

Pinhole cameras



15-463, 15-663, 15-862
Computational Photography
Fall 2017, Lecture 14

Course announcements

- Homework 4 is out.
 - Due October 26th.
 - Bilateral filter will take a very long time to run.
 - Make sure to sign up for a camera and a team.
 - Drop by Yannis' office to pick up cameras any time.
- Yannis has extra office hours on Wednesday, 2-4pm.
 - You can come to ask questions about HW4 (e.g., "how do I use a DSLR camera?").
 - You can come to ask questions about final project.
- Project ideas are due on Piazza on Friday 20th.

Overview of today's lecture

- Some motivational imaging experiments.
- Pinhole camera.
- Accidental pinholes.
- Camera matrix.
- Perspective.
- Orthographic camera.

Slide credits

Most of these slides were adapted from:

- Kris Kitani (15-463, Fall 2016).

Some slides inspired from:

- Fredo Durand (MIT).

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

real-world
object




digital sensor
(CCD or CMOS)




Bare-sensor imaging


real-world
object



digital sensor
(CCD or CMOS)



Bare-sensor imaging




Bare-sensor imaging



All scene points contribute to all sensor pixels

Let's add something to this scene



What would an image taken like this look like?

Pinhole imaging


real-world
object




most rays
are blocked

one makes
it through


digital sensor
(CCD or CMOS)



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

real-world
object



barrier (diaphragm)



pinhole
(aperture)



camera center
(center of projection)


image plane



digital sensor
(CCD or CMOS)

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?



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?

real-world
object




pinhole
diameter



Pinhole size

What happens as we change the pinhole diameter?


real-world
object



Pinhole size

What happens as we change the pinhole diameter?

real-world
object




Pinhole size

What happens as we change the pinhole diameter?


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.



small pinhole


wide
diffraction
pattern



large pinhole

narrow
diffraction
pattern

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

Can we do better than pinhole imaging?



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

Camera matrix

The camera as a coordinate transformation



The camera as a coordinate transformation

A camera is a mapping from:

the 3D world

to:

a 2D image

homogeneous coordinates

The diagram shows the equation $x = PX$. Above the equation, the text "homogeneous coordinates" has two arrows pointing to the x and X terms. Below the equation, the terms are labeled: "2D image point" under x , "camera matrix" under P , and "3D world point" under X .

$$x = PX$$

2D image point camera matrix 3D world point

What are the dimensions of each variable?

The camera as a coordinate transformation

$$\boldsymbol{x} = \mathbf{P}\mathbf{X}$$


$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

homogeneous
image coordinates
3 x 1


camera
matrix
3 x 4

homogeneous
world coordinates
4 x 1


The pinhole camera



The (rearranged) pinhole camera




The (rearranged) pinhole camera




What is the equation for image coordinate $\mathbf{x}$ in terms of $\mathbf{X}$?

The 2D view of the (rearranged) pinhole camera




What is the equation for image coordinate $\mathbf{x}$ in terms of $\mathbf{X}$?

The 2D view of the (rearranged) pinhole camera



The (rearranged) pinhole camera



What is the camera matrix $\mathbf{P}$ for a pinhole camera?

$$\mathbf{x} = \mathbf{P}\mathbf{X}$$

The pinhole camera matrix

Relationship from similar triangles:

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

General camera model:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

What does the pinhole camera projection look like?

$$\mathbf{P} = \begin{bmatrix} ? & ? & ? & ? \\ ? & ? & ? & ? \\ ? & ? & ? & ? \end{bmatrix}$$

The pinhole camera matrix

Relationship from similar triangles:

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

General camera model:


$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

What does the pinhole camera projection look like?

$$\mathbf{P} = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Generalizing the camera matrix

Camera origin and image origin might be different




How does the camera matrix change?

$$\mathbf{P} = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Generalizing the camera matrix

Camera origin and image origin might be different



How does the camera matrix change?

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x & 0 \\ 0 & f & p_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Camera matrix decomposition

We can decompose the camera matrix like this:

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & | & 0 \\ 0 & 1 & 0 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix}$$

intrinsic (3 x 3) extrinsic (3 x 4)

$$\mathbf{P} = \mathbf{K}[\mathbf{I}|\mathbf{0}]$$

$$\mathbf{K} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \quad \text{calibration matrix}$$



Extrinsic camera parameters

We can decompose the camera matrix like this:

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & | & 0 \\ 0 & 1 & 0 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix} \leftarrow \begin{array}{l} \text{assumes camera and} \\ \text{world share the same} \\ \text{coordinate system} \end{array}$$

intrinsic (3 x 3) extrinsic (3 x 4)

What if world and camera coordinate systems are different?




Extrinsic camera parameters

We can decompose the camera matrix like this:

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & | & 0 \\ 0 & 1 & 0 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix} \leftarrow \begin{array}{l} \text{assumes camera and} \\ \text{world share the same} \\ \text{coordinate system} \end{array}$$

intrinsic (3 x 3) extrinsic (3 x 4)

What if world and camera coordinate systems are different?



