Announcements

- Project 3 code & artifact due Tuesday
- Final project proposals due noon Wed (by email)
  - One-page writeup (from project web page), specifying:
    - » Your team members
    - » Project goals. Be specific. Describe the input and output.
    - » Brief description of your approach. If you are implementing or extending a previous method, give the reference and web link to the paper.
    - » Will you be using helper code (e.g., available online) or will you implement it all yourself?
    - » Evaluation method. How will you test it? Which test cases will you use?
    - » Breakdown--what will each team-member do? Ideally, everyone should do something imaging/vision related (it's not good for one team member to focus purely on user-interface, for instance).
    - » Special equipment that will be needed. We may be able to help with cameras, tripods, etc.

Stereo

Single image stereogram, by Niklas Een

Readings

- Szeliski, Chapter 10 (through 10.5)

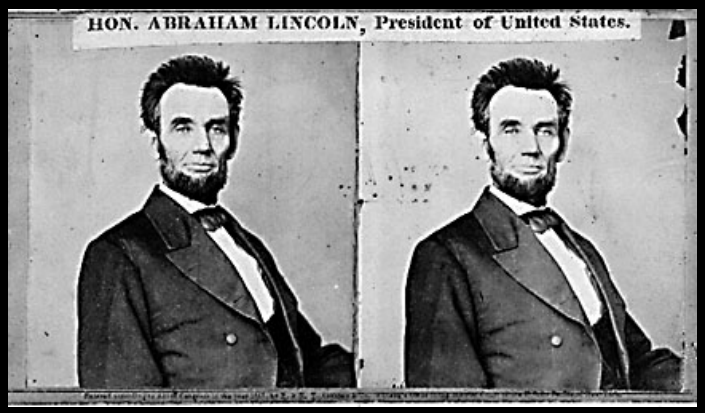

Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923

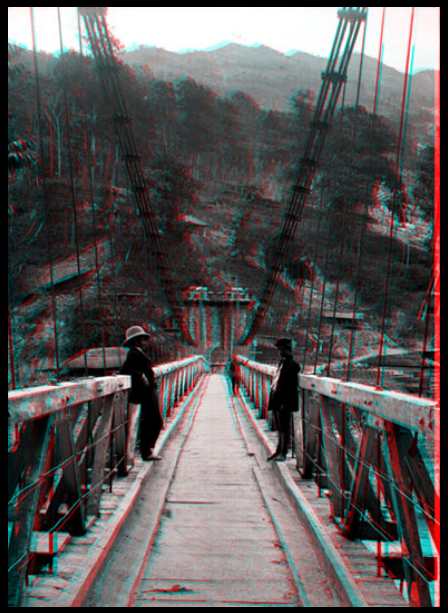

Teesta suspension bridge-Darjeeling, India

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

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

Anaglyphs online

I used to maintain of list of sites, but too hard to keep up to date. Instead, see wikipedia page:

http://en.wikipedia.org/wiki/Anaglyph_image

A free pair of red-blue stereo glasses can be ordered from Rainbow Symphony Inc

- http://www.rainbowsymphony.com/freestuff.html

Stereo

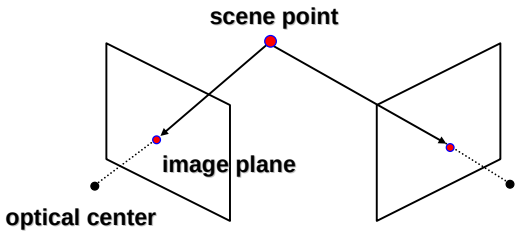

Stereo

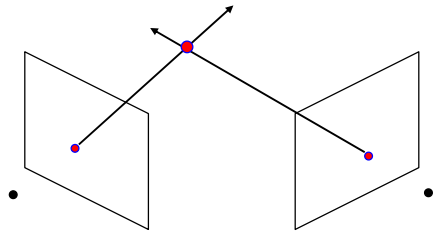

Basic Principle: Triangulation

- Gives reconstruction as intersection of two rays
- Requires
  - camera pose (calibration)
  - ***point correspondence***

