Digital Image Processing

Stereo Images

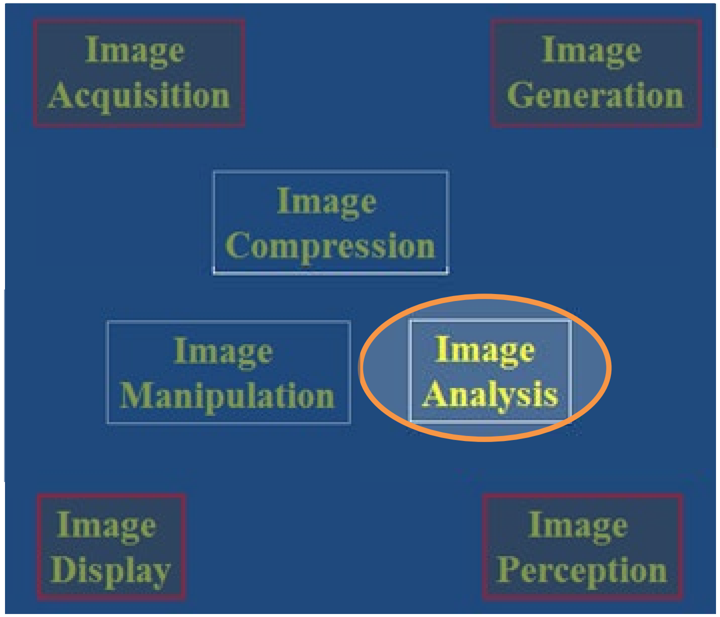

DRAFT

Brent M. Dingle, Ph.D. 2015
Game Design and Development Program
Mathematics, Statistics and Computer Science
University of Wisconsin - Stout

Material in this presentation is largely based on/derived from presentations by: Sventlana Lazebnik, and Noah Snavely

Lecture Objectives

- Previously
  - Image Manipulation and Enhancement
    - Filtering
    - Interpolation
    - Warping
    - Morphing
  - Image Compression
  - Image Analysis
    - Edge Detection
    - Smart Scissors
- Today
  - Stereo Images

Stereo Images, Relation

- Stereo Image Processing
  - Builds from previous topics
  - Requires methods derived from image manipulation and enhancement
  - Relates to feature detection and extraction
  - Relates to human stereoscopic vision
    - Useful in navigating 3D environments
      - i.e. computer vision for robots

