Pinholes and lenses

15-463, 15-663, 15-862
Computational Photography
Fall 2024, Lecture 3

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

Course announcements

- Homework assignment 1 is out.
  - Due September 13th.
  - Start early! Second part takes a lot of time and requires considerable handiwork.
  - You can get some (few) boxes from my office.
  - Any issues with homework assignment 1?

- Office hours for rest of the semester:
  - Wednesday 3 – 5 pm, Yannis.
  - Thursday 3 – 5 pm, Dorian.

- Camera distribution this week: drop by my office during office hours today, and whenever is convenient on Thursday and Friday.

Go to this talk on Thursday!

Details on Slack.

Carnegie Mellon Graphics Colloquium

Thursday, 5 September 2024
4:30–5:30pm
Rashid Auditorium, Gates Hillman 4401

Sampling and Signal-Processing for High-Dimensional Visual Appearance in Computer Graphics and Vision

Ravi Ramamoorthi
Ronald L. Graham Professor of Computer Science, UC San Diego

Many problems in computer graphics and vision, such as acquiring images of a scene to enable synthesis of novel views from many directions for virtual reality, computing realistic images by integrating lighting from many different incident directions across a range of scene pixels and viewing angles, or acquiring and modeling the appearance of realistic materials like fur or skin, require sampling and signal-processing on high-dimensional visual appearance spaces involving changes in lighting, viewpoint, spatial location and other parameters. Over my career, my group has developed a number of novel mathematical and signal-processing tools to address these challenges, significantly reducing the cost of acquisition and computation. In this talk, we describe significant theoretical and practical advances in real-time high quality precomputed rendering, Monte Carlo rendering with orders of magnitude fewer samples, and realistic novel view synthesis. In all cases, the methods are now widely deployed in production, and we discuss new computational and signal-processing tools we have developed, including reflection as convolution, sheared and multiple axis-aligned filtering, plenoptic light field sampling and neural radiance fields.

Bio: Ravi Ramamoorthi is the Ronald L. Graham Professor of Computer Science at UCSD and founding director of the UC San Diego Center for Visual Computing. He earlier held tenured faculty positions at UC Berkeley and Columbia University, in all of which he played a key leadership role in building multi-faculty research groups recognized as leaders in computer vision and graphics. He has authored more than 200 refereed publications in computer graphics and vision, including 100+ ACM SIGGRAPH/TOG papers. He has consulted with Pixar and startups in computational imaging, and currently holds a part-time appointment as a Distinguished Research Scientist at NVIDIA. Prof. Ramamoorthi has received about twenty major honors including the ACM SIGGRAPH Significant New Researcher Award for his research in computer graphics, and the Presidential Early Career Award for Scientists and Engineers for his work on physics-based computer vision. He is a fellow of IEEE, ACM and the SIGGRAPH Academy, received two inaugural Frontiers of Science Awards, and has twice been honored with the edX Prize certificate for exceptional contributions in online teaching and learning. He has graduated more than 30 postdoctoral and Ph.D. students, whose theses have been recognized by the ACM Dissertation Award honorable mention, the SIGGRAPH outstanding dissertation award and the UCSD Chancellor's Dissertation Medal.

Web Page: https://cseweb.ucsd.edu/~ravir/

Carnegie Mellon GRAPHICS

The Carnegie Mellon Graphics Colloquium is hosted by the Carnegie Mellon Graphics Lab and supported by Meta and Adobe.
→ http://graphics.cs.cmu.edu

Overview of today's lecture

- 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)

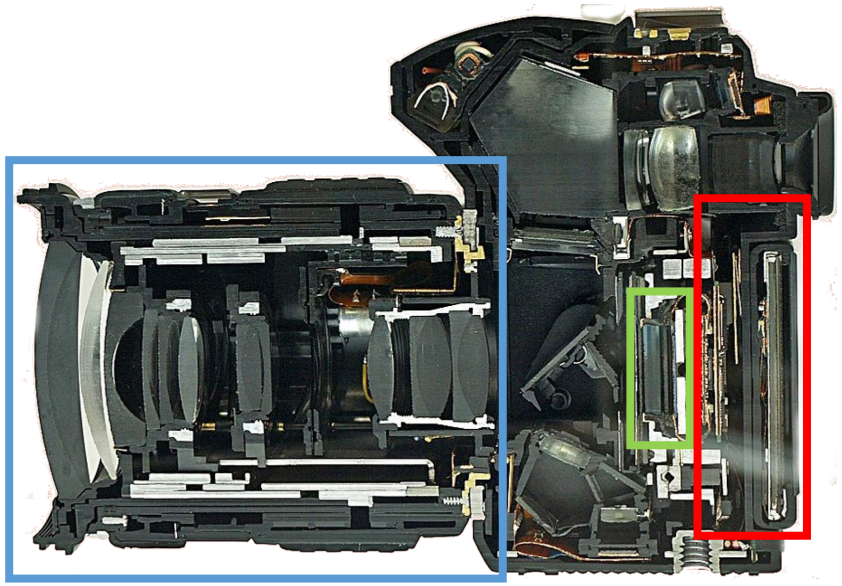

optics and optical controls
(lectures 2-3, 11-20)

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

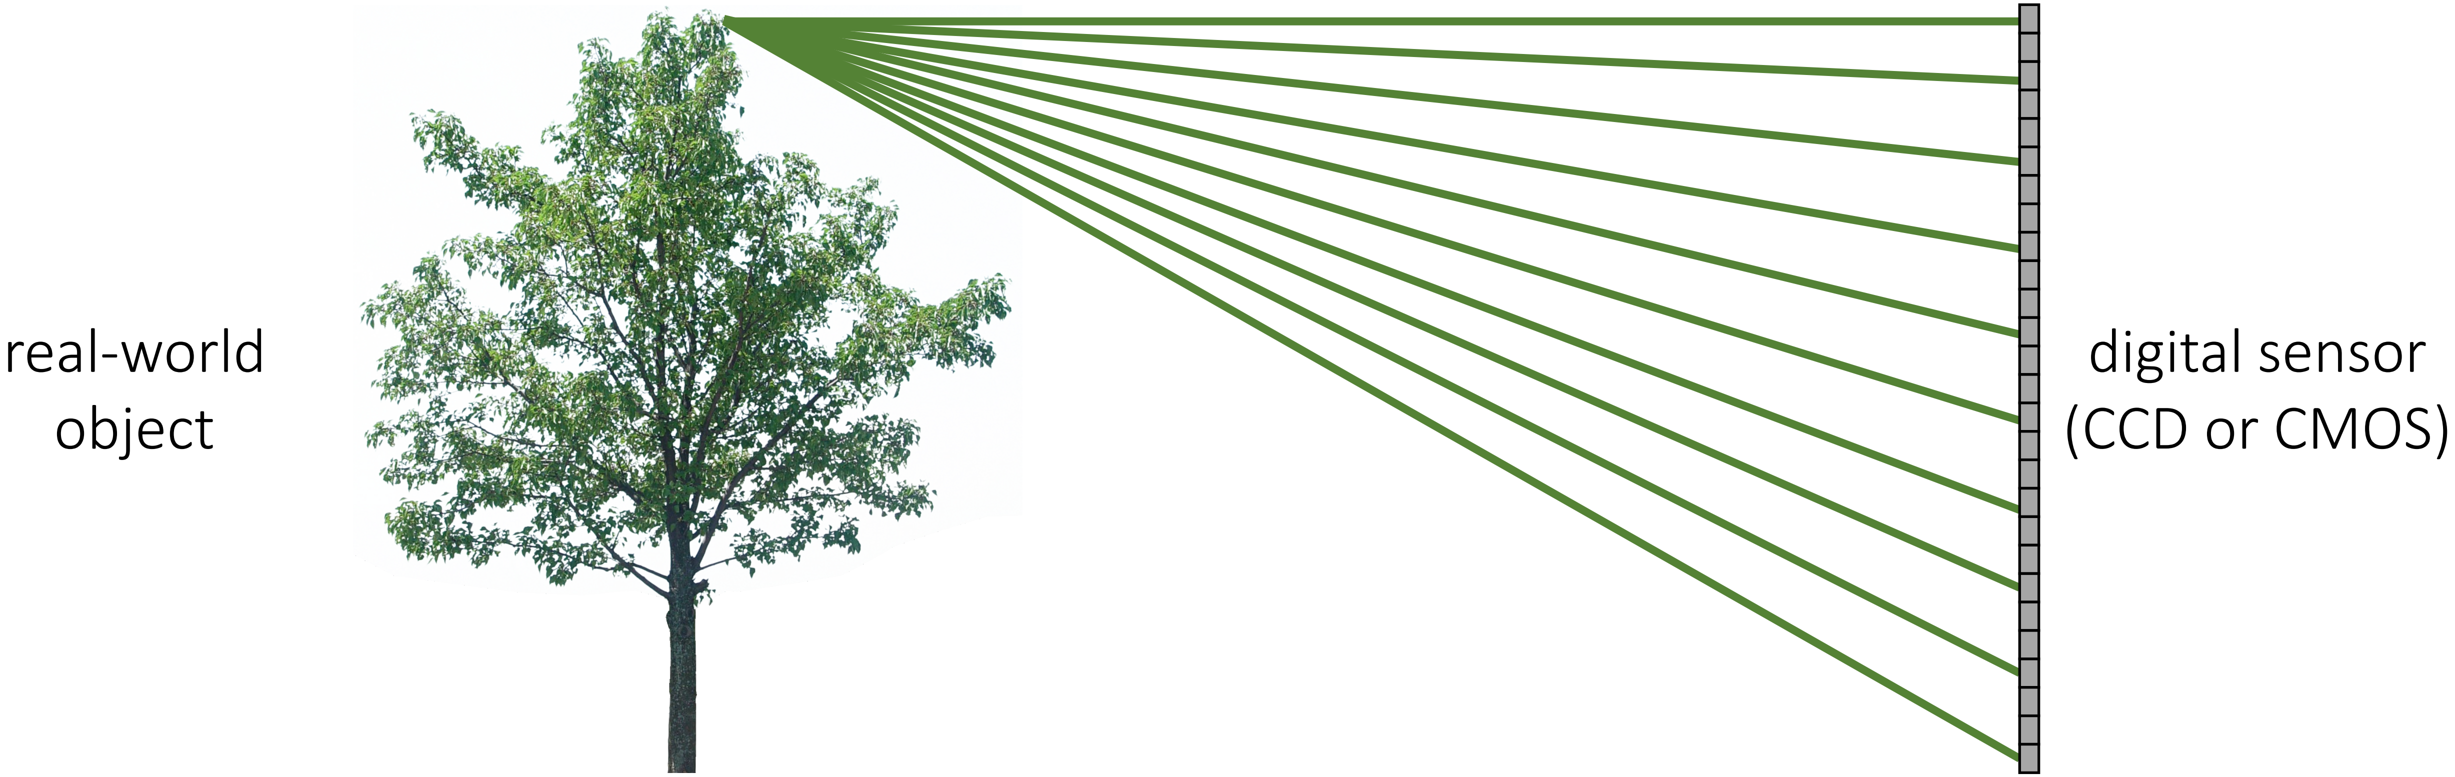

Bare-sensor imaging

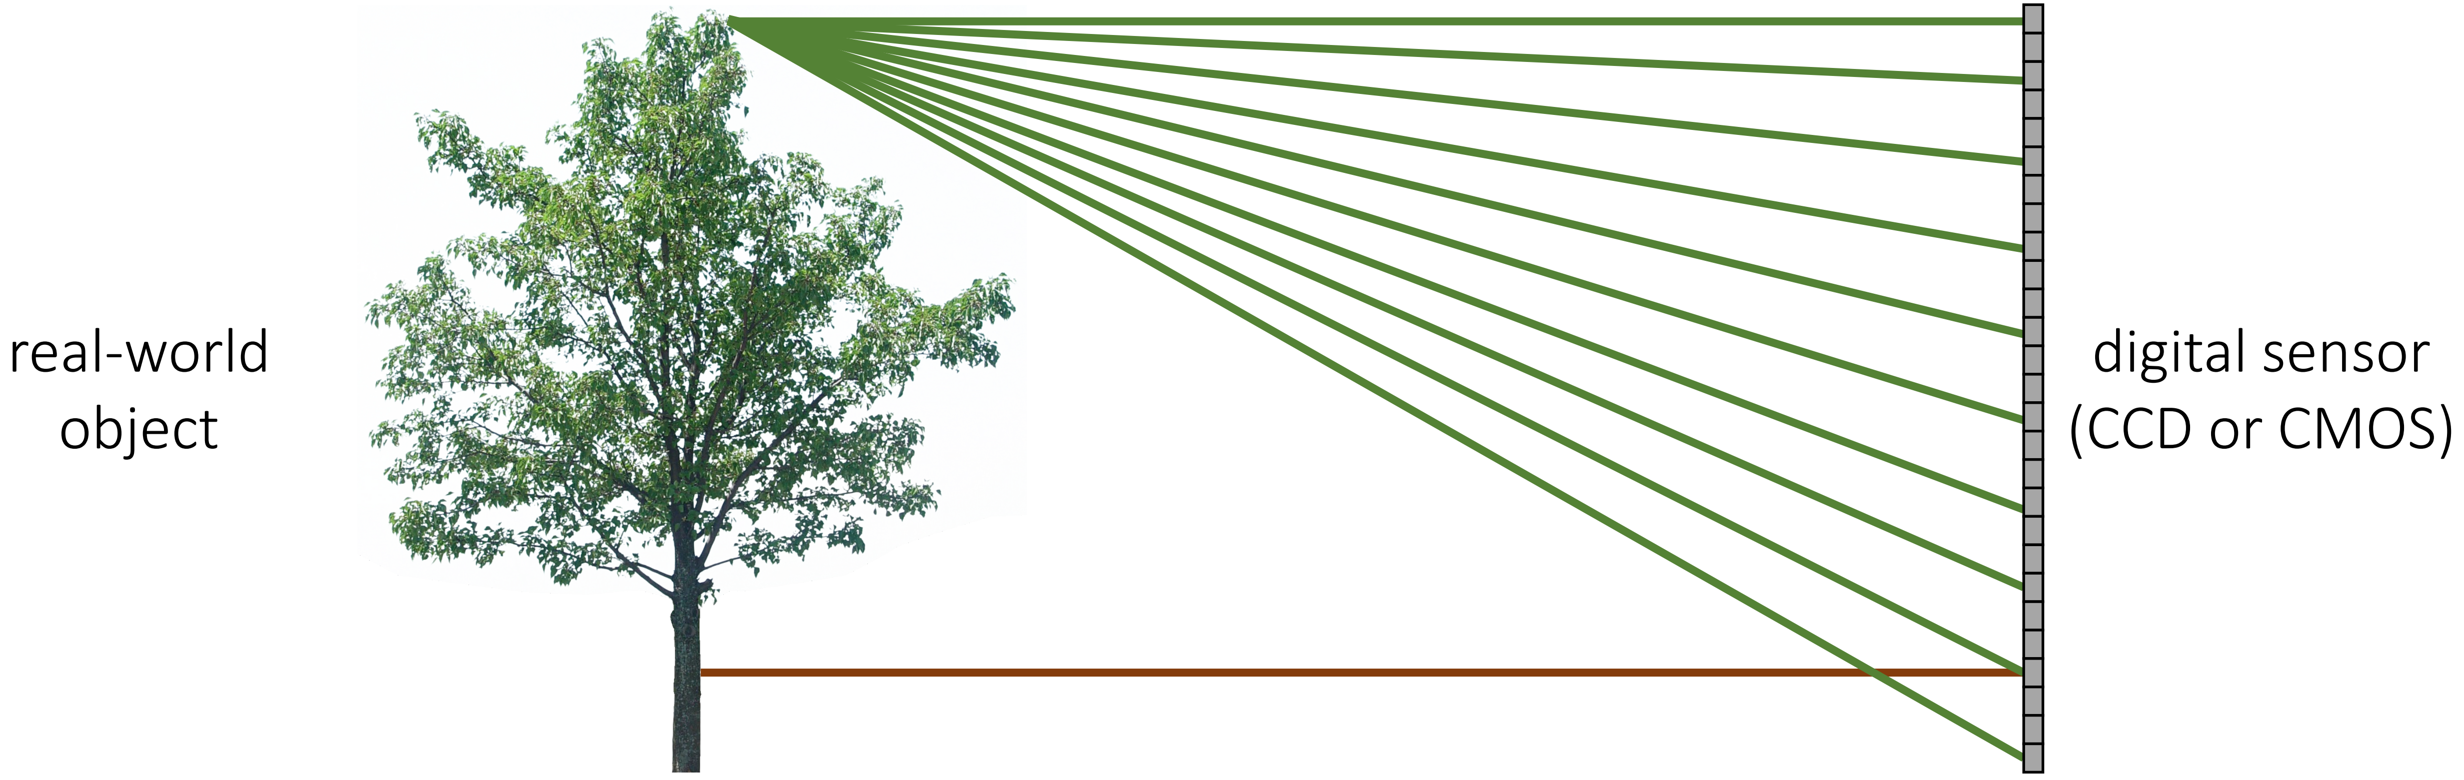

Bare-sensor imaging

real-world
object

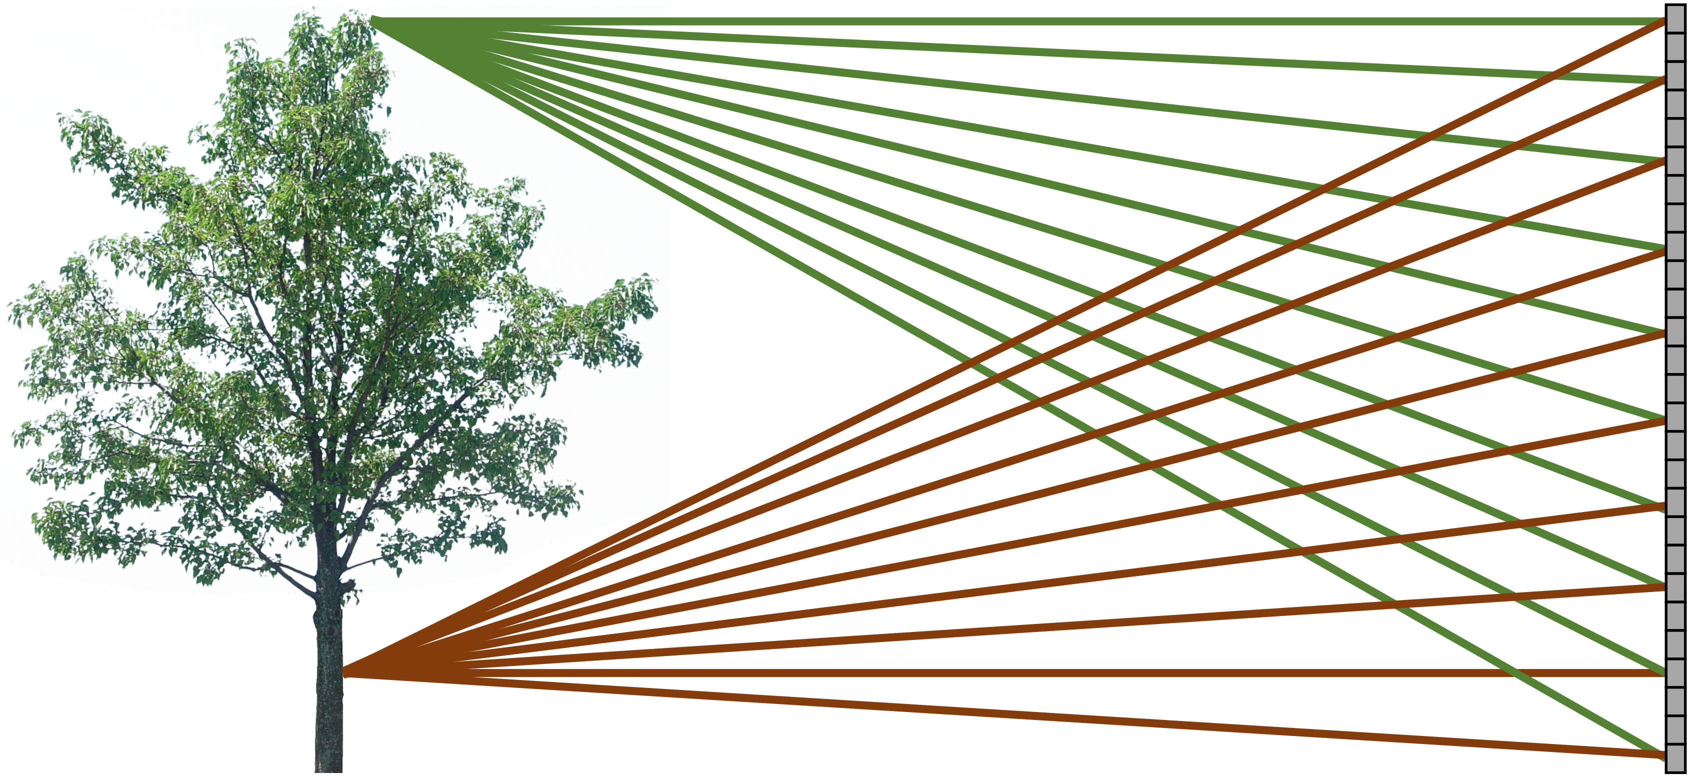

digital sensor
(CCD or CMOS)

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

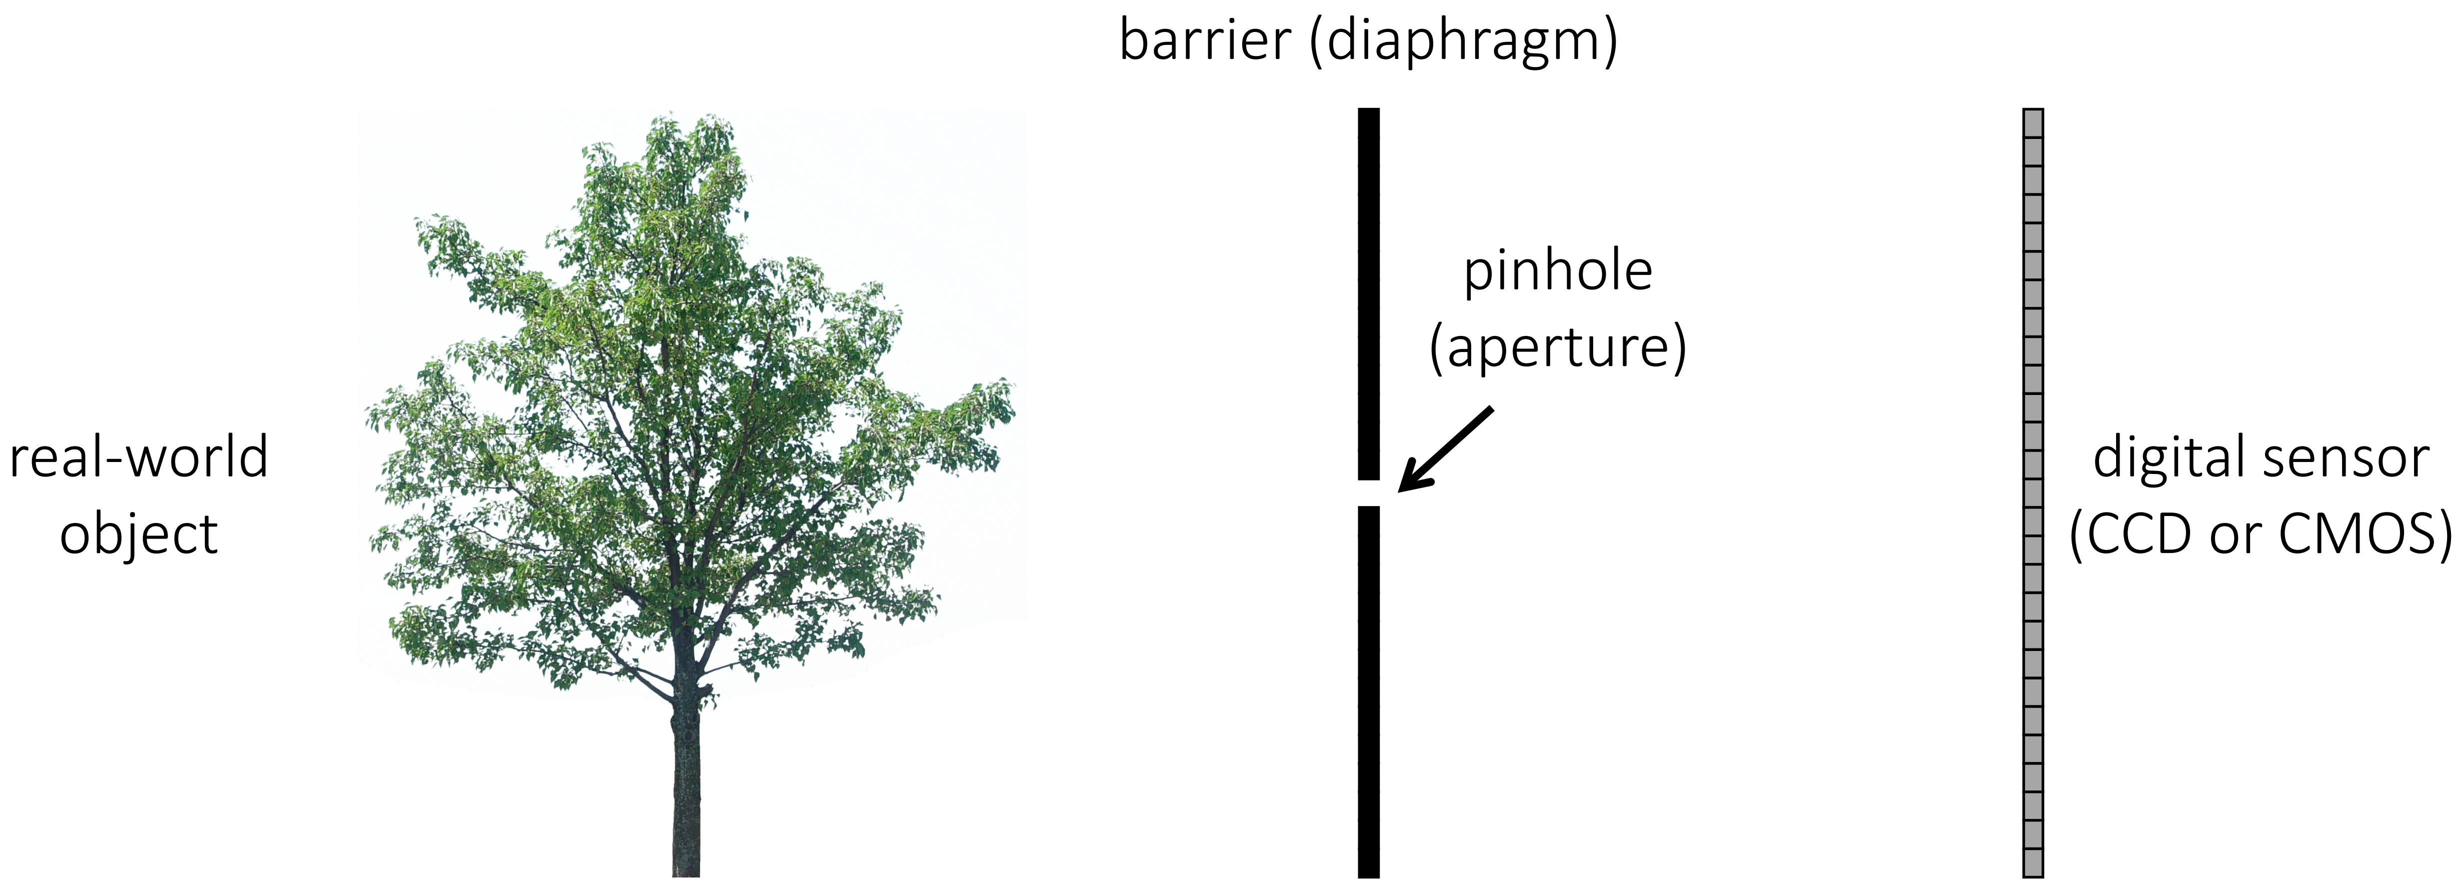

What would an image taken like this look like?

Pinhole imaging

real-world
object

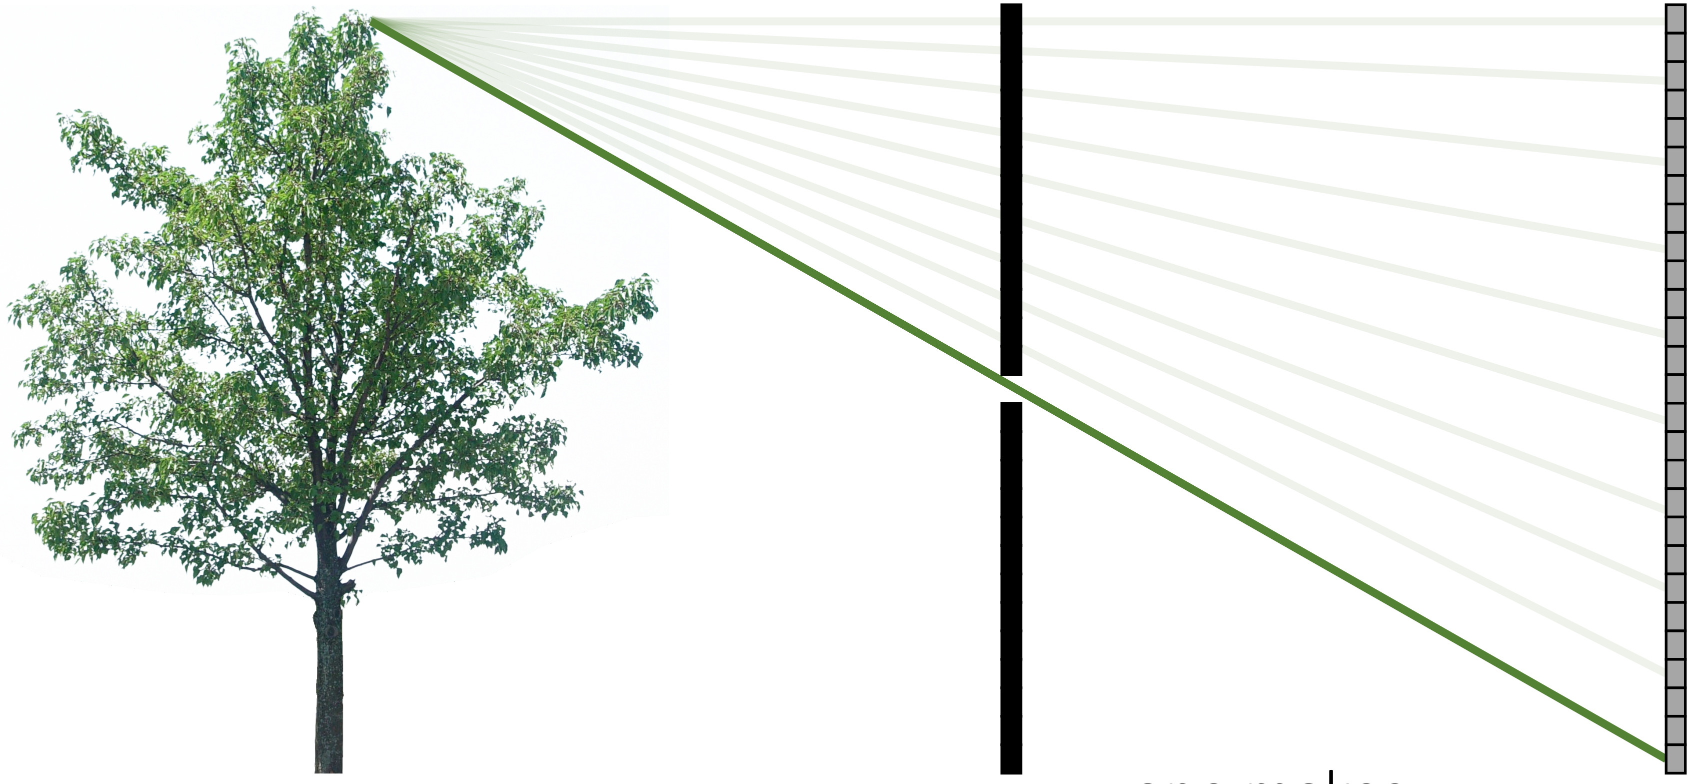

digital sensor
(CCD or CMOS)

Pinhole imaging

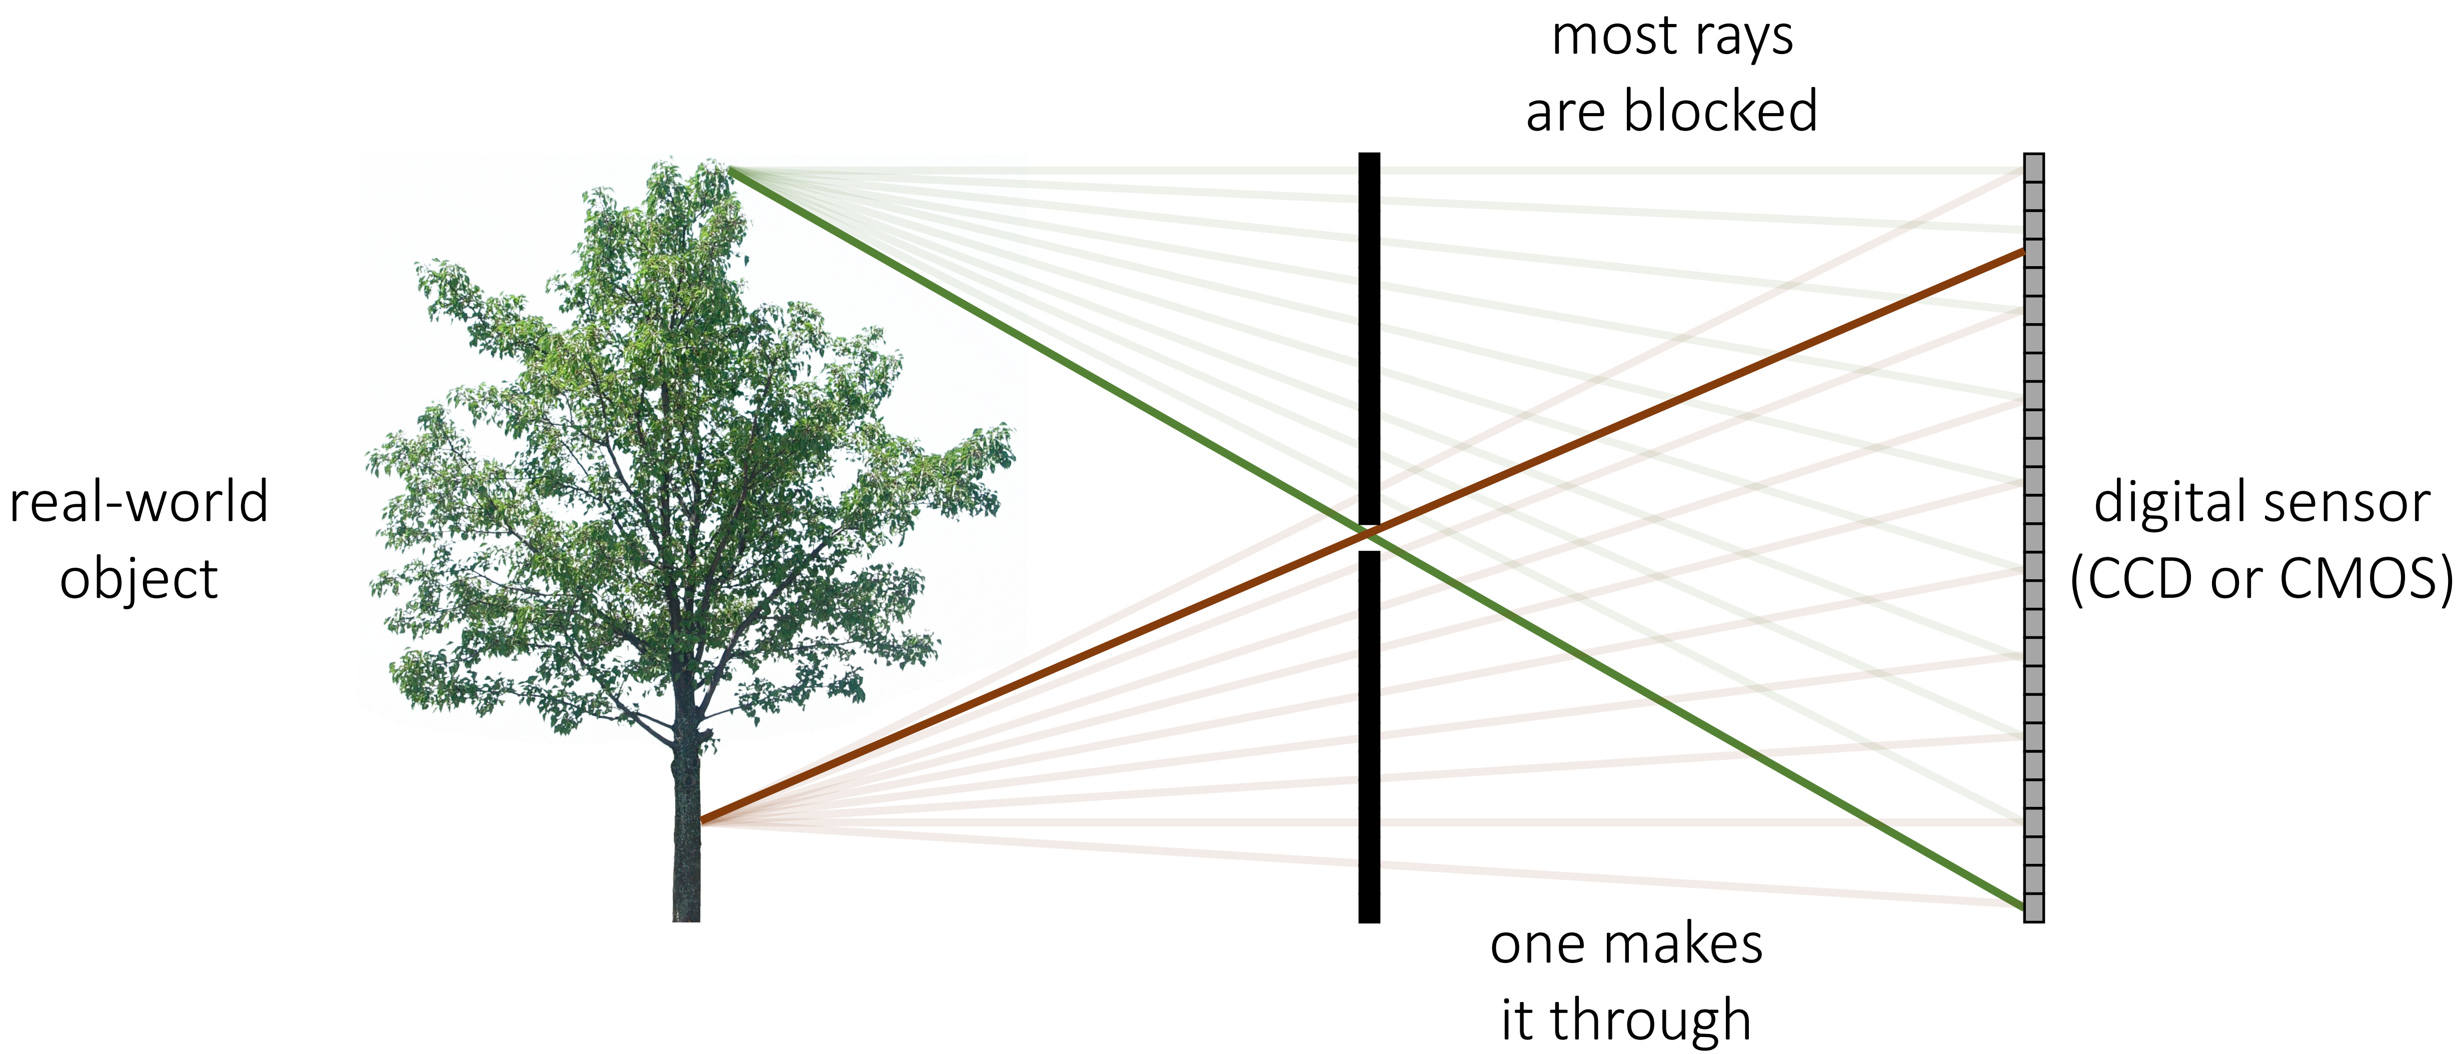

Pinhole imaging

real-world
object

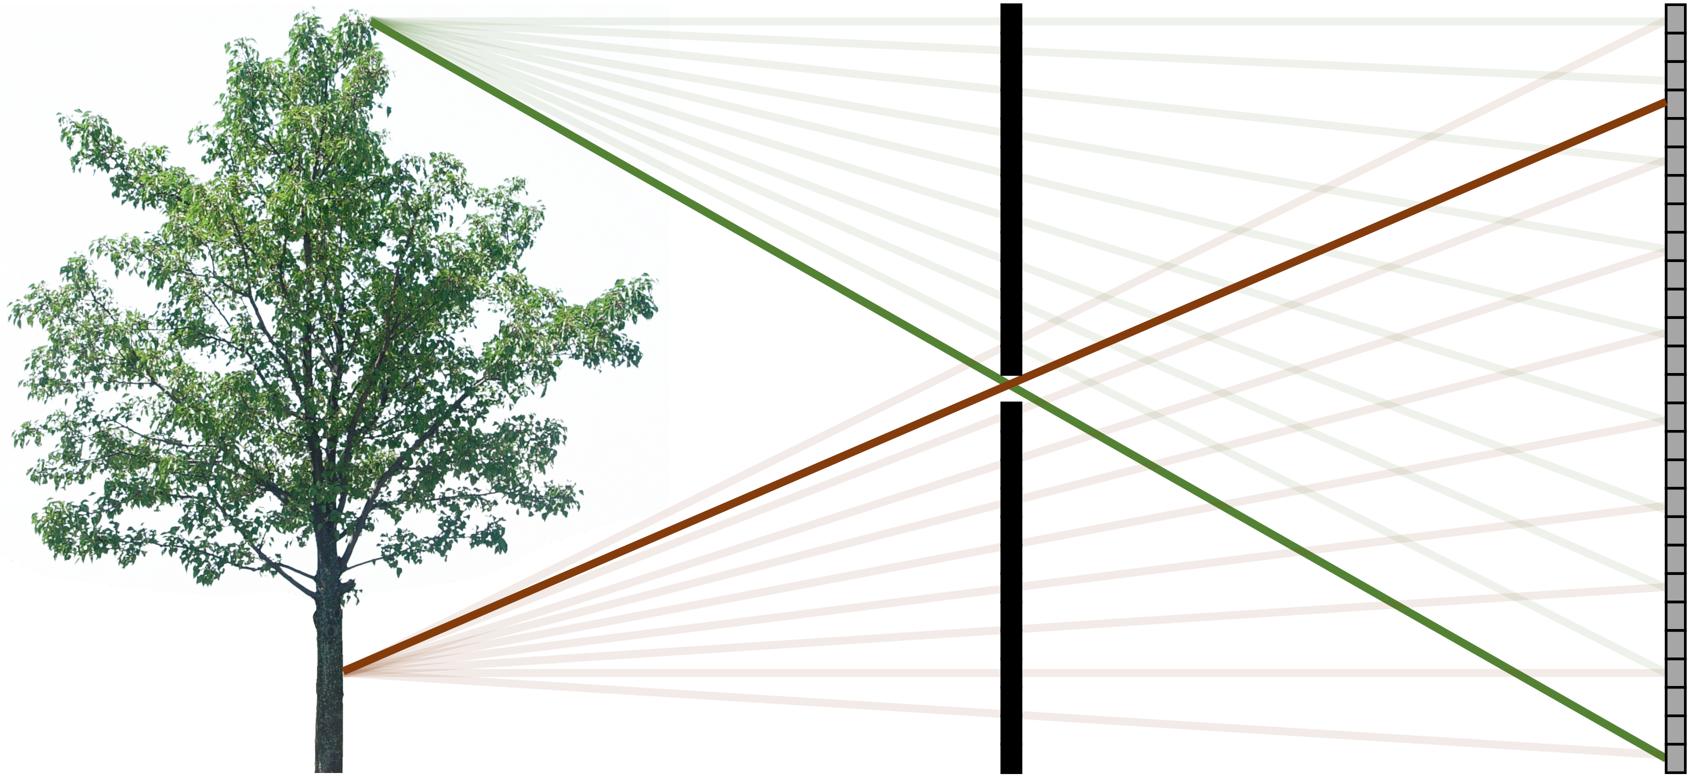

digital sensor
(CCD or CMOS)

What does the
image on the
sensor look like?

Each scene point contributes to only one sensor pixel

Pinhole imaging

real-world
object

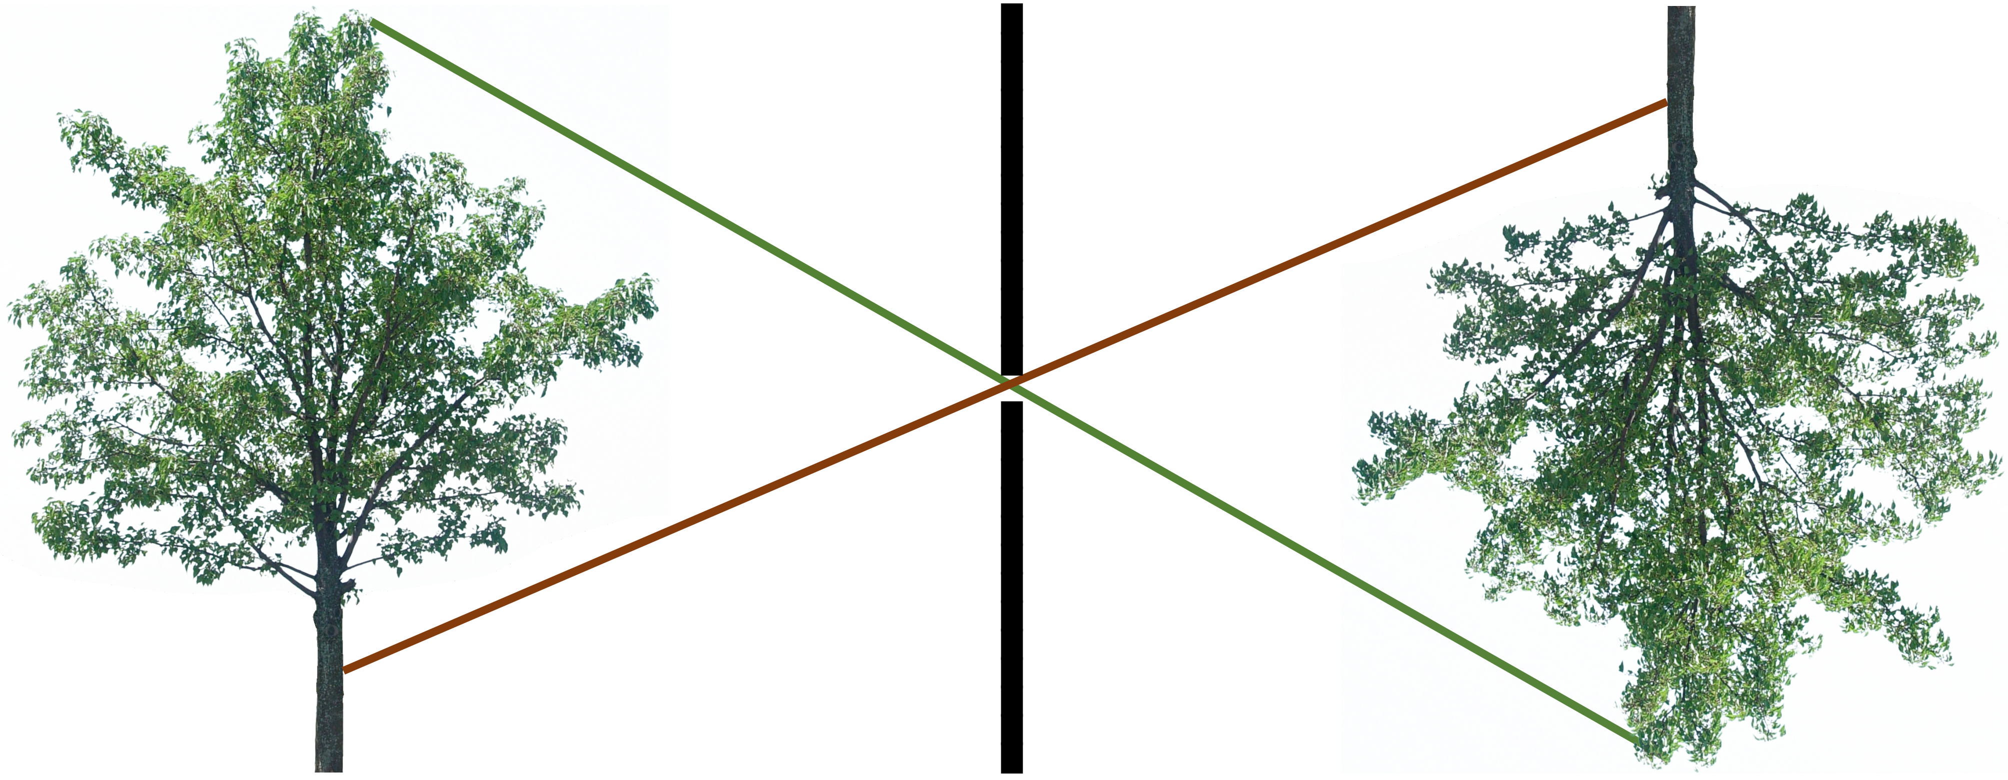

copy of real-world object
(inverted and scaled)

Pinhole camera

Pinhole camera a.k.a. camera obscura

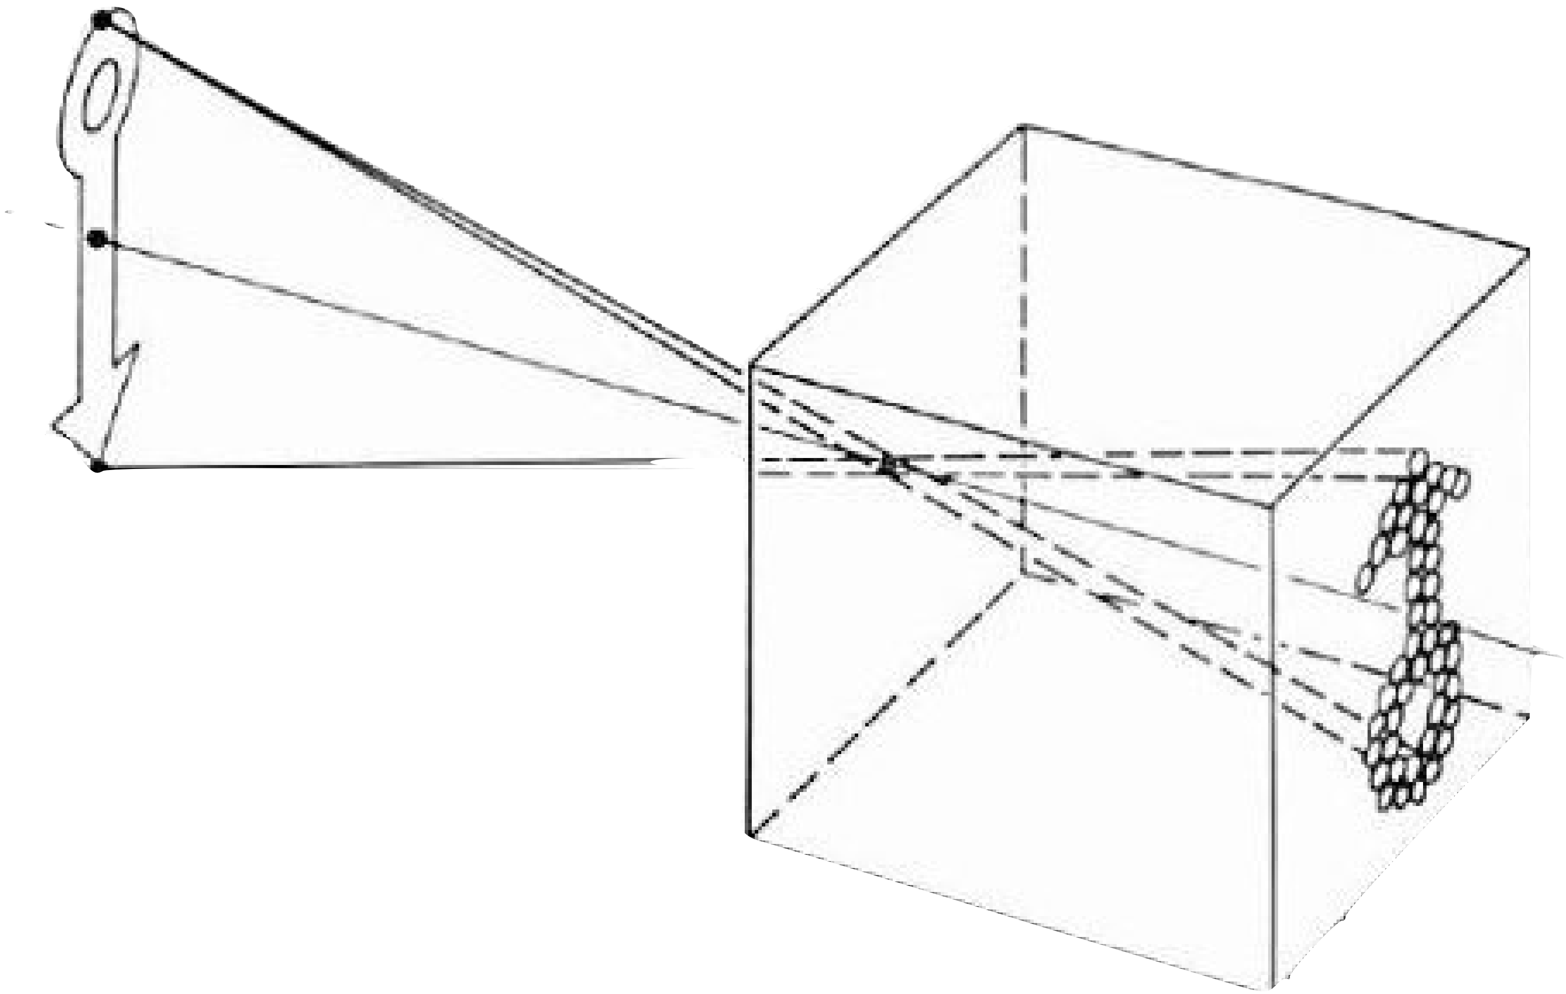

Pinhole camera a.k.a. camera obscura

First mention ...

Chinese philosopher Mozi
(470 to 390 BC)

First camera ...

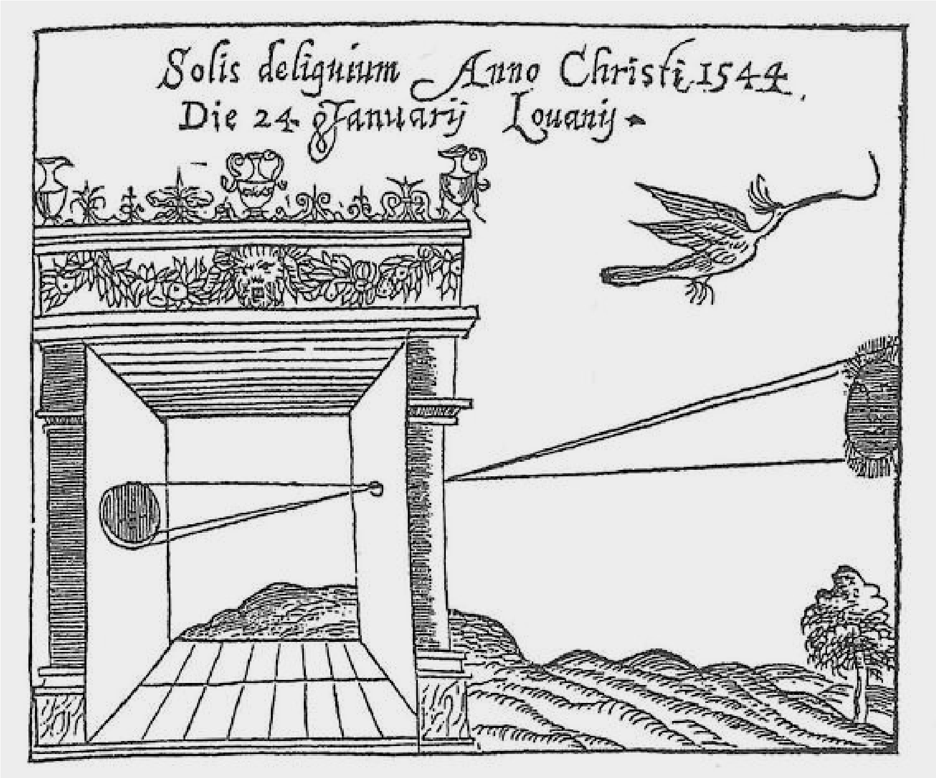

Greek philosopher Aristotle
(384 to 322 BC)

Pinhole camera terms

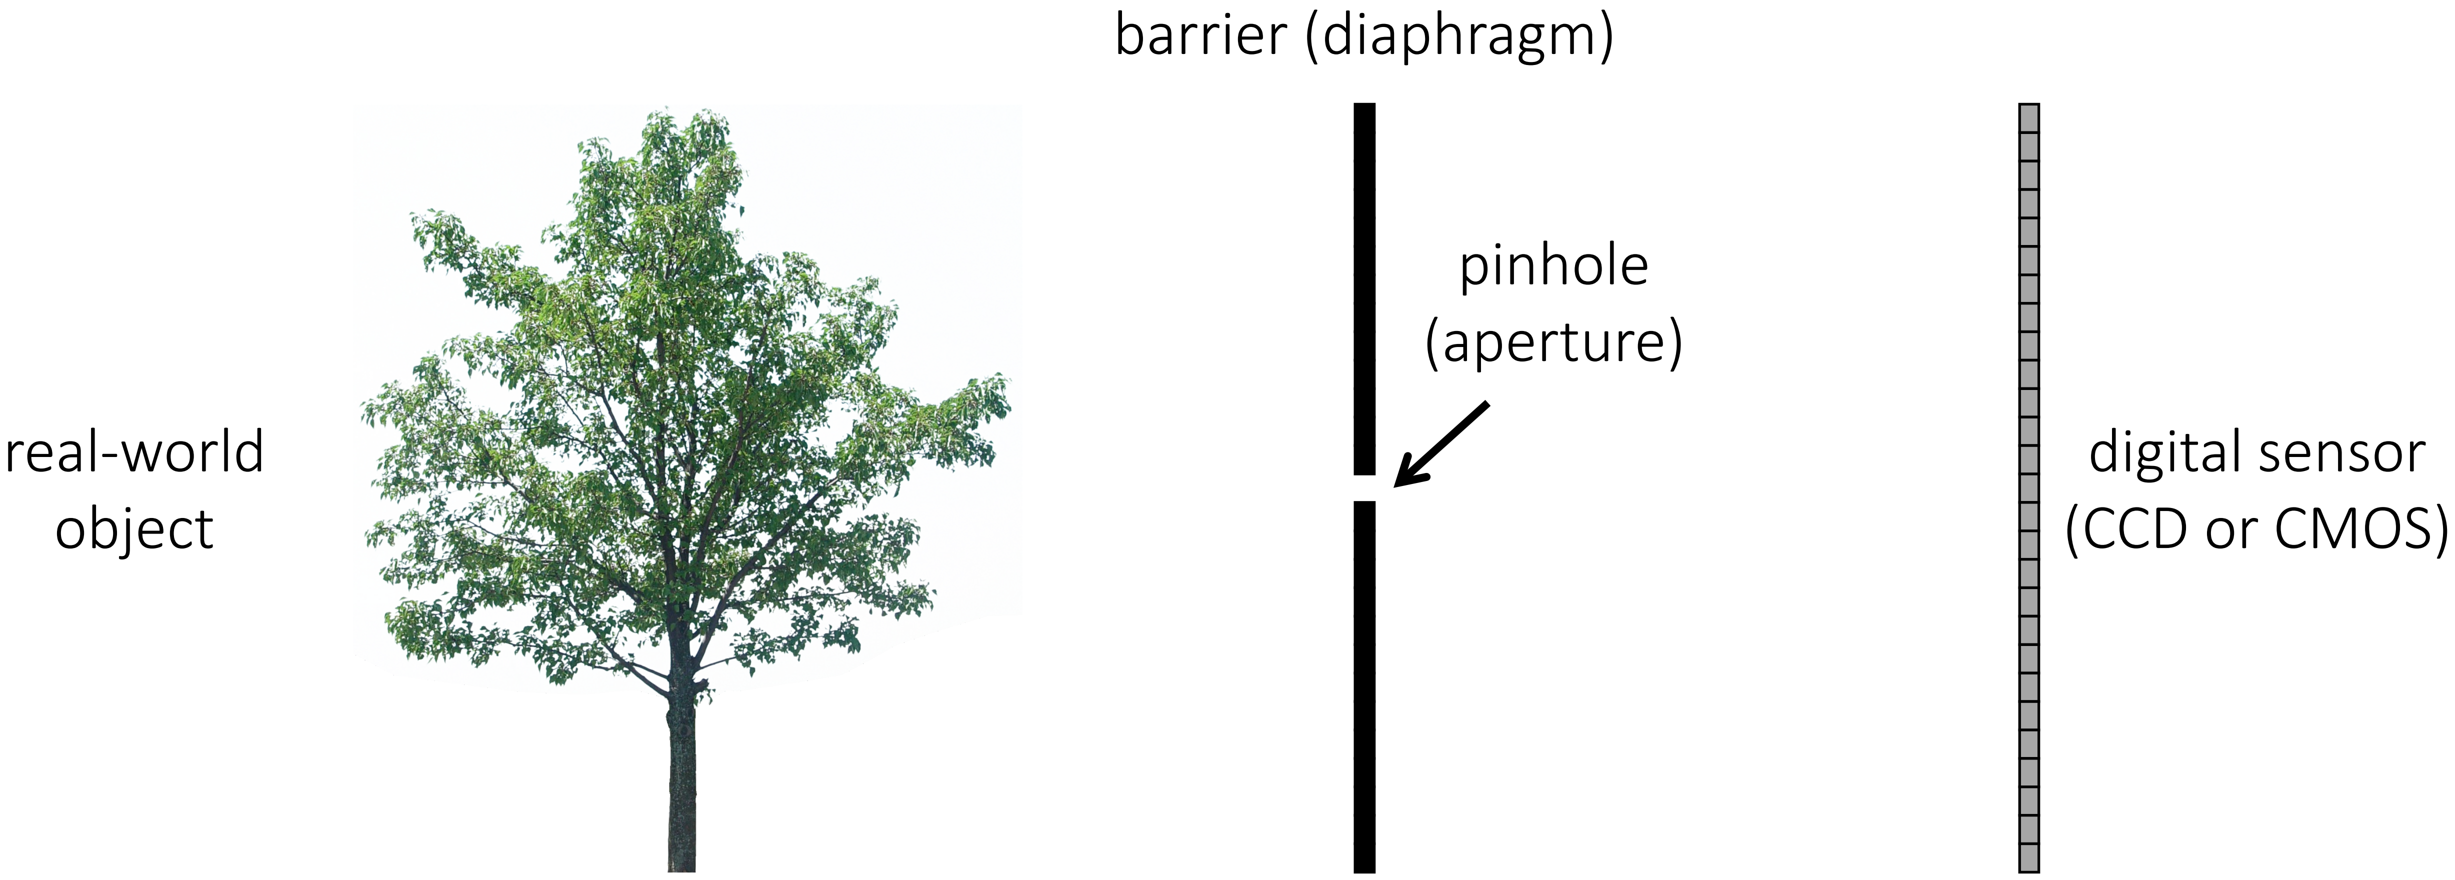

Pinhole camera terms

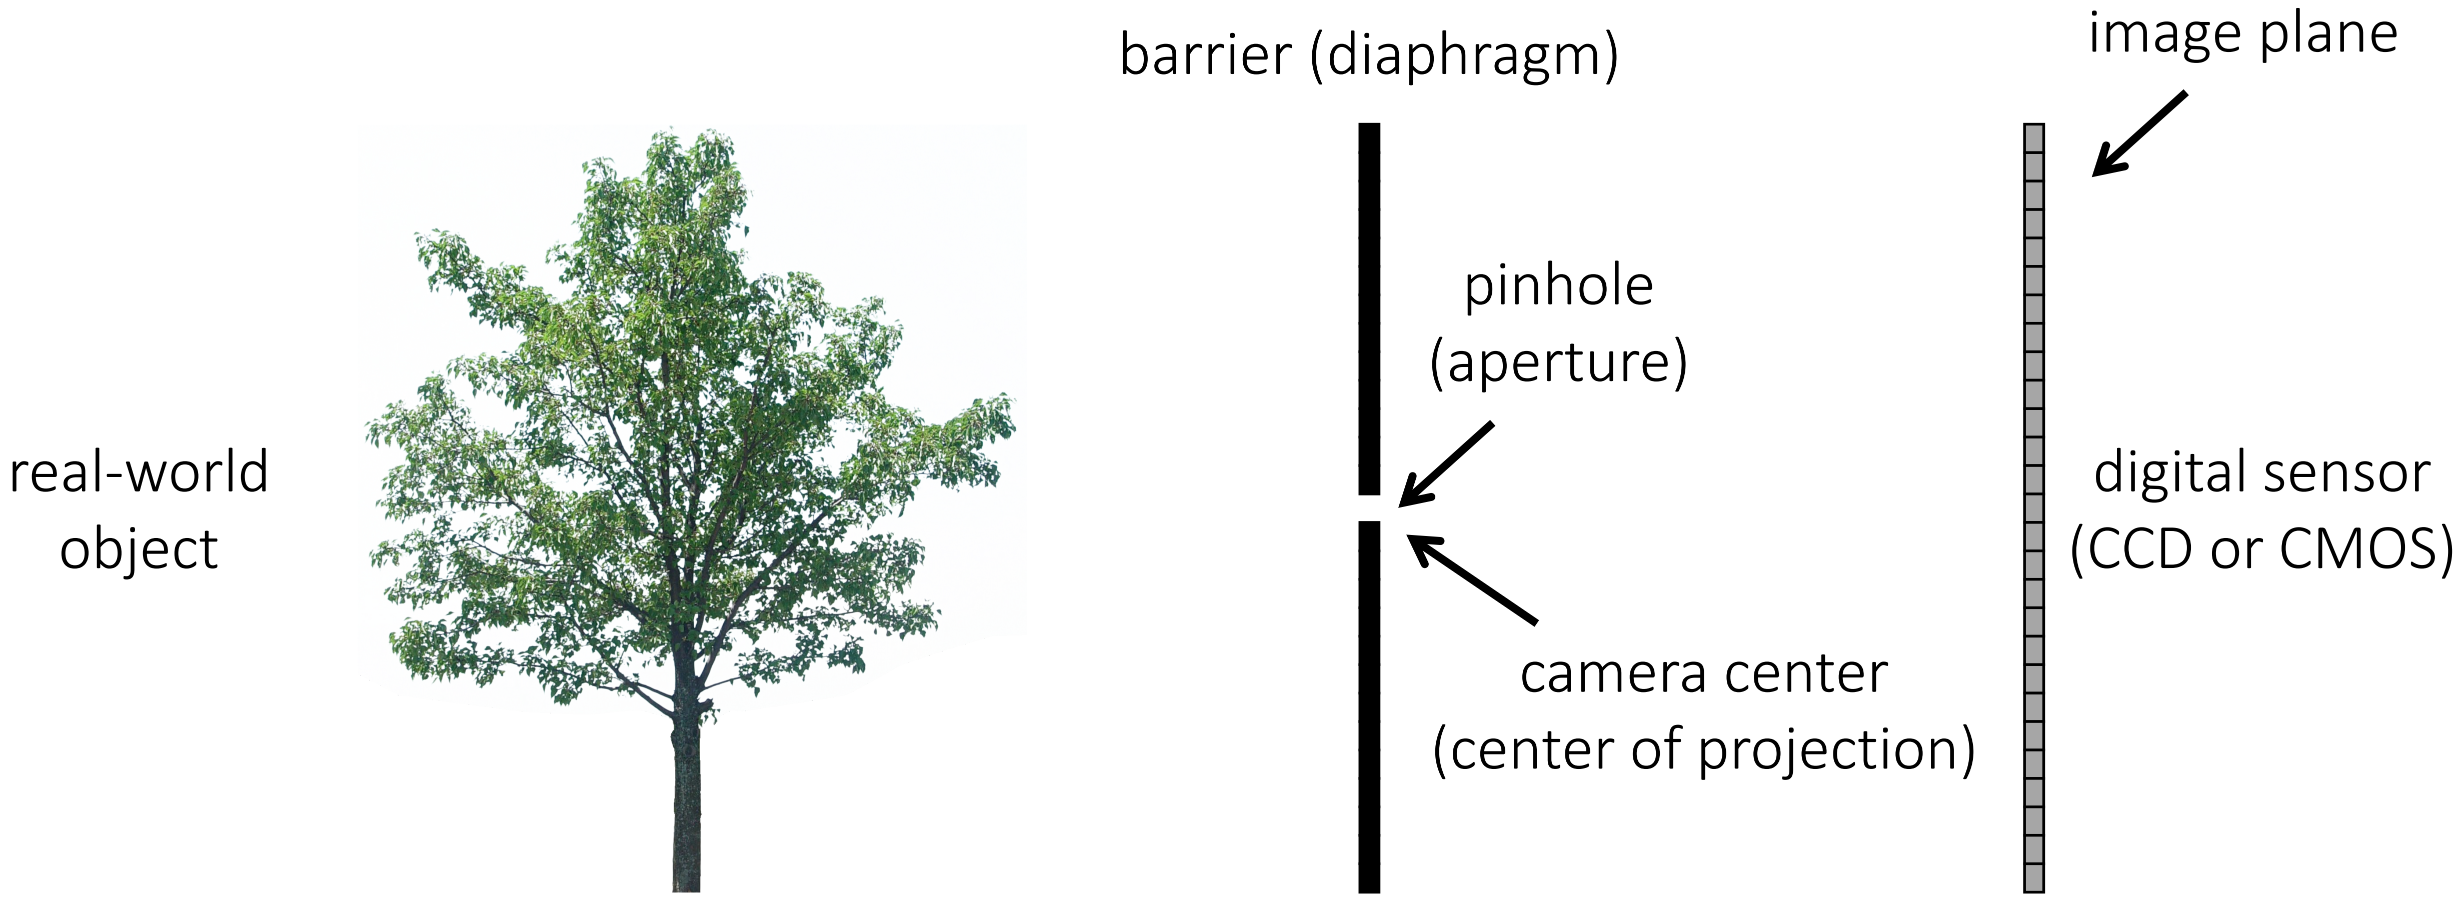

Focal length

real-world object

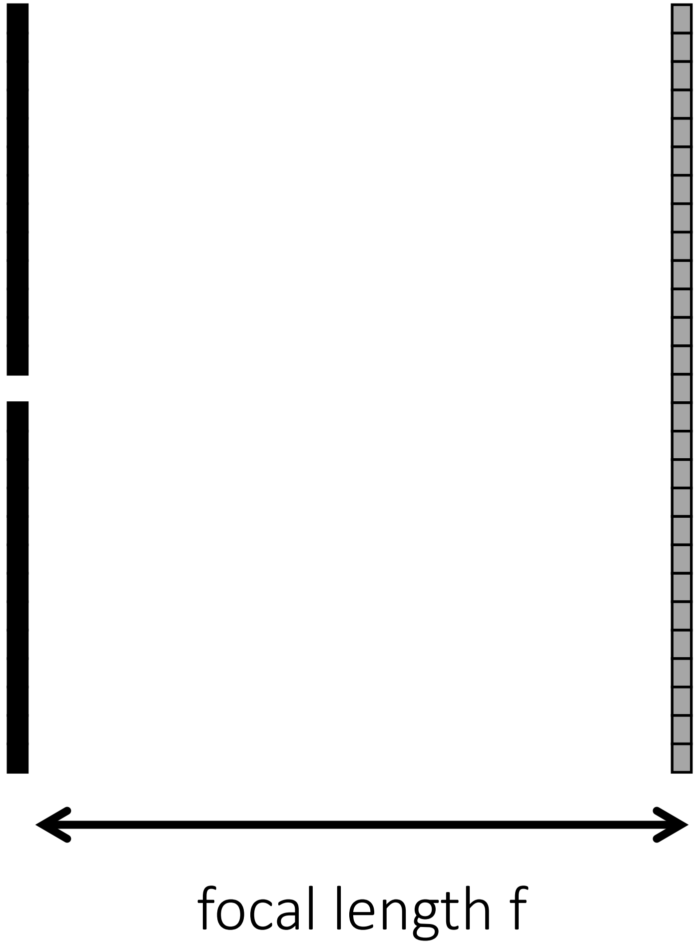

Focal length

What happens as we change the focal length?

real-world
object

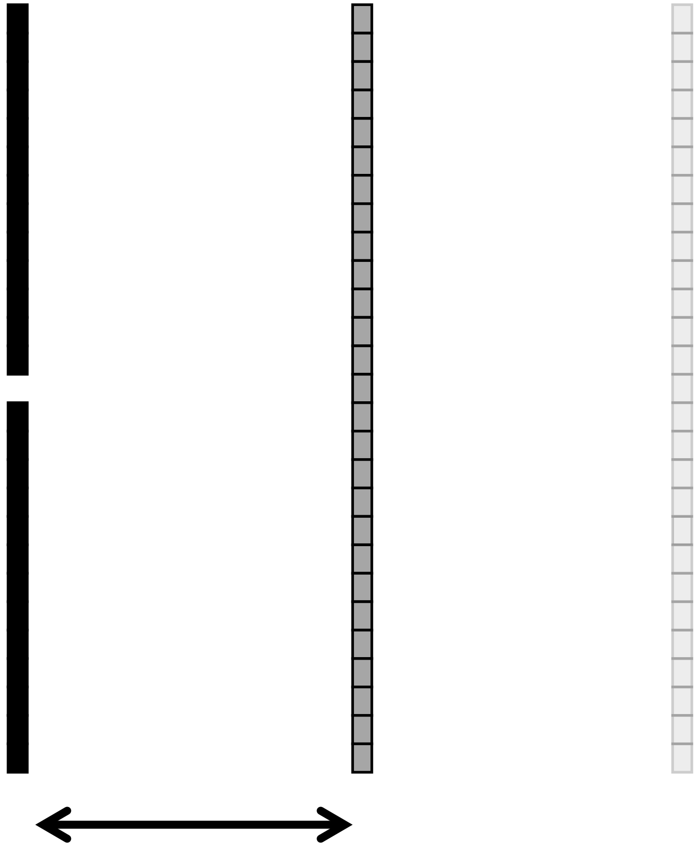

Focal length

What happens as we change the focal length?

real-world object

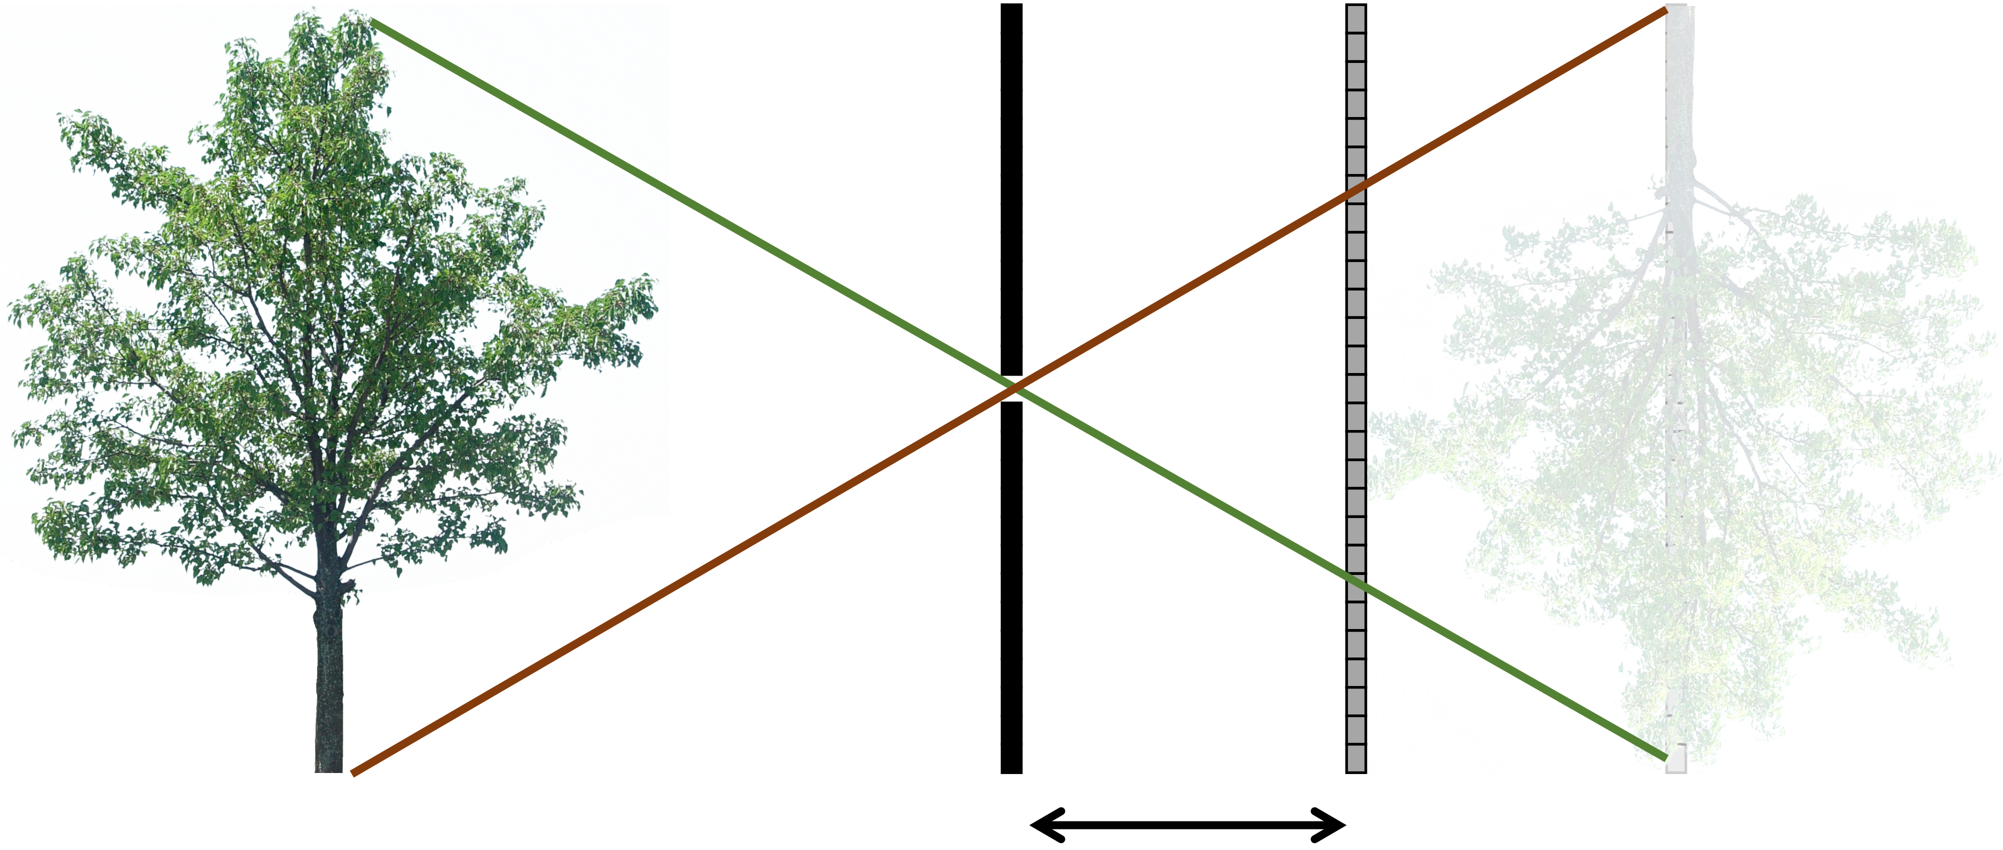

focal length 0.5 f

Focal length

What happens as we change the focal length?

