Digital Photography I

optics and sensors

CSC2529
David Lindell
University of Toronto
cs.toronto.edu/~lindell/teaching/2529

*slides adapted from Gordon Wetzstein, Fredo Durand, Ioannis Gkioulekas, Marc Levoy

Announcements

- HW 1 is due Friday at 11:59pm
- HW 2 is out (due next Friday 9/27)

- TA office hours Wed 1-2pm BA 5256

https://web.cs.toronto.edu/news-events/events/colloquium-in-applied-ai

Michal Irani

Title: Reading Minds & Machines

1. Can we reconstruct images that a person saw, directly from his/her fMRI brain recordings?

2. Can we reconstruct the training data that a deep-network trained on, directly from the parameters of the network?

The answer to both of these intriguing questions is “Yes!”

Toronto Vision Seminar

FIELDS

Scan to register and join the meeting

Michal Irani

Professor of Computer Science and Applied Mathematics

"Deep Internal Learning" – Deep Learning Without Prior Examples

Friday, Sep 13th

12:00pm–1:00pm

BA3200

Reception to follow

Let’s say we have a sensor…

digital sensor
(CCD or
CMOS)

… and an object we like to photograph

What would an image taken like this look like?

Bare-sensor imaging

real-world object

digital sensor (CCD or CMOS)

Bare-sensor imaging

real-world
object

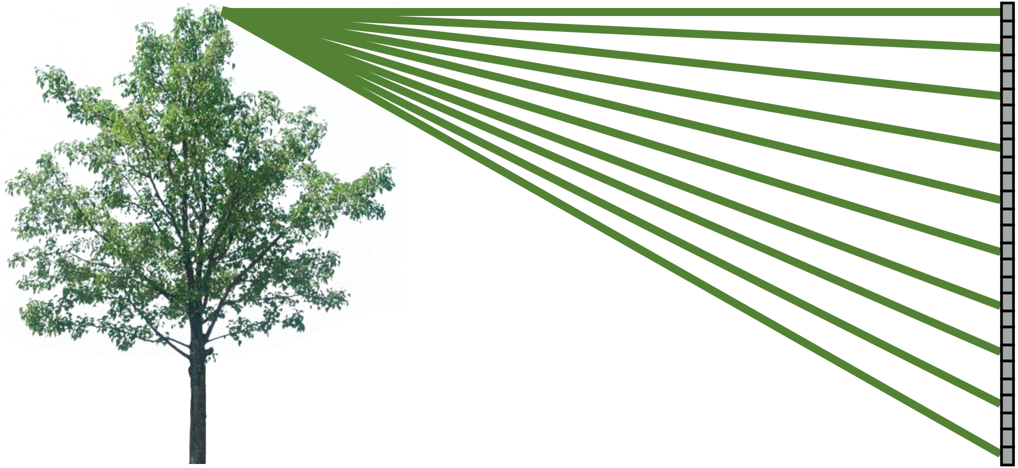

digital sensor
(CCD or
CMOS)

Bare-sensor imaging

real-world object

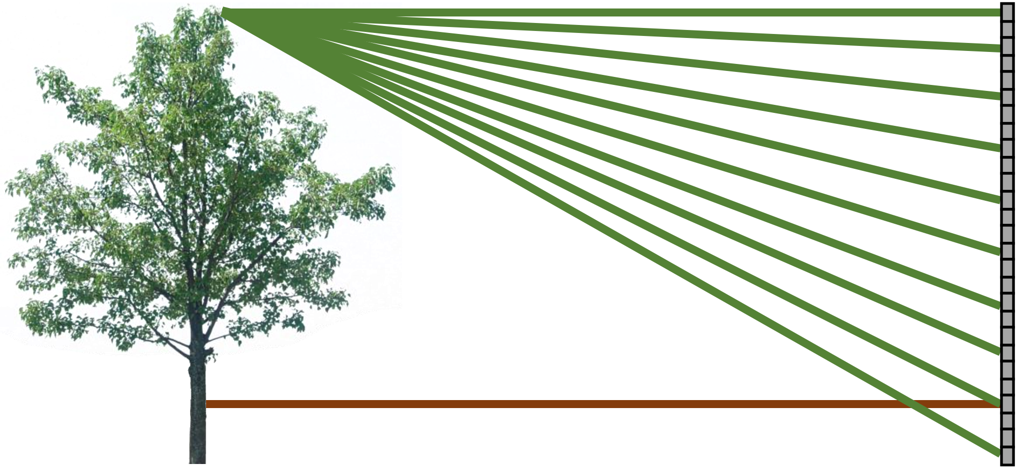

digital sensor (CCD or CMOS)

Bare-sensor imaging

real-world
object

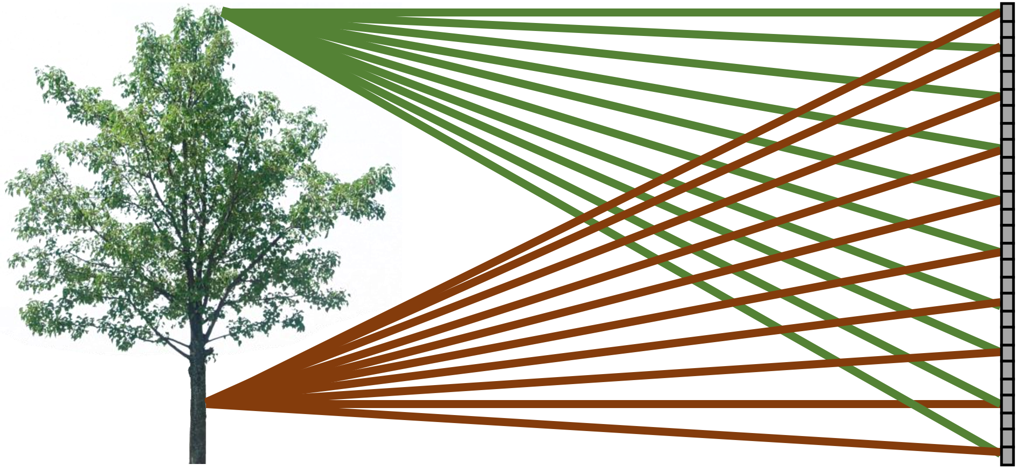

digital sensor
(CCD or
CMOS)

All scene points contribute to all sensor pixels

What does the
image on the
sensor look like?

Bare-sensor imaging

All scene points contribute to all sensor pixels

What can we do to make our image look better?

real-world
object

digital sensor
(CCD or
CMOS)

Let's add something to this scene

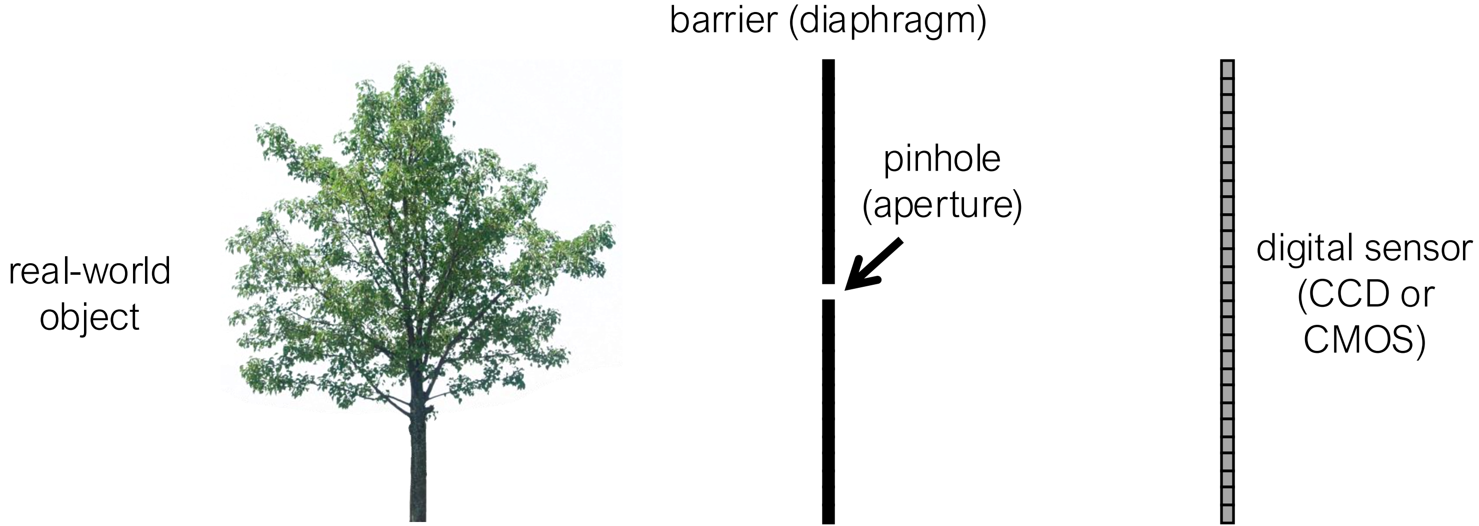

What would an image taken like this look like?

Pinhole imaging

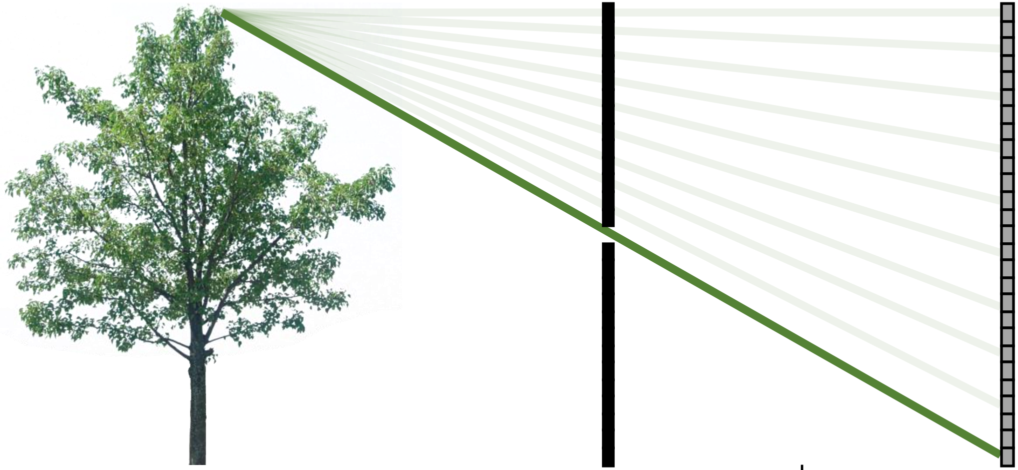

Pinhole imaging

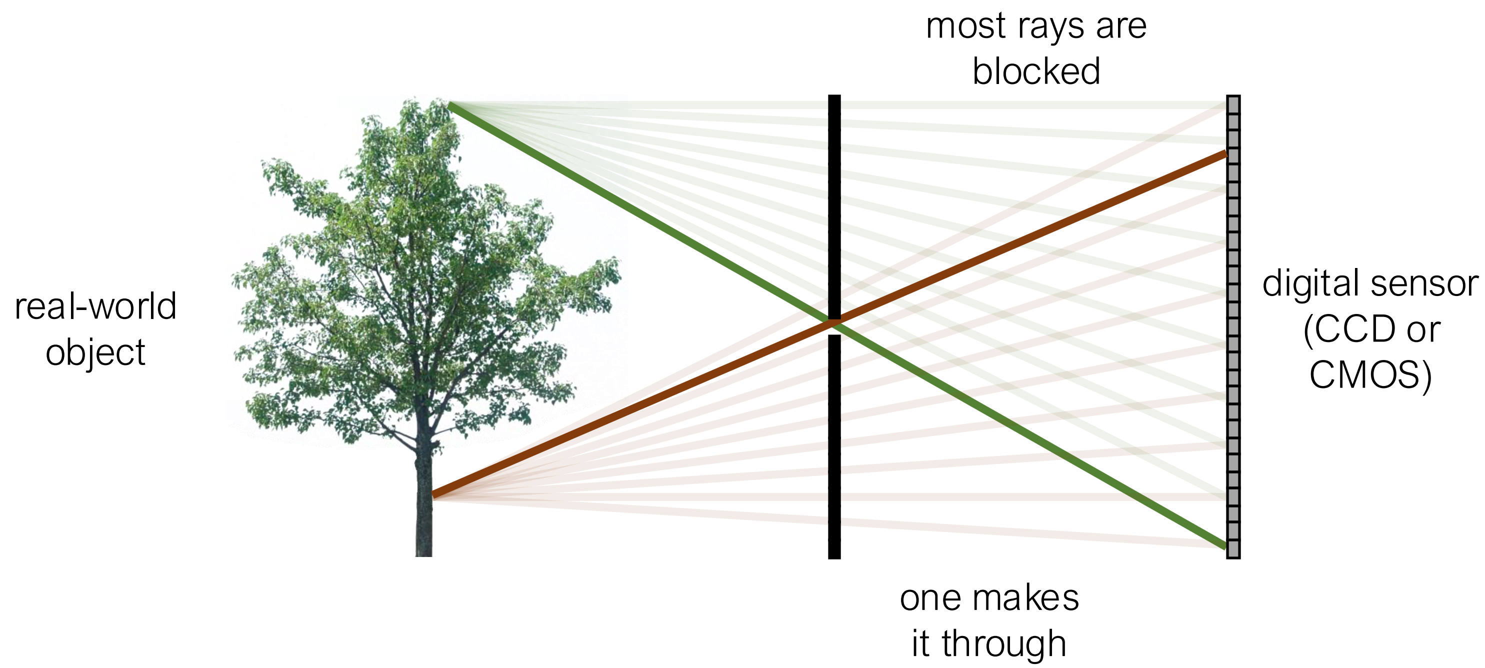

Pinhole imaging

real-world
object

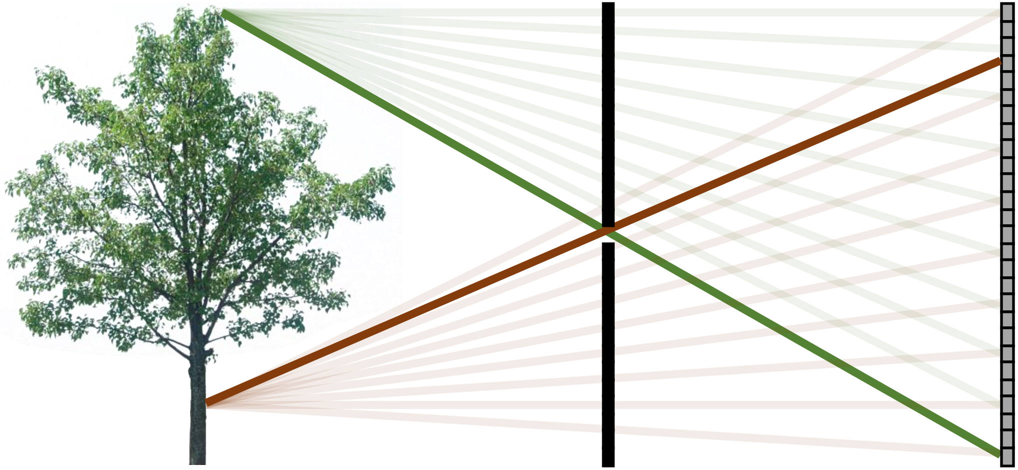

digital sensor
(CCD or
CMOS)

Each scene point contributes to only one sensor pixel

What does the
image on the
sensor look like?

Pinhole imaging

Pinhole camera terms

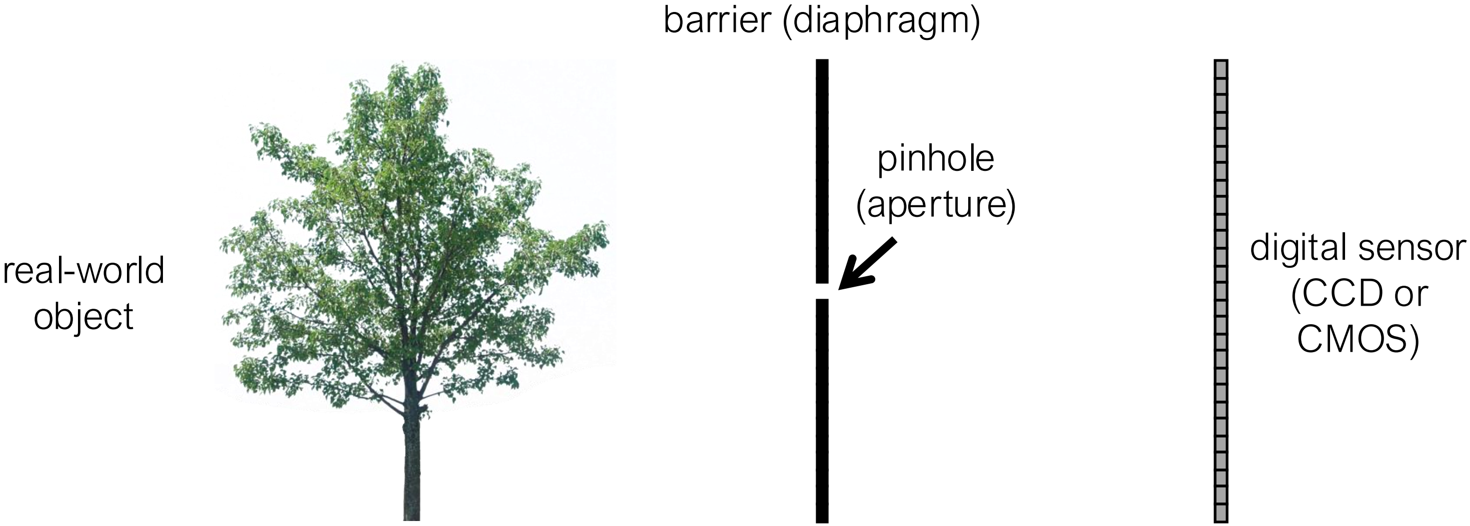

Pinhole camera terms

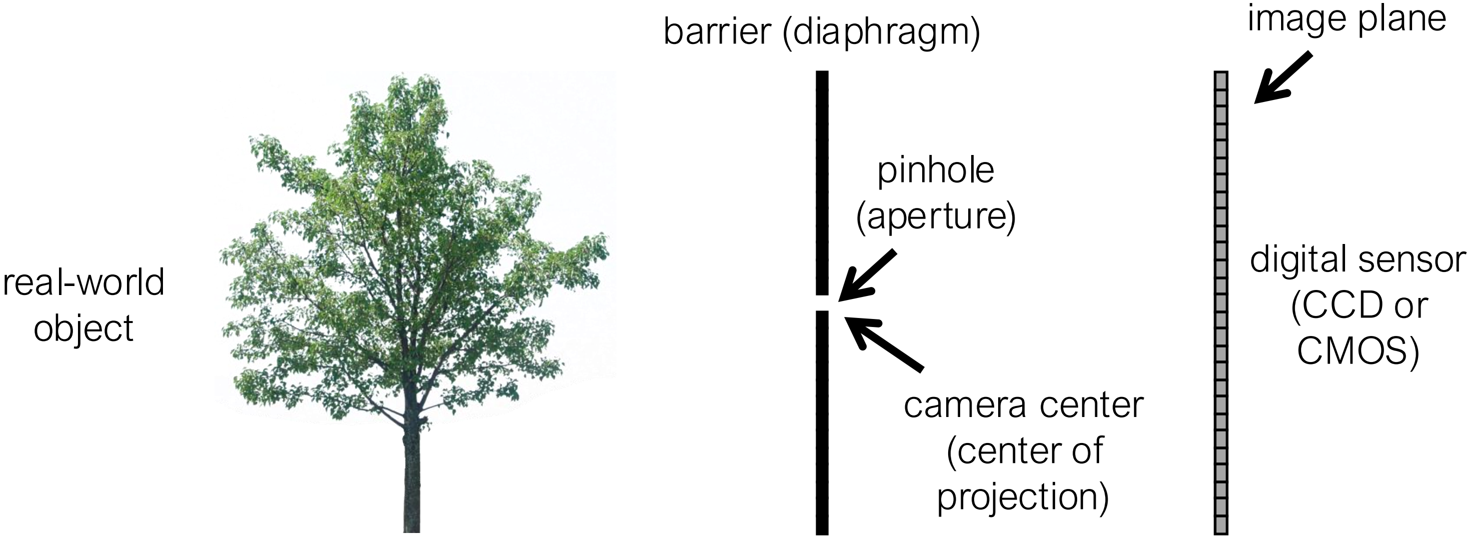

Focal length

real-world object

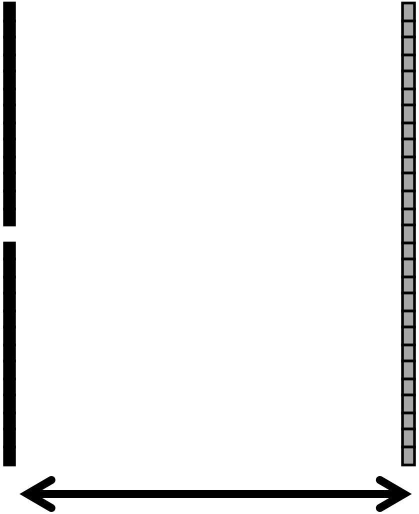

focal length f

Focal length

What happens as we change the focal length?

real-world object

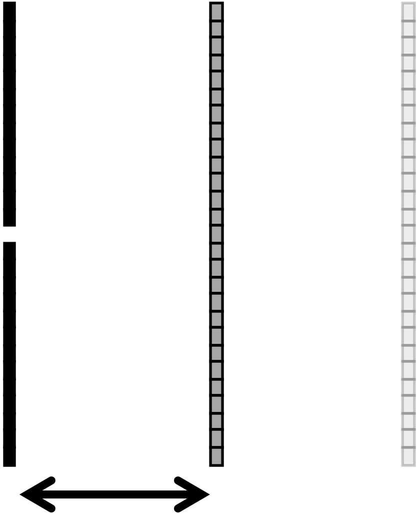

focal length 0.5 f

Focal length

What happens as we change the focal length?

real-world object

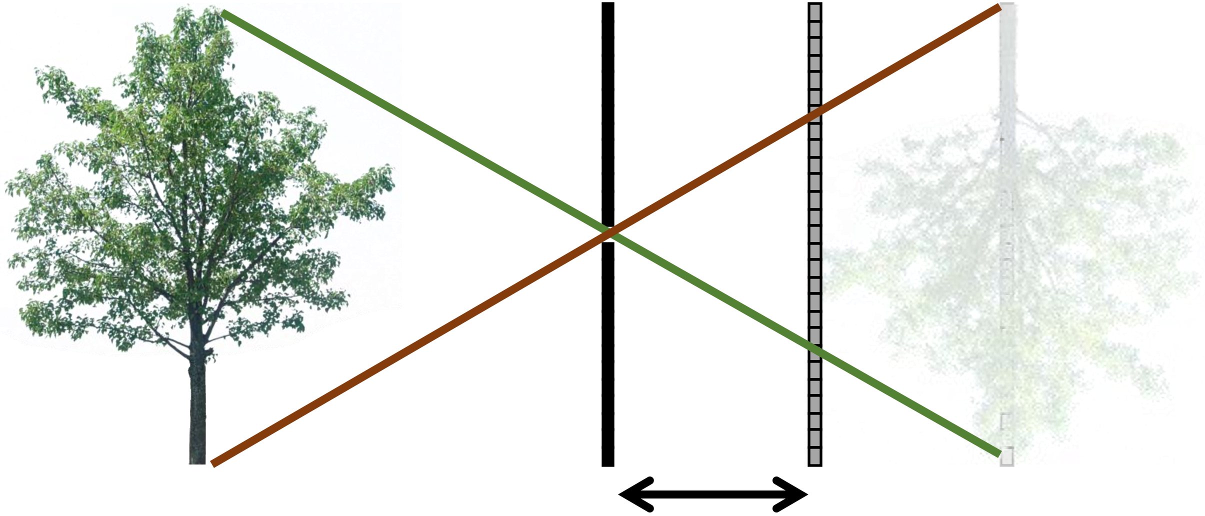

focal length 0.5 f

Focal length

What happens as we change the focal length?