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

slide by: Noah Snavely

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

slide by: Noah Snavely

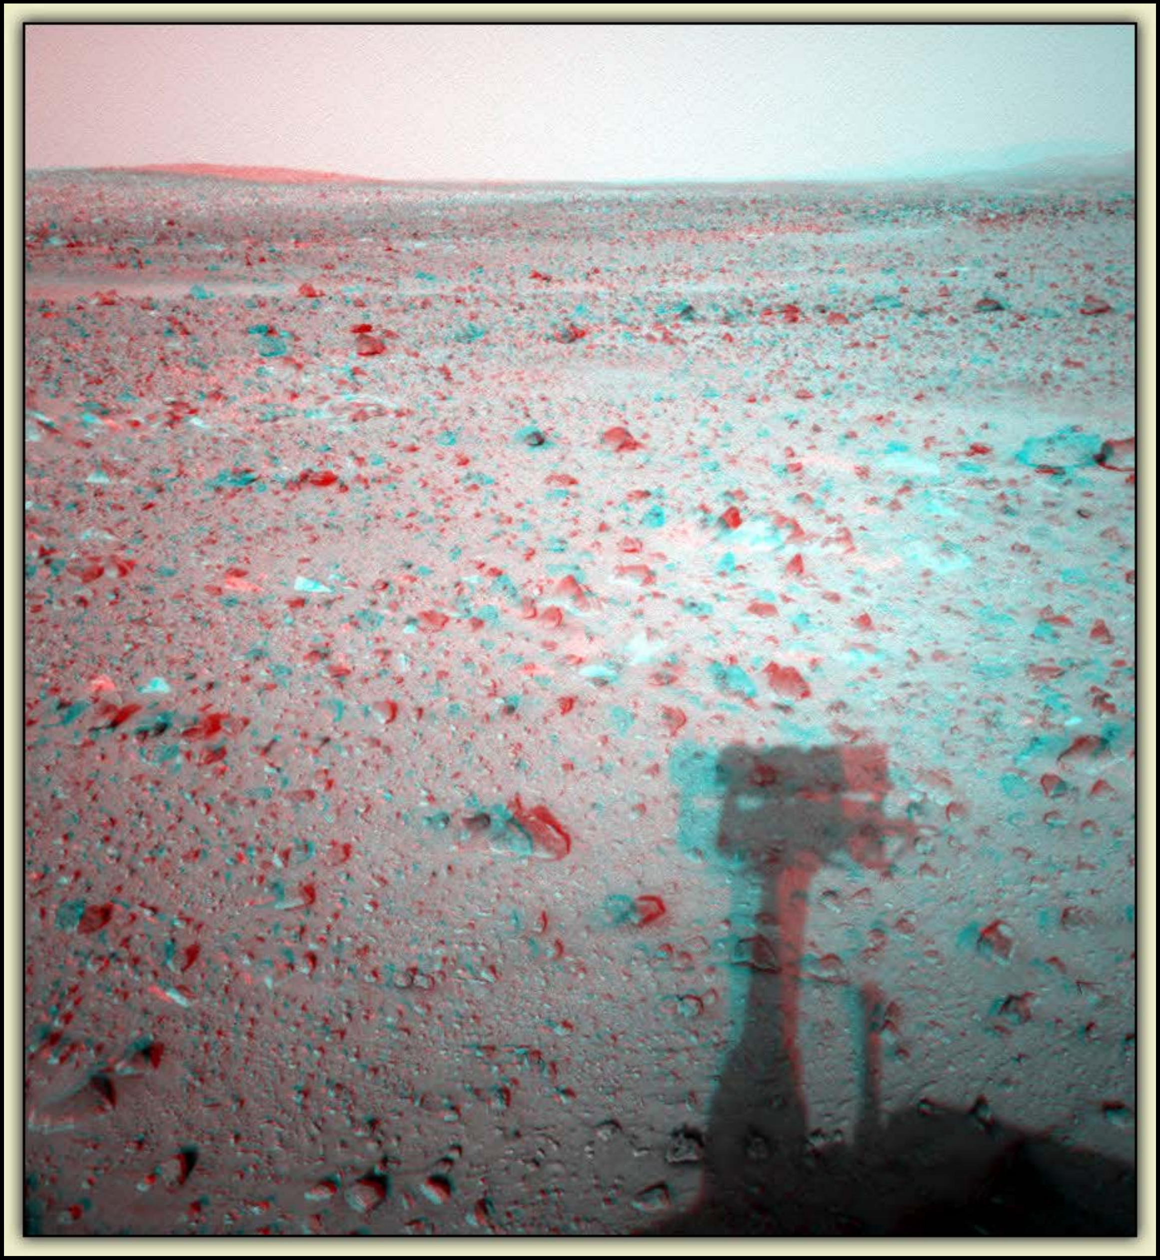

slide by: Noah Snavely

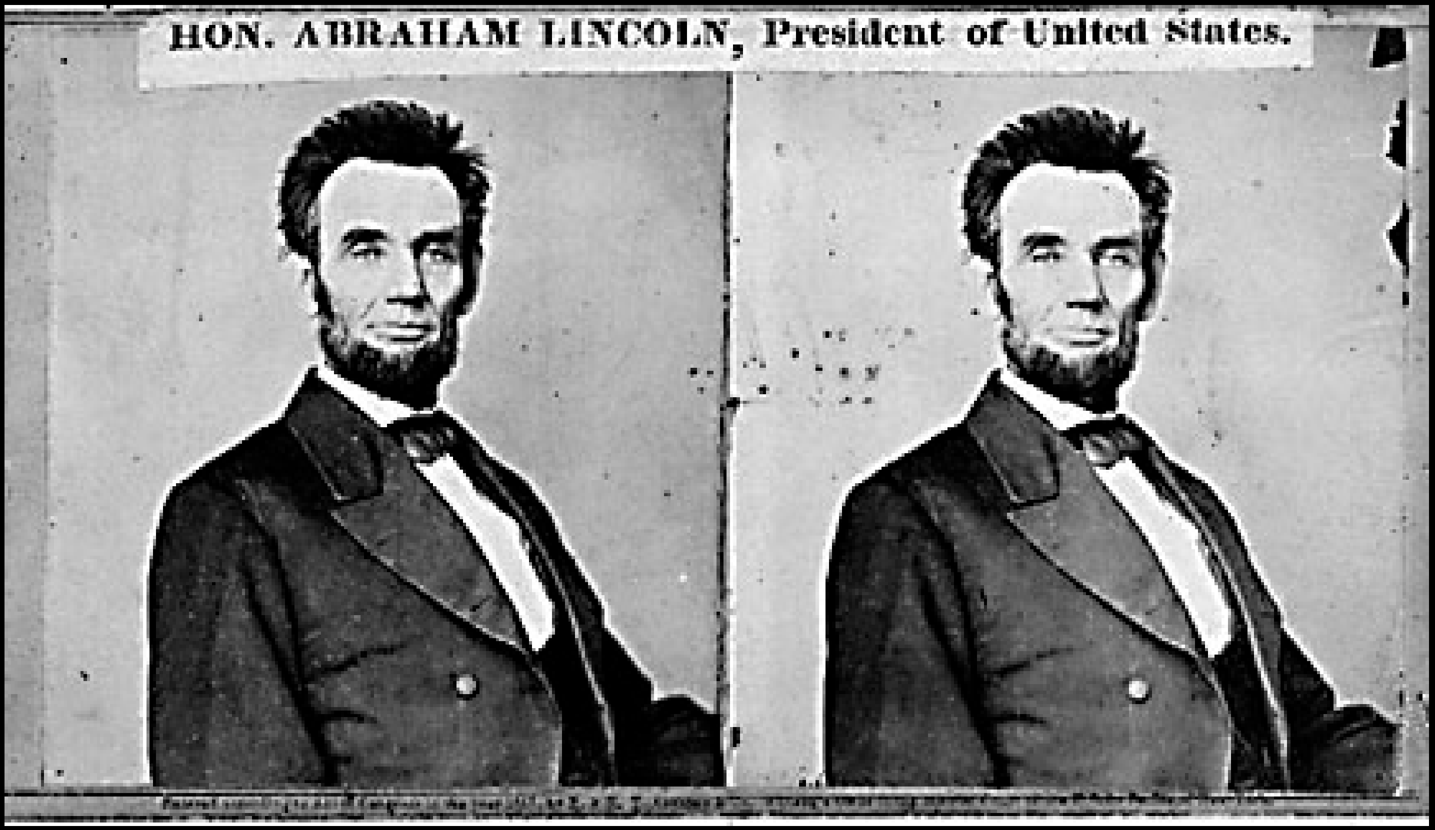

slide by: Noah Snavely

Stereo Image Goal

