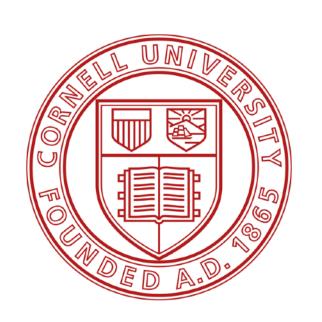
Lecture 18: Image formation

CS 5670: Introduction to Computer Vision



Physically-based computer vision

Use physical constraints to solve perception problems.

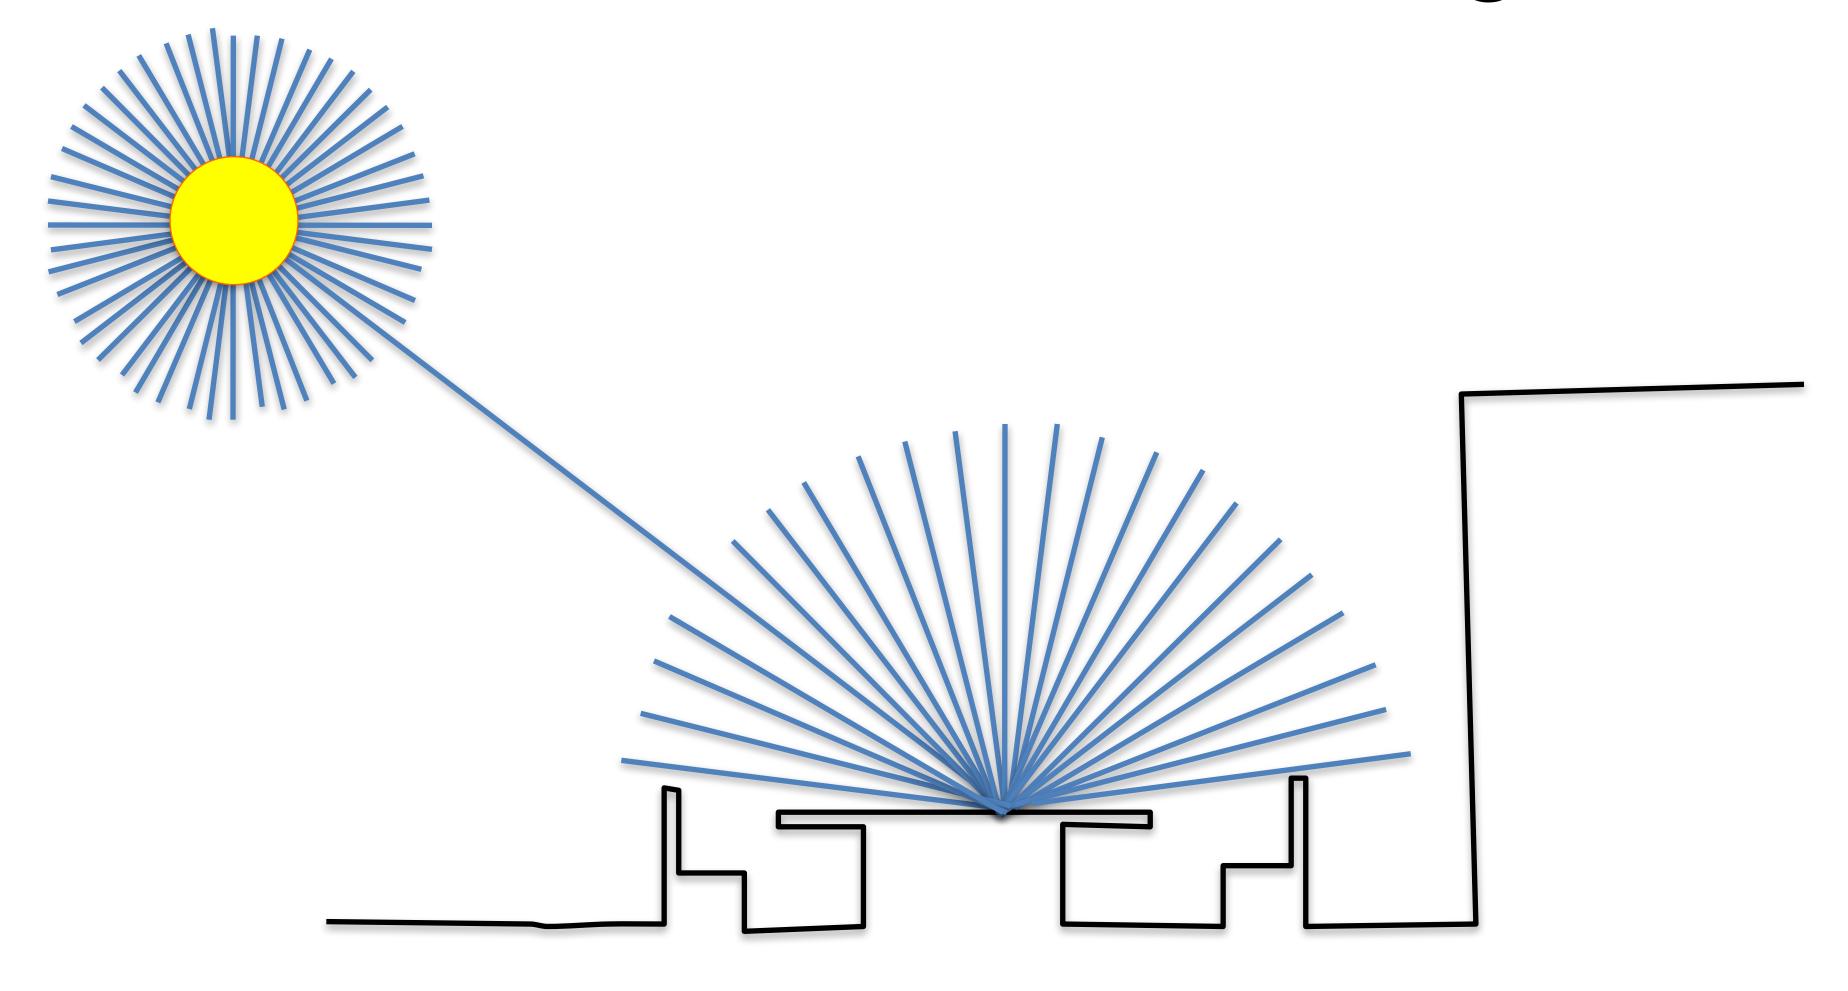
Why study this?

- Physically-based methods often (still!) outperform learned systems
- Hybrid approaches: learning methods that exploit physical constraints
- Helps us understand limitations of visual perception
- Many domains where you do not want to rely much on prior knowledge (e.g., computational imaging for astronomy)

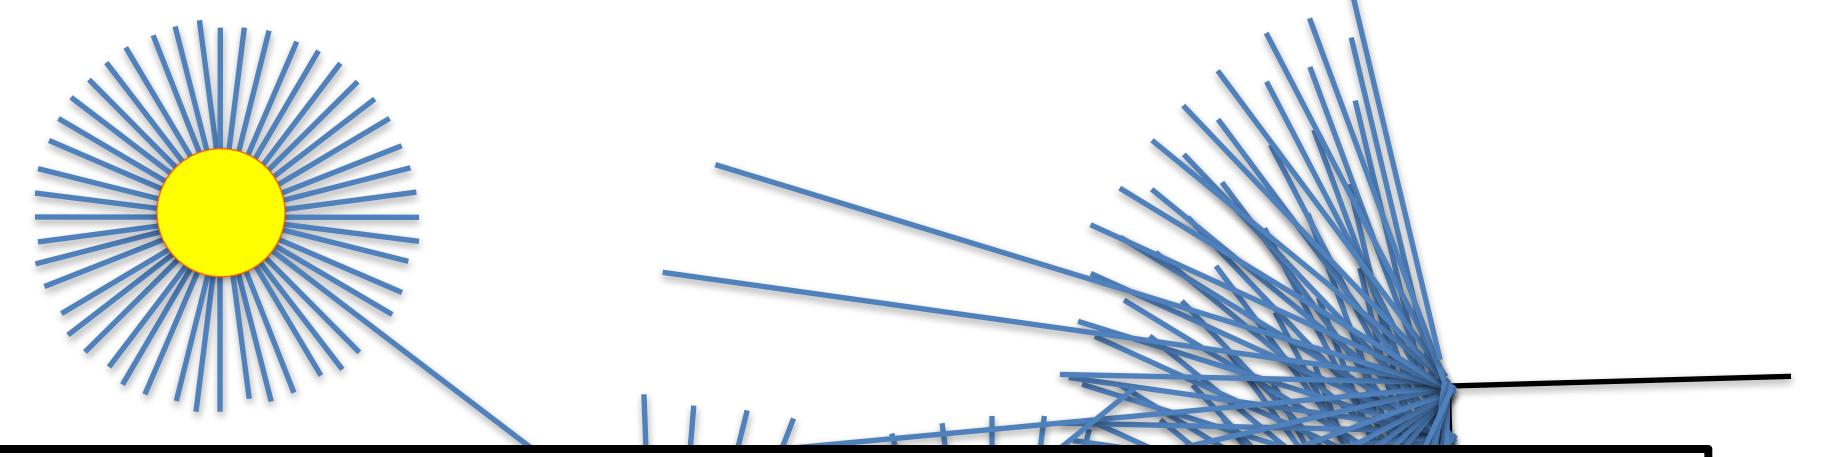
Today

- Camera models
- Projection equations

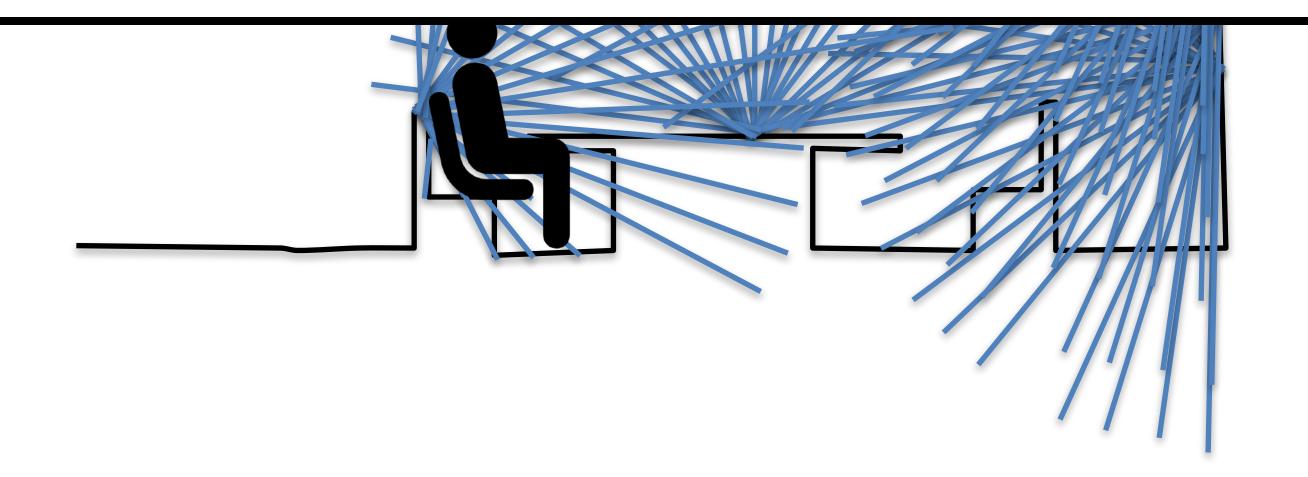
The structure of ambient light

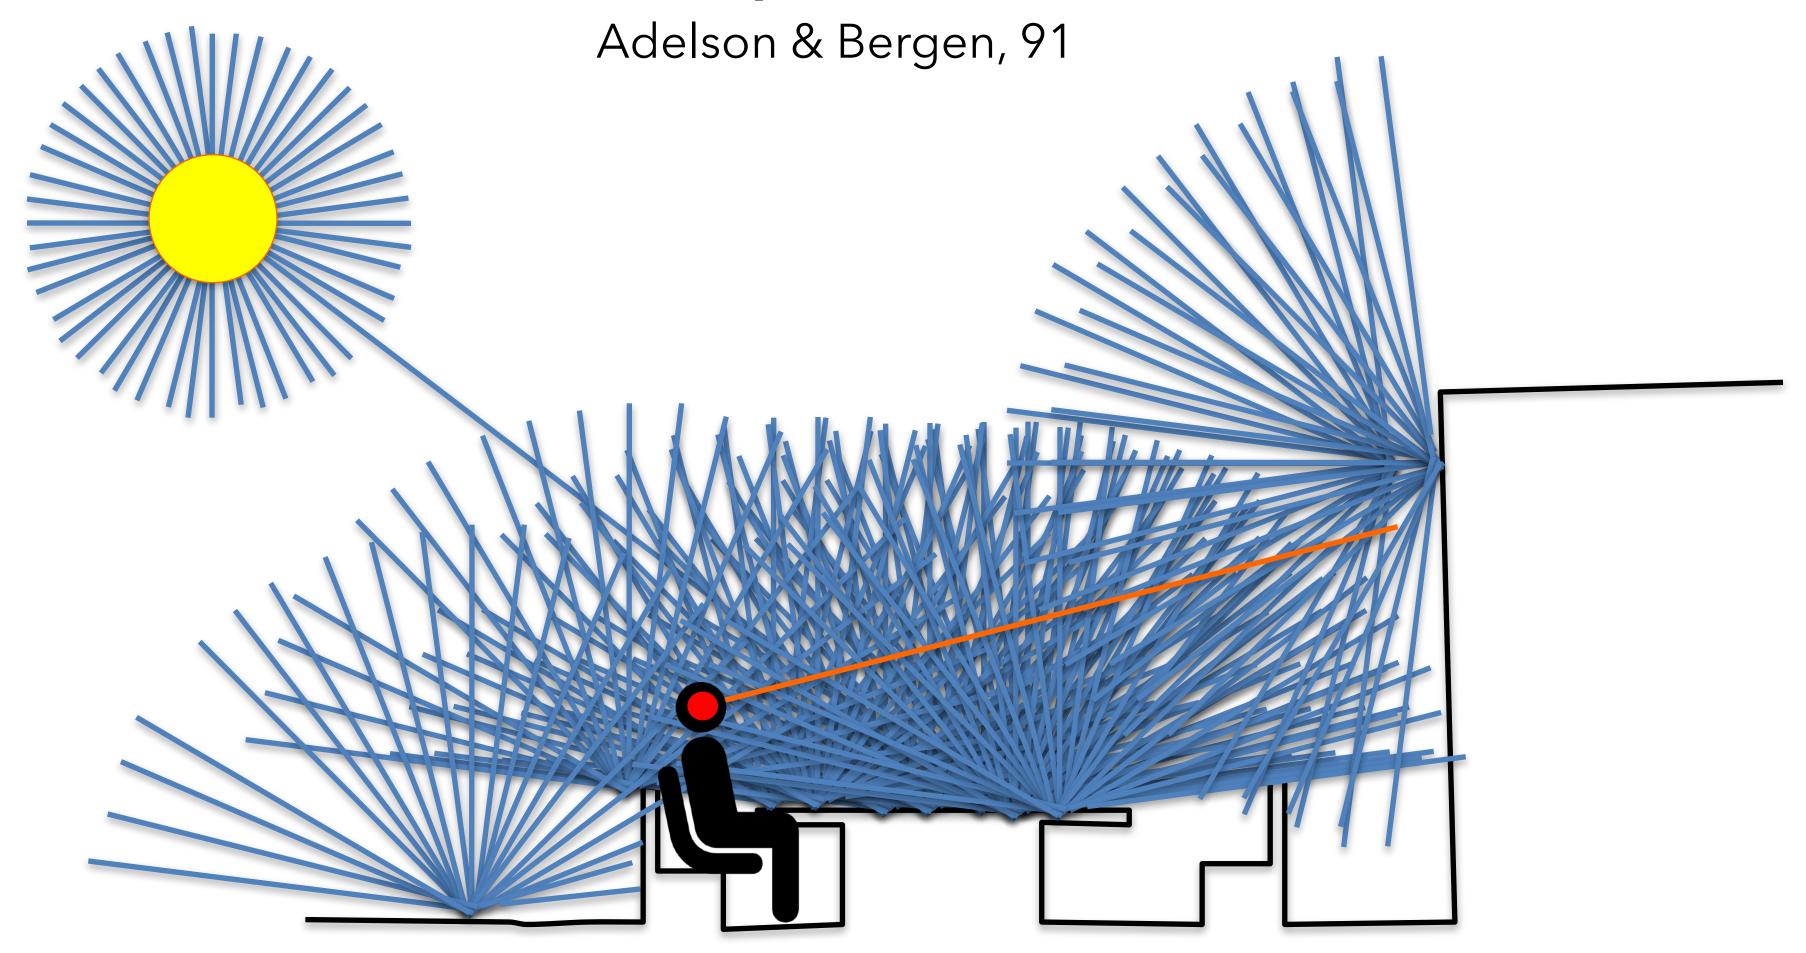


The structure of ambient light

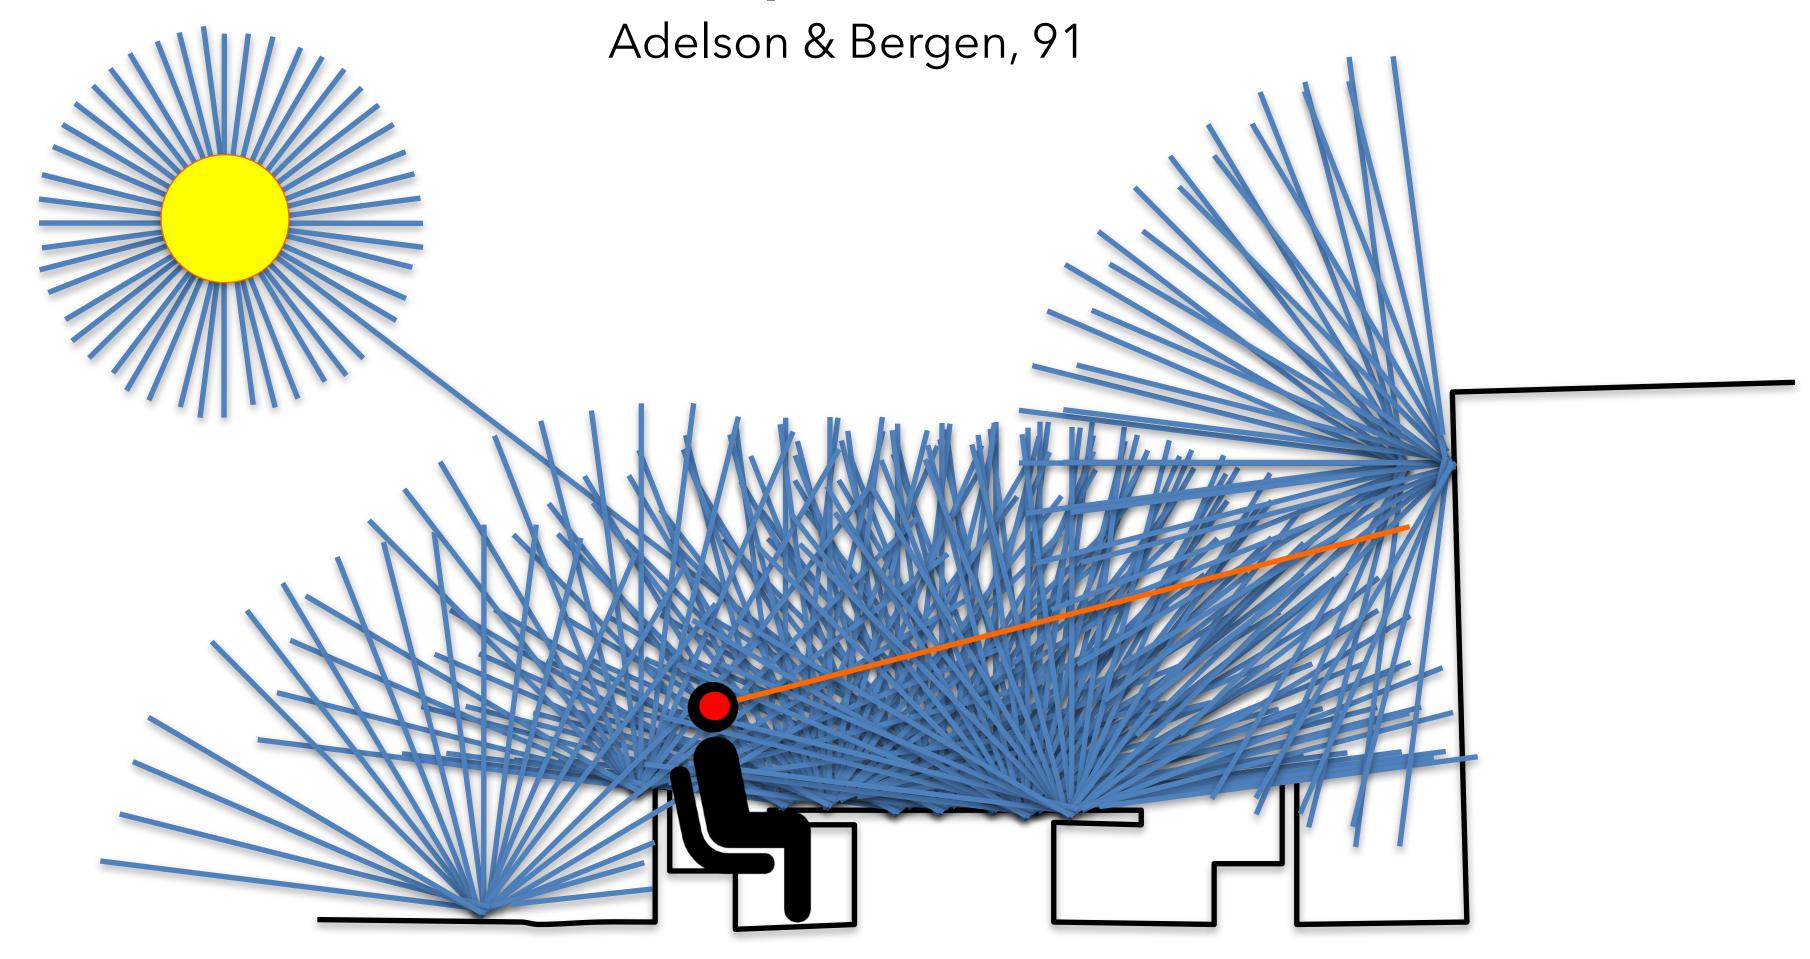


What information does this light provide?