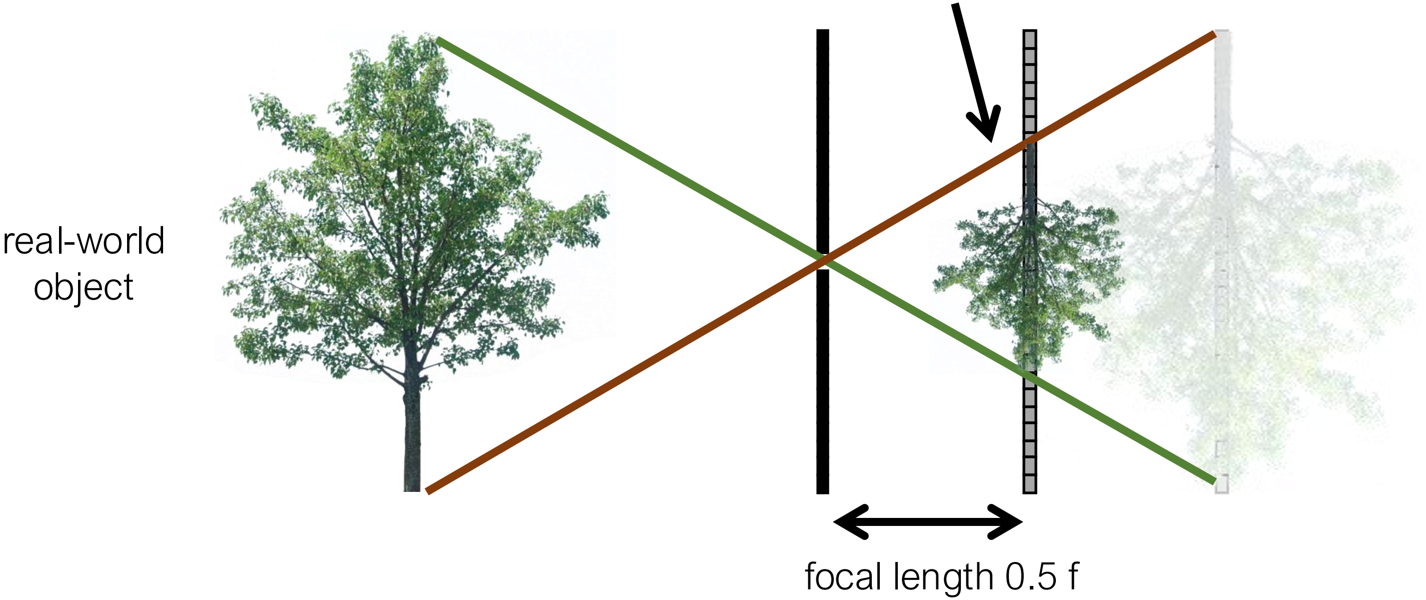

Pinhole size

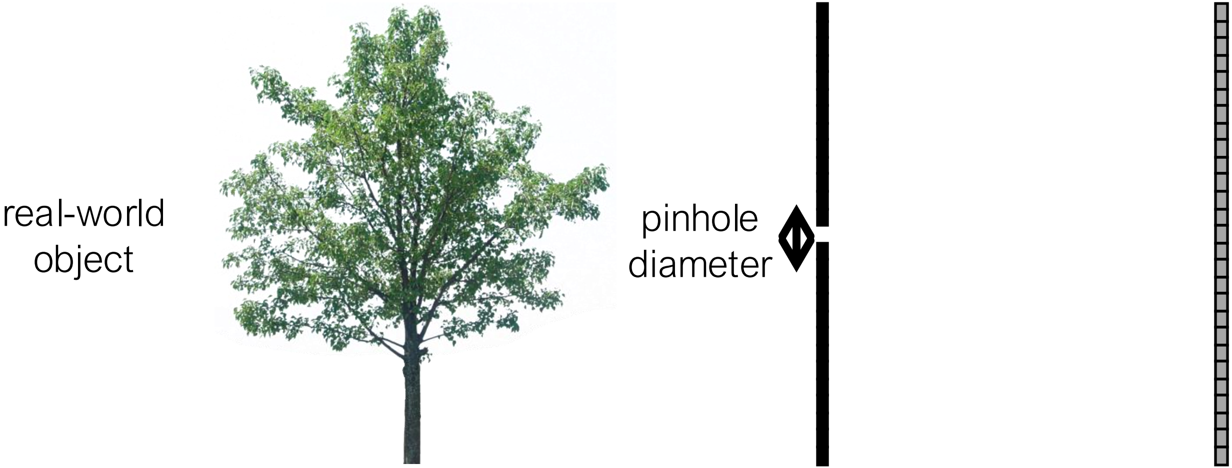

Ideal pinhole has infinitesimally small size
- In practice that is impossible.

Pinhole size

What happens as we change the pinhole diameter?

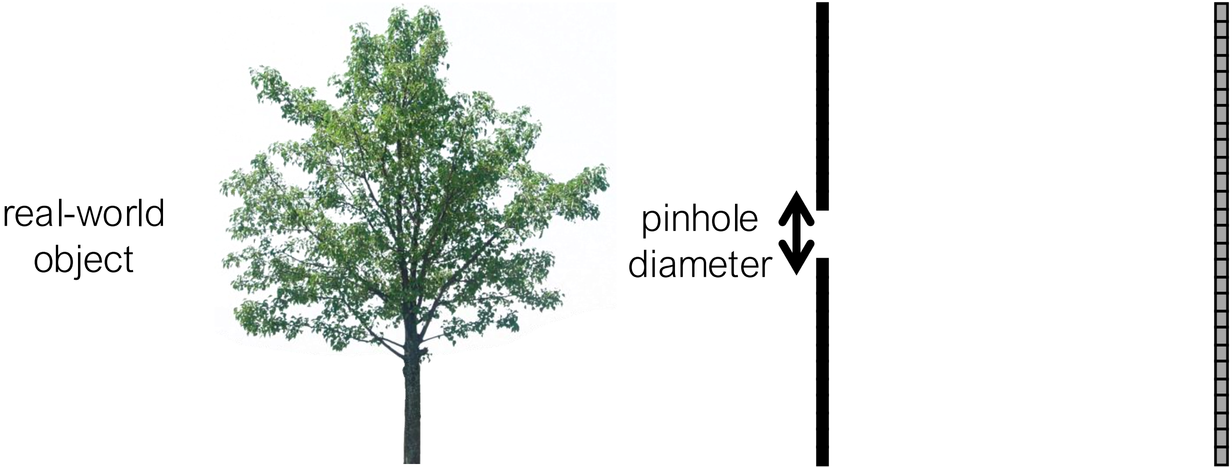

Pinhole size

What happens as we change the pinhole diameter?

real-world object

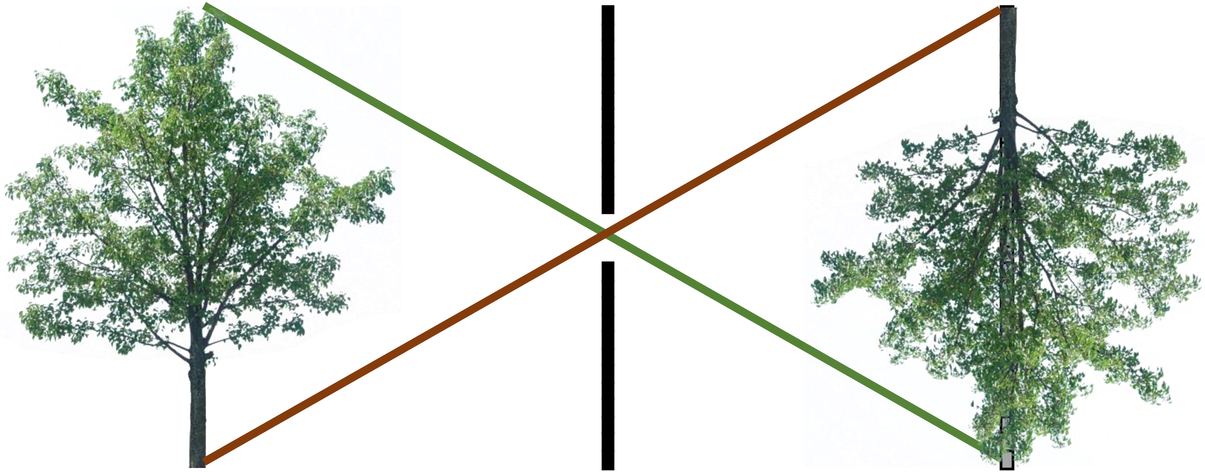

Pinhole size

What happens as we change the pinhole diameter?

real-world object

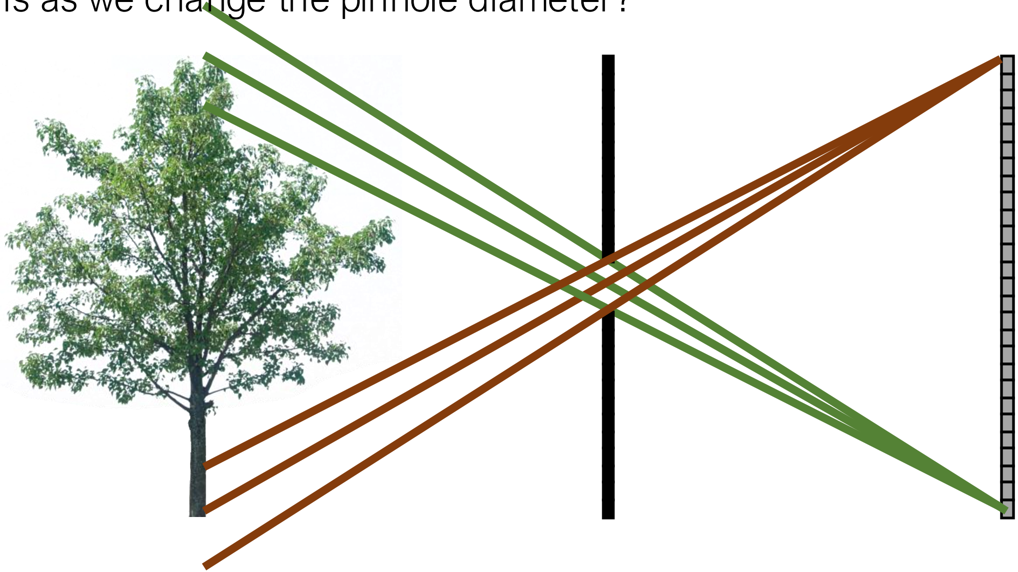

Pinhole size

What happens as we change the pinhole diameter?

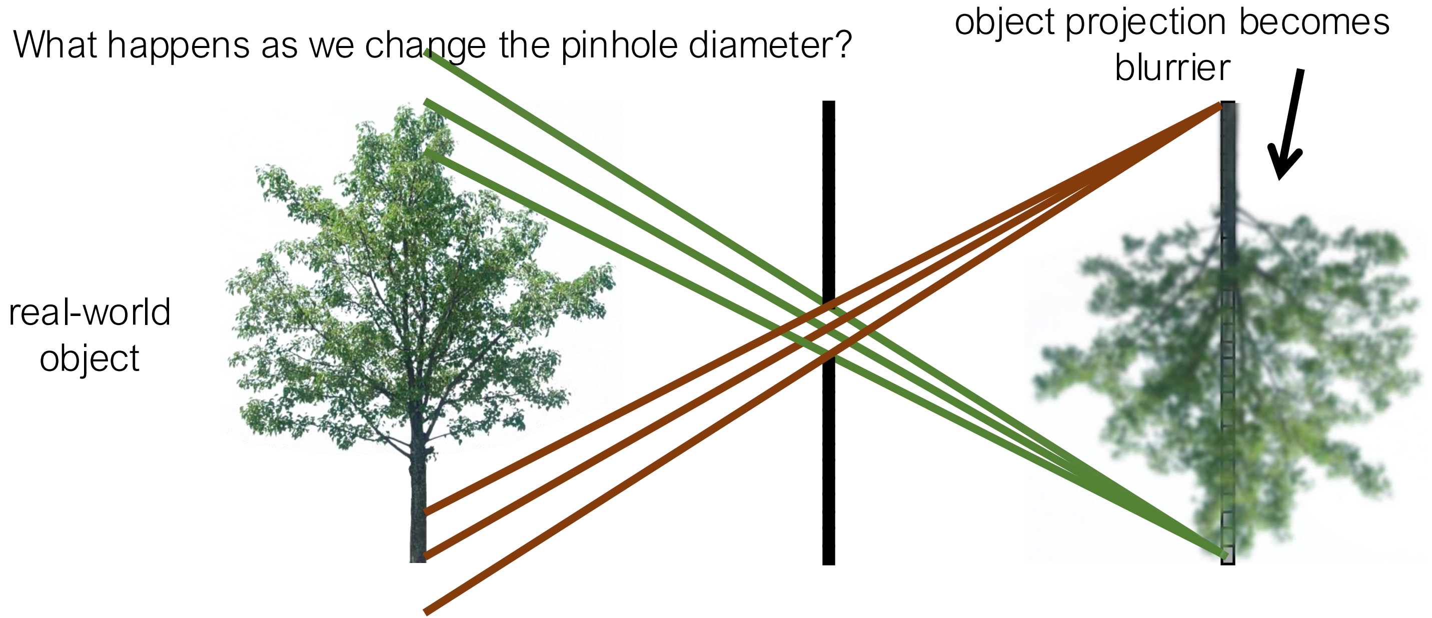

Pinhole size

What happens as we change the pinhole diameter?

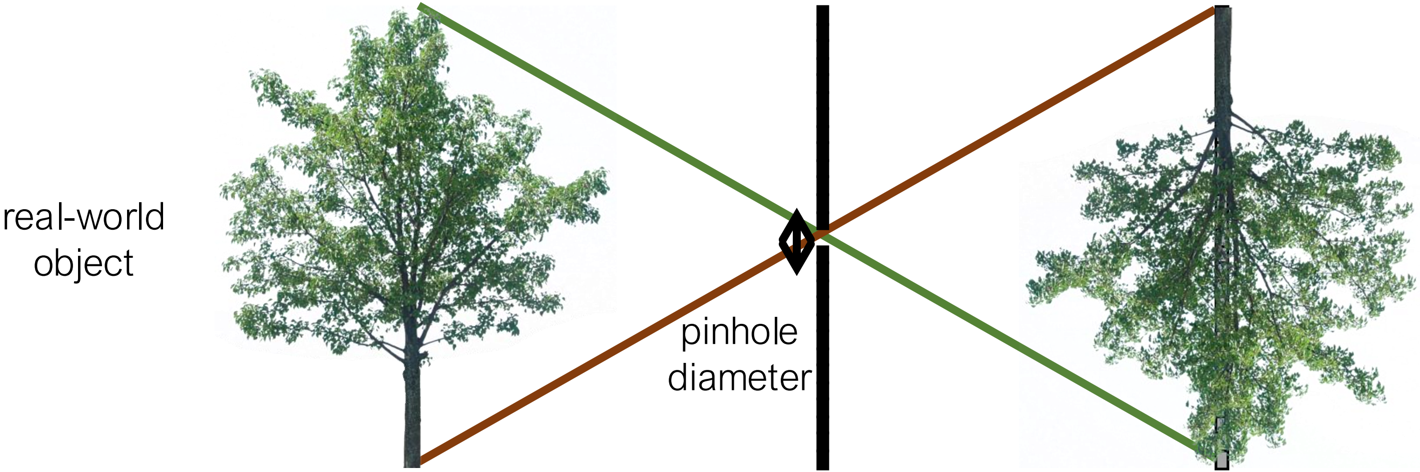

Will the image keep getting sharper the smaller we make the pinhole?

Diffraction limit

A consequence of the wave nature of light

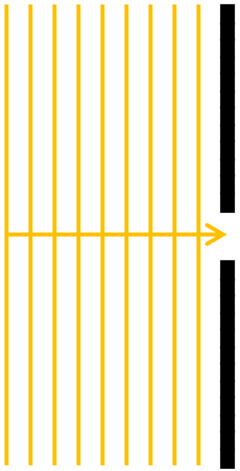

What do geometric optics predict will happen?

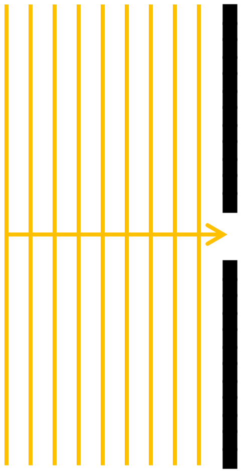

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

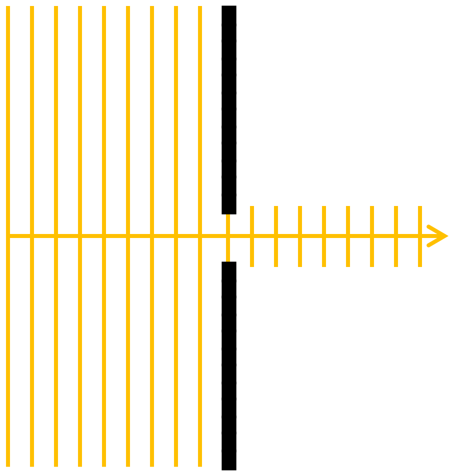

What do geometric optics predict will happen?

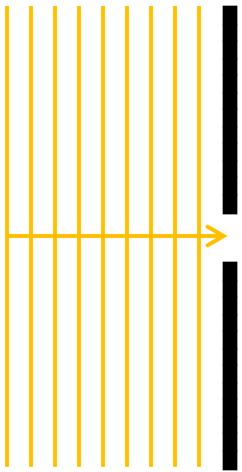

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

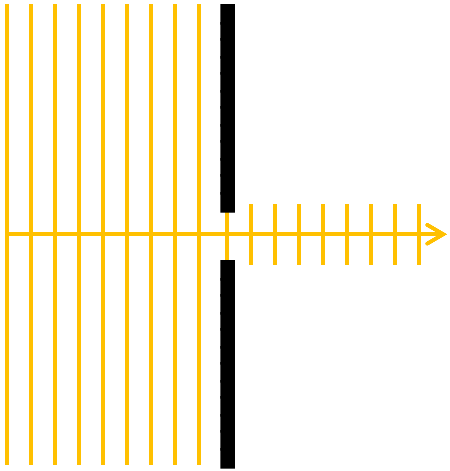

What do geometric optics predict will happen?

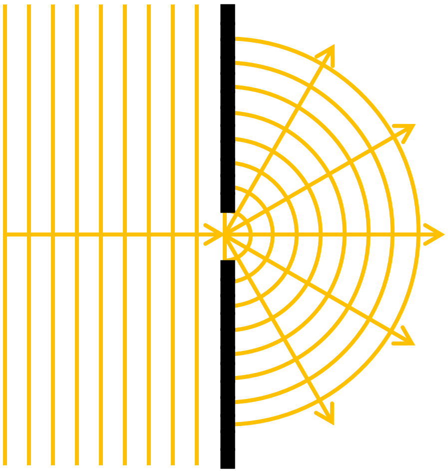

What do wave optics predict will happen?

Diffraction limit

Diffraction pattern = Fourier transform of the pinhole.

- Smaller pinhole means bigger Fourier spectrum.
- Smaller pinhole means more diffraction.

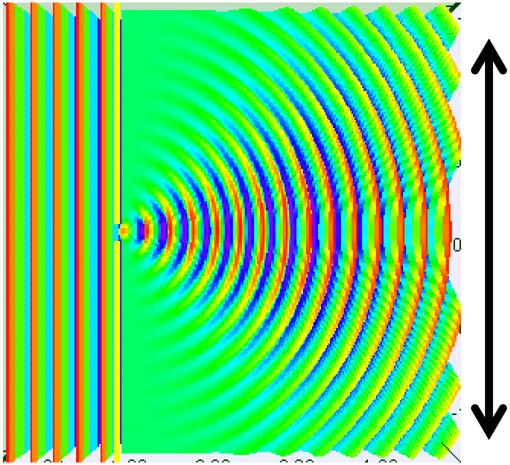

wide diffraction pattern

small pinhole

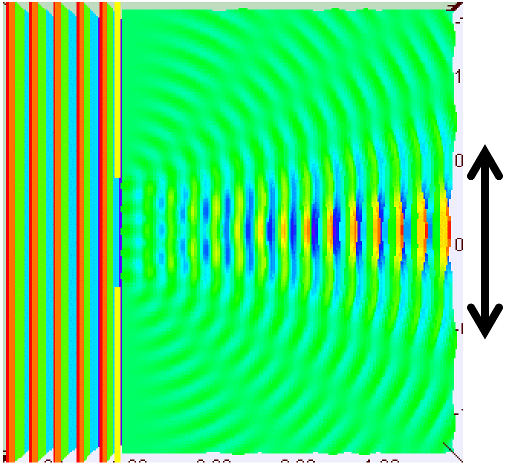

narrow diffraction pattern

large pinhole

What about light efficiency?

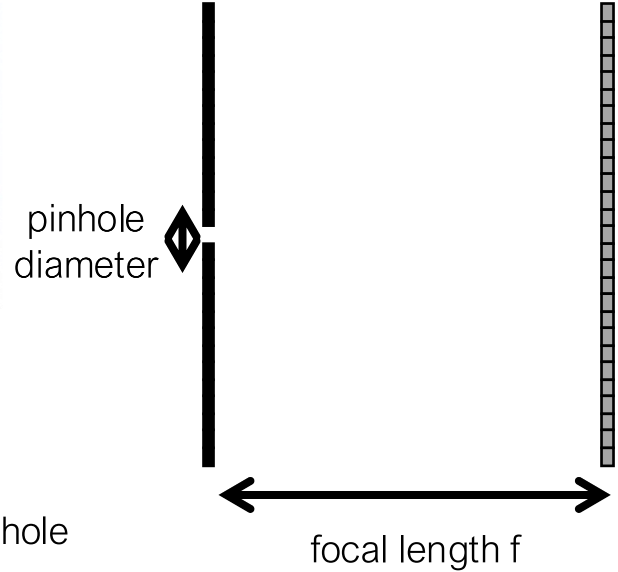

- What is the effect of doubling the pinhole diameter?
- What is the effect of doubling the focal length?

What about light efficiency?