- Given 2 images from different viewpoints
  - Compute the depth of each point in the image
  - Suggested Method
    - Perform the computation based on how much each pixel moves between the two images

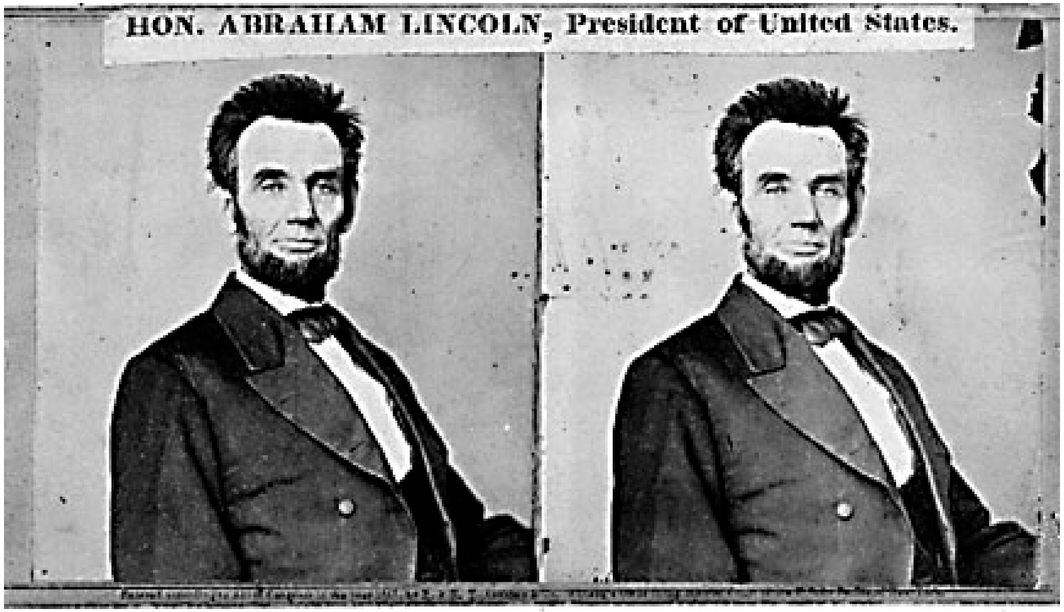

Epipolar Geometry

- x2 – x1 = the disparity of pixel (x1, y1)

