


Lecture 11:

Camera Models

COMP 590/776: Computer Vision

Instructor: Soumyadip (Roni) Sengupta

TA: Mykhailo (Misha) Shvets



Course Website:
Scan Me!

Breaking out of 2D

...now we are ready to break out of 2D



And enter the real world!



NeRF in the wild (will get to it towards the end!)




Nerfstudio: Colab friendly github repo for NeRF

(some of you might wanna try this for your project!)





Captured Images: Right




“Shape & Material Capture at Home”, Lichy, Wu, Sengupta, Jacobs, CVPR 2021

Single iPhone Image with Built-In Flash

Image 1/1



Mesh



Lectures in 3D Vision

- Fundamental Concepts (4 lectures)
 - Modeling camera and 3D->2D projection (2 lectures)
 - 2-view geometry & Stereo Vision (2 lectures)
- 3D Reconstruction techniques (2 lectures)
 - Multiview Stereo (MVS)
 - Structure from Motion (SfM) + SLAM
 - Photometric Stereo (PS)
- Deep Learning + 3D Vision (2 lectures)
 - Deep Learning + MVS, SfM, PS
 - Neural Radiance Fields (NeRFs)


Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- Perspective Distortion
- Other Projection models
- How to calibrate camera, i.e. estimate camera parameters (next class)

Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- Perspective Distortion
- Other Projection models
- How to calibrate camera, i.e. estimate camera parameters (next class)

Pinhole Camera



Focal length

- Can think of as “zoom”



24mm



50mm



200mm




800mm

- Also related to *field of view* (*inversely*)



Focal length in practice



24mm




50mm




135mm



Focal length = cropping



24mm




50mm



135mm




The lens camera




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

The pinhole camera




Central rays propagate in the same way for both models!

Important Difference: focal length




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

Important Difference: focal length



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

Describing both lens and pinhole cameras




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

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

From now on, we will describe cameras as pinhole cameras!
Focal length will mean aperture-sensor distance.

Camera parameters


- How can we model the geometry of a camera?



