Optics I: lenses and apertures

CS 178, Spring 2011

Begun 4/5/11, finished 4/7. Error on slide 63 corrected 4/12.

Marc Levoy
Computer Science Department
Stanford University

Outline

- why study lenses?
- thin lenses
  - graphical constructions, algebraic formulae
- thick lenses
  - center of perspective, 3D perspective transformations
- depth of field

- aberrations & distortion
- vignetting, glare, and other lens artifacts
- diffraction and lens quality
- special lenses
  - telephoto, zoom

Cameras and their lenses

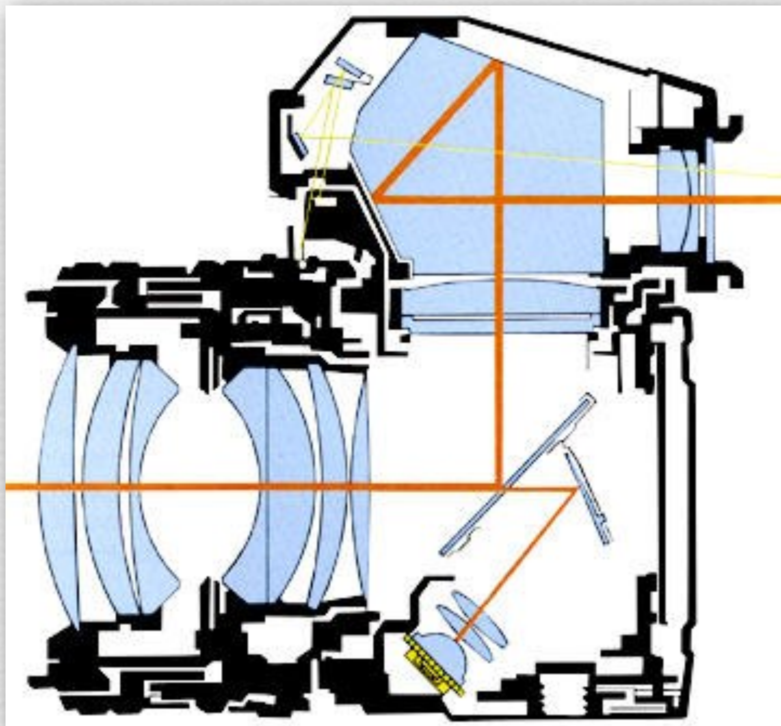

single lens reflex (SLR) camera

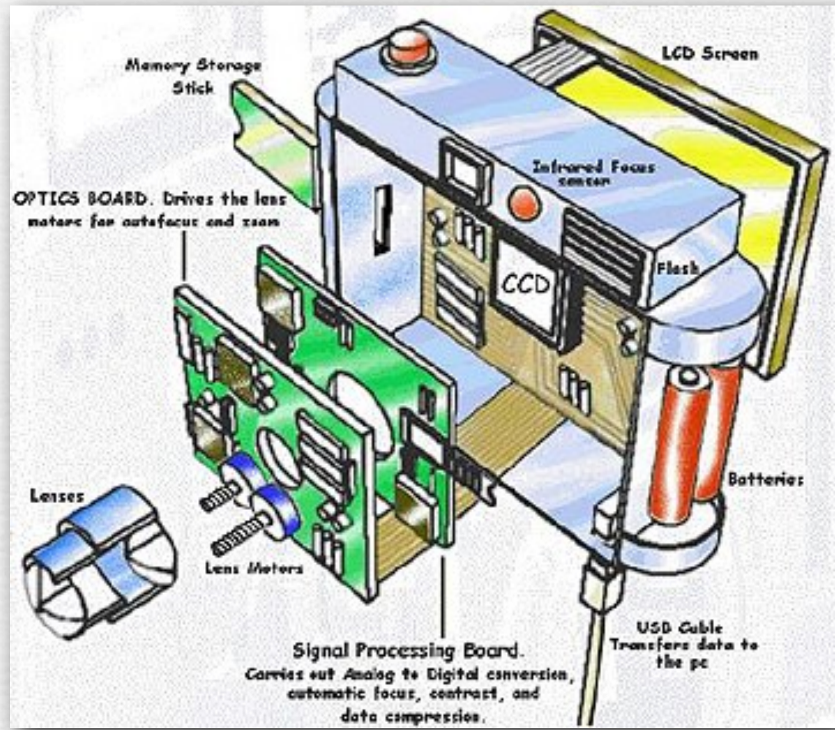

digital still camera (DSC), i.e. point-and-shoot

Cutaway view of a real lens

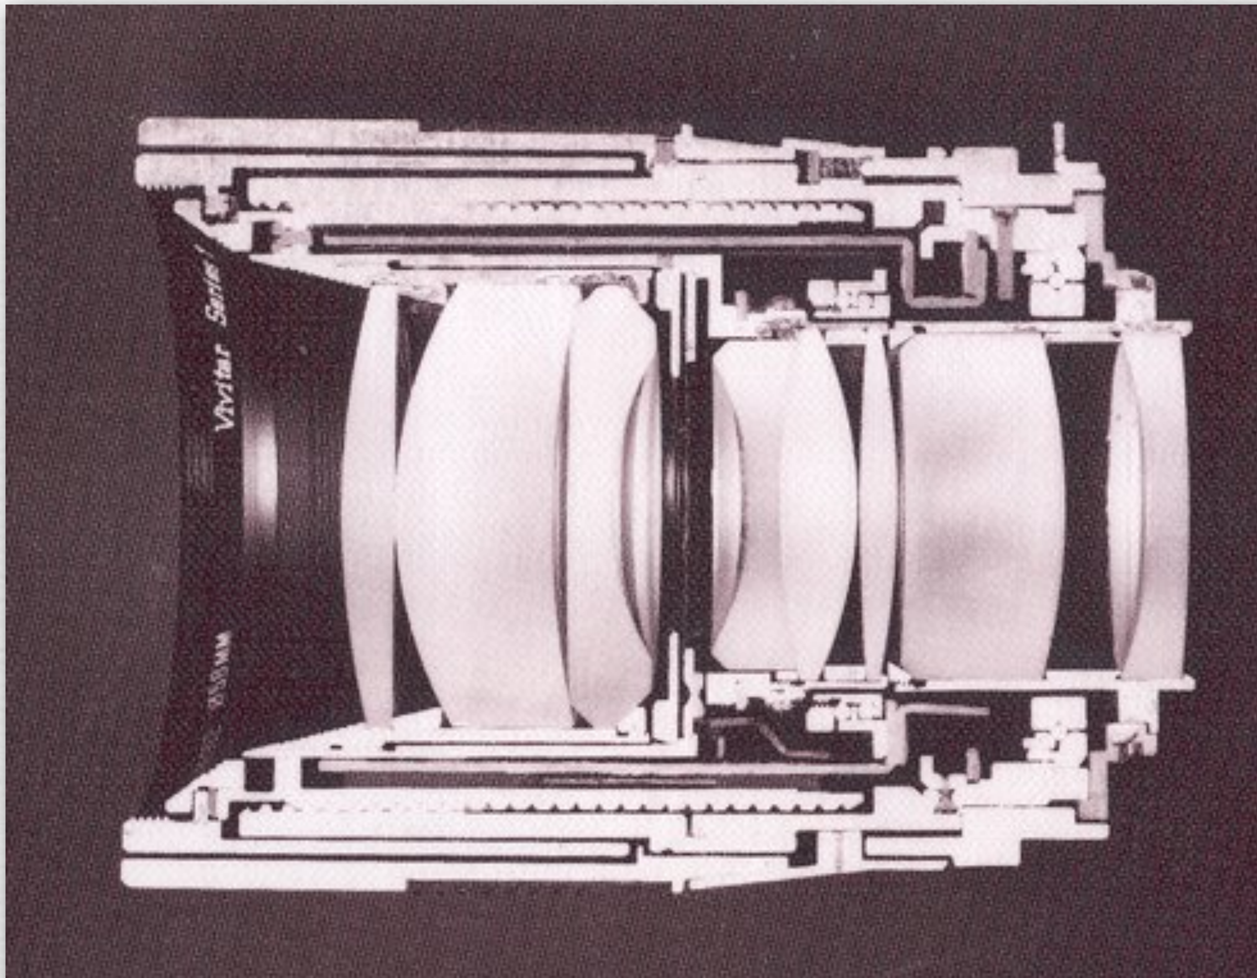

Vivitar Series 1 90mm f/2.5

Cover photo, Kingslake, *Optics in Photography*

Lens quality varies

- Why is this toy so expensive?
  - EF 70-200mm f/2.8L IS USM
  - $1700

- Why is it better than this toy?
  - EF 70-300mm f/4-5.6 IS USM
  - $550

- Why is it so complicated?

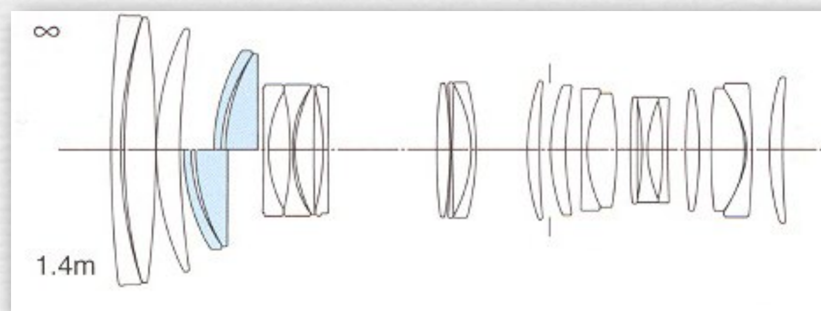

(Canon)

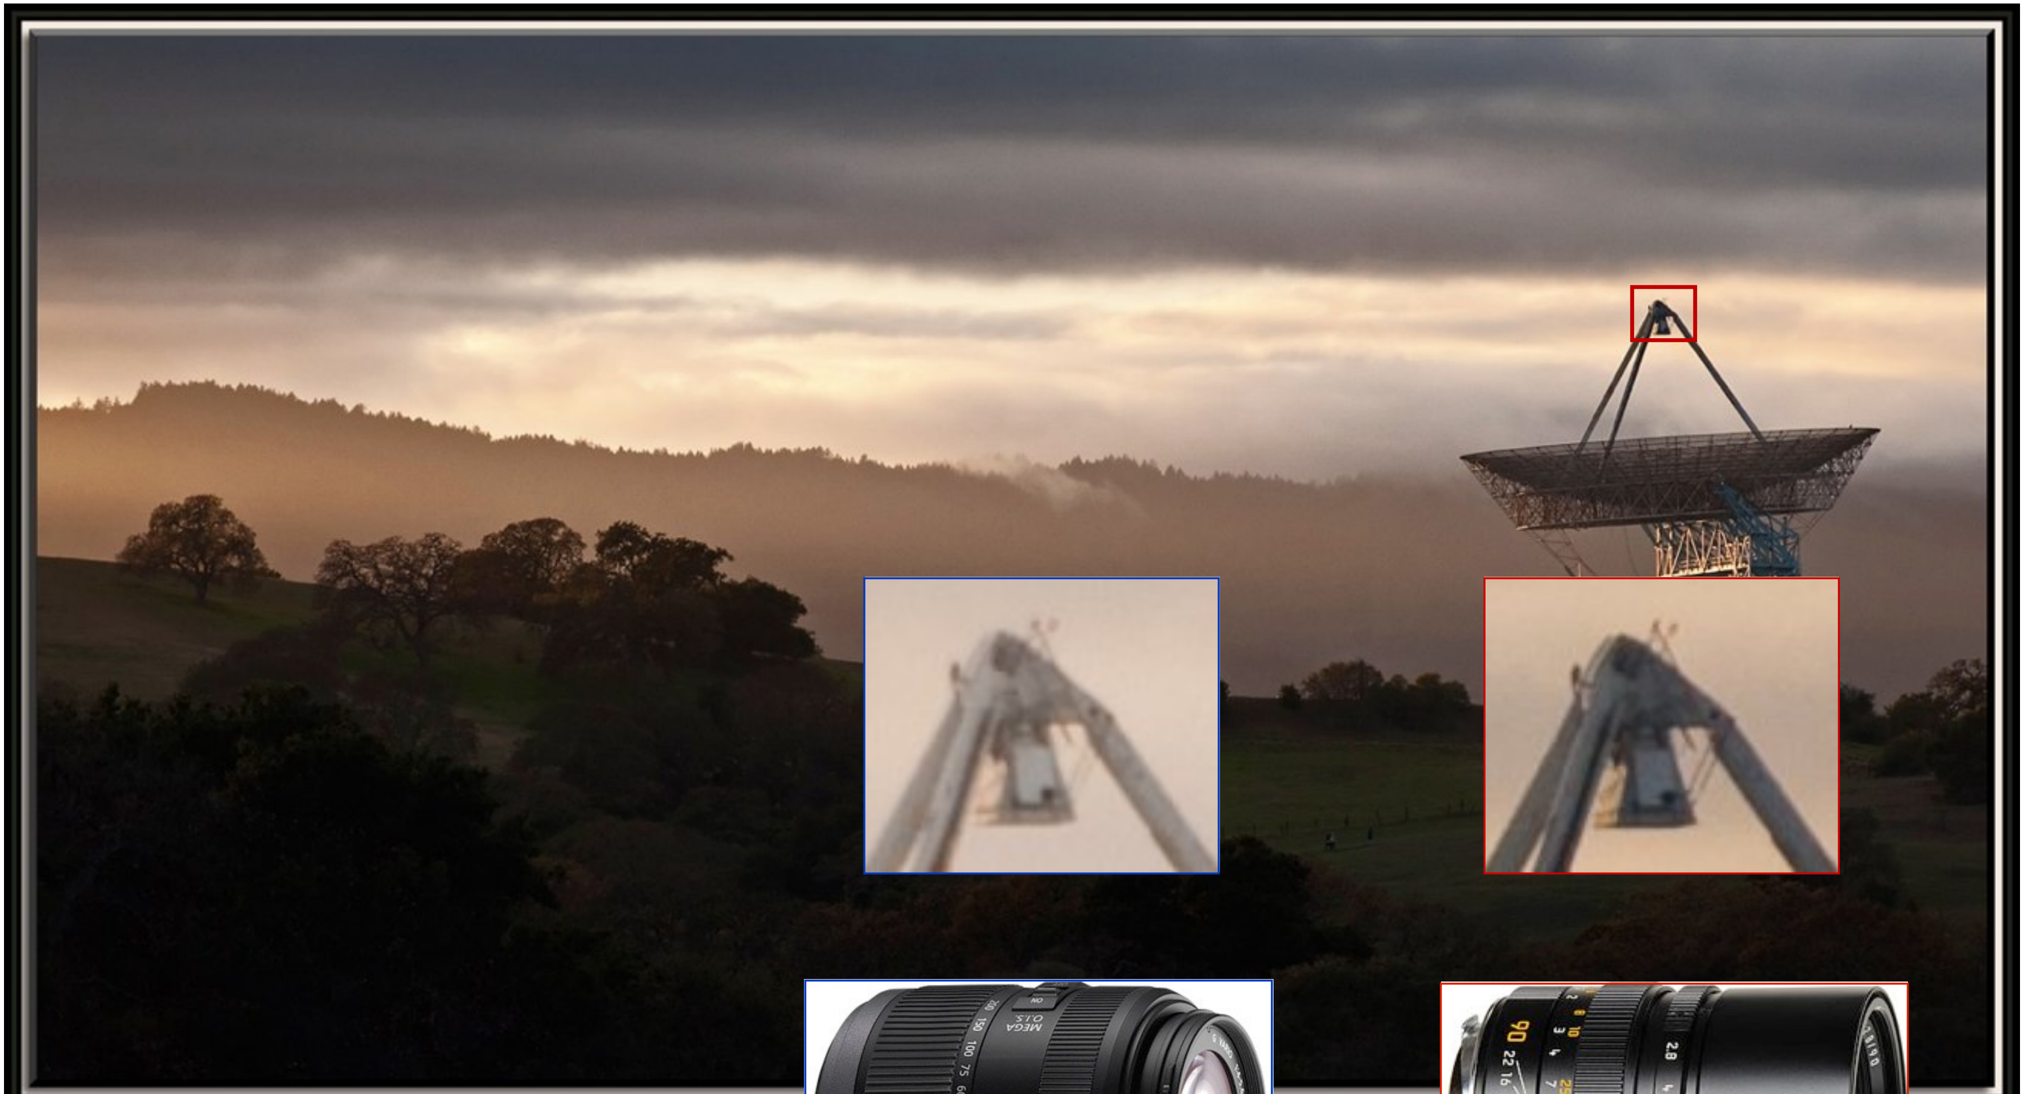

Panasonic 45-200/4-5.6 zoom, at 200mm f/4.6
$300

Leica 90mm/2.8 Elmarit-M prime, at f/4
$2000

Stanford Big Dish

Panasonic GF1

Zoom lens versus prime lens

Canon 100-400mm/4.5-5.6
zoom, at 300mm and f/5.6
$1600

Canon 300mm/2.8
prime, at f/5.6
$4300

Physical versus geometrical optics

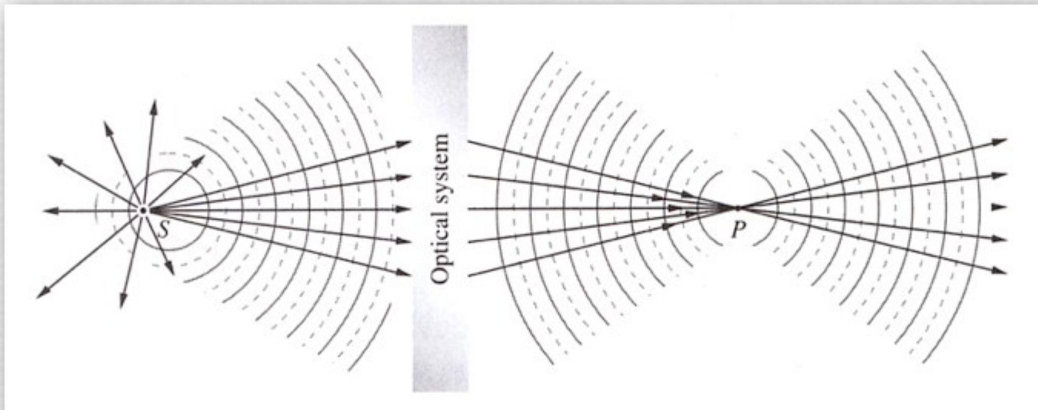

(Hecht)

- light can be modeled as traveling waves
- the perpendiculars to these waves can be drawn as rays
- diffraction causes these rays to bend, e.g. at a slit
- *geometrical optics* assumes
  - $\lambda \rightarrow 0$
  - no diffraction
  - in free space, rays are straight (a.k.a. rectilinear propagation)

Physical versus geometrical optics (contents of whiteboard)

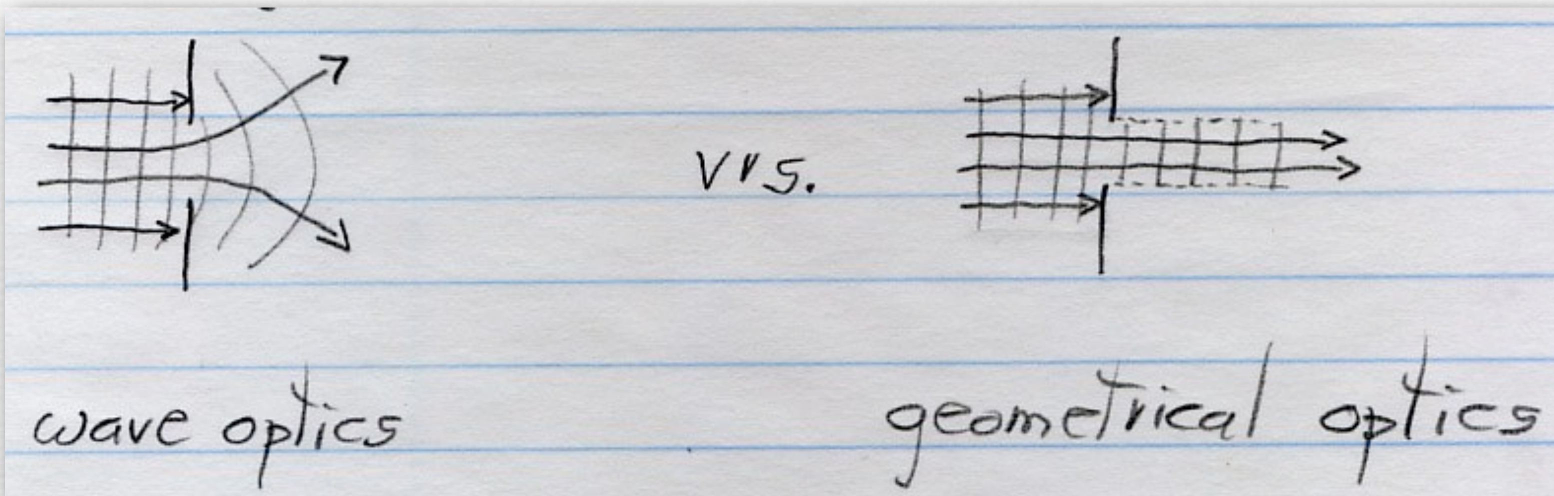

- in geometrical optics, we assume that rays do not bend as they pass through a narrow slit
- this assumption is valid if the slit is much larger than the wavelength
- physical optics is a.k.a. wave optics

Snell's law of refraction

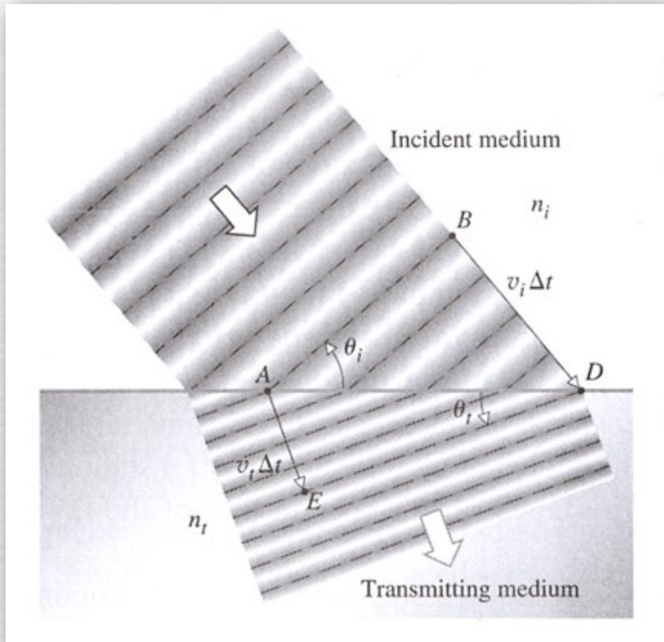

(Hecht)

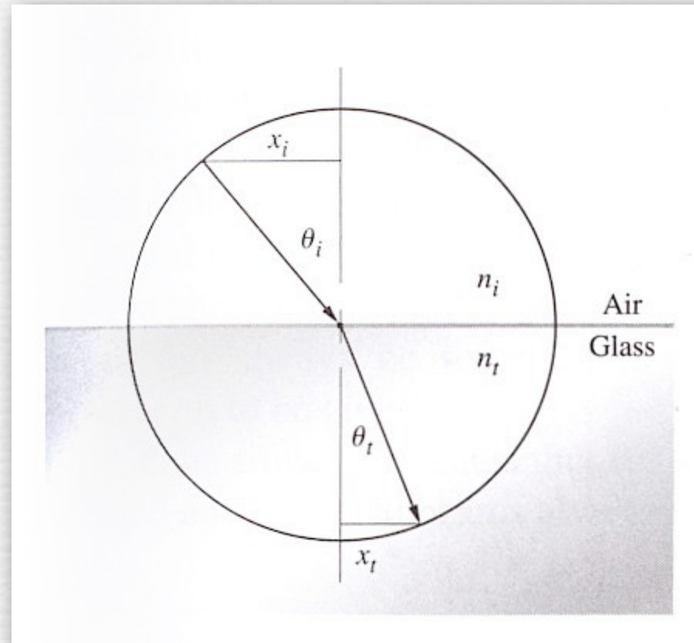

$$\frac{x_i}{x_t} = \frac{\sin\theta_i}{\sin\theta_t} = \frac{n_t}{n_i}$$

- as waves change speed at an interface, they also change direction
- index of refraction n_t is defined as $\dfrac{\text{speed of light in a vacuum}}{\text{speed of light in medium } t}$

Typical refractive indices (*n*)