Two images captured by a purely horizontal translating camera
(*rectified* stereo pair)

epipolar lines

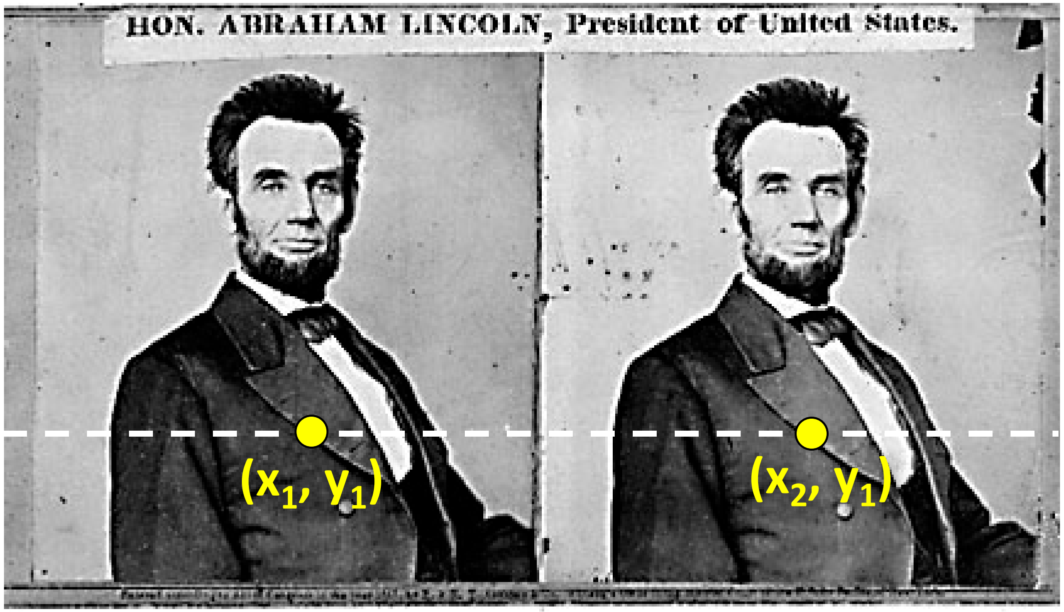

Simple Stereo Algorithm

- For each epipolar line
  - For each pixel in the ‘left’ image
    - Compare it with every pixel in the same epipolar line of the ‘right’ image
    - Select the pixel with minimum cost match

