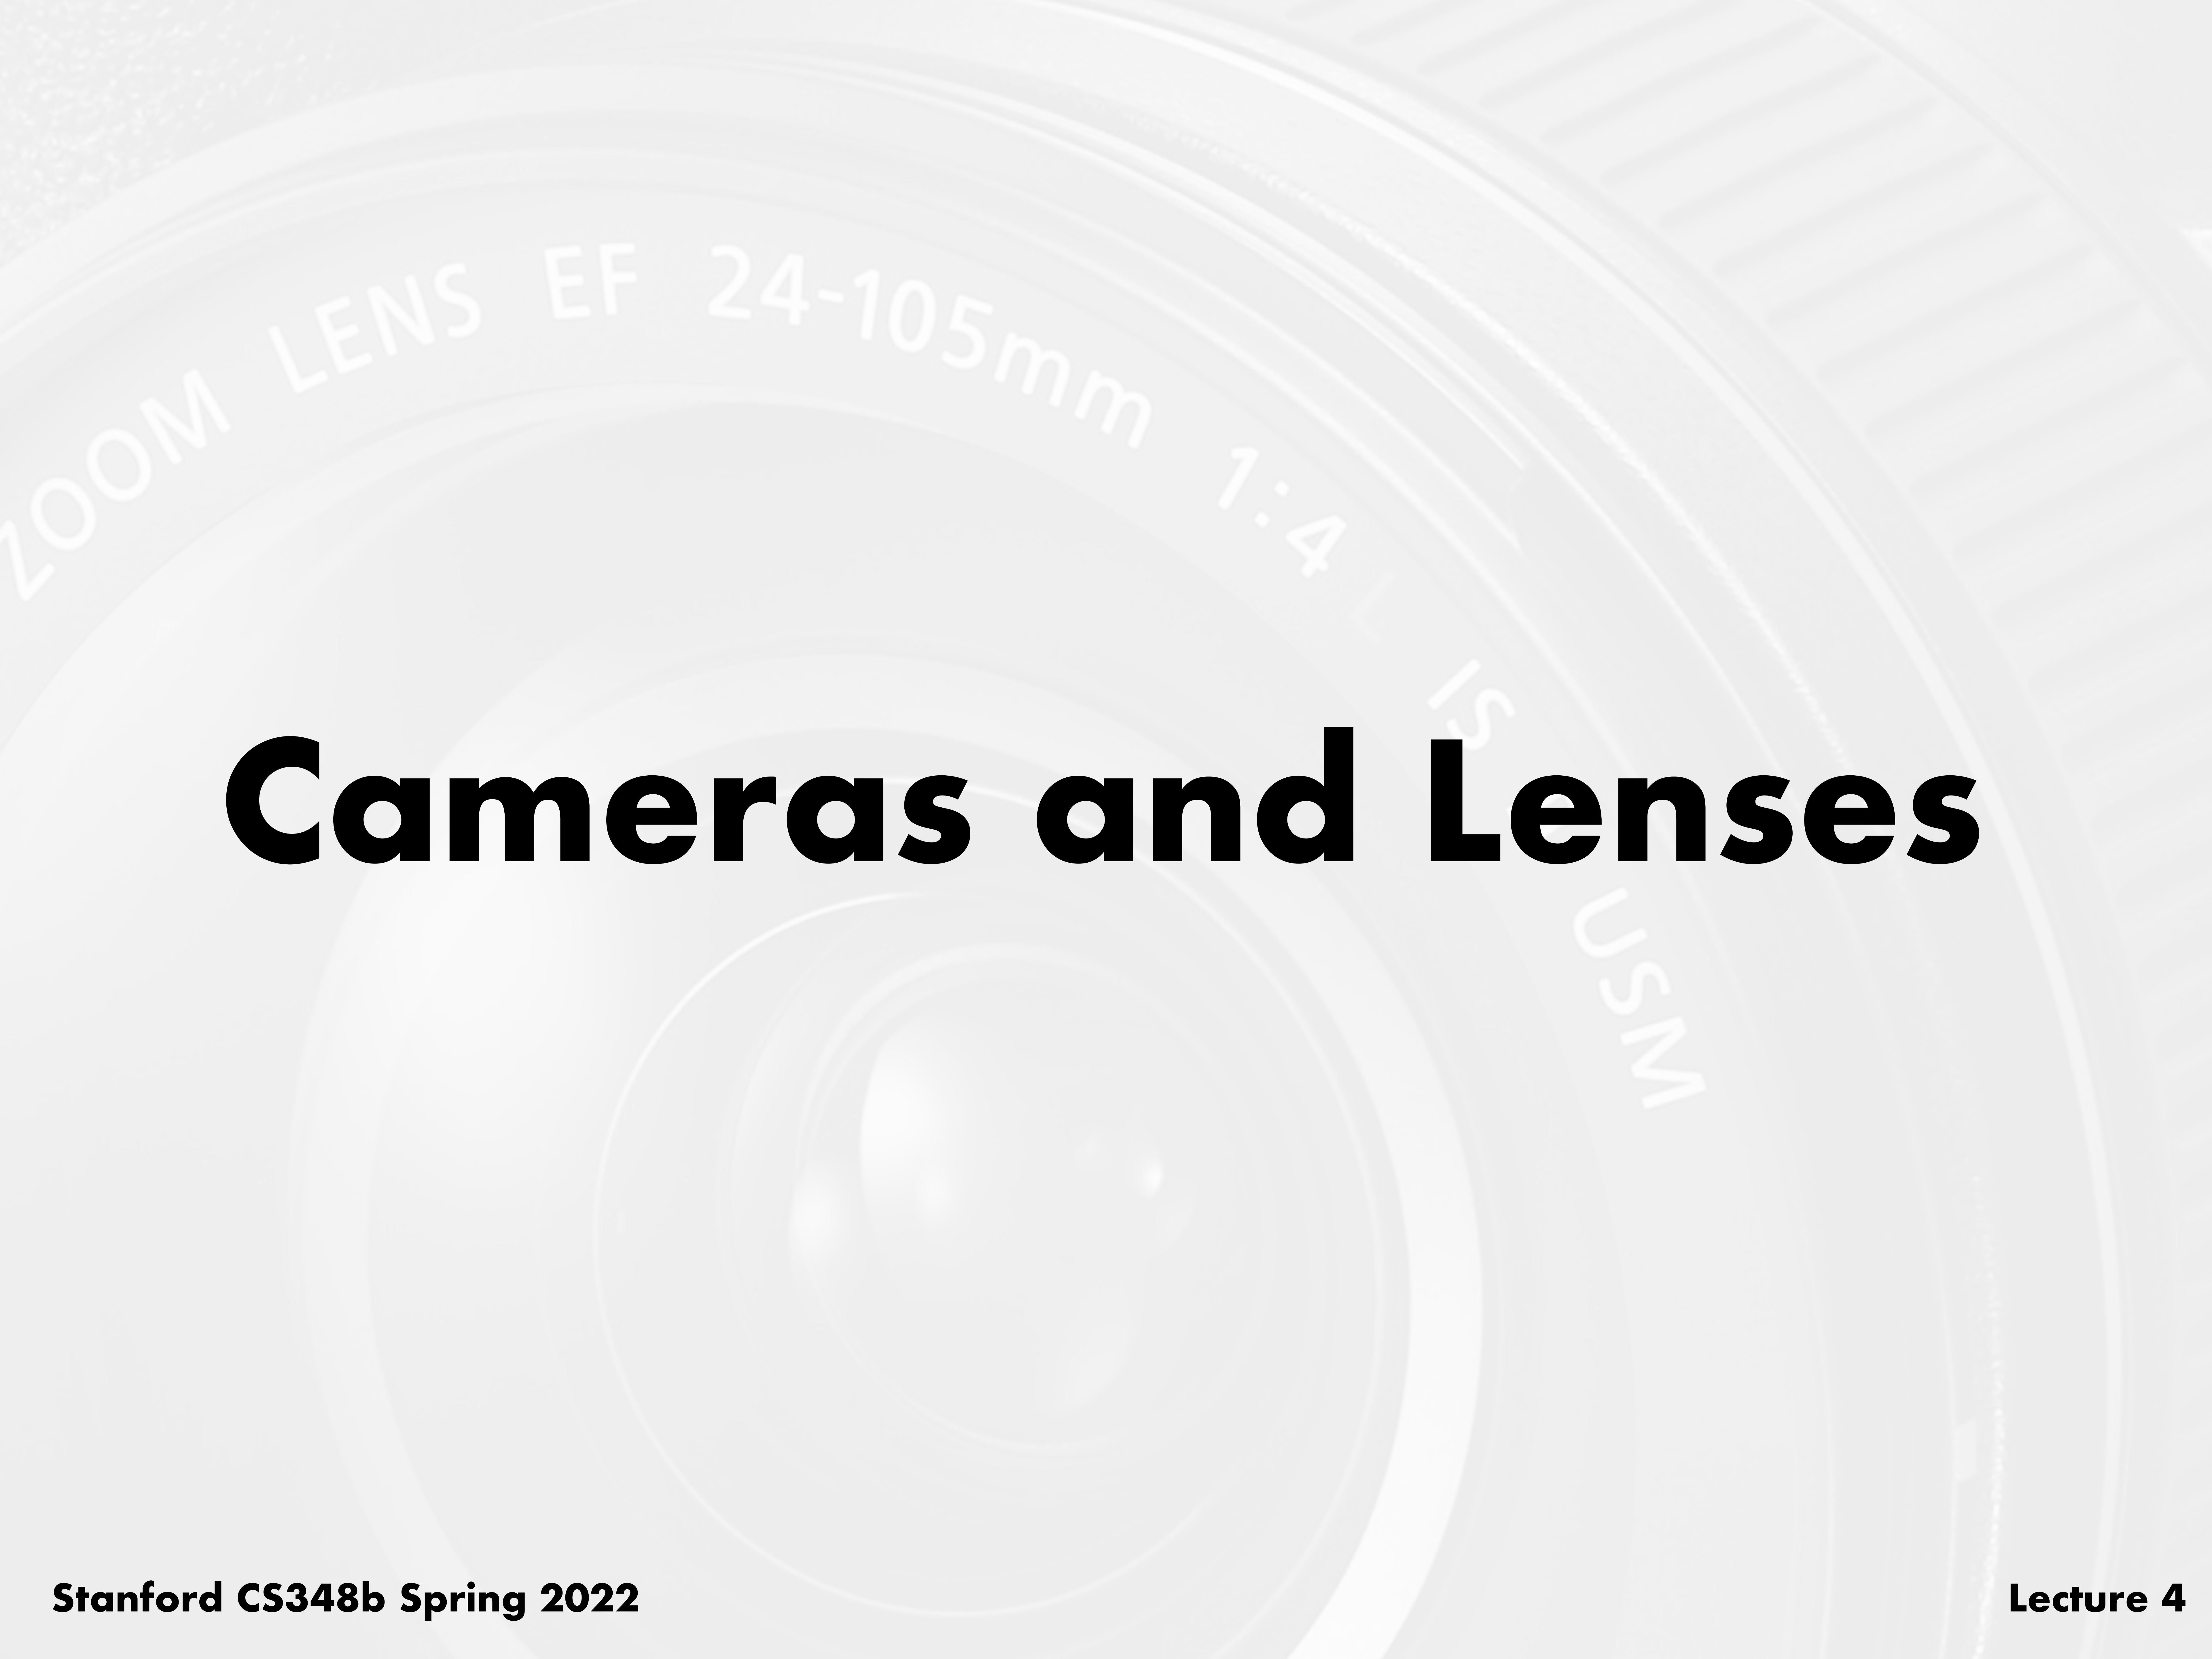

Cameras and Lenses


Image credit: Canon

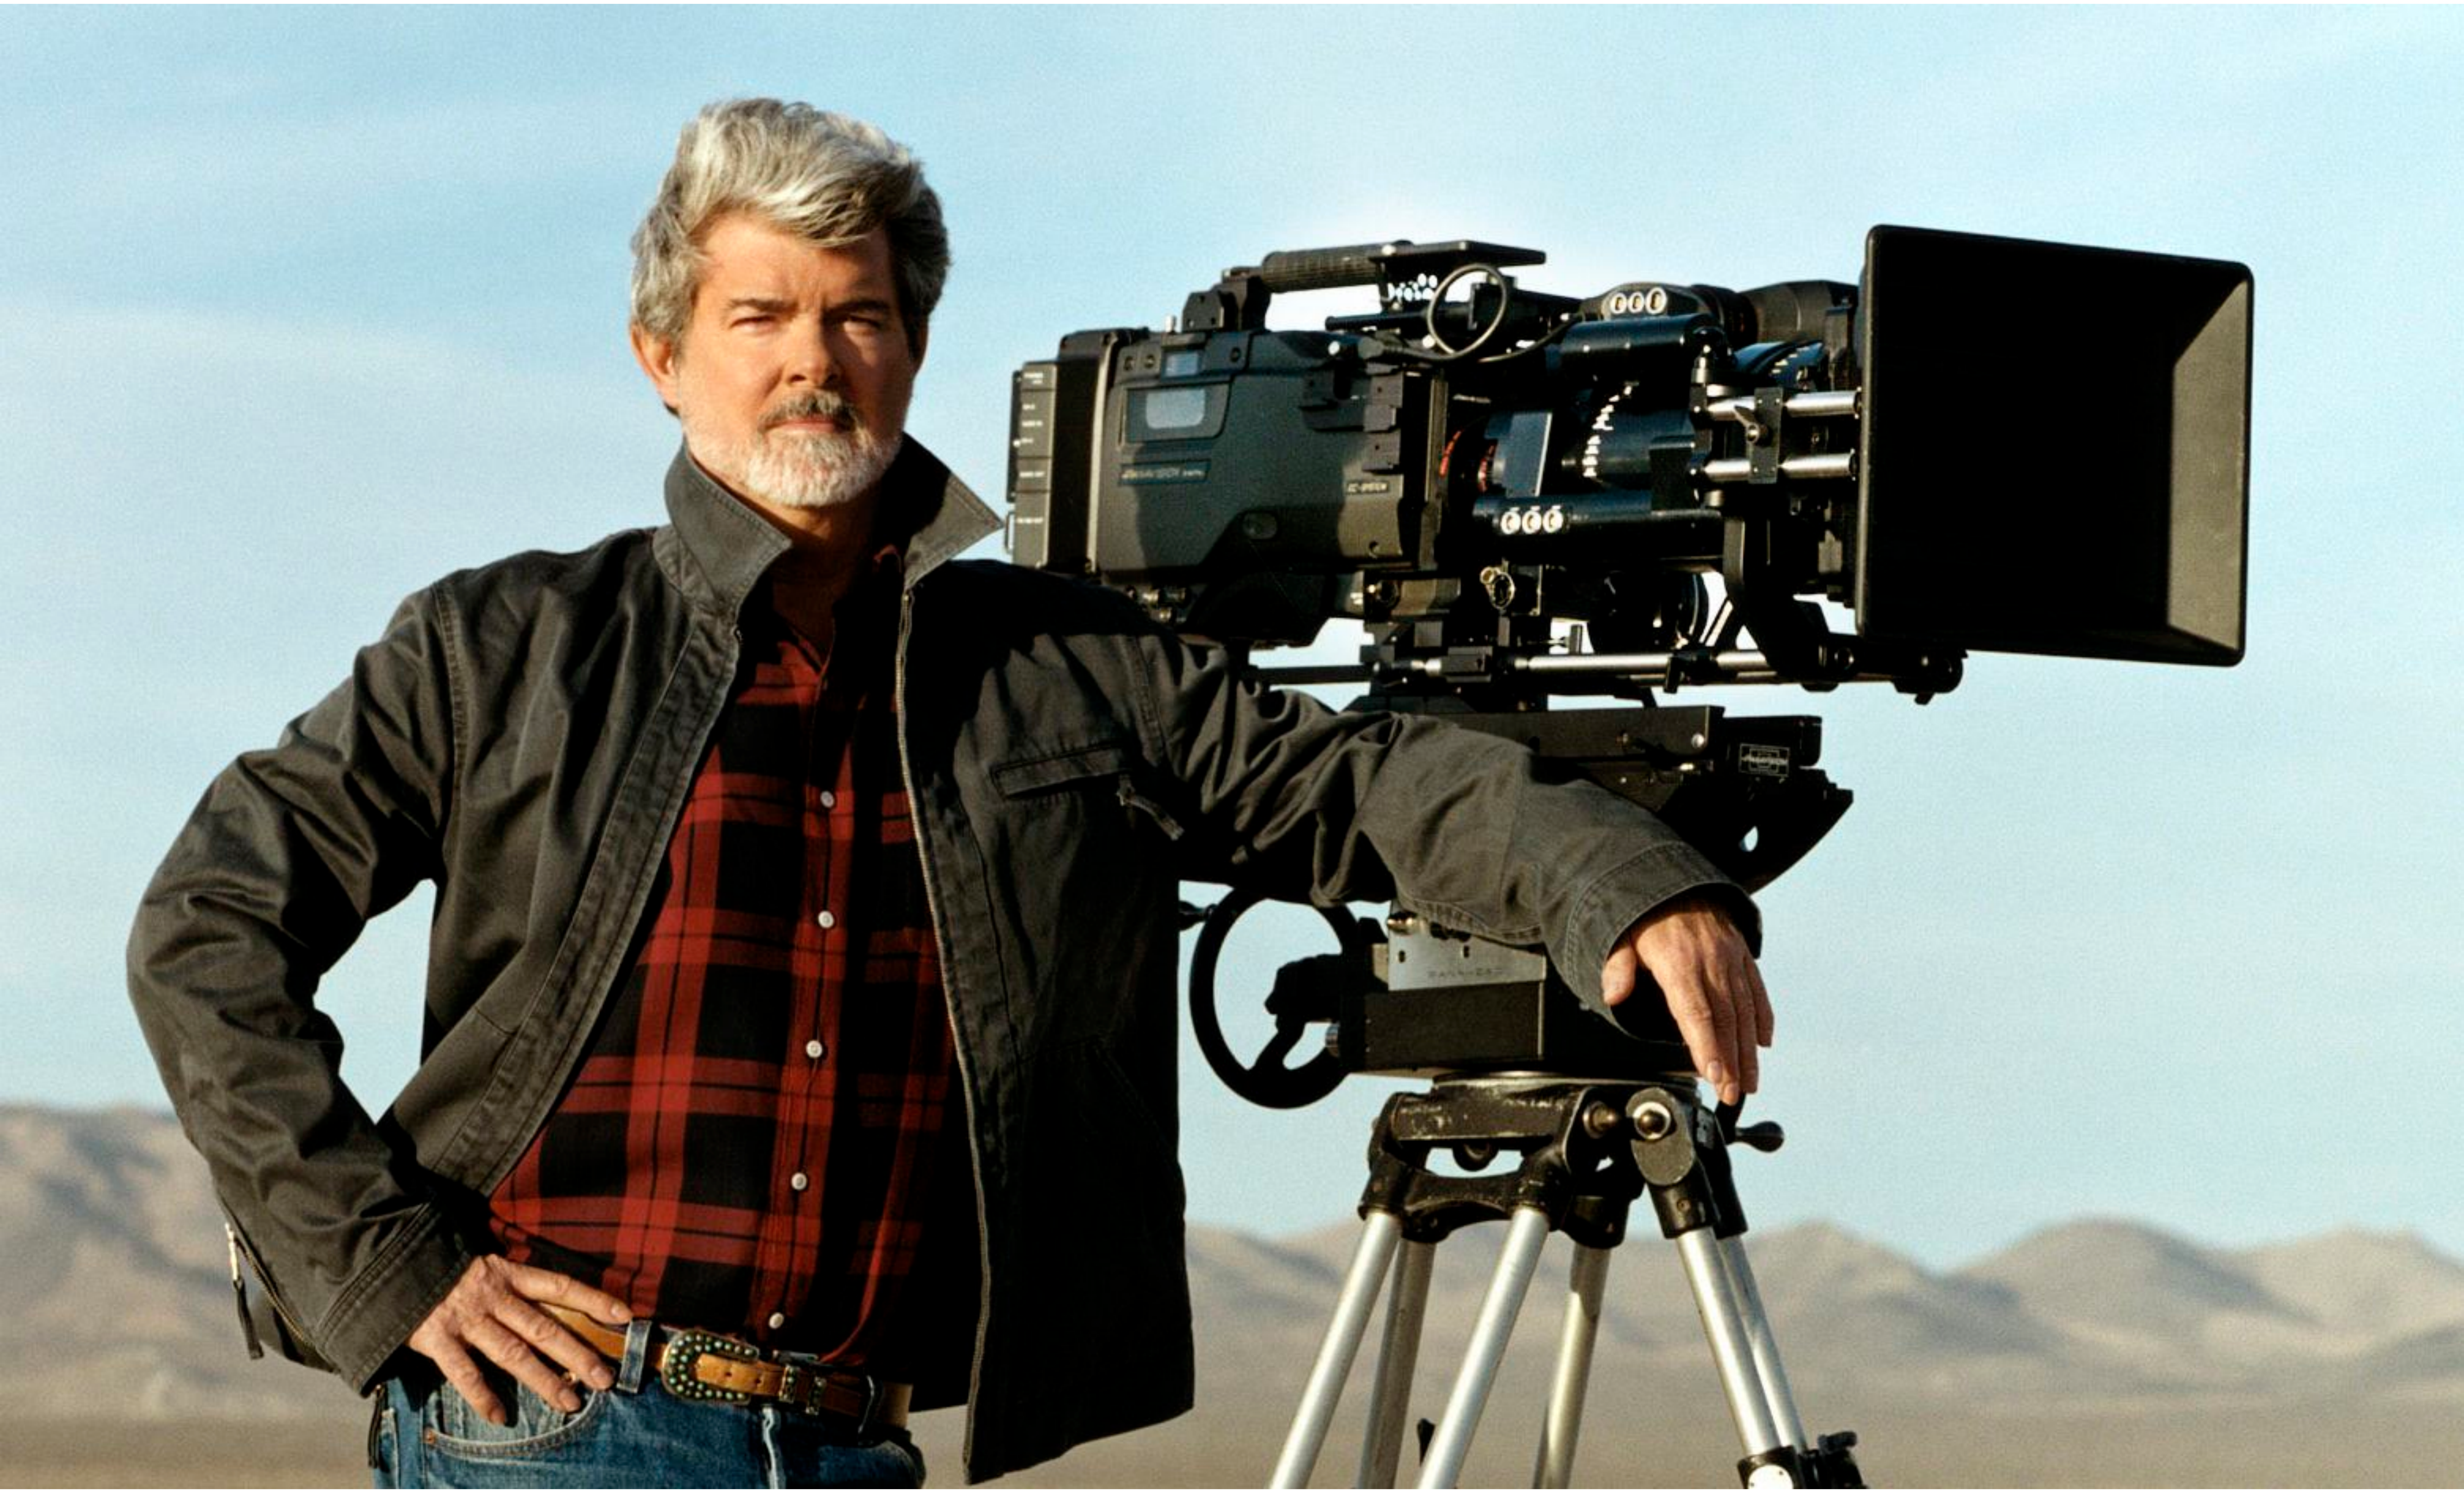


"Tangerine" (2015)
Shot with iPhone 5

Today

Camera basics

- Lenses
- Exposure
- Depth of Field
- Light field cameras

Simulating these effects in a ray tracer

Rendering with lens defocus blur

Credit: Pixar

Rendering with motion blur

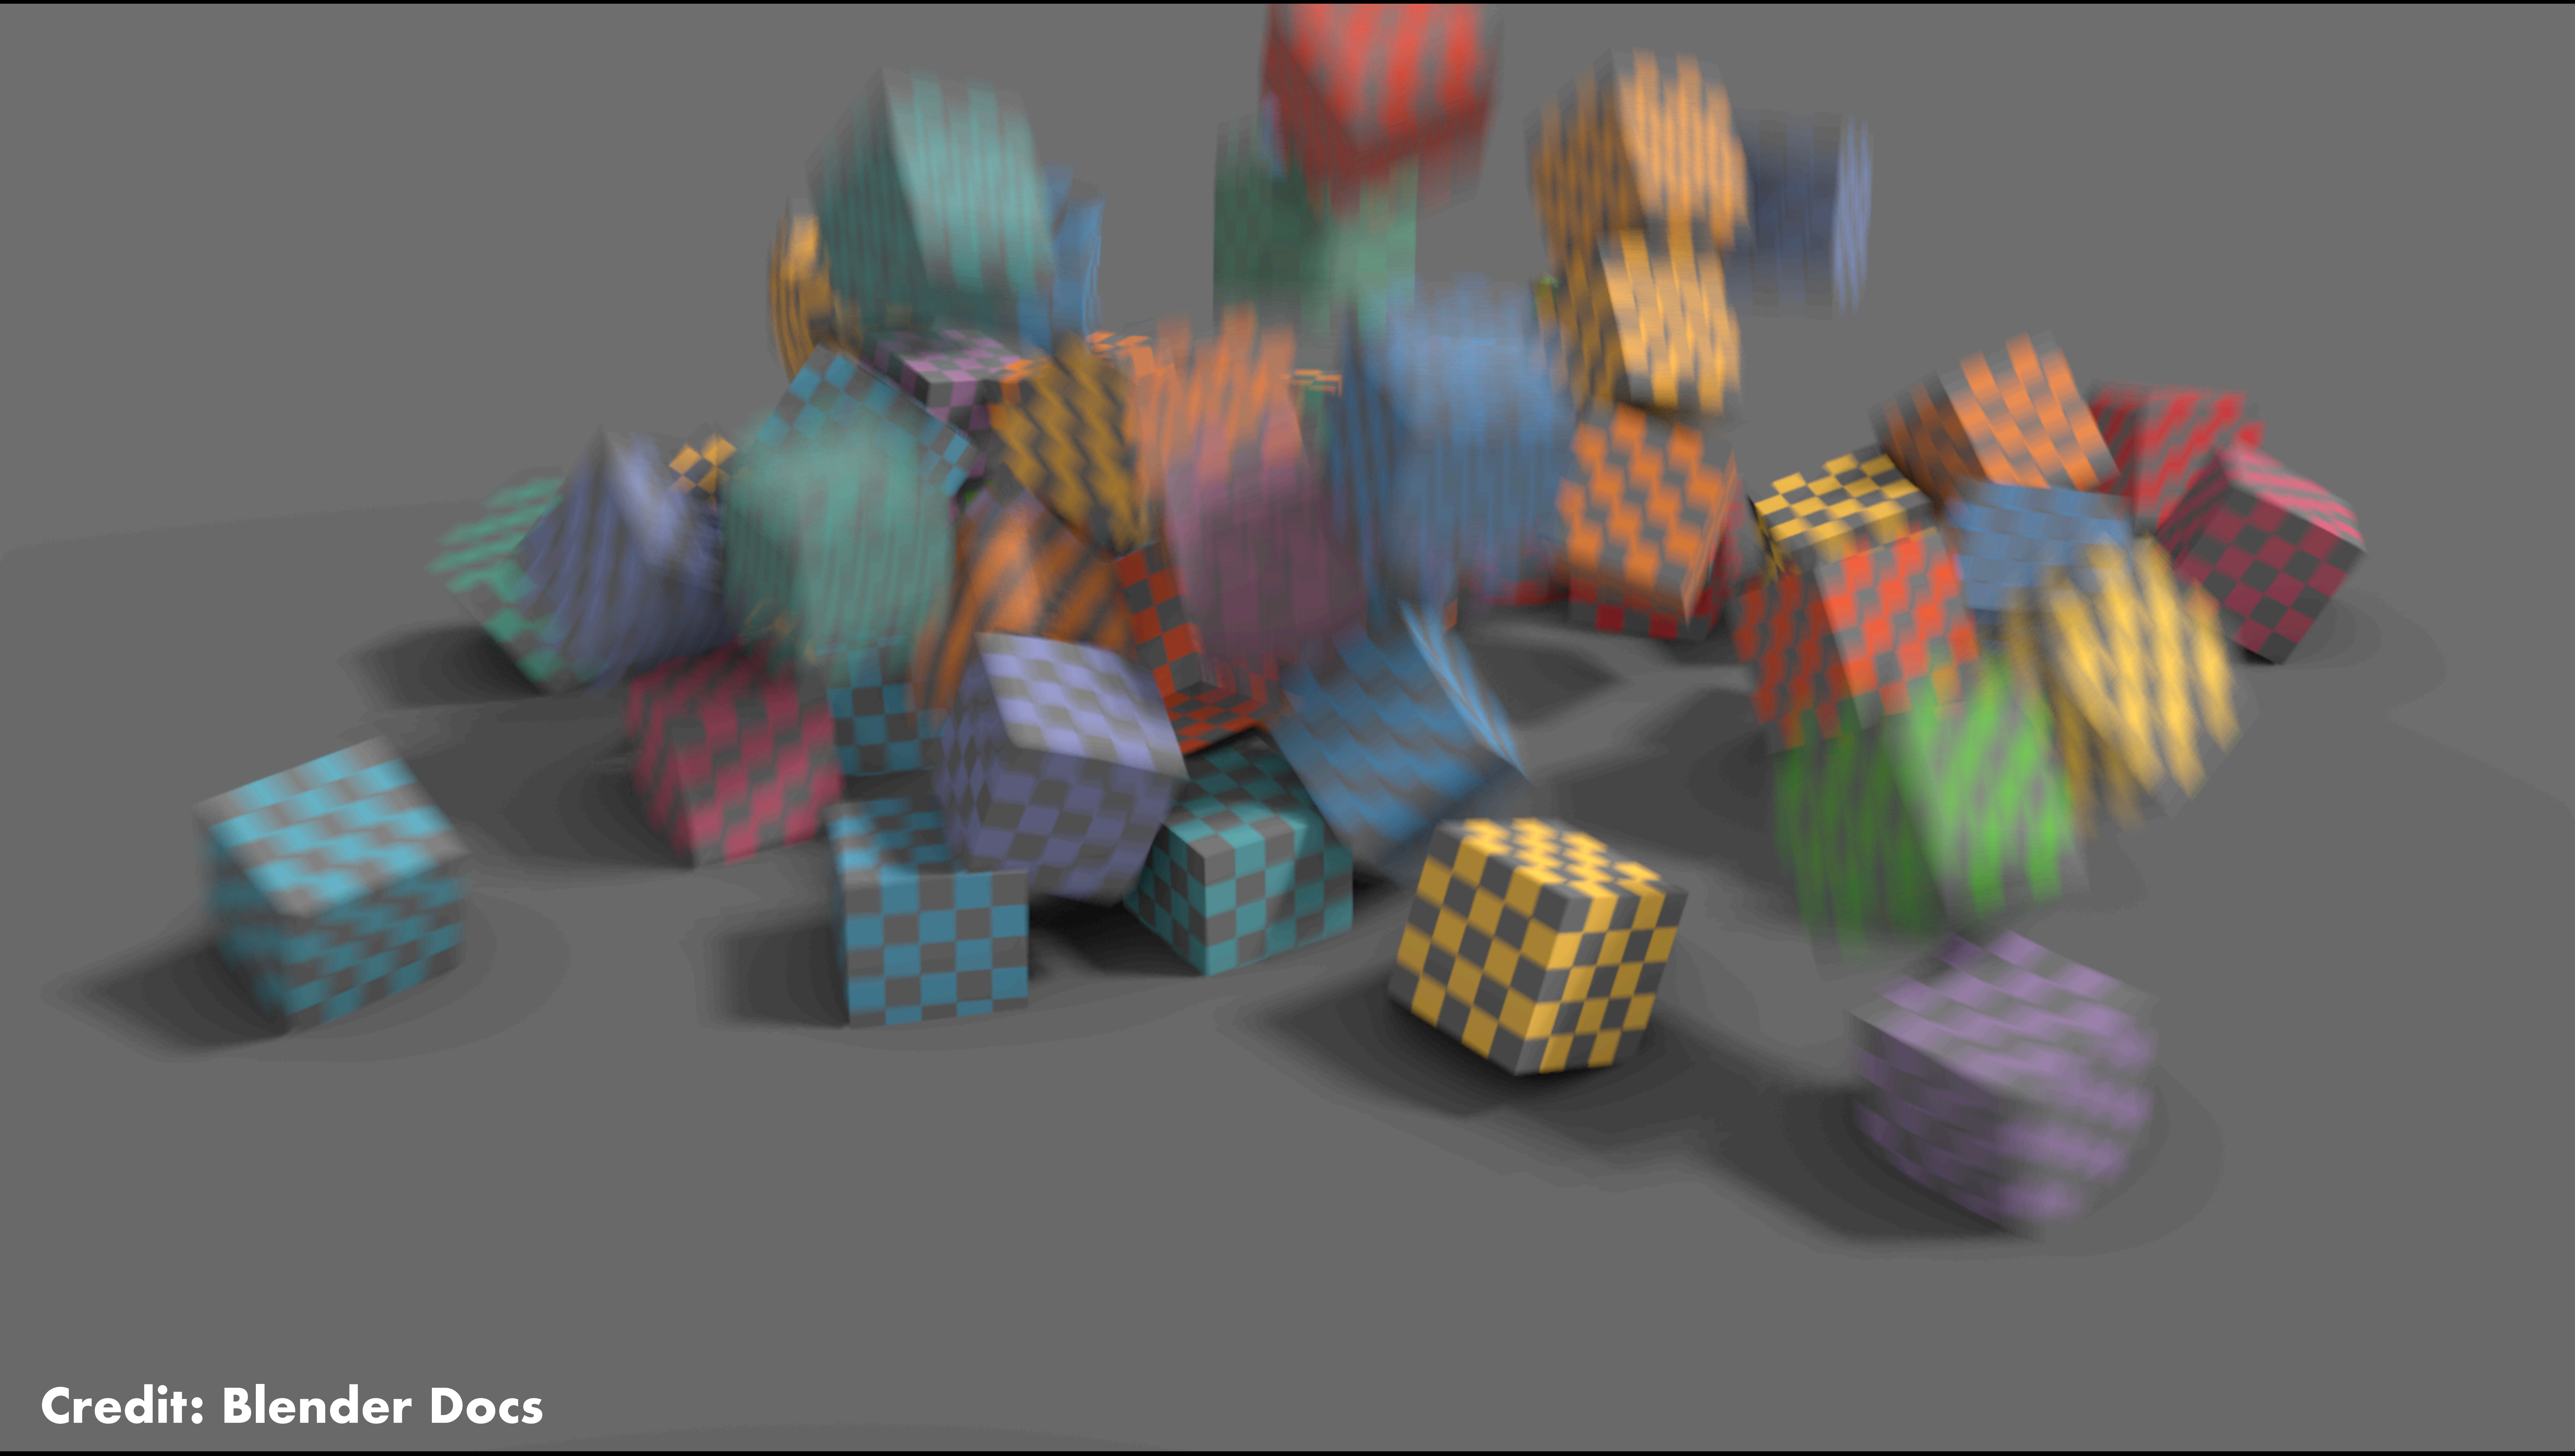

Credit: Blender Docs

What's Happening Inside the Camera?

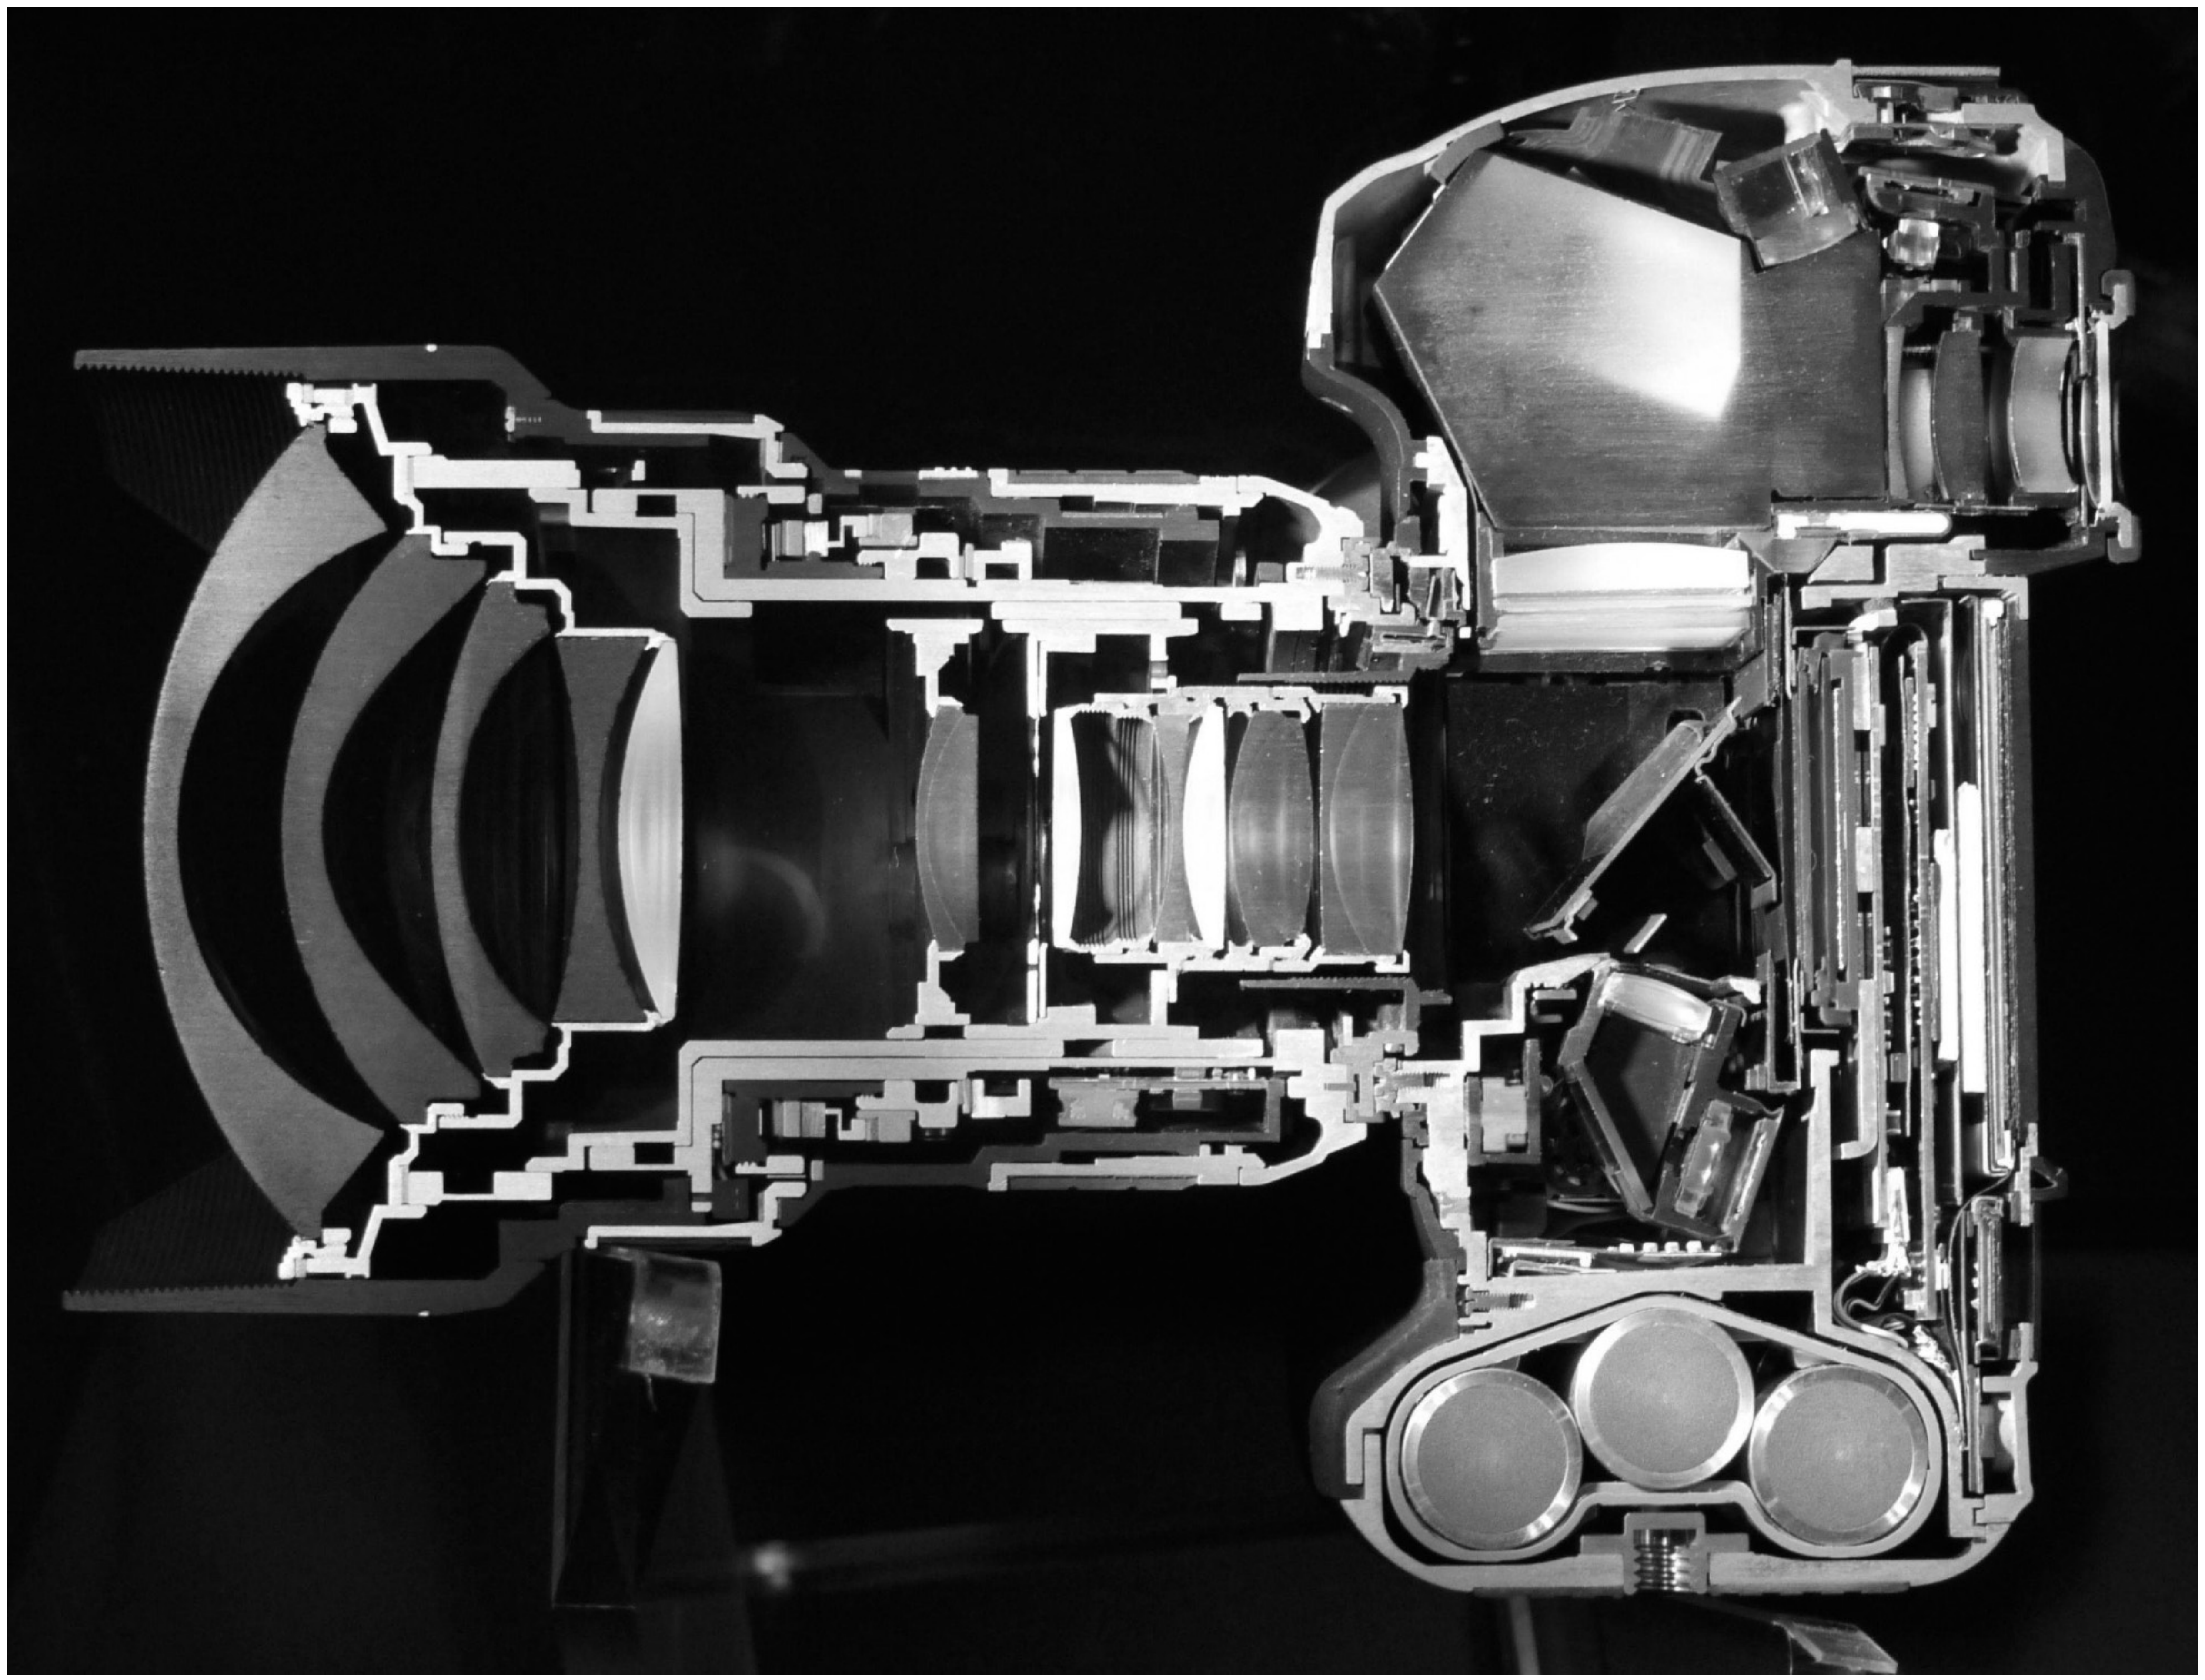

Cross-section of Nikon D3, 14-24mm F2.8 lens

A much simpler camera

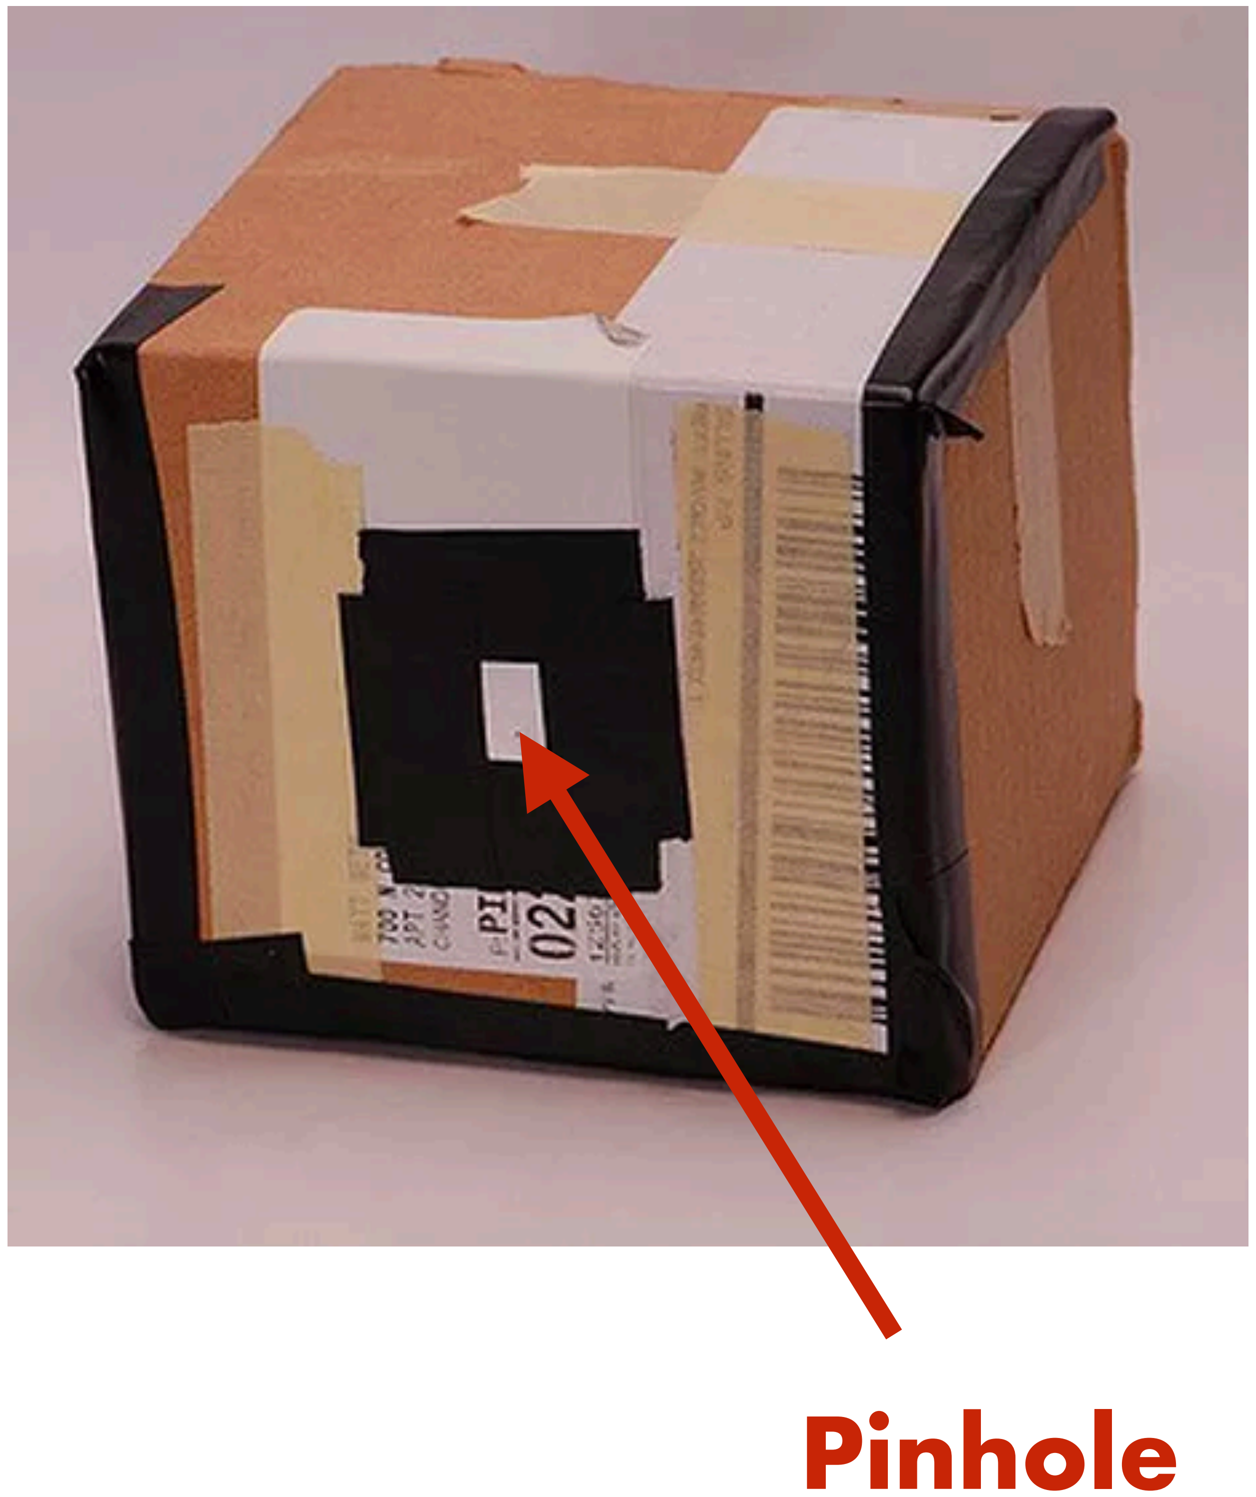

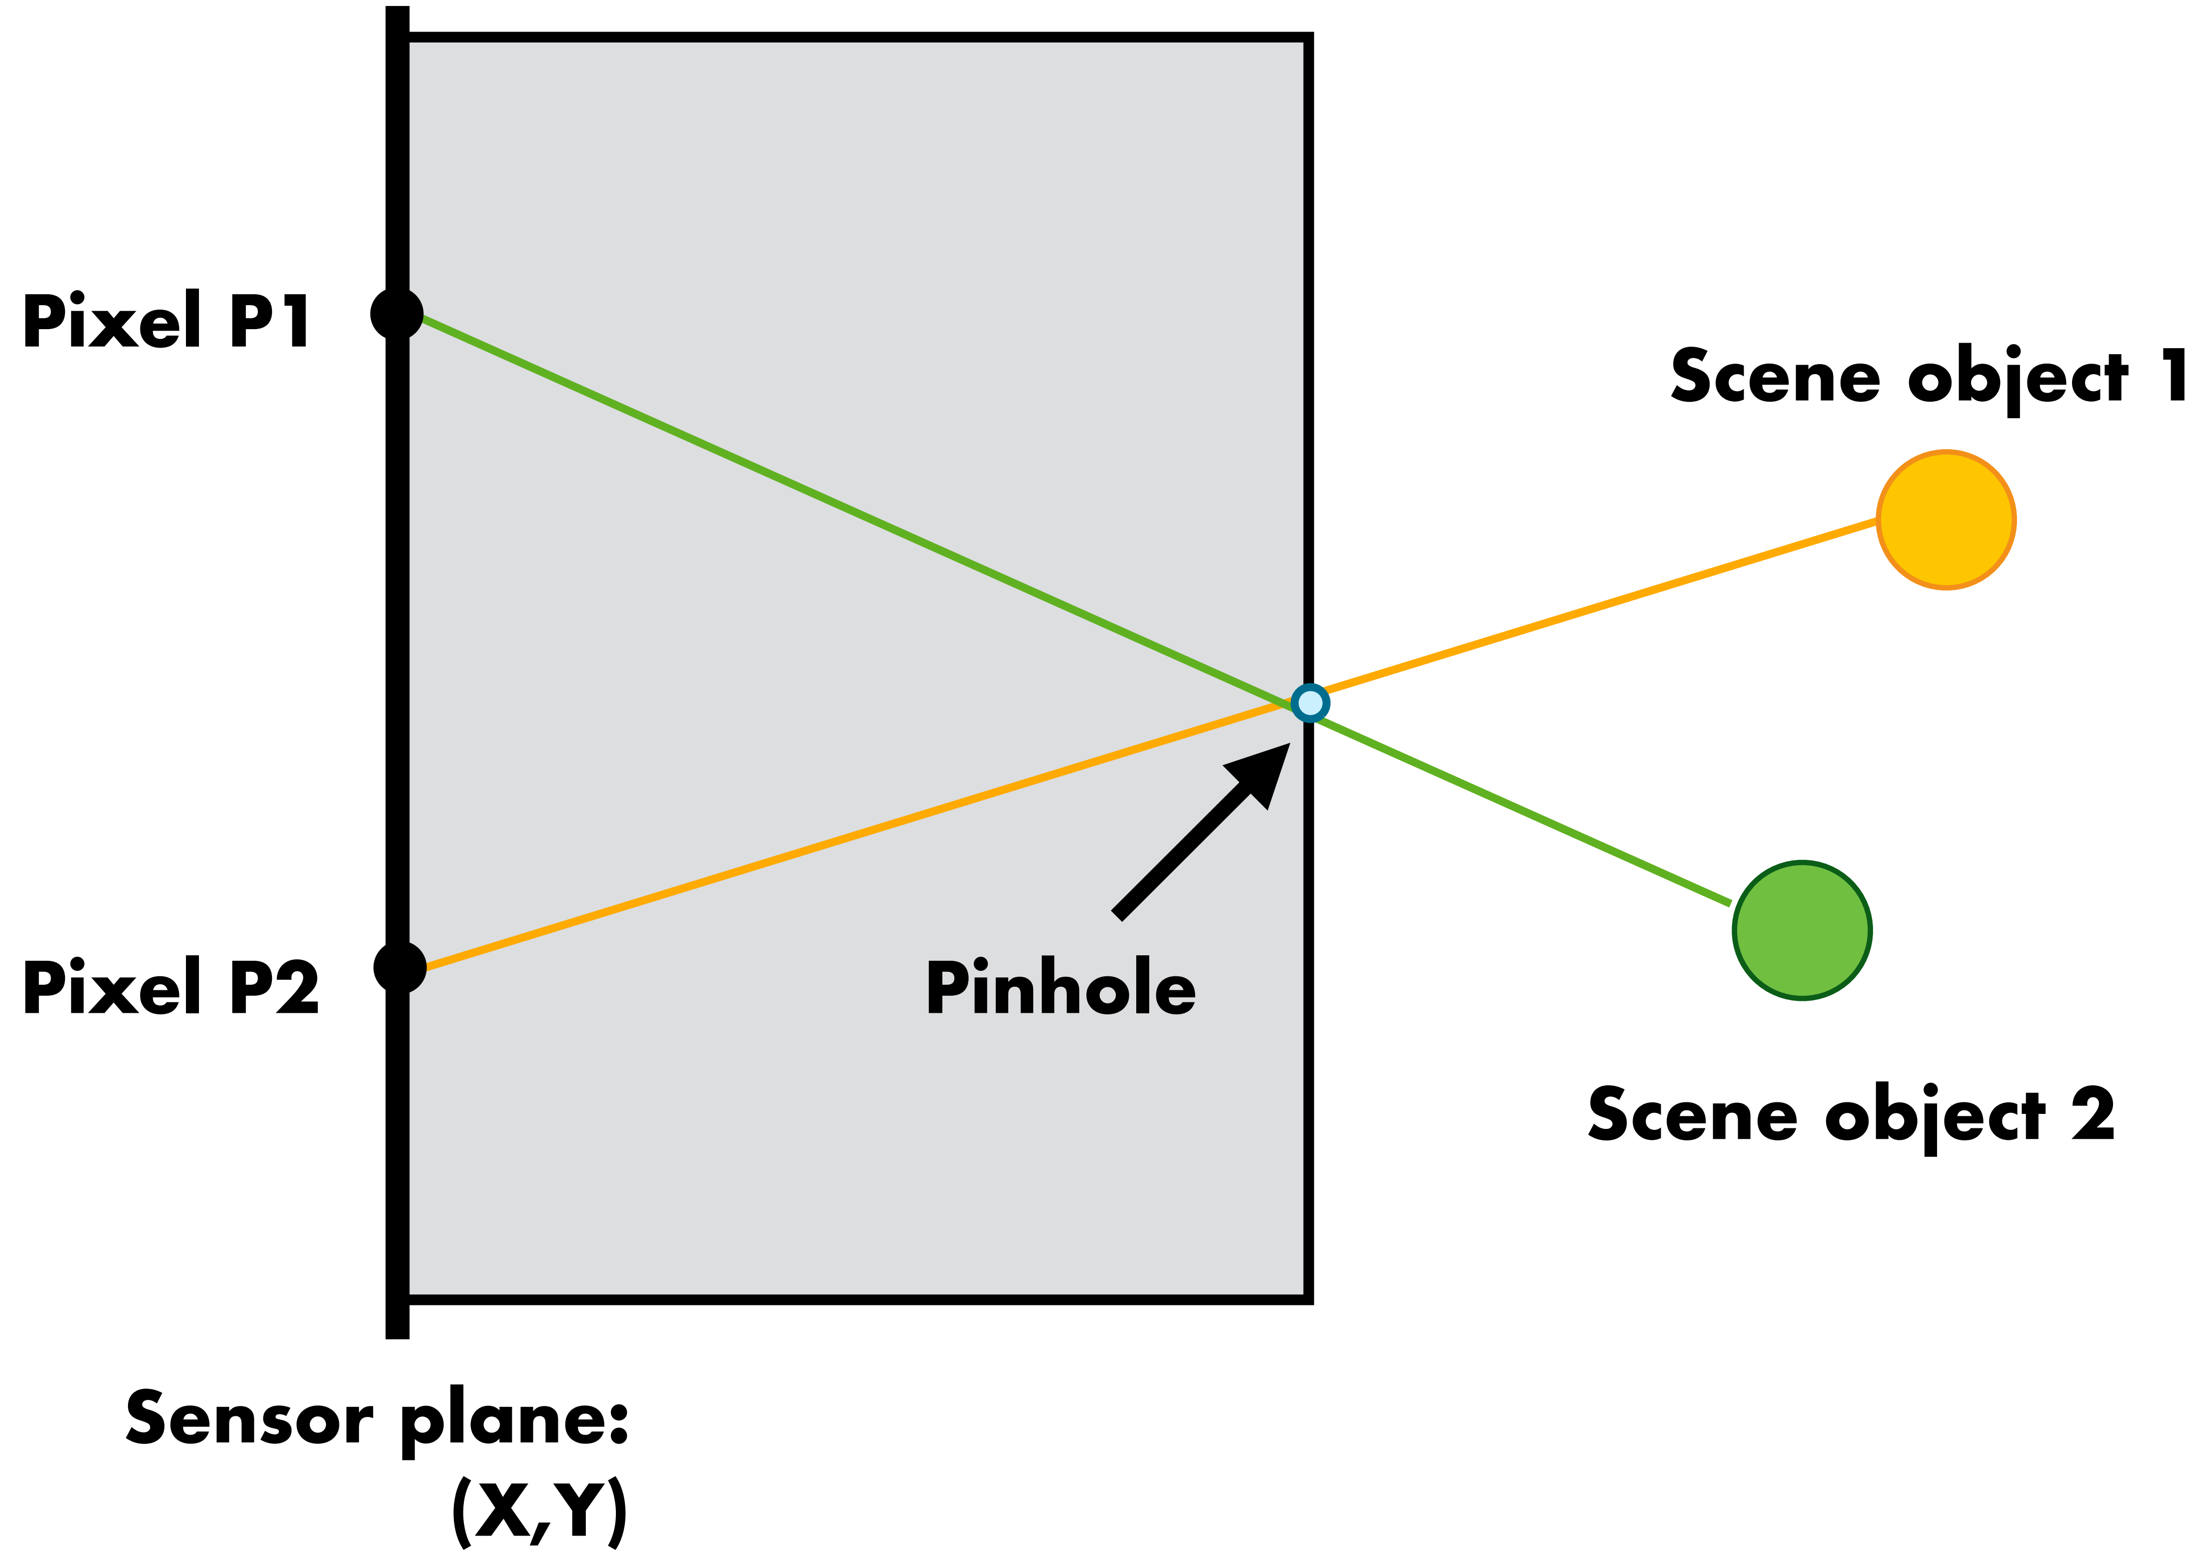

Diagram in 3D...