- air = ~1.0
- water = 1.33
- glass = 1.5 - 1.8

mirage due to changes in the index of refraction of air with temperature

Refraction in glass lenses

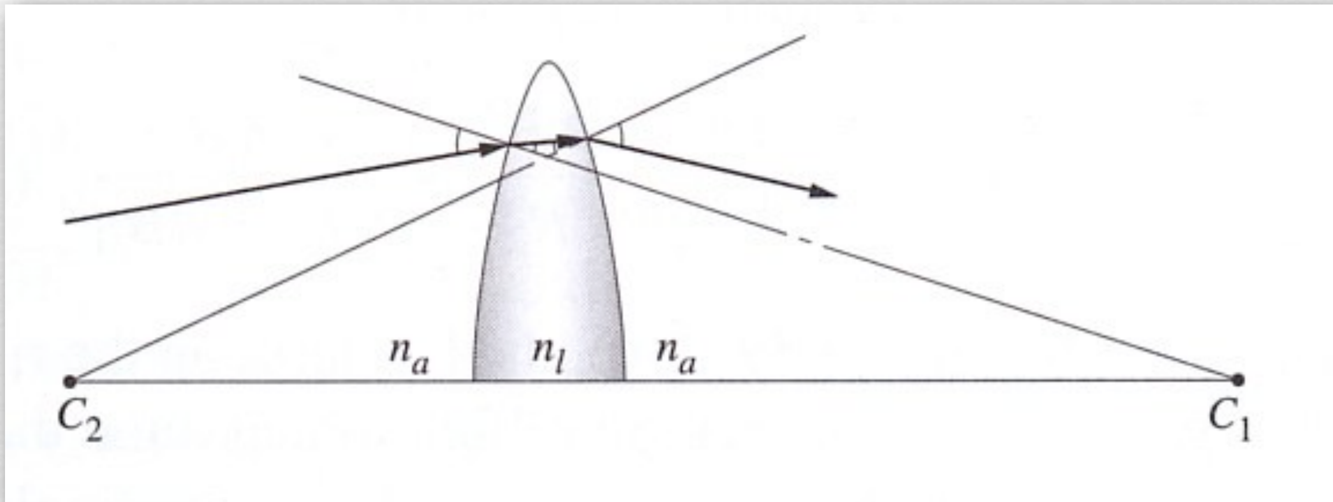

(Hecht)

- when transiting from air to glass, light bends towards the normal
- when transiting from glass to air, light bends away from the normal
- light striking a surface perpendicularly does not bend

Q. What shape should an interface be to make parallel rays converge to a point?

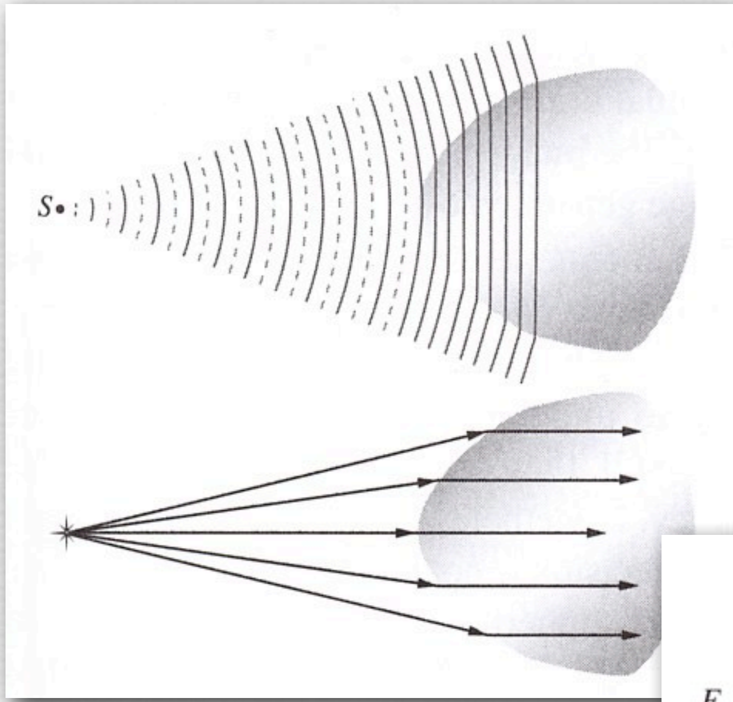

(Hecht)

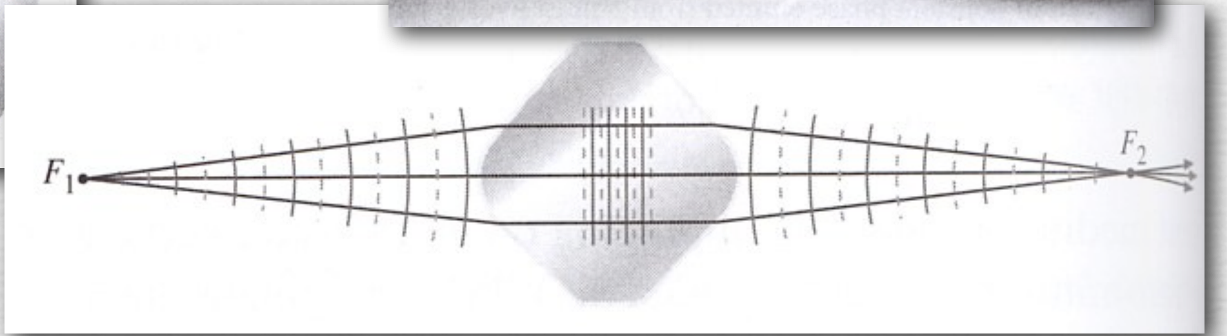

A. a hyperbola

- so lenses should be hyperbolic!

Spherical lenses

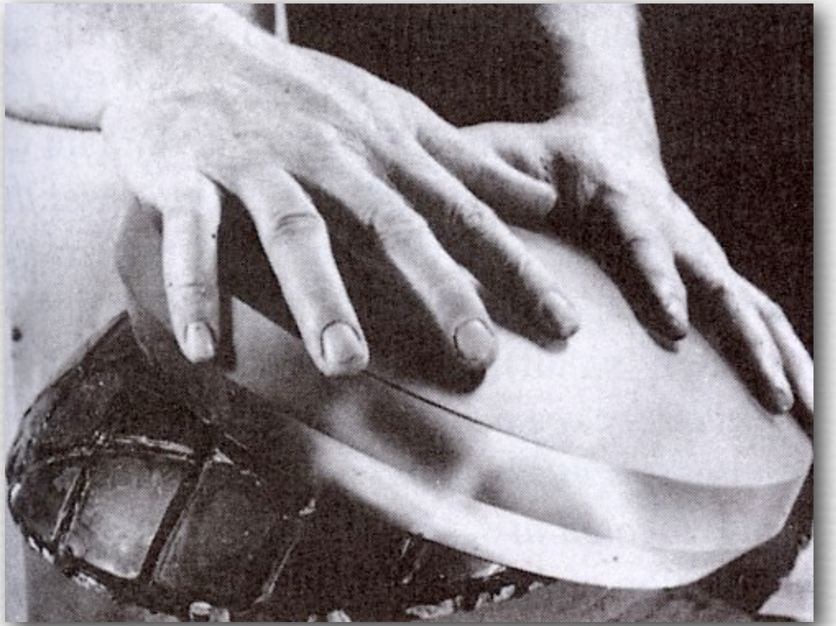

(Hecht)

hyperbolic lens

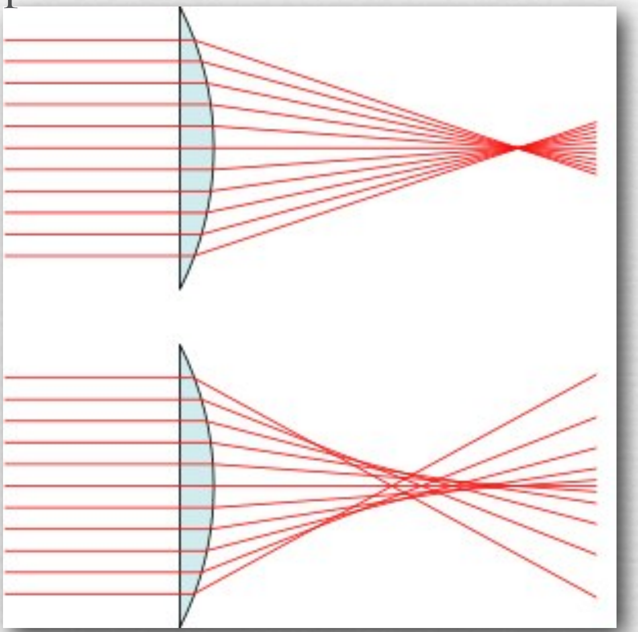

spherical lens

(wikipedia)

- two roughly fitting curved surfaces ground together will eventually become spherical
- spheres don't bring parallel rays to a point
  - this is called *spherical aberration*
  - nearly axial rays *(paraxial rays)* behave best

Examples of spherical aberration

(gtmerideth)

(Canon)

Canon 135mm
soft focus lens

Paraxial approximation

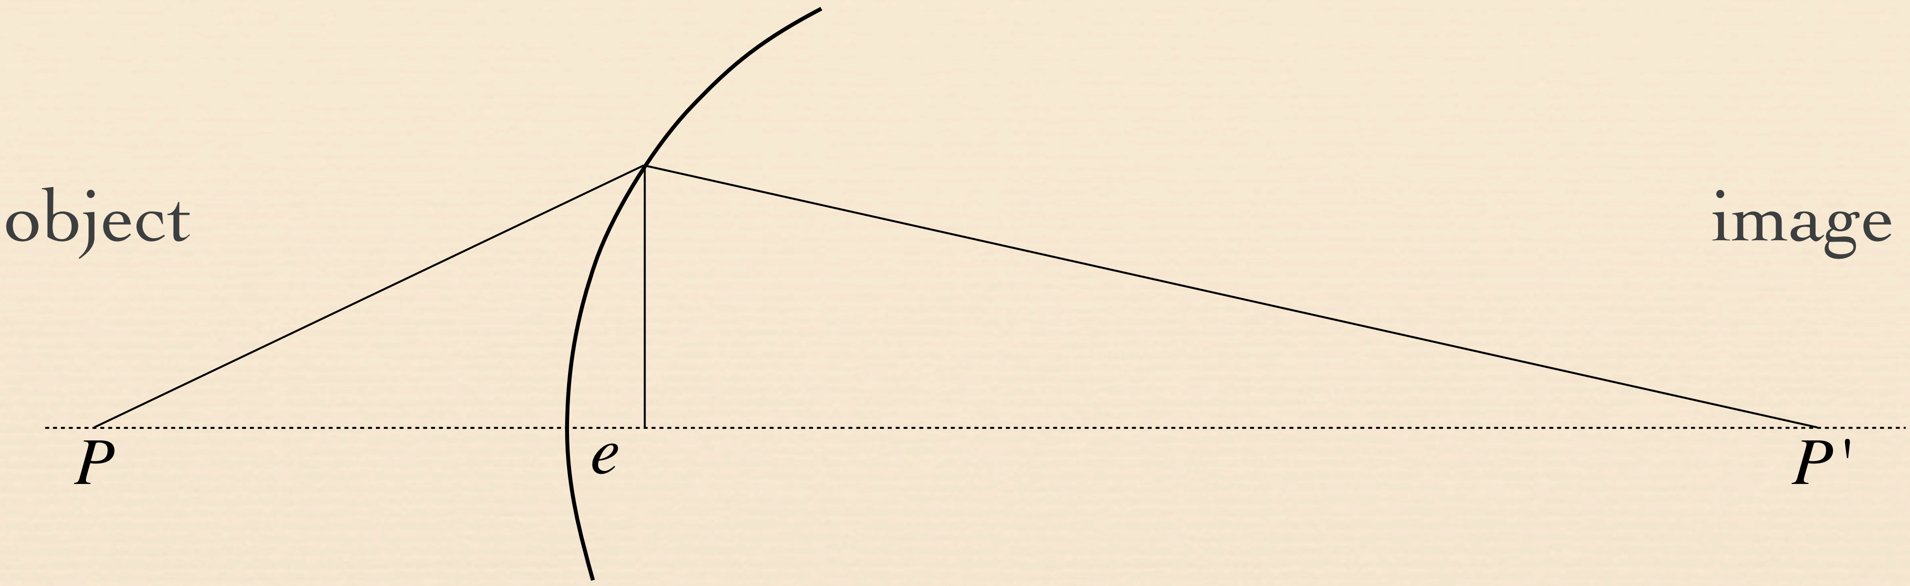

- assume $e \approx 0$

Not responsible on exams
for orange-tinted slides

Paraxial approximation

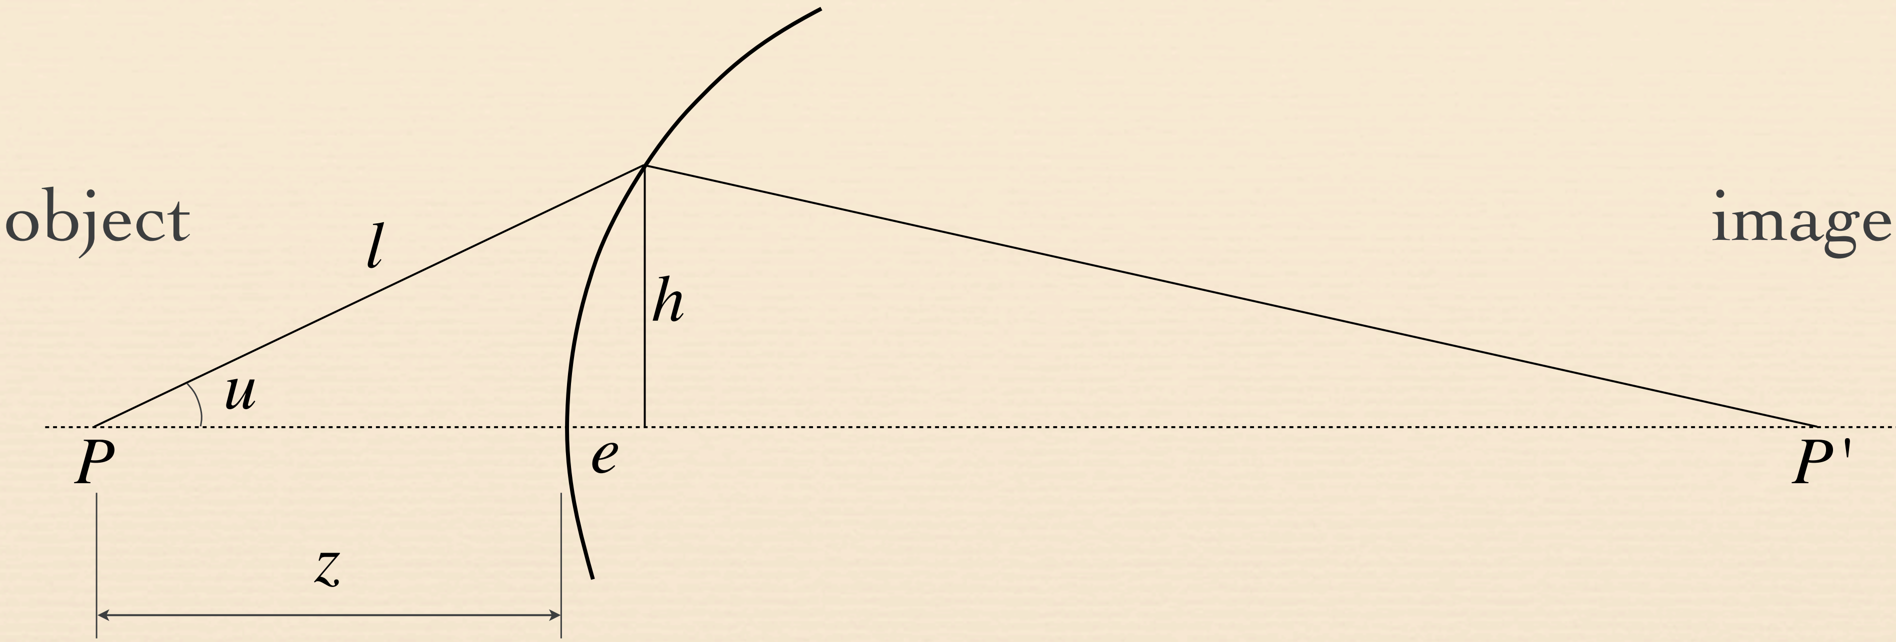

- assume $e \approx 0$
- assume $\sin u = h / l \approx u$ (for u in radians)
- assume $\cos u \approx z / l \approx 1$
- assume $\tan u \approx \sin u \approx u$

The paraxial approximation is a.k.a. first-order optics

- assume first term of $\sin\phi \quad = \quad \phi - \frac{\phi^3}{3!} + \frac{\phi^5}{5!} - \frac{\phi^7}{7!} + \ldots$
  - i.e. $\sin\phi \approx \phi$
- assume first term of $\cos\phi \quad = \quad 1 - \frac{\phi^2}{2!} + \frac{\phi^4}{4!} - \frac{\phi^6}{6!} + \ldots$
  - i.e. $\cos\phi \approx 1$
  - so $\tan\phi \approx \sin\phi \approx \phi$

these are the Taylor series for $\sin\phi$ and $\cos\phi$

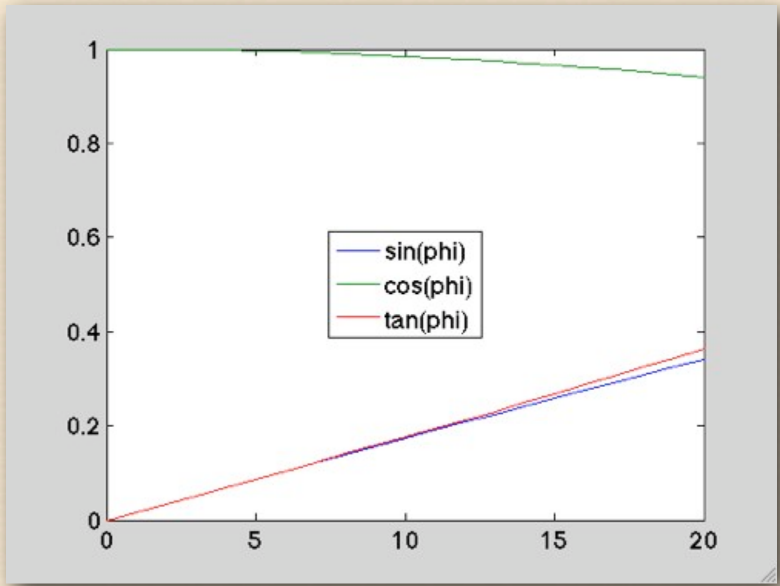

(phi in degrees)

Paraxial focusing

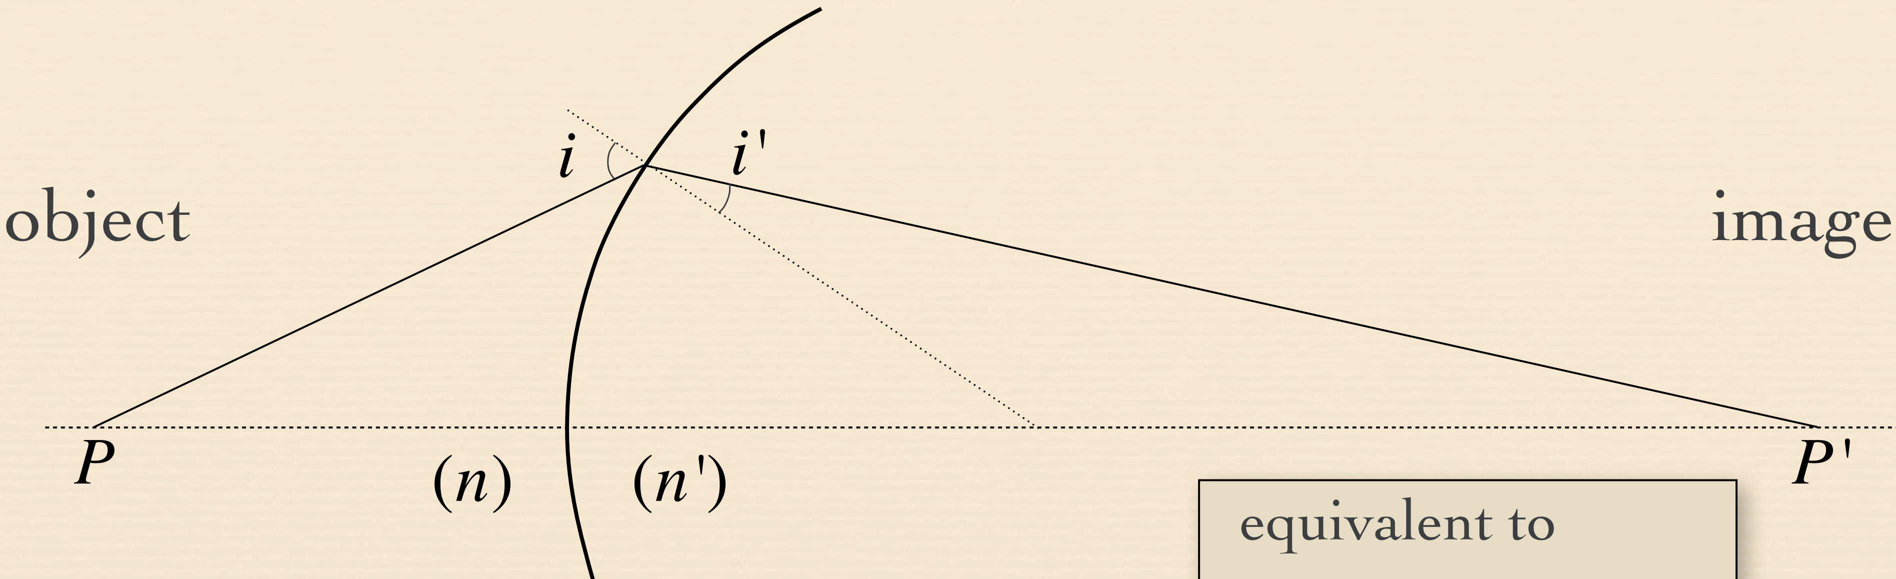

Snell's law:

$$n \sin i = n' \sin i'$$

paraxial approximation:

$$n\, i \approx n'\, i'$$

equivalent to

$$\frac{\sin\theta_i}{\sin\theta_t} = \frac{n_t}{n_i}$$

with

$n = n_i$ for air

$n' = n_t$ for glass

i, i' in radians

$\theta_i,\ \theta_t$ in degrees

Paraxial focusing

$i = u + a$

$u \approx h / z$

$u' \approx h / z'$

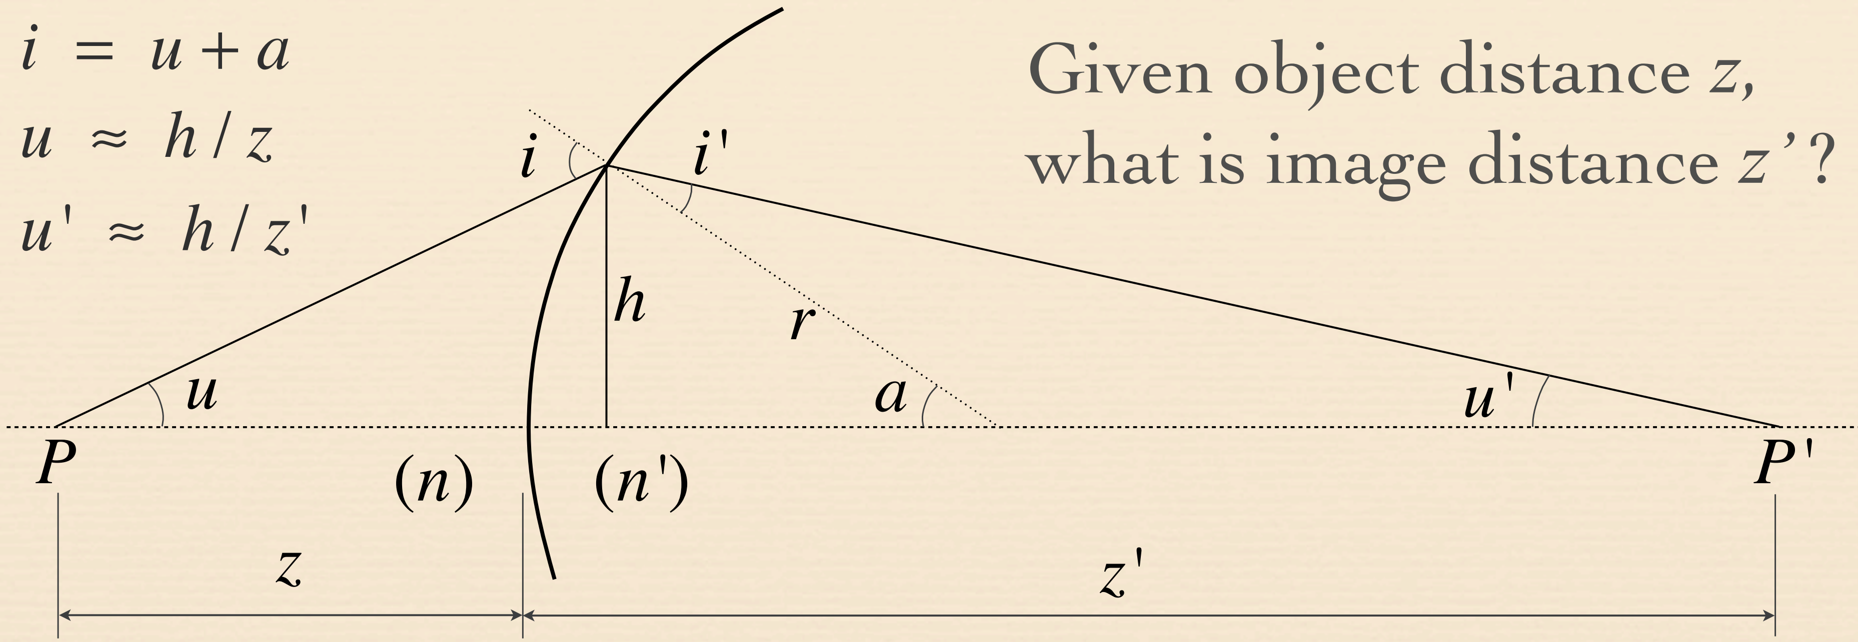

Given object distance z, what is image distance z'?