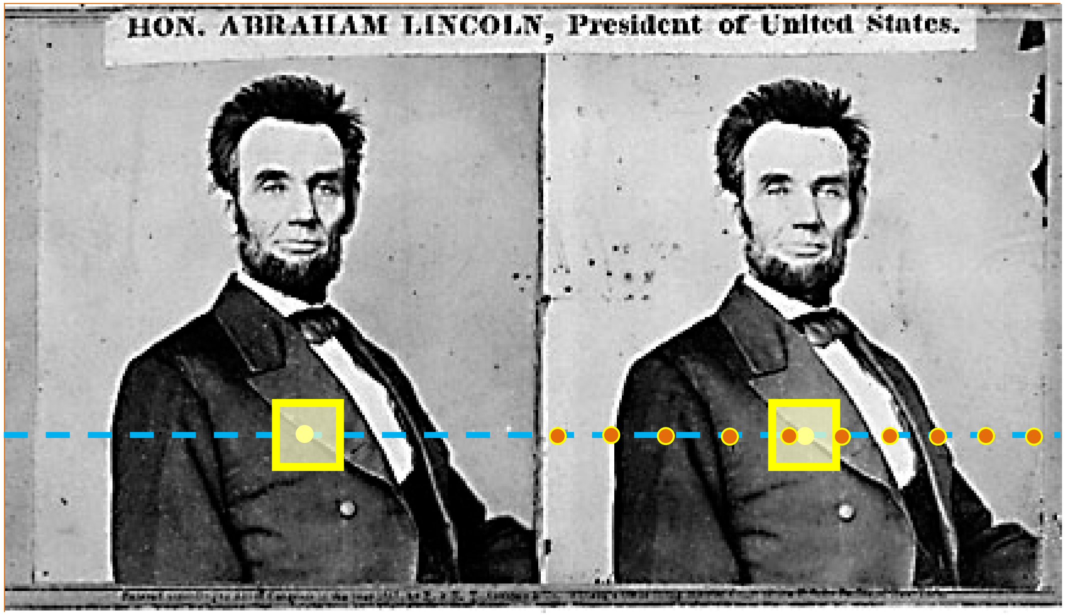

Possible improvement: use matching windows instead of pixels

Window size

W = 3

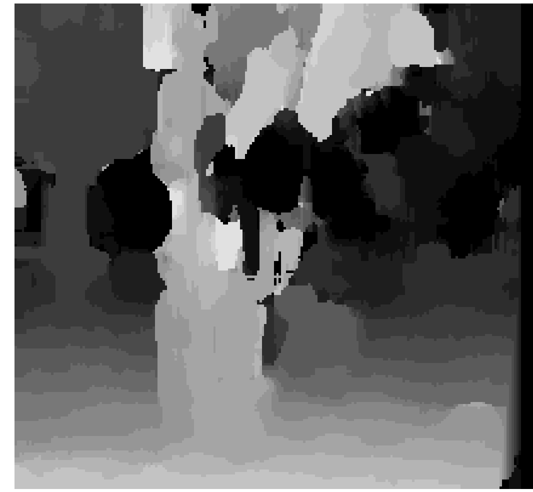

W = 20

Effect of window size

- Smaller window
  - more detail
  - more noise
- Larger window
  - less noise
  - more detail

