Lecture 17: Multi-view Geometry

Announcements

New IA office hours, starting next week

Name

Andrew Owens	Fri.
Haozhu Wang	Thu
Bingqi Sun	Mor
Anthony Liang	Tue

Instead of raising your hand: send a message, then (after I answer) unmute and ask your question Please send us feedback!

Office hour times

3:00pm - 4:00pm

J. 7:30pm - 8:30pm

n. 4:30 - 5:30

. 1:30 - 2:30pm

Today

- Review image formation • Epipolar geometry
- Image alignment

Recall: homogeneous coordinates

Representing translations:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

nomogeneous image coordinates

Converting from homogeneous coordinates:

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$



Source: N. Shavely





Recall: camera parameters



Three important coordinate systems:

- *World* coordinates
- *Camera* coordinates 2.
- *Image* coordinates З.

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



"The World"

Source: N. Snavely



Recall: camera parameters

 $\mathbf{\Pi} = \begin{bmatrix} f & s & c_x \\ 0 & \alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{R}_{3\times3} & \mathbf{0}_{3\times1} \\ \mathbf{0}_{1\times3} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3\times3} & \mathbf{T}_{3\times1} \\ \mathbf{0}_{1\times3} & 0 \end{bmatrix}$ intrinsics

rotation

translation

Source: N. Snavely



Estimating depth from multiple views



Stereo vision









1, 2, N eyes





1, 2, N eyes







1, 2, N eyes









Depth without objects Random dot stereograms (Bela Julesz)



FIGURE 8.13

1	0	1	0	0	1	0	1	0	1
0	٥	1	0	1	0	1	0	0	1
0	1	0	1	1	0	1	1	0	0
1	0	8	8	A	A	Y	0	1	٥
1	0	А.	₿	A	8	x	1	1	1
0	1	Ą	8	A	A	x	1	0	٥
1	0	8	A	8	8	Y	1	1	1
1	0	1	1	0	1	1	0	0	1
1	1	1	0	1	1	0	0	1	1
0	1	1	1	1	0	0	0	1	0

		_								
ſ	1	0	1	0	1	0	0	1	0	1
ľ	1	0	0	1	٥	1	0	1	0	٥
ſ	0	0	1	1	0	1	1	0	1	0
ľ	0	1	٥	A	A	8	8	X	0	1
ľ	1	1	1	9	А	8	A	Y	0	1
ľ	0	0	1	А	A	8	A	Y	1	0
ľ	1	1	1	в	в	A	8	x	0	1
Ī	1	0	0	1	1	0	1	1	Ö	1
Ì	1	1	0	0	1	1	0	1	1	1
İ	0	1	0	0	0	1	1	1	1	0







Brewster-type stereoscope, 1870

Visore stereoscopico portatile di tipo Brewster, J. Fleury - Hermagis, 1870, con messa a fuoco manuale. Per la visione di lastre e stampe stereoscopiche 8,5x17cm. Museo nazionale della scienza e della tecnologia Leonardo da Vinci, Milano.



View of Boston, c. 1860; an early stereoscopic card for viewing a scene from nature

Soule, John P., 1827-1904 -- Photographer - This image is available from the New York Public Library's Digital Library under the digital ID G90F336_113F: digitalgallery.nypl.org \rightarrow digitalcollections.nypl.org

🕲 Public Domain File: Charles Street Mall, Boston Common, by Soule, John P., 1827-1904 3.jpg Created: Coverage: 1860?-1890?. Source Imprint: 1860?-1890?. Digital item published 7-28-2005; updated 4-23-2009.

E Alessandro Nassiri - Museo della Scienza e della Tecnologia "Leonardo da Vinci"

More details

© CC BY-SA 4.0

File: IGB 006055 Visore stereoscopico portatile Museo scienza e tecnologia Milano.jpg Created: 1 July 2014

b More details







Z?



















Similar triangles: $\frac{T+X_L-X_R}{Z-f} = \frac{T-T+X_R}{Z-f}$





Similar triangles: $\frac{T+X_L-X_R}{Z-f} = \frac{T}{Z}$









21 Source: Torralba, Isola, Freeman



In 3D







Left image

Second picture is ~1m to the right

Right image





Left image

Right image





Left image

Right image





D(x,y)

$$Z(x,y) = \frac{a}{D(x,y)}$$



Finding correspondences



We only need to search for matches along horizontal lines.



Basic stereo algorithm



For each "epipolar line"

For each pixel in the left image

- compare with every pixel on same epipolar line in right image \bullet
- pick pixel with minimum match cost \bullet



Computing disparity





But you can learn depth from a single image



MegaDepth: Learning Single-View Depth Prediction from Internet Photos

Zhengqi LiNoah SnavelyDepartment of Computer Science & Cornell Tech, Cornell University





General case



• The two cameras need not have parallel optical axes.











Do we need to search for matches only along horizontal lines?





Do we need to search for matches only along horizontal lines?





Do we need to search for matches only along horizontal lines?

It looks like we need to search everywhere... are there any constraints that can guide the search?

Stereo correspondence constraints

If we see a point in camera 1, are there any constraints on where we will find it on camera 2?

