Background

- Sales of digital cameras surpassed sales of film cameras in 2004.

Digital cameras are boring

- The most common type of digital camera today: cellphone camera.

Can we leverage the computational power?

Why not use sensors without optics?

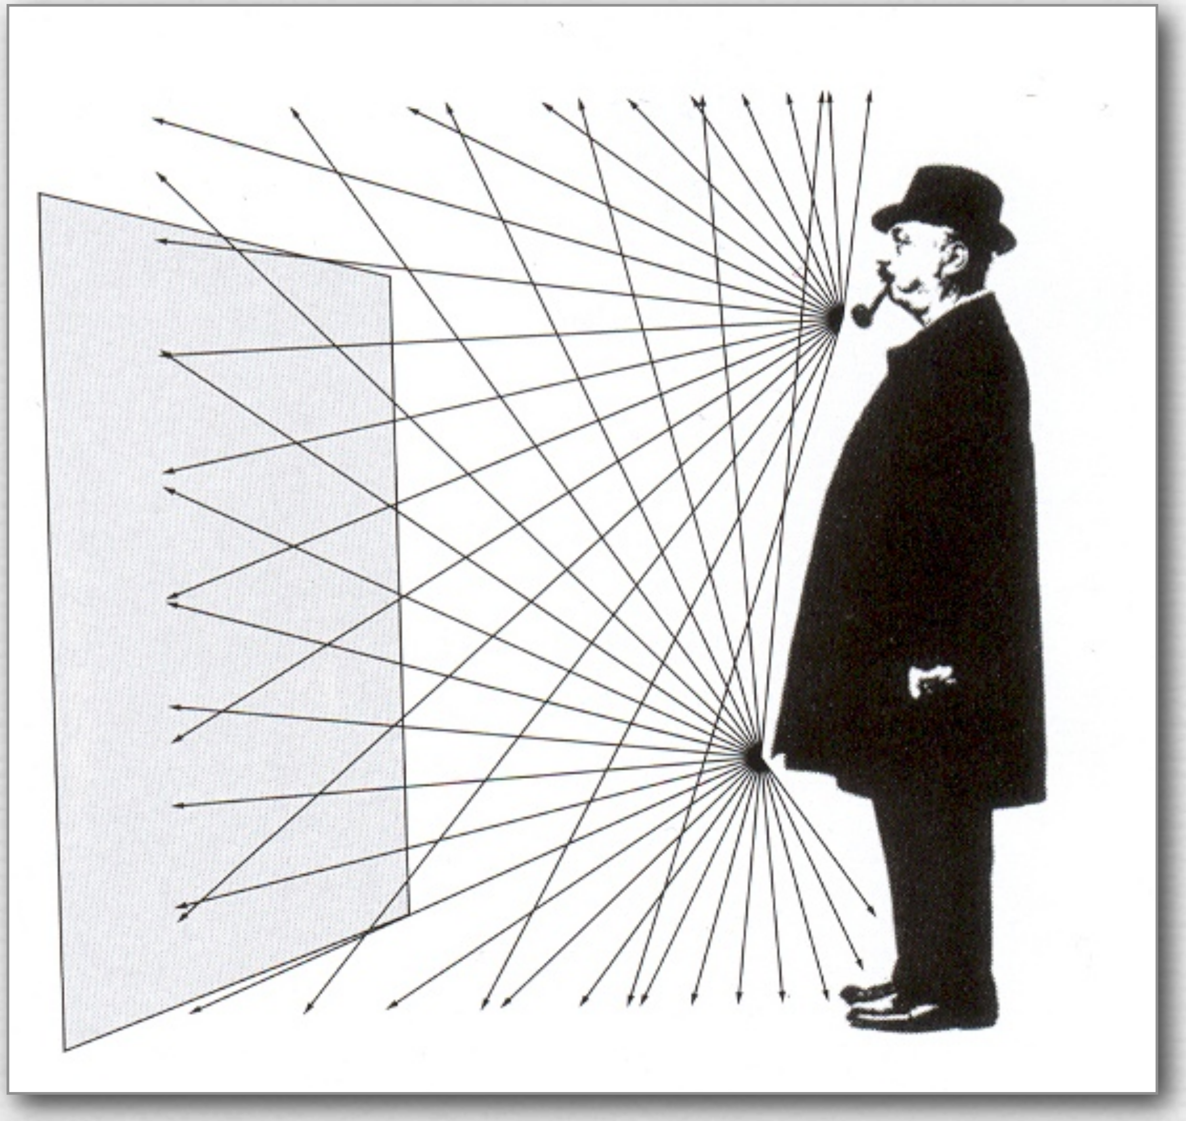

(London)

- each point on sensor would record the integral of light arriving from every point on subject
- all sensor points would record similar colors

Pinhole camera
(a.k.a. *camera obscura*)

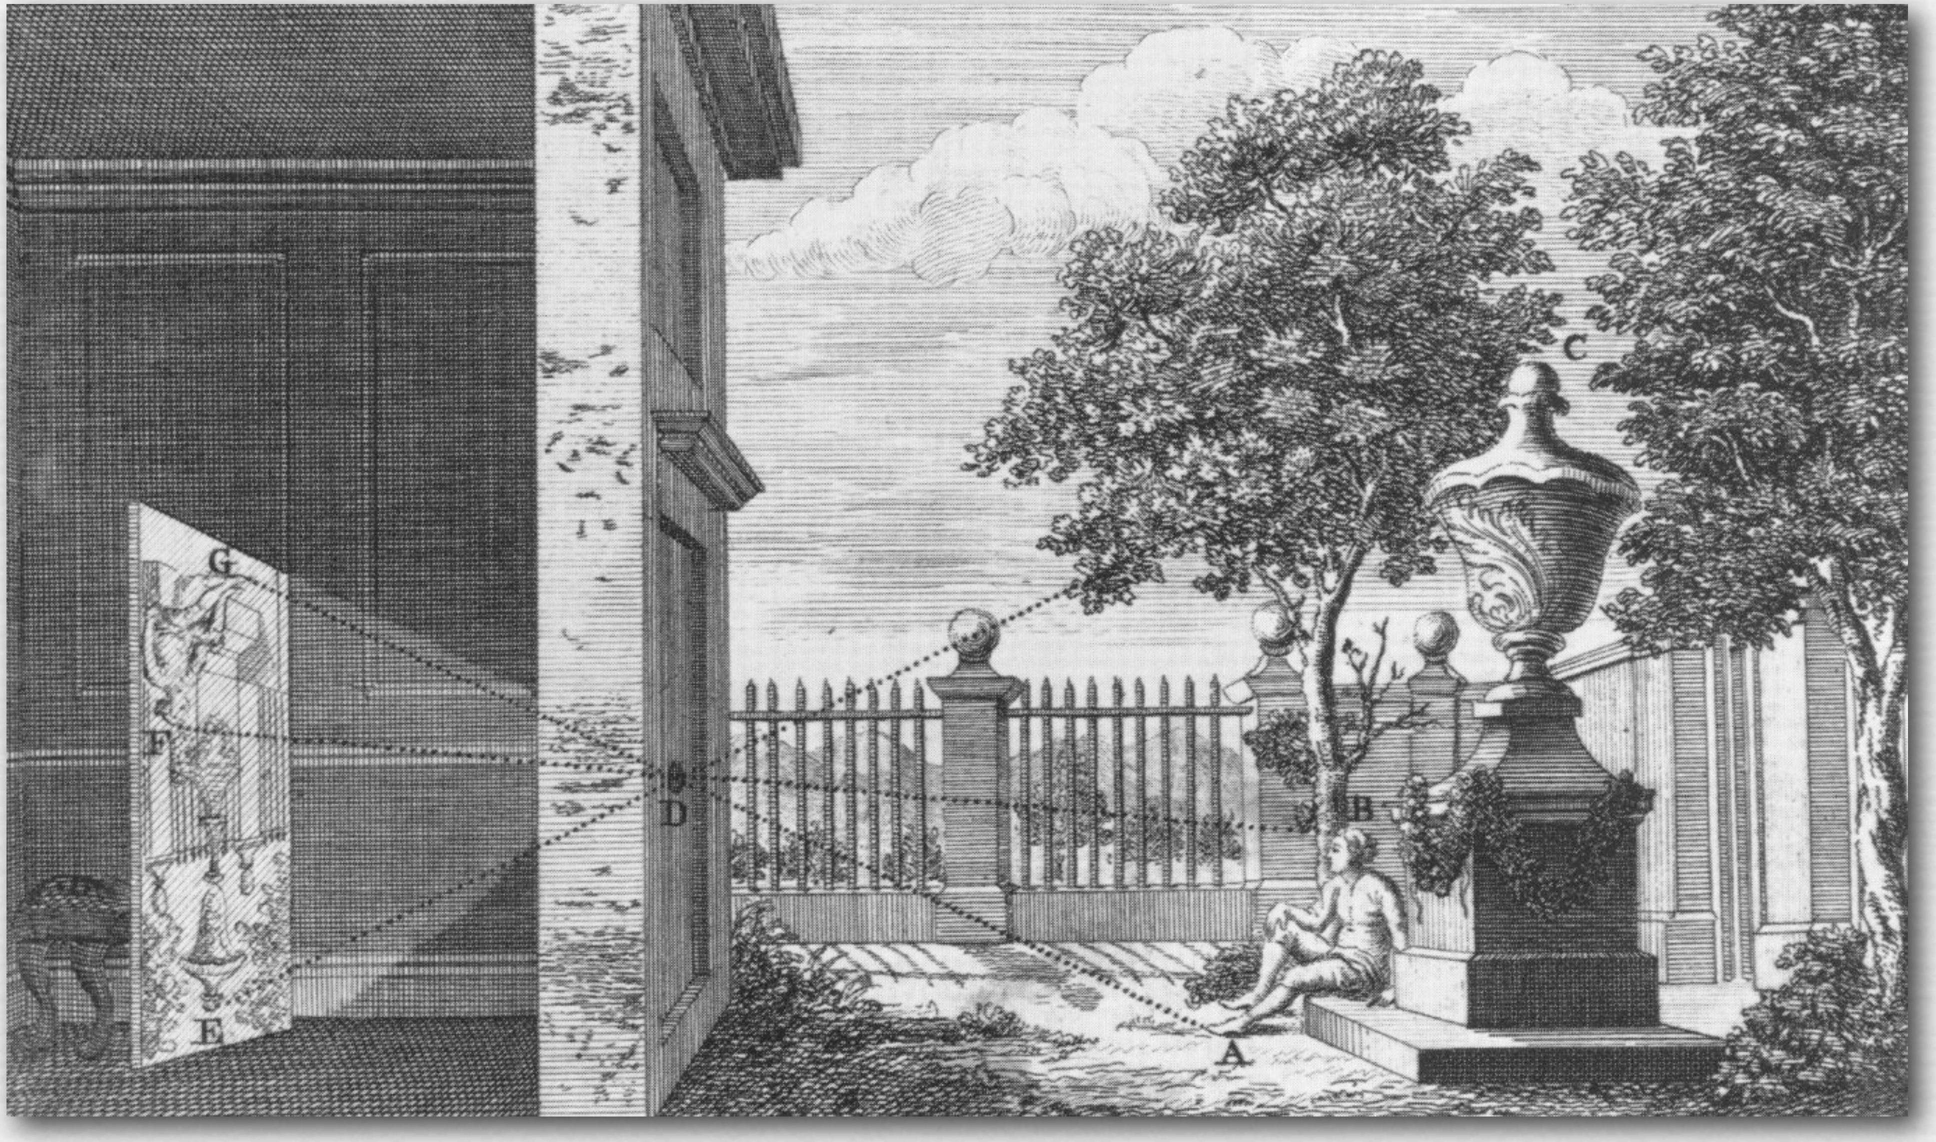

- ✦ linear perspective with viewpoint at pinhole

Pinhole photography

- no distortion
  - straight lines remain straight
- infinite depth of field
  - everything is in focus

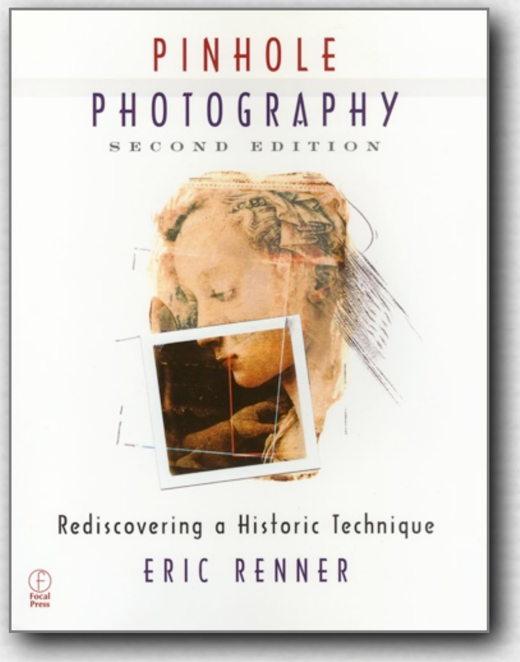
PINHOLE
PHOTOGRAPHY
SECOND EDITION
Rediscovering a Historic Technique
ERIC RENNER


(Bami Adedoyin)

Effect of pinhole size

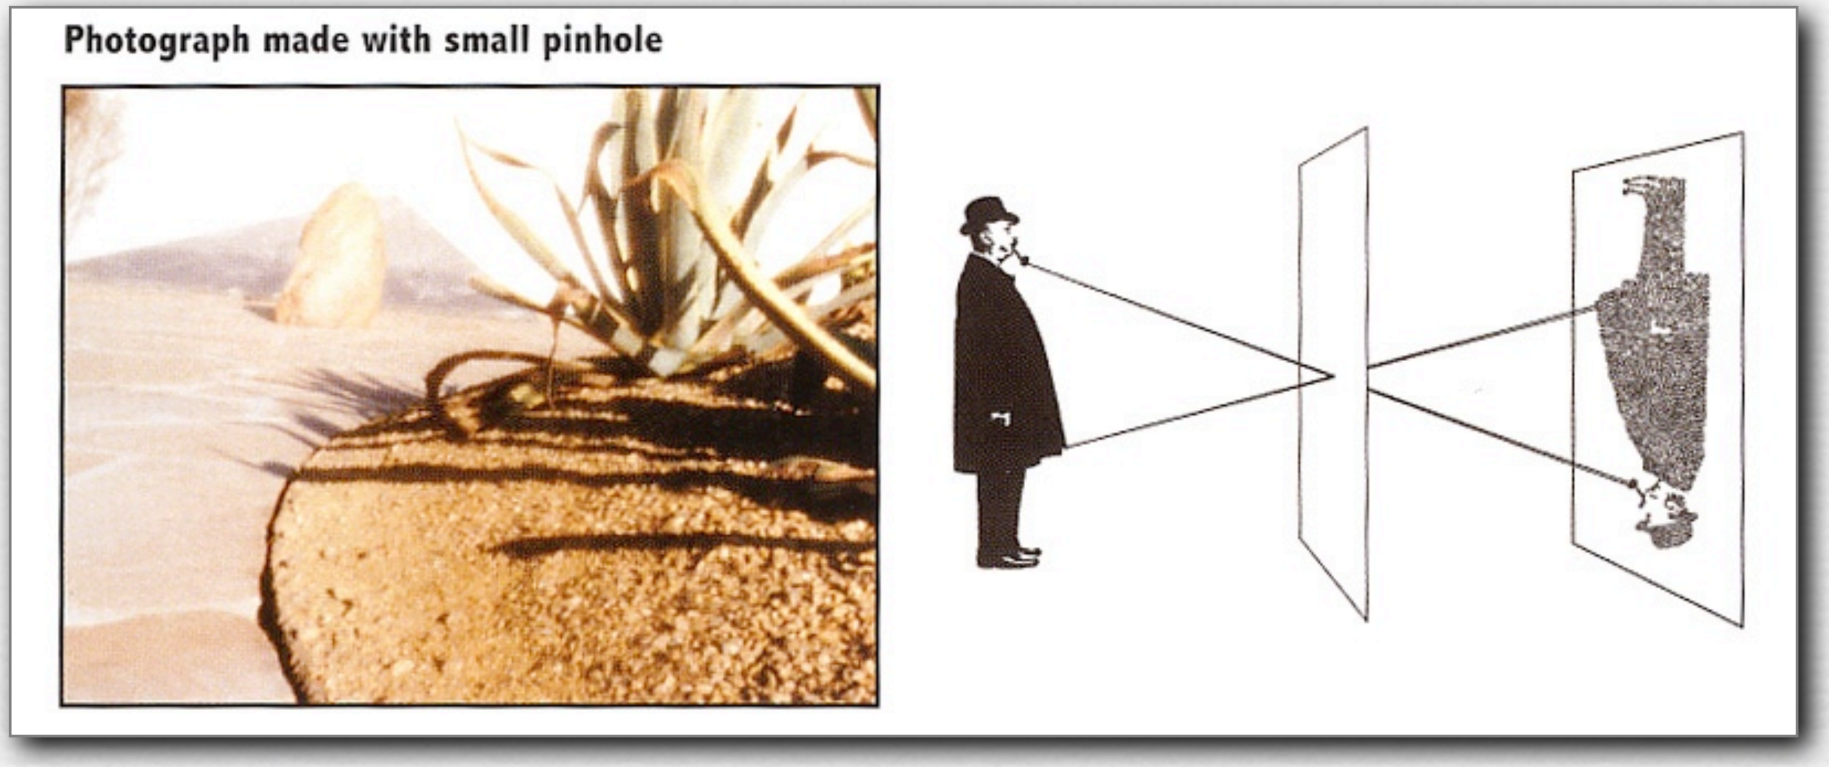
Photograph made with small pinhole

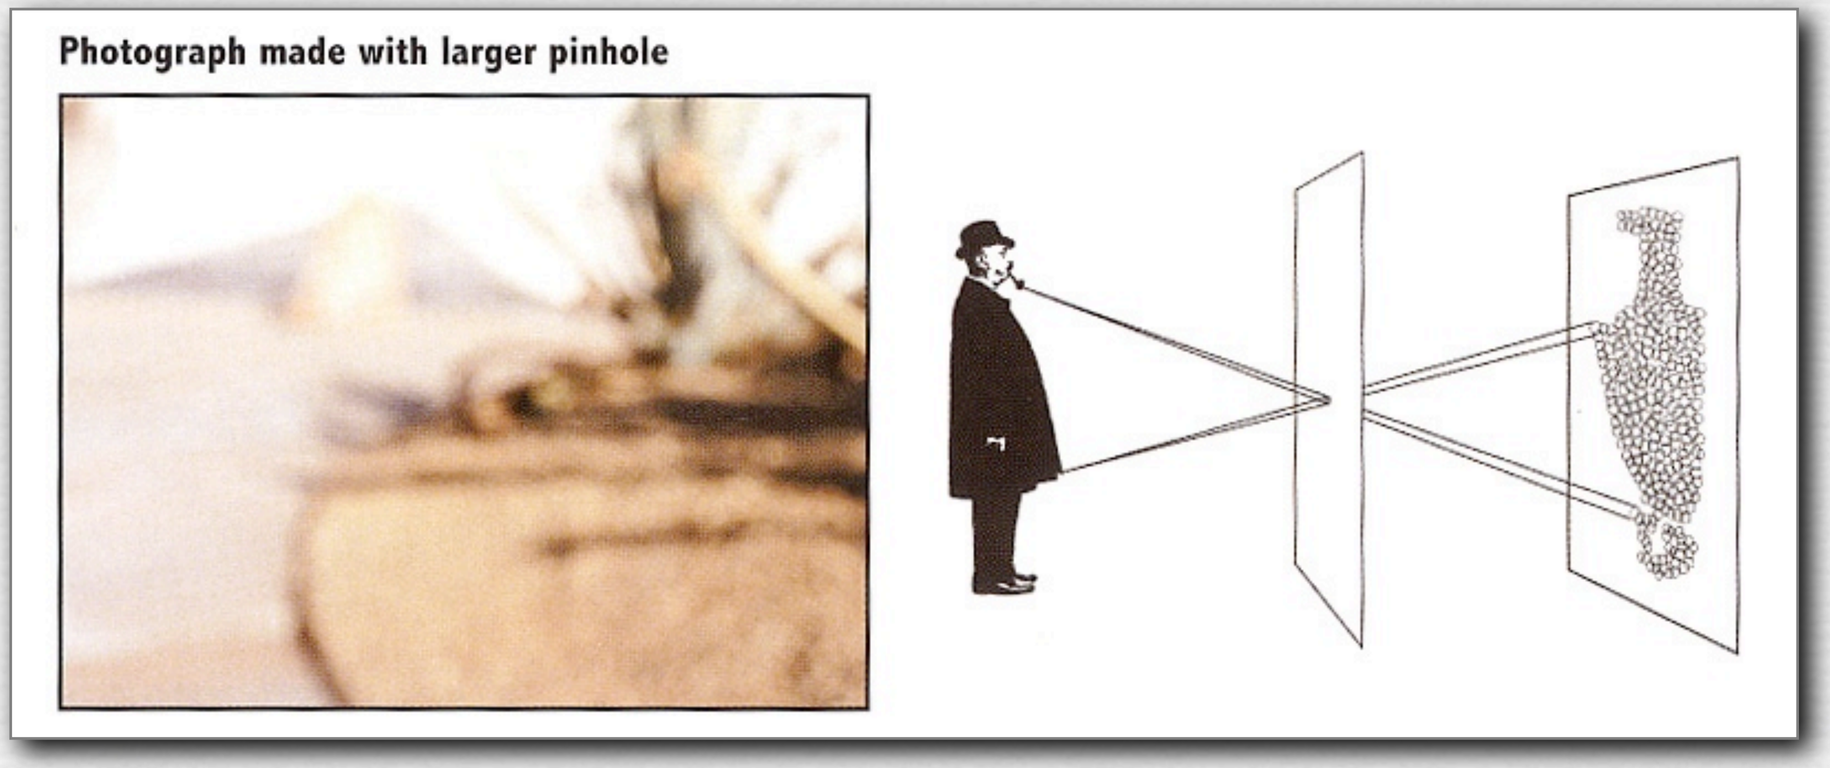
Photograph made with larger pinhole

(London)

Effect of pinhole size

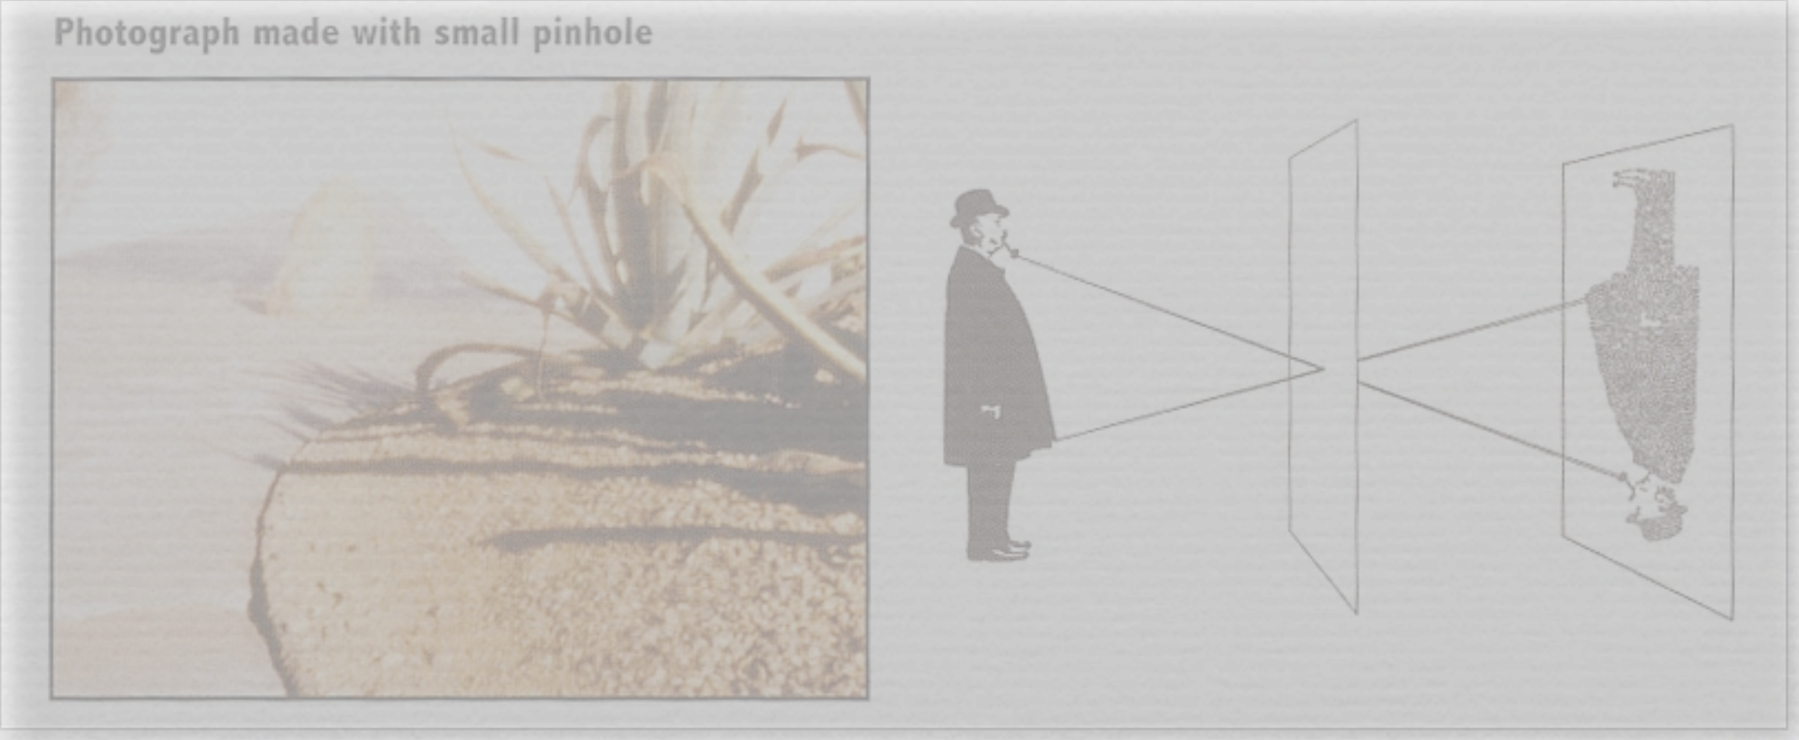

Photograph made with small pinhole

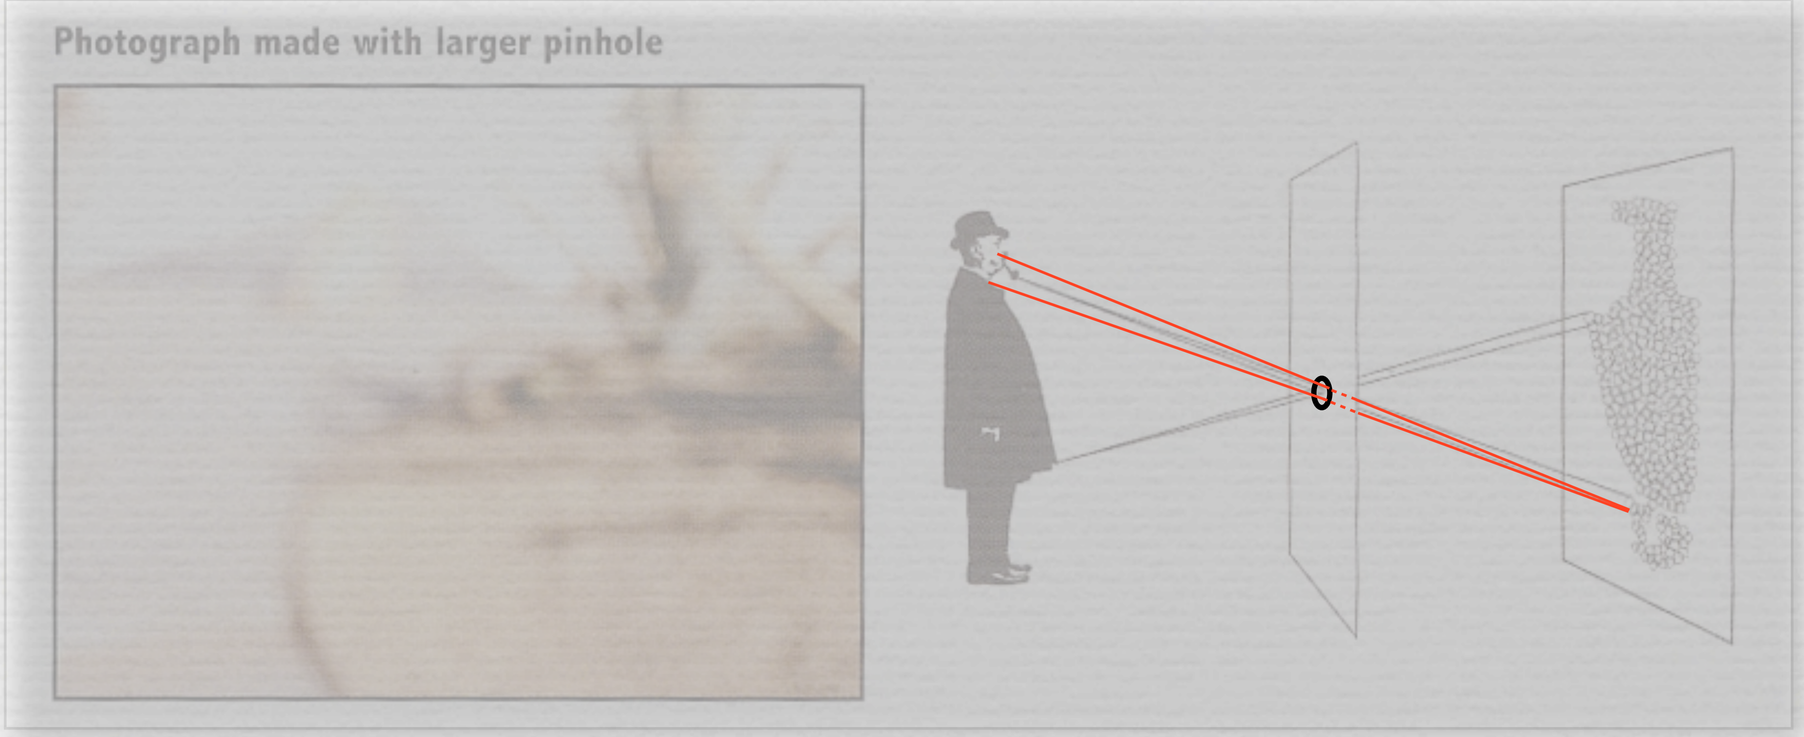

Photograph made with larger pinhole

(London)

Replacing the pinhole with a lens

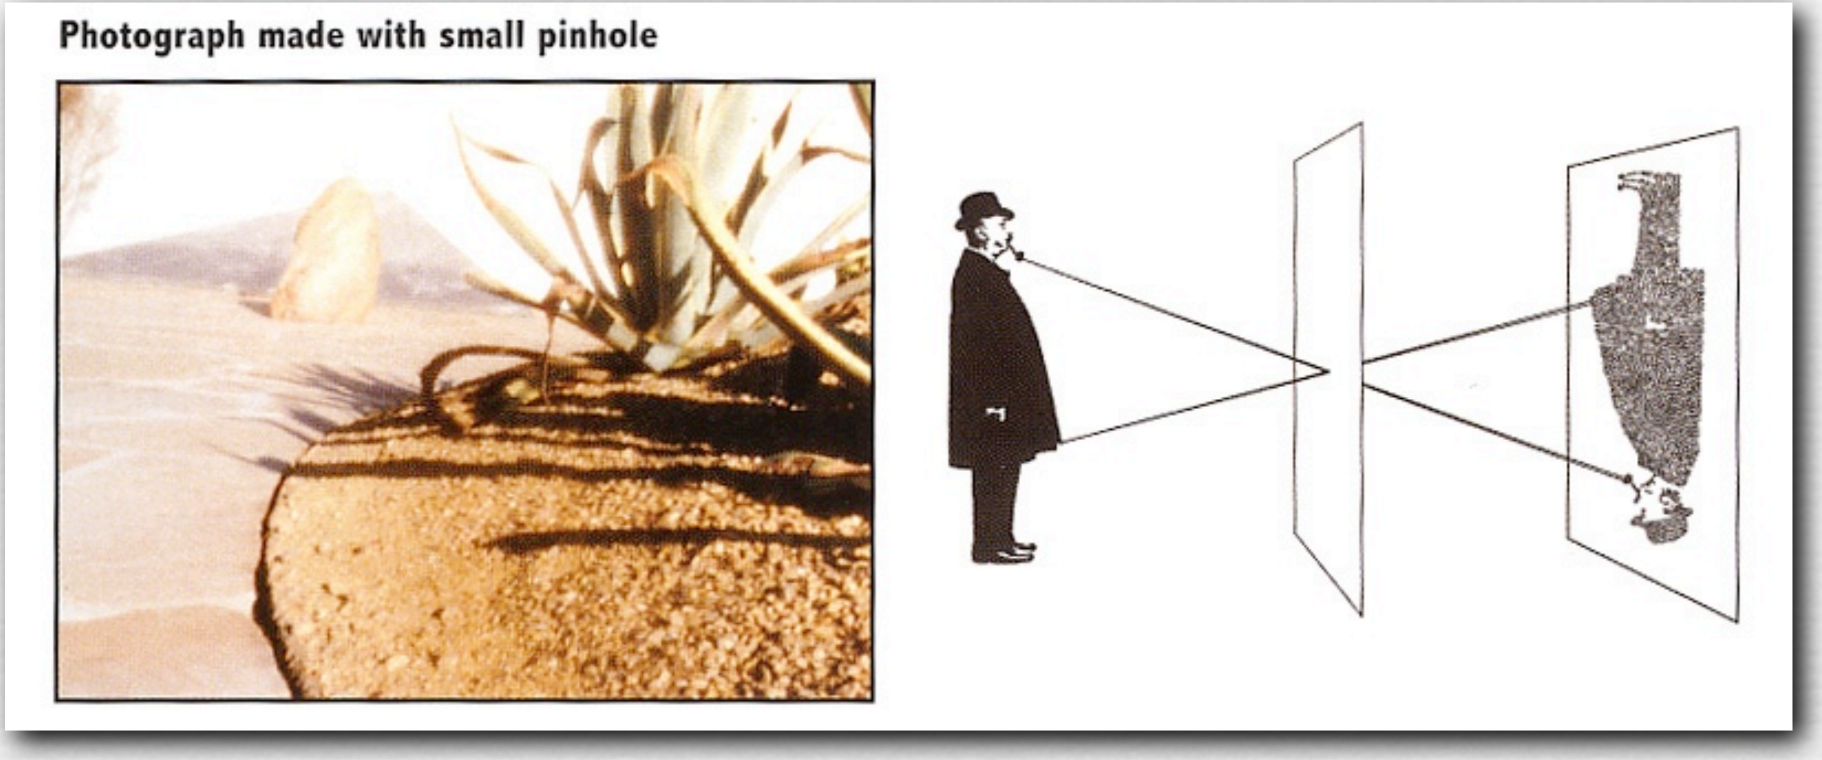

Photograph made with small pinhole

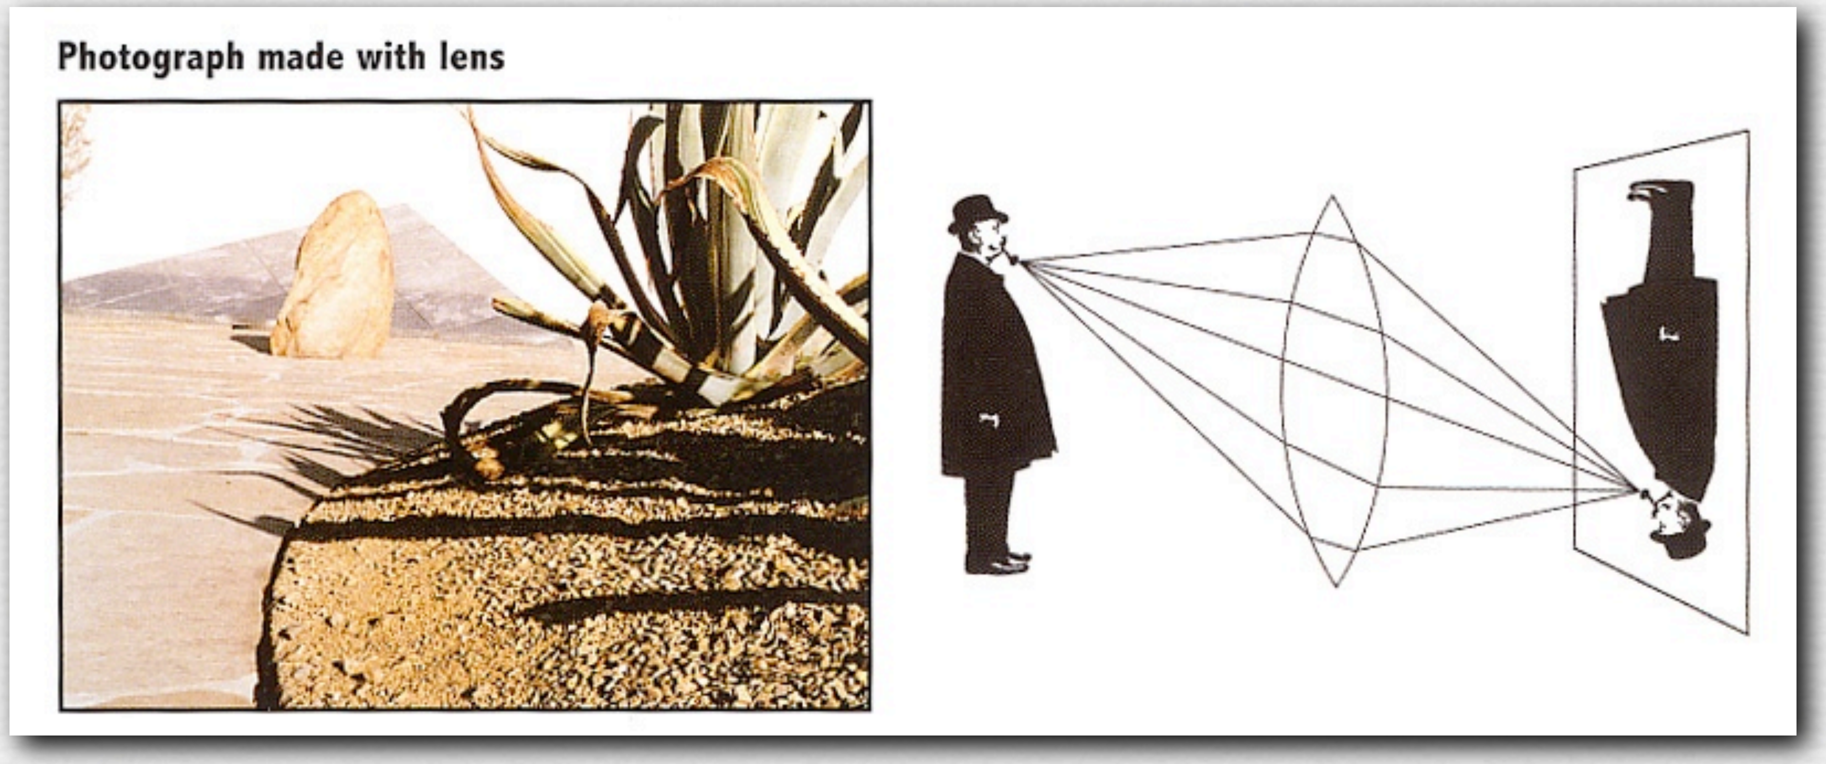

Photograph made with lens

(London)