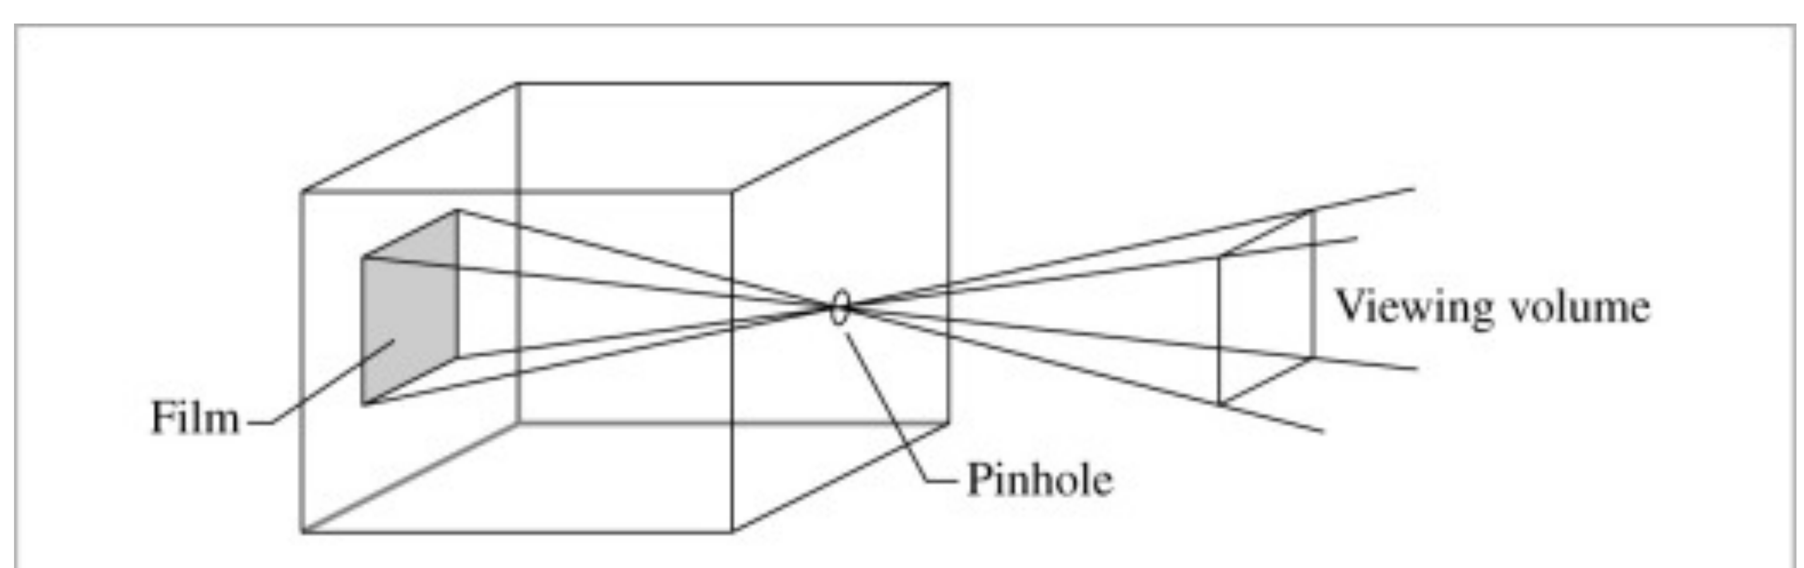

Pinholes & Lenses Form Image on Sensor

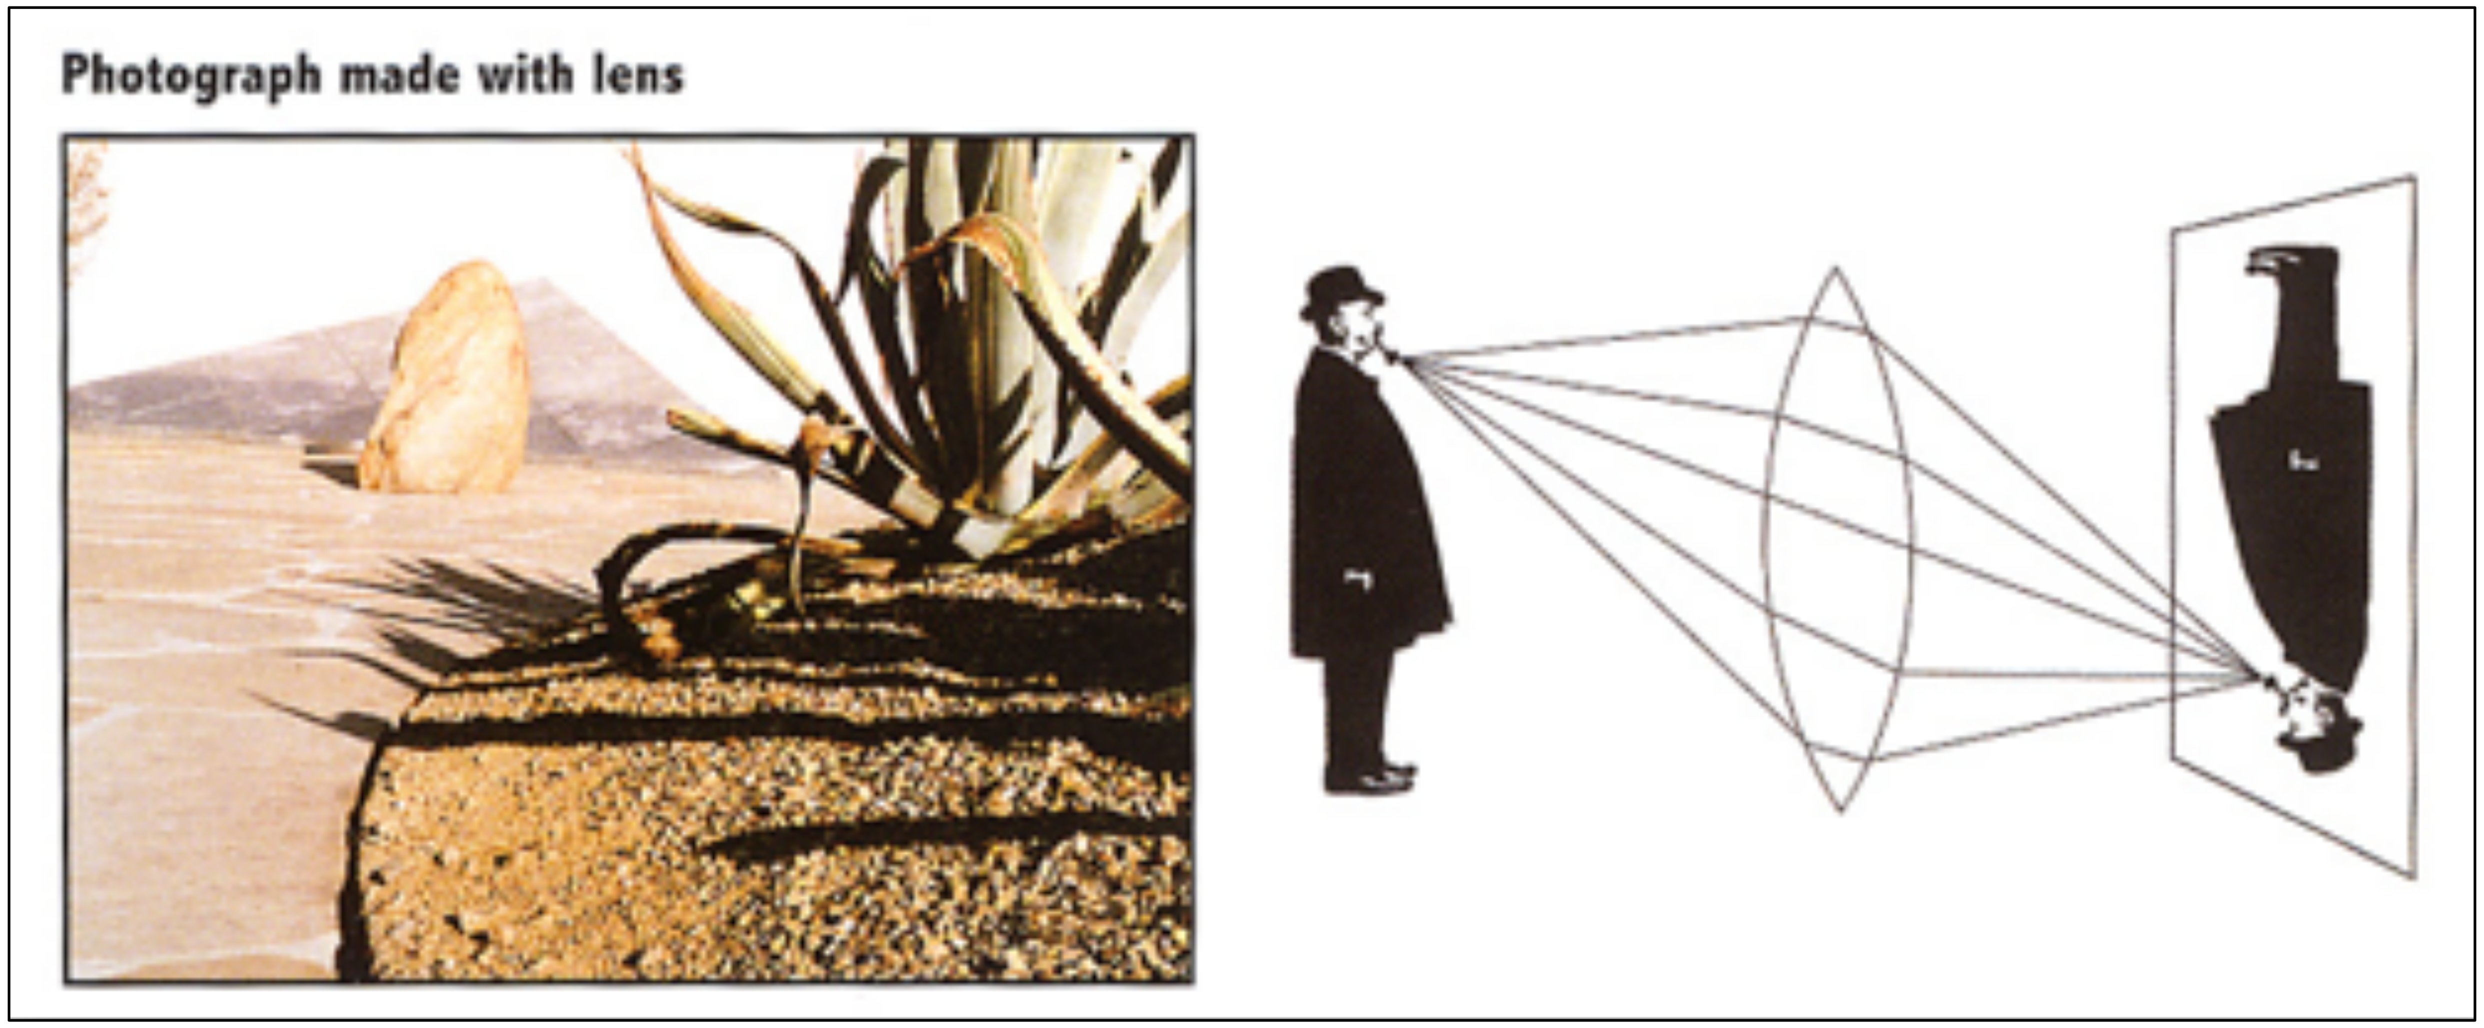

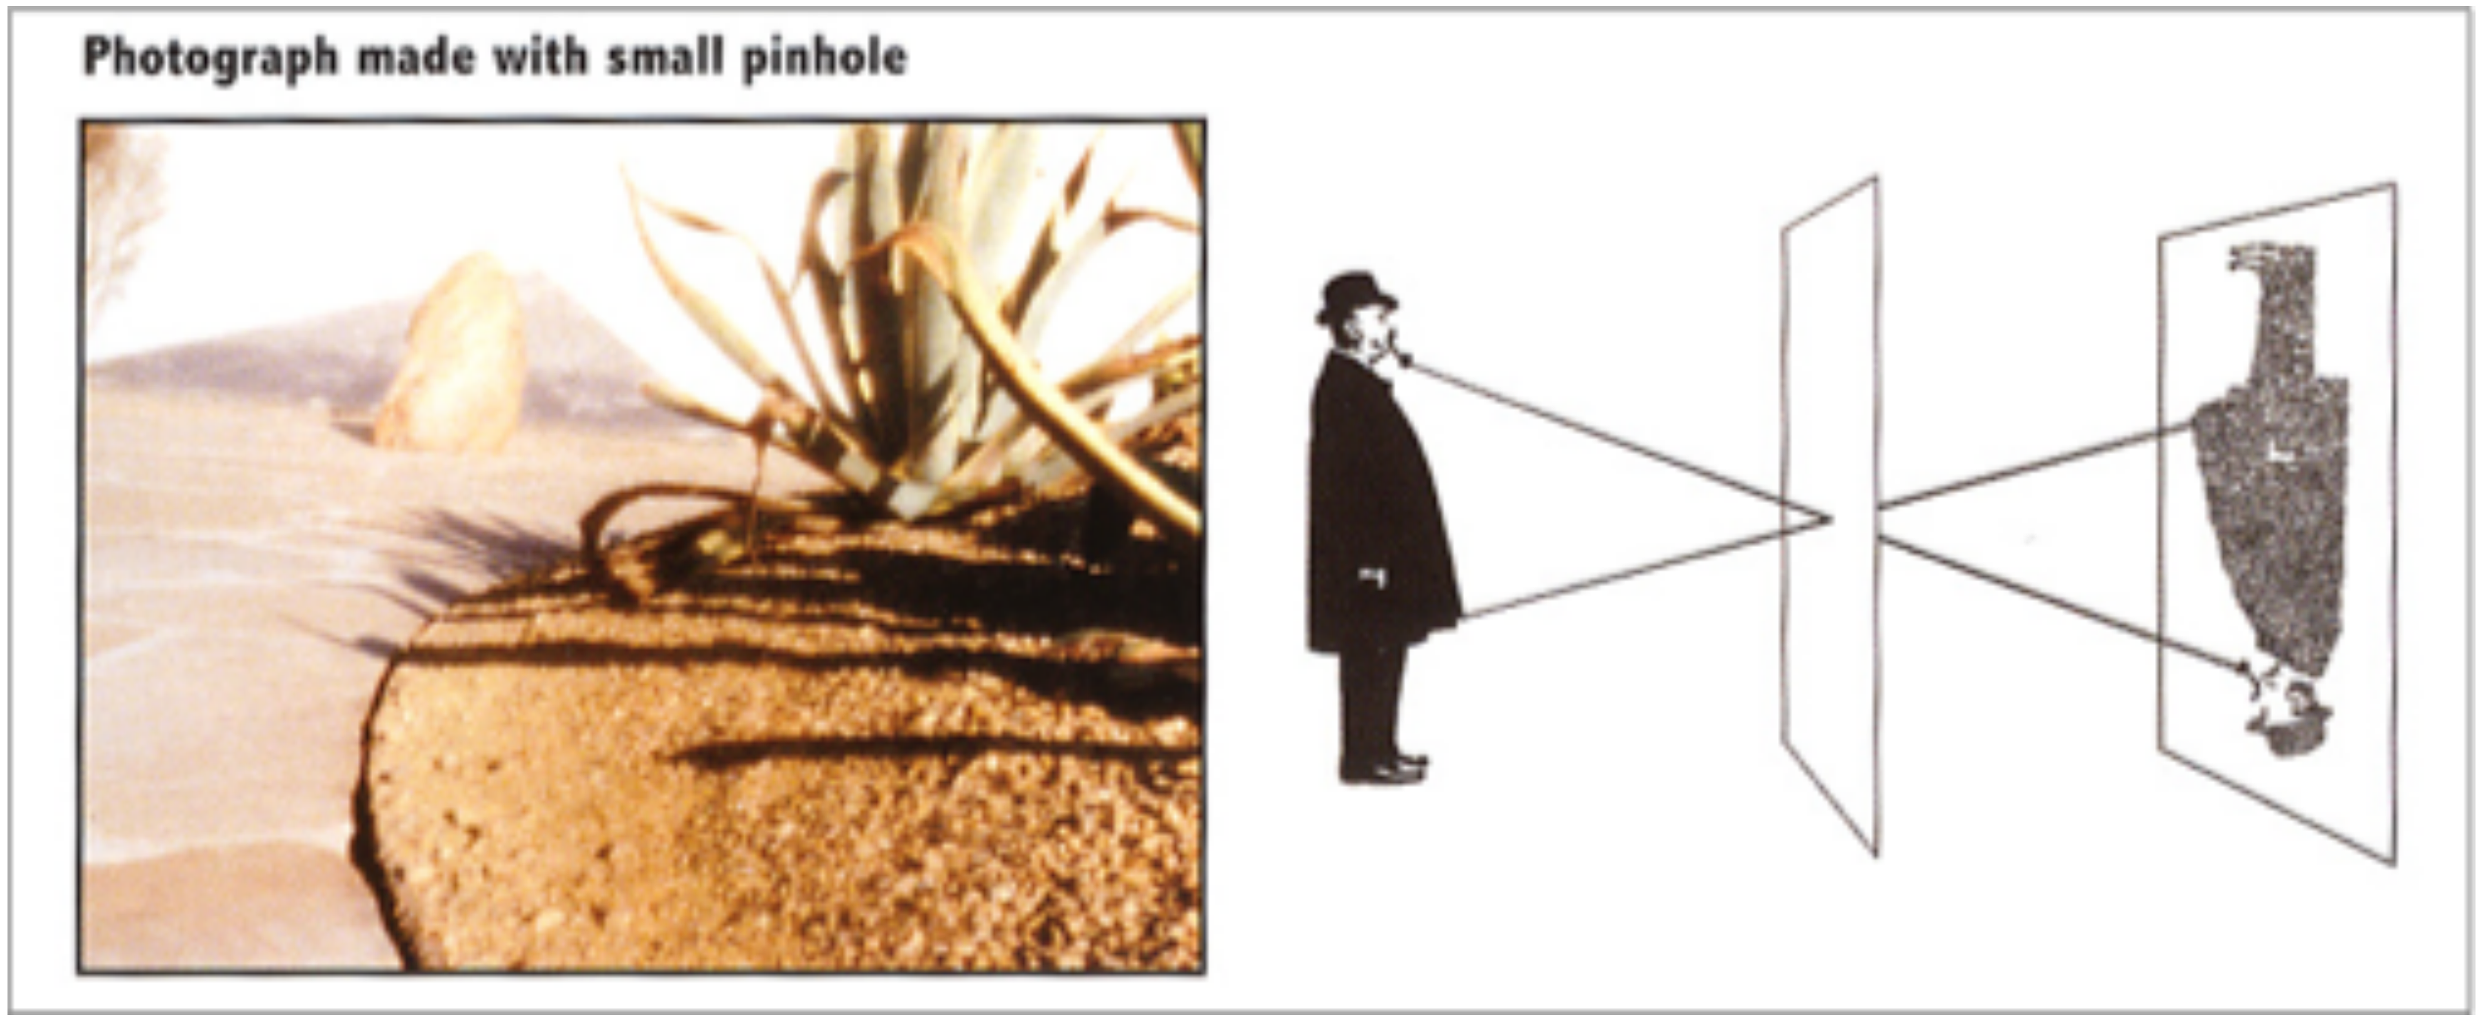

London and Upton

Shutter Exposes Sensor For Precise Duration

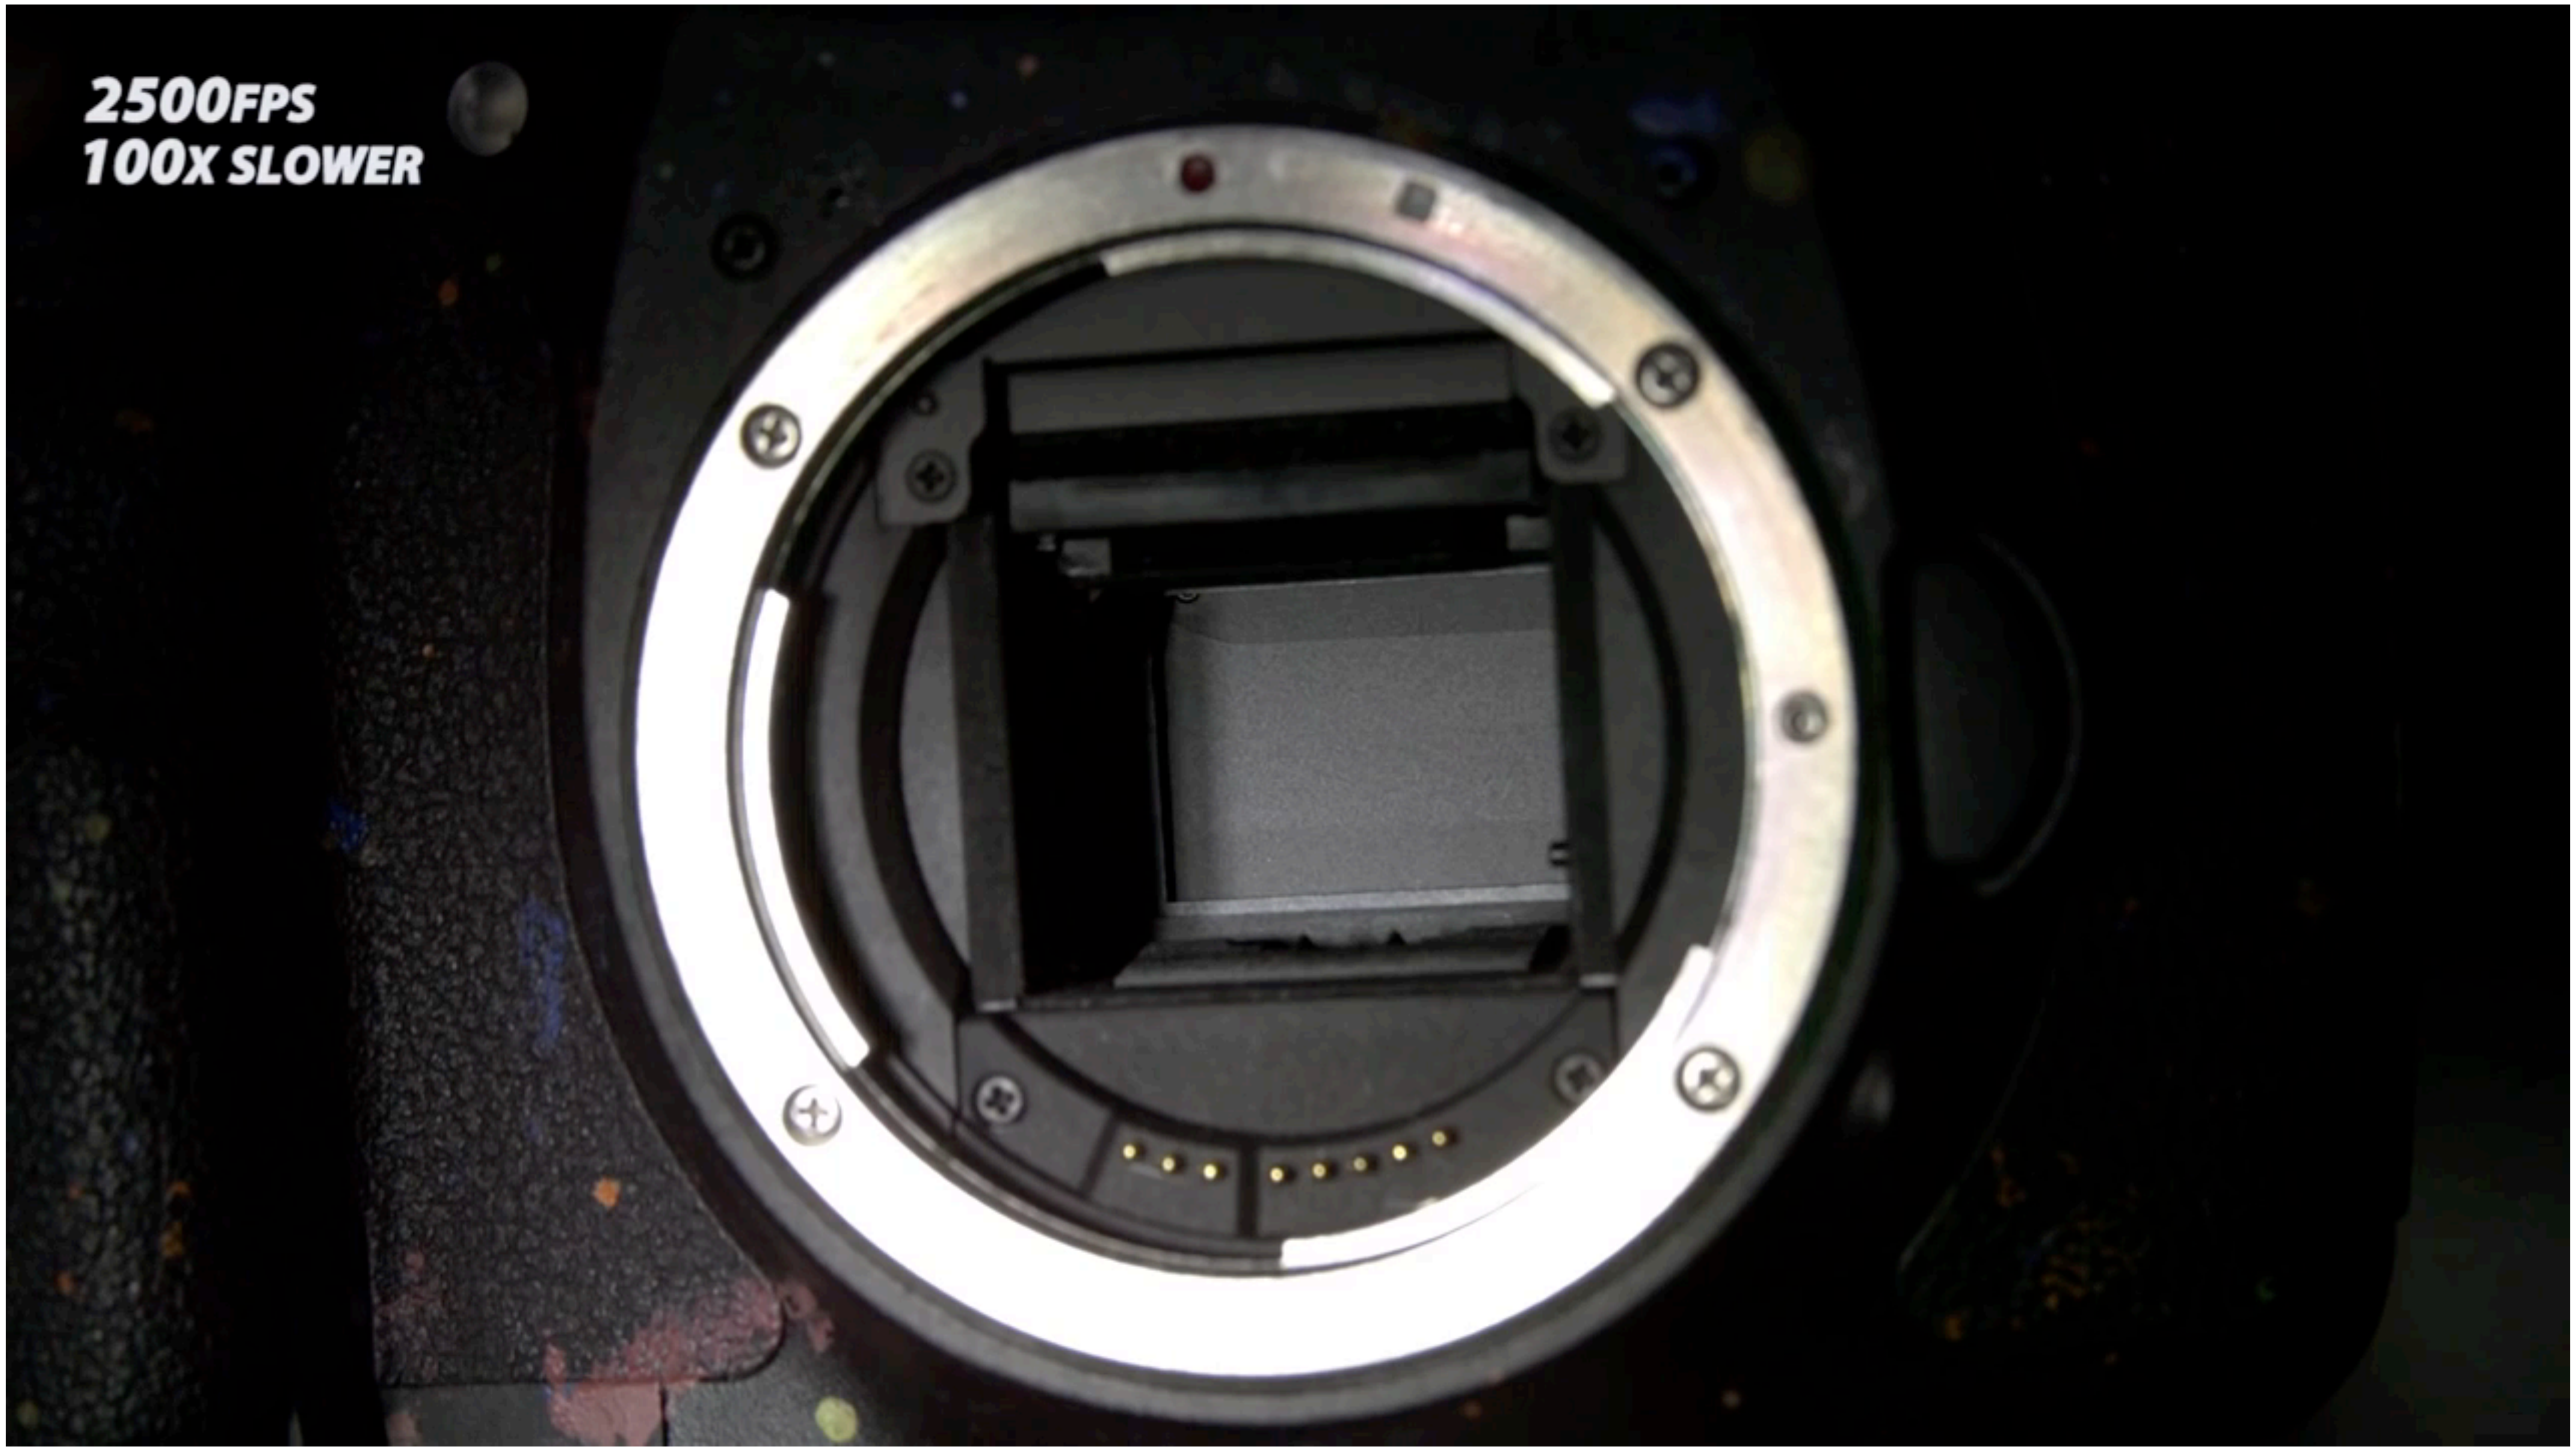

The Slow Mo Guys, https://youtu.be/CmjeCchGRQo

Sensor Accumulates Irradiance During Exposure

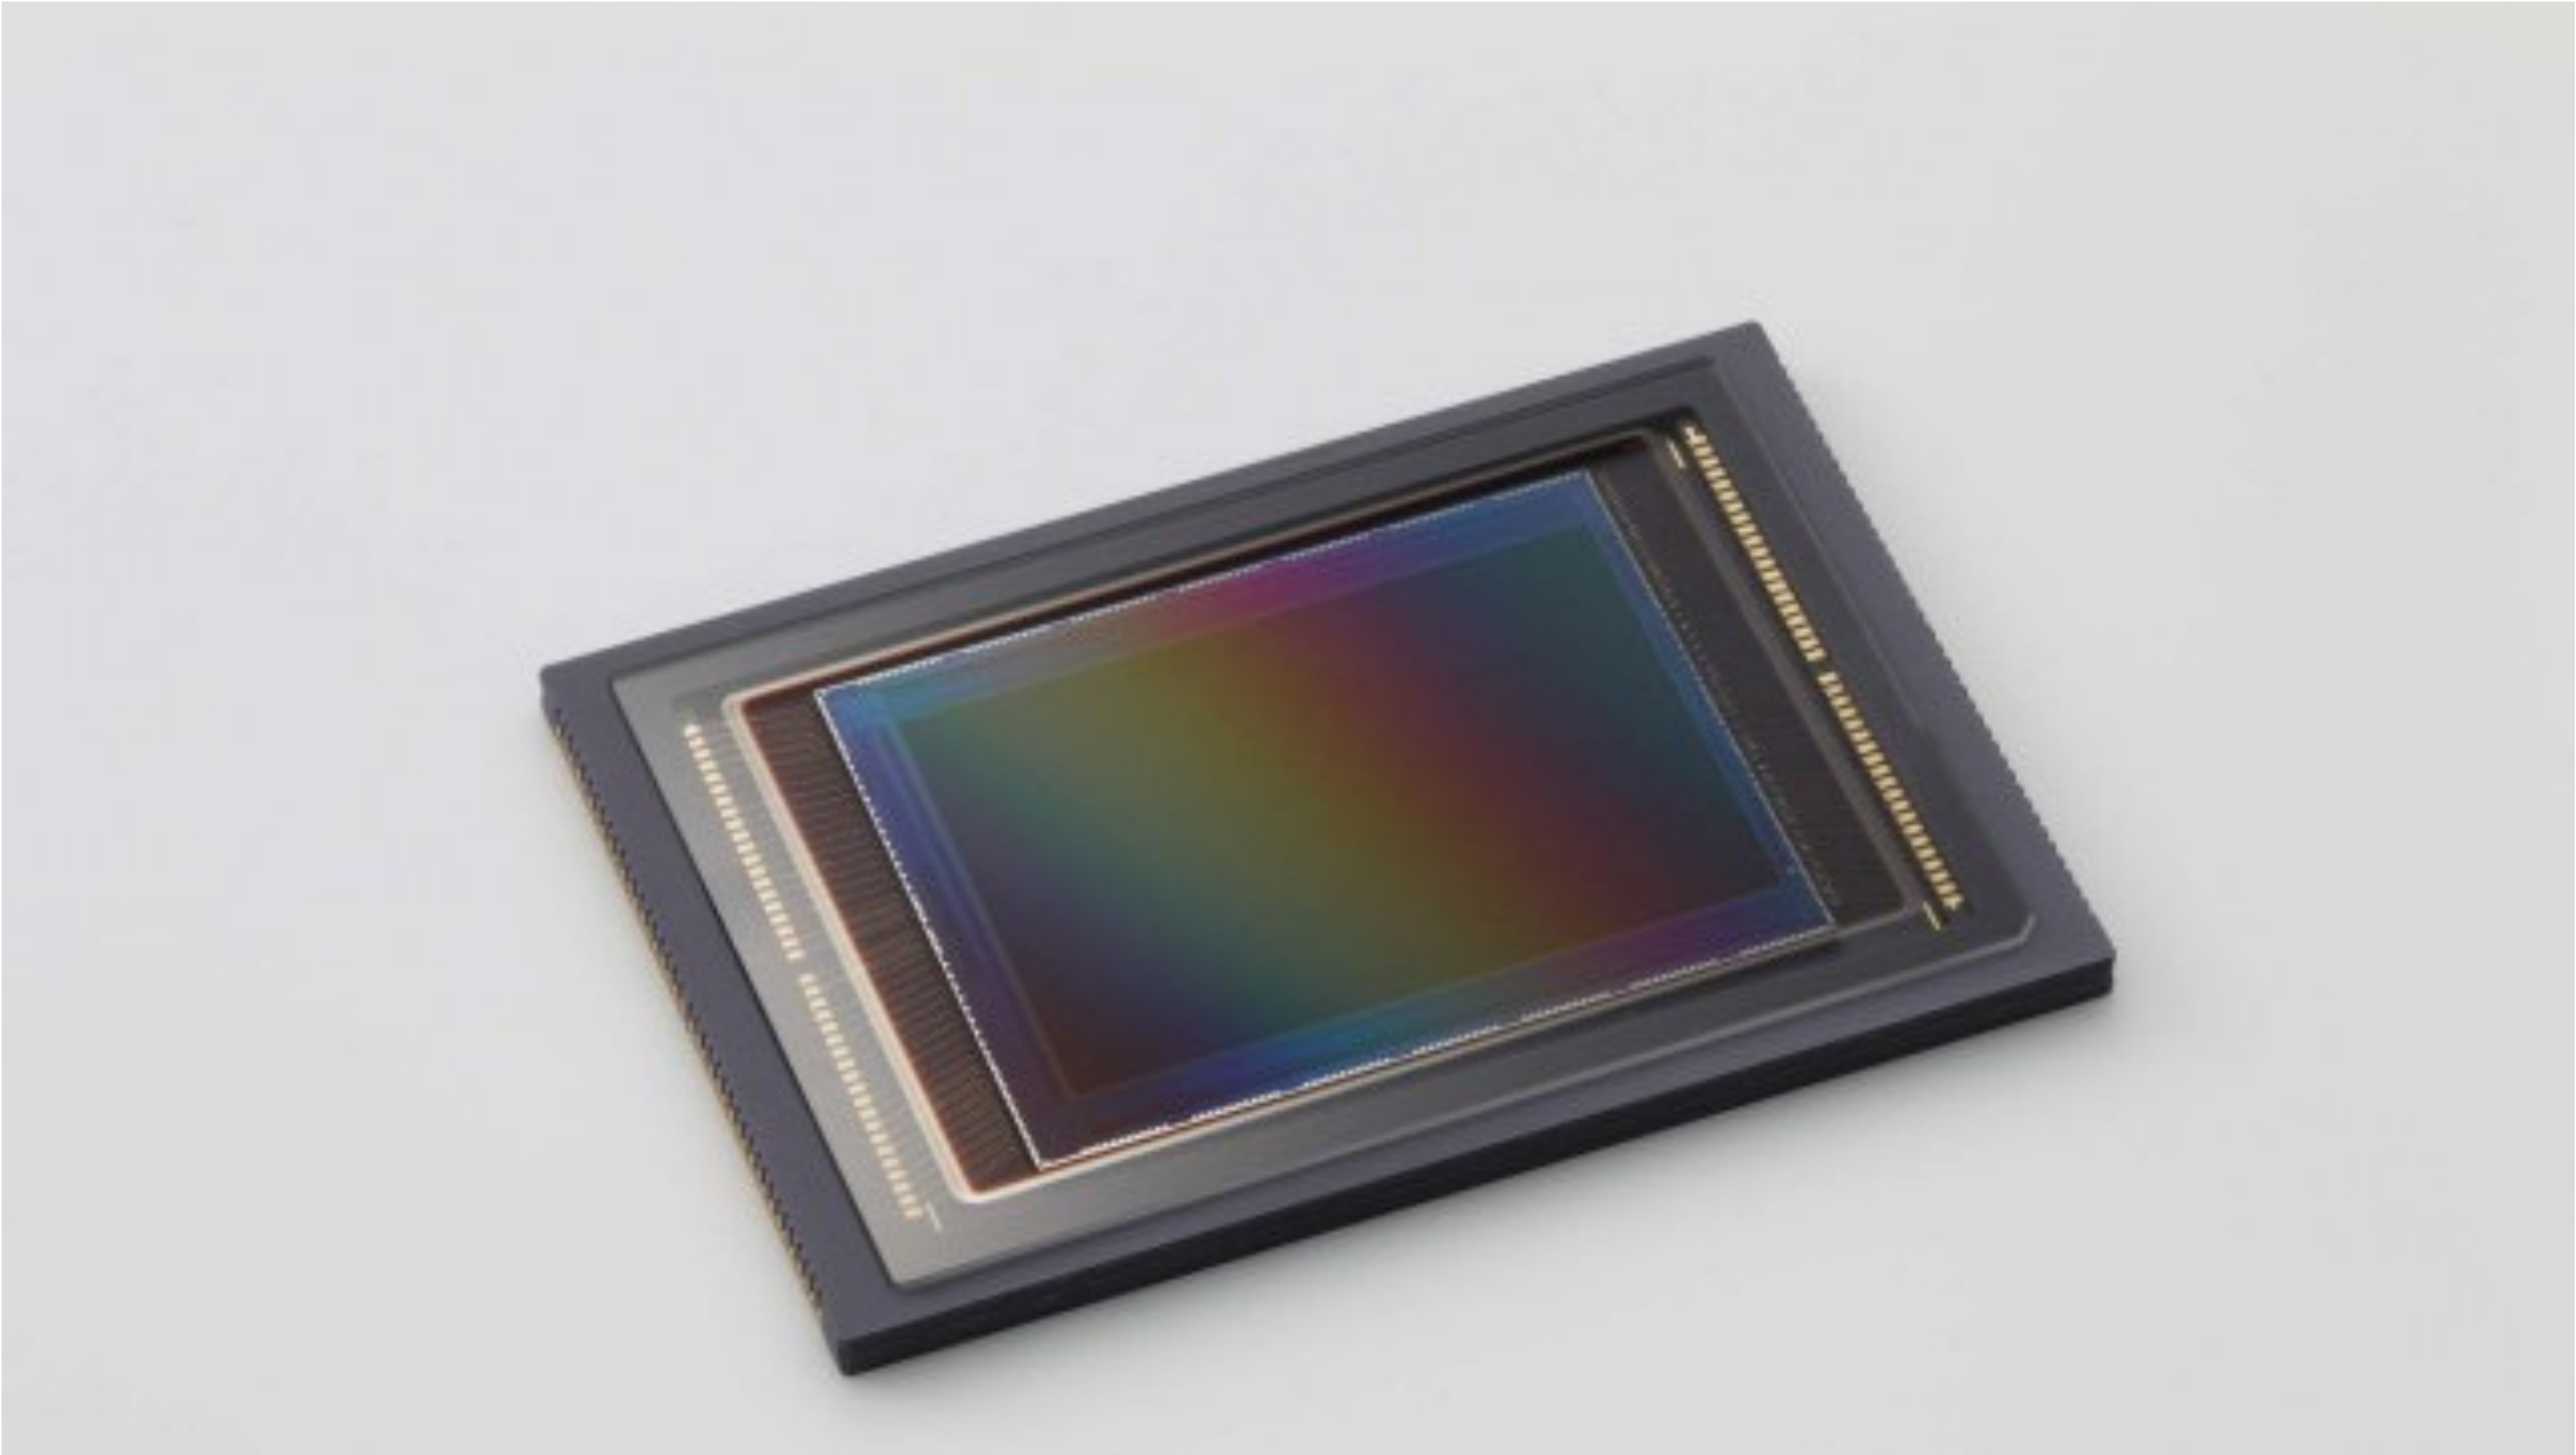

Recall: irradiance = power per unit area
Sensor pixels measure integral of irradiance over time

Lenses

Modern Lens Designs Are Complex

First: Thin Lens Approximation

Ideal Thin Lens – Focal Point

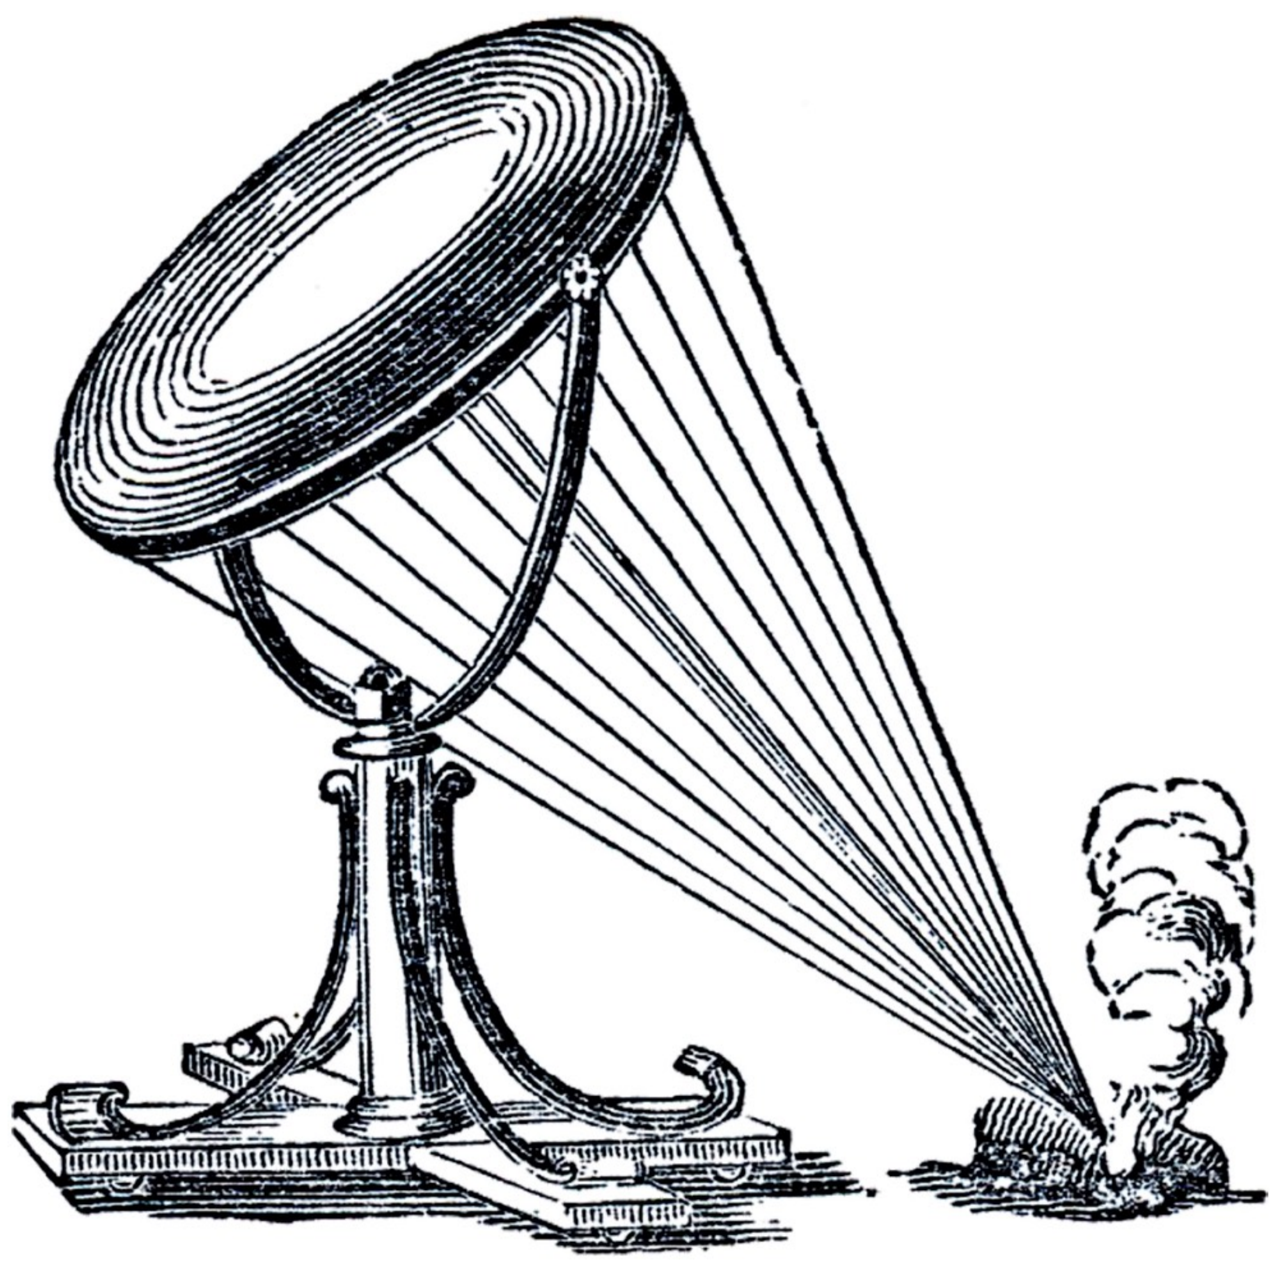
Credit: Karen Watson

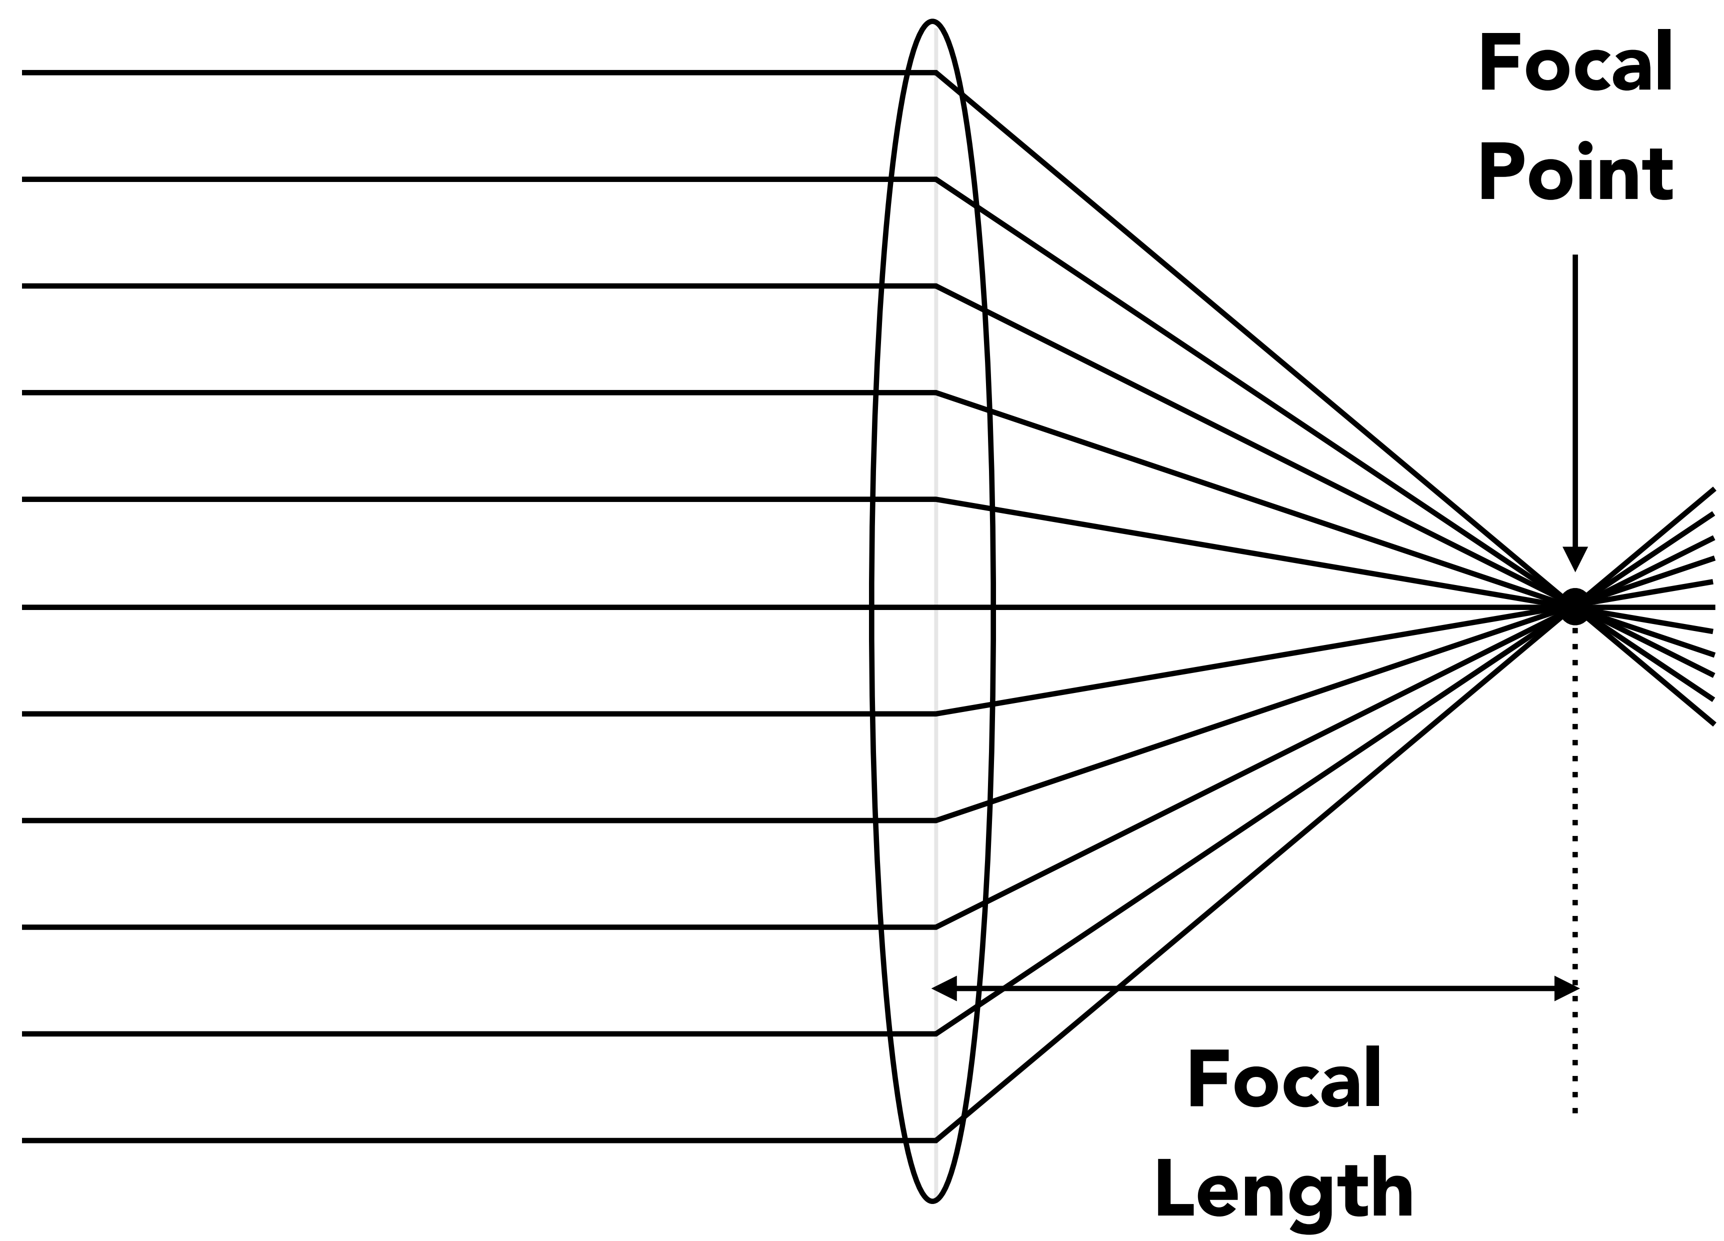

Assume all parallel rays entering a lens pass through its focal point.

Keep in mind: Real Lens Elements Are Not Ideal – Aberrations

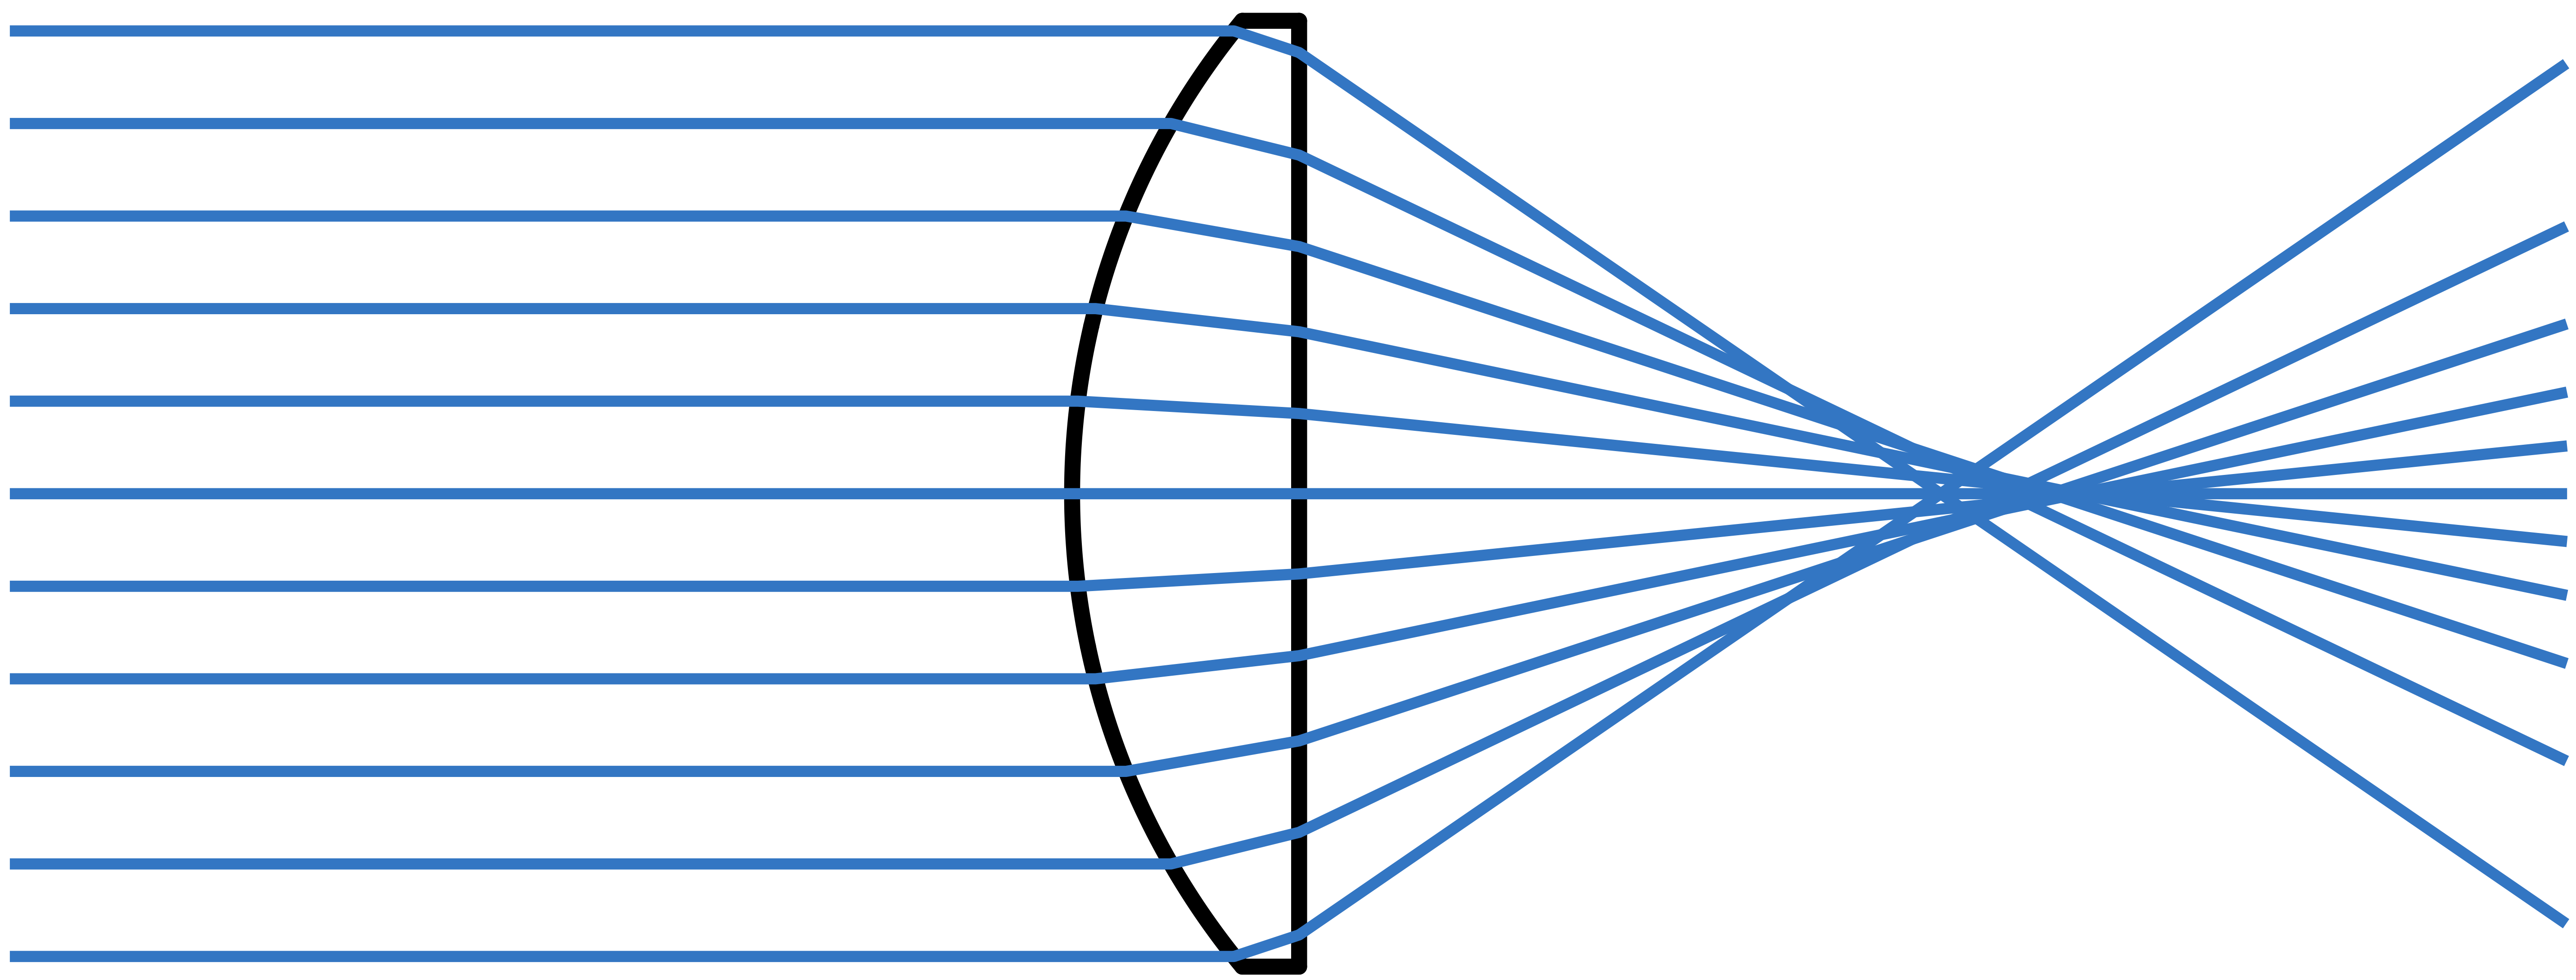

Real plano-convex lens (spherical surface shape).
Lens does not converge rays to a point anywhere.

Gauss’ Ray Diagrams

Gauss' Ray Tracing Construction

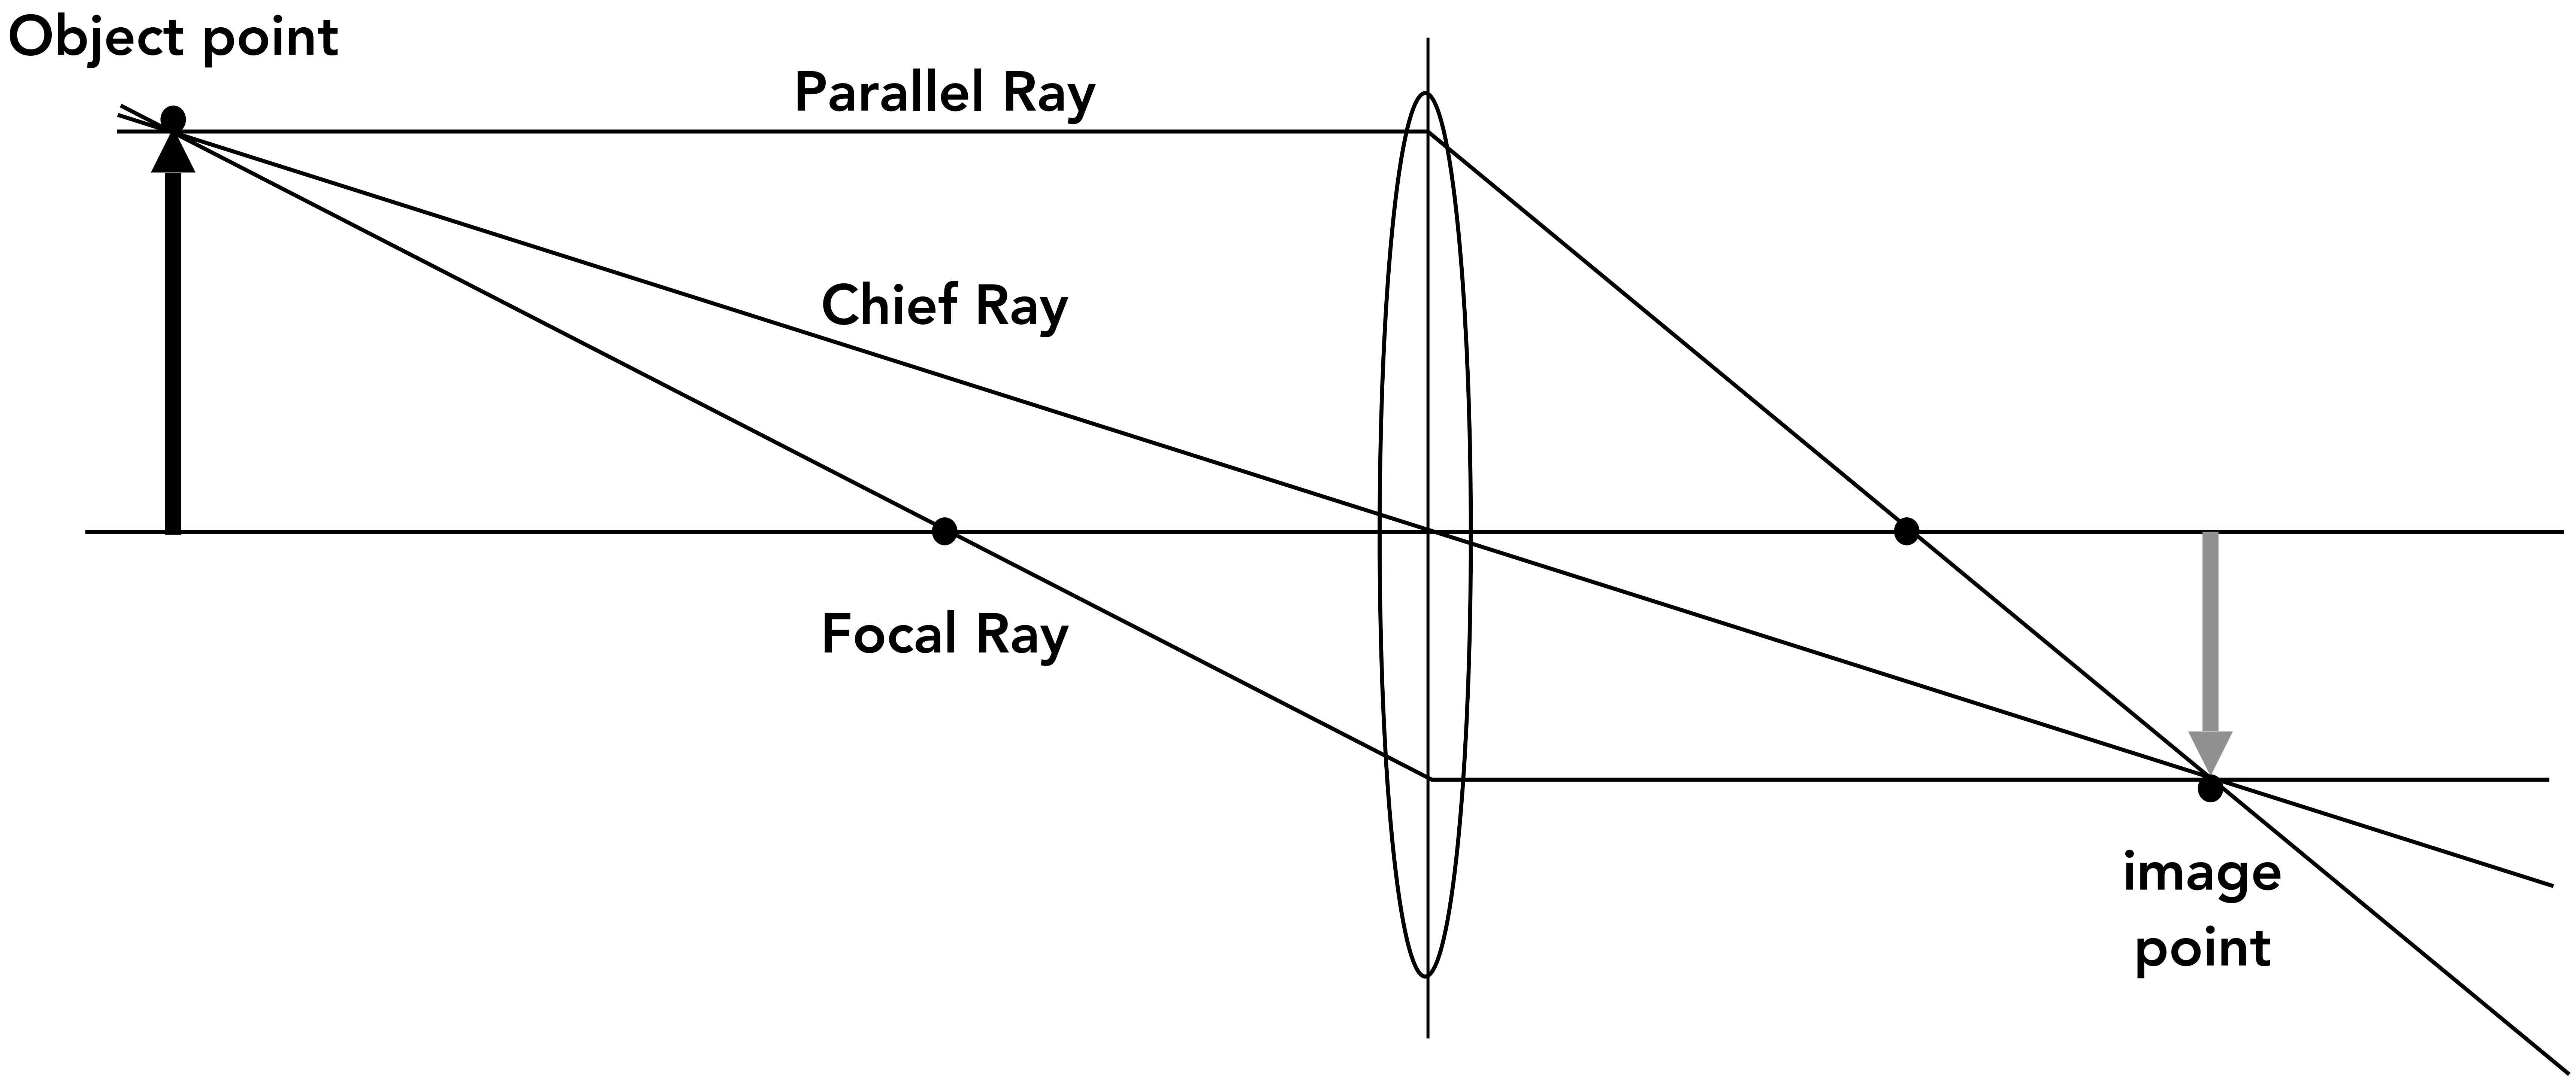

Object
(Outside the camera)

Image
("Inside the camera")

Lenses focus light!

