Pinholes and lenses

15-463, 15-663, 15-862
Computational Photography
Fall 2017, Lecture 3

http://graphics.cs.cmu.edu/courses/15-463

Course announcements

- Homework 1 is out.
  - Due September 14th.
  - Make sure to sign up for a camera if you need one.
  - Drop by Yannis' office to pick up cameras any time.

- Office hours for the semester:
  - Alankar: Tuesdays 5-7 pm, Smith 220.
  - Yannis: Fridays 5-7 pm, Smith 225.

- Those on the waitlist: Class is currently capped, but some seats will likely become available after the drop deadline.

Overview of today's lecture

- Leftover from lecture 2: thoughts on the image processing pipeline.
- Some motivational imaging experiments.
- Pinhole camera.
- Accidental pinholes.
- The thin lens model.
- Lens camera and pinhole camera.
- Perspective.
- Field of view.
- Orthographic camera and telecentric lenses.

Slide credits

Many of these slides were adapted from:

- Kris Kitani (15-463, Fall 2016).
- Fredo Durand (MIT).
- Gordon Wetzstein (Stanford).

The modern photography pipeline

post-capture processing
(lectures 5-10)

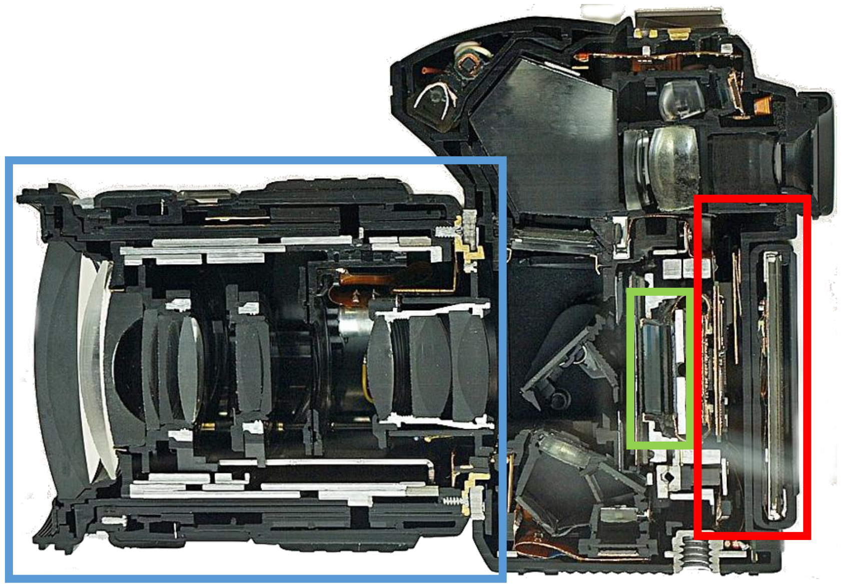

optics and optical controls
(lectures 2-3, 11-20)

sensor, analog front-end, and color filter array
(lectures 2, 23)

in-camera image processing pipeline
(lecture 2)

Some motivational imaging experiments

Let’s say we have a sensor…

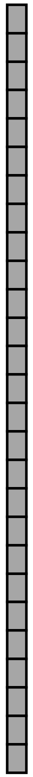

digital sensor
(CCD or CMOS)

… and an object we like to photograph

real-world object

digital sensor (CCD or CMOS)

What would an image taken like this look like?

Bare-sensor imaging

Bare-sensor imaging

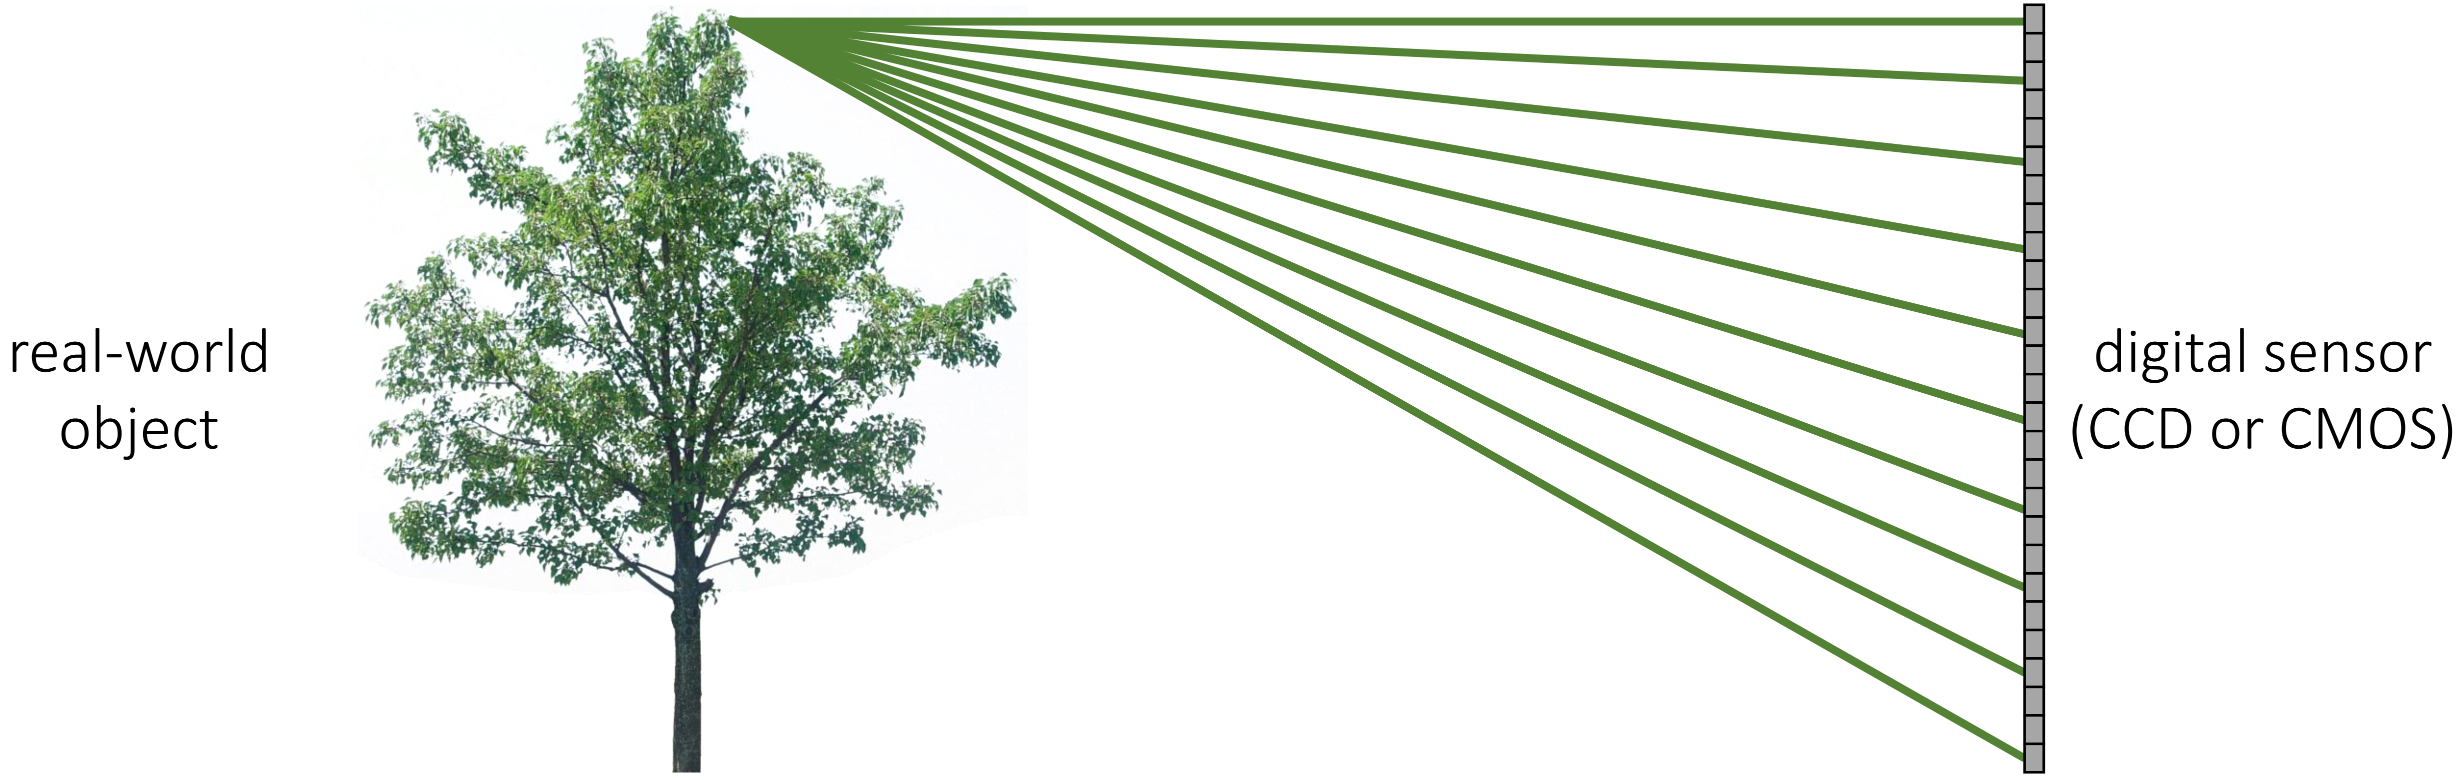

Bare-sensor imaging

real-world object

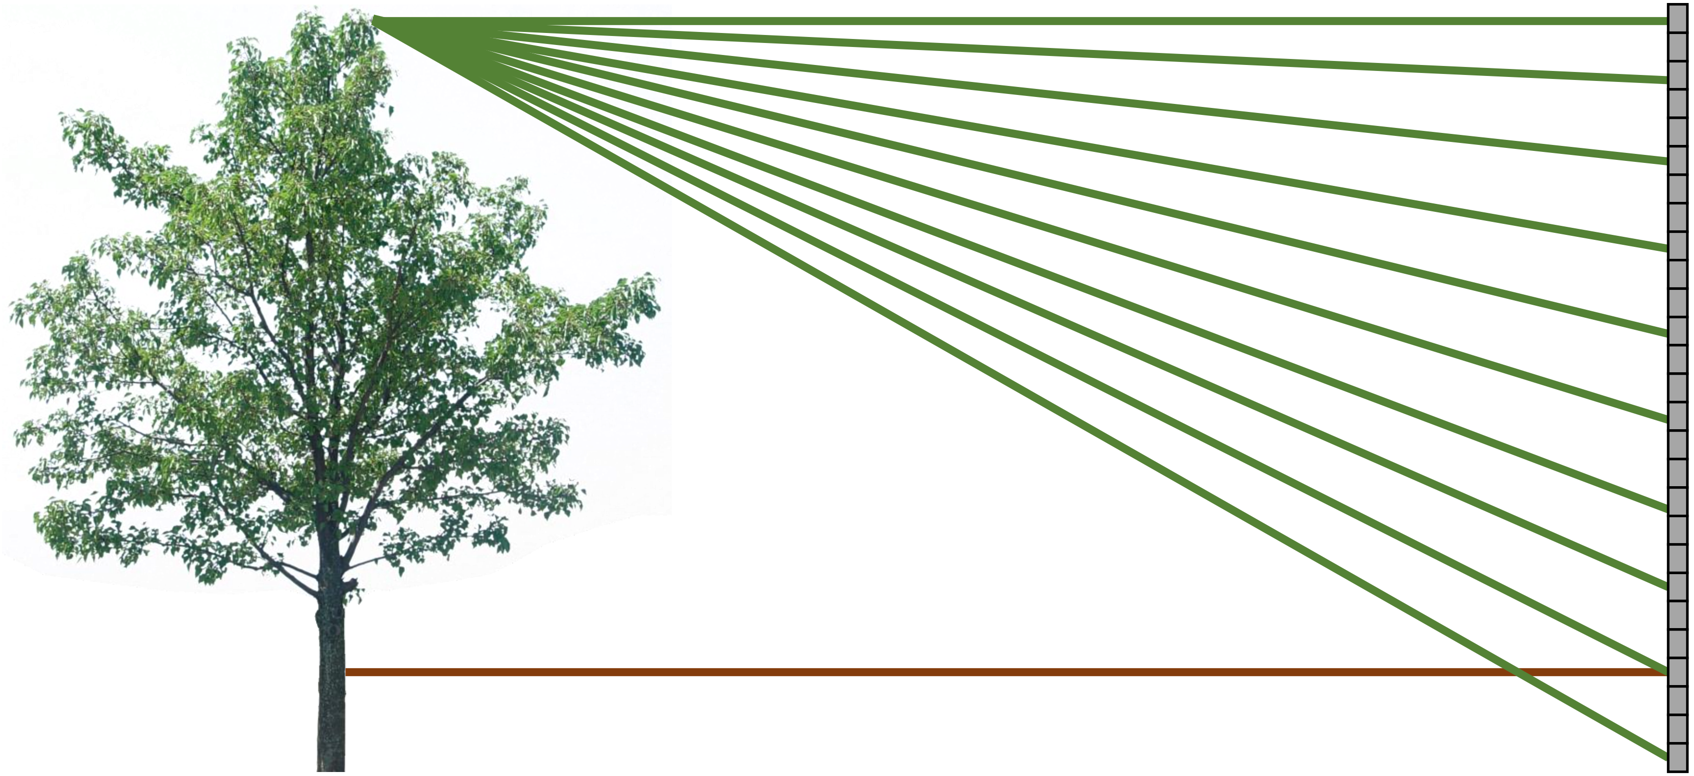

digital sensor (CCD or CMOS)

Bare-sensor imaging

real-world object

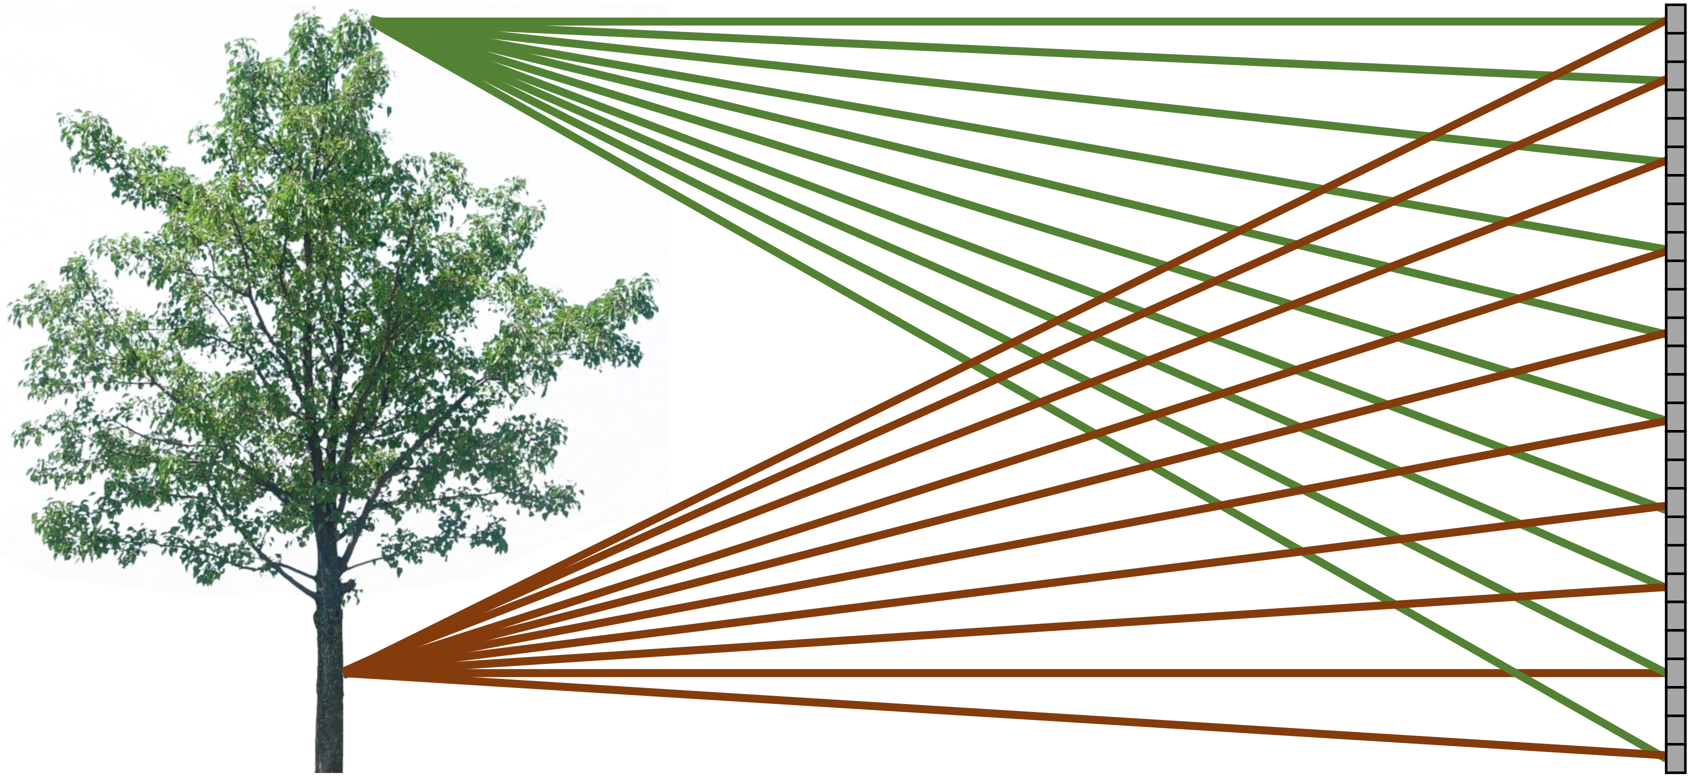

digital sensor (CCD or CMOS)

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

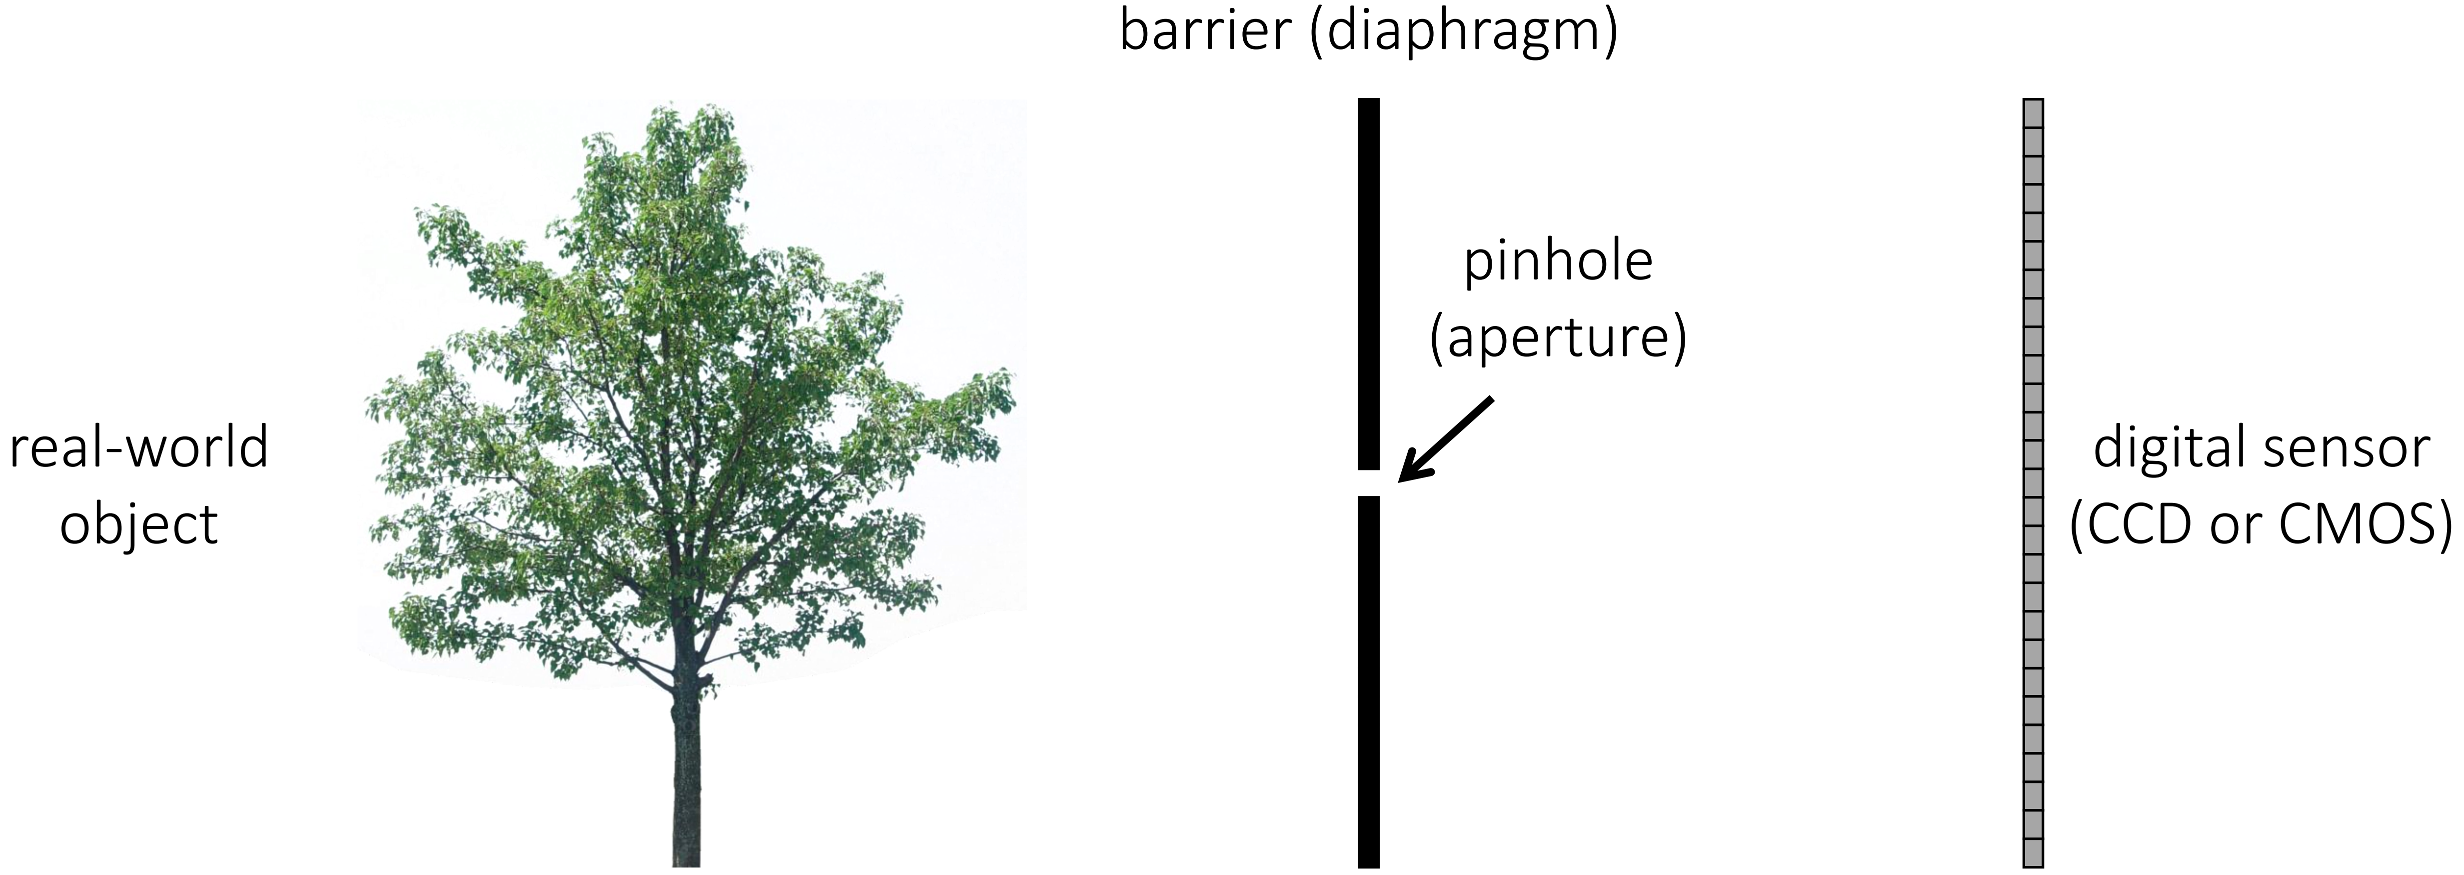

What would an image taken like this look like?

Pinhole imaging

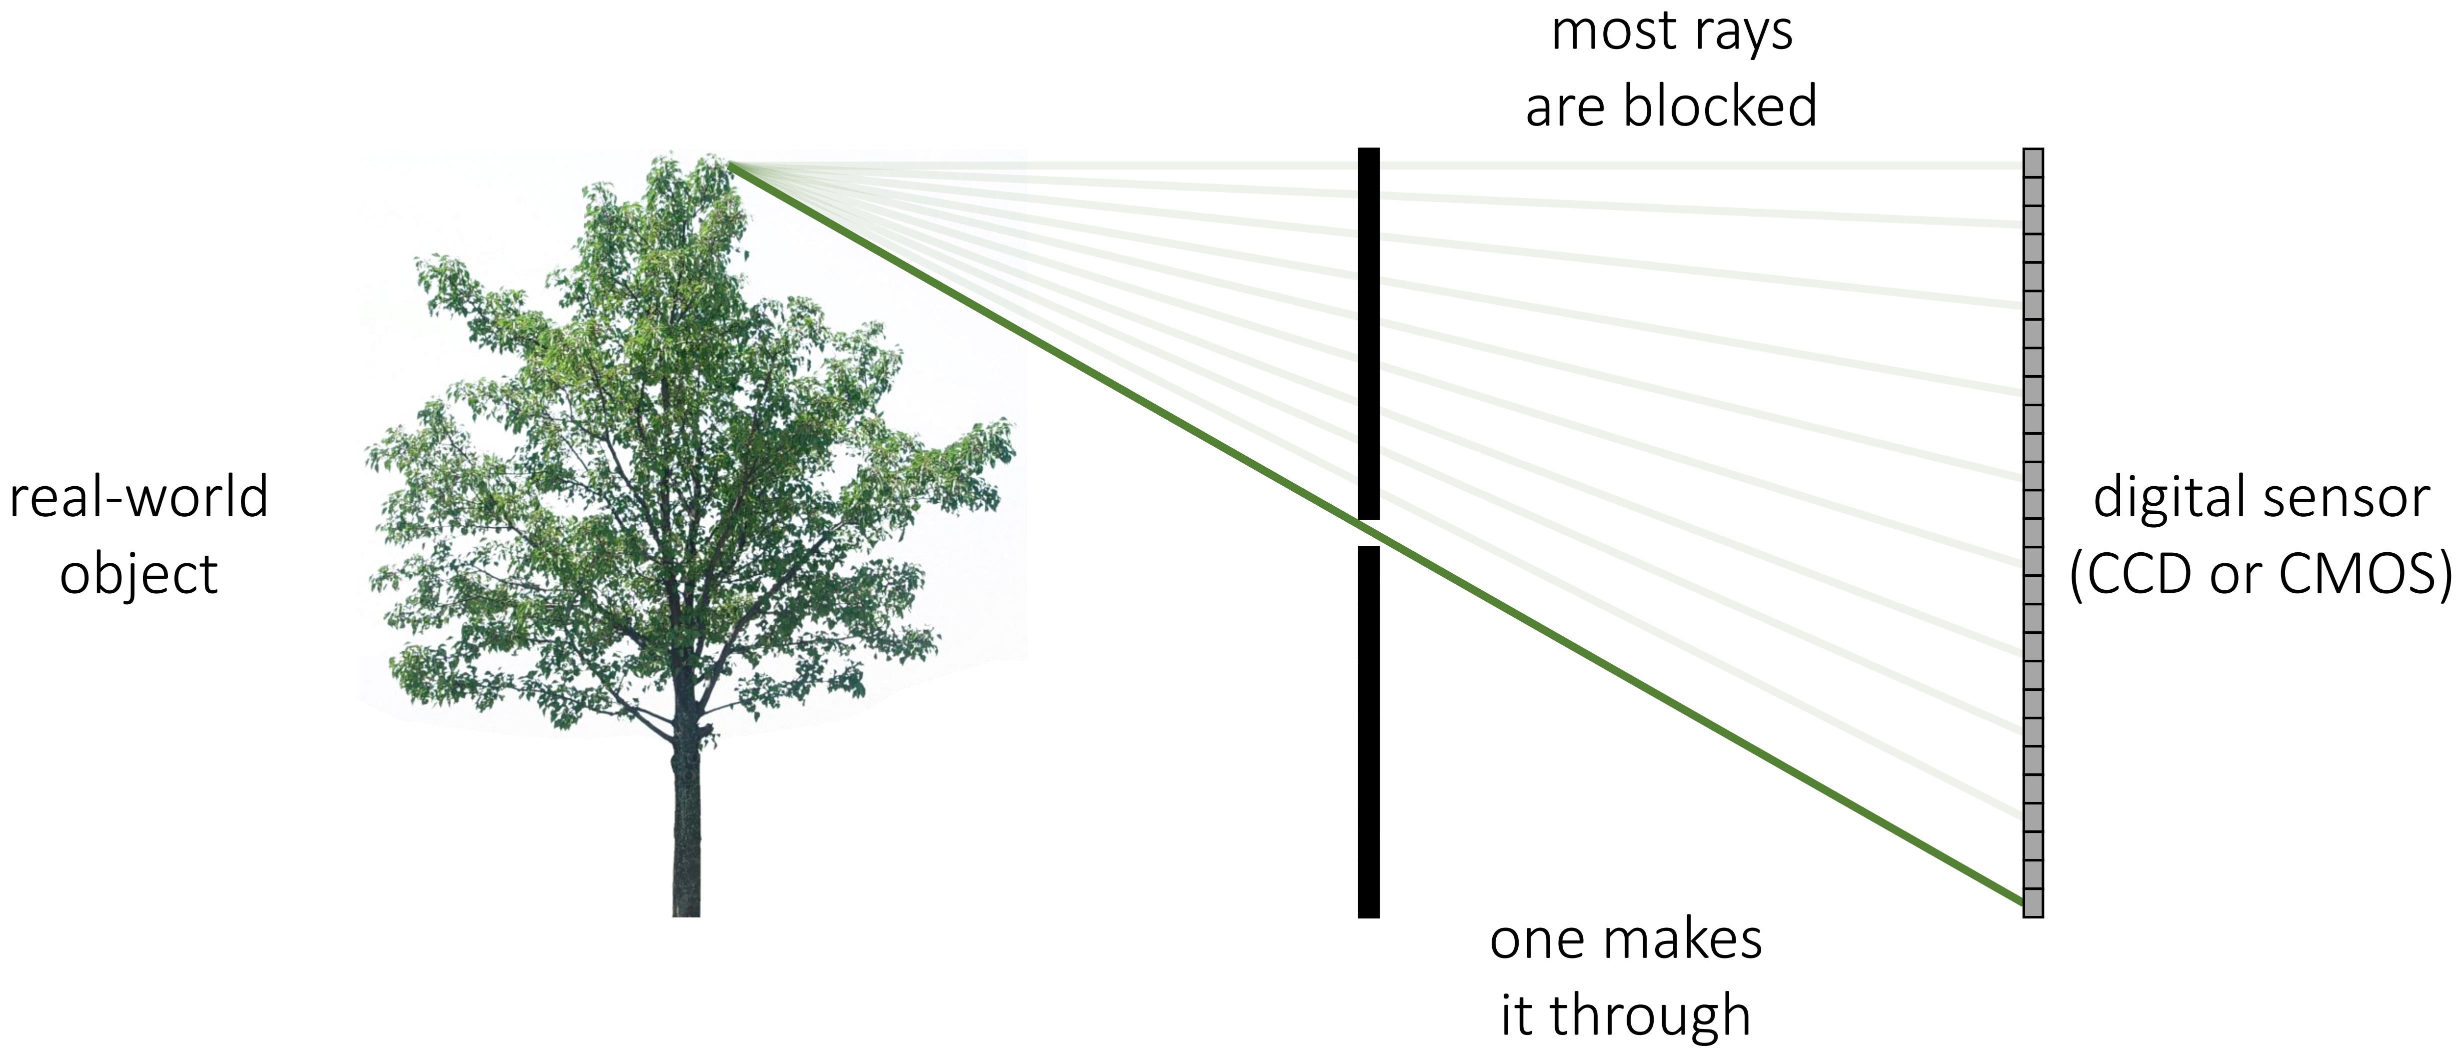

Pinhole imaging

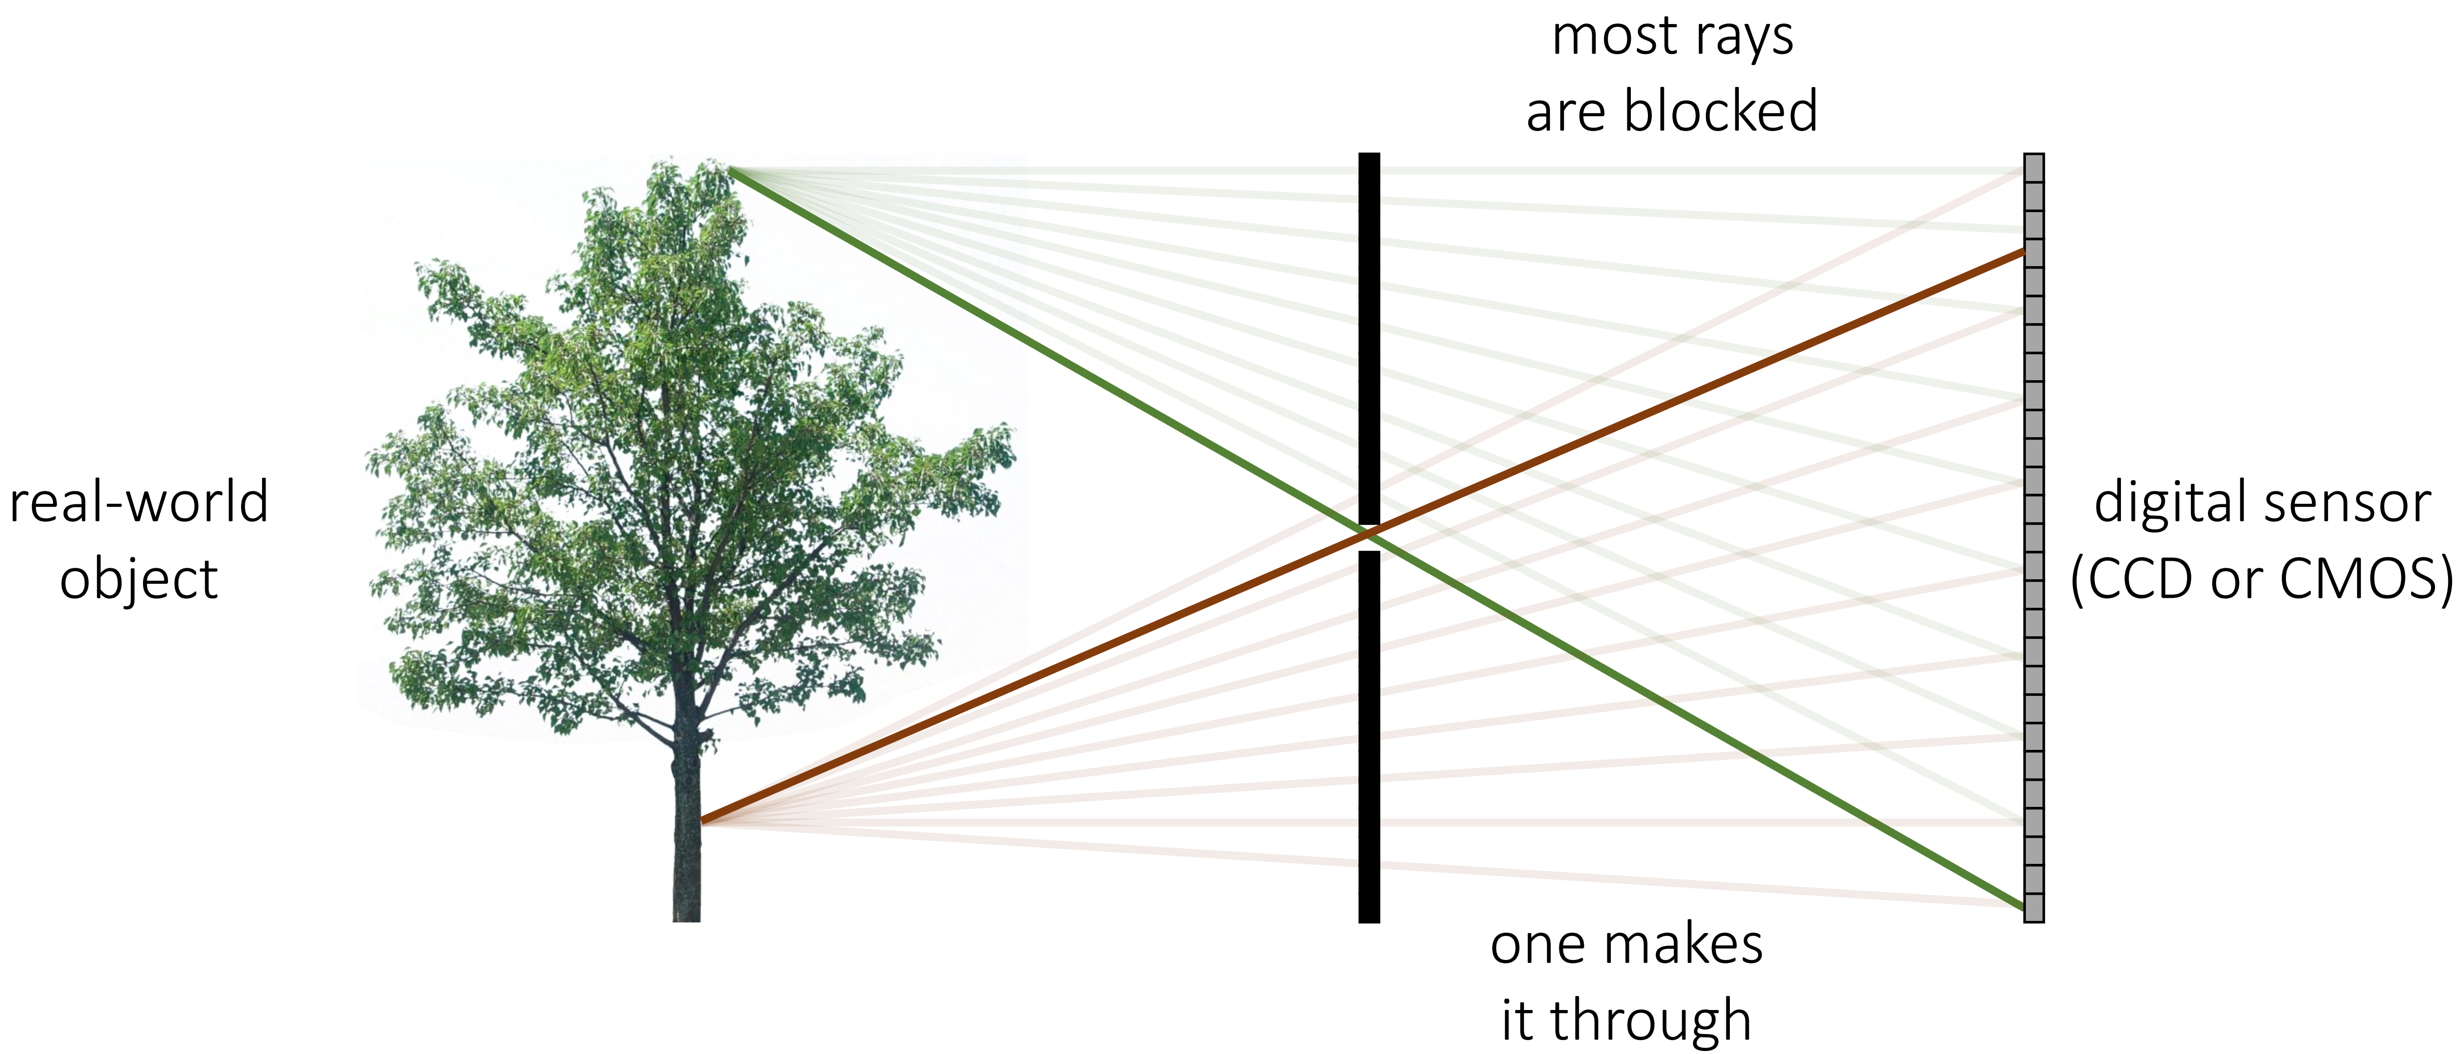

Pinhole imaging

real-world object

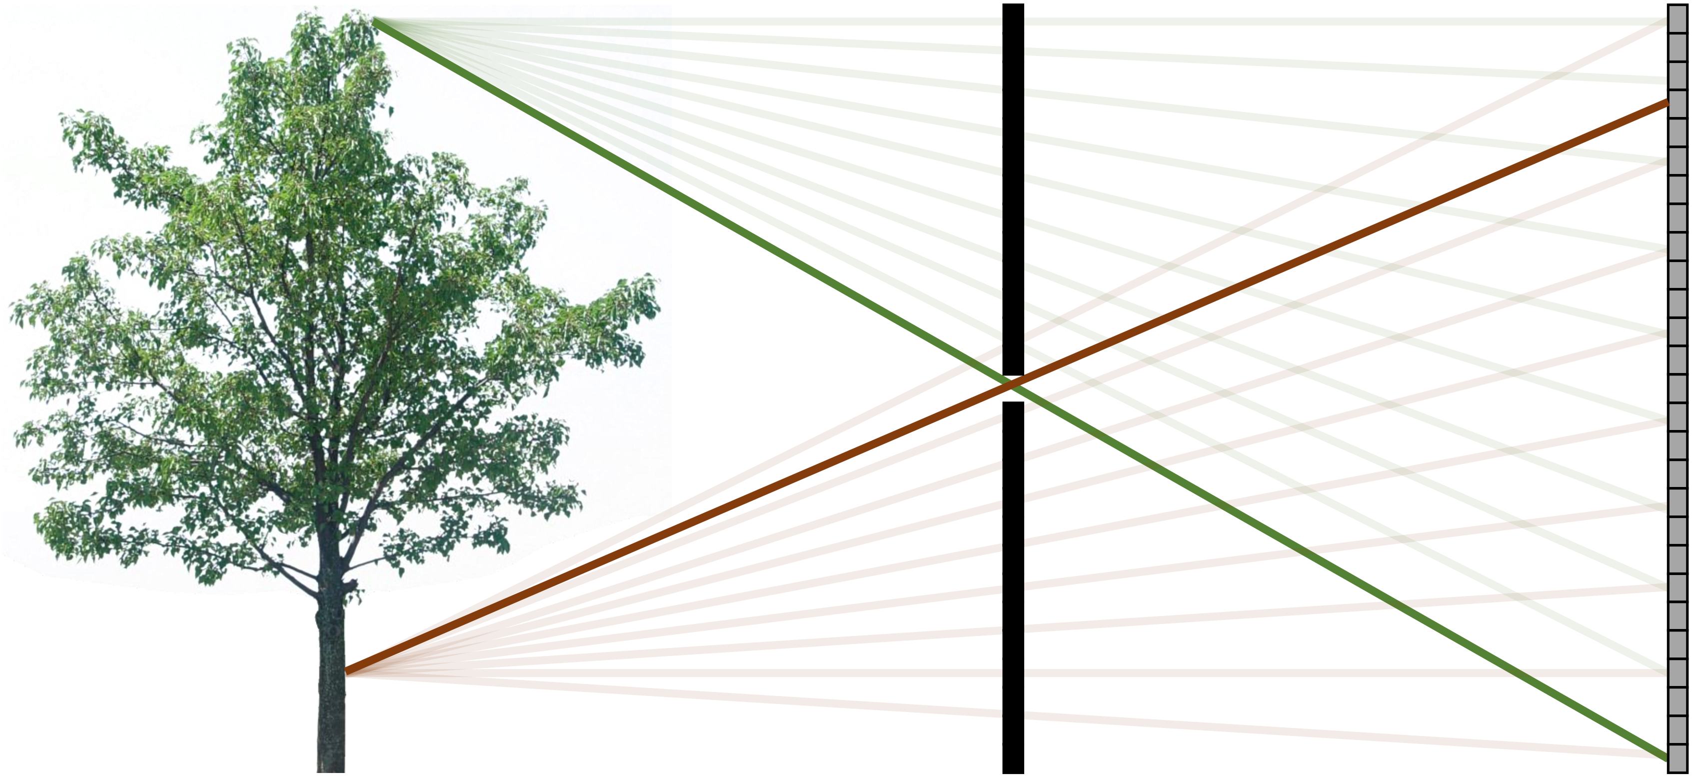

digital sensor (CCD or CMOS)

Each scene point contributes to only one sensor pixel

What does the image on the sensor look like?

Pinhole imaging

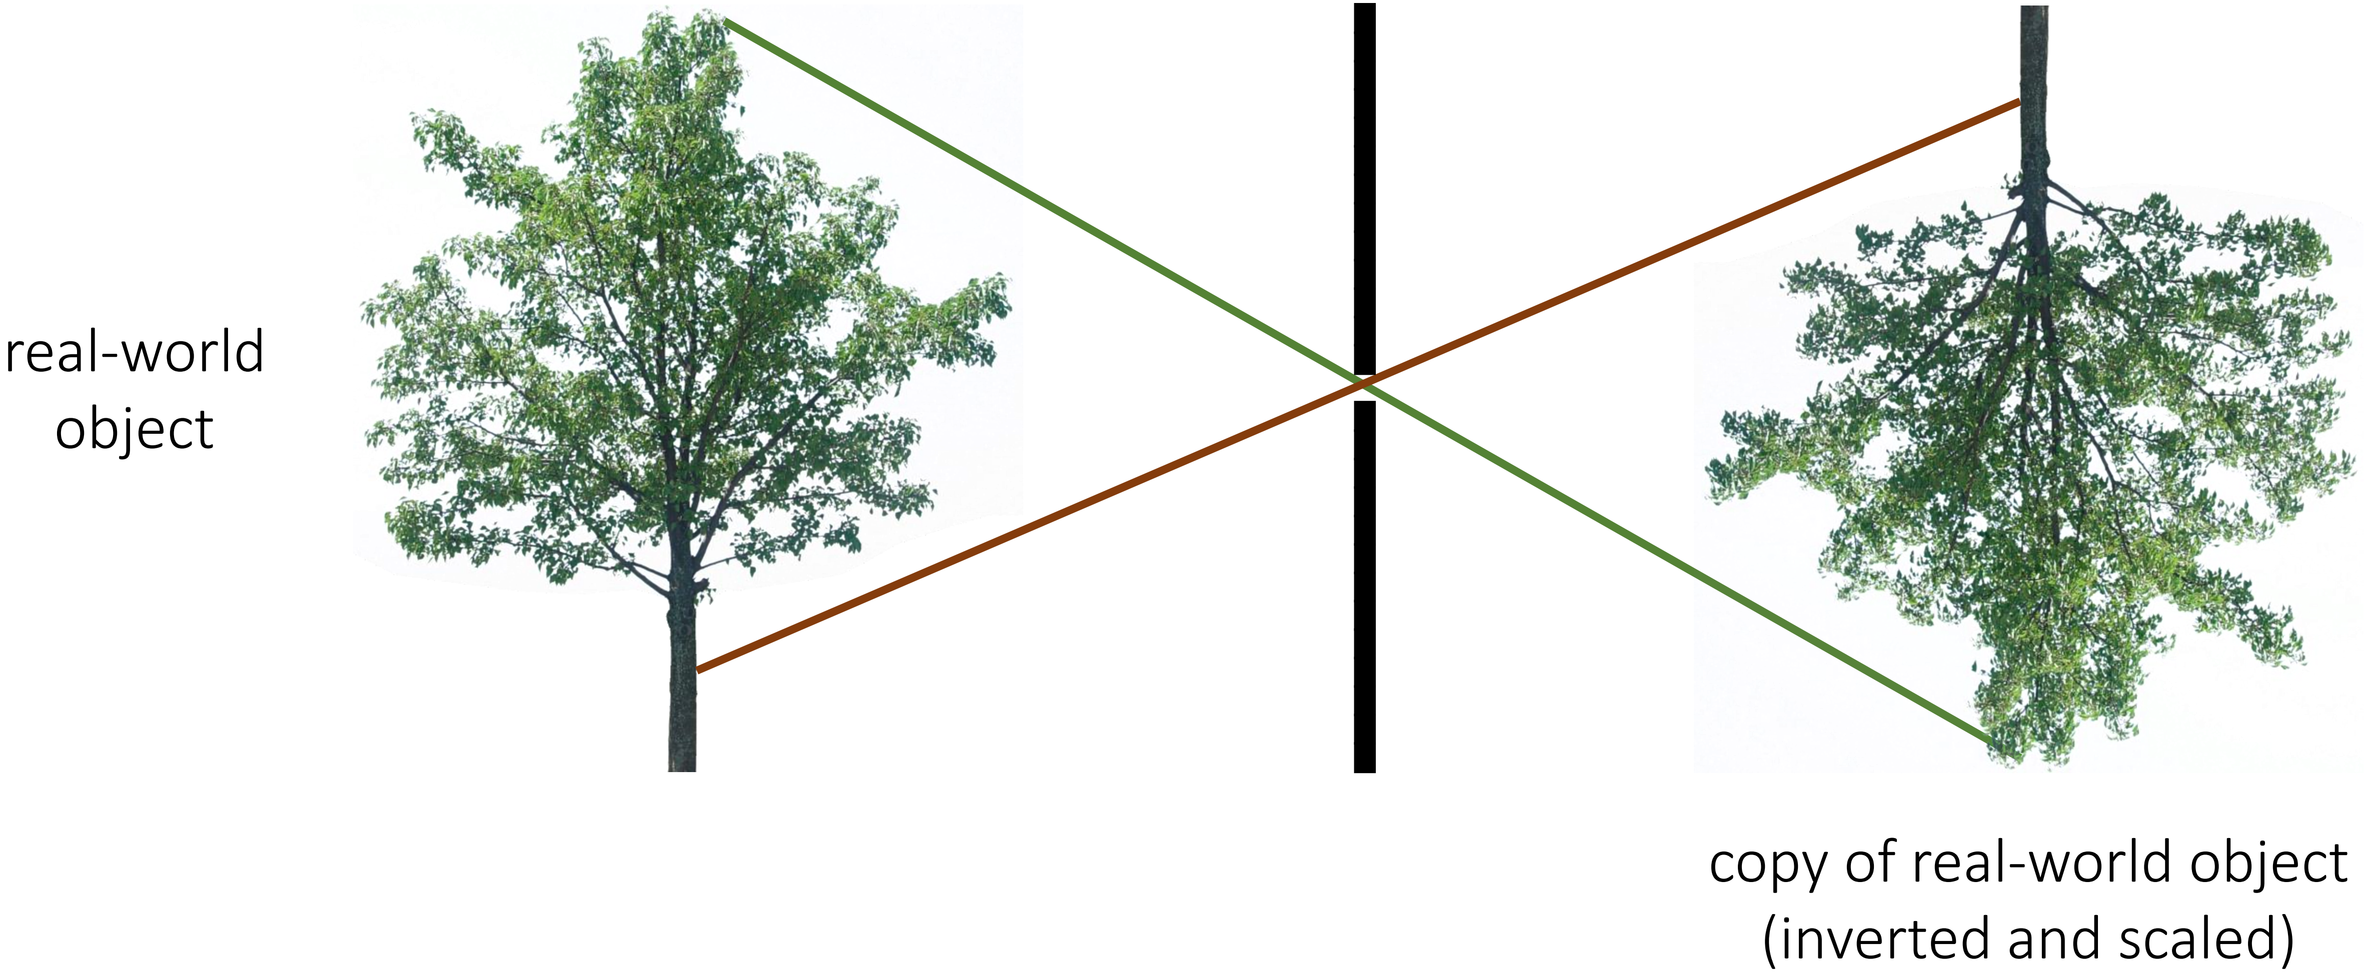

Pinhole camera

Pinhole camera a.k.a. camera obscura

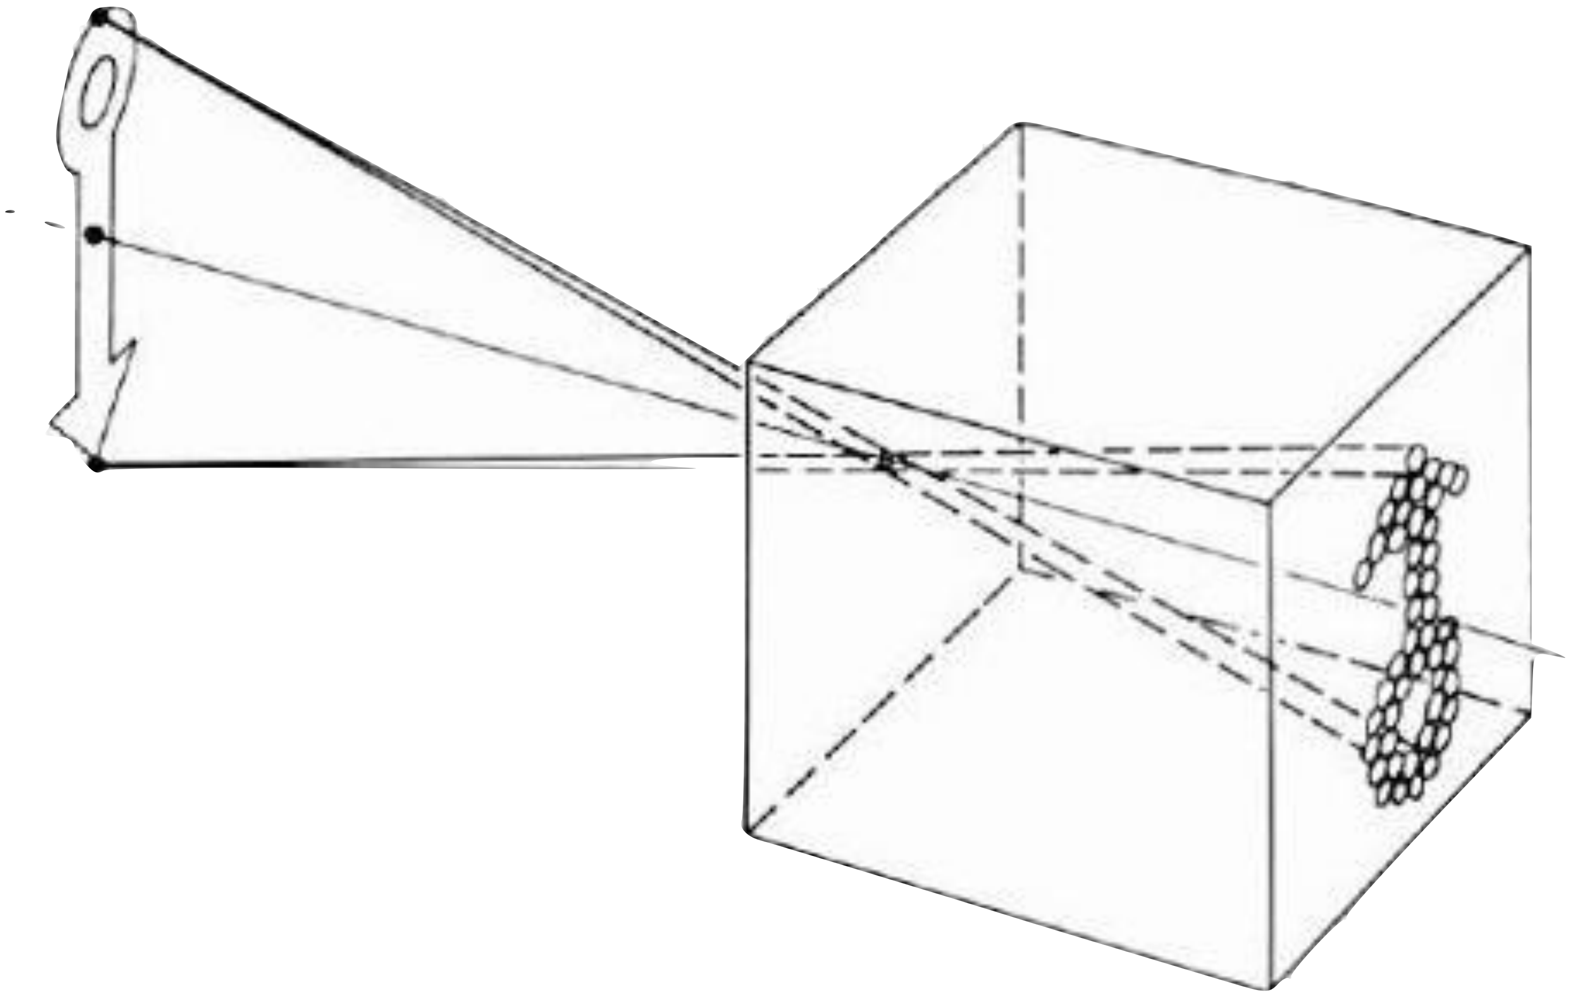

Pinhole camera a.k.a. camera obscura

First mention ...

Chinese philosopher Mozi
(470 to 390 BC)

First camera ...