Snell's law of refraction

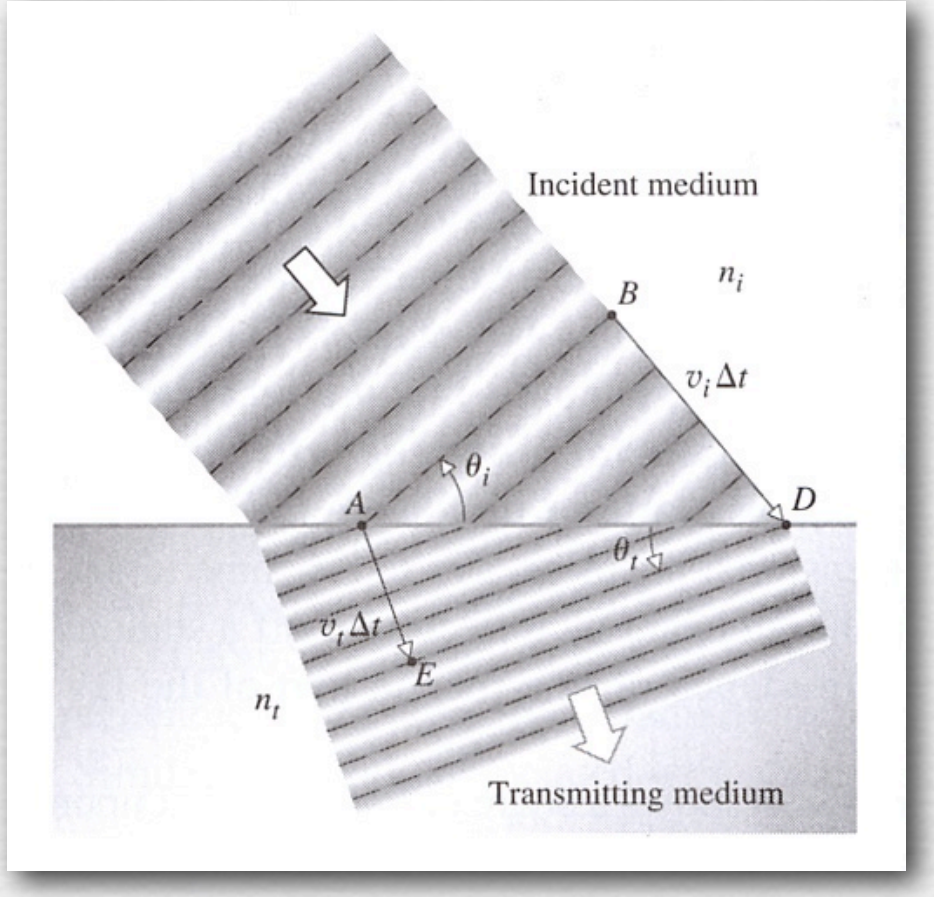

(Hecht)

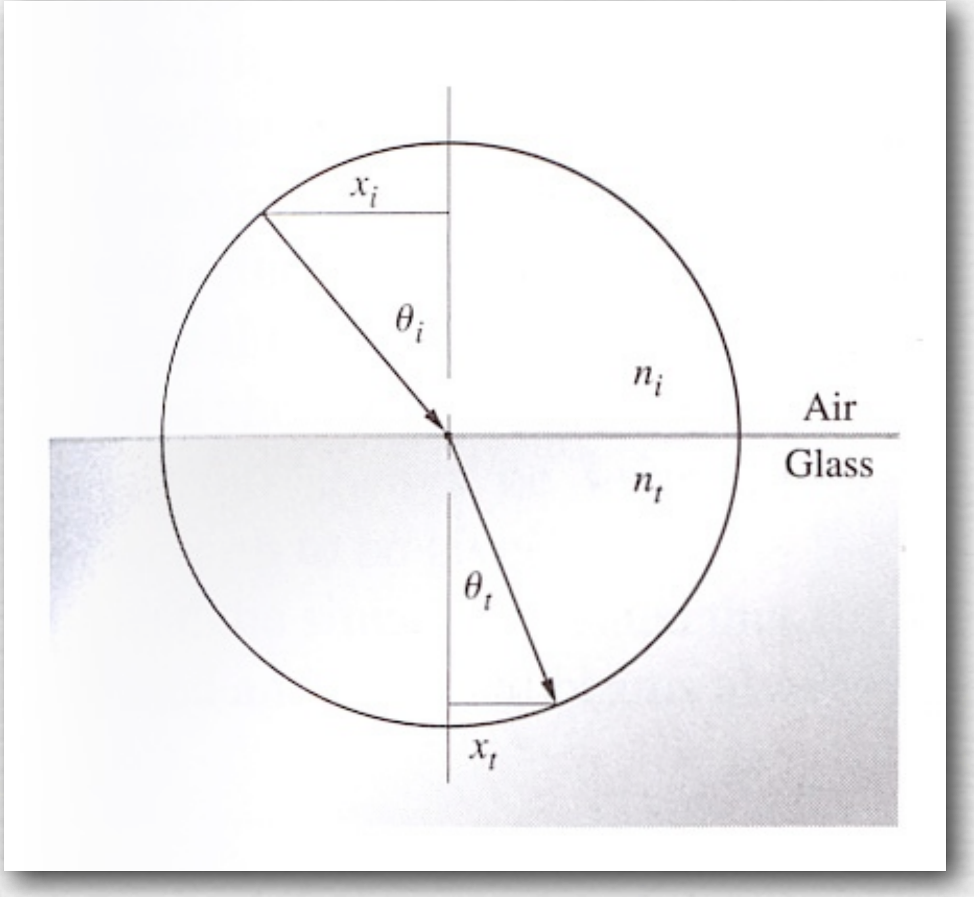

$$\frac{x_i}{x_t} = \frac{\sin\theta_i}{\sin\theta_t} = \frac{n_t}{n_i}$$

- 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

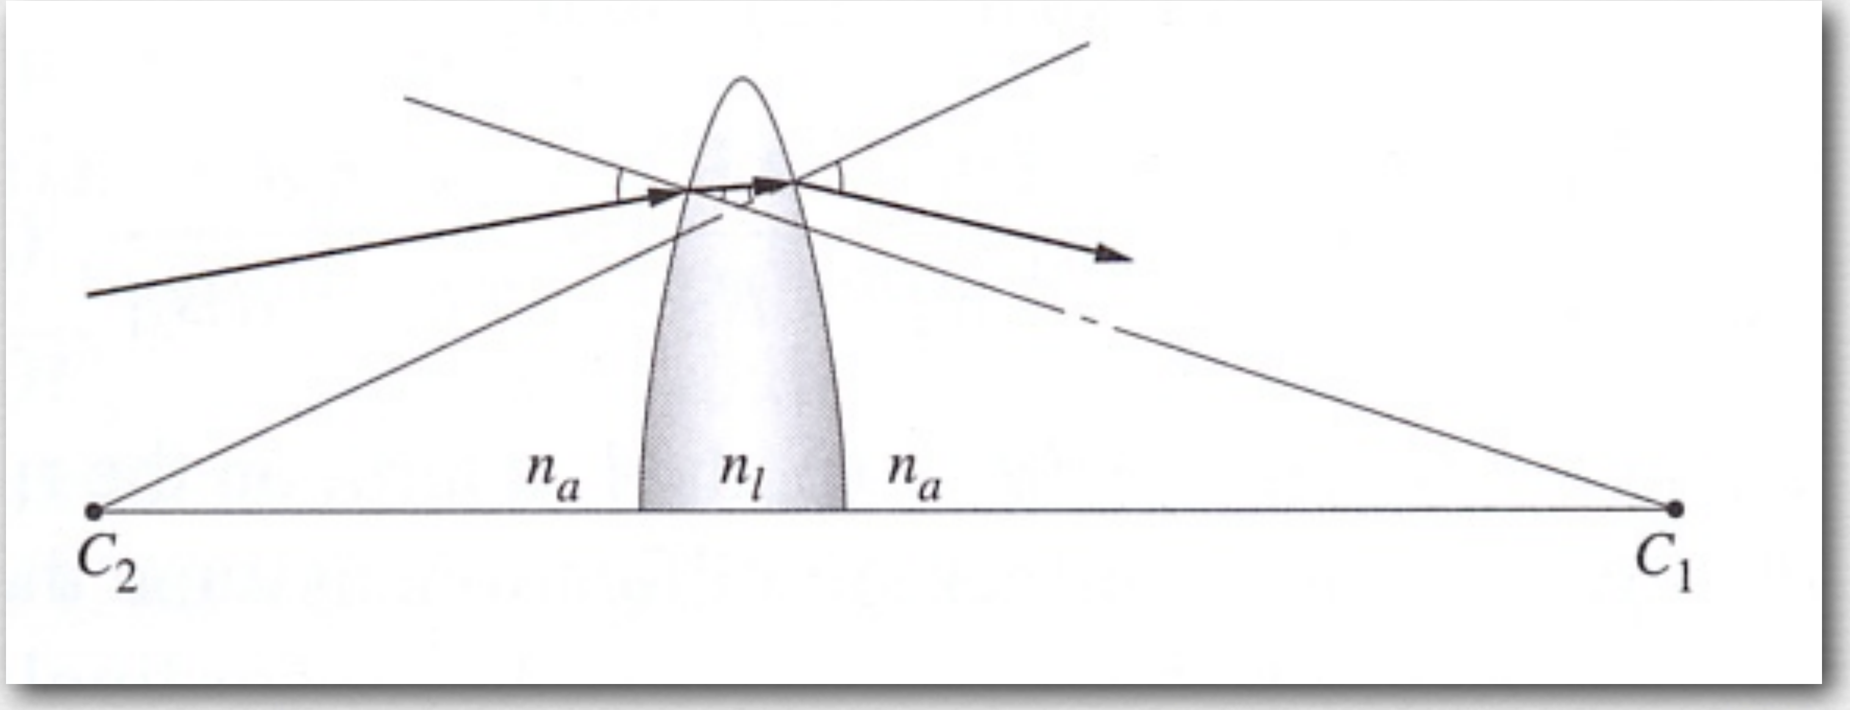

(Hecht)

- 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?

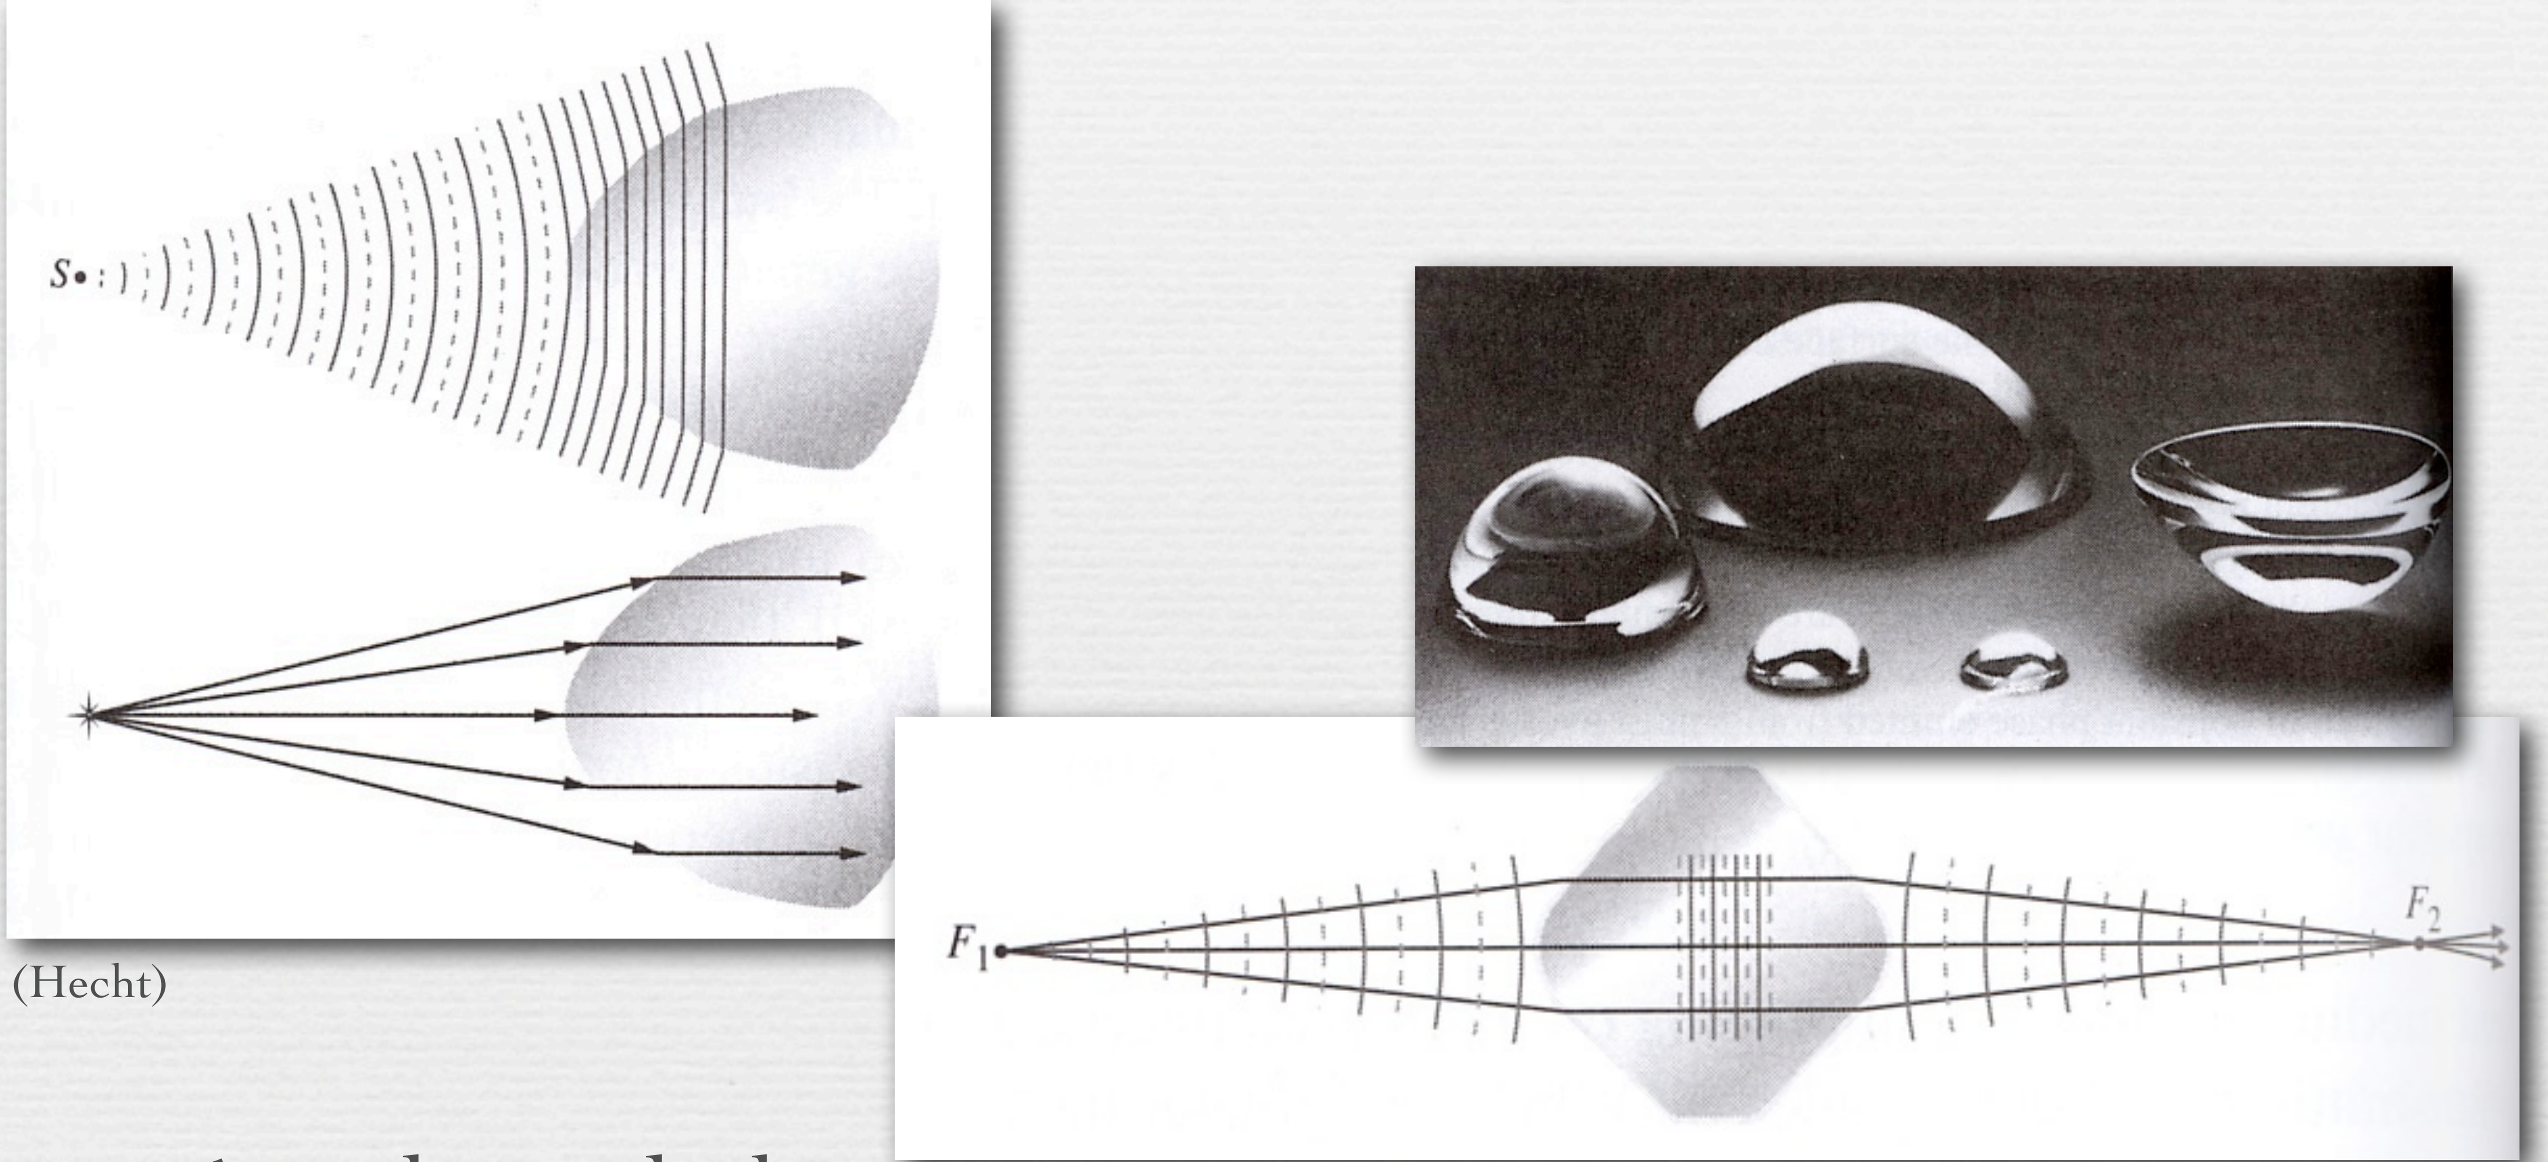

(Hecht)

A. a hyperbola

- so lenses should be hyperbolic!

Spherical lenses

(Hecht)

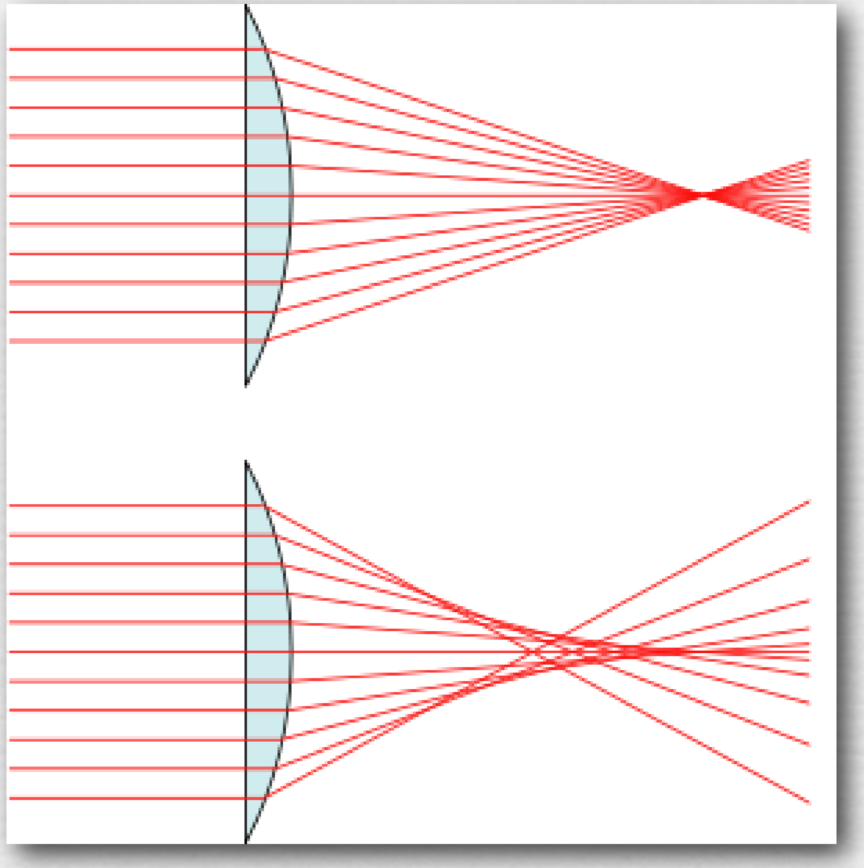

(wikipedia)

- 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


(Canon)

sharp

soft focus

Canon 135mm f/2.8 soft focus lens

Focus shift

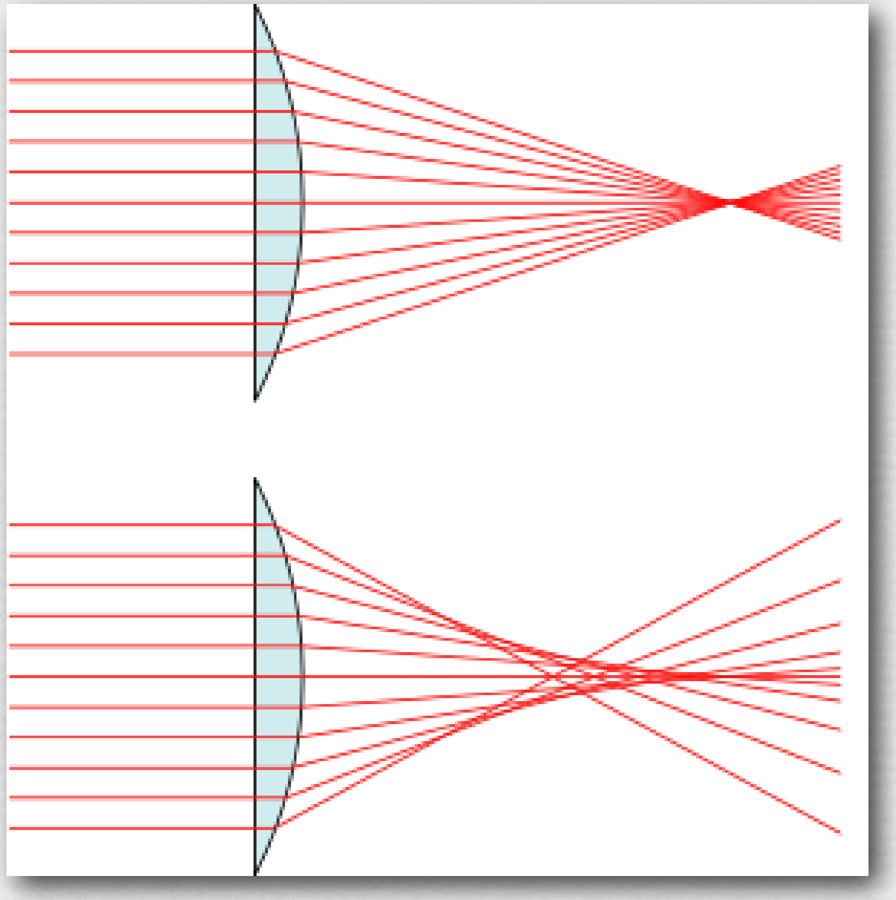

(wikipedia)

(diglloyd.com)

focused at f/1.2

- Canon 50mm f/1.2 L

Focus shift

(diglloyd.com)

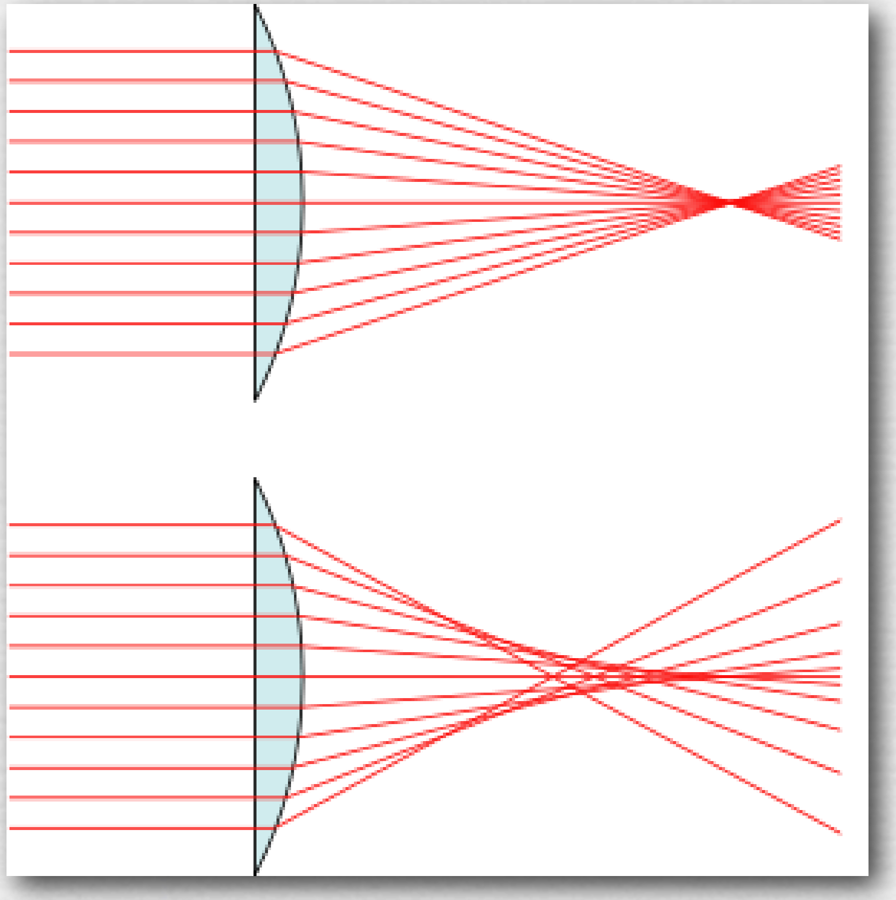

(wikipedia)

shot at f/1.8

- Canon 50mm f/1.2 L
- narrowing the aperture pushed the focus deeper

Paraxial approximation

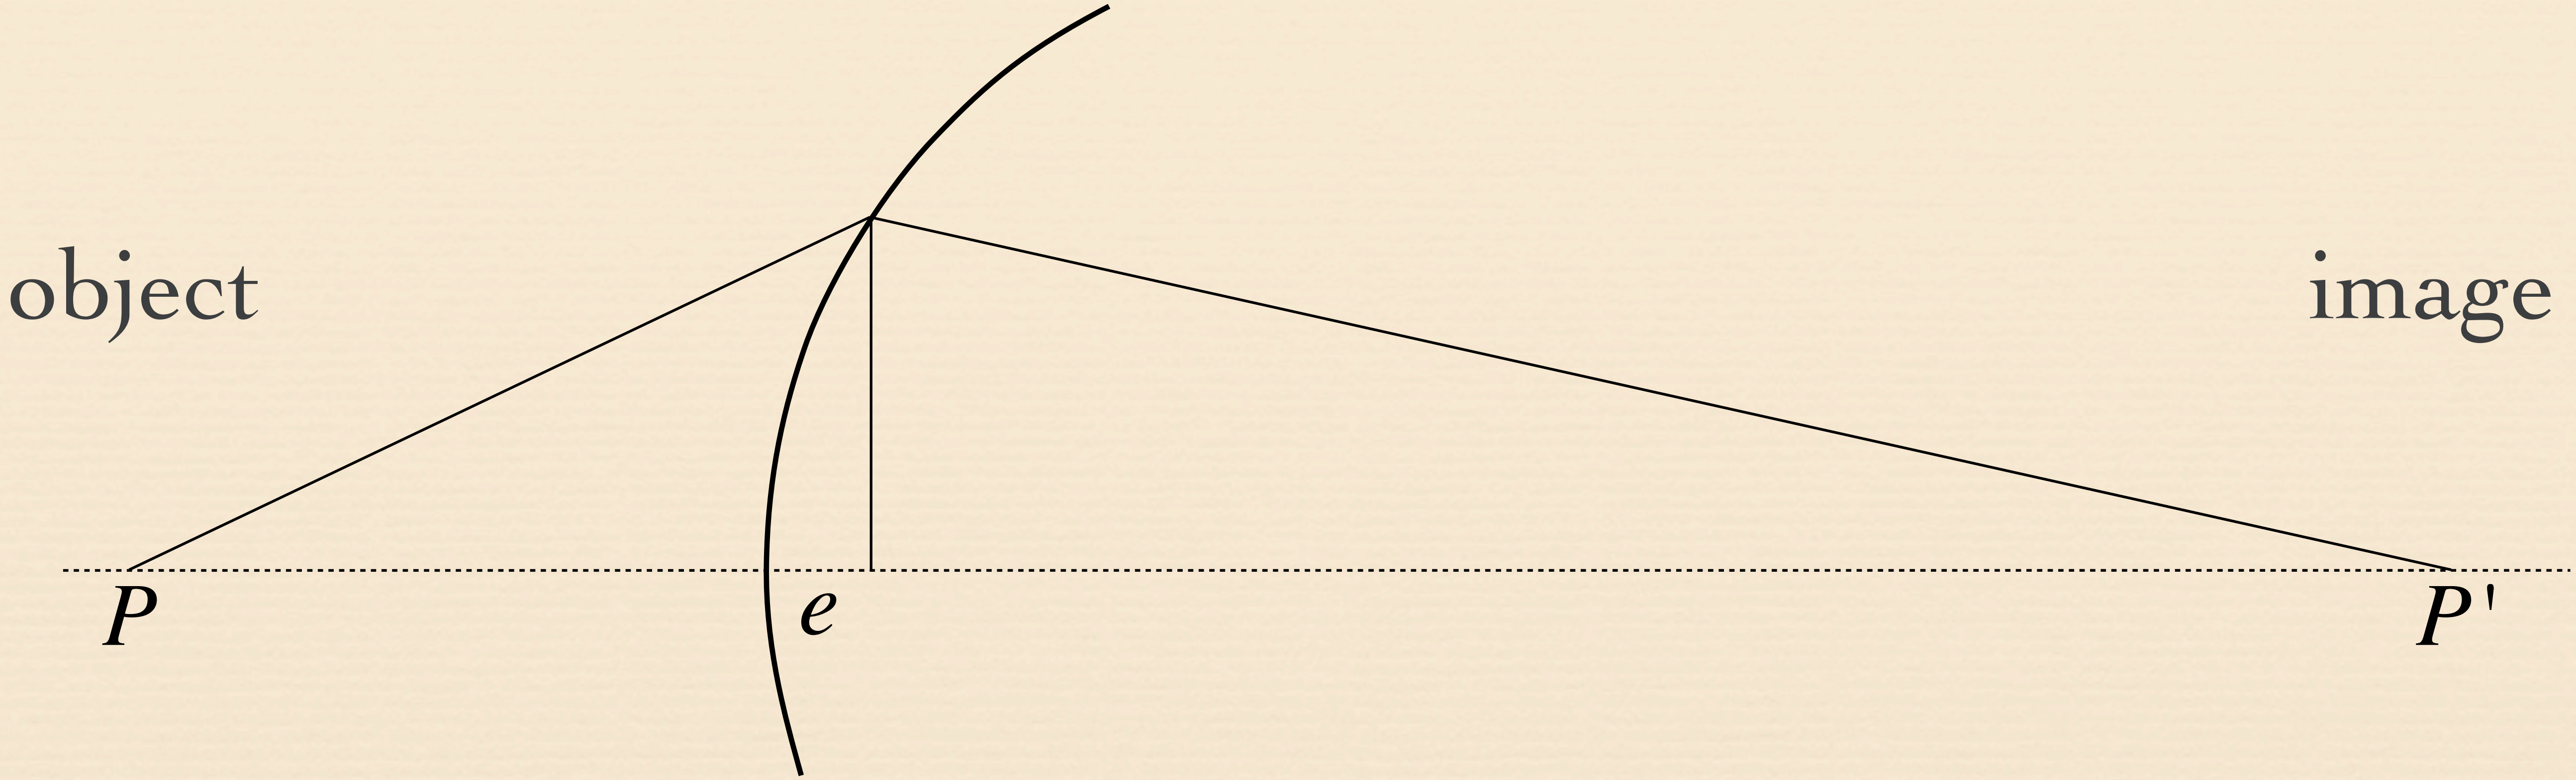

- assume $e \approx 0$

Paraxial approximation

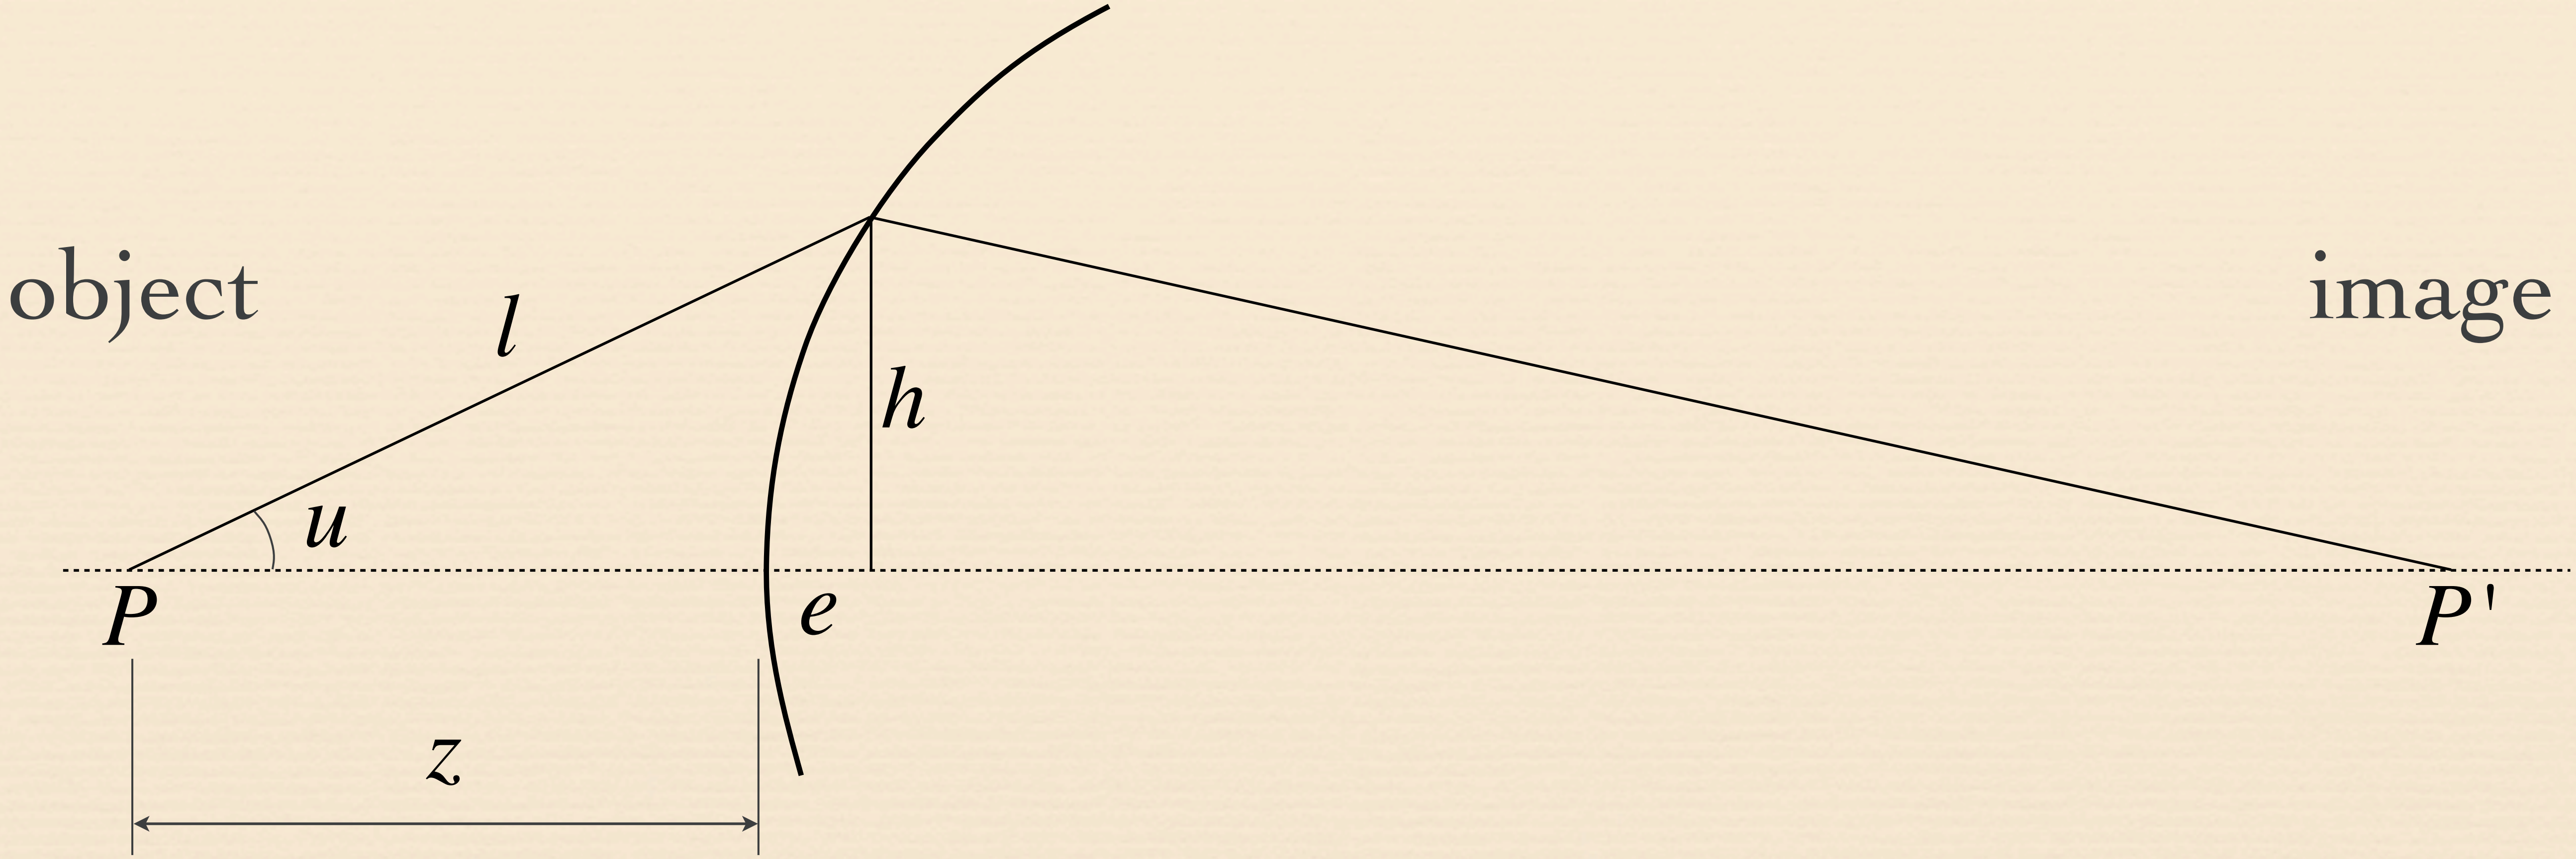

- assume $e \approx 0$
- assume $\sin u = h / l \approx u$ (for u in radians)
- assume $\cos u \approx z / l \approx 1$
- assume $\tan u \approx \sin u \approx u$