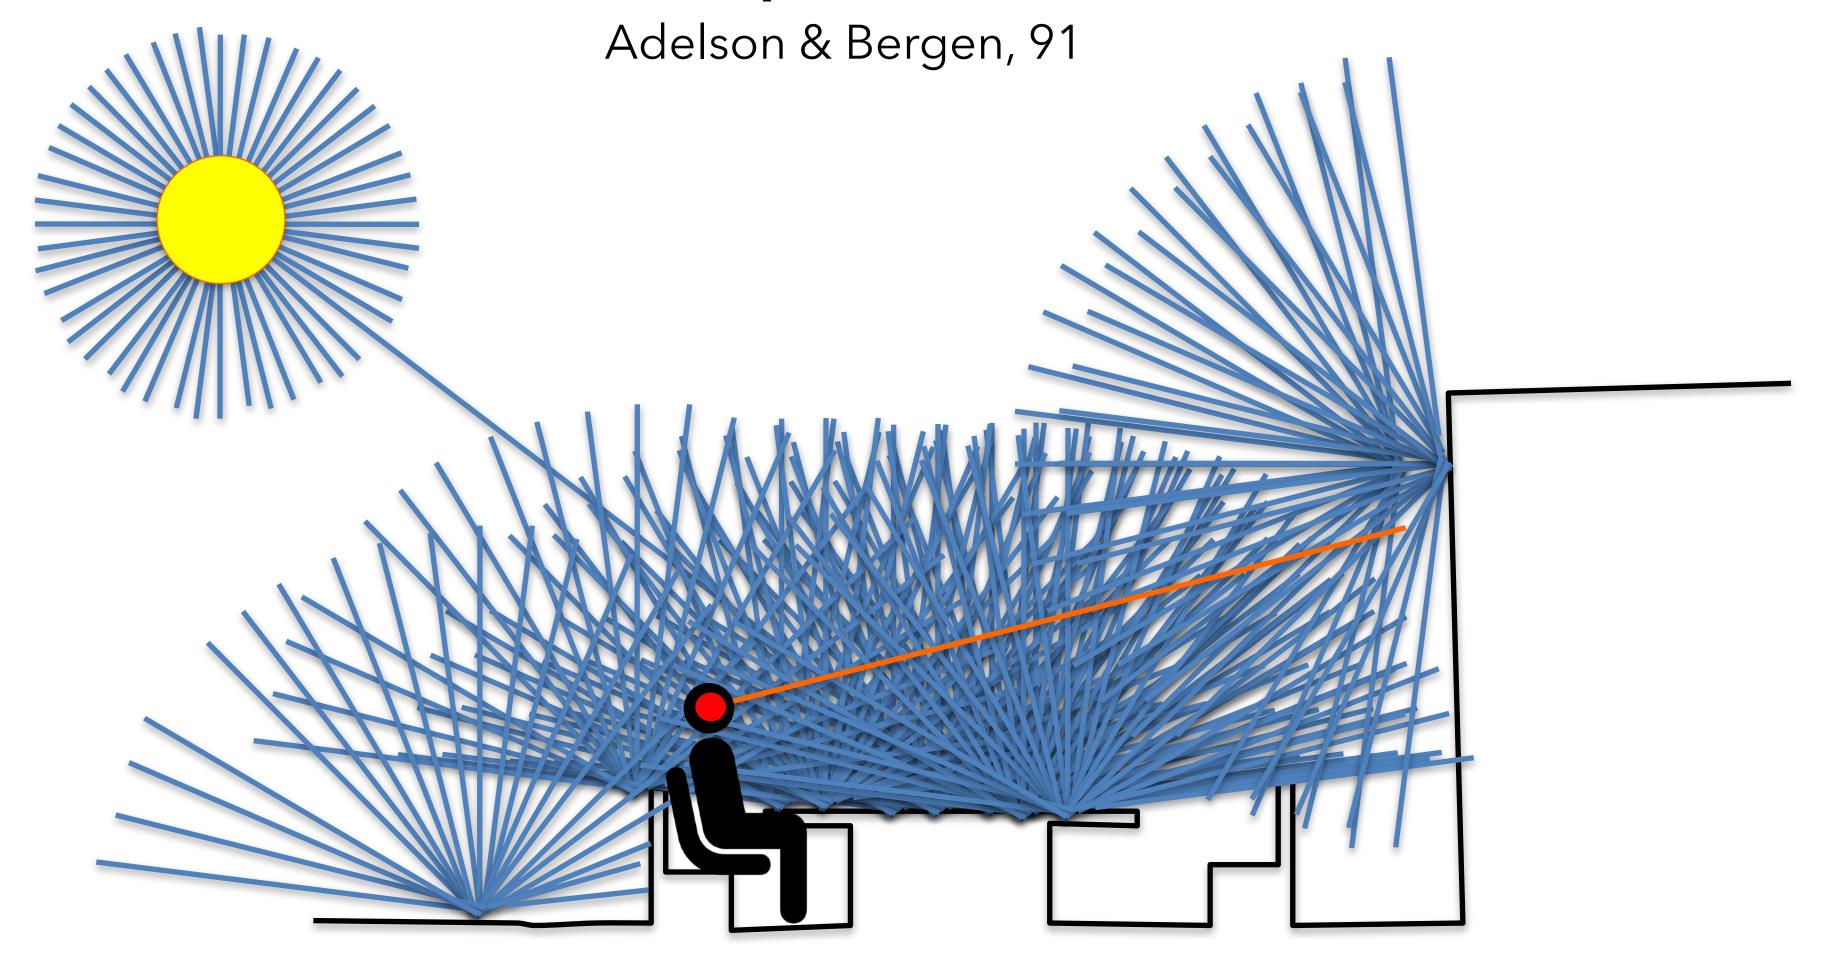


The intensity P can be parameterized as:



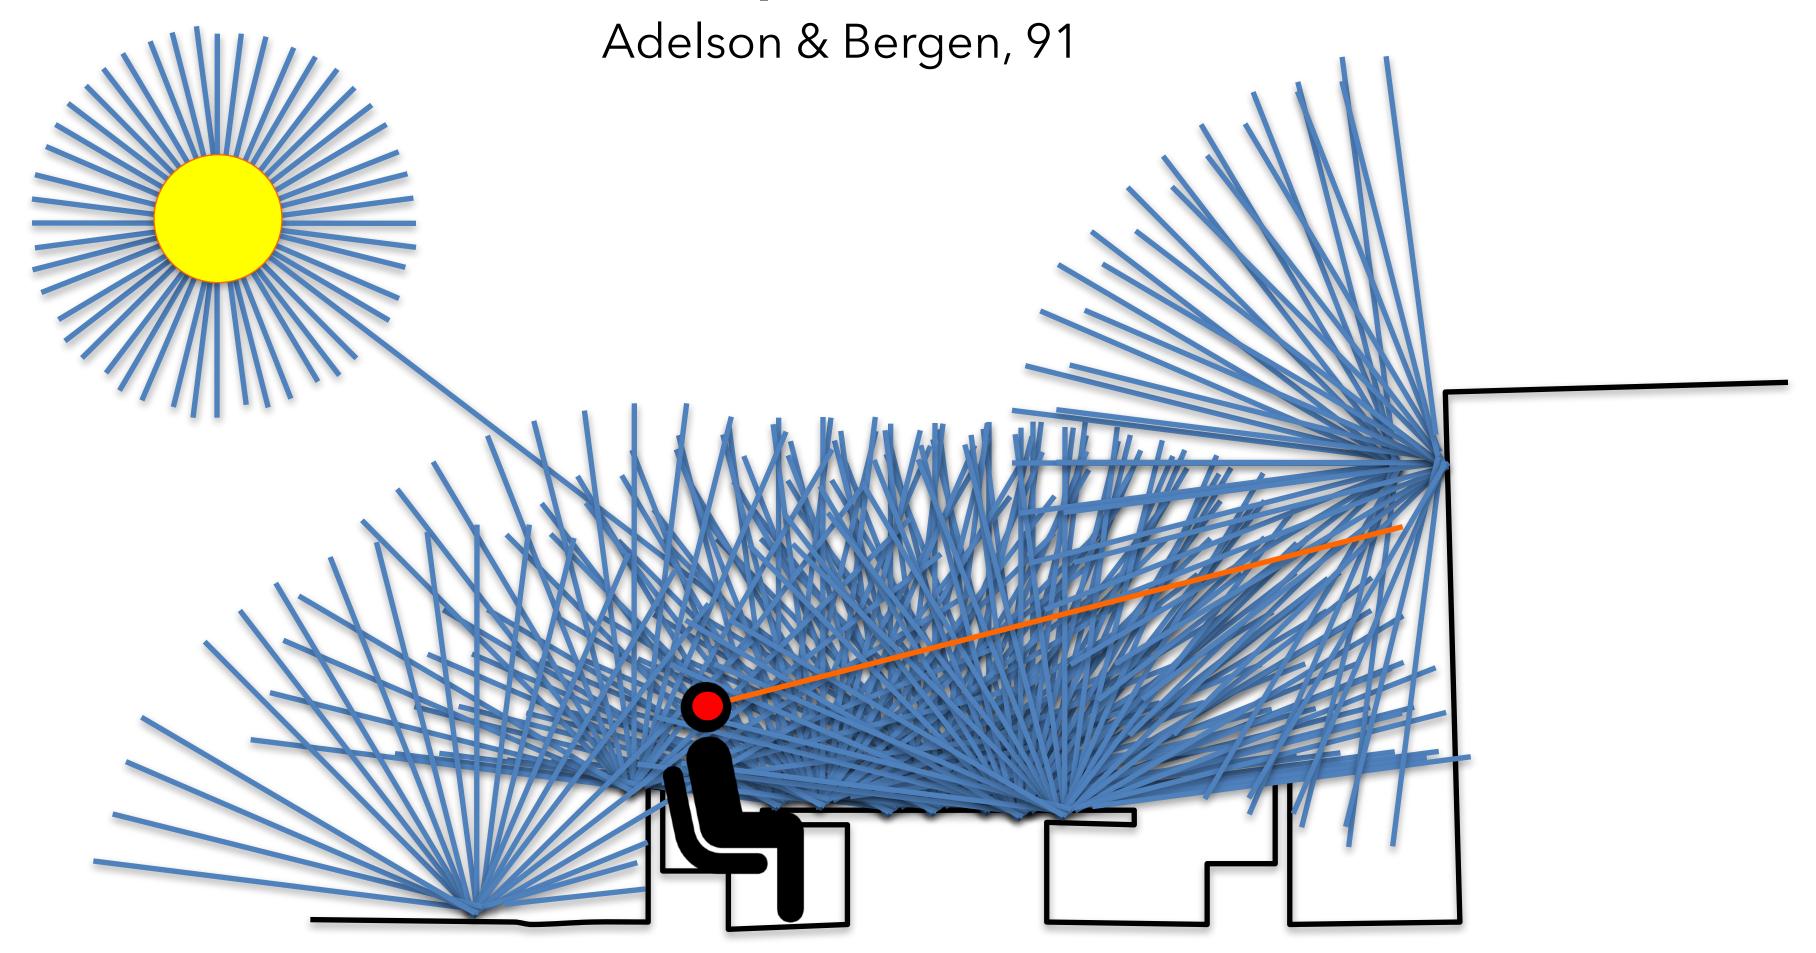
The intensity P can be parameterized as:

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



The intensity P can be parameterized as:

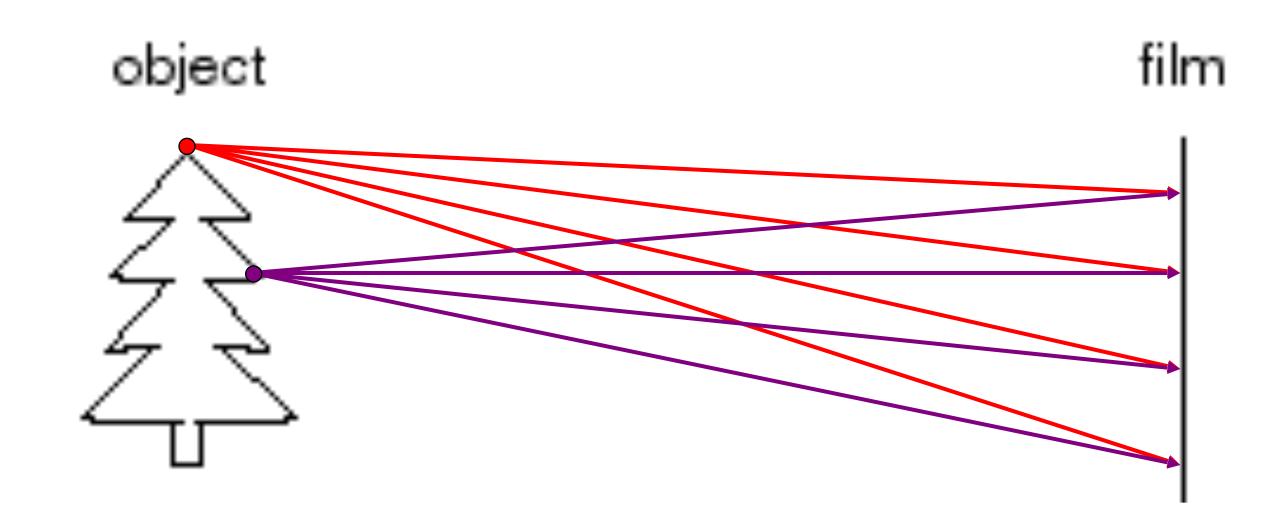
P (θ, φ, λ, t, X, Y, Z) Wavelength, time



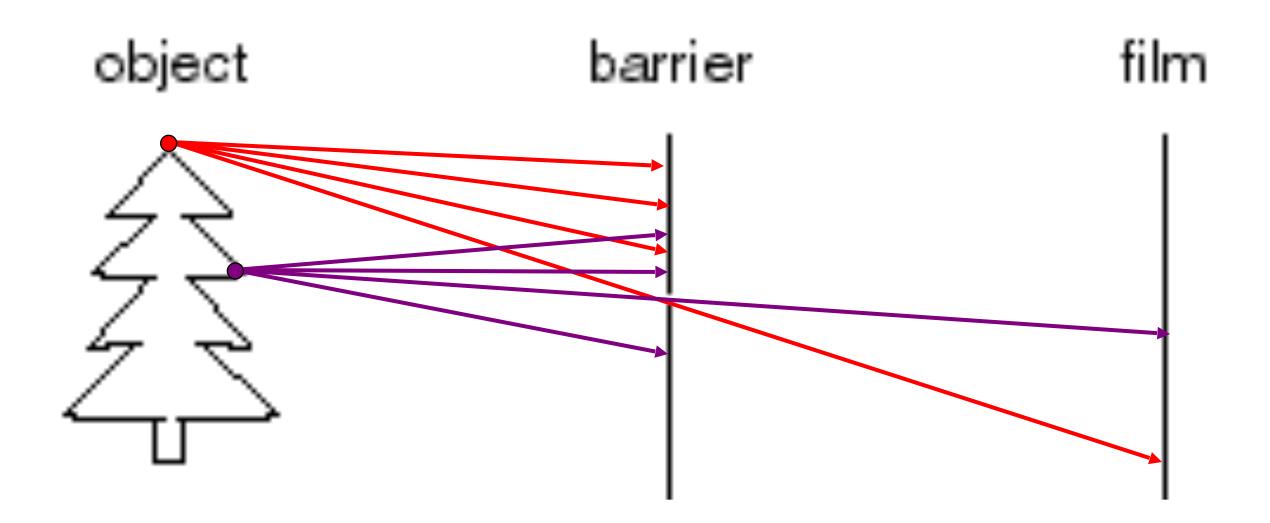
The intensity P can be parameterized as:

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

Making a camera



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

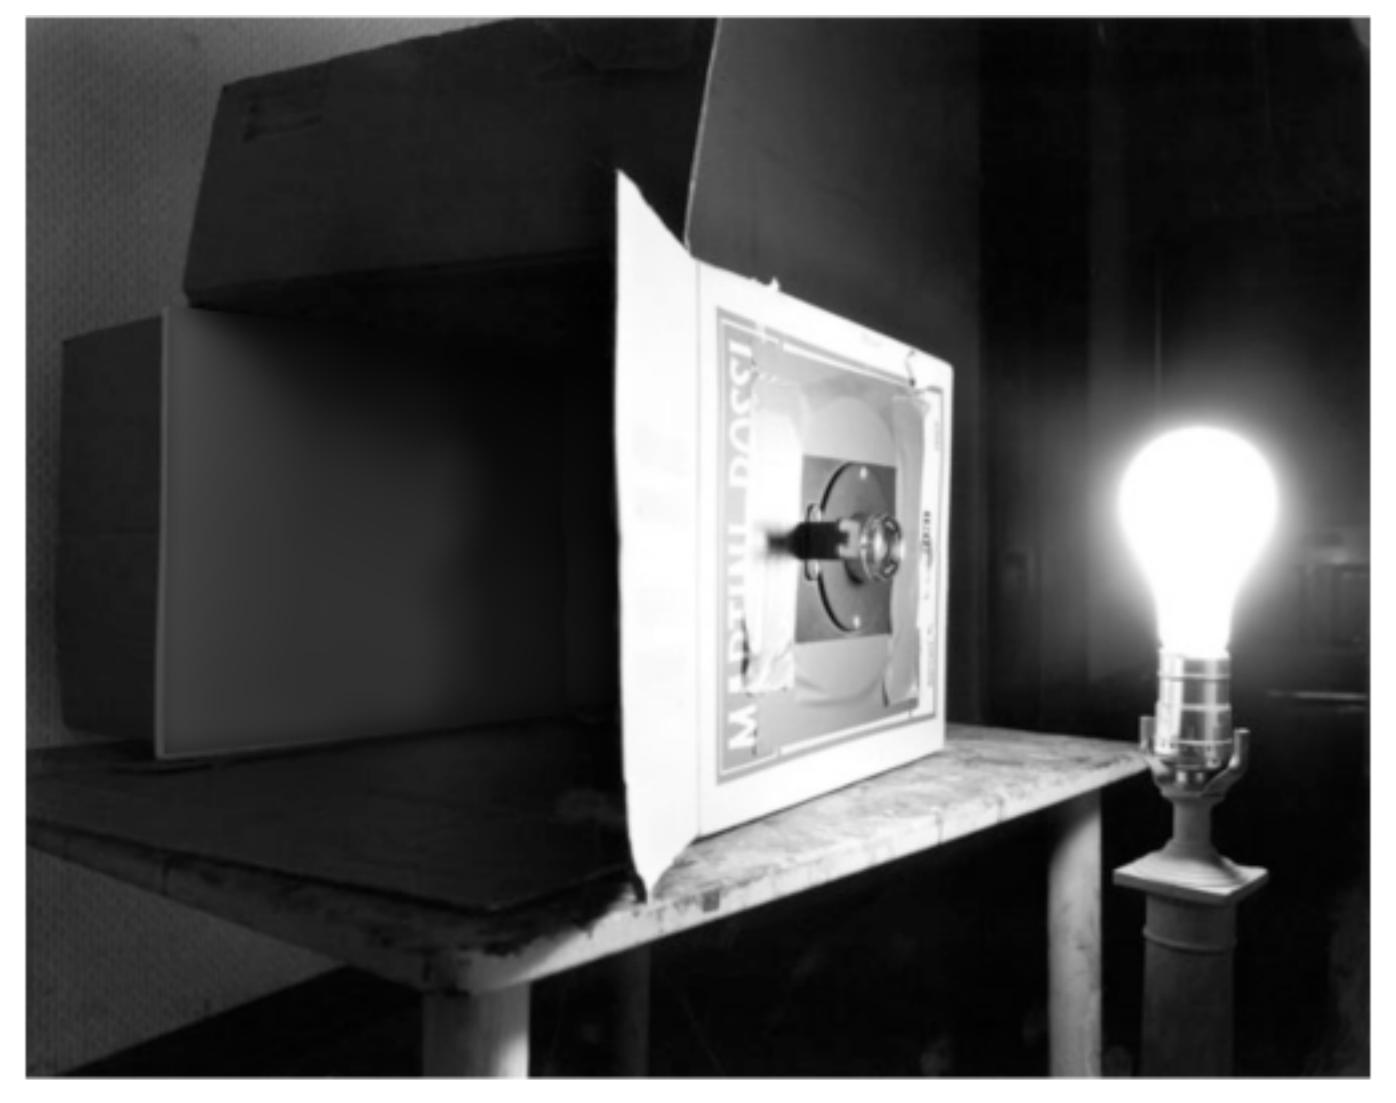


Add a barrier to block off most of the rays

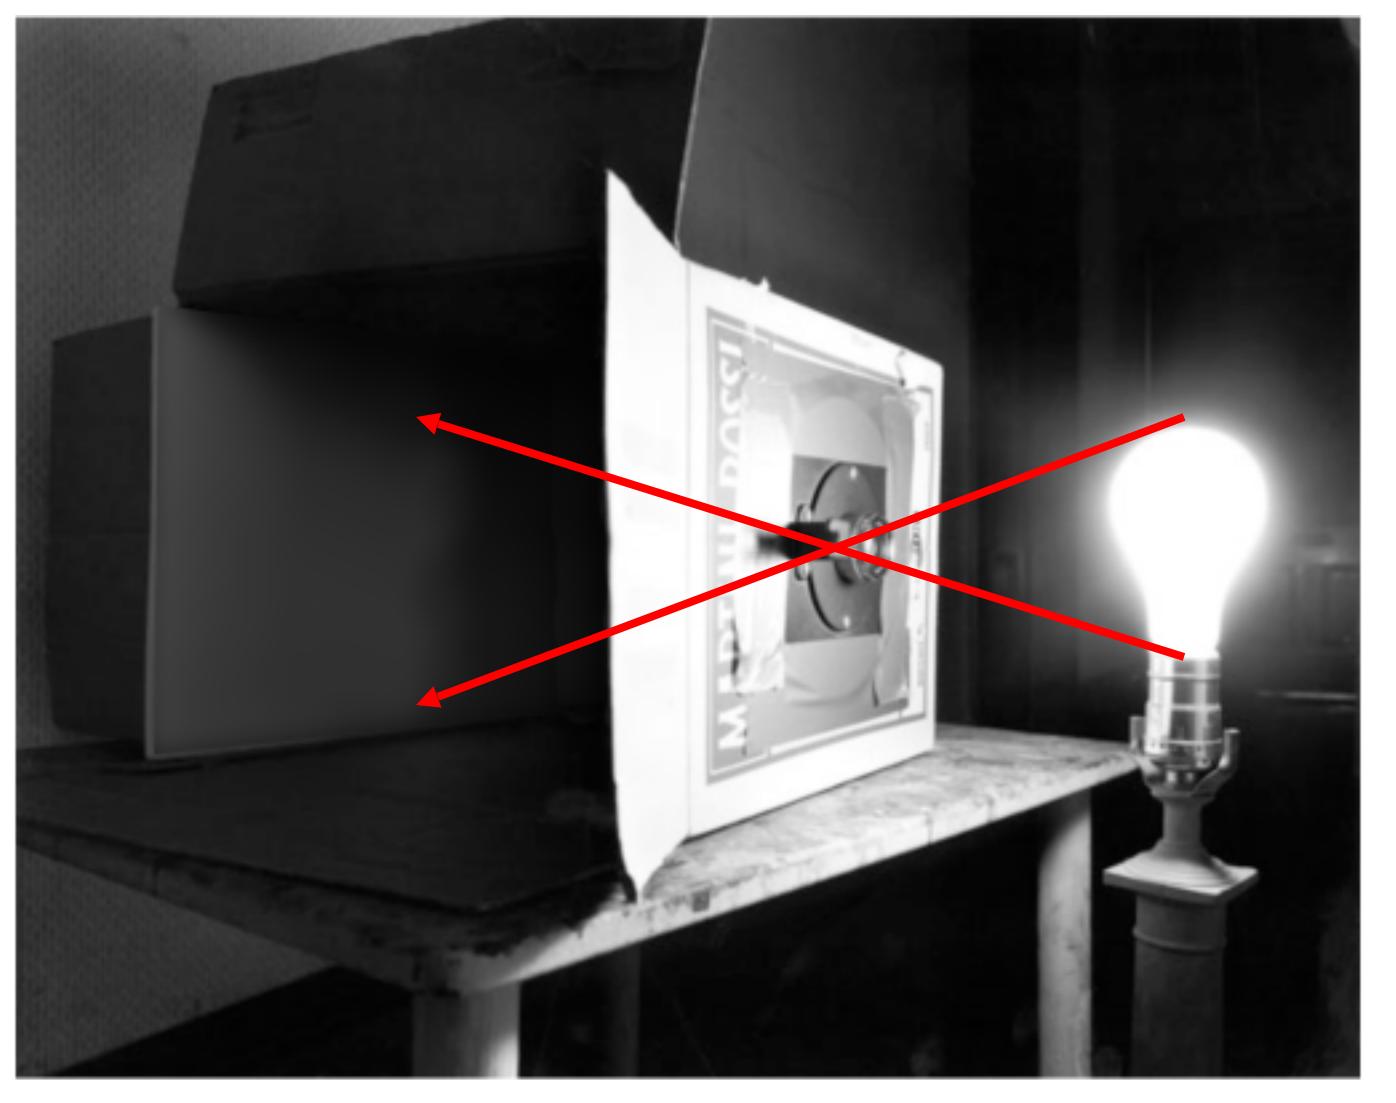
- This reduces blurring
- The opening known as the aperture



