Announcements

- Project 2 artifacts—vote now!!
- Project 3 questions?
- Start thinking about final project ideas, partners

Recovering 3D from images

So far, we've relied on a human to provide depth cues

- parallel lines, reference points, etc.

How might we do this automatically?

- What cues in the image provide 3D information?

Visual cues

Shading

Merle Norman Cosmetics, Los Angeles

Visual cues

Shading

Texture

***The Visual Cliff*, by William Vandivert, 1960**

Visual cues

Shading

Texture

Focus

From ***The Art of Photography***, **Canon**

Visual cues

Shading

Texture

Focus

Motion

Visual cues

Shading

Texture

Focus

Motion

Others:

- Highlights
- Shadows
- Silhouettes
- Inter-reflections
- Symmetry
- Light Polarization
- ...

Shape From X

- X = shading, texture, focus, motion, ...
- In this class we'll focus on the motion cue

Stereo

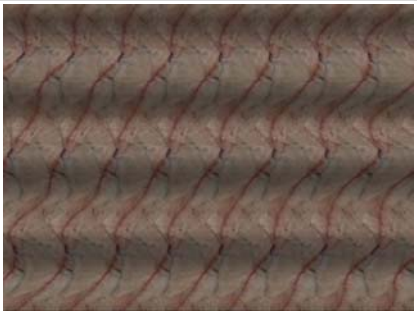

Single image stereogram, by Niklas Een

Readings

- Forsyth, Chapters 10.1, 11

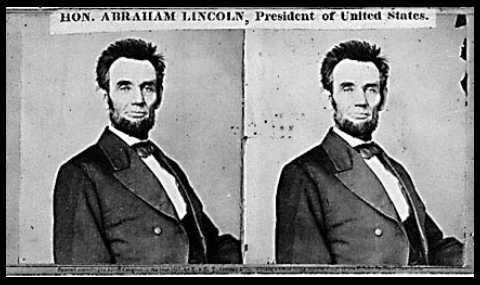

Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923

Teesta suspension bridge-Darjeeling, India

Mark Twain at Pool Table", no date, UCR Museum of Photography

Woman getting eye exam during immigration procedure at Ellis Island, c. 1905 - 1920 , UCR Museum of Phography

Stereograms online

UCR stereographs
- http://www.cmp.ucr.edu/site/exhibitions/stereo/

The Art of Stereo Photography
- http://www.photostuff.co.uk/stereo.htm

History of Stereo Photography
- http://www.rpi.edu/~ruiz/stereo_history/text/historystereog.html

Double Exposure
- http://home.centurytel.net/s3dcor/index.html

Stereo Photography
- http://www.shortcourses.com/book01/chapter09.htm

3D Photography links
- http://www.studyweb.com/links/5243.html

National Stereoscopic Association
- http://204.248.144.203/3dLibrary/welcome.html

Books on Stereo Photography
- http://userwww.sfsu.edu/~hl/3d.biblio.html

A free pair of red-blue stereo glasses can be ordered from Rainbow Symphony Inc
- http://www.rainbowsymphony.com/freestuff.html

Stereo

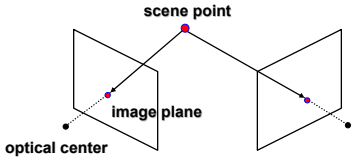

Stereo

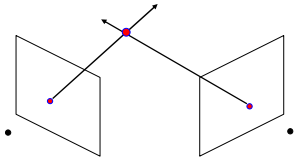

Basic Principle: Triangulation
- Gives reconstruction as intersection of two rays
- Requires
  - calibration
  - ***point correspondence***

Stereo correspondence

Determine Pixel Correspondence

- Pairs of points that correspond to same scene point

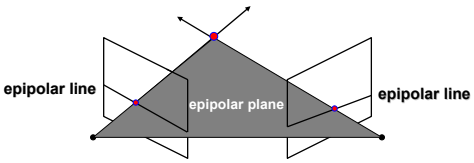

Epipolar Constraint

- Reduces correspondence problem to 1D search along *conjugate epipolar lines*
- Java demo: http://www.ai.sri.com/~luong/research/Meta3DViewer/EpipolarGeo.html