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

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

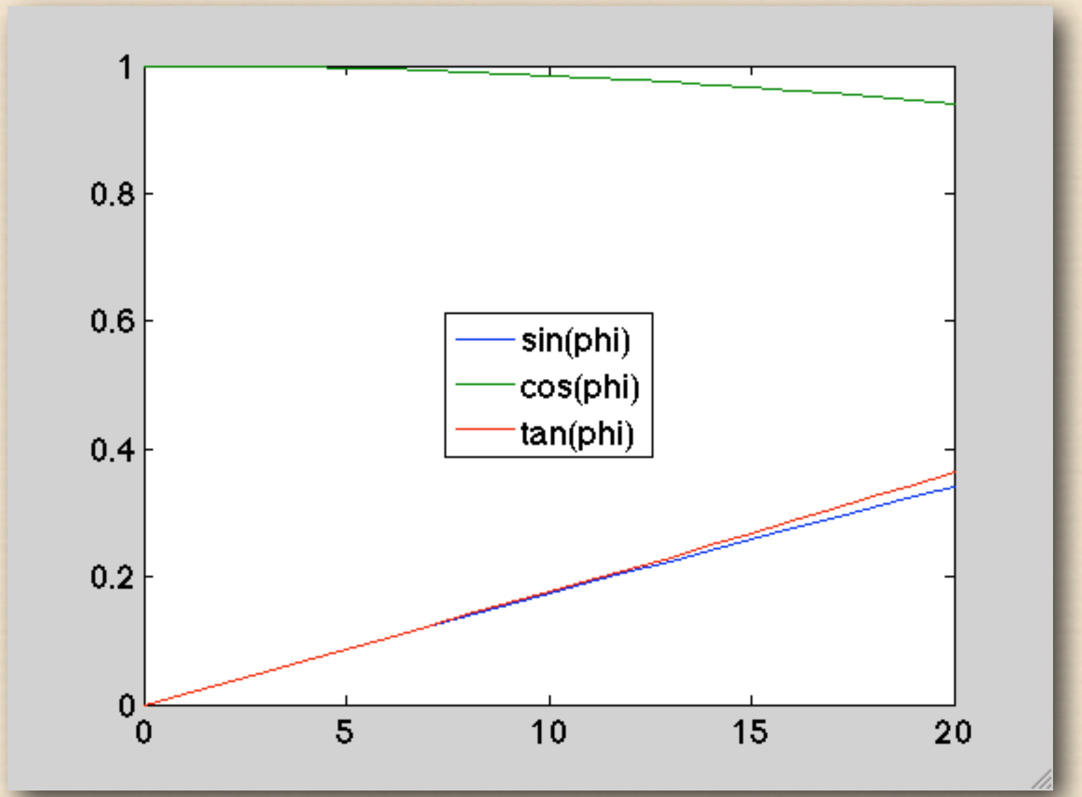

Paraxial focusing

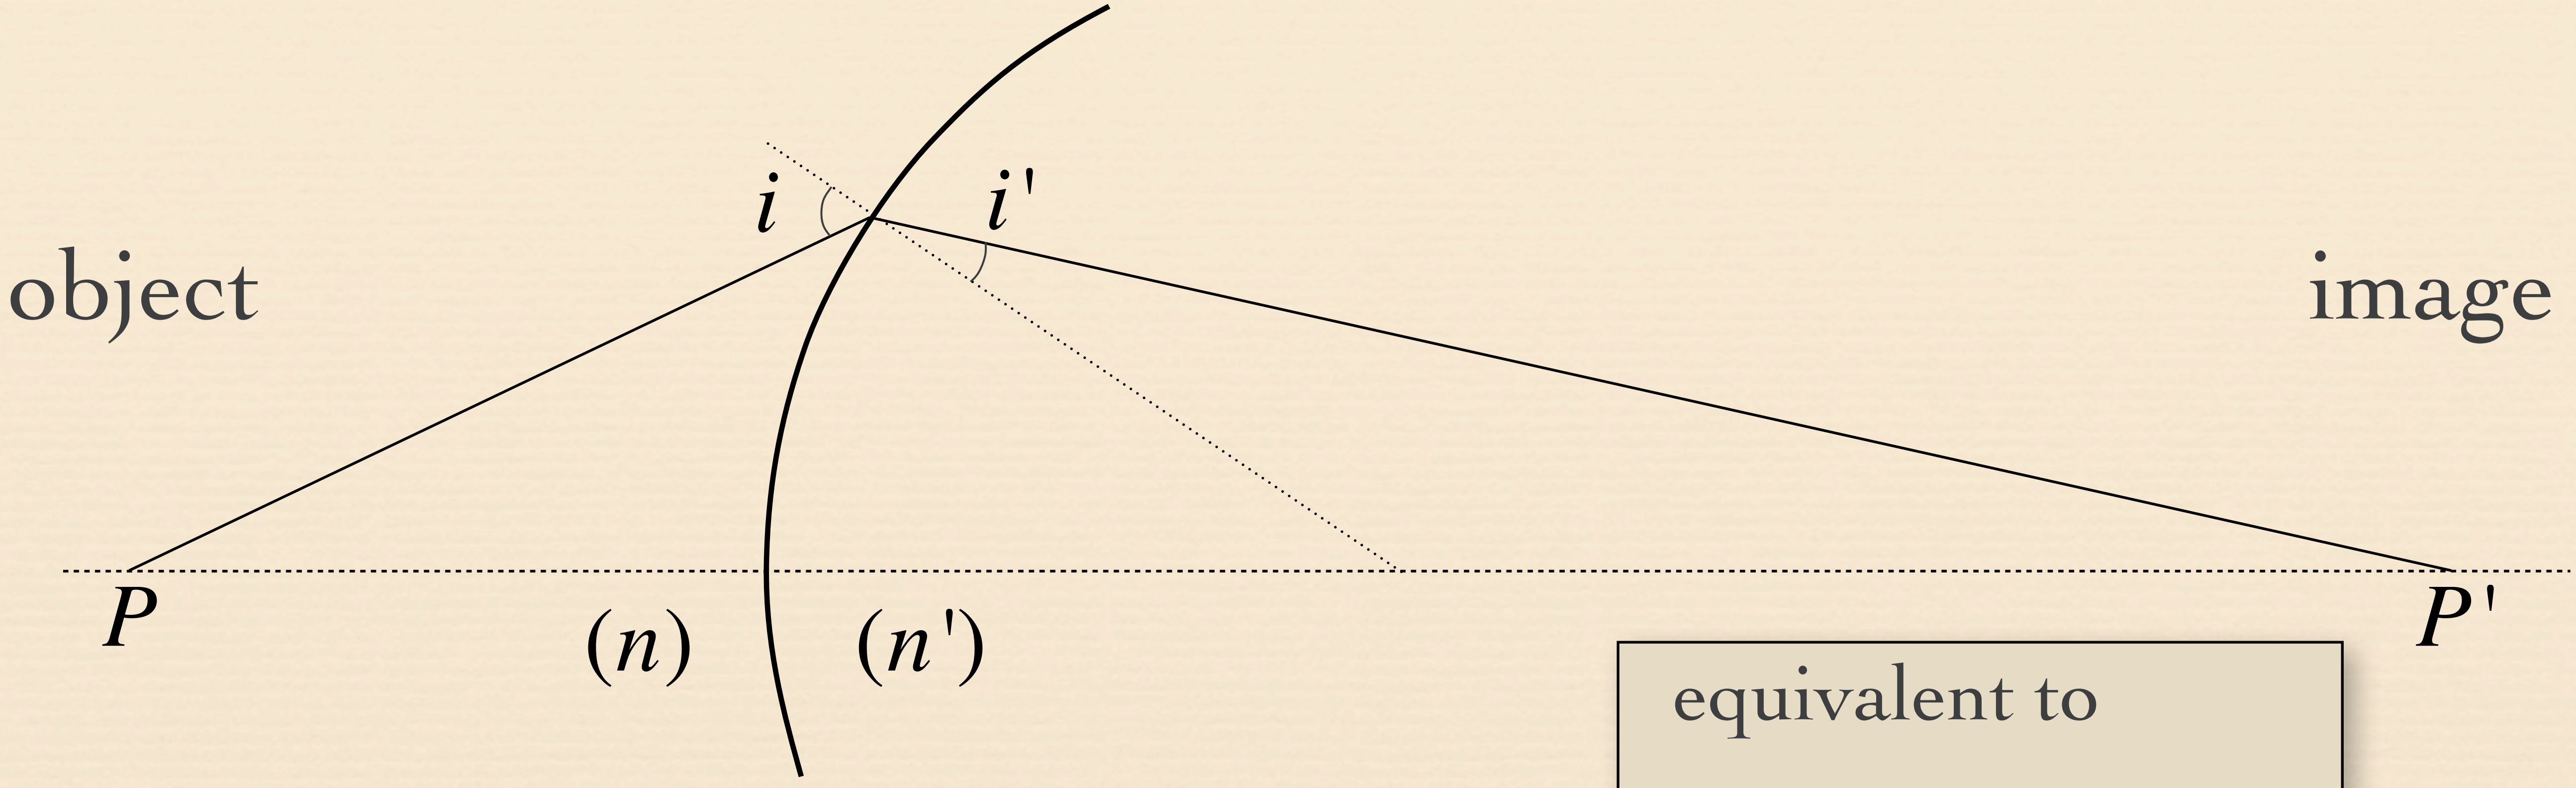

Snell's law:

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

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

$i = u + a$

$u \approx h / z$

$u' \approx h / z'$

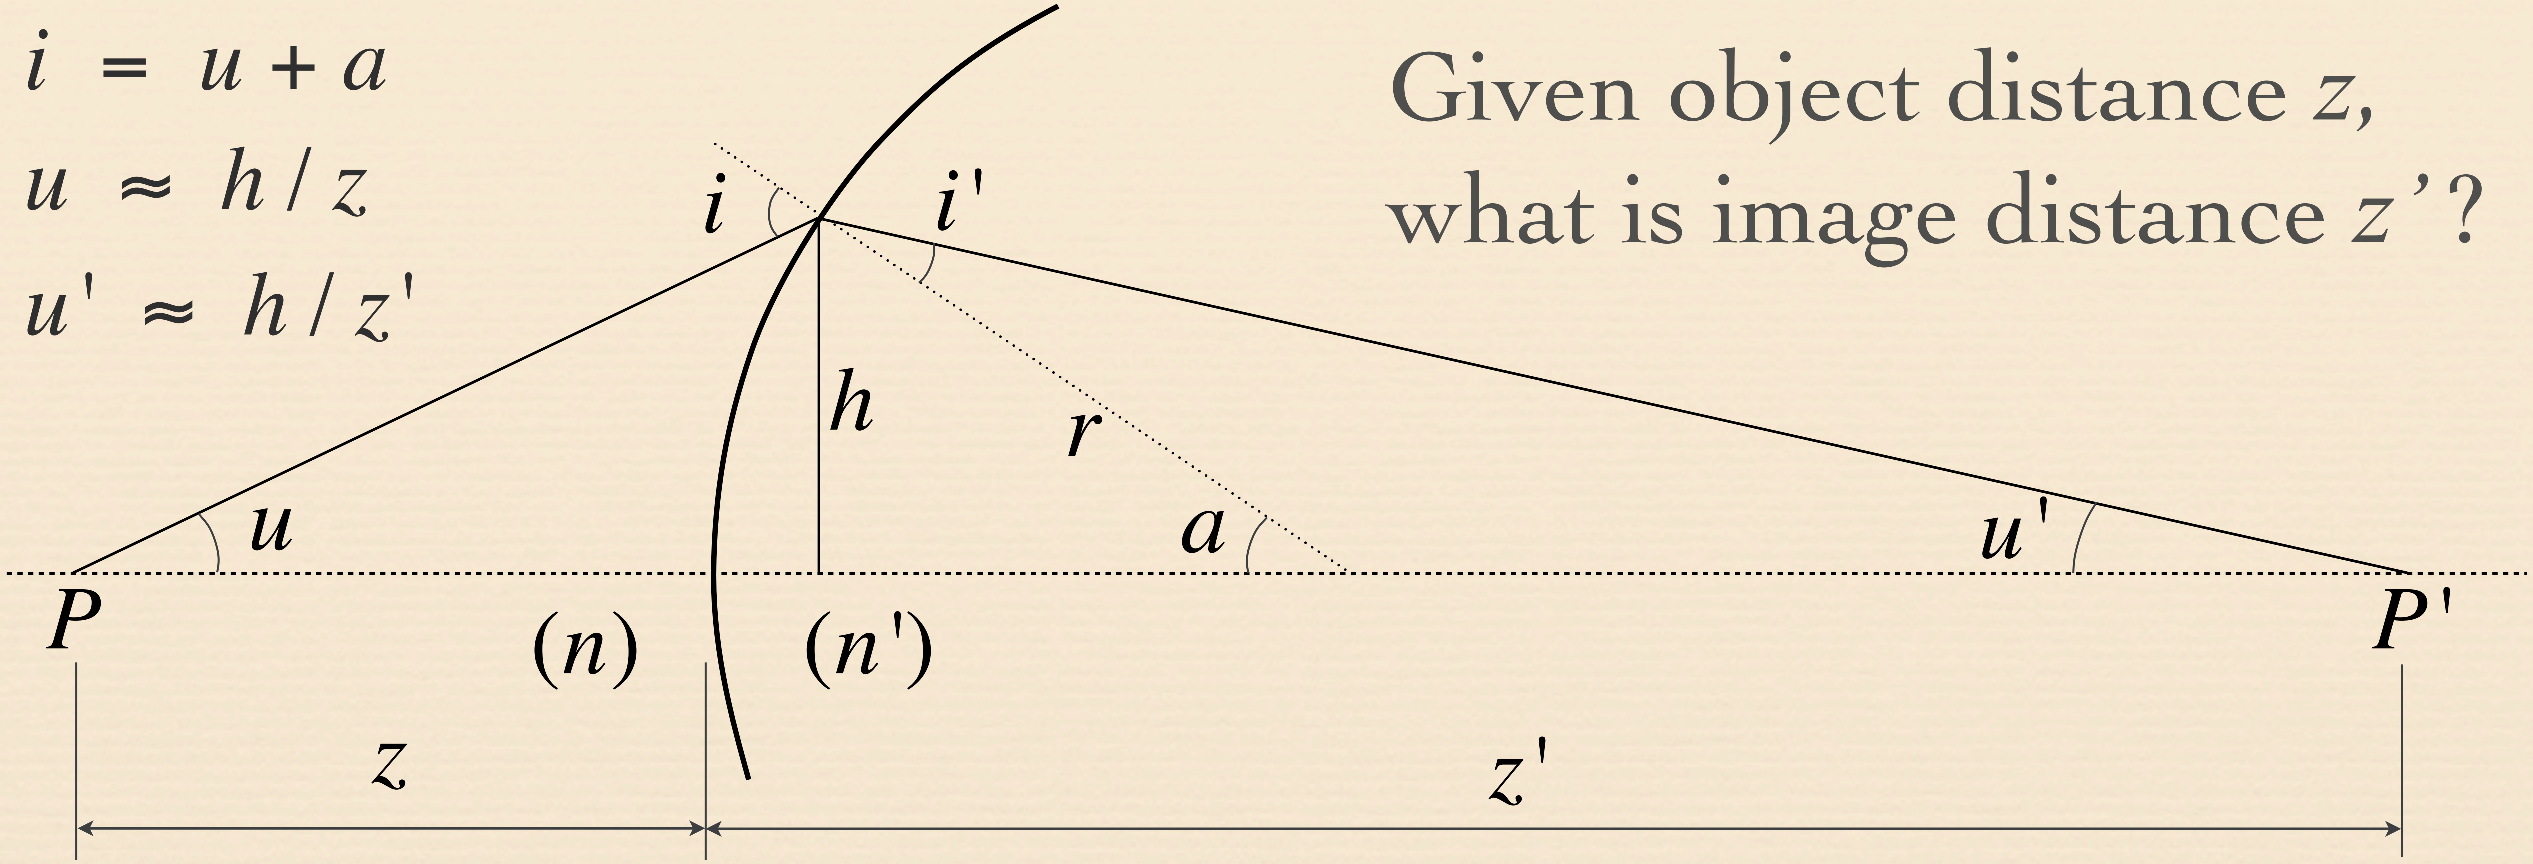

Given object distance z, what is image distance z'?

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

Paraxial focusing

$$i = u + a$$

$$u \approx h / z$$

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

$$a = u' + i'$$

$$a \approx h / r$$

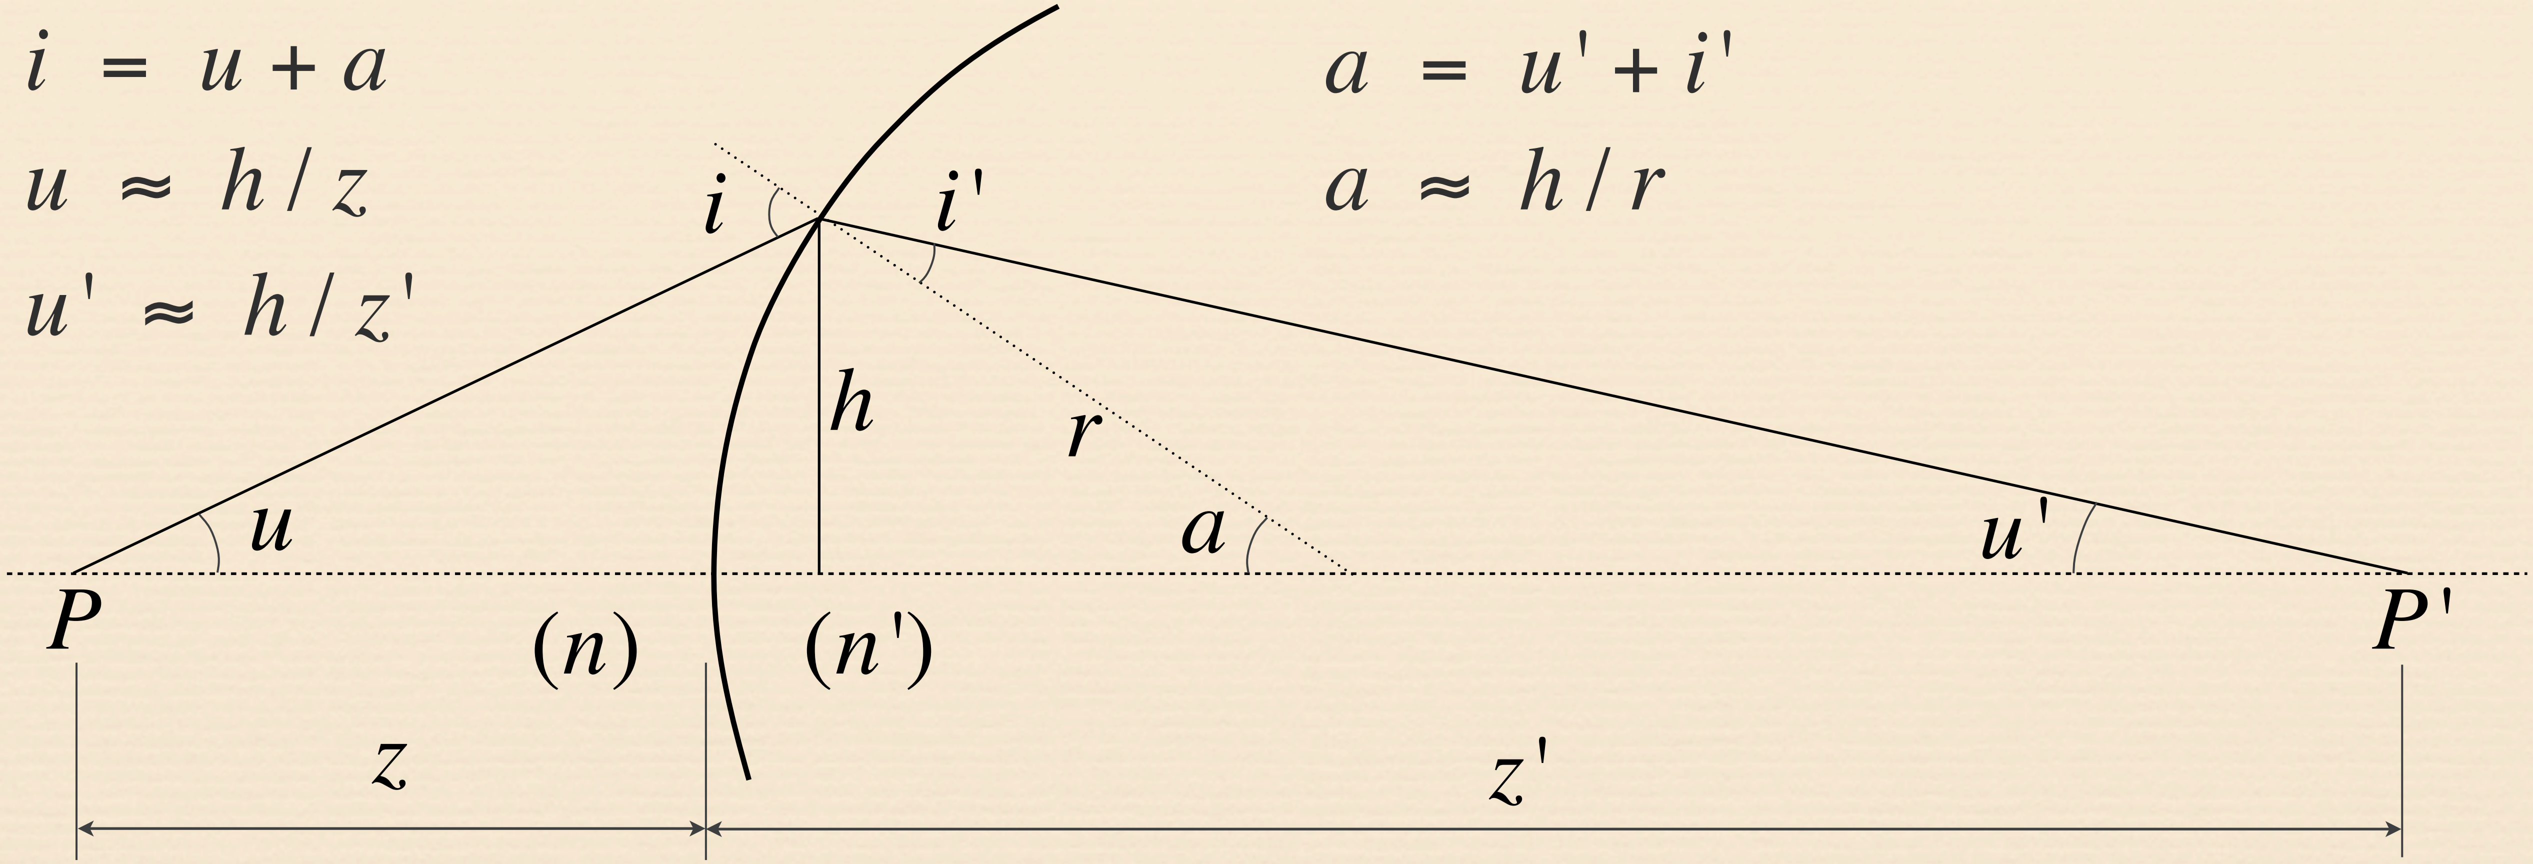

$$n\,(u + a) \approx n'\,(a - u')$$

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

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

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

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

Focal length

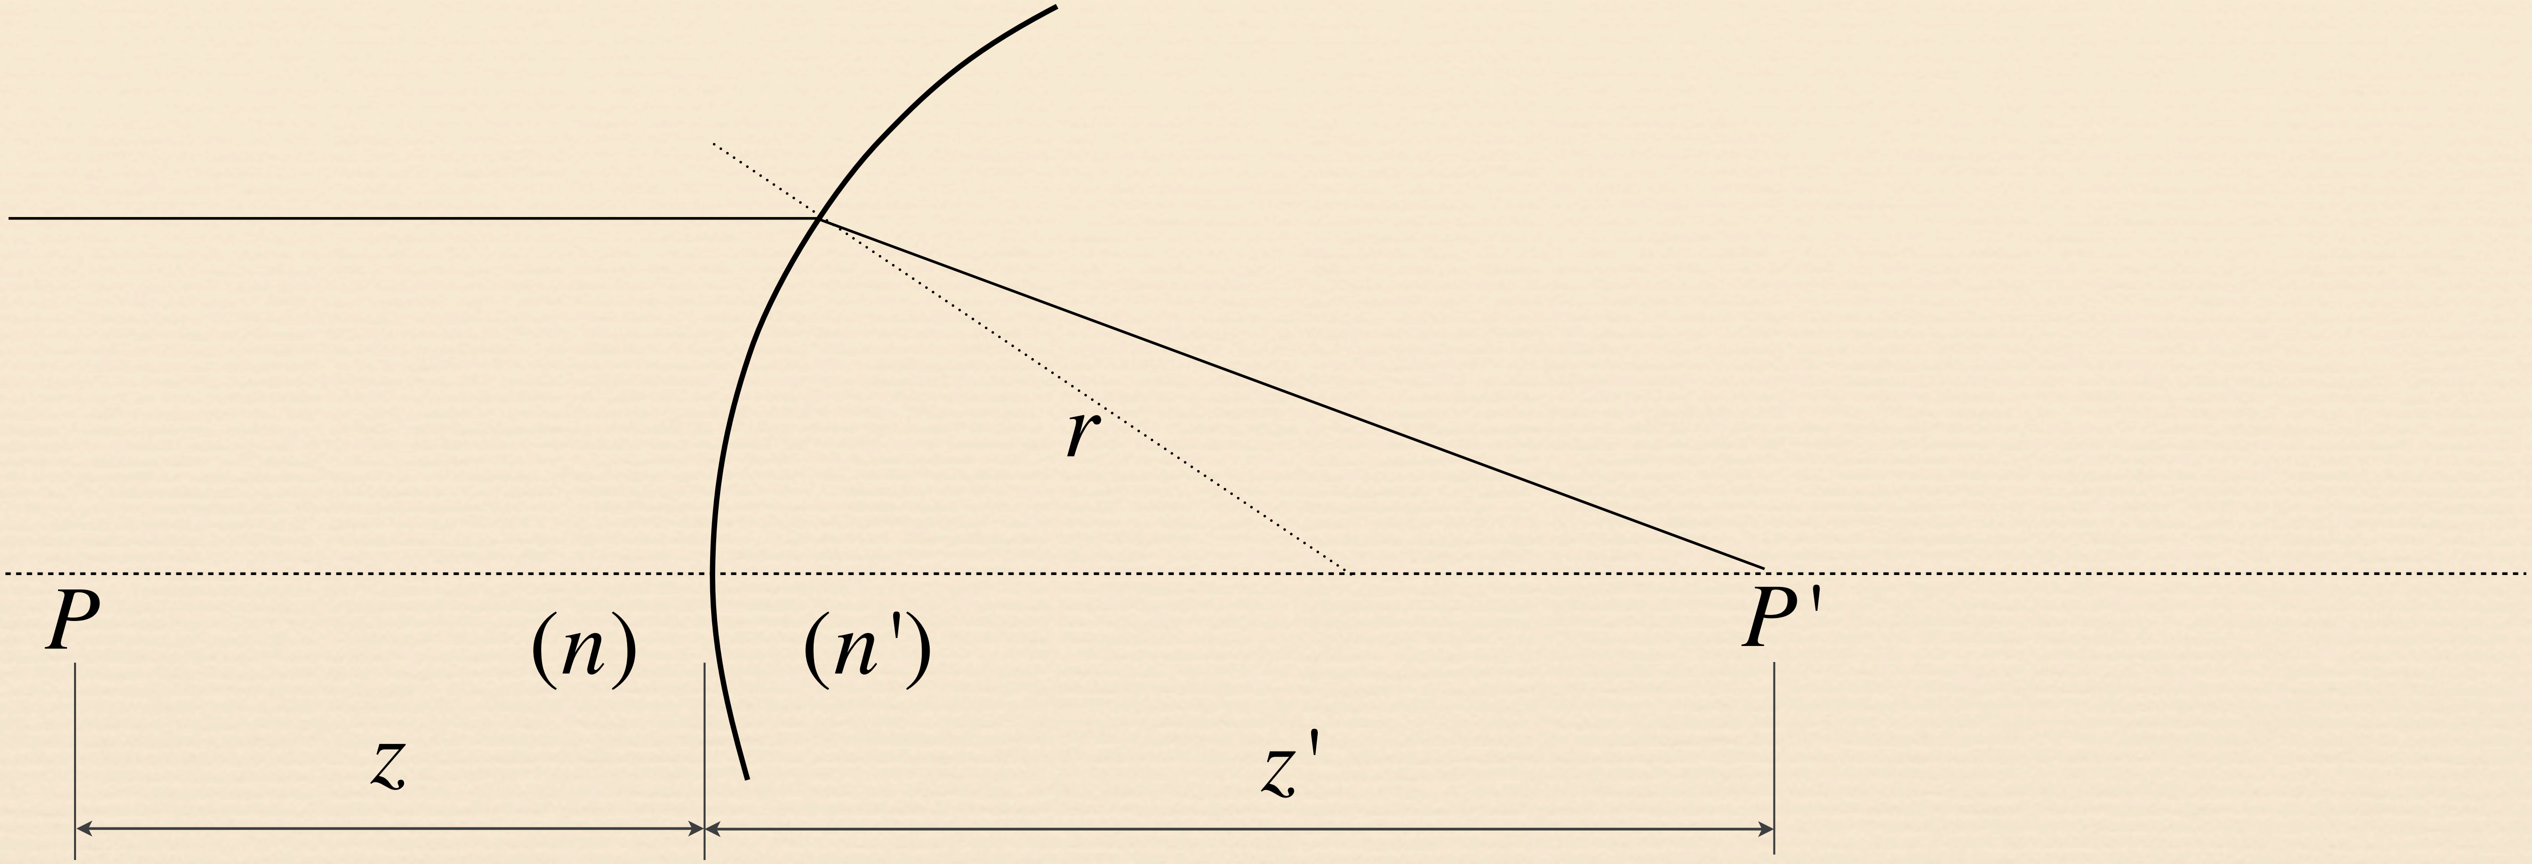

What happens if z is ∞ ?

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

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

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

- $f \triangleq$ focal length = z'

Lensmaker's formula

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

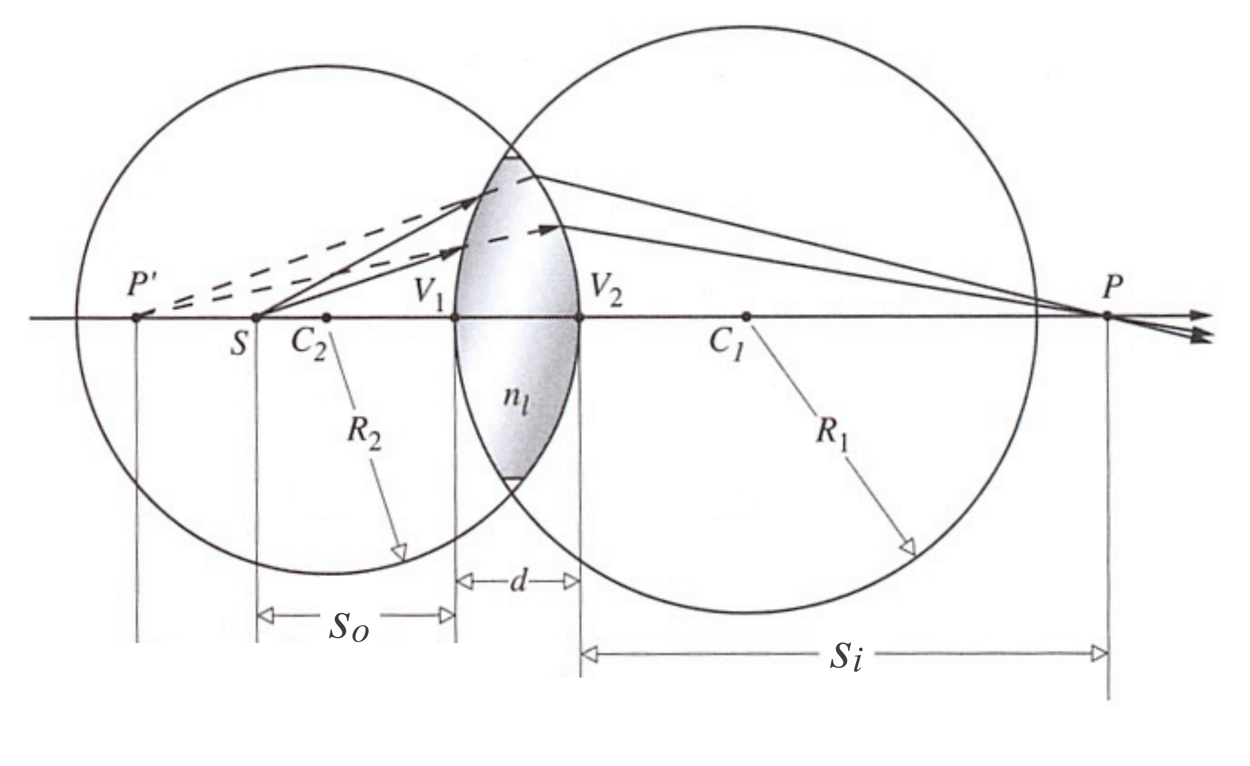

(Hecht, edited)

