MIT CSAIL

6.869: Advances in Computer Vision

Slides by Bill Freeman and Antonio Torralba
October 6, 2015

Image formation

Image formation

3D world

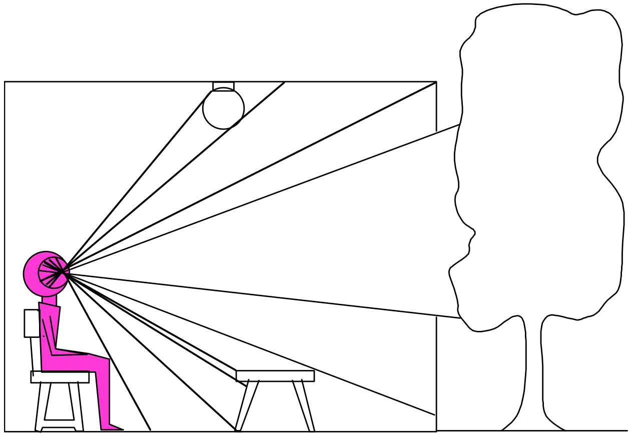

Point of observation

2D image

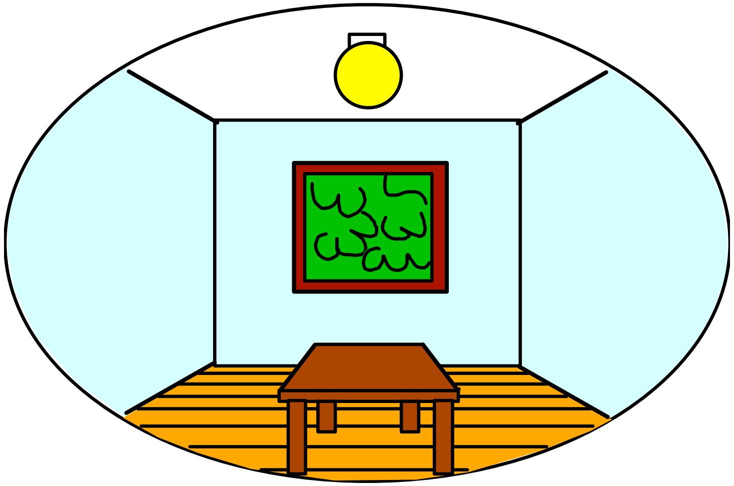

Cameras and lenses

- Camera models
- Projection equations

Images are projections of the 3-D world onto a 2-D plane…

The structure of ambient light

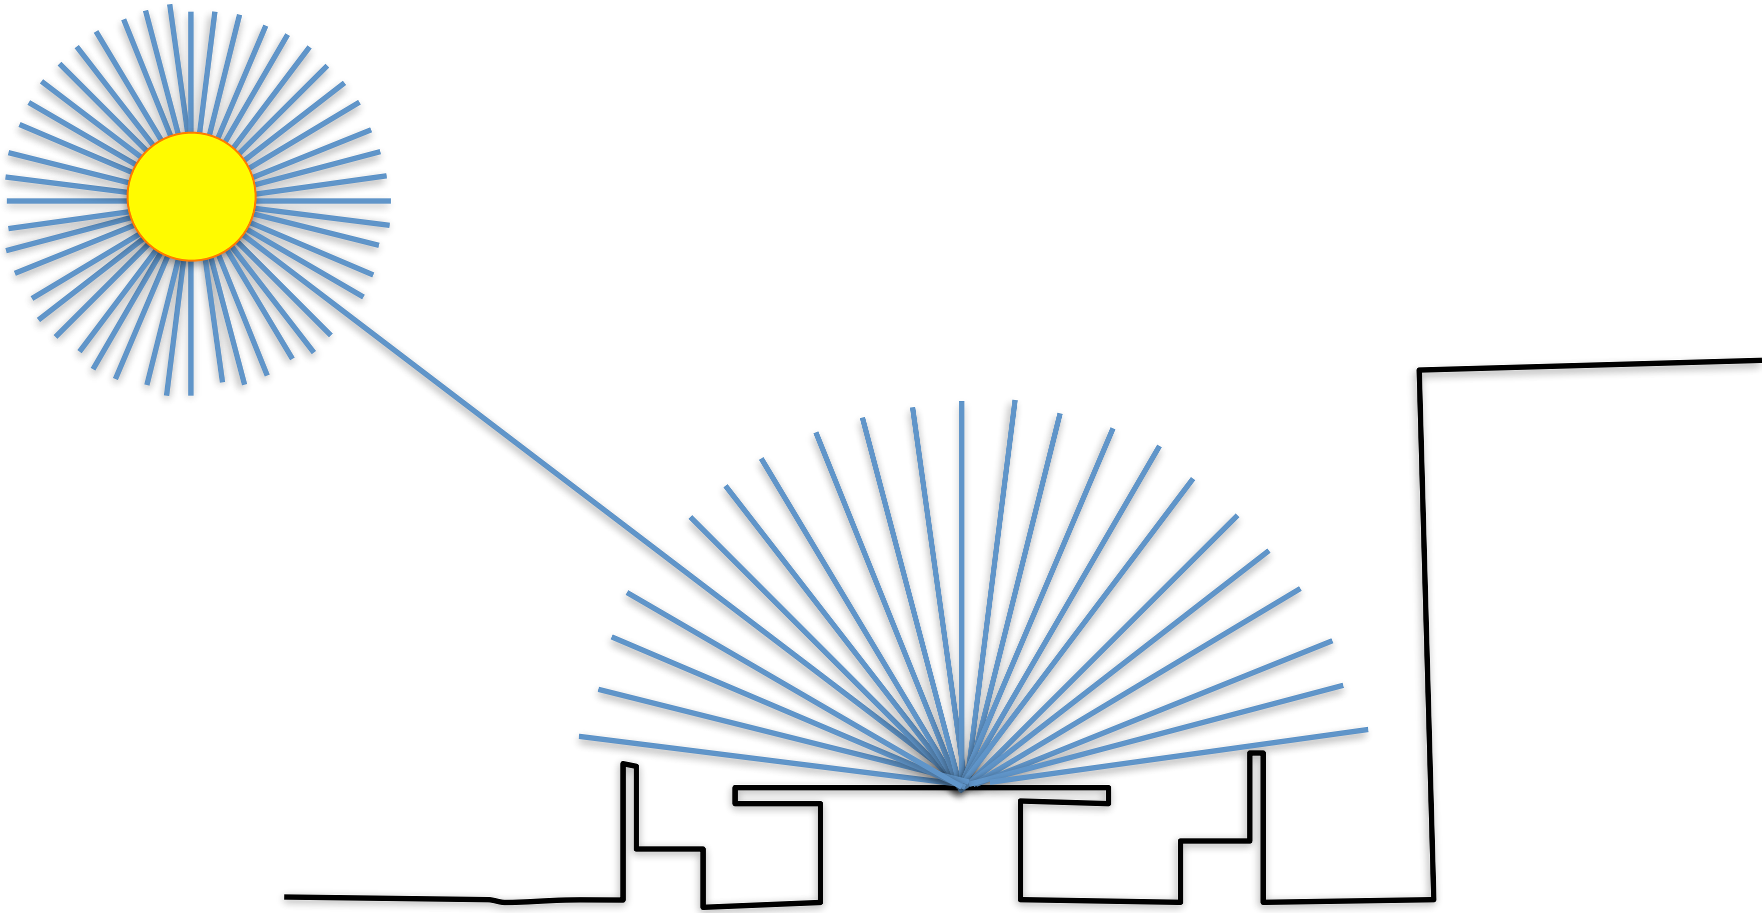

The structure of ambient light

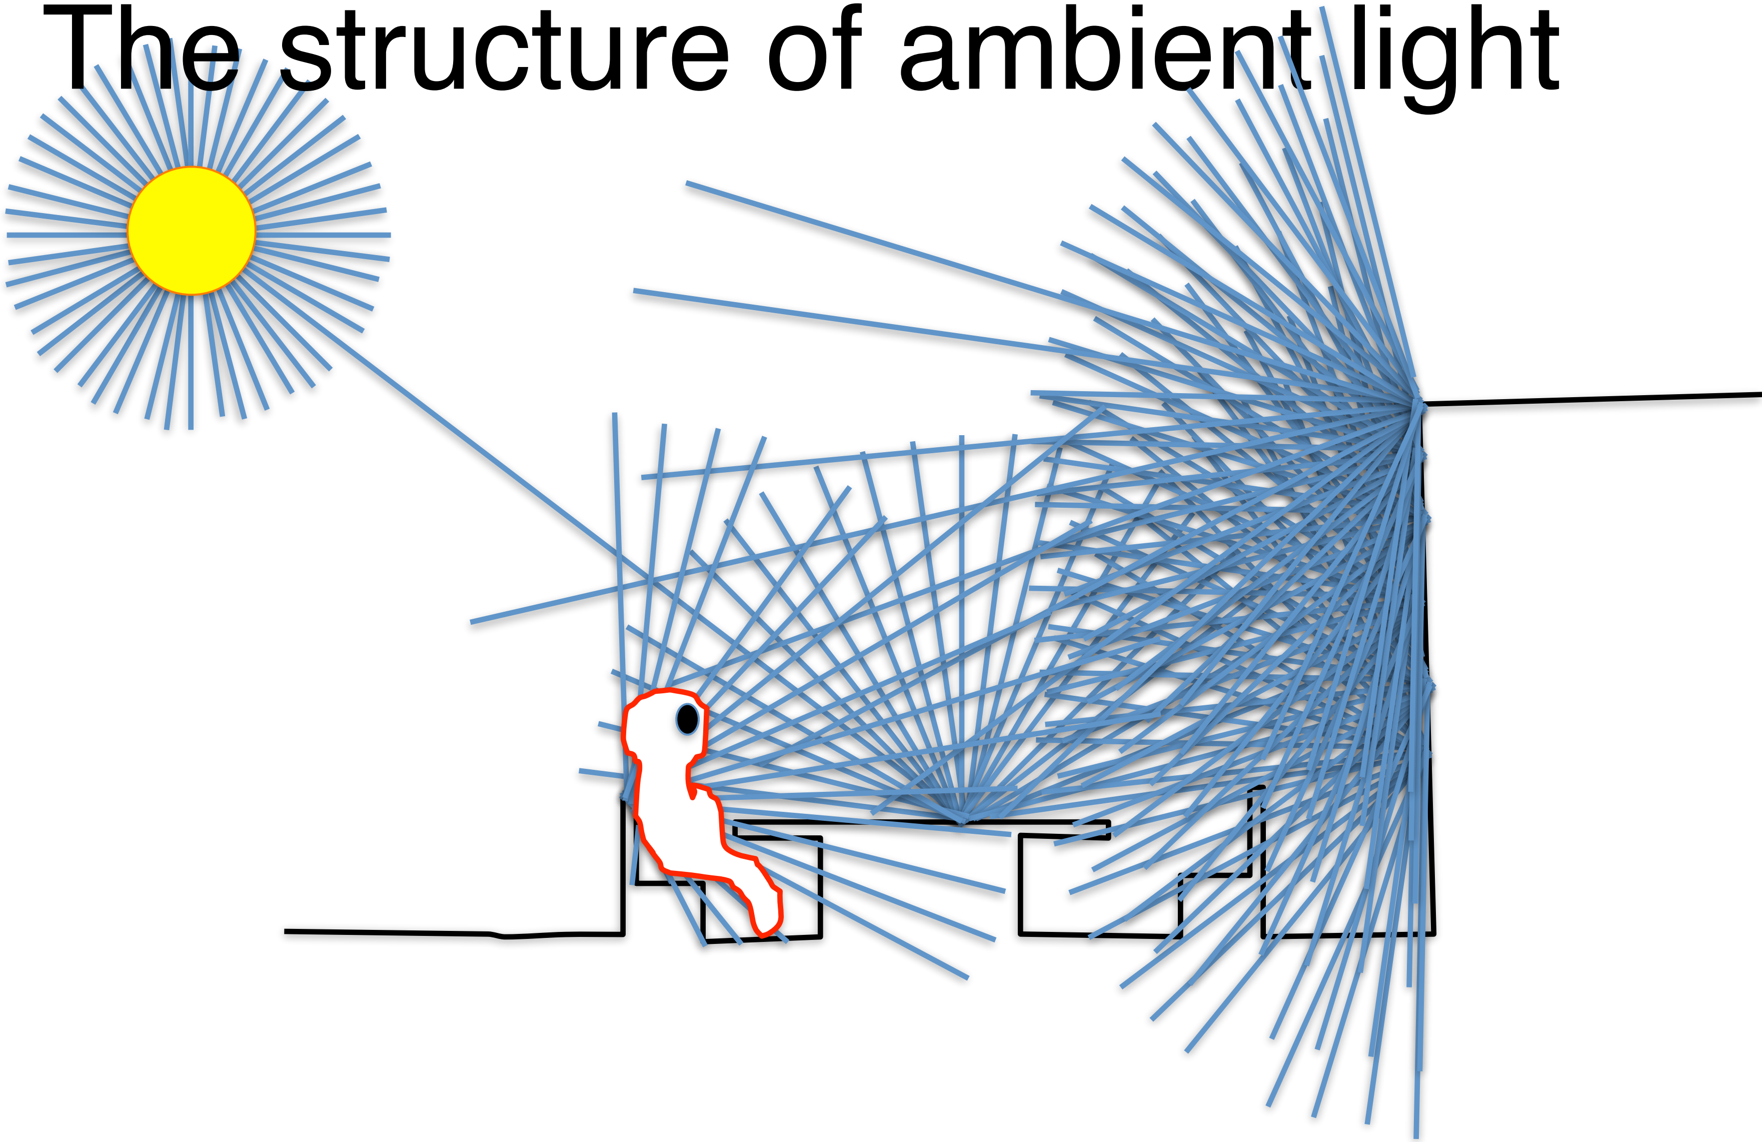

The Plenoptic Function

Adelson & Bergen, 91

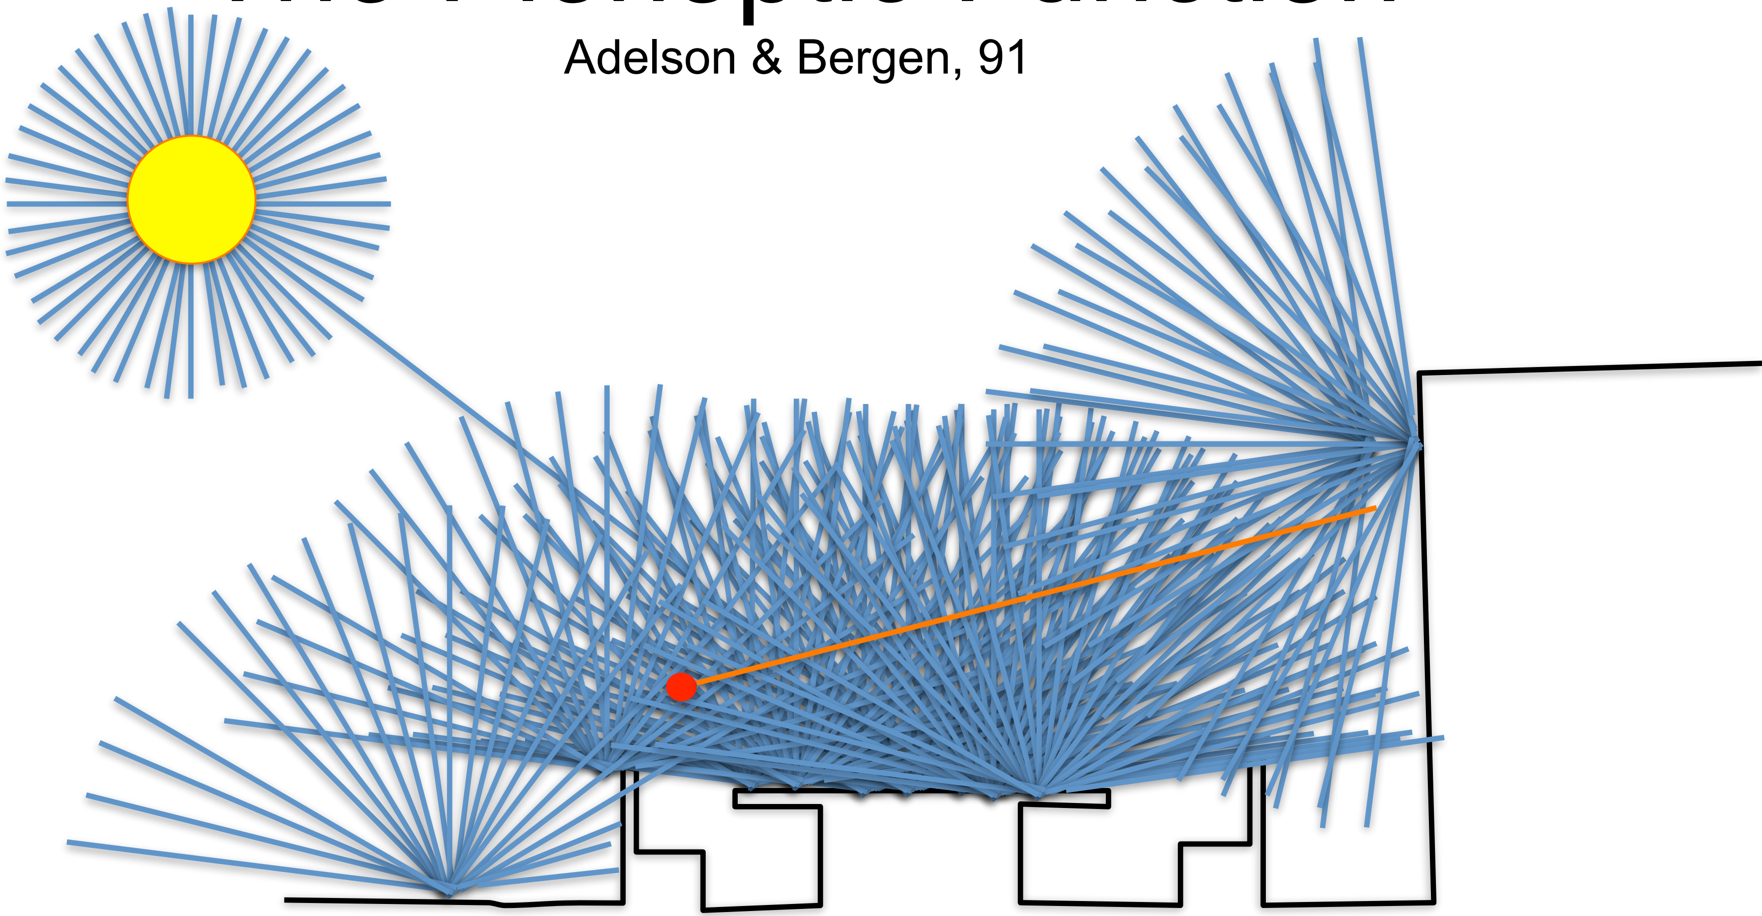

The intensity P can be parameterized as:

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

Measuring the Plenoptic function

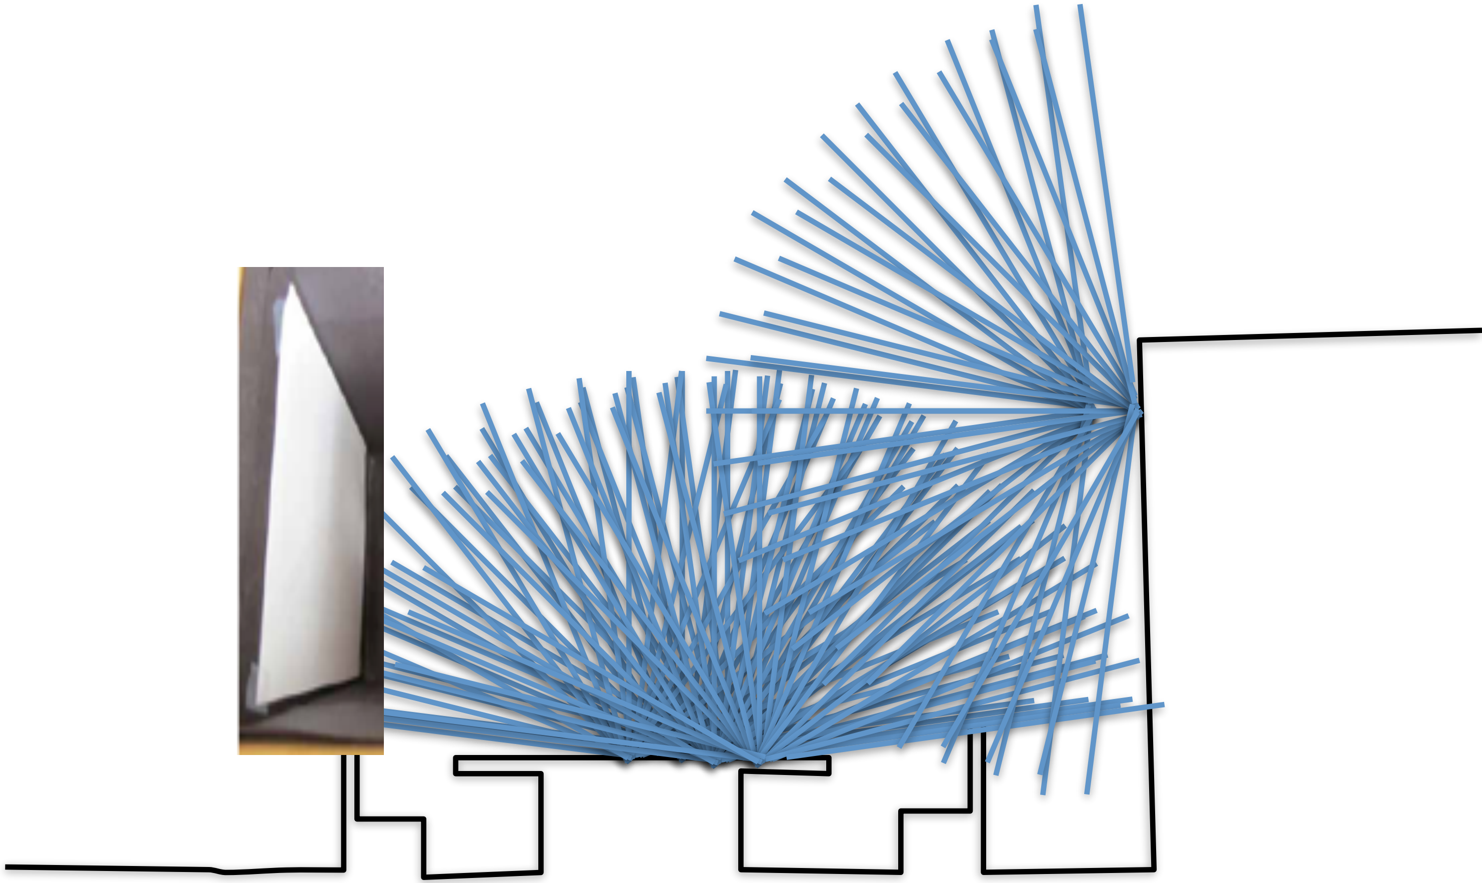

Why is there no picture appearing on the paper?

Light rays from many different parts of the scene strike the same point on the paper.

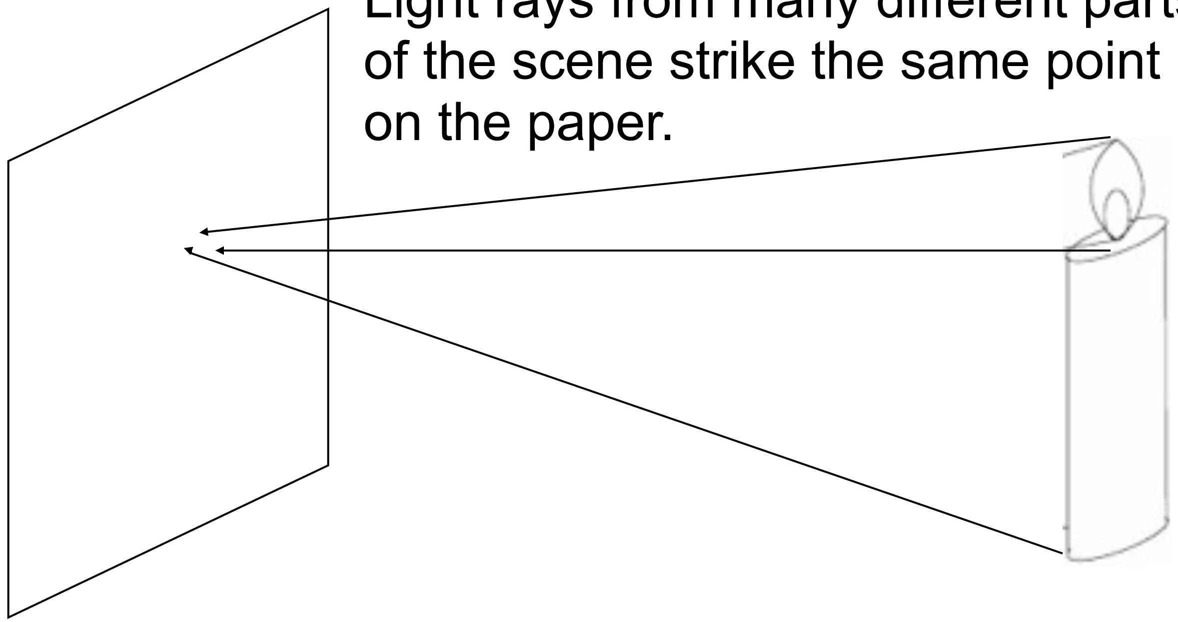

Forsyth & Ponce

Measuring the Plenoptic function

The camera obscura
The pinhole camera

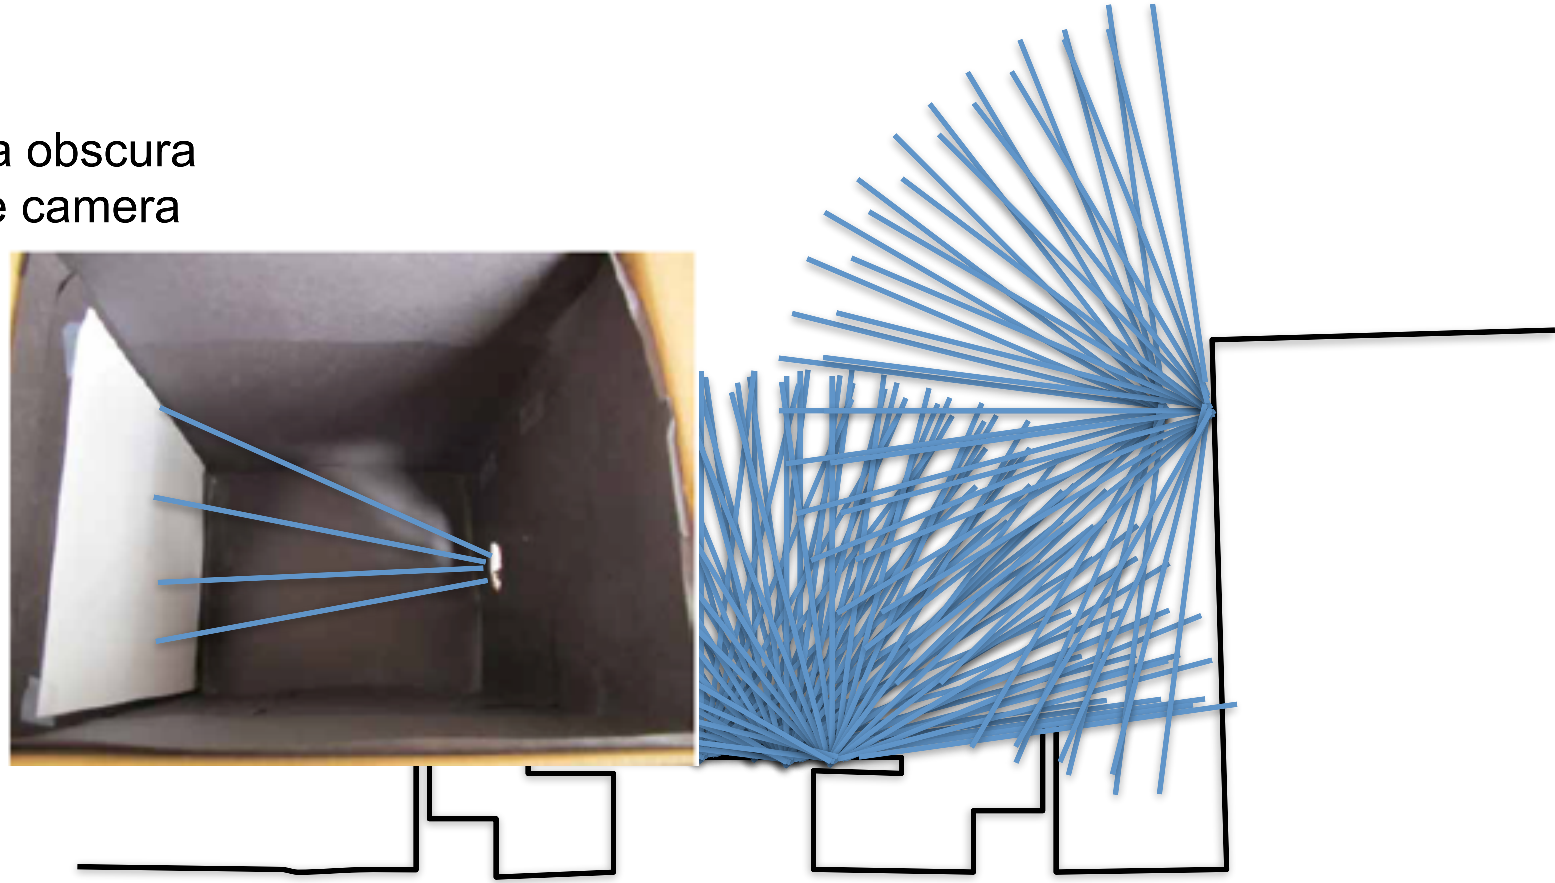

Light rays from many different parts of the scene strike the same point on the paper.