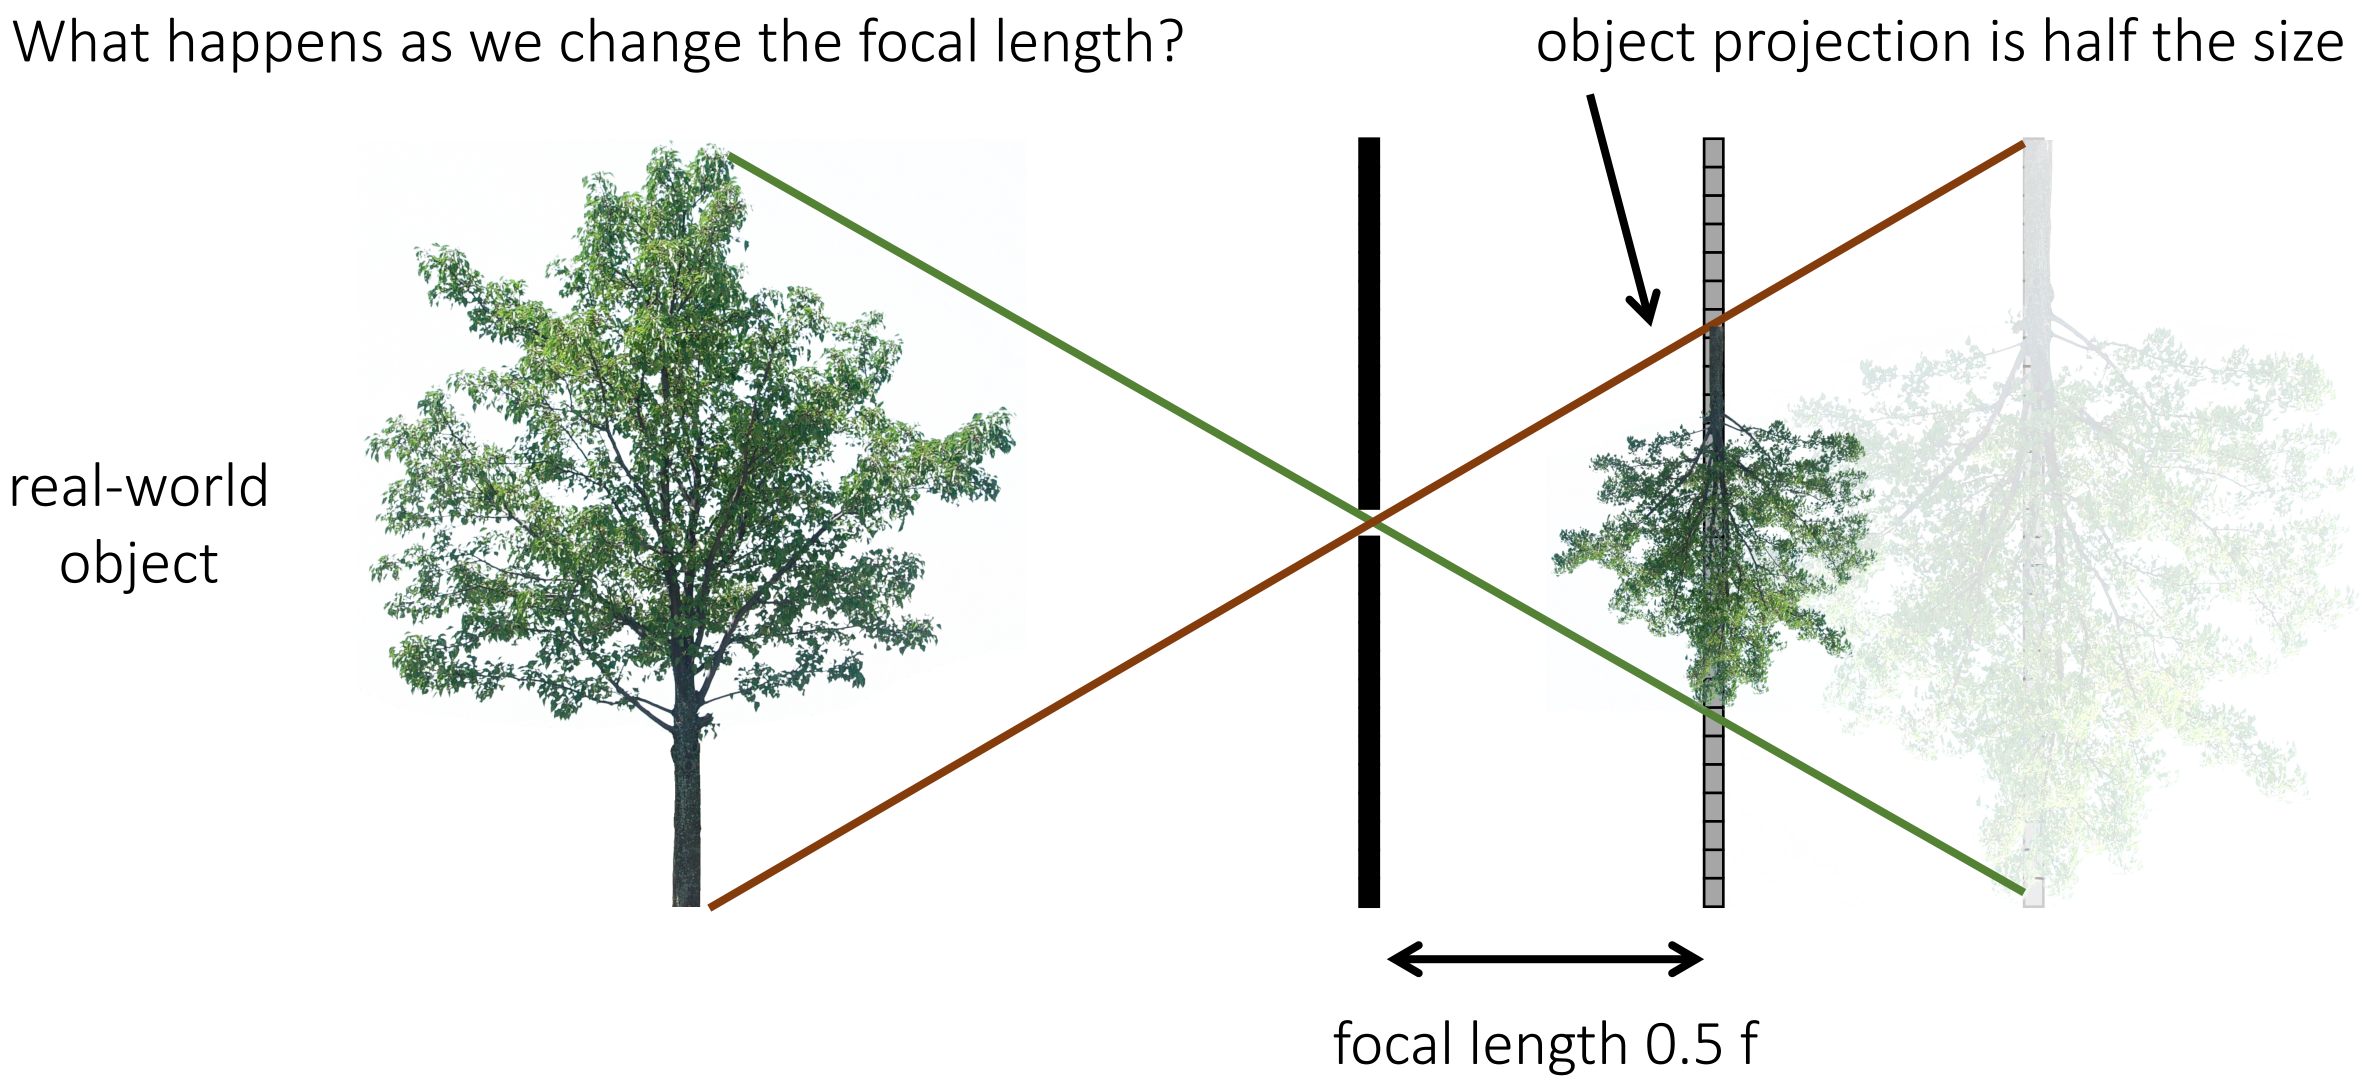

Pinhole size

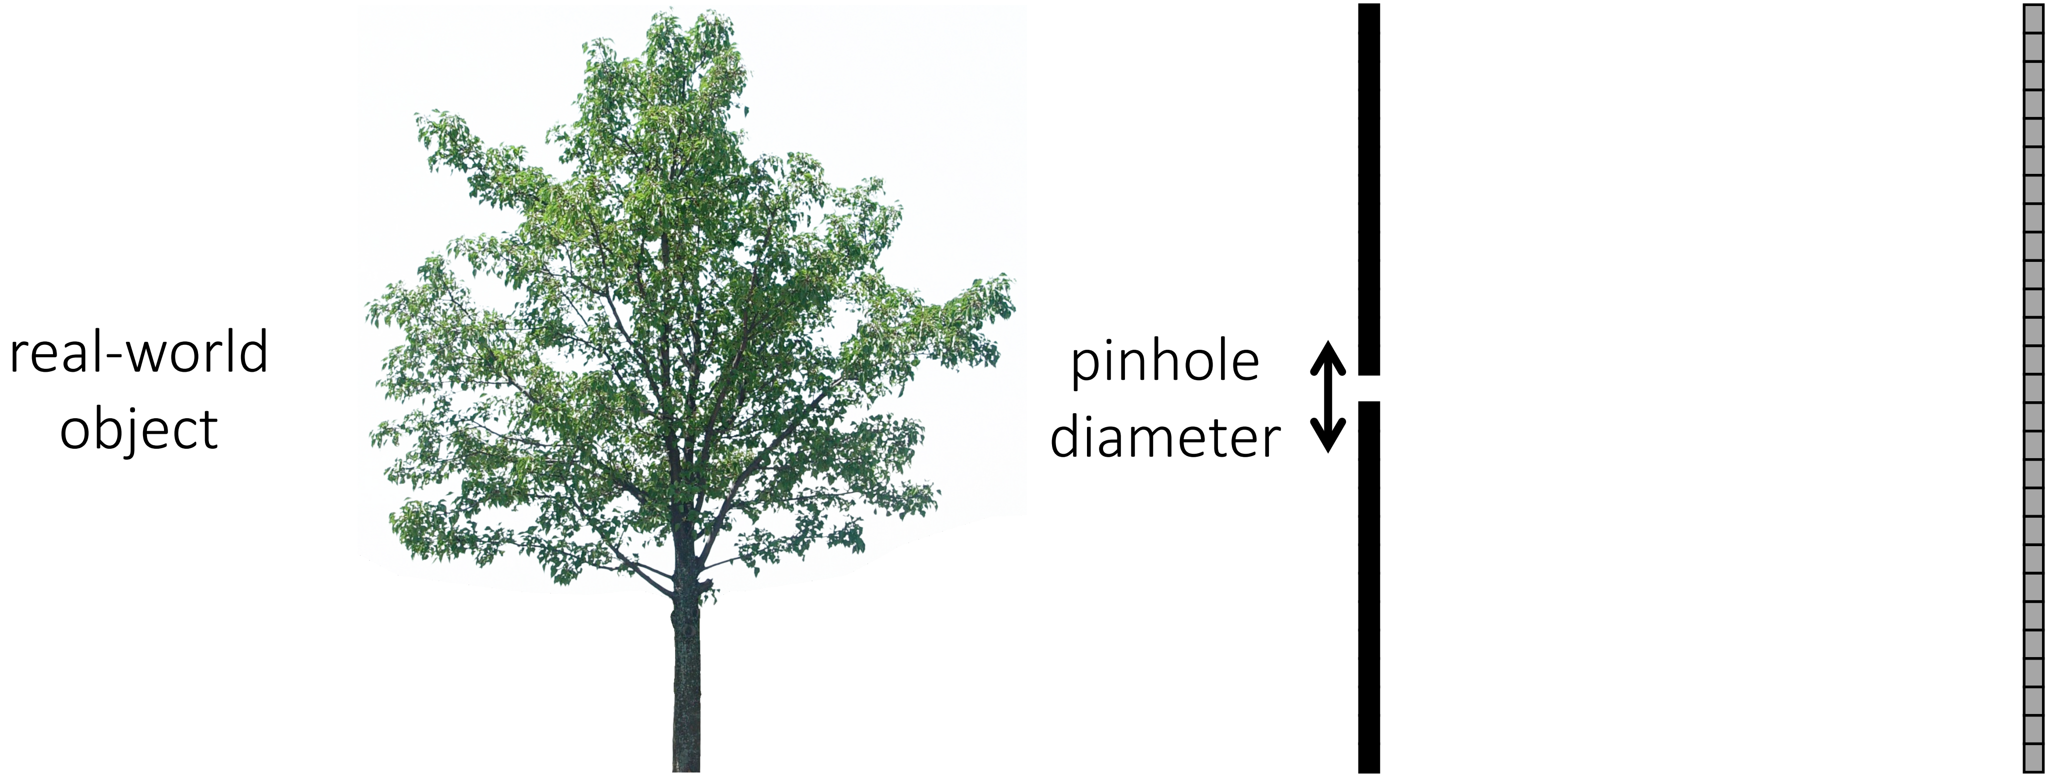

Ideal pinhole has infinitesimally small size

- In practice that is impossible.

Pinhole size

What happens as we change the pinhole diameter?

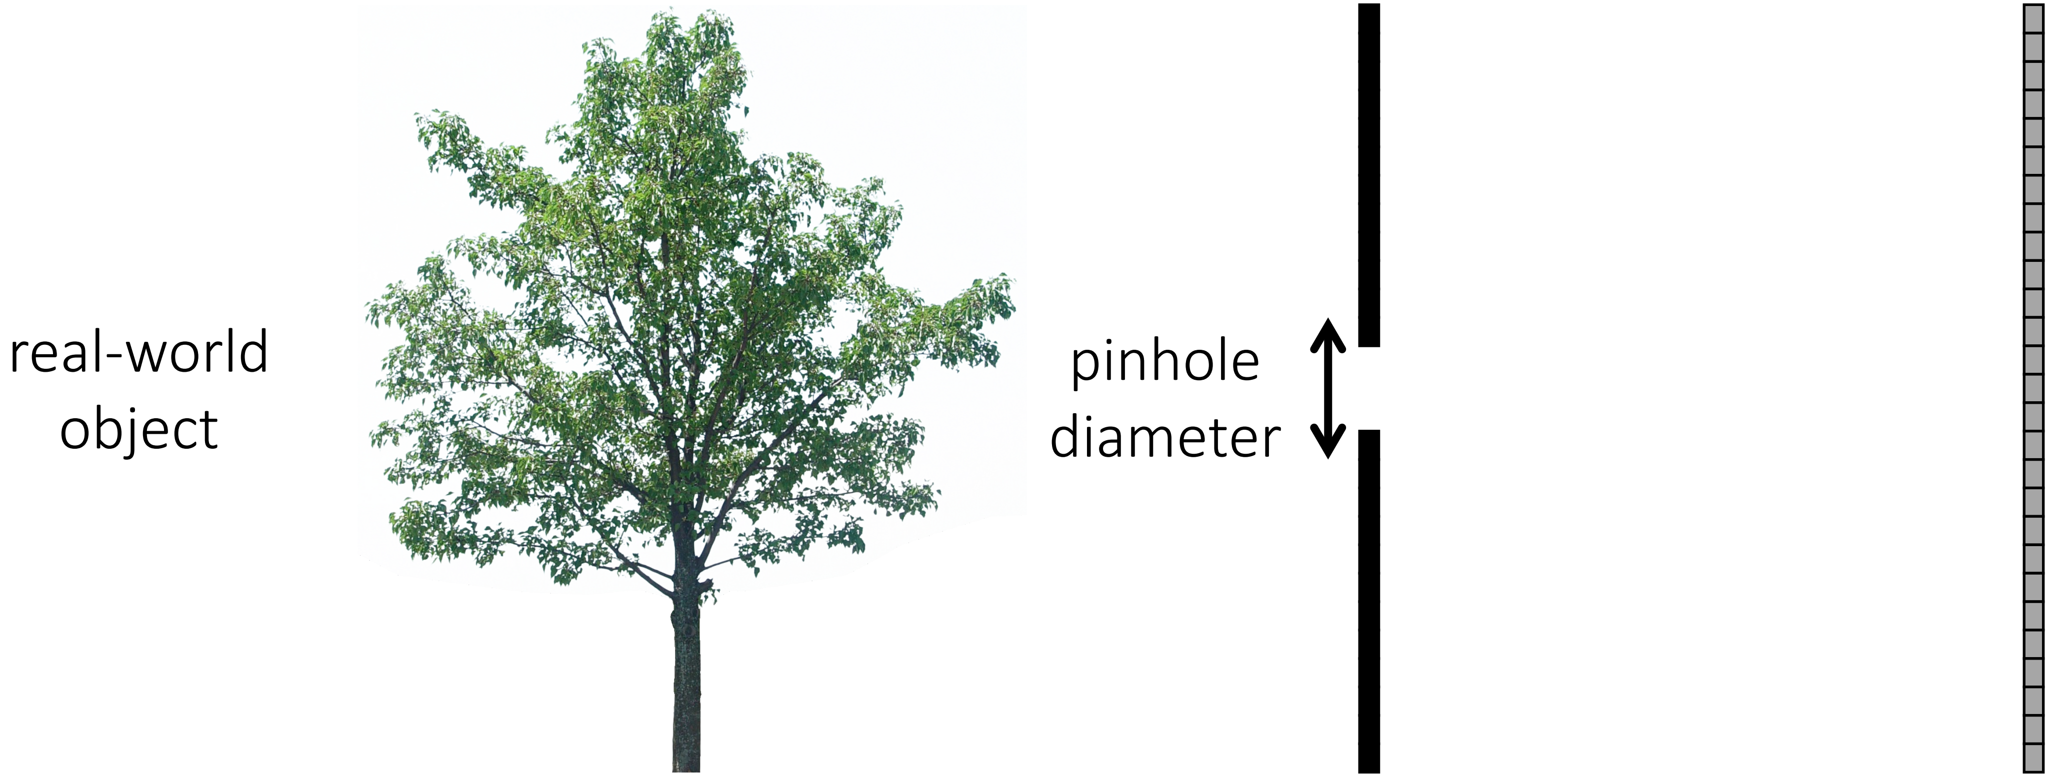

Pinhole size

What happens as we change the pinhole diameter?

real-world
object

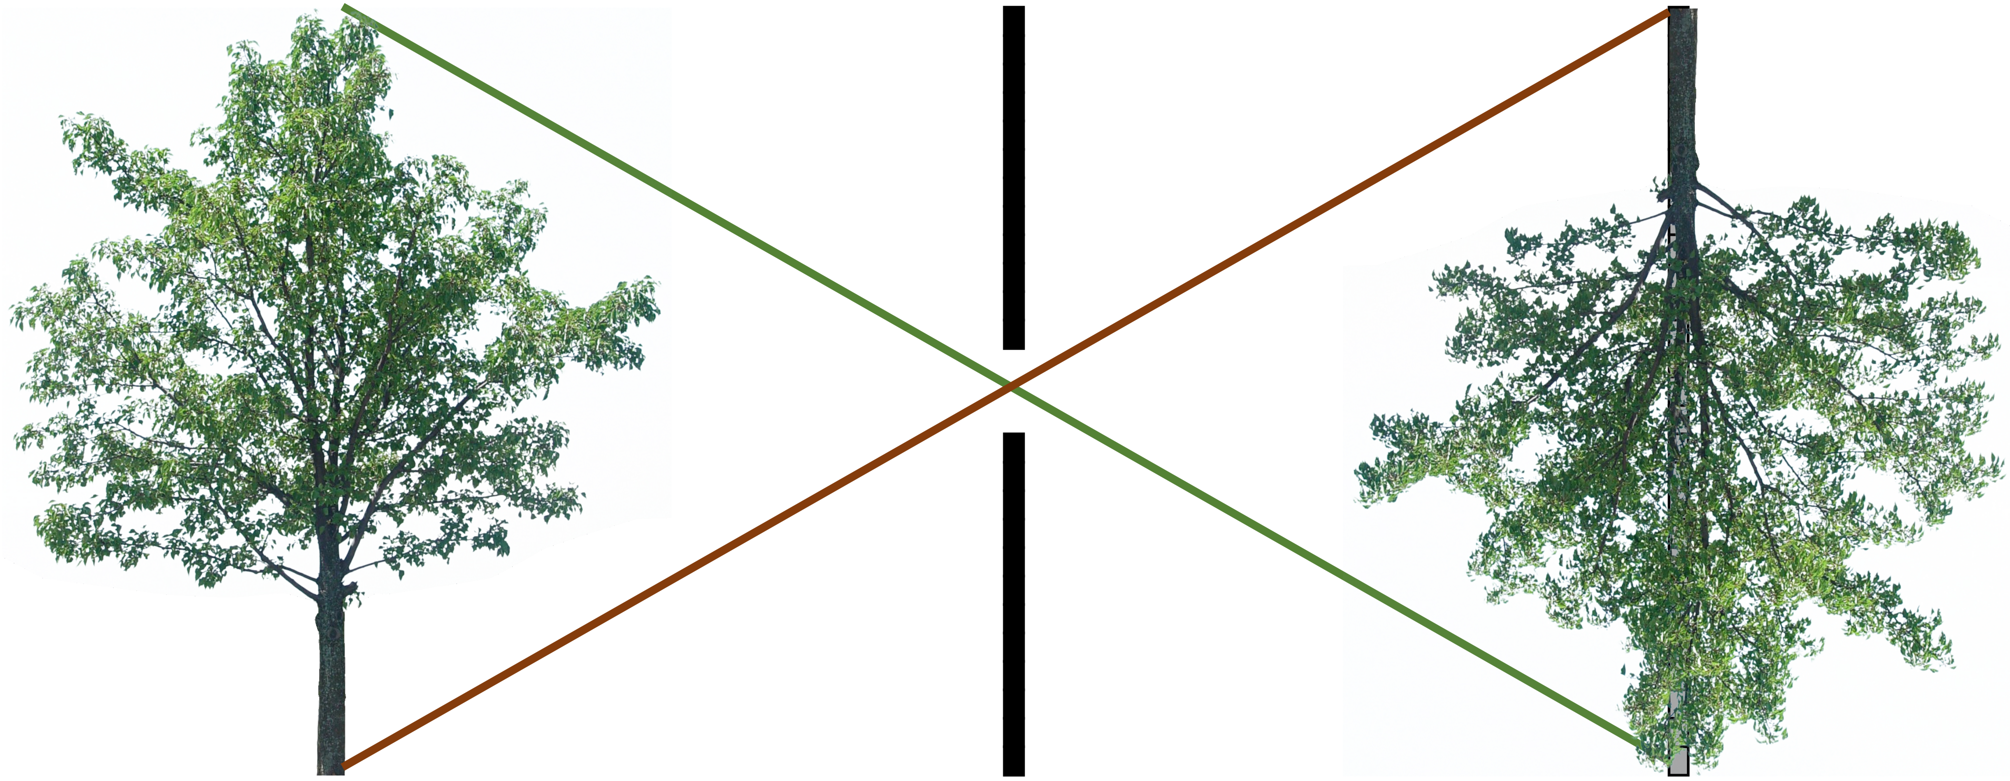

Pinhole size

What happens as we change the pinhole diameter?

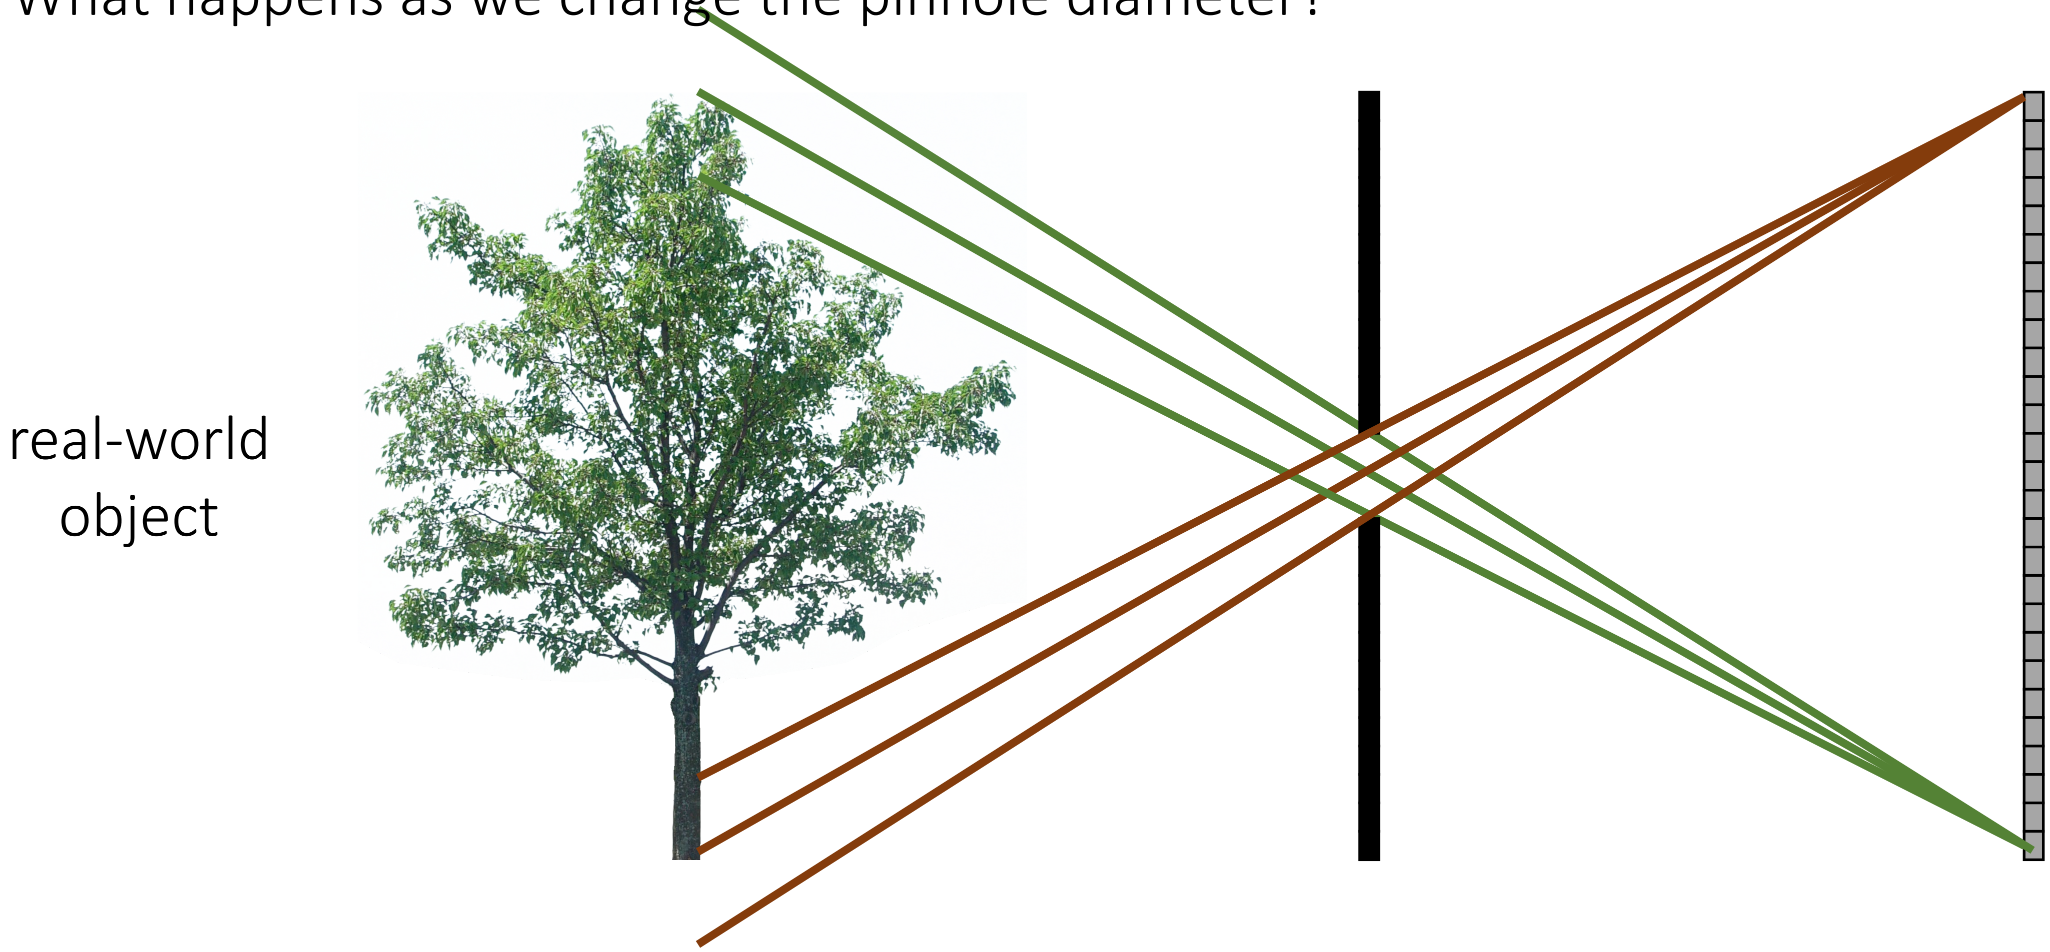

Pinhole size

What happens as we change the pinhole diameter?

object projection becomes blurrier

real-world object

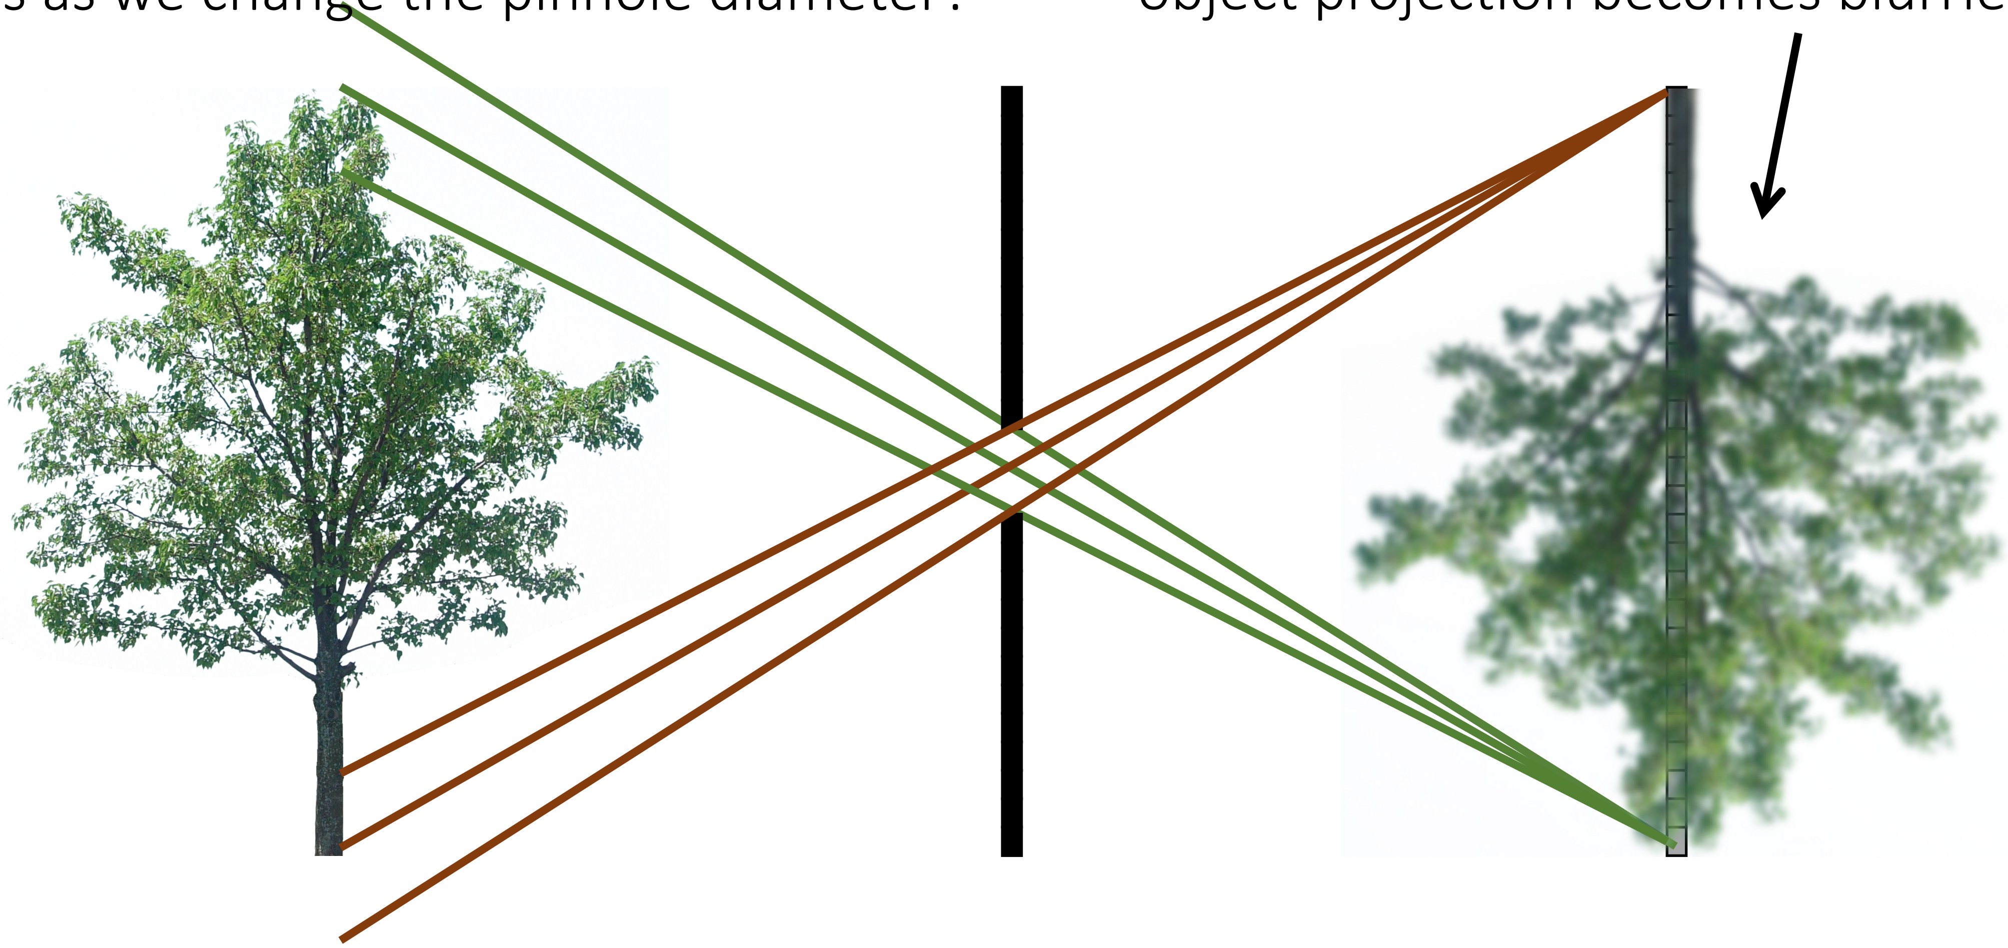

Pinhole size

What happens as we change the pinhole diameter?

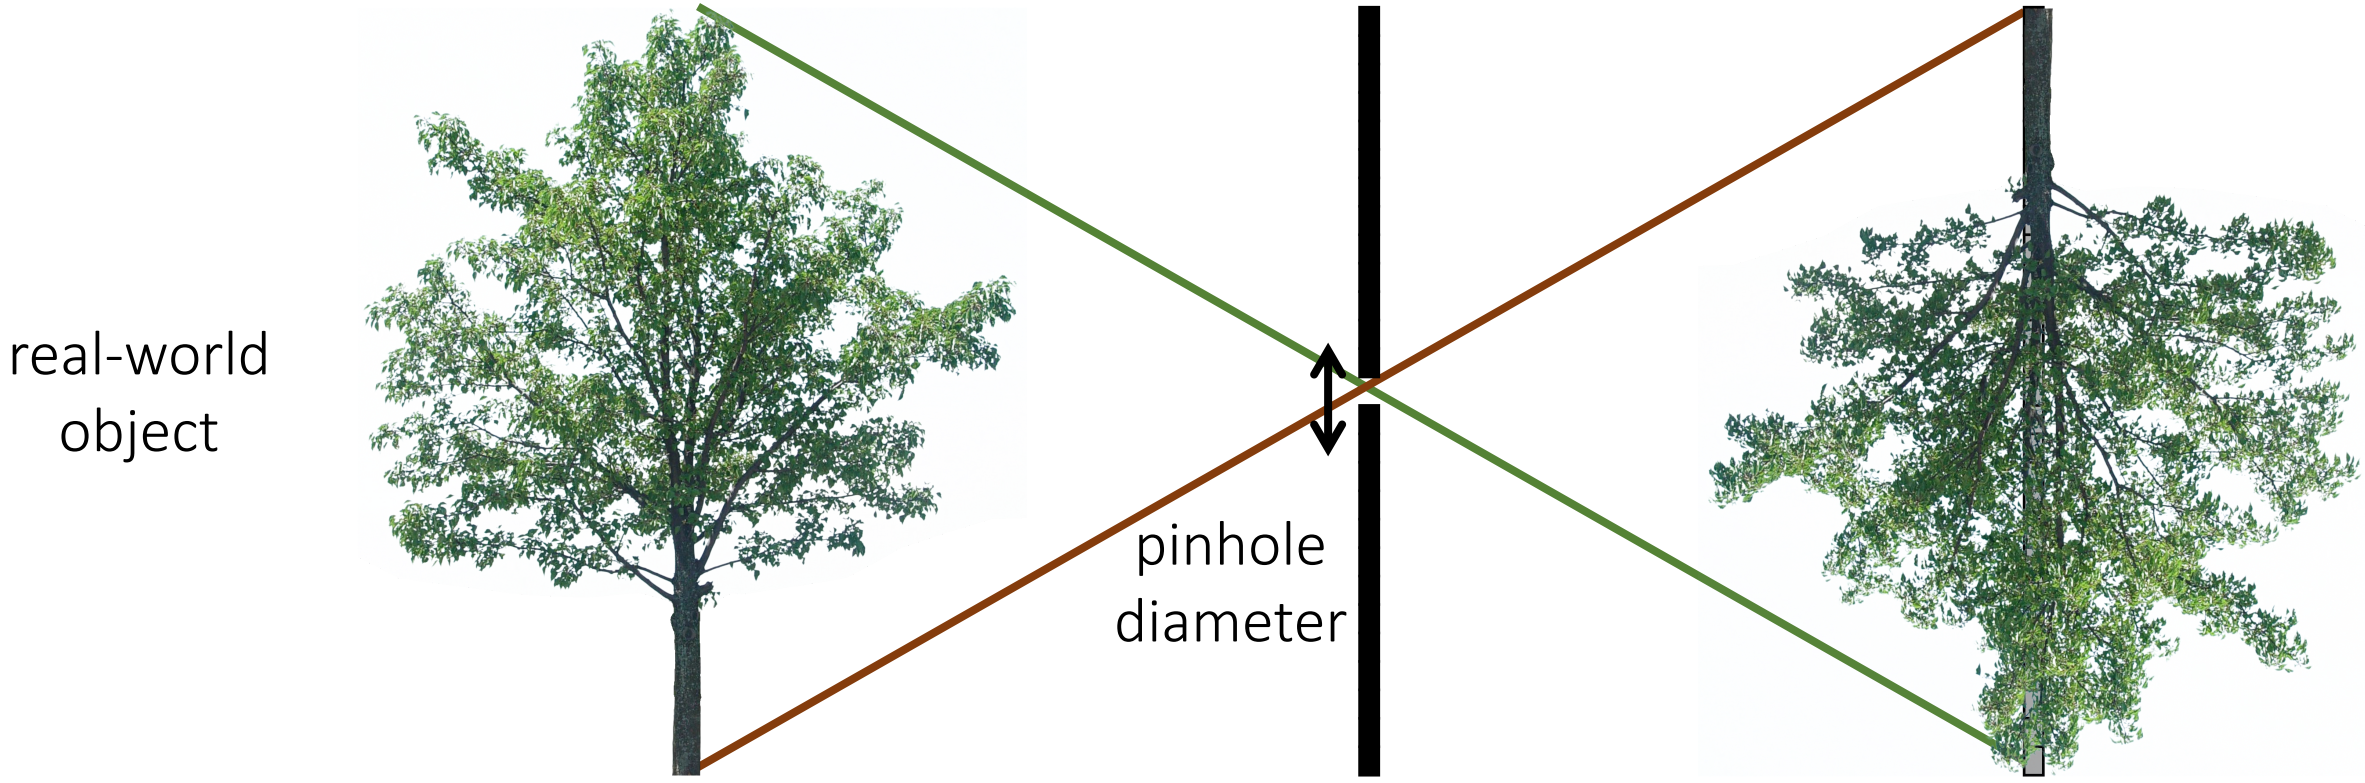

Will the image keep getting sharper the smaller we make the pinhole?

Diffraction limit

A consequence of the wave nature of light

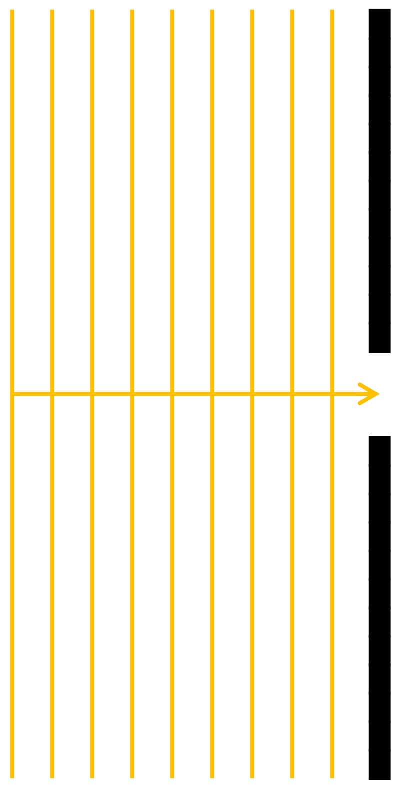

What do geometric optics predict will happen?

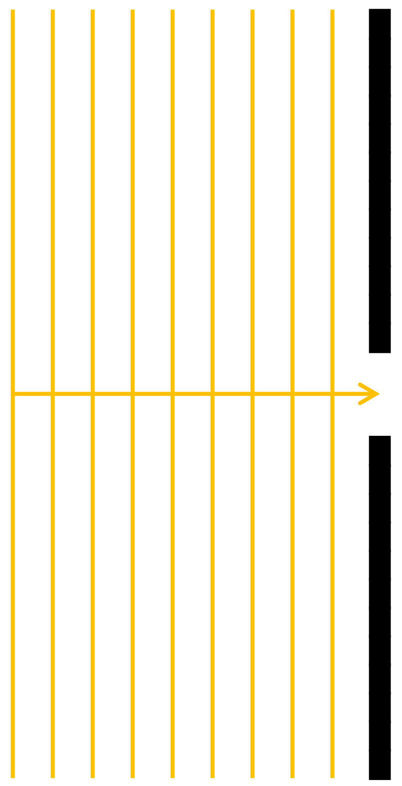

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

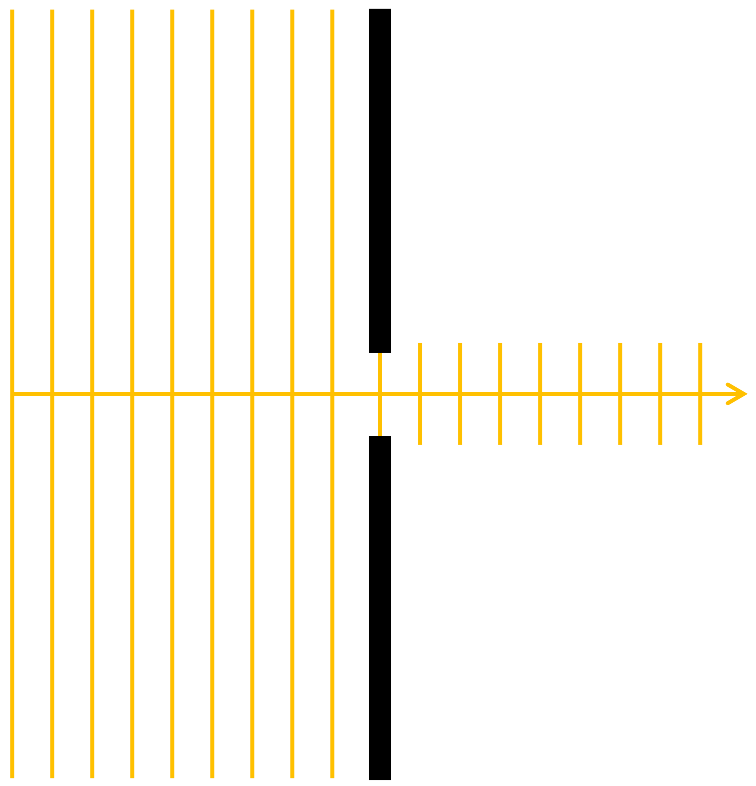

What do geometric optics predict will happen?

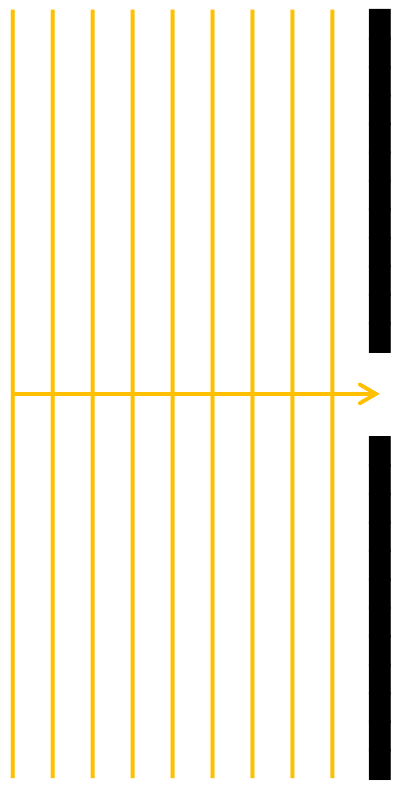

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

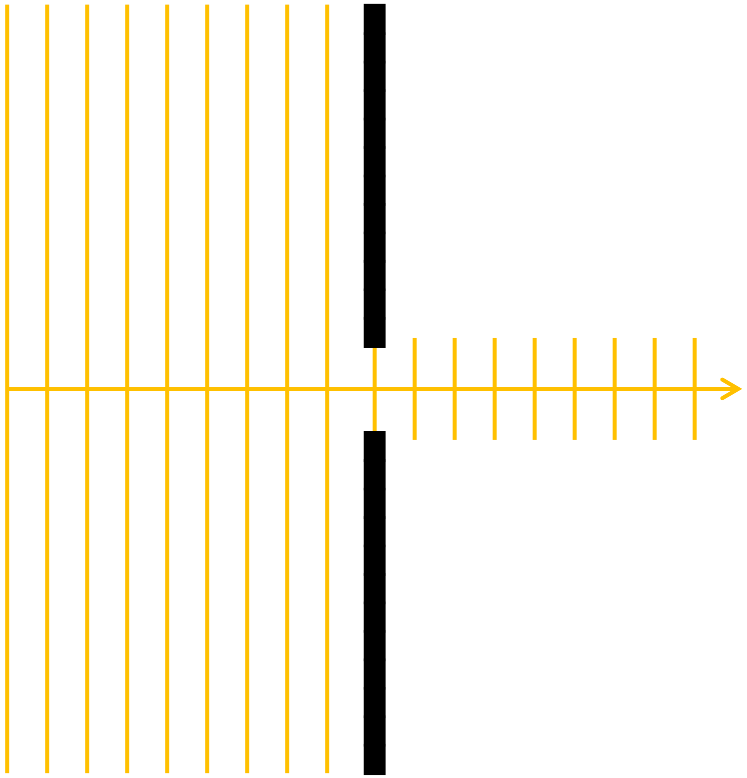

What do geometric optics predict will happen?

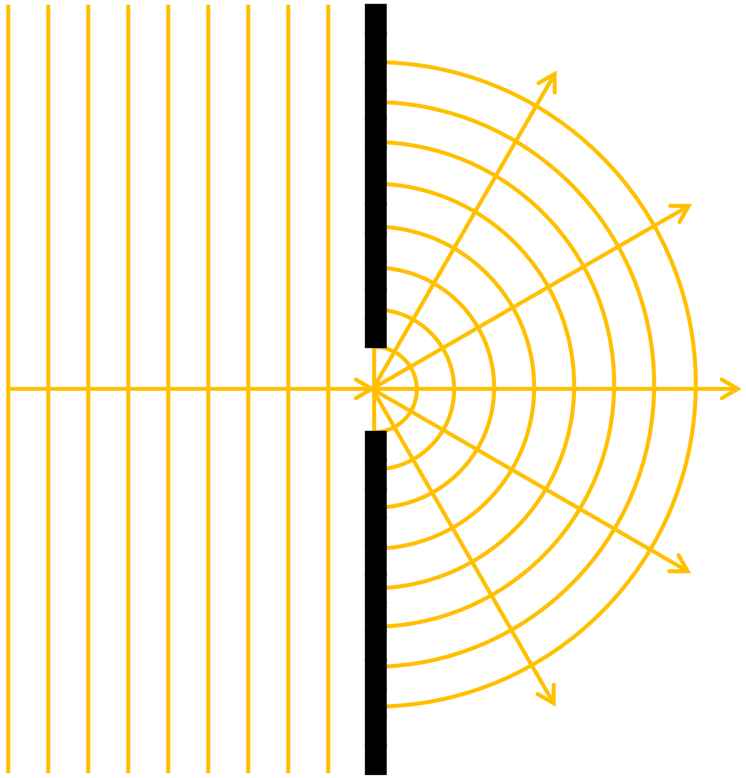

What do wave optics predict will happen?

Diffraction limit

Diffraction pattern = Fourier transform of the pinhole.

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

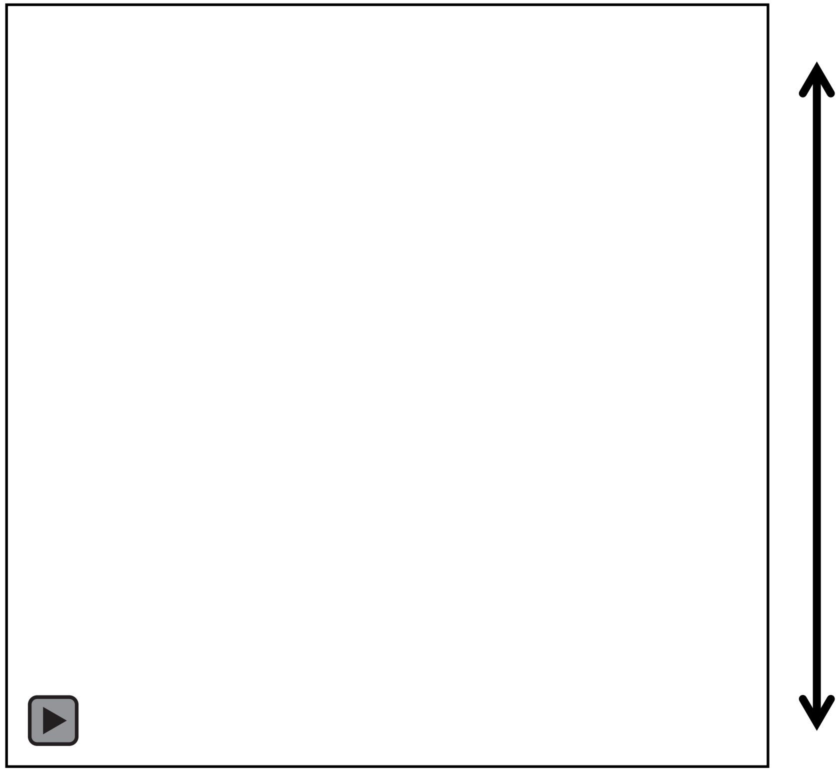

wide diffraction pattern

small pinhole

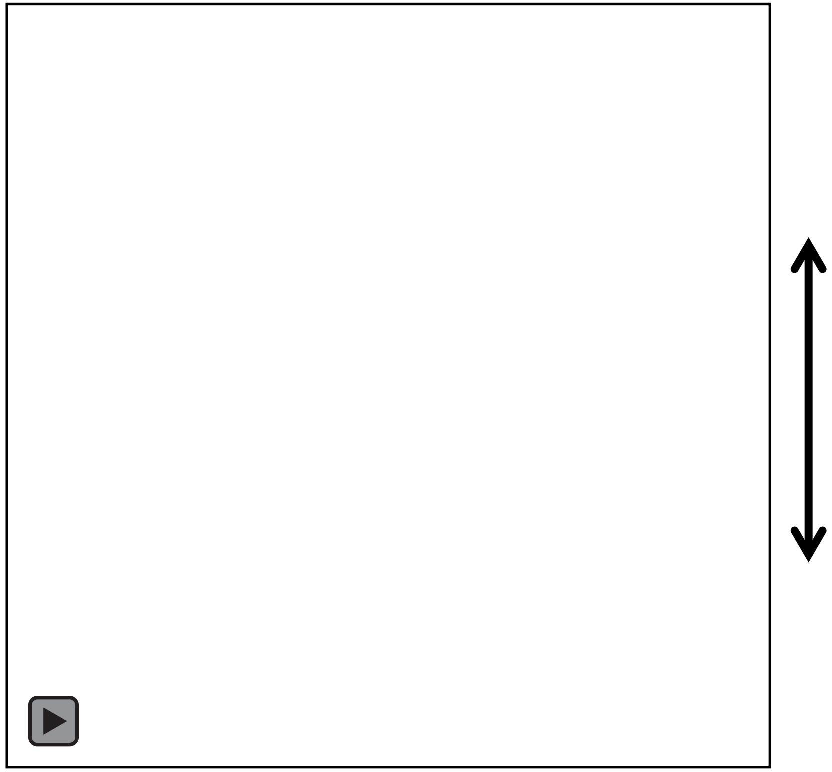

narrow diffraction pattern

large pinhole

What about light efficiency?

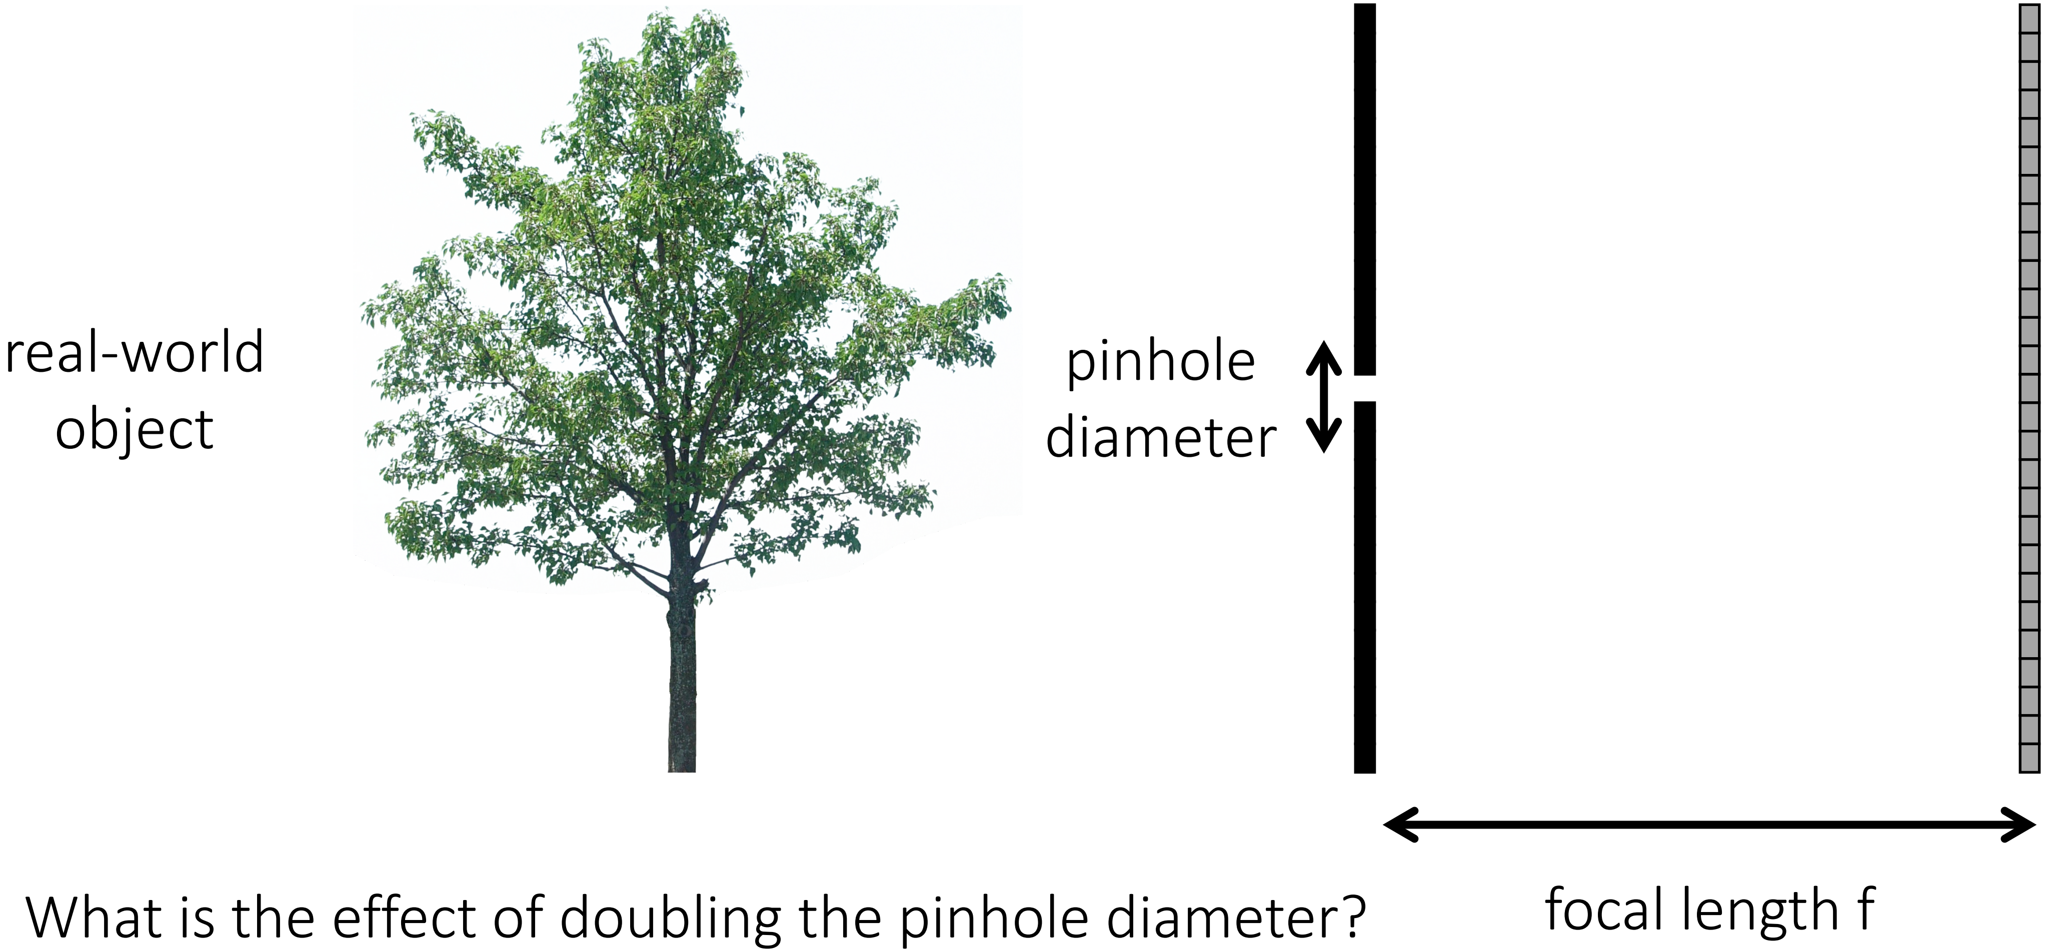

- What is the effect of doubling the pinhole diameter?
- What is the effect of doubling the focal length?

What about light efficiency?

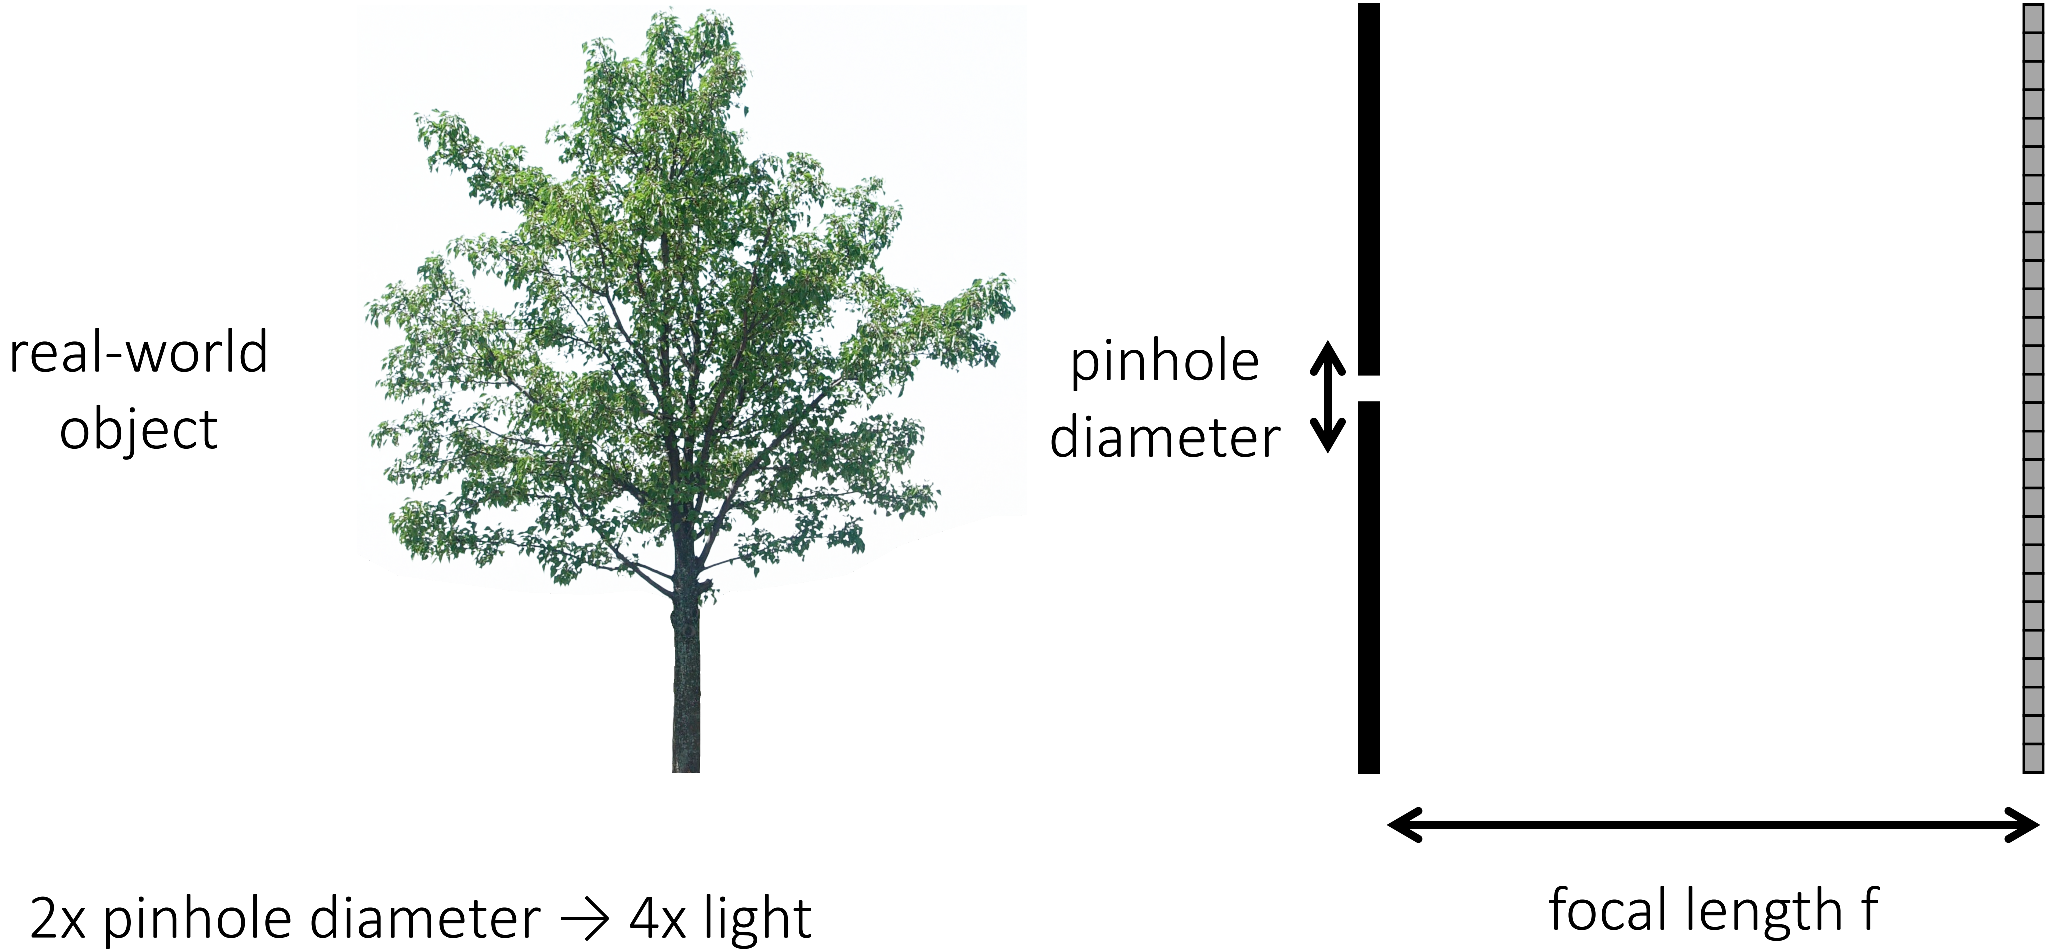

- 2x pinhole diameter → 4x light
- 2x focal length → ¼x light

Some terminology notes

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

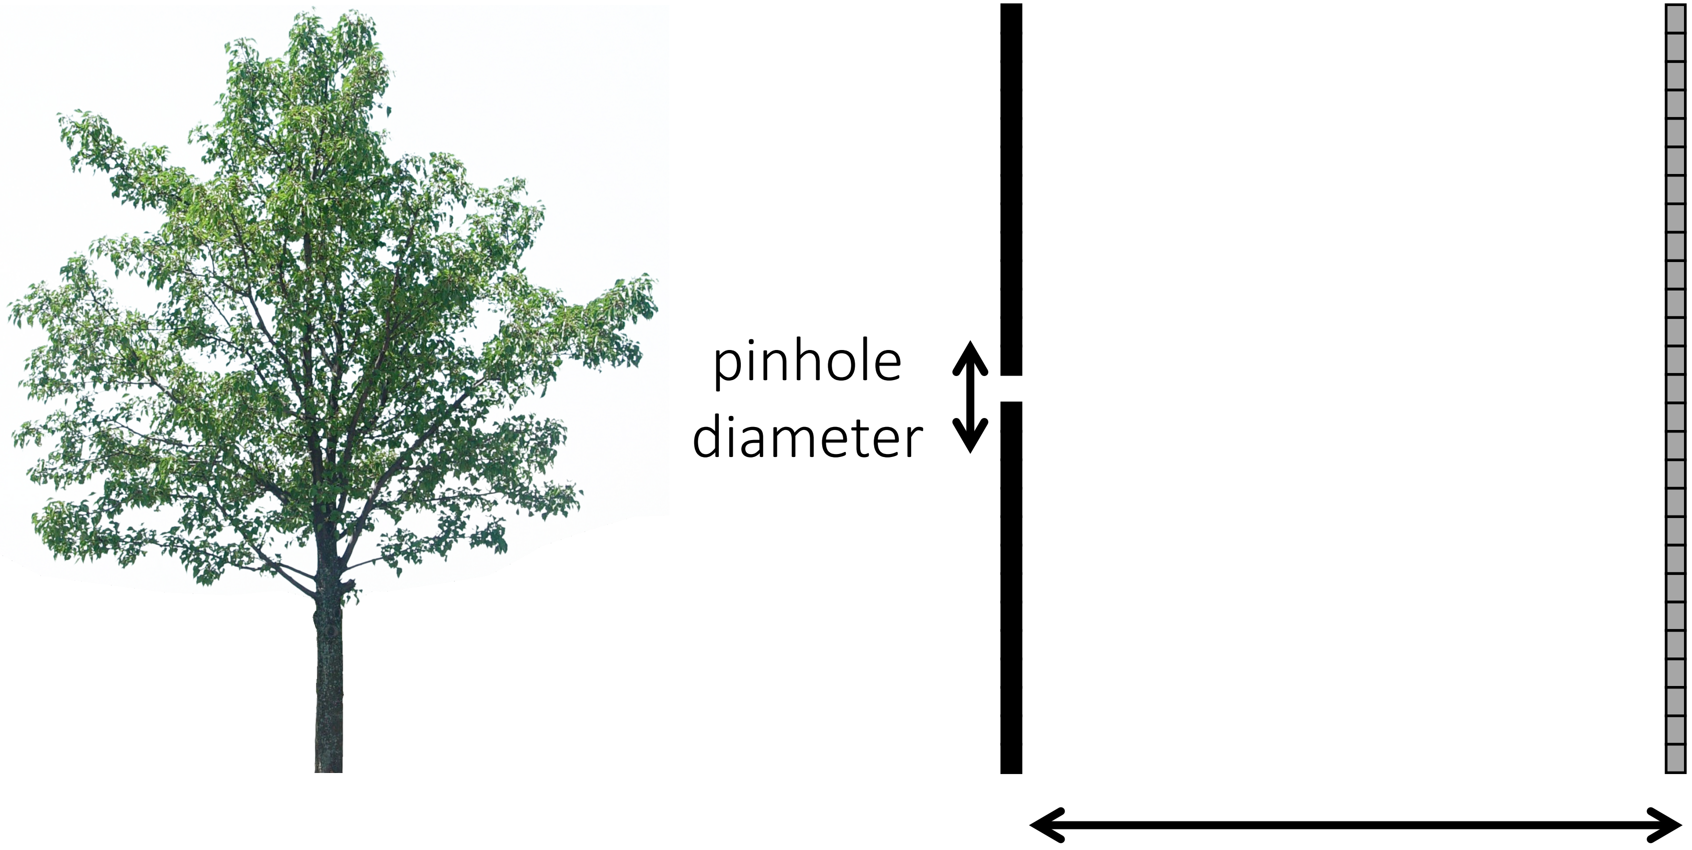

The “f-number” is the ratio: focal length / pinhole diameter

Accidental pinholes

What does this image say about the world outside?

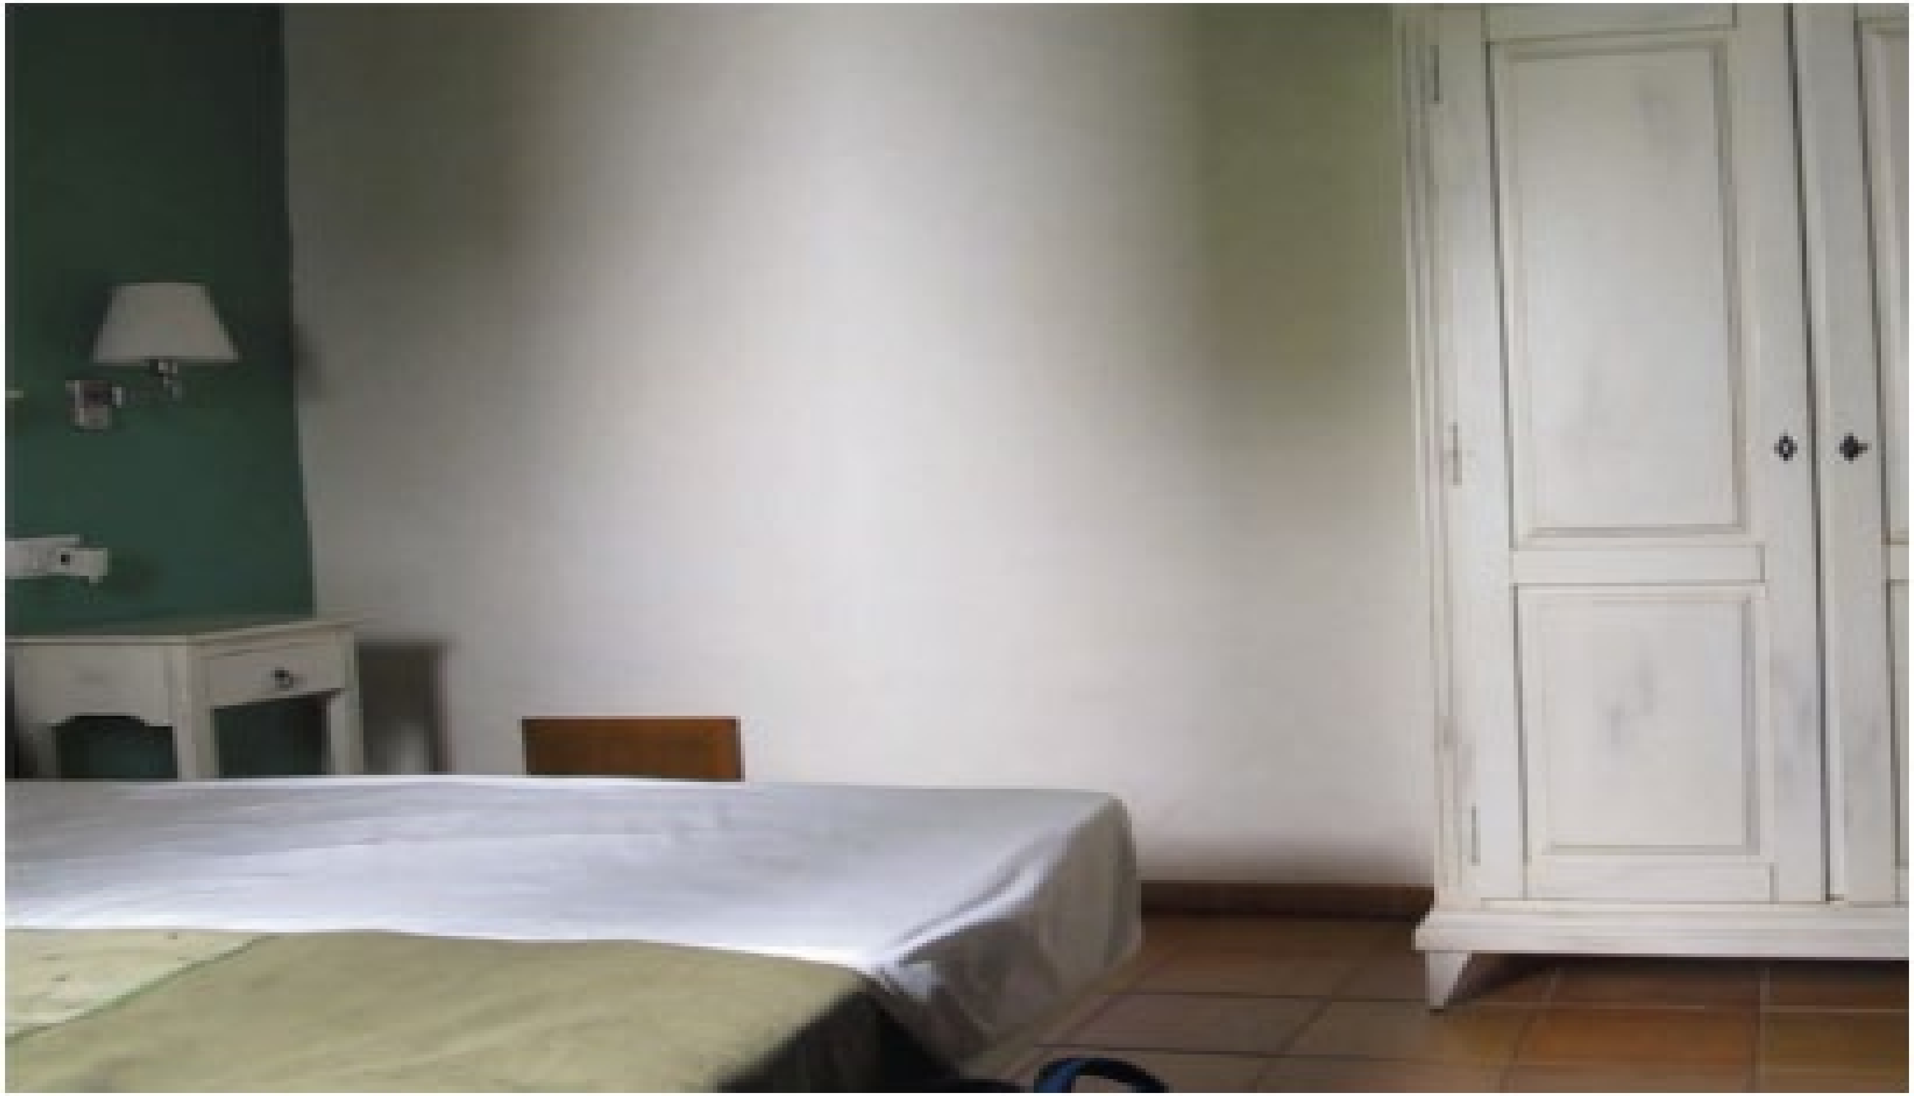

Accidental pinhole camera

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

Accidental pinhole camera

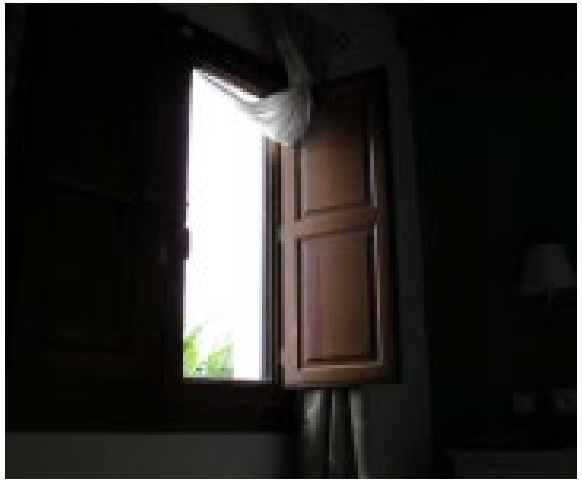

window is an aperture

projected pattern on the wall

upside down

window with smaller gap

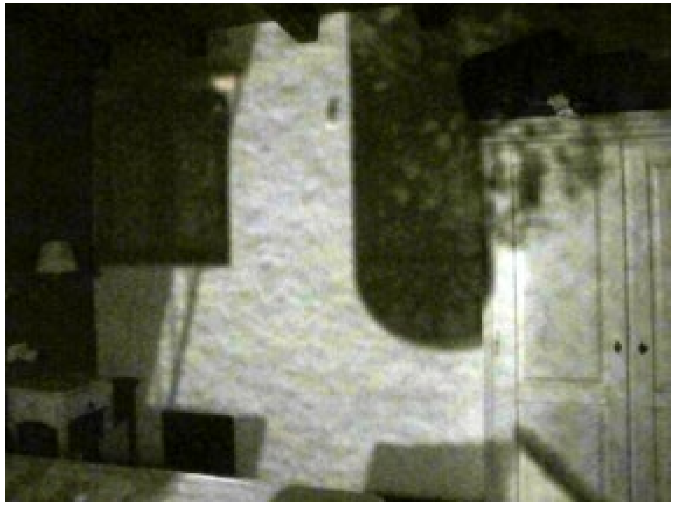

view outside window

Accidental pinspeck camera

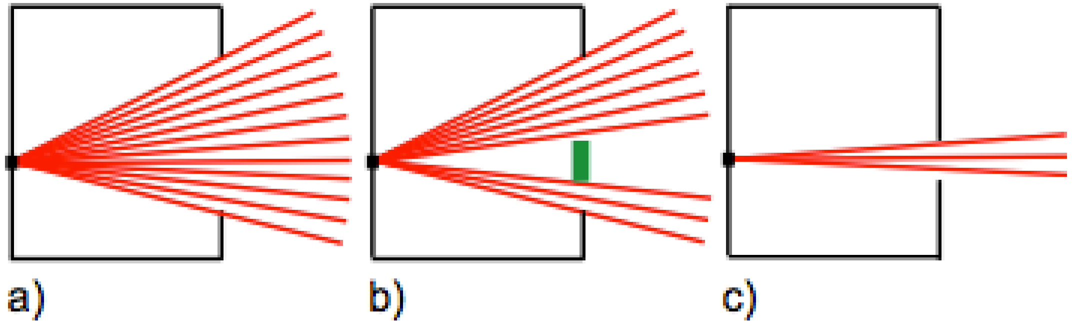

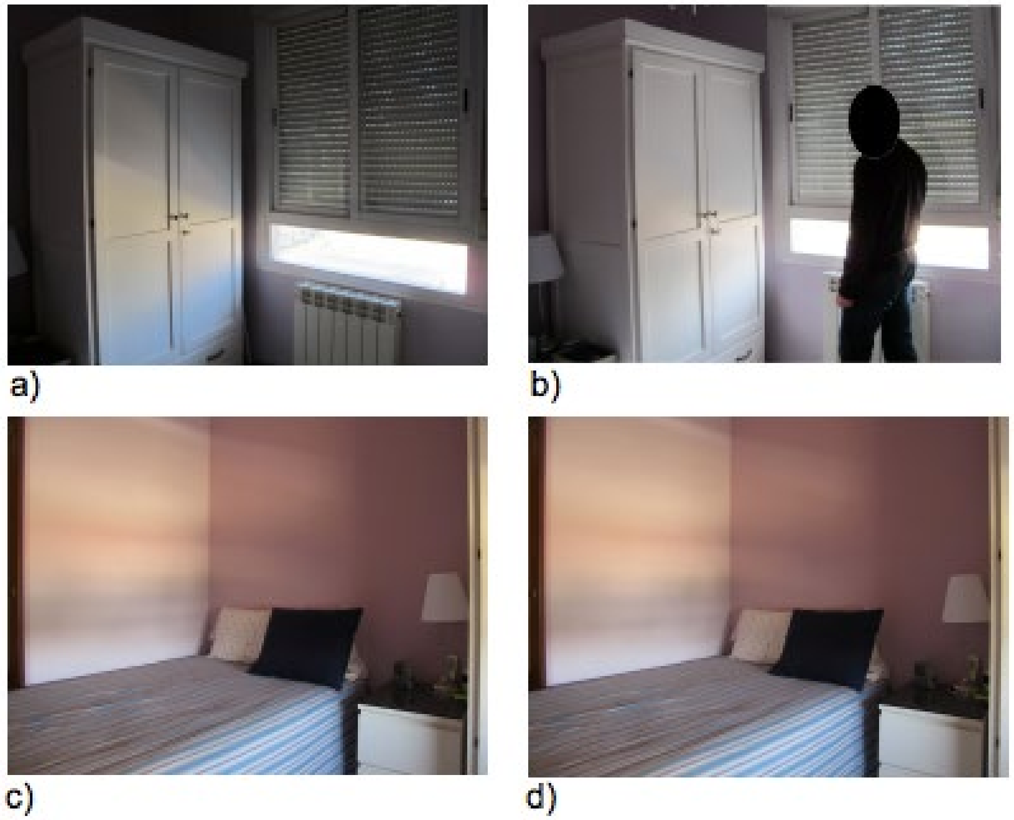

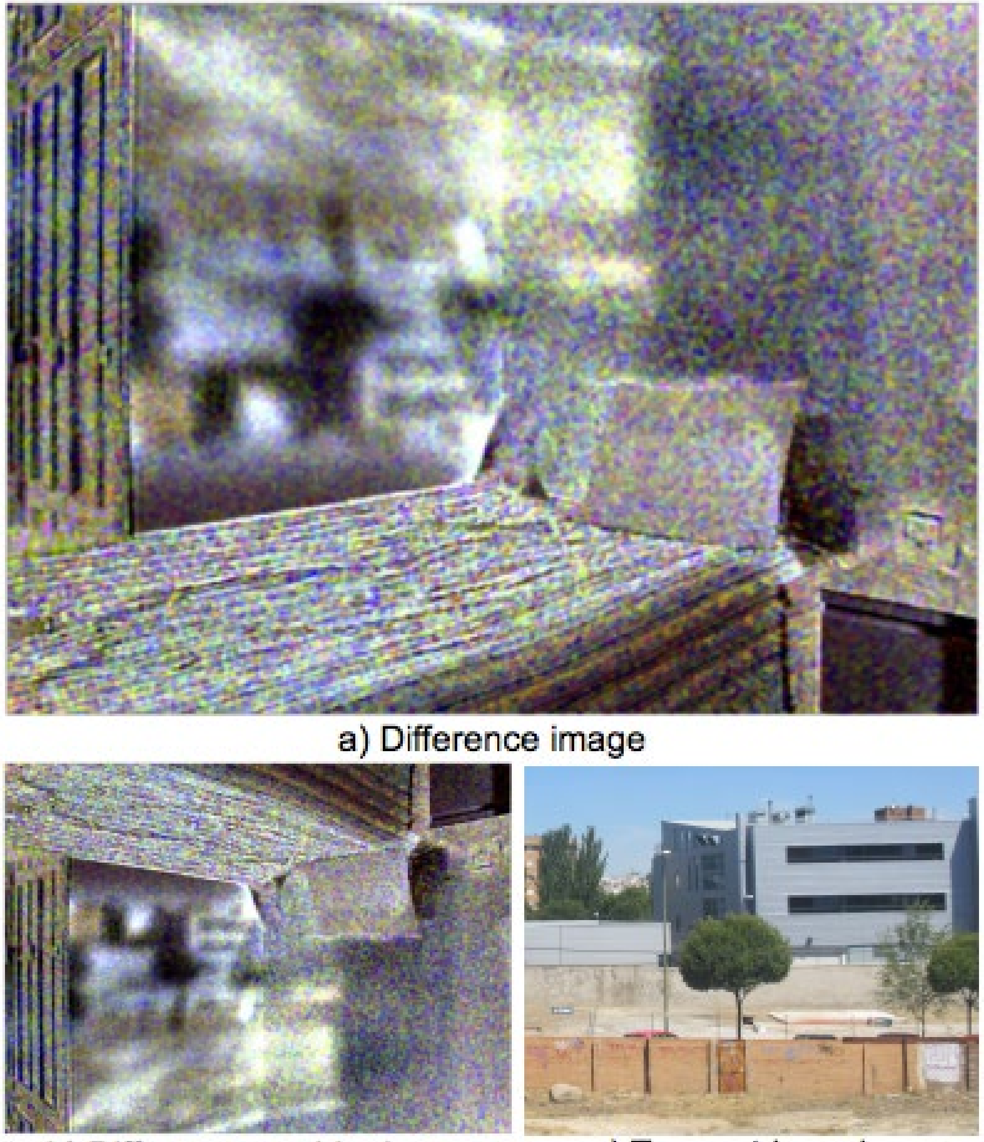
a) Difference image
b) Difference upside down
c) True outdoor view

Pinhole camera trade-off

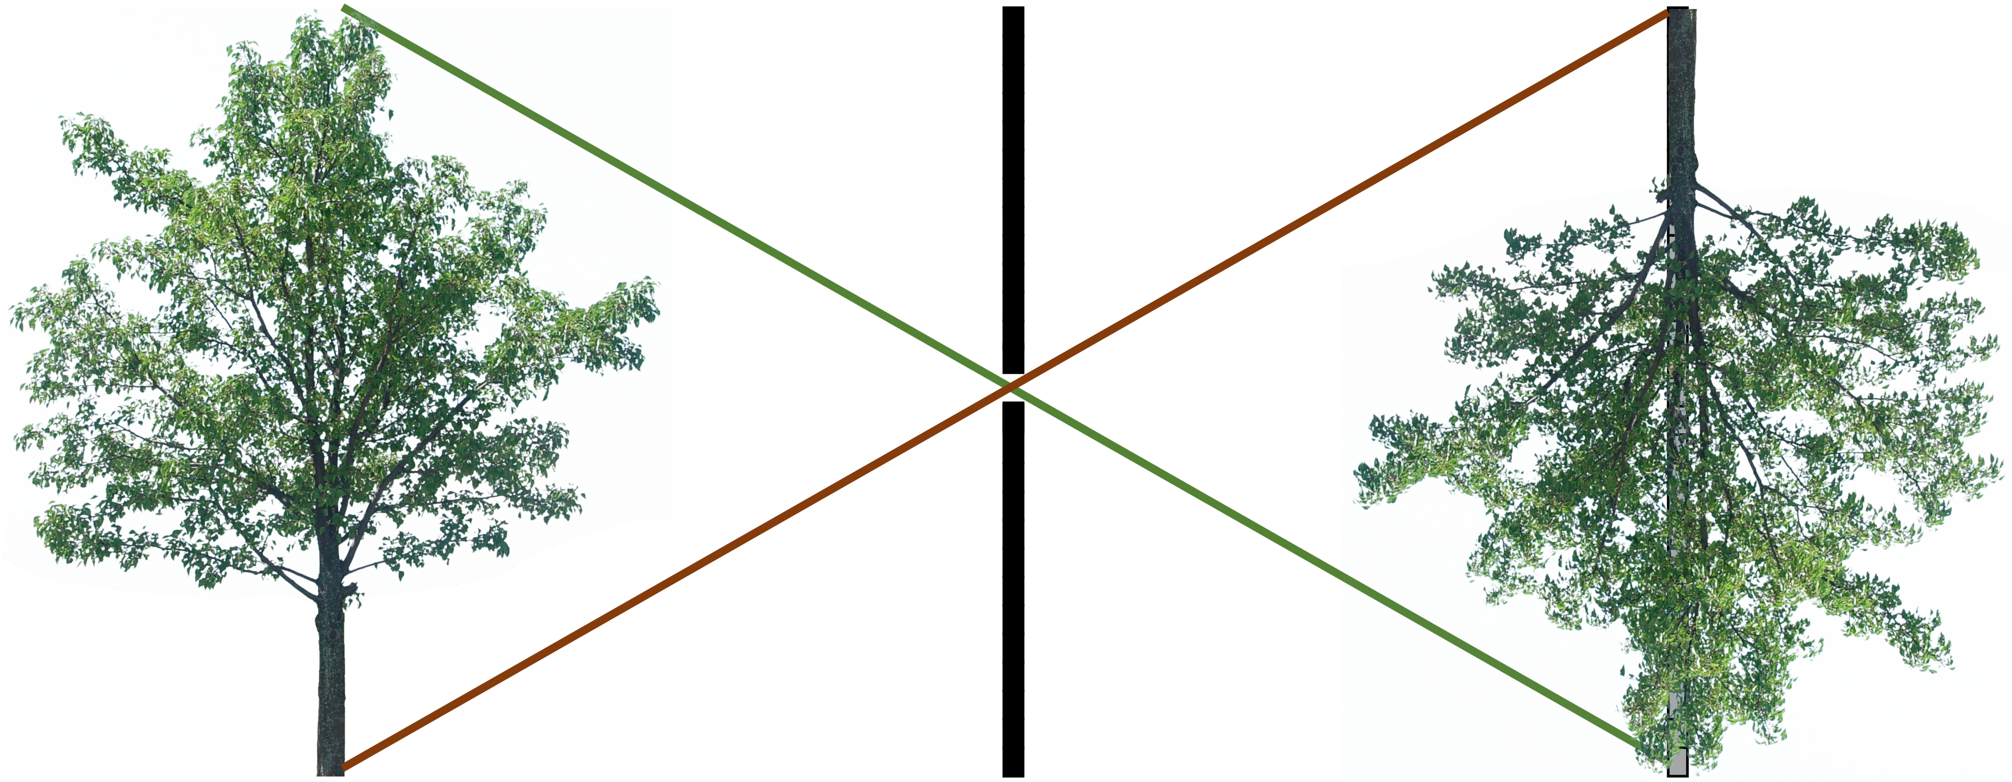

Small (ideal) pinhole:

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

Pinhole camera trade-off

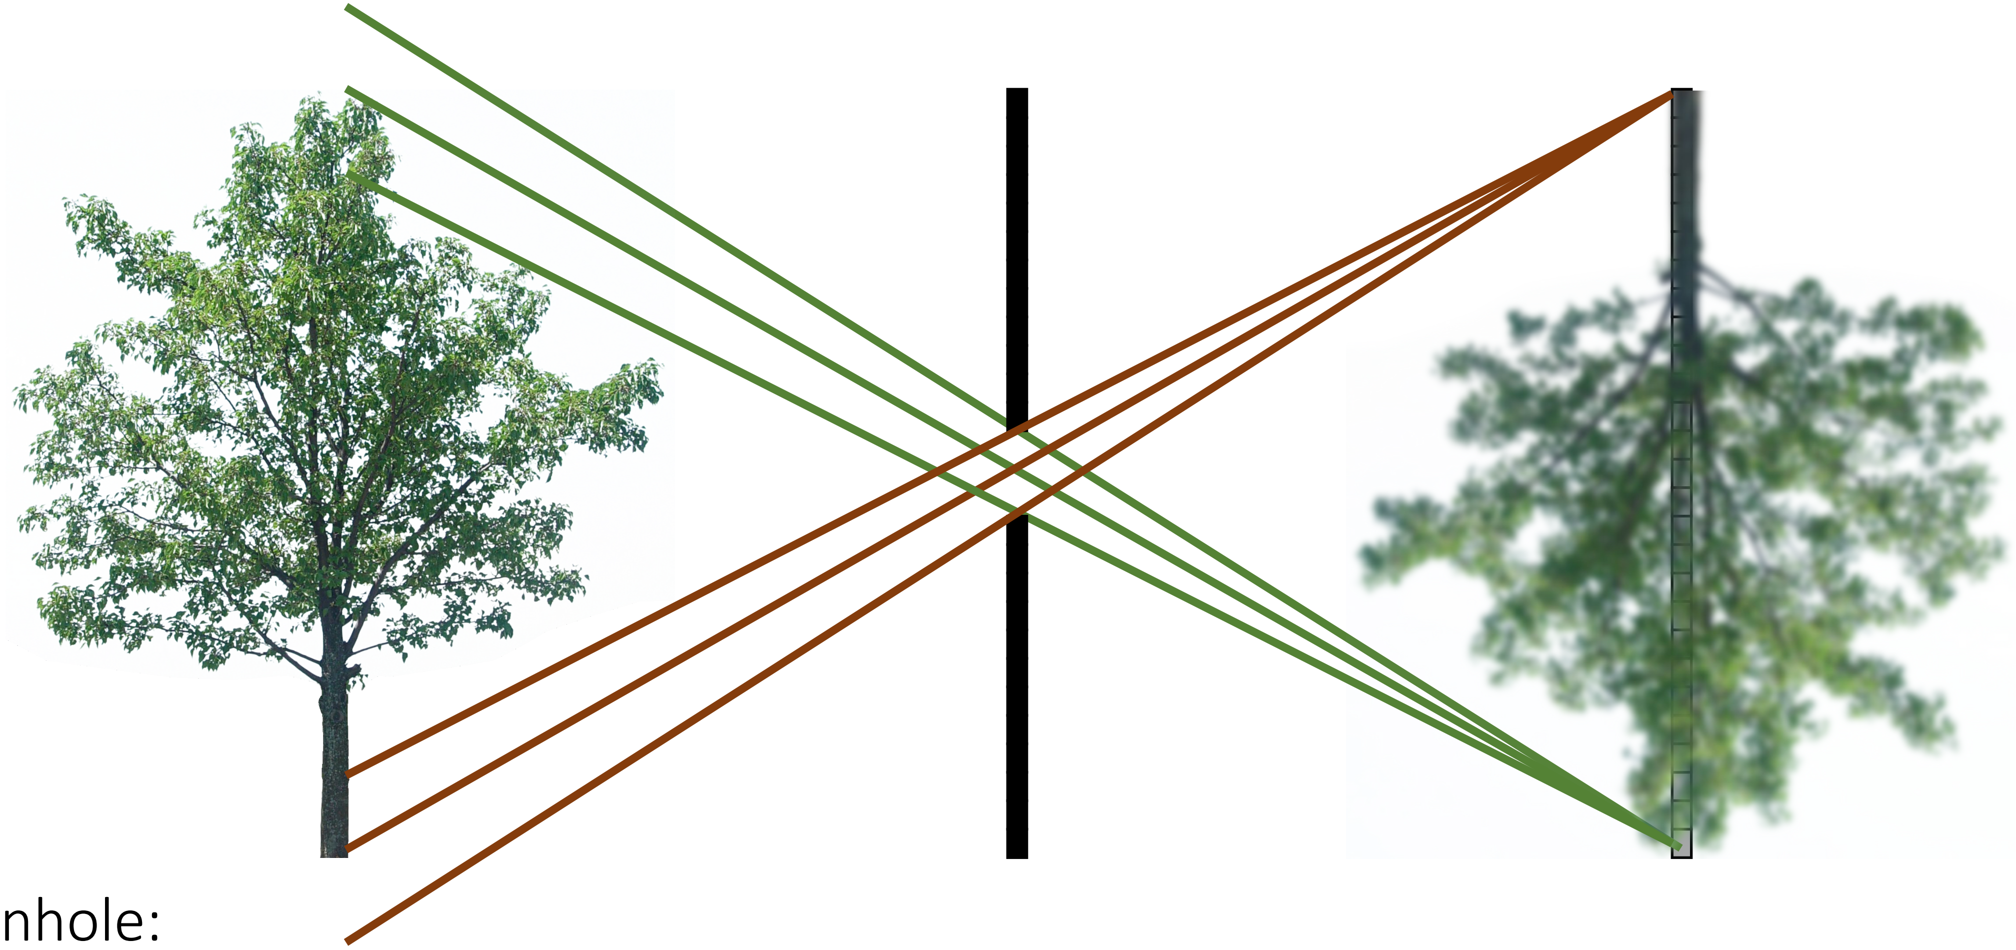

Large pinhole:
1. Image is blurry.
2. Signal-to-noise ratio is high.