- as $d \to 0$ *(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

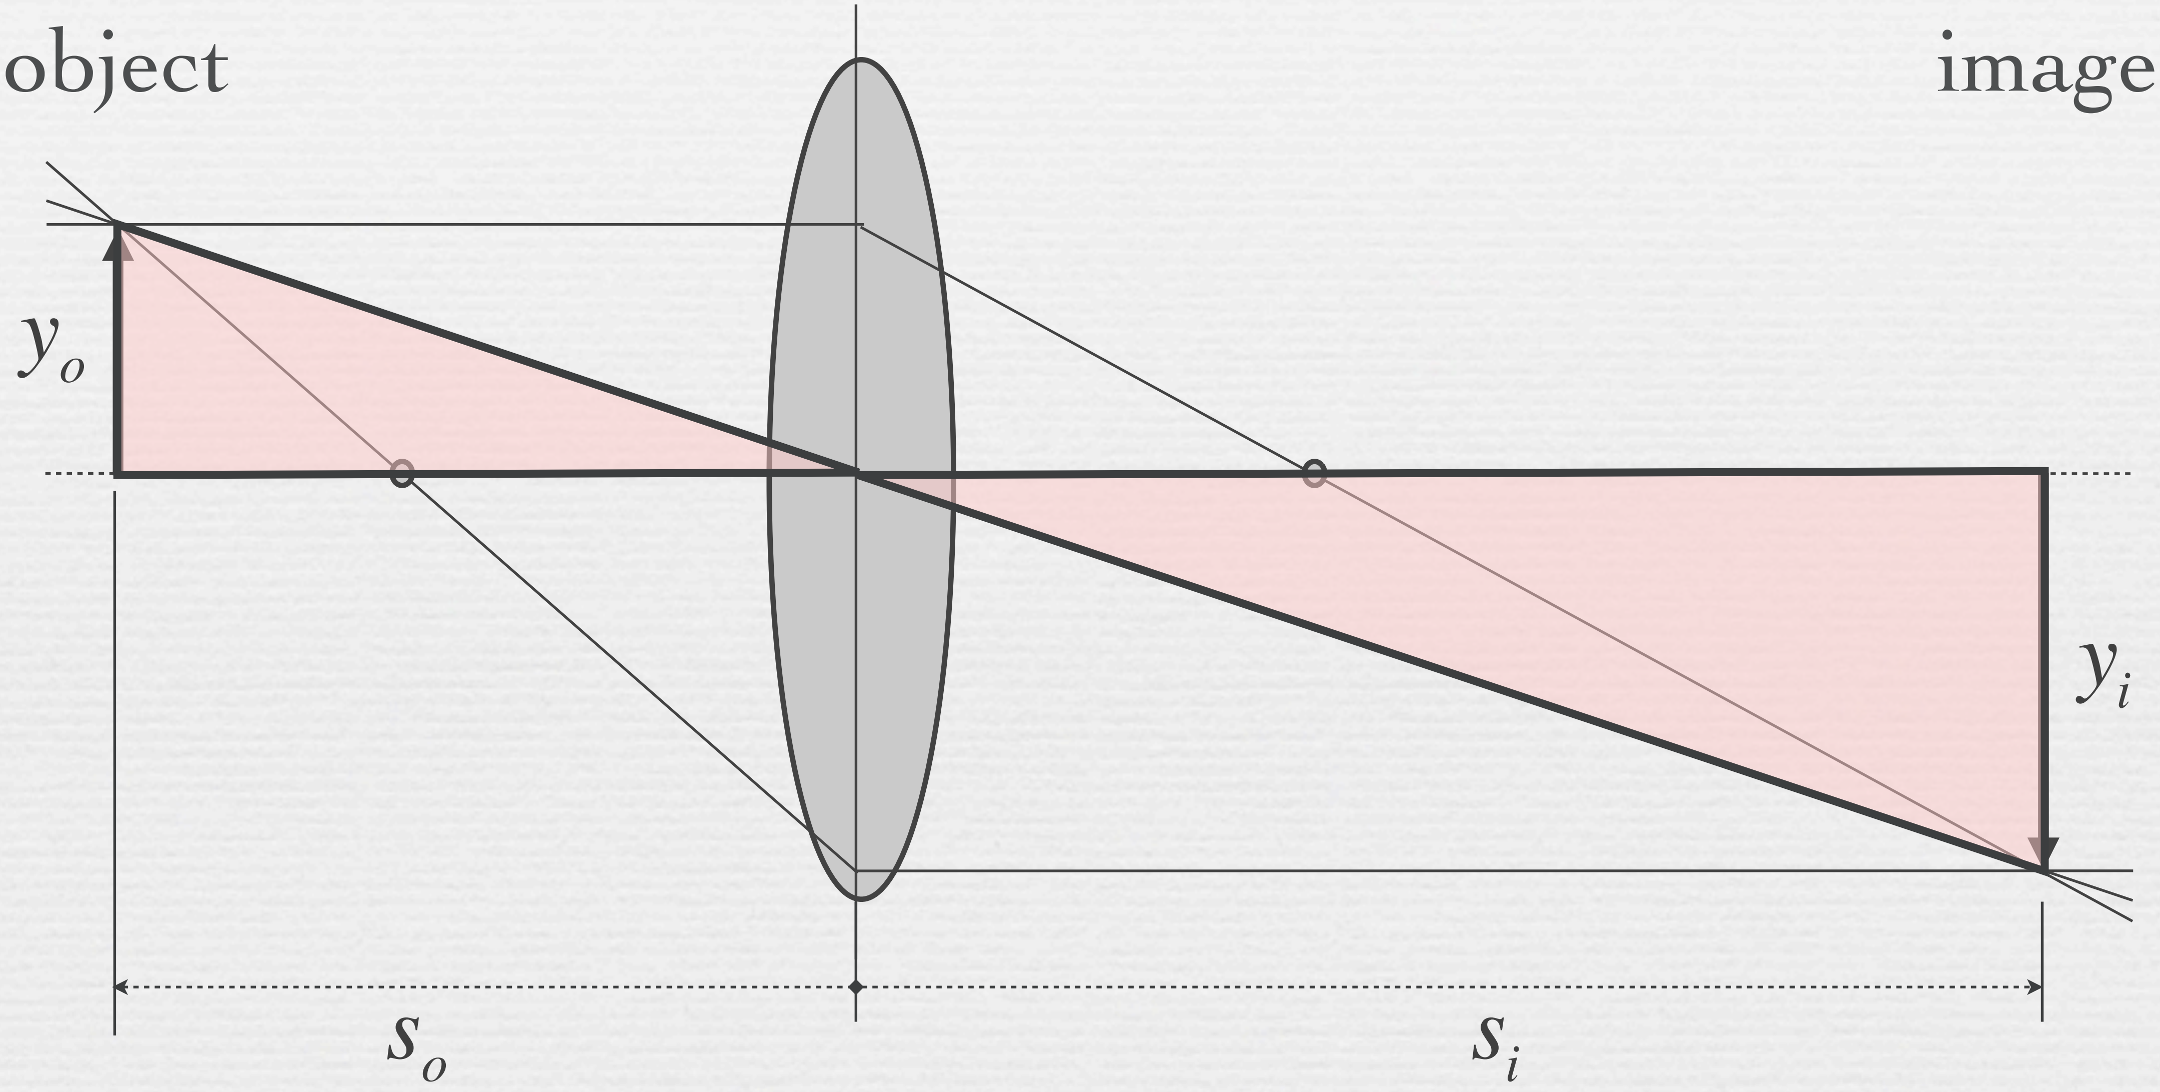

- positive s_i is rightward, positive s_o is leftward
- positive y is upward

From Gauss's ray construction to the Gaussian lens formula

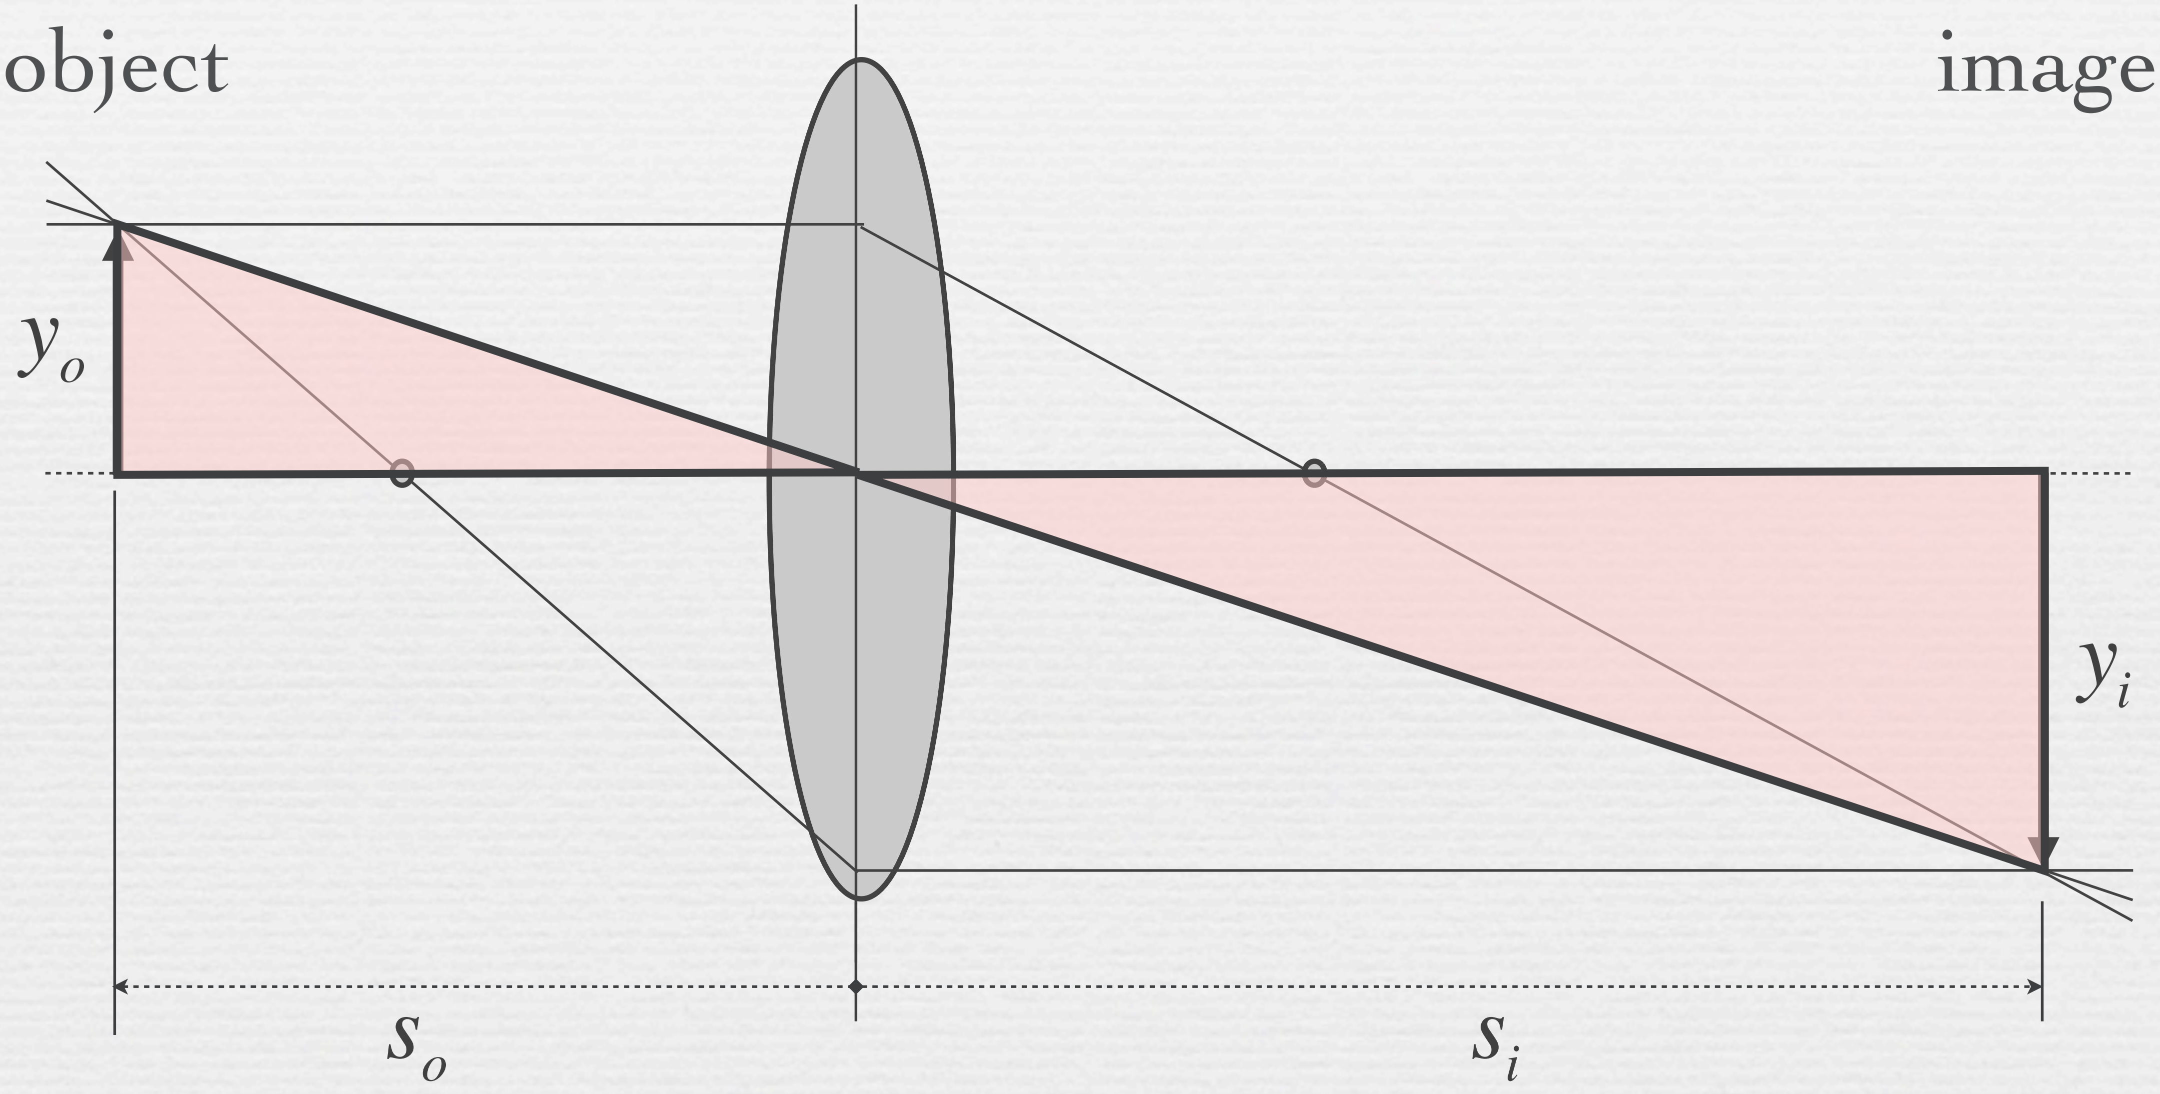

$$\frac{|y_i|}{y_o} = \frac{s_i}{s_o}$$

From Gauss's ray construction to the Gaussian lens formula

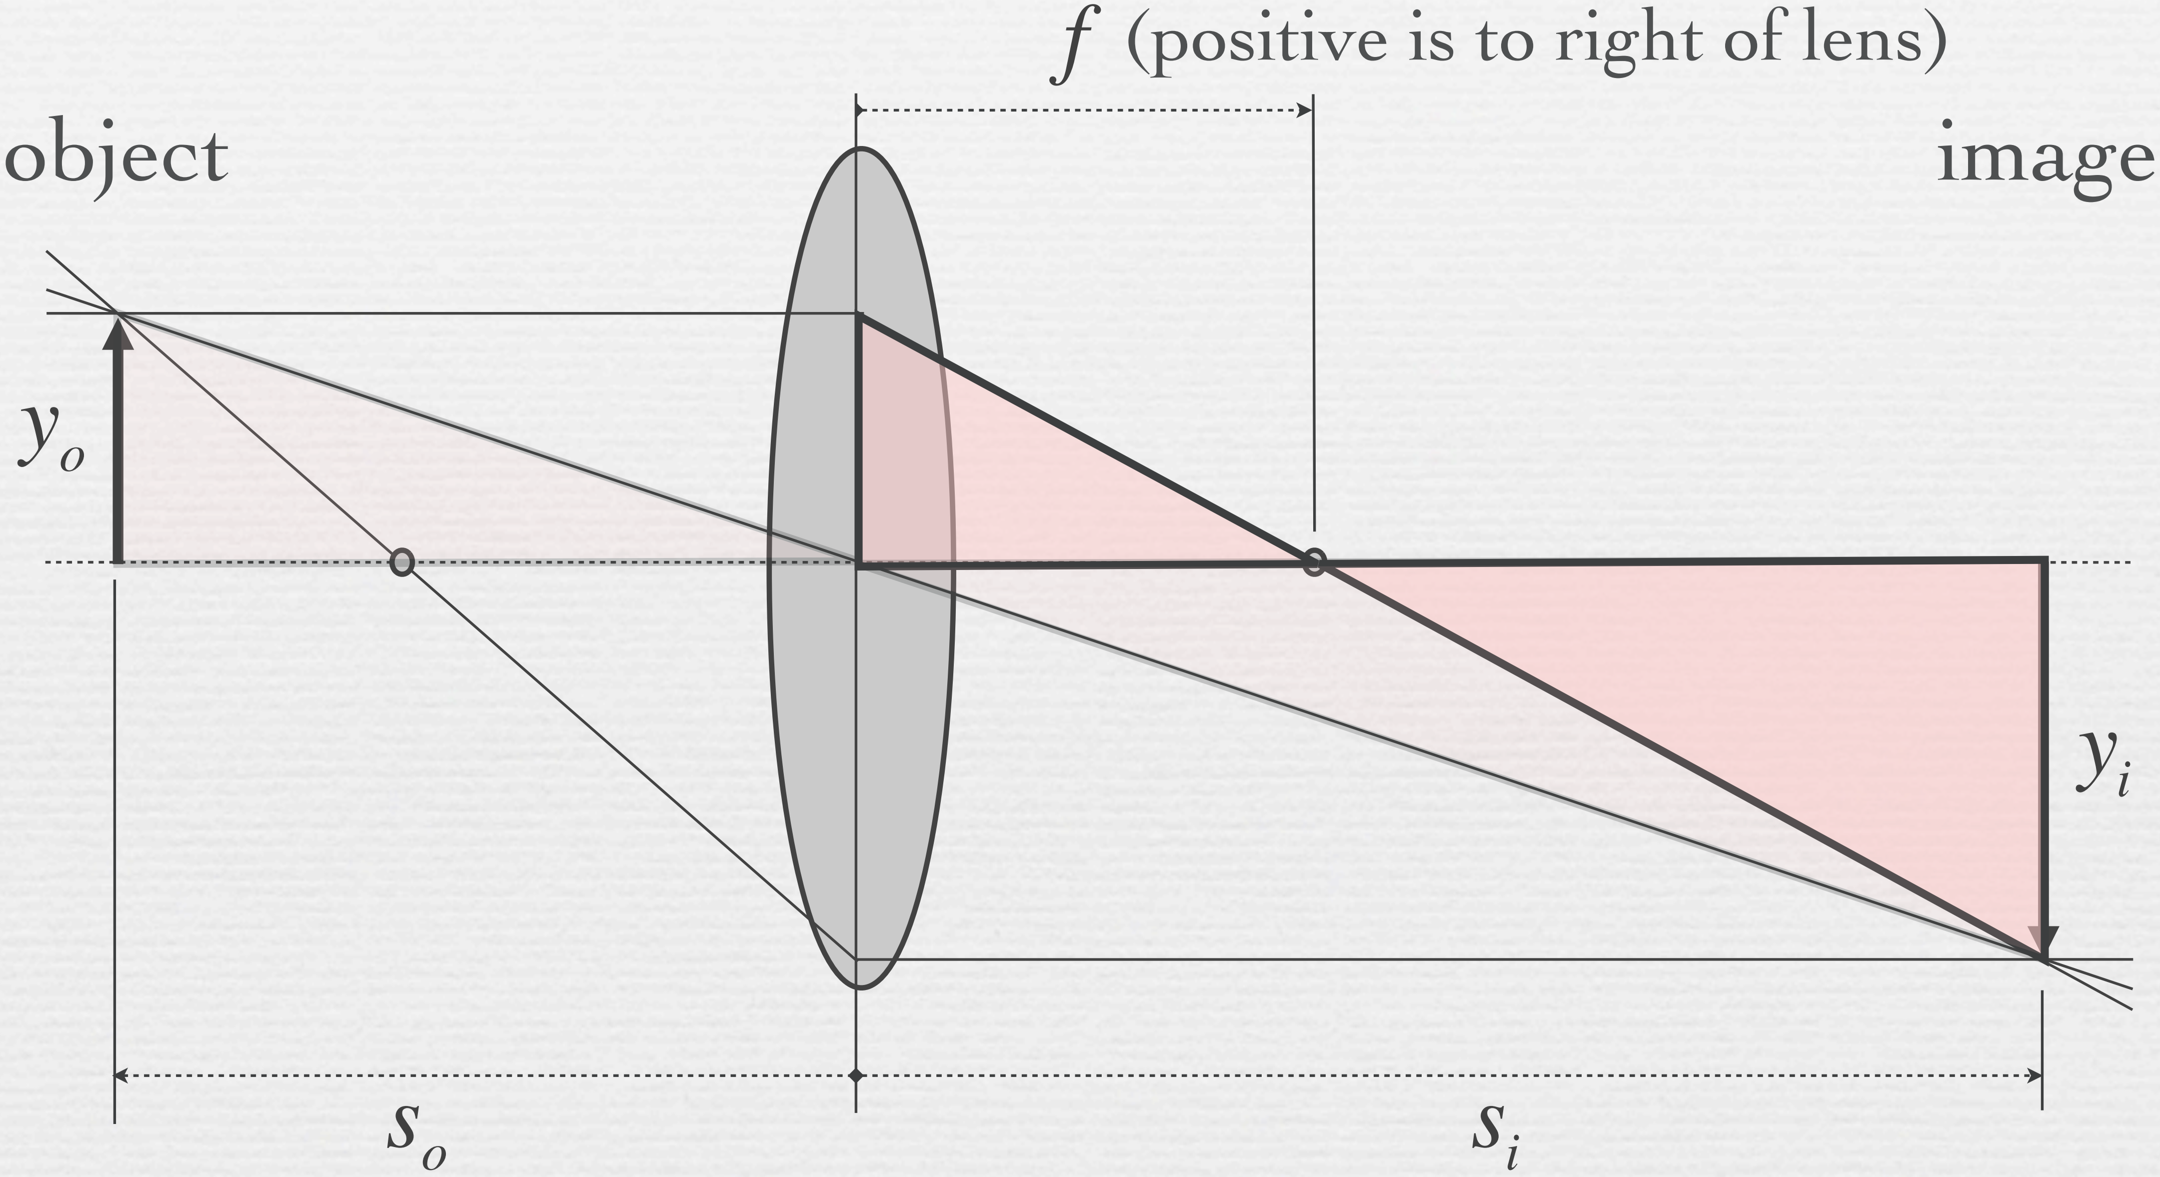

$$\frac{|y_i|}{y_o} = \frac{s_i}{s_o} \quad \text{and} \quad \frac{|y_i|}{y_o} = \frac{s_i - f}{f} \quad \ldots\ldots \quad \boxed{\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}}$$

Changing the focus distance

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

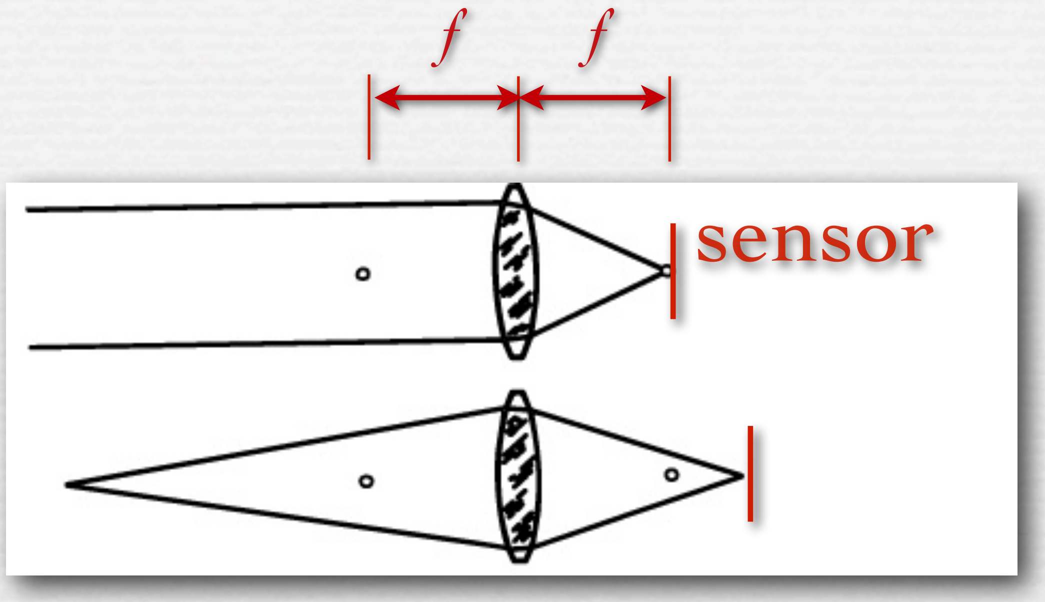

(FLASH DEMO)

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

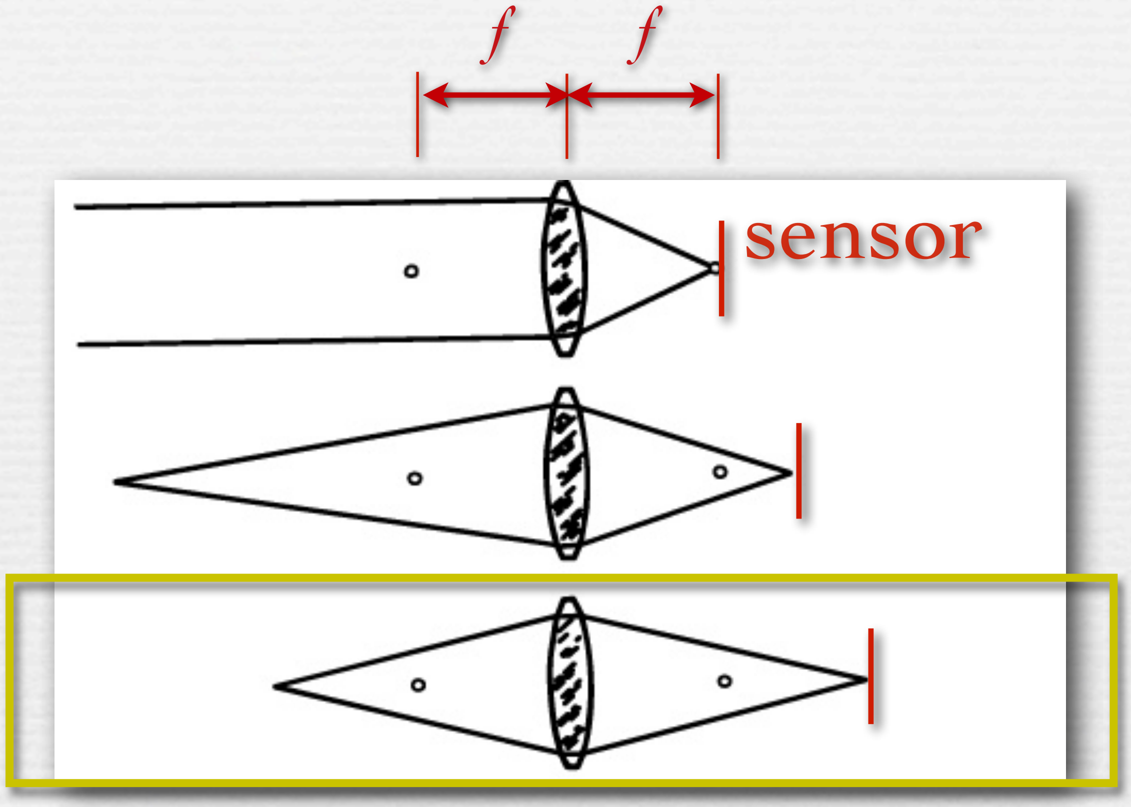

- at $s_o = s_i = 2f$ we have 1:1 imaging, because

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

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

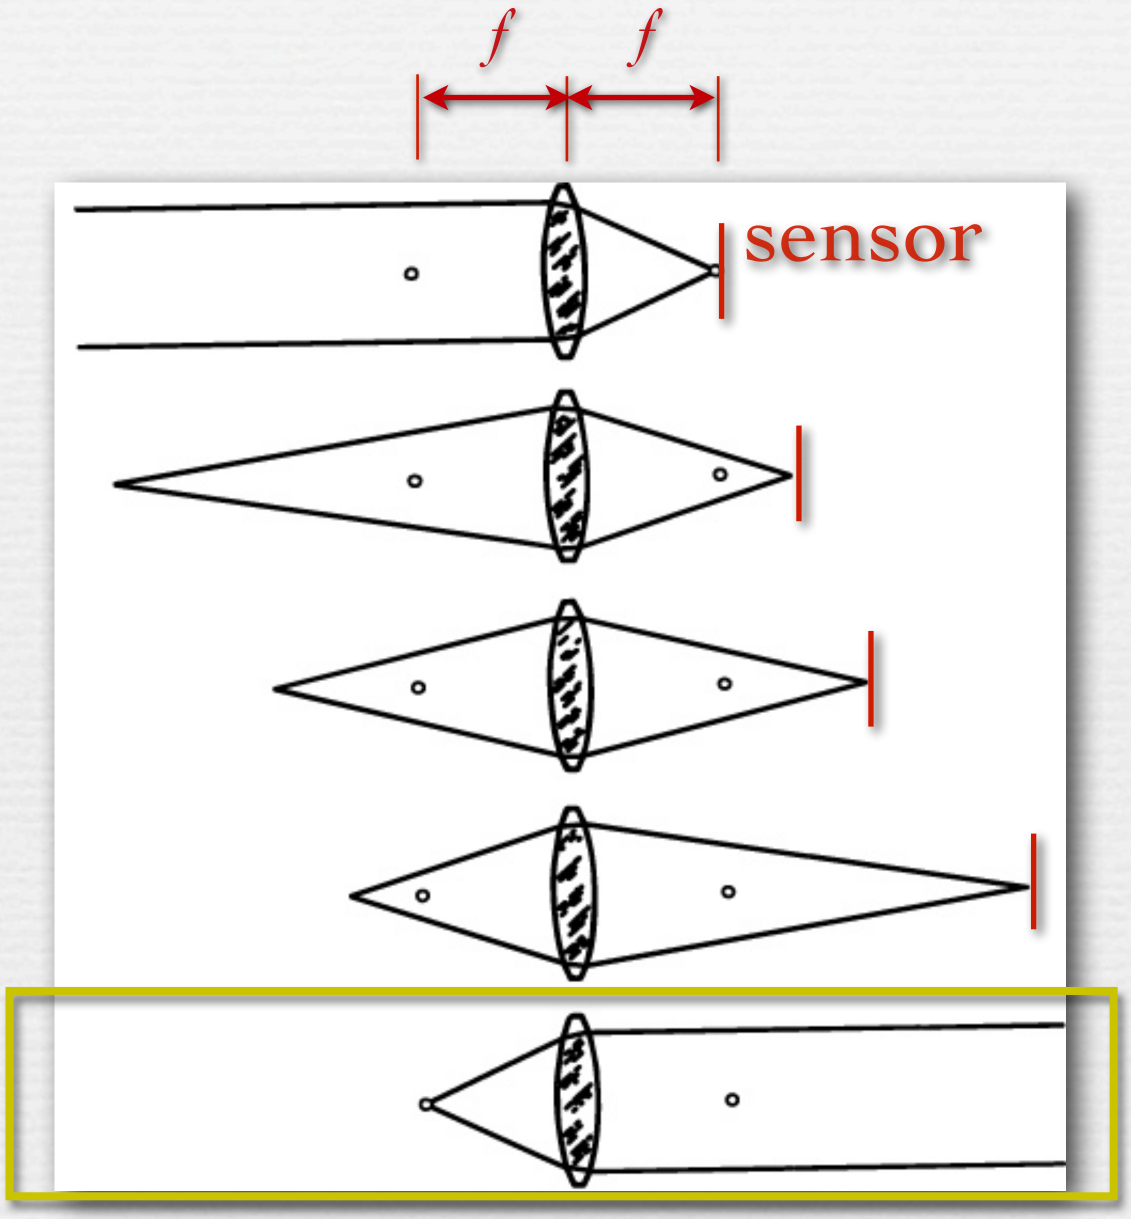

$$\frac{1}{s_o} + \frac{1}{s_i} = \frac{1}{f}$$

Changing the focal length

- weaker lenses have longer focal lengths
- to stay in focus, move the sensor further back
- focused image of tree is located slightly beyond the focal length

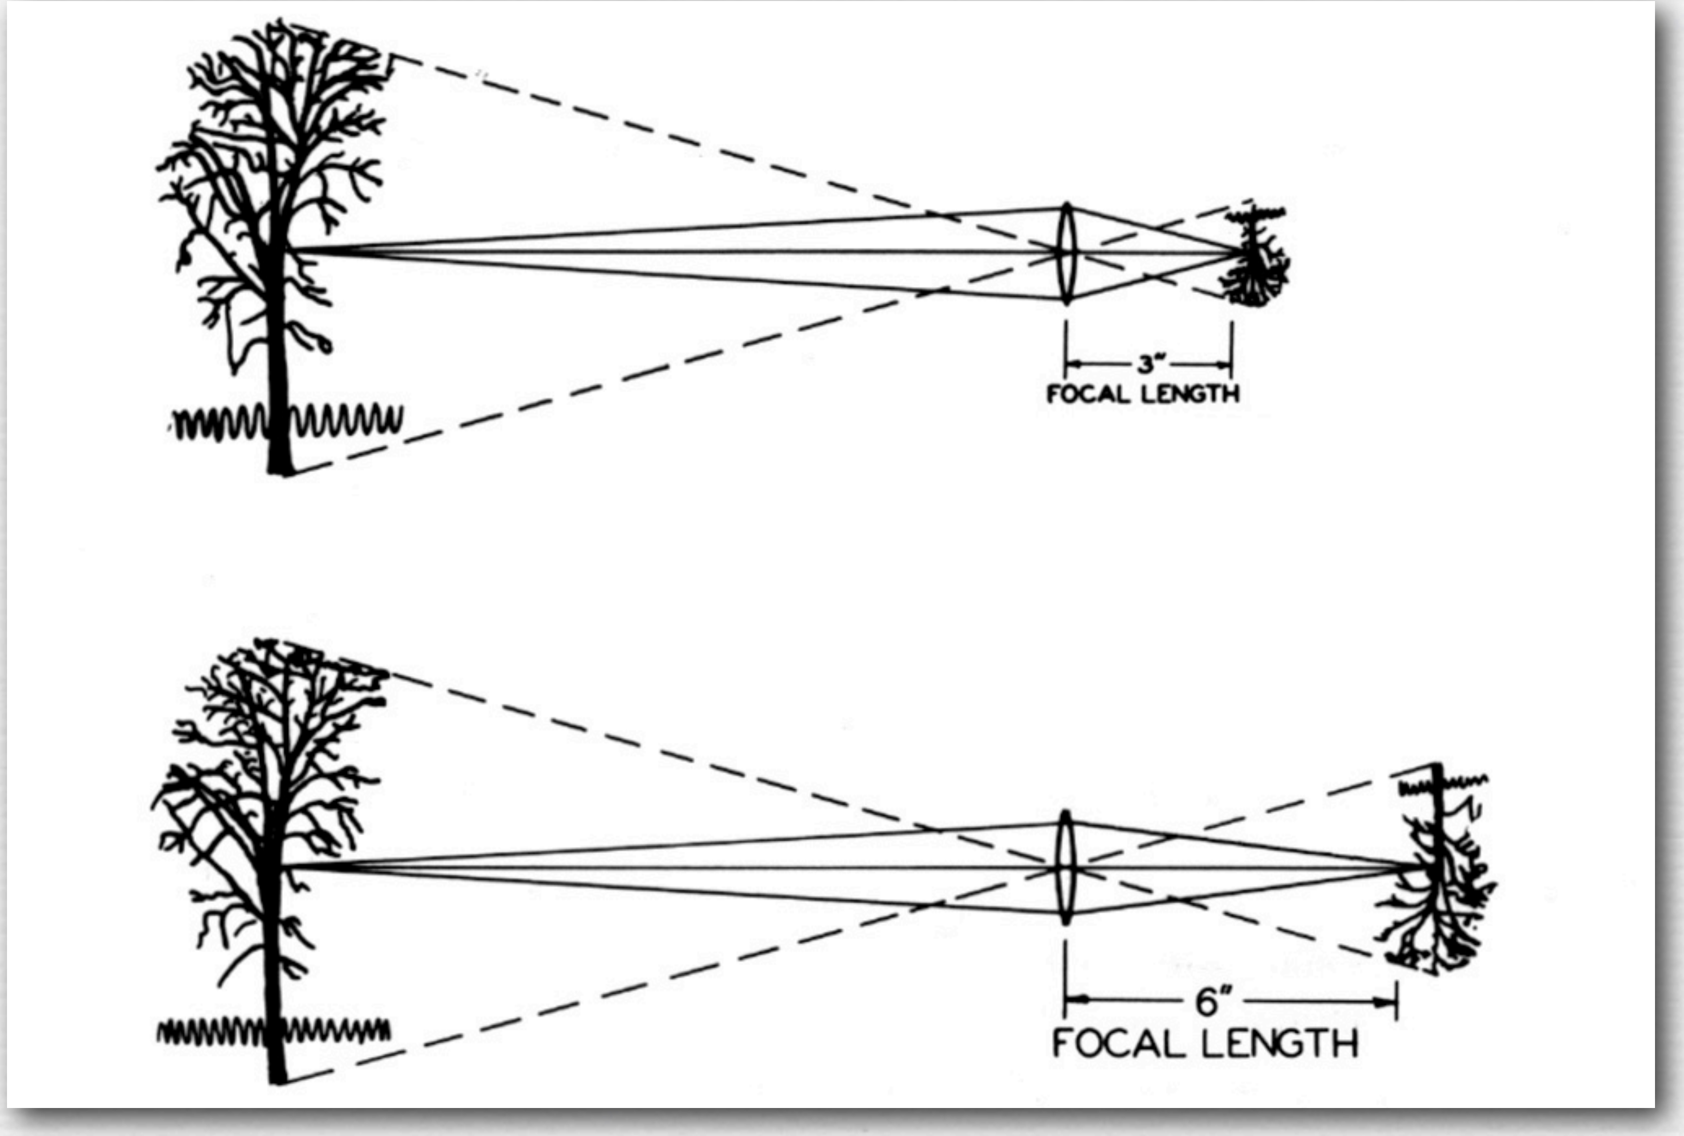

(Kingslake)

The tree would be in focus at the lens focal length only if it were infinitely far away.