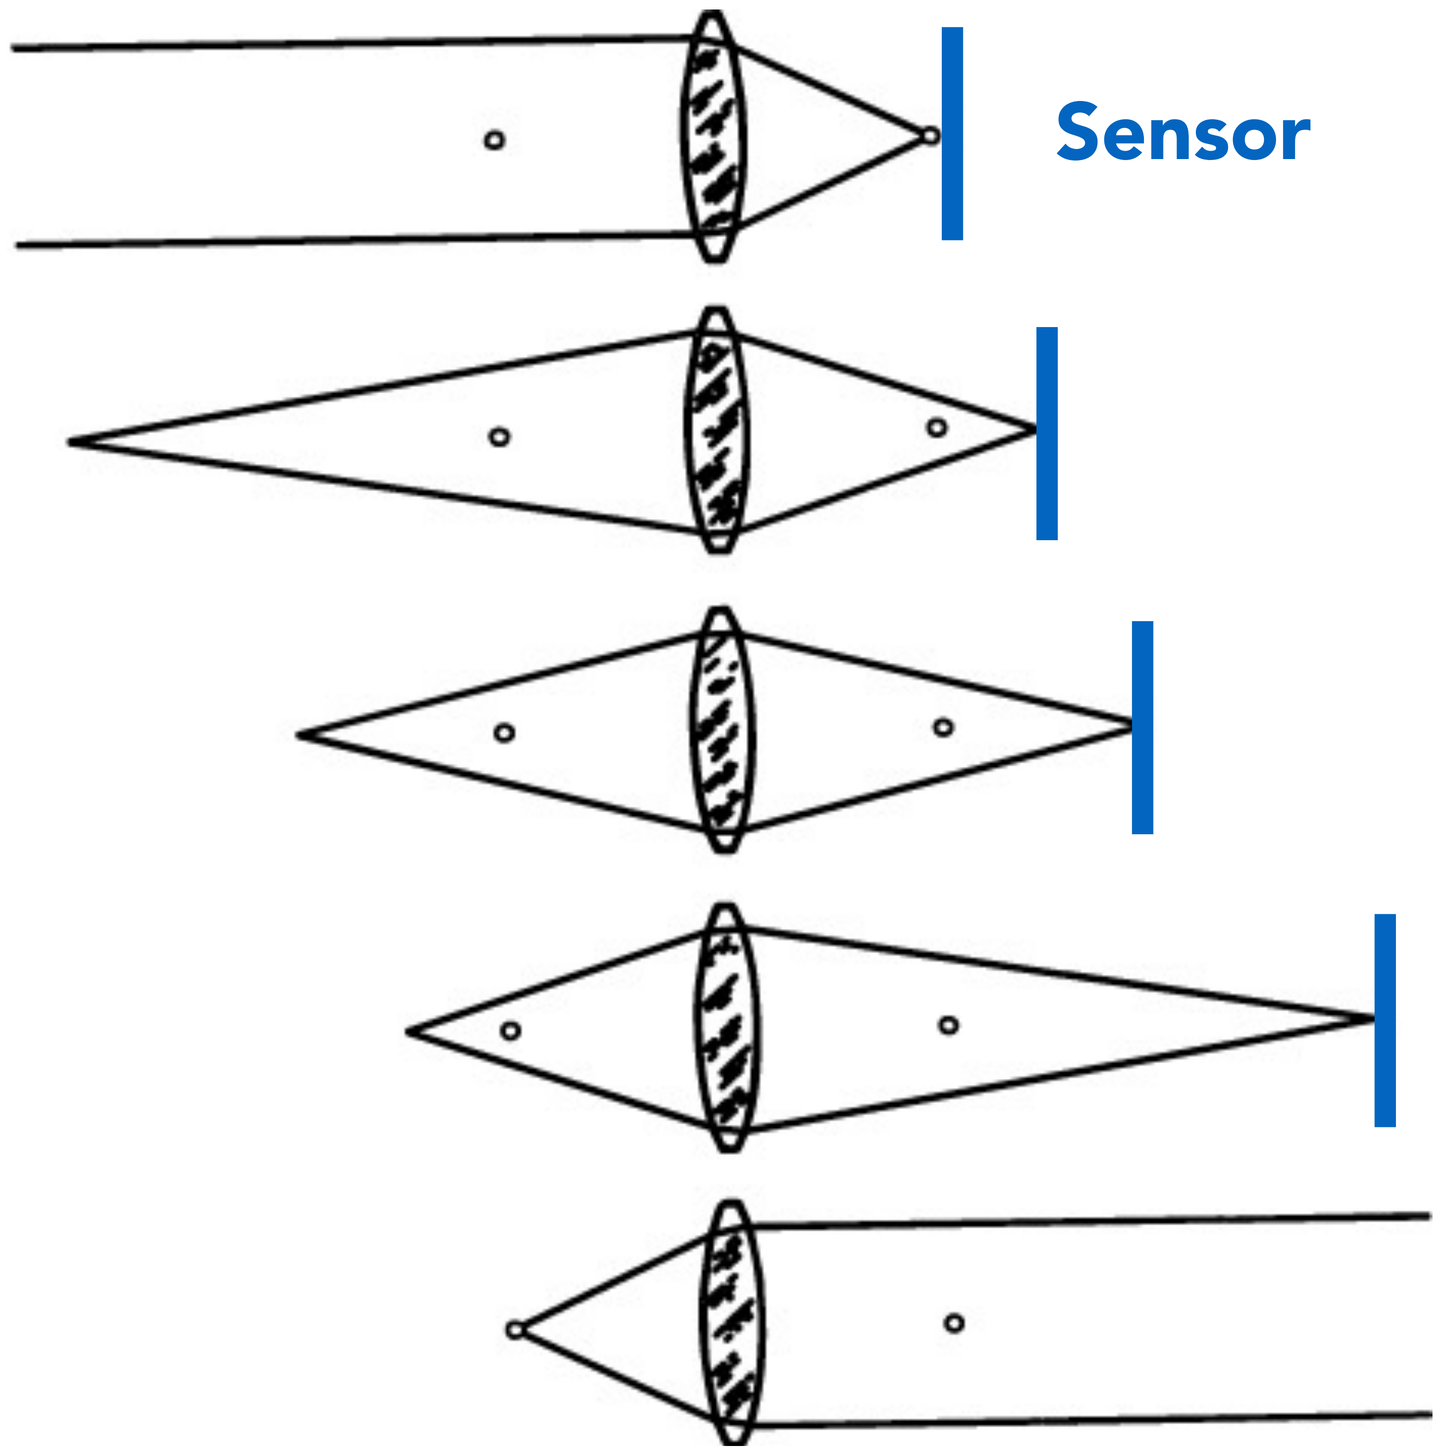

Rays from a point in object space intersect at a point in image space

These are conjugate points

Points at infinity converge at the focal point or focus

Points on different object planes focus on different image planes

The position of the point of focus can be changed by changing the distance between the lens and the sensor

Gauss' Ray Tracing Construction

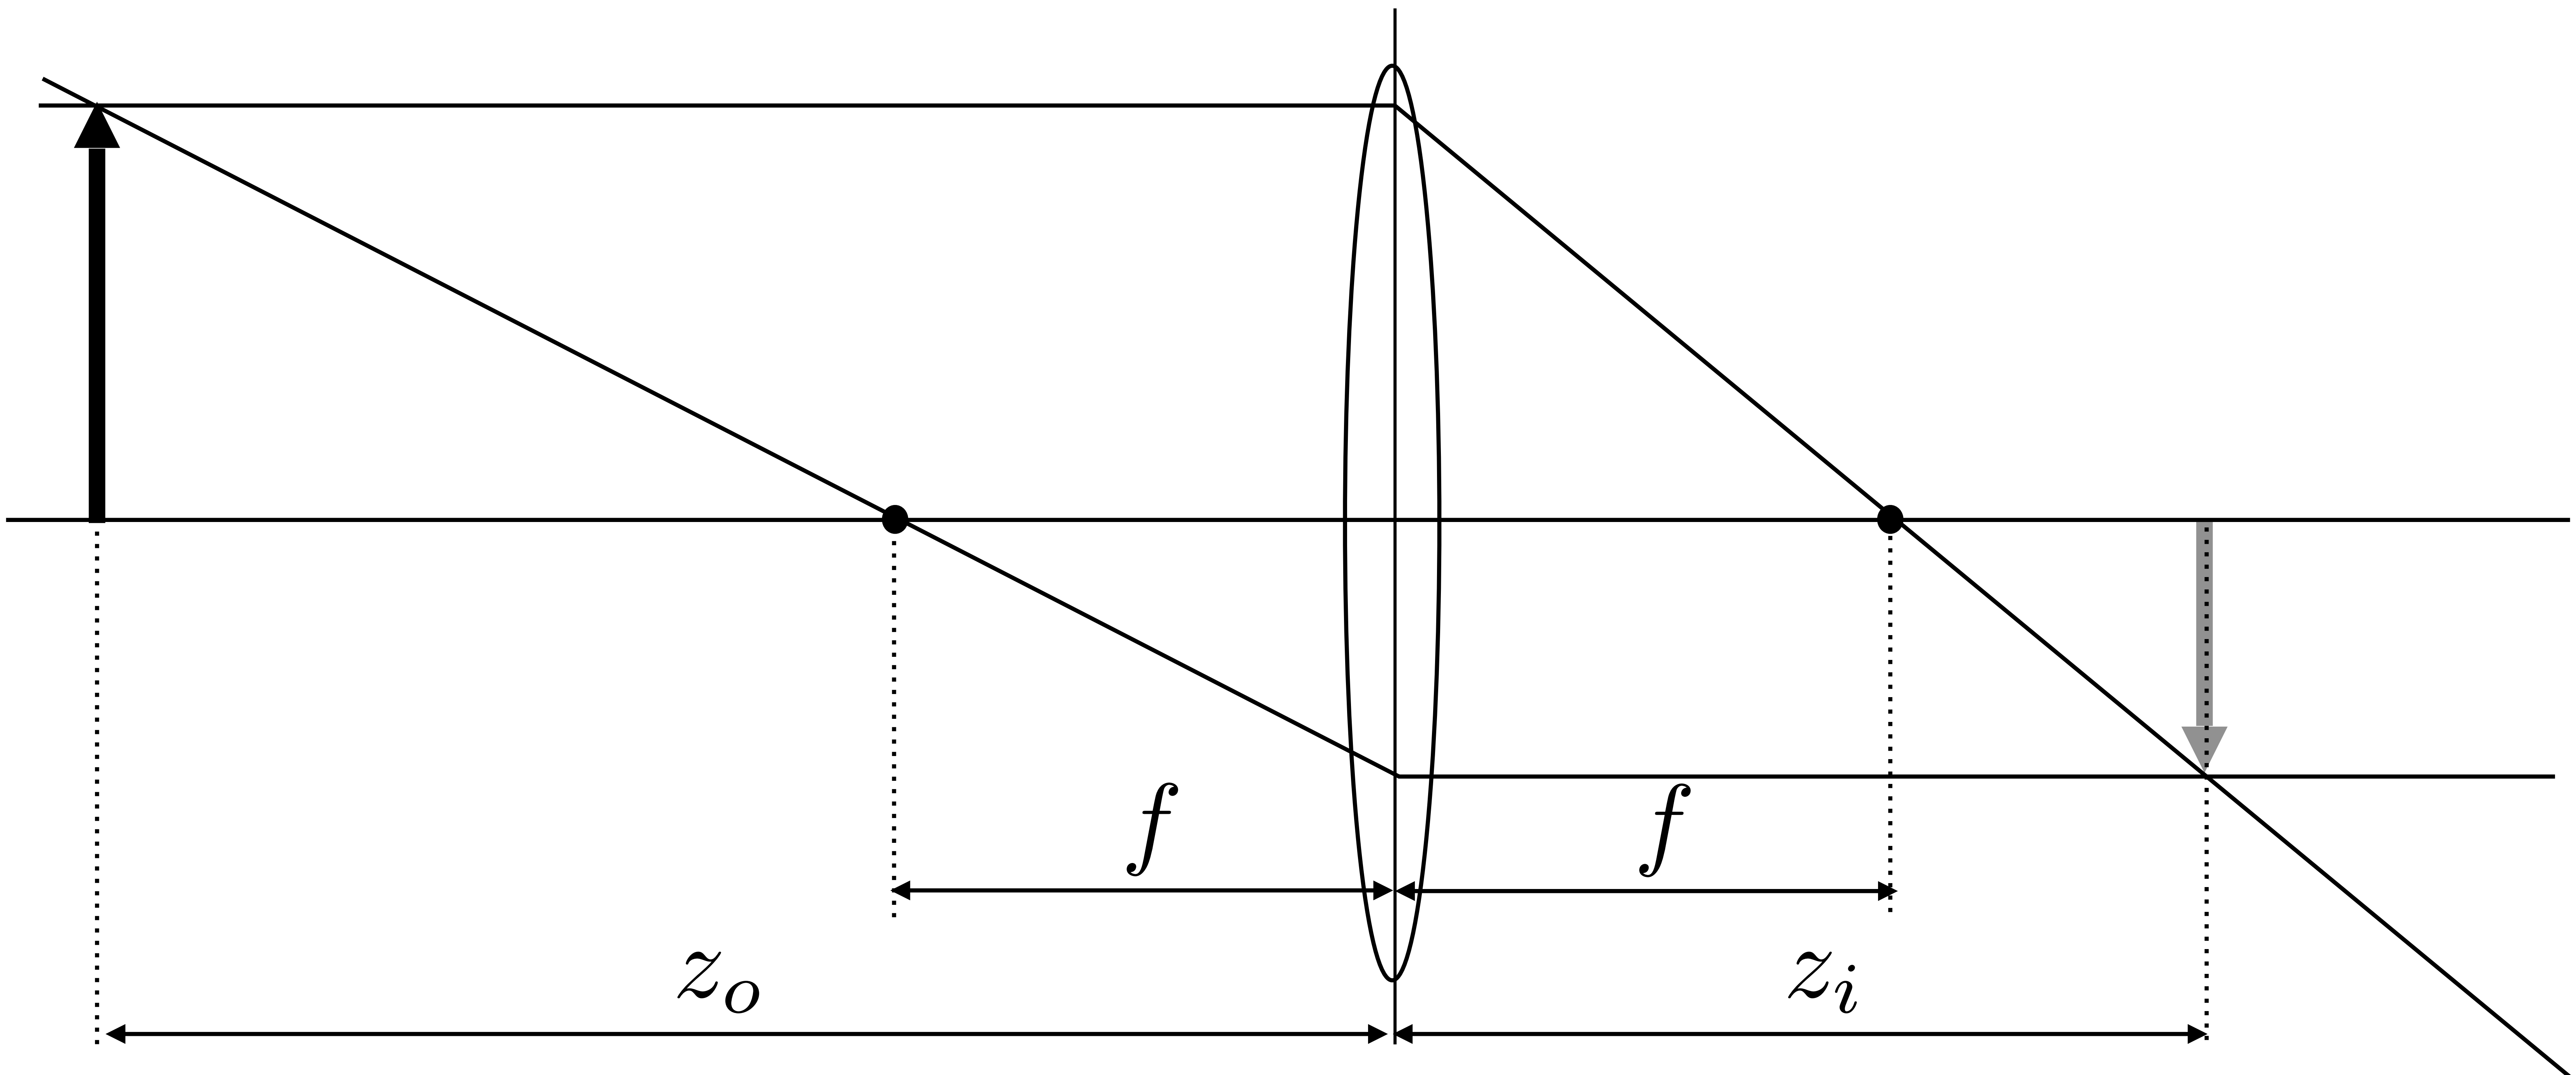

What is the relationship between conjugate depths z_o, z_i?

Gauss' Ray Tracing Construction

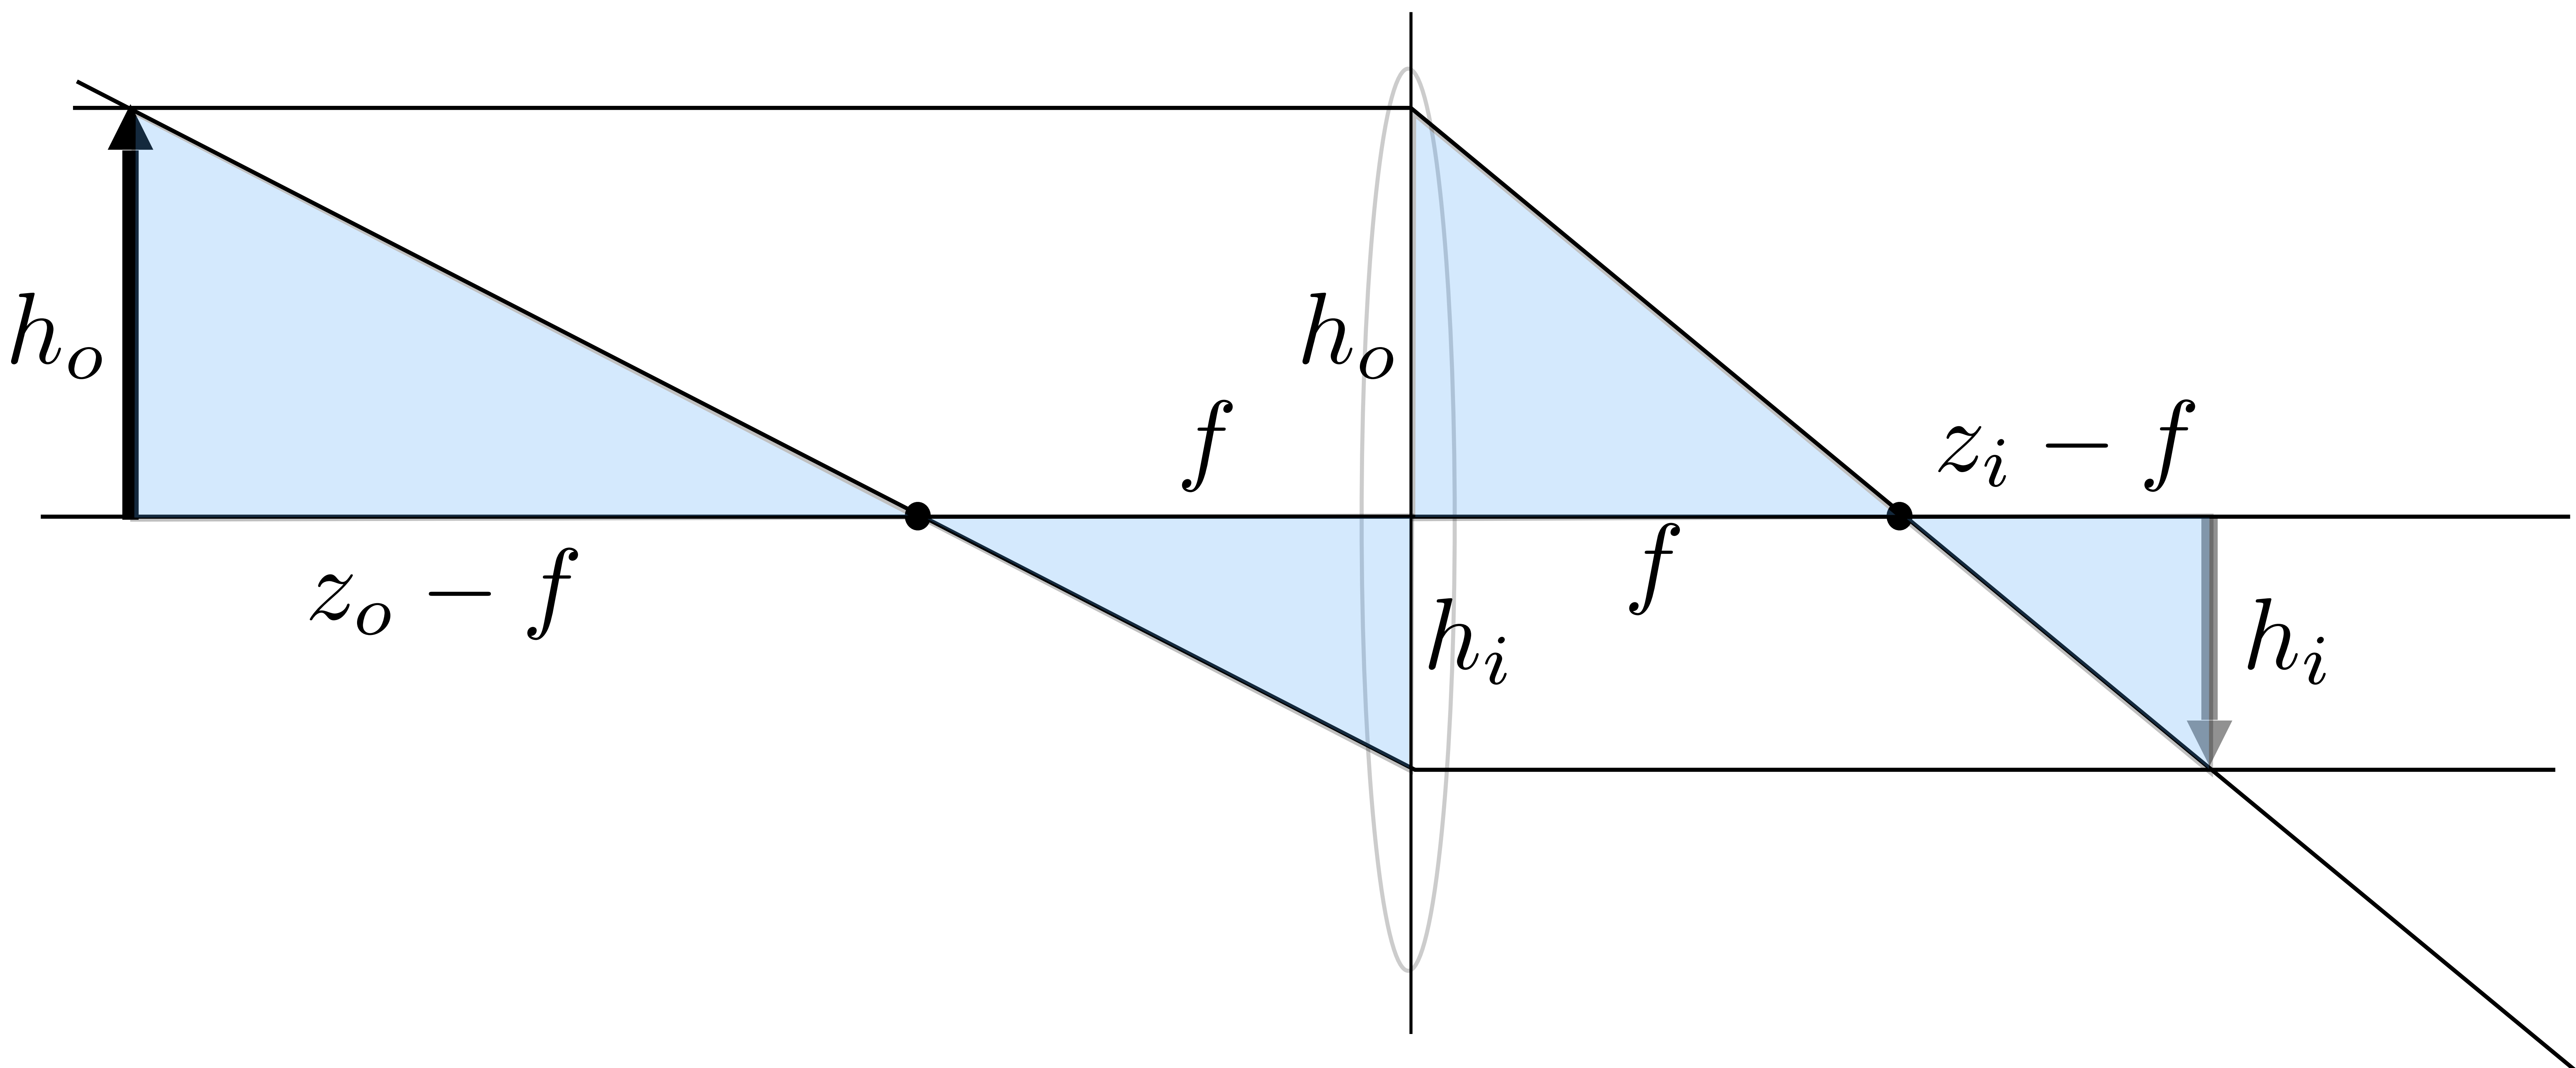

$$\frac{h_o}{z_o - f} = \frac{h_i}{f} \qquad \frac{h_o}{f} = \frac{h_i}{z_i - f}$$

Gauss' Ray Tracing Construction

$$\frac{h_o}{z_o - f} = \frac{h_i}{f} \qquad \frac{h_o}{f} = \frac{h_i}{z_i - f}$$

$$\frac{h_o}{h_i} = \frac{z_o - f}{f} \qquad \frac{h_o}{h_i} = \frac{f}{z_i - f}$$

$$\frac{z_o - f}{f} = \frac{f}{z_i - f}$$

Object / image heights factor out - applies to all rays

$$(z_o - f)(z_i - f) = f^2$$

Newtonian Thin Lens Equation

$$z_o z_i - (z_o + z_i) f + f^2 = f^2$$

$$z_o z_i = (z_o + z_i) f$$

$$\frac{1}{f} = \frac{1}{z_i} + \frac{1}{z_o}$$

Gaussian Thin Lens Equation

The Thin Lens Equation

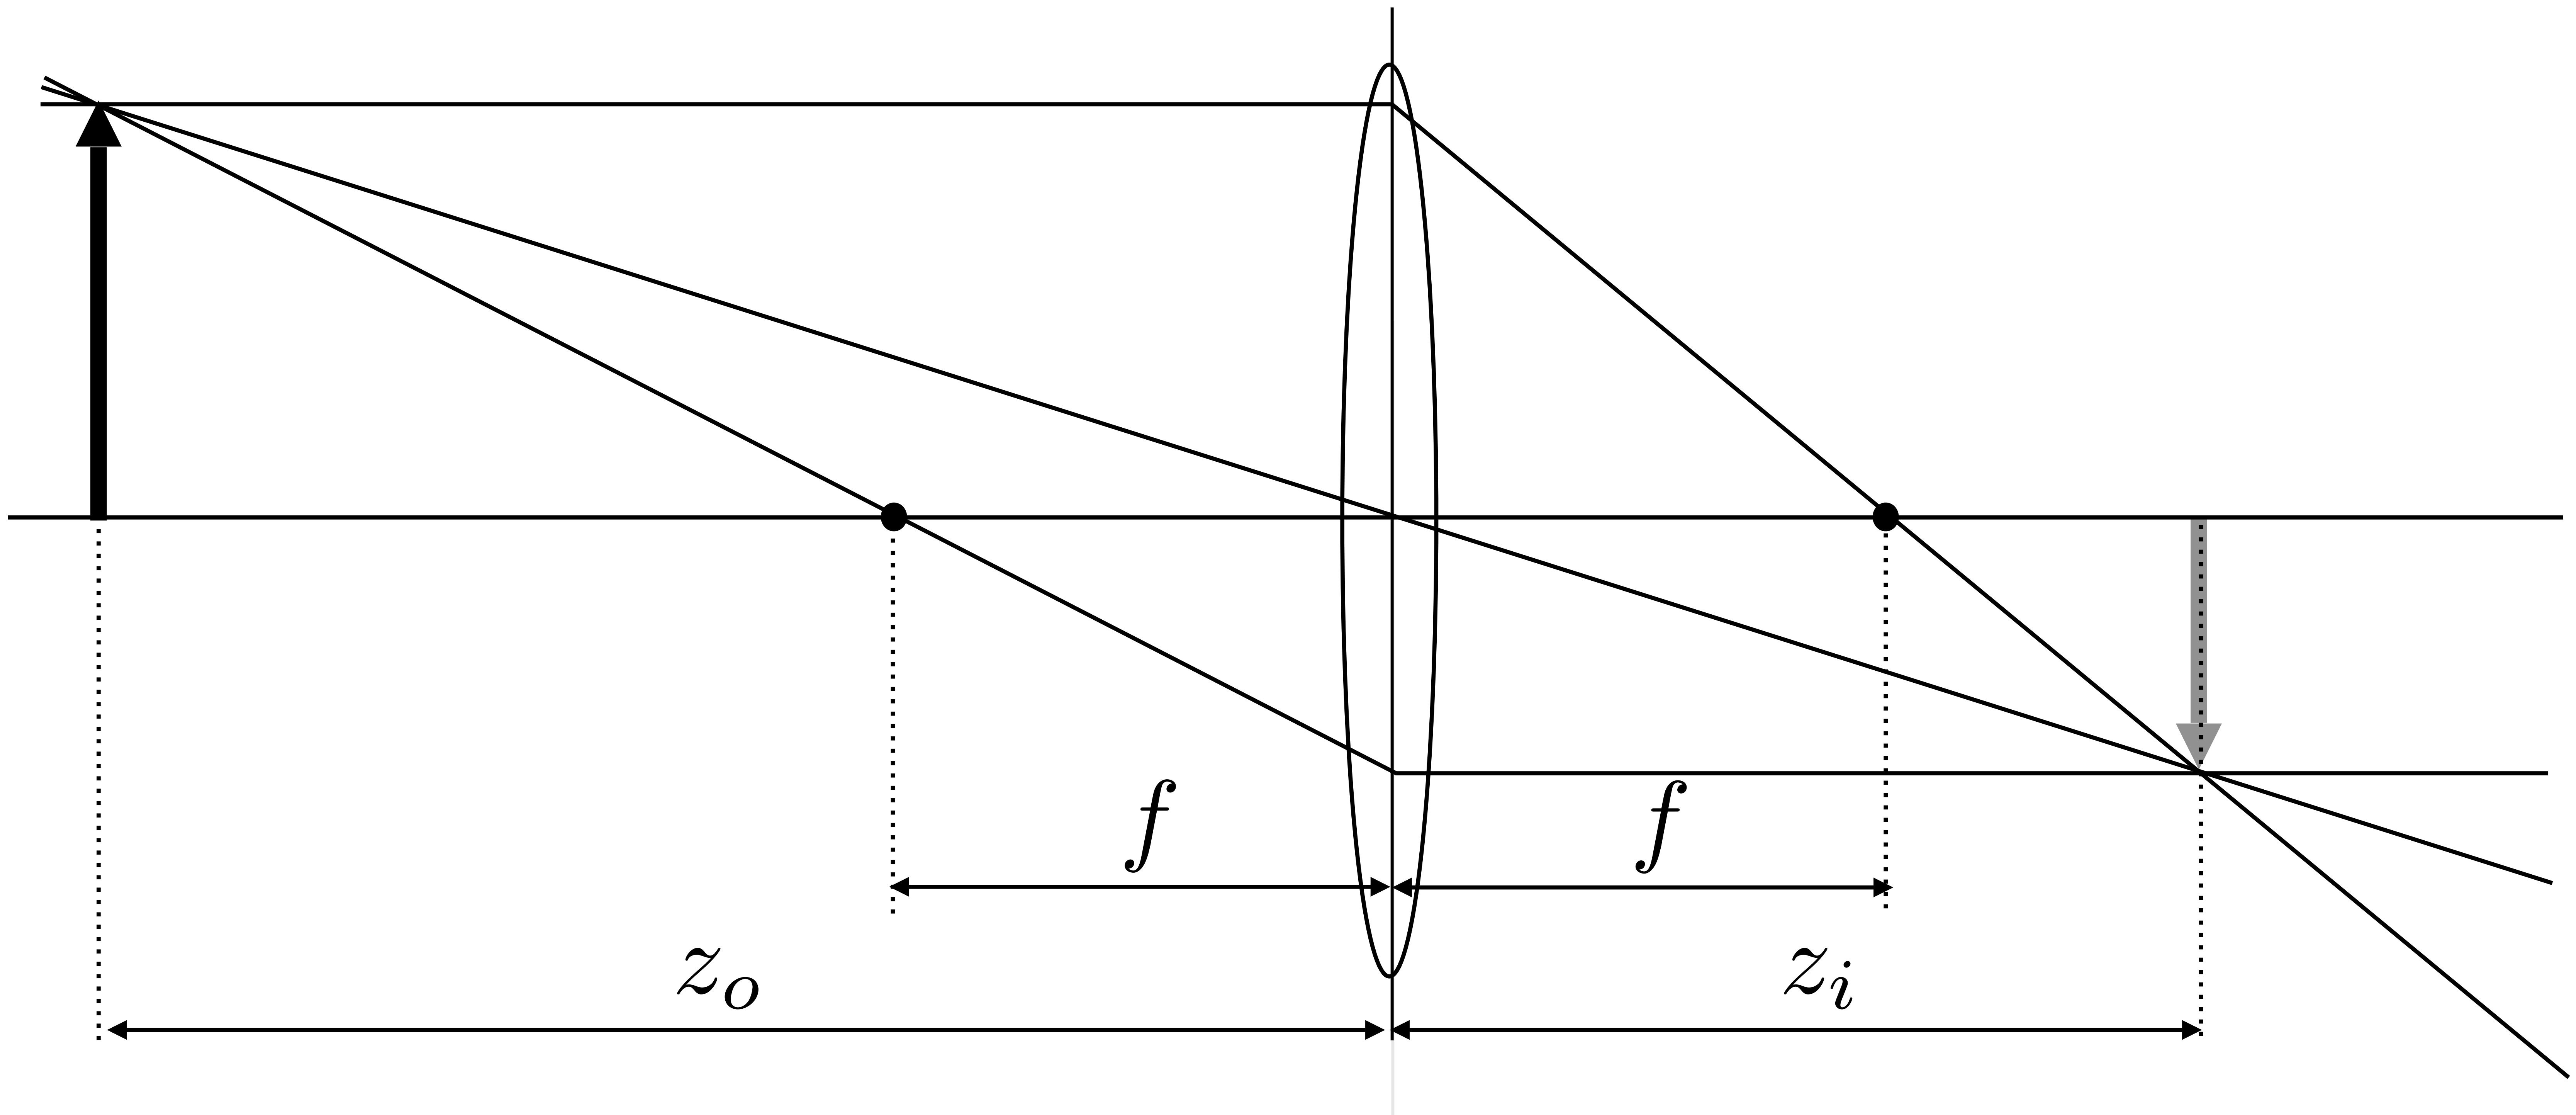

$$\frac{1}{f} = \frac{1}{z_i} + \frac{1}{z_o}$$

Changing the Focus Distance

$$\frac{1}{f} = \frac{1}{z_i} + \frac{1}{z_o}$$

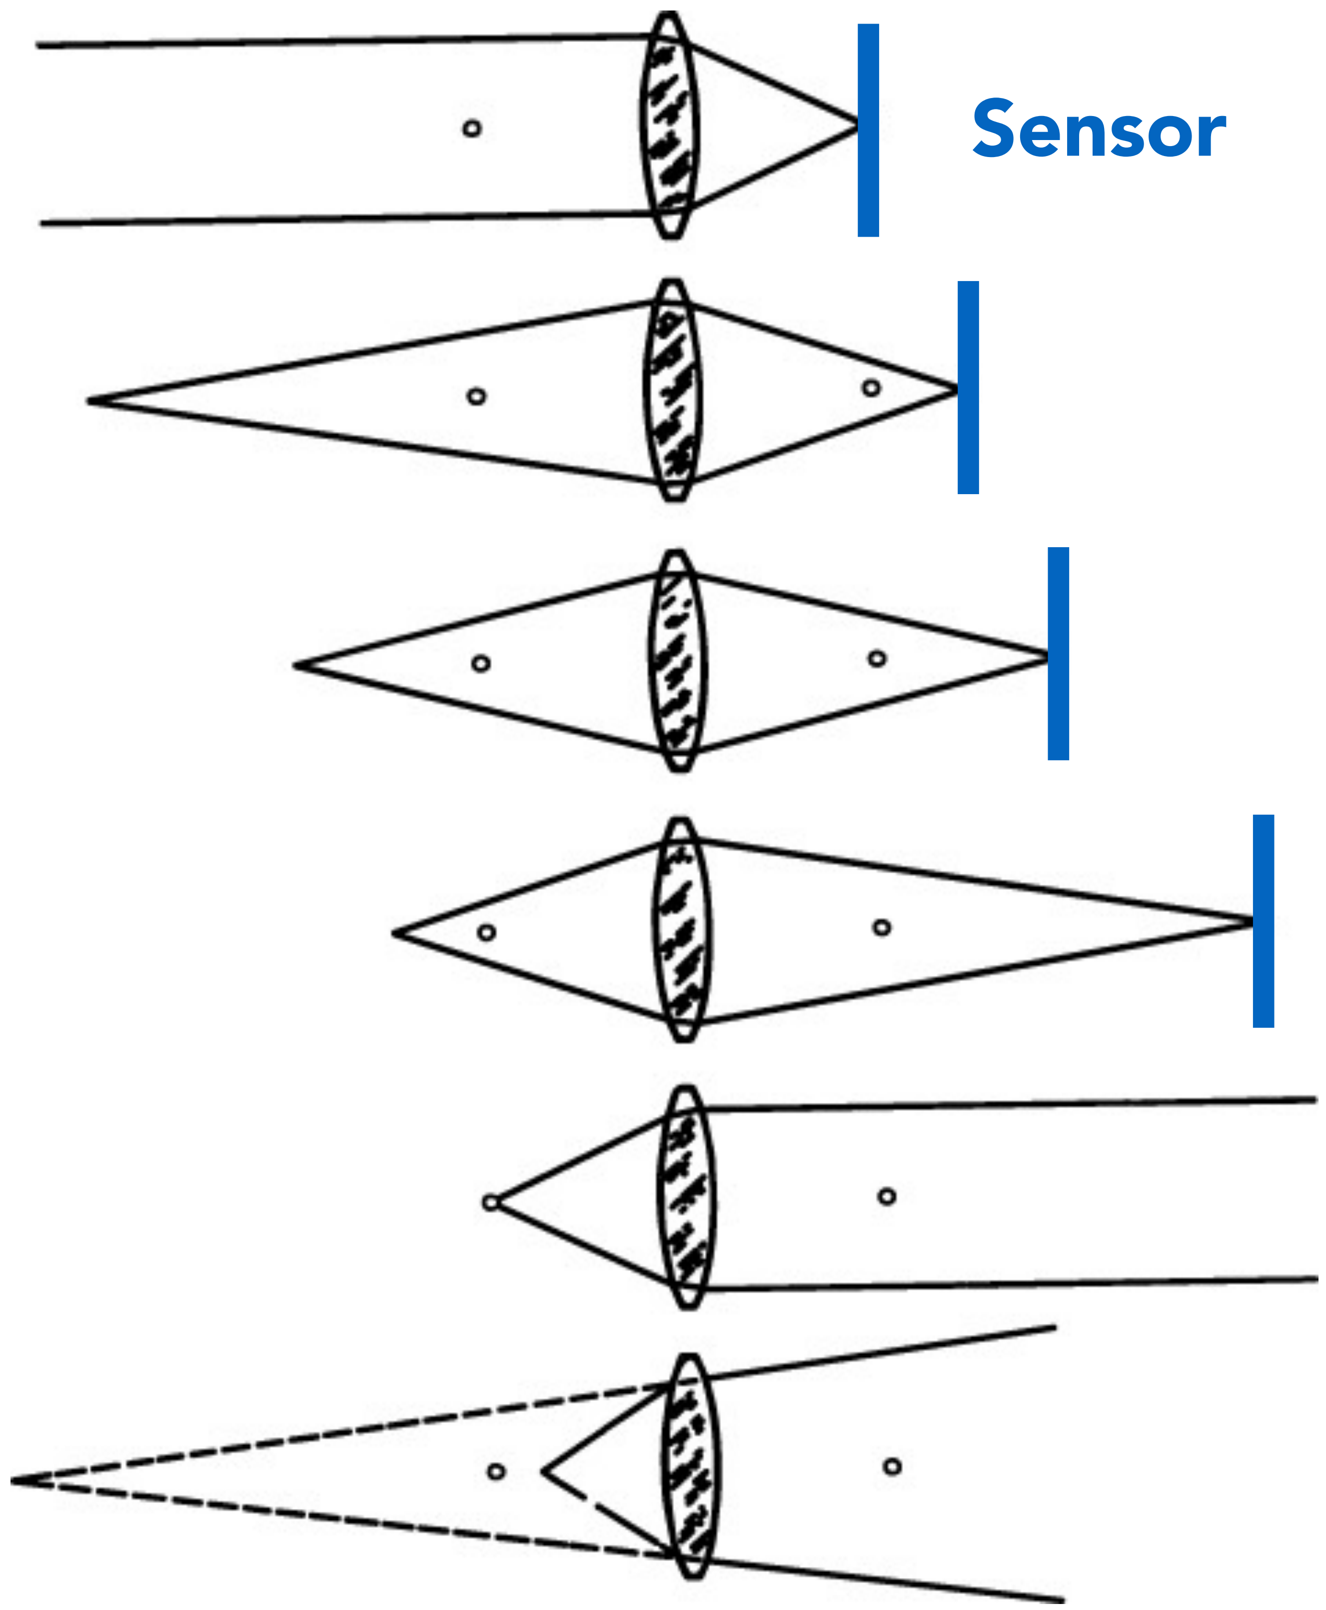

- z_i and z_o are called <u>conjugate points</u>
- To focus on objects at different distances, move the sensor relative to the lens
- For $z_i < z_o$ the object is larger than the image
- At $z_i = z_o$ we have 1:1 macro imaging
- For $z_i > z_o$ the image is larger than the object (magnified)
- Can't focus on objects closer than the lens' focal length

Magnification

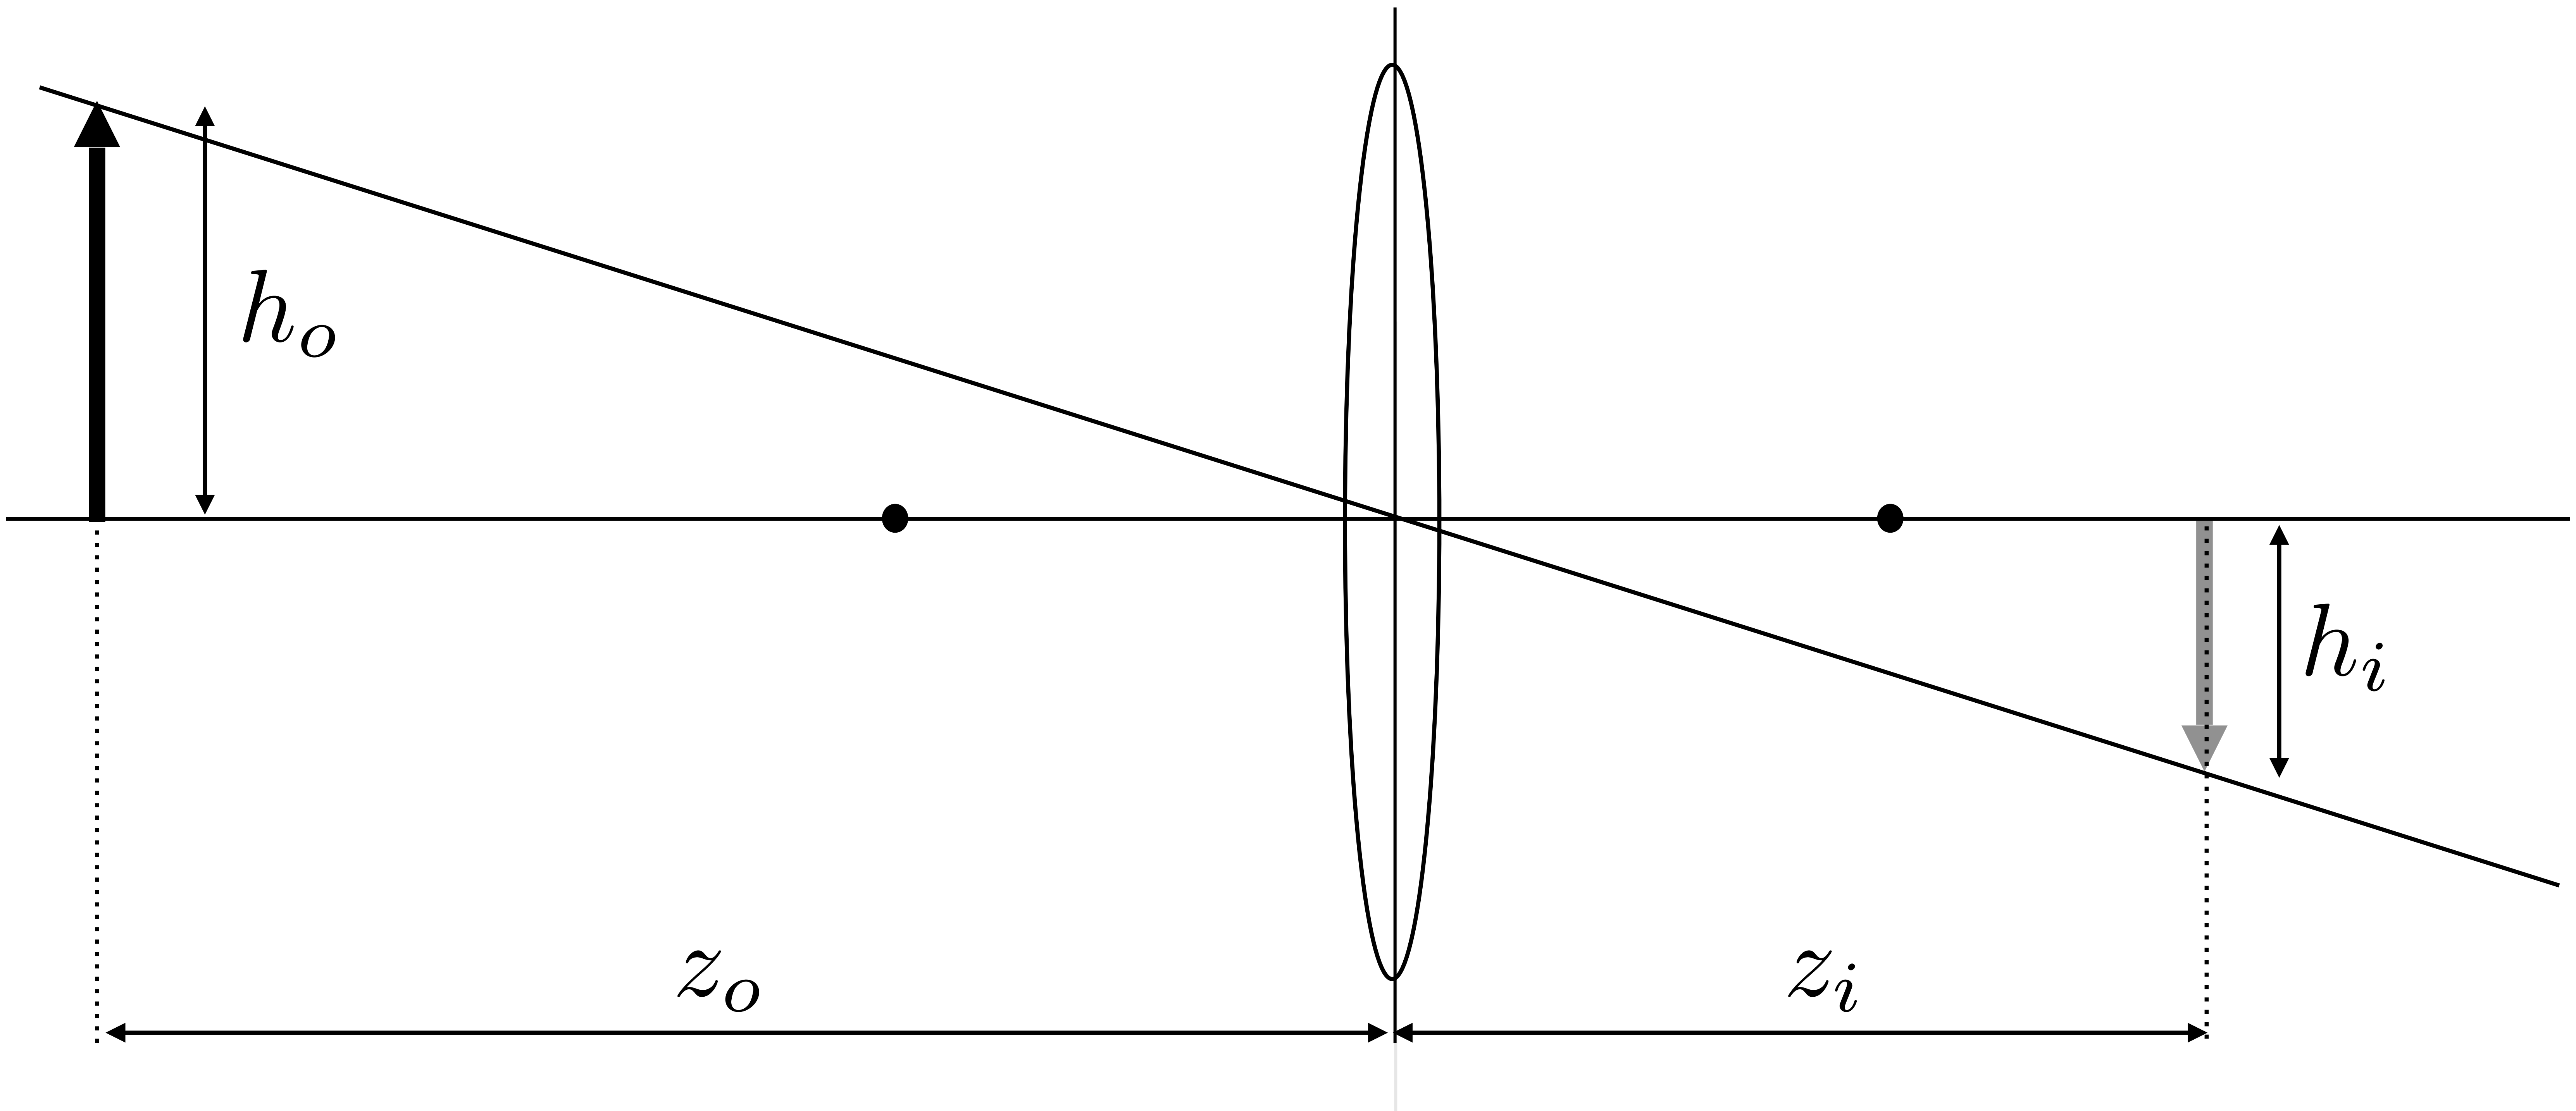

$$m = \frac{h_i}{h_o} = \frac{z_i}{z_o}$$

Magnification Example – Focus at Infinity

$$\frac{1}{f} = \frac{1}{z_i} + \frac{1}{z_o} \qquad m = \frac{z_i}{z_o}$$

If focused on a distant mountain

- $z_o \approx \infty$, so $z_i = f$
- sensor at focal point
- magnification ≈ 0

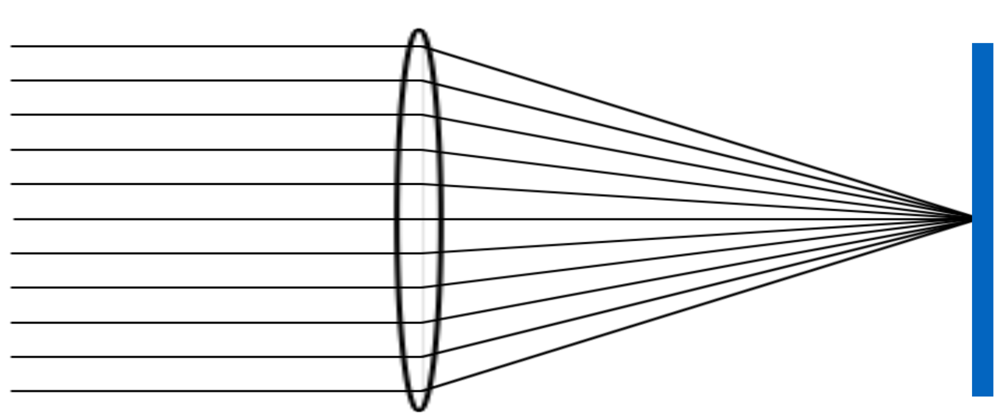

Magnification Example – Focus at 1:1 Macro

$$\frac{1}{f} = \frac{1}{z_i} + \frac{1}{z_o} \qquad m = \frac{z_i}{z_o}$$

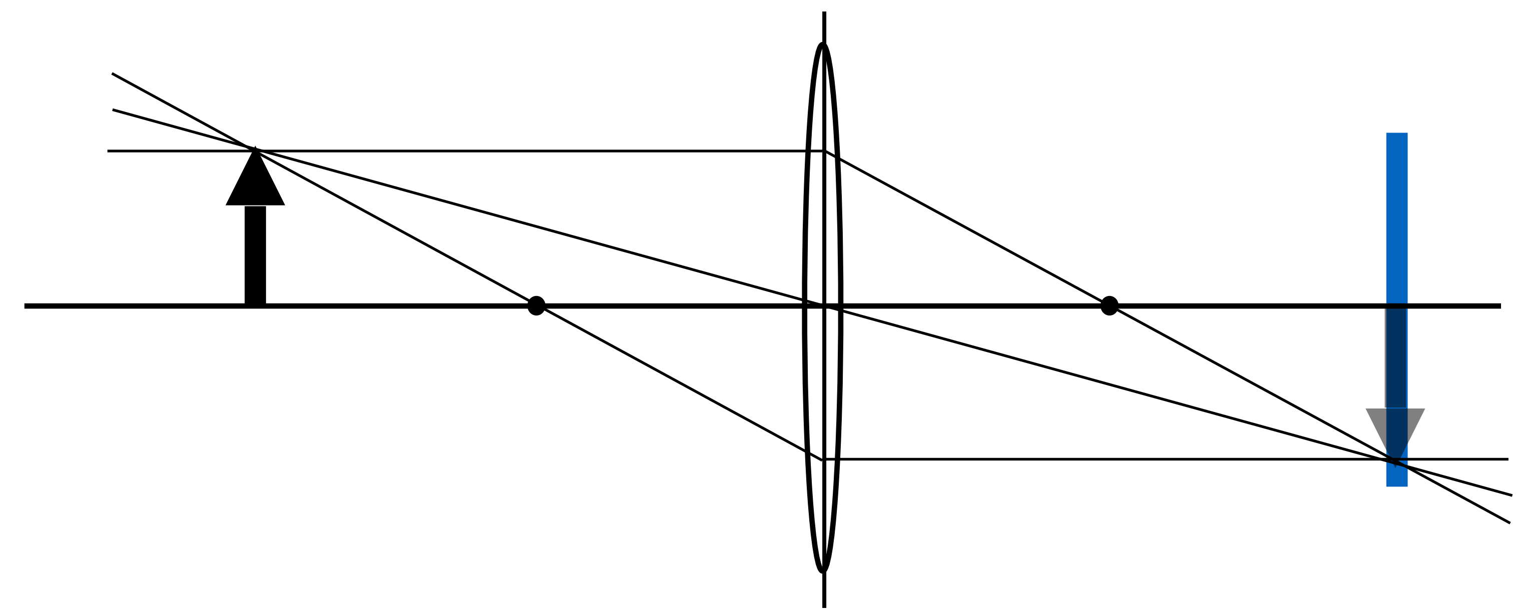

What configuration do we need to achieve a magnification of 1 (i.e. image and object the same size, a.k.a. 1:1 macro)?

- Need $z_i = z_o$, so $z_i = z_o = 2f$ —- sensor at twice focal length
- In 1:1 imaging, if the sensor is 36 mm wide, an object 36 mm wide will fill the frame

Relationship between focal length and field of view

Effect of Focal Length on FOV

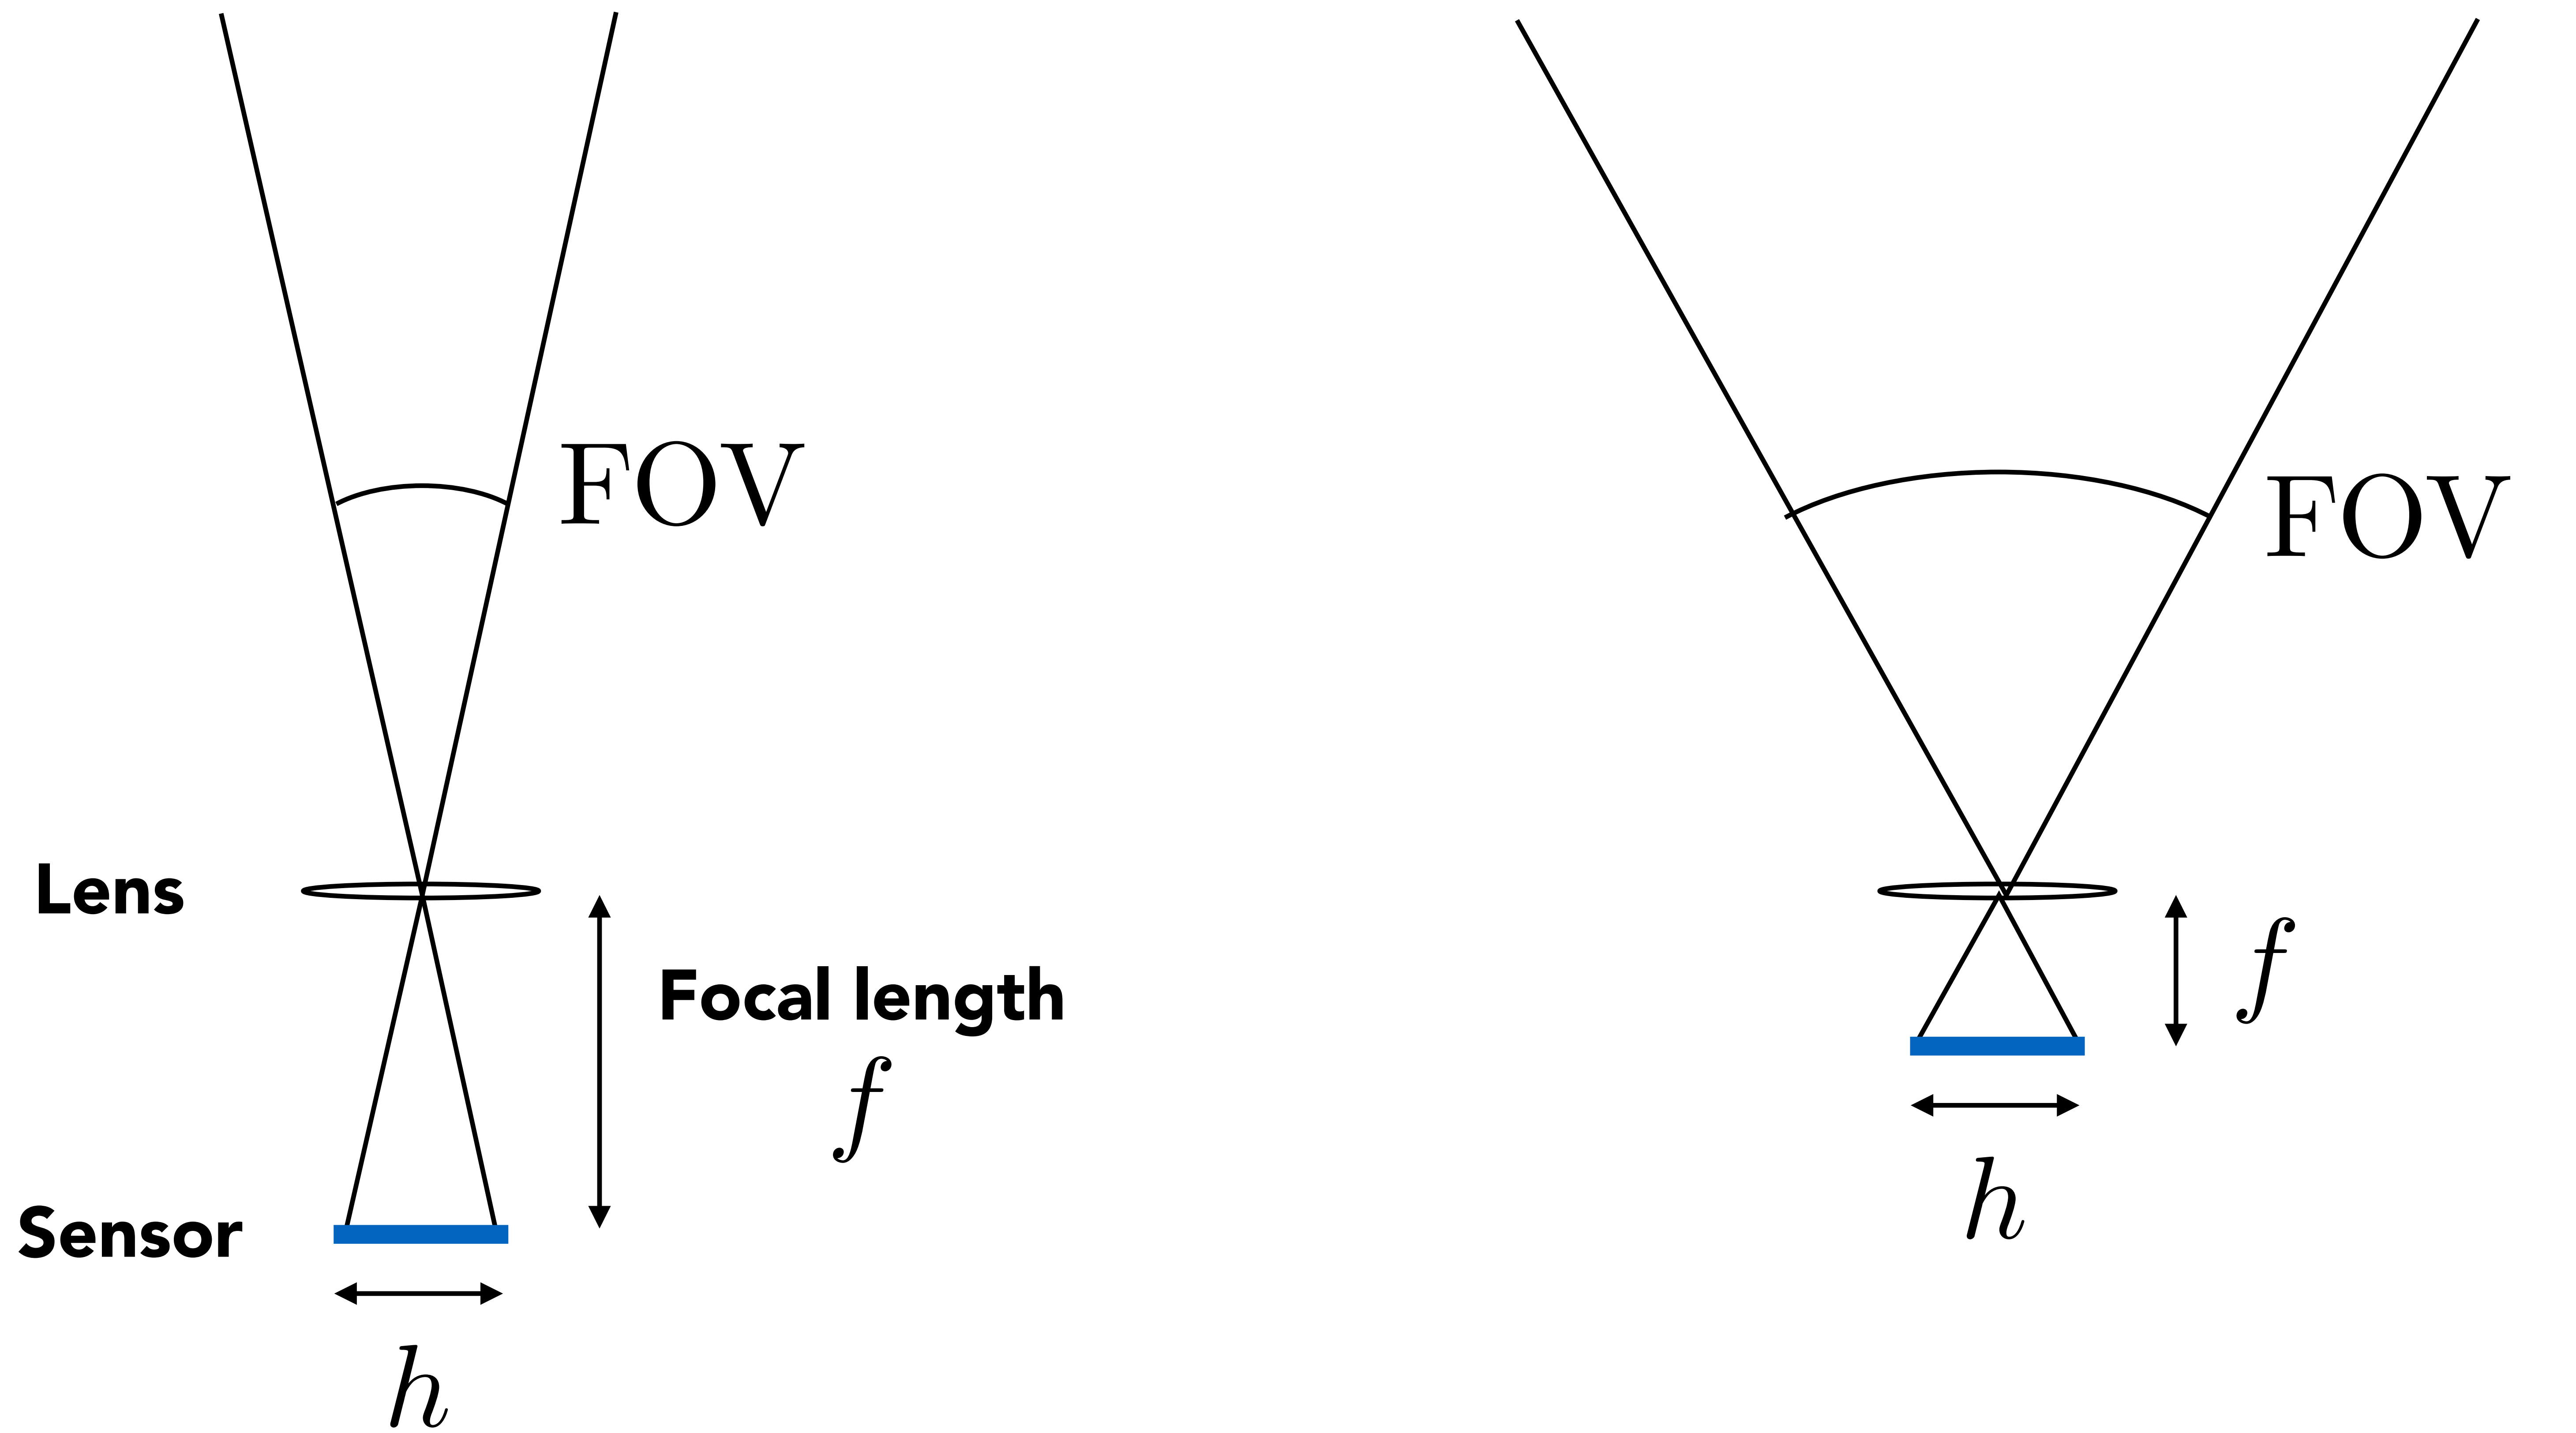

For a fixed sensor size, decreasing the focal length increases the field of view.