Can we get best of both worlds?

Almost, by using lenses

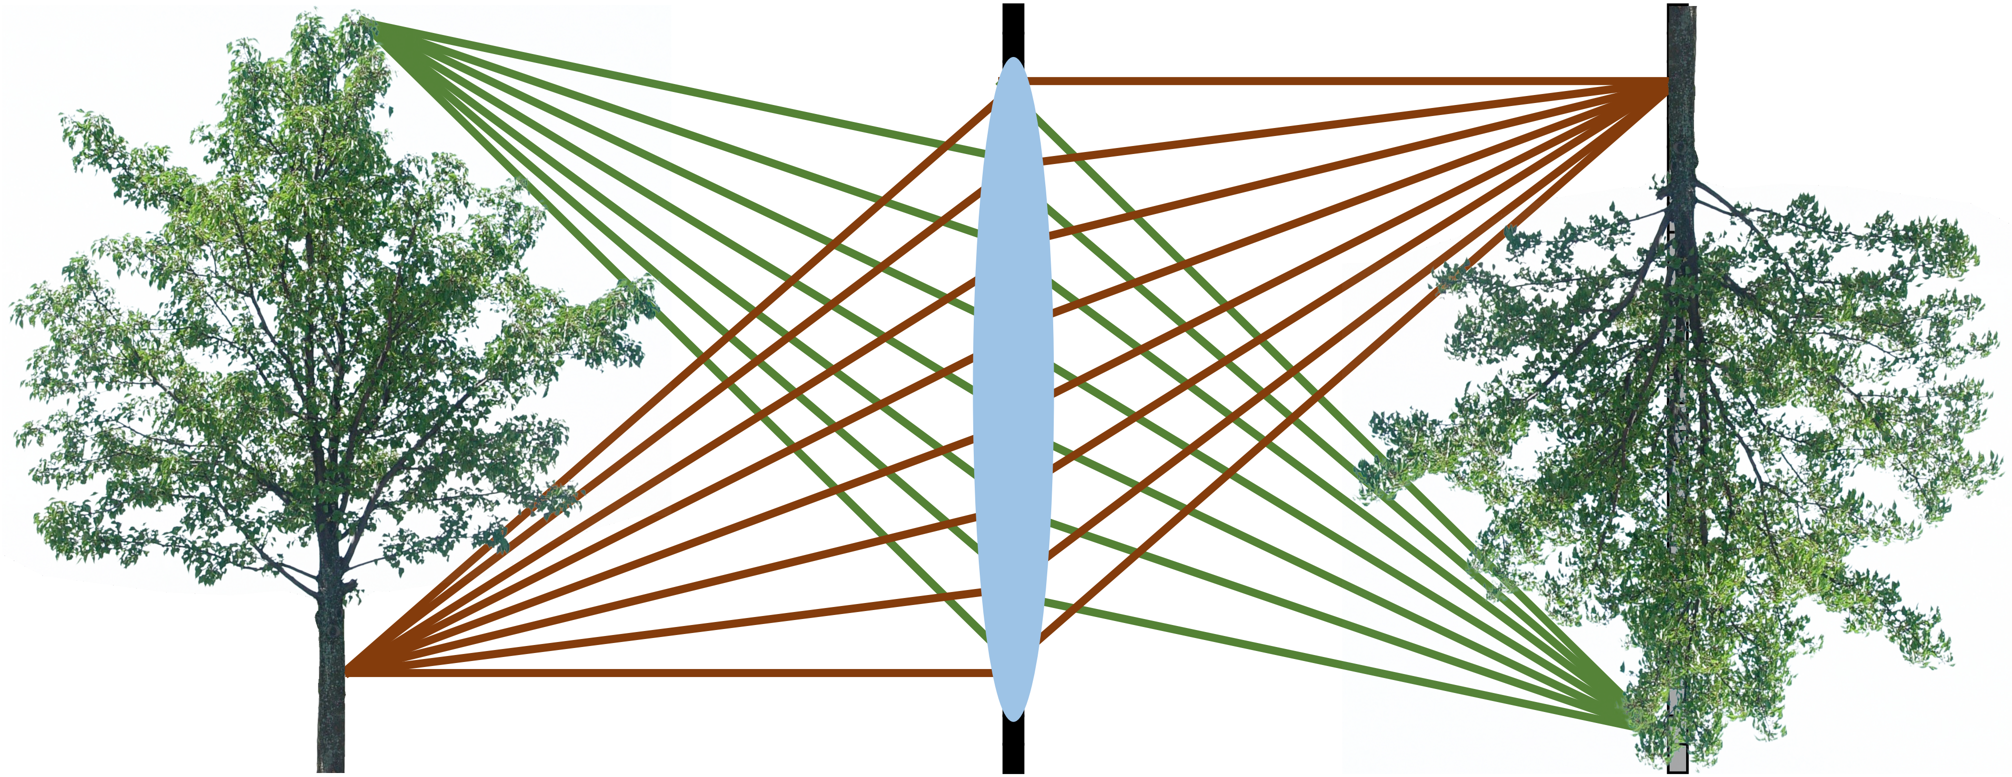

Lenses map “bundles” of rays from points on the scene to the sensor.

How does this mapping work exactly?

Lens (very) basics

What is a lens?

A piece of glass manufactured to have a specific shape

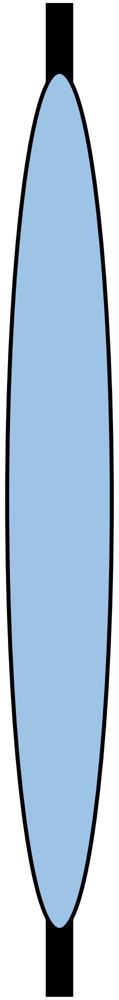

What is a lens?

A piece of glass manufactured to have a specific shape

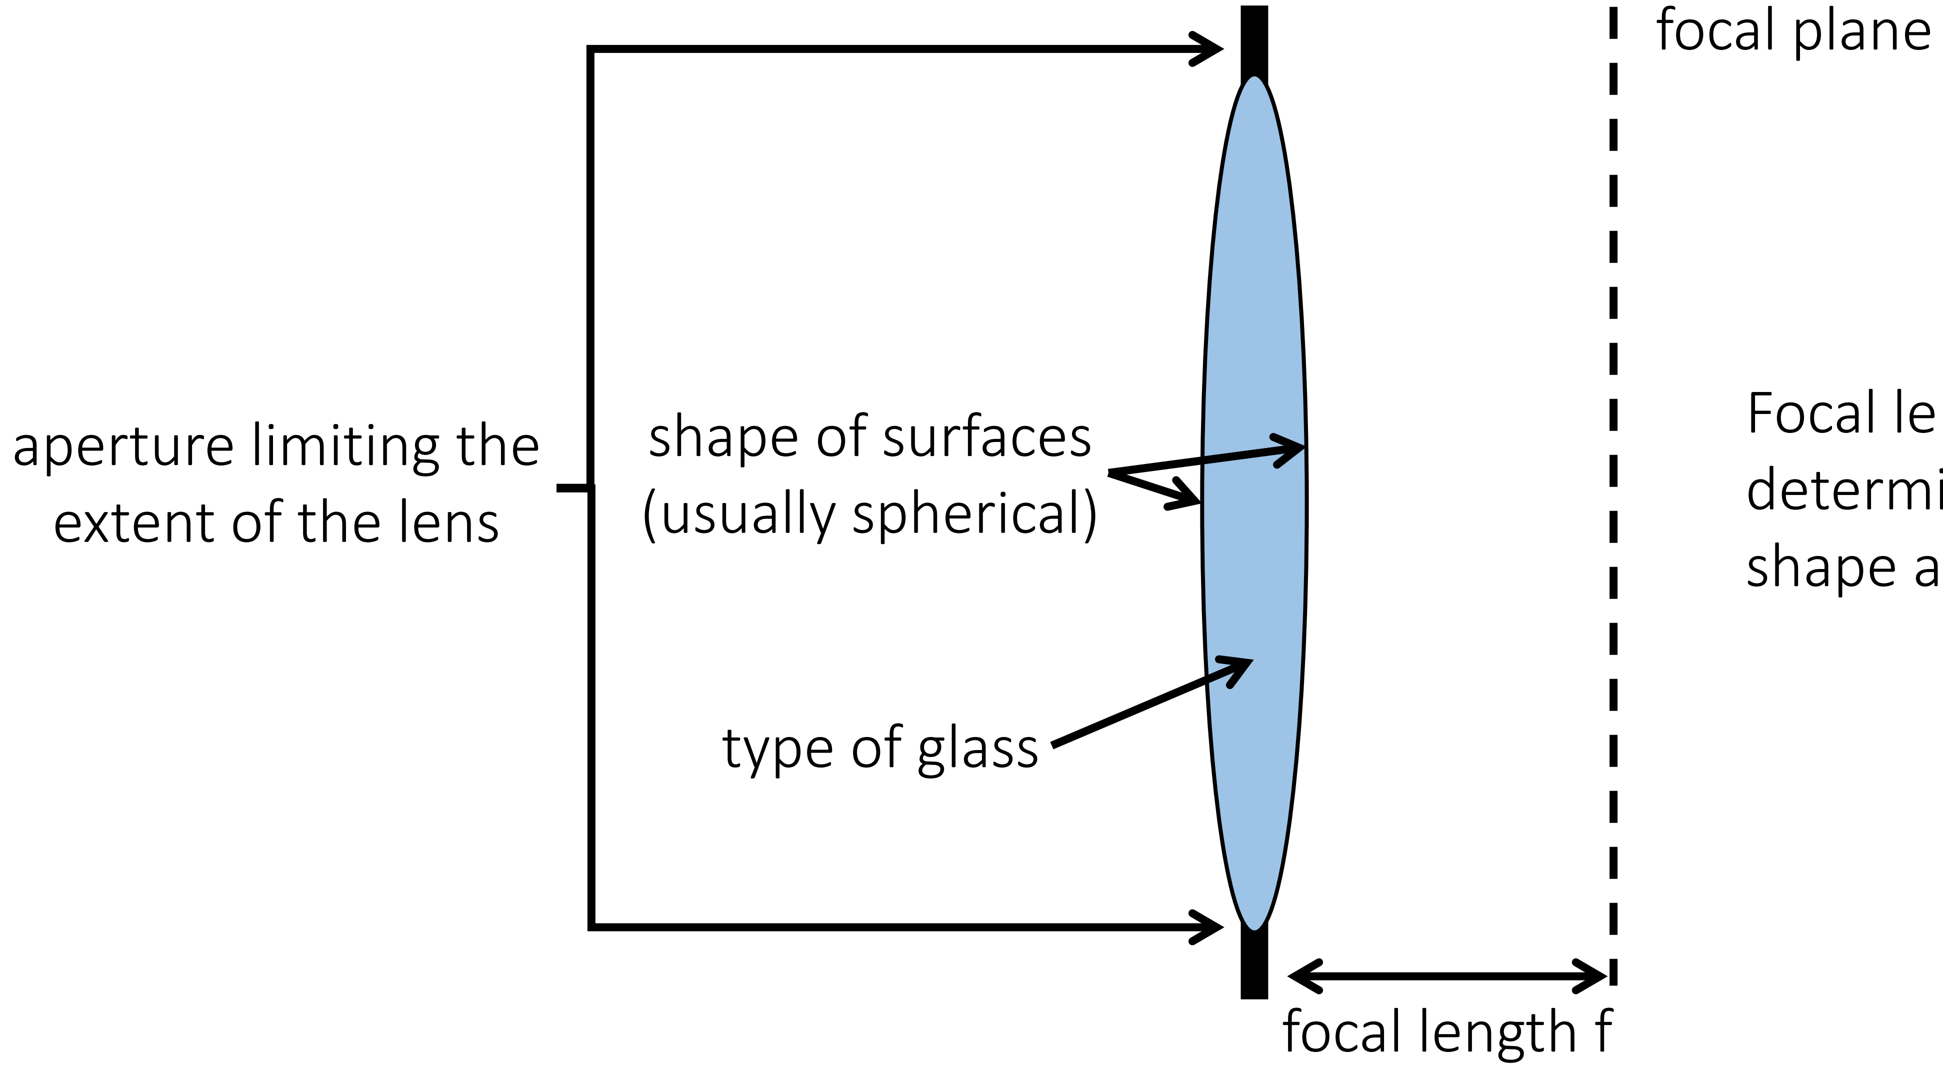

Focal length is determined by the lens' shape and material

The lens on your camera

The reality: compound lenses

Many pieces of glass manufactured to have a specific shape, and placed in a specific configuration

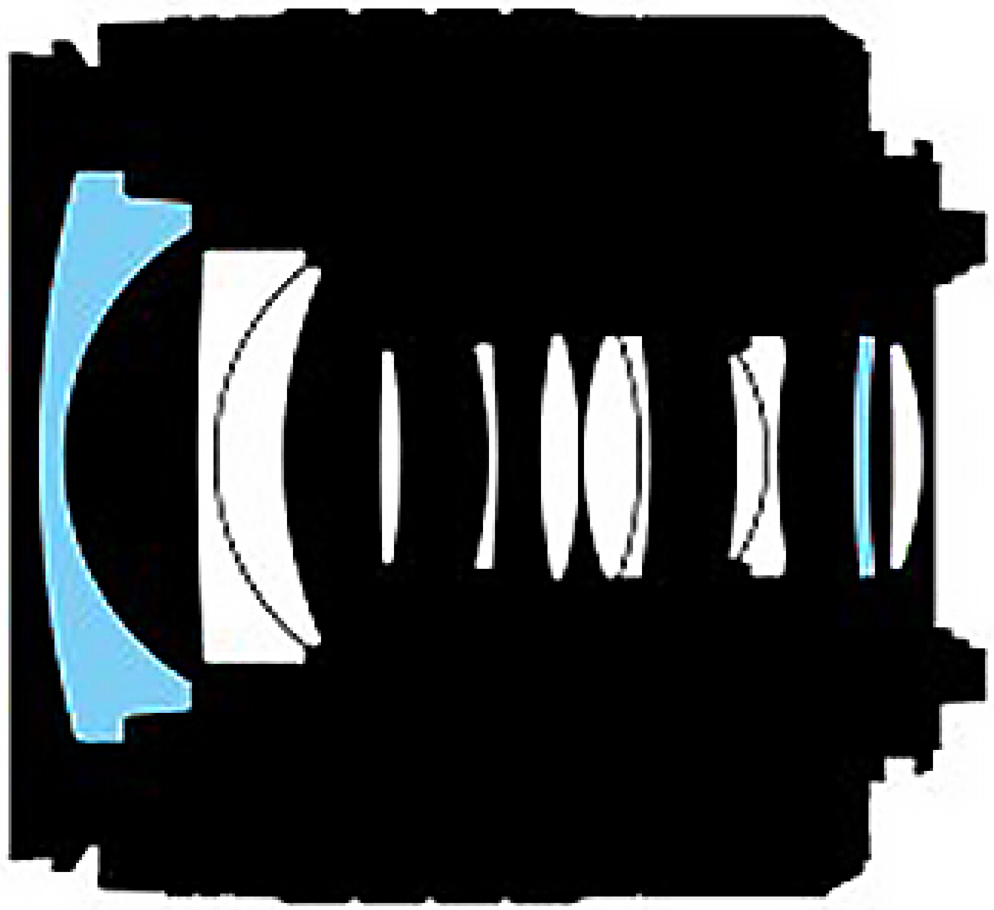

Cross-section of Nikon 18-55 mm lens

The *effective* aperture size and focal length are determined by the lenses' shape, material, and relative placement.

In this lecture, we will deal with single-element "thin" lenses. More on compound lenses in the next lecture.

How does a lens work?

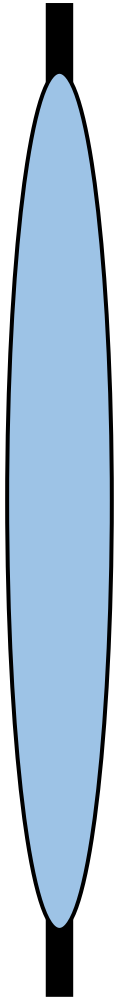

Refraction

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

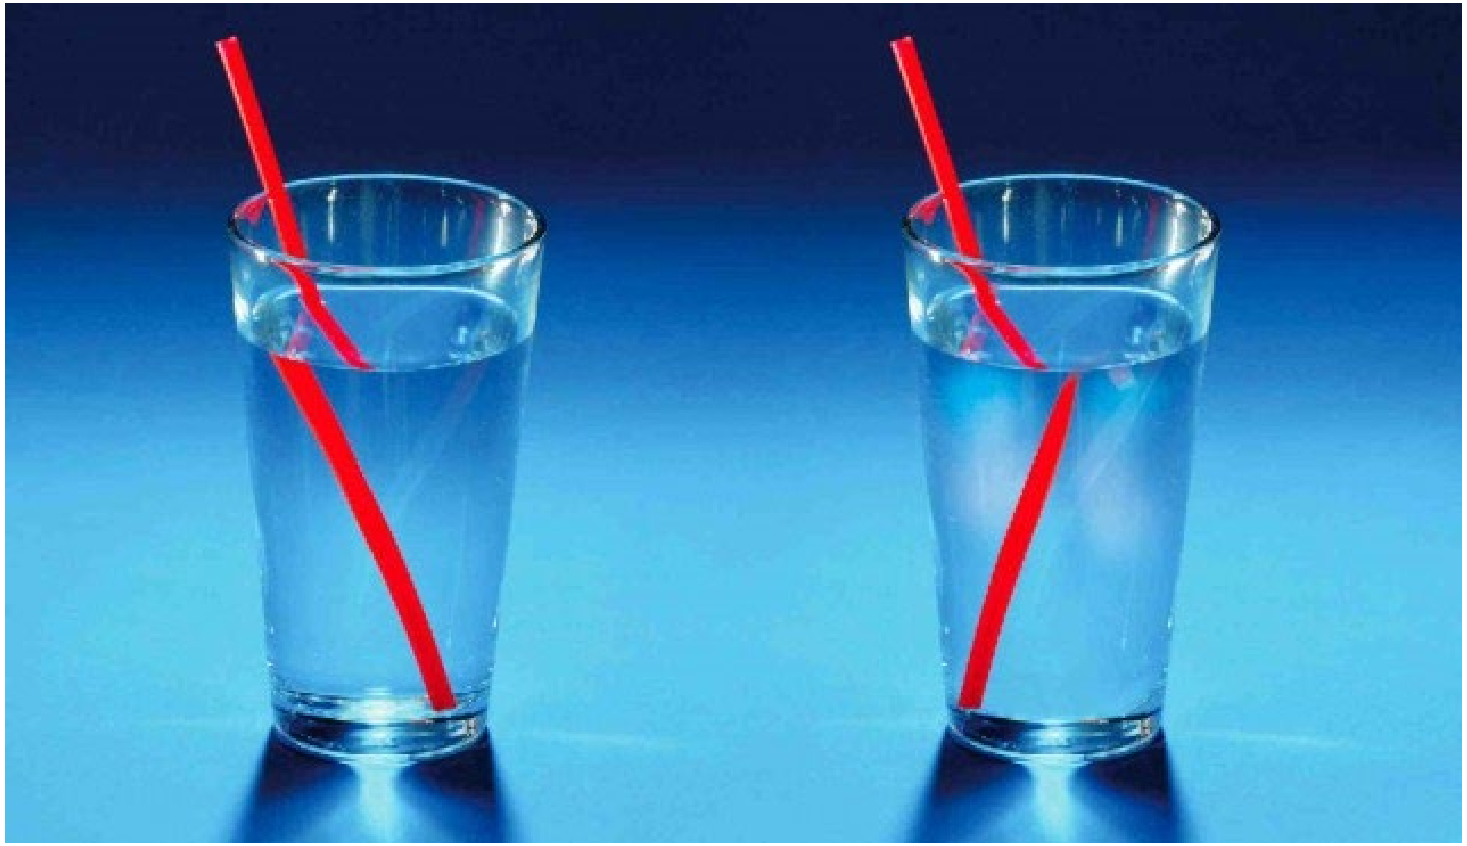

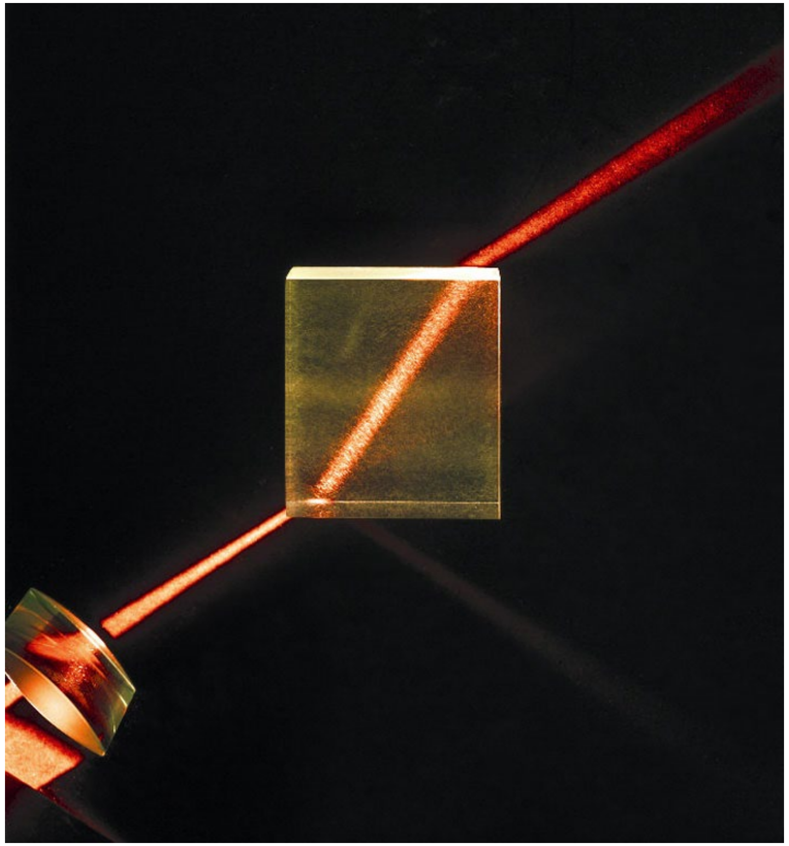

How does a lens work?

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

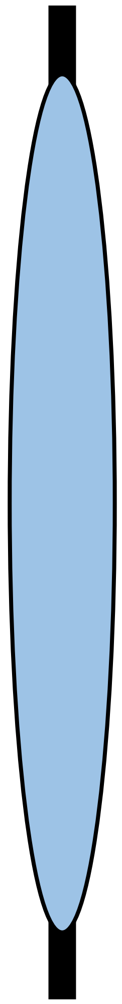

The thin lens model

Thin lens model

Idealized geometric optics model for <u>well-designed</u> lenses.

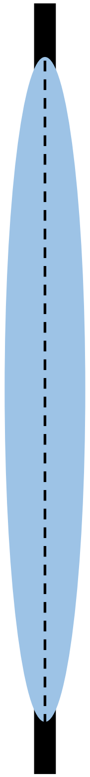

The reality: compound lenses

Remember, real lenses are not “thin” lenses.

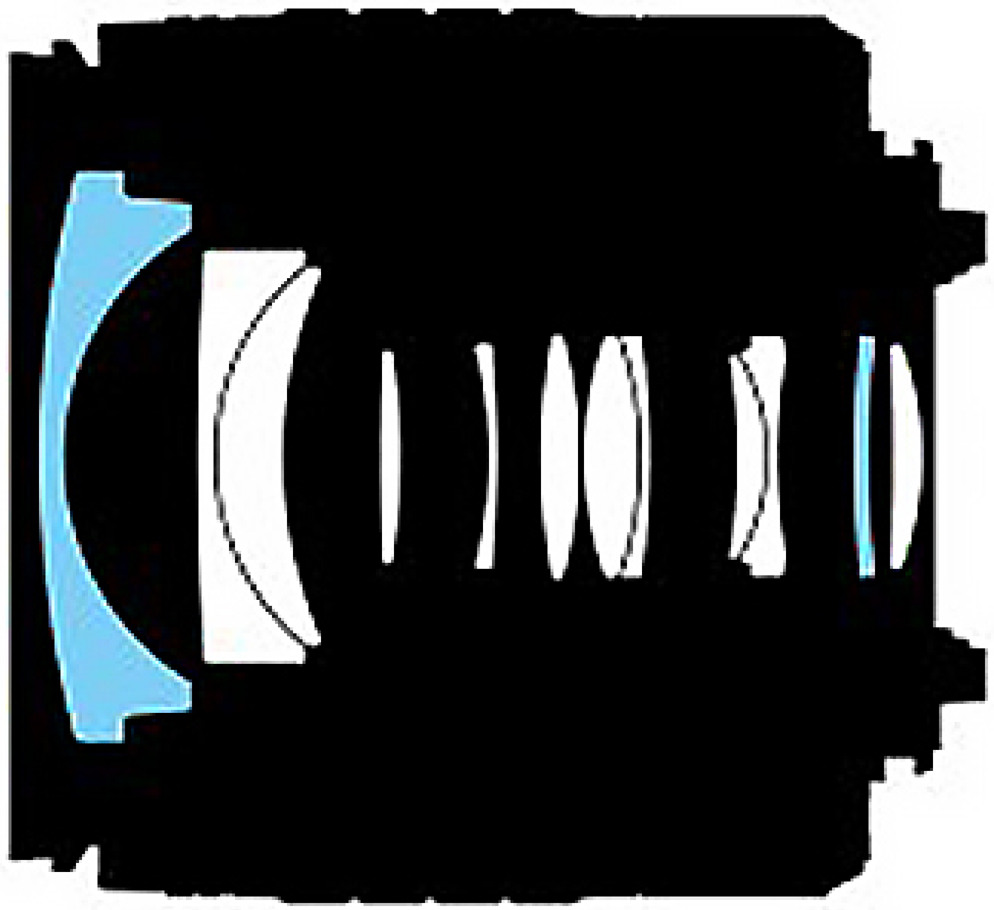

Cross-section of Nikon 18-55 mm lens

We will see in the next lecture why we can approximate a thin lenses with such a compound lens.

Thin lens model

Simplification of geometric optics for well-designed lenses.

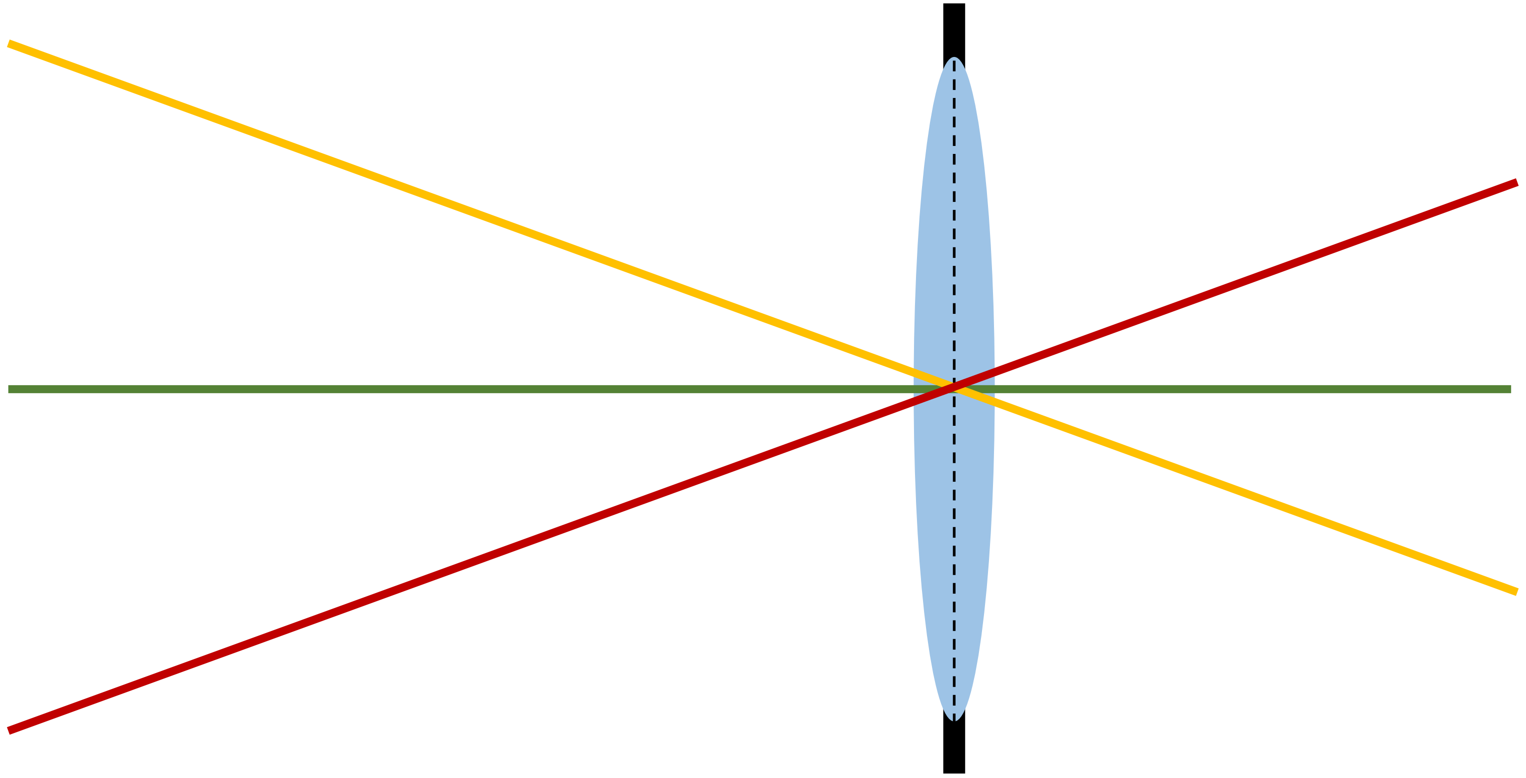

Two assumptions:

1. Rays passing through lens center are unaffected.

Thin lens model

Simplification of geometric optics for well-designed lenses.

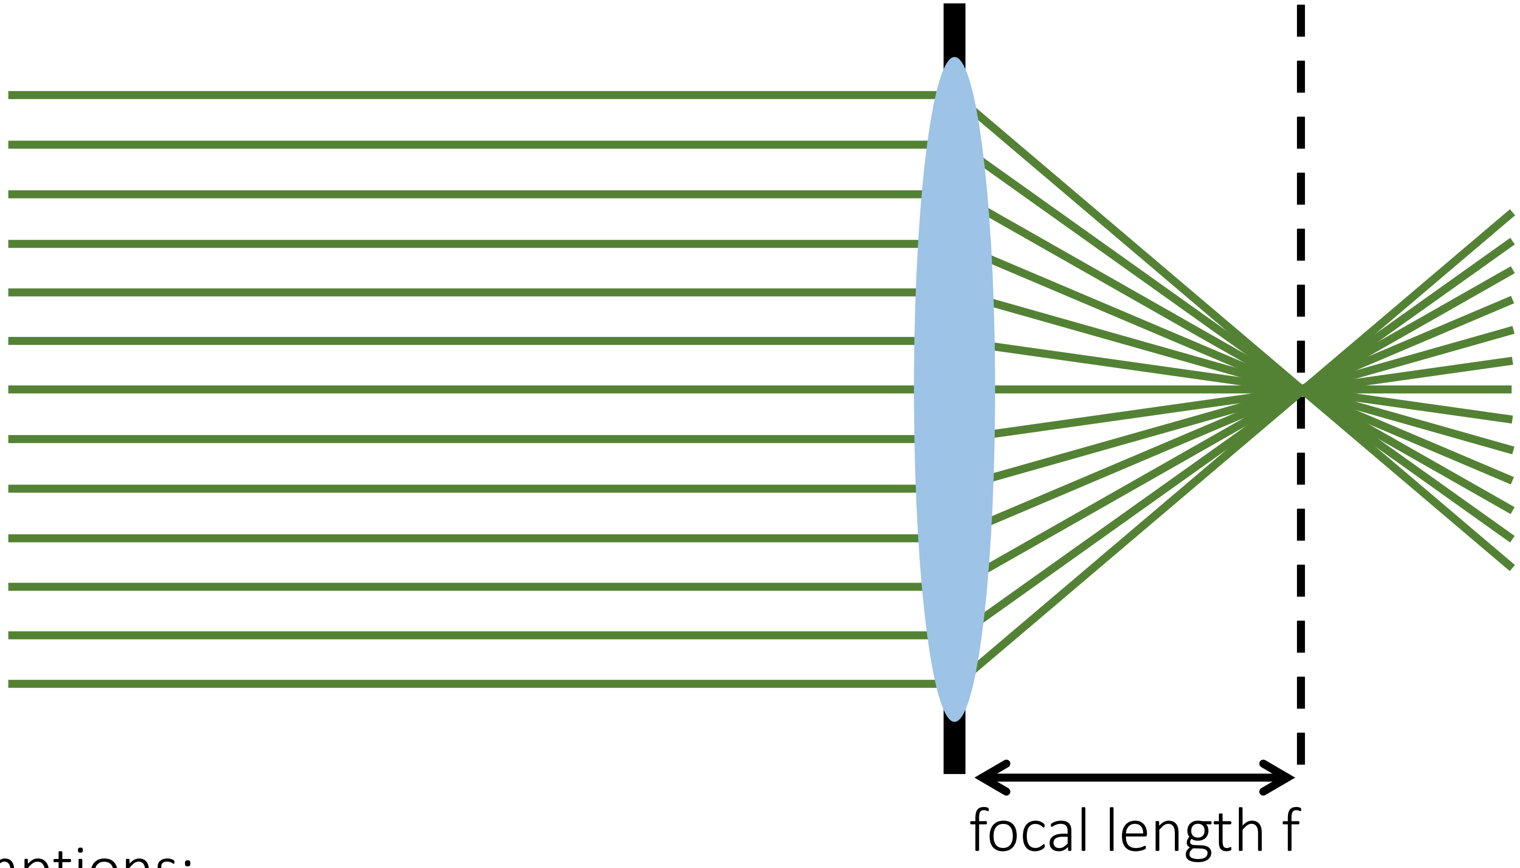

Two assumptions:

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.

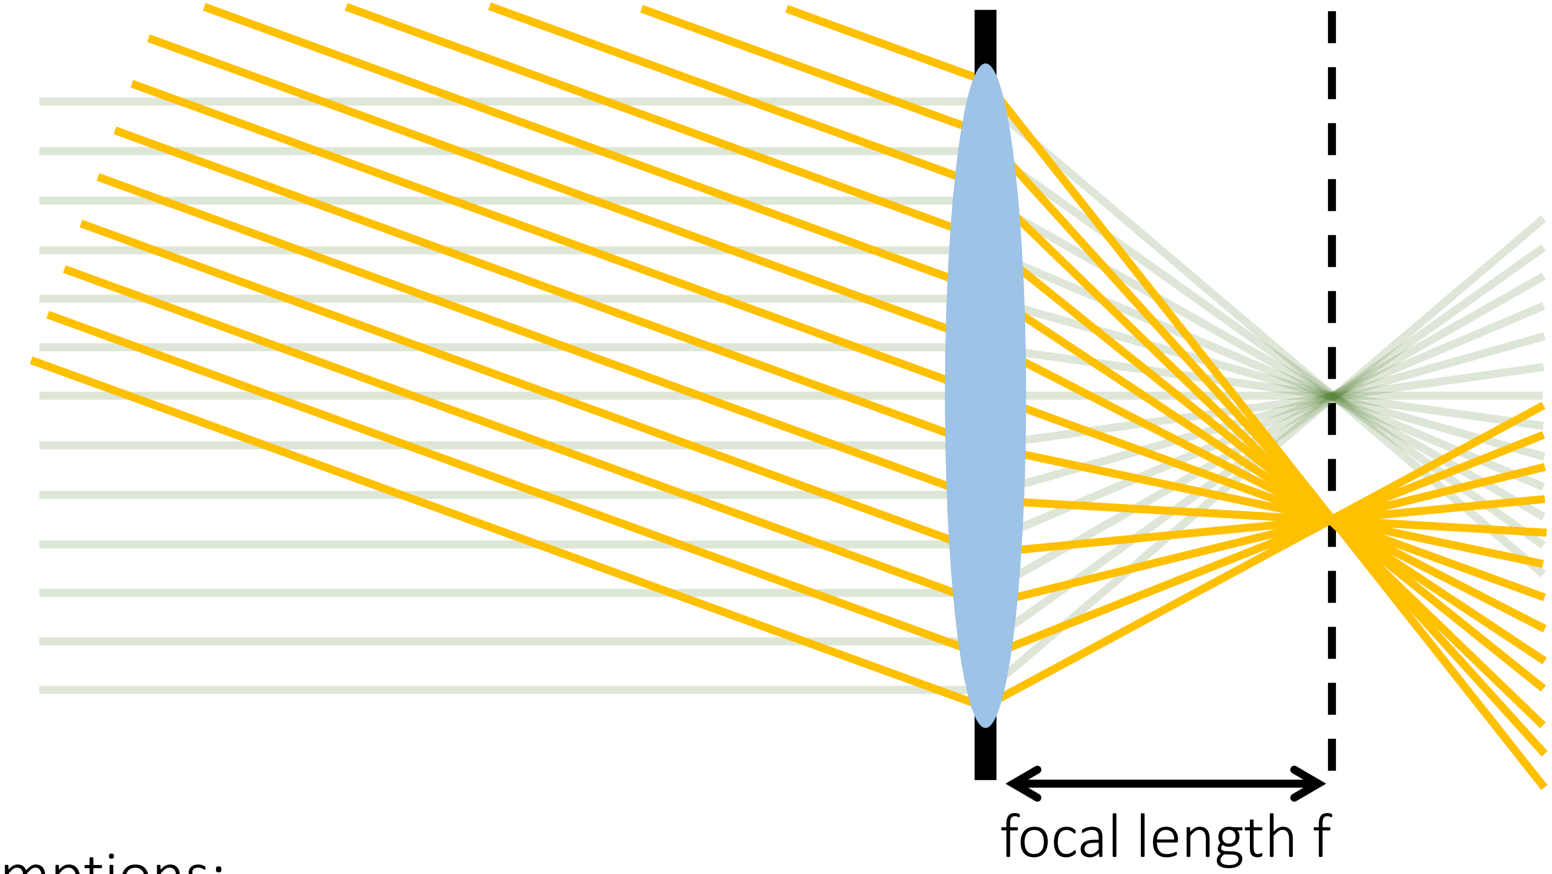

Two assumptions:

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.

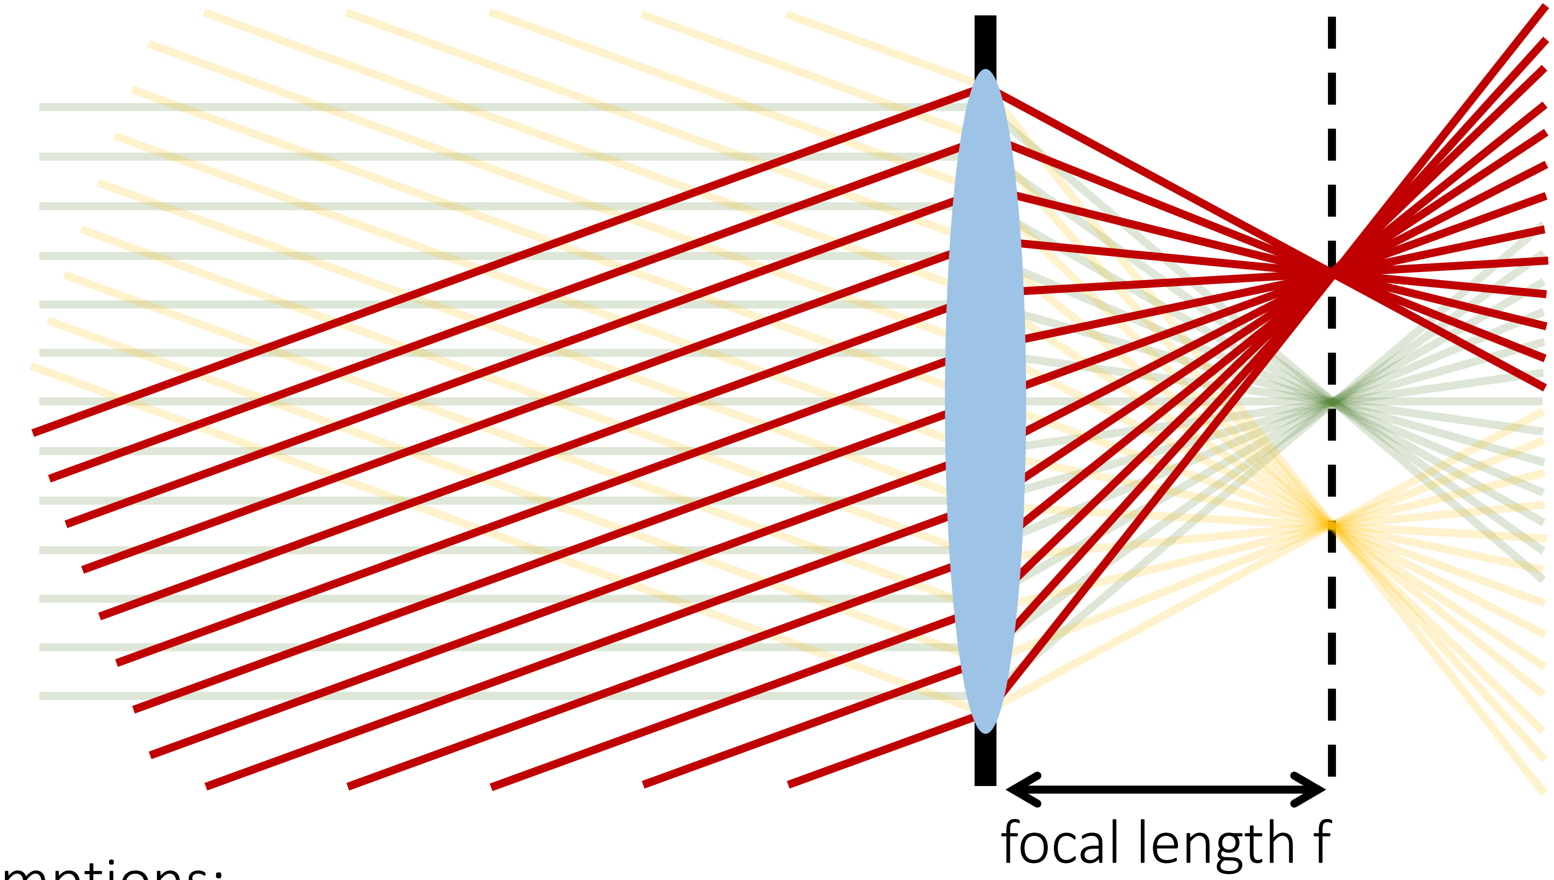

Two assumptions:

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

Tracing rays through a thin lens

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

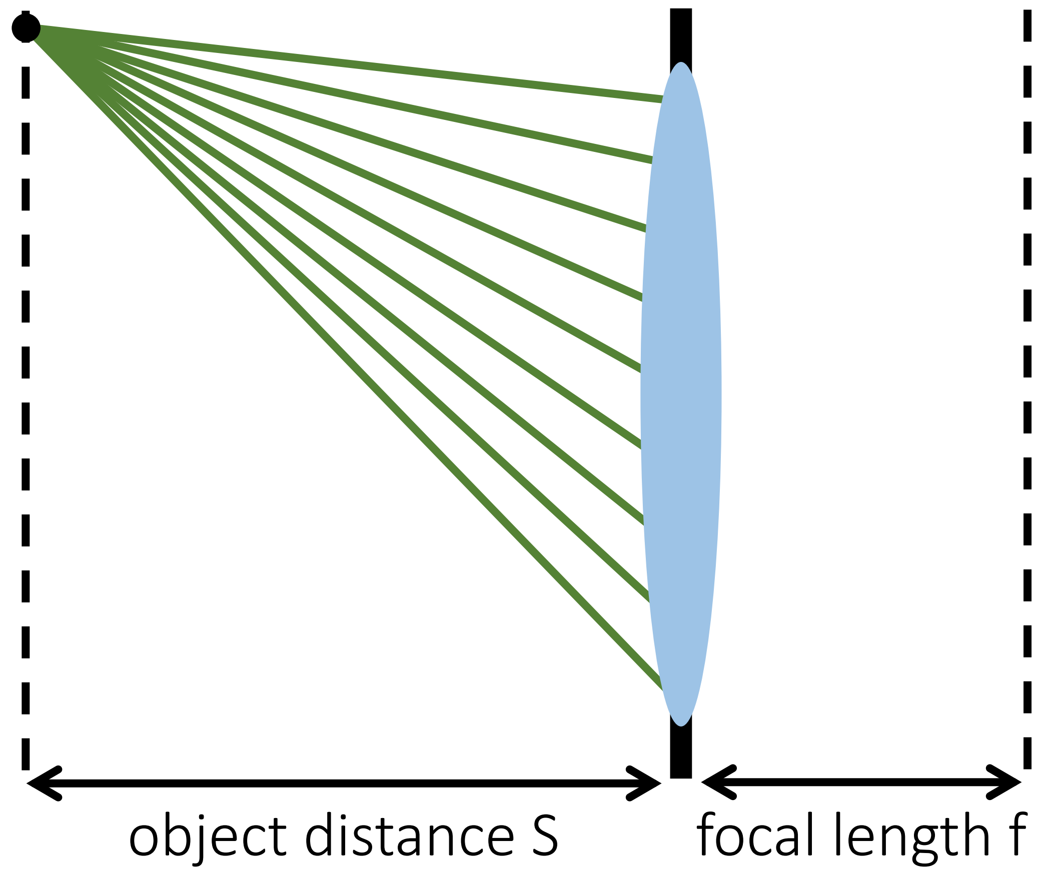

Tracing rays through a thin lens

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

1. Trace rays through lens center.

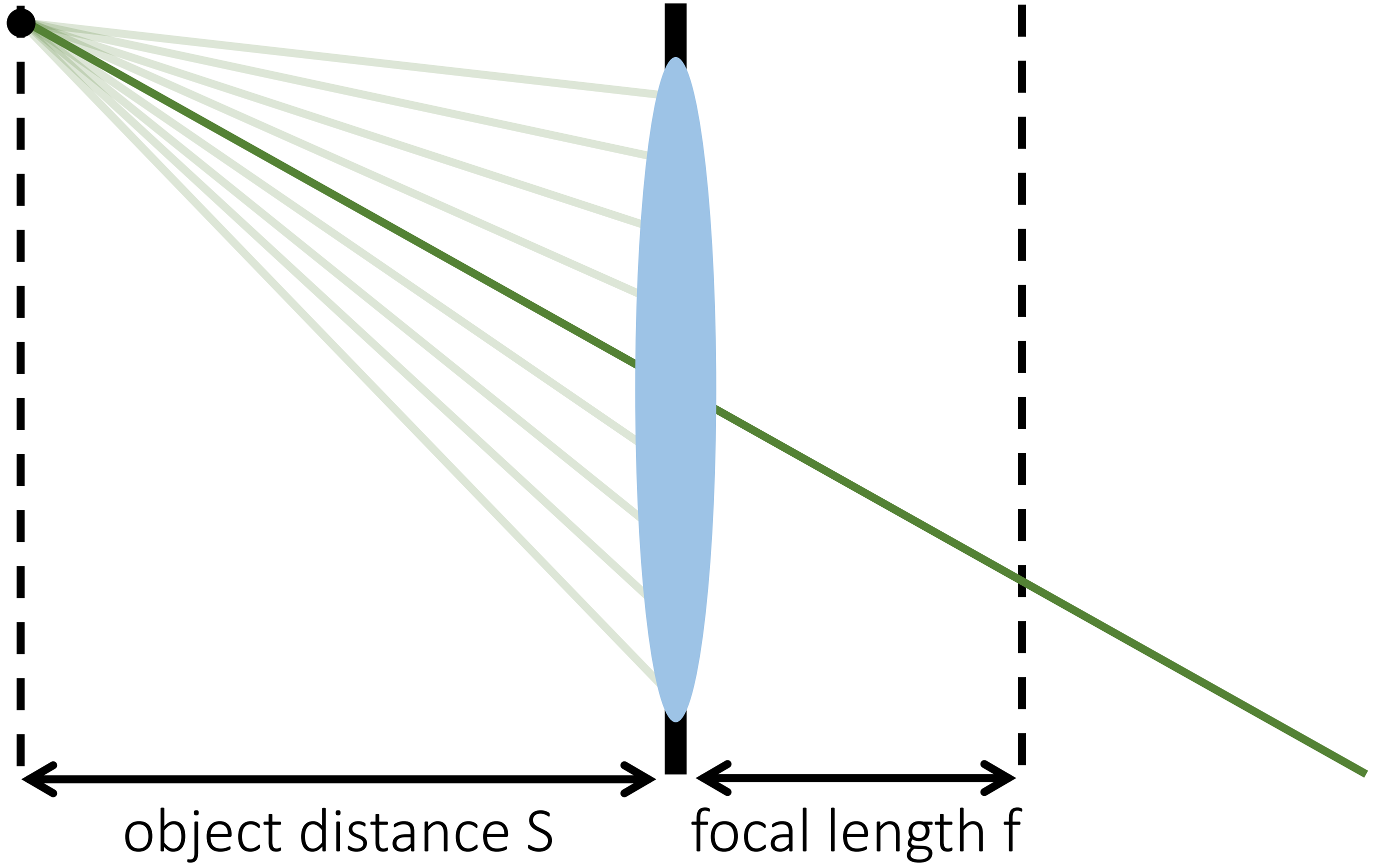

Tracing rays through a thin lens

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:

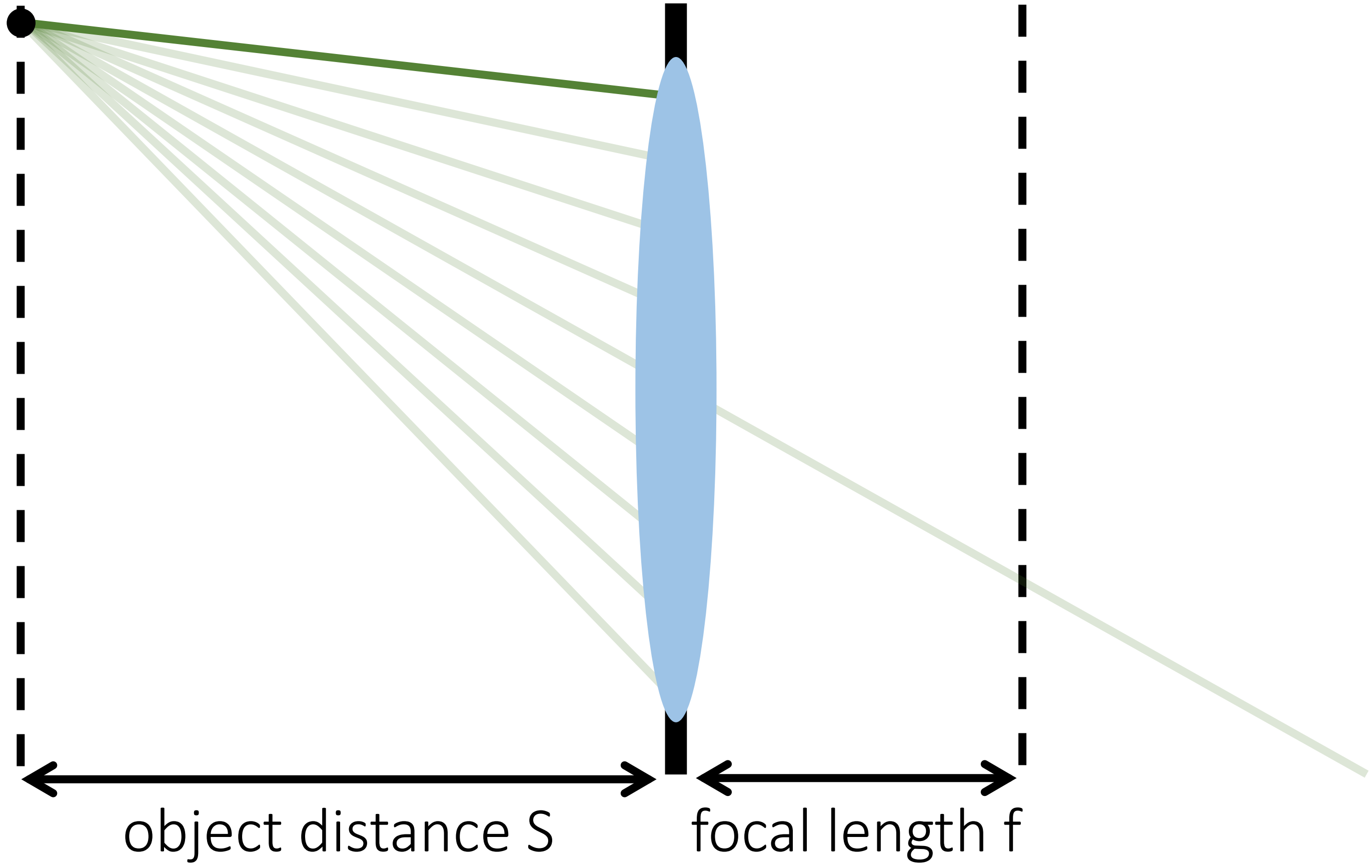

Tracing rays through a thin lens

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.

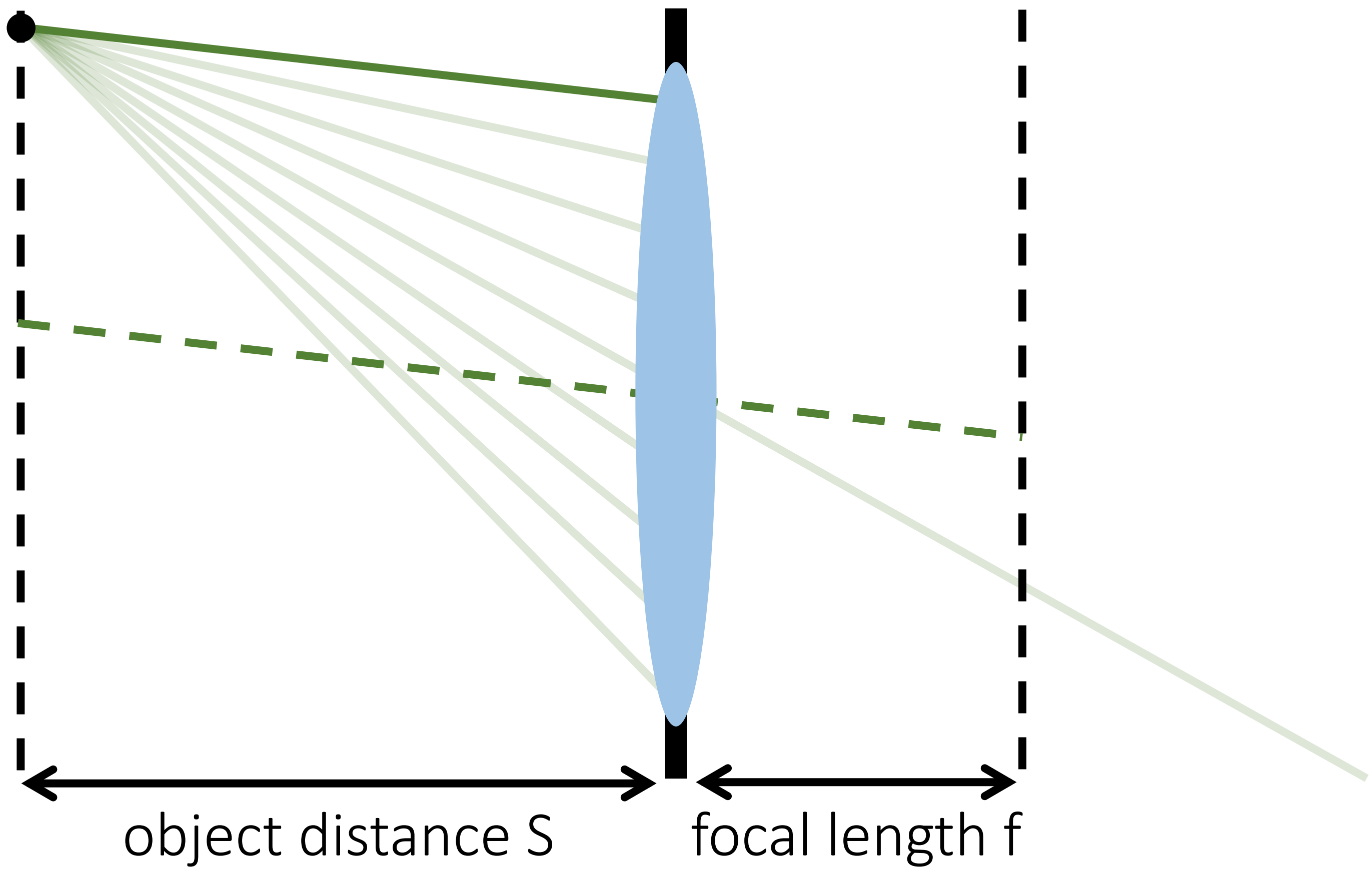

Tracing rays through a thin lens

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.

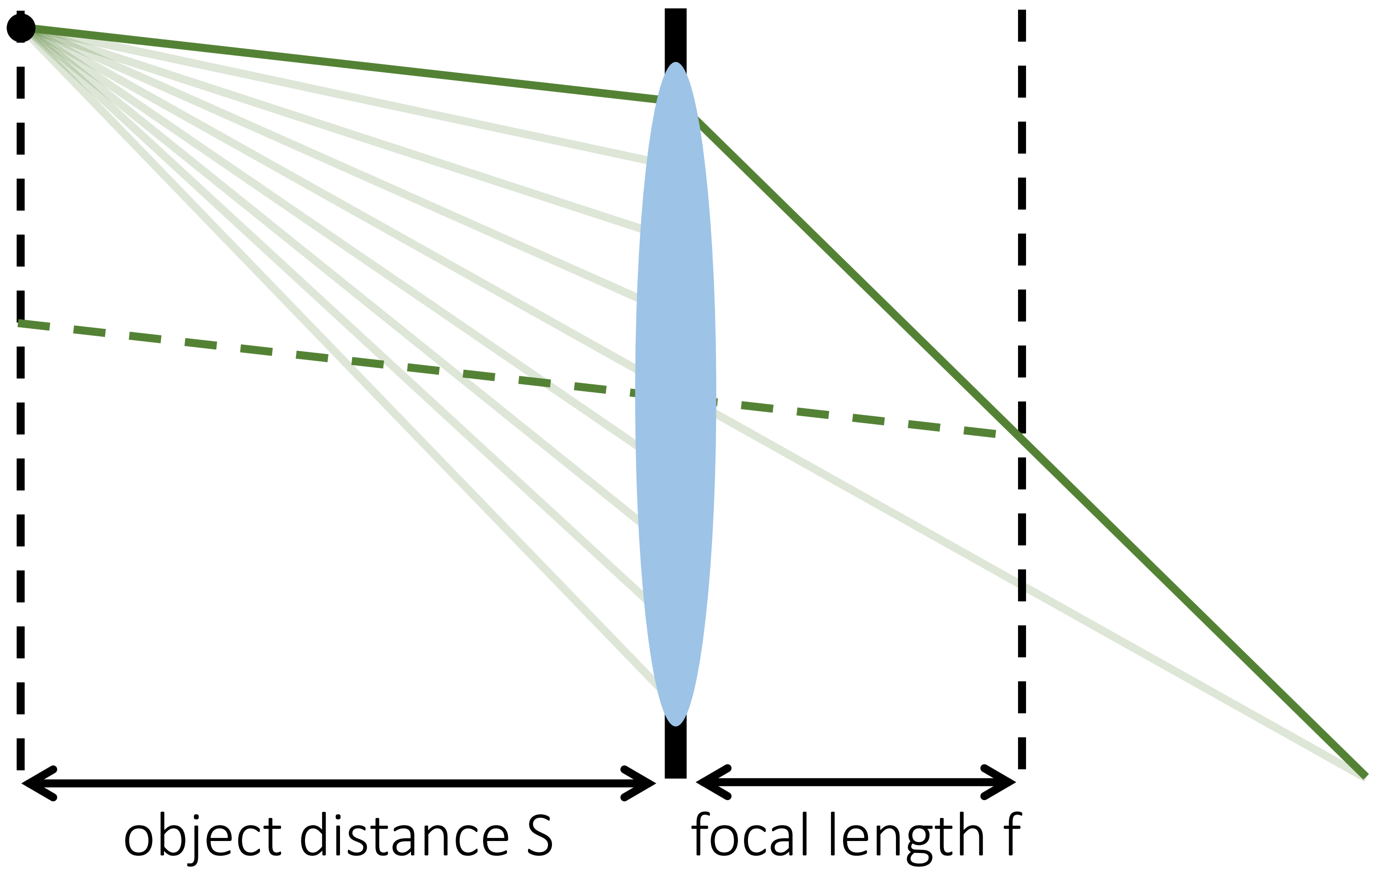

Tracing rays through a thin lens

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.

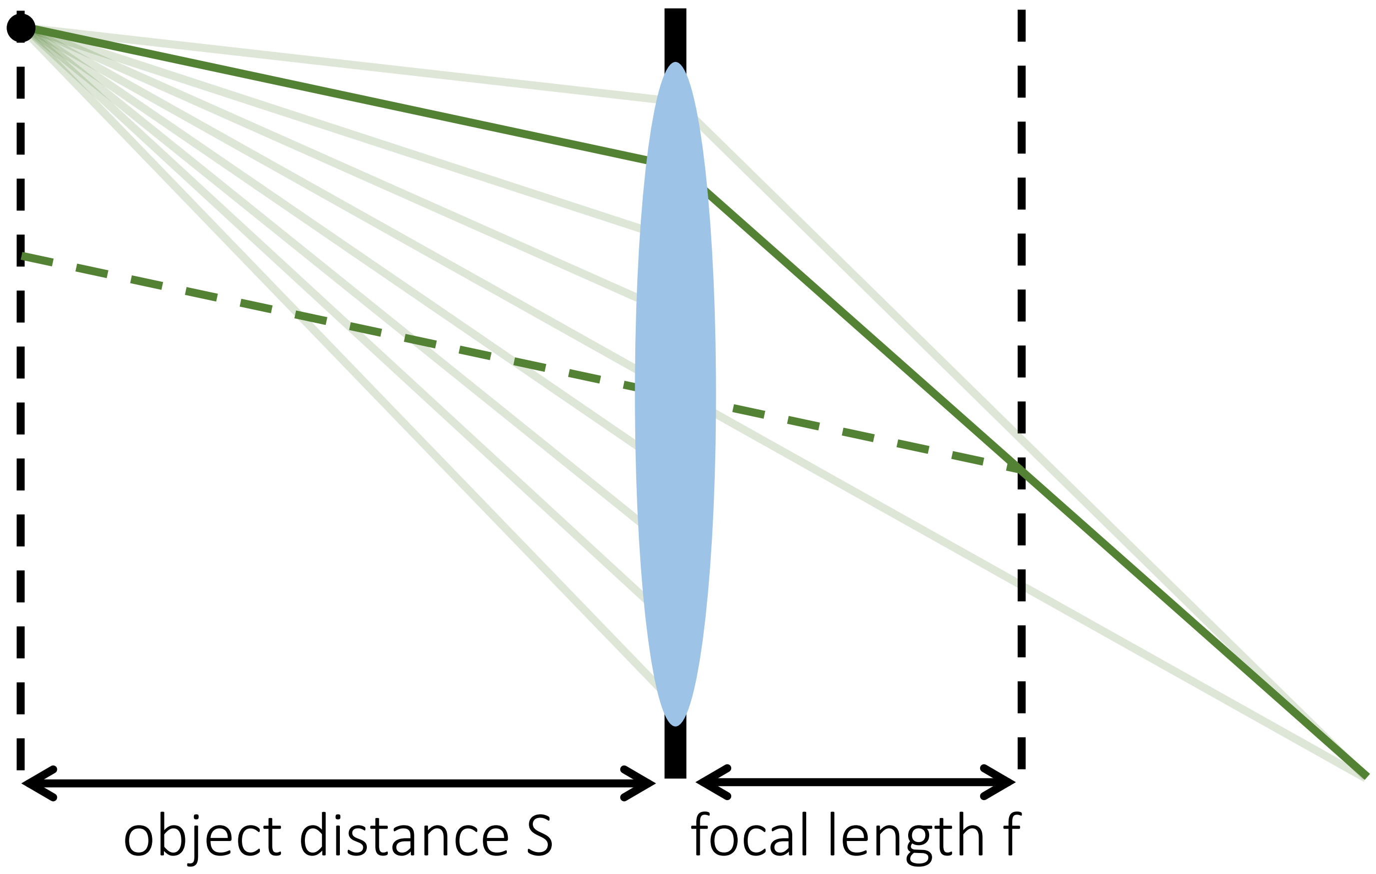

Tracing rays through a thin lens

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.

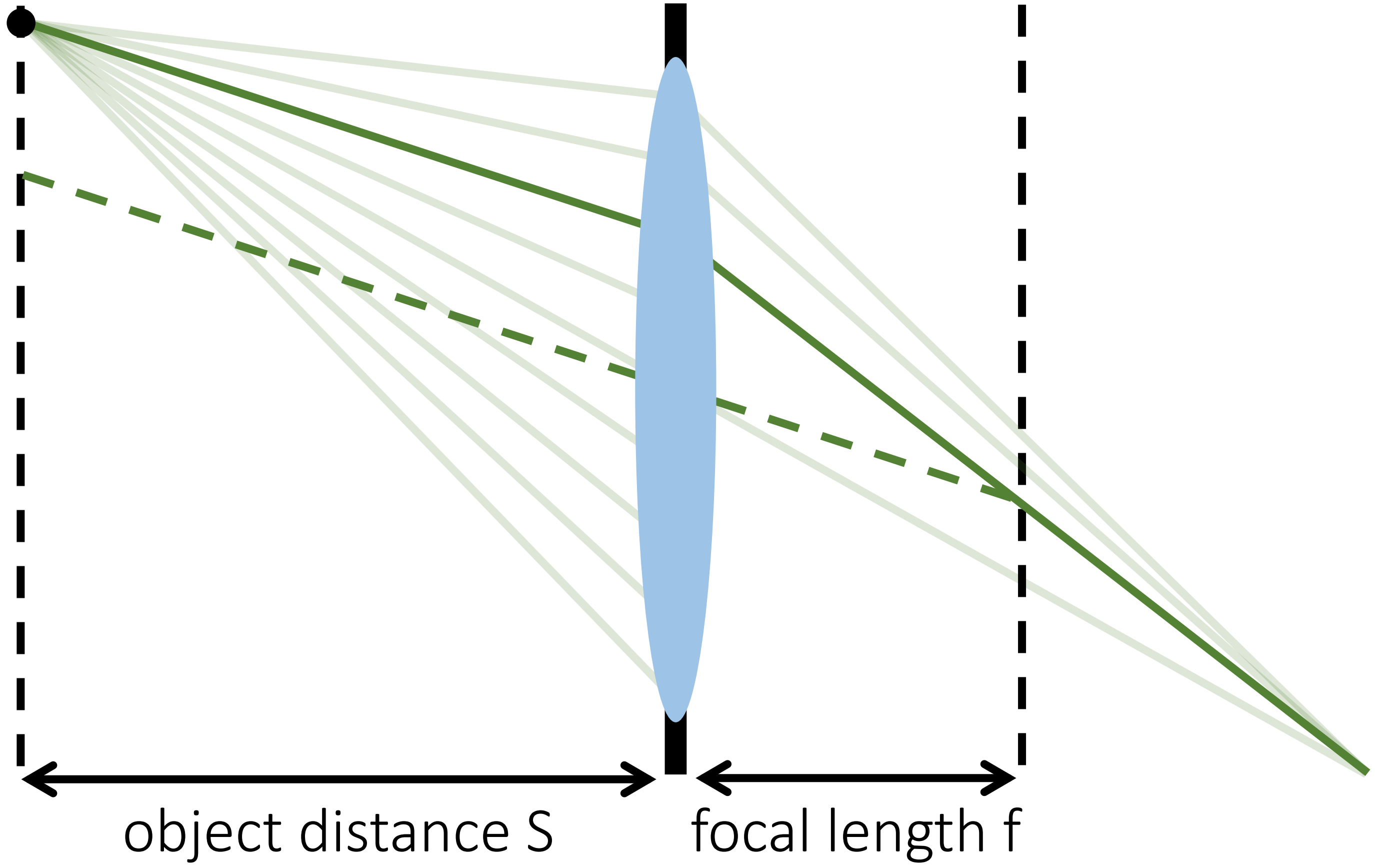

Tracing rays through a thin lens

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.

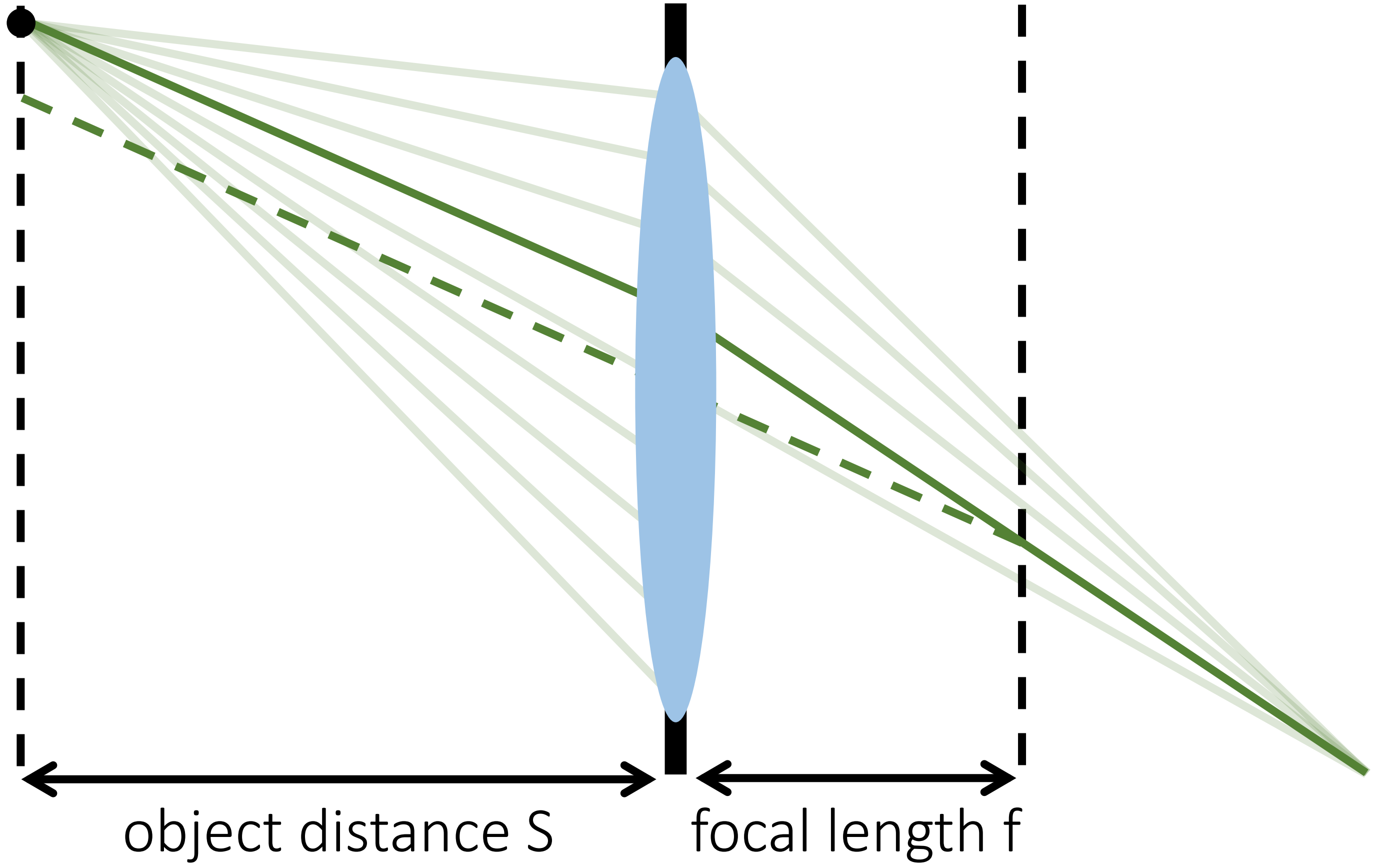

Tracing rays through a thin lens

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.

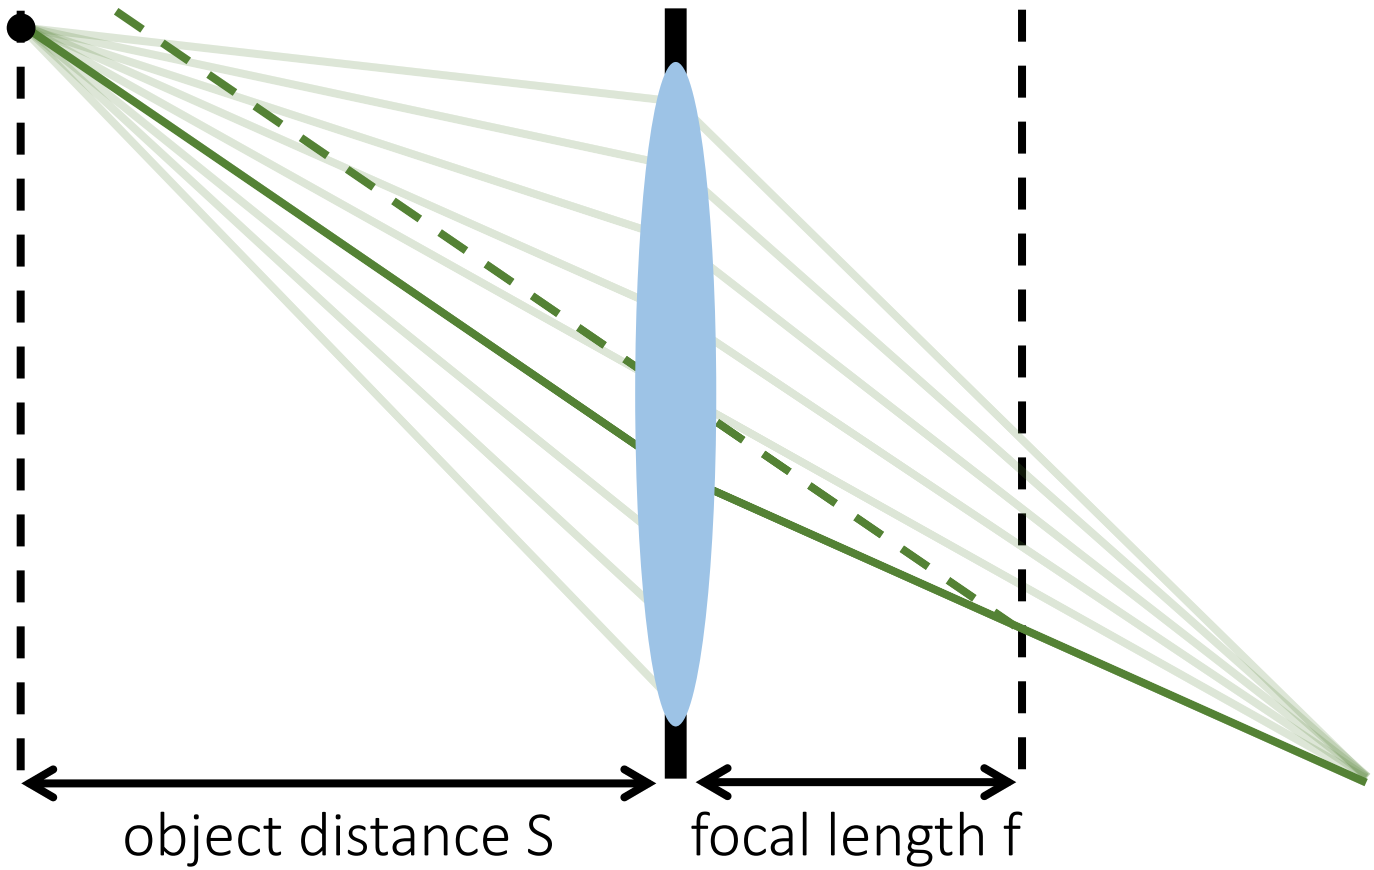

Tracing rays through a thin lens

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.

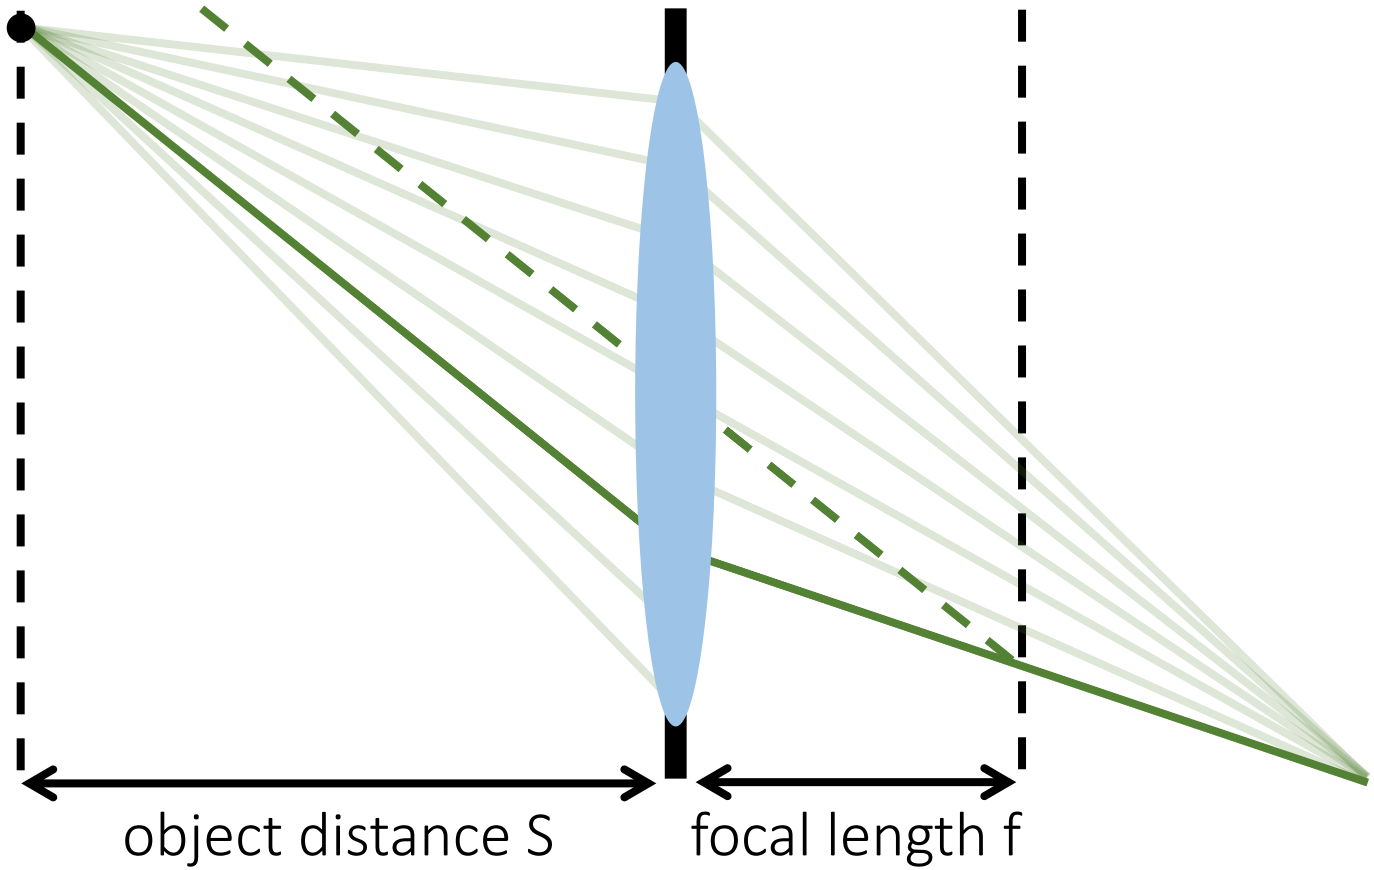

Tracing rays through a thin lens

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.