Photograph by Abelardo Morell, 1991

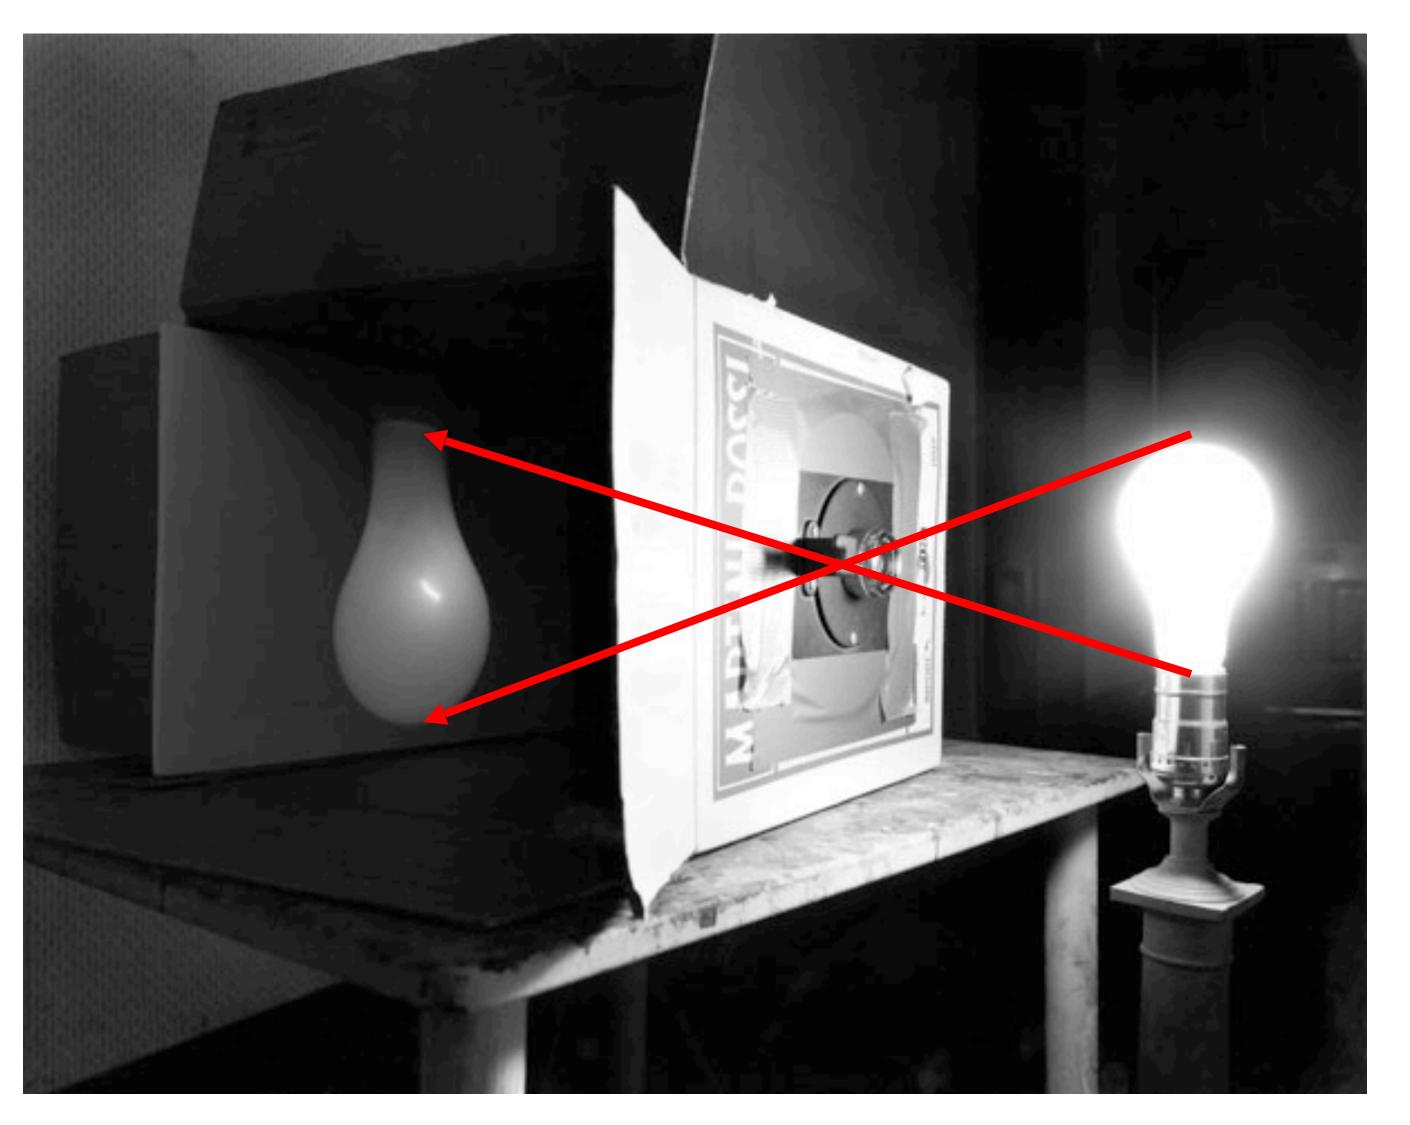


Photograph by Abelardo Morell, 1991



Photograph by Abelardo Morell, 1991

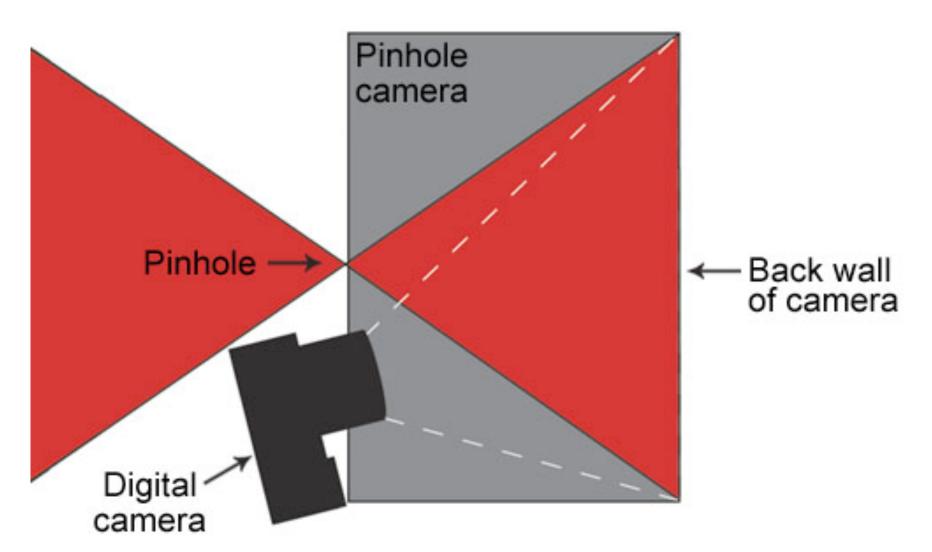
14 Source: Freeman, Torralba, Isola



Photograph by Abelardo Morell, 1991





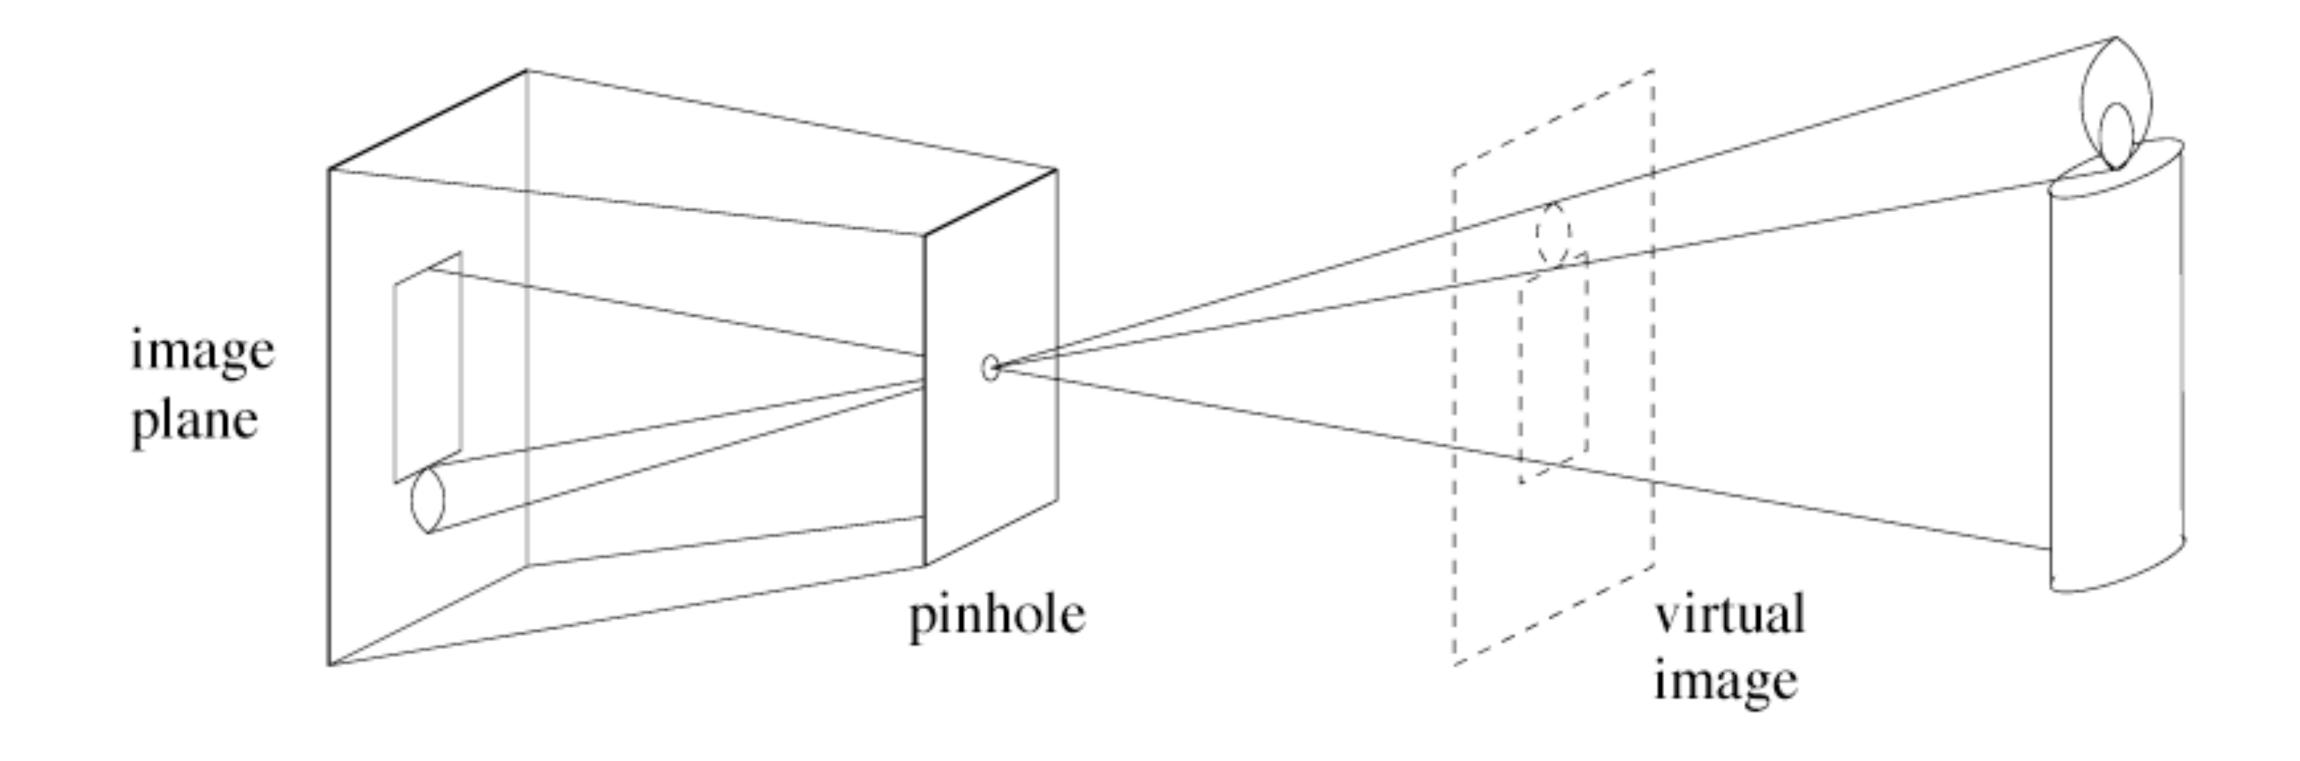


http://www.foundphotography.com/PhotoThoughts/archives/2005/04/pinhole_camera_2.html 16 Source: Freeman, Torralba, Isola



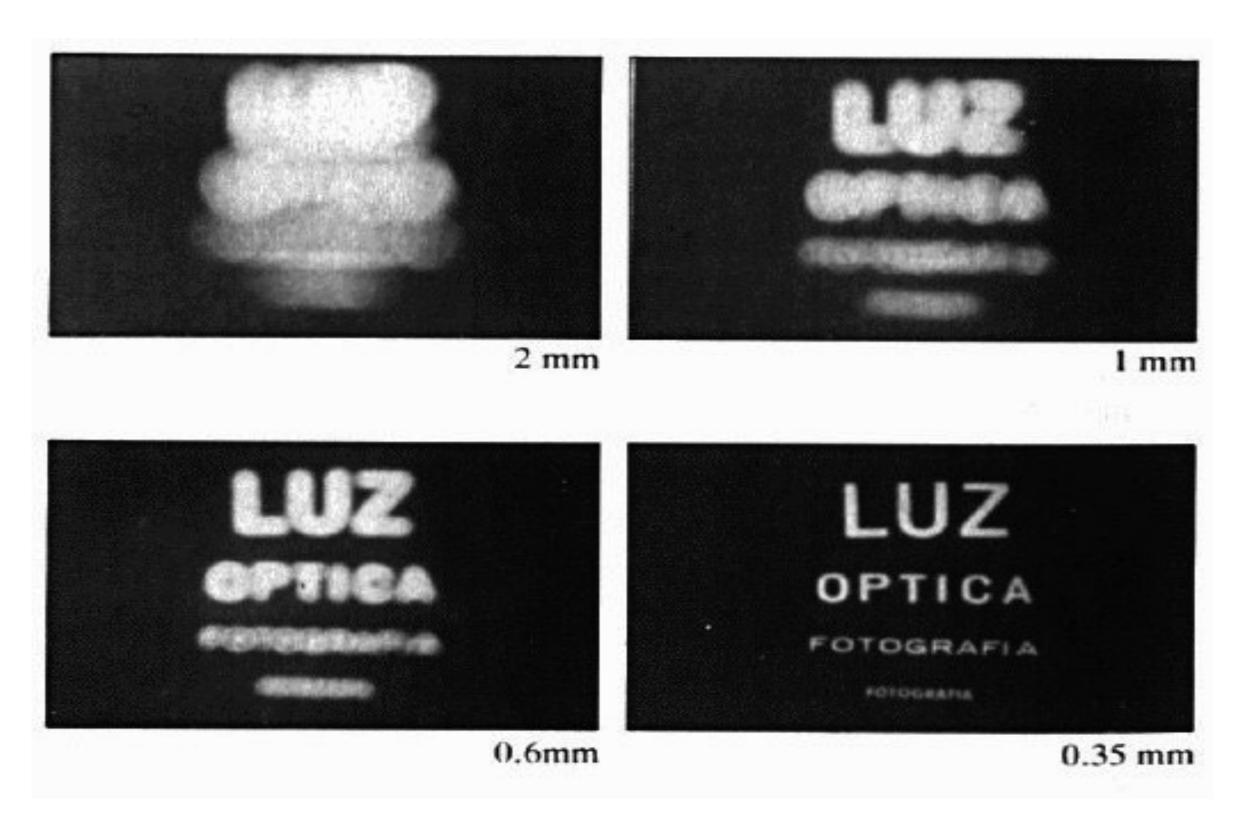


Upside down images!



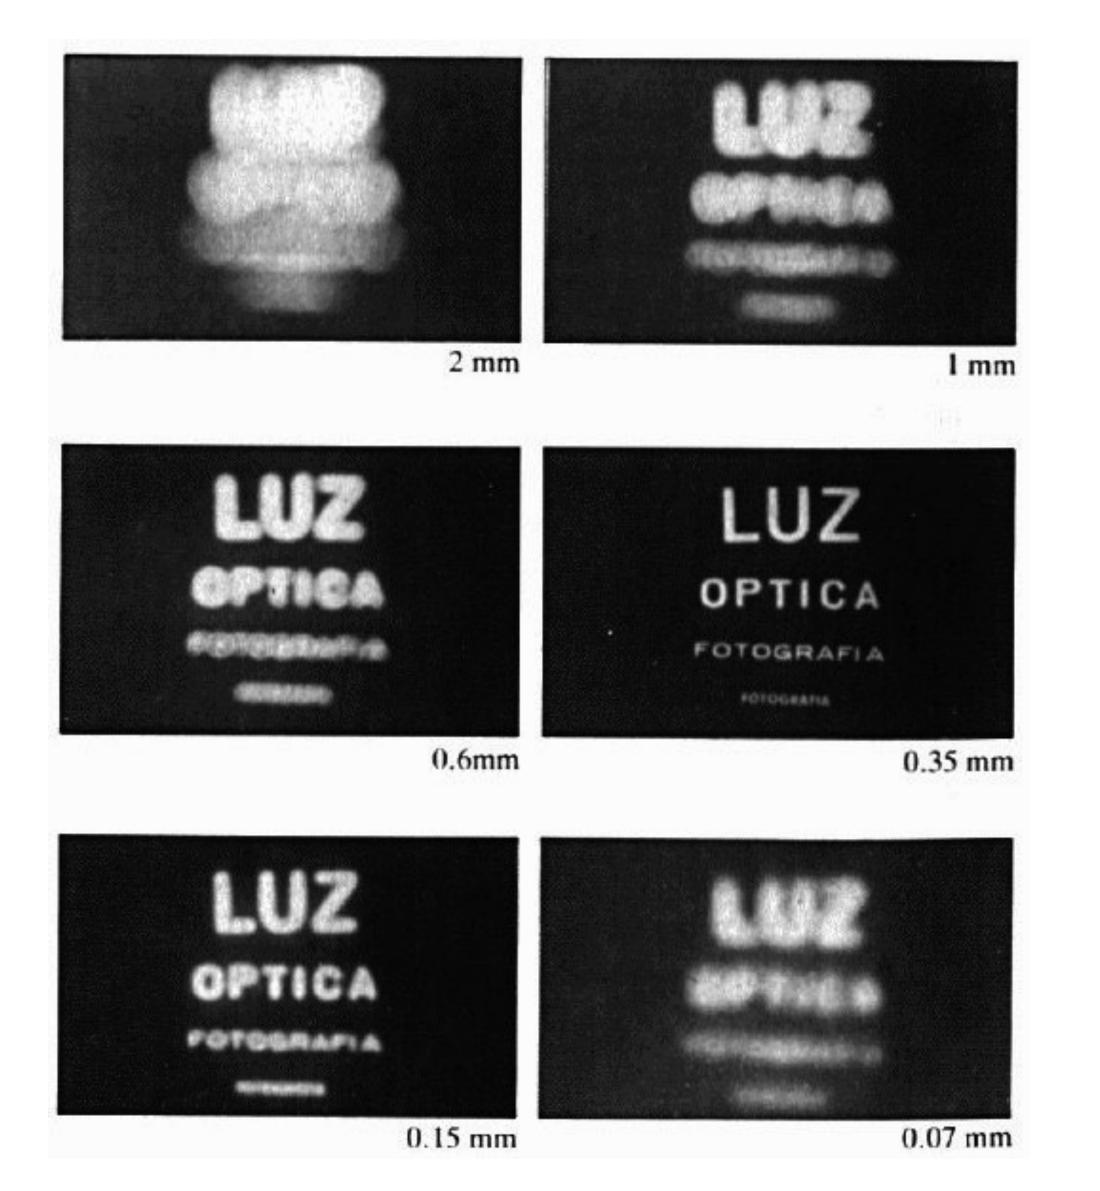
A useful concept: virtual image

Shrinking the aperture

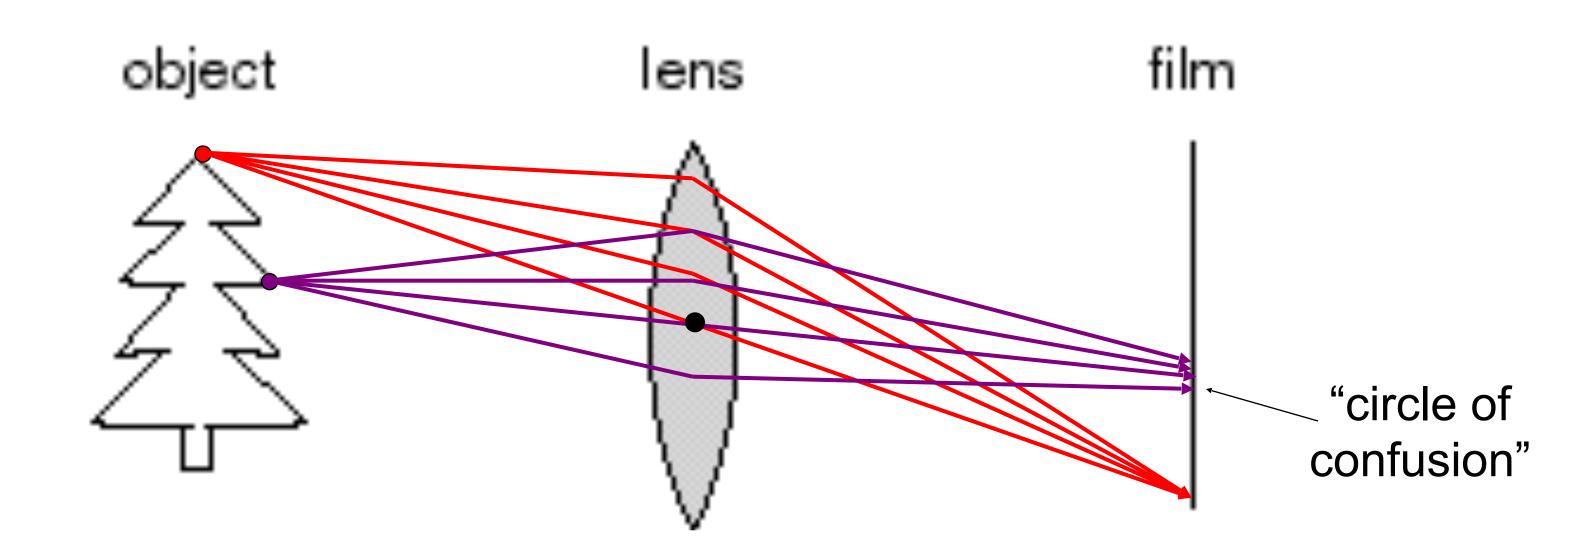


- Why not make the aperture as small as possible?
 - Less light gets through
 - Diffraction effects...

Shrinking the aperture



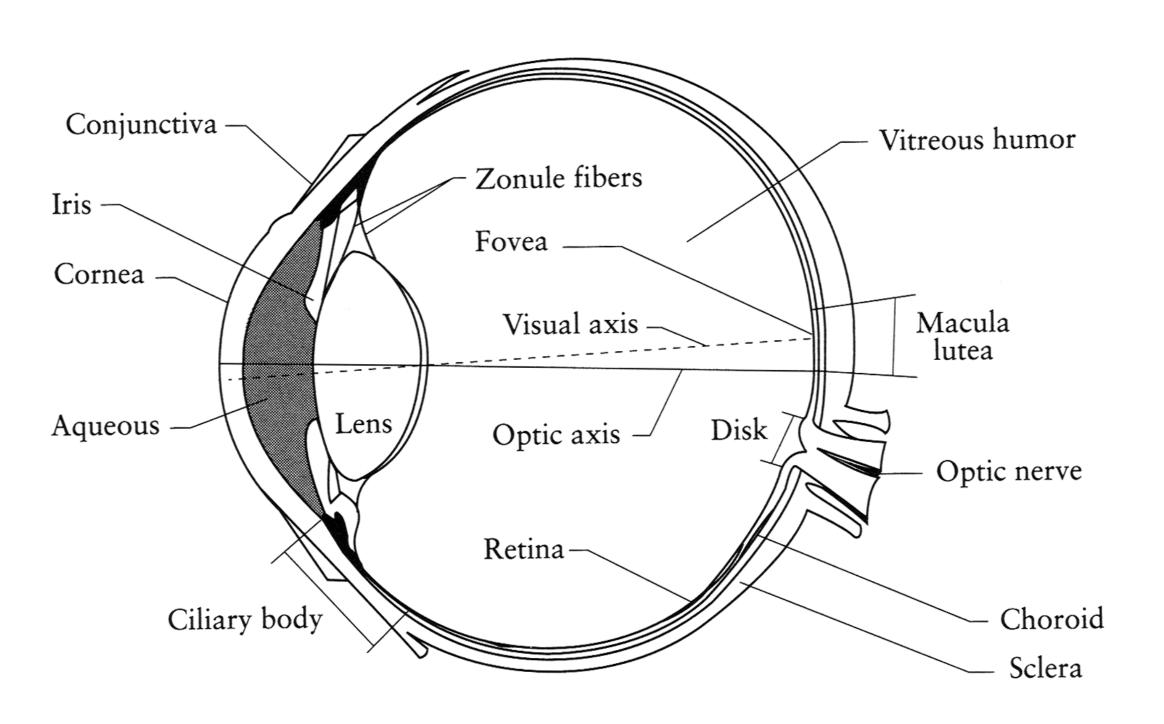
Adding a lens



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

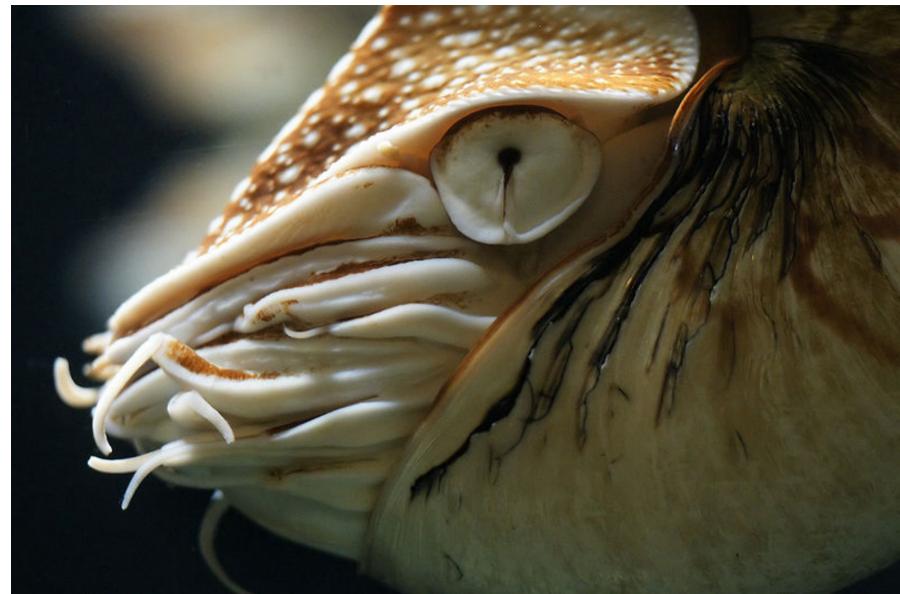
The eye



- The human eye is a camera:
 - Iris: colored annulus with radial muscles
 - Pupil: the hole (aperture) whose size is controlled by the iris
 - Photoreceptor cells (rods and cones) in the retina are analogous to film

