

Miniature faking



In close-up photo, the depth of field is limited.

http://en.wikipedia.org/wiki/File:Jodhpur_tilt_shift.jpg

Miniature faking



Miniature faking







http://en.wikipedia.org/wiki/File:Oregon_State_Beavers_Tilt-Shift_Miniature_Greg_Keene.jpg


Review

- Previous section:
 - Model fitting and outlier rejection

Review: Hough transform



Review: RANSAC




Algorithm:

1. **Sample** (randomly) the number of points required to fit the model ($\#=2$)
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model

Repeat 1-3 until the best model is found with high confidence

Review: 2D image transformations




Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	

This section – multiple views

- Today – Intro to multiple views and Stereo. Camera Calibration (if we have time).
- Next Lecture – Epipolar Geometry and Fundamental Matrix. Stereo Matching (if there is time).
- Both lectures are relevant for project 3.

Recap: Oriented and Translated Camera



Recap: Degrees of freedom

$$\mathbf{x} = \mathbf{K} \begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix} \mathbf{X}$$







$$w \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{matrix} 5 \\ \begin{bmatrix} \alpha & s & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix} \end{matrix} \begin{matrix} 6 \\ \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \end{matrix} \begin{matrix} t_x \\ t_y \\ t_z \end{matrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$



This Lecture: How to calibrate the camera?

$$\mathbf{x} = \mathbf{K} \begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix} \mathbf{X}$$

$$\begin{bmatrix} su \\ sv \\ s \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$


What can we do with camera calibration?






Stereo: Intro

Computer Vision
James Hays




Slides by
Kristen Grauman

Multiple views




a



Hartley and Zisserman

stereo vision
structure from motion
optical flow



Why multiple views?


- Structure and depth are inherently ambiguous from single views.





Why multiple views?

- Structure and depth are inherently ambiguous from single views.




- What cues help us to perceive 3d shape and depth?


Shading



a)



b)




c)

[Figure from Prados & Faugeras 2006]

Focus/defocus





Images from
same point of
view, different
camera
parameters



3d shape / depth
estimates

Texture



[From [A.M. Loh. The recovery of 3-D structure using visual texture patterns.](#) PhD thesis]



Perspective effects




NATIONALGEOGRAPHIC.COM

© 2003 National Geographic Society. All rights reserved.

Motion



Occlusion




Rene Magritte's famous painting *Le Blanc-Seing* (literal translation: "The Blank Signature") roughly translates as "free hand". 1965



If stereo were critical for depth perception, navigation, recognition, etc., then this would be a problem

Stereo photography and stereo viewers

Take two pictures of the same subject from two slightly different viewpoints and display so that each eye sees only one of the images.



Invented by Sir Charles Wheatstone, 1838




Image from fisher-price.com






© Copyright 2001 Johnson-Shaw Stereoscopic Museum



<http://www.johnsonshawmuseum.org>



© Copyright 2001 Johnson-Shaw Stereoscopic Museum



<http://www.johnsonshawmuseum.org>



Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923





http://www.well.com/~jimmg/stereo/stereo_list.html



http://www.well.com/~jimmg/stereo/stereo_list.html

Autostereograms



Exploit disparity as depth cue using single image.


(Single image random dot stereogram, Single image stereogram)

Autostereograms



Images from magiceye.com

Parallax and our universe




Look again at that dot. That's here. That's home. That's us. On it everyone you love, everyone you know, everyone you ever heard of, every human being who ever was, lived out their lives. The aggregate of our joy and suffering, thousands of confident religions, ideologies, and economic doctrines, every hunter and forager, every hero and coward, every creator and destroyer of civilization, every king and peasant, every young couple in love, every mother and father, hopeful child, inventor and explorer, every teacher of morals, every corrupt politician, every "superstar," every "supreme leader," every saint and sinner in the history of our species lived there--on a mote of dust suspended in a sunbeam.


— Carl Sagan




Nicolaus Copernicus



Motion of Sun (yellow), Earth (blue), and Mars (red). At left, Copernicus' heliocentric motion. At right, traditional geocentric motion, including the retrograde motion of Mars.



geocentric model (often exemplified specifically by the **Ptolemaic system**)





Tycho Brahe

If the apparent motion of the planets is caused by parallax, why aren't we seeing parallax for stars?

it was one of Tycho Brahe's principal objections to Copernican heliocentrism that for it to be compatible with the lack of observable stellar parallax, there would have to be an enormous and unlikely void between the orbit of Saturn and the eighth sphere (the fixed stars).