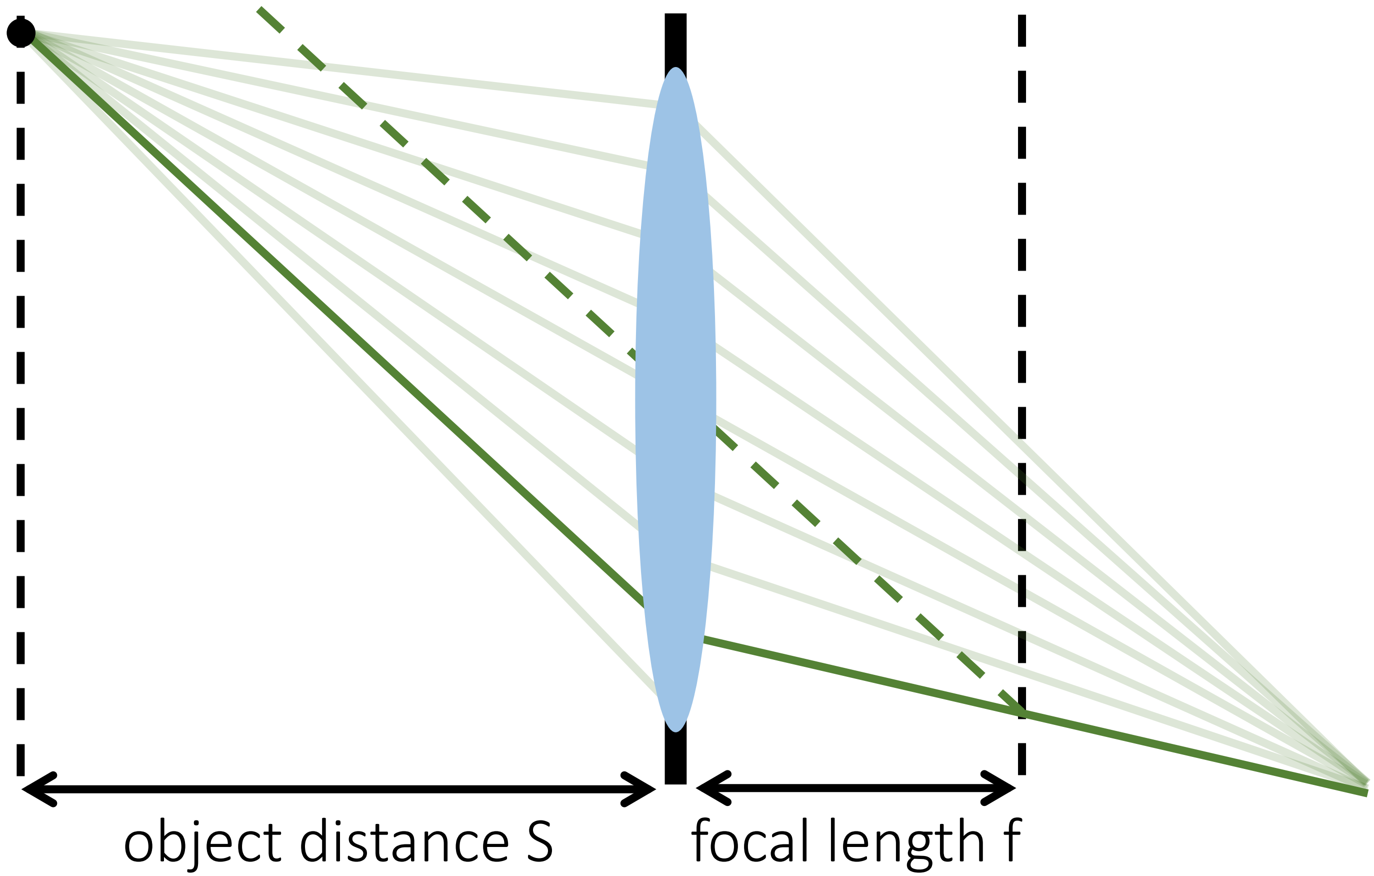

Tracing rays through a thin lens

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.

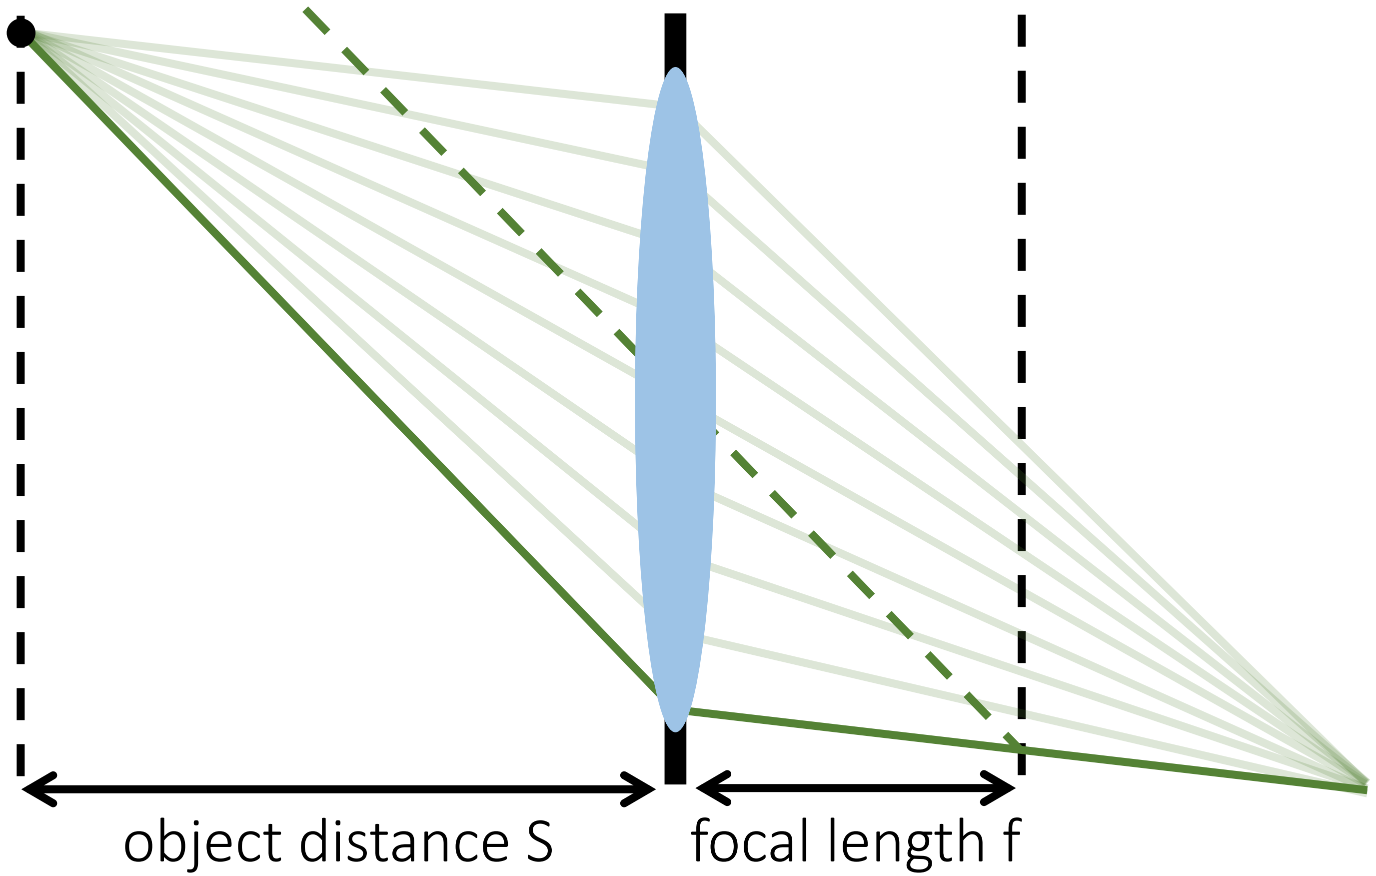

Tracing rays through a thin lens

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.

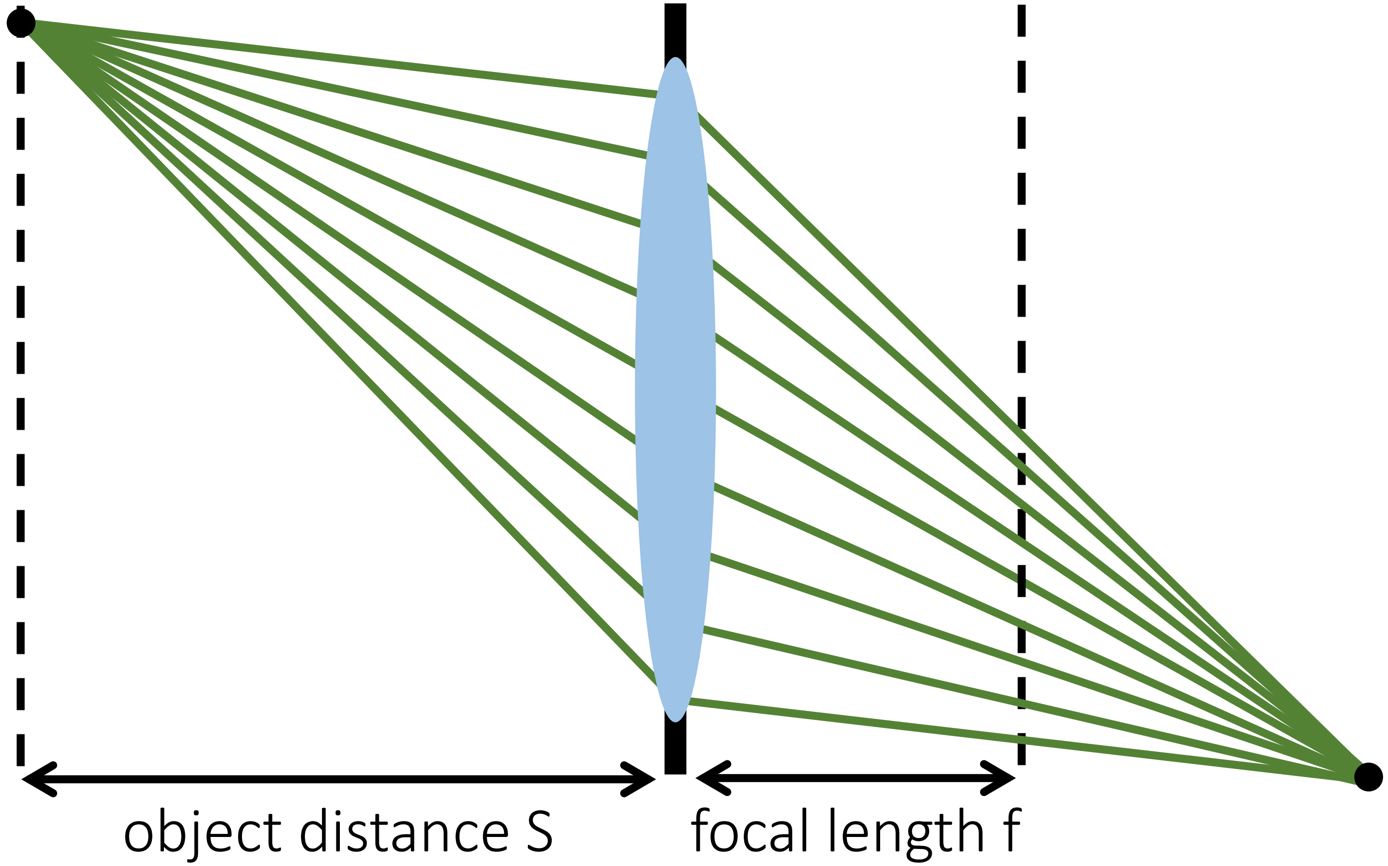

Focusing property:

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.

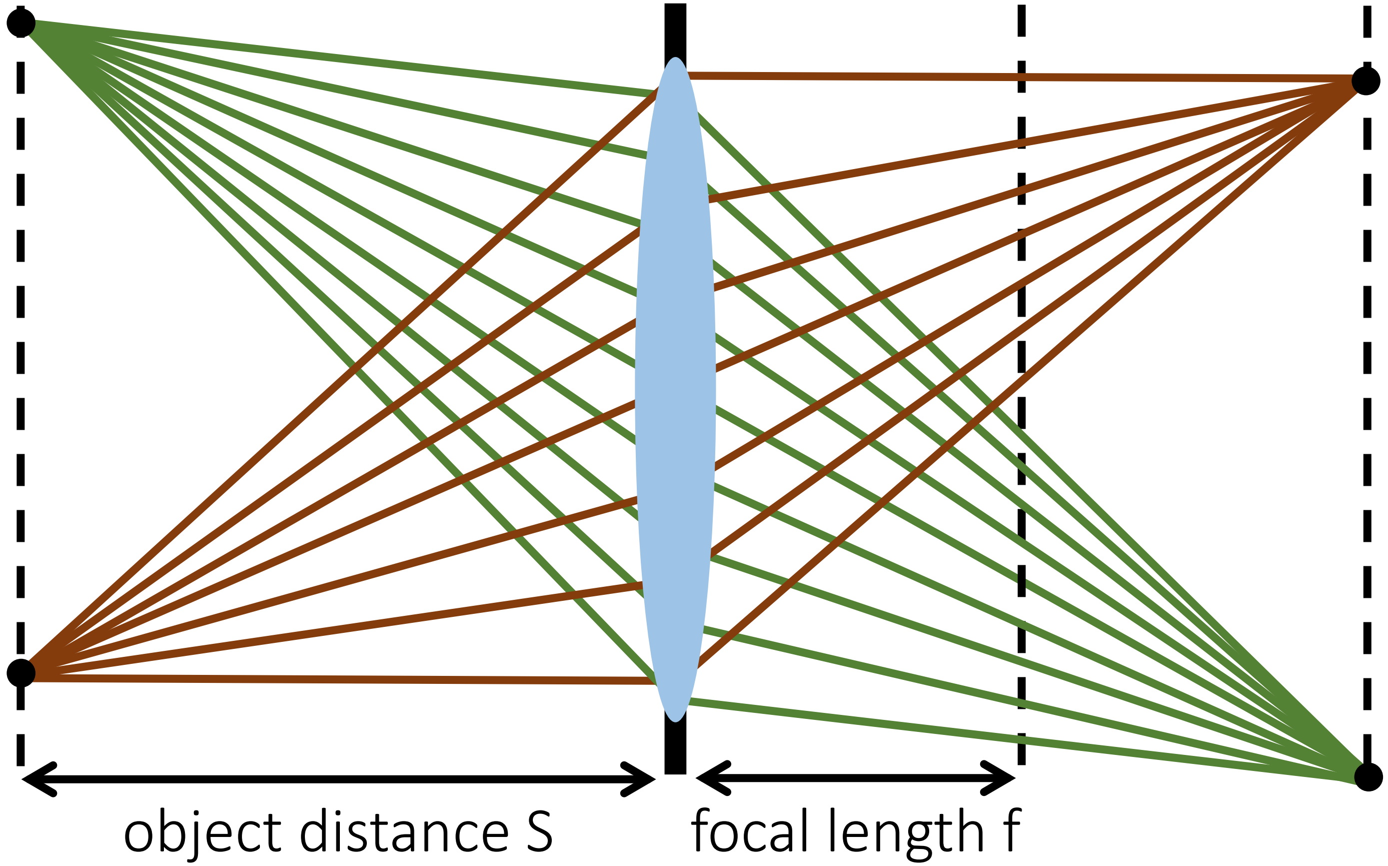

Focusing property:

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.

Gaussian lens formula

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

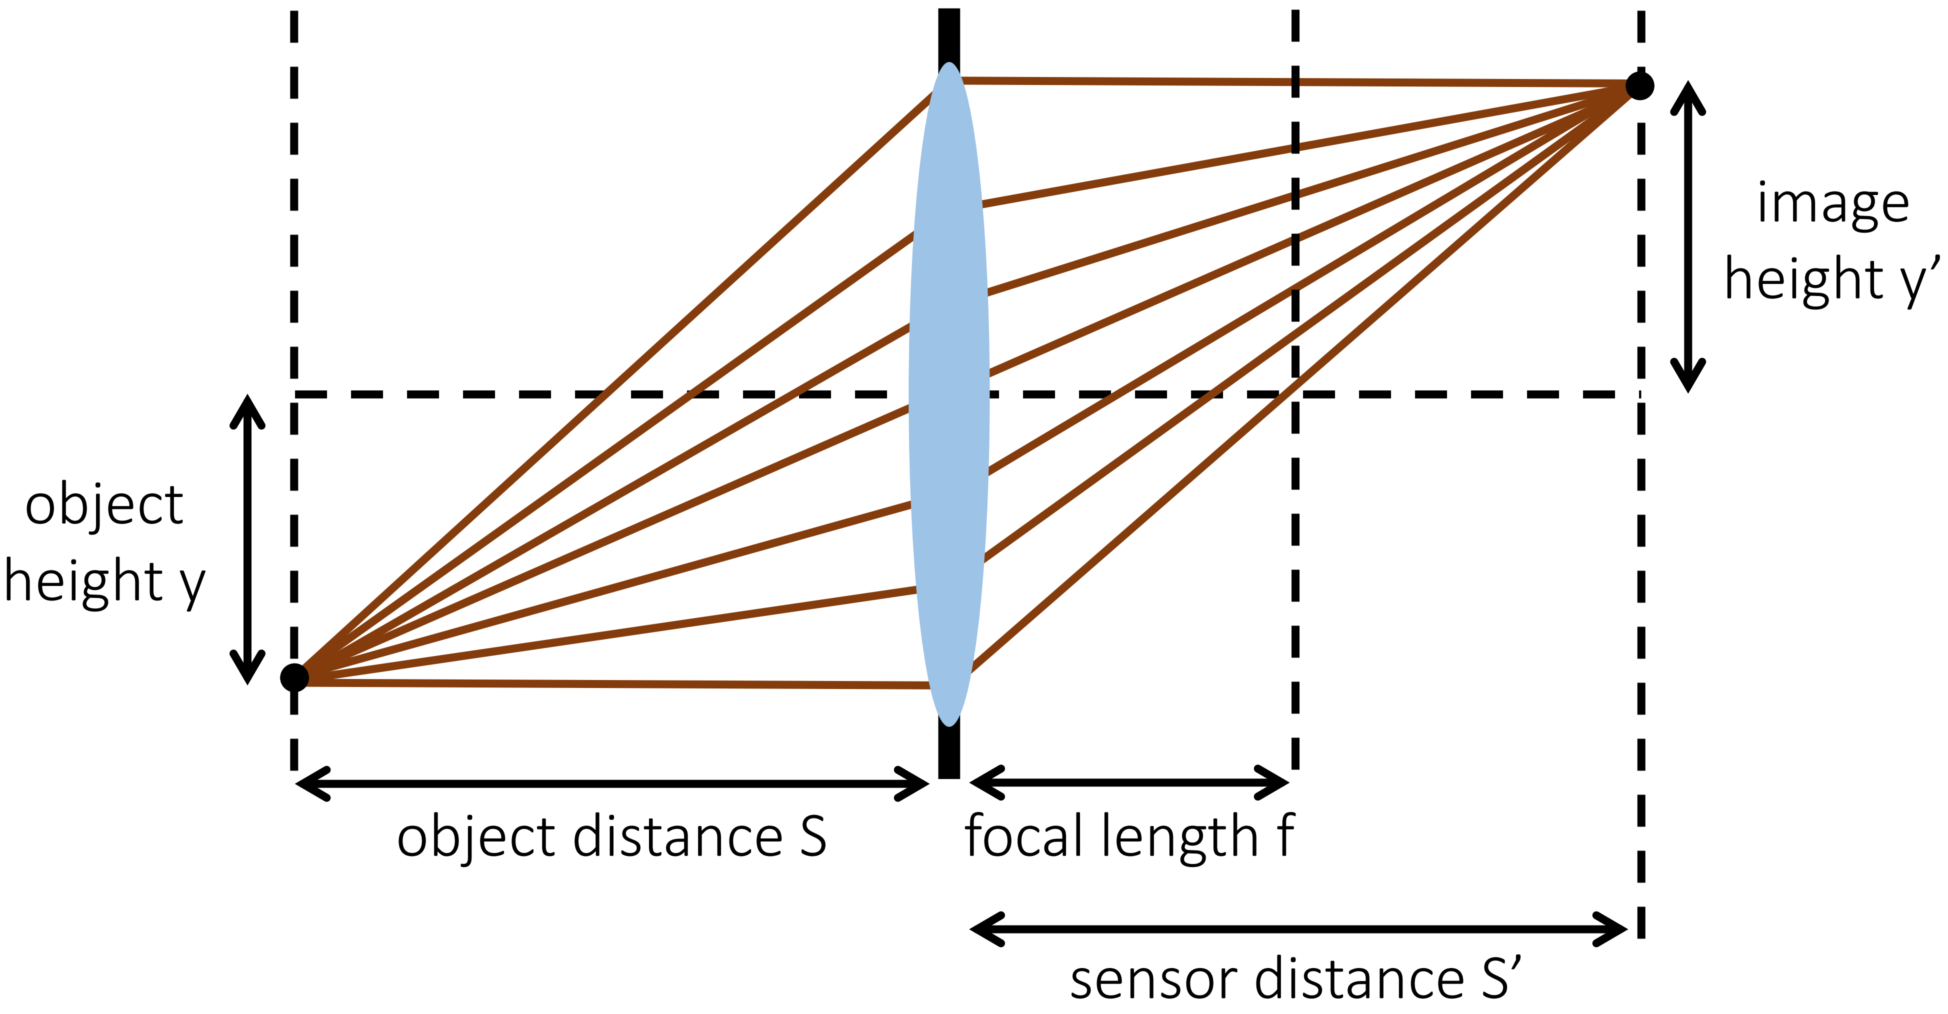

Gaussian lens formula

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

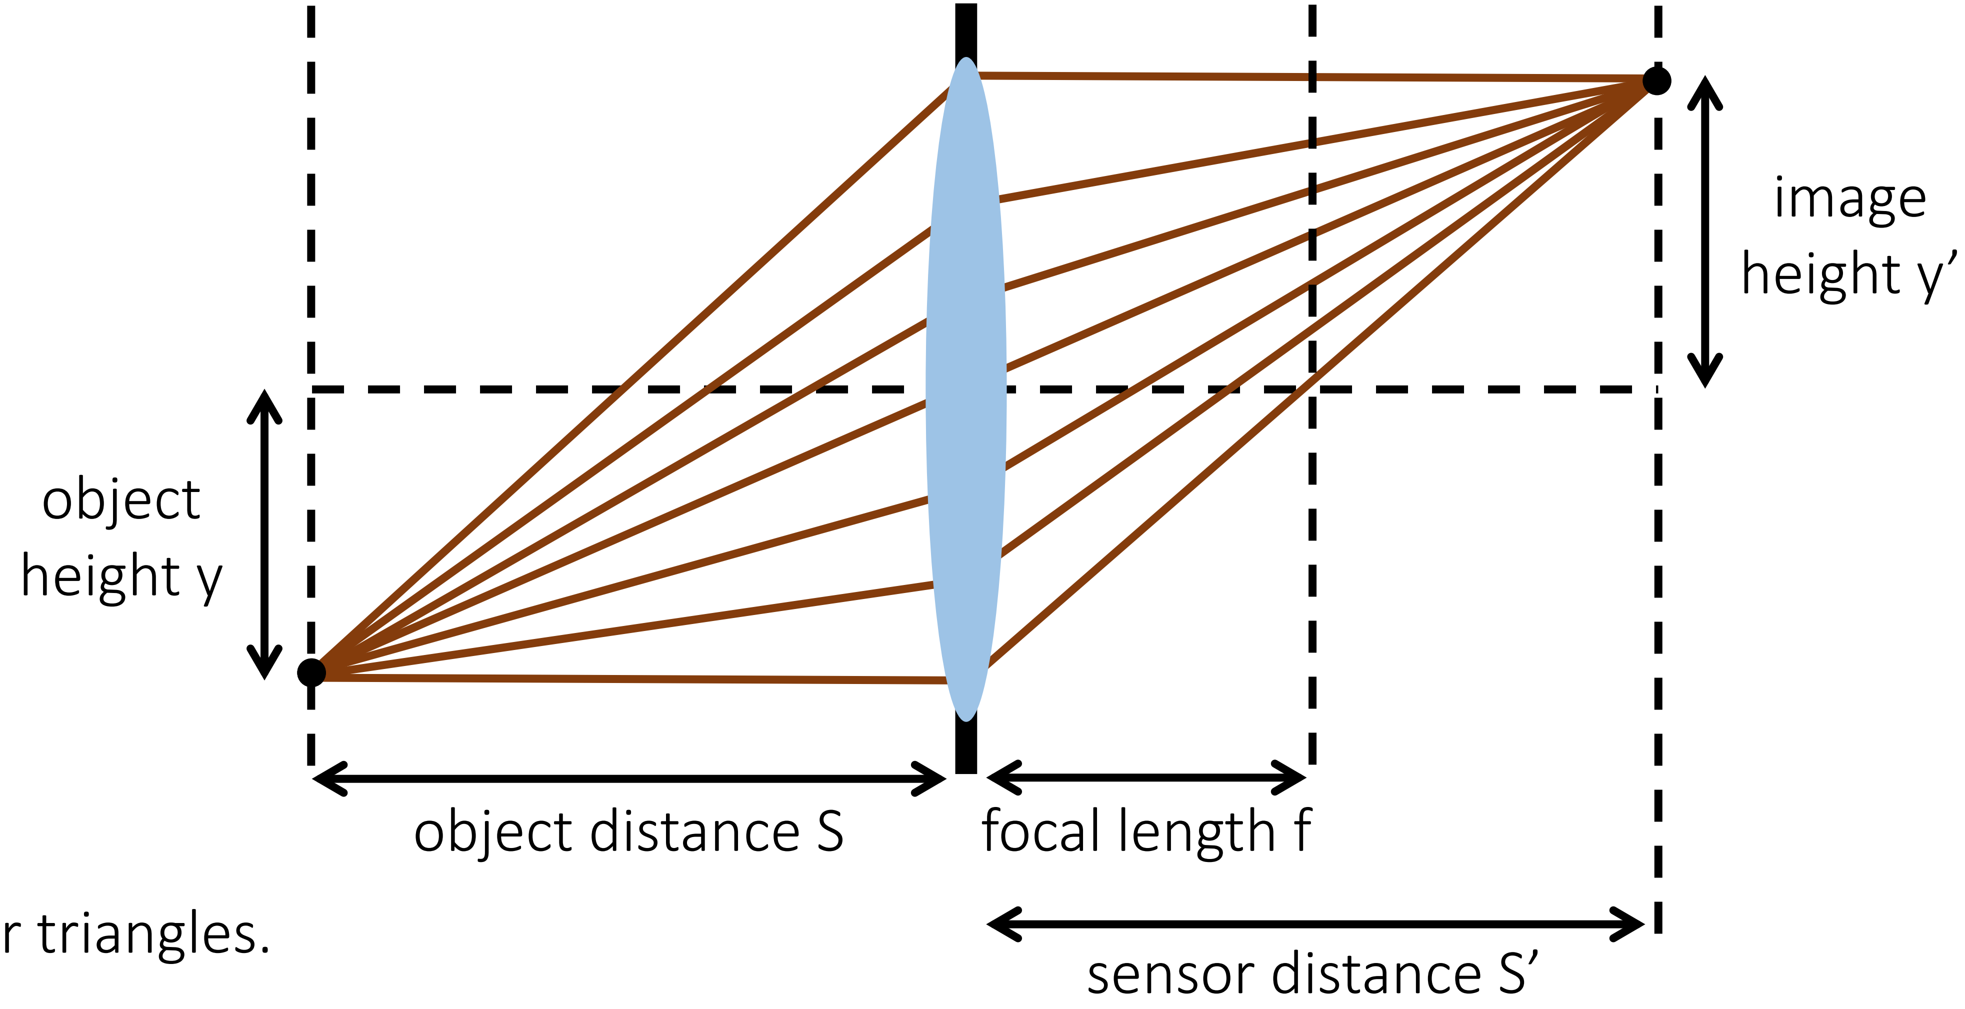

Use similar triangles.

Gaussian lens formula

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

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

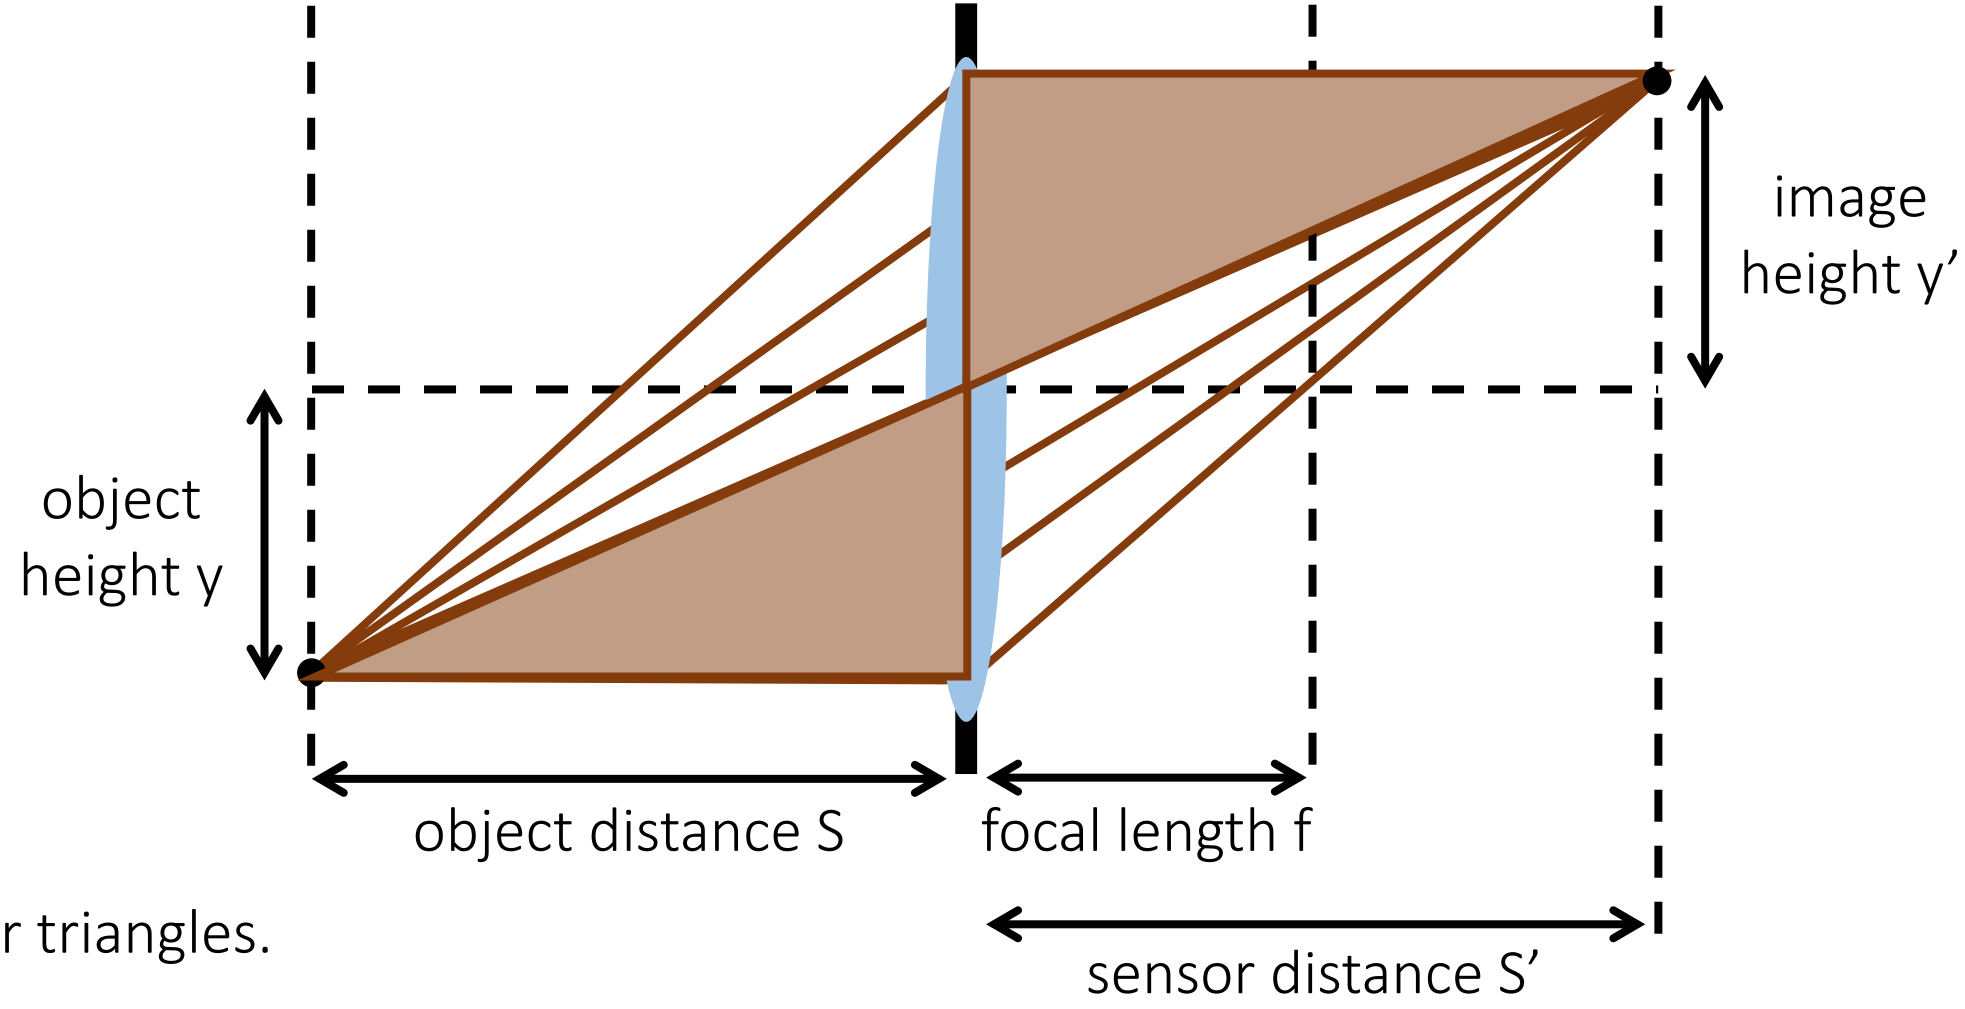

Use similar triangles.

Gaussian lens formula

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

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

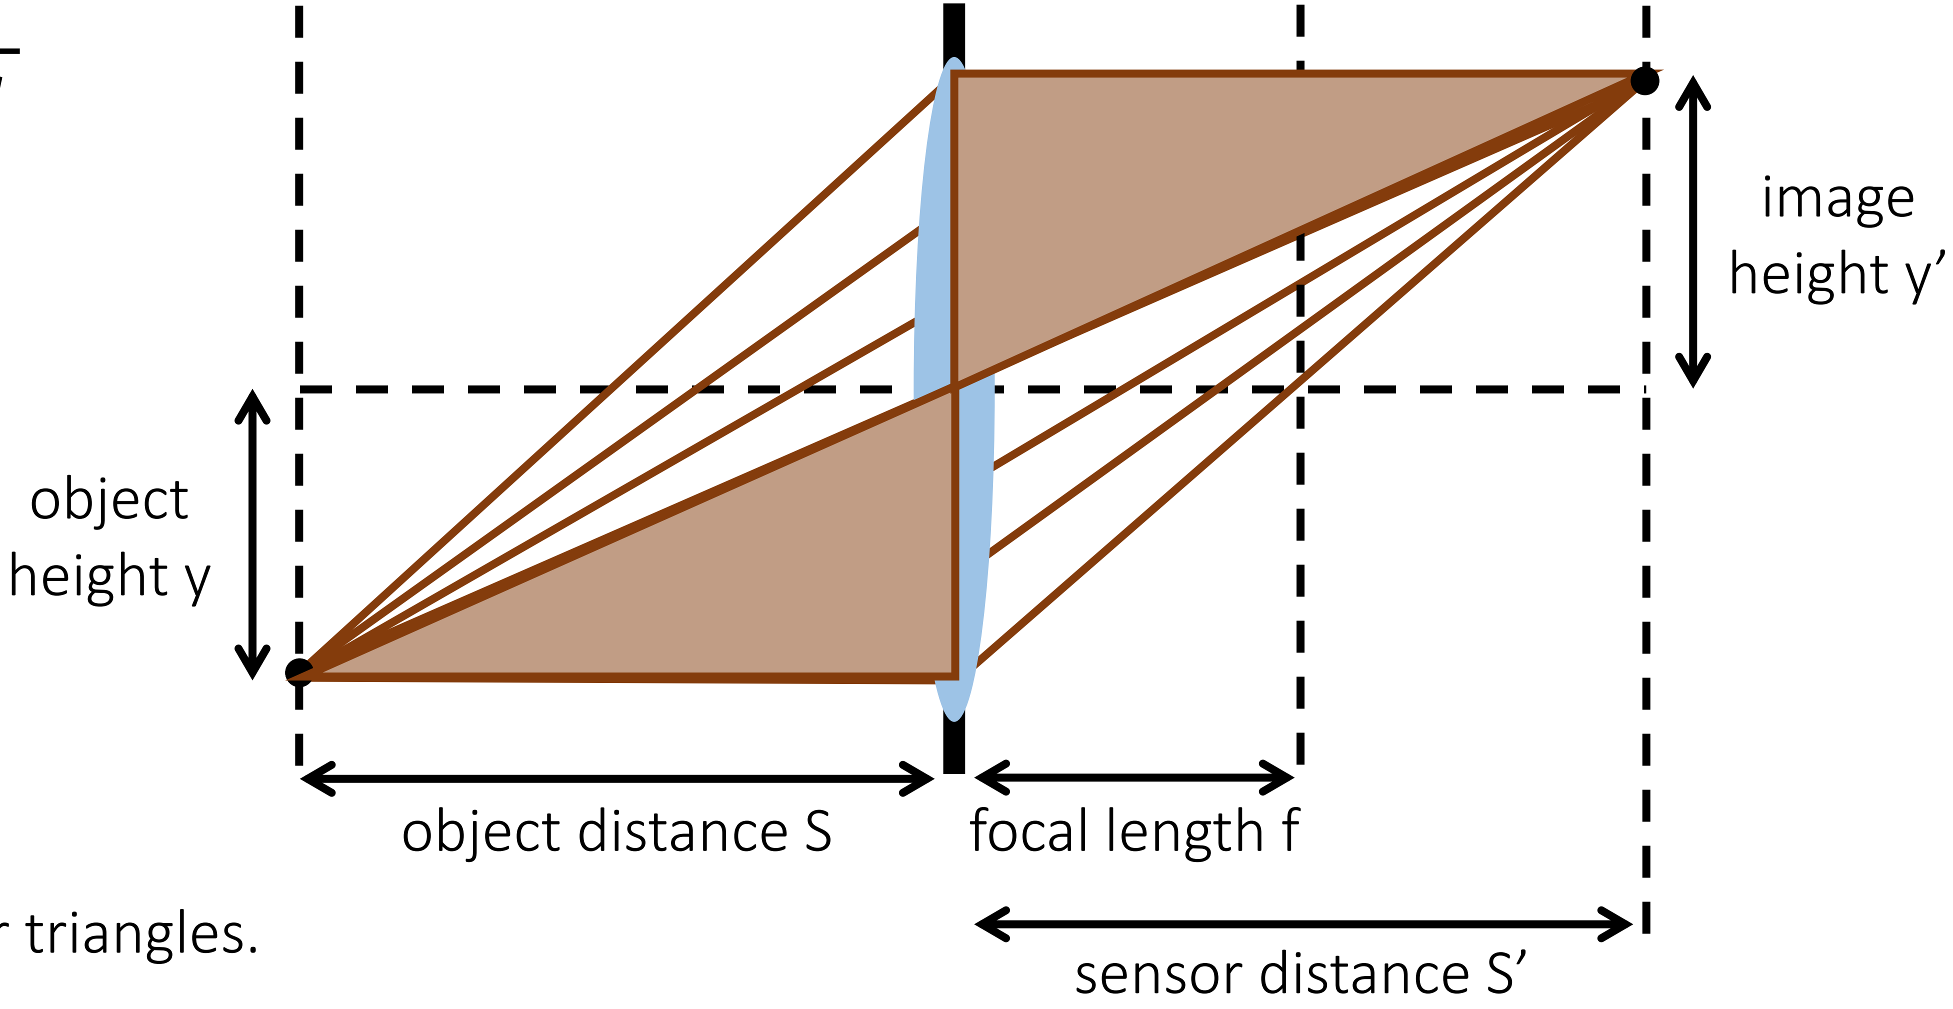

Use similar triangles.

Gaussian lens formula

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

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

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

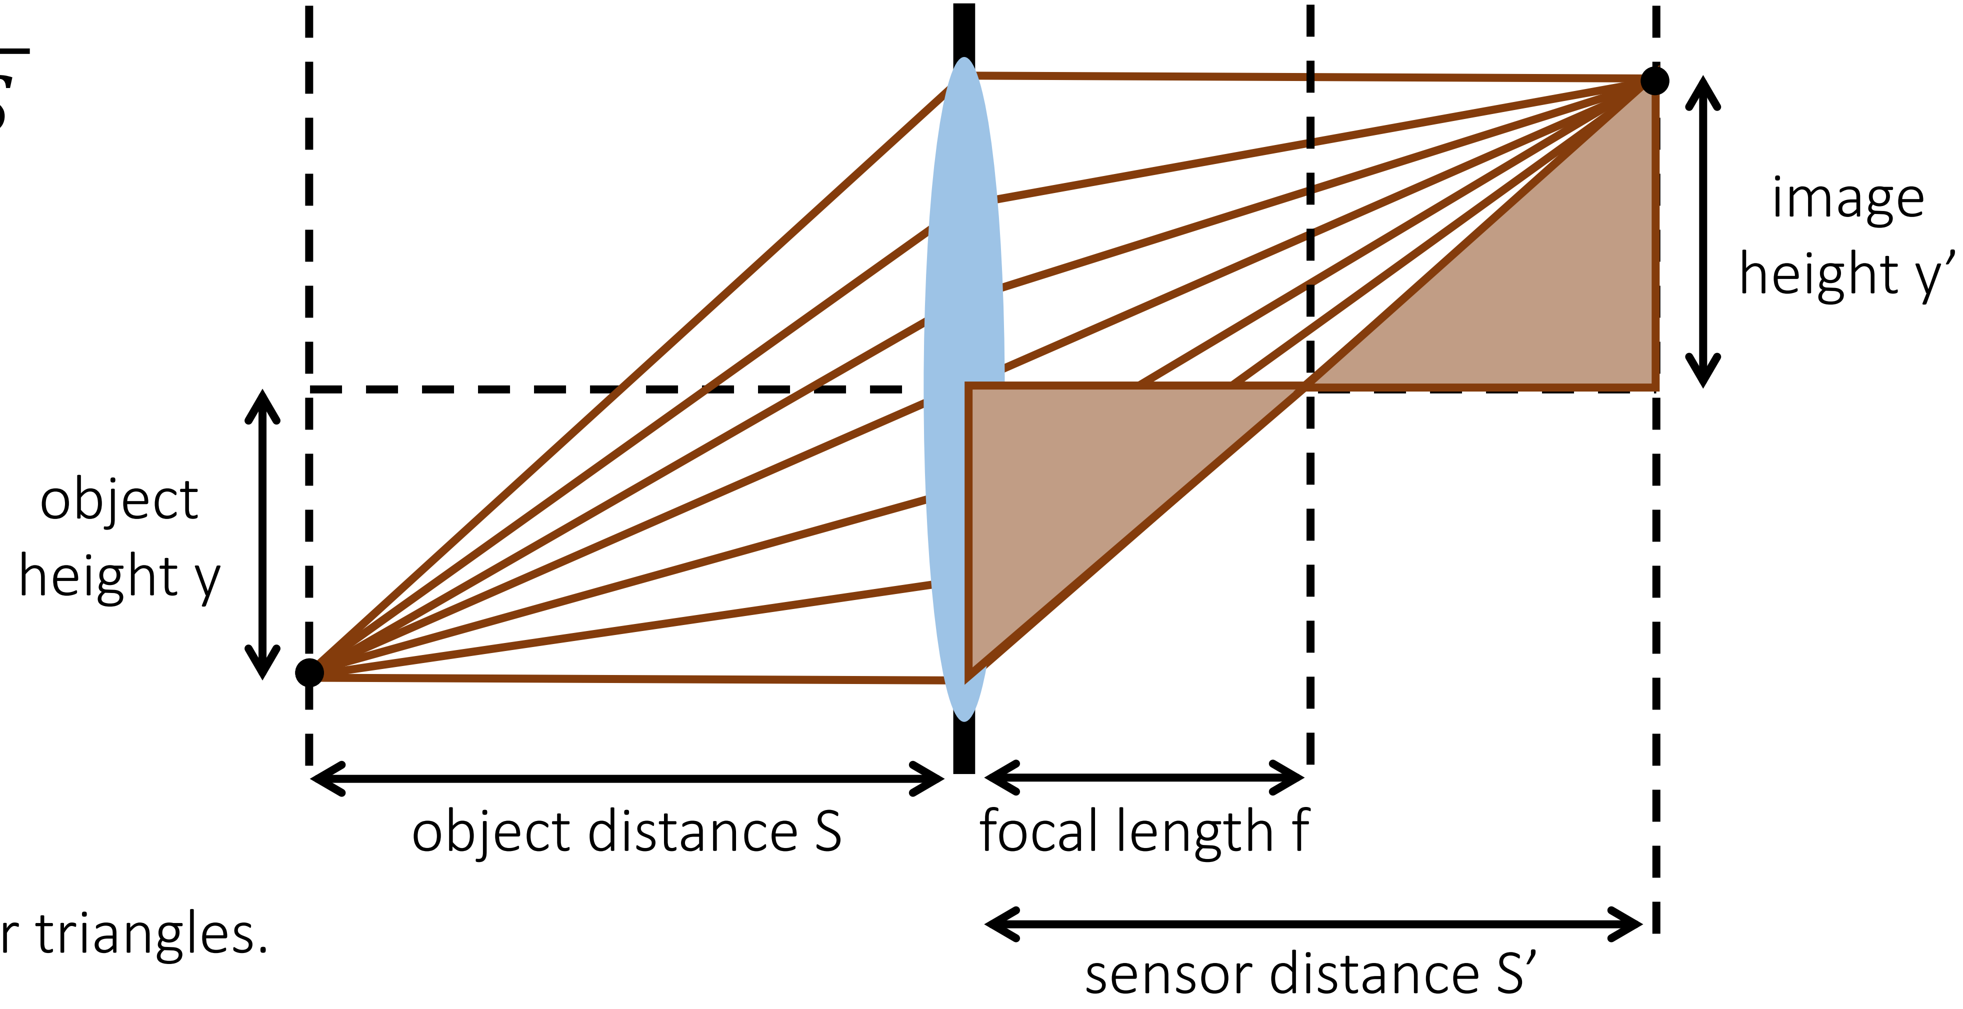

Use similar triangles.

Gaussian lens formula

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

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

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

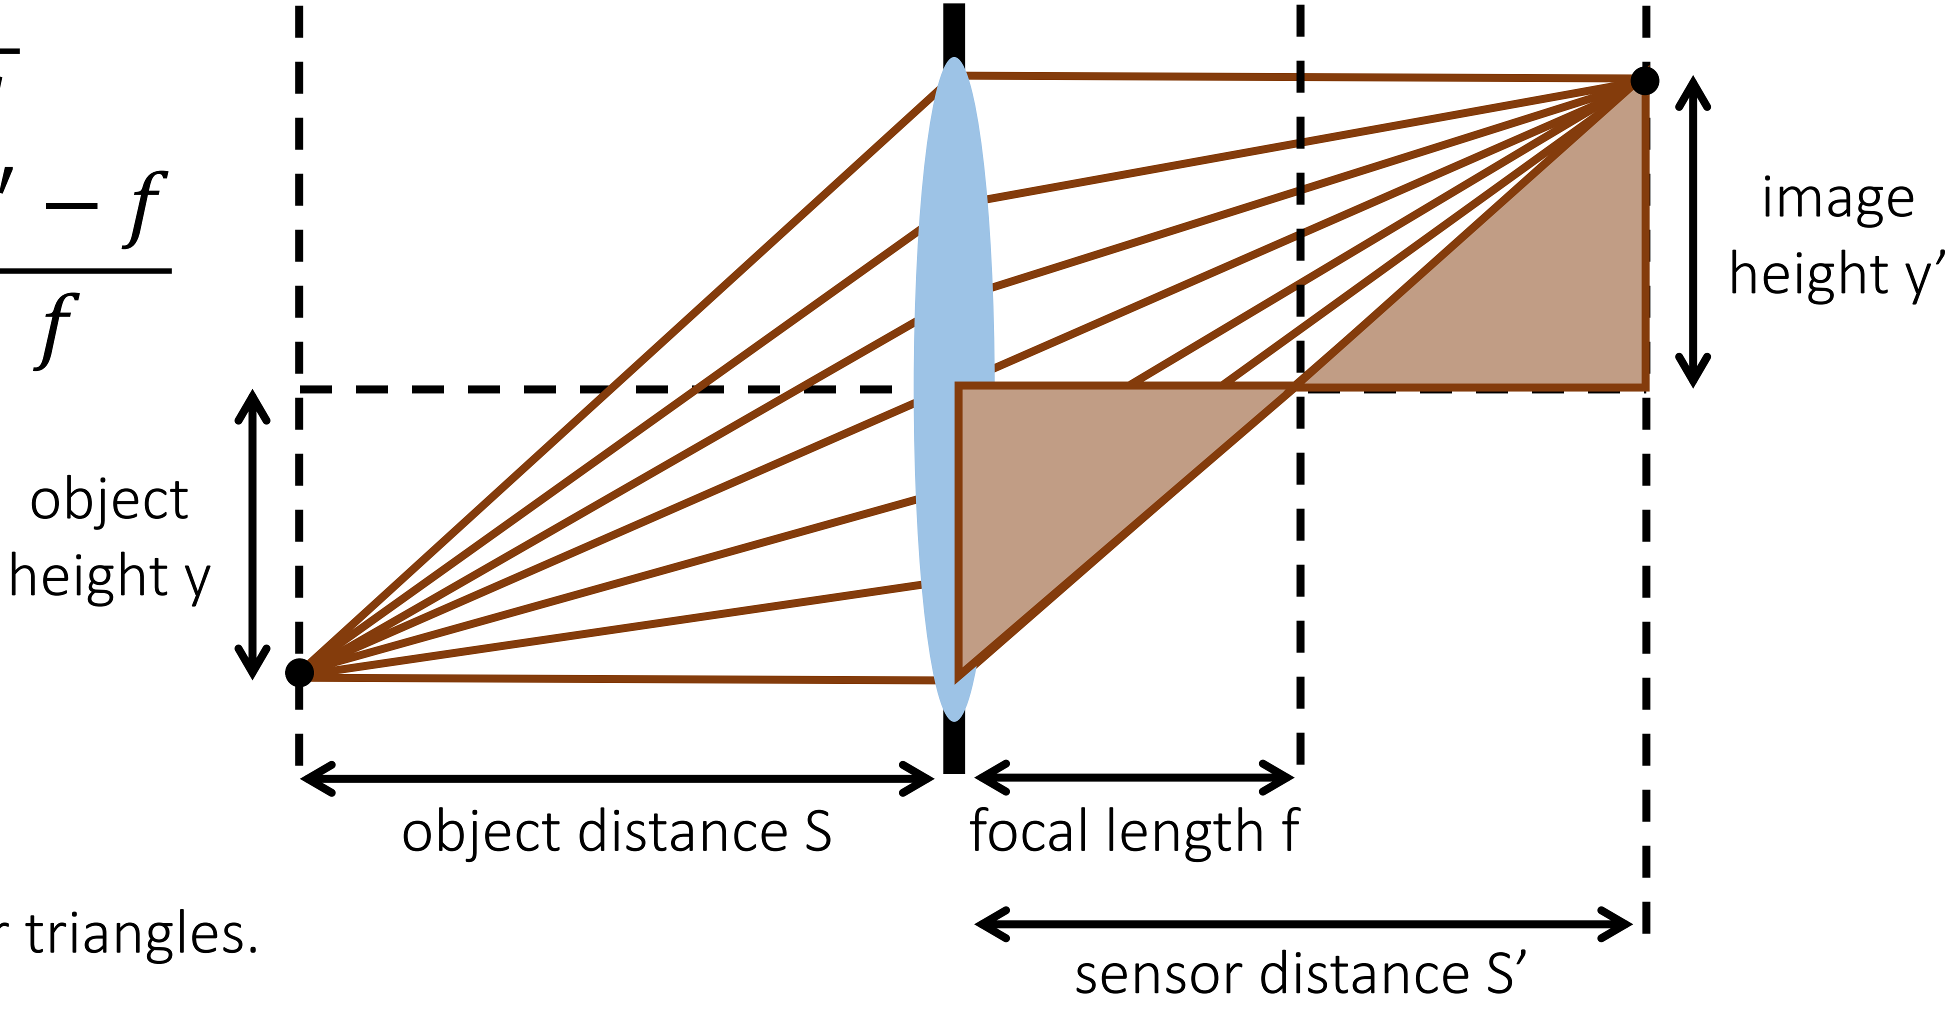

Use similar triangles.

Gaussian lens formula

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

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

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

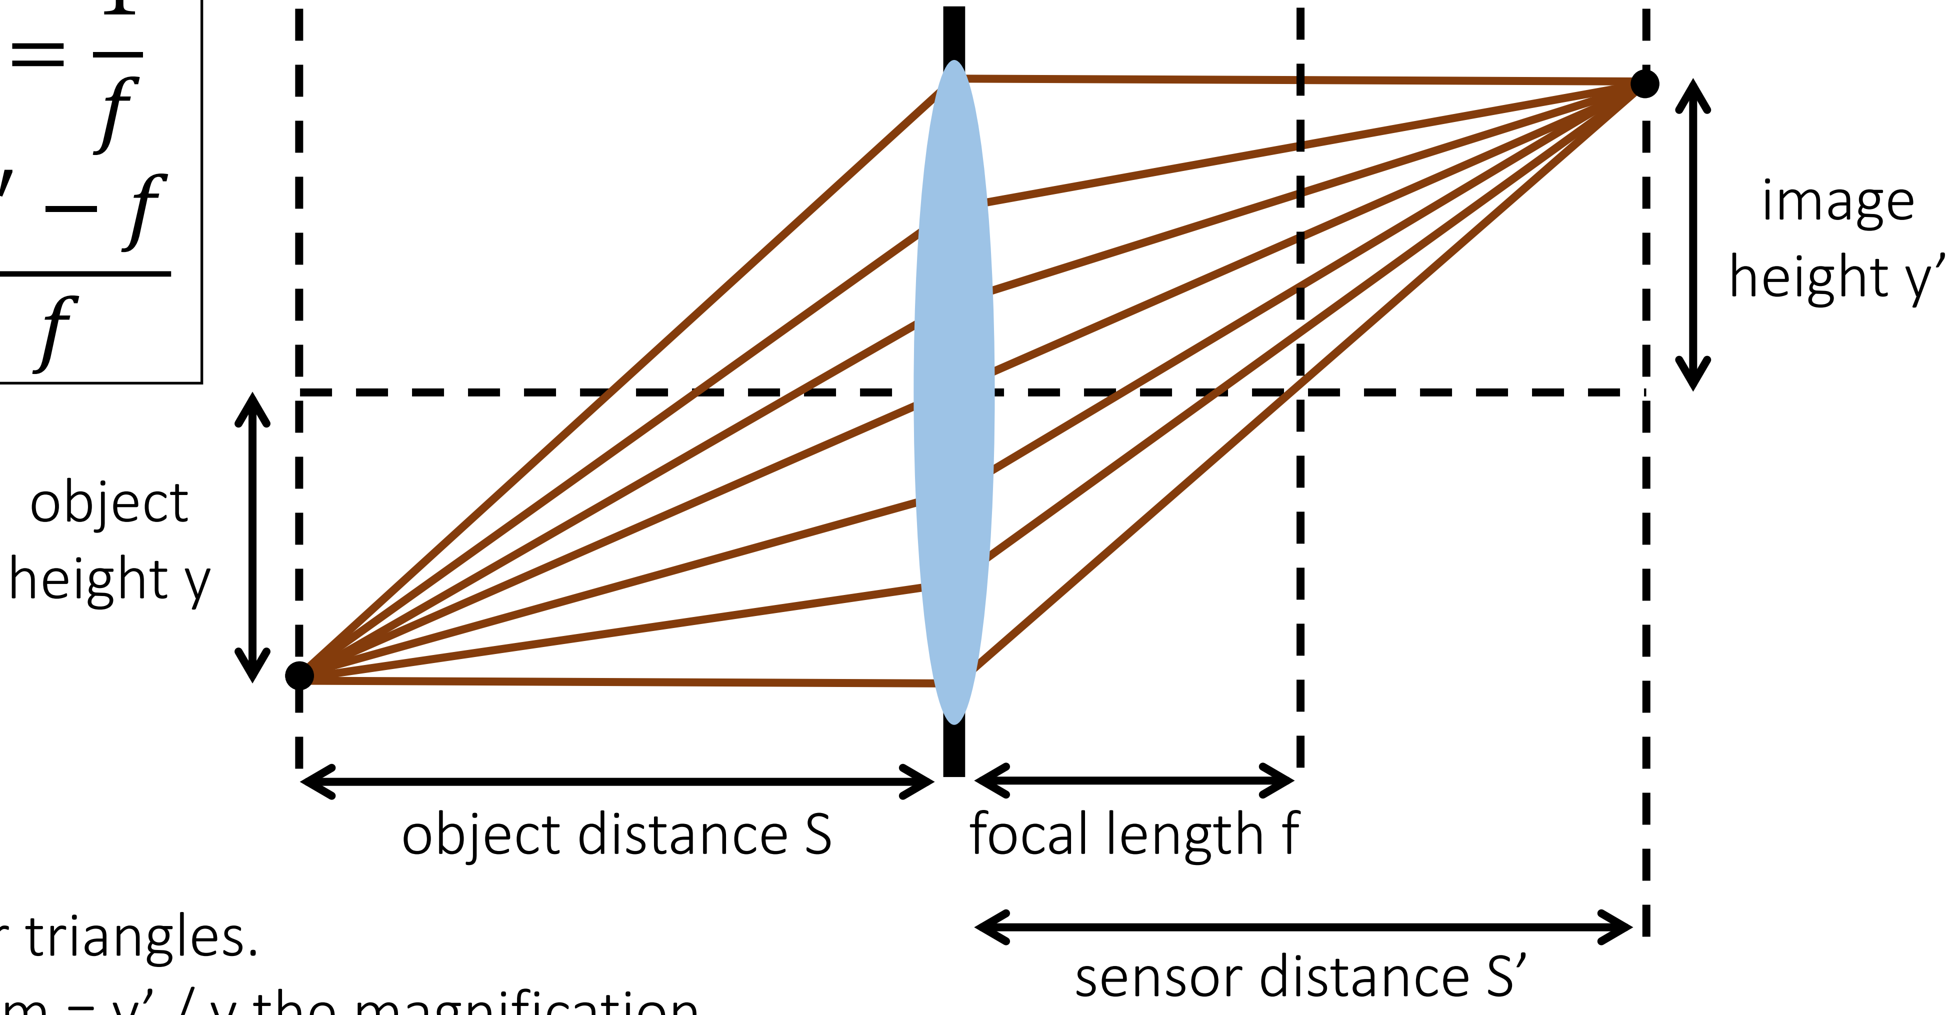

Use similar triangles.

- We call m = y’ / y the magnification.

Gaussian lens formula

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

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

$$m = \frac{f}{S - f}$$

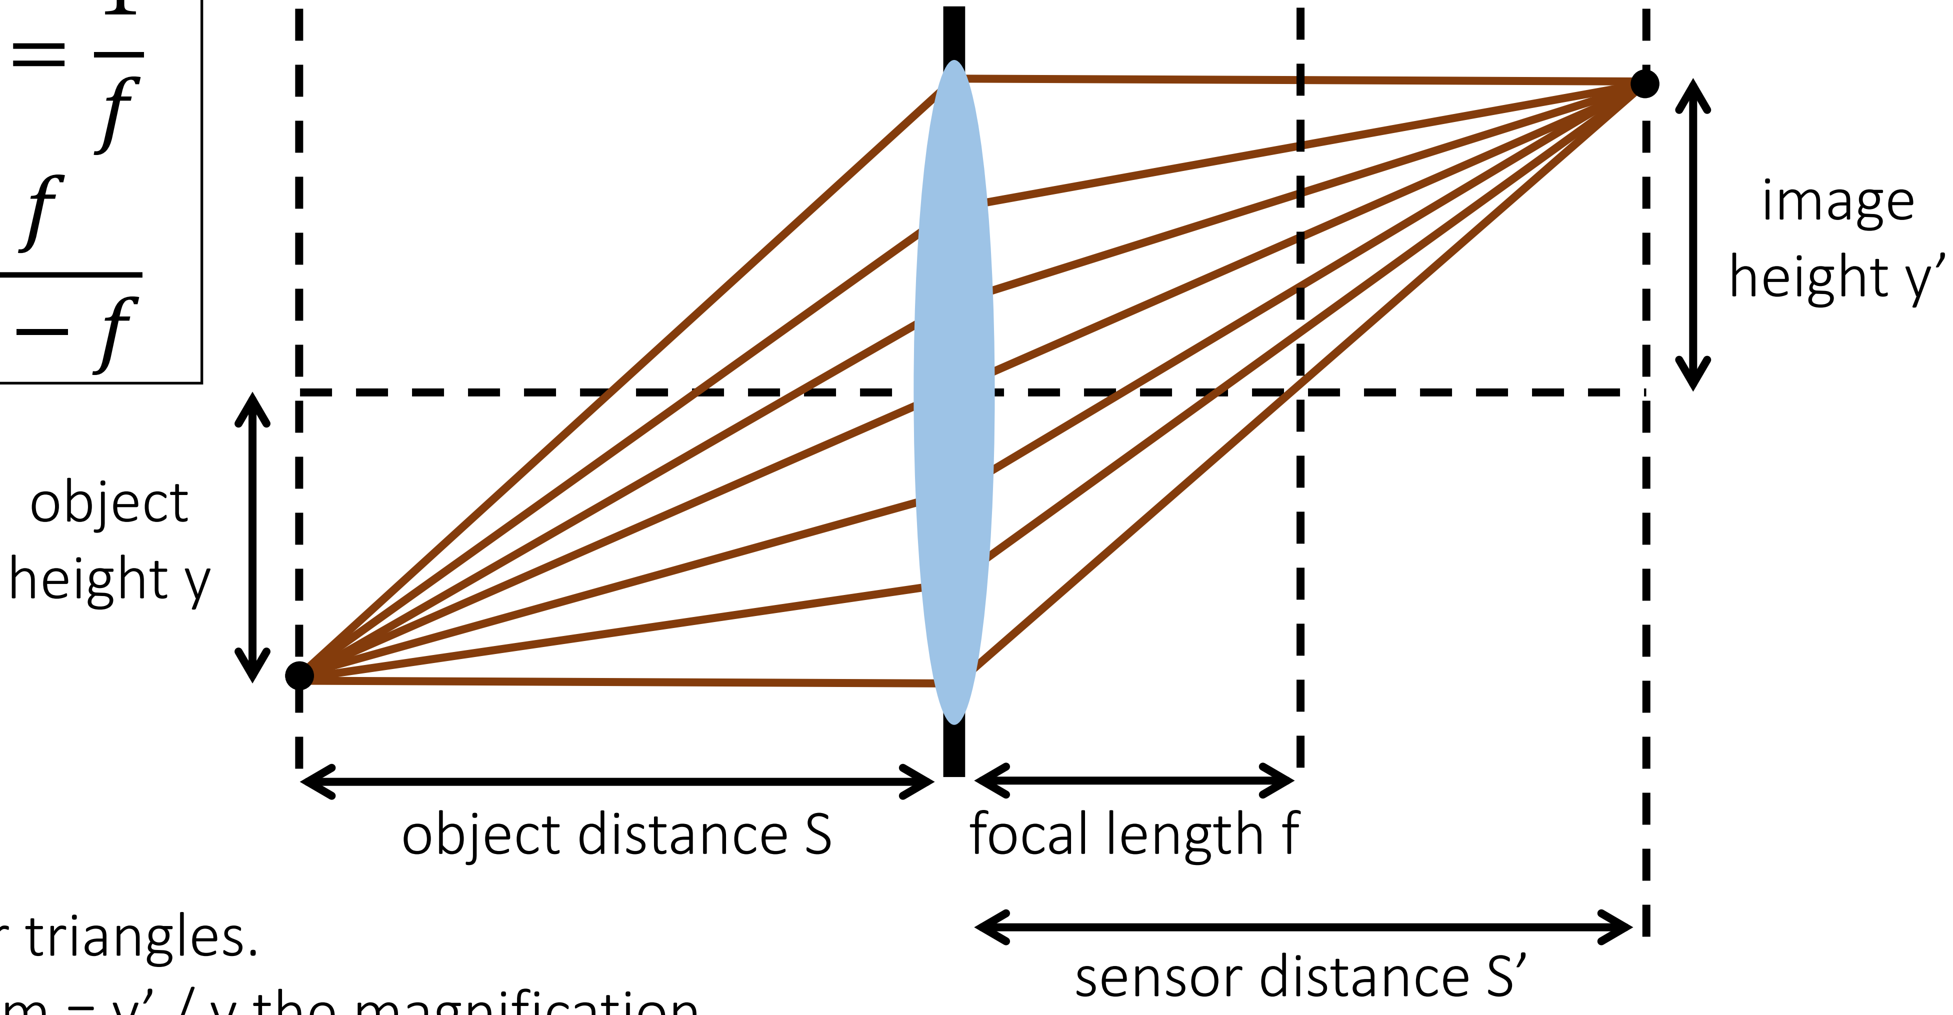

Use similar triangles.

- We call m = y' / y the magnification.

Special focus distances

S' = f, S = ?, m = ?

$$\boxed{\begin{gathered}\frac{1}{S'} + \frac{1}{S} = \frac{1}{f} \\ m = \frac{f}{S - f}\end{gathered}}$$

Special focus distances

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

$$m = \frac{f}{S - f}$$

S’ = f, S = ∞, m = 0 → infinity focus (parallel rays)

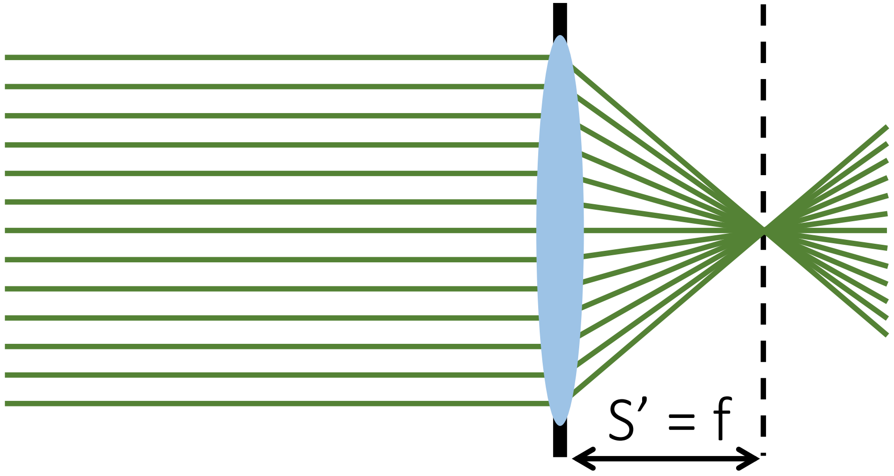

S’ = S = ?, m = ?

Special focus distances

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

$$m = \frac{f}{S - f}$$

S’ = f, S = ∞, m = 0 → infinity focus (parallel rays)

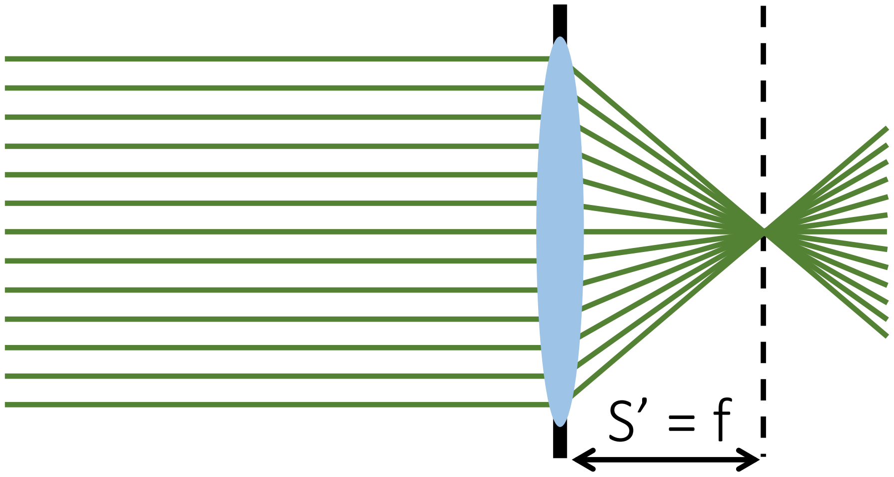

S’ = S = 2 f, m = 1 → object is reproduced in real-life size

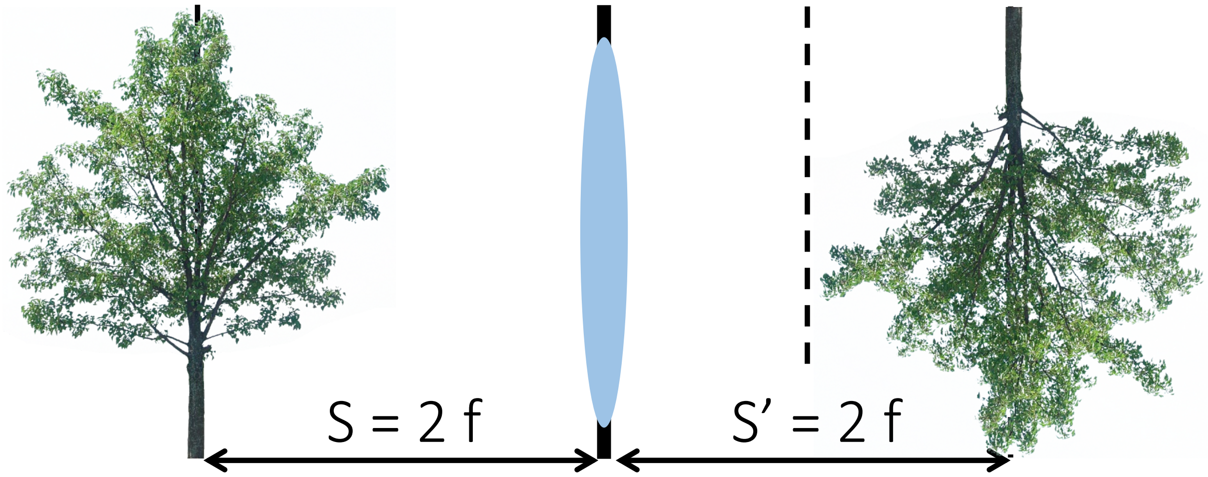

Free lunch?

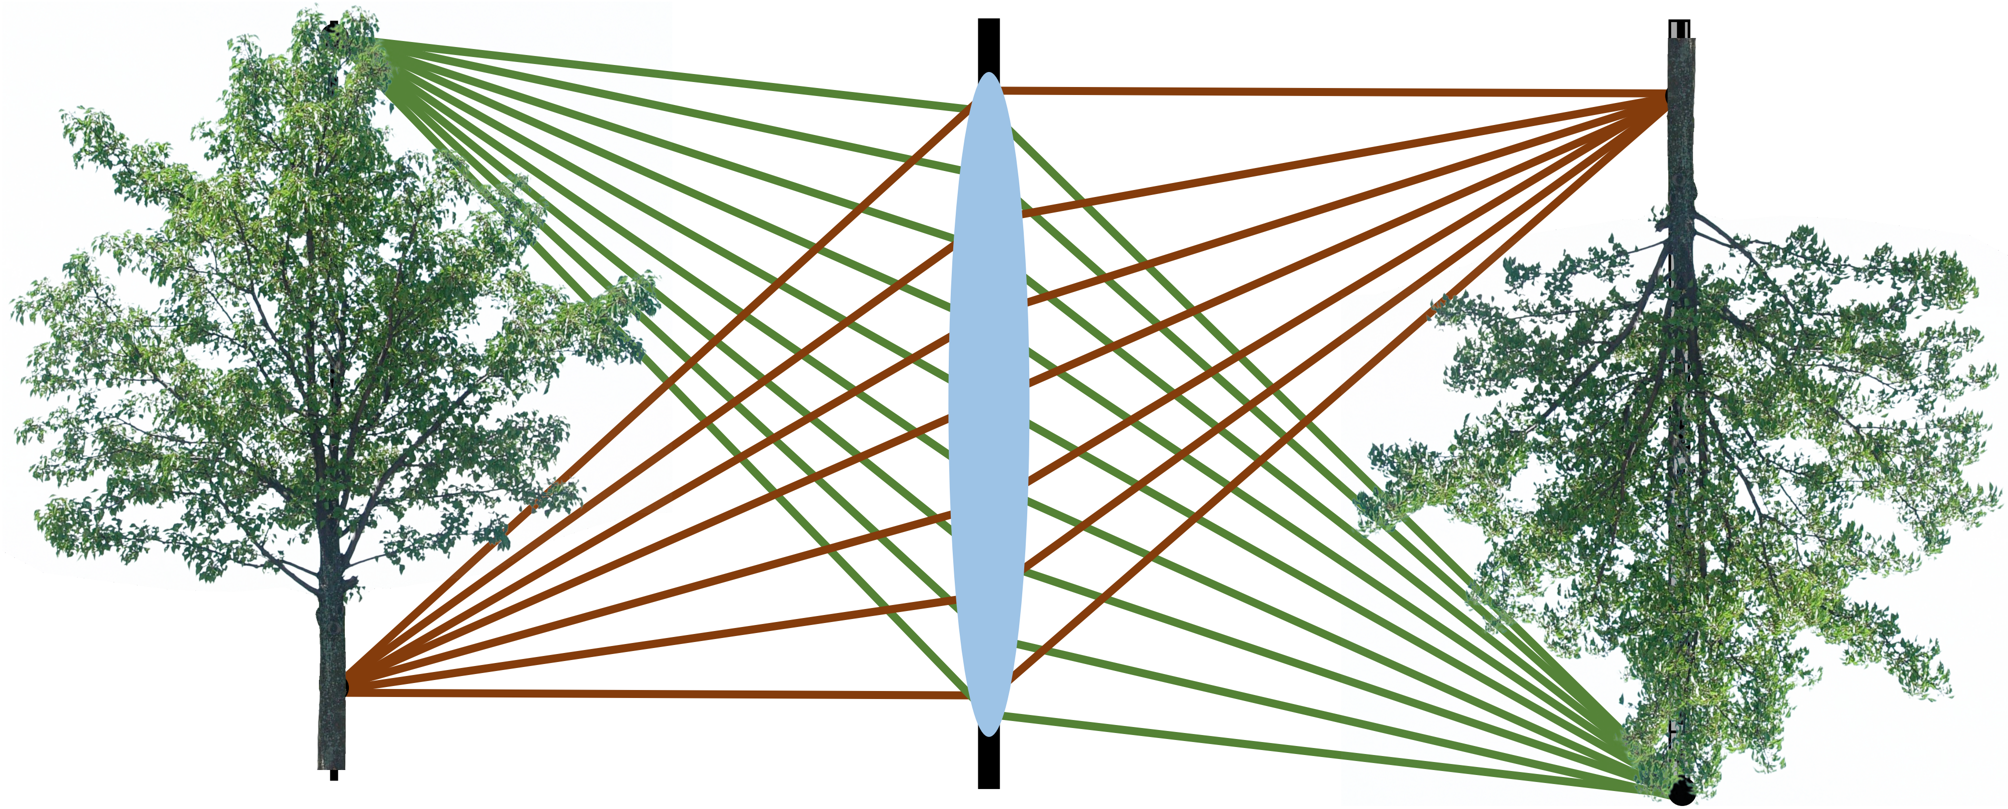

By using a lens, we simultaneously achieve:

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

Do we lose anything by using a lens?

Defocus

What happens if we don't place the sensor at the focus distance?

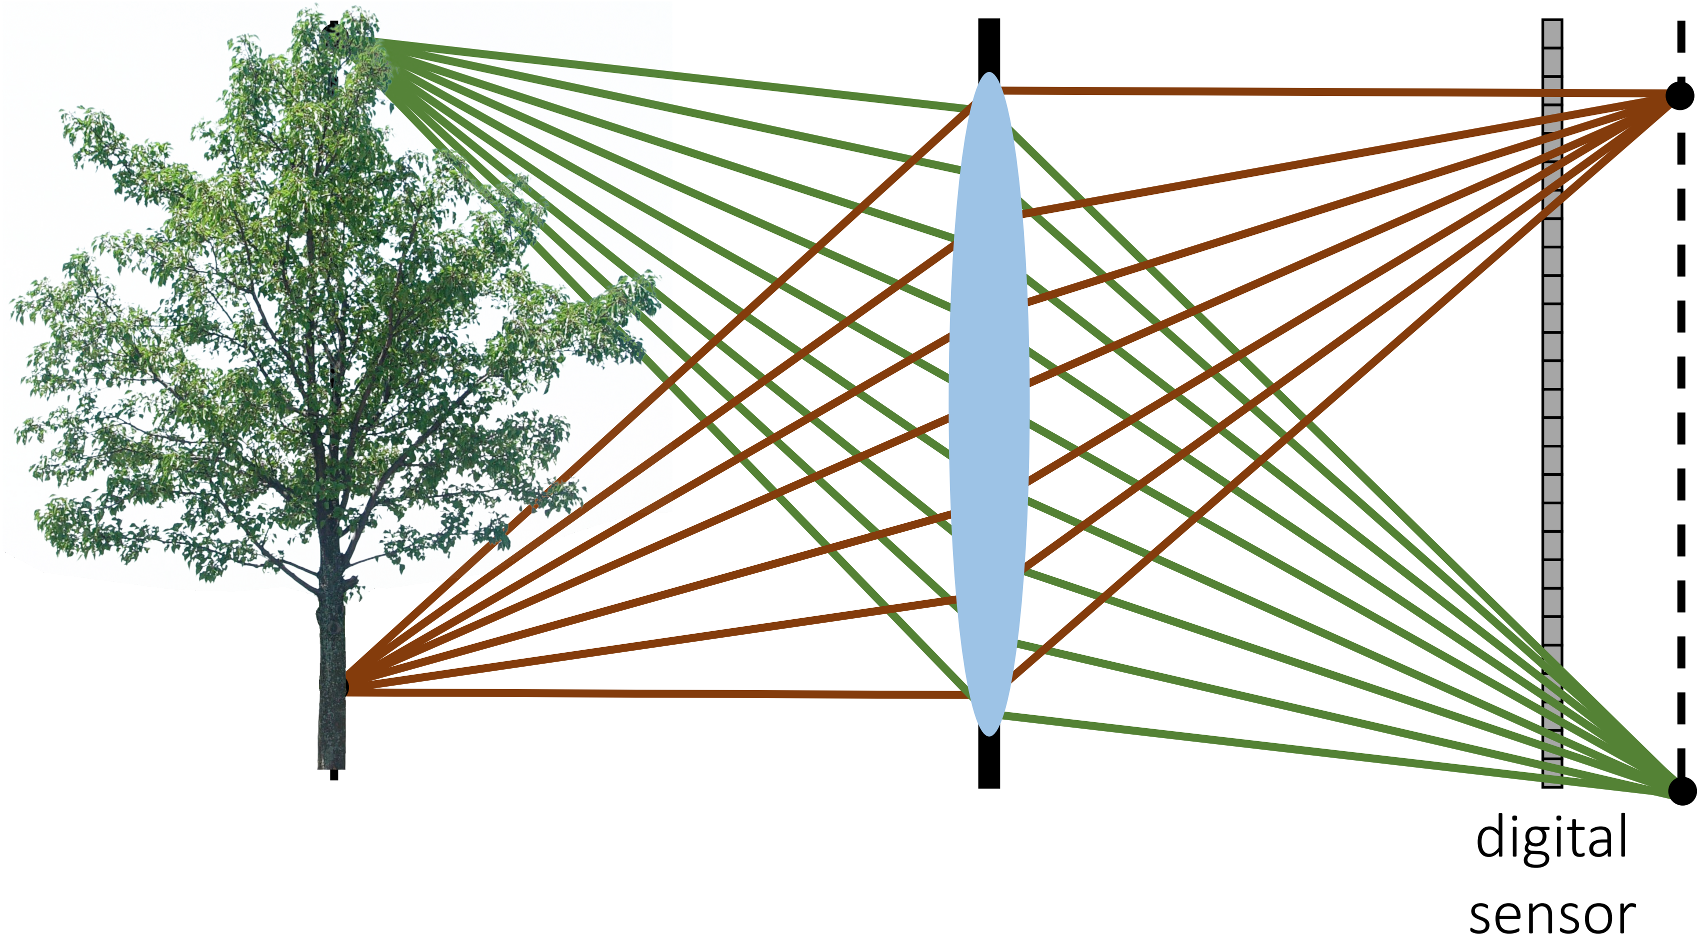

Defocus

What happens if we don’t place the sensor at the focus distance?

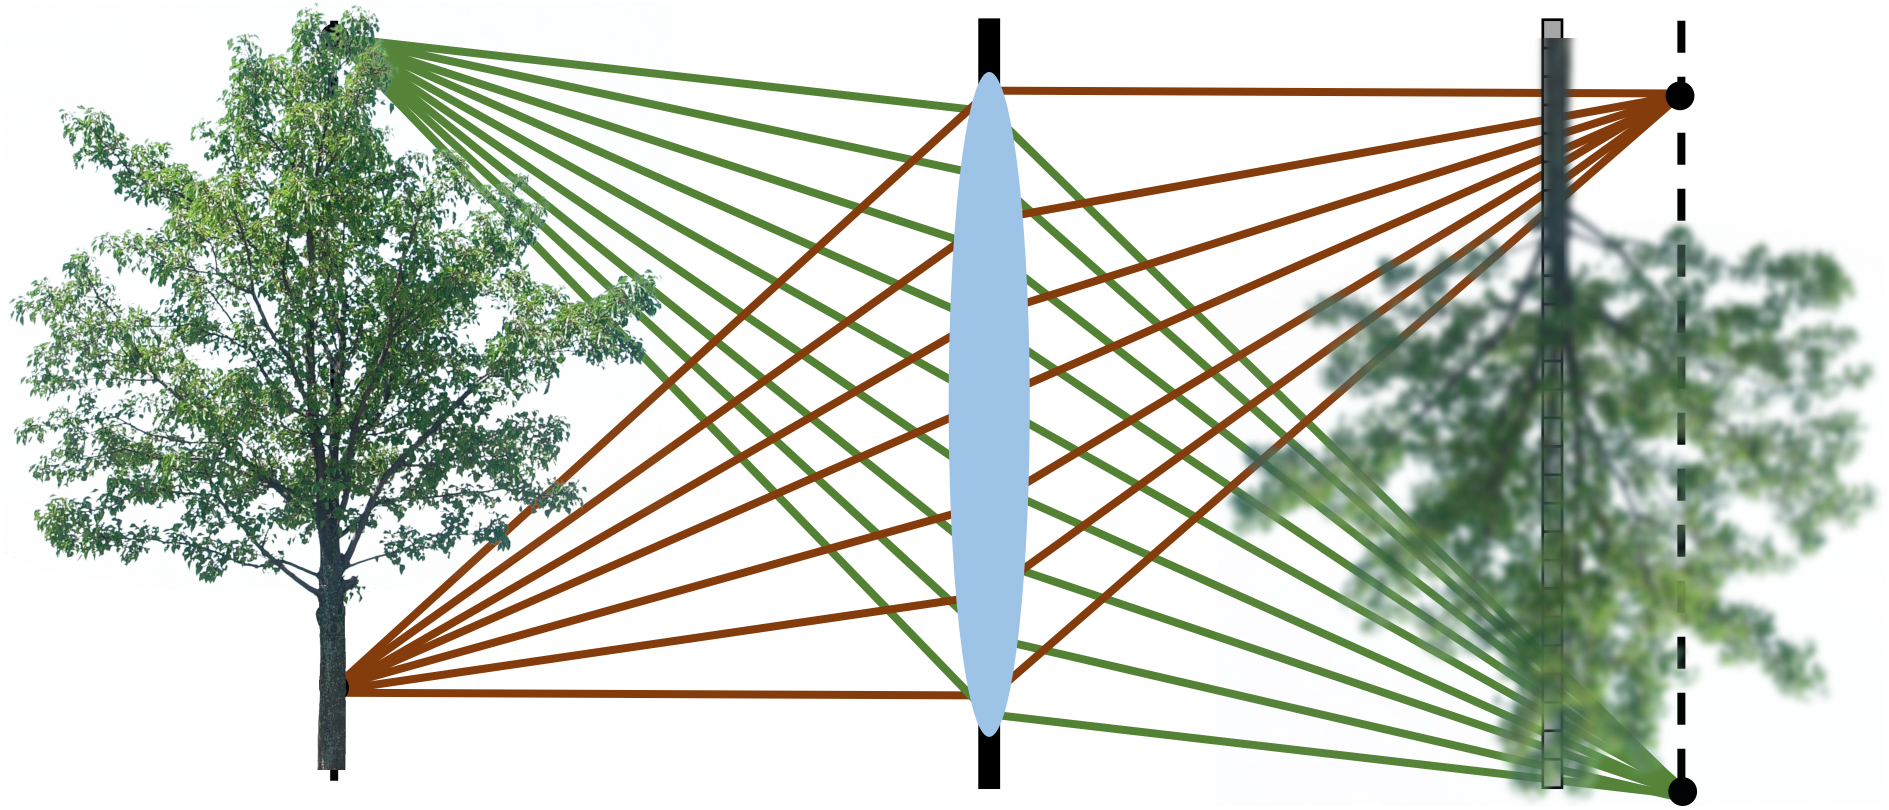

We get a blurry image. This is called defocus.

- Defocus never happens with an ideal pinhole camera.

Defocus

Can’t we just move the sensor to the correct distance?

Defocus

Can't we just move the sensor to the correct distance?