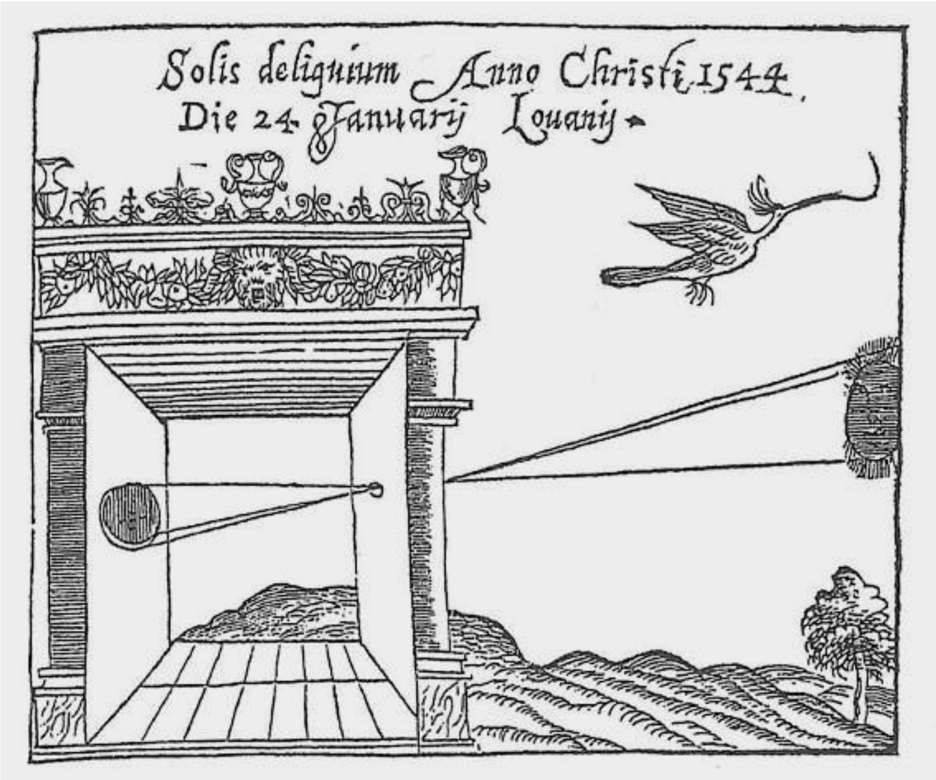

Greek philosopher Aristotle
(384 to 322 BC)

Pinhole camera terms

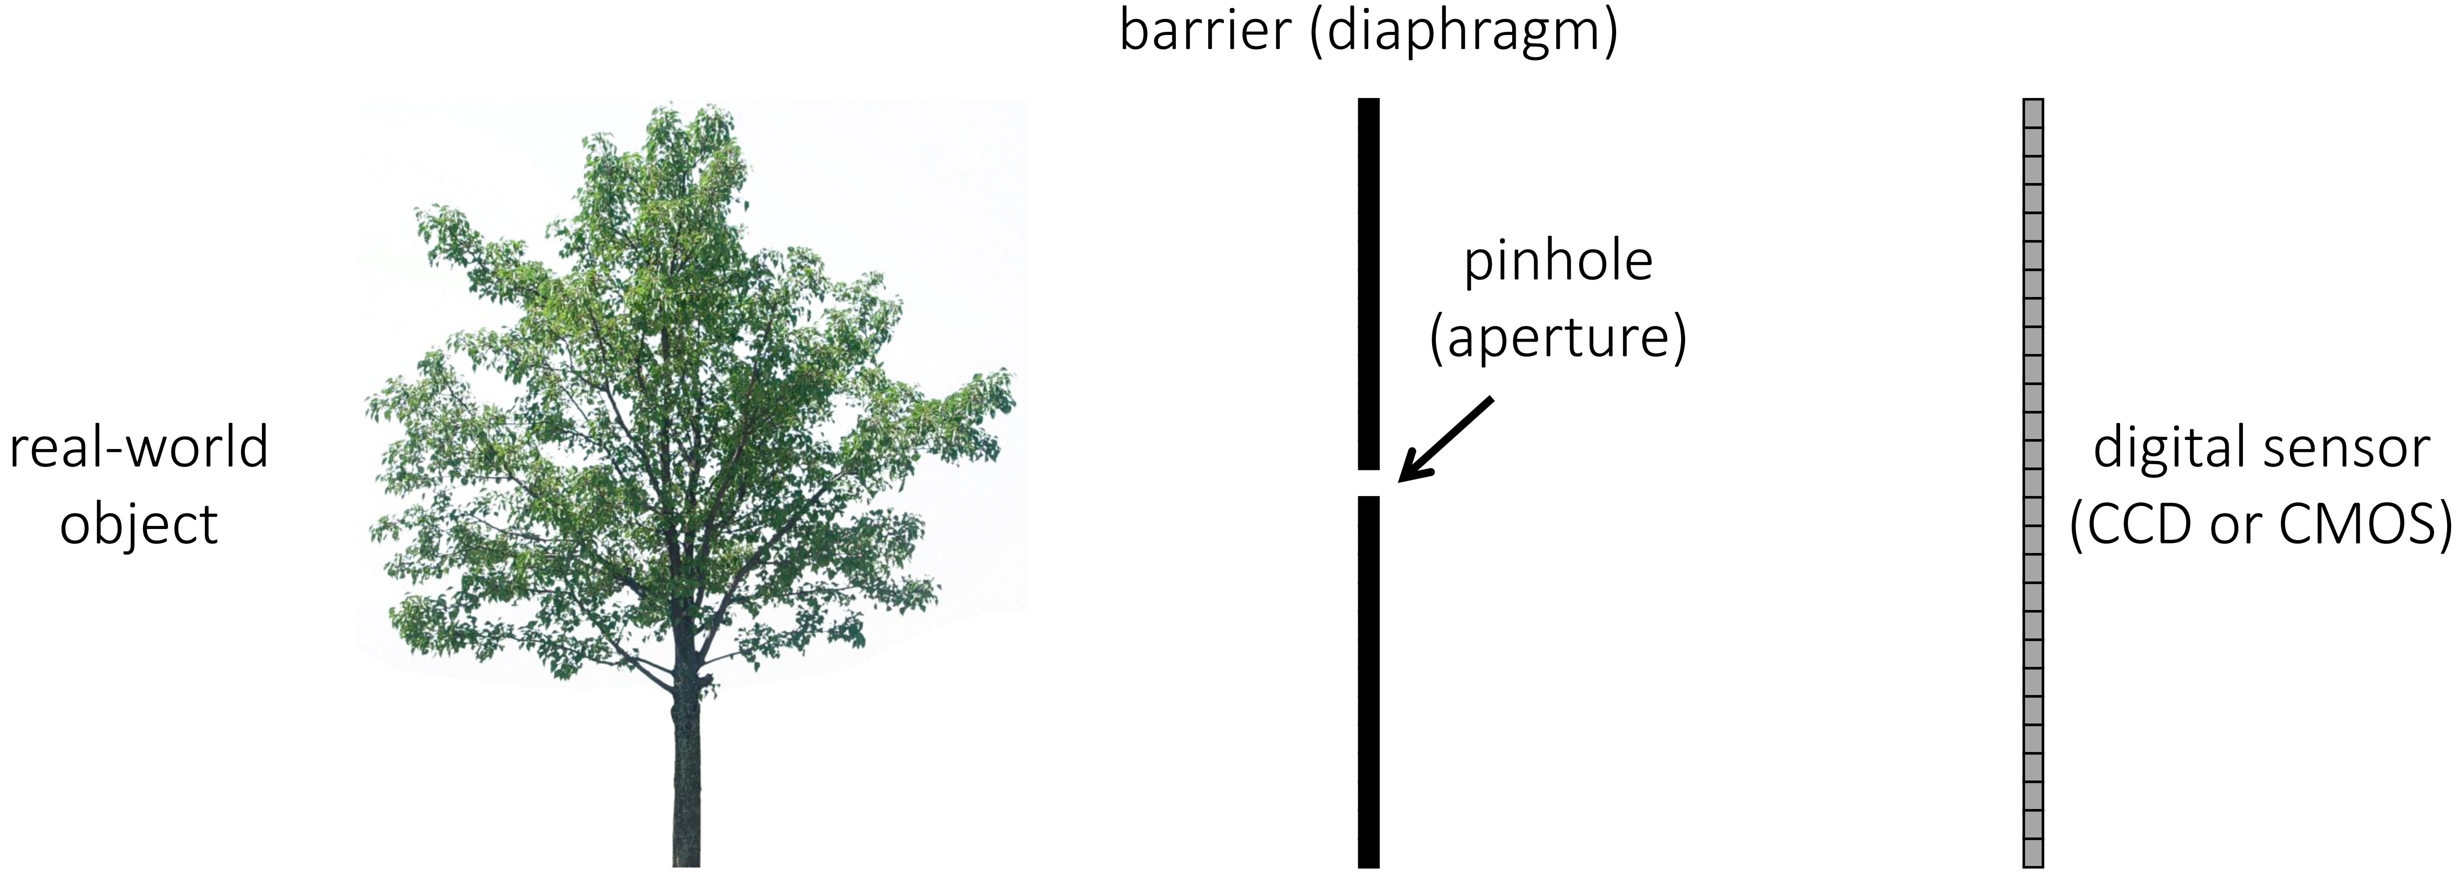

Pinhole camera terms

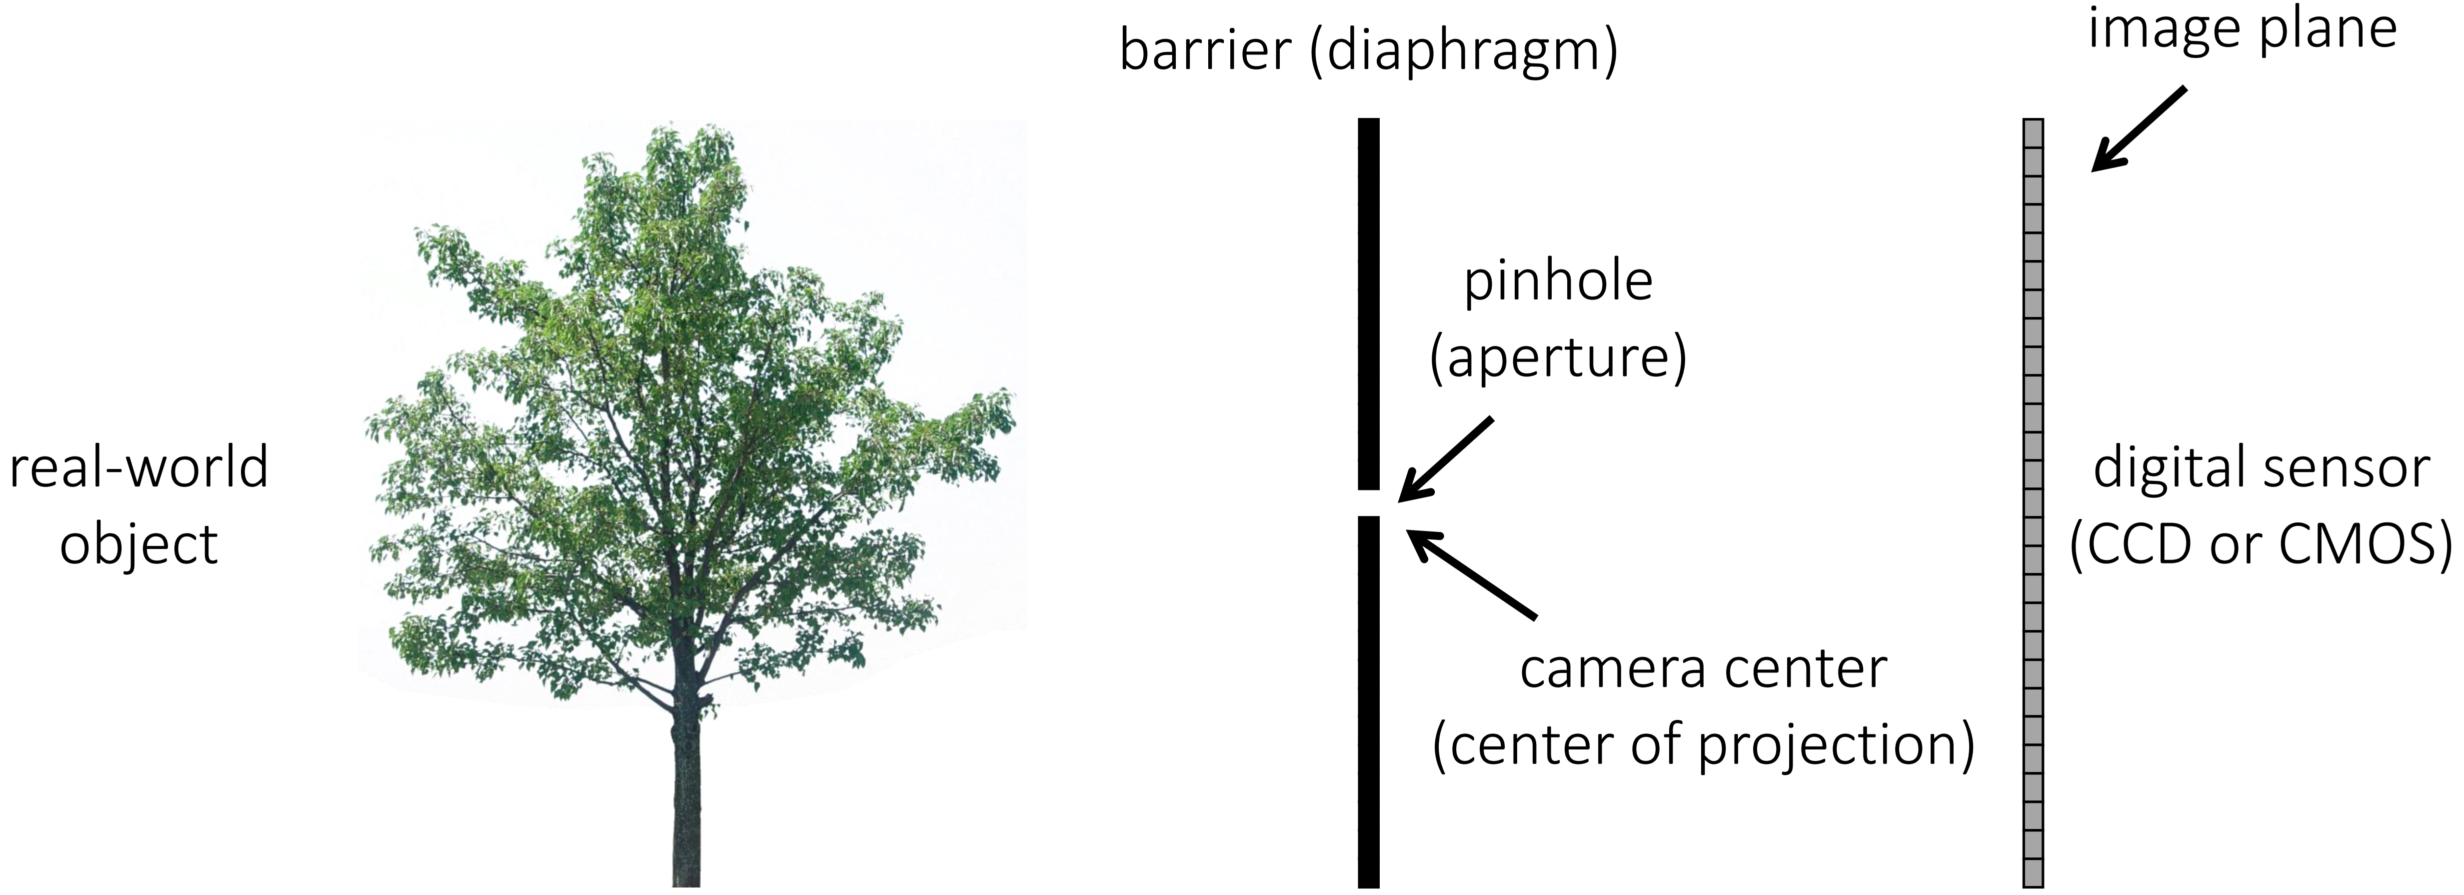

Focal length

real-world
object

Focal length

What happens as we change the focal length?

real-world object

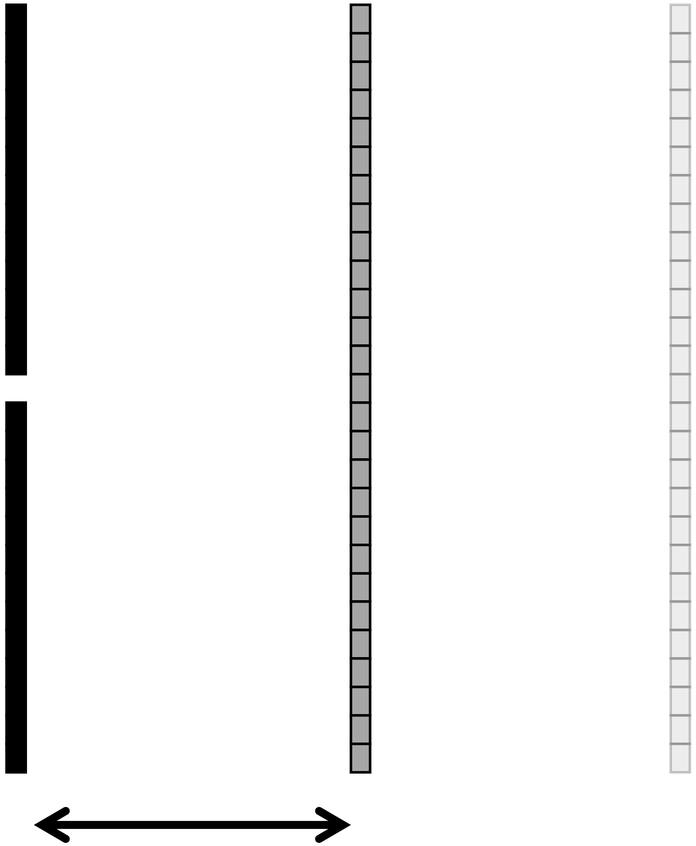

focal length 0.5 f

Focal length

What happens as we change the focal length?

real-world object

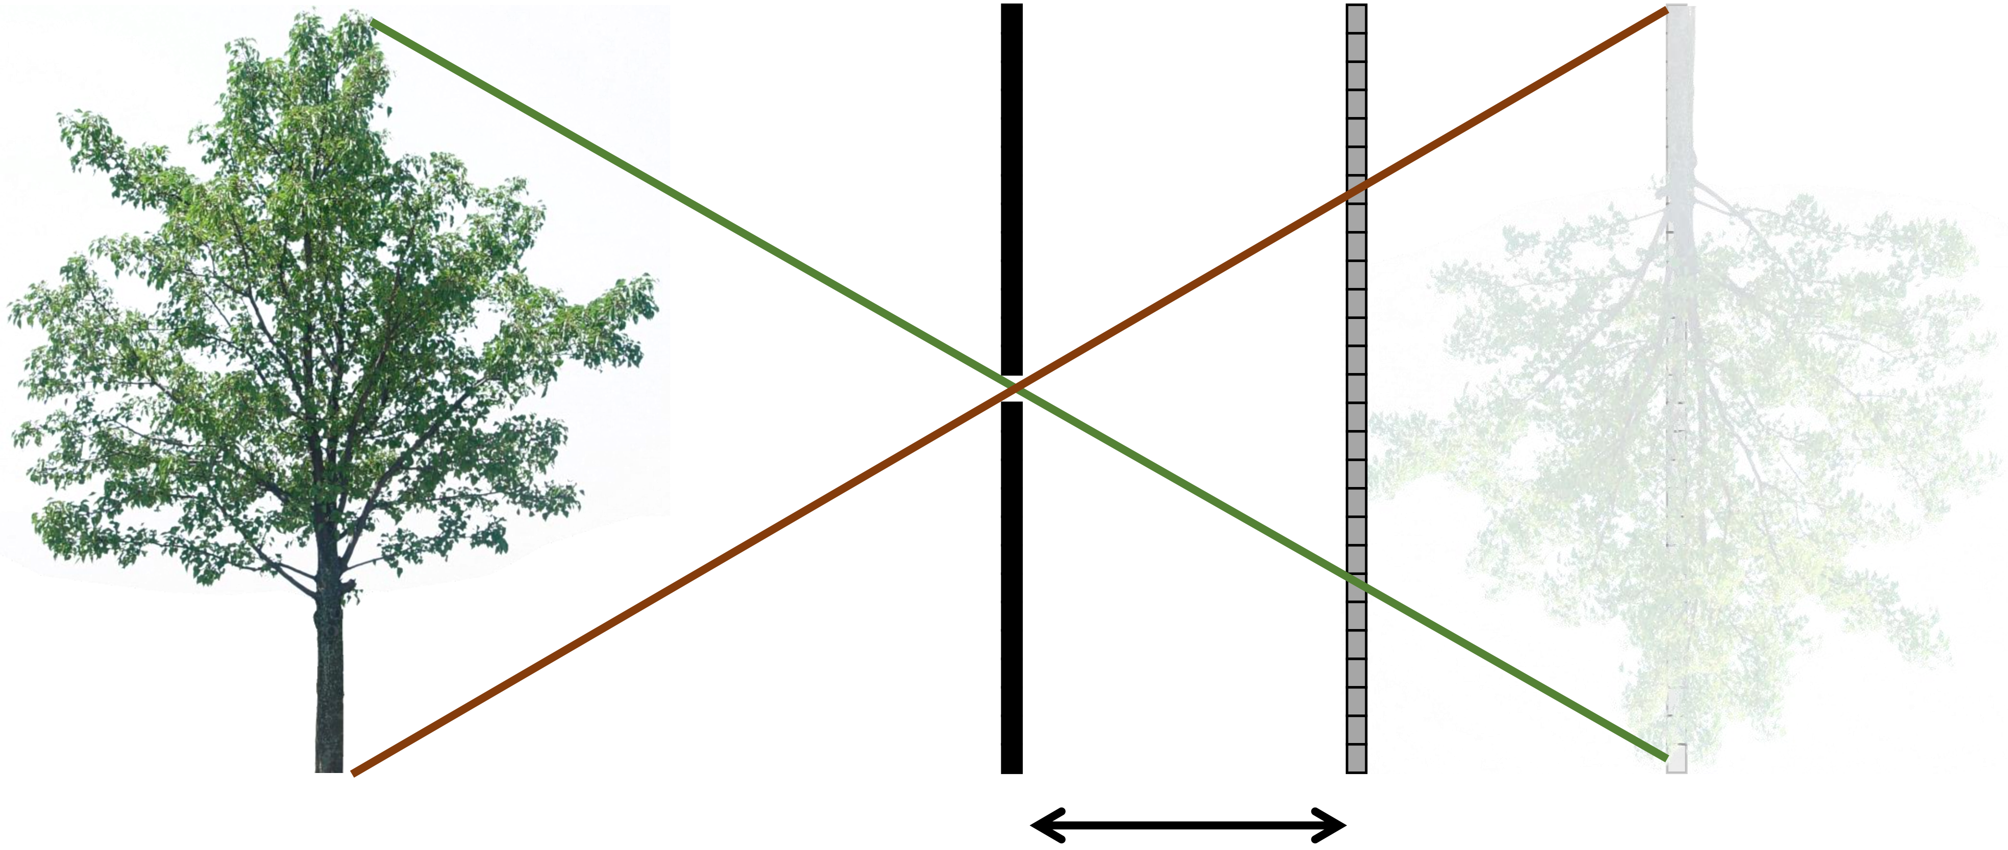

focal length 0.5 f

Focal length

What happens as we change the focal length?

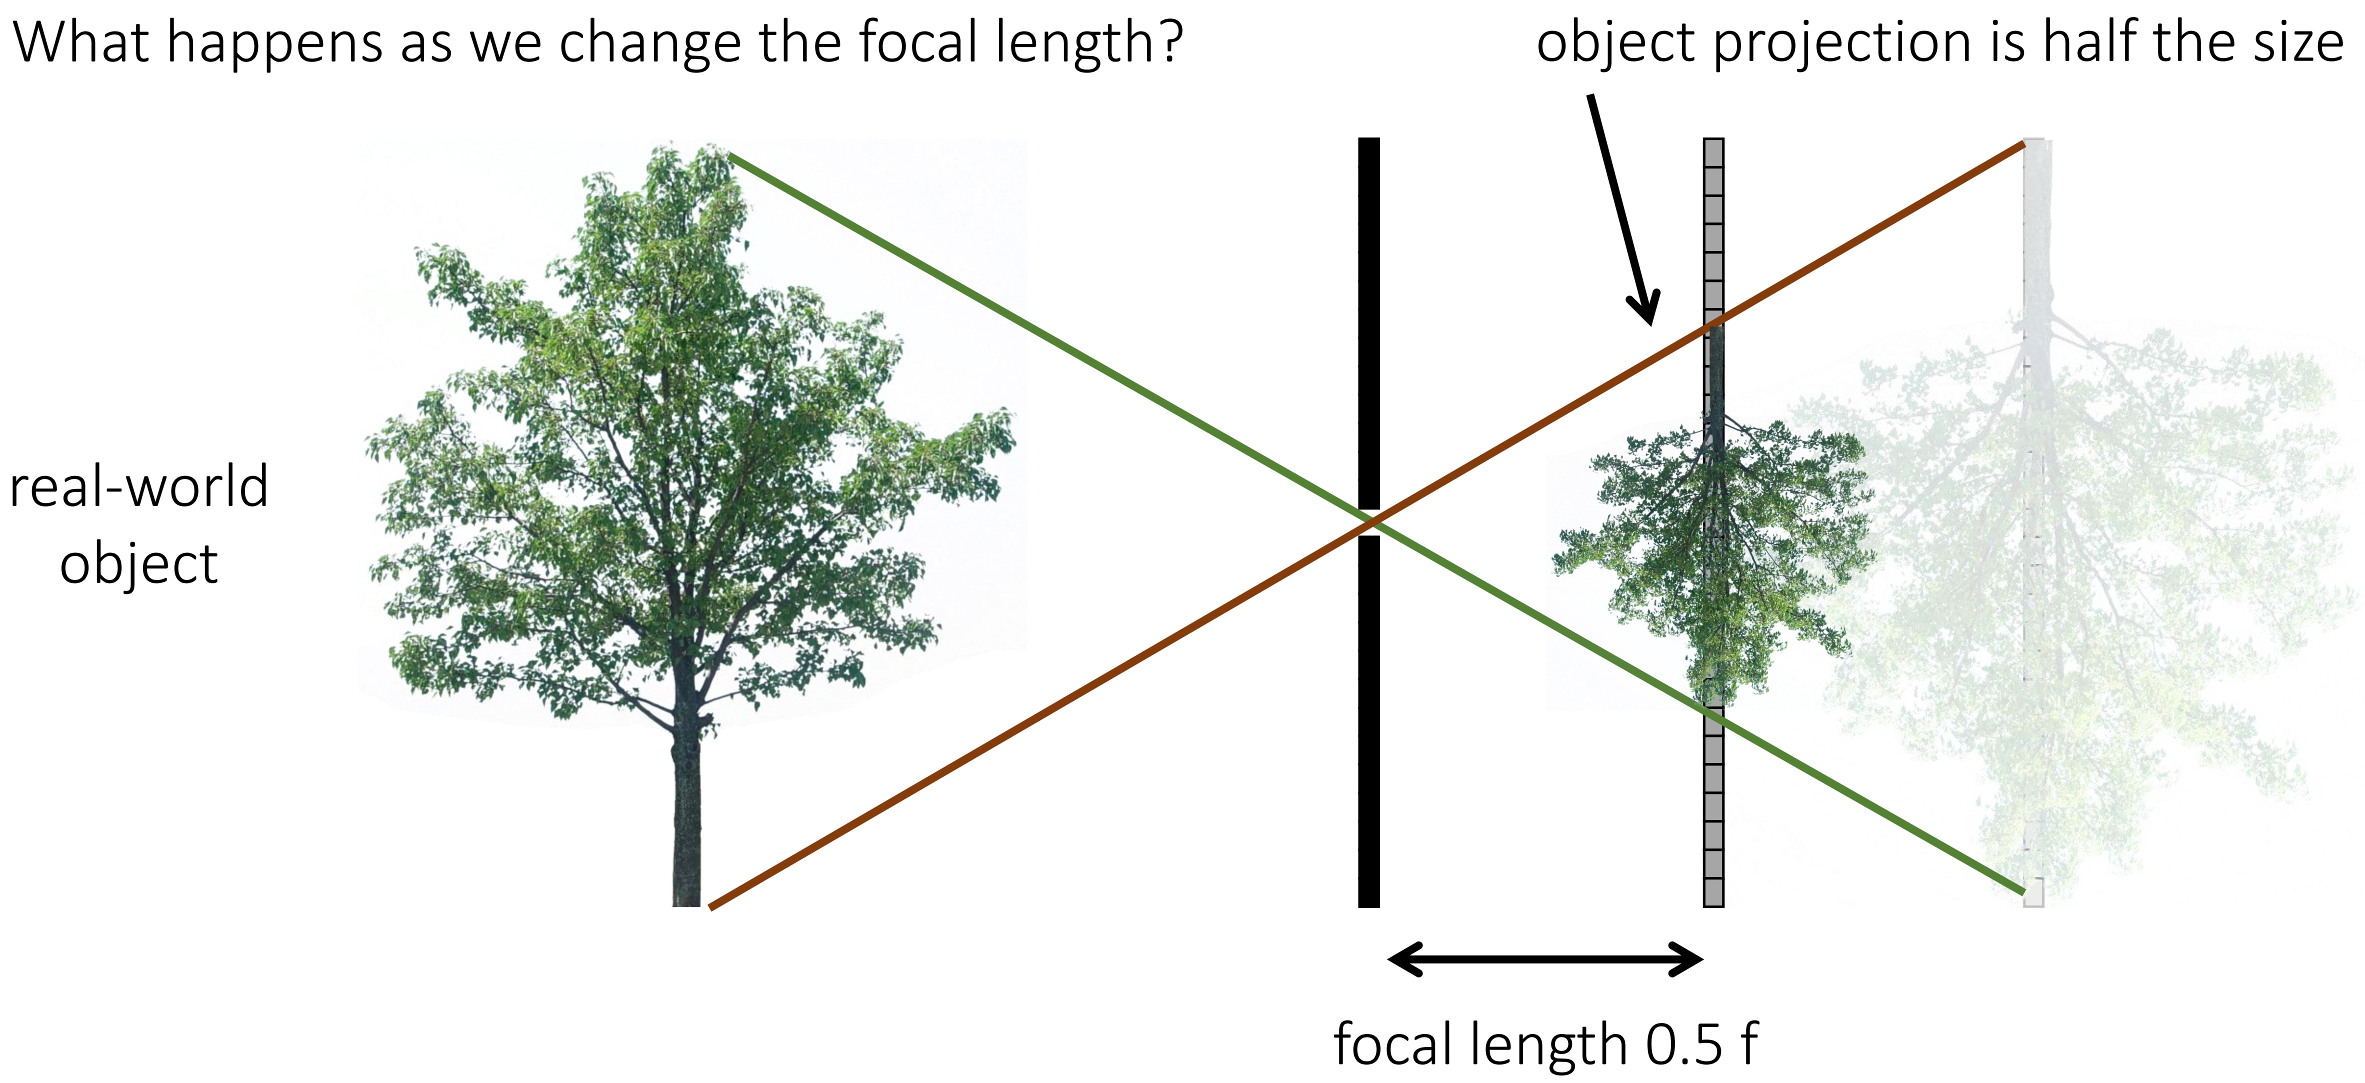

Pinhole size

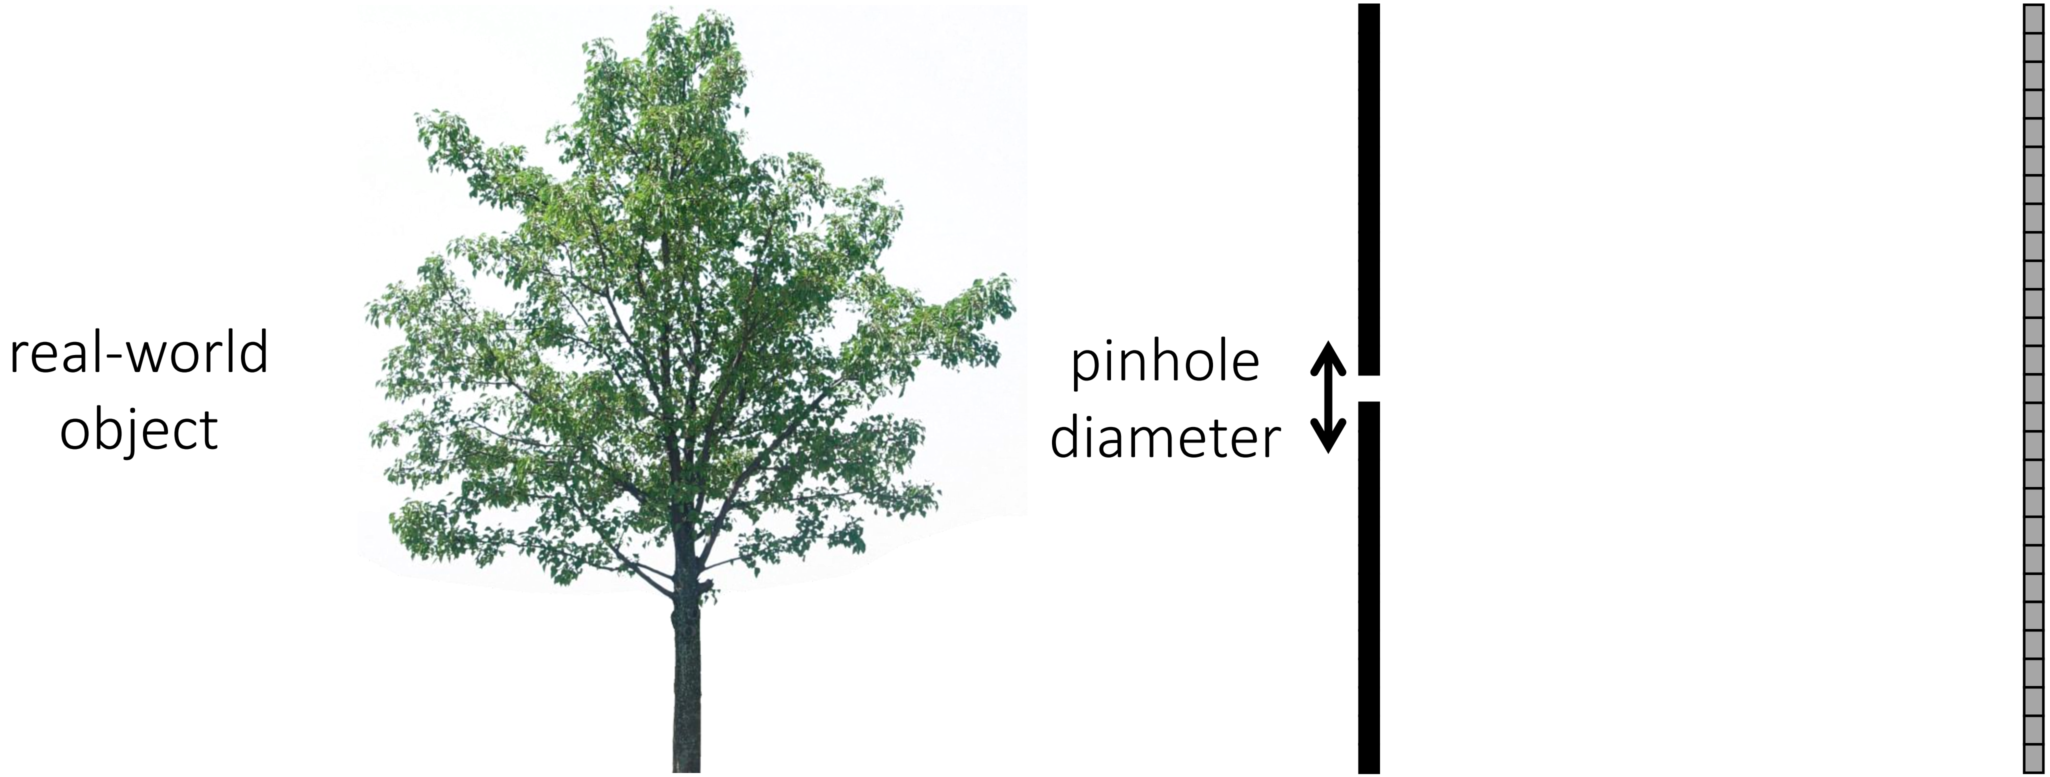

Ideal pinhole has infinitesimally small size

- In practice that is impossible.

Pinhole size

What happens as we change the pinhole diameter?

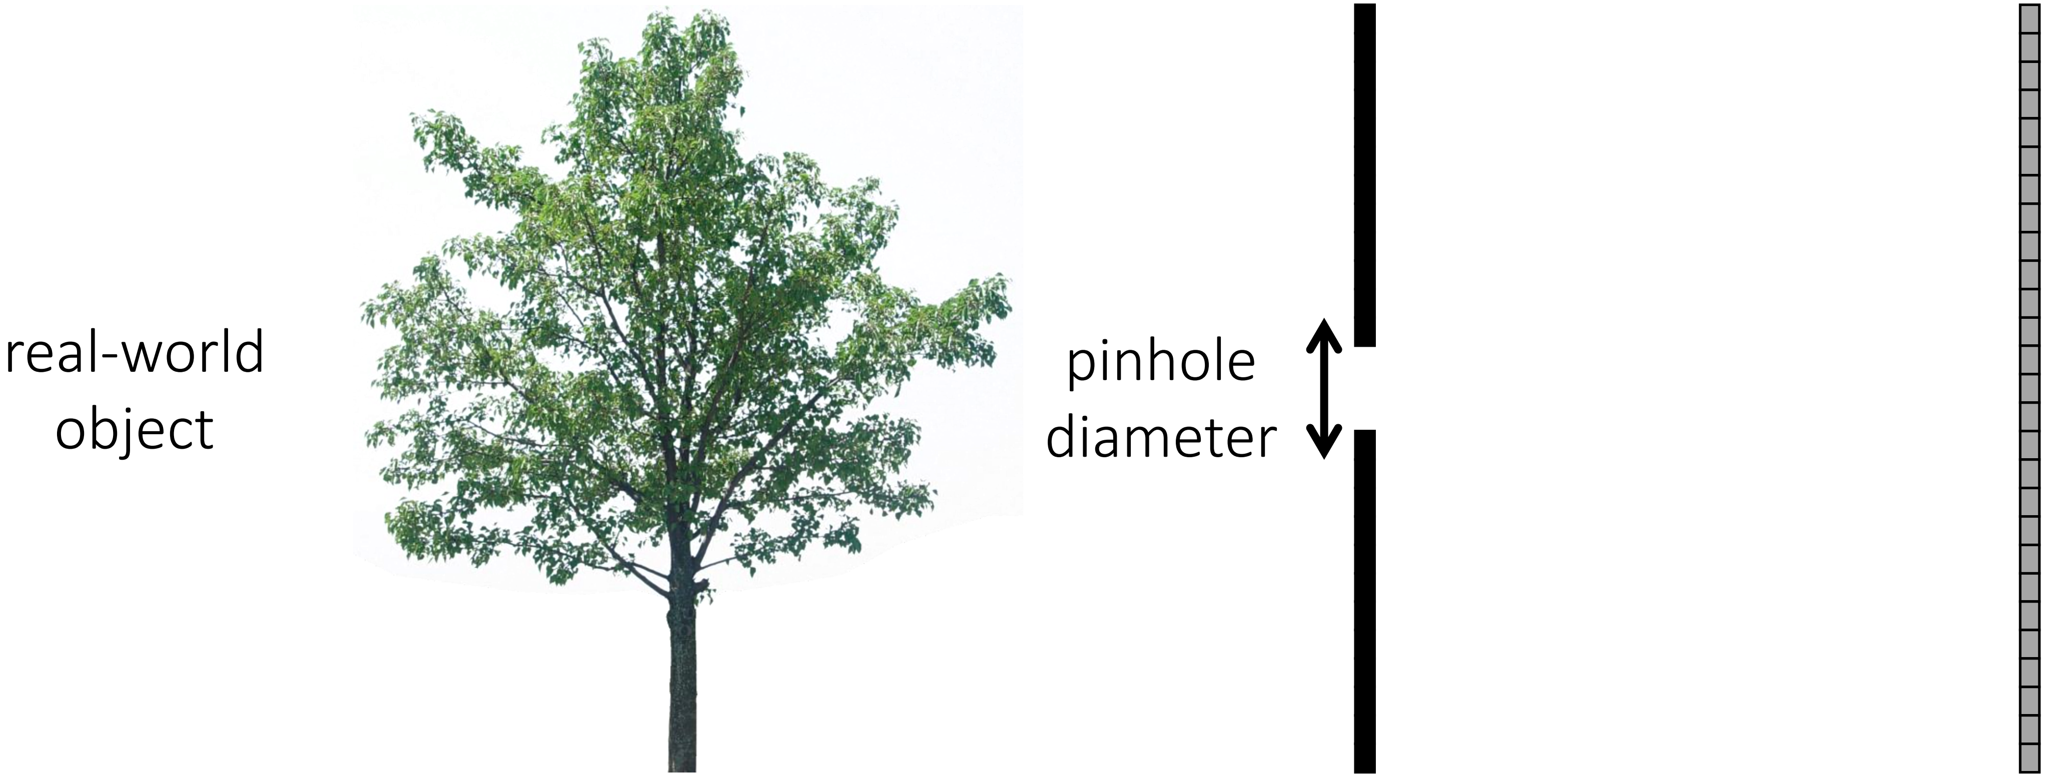

Pinhole size

What happens as we change the pinhole diameter?

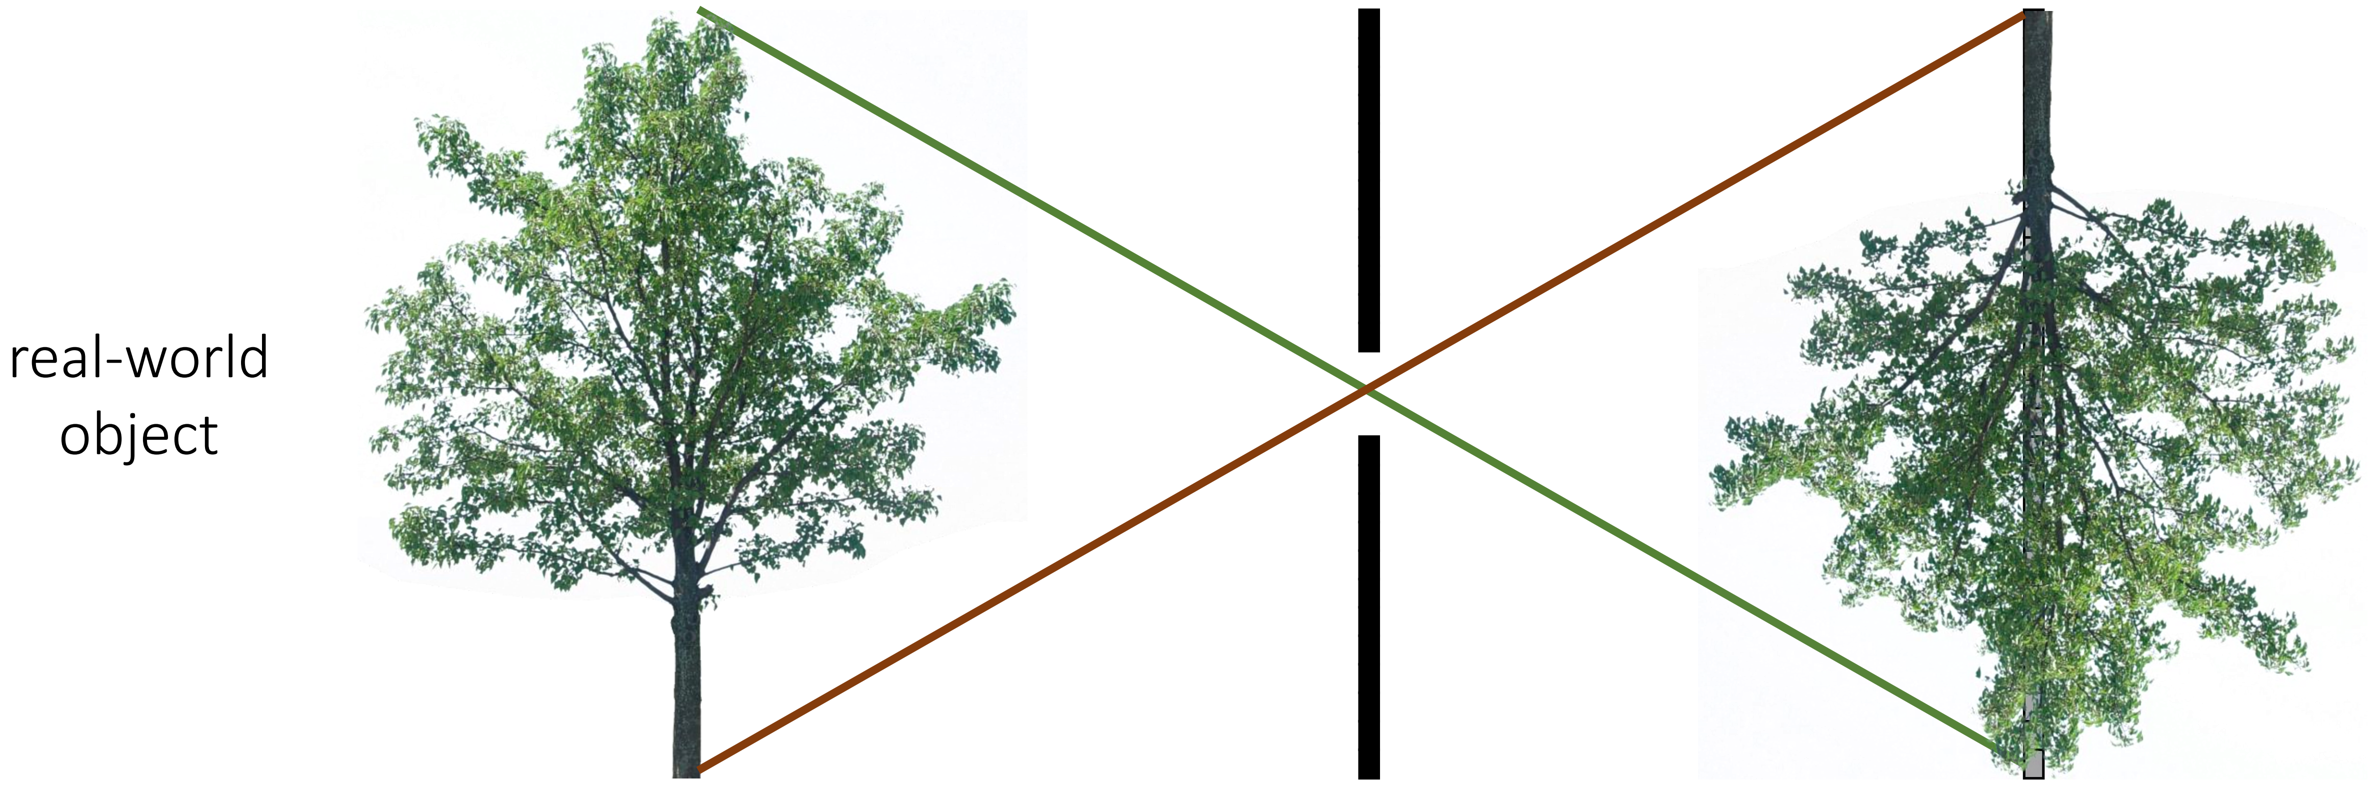

Pinhole size

What happens as we change the pinhole diameter?

Pinhole size

What happens as we change the pinhole diameter?

object projection becomes blurrier

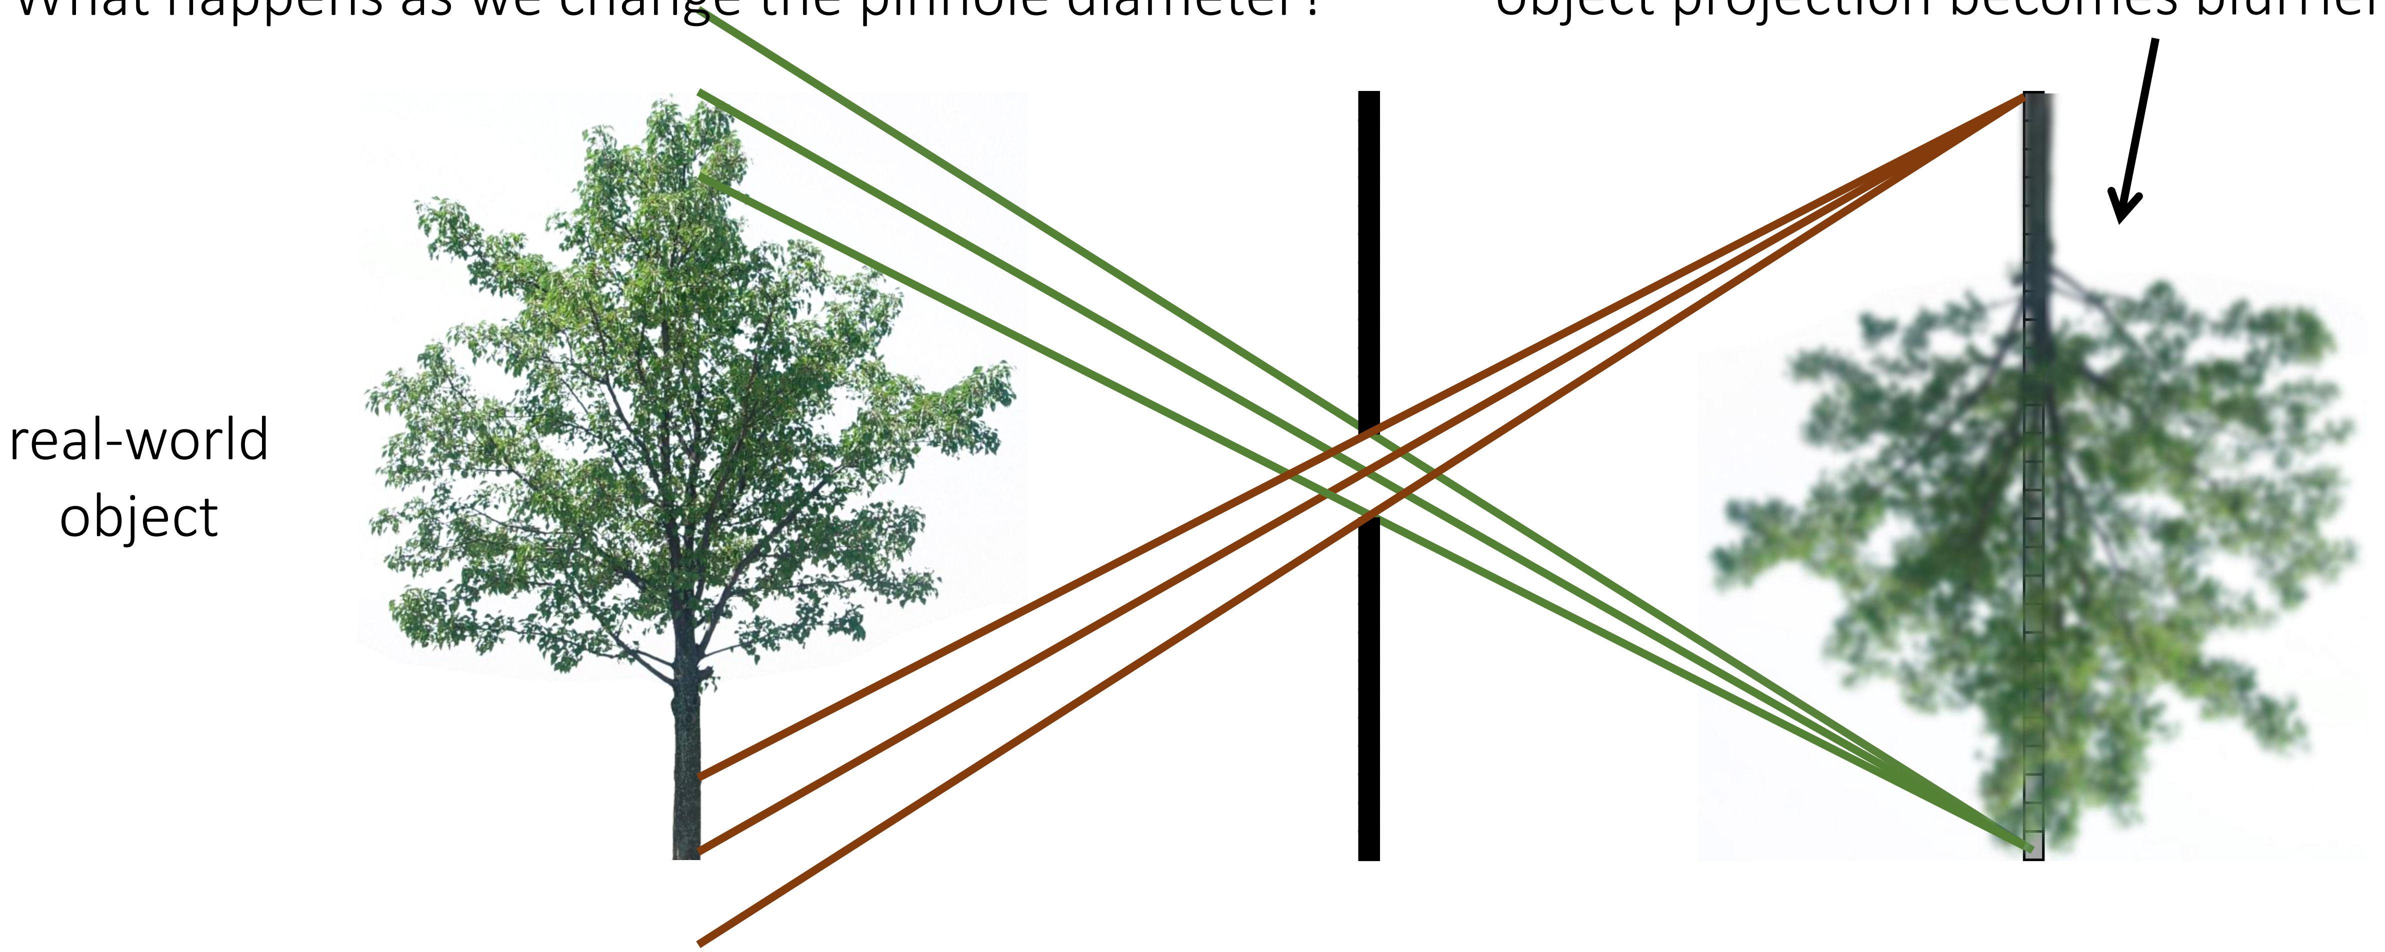

Pinhole size

What happens as we change the pinhole diameter?

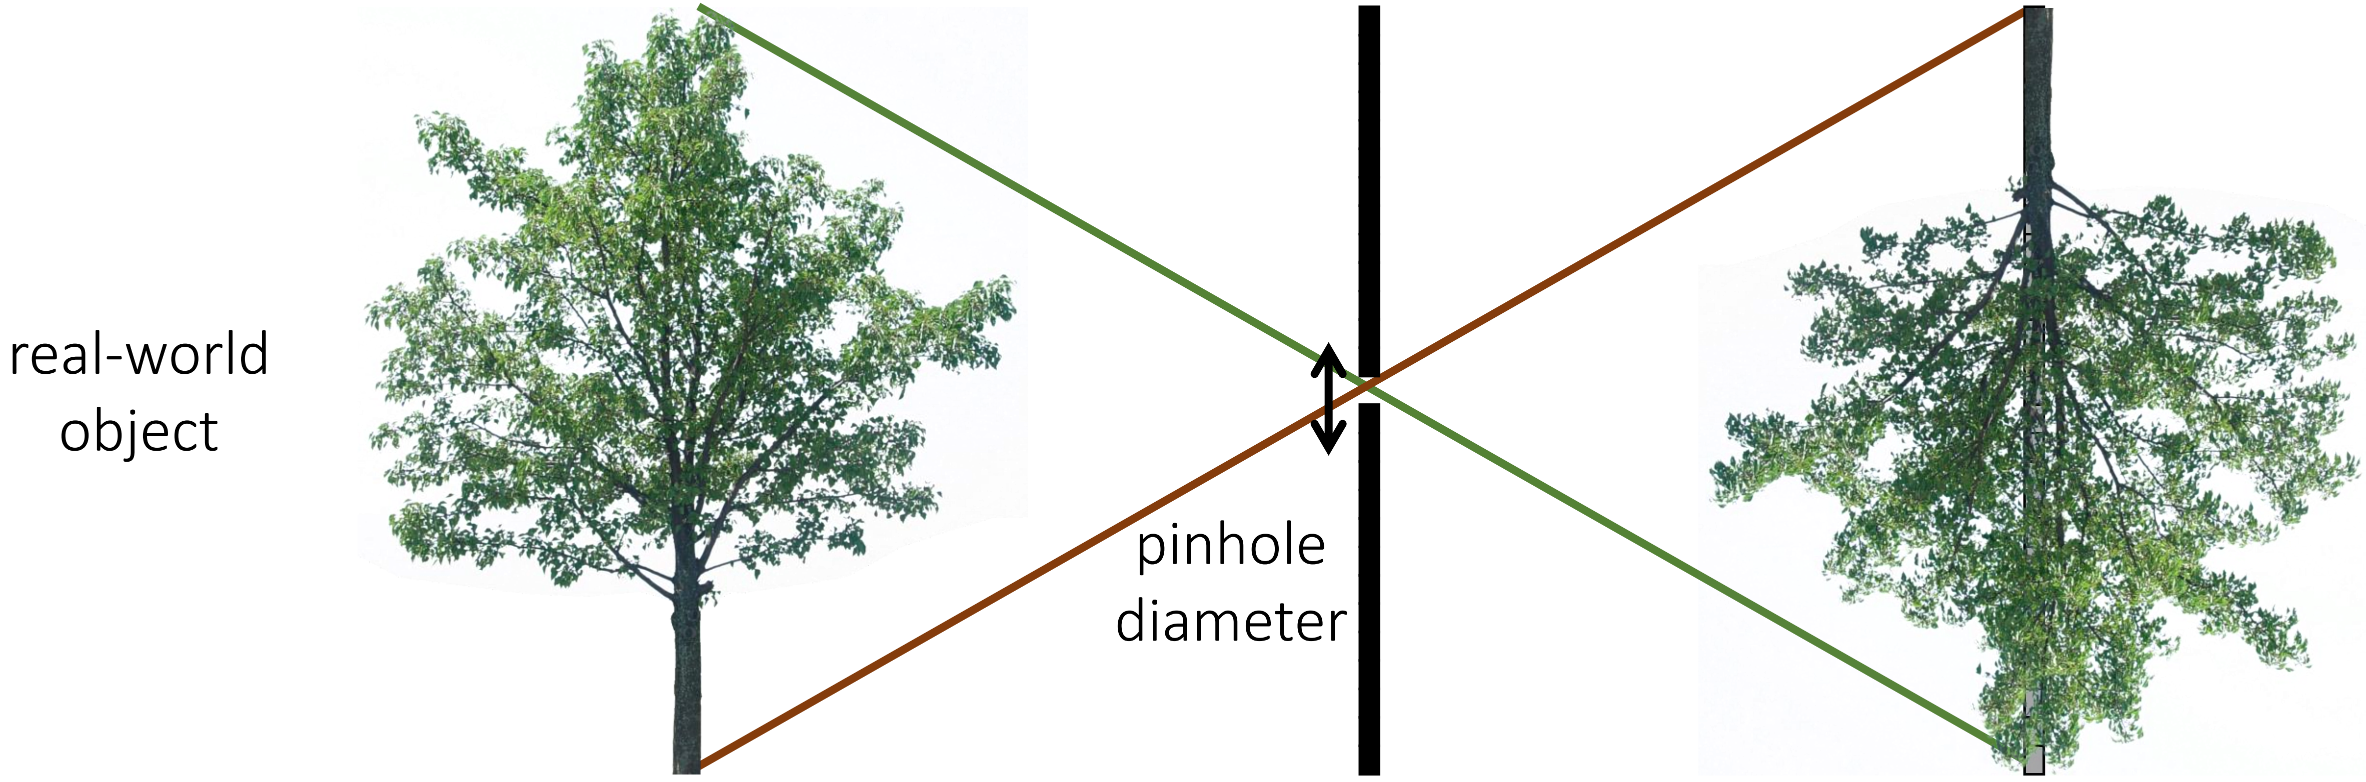

Will the image keep getting sharper the smaller we make the pinhole?

Diffraction limit

A consequence of the wave nature of light

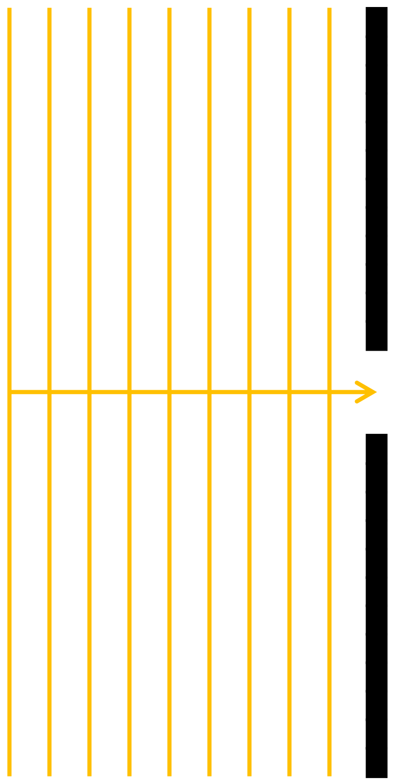

What do geometric optics predict will happen?

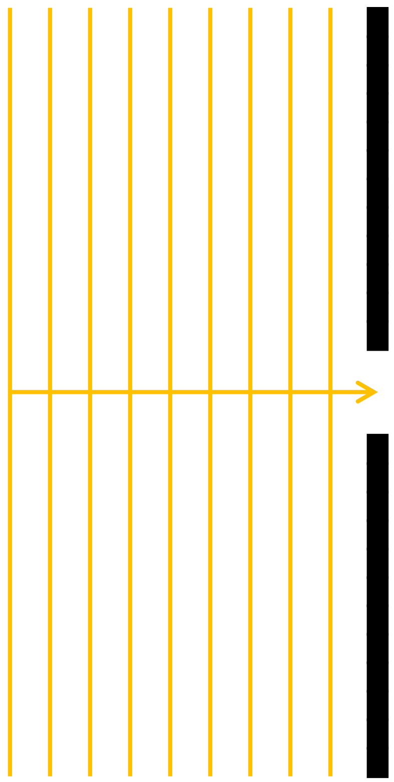

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

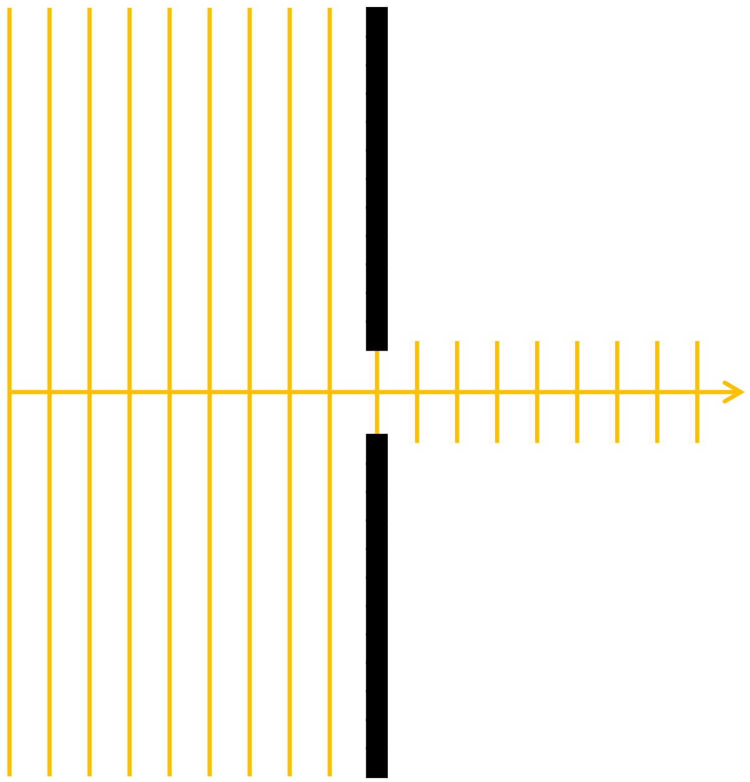

What do geometric optics predict will happen?

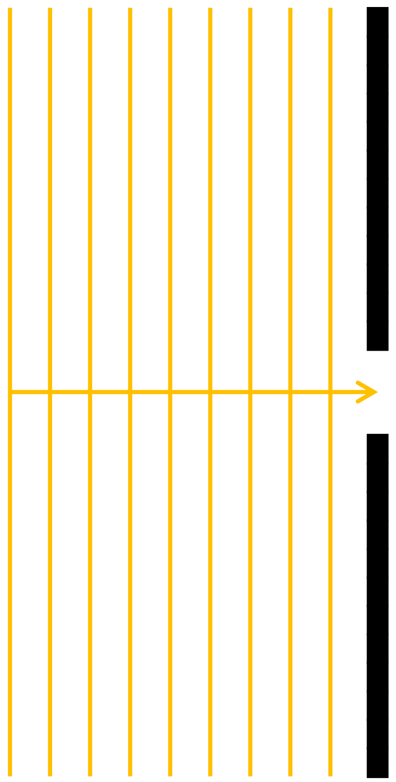

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

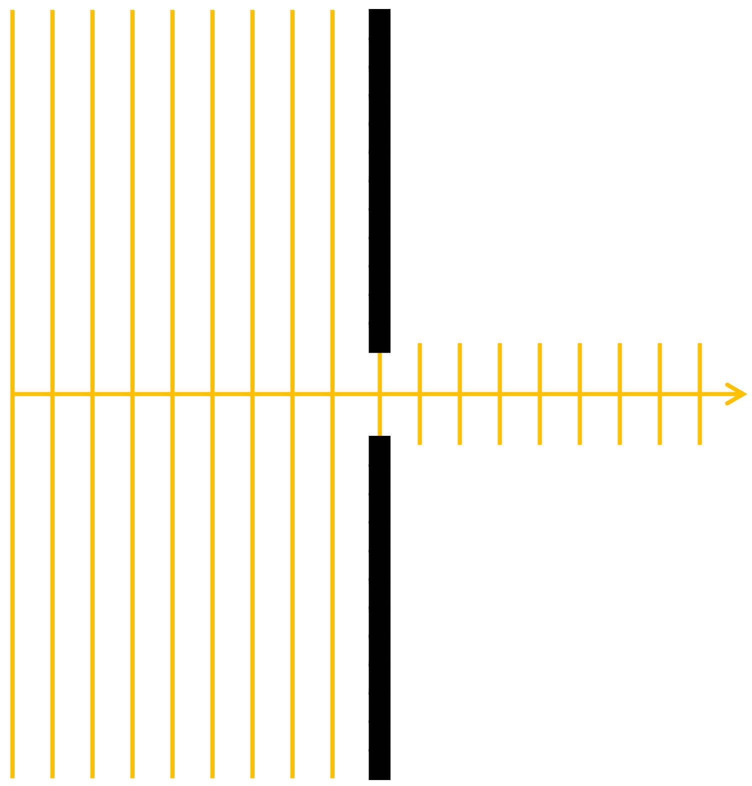

What do geometric optics predict will happen?