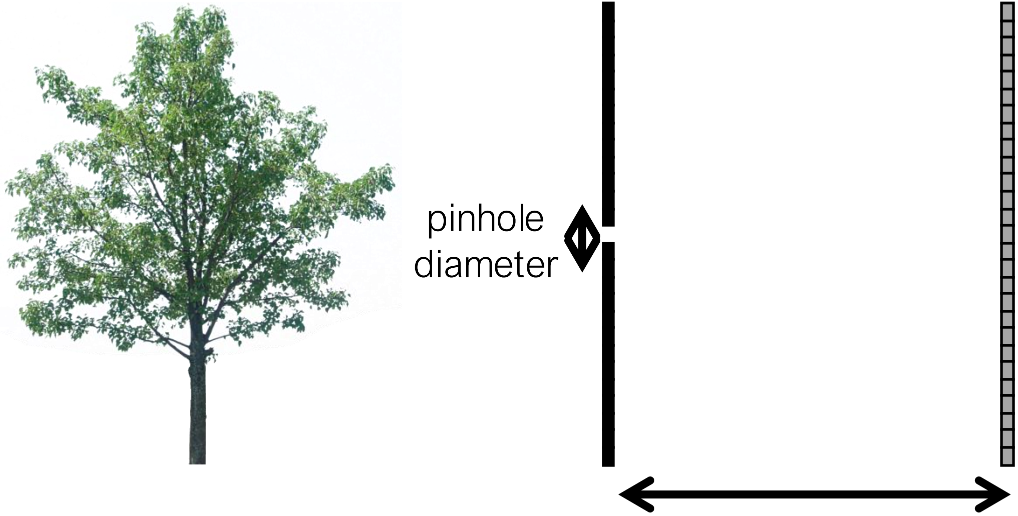

- 2x pinhole diameter → 4x light
- 2x focal length → ¼x light

Pinhole Camera / Camera Obscura

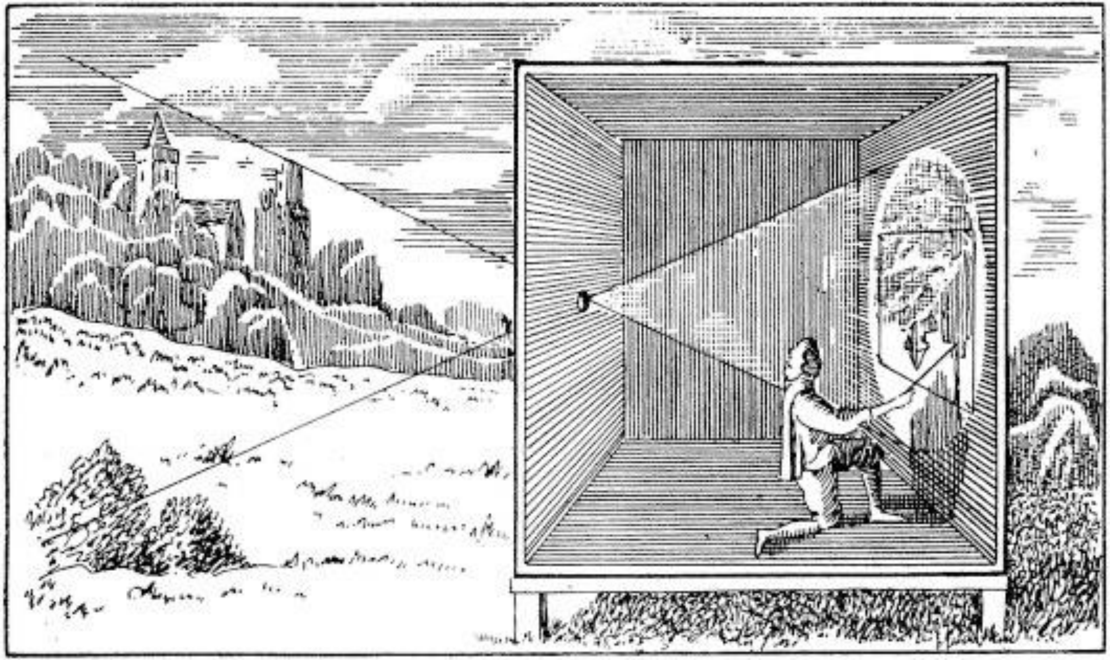

Mo-Ti (Chinese Philosopher) 470-390 BC

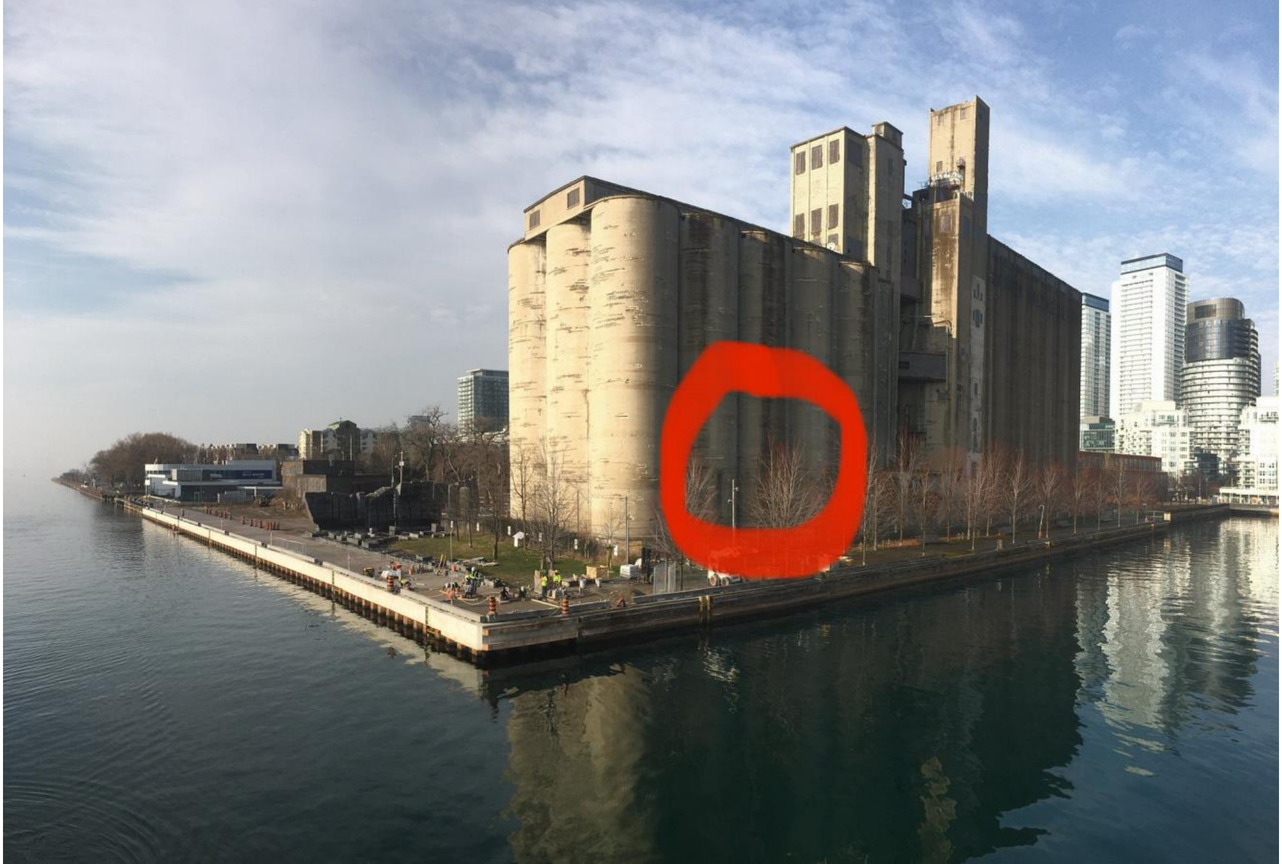

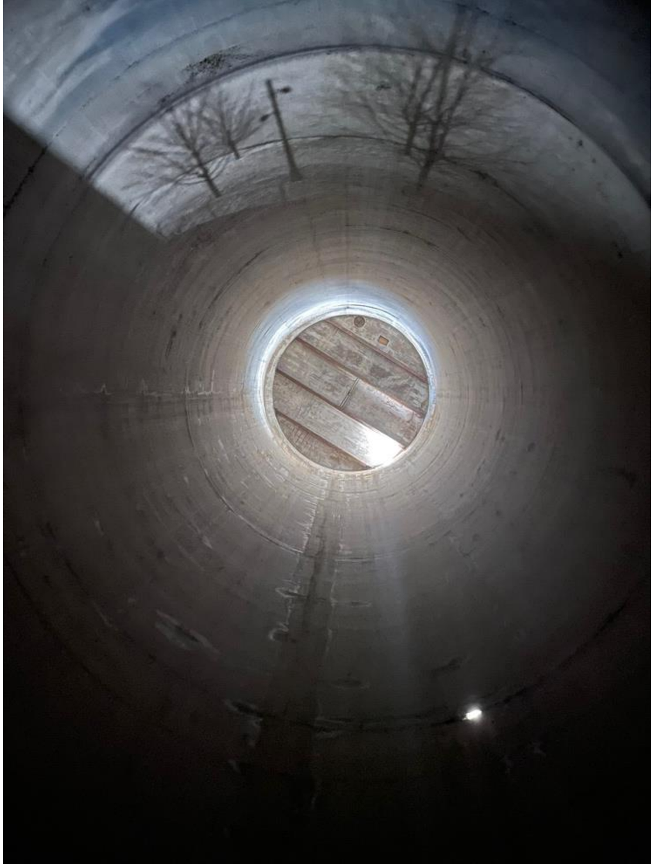

Bryan Bowen
@bryanmbowen

Fun discovery - a small crack in the eastern facade of the Canada Malting Co silos has created a perfect pinhole camera. The result: real time projection of Toronto's waterfront on the silo's interior curved surfaces. An unplugged projection show!

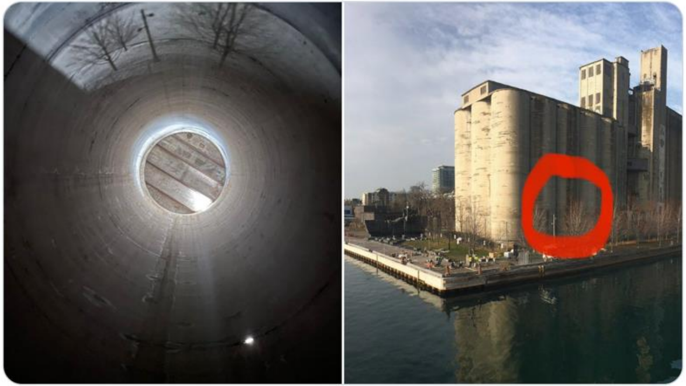

9:37 AM · Jan 27, 2022 · Twitter for iPhone

656 Retweets **70** Quote Tweets **2,836** Likes

Abelardo Morell

Pinhole Camera / Camera Obscura

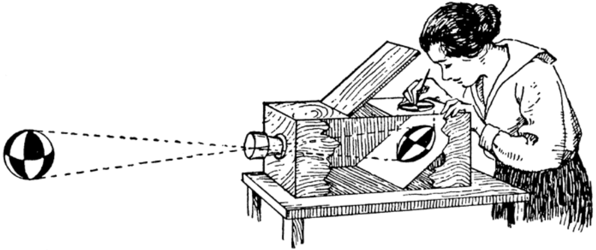

J. Vermeer "The Milkmaid", 1658


同サイズの《牛乳を注ぐ女》が投影される額縁の前に立ち

Digital Photography - Overview

- optics
- aperture
- depth of field
- field of view
- exposure
- noise
- color filter arrays
- image processing pipeline

SLR Camera

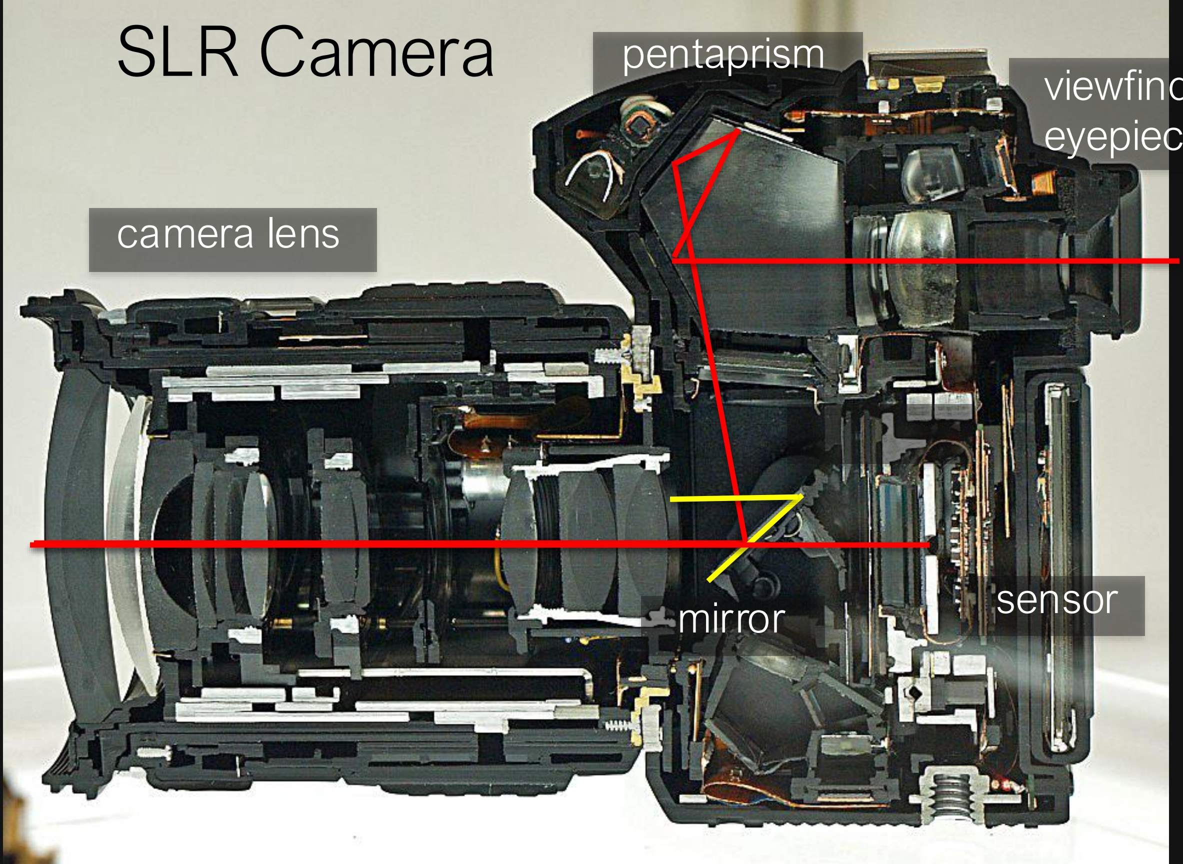
pentaprism
viewfinder eyepiece
camera lens
mirror
sensor

SLR Camera

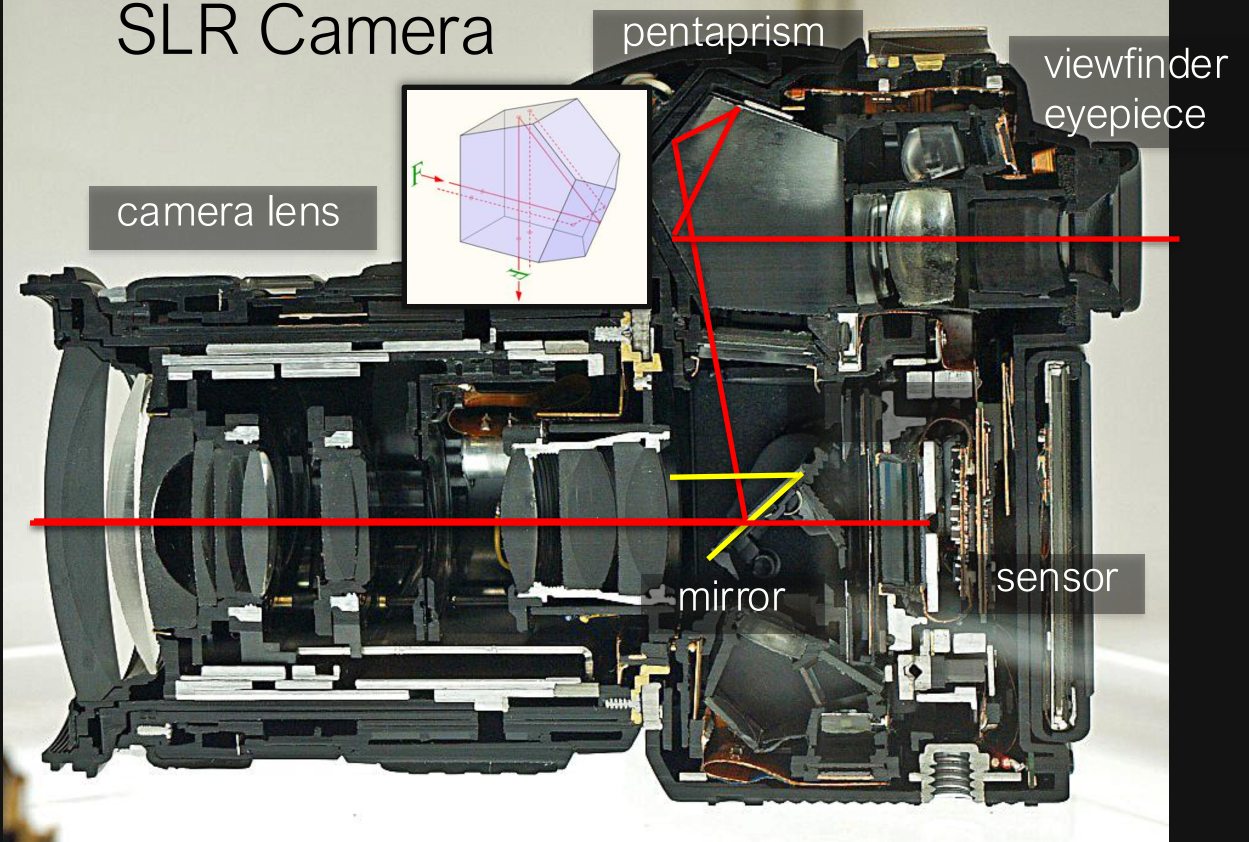
pentaprism
viewfinder eyepiece
camera lens
mirror
sensor

Camera Optics

Niepce “View from the Window at Le Gras”, 1826

1826
8h exp

Daguerrotype

- invented in 1836 by Louis Daguerre
- lenses focus light, better chemicals!

Daguerre “Boulevard du Temple”, 1838

exposure
10-12 mins

Lenses

- focus light
- magnify objects

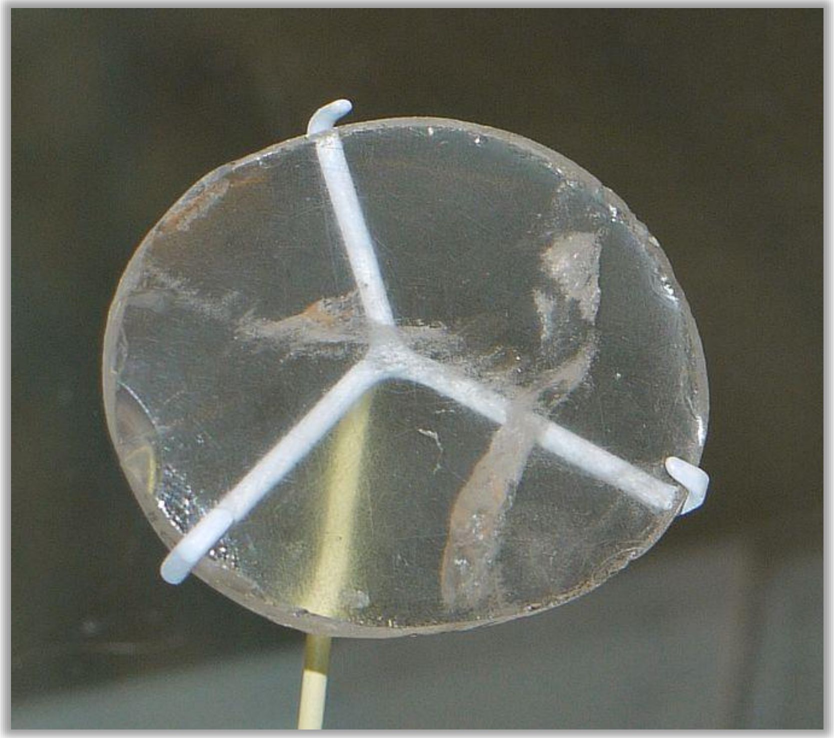

Nimrud lens - 2700 years old

What is a lens?

A piece of glass manufactured to have a specific shape

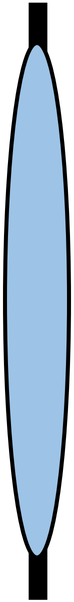

What is a lens?

A piece of glass manufactured to have a specific shape

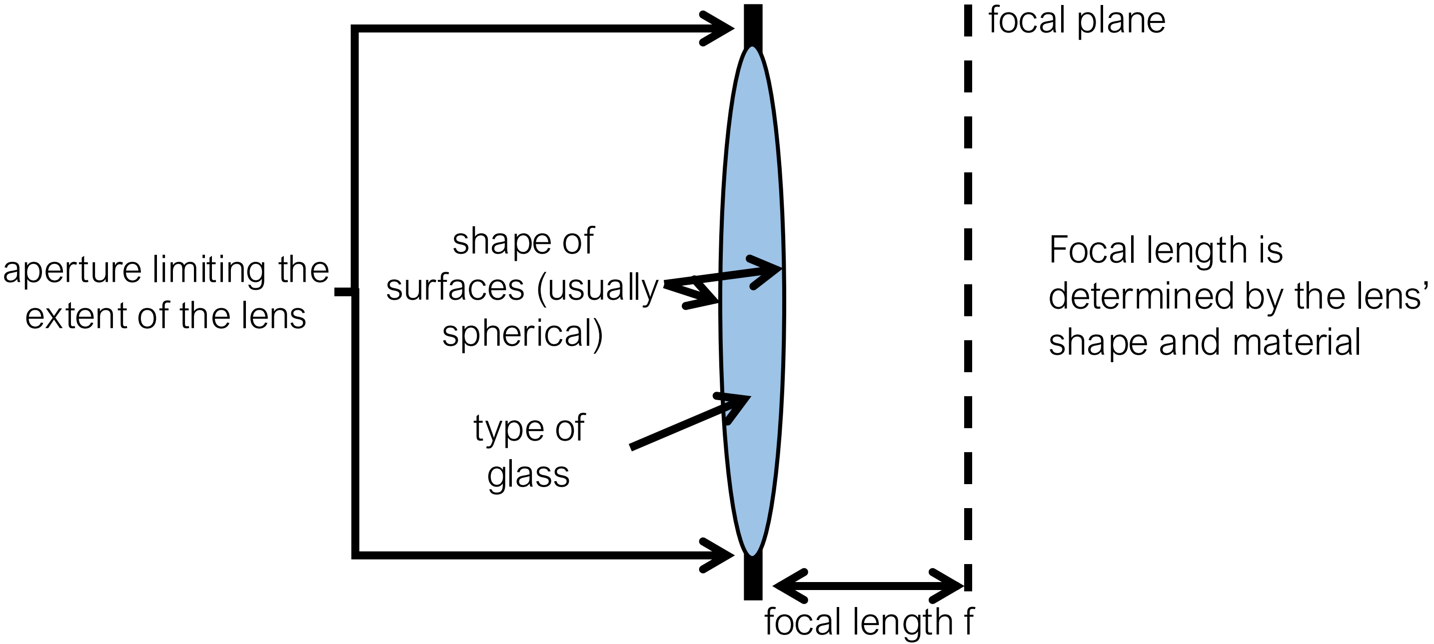

How does a lens work?

