Announcements

- Project 2 artifact due Thursday
- Project 3 out today (demo session end of class)

Stereo

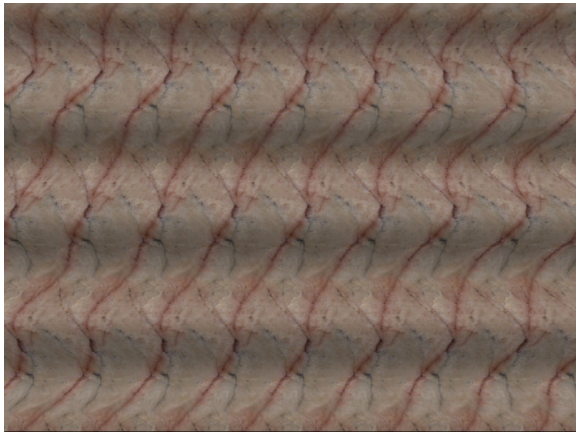

Single image stereogram, by Niklas Een

Readings

- Trucco & Verri, Chapter 7
  - Read through 7.1, 7.2.1, 7.2.2, 7.3.1, 7.3.2, 7.3.7 and 7.4, 7.4.1.
  - The rest is optional.

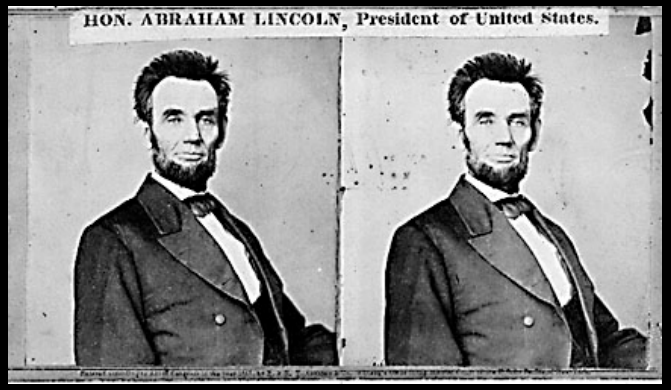

Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923

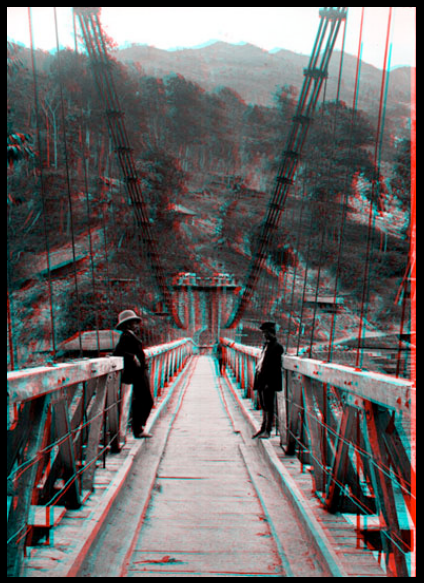

Teesta suspension bridge-Darjeeling, India

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

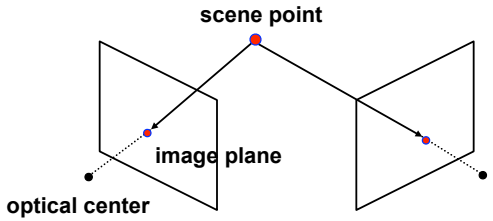

Stereo

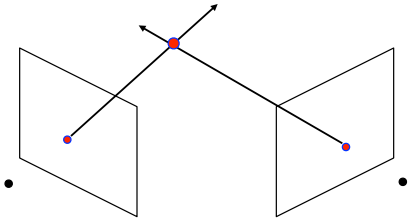

Basic Principle: Triangulation

- 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