$$\text{FOV} = 2\,\arctan\left(\frac{h}{2f}\right)$$

Effect of Sensor Size on FOV

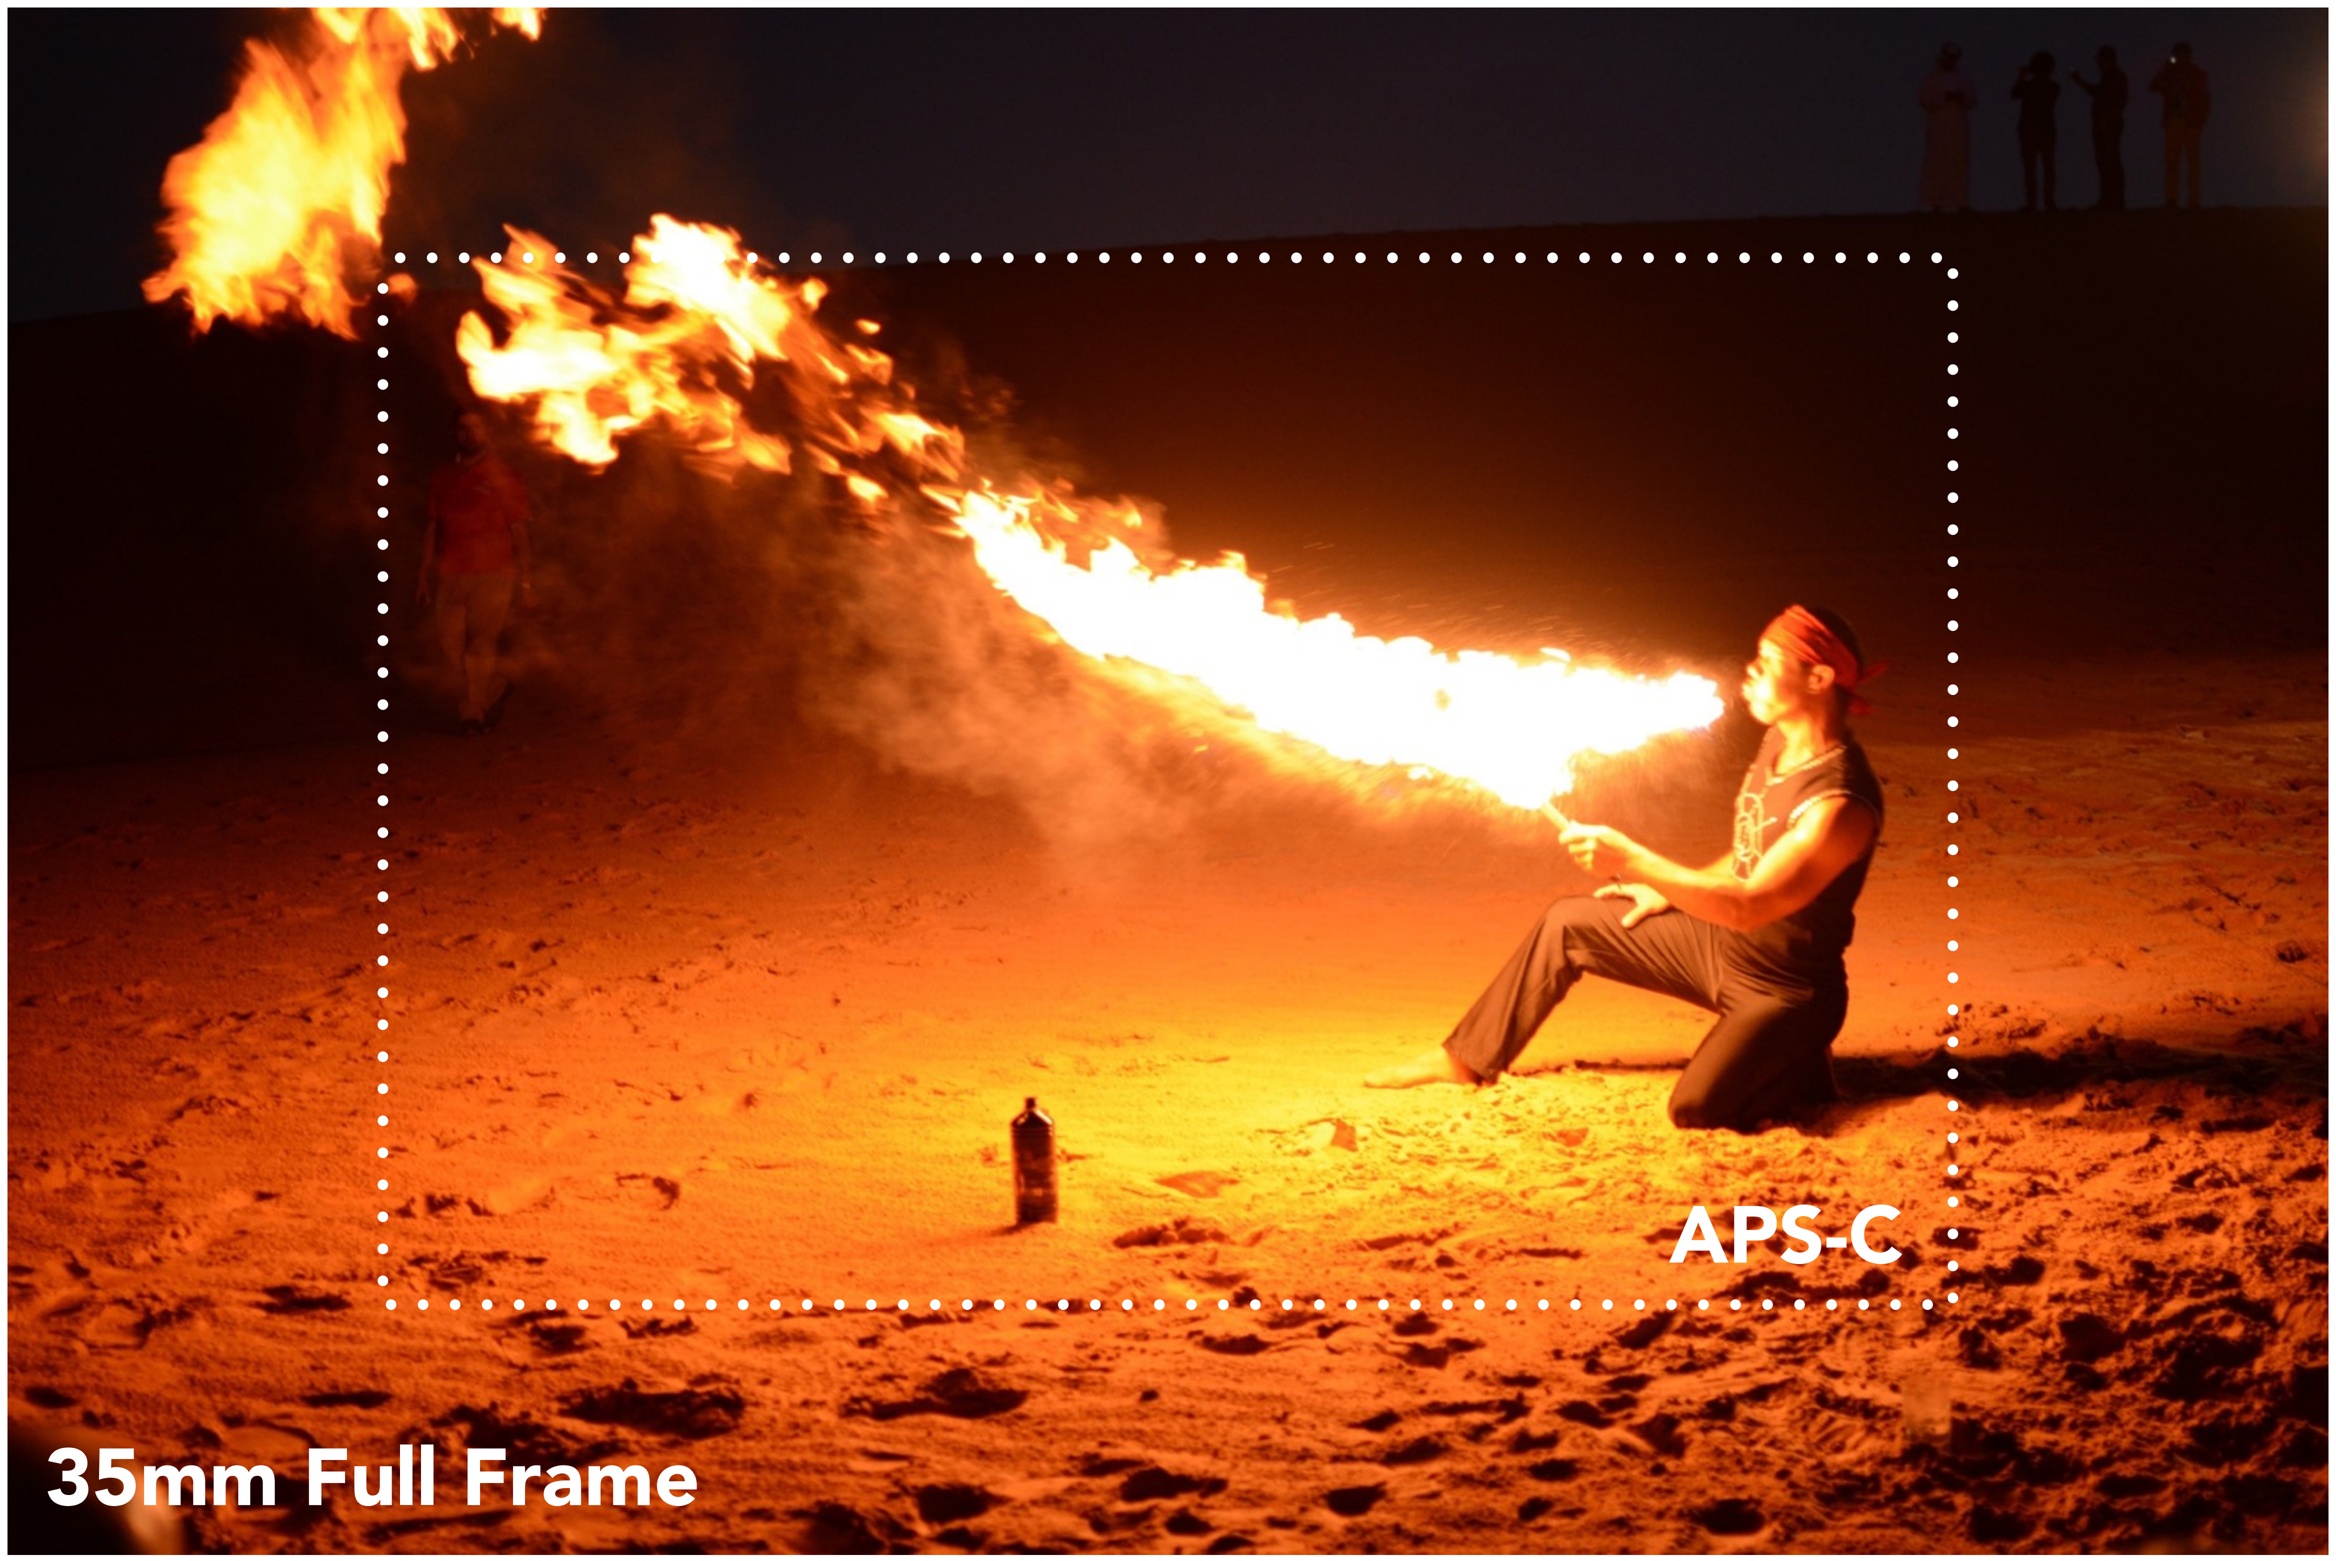

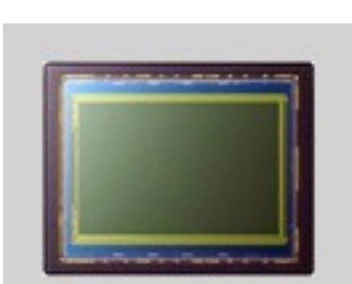

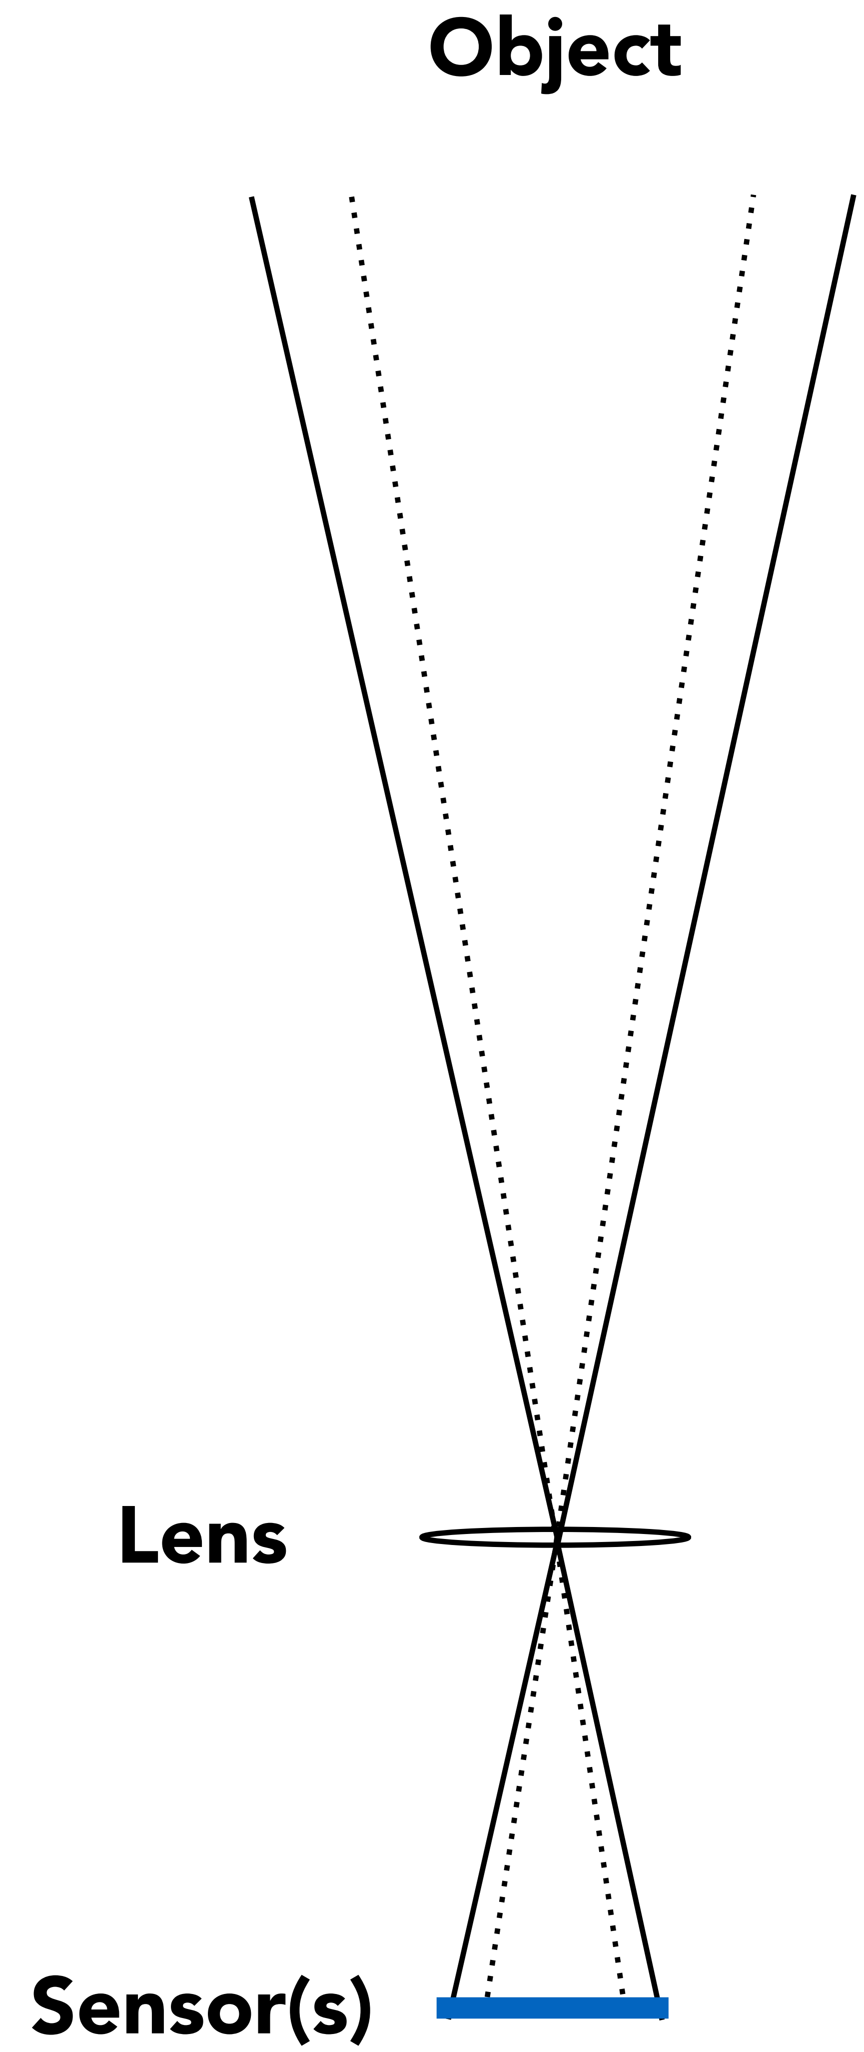

Sensor Sizes

Sensor Name	Medium Format	Full Frame	APS-H	APS-C	4/3	1"	1/1.63"	1/2.3"	1/3.2"
Sensor Size	53.7 x 40.2mm	36 x 23.9mm	27.9x18.6mm	23.6x15.8mm	17.3x13mm	13.2x8.8mm	8.38x5.59mm	6.16x4.62mm	4.54x3.42mm
Sensor Area	21.59 cm²	8.6 cm²	5.19 cm²	3.73 cm²	2.25 cm²	1.16 cm²	0.47 cm²	0.28 cm²	0.15 cm²
Crop Factor	0.64	1.0	1.29	1.52	2.0	2.7	4.3	5.62	7.61
Image				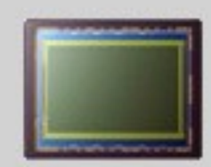				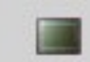	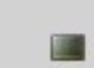
Example									

Credit: lensvid.com

Maintain FOV on Smaller Sensor?

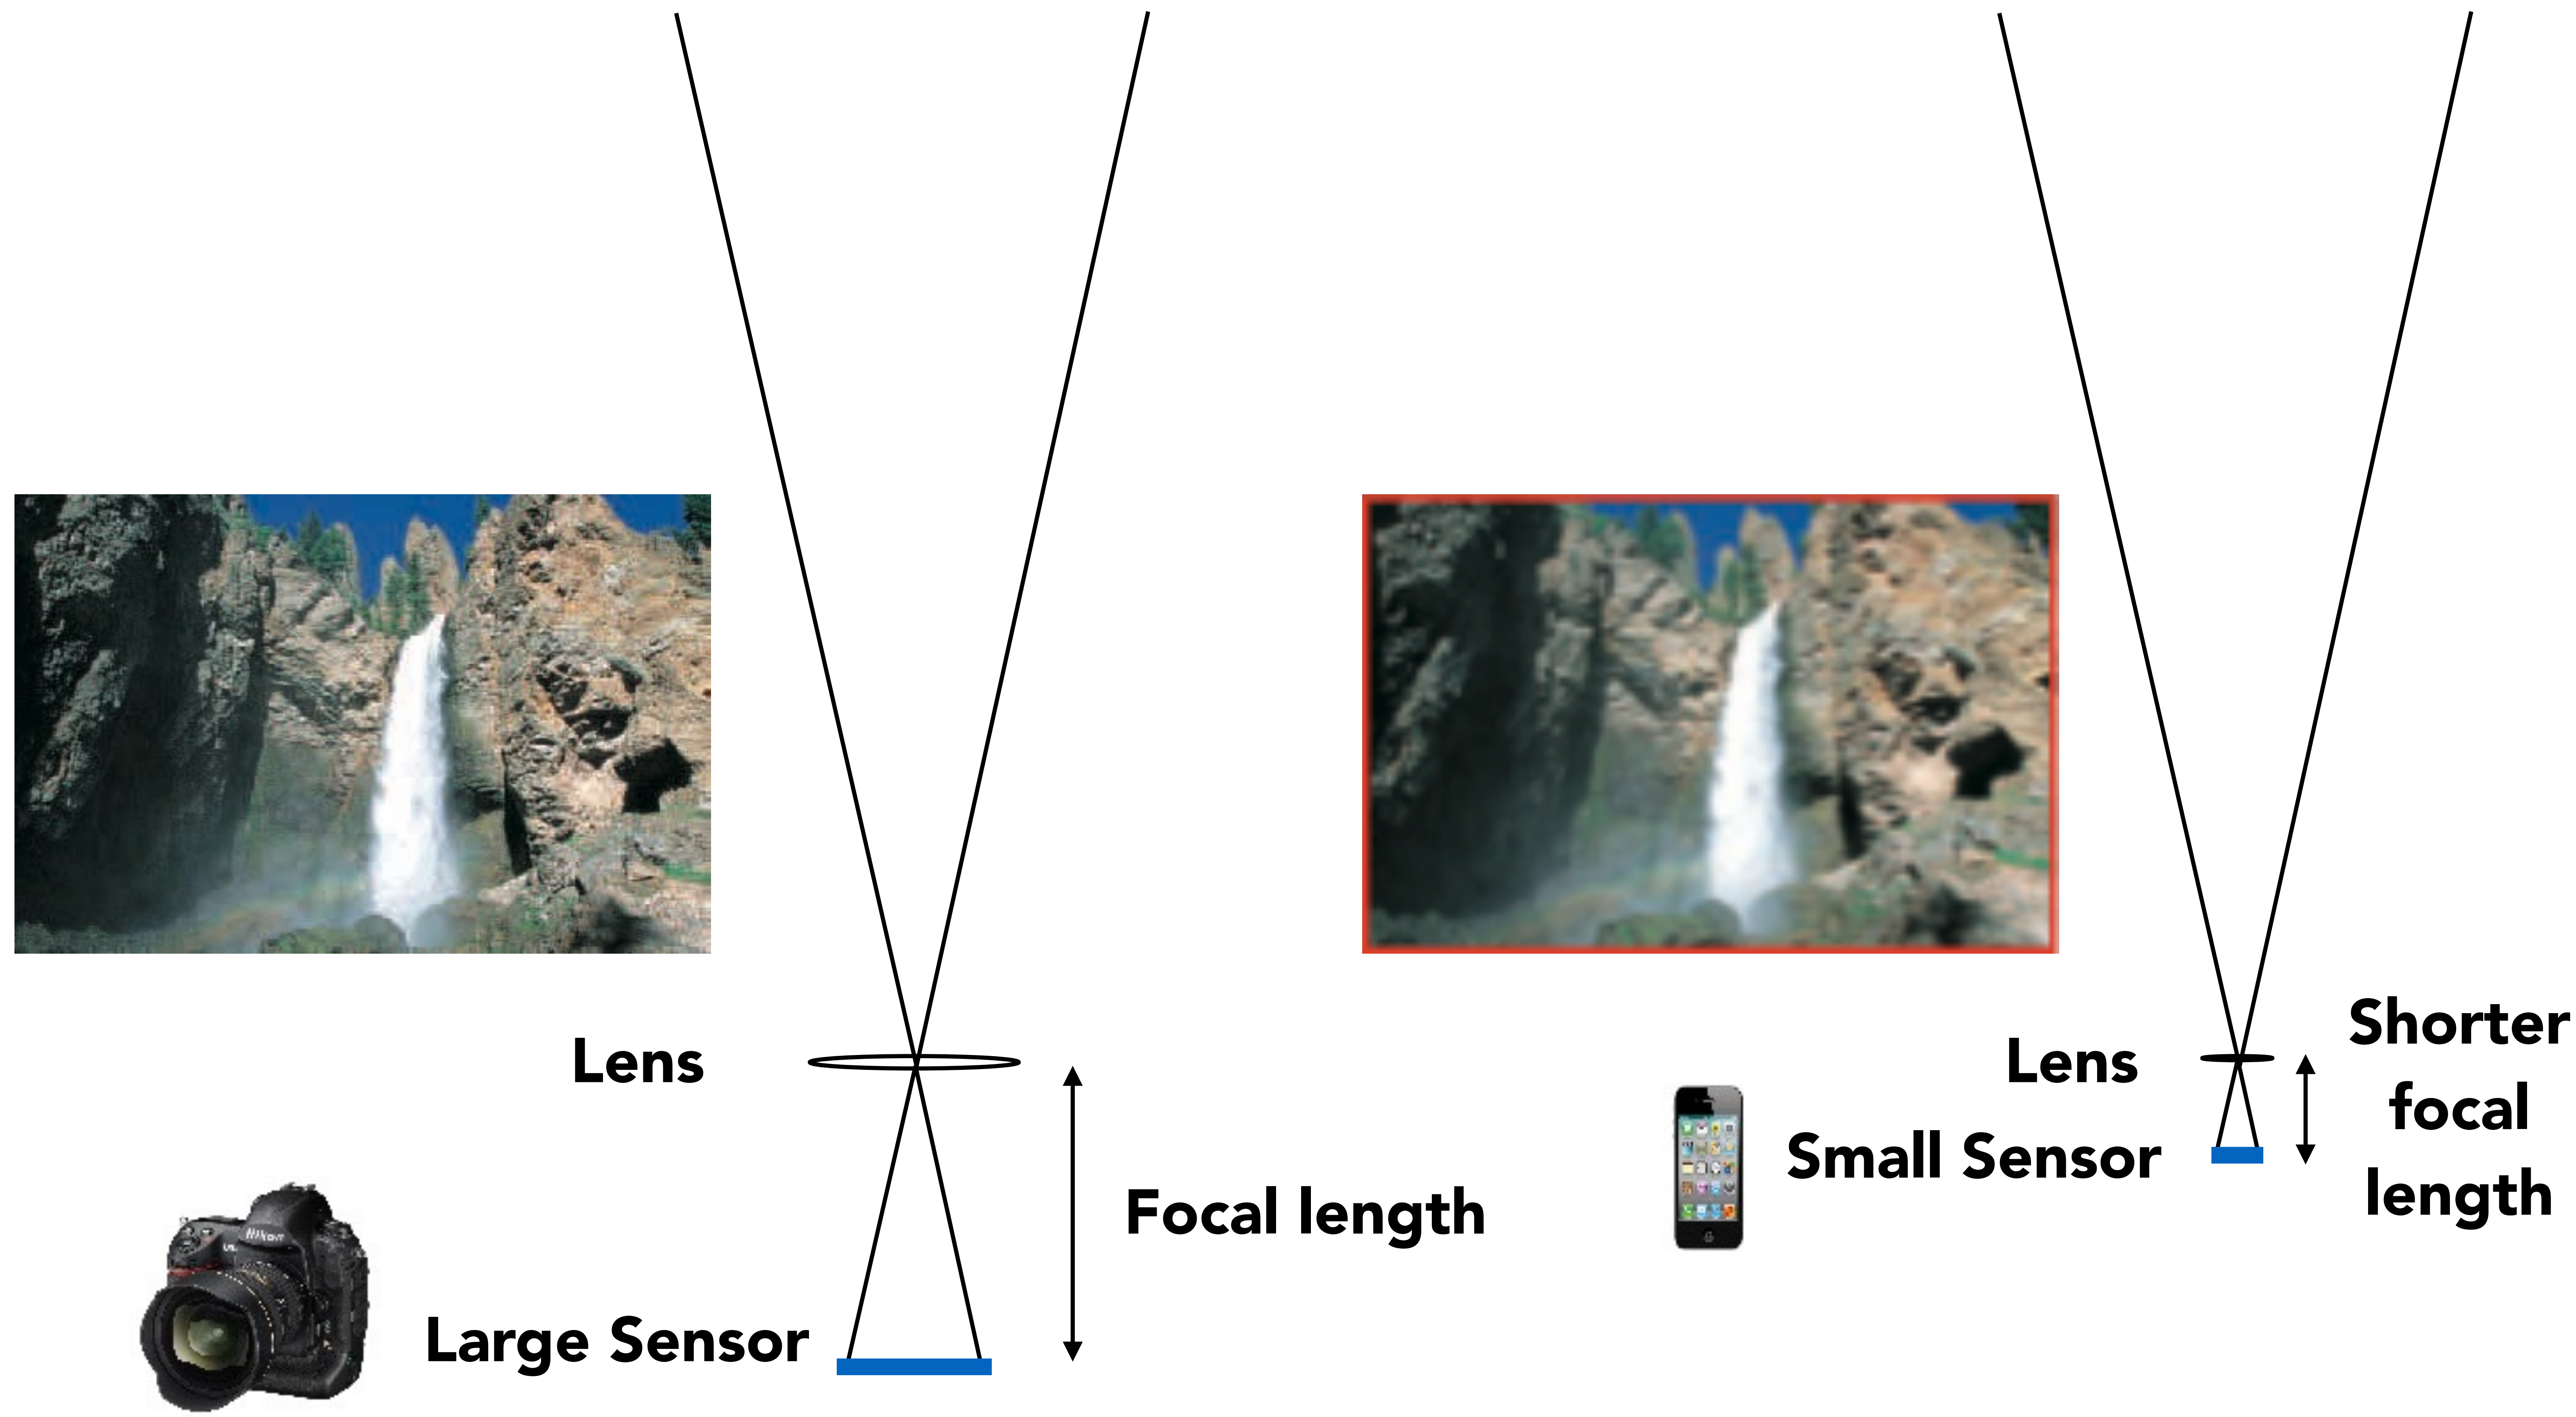

To maintain FOV, decrease focal length of lens in proportion to width/height of sensor

Focal Length v. Field of View

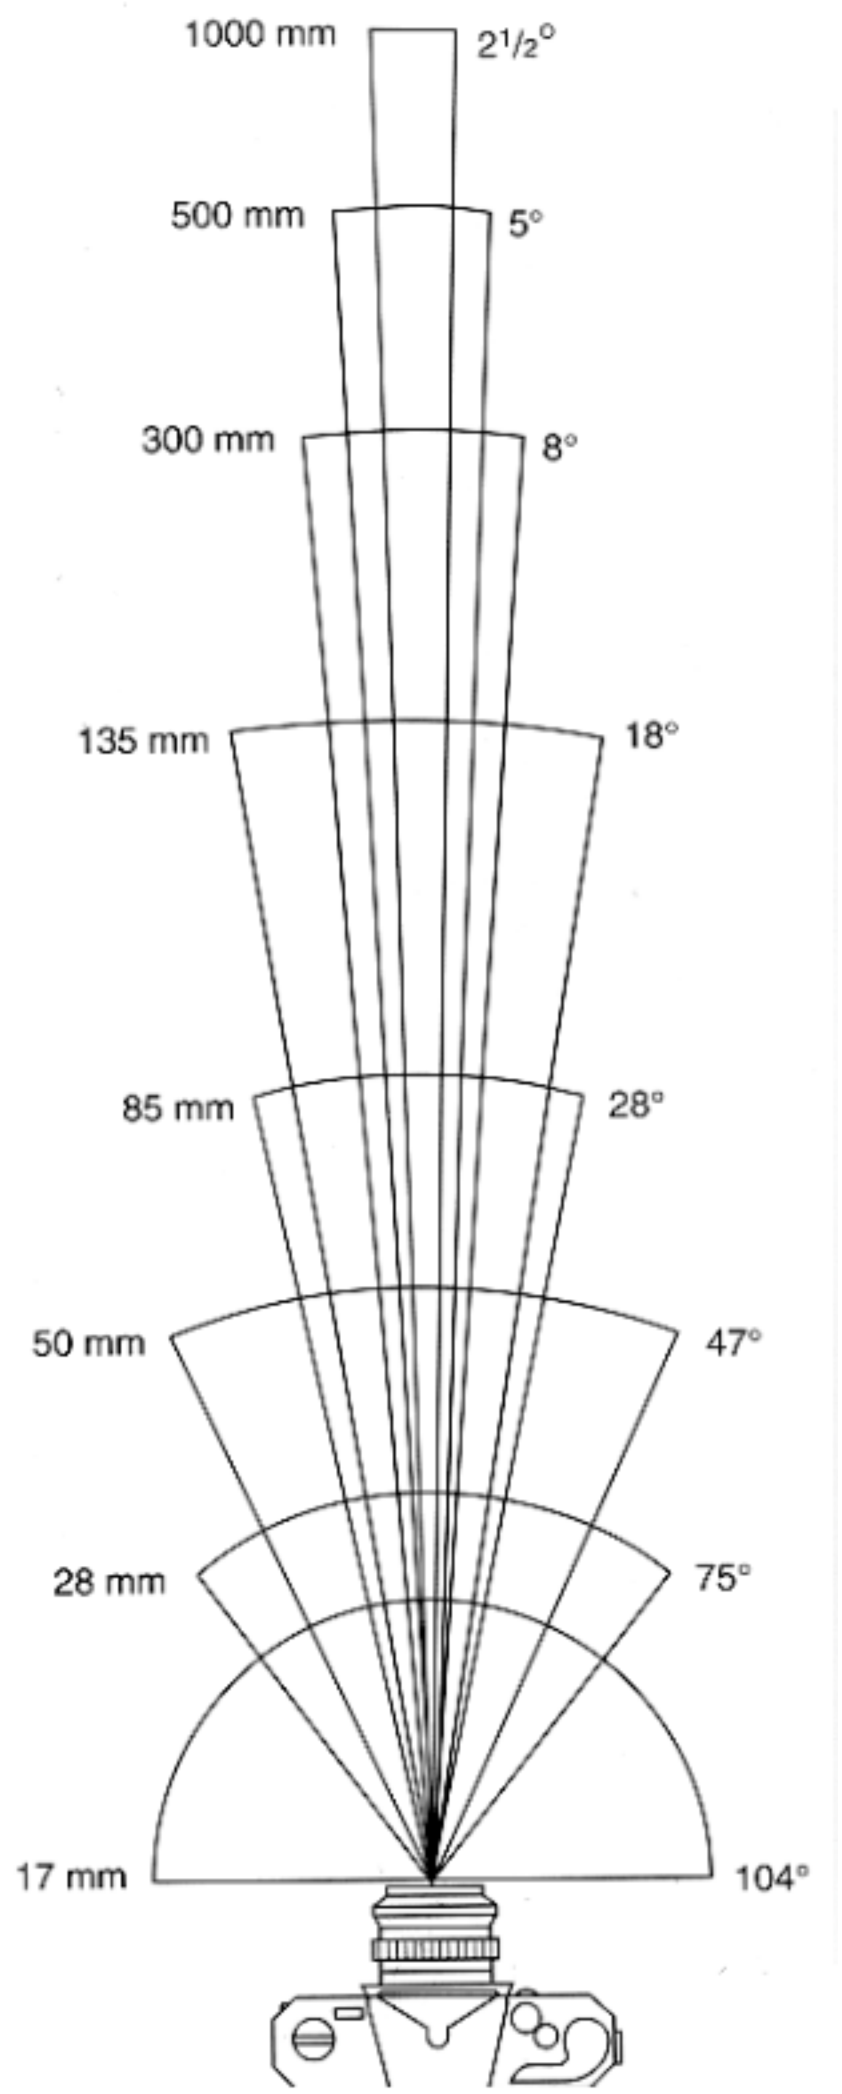

- For historical reasons, it is common to refer to angular field of view by focal length of a lens used on a 35mm-format film (36 x 24mm)
- Examples of focal lengths on 35mm format:
  - 17mm is wide angle 104°
  - 50mm is a "normal" lens 47°
  - 200mm is telephoto lens 12°
- Careful! When we say current cell phones have approximately 28mm "equivalent" focal length, this uses the above convention. The physical focal length is often 5-6 times shorter, because the sensor is correspondingly smaller

Focal Length v. Field of View

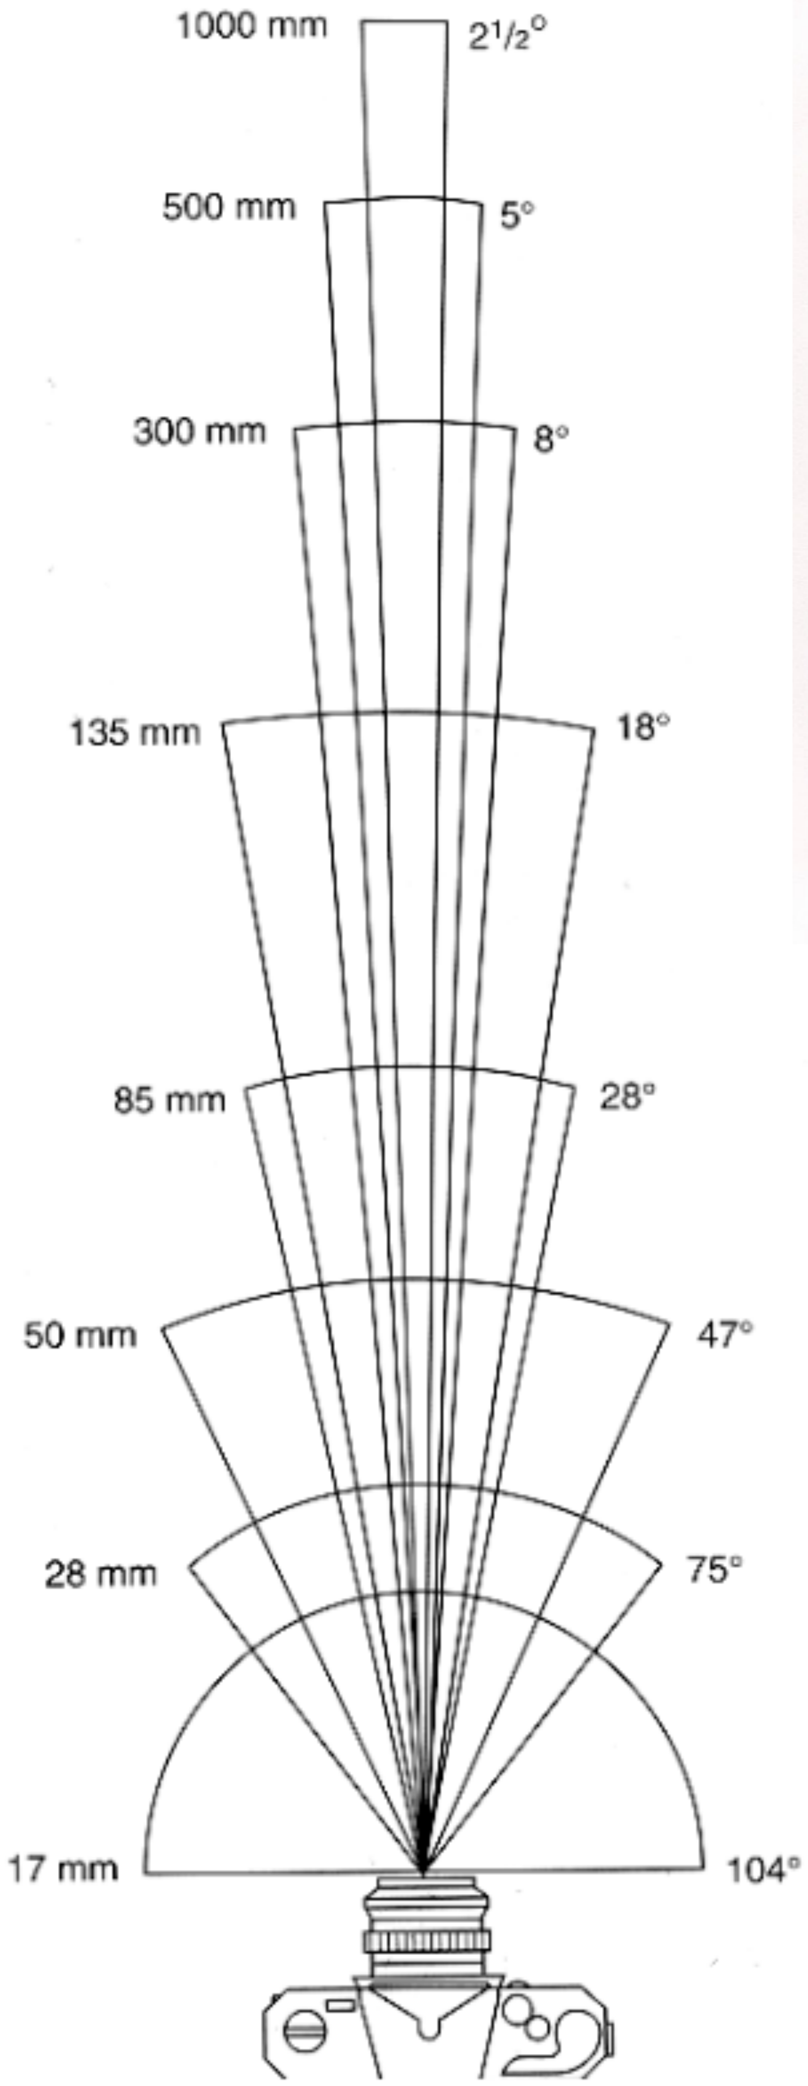

From London and Upton, and Canon EF Lens Work III

Focal Length v. Field of View

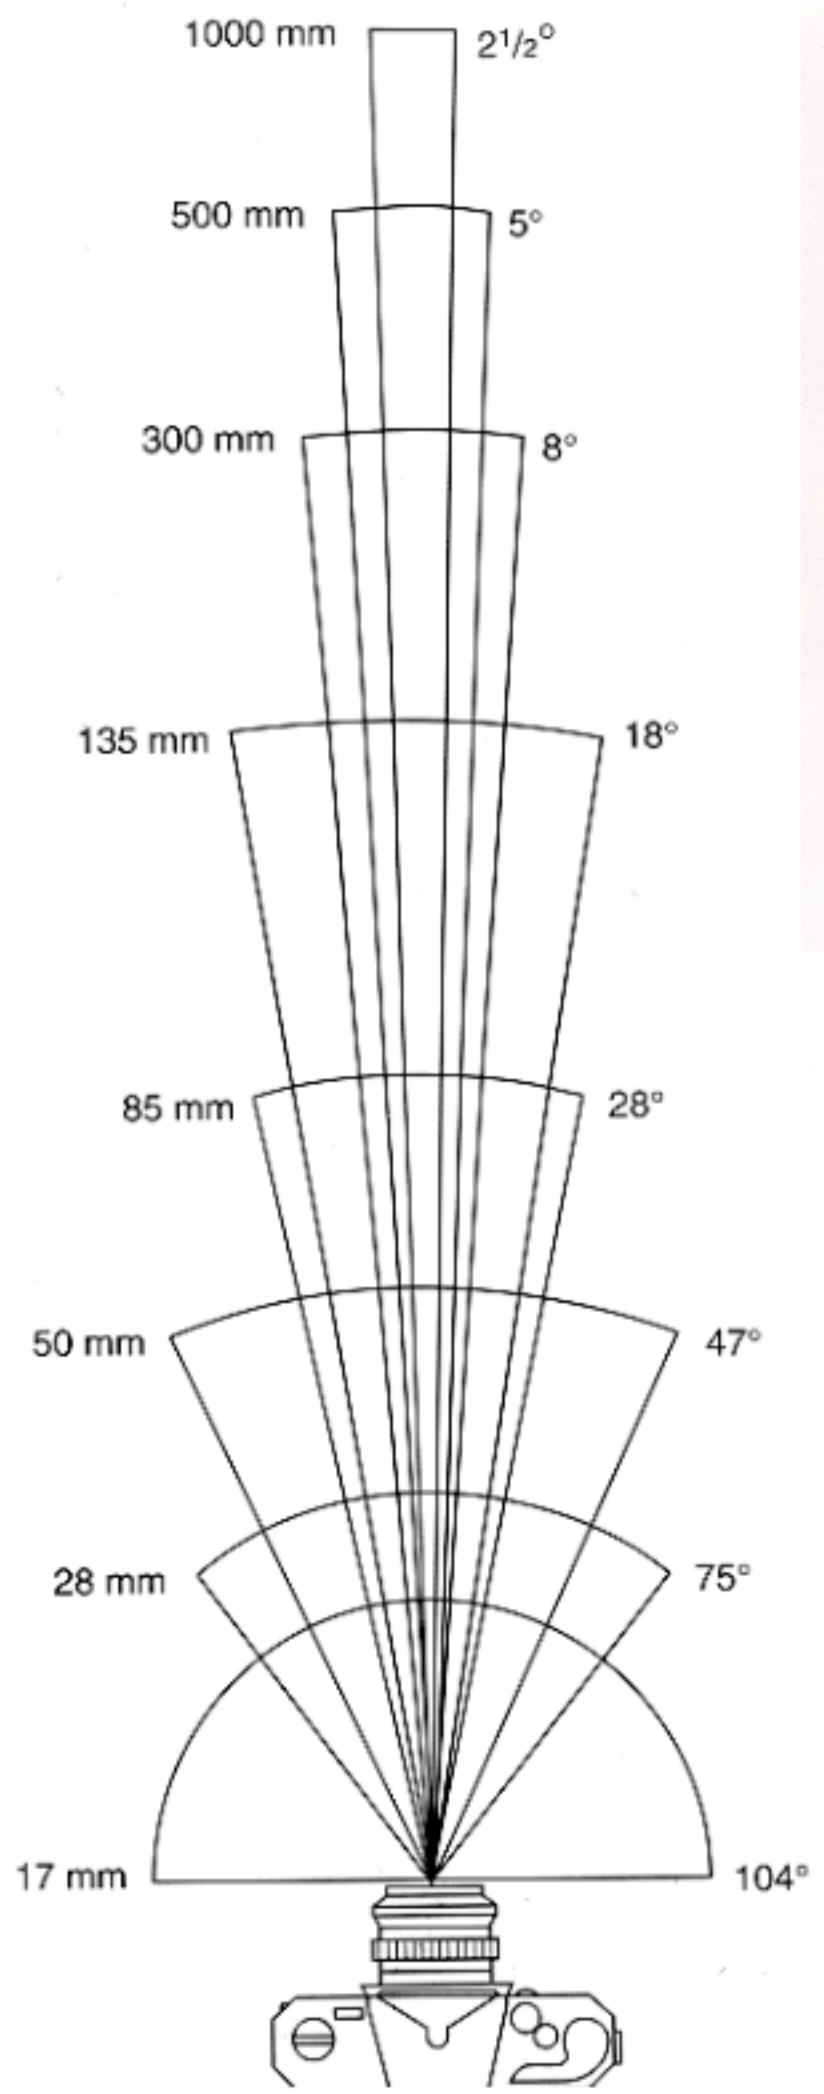

From London and Upton, and Canon EF Lens Work III

Focal Length v. Field of View

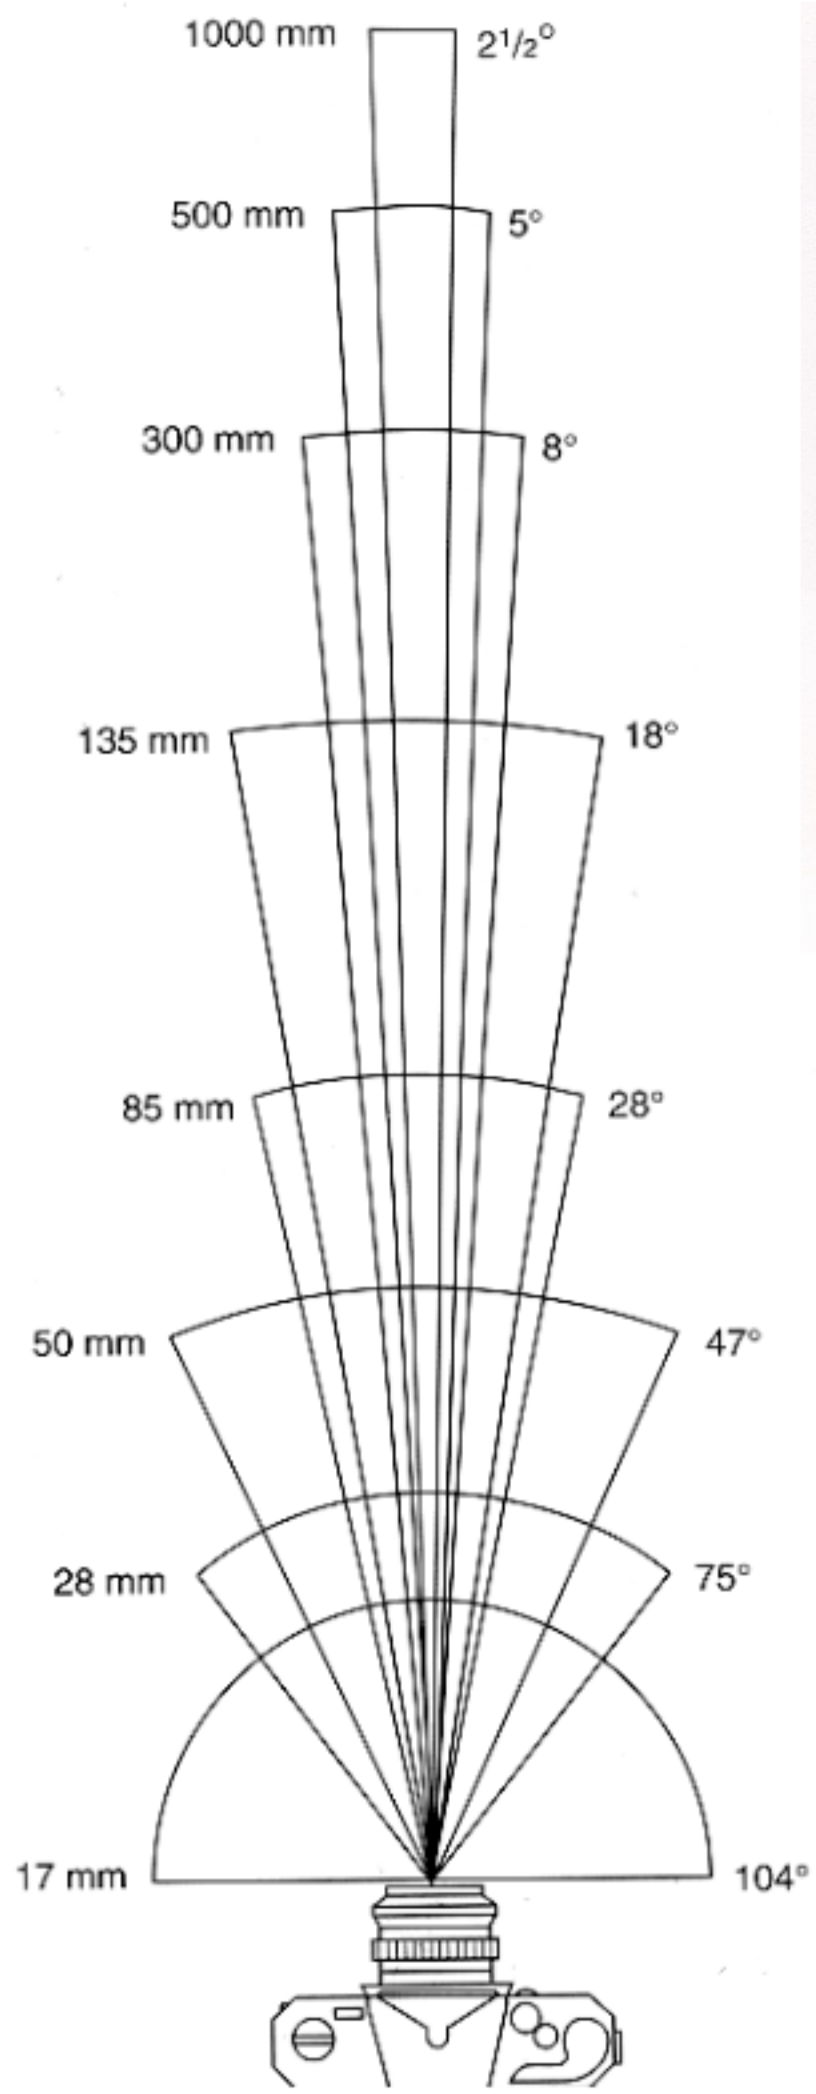

From London and Upton, and Canon EF Lens Work III

Defocus Blur

Defocus blur

Image credit: Gary Dineen/NBAE via Getty Images

Computing Circle of Confusion Diameter (C)

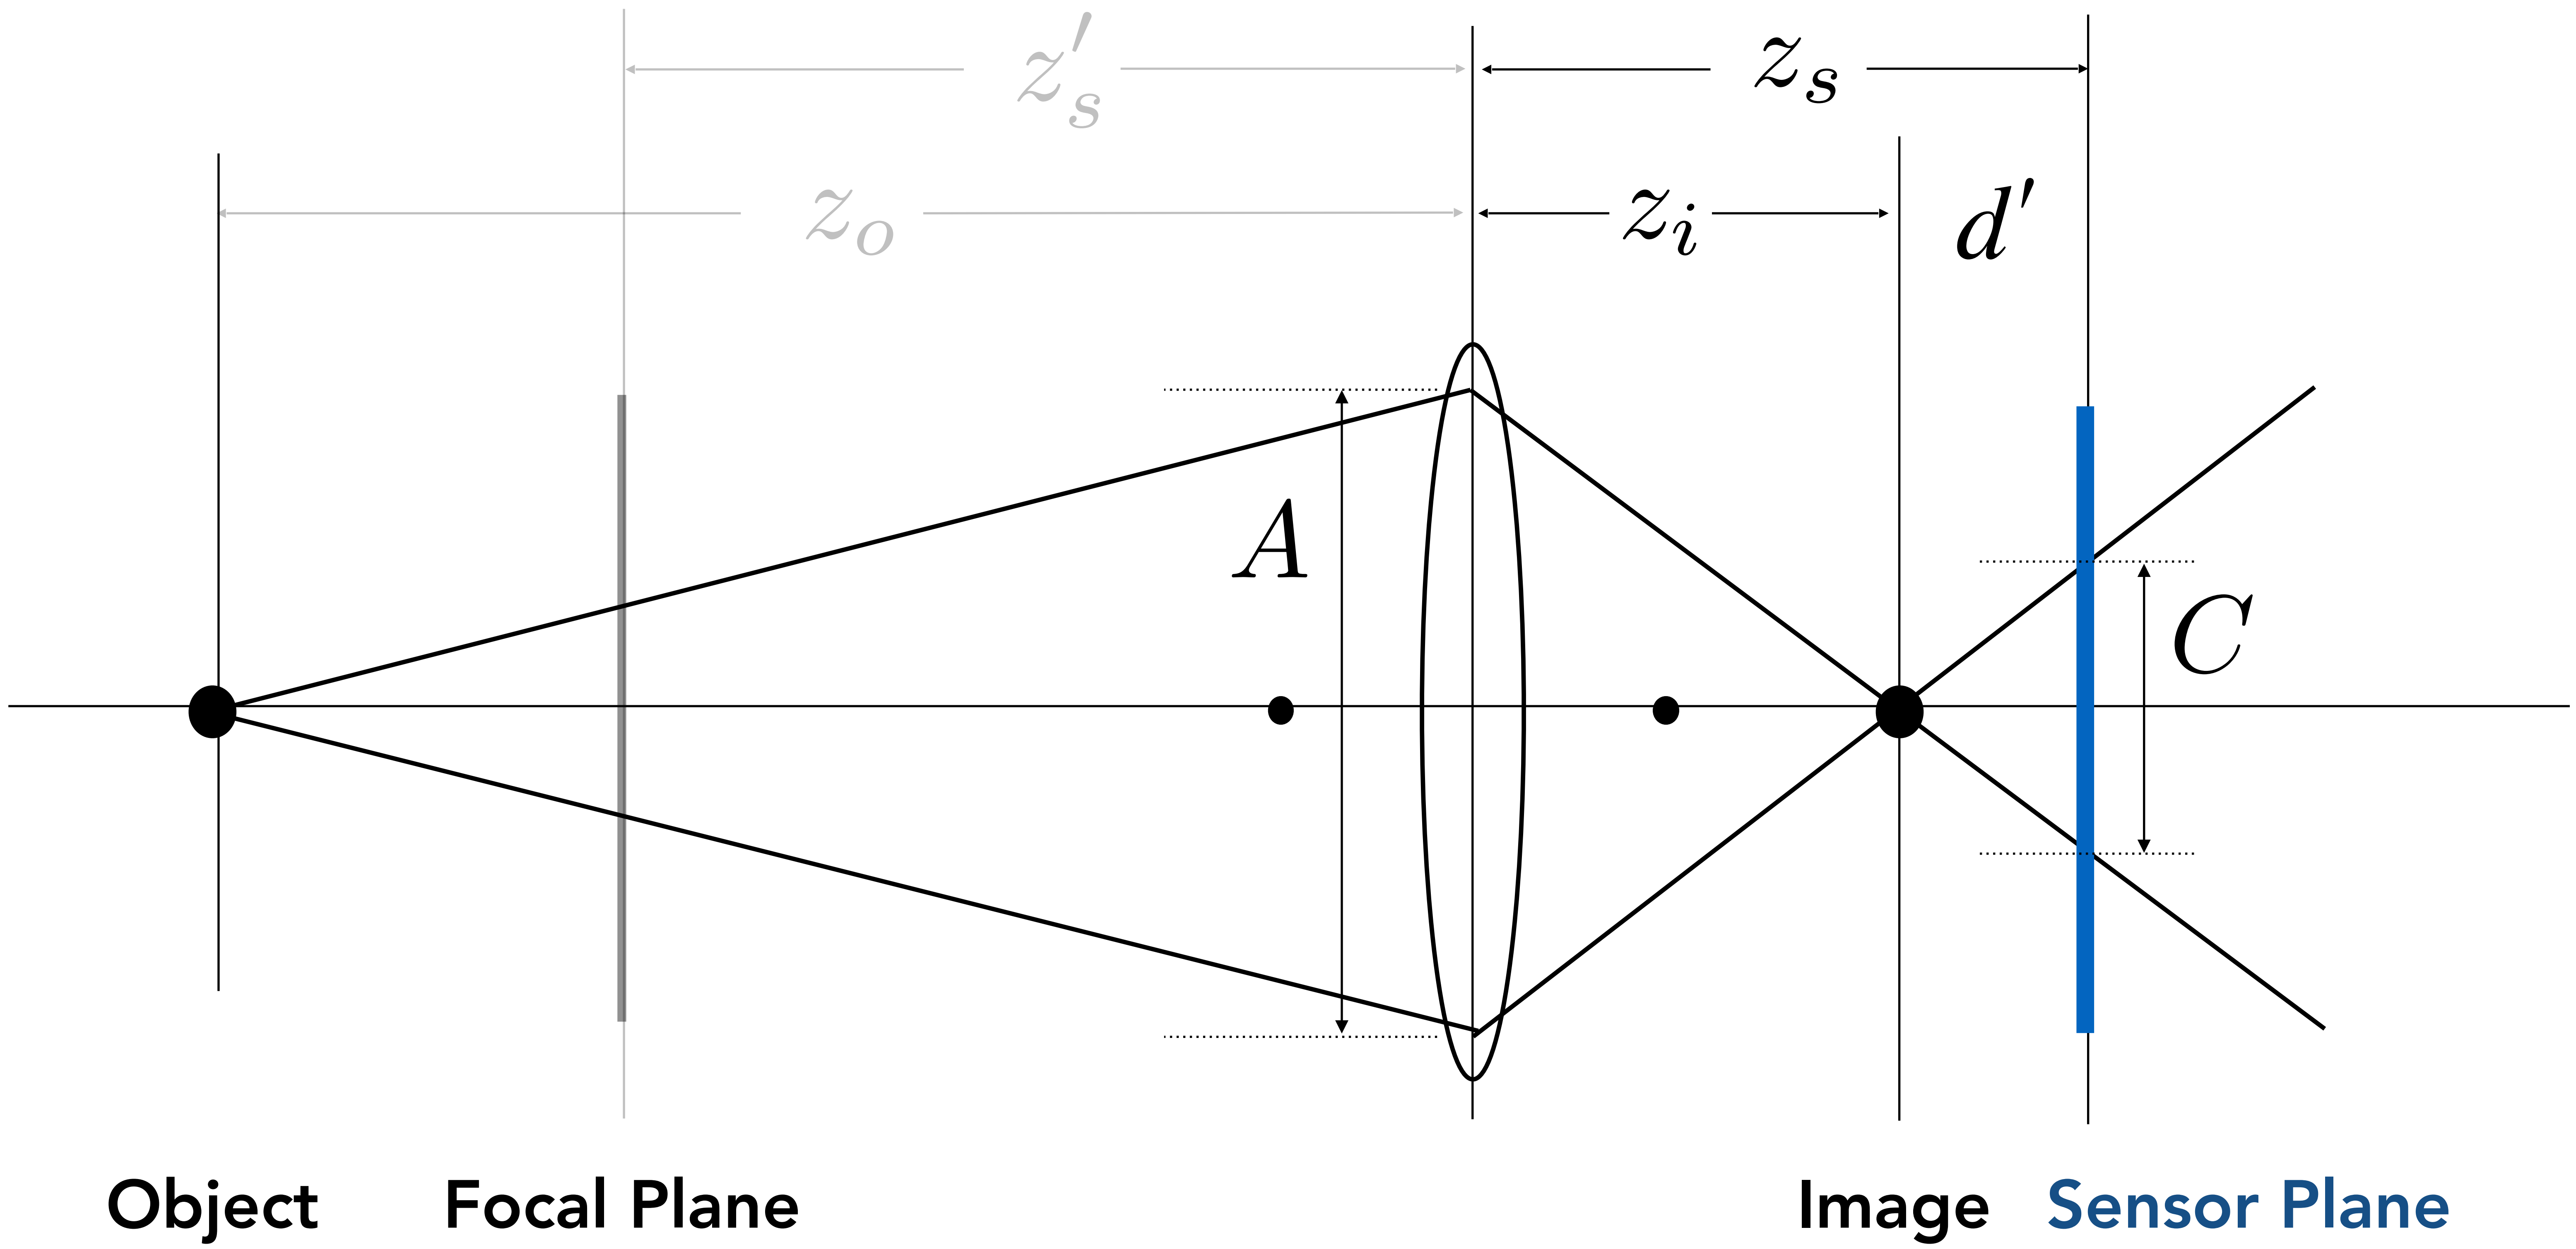

Circle of confusion is proportional to the size of the aperture

$$\frac{C}{A} = \frac{d'}{z_i} = \frac{|z_s - z_i|}{z_i}$$

$$C = A\frac{|z_s - z_i|}{z_i}$$

Bokeh

Photo credit: Wikimedia's The Photographer

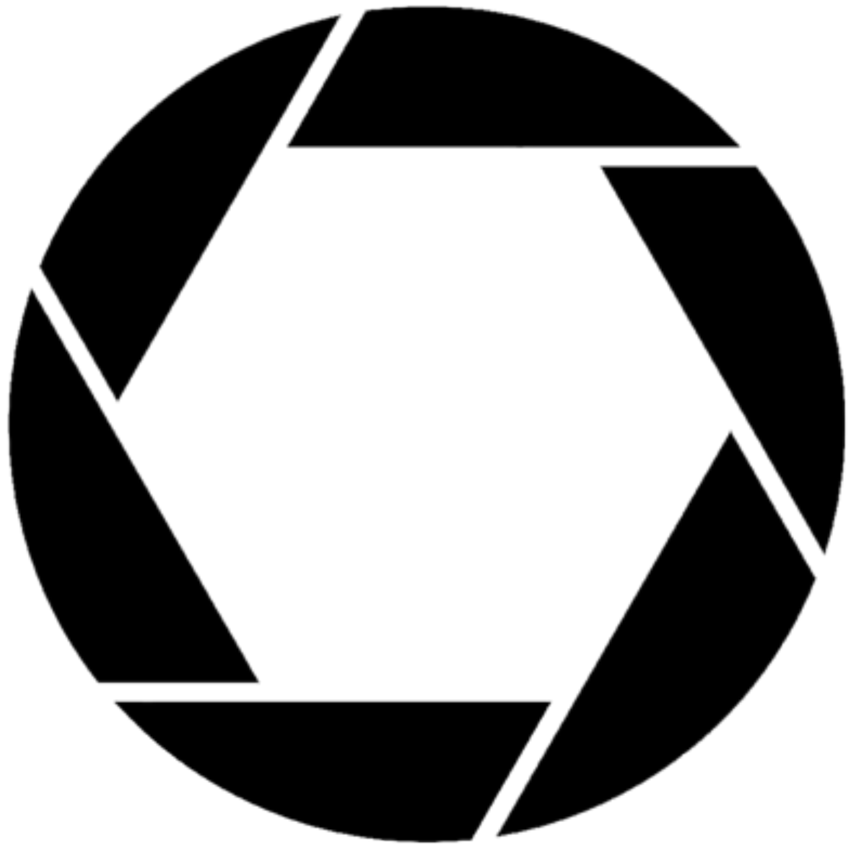

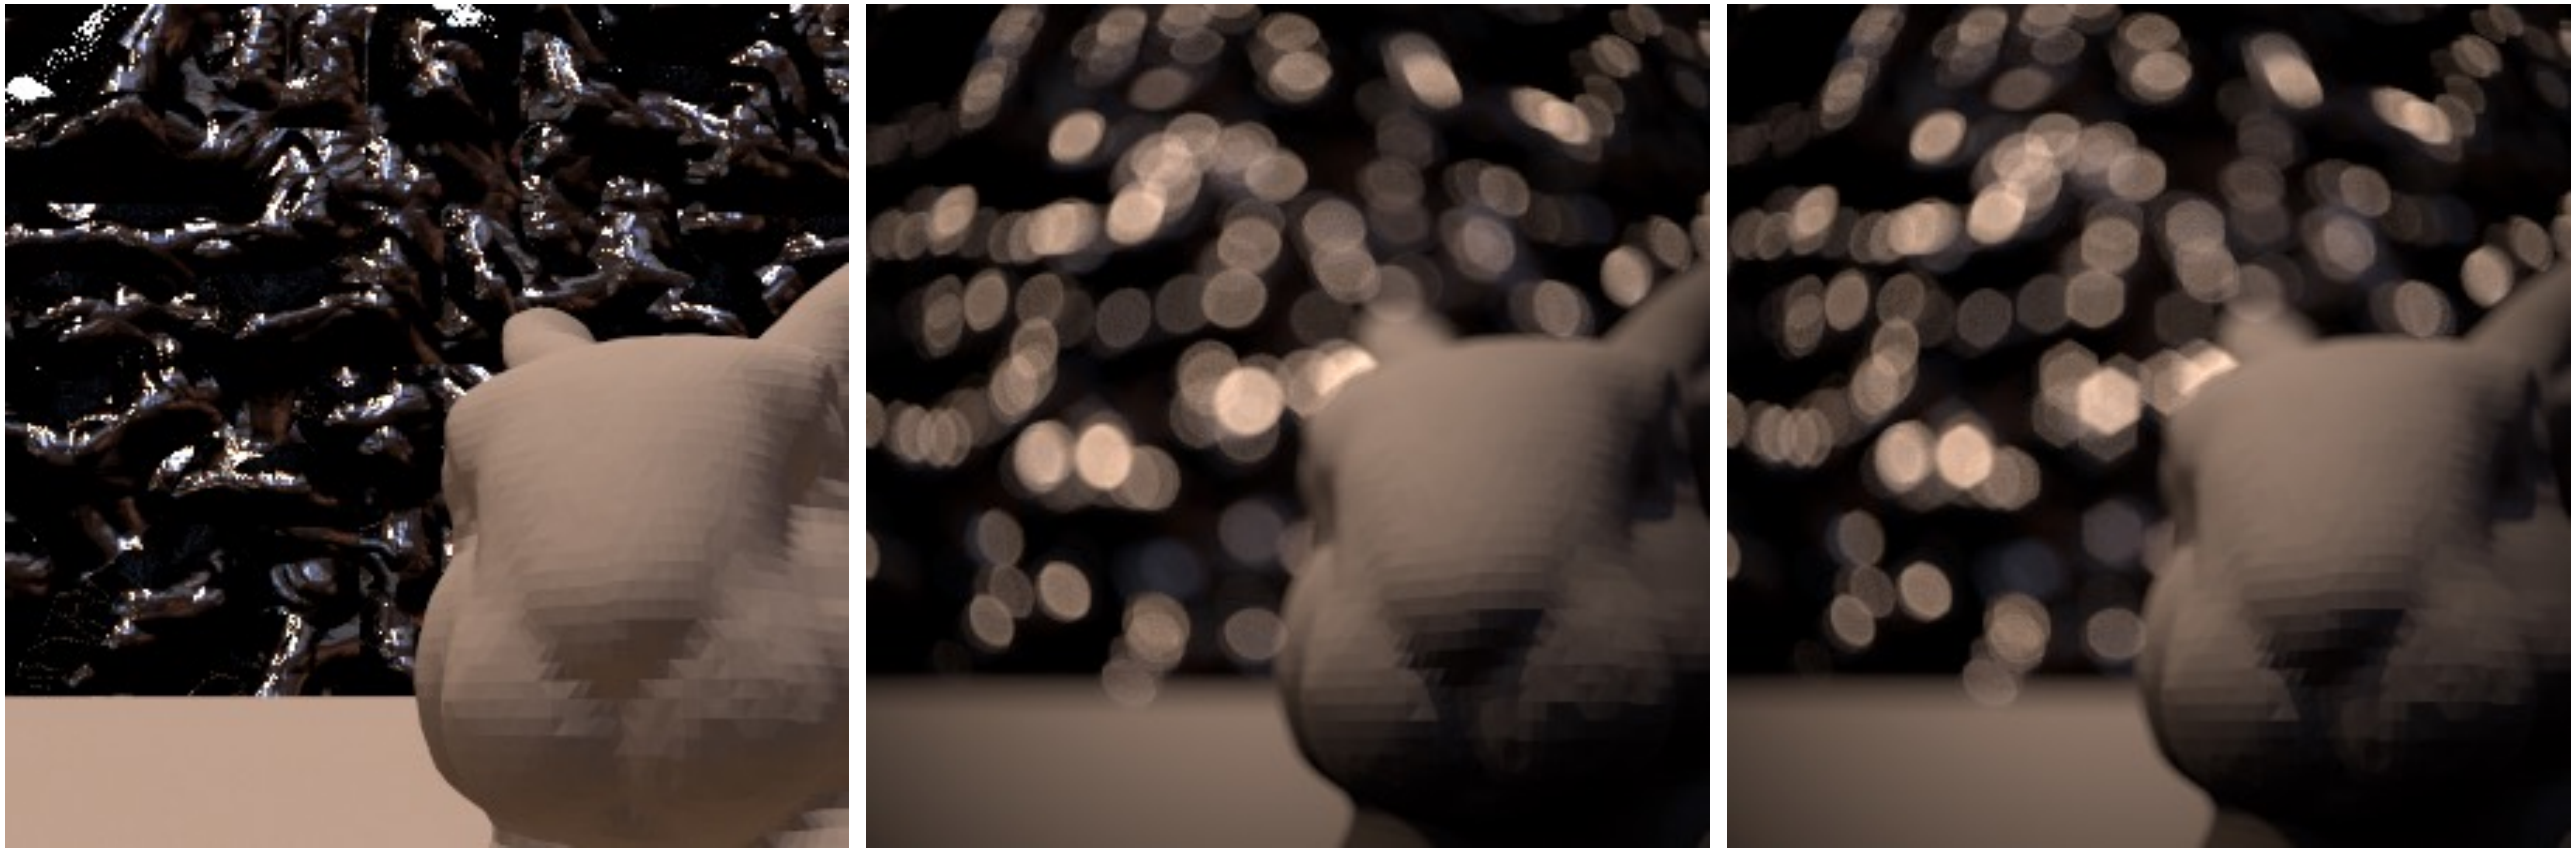

Bokeh

Bokeh is the shape and quality of out-of-focus blur

- **For small, out-of-focus lights, bokeh takes on the shape of the lens aperture**

M Yashna, flickr, 40mm f/3.0

Bokeh

Dino Quinzani, Leica Noctilux 50mm, f/0.95

Why does the bokeh vary across the image?