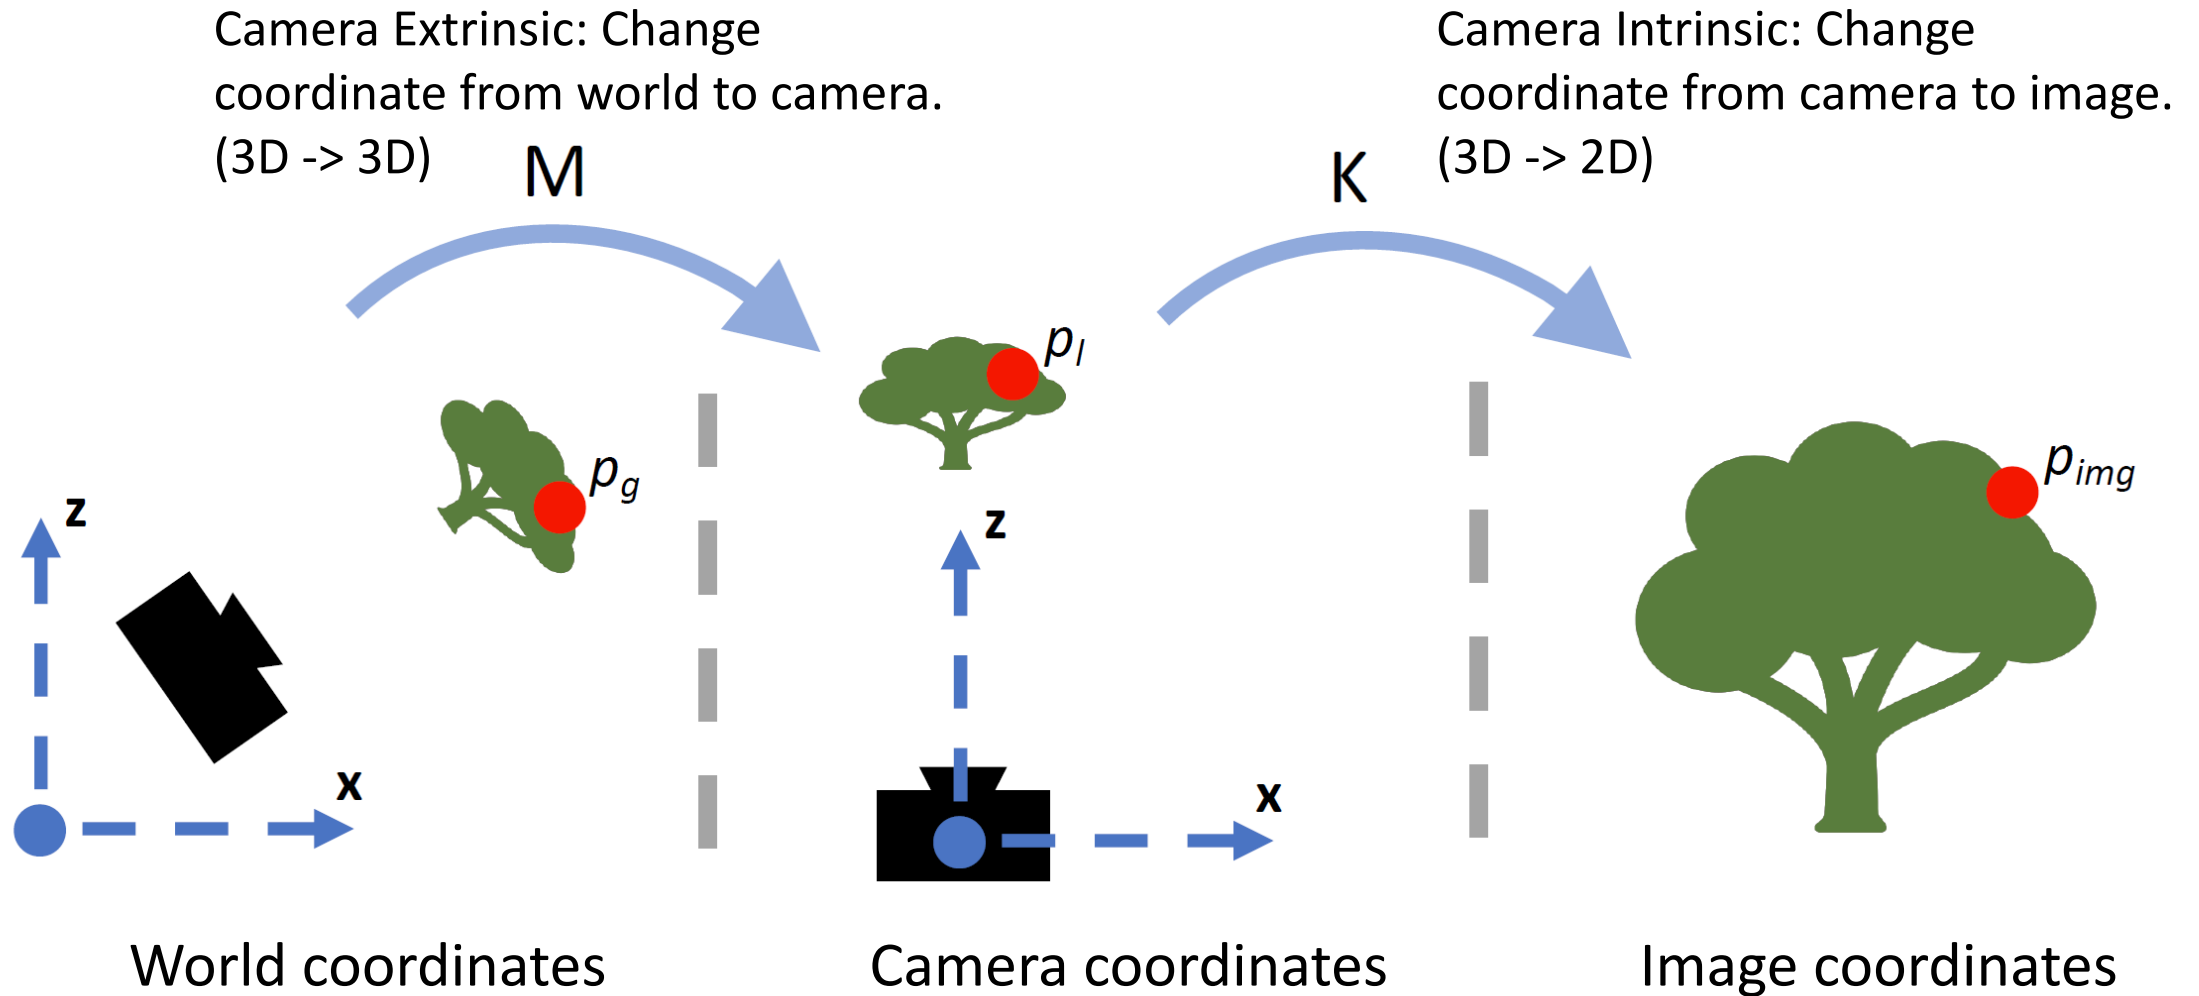
Three important coordinate systems:

1. *World* coordinates
2. *Camera* coordinates
3. *Image* coordinates

How do we project a given world point (x, y, z) to an image point?



Coordinate frames



Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- Perspective Distortion
- Other Projection models
- How to calibrate camera, i.e. estimate camera parameters (next class)

Camera parameters

To project a point (x, y, z) in *world* coordinates into a camera

- First transform (x, y, z) into *camera* coordinates
- Need to know
 - Camera position (in world coordinates)
 - Camera orientation (in world coordinates)
 - Together they form *Camera Extrinsics*
- Then project into the image plane to get *image (pixel) coordinates*
 - Need to know *Camera Intrinsics*

A camera is a mapping between the 3D world and a 2D image

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

2D image point camera matrix 3D world point

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

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

homogeneous
image
3 x 1

Camera
matrix
3 x 4

homogeneous
world point
4 x 1



Extrinsics

- How do we get the camera to “canonical form”?
 - Canonical form: Center of projection at the origin, x-axis points right, y-axis points down, z-axis points forwards

\mathbf{X}_w -> location of a point in world coordinate.


\mathbf{X}_c -> location of a point in camera coordinate.

Step 1: Translate by $-\mathbf{c}$



Extrinsics


- How do we get the camera to “canonical form”?
 - Canonical form: Center of projection at the origin, x-axis points right, y-axis points down, z-axis points forwards



Step 1: Translate by $-c$

Extrinsics


- How do we get the camera to “canonical form”?
 - Canonical form: Center of projection at the origin, x-axis points right, y-axis points down, z-axis points forwards



Step 1: Translate by $-c$
Step 2: Rotate by R

Extrinsics



- How do we get the camera to “canonical form”?
 - Canonical form: Center of projection at the origin, x-axis points right, y-axis points down, z-axis points forwards



Step 1: Translate by $-c$


Step 2: Rotate by \mathbf{R}

How do we represent 3D rotation?



Euler Angles


What did we do with 2D rotation?



$$\mathbf{R} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \underbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}}_{\mathbf{R}} \begin{bmatrix} x \\ y \end{bmatrix}$$

Rotation matrix in 3D



$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}$$

$$R_y(\theta) = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}$$

$$R_z(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Order of applying rotation matters
(composition of 3D rotations is not commutative.)

$$R = R_z(\alpha) R_y(\beta) R_x(\gamma) = \begin{bmatrix} \cos \alpha & \overset{\text{yaw}}{-\sin \alpha} & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \overset{\text{pitch}}{\cos \beta} & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix} \begin{bmatrix} \overset{\text{roll}}{1} & 0 & 0 \\ 0 & \cos \gamma & -\sin \gamma \\ 0 & \sin \gamma & \cos \gamma \end{bmatrix}$$

How to derive camera extrinsics? [*]