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

slide by: Noah Snavely

Stereo results

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

Scene

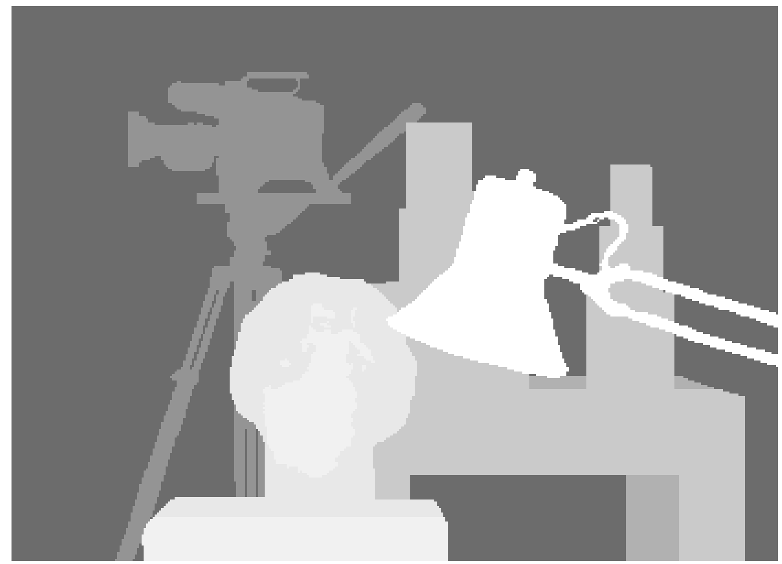

Ground truth

slide by: Noah Snavely

Results with window search

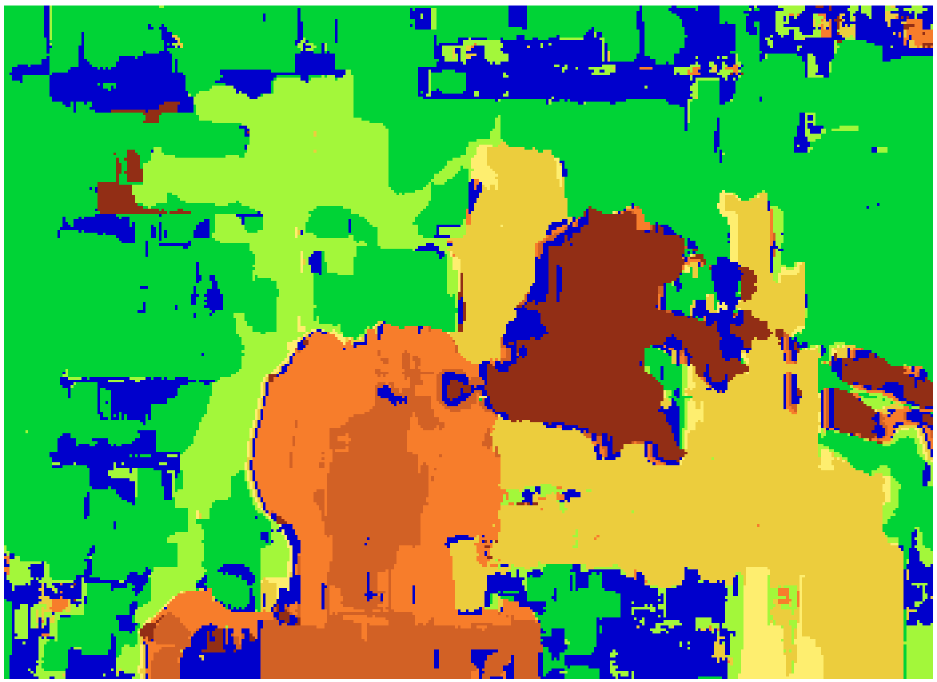

Window-based matching
(best window size)

Ground truth

slide by: Noah Snavely

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

slide by: Noah Snavely

Stereo as Energy Minimization

- What defines a good stereo correspondence?
  - Match quality
    - Each pixel needs a good match in the other image
  - Smoothness
    - Two adjacent pixels in one image should move about the same amount

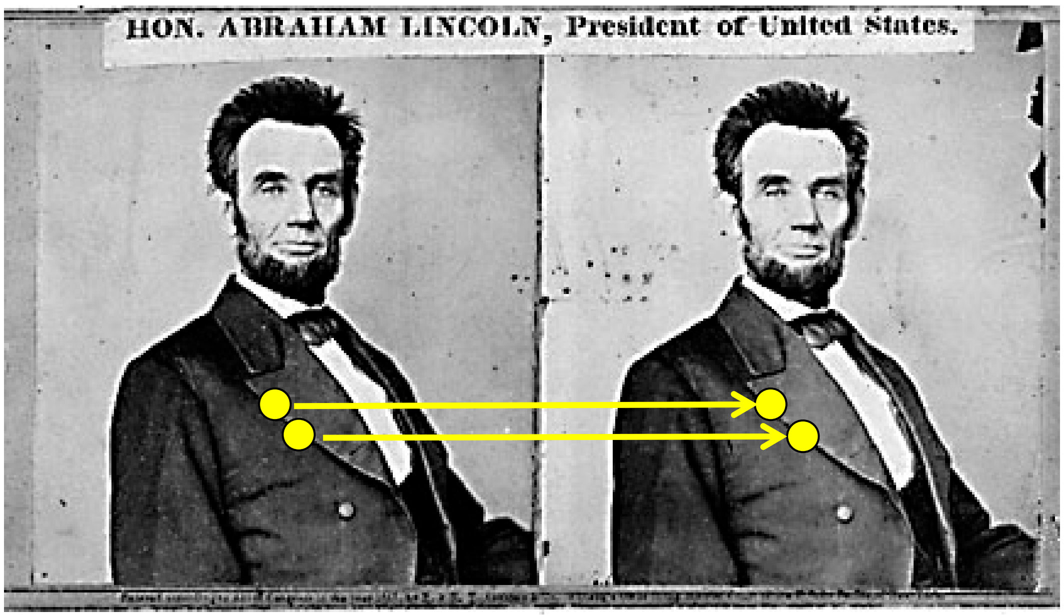

Stereo as Energy Minimization

- Find a disparity map *d* that minimizes an energy function: $E(d)$

- Simple pixel / window matching

$$E(d) = \sum_{(x,y) \in I} C(x, y, d(x, y))$$

$C(x, y, d(x, y)) =$ SSD distance between windows $I(x, y)$ and $J(x + d(x,y), y)$

Stereo as energy minimization

$I(x, y)$

$J(x, y)$

$C(x, y, d)$; the *disparity space image* (DSI)

slide by: Noah Snavely

Stereo as energy minimization

Simple pixel / window matching: choose the minimum of each column in the DSI independently:

$$d(x, y) = \arg\min_{d'} \; C(x, y, d')$$

slide by: Noah Snavely

Stereo as energy minimization

- Better objective function

$$E(d) = \underbrace{E_d(d)}_{\text{match cost}} + \lambda \underbrace{E_s(d)}_{\text{smoothness cost}}$$