Fundamental matrix

Let *p* be a point in left image, *p'* in right image

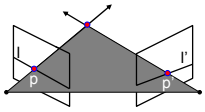

Epipolar relation

- *p* maps to epipolar line *l'*
- *p'* maps to epipolar line *l*

Epipolar mapping described by a 3x3 matrix *F*

$$l' = Fp$$

$$l = p'F$$

It follows that

$$p'Fp = 0$$

Fundamental matrix

This matrix F is called

- the "Essential Matrix"
  - when image intrinsic parameters are known
- the "Fundamental Matrix"
  - more generally (uncalibrated case)

Can solve for F from point correspondences

- Each (p, p') pair gives one linear equation in entries of F

$$p'Fp = 0$$

- 8 points give enough to solve for F (8-point algorithm)
- see readings (Forsyth chapter 10.1) for more on this

Stereo image rectification

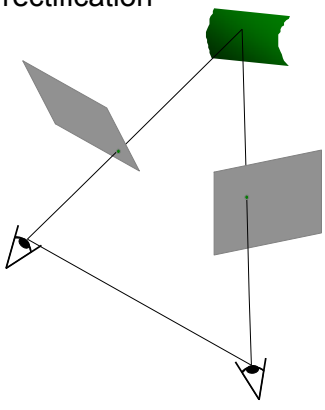

Stereo image rectification

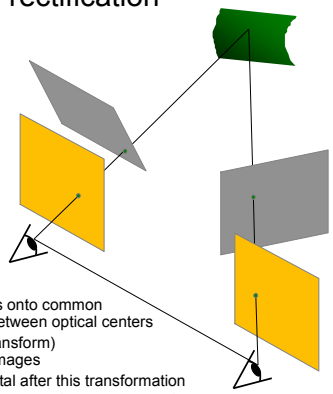

Image Reprojection

- reproject image planes onto common plane parallel to line between optical centers
- a homography (3x3 transform) applied to both input images
- pixel motion is horizontal after this transformation
- C. Loop and Z. Zhang. Computing Rectifying Homographies for Stereo Vision. IEEE Conf. Computer Vision and Pattern Recognition, 1999.

Stereo matching algorithms

Match Pixels in Conjugate Epipolar Lines

- Assume brightness constancy
- This is a tough problem
- Numerous approaches
  - dynamic programming [Baker 81,Ohta 85]
  - smoothness functionals
  - more images (trinocular, N-ocular) [Okutomi 93]
  - graph cuts [Boykov 00]
- A good survey and evaluation: http://www.middlebury.edu/stereo/

Your basic stereo algorithm

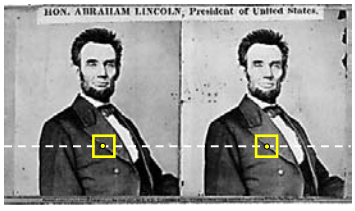

For each epipolar line

For each pixel in the left image

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

Improvement: match ***windows***

- This should look familar...
- Can use Lukas-Kanade or discrete search (latter more common)

Window size

W = 3 W = 20

Effect of window size

- Smaller window
  - \+
  - –
- Larger window
  - \+
  - –

Better results with *adaptive window*

- T. Kanade and M. Okutomi, *A Stereo Matching Algorithm with an Adaptive Window: Theory and Experiment*., Proc. International Conference on Robotics and Automation, 1991.
- D. Scharstein and R. Szeliski. Stereo matching with nonlinear diffusion. International Journal of Computer Vision, 28(2):155-174, July 1998

Stereo results

- Data from University of Tsukuba
- Similar results on other images without ground truth

Scene

Ground truth

Results with window search

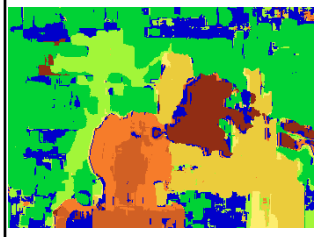

Window-based matching
(best window size)

Ground truth

Better methods exist...