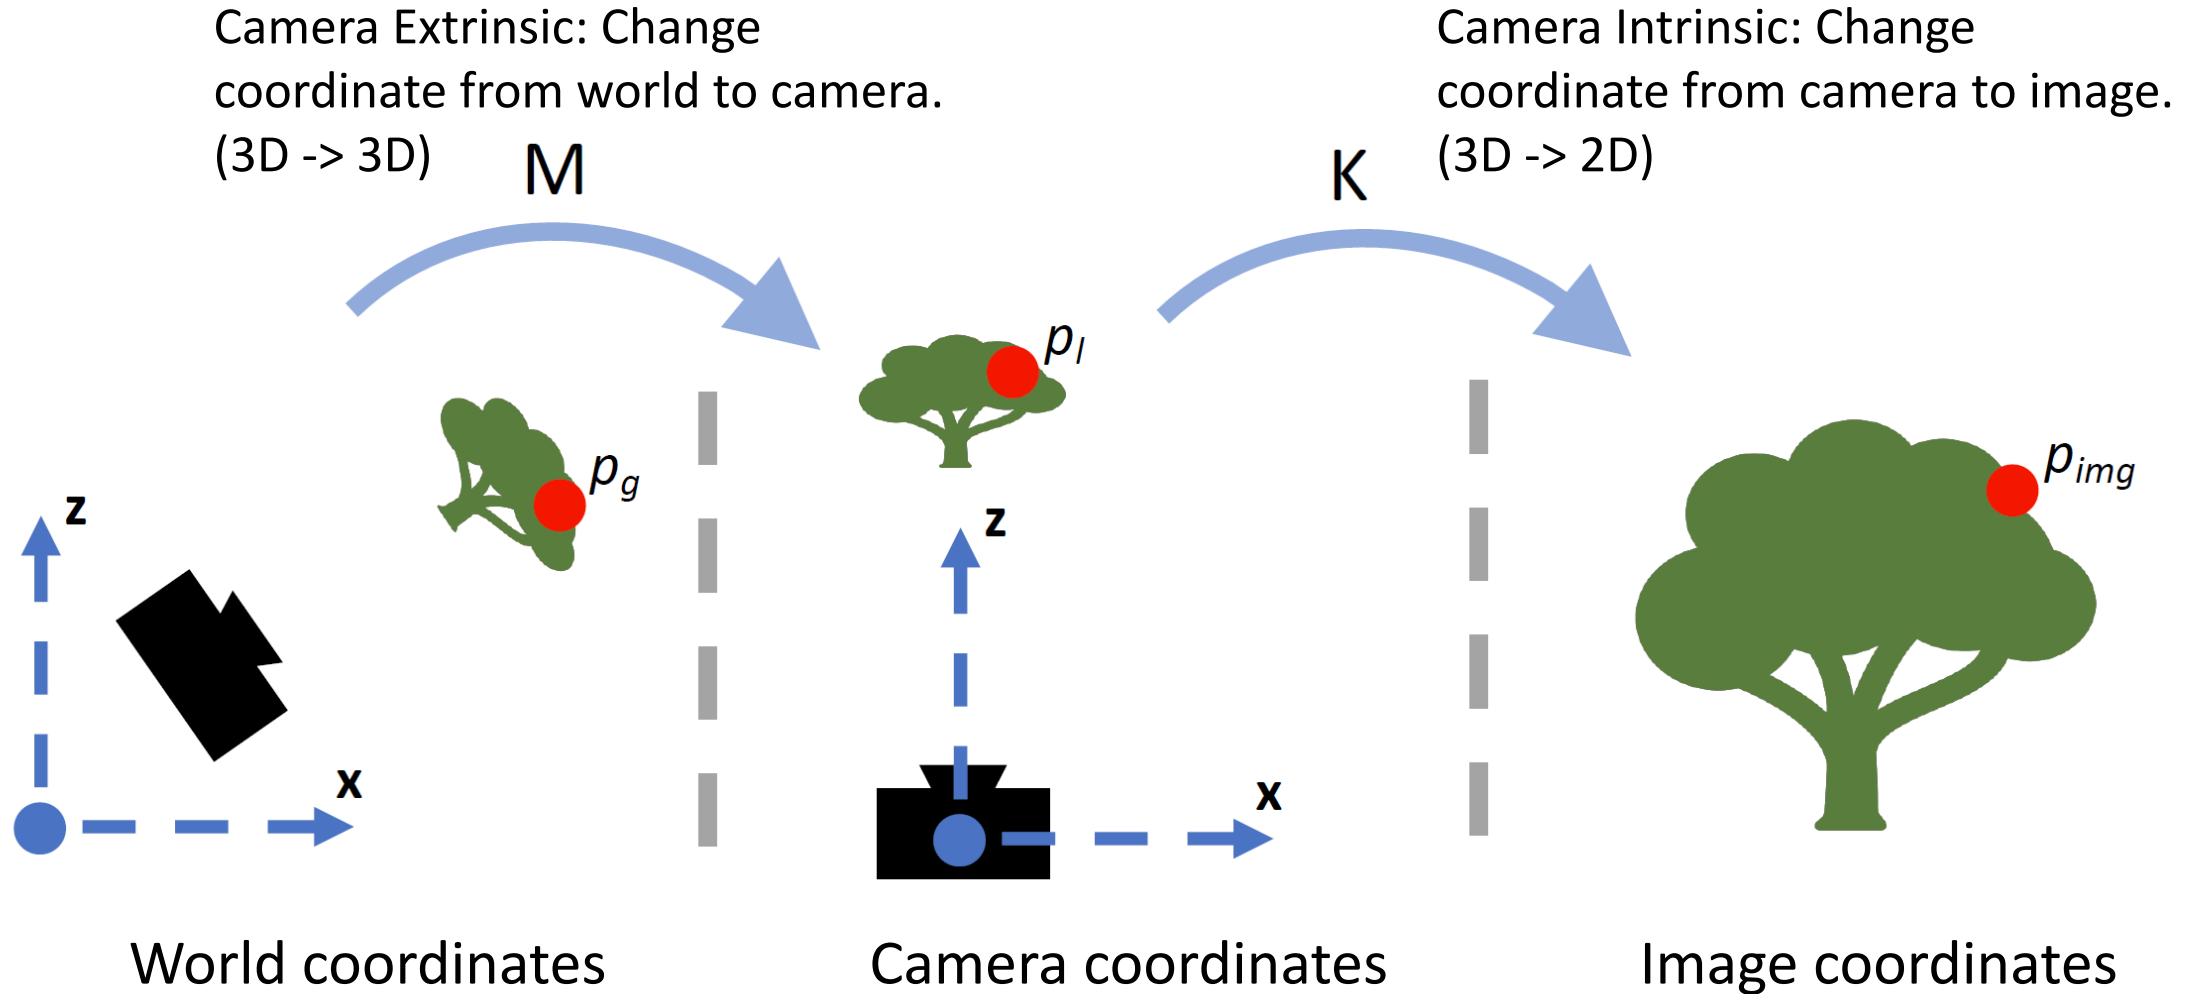
Show on board.