The angles involved in these calculations are very small and thus difficult to measure. The nearest star to the Sun (and also the star with the largest parallax), Proxima Centauri, has a parallax of 0.7685 ± 0.0002 arcsec.[1] This angle is approximately that subtended by an object 2 centimeters in diameter located 5.3 kilometers away. First reliable measurements of parallax were not made until 1838, by Friedrich Bessel



Stereo vision



Two cameras, simultaneous views




Single moving camera and static scene

Modern stereo depth estimation example




Left stereo image

all:10 = 1.17	fg:10 = 1.23	bg:10 = 1.14
all:5 = 3.79	fg:5 = 2.59	bg:5 = 4.39
all:3 = 10.45	fg:3 = 9.95	bg:3 = 10.71




Ground truth (GT)

all:10 = 1.20	fg:10 = 1.21	bg:10 = 1.20
all:5 = 4.37	fg:5 = 3.02	bg:5 = 5.06
all:3 = 11.21	fg:3 = 10.47	bg:3 = 11.58




GT overlaid on image

all:10 = 1.75	fg:10 = 1.41	bg:10 = 1.91
all:5 = 2.93	fg:5 = 2.58	bg:5 = 3.10
all:3 = 6.45	fg:3 = 3.03	bg:3 = 8.19




msmdnet


all:10 = 1.04	fg:10 = 0.98	bg:10 = 1.08
all:5 = 4.31	fg:5 = 3.91	bg:5 = 4.51
all:3 = 9.77	fg:3 = 8.06	bg:3 = 10.64