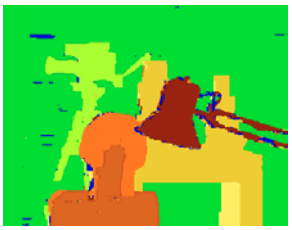

State of the art method

Ground truth

Boykov et al., Fast Approximate Energy Minimization via Graph Cuts, International Conference on Computer Vision, September 1999.

Stereo as energy minimization

Matching Cost Formulated as Energy

- "data" term penalizing bad matches

$$D(x,y,d) = \left|\mathbf{I}(x,y) - \mathbf{J}(x+d,y)\right|$$

- "neighborhood term" encouraging spatial smoothness

$$V(d_1,d_2) = \text{cost of adjacent pixels with labels d1 and d2}$$
$$= \left|d_1 - d_2\right| \quad \text{(or something similar)}$$

$$E = \sum_{(x,y)} D(x,y,d_{x,y}) + \sum_{neighbors\ (x1,y1),(x2,y2)} V(d_{x1,y1}, d_{x2,y2})$$

Stereo as a graph problem [Boykov, 1999]

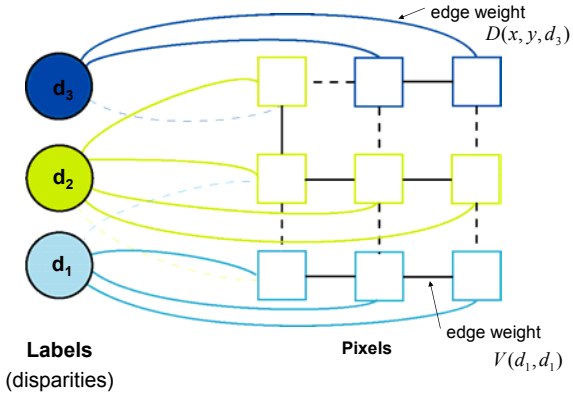

Graph definition

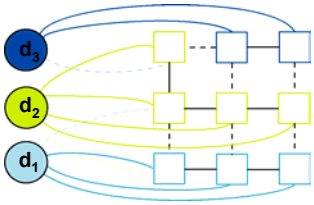

Initial state

- Each pixel connected to it's immediate neighbors
- Each disparity label connected to all of the pixels

Stereo matching by graph cuts

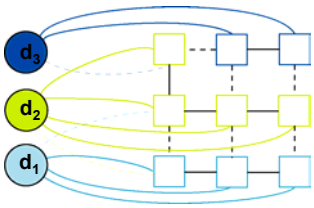

Graph Cut

- Delete enough edges so that
  - each pixel is (transitively) connected to exactly one label node
- Cost of a cut: sum of deleted edge weights
- Finding min cost cut equivalent to finding global minimum of the energy function

Computing a multiway cut

With two labels: classical min-cut problem

- Solvable by standard network flow algorithms
  - polynomial time in theory, nearly linear in practice

More than 2 labels: NP-hard [Dahlhaus *et al.*, STOC '92]

- But efficient approximation algorithms exist
  - Within a factor of 2 of optimal
  - Computes local minimum in a strong sense
    - » even very large moves will not improve the energy
  - Yuri Boykov, Olga Veksler and Ramin Zabih, Fast Approximate Energy Minimization via Graph Cuts, International Conference on Computer Vision, September 1999.
- Basic idea
  - reduce to a series of 2-way-cut sub-problems, using one of:
    - » swap move: pixels with label l1 can change to l2, and vice-versa
    - » expansion move: any pixel can change it's label to l1