$M = R [I \mid -C]$ (translate first then rotate)


or

$M = [R \mid t]$, where $t = -RC$ (rotate first then translate)


Coordinate frames




The pinhole camera




The (rearranged) pinhole camera



Geometric Model: A Pinhole Camera




The (rearranged) pinhole camera



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

The 2D view of the (rearranged) pinhole camera



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

Derive this on board

The pinhole camera matrix for arbitrary focal length

Relationship from similar triangles:

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


General camera model *in homogeneous coordinates*:

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

What does the pinhole camera projection look like?


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

Geometric Model: A Pinhole Camera



Generalizing the camera matrix

In particular, the camera origin and image origin may 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

In particular, the camera origin and image origin may be different:



How does the camera matrix change?

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

shift vector
transforming
camera origin to
image origin

Show on board, why?

Typical Intrinsics matrix

$$\mathbf{K} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

2D affine transform corresponding to a scale by f (focal length) and a translation by (c_x, c_y) (principal point)

Maps 3D rays to 2D pixels

General case


$$\mathbf{K} = \begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

α : **aspect ratio** (1 unless pixels are not square)

s : **skew** (0 unless pixels are shaped like rhombi/parallelograms)

(c_x, c_y) : **principal point** ((w/2,h/2) unless optical axis doesn't intersect projection plane at image center)

Coordinate frames




Camera parameters

A camera is described by several parameters

- Translation \mathbf{T} of the optical center from the origin of world coords
- Rotation \mathbf{R} of the image plane
- focal length f , principal point (c_x, c_y) , pixel aspect size α
- blue parameters are called “extrinsics,” red are “intrinsics”

Projection equation

$$\mathbf{x} = \begin{bmatrix} sx \\ sy \\ s \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{\Pi} \mathbf{X}$$



- The projection matrix models the cumulative effect of all parameters
- Useful to decompose into a series of operations

$$\mathbf{\Pi} = \underbrace{\begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\text{intrinsics}} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{T}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{\text{translation}}$$

identity matrix

- The definitions of these parameters are **not** completely standardized
 - especially intrinsics—varies from one book to another

General pinhole camera matrix

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


$$\mathbf{P} = \underbrace{\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}}_{\substack{\text{intrinsic} \\ \text{parameters}}} \underbrace{\begin{bmatrix} r_1 & r_2 & r_3 & | & t_1 \\ r_4 & r_5 & r_6 & | & t_2 \\ r_7 & r_8 & r_9 & | & t_3 \end{bmatrix}}_{\substack{\text{extrinsic} \\ \text{parameters}}}$$

$$\mathbf{R} = \begin{bmatrix} r_1 & r_2 & r_3 \\ r_4 & r_5 & r_6 \\ r_7 & r_8 & r_9 \end{bmatrix} \quad \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix}$$

3D rotation 3D translation

Recap

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

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



The diagram shows four arrows pointing from labels below to terms in the equation above. 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 '3x1 3D translation' to $-\mathbf{C}$.

3x3	3x3	3x3	3x1
intrinsics	3D rotation	identity	3D translation

Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- **Perspective Distortion**
- Other Projection models
- How to calibrate camera, i.e. estimate camera parameters (next class)

The 2D view of the (rearranged) pinhole camera



Perspective distortion



long focal length



mid focal length



short focal length

Perspective distortion





<http://petapixel.com/2013/01/11/how-focal-length-affects-your-subjects-apparent-weight-as-seen-with-a-cat/>



Forced perspective



The Ames room illusion




The Ames room illusion



Magnification depends on depth

real-world
object



Forced Perspective in movies (VFX)



Forced Perspective in displays



Copyright © 2020 Daktronics Inc.


All rights reserved. Manuscript and artwork in this proposal is protected by copyright and may not be reproduced in any form without permission of Daktronics Inc.

No part of this proposal may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording or any other

information storage and retrieval system, without prior permission in writing from Daktronics Inc. 072709 Revision 1

Projective geometry—what's it good for?

- Uses of projective geometry
 - Drawing
 - Measurements
 - Mathematics for projection
 - Undistorting images
 - Camera pose estimation



[Paolo Uccello](#)

Dolly Zoom aka Vertigo Effect





Fredo Durand


Other camera models

Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- Perspective Distortion
- **Other Projection models**
- How to calibrate camera, i.e. estimate camera parameters (next class)

What if...


real-world
object



... we continue increasing Z
and f while maintaining
same magnification?

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

camera is *close*
to object and has
small focal length



perspective

weak perspective


camera is *far* from
object and has
large focal length

increasing focal length →


increasing distance from camera →



Different cameras




perspective camera



weak perspective camera

Weak perspective vs perspective camera



$$\begin{bmatrix} X & Y & Z \end{bmatrix}^T \mapsto \begin{bmatrix} fX/Z_0 & fY/Z_0 \end{bmatrix}^T$$

- magnification does not change with depth
- *constant* magnification depending on f and Z_0

magnification
changes with depth

$$\begin{bmatrix} X & Y & Z \end{bmatrix}^T \mapsto \begin{bmatrix} fX/Z & fY/Z \end{bmatrix}^T$$

Comparing camera matrices

Let's assume that the world and camera coordinate systems are the same.

- The *perspective* camera matrix can be written as:

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

- The *weak perspective* camera matrix can be written as:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1/d \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1/d \end{bmatrix} \Rightarrow (dx, dy) \quad \mathbf{P} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & Z_o \end{bmatrix}$$

When can we assume a weak perspective camera?