match cost ← Want each pixel to find a good match in the other image

smoothness cost ← Adjacent pixels should (usually) move about the same amount

slide by: Noah Snavely

Stereo as energy minimization

$$E(d) = E_d(d) + \lambda E_s(d)$$

match cost: $$E_d(d) = \sum_{(x,y)\in I} C(x,y,d(x,y))$$

smoothness cost: $$E_s(d) = \sum_{(p,q)\in \mathcal{E}} V(d_p, d_q)$$

$\mathcal{E}$: set of neighboring pixels

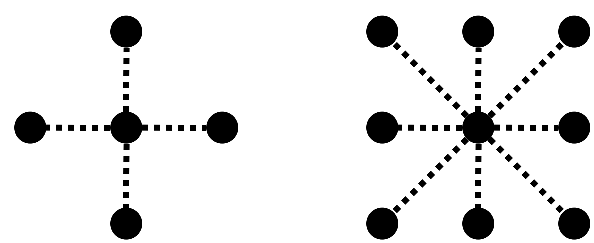

4-connected neighborhood

8-connected neighborhood

slide by: Noah Snavely

Smoothness cost

$$E_s(d) = \sum_{(p,q)\in\mathcal{E}} V(d_p, d_q)$$

How do we choose *V*?

$$V(d_p, d_q) = |d_p - d_q|$$

L_1 distance

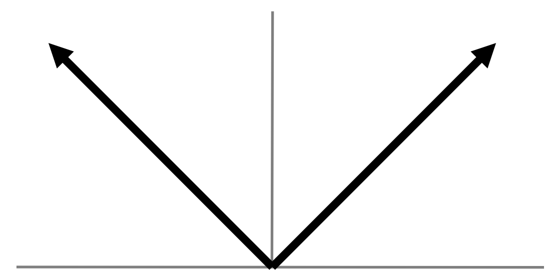

$$V(d_p, d_q) = \begin{cases} 0 & \text{if } d_p = d_q \\ 1 & \text{if } d_p \neq d_q \end{cases}$$

“Potts model”

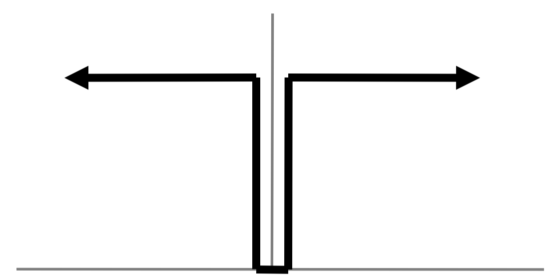

slide by: Noah Snavely

Dynamic programming

$$E(d) = E_d(d) + \lambda E_s(d)$$

- Can minimize this independently per scanline using dynamic programming (DP) 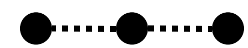

$D(x, y, d)$: minimum cost of solution such that *d*(*x*,*y*) = *d*

$$D(x, y, d) = C(x, y, d) + \min_{d'} \{D(x-1, y, d') + \lambda |d - d'|\}$$

slide by: Noah Snavely

Dynamic programming

- Finds “smooth” path through DPI from left to right

slide by: Noah Snavely

Dynamic Programming

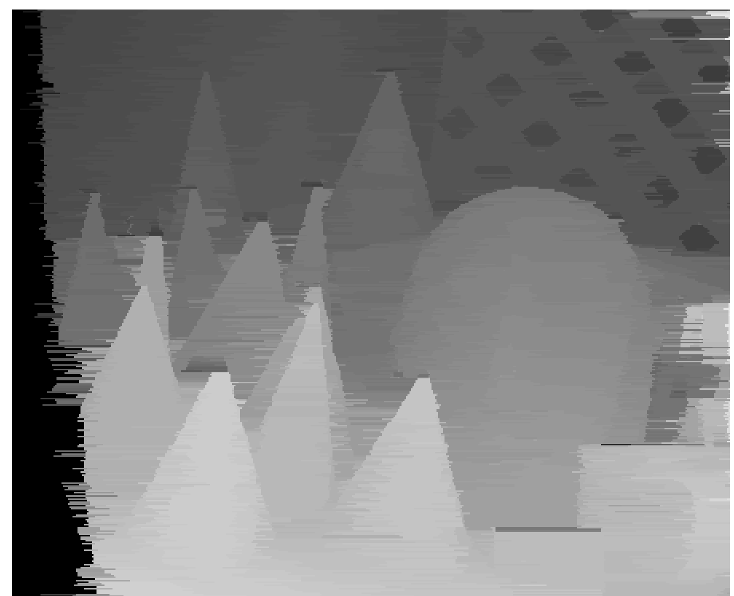

slide by: Noah Snavely

Dynamic programming

- Can we apply this trick in 2D as well?