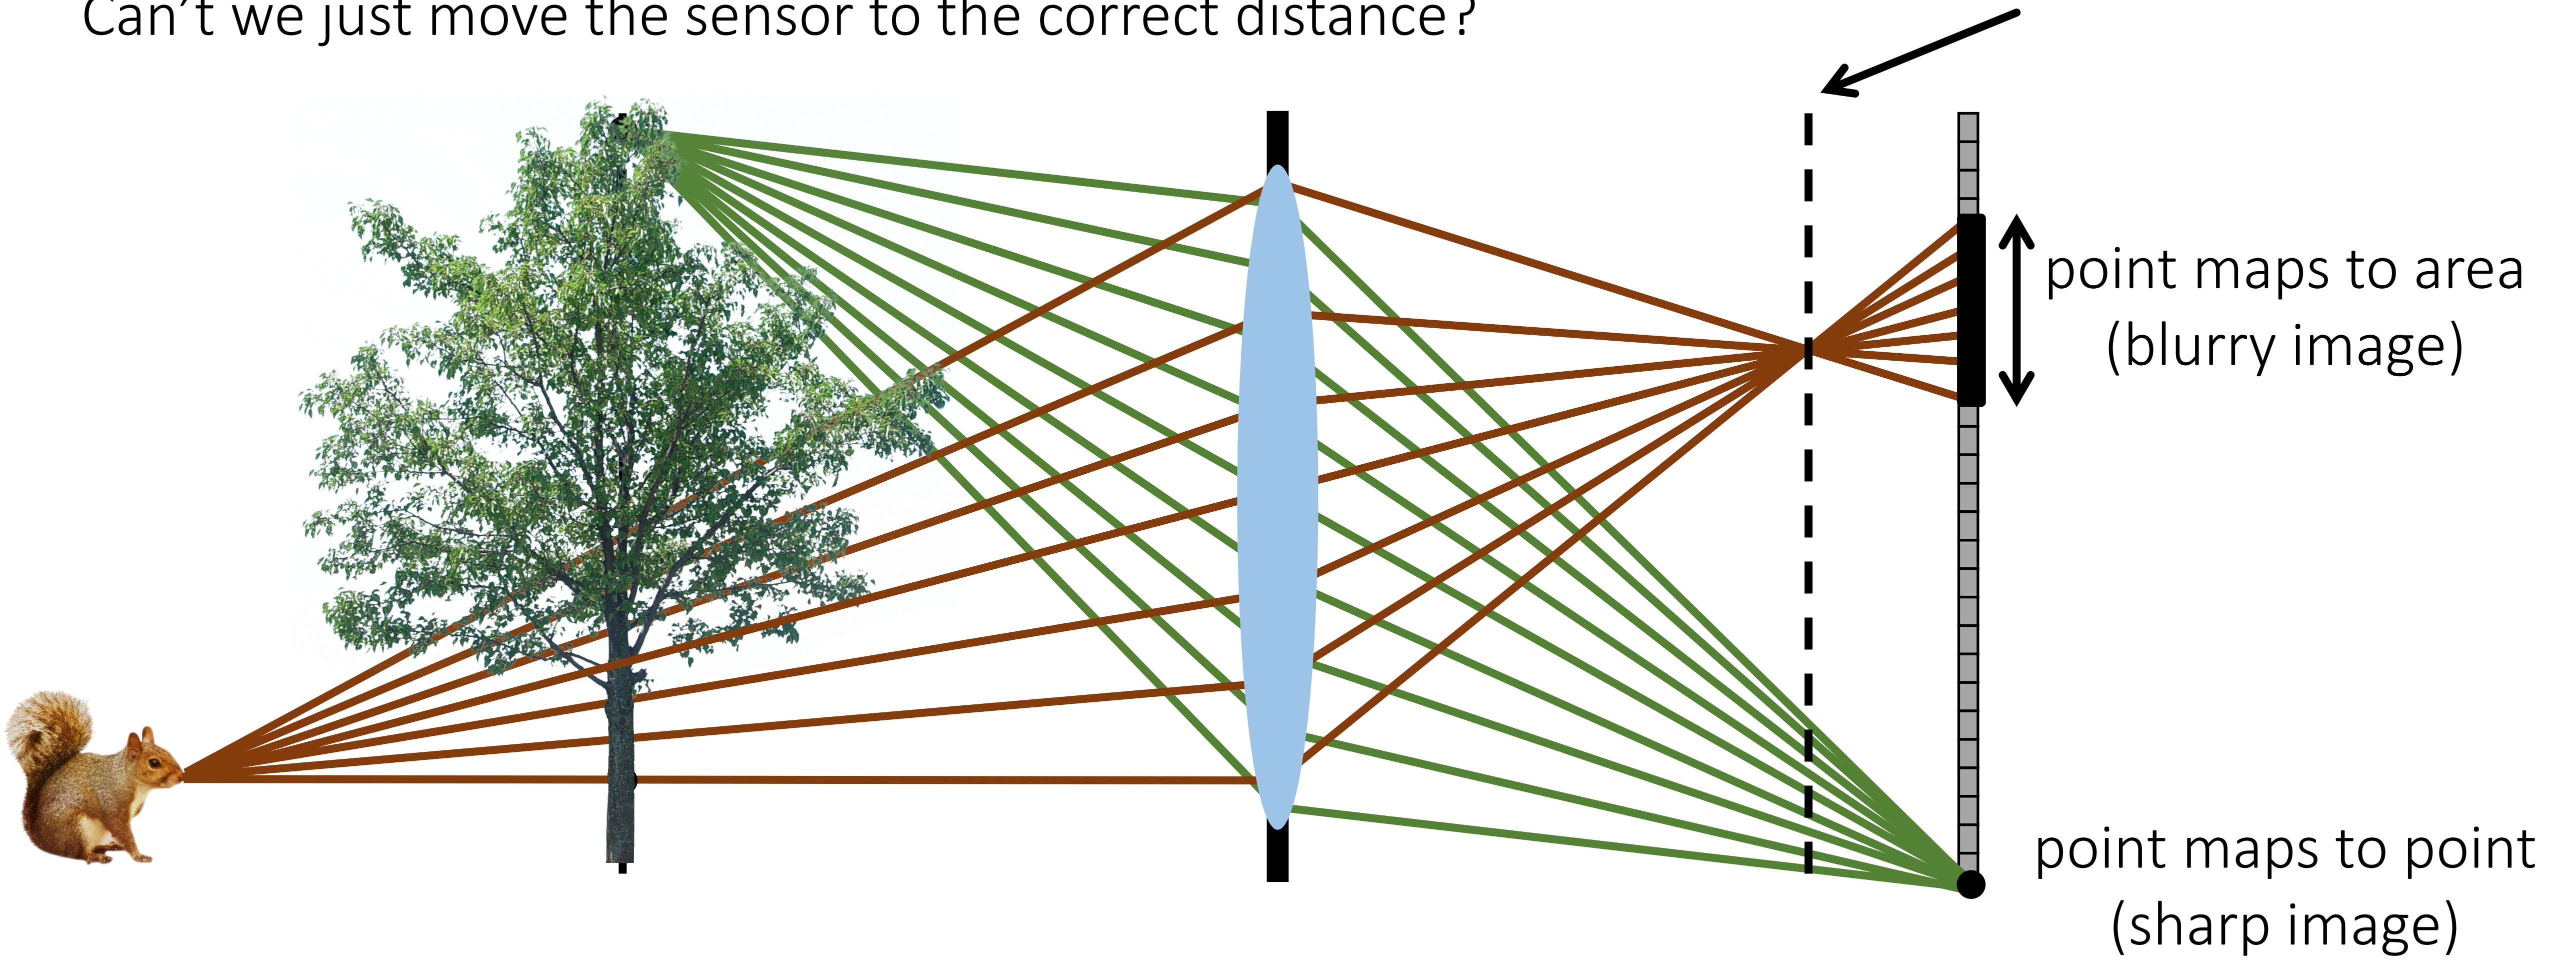

Unless our scene is just one plane, part of it will always be out of focus.

How do we control what is in focus?

How do we control what is in focus?

We change the distance between the sensor and the lens

How do we control what is in focus?

We change the distance between the sensor and the lens

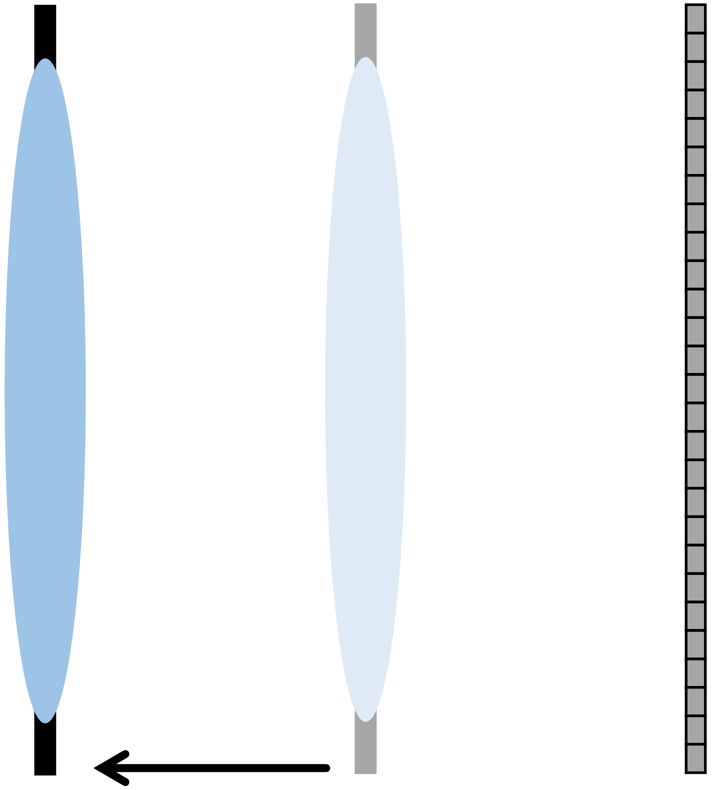

- What happens to plane in focus?

How do we control what is in focus?

We change the distance between the sensor and the lens

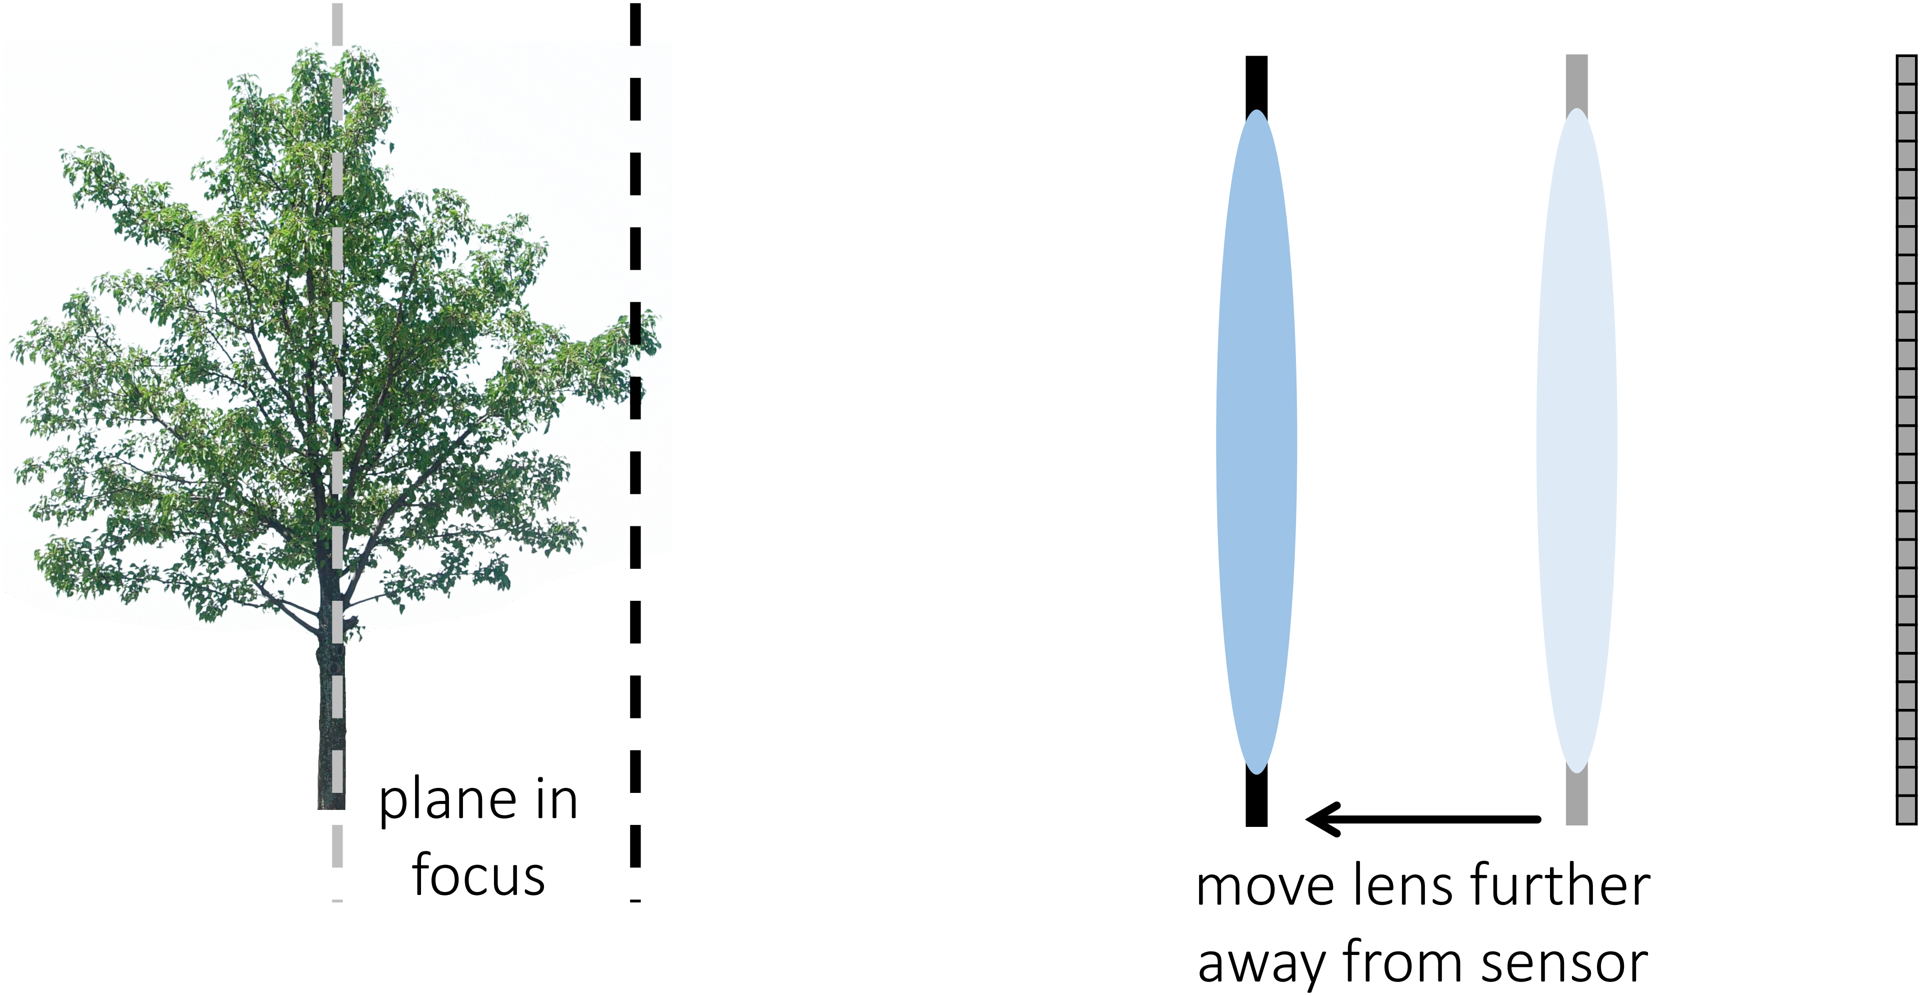

- What happens to plane in focus? → It moves closer.

The lens on your camera

Focus ring: controls distance of lens from sensor

Demonstration

Defocus

Does the mean that lenses are only good for planar scenes?

Circle of confusion

How do we find where the point will focus?

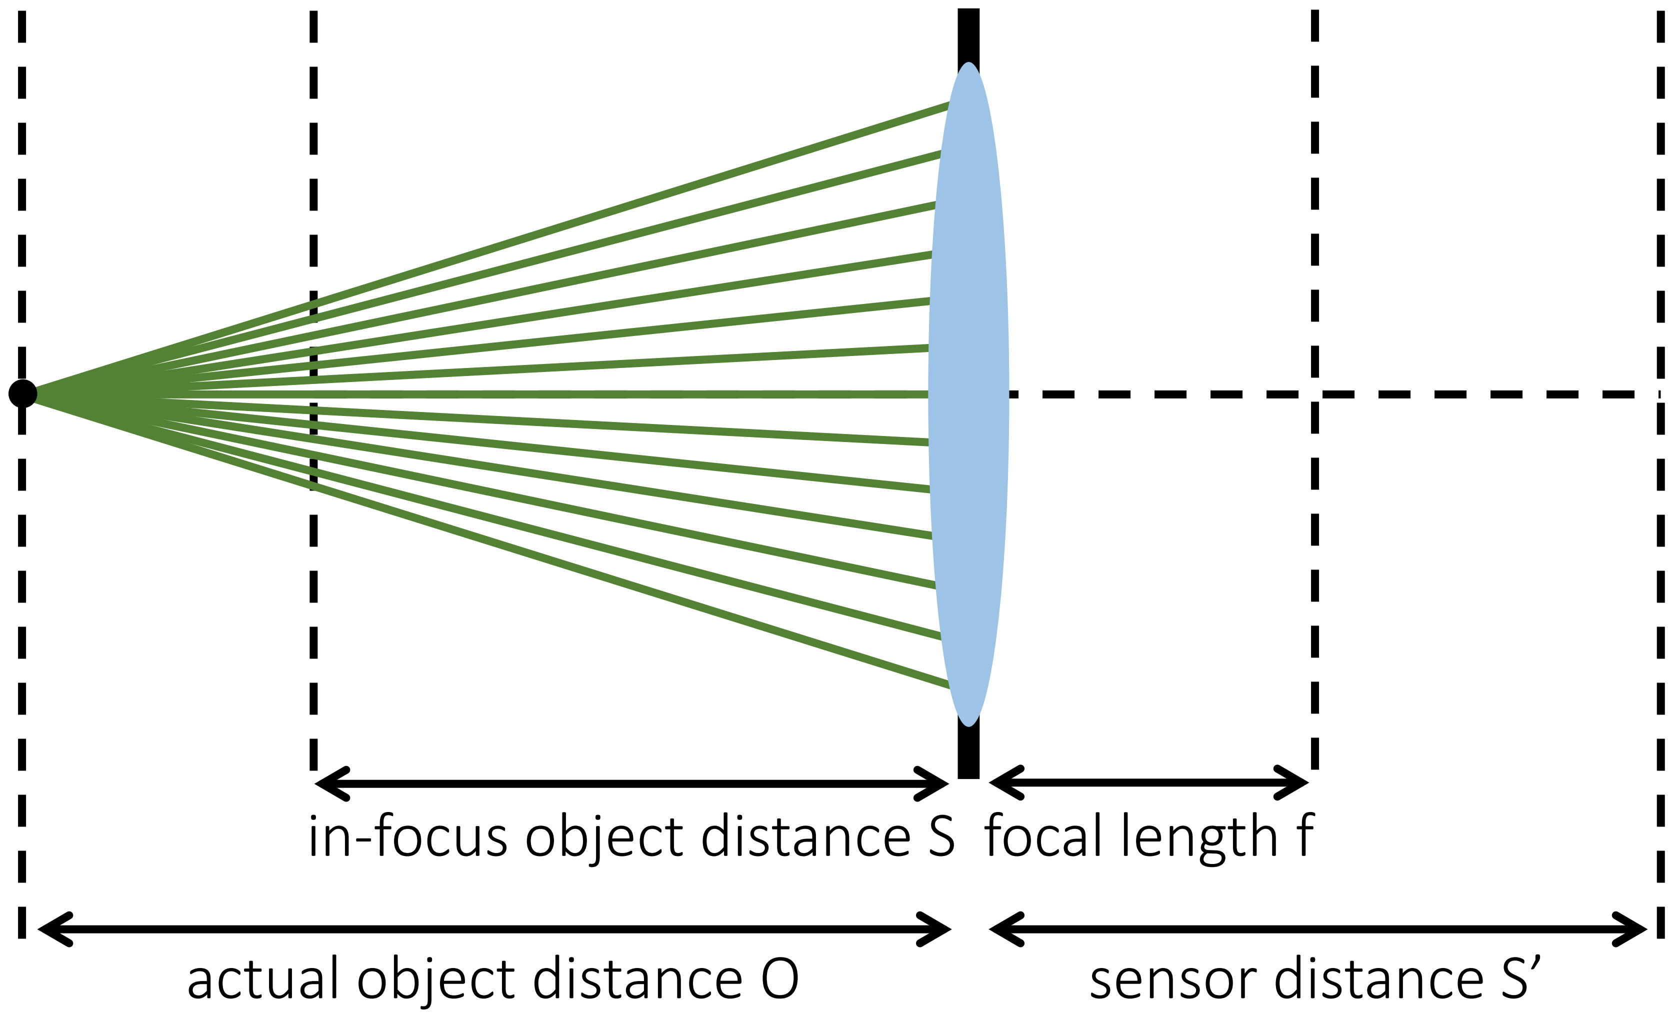

Circle of confusion

Will the point focus at a distance smaller or larger than S'?

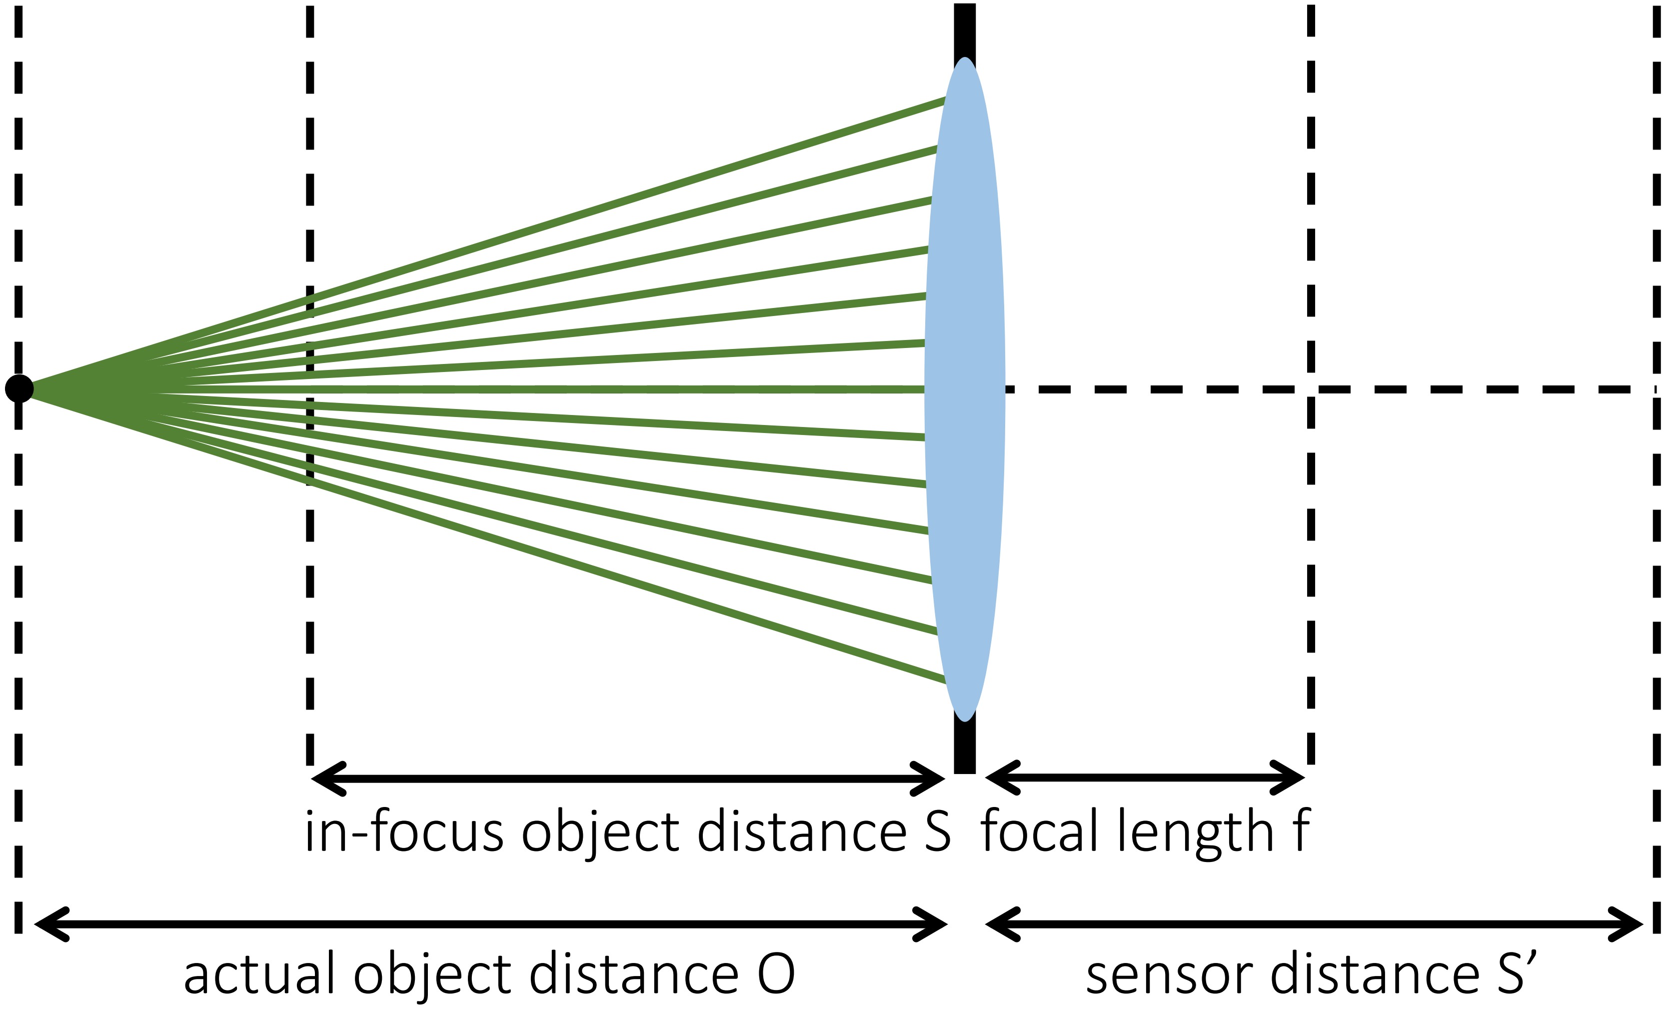

Circle of confusion

How can we compute the diameter of the circle of confusion?

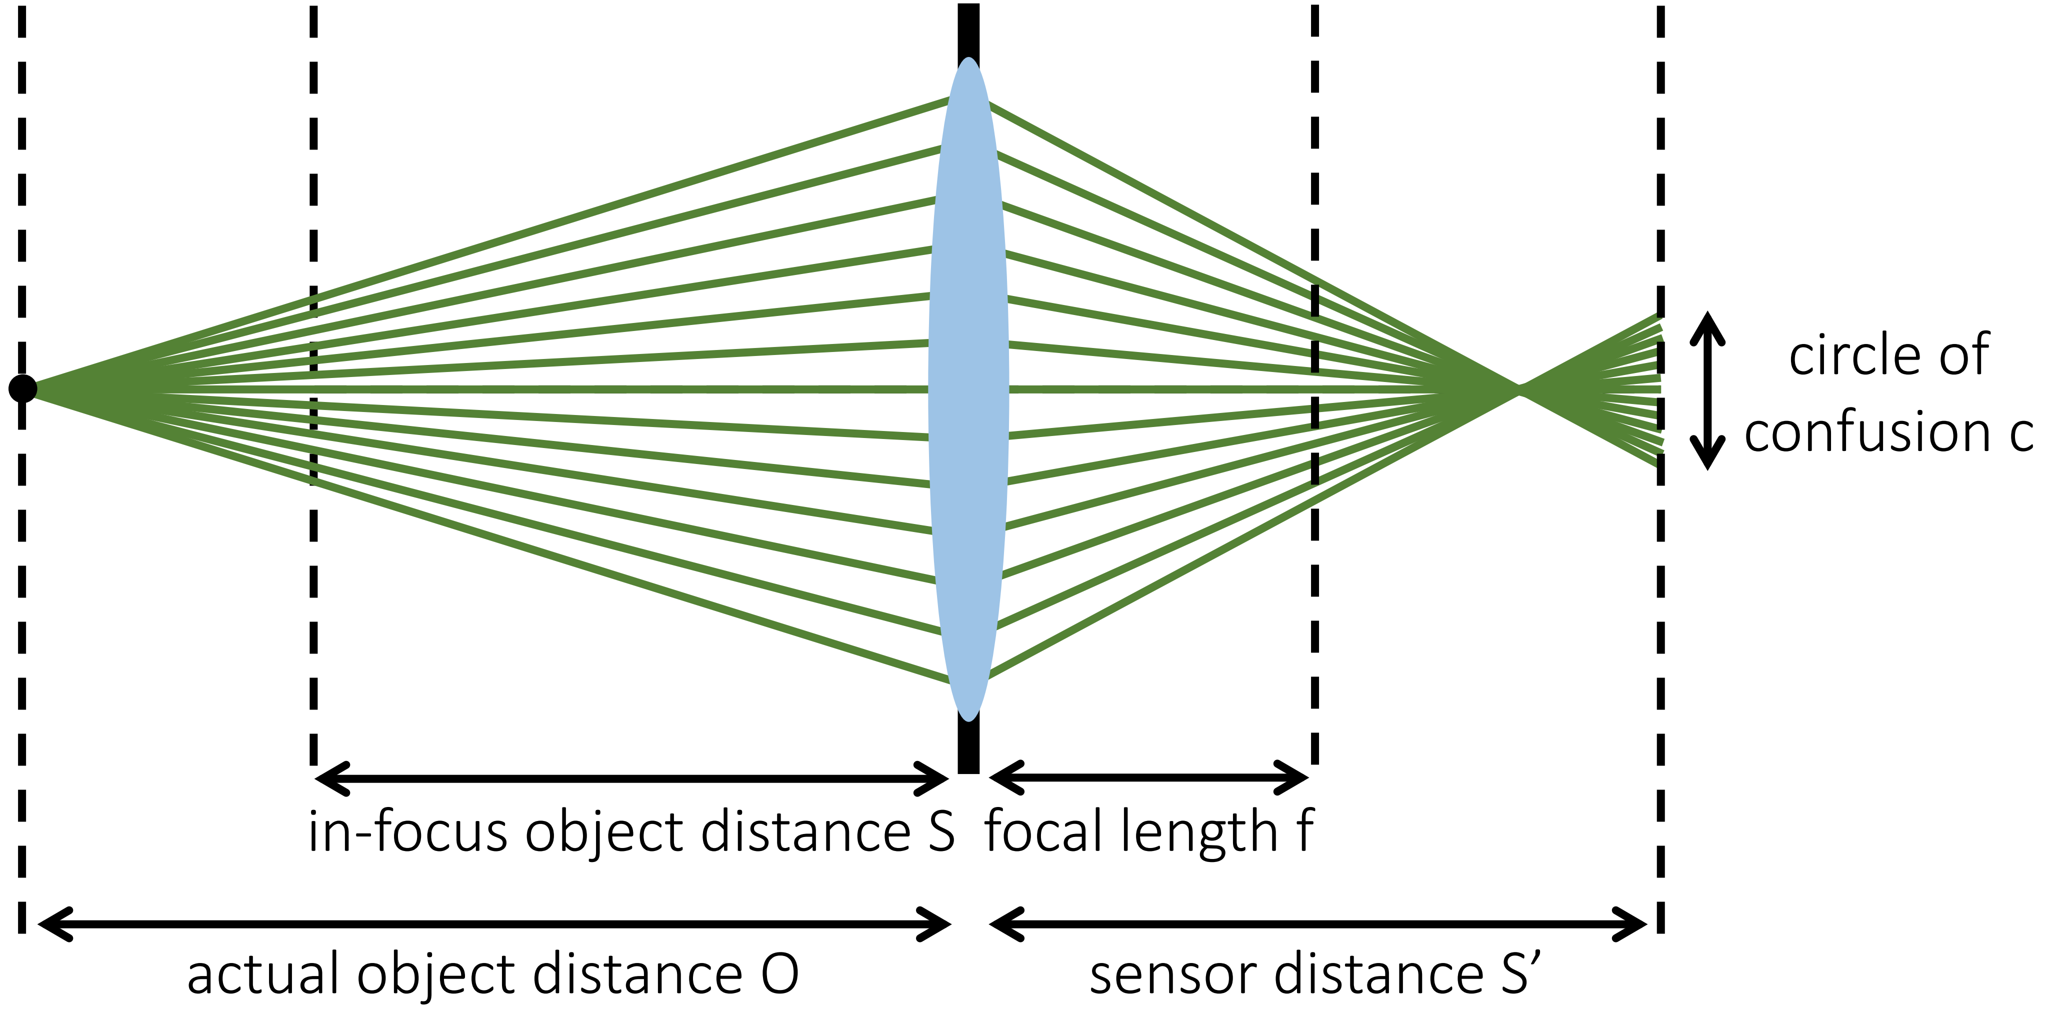

Circle of confusion

How can we compute the diameter of the circle of confusion? → Use similar triangles.

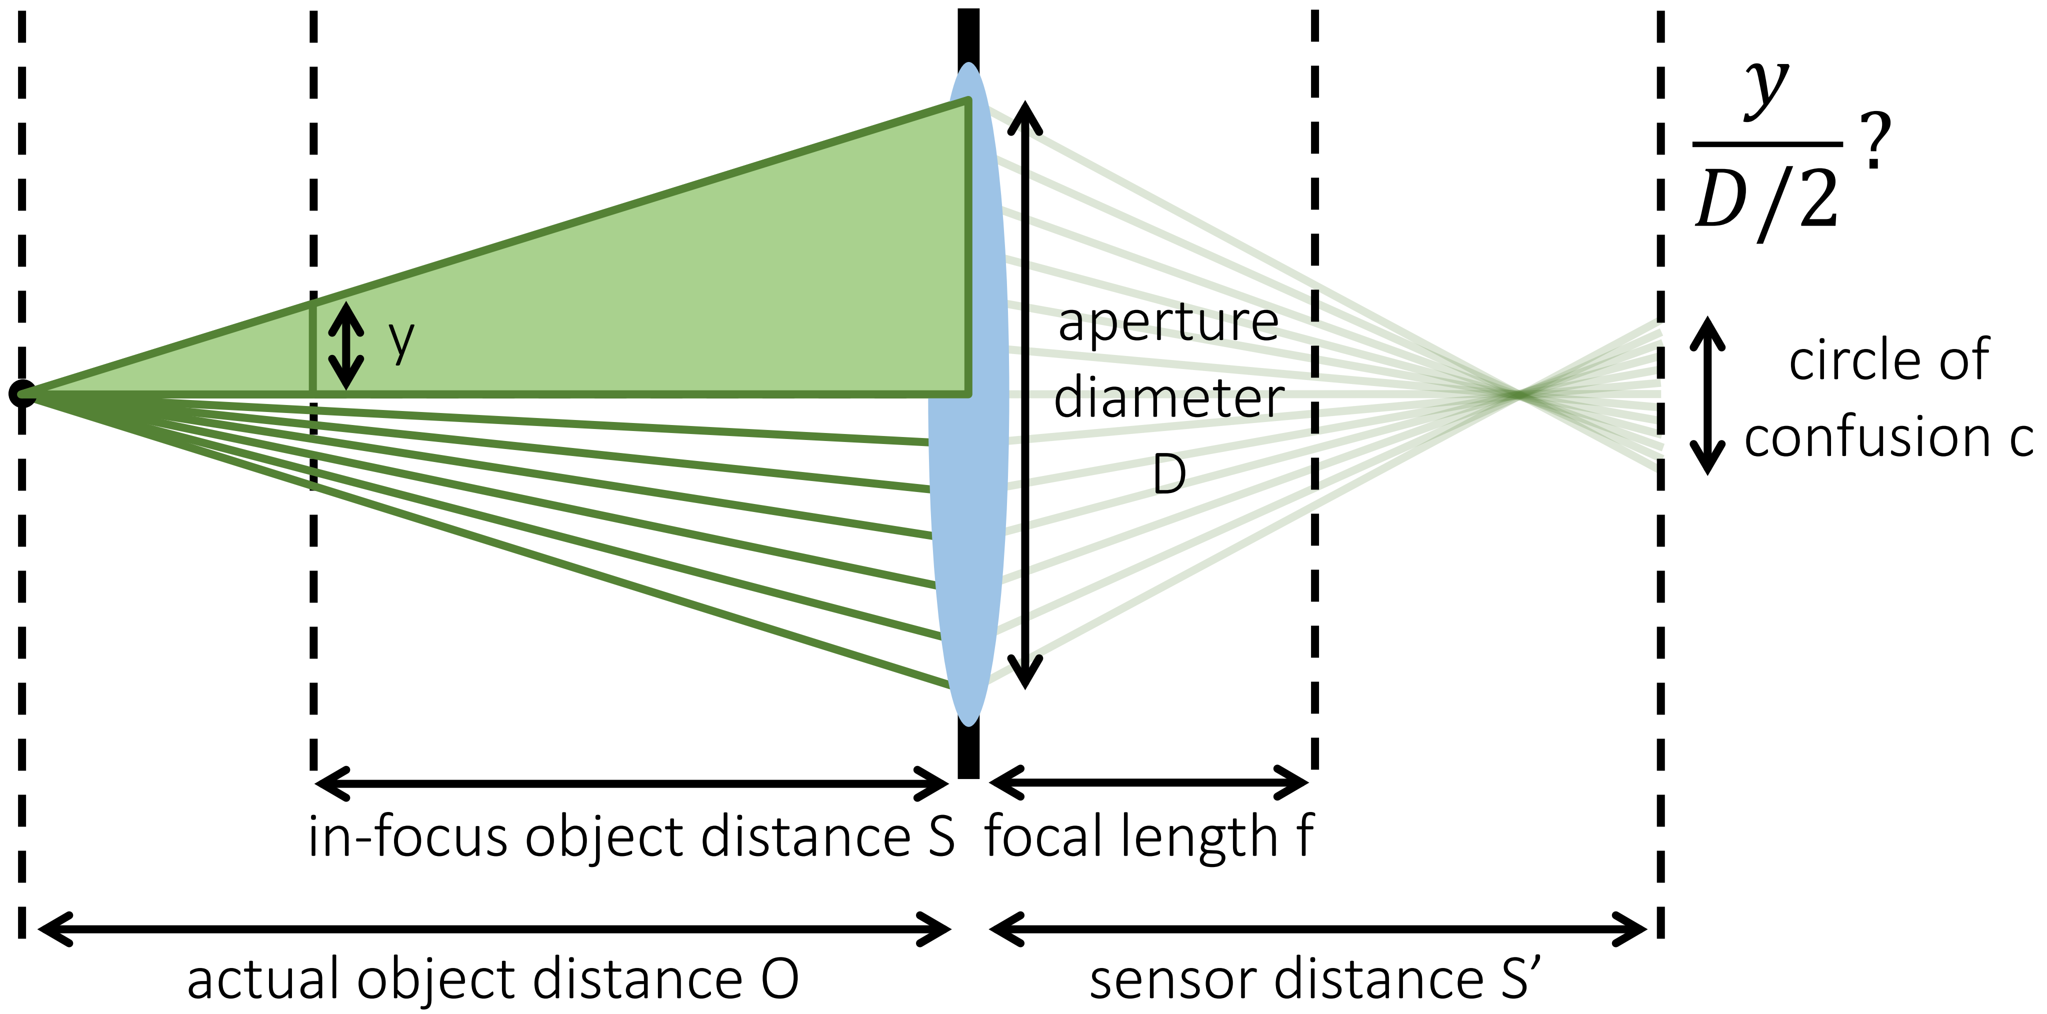

Circle of confusion

How can we compute the diameter of the circle of confusion? → Use similar triangles.

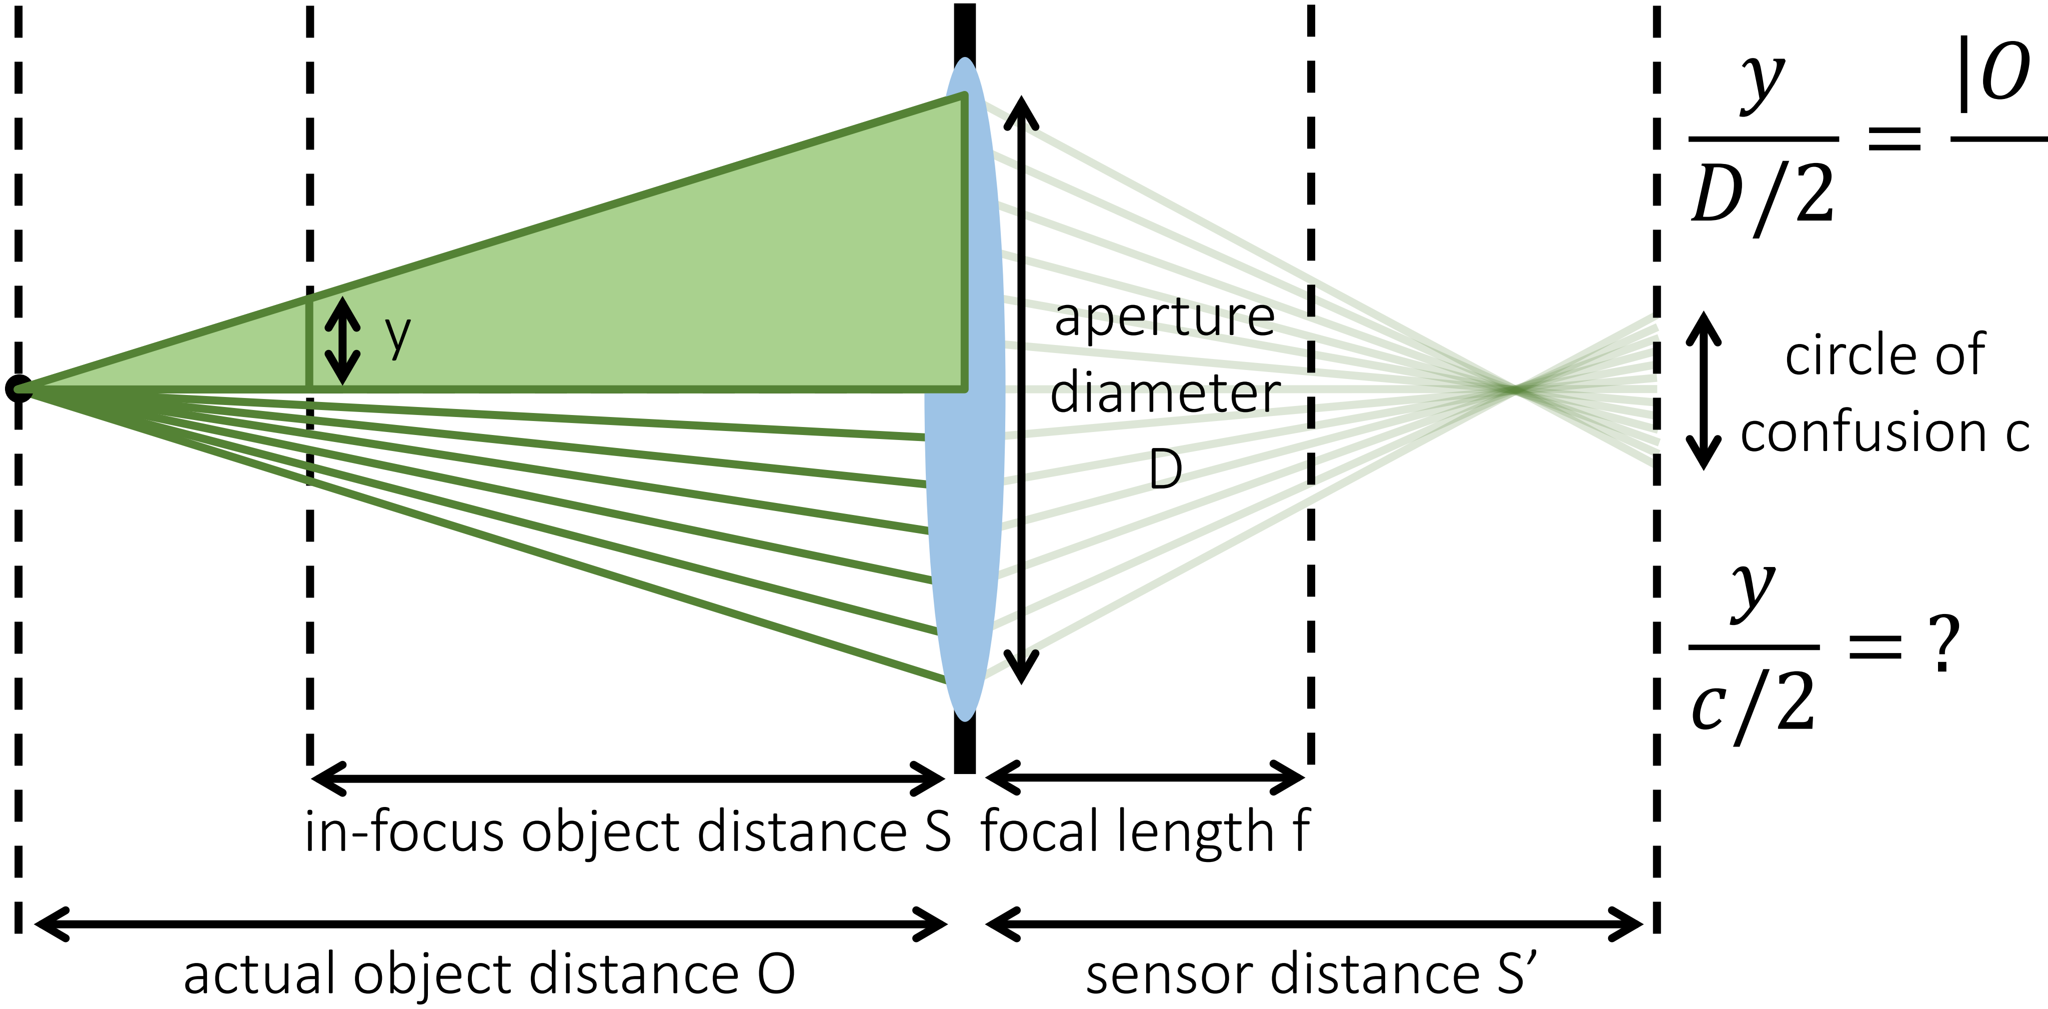

$$\frac{y}{D/2} = \frac{|O - S|}{O}$$

$$\frac{y}{c/2} = ?$$

Circle of confusion

How can we compute the diameter of the circle of confusion? → Use similar triangles.

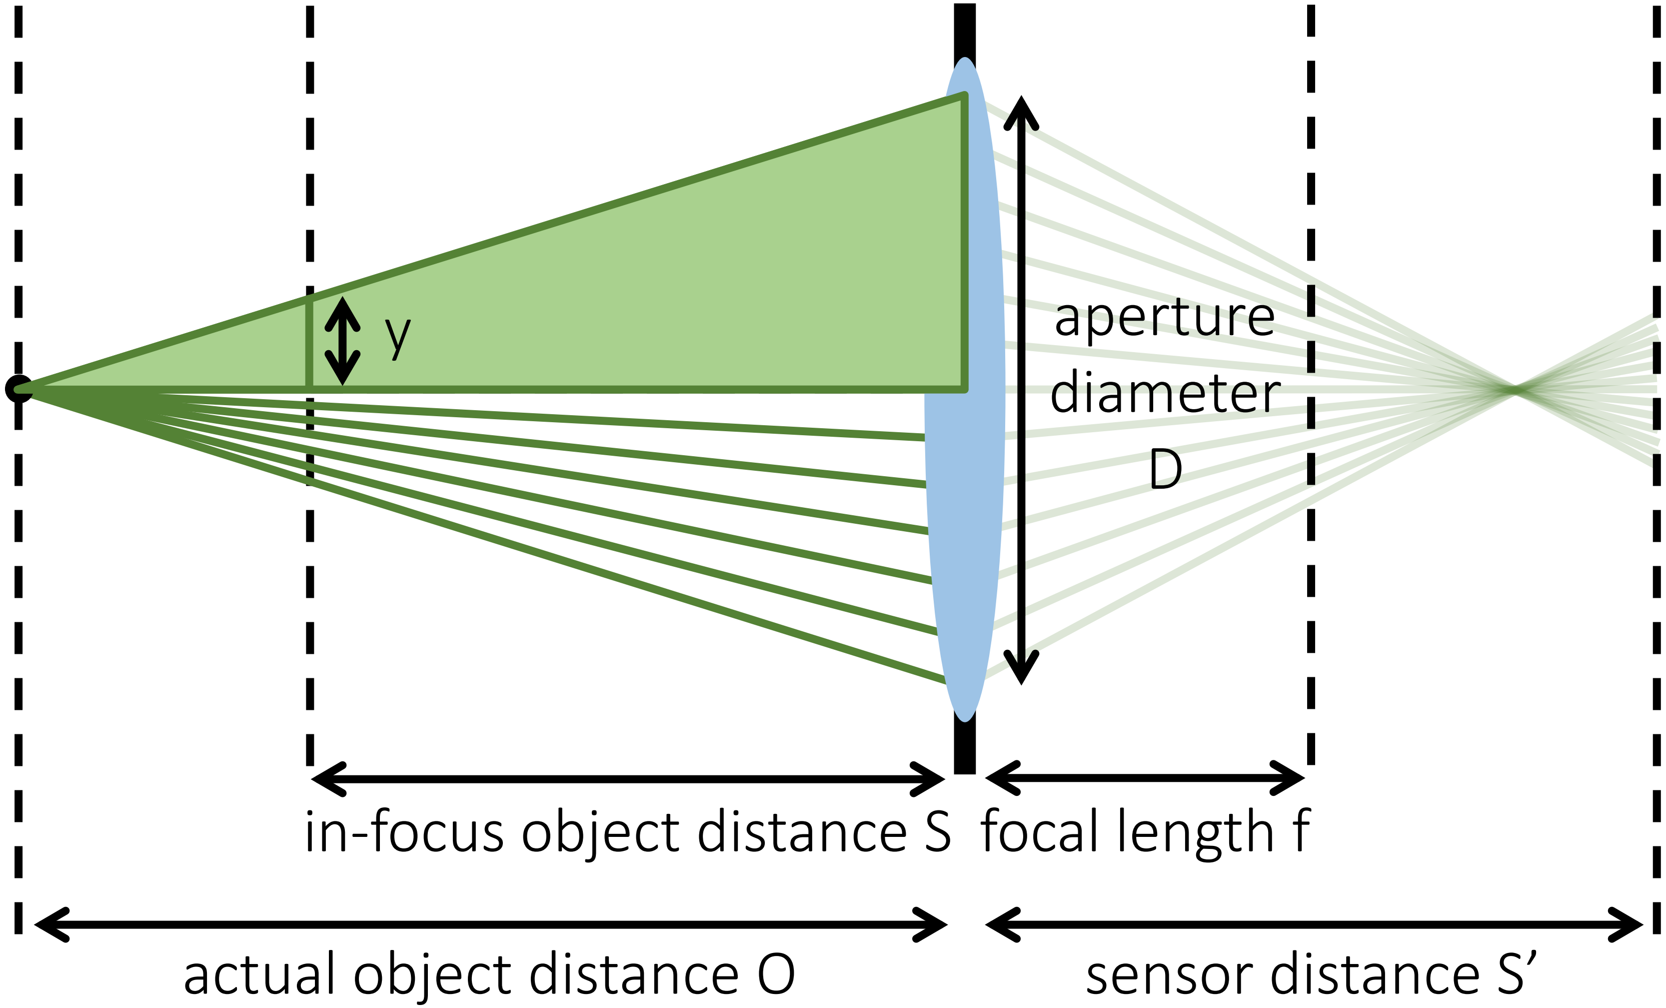

$$\frac{y}{D/2} = \frac{|O - S|}{O}$$

$$\frac{y}{c/2} = \frac{1}{m}$$

Circle of confusion

How can we compute the diameter of the circle of confusion? → Use similar triangles.

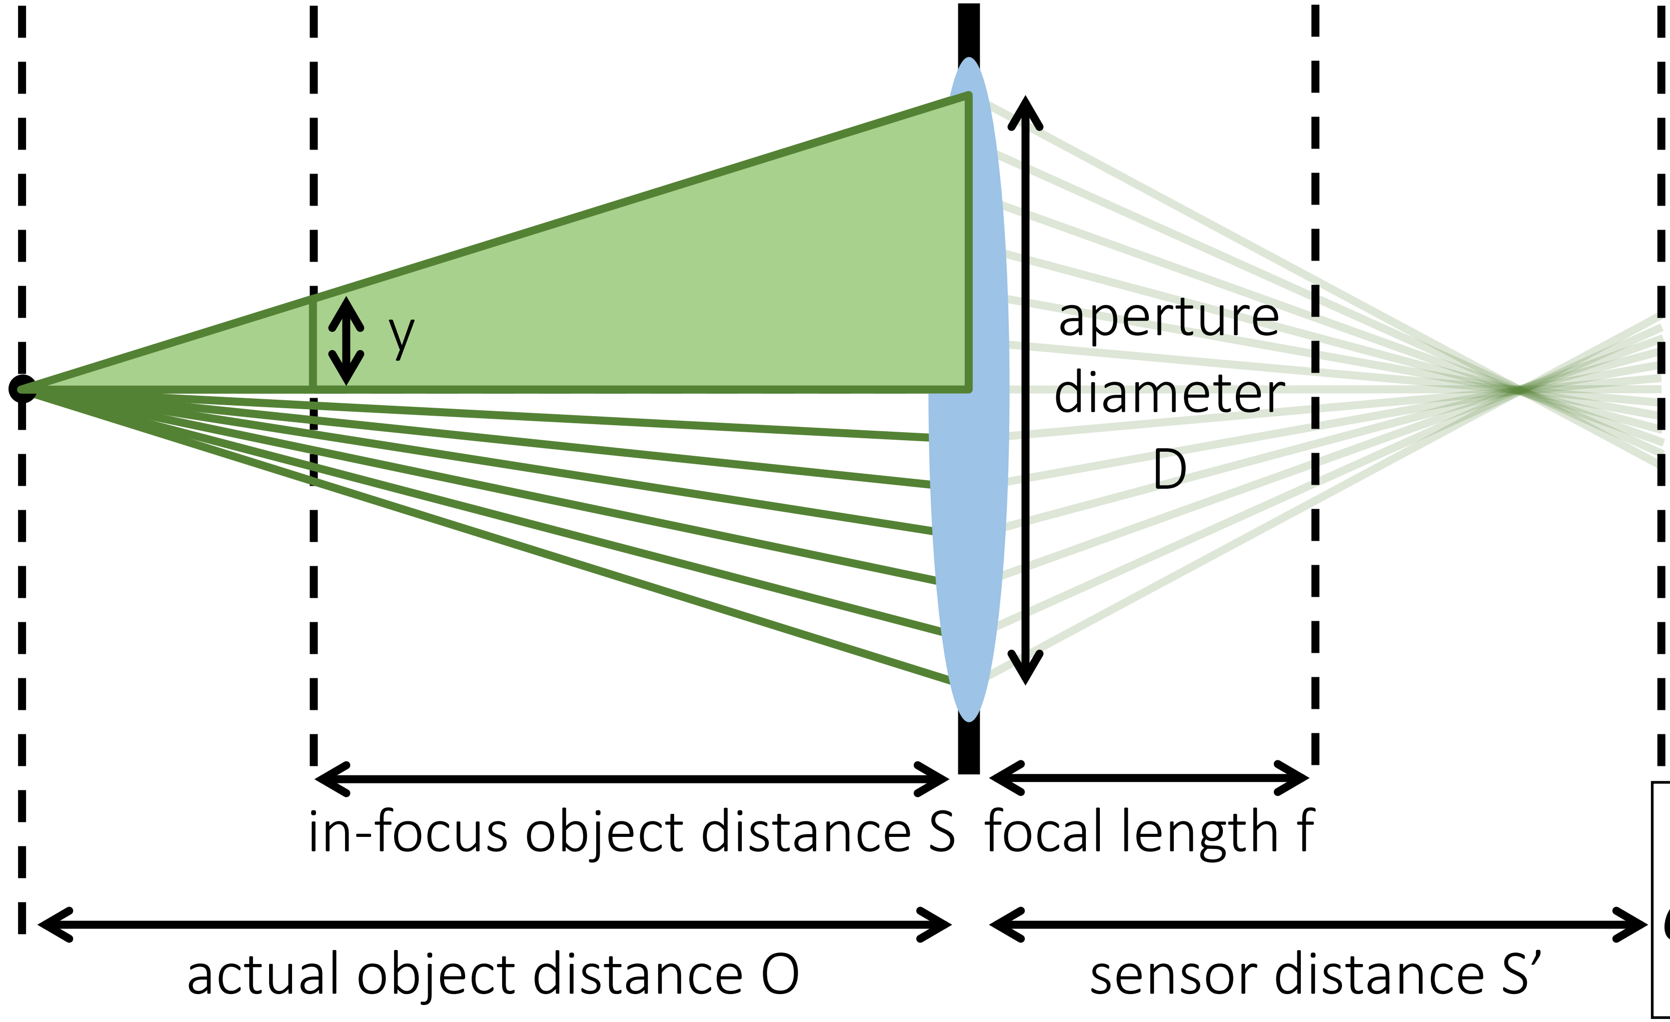

$$\frac{y}{D/2} = \frac{|O-S|}{O}$$

$$\frac{y}{c/2} = \frac{1}{m}$$

$$\boxed{c = mD\frac{|O-S|}{O}}$$

Depth of field

Distance from the in-focus object plane where the circle of confusion is *acceptably small.*

e.g., equal to 4-5 pixels

$$c < \varepsilon \Rightarrow$$

$$\mathrm{DOF} = \frac{2\varepsilon O}{mD}$$

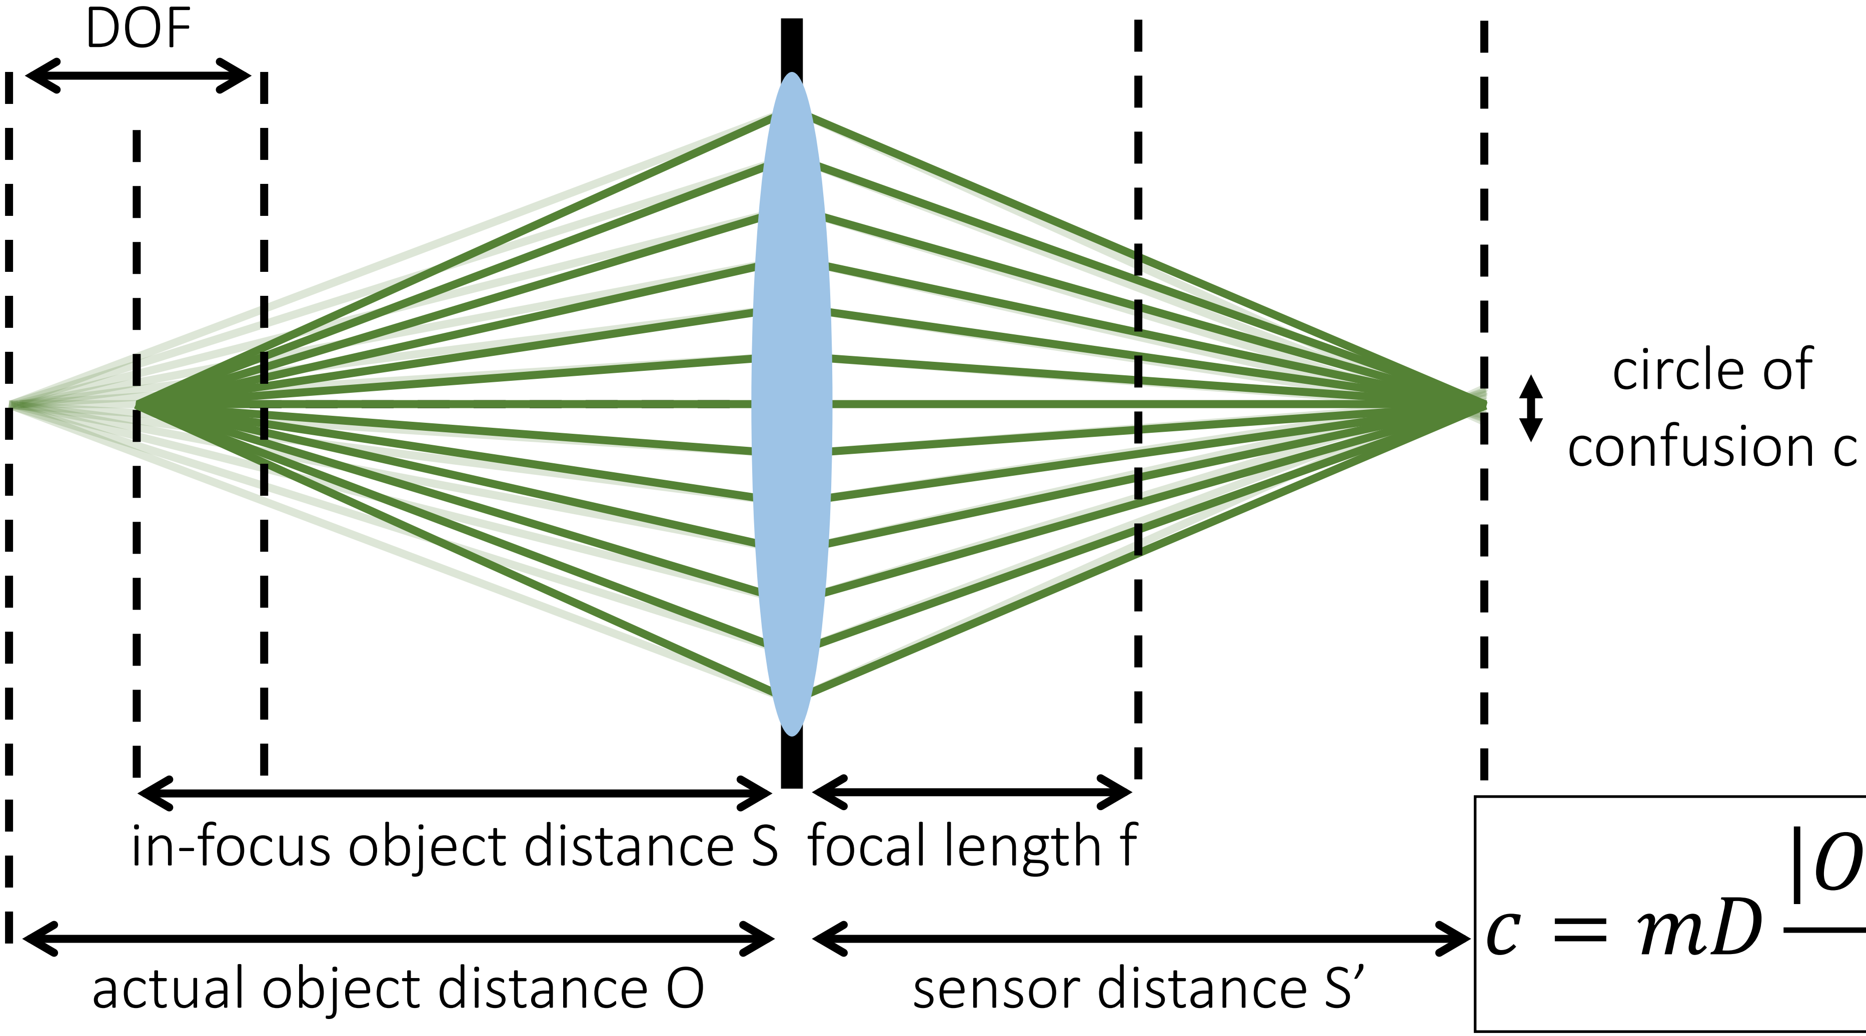

$$\boxed{c = mD\,\frac{|O - S|}{O}}$$

Note: in reality, DOF is slightly asymmetrical.

Depth of field

depth of field

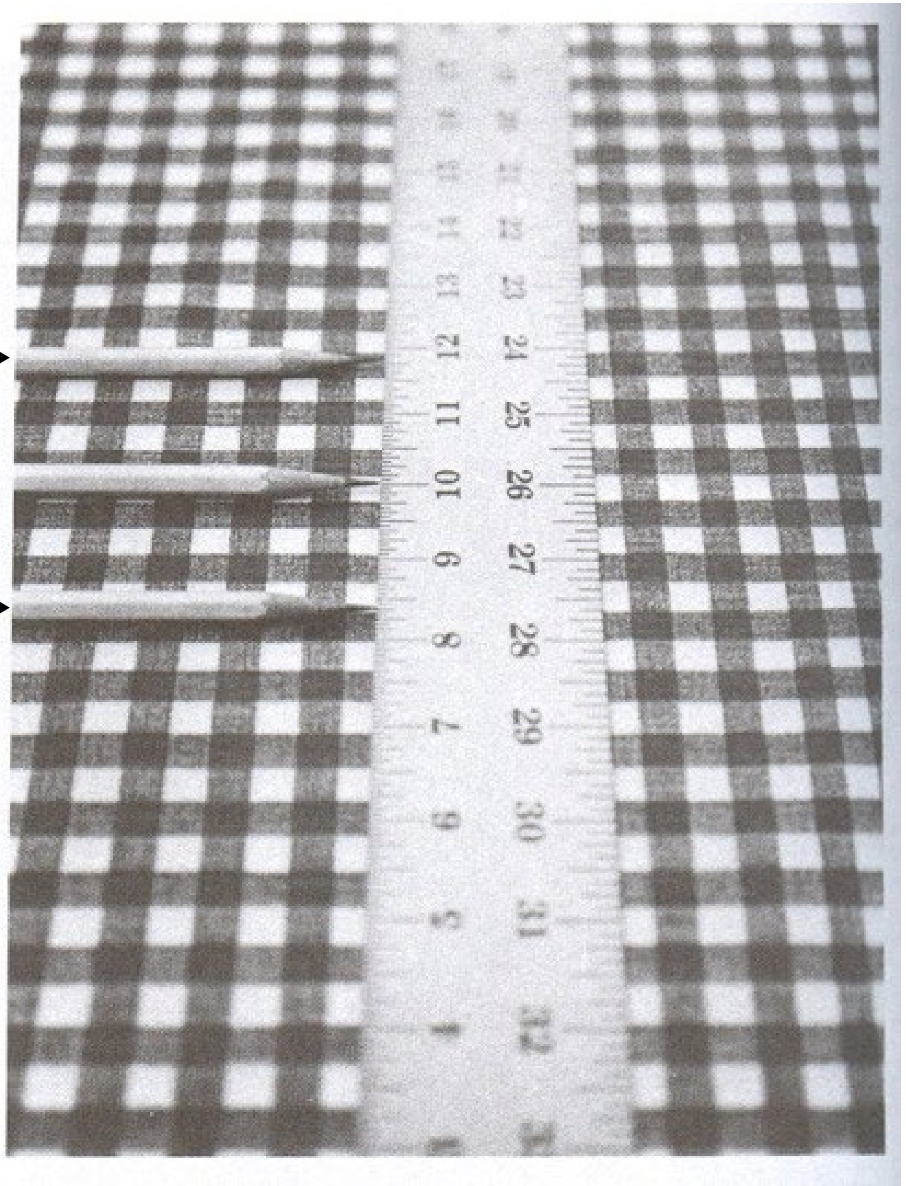

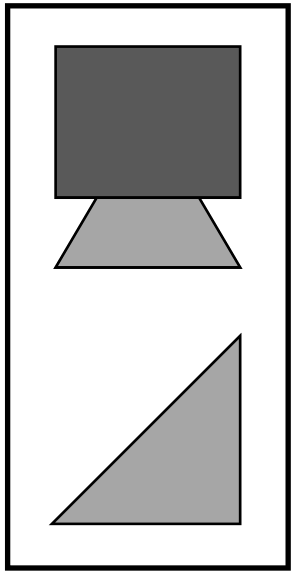

scene

Circle of Confusion

$$\boxed{c = mD\frac{|O - S|}{O}}$$

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

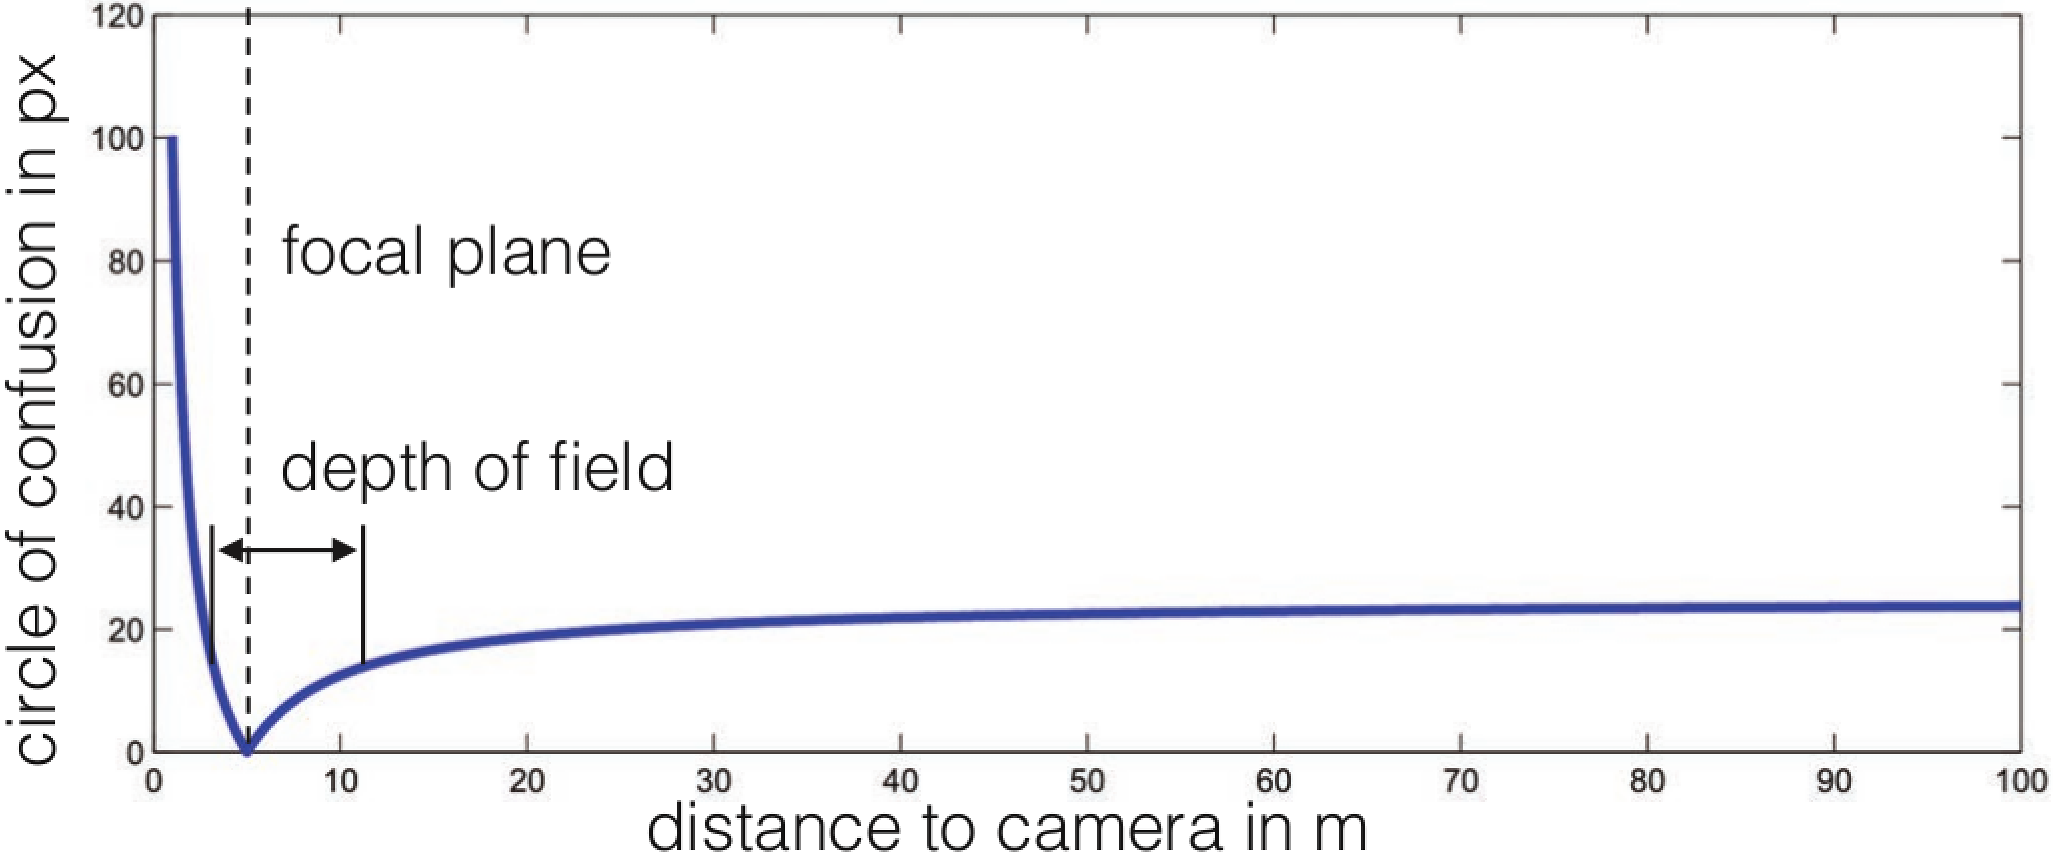

Defocus depends on aperture

What happens to the circle of confusion as the aperture diameter is reduced?

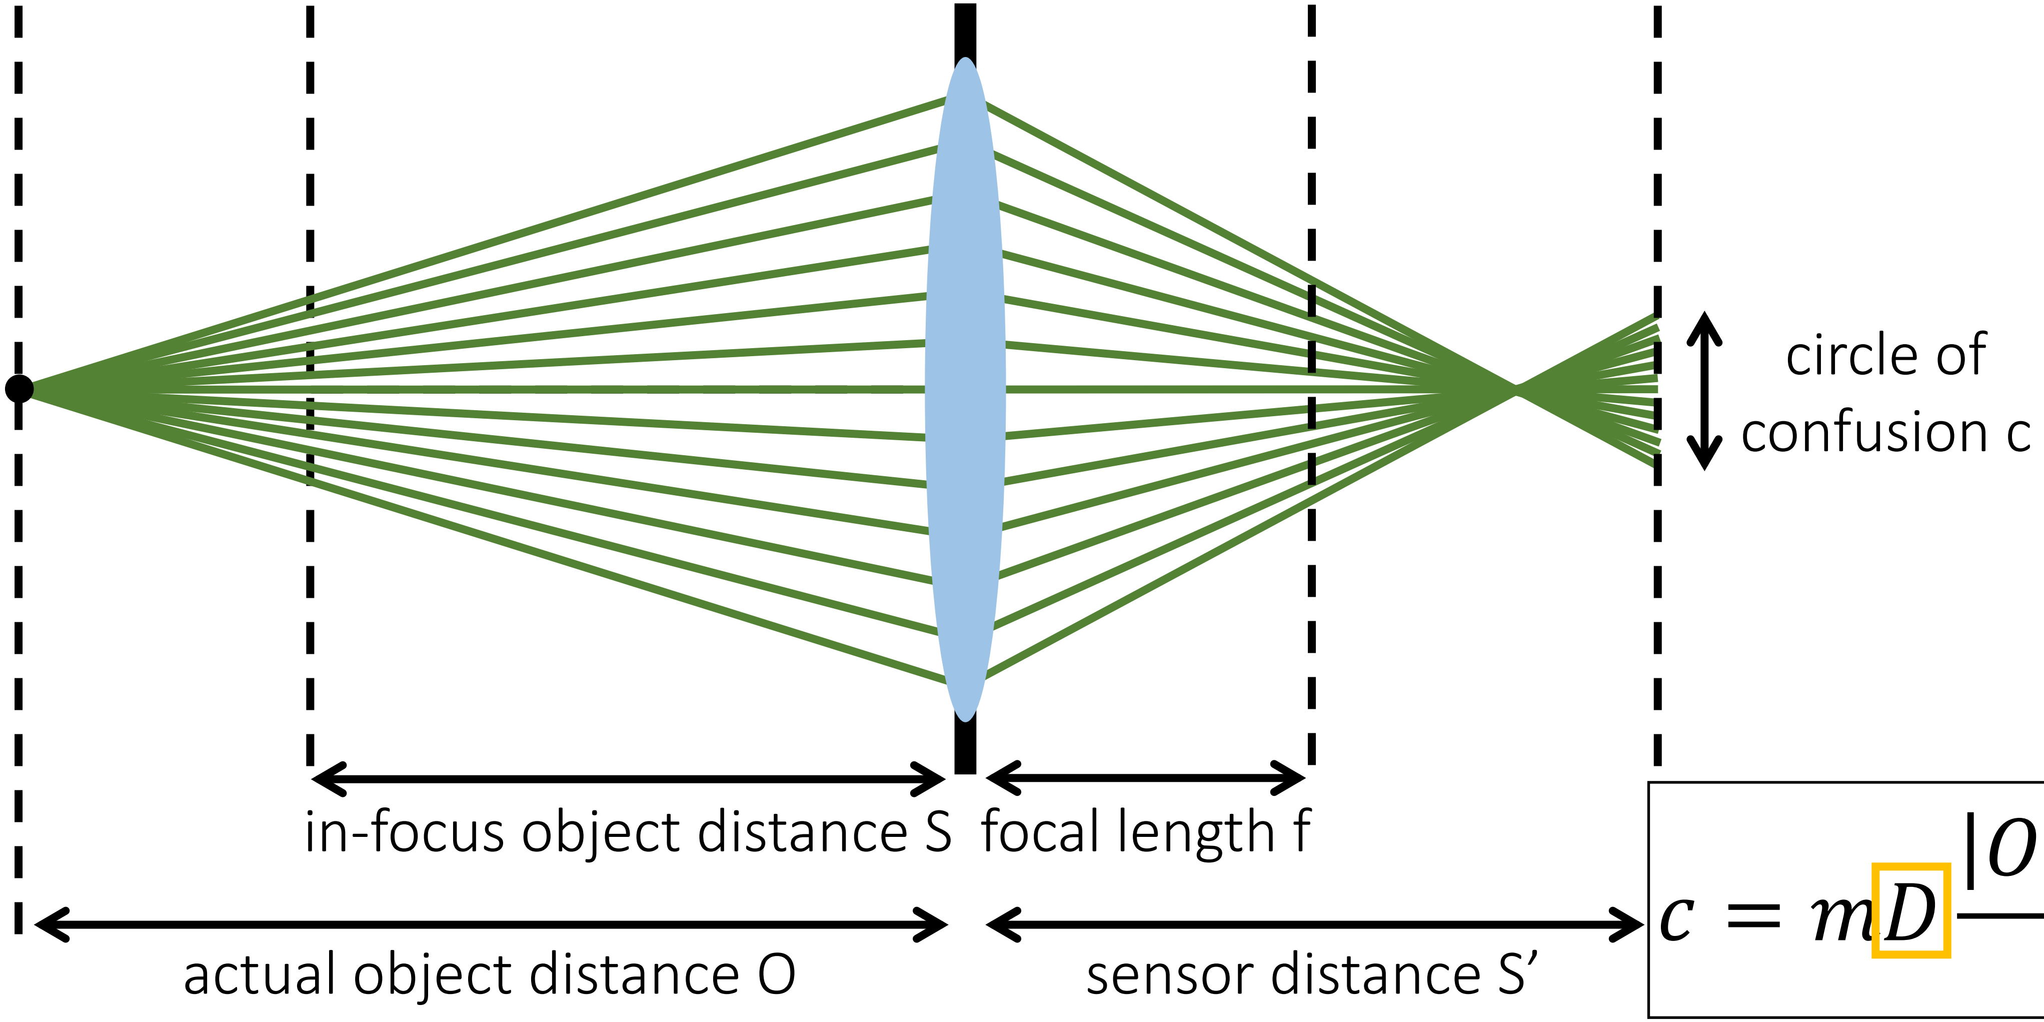

$$c = m\boxed{D}\frac{|O - S|}{O}$$

Defocus depends on aperture diameter

What happens to the circle of confusion as the aperture diameter is reduced? → It shrinks.

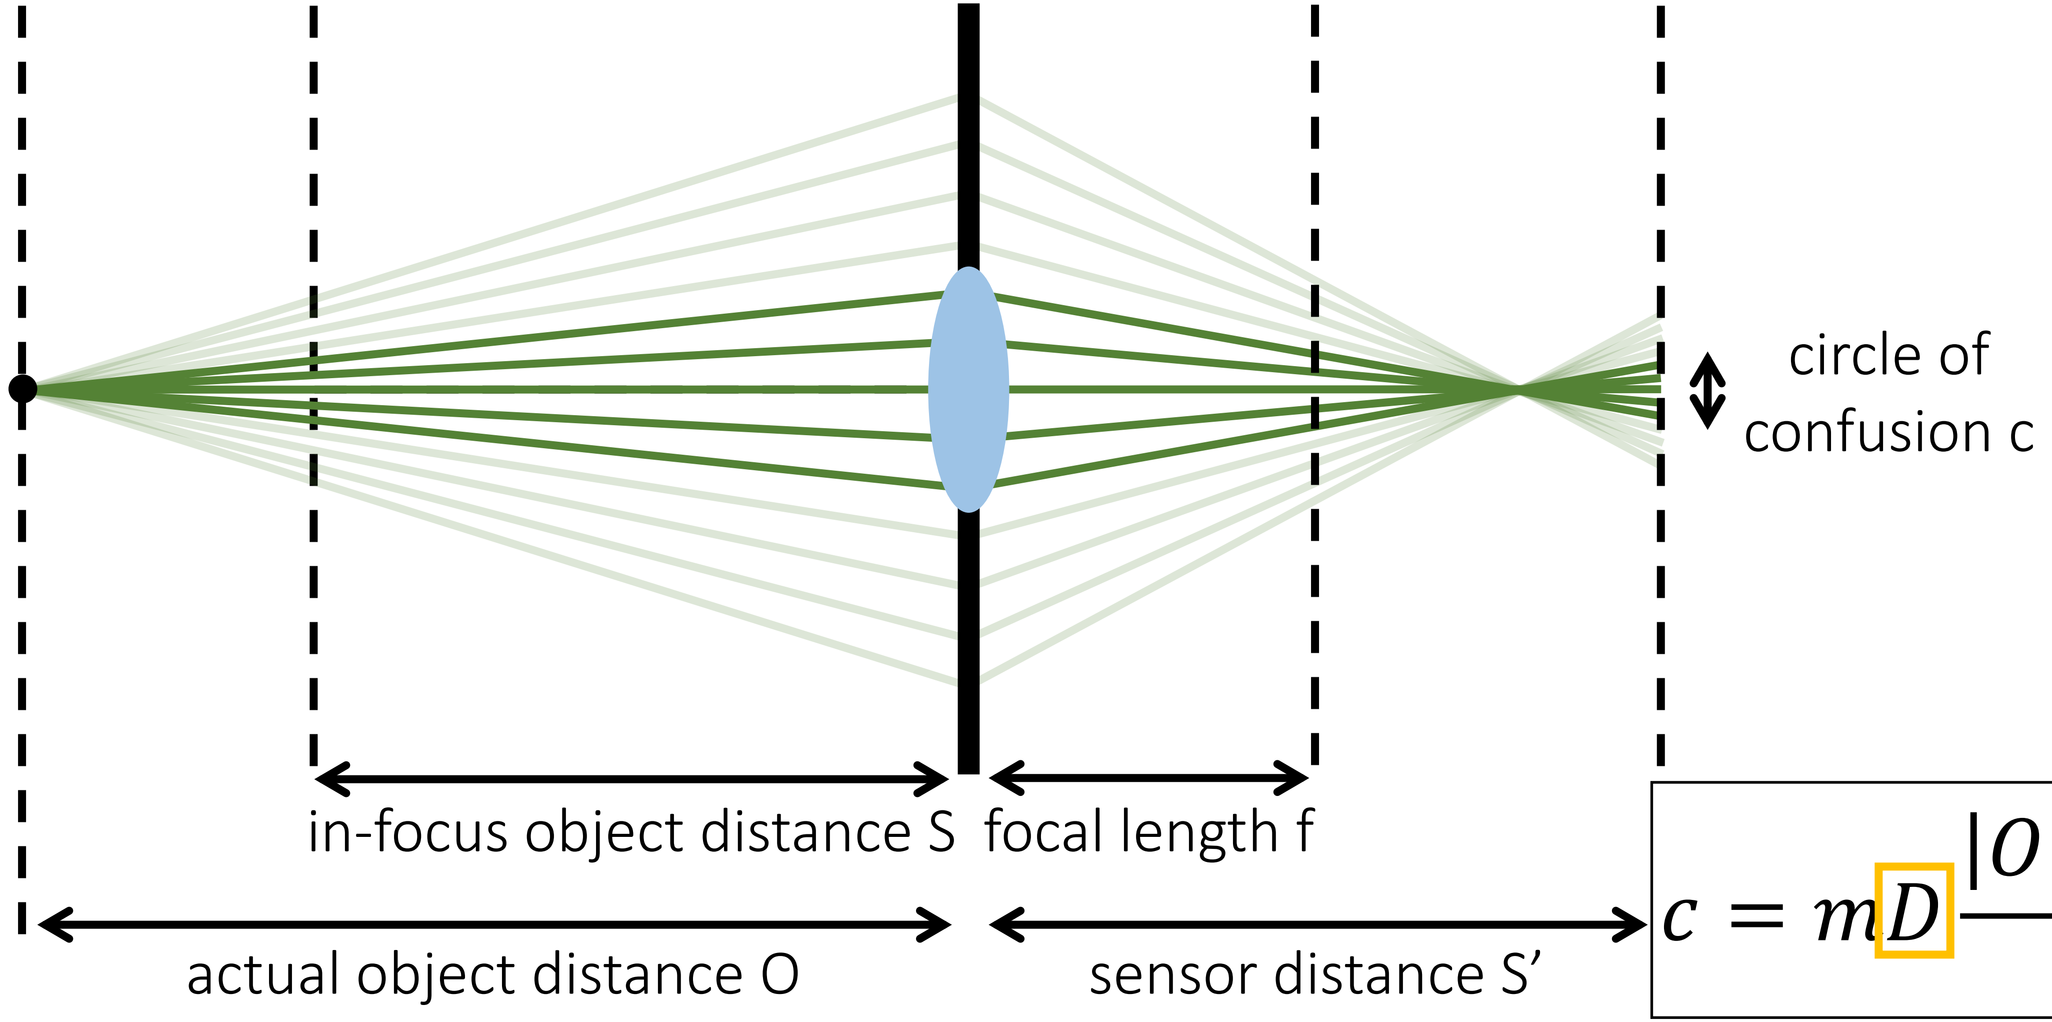

$$c = m\boxed{D}\frac{|O - S|}{O}$$

Defocus depends on aperture diameter

What happens to the depth of field as the aperture diameter is reduced?

