

Computer Vision Fall 2018 Columbia University

## Homework

Homework 2 grades are back

- Median 37/40, std 7.2
- Homework 3 due now
- Homework 4 out today

# My Office Hours

### Now Mondays 5pm-6pm

## **Course Evaluations**

• 60% response rate so far

Please respond by tomorrow

• We read all feedback!

# Image Stitching

### Image alignment



Why don't these image line up exactly?

# **Transformation Models**

- Translation only
- Rigid body (translate+rotate)
- Similarity (translate+rotate+scale)
- Affine
- Homography (projective)





# **Camera Projection**



## Camera Matrix

Mapping points from the world to image coordinates is matrix multiplication in homogenous coordinates

$$\begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix} = \begin{pmatrix} C_{11} & C_{12} & C_{13} & C_{14} \\ C_{21} & C_{22} & C_{23} & C_{24} \\ C_{31} & C_{32} & C_{33} & C_{34} \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}$$





## **Two-views of Plane**



Slide credit: Deva Ramanan

## Image Alignment Algorithm

Given images A and B

- 1. Compute image features for A and B
- 2. Match features between A and B
- 3. Compute homography between A and B using least squares on set of matches

What could go wrong?



Slide credit: Noah Snavely

### Robustness

• Let's consider a simpler example... linear regression



Problem: Fit a line to these datapoints

How can we fix this?

Slide credit: Noah Snavely

### We need a better cost function...

• Suggestions?

## **Counting inliers**



### **Counting inliers**



**Inliers: 3** 

### **Counting inliers**



## Idea

- Given a hypothesized line
- Count the number of points that "agree" with the line
  - "Agree" = within a small distance of the line
  - I.e., the inliers to that line
- For all possible lines, select the one with the largest number of inliers

### How do we find the best line?

Unlike least-squares, no simple closed-form solution

- Hypothesize-and-test
  - Try out many lines, keep the best one
  - Which lines?

(RANdom SAmple Consensus) :

Fischler & Bolles in '81.



#### Algorithm:

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

Line fitting example



Algorithm:

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

Line fitting example



#### Algorithm:

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

Line fitting example



#### Algorithm:

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



Algorithm:

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

### RANSAC for alignment



Slide credit: Deva Ramanan

### RANSAC for alignment



Slide credit: Deva Ramanan

### **RANSAC** for alignment



Slide credit: Deva Ramanan

### Implementing image warping

Given a coordinate xform (x',y') = T(x,y) and a source image f(x,y), how do we compute an xformed image g(x',y') = f(T(x,y))?



### Forward Warping

- Send each pixel f(x) to its corresponding
  location (x',y') = T(x,y) in g(x',y')
  - What if pixel lands "between" two pixels?



### Inverse Warping

- Get each pixel g(x',y') from its corresponding location (x,y) = T<sup>-1</sup>(x,y) in f(x,y)
  - Requires taking the inverse of the transform
  - What if pixel comes from "between" two pixels?



### Inverse Warping

- Get each pixel g(x') from its corresponding location x' = h(x) in f(x)
  - What if pixel comes from "between" two pixels?
  - Answer: *resample* color value from *interpolated* (*prefiltered*) source image



## Blending

• We've aligned the images – now what?



## Blending

• Want to seamlessly blend them together



### Image Blending


## Feathering



#### Effect of window size









### Effect of window size









## Blending



#### Moving object, simple blending => blur

# Blending

Instead of blending high frequencies along a straight line, blend along line of minimum differences in image intensities



## Blending



#### Minimum-cost cut 📰 no blur



# Stereo vision







# Why not put our second eye here?

#### Stereoscopes: A 19<sup>th</sup> Century Pastime







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

#### **3D** Movies



#### (WALT DISNEP

## JONAS BROTHERS

HOLE AND A CALL AND A CONTRACT OF THE ADDRESS OF



#### Depth without objects Random dot stereograms (Bela Julesz)



1	Ð	1	0	1	0	0	1	0	1
1	0	0	1	0	1	Q	1	0	0
0	0	1	1	0	1	1	0	1	0
٥	1	0	Y	A	A	8	8	0	1
1	1	1	х	8	A	₿	А.	0	1
¢	0	1	х	A	А	в	A	1	0
1	1	1	Y	8	8	А	ß	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

									_
1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	A	A	8	8	×	٥	1
1	1	1	9	А	8	A	Y	0	-
0	0	1	А	A	8	A	Y	1	0
1	1	1	в	в	A	8	×	0	1
1	0	0	1	1	0	1	1	Ö	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