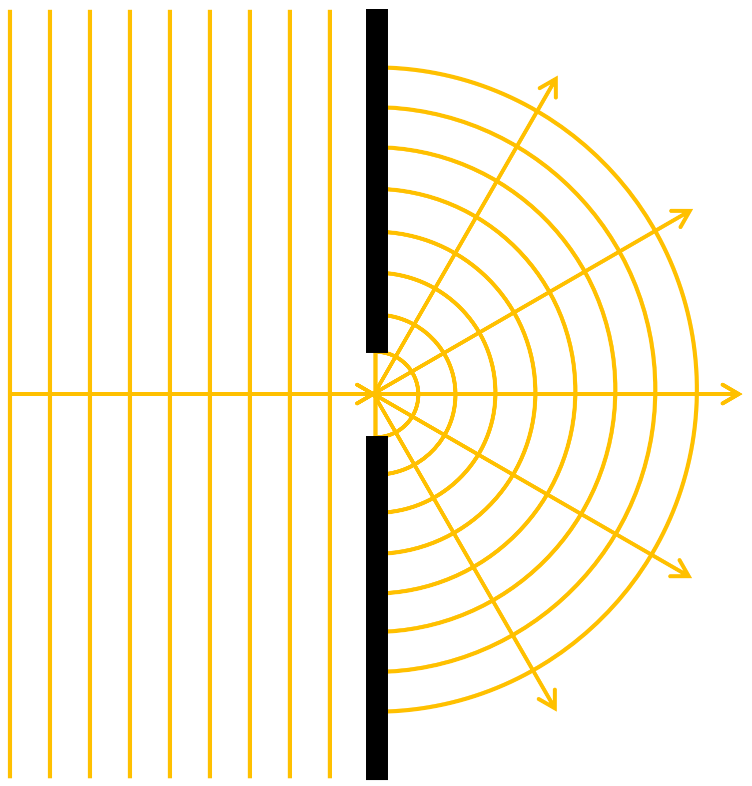

What do wave optics predict will happen?

Diffraction limit

Diffraction pattern = Fourier transform of the pinhole.

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

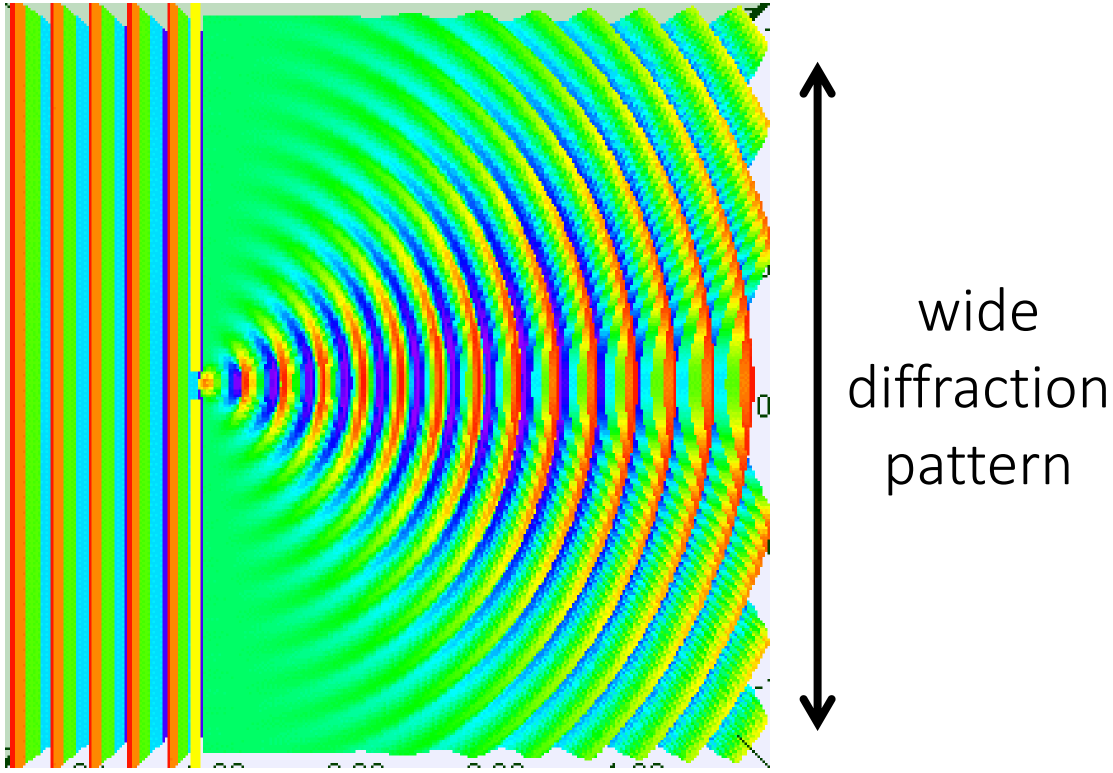

small pinhole

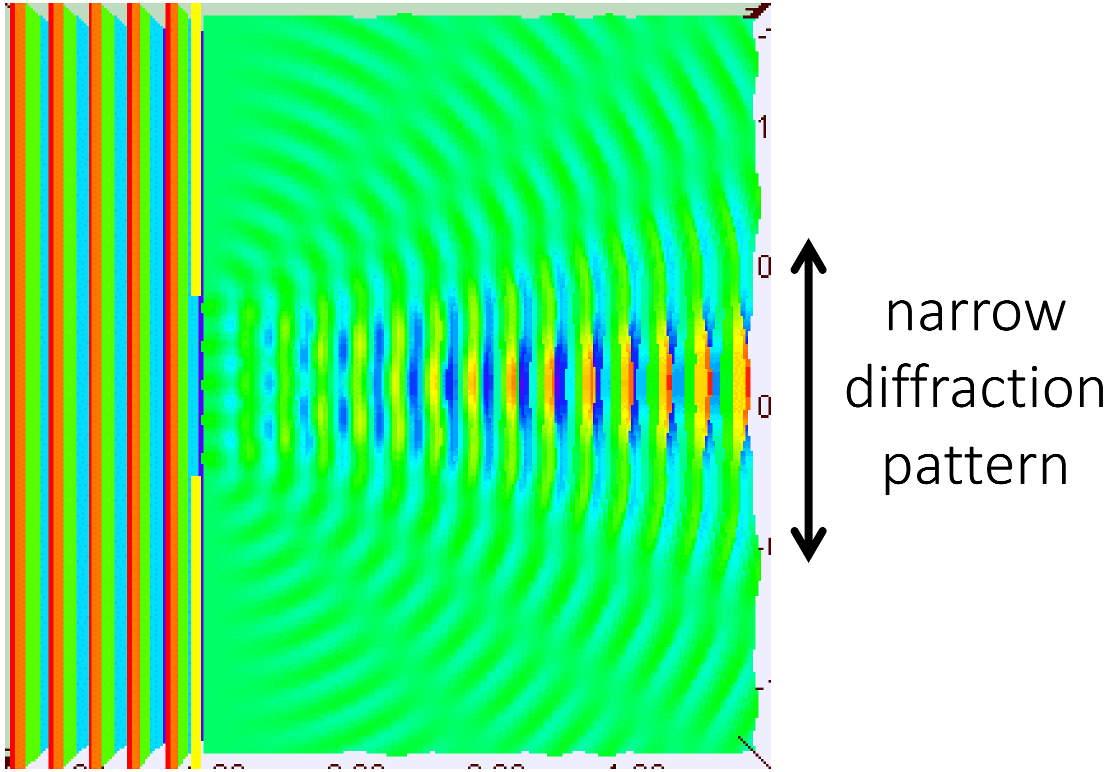

large pinhole

What about light efficiency?

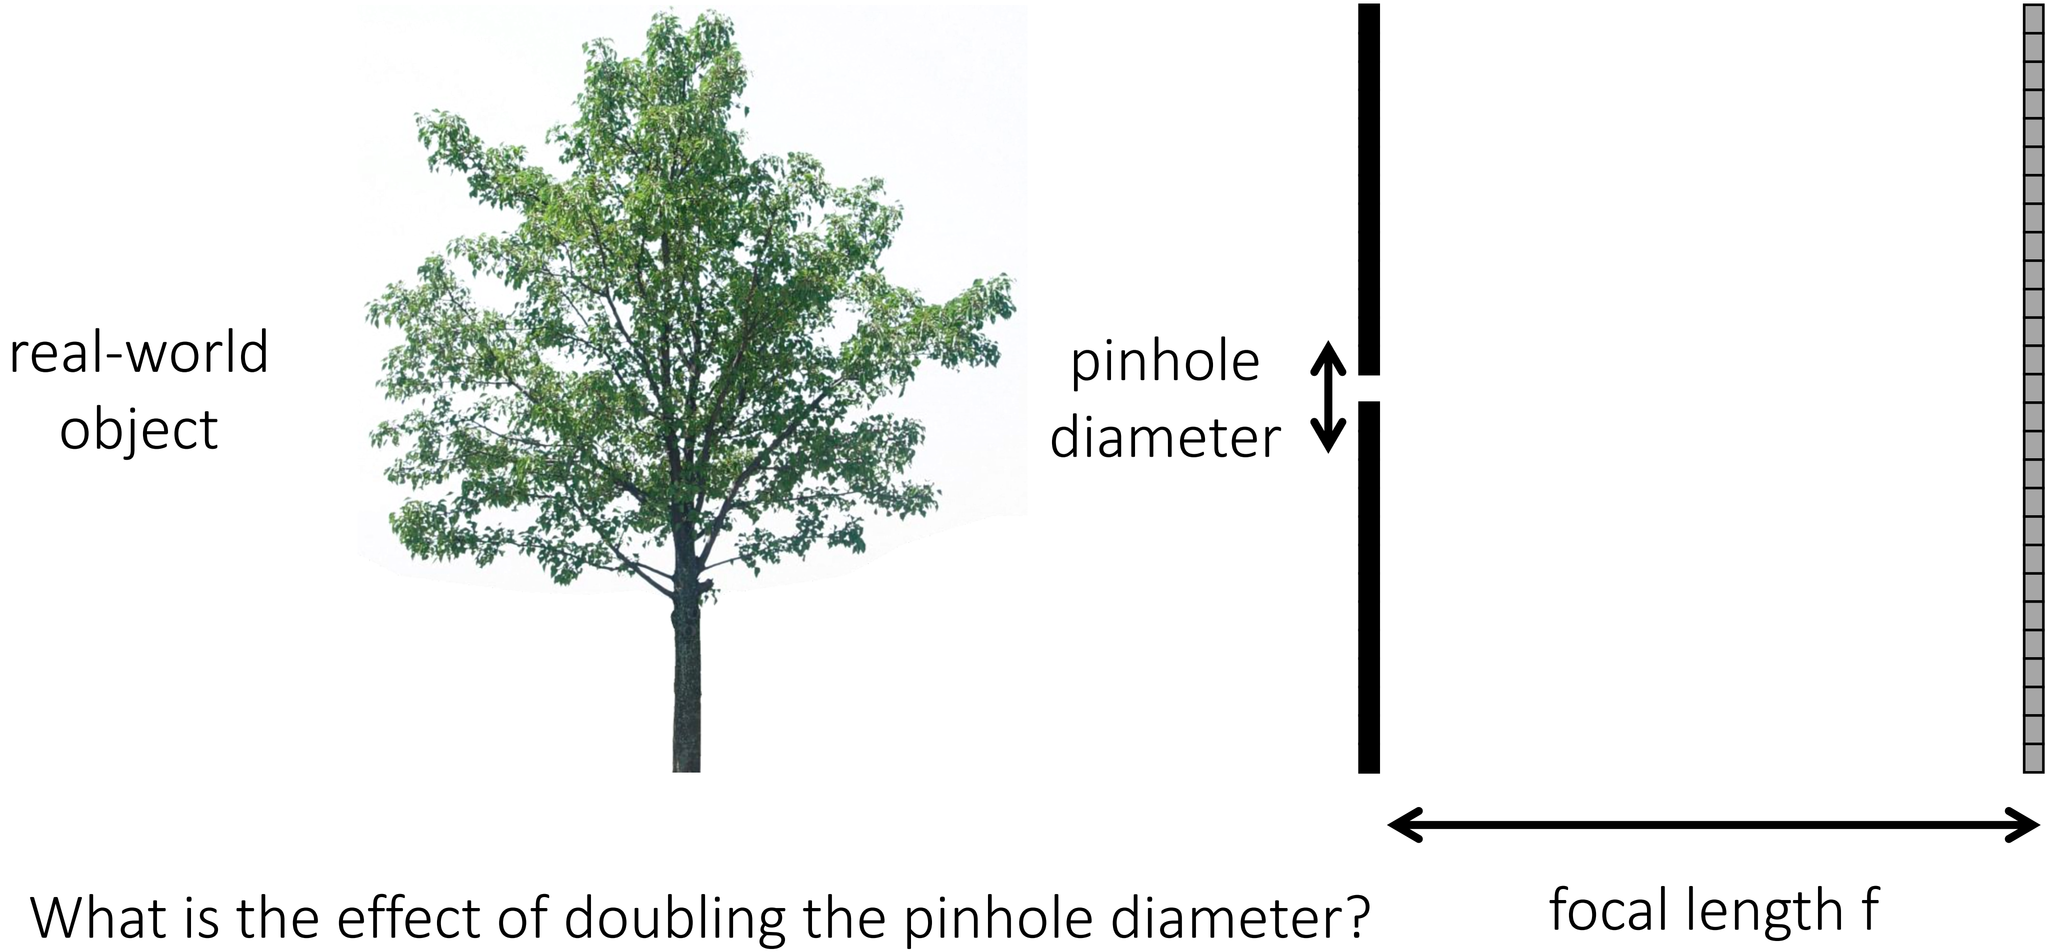

- What is the effect of doubling the pinhole diameter?
- What is the effect of doubling the focal length?

What about light efficiency?

real-world
object

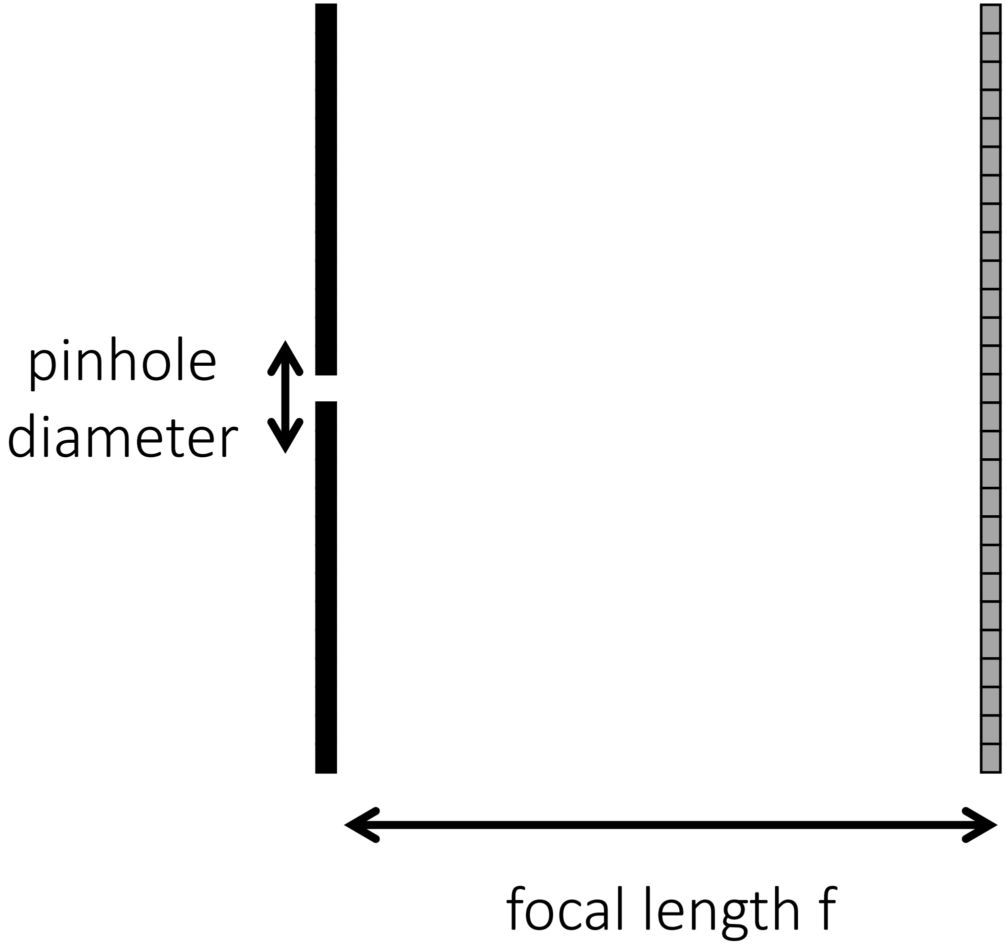

- 2x pinhole diameter → 4x light
- 2x focal length → ¼x light

Some terminology notes

A “stop” is a change in camera settings that changes amount of light by a factor of 2

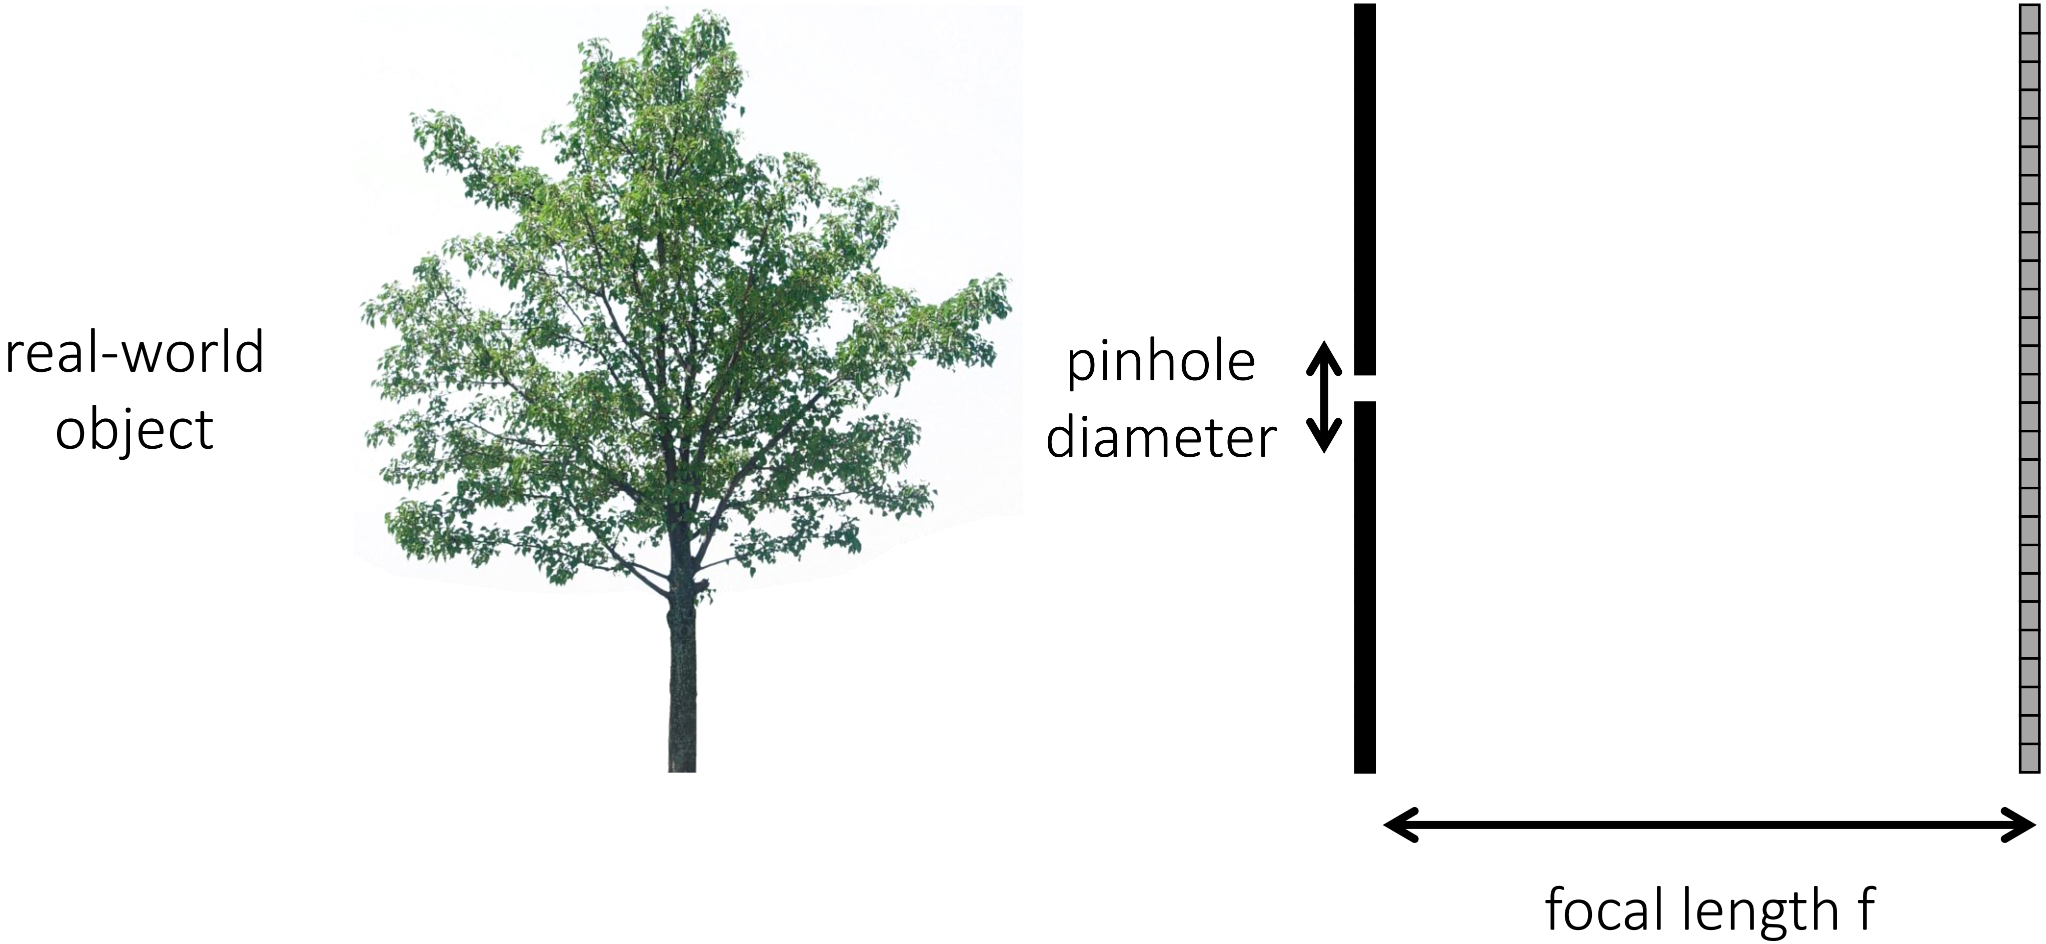

The “f-number” is the ratio: focal length / pinhole diameter

Accidental pinholes

What does this image say about the world outside?

Accidental pinhole camera

Antonio Torralba, William T. Freeman
Computer Science and Artificial Intelligence Laboratory (CSAIL)
MIT
torralba@mit.edu, billf@mit.edu

Accidental pinhole camera

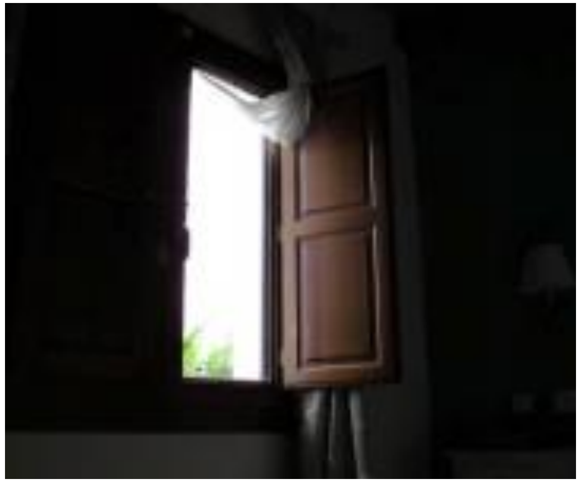

window is an aperture

projected pattern on the wall

upside down

window with smaller gap

view outside window

Accidental pinspeck camera

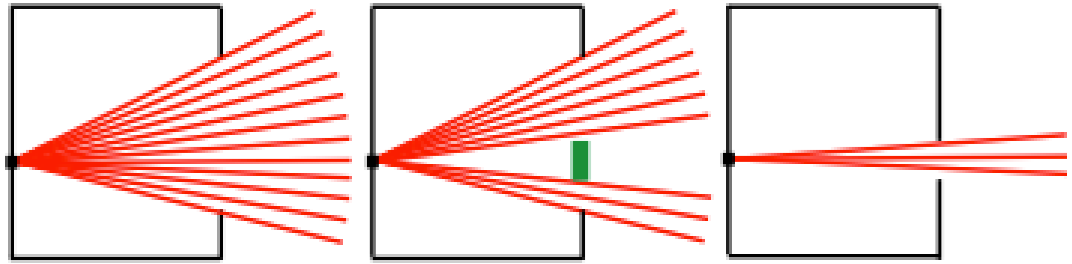

a) b) c)

a)

b)

c)

d)

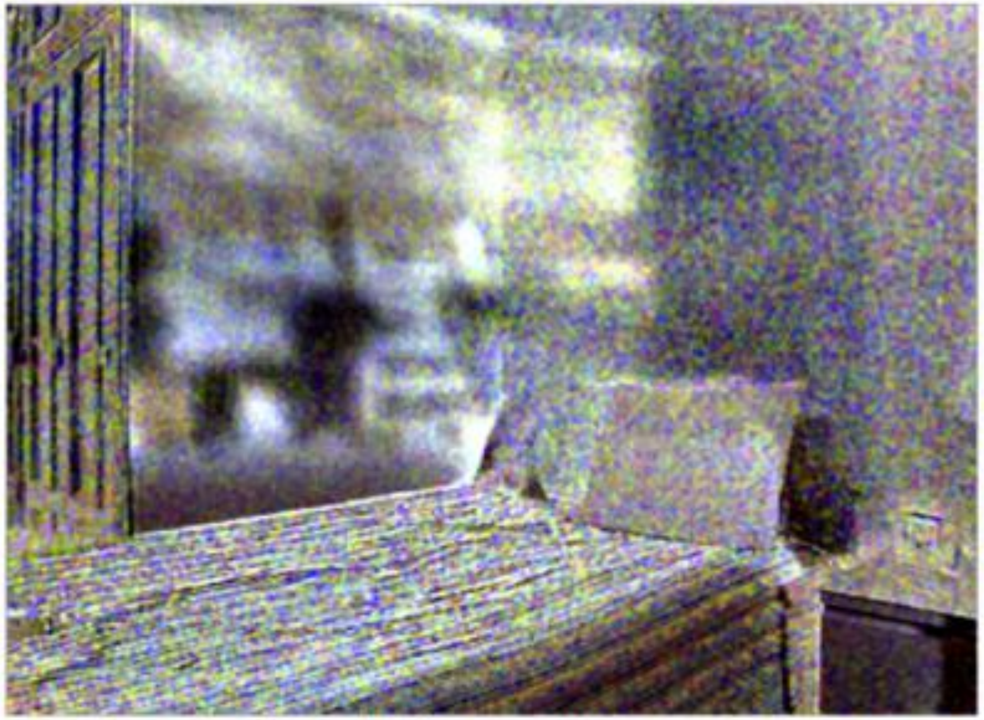

a) Difference image

b) Difference upside down

c) True outdoor view

Pinhole camera trade-off

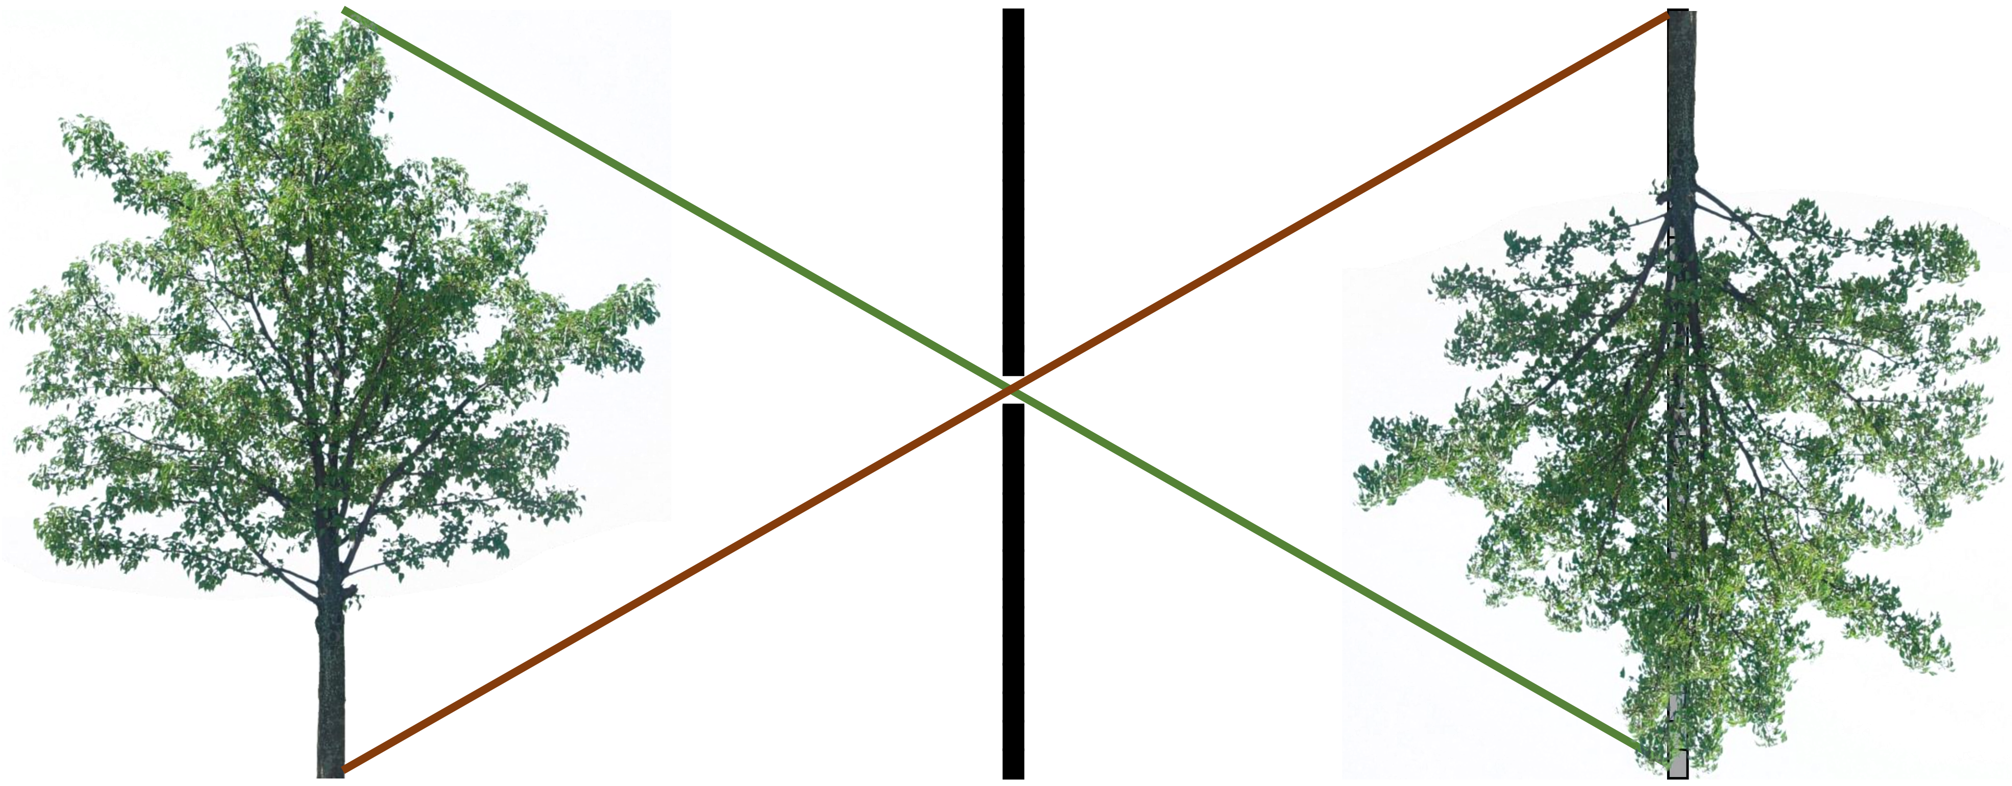

Small (ideal) pinhole:

1. Image is sharp.
2. Signal-to-noise ratio is low.

Pinhole camera trade-off

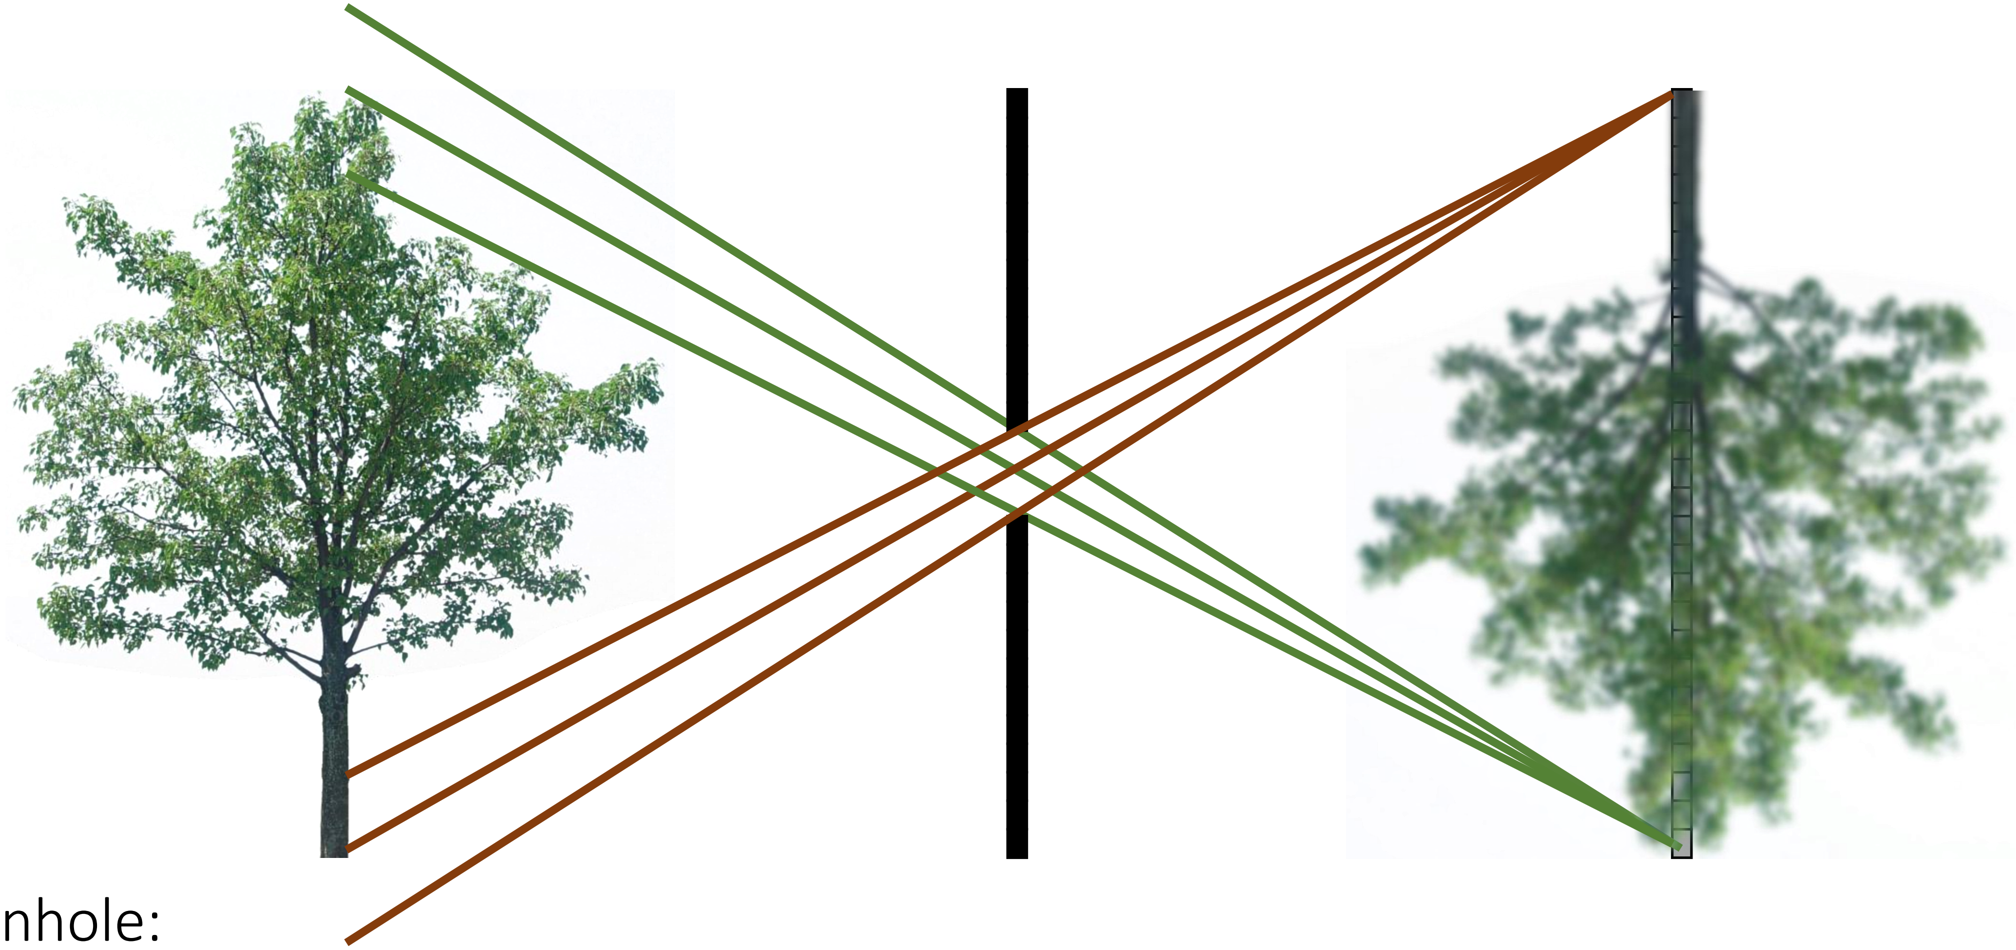

Large pinhole:

1. Image is blurry.
2. Signal-to-noise ratio is high.

Can we get best of both worlds?

Almost, by using lenses

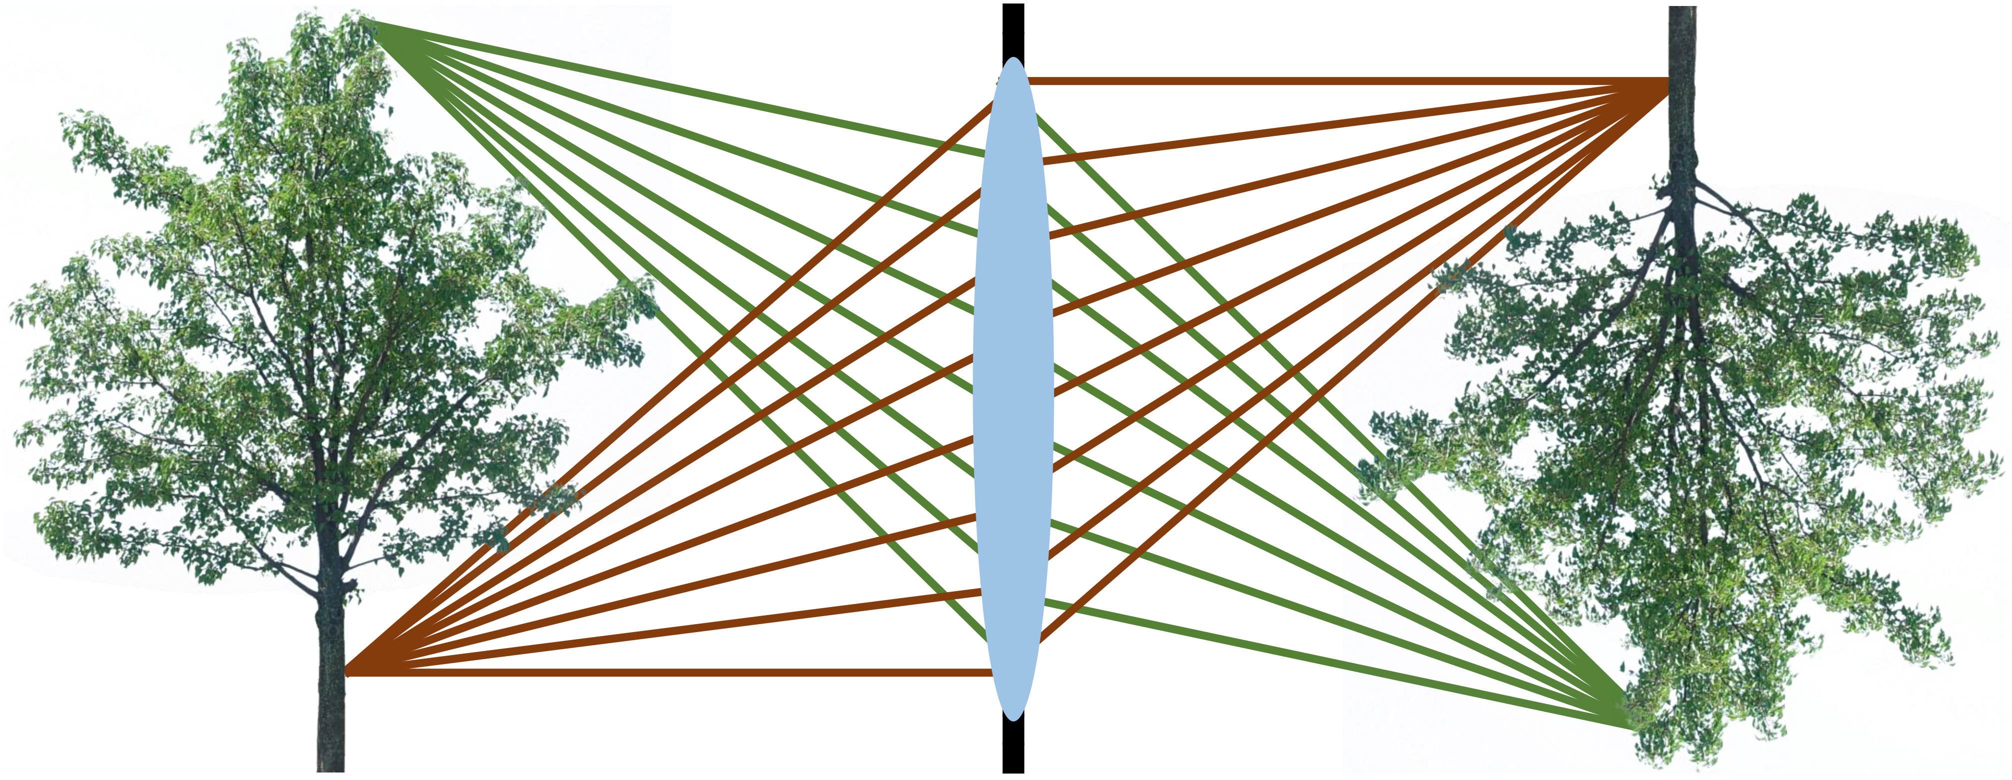

Lenses map "bundles" of rays from points on the scene to the sensor.

How does this mapping work exactly?

The thin lens model

Thin lens model

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

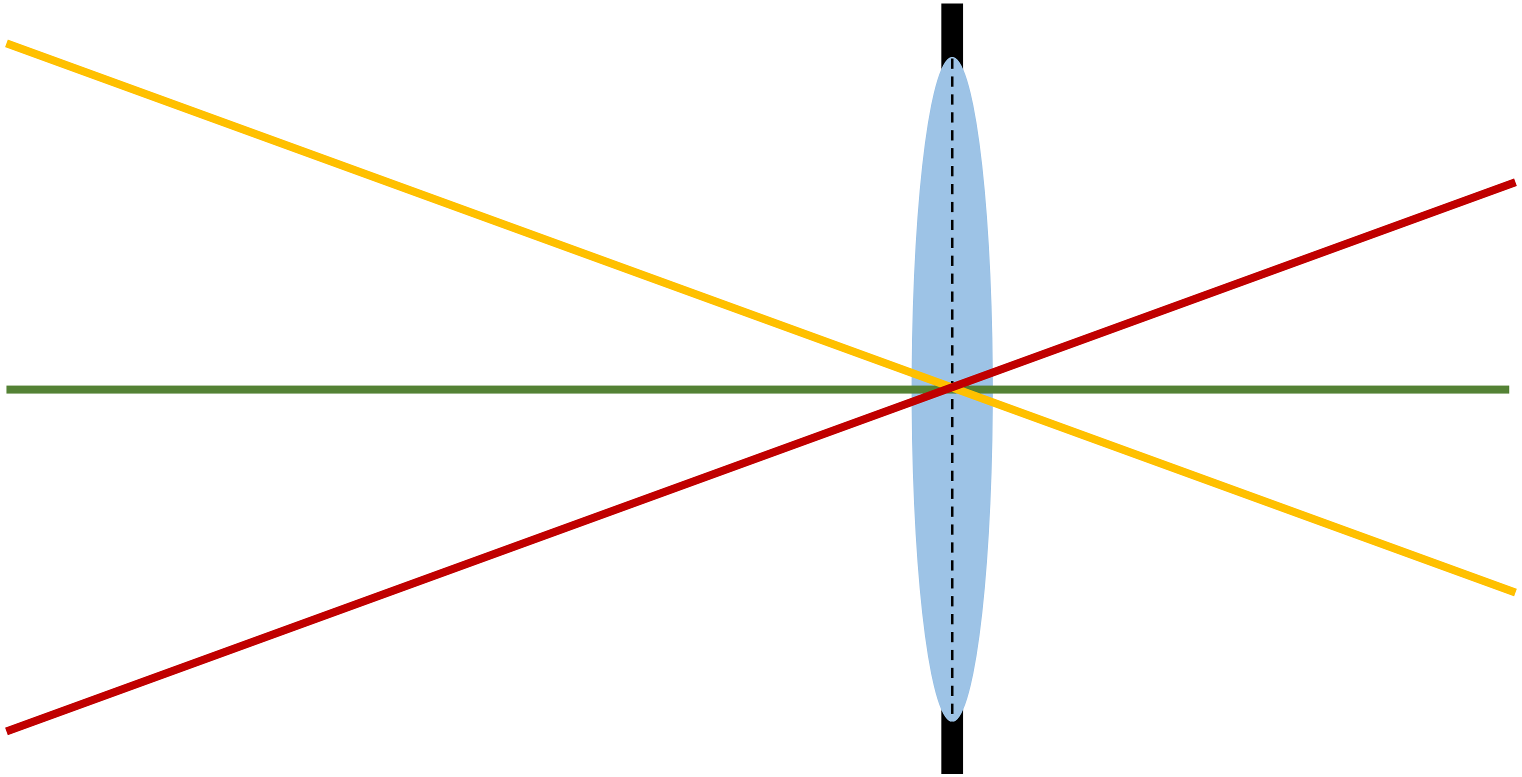

Two assumptions:

1. Rays passing through lens center are unaffected.

Thin lens model

Simplification of geometric optics for well-designed lenses.

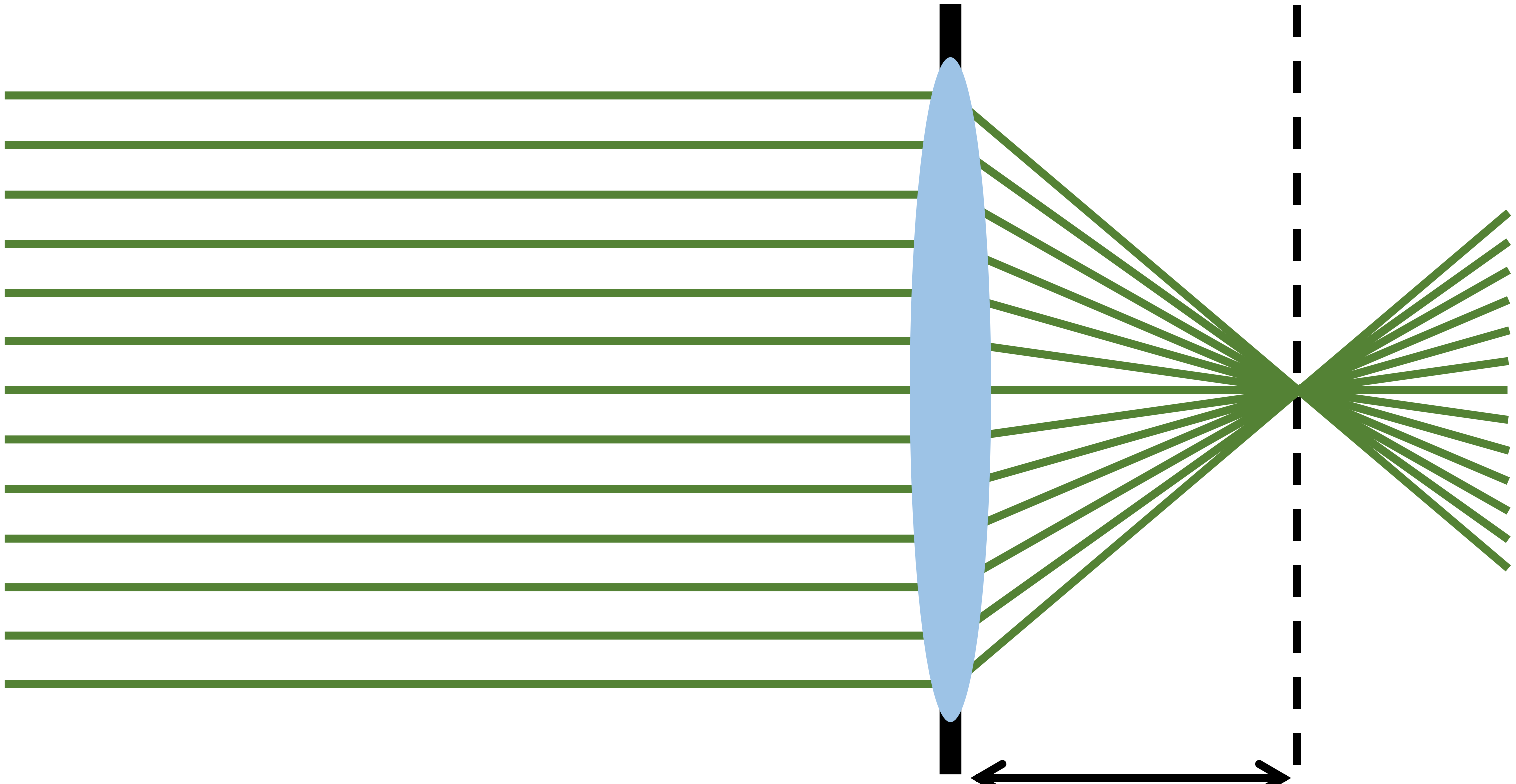

Two assumptions:

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.

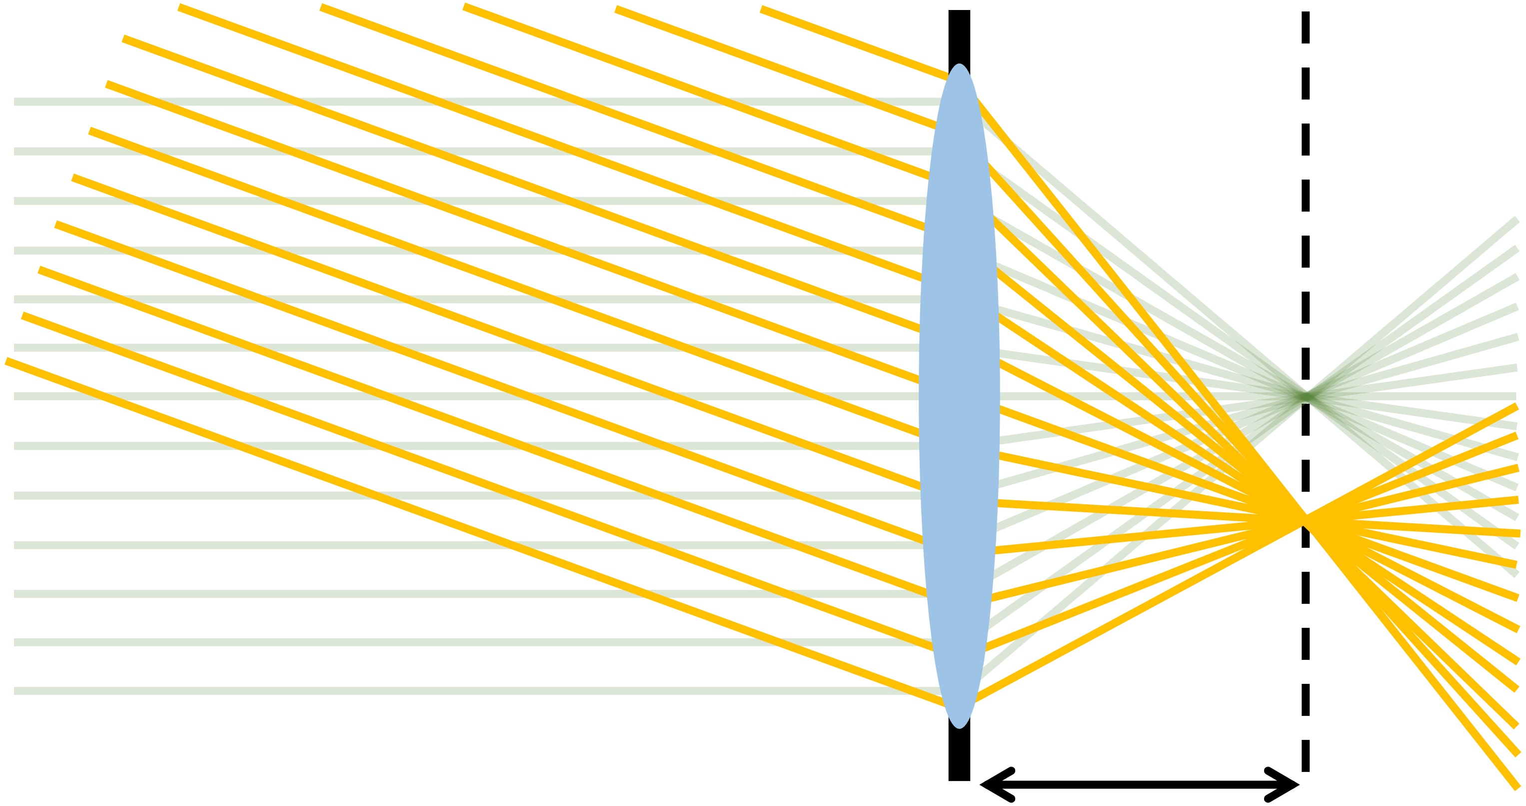

Two assumptions:

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.

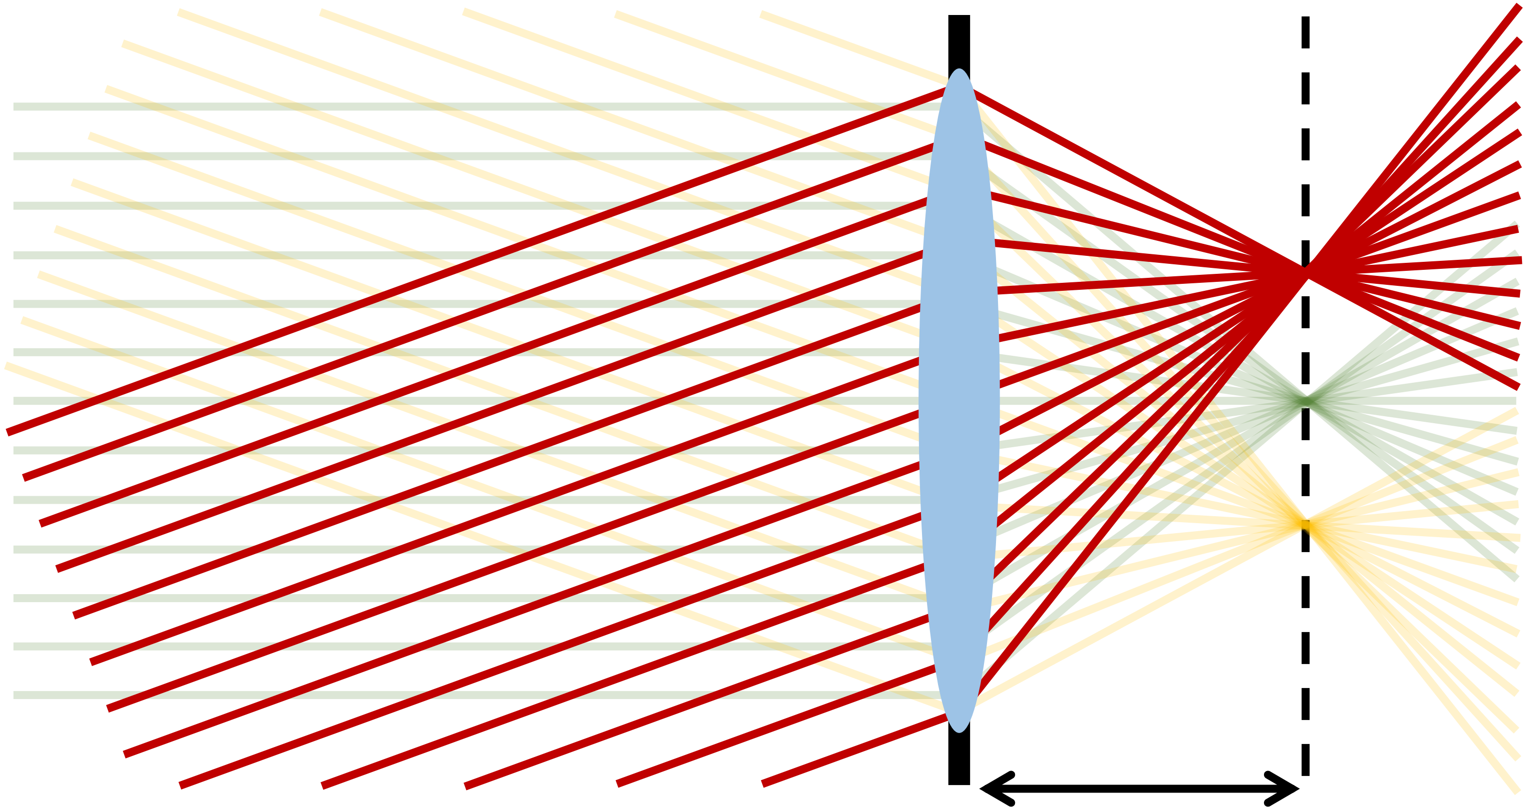

Two assumptions:

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

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

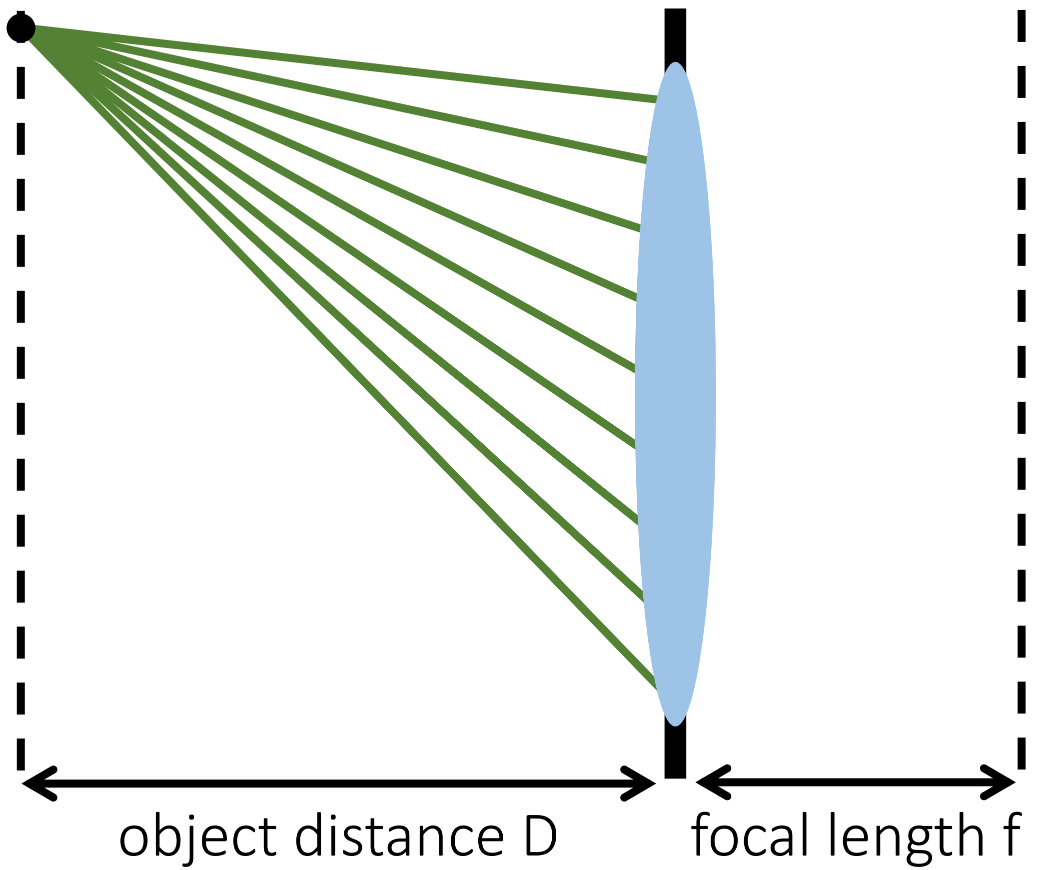

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.