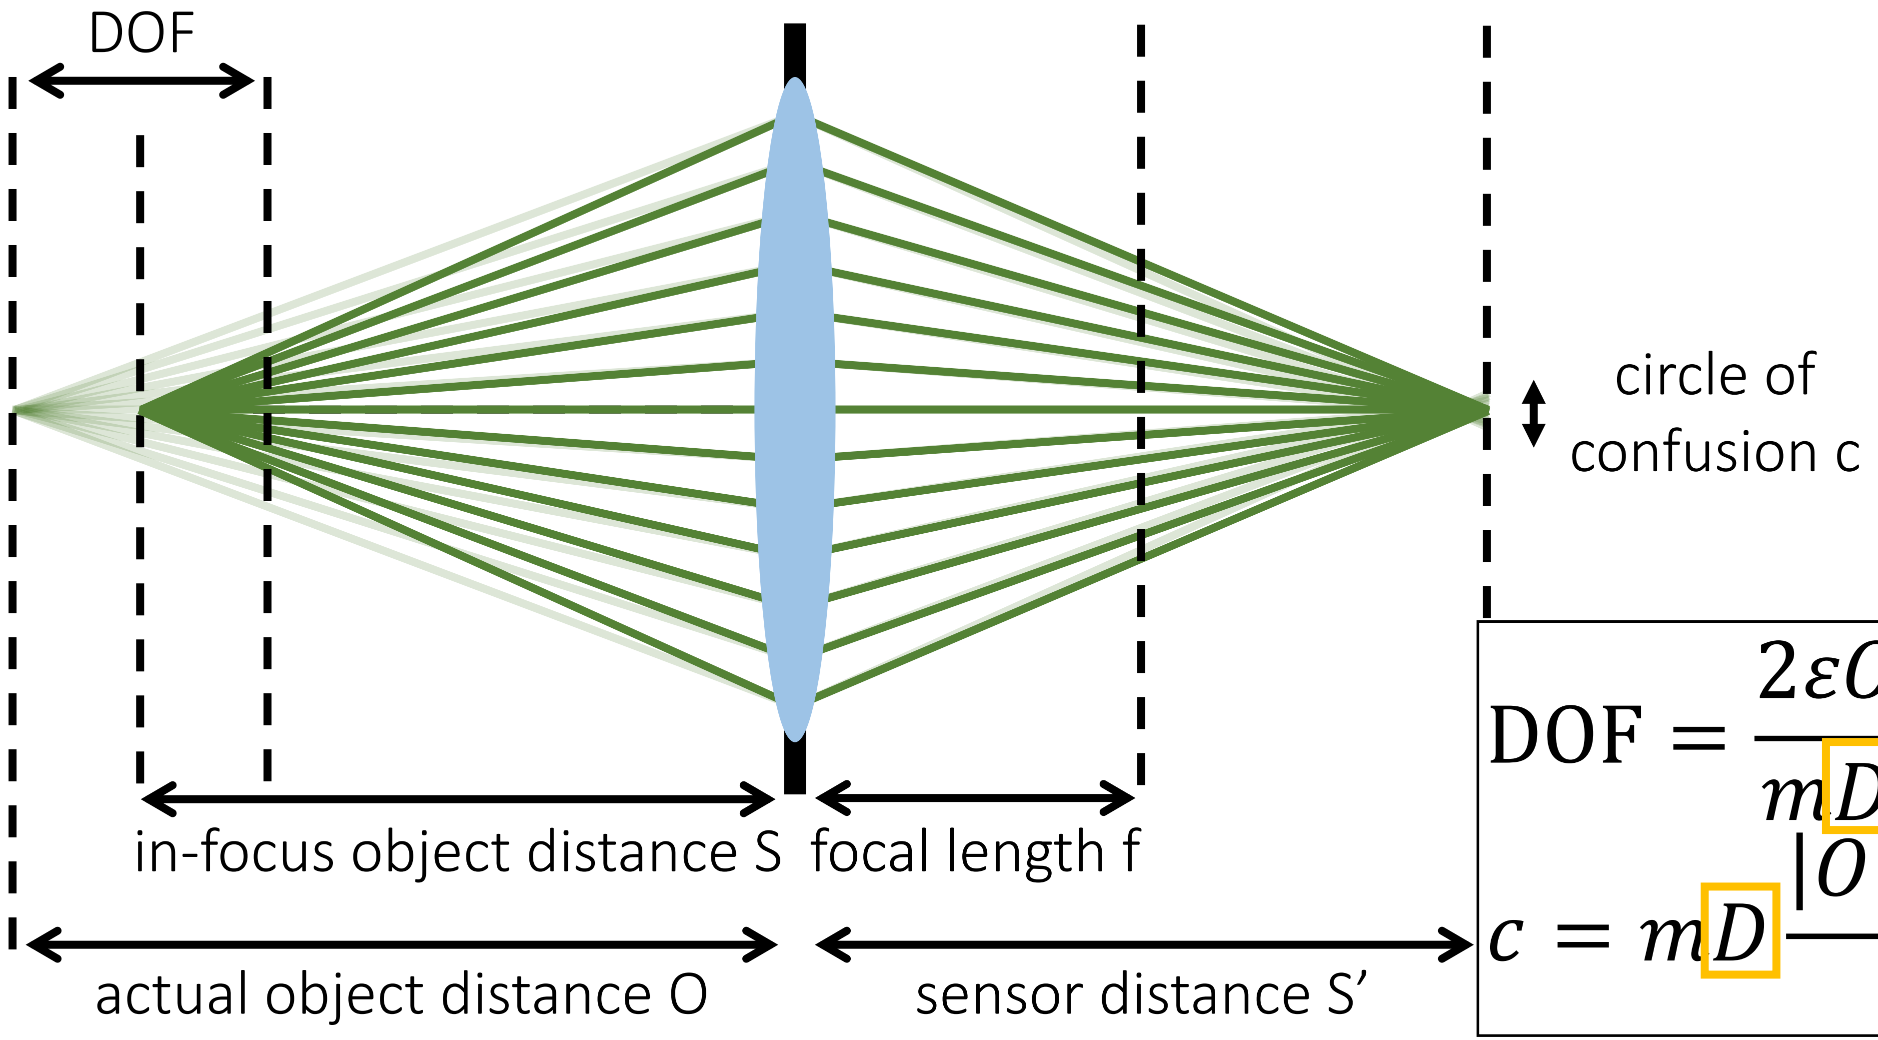

$$\mathrm{DOF} = \frac{2\varepsilon O}{m\boxed{D}}$$

$$c = m\boxed{D}\frac{|O - S|}{O}$$

Defocus depends on aperture diameter

What happens to the depth of field as the aperture diameter is reduced? → It expands.

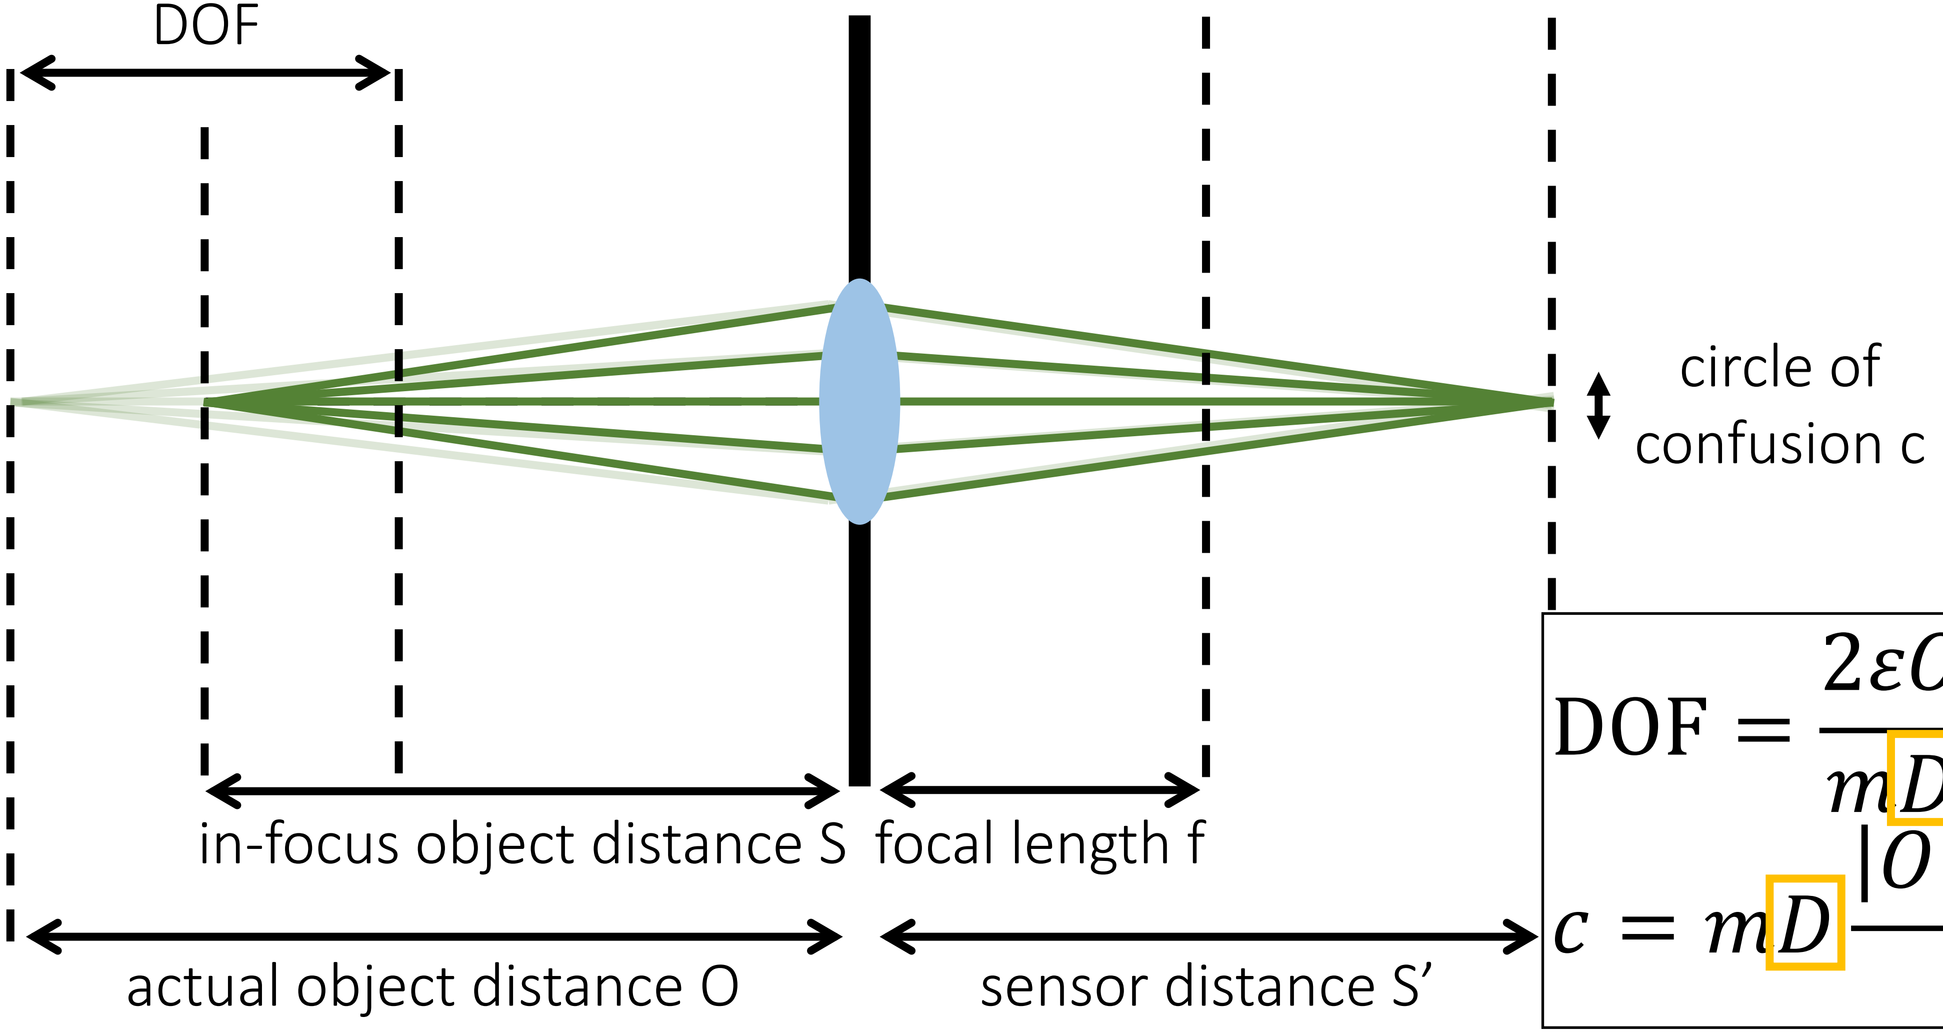

$$\mathrm{DOF} = \frac{2\varepsilon O}{m\boxed{D}}$$

$$c = m\boxed{D}\frac{|O - S|}{O}$$

Aperture size

Most lenses have apertures of variable size.

- The size of the aperture is expressed as the “f-number”: The bigger this number, the smaller the aperture.

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

You can see the aperture by removing the lens and looking inside it.

Demonstration

Depth of field

Form of defocus blur is determined by shape of aperture.

Defocus depends on aperture diameter

If small aperture sizes reduce defocus blur, should we always use the smallest aperture?

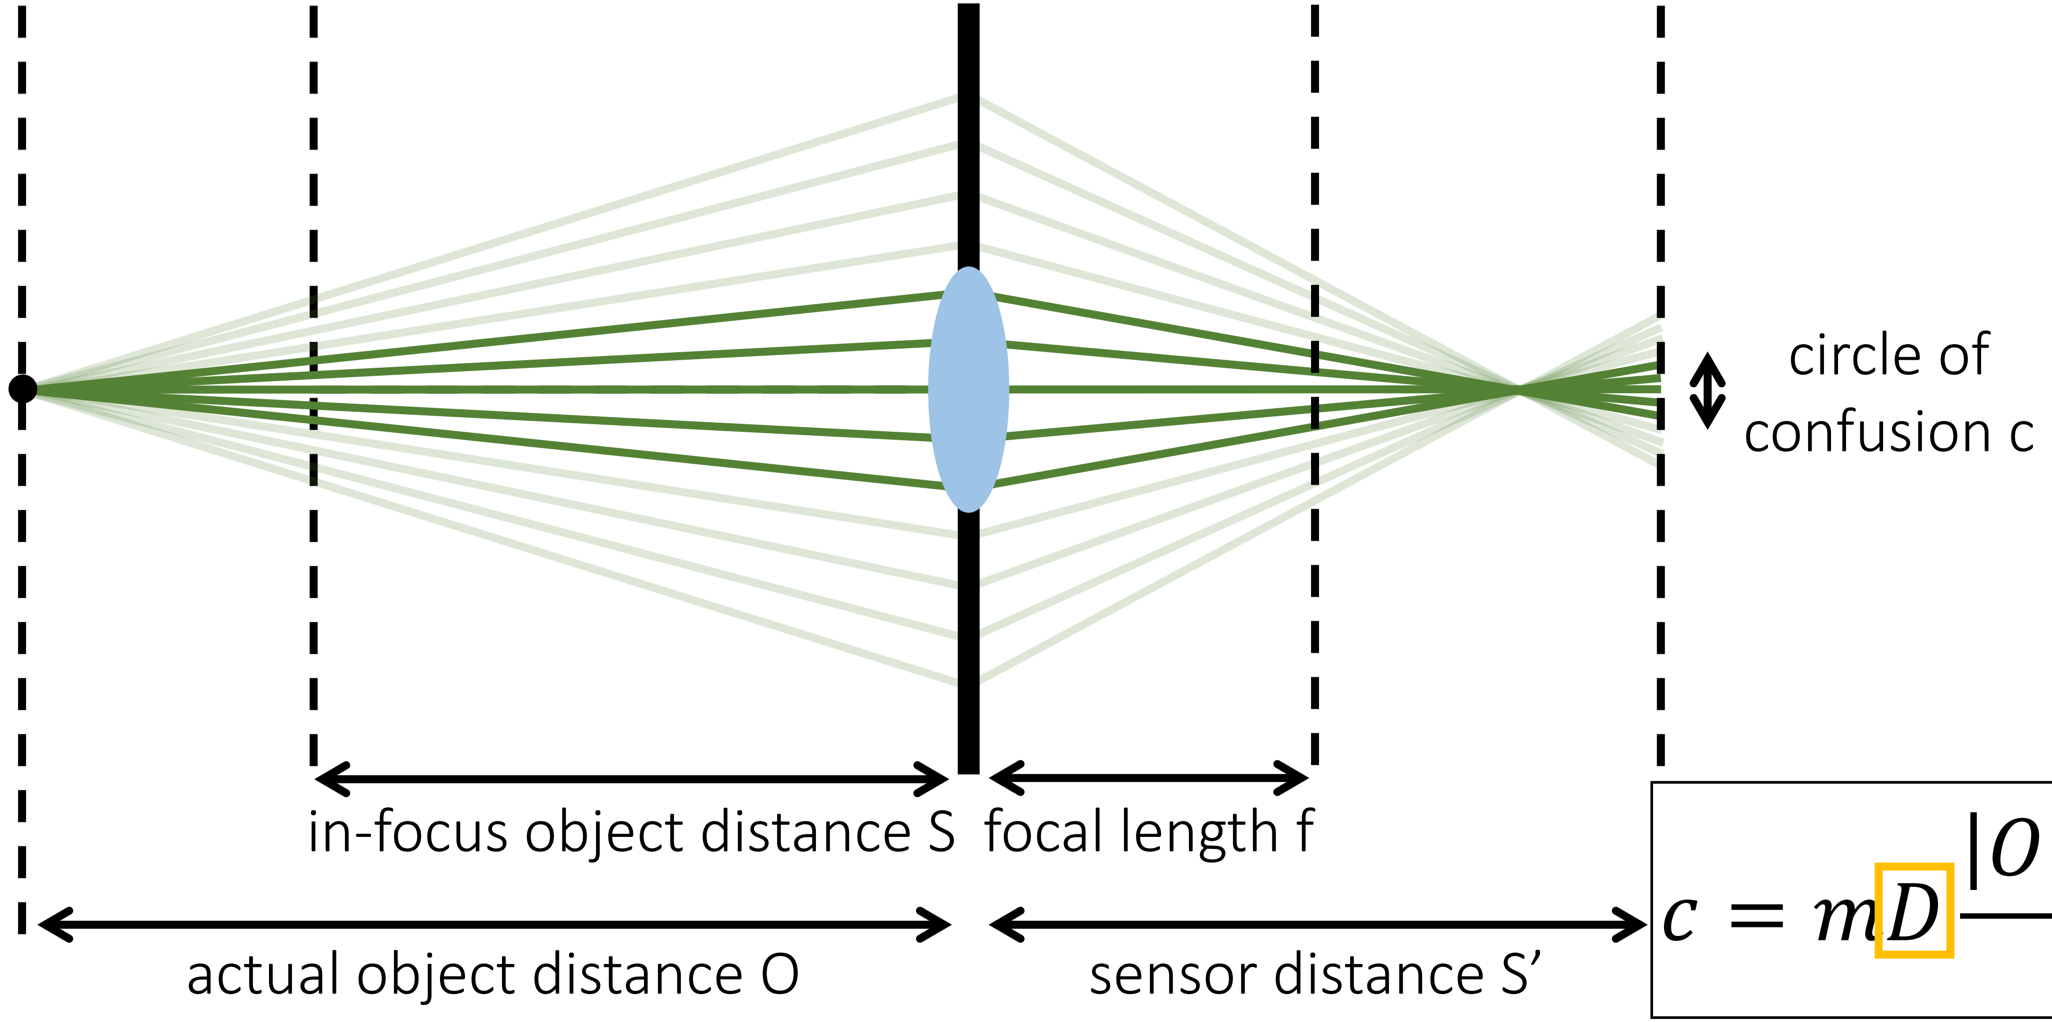

$$c = m\boxed{D}\frac{|O - S|}{O}$$

Bokeh

Sharp depth of field ("bokeh") is often desirable.

What happens as the aperture keeps getting smaller?

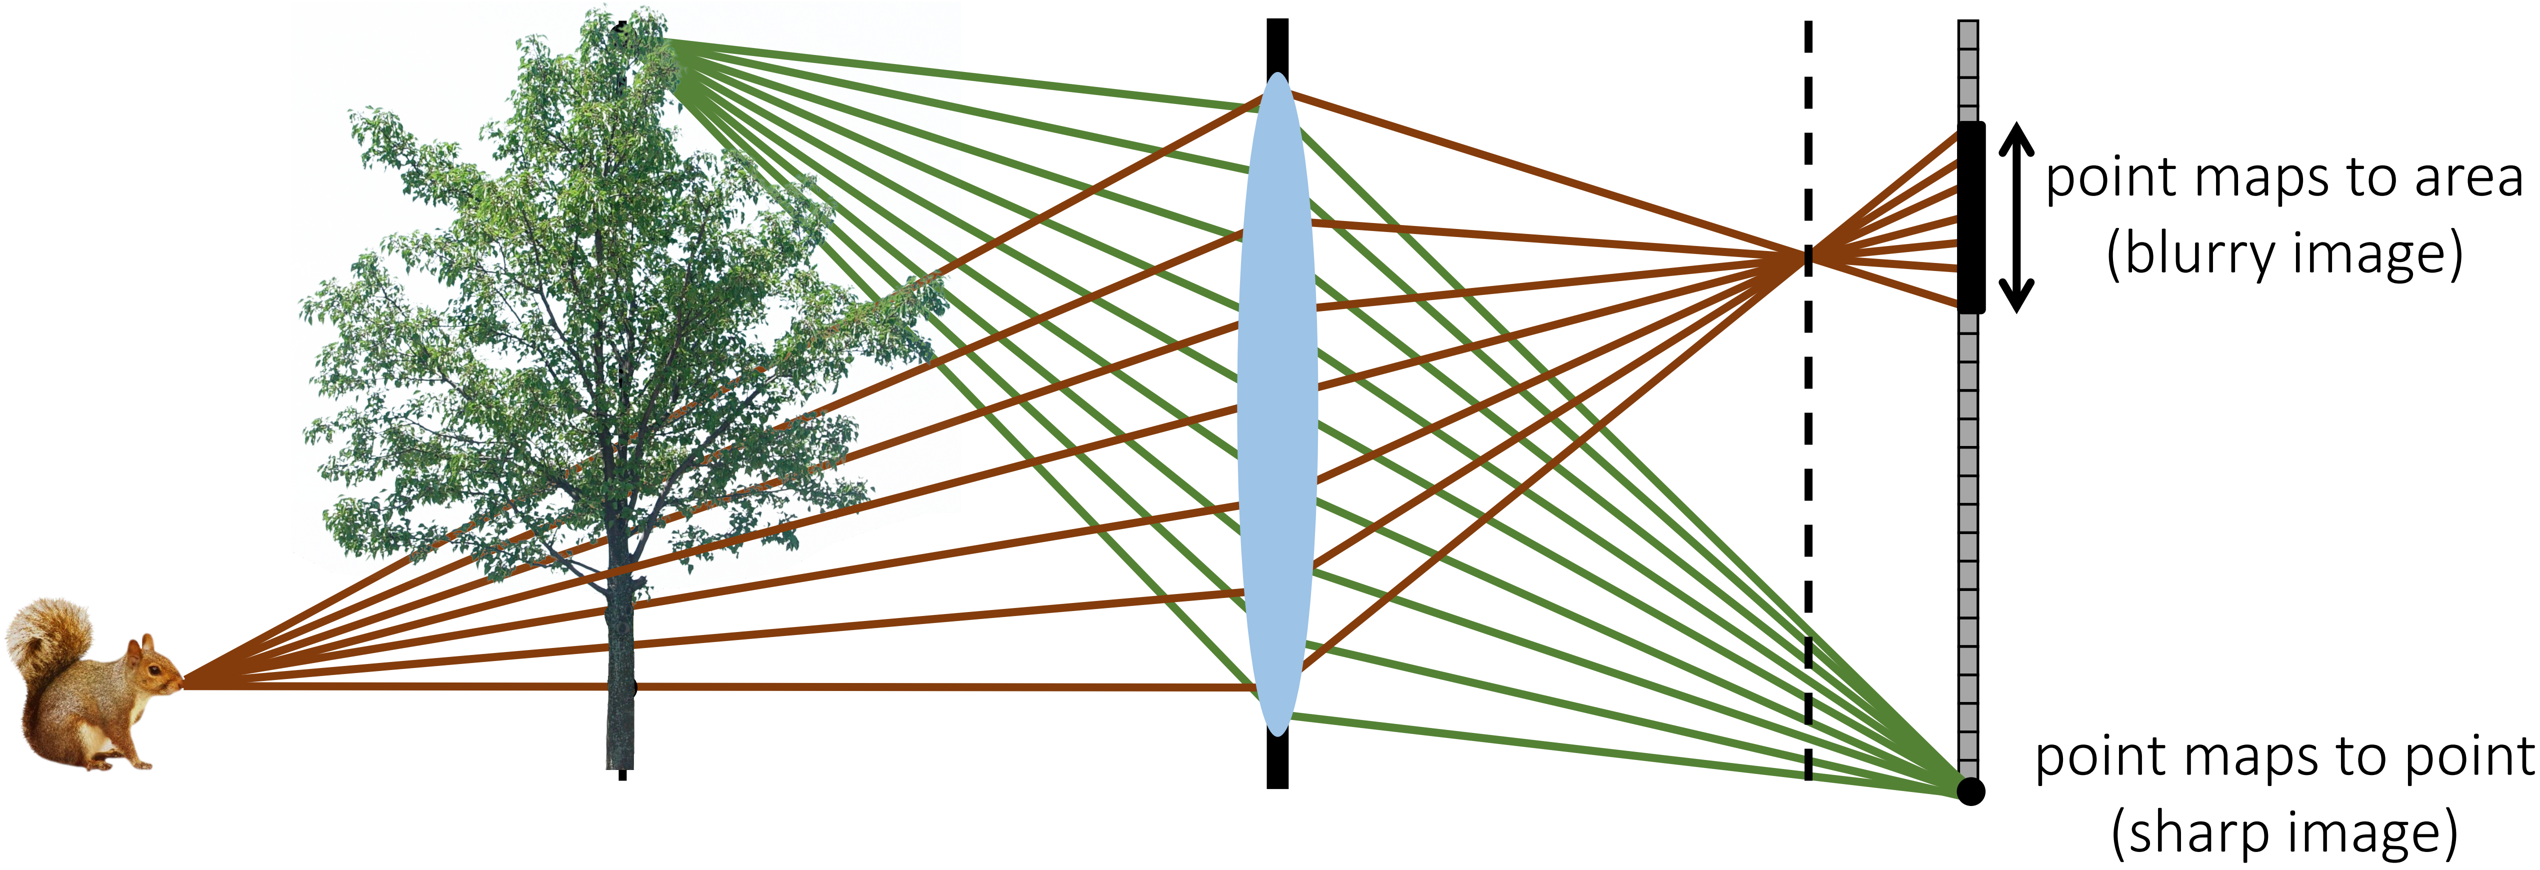

What happens as the aperture keeps getting smaller?

Lens becomes equivalent to a pinhole.

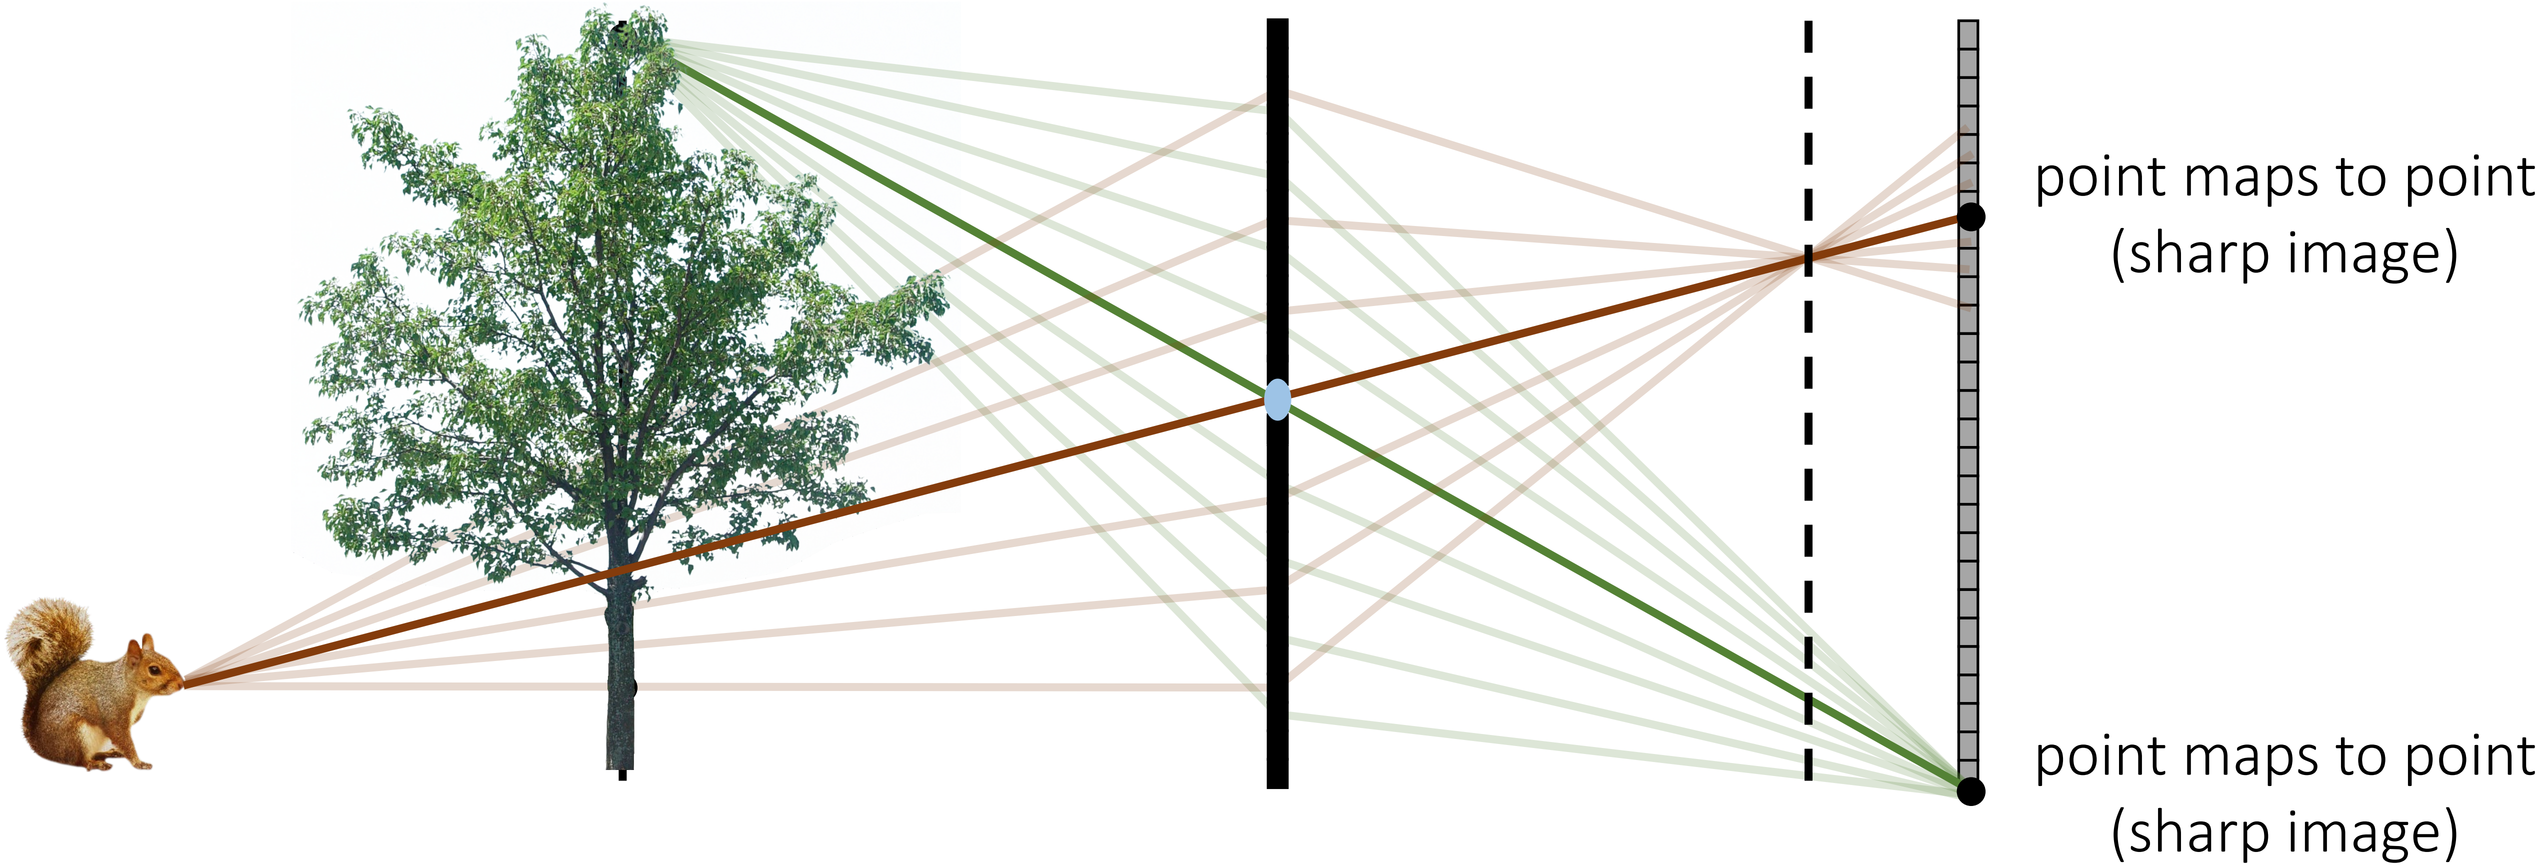

- No defocus, everything is sharp regardless of depth.
- Very little light, signal-to-noise ratio is just as bad as pinhole.

Lens camera and pinhole camera

The lens camera

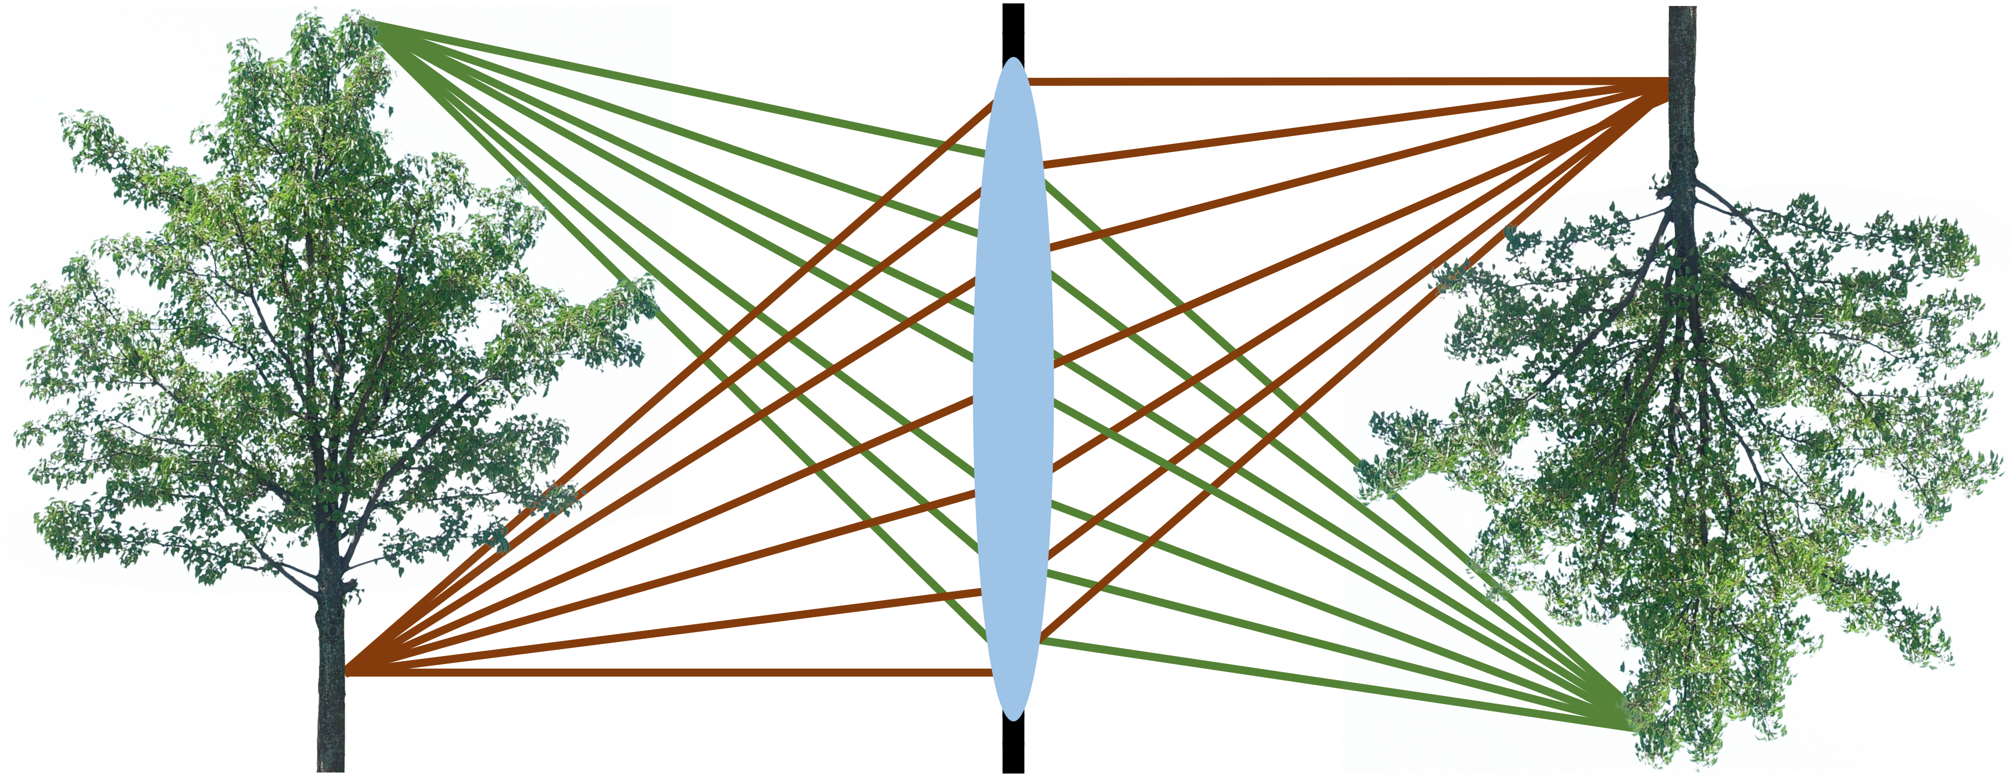

The pinhole camera

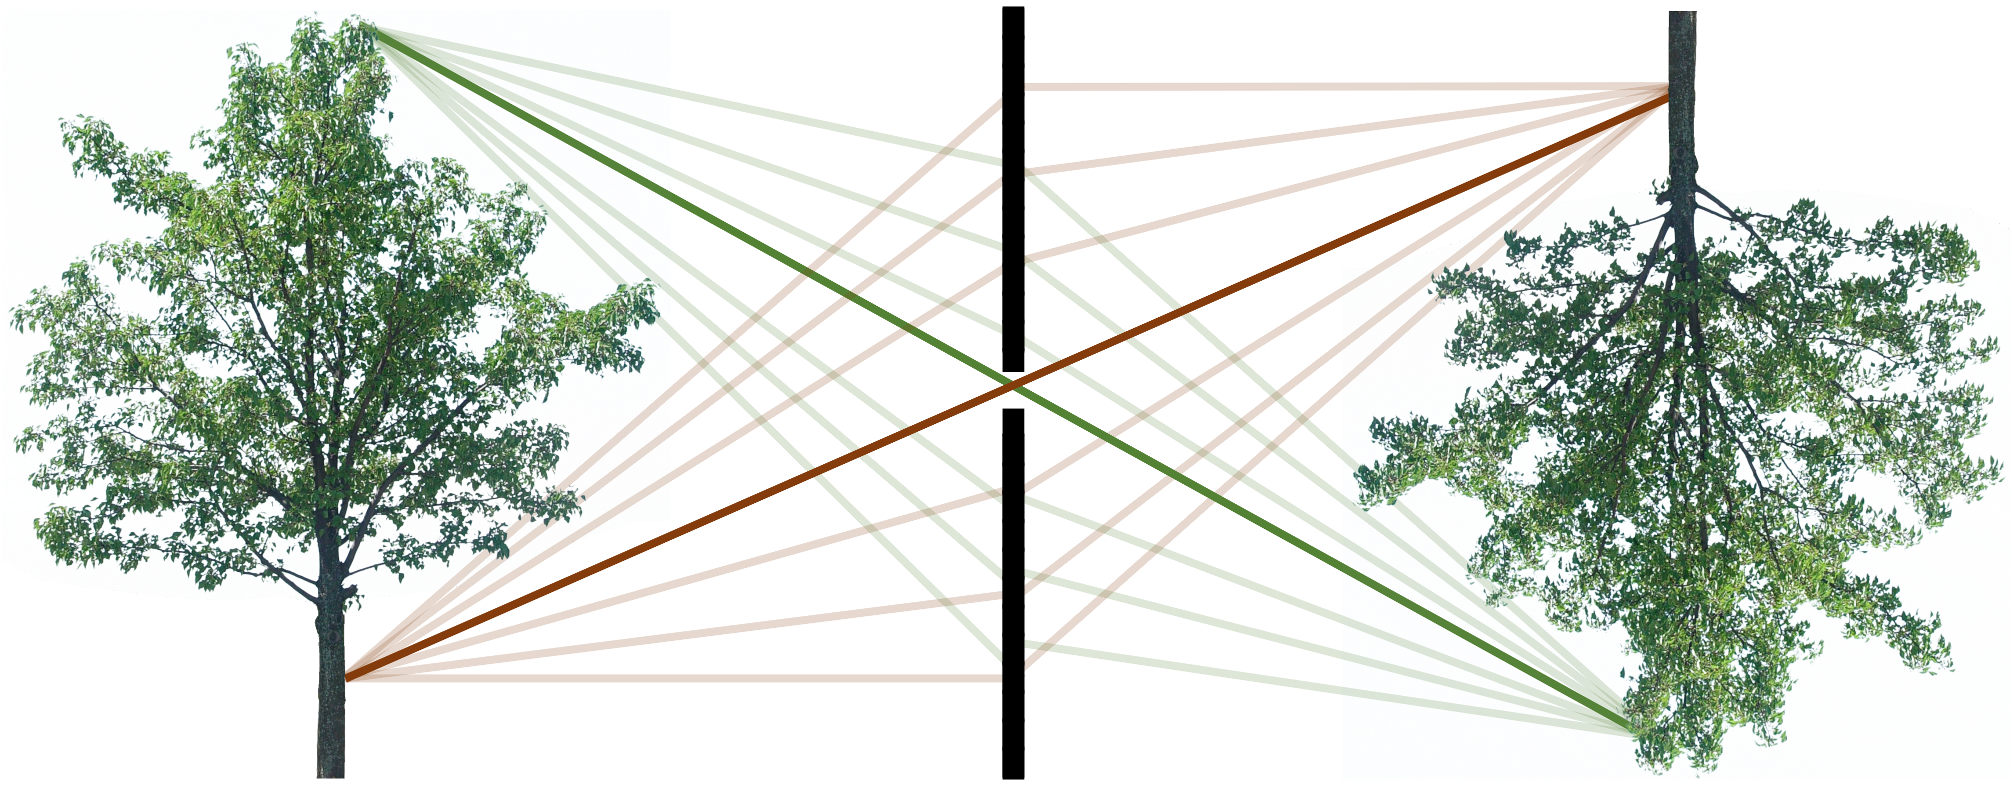

The pinhole camera

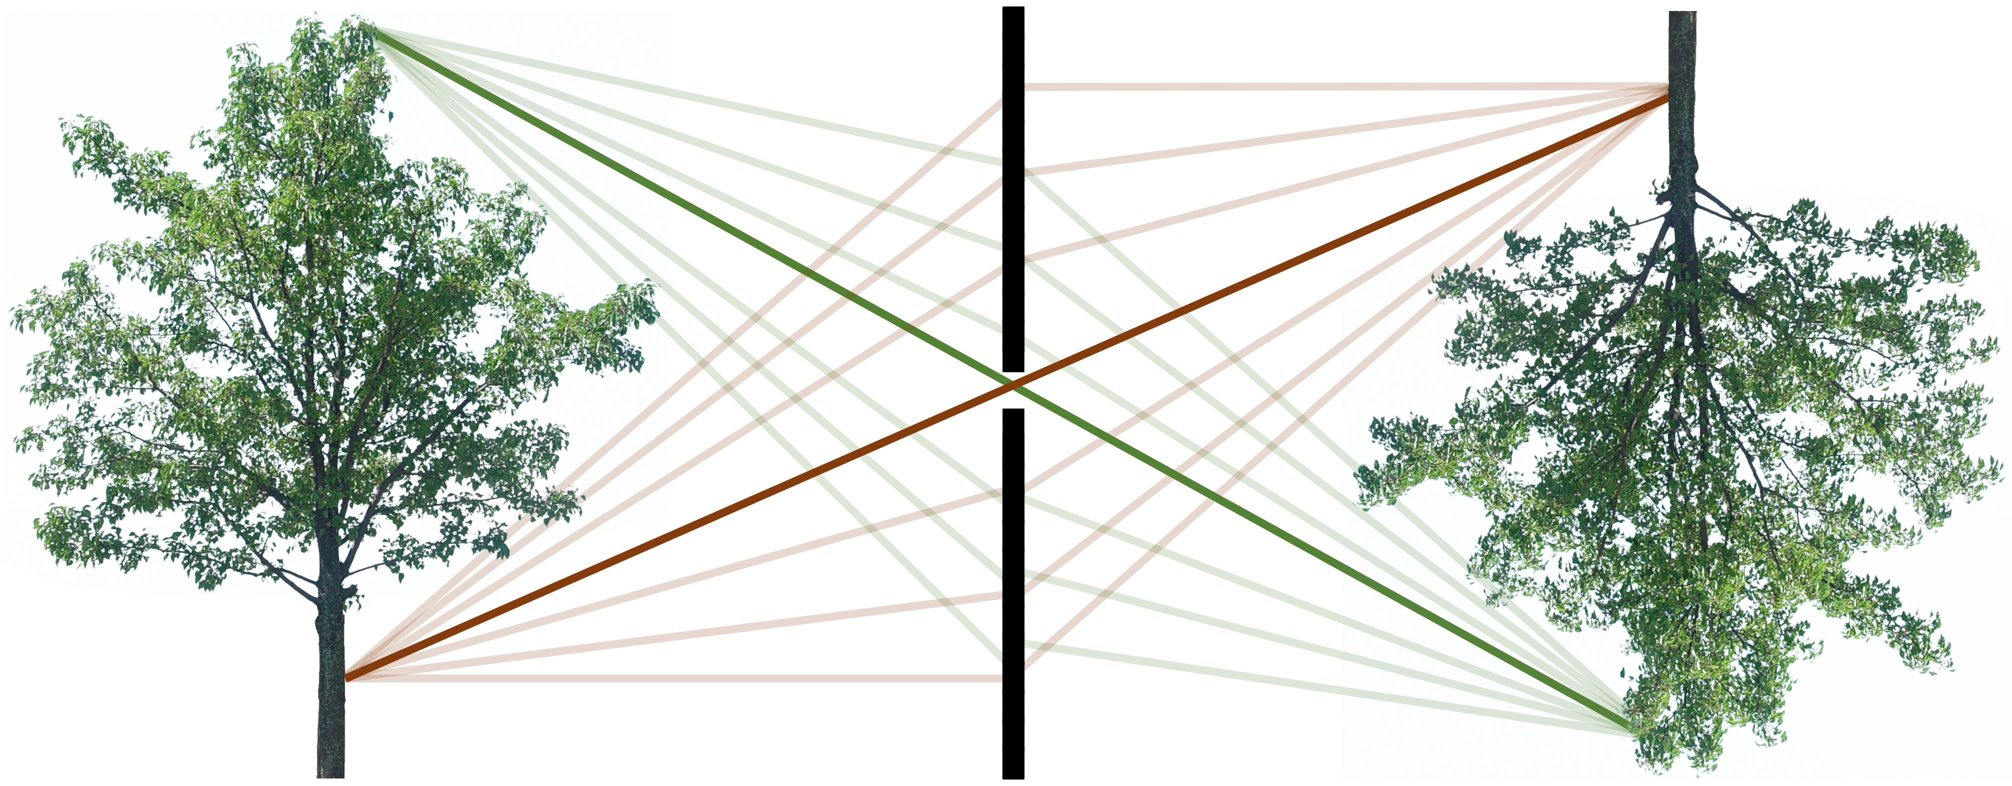

Central rays propagate in the same way for both models!

Describing both lens and pinhole cameras

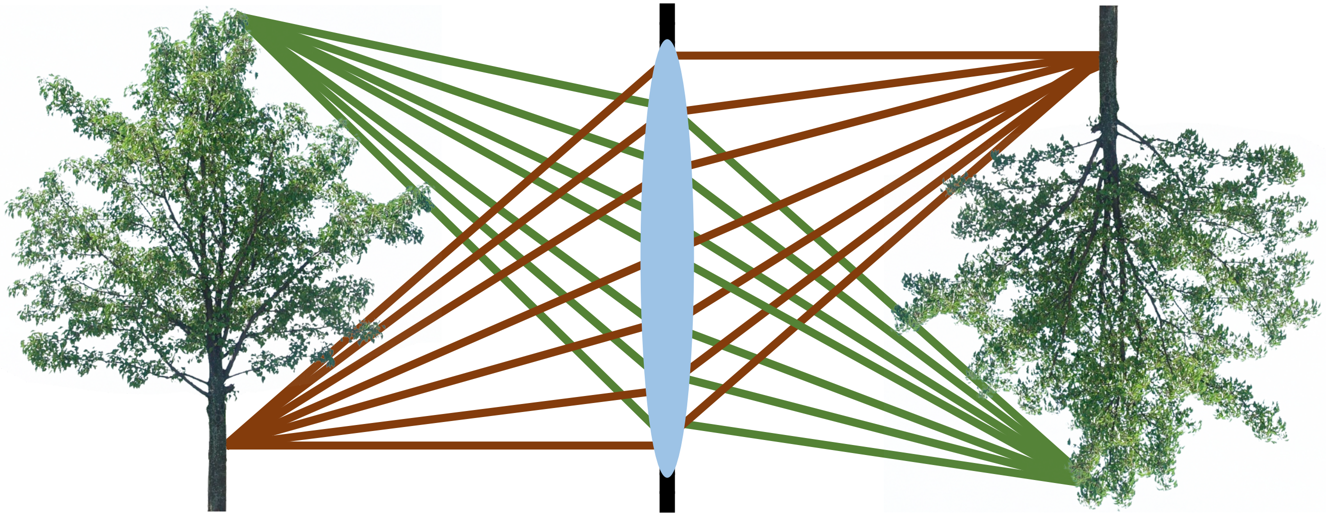

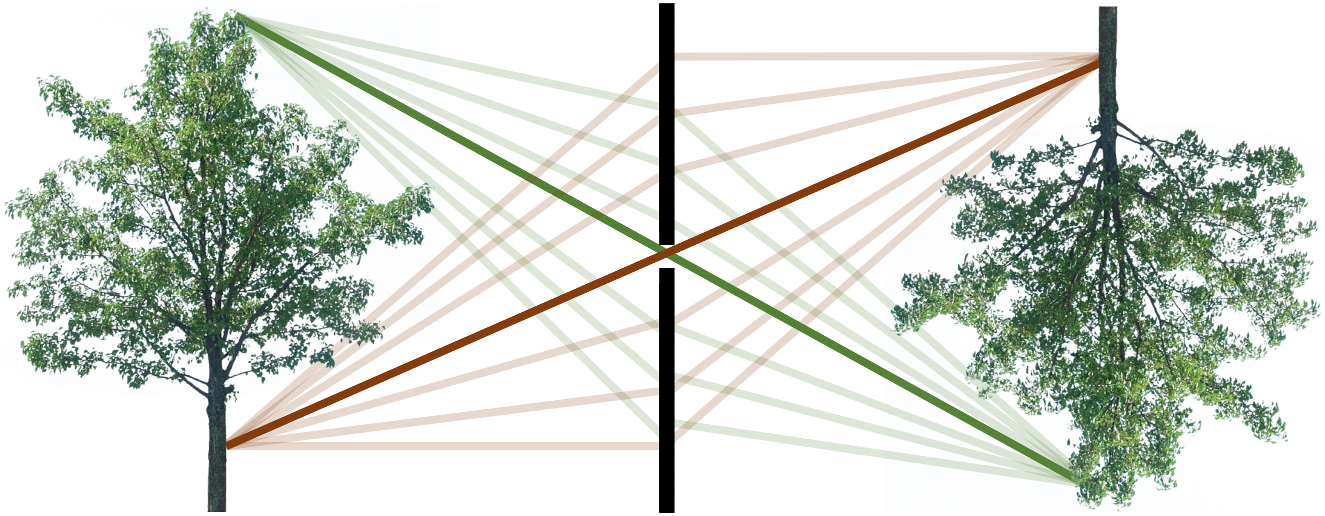

We can derive properties and descriptions that hold for both camera models if:

- We consider only central rays.
- We assume that everything of interest in the scene is within the depth of field.

Important difference: focal length

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

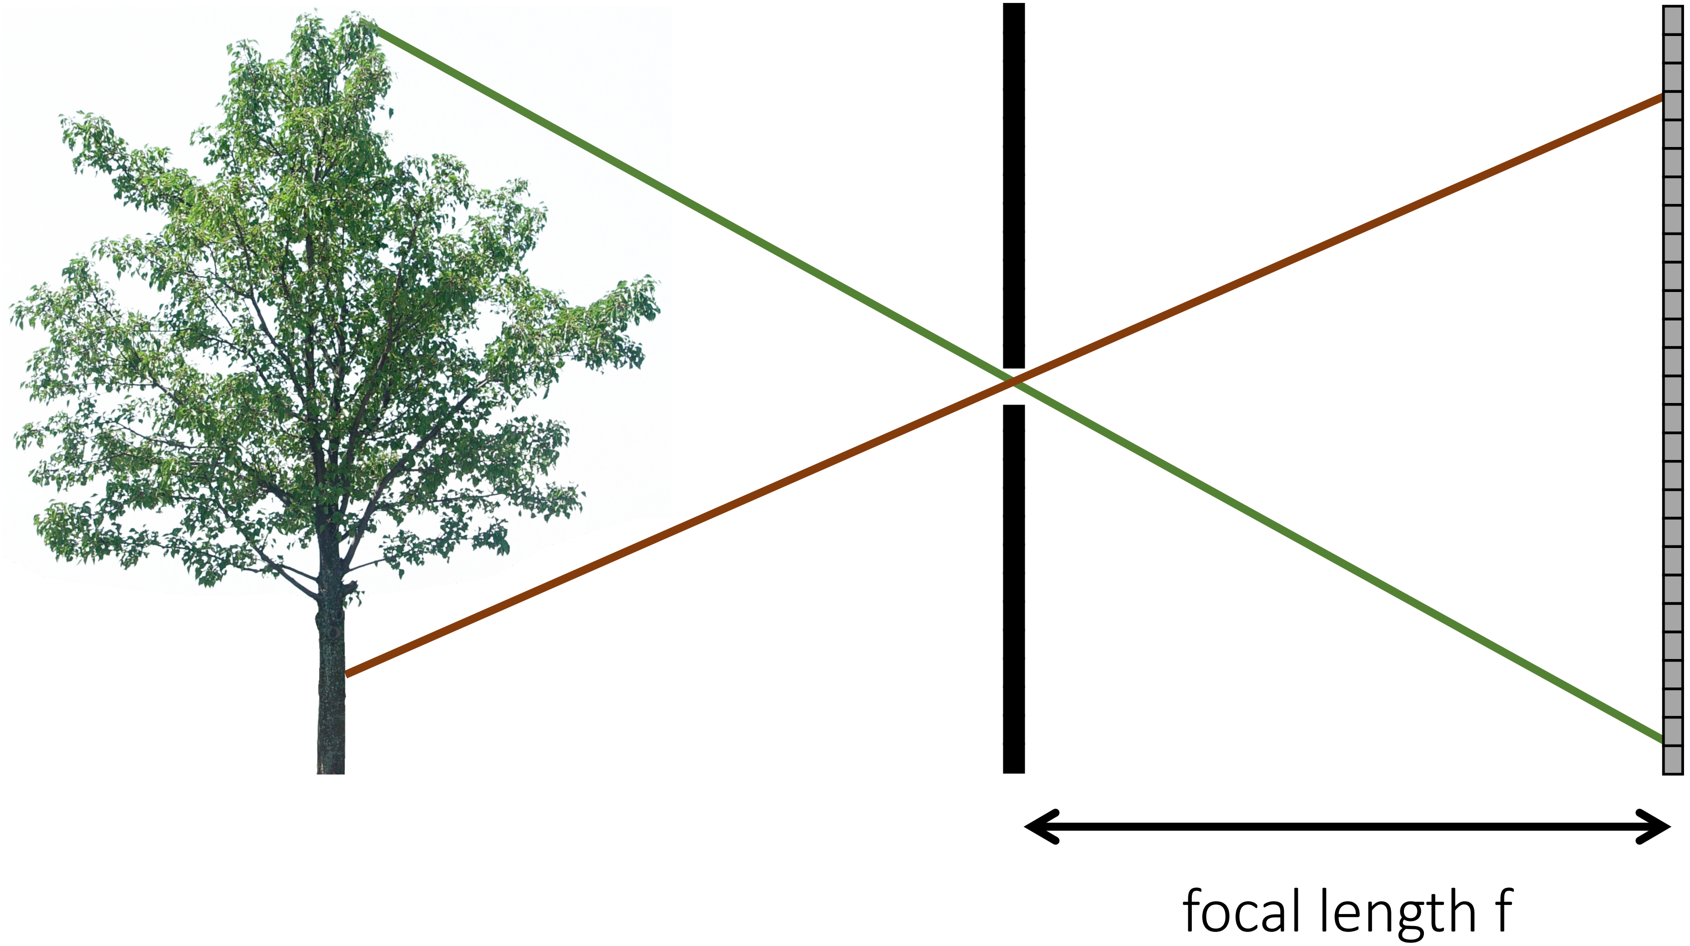

Important difference: focal length

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

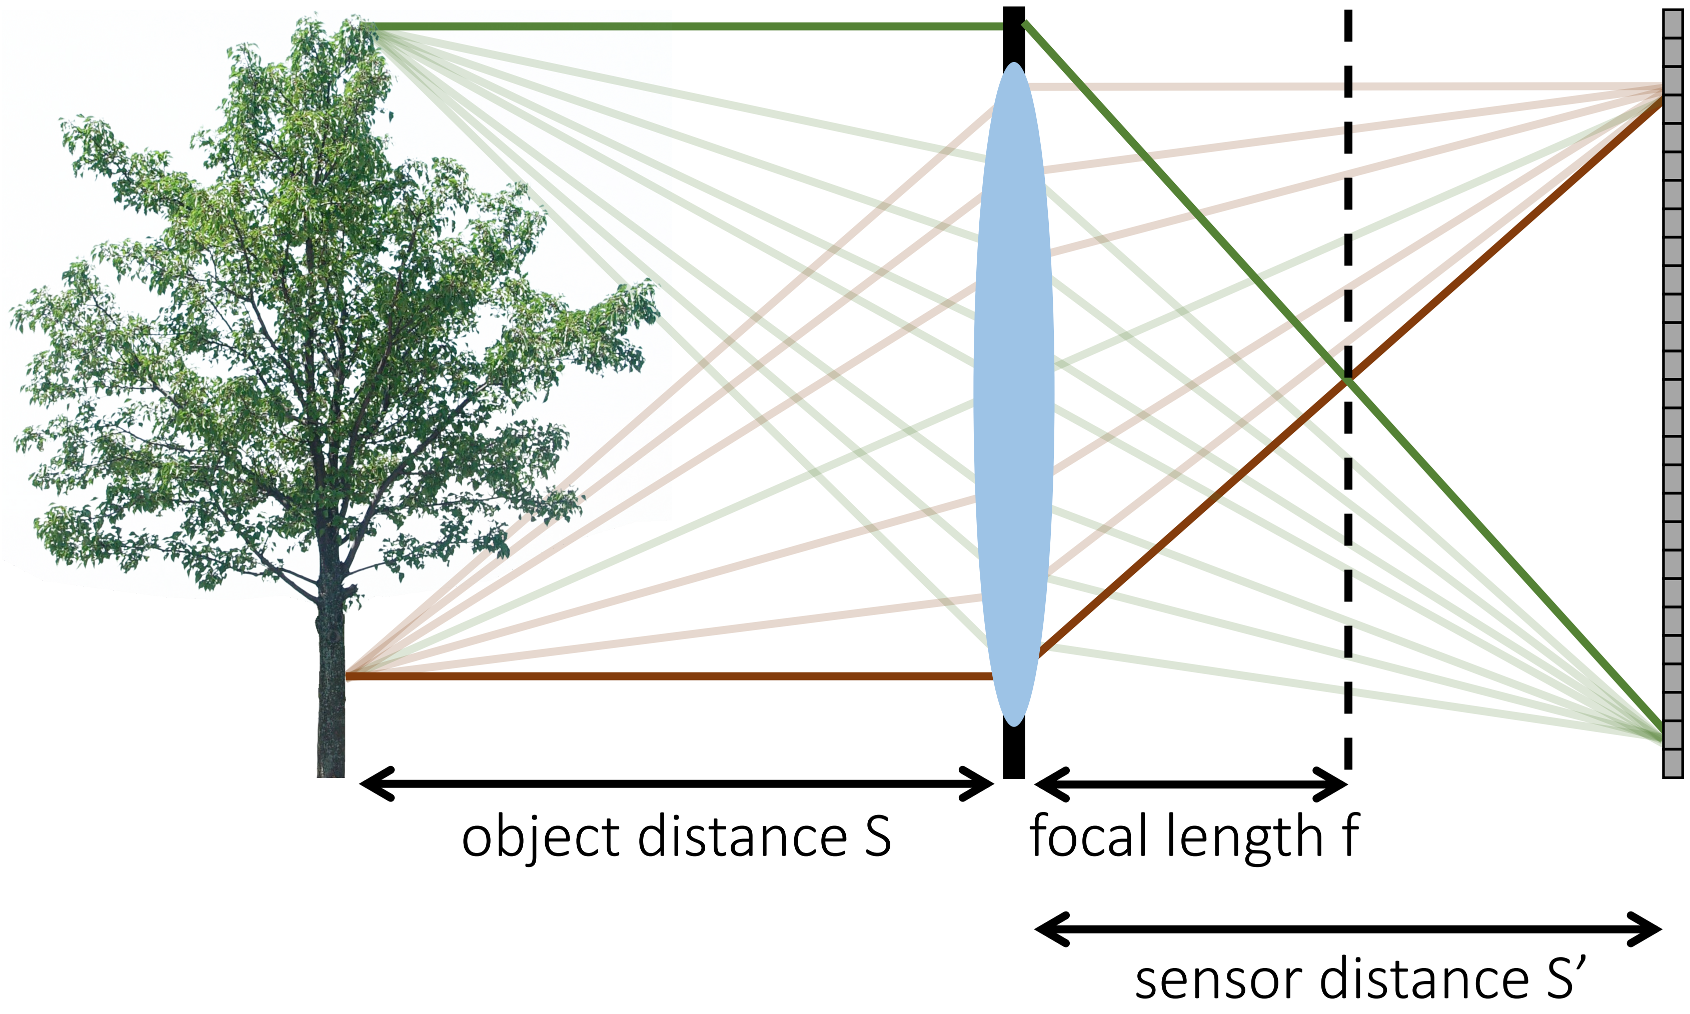

Describing both lens and pinhole cameras

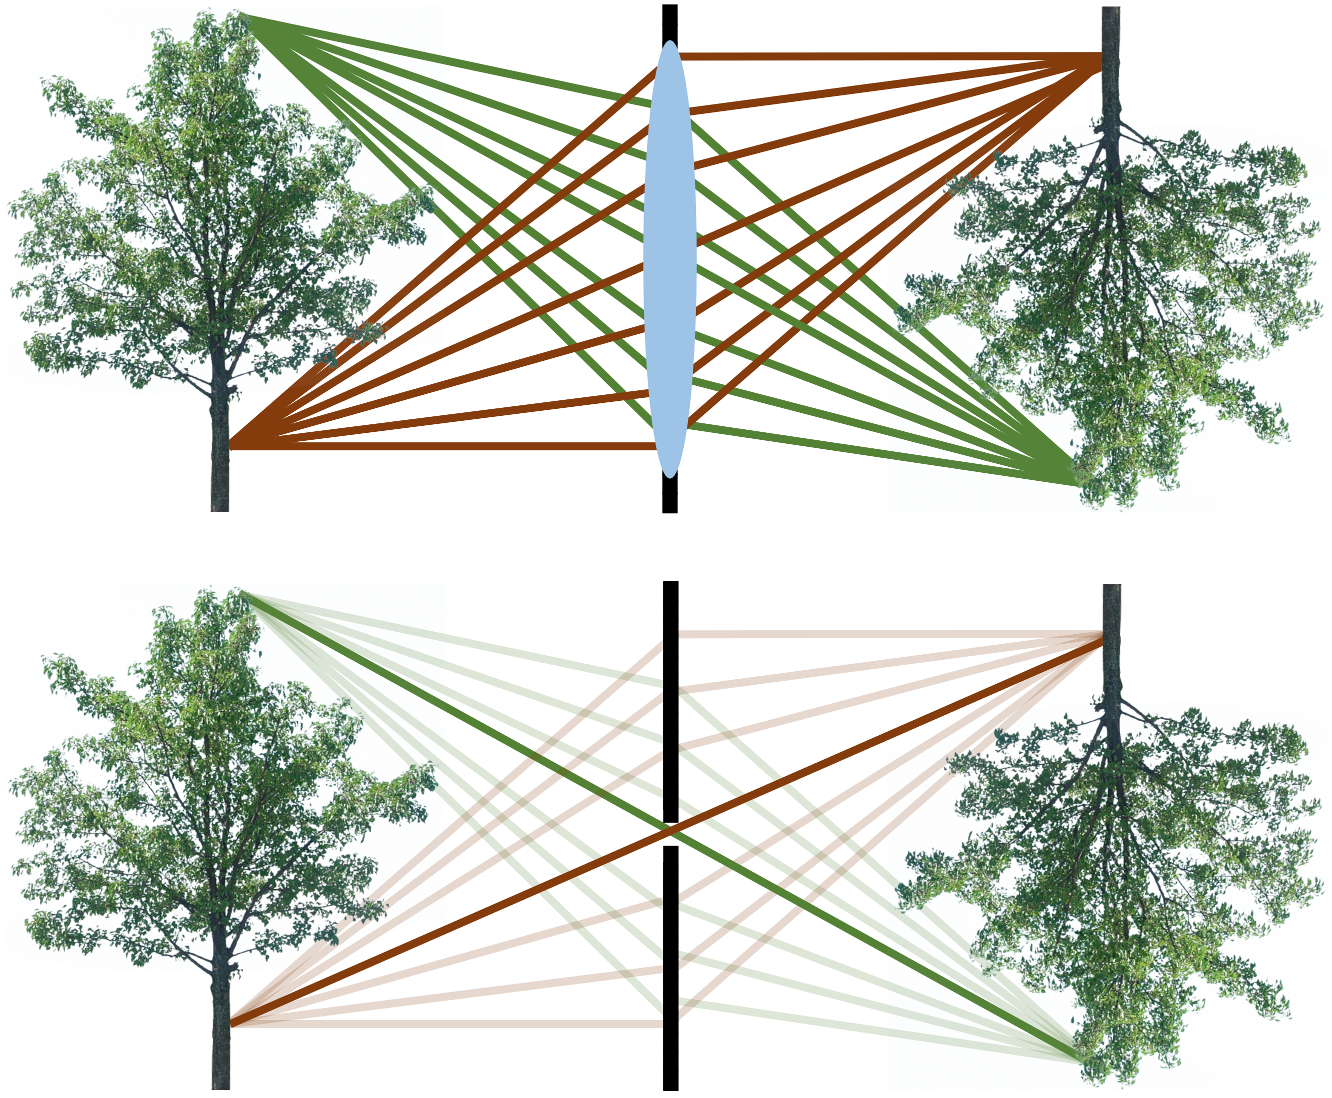

We can derive properties and descriptions that hold for both camera models if:

- We consider only central rays.
- We assume everything of interest in the scene is within the depth of field.
- We assume that the focus distance of the lens camera is equal to the focal length of the pinhole camera.

Effect of aperture size on lens and pinhole cameras

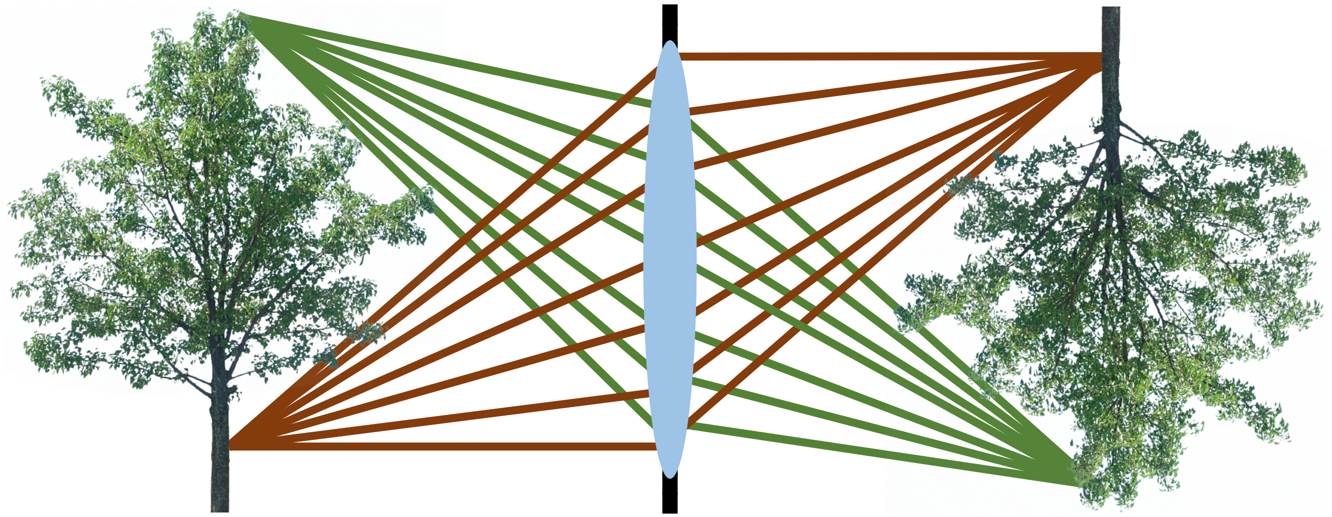

Doubling the aperture diameter:
- Increases light throughput by four times.
- Increases circle of confusion for out-of-focus plane by two times.
- Decreases depth of field by two times.

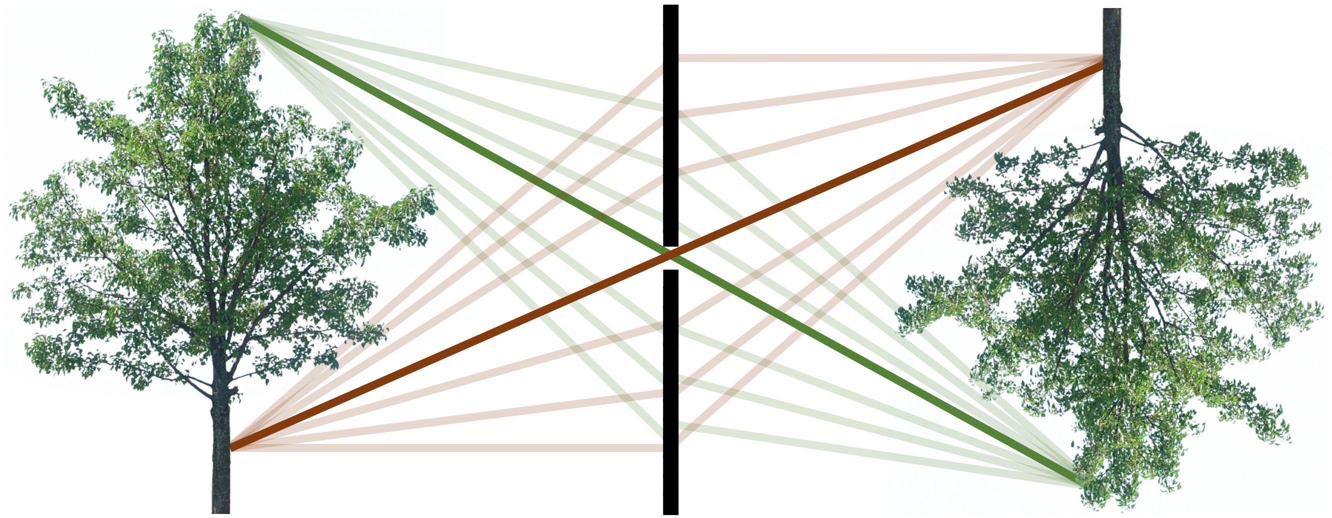

Doubling the aperture diameter:
- Increases light throughput by four times.
- Increases circle of confusion for all planes by two times.

Field of view

Field of view

also described using angle of view

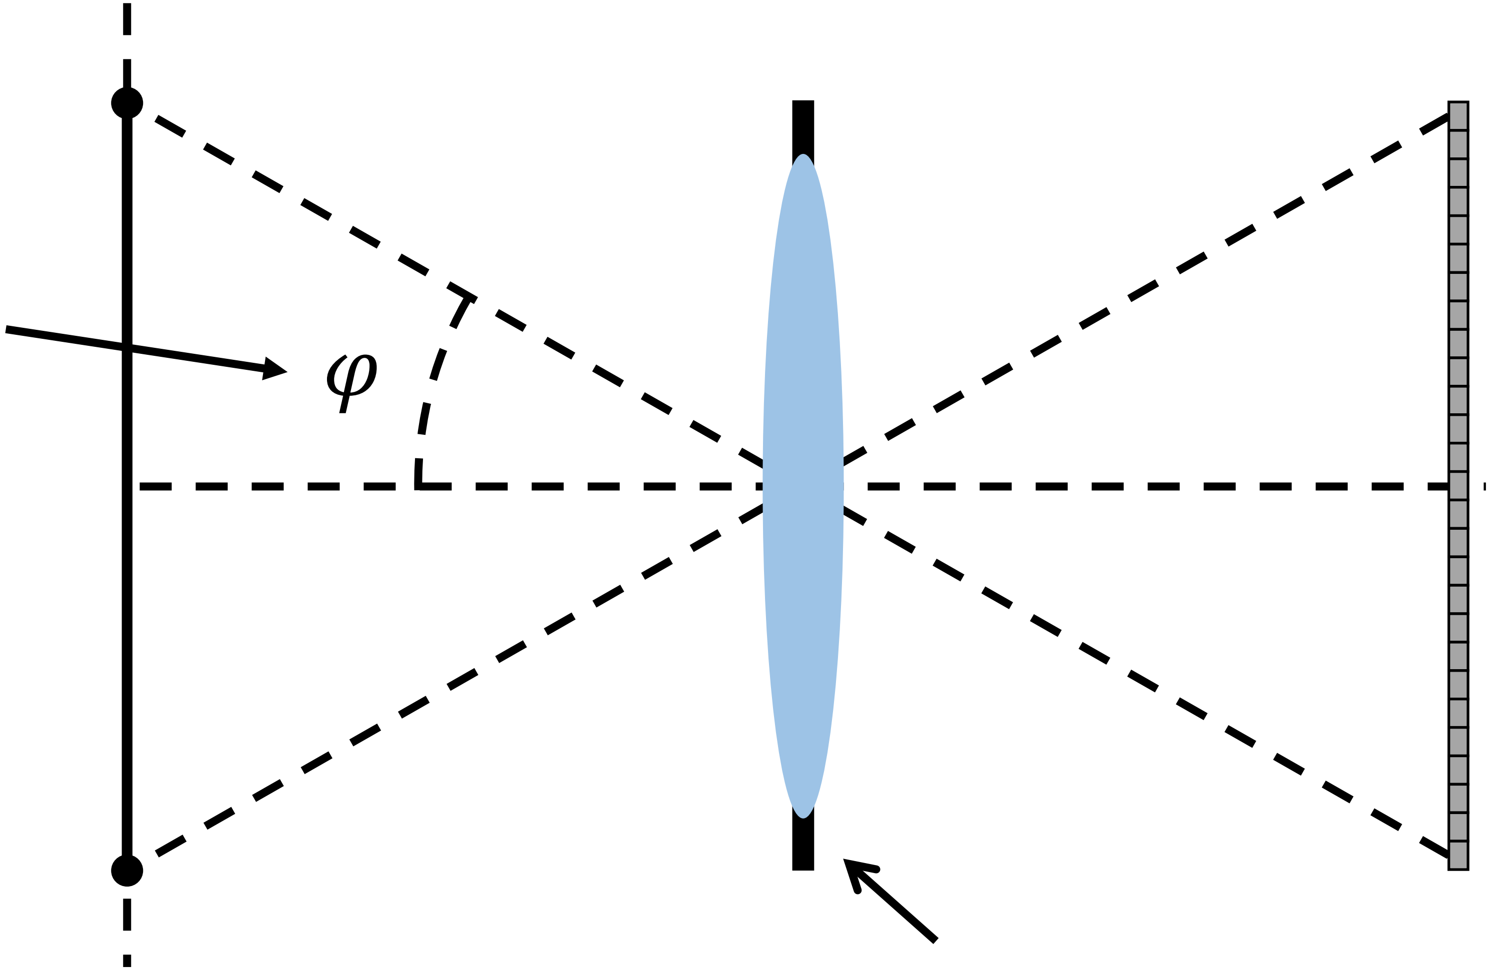

Note: here I drew a lens, but I could have just as well drawn a pinhole

Field of view

The part of the in-focus plane that gets mapped on the sensor.

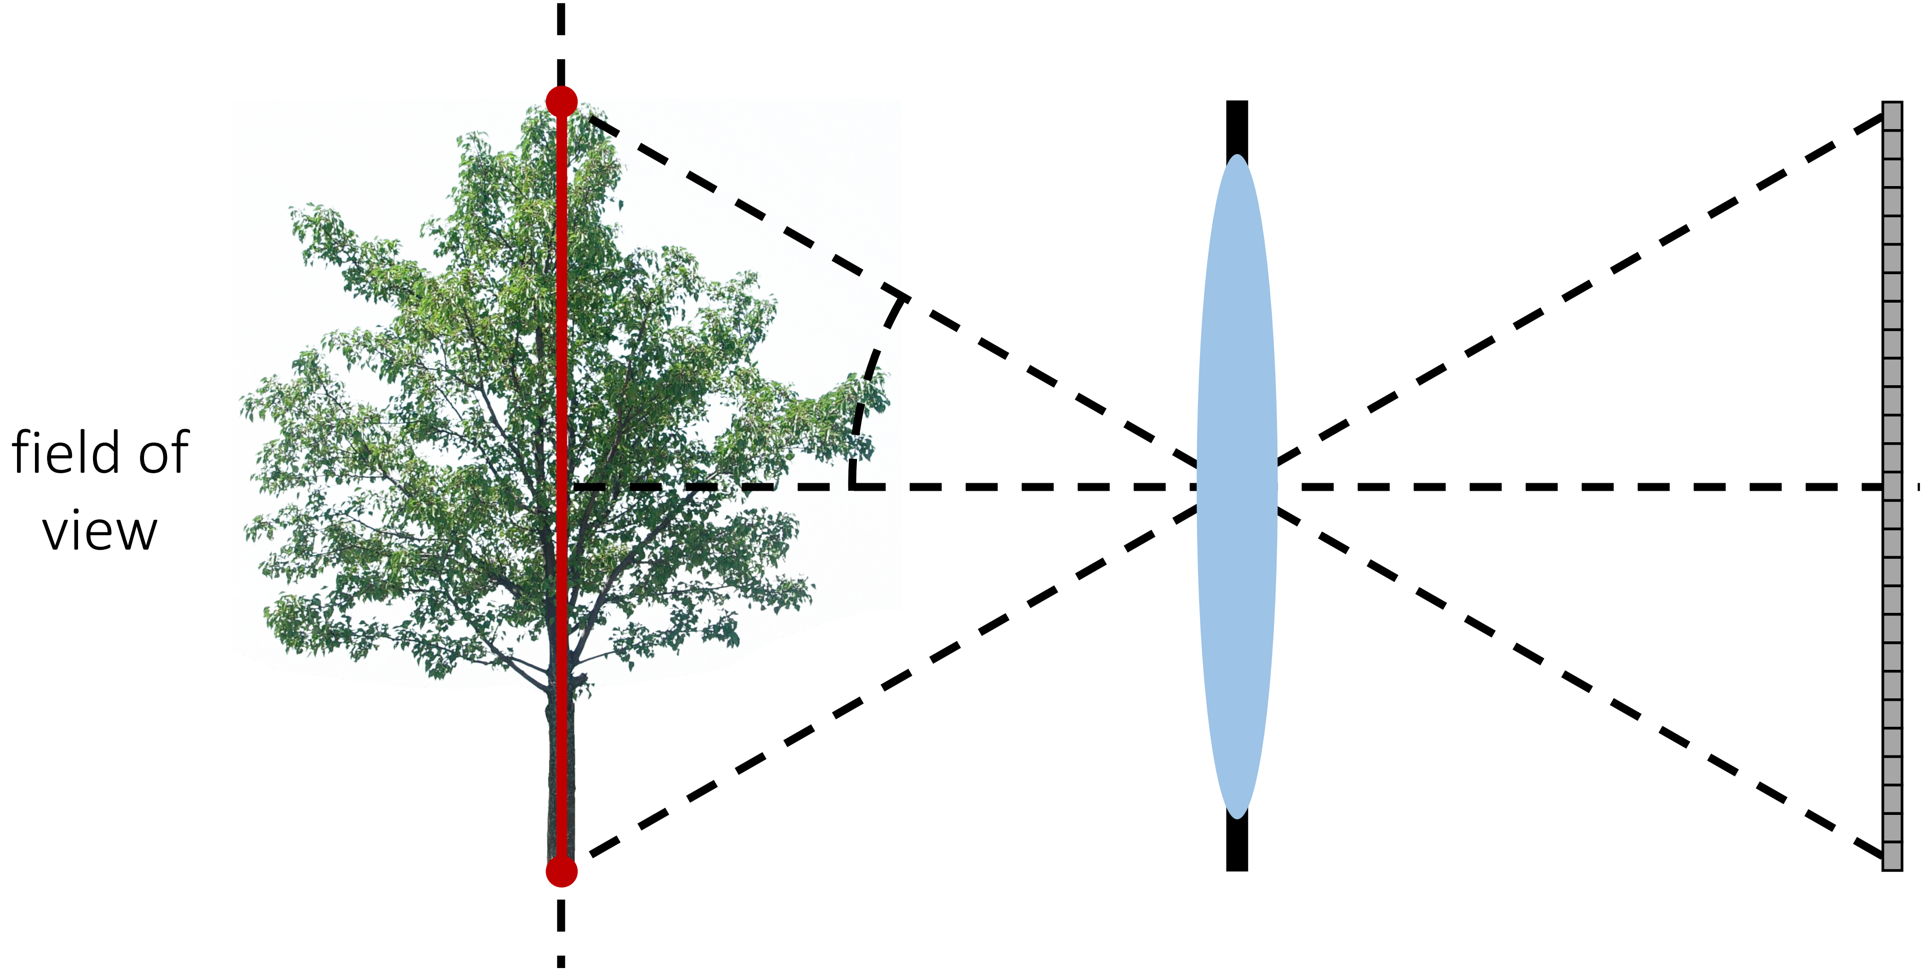

- What happens to field of view as we focus closer?

Field of view

The part of the in-focus plane that gets mapped on the sensor.