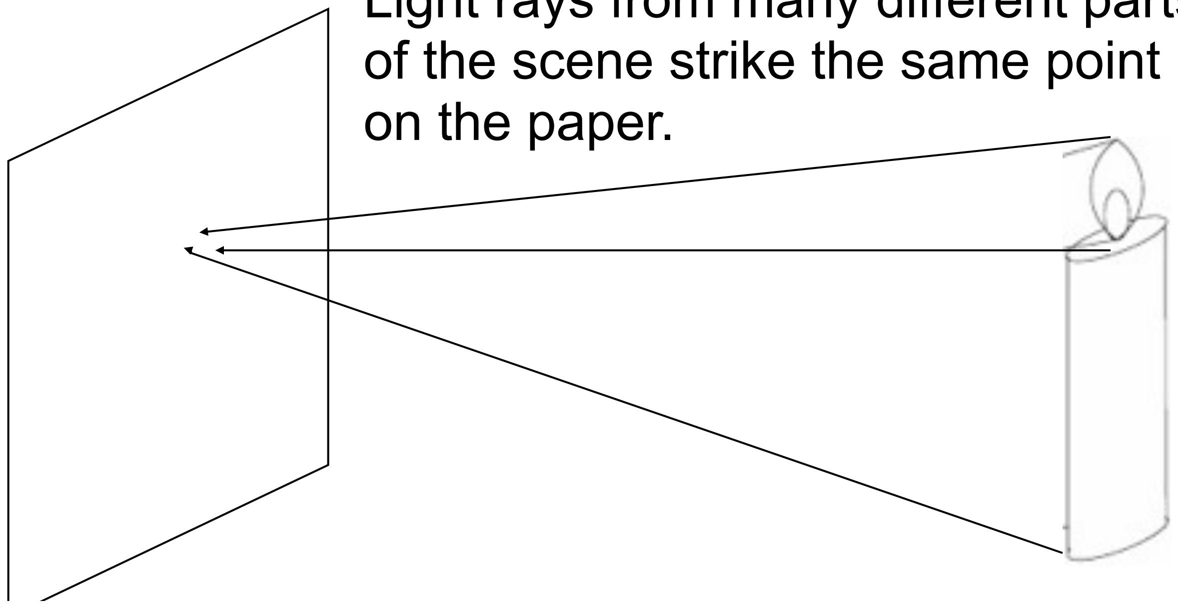

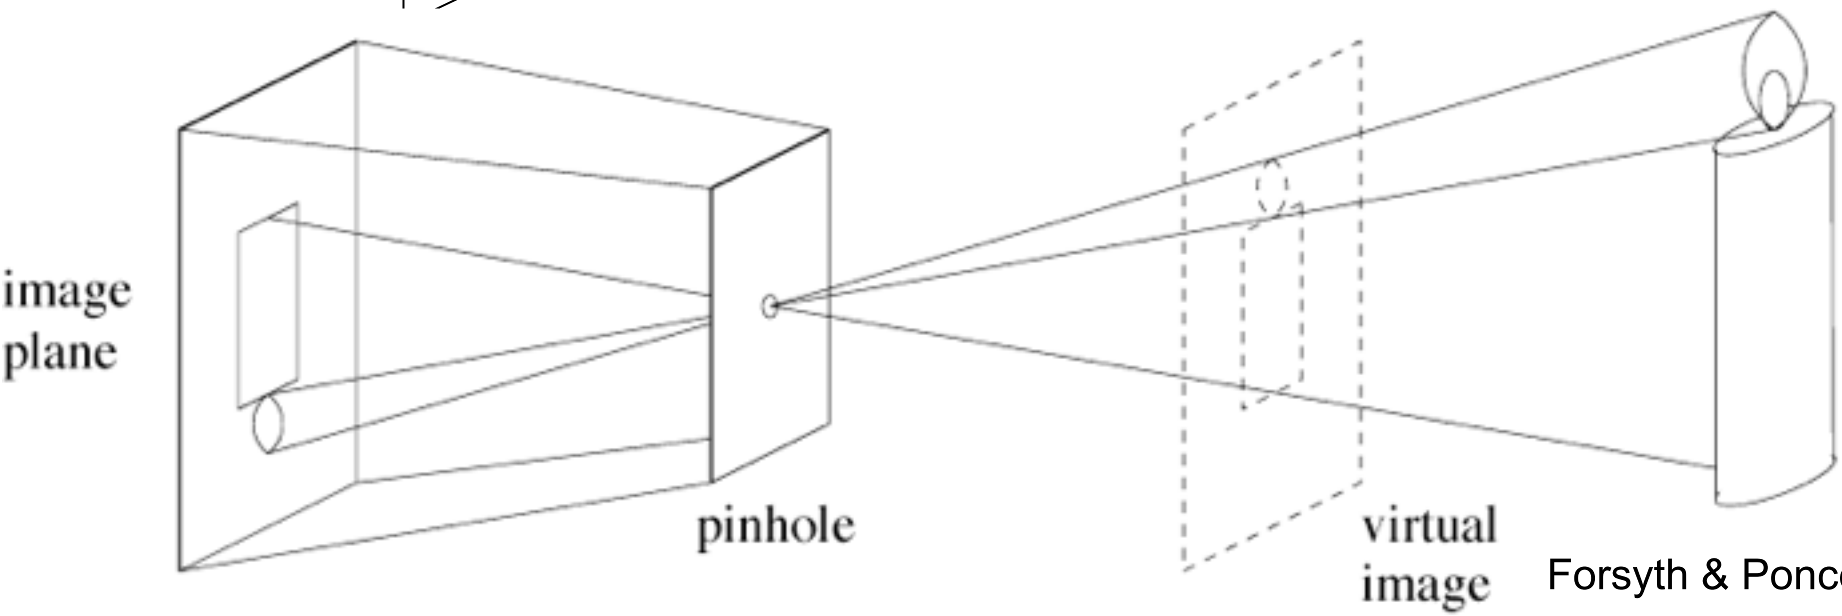

Forsyth & Ponce

The pinhole camera only allows rays from one point in the scene to strike each point of the paper.

Pinhole camera

Photograph by Abelardo Morell, 1991

Pinhole camera

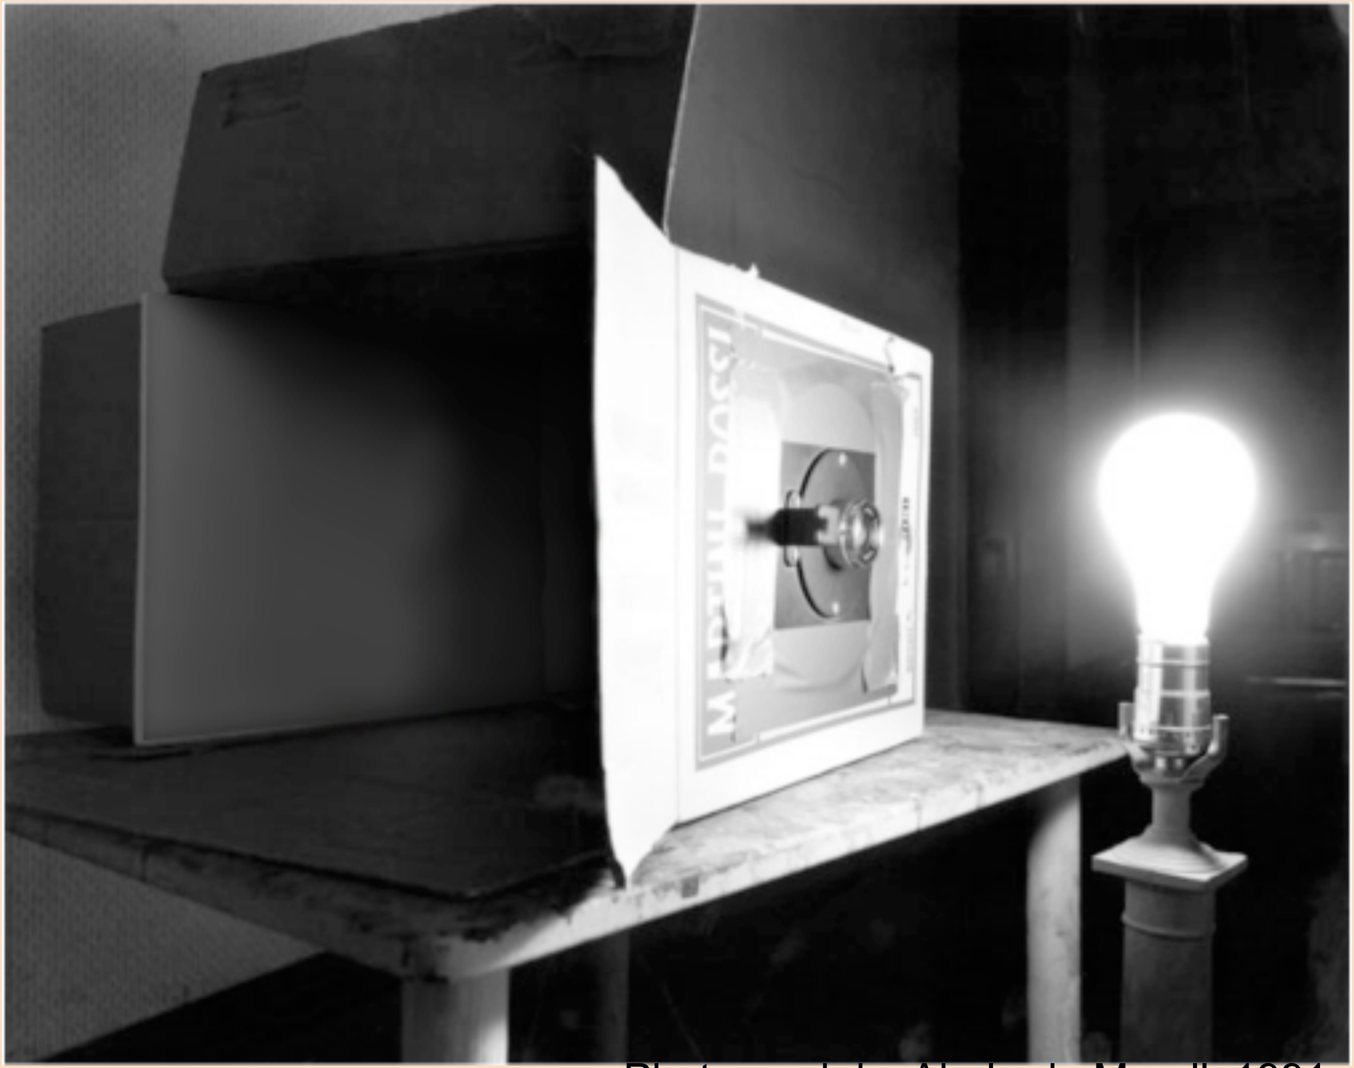

Photograph by Abelardo Morell, 1991

Pinhole camera

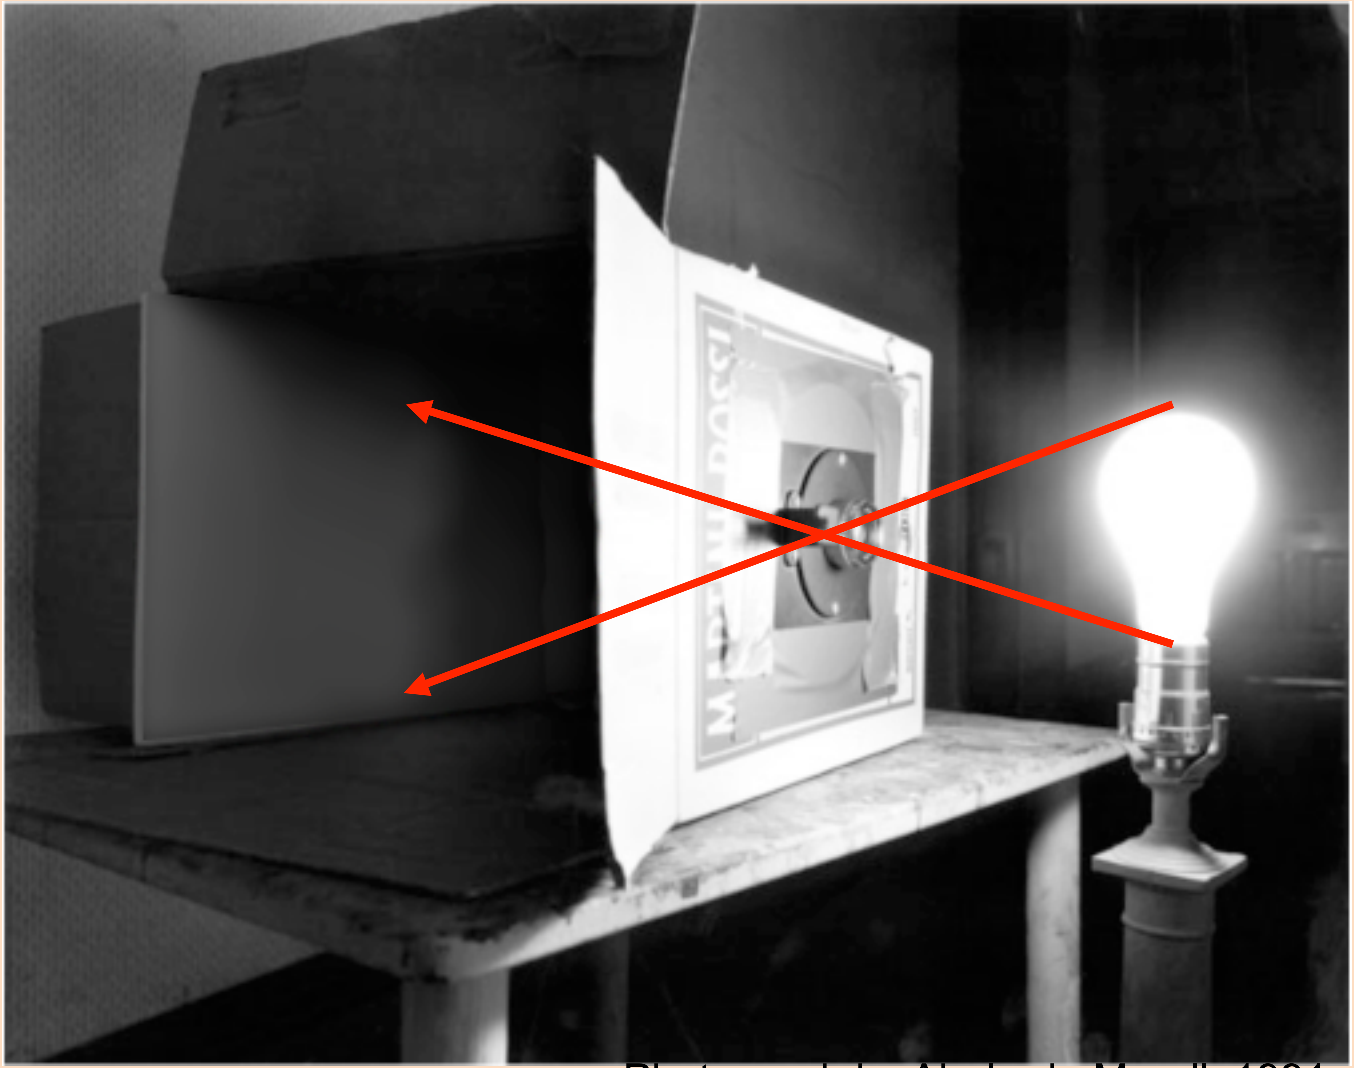

Photograph by Abelardo Morell, 1991

Pinhole camera

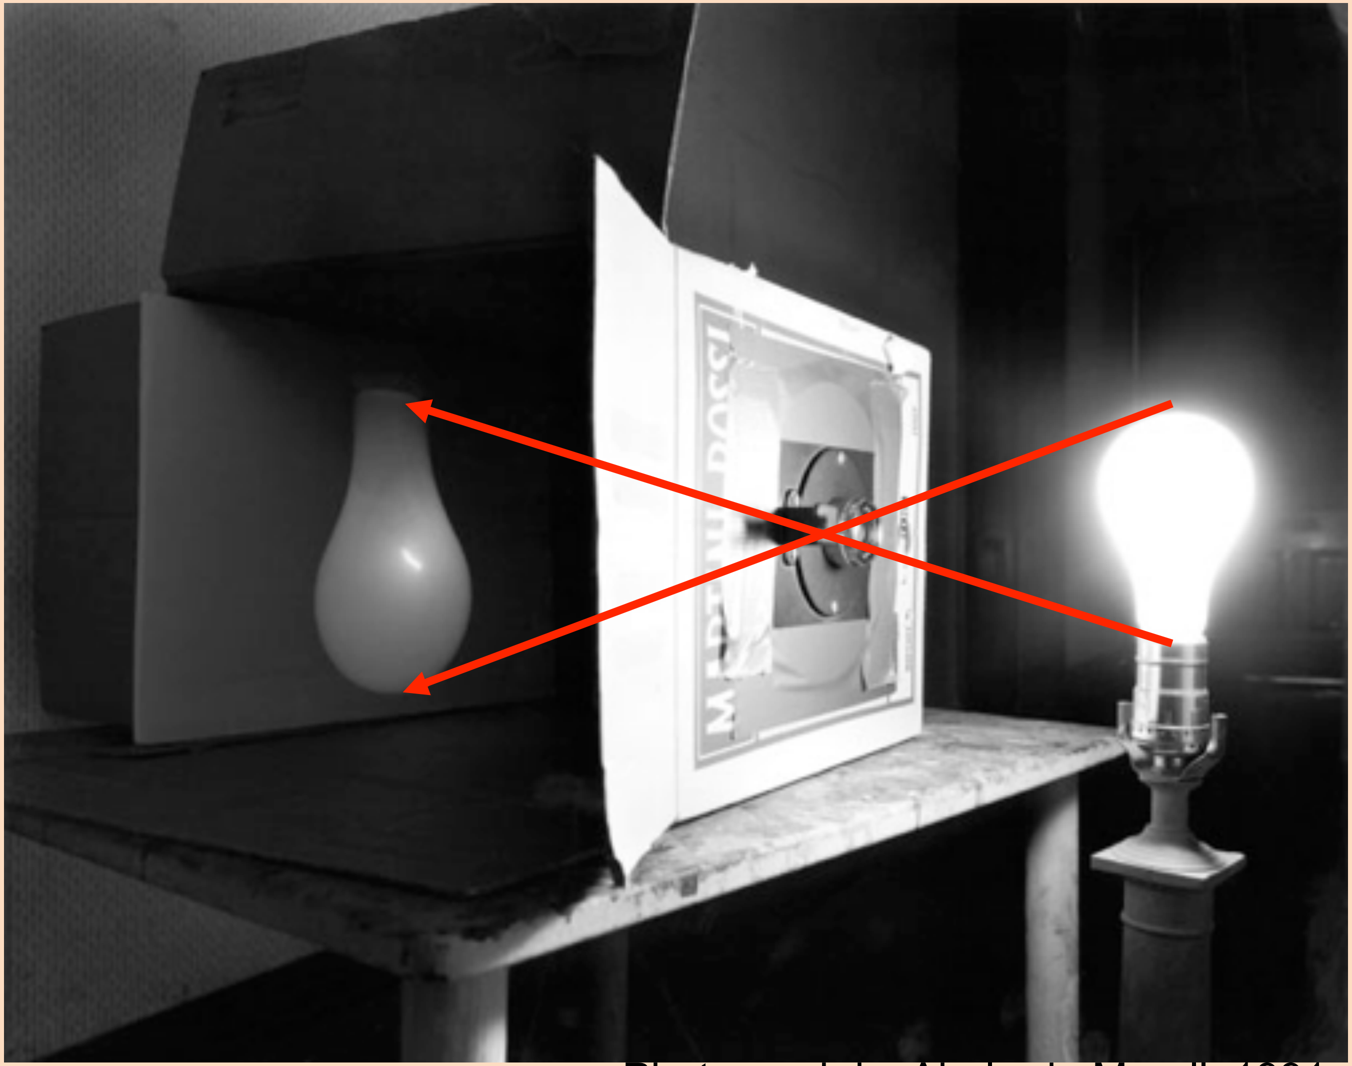

Photograph by Abelardo Morell, 1991

grocery bag pinhole camera

view from outside the bag

http://www.youtube.com/watch?v=FZyCFxsyx8o

view from inside the bag

http://youtu.be/-rhZaAM3F44

Optional Problem Set Problem

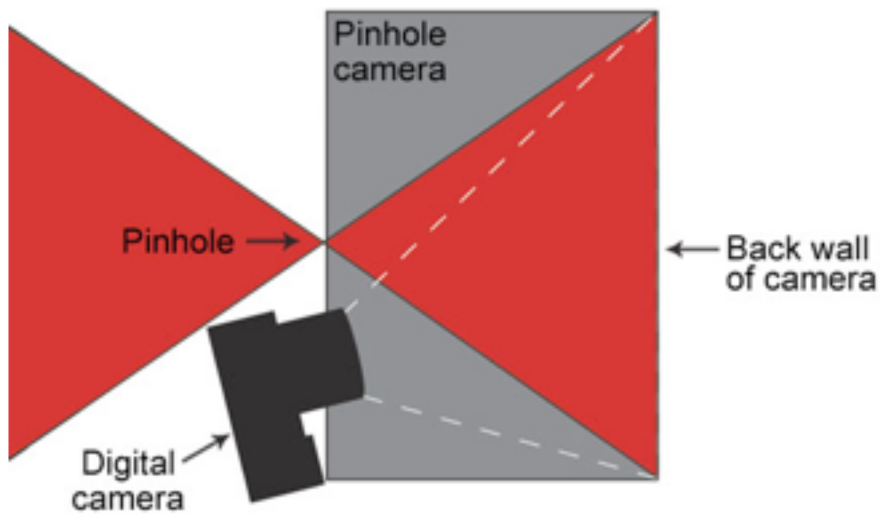

http://www.foundphotography.com/PhotoThoughts/archives/2005/04/pinhole_camera_2.html