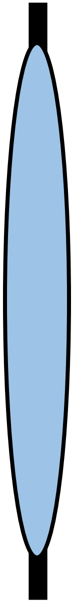

Refraction

Refraction is the bending of rays of light when they move from one material to another

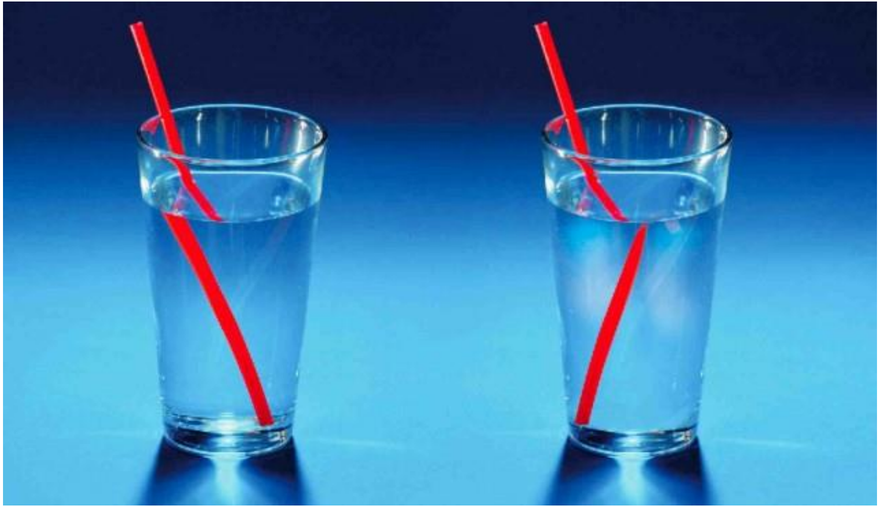

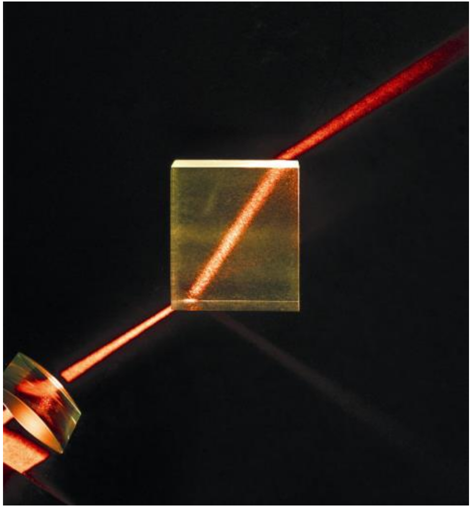

How does a lens work?

Lenses are designed so that their refraction makes light rays bend in a very specific way.

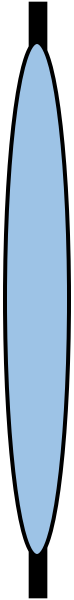

Thin lens model

Simplification of geometric optics for <u>well-designed</u> lenses.

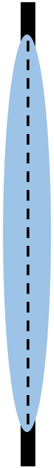

Thin lens model

Simplification of geometric optics for well-designed lenses.

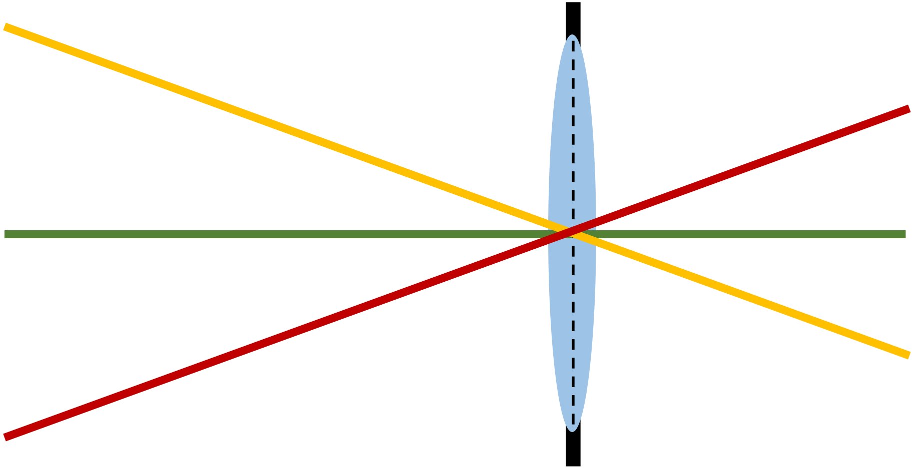

Two assumptions:

1. Rays passing through lens center are unaffected.

Thin lens model

Simplification of geometric optics for well-designed lenses.

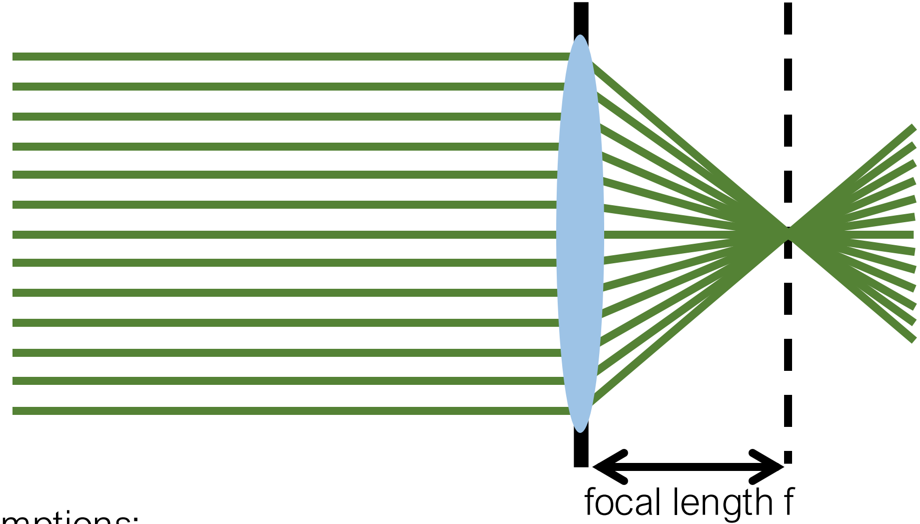

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin lens model

Simplification of geometric optics for well-designed lenses.

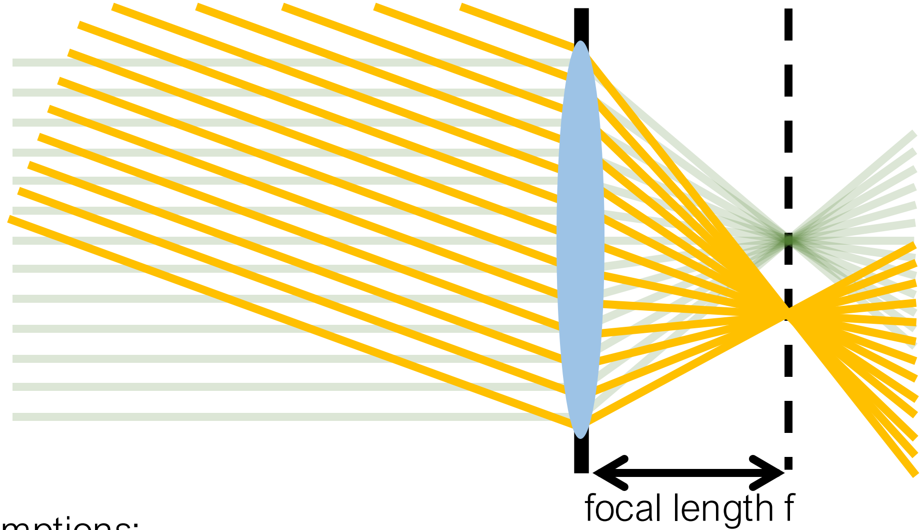

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin lens model

Simplification of geometric optics for well-designed lenses.

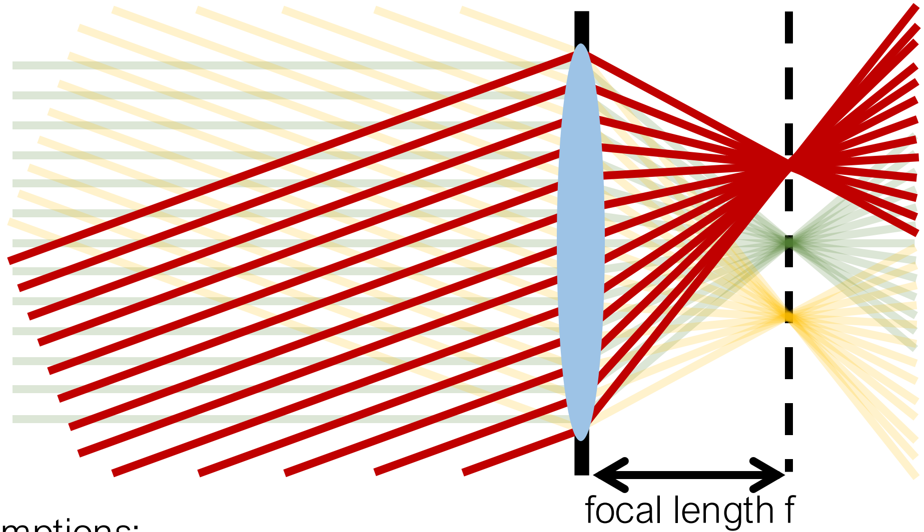

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane

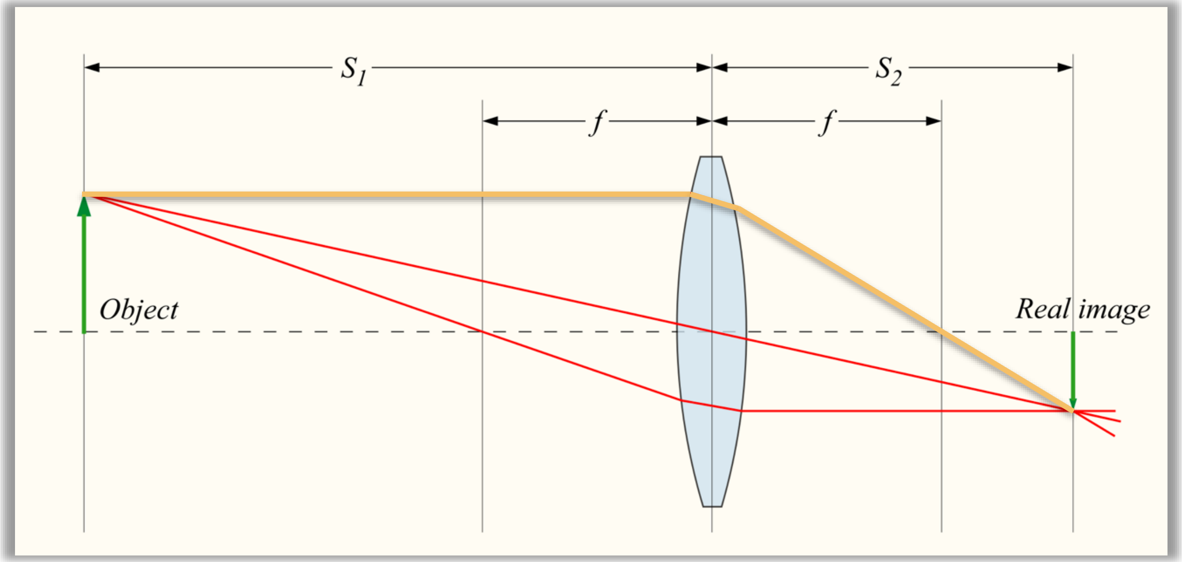

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane
- The chief ray passes straight through the center

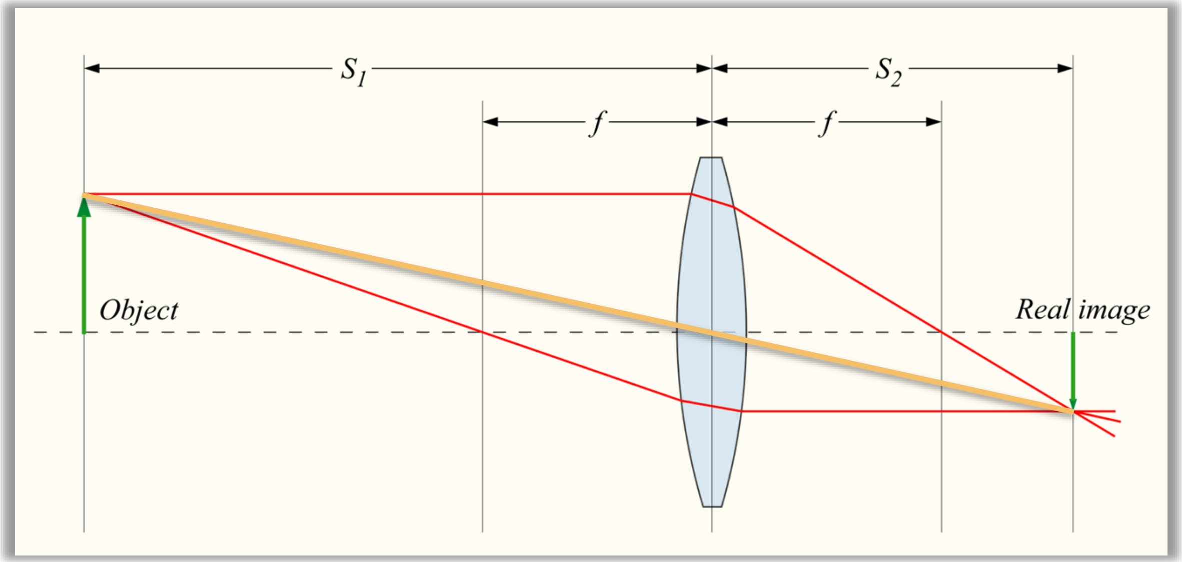

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane
- The chief ray passes straight through the center
- The ray that passes through the near focal plane becomes parallel

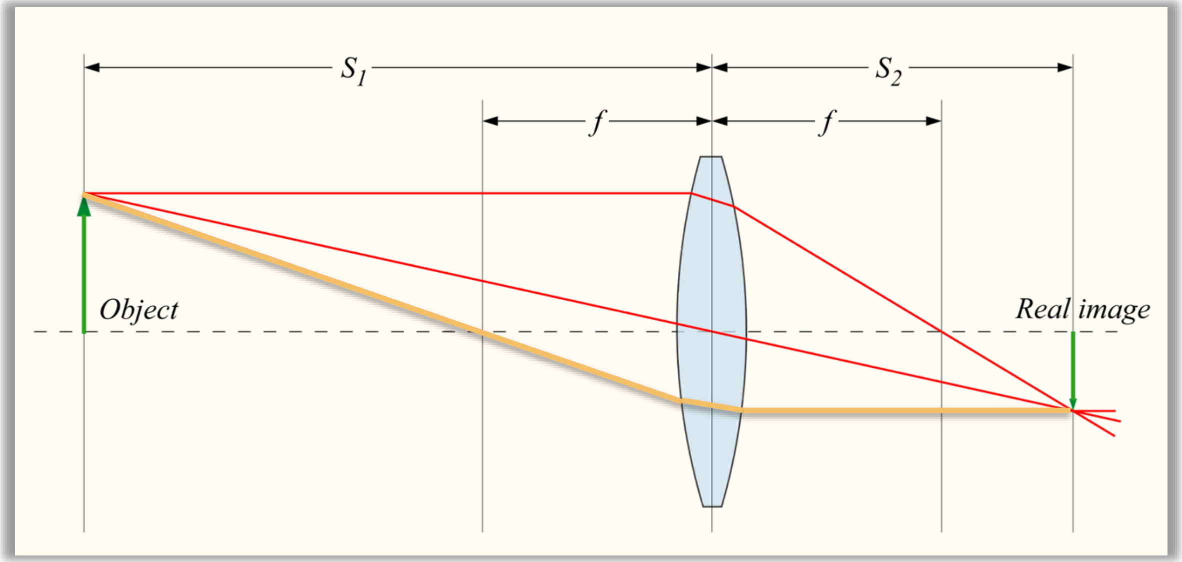

Thin Lens Model

n lens formula $$\frac{1}{f} = \frac{1}{S_1} + \frac{1}{S_2}$$

magnification: $$M = -\frac{S_2}{S_1} = \frac{f}{f - S_1}$$

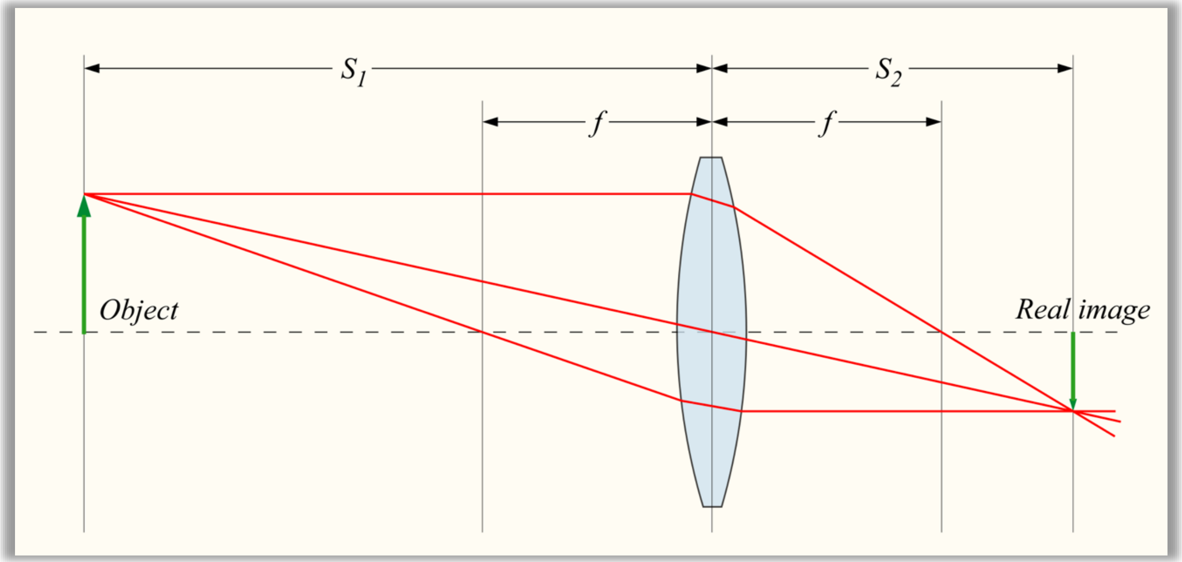

Lenses

S1<f: magnifying glass

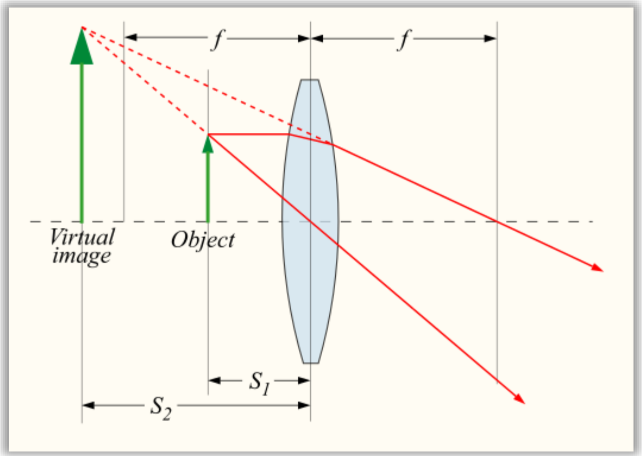

Lenses

S1<f: magnifying glass

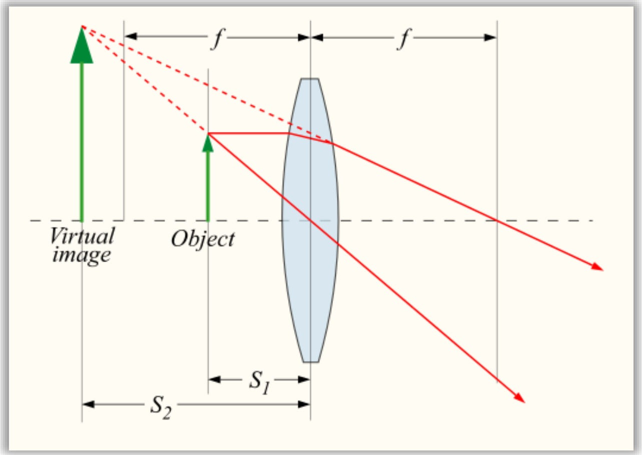

minification

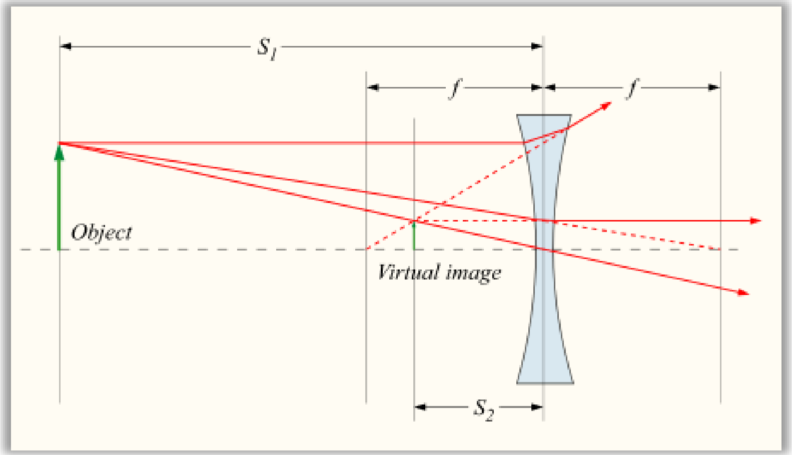

Yes, but…

Thin lenses are a fiction

The thin lens model assumes that the lens has no thickness, but this is rarely true…

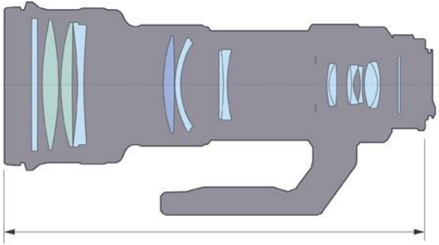

To make real lenses behave like ideal thin lenses, we have to use combinations of multiple lens elements (compound lenses).