Changing the focal length

- if the sensor size is constant, the field of view becomes smaller

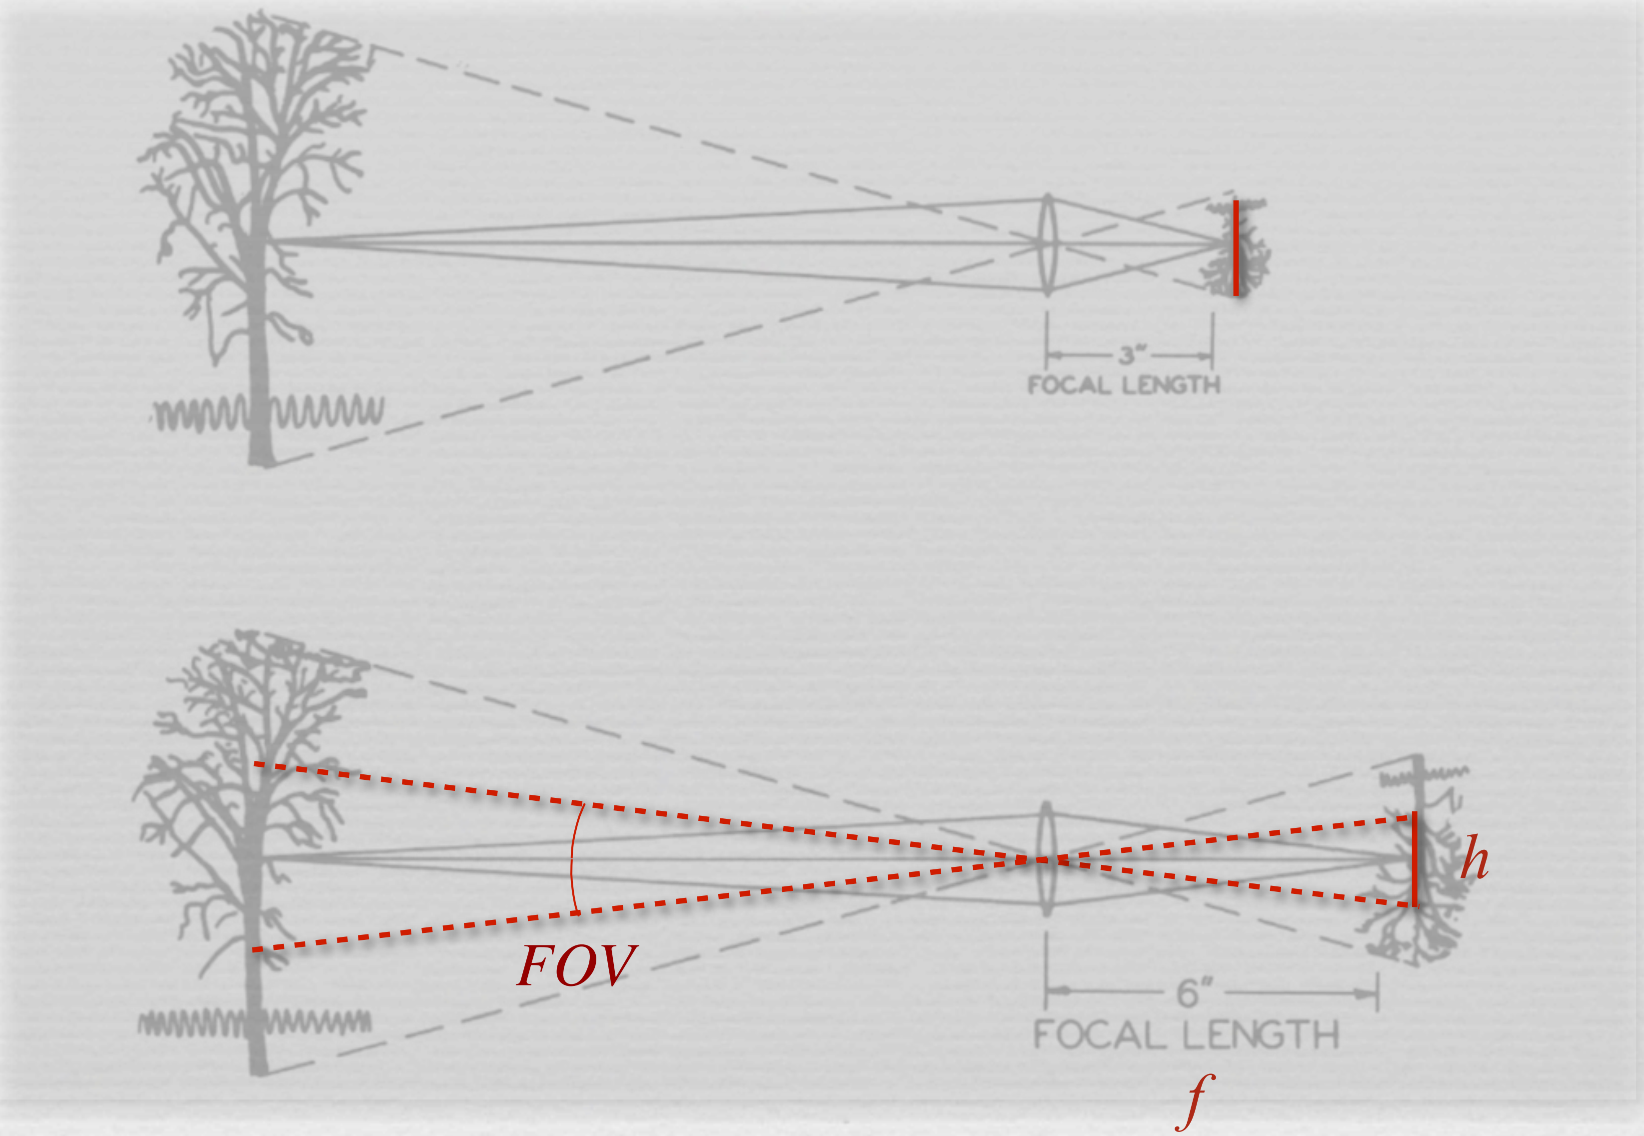

(Kingslake)

$$FOV = 2 \arctan (h / 2 f)$$

Focal length and field of view

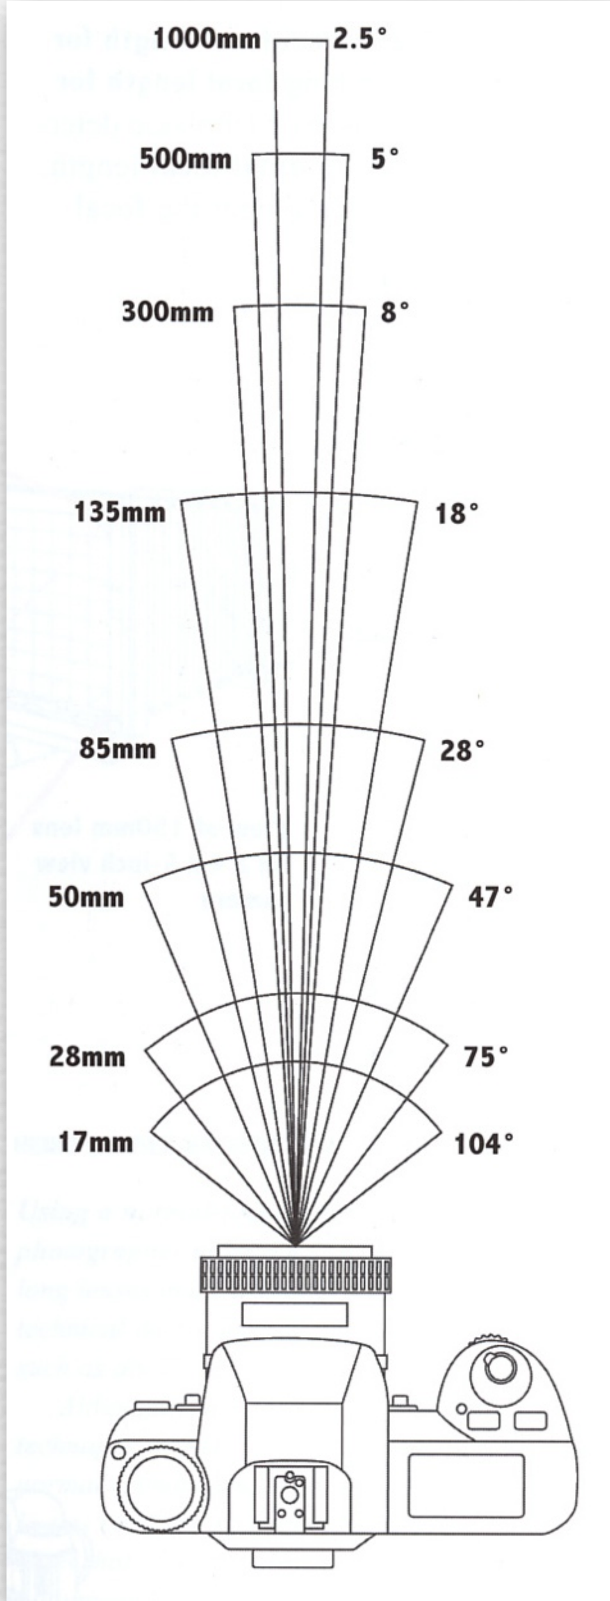

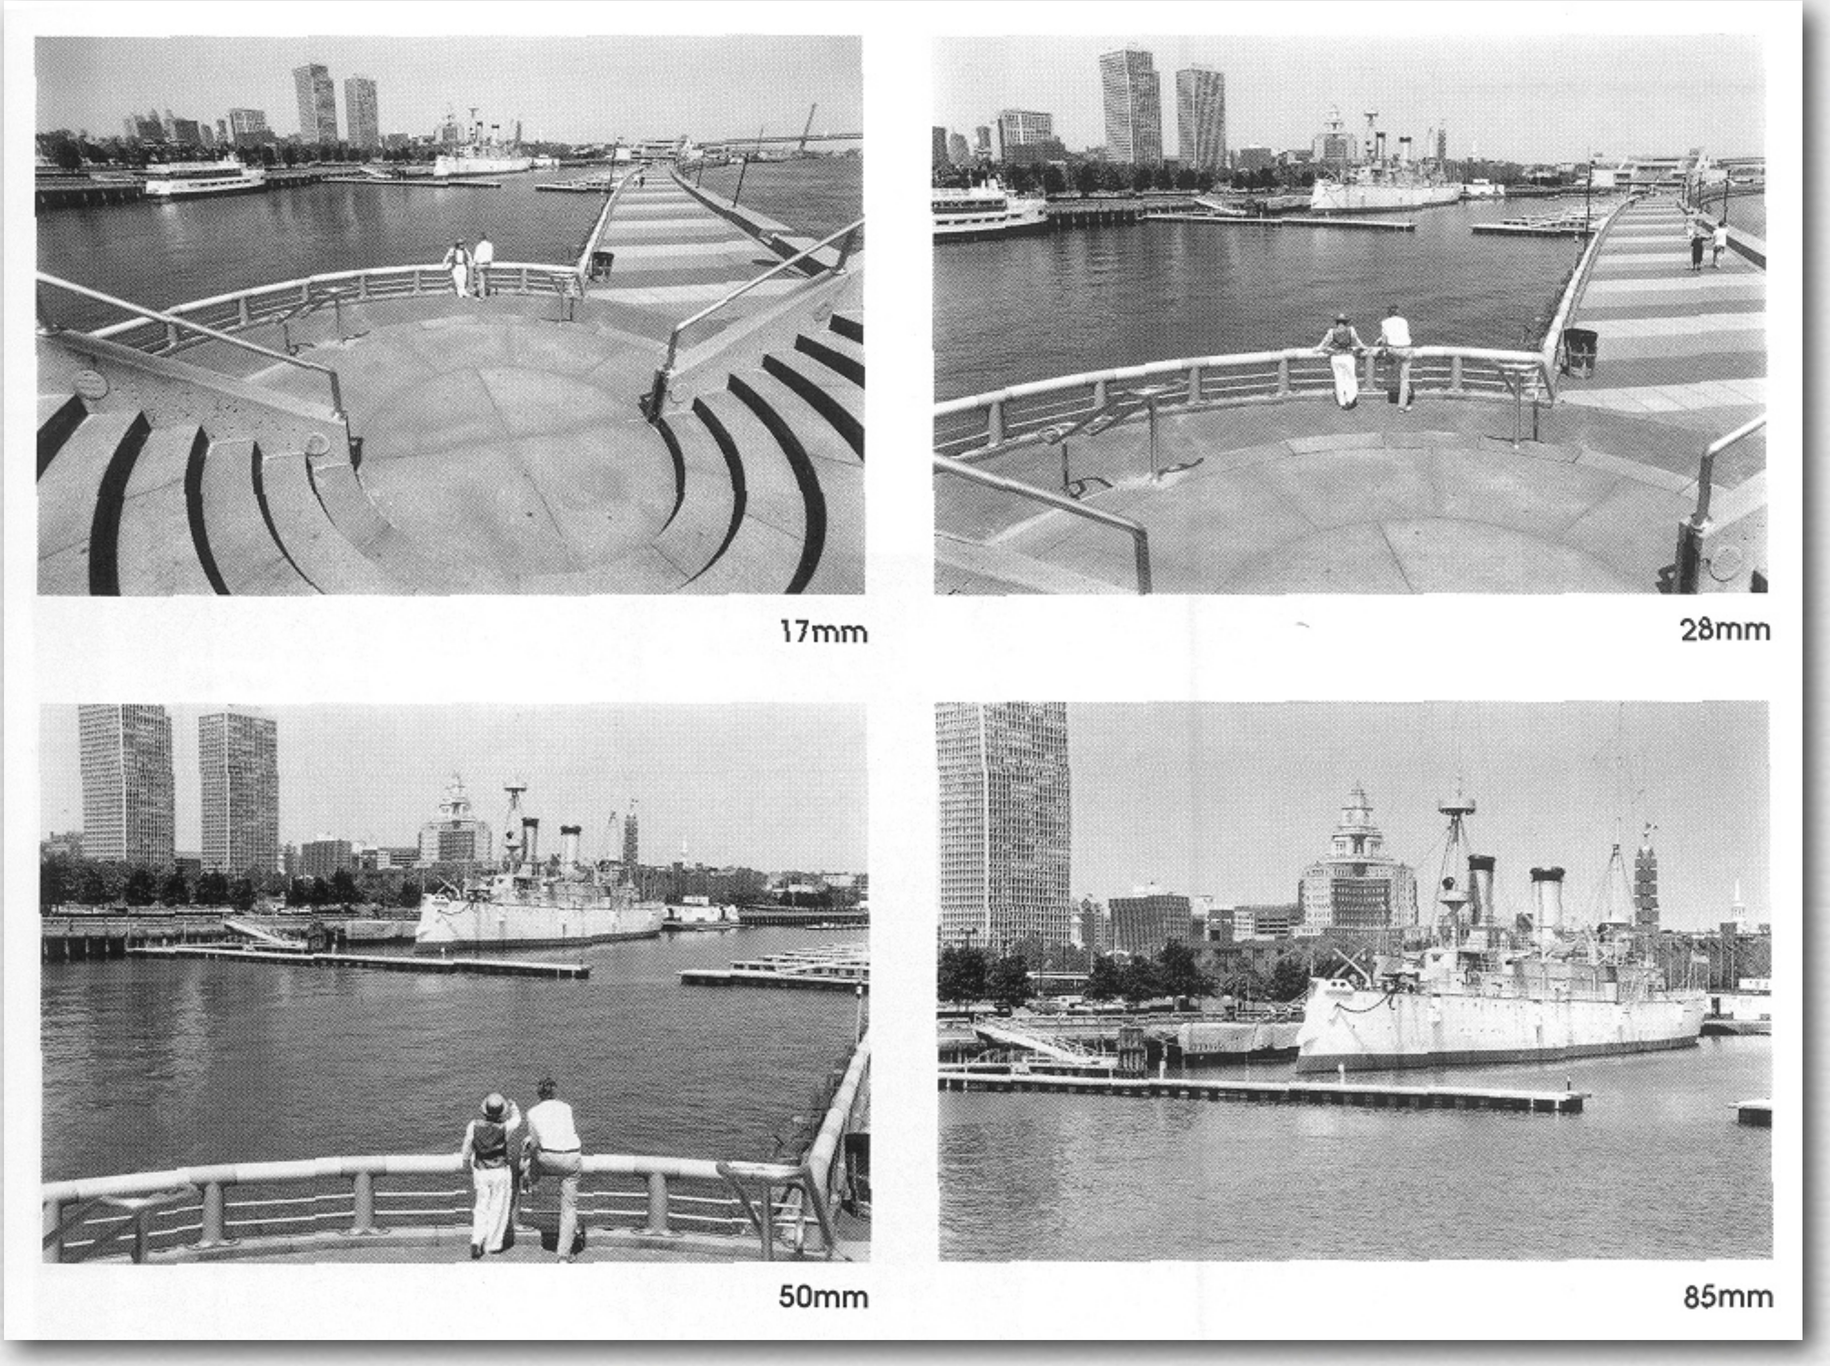

(London)

FOV measured diagonally on a
35mm full-frame camera (24 × 36mm)

Focal length and field of view

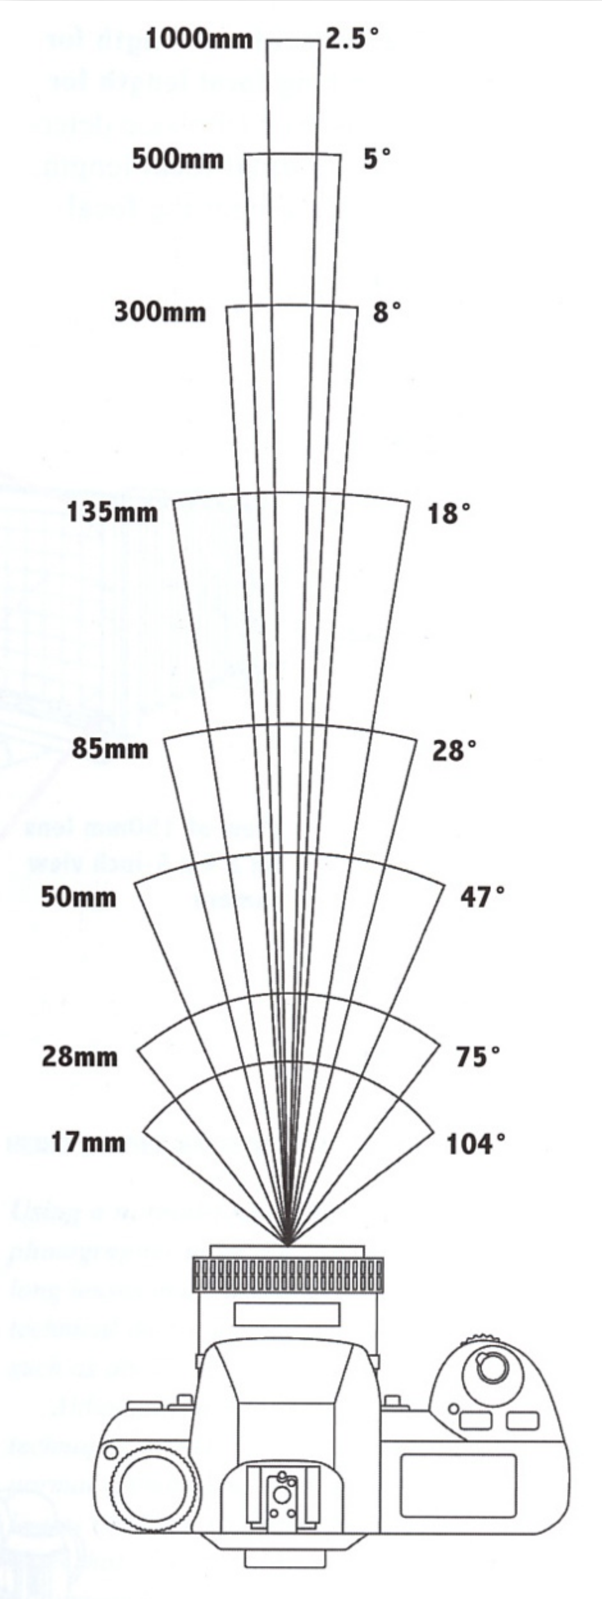

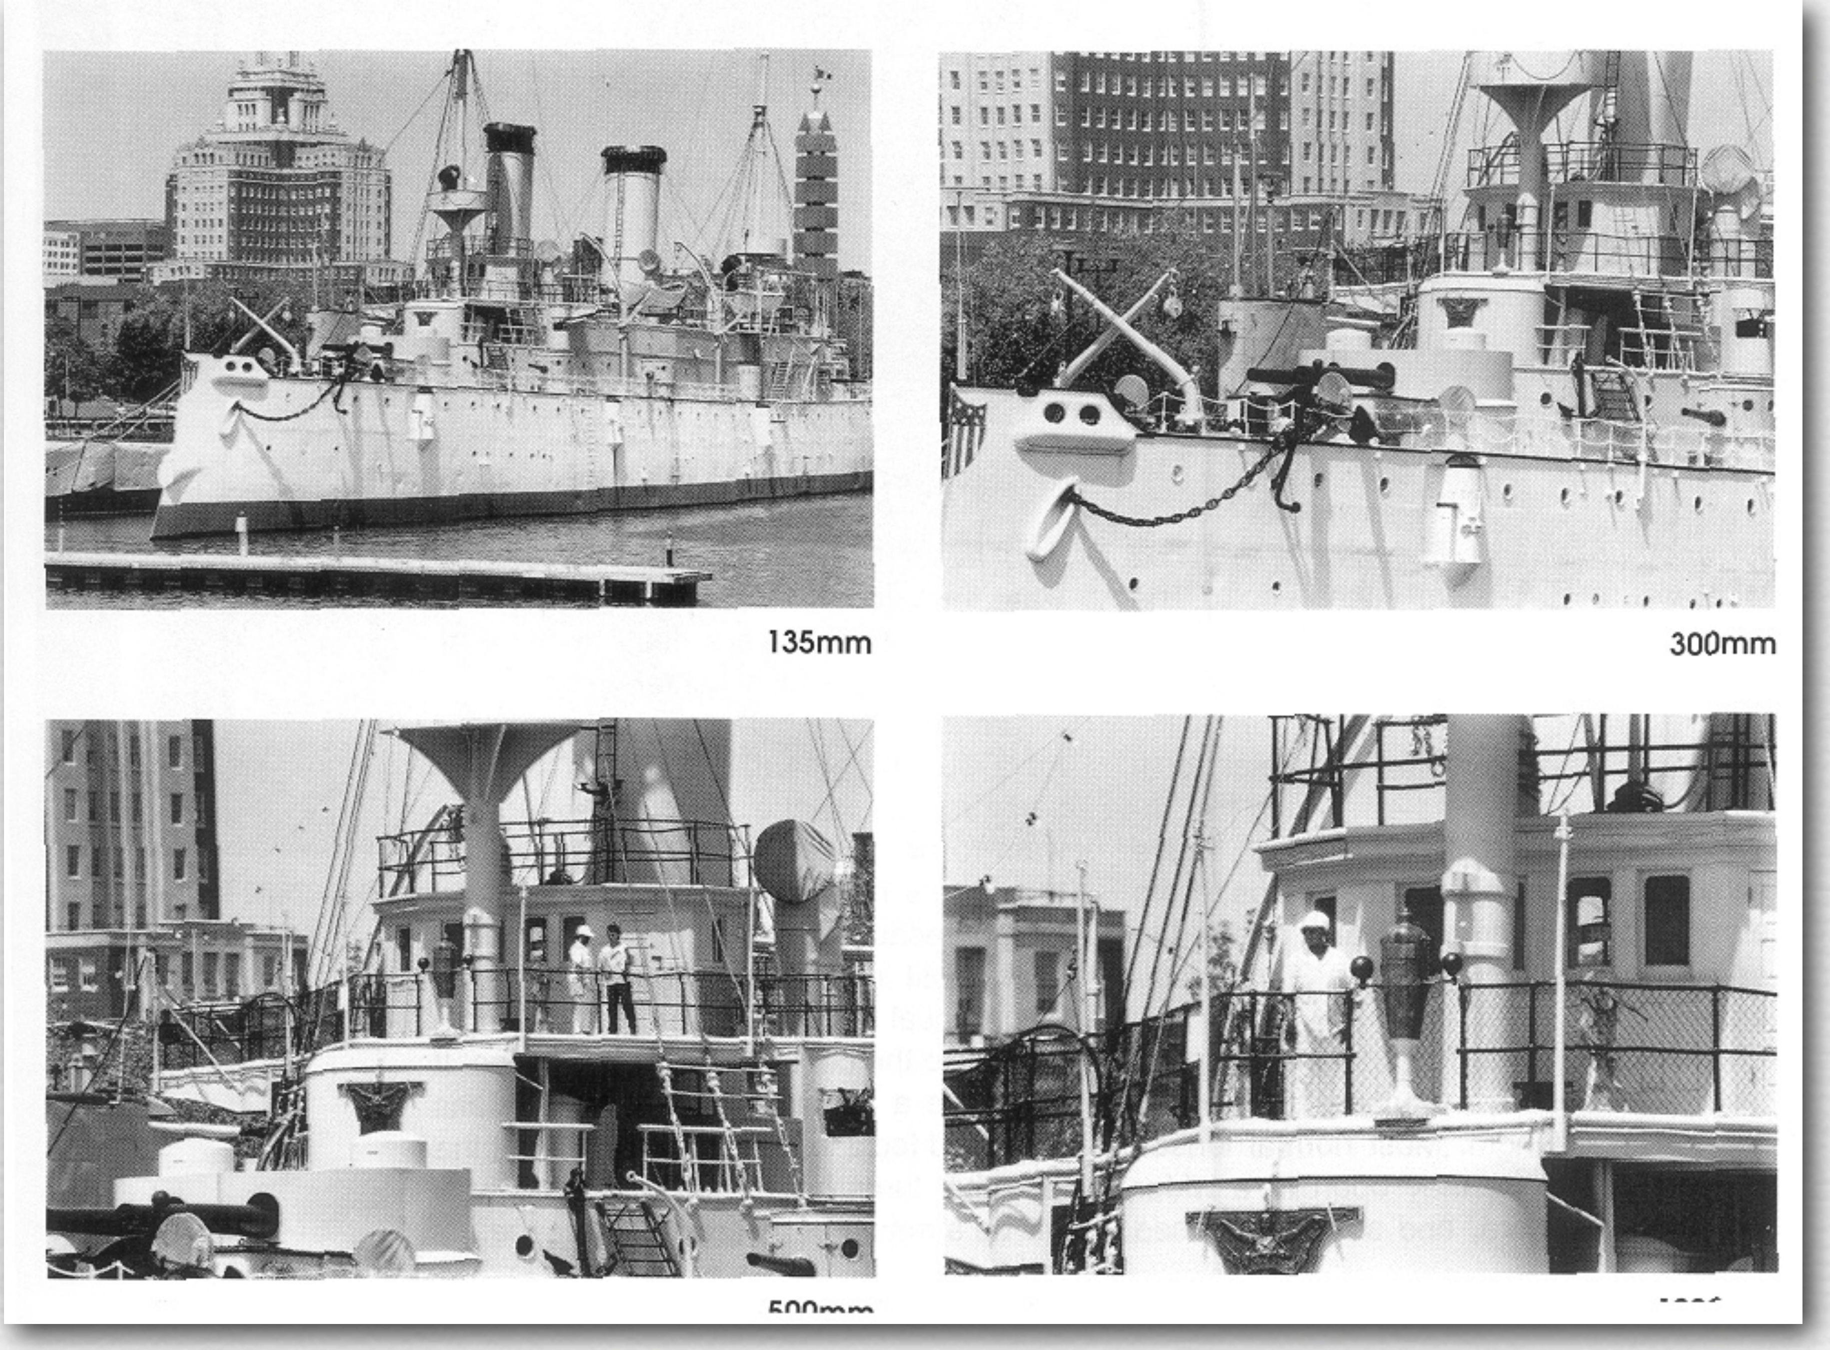

(London)

FOV measured diagonally on a
35mm full-frame camera (24 × 36mm)

Changing the sensor size

- if the sensor size is smaller, the field of view is smaller too
- smaller sensors either have fewer pixels, or noiser pixels

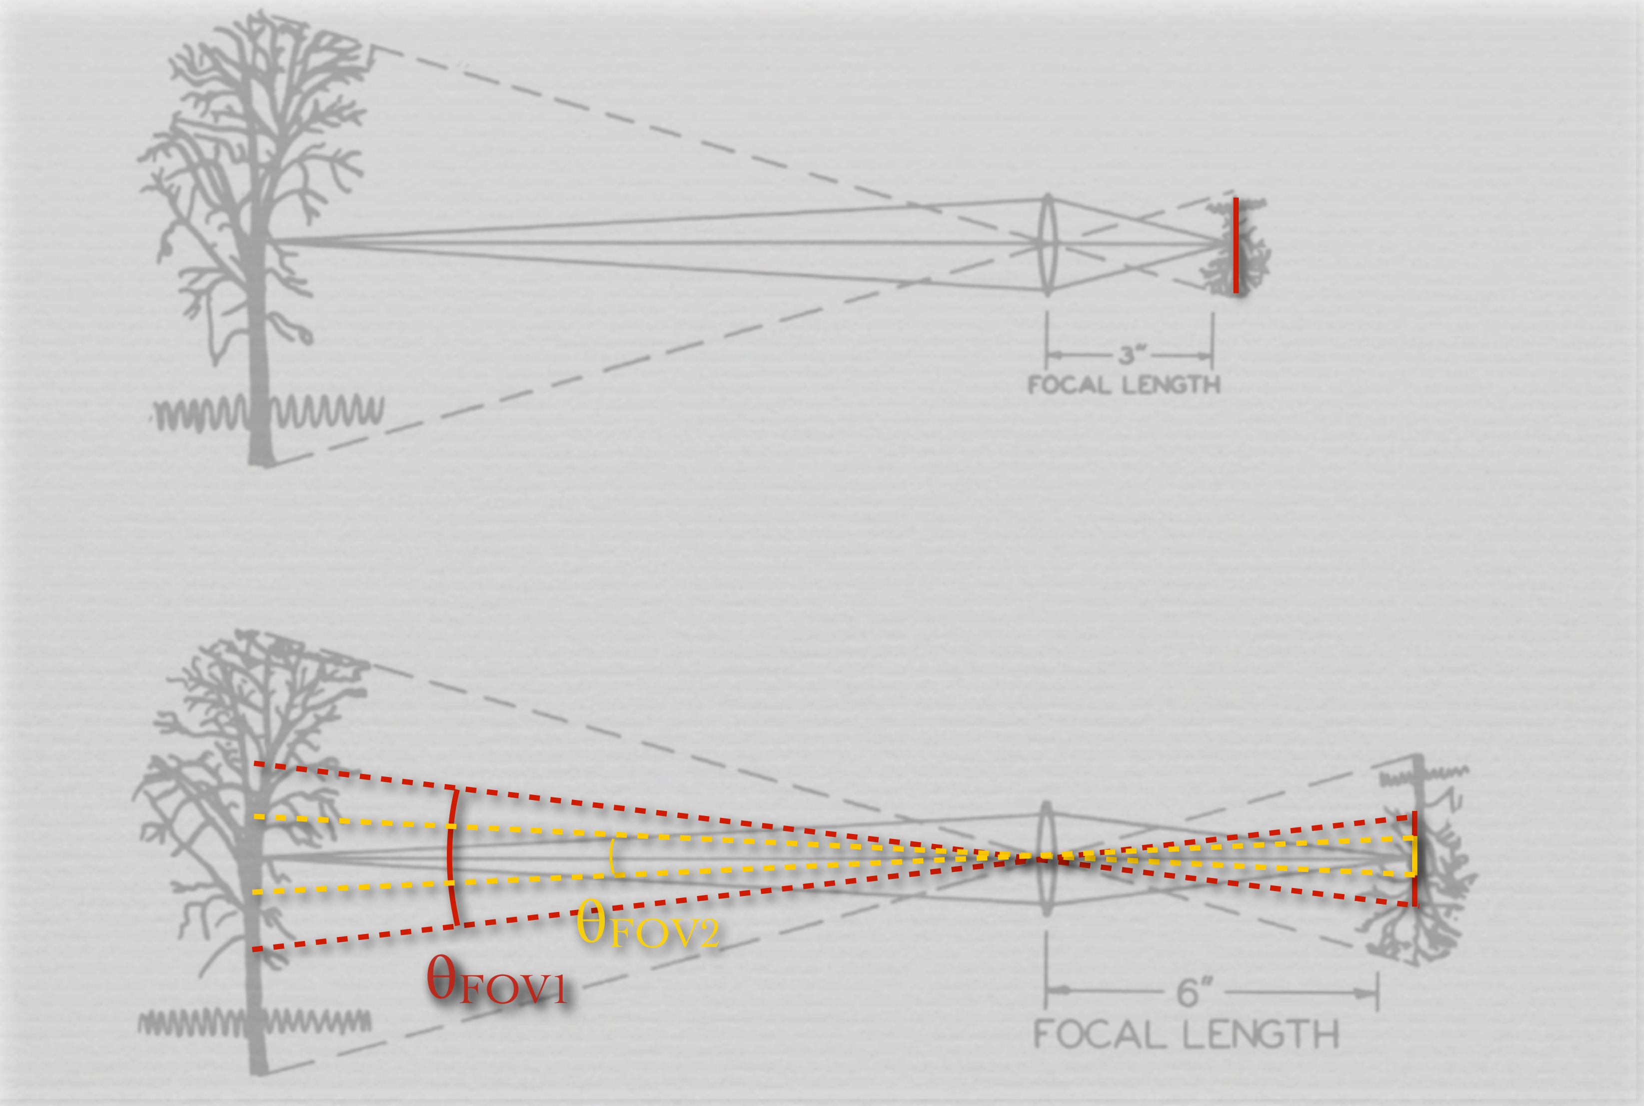

(Kingslake)

Sensor sizes

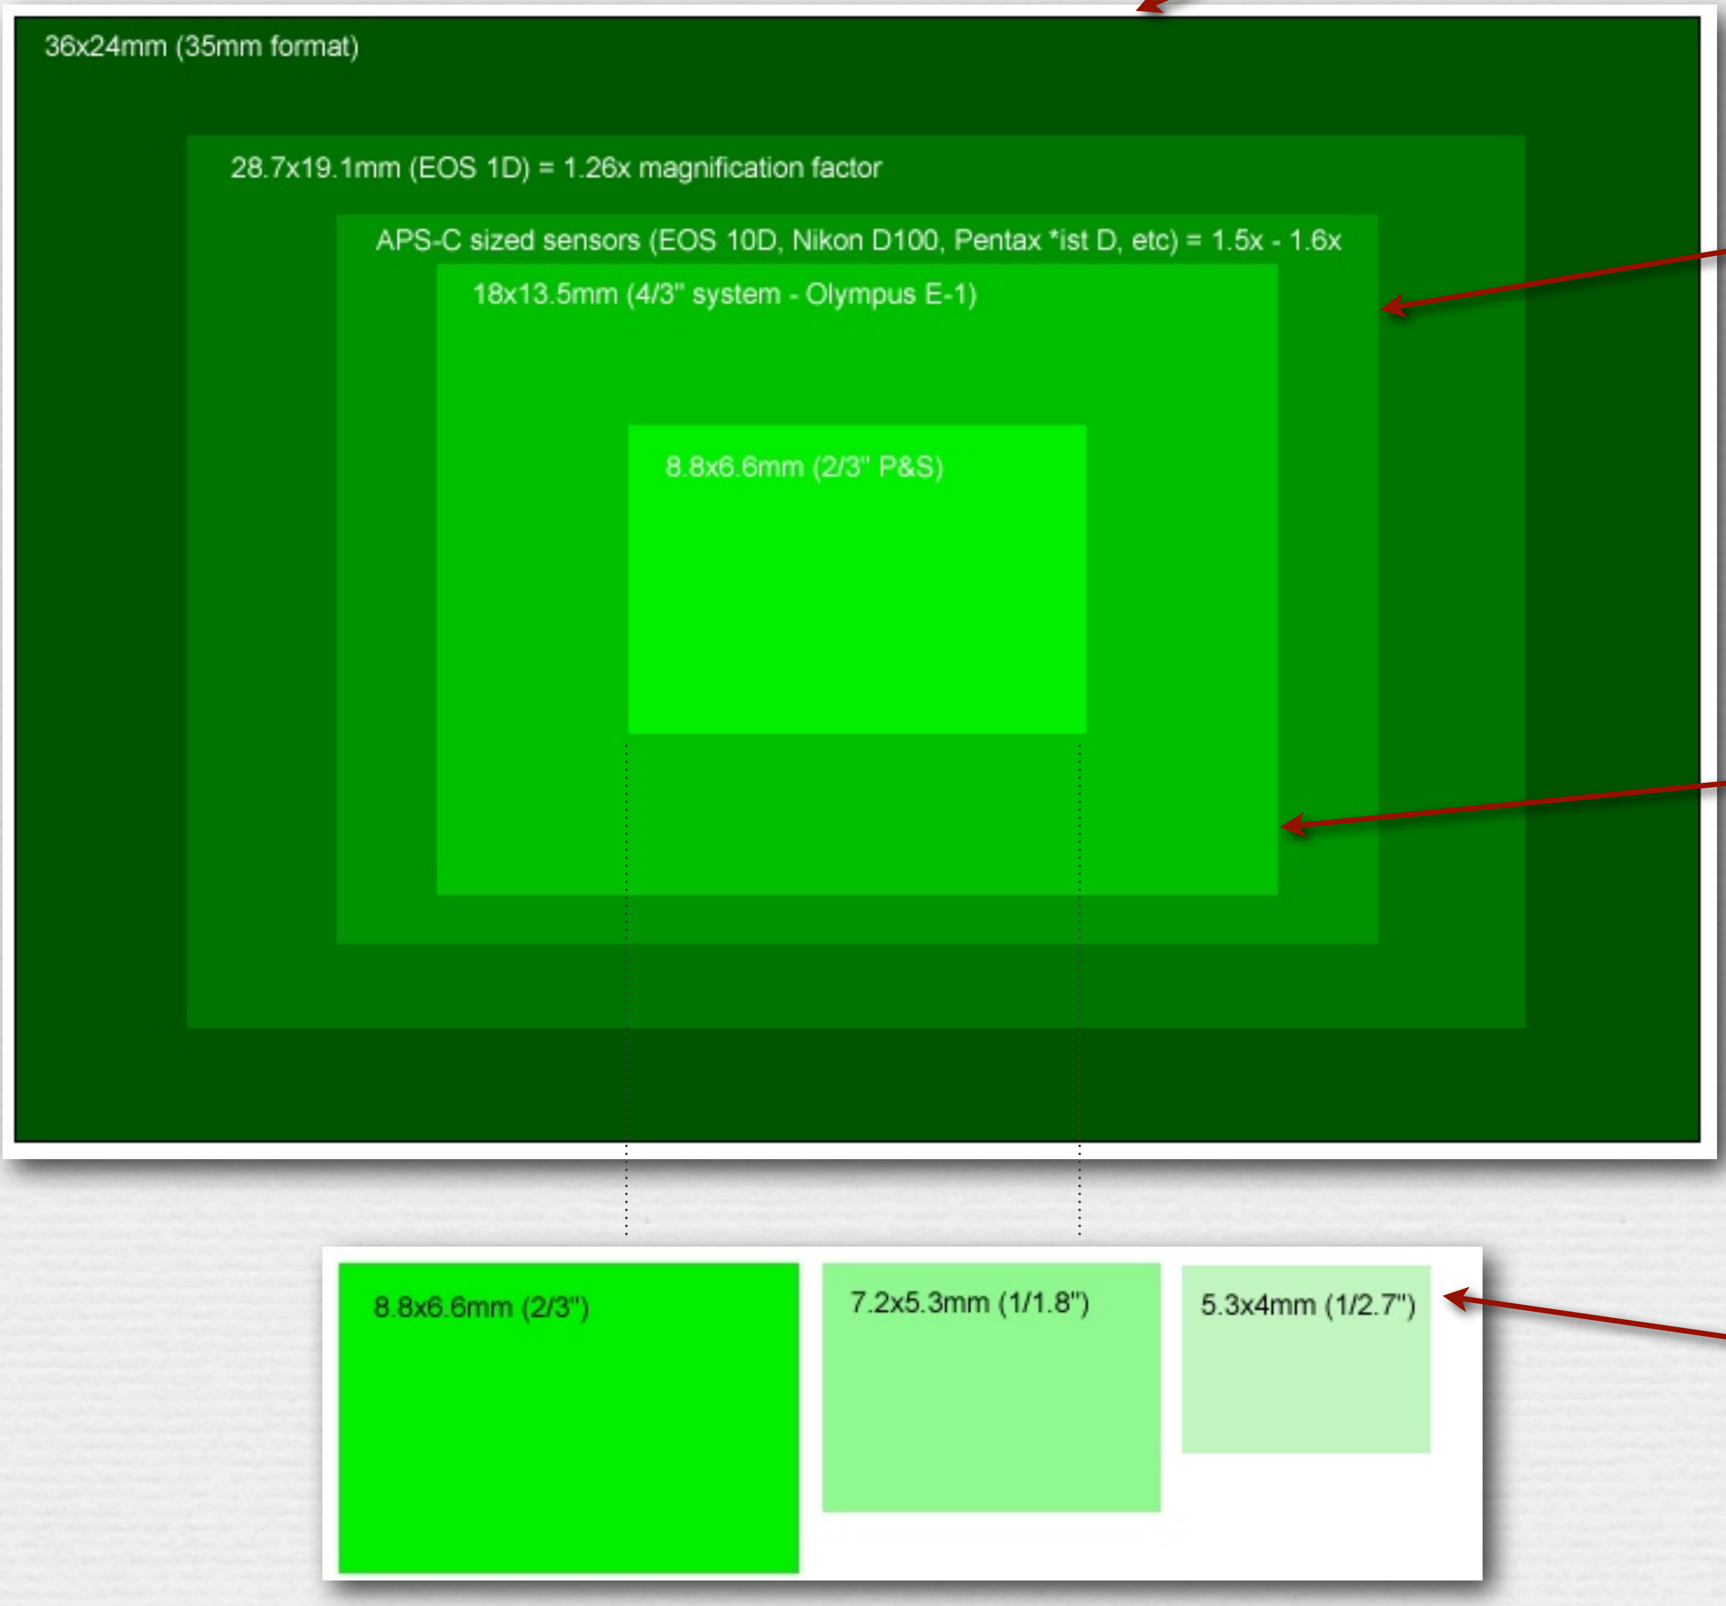

"full frame"
Canon 5D Mark II
(24mm × 36mm)

"APS-C"
Nikon D40
(15.5mm × 23.7mm)
(~1.5× crop factor)

"micro 4/3"
Panasonic GF1
(13mm × 17.3mm)
(~2× crop factor)

"point-and-shoot"
Canon A590
(5.75mm × 4.31mm)
(~8× crop factor)