Problem Set 1

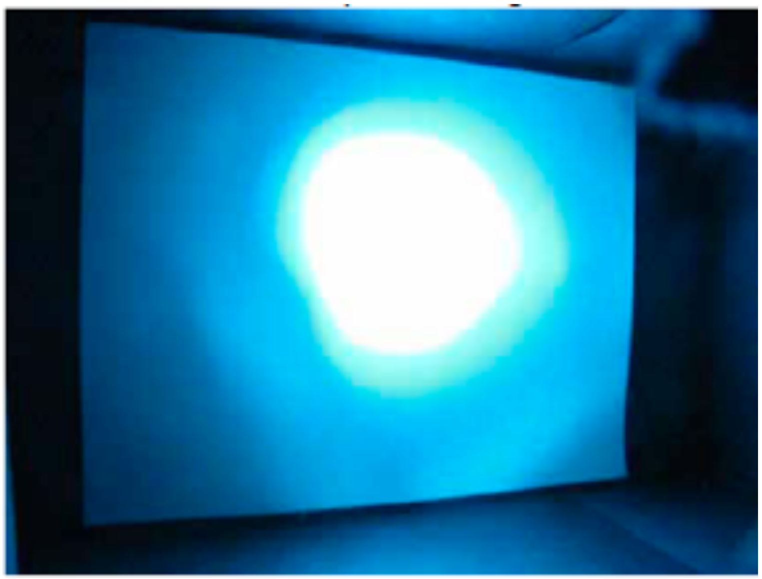

Effect of pinhole size

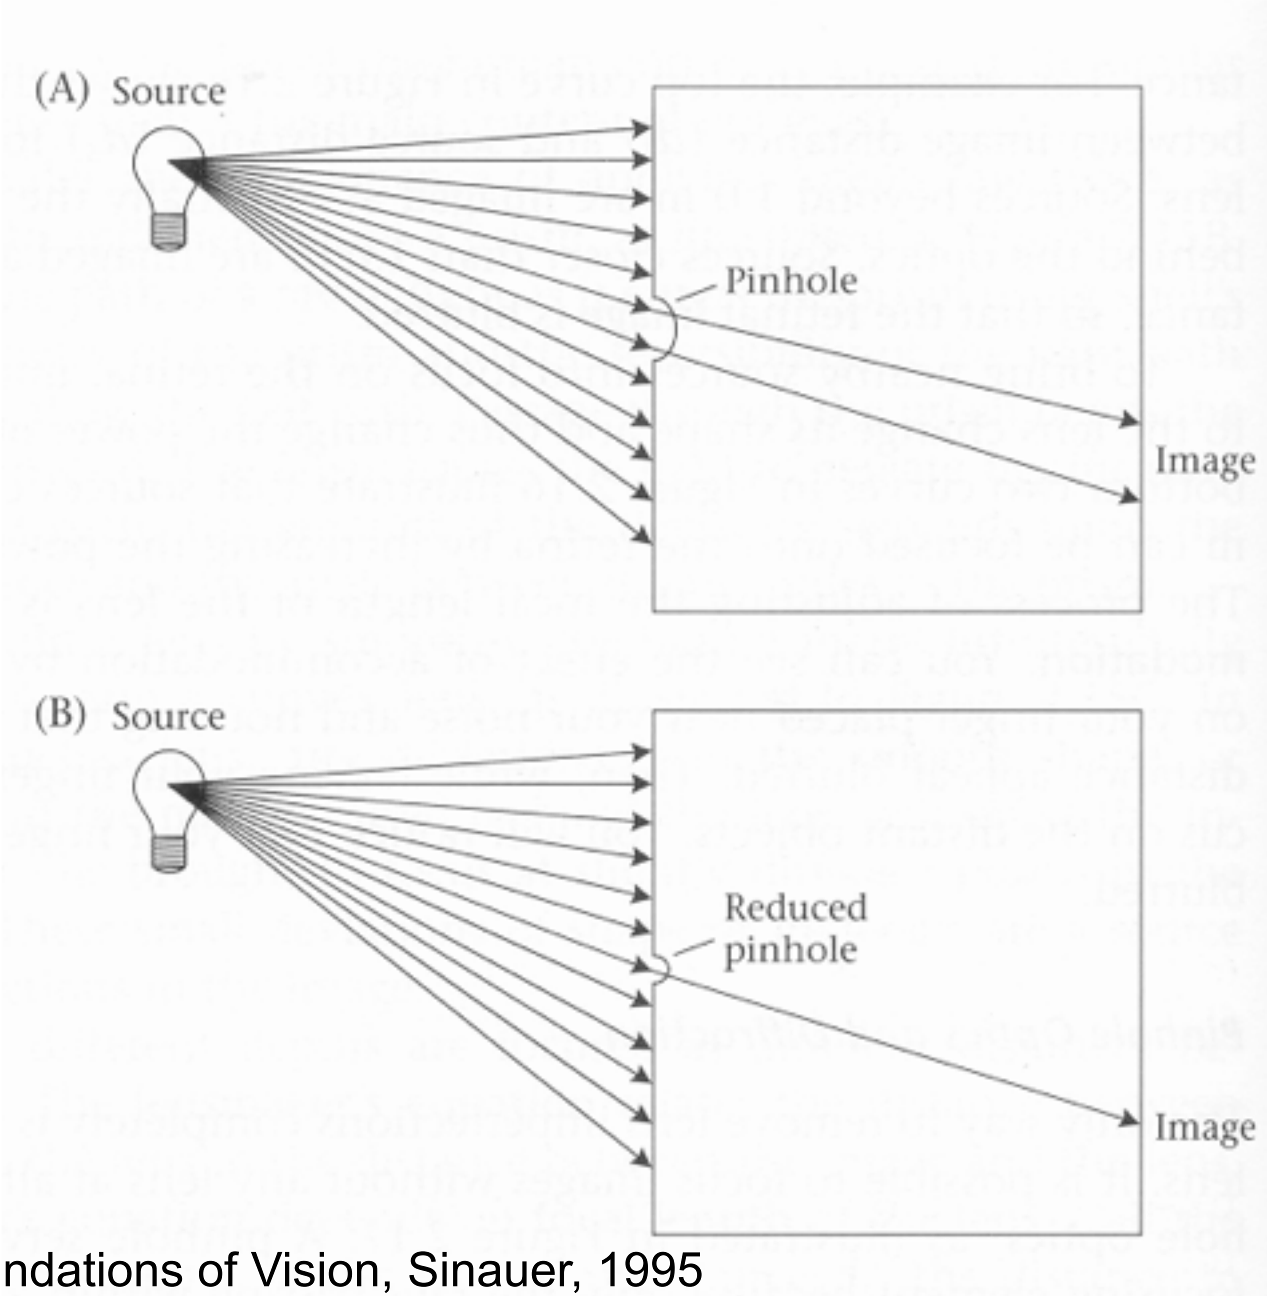

Wandell, Foundations of Vision, Sinauer, 1995

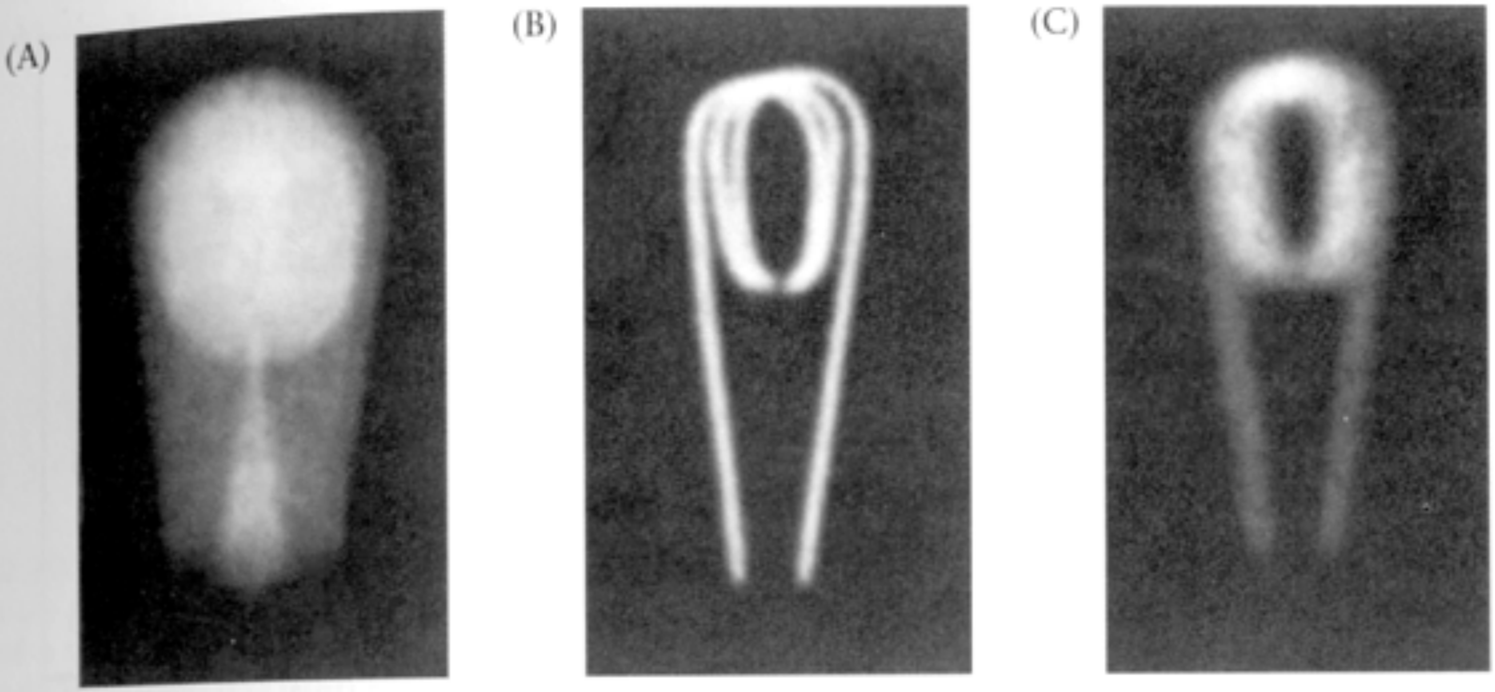

2.18 DIFFRACTION LIMITS THE QUALITY OF PINHOLE OPTICS. These three images of a bulb filament were made using pinholes with decreasing size. (A) When the pinhole is relatively large, the image rays are not properly converged, and the image is blurred. (B) Reducing the size of the pinhole improves the focus. (C) Reducing the size of the pinhole further worsens the focus, due to diffraction. From Ruechardt, 1958.

Wandell, Foundations of Vision, Sinauer, 1995

Measuring distance

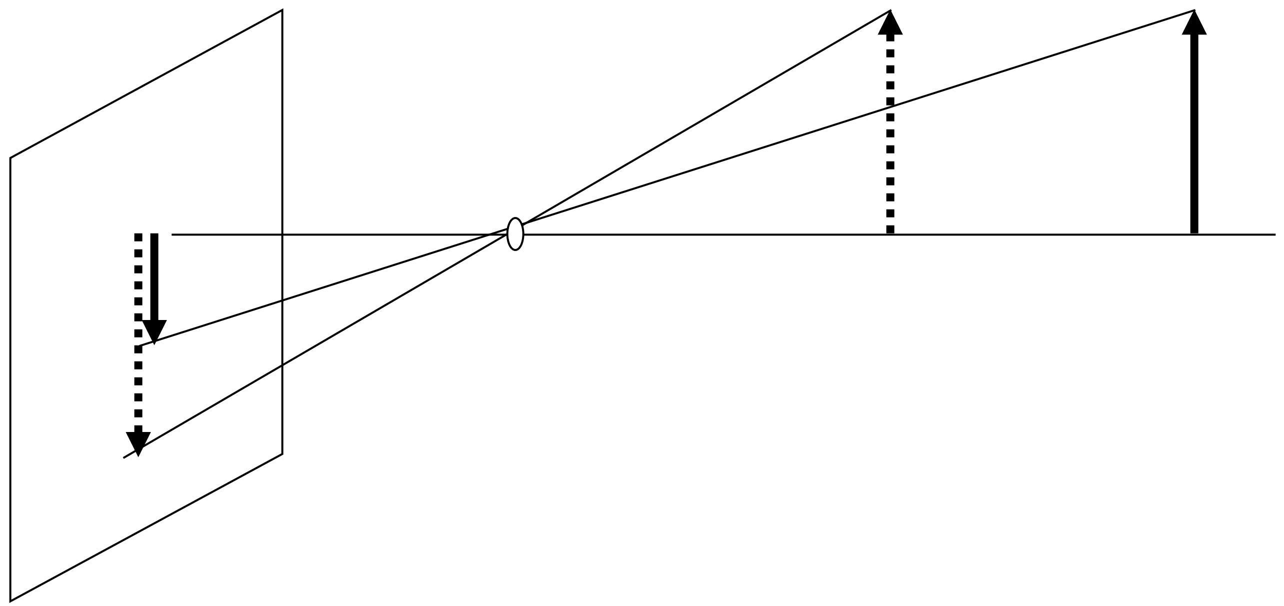

- Object size decreases with distance to the pinhole
- There, given a single projection, if we know the size of the object we can know how far it is.
- But for objects of unknown size, the 3D information seems to be lost.

Playing with pinholes

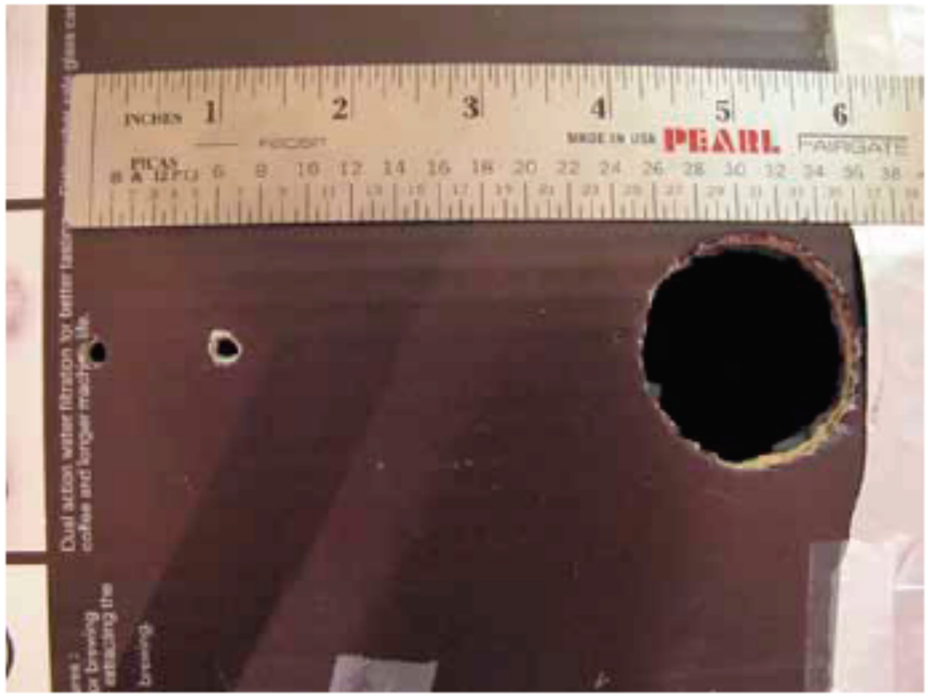

Two pinholes

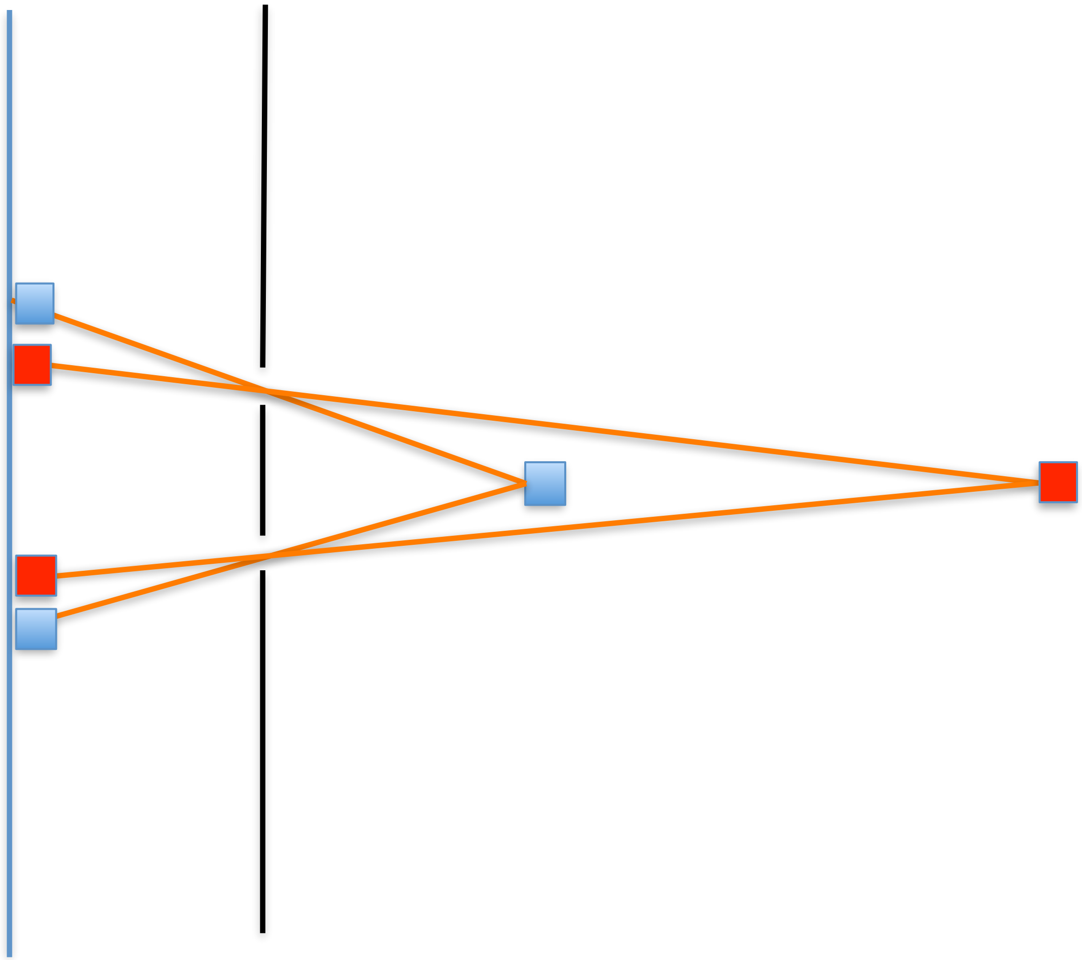

Two pinholes

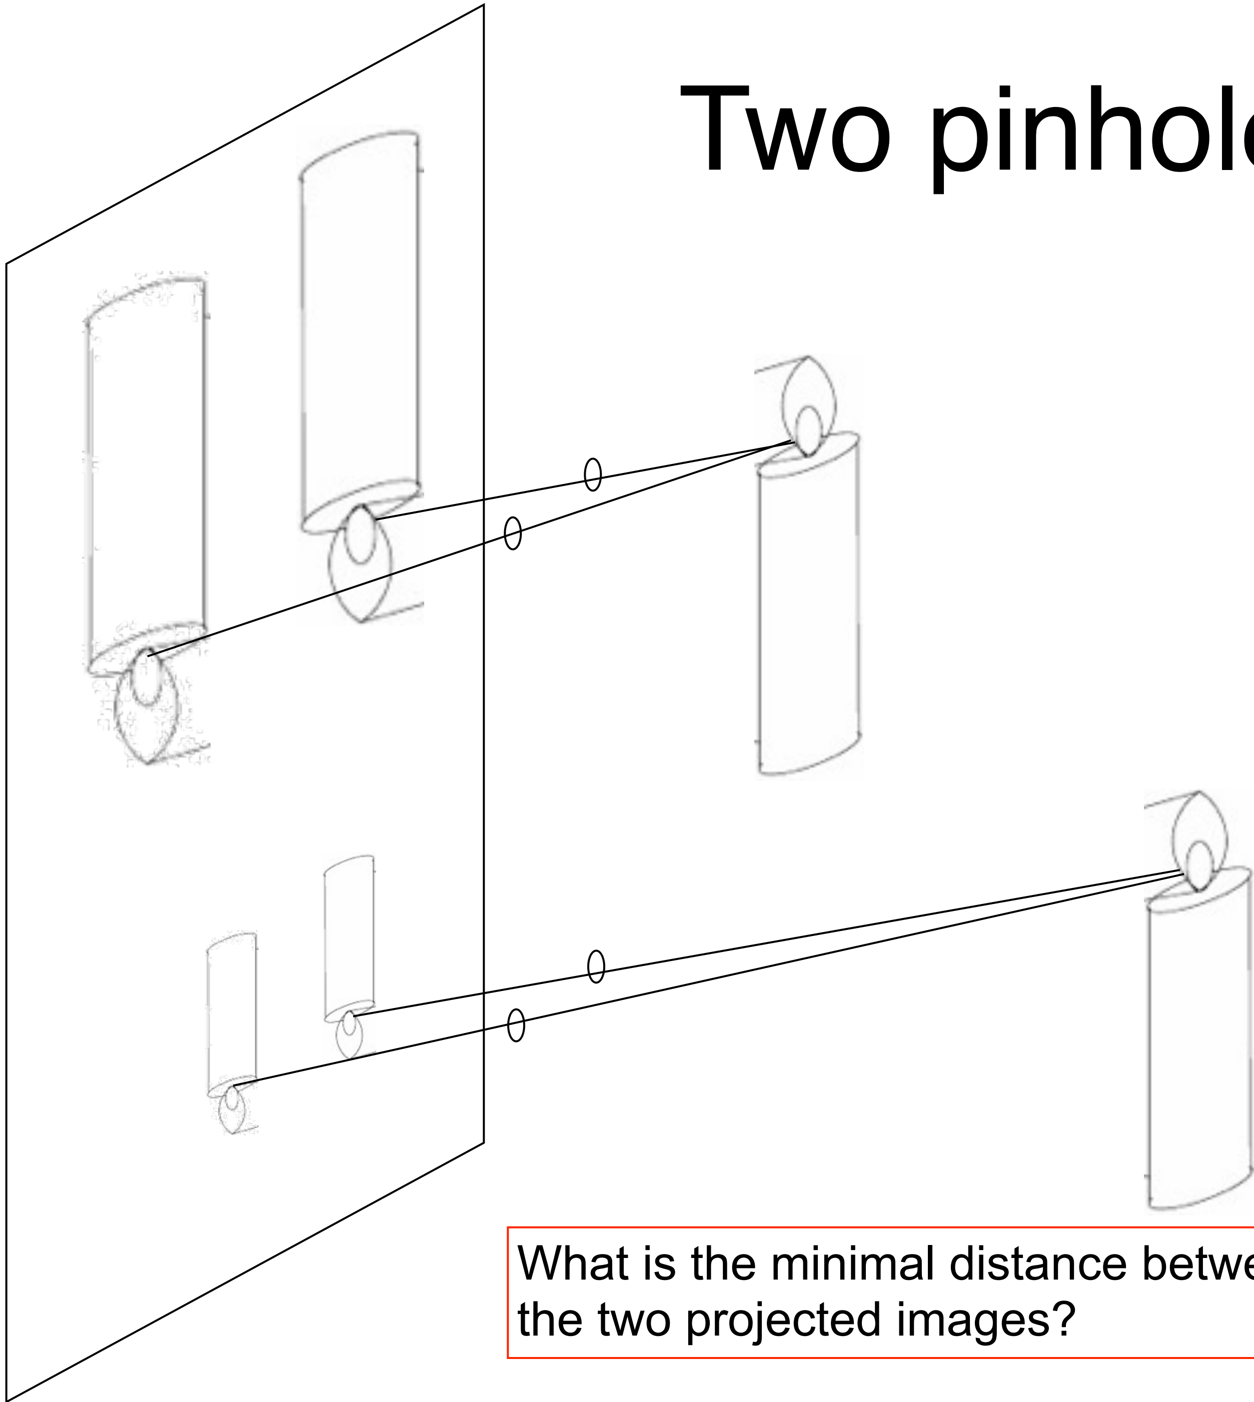

What is the minimal distance between the two projected images?

Anaglyph pinhole camera

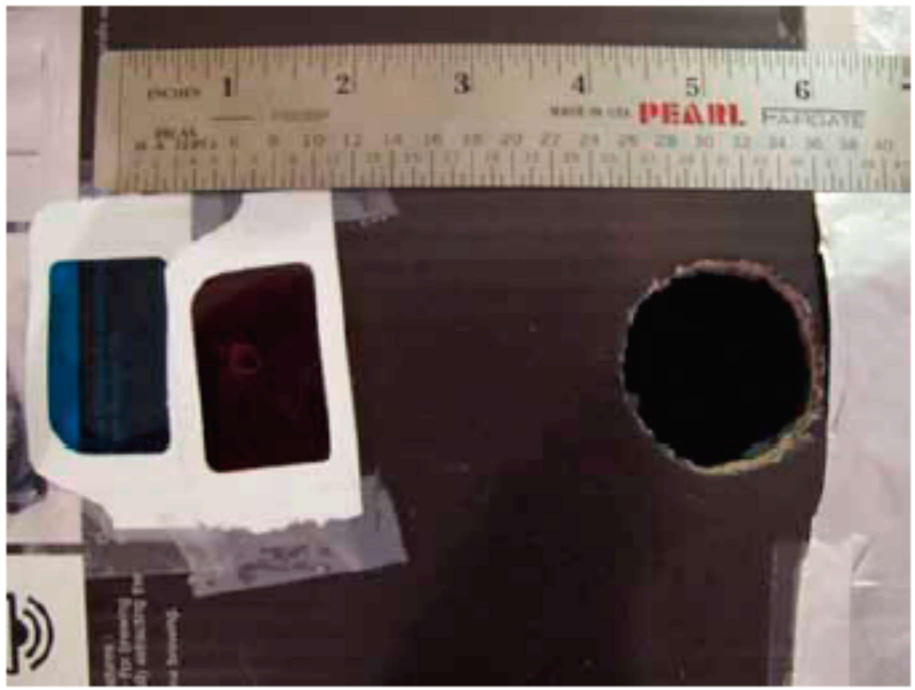

Anaglyph pinhole camera

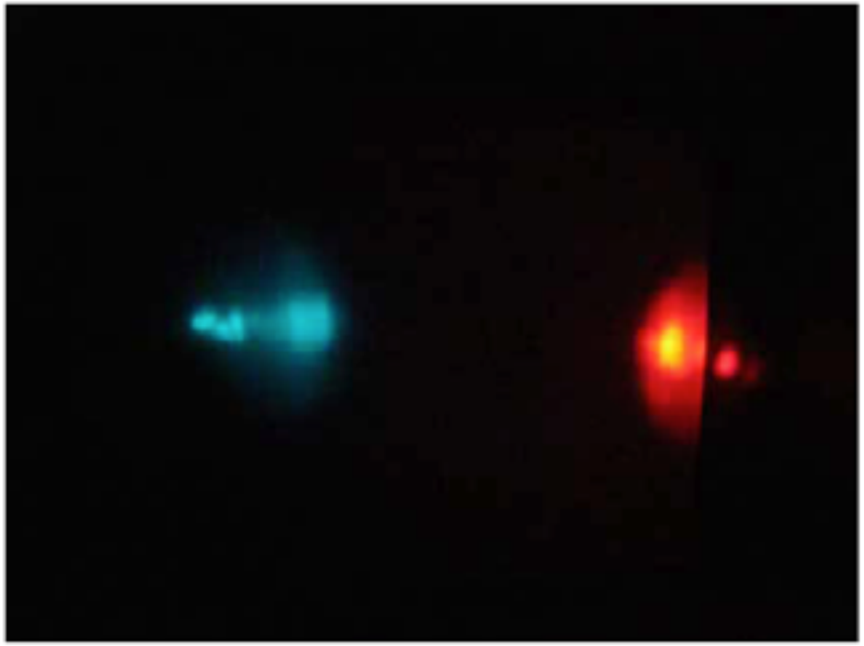

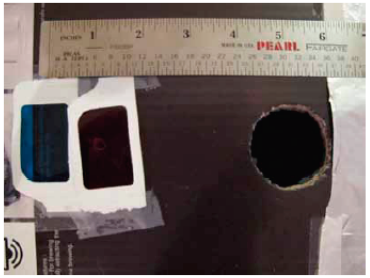

Anaglyph pinhole camera

Anaglyph

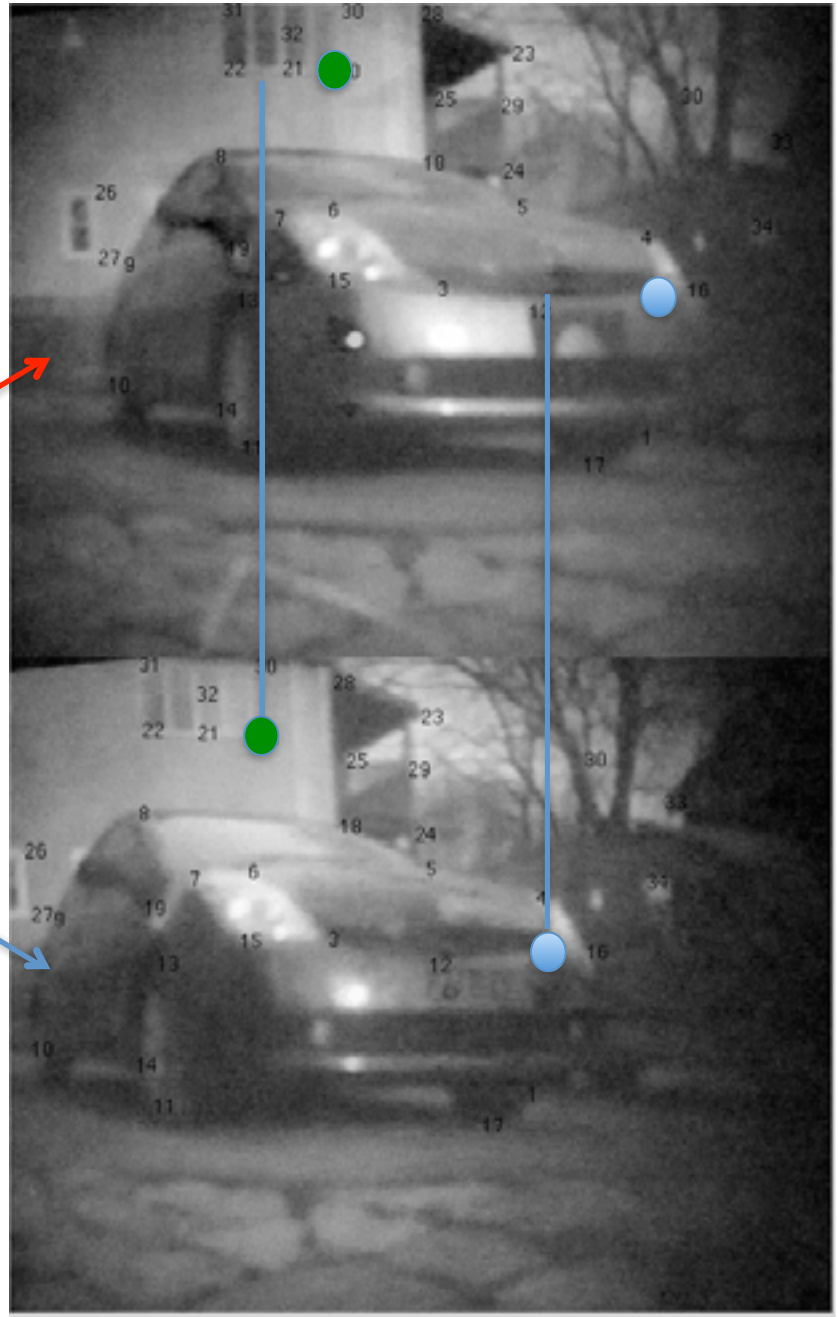

Synthesis of new views

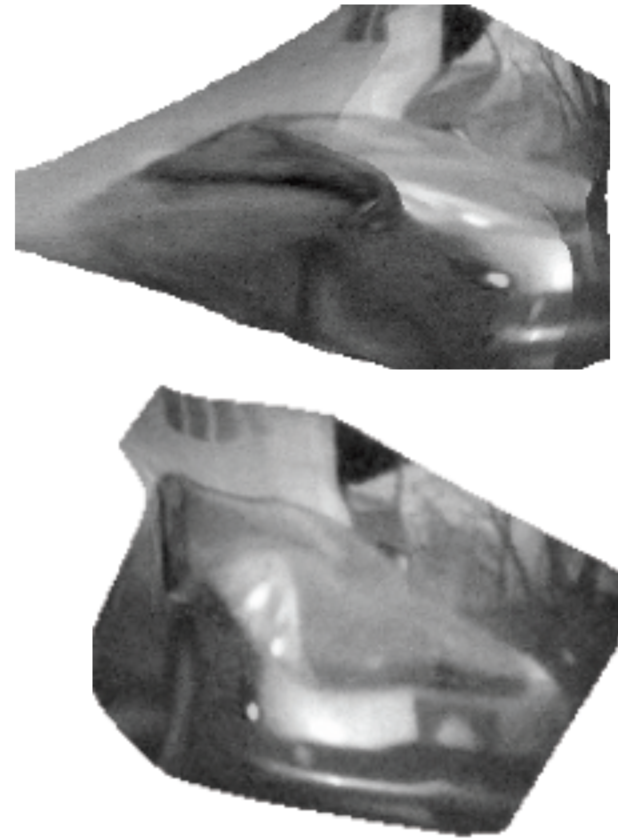

Problem set

- Build the device
- Take some pictures and put them in the report
- Take anaglyph images
- Work out the geometry
- Recover depth for some points in the image