**Julesz**, 1971



#### Stereo



- Given two images from different viewpoints
  - How can we compute the depth of each point in the image?
  - Based on how much each pixel moves between the two images







55 Slide credit: Antonio Torralba



Slide credit: Antonio Torralba

56



Slide credit: Antonio Torralba

57



Slide credit: Antonio Torralba



Similar triangles:

$$\frac{T+X_L-X_R}{Z-f} = \frac{T}{Z}$$



Slide credit: Antonio Torralba

60

## Epipolar geometry



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

#### $x_2 - x_1 =$ the *disparity* of pixel ( $x_1, y_1$ )

## 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* 

#### Stereo matching based on SSD



### Window size







$$N = 3$$

W = 20

#### Effect of window size

- Smaller window
  - +
  - •
- Larger window
  - +

#### Better results with adaptive window

- T. Kanade and M. Okutomi, <u>A Stereo Matching Algorithm</u> <u>with an Adaptive Window: Theory and Experiment</u>,, Proc. International Conference on Robotics and Automation, 1991.
- D. Scharstein and R. Szeliski. <u>Stereo matching with</u> <u>nonlinear diffusion</u>. 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





Ground truth

Scene

#### 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., <u>Fast Approximate Energy Minimization via Graph Cuts</u>, International Conference on Computer Vision, September 1999.

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



- 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

- Find disparity map d that minimizes an energy function  ${\cal E}(d)$
- Simple pixel / window matching

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

 $C(x, y, d(x, y)) = \frac{\text{SSD distance between windows}}{I(x, y) \text{ and } J(x + d(x, y), y)}$ 





J(x, y)





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

$$d(x, y) = \underset{d'}{\operatorname{arg\,min}} C(x, y, d')$$

#### Greedy selection of best match


## Stereo as energy minimization

Better objective function



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

$$\begin{aligned} \text{Smoothness cost} \\ E_s(d) &= \sum_{(p,q)\in\mathcal{E}} V(d_p, d_q) \\ \text{How do we choose V?} \\ V(d_p, d_q) &= |d_p - d_q| \\ L_1 \text{ distance} \\ V(d_p, d_q) &= \begin{cases} 0 & \text{if } d_p = d_q \\ 1 & \text{if } d_p \neq d_q \end{cases} \end{aligned}$$

Slide credit: Noah Snavely

## Dynamic programming $E(d) = E_d(d) + \lambda E_s(d)$

- Can minimize this independently per scanline using dynamic programming (DP)
- Basic idea: incrementally build a table of costs
   D one column at a time

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

Base case: D(0, y, i) = C(0, y, i), i = 0, ..., L (L = max disparity) Recurrence:  $D(x, y, i) = C(x, y, i) + \min_{i \in \{0, 1, ..., L\}} D(x - 1, y, j) + \lambda |i - j|$ 

Slide credit: Noah Snavely

#### Dynamic programming



 Finds "smooth", low-cost path through DPI from left to right

#### **Dynamic Programming**









# 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
  - $-n \ge m$  image w/k disparities has  $k^{nm}$  possible solutions
  - Finding the global minimum is NP-hard in general

## Stereo correspondence constraints



If we see a point in camera 1, are there any constraints on where we will find it on camera 2?

## **Epipolar constraint**







Baseline: the line connecting the two camera centers

Epipole: point of intersection of baseline with the image plane



Baseline: the line connecting the two camera centers

Epipole: point of intersection of *baseline* with the image plane



Baseline: the line connecting the two camera centers

**Epipole**: point of intersection of *baseline* with the image plane

Epipolar plane: the plane that contains the two camera centers and a 3D point in the world



Baseline: the line connecting the two camera centers

**Epipole**: point of intersection of *baseline* with the image plane

Epipolar plane: the plane that contains the two camera centers and a 3D point in the world

**Epipolar line**: intersection of the *epipolar plane* with each image plane

## Epipolar constraint



We can search for matches across epipolar lines

All epipolar lines intersect at the epipoles

## The essential matrix



If we observe a point in one image, its position in the other image is constrained to lie on line defined by above.



E: essential matrix

p, p': image points in homogeneous coordinates

Slide credit: Antonio Torralba

#### Real-time stereo



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

- Used for robot navigation (and other tasks)
  - Several 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
  - basis for active depth sensors, such as Kinect and iPhone X (using IR)

#### Active stereo with structured light



https://ios.gadgethacks.com/news/watch-iphone-xs-30k-ir-dots-scan-your-face-0180944/

#### 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







