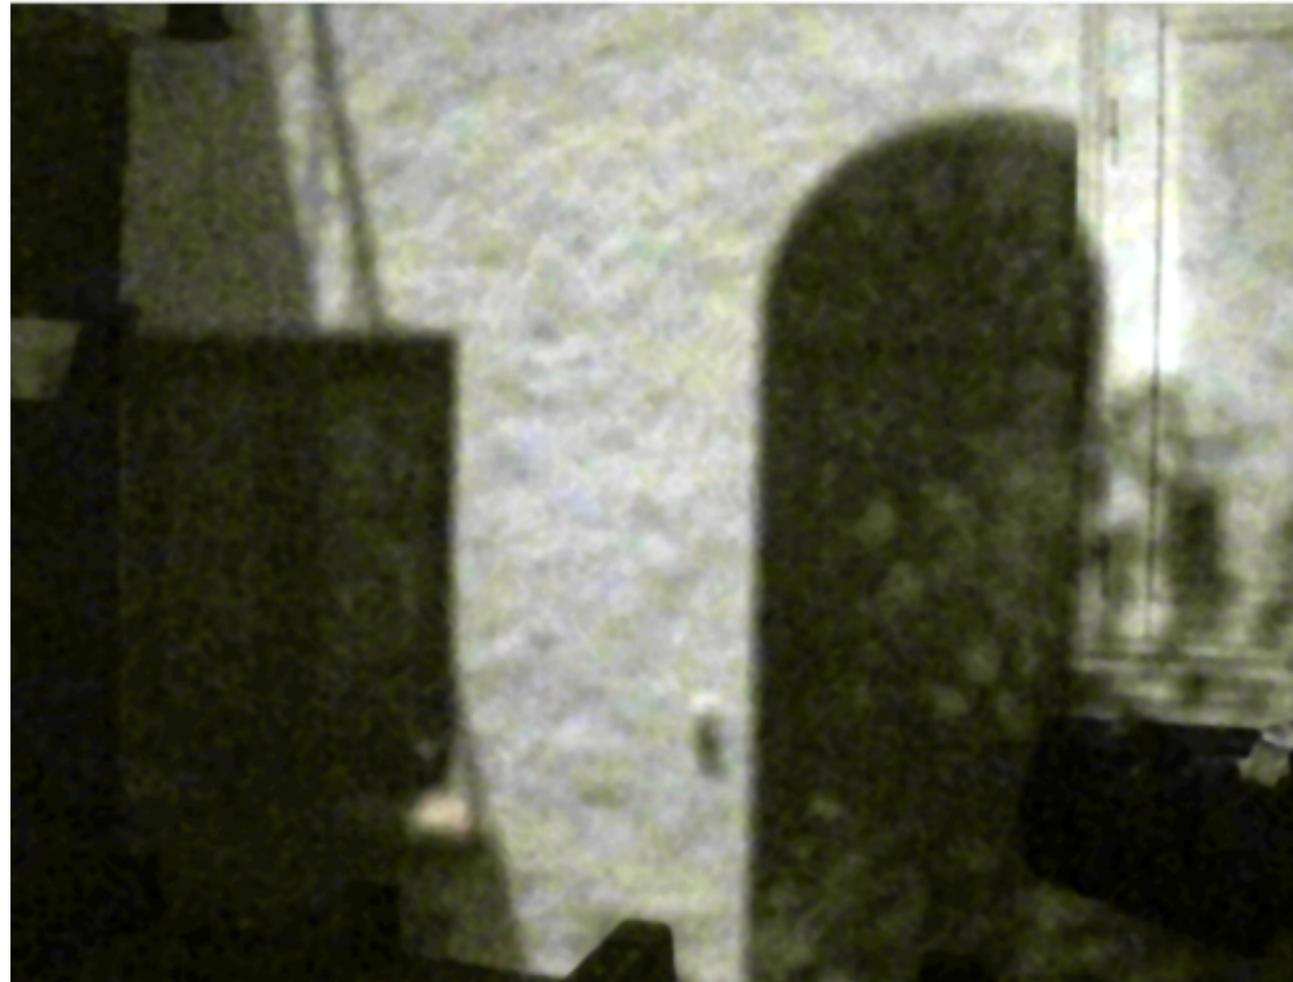
# Accidental pinhole camera







Window turned into a pinhole



View outside







Window open

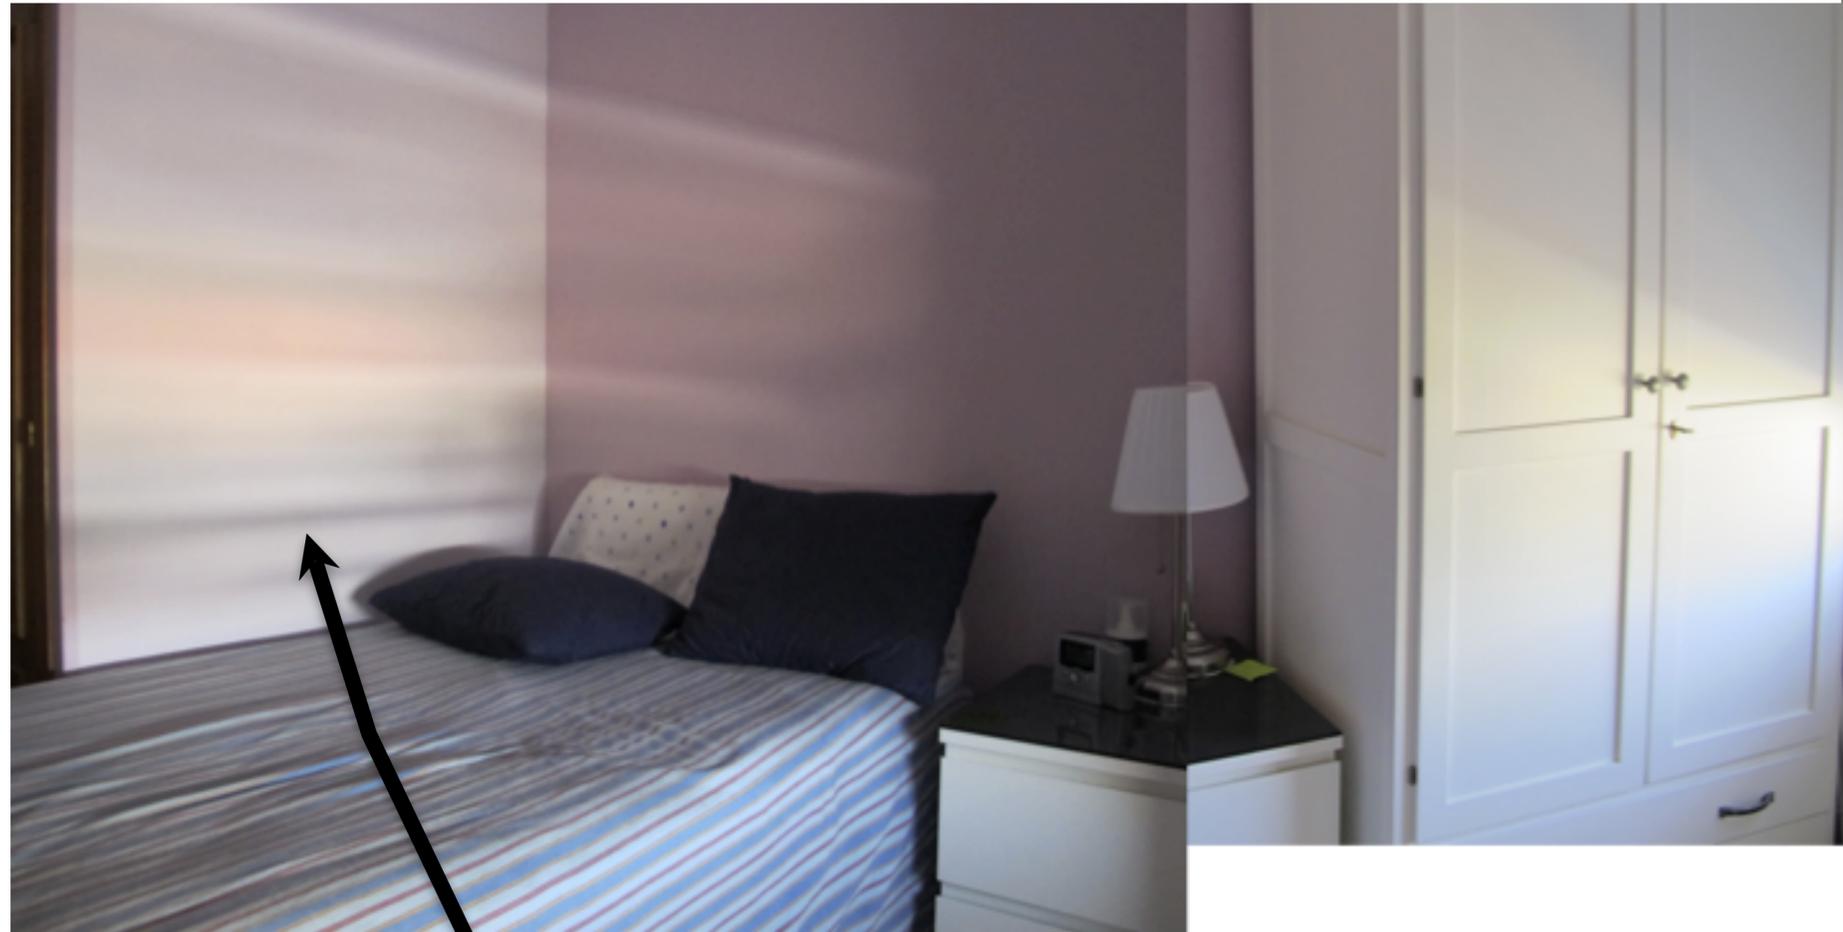


Window turned into a pinhole





# Accidental pinhole camera



Outside scene

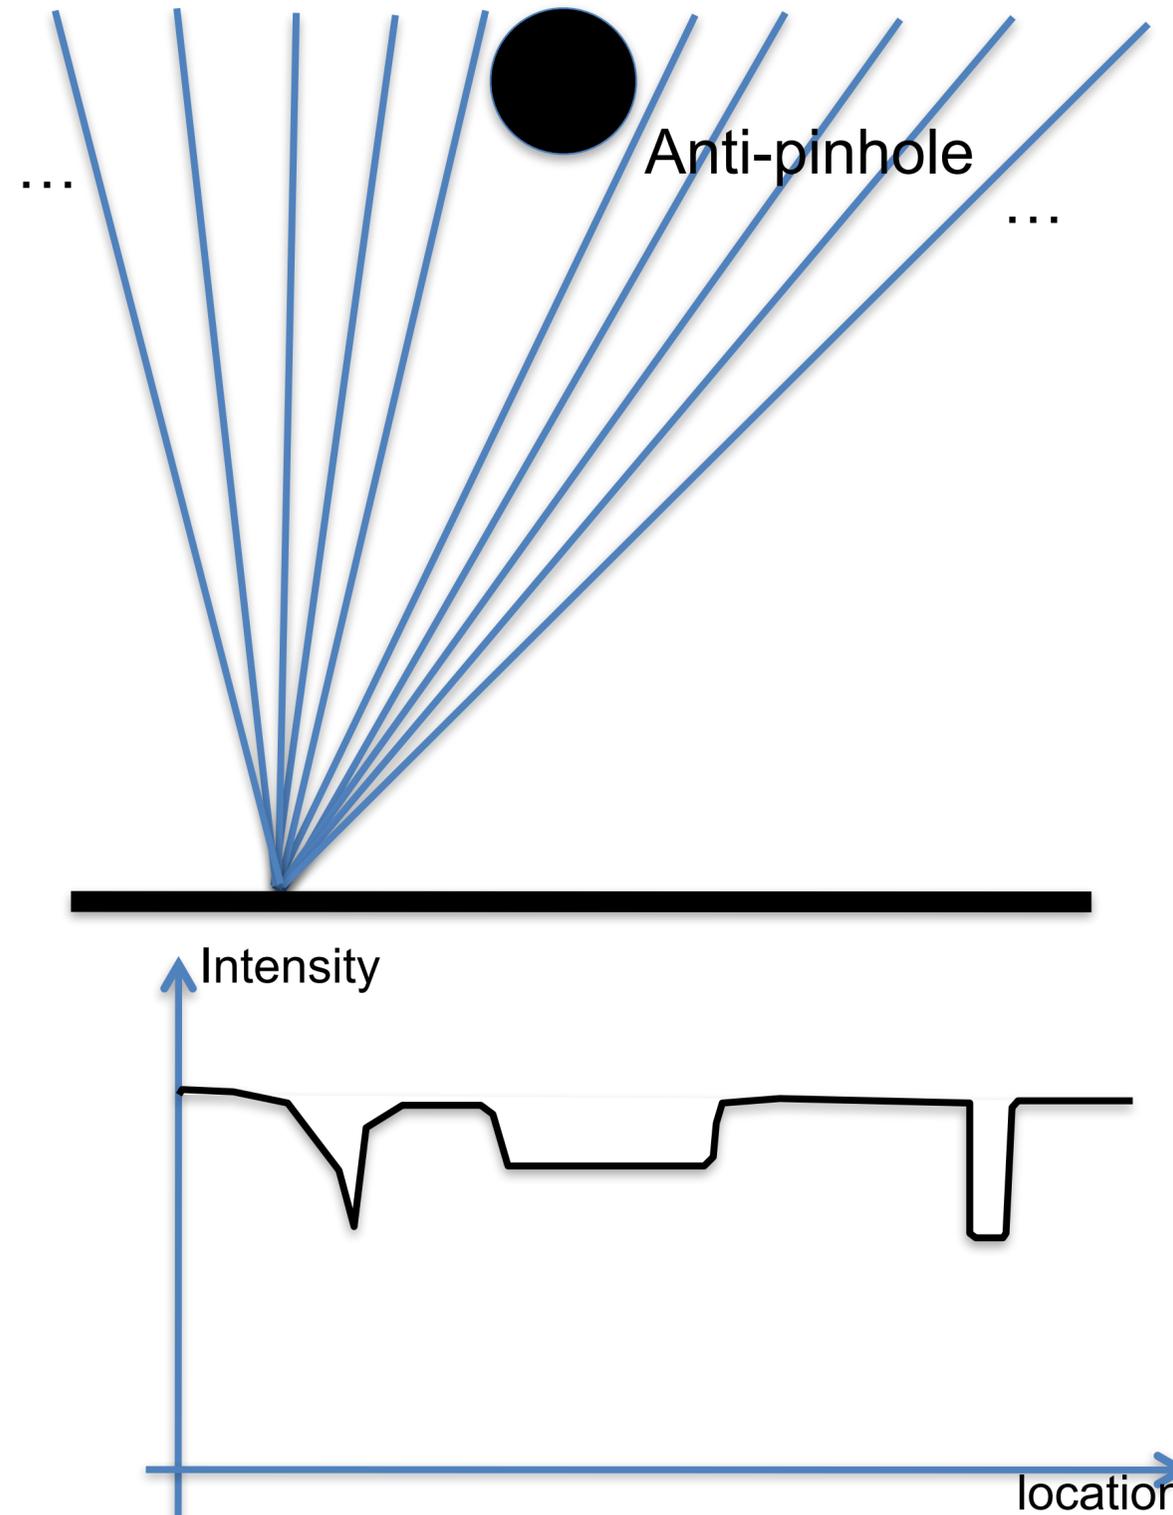
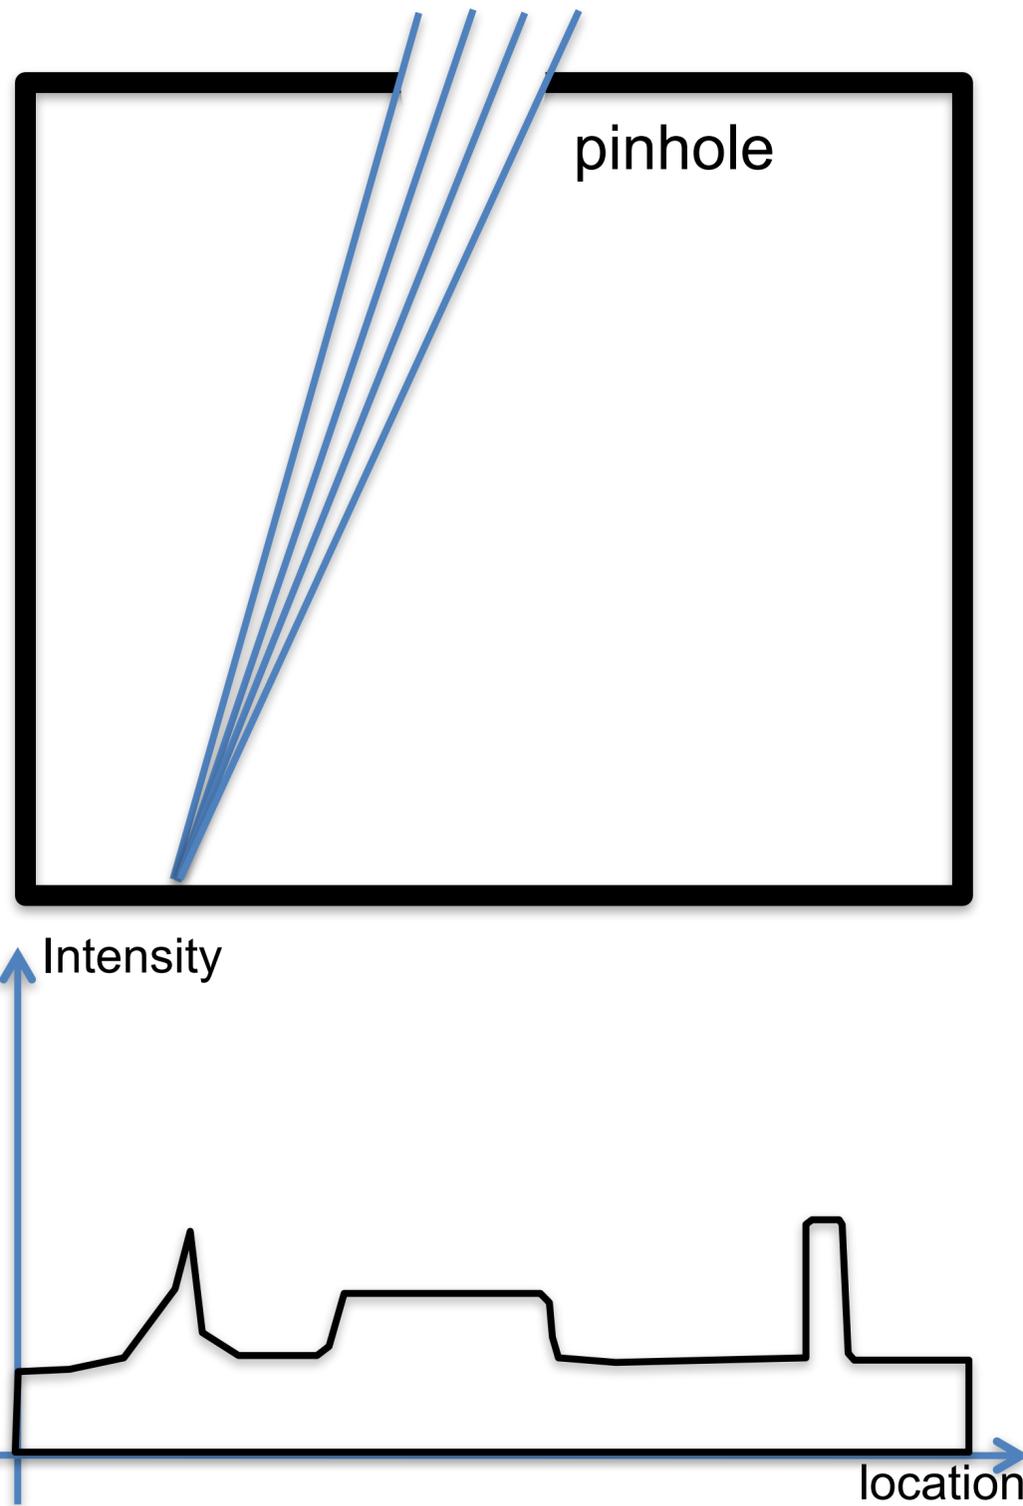
\*