Accidental pinhole and pinspeck cameras: Revealing the scene outside the picture

Antonio Torralba
William T. Freeman

See project page for videos:
http://people.csail.mit.edu/torralba/research/accidentalcameras/

CSAIL

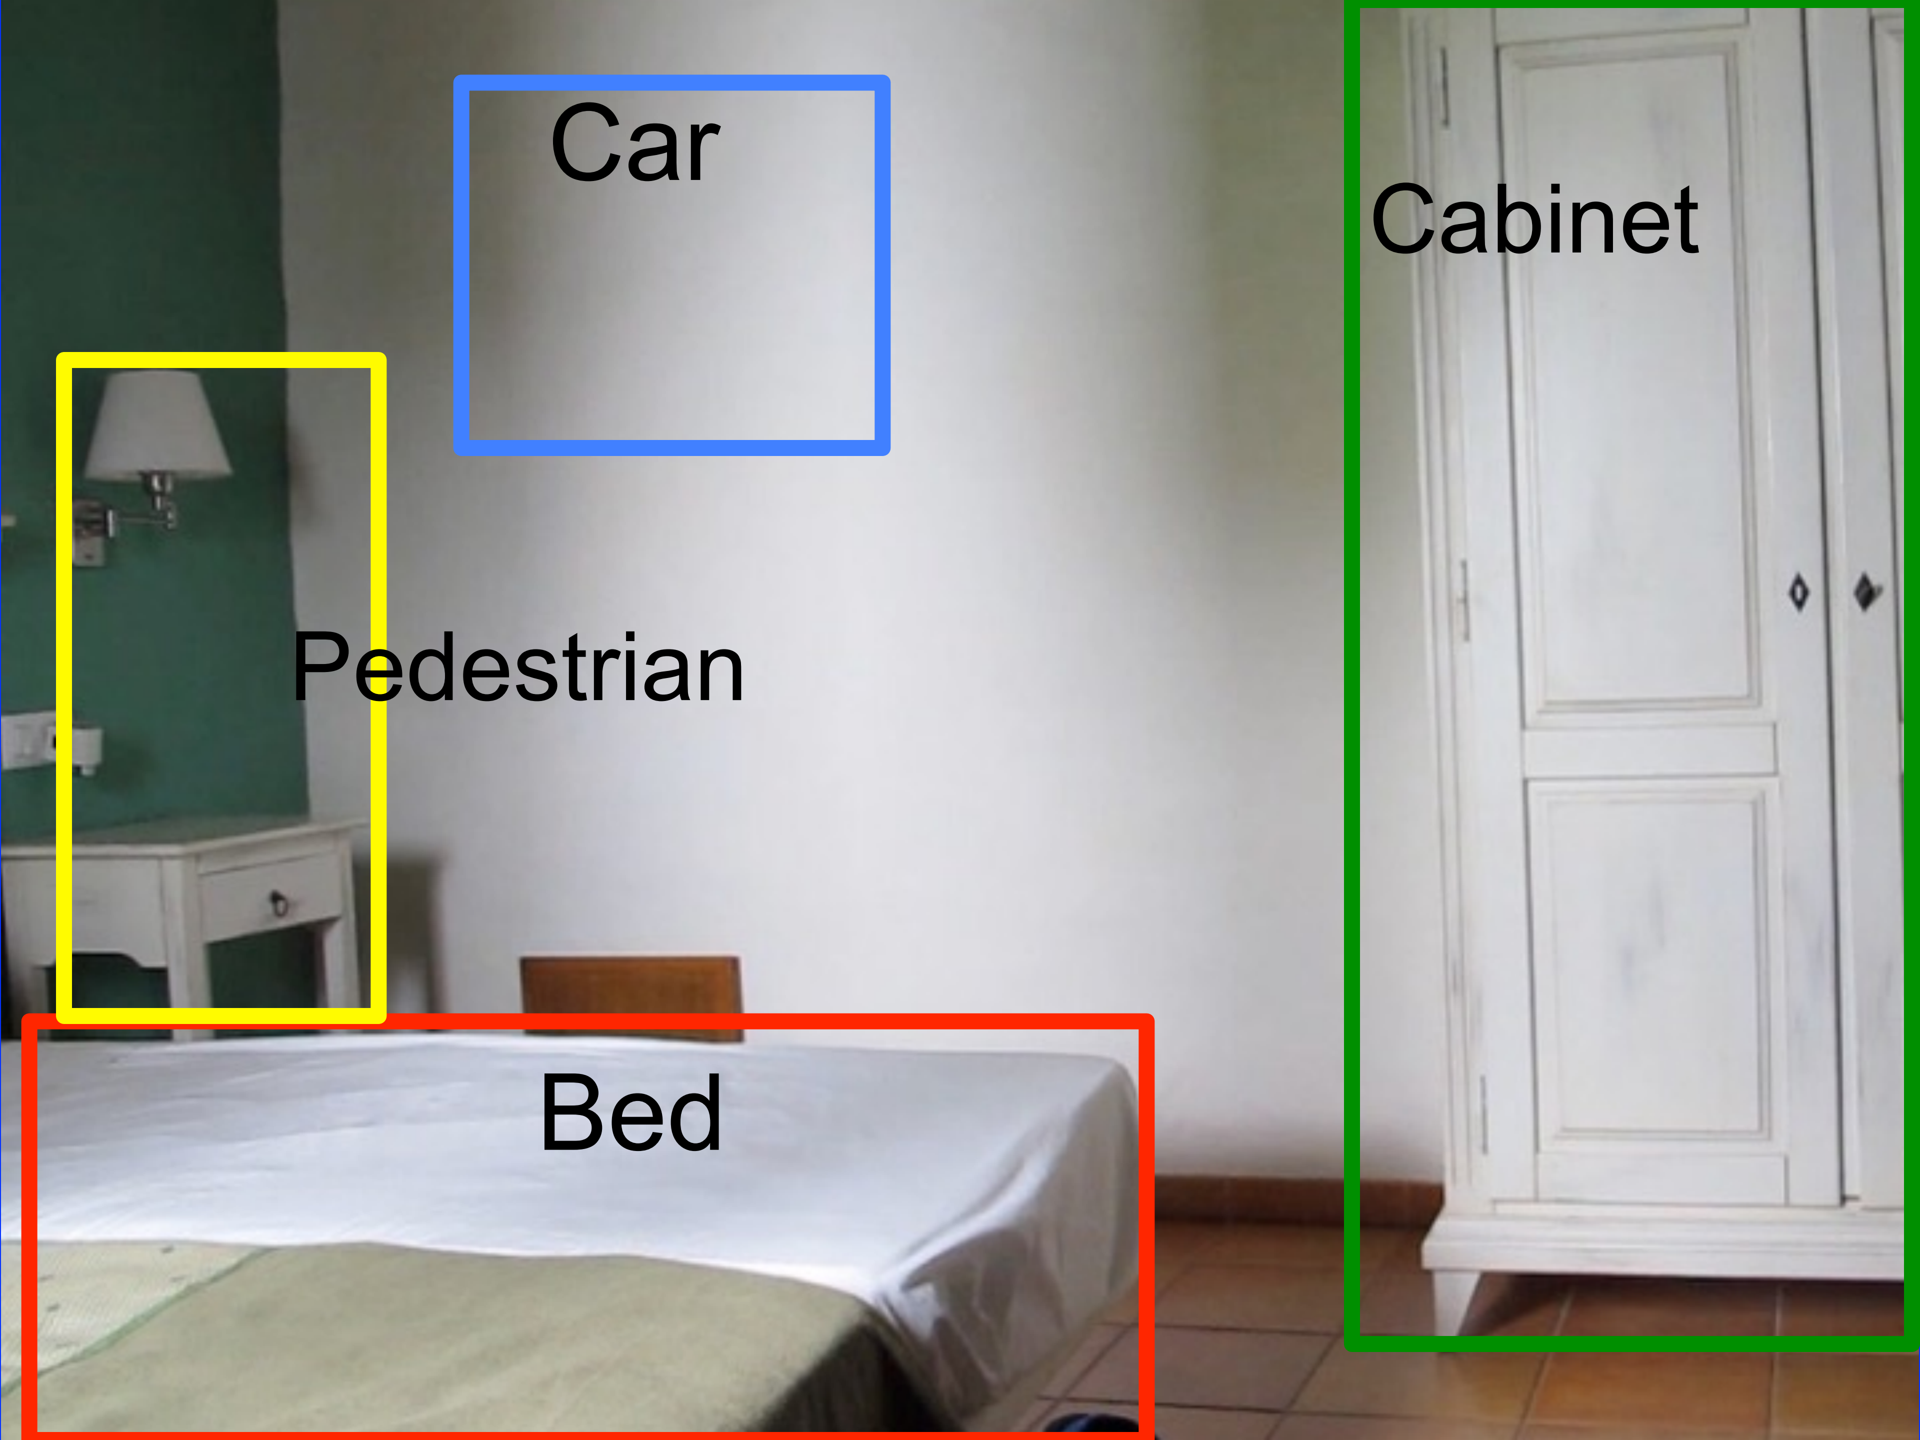
Car
Cabinet
Pedestrian
Bed


Shadows?

Accidental pinhole camera

Window turned into a pinhole

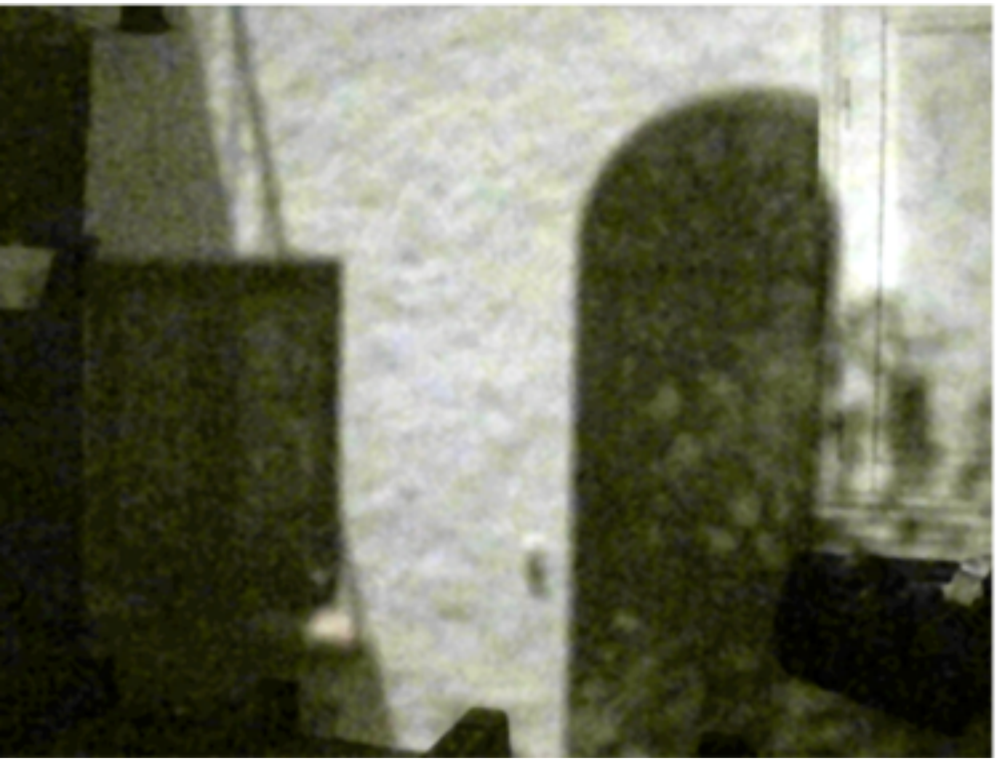

View outside

Accidental pinhole camera

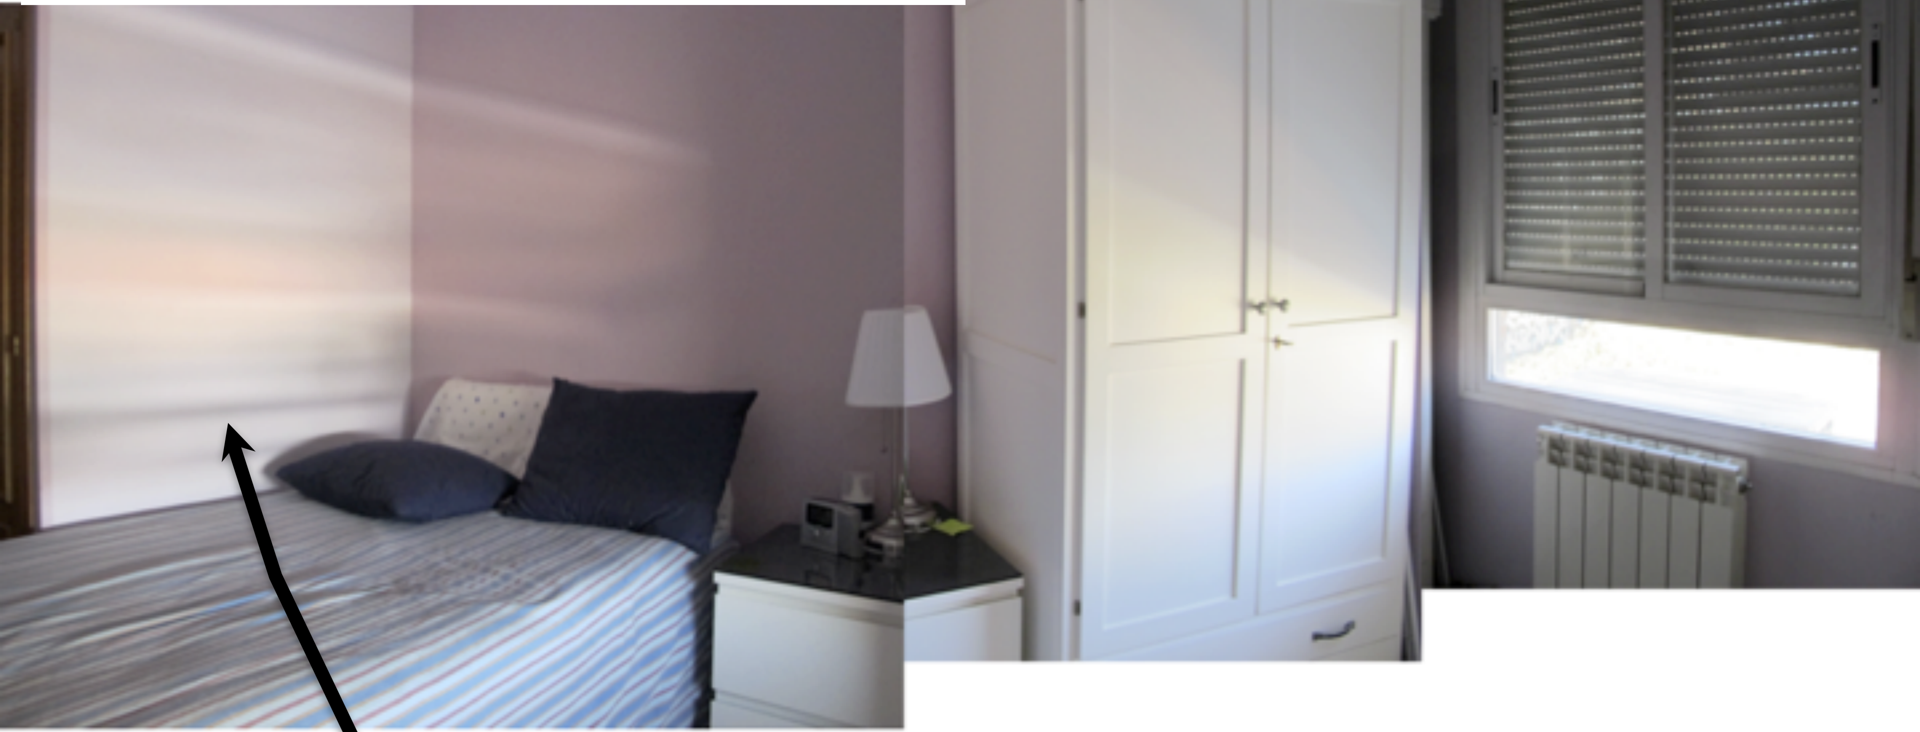

Outside scene

*

Aperture

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

Anti-pinhole or Pinspeck cameras

Adam L. Cohen, 1982

OPTICA ACTA, 1982, VOL. 29, NO. 1, 63–67

Anti-pinhole imaging

ADAM LLOYD COHEN
Parmly Research Institute, Loyola University of Chicago, Chicago, Illinois 60626, U.S.A.

(*Received 16 April 1981; revision received 8 July 1981*)

Abstract. By complementing a pinhole to produce an isolated opaque spot, the light ordinarily blocked from the pinhole image is transmitted, and the light ordinarily transmitted is blocked. A negative geometrical image is formed, distinct from the familiar 'bright-spot' diffraction image. Anti-pinhole, or 'pinspeck' images are visible during a solar eclipse, when the shadows of objects appear crescent-shaped. Pinspecks demonstrate unlimited depth of field, freedom from distortion and large angular field. Images of different magnification may be formed simultaneously. Contrast is poor, but is improvable by averaging to remove noise and subtraction of a d.c. bias. Pinspecks may have application in X-ray space optics, and might be employed in the eyes of simple organisms.