4Fun



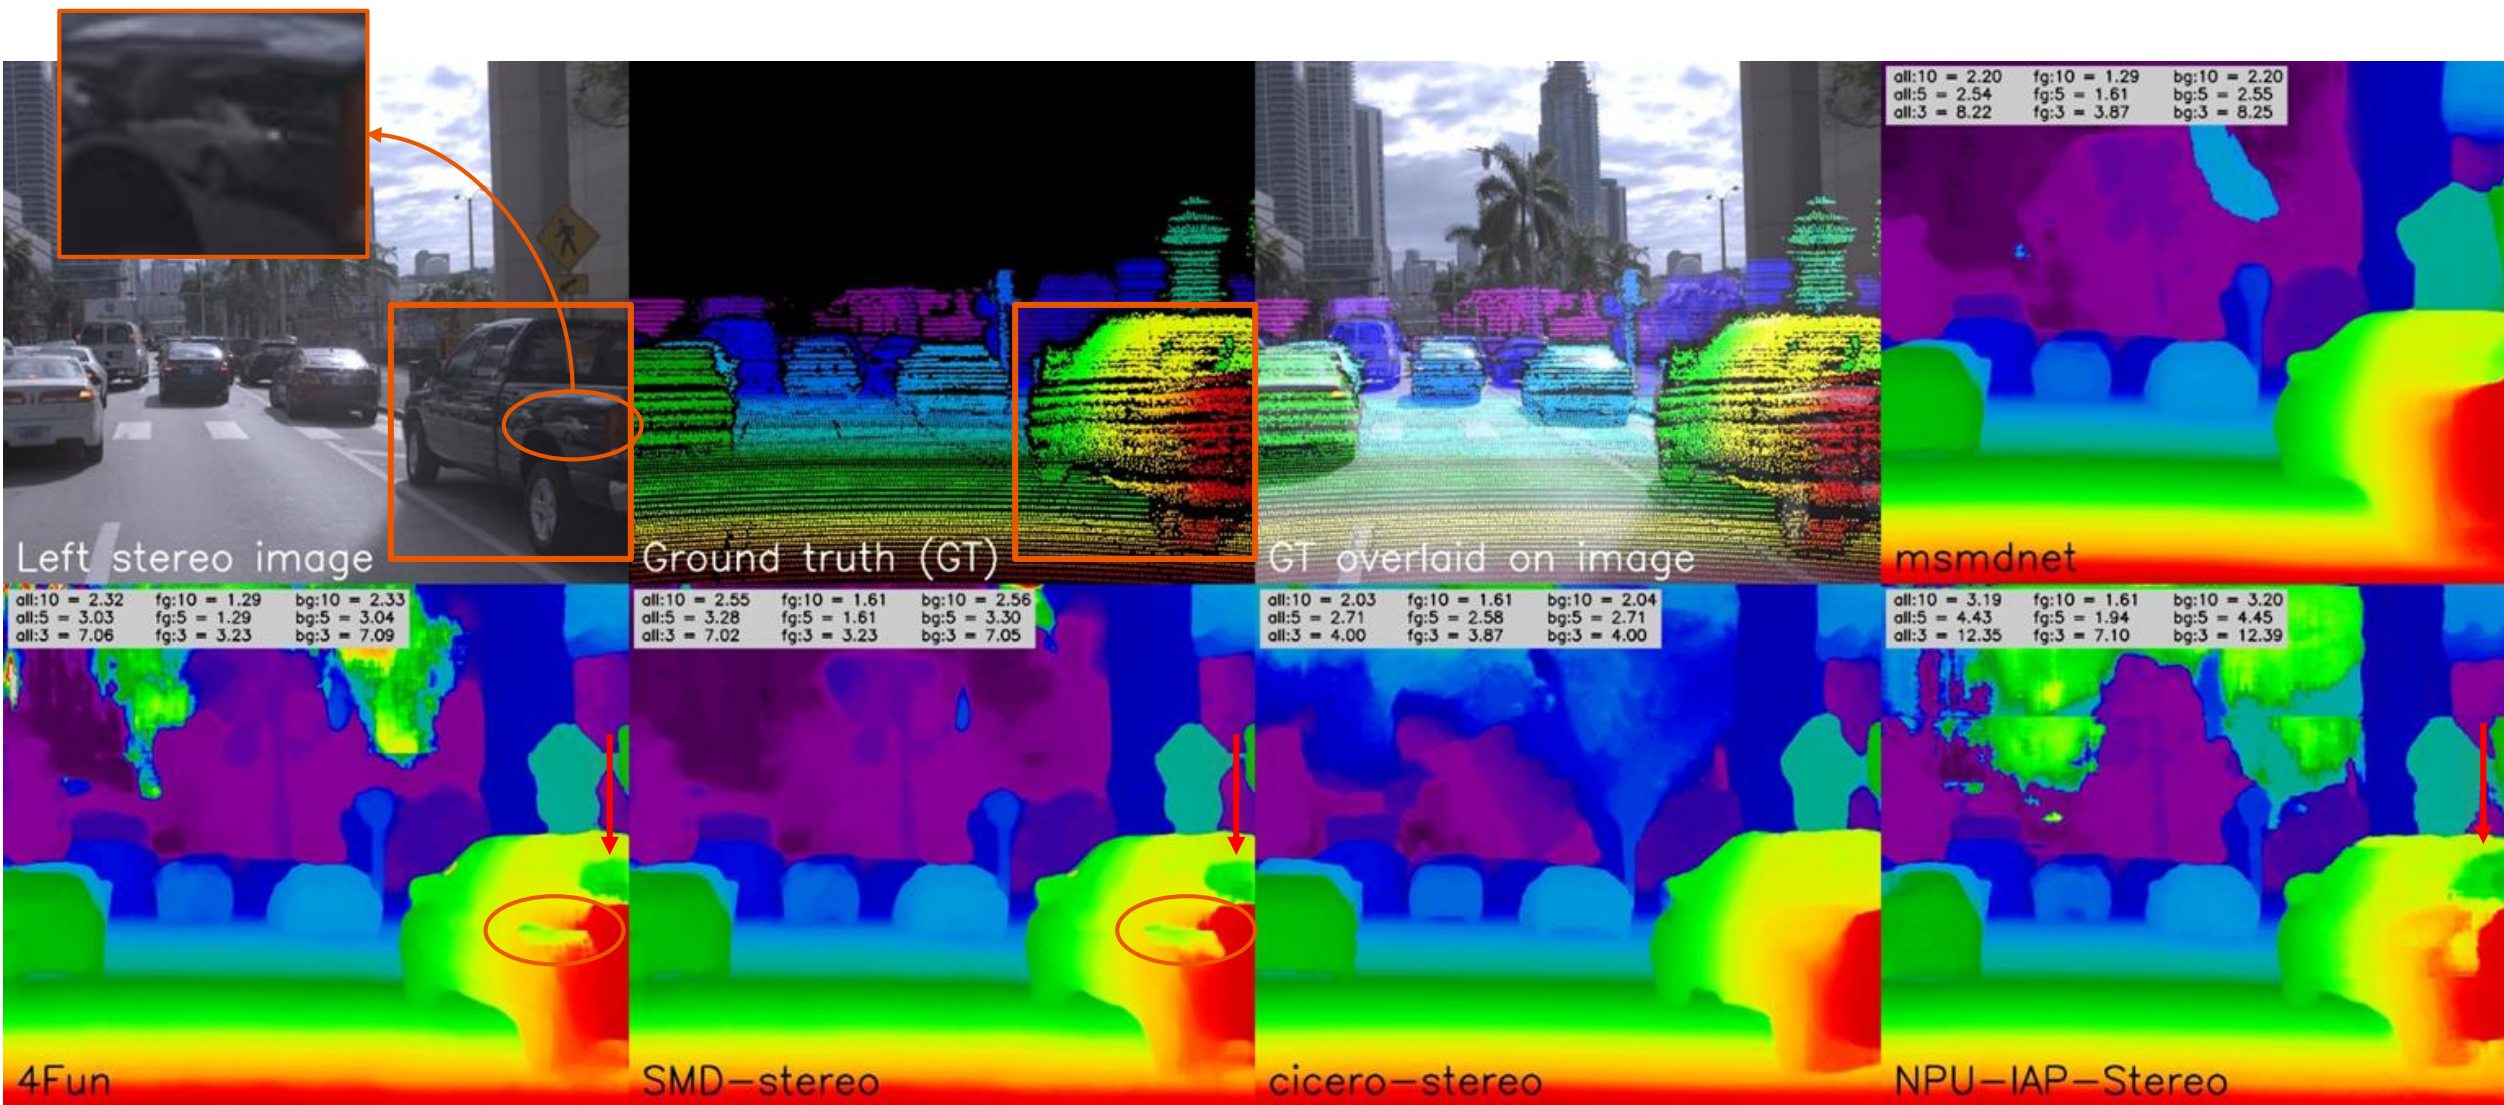
SMD-stereo










cicero-stereo



NPU-IAP-Stereo









Multi-view geometry problems

- **Stereo correspondence:** Given a point in one of the images, where could its corresponding points be in the other images?




Multi-view geometry problems

- **Structure:** Given projections of the same 3D point in two or more images, compute the 3D coordinates of that point





Multi-view geometry problems

- **Motion:** Given a set of corresponding points in two or more images, compute the camera parameters




Estimating depth with stereo

- **Stereo:** shape from “motion” between two views
- We’ll need to consider:
 - Info on camera pose (“calibration”)
 - Image point correspondences



Camera parameters



Extrinsic parameters:

Camera frame 1 \leftrightarrow Camera frame 2

Intrinsic parameters:


Image coordinates relative to camera \leftrightarrow Pixel coordinates

- *Extrinsic* params: rotation matrix and translation vector
- *Intrinsic* params: focal length, pixel sizes (mm), image center point, radial distortion parameters

We'll assume for now that these parameters are given and fixed.


Geometry for a simple stereo system

- First, assuming parallel optical axes, known camera parameters (i.e., calibrated cameras):



Geometry for a simple stereo system

- Assume parallel optical axes, known camera parameters (i.e., calibrated cameras). **What is expression for Z?**



Similar triangles (p_l, P, p_r) and (O_l, P, O_r) :

$$\frac{T - x_l + x_r}{Z - f} = \frac{T}{Z}$$

$$Z = f \frac{T}{x_l - x_r}$$

disparity

Depth from disparity

image $I(x,y)$



Disparity map $D(x,y)$





image $I'(x',y')$