$n\,i \approx n'\,i'$

Paraxial focusing

$$i = u + a$$

$$u \approx h / z$$

$$u' \approx h / z'$$

$$a = u' + i'$$

$$a \approx h / r$$

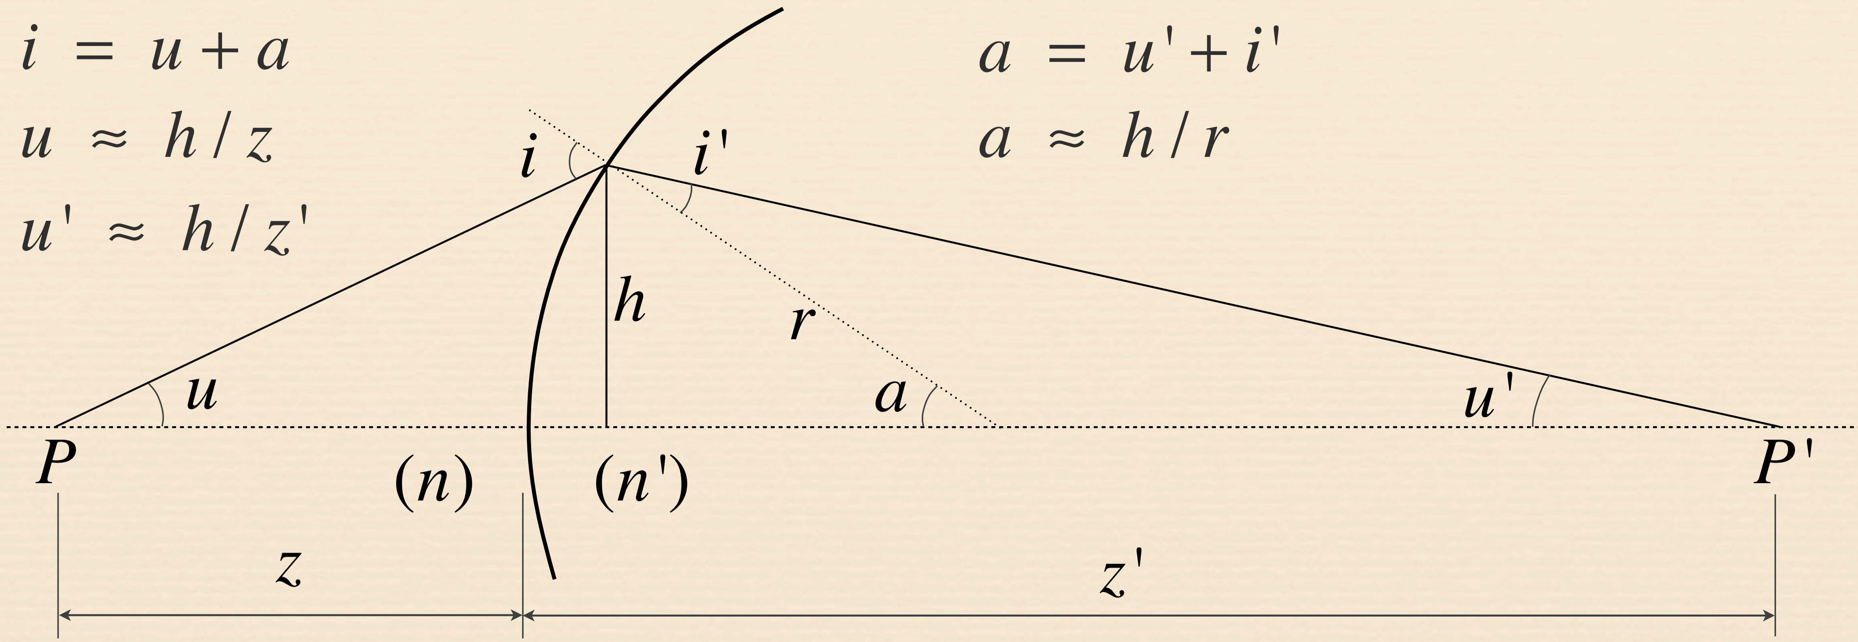

$$n(u + a) \approx n'(a - u')$$

$$n(h/z + h/r) \approx n'(h/r - h/z')$$

$$n/z + n/r \approx n'/r - n'/z'$$

$$n\, i \approx n'\, i'$$

- *h* has canceled out, so any ray from *P* will focus to *P'*

Focal length

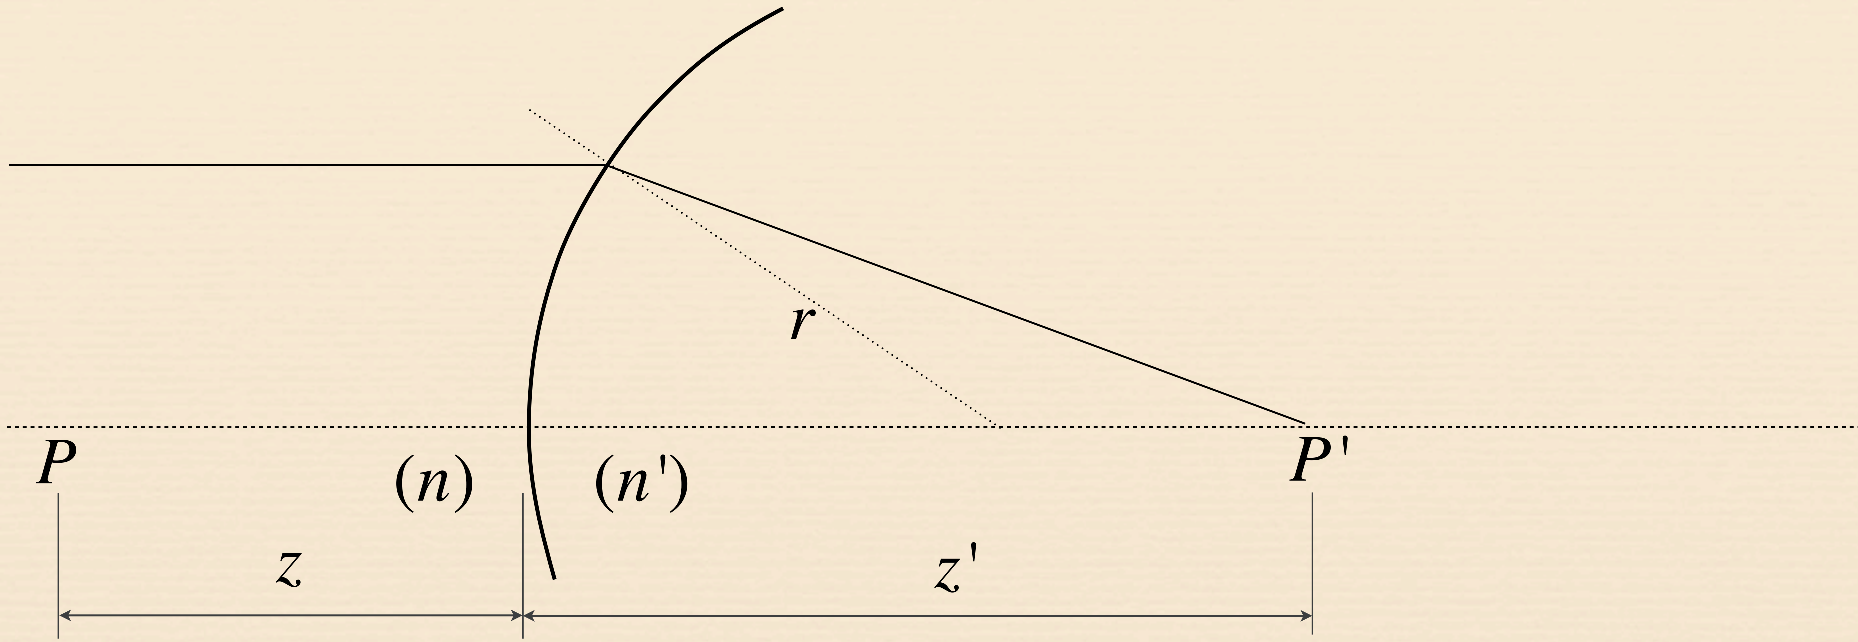

What happens if z is ∞ ?

$$n / z + n / r \approx n' / r - n' / z'$$

$$n / r \approx n' / r - n' / z'$$

$$z' \approx (r\, n') / (n' - n)$$

- $f \triangleq$ focal length = z'

Lensmaker's formula

- using similar derivations, one can extend these results to two spherical interfaces forming a lens in air

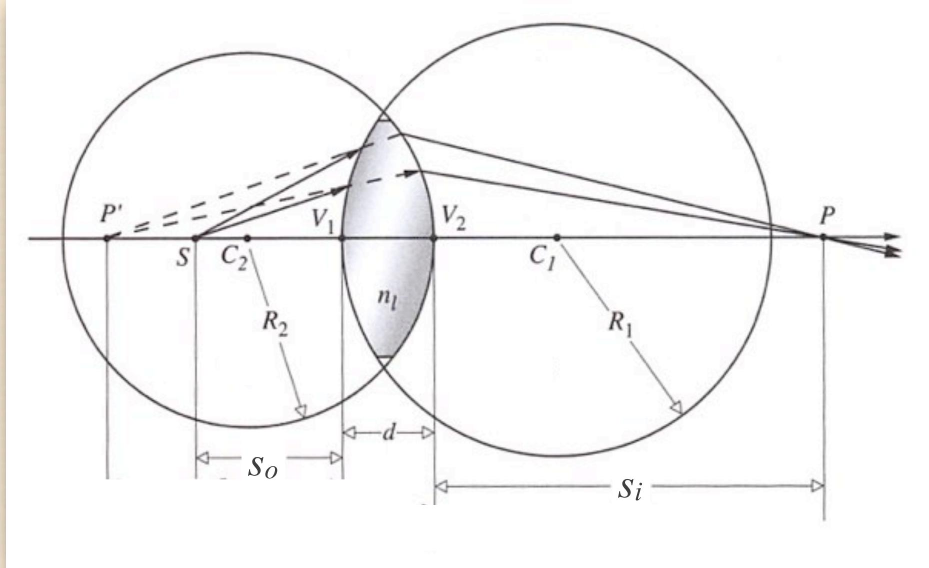

(Hecht, edited)

- as $d \rightarrow 0$ *(thin lens approximation)*, we obtain the lensmaker's formula

$$\frac{1}{s_o} + \frac{1}{s_i} = (n_l - 1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right)$$

Gaussian lens formula

- Starting from the lensmaker's formula

$$\frac{1}{s_o} + \frac{1}{s_i} = (n_l - 1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right),$$ (Hecht, eqn 5.15)

- and recalling that as object distance s_o is moved to infinity, image distance s_i becomes focal length f_i , we get

$$\frac{1}{f_i} = (n_l - 1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right).$$ (Hecht, eqn 5.16)

- Equating these two, we get the Gaussian lens formula

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f_i}.$$ (Hecht, eqn 5.17)

From Gauss's ray construction to the Gaussian lens formula

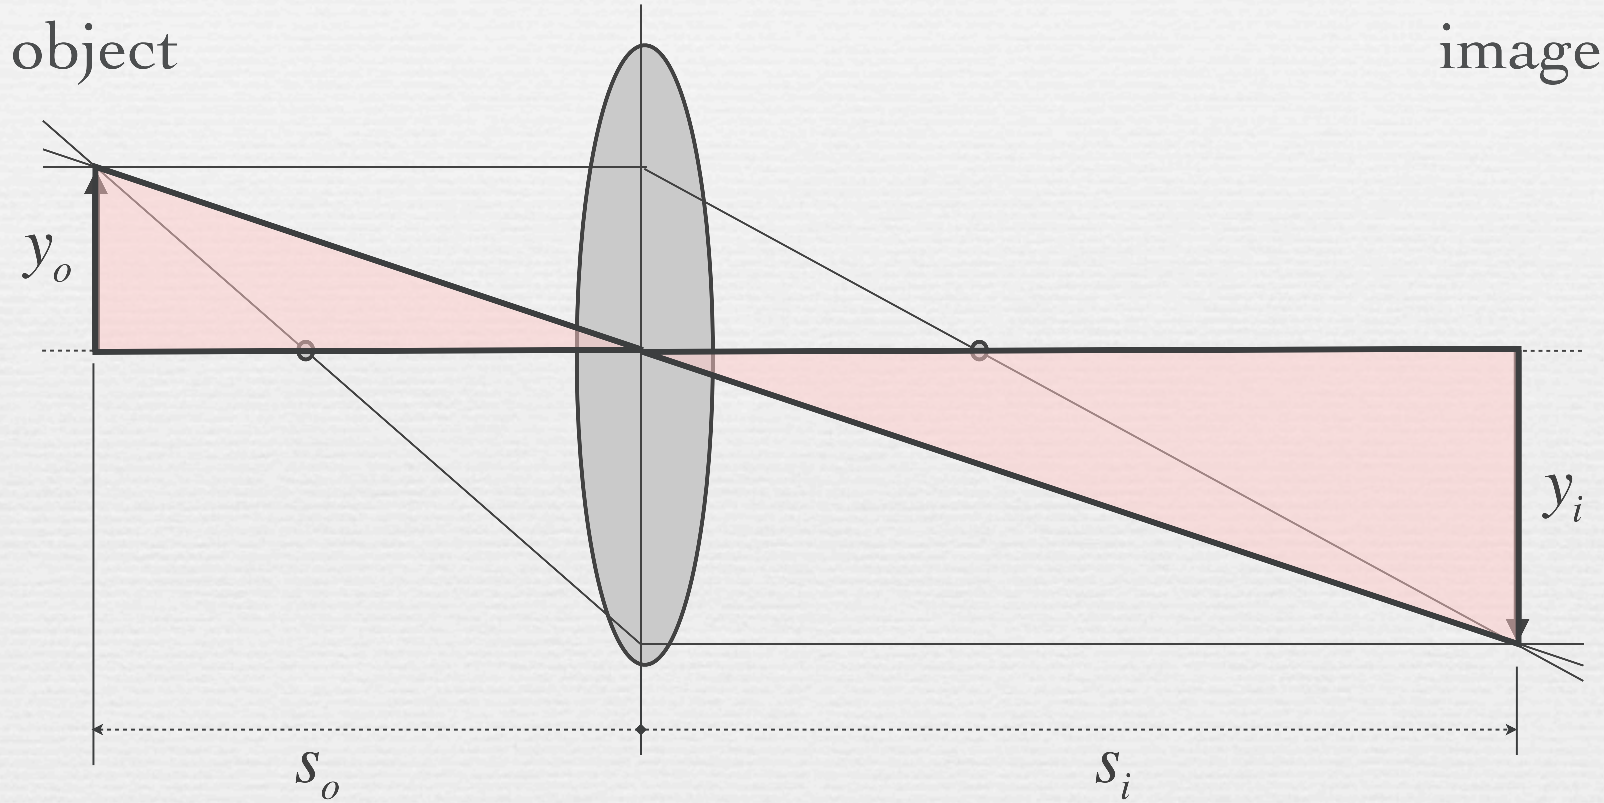

- positive s_i is rightward, positive s_o is leftward
- positive y is upward

From Gauss's ray construction to the Gaussian lens formula

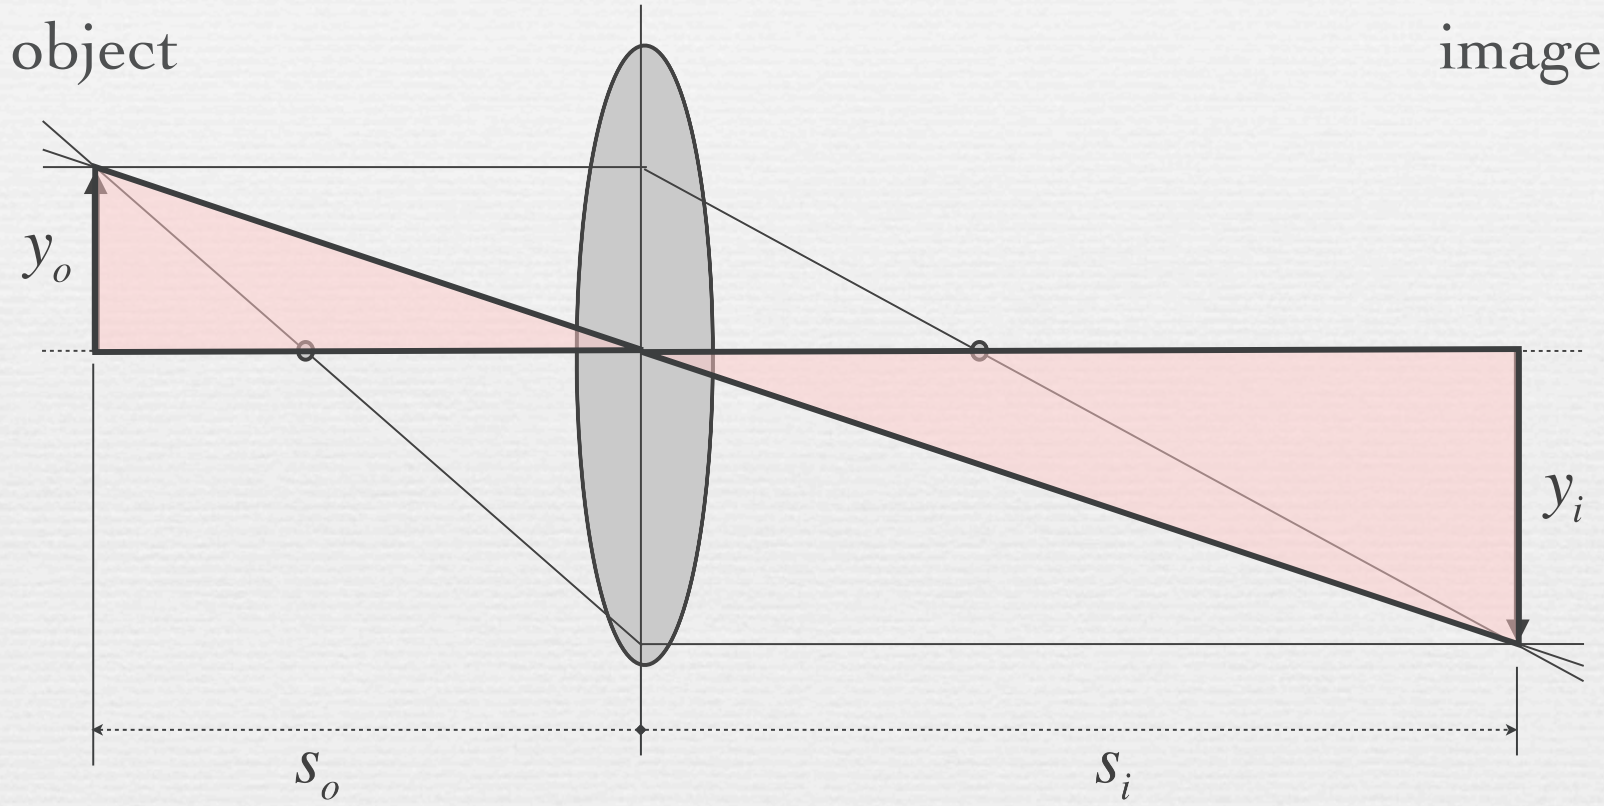

$$\frac{|y_i|}{y_o} = \frac{s_i}{s_o}$$

From Gauss's ray construction to the Gaussian lens formula

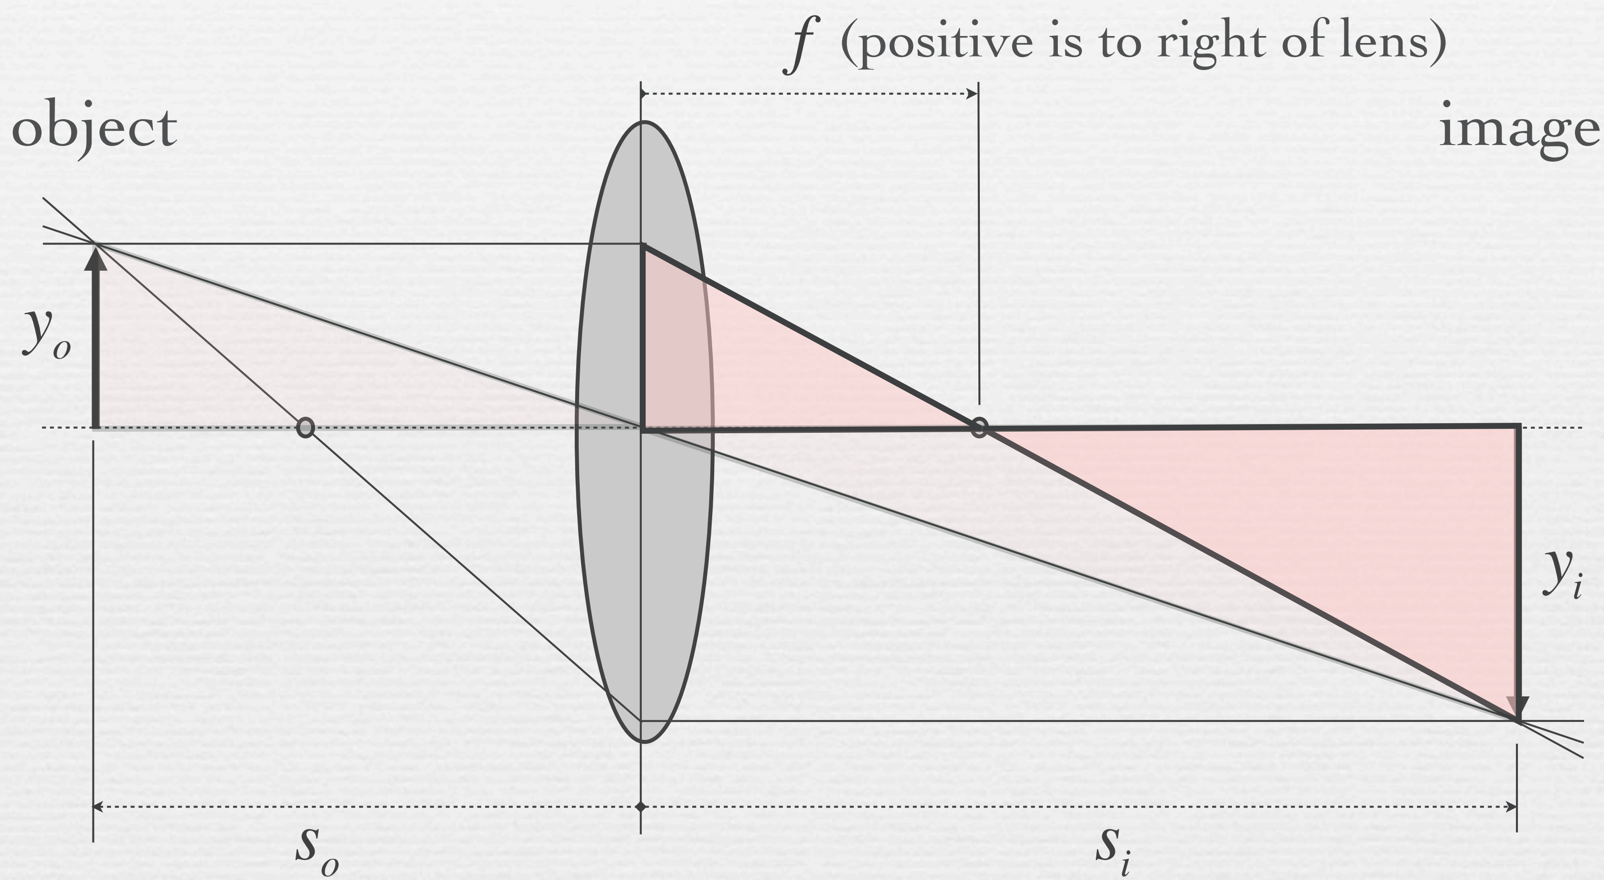

$$\frac{|y_i|}{y_o} = \frac{s_i}{s_o} \quad \text{and} \quad \frac{|y_i|}{y_o} = \frac{s_i - f}{f} \quad \ldots\ldots \quad \boxed{\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}}$$