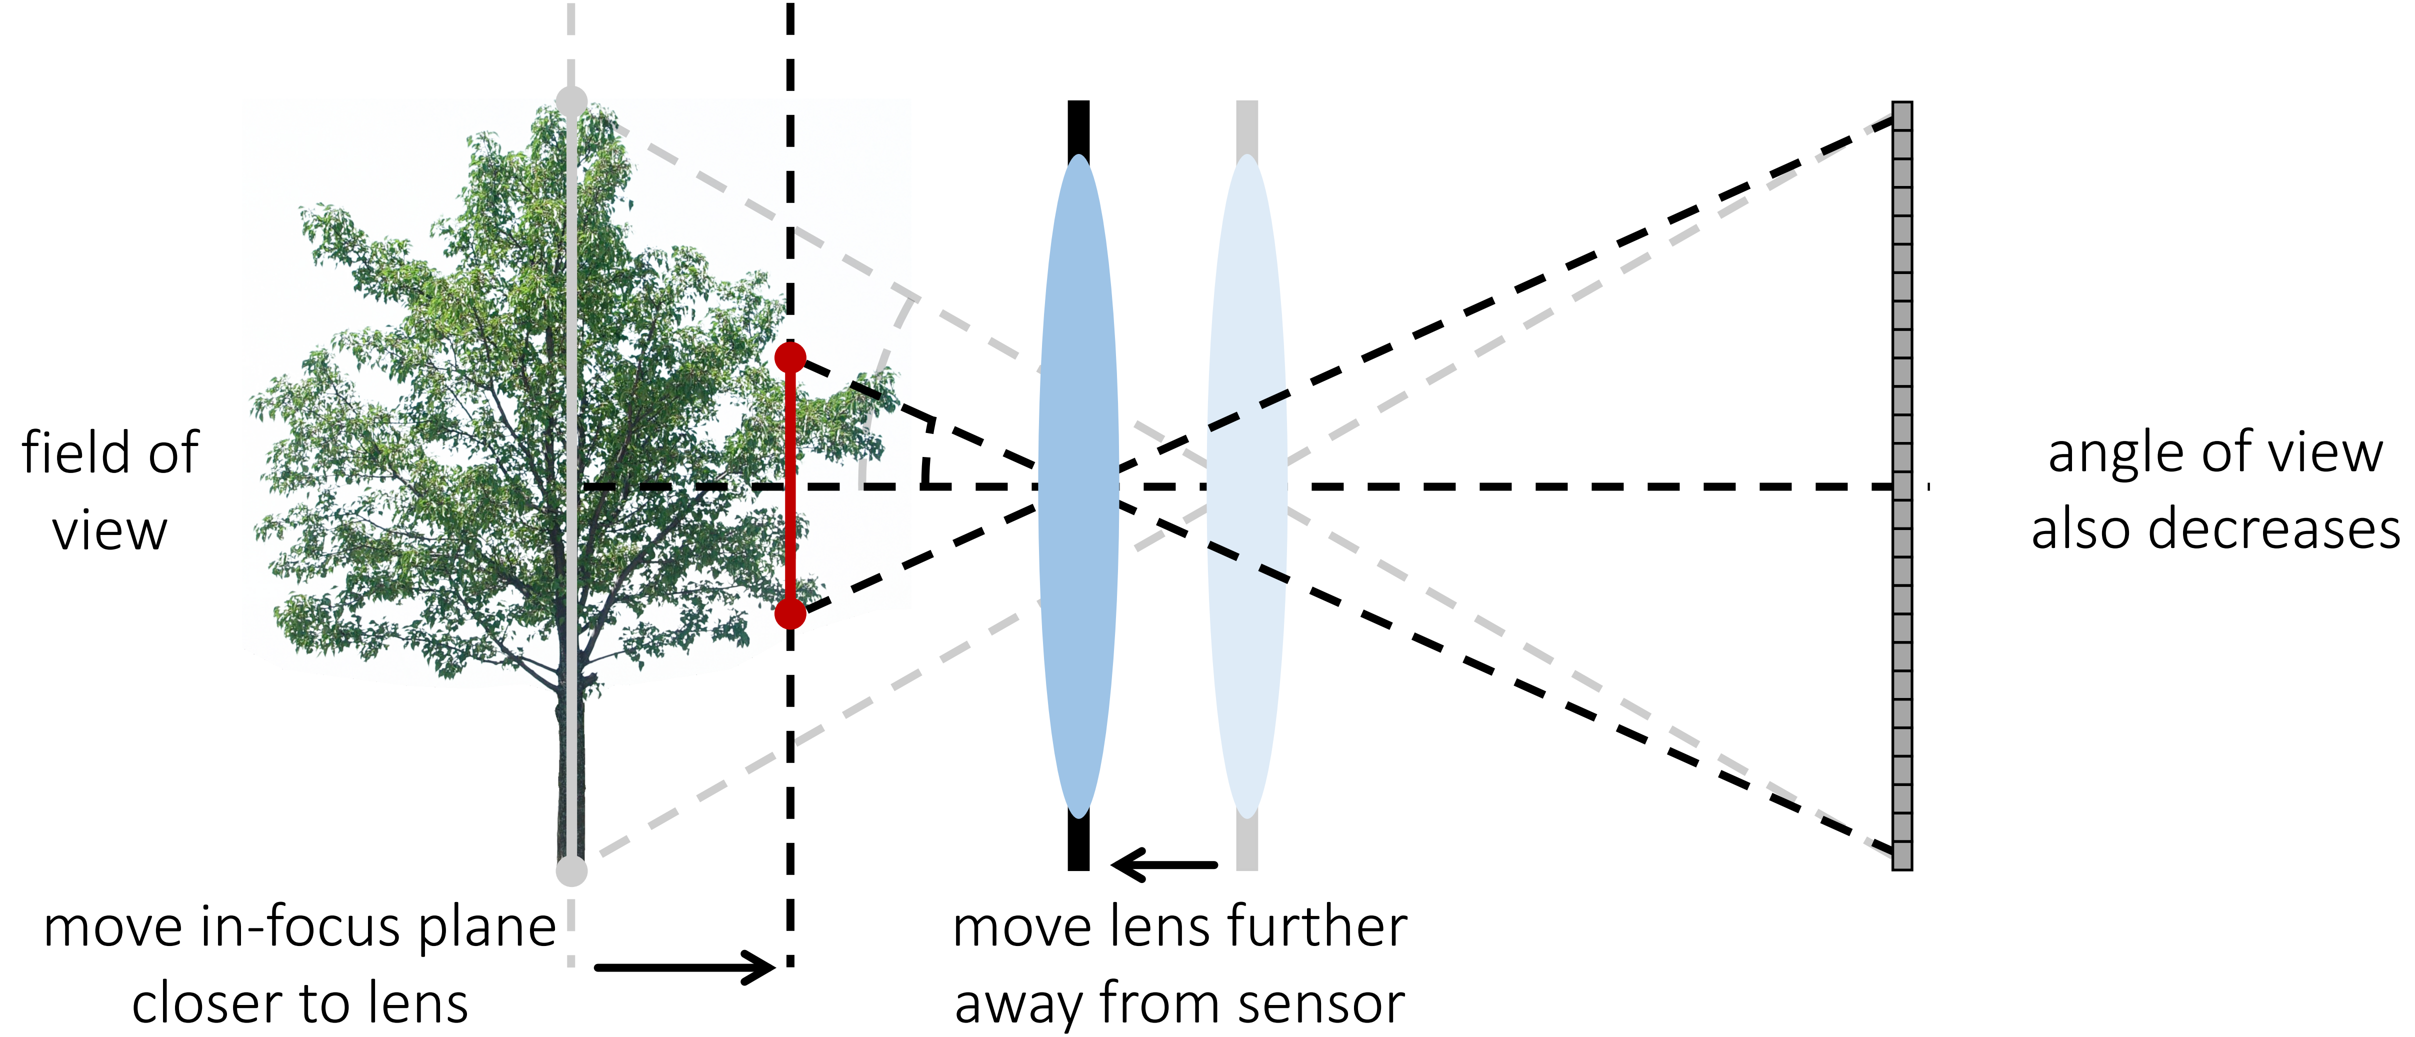

- What happens to field of view as we focus closer? → It becomes smaller.

Comparison with pinhole camera

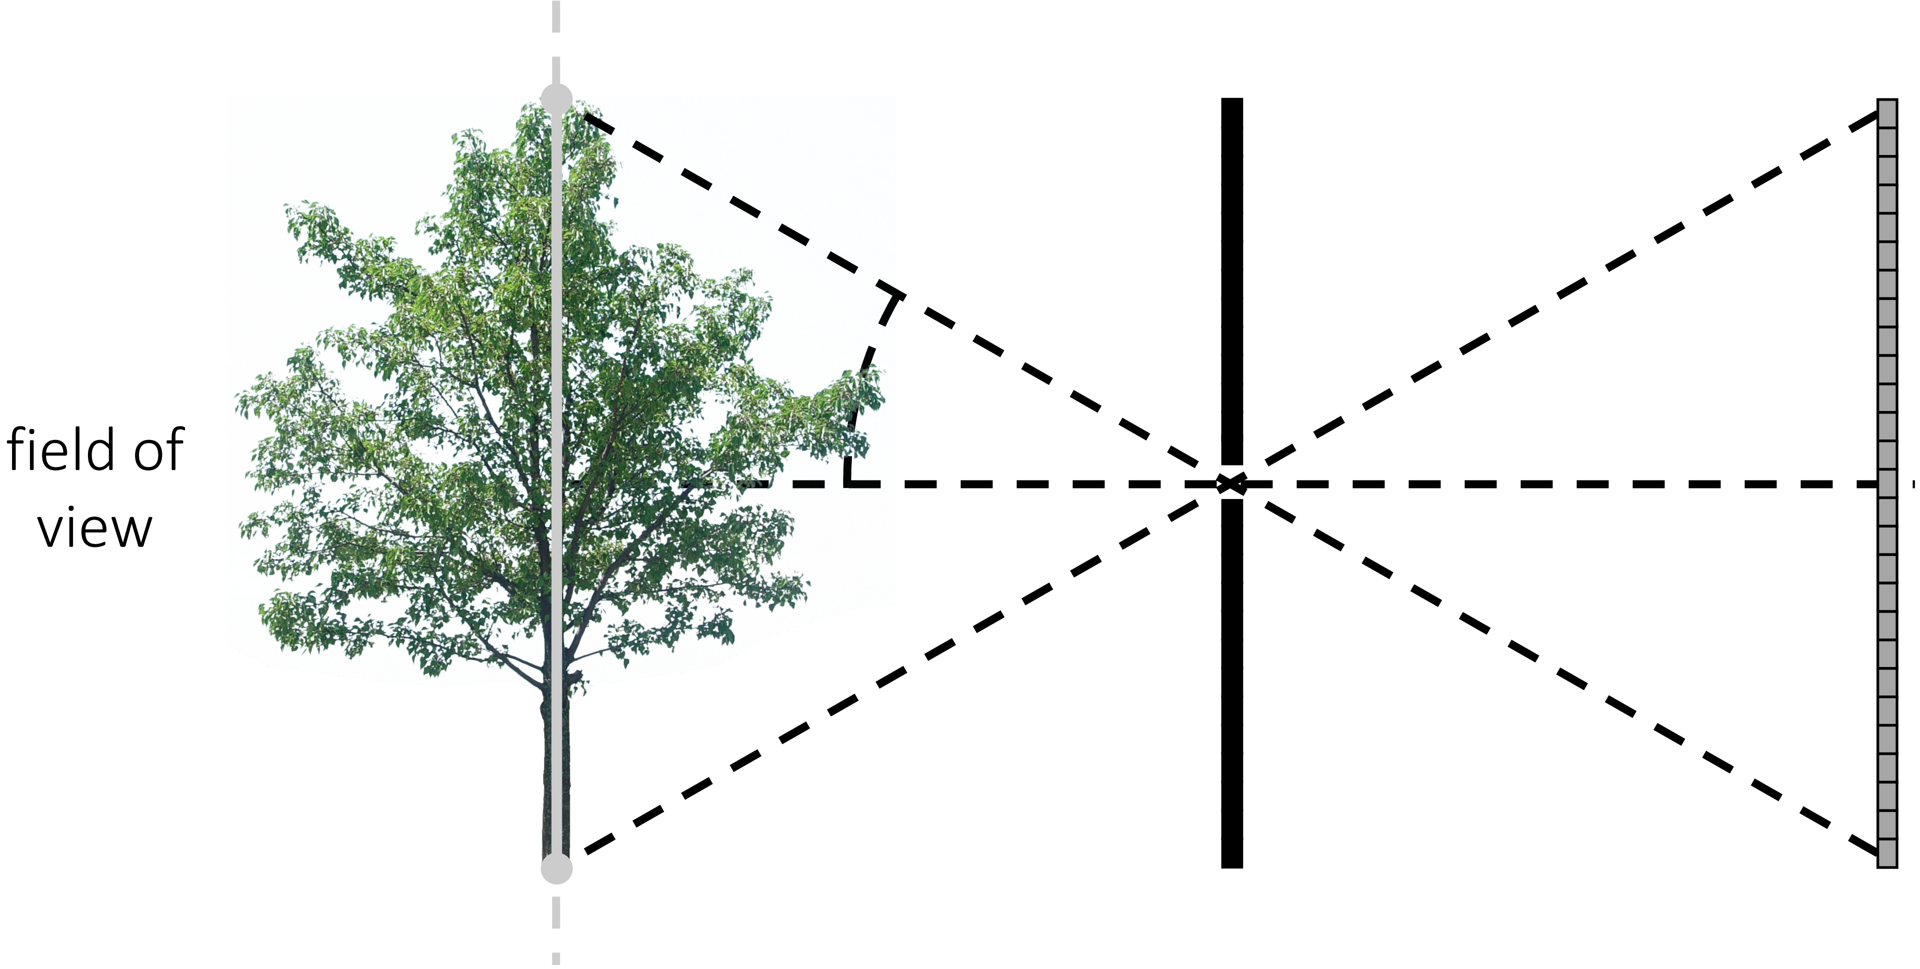

- What happens to field of view as we move closer?

Comparison with pinhole camera

No need to refocus: we can move object closer without changing aperture-sensor distance.

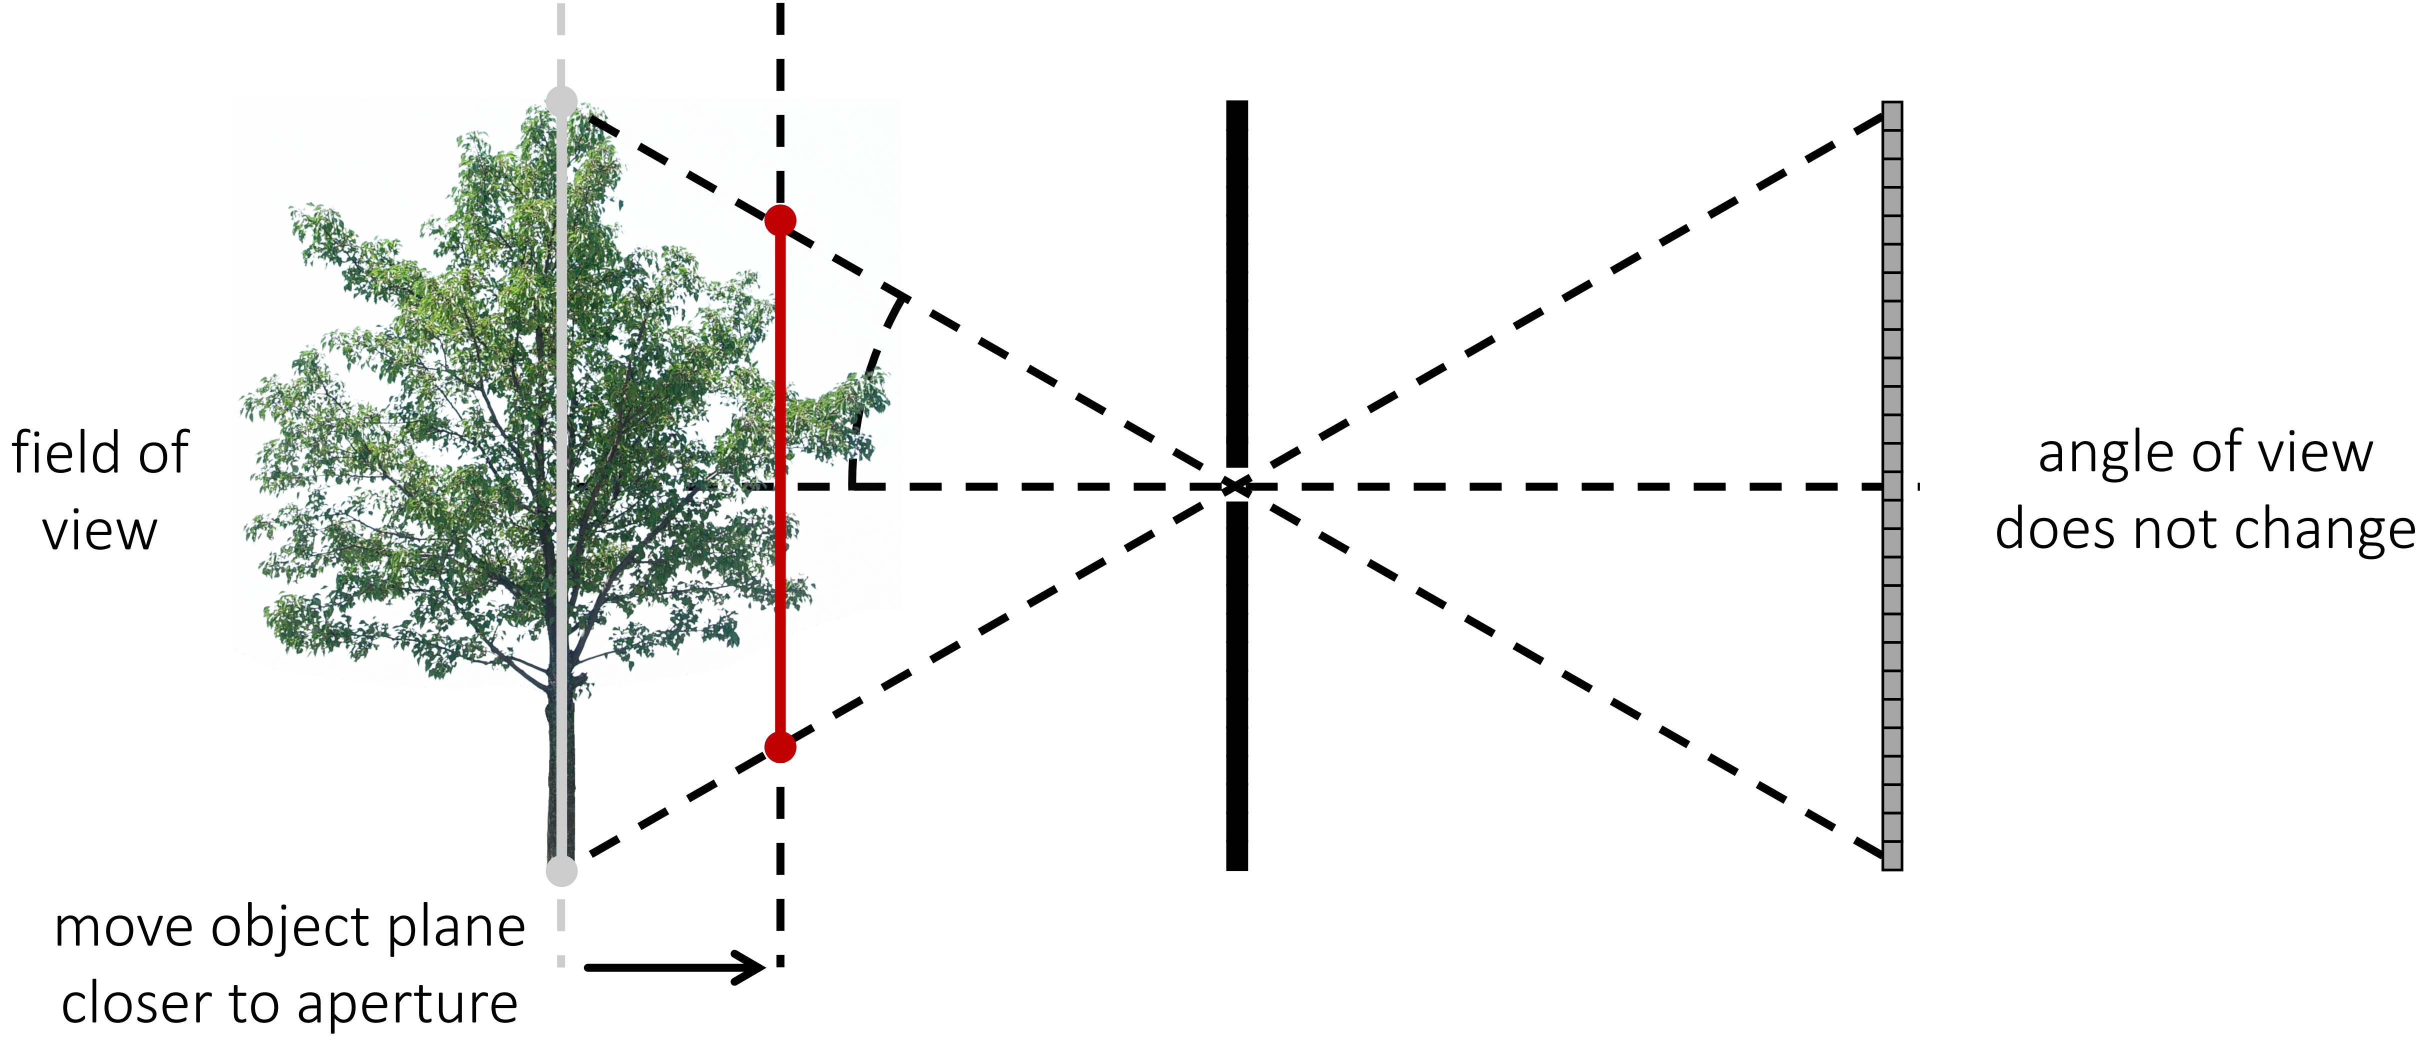

- What happens to field of view as we move closer? → It becomes smaller, but amount differs.

Comparison with pinhole camera

No need to refocus: we can move object closer without changing aperture-sensor distance.

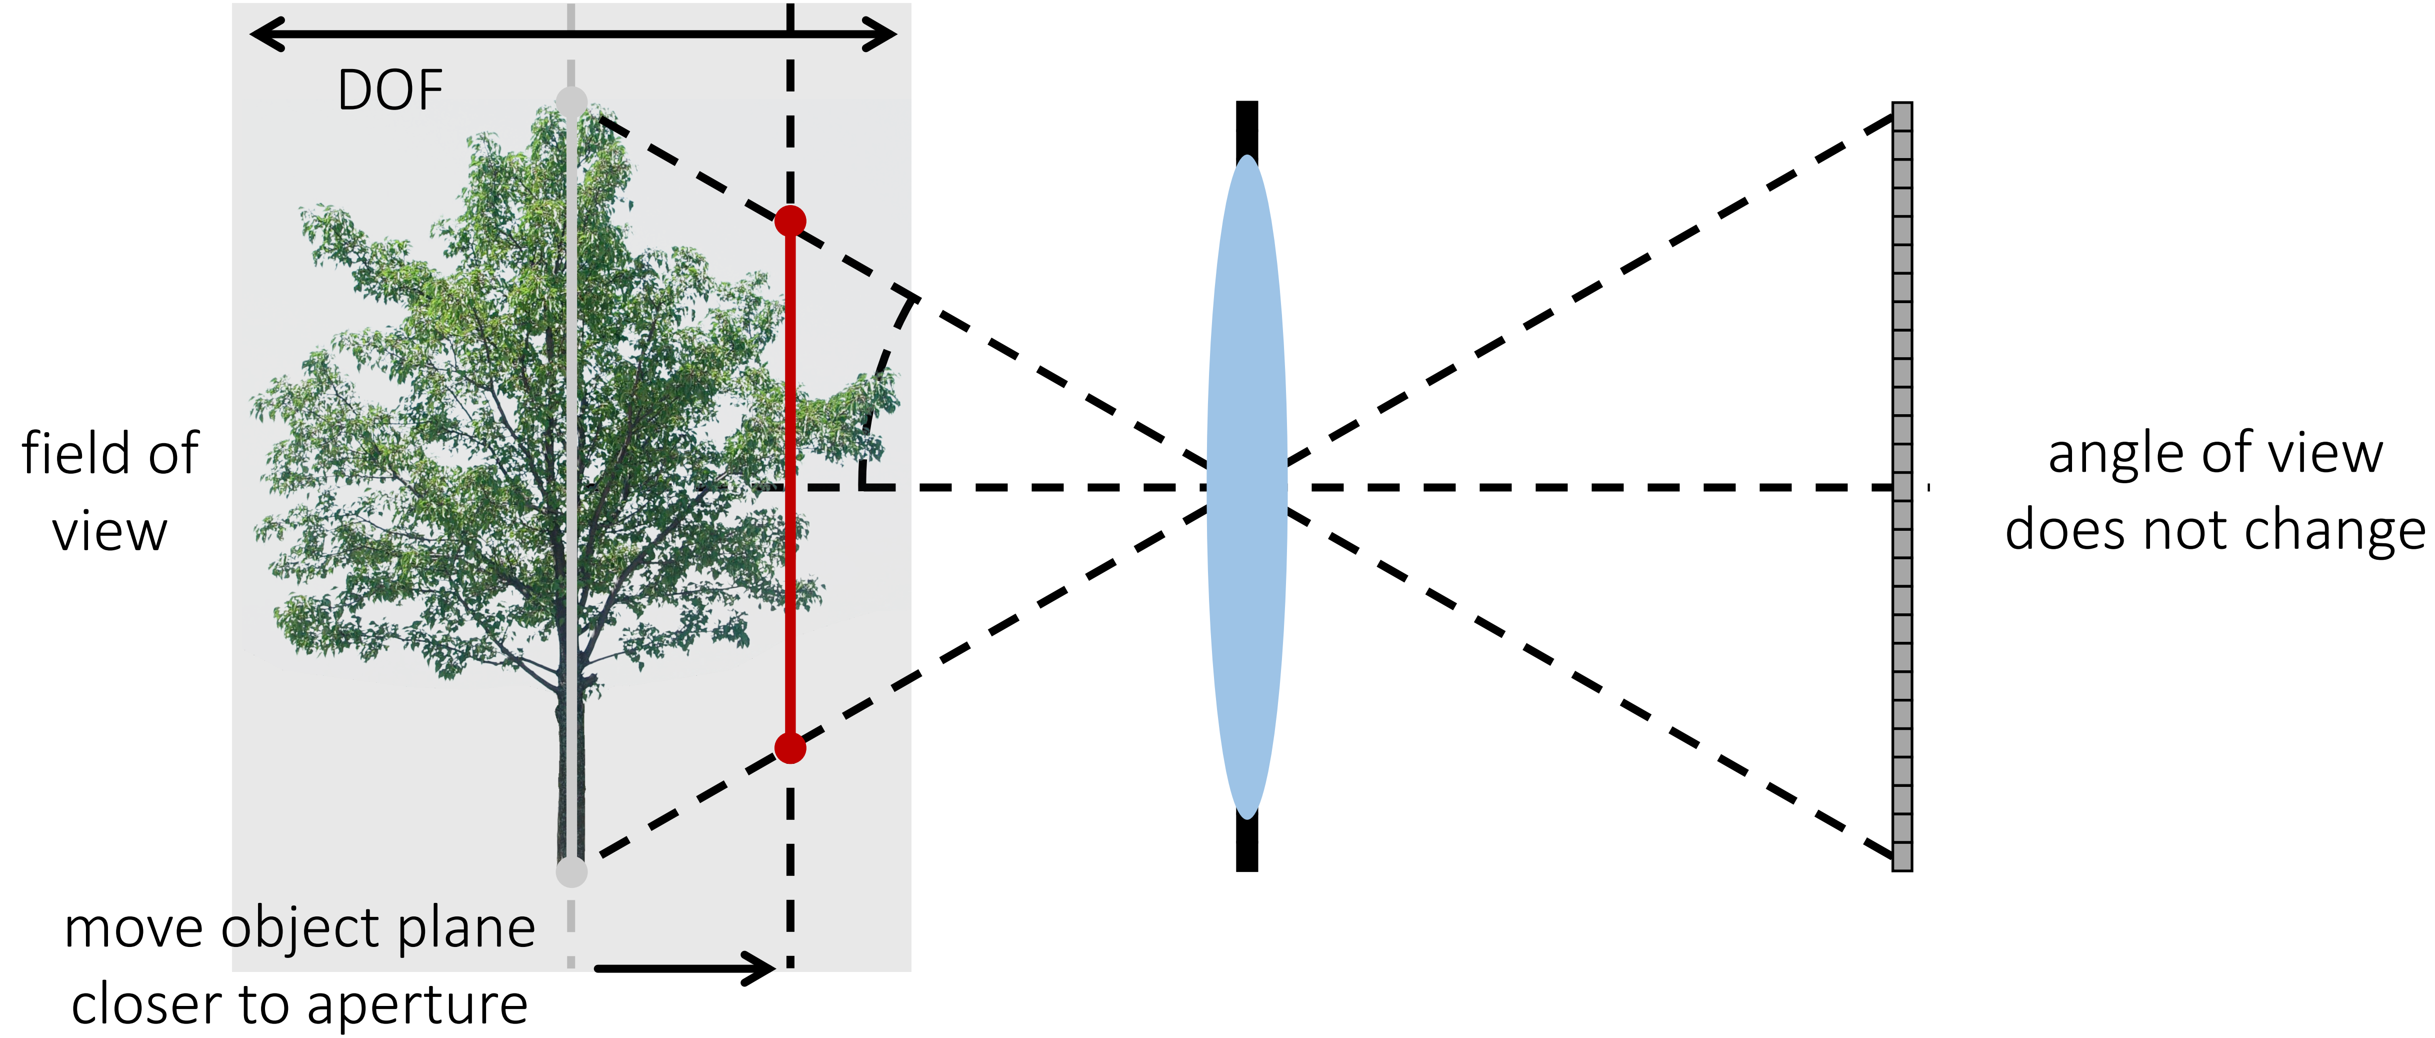

This can be done with a lens *only if* depth of field is large enough. Then the two behave the same.

Field of view also depends on sensor size

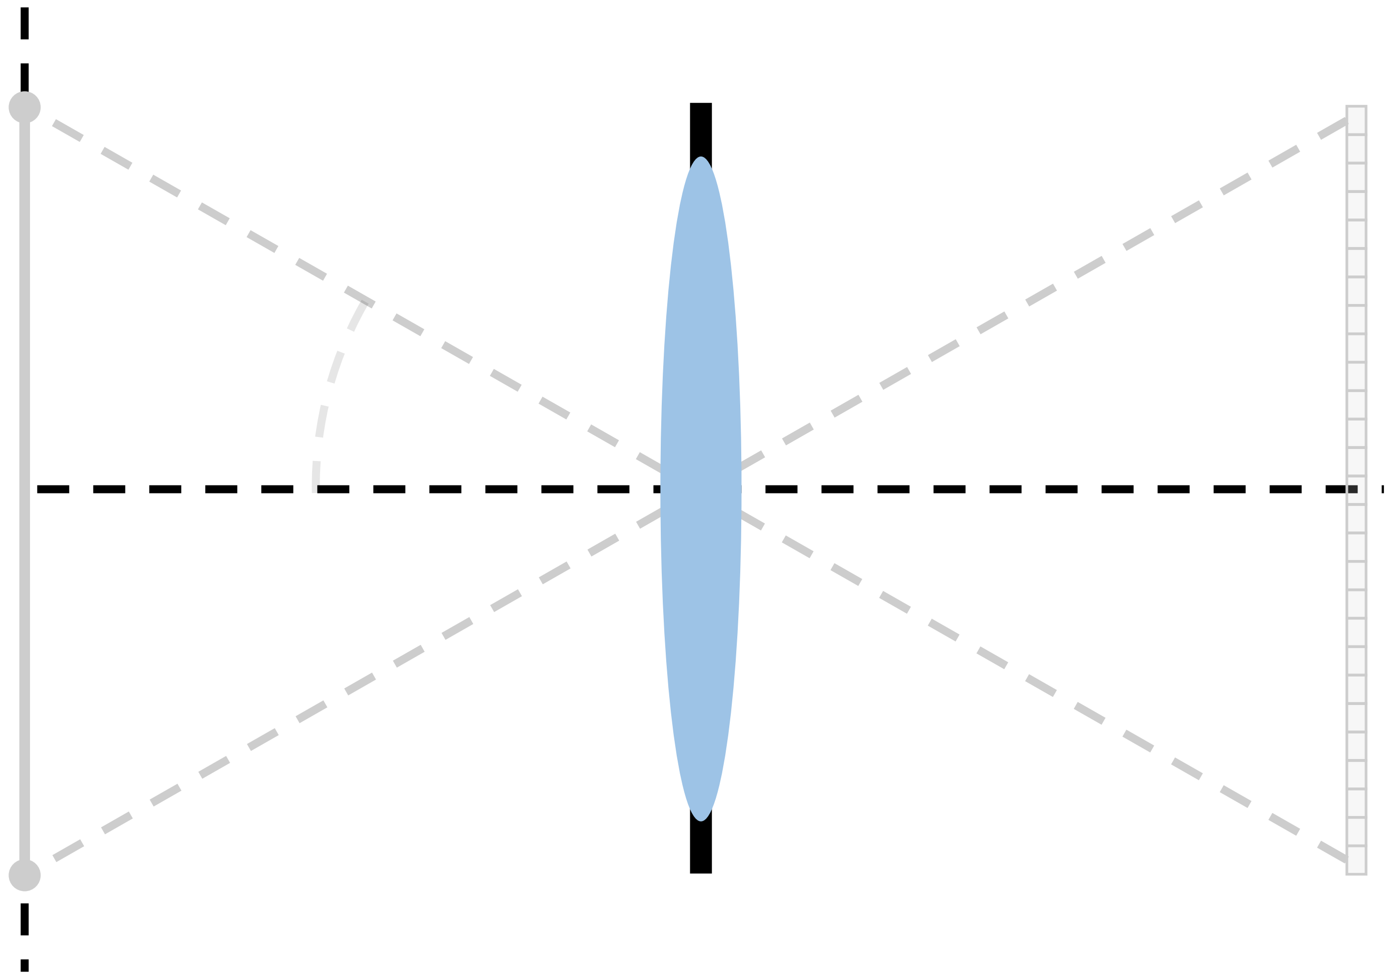

- What happens to field of view when we reduce sensor size?

Field of view also depends on sensor size

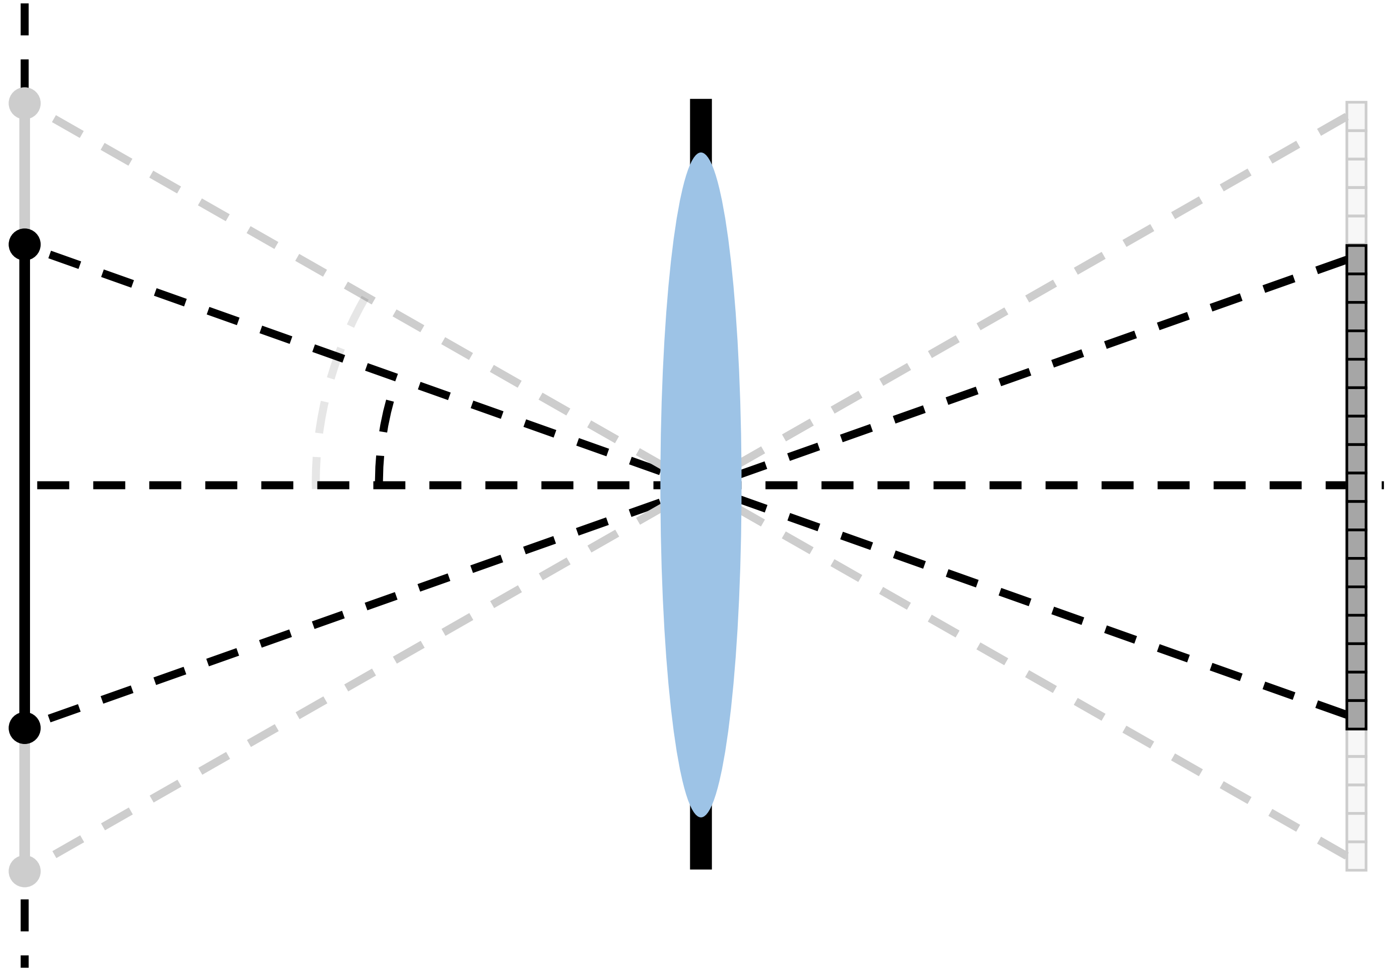

Lens and pinhole cameras behave the same in this case.

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

Field of view also depends on sensor size

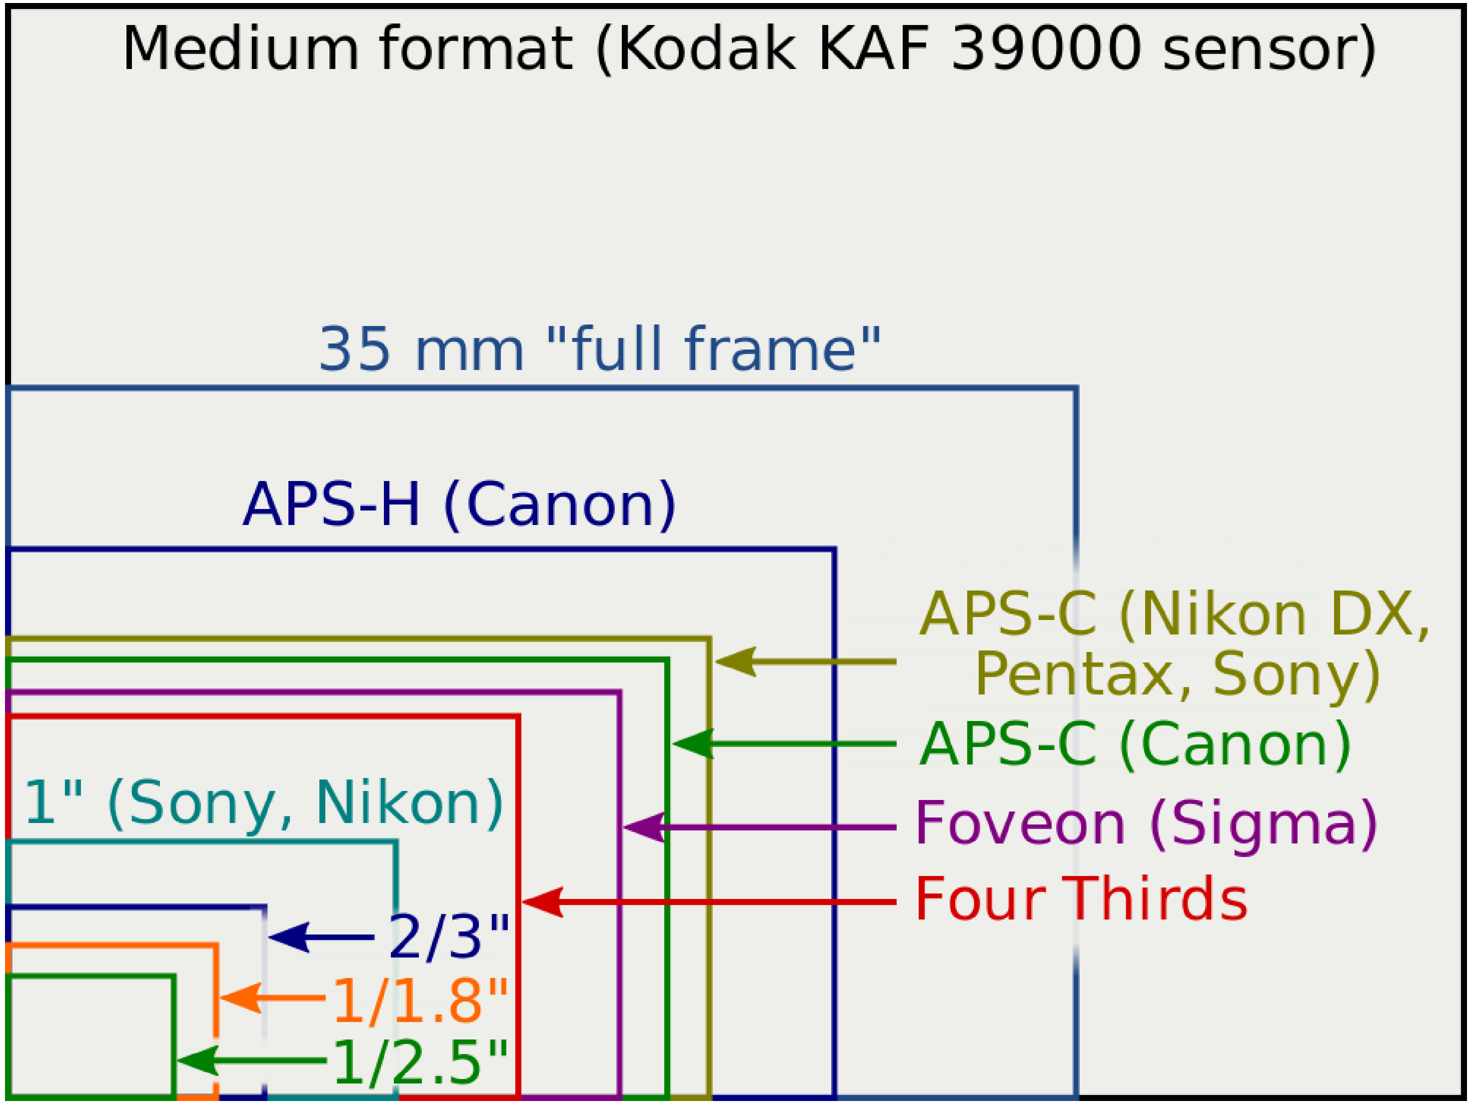

- “Full frame” corresponds to standard film size.
- Digital sensors come in smaller formats due to manufacturing limitations (now mostly overcome).
- Lenses are often described in terms of field of view on film instead of focal length.
- These descriptions are invalid when not using full-frame sensor.

Crop factor

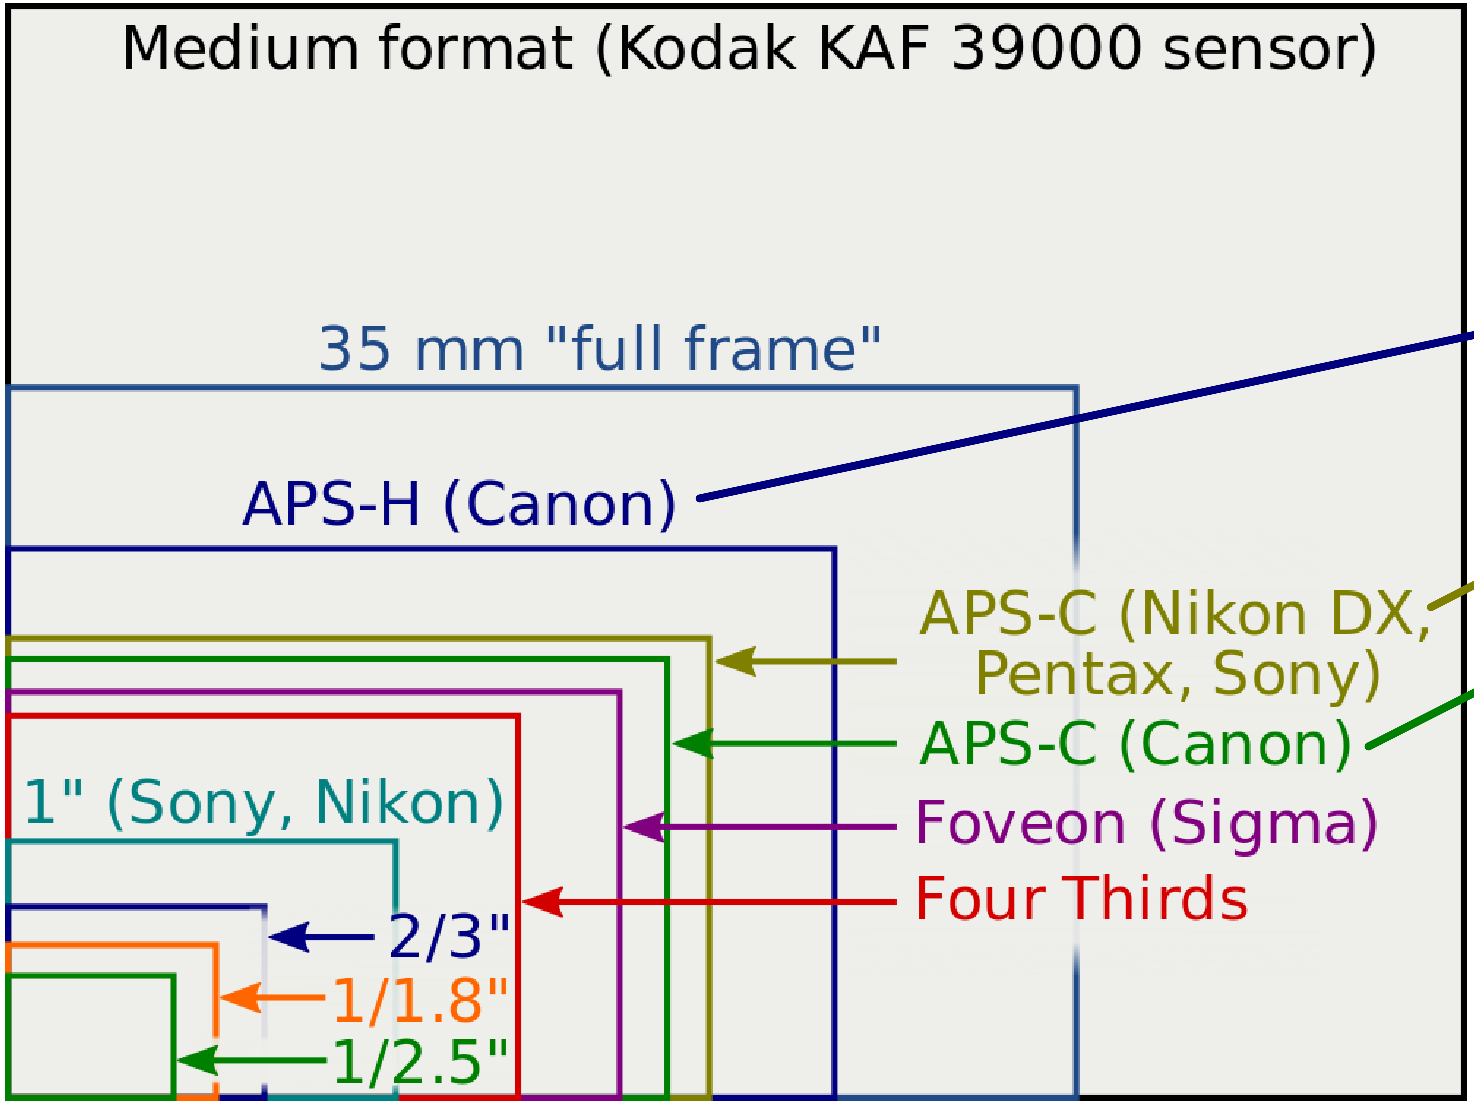

How much field of view is cropped when using a sensor smaller than full frame.

Magnification and perspective

Magnification depends on depth

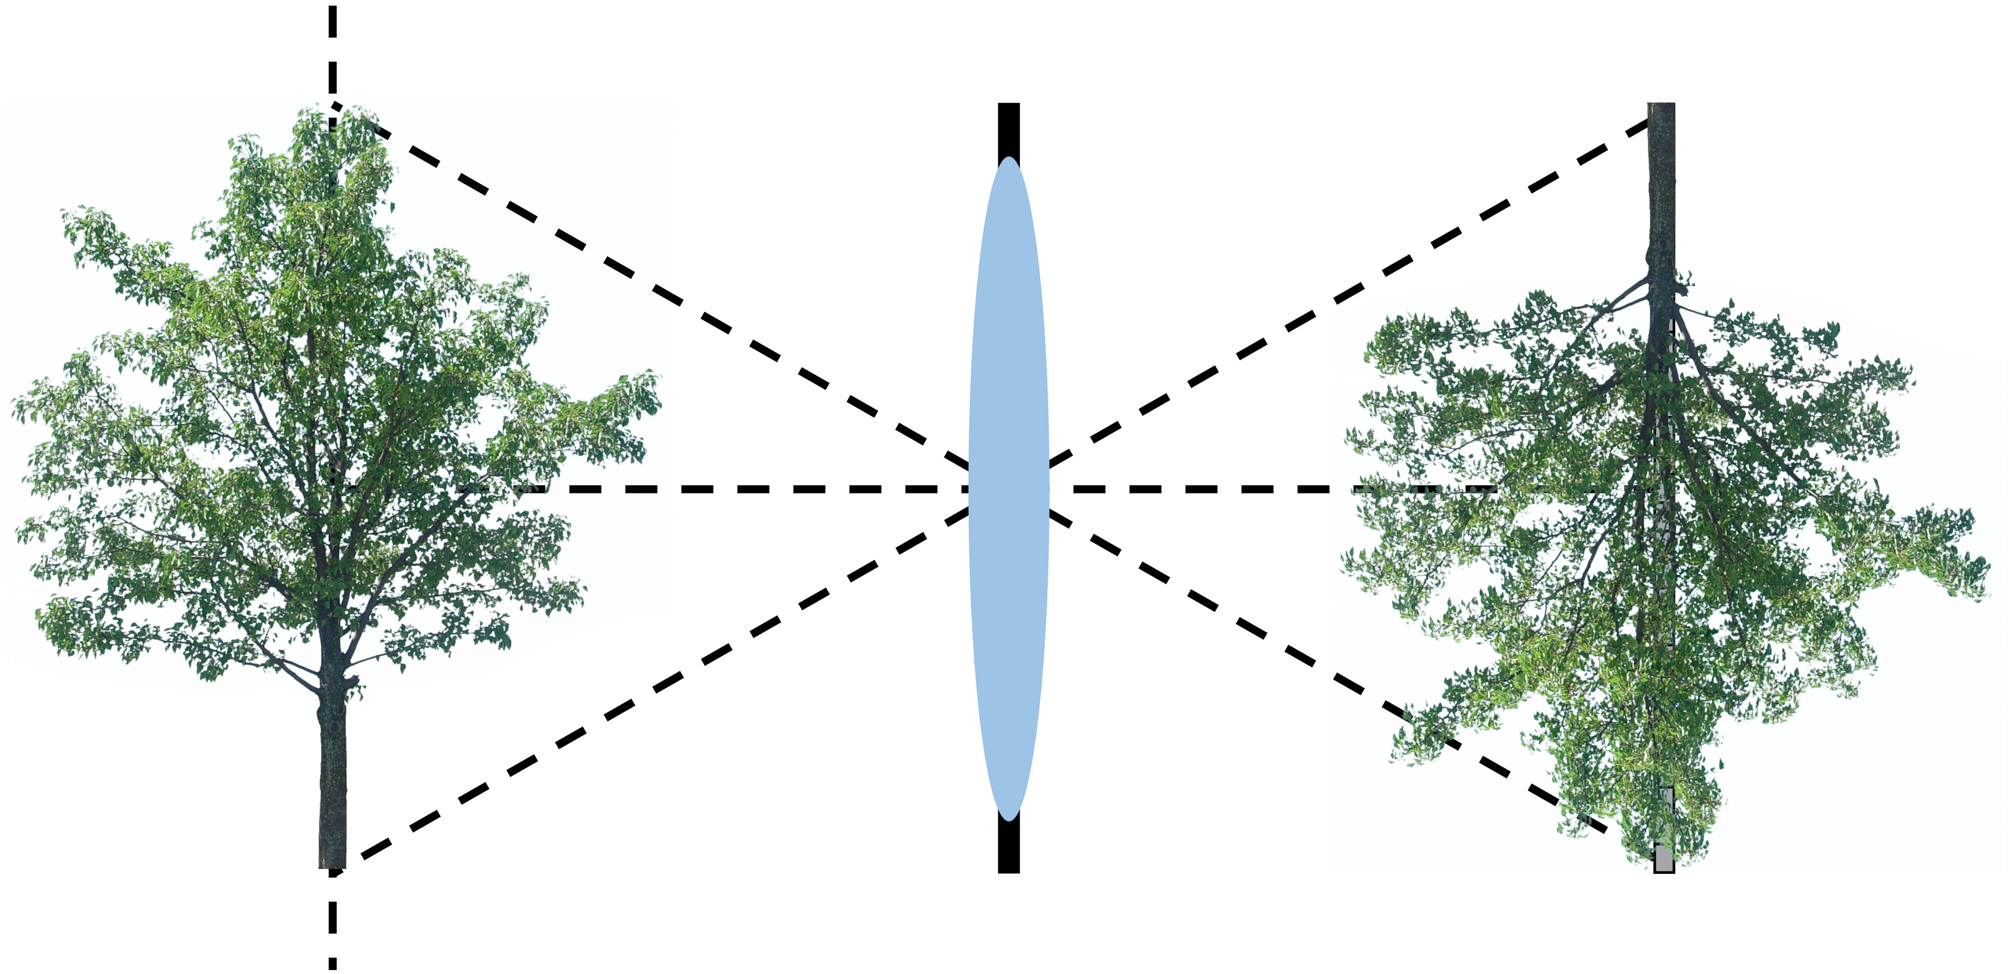

- What happens to magnification as we focus further away?

Magnification depends on depth

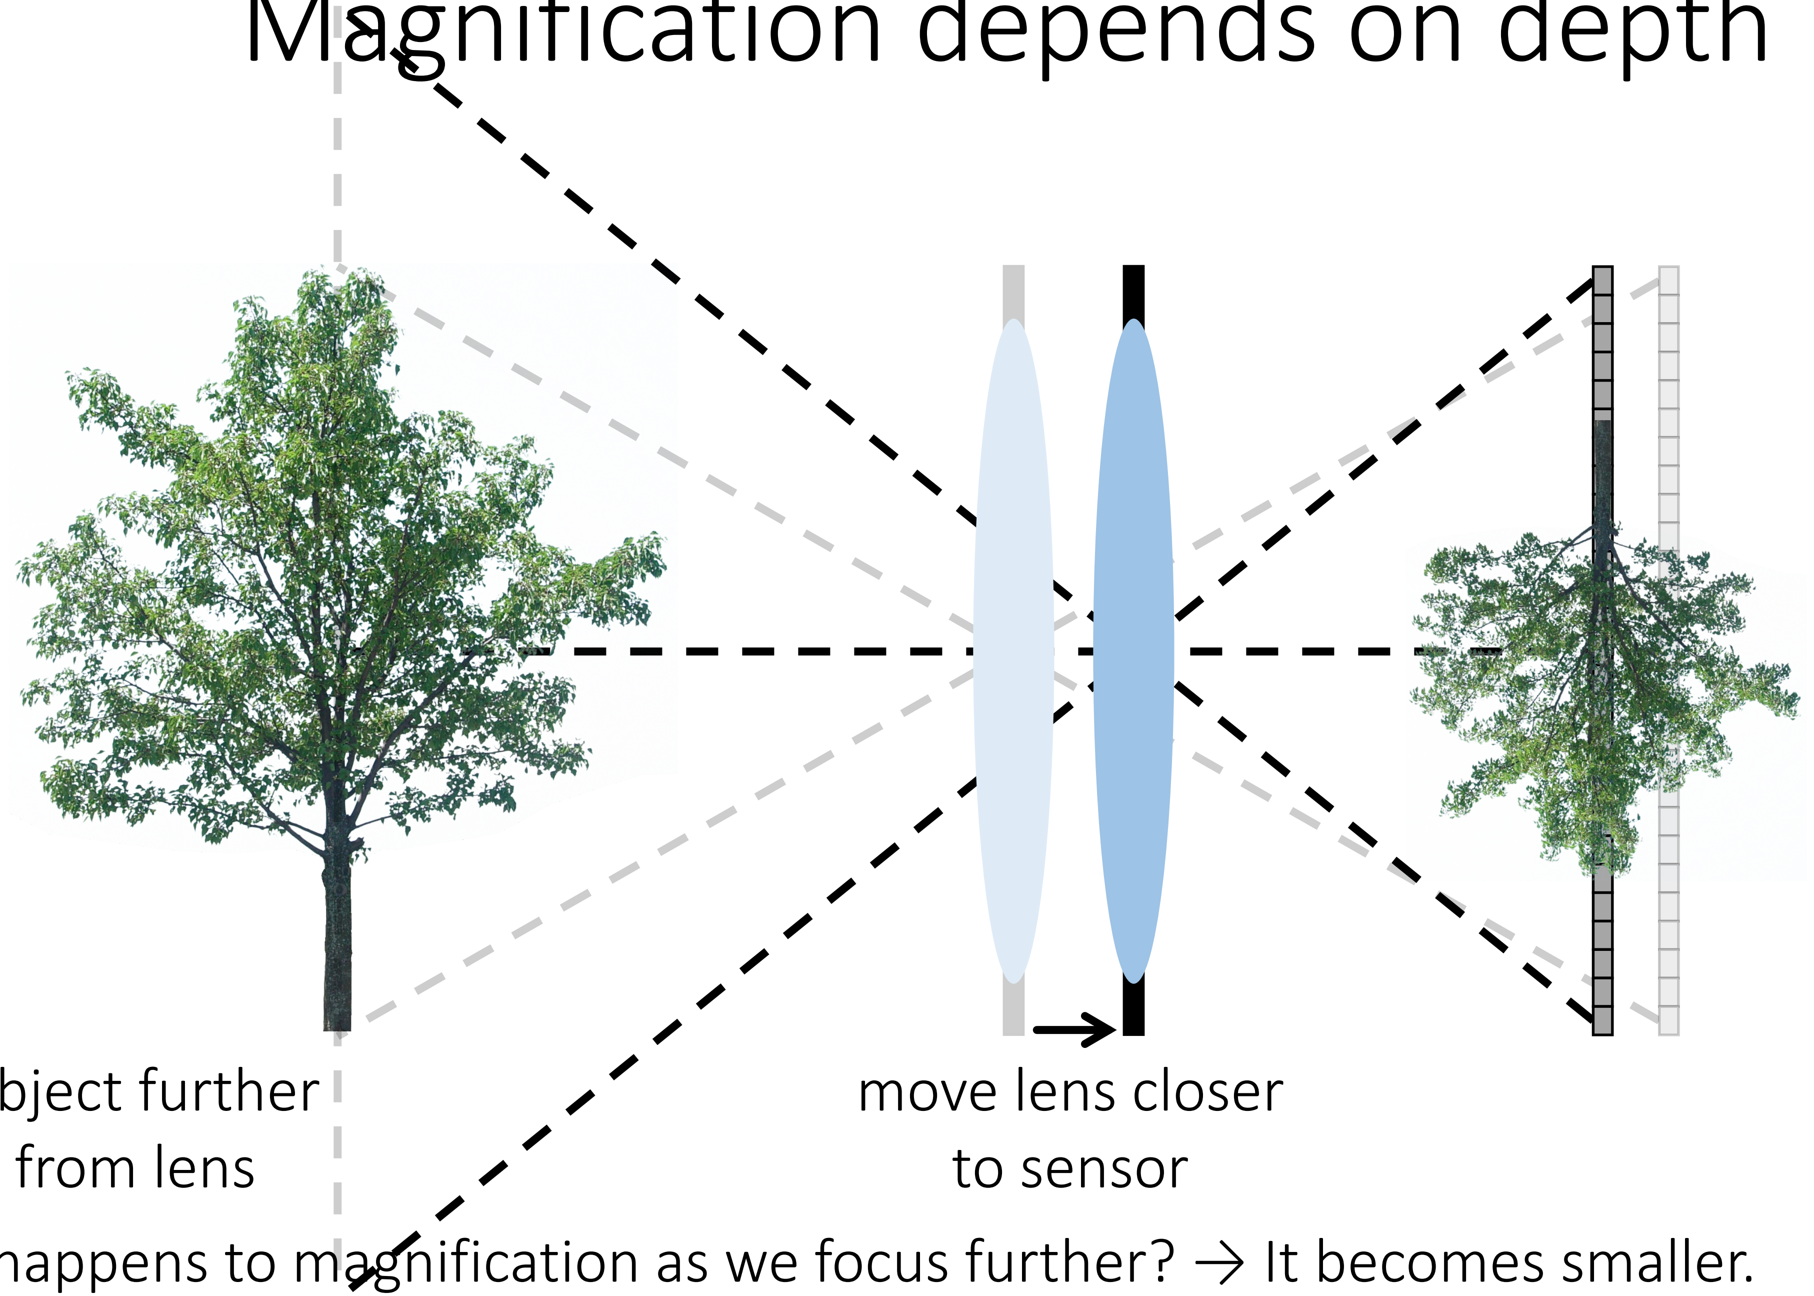

- What happens to magnification as we focus further? → It becomes smaller.

Magnification depends on depth

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

$$m = \frac{f}{S - f}$$

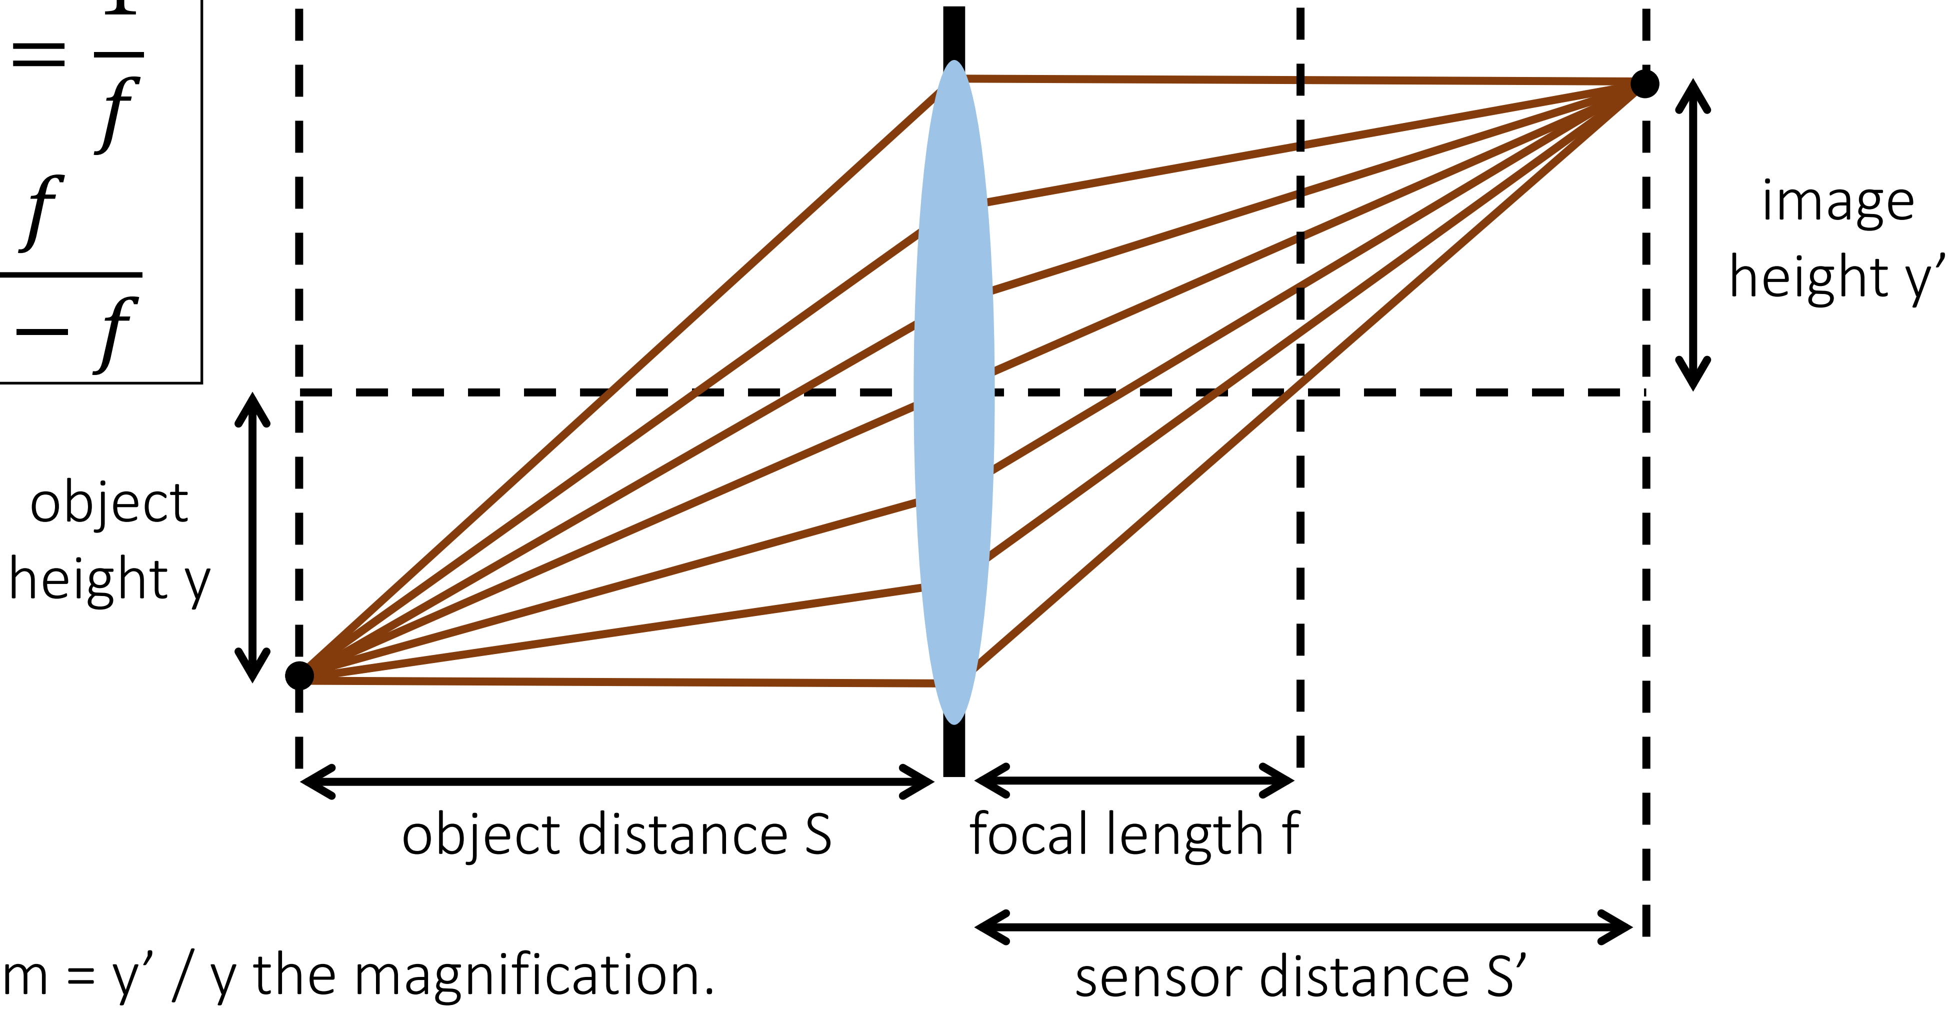

- We call m = y’ / y the magnification.

Comparison with pinhole camera

No need to refocus: we can move object further without changing aperture-sensor distance.

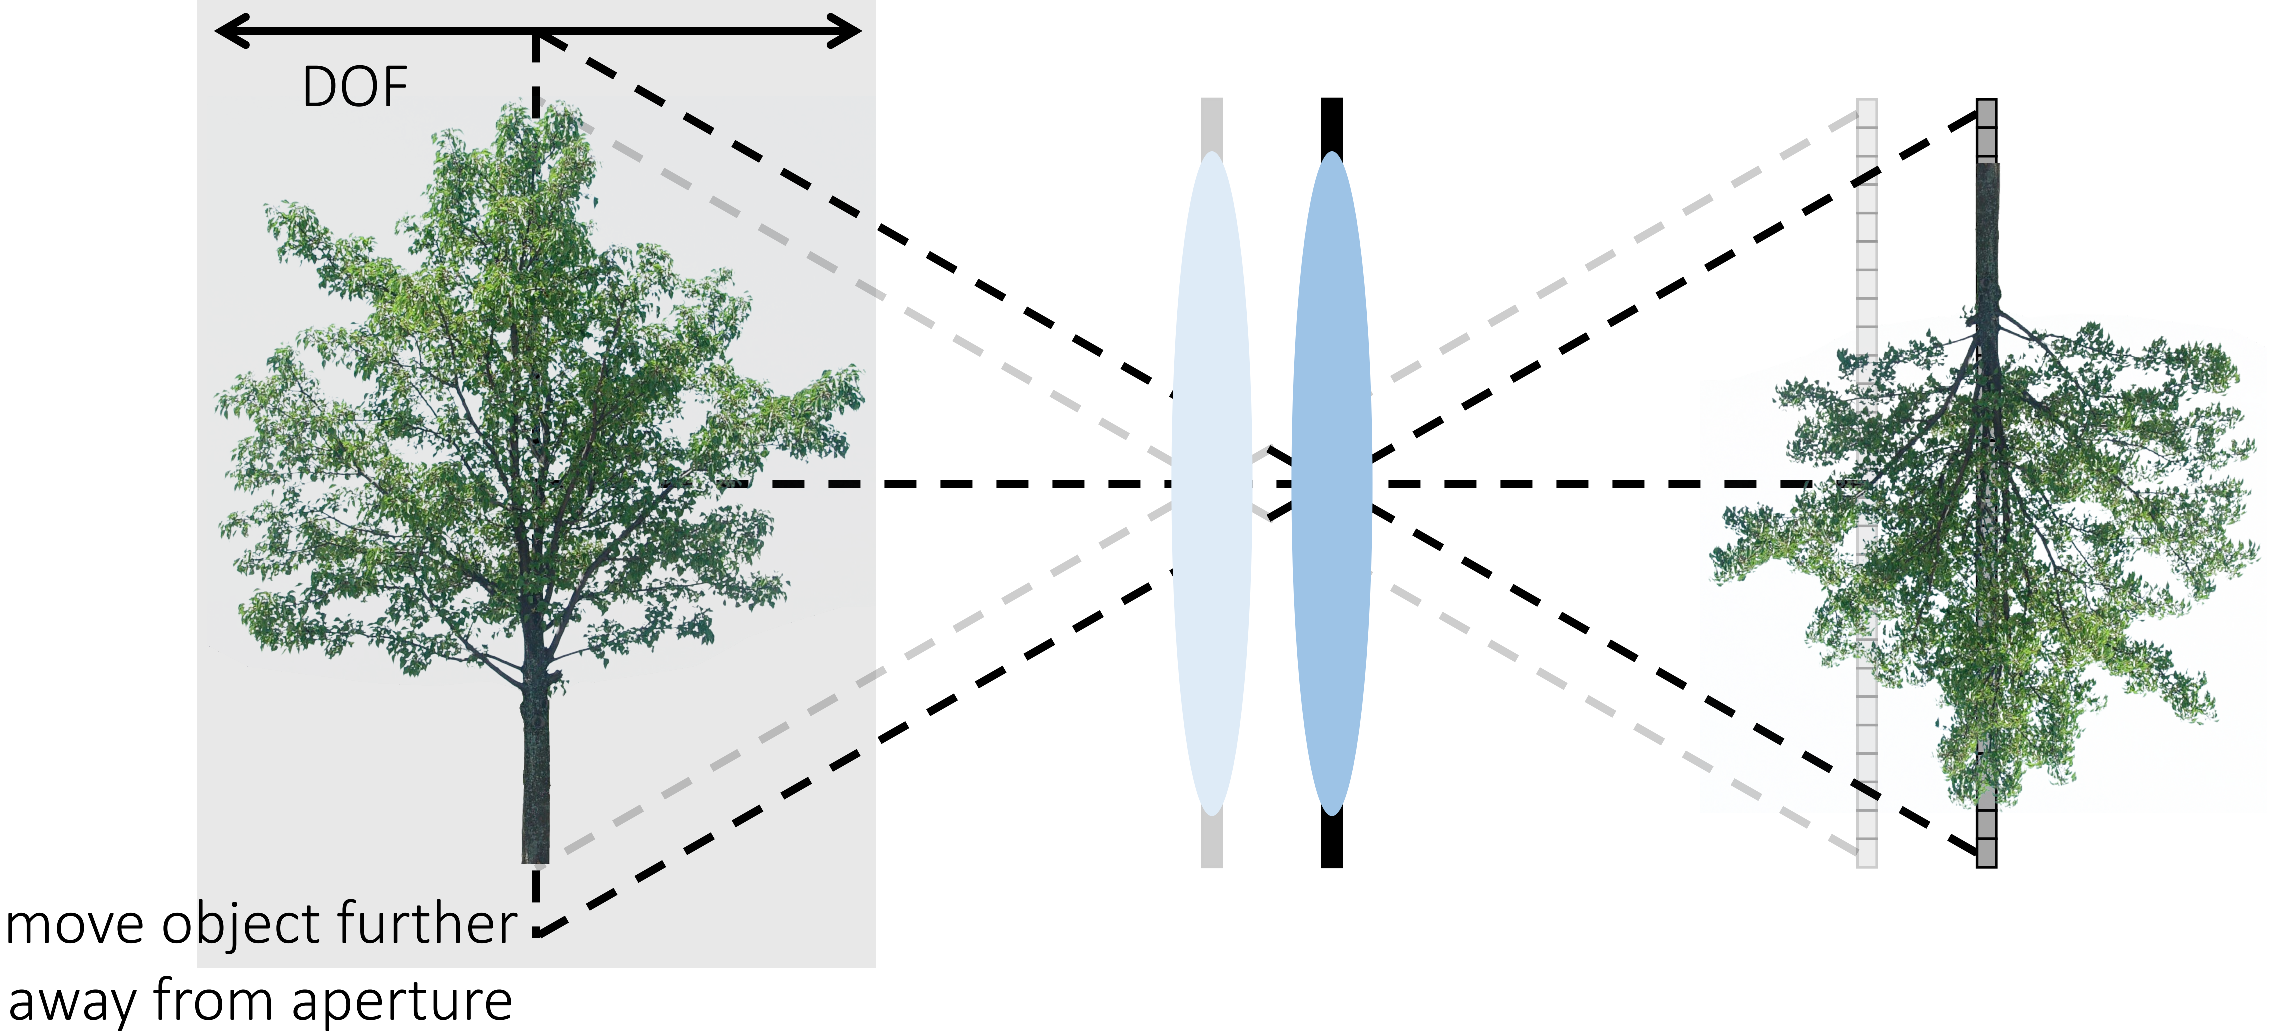

This can be done with a lens *only if* depth of field is large enough. Then the two behave the same.

Forced perspective

The Ames room illusion

The Ames room illusion

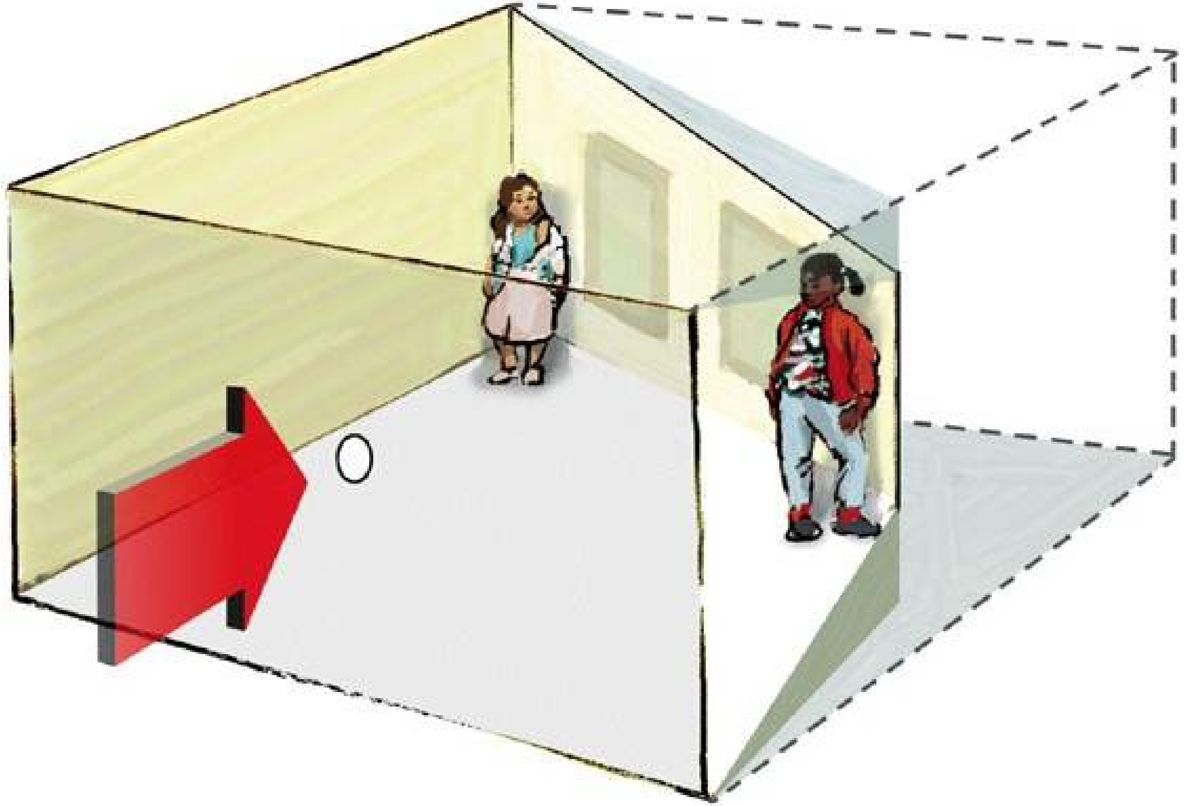

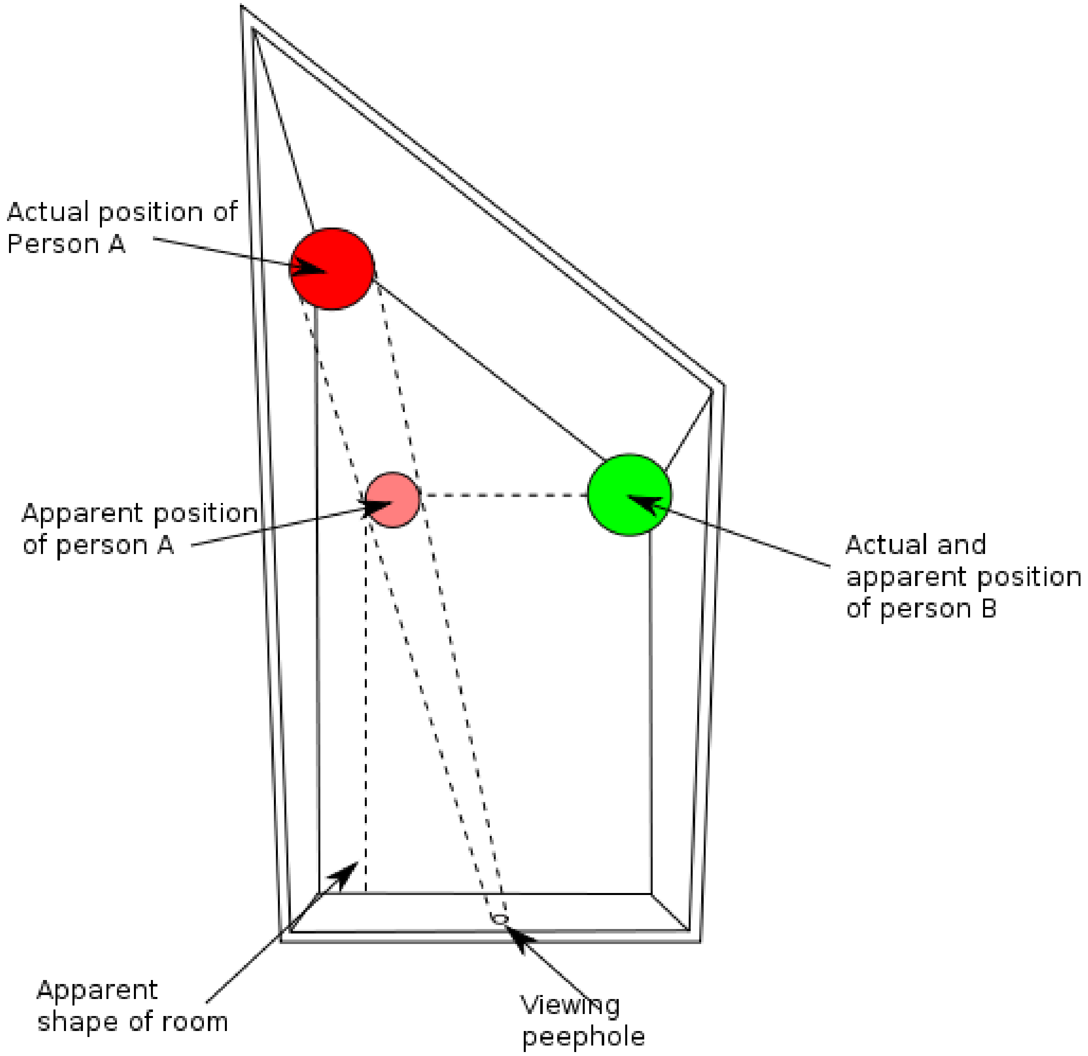

The arrow illusion

Prof. Kokichi Sugihara has many other amazing illusions involving perspective distortion, check them out on YouTube or on his website: http://www.isc.meiji.ac.jp/~kokichis/

Zooming

Zooming means changing the focal length

Very different process from refocusing

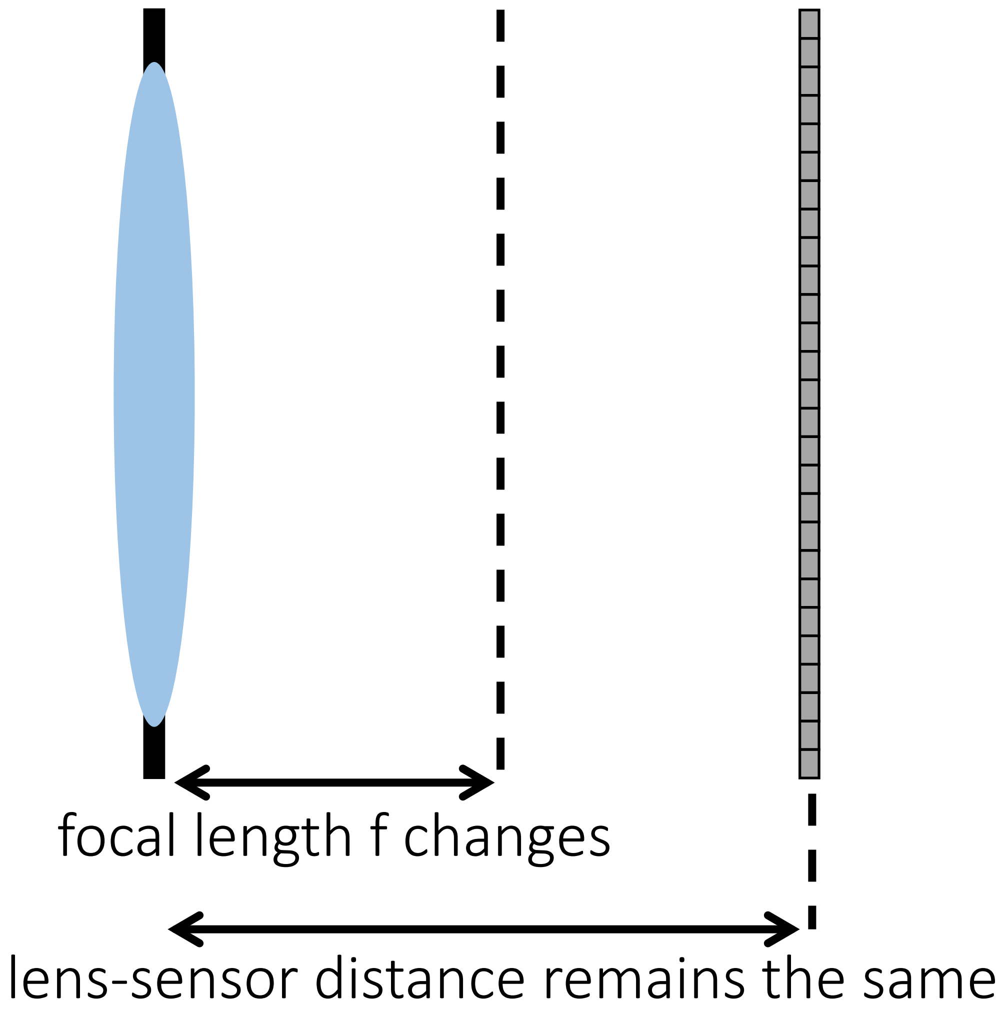

Zooming and field of view

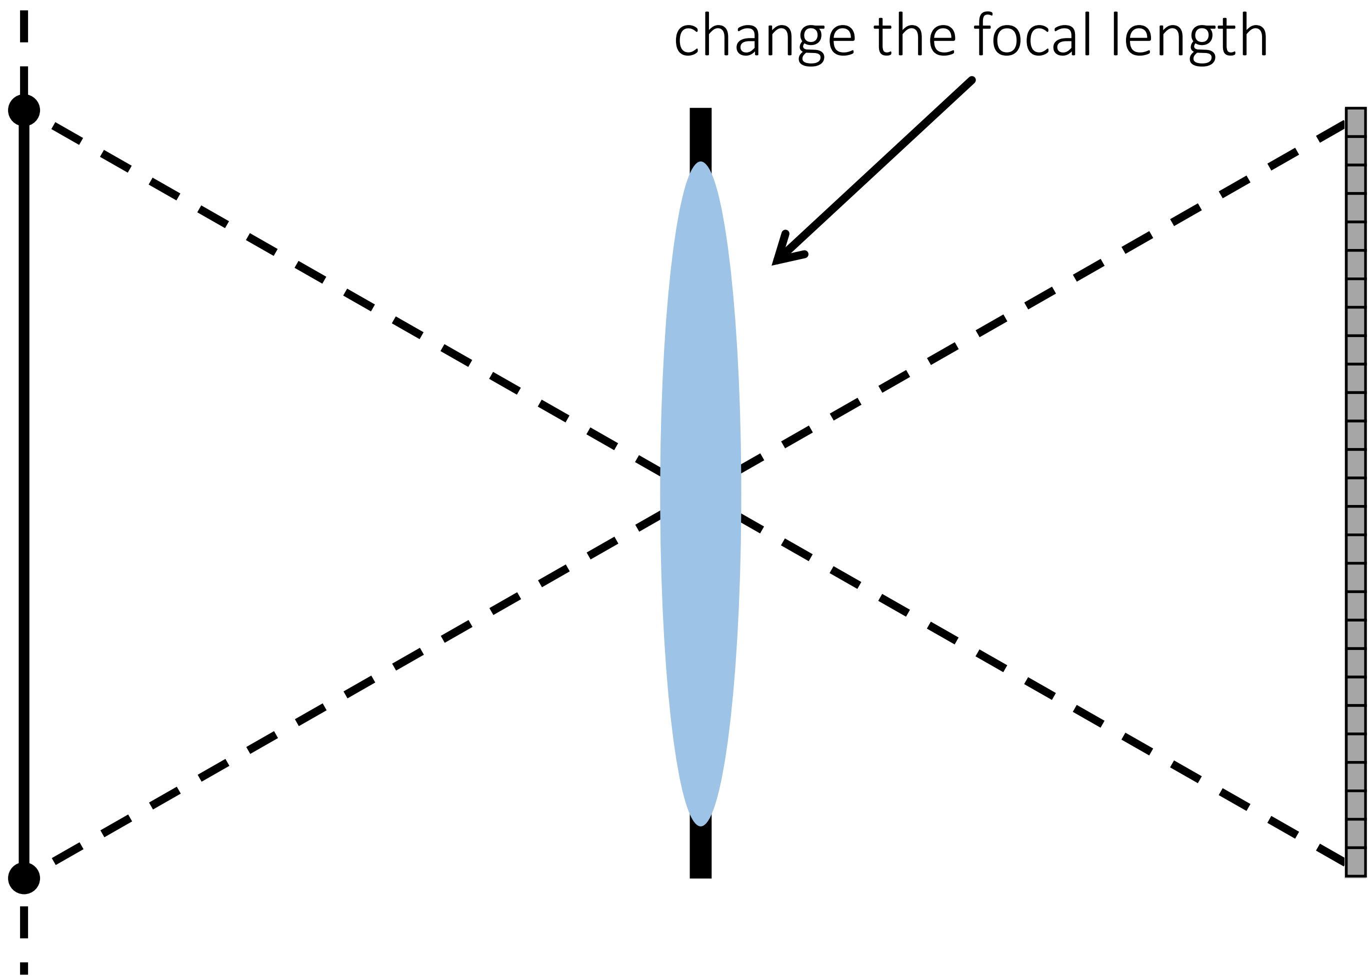

- What happens to field of view when we focus closer? → It decreases.
- What happens to field of view when we increase <u>lens</u> focal length?