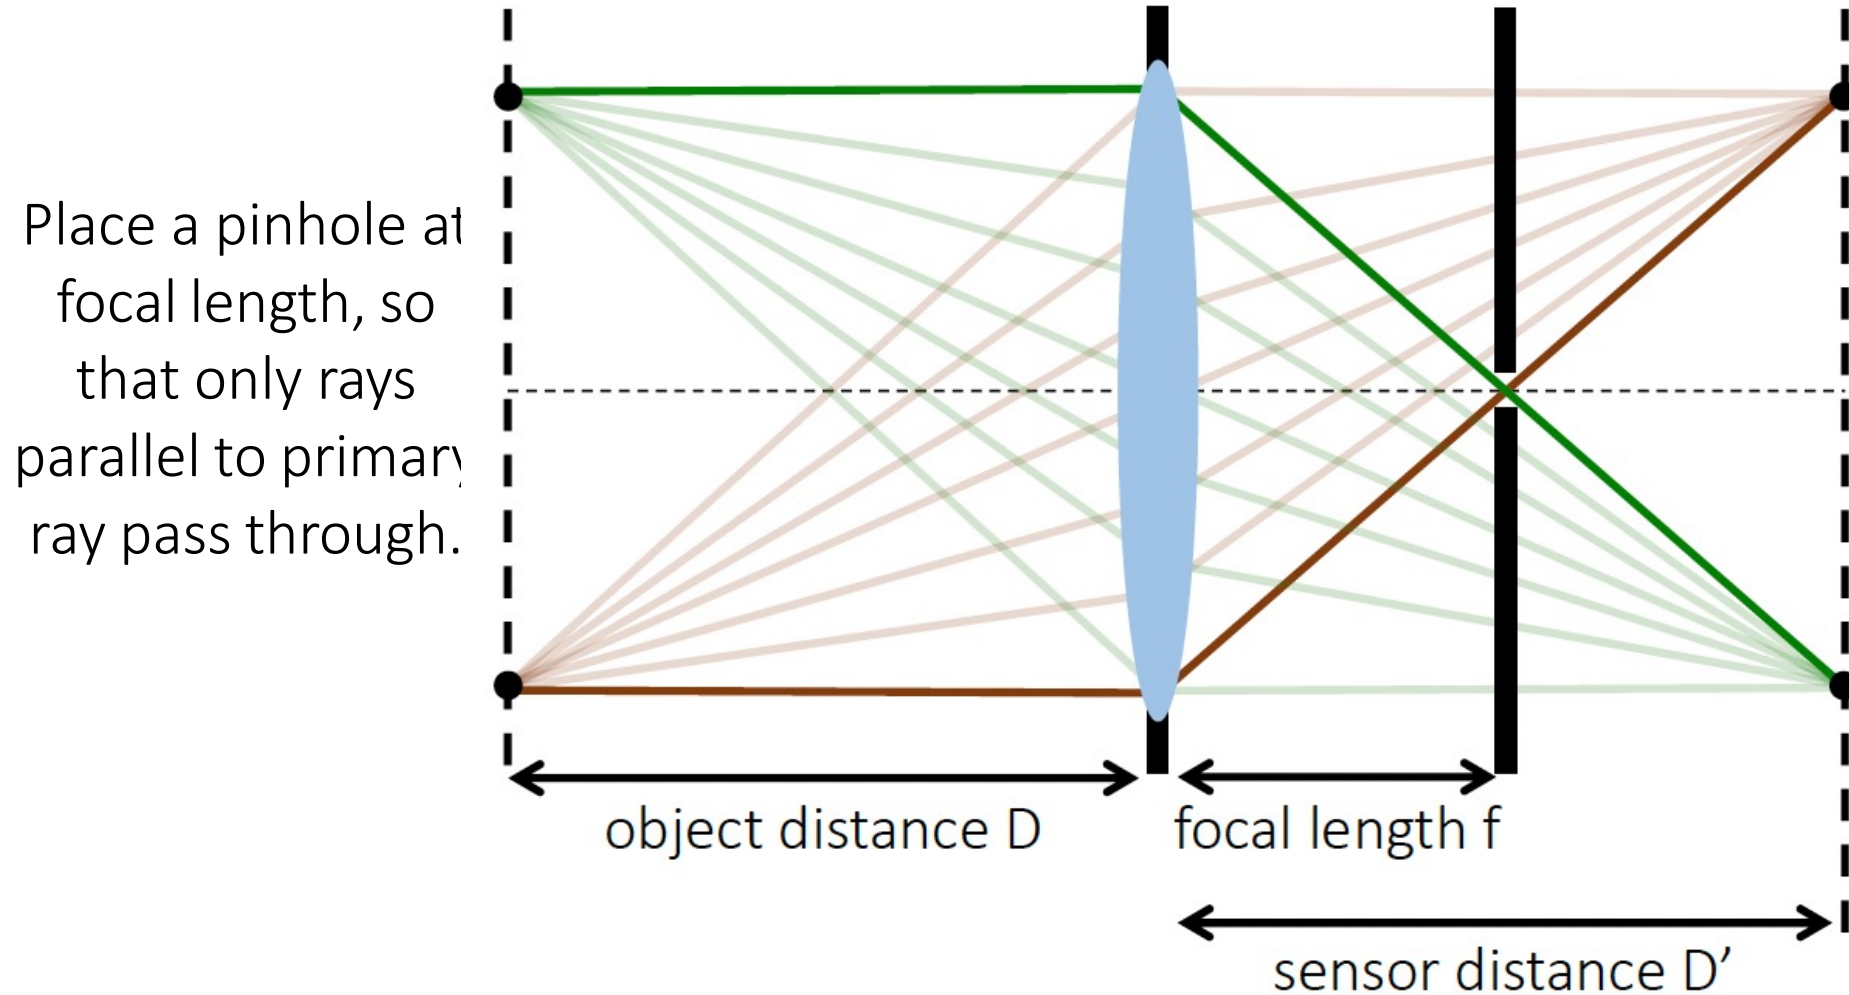
When the scene (or parts of it) is very far away.