Thin lenses are a fiction

The thin lens model assumes that the lens has no thickness, but this is rarely true…

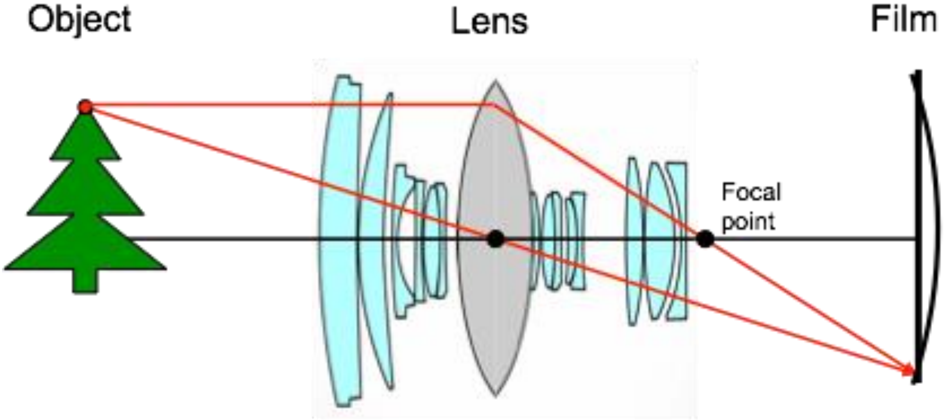

- Even though we have multiple lenses, the entire optical system can be (paraxially) described using a single thin lens of some equivalent focal length and aperture number.

To make real lenses behave like ideal thin lenses, we have to use combinations of multiple lens elements (compound lenses).

Lenses - Aberrations

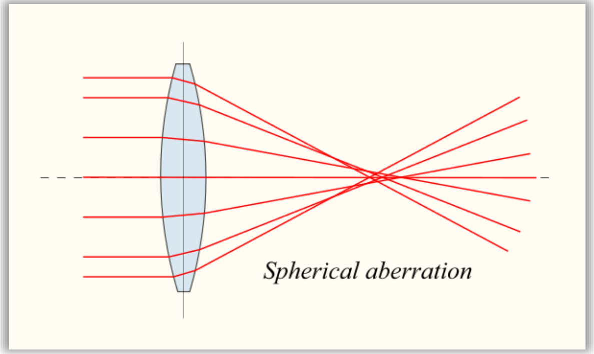

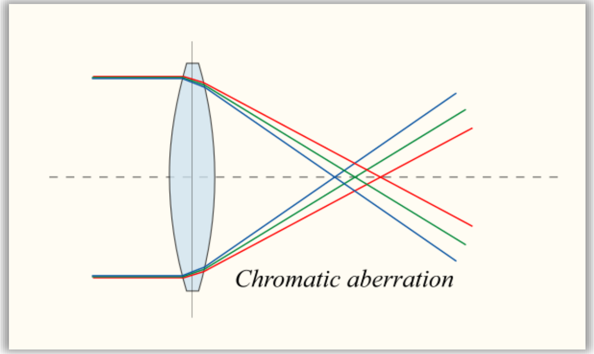

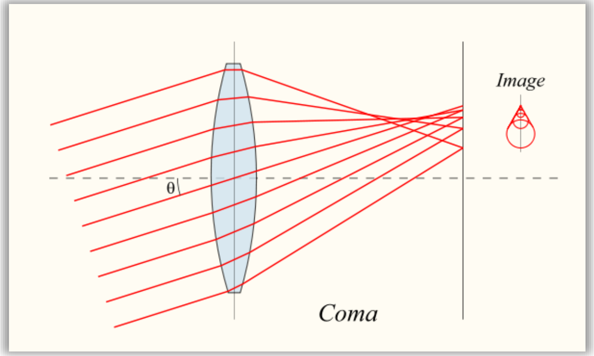

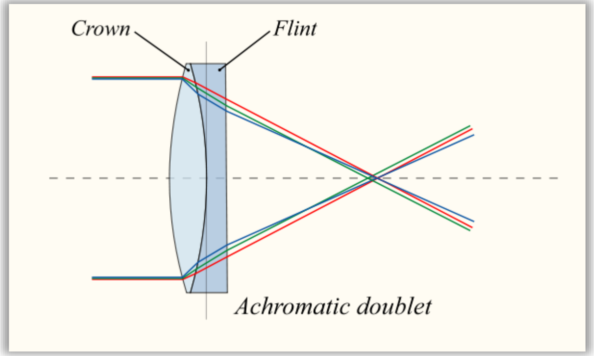

Lenses - Aberrations

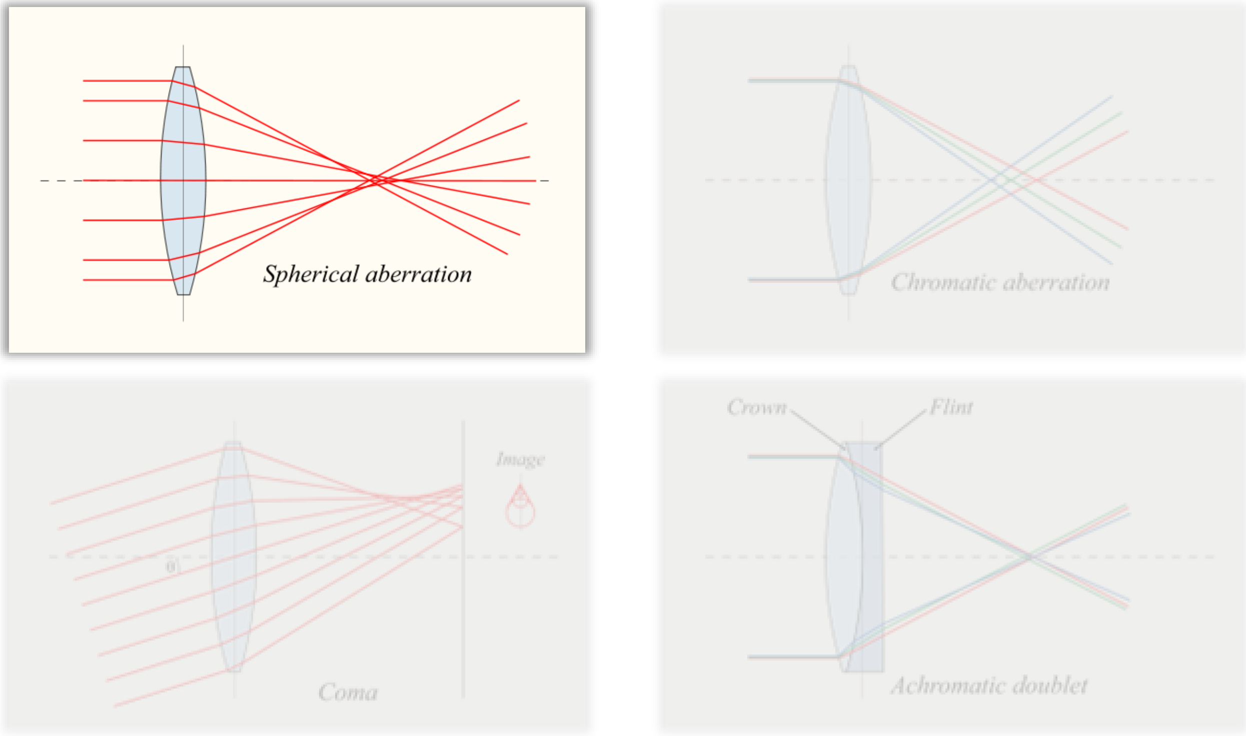

Refraction at interfaces of complicated shapes

What shape should an interface have to make parallel rays converge to a point?

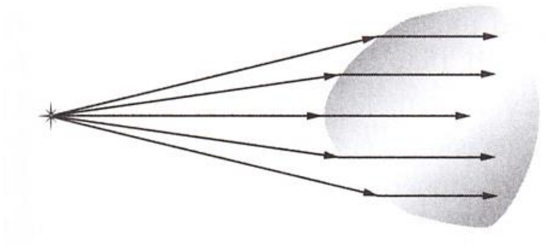

Refraction at interfaces of complicated shapes

What shape should an interface have to make parallel rays converge to a point?

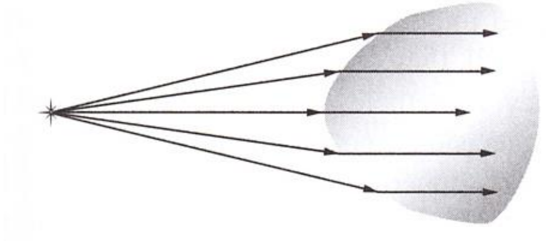

Single hyperbolic interface:
point to parallel rays

Double hyperbolic interface:
point to point rays

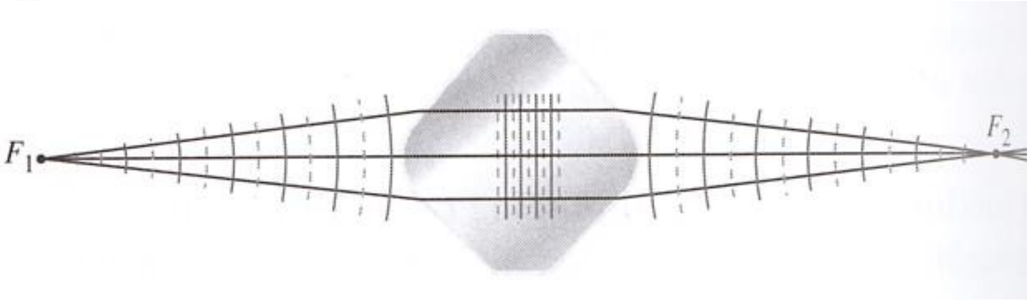

Therefore, lenses should also have hyperbolic shapes.

Spherical lenses

In practice, lenses are often made to have spherical interfaces for ease of fabrication.
- Two roughly fitting curved surfaces ground together will eventually become spherical.

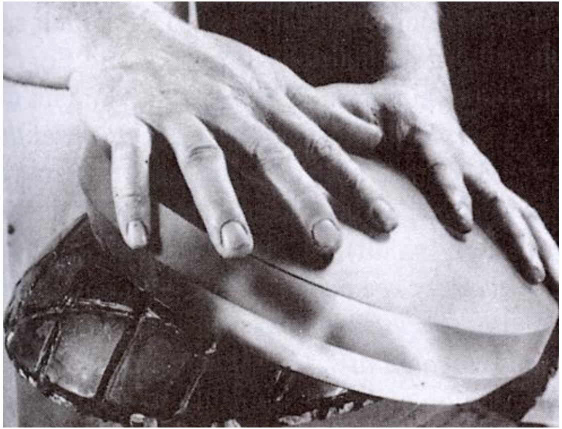

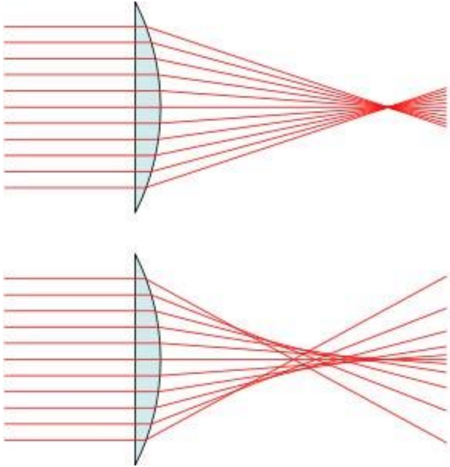

Spherical lenses don't bring parallel rays to a point.
- This is called spherical aberration.
- Approximately axial (i.e., paraxial) rays behave better.

Aberrations

Deviations from ideal thin lens behavior (e.g., imperfect focus).

- Example: spherical aberration.

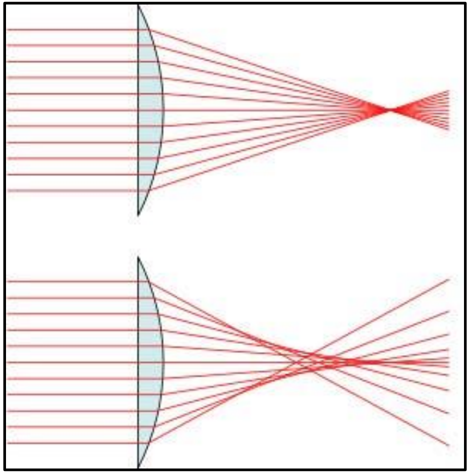

Lenses - Aberrations

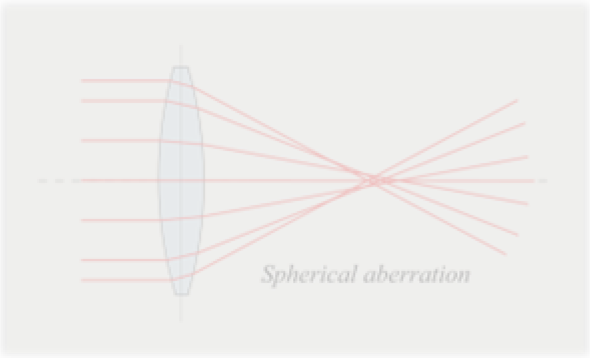

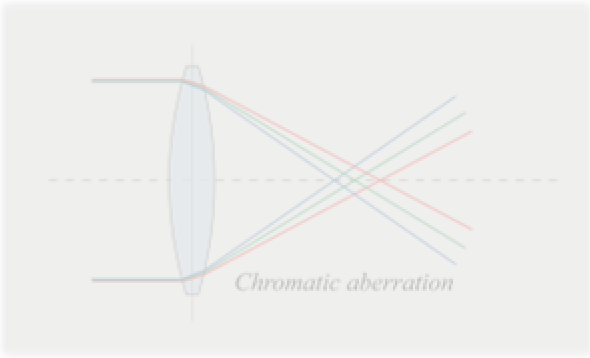

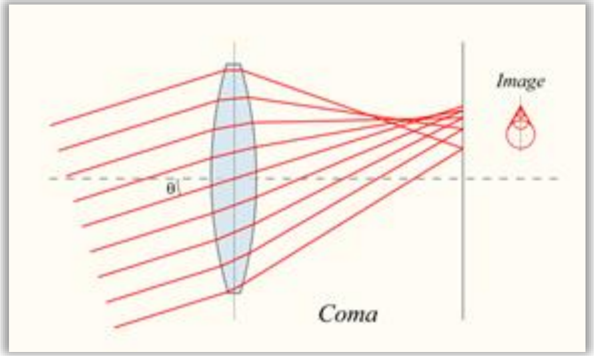

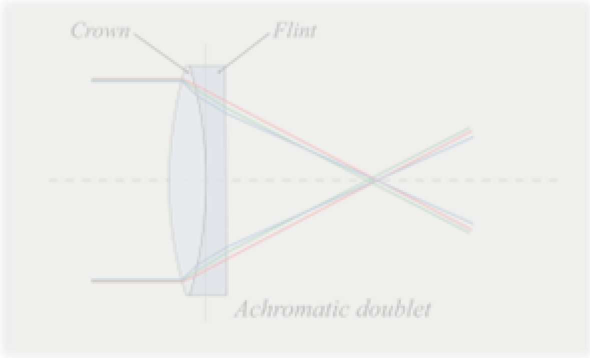

Oblique aberrations

These appear only as we move further from the center of the field of view.

- Contrast with spherical and chromatic, which appear everywhere.
- Many other examples (astigmatism, field curvature, etc.).

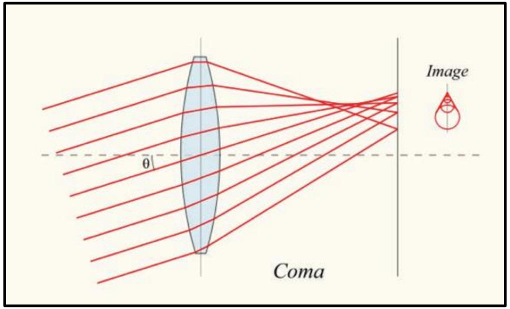

Coma

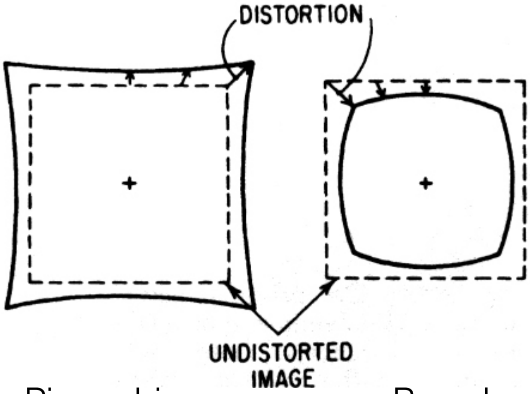

Pincushion Barrel

Distortion example

Lenses - Aberrations

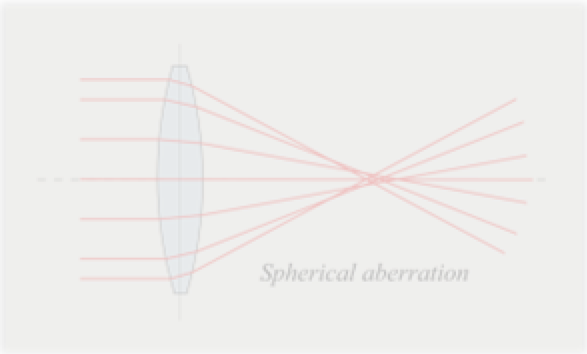

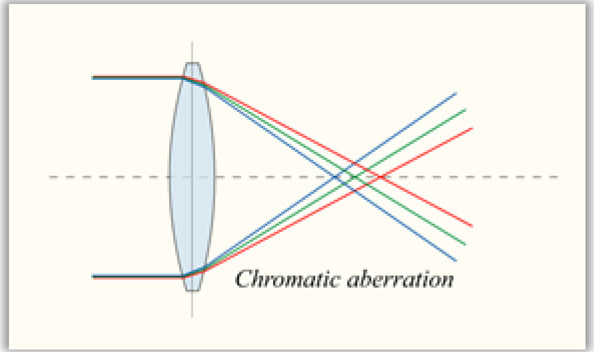

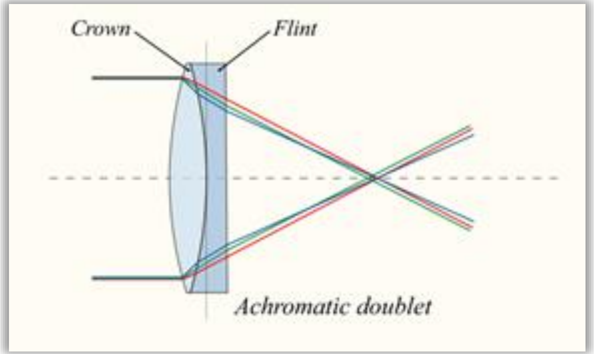

Aberrations

Deviations from ideal thin lens behavior (e.g., imperfect focus).

- Example: chromatic aberration.

focal length shifts with wavelength

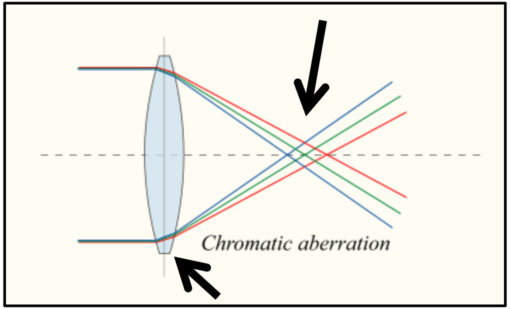

glass has dispersion (refractive index changes with wavelength)

one lens cancels out dispersion of other

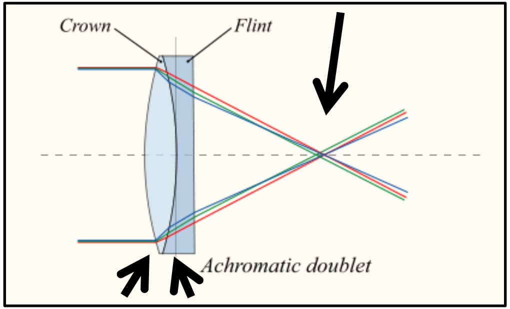

glasses of different refractive index

Using a doublet (two-element compound lens), we can reduce chromatic aberration.

Chromatic aberration examples

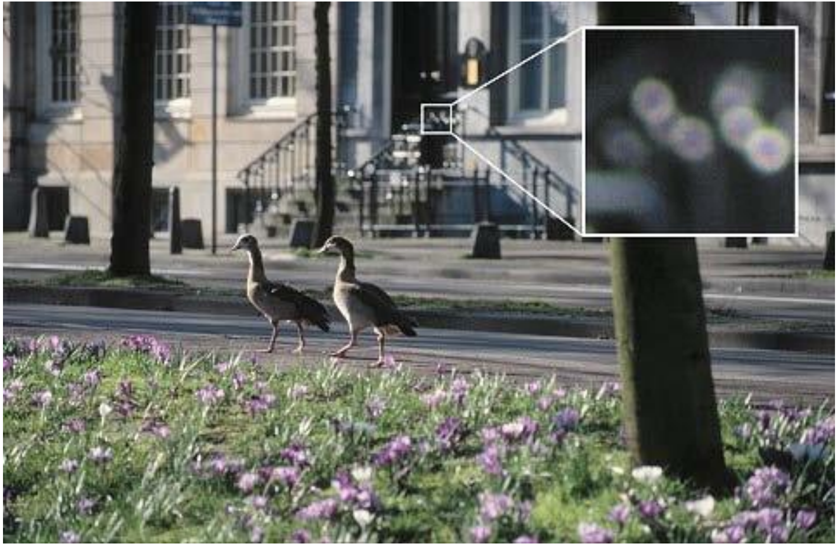

Field of View

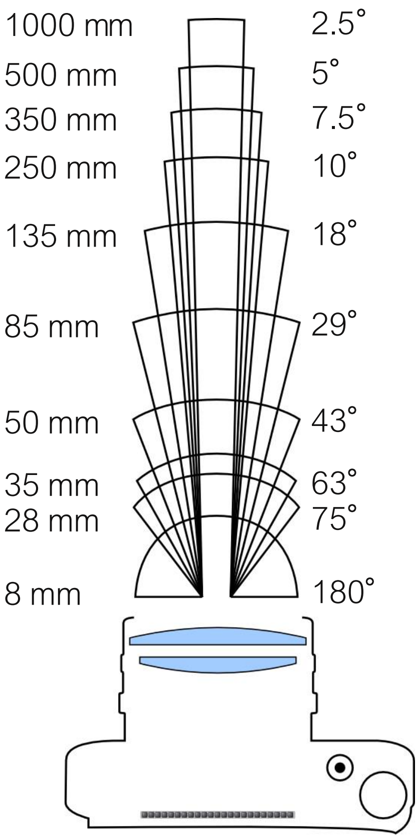

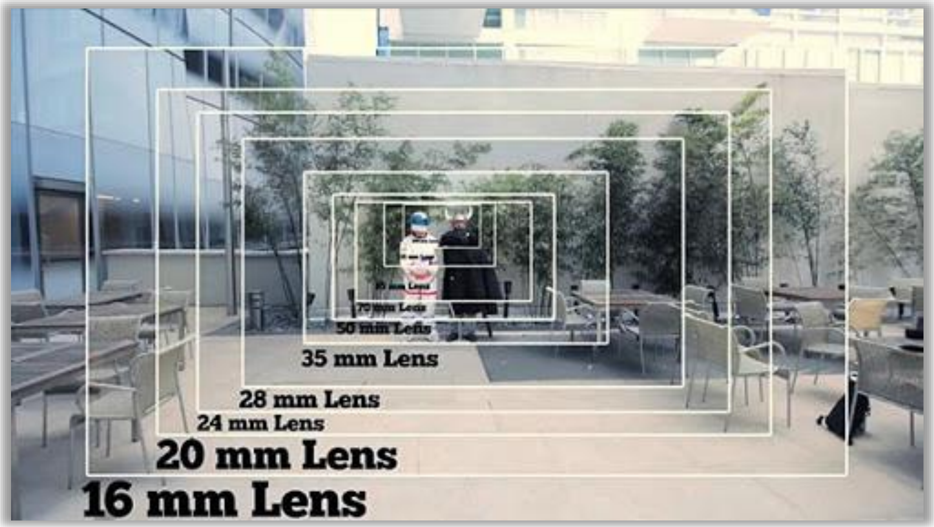

Andrew McWilliams

Field of View

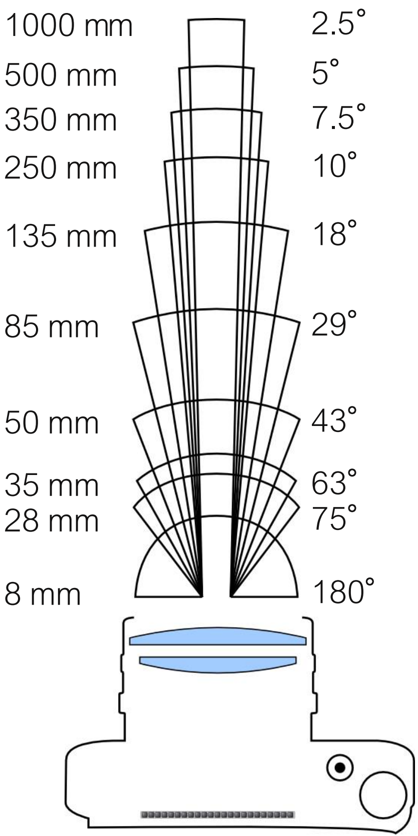

Hubble – what's the focal length?