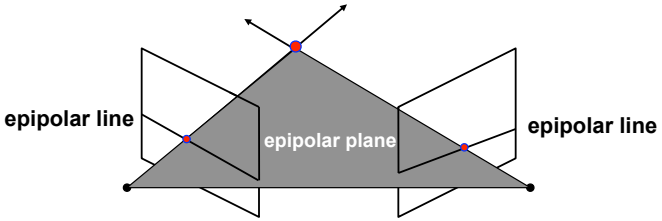

Epipolar Constraint

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

Stereo image rectification

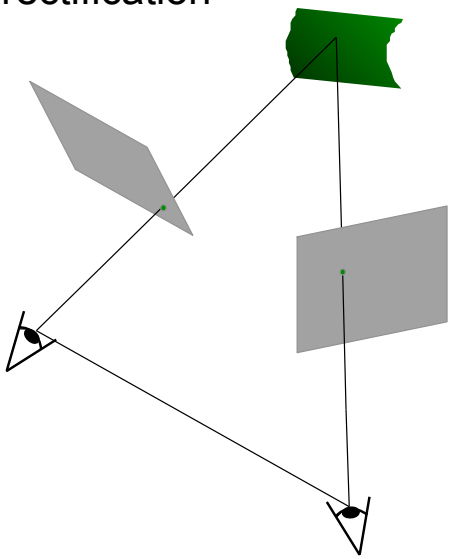

Stereo image rectification

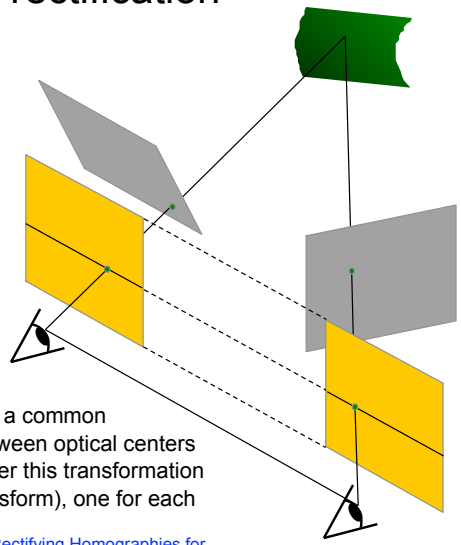

- reproject image planes onto a common plane parallel to the line between optical centers
- pixel motion is horizontal after this transformation
- two homographies (3x3 transform), one for each input image reprojection

➢ C. Loop and Z. Zhang. Computing Rectifying Homographies for Stereo Vision. IEEE Conf. Computer Vision and Pattern Recognition, 1999.

Stereo Matching

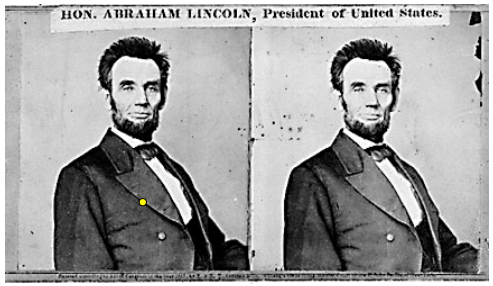

Given a pixel in the left image, how to find its match?