Zooming and field of view

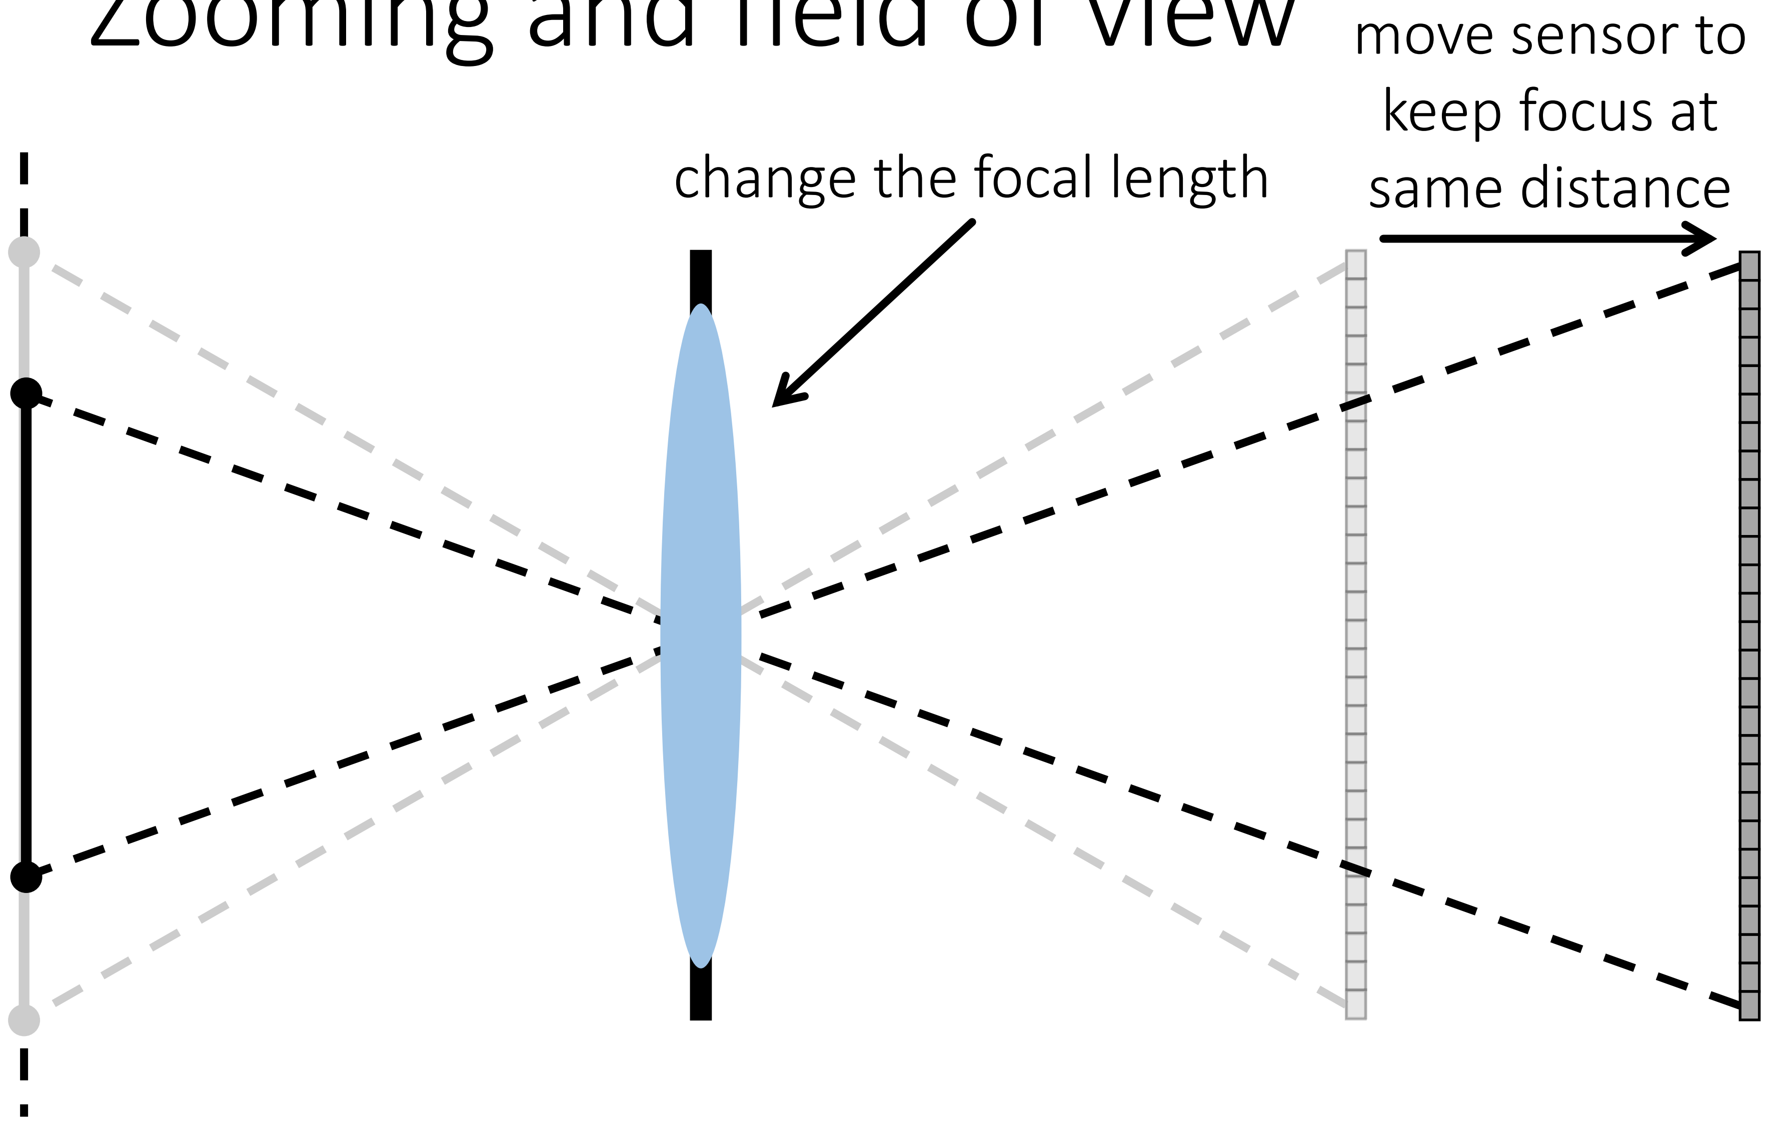

- When we increase lens focal length, field of view decreases (we “zoom in”).

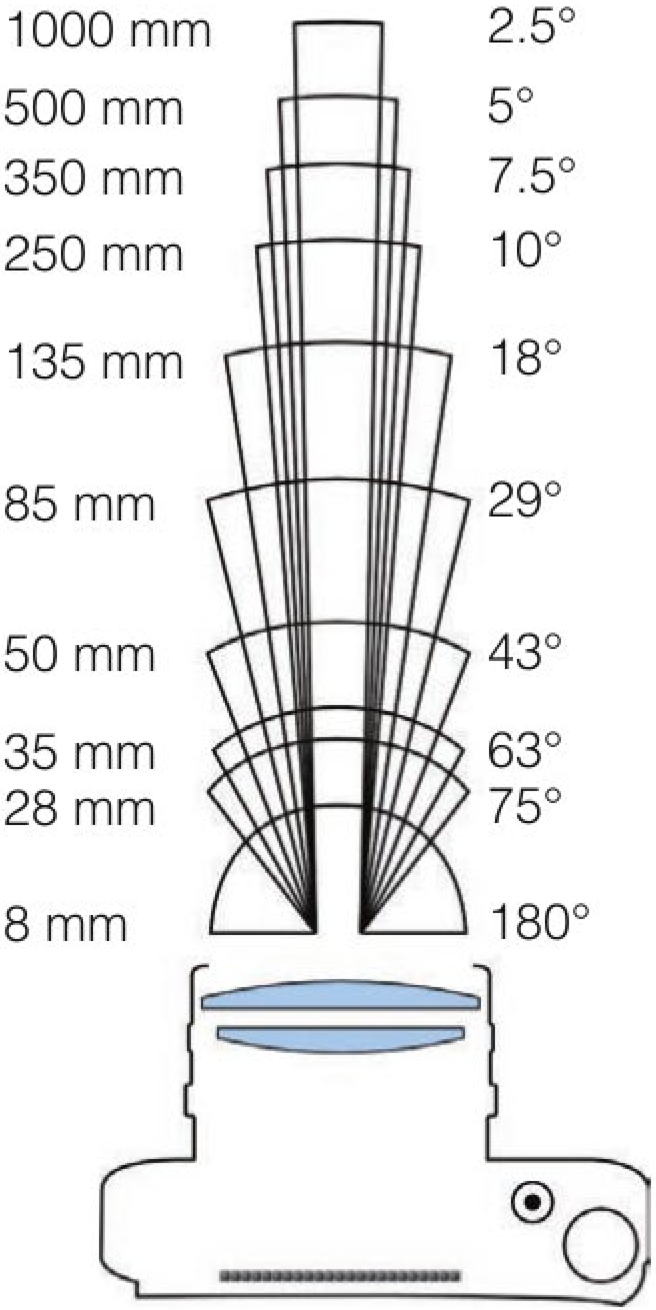

Field of view

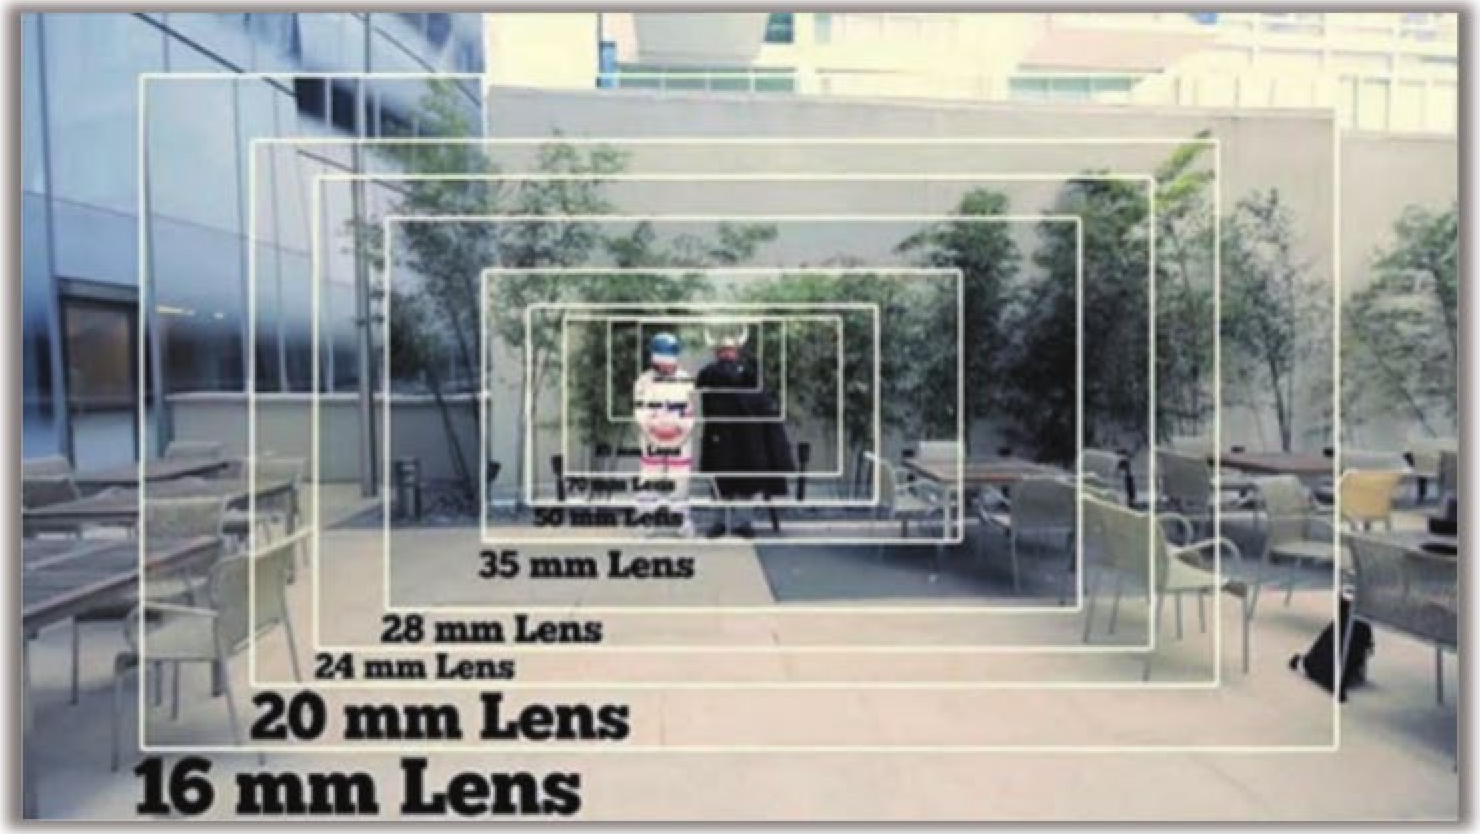

Andrew McWilliams

Field of view

Increasing the lens focal length is similar to cropping

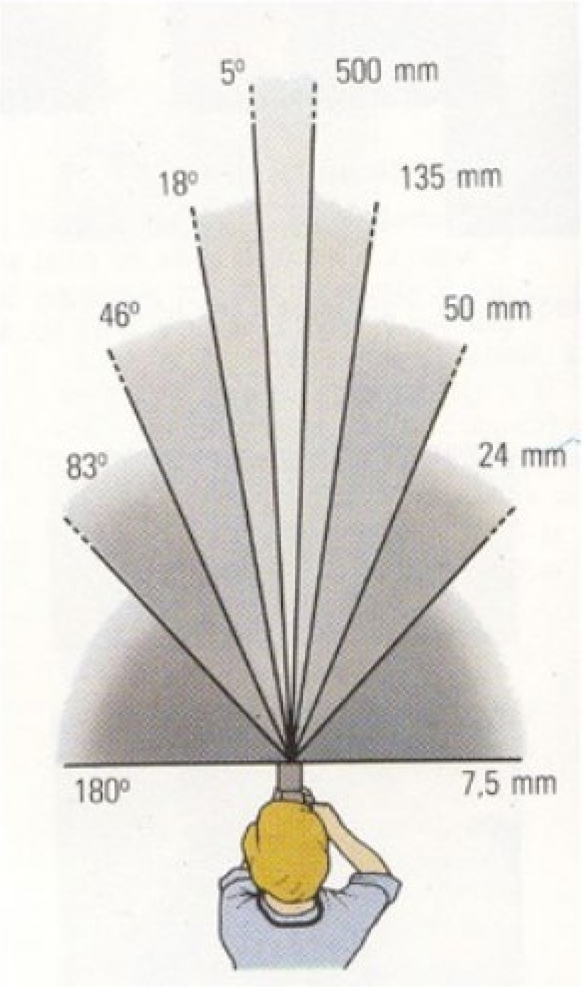

f = 25 mm

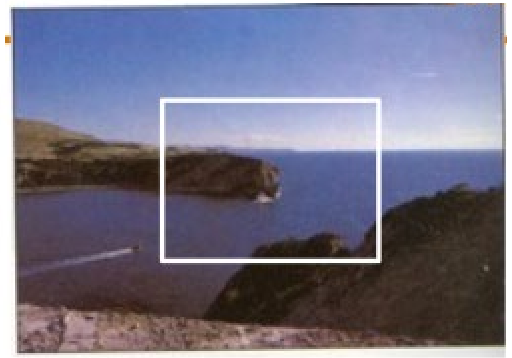

f = 50 mm

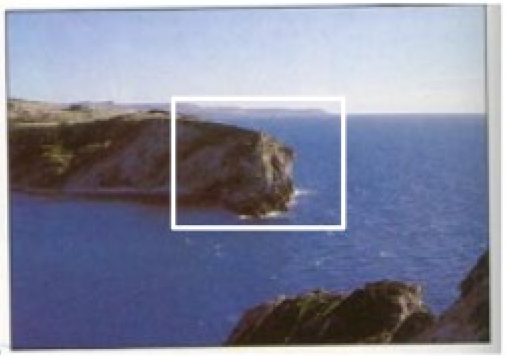

f = 135 mm

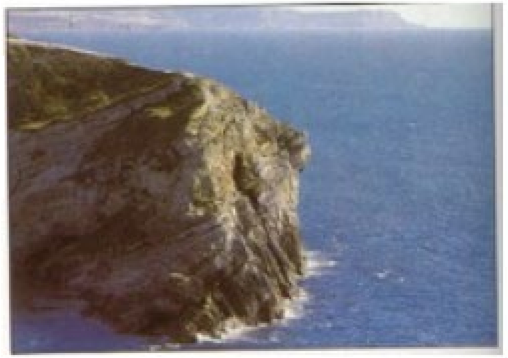

Is this effect identical to cropping?

How is zoom actually implemented?

Compound lenses

Many pieces of glass manufactured to have a specific shape, and placed in a specific configuration

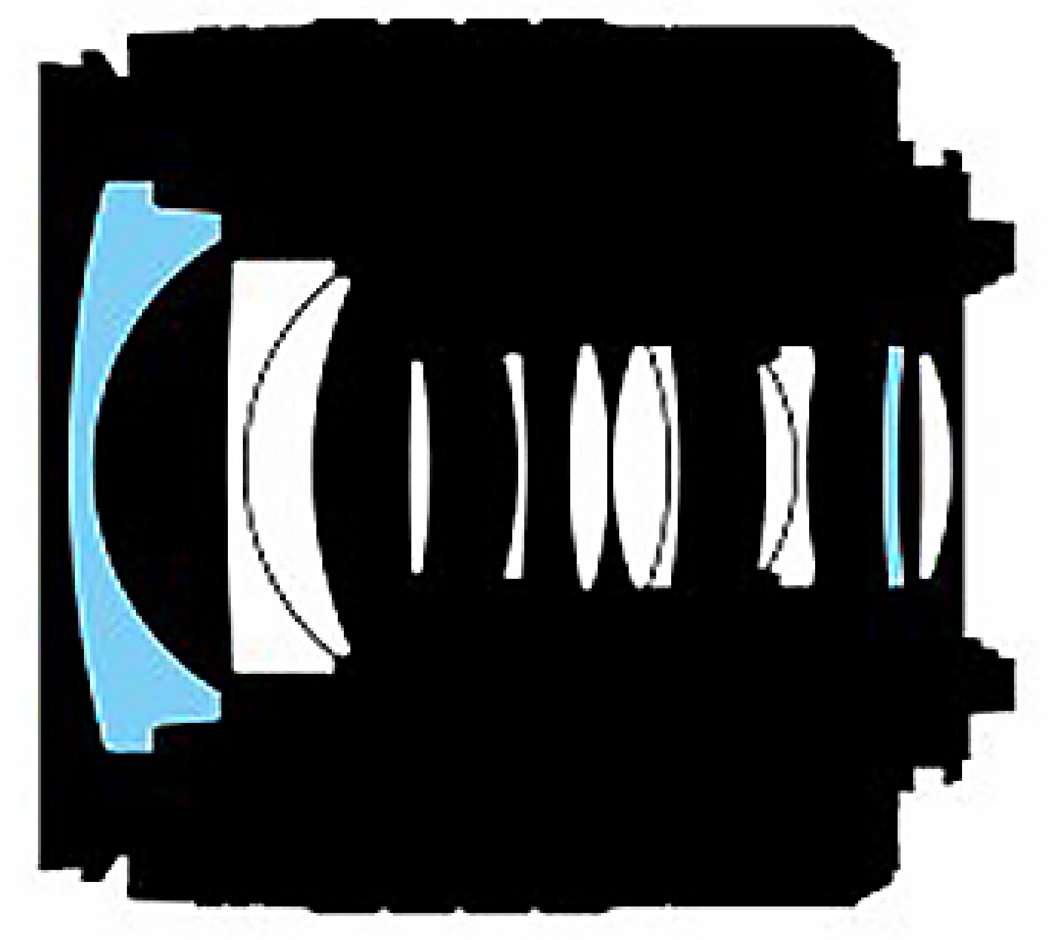

Cross-section of Nikon 18-55 mm lens

You can change the effective focal length of the overall compound lens by changing the relative placement of the individual lenses inside the lens tube.

The lens on your camera

Focus ring: controls distance of lens from sensor

Zoom ring: controls focal length of lens

Tunable (or deformable) lenses

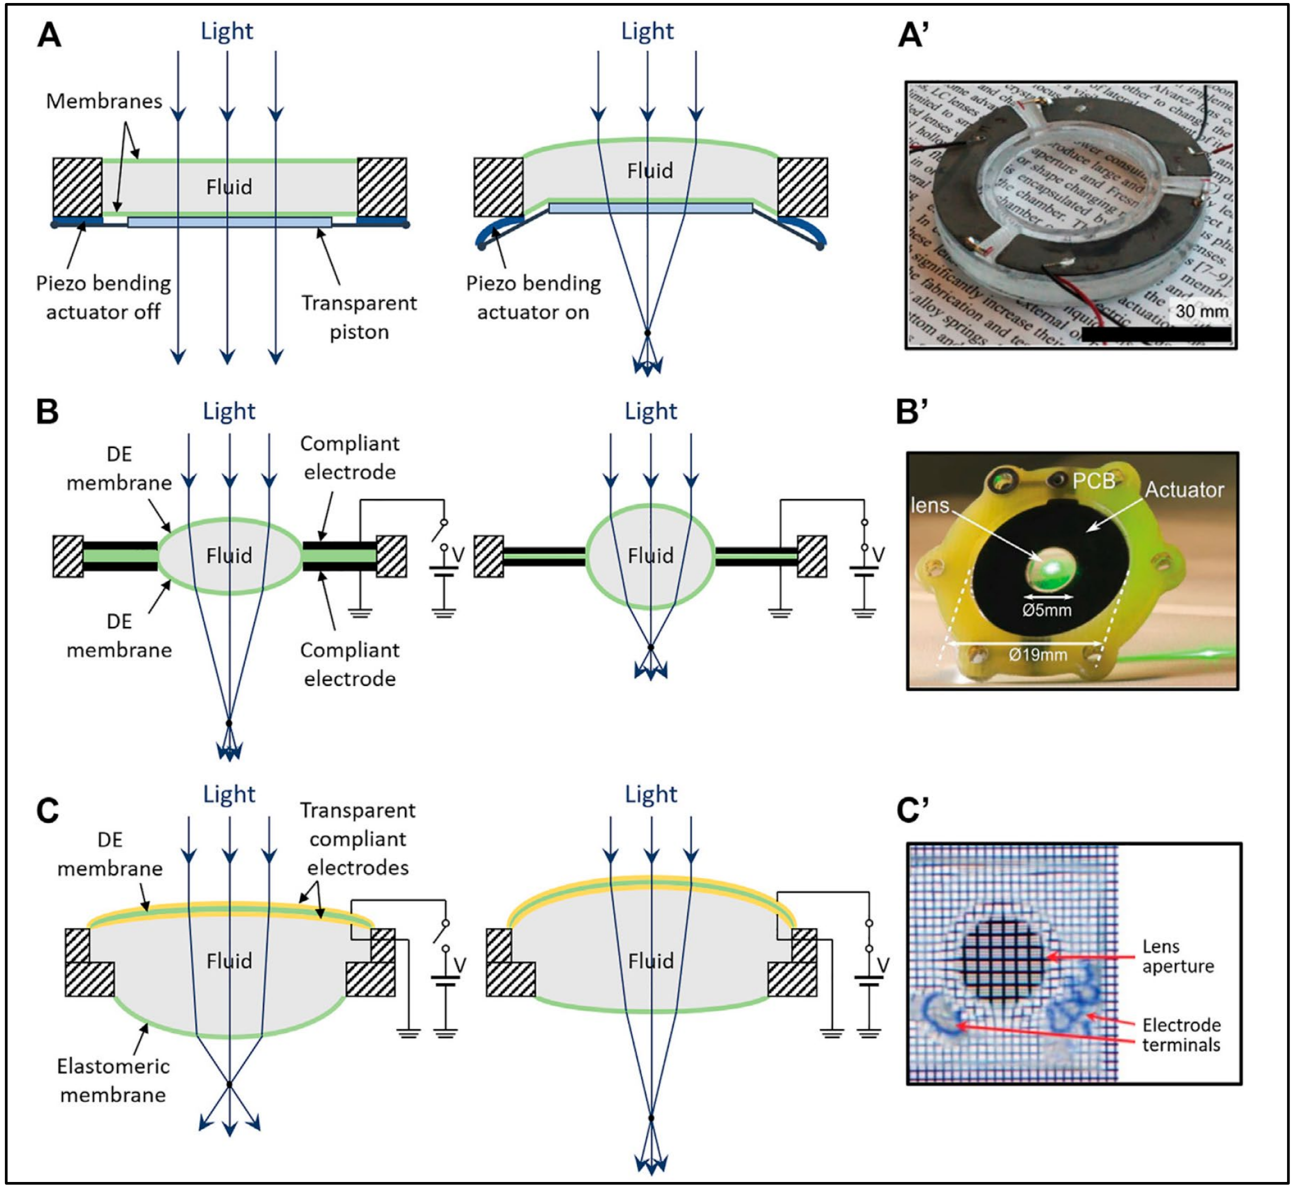

Use different processes (electric, magnetic, acoustic) to change the shape of a liquid lens (e.g., water droplet).

Programmable lenses!

Focusing versus zooming

When you turn the focus ring to bring lens further-away from the sensor:

1. The in-focus distance decreases (you need to get closer to object).
2. The field of view decreases (you see a smaller part of the object).
3. The magnification increases (same part of the object is bigger on sensor).

When you turn the zoom ring to decrease the focal length of the lens:

1. The in-focus distance increases (you need to move away from the object).
2. The field of view increases (you see a larger part of the object).
3. The magnification decreases (same part of the object is smaller on sensor).

Focusing versus zooming

When you turn the focus ring to bring lens further-away from the sensor:

1. The in-focus distance decreases (you need to get closer to object).
2. The field of view decreases (you see a smaller part of the object).
3. The magnification increases (same part of the object is bigger on sensor).

When you turn the zoom ring to decrease the focal length of the lens:

1. The in-focus distance increases (you need to move away from the object).
2. The field of view increases (you see a larger part of the object).
3. The magnification decreases (same part of the object is smaller on sensor).

We can use both focus and zoom to cancel out their effects.

What if…

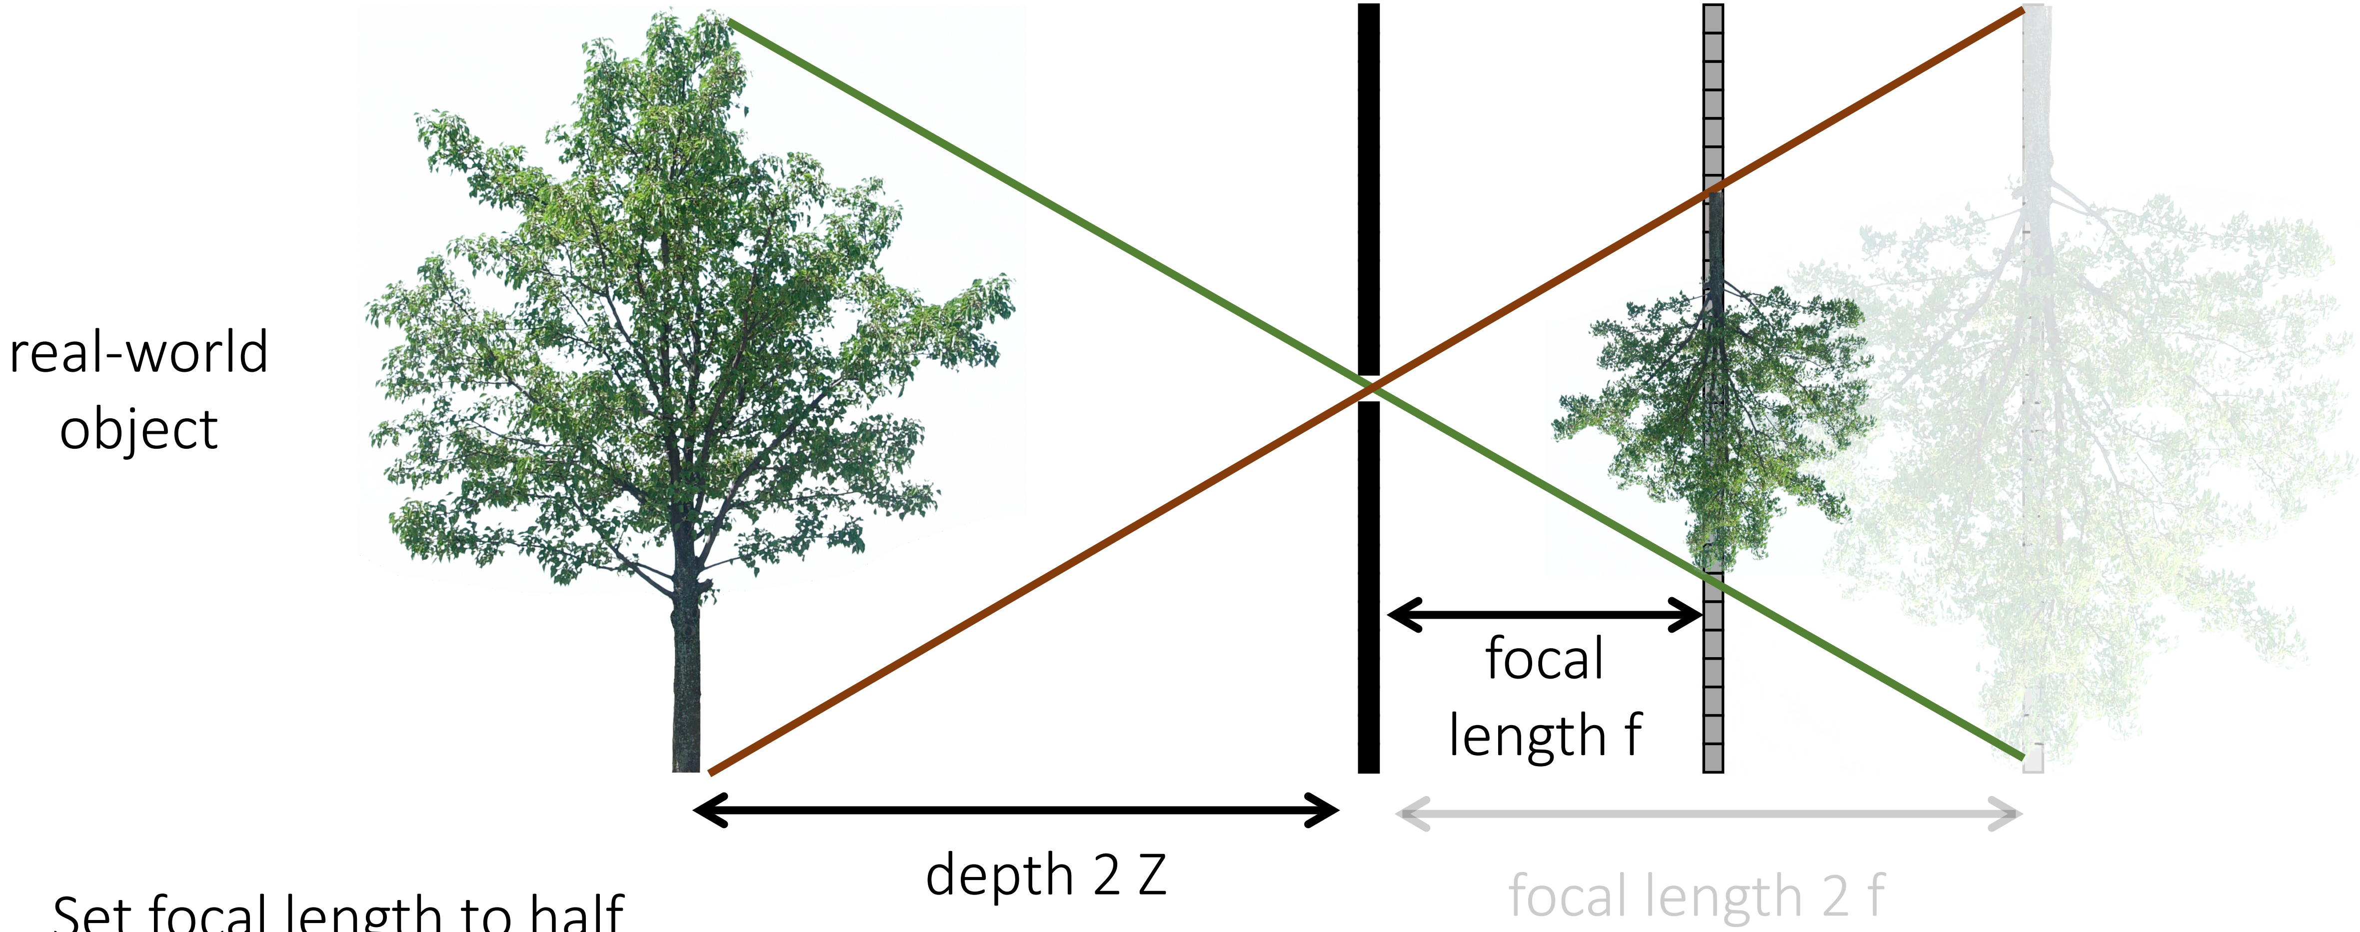

1. Set focal length to half

What if…

real-world object

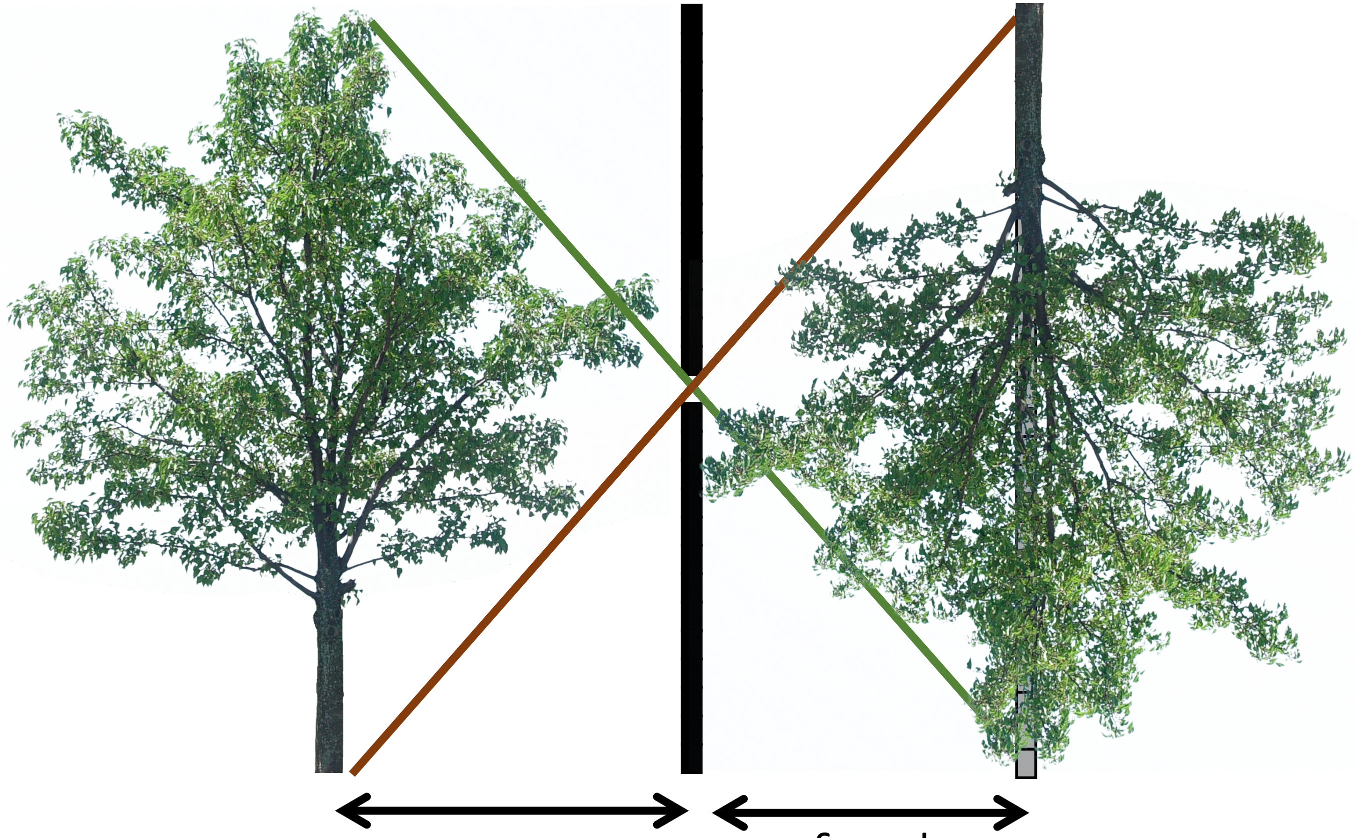

Is this the same image as the one we had at focal length 2f and distance 2Z?

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

Similar construction can be done with lenses, after taking care of refocusing.

Perspective distortion

long focal length

mid focal length

short focal length

Perspective distortion

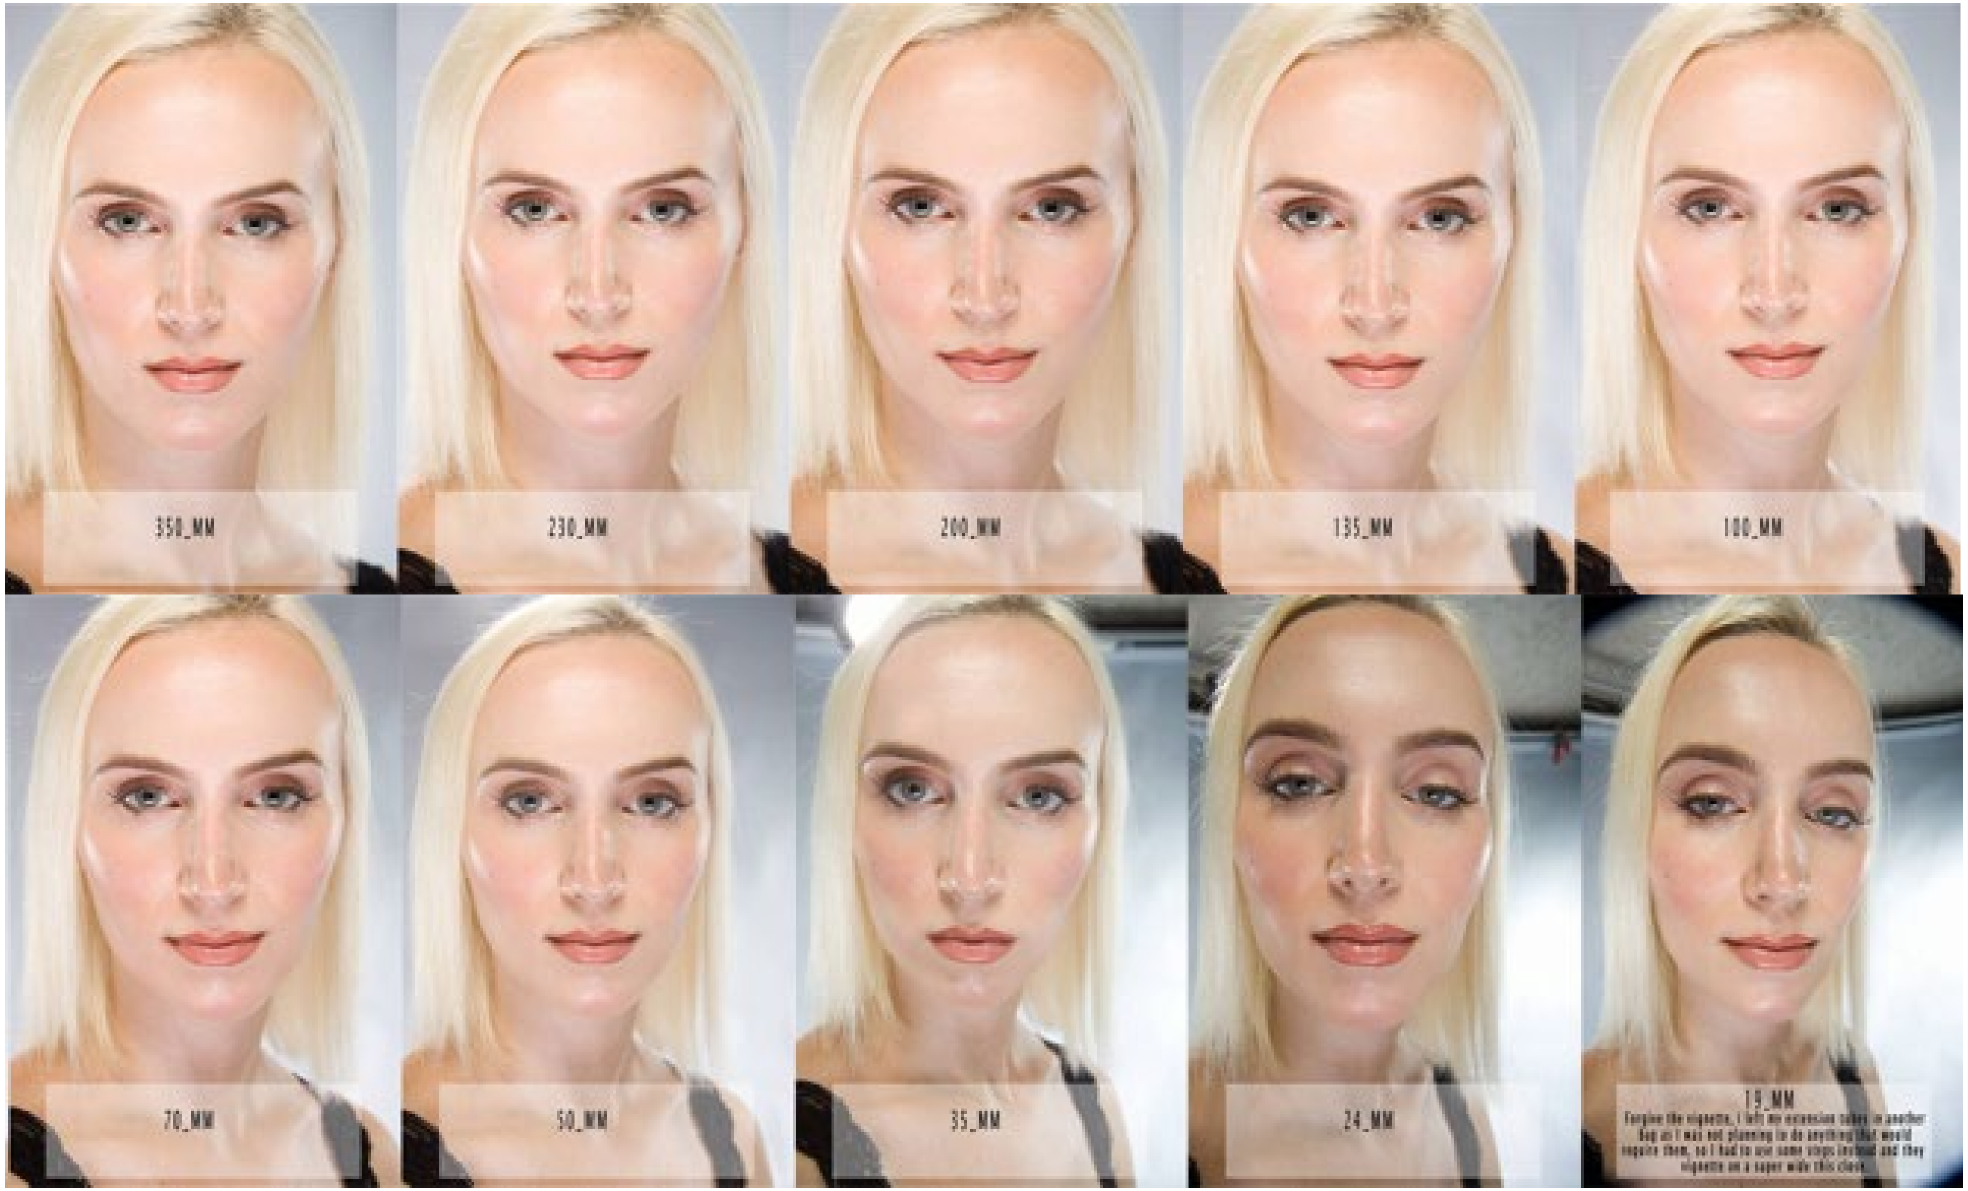

What is the best focal length for portraits?

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"

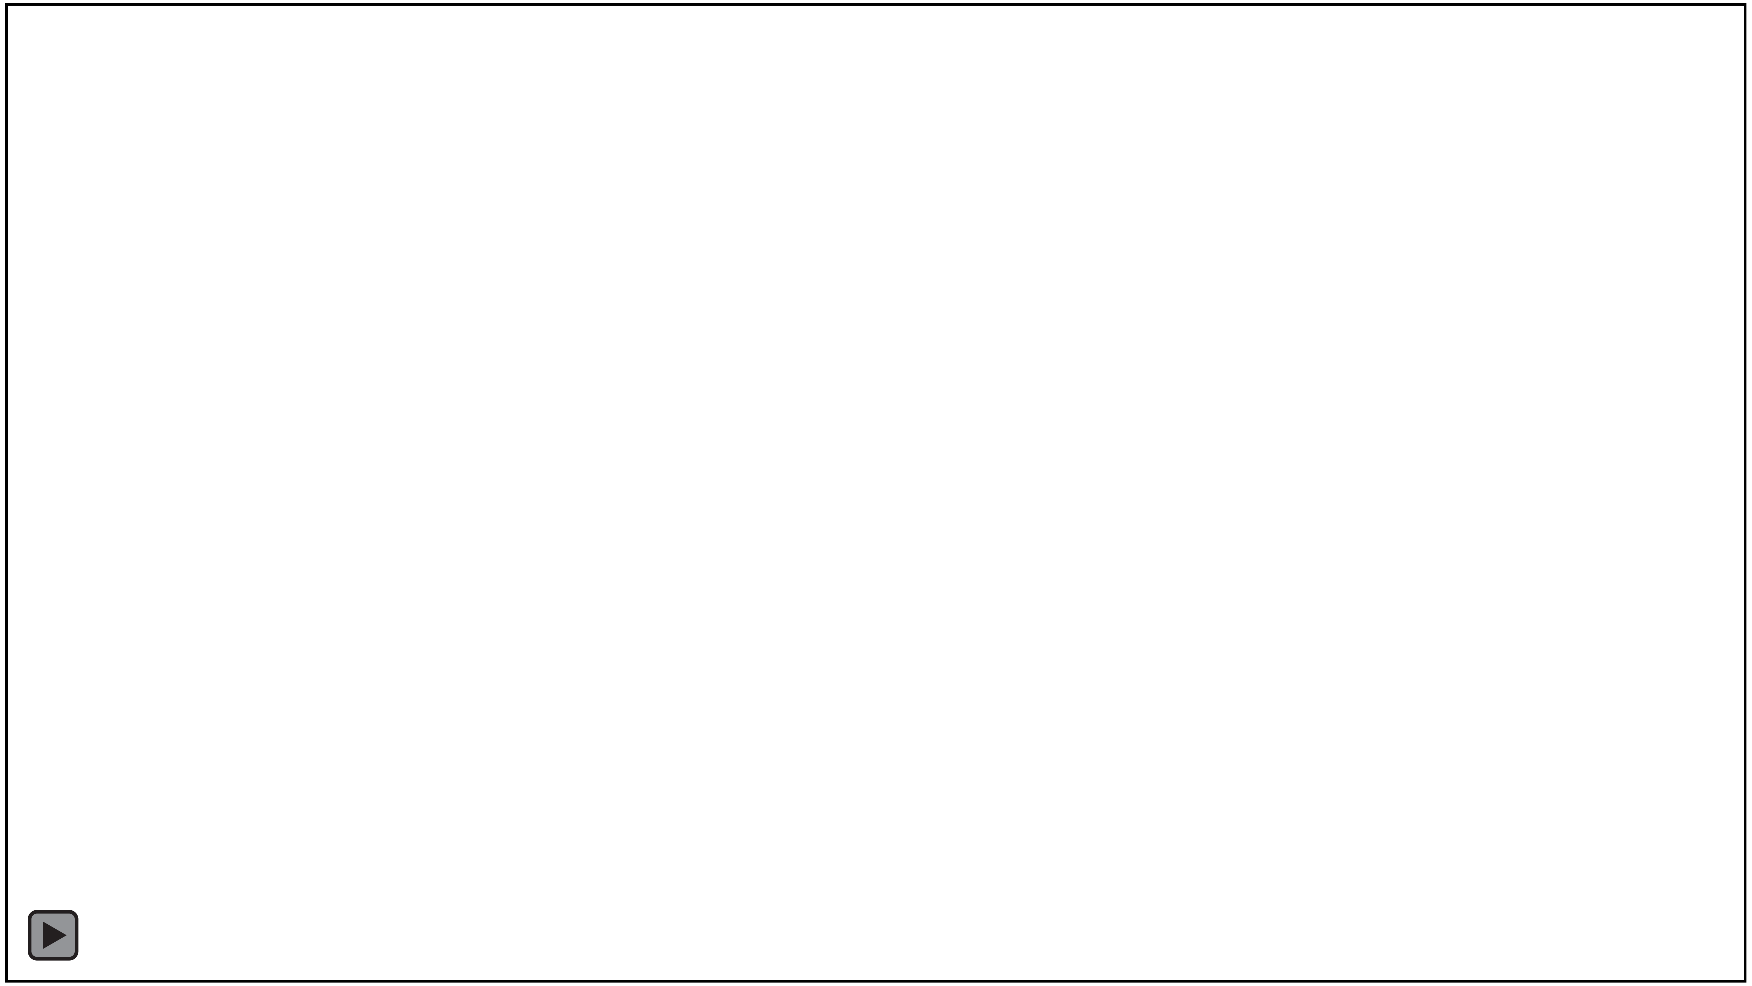

Vertigo effect

How would you
create this effect?

Orthographic camera and telecentric lenses

Can we make magnification depth-independent?

Orthographic camera

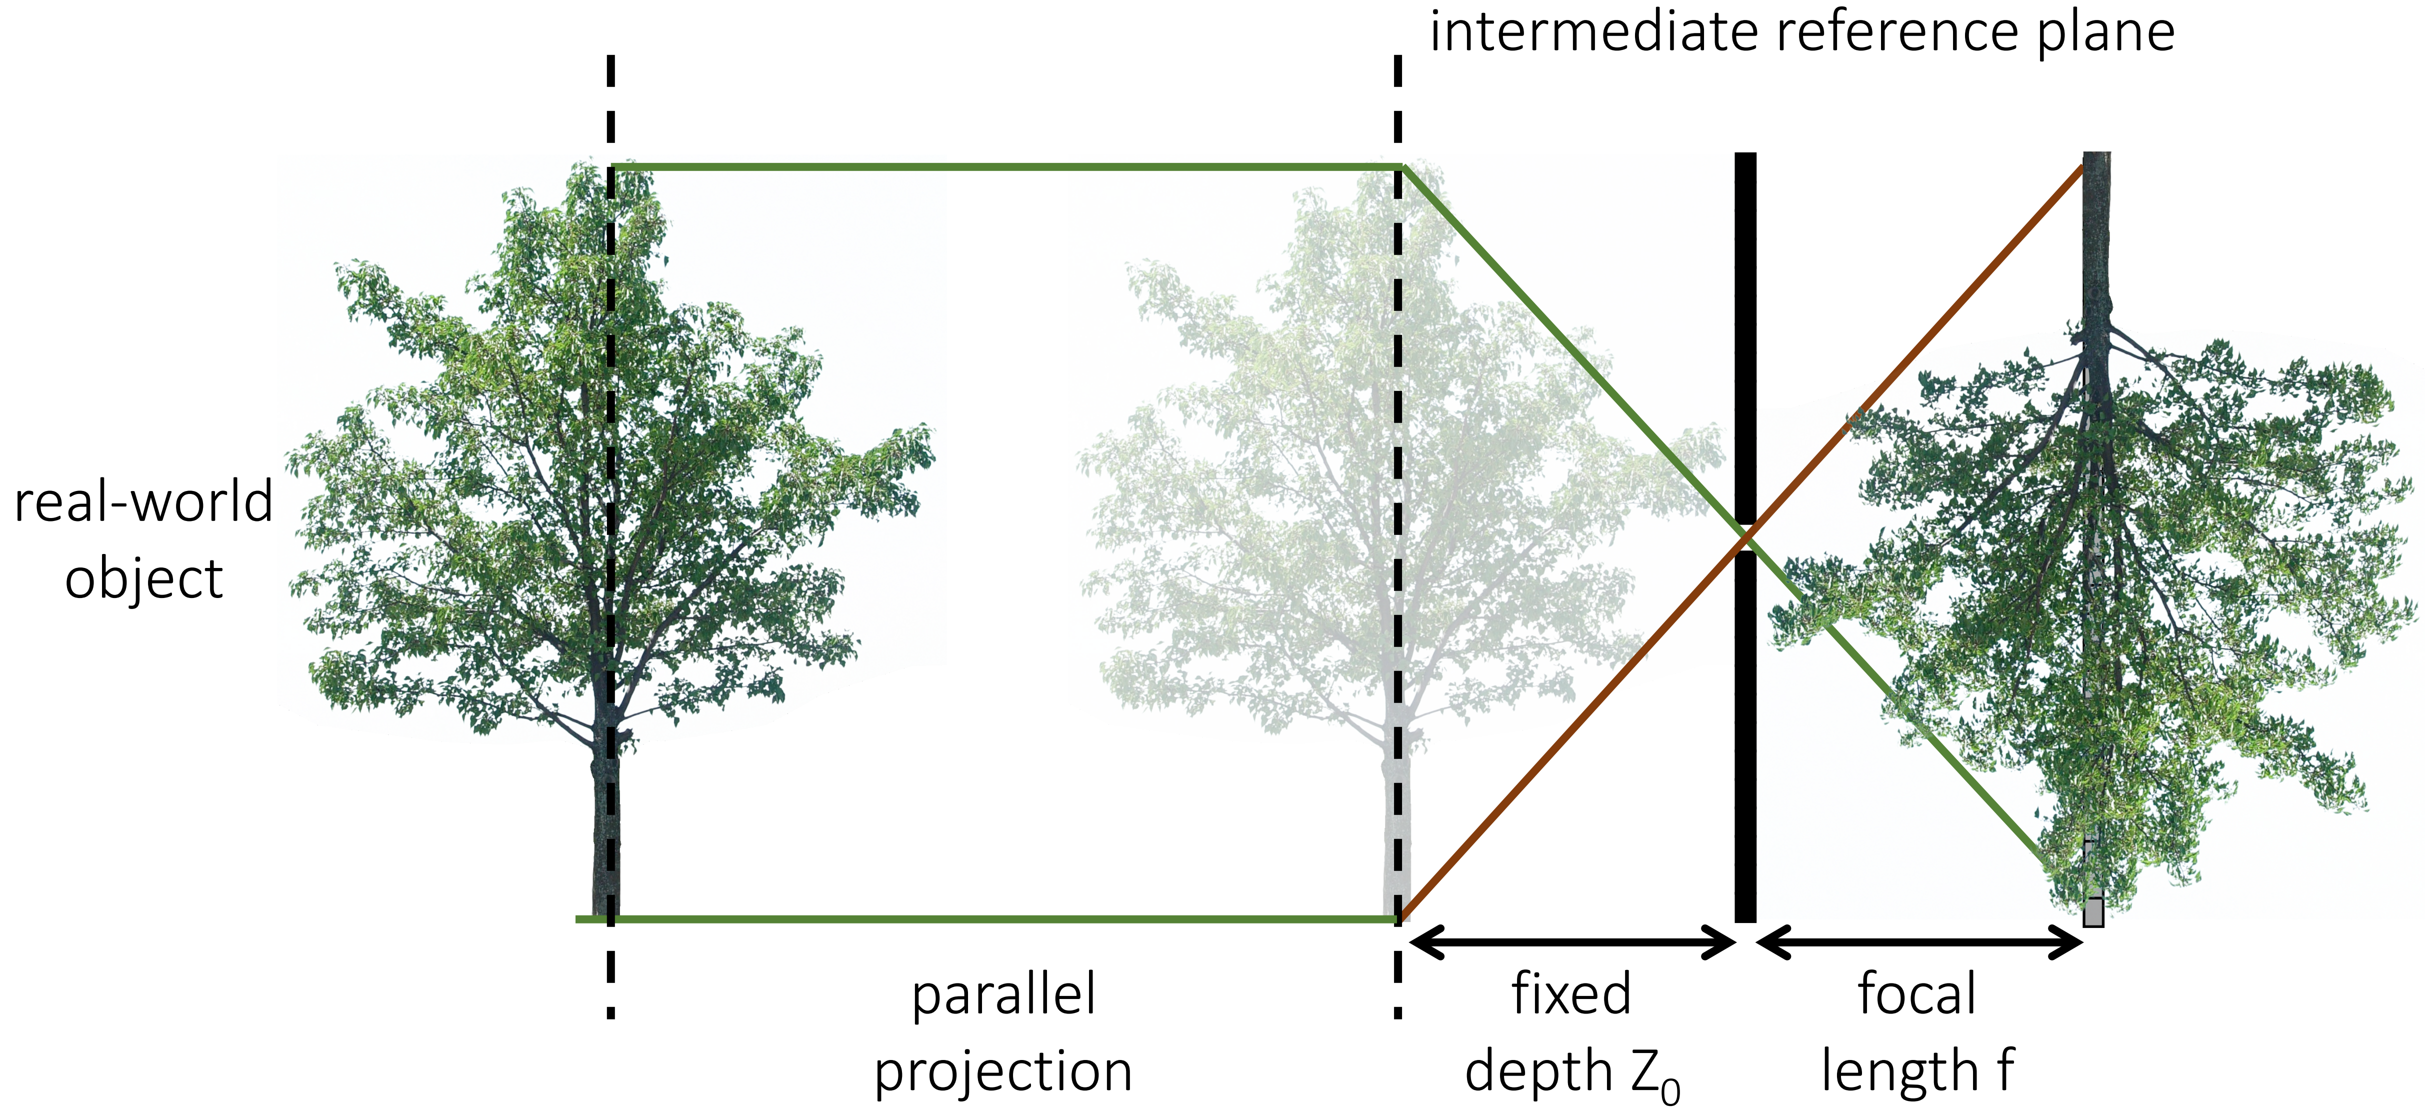

Depth-independent magnification m = 1 (real-life size).

Weak-perspective camera

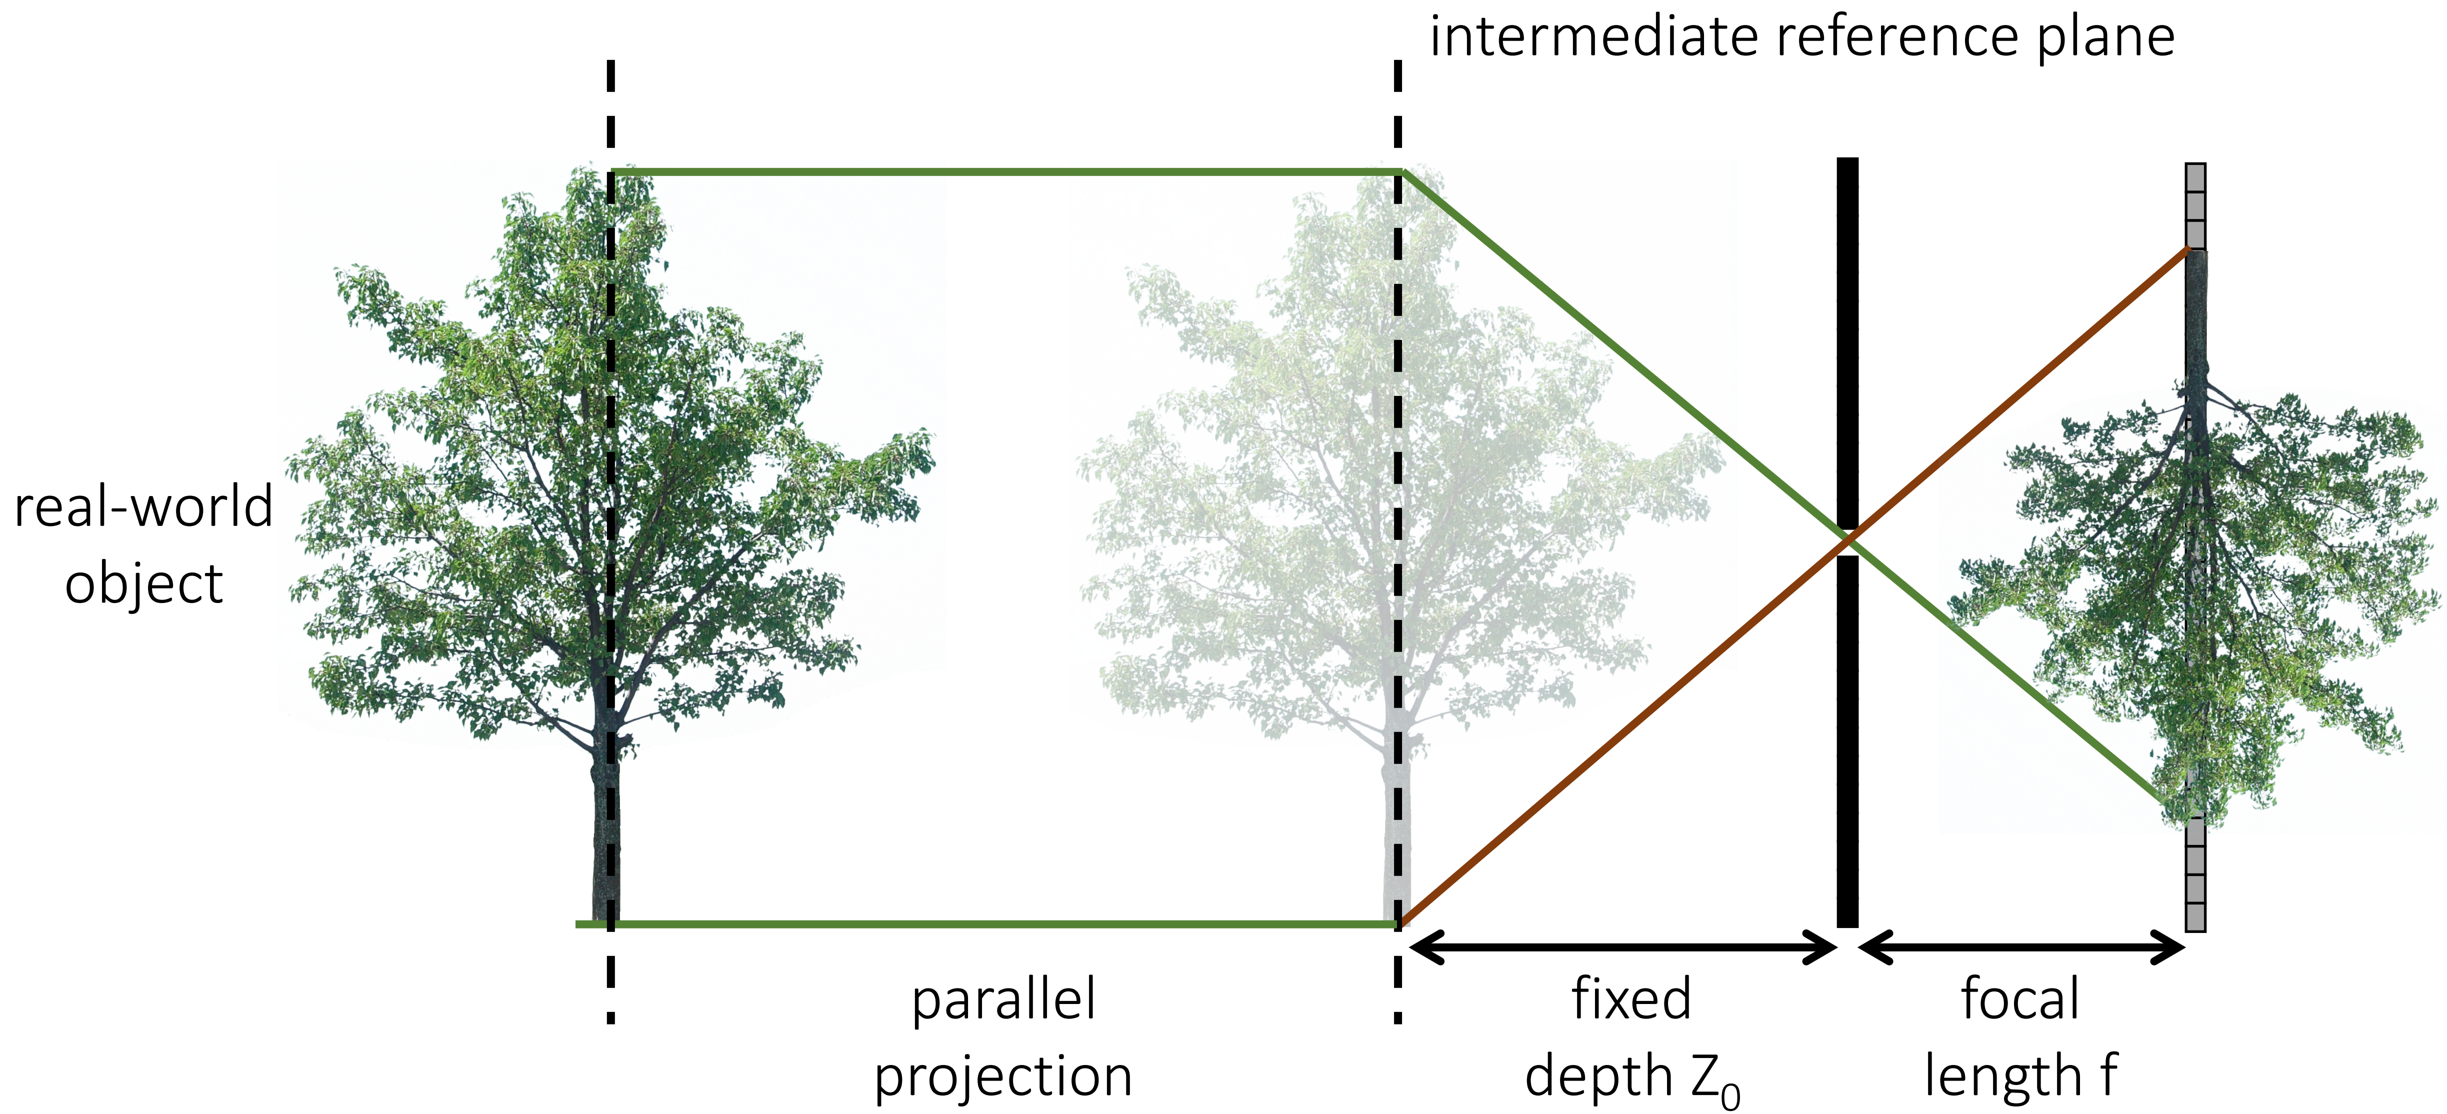

Depth-independent magnification $m = f / Z_0$.

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.

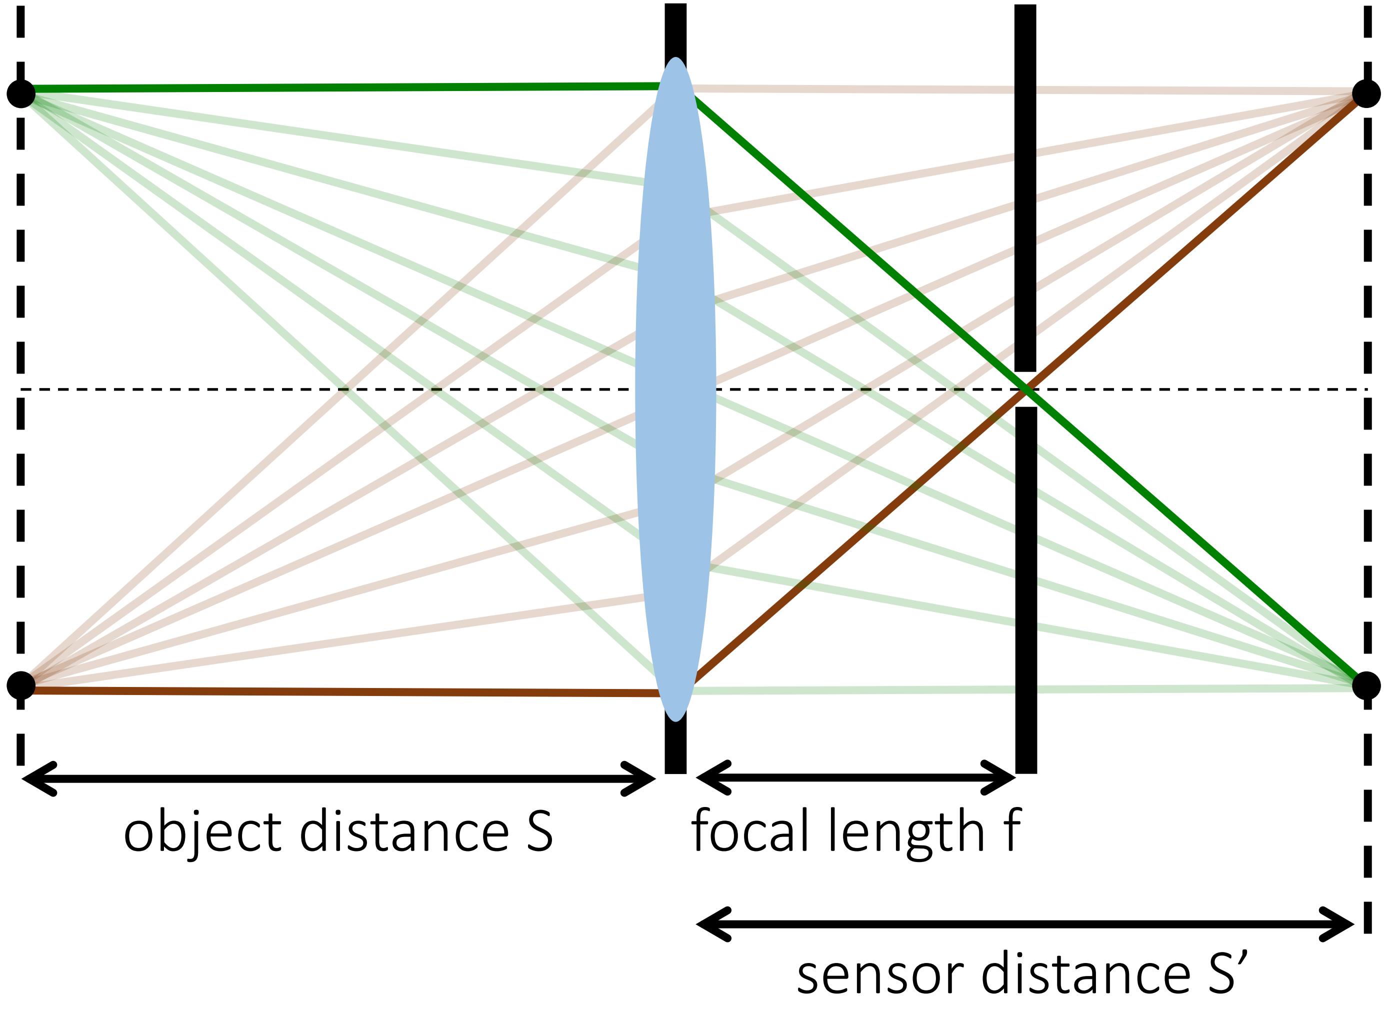

Telecentric lens

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

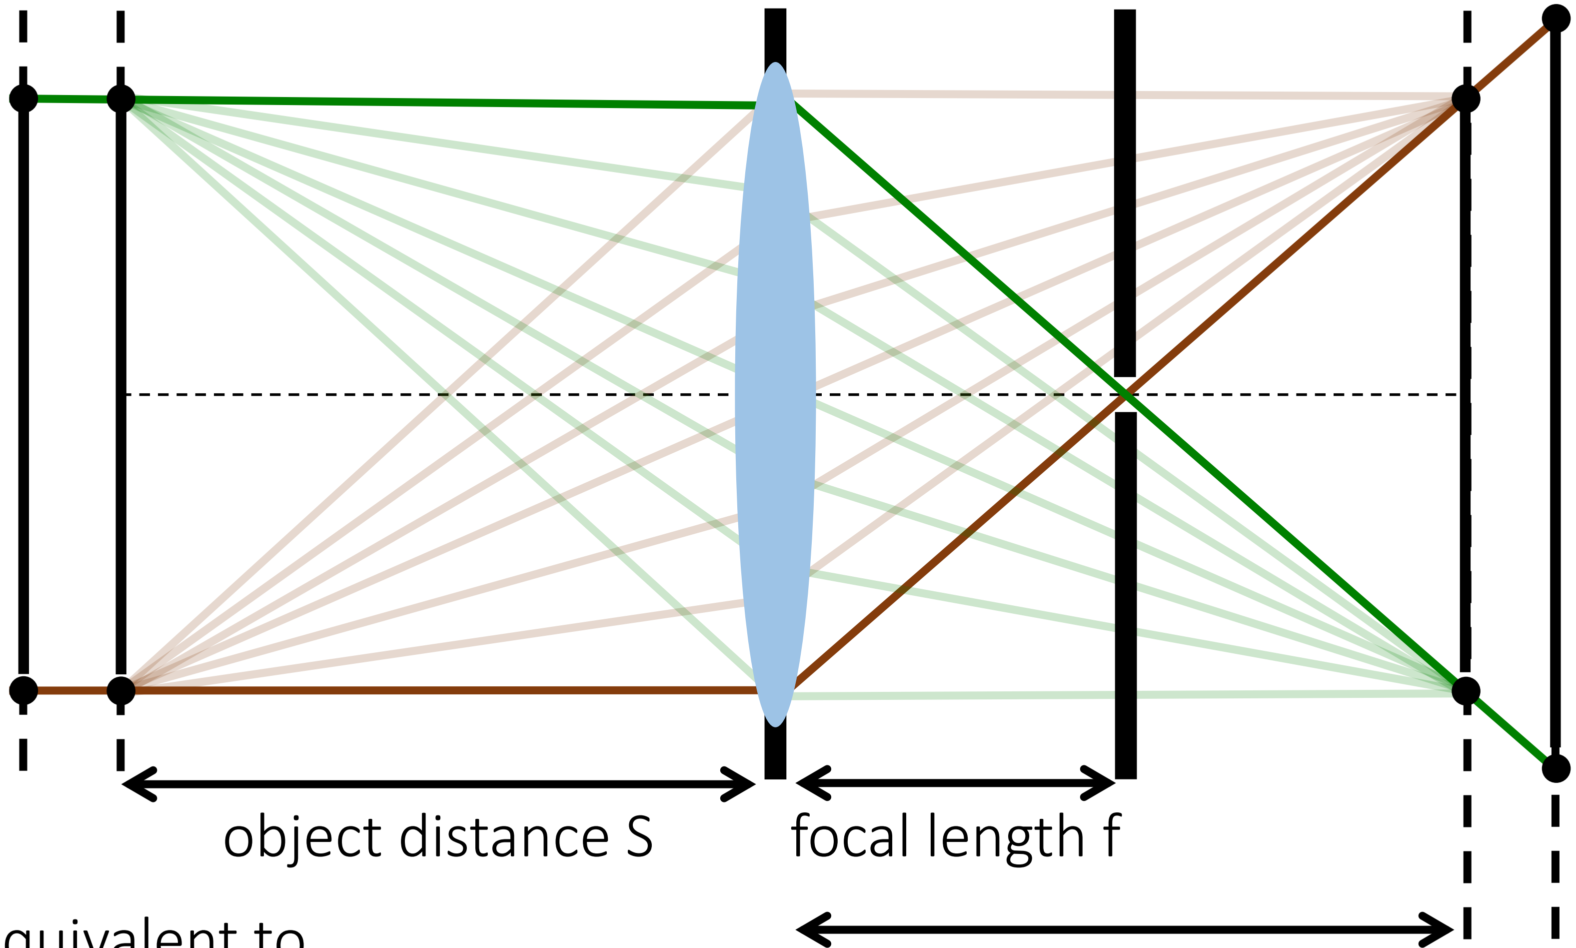

Magnification independent of object depth.

Magnification depends only on sensor-lens distance S'.

When is this lens equivalent to an orthographic camera?

Regular vs telecentric lens

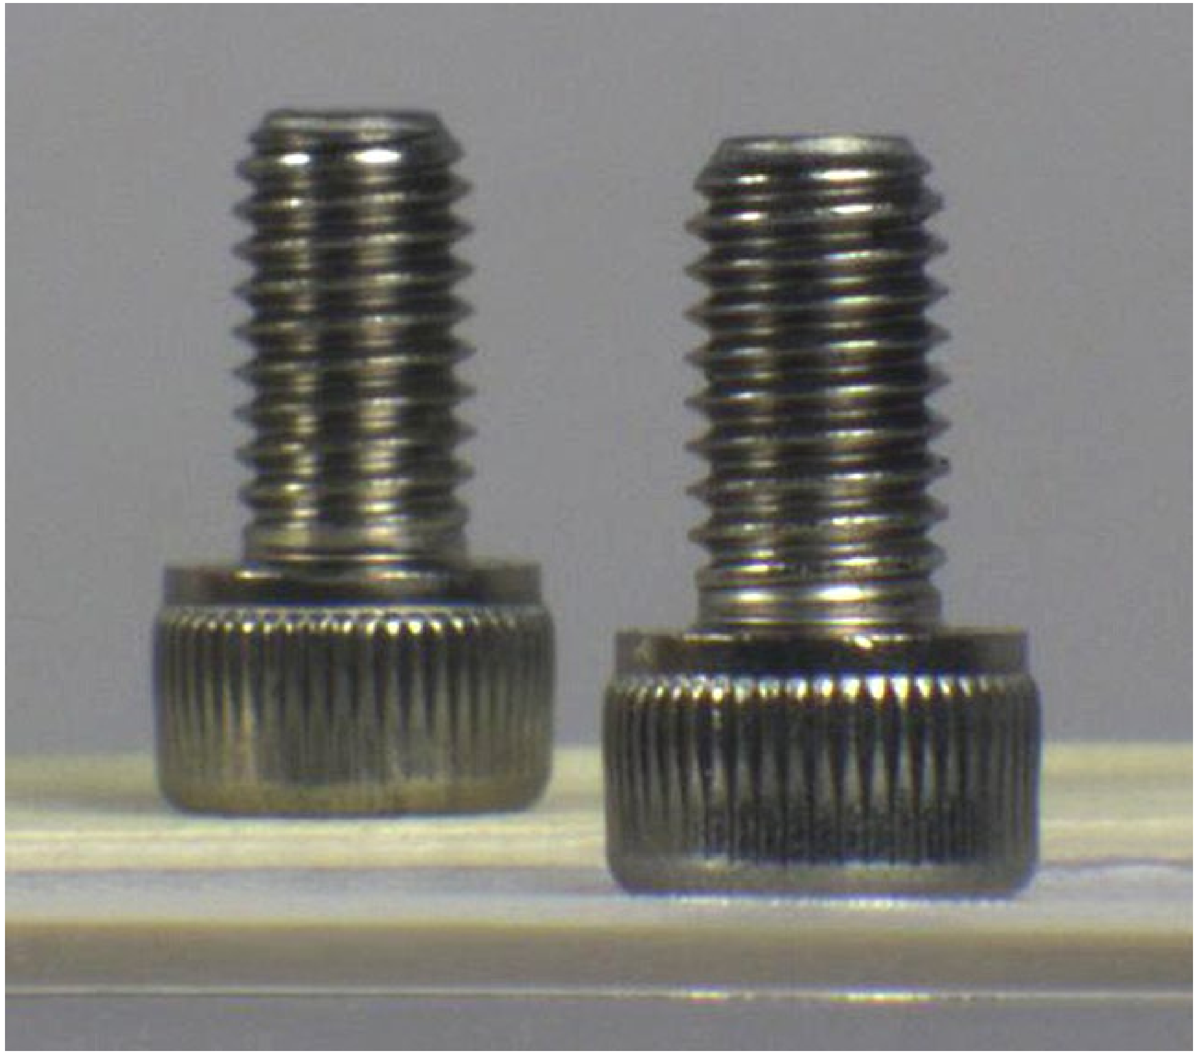

regular lens

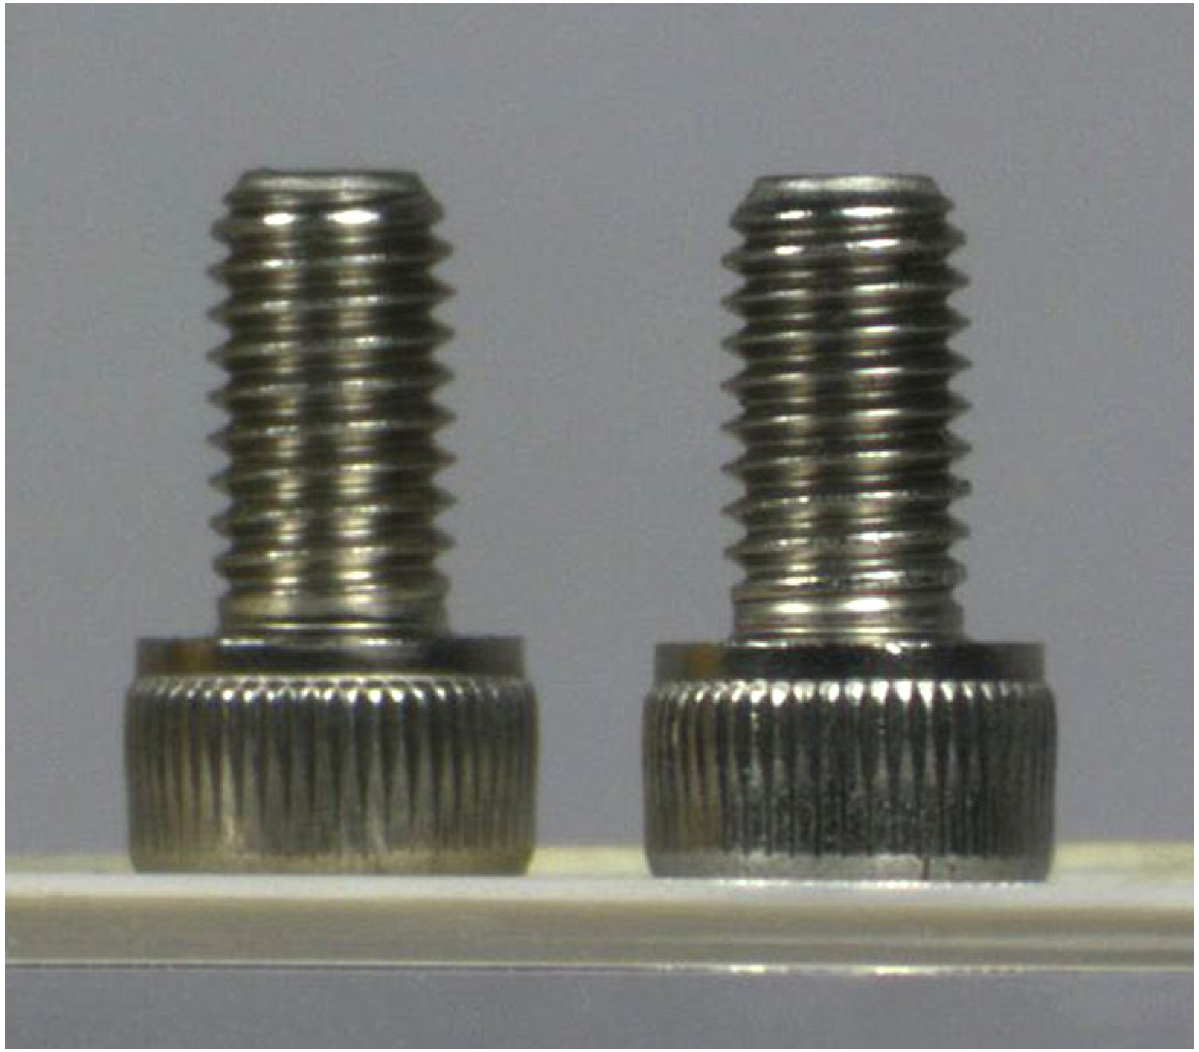

telecentric lens

References

Basic reading:

- Szeliski textbook, Section 2.1.5, 2.2.3.
- Pedrotti, Pedrotti, and Pedrotti, Introduction to Optics.
  Chapters 2 and 3 have a detailed overview of basic geometric optics and lenses.

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.
- Watanabe and Nayar, "Telecentric Optics for Focus Analysis," PAMI 1997.
  An early computational photography paper analyzing the ray optics and explaining the advantages and disadvantages of telecentric lenses relative to conventional lenses.
- Chen et al., "Electrically Tunable Lenses: A Review," Frontiers in Robotics and AI 2021.
  A great review of electrically tunable lenses.