Depth from disparity

input image (1 of 2)

depth map
[Szeliski & Kang '95]

3D rendering

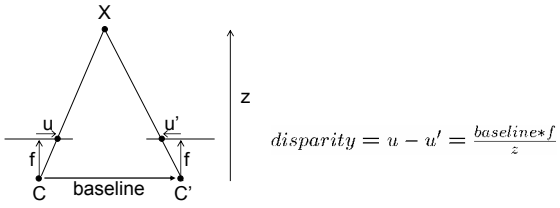

$$disparity = u - u' = \frac{baseline * f}{z}$$

Image-based rendering

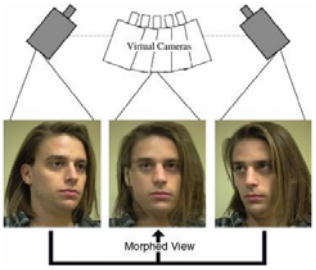

Render new views from raw disparity

- S. M. Seitz and C. R. Dyer, View Morphing, *Proc. SIGGRAPH 96*, 1996, pp. 21-30.
- L. McMillan and G. Bishop. Plenoptic Modeling: An Image-Based Rendering System, *Proc. of SIGGRAPH 95*, 1995, pp. 39-46.

Stereo reconstruction pipeline

Steps

- Calibrate cameras
- Rectify images
- Compute disparity
- Estimate depth

What will cause errors?

- Camera calibration errors
- Poor image resolution
- Occlusions
- Violations of brightness constancy (specular reflections)
- Large motions
- Low-contrast image regions

Stereo matching

Features vs. Pixels?

- Do we extract features prior to matching?

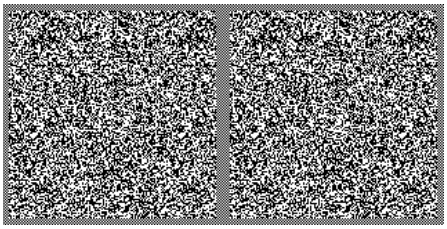

Julesz-style Random Dot Stereogram

Active stereo with structured light

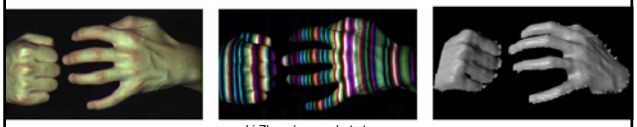

Li Zhang's one-shot stereo

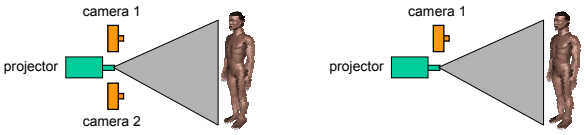

Project "structured" light patterns onto the object

- simplifies the correspondence problem

Active stereo with structured light

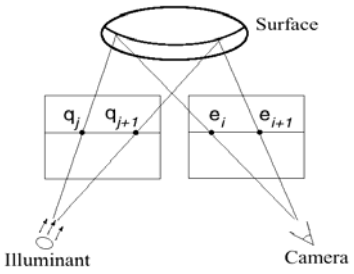

Laser scanning

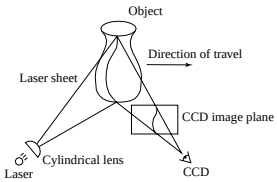

Digital Michelangelo Project
http://graphics.stanford.edu/projects/mich/

Optical triangulation

- Project a single stripe of laser light
- Scan it across the surface of the object
- This is a very precise version of structured light scanning

Portable 3D laser scanner (this one by Minolta)

Real-time stereo

Nomad robot searches for meteorites in Antartica
http://www.frc.ri.cmu.edu/projects/meteorobot/index.html

real-time stereo video

Used for robot navigation (and other tasks)

- Several software-based real-time stereo techniques have been developed (most based on simple discrete search)