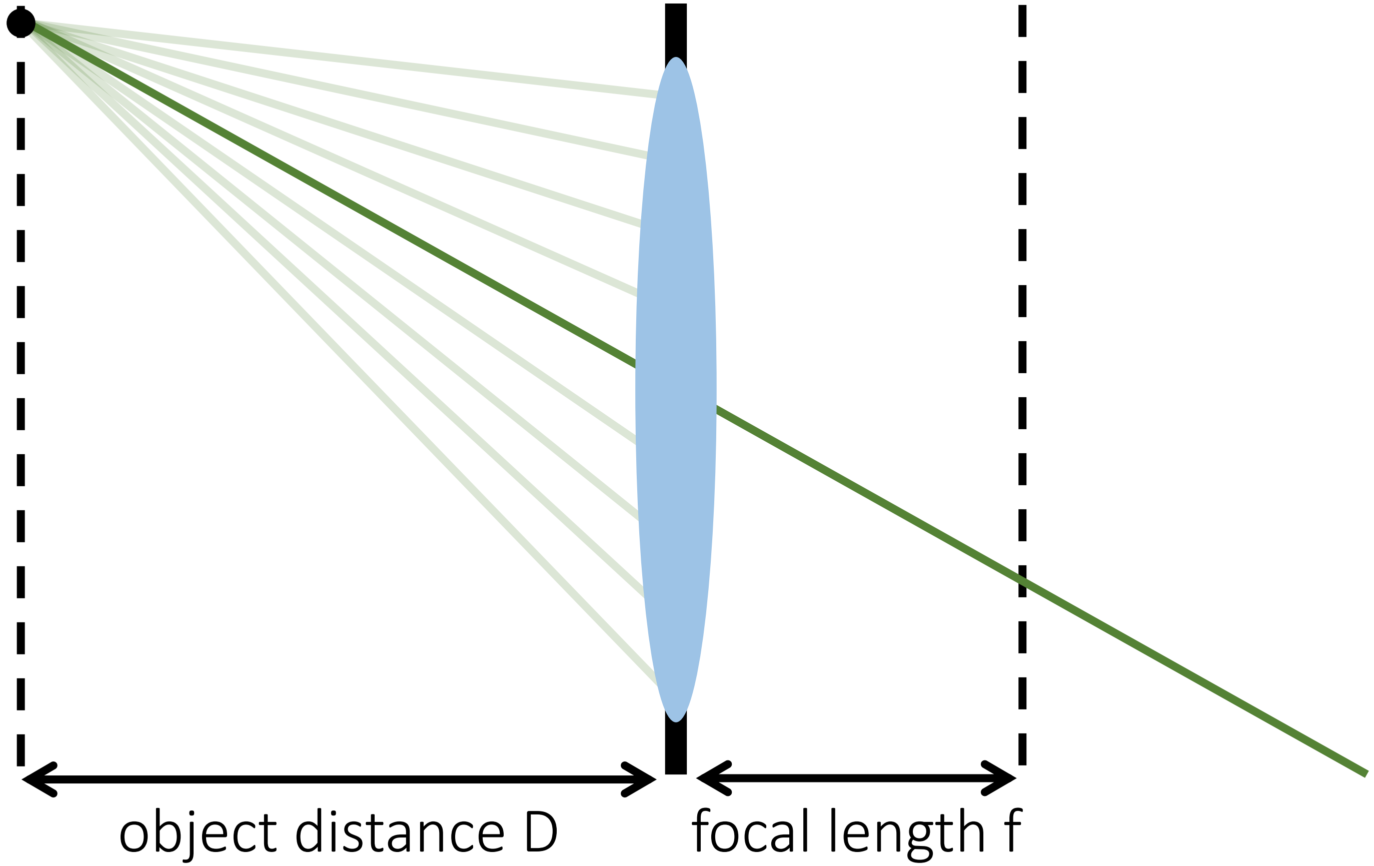

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:

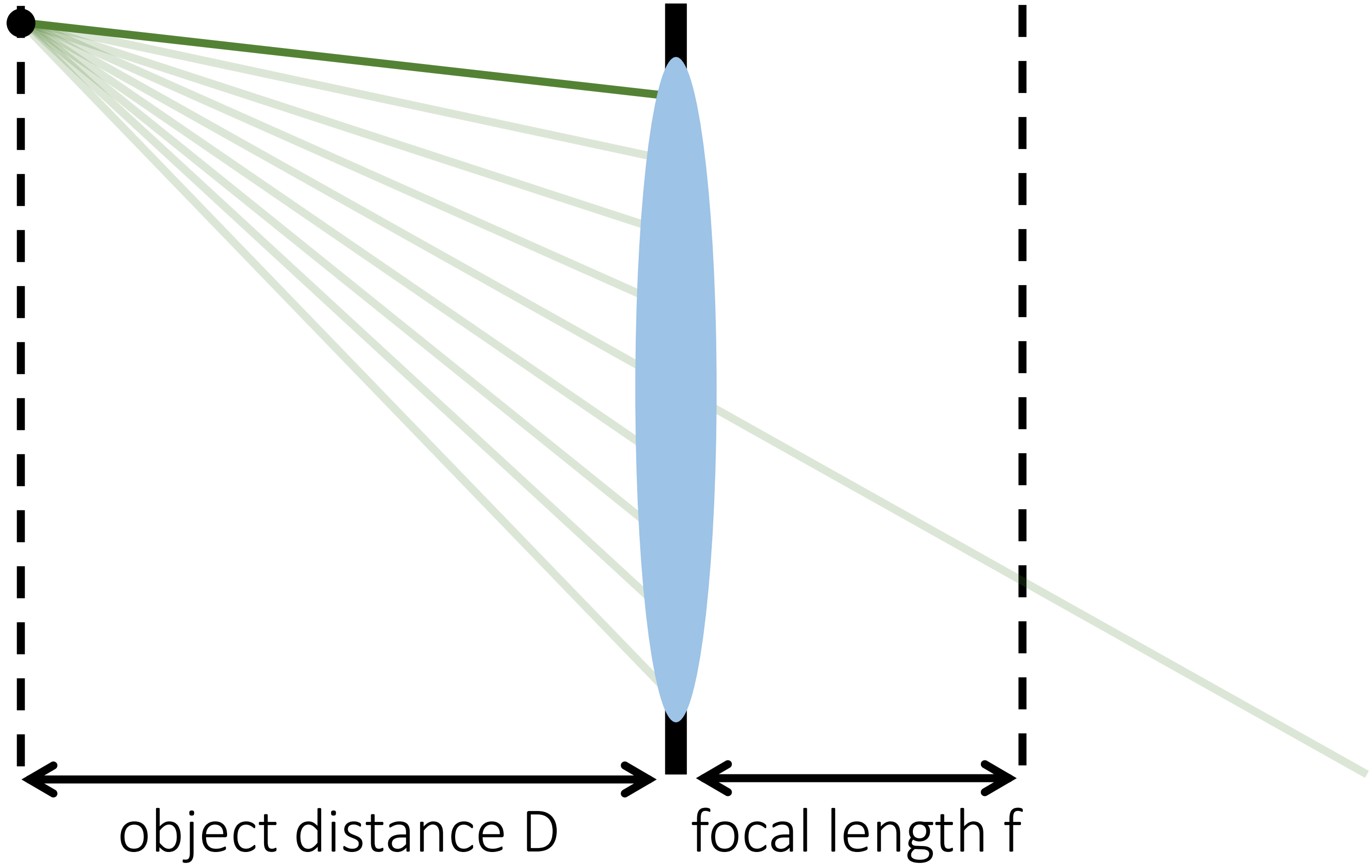

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.

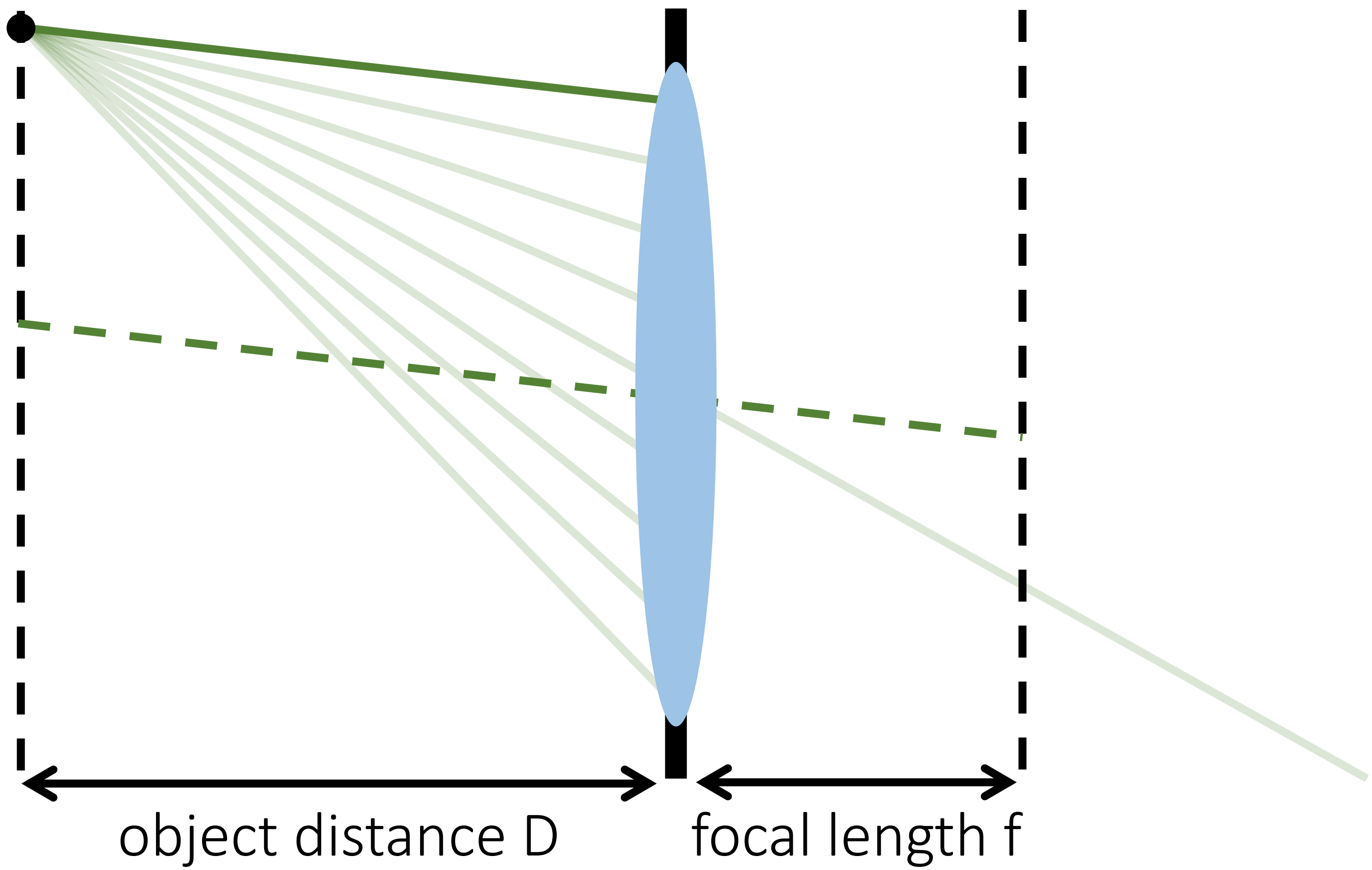

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

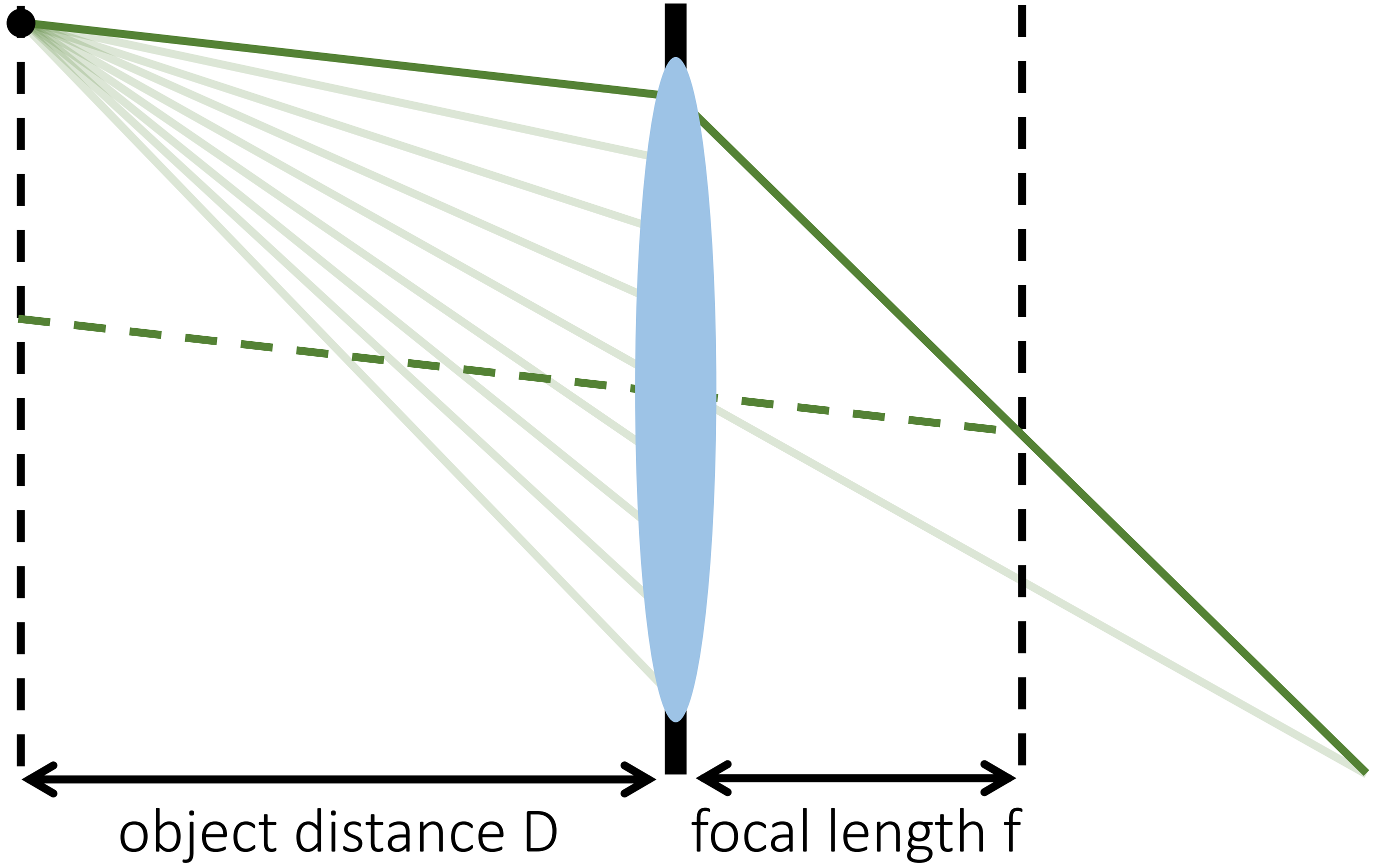

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

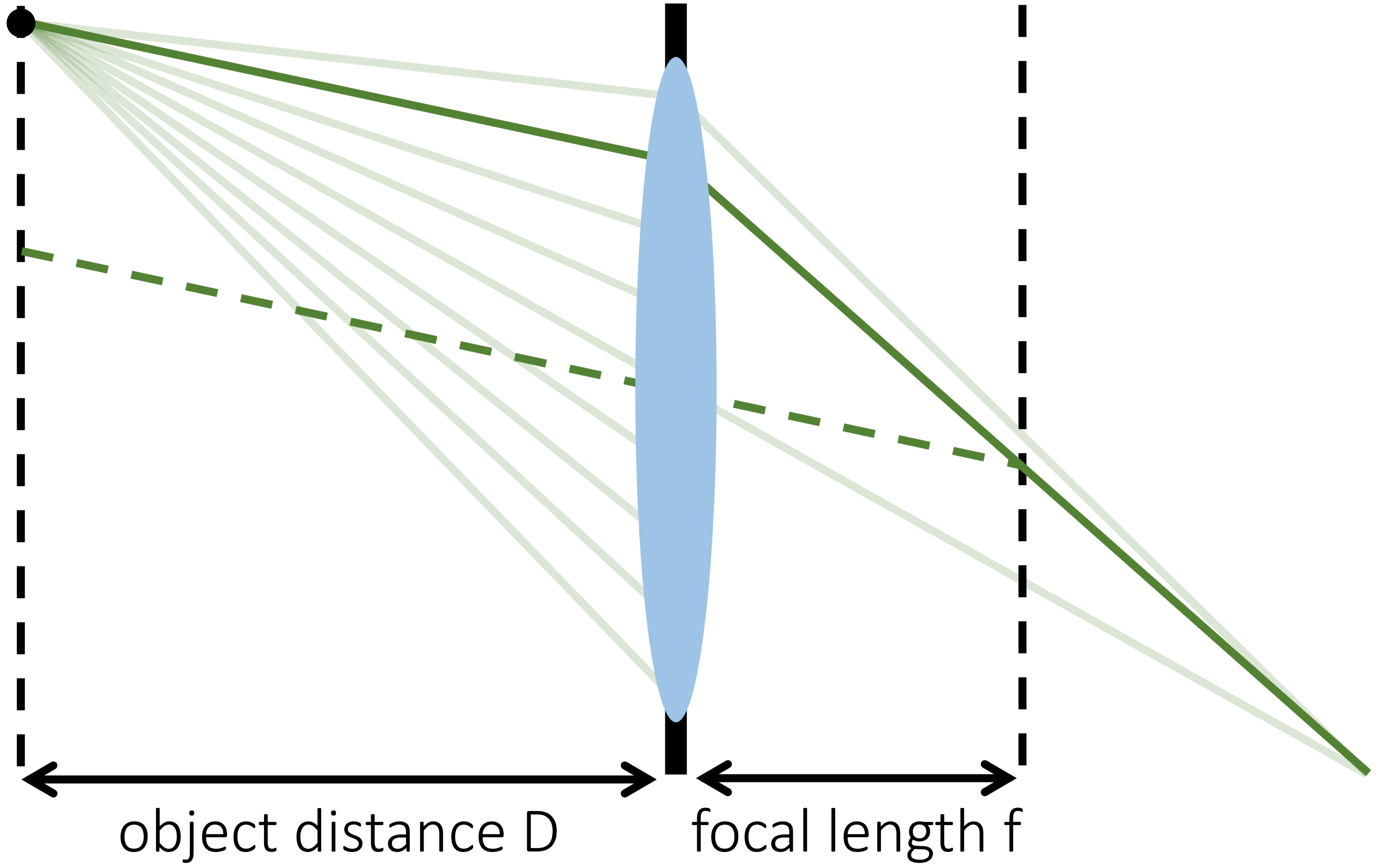

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

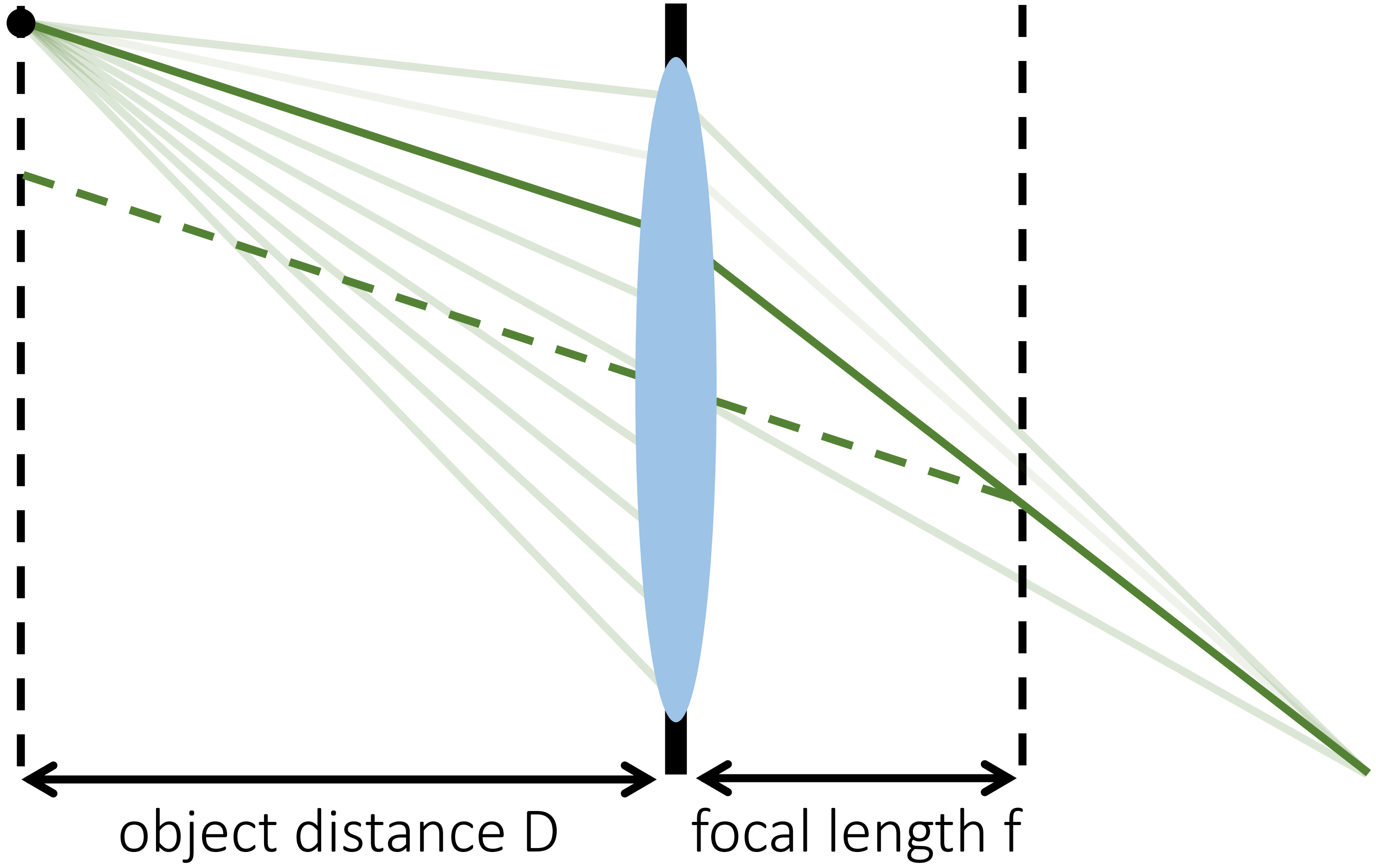

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

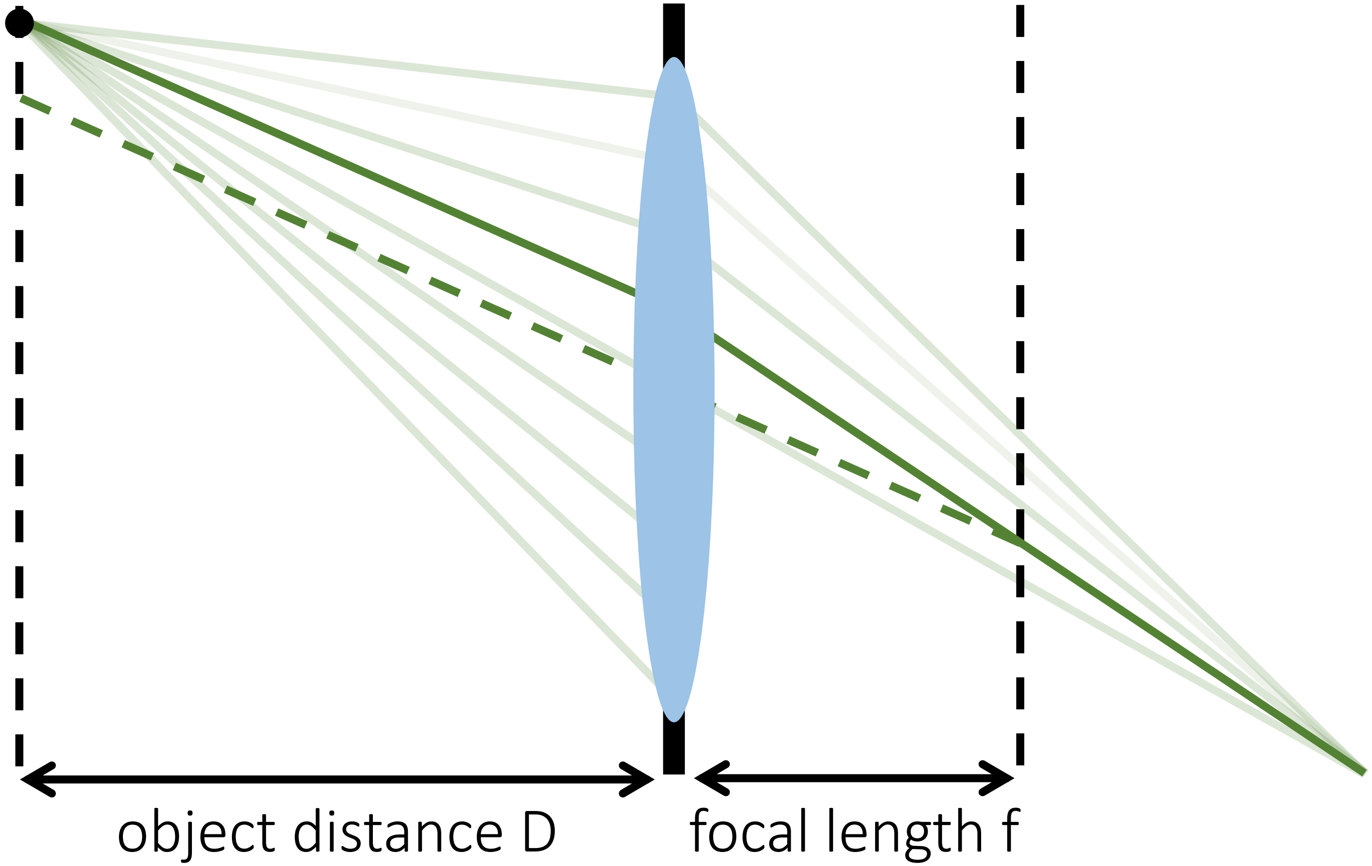

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

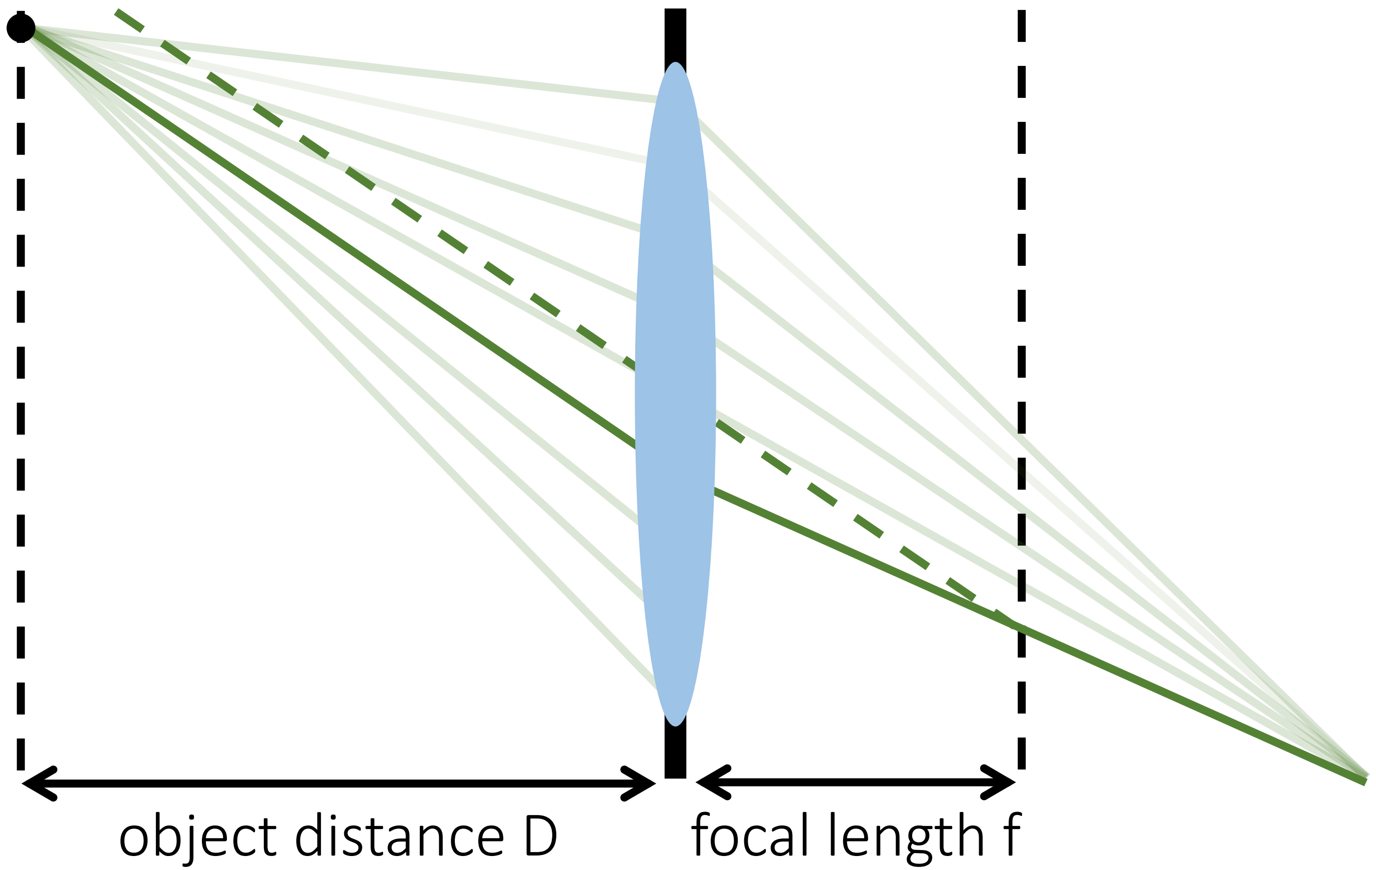

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

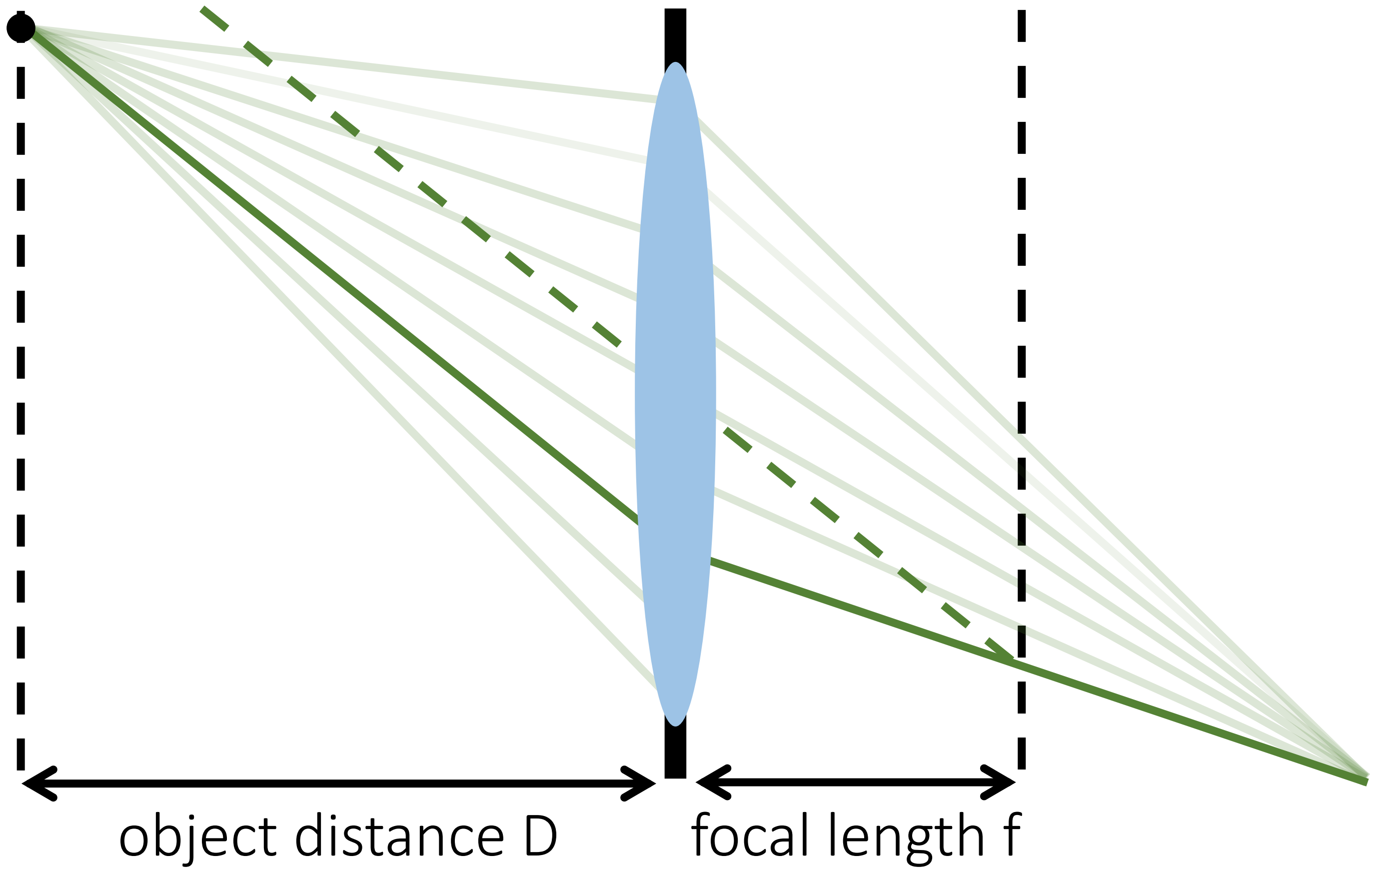

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

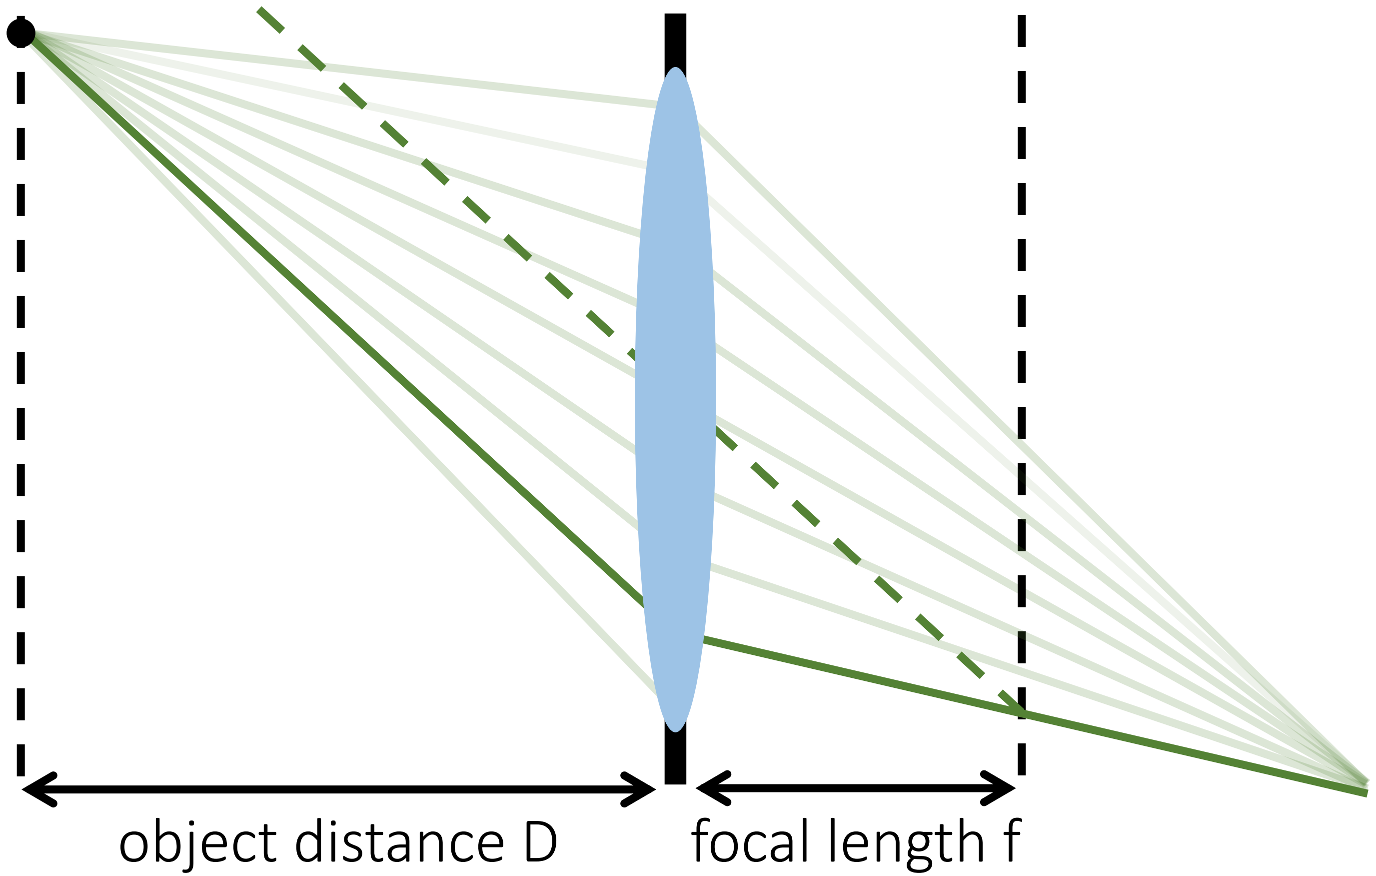

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

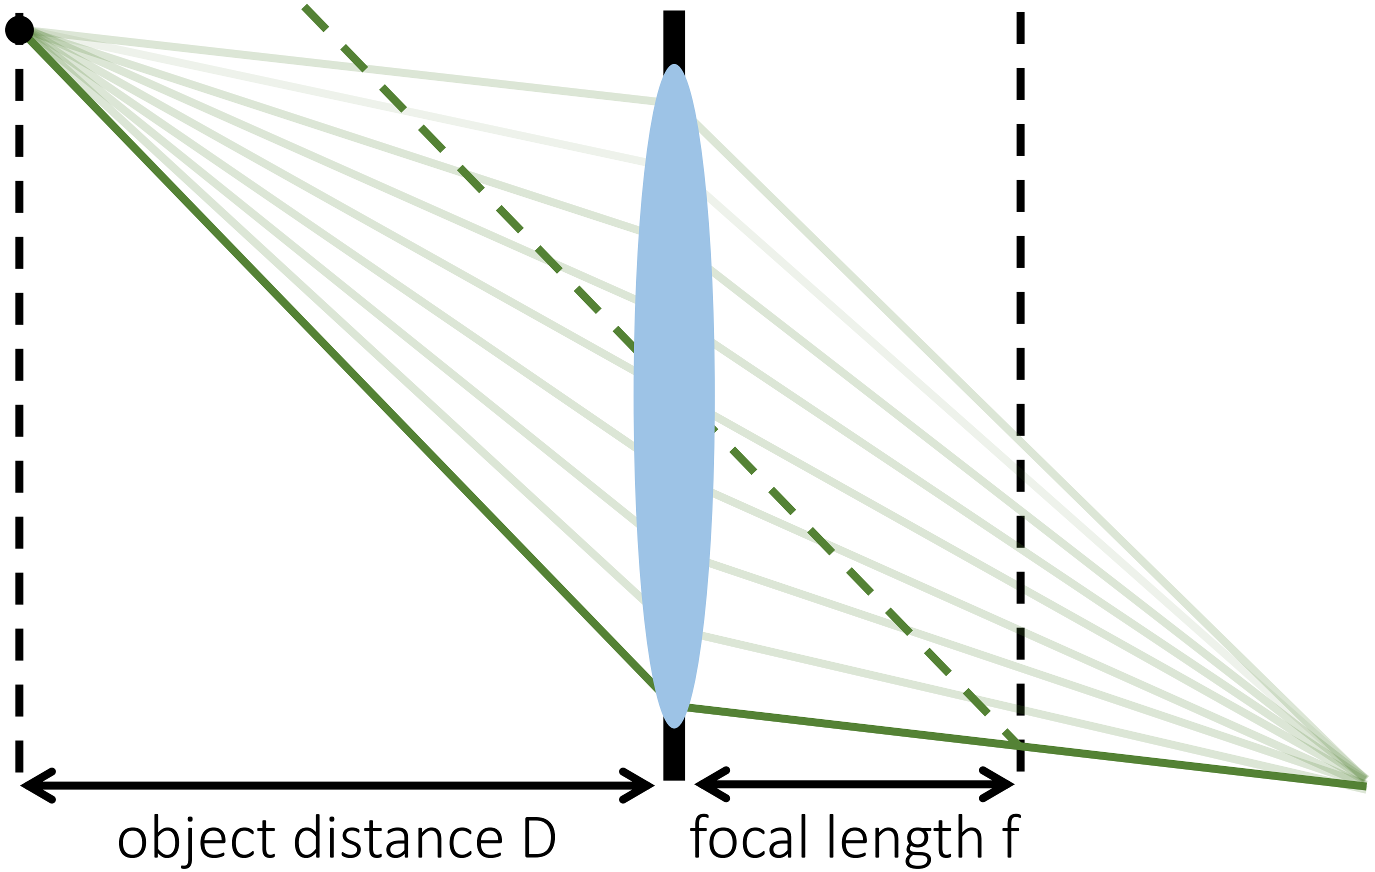

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

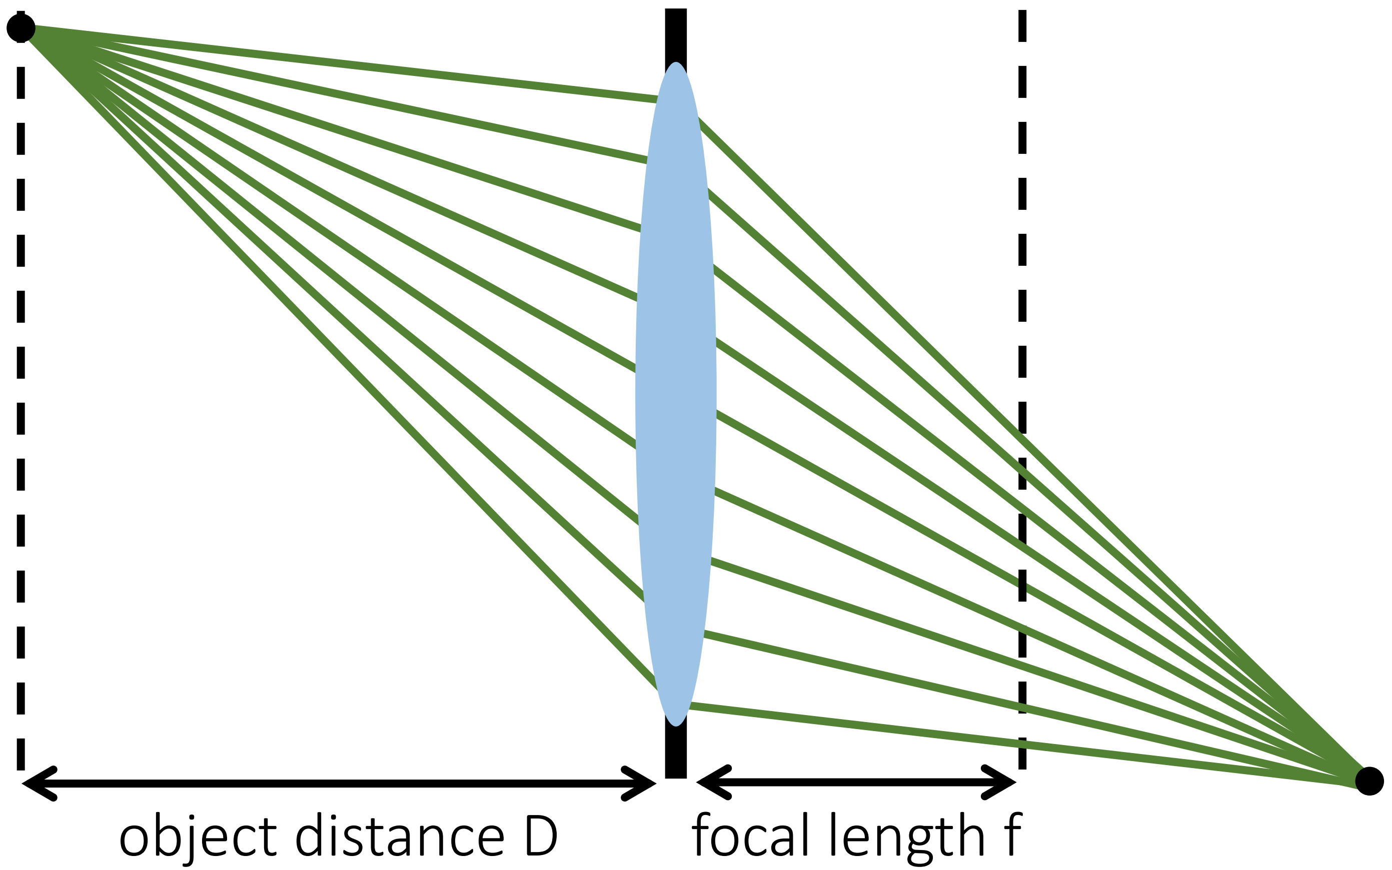

Focusing property:

1. Rays emitted from a point on one side converge to a point on the other side.

Tracing rays through a thin lens

Consider an object emitting a bundle of rays. How do they propagate through the lens?

1. Trace rays through lens center.
2. For all other rays:
   a. Trace their parallel through lens center.
   b. Connect on focal plane.

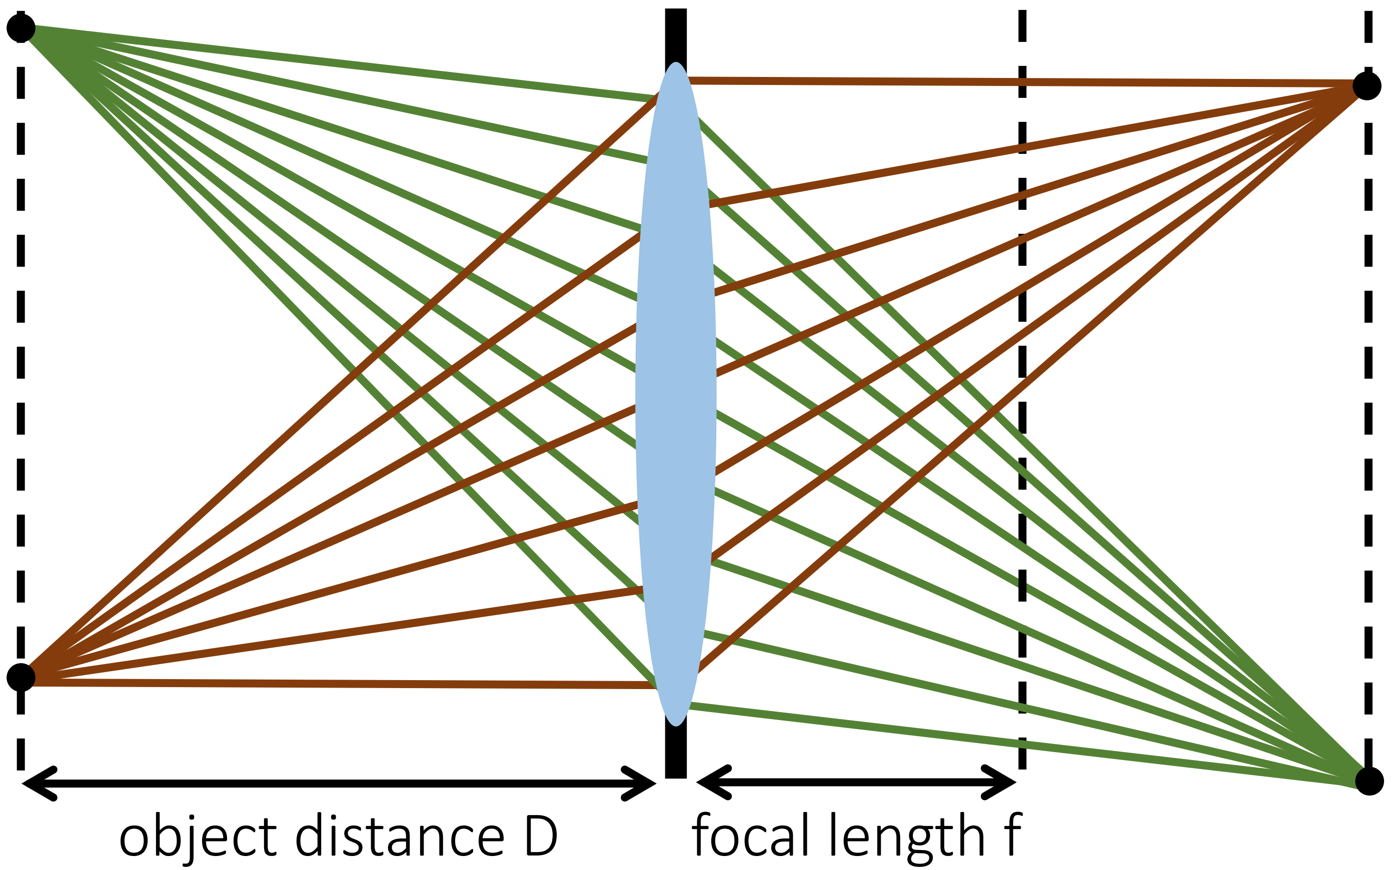

Focusing property:

1. Rays emitted from a point on one side converge to a point on the other side.
2. Bundles emitted from a plane parallel to the lens converge on a common plane.

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

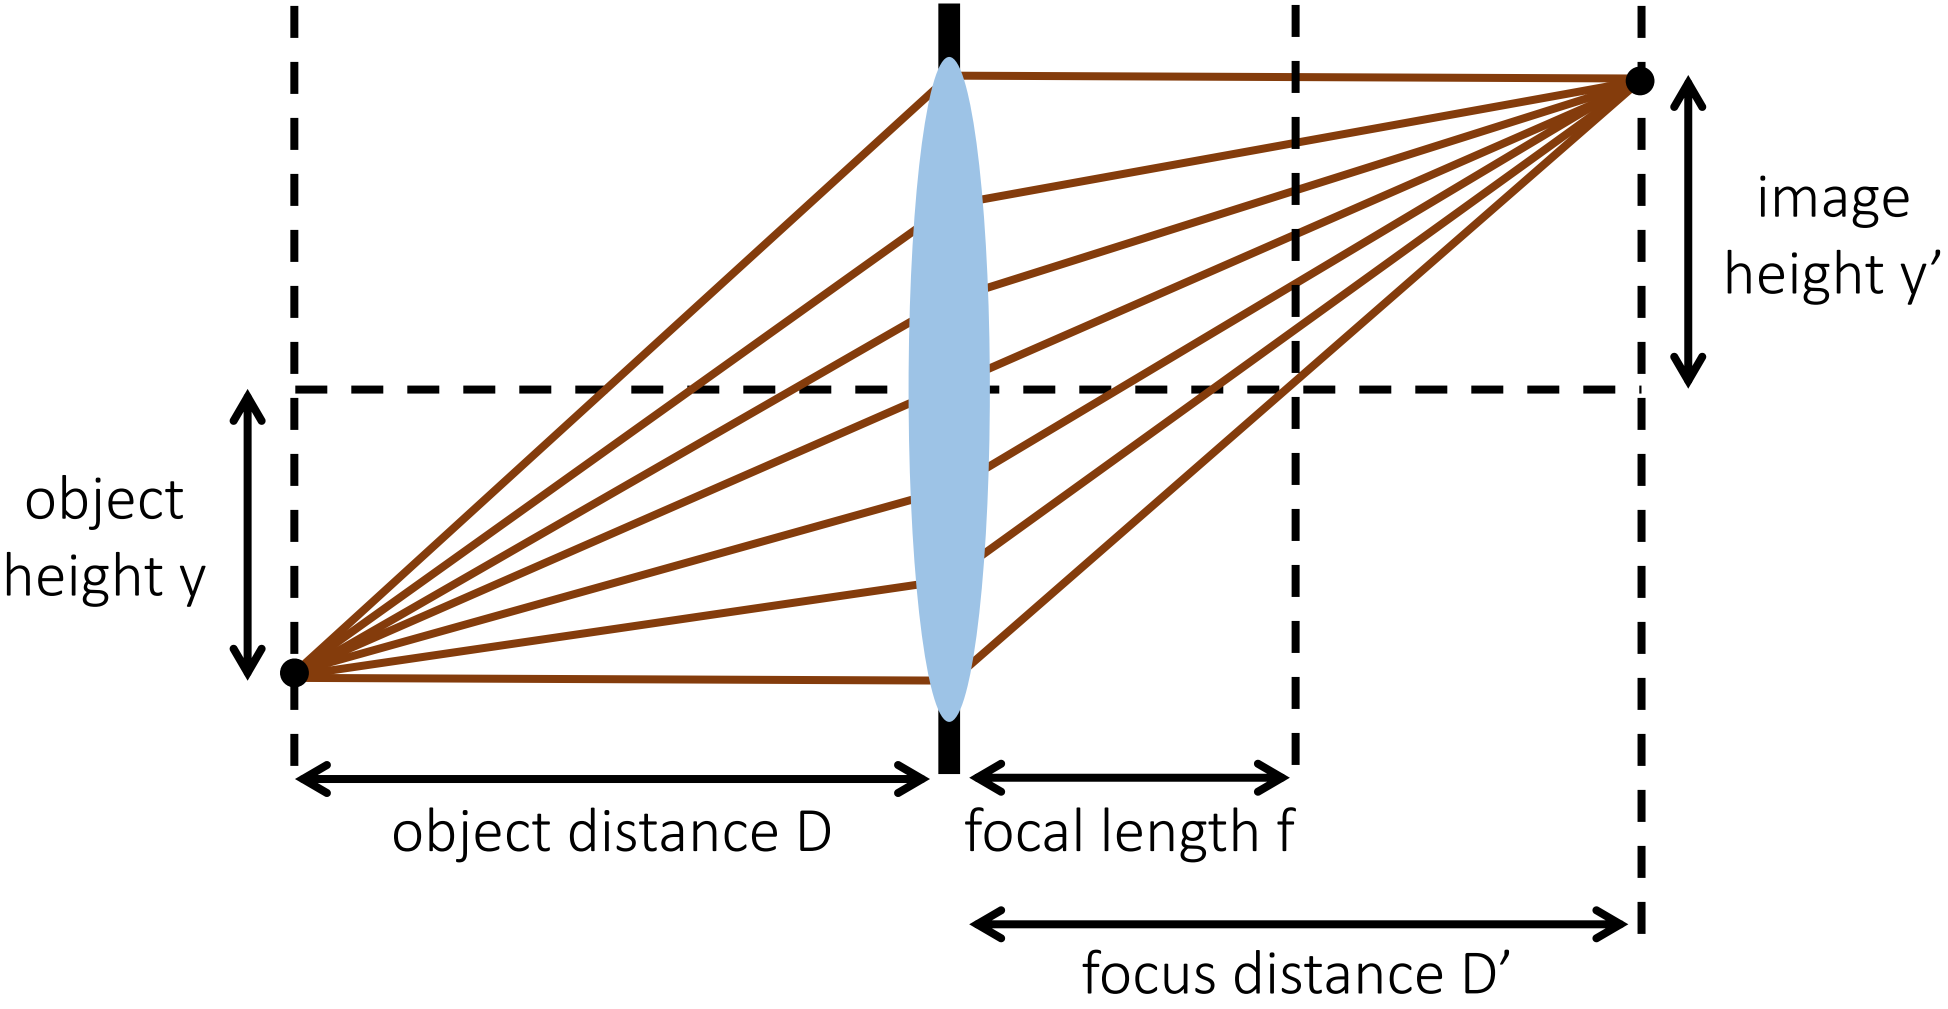

Thin lens formula

How can we relate scene-space (D, y) and image space (D', y') quantities?