- Assume the photos have been rectified

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

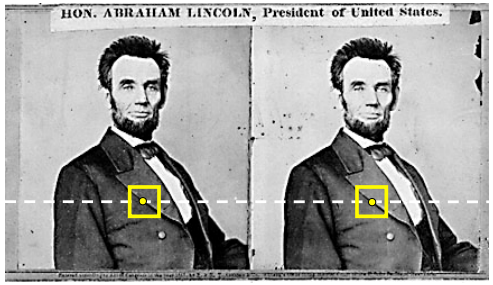

For each epipolar line

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

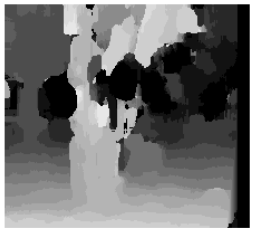

W = 20

Effect of window size

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

Stereo results

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

Scene

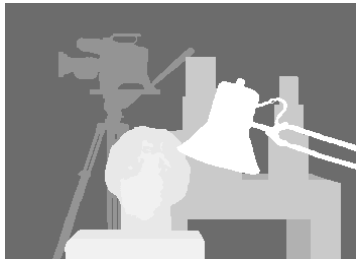

Ground truth

Results with window search

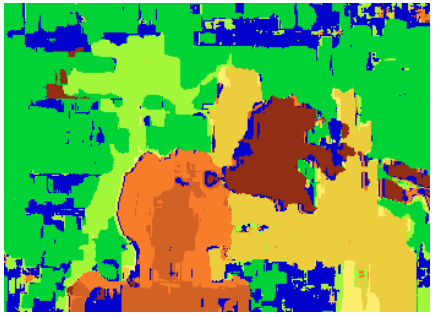

Window-based matching
(best window size)

Ground truth

Better methods exist...

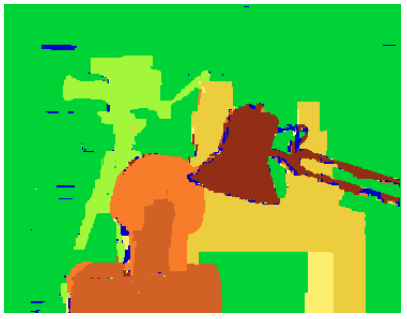

State of the art method

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

Ground truth

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

Stereo as energy minimization

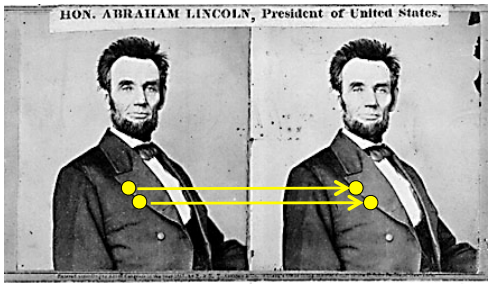

What defines a good stereo correspondence?

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 global optimization

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 sum of these two cost terms

- This is a special type of cost 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/

Depth from disparity

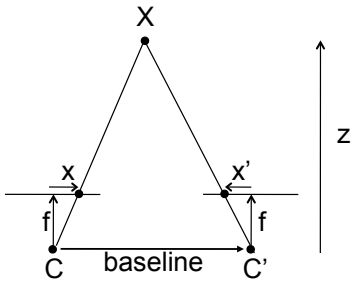

$$disparity = x - x' = \frac{baseline * f}{z}$$

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

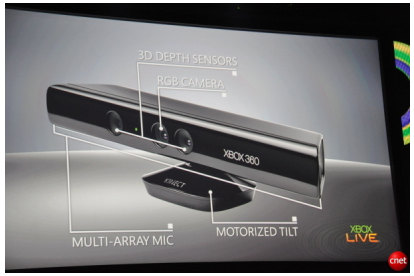

Microsoft's Kinect

http://www.youtube.com/watch?v=7QrnwoO1-8A

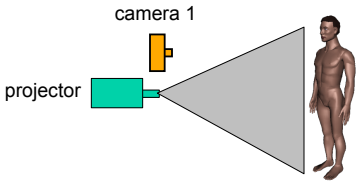

Project "structured" light patterns onto the object

- simplifies the correspondence problem
- can remove one of the cameras (replace with projector)