F Number

A lens' F-number is the ratio of focal length (f) to the maximum aperture diameter (A) of the lens

Example: two lenses with the same F-number (Denoted as F/2 or F:2)

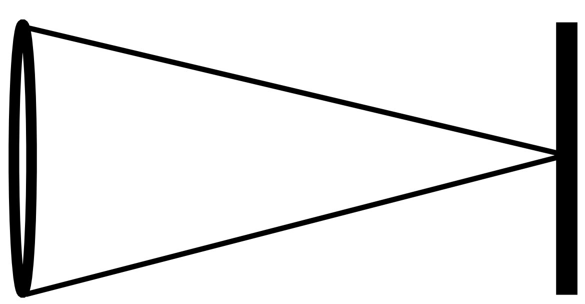

$$A = 50 \text{ mm}$$
$$f = 100 \text{ mm}$$
$$N = f/A = 2$$

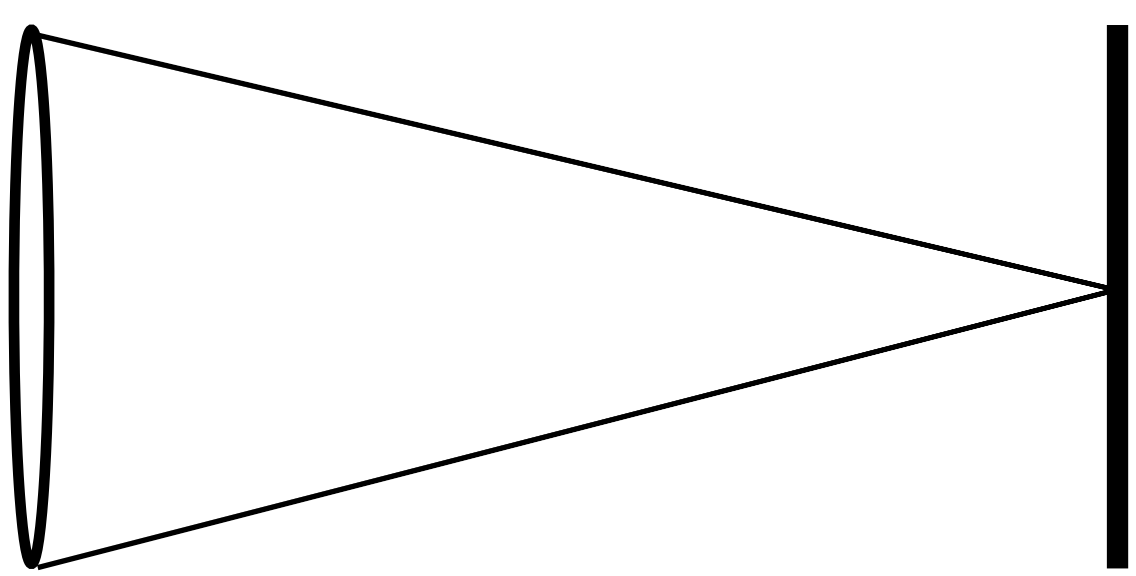

$$A = 100 \text{ mm}$$
$$f = 200 \text{ mm}$$
$$N = f/A = 2$$

Lens's F-Number vs F-Number for Photo

But for an individual photo, the lens aperture may be "stopped down" to a smaller size

- E.g. 50 mm F/1.4 lens stopped down to F/4
- Aperture is closed down with an iris to 50/4 = 12.5 mm

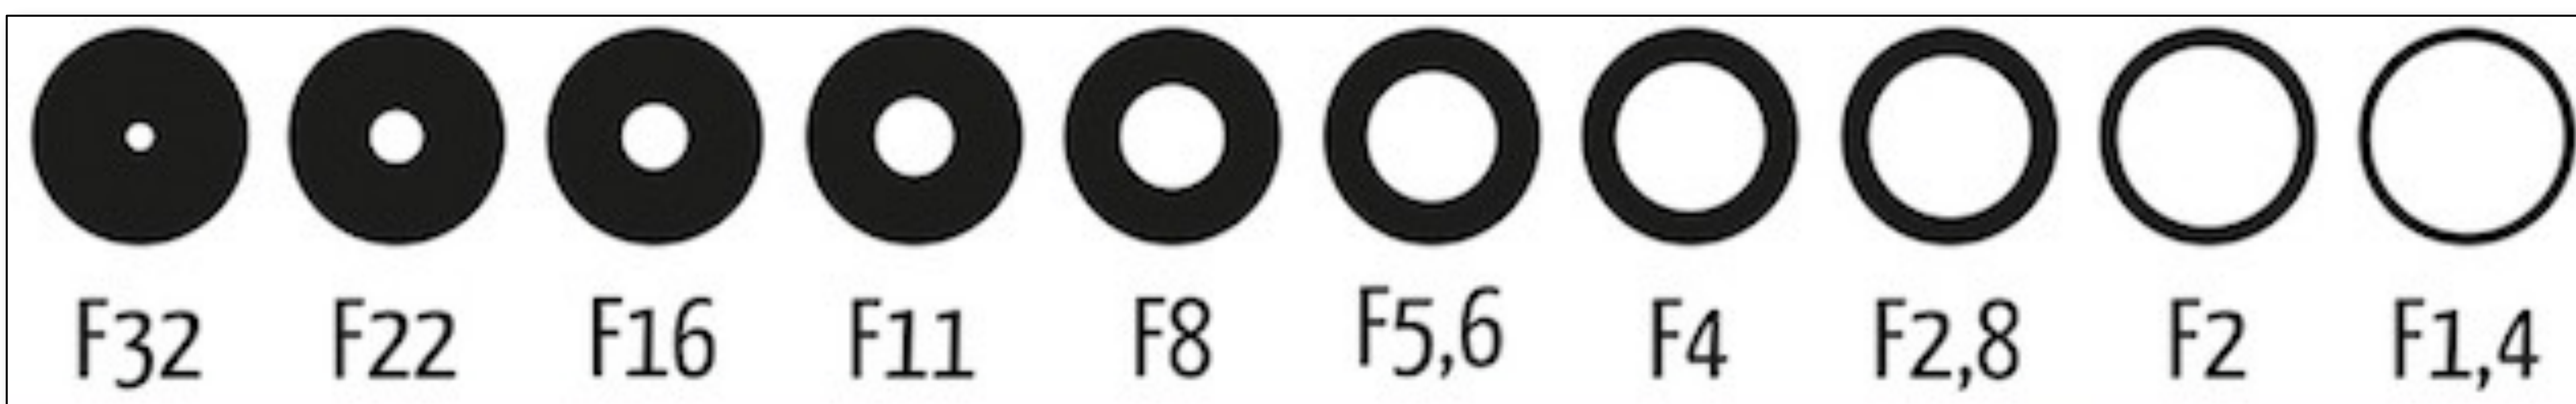

Size of Circle of Confusion is Inversely Proportional to F-Number for Photo

R. Berdan, canadiannaturephotographer.com

$$C = A\frac{|z_s - z_i|}{z_i} = \frac{f}{N}\frac{|z_s - z_i|}{z_i}$$

Ray Tracing Ideal Thin Lenses

Rendering with Lens Focus

Pharr and Humphreys

Ray Tracing for Defocus Blur (Thin Lens)

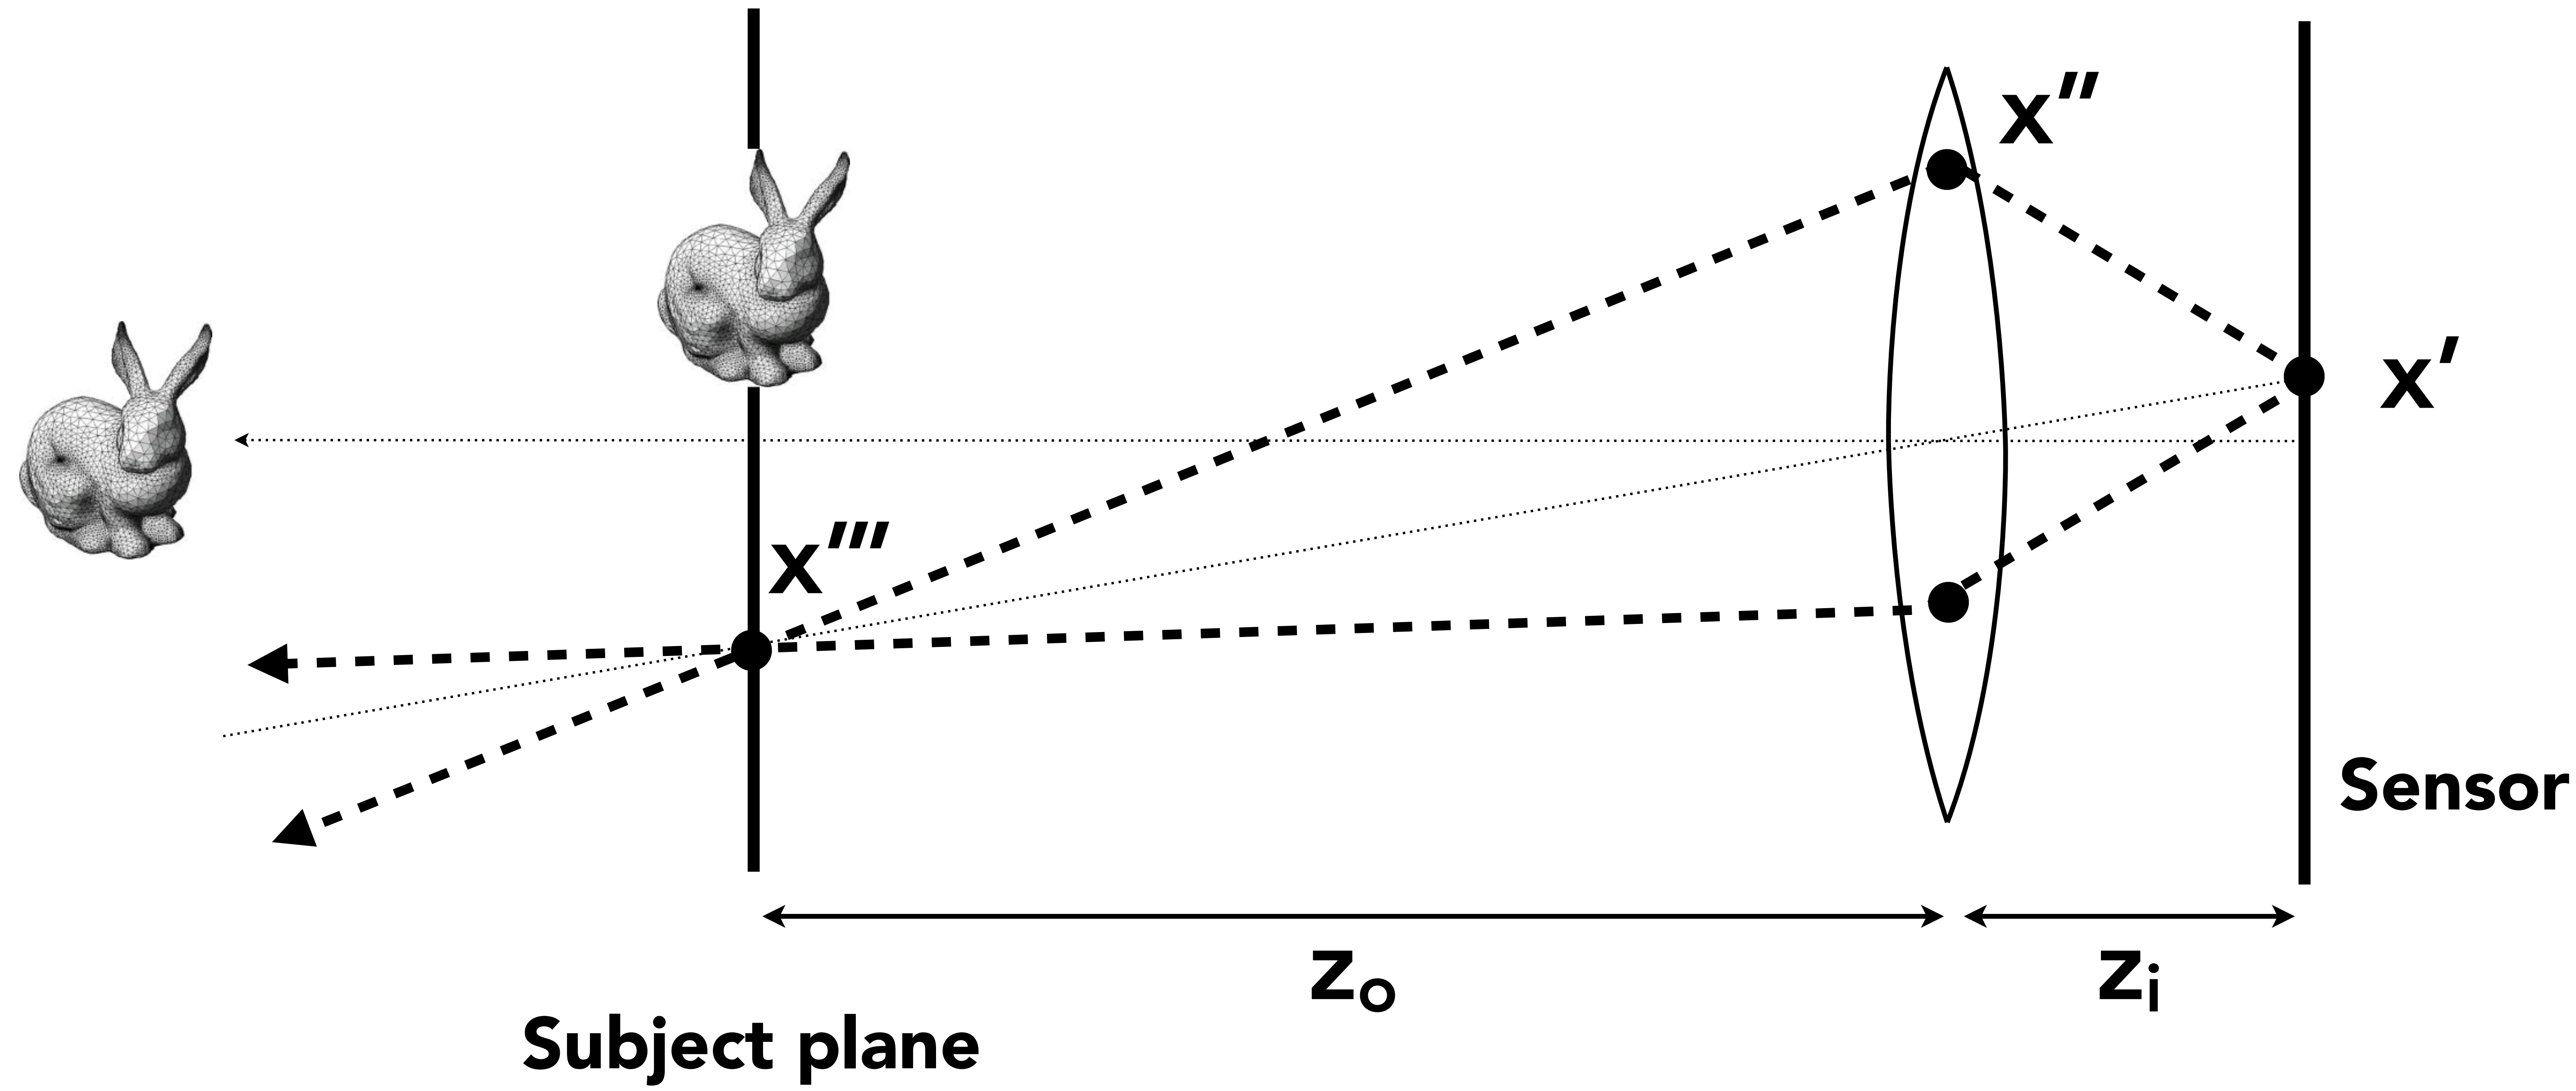

Setup:

- **Choose sensor size, lens focal length and aperture size**
- **Choose depth of subject of interest z_o**
- **Calculate corresponding depth of sensor z_i from thin lens equation (focusing)**

Ray Tracing for Defocus Blur (Thin Lens)

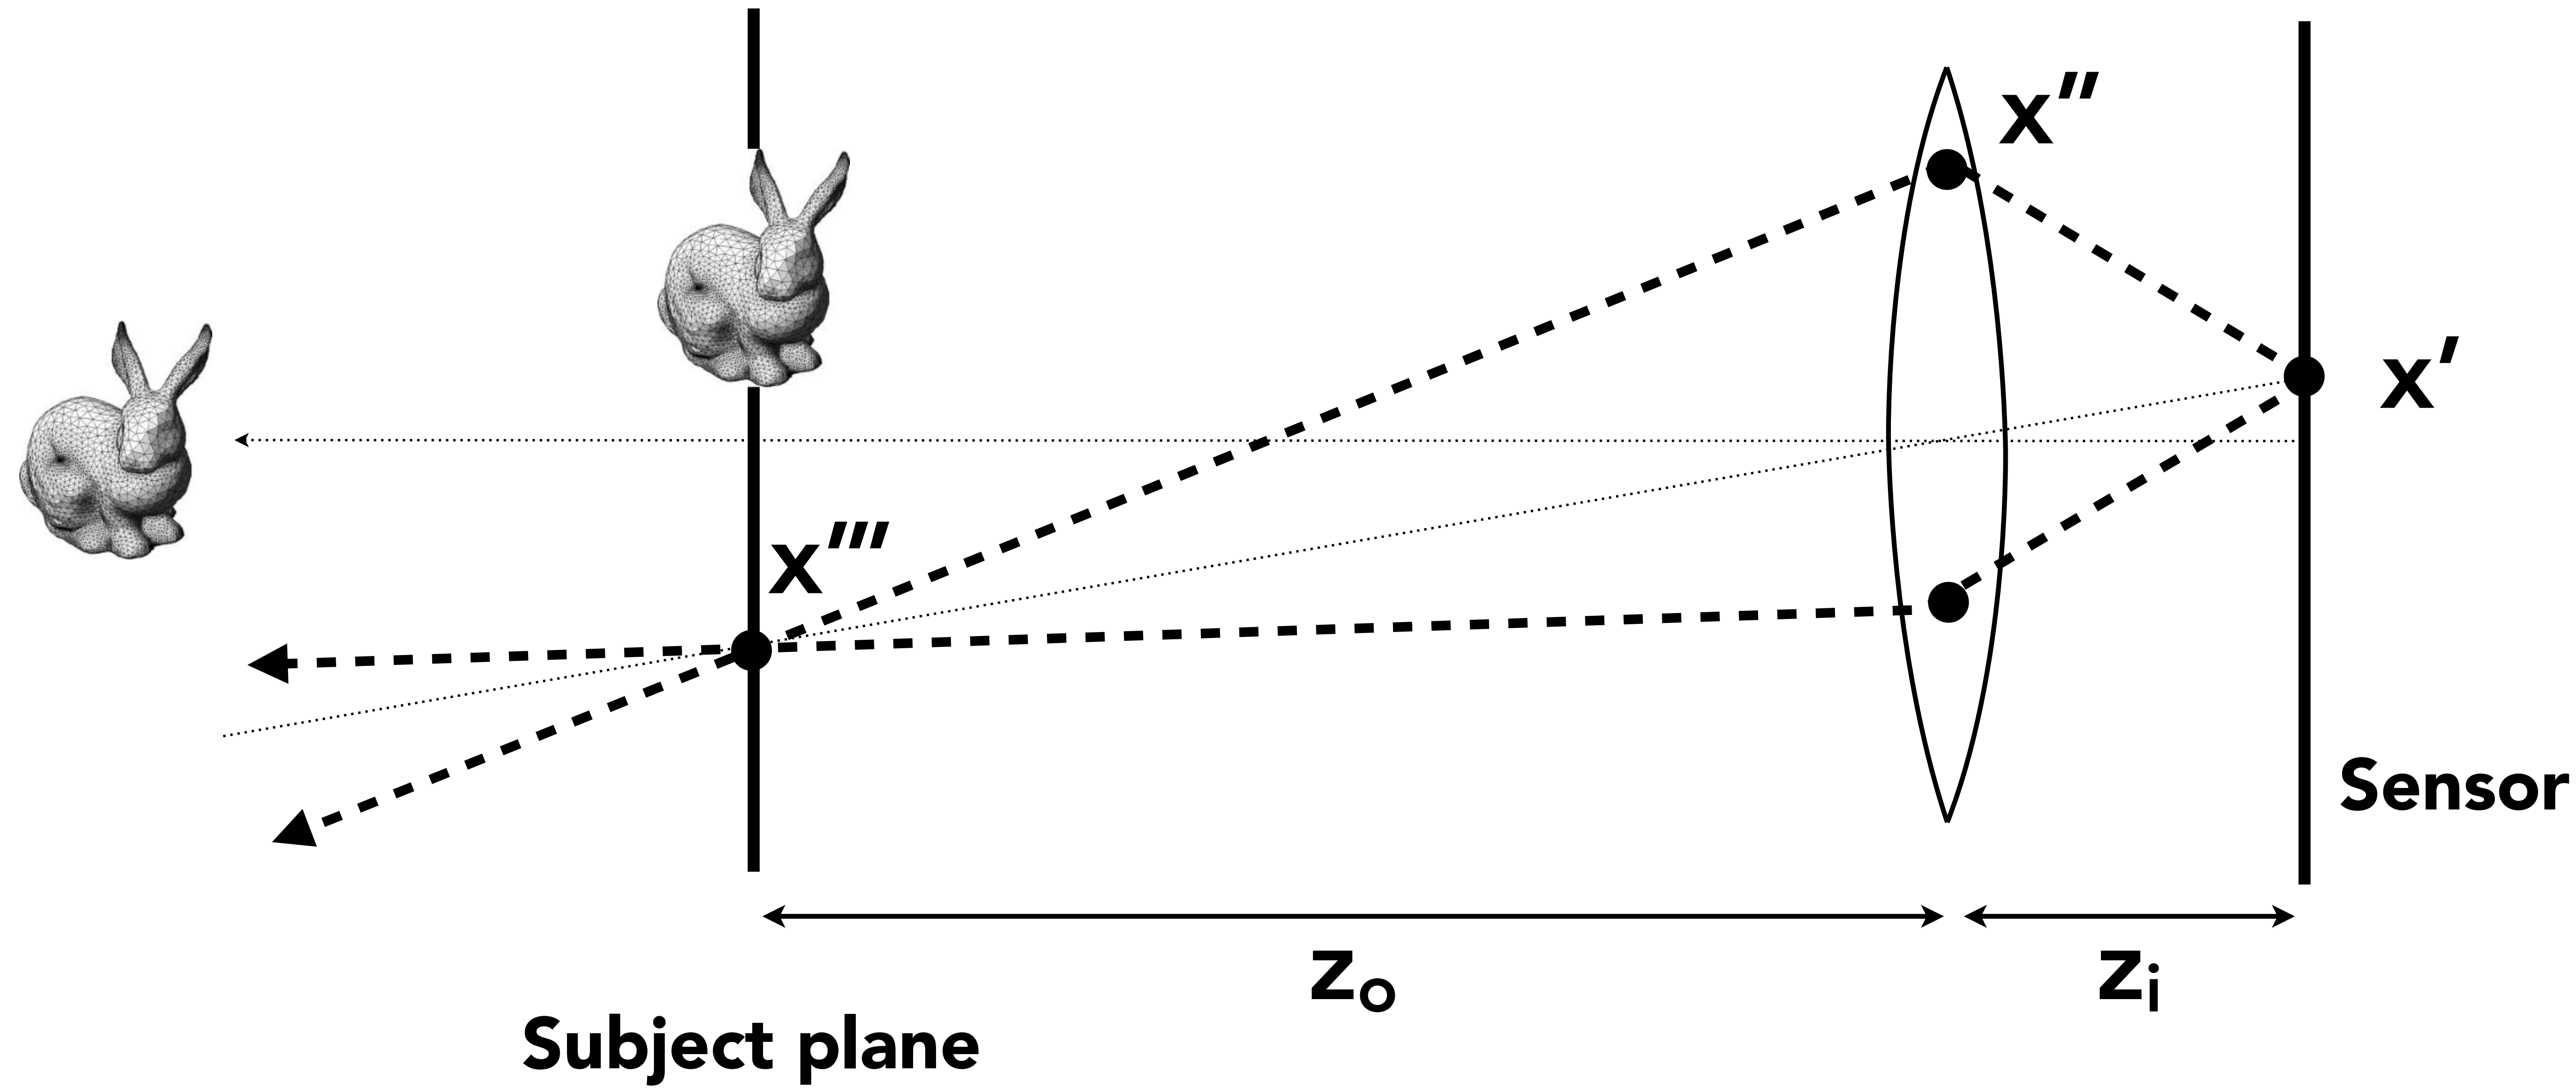

To compute value of pixel at position x′ by Monte Carlo integration:

- **Select random points x′′ on lens plane**
- **Rays pass from point x′ on image plane z_i through points x′′ on lens**
- **Each ray passes through conjugate point x′′′ on the plane of focus z_o**
  Can determine x′′′ from Gauss′ ray diagram
  So just trace ray from x′′ to x′′′
- **Estimate radiance on rays using path-tracing, and sum over all points x′′**

Real Compound Lenses

Modern Lens Designs Are Complex

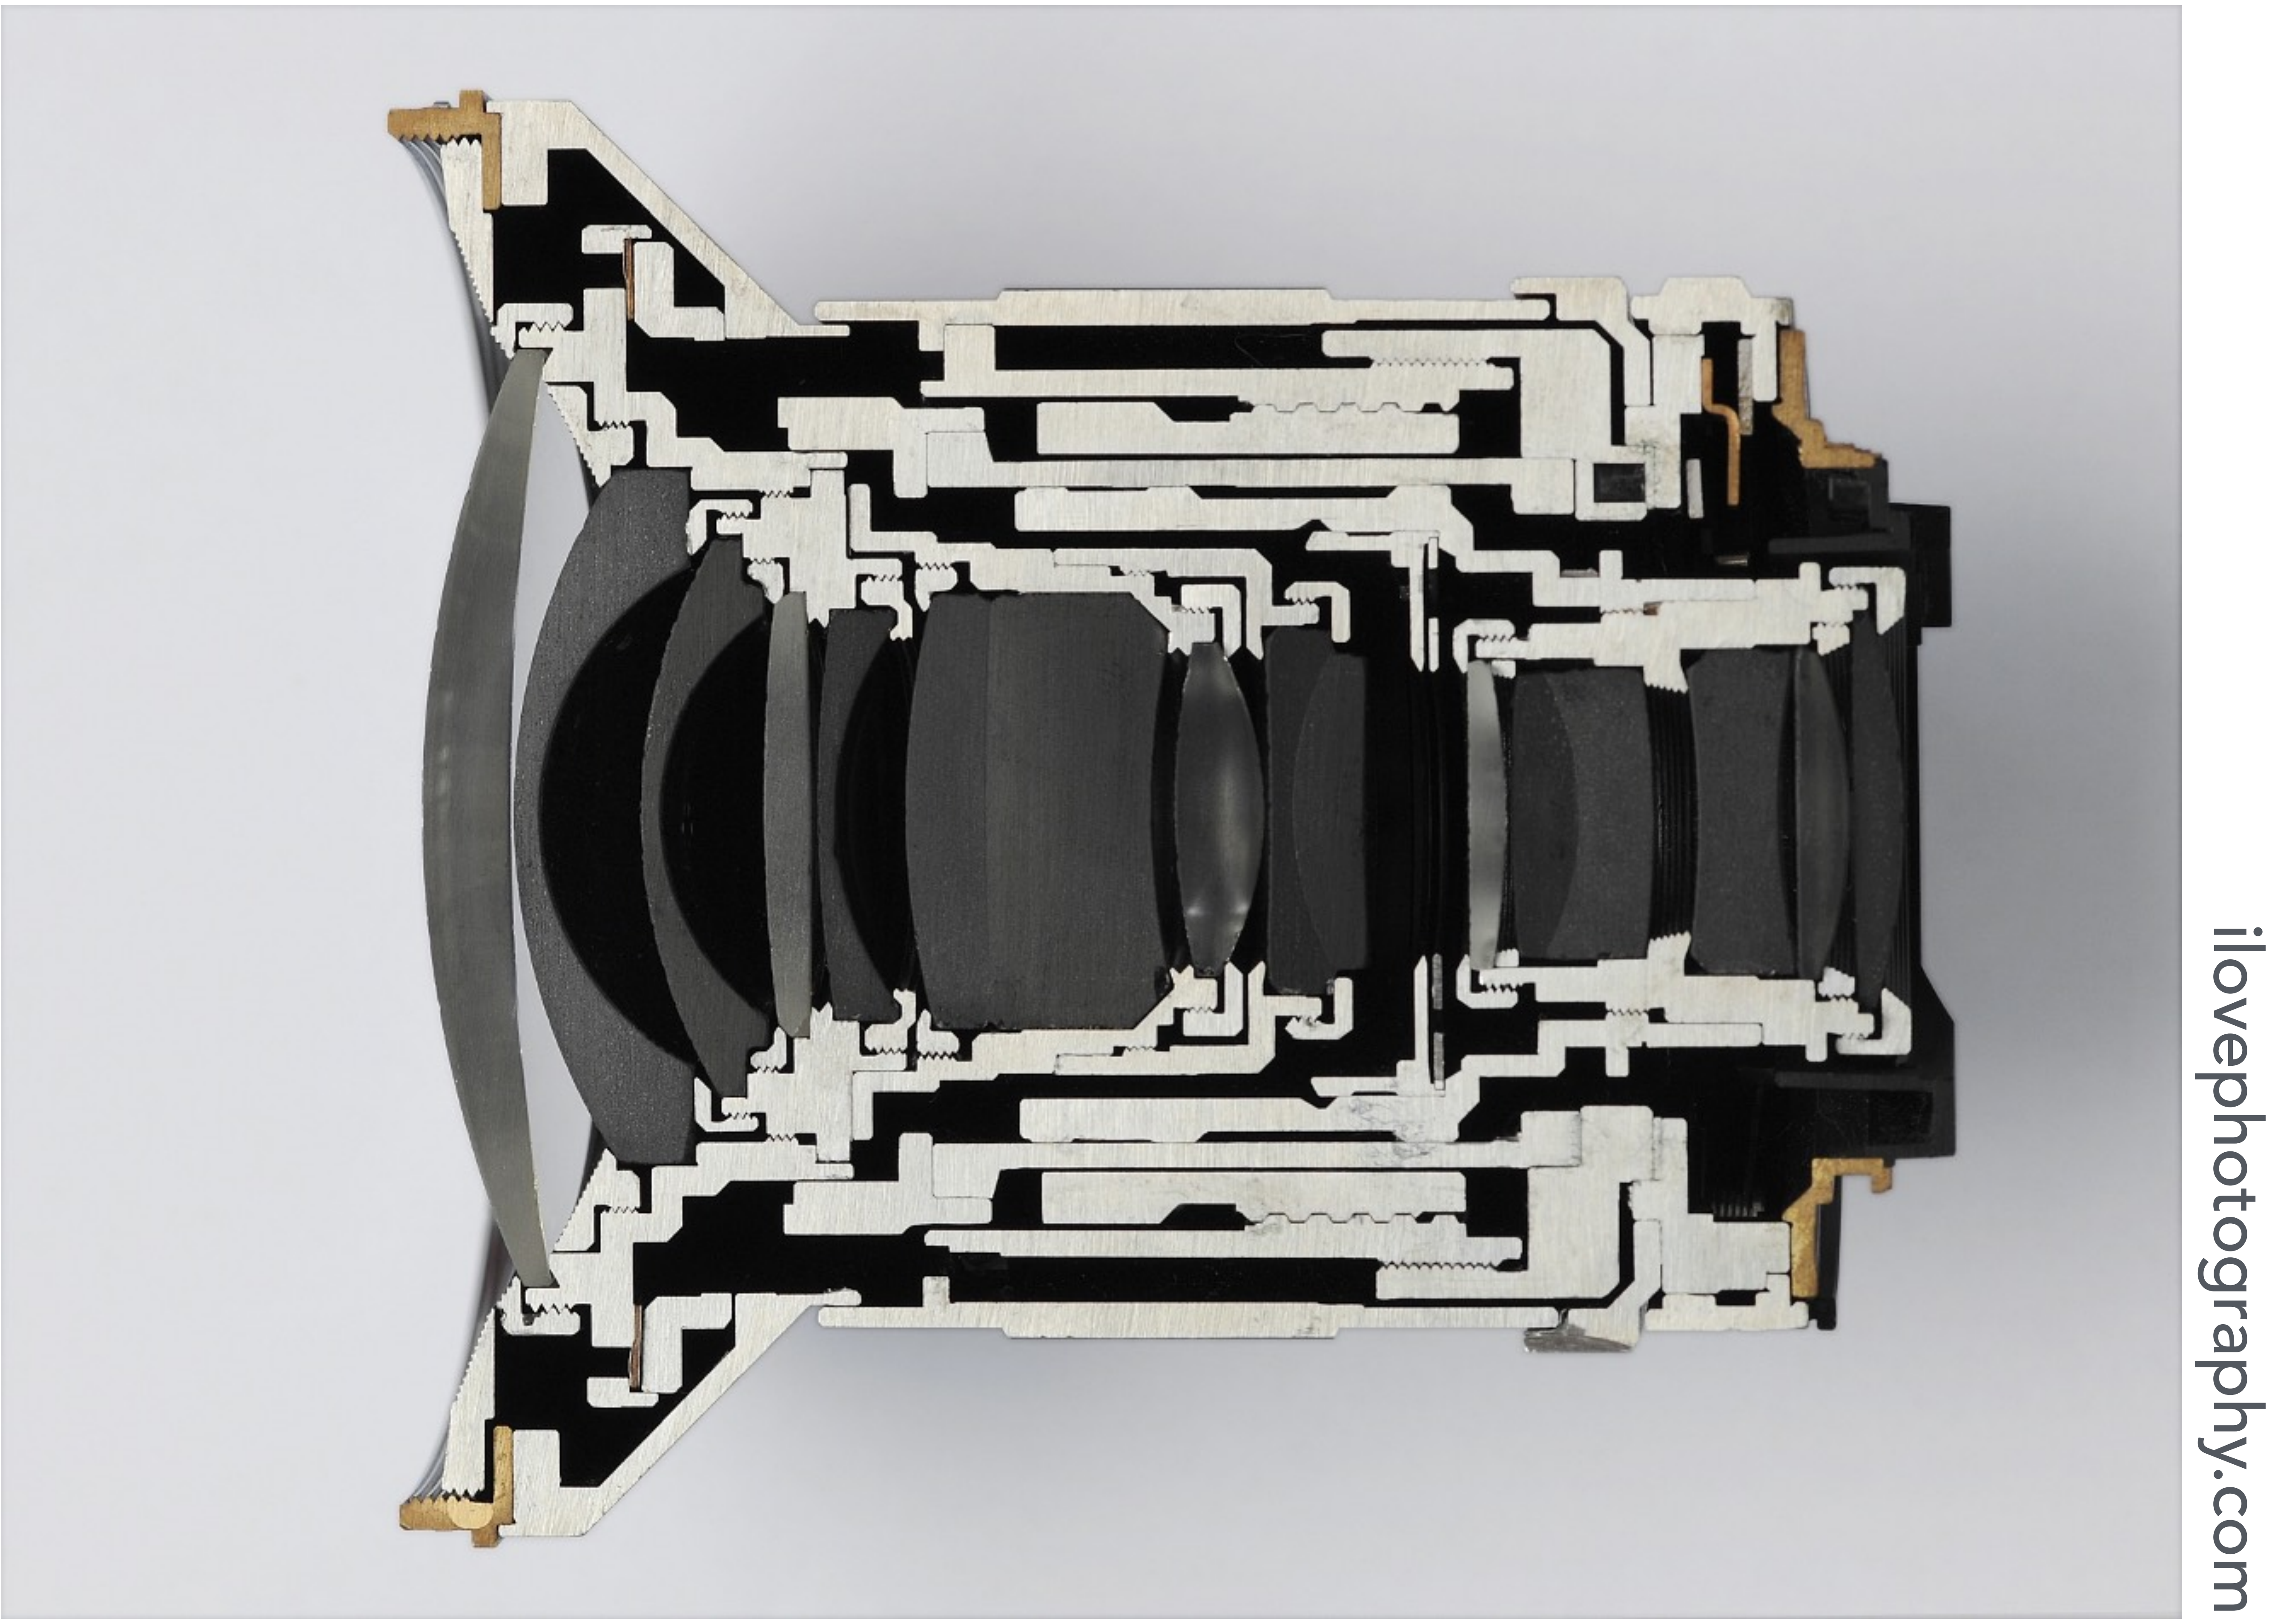

Photographic lens cross section

Modern Lens Designs Are Complex

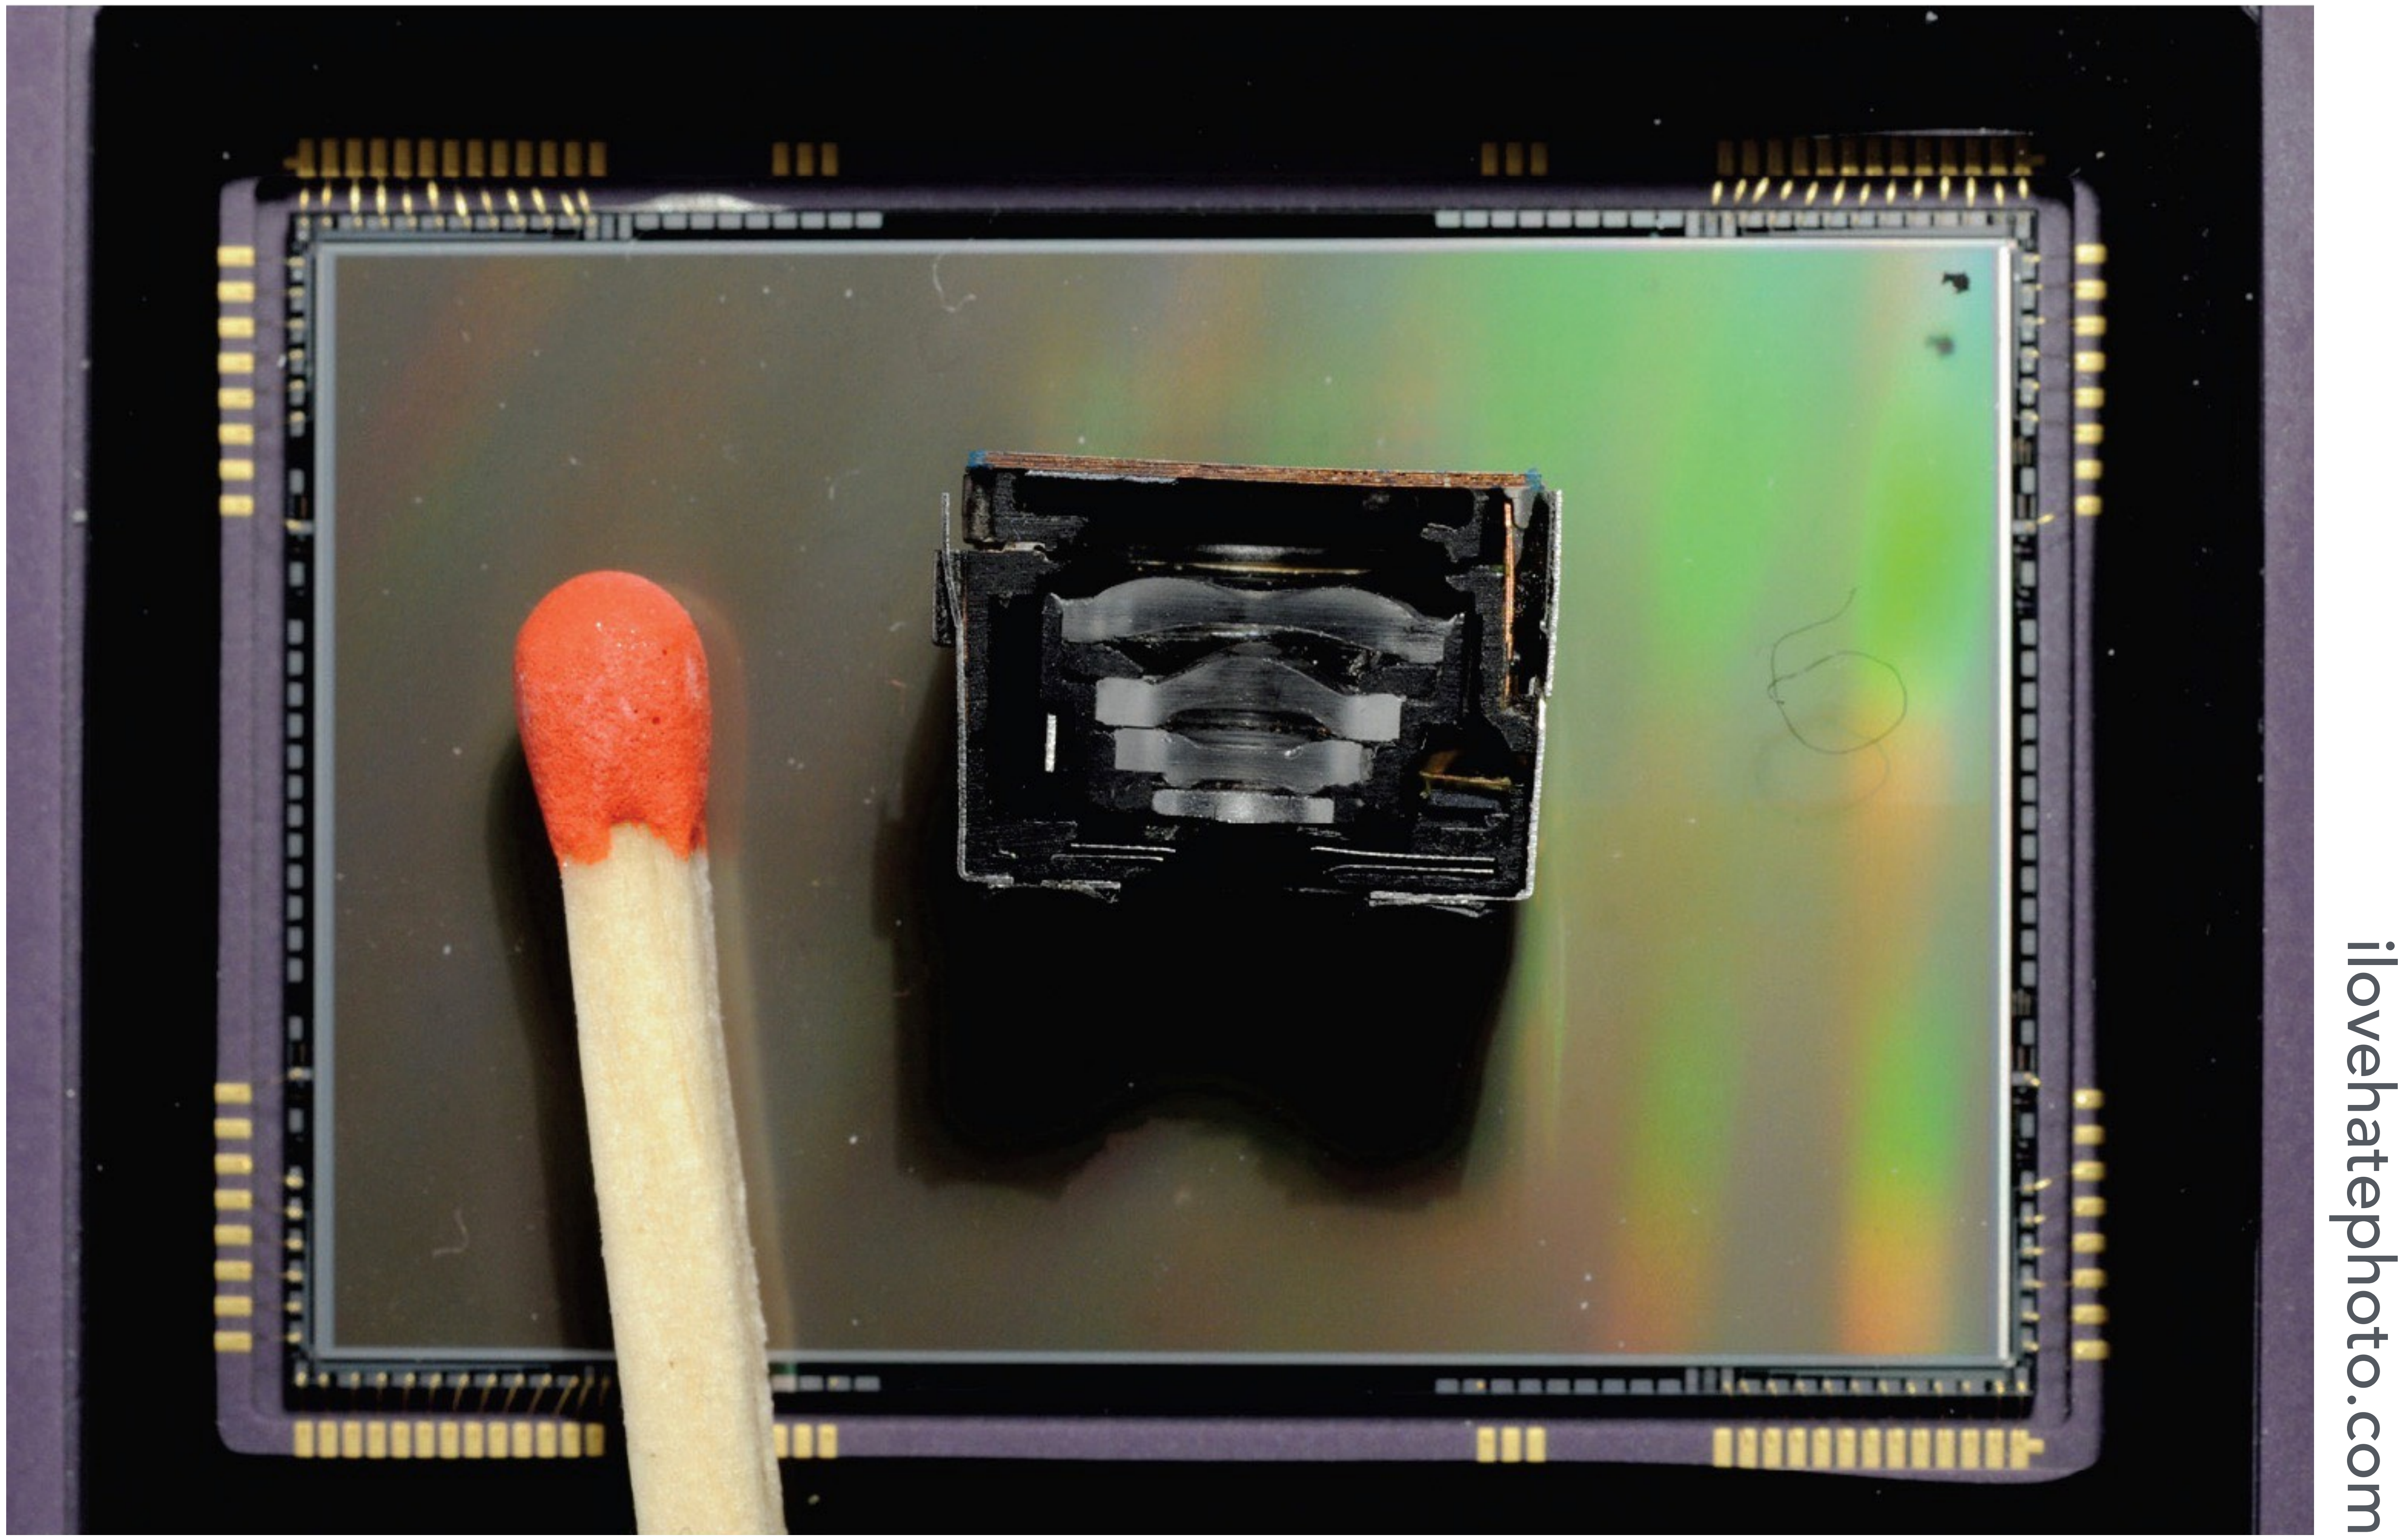

4 element mobile phone lens (on 24x36mm sensor)

Real Lenses vs Ideal Thin Lenses

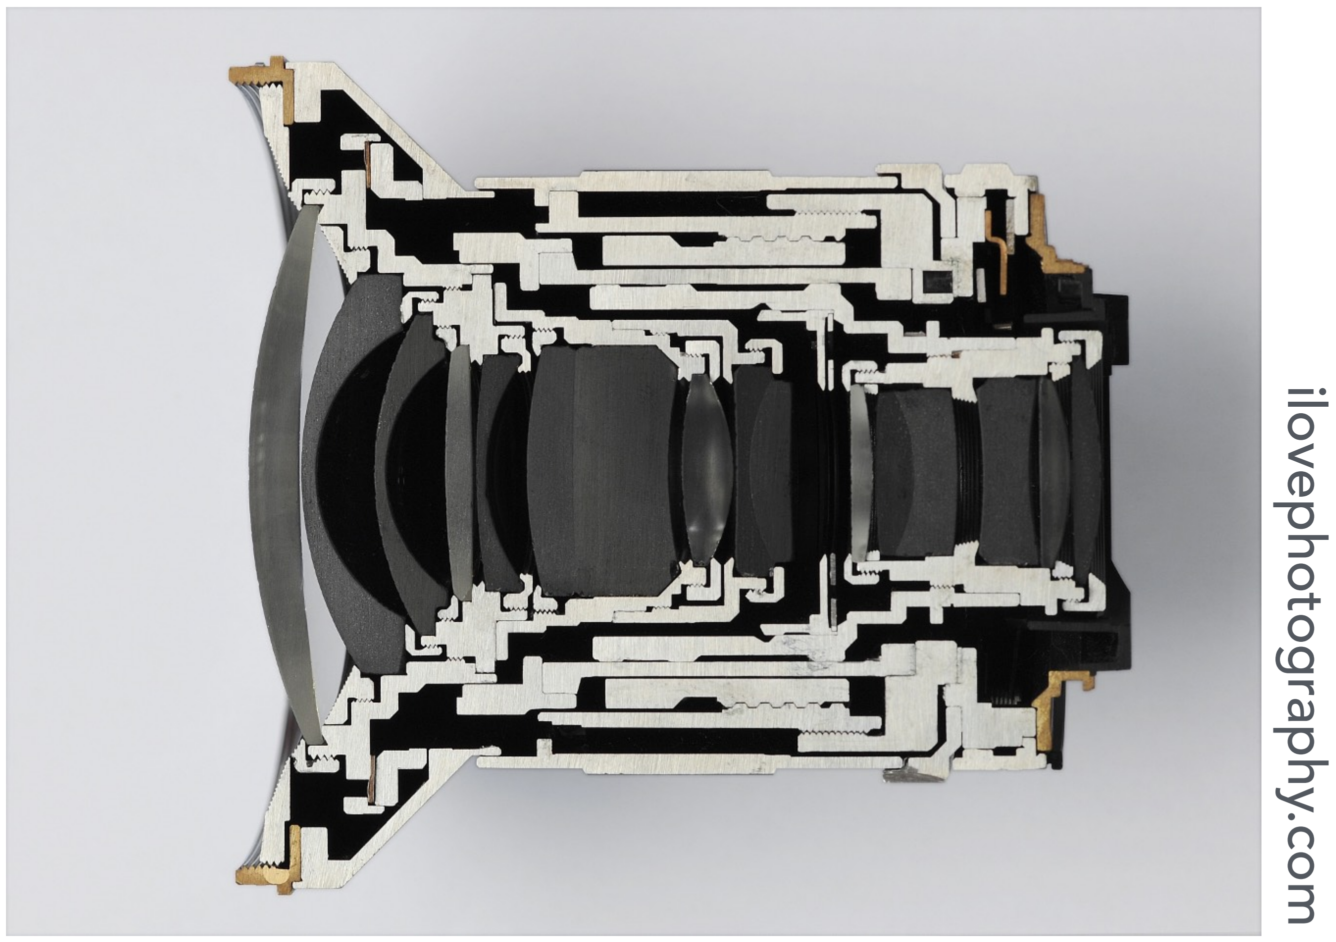

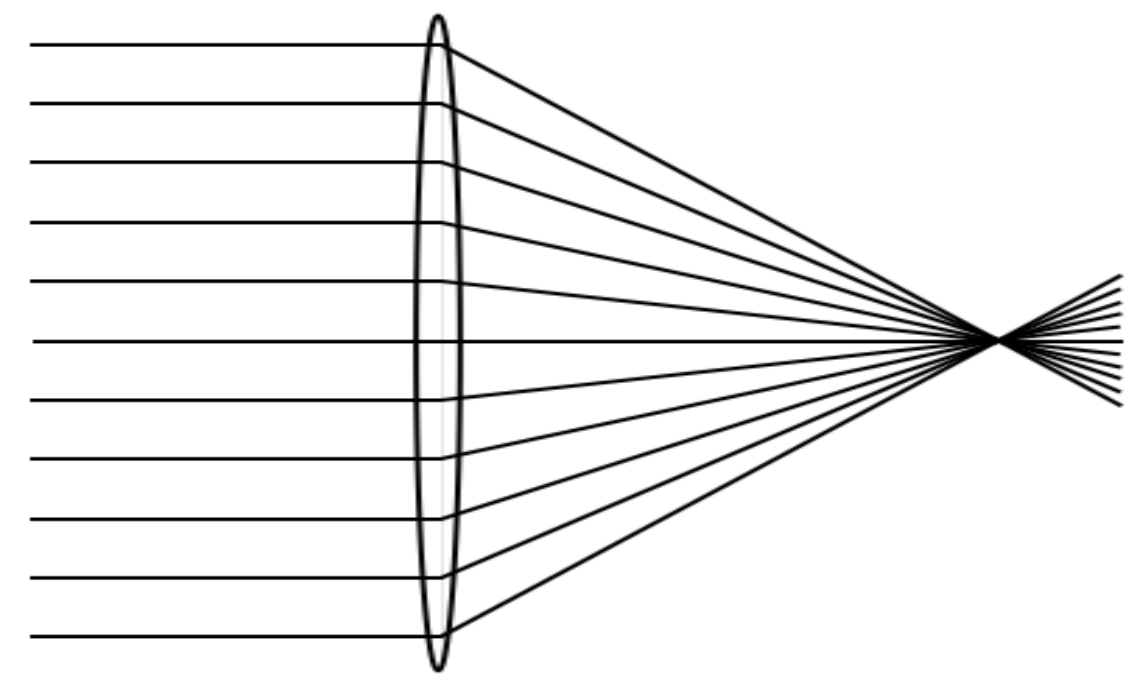

- Real optical system
- Multiple physical elements in compound design
- Optical aberrations prevent rays from converging perfectly

- Theoretical abstraction
- Assume all rays refract at a plane & converge to a point
- Quick and intuitive calculation of main imaging effects

Double Gauss

Data from W. Smith, Modern Lens Design, p 312

Radius	Thick	nd	V-no	aperture
58.950	7.520	1.670	47.1	50.4
169.660	0.240			50.4
38.550	8.050	1.670	47.1	46.0
81.540	6.550	1.699	30.1	46.0
25.500	11.410			36.0
	9.000			34.2
-28.990	2.360	1.603	38.0	34.0
81.540	12.130	1.658	57.3	40.0
-40.770	0.380			40.0
874.130	6.440	1.717	48.0	40.0
-79.460	72.228			40.0

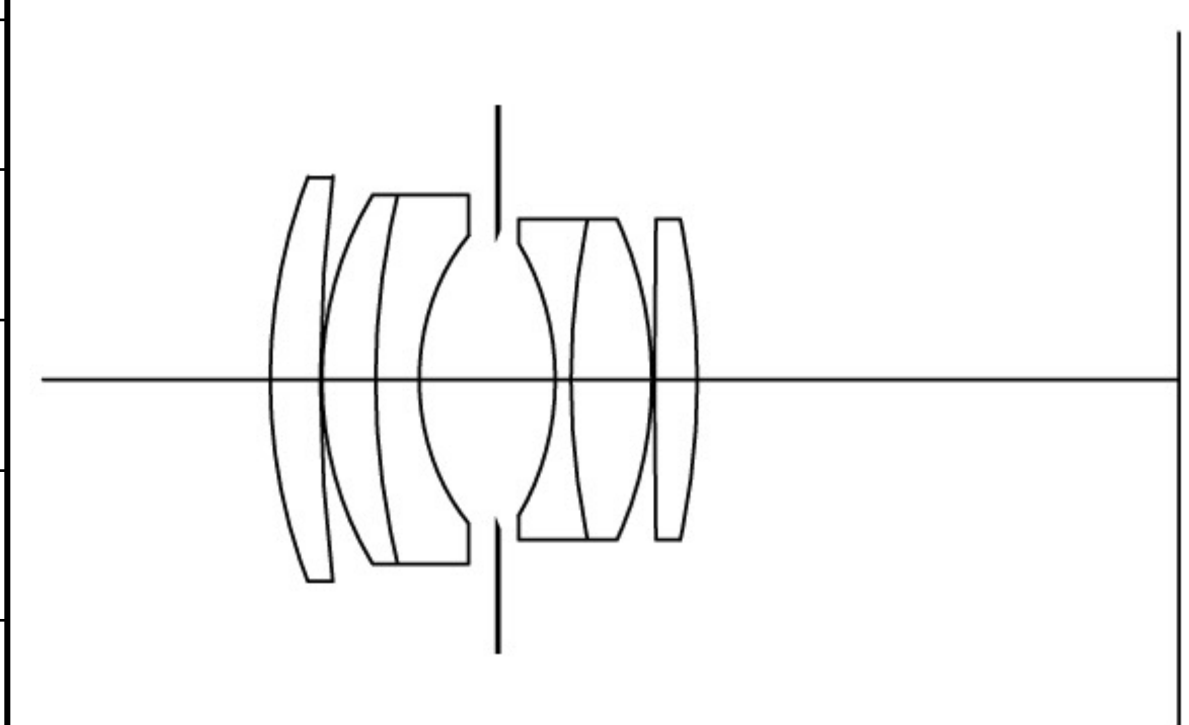

Refraction

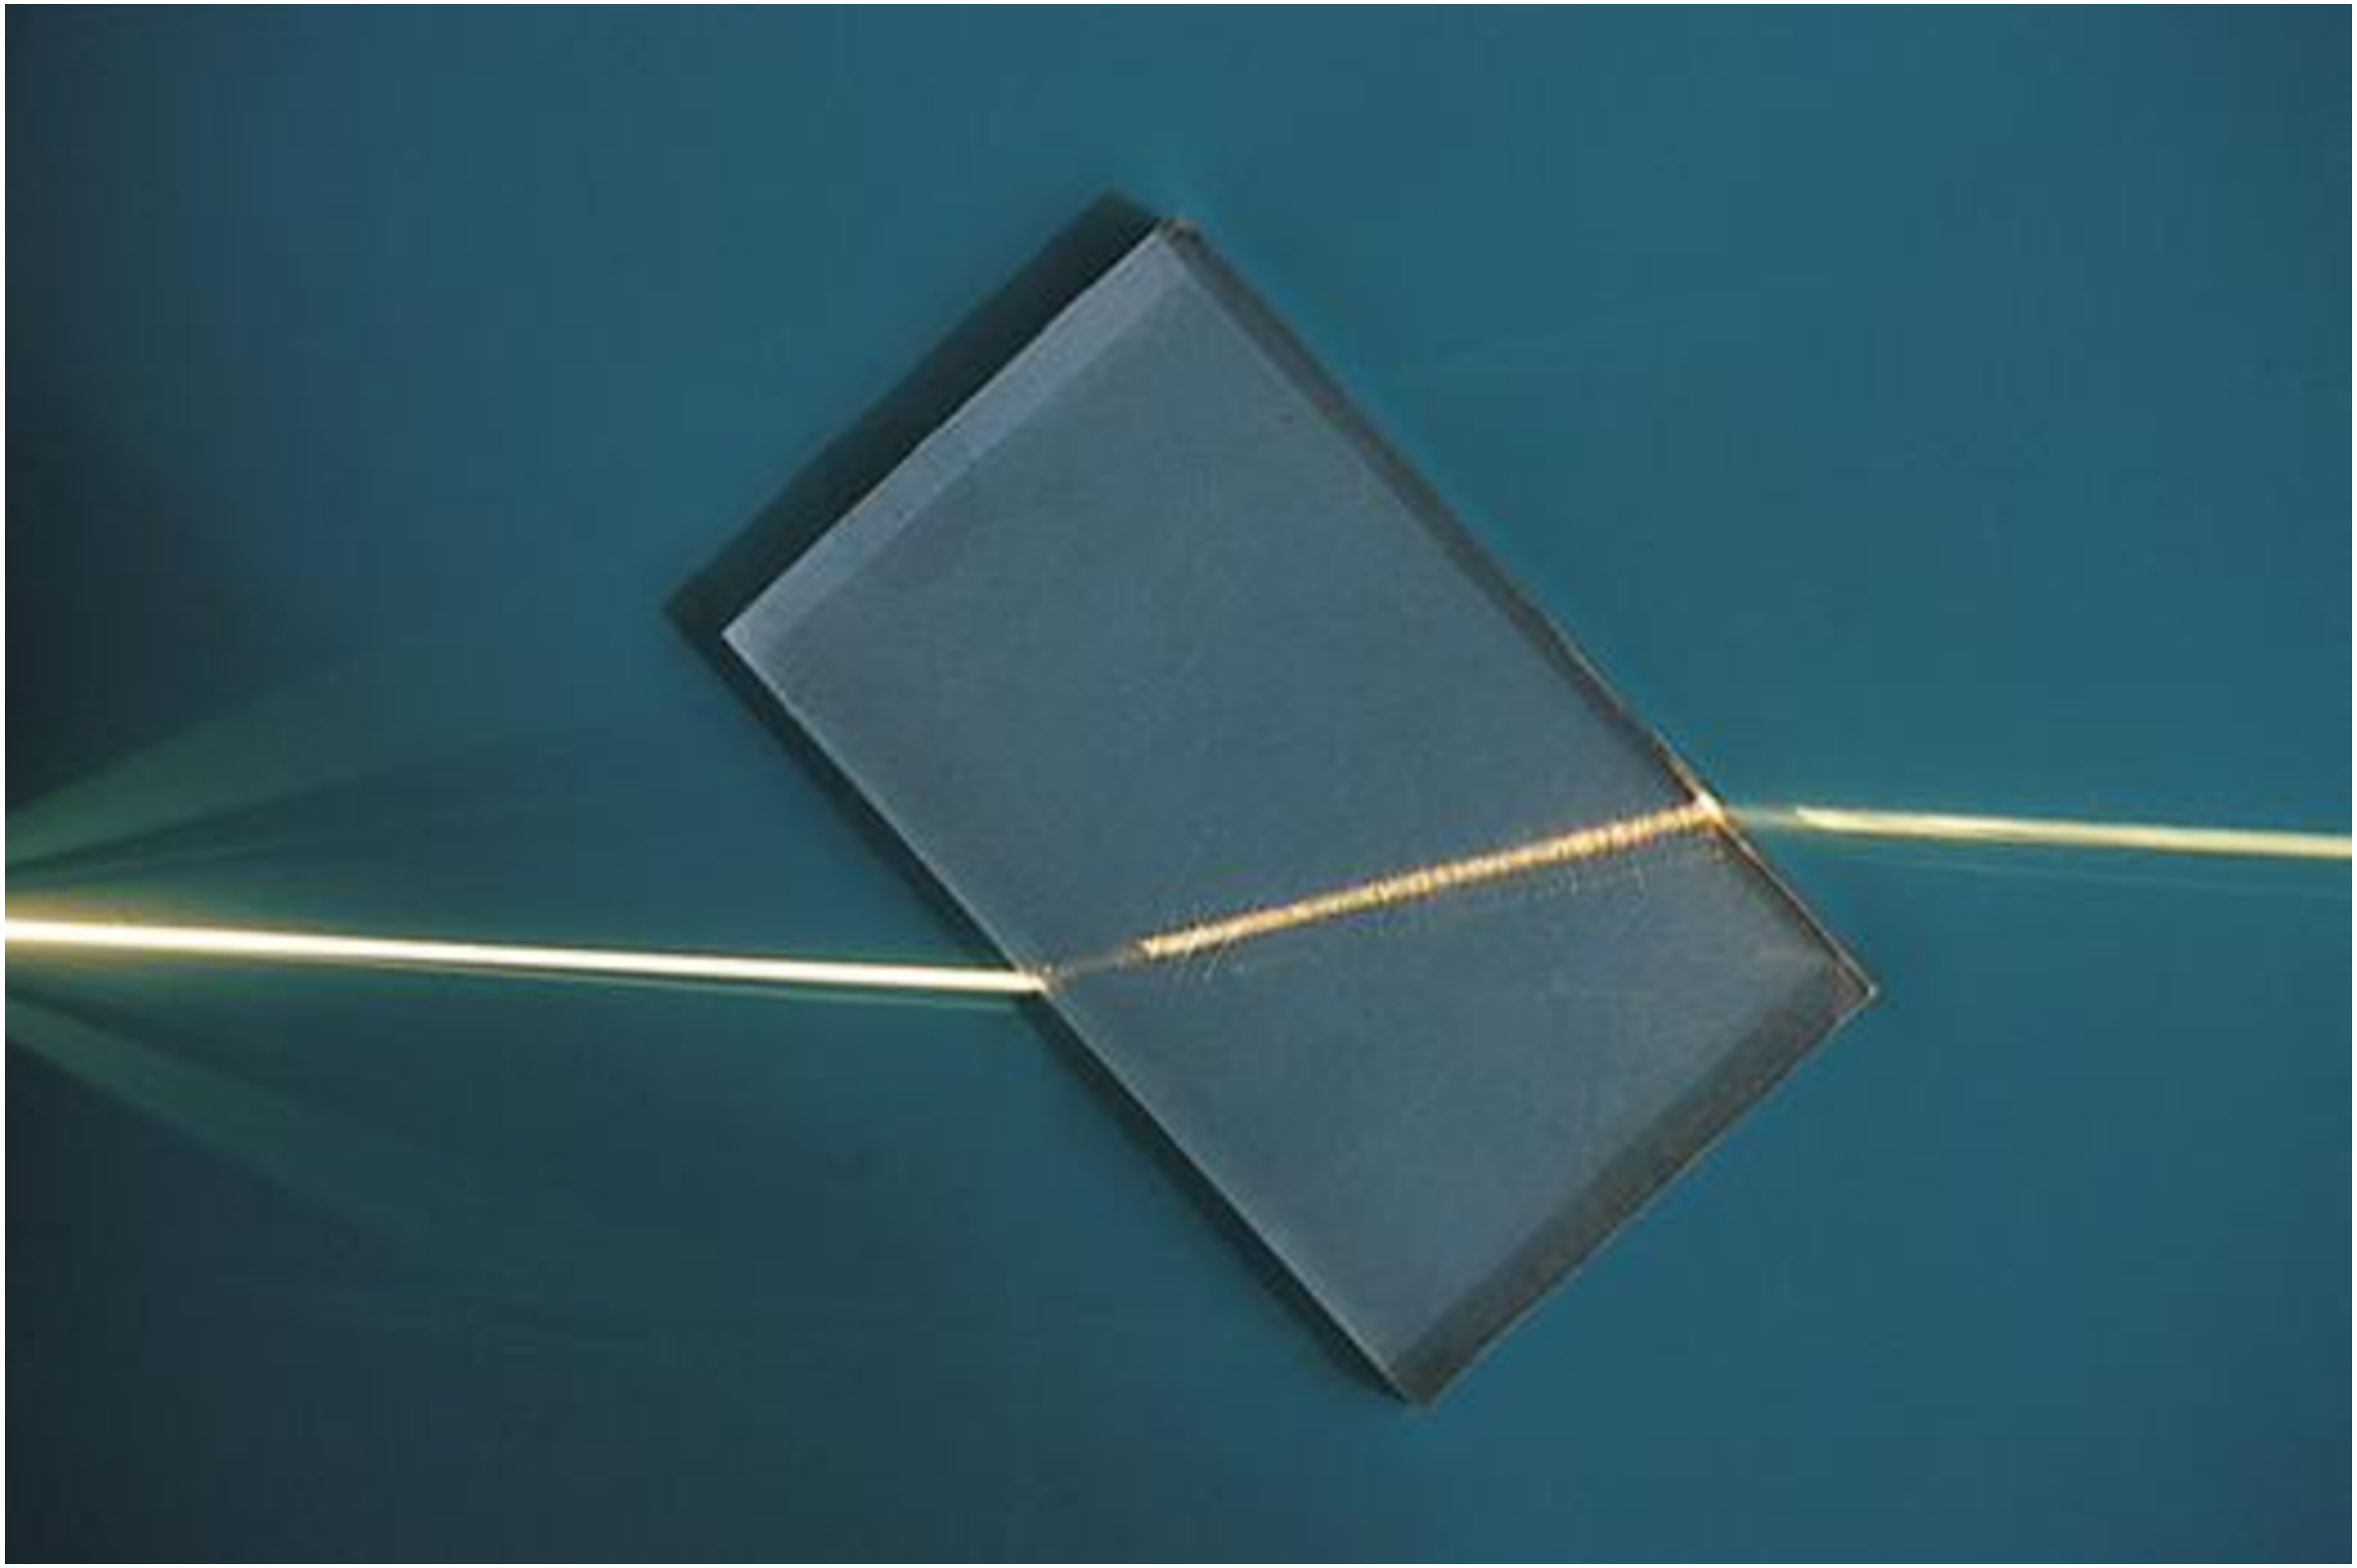

Snell's Law

Transmitted angle depends on index of refraction of medium incident ray is in and index of refraction of medium light is entering.