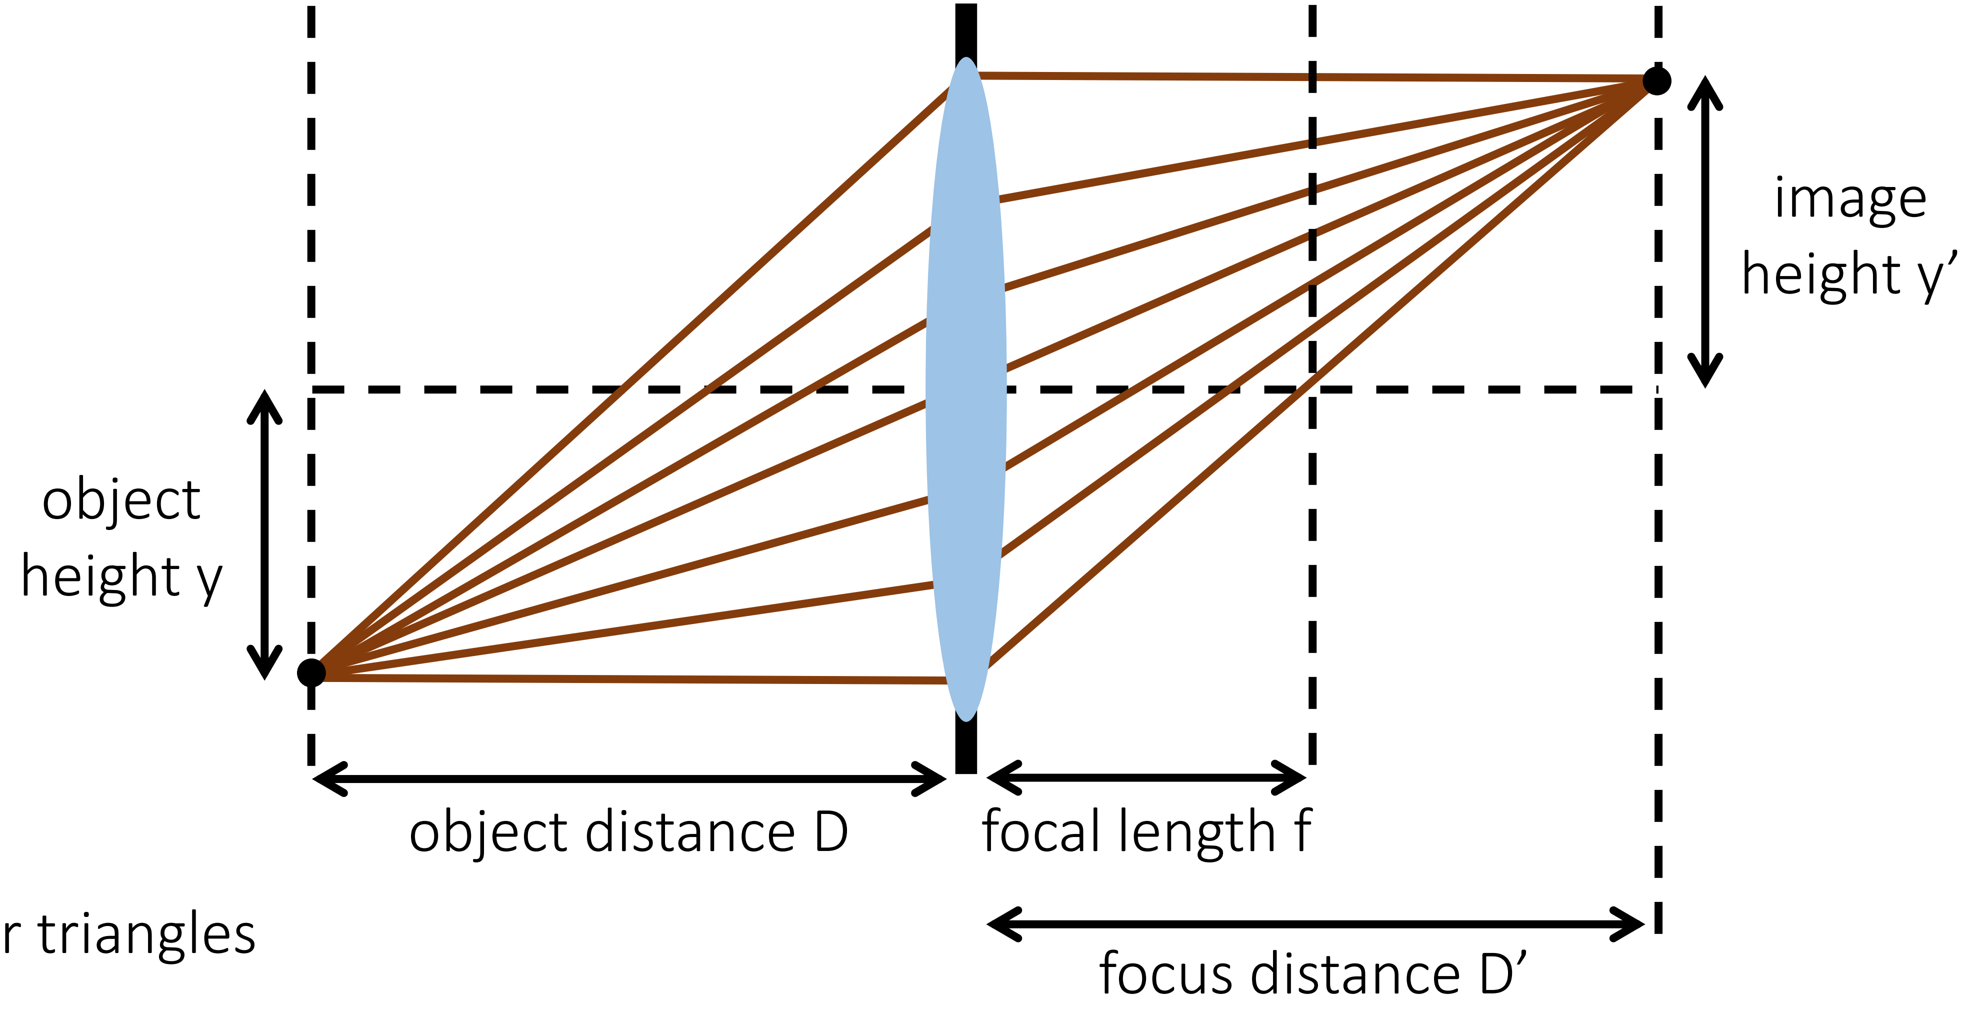

Use similar triangles

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

$$\frac{y}{y'} = ?$$

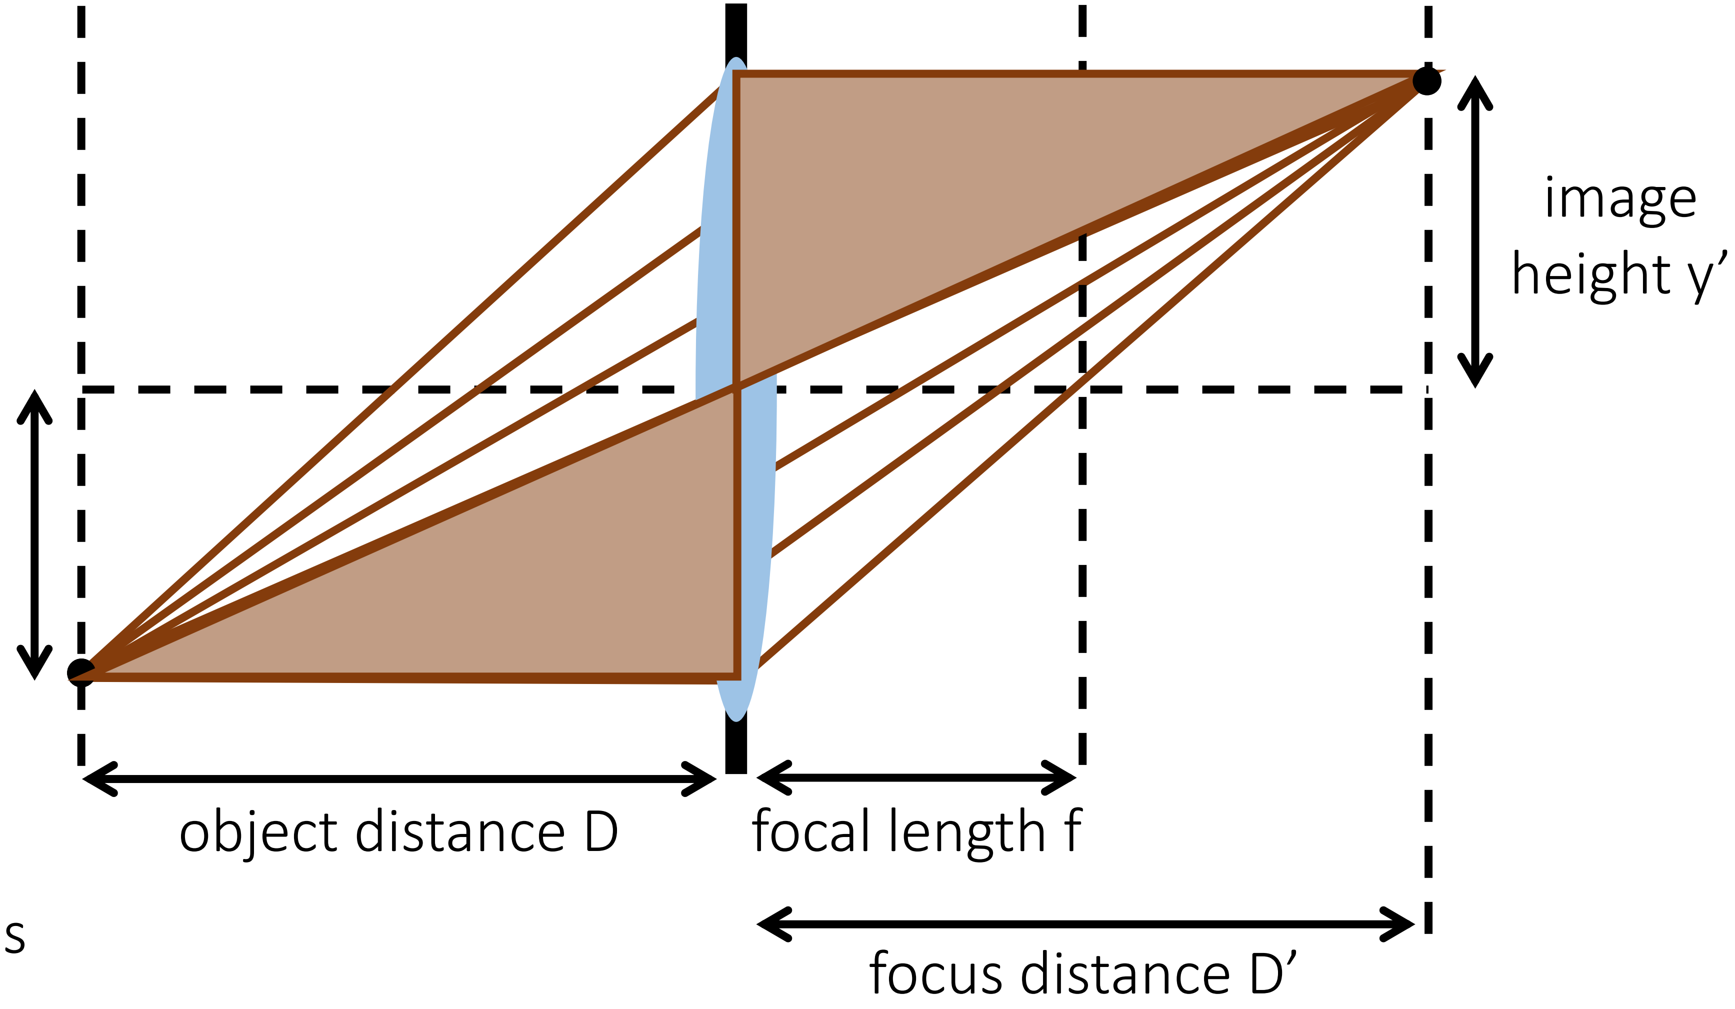

Use similar triangles

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

$$\frac{y}{y'} = \frac{D}{D'}$$

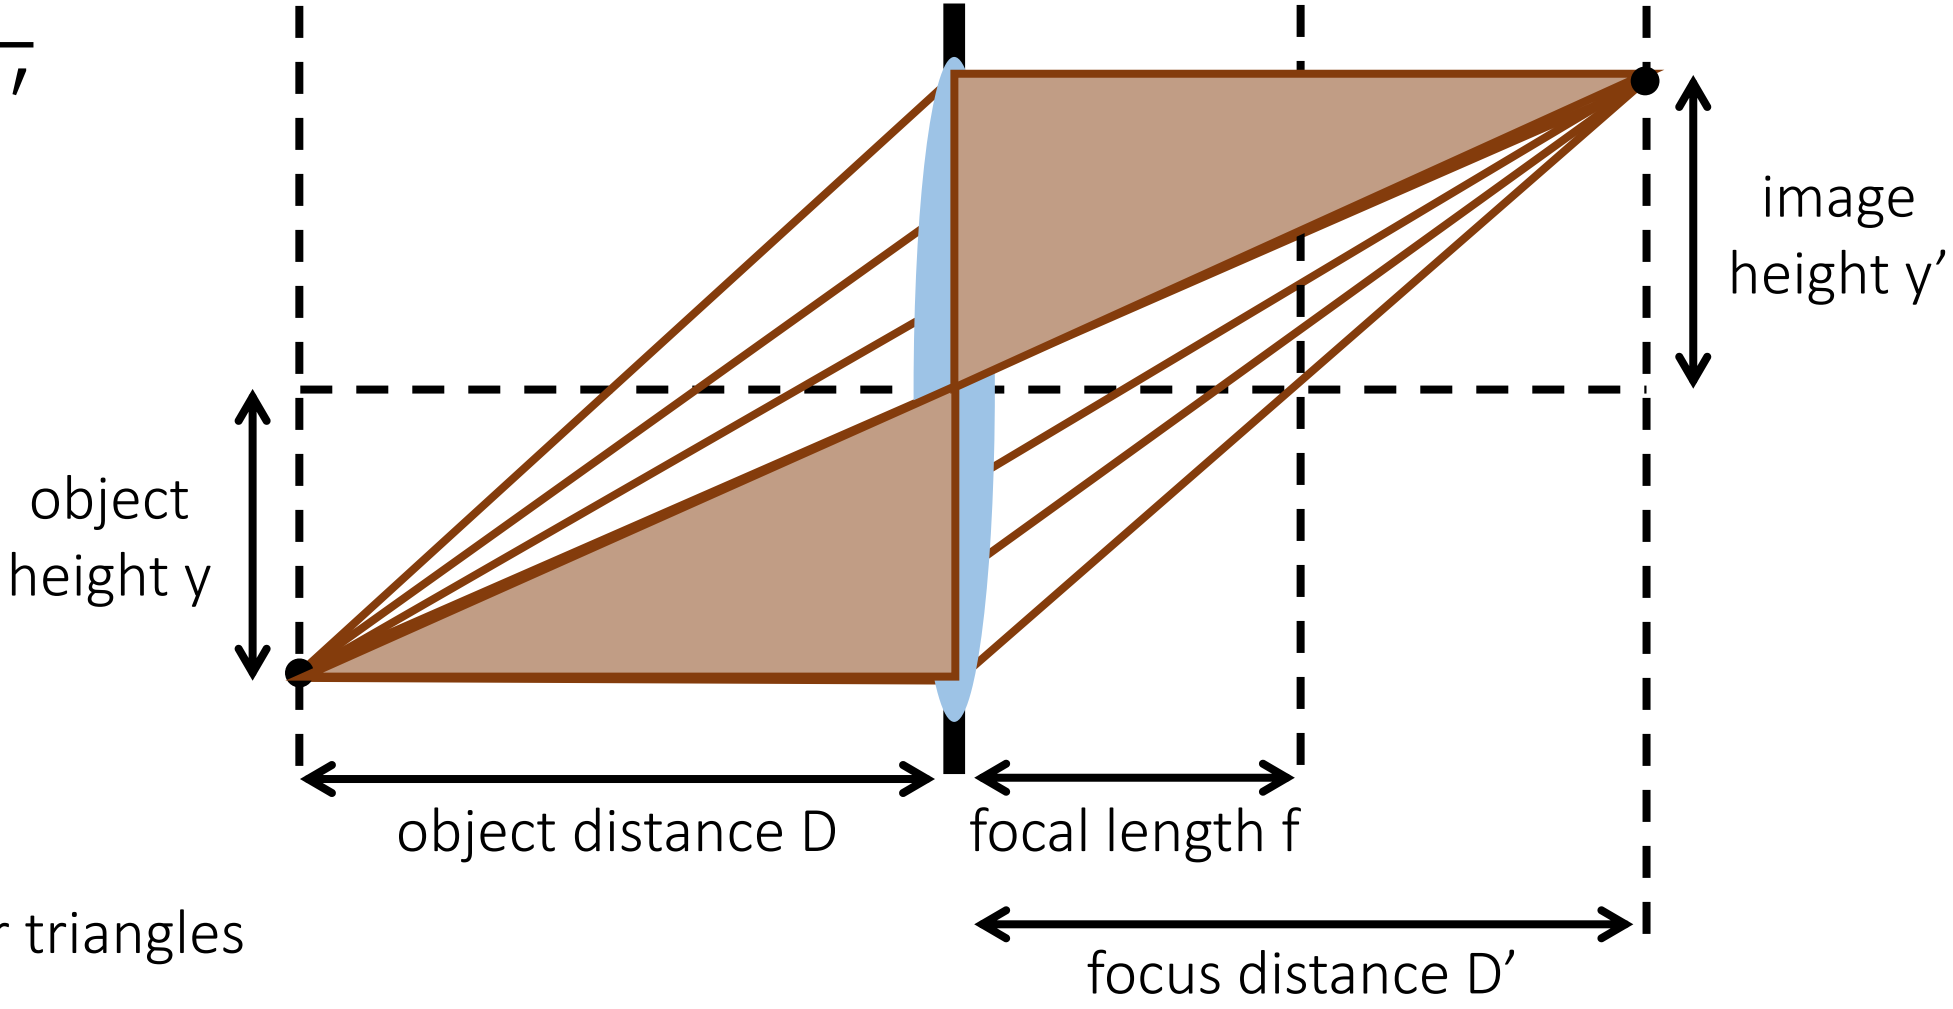

Use similar triangles

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

$$\frac{y}{y'} = \frac{D}{D'}$$

$$\frac{y}{y'} = ?$$

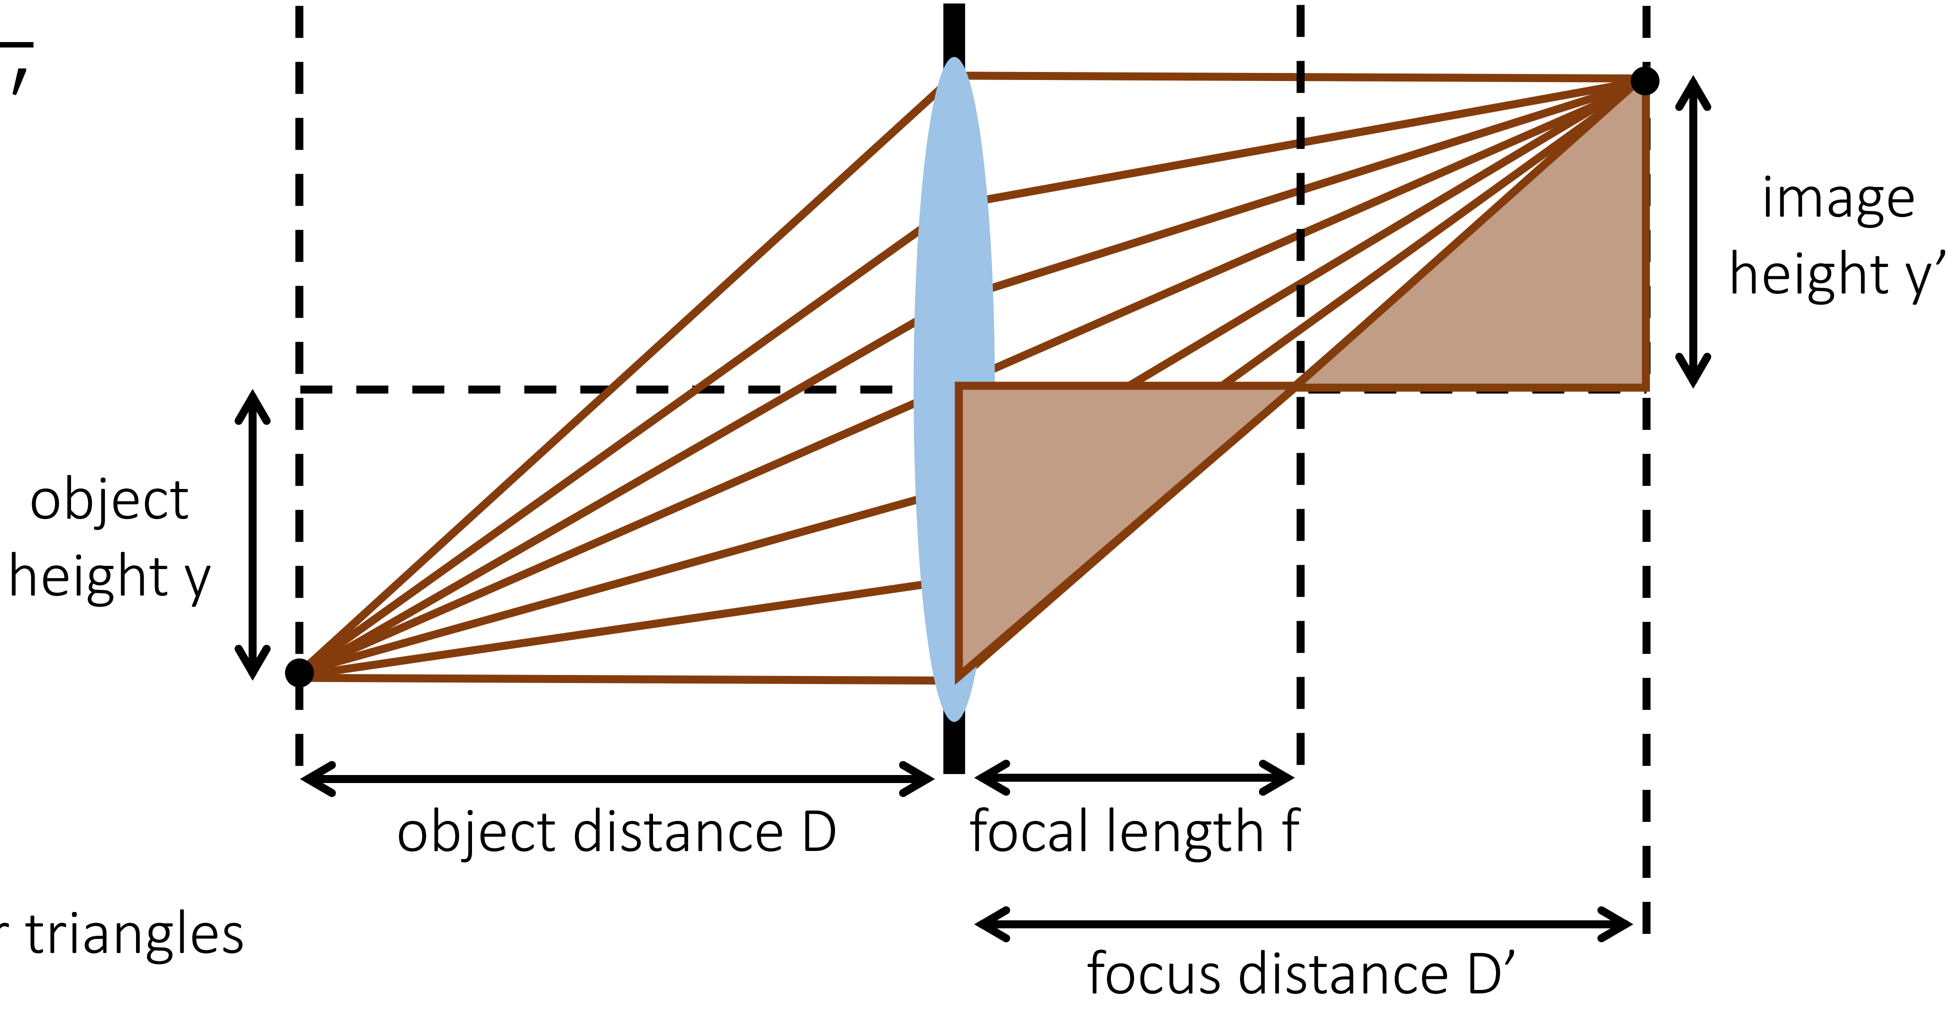

Use similar triangles

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

$$\frac{y}{y'} = \frac{D}{D'}$$

$$\frac{y}{y'} = \frac{f}{D' - f}$$

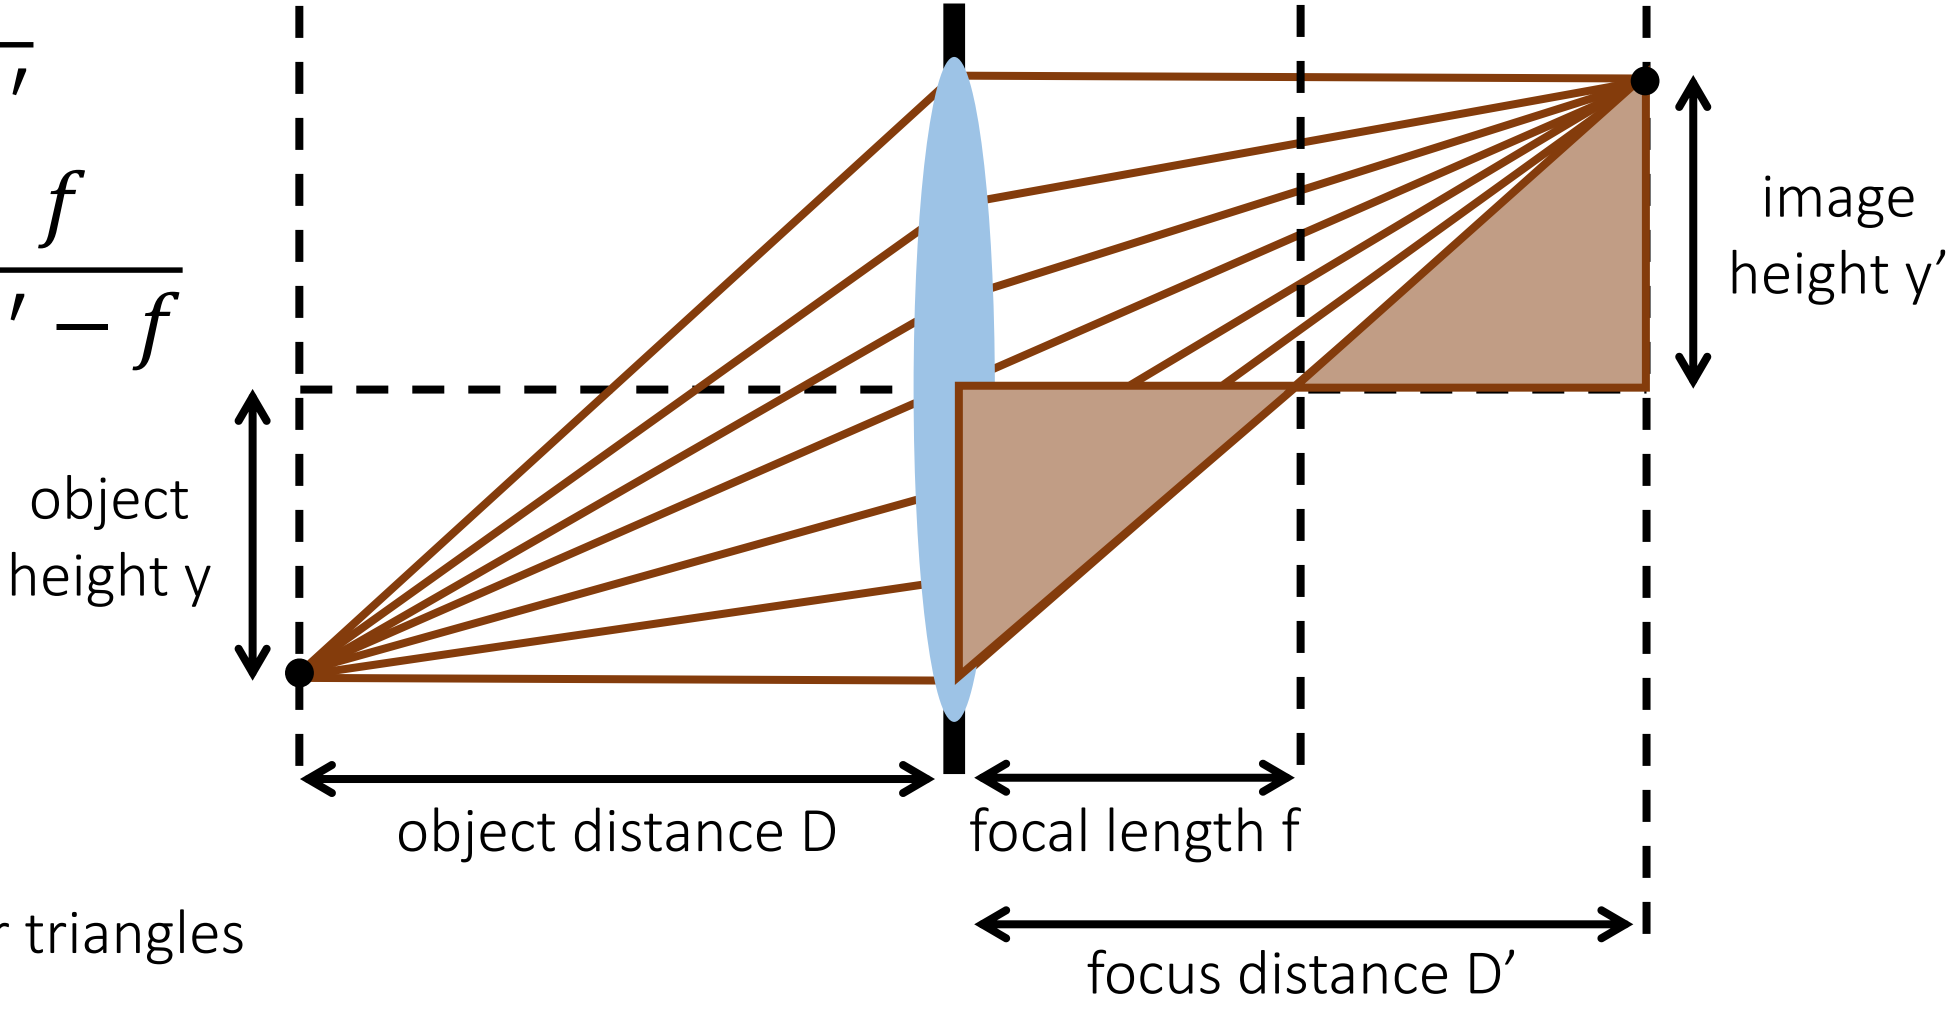

Use similar triangles

Thin lens formula

How can we relate scene-space (D, y) and image space (D’, y’) quantities?

$$m = \frac{f}{D' - f}$$

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

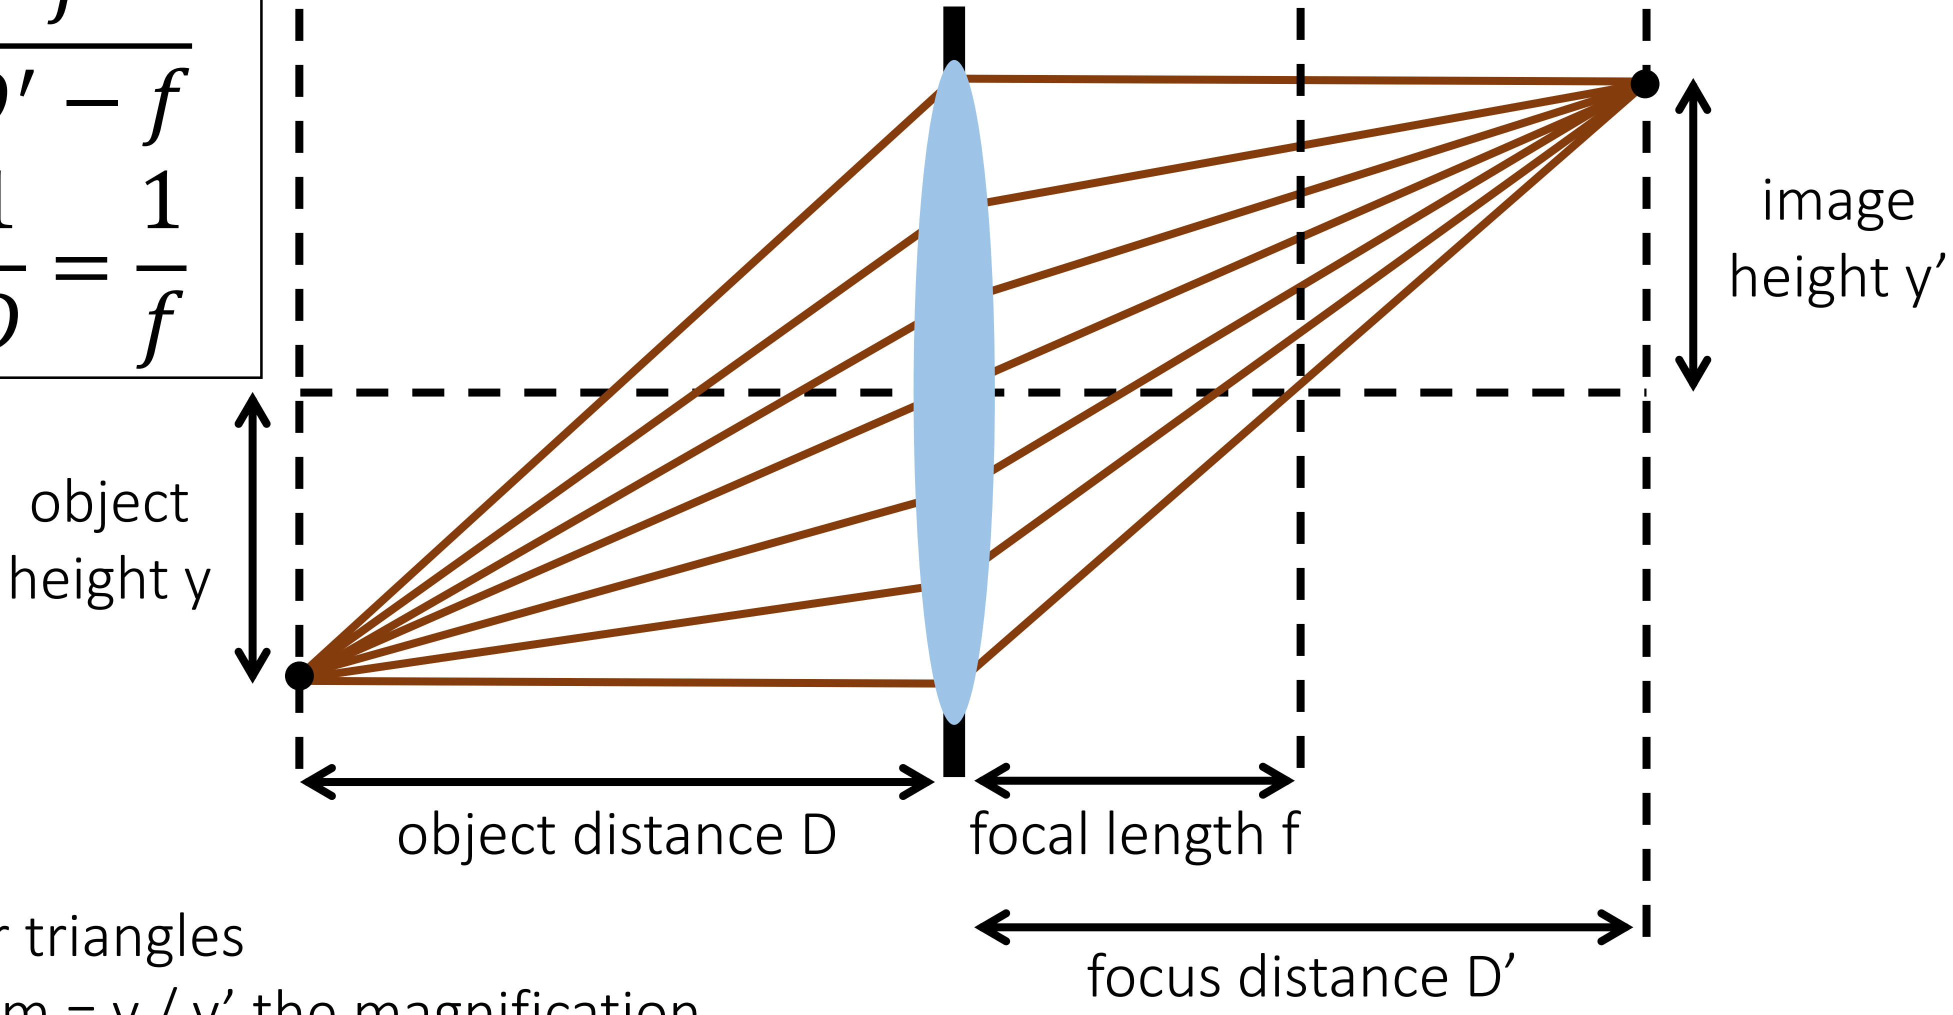

Use similar triangles

- We call m = y / y’ the magnification

Special focus distances

D’ = f, D = ?, m = ?

$$m = \frac{f}{D' - f}$$

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

Special focus distances

$$m = \frac{f}{D' - f}$$

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

D’ = f, D = ∞, m = ∞ → infinity focus (parallel rays)

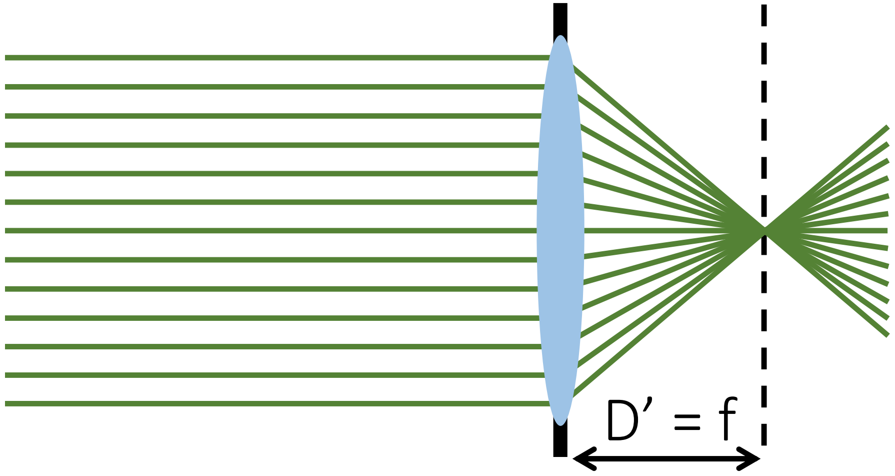

D’ = D = ?, m = ?

Special focus distances

$$m = \frac{f}{D' - f}$$

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

D’ = f, D = ∞, m = ∞ → infinity focus (parallel rays)

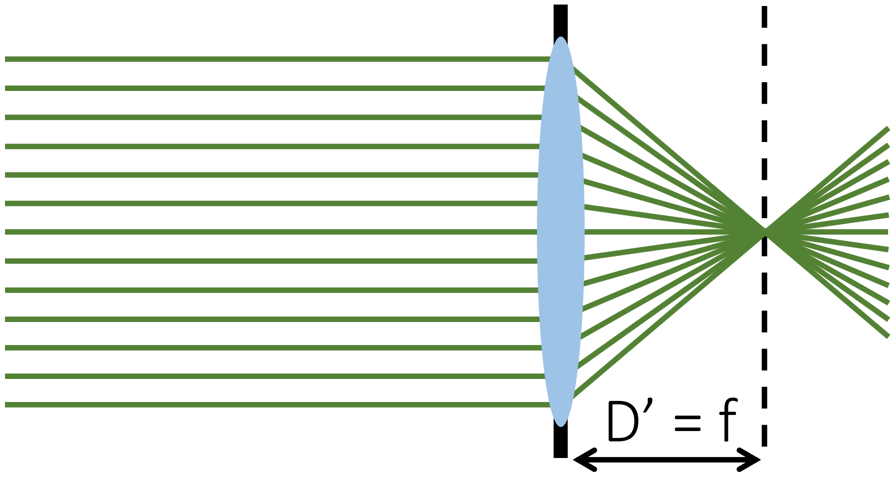

D’ = D = 2 f, m = 1 → object is reproduced in real-life size

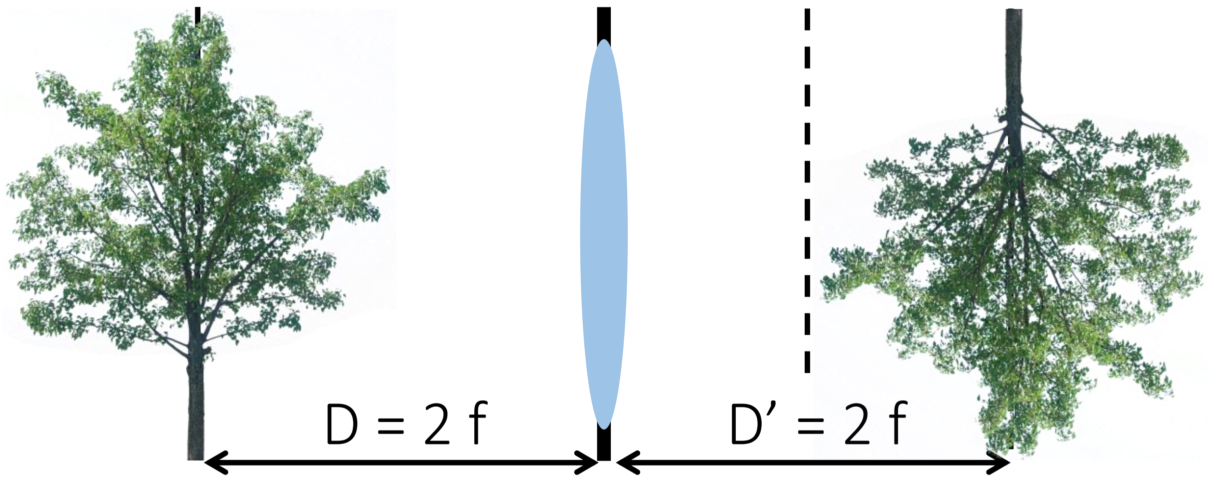

Free lunch?

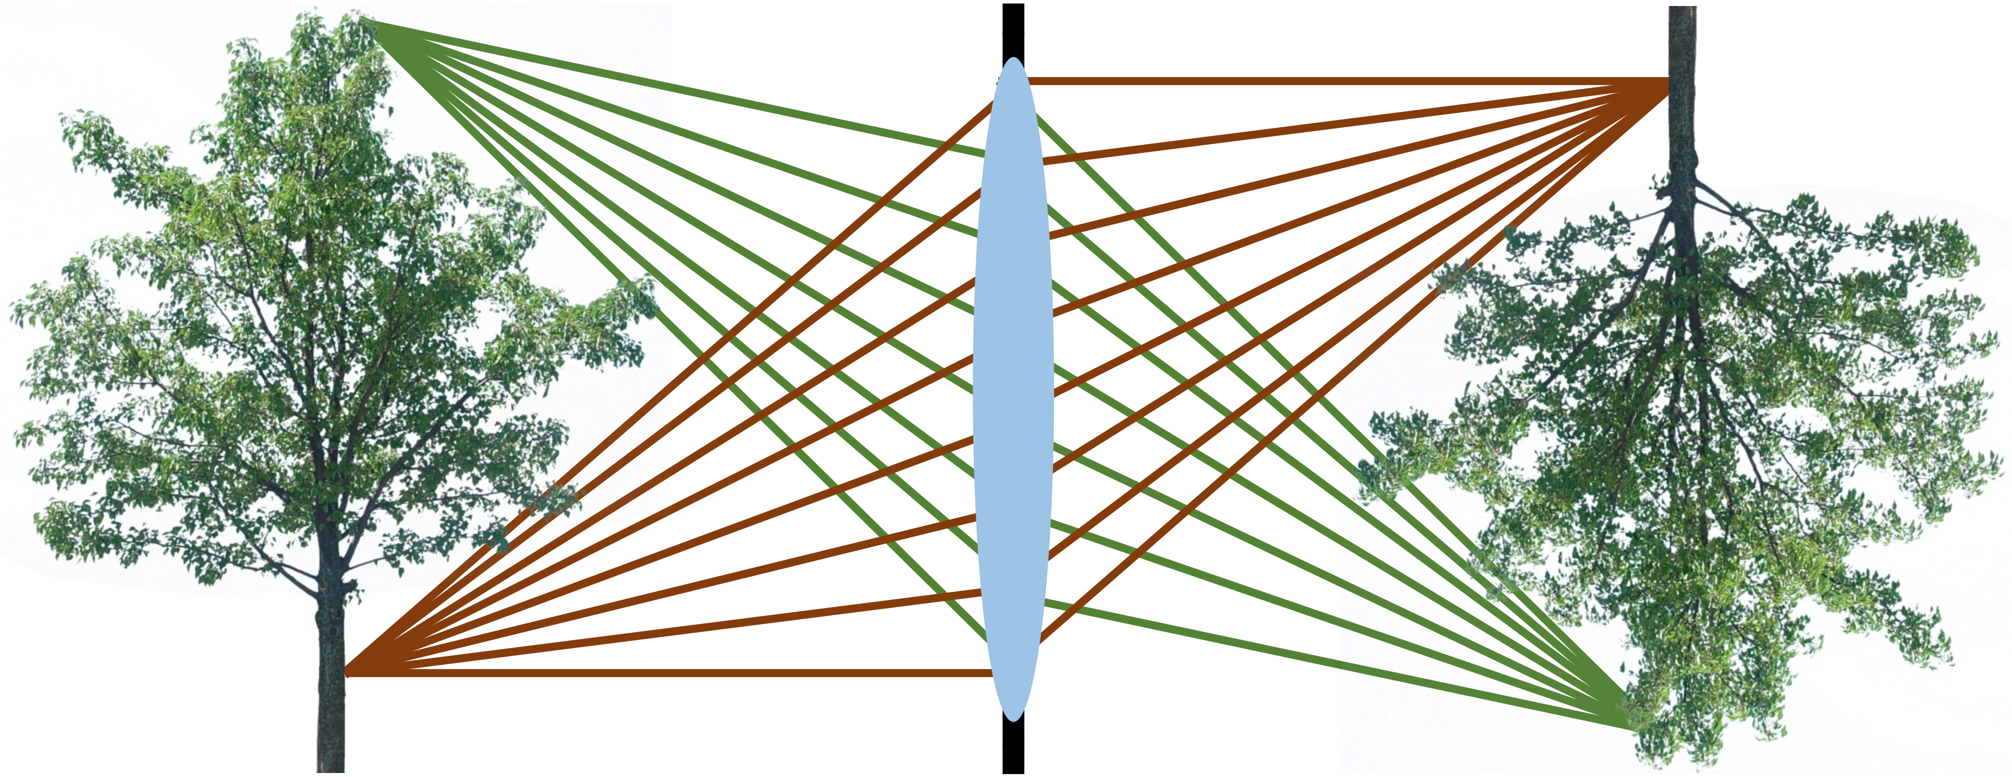

By using a lens we simultaneously achieve:

1. Image is sharp.
2. Signal-to-noise ratio is high.

Do we lose anything by using a lens?

Defocus

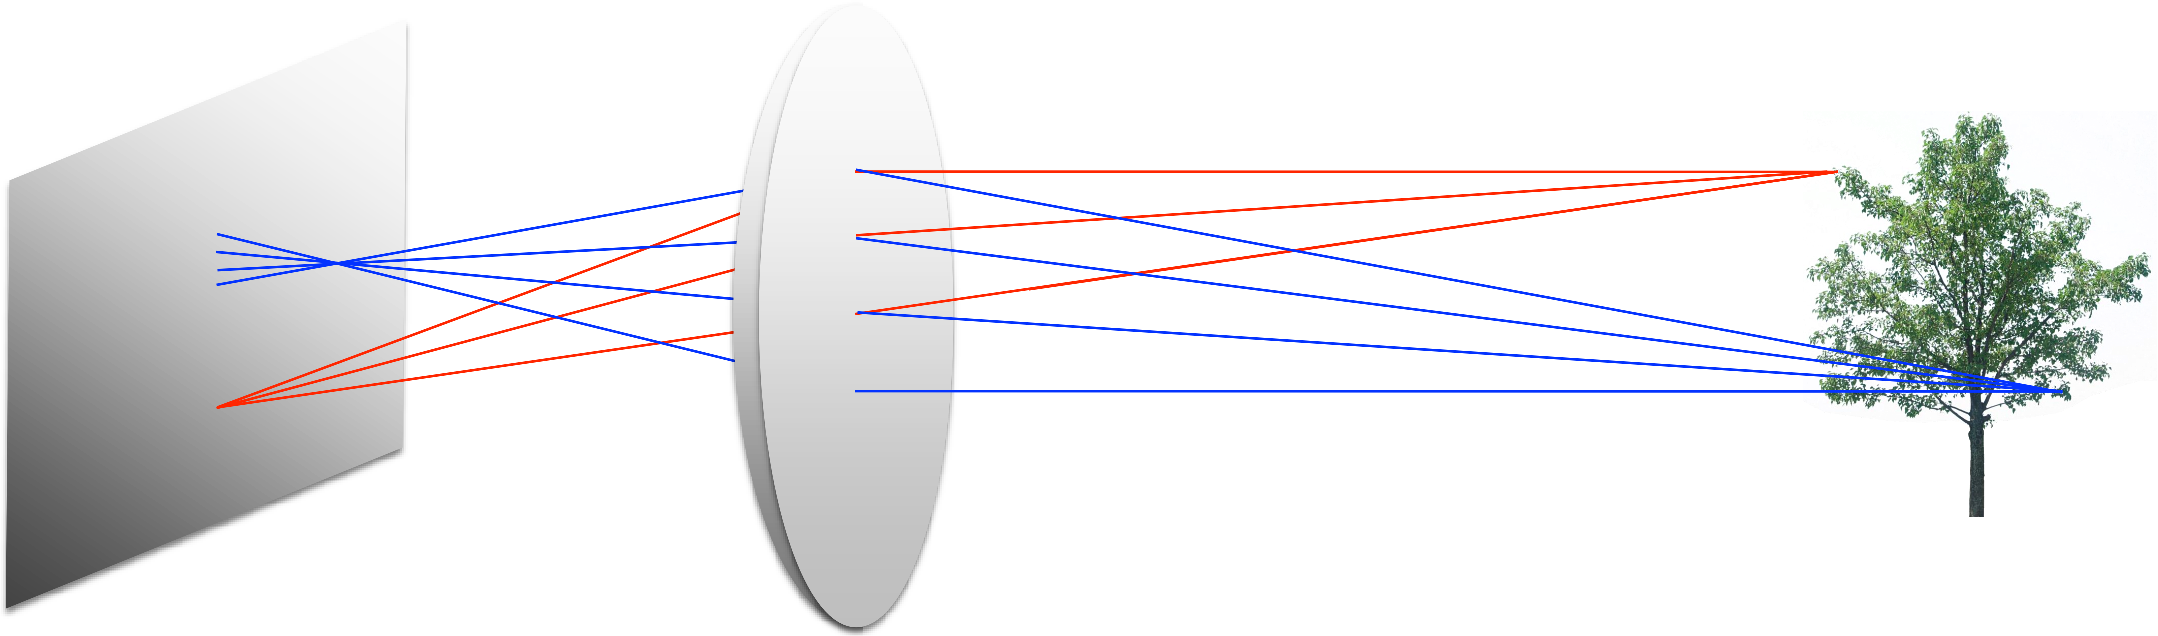

Defocus

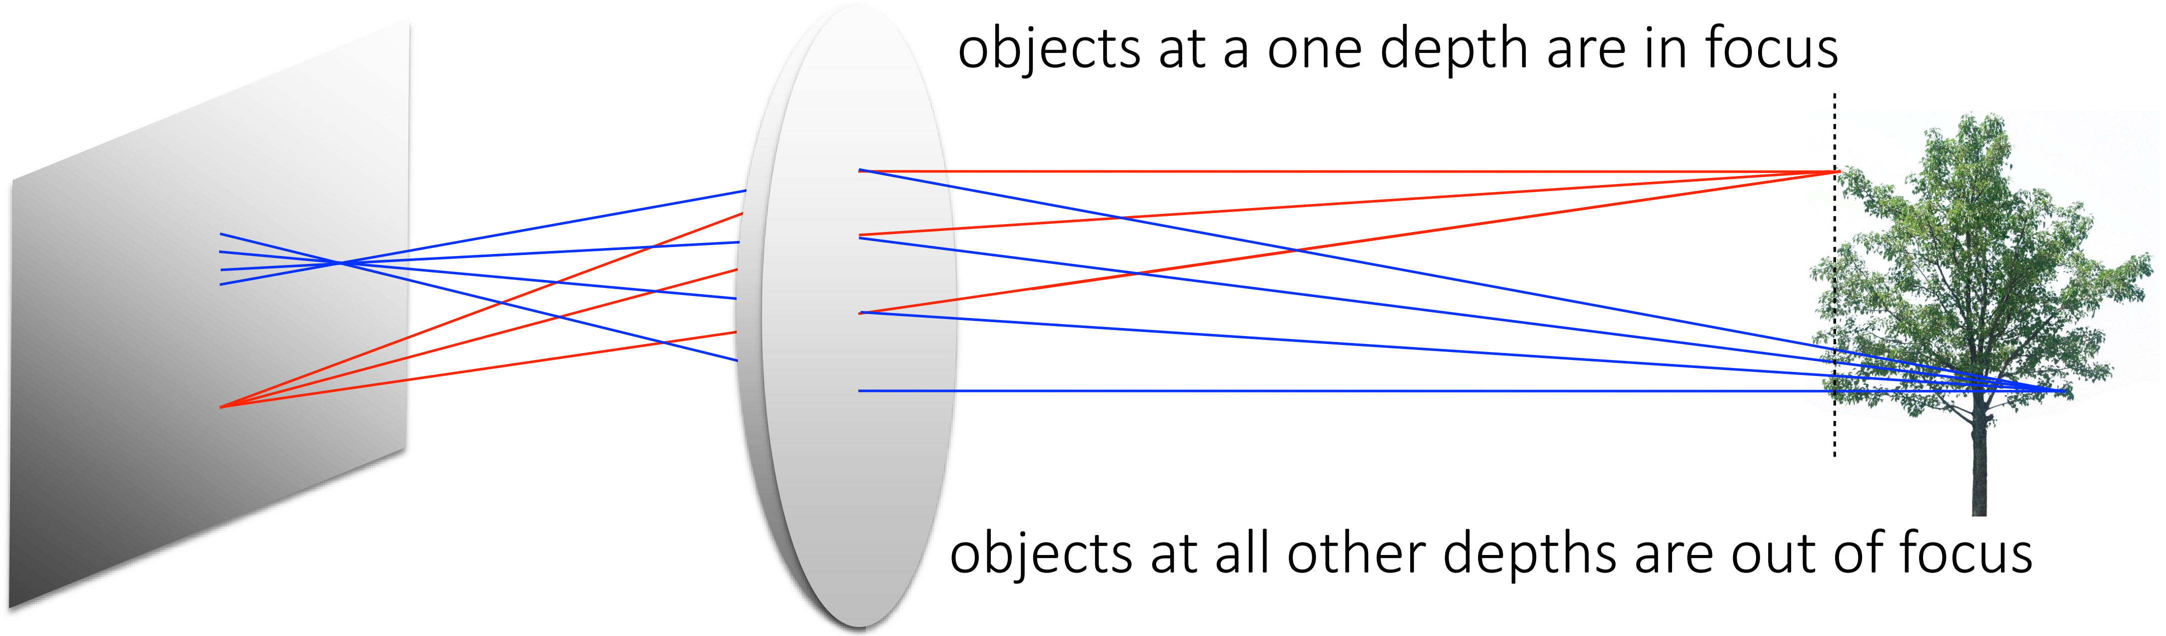
objects at a one depth are in focus
objects at all other depths are out of focus

Defocus

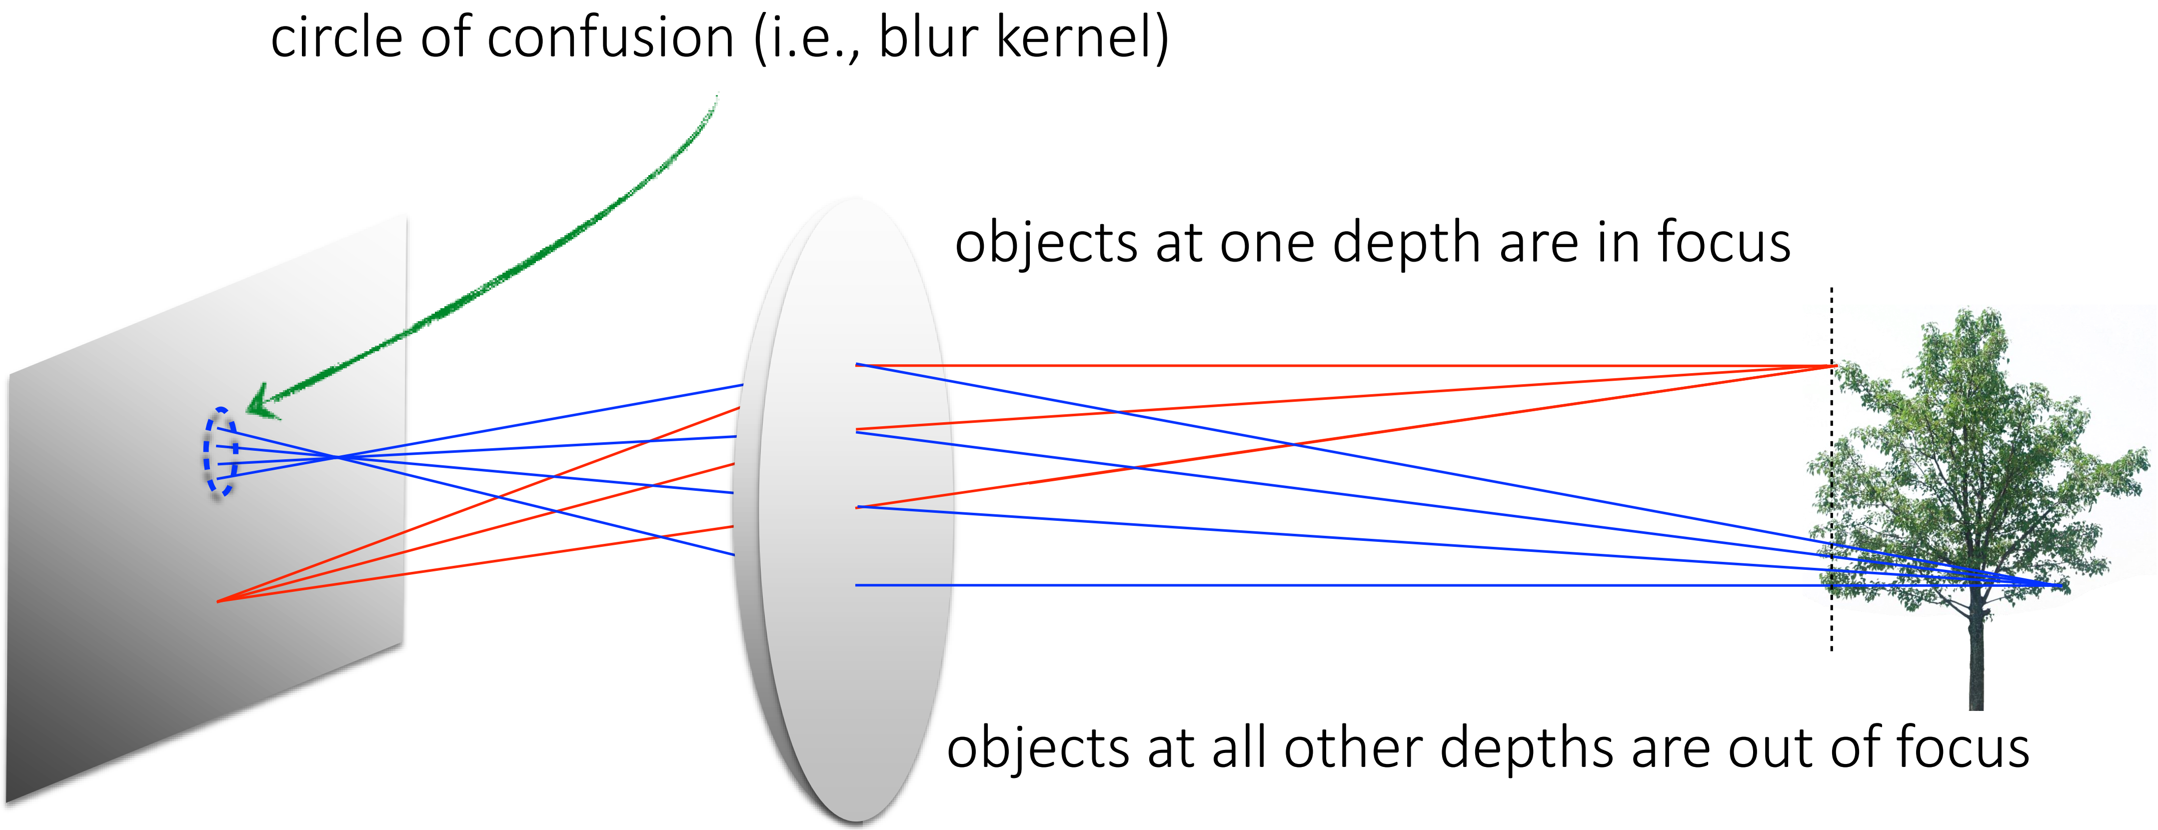

Is the circle of confusion constant?

Defocus

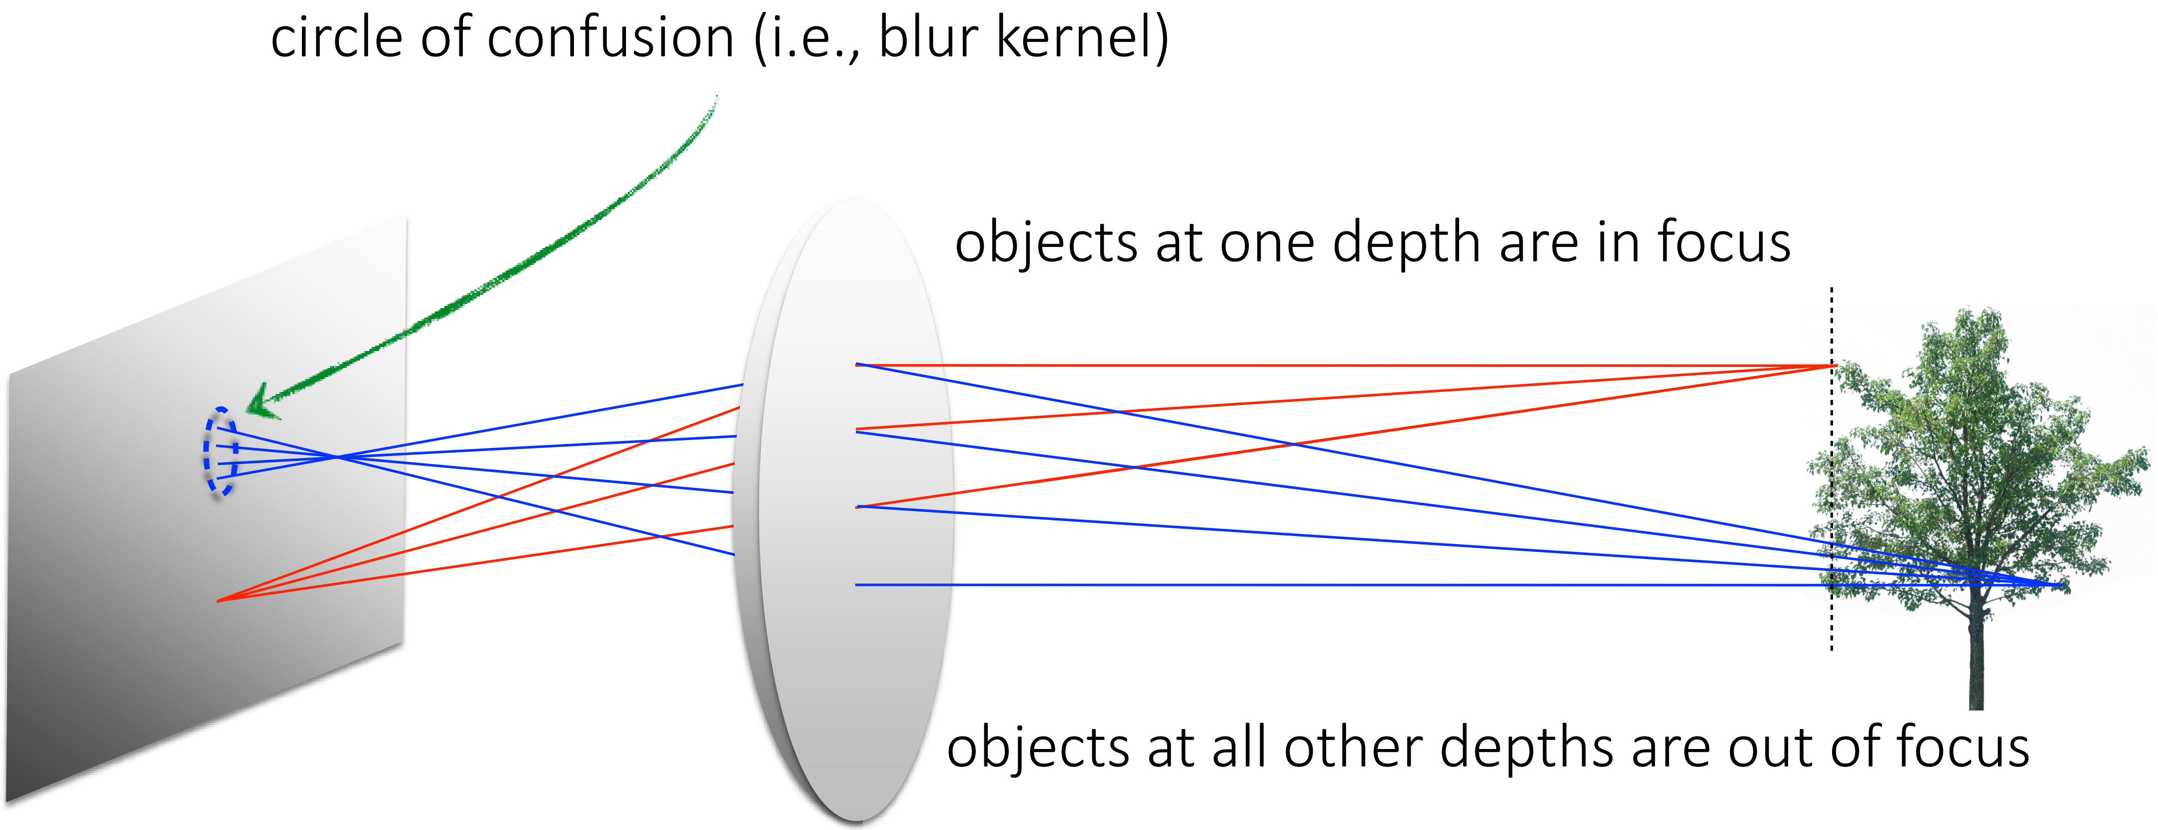

How do we change the depth where objects are in focus?