Pinhole and Anti-pinhole cameras

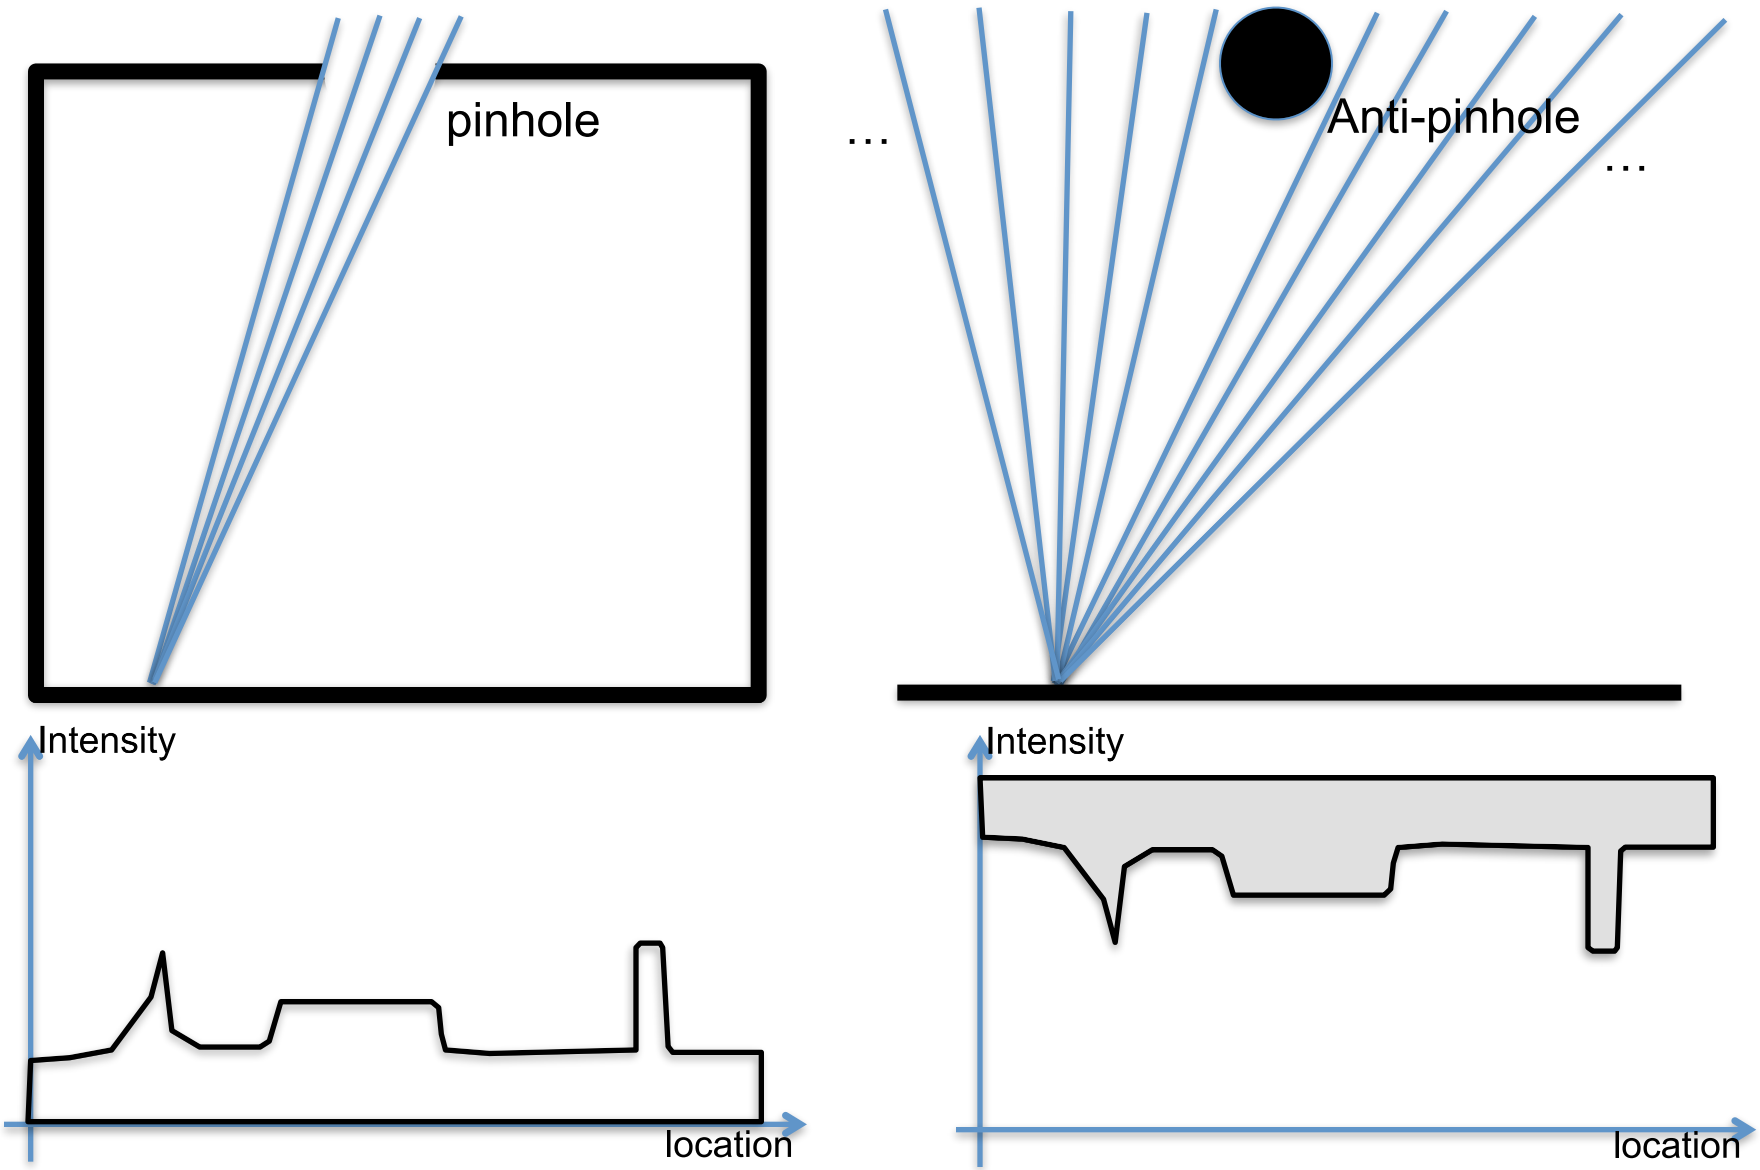

Adam L. Cohen, 1982

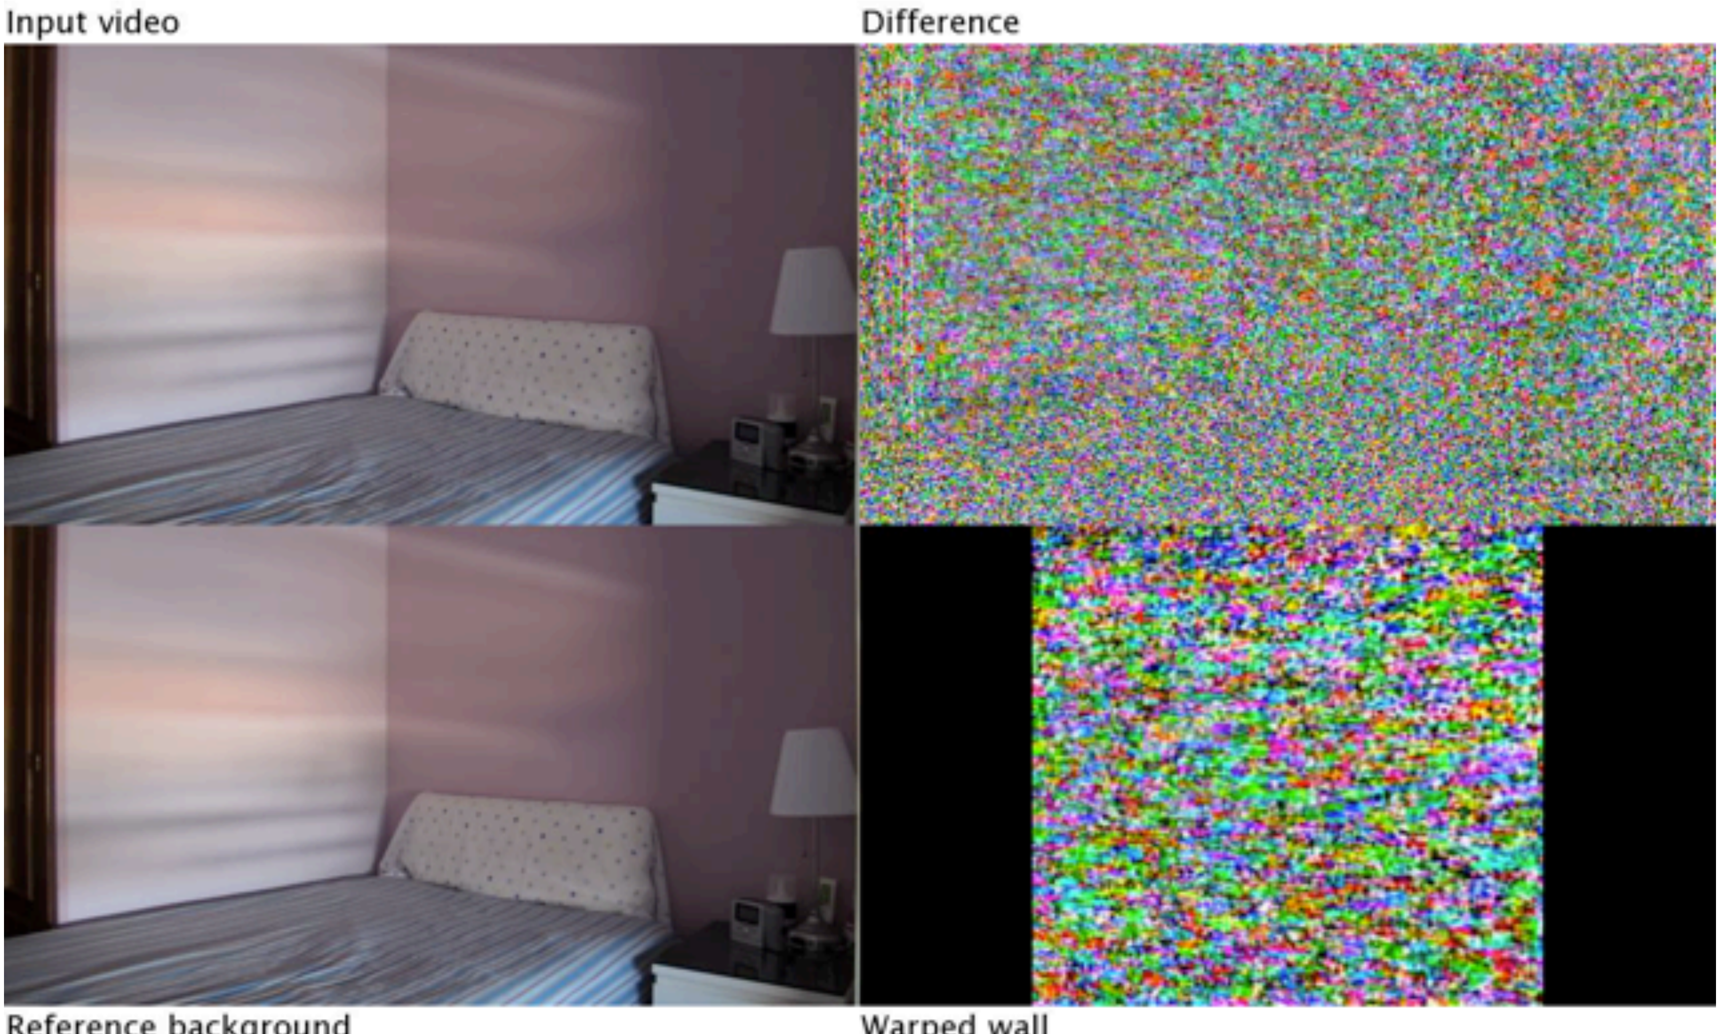
Input video
Difference
Reference background
Warped wall

Shadows
Accidental anti-pinhole cameras

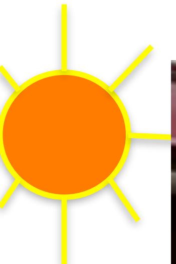

Background image

-

Input video

=

Negative of the shadow

Background image

-

Input video

= Negative of the shadow

Input video

Negative of the shadow

Inverted difference image

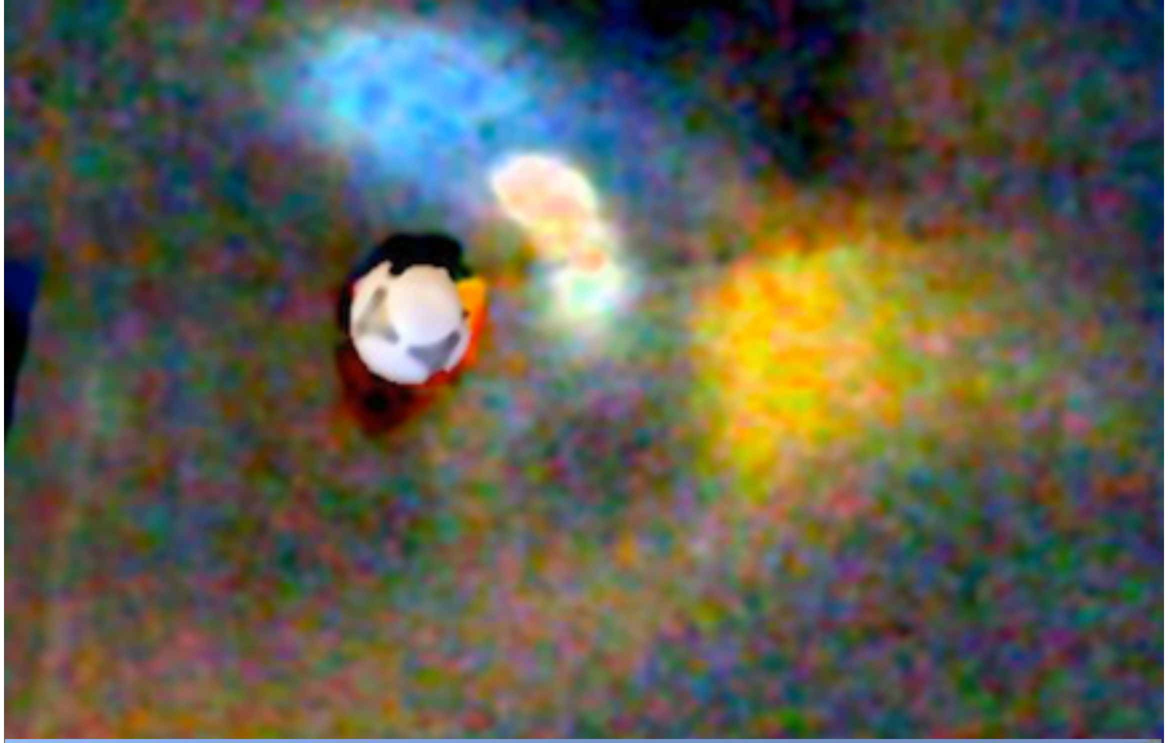

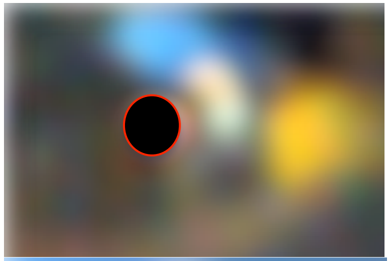

View behind the ball

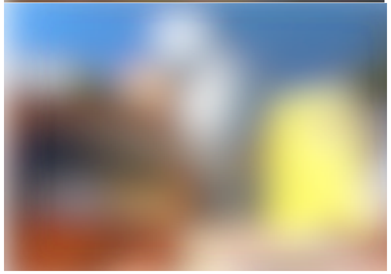

The importance of the size of the occluder

Input video

Negative of the shadow

Size of the occluder

Antonio

Ball

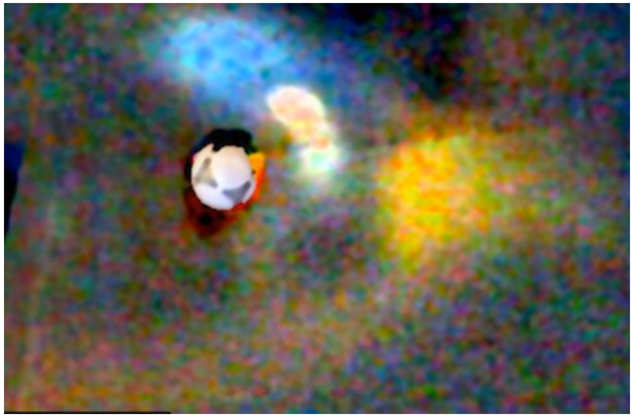

Input video

Negative of the shadow

Using some single view metrology. A. Criminisi, I. Reid, and A. Zisserman 1999

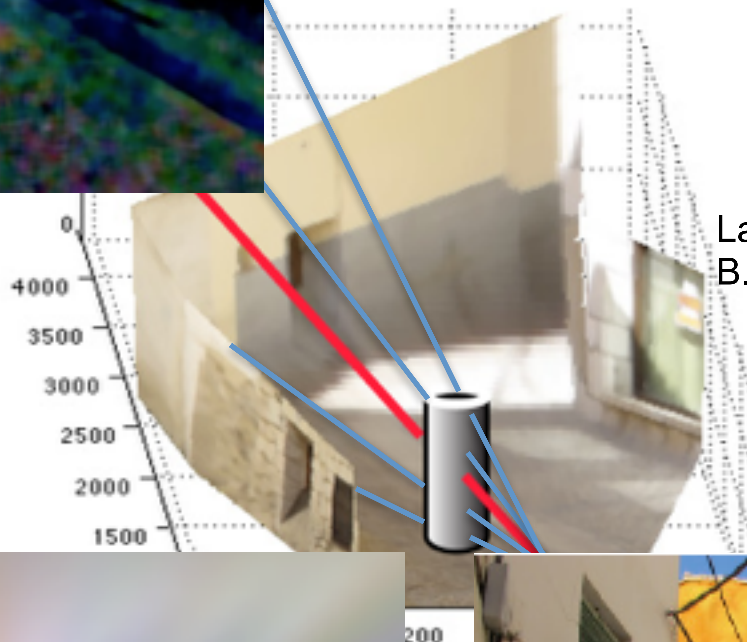

Labelme 3D toolbox.
B. Russell, A. Torralba. CVPR 2009

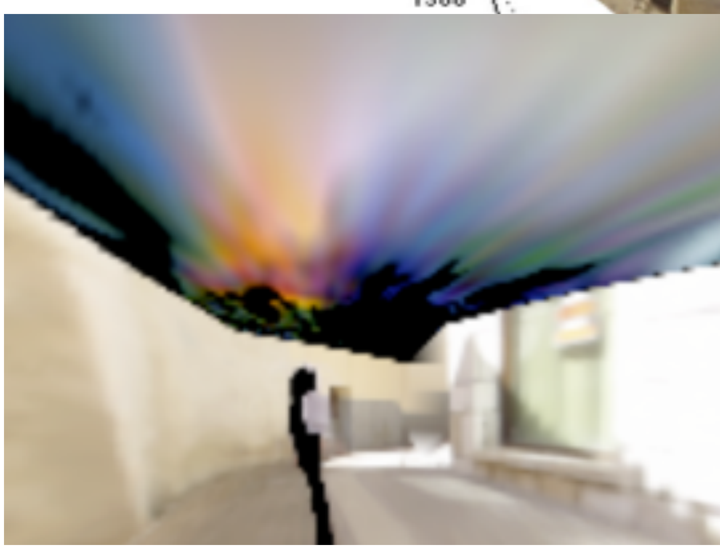

summary of accidental cameras

Shadows and apertures produce accidental images that are unnoticed most of the time. Accidental cameras can reveal the scene outside the picture.

Applications:

- Image forensics (J. O'Brian & H. Farid, 2012)

- Computer graphics providing better light models

Funding for this work was provided by NSF Career award 0747120 to A.T, and NSF CGV 1111415 and NSF CGV 0964004 to W.T.F.

Why do we need lenses?

Two ways to make a pinhole camera image brighter: enlarge the hole, or add a lens

Sum over a cone of rays impinging from the world onto a point. This blurs the image.

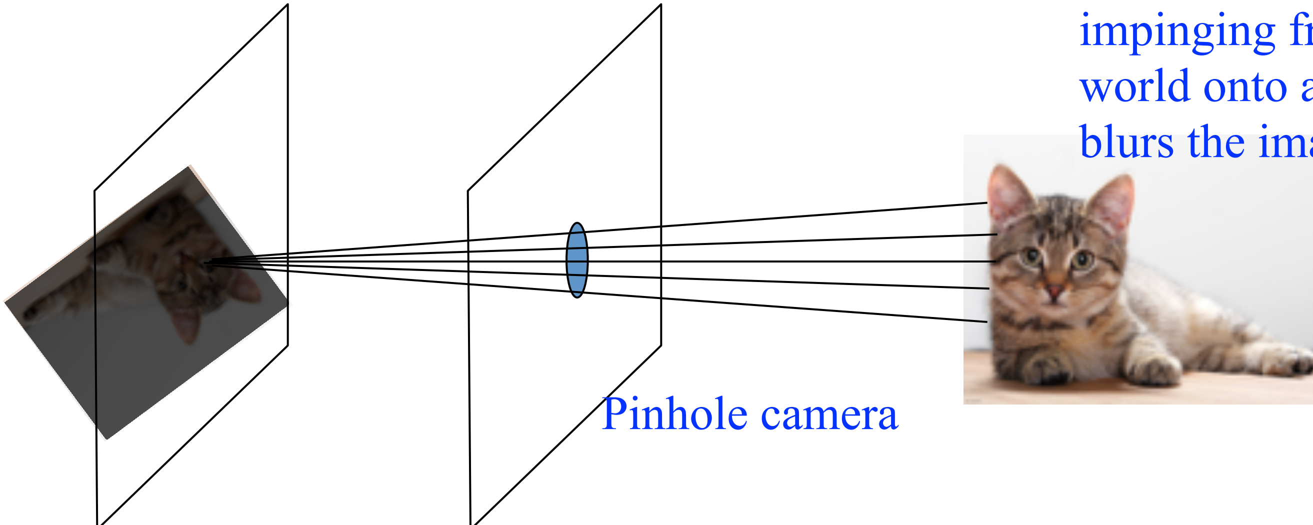

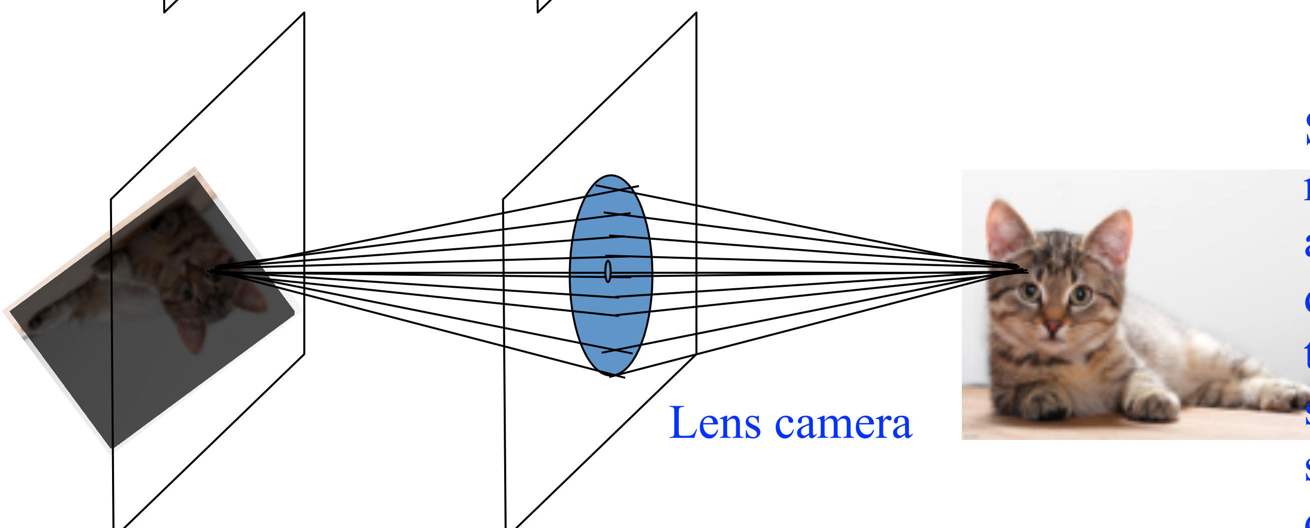

Sum over a cone of rays reflecting off of a surface. For diffuse surfaces, those rays have the same intensities, and summing over them does not blur the image.

Derivation of Snell's law

$$\lambda_1 = \frac{c}{\omega\, n_1}$$

$$\lambda_1 = L \sin(\alpha_1)$$

$$n_1 \sin(\alpha_1) = \frac{c}{\omega L} = n_2 \sin(\alpha_2)$$

Speed, and thus wavelength of light, scales inversely with n. This requires that plane waves bend, according to

$$n_1 \sin(\alpha_1) = n_2 \sin(\alpha_2)$$

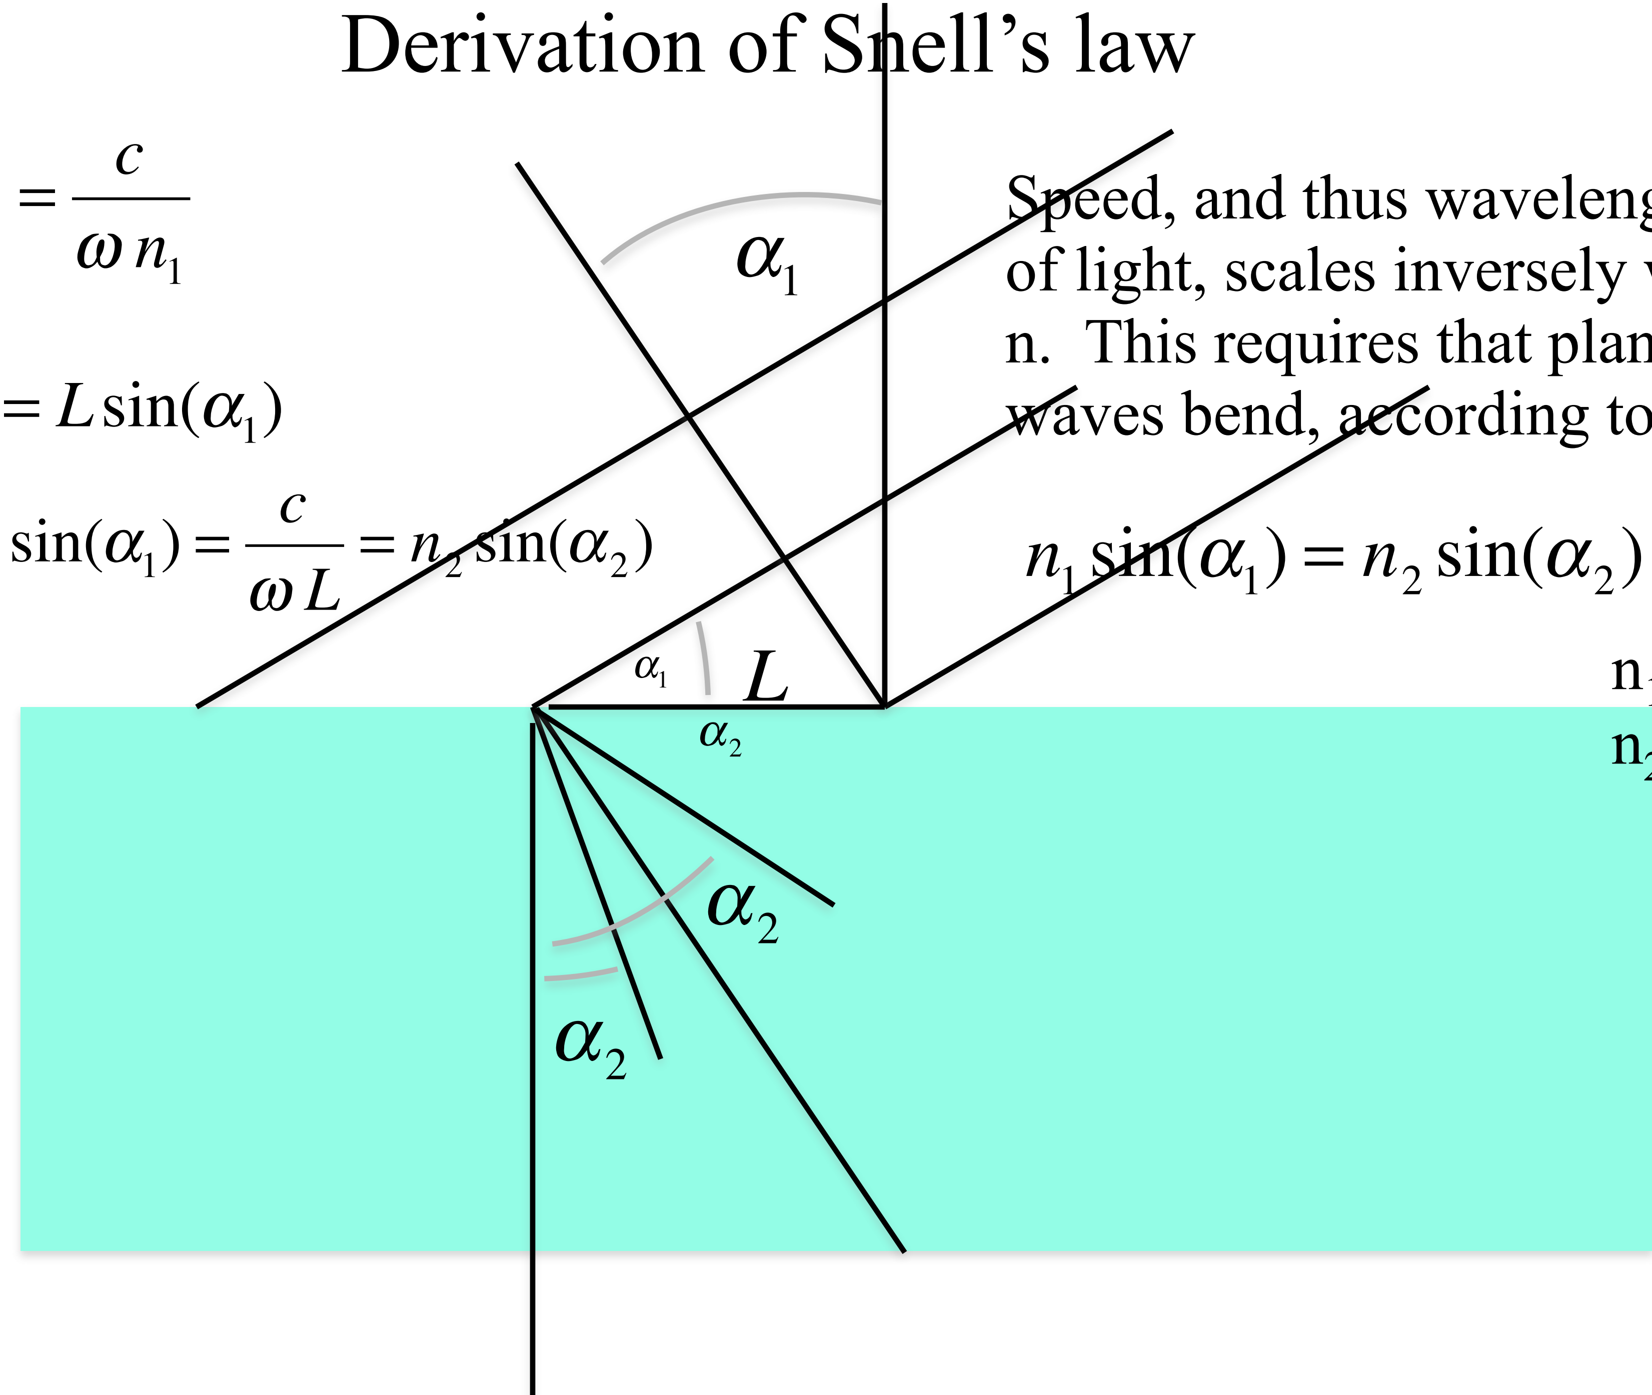

Derivation of Snell's law

$$\lambda_1 = \frac{c}{\omega n_1}$$

$$\lambda_1 = L\sin(\alpha_1)$$

$$n_1\sin(\alpha_1) = \frac{c}{\omega L} = n_2\sin(\alpha_2)$$

Speed, and thus wavelength of light, scales inversely with n. This requires that plane waves bend, according to