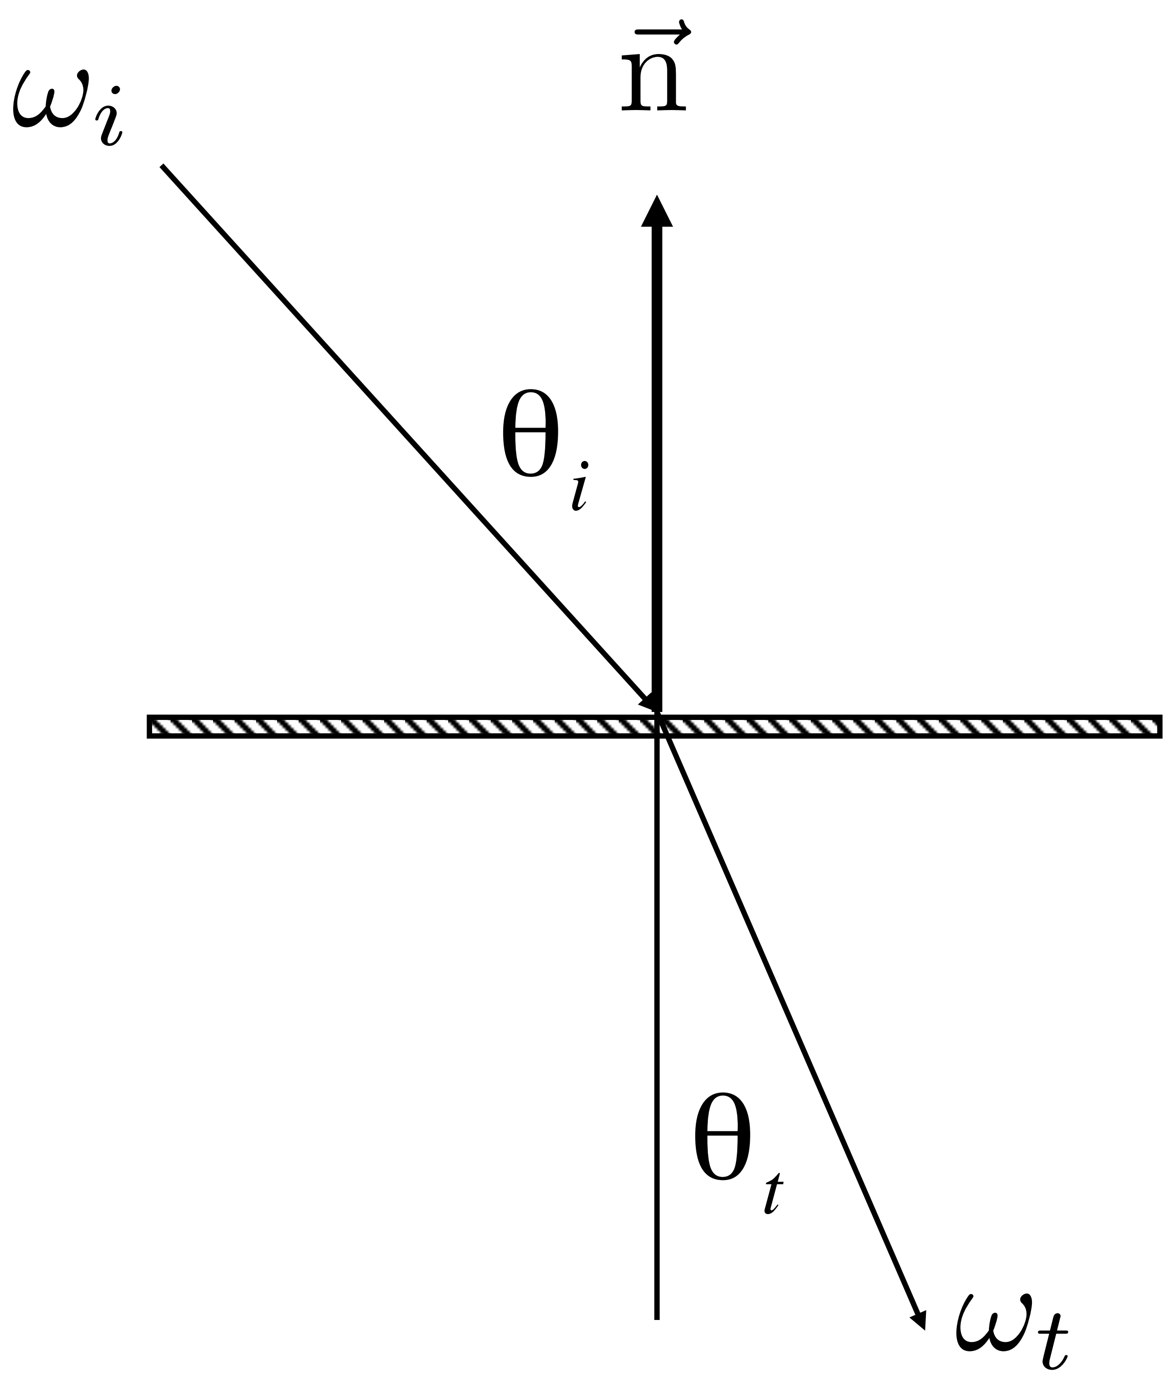

$$\eta_i \sin\theta_i = \eta_t \sin\theta_t$$

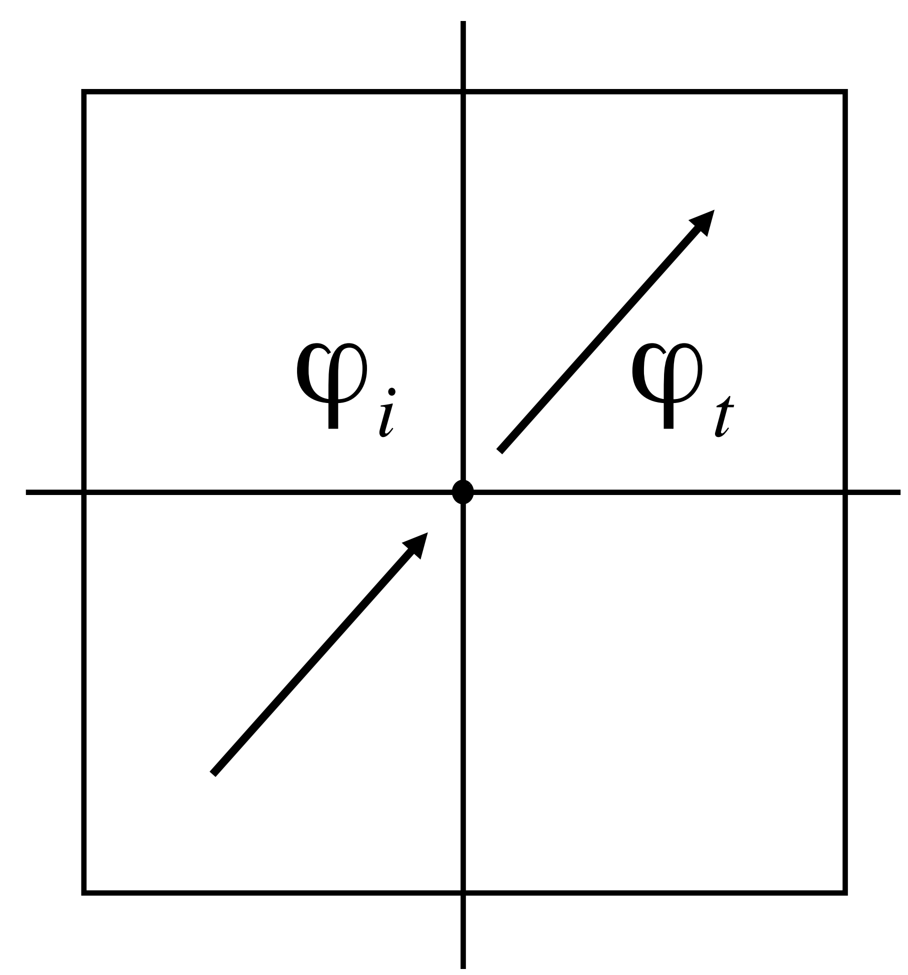

$$\varphi_t = \varphi_i \pm \pi$$

Medium	η*
Vacuum	1.0
Air (sea level)	1.0002
Water (20°C)	9
Glass	1.333
Diamond	1.5-1.6
	2.42

* index of refraction is wavelength dependent (these are averages)

Ray Tracing Through Lenses

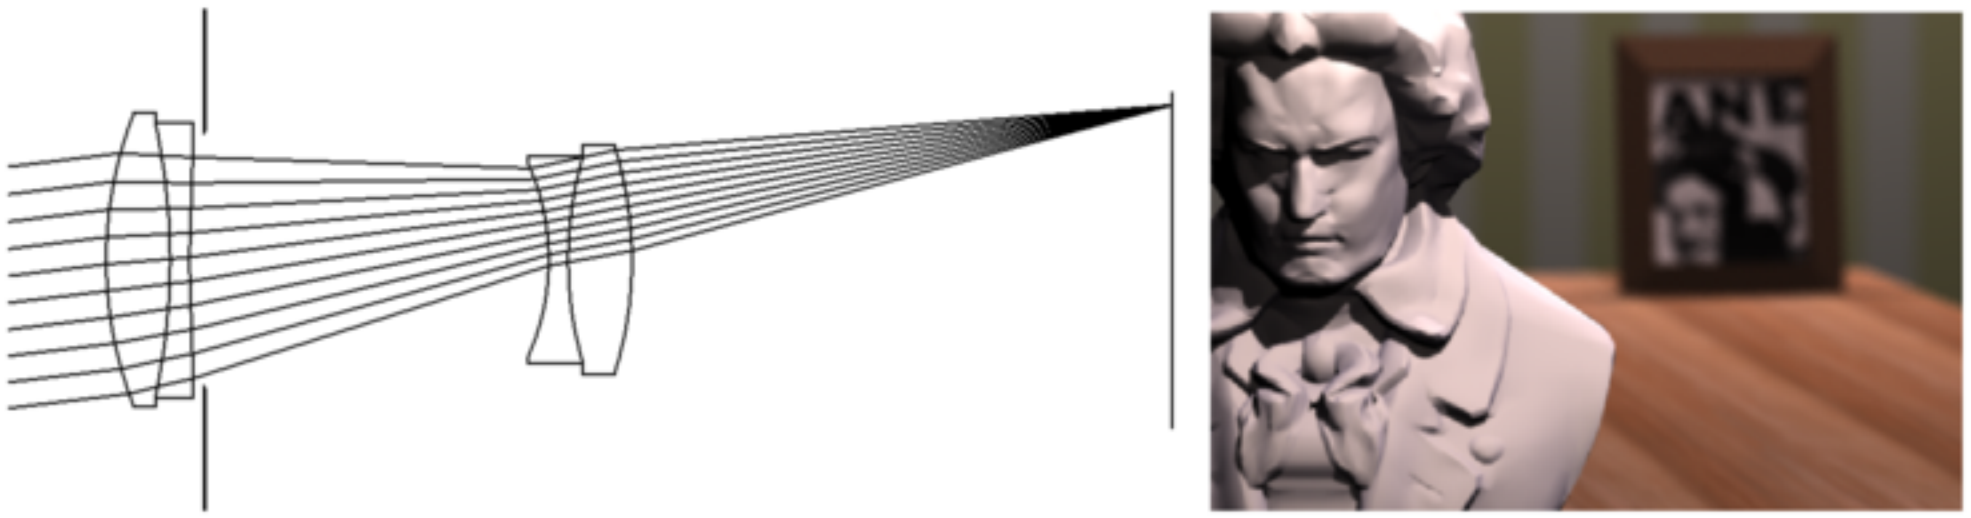

200 mm telephoto

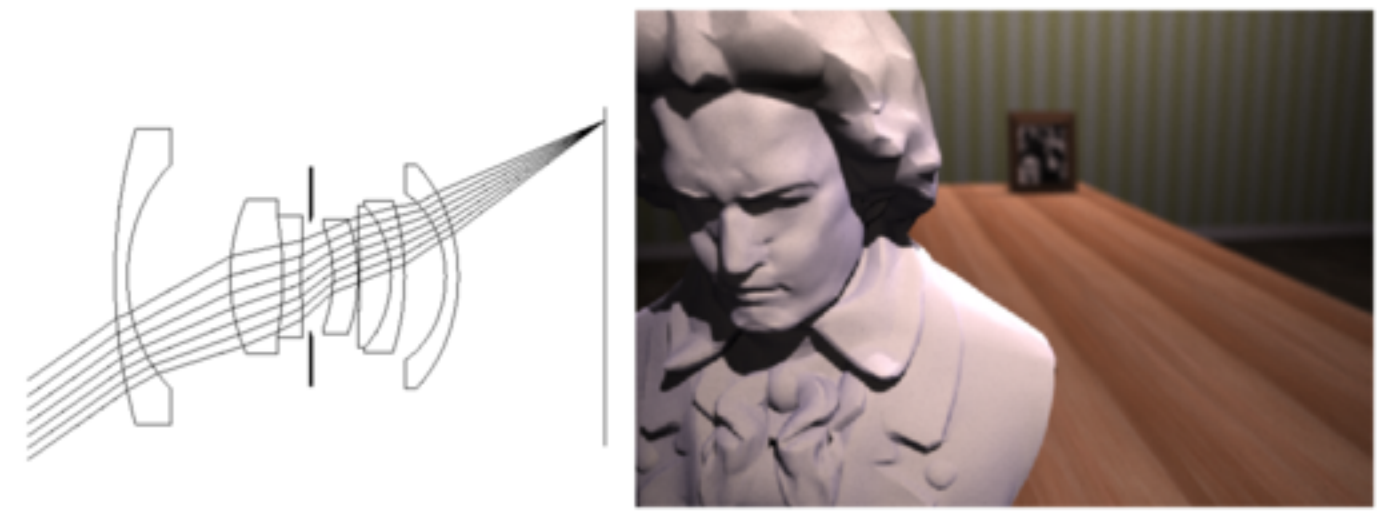

35 mm wide-angle

50 mm double-gauss

16 mm fisheye

From Kolb, Mitchell and Hanrahan (1995)

Ray Tracing Through Real Lens Designs

200 mm telephoto

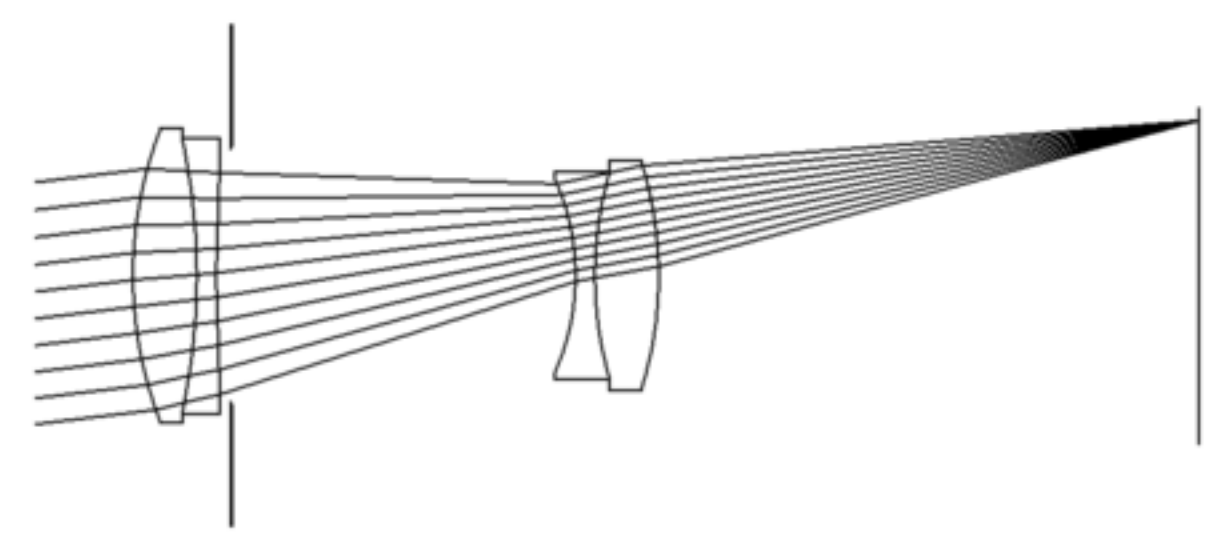

Notice shallow depth of field (out of focus background)

Ray Tracing Through Real Lens Designs

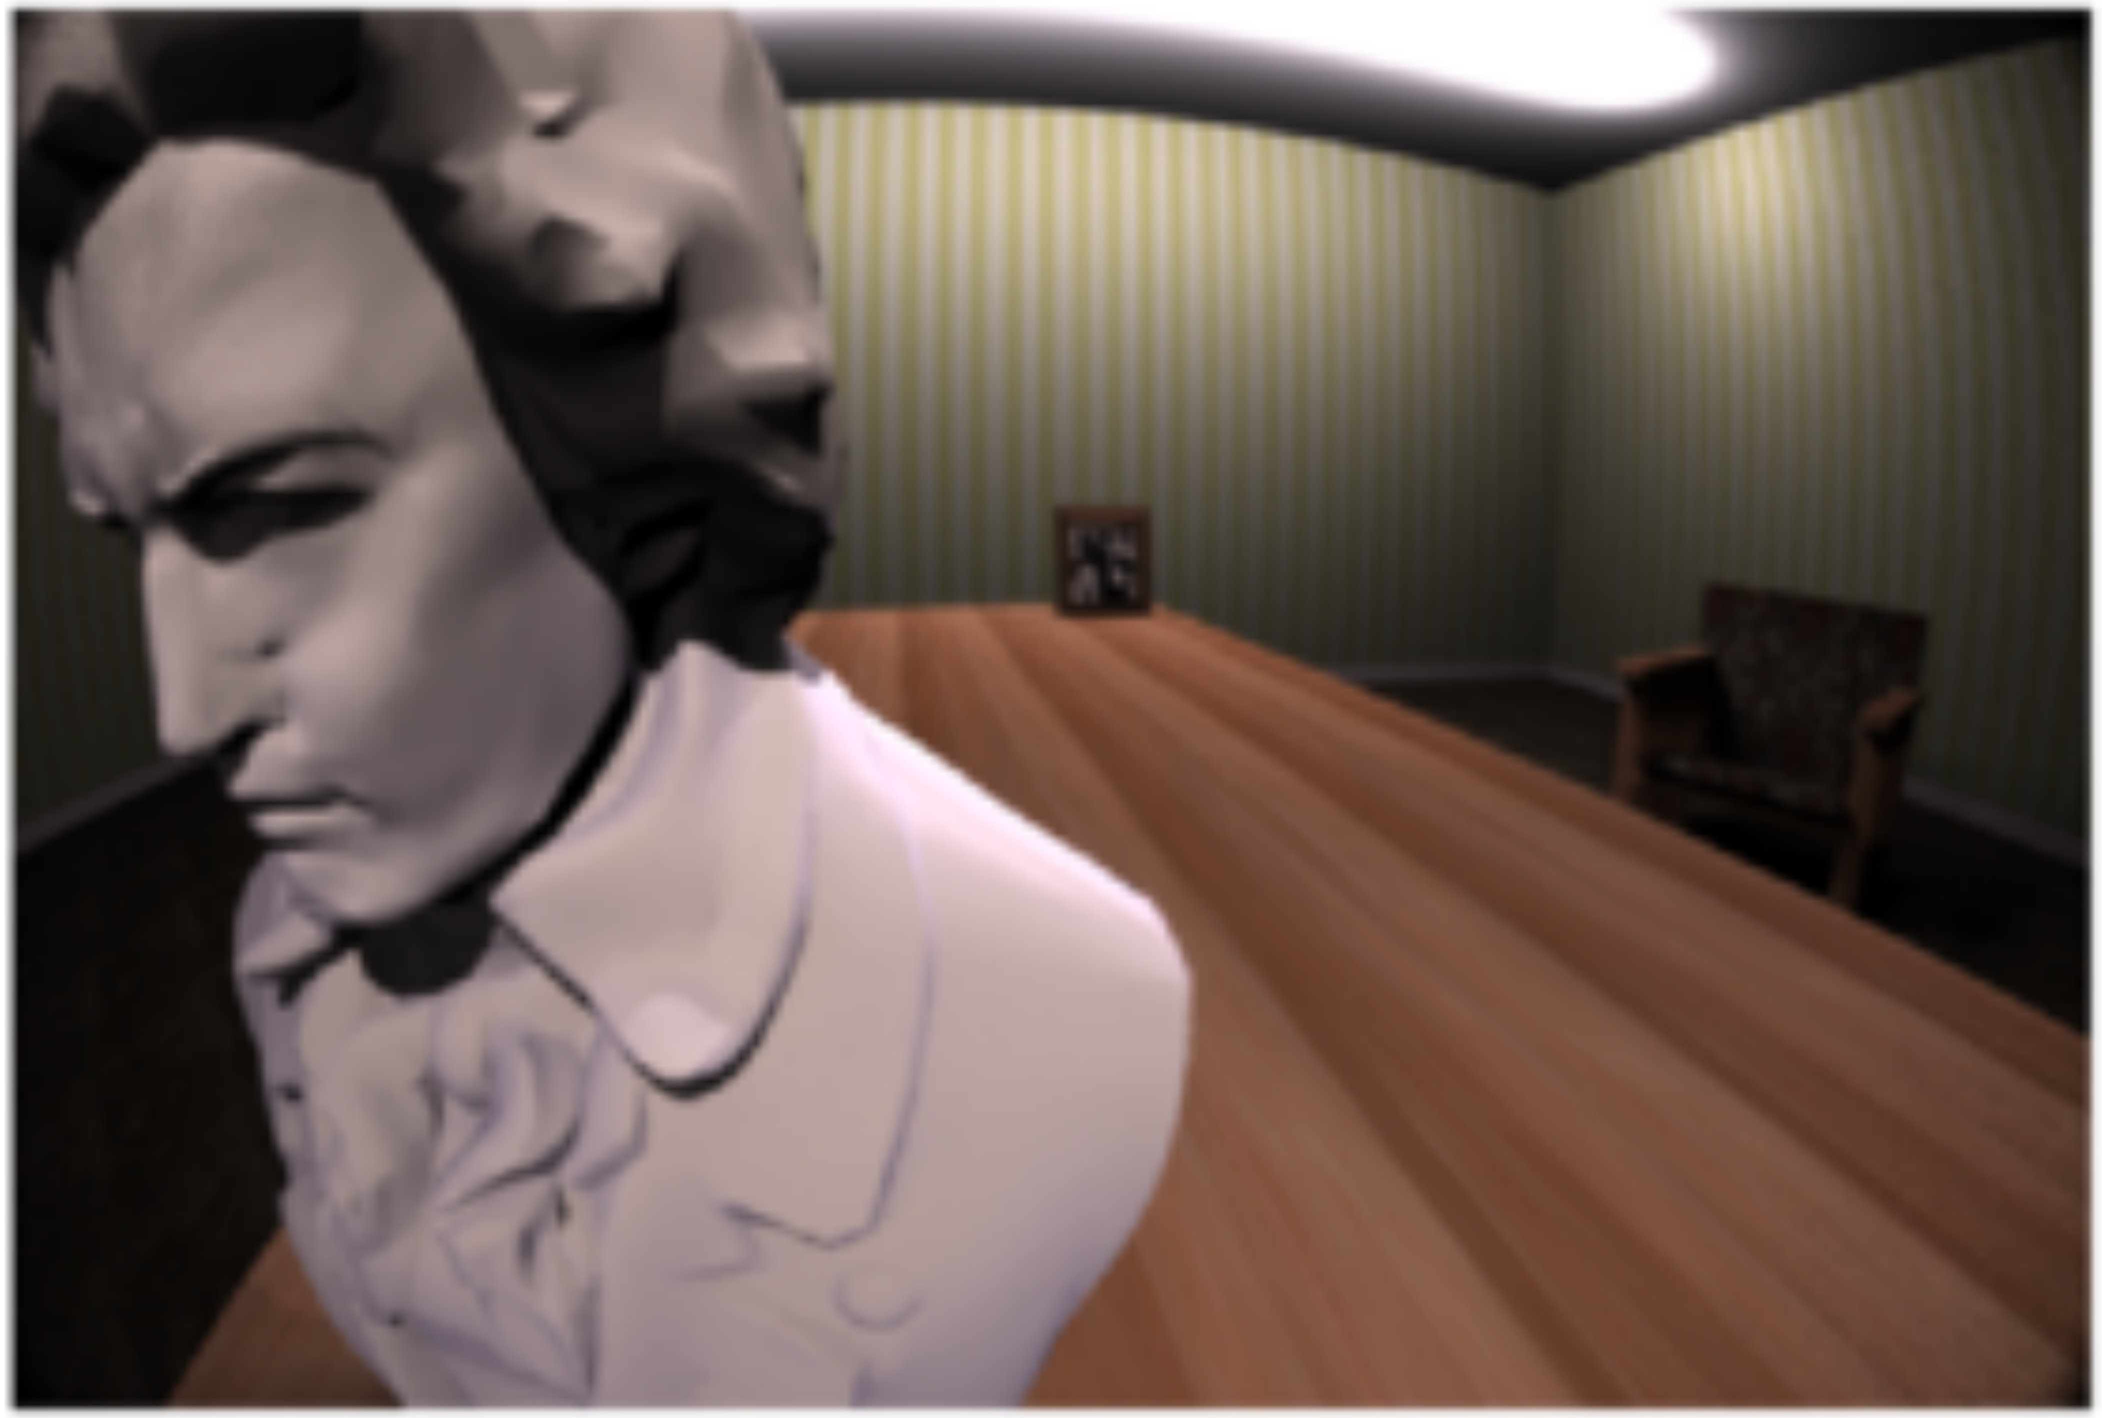

16 mm fisheye

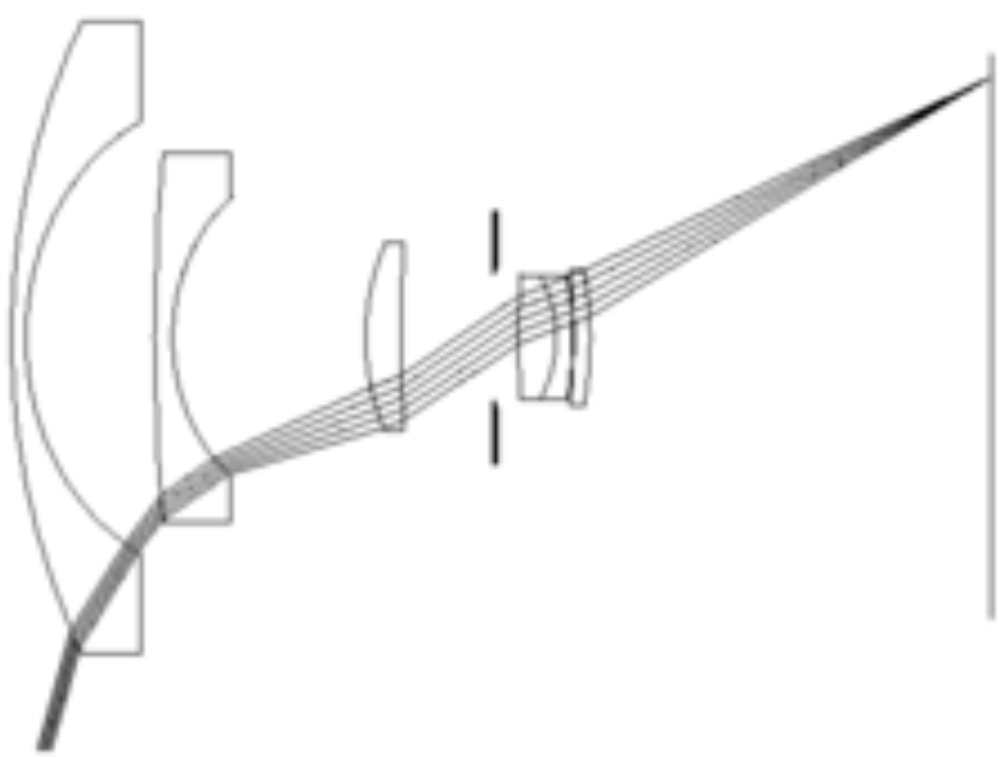

Notice distortion in the corners (straight lines become curved)

Ray Tracing Through Real Lens Designs

Monte Carlo approach

- **At every sensor pixel, compute integral of rays incident on pixel area arriving from all paths through the lens**

Algorithm (for a pixel)

- **Choose N random positions in pixel**
- **For each position x′, choose a random position on the back element of the lens x′′**
- **Trace a ray through from x′ to x′′, trace refractions through lens elements until it misses the next element (kill ray) or exits the lens (path trace through the scene)**
- **Weight each ray according to radiometric calculation on next slide to estimate irradiance E(x′)**

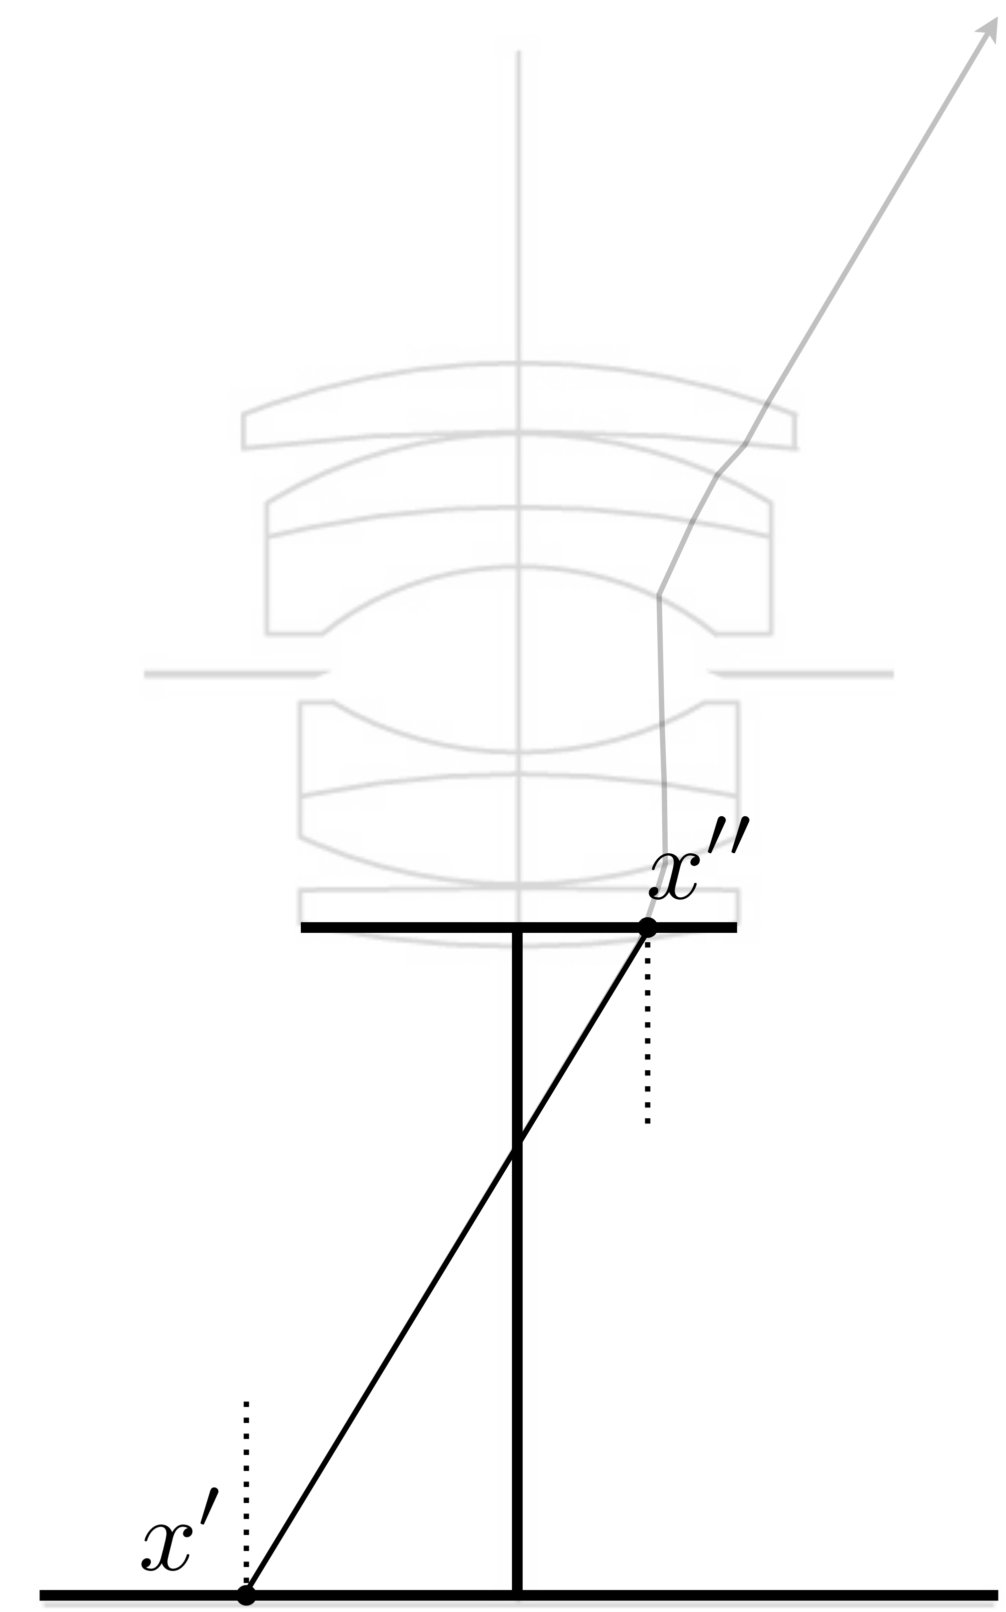

Radiometry for Tracing Lens Designs

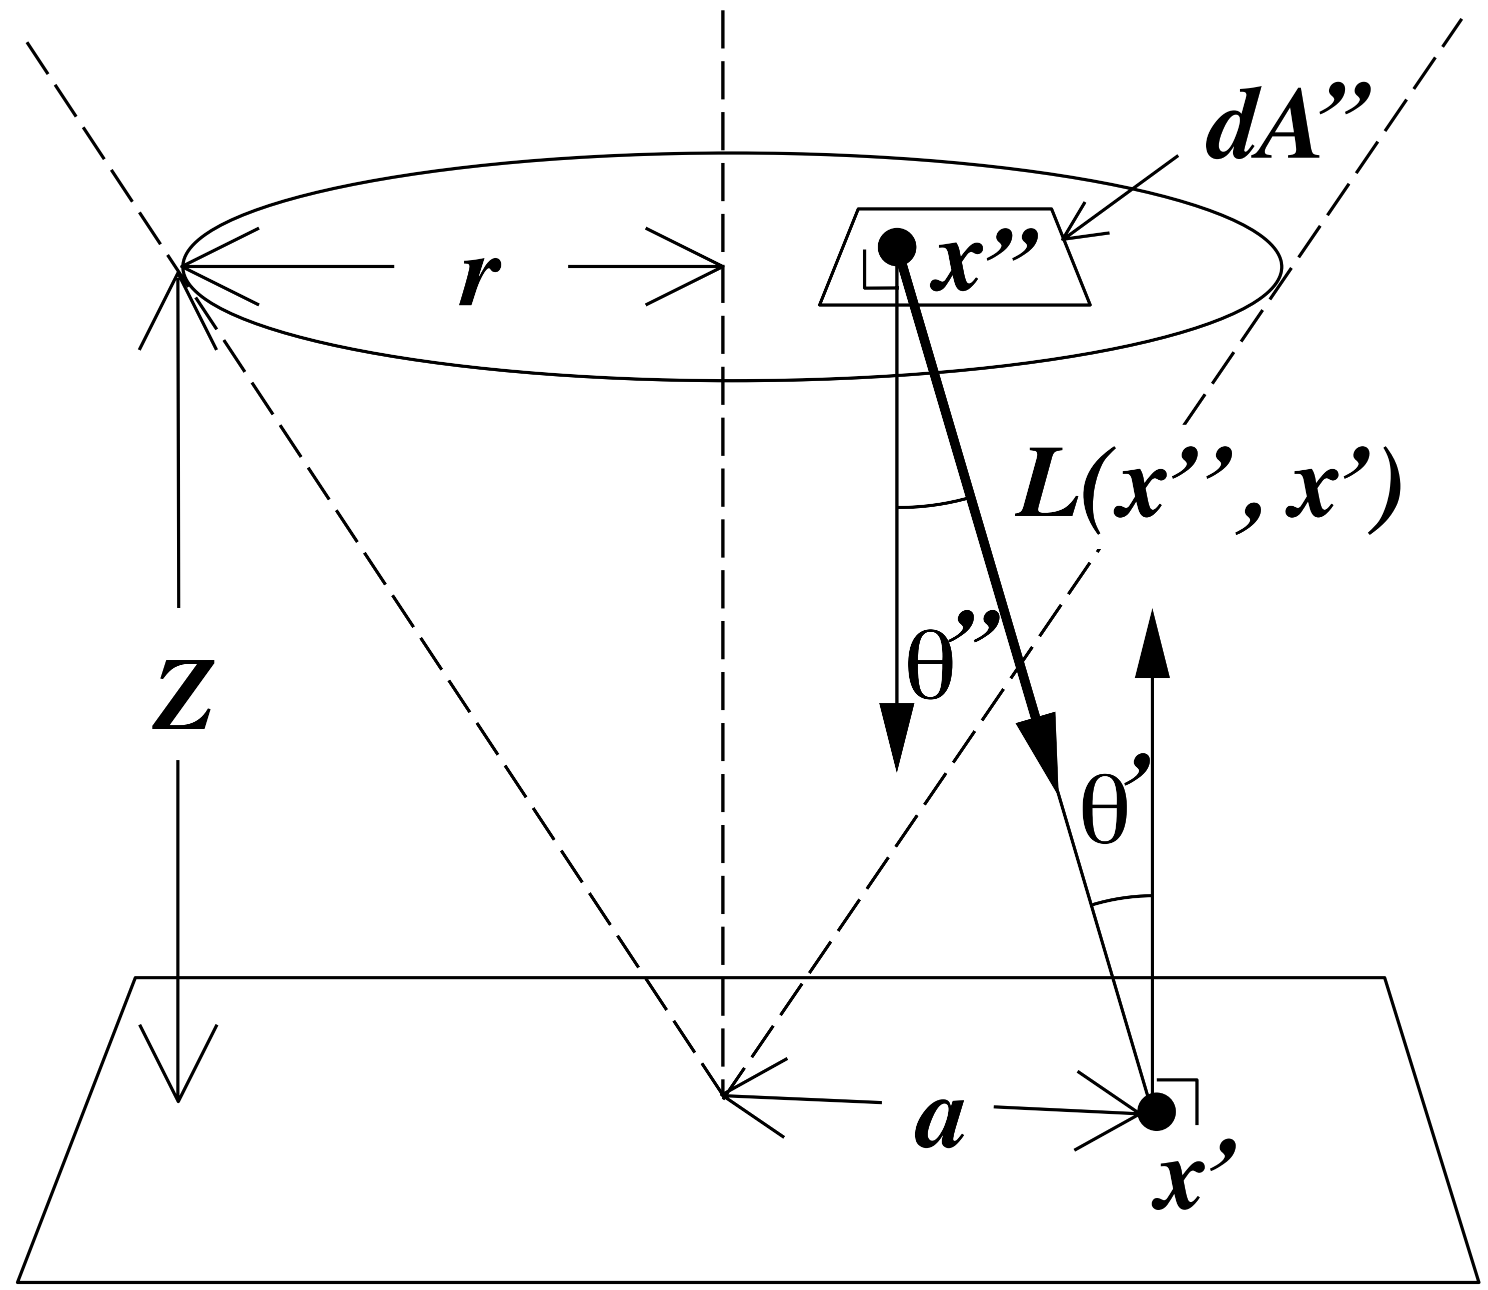

$$E(x') = \int_{x'' \in D} L(x'' \to x') \frac{\cos\theta' \cos\theta''}{||x'' - x'||^2} dA''$$

$$= \frac{1}{Z^2} \int_{x'' \in D} L(x'' \to x') \cos^4 \theta dA''$$

D = region of lens aperture

Real sensor: pixel measurement integrates irradiance over area of pixel

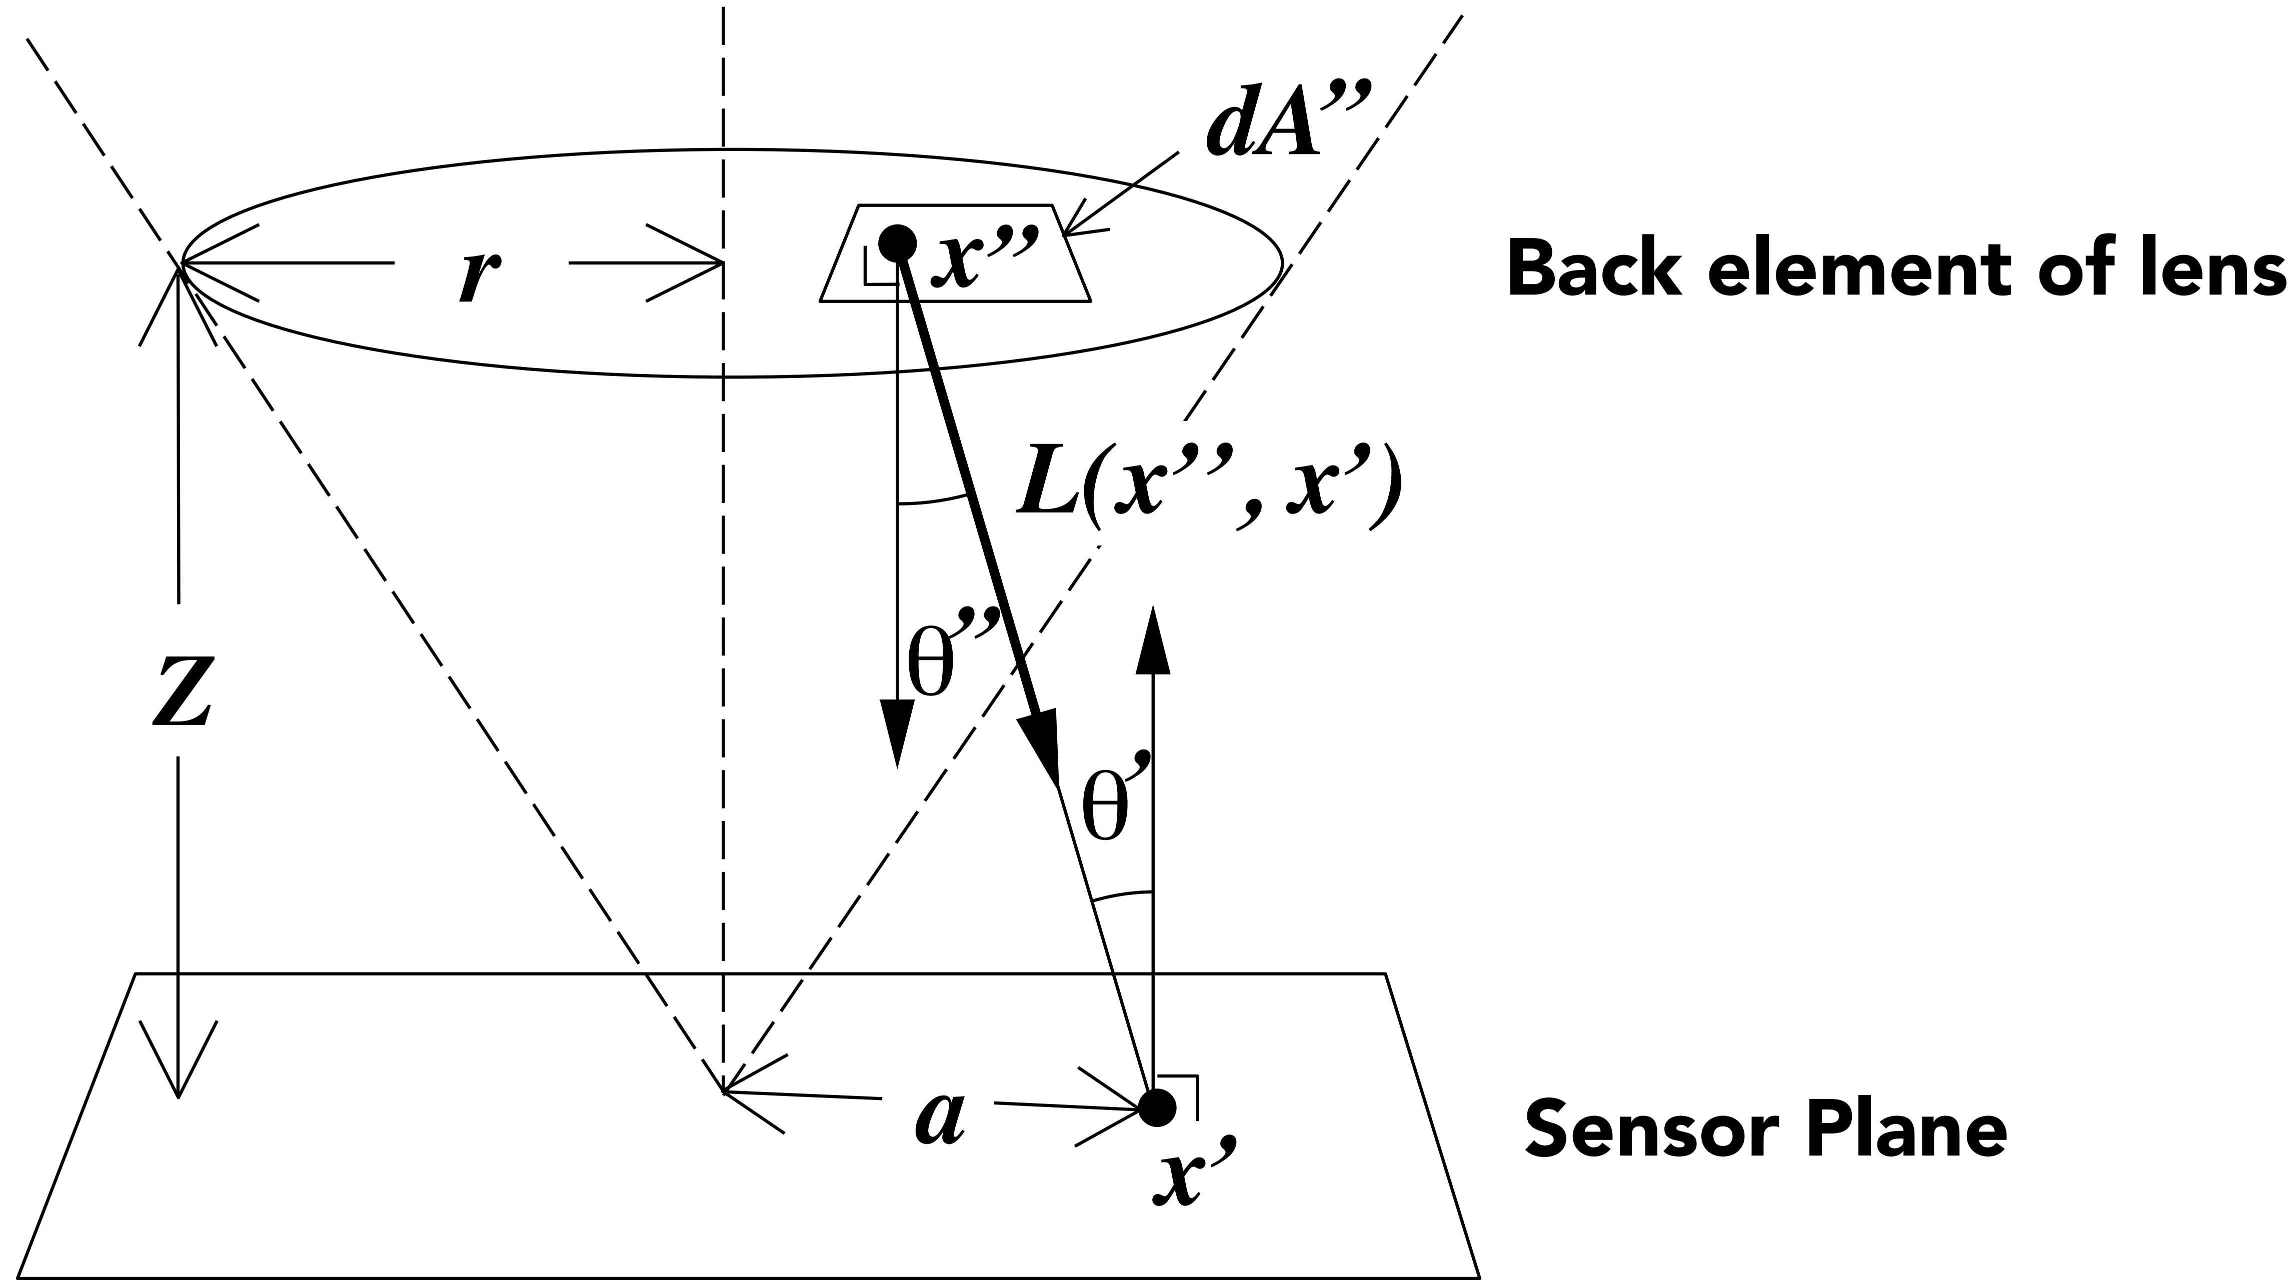

$$E(\text{pixel}) = \int_{x' \in P} \int_{x'' \in D} L(x'' \to x') \frac{\cos\theta' \cos\theta''}{||x'' - x'||^2} dA'' dA'$$

P = region of pixel

D = region of lens aperture

Exposure

High-Speed Photography (short exposure)