Stereo correspondence

Determine Pixel Correspondence

- Pairs of points that correspond to same scene point

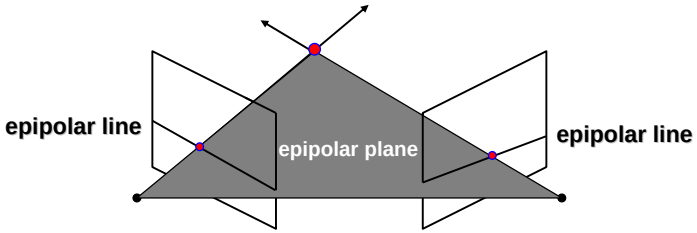

Epipolar Constraint

- 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

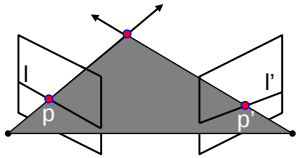

Epipolar relation

- *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 Marc Pollefey’s notes for a nice tutorial

Stereo image rectification

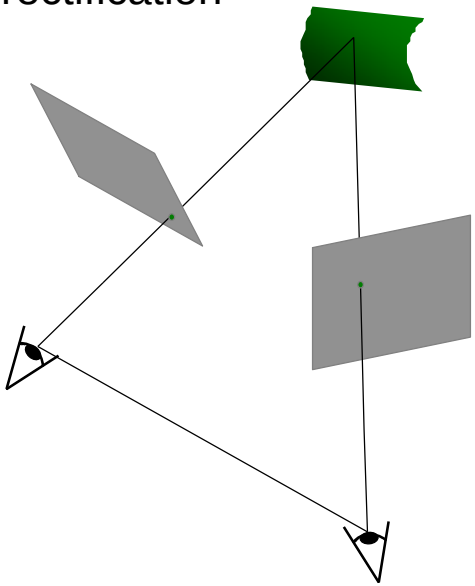

Stereo image rectification

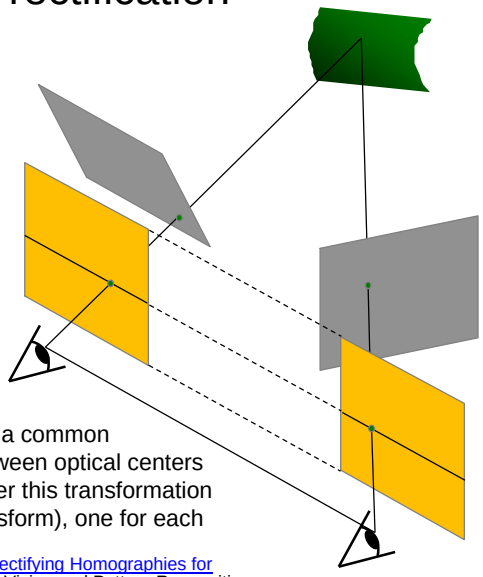

- reproject image planes onto a common plane parallel to the line between optical centers
- pixel motion is horizontal after this transformation
- two homographies (3x3 transform), one for each input image reprojection

➢ 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
  - A good survey and evaluation: http://www.middlebury.edu/stereo/

Your basic stereo algorithm

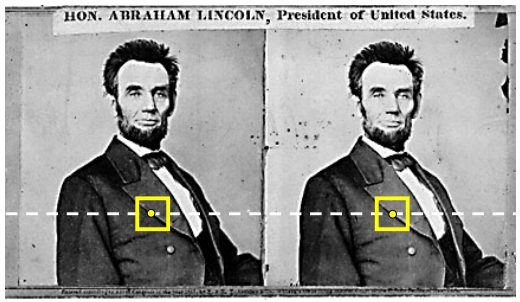

For each epipolar line

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...