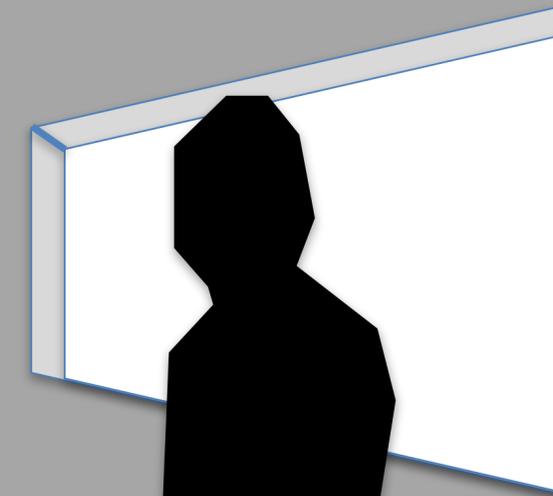
Aperture

See Zomet, A.; Nayar, S.K. CVPR 2006 for a detailed analysis.

# Pinhole and Anti-pinhole cameras



# Mixed accidental pinhole and anti-pinhole cameras

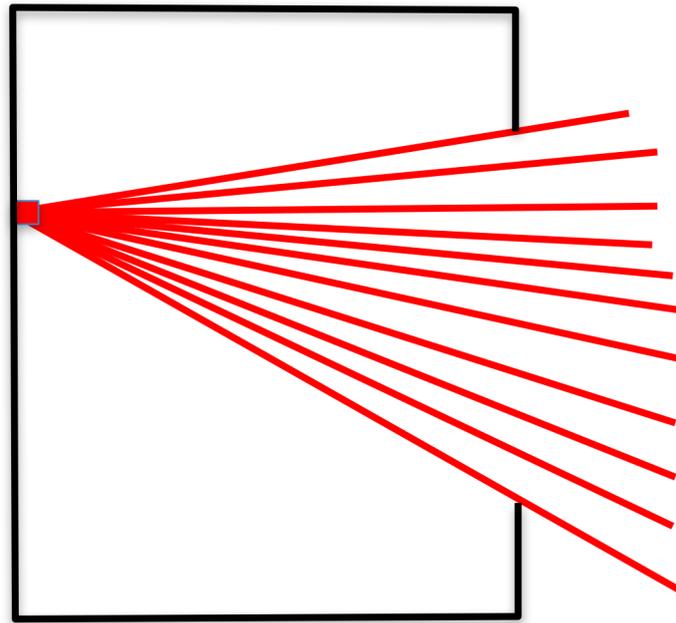


# Mixed accidental pinhole and anti-pinhole cameras

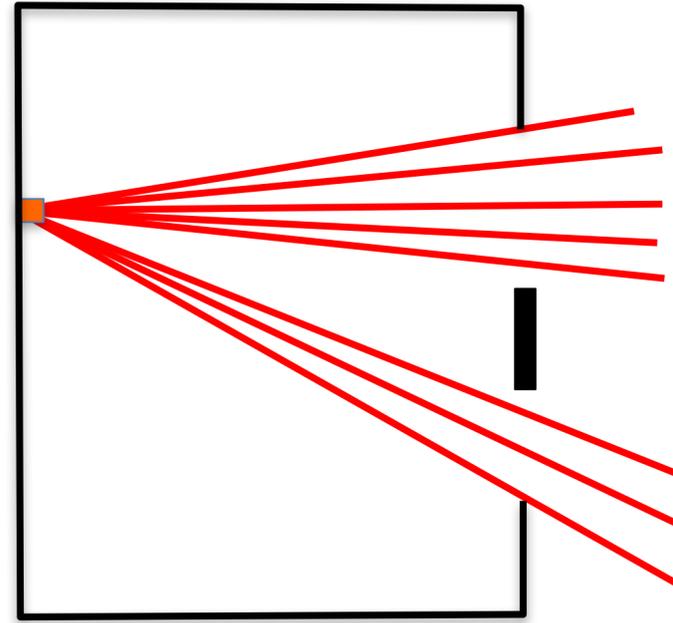


# Mixed accidental pinhole and anti-pinhole cameras

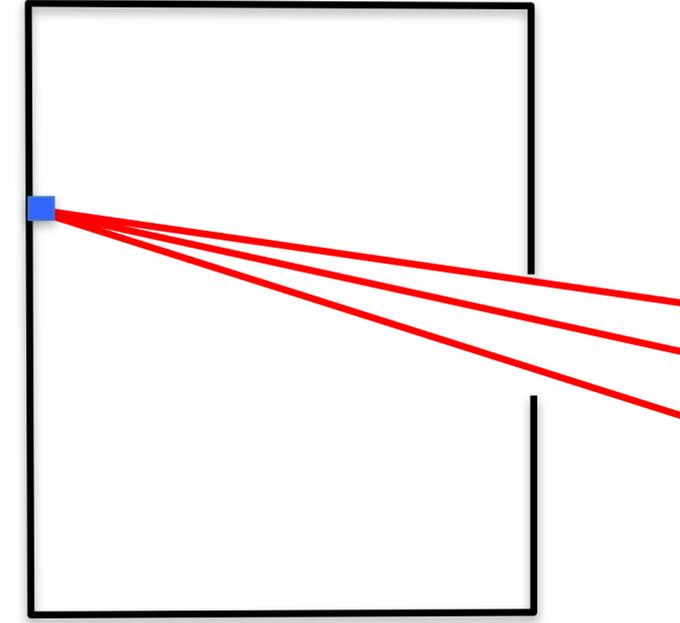
Room with a window



Person in front of the window



Difference image



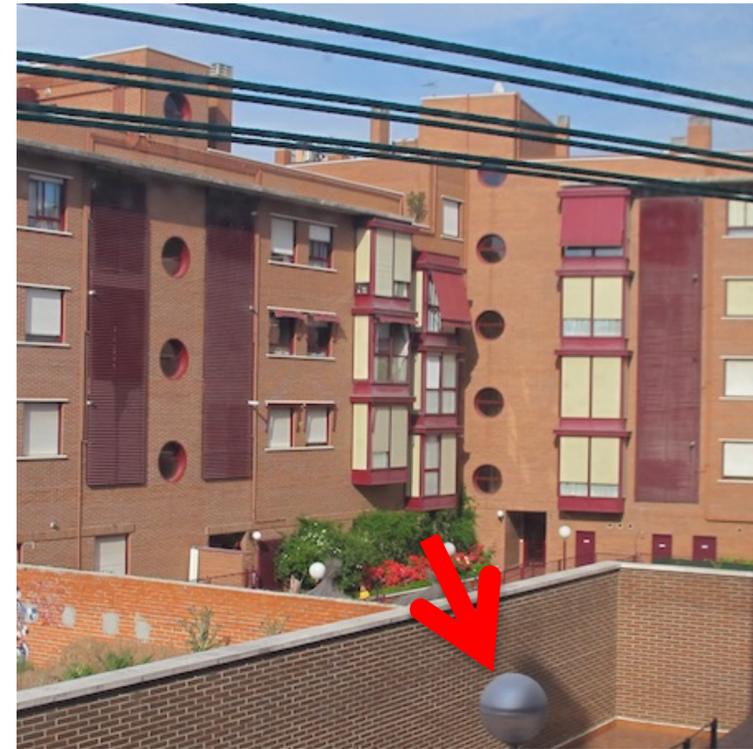
= ?