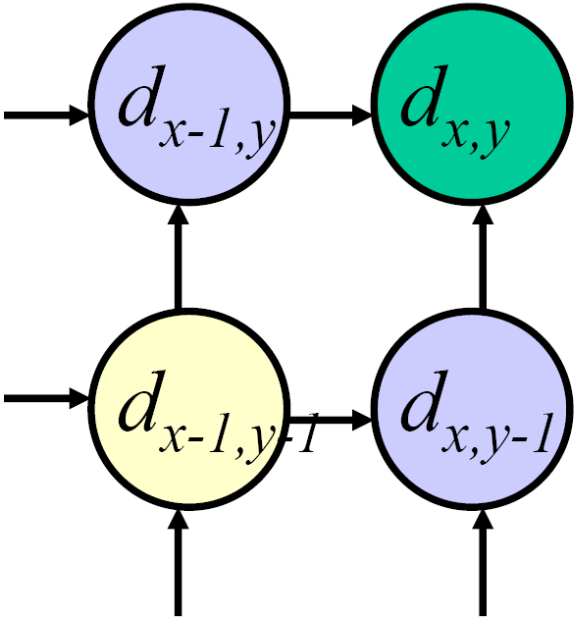

- No: $d_{x,y-1}$ and $d_{x-1,y}$ may depend on different values of $d_{x-1,y-1}$

Slide by: D. Huttenlocher

Stereo as a minimization problem

$$E(d) = E_d(d) + \lambda E_s(d)$$

- The 2D problem has many local minima
  - Gradient descent doesn't work well
- And a large search space
  - *n* x *m* image w/ *k* disparities has k^{nm} possible solutions
  - Finding the global minimum is NP-hard in general
- Good approximations exist…

slide by: Noah Snavely

Depth From Disparity

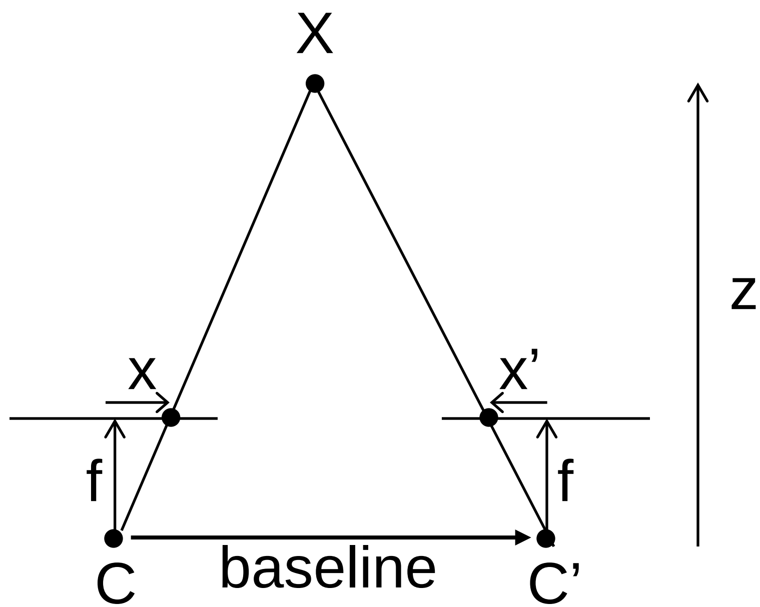

$$disparity = x - x' = \frac{baseline * f}{z}$$

slide by: Noah Snavely

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

slide by: Noah Snavely

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)

slide by: Noah Snavely

Questions?

- Beyond D2L
  - Examples and information can be found online at:
    - *http://docdingle.com/teaching/cs.html*

    - *Continue to more stuff as needed*

Extra Reference Stuff Follows

Credits

- Much of the content derived/based on slides for use with the book:
  - *Digital Image Processing,* Gonzalez and Woods
- Some layout and presentation style derived/based on presentations by
  - Donald House, Texas A&M University, 1999
  - Sventlana Lazebnik, UNC, 2010
  - Noah Snavely, Cornell University, 2012
  - Xin Li, WVU, 2014
  - George Wolberg, City College of New York, 2015
  - Yao Wang and Zhu Liu, NYU-Poly, 2015