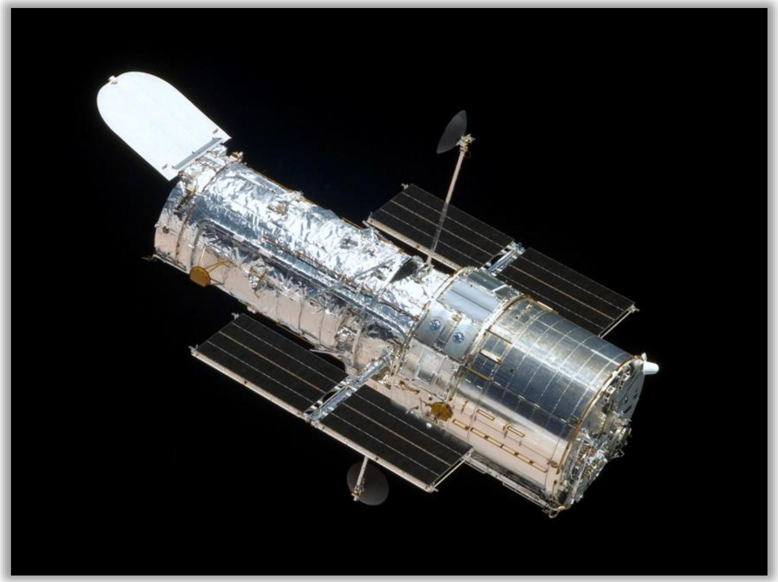

Field of View

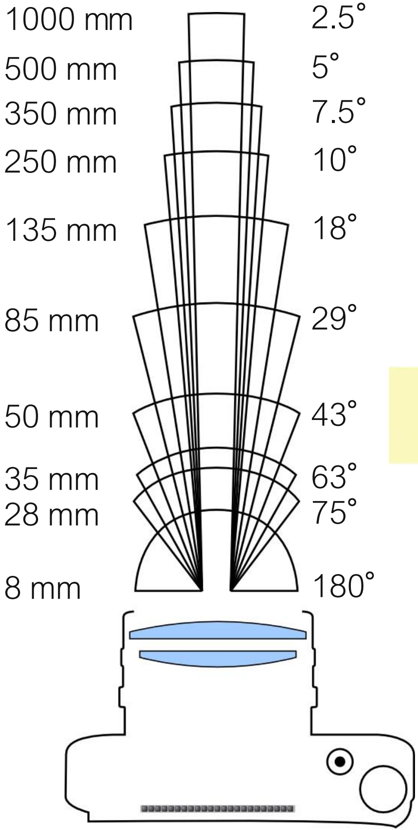

Hubble – what's the focal length?

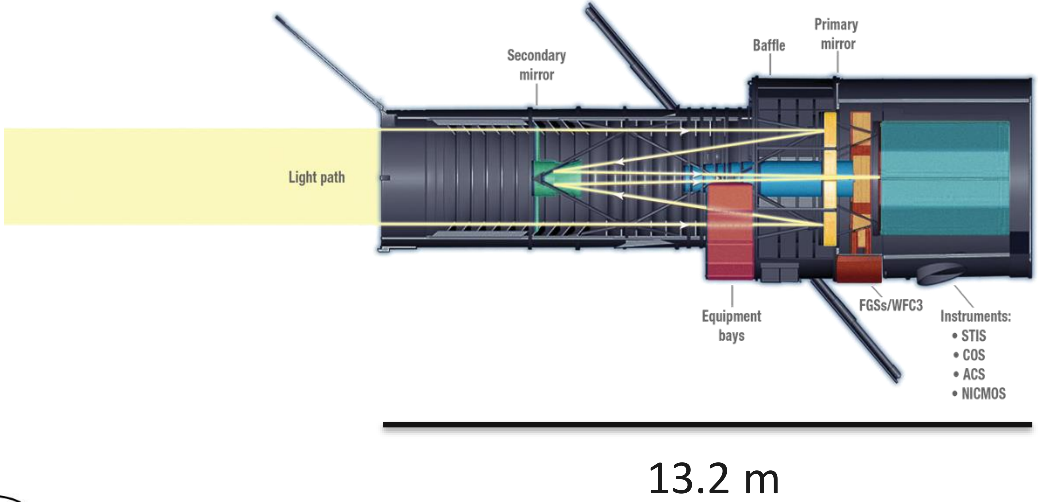

A costly aberration

Hubble telescope originally suffered from severe spherical aberration.

- COSTAR mission inserted optics to correct the aberration.

Aperture

- aperture controls amount of light

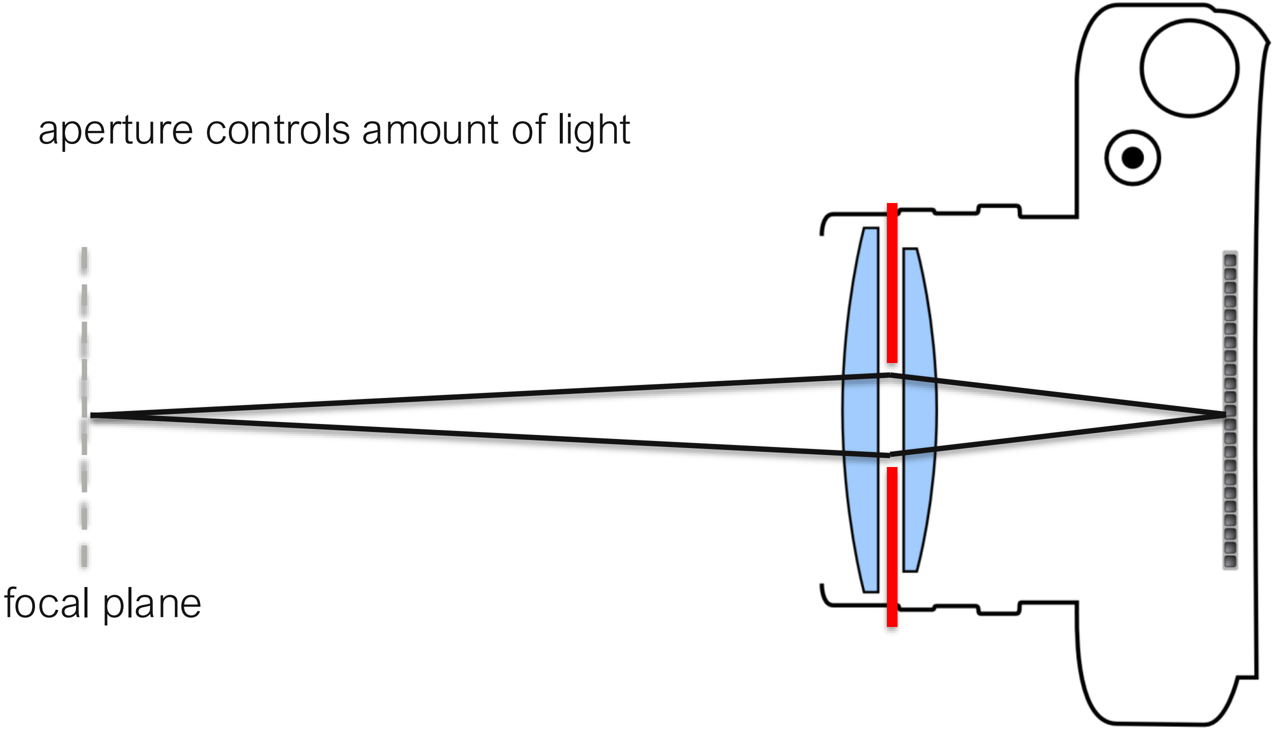

Aperture

- aperture controls amount of light

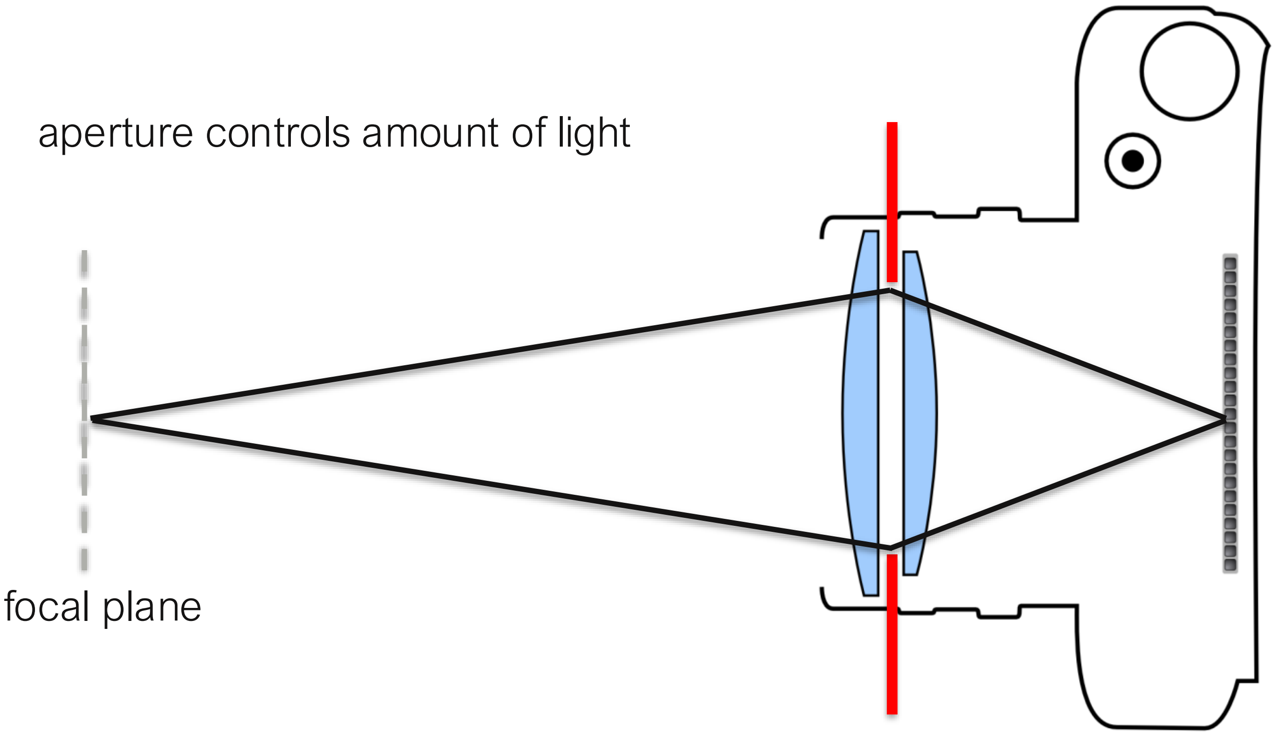

Aperture size

Most lenses have variable aperture size.

- F-number notation: “f/1.4” means $f/ = 1.4$ (**focal length / diameter**).
- Usually aperture sizes available at steps of one-half or one-third stops.
- Older lenses have separate manual aperture ring.
- Modern lenses control the aperture through a dial on the camera body (“gelded” lenses).

f/1.4 f/2.8 f/4 f/8 f/16

Aperture size

Most lenses have variable aperture size.

- F-number notation: “f/1.4” means $f/=1.4$ (**focal length / diameter**).
- Usually aperture sizes available at steps of one-half or one-third stops.
- Older lenses have separate manual aperture ring.
- Modern lenses control the aperture through a dial on the camera body (“gelded” lenses).

f/1.4

f/2.8

f/4

f/8

f/16

Reminder: A “stop” changes the amount of light by a factor of 2.