Weak perspective projection applies to the mountains.


When can we assume a weak perspective camera?

When we use a telecentric lens.



Orthographic camera


- Special case of weak perspective camera where:
- constant magnification is equal to 1.
- there is no shift between camera and image origins.
- the world and camera coordinate systems are the same.



What is the camera matrix in this case?

Orthographic camera


- Special case of weak perspective camera where:
- constant magnification is equal to 1.
- there is no shift between camera and image origins.
- the world and camera coordinate systems are the same.



$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \Rightarrow (x, y)$$

Orthographic projection using a telecentric lens

How do we make the telecentric lens act as an orthographic camera?





We set the sensor distance as:


$$D' = 2f$$

in order to achieve unit magnification.


Orthographic projection



Perspective projection




Distortion



- Radial distortion of the image
 - Caused by imperfect lenses
 - Deviations are most noticeable for rays that pass through the edge of the lens



Radial distortion



- Arrows show motion of projected points relative to an ideal (distortion-free lens)

[Image credit: J. Bouguet http://www.vision.caltech.edu/bouguetj/calib_doc/htmls/example.html]

Correcting radial distortion



from [Helmut Dersch](#)

Today's Class

- Pinhole & Lens Camera
- Camera Parameters
 - Extrinsic
 - Intrinsic
- Perspective Distortion
- Other Projection models
- How to calibrate camera, i.e. estimate camera parameters (next class)

Slide Credits

- [CS5670, Introduction to Computer Vision](#), **Cornell Tech**, by Noah Snavely.
- [CS 194-26/294-26: Intro to Computer Vision and Computational Photography](#), **UC Berkeley**, by Angjoo Kanazawa.
- [CS 16-385: Computer Vision](#), **CMU**, by Matthew O'Toole

Additional Reading

- Multiview Geometry, Hartley & Zisserman, Chapter 6.1, 6.2.

Related Readings from the past

- Multiview Geometry, Hartley & Zisserman, Chapter 2.1, 2.2, 2.3, 2.4, 2.5, 2.6, 2.7, 3.1, 3.2, 4.1, 4.2