Stereo correspondence constraints





Epipolar constraint











Baseline: the line connecting the two camera centers **Epipole**: point of intersection of *baseline* with the image plane





Baseline: the line connecting the two camera centers **Epipole**: point of intersection of *baseline* with the image plane





Baseline: the line connecting the two camera centers **Epipole**: point of intersection of *baseline* with the image plane

- Epipolar plane: the plane that contains the two camera centers and a 3D point in the world





Baseline: the line connecting the two camera centers **Epipole**: point of intersection of *baseline* with the image plane **Epipolar line**: intersection of the *epipolar plane* with each image plane

- **Epipolar plane:** the plane that contains the two camera centers and a 3D point in the world



Epipolar constraint



We can search for matches across epipolar lines

All epipolar lines intersect at the epipoles



The fundamental matrix



If we observe a point in one image, its position in the other image is constrained to lie on line defined by above.

- F: fundamental matrix

 $p^{T} F p' = 0$

p, p': image points in homogeneous coordinates



The fundamental matrix



$(p^{T} F) p' = 0$ $u^{T} p' = 0$

u: a line induced by p p, p': image points in homogeneous coordinates Closely related to projection matrix:

 $\mathbf{F} = \mathbf{K'}^{-\mathsf{T}}[\mathbf{t}]_{\mathsf{X}} \mathbf{R} \mathbf{K}^{-1}$

K, K': intrinsics matrices

R, t: relative pose

See <u>Hartley and Zisserman</u> for derivation



Example: converging cameras





Figure from Hartley & Zisserman



Source: Kristen Grauman



Image rectification





Source: A. Efros



Active stereo with structured light







Li Zhang, Brian Curless, and Steven M. Seitz. Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming. In Proceedings of the 1st International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT), Padova, Italy, June 19-21, 2002, pp. 24-36.

Li Zhang's one-shot stereo





49 Source: R. Szeliski











Future problem set!

Making panoramas





Image alignment

Why don't these image line up exactly?



$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \longleftrightarrow$ what happens when we change this row?

affine transformation

Recall: affine transformations





Projective Transformations aka Homographies aka Planar Perspective Maps

$\mathbf{H} = \left[\begin{array}{cccc} a & b & c \\ d & e & f \\ g & h & 1 \end{array} \right]$

Called a **homography** (or planar perspective map)











$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$

Note that this can be 0! A "point at infinity"







Points at infinity





Homography

Example: two pictures taken by rotating the camera:



If we try to build a panorama by overlapping them:





60



Homography

Example: two pictures taken by rotating the camera:







With a homography you can map both images into a single camera:

61



Why does this work?



How do we map points in image 2 into image 1?

	image 1	image 2
intrinsics	\mathbf{K}_1	\mathbf{K}_{2}
extrinsics (rotation only	$\mathbf{R}_1 = \mathbf{I}_{3 \times 3}$	\mathbf{R}_2

Step 1: Convert pixels in image 2 to rays in camera 2's coordinate system.



Step 2: Convert rays in camera 2's coordinates to rays in camera 1's coordinates.

 $\begin{bmatrix} X_1 \\ Y_1 \\ Z_1 \end{bmatrix} = \mathbf{R}_2^T \mathbf{K}_2^{-1} \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix}$

Step 3: Convert rays in camera 1's coordinates to pixels in image 1's coordinates.









Homographies

- Projective warps
- Properties of projective transformations: Origin does not necessarily map to origin

 - Lines map to lines
 - Parallel lines do not necessarily remain parallel
 - Ratios are not preserved
 - Closed under composition

• Homographies ... – Affine transformations, and $\begin{vmatrix} x' \\ y' \\ w' \end{vmatrix} = \begin{vmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{vmatrix} \begin{vmatrix} x \\ y \\ w \end{vmatrix}$



2D image transformations



Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$igg[egin{array}{c c} I & t \end{array} igg]_{2 imes 3} \end{array}$	2	orientation $+ \cdots$	
rigid (Euclidean)	$\left[egin{array}{c c} m{R} & t \end{array} ight]_{2 imes 3}$	3	lengths $+\cdots$	$\langle \rangle$
similarity	$\left[\left. s oldsymbol{R} \right oldsymbol{t} ight]_{2 imes 3}$	4	angles $+ \cdots$	\bigcirc
affine	$\left[egin{array}{c} oldsymbol{A} \end{array} ight]_{2 imes 3}$	6	parallelism $+\cdots$	
projective	$\left[egin{array}{c} ilde{oldsymbol{H}} ight]_{3 imes 3}$	8	straight lines	



Image warping

Given a coordinate transformation (x',y') = T(x,y) and a source image *f(x,y)*, how do we compute a transformed image g(x',y') = f(T(x,y))?





Forward warping

- location (x',y') = T(x,y) in g(x',y')



Send each pixel f(x) to its corresponding

What if a pixel lands "between" two pixels?



Forward warping

- Send each pixel f(x) to its corresponding location (x',y') = T(x,y) in g(x',y')
- What if a pixel lands "between" two pixels?
 - Answer: add "contribution" to several pixels, normalize later (splatting)
 - Can still result in holes







Inverse warping

- Get each pixel g(x',y') from its corresponding location $(\mathbf{x}, \mathbf{y}) = \mathbf{T}^{-1}(\mathbf{x}, \mathbf{y})$ in $f(\mathbf{x}, \mathbf{y})$
 - Requires taking the inverse of the transform
 - What if pixel comes from "between" two pixels?





Inverse warping

- Get each pixel g(x') from its corresponding location x' = h(x) in f(x)
 - What if pixel comes from "between" two pixels?
 - Answer: resample color value from interpolated (prefiltered) source image





Interpolation

- Possible interpolation filters:
 - nearest neighbor
 - bilinear
 - bicubic (interpolating)
 - sinc
- Needed to prevent "jaggies" and aliasing artifacts

(with prefiltering)









Next lecture: estimating geometry from images