Magnification

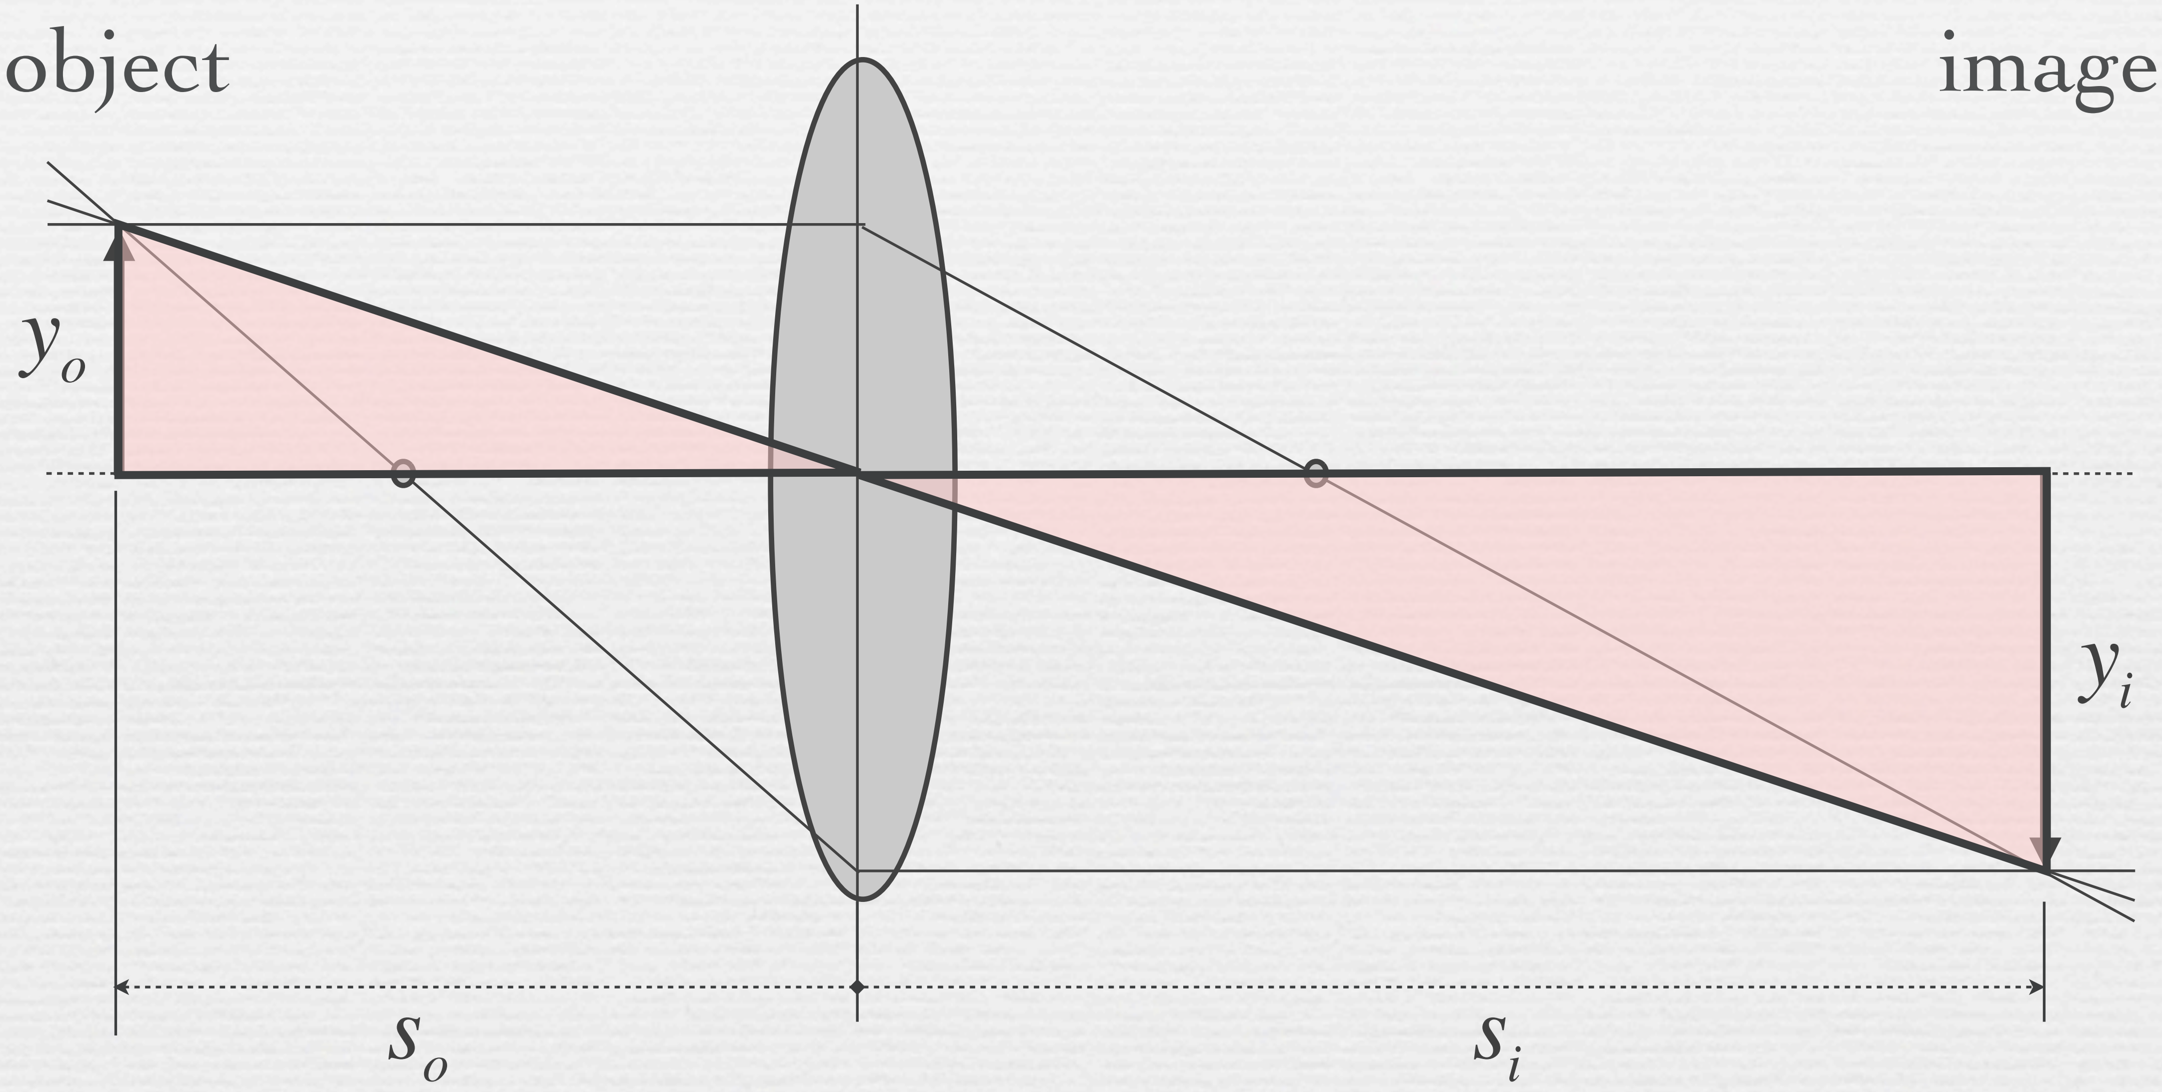

$$M_T \; @ \; \frac{y_i}{y_o} = -\frac{s_i}{s_o}$$

Lenses perform a 3D perspective transform

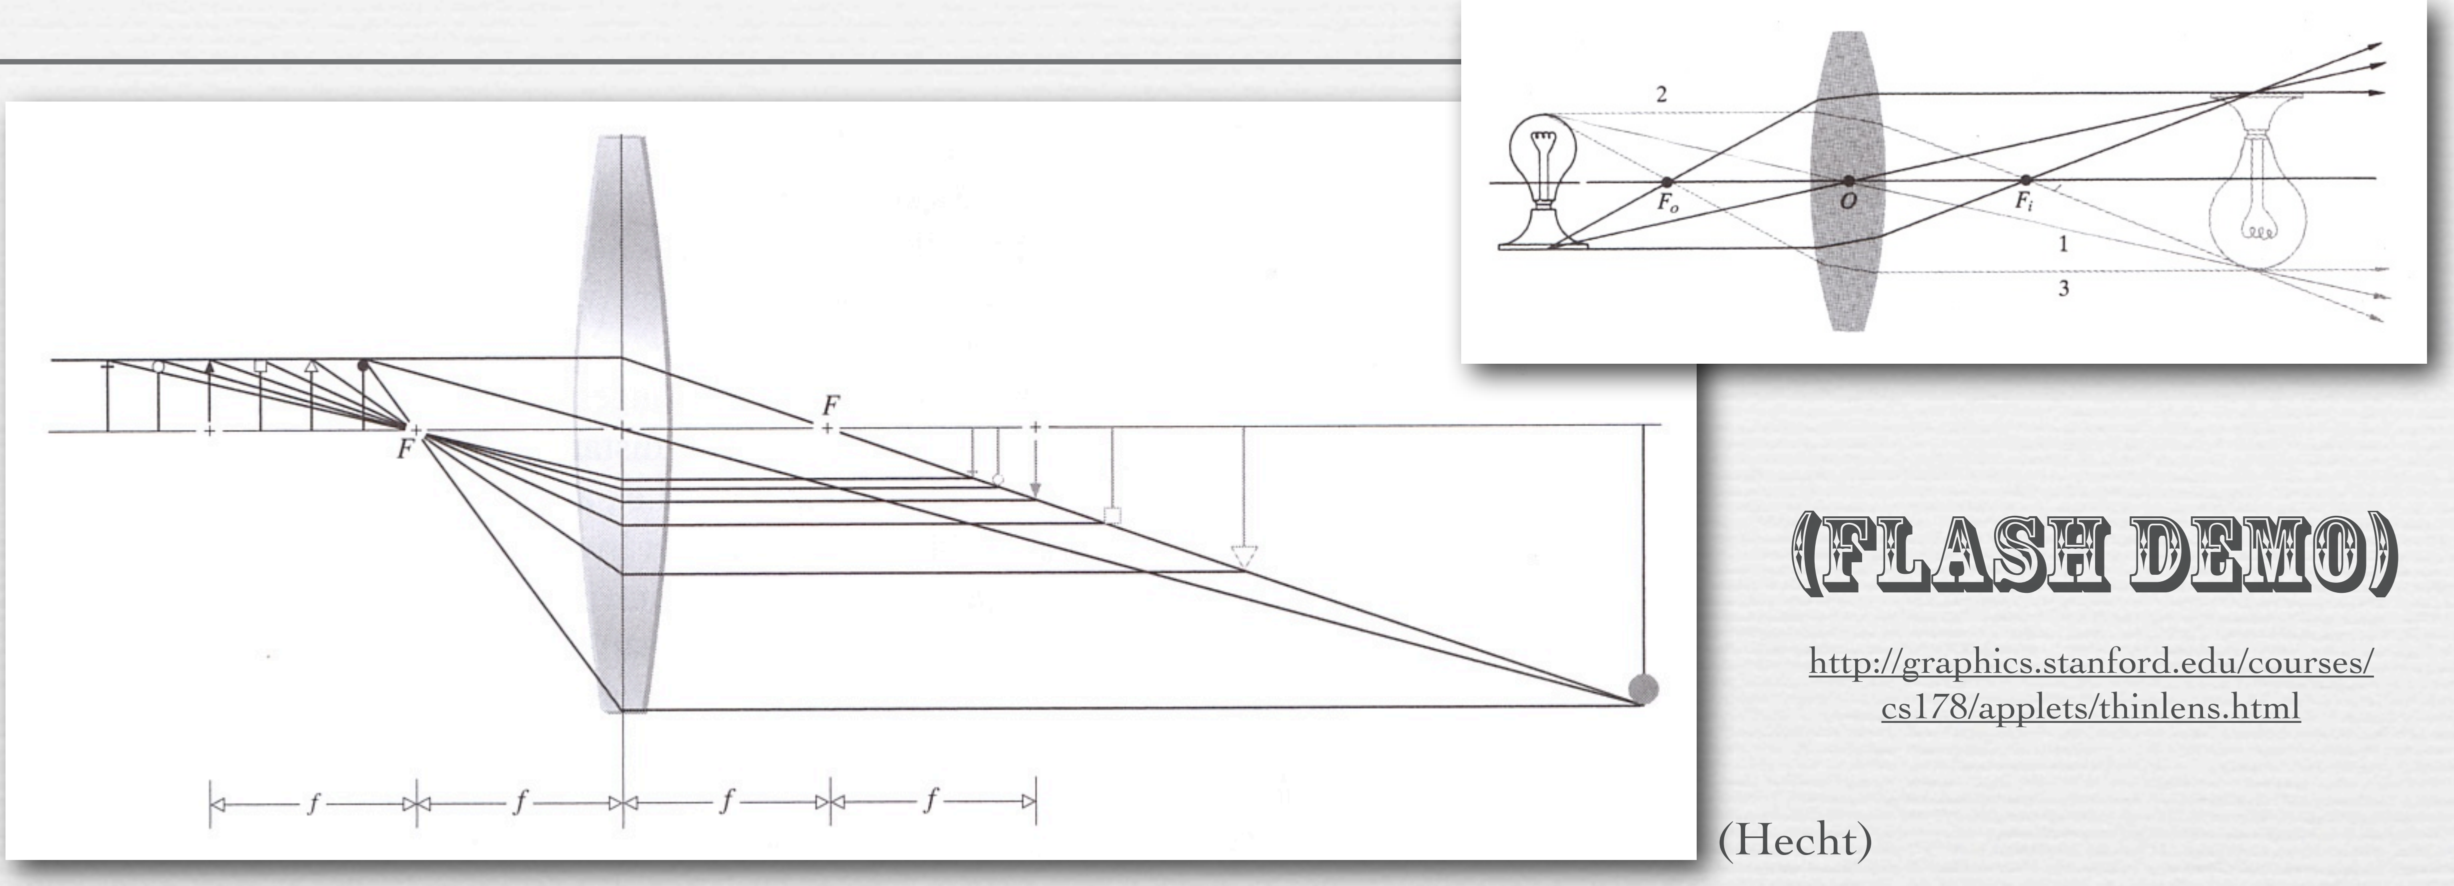

(Hecht)

(FLASH DEMO)

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

- 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)

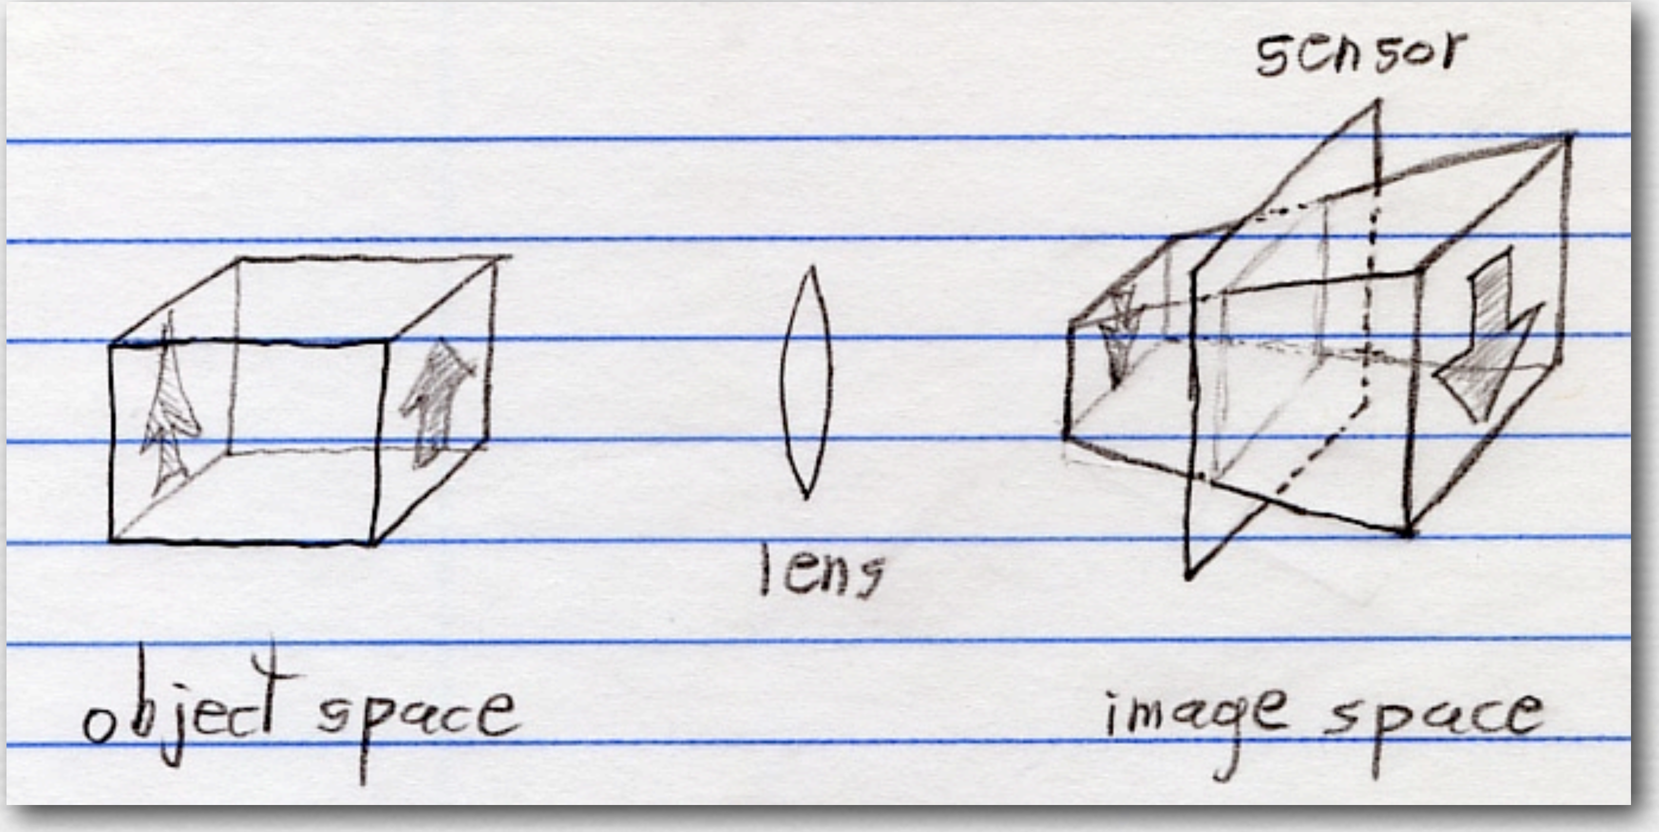

- a cube in object space is transformed by a lens into a 3D frustum in image space, with the orientations shown by the arrows
- in computer graphics this transformation is modeled as a 4 × 4 matrix multiplication of 3D points expressed in 4D homogenous coordinates
- in photography a sensor extracts a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

Depth of field

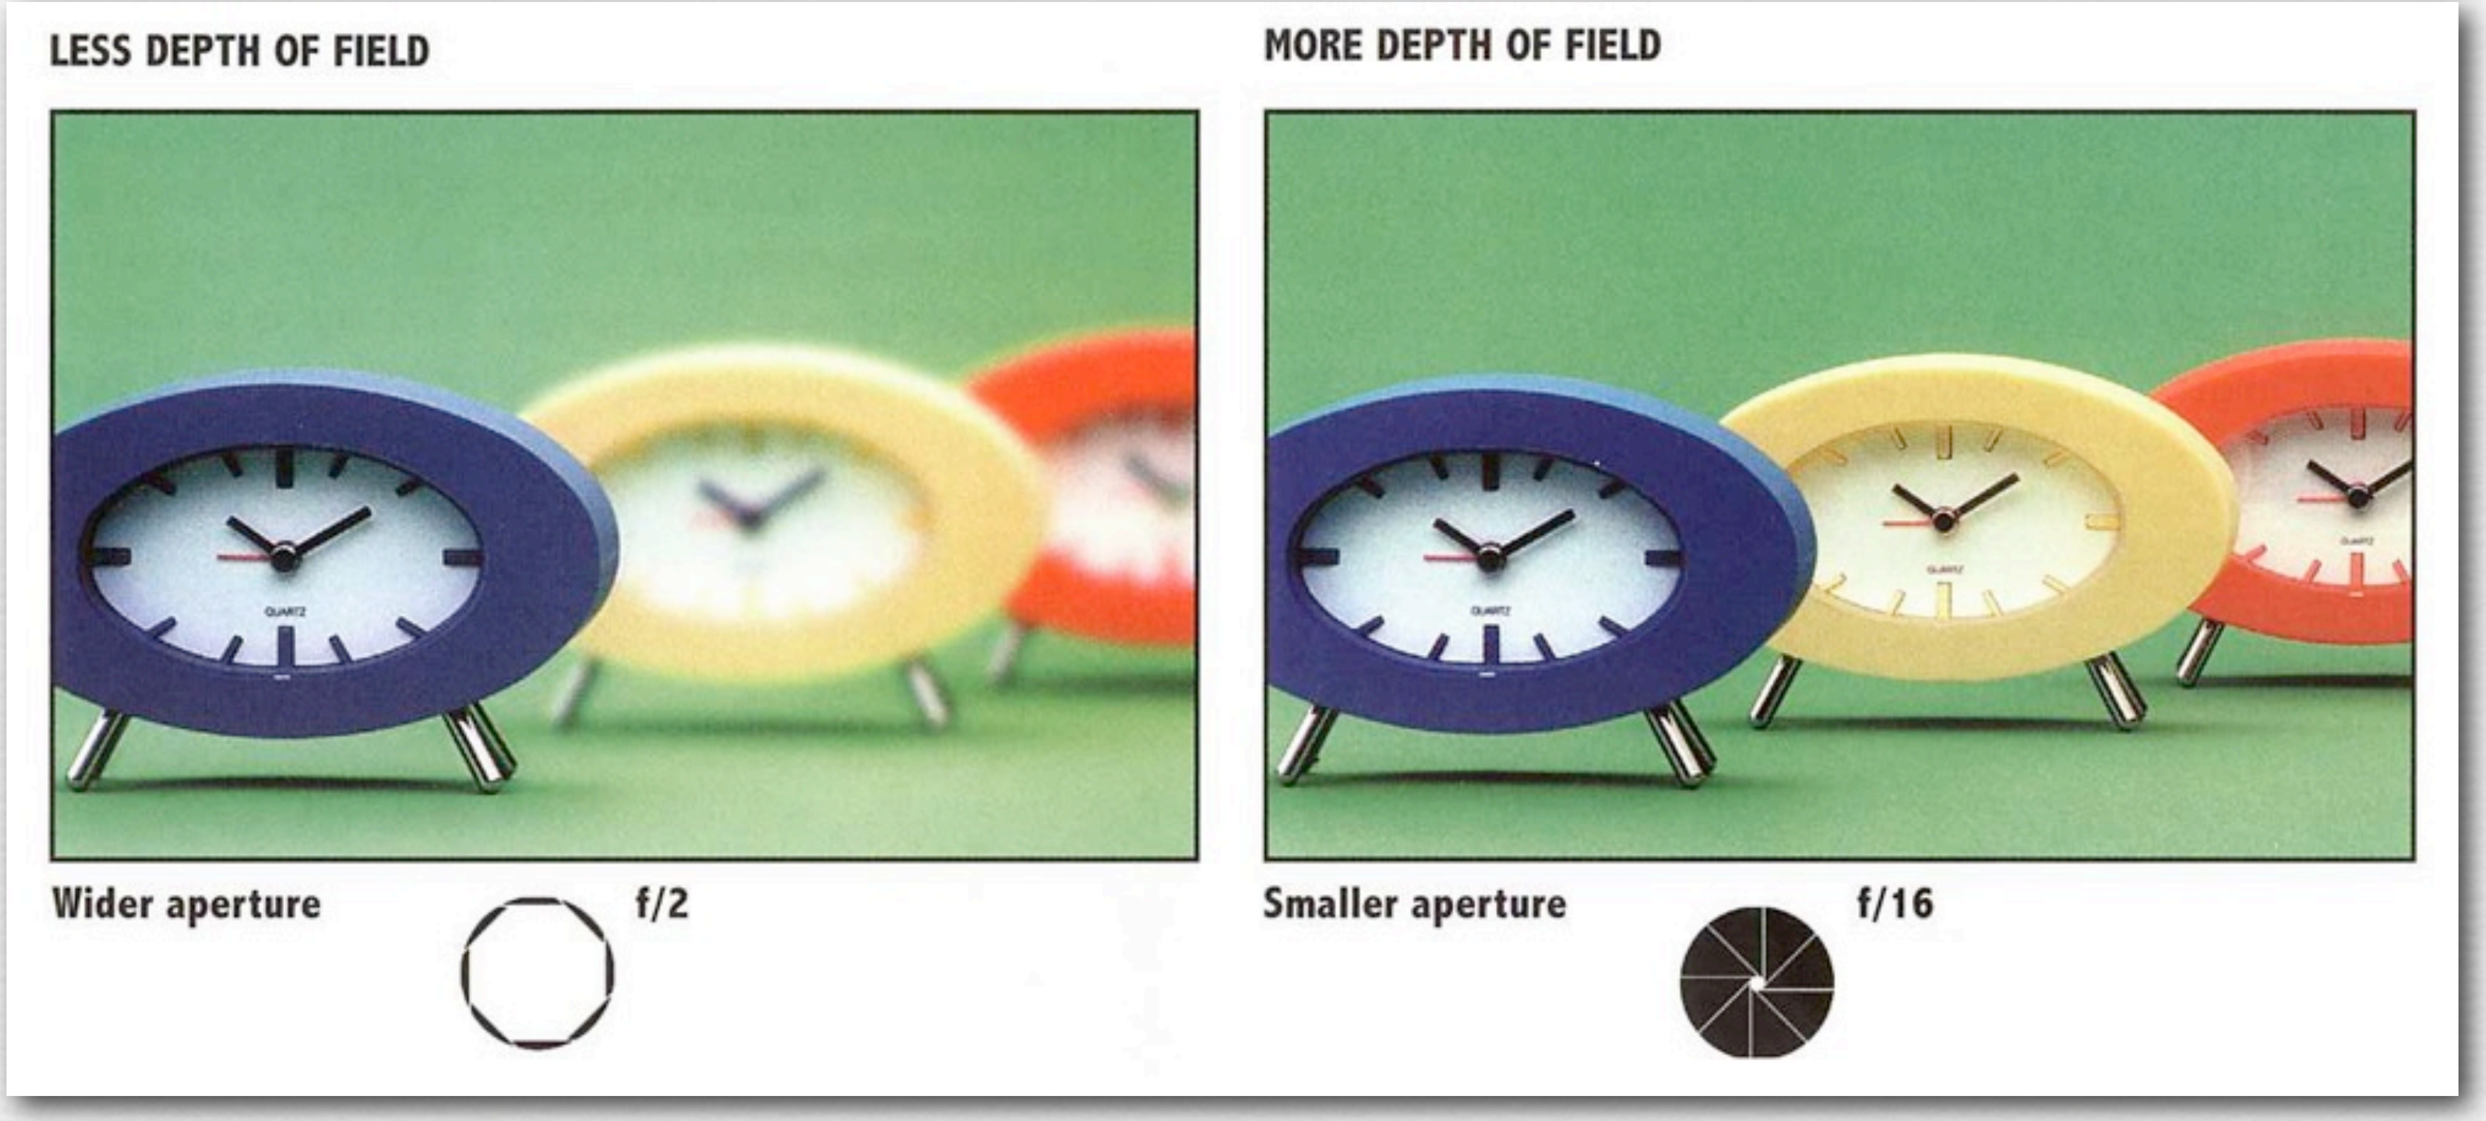

(London)

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

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

Circle of confusion (C)

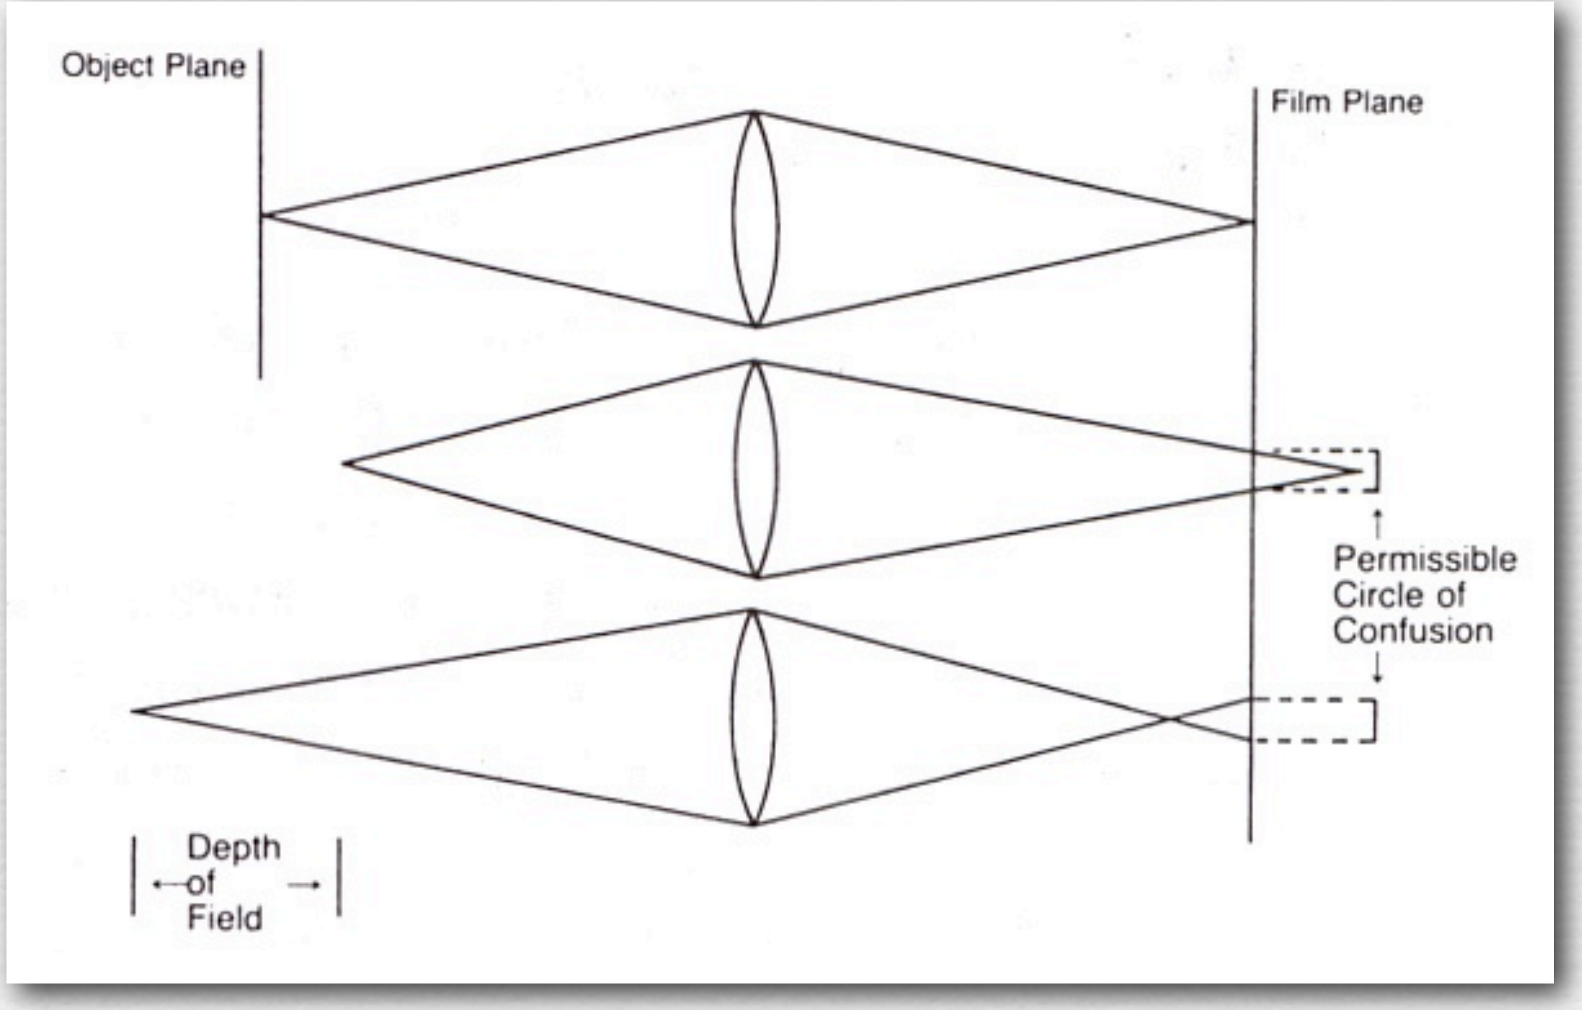

- C depends on sensing medium, reproduction medium, viewing distance, human vision,...
    - for print from 35mm film, 0.02mm (on negative) is typical
    - for high-end SLR, 6μ is typical (1 pixel)
    - larger if downsizing for web, or lens is poor

Chromatic aberration

red and blue have the same focal length

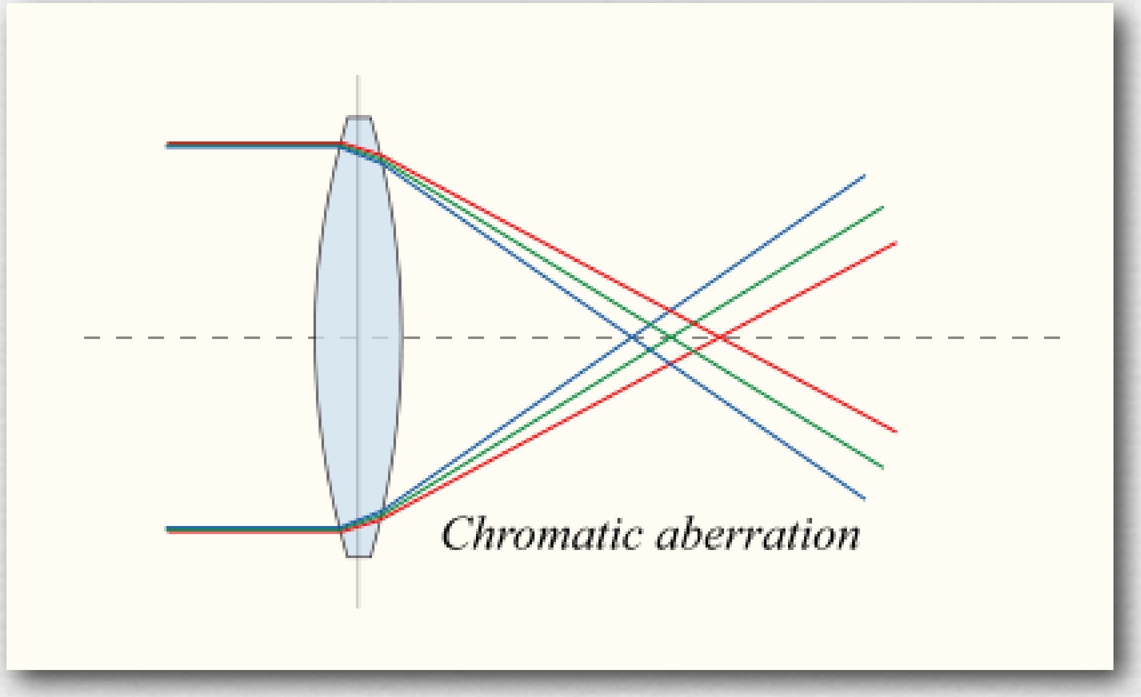

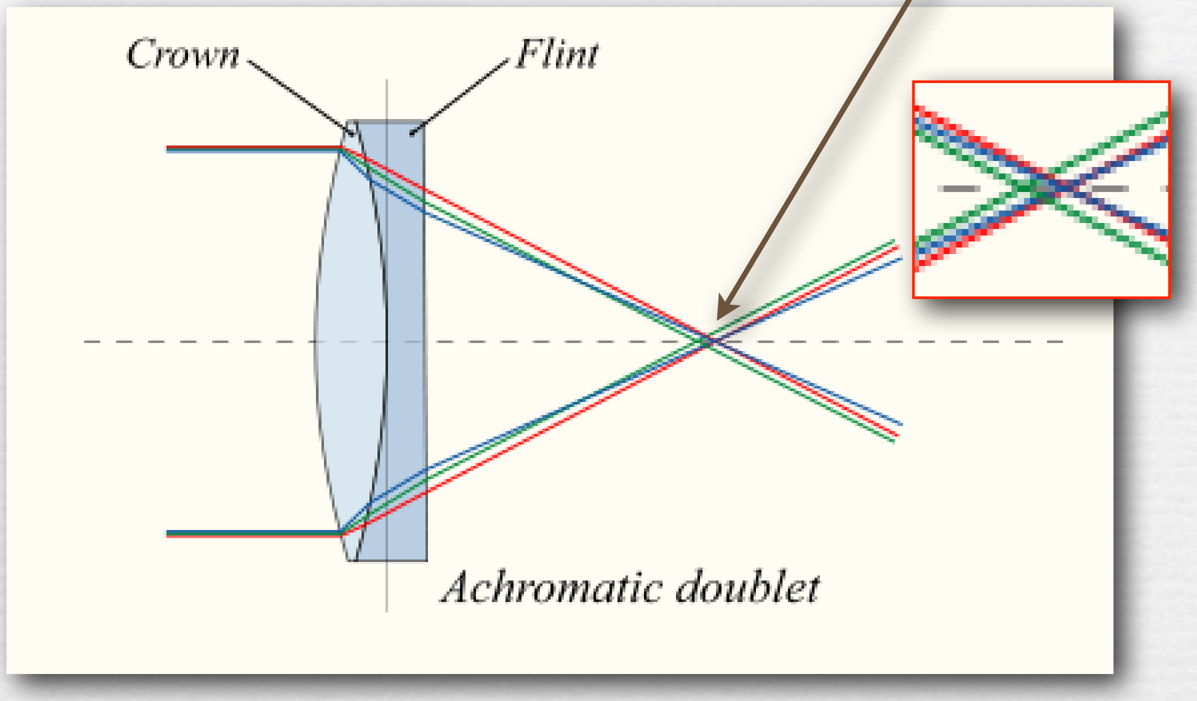

(wikipedia)

- dispersion causes focal length to vary with wavelength
  - for convex lens, blue focal length is shorter
- correct using *achromatic doublet*
  - strong positive lens + weak negative lens = weak positive compound lens
  - by adjusting dispersions, can correct at two wavelengths

Visible Light

(wikipedia)