# Mixed accidental pinhole and anti-pinhole cameras

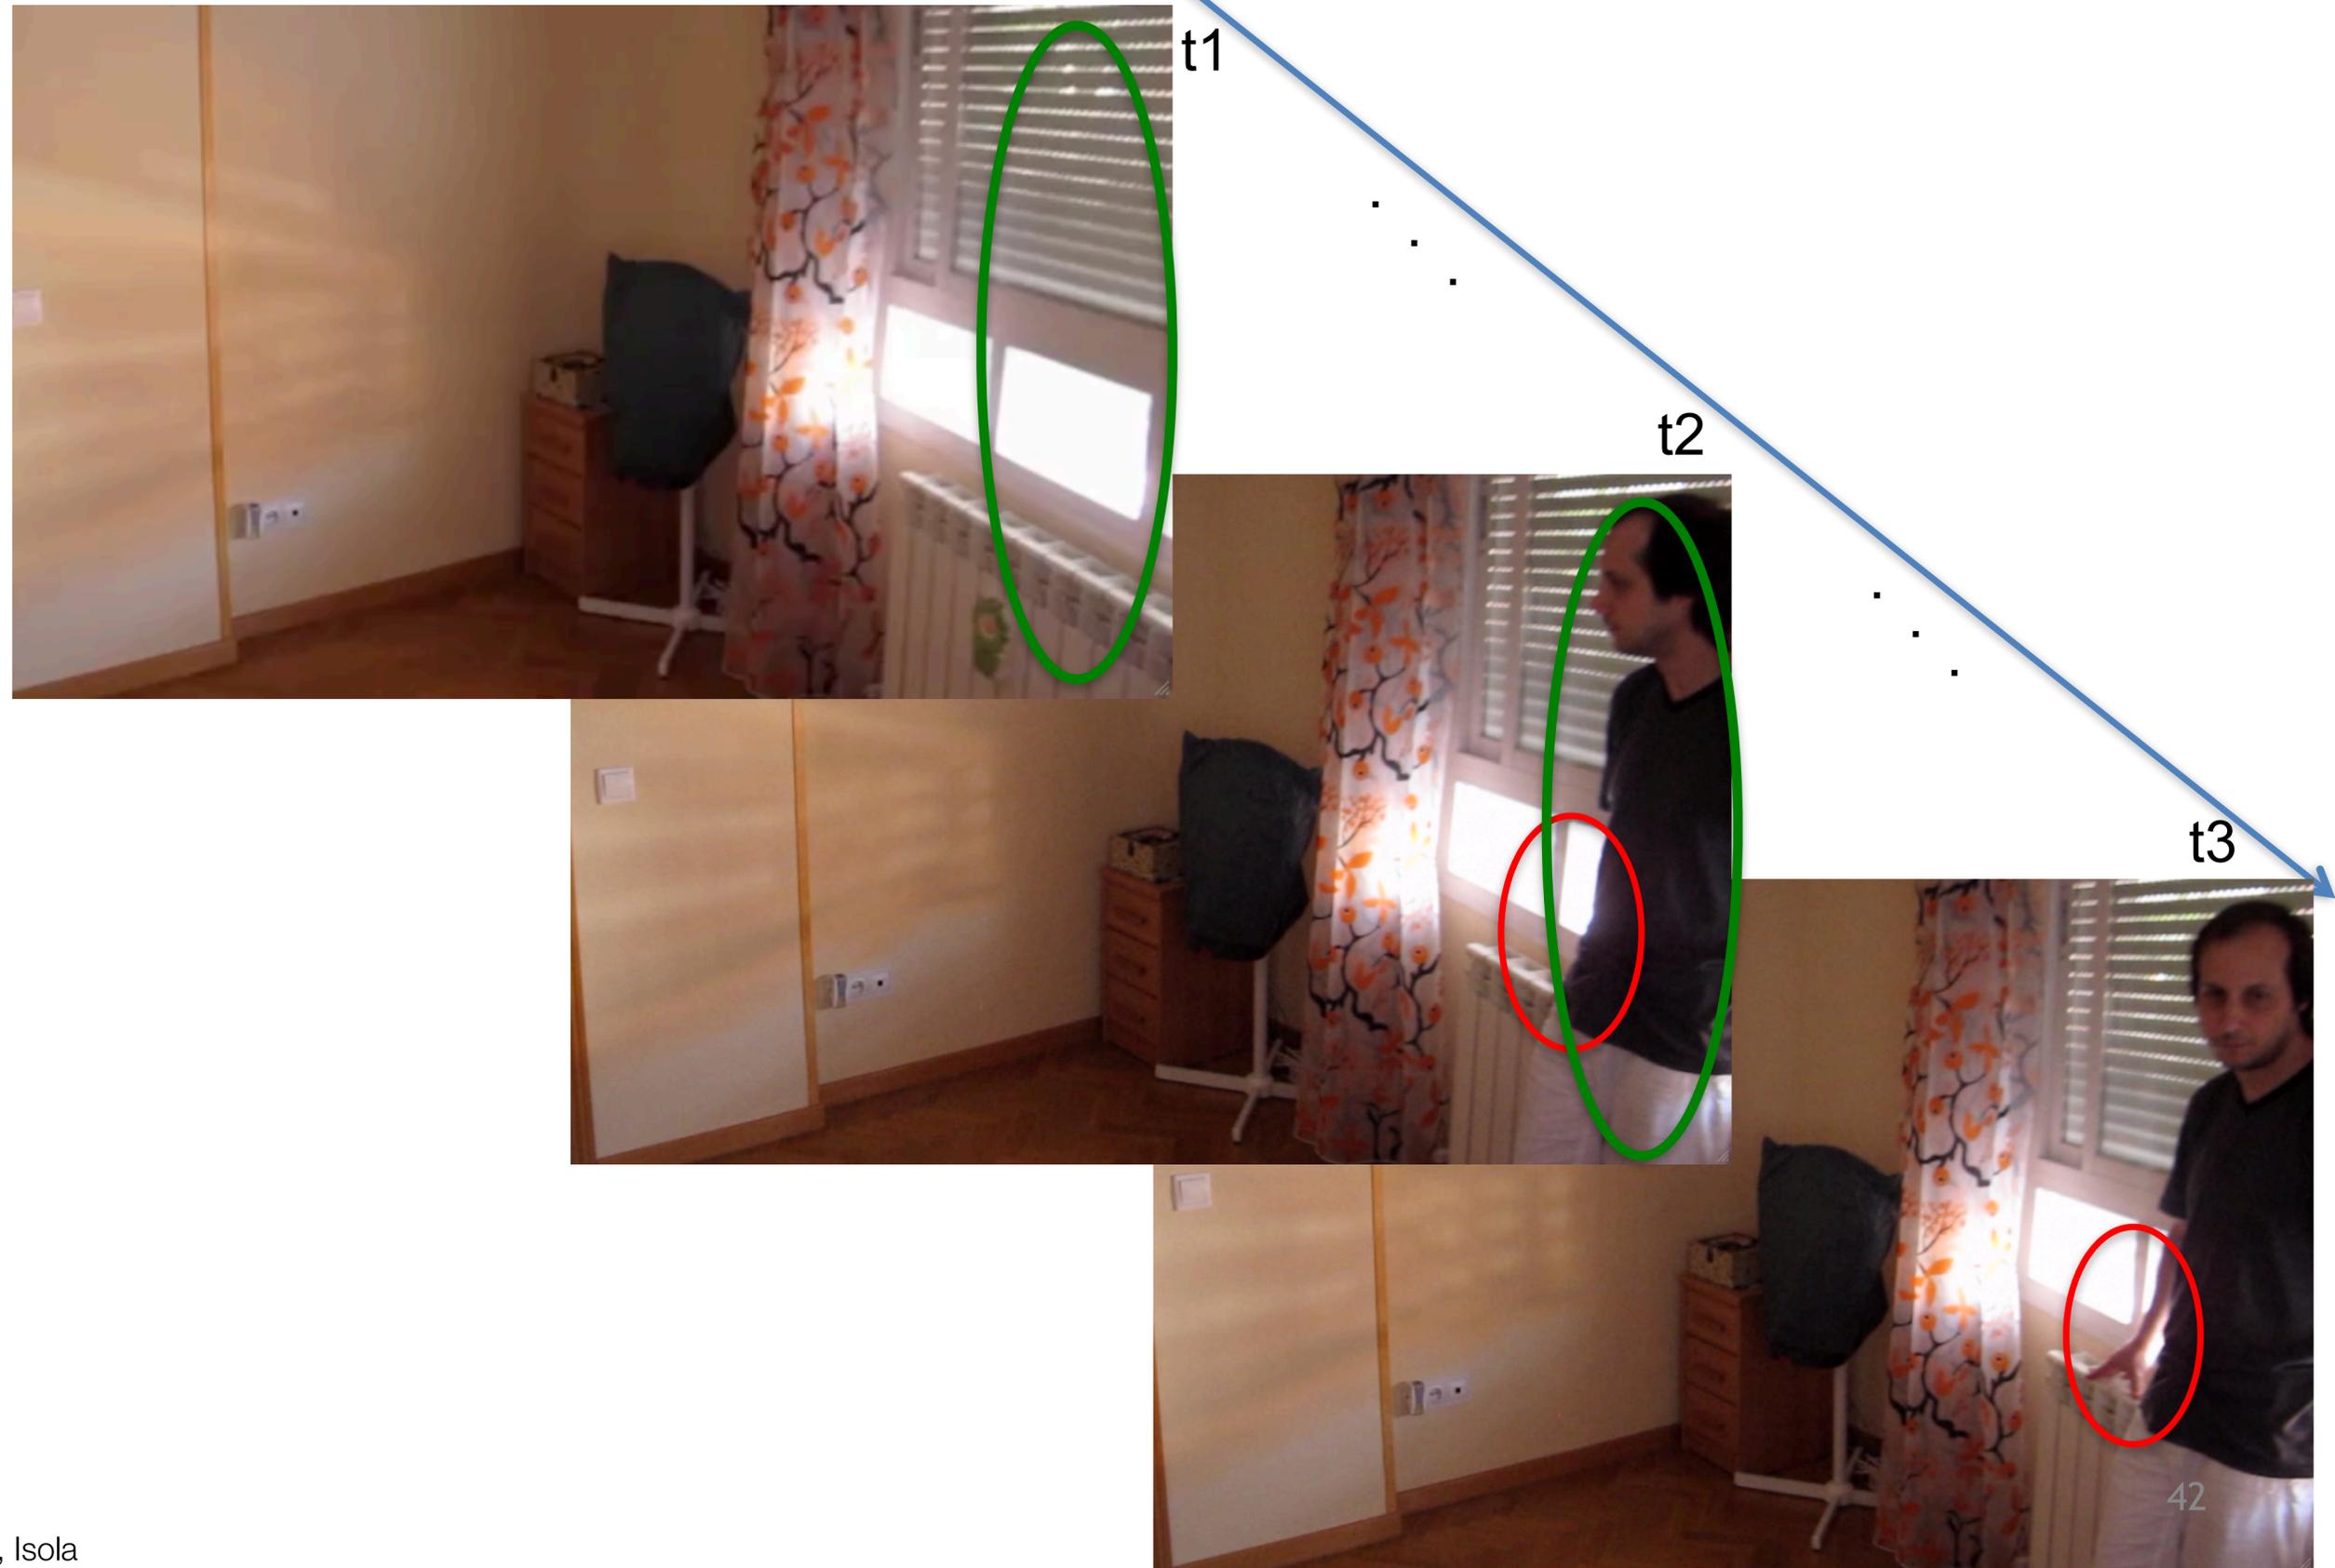
Body as the occluder



View outside the window



# Looking for a small accidental occluder



# Looking for a small accidental occluder

Body as the occluder



Hand as the occluder

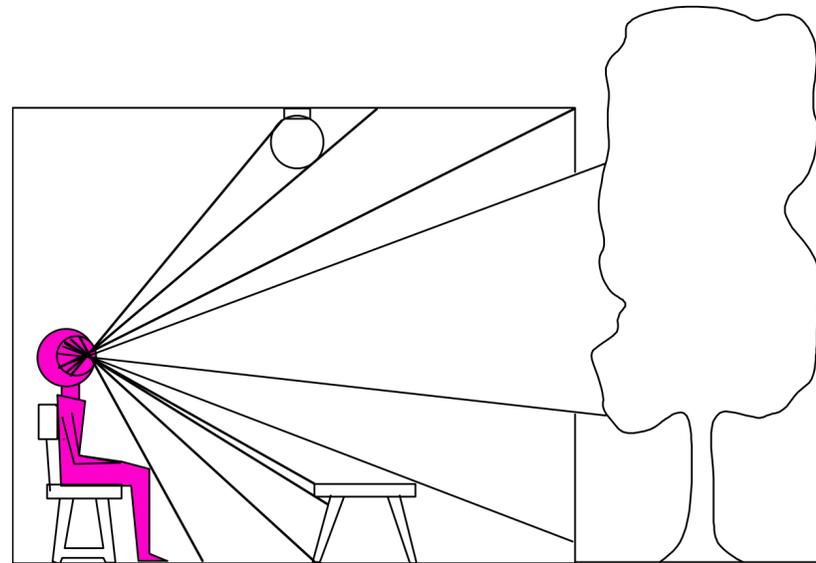


View outside the window



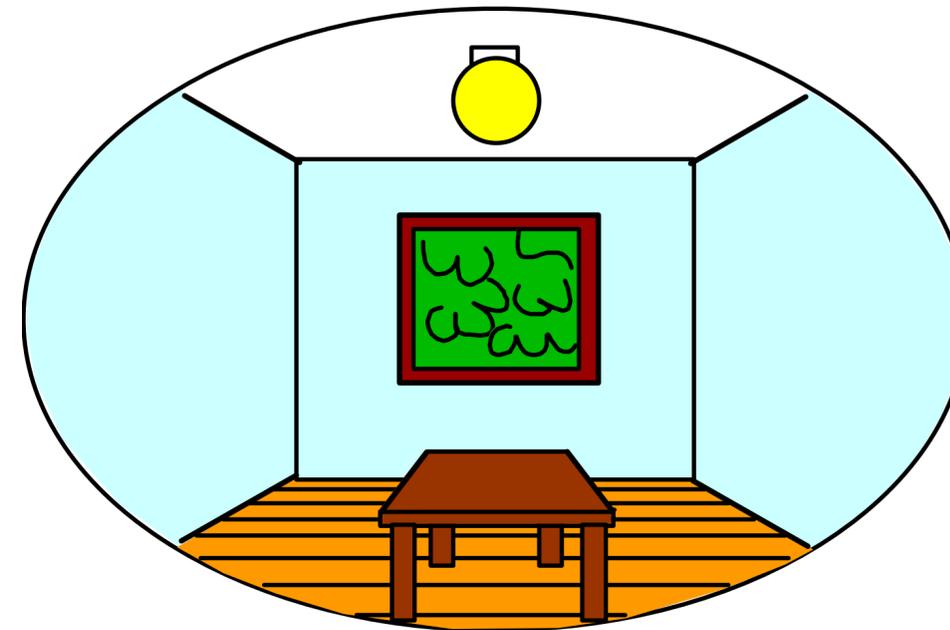
# Projection from 3D to 2D

*3D world*



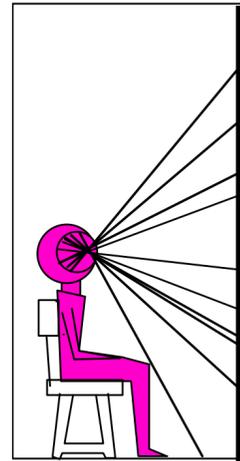
Point of observation

*2D image*



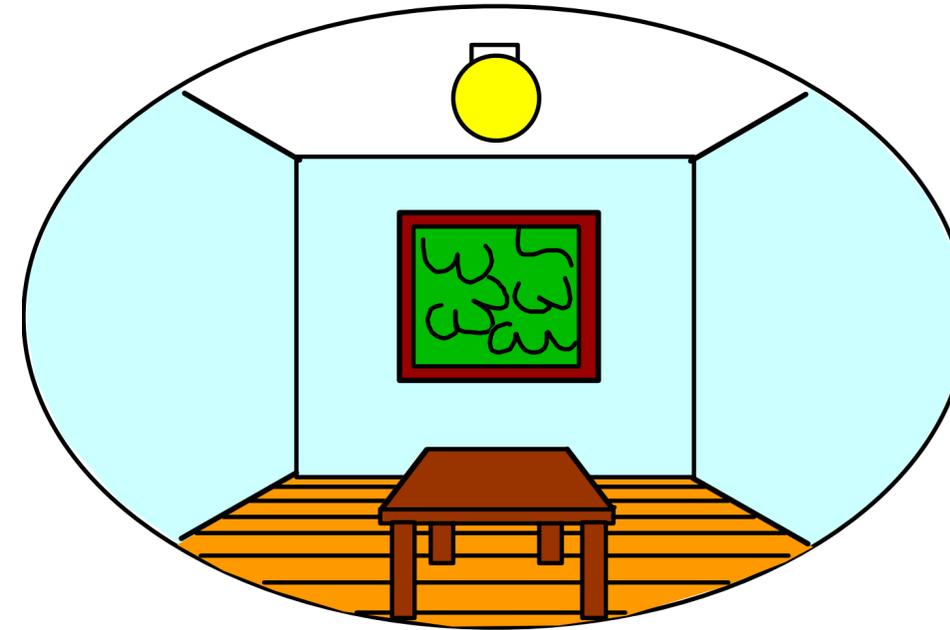
# Projection from 3D to 2D

*3D world*



Painted  
backdrop

*2D image*



# Fooling the eye



# Fooling the eye

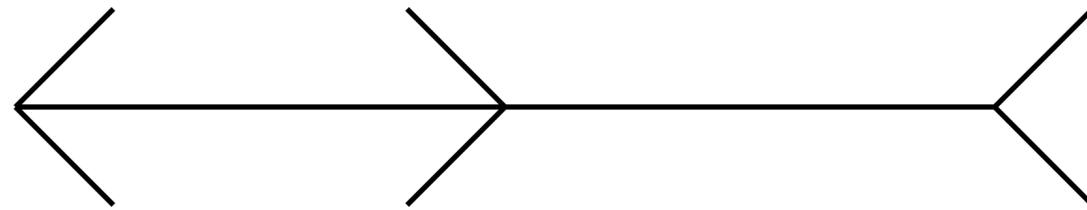
---



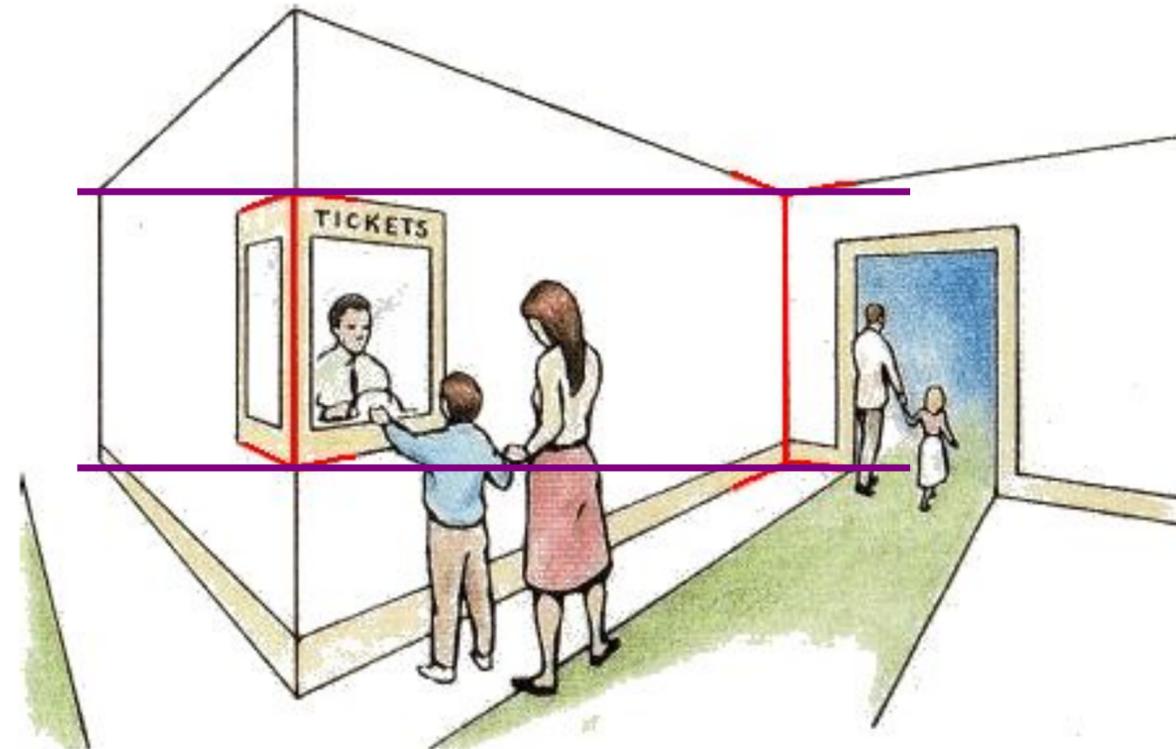
Making of 3D sidewalk art: <http://www.youtube.com/watch?v=3SNYtd0Ayt0>