$$n_1\sin(\alpha_1) = n_2\sin(\alpha_2)$$

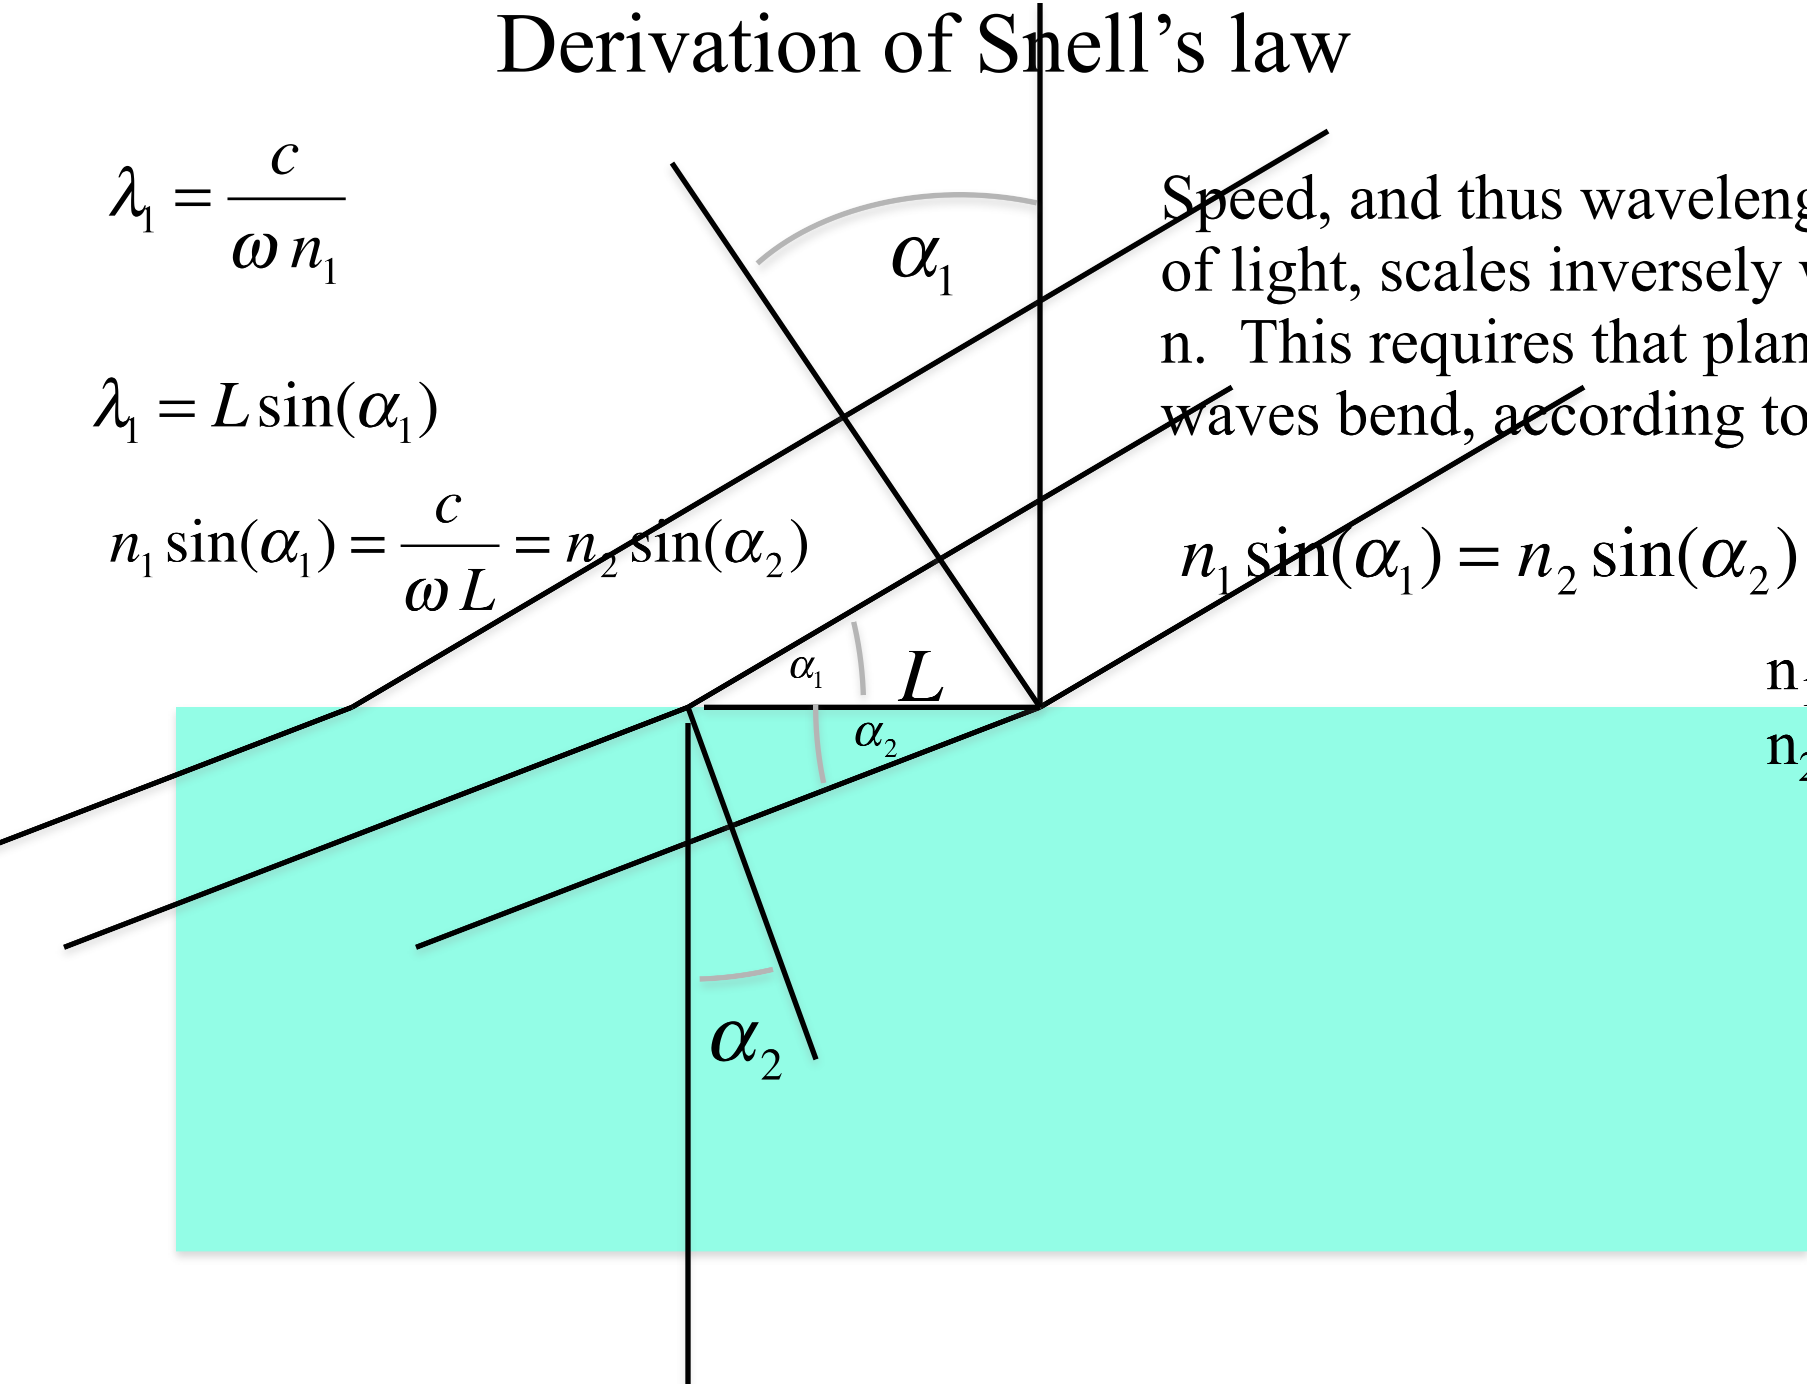

Refraction: Snell's law

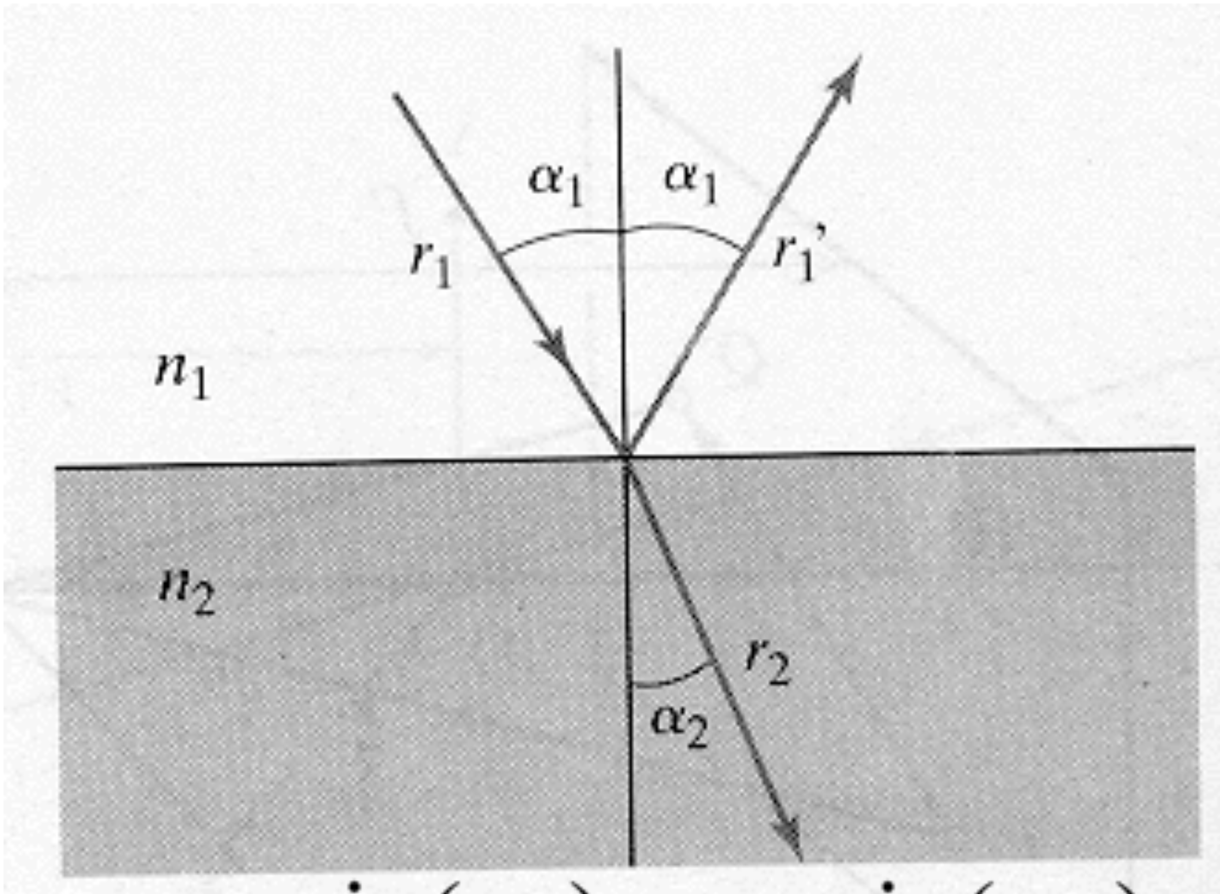

$$n_1 \sin(\alpha_1) = n_2 \sin(\alpha_2)$$

For small angles, $n_1\alpha_1 \approx n_2\alpha_2$

Spherical lens

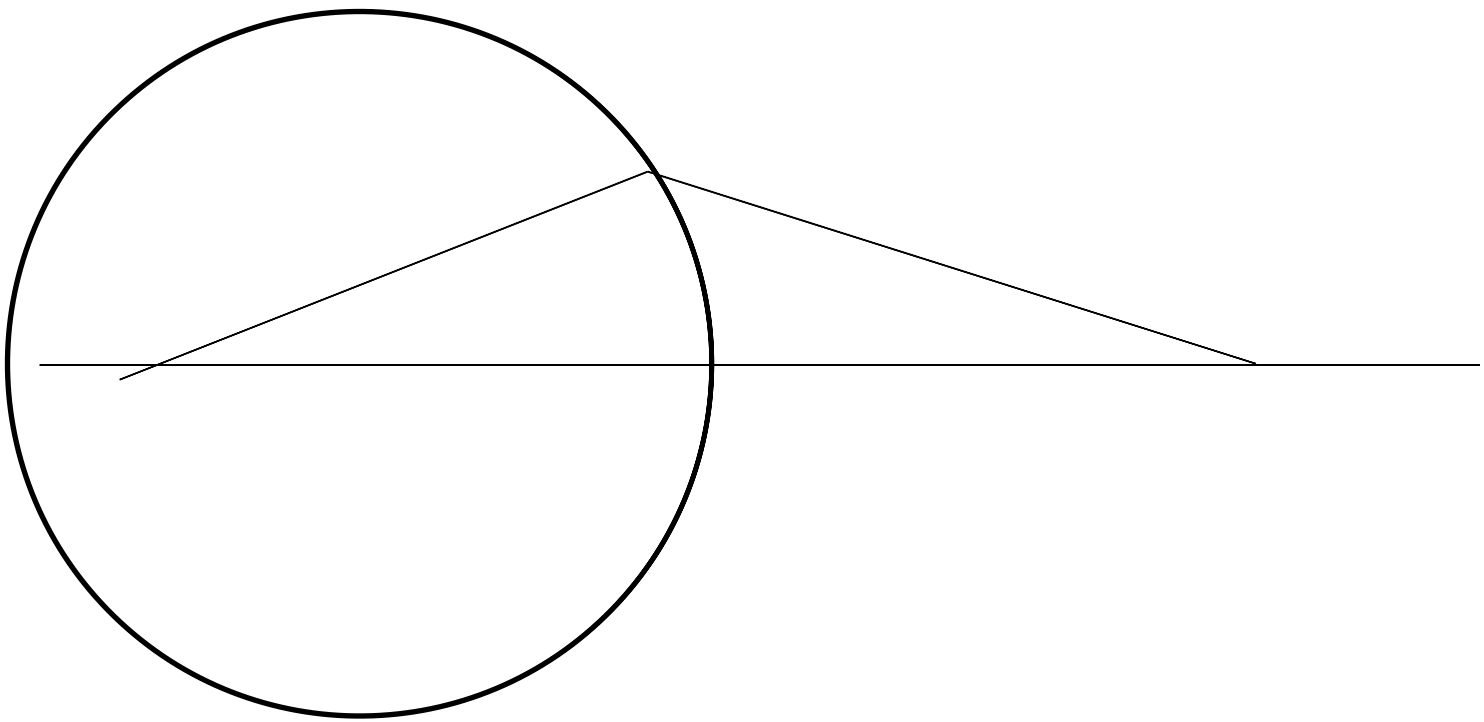

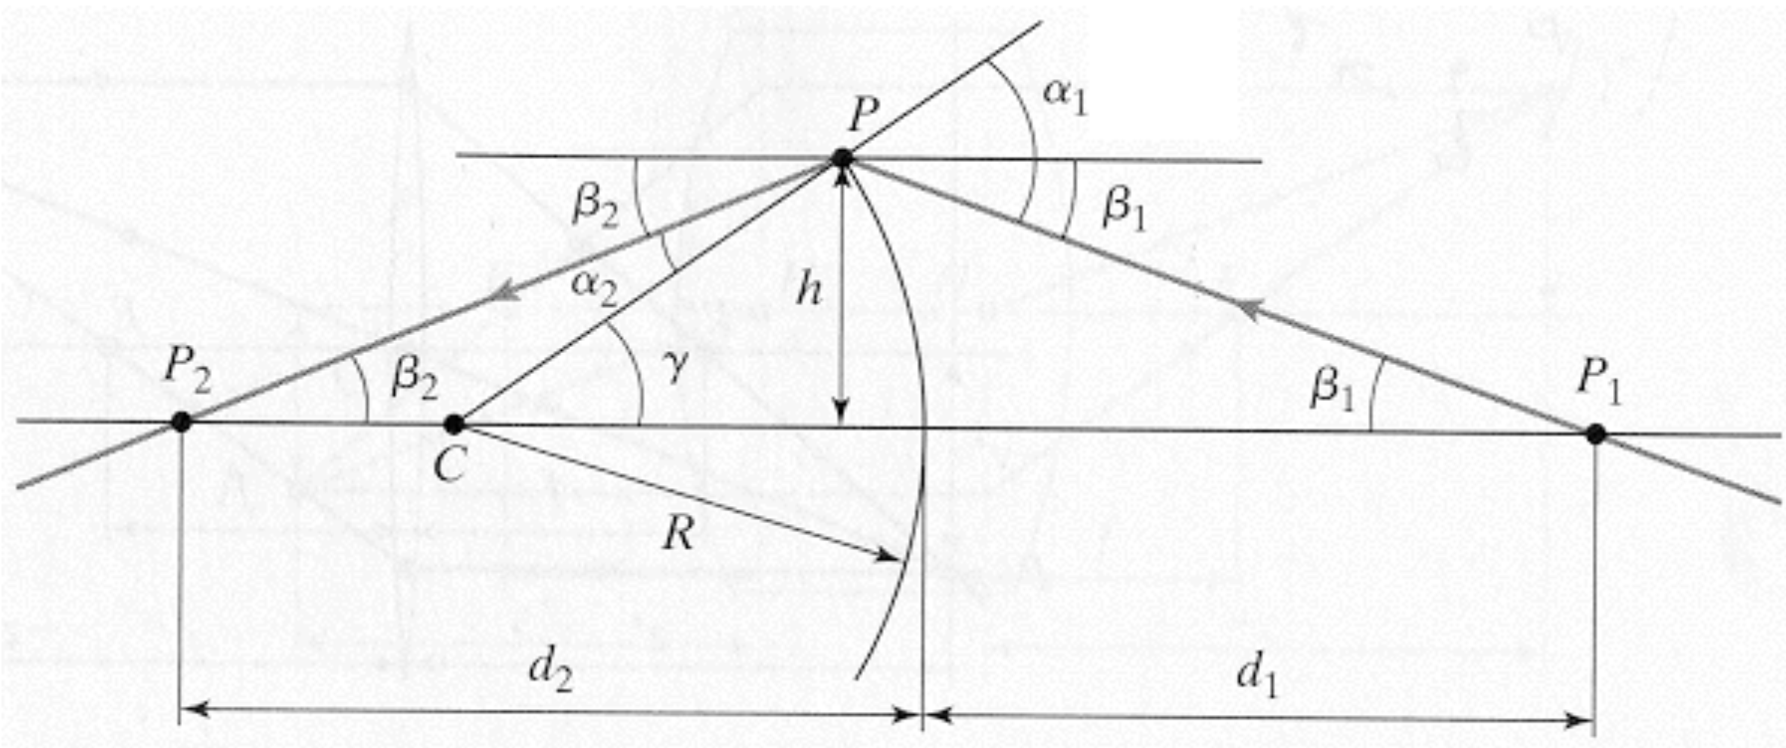

Forsyth and Ponce

First order optics

$$\sin(\theta) \approx \theta$$

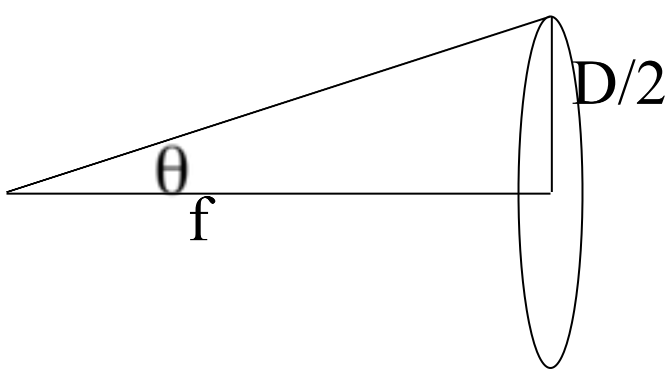

$$\theta \approx \frac{D/2}{f}$$

Paraxial refraction equation

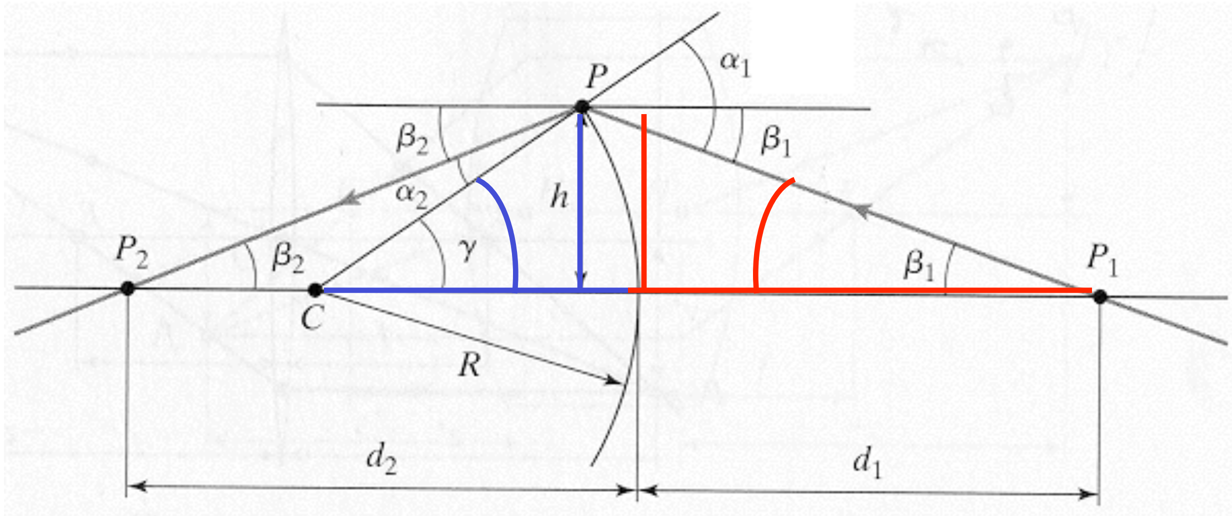

$$\alpha_1 = \boxed{\gamma} + \boxed{\beta_1} \approx h\left(\frac{1}{R} + \frac{1}{d_1}\right)$$

Relates distance from sphere and sphere radius to bending angles of light ray, for lens with one spherical surface.

$$n_1\alpha_1 \approx n_2\alpha_2$$

Deriving the lensmaker's formula

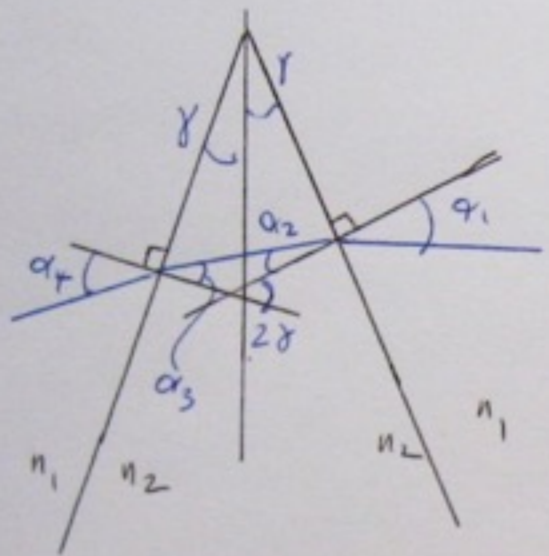

$\alpha_1 = h\left(\frac{1}{R} + \frac{1}{d_1}\right)$ small angle approx

$n_1\alpha_1 \cong n_2\alpha_2$ snell's law

$\alpha_2 = 2\gamma - \alpha_3$ geometry

$n_2\alpha_3 \cong n_1\alpha_4$ snell's law

$\alpha_4 = h_1\left(\frac{1}{R} + \frac{1}{d_2}\right)$ small angle

$\gamma = \frac{h}{R}$ small angle

$$n_1\alpha_1 = n_2\left(\frac{2h}{R} - \frac{n_1}{n_2}\alpha_4\right) = h\left(\frac{1}{R} + \frac{1}{d_1}\right)$$

let $n_1 = 1$, $n_2 = n$

cancel h's $\quad n\left(\frac{2}{R} - \frac{1}{n}\left(\frac{1}{R} + \frac{1}{d_2}\right)\right) = \frac{1}{R} + \frac{1}{d_1}$

$$\frac{2n}{R} - \frac{1}{R} - \frac{1}{d_2} = \frac{1}{R} + \frac{1}{d_1}$$

$$\frac{2(n-1)}{R} = \frac{1}{d_1} + \frac{1}{d_2} \qquad \text{"Lens maker's formula"}$$

The thin lens, first order optics

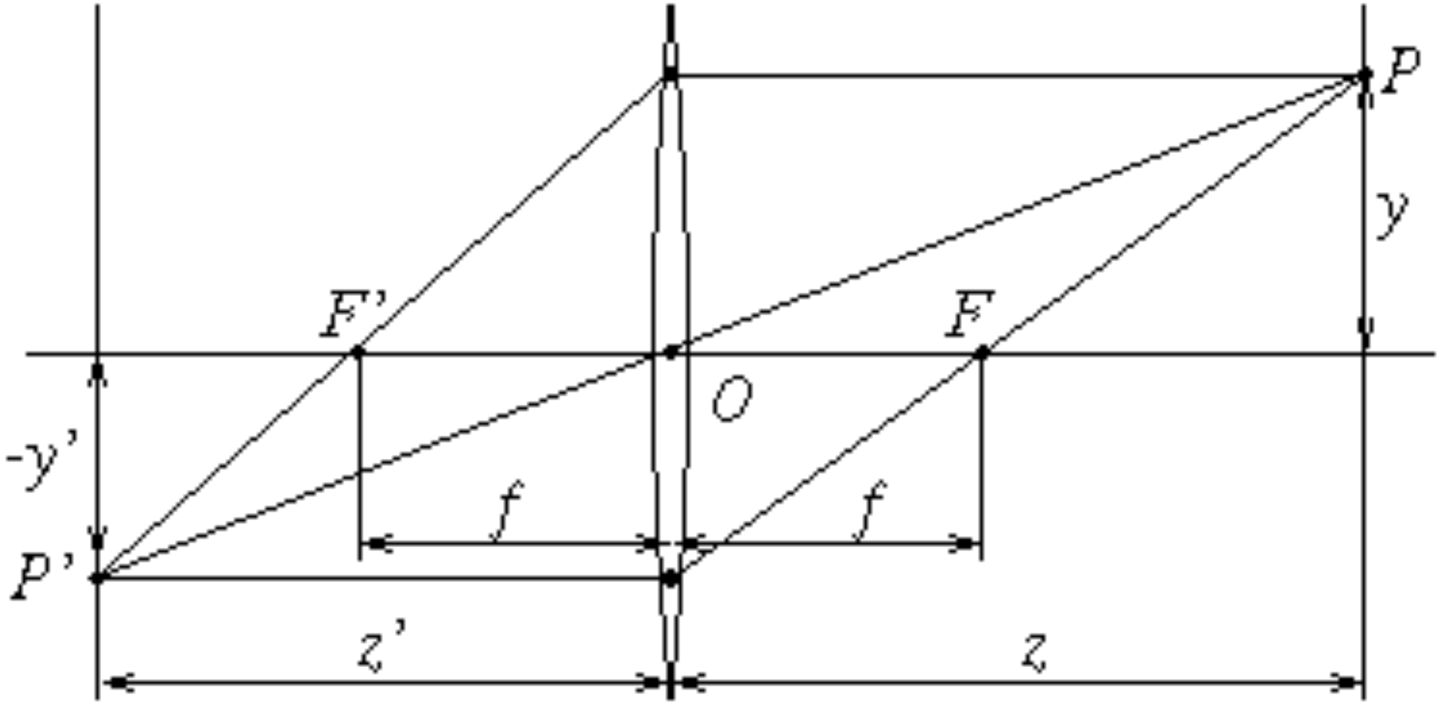

The lensmaker's equation:

$$\frac{1}{z'} + \frac{1}{z} = \frac{1}{f} \qquad f = \frac{R}{2(n-1)}$$

Forsyth&Ponce

How does the mapping from the 3-d world to the image plane compare for a lens and for a pinhole camera?