Window size

W = 3

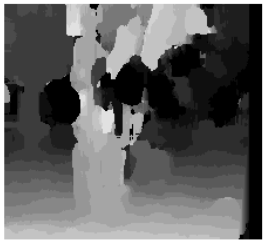

W = 20

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

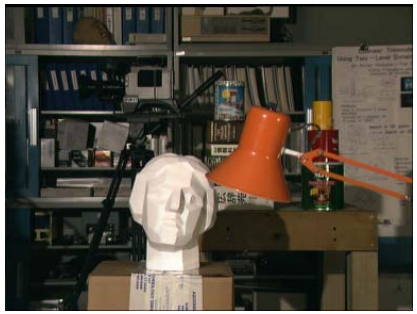

Scene

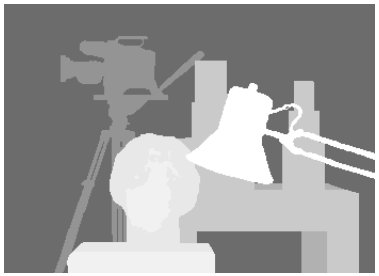

Ground truth

Results with window search

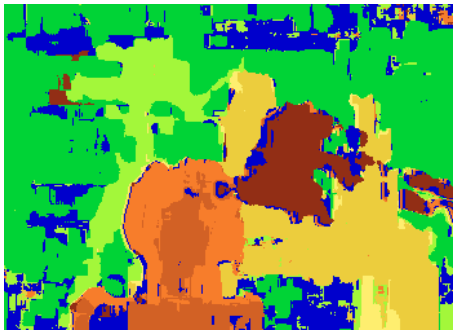

Window-based matching
(best window size)

Ground truth

Better methods exist...

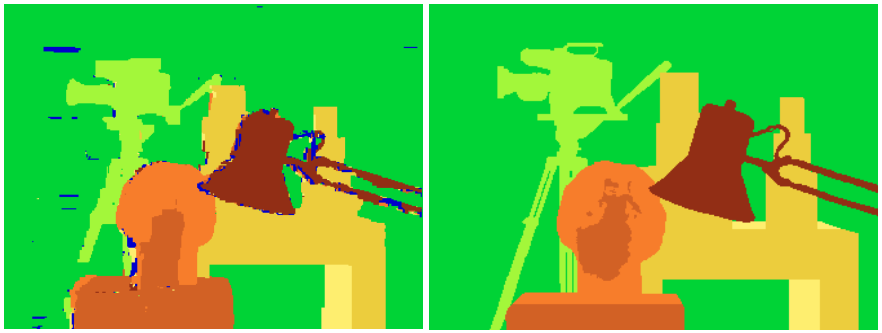

State of the art method

Ground truth

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

For the latest and greatest: http://www.middlebury.edu/stereo/

Stereo as energy minimization

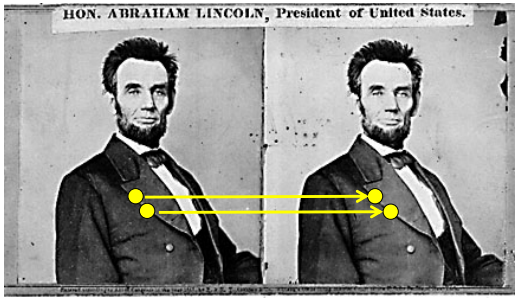

What defines a good stereo correspondence?