Change of focus for different depths

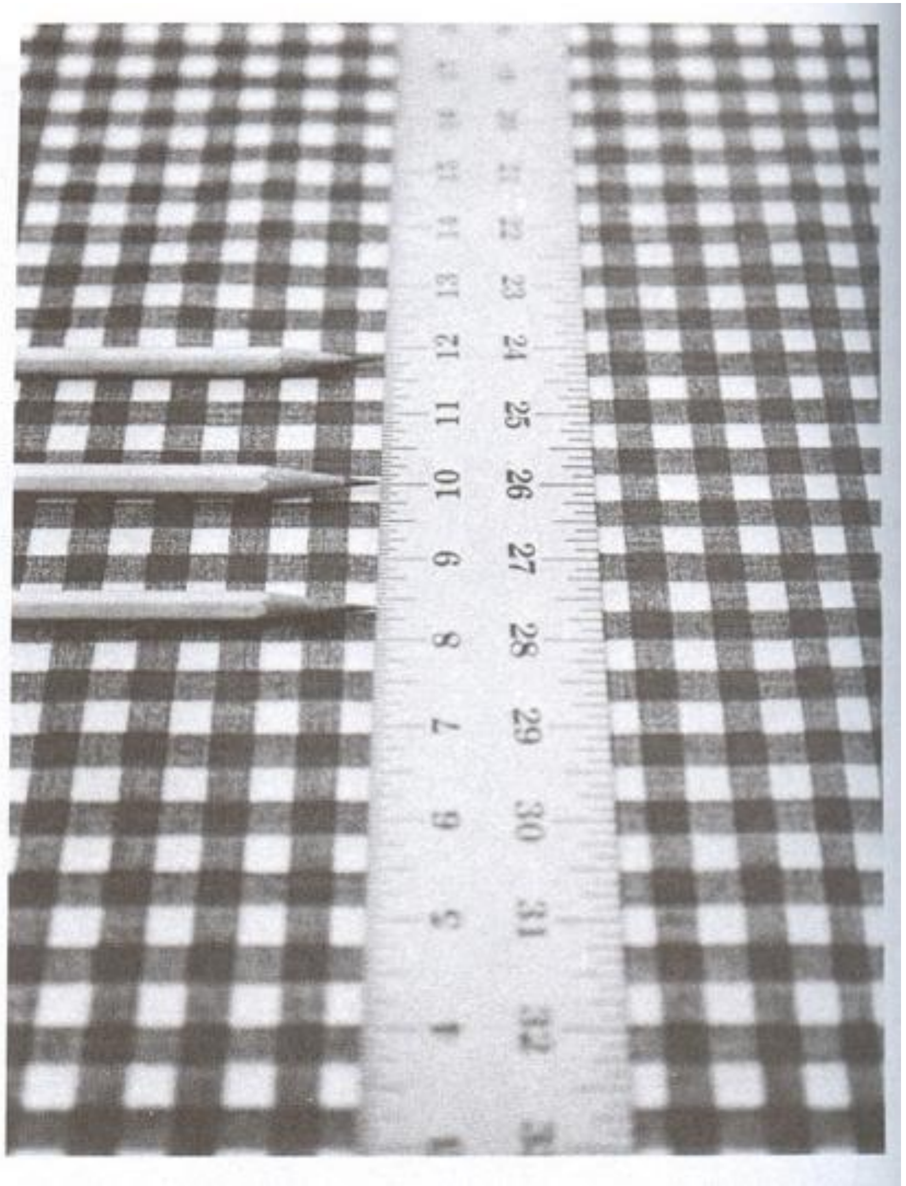

Lens camera and pinhole camera

The lens camera

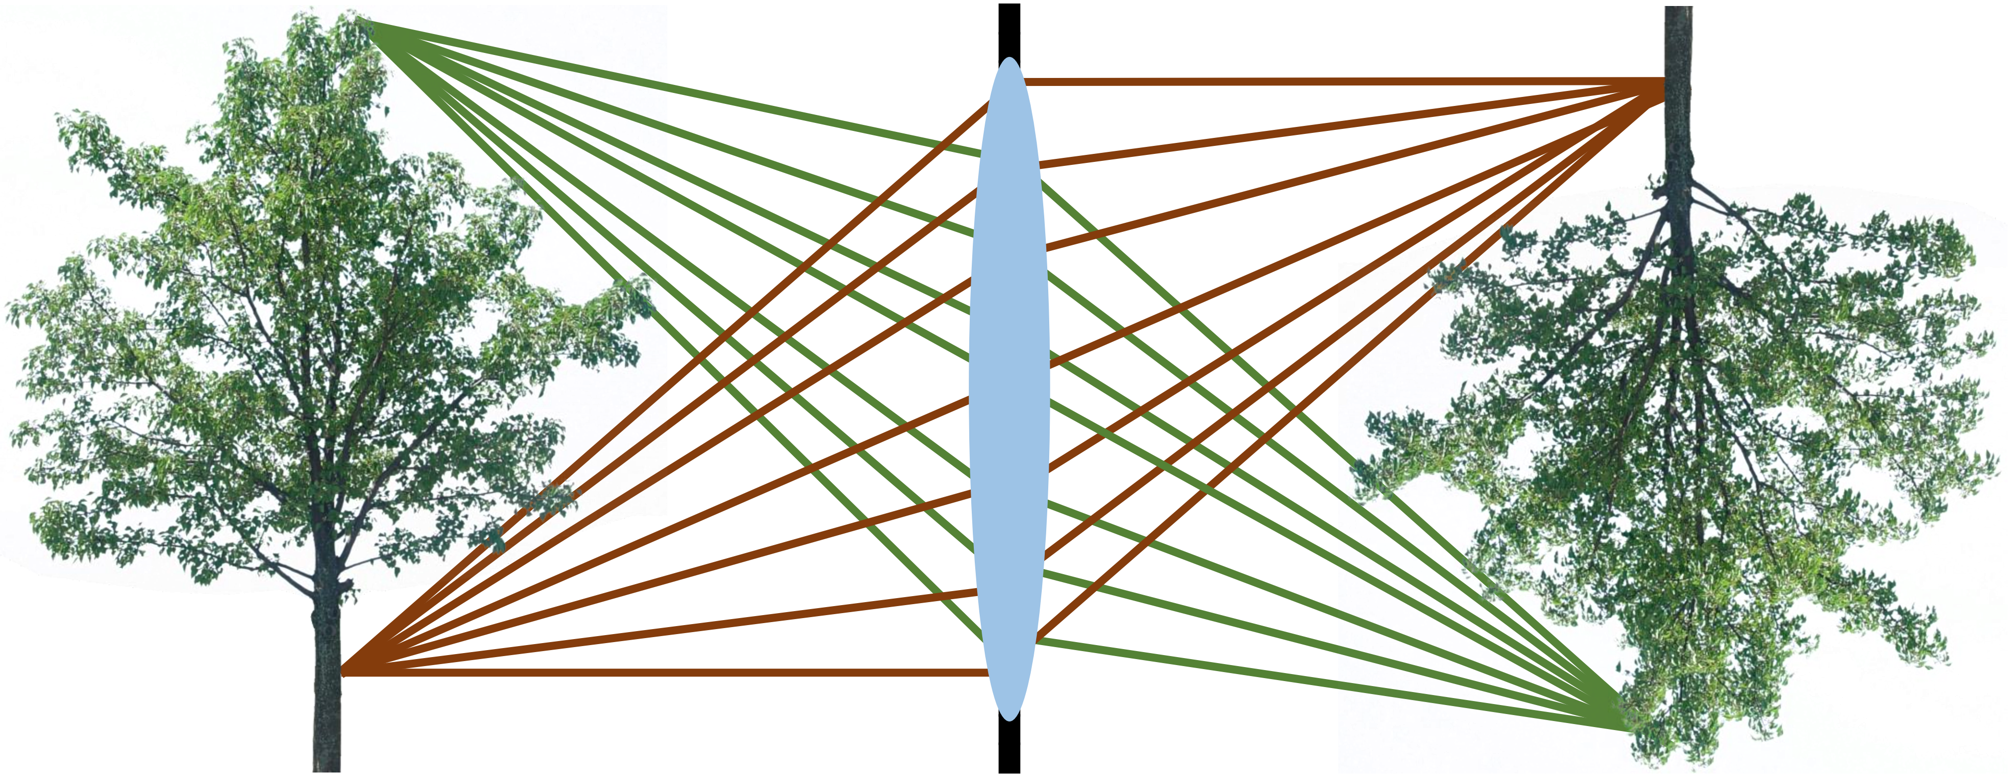

The pinhole camera

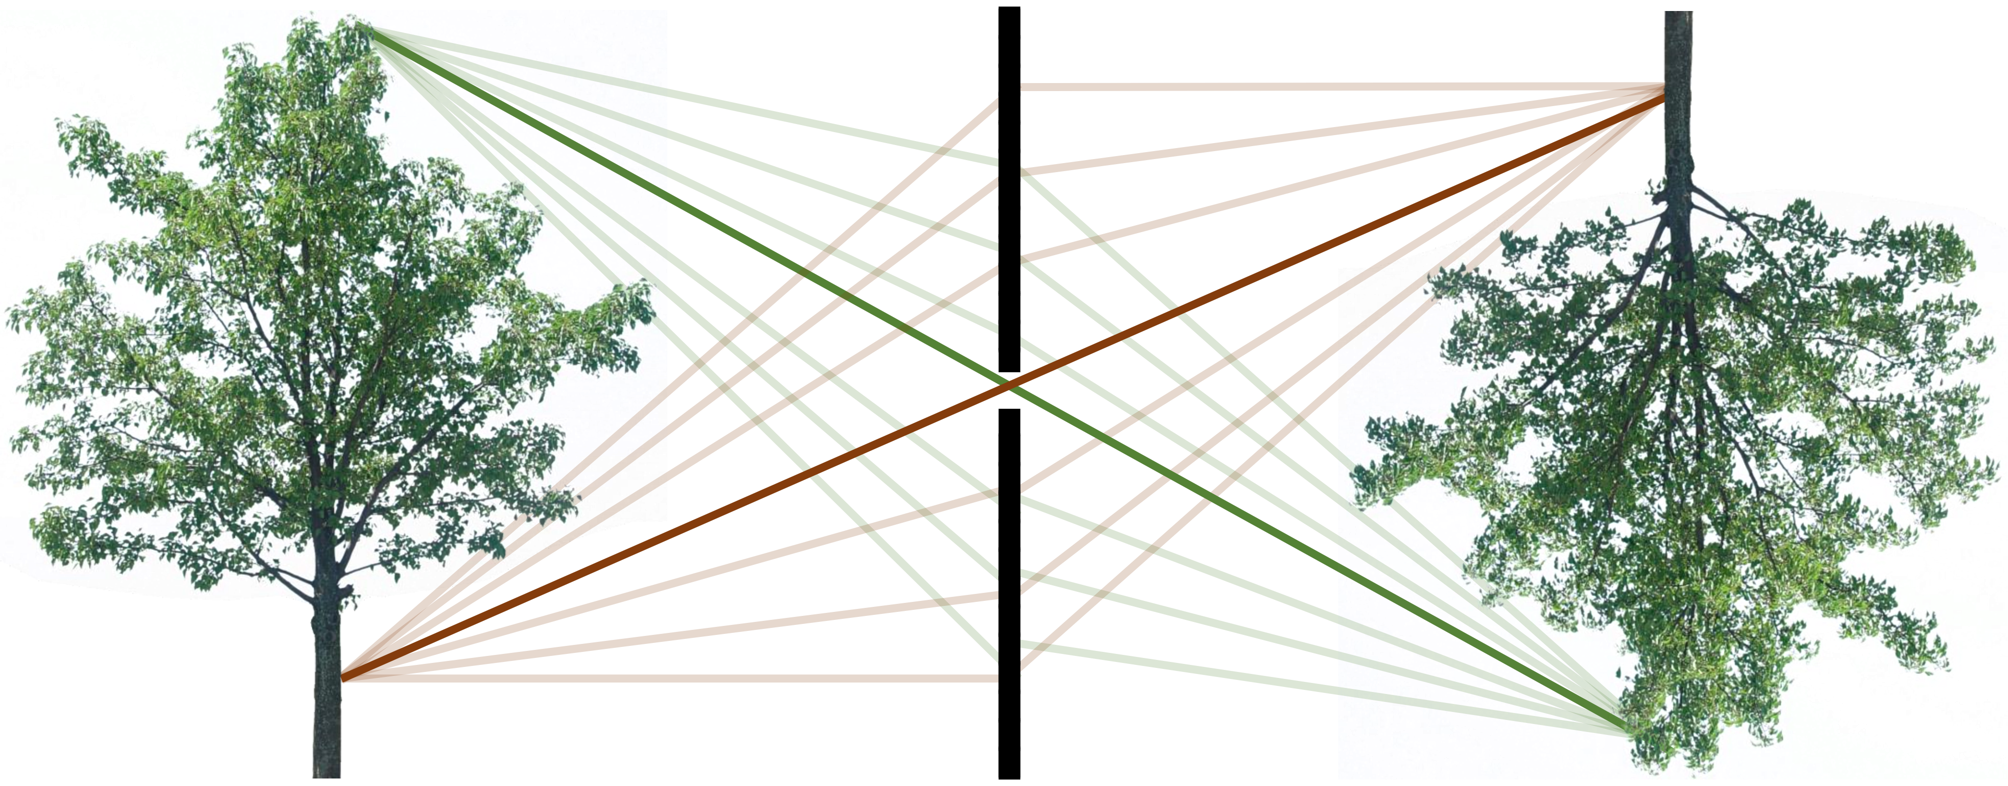

The pinhole camera

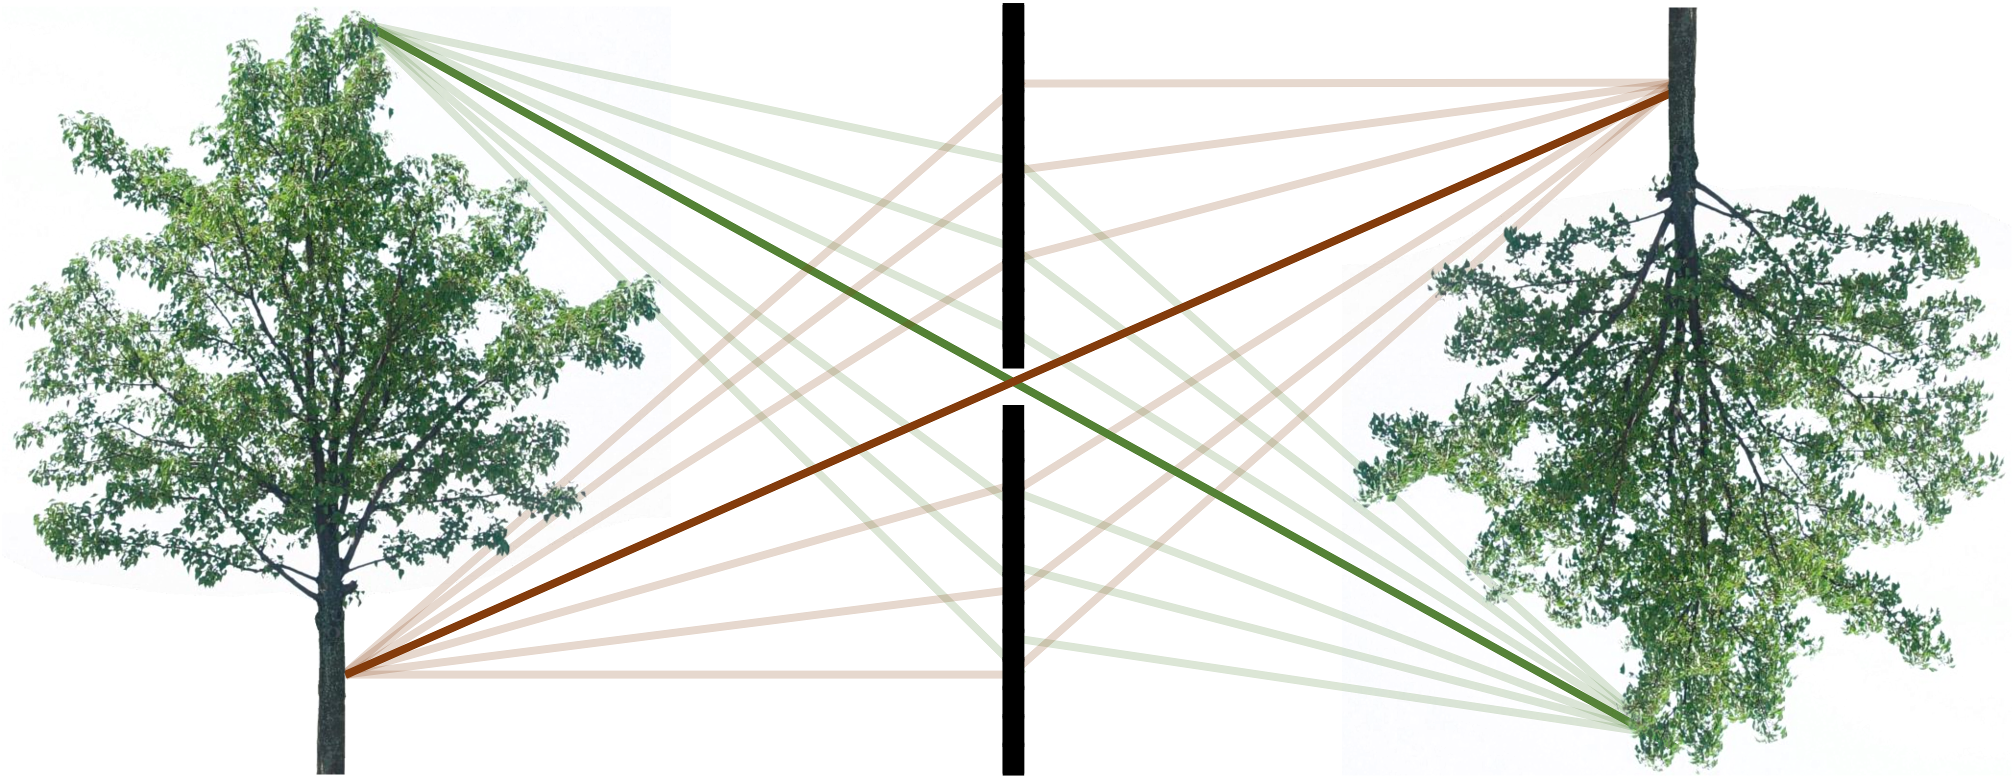

Central rays propagate in the same way for both models!

Describing both lens and pinhole cameras

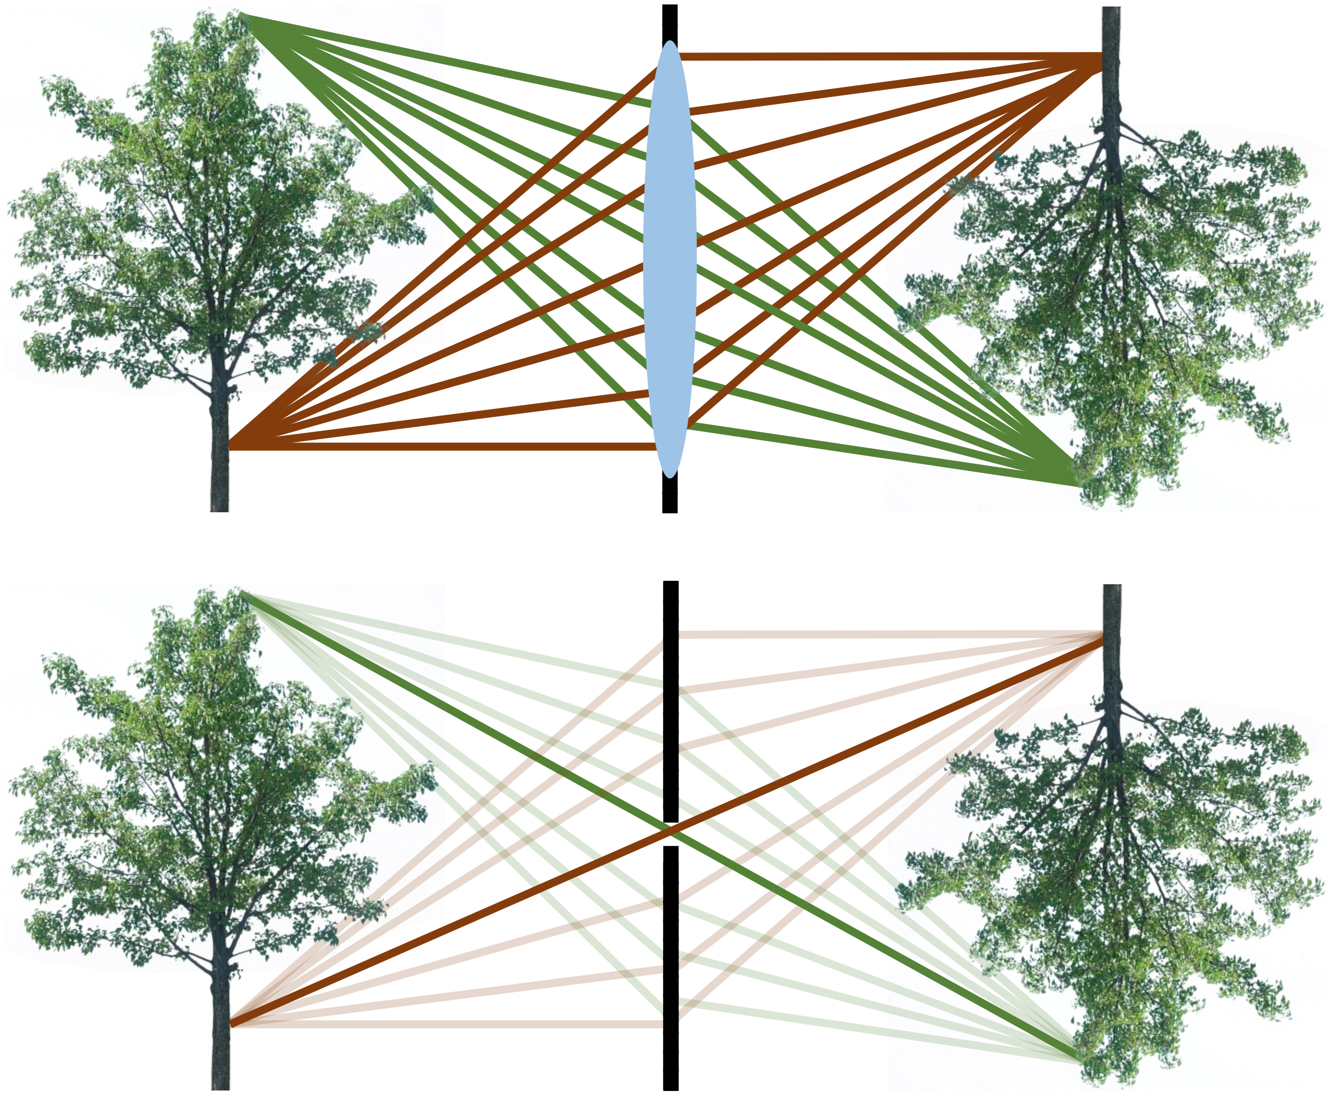

We can derive properties and descriptions that hold for both camera models if:

- We use only central rays.
- We assume the lens camera is in focus.

Important difference: focal length

In a pinhole camera, focal length is distance between aperture and sensor

Important difference: focal length

In a lens camera, focal length is distance where parallel rays intersect

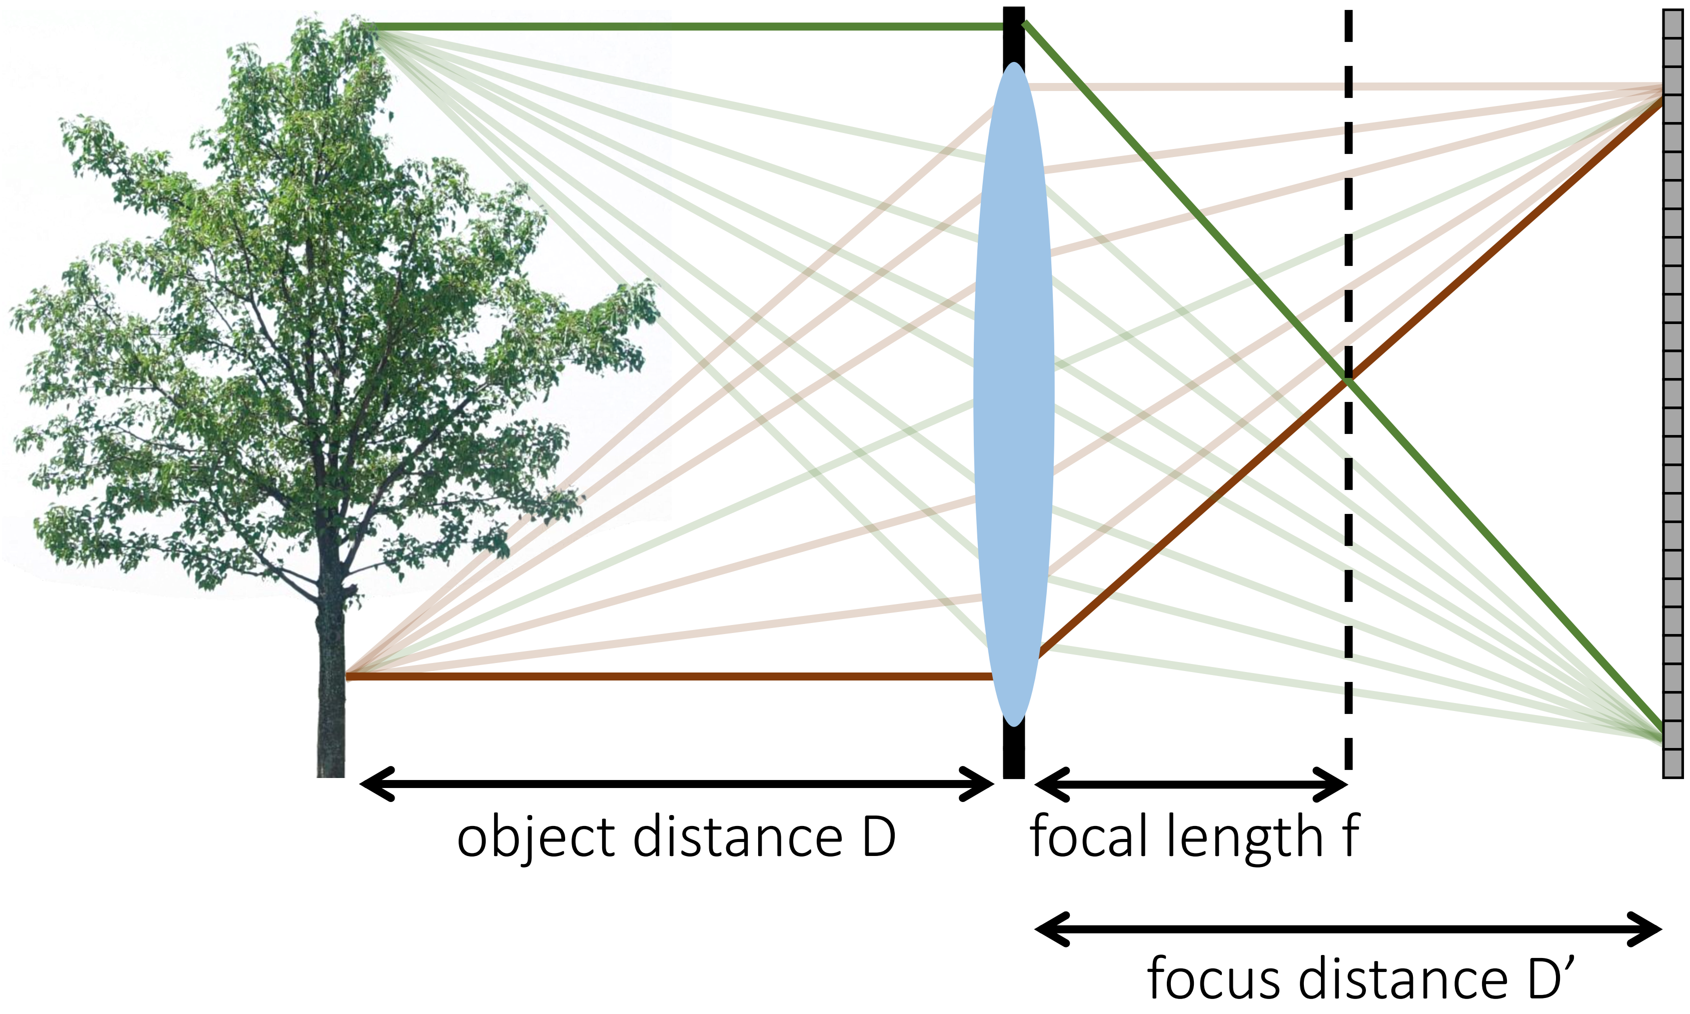

Describing both lens and pinhole cameras

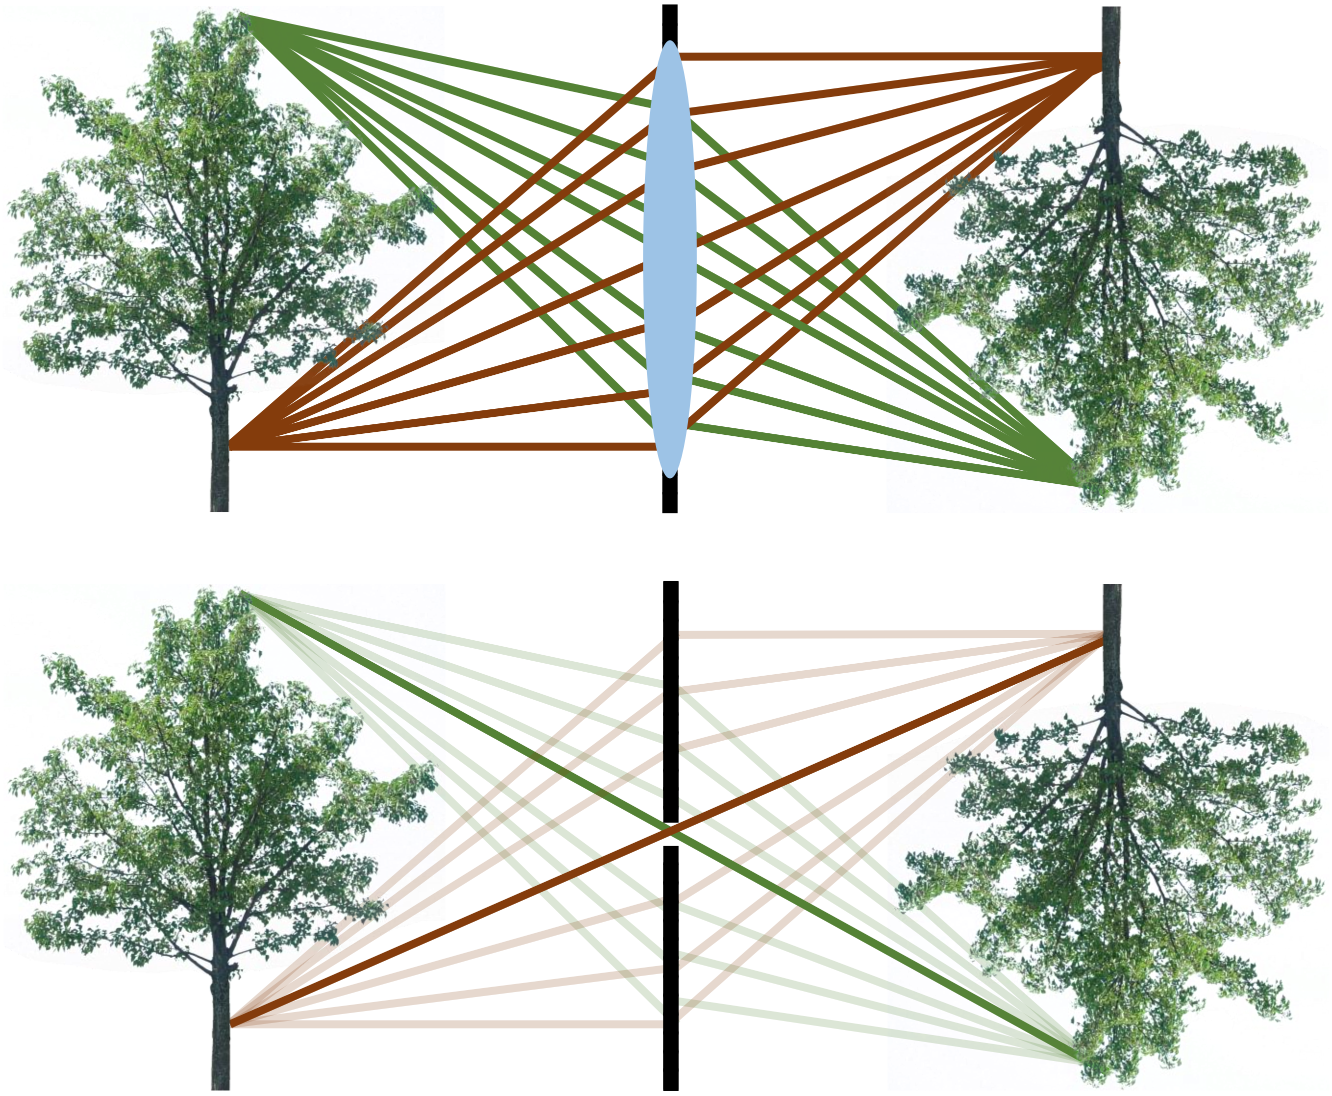

We can derive properties and descriptions that hold for both camera models if:

- We use only central rays.
- We assume the lens camera is in focus.
- We assume that the focus distance of the lens camera is equal to the focal length of the pinhole camera.

Field of view

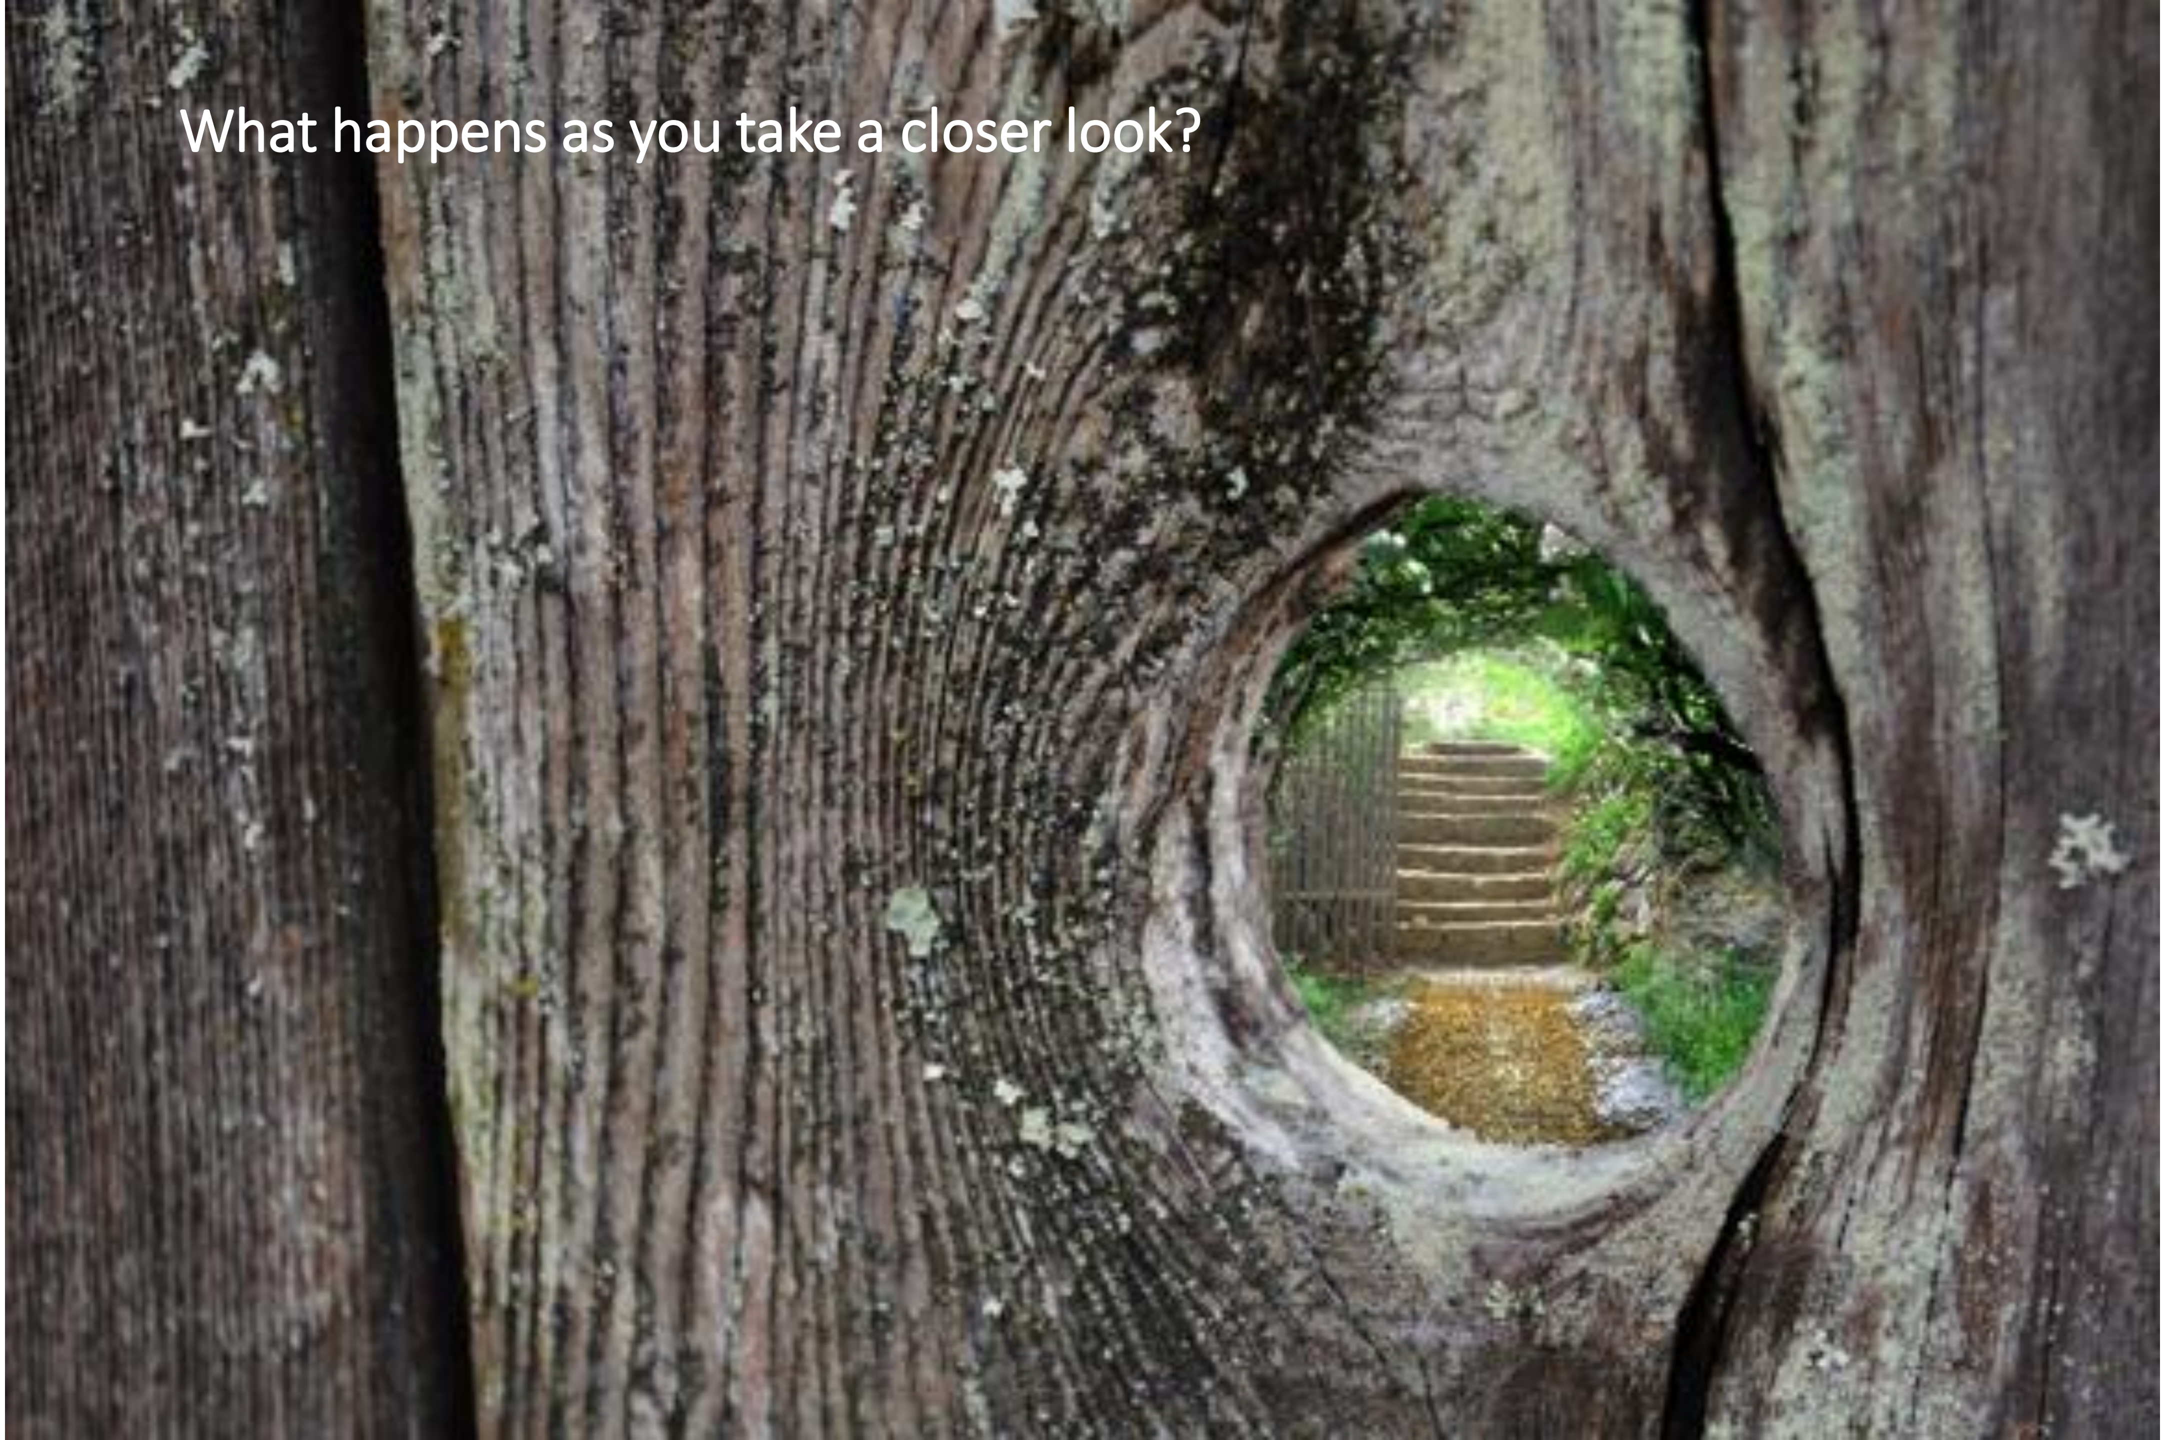
What happens as you take a closer look?

Field of view

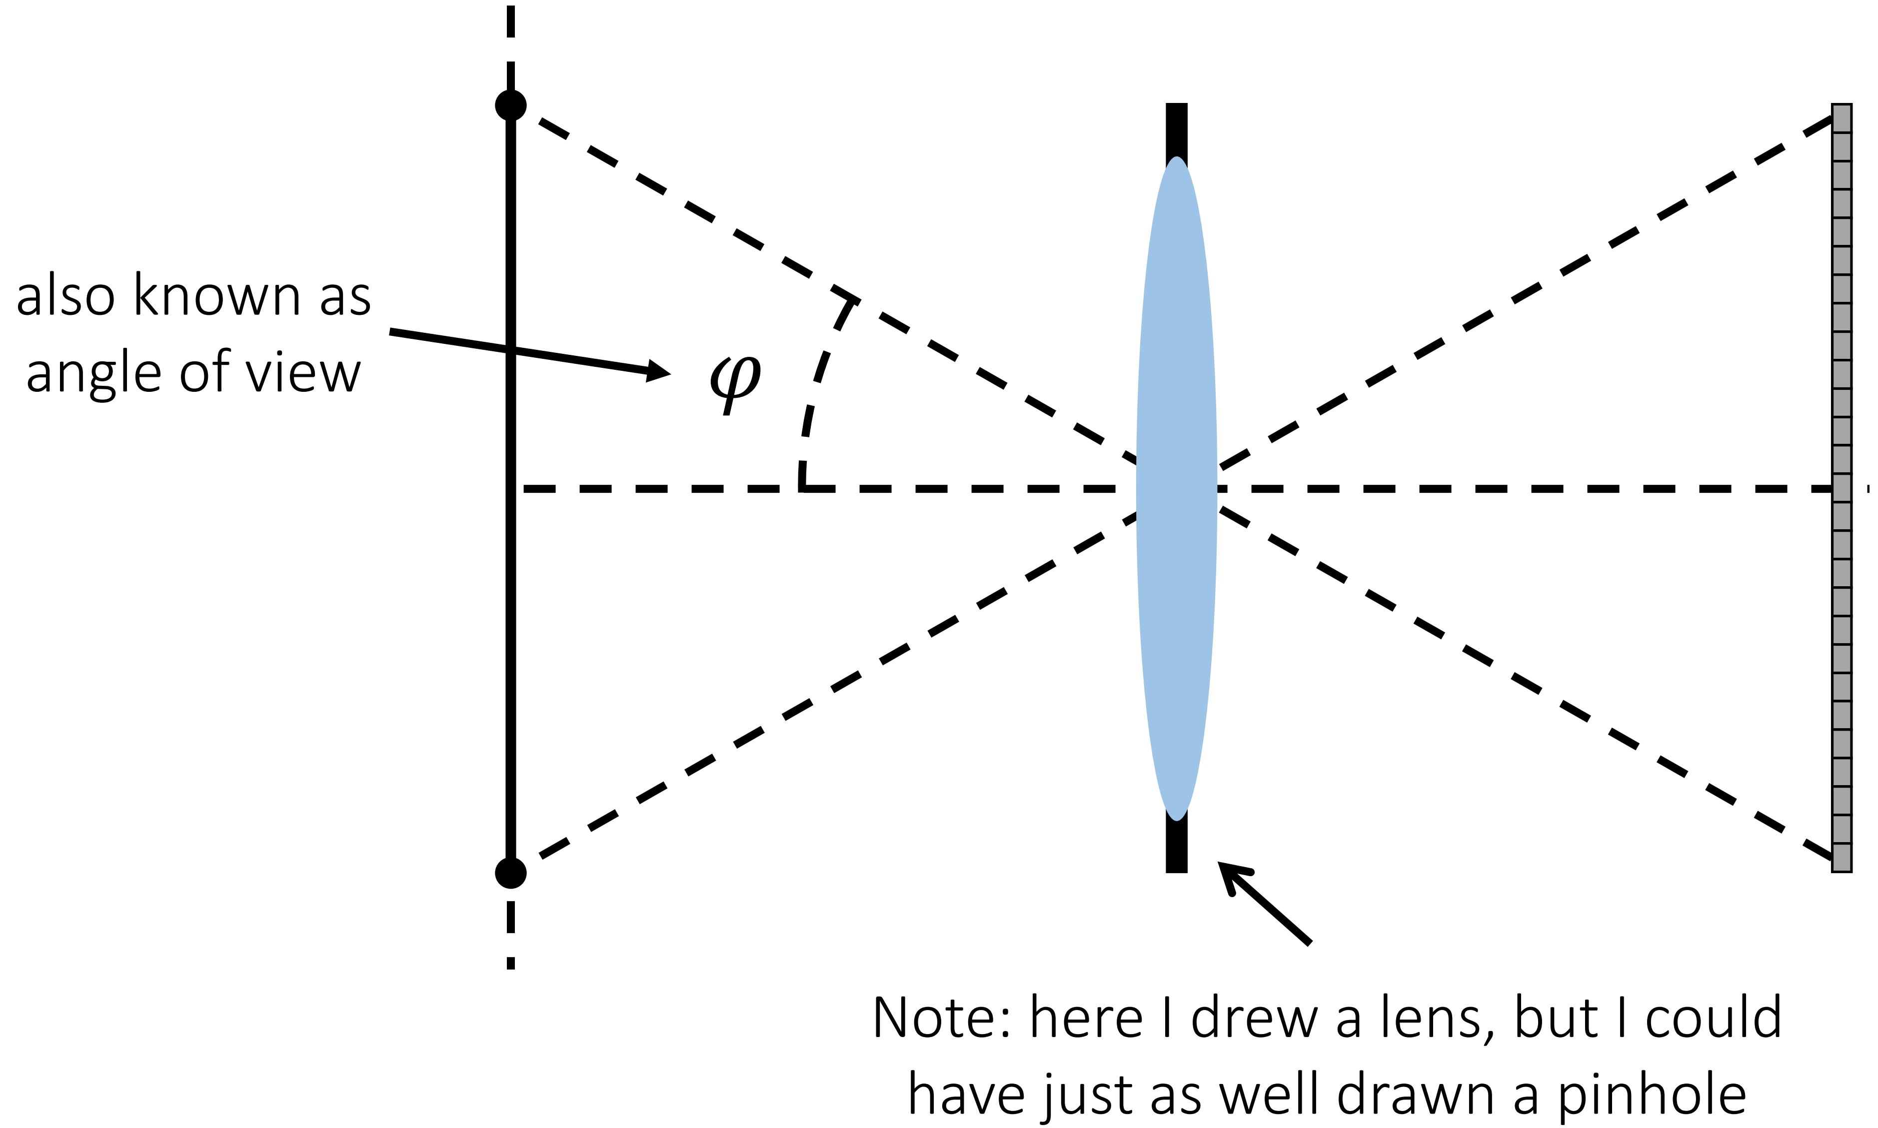

Field of view

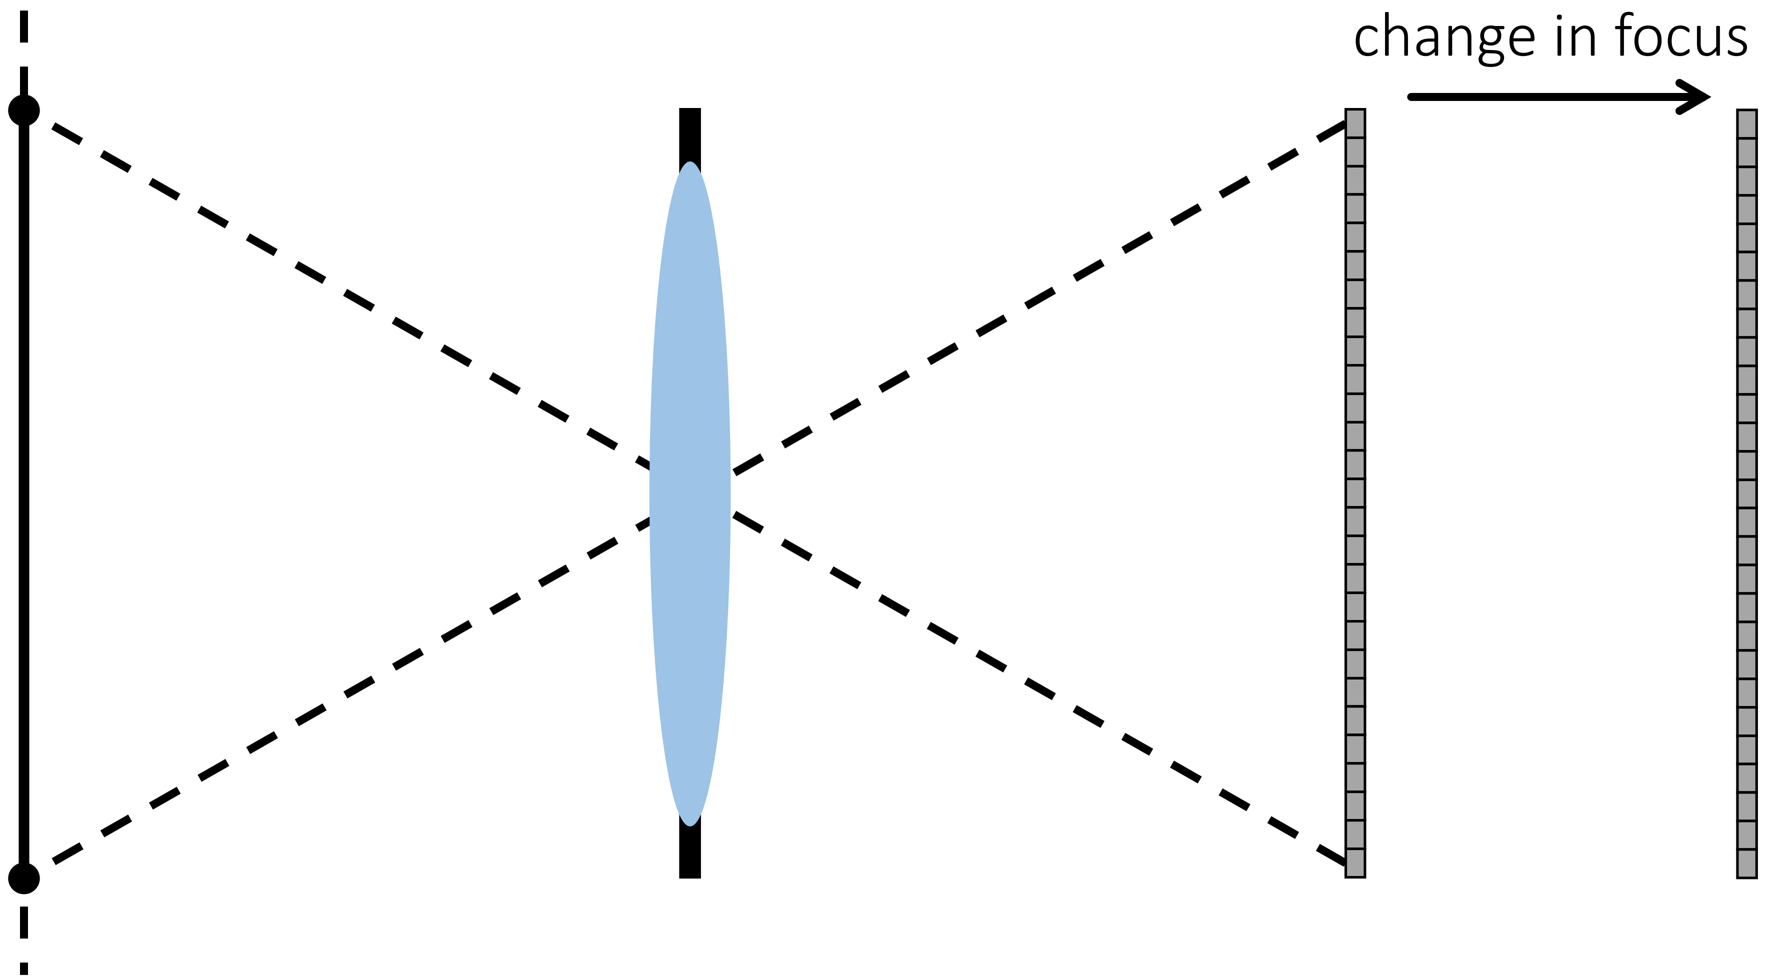

- What happens to field of view when we focus closer?

Field of view

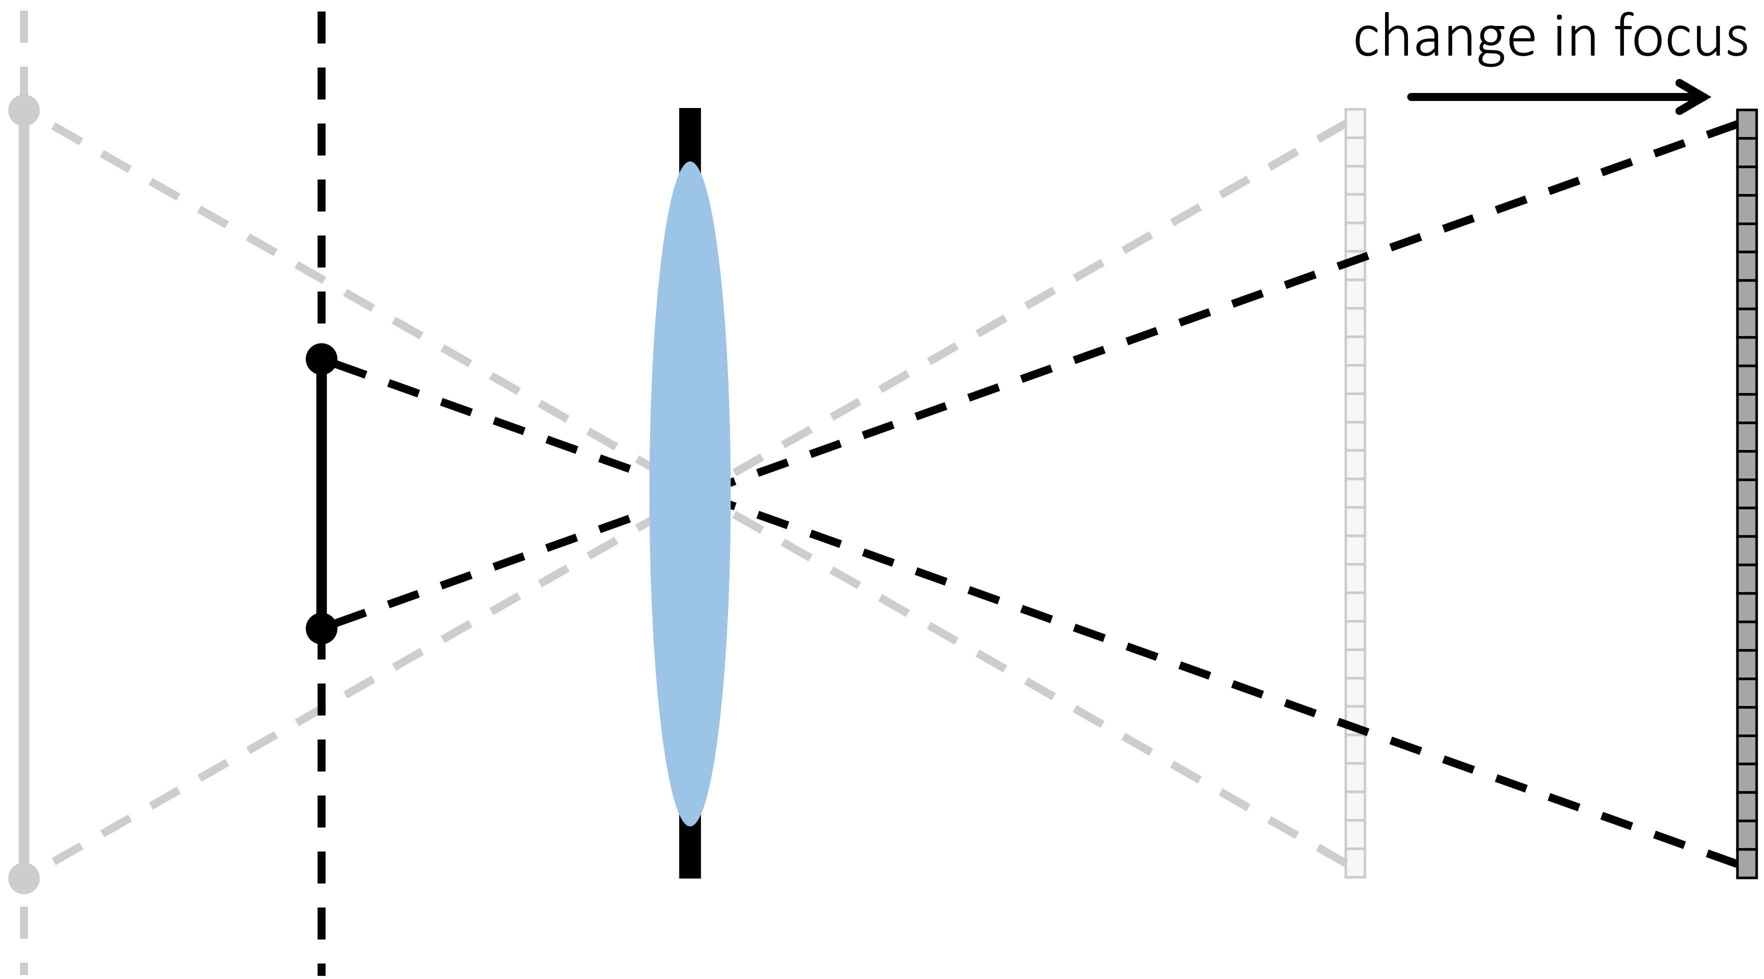

- What happens to field of view when we focus closer? → It decreases.

Field of view

Note: "zooming" means changing focal length, which is different from refocusing

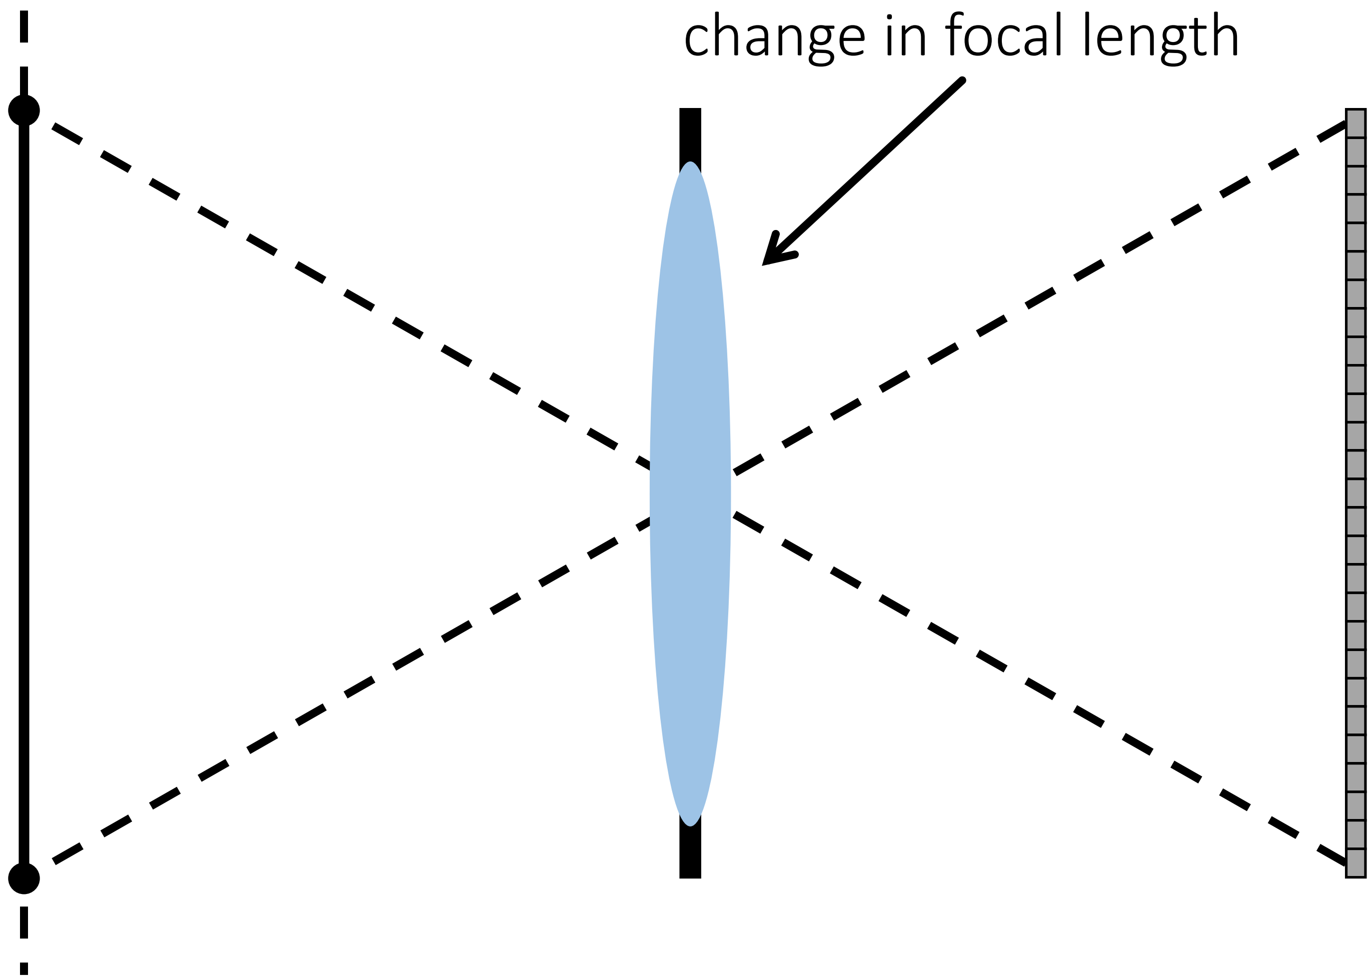

- What happens to field of view when we focus closer? → It decreases.
- What happens to field of view when we increase <u>lens</u> focal length?