The perspective projection of a pinhole camera. But note that many more of the rays leaving from P arrive at P'

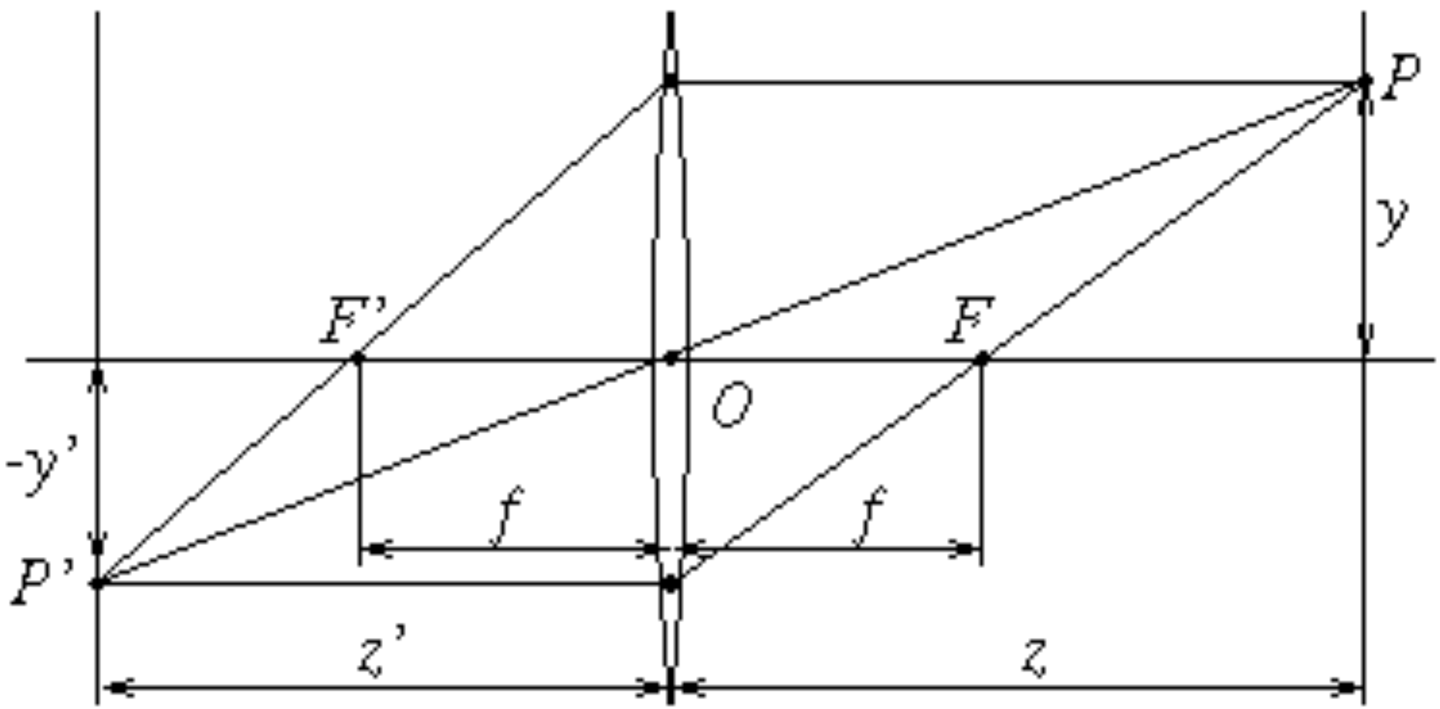

Lens demonstration

- Verify:
  - Focusing property
  - Lens maker's equation

Some animal imaging systems

Lenses

Pinholes

Anti-pinholes

Like other Euglenoids, *Euglena* possess a red eyespot, an organelle composed of carotenoid pigment granules. The red spot itself is not thought to be photosensitive. Rather, it filters the sunlight that falls on a light-detecting structure at the base of the flagellum (a swelling, known as the paraflagellar body), allowing only certain wavelengths of light to reach it. As the cell rotates with respect to the light source, the eyespot partially blocks the source, permitting the *Euglena* to find the light and move toward it (a process known as phototaxis).[11]

Animal Eyes

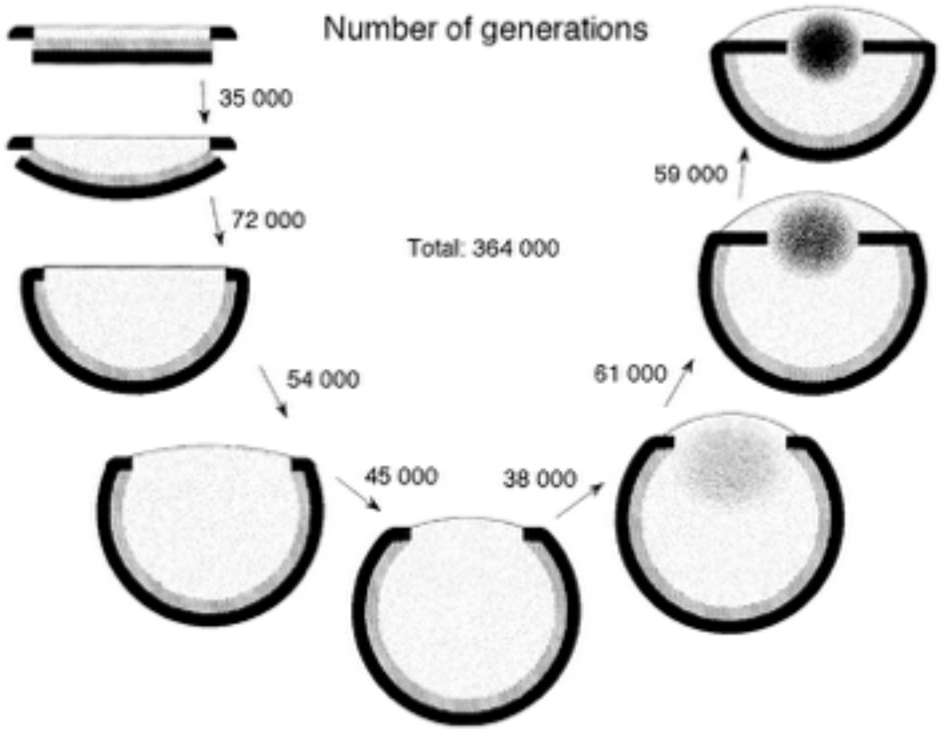

Fig. 1.6 A patch of light sensitive epithelium can be gradually turned into a perfectly focussed camera-type eye if there is a continuous selection for improved spatial vision. A theoretical model based on conservative assumptions about selection pressure and the amount of variation in natural populations suggest that the whole sequence can be accomplished amazingly fast, in less than 400 000 generations. The number of generations is also given between each of the consecutive intermediates that are drawn in the figure. The starting point is a flat piece of epithelium with an outer protective layer, an intermediate layer of receptor cells, and a bottom layer of pigment cells. The first half of the sequence is the formation of a pigment cup eye. When this principle cannot be improved any further, a lens gradually evolves. Modified from Nilsson and Pelger (1994).

Animal Eyes. Land & Nilsson. Oxford Univ. Press

Camera Models

Right - handed system

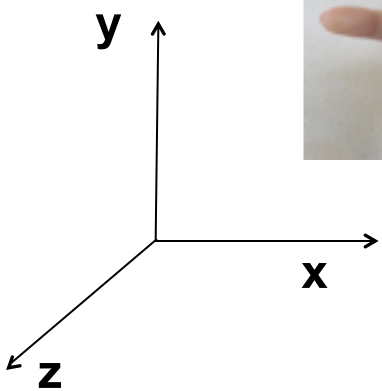

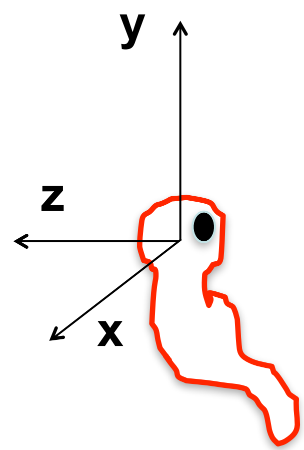

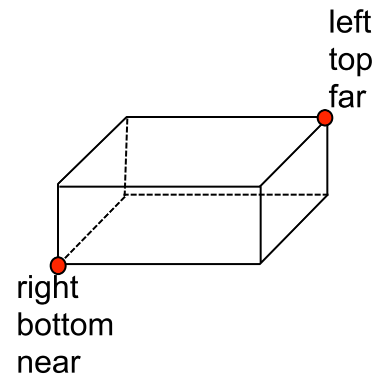

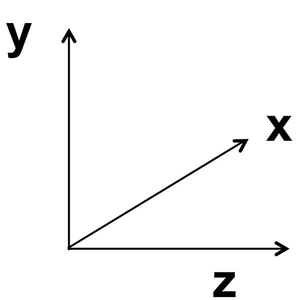

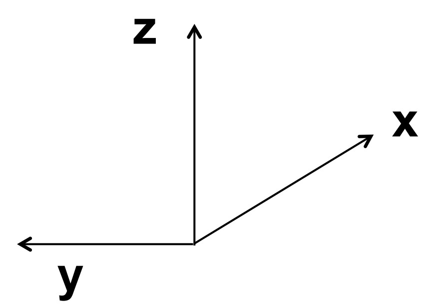

Perspective projection

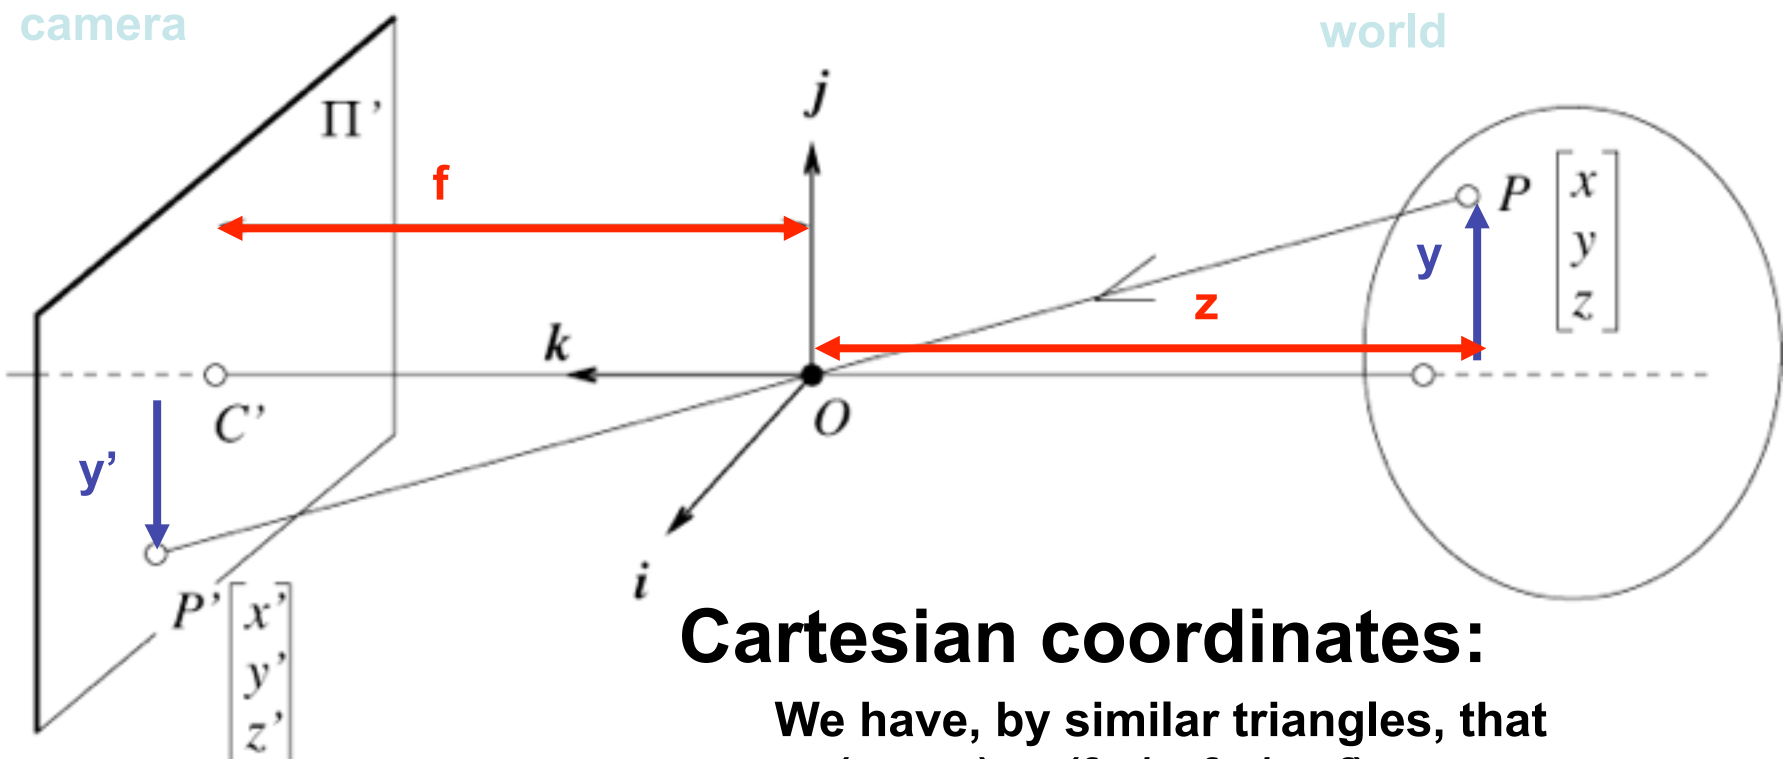

Cartesian coordinates:

We have, by similar triangles, that
(x, y, z) -> (f x/z, f y/z, -f)

Ignore the third coordinate, and get

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

Geometric properties of projection

- Points go to
- Lines go to
- Planes go to
- Polygons go to
- Degenerate cases
  - line through focal point to point
  - plane through focal point to line

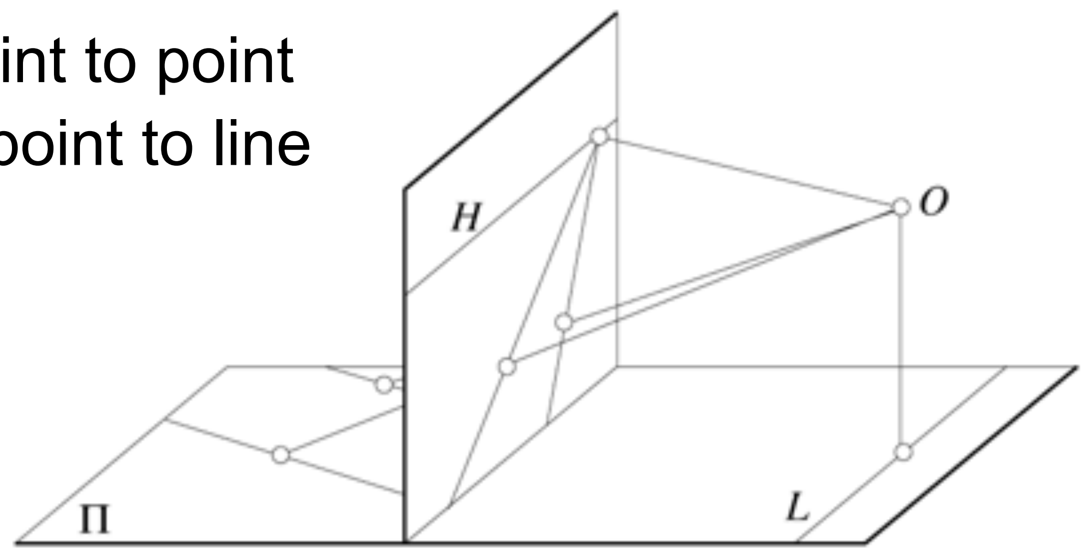

Vanishing point

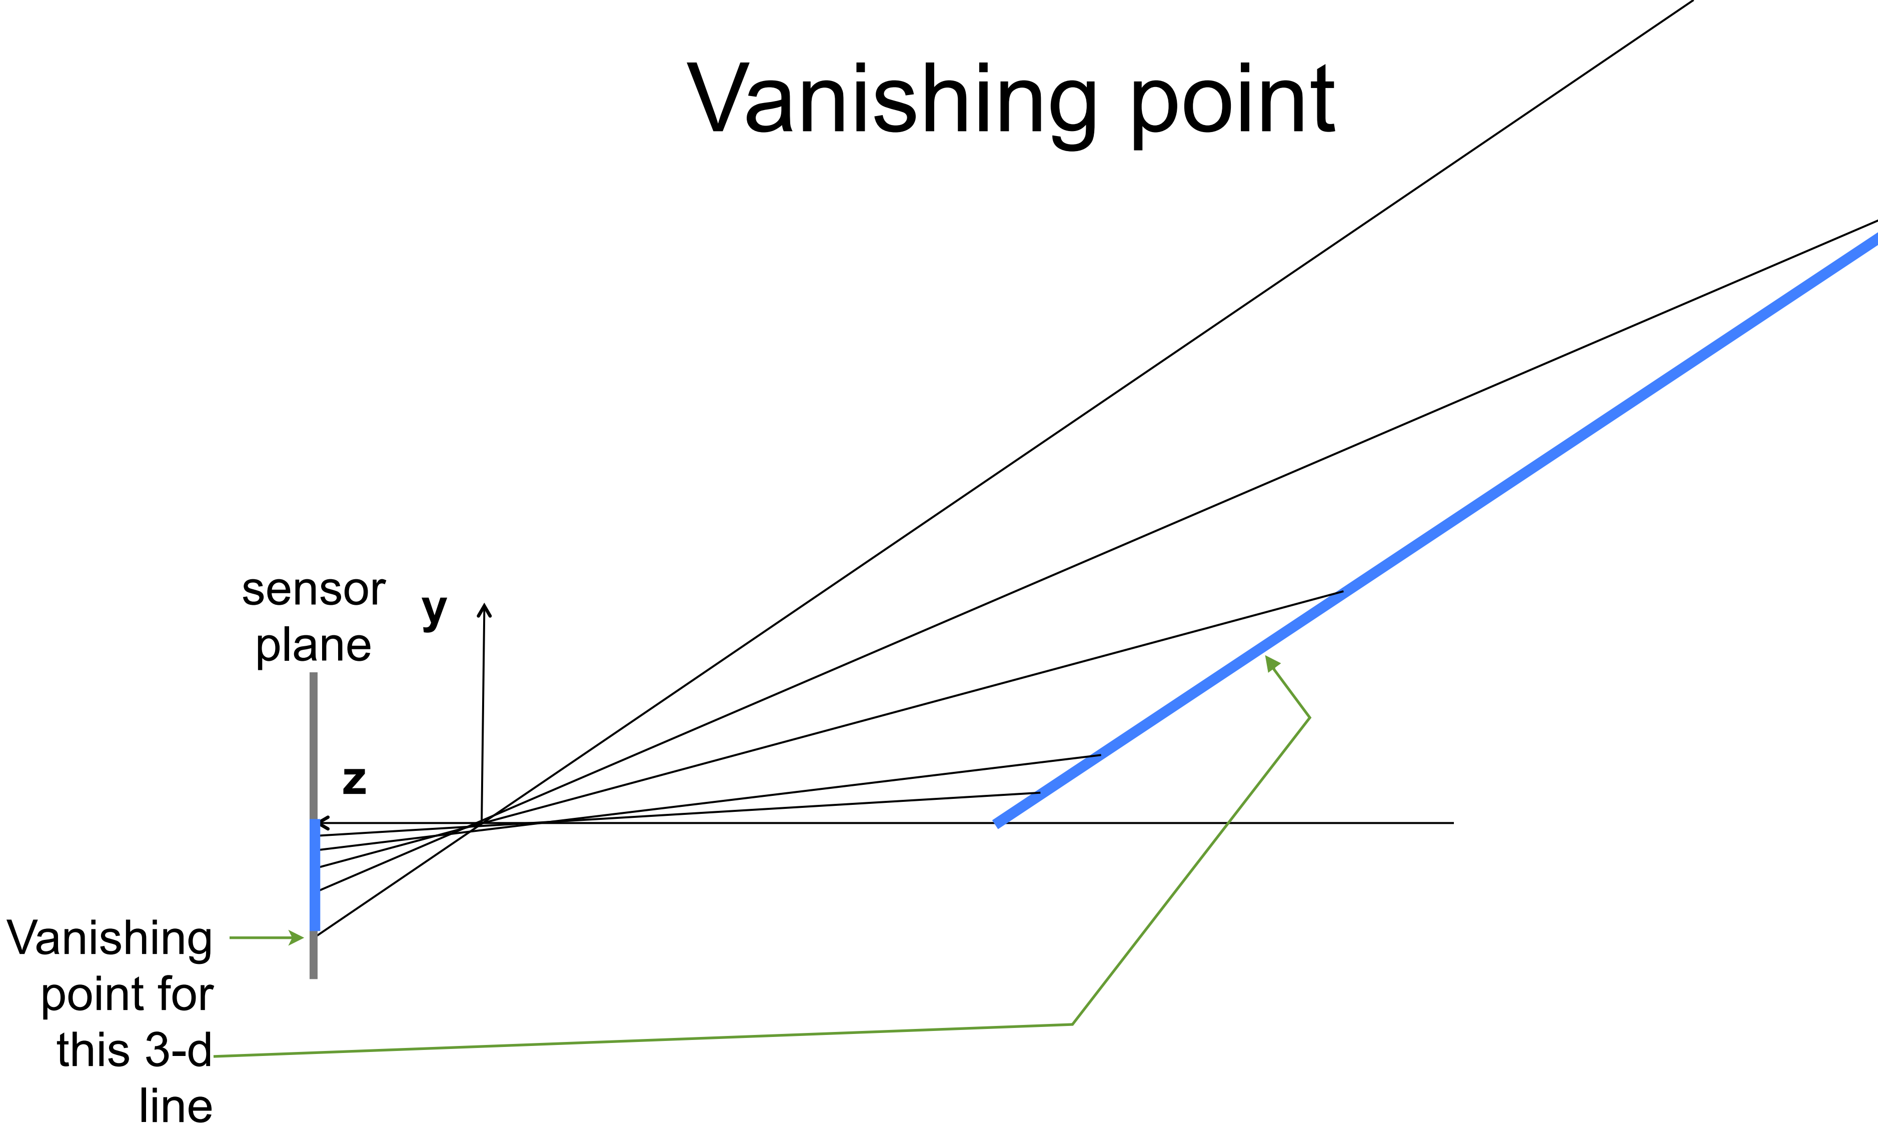

Line in 3-space

$$x(t) = x_0 + at$$
$$y(t) = y_0 + bt$$
$$z(t) = z_0 + ct$$

Perspective projection of that line

$$x'(t) = \frac{fx}{z} = \frac{f(x_0 + at)}{z_0 + ct}$$
$$y'(t) = \frac{fy}{z} = \frac{f(y_0 + bt)}{z_0 + ct}$$

In the limit as $t \longrightarrow \pm\infty$ **we have (for** $c \neq 0$**):**

$$x'(t) \longrightarrow \frac{fa}{c}$$
$$y'(t) \longrightarrow \frac{fb}{c}$$

This tells us that any set of parallel lines (same a, b, c parameters) project to the same point (called the vanishing point).

Vanishing points

- Each set of parallel lines (=direction) meets at a different point
  - The *vanishing point* for this direction
- Sets of parallel lines on the same plane lead to *collinear* vanishing points.
  - The line is called the *horizon* for that plane

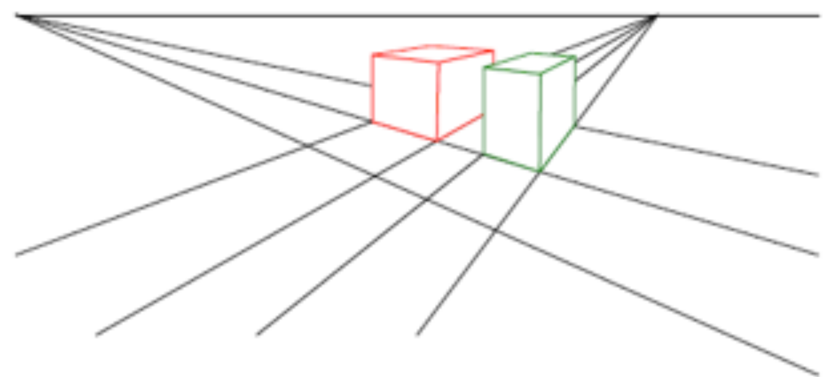

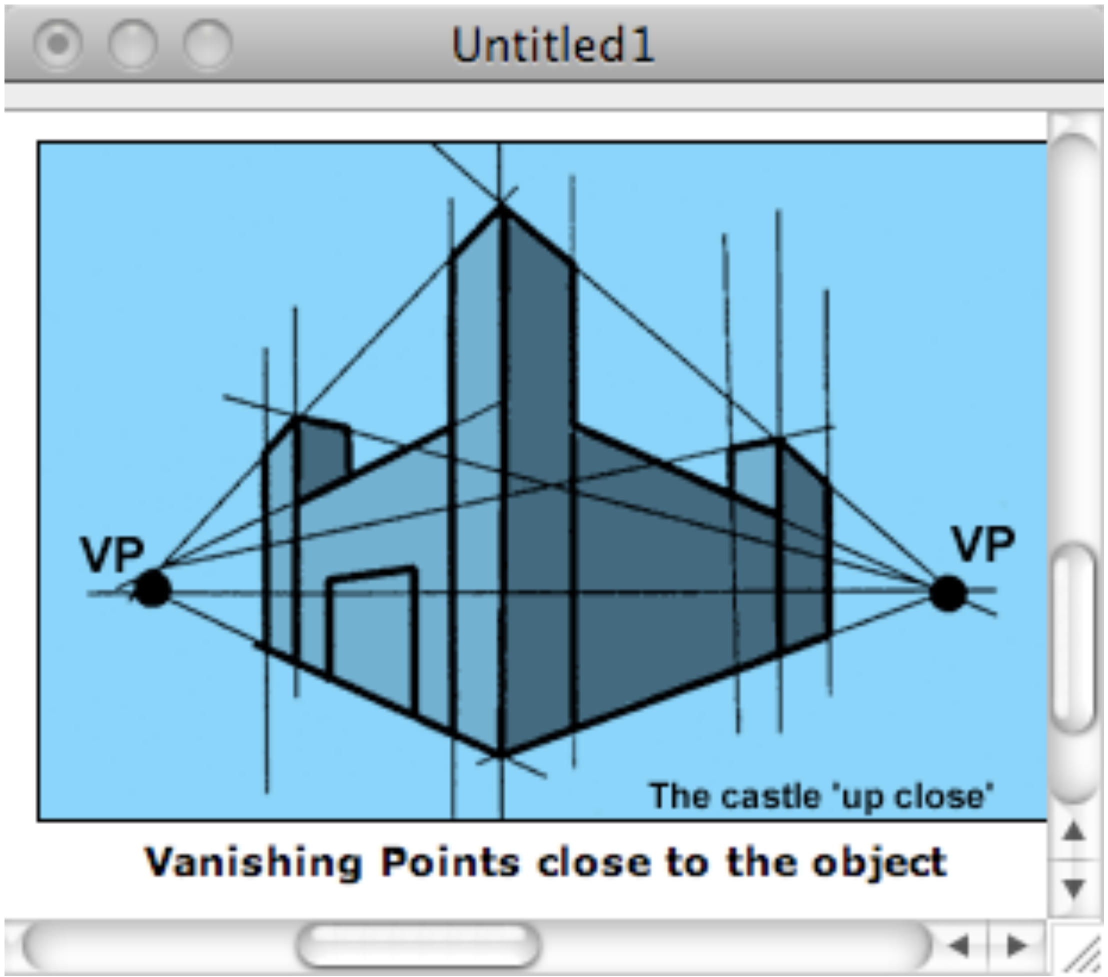

http://www.ider.herts.ac.uk/school/courseware/graphics/two_point_perspective.html

What if you photograph a brick wall head-on?