Eyes in nature: eyespots to pinhole camera







Pinhole cameras in unexpected places



Tree shadow during a solar eclipse

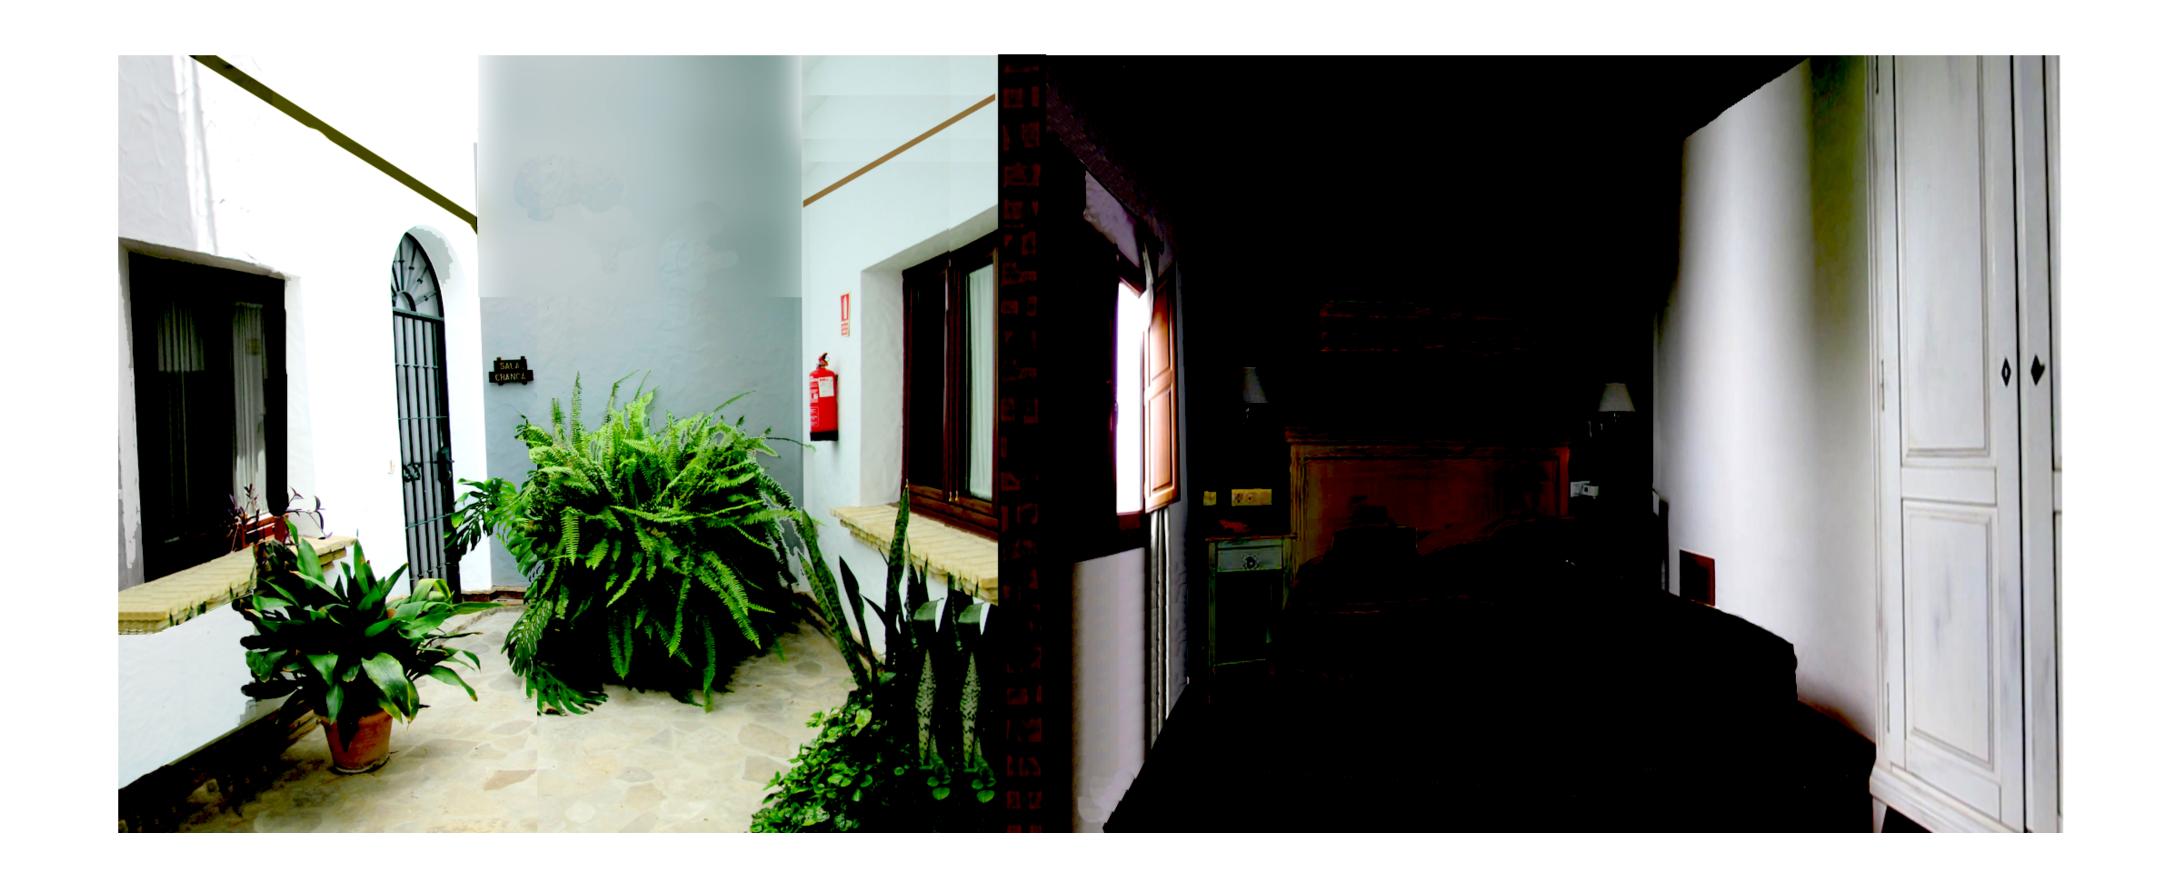
photo credit: Nils van der Burg

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

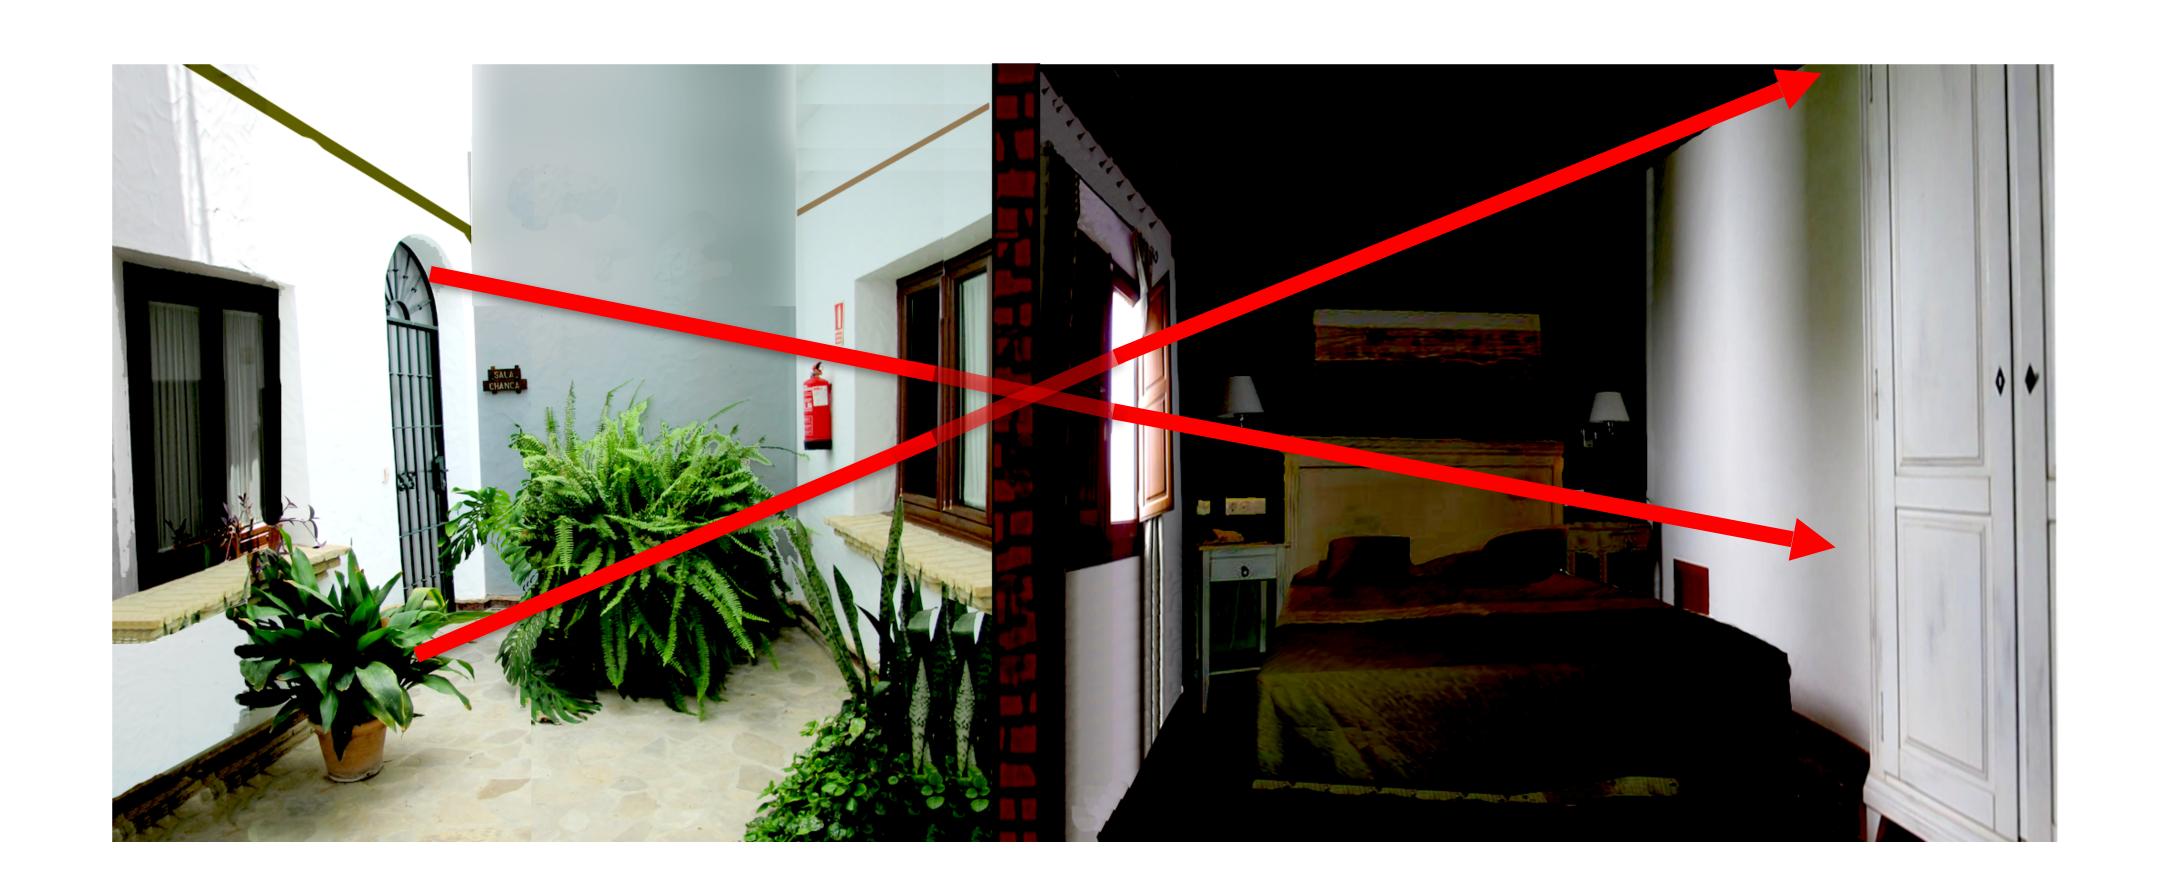
Source: Steve Seitz







Accidental pinhole camera







Window turned into a pinhole

View outside







Window open

Window turned into a pinhole







Accidental pinhole camera

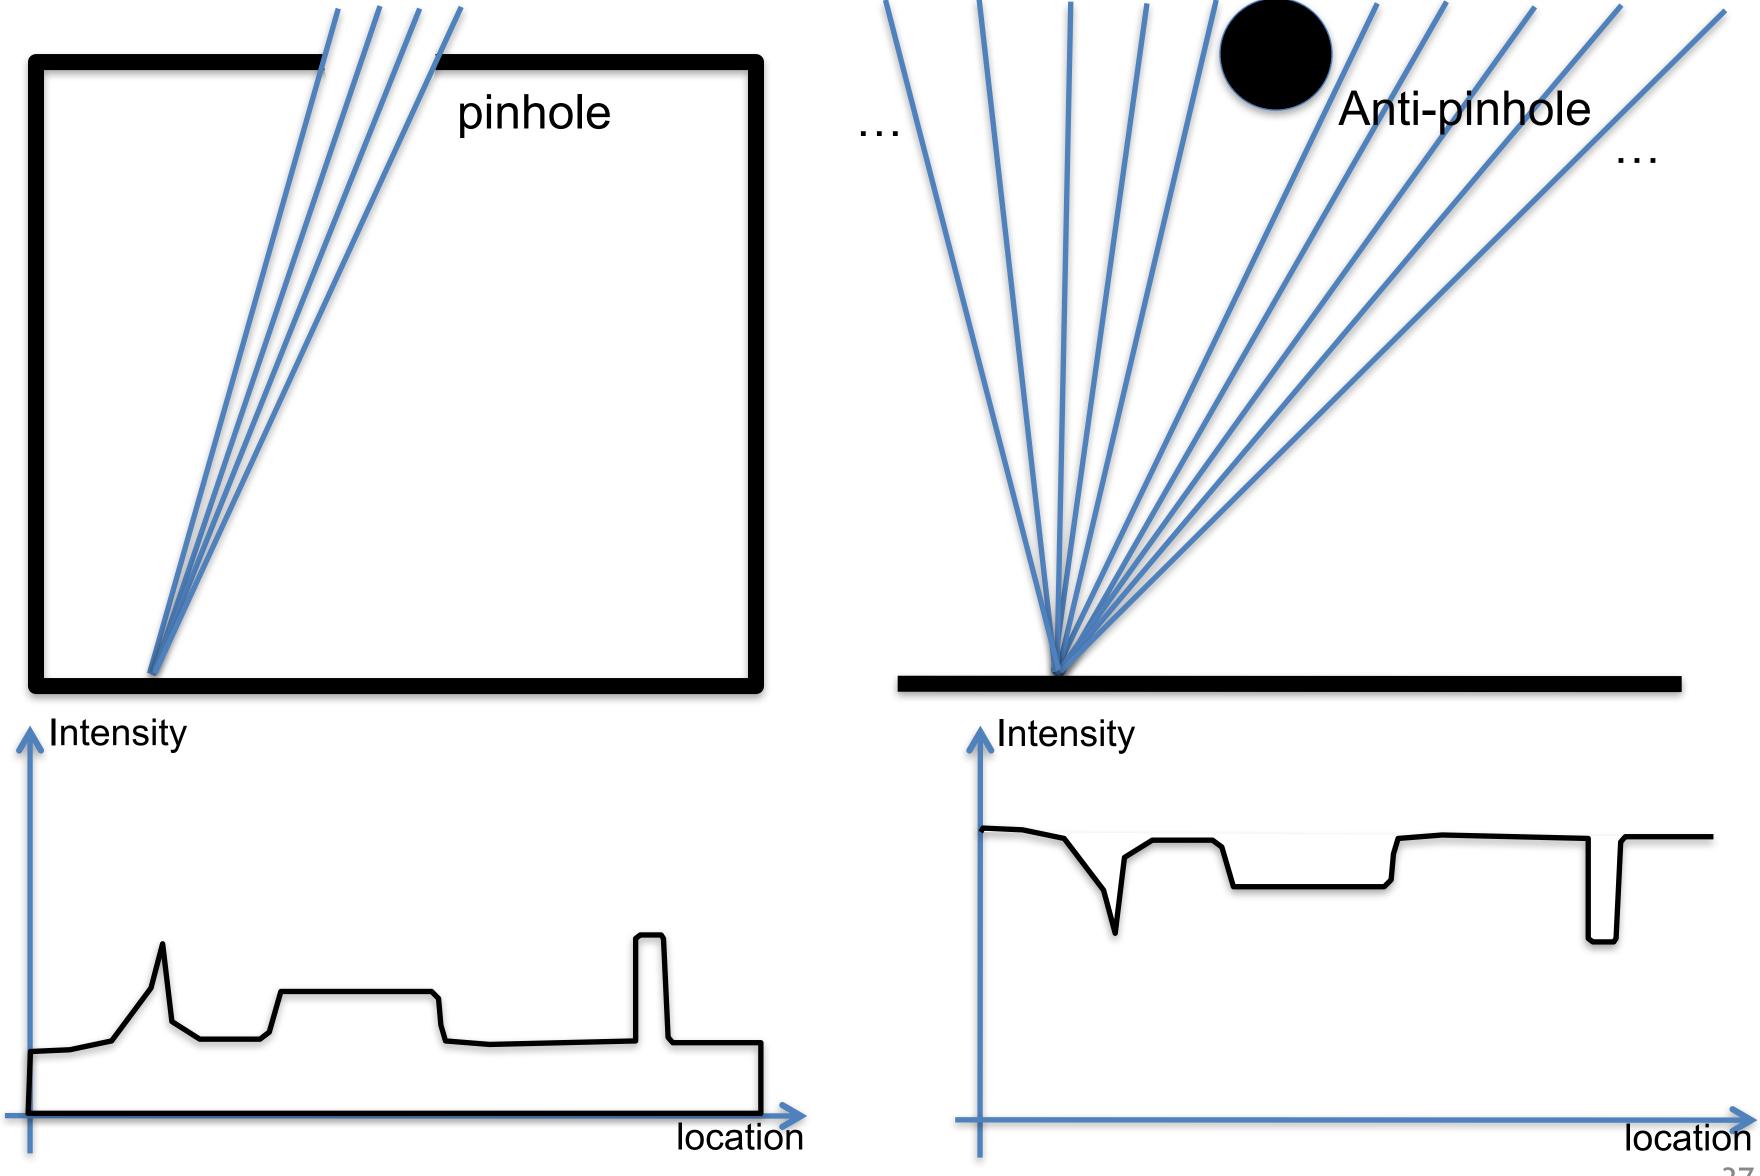


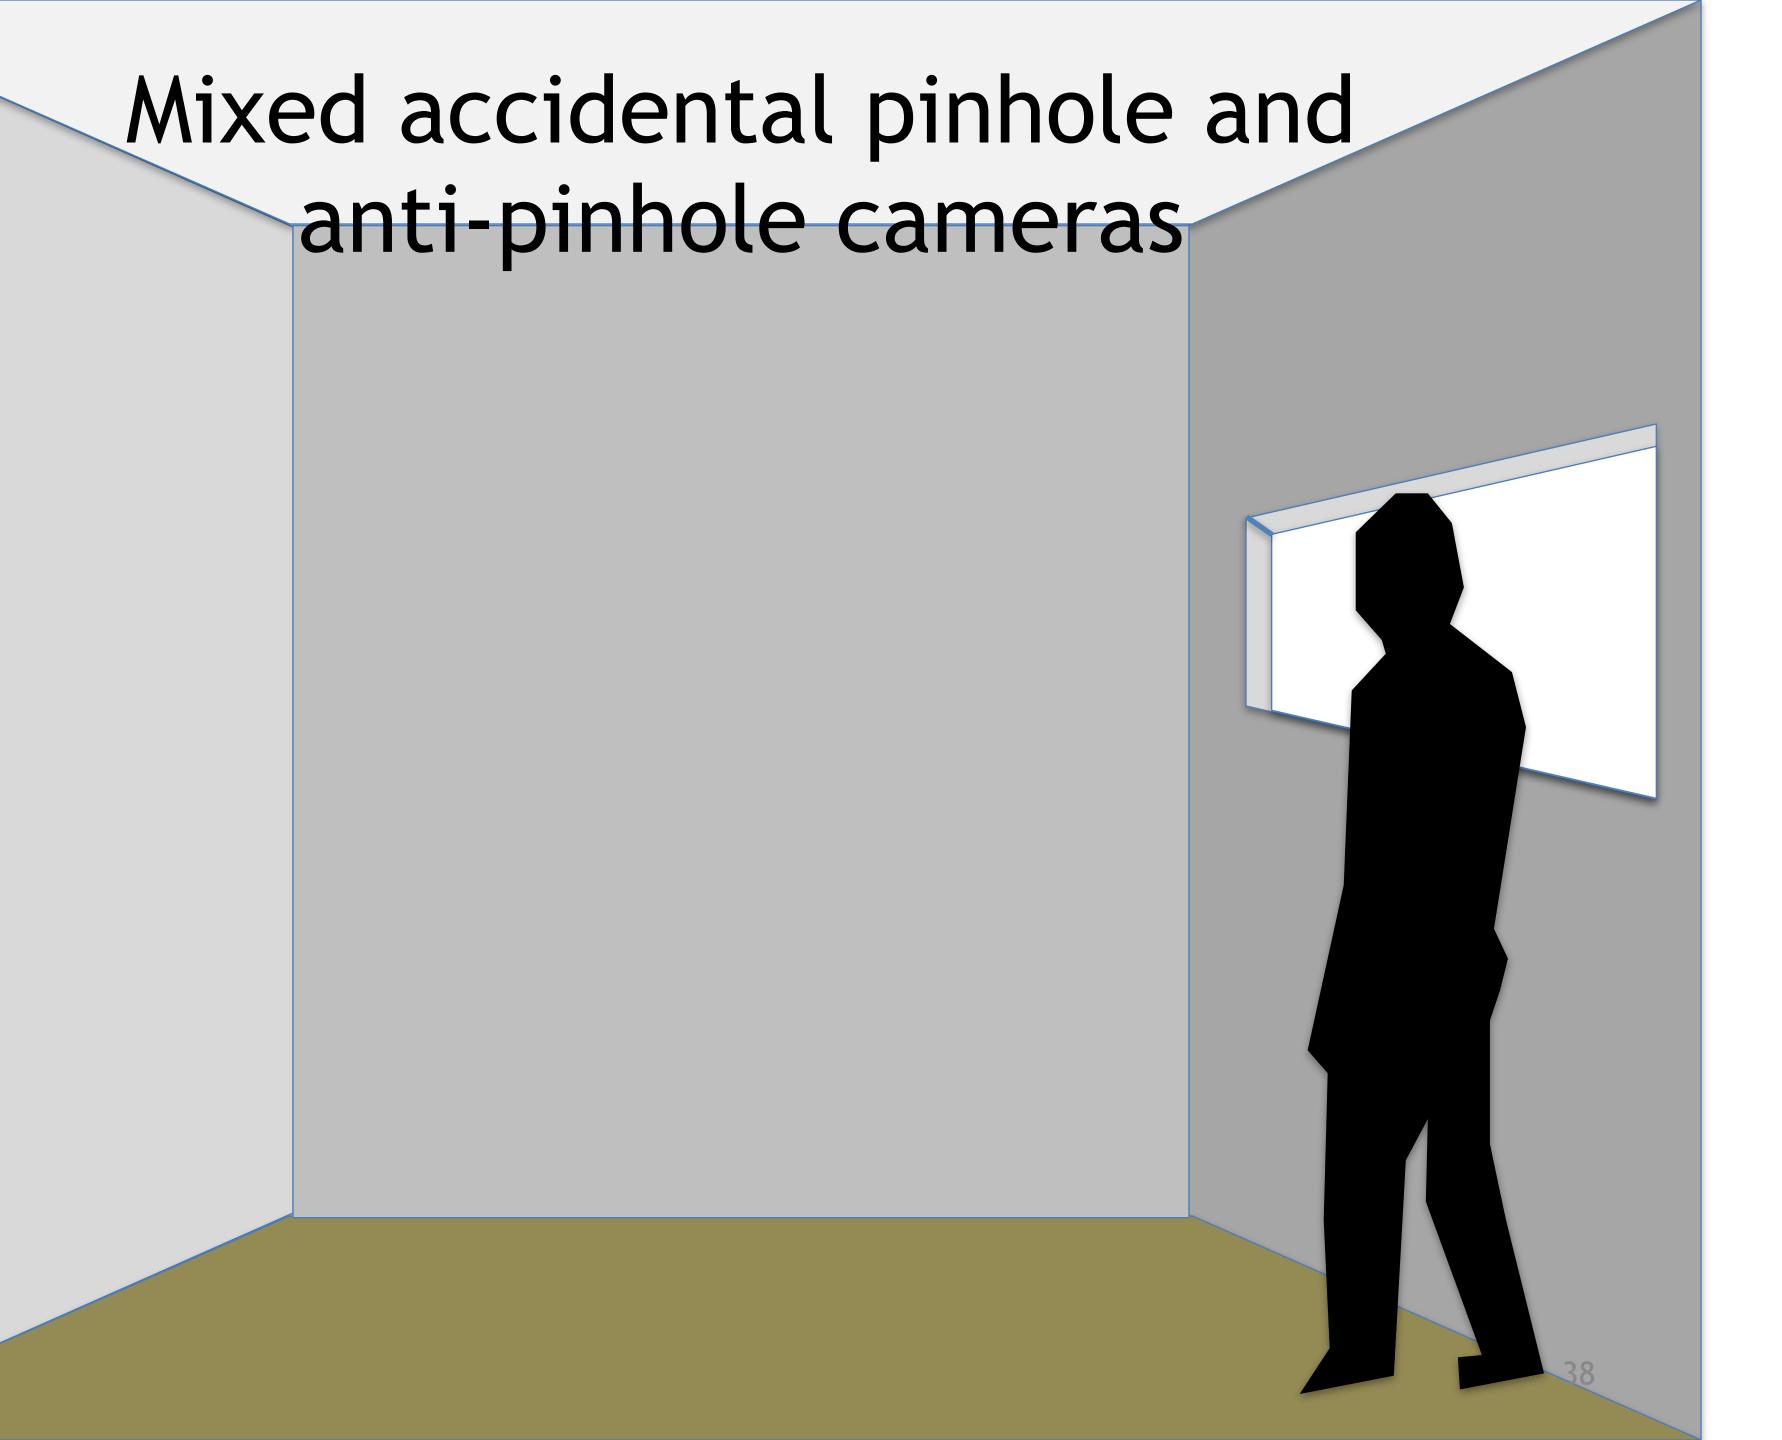
Outside scene

36

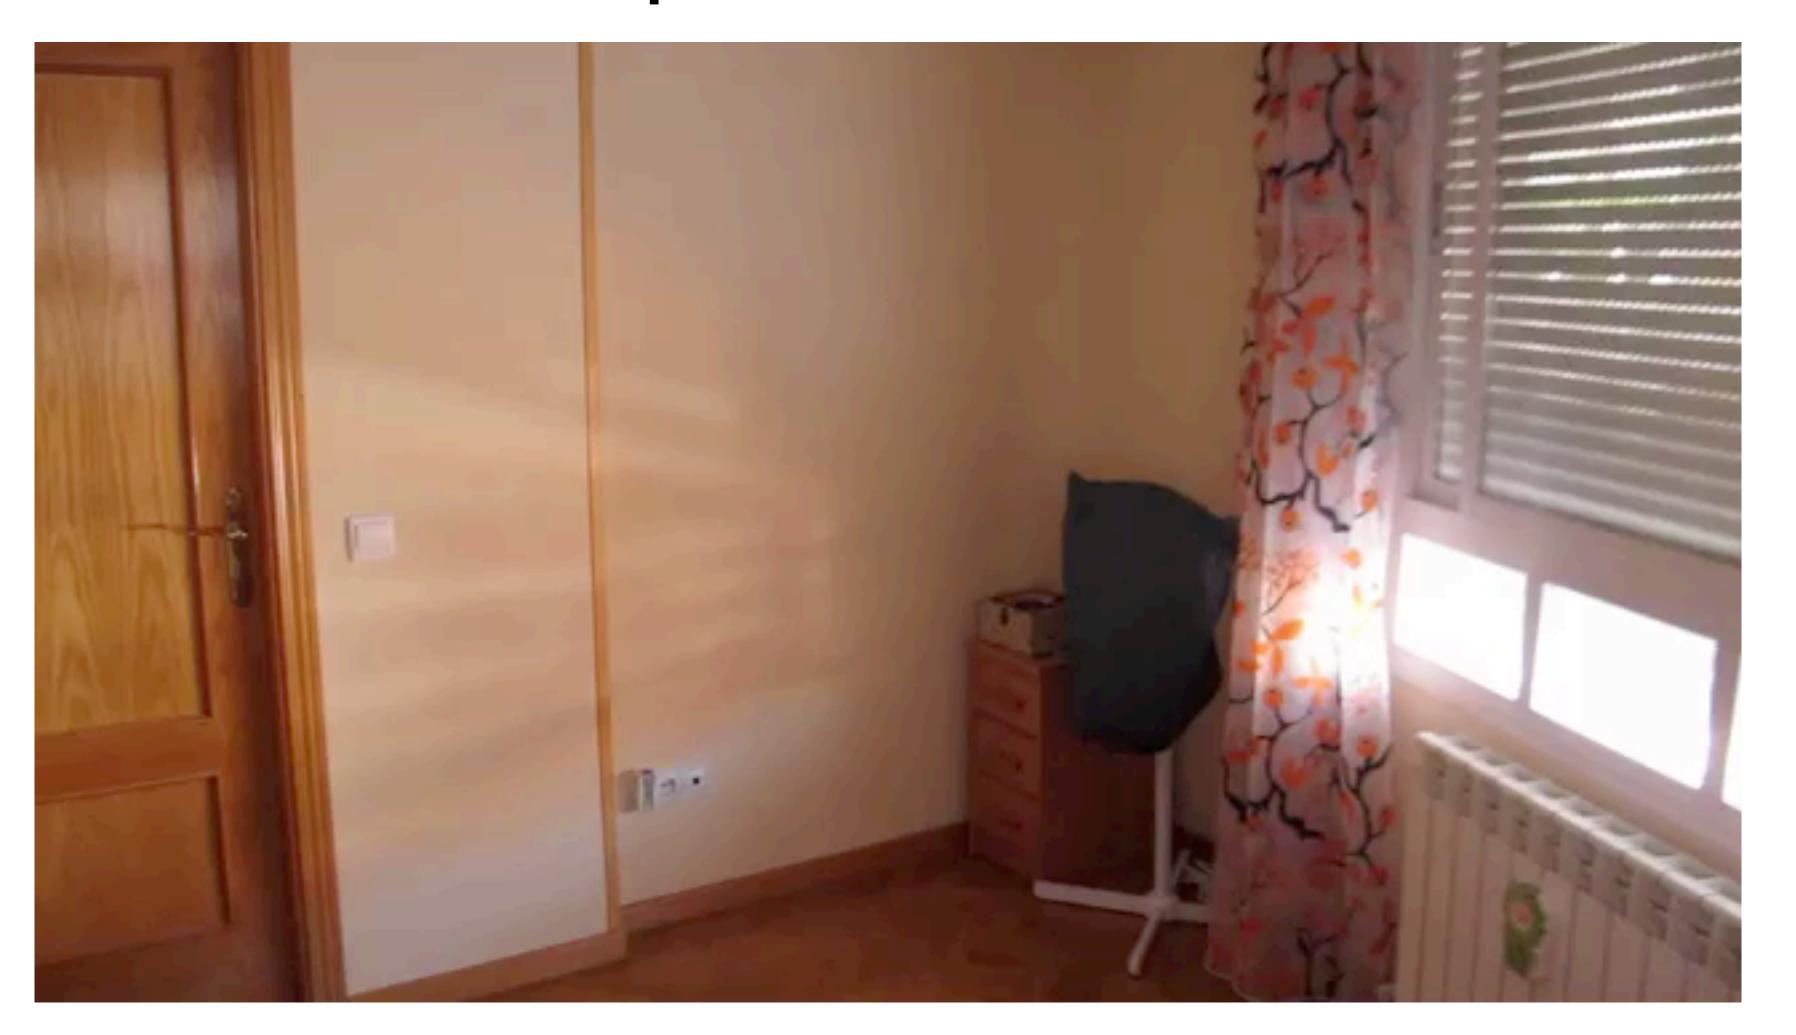
Aperture

Pinhole and Anti-pinhole cameras

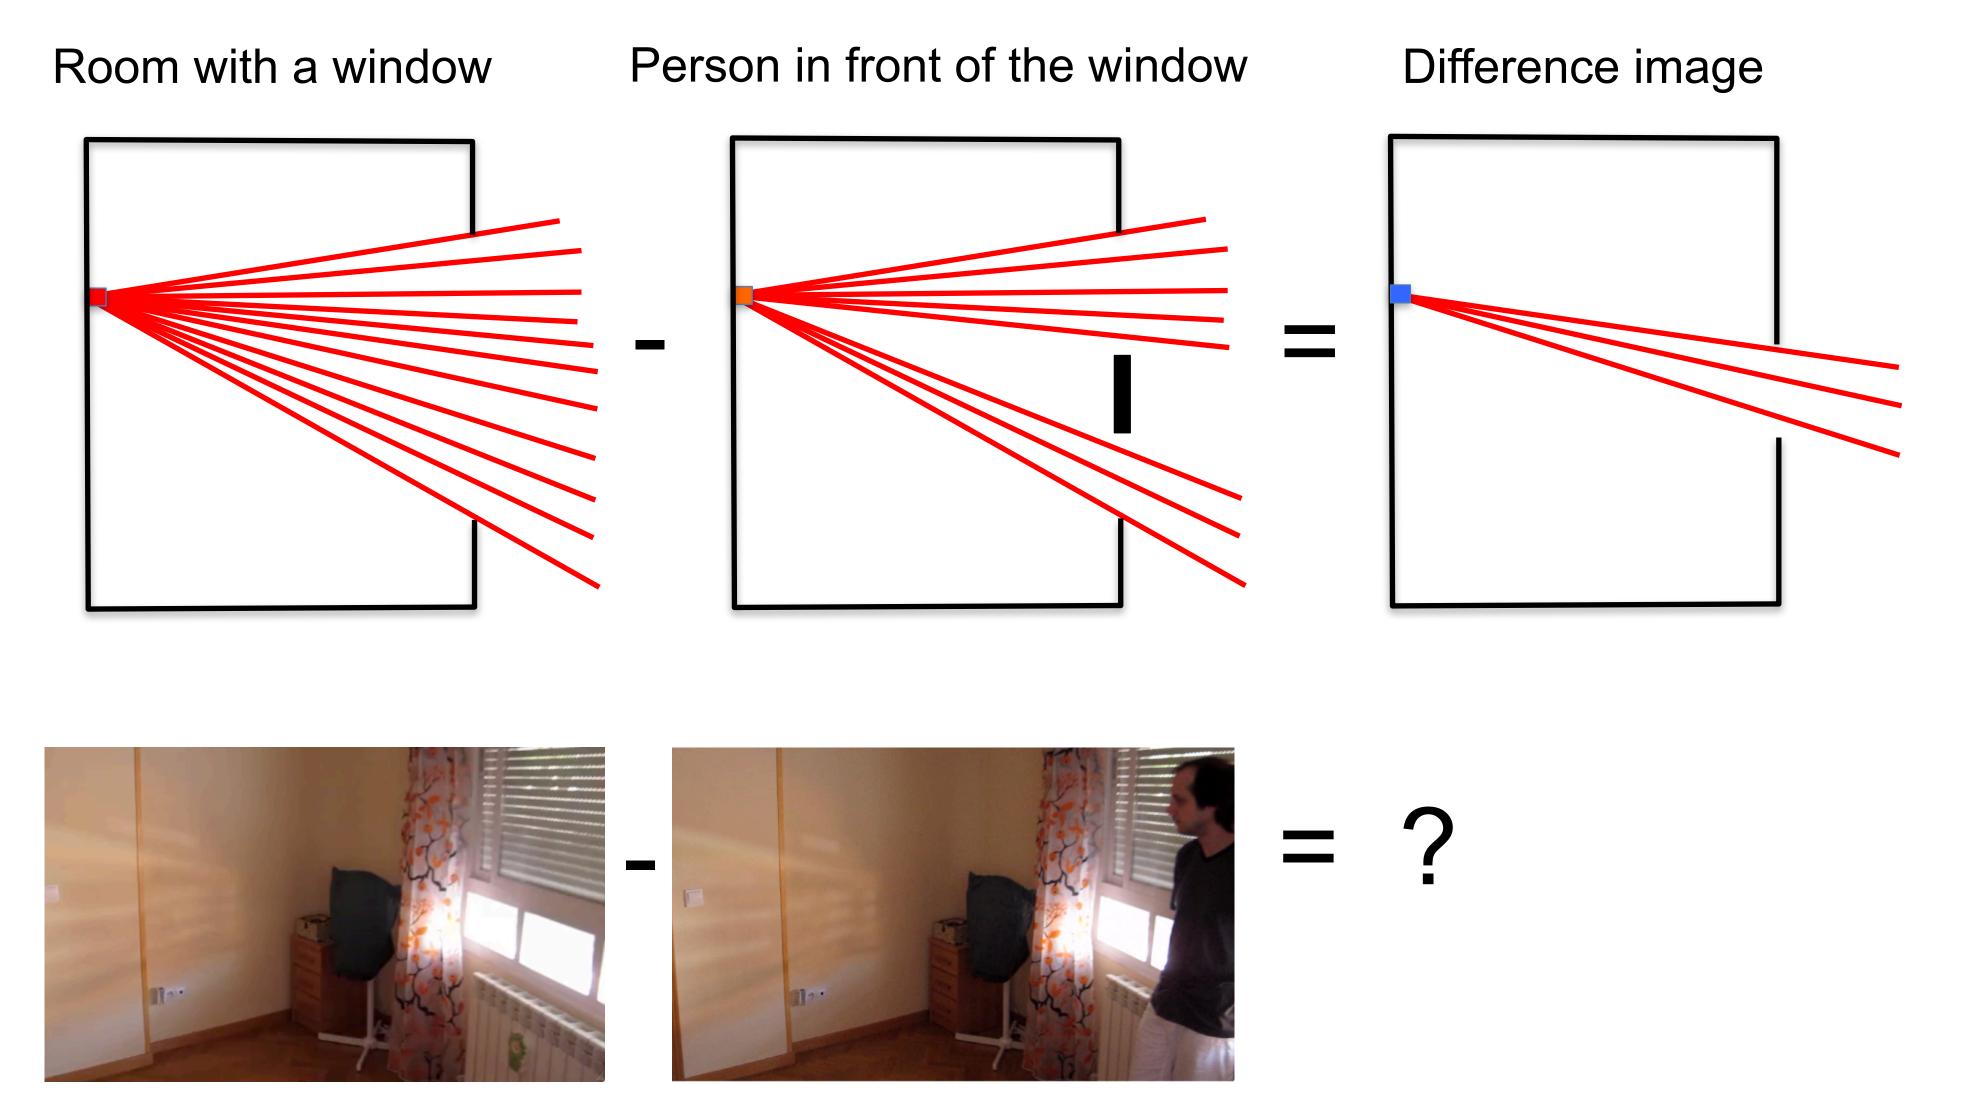




Mixed accidental pinhole and anti-pinhole cameras

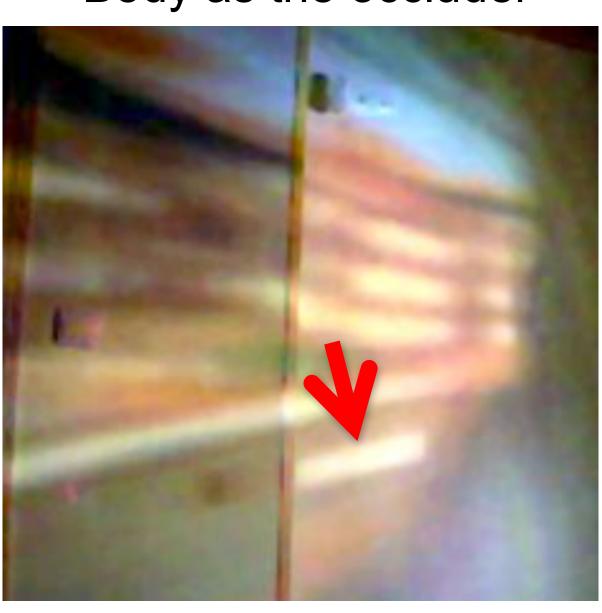


Mixed accidental pinhole and anti-pinhole cameras

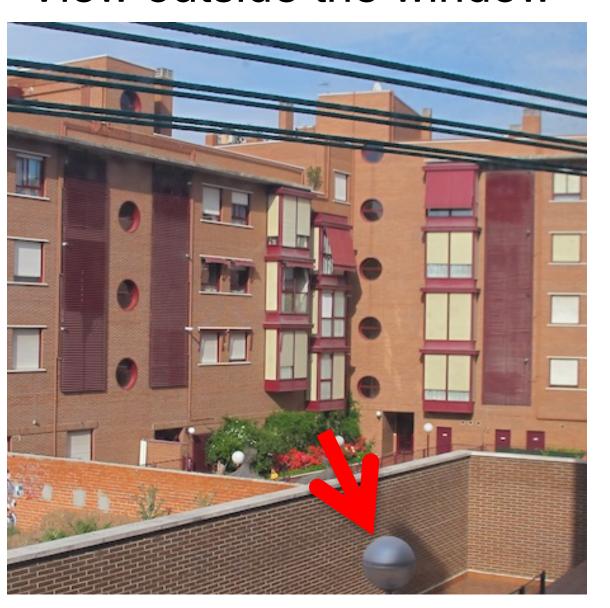


Mixed accidental pinhole and anti-pinhole cameras