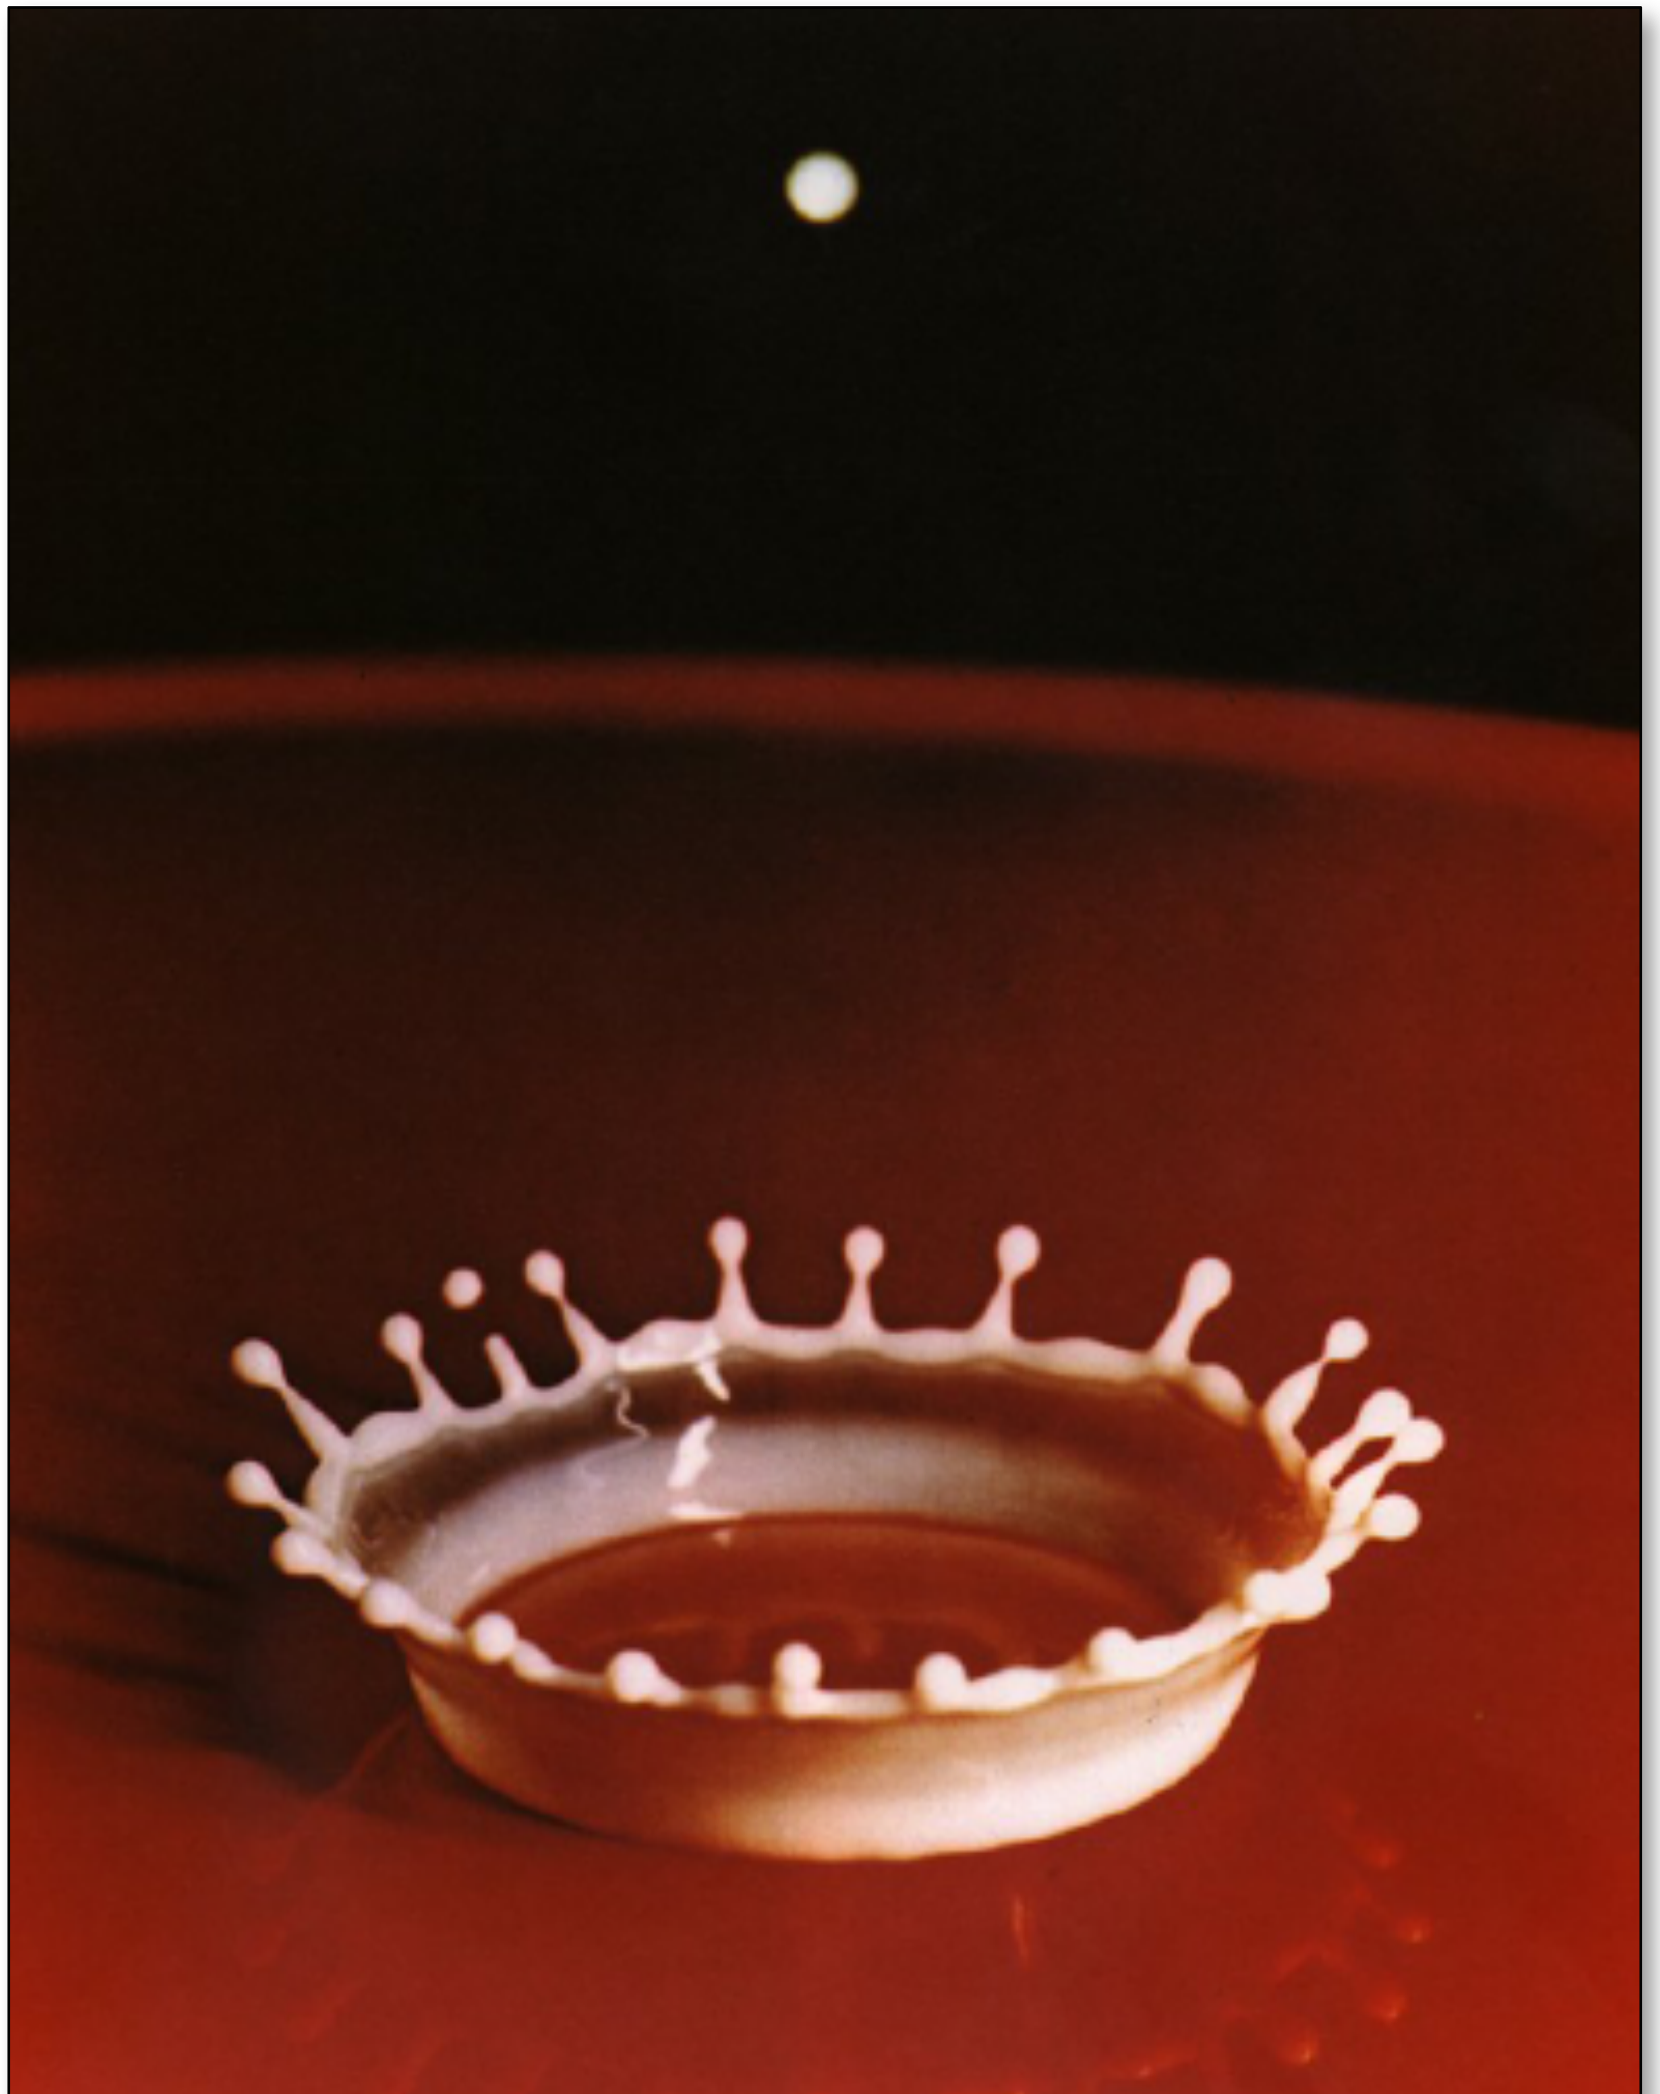

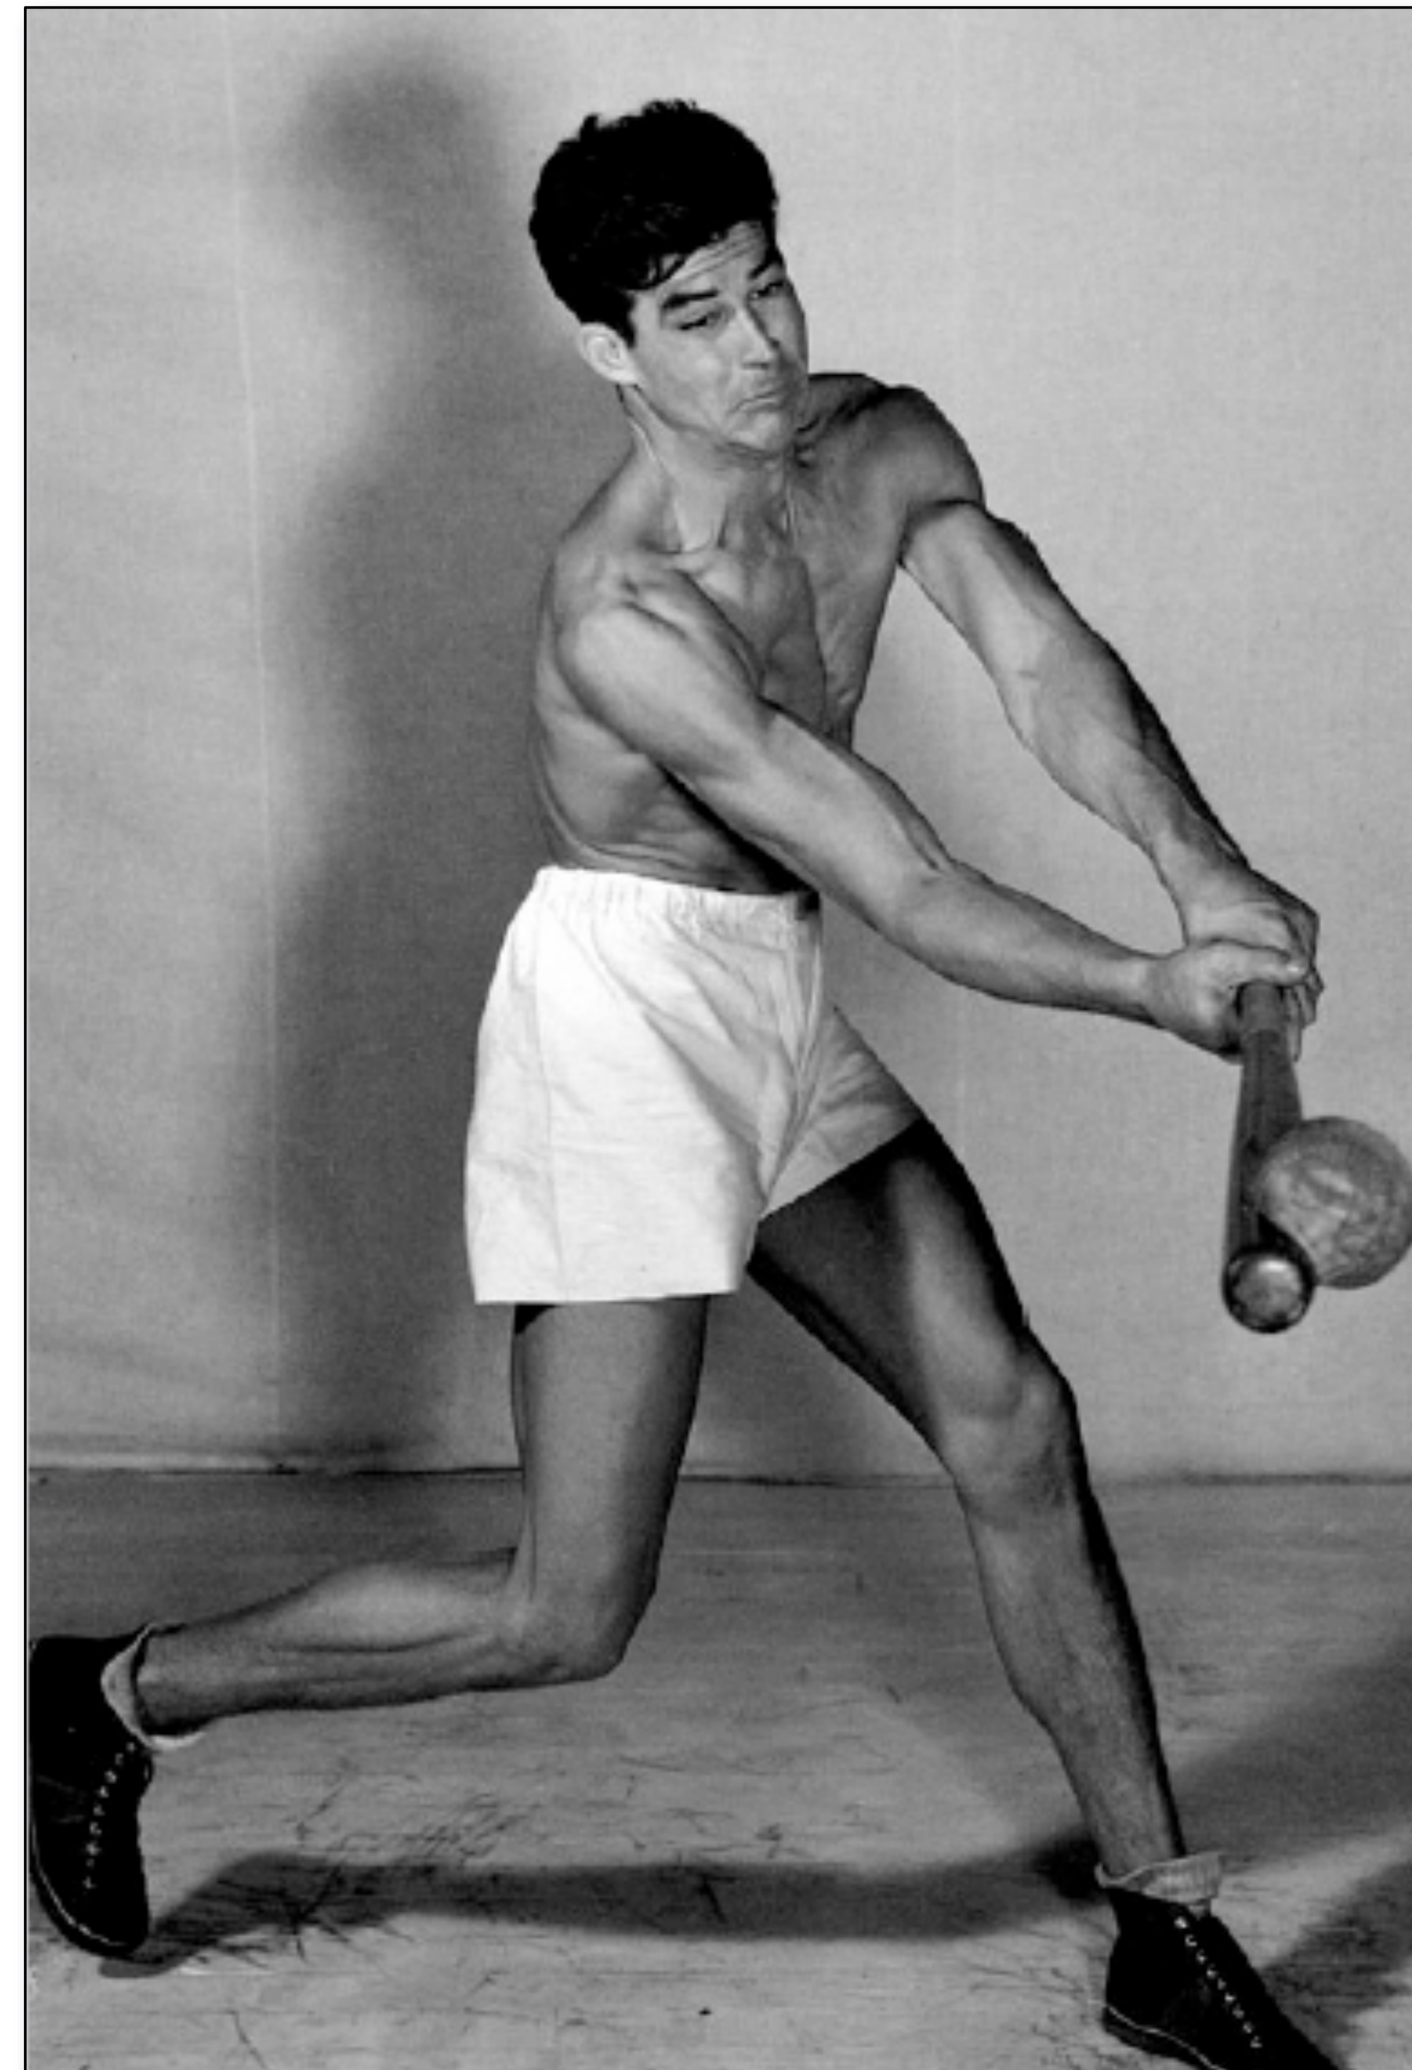

Harold Edgerton

Long exposure
(a bad photo by Kayvon)

Long exposure
(intentional for artistic effect)

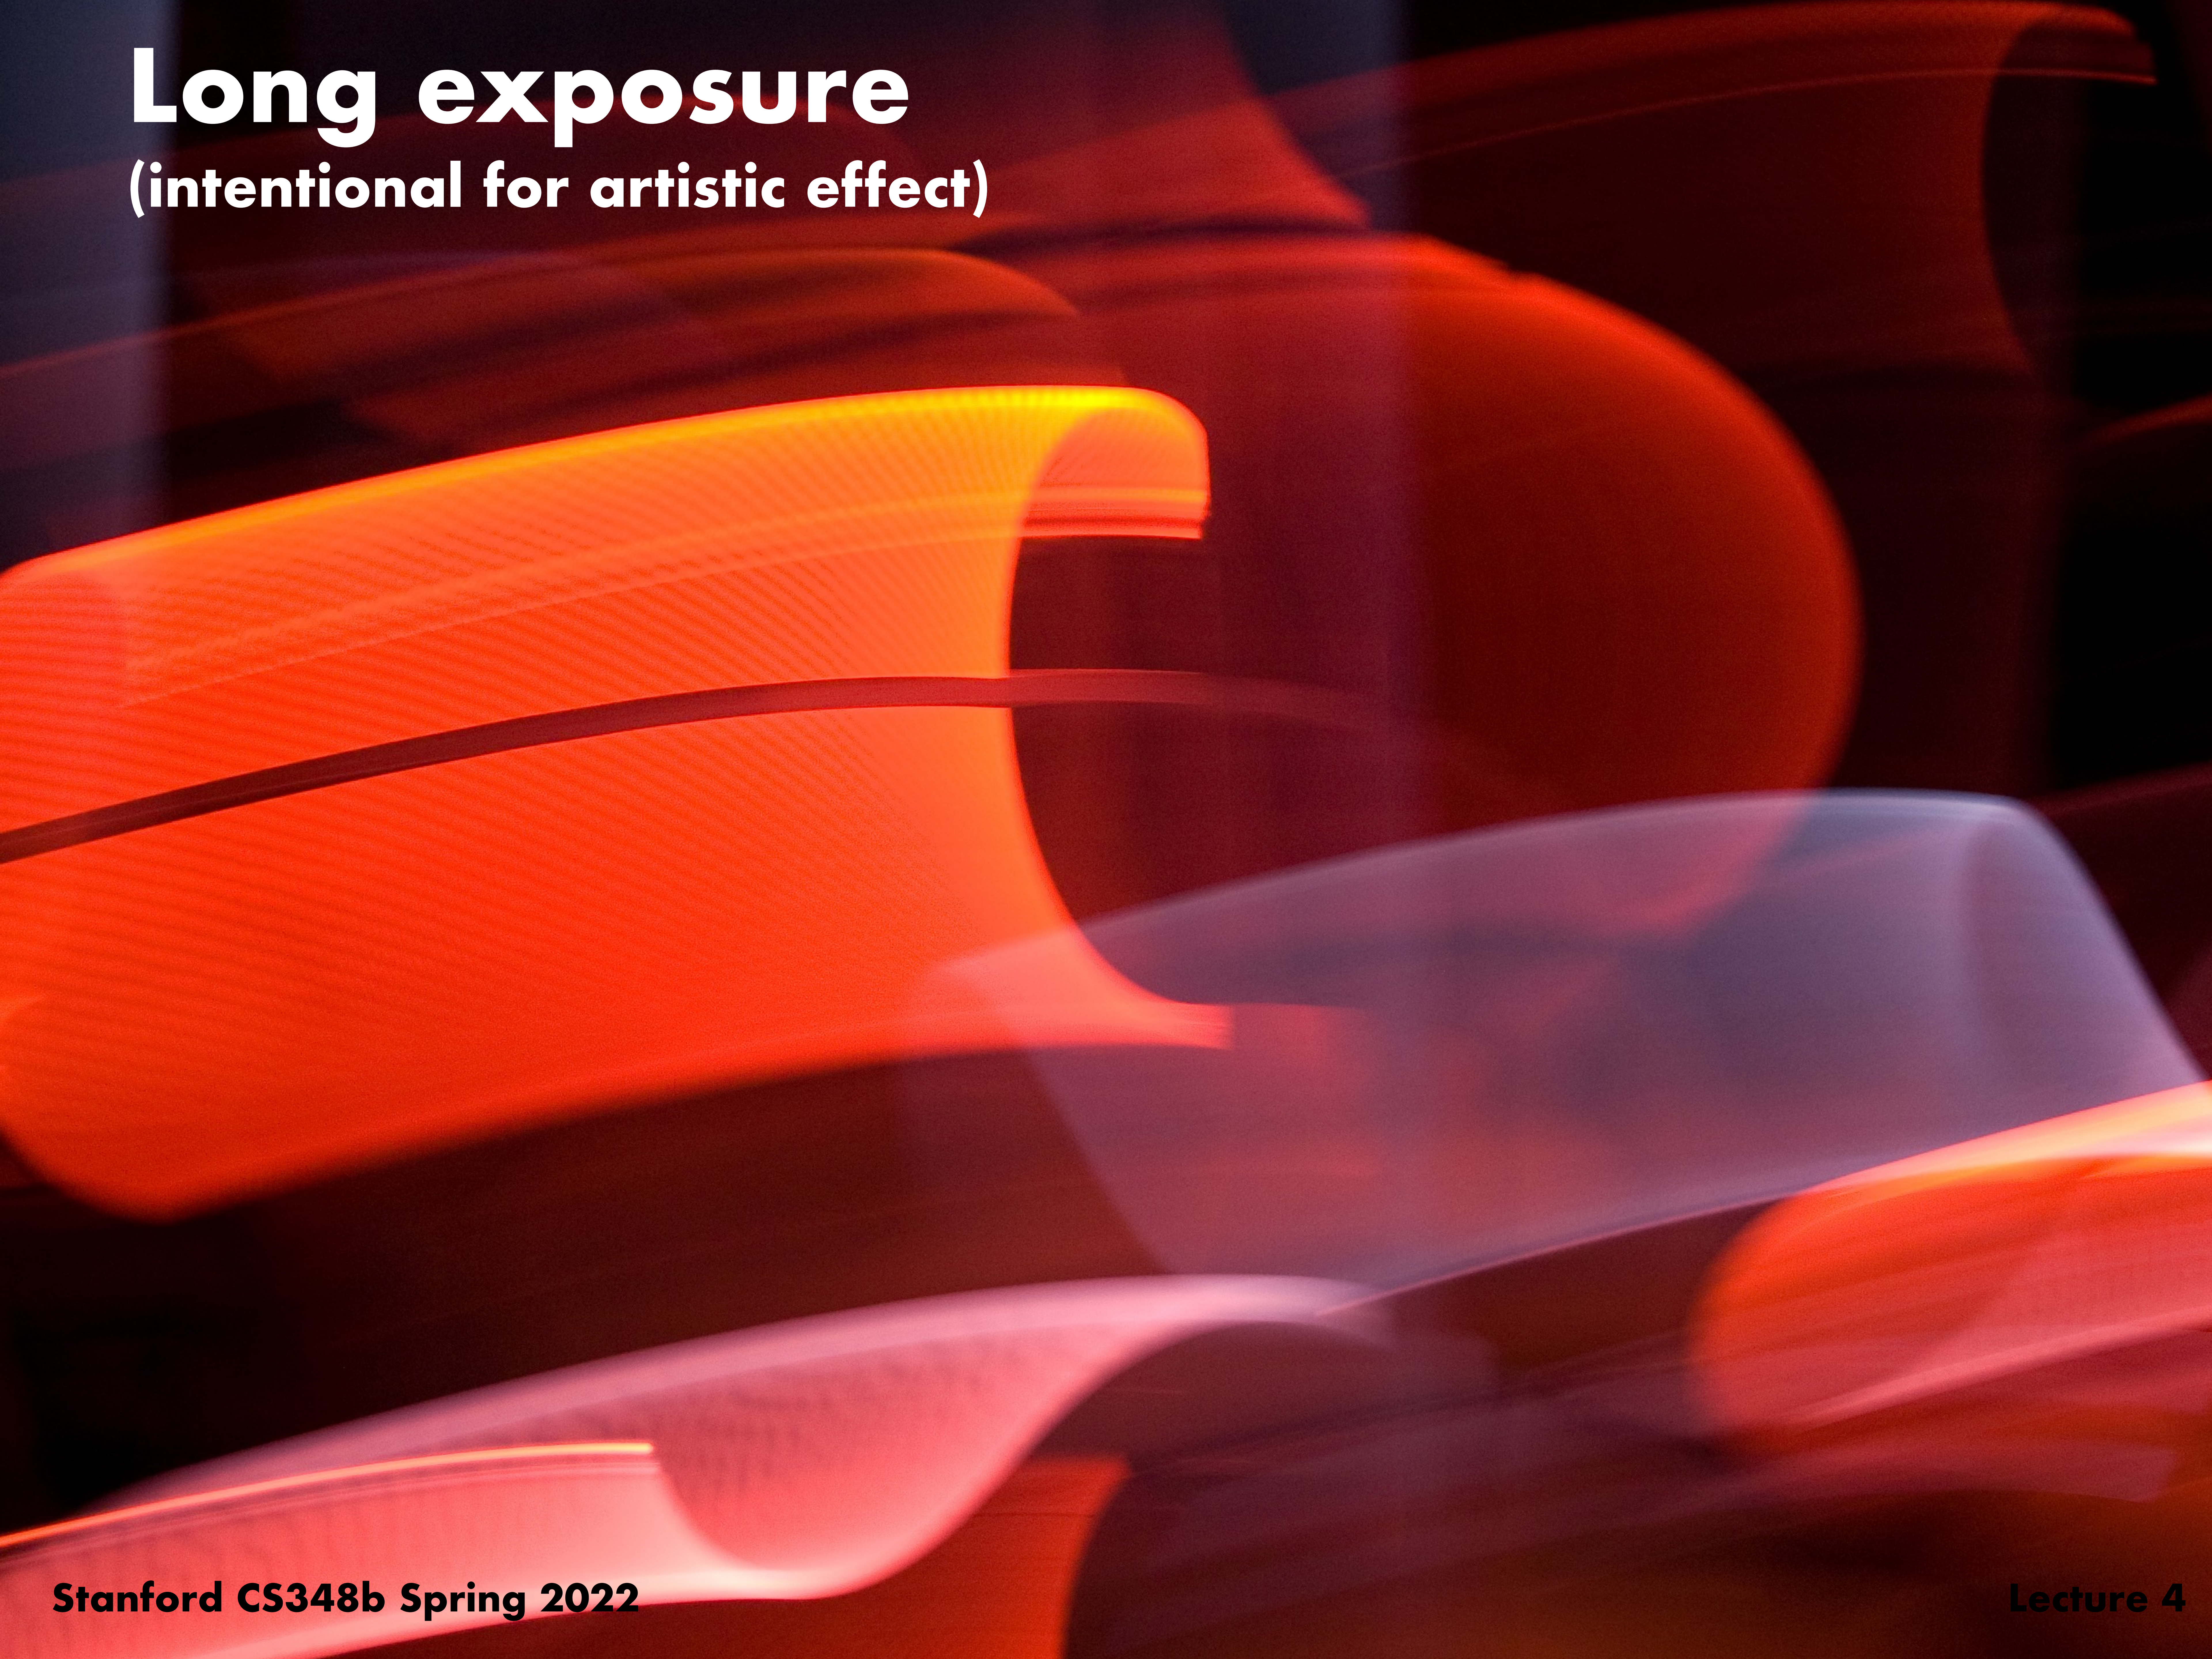

Physical Shutter (1/25 Sec Exposure)

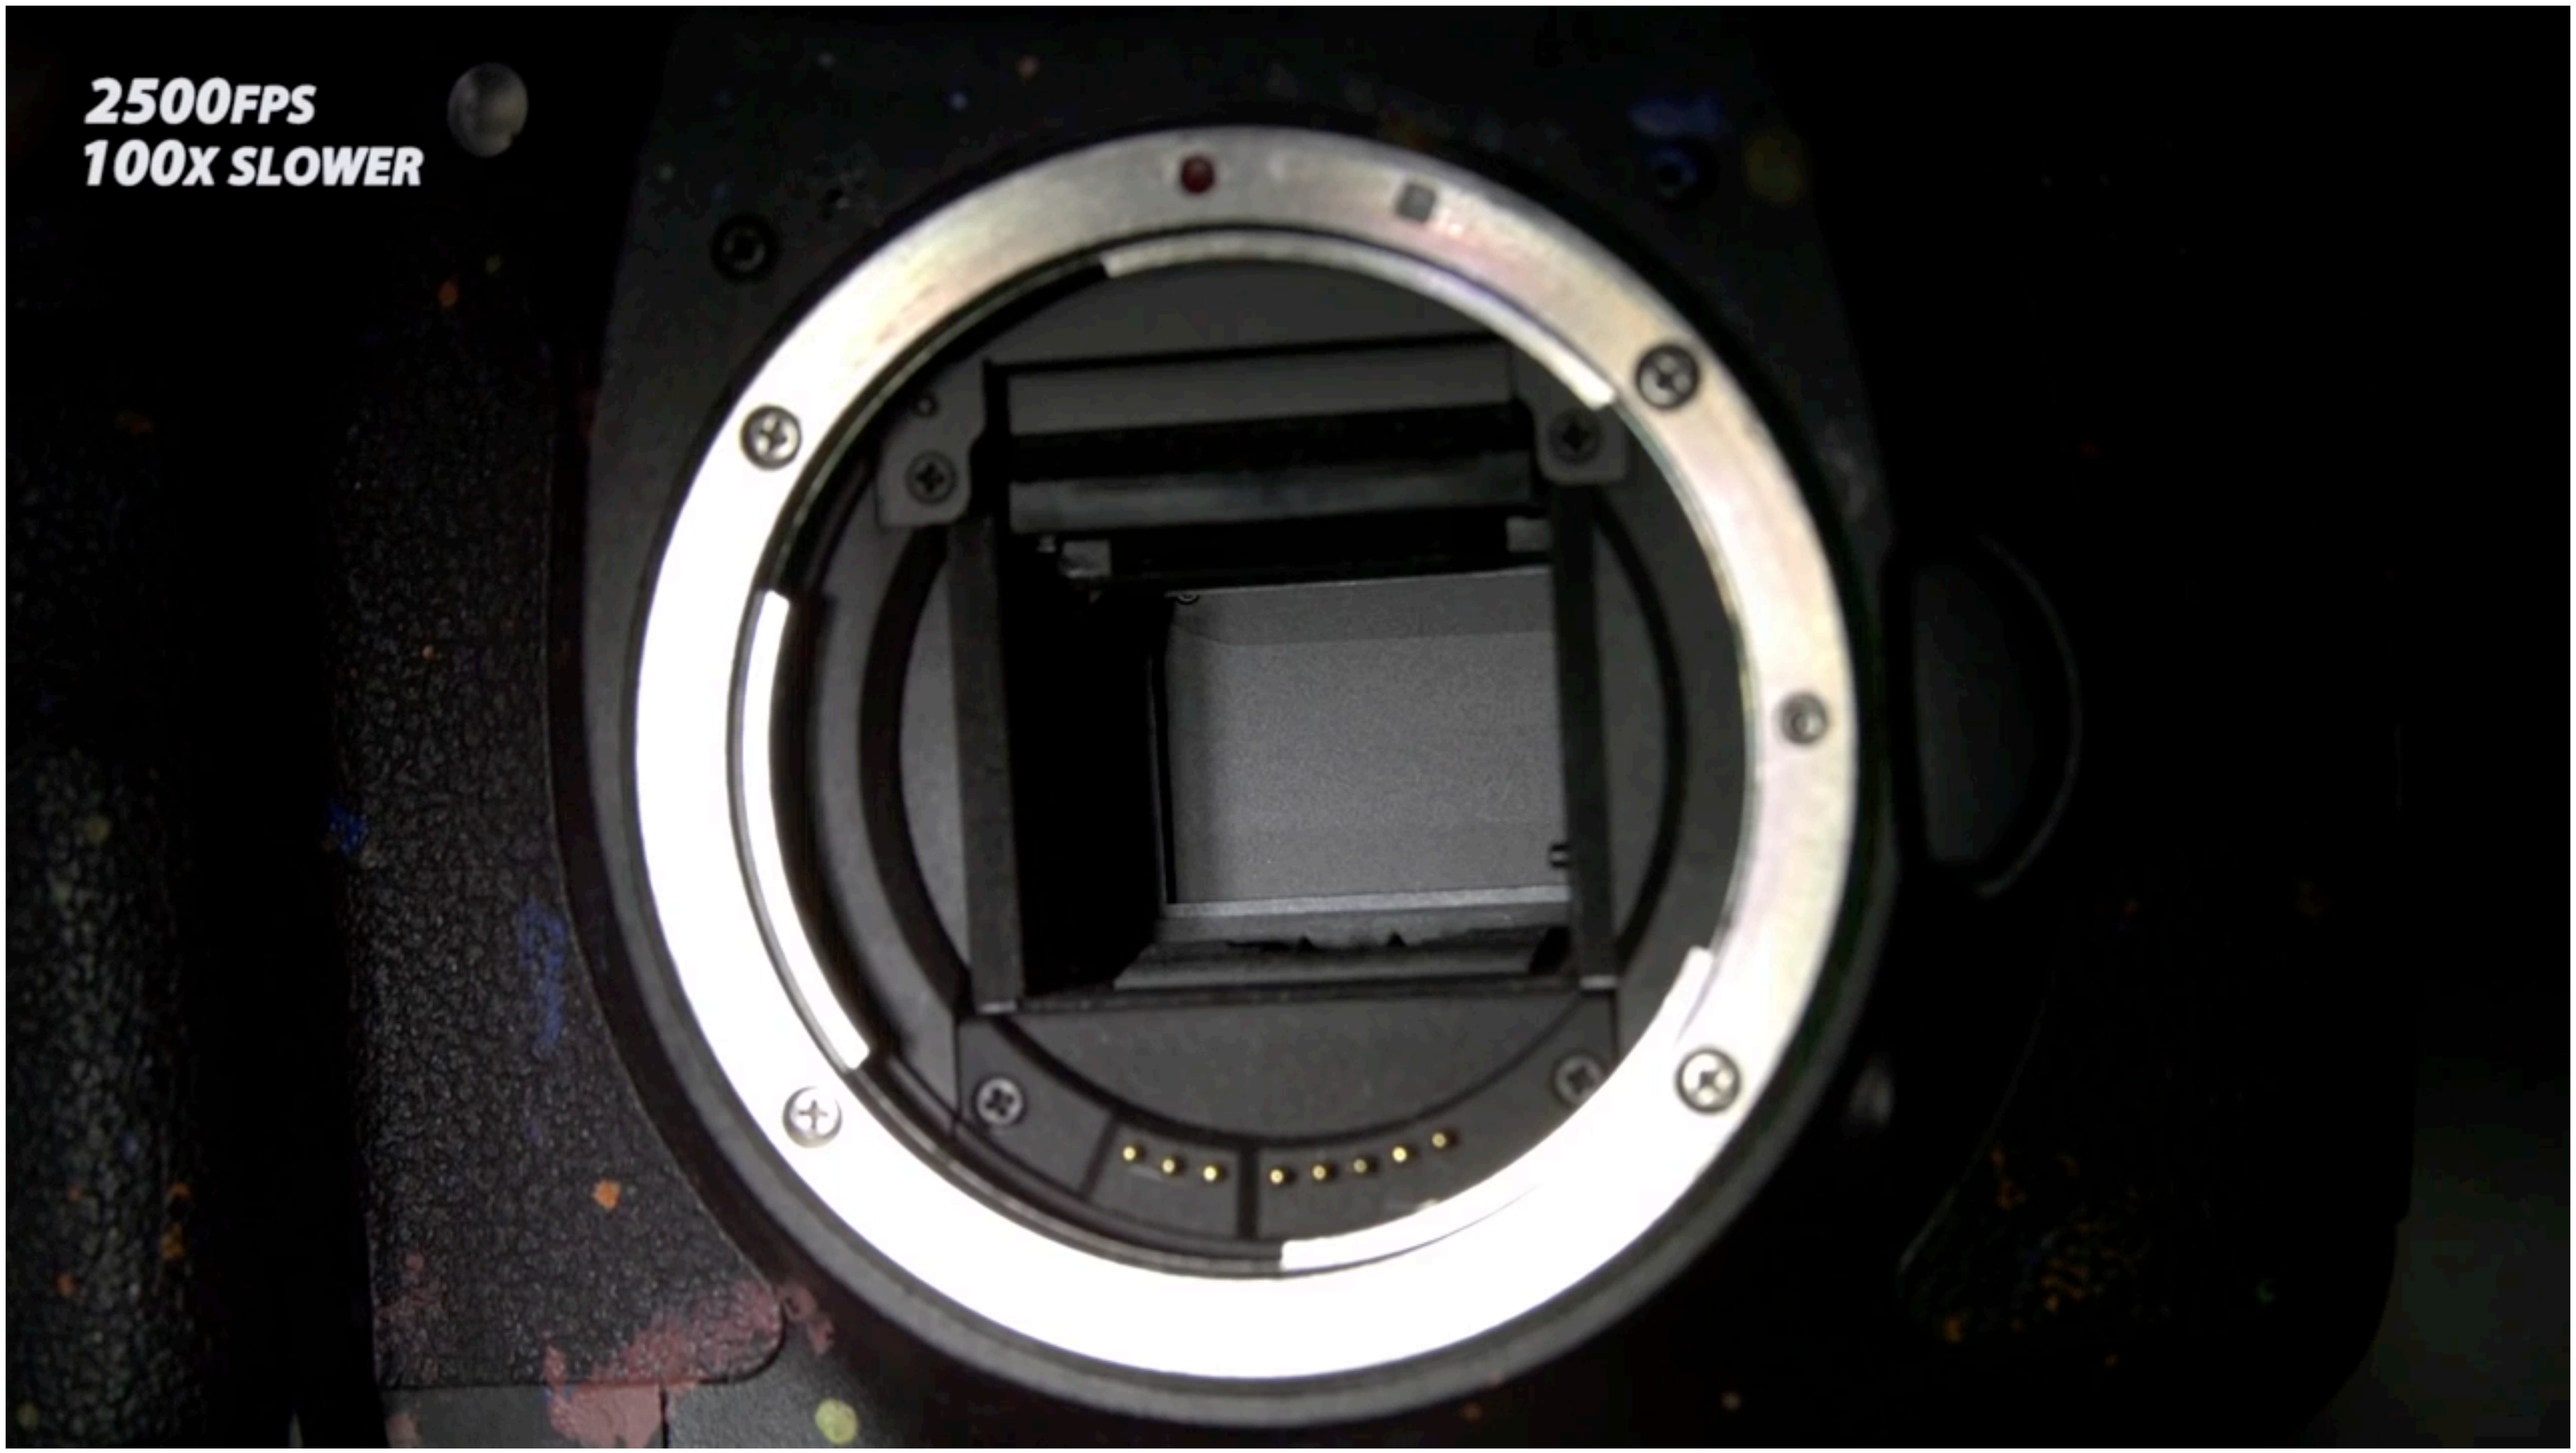

The Slow Mo Guys, https://youtu.be/CmjeCchGRQo

Main Side Effect of Shutter Speed

Motion blur: handshake, subject movement

Doubling shutter time doubles motion blur

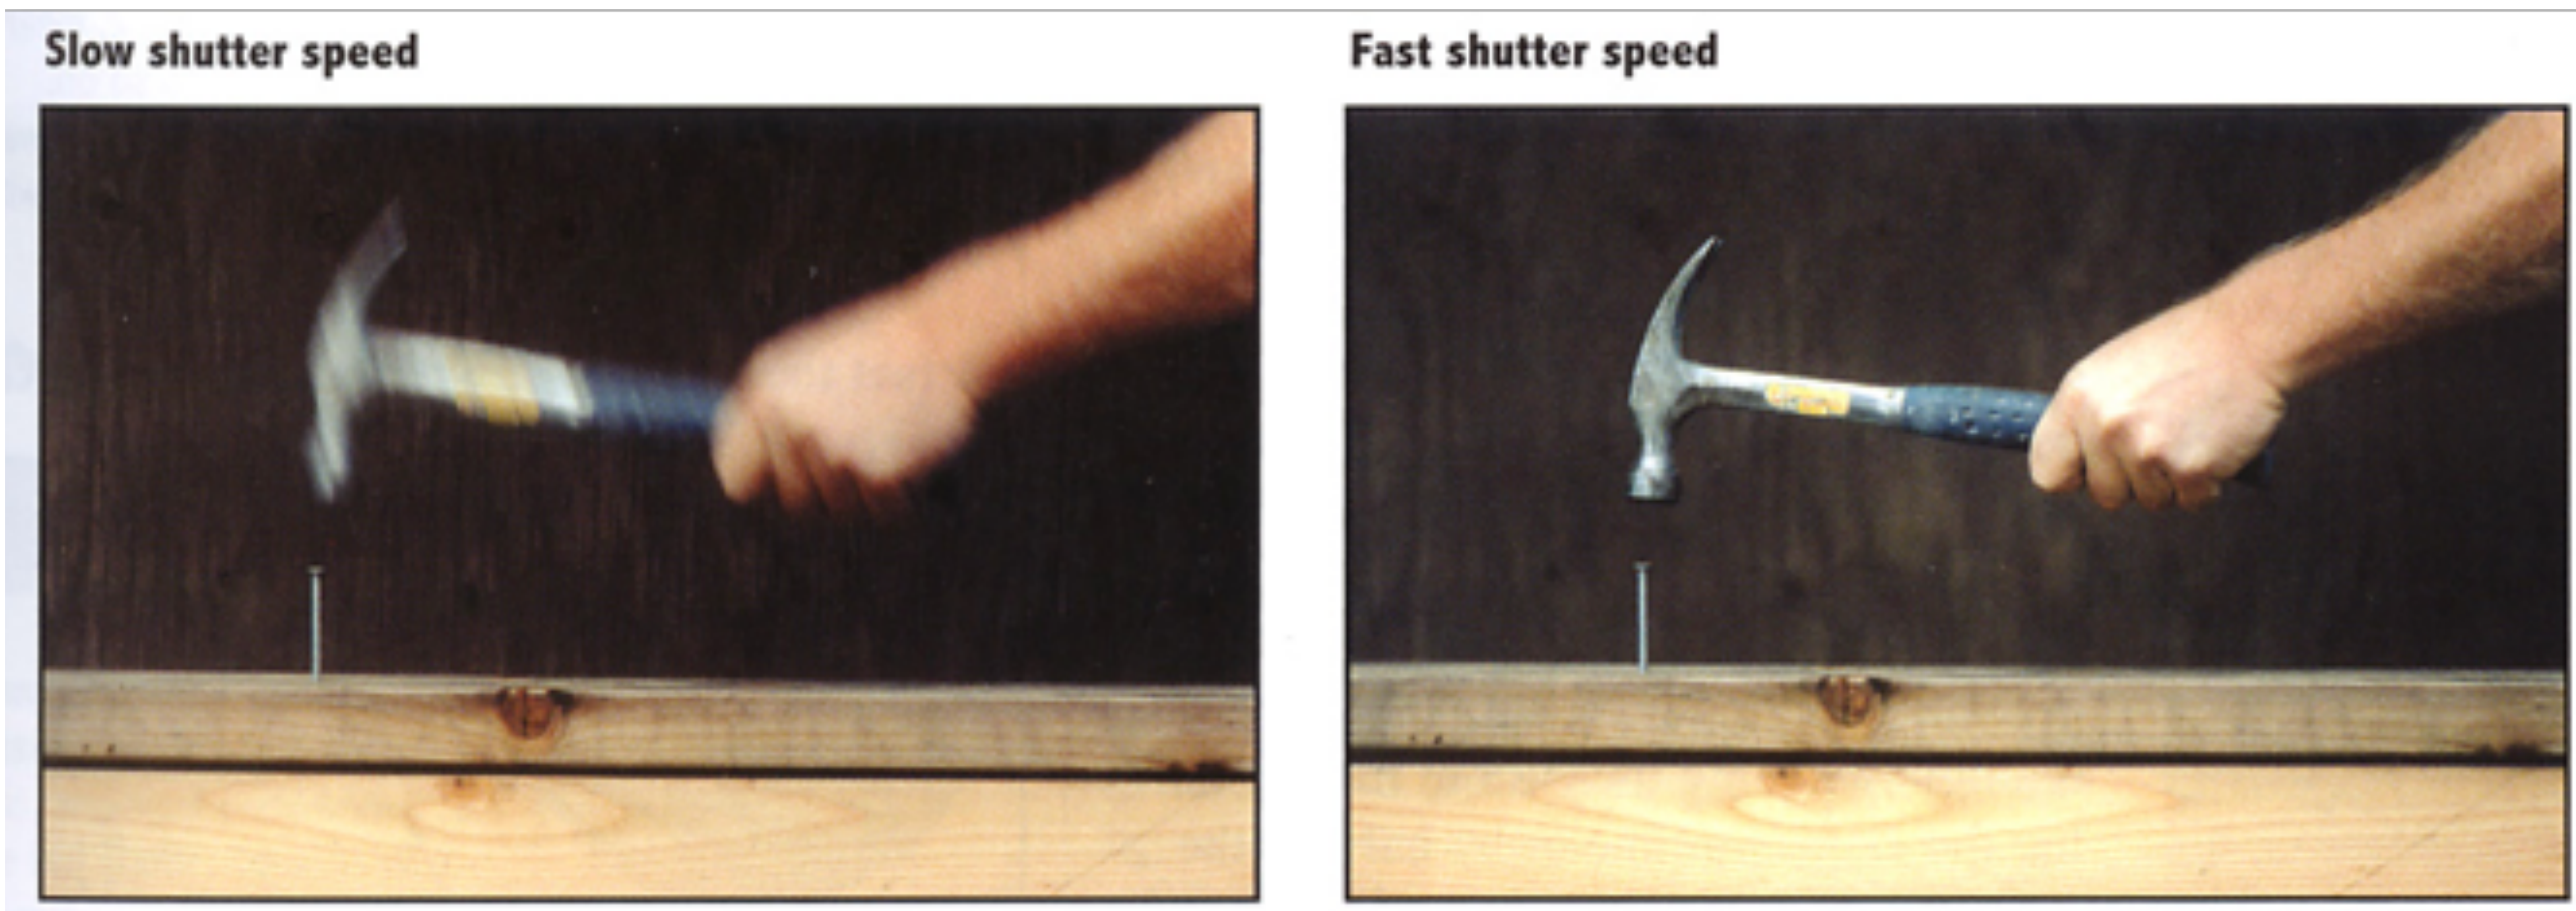

London

Electronic (Rolling) Shutter

- Pixel is electronically reset to start exposure
- Fills with photoelectrons as light falls on sensor
- Reading out pixel electronically "ends" exposure
- Problem: most sensors read out pixels sequentially, takes time (e.g. 1/30 sec) to read entire sensor
  - If reset all pixels at the same time, last pixel read out will have longer exposure
  - So, usually stagger reset of pixels to ensure uniform exposure time
  - Problem: rolling shutter artifact

Electronic Rolling Shutter

Can you explain these images?

Credit: David Adler, B&H Photo Video
https://www.bhphotovideo.com/explora/video/tips-and-solutions/rolling-shutter-versus-global-shutter

Credit: Soren Ragsdale
https://flic.kr/p/5S6rKw

Exposure (Q)

Pixels integrate (time varying) irradiance over time (5D integral)

$$Q(\text{pixel}) = \int_{t \in T} E(\text{pixel}, t) dt$$

$$Q(\text{pixel}) = \int_{t \in T} \int_{x' \in P} \int_{x'' \in D} L(x'' \to x', t) \frac{\cos\theta' \cos\theta''}{||x'' - x'||^2} dA'' dA' dt$$

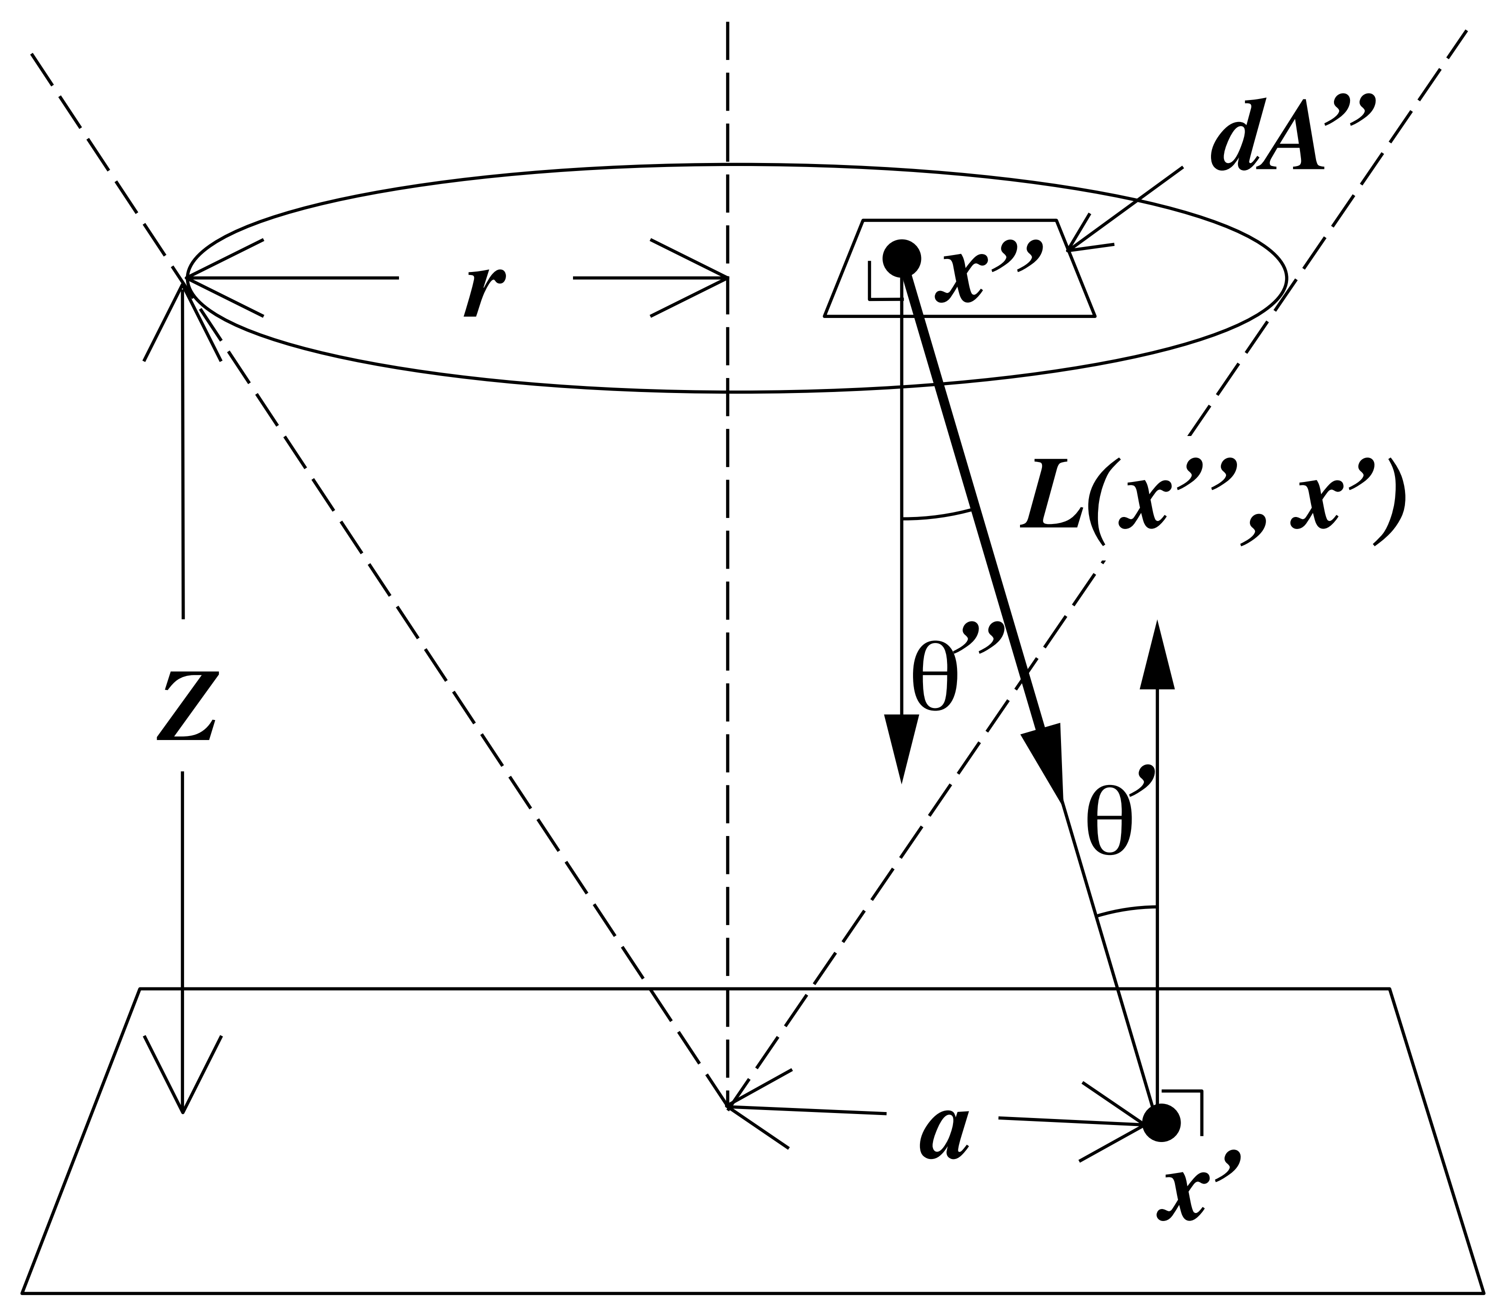

Exposure

- Exposure (Q) is an integral of irradiance over time
- Exposure time (T)
    - Controlled by shutter
- Irradiance (E)
    - Controlled by lens aperture
- Implication: can increase exposure by increasing exposure time or by increasing aperture
    - Aperture: f-stop - 1 “stop" doubles Q

        Note: Also decreases depth of field
    - Shutter: Doubling the open time doubles Q

        Note: Increases motion blur

Aperture vs Shutter

Constant Exposure

f/16
1/8s

f/4
1/125s

f/2
1/500s

From London and Upton

Ray tracing with motion blur

Ray Tracing for Defocus Blur (Thin Lens)

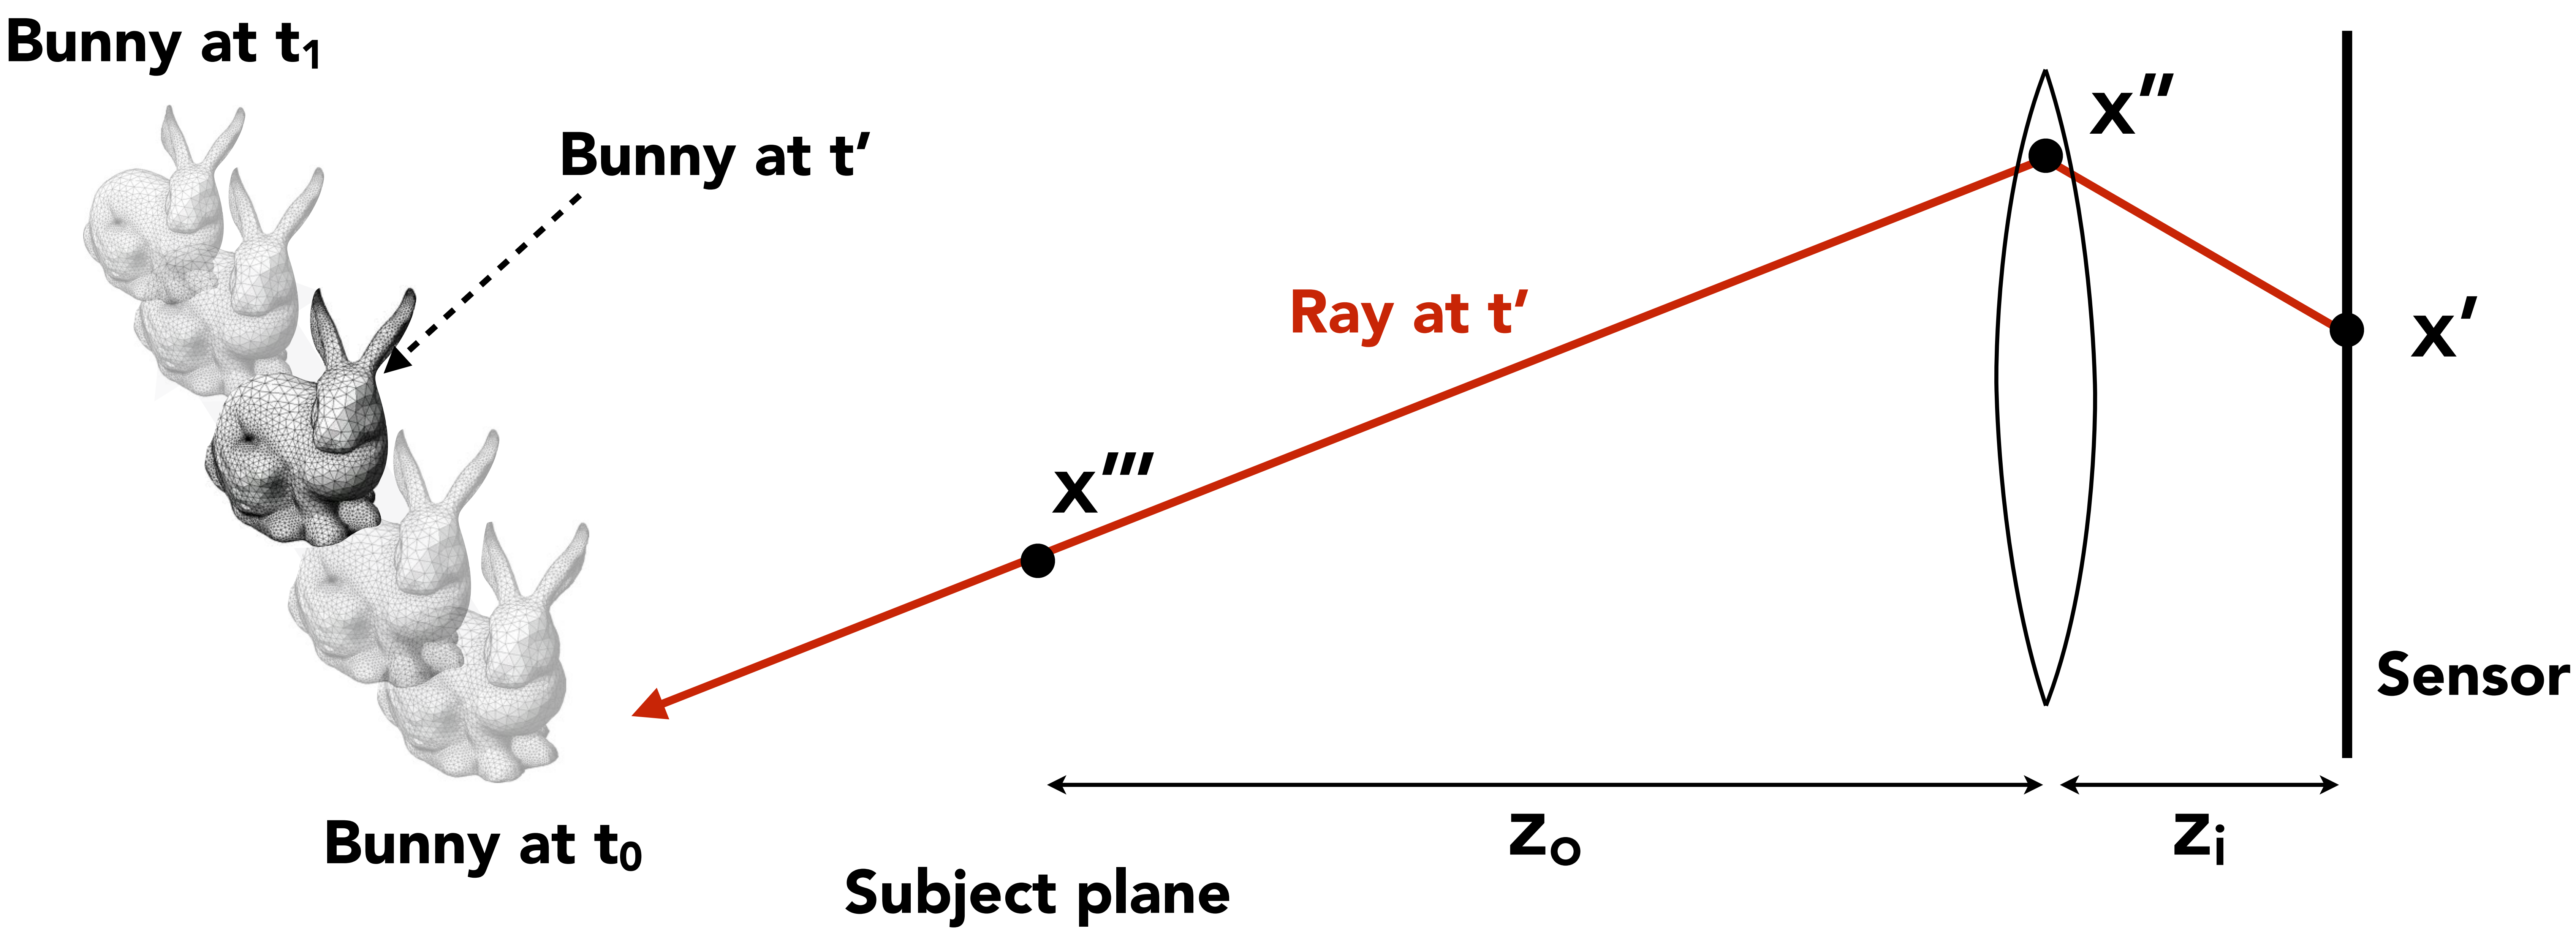

To compute value of pixel at position x′ by Monte Carlo integration:

- Select random time t′ for ray (in addition to x′′ on lens aperture)
- Trace ray through lens system (or use thin lens approximation)
- Update scene geometry to its position at t′
- Ray trace through BVH of geometry at t′

BVH traversal with motion blur

Assume: all primitives provide vertex positions at t_0 (open shutter time) and t_1 (close shutter time)

Assume vertices move linearly during shutter interval

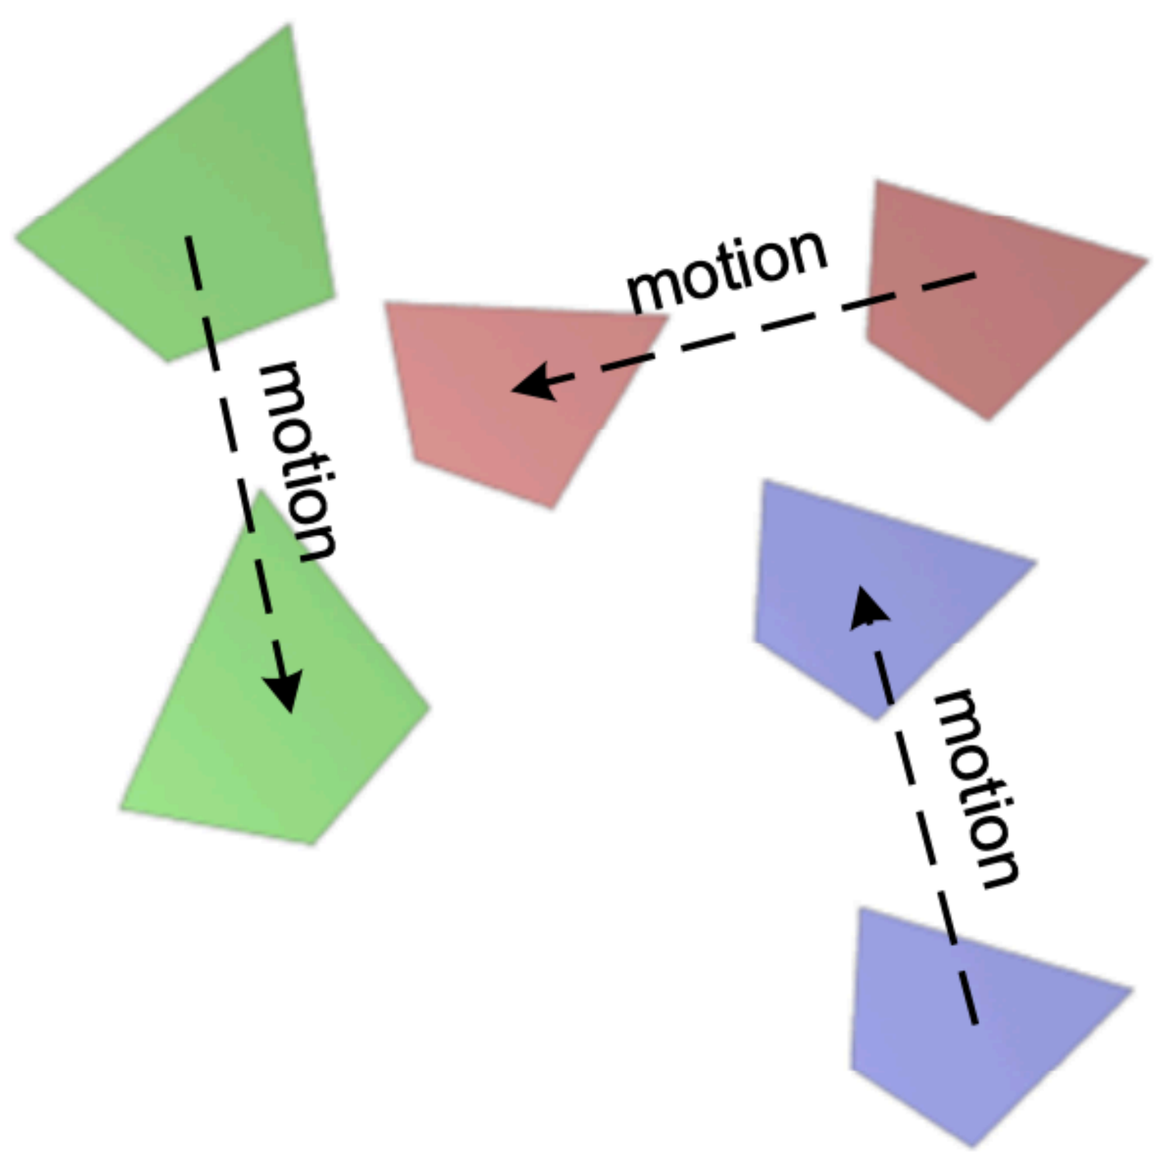

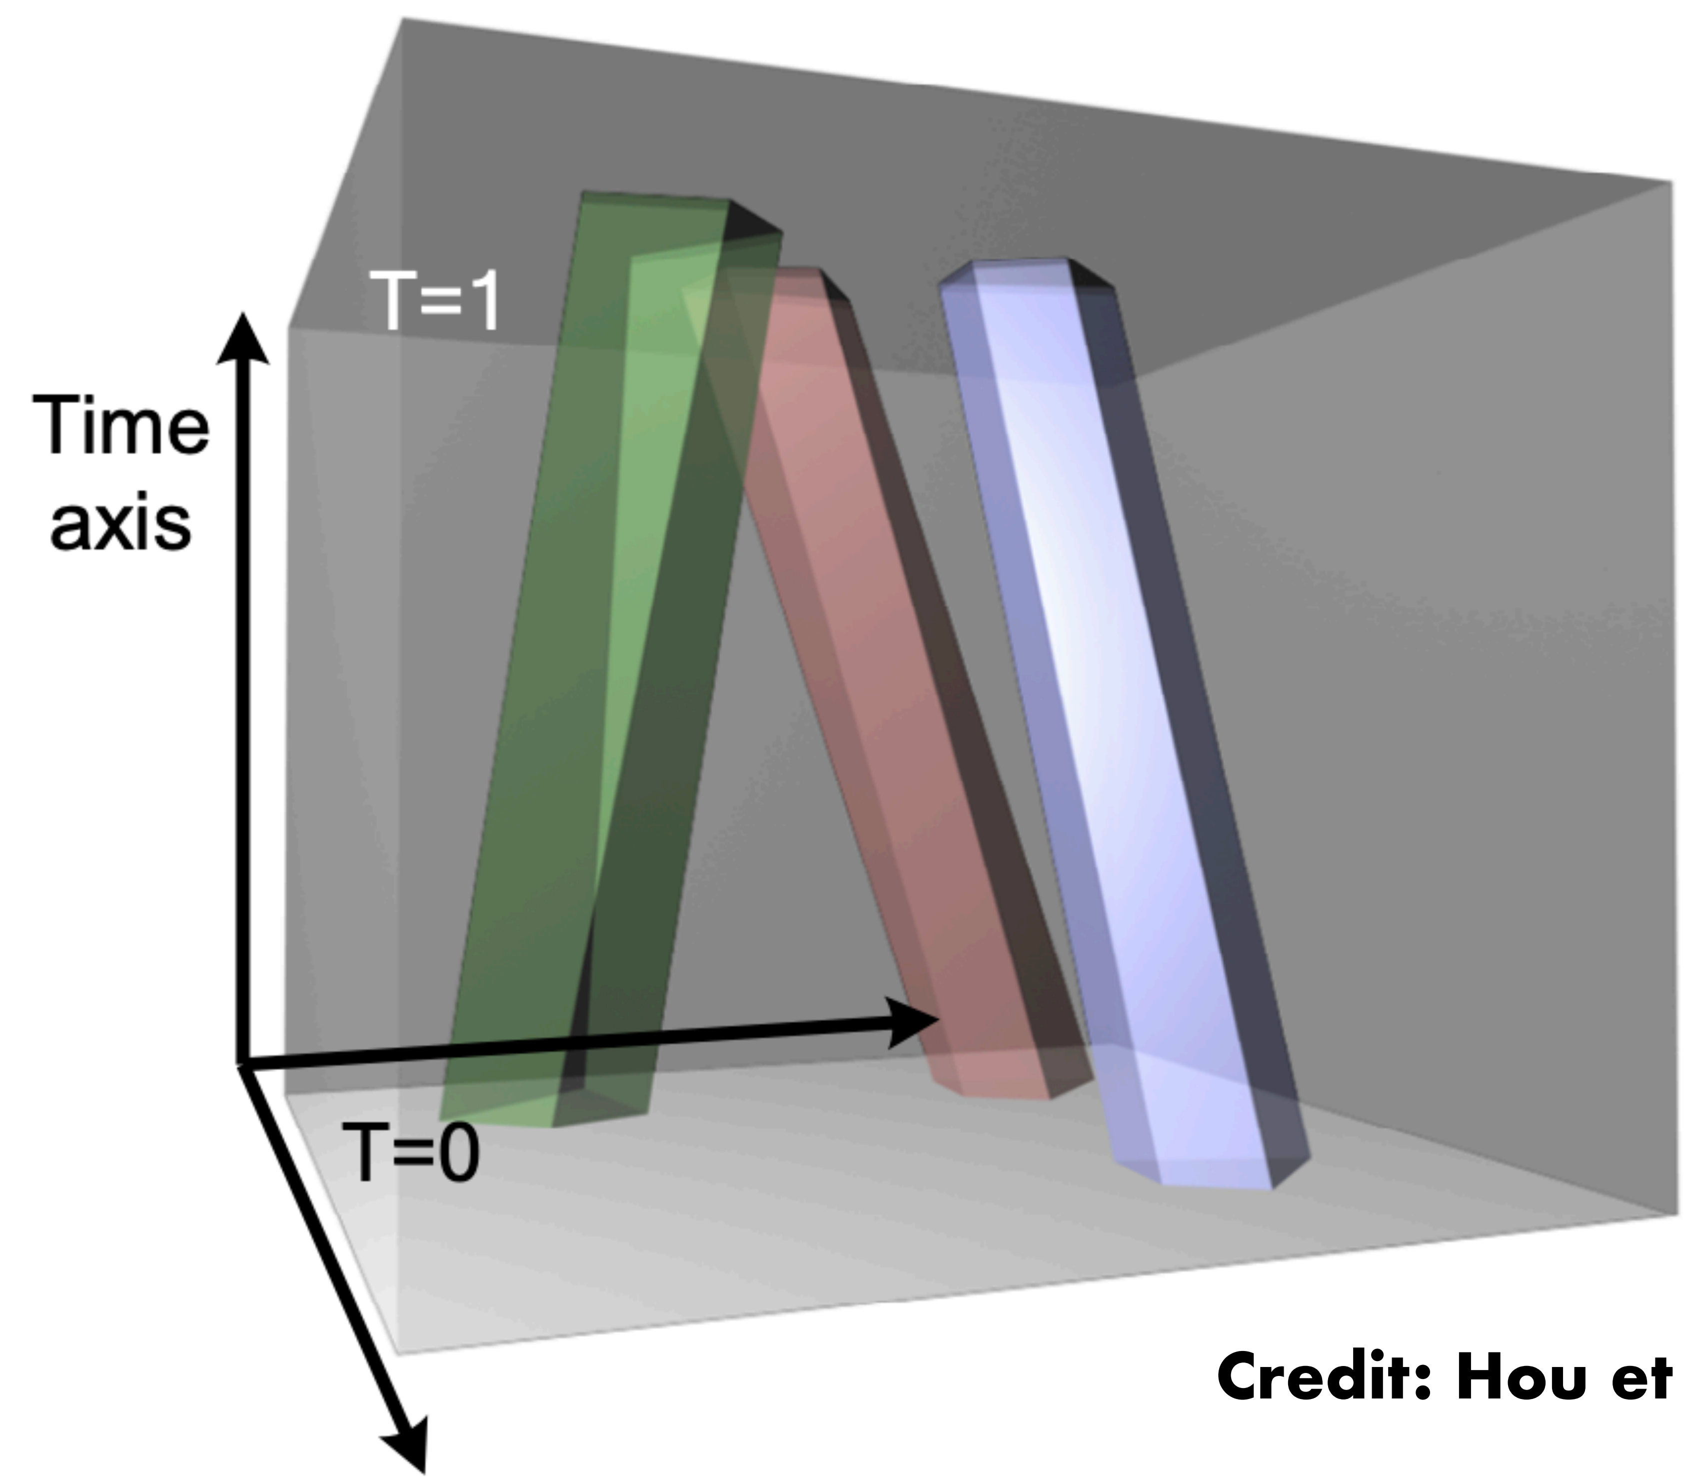

Credit: Hou et al.