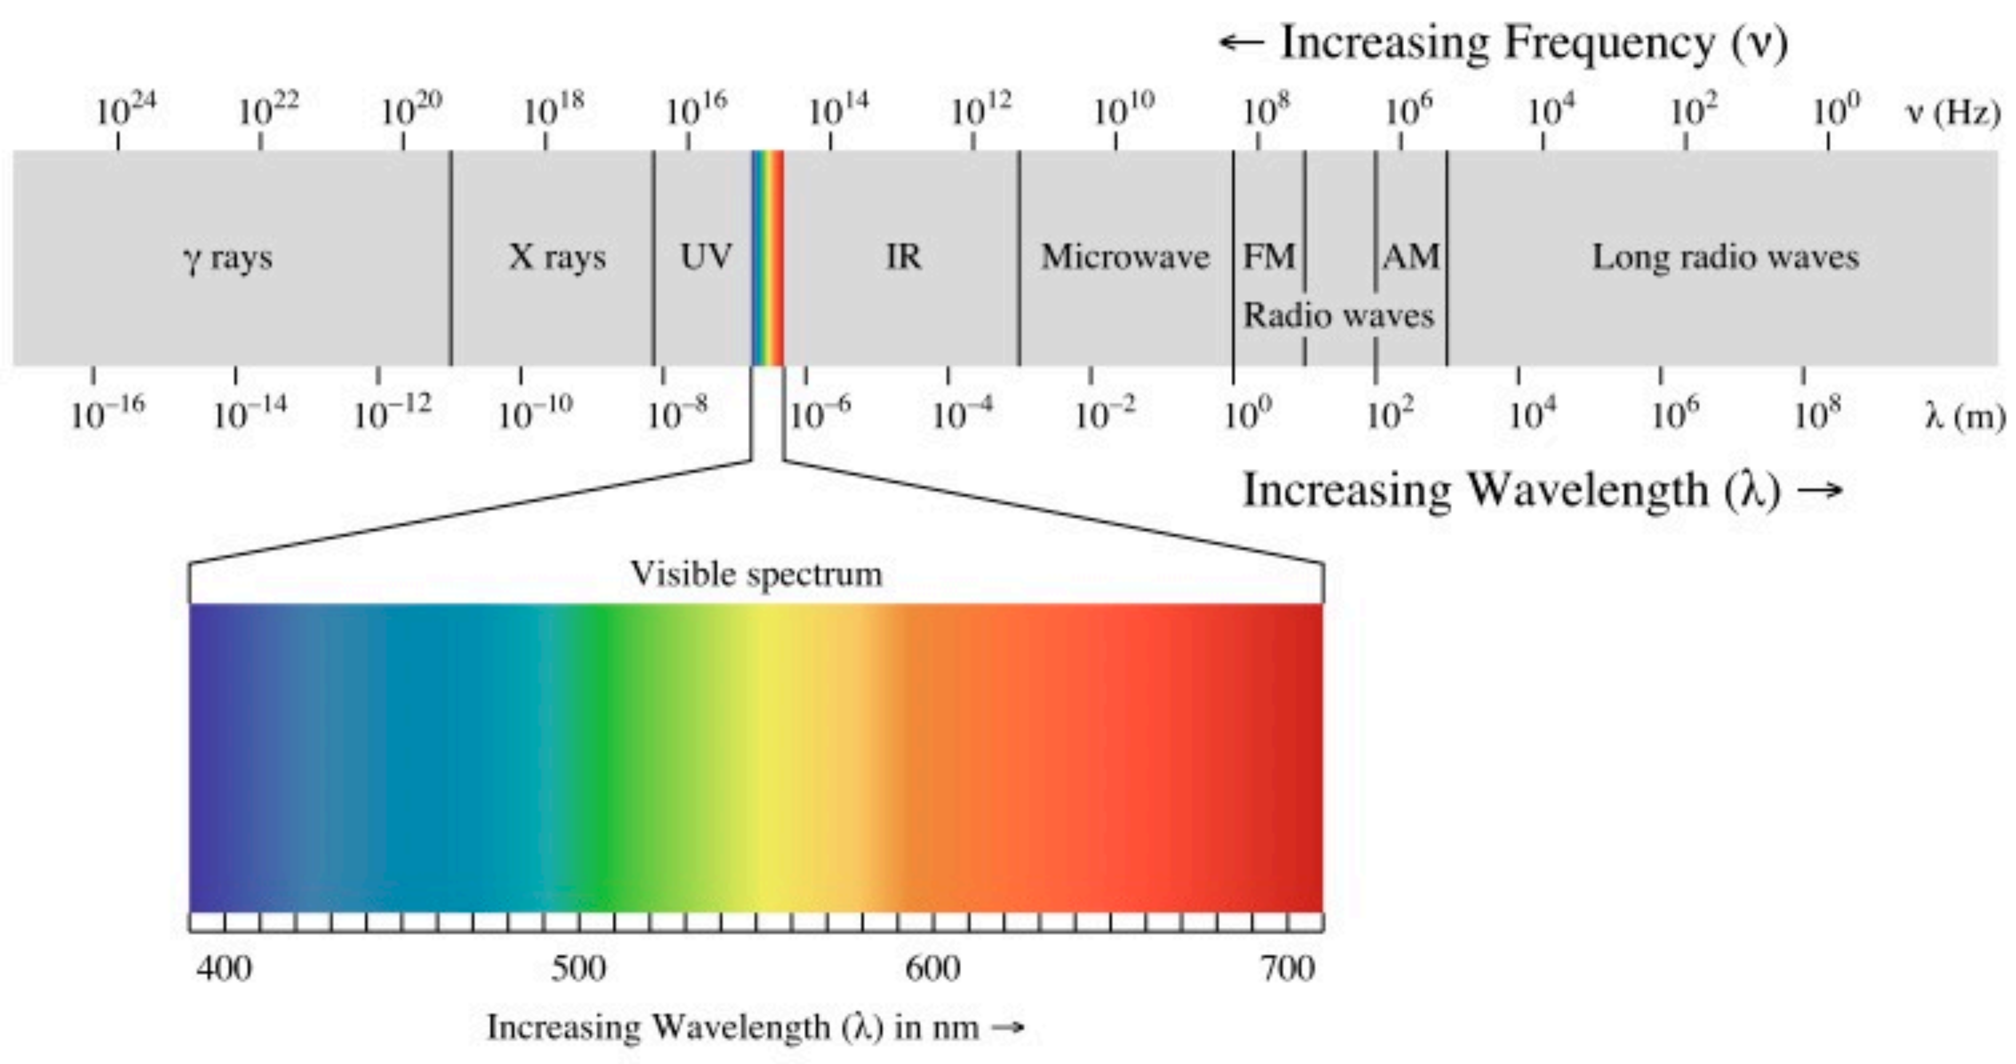

- wavelengths between 400nm and 700nm

Illumination

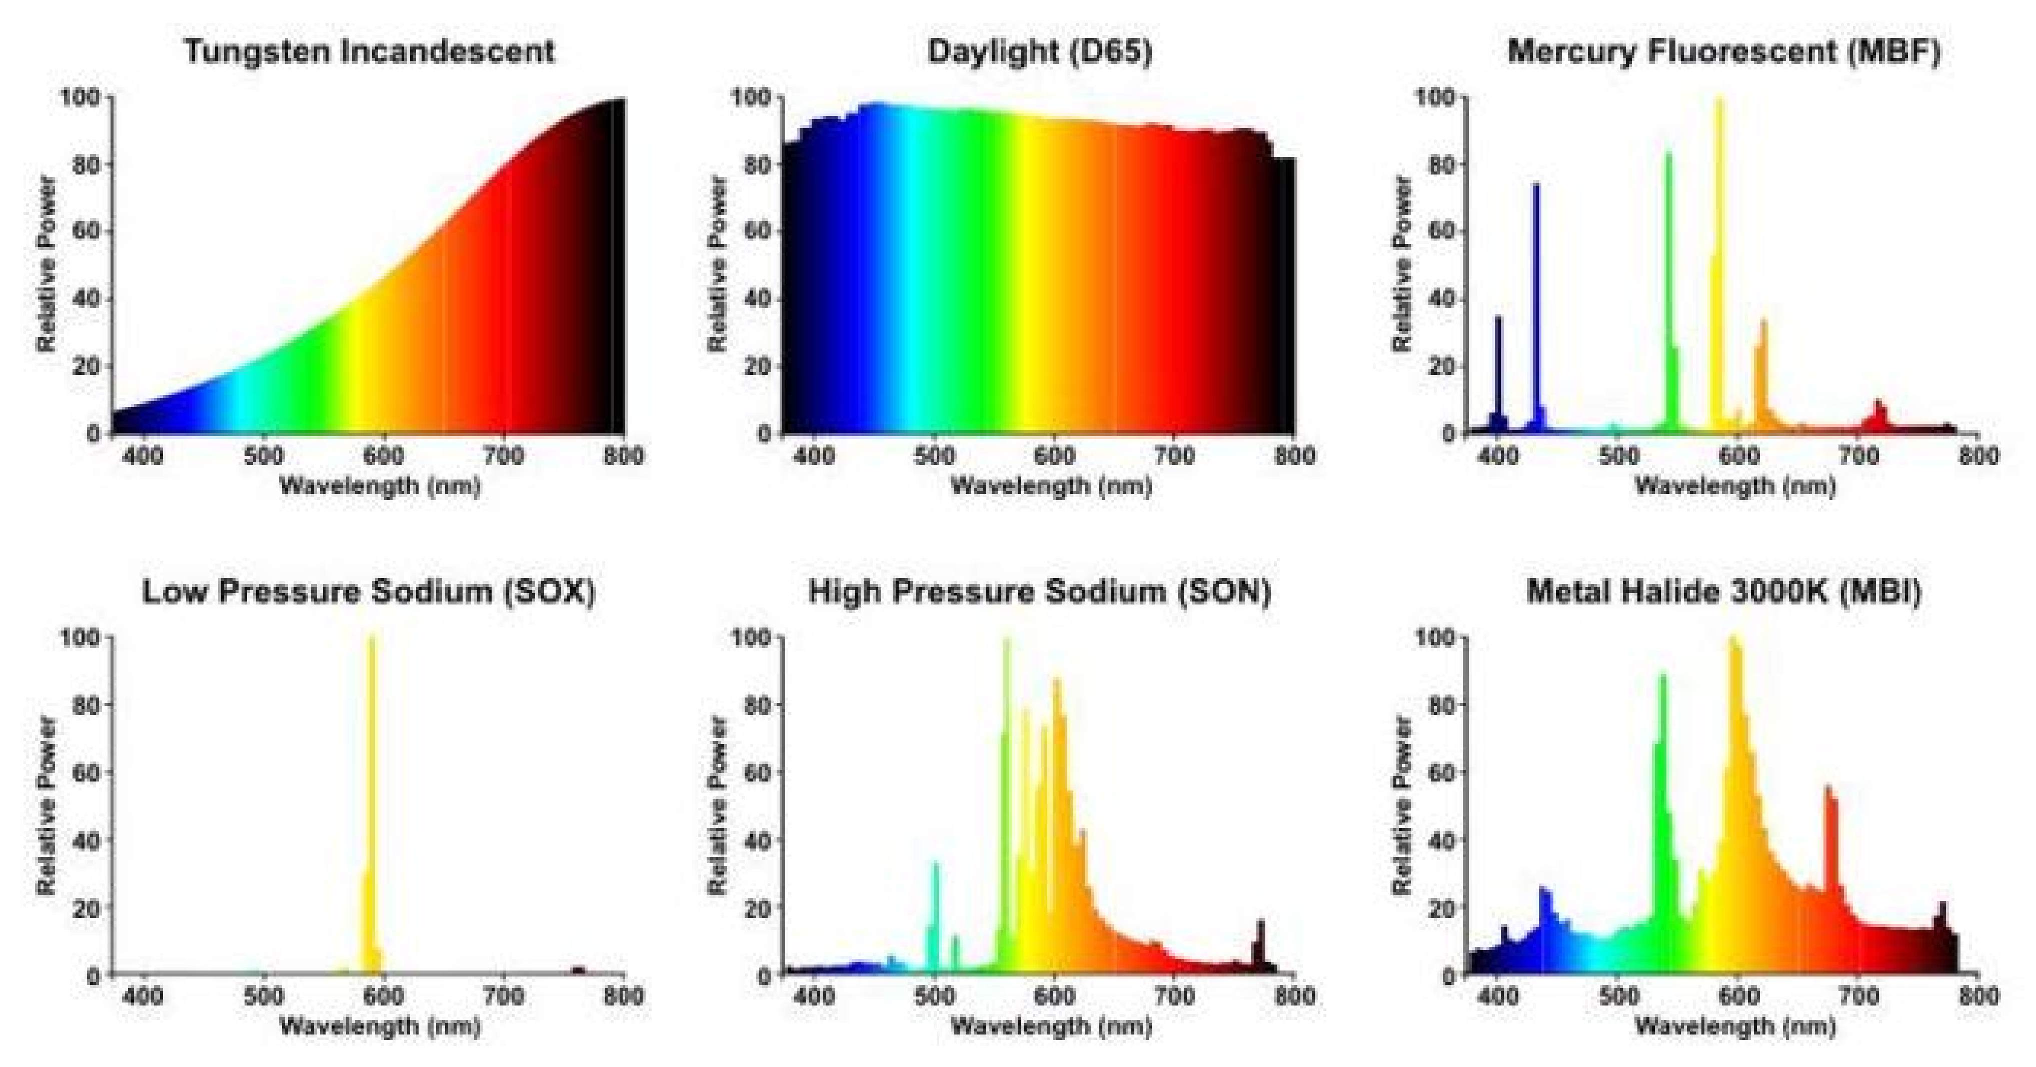

(LampTech)

- wavelengths between 400nm and 700nm

Trichromatic Vision

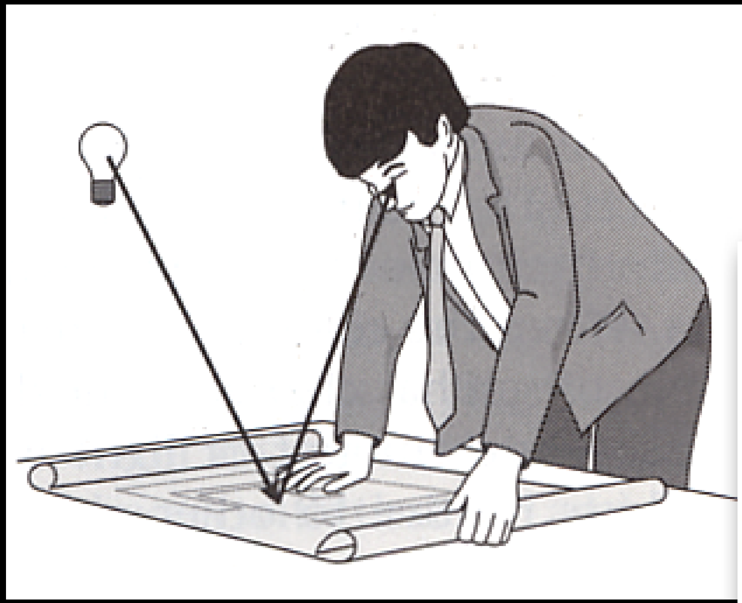

light is reflected by an object

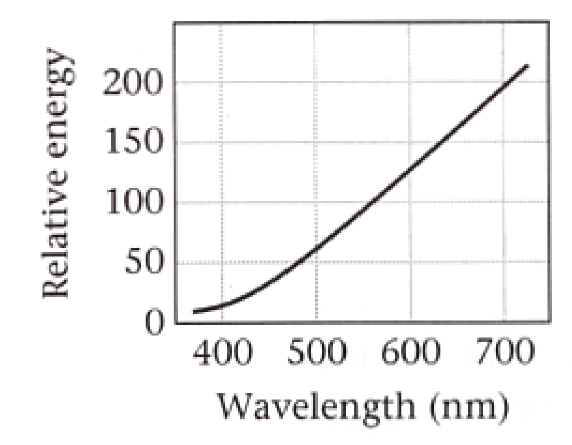

illumination
(Light)

×

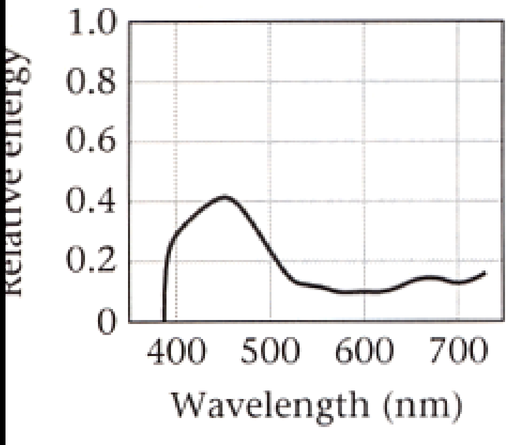

reflectance
(Object)

=

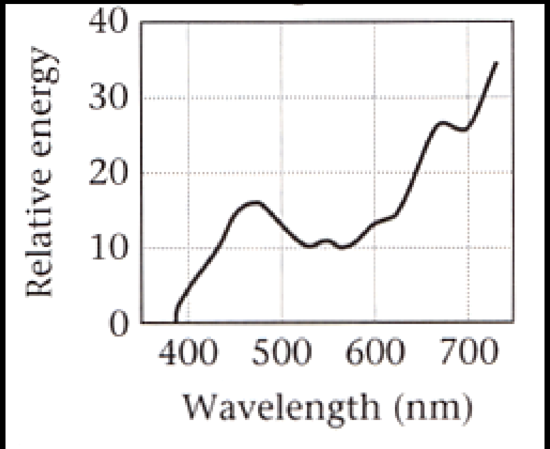

stimulus that enters your eye

Color Gamut: Consequences

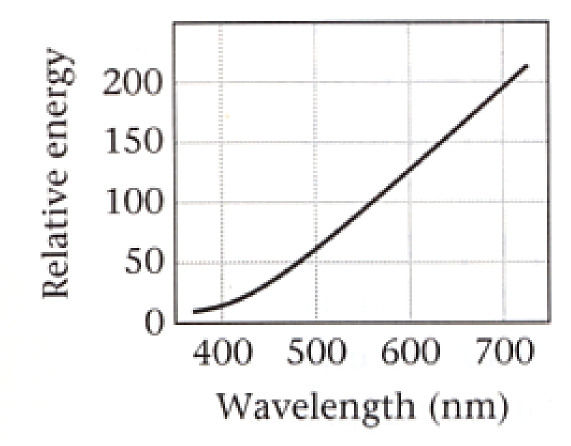

×

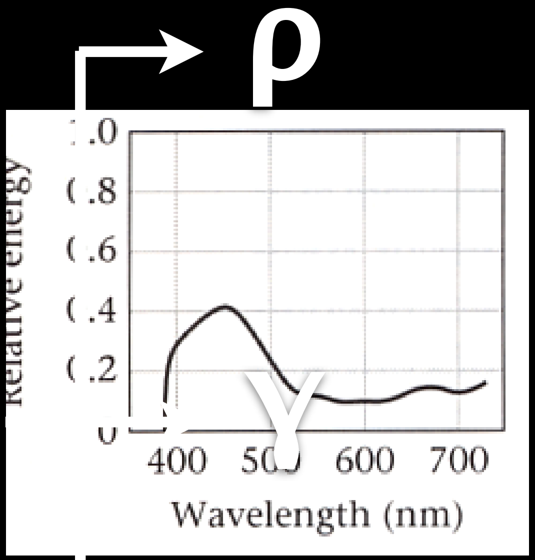

illumination
(Light)

reflectance
(Object)

β

- Goal of photography: reproduce the sensation of seeing a scene.

Color Primaries

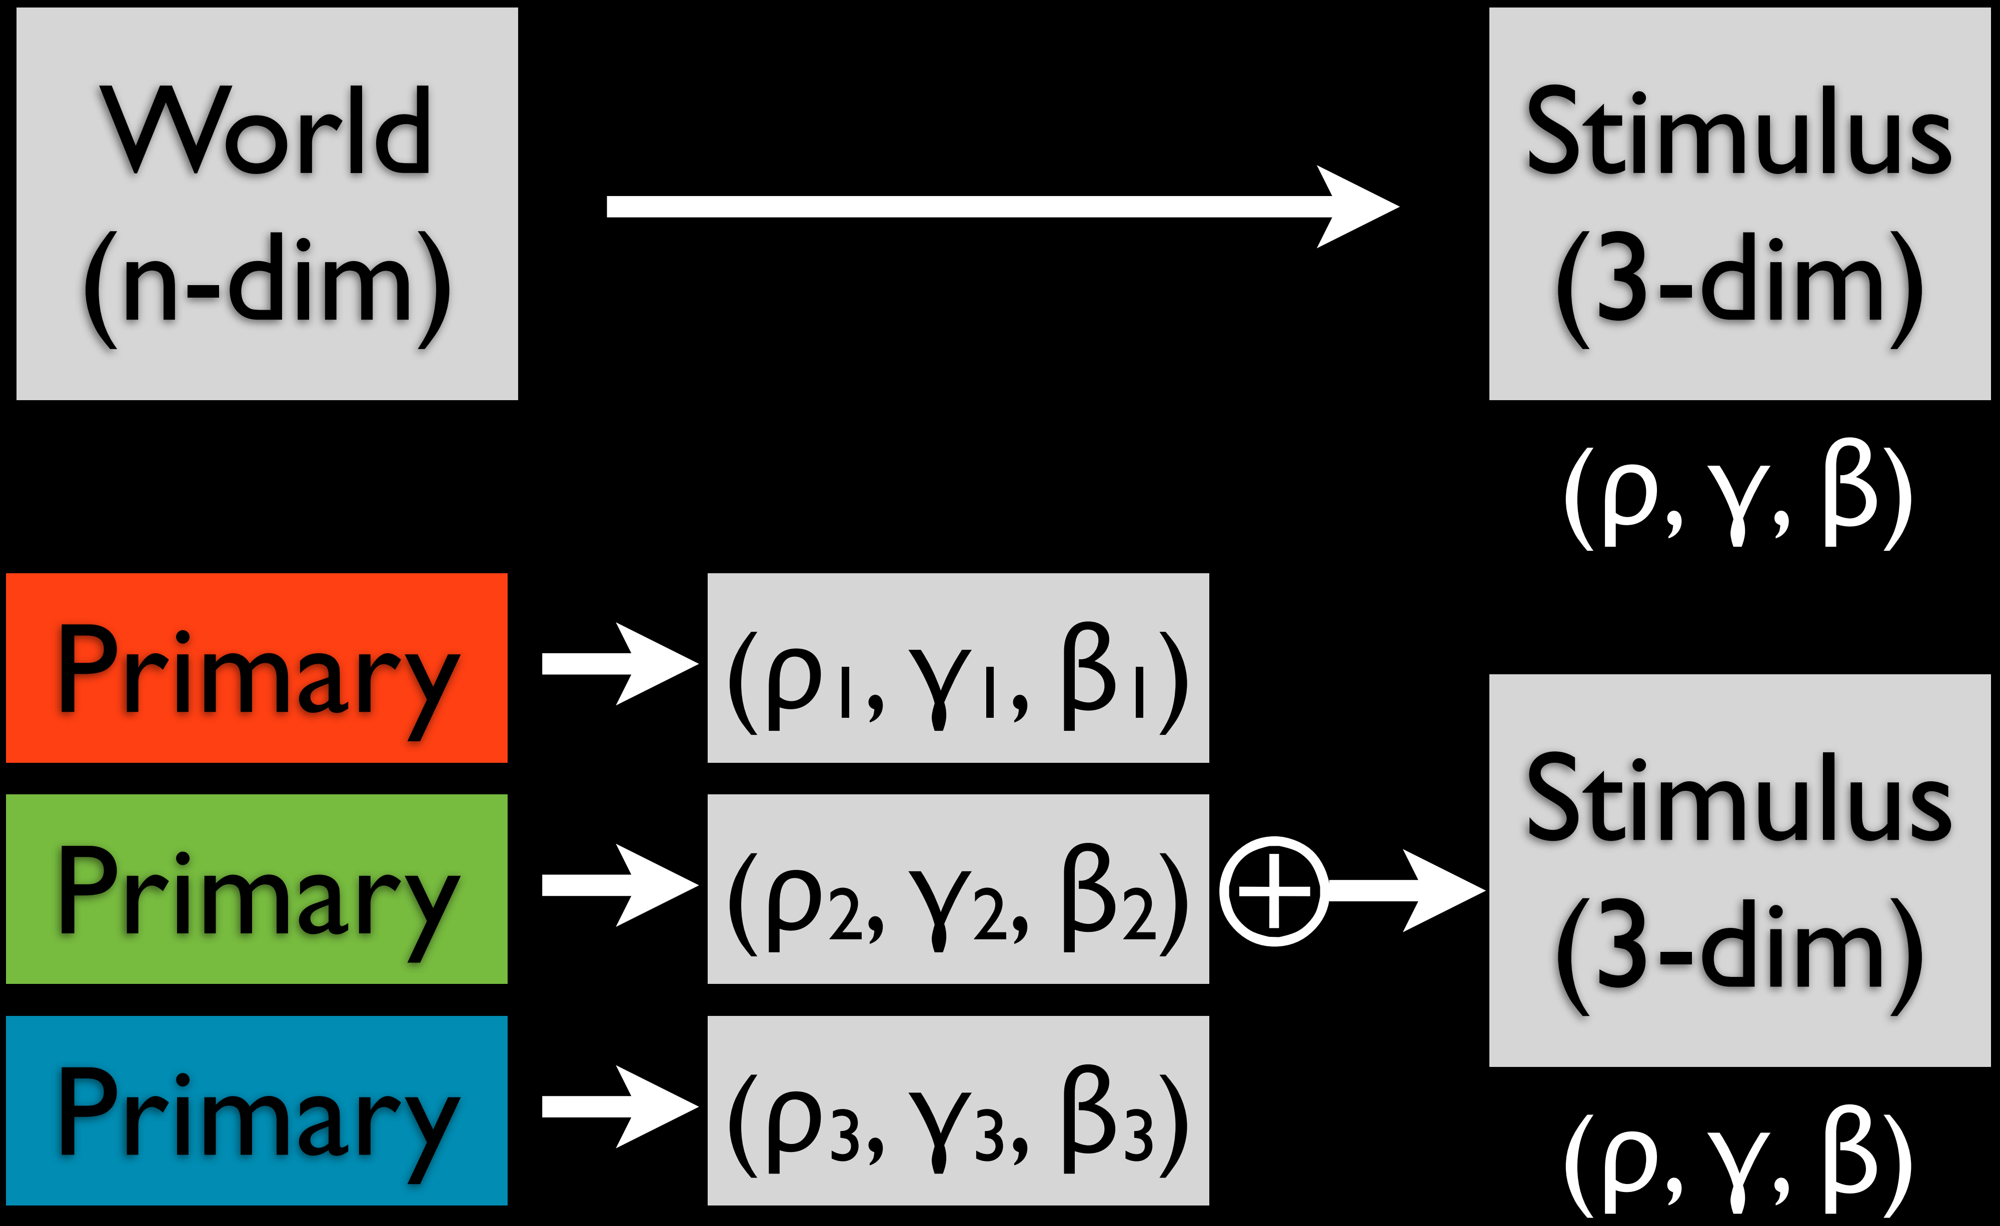

World (n-dim) → Stimulus (3-dim)

(ρ, γ, β)

Primary → $(\rho_1, \gamma_1, \beta_1)$

Primary → $(\rho_2, \gamma_2, \beta_2)$

Primary → $(\rho_3, \gamma_3, \beta_3)$

$\oplus \rightarrow$ Stimulus (3-dim)

(ρ, γ, β)

Color Primaries

- Choose three primaries R, G, B.
  - Does not have to be pure wavelengths.
- Normalize to obtain a desired *reference white*
  - This yields an *RGB cube*

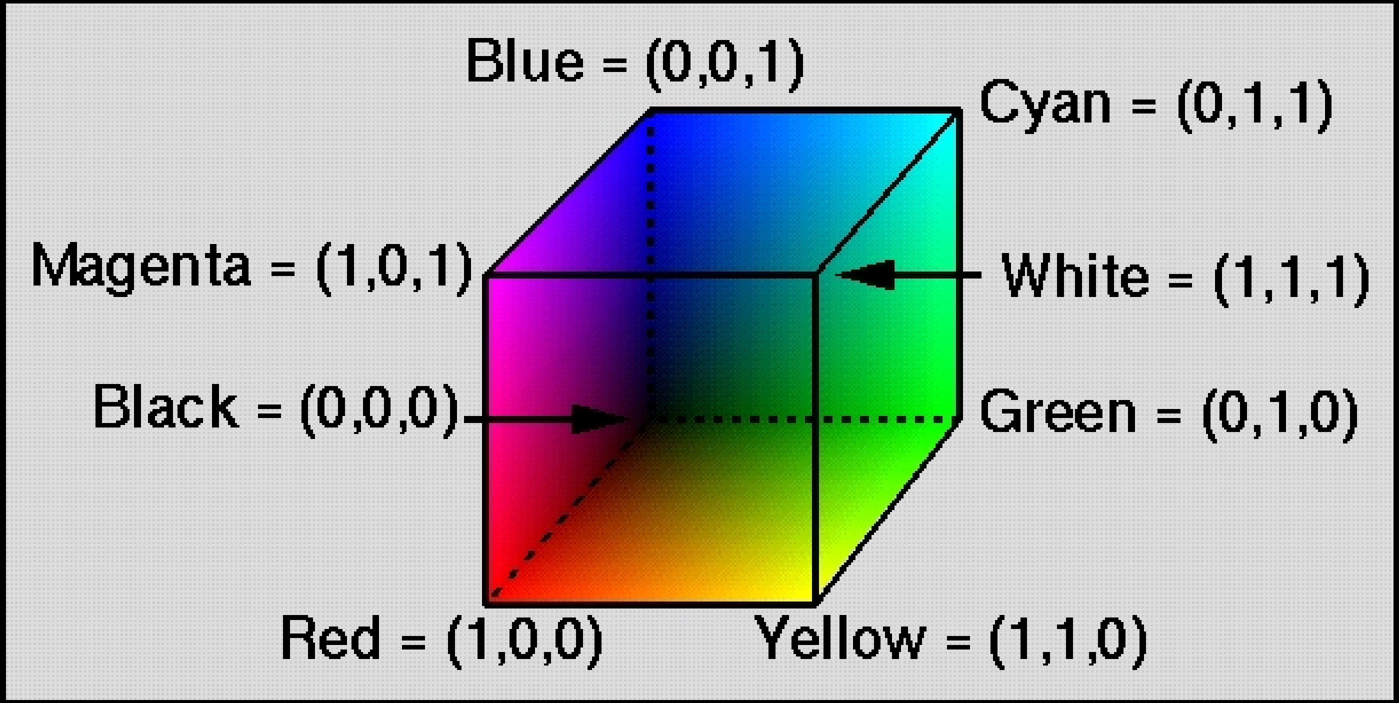

Color Primaries

- What exactly is R, G, B each?
  - Is there a specfic wavelength for each? No.
  - Is there a specific spectrum for each? Yes, but you can pick your own.

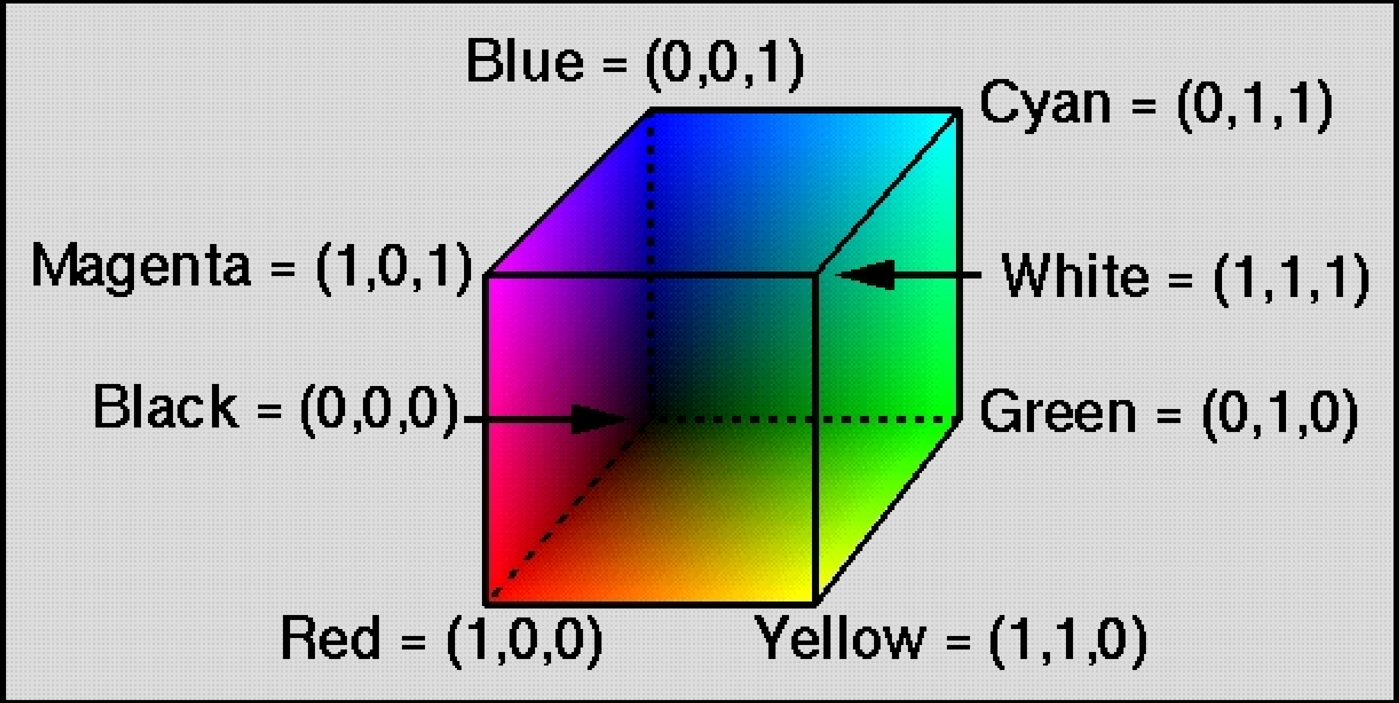

Choice of Primaries

- sRGB (HP, Microsoft, 1996)
- Adobe RGB (Adobe, 1998)
- Adobe Wide-Gamut RGB
- ...

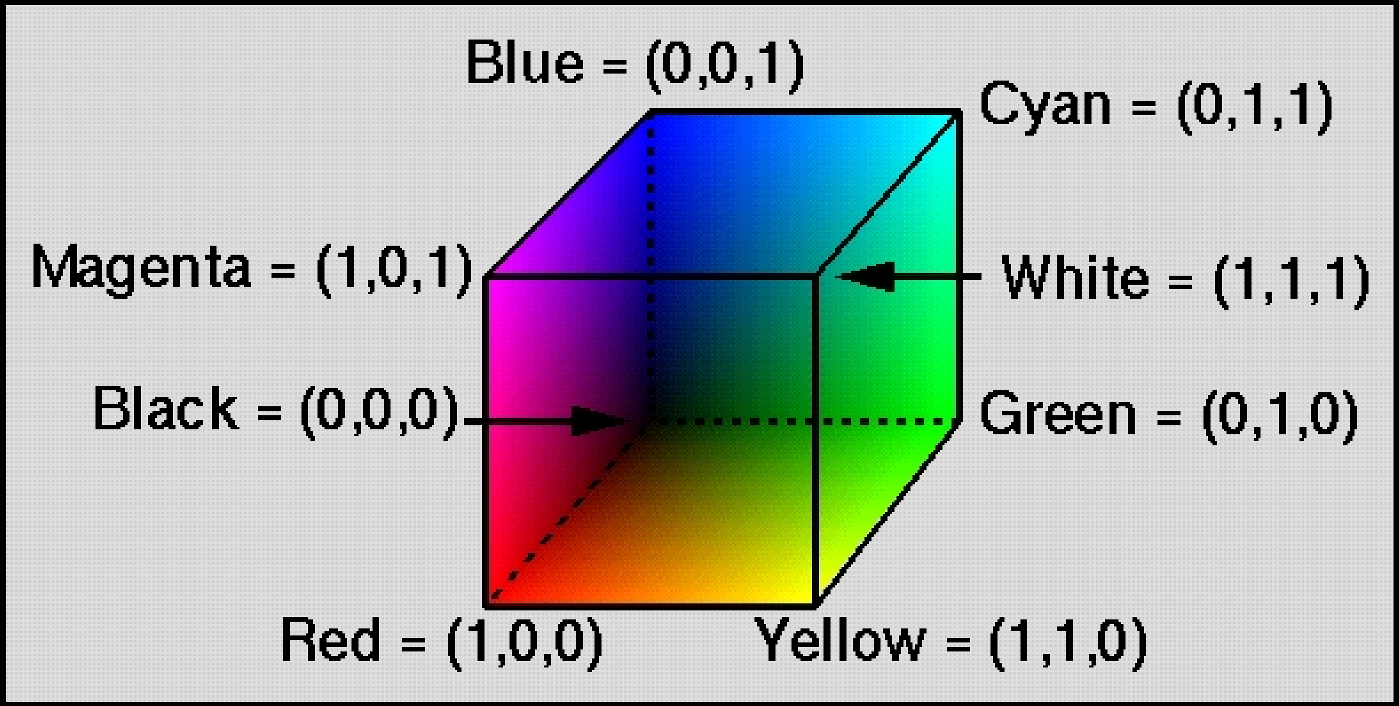

Example Pipeline

sensor → analog to digital conversion (ADC) → processing: demosaicing, tone mapping & white balancing, denoising & sharpening, compression → storage

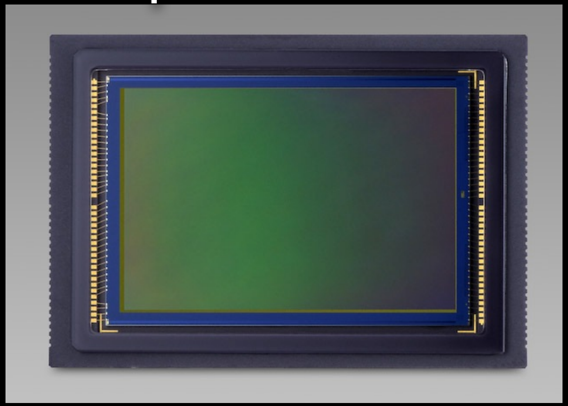

Canon 21 Mpix CMOS sensor

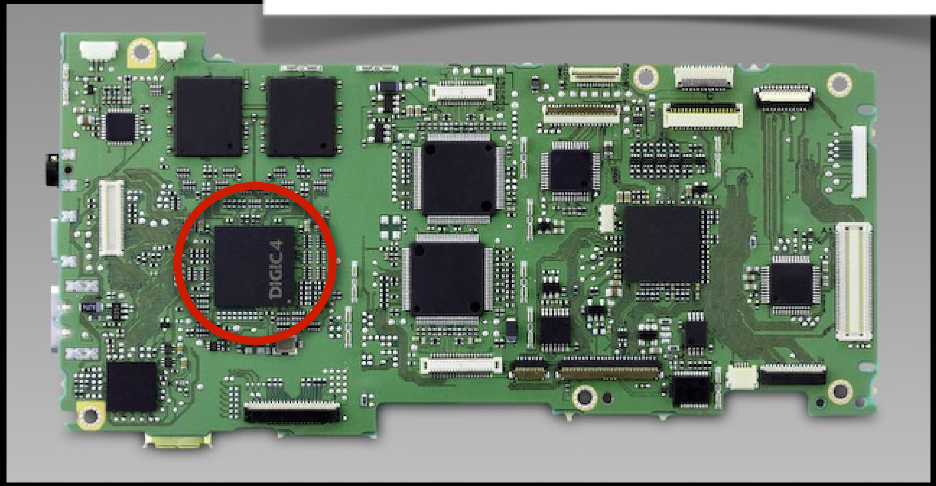

Canon DIGIC 4 processor

Compact Flash card

The Science

- Photoelectric Effect
  - Materials may generate electrons upon being hit by a photon.
- Quantum Efficiency
  - Not all photons will produce an electron.

back-illuminated CMOS (Sony)

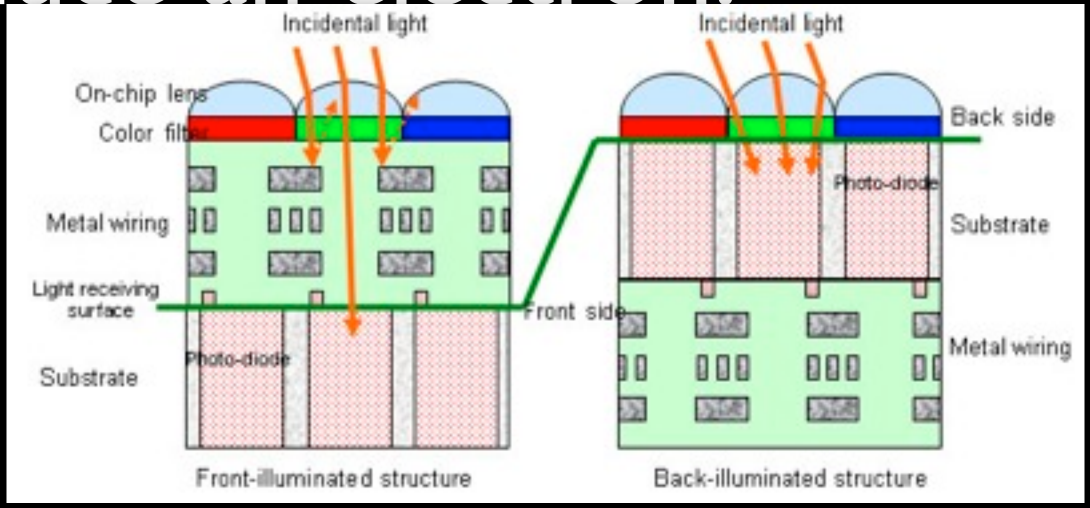

The Pixel

- Size matters
  - Casio EX-F1: 2.5µ x 2.5µ
  - Nokia N900: 3.1µ x 3.1µ
  - Canon 5D II: 6.4µ x 6.4µ
- Capacity matters

CMOS vs. CCD

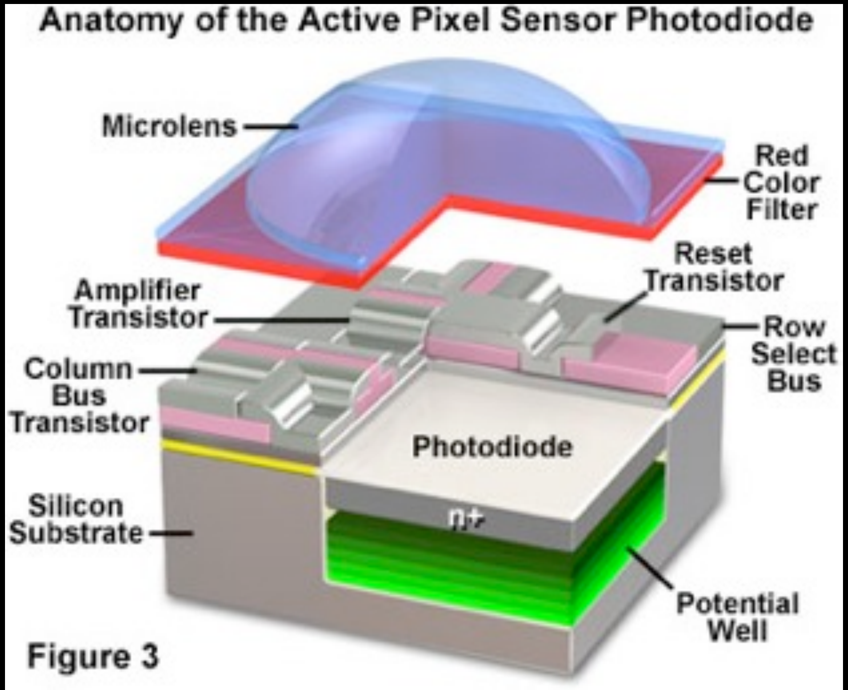

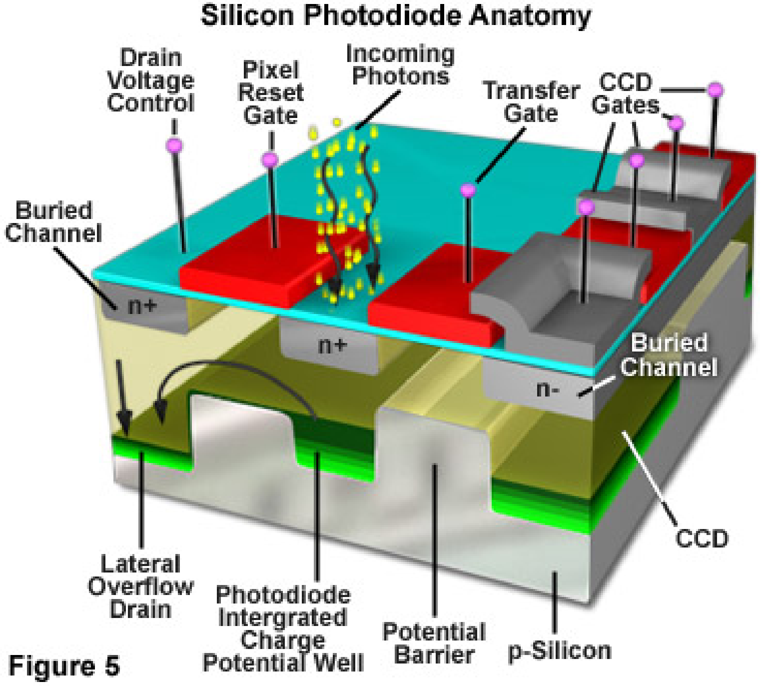

- Complimentary Metal-Oxide Semiconductor
  - per-pixel amplifier converts charges to voltage.
  - cheap, low-power but noisy

- Charge-Coupled Device
  - charge shifts along column to an amplifier
  - good but not as cheap.