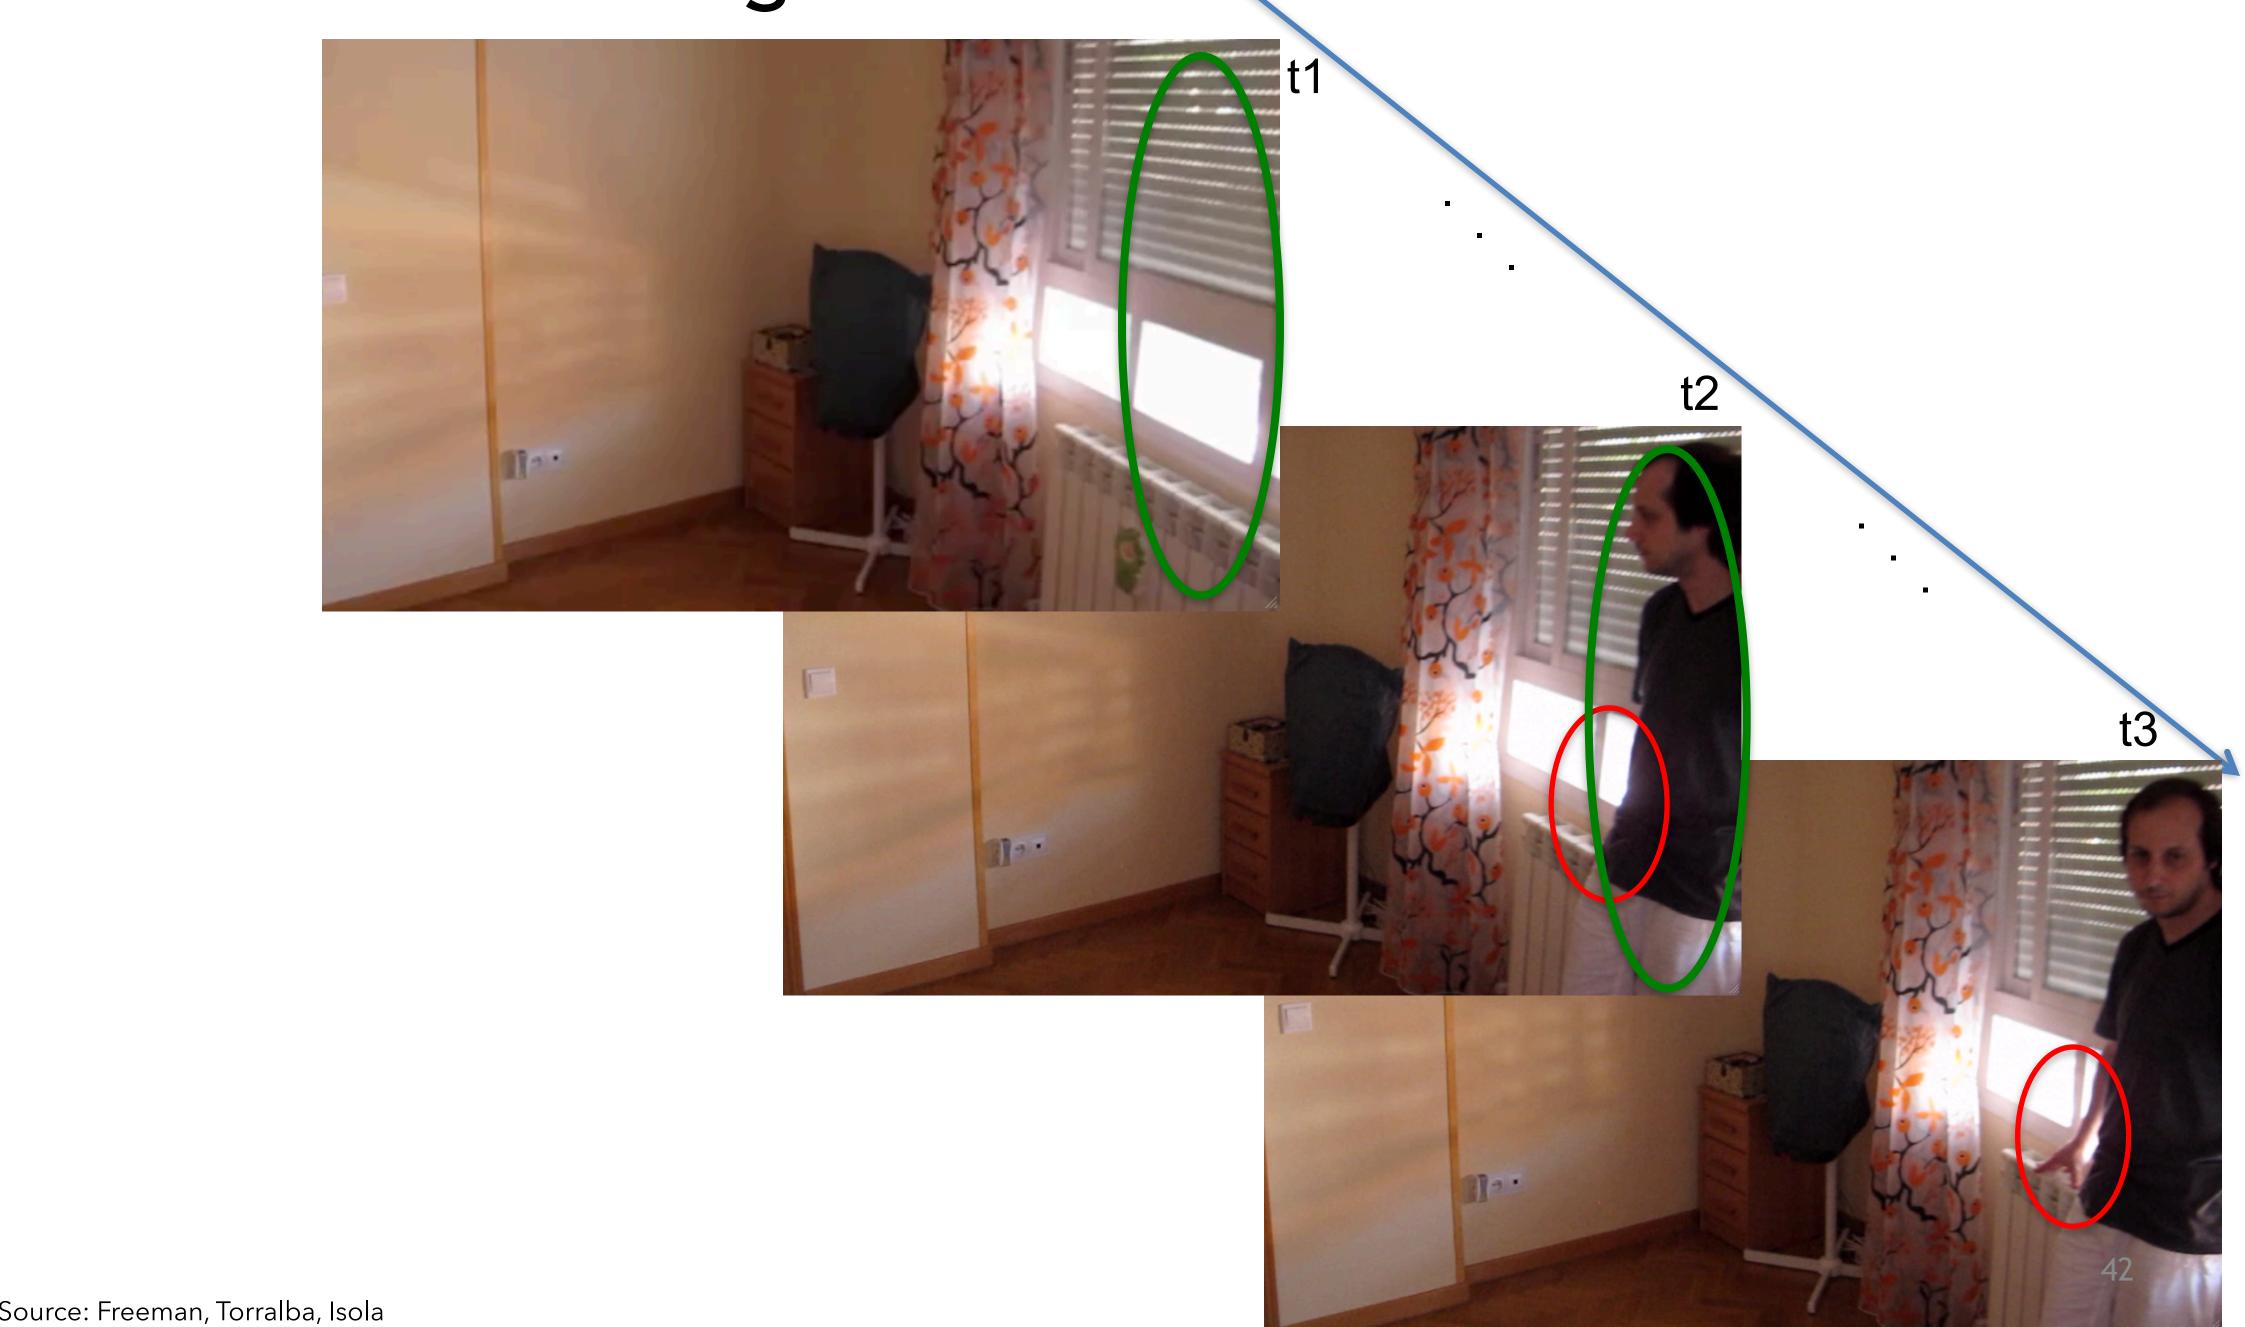
Body as the occluder



View outside the window

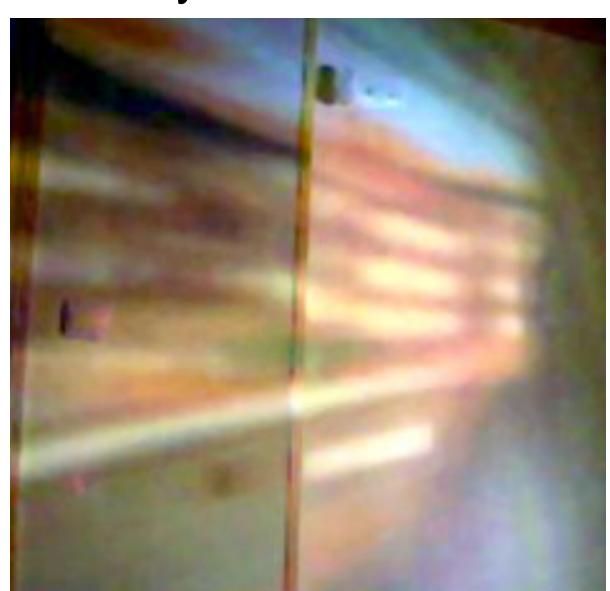


Looking for a small accidental occluder



Looking for a small accidental occluder

Body as the occluder



Hand as the occluder

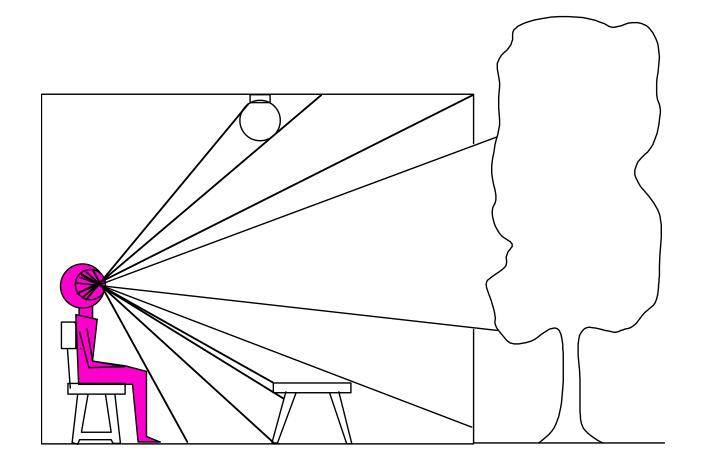


View outside the window

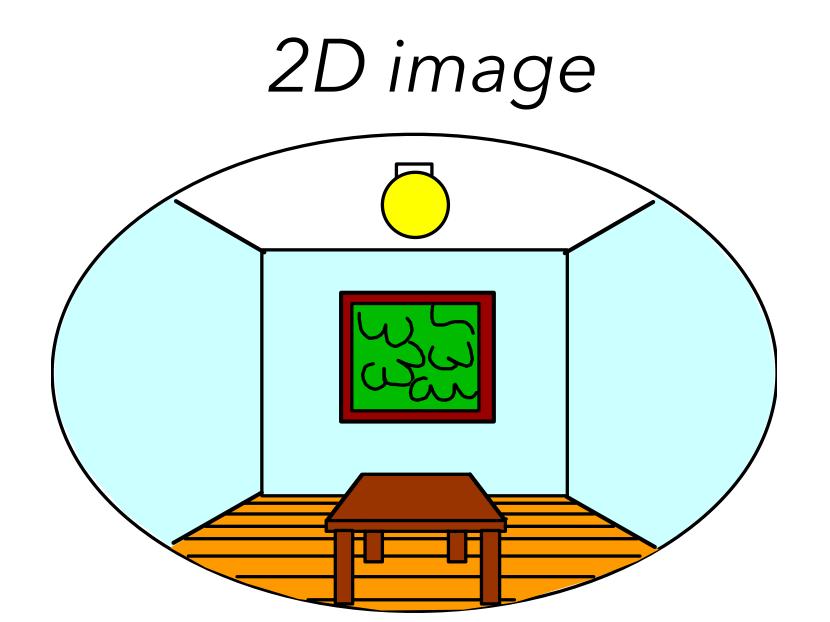


Projection from 3D to 2D

3D world



Point of observation

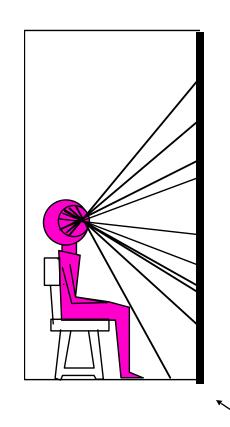


Slide by A. Efros

Figure source: Stephen E. Palmer, 2002

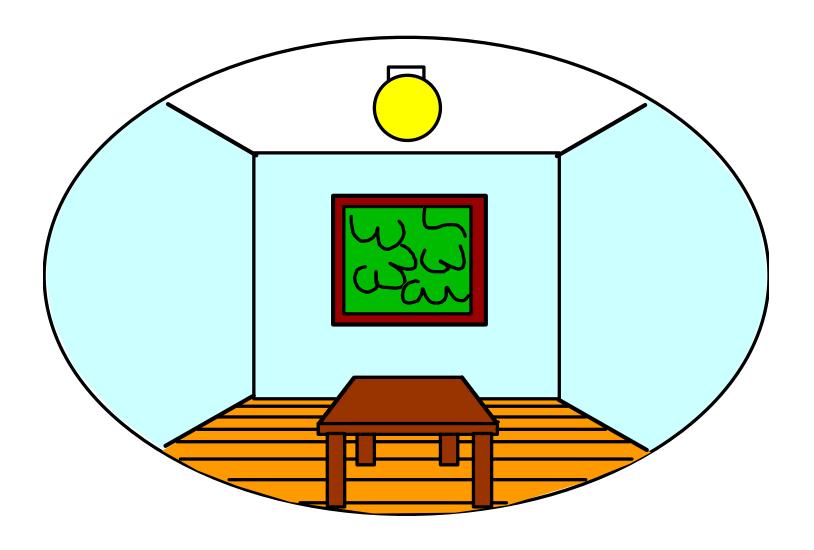
Projection from 3D to 2D

3D world



Painted backdrop?

2D image



Source: S. Lazebnik

Fooling the eye



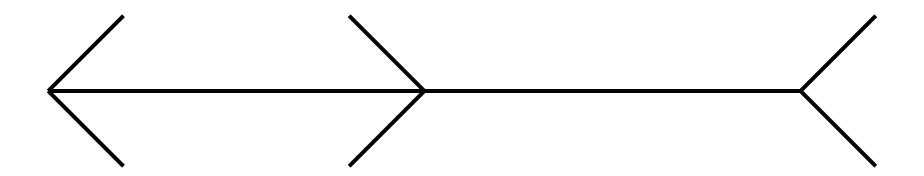
46

Source: S. Lazebnik

Fooling the eye

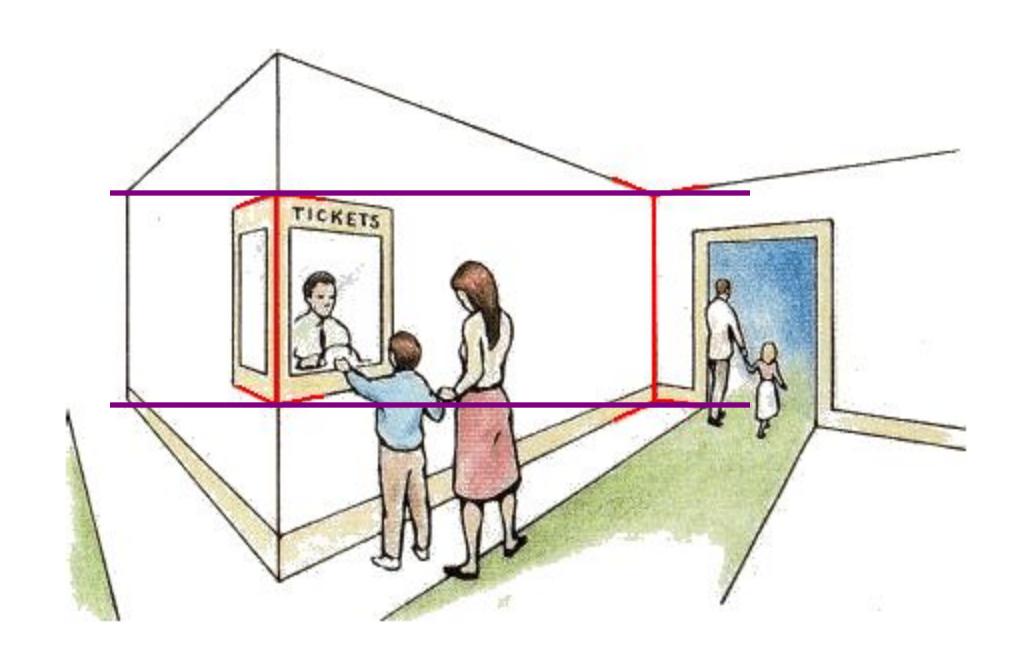


Müller-Lyer Illusion

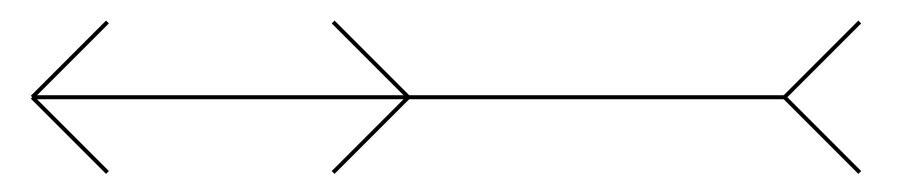


Source: N. Snavely, S. Lazebnik

Müller-Lyer Illusion