Extrinsic camera parameters

We can decompose the camera matrix like this:

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & | & 0 \\ 0 & 1 & 0 & | & 0 \\ 0 & 0 & 1 & | & 0 \end{bmatrix} \leftarrow \begin{array}{l} \text{assumes camera and} \\ \text{world share the same} \\ \text{coordinate system} \end{array}$$

intrinsic (3 x 3) extrinsic (3 x 4)

What if world and camera coordinate systems are different?

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R} & -\mathbf{RC} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} X_w \\ Y_w \\ Z_w \\ 1 \end{bmatrix}$$

Extrinsic camera parameters

We can decompose the camera matrix like this:

$$\mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R} & | & -\mathbf{RC} \\ \mathbf{0} & | & 1 \end{bmatrix}$$

intrinsic (3 x 3) extrinsic (3 x 4)

What if world and camera coordinate systems are different?

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R} & -\mathbf{RC} \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} X_w \\ Y_w \\ Z_w \\ 1 \end{bmatrix}$$

General pinhole camera matrix

We can decompose the camera matrix like this:

$$\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$$

(translate first then rotate)

Another way to write the mapping:


$$\mathbf{P} = \mathbf{K}[\mathbf{R} | \mathbf{t}]$$

where $\mathbf{t} = -\mathbf{R}\mathbf{C}$

(rotate first then translate)

Recap


What is the size and meaning of each term in the camera matrix?

$$\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$$


The diagram illustrates the components of the camera matrix equation $\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$. Four arrows point from question marks below to the terms $\mathbf{K}$, $\mathbf{R}$, $[\mathbf{I} | -\mathbf{C}]$, and the overall matrix $\mathbf{P}$, indicating a query about their sizes and meanings.

Recap

What is the size and meaning of each term in the camera matrix?

$$\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$$


The diagram shows four arrows pointing from labels below to terms in the equation $\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$. The first arrow points from '3x3 intrinsics' to $\mathbf{K}$. The second arrow points from '?' to $\mathbf{R}$. The third arrow points from '?' to $\mathbf{I}$. The fourth arrow points from '?' to $\mathbf{C}$.

3x3
intrinsics


?

?

?

Recap

What is the size and meaning of each term in the camera matrix?

$$\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$$


The diagram shows four arrows pointing from labels below to terms in the equation $\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$. The first arrow points from '3x3 intrinsics' to $\mathbf{K}$. The second arrow points from '3x3 3D rotation' to $\mathbf{R}$. The third arrow points from a '?' to the $\mathbf{I}$ in the vector $[\mathbf{I} | -\mathbf{C}]$. The fourth arrow points from a '?' to $-\mathbf{C}$.

3x3
intrinsics


3x3
3D rotation

?

?

Recap

What is the size and meaning of each term in the camera matrix?

$$\mathbf{P} = \mathbf{K} \mathbf{R} [\mathbf{I} | \mathbf{C}]$$


The diagram shows four arrows pointing upwards from labels below to terms in the equation $\mathbf{P} = \mathbf{K} \mathbf{R} [\mathbf{I} | \mathbf{C}]$. The first arrow points from '3x3 intrinsics' to $\mathbf{K}$. The second arrow points from '3x3 3D rotation' to $\mathbf{R}$. The third arrow points from '3x3 identity' to $\mathbf{I}$. The fourth arrow points from '?' to $\mathbf{C}$.

3x3
intrinsics

3x3
3D rotation

3x3
identity

?

Recap

What is the size and meaning of each term in the camera matrix?


$$\mathbf{P} = \mathbf{K}\mathbf{R}[\mathbf{I} | -\mathbf{C}]$$


Diagram illustrating the components of the camera matrix $\mathbf{P}$:

- $\mathbf{K}$: 3x3 intrinsics
- $\mathbf{R}$: 3x3 3D rotation
- $[\mathbf{I}]$: 3x3 identity
- $-\mathbf{C}$: 3x1 3D translation

Perspective



Forced perspective




The Ames room illusion



The Ames room illusion




The 2D view of the (rearranged) pinhole camera




Magnification depends on depth

What happens as we change the focal length?




Magnification depends on focal length



What if...


real-world
object



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

What if...


real-world
object



Continue increasing Z and
 f while maintaining same
magnification?

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

Orthographic vs pinhole camera



Orthographic vs pinhole camera

General pinhole camera:

We also call these cameras:

$$\mathbf{P} = \mathbf{KR}[\mathbf{I} \mid -\mathbf{C}]$$

Projective camera

General orthographic camera:

$$\mathbf{P} = \mathbf{KR}[\mathbf{I} \mid -\mathbf{C}]$$

Affine camera

Bottom row is always $[0 \ 0 \ 0 \ 1]$

What is the rationale
behind these names?

References

Basic reading:

- Szeliski textbook, Section 2.1.5.

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.
- Torralba and Freeman, “Accidental Pinhole and Pinspeck Cameras,” CVPR 2012.
the eponymous paper discussed in the slides.