Changing the focus distance

- to focus on objects at different distances, move sensor relative to lens

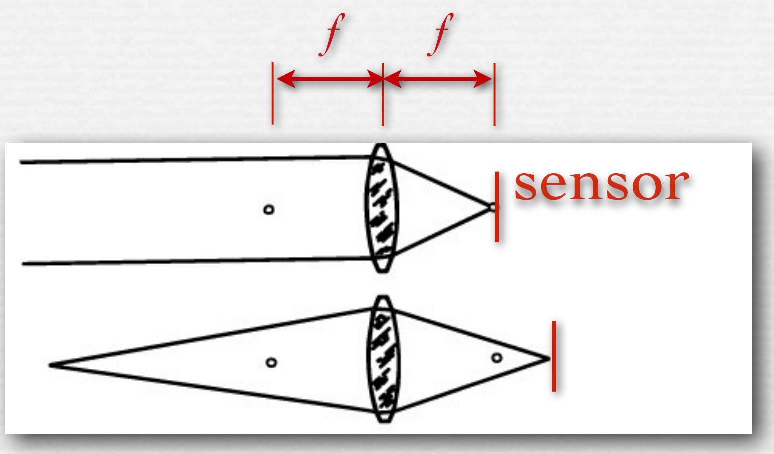

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/gaussian.html

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}$$

Changing the focus distance

- to focus on objects at different distances, move sensor relative to lens

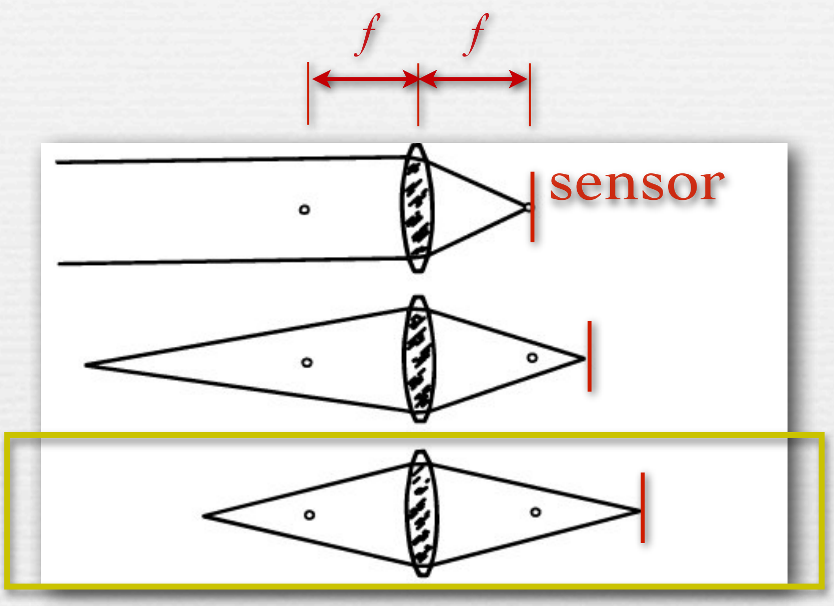

- at $s_o = s_i = 2f$ we have 1:1 imaging, because

$$\frac{1}{2f} + \frac{1}{2f} = \frac{1}{f}$$

In 1:1 imaging, if the sensor is 36mm wide, an object 36mm wide will fill the frame.

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}$$

Changing the focus distance

- to focus on objects at different distances, move sensor relative to lens

- at $s_o = s_i = 2f$ we have 1:1 imaging, because

$$\frac{1}{2f} + \frac{1}{2f} = \frac{1}{f}$$

- can't focus on objects closer to lens than its focal length f

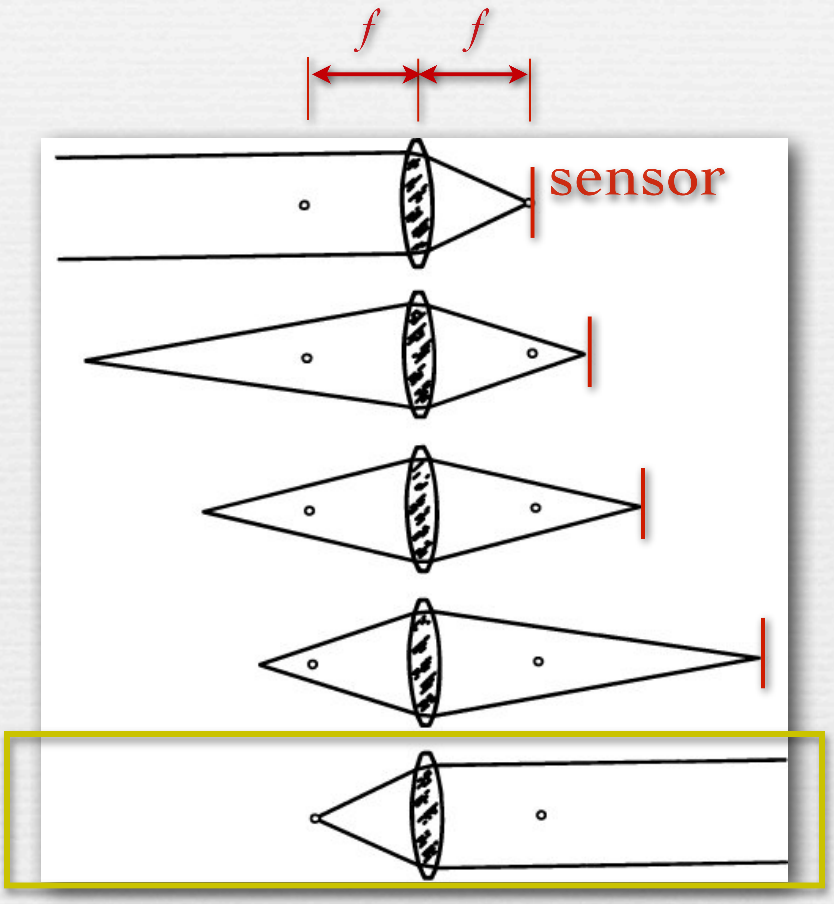

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}$$

Convex versus concave lenses

(Hecht)

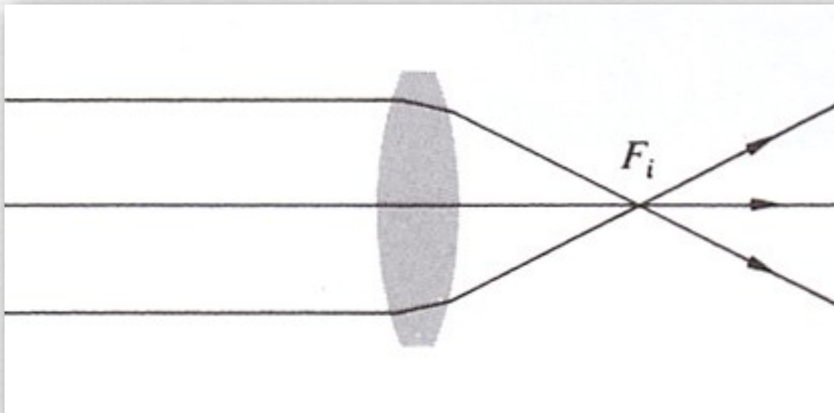

rays from a convex lens converge

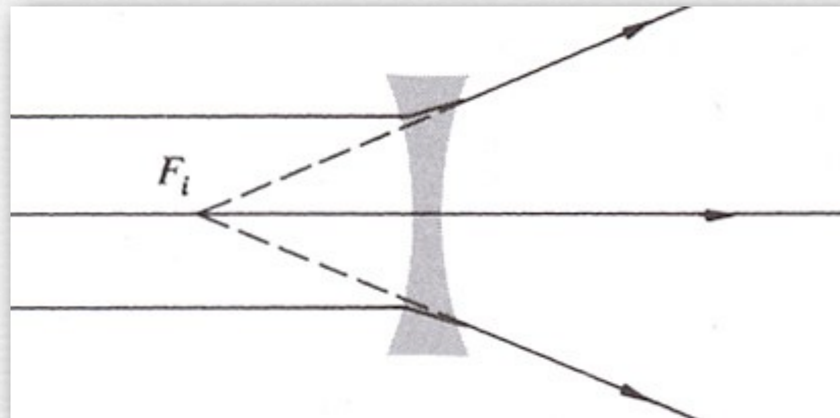

rays from a concave lens diverge

- positive focal length *f* means parallel rays from the left converge to a point on the right
- negative focal length *f* means parallel rays from the left converge to a point on the left (dashed lines above)

Convex versus concave lenses

(Hecht)

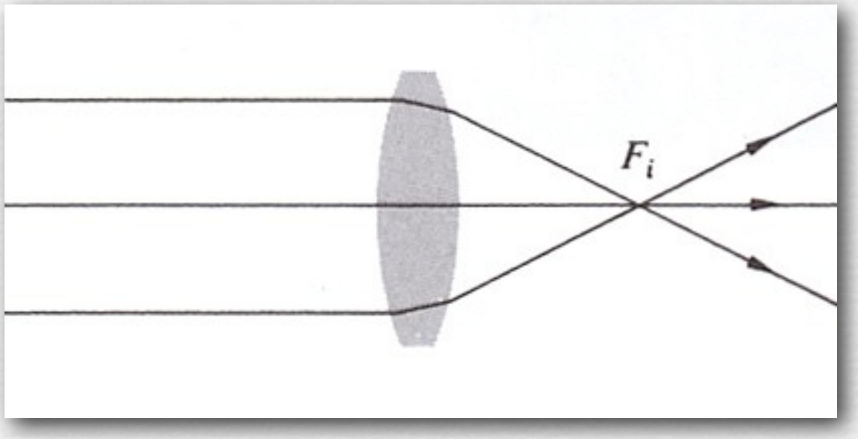

rays from a convex lens converge

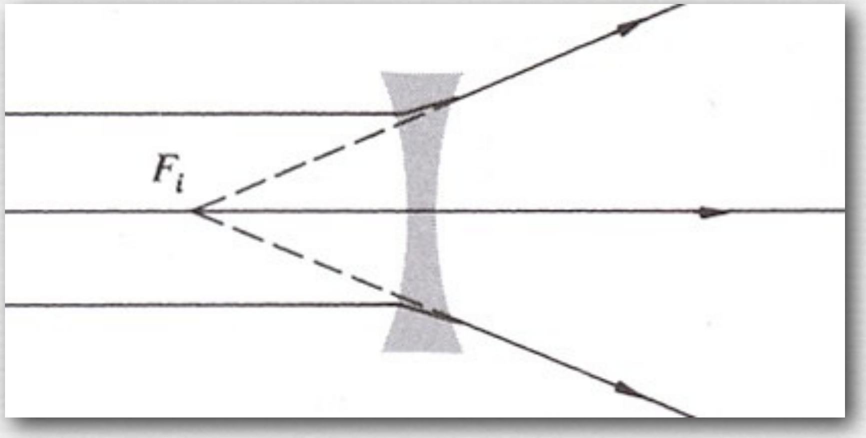

rays from a concave lens diverge

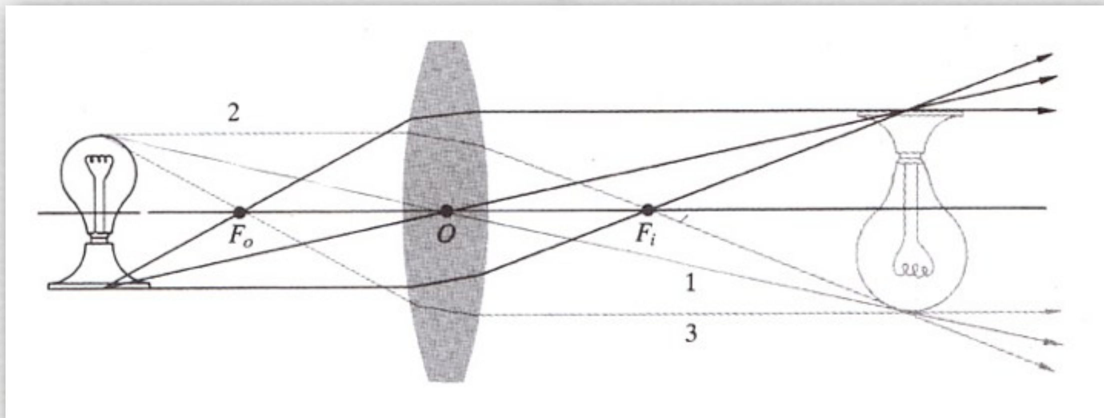

...producing a real image

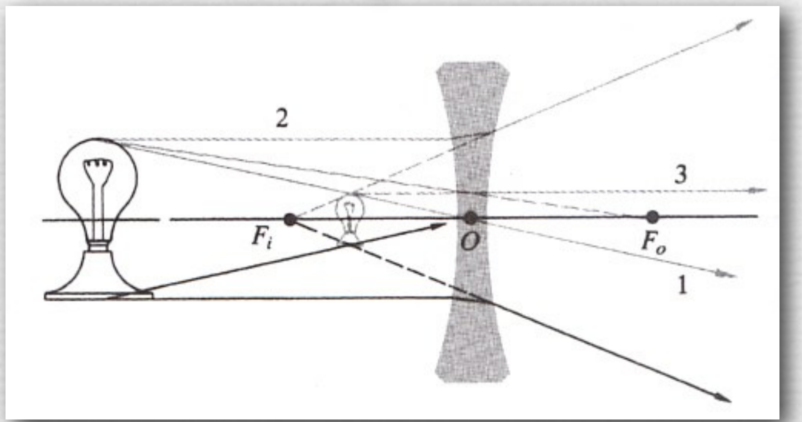

...producing a virtual image

Convex versus concave lenses

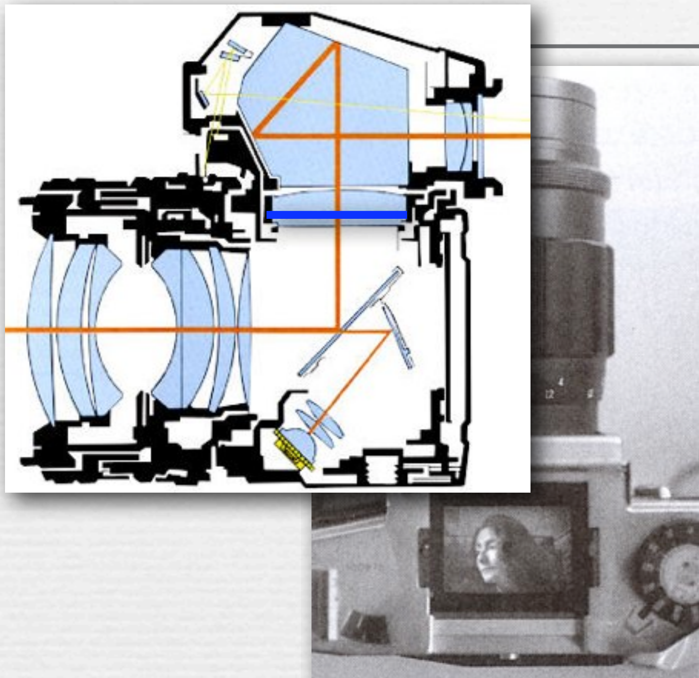

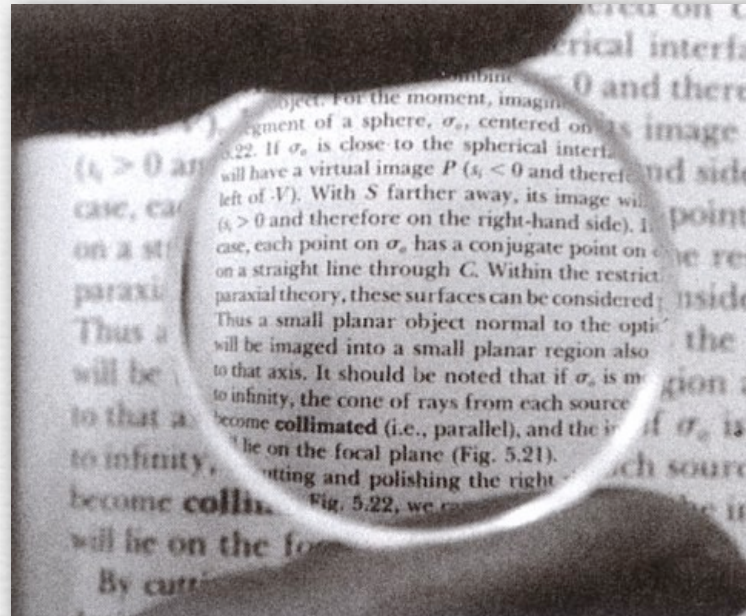

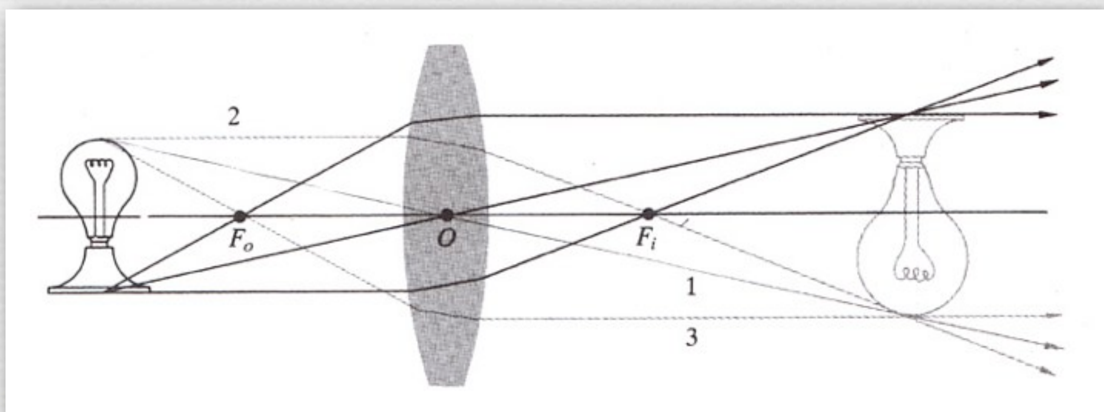

...producing a real image

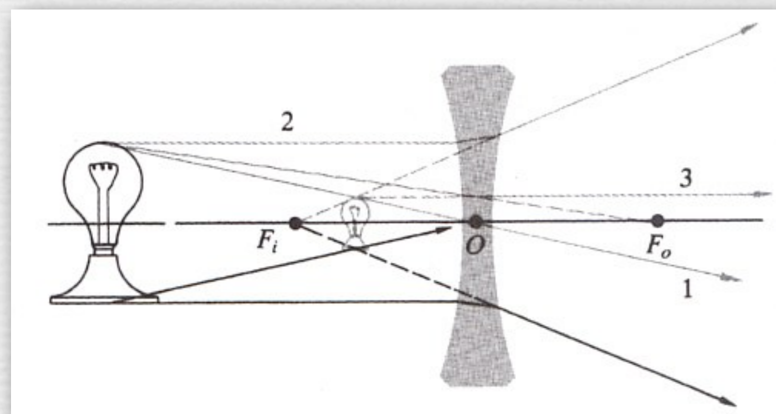

...producing a virtual image

The power of a lens

$$P = \frac{1}{f}$$

units are meters^{-1}

a.k.a. diopters

- my eyeglasses have the prescription
  - right eye: -0.75 diopters
  - left eye: -1.00 diopters

Q. What's wrong with me?
A. Myopia (nearsightedness)

(Pamplona)

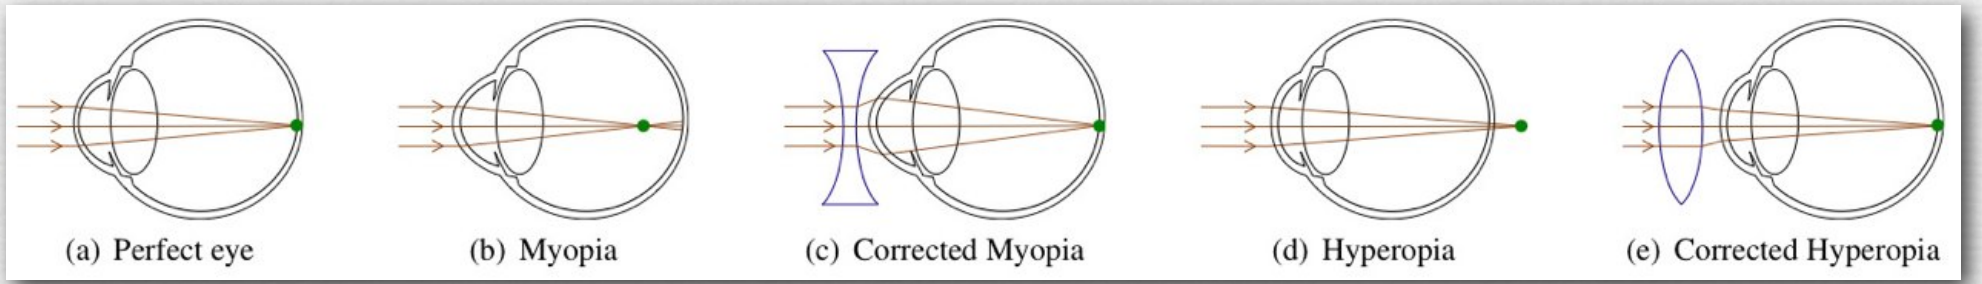

Combining two lenses

- using focal lengths

$$\frac{1}{f_{tot}} = \frac{1}{f_1} + \frac{1}{f_2}$$

- using diopters

$$P_{tot} = P_1 + P_2$$

- example

$$\frac{1}{70mm} + \frac{1}{500mm} = \frac{1}{61.4mm} \quad \text{-or-} \quad 14.28 + 2.0 = 16.28$$

Close-up filters

1/500mm = +2 diopters

(wikipedia)

- screw on to end of lens
- power is designated in diopters (usually)