http://www.michaelbach.de/ot/sze_muelue/index.html



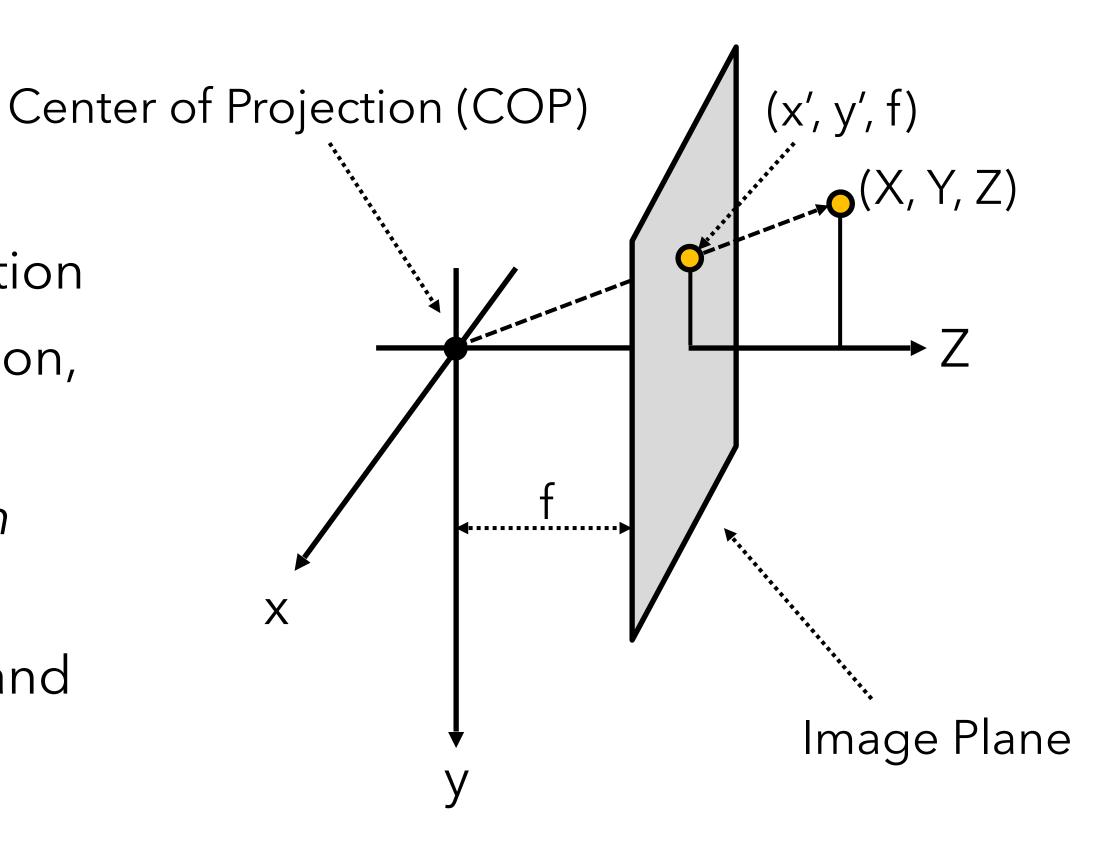
Source: N. Snavely, S. Lazebnik

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 virtual image idea)
- The camera looks down the positive z-axis, and the y-axis points down



Modeling projection

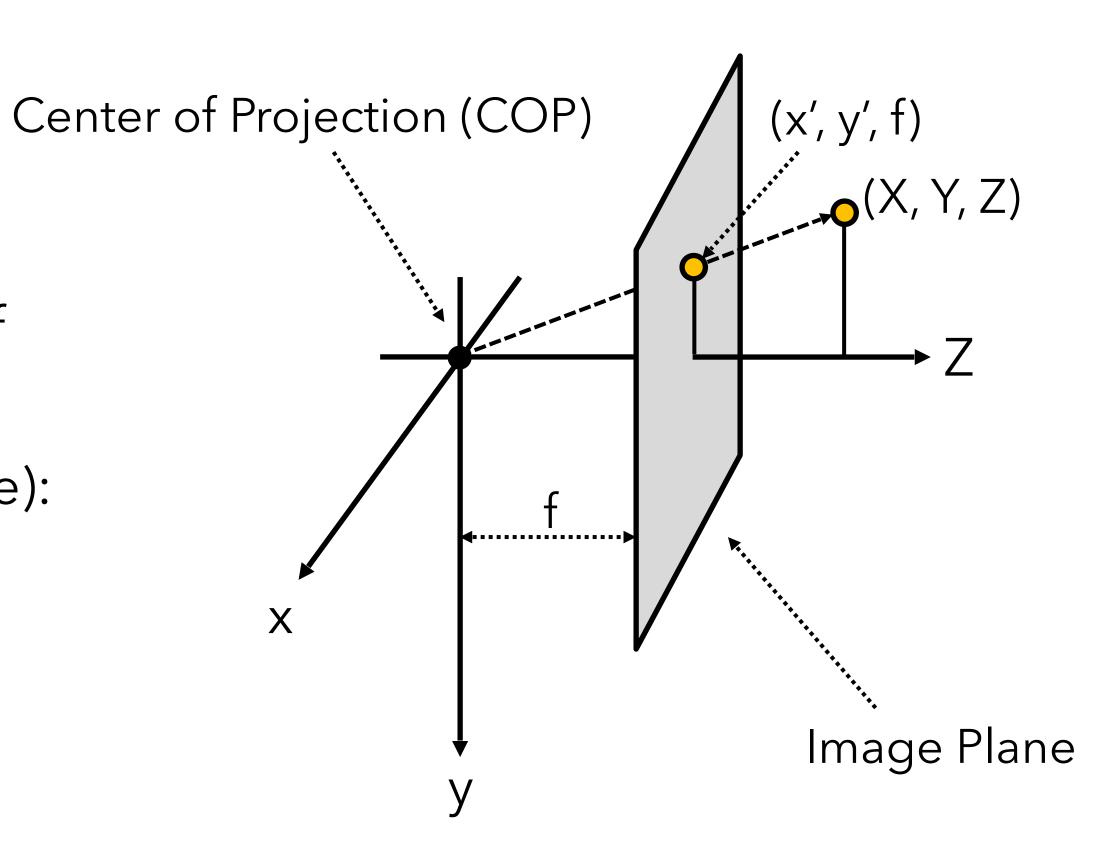
• Projection equations

- Compute intersection with image plane of ray from (X,Y,Z) to COP
- Derived using similar triangles (as we'll see):

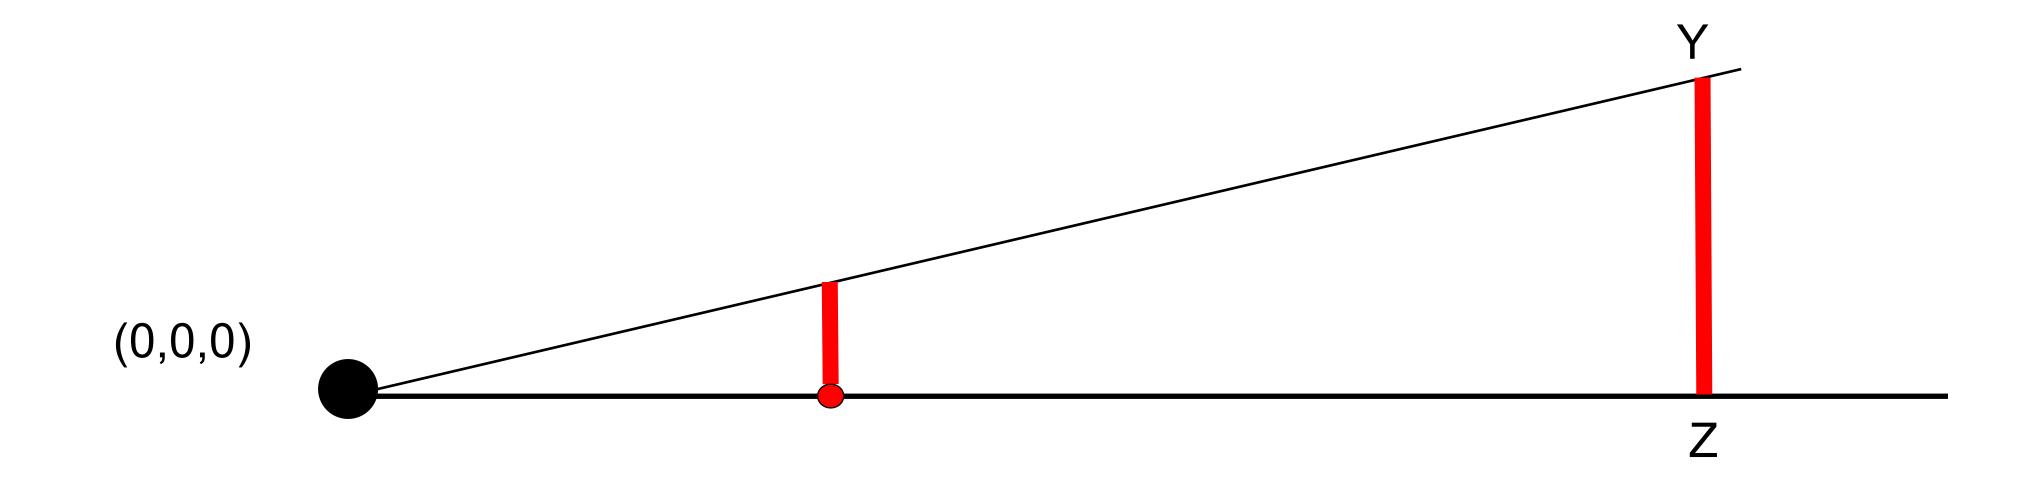
$$(x,y,z) \to (f\frac{x}{z},f\frac{y}{z},f)$$

 We get the projection by throwing out the last coordinate:

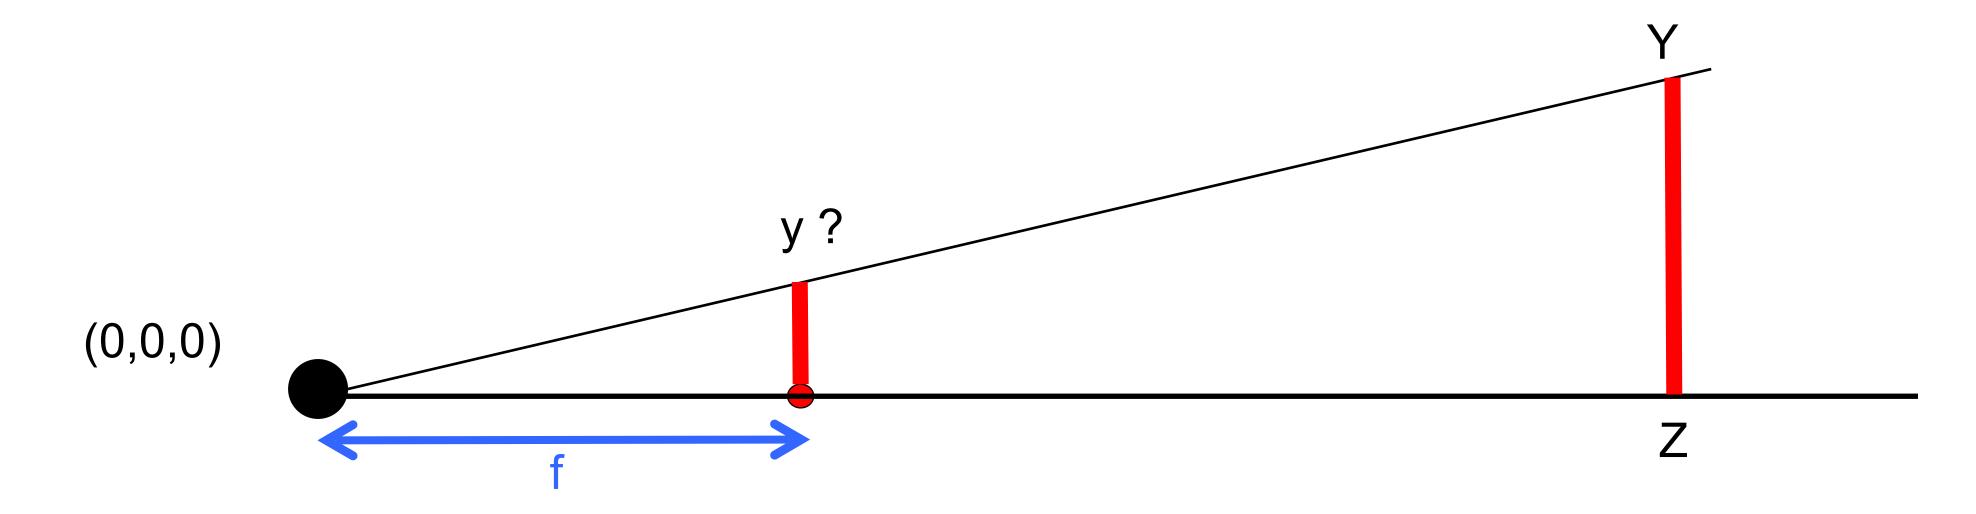
$$(x,y,z) \to (f\frac{x}{z},f\frac{y}{z})$$



Perspective projection (in 1D)



Perspective projection



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

 $y = f Y / Z$

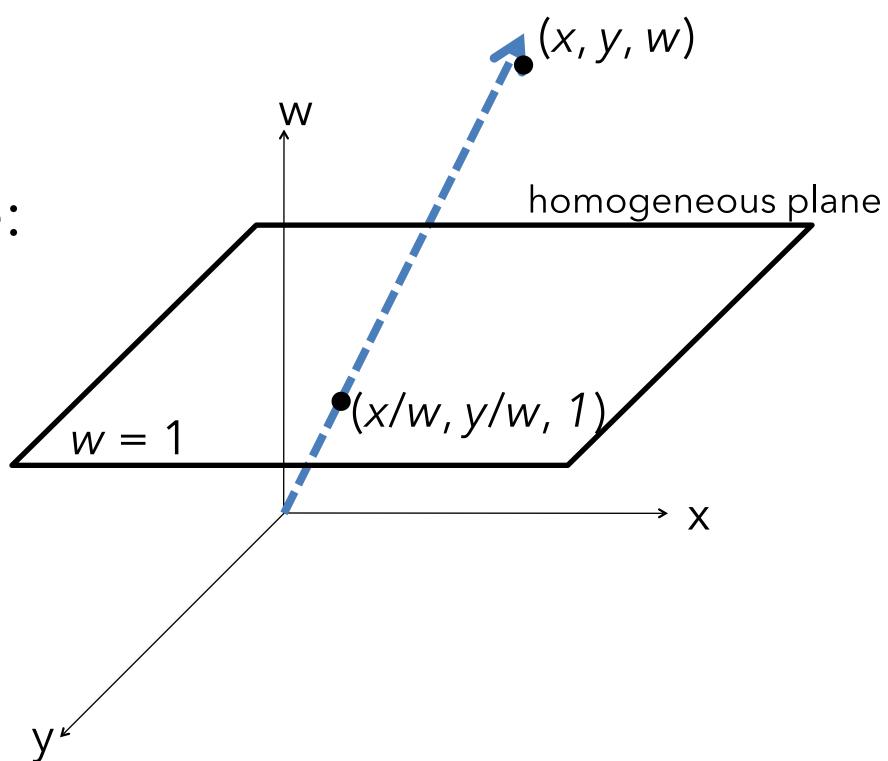
How can we represent this more compactly?