Active stereo with structured light

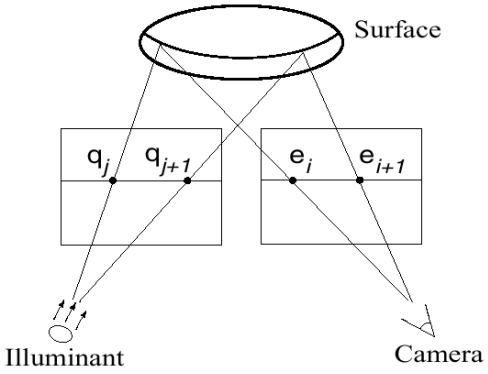

Laser scanning

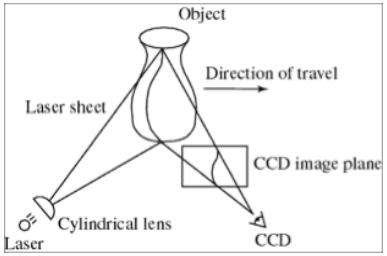

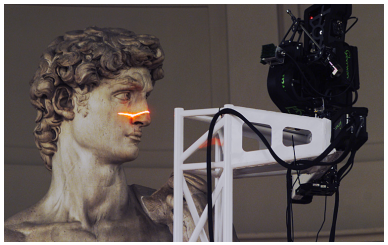

Digital Michelangelo Project
http://graphics.stanford.edu/projects/mich/

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

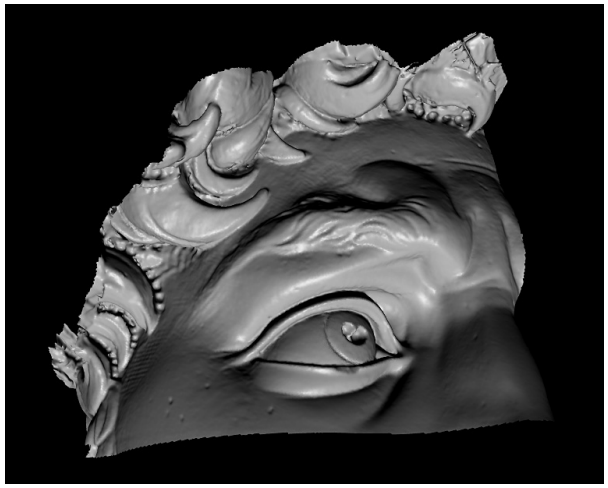

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

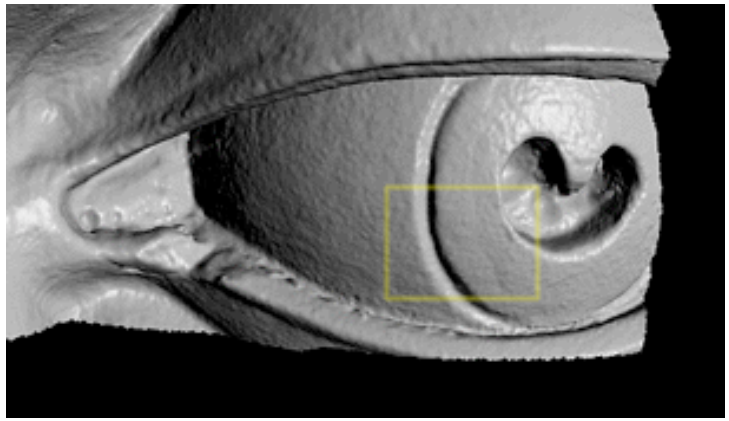

The Digital Michelangelo Project, Levoy et al.

Laser scanned models

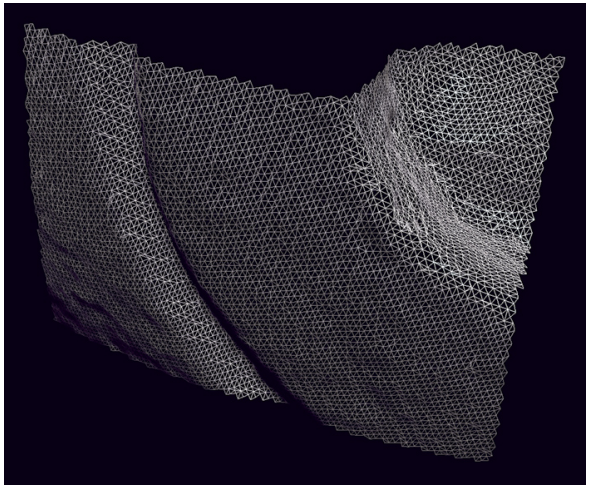

The Digital Michelangelo Project, Levoy et al.

Stereo on Internet photo collections

Furukawa et al.,
http://www.cs.washington.edu/homes/furukawa/