BVH traversal with motion blur

BVH construction: compute triangle bbox for t_0 and t_1. Use union($bbox_{t0}$, $bbox_{t1}$) as the triangle's bbox during regular SAH-based BVH construction.

Each BVH node stores its $bbox_{t0}$ and $bbox_{t1}$.

During traversal: given ray at time t:

- Compute $bbox(n)_t$ for node n at time t via linear interpolation of $bbox(n)_{t0}$ and $bbox(n)_{t1}$.
- Intersect ray with $bbox(n)_t$ as normal

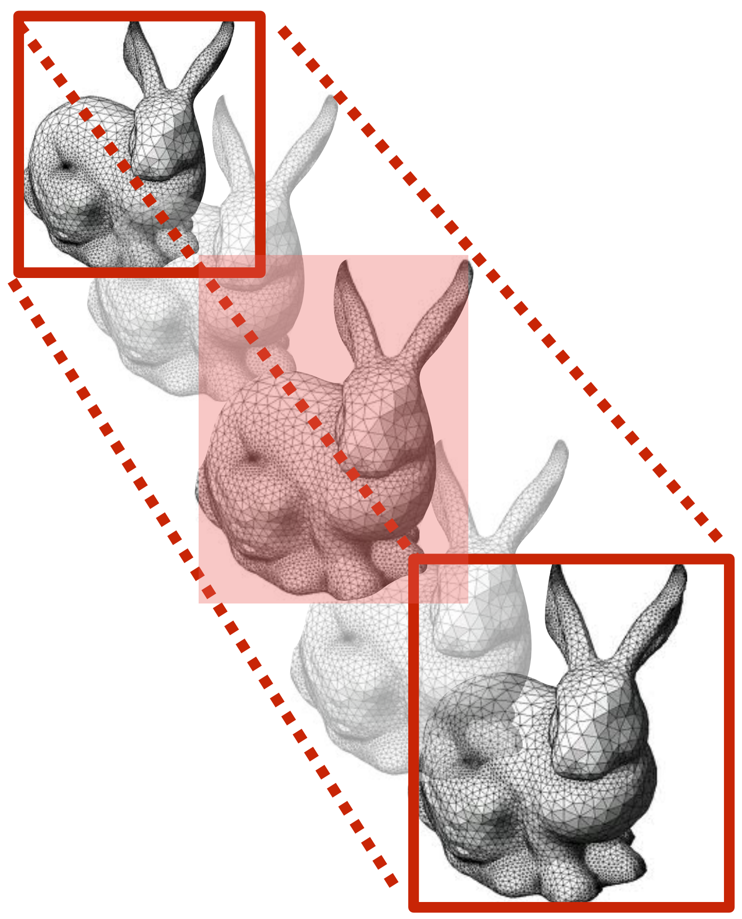

Challenges

Linear motion over a frame may be insufficient to capture fast non-linear motions. (e.g, walk cycles, rotating wheels)

Credit: Houdini docs

1 linear segment — 4 segments — 9 segments

Solution: piecewise linear motion approximation — store multiple sets of vertex positions over shutter interval

Memory cost: must store many bboxes for each node in BVH

Challenges

Fast motion can significantly bloat primitive bounds used for SAH, resulting in inefficient BVH trees

Idea: enhance splitting algorithm to partition time as well as objects

BVH nodes now have both space and time bounds.

(Ray misses box if ray's t does not overlap box's time bound)

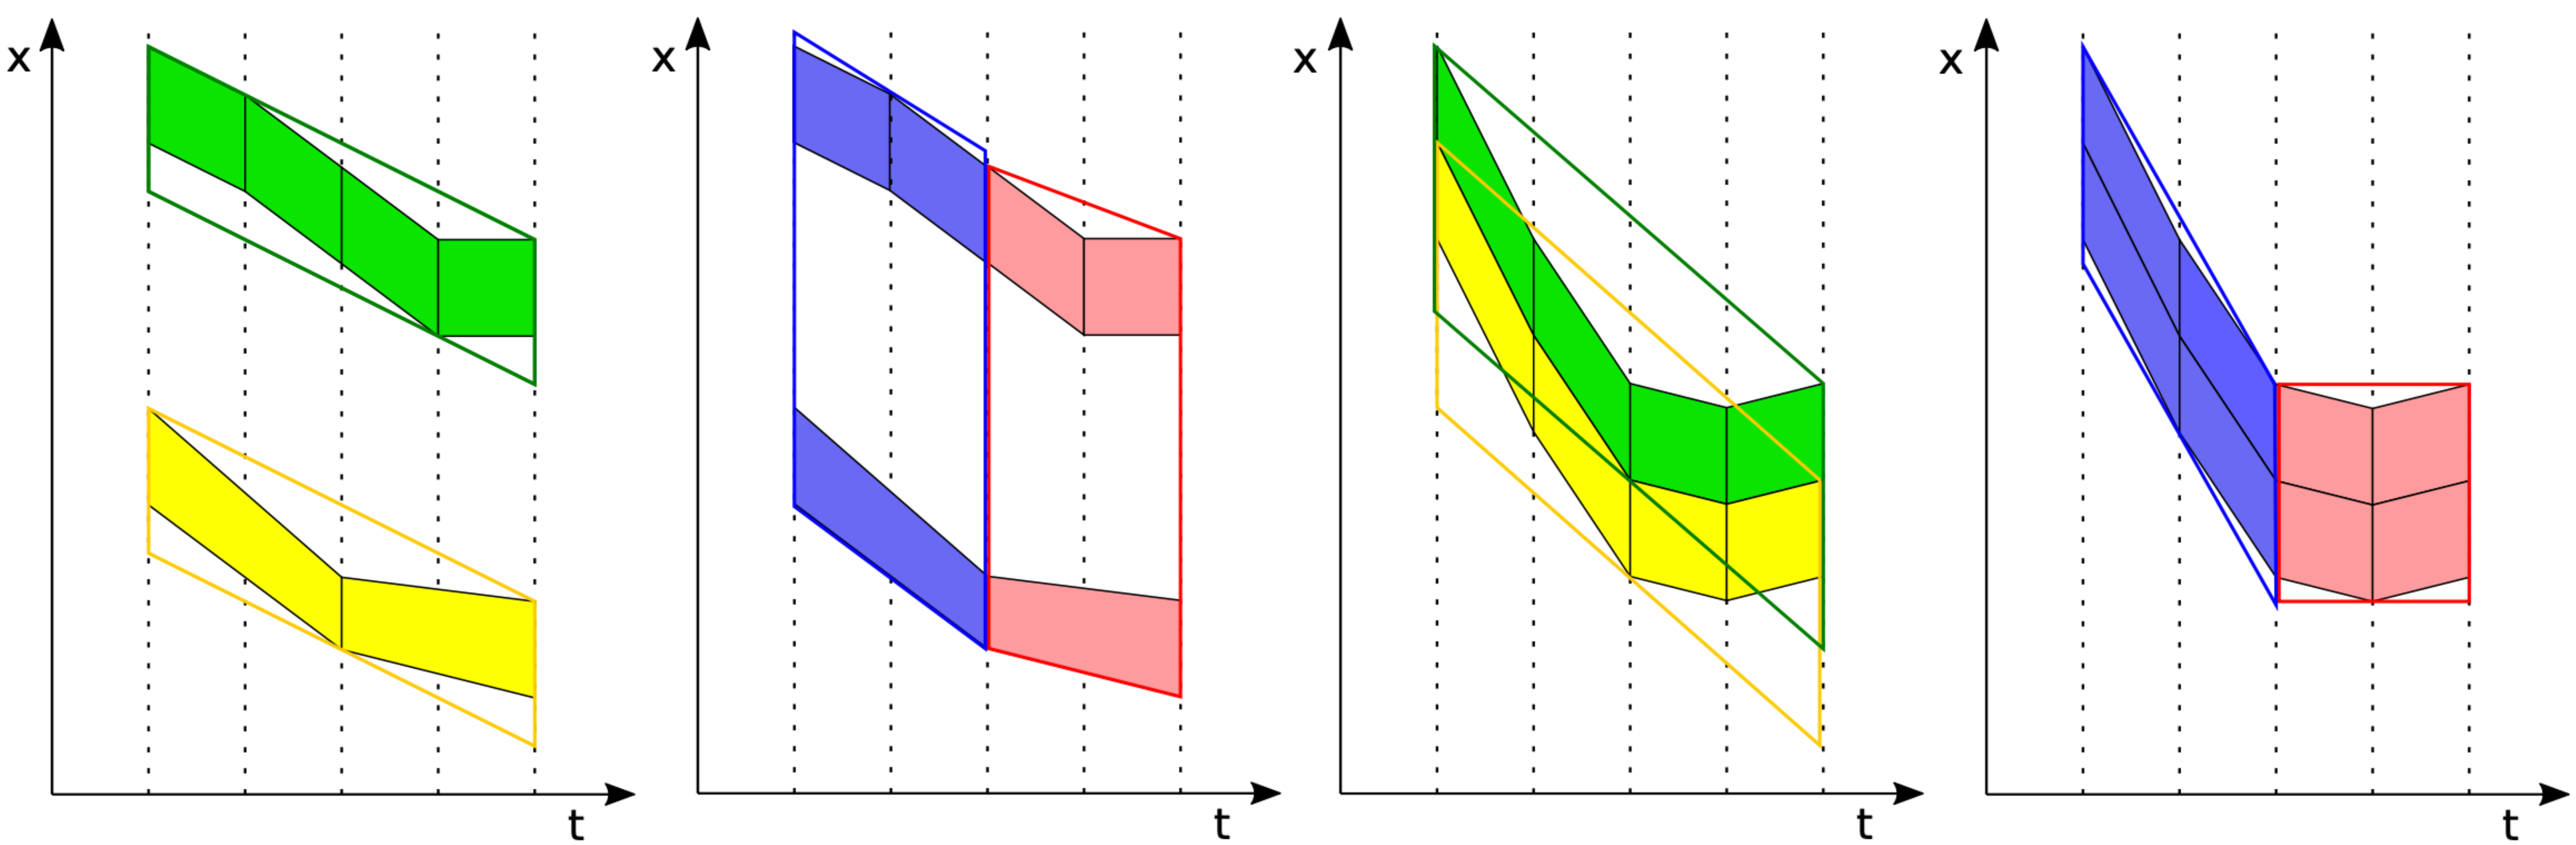

(a) Example 1: Object partitioning (b) Example 1: Temporal partitioning (c) Example 2: Object partitioning (d) Example 2: Temporal partitioning

See Woop et al. 2017 (STBVH)

Frame from “Turning Red” trailer

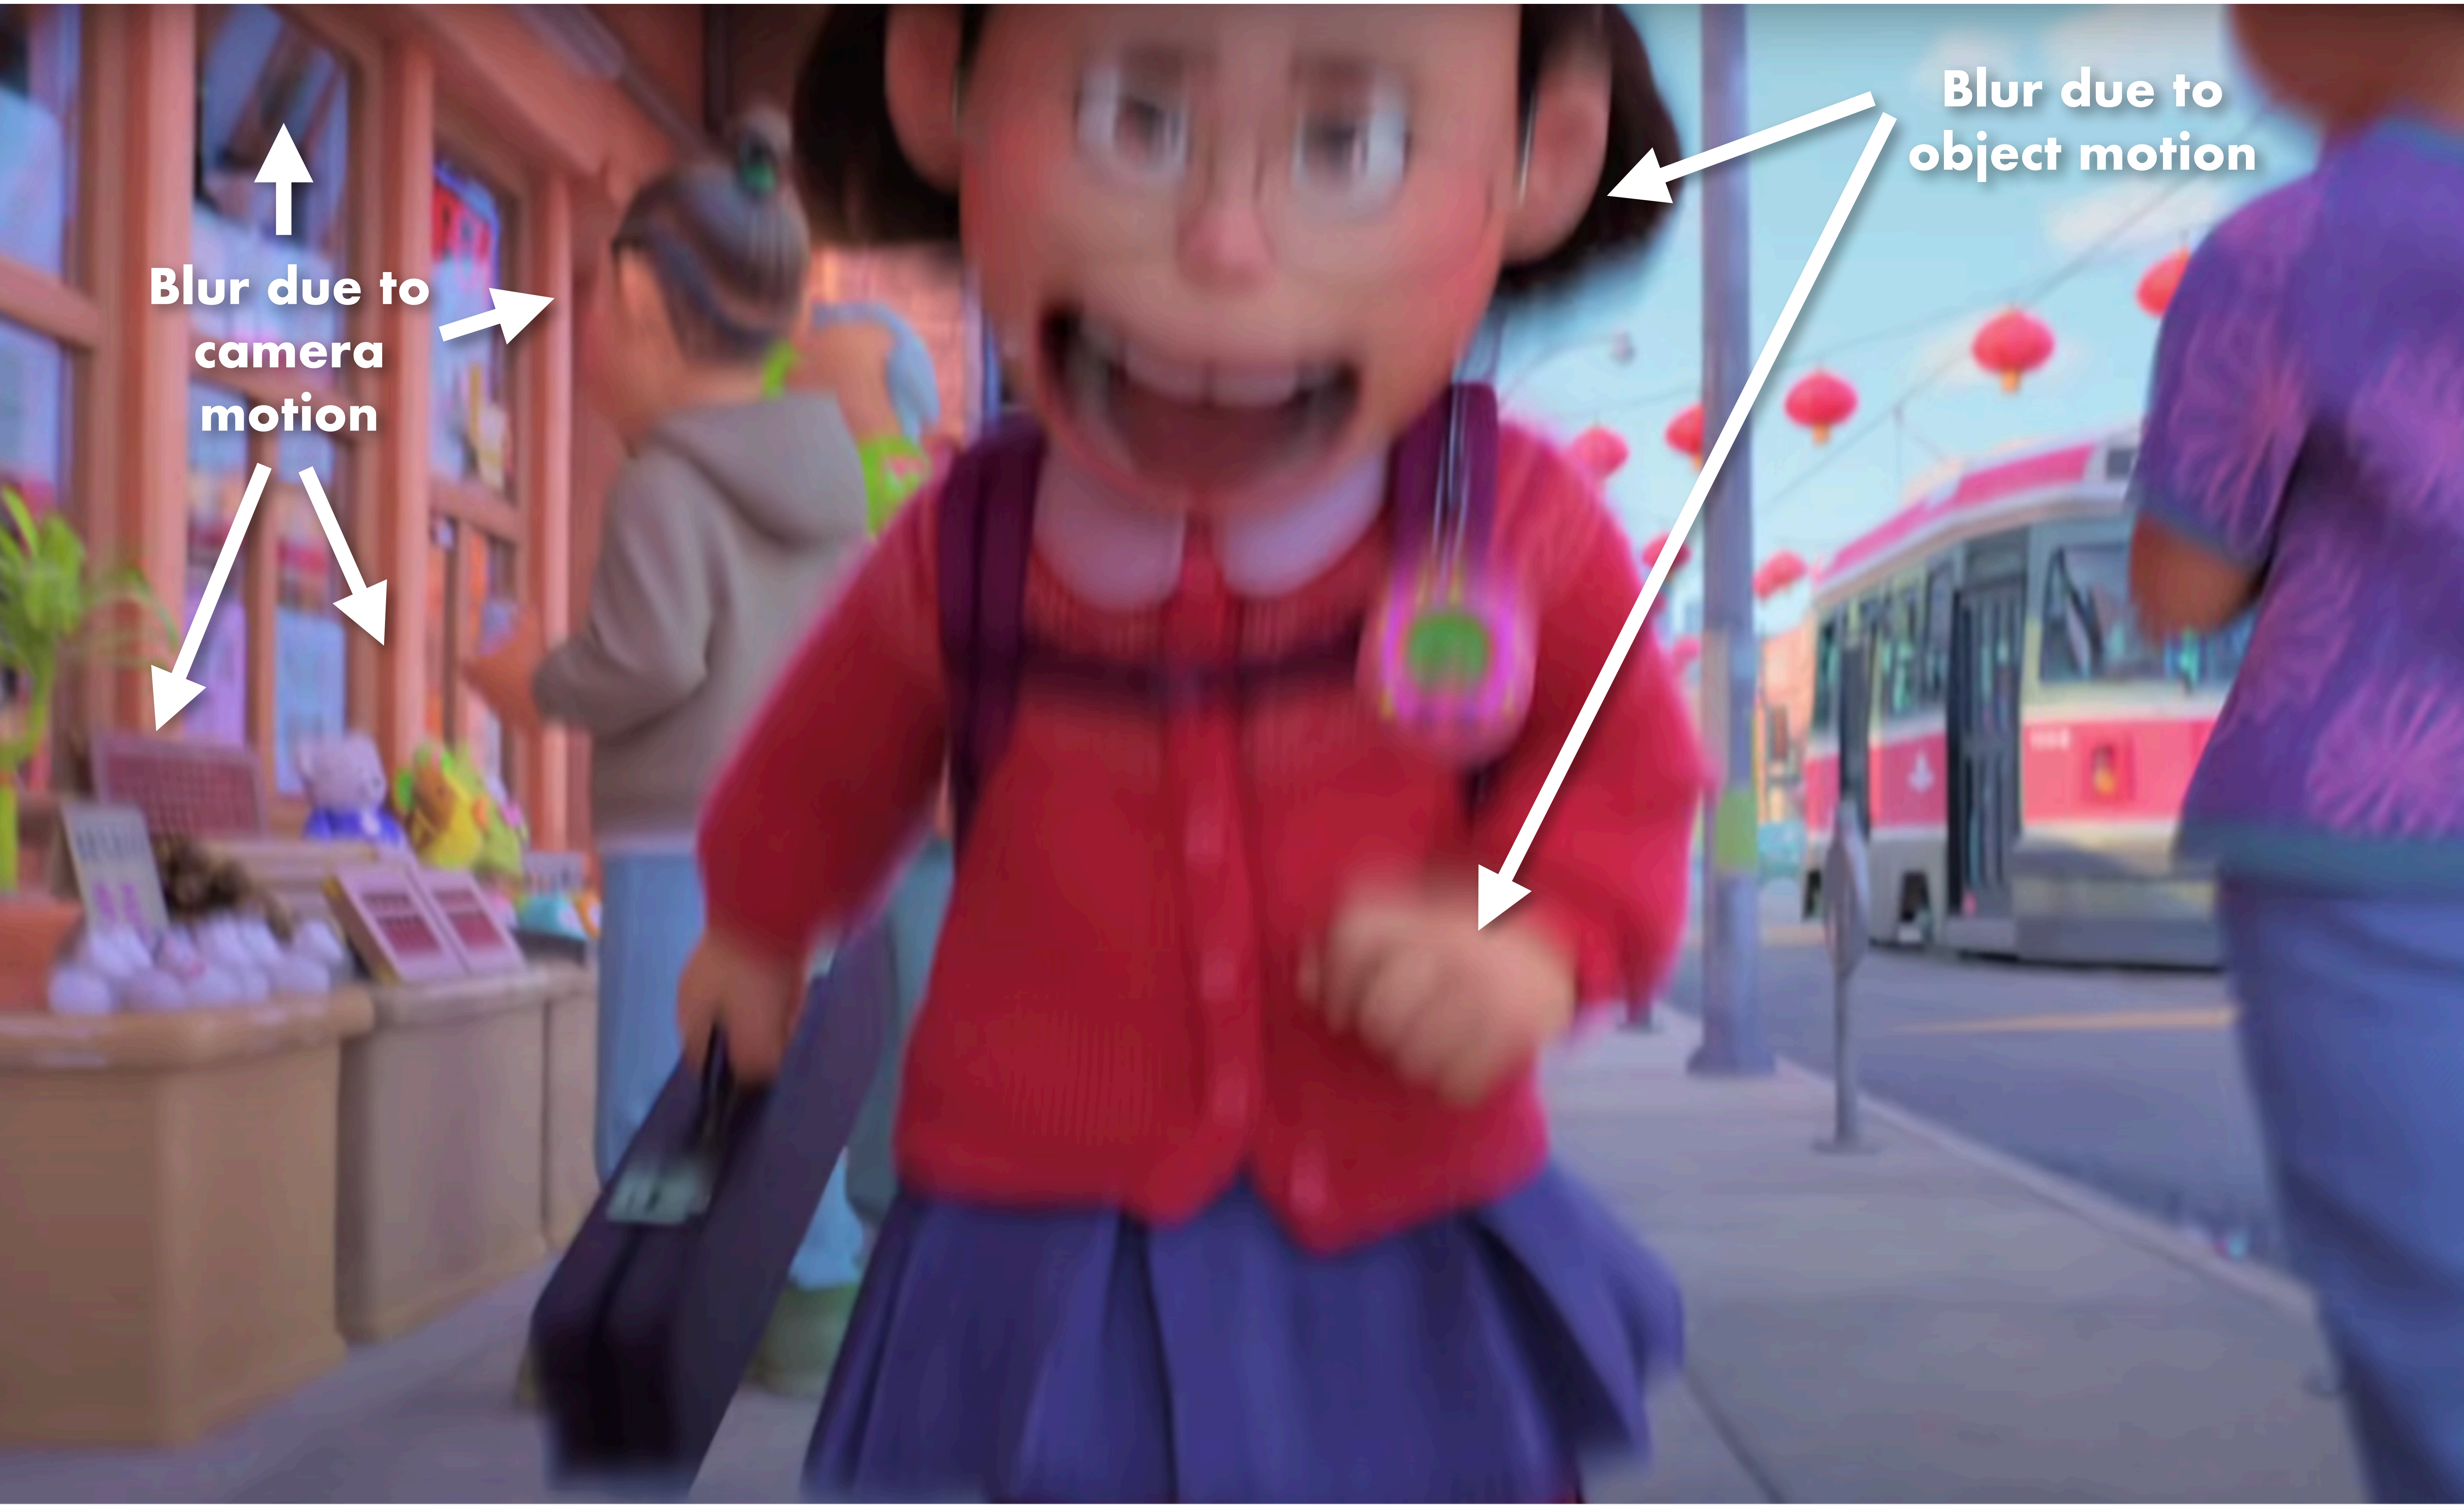

Light Field Cameras

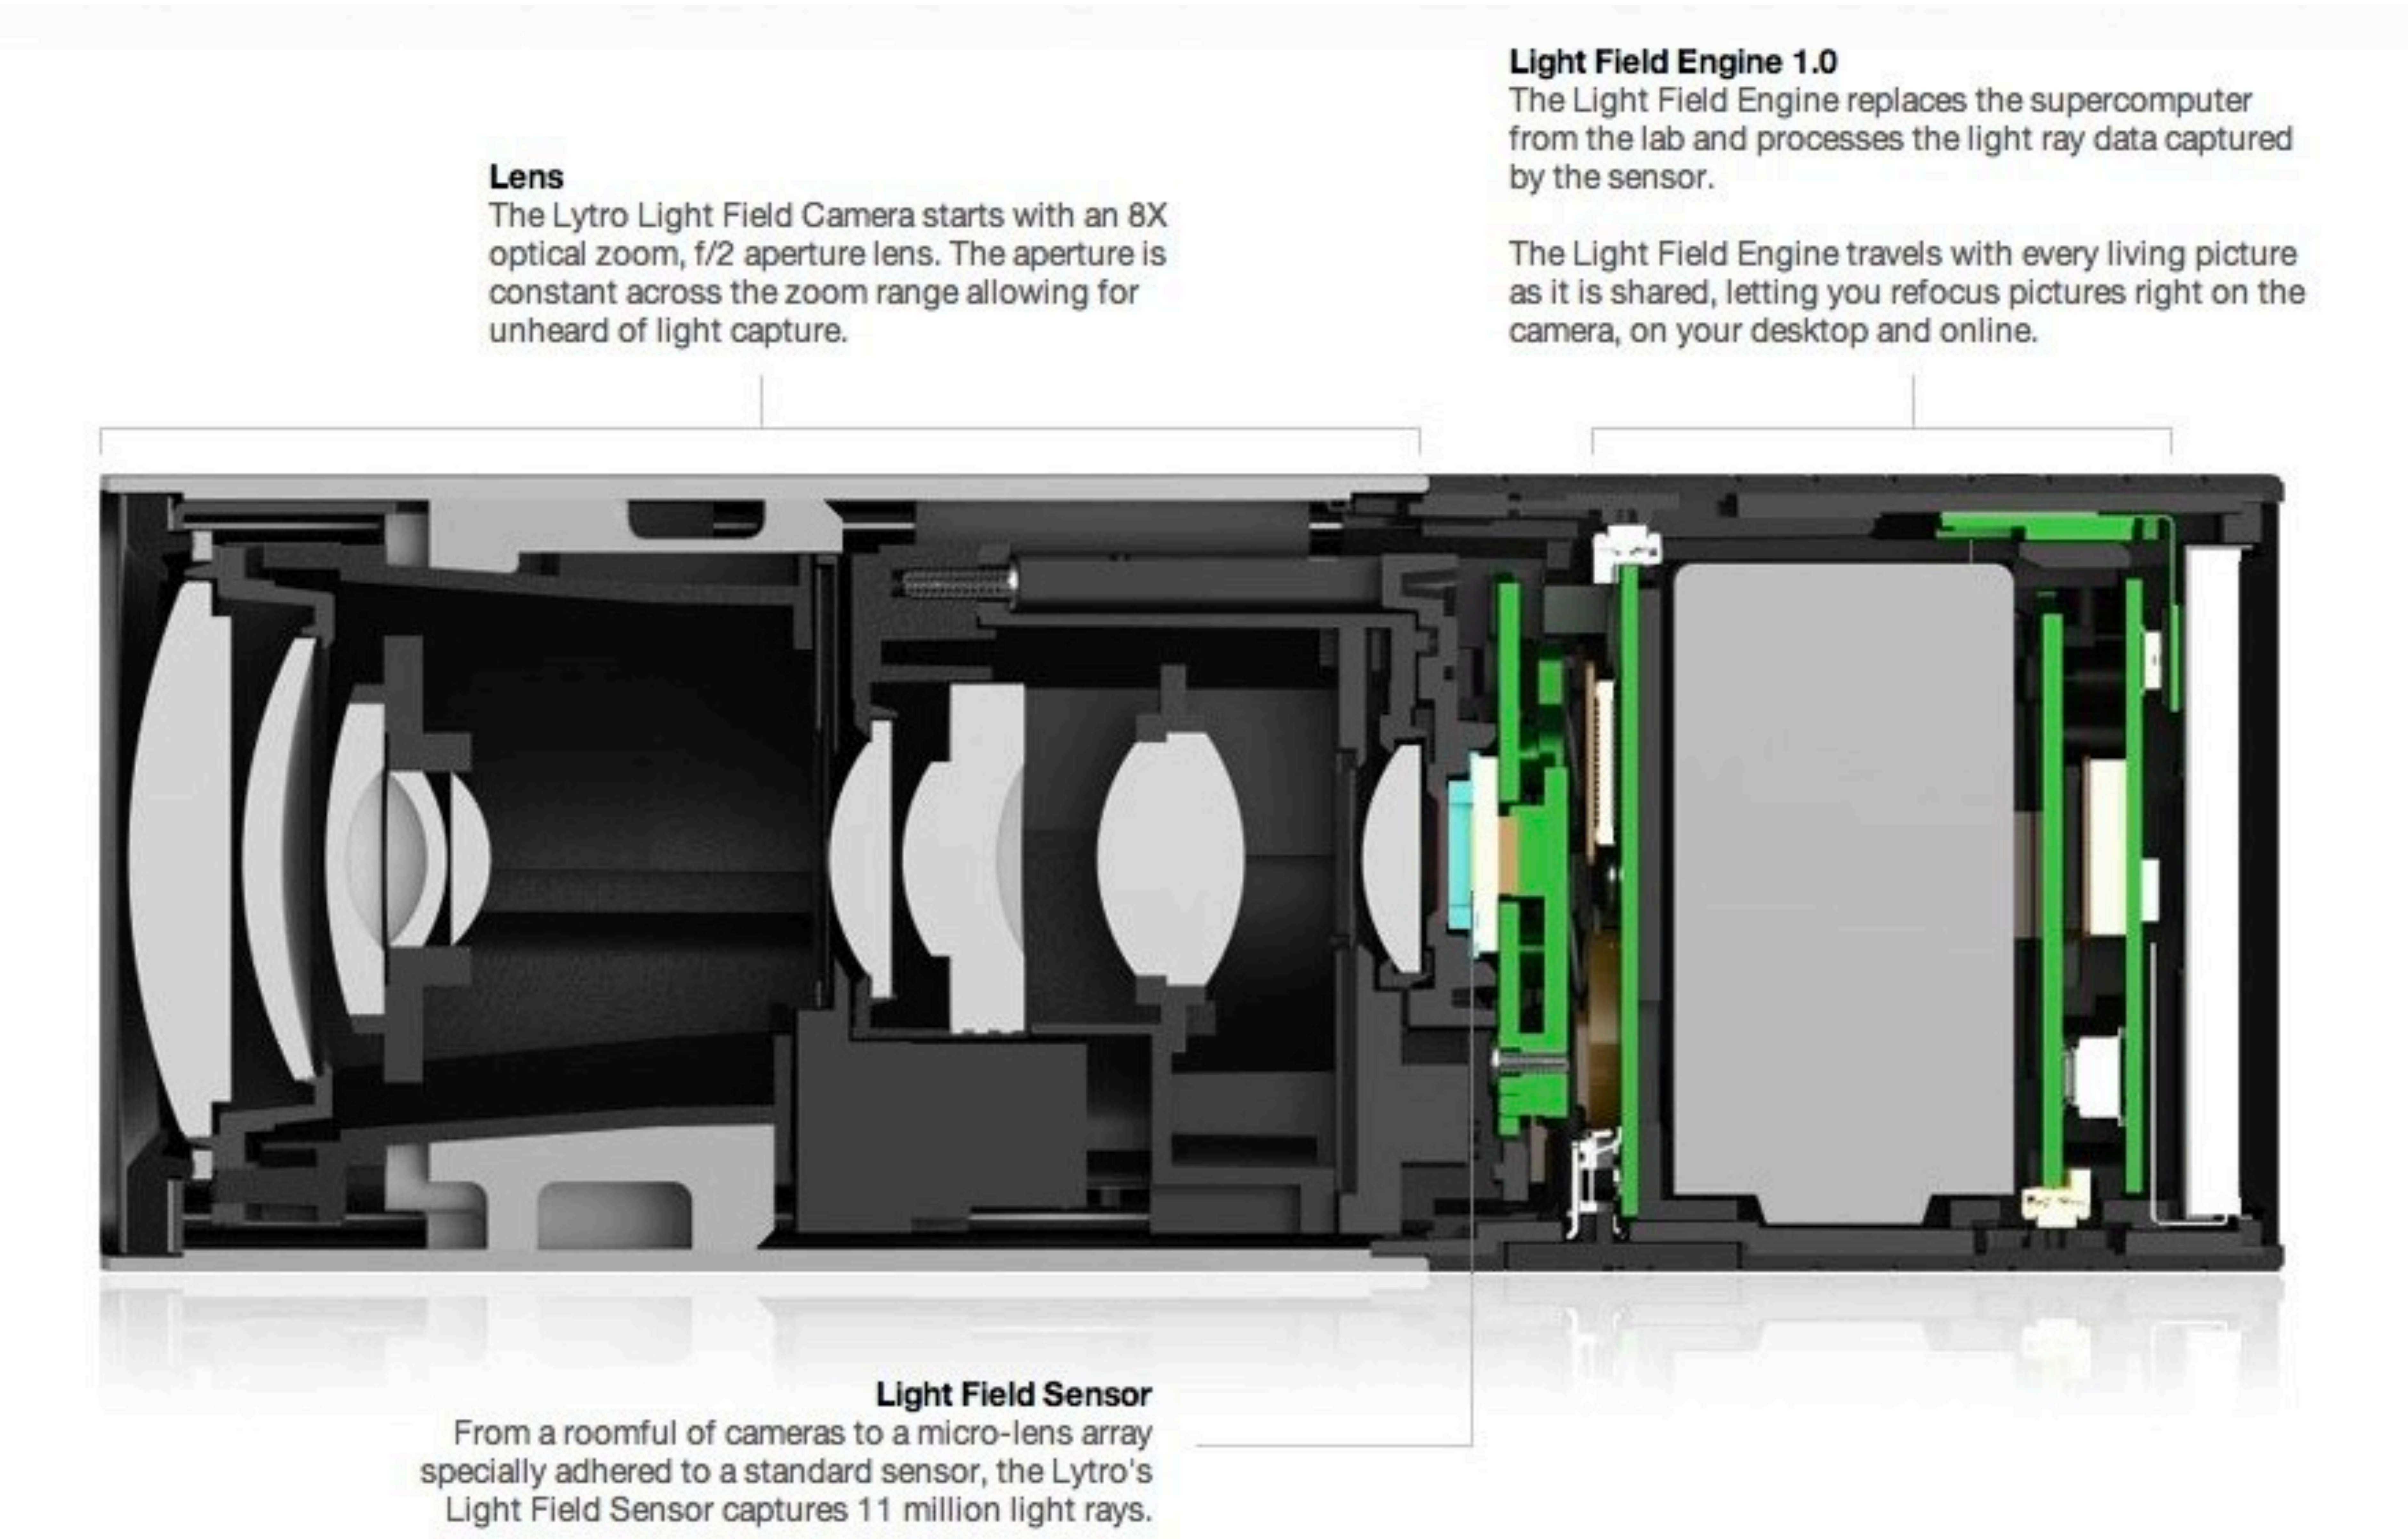

https://www.wired.com/2011/10/cross-section-shows-lytro-light-field-cameras-insides/

PIXAR Super Lightfield Lens

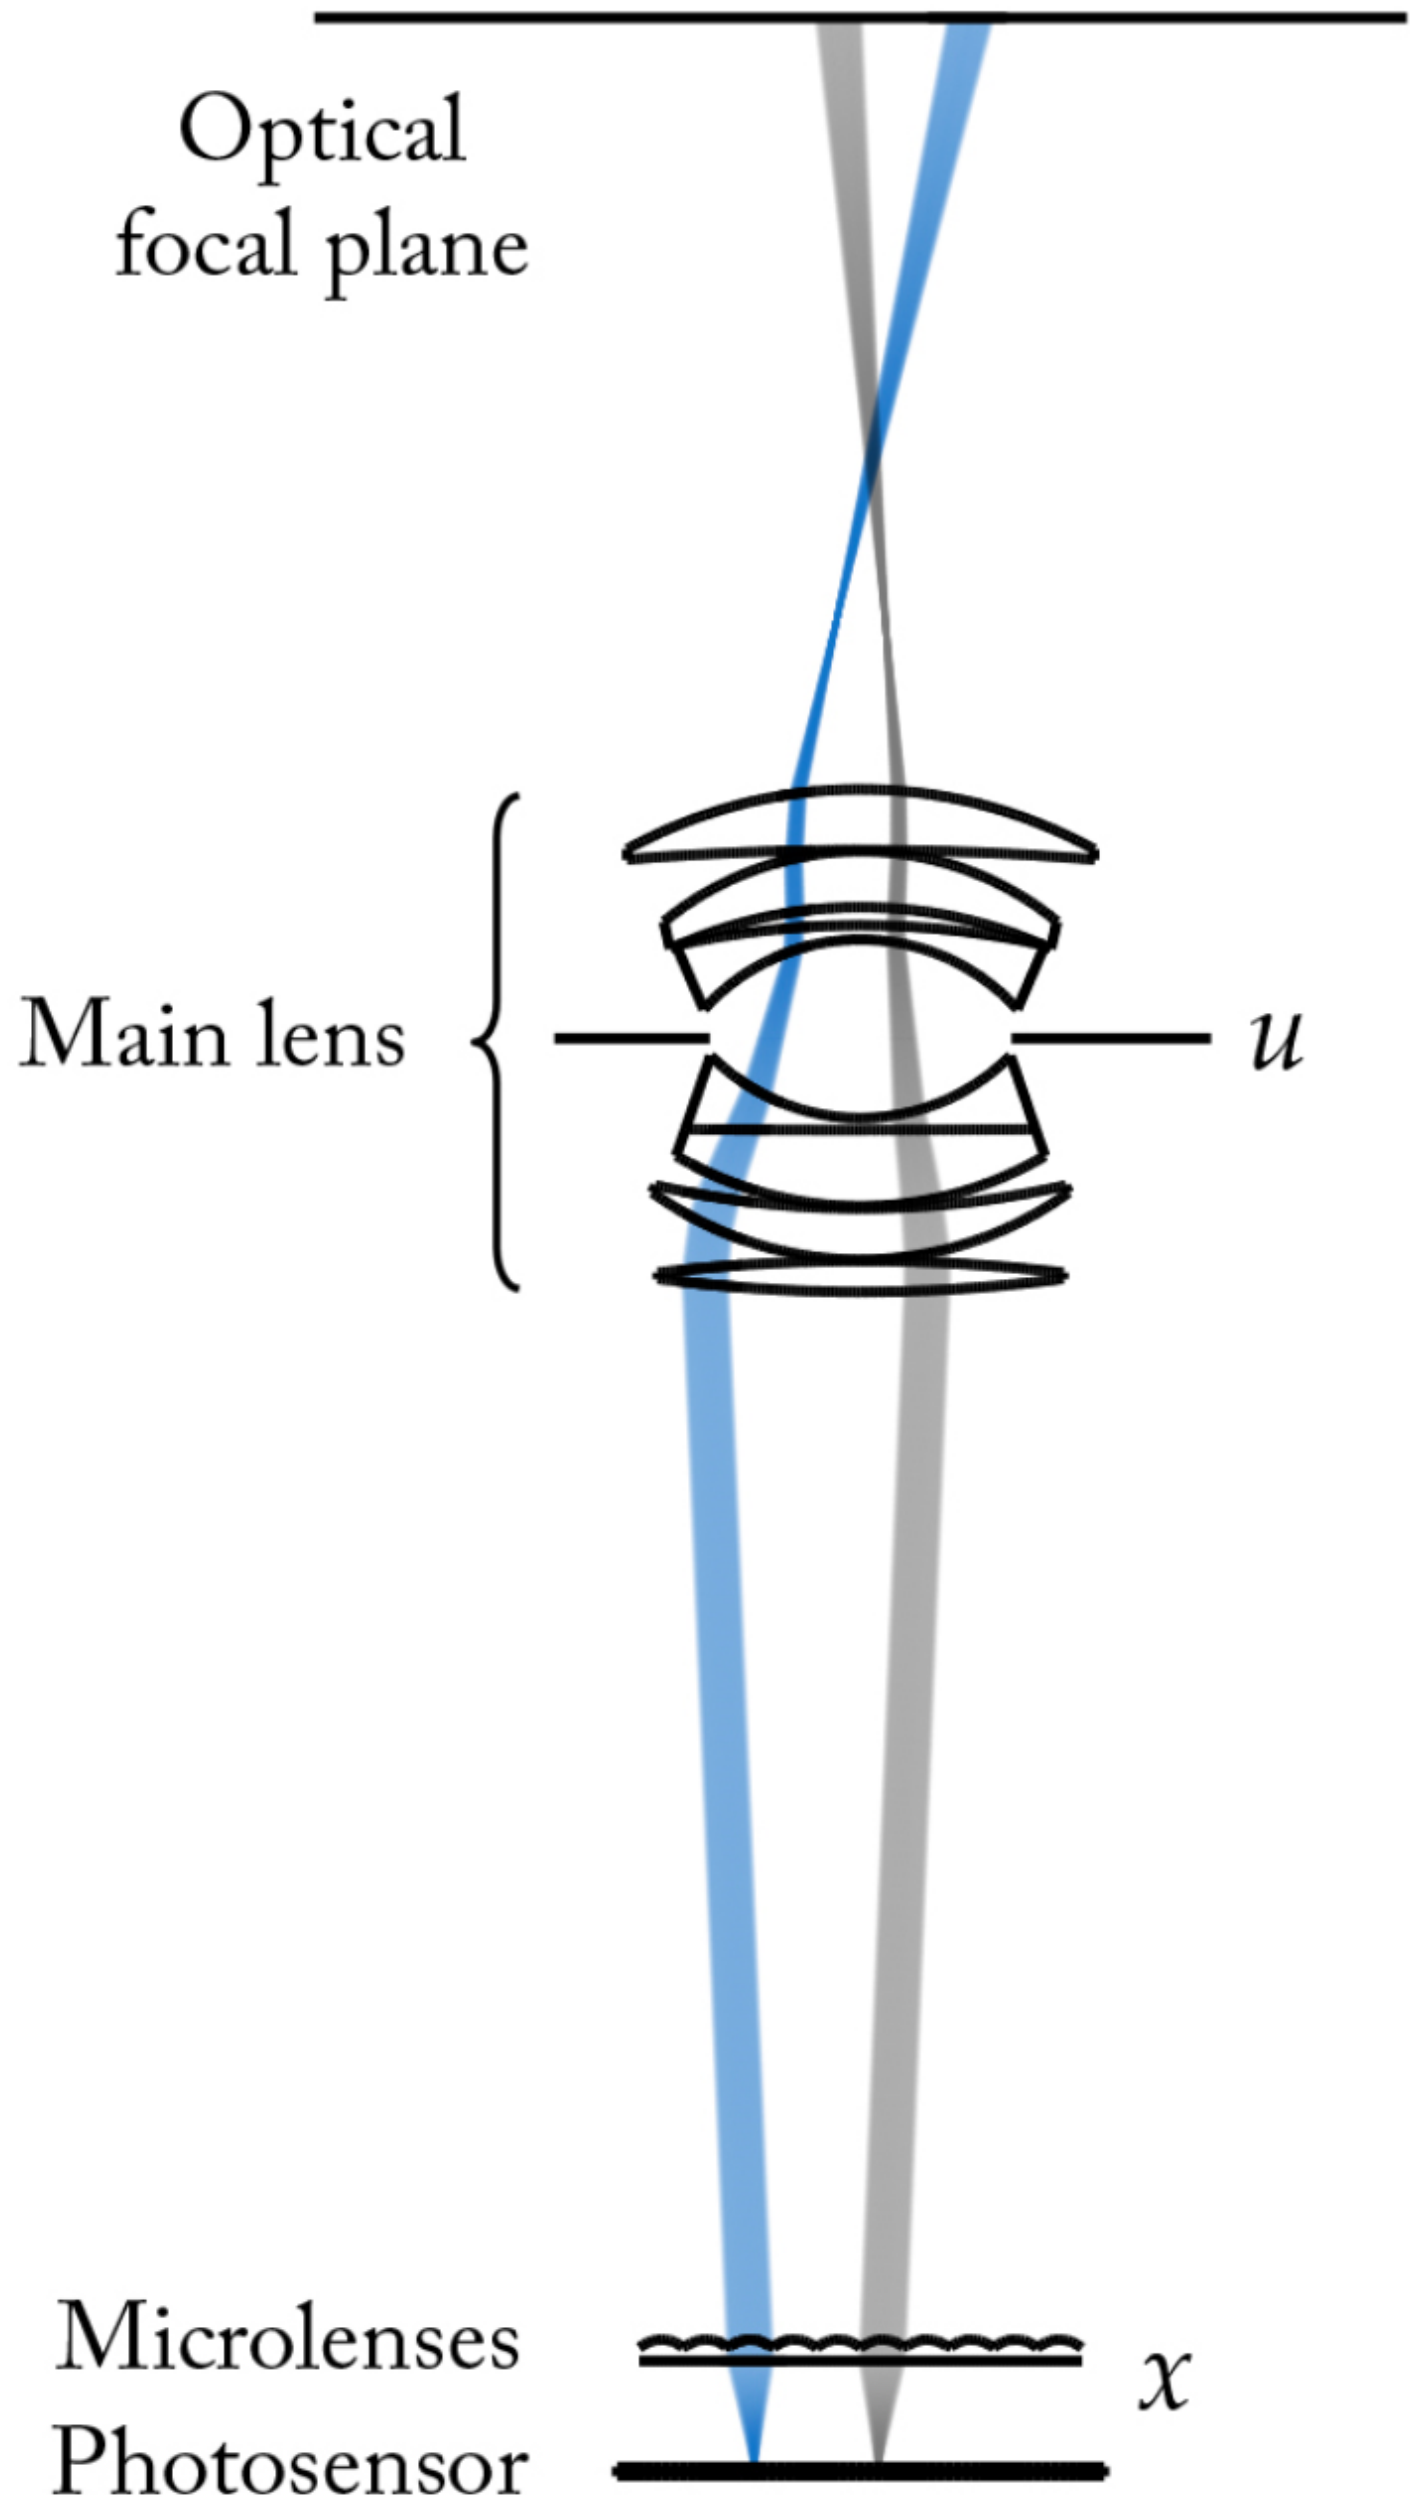

Each microlens (and the pixels underneath it) form a "mini" camera taking a picture of the back of the main lens

Stanford camera array

Assignment 3

Part 1. Create light field image using pbrt

Assignment 3

Part 1. Create light field image using pbrt

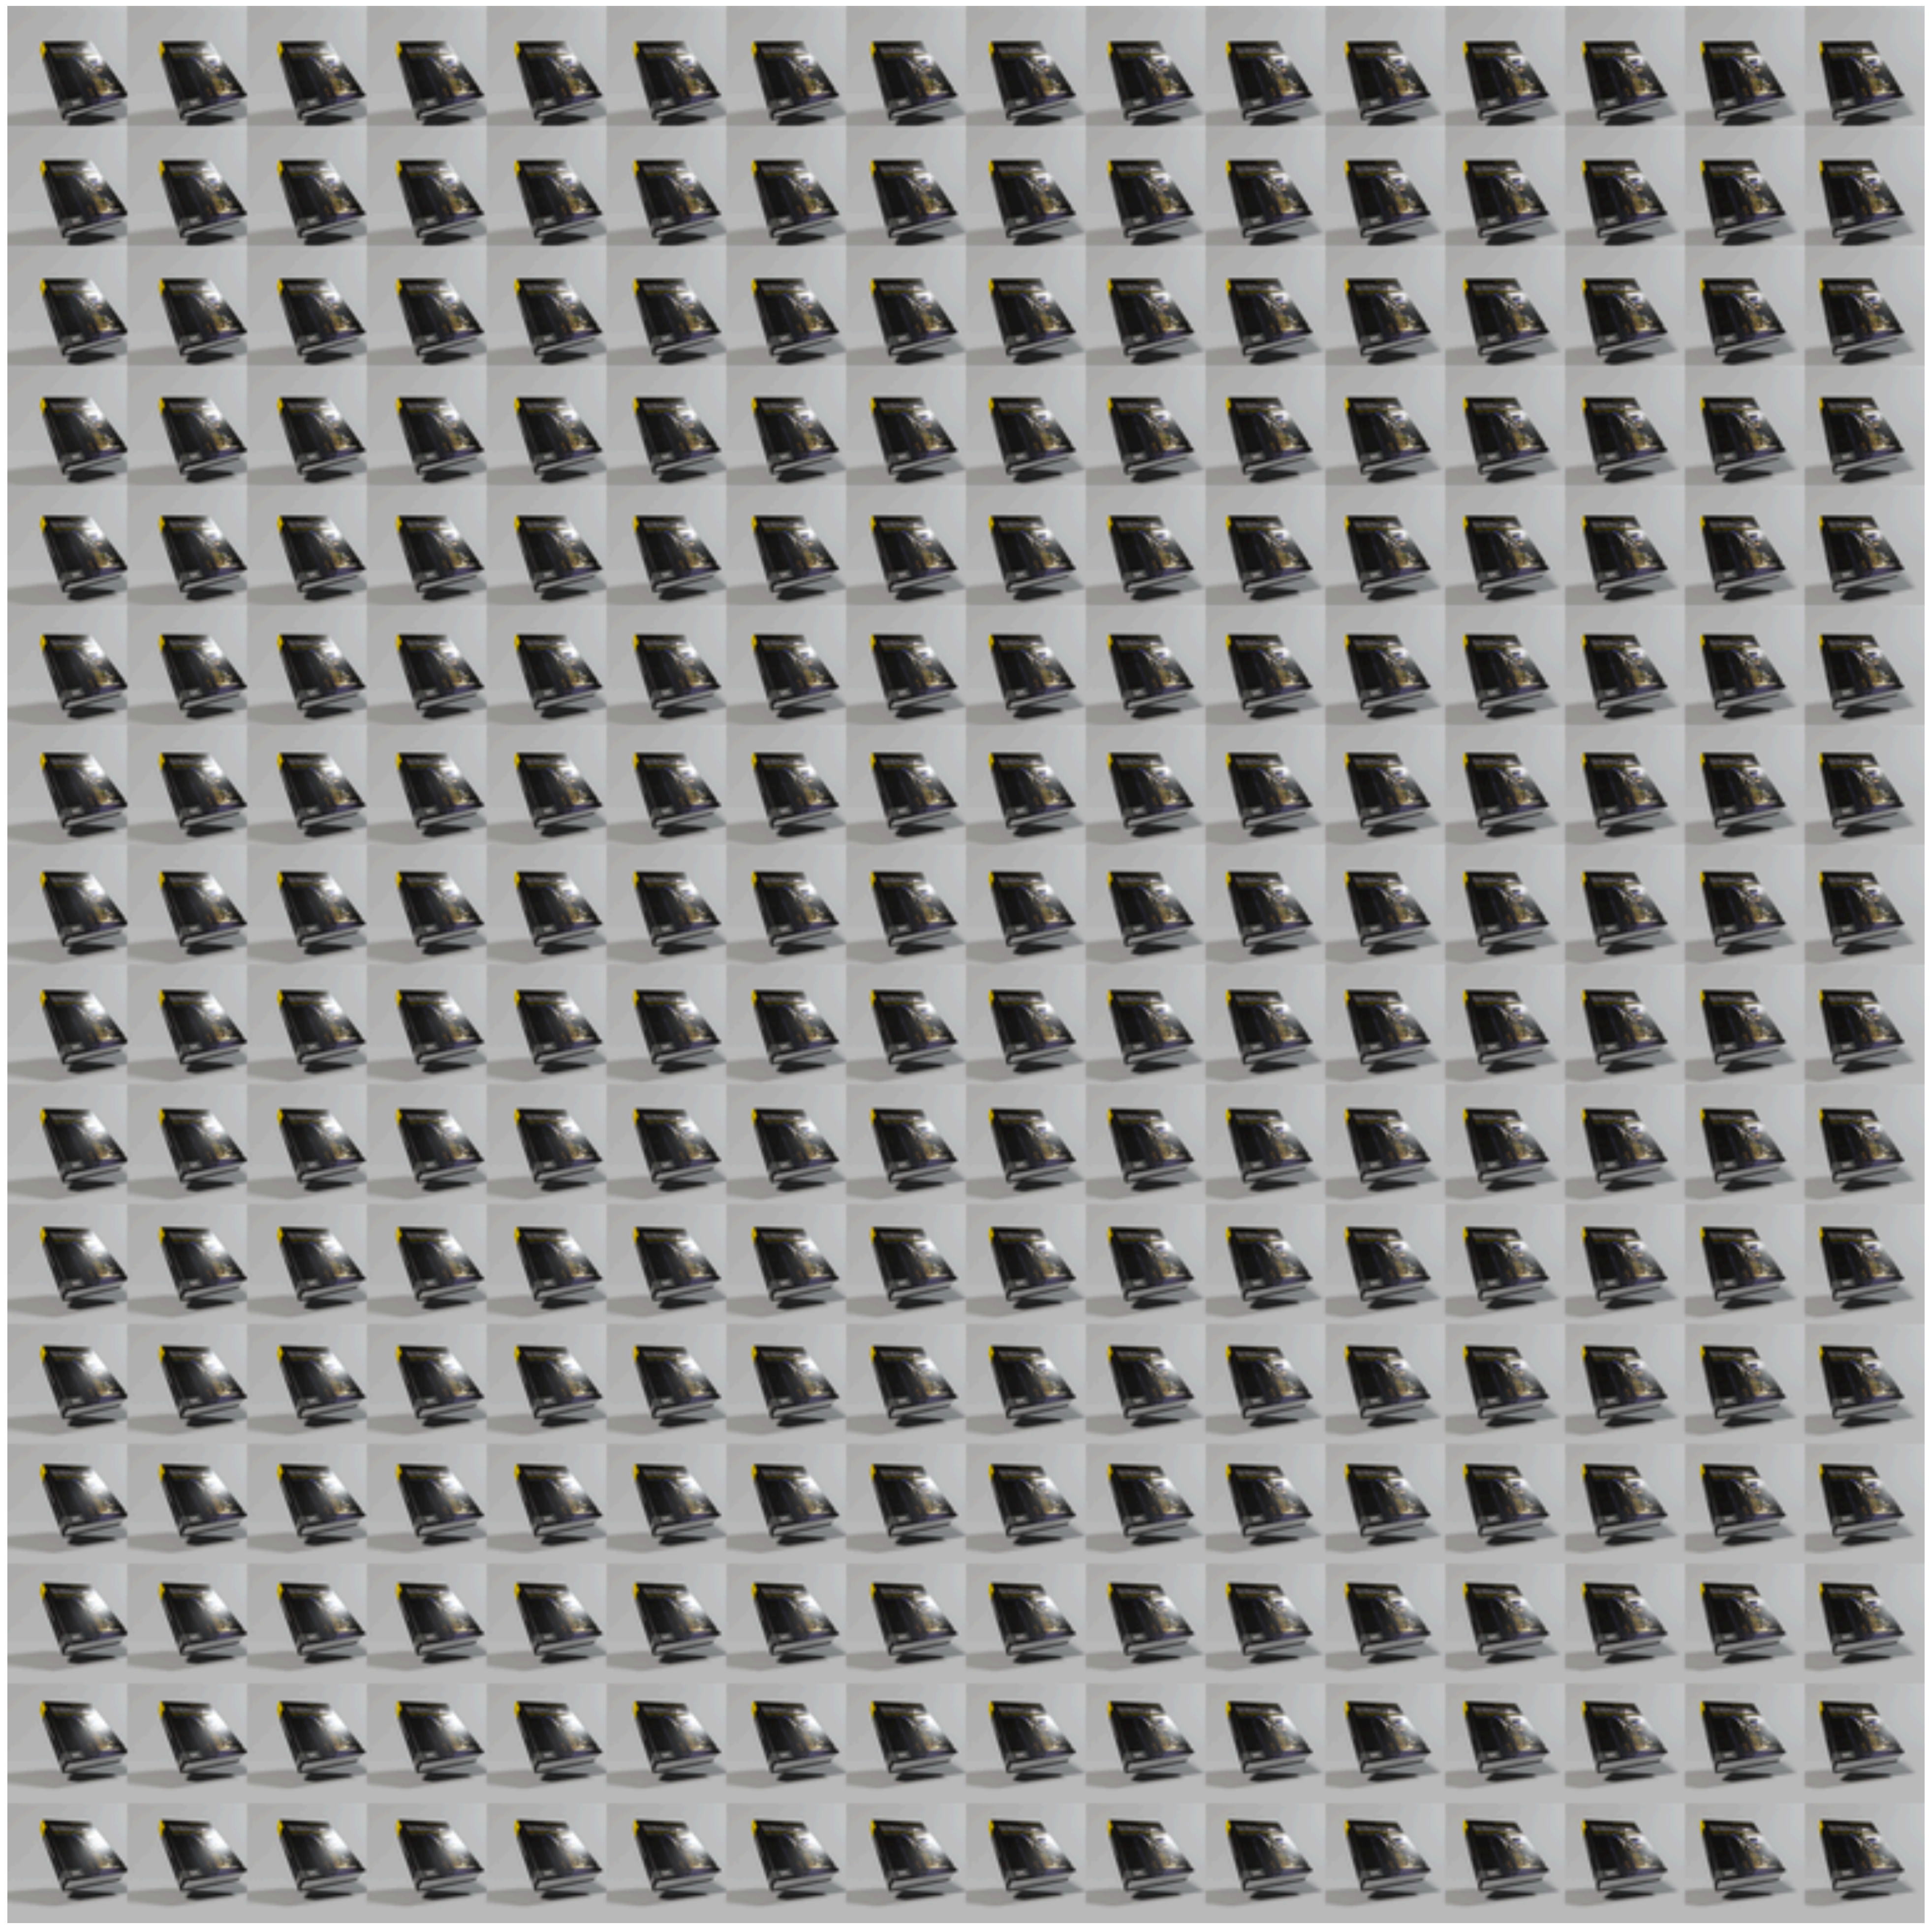

Assignment 3

Part 2. Compute image I(x,y) from another viewpoint using the light field image

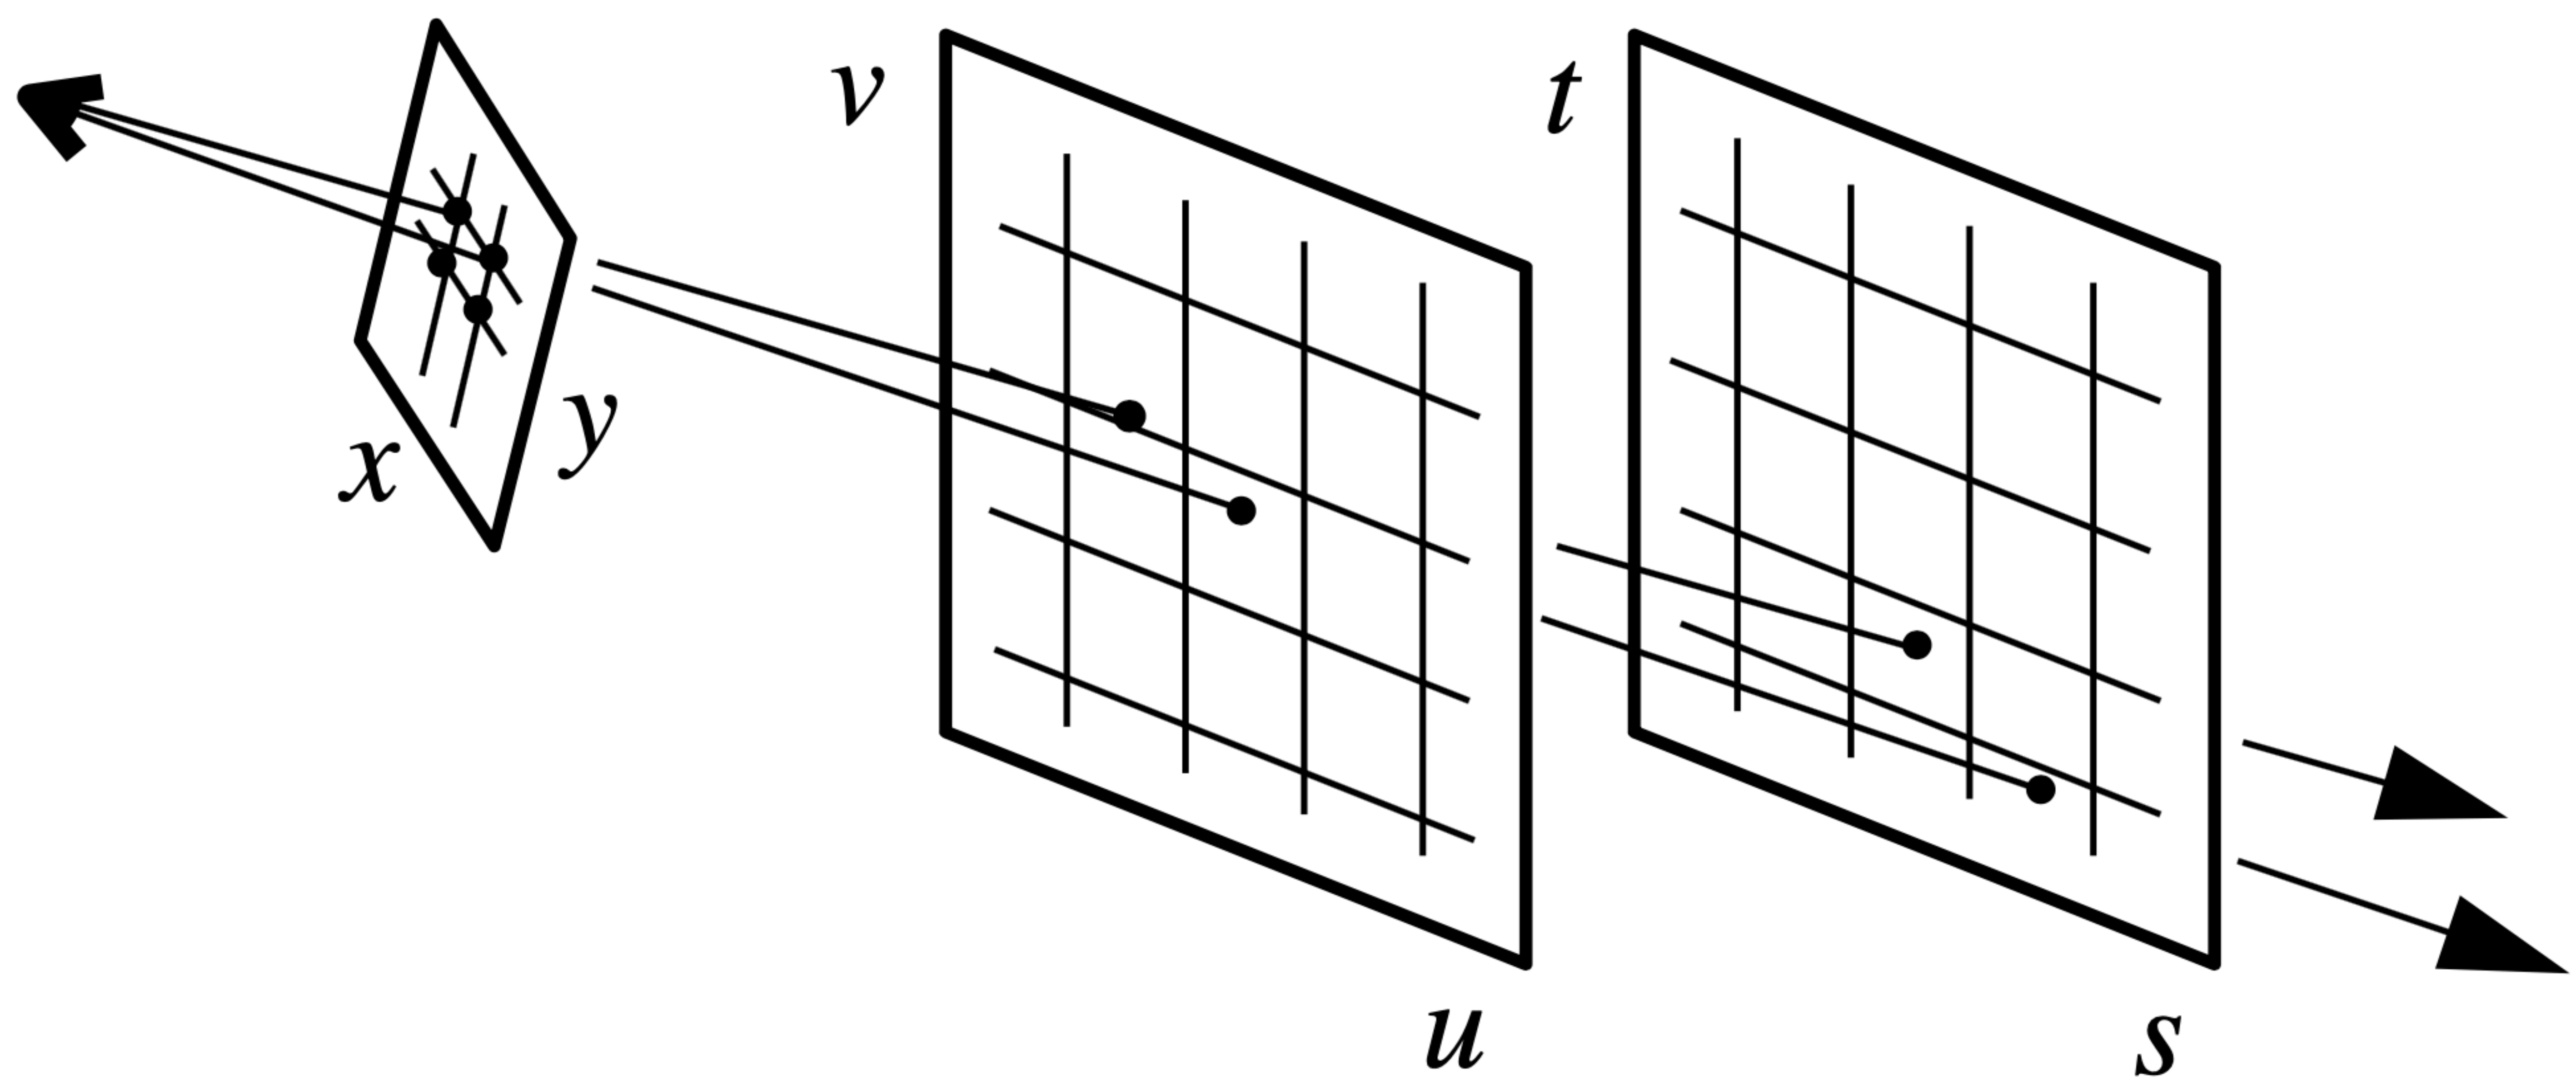

Acknowledgments

Many thanks to Marc Levoy, Pat Hanrahan, and Ren Ng