Close-up filters

Canon 70-300mm

+

- changes minimum focal length from 70mm to 61.4mm

$$\frac{1}{70mm} + \frac{1}{500mm} = \frac{1}{61.4mm} \quad \text{-or-} \quad 14.28 + 2.0 = 16.28$$

poor man's macro lens

- for a fixed image distance, it reduces the object distance

Magnification

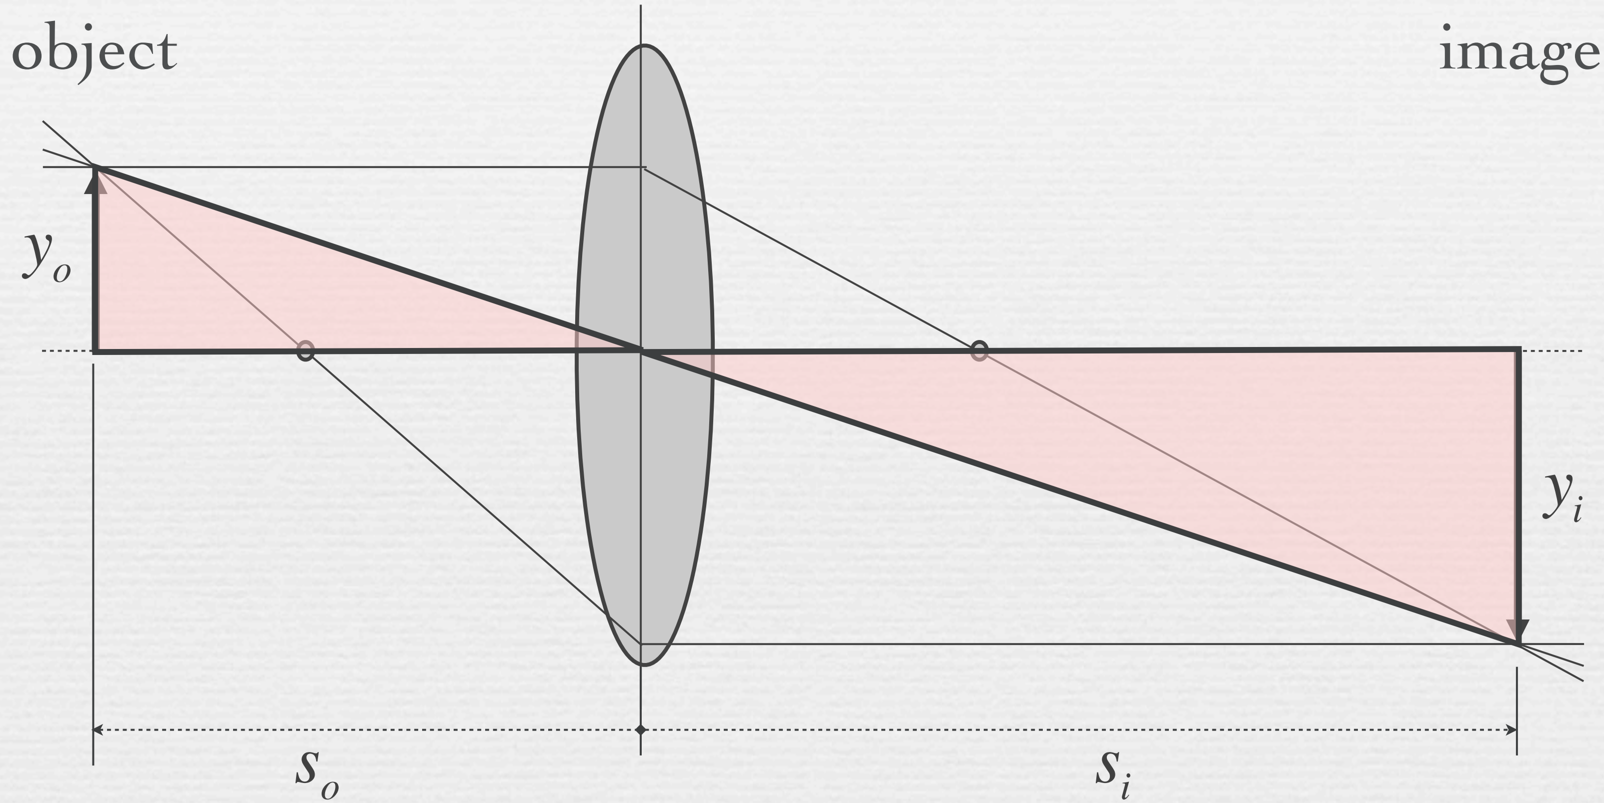

$$M_T \triangleq \frac{y_i}{y_o} = -\frac{s_i}{s_o}$$

Thick lenses

- an optical system may contain many lenses, but can be characterized by a few numbers

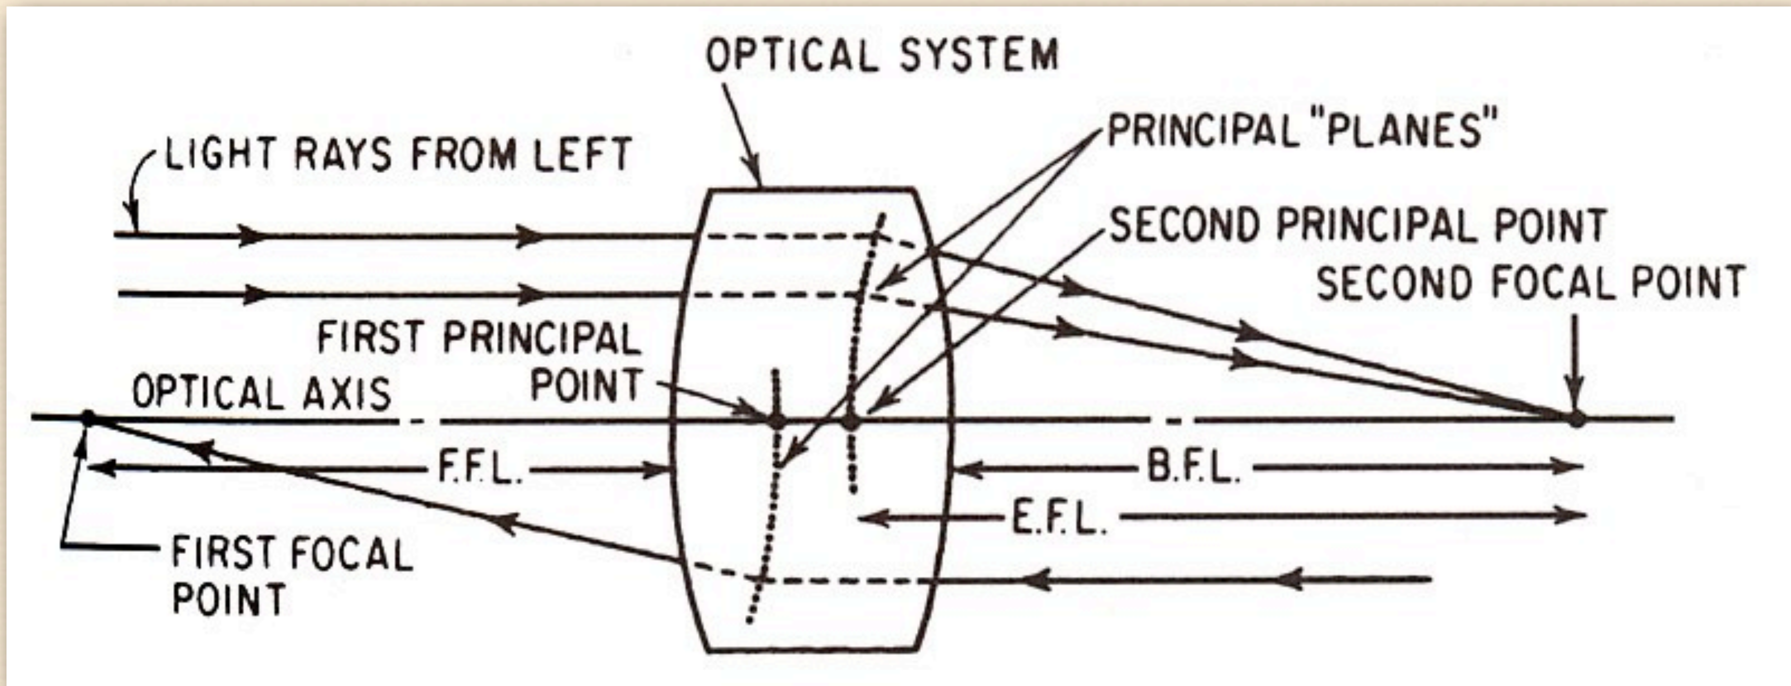

(Smith)

Center of perspective

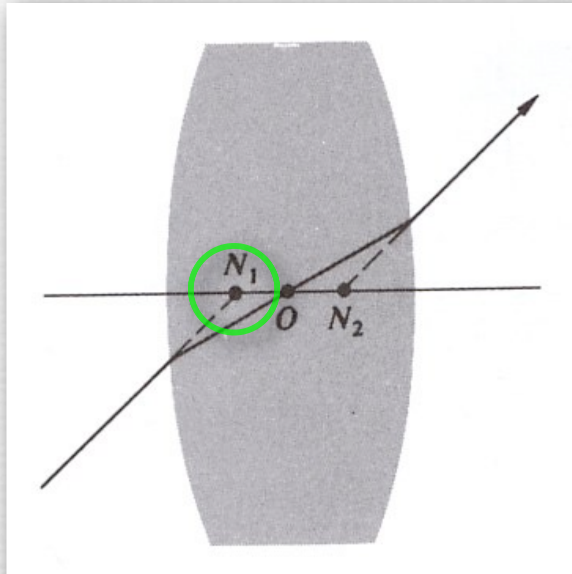

(Hecht)

- in a thin lens, the *chief ray* traverses the lens (through its optical center) without changing direction
- in a thick lens, the intersections of this ray with the optical axis are called the *nodal points*
- for a lens in air, these coincide with the *principal points*
- the first nodal point is the *center of perspective*

Lenses perform a 3D perspective transform

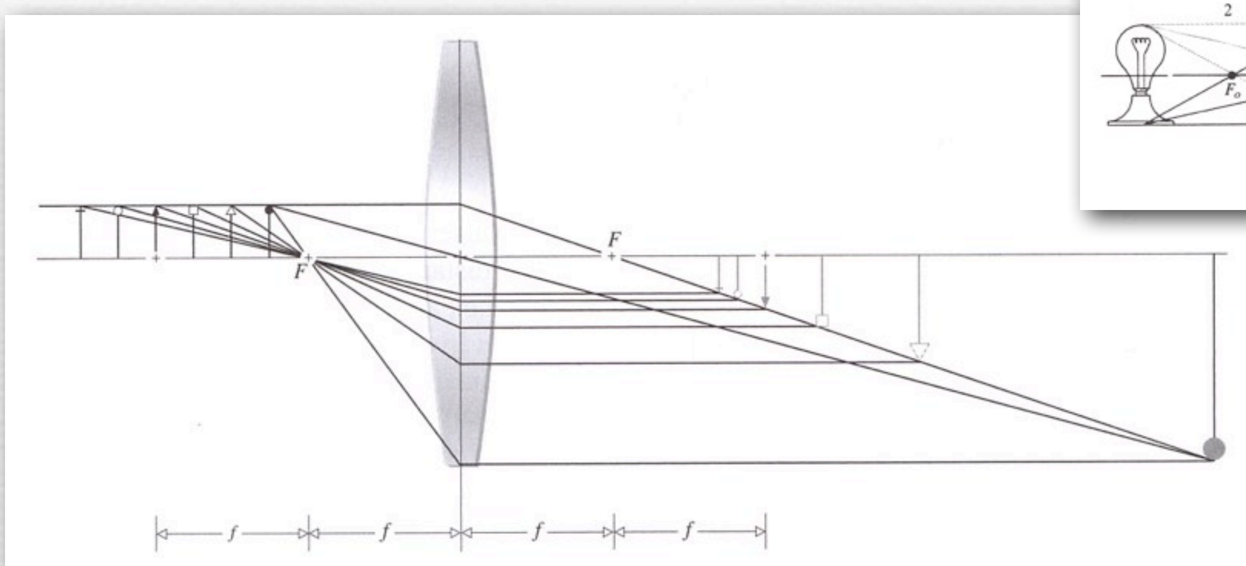

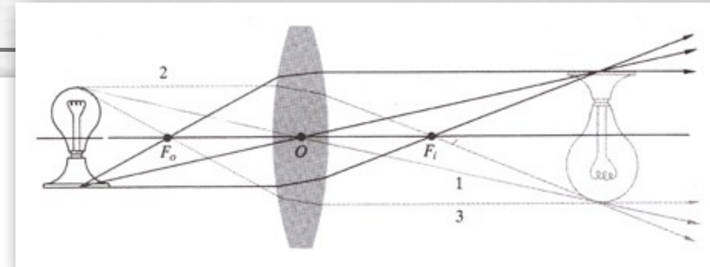

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/thinlens.html

(Hecht)

- lenses transform a 3D object to a 3D image; the sensor extracts a 2D slice from that image
- as an object moves linearly (in Z), its image moves non-proportionately (in Z)
- as you move a lens linearly relative to the sensor, the in-focus object plane moves non-proportionately
- as you refocus a camera, the image changes size !

Lenses perform a 3D perspective transform (contents of whiteboard)

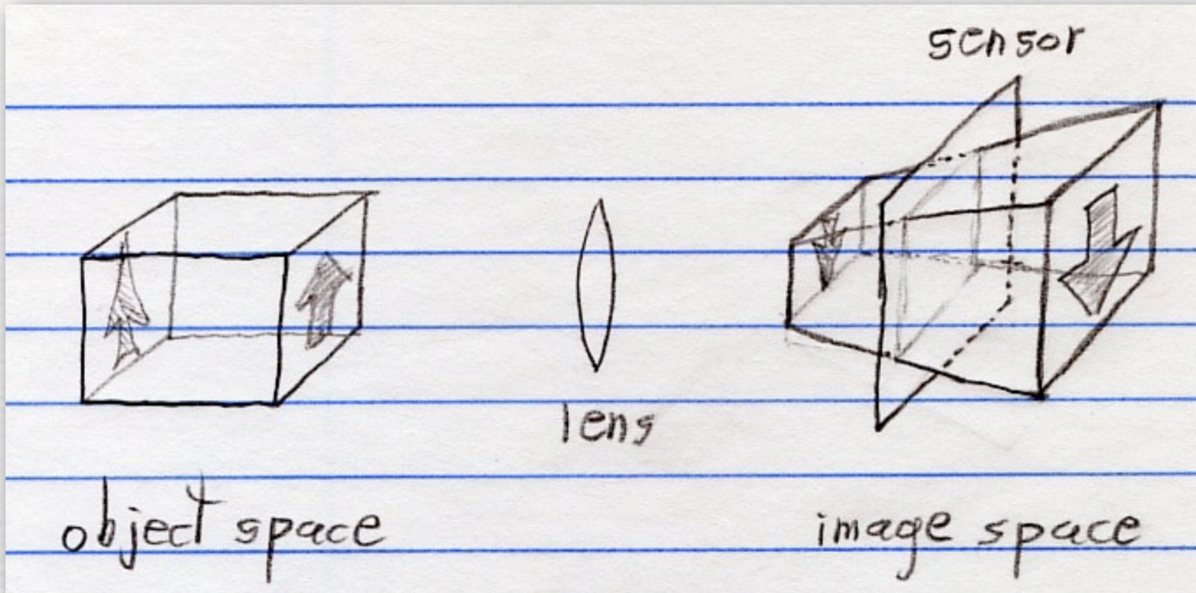

- a cube in object space is transformed by a lens into a 3D frustum in image space, with the orientations shown by the arrows
- in computer graphics this transformation is modeled as a 4 × 4 matrix multiplication of 3D points expressed in 4D homogenous coordinates
- in photography a sensor extracts a 2D slice from the 3D frustum; on this slice some objects may be sharply focused; others may be blurry

Recap

- approximations we sometimes make when analyzing lenses
  - geometrical optics instead of physical optics
  - spherical lenses instead of hyperbolic lenses
  - thin lens representation of thick optical systems
  - paraxial approximation of ray angles
- the Gaussian lens formula relates focal length, object distance, and image distance
  - changing these settings also changes magnification
  - these settings and sensor size determine field of view
  - convex lenses make real images; concave make virtual images
  - lenses perform a 3D perspective transform of object space

Questions?

Depth of field

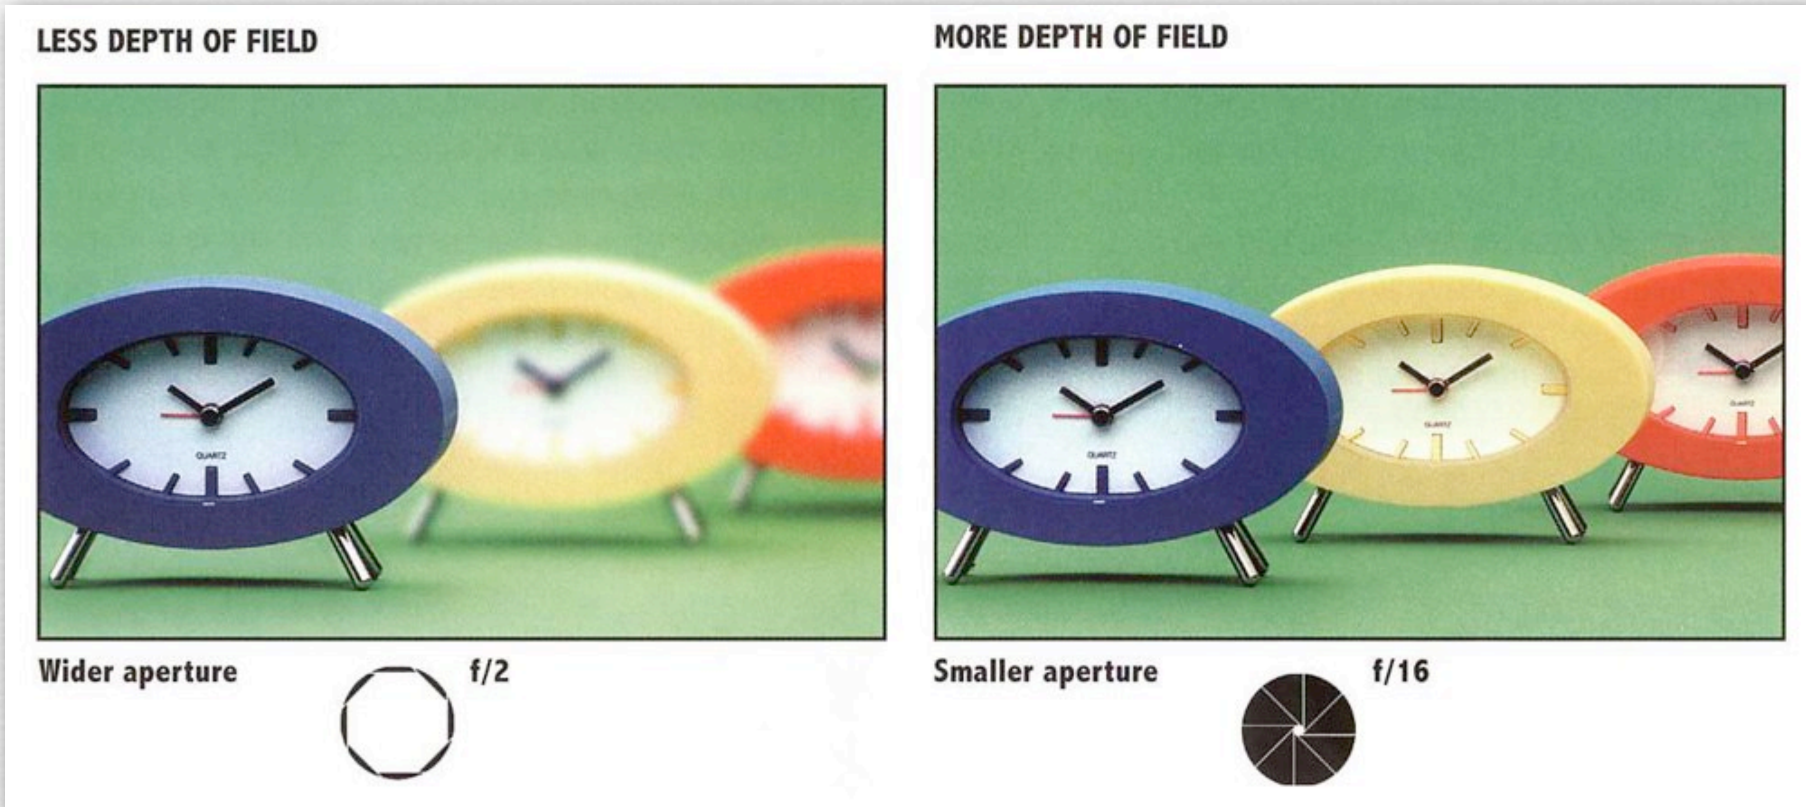

(London)

$$N = \frac{f}{A}$$

- lower N means a wider aperture and less depth of field

How low can N be?

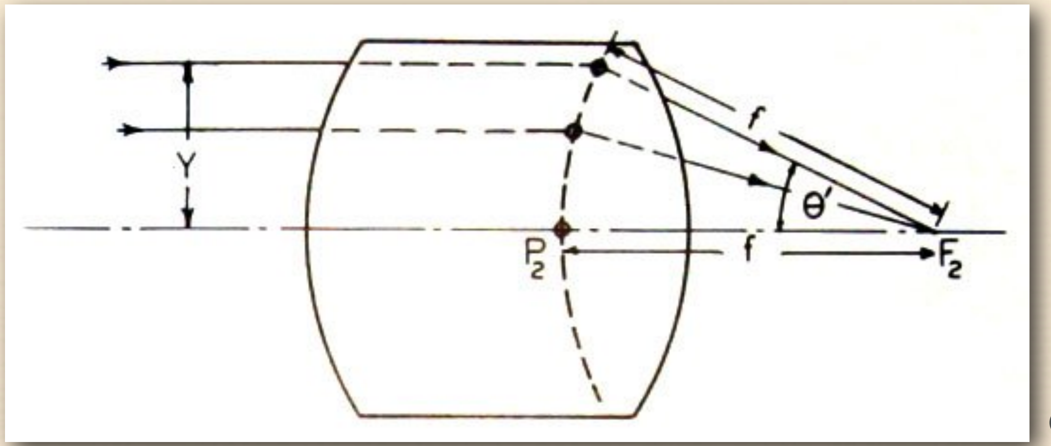

(Kingslake)

- principal planes are the paraxial approximation of a spherical "equivalent refracting surface"

$$N = \frac{1}{2 \sin \theta'}$$

- lowest possible N in air is f/0.5
- lowest N in SLR lenses is f/1.0

Canon EOS 50mm f/1.0
(discontinued)

Cinematography by candlelight

Stanley Kubrick,
Barry Lyndon,
1975

- Zeiss 50mm f/0.7 Planar lens
  - originally developed for NASA's Apollo missions
  - very shallow depth of field in closeups (small object distance)

Cinematography by candlelight

Stanley Kubrick,
Barry Lyndon,
1975

- Zeiss 50mm f/0.7 Planar lens
  - originally developed for NASA's Apollo missions
  - very shallow depth of field in closeups (small object distance)

Circle of confusion (C)

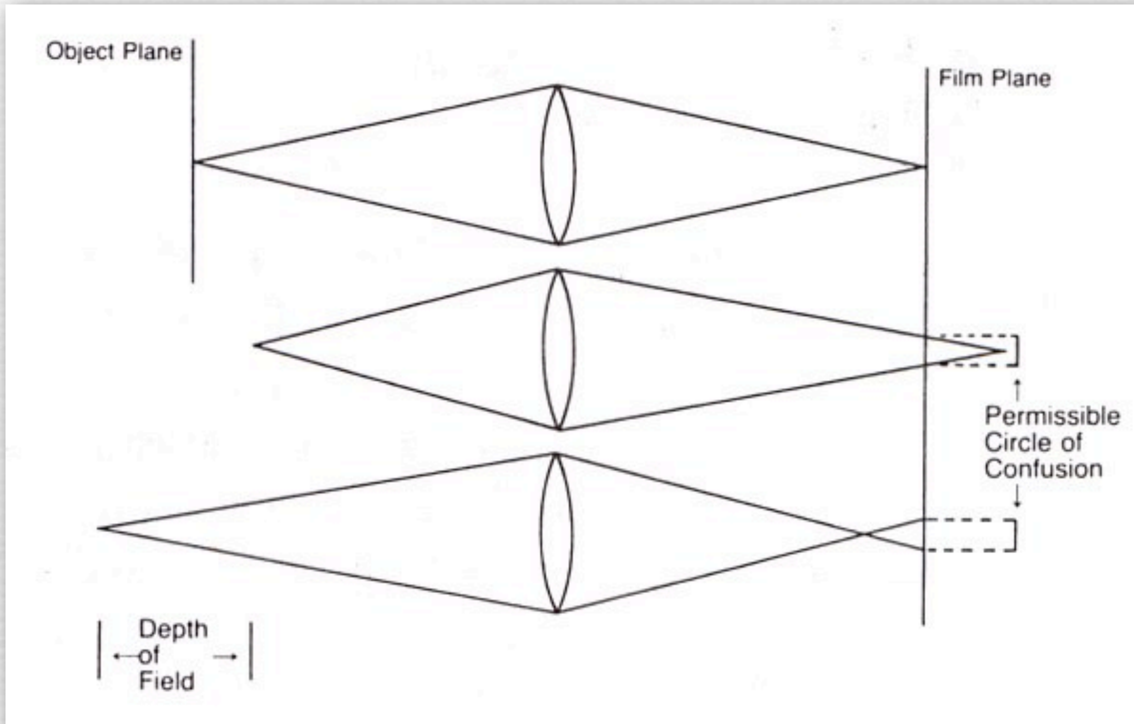

- C depends on sensing medium, reproduction medium, viewing distance, human vision,...
  - for print from 35mm film, 0.02mm (on negative) is typical
  - for high-end SLR, 6μ is typical (1 pixel)
  - larger if downsizing for web, or lens is poor