Field of view

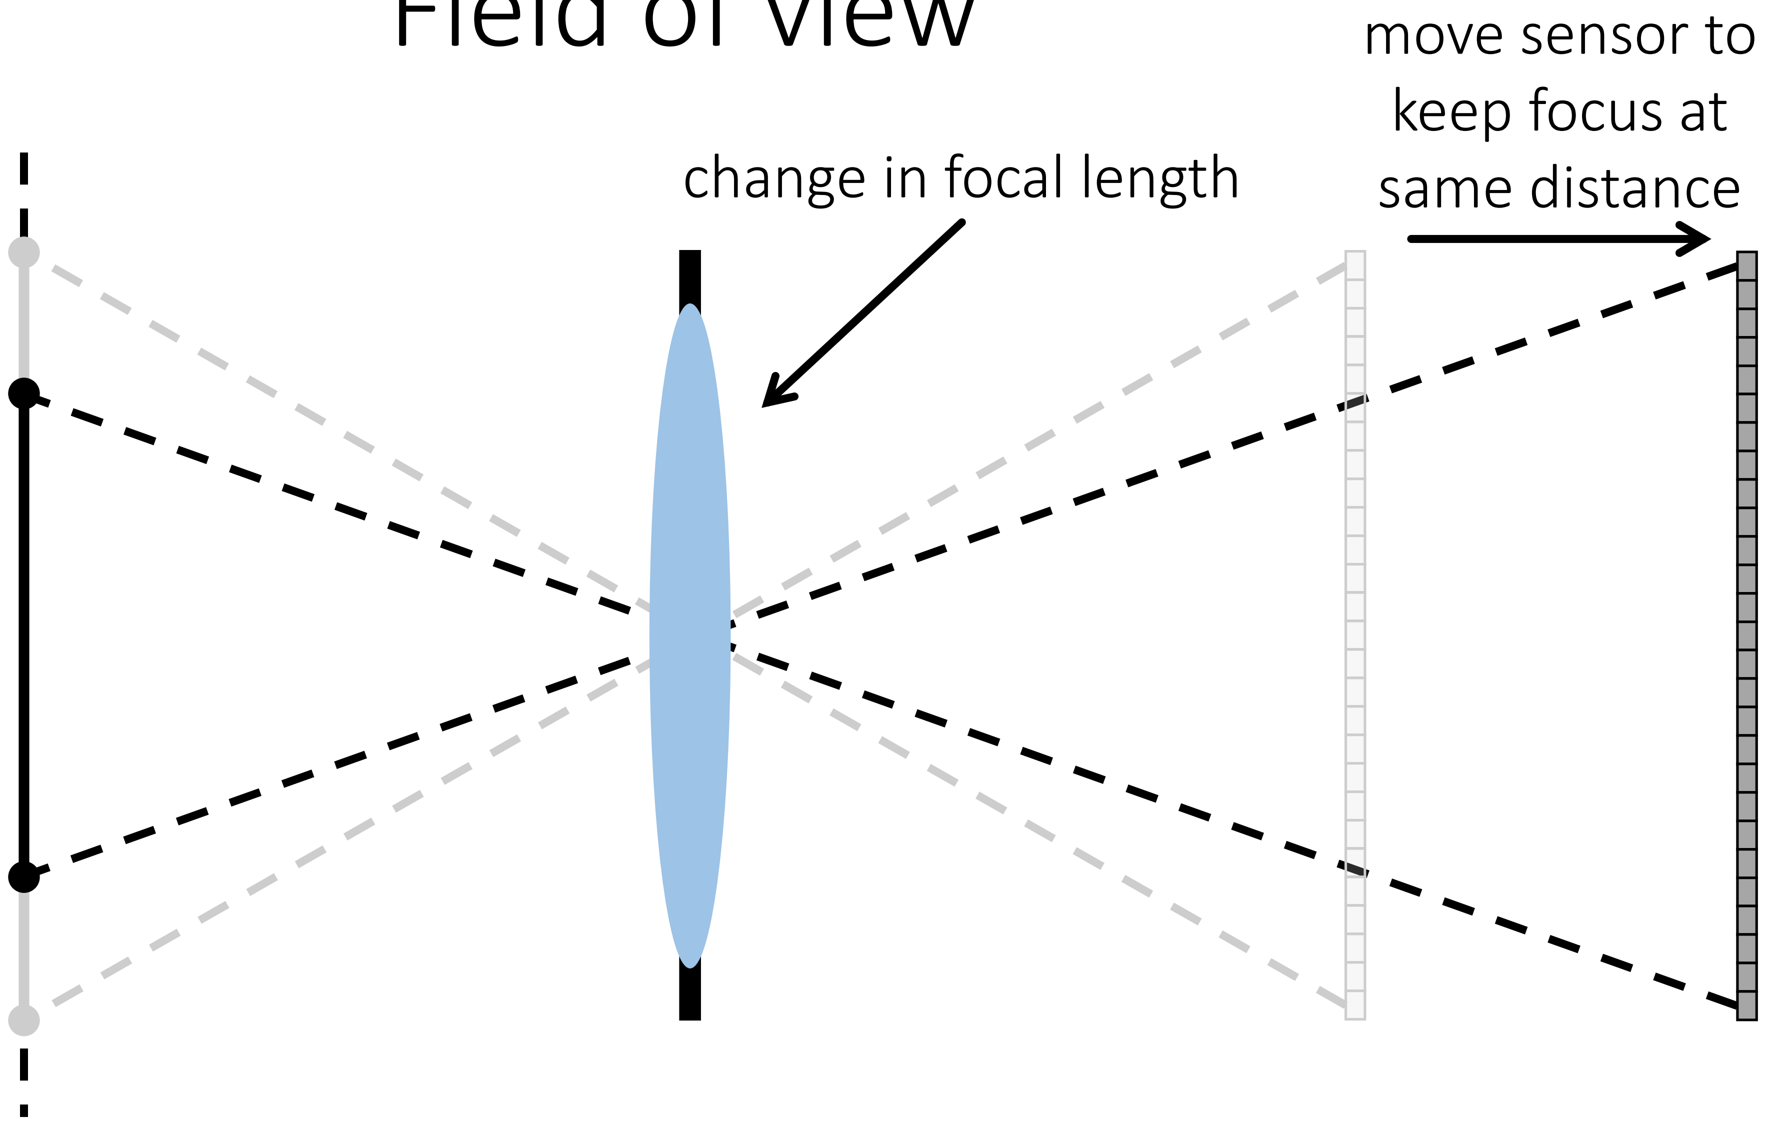

Note: "zooming" means changing focal length, which is different from refocusing

- What happens to field of view when we focus closer? → It decreases.
- What happens to field of view when we increase <u>lens</u> focal length? → It decreases.

Field of view

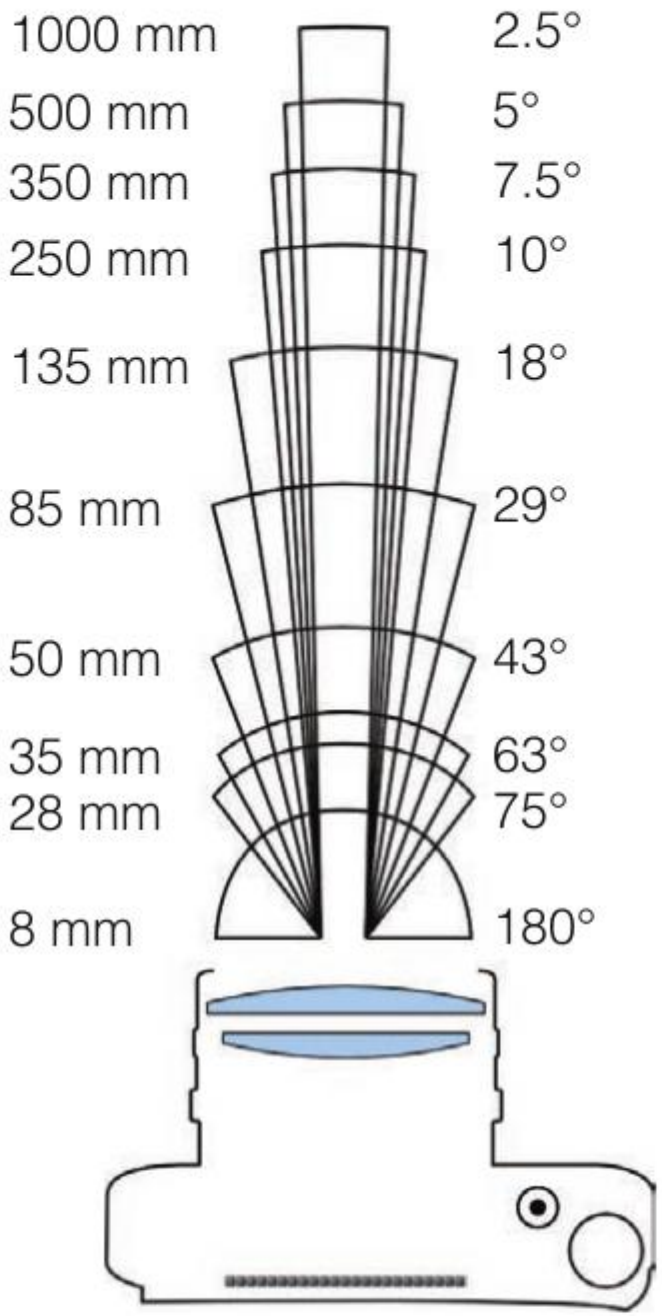

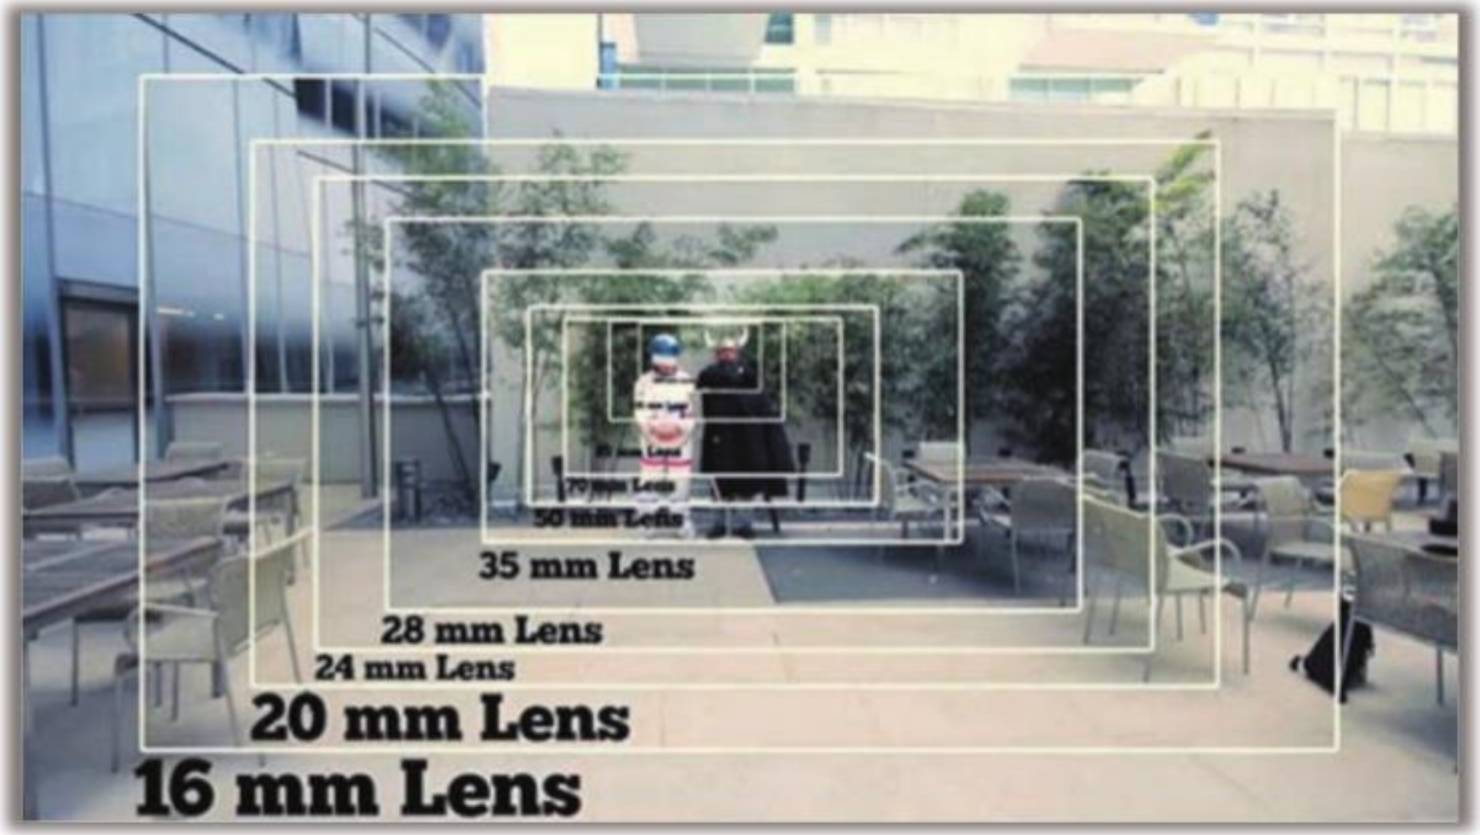

Andrew McWilliams

Field of view

Increasing the <u>lens</u> focal length is similar to cropping

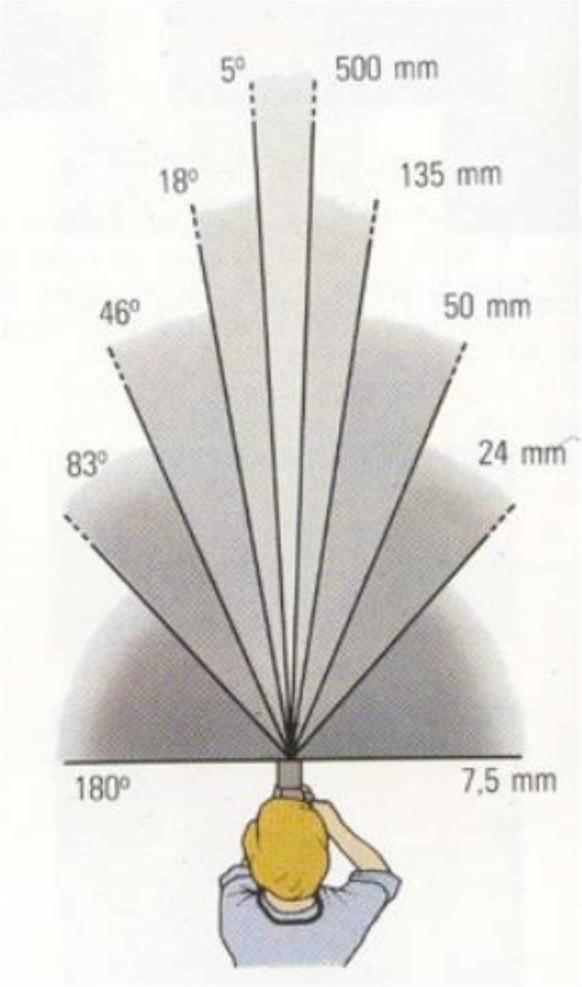

f = 25 mm

f = 50 mm

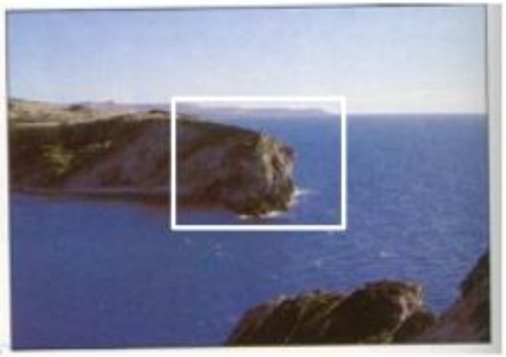

f = 135 mm

Is this effect identical to cropping?

Field of view also depends on sensor size

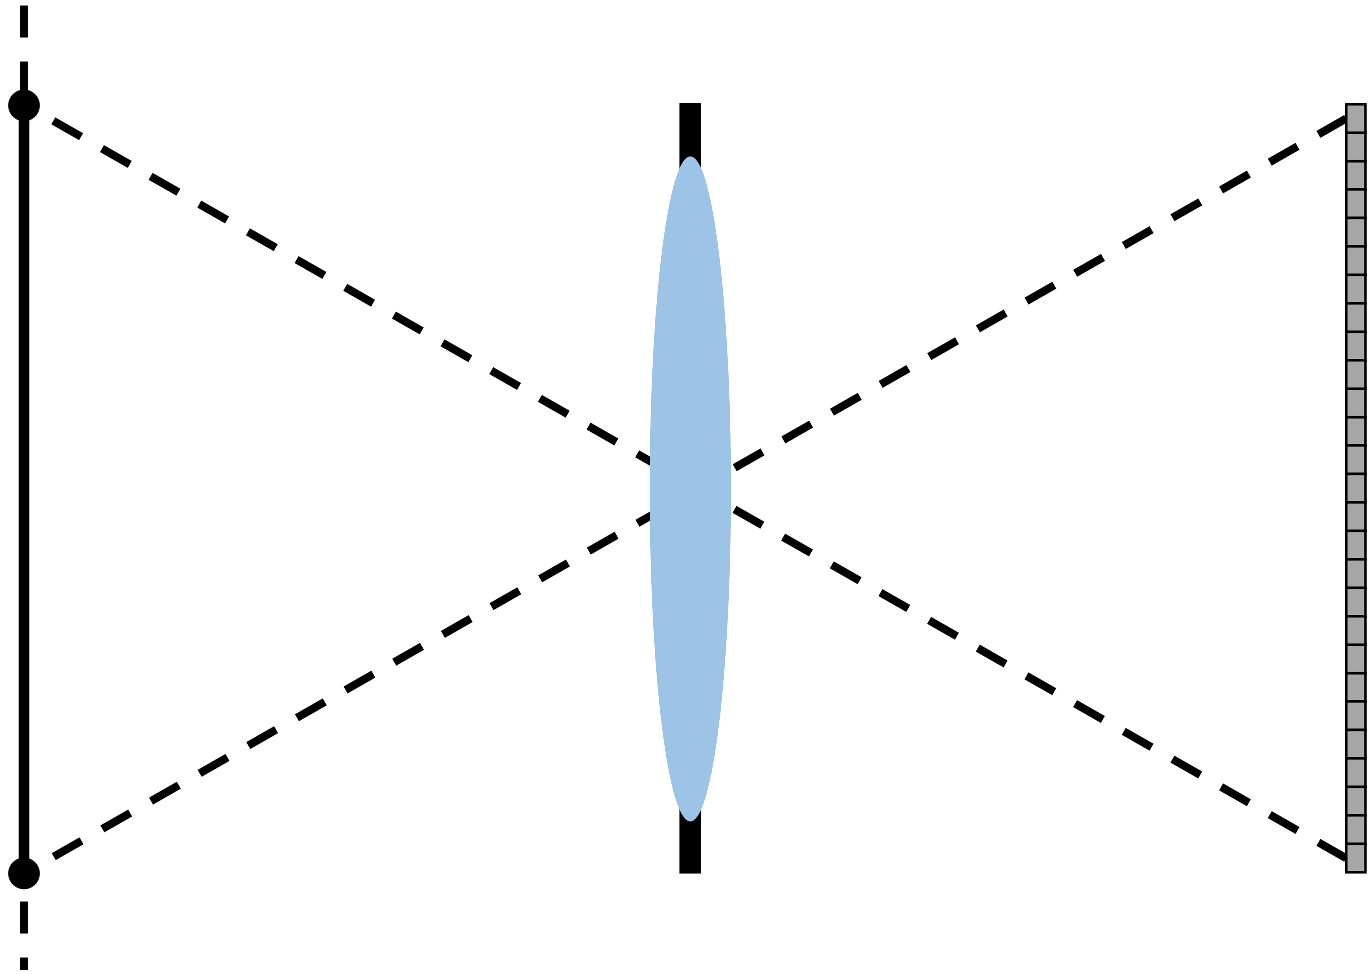

- What happens to field of view when we reduce sensor size?

Field of view also depends on sensor size

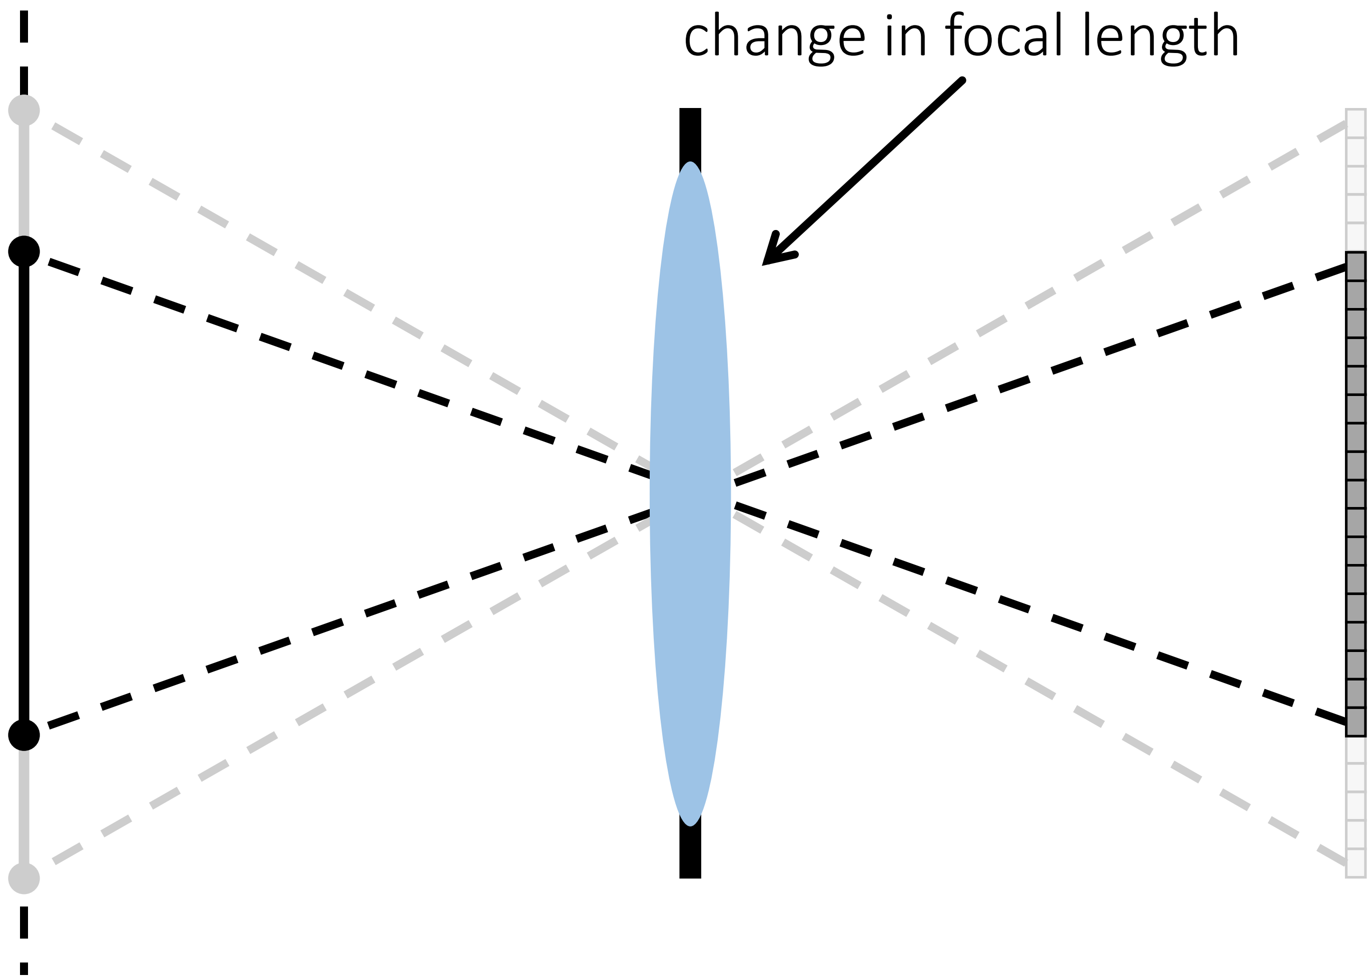

- What happens to field of view when we reduce sensor size? → It decreases.

Field of view also depends on sensor size

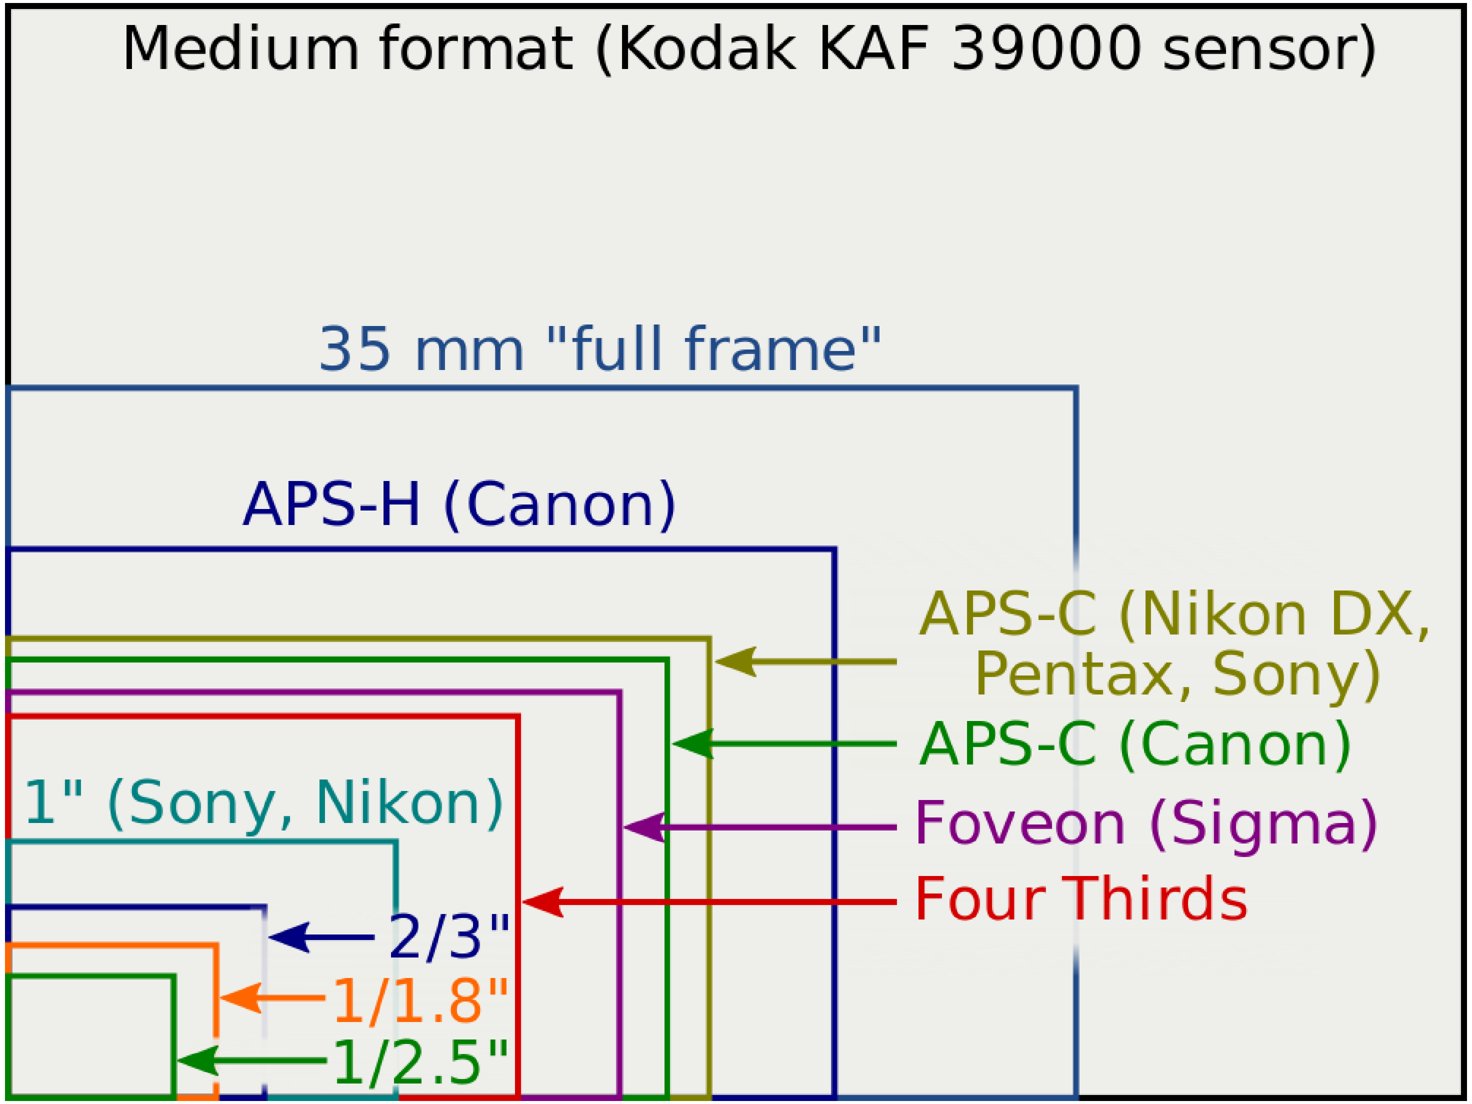

- "Full frame" corresponds to standard film size.
- Digital sensors come in smaller formats due to manufacturing limitations (now mostly overcome).
- Lenses are often described in terms of field of view on film instead of focal length.
- These descriptions are invalid when not using full-frame sensor.

Crop factor

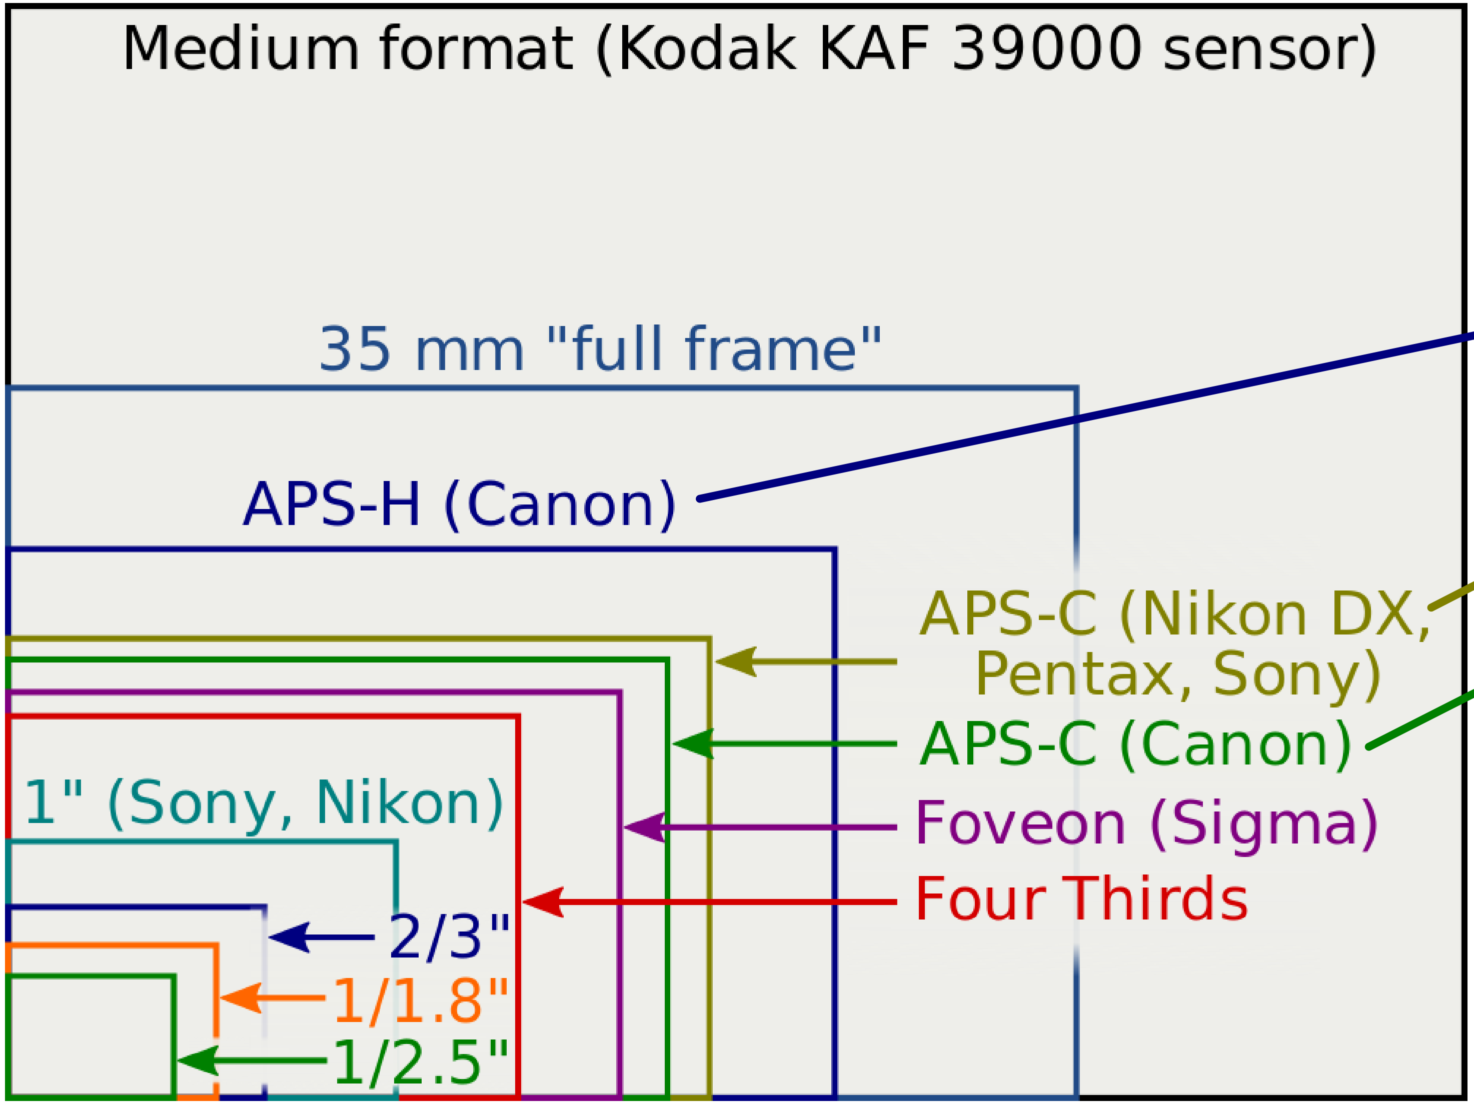

How much field of view is cropped when using a sensor smaller than full frame.

Perspective

The pinhole camera

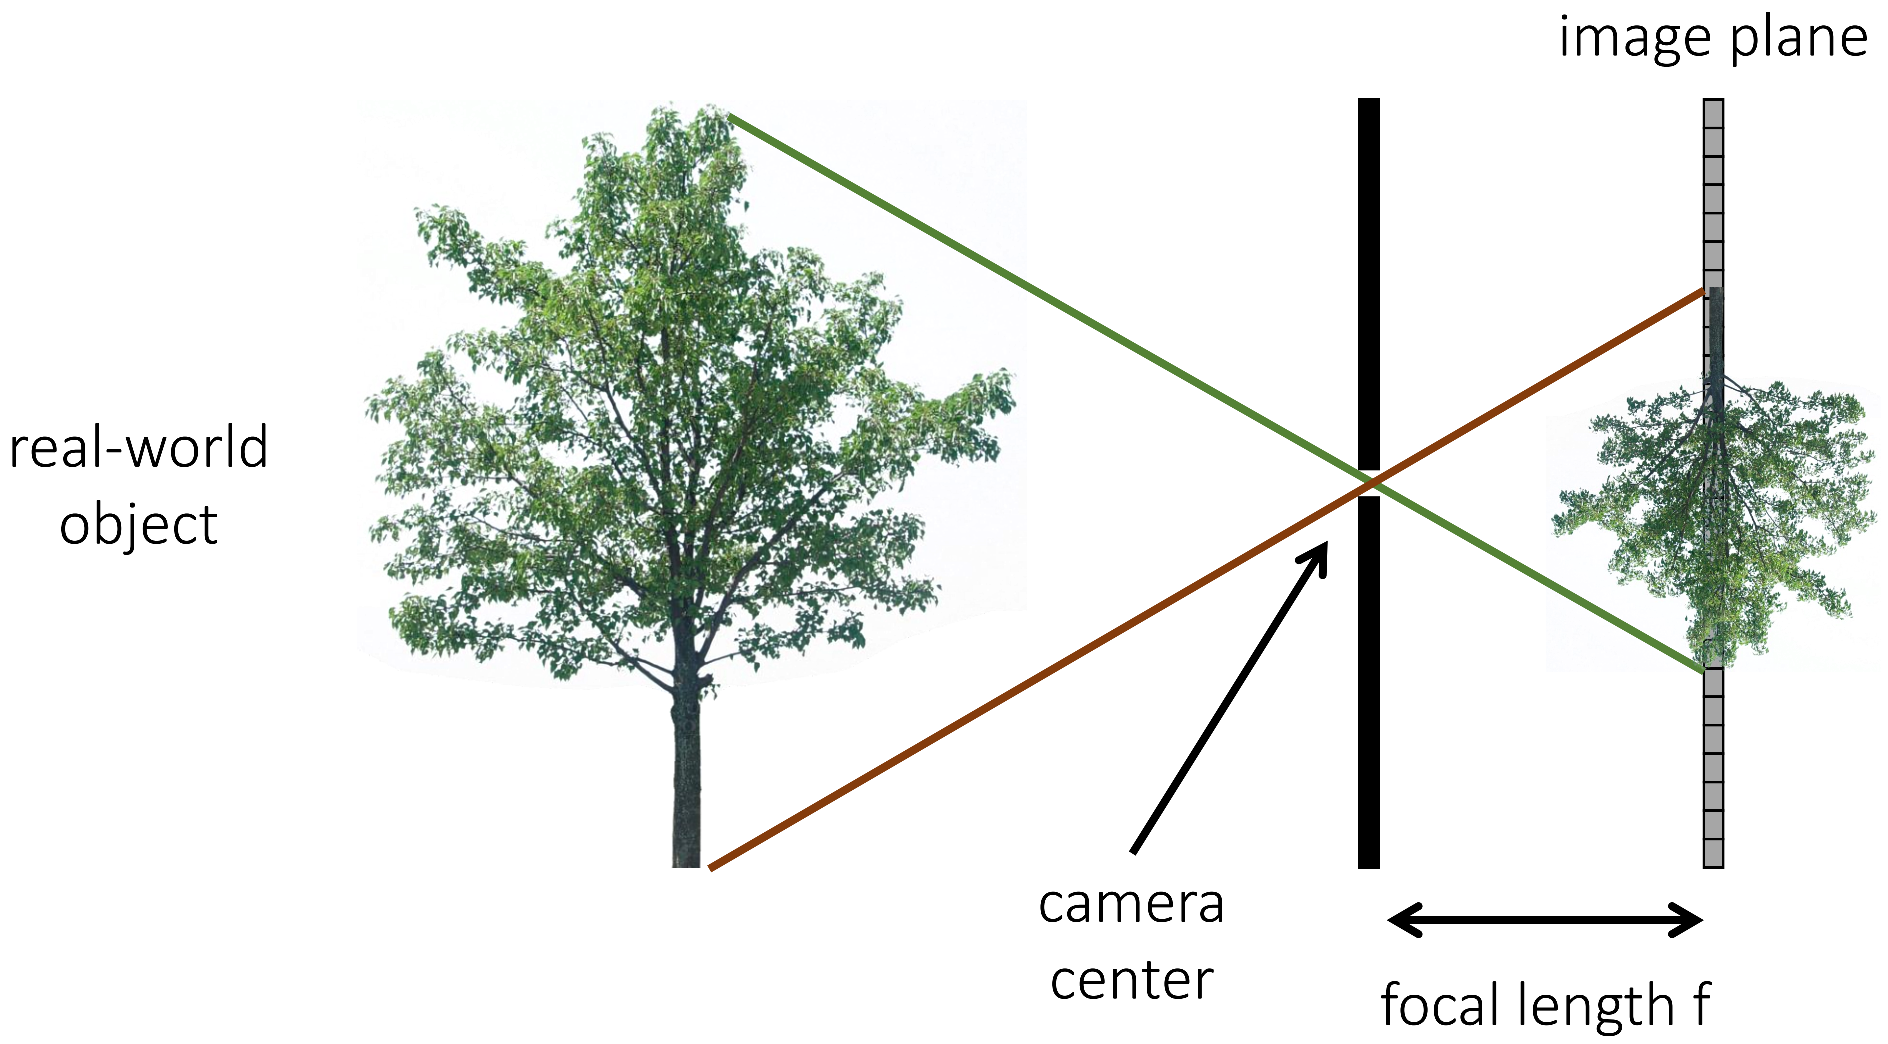

The (rearranged) pinhole camera

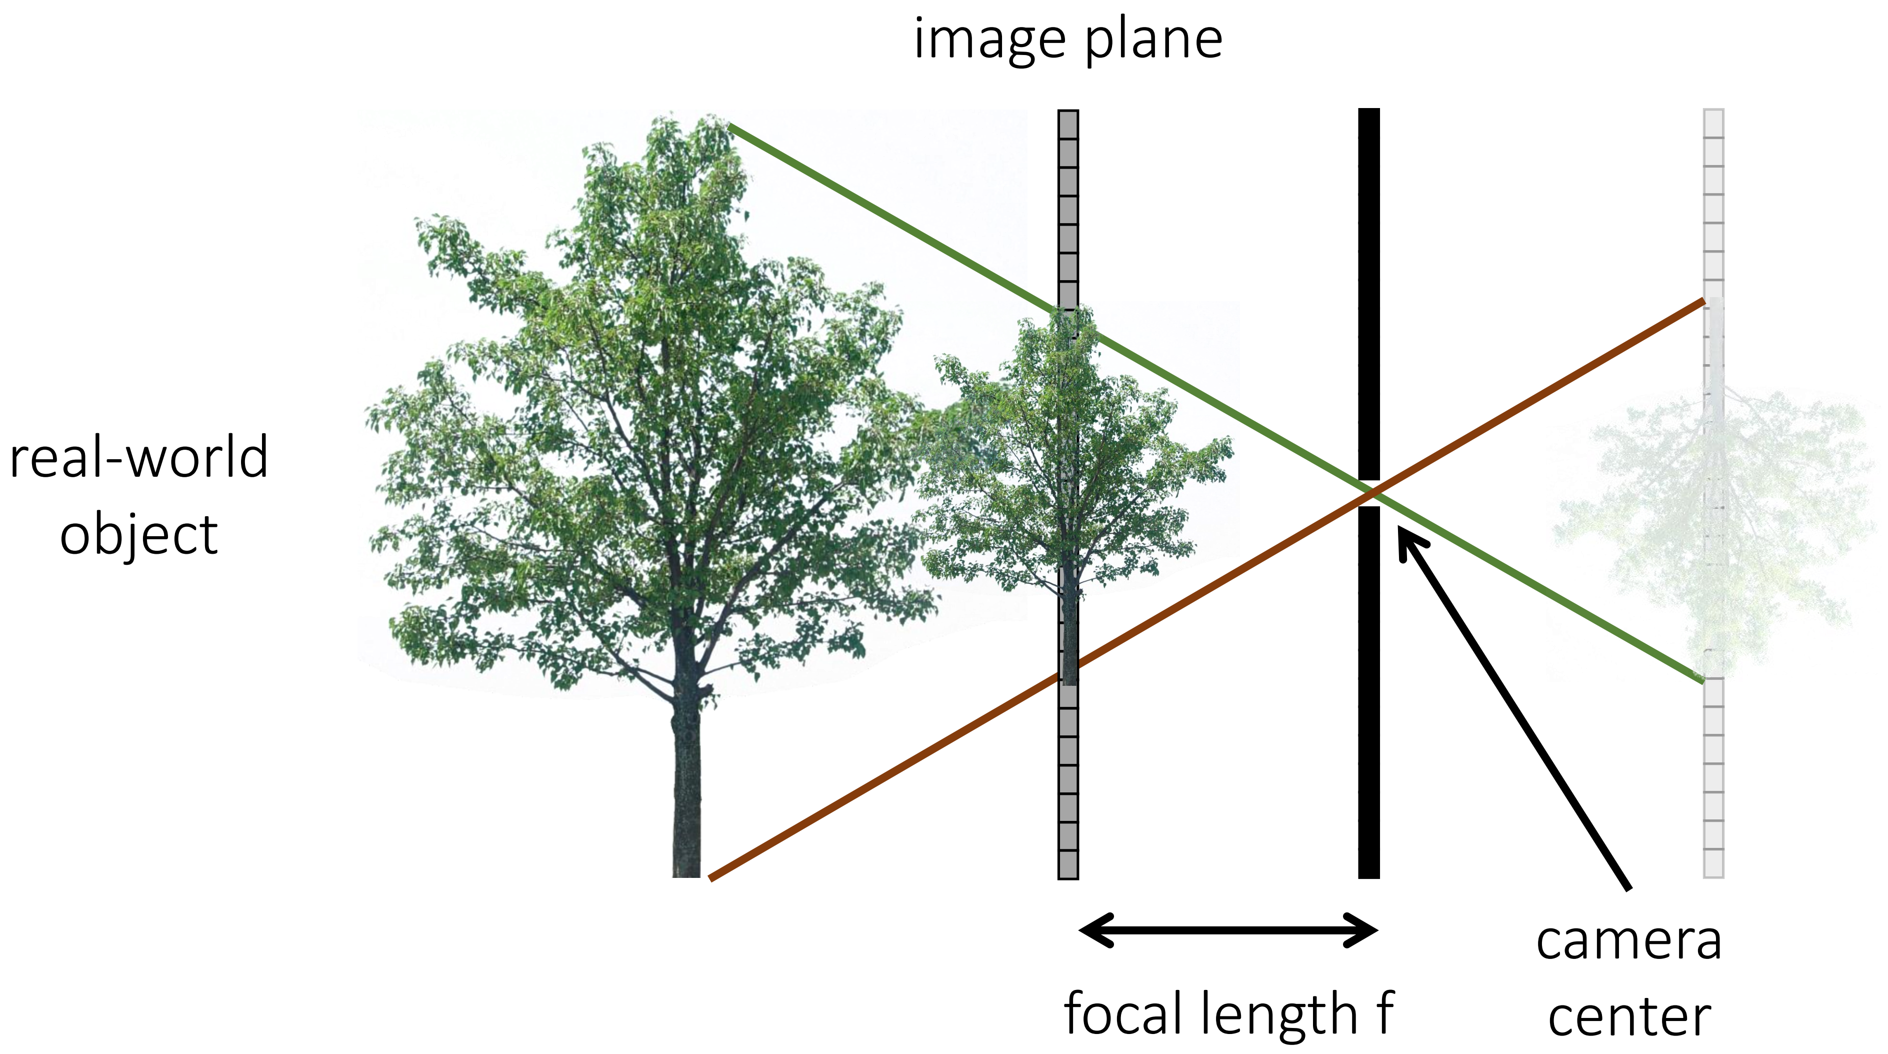

The 2D view of the (rearranged) pinhole camera

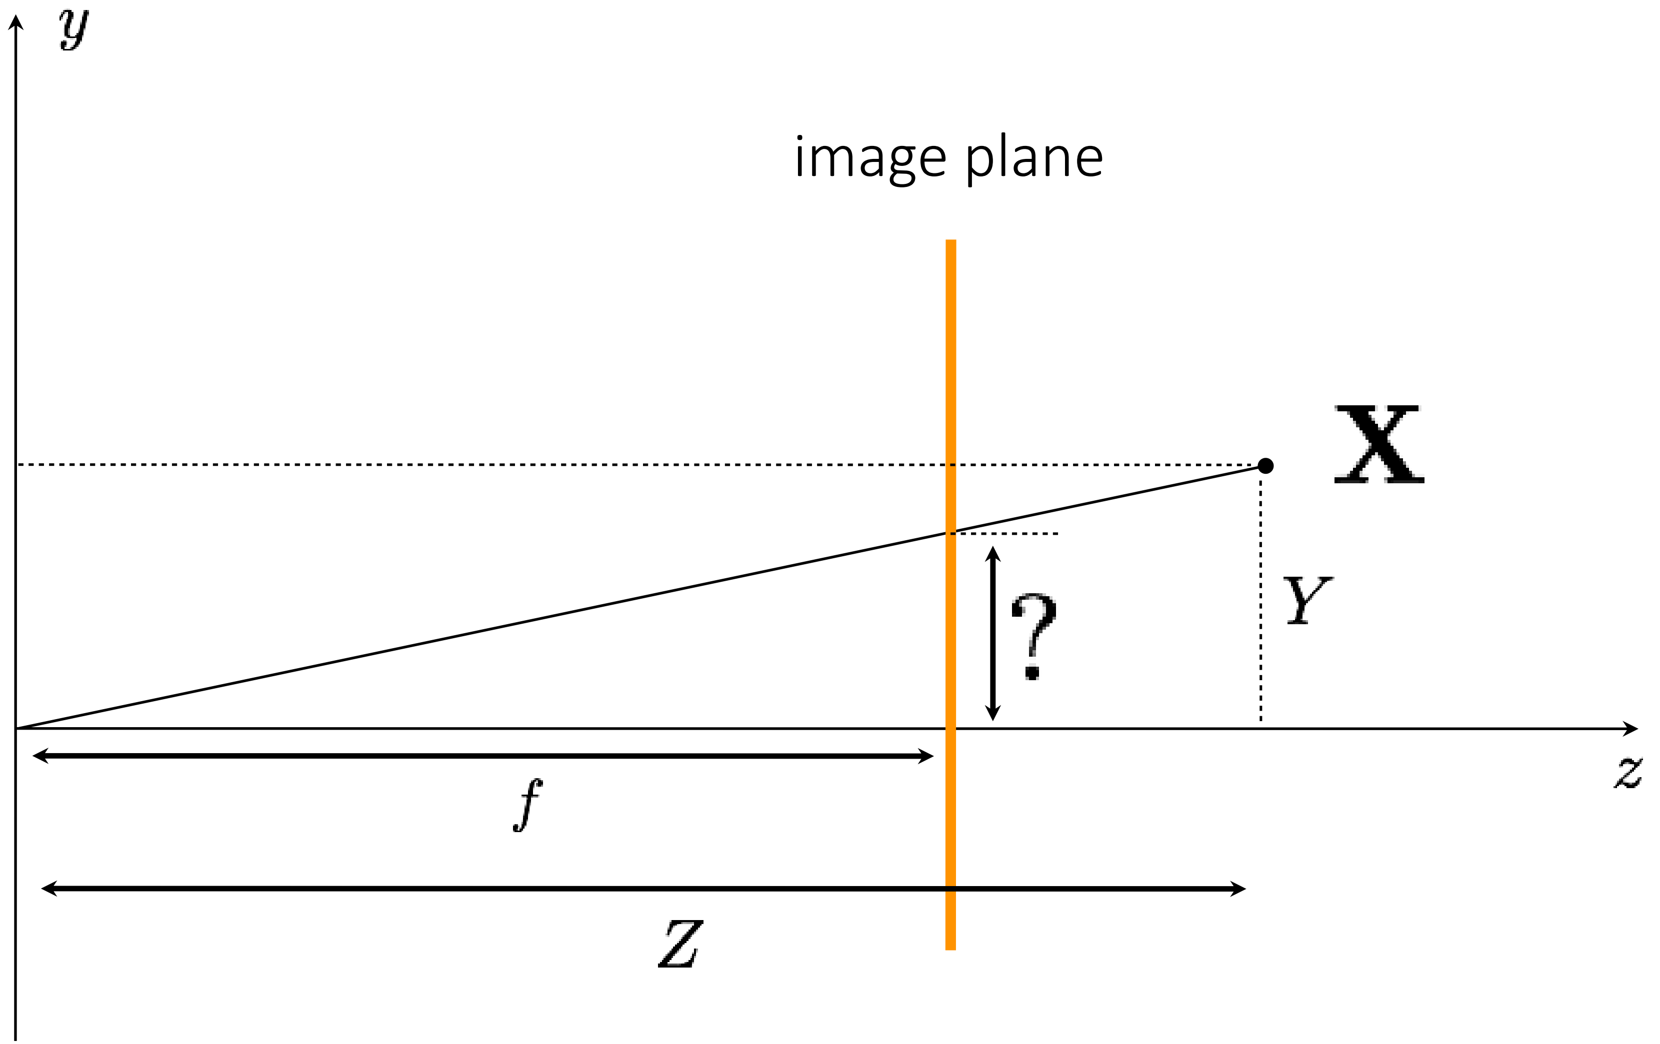

What is the equation for image coordinate **x** in terms of **X**?

The 2D view of the (rearranged) pinhole camera

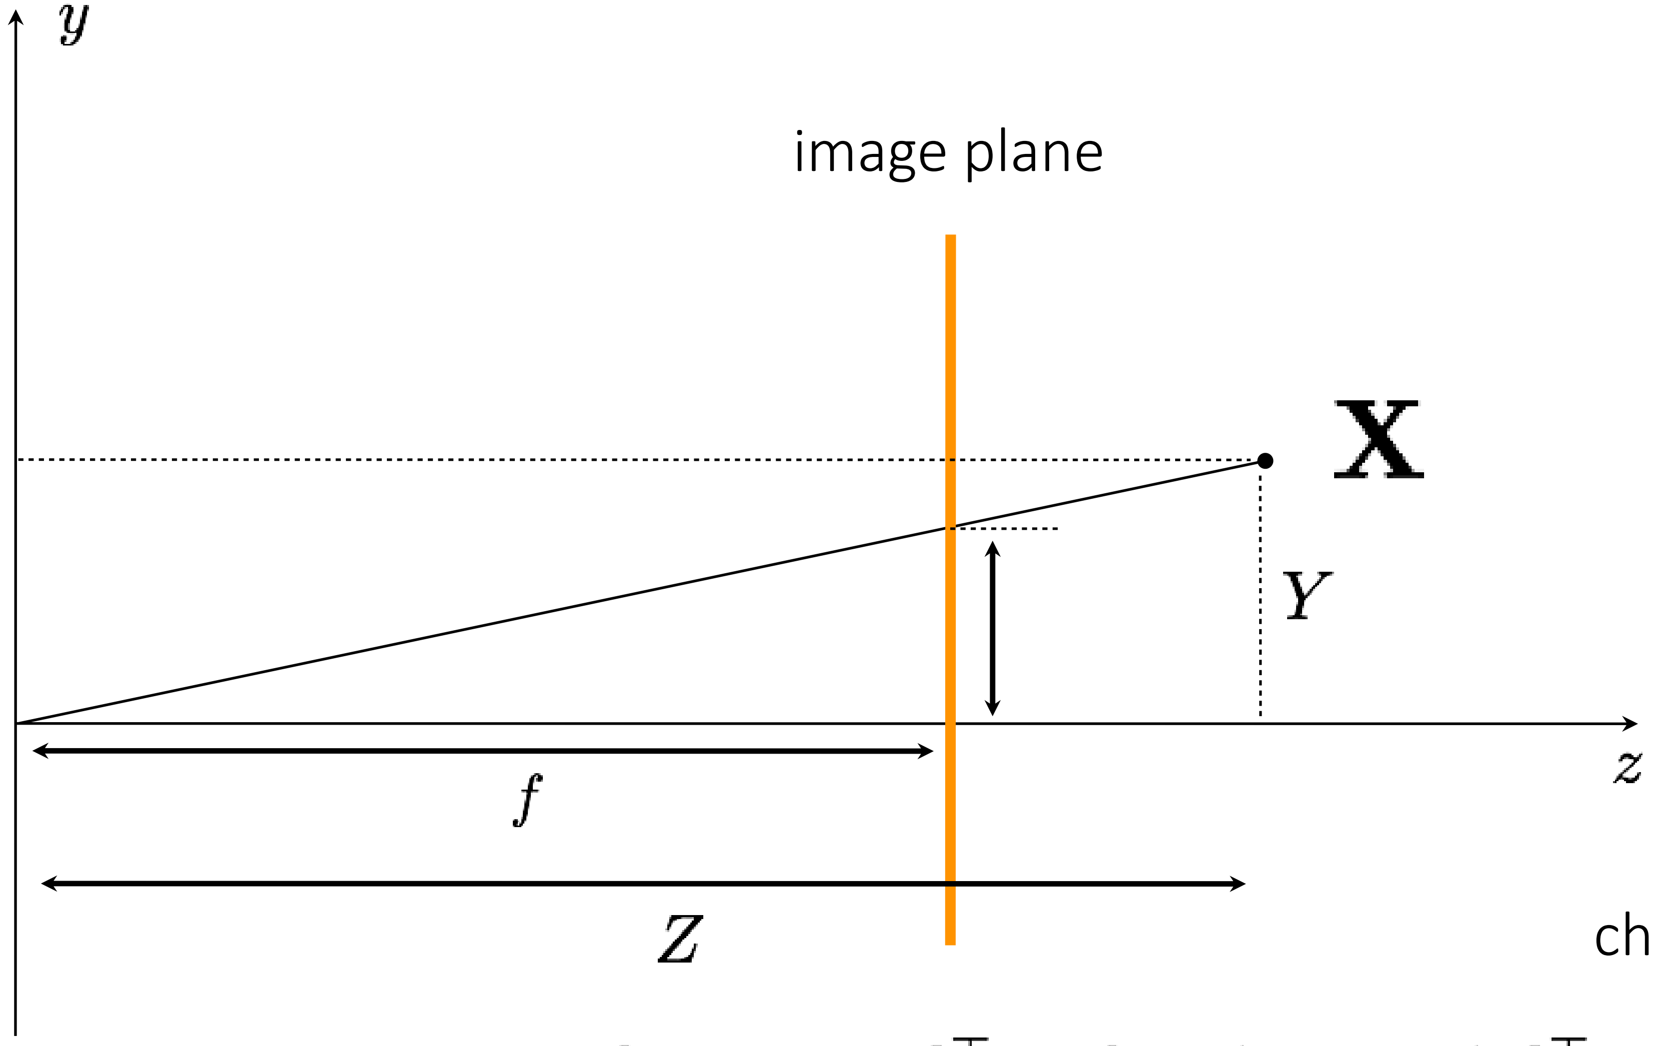

magnification changes with depth

$$[X \quad Y \quad Z]^\top \mapsto [fX/Z \quad fY/Z]^\top$$

Forced perspective

The Ames room illusion

The Ames room illusion

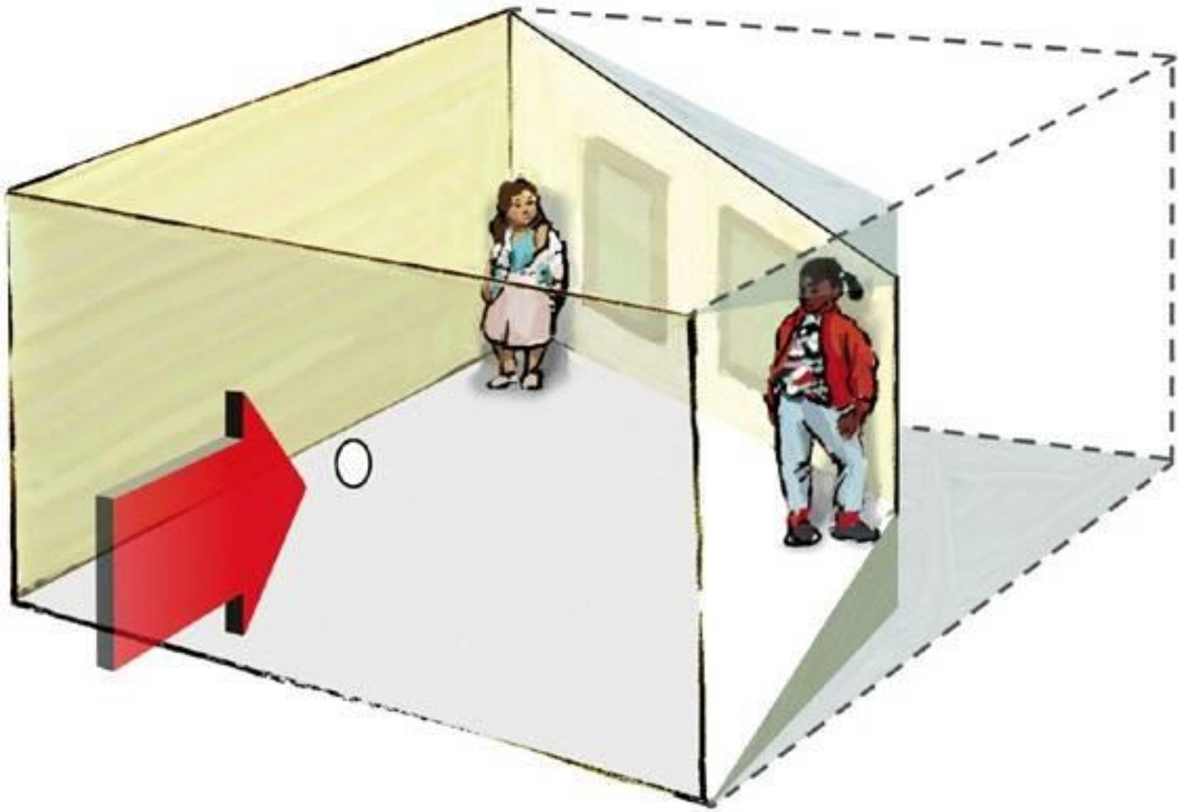

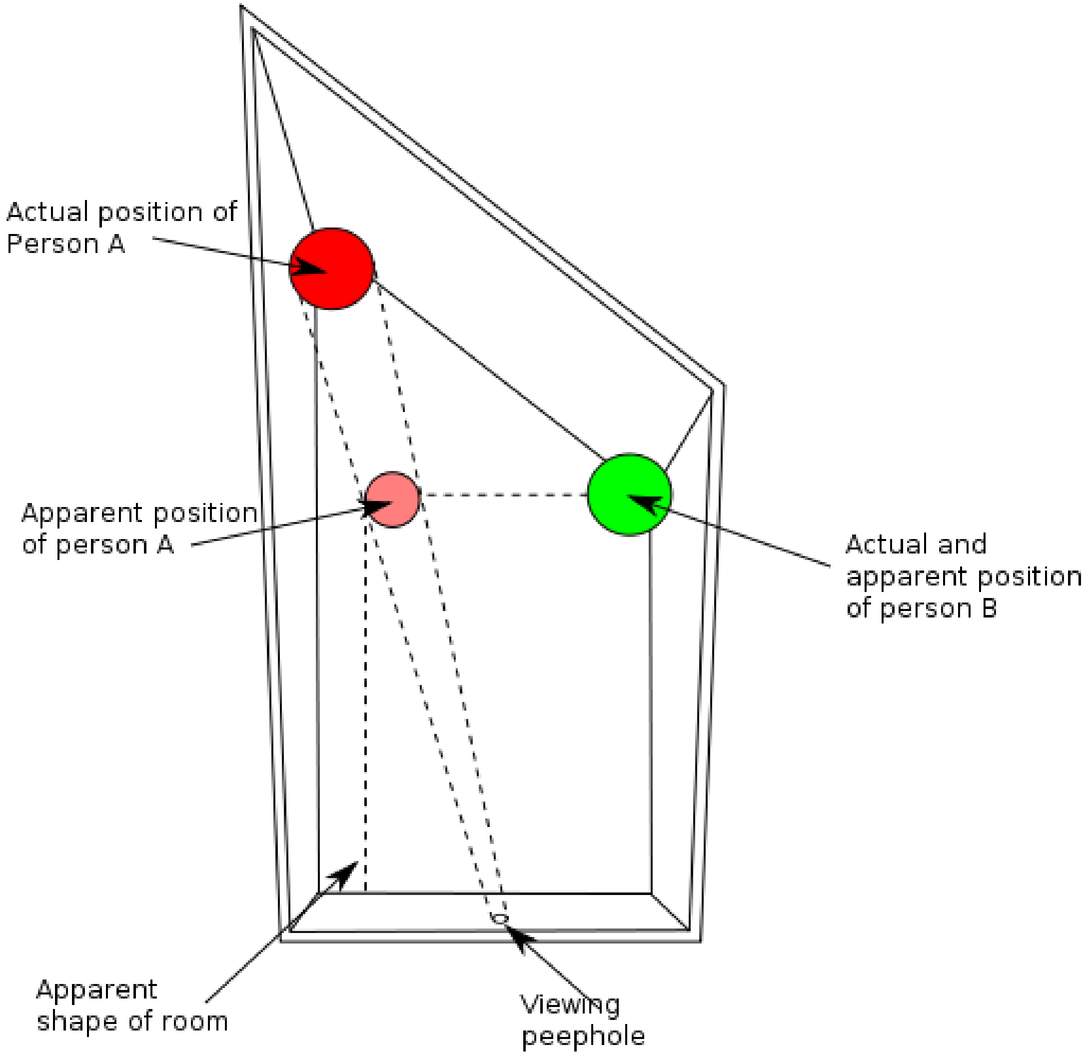

Magnification depends on depth

What happens as we change the focal length?