$$(x',y')=(x+D(x,y), y)$$


So if we could find the **corresponding points** in two images, we could **estimate relative depth**...

Where do we need to search?



Key idea: Epipolar constraint

Key idea: Epipolar constraint




Potential matches for x have to lie on the corresponding line l' .


Potential matches for x' have to lie on the corresponding line l .

Wouldn't it be nice to know where matches can live? To constrain our 2d search to 1d.

VLFeat's 800 most confident matches
among 10,000+ local features.




Epipolar geometry: notation




- **Baseline** – line connecting the two camera centers
- **Epipoles**
= intersections of baseline with image planes
= projections of the other camera center
- **Epipolar Plane** – plane containing baseline (1D family)

Epipolar geometry: notation






- **Baseline** – line connecting the two camera centers
- **Epipoles**
= intersections of baseline with image planes
= projections of the other camera center
- **Epipolar Plane** – plane containing baseline (1D family)
- **Epipolar Lines** - intersections of epipolar plane with image planes (always come in corresponding pairs)

Example: Converging cameras






Example: Motion parallel to image plane



Example: Forward motion

What would the epipolar lines look like if the camera moves directly forward?


Example: Forward motion



Epipole has same coordinates in both images.

Points move along lines radiating from e :
“Focus of expansion”

Epipolar constraint: Calibrated case



$$\hat{x} = K^{-1} x = X$$

$$\hat{x}' = K'^{-1} x' = X'$$

$$\hat{x} \cdot [t \times (R\hat{x}')] = 0$$

(because \hat{x} , $R\hat{x}'$, and t are co-planar)

To be continued