Depth of field formula

$$M_T \triangleq \frac{y_i}{y_o} = -\frac{s_i}{s_o}$$

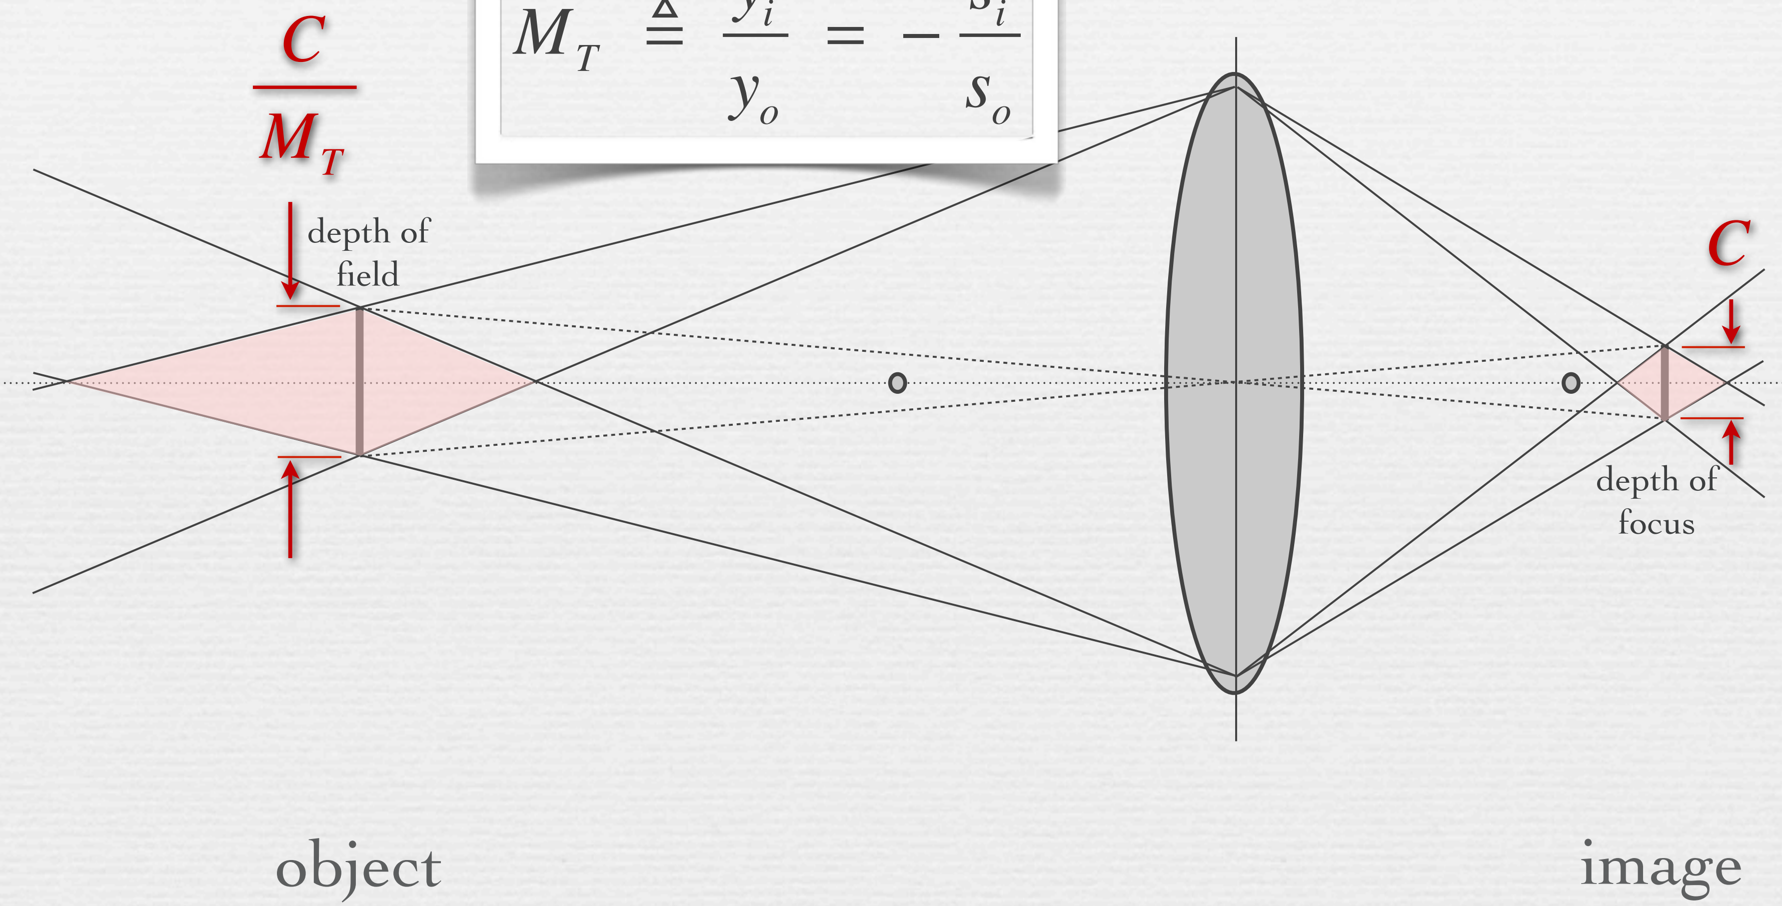

- DoF is asymmetrical around the in-focus object plane
- conjugate in object space is typically bigger than C

Depth of field formula

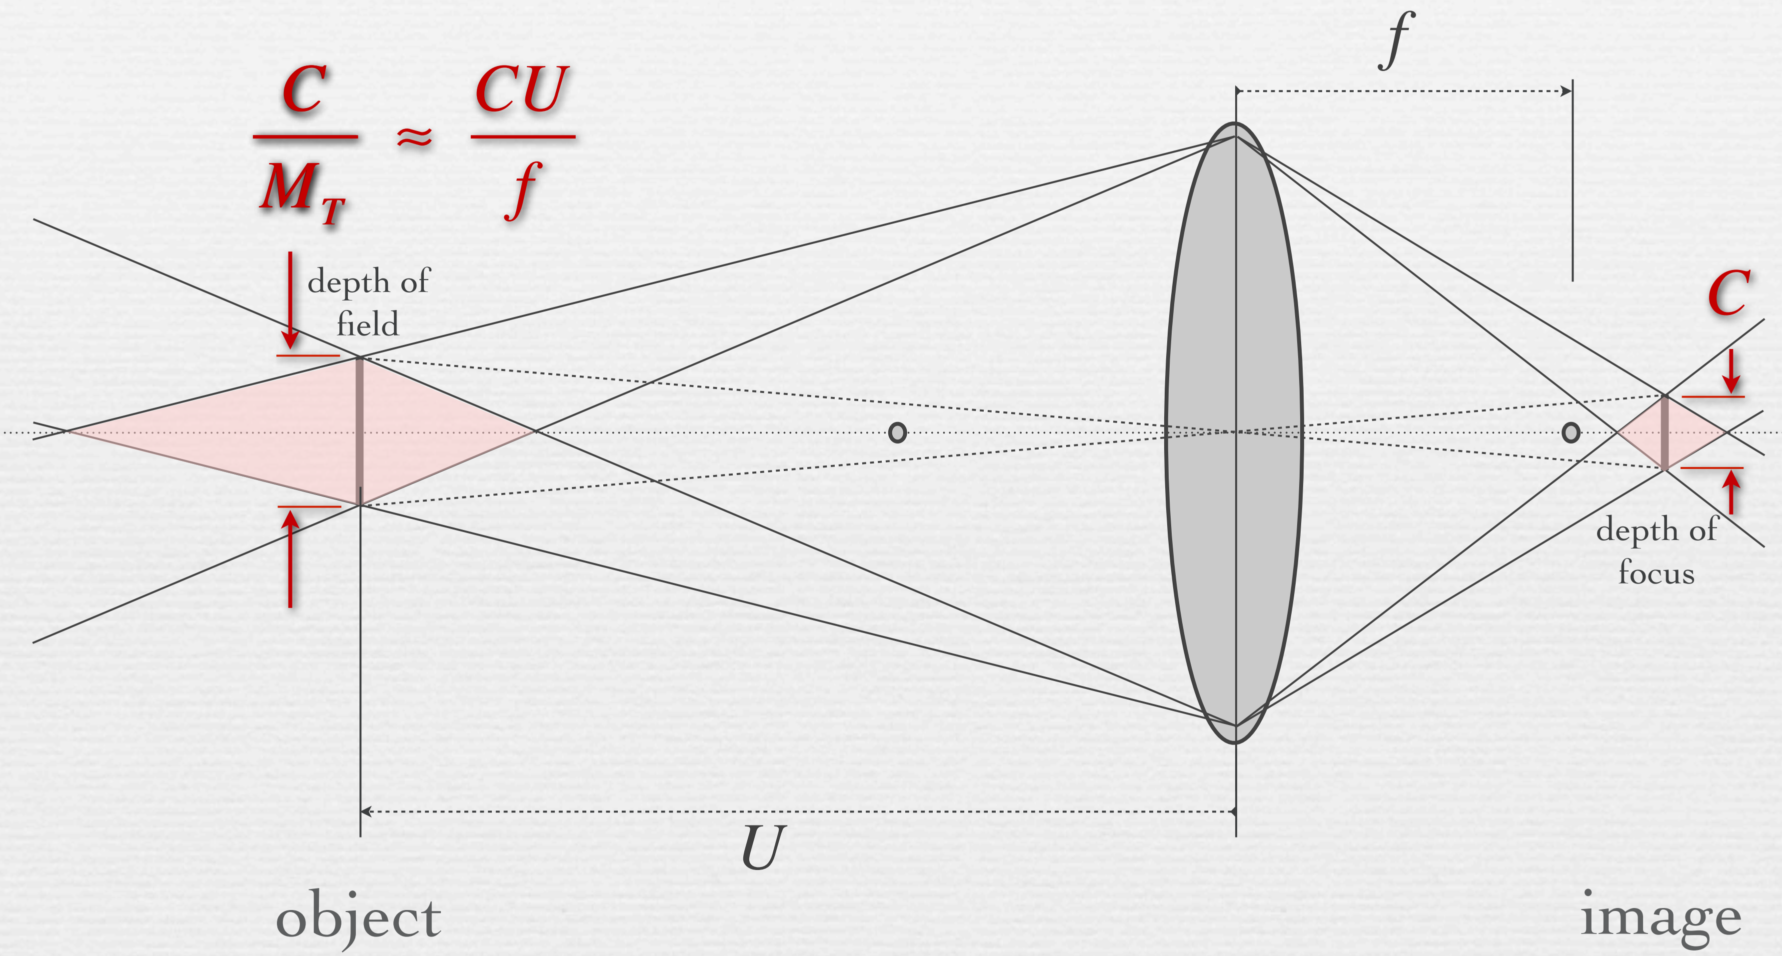

- DoF is asymmetrical around the in-focus object plane
- conjugate in object space is typically bigger than C

Depth of field formula

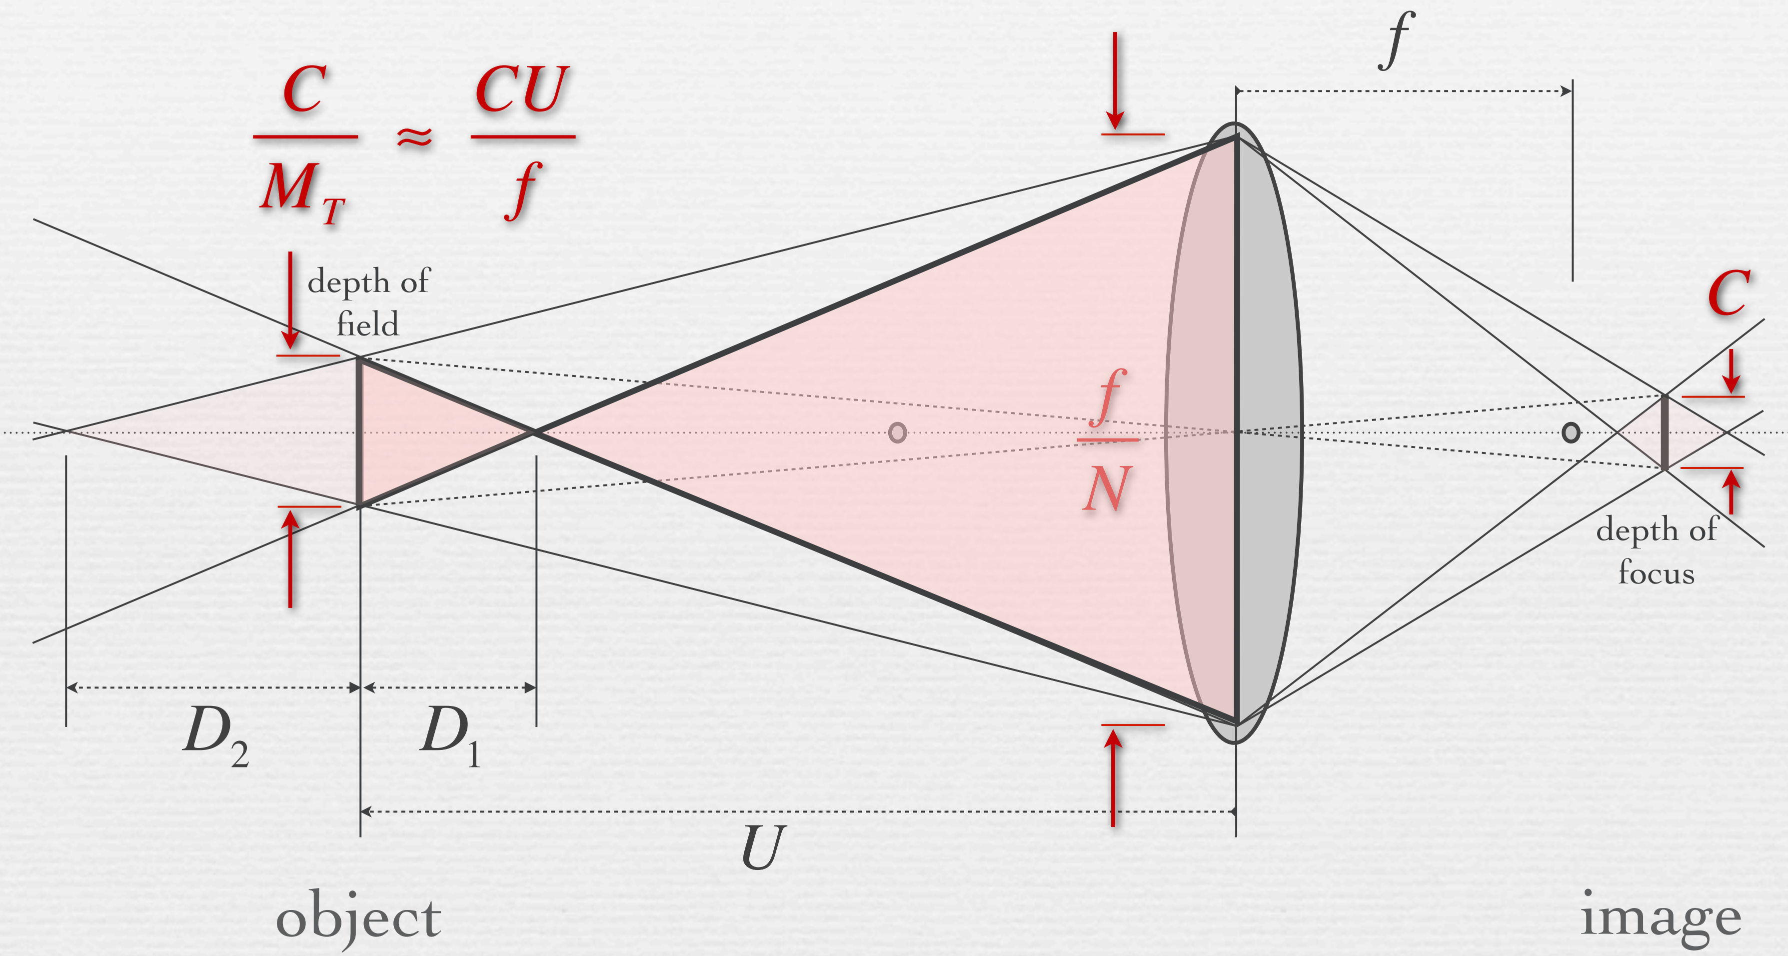

$$\frac{D_1}{CU/f} = \frac{U - D_1}{f/N} \quad \ldots\ldots \quad D_1 = \frac{NCU^2}{f^2 + NCU} \qquad D_2 = \frac{NCU^2}{f^2 - NCU}$$

Depth of field formula

$$D_{TOT} = D_1 + D_2 = \frac{2NCU^2 f^2}{f^4 - N^2C^2U^2}$$

- $N^2C^2U^2$ can be ignored when conjugate of circle of confusion is small relative to the aperture

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

- where
  - N is F-number of lens
  - C is circle of confusion (on image)
  - U is distance to in-focus plane (in object space)
  - f is focal length of lens

$D_{TOT} \approx \frac{2NCU^2}{f^2}$
✦ N = f/4.1
C = 2.5µ
U = 5.9m (19’)
f = 73mm (equiv to 362mm)
D_{TOT} = 132mm
✦ 1 pixel on this video projector
C = 2.5µ × 2816 / 1024 pixels
D_{EFF} = 363mm

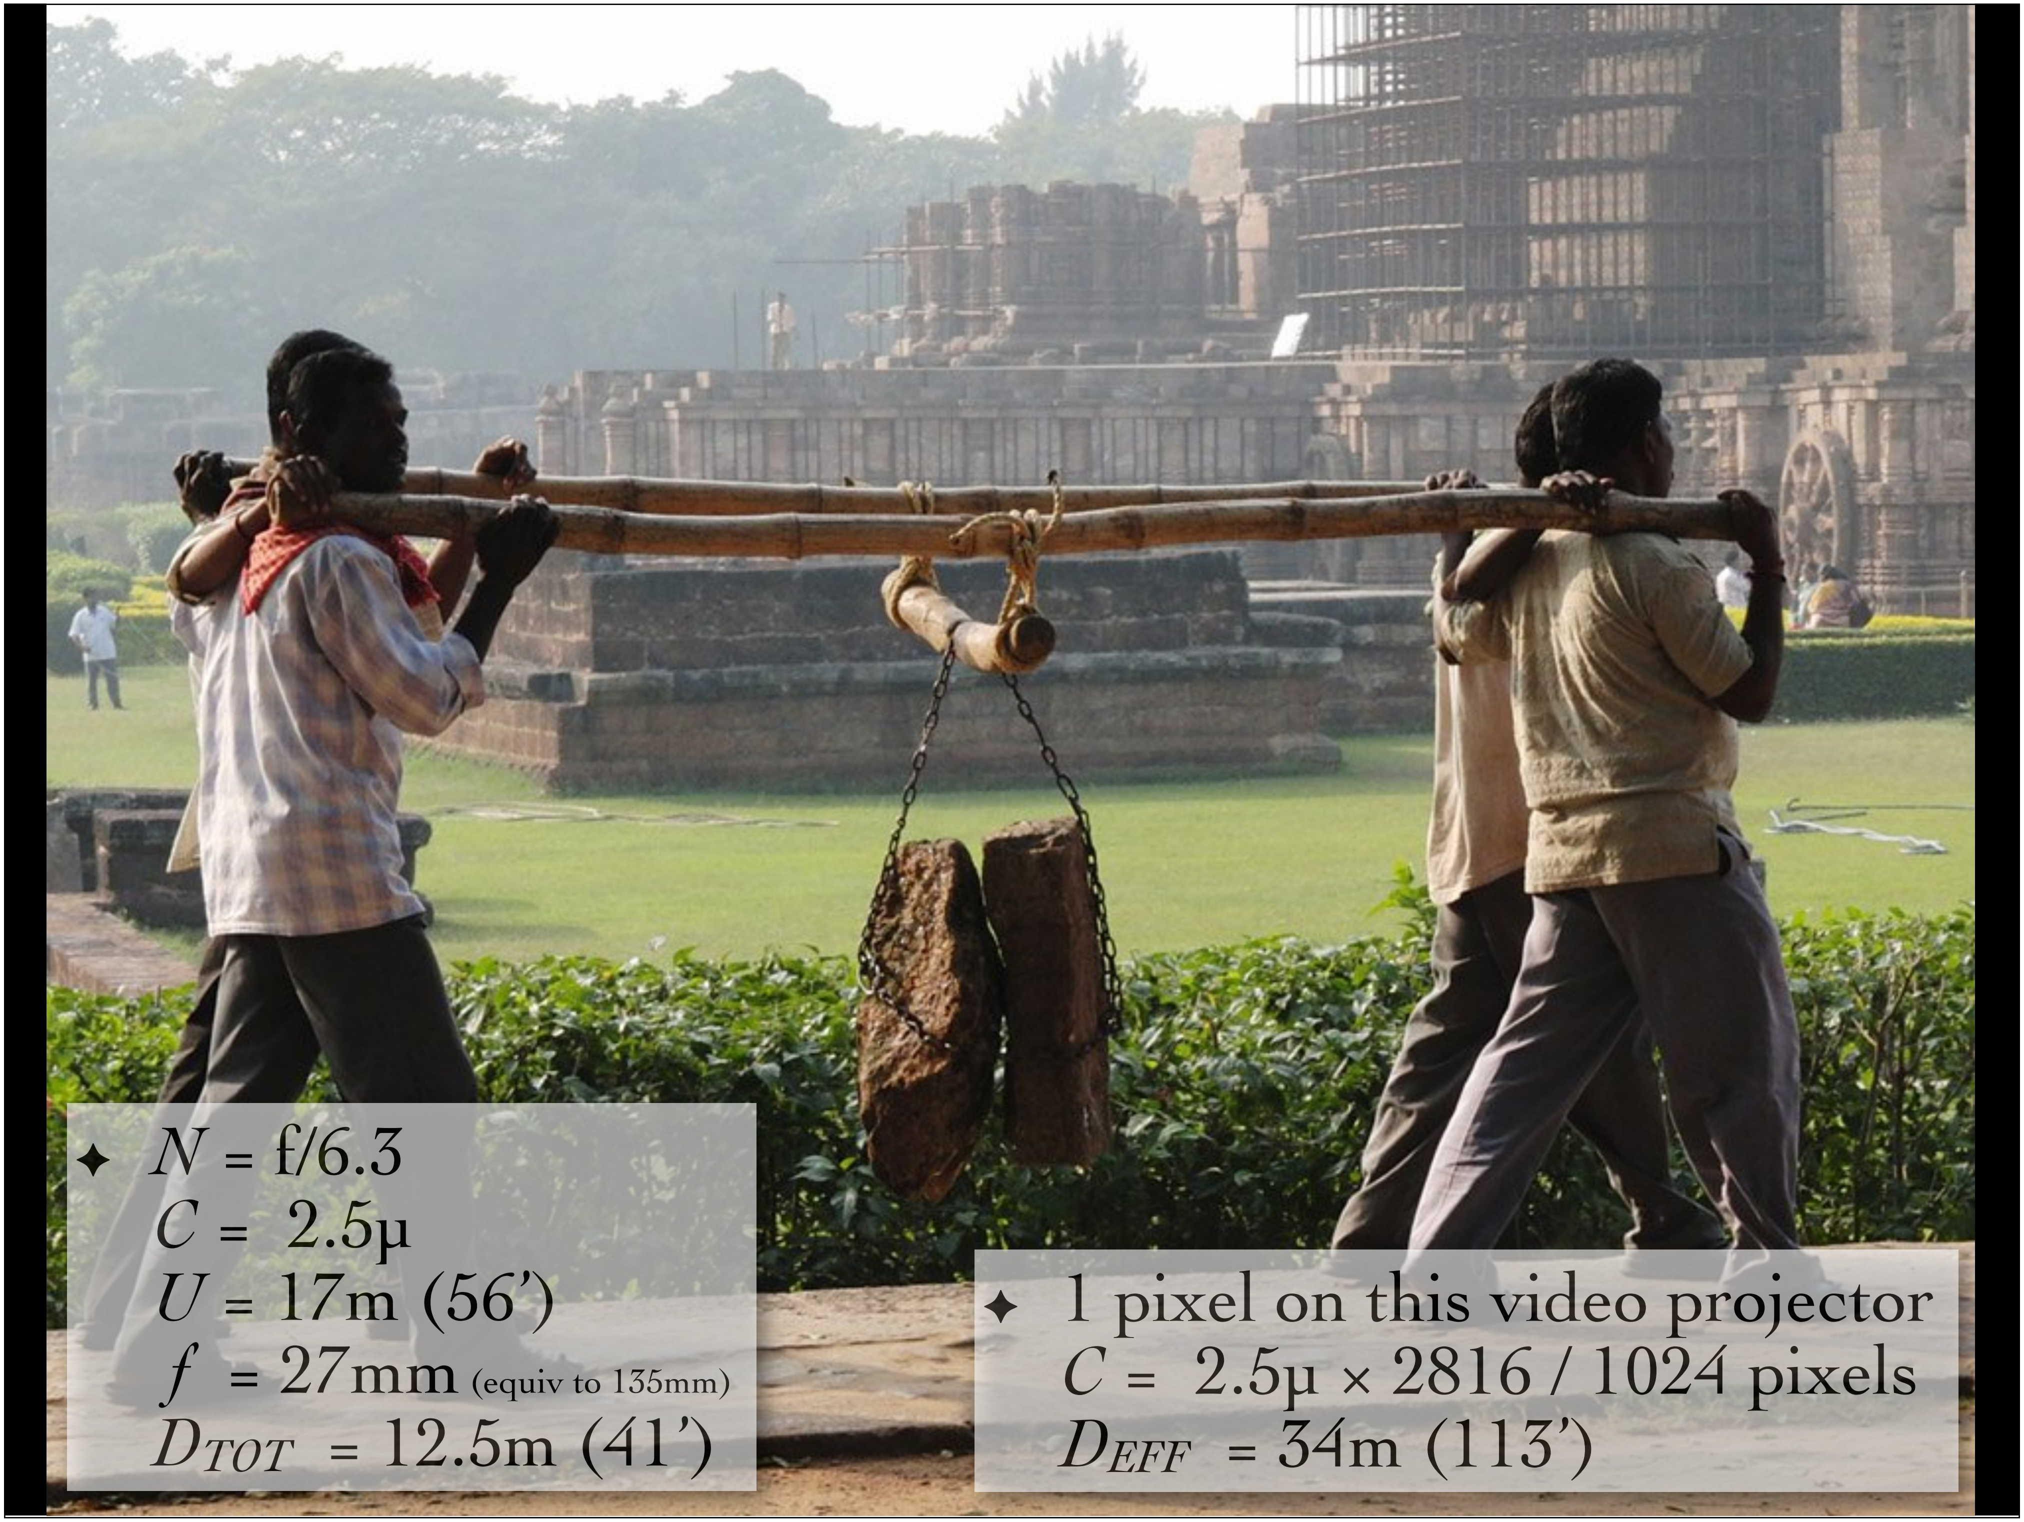

✦ N = f/6.3
C = 2.5μ
U = 17m (56’)
f = 27mm (equiv to 135mm)
D_{TOT} = 12.5m (41’)
✦ 1 pixel on this video projector
C = 2.5μ × 2816 / 1024 pixels
D_{EFF} = 34m (113’)

