

# Pinholes and lenses



15-463, 15-663, 15-862  
Computational photography  
Fall 2026, lecture 3

# Course announcements

- Homework assignment 1 is out.
  - Due September 11<sup>th</sup>.
  - Start early! Second part takes a lot of time and requires considerable handiwork.
  - Any issues with homework assignment 1?
- Office hours for rest of the semester on course website.

# 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


real-world  
object



digital sensor  
(CCD or CMOS)




# Bare-sensor imaging






# Bare-sensor imaging



# Bare-sensor imaging





# 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





Each scene point contributes to only one sensor pixel

What does the image on the sensor look like?



# 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

real-world  
object



barrier (diaphragm)




pinhole  
(aperture)



digital sensor  
(CCD or CMOS)




# Pinhole camera terms





# Focal length

real-world  
object




focal length  $f$

# Focal length


What happens as we change the focal length?

real-world  
object



# Focal length

What happens as we change the focal length?




# Focal length

What happens as we change the focal length?

object projection is half the size

real-world  
object





# Pinhole size

real-world  
object



pinhole  
diameter




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


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.



# 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



pinhole  
diameter




focal length  $f$

- 2x pinhole diameter  $\rightarrow$  4x light
- 2x focal length  $\rightarrow$   $\frac{1}{4}$ x 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

projected pattern on the wall




upside down



window with smaller gap




view outside window



window is an  
aperture


# Accidental pinspeck camera




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?

# 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




# 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




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.



Cross-section of Nikon 18-55 mm lens

# 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 well-designed lenses.



# The reality: compound lenses

Remember, real lenses are not “thin” lenses.




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

Cross-section of Nikon 18-55 mm 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?




# 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?



# 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?




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?



# 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

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


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

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



# Special focus distances

$S' = f, S = \infty, m = 0 \rightarrow$  infinity focus (parallel rays)




$S' = S = ?, m = ?$

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

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

# Special focus distances


$S' = f, S = \infty, m = 0 \rightarrow$  infinity focus (parallel rays)




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

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

$S' = S = 2f, m = 1 \rightarrow$  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?

where we need to move sensor  
for squirrel to be in focus




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.






# 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.




# 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




$$c < \varepsilon \Rightarrow \text{DOF} = \frac{2\varepsilon O}{mD}$$

Note: in reality,  
DOF is slightly  
asymmetrical.



# Depth of field

depth of field




scene

# Circle of Confusion

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

Canon 5D Mark III:  $f=50\text{mm}$ ,  $f/2.8$  ( $N=2.8$ ),

focused at  $5\text{m}$ , pixel size= $7.5\mu\text{m}$



# Defocus depends on aperture

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





Diagram illustrating the circle of confusion in a camera model. A point source (black dot) emits rays (green lines) that pass through a lens (blue oval) and converge at a point on the right. The distance from the source to the lens is the in-focus object distance  $S$ . The distance from the lens to the sensor is the focal length  $f$ . The distance from the source to the sensor is the actual object distance  $O$ . The sensor is at a distance  $S'$  from the lens. The circle of confusion  $c$  is the diameter of the blur spot on the sensor. The diagram shows that when the object is not at the focal distance, the rays do not converge exactly at the sensor, creating a blur spot.

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




# Defocus depends on aperture diameter

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



# Defocus depends on aperture diameter

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



# 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?






# 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




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



move object further  
away from lens


move lens closer  
to sensor

- 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 lens 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



$f = 25 \text{ mm}$



$f = 50 \text{ mm}$



$f = 135 \text{ mm}$




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



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.

Cross-section of Nikon 18-55 mm lens


# 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...


real-world  
object



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  $2Z$ ?

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

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


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.




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.