Color Filter Arrays

- Recall: we need information on (ρ, γ, β).
- Need discrimination among multiple wavelengths
  - Three types (of spectral sensitivity) of pixels would be sufficient.
- Color filter array: turns pixels into one of three types.

Bayer Pattern

- Checkered pattern of green and alternating red/blue
- Pretty much everywhere

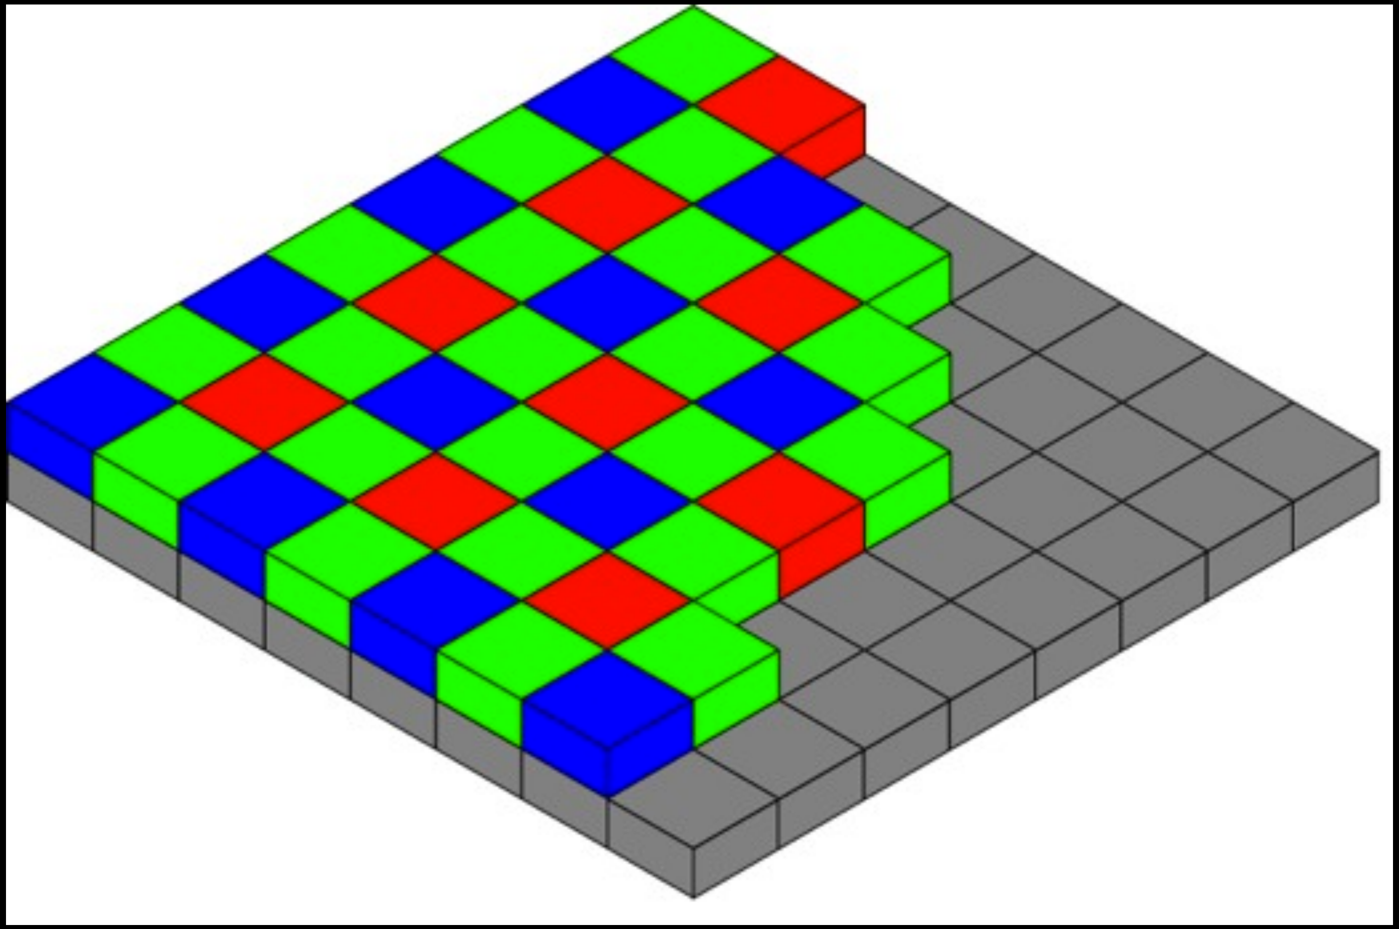

Bayer Pattern

- Checkered pattern of green and alternating red/blue
- Pretty much everywhere

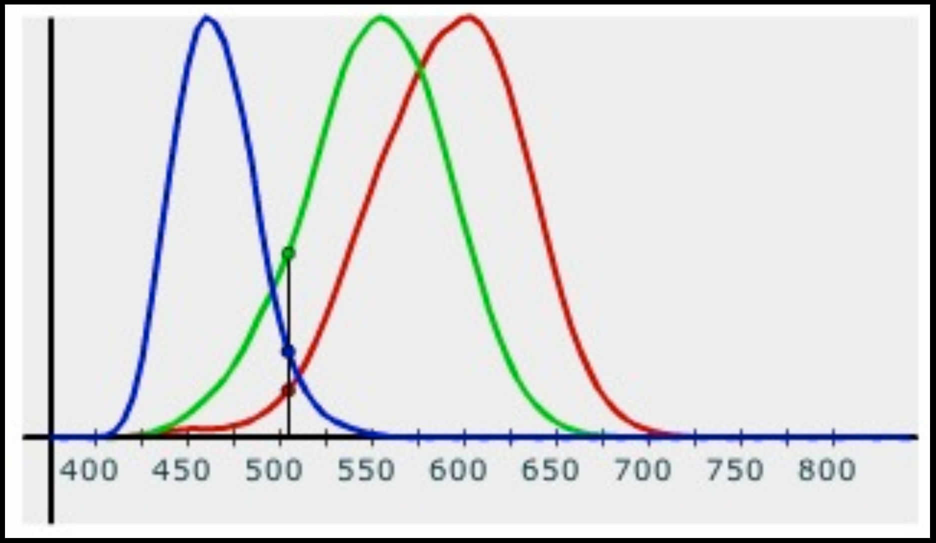

Cone cells

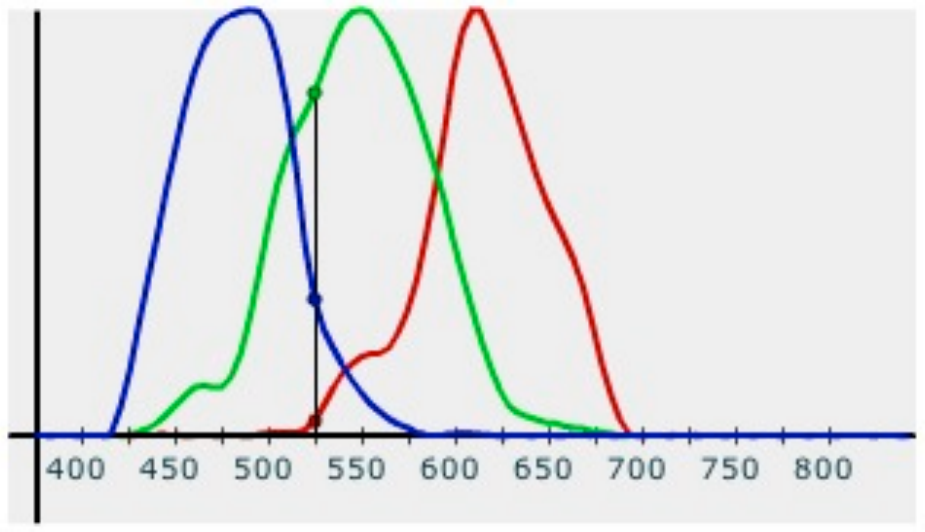

Color filters in Canon 30D

Foveon Sensor

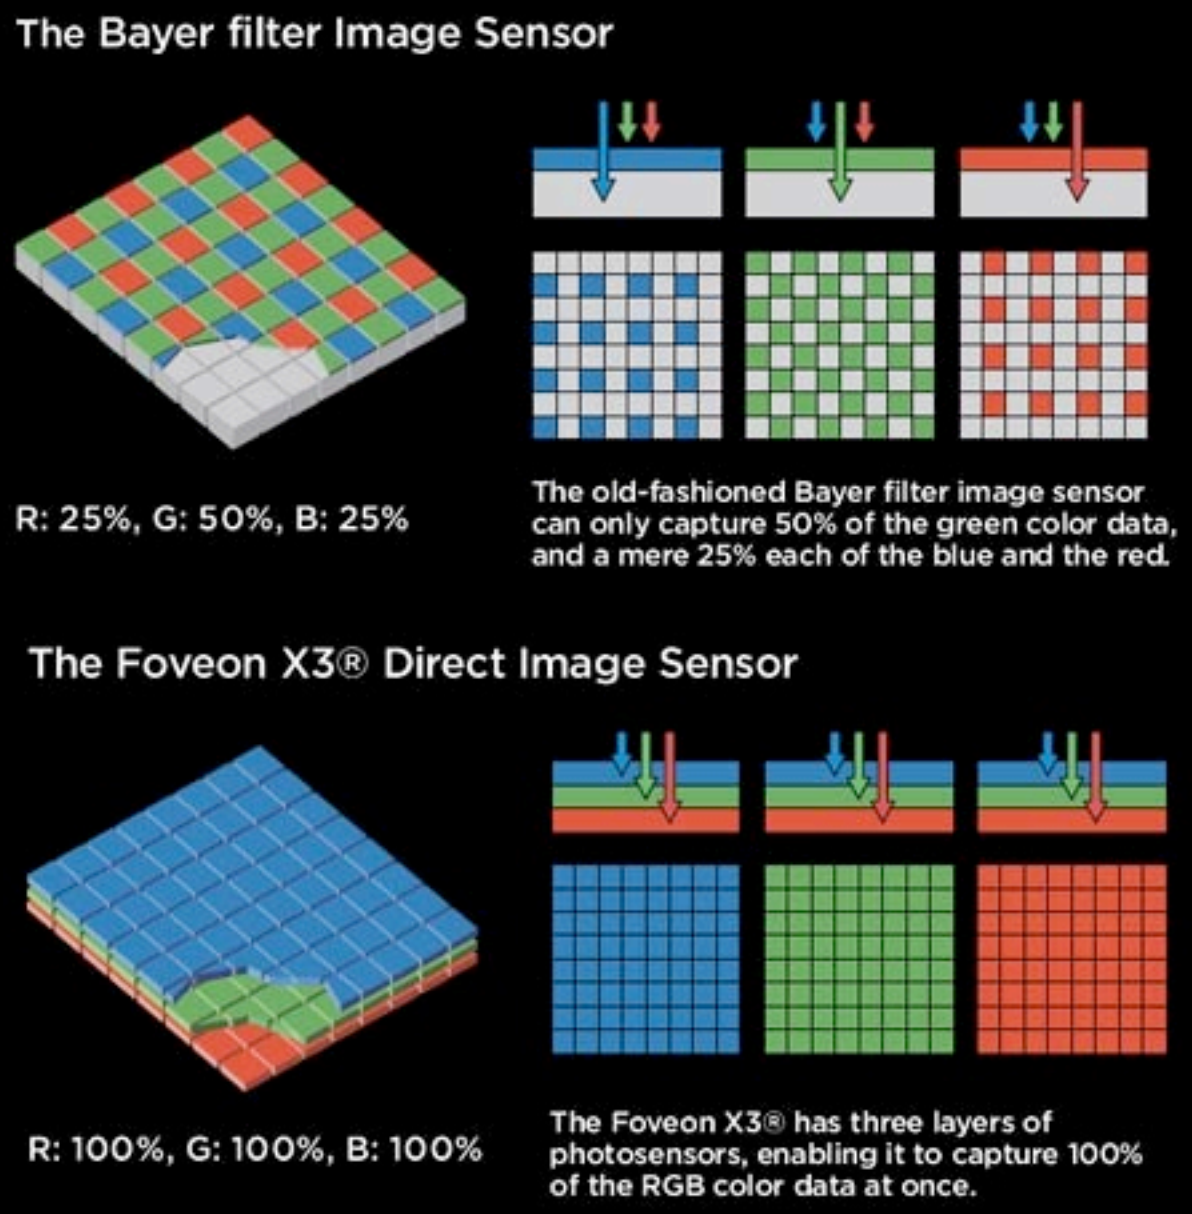

The Bayer filter Image Sensor

R: 25%, G: 50%, B: 25%

The old-fashioned Bayer filter image sensor can only capture 50% of the green color data, and a mere 25% each of the blue and the red.

The Foveon X3® Direct Image Sensor

R: 100%, G: 100%, B: 100%

The Foveon X3® has three layers of photosensors, enabling it to capture 100% of the RGB color data at once.

Analog-to-Digital Conversion

- Convert analog voltage to discrete values.

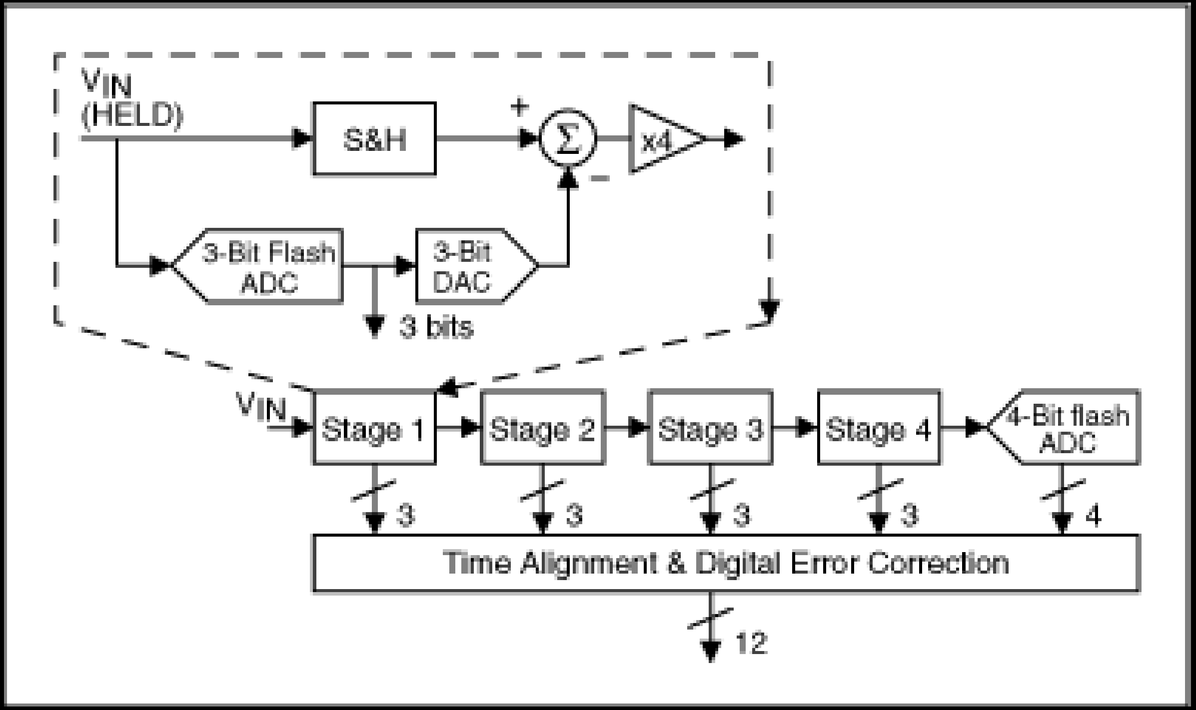

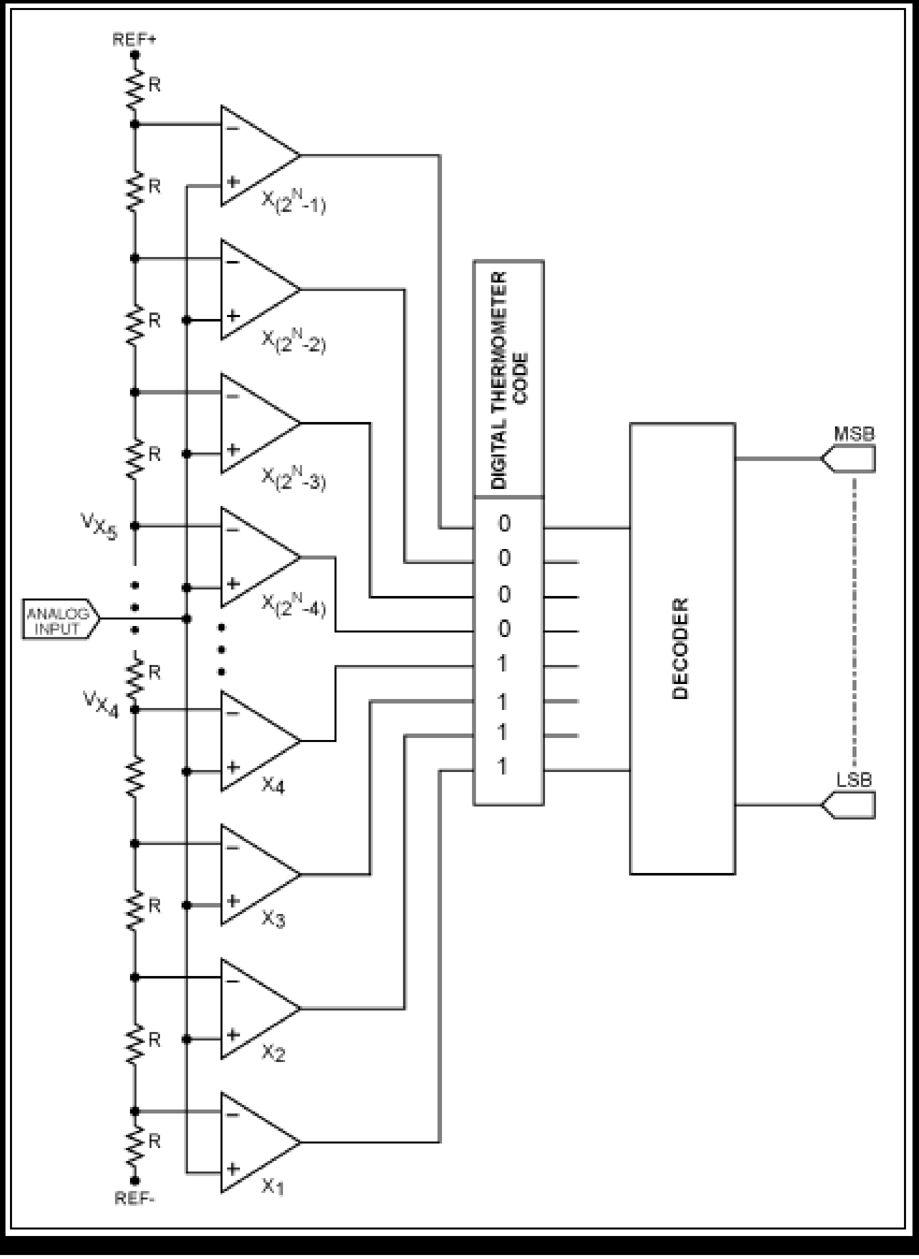

Noise: Summary

- Photon shot noise
- Hot pixels
- Dark current
- Fixed pattern noise
- Read noise
- Pixel non-uniformity
- ...

Much of the literature treats these altogether as a Gaussian noise

Signal v. Noise

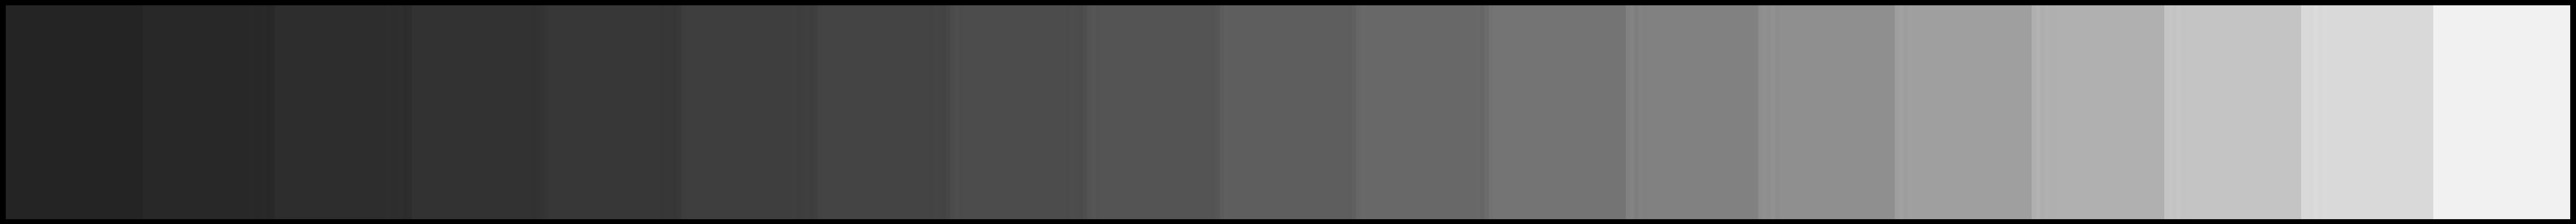

Test Chart

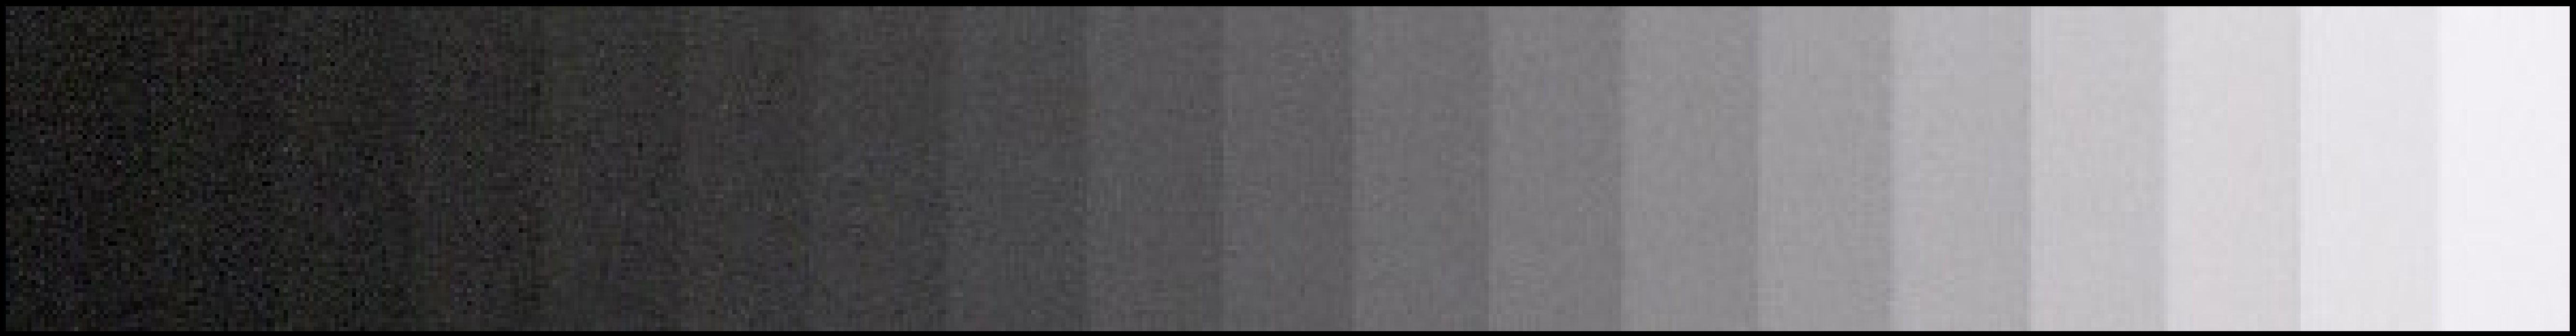

Captured by Canon 10D (ISO 1600)

Photon Shot Noise

- Pixels measure the # of incident photons.
    - Upon a fixed area, during a fixed time.

- Varies from time to time.
- Varies from pixel to pixel.
    - Follows the Poisson distribution.

Dark Current

- Electrons dislodged by random thermal activity.
  - Increases linearly with exposure time.
  - Increases exponentially with temperature..

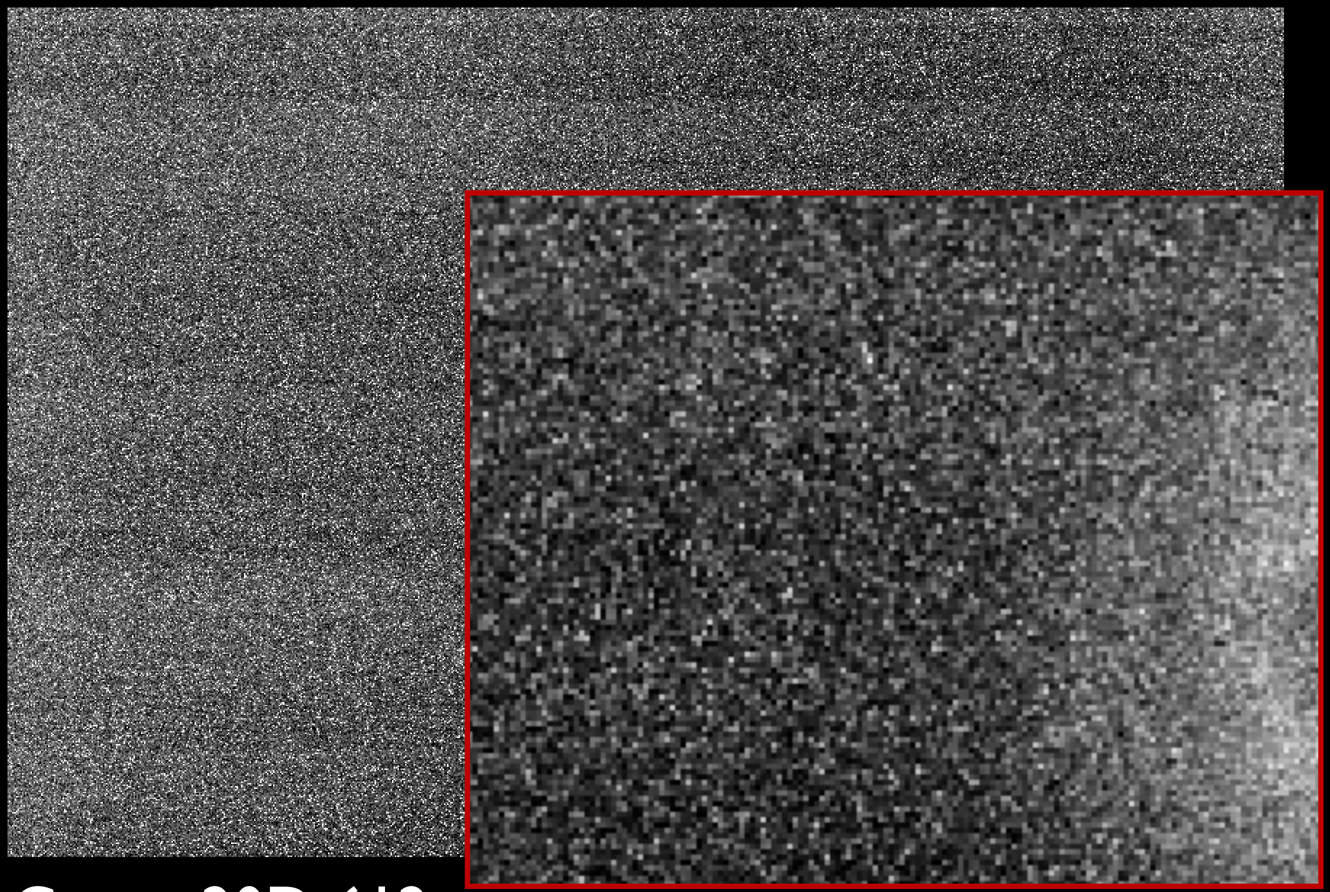

Canon 20D, 612s exposure

Hot Pixels

- Electrons leaking into wells because of manufacturing defects
  - Increases linearly with exposure time.

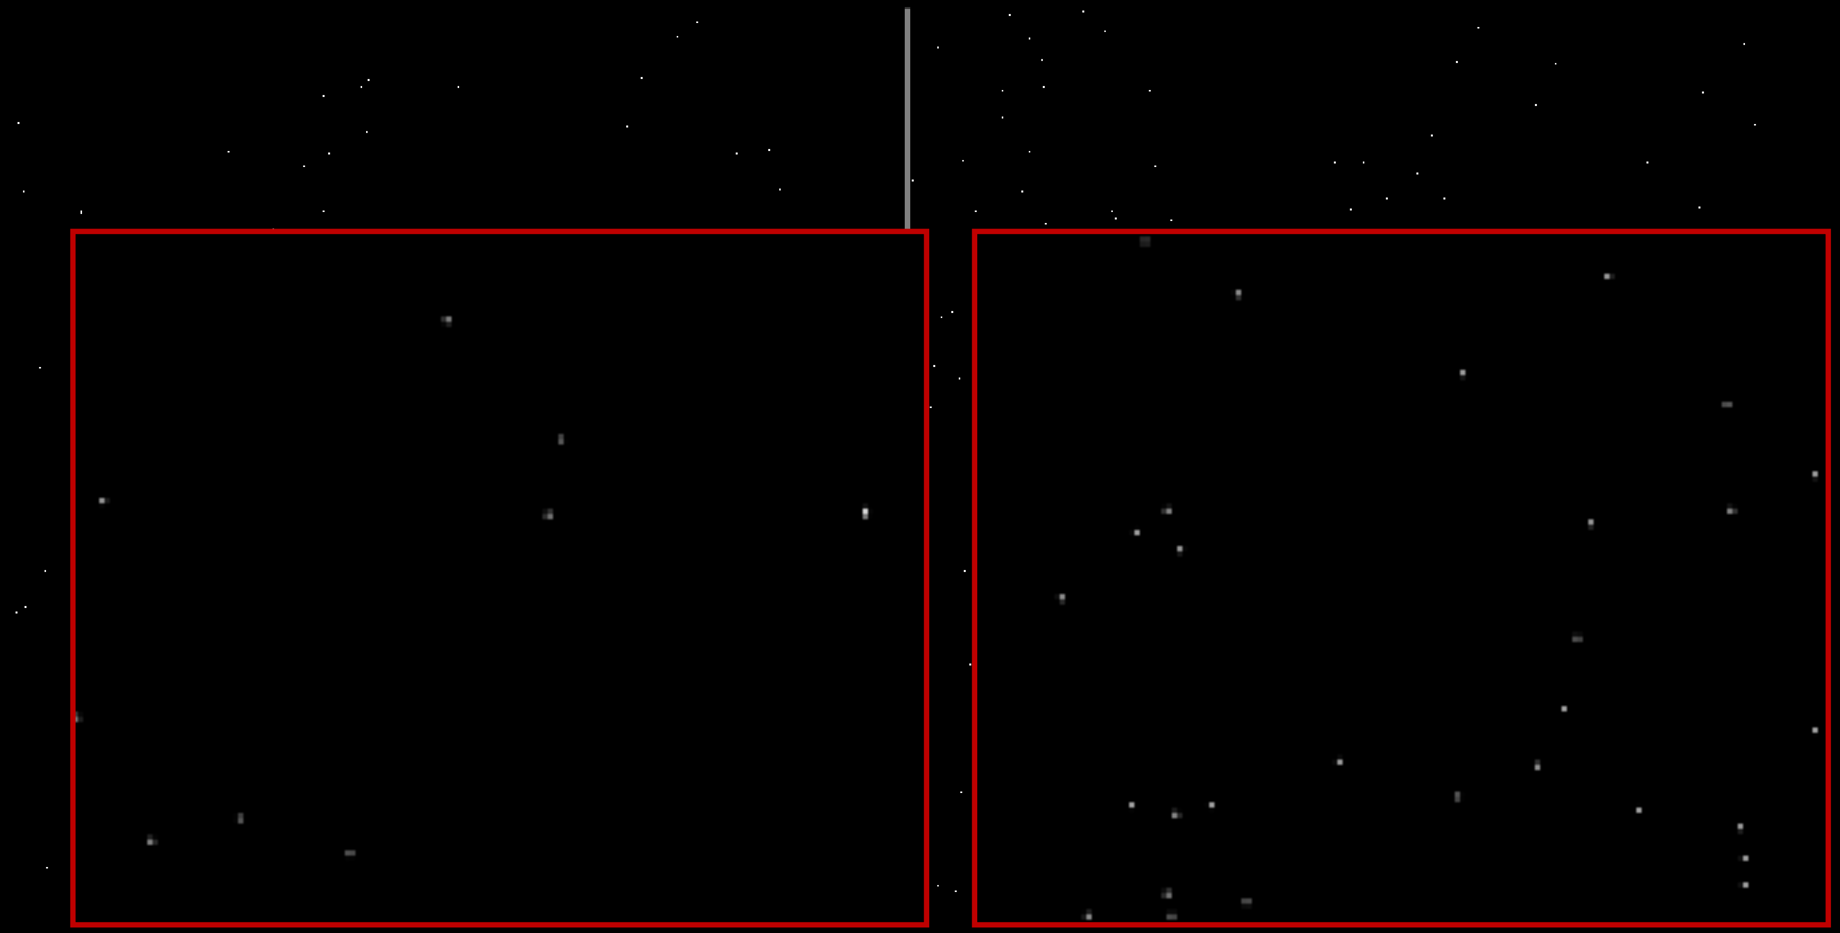

Canon 20D, 15s/30s exposure

Fixed Pattern Noise

- Manufacturing variations across pixels, columns, etc
  - Constant over time

Canon 20D, ISO 800, cropped

Read Noise

- Thermal noise in readout circuitry
  - Mainly in CMOS

Canon 1D Mk III, cropped

Pixel Response Non-Uniformity

- ~1% variance in the sensivity of pixels
  - Think about it as a per-pixel vignetting issue.

Quantization Error

- Any ADC process has quantization errors.
    - Depends on the bitdepth of the ADC.

Electronic Interference

- Interference from other circuitry
  - Exacerbated by poor insulation