1. Match quality
   - Want each pixel to find a good match in the other image
2. Smoothness
   - If two pixels are adjacent, they should (usually) move about the same amount

Stereo as energy minimization

Expressing this mathematically

1. Match quality
   - Want each pixel to find a good match in the other image

$$matchCost = \sum_{x,y} \|I(x,y) - J(x + d_{xy}, y)\|$$

2. Smoothness
   - If two pixels are adjacent, they should (usually) move about the same amount

$$smoothnessCost = \sum_{neighbor\ \ pixels\ \ p,q} |d_p - d_q|$$

We want to minimize $Energy = matchCost + smoothnessCost$

- This is a special type of energy function known as an MRF (Markov Random Field)
  - Effective and fast algorithms have been recently developed:
    - » Graph cuts, belief propagation….
    - » for more details (and code): http://vision.middlebury.edu/MRF/
    - » Great tutorials available online (including video of talks)

Depth from disparity

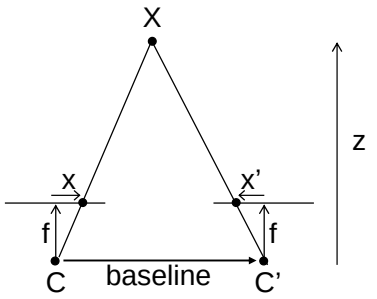

$$disparity = x - x' = \frac{baseline * f}{z}$$

Video View Interpolation

http://research.microsoft.com/users/larryz/videoviewinterpolation.htm

Real-time stereo

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

Used for robot navigation (and other tasks)