Aperture

- out of focus blur

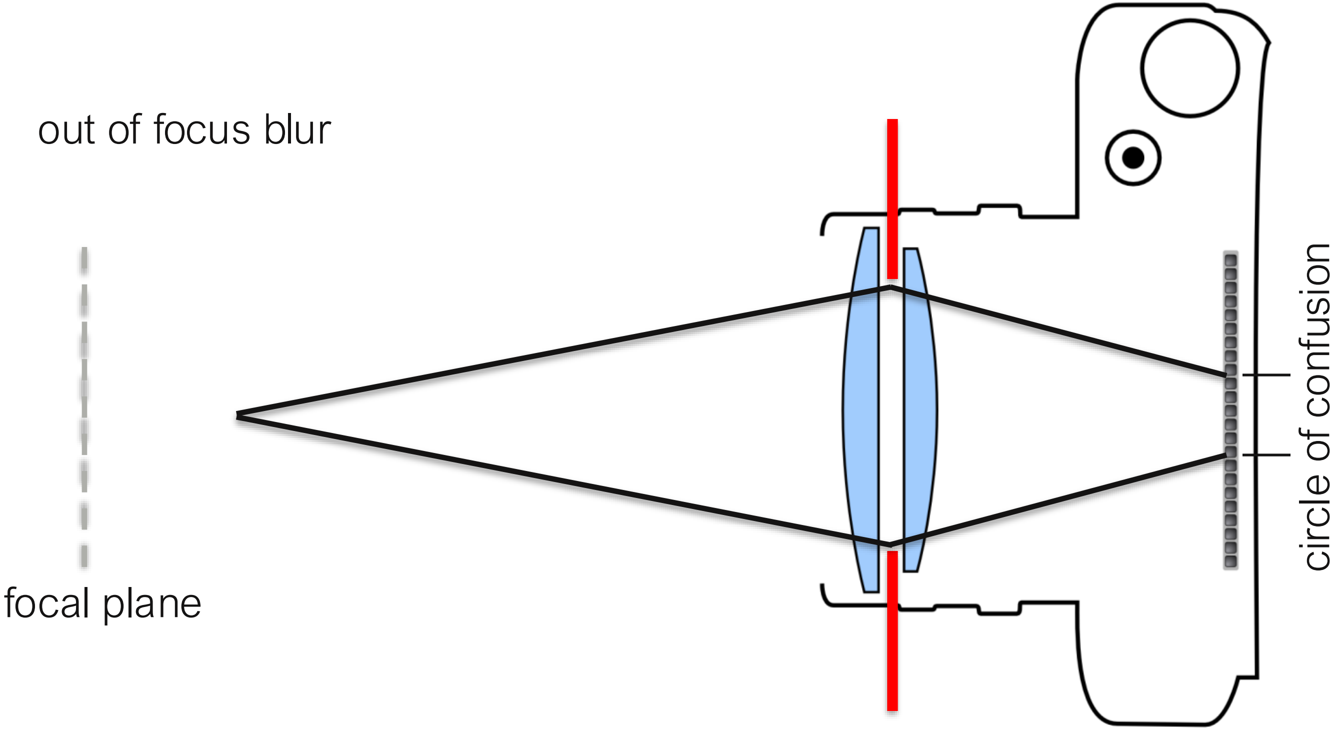

Aperture

- out of focus blur

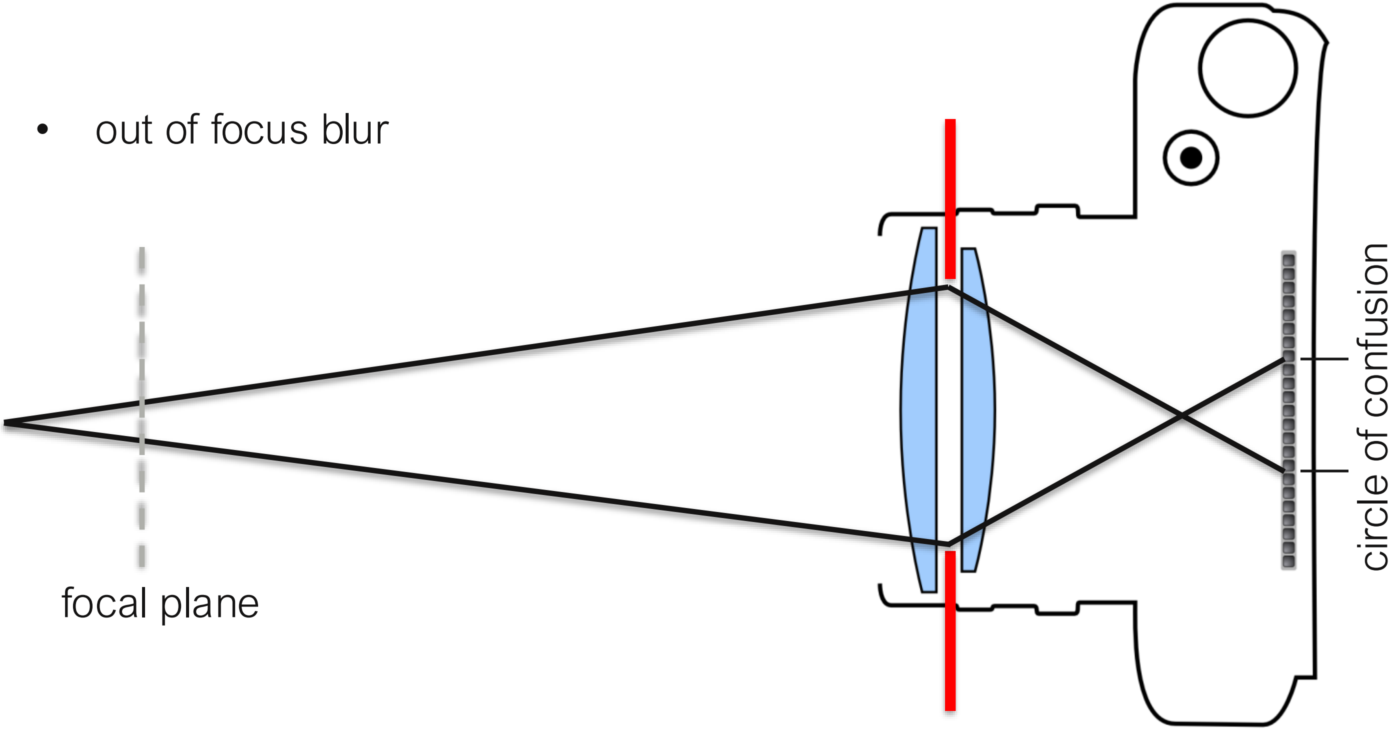

Depth of Field

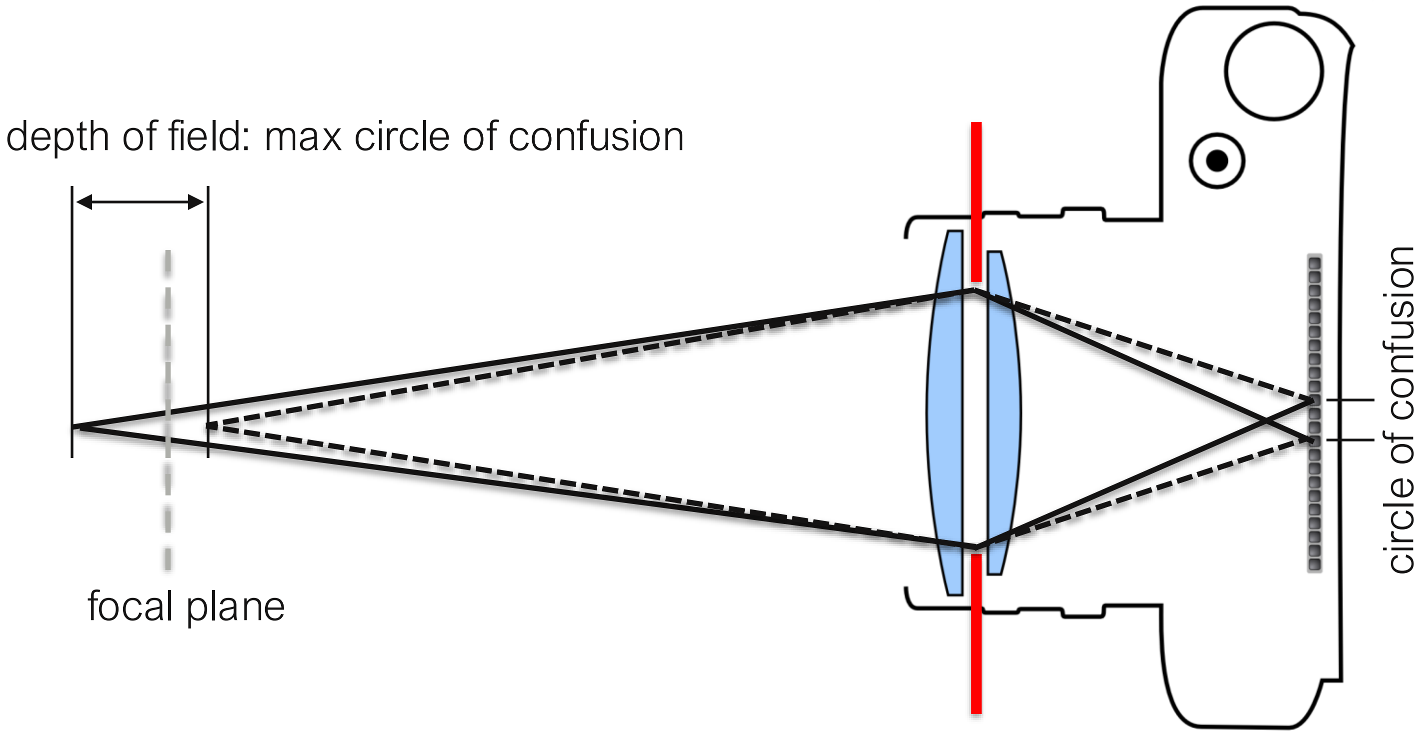

Circle of Confusion

$$c = M \times D \times \frac{|S - S_1|}{S}$$

$$M = \frac{f}{S_1 - f}$$

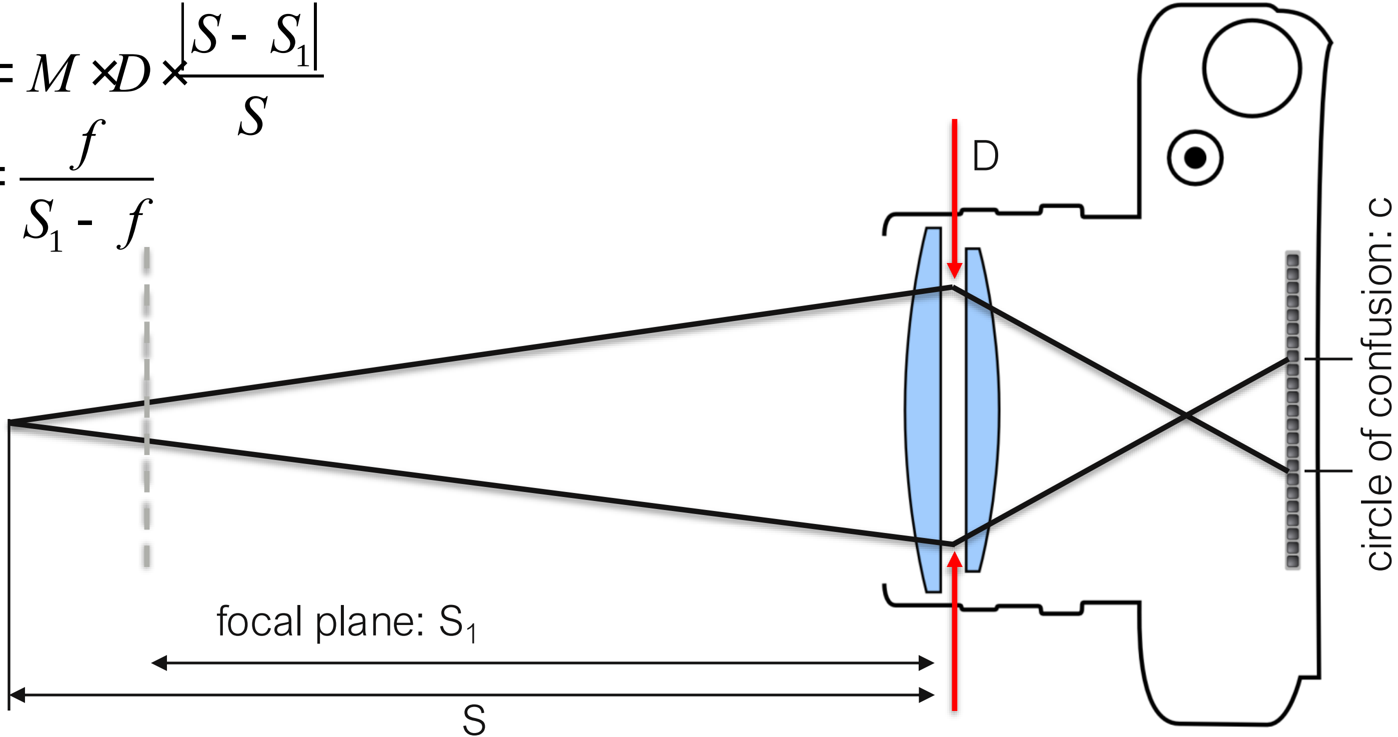

Circle of Confusion

$$c = M \times D \times \frac{|S - S_1|}{S}$$

Canon 5D Mark III: f=50mm, f/2.8 (N=2.8), focused at 5m, pixel size=7.5um

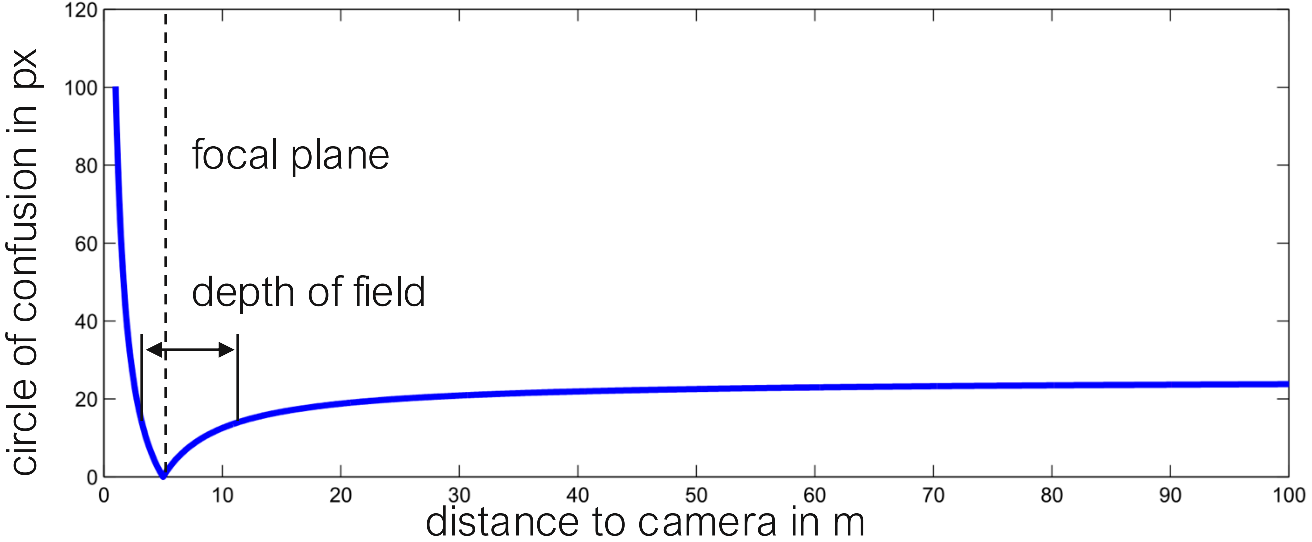

Hyperfocal Distance

$$H = \frac{f^2}{Nc}$$

Canon 5D Mark III: f=50mm, f/2.8 (N=2.8), focused at 5m, pixel size=7.5um

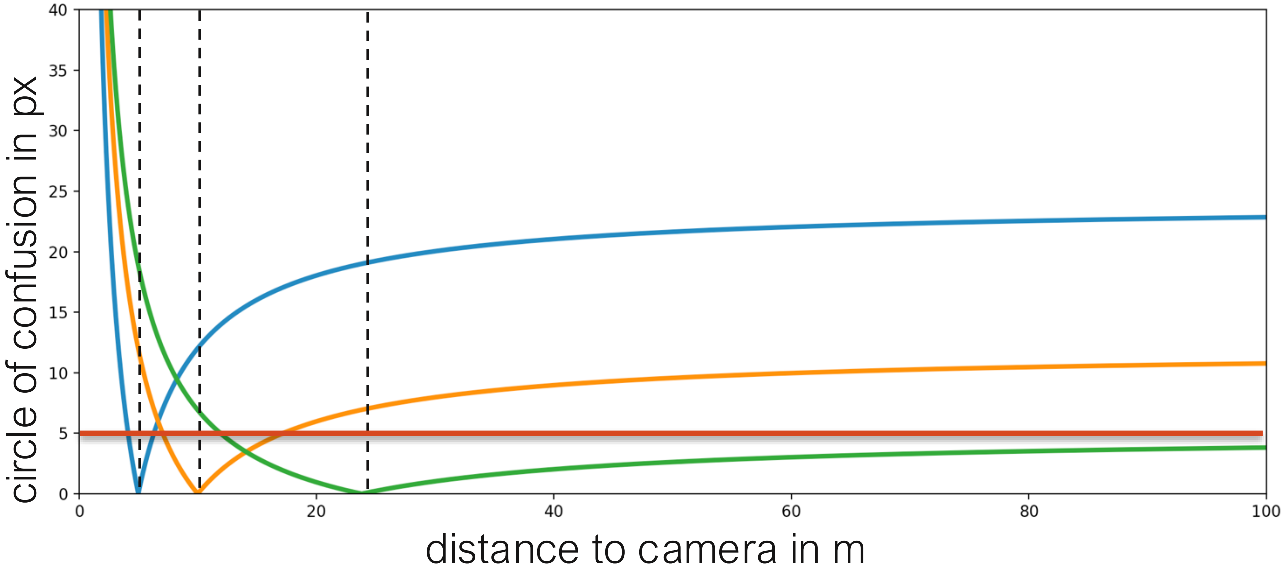

Hyperfocal Distance

$$H = \frac{f^2}{Nc}$$

Canon 5D Mark III: f=50mm, f/2.8 (N=2.8), focused at 5m, pixel size=7.5um

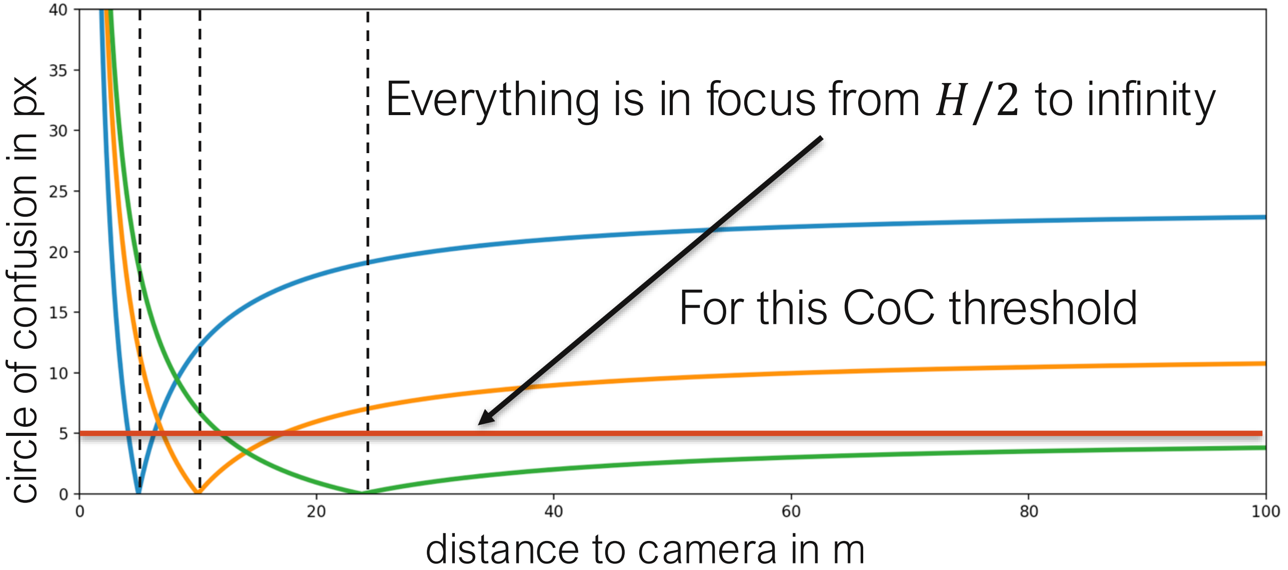

Depth of Field

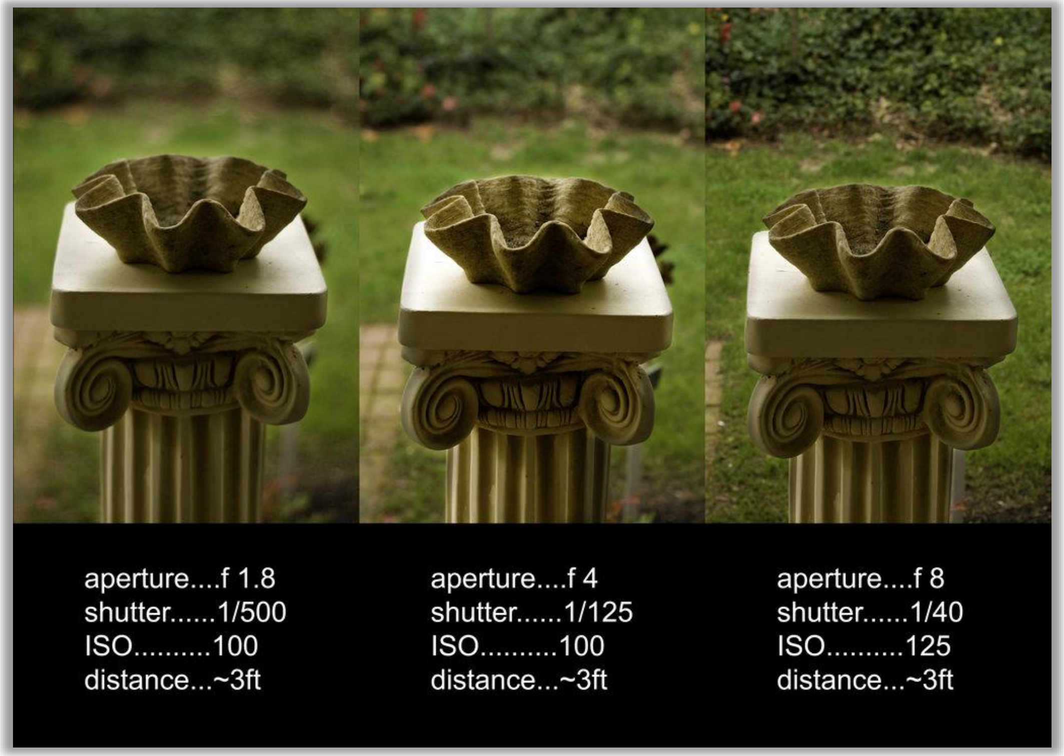

http://photographywisdom.com/

Depth of Field & Motion Blur

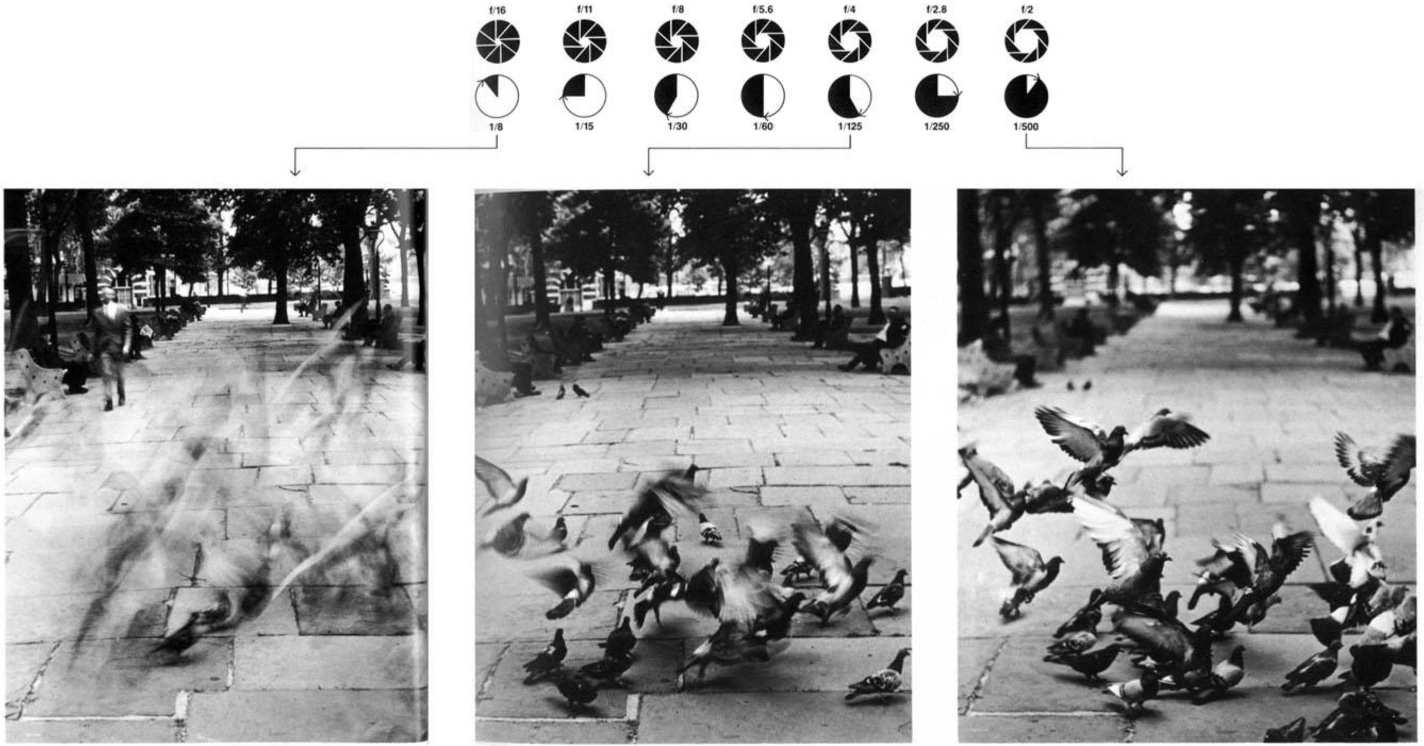

London, Photography

Bokeh

artistic use

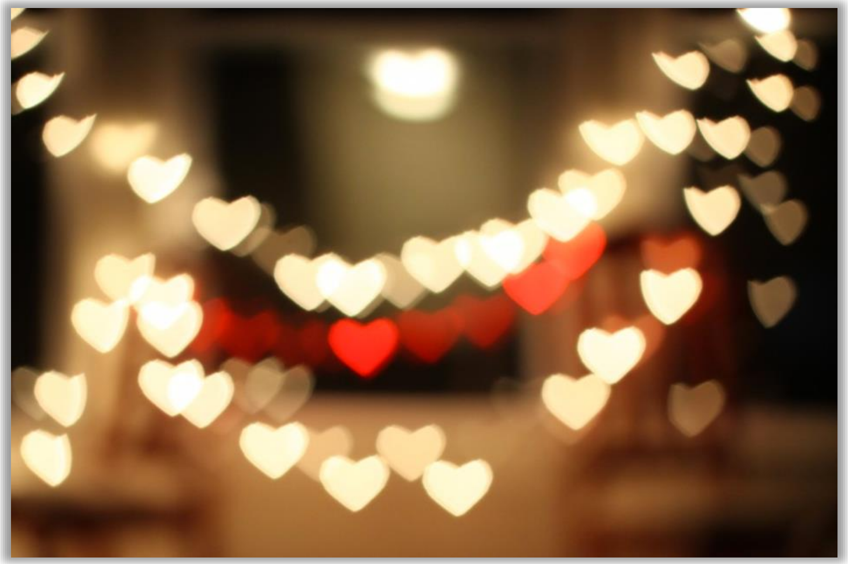

two delighted blog

coded aperture

Levin et al., SIGGRAPH 2007

Diffraction Limit

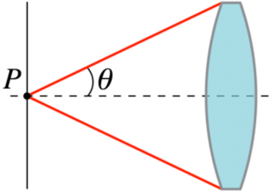

- Ernst Abbe 1873: $d = \frac{\lambda}{2n\sin\theta}$

spot radius (image space)

diffraction

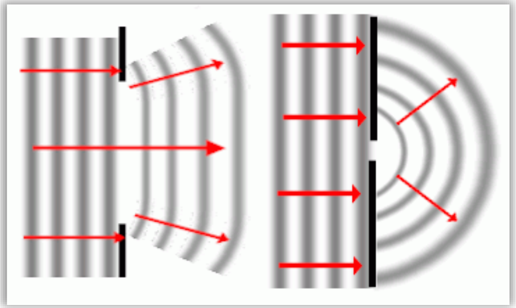

Airy pattern

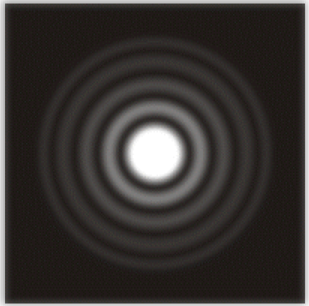

wikipedia

Diffraction Limit

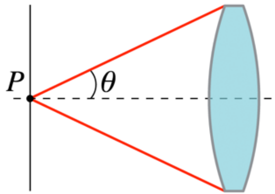

f-number

- Ernst Abbe 1873: $d = \dfrac{\lambda}{2n\sin\theta} = \dfrac{\lambda}{2NA} \approx \lambda N$

numerical aperture

- microscope objectives today: NA 1.4-1.6 → d=λ/2.8
- small f-number (large NA) = high resolution but shallow depth of field
  - inherent tradeoff between “3D” information and 2D resolution
  - space-bandwidth product (uncertainty principle)

Fastest lens ever made?

Zeiss 50 mm f / 0.7 Planar lens

- Originally developed for NASA's Apollo missions.
- Stanley Kubrick somehow got to use the lens to shoot Barry Lyndon under only candlelight.

Fastest lens ever made?

Zeiss 50 mm f / 0.7 Planar lens

- Originally developed for NASA's Apollo missions.
- Stanley Kubrick somehow got to use the lens to shoot Barry Lyndon under only candlelight.

Sensors

What's a Pixel?

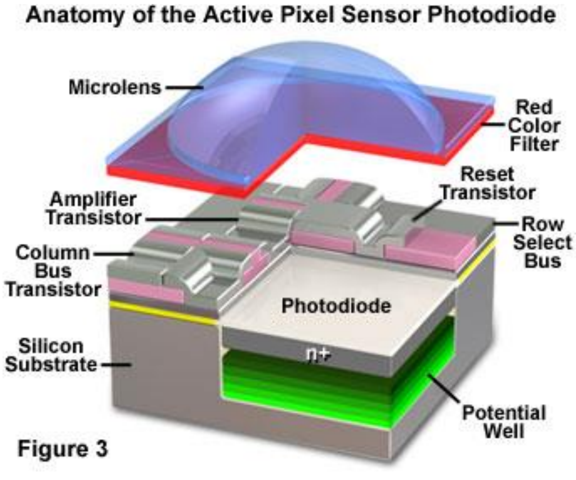

source: Molecular Expressions

photon to electron converter

→ photoelectric effect!

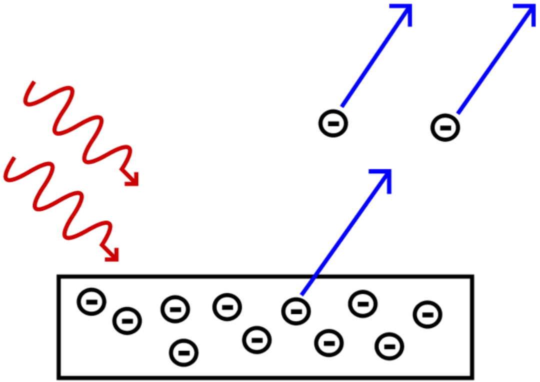

wikipedia

What’s a Pixel?