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

• Implement translation using a matrix multiplication

$$\mathbf{T} = \left[egin{array}{cccc} 1 & 0 & t_x \ 0 & 1 & t_y \ 0 & 0 & 1 \end{array}
ight]$$

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

Affine transformations

$$\mathbf{T} = egin{bmatrix} 1 & 0 & t_x \ 0 & 1 & t_y \ 0 & 0 & 1 \end{bmatrix}$$
 any transformation represented by a 3x3 matrix with last row [0 0 1] we call an affine transformation



we call an affine transformation

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

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

2D in-plane rotation

$$\begin{bmatrix} \mathbf{x'} \\ \mathbf{y'} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{s}_{x} & 0 & 0 \\ 0 & \mathbf{s}_{y} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ 1 \end{bmatrix}$$

Scale

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

Shear

Perspective Projection

Projection is a matrix multiply using homogeneous coordinates:

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

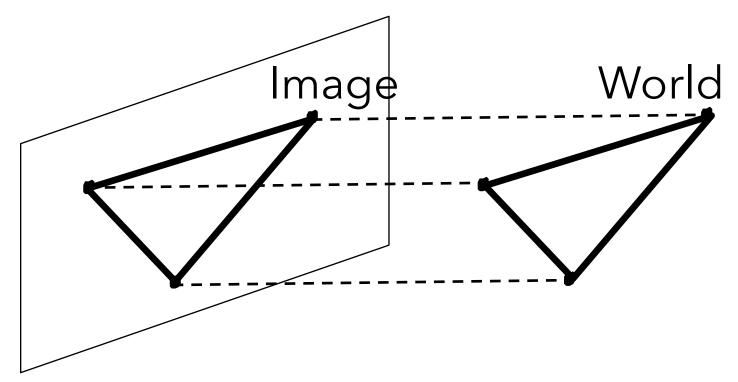
$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/f & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z/f \end{bmatrix} \Rightarrow (f\frac{x}{z}, f\frac{y}{z})$$

What if we scale by
$$f$$
?
$$\begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} fx \\ fy \\ z \end{bmatrix} \Rightarrow (f\frac{x}{z}, f\frac{y}{z})$$

Scaling a projection matrix produces an equivalent projection matrix!

Orthographic projection

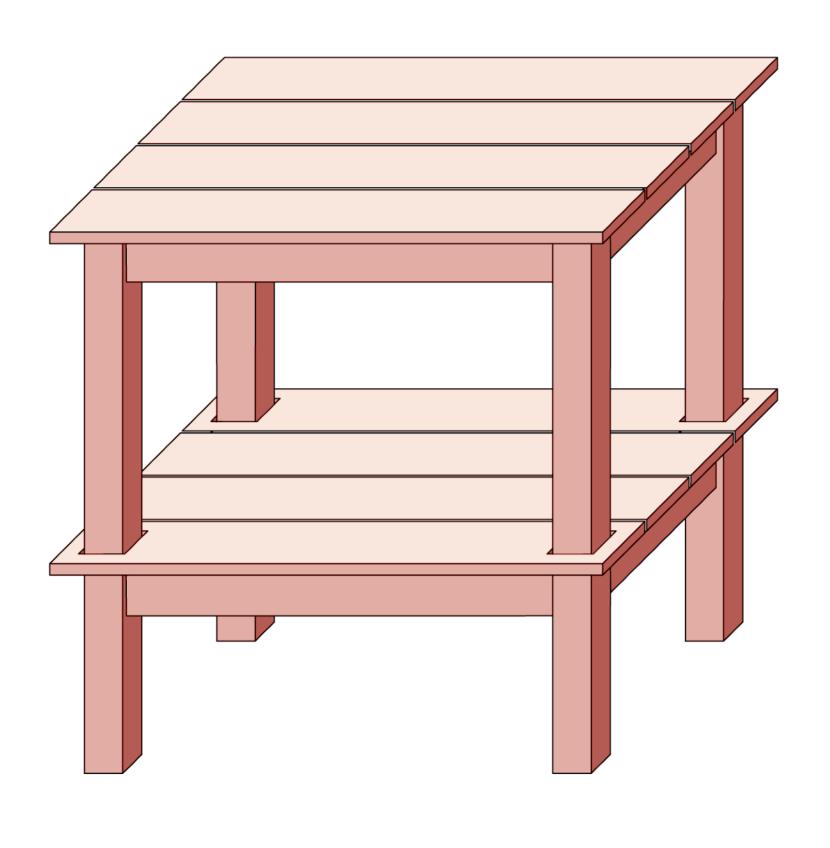
- Special case of perspective projection
 - Distance from the COP to the PP is infinite



- Good approximation for telephoto optics
- Also called "parallel projection": $(x, y, z) \rightarrow (x, y)$

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

Orthographic projection







Perspective projection





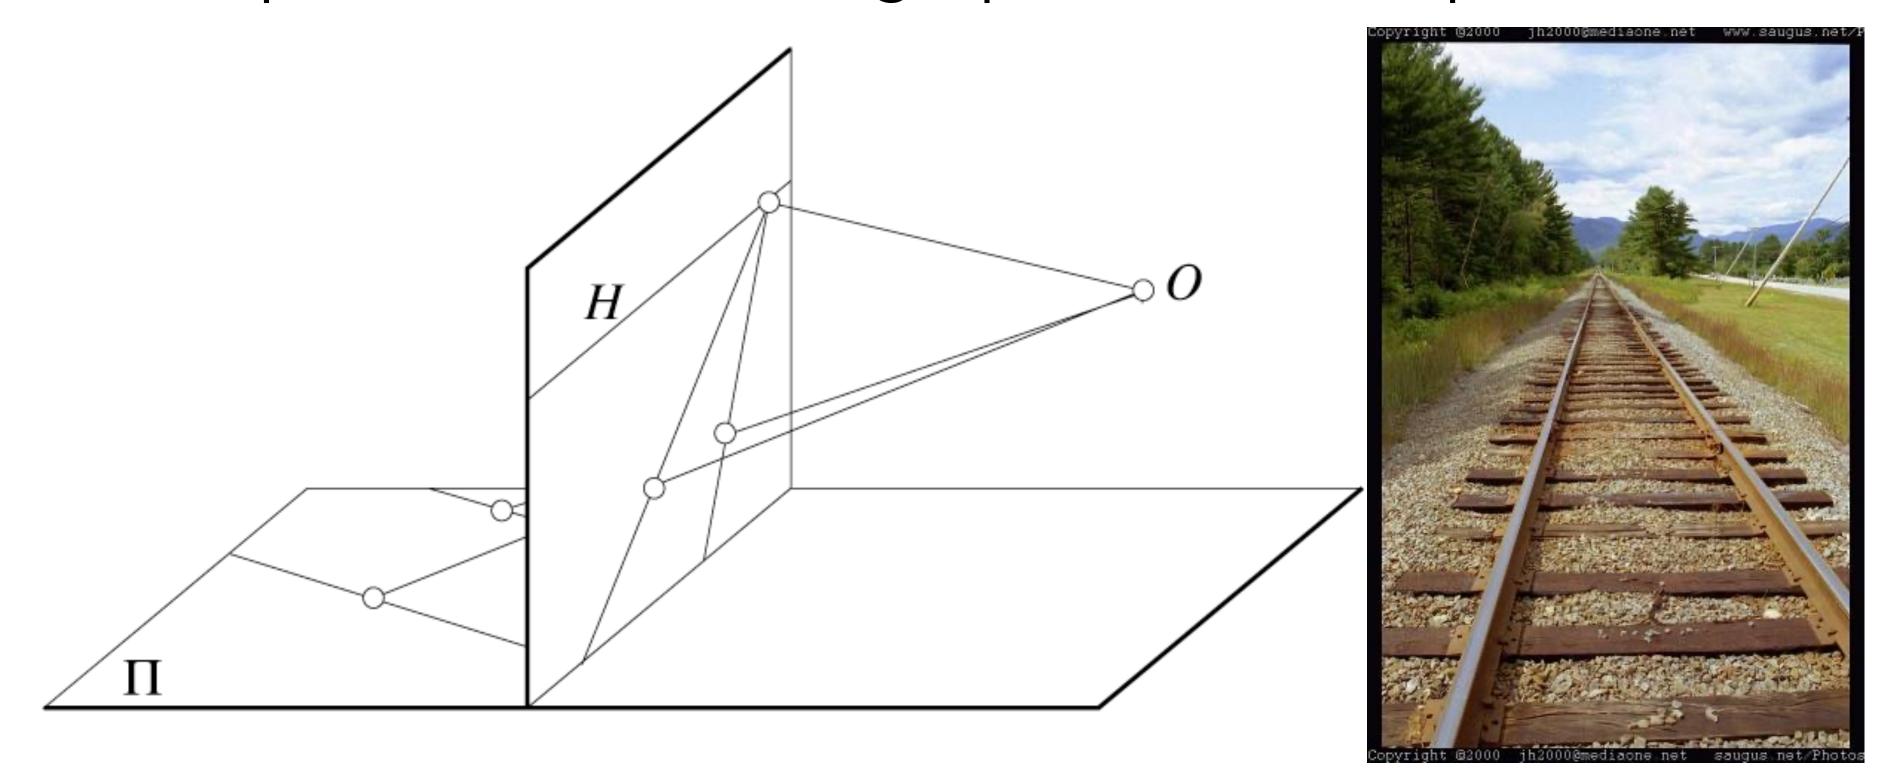


Projection properties

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

Projection properties

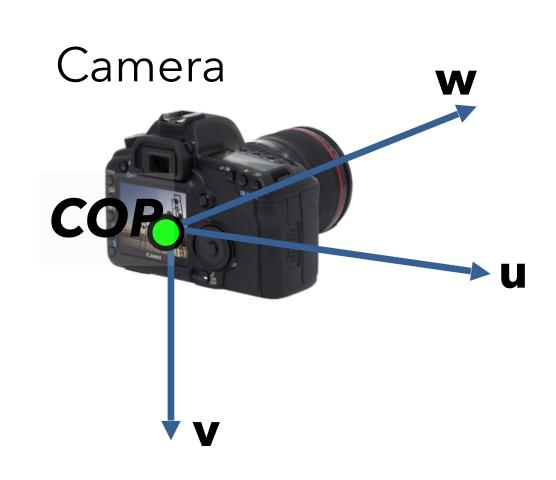
- Parallel lines converge at a vanishing point
 - Each direction in space has its own vanishing point
 - But lines parallel to the image plane remain parallel



Source: N.⁶\$navely

Camera parameters

How can we model the geometry of a camera?



Three important coordinate systems:

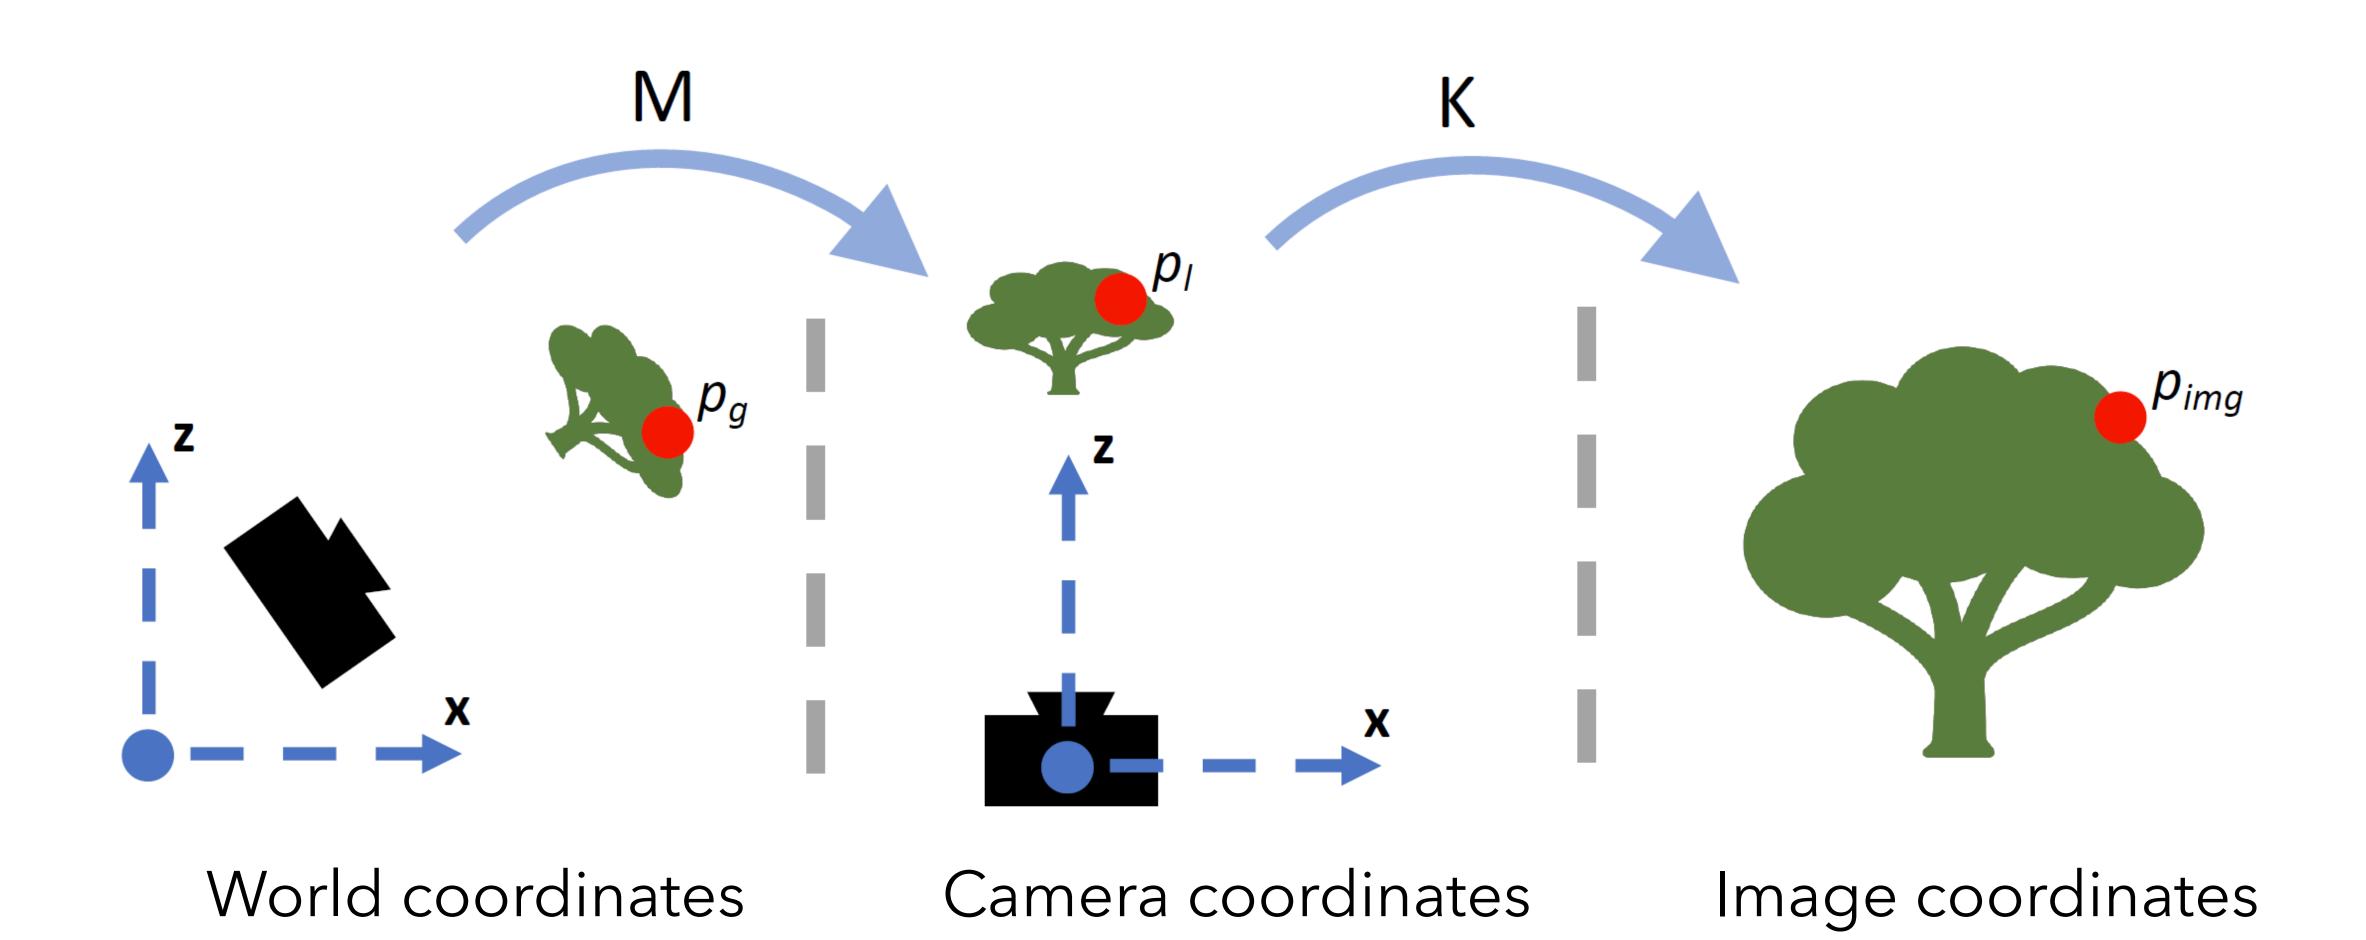
- 1. World coordinates
- Camera coordinates
- Image coordinates



"The World"

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

Coordinate frames



Source: N. Snavely Figure credit: Peter Hedman

Camera parameters

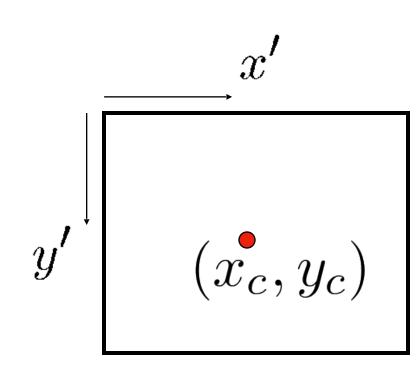
To project a point (x, y, z) in world coordinates into a camera

- 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 intrinsics

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 (x_c, y_c) , pixel aspect size α
- blue parameters are called "extrinsics," red are "intrinsics" Projection equation:



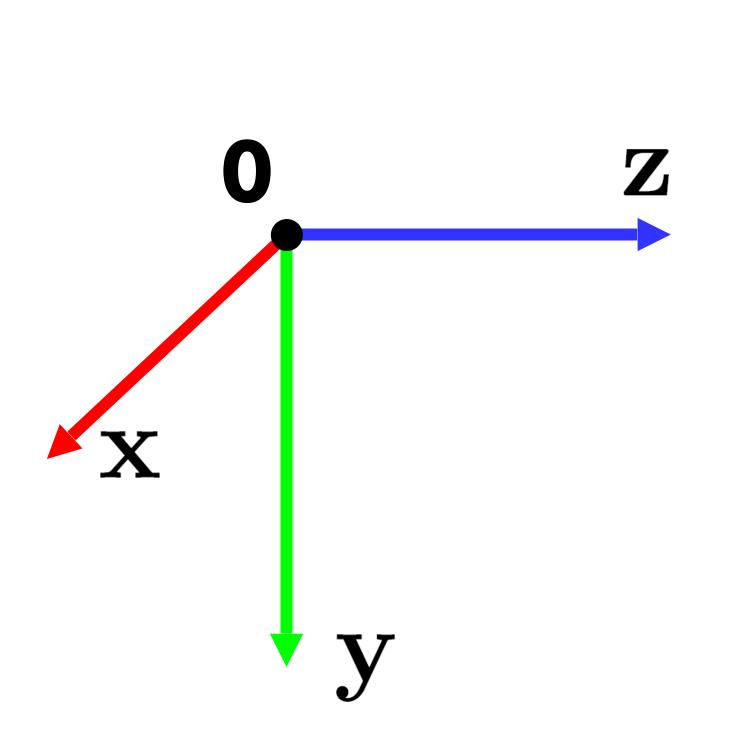
identity matrix

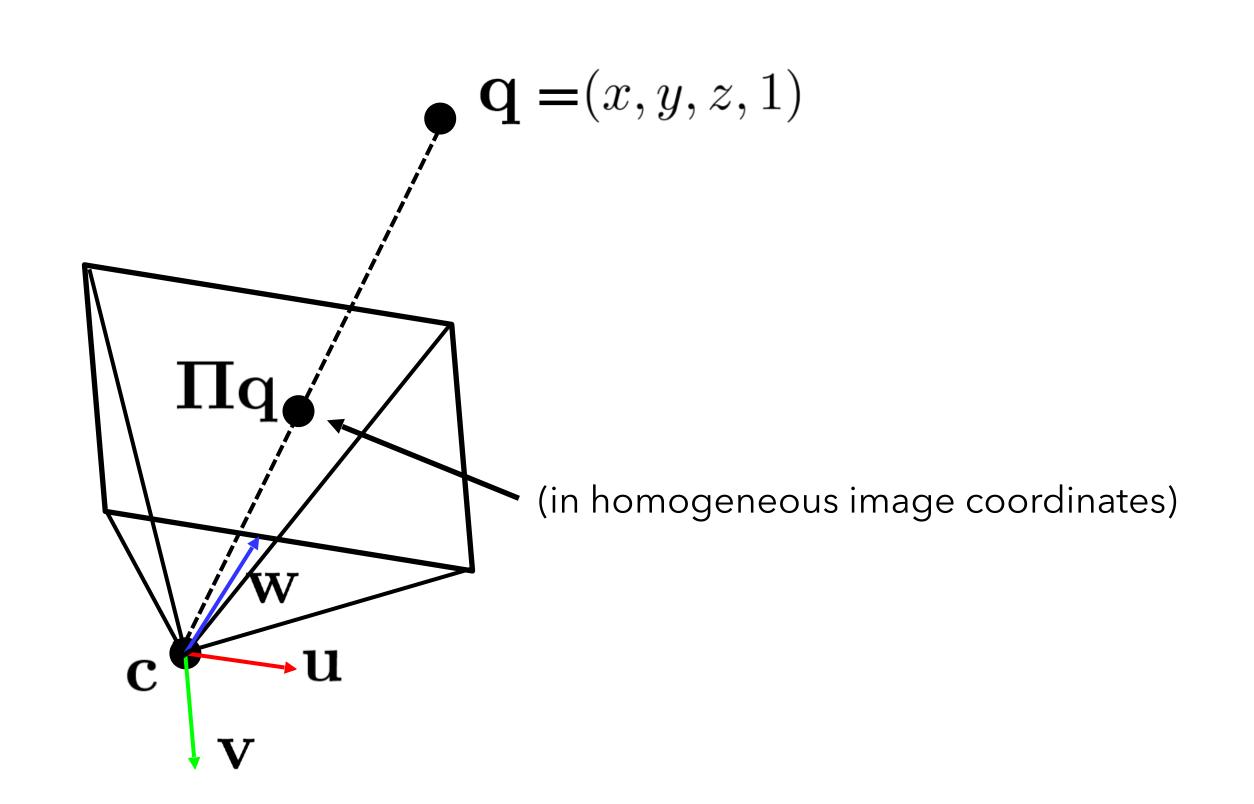
- The projection matrix models the cumulative effect of all parameters
- Useful to decompose into a series of operations

$$\boldsymbol{\Pi} = \begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R}_{3\times3} & \mathbf{0}_{3\times1} \\ \mathbf{0}_{1\times3} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3\times3} & \mathbf{T}_{3\times1} \\ \mathbf{0}_{1\times3} & 0 \end{bmatrix}$$
intrinsics projection rotation translation

• The definitions of these parameters are **not** completely standardized

Projection matrix

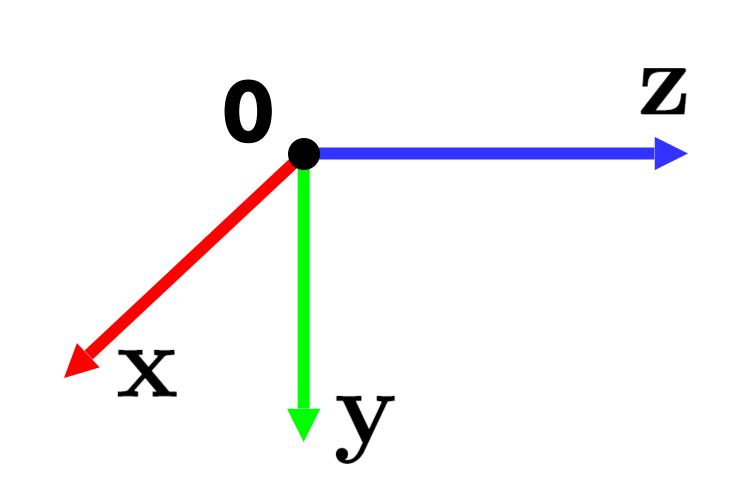


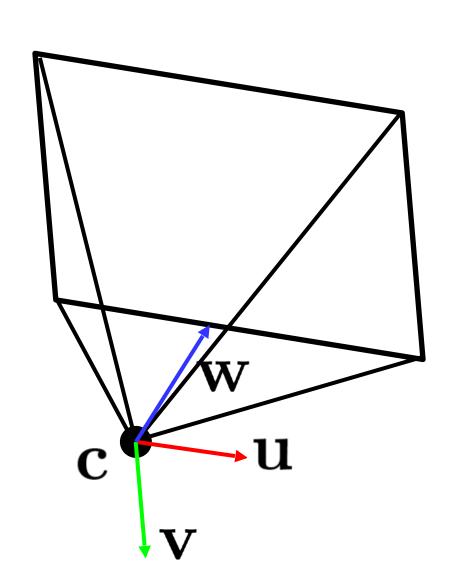


Extrinsics

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)