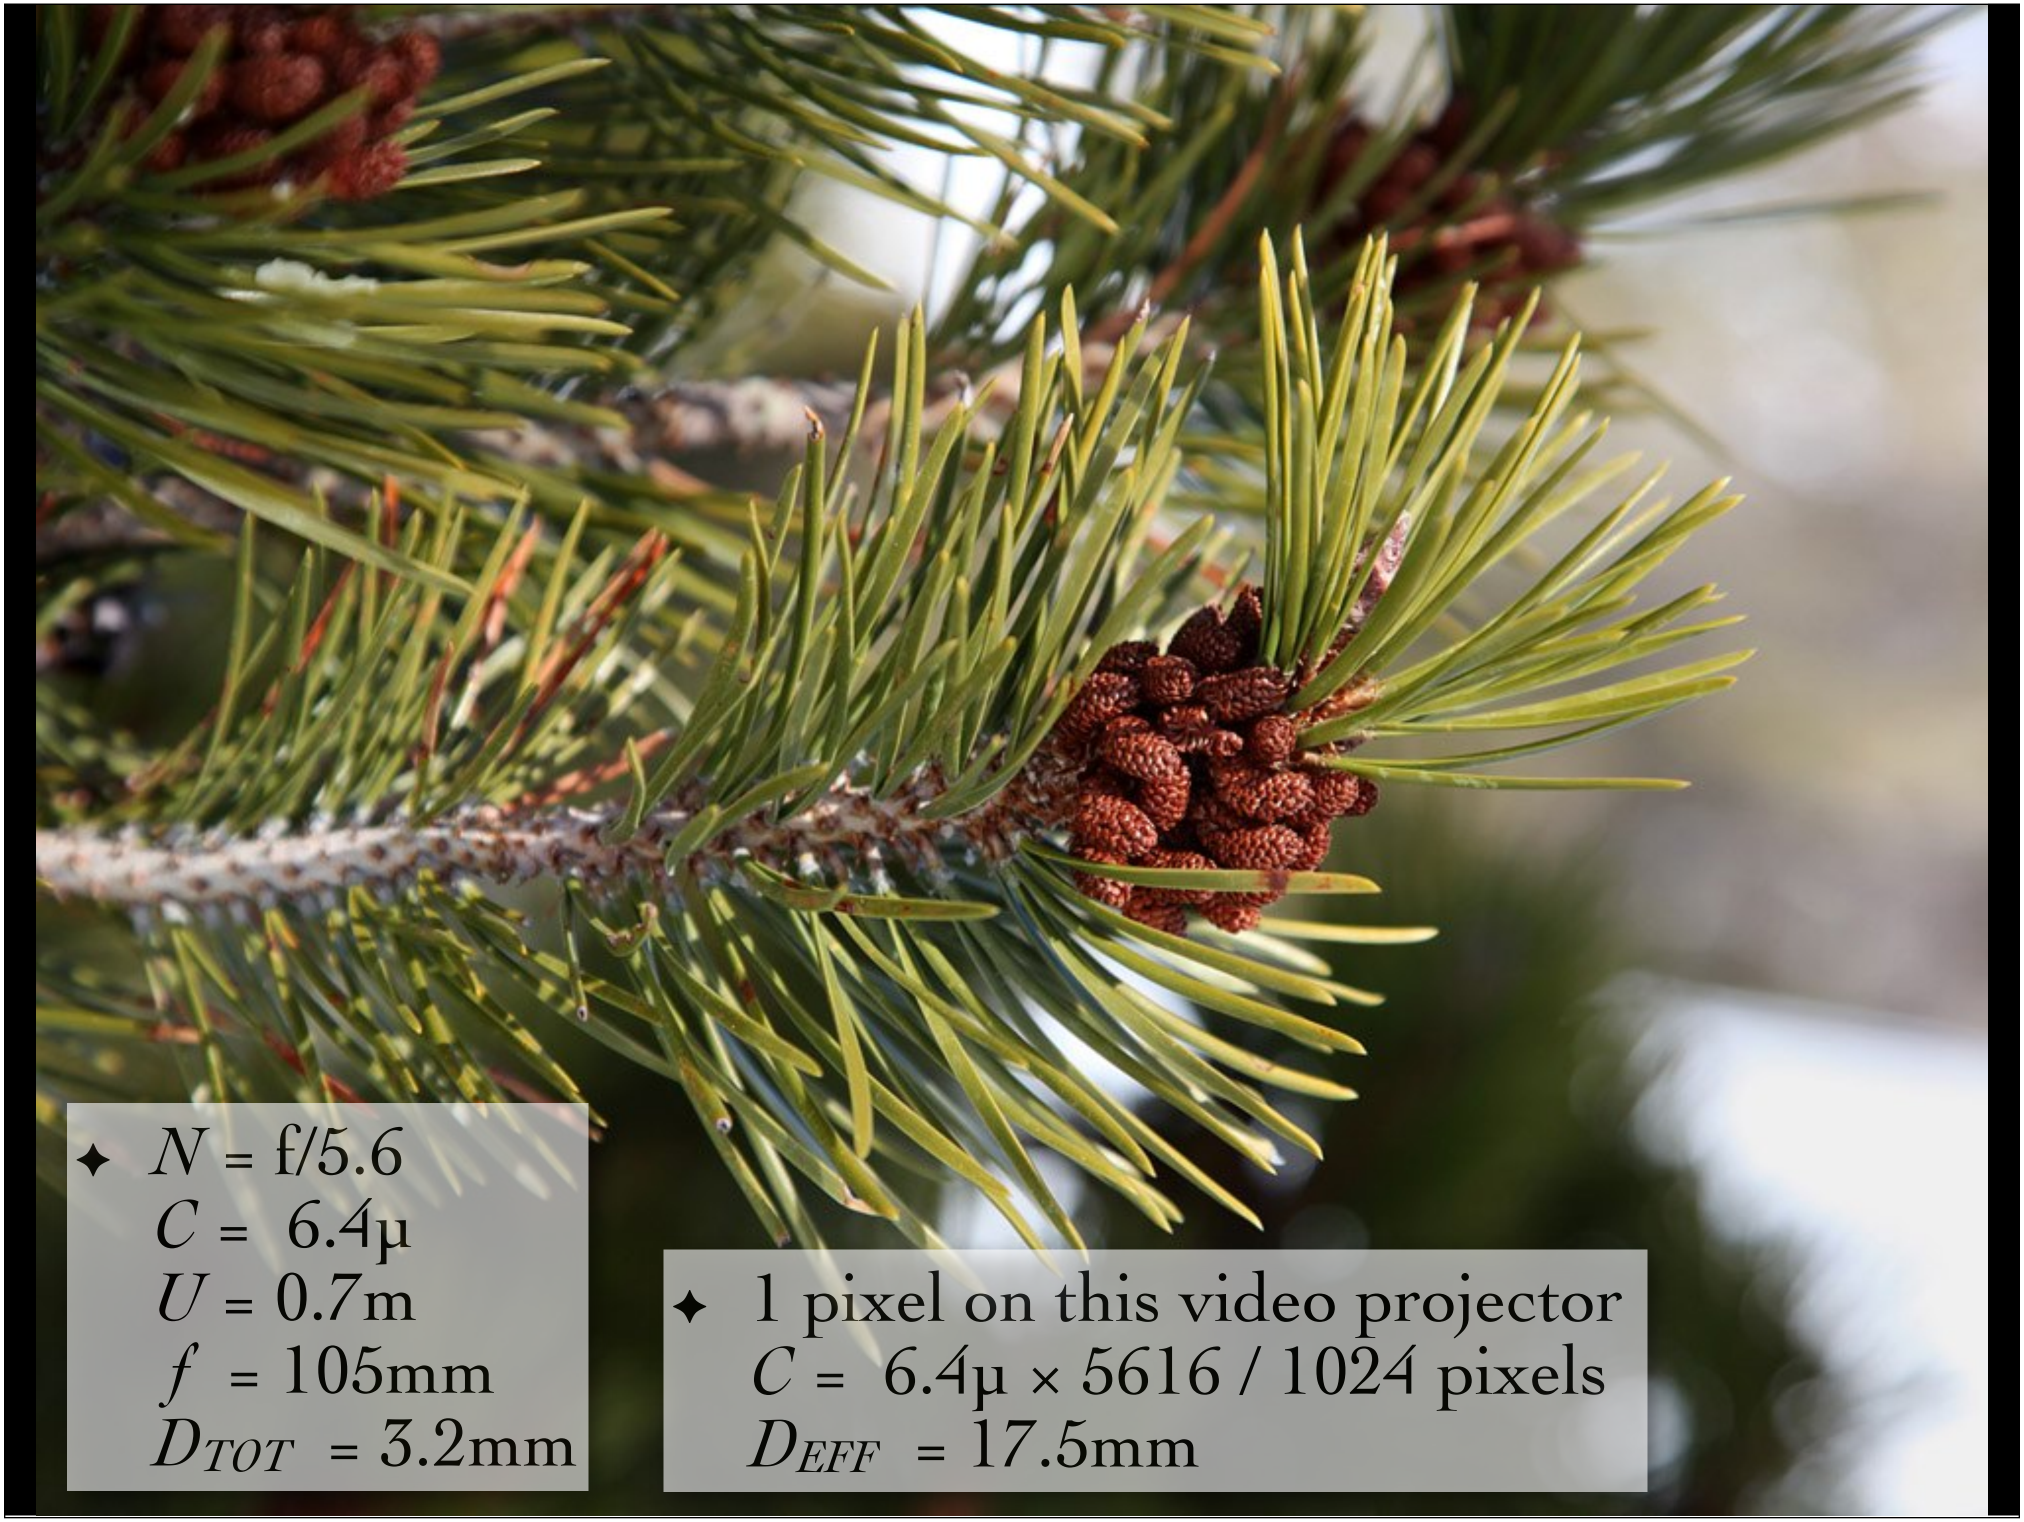
✦ N = f/5.6
C = 6.4µ
U = 0.7m
f = 105mm
D_{TOT} = 3.2mm
✦ 1 pixel on this video projector
C = 6.4µ × 5616 / 1024 pixels
D_{EFF} = 17.5mm

Canon MP-E
65mm 5:1 macro

(Mikhail Shlemov)

- N = f/2.8
  C = 6.4μ
  U = 78mm
  f = 65mm
  (use $N' = (1+M_T)N$ at short conjugates (M_T=5 here)) = f/16
  D_{TOT} = 0.29mm!

Sidelight: macro lenses

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}$$

Q. How can the Casio EX-F1 at 73mm and the Canon MP-E 65mm macro, which have similar *f* 's, have such different focusing distances?

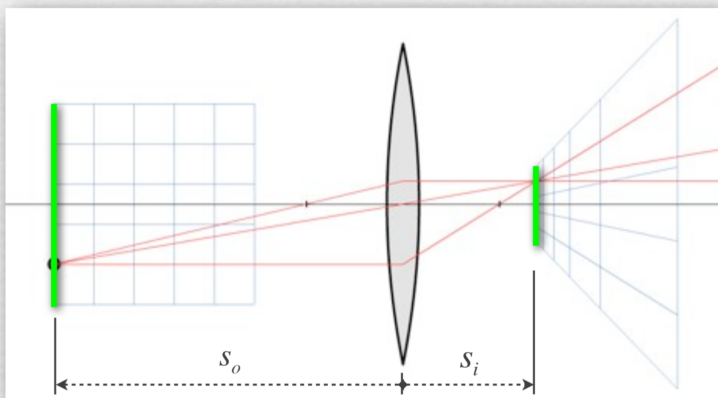

normal

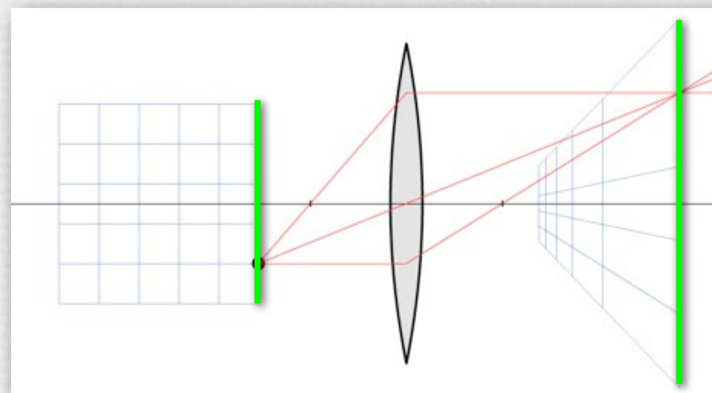

macro

- A. Because macro lenses are built to allow long s_i
  - this changes s_o, which changes magnification $M_T \triangleq -s_i / s_o$
  - macro lenses are also well corrected for aberrations at short s_o

Extension tube: fits between camera and lens, converts a normal lens to a macro lens

- toilet paper tube, black construction paper, masking tape
- camera hack by Katie Dektar (CS 178, 2009)

Extension tubes versus close-up filters

Canon 25mm

Canon f = 500mm

- both allow closer focusing, hence greater magnification
- both degrade image quality relative to a macro lens
- extension tubes work best with wide-angle lenses; close-up filters work best with telephoto lenses
- extension tubes raise F-number, reducing light
- need different close-up filter for each lens filter diameter

Extension tubes versus close-up filters versus teleconverters

Canon 25mm

Canon f = 500mm

Nikon 1.4×

- a teleconverter fits between the camera and lens, like an extension tube
- they increase *f*, narrowing FOV & increasing magnification, but they don't change the focusing range
- like extension tubes, they raise F-number, reducing light, and they are awkward to add or remove
- see http://www.cambridgeincolour.com/tutorials/macro-extension-tubes-closeup.htm

DoF is linear with F-number

(juzaphoto.com)

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/dof.html

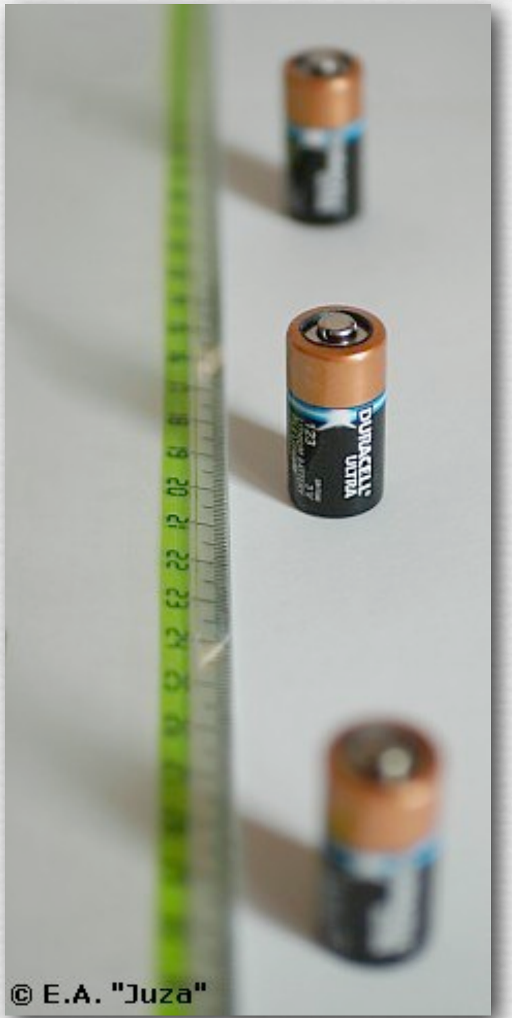

f/2.8

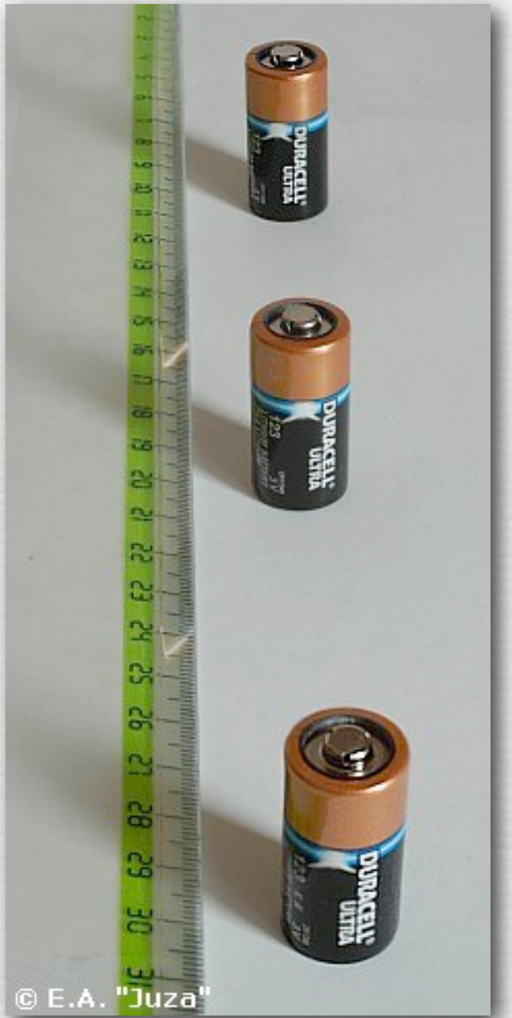

f/32

DoF is quadratic with subject distance

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/dof.html

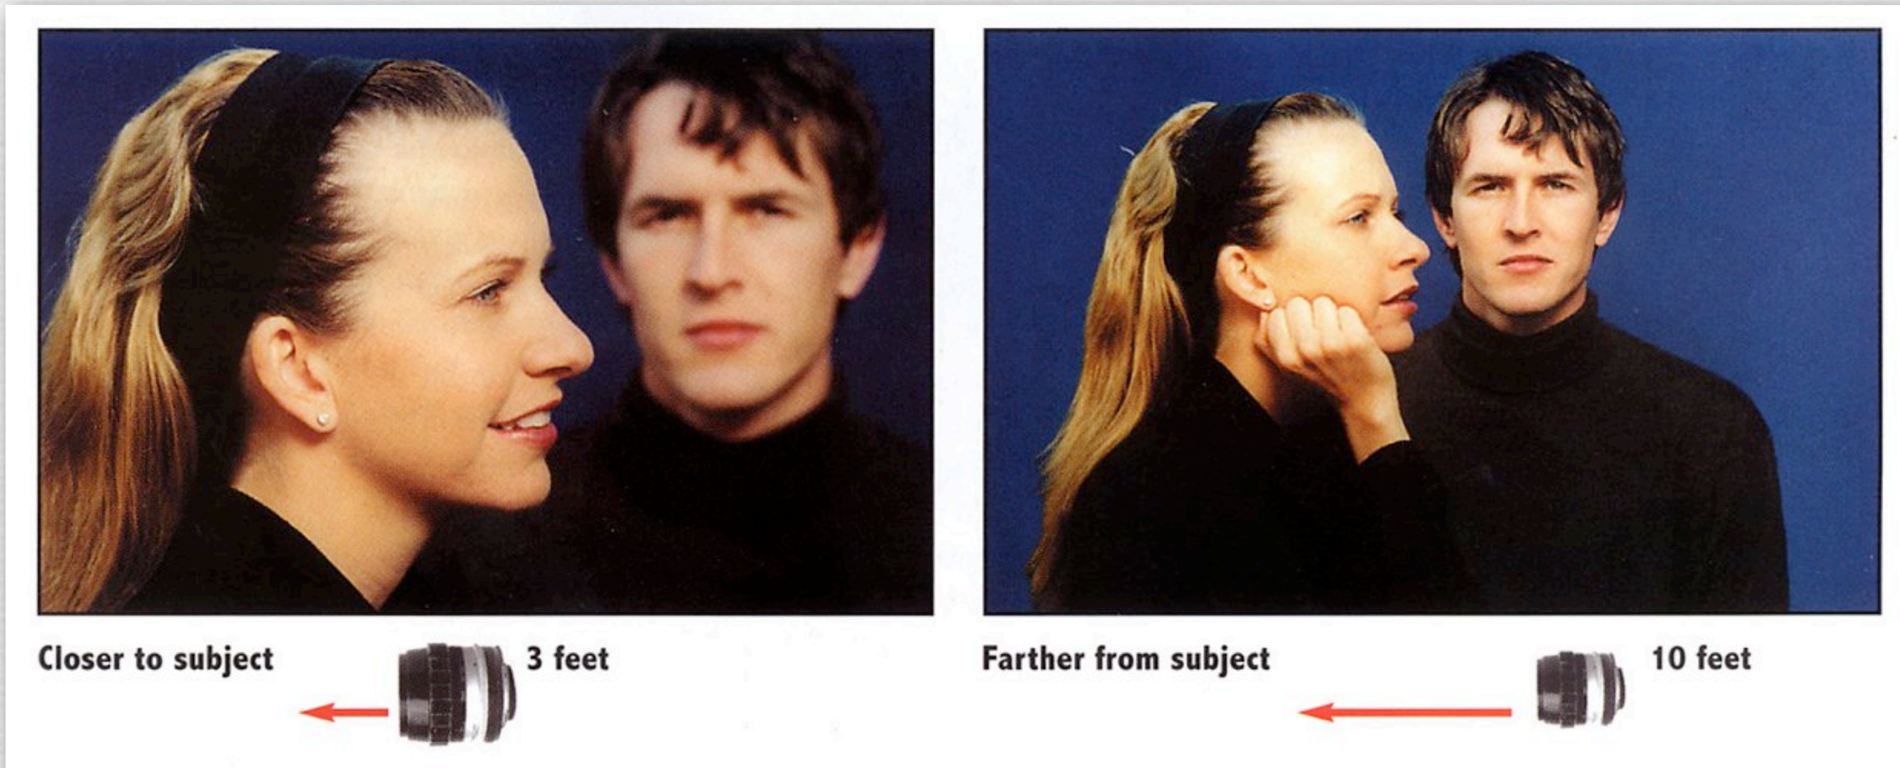

(London)

Hyperfocal distance

An observant student discovered an error in my calculation of H for this photograph. I believe the values below are now correct. I've changed the projector to an "HD projector" (1920 pixels wide) to make the numbers work out conveniently.