Source: S. Lazebnik

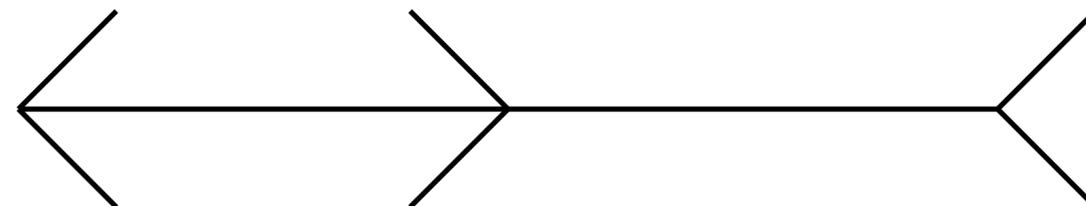
# Müller-Lyer Illusion



# Müller-Lyer Illusion



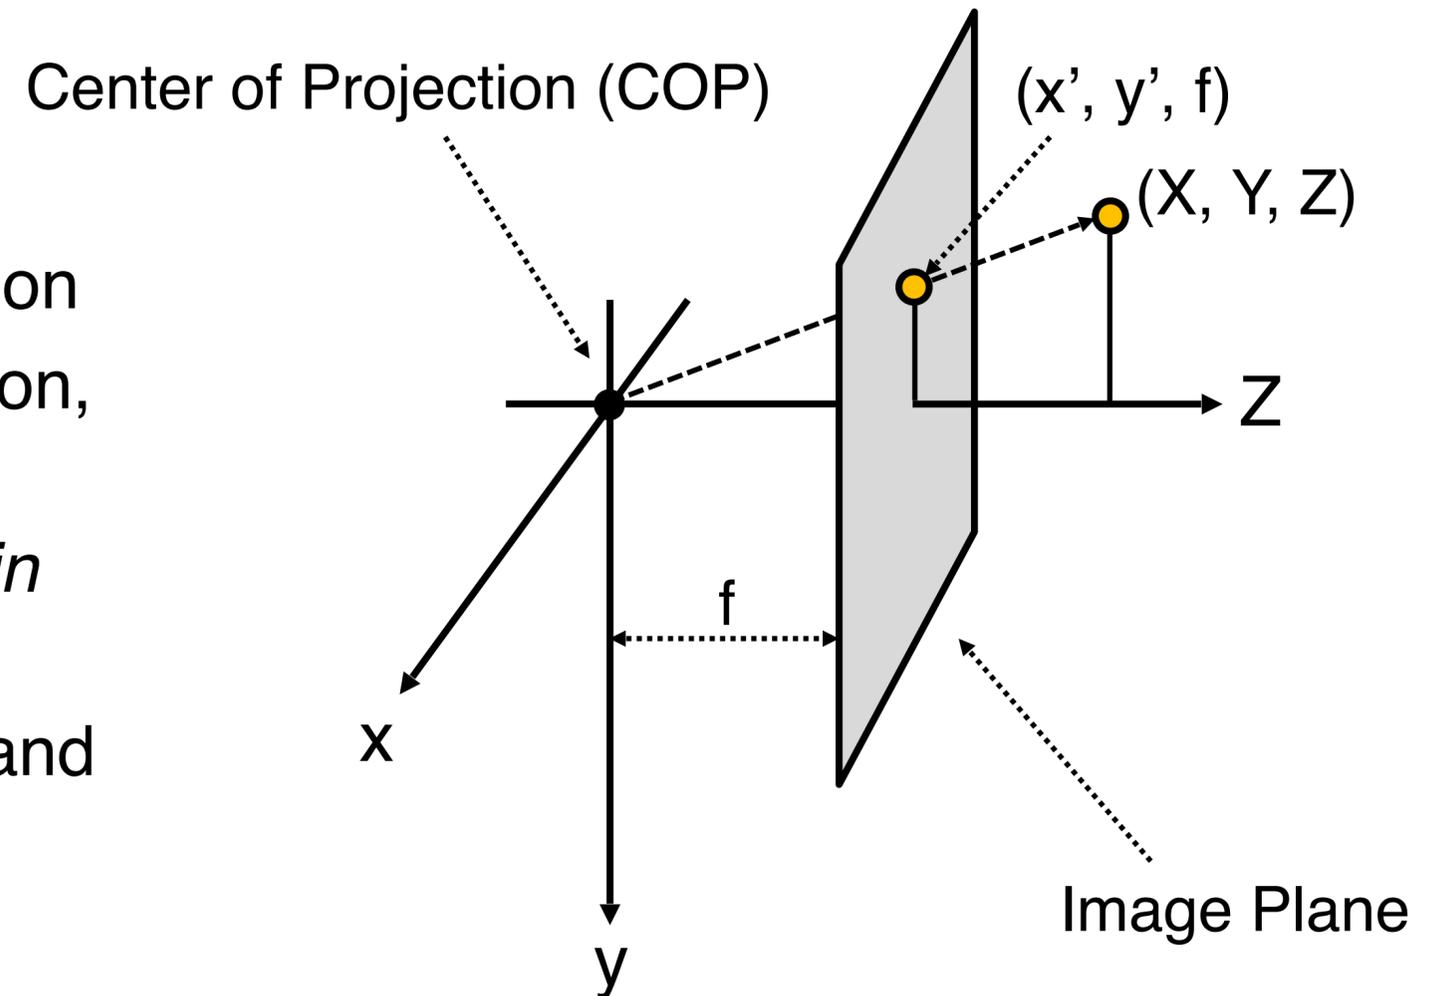
[http://www.michaelbach.de/ot/sze\\_muelue/index.html](http://www.michaelbach.de/ot/sze_muelue/index.html)



# Modeling projection

- The coordinate system

- We use the pinhole model as an approximation
- Put the optical center (aka Center of Projection, or COP) at the origin
- Put the Image Plane (aka Projection Plane) *in front* of the COP
- The camera looks down the *positive* z-axis, and the y-axis points down



# Modeling projection

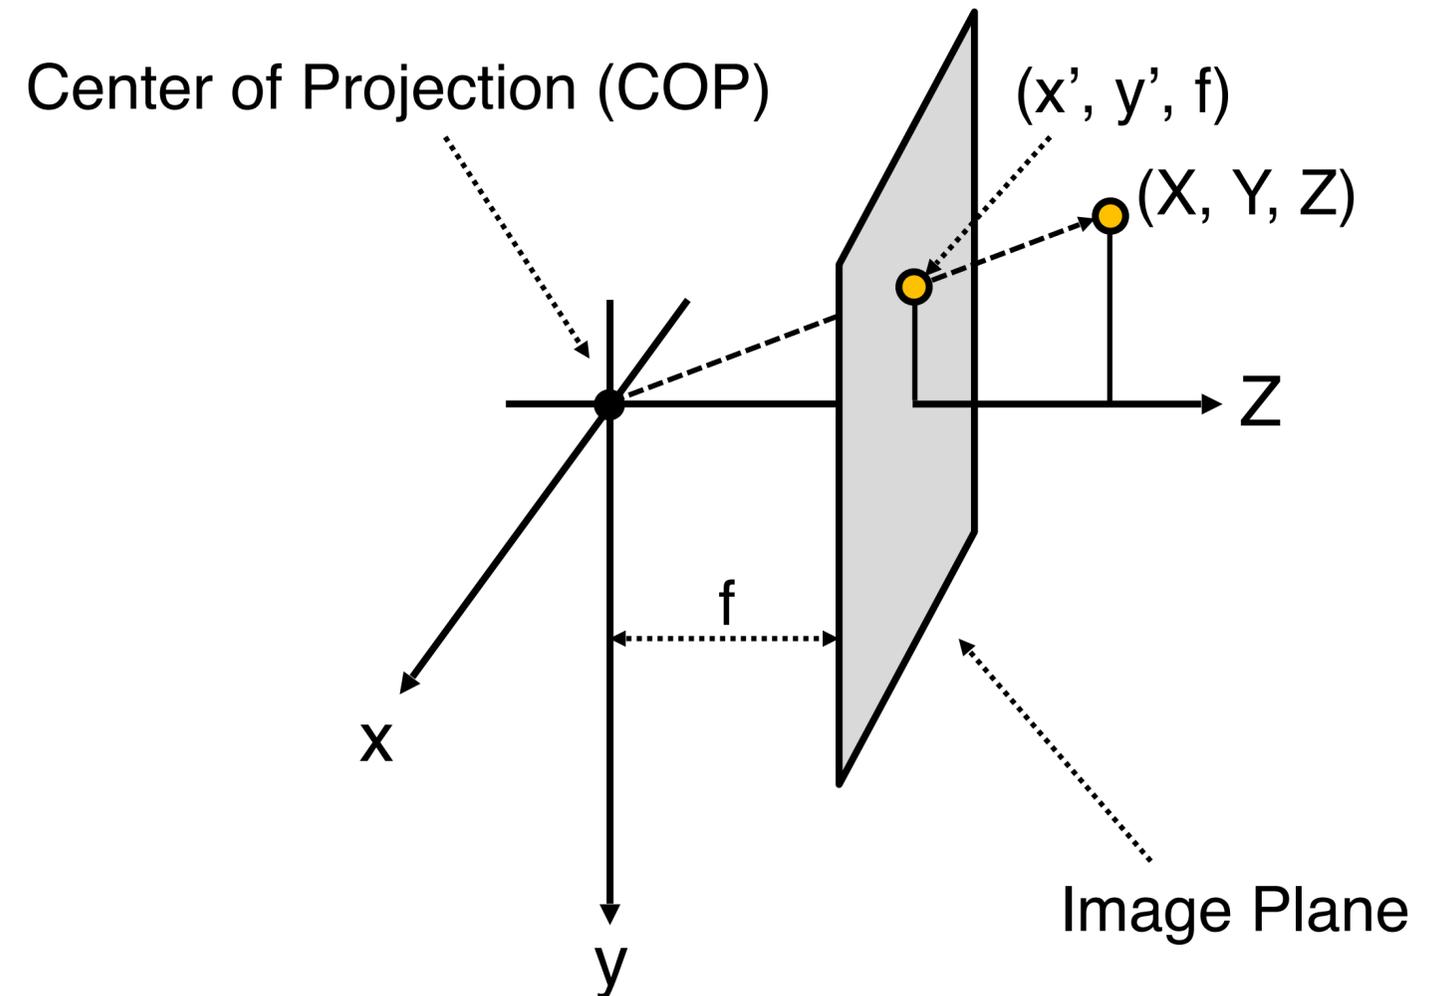
- Projection equations

- Compute intersection with image plane of ray from  $(X, Y, Z)$  to COP
- Derived using similar triangles

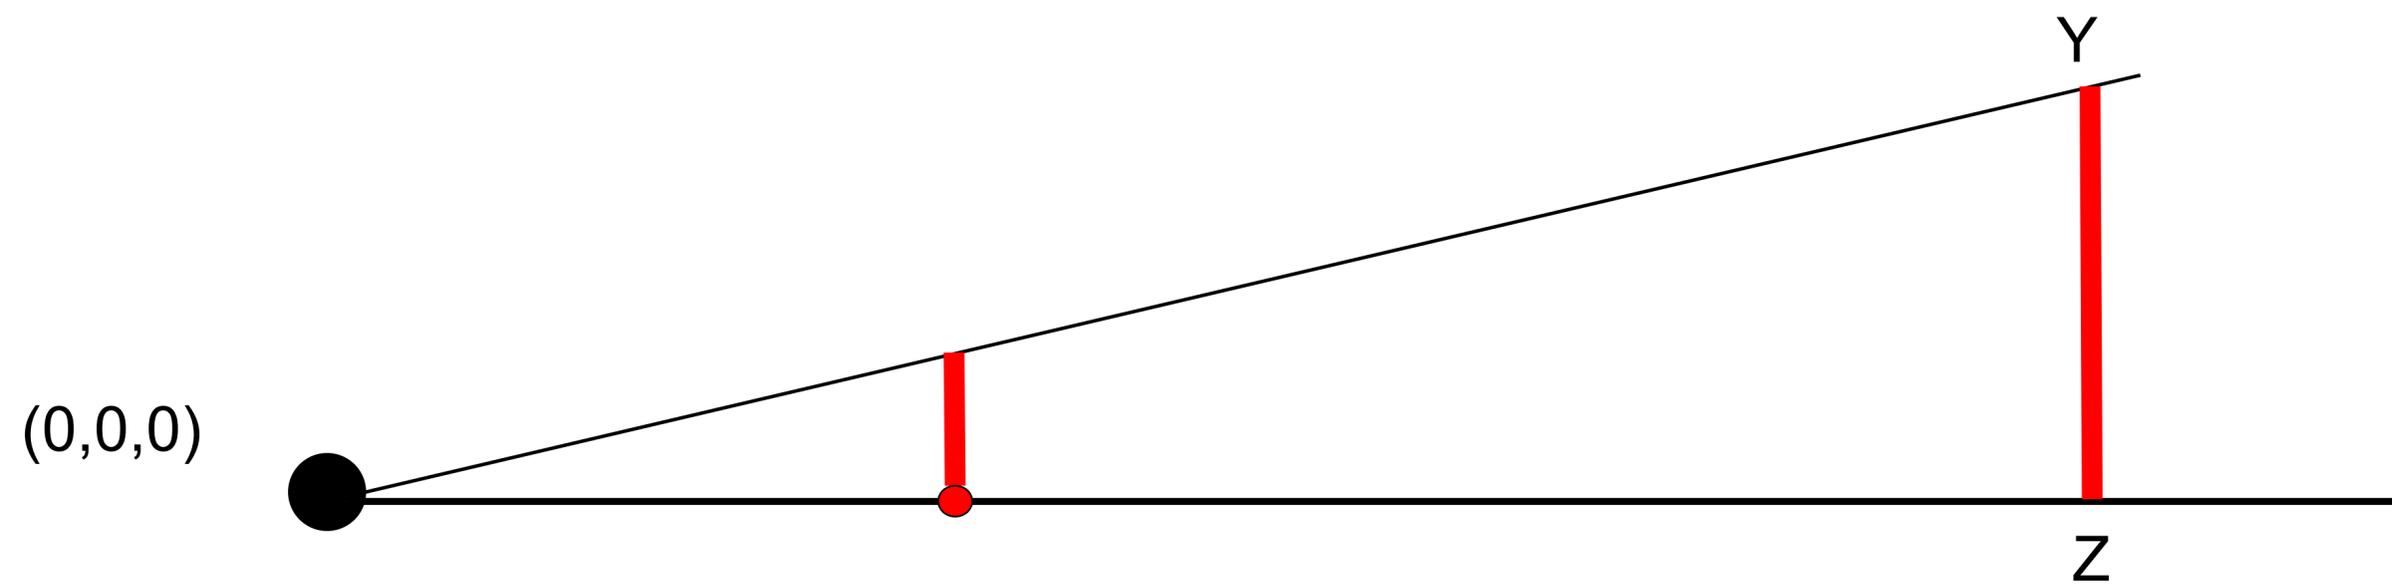
$$(x, y, z) \rightarrow \left( f \frac{x}{z}, f \frac{y}{z}, f \right)$$

- We get the projection by throwing out the last coordinate:

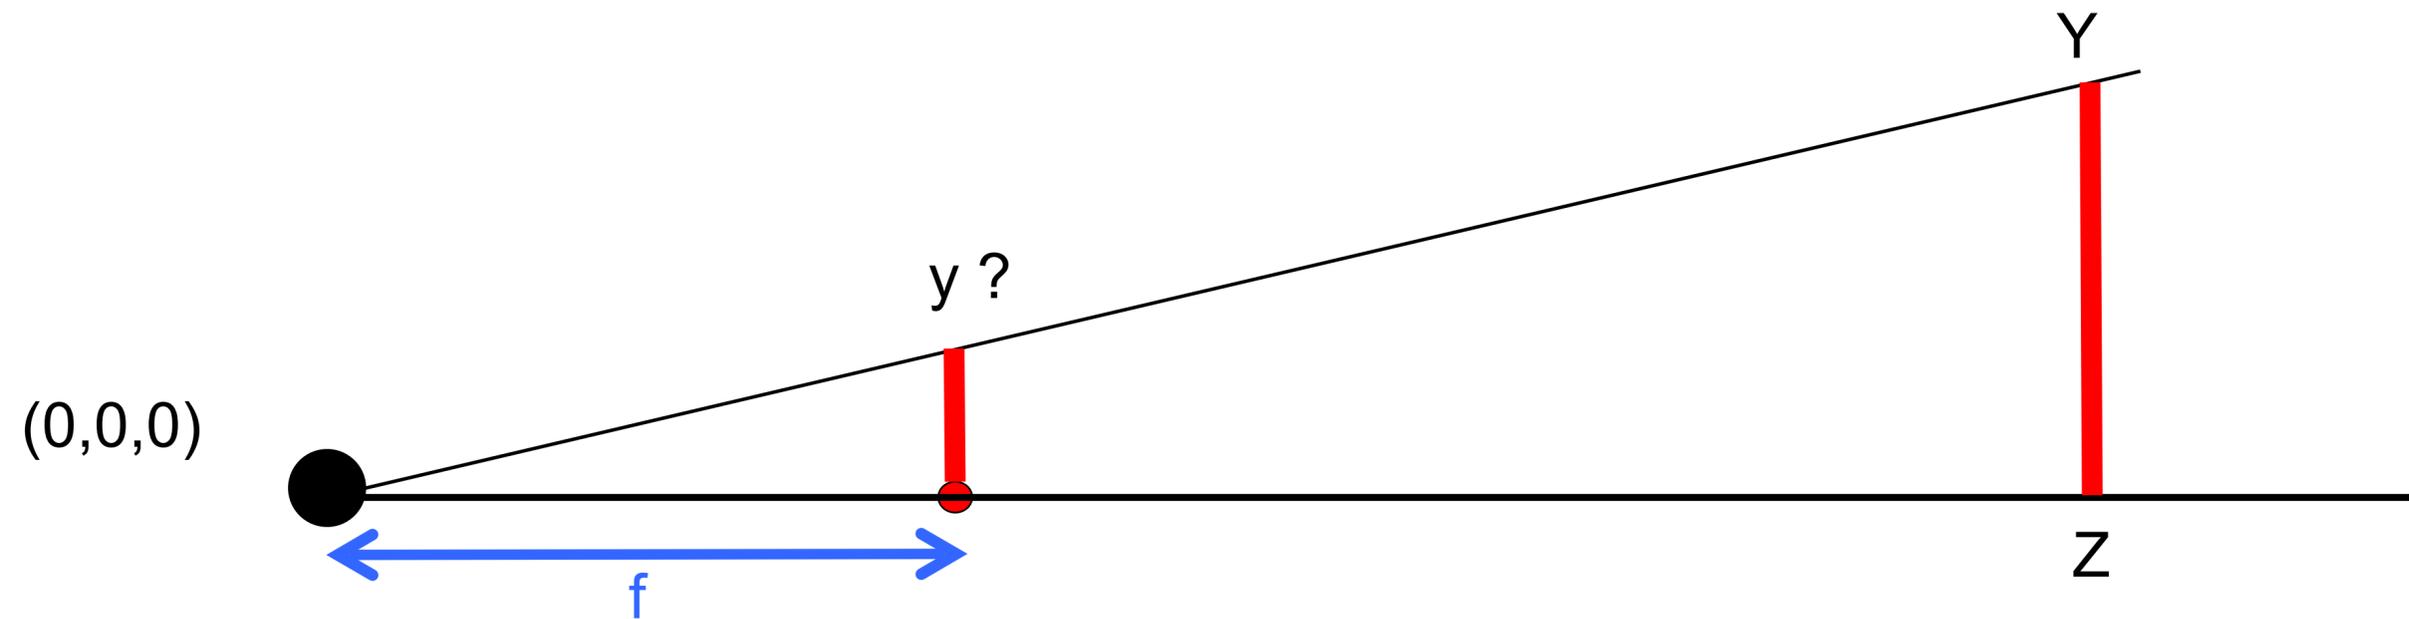
$$(x, y, z) \rightarrow \left( f \frac{x}{z}, f \frac{y}{z} \right)$$



# Perspective projection



# Perspective projection



Similar triangles:  $y / f = Y / Z$

$$y = f Y/Z$$

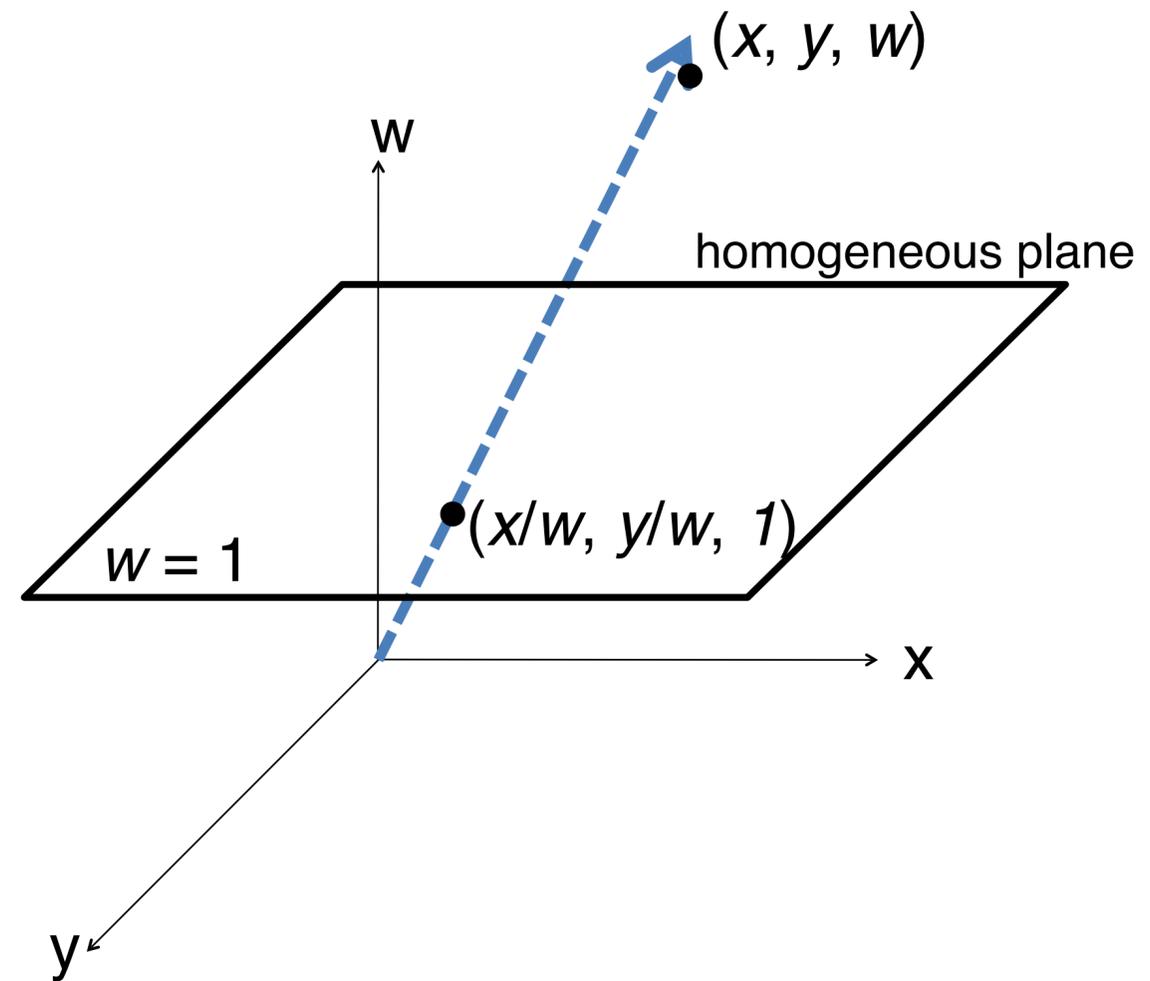
How can we represent this more compactly?

# Homogeneous coordinates

Trick: add one more coordinate:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

homogeneous image  
coordinates



Converting *from* homogeneous coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

# Application: translation with homogeneous coordinates

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix}$$

# Affine transformations

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$



any transformation represented by a 3x3 matrix with last row  $[0 \ 0 \ 1]$  we call an *affine* transformation

$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

# Examples of Affine Transformations

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Translate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Scale

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

2D *in-plane* rotation

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & sh_x & 0 \\ sh_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Shear

# Perspective Projection

Projection is a matrix multiply using homogeneous coordinates:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/f & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z/f \end{bmatrix} \Rightarrow \left( f \frac{x}{z}, f \frac{y}{z} \right)$$

divide by third coordinate

This is known as **perspective projection**

- The matrix is the **projection matrix**

# Perspective Projection

How does scaling the projection matrix change the transformation?

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/f & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z/f \end{bmatrix} \Rightarrow \left( f \frac{x}{z}, f \frac{y}{z} \right)$$

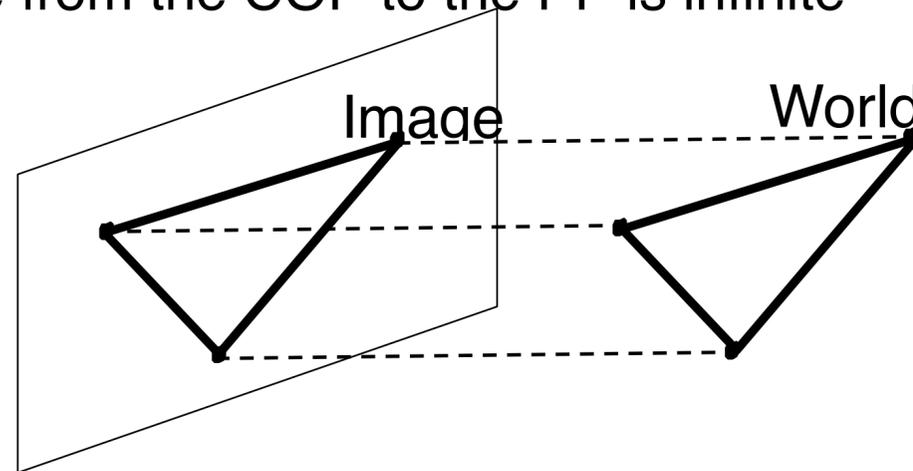
What if we scale by  $f$ ?

$$\begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} fx \\ fy \\ z \end{bmatrix} \Rightarrow \left( f \frac{x}{z}, f \frac{y}{z} \right)$$

Scaling a projection matrix produces an equivalent projection matrix!

# Orthographic projection

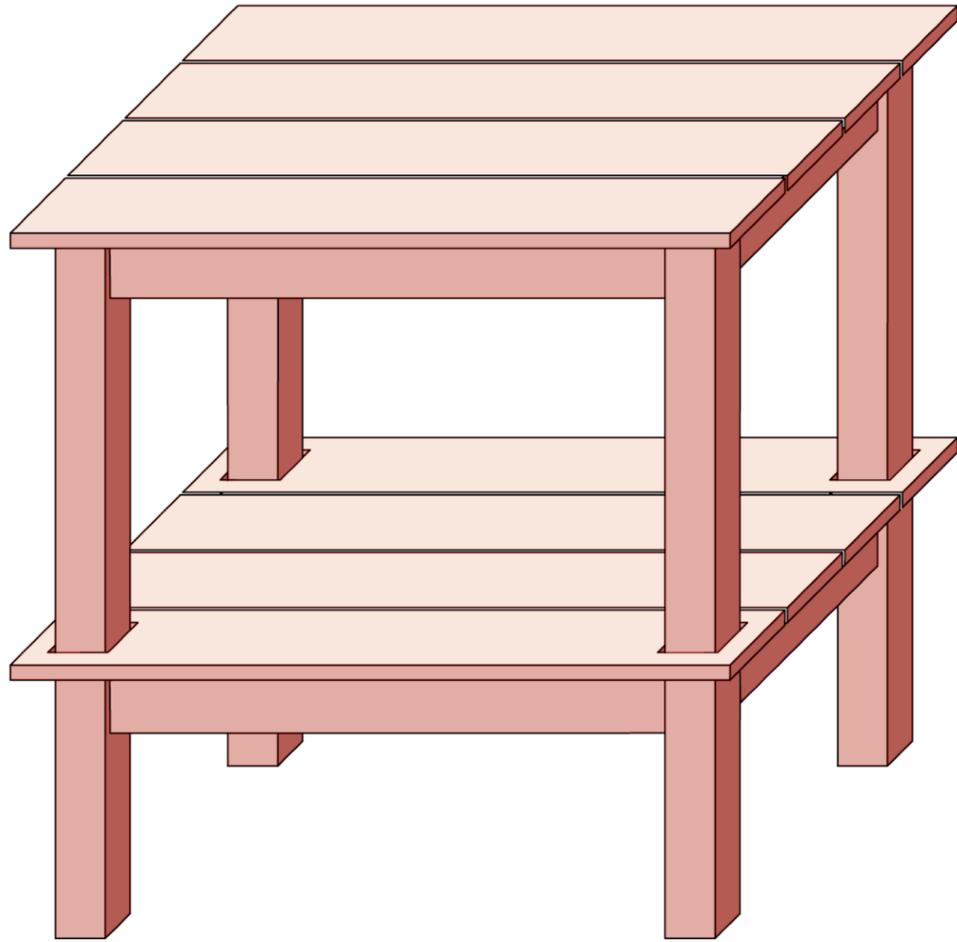
- Special case of perspective projection
  - Distance from the COP to the PP is infinite



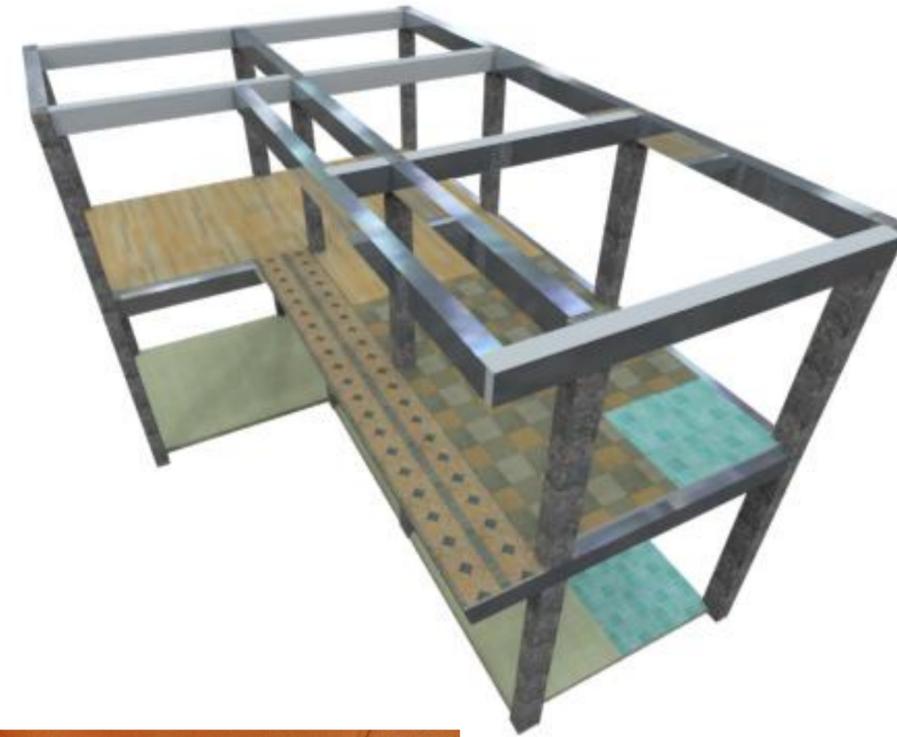
- Good approximation for telephoto optics
- Also called “parallel projection”:  $(x, y, z) \rightarrow (x, y)$
- What’s the projection matrix?