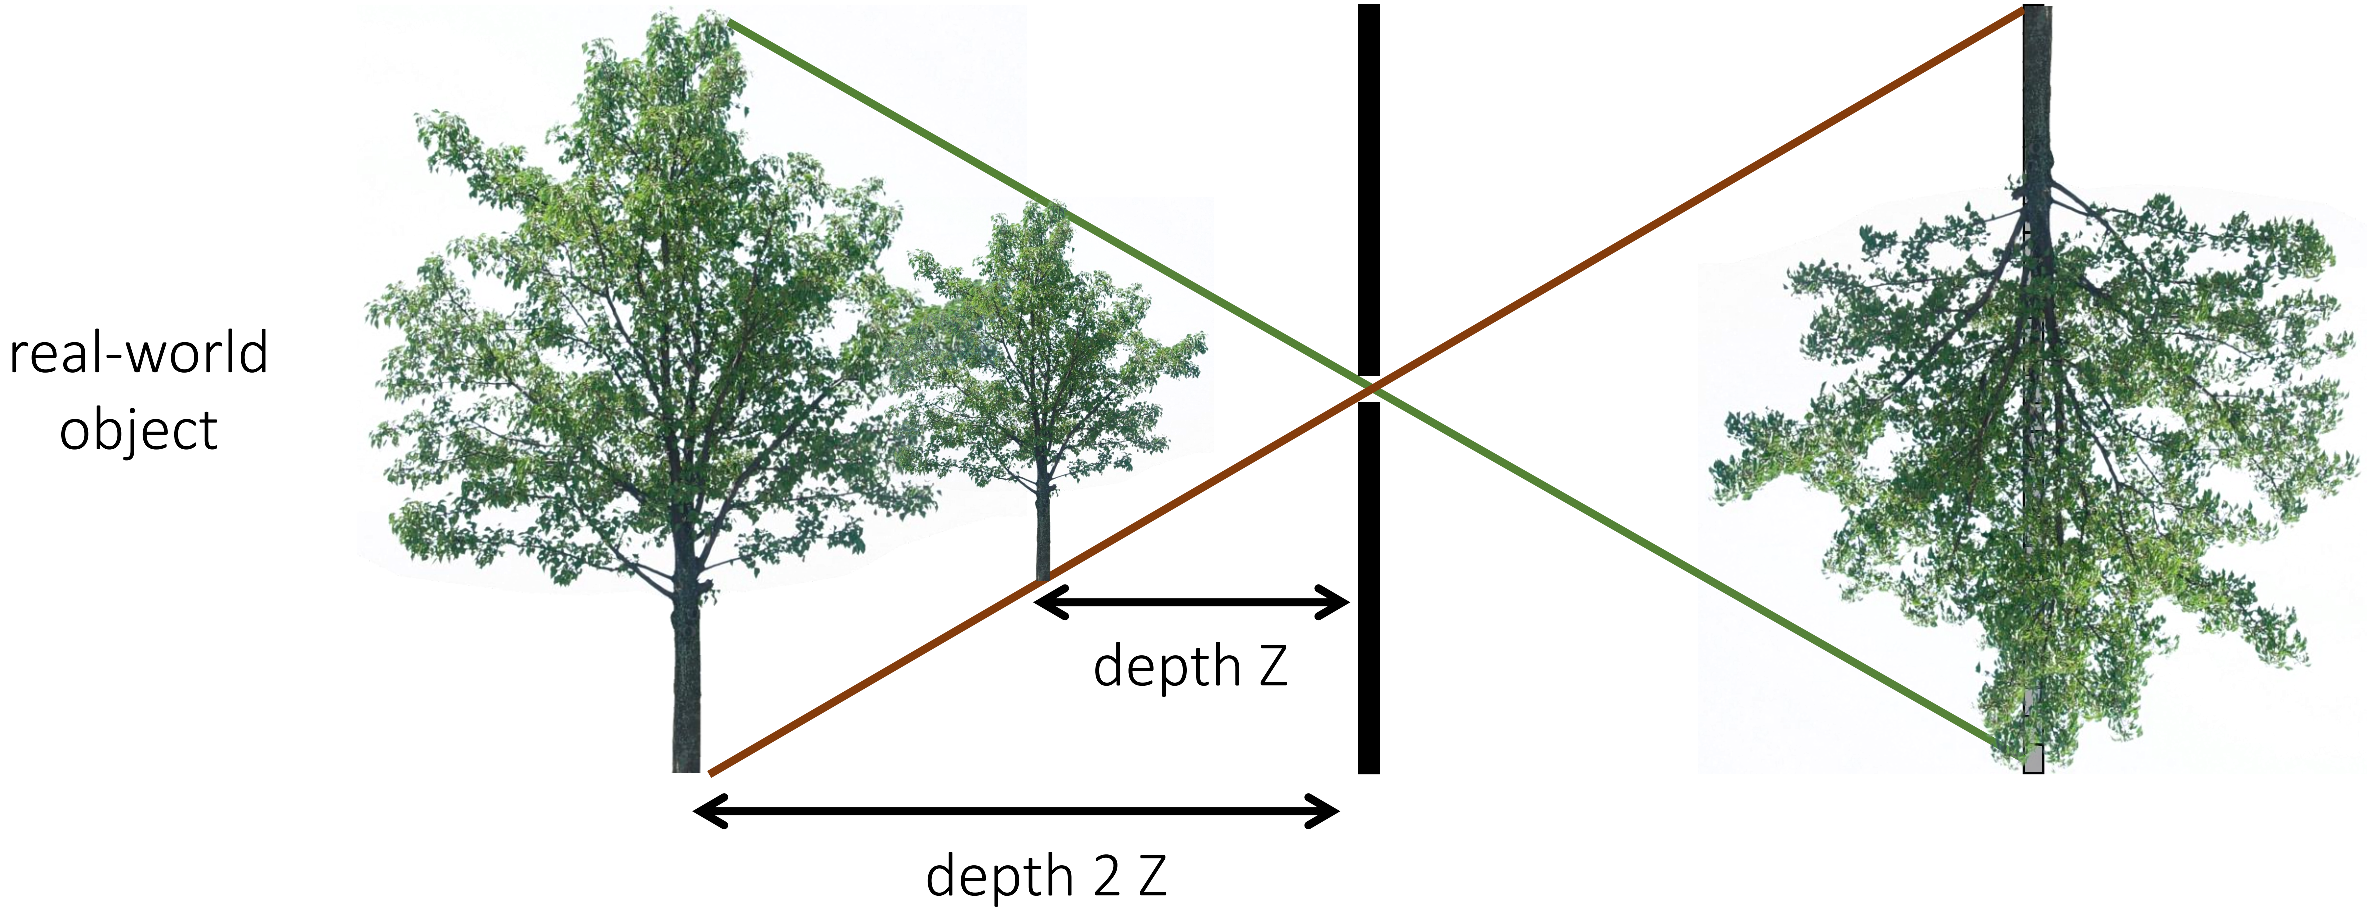

Magnification depends on focal length

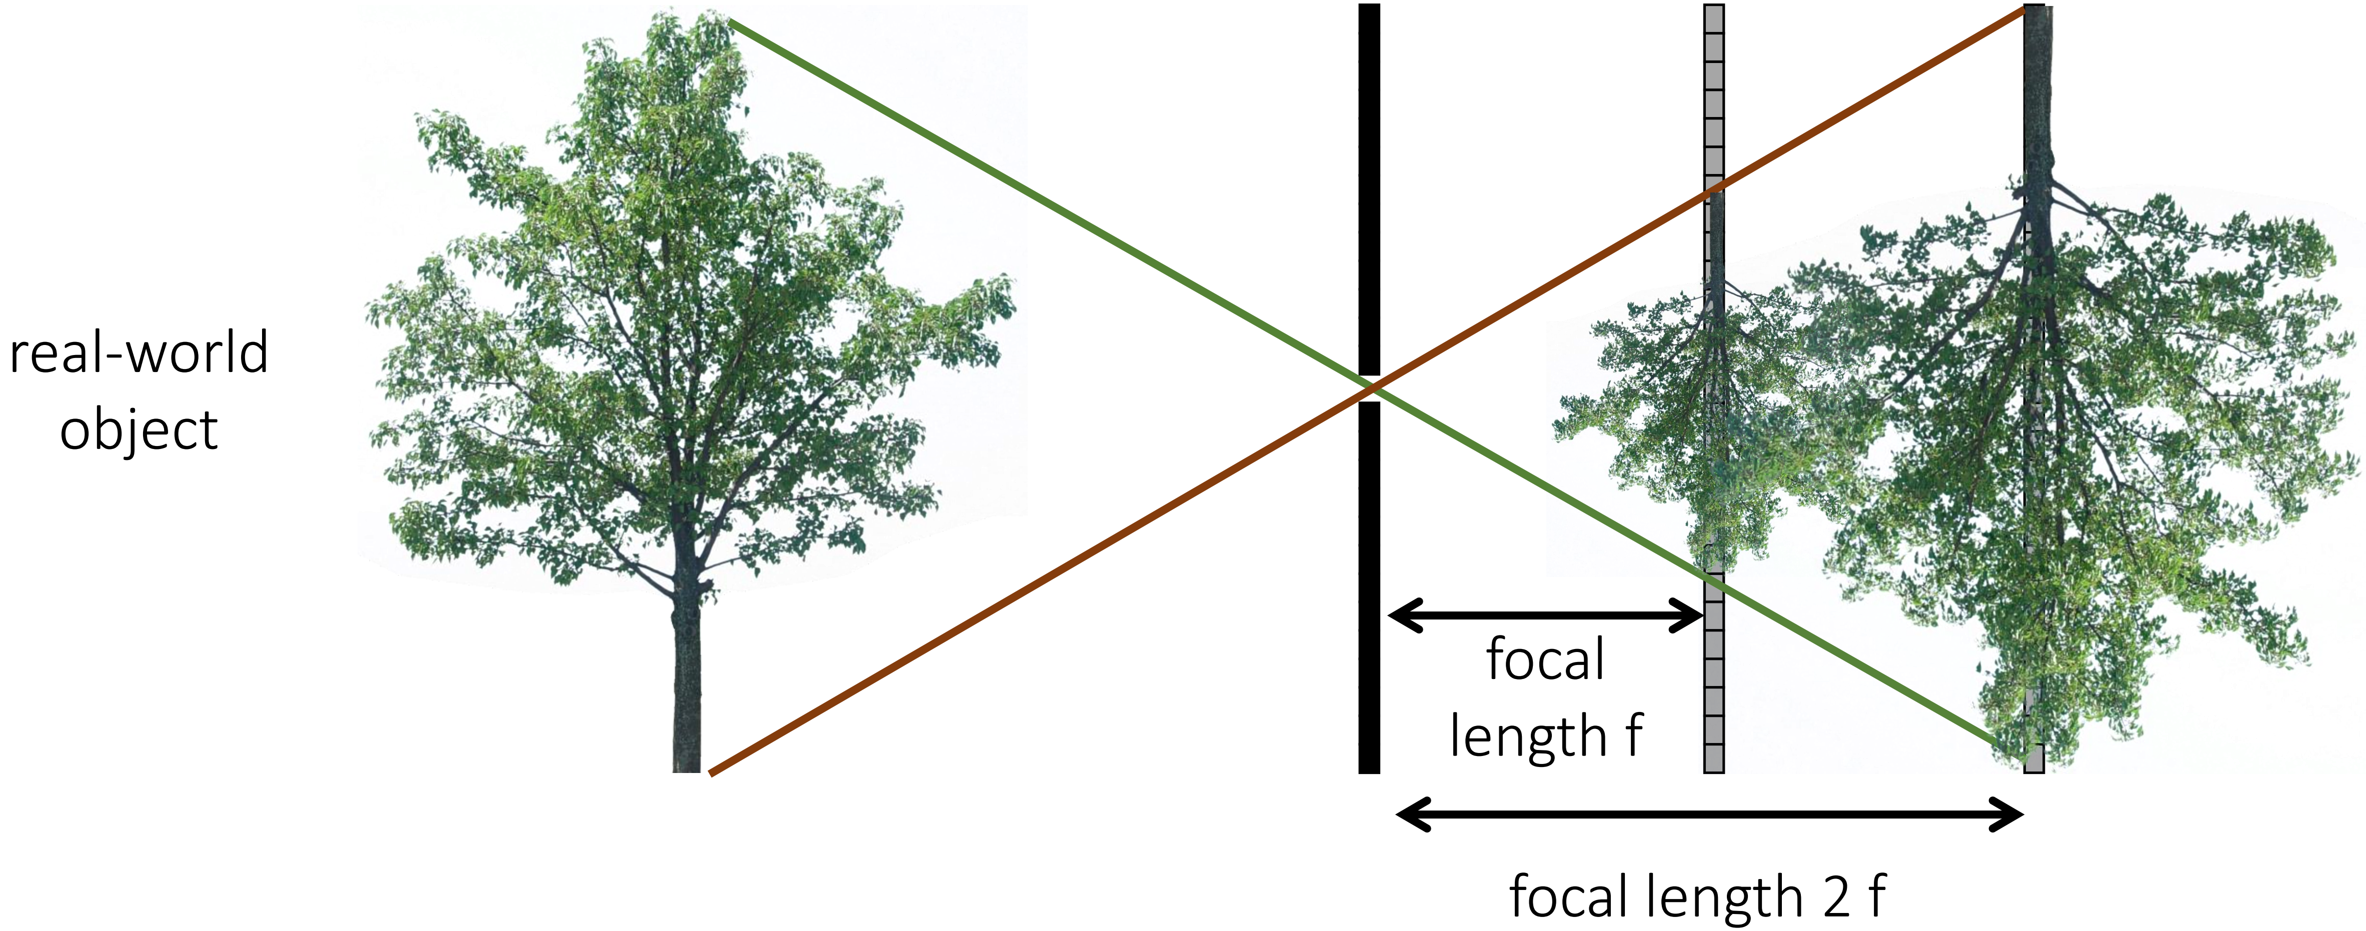

What if...

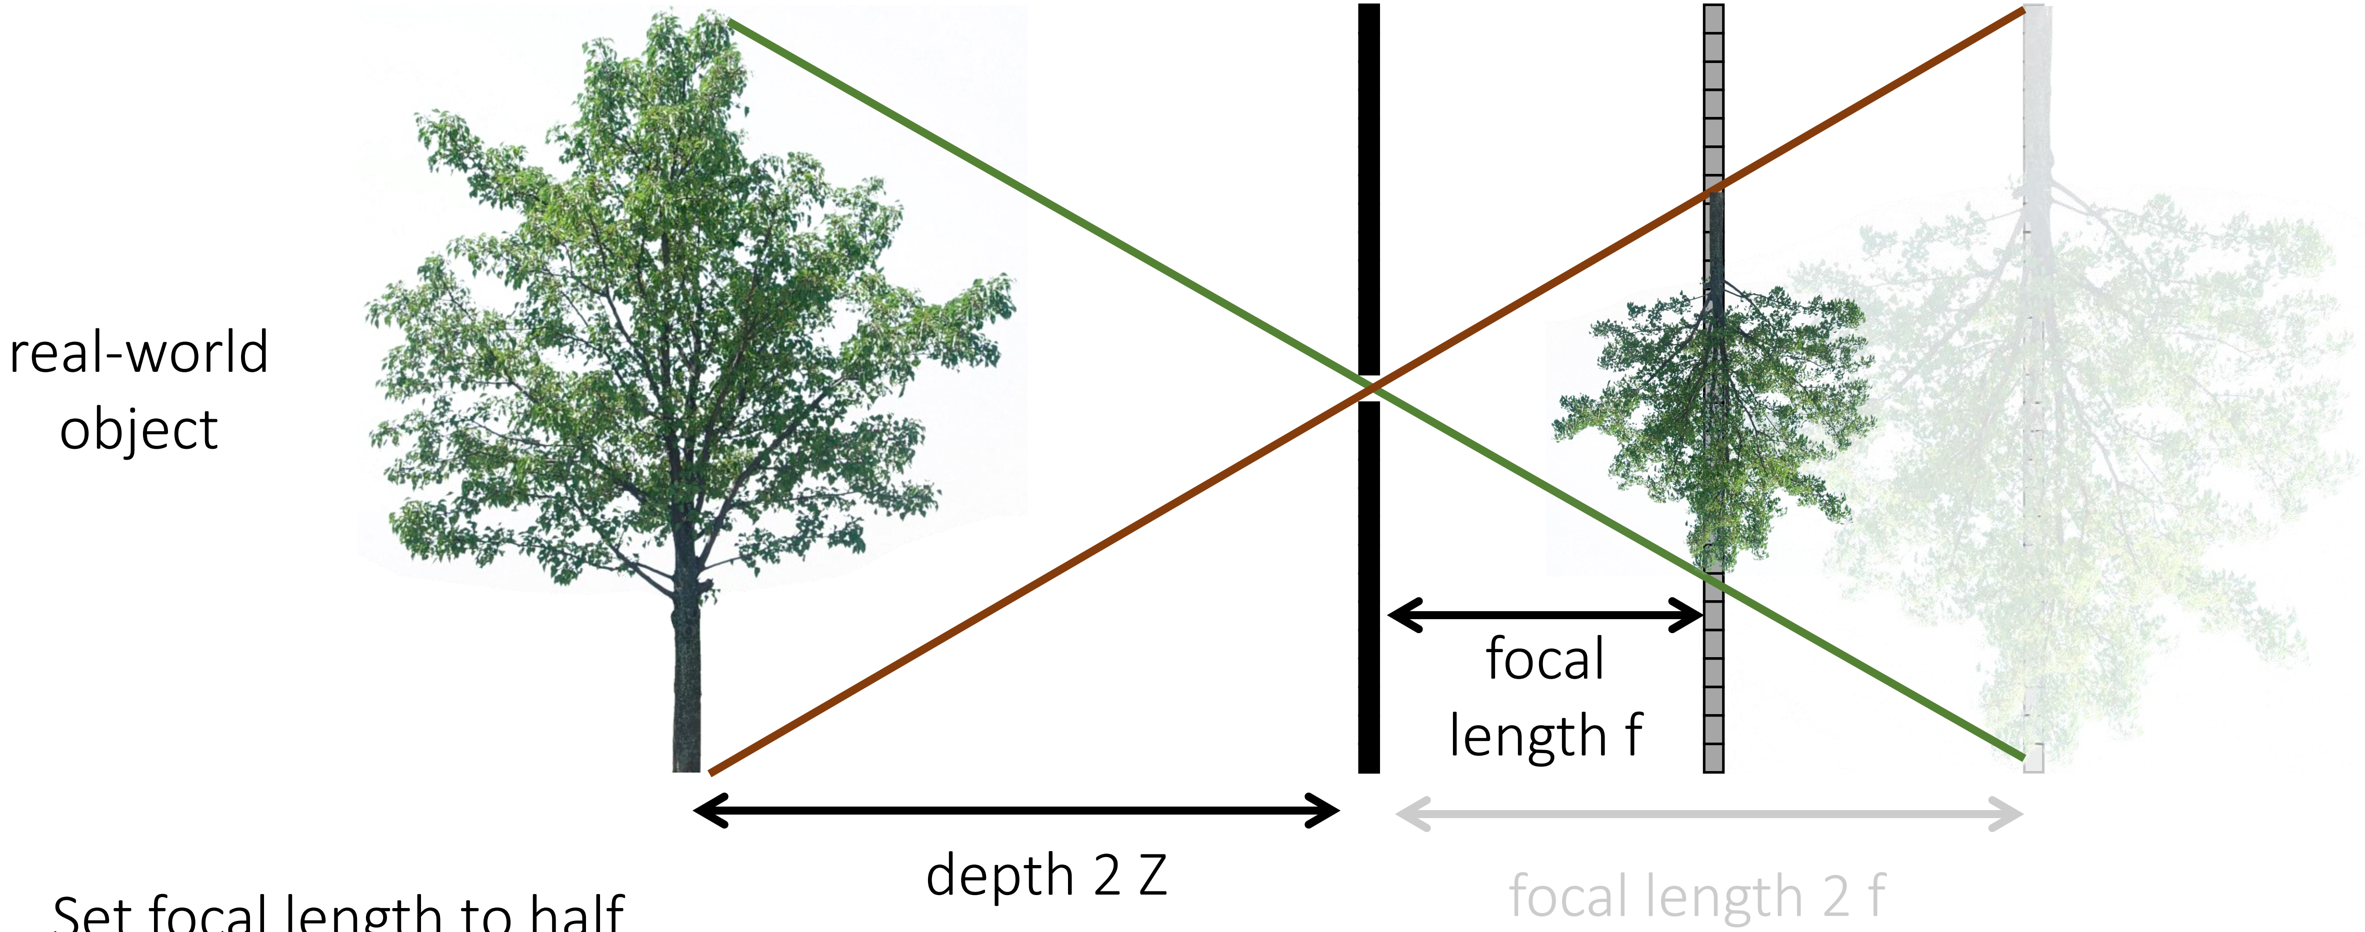

1. Set focal length to half

What if...

real-world
object

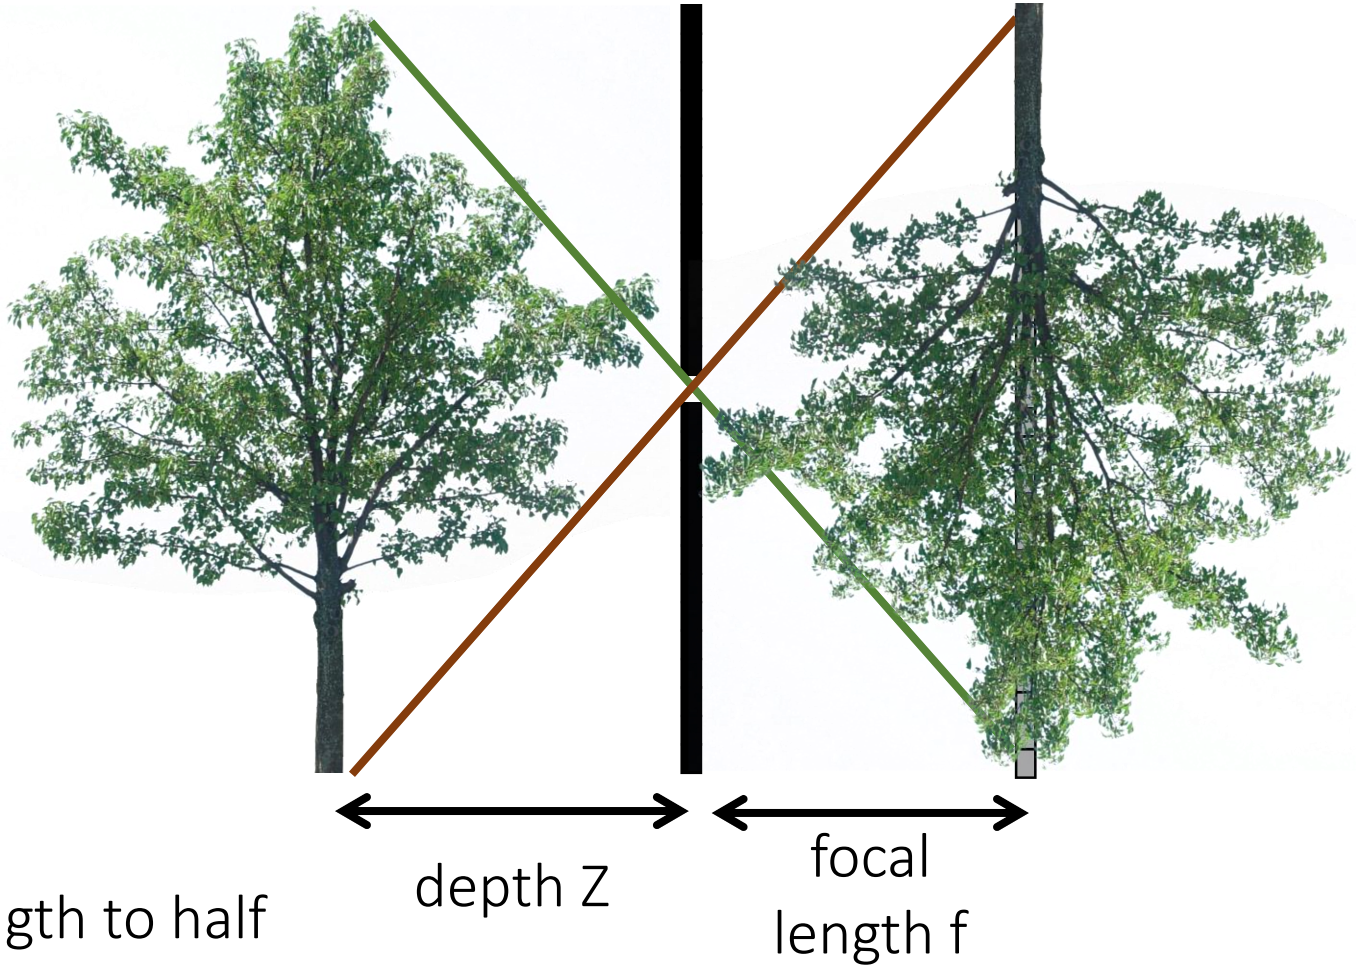

Is this the same image as
the one I had at focal
length 2f and distance 2Z?

1. Set focal length to half
2. Set depth to half

Perspective distortion

long focal length

mid focal length

short focal length

Perspective distortion

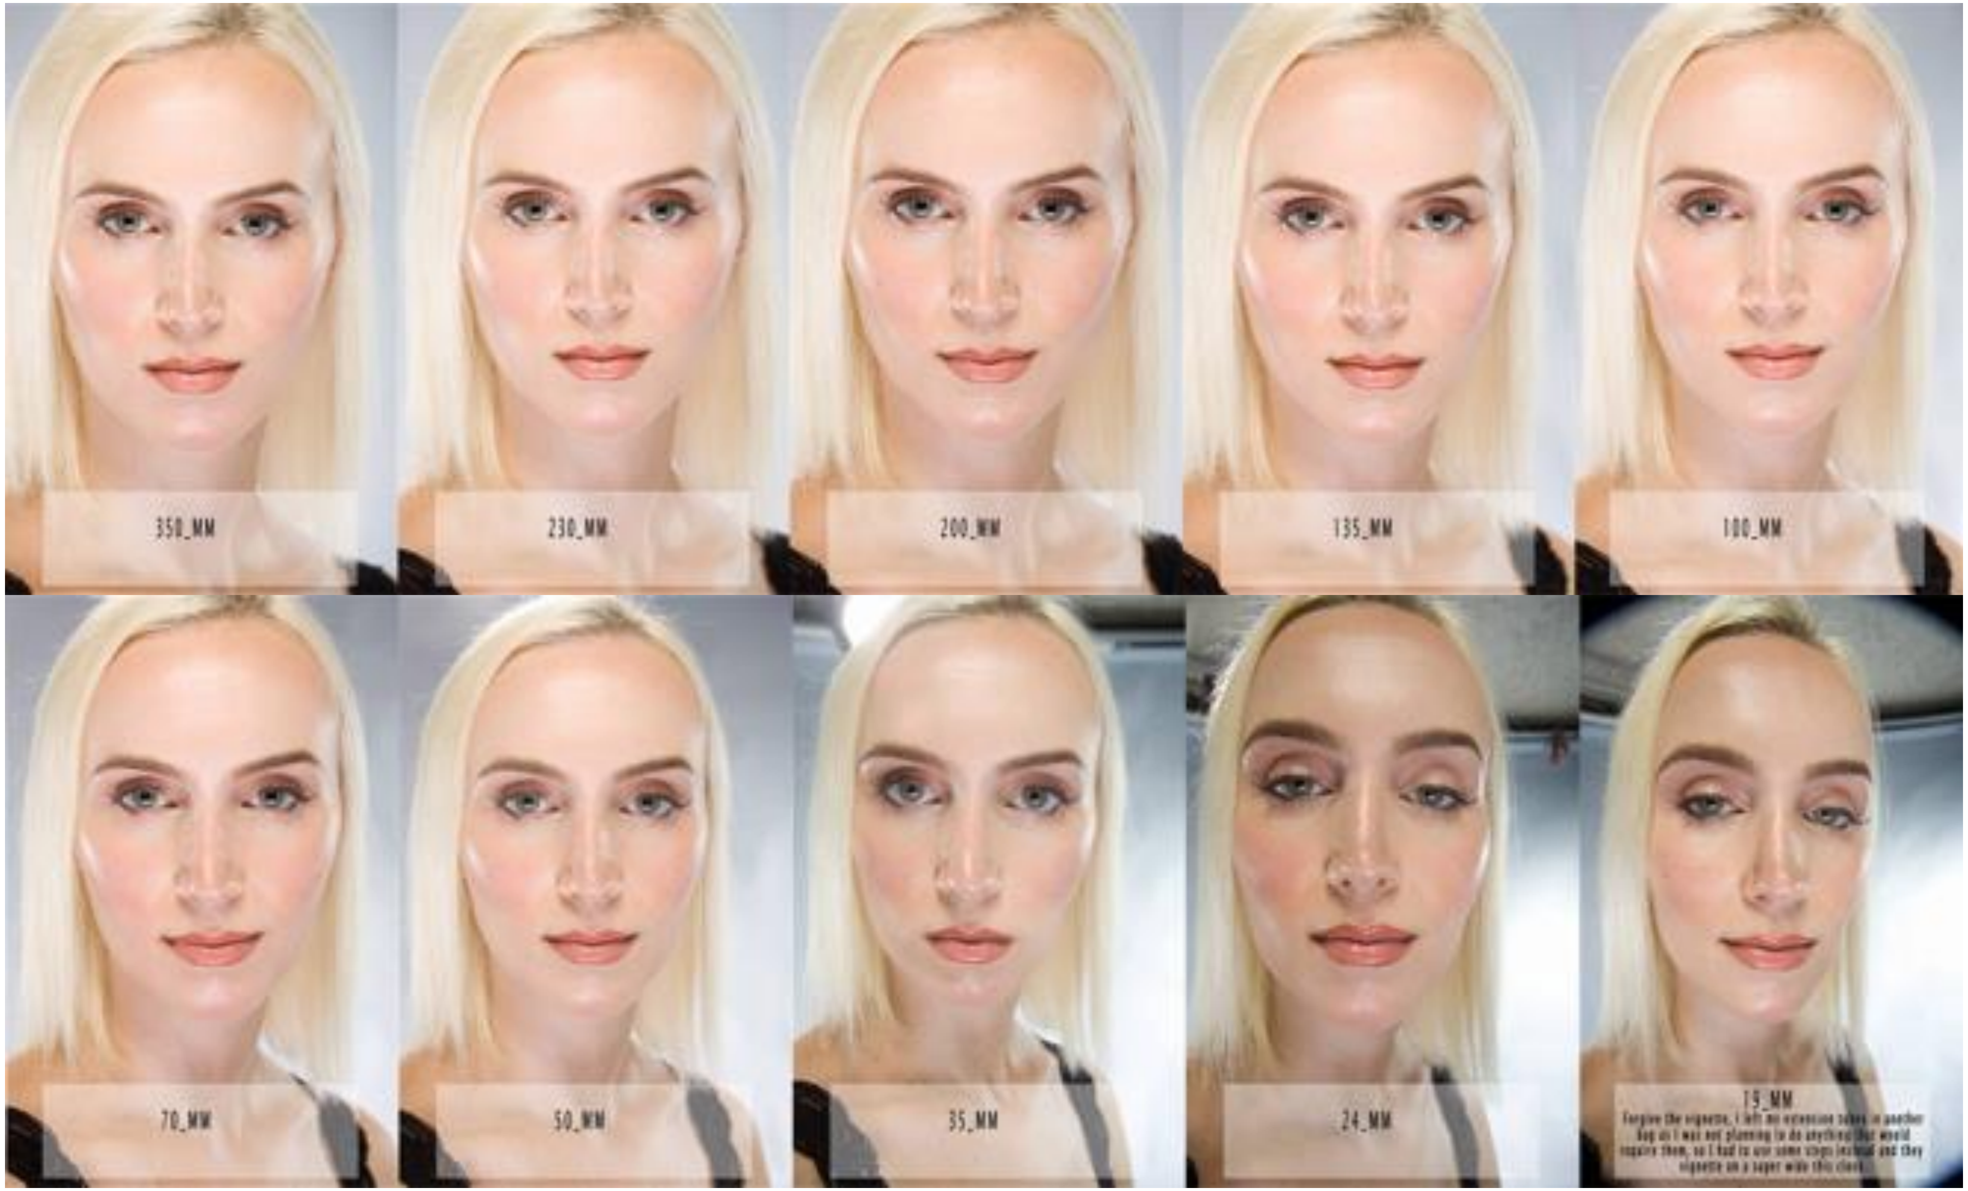

What is the best focal length for portraits?

That’s like asking which is better, vi or emacs…

long focal length

mid focal length

short focal length

Vertigo effect

Named after Alfred Hitchcock’s movie

- also known as “dolly zoom”

Vertigo effect

How would you create this effect?

Orthographic camera and telecentric lenses

What if...

real-world object

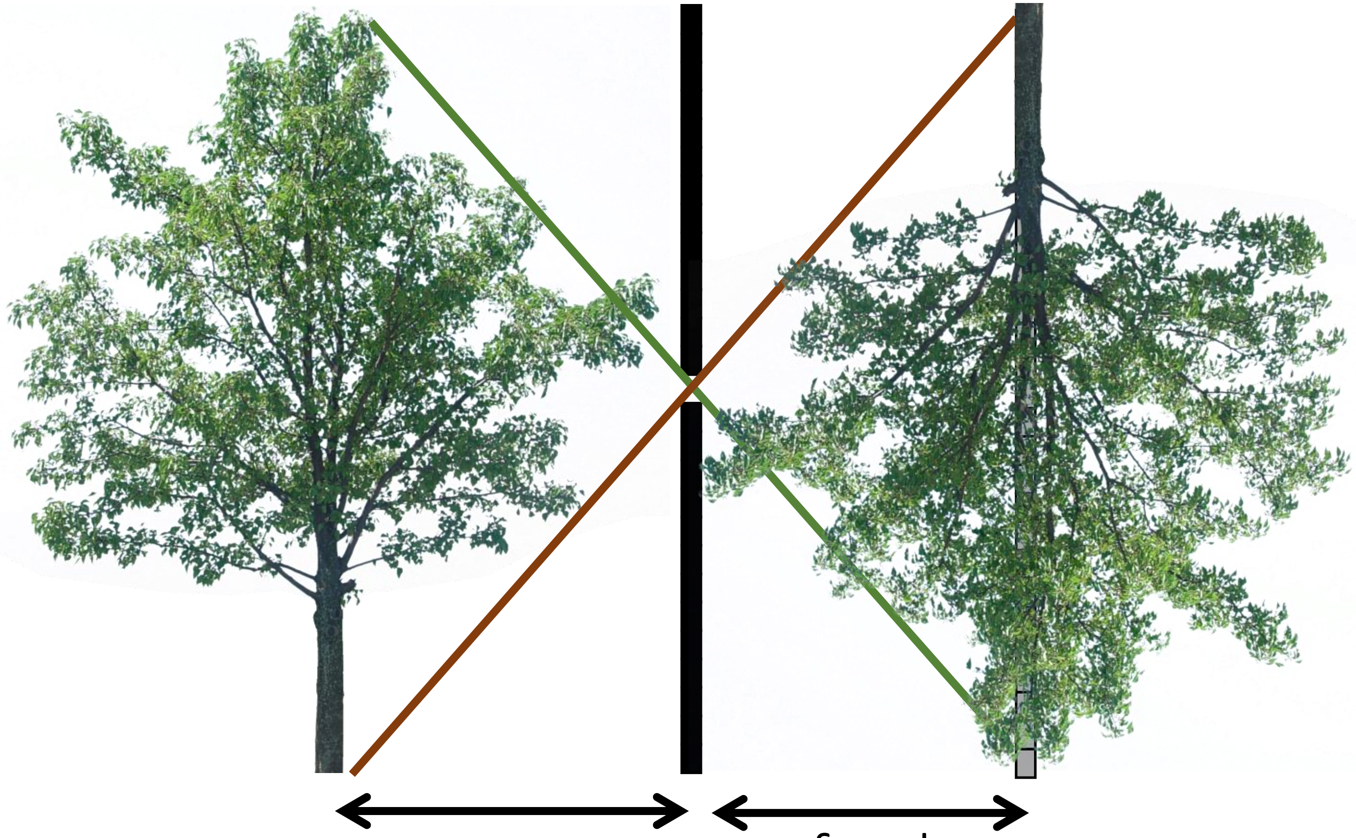

Continue increasing Z and f while maintaining same magnification?

$$f \to \infty \text{ and } \frac{f}{Z} = \text{constant}$$

The 2D view of the (rearranged) pinhole camera

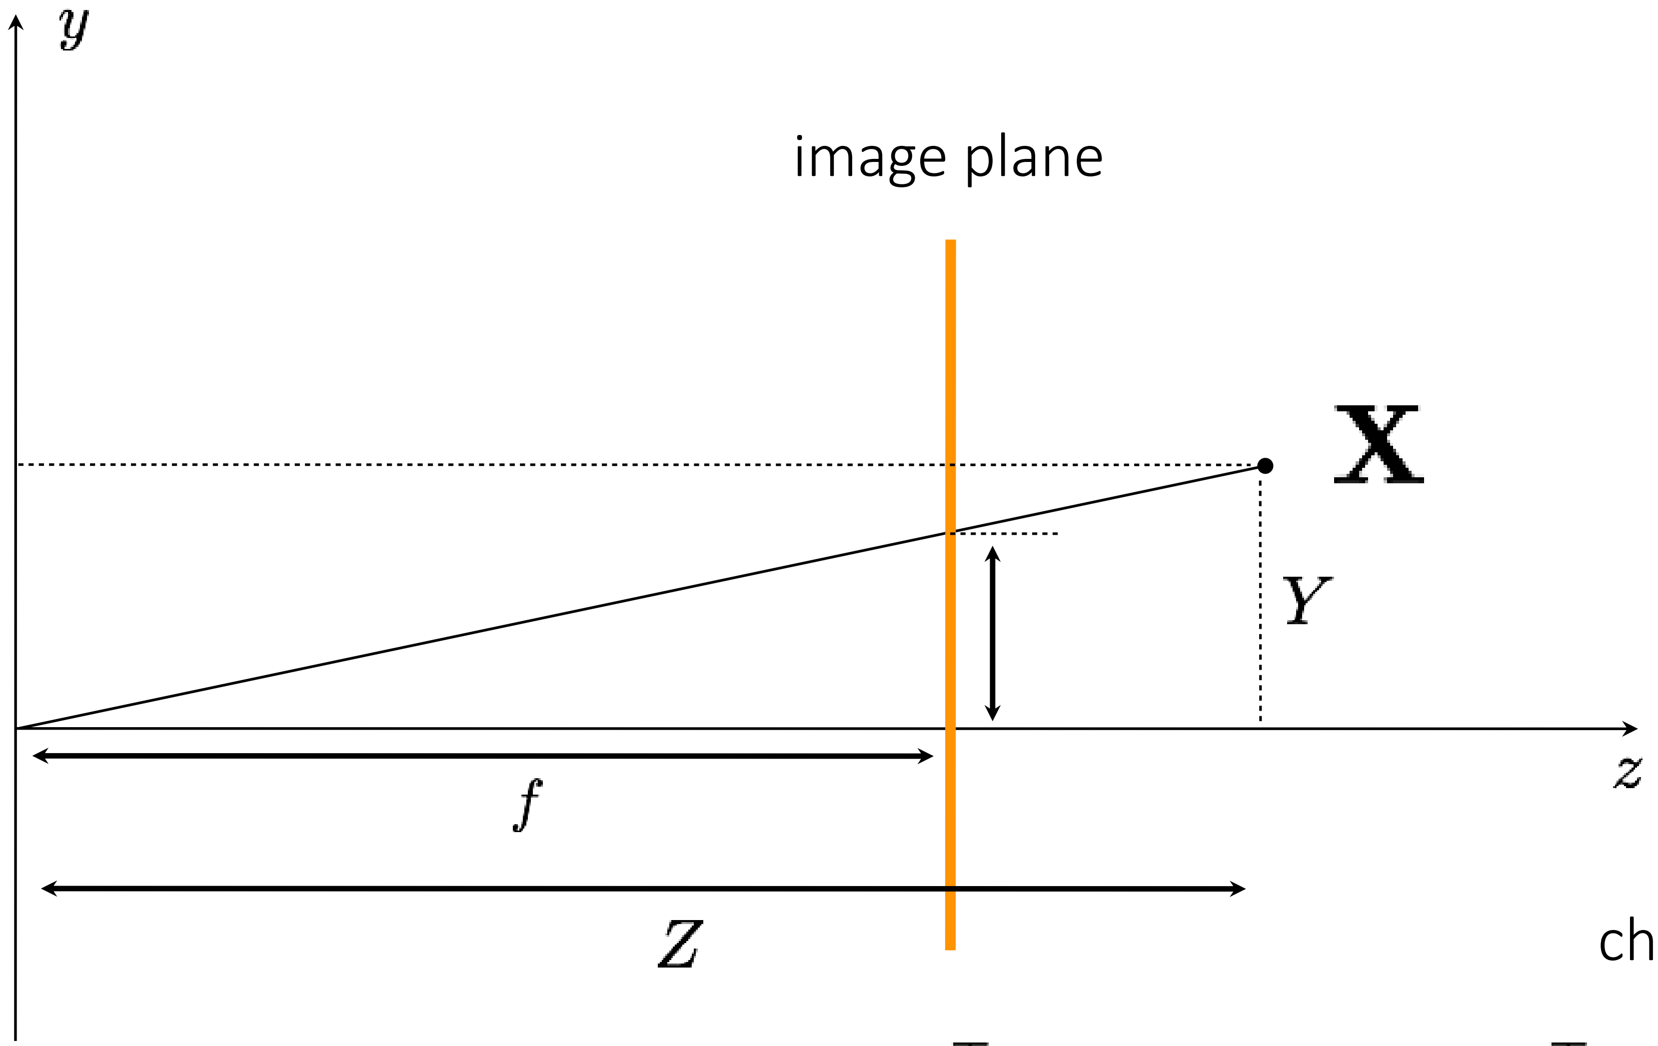

magnification changes with depth

$$[X \quad Y \quad Z]^\top \mapsto [fX/Z \quad fY/Z]^\top$$

Orthographic vs pinhole camera

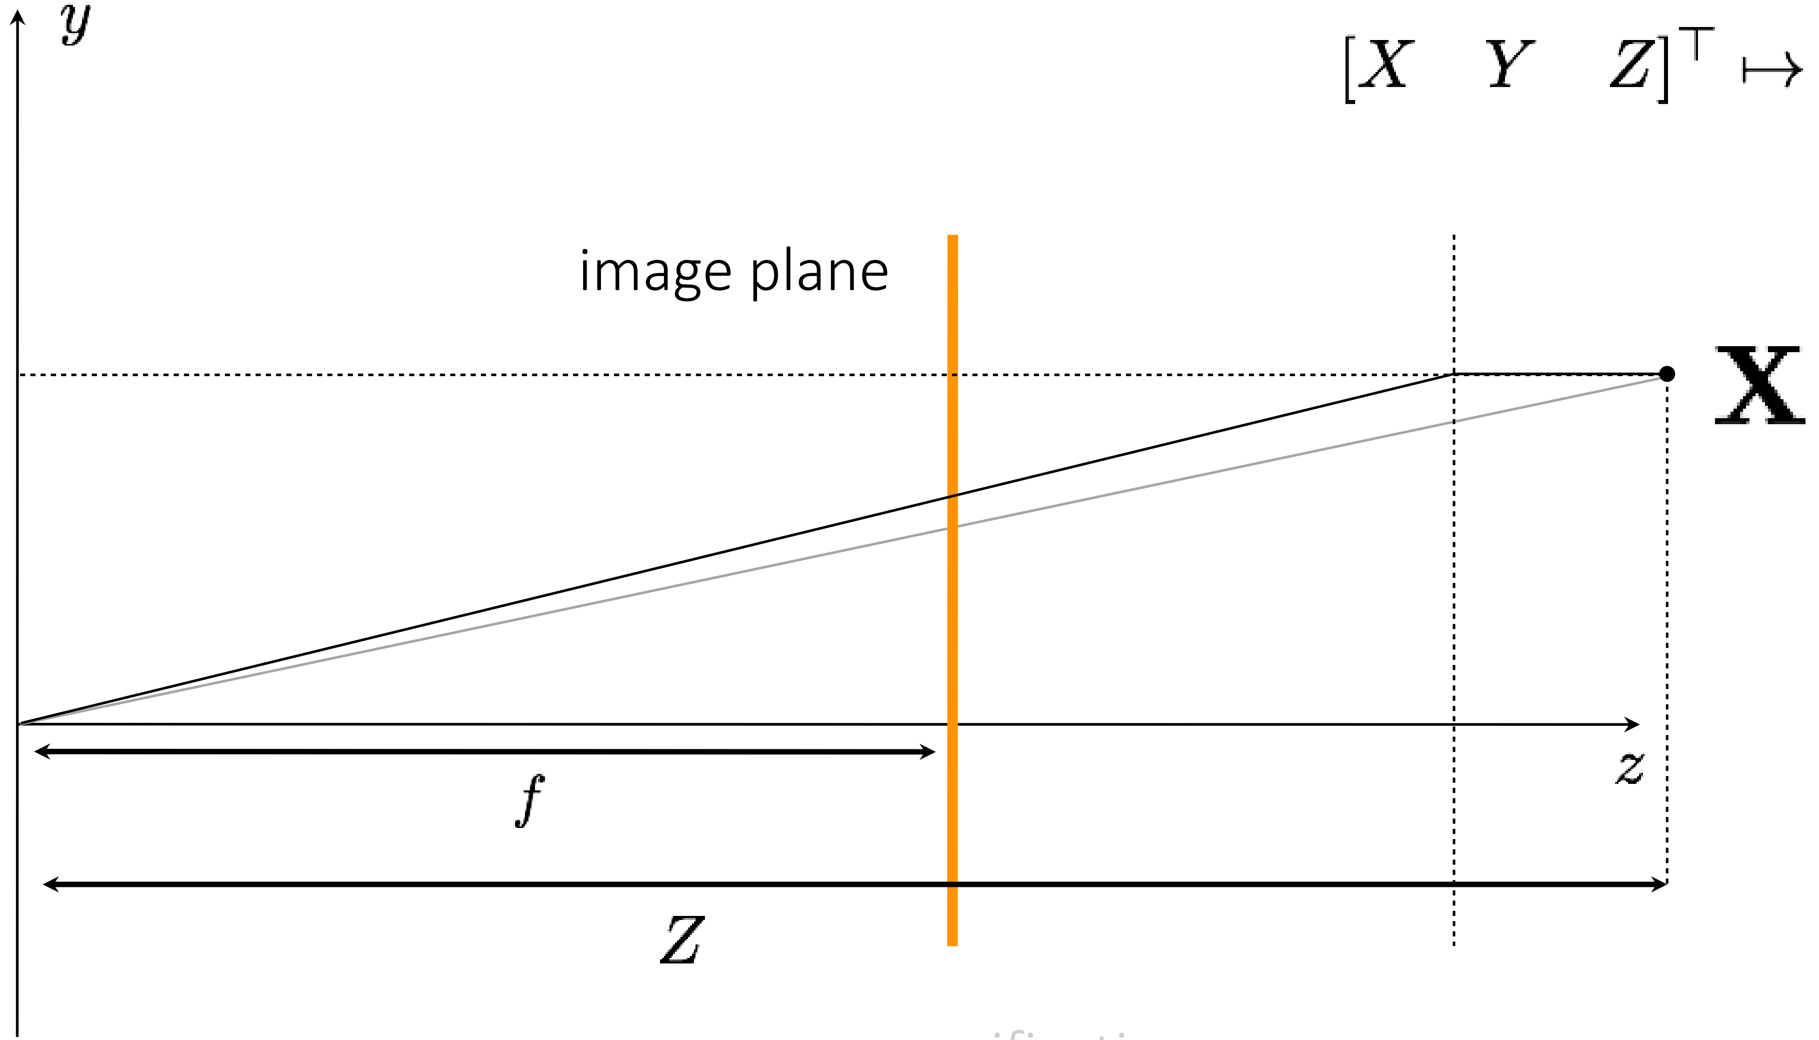

$$[X \quad Y \quad Z]^\top \mapsto [fX \quad fY]^\top$$

magnification does not change with depth

magnification changes with depth

$$[X \quad Y \quad Z]^\top \mapsto [fX/Z \quad fY/Z]^\top$$

How can we implement such a camera with lenses?

Telecentric lens

Place a pinhole at focal length, so that only rays parallel to primary ray pass through.

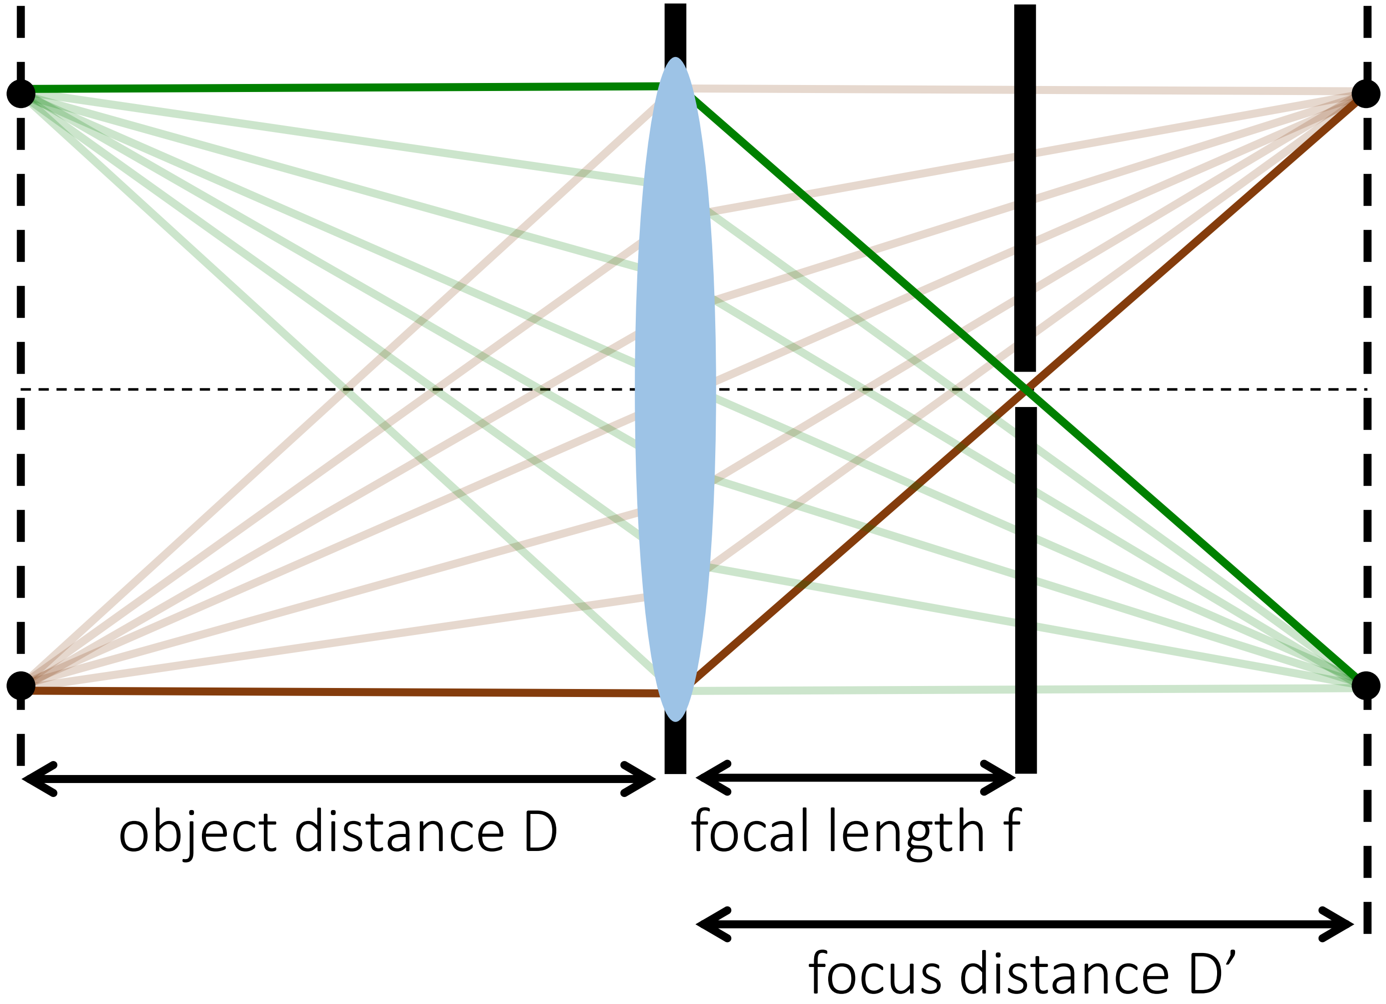

Regular vs telecentric lens

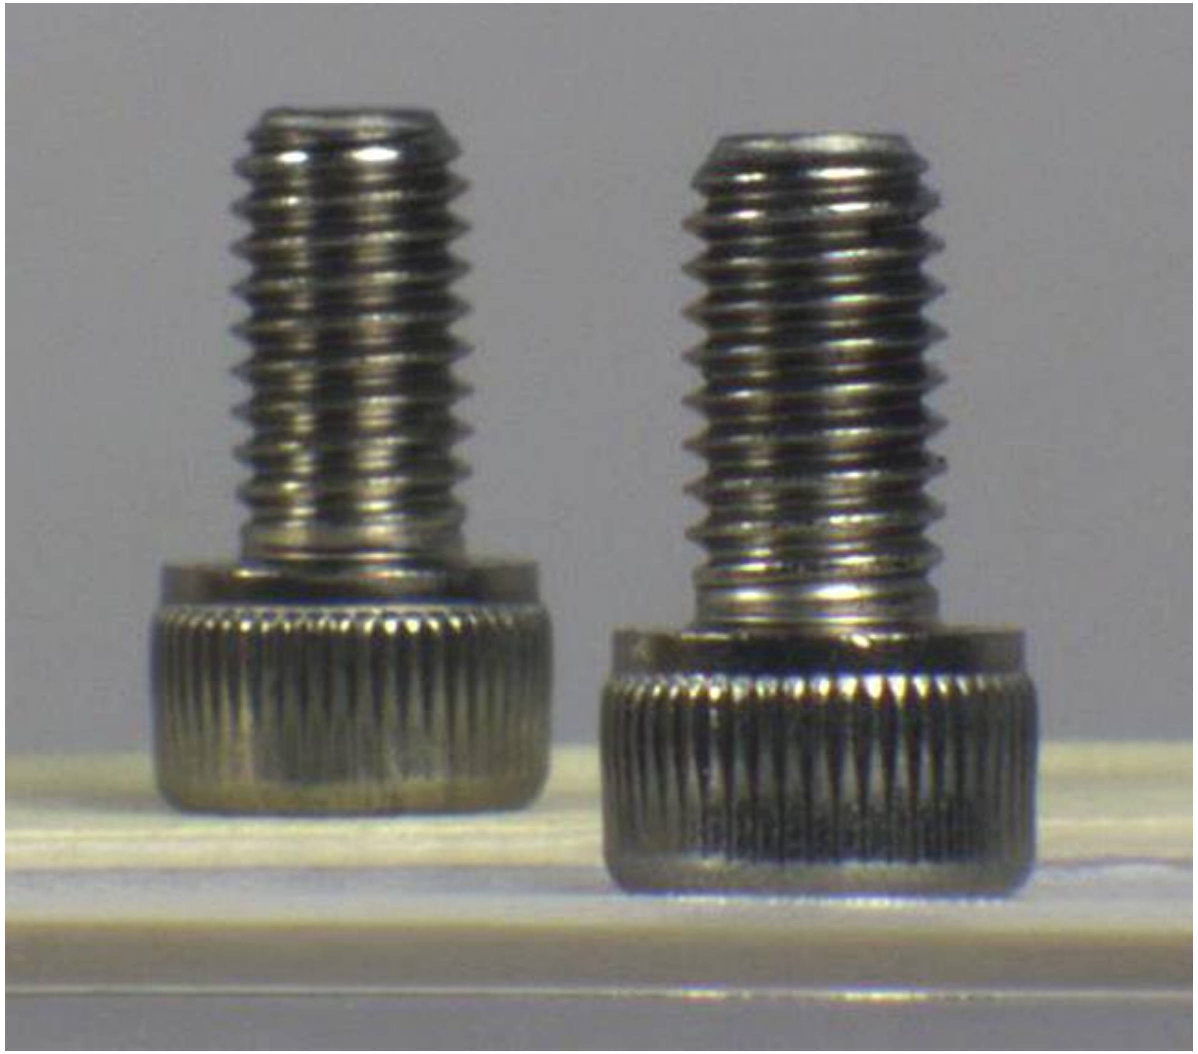

regular lens

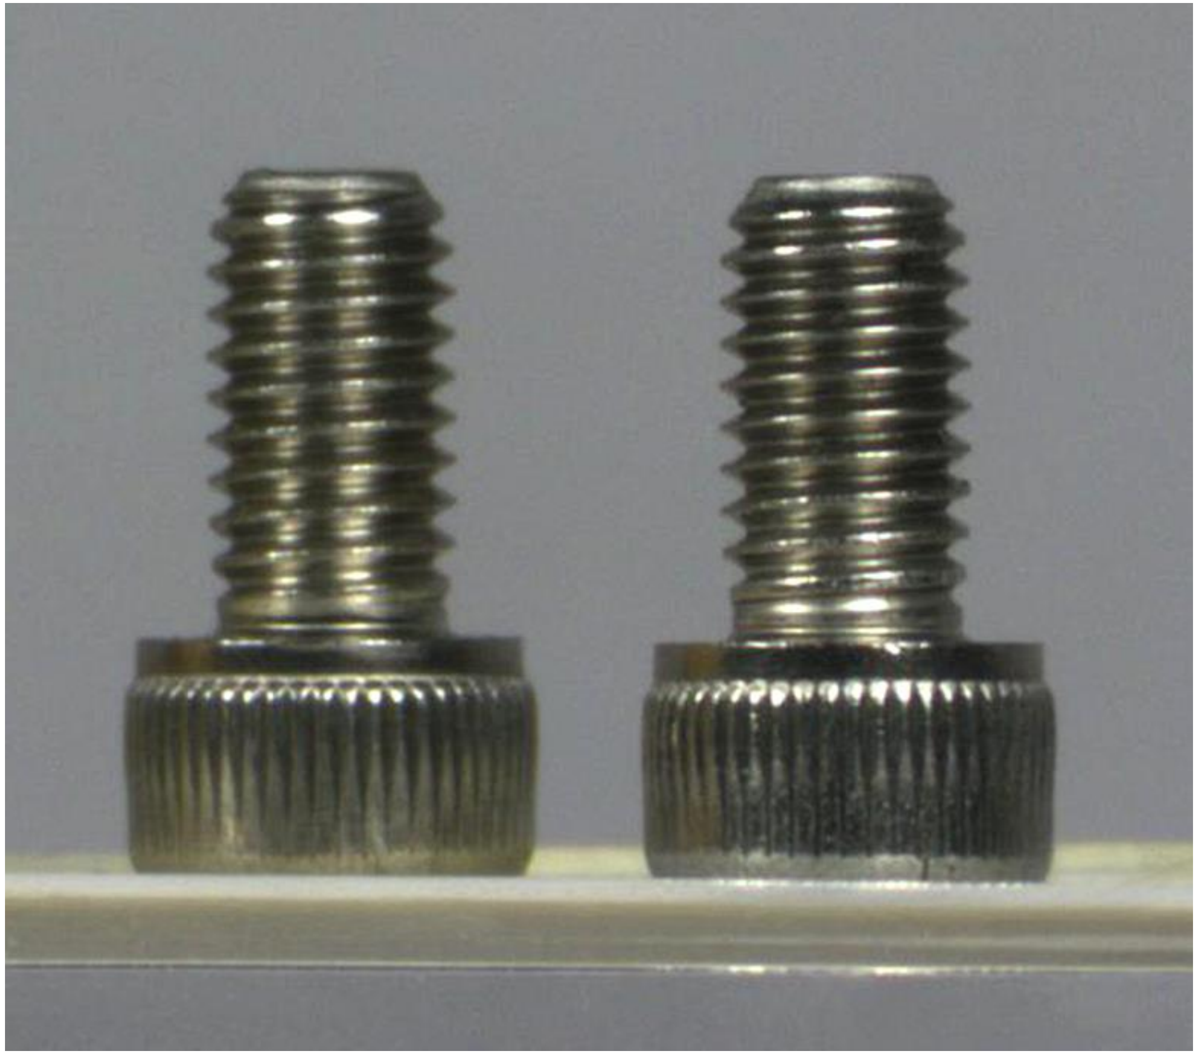

telecentric lens

References

Basic reading:

- Szeliski textbook, Section 2.1.5, 2.2.3.

Additional reading:

- Hartley and Zisserman, “Multiple View Geometry in Computer Vision,” Cambridge University Press 2004.
  Chapter 6 of this book is a very thorough treatment of camera models.
- Goodman, “Introduction to Fourier Optics,” W.H. Freeman 2004.
  The standard reference on Fourier optics, chapter 4 covers aperture diffraction.
- Ray, “Applied Photographic Optics,” Focal Press 2002.
  A great book covering everything about photographic optics.
- Torralba and Freeman, “Accidental Pinhole and Pinspeck Cameras,” CVPR 2012.
  The eponymous paper discussed in the slides.