- the back depth of field

$$D_2 = \frac{NCU^2}{f^2 - NCU}$$

- becomes infinite if

$$U \geq \frac{f^2}{NC} \triangleq H$$

- In that case, the front depth of field becomes

$$D_1 = \frac{H}{2}$$

- N = f/6.3
  C = 2.5μ × 2816 / 1920 pixels
  U = 17m (56')
  f = 27mm (equiv to 135mm)
  D_{TOT} = 18.3m on HD projector
  H = 31.6m (104')

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/dof.html

- so if I had focused at 32m, everything from 16m to infinity would be in focus on a video projector, including the men at 17m

DoF is inverse quadratic with focal length

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/dof.html

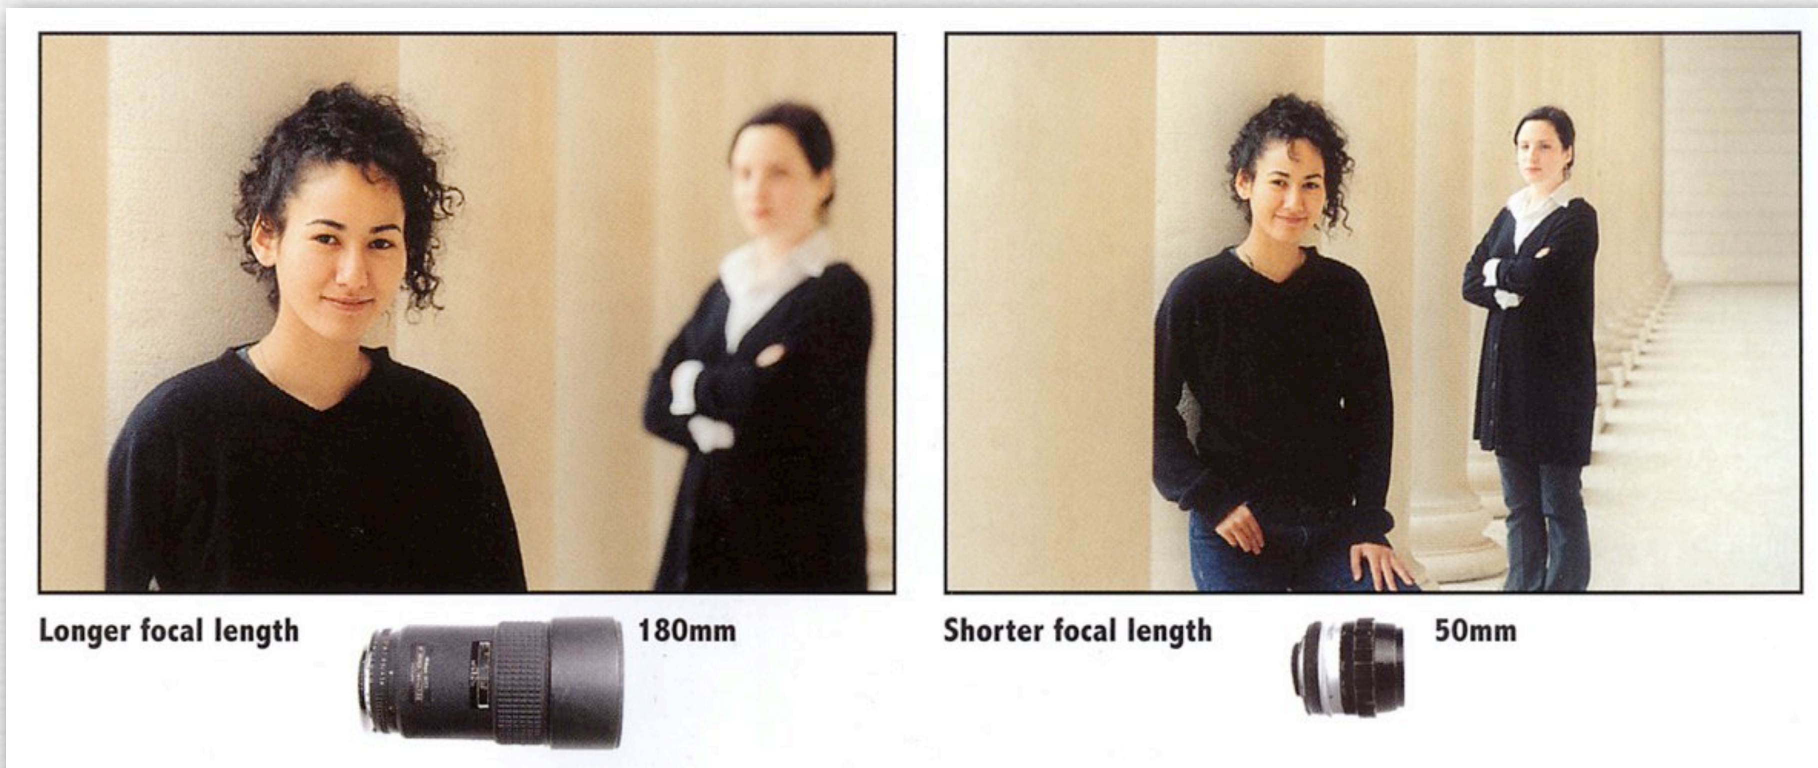

(London)

Q. Does sensor size affect DoF?

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

- as sensor shrinks, lens focal length f typically shrinks to maintain a comparable field of view
- as sensor shrinks, pixel size C typically shrinks to maintain a comparable number of pixels in the image
- thus, depth of field D_{TOT} increases linearly with decreasing sensor size
- this is why amateur cinematographers are drawn to SLRs
  - their chips are larger than even pro-level video camera chips
  - so they provide unprecedented control over depth of field

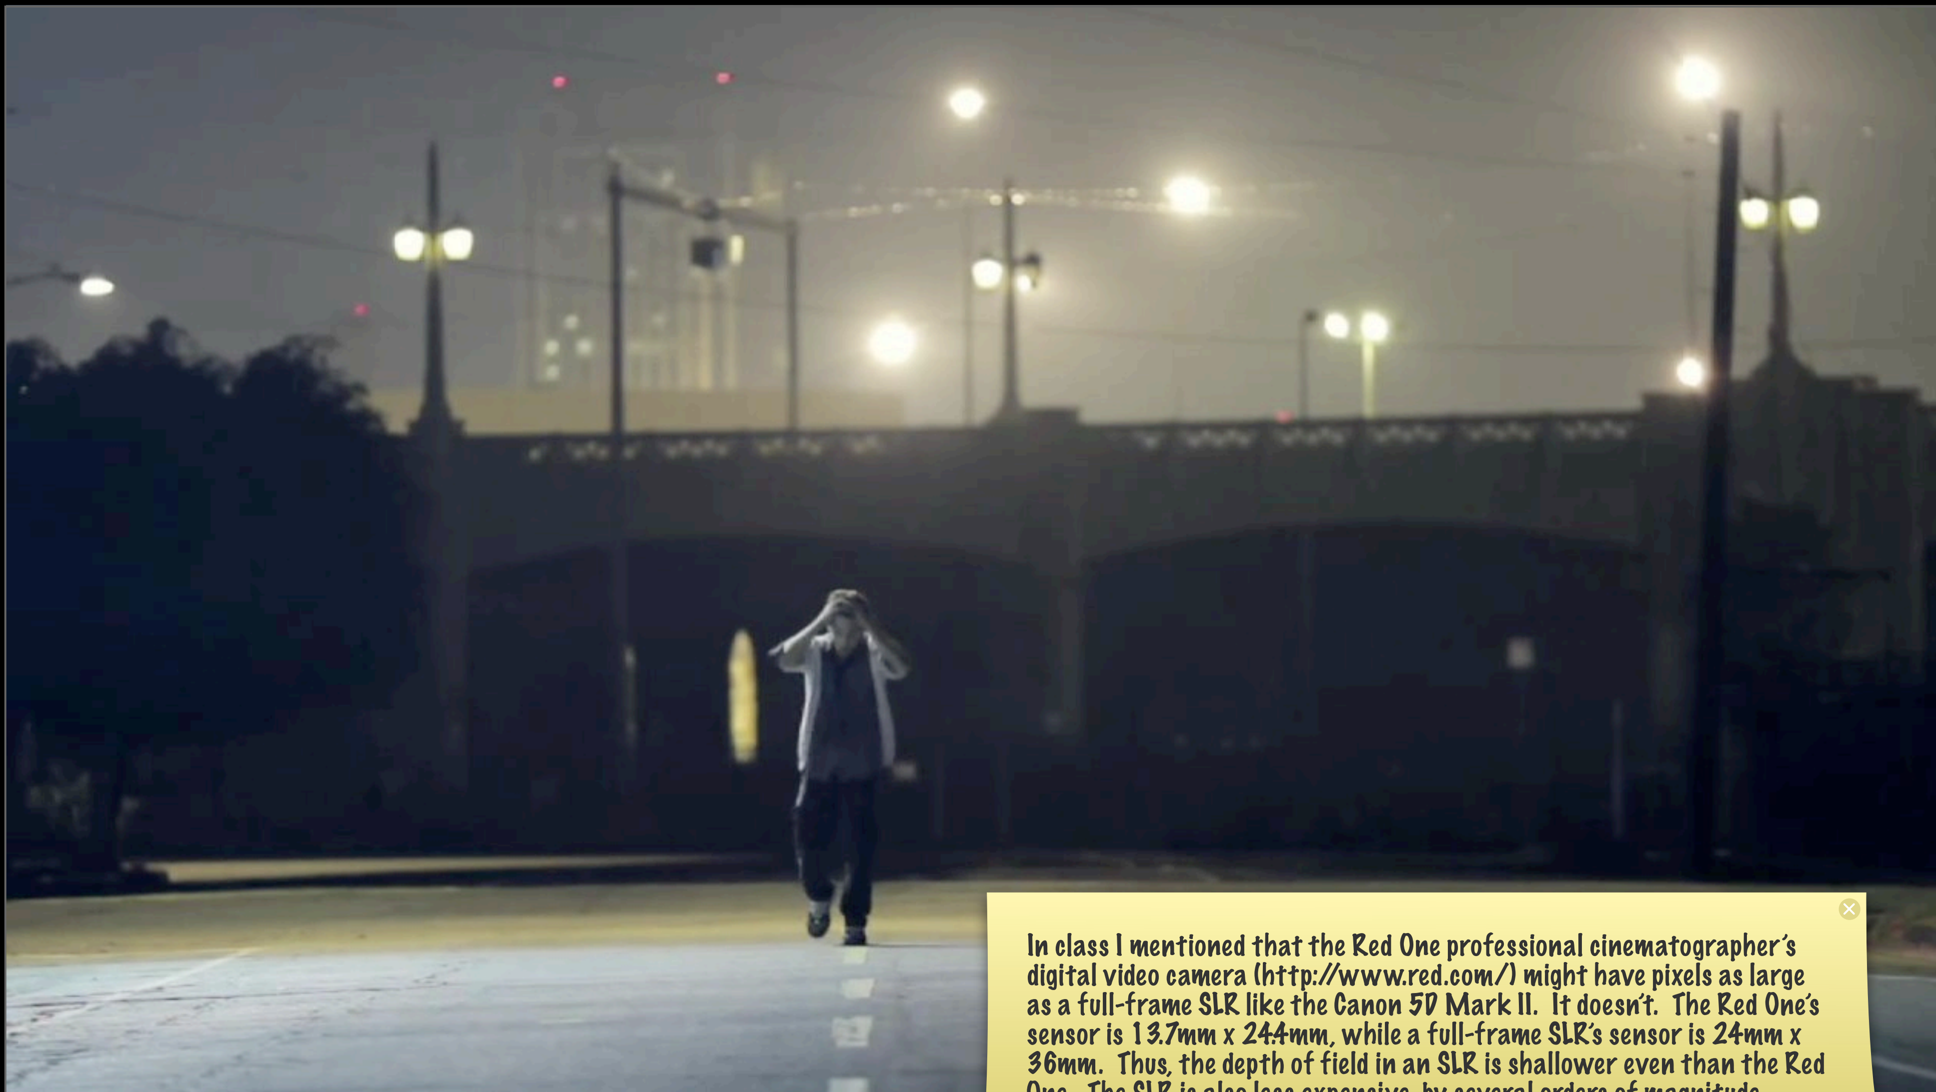

In class I mentioned that the Red One professional cinematographer's digital video camera (http://www.red.com/) might have pixels as large as a full-frame SLR like the Canon 5D Mark II. It doesn't. The Red One's sensor is 13.7mm x 24.4mm, while a full-frame SLR's sensor is 24mm x 36mm. Thus, the depth of field in an SLR is shallower even than the Red One. The SLR is also less expensive, by several orders of magnitude.

Vincent Laforet, Nocturne (2009)
Canon 1D Mark IV

DoF and the dolly-zoom

- if we zoom in (increase f)
  and stand further back (decrease U) by the same factor

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

- the depth of field at the subject stays the same!
  - useful for macro when you can't get close enough

50mm f/4.8

(juzaphoto.com)

200mm f/4.8,
moved back 4× from subject

Macro photography using a telephoto lens (contents of whiteboard)

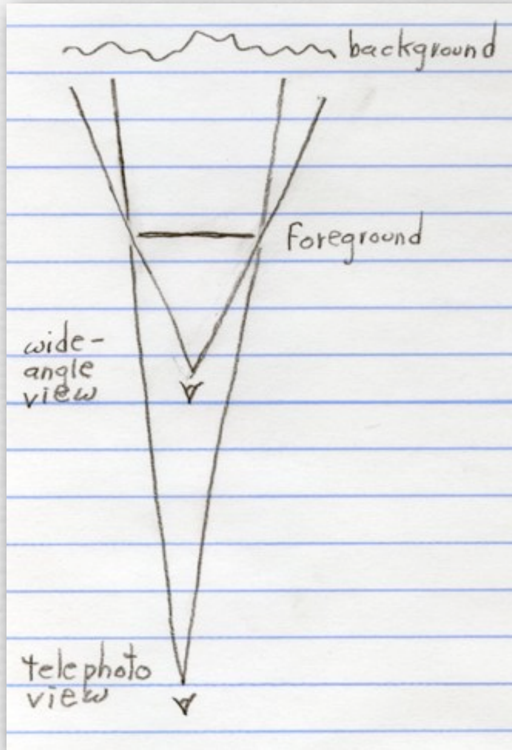

- changing from a wide-angle lens to a telephoto lens and stepping back, you can make a foreground object appear the same size in both lenses
- and both lenses will have the same depth of field on that object
- but the telephoto sees a smaller part of the background (which it blows up to fill the field of view), so the background will appear blurrier

Parting thoughts on DoF: the zen of *bokeh*

Canon 85mm prime f/1.8 lens

- the appearance of small out-of-focus features in a photograph with shallow depth of field
  - determined by the shape of the aperture
  - people get religious about it
  - but not every picture with shallow DoF has evident bokeh...

As I was presenting this slide, it occurred to me that it may not be obvious why the bokeh takes on the shape of the aperture.

If you look back at slide 48 (on the Circle of Confusion), it's not hard to see that for scene points closer to the lens than the in-focus plane (middle diagram on that slide), they will converge to a "splat" on the sensor, rather than converging to a sharp point.

If you think about where the rays come from that form the outer boundary of this splat, you will see that they pass through the outer edge of the lens. If the lens is circular, then the splat will be circular. If the lens is "stopped down" using an aperture composed of 8 leaves as in the Canon 85/1.8 pictured at upper-right on this slide, then the splat will be octagonal. This determines the shape of the bokeh.

Natasha Gelfand (Canon 100mm f/2.8 prime macro lens)

Games with bokeh

- picture by Alice Che (CS 178, 2010)
  - heart-shaped mask in front of lens
  - subject was Christmas lights
  - photograph was misfocused and under-exposed

Parting thoughts on DoF: seeing through occlusions

(Fredo Durand)

- depth of field is not a convolution of the image
  - i.e. not the same as blurring in Photoshop
  - DoF lets you eliminate occlusions, like a chain-link fence

Seeing through occlusions using a large aperture (contents of whiteboard)

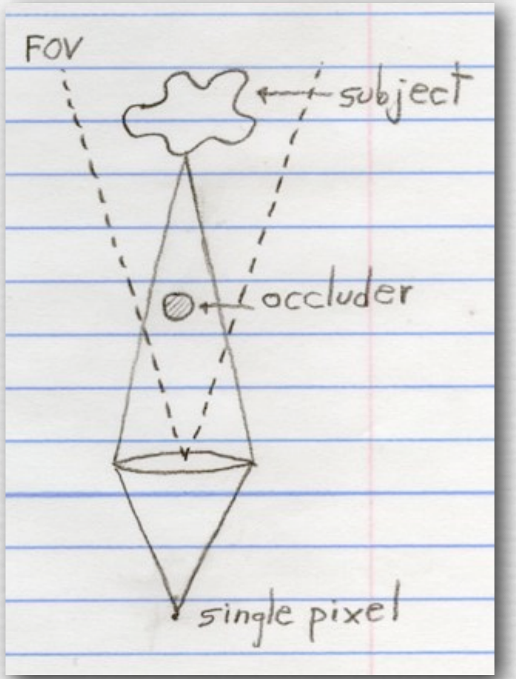

- for a pixel focused on the subject, some of its rays will strike the occluder, but some will pass to the side of it, if the occluder is small enough
- the pixel will then be a mixture of the colors of the subject and occluder
- thus, the occluder reduces the contrast of your image of the subject, but it doesn't actually block your view of it

Tradeoffs affecting depth of field

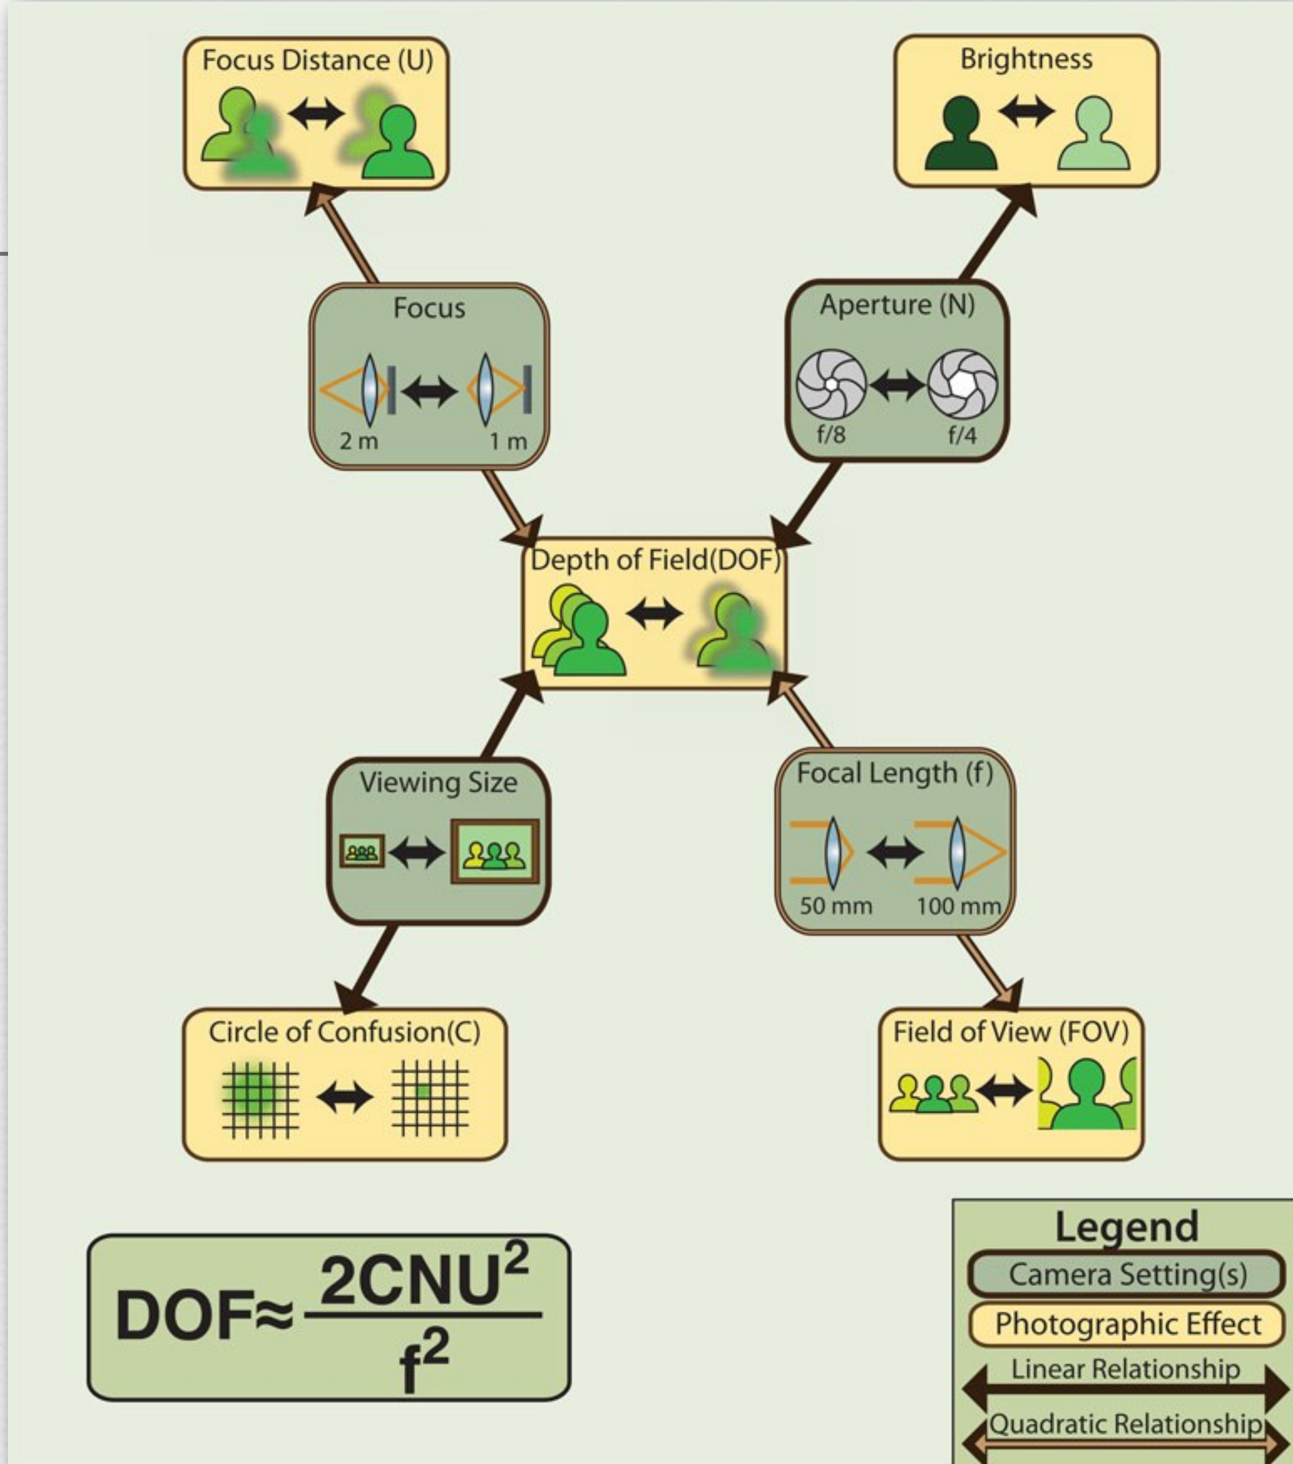

$$DOF \approx \frac{2CNU^2}{f^2}$$

(Eddy Talvala)

Recap

- depth of field (D_{TOT}) is governed by circle of confusion (C), aperture size (N), subject distance (U), and focal length (f)

$$D_{TOT} \approx \frac{2NCU^2}{f^2}$$

  - depth of field is linear in some terms and quadratic in others
  - if you focus at the hyperfocal distance $H = f^2 / NC$, everything from $H / 2$ to infinity will be in focus
  - depth of field increases linearly with decreasing sensor size
- useful sidelights
  - bokeh refers to the appearance of small out-of-focus features
  - you can take macro photographs using a telephoto lens
  - depth of field blur is not the same as blurring an image

Questions?