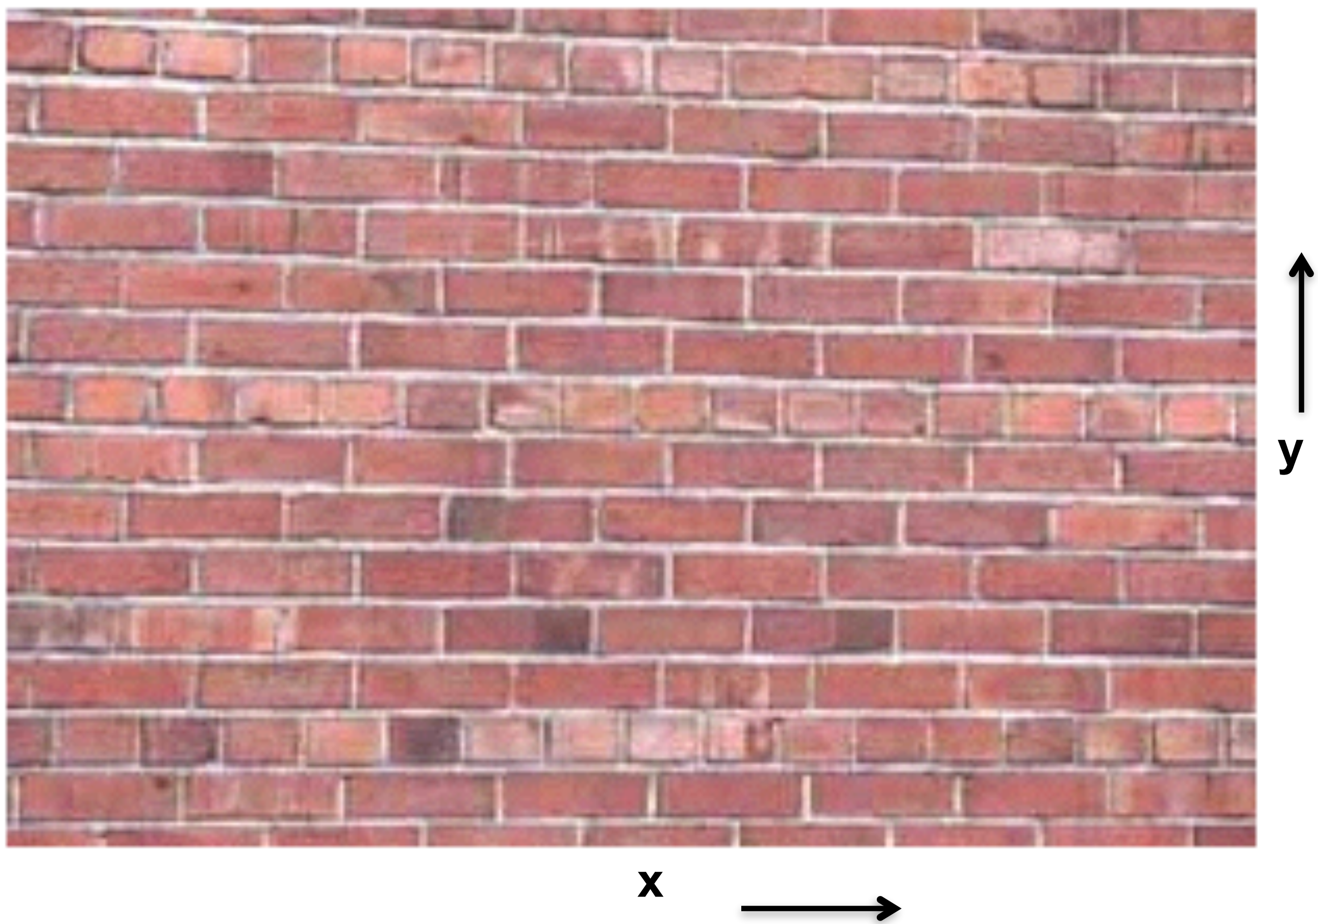

Brick wall line in 3-space

$$x(t) = x_0 + at$$

$$y(t) = y_0$$

$$z(t) = z_0$$

Perspective projection of that line

$$x'(t) = \frac{f \cdot (x_0 + at)}{z_0}$$

$$y'(t) = \frac{f \cdot y_0}{z_0}$$

All bricks have same z_0. Those in same row have same y_0

Thus, a brick wall, photographed head-on, gets rendered as set of parallel lines in the image plane.

Other projection models: Orthographic projection

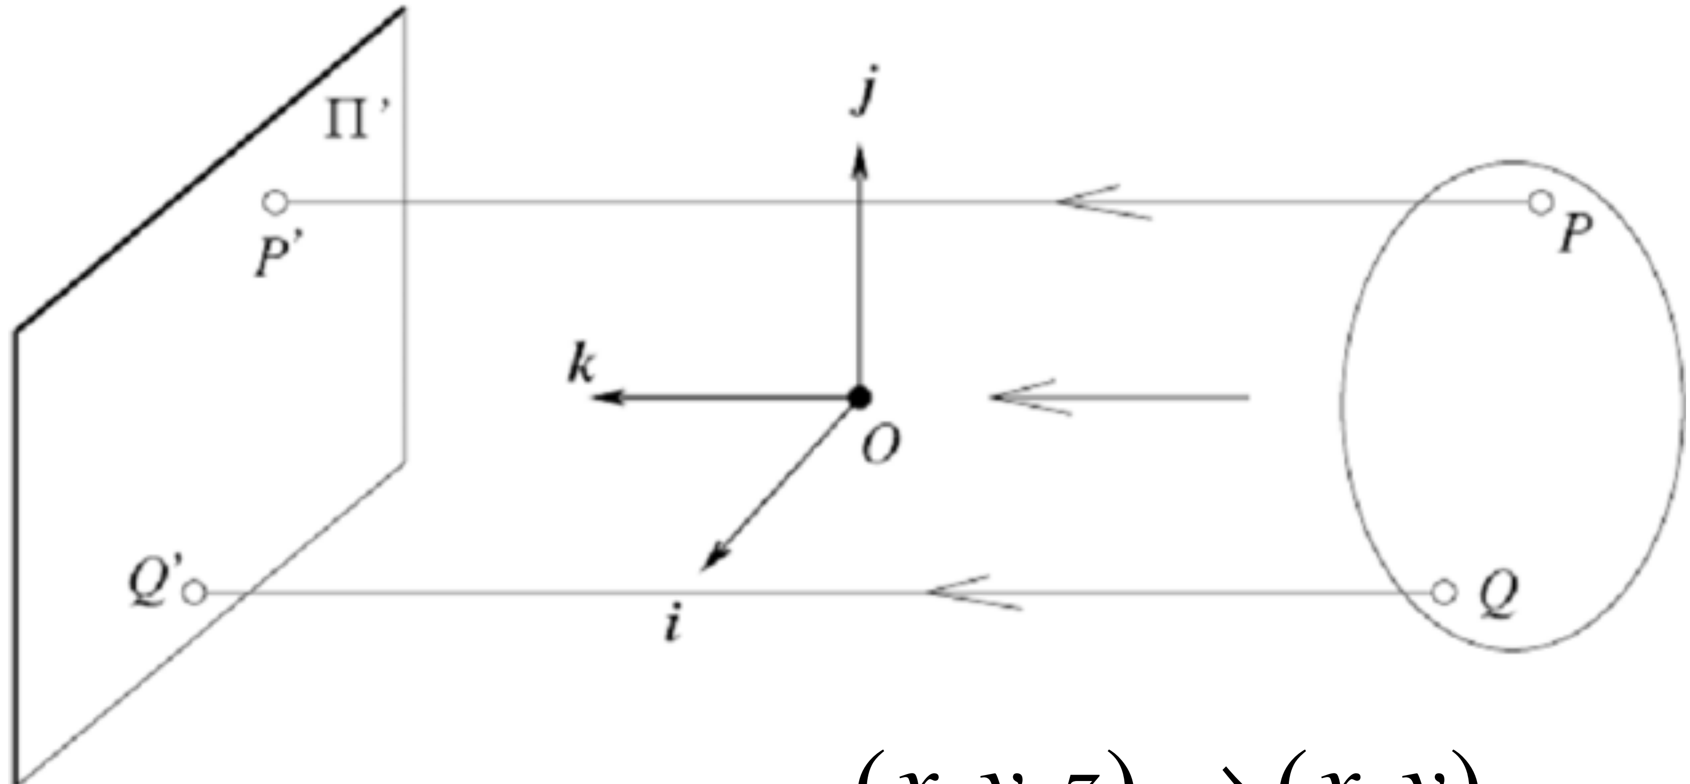

$$(x, y, z) \rightarrow (x, y)$$

Other projection models: Weak perspective

- Issue
  - perspective effects, but not over the scale of individual objects
  - collect points into a group at about the same depth, then divide each point by the depth of its group
  - Adv: easy
  - Disadv: only approximate

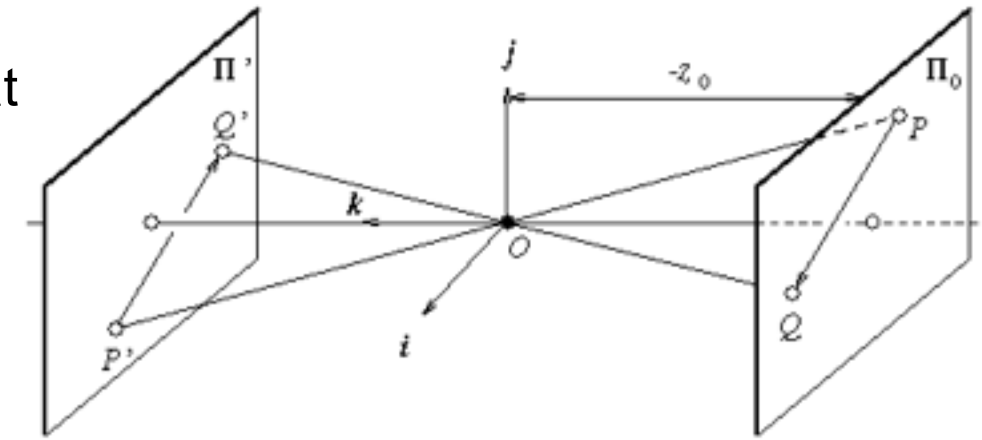

$$(x, y, z) \rightarrow \left(\frac{fx}{z_0}, \frac{fy}{z_0} \right)$$

Three camera projections

3-d point **2-d image position**

(1) Perspective: $(x, y, z) \rightarrow \left(\frac{fx}{z}, \frac{fy}{z} \right)$

(2) Weak perspective: $(x, y, z) \rightarrow \left(\frac{fx}{z_0}, \frac{fy}{z_0} \right)$

(3) Orthographic: $(x, y, z) \rightarrow (x, y)$

Three camera projections

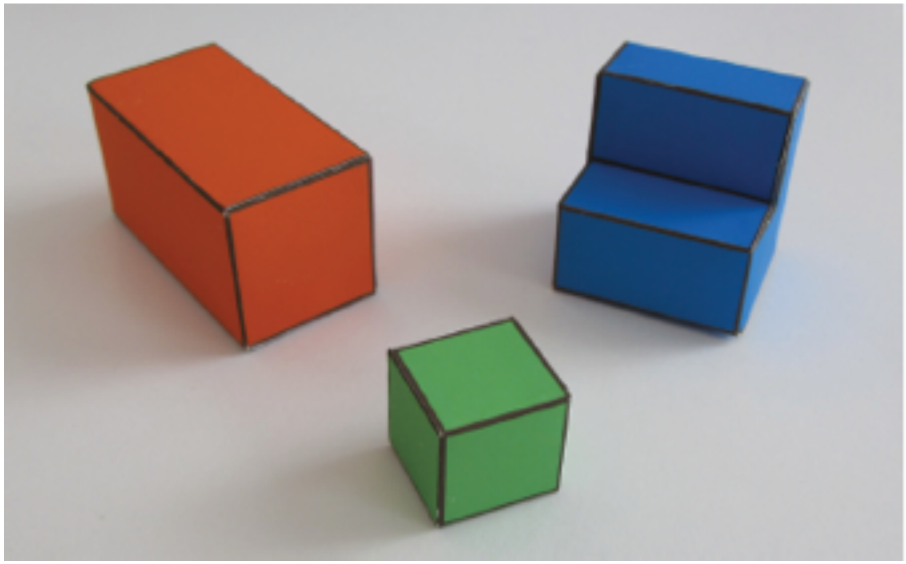

Perspective projection

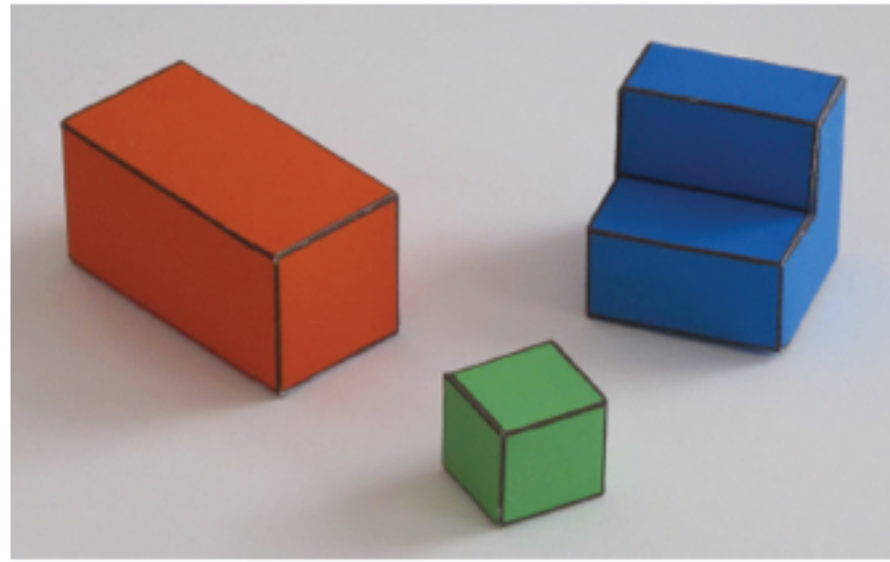

Parallel (orthographic) projection

Weak perspective?

More accurate models of real lenses

- Finite lens thickness
- Higher order approximation to $\sin(\theta)$
- Chromatic aberration
- Vignetting

Thick lens

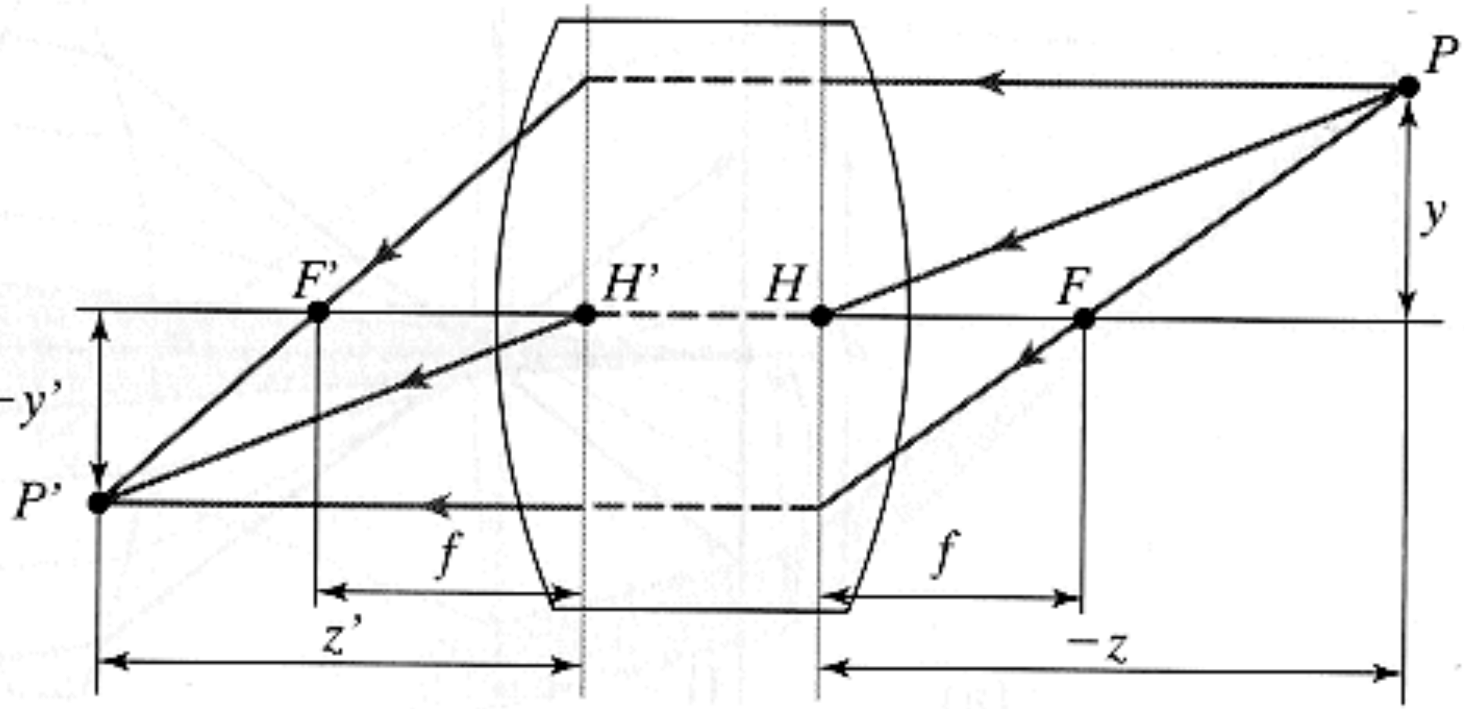

Figure 1.11 A simple thick lens with two spherical surfaces.

Forsyth&Ponce

Third order optics

$$\sin(\theta) \approx \theta - \frac{\theta^3}{6}$$

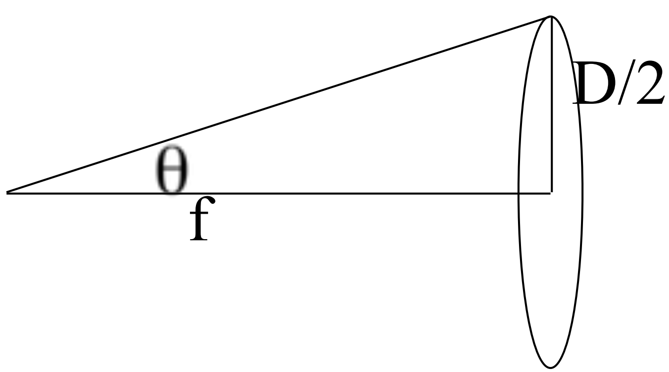

$$\theta \approx \frac{D/2}{f} - \frac{\left(\frac{D/2}{f}\right)^3}{6}$$

Paraxial refraction equation, 3rd order optics

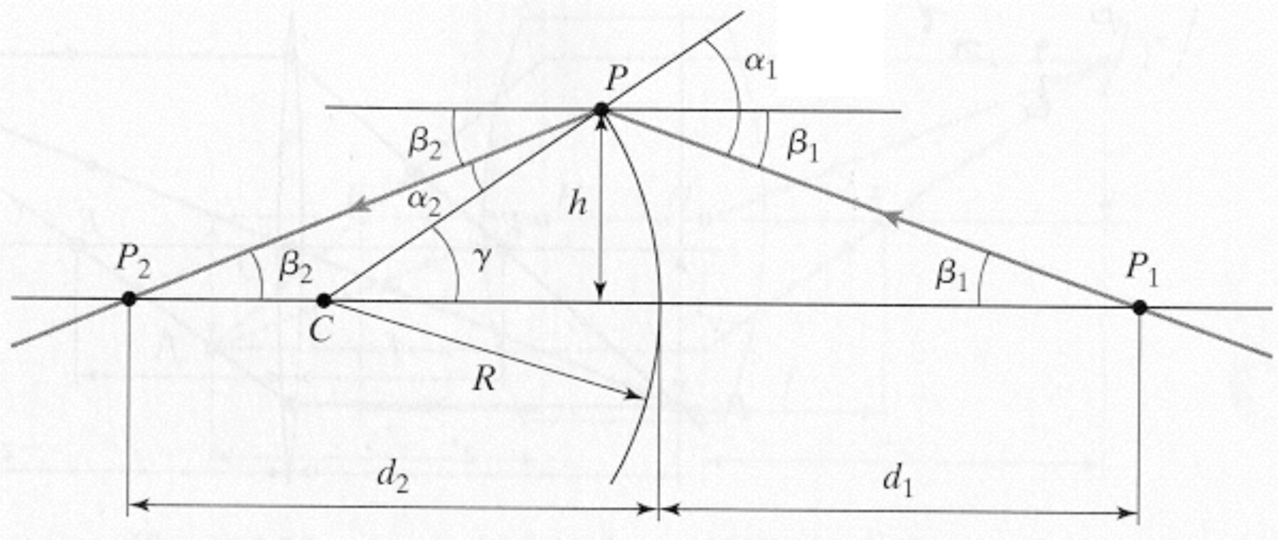

$$\frac{n_1}{d_1} + \frac{n_2}{d_2} = \frac{n_2 - n_1}{R} + h^2 \left[\frac{n_1}{2d_1} \left(\frac{1}{R} + \frac{1}{d_1} \right)^2 + \frac{n_2}{2d_2} \left(\frac{1}{R} - \frac{1}{d_2} \right)^2 \right]$$

Forsyth&Ponce

Spherical aberration (from 3rd order optics

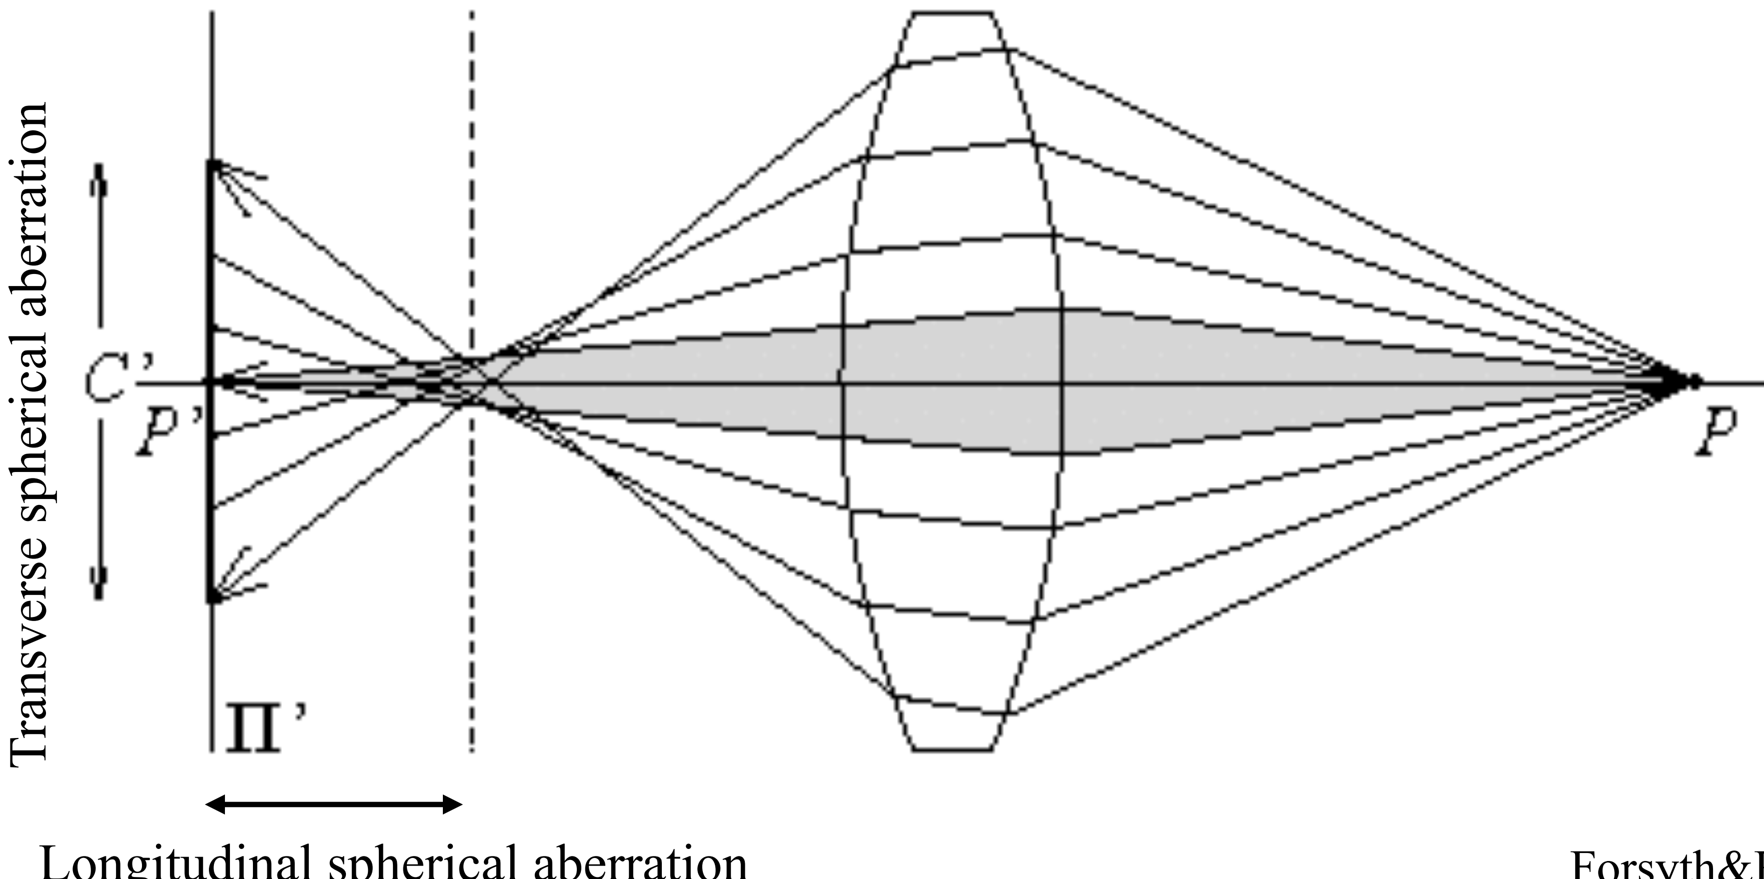

Forsyth&Ponce

Other 3rd order effects

- Coma, astigmatism, field curvature, distortion.

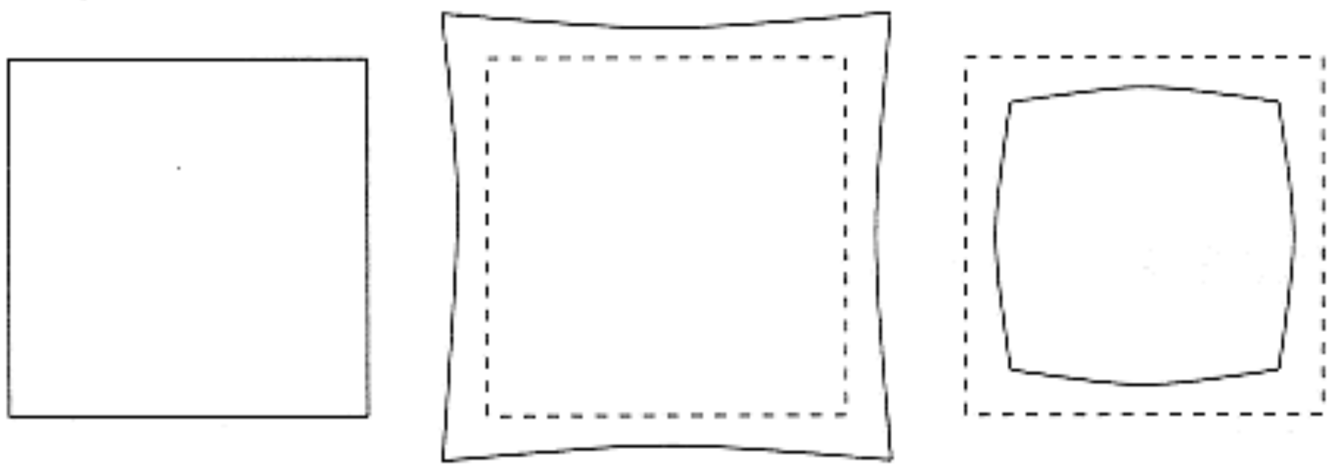

no distortion

pincushion distortion

barrel distortion

Forsyth&Ponce

Chromatic aberration

(desirable for prisms, bad for lenses)

Other (possibly annoying) phenomena

- Chromatic aberration
  - Light at different wavelengths follows different paths; hence, some wavelengths are defocussed
  - Machines: coat the lens
  - Humans: live with it
- Scattering at the lens surface
  - Some light entering the lens system is reflected off each surface it encounters (Fresnel's law gives details)
  - Machines: coat the lens, interior
  - Humans: live with it (various scattering phenomena are visible in the human eye)

Summary

- Want to make images
- Pinhole camera models the geometry of perspective projection
- Lenses make it work in practice
- Models for lenses
  - Thin lens, spherical surfaces, first order optics
  - Thick lens, higher-order optics, vignetting.