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \Rightarrow (x, y)$$

# Orthographic projection



# Perspective projection

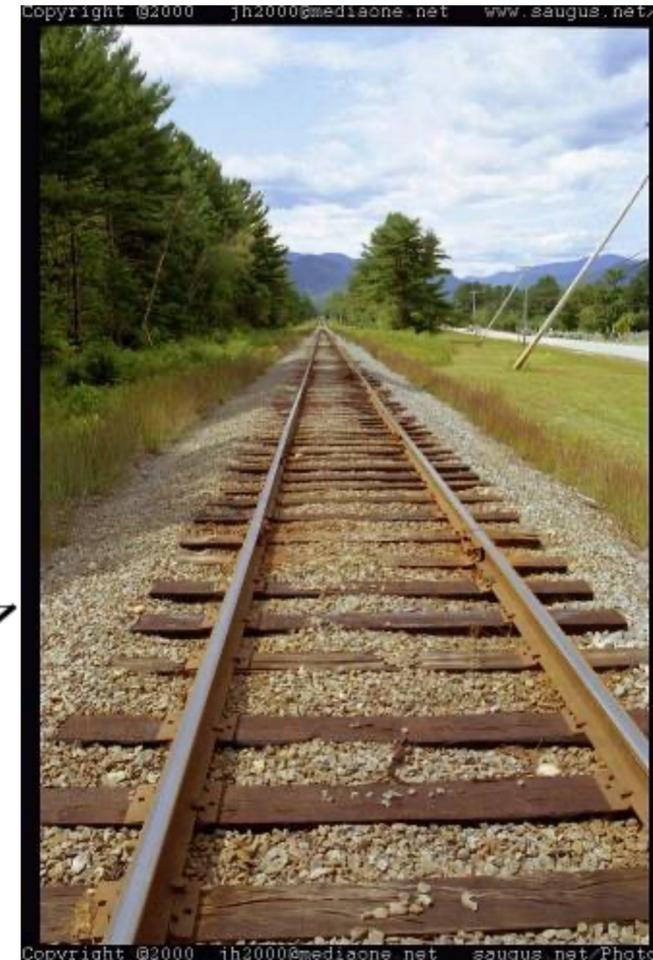
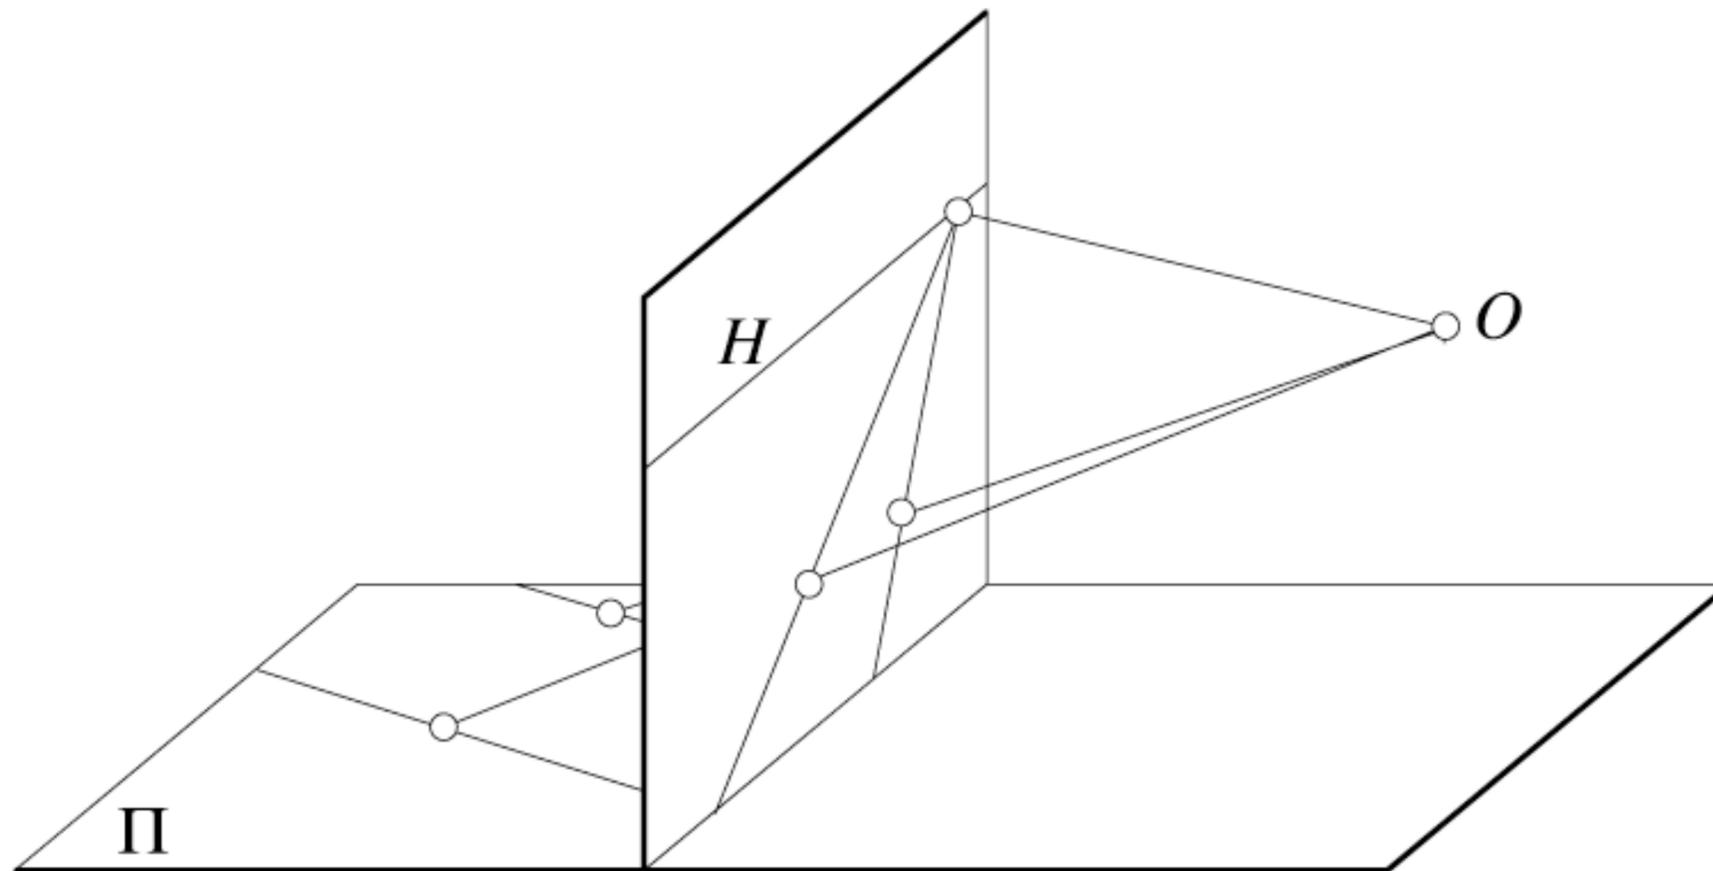


# Projection properties

- Many-to-one: any points along same ray map to same point in image
- Points  $\rightarrow$  points
- Lines  $\rightarrow$  lines (collinearity is preserved)
  - But line through focal point projects to a point
- Planes  $\rightarrow$  planes (or half-planes)
  - But plane through focal point projects to line

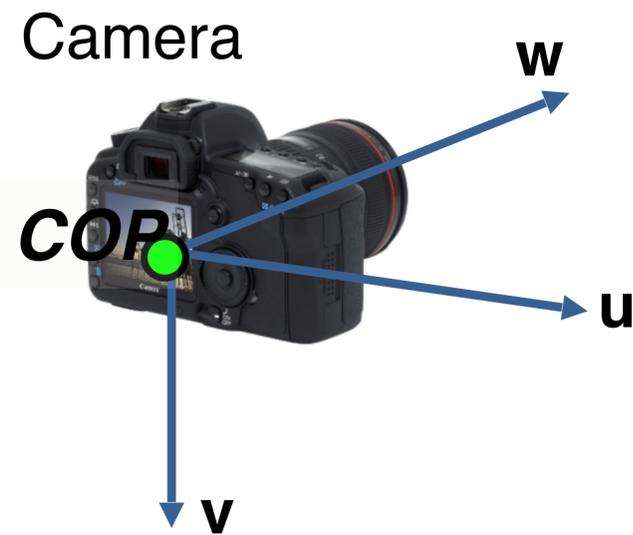
# Projection properties

- Parallel lines converge at a vanishing point
  - Each direction in space has its own vanishing point
  - But lines parallel to the image plane remain parallel



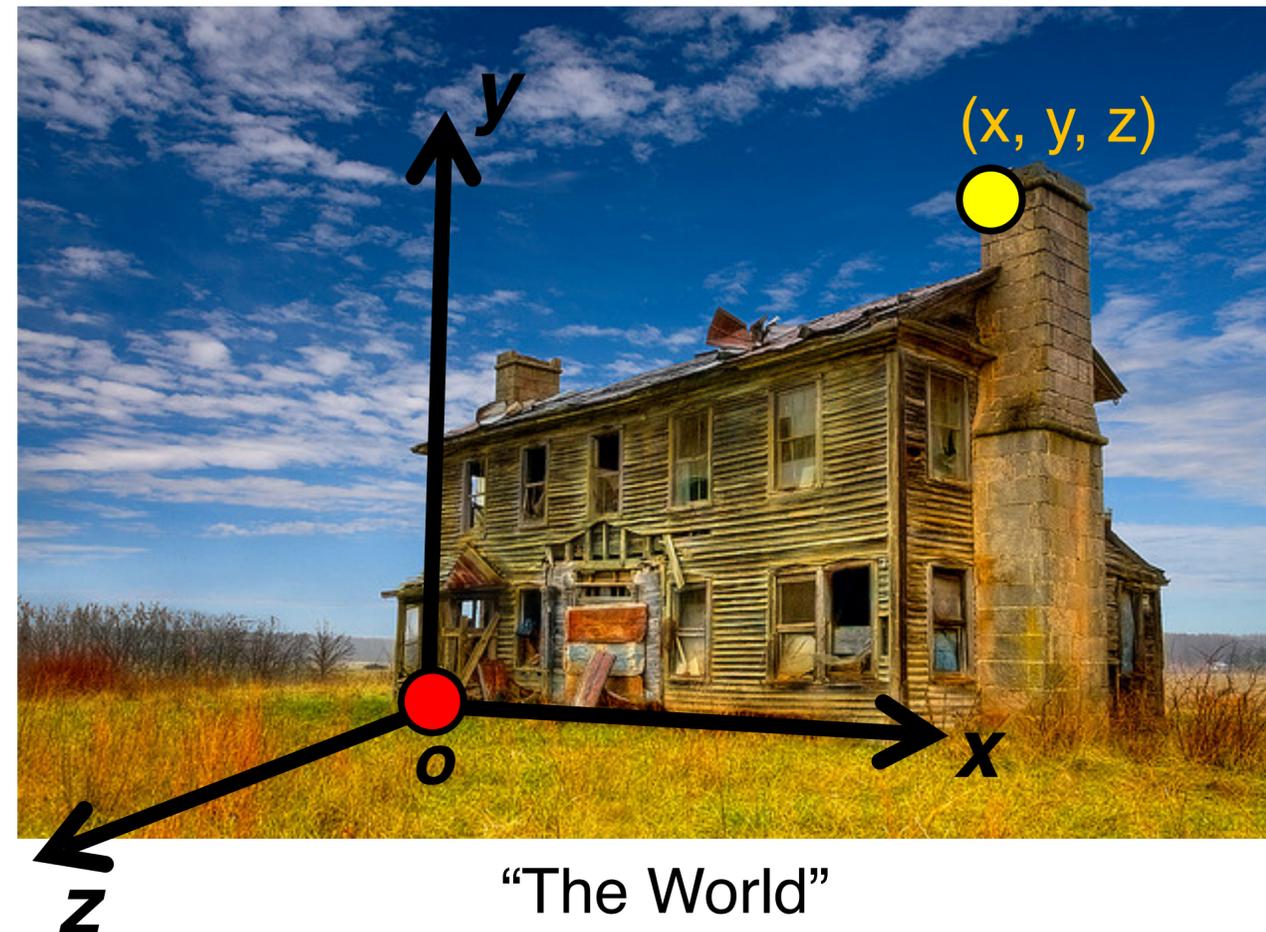
# Camera parameters

- How can we model the geometry of a camera?



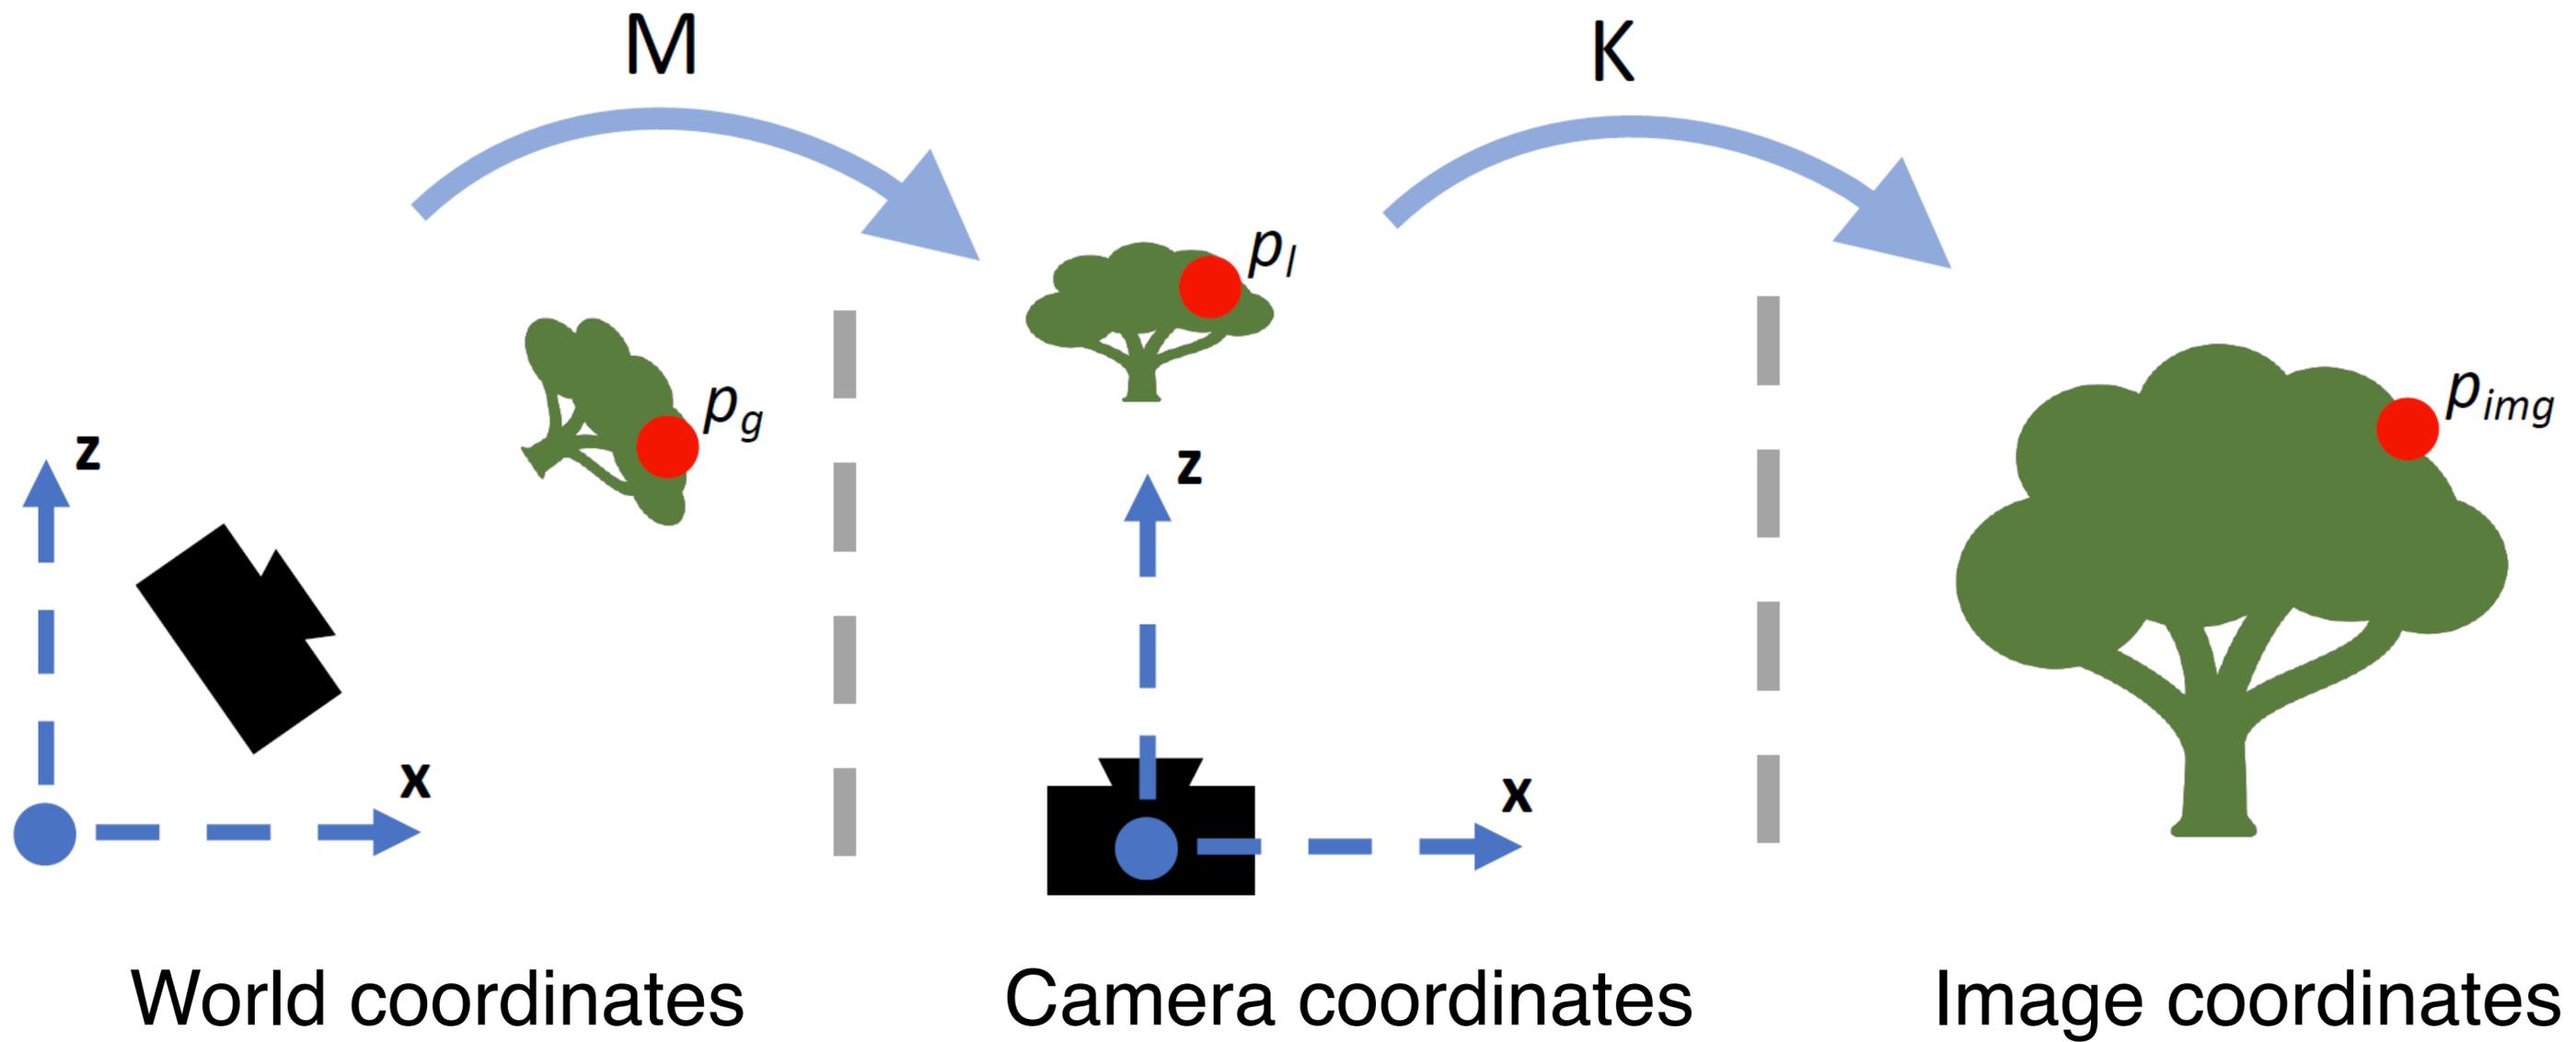
Three important coordinate systems:

1. *World* coordinates
2. *Camera* coordinates
3. *Image* coordinates



How do we project a given world point  $(x, y, z)$  to an image point?