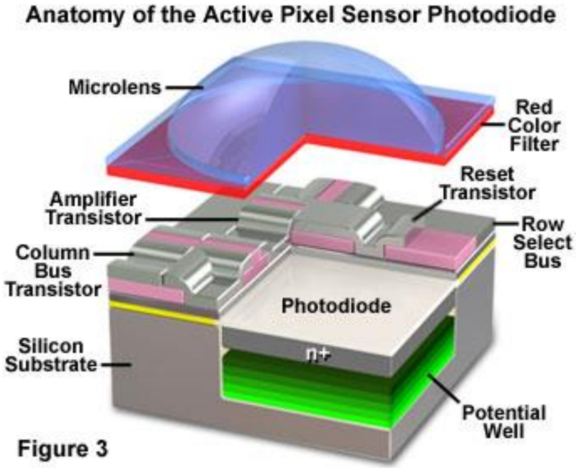

source: Molecular Expressions

- microlens: focus light on photodiode
- color filter: select color channel
- quantum efficiency: ~50%
- fill factor: fraction of surface area used for light gathering

Two main types of imaging sensors

Two main types of imaging sensors

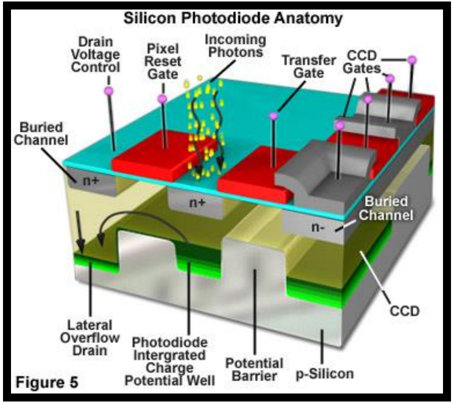

Figure 5

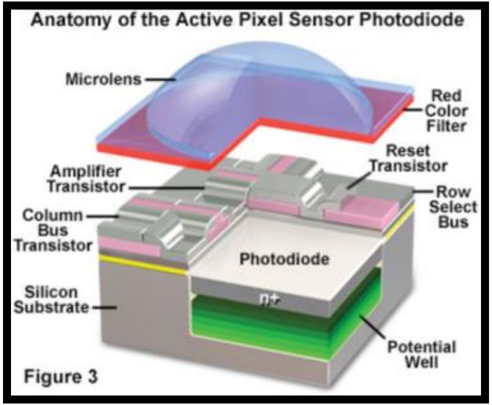

Figure 3

Charged coupled device (CCD):

- row brigade shifts charges row-by-row
- amplifiers convert charges to voltages row-by-row

Complementary metal oxide semiconductor (CMOS):

- per-pixel amplifiers convert charges to voltages
- multiplexer reads voltages row-by-row

Can you think of advantages and disadvantages of each type?

Two main types of imaging sensors

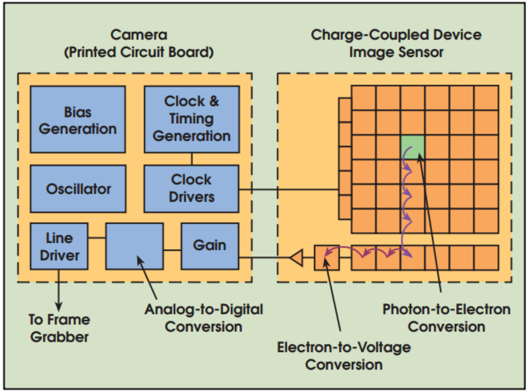

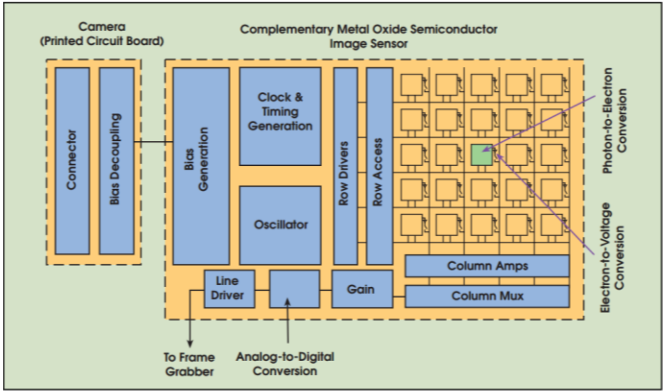

<u>C</u>harged <u>c</u>oupled <u>d</u>evice (CCD):

- row brigade shifts charges row-by-row
- amplifiers convert charges to voltages row-by-row

<u>C</u>omplementary <u>m</u>etal <u>o</u>xide <u>s</u>emiconductor (CMOS):

- per-pixel amplifiers convert charges to voltages
- multiplexer reads voltages row-by-row

Can you think of advantages and disadvantages of each type?

Two main types of imaging sensors

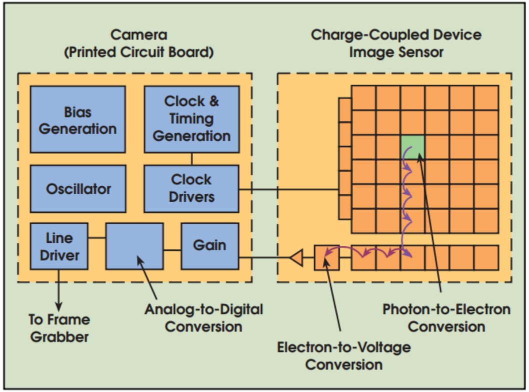

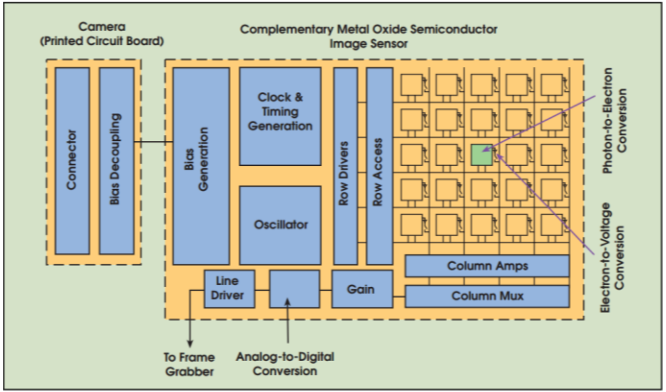

<u>C</u>harged <u>c</u>oupled <u>d</u>evice (CCD):

- row brigade shifts charges row-by-row
- amplifiers convert charges to voltages row-by-row

<u>C</u>omplementary <u>m</u>etal <u>o</u>xide <u>s</u>emiconductor (CMOS):

- per-pixel amplifiers convert charges to voltages
- multiplexer reads voltages row-by-row

✓ higher sensitivity
✓ lower noise

✓ faster read-out
✓ lower cost

What's a Pixel?

What's a Pixel?

What's a Pixel?

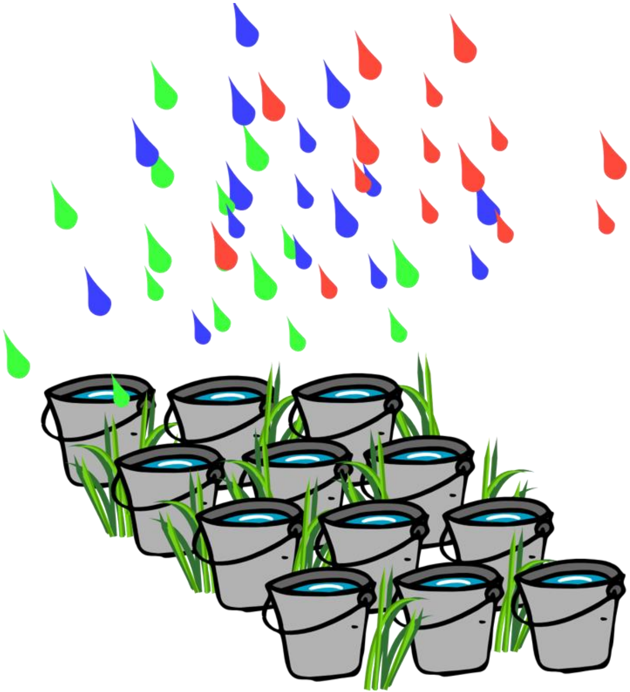

What’s a Pixel?

Most Common: Color Filter Arrays

Bayer pattern

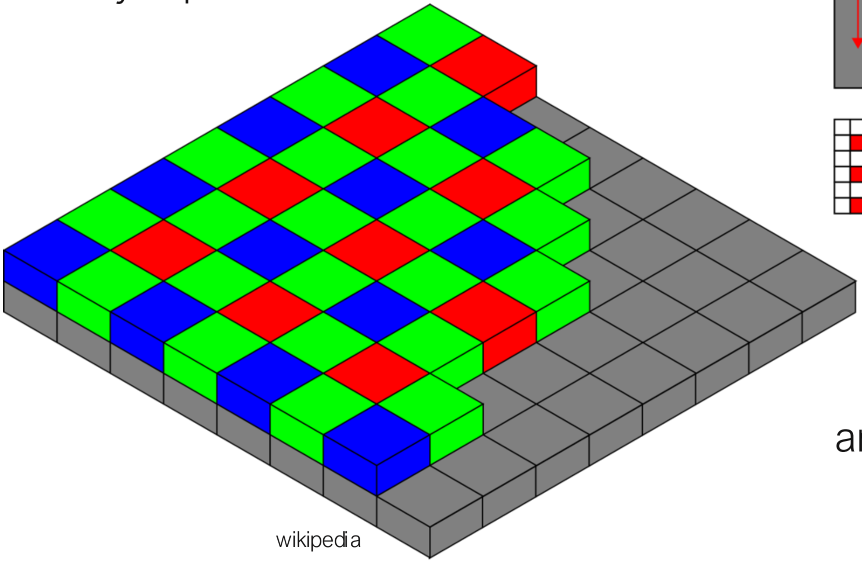

wikipedia

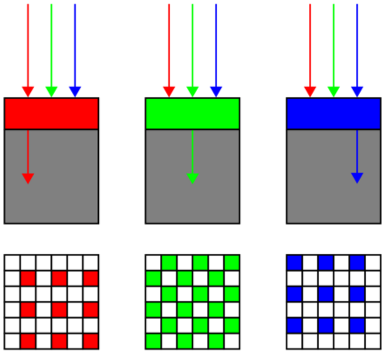

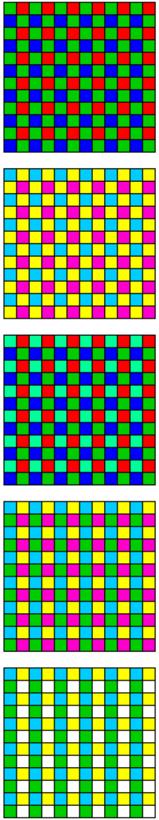

any combination possible

tradeoffs?

Assorted Pixels

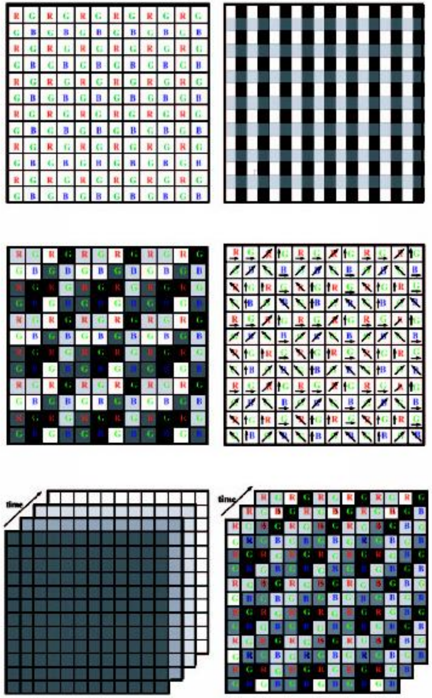

- Narasimhan & Nayar @ Columbia
- multiplex anything: polarization, color, time, ND, …

Exposure (shutter speed)

- exposure = time (e.g. 1/250, 1/60, 1, 15, bulb)

wikipedia

¼ sec, f/3.3,

2 sec, f/6.3,

ISO (“film speed”)

sensor

sensitivity

–

nalog gain applied

before ADC!

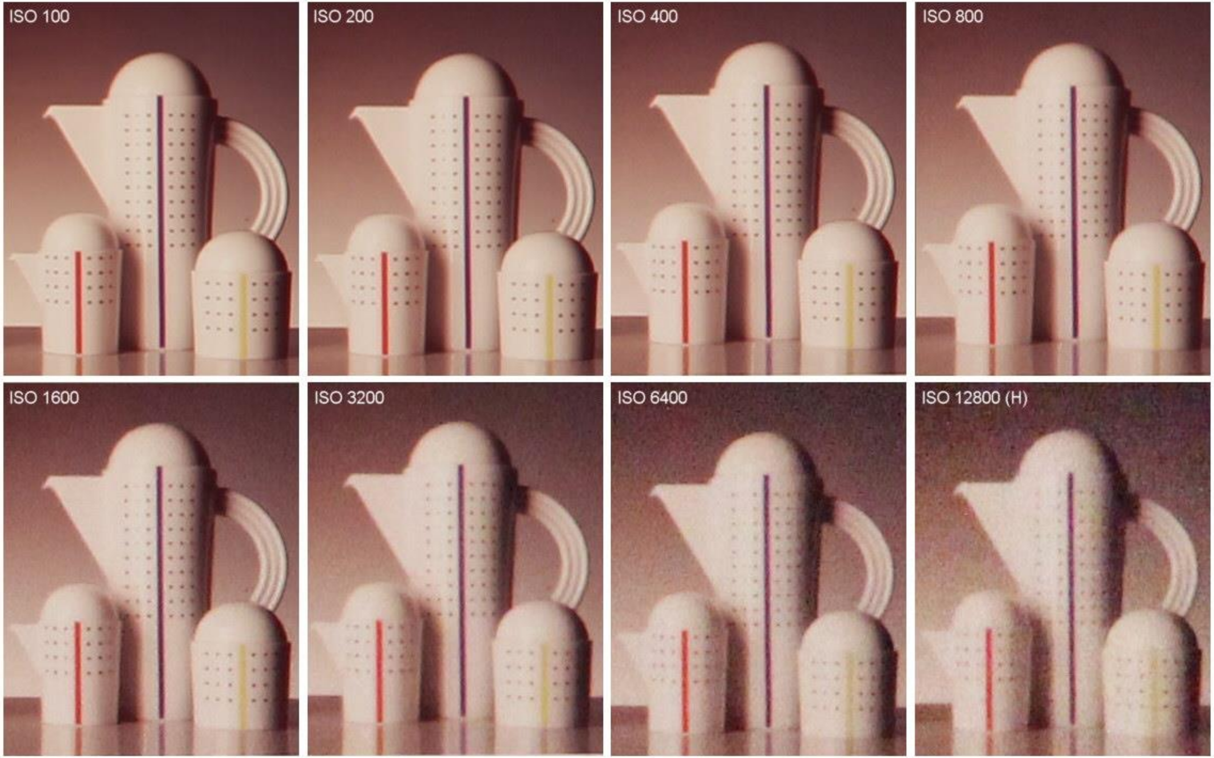

bobatkins.com

Dynamic Range

- ratio between largest and smallest possible value
- bit depth also important! common bit depths: 12-14 bits RAW / 8 bits JPEG

high dynamic range

Kevin McCoy

Global Shutter vs. Rolling Shutter

All sensor pixels exposed at same time

Row-by-row readout of image

- shorter exposure times per pixel
- motion artifacts

http://lfa.mobivap.uva.es/~fradelg/phd/notes/global-shutter.html

What are these dark bands?

60 Hz AC power results in 120 Hz flicker!

YouTube: user cameratest

[Sheinin et al. ‘17]

26 frames over 10 ms

[Sheinin et al. ‘17]


26 frames over 10 ms
[Sheinin et al. '17]

Photons to RAW Image

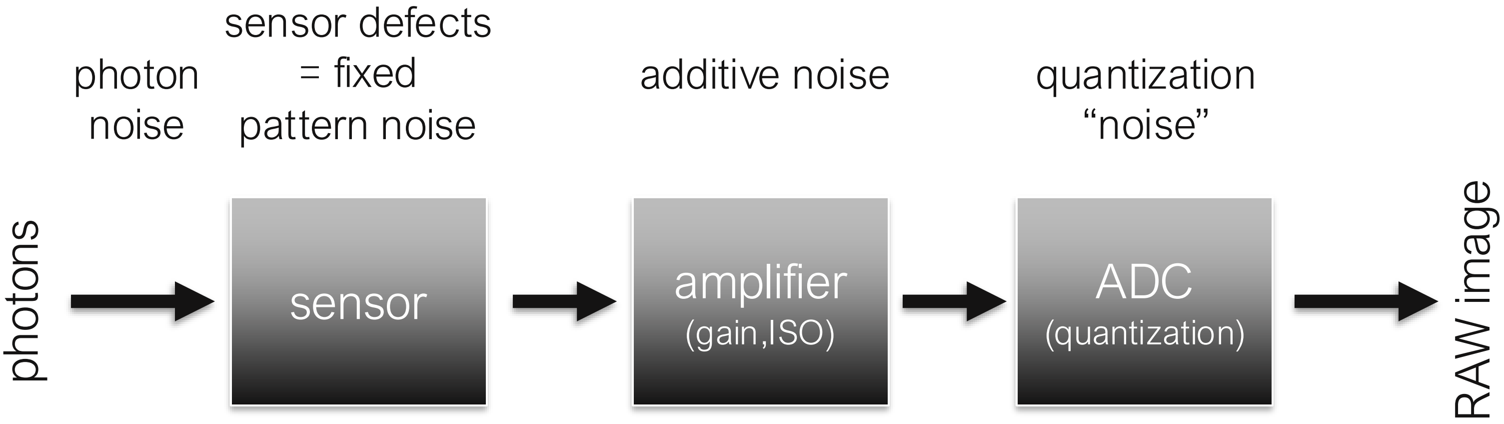

Sensor Noise

- noise is (usually) bad!
- many sources of noise: heat, electronics, amplifier gain, photon to electron conversion, pixel defects, read, …
- different noise follows different statistical distributions, two crucial ones:
  - Gaussian
  - Poisson

Gaussian Noise

- thermal, read, amplifier
- additive, signal-independent!

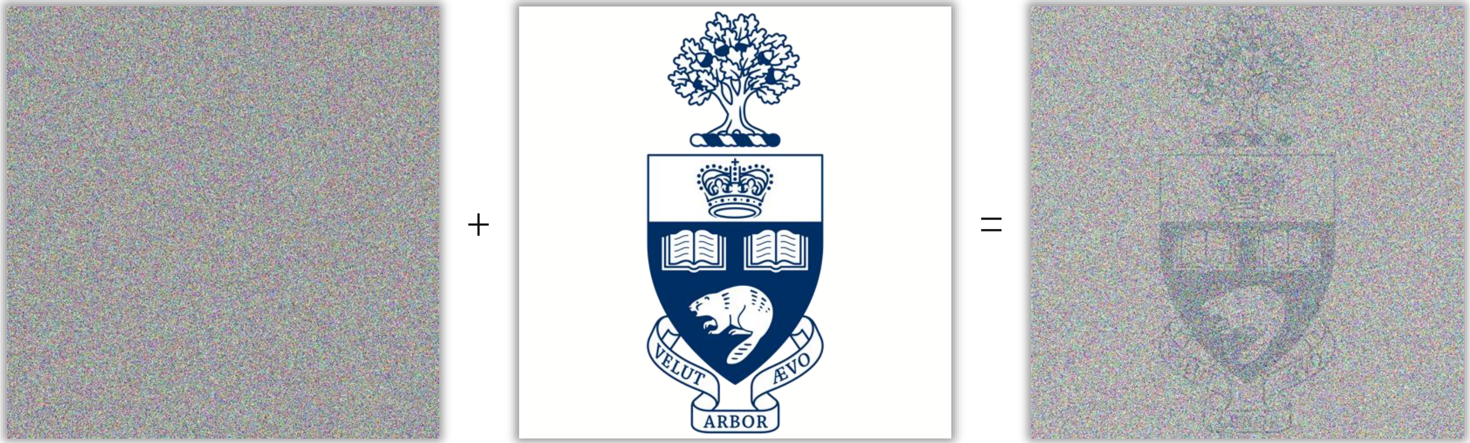

Photon or Shot Noise

- signal dependent

- Poisson distribution:

$$f(k; /) = \frac{/^{k} e^{- /}}{k!}$$

$$s = \sqrt{/}$$

wikipedia

N photons: $s = \sqrt{N}$

2N photons: $s = \sqrt{2}\sqrt{N}$

nonlinear!

Signal-to-Noise Ratio (SNR)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{m}{s}$$

$m \leftarrow$ signal

$s \leftarrow$ noise

$$= \frac{PQ_e t}{\sqrt{PQ_e t + Dt + N_r^2}}$$

P = incident photon flux (photons/pixel/sec)

Q_e = quantum efficiency

t = eposure time (sec)

D = dark current (electroncs/pixel/sec), including hot pixels

N_r = read noise (rms electrons/pixel), including fixed pattern noise

Scientific Sensors

- e.g., Andor iXon Ultra 897: cooled to -100° C
- scientific CMOS & CCD
- reduce pretty much all noise, except for photon noise

Digital Photography

- optics
- aperture
- depth of field
- field of view
- exposure
- noise
- color filter arrays
- image processing pipeline

Digital Photography – Additional Resources

- What we left out: metering, autofocus, autoexposure, anti-aliasing filter, IR filter (and probably much more)
- Stanford CS 178 – Digital Photography: slides, applets, and other material online
- CMU Computational Photography 15-862
- looking for a camera? check dpreview.com

Next: The Image Processing Pipeline

- RAW images
- demosaicking
- denoising
- deblurring
- white balancing
- gamma correction
- compression

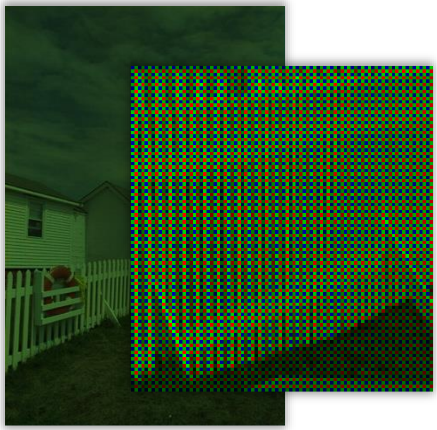

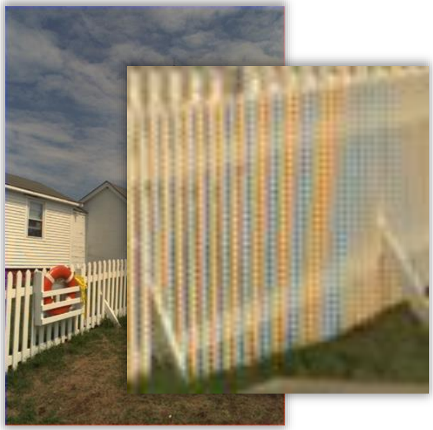

References and Further Reading

- London, Upton, Stone, “Photography”, Pearson, 11th edition, 2013
- Stanford CS 178, “Digital Photography”, Course Notes
- CMU Computational Photography course
- wikipedia