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

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

Active stereo with structured light

Li Zhang's one-shot stereo

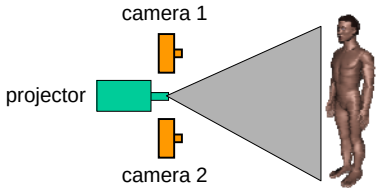

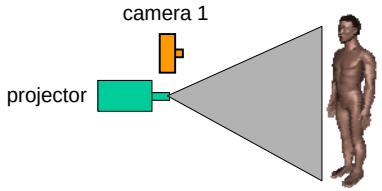

Project "structured" light patterns onto the object

- simplifies the correspondence problem

Active stereo with structured light

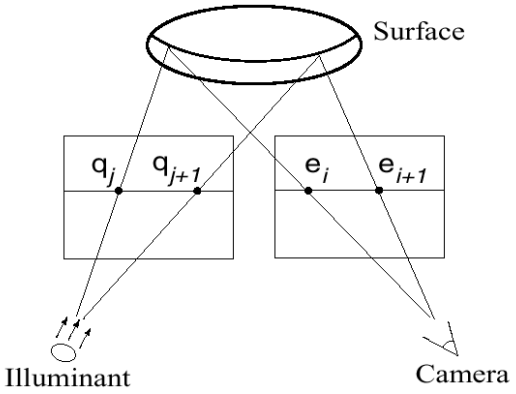

Laser scanning

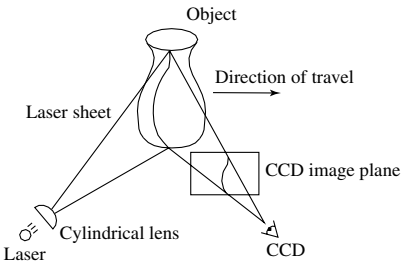

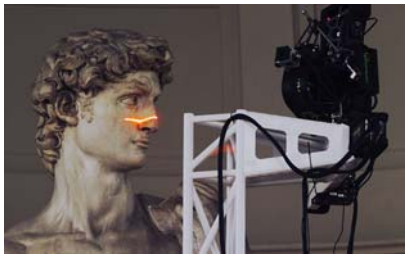

Digital Michelangelo Project
http://graphics.stanford.edu/projects/mich/

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

Laser scanned models

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

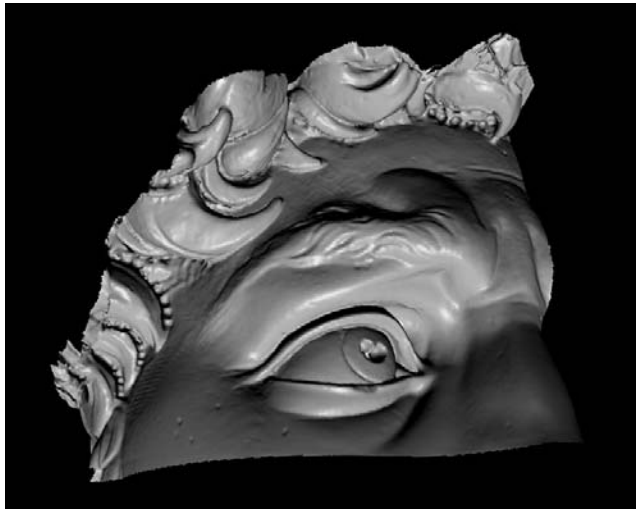

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

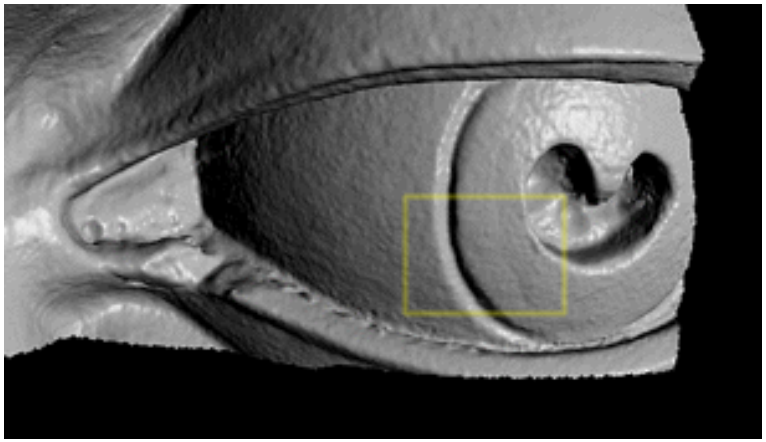

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

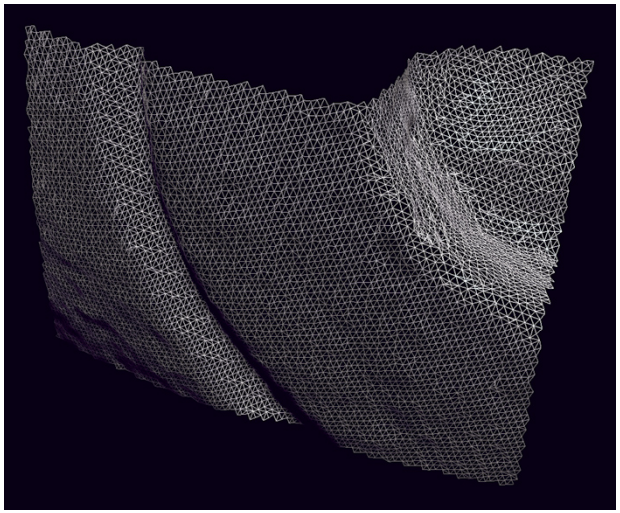

The Digital Michelangelo Project, Levoy et al.

Spacetime Stereo

Li Zhang, Noah Snavely, Brian Curless, Steve Seitz

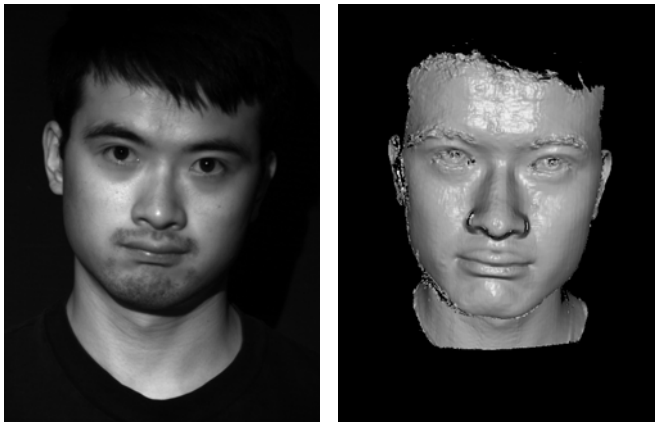

http://grail.cs.washington.edu/projects/stfaces/