# Coordinate frames



# Camera parameters

To project a point  $(x, y, z)$  in *world* coordinates into a camera

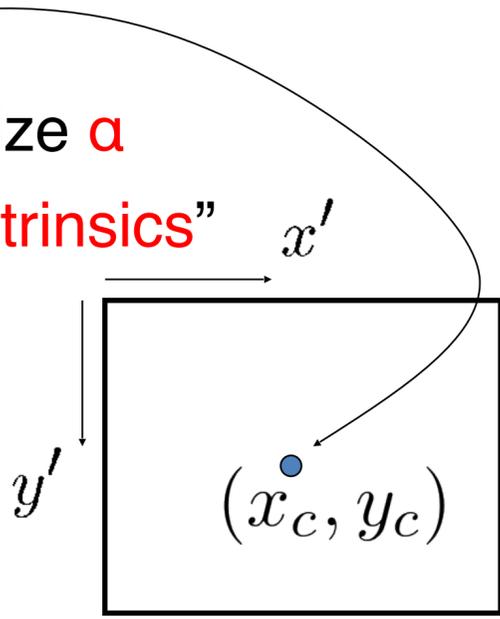
- First transform  $(x, y, z)$  into *camera* coordinates
- Need to know
  - Camera position (in world coordinates)
  - Camera orientation (in world coordinates)
- Then project into the image plane to get *image (pixel) coordinates*
  - Need to know camera *intrinsic*s

# Camera parameters

A camera is described by several parameters

- Translation **T** of the optical center from the origin of world coords
- Rotation **R** of the image plane
- focal length **f**, principal point **(c<sub>x</sub>, c<sub>y</sub>)**, pixel aspect size **α**
- blue parameters are called “**extrinsics**,” red are “**intrinsics**”

Projection equation

$$\mathbf{x} = \begin{bmatrix} sx \\ sy \\ s \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{\Pi X}$$


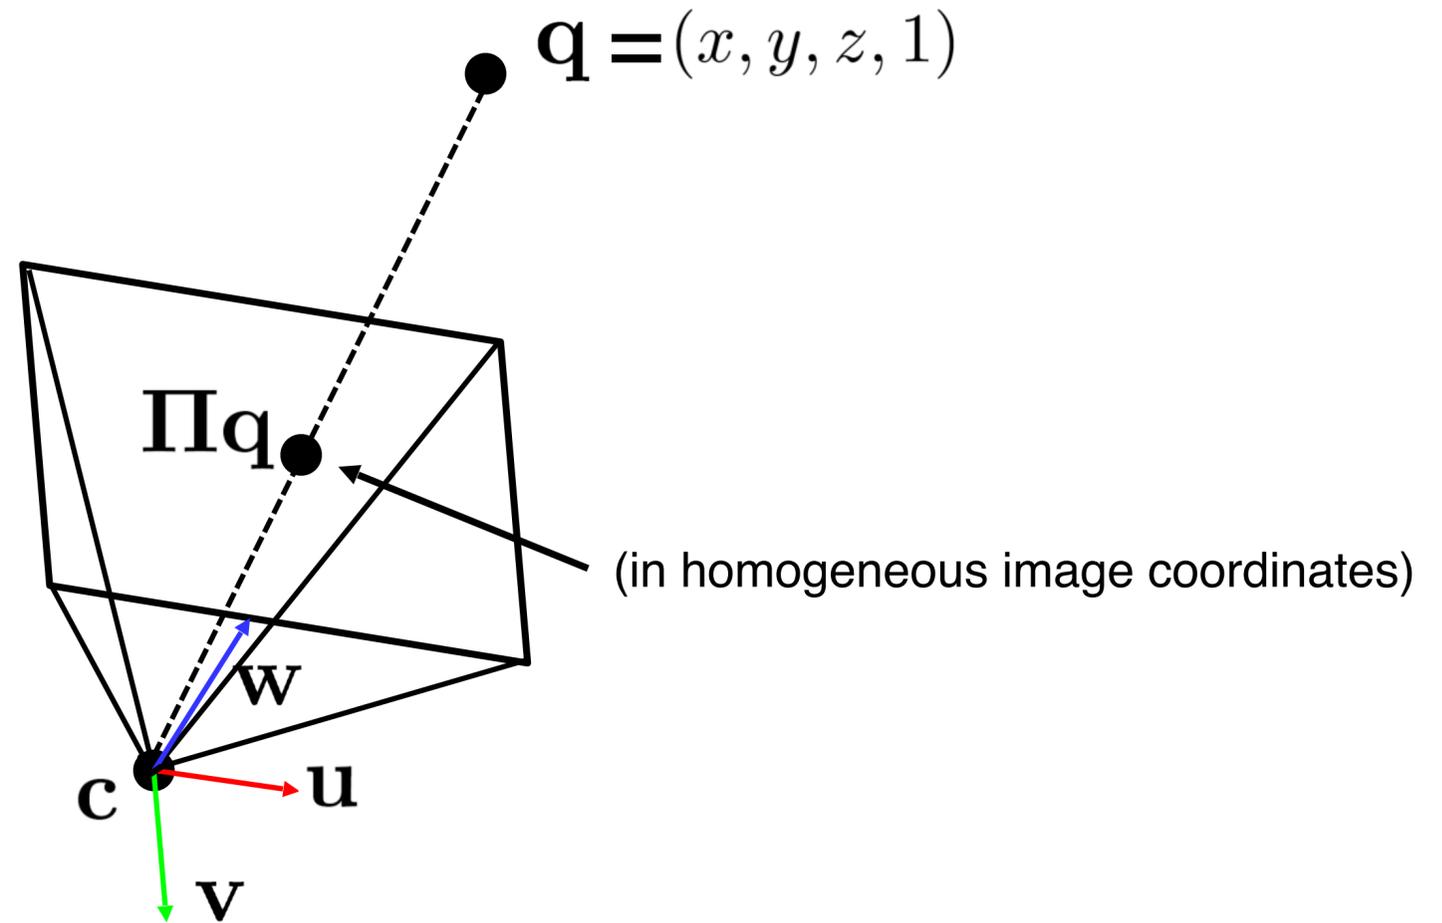
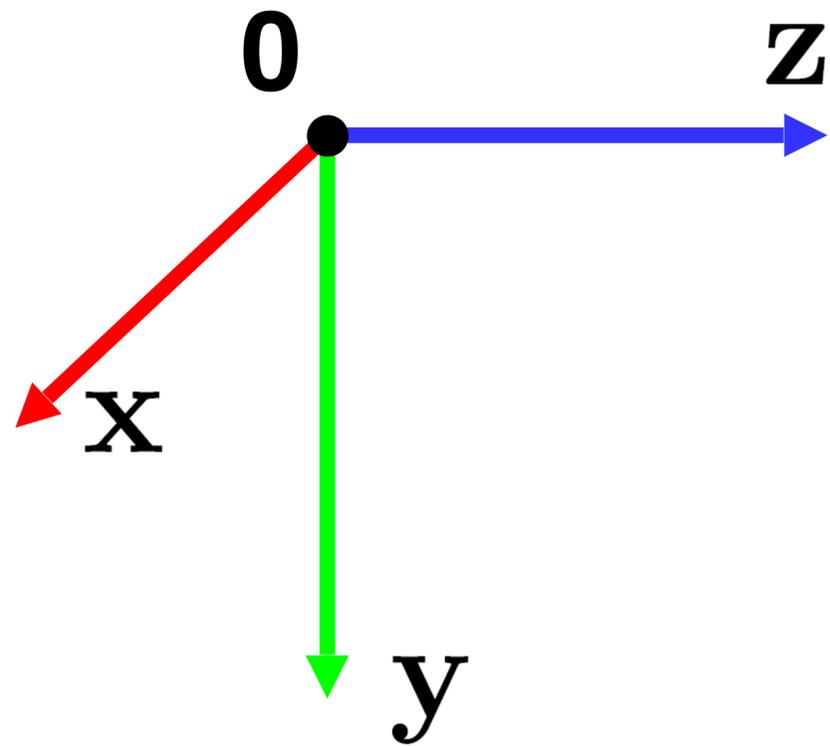
- The projection matrix models the cumulative effect of all parameters
- Useful to decompose into a series of operations

$$\mathbf{\Pi} = \underbrace{\begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\text{intrinsics}} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{T}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{translation}}$$

identity matrix

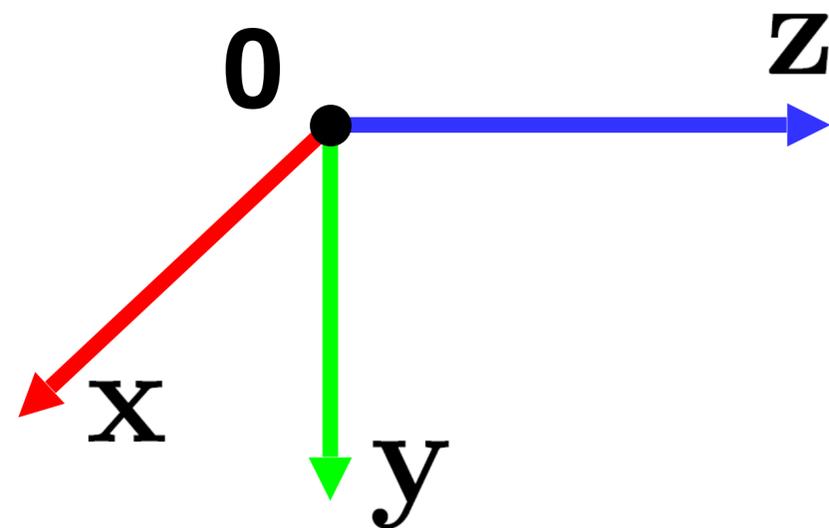
- The definitions of these parameters are **not** completely standardized
  - especially intrinsics—varies from one book to another

# Projection matrix

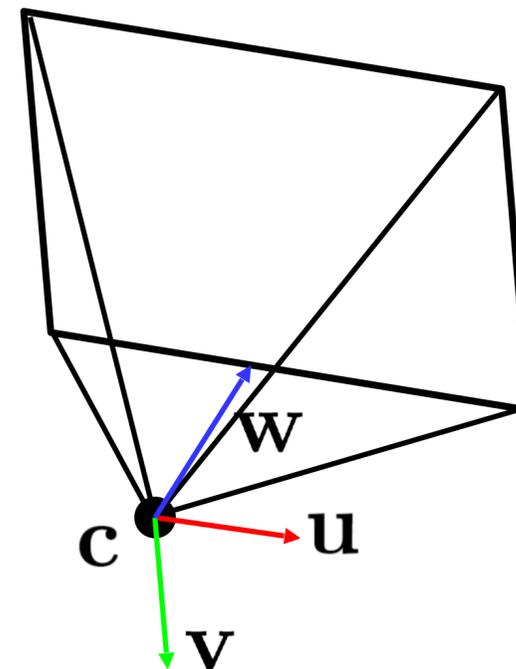


# Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

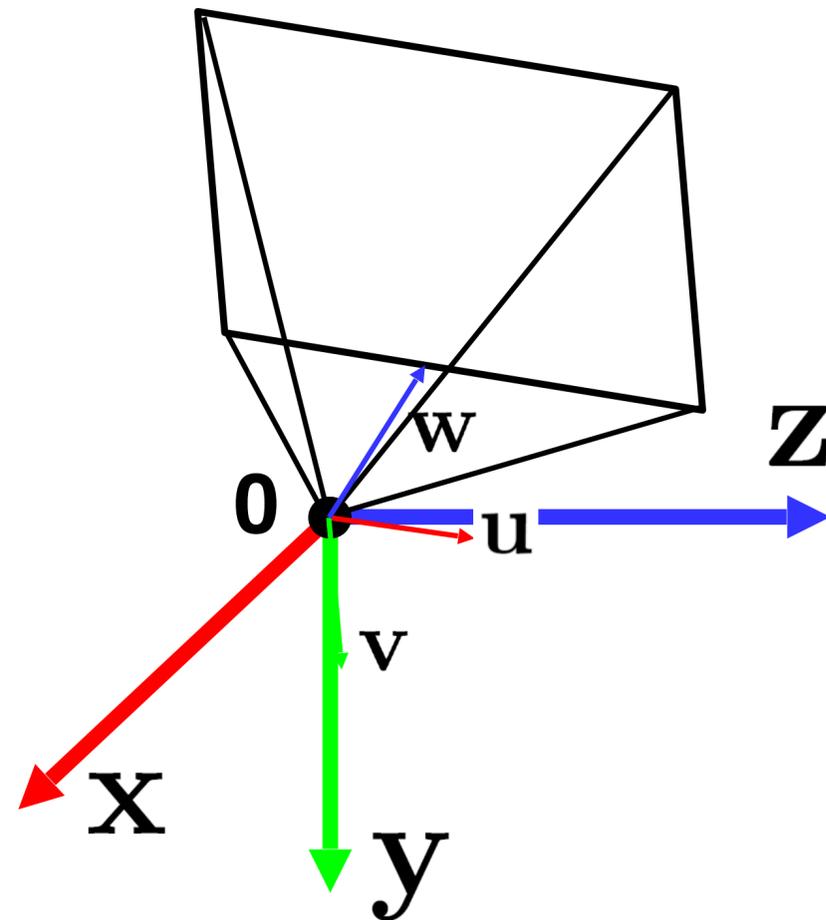


Step 1: Translate by  $-\mathbf{c}$



# Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



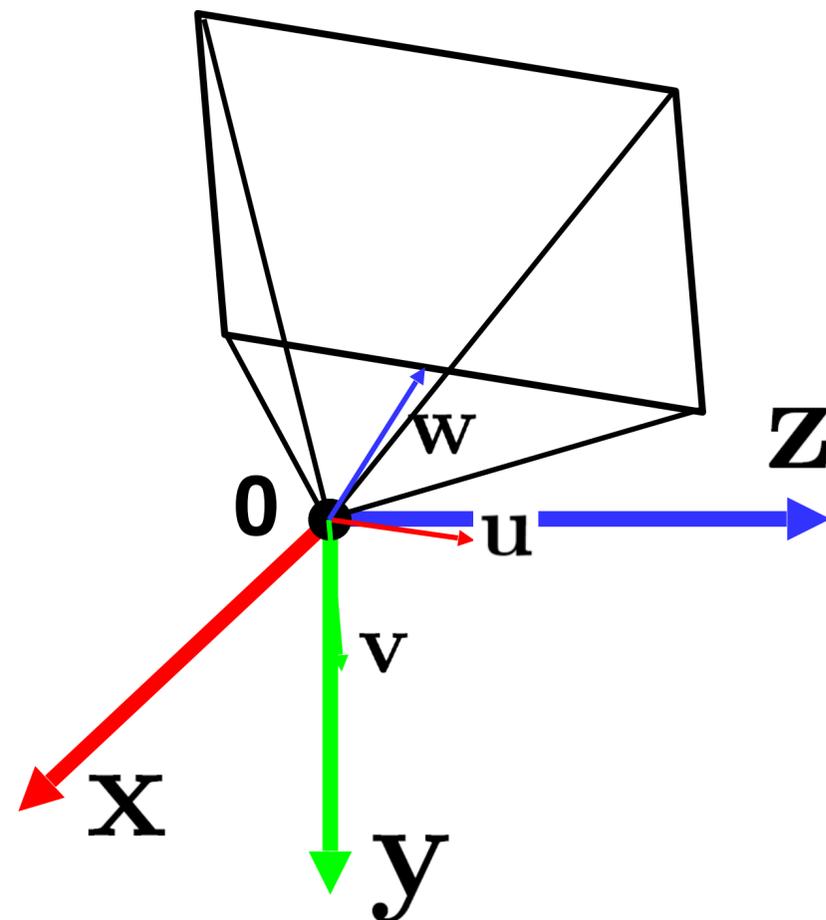
Step 1: Translate by  $-\mathbf{c}$

How do we represent translation as a matrix multiplication?