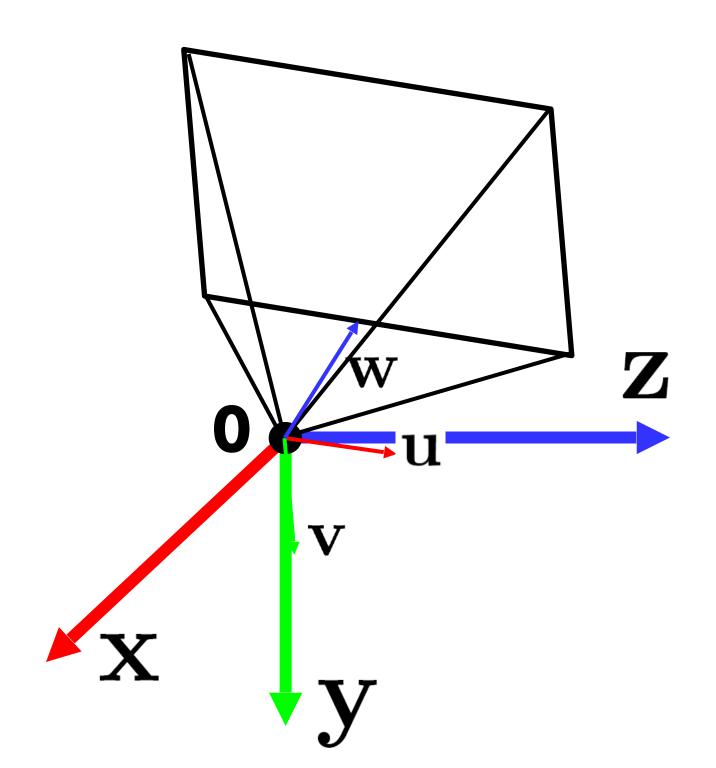
Step 1: Translate by -c





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)



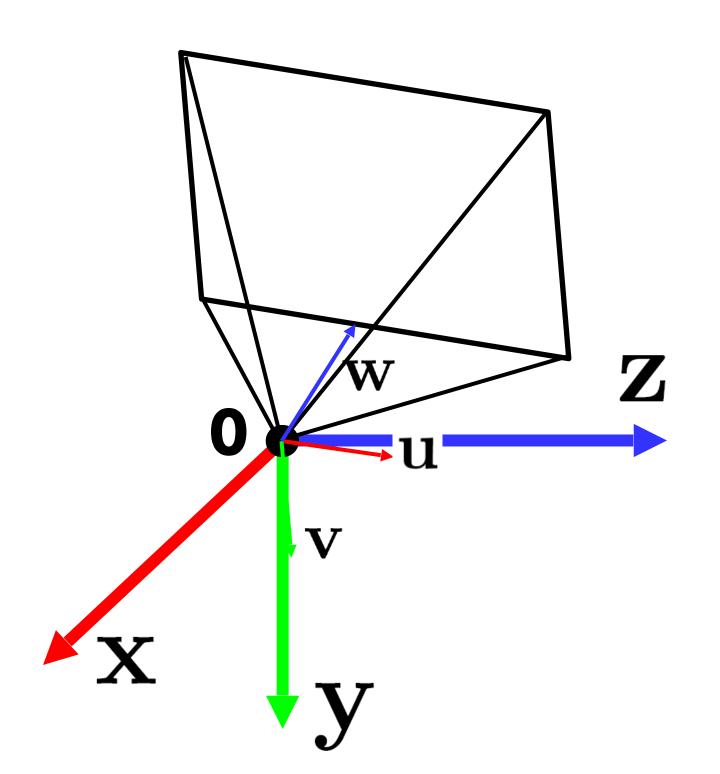
Step 1: Translate by -c

How do we represent translation as a matrix multiplication?

$$\mathbf{T} = \begin{bmatrix} \mathbf{I}_{3\times3} & -\mathbf{c} \\ 0 & 0 & 0 \end{bmatrix}$$

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)



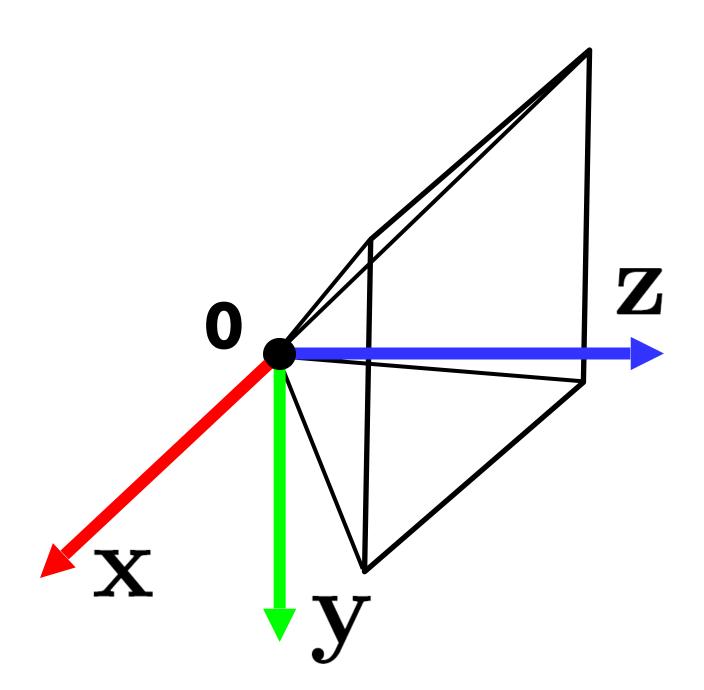
Step 1: Translate by -c

Step 2: Rotate by **R**

$$\mathbf{R} = \left[egin{array}{c} \mathbf{u}^T \ \mathbf{v}^T \ \mathbf{w}^T \end{array}
ight]$$
 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 -c

Step 2: Rotate by R

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

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

Perspective projection

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

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

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

lpha: aspect ratio (1 unless pixels are not square)

 $S: \mathbf{skew}$ (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"





50mm



200mm



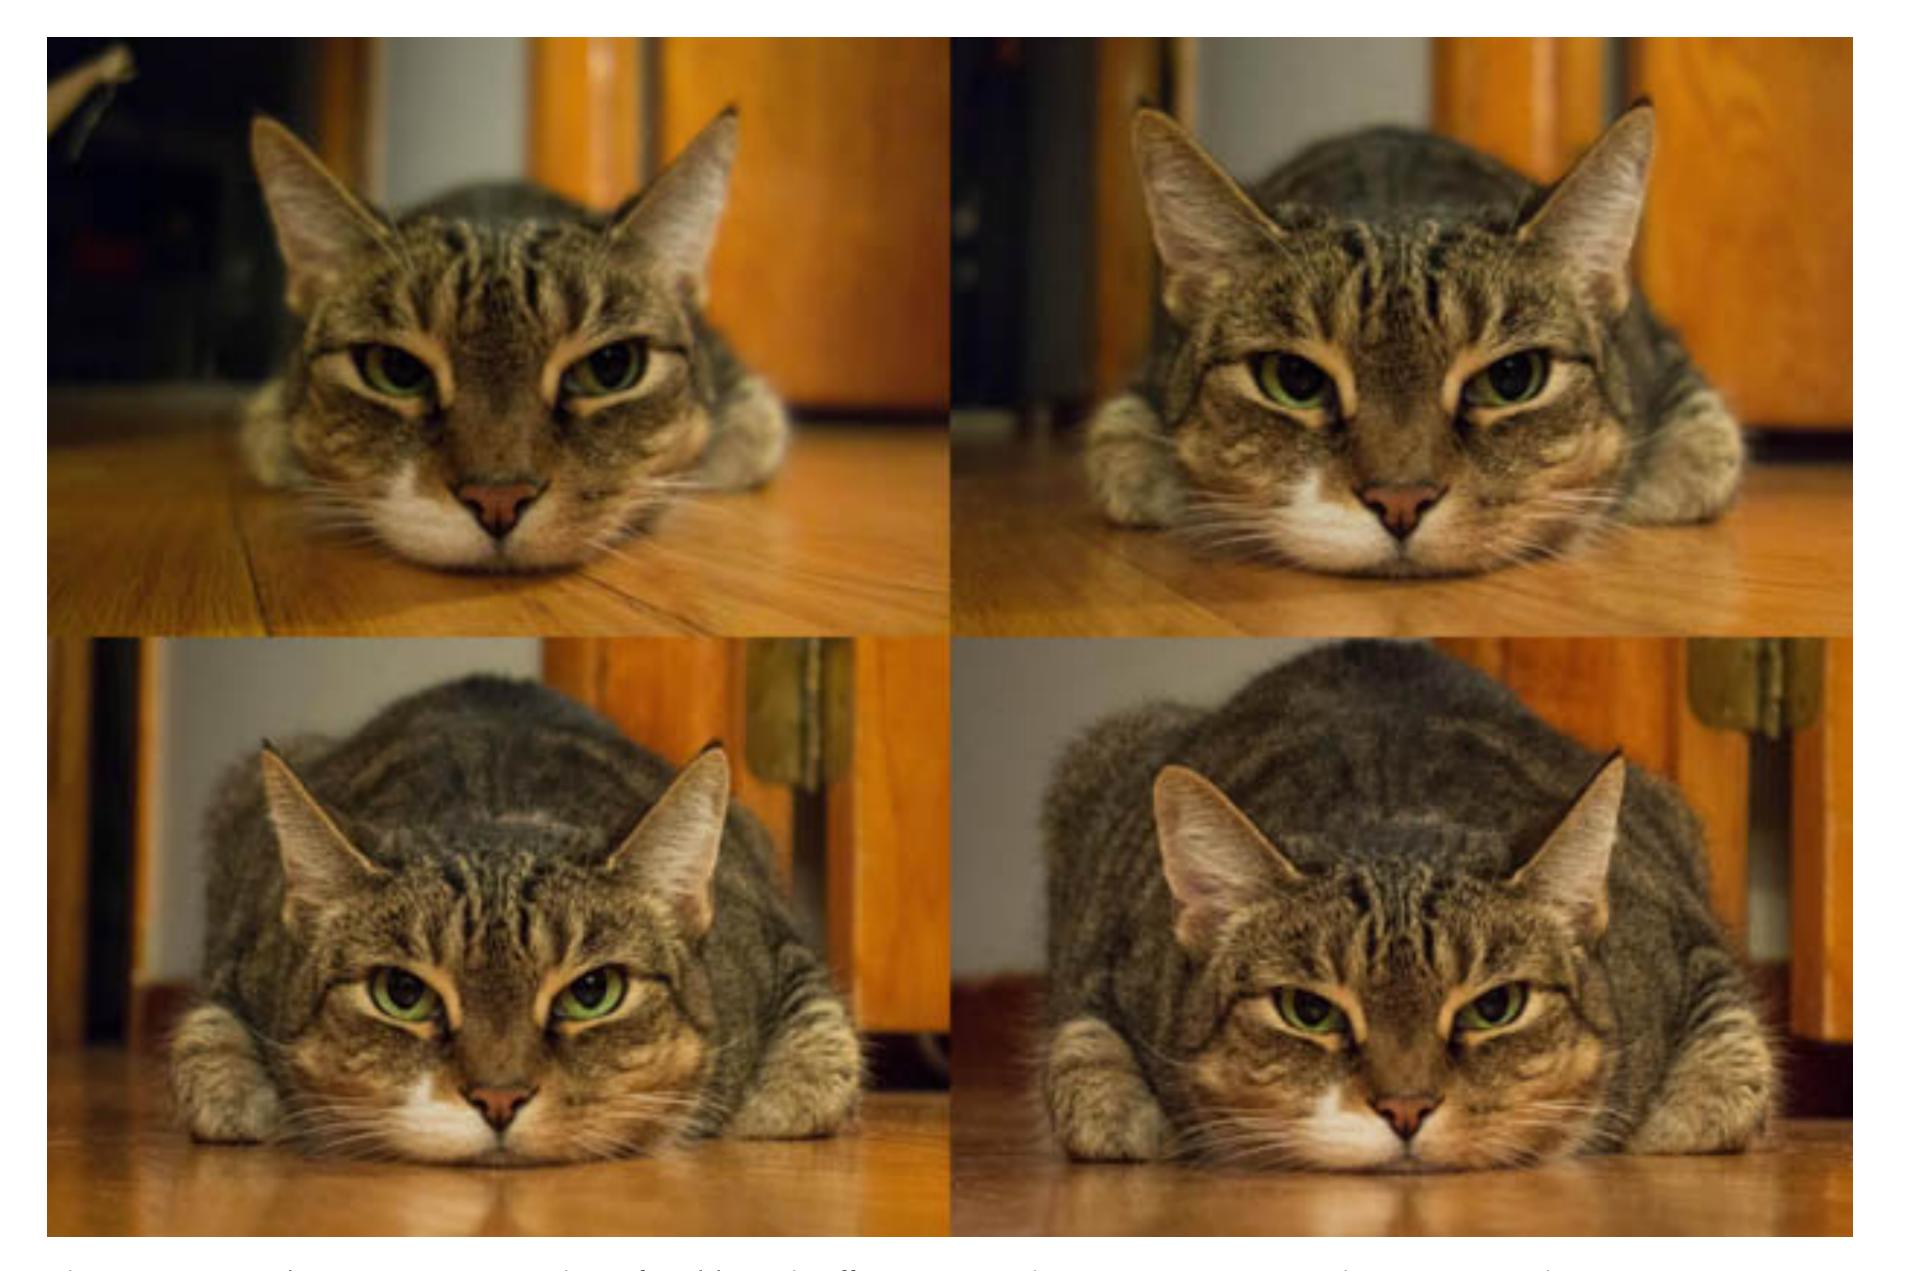
Also related to field of view

Changing focal length



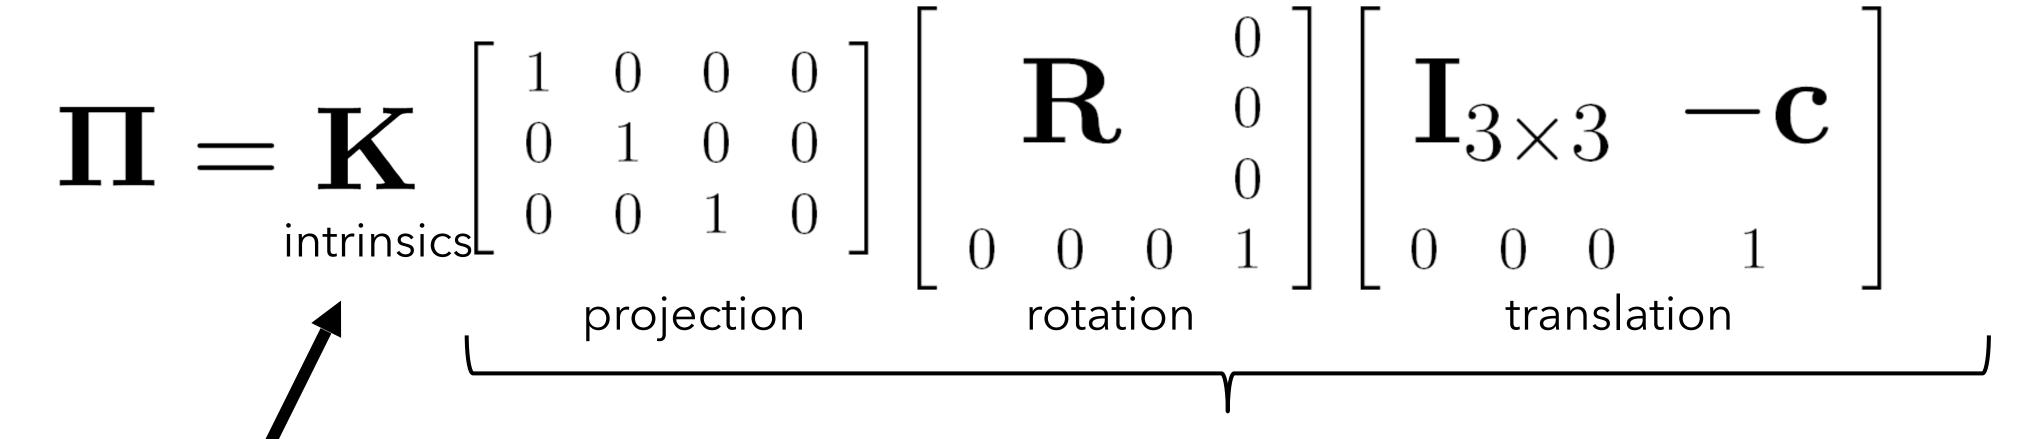
77

Source: N. Snavely



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

Projection matrix



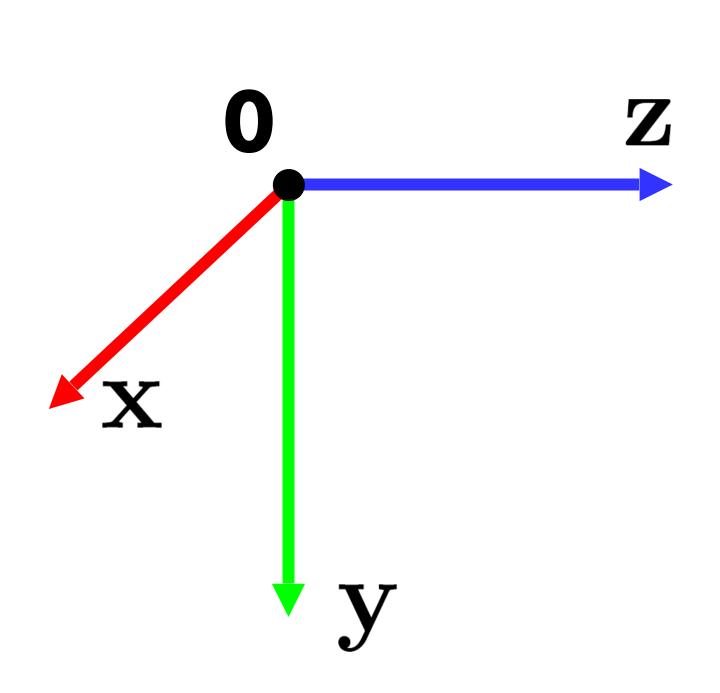
The **K** matrix converts 3D rays in the camera's coordinate system to 2D image points in image (pixel) coordinates.

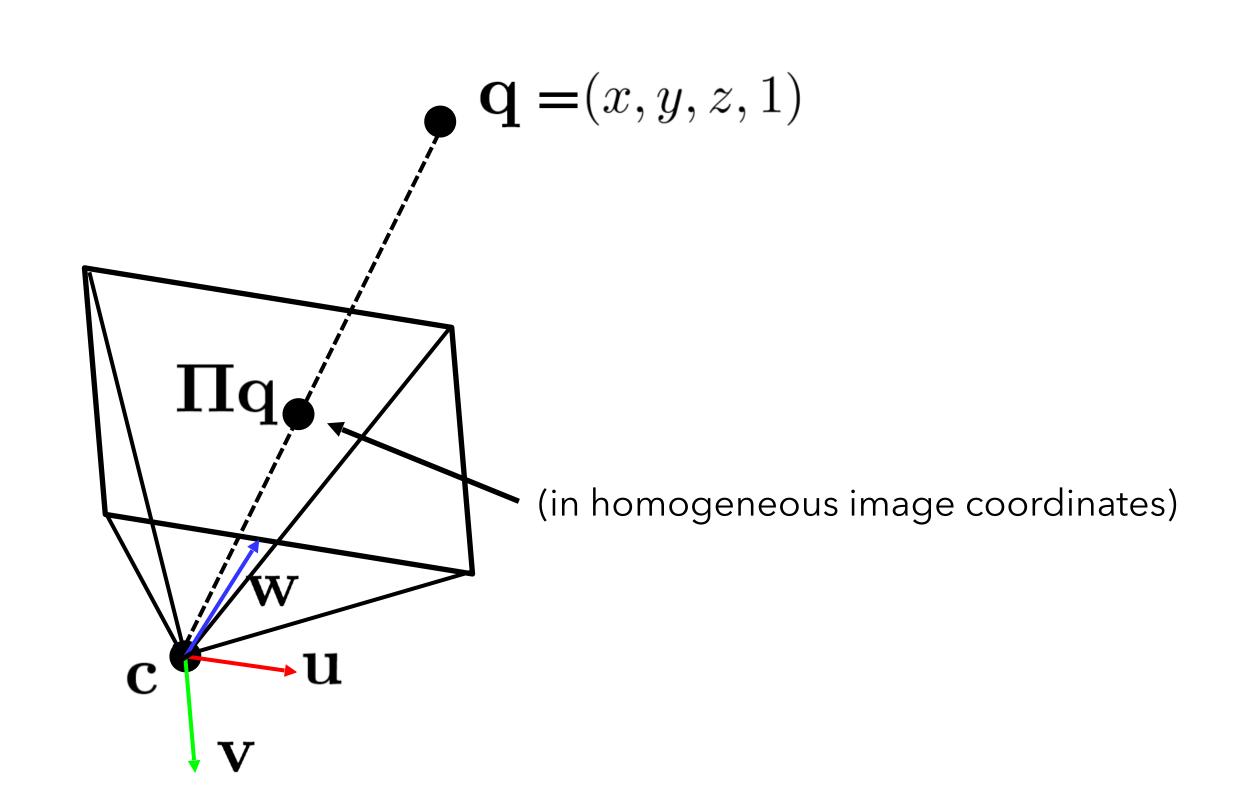
This part converts 3D points in world coordinates to 3D rays in the camera's coordinate system. There are 6 parameters represented (3 for position/translation, 3 for rotation).

Projection matrix

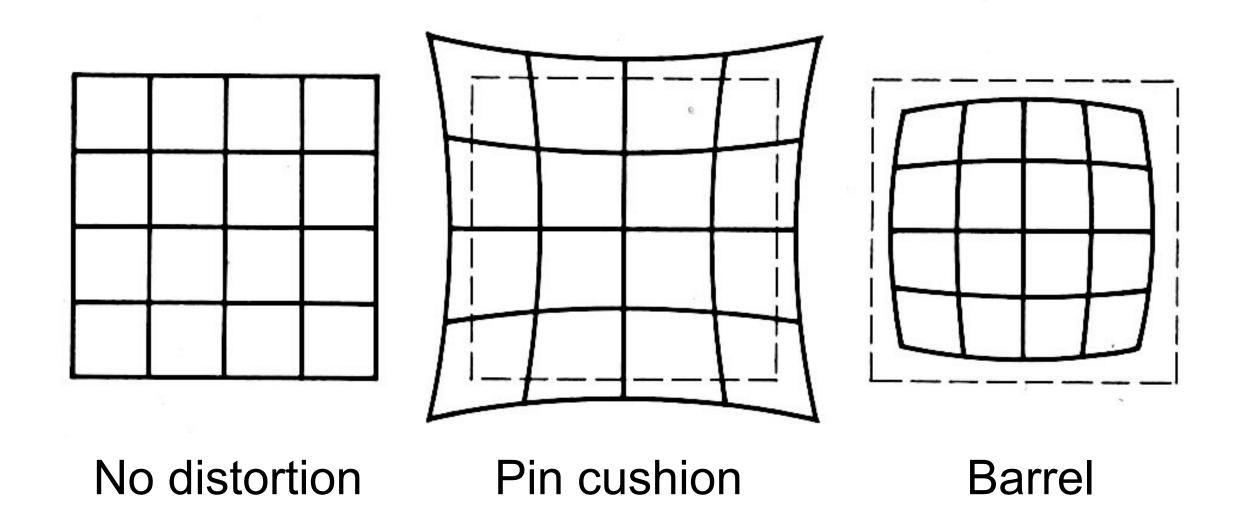
$$\boldsymbol{\Pi} = \mathbf{K} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R} & 0 \\ 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3\times3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
projection rotation translation
$$\begin{bmatrix} \mathbf{R} & -\mathbf{Rc} \end{bmatrix}$$
(sometimes called t)
$$\boldsymbol{\Pi} = \mathbf{K} \begin{bmatrix} \mathbf{R} & -\mathbf{Rc} \end{bmatrix}$$

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

Project
$$x_n' = \widehat{x}/\widehat{z}$$
 to "normalized" $y_n' = \widehat{y}/\widehat{z}$ image coordinates
$$r^2 = x_n'^2 + y_n'^2$$
 Apply radial distortion
$$x_d' = x_n'(1 + \kappa_1 r^2 + \kappa_2 r^4)$$

$$y_d' = y_n'(1 + \kappa_1 r^2 + \kappa_2 r^4)$$
 Apply focal length translate image center
$$x_d' = f x_d' + x_c$$

$$y' = f y_d' + y_c$$

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

Correcting radial distortion





from Helmut Dersch

Next class: More geometry!