$$\mathbf{T} = \begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

# Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



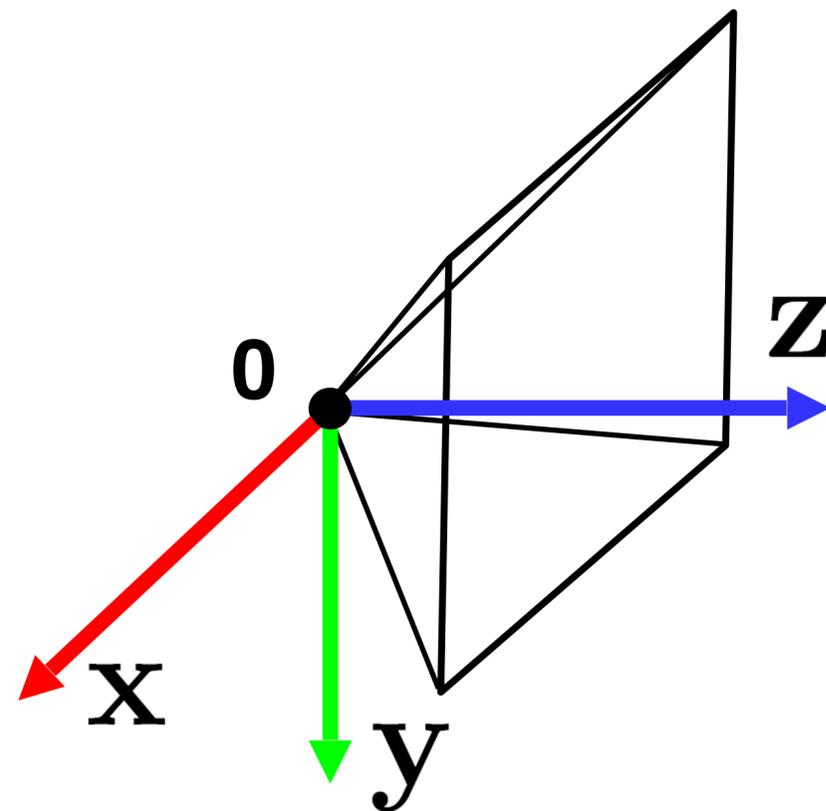
Step 1: Translate by  $-\mathbf{c}$   
Step 2: Rotate by  $\mathbf{R}$

$\mathbf{R} = \begin{bmatrix} \mathbf{u}^T \\ \mathbf{v}^T \\ \mathbf{w}^T \end{bmatrix}$

3x3 rotation matrix

# Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



Step 1: Translate by  $-\mathbf{c}$

Step 2: Rotate by  $\mathbf{R}$

$$\mathbf{R} = \begin{bmatrix} \mathbf{u}^T \\ \mathbf{v}^T \\ \mathbf{w}^T \end{bmatrix}$$

(with extra row/column of  $[0 \ 0 \ 0 \ 1]$ )

# Perspective projection

$$\underbrace{\begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{K}} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

**K**  
(intrinsics) (converts from 3D rays in camera coordinate system to pixel coordinates)

in general,  $\mathbf{K} = \begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}$  (upper triangular matrix)

$\alpha$ : **aspect ratio** (1 unless pixels are not square)

$s$ : **skew** (0 unless pixels are shaped like rhombi/parallelograms)

$(c_x, c_y)$ : **principal point** ((w/2, h/2) unless optical axis doesn't intersect projection plane at image center)

# Typical intrinsics matrix

$$\mathbf{K} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

- **2D affine transform** corresponding to a scale by  $f$  (focal length) and a translation by  $(c_x, c_y)$  (principal point)
- Maps 3D rays to 2D pixels

# Focal length

- Can think of as “zoom”



24mm



50mm



200mm



800mm



- Also related to *field of view*

# Changing focal length



Wide angle



Standard



Telephoto



<http://petapixel.com/2013/01/11/how-focal-length-affects-your-subjects-apparent-weight-as-seen-with-a-cat/>

# Projection matrix

$$\mathbf{\Pi} = \mathbf{K} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{translation}}$$

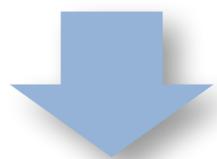
The **K** matrix converts 3D rays in the camera's coordinate system to 2D image points in image (pixel) coordinates.

This part converts 3D points in world coordinates to 3D rays in the camera's coordinate system. There are 6 parameters represented (3 for position/translation, 3 for rotation).

# Projection matrix

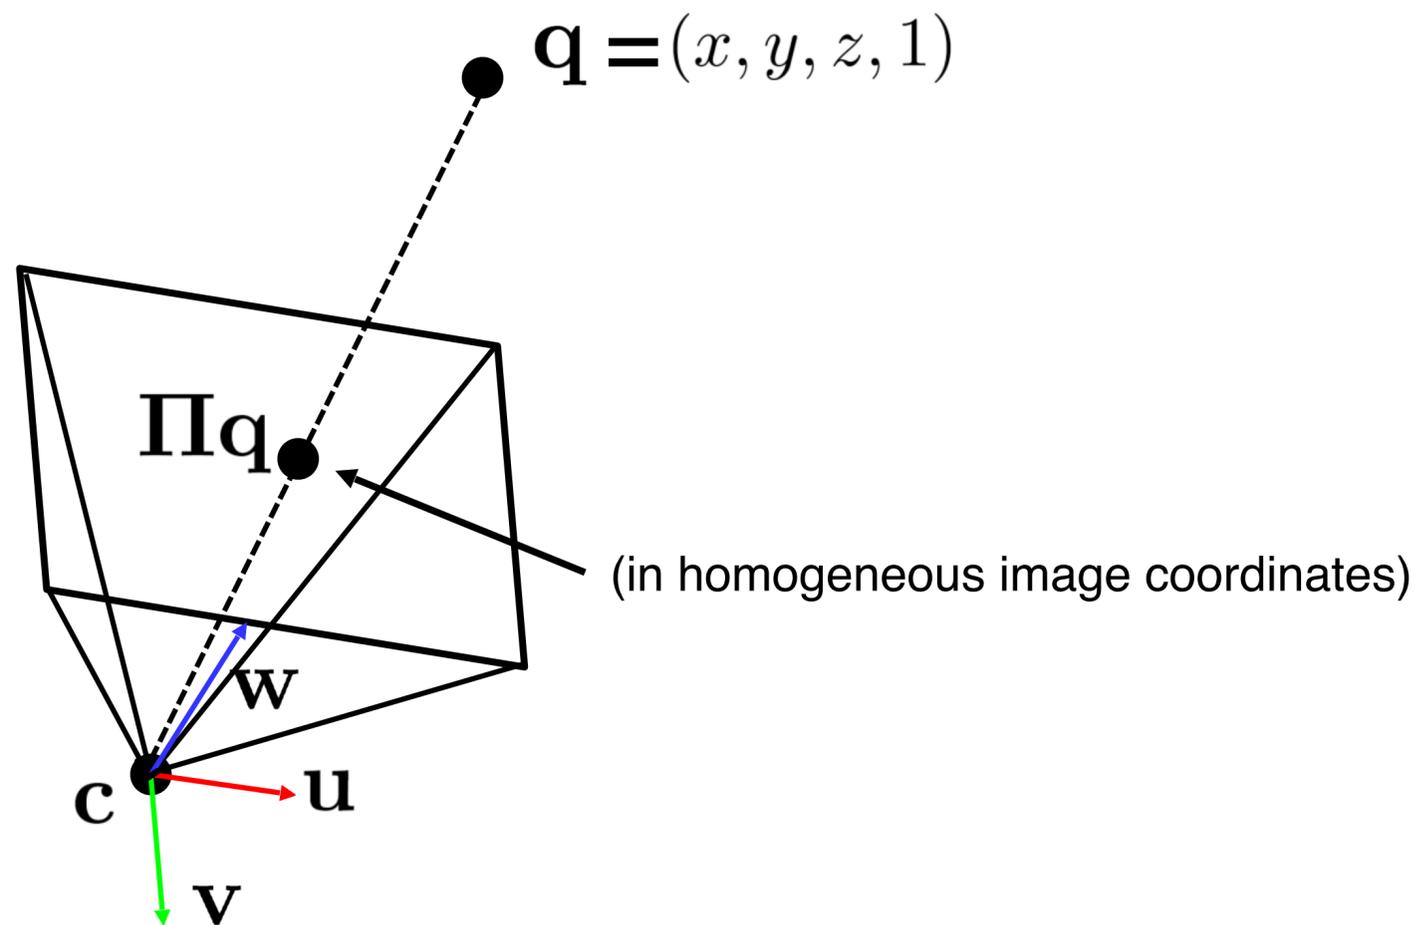
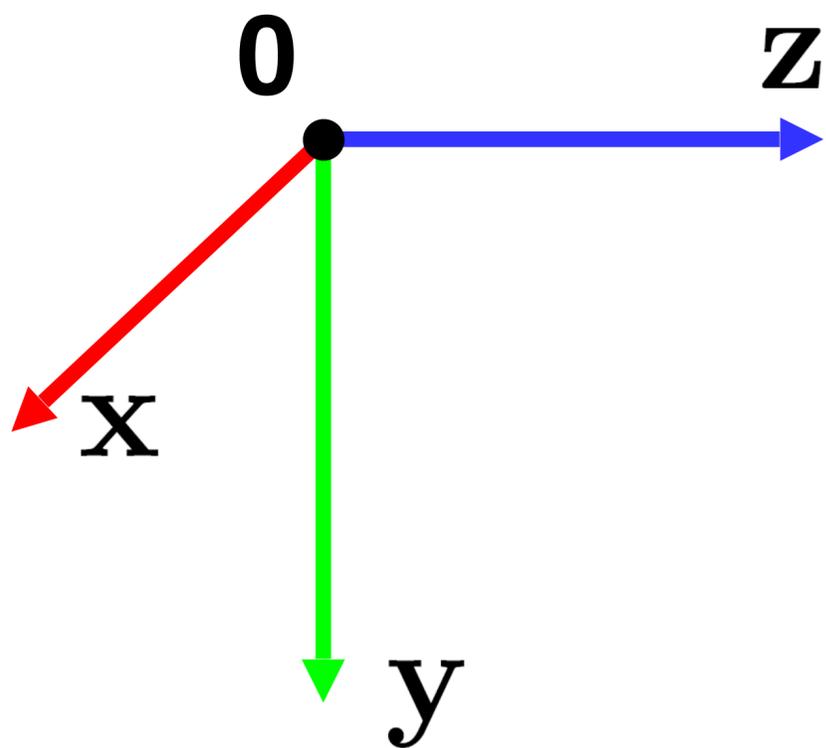
$$\mathbf{\Pi} = \mathbf{K} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{translation}}$$

$$\left[ \mathbf{R} \mid \underbrace{-\mathbf{R}\mathbf{c}}_{\text{(sometimes called } \mathbf{t})} \right]$$

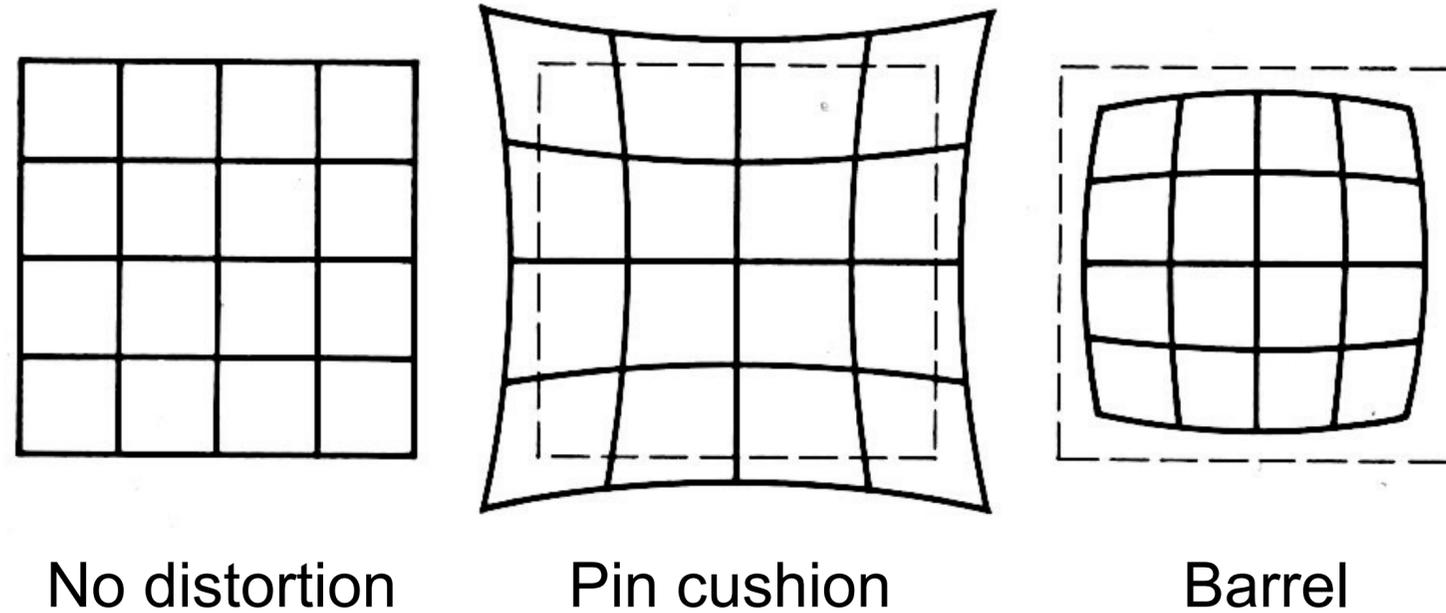


$$\mathbf{\Pi} = \mathbf{K} \left[ \mathbf{R} \mid -\mathbf{R}\mathbf{c} \right]$$

# Projection matrix



# Distortion



- Radial distortion of the image
  - Caused by imperfect lenses
  - Deviations are most noticeable for rays that pass through the edge of the lens



# Modeling distortion

$$\begin{array}{l} (\hat{x}, \hat{y}, \hat{z}) \\ \text{Project} \\ \text{to "normalized"} \\ \text{image coordinates} \end{array} \quad \begin{array}{l} x'_n = \hat{x} / \hat{z} \\ y'_n = \hat{y} / \hat{z} \end{array}$$

$$\begin{array}{l} \text{Apply radial distortion} \end{array} \quad \begin{array}{l} r^2 = x'^2_n + y'^2_n \\ x'_d = x'_n (1 + \kappa_1 r^2 + \kappa_2 r^4) \\ y'_d = y'_n (1 + \kappa_1 r^2 + \kappa_2 r^4) \end{array}$$

$$\begin{array}{l} \text{Apply focal length} \\ \text{translate image center} \end{array} \quad \begin{array}{l} x' = f x'_d + x_c \\ y' = f y'_d + y_c \end{array}$$

- To model lens distortion
  - Use above projection operation instead of standard projection matrix multiplication

# Correcting radial distortion



from [Helmut Dersch](#)

**Next class: More geometry!**