

## Announcements

---

### Project 2

- Goes out today! James will do a 10 min. demo session at the end of the class.

## Mosaics

---




[StreetView](#)  
[Photosynth App](#)

### Today's Readings

- Szeliski and Shum paper (sections 1 and 2, skim the rest)  
– <http://www.cs.washington.edu/education/courses/455/08wi/readings/szeliskiShum07.pdf>

## Image Mosaics

---



### Goal

- Stitch together several images into a seamless composite

## How to do it?

---

### Basic Procedure

- Take a sequence of images from the same position
  - Rotate the camera about its optical center
- Compute transformation between second image and first
- Shift the second image to overlap with the first
- Blend the two together to create a mosaic
- If there are more images, repeat

## Aligning images




How to account for warping?

- Translations are not enough to align the images

## Alignment Demo



## Image reprojection



The mosaic has a natural interpretation in 3D

- The images are reprojected onto a common plane
- The mosaic is formed on this plane


## Image reprojection

Basic question

- How to relate two images from the same camera center?
  - how to map a pixel from PP1 to PP2


Answer

- Cast a ray through each pixel in PP1
- Draw the pixel where that ray intersects PP2



Don't need to know what's in the scene!

## Image reprojection



### Observation

- Rather than thinking of this as a 3D reprojection, think of it as a 2D image warp from one image to another

## Homographies

### Perspective projection of a plane

- Lots of names for this:
  - **homography**, texture-map, colineation, planar projective map
- Modeled as a 2D warp using homogeneous coordinates


$$\begin{bmatrix} wx' \\ wy' \\ w \end{bmatrix} = \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$\mathbf{p}' \quad \mathbf{H} \quad \mathbf{p}$

### To apply a homography $\mathbf{H}$

- Compute  $\mathbf{p}' = \mathbf{H}\mathbf{p}$  (regular matrix multiply)
- Convert  $\mathbf{p}'$  from homogeneous to image coordinates
  - divide by  $w$  (third) coordinate

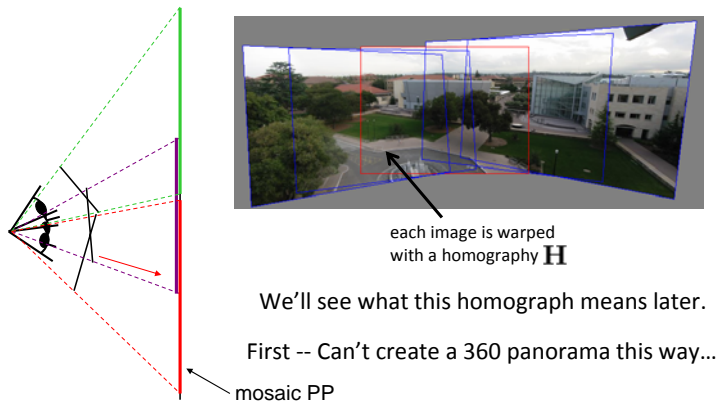
## Image warping with homographies



## Homographies




## Idea: projecting images onto a common plane




## Panoramas


What if you want a 360° field of view?




## Panoramas





## Spherical projection systems



## Spherical projection





- Map 3D point (X,Y,Z) onto sphere
 
$$(\hat{x}, \hat{y}, \hat{z}) = \frac{1}{\sqrt{X^2+Y^2+Z^2}}(X, Y, Z)$$
- Convert to spherical coordinates
 
$$(\sin\theta\cos\phi, \sin\phi, \cos\theta\cos\phi) = (\hat{x}, \hat{y}, \hat{z})$$
- Convert to spherical image coordinates
 
$$(\tilde{x}, \tilde{y}) = (s\theta, s\phi) + (\tilde{x}_c, \tilde{y}_c)$$
  - s defines size of the final image
  - » often convenient to set s = camera focal length



## Spherical reprojection

How to map sphere onto a flat image?



- $(\hat{x}, \hat{y}, \hat{z})$  to  $(x', y')$




## Spherical reprojection

How to map sphere onto a flat image?

- $(\hat{x}, \hat{y}, \hat{z})$  to  $(x', y')$
- Use image projection matrix!
- or use the version of projection that properly accounts for radial distortion, as discussed in projection slides. This is what you'll do for project 2.



## Spherical reprojection



Map image to spherical coordinates

- need to know the focal length

## Aligning spherical images



Suppose we rotate the camera by  $\theta$  about the vertical axis

- How does this change the spherical image?

## Aligning spherical images



Suppose we rotate the camera by  $\theta$  about the vertical axis

- How does this change the spherical image?
- Translation by  $\theta$
- This means that we can align spherical images by translation

## Spherical image stitching



What if you don't know the camera rotation?

- Solve for the camera rotations
  - Note that a pan (rotation) of the camera is a **translation** of the sphere!
  - Use **feature matching to solve for translations of spherical-warped images**


## Computing transformations

- Given a set of matches between images A and B
  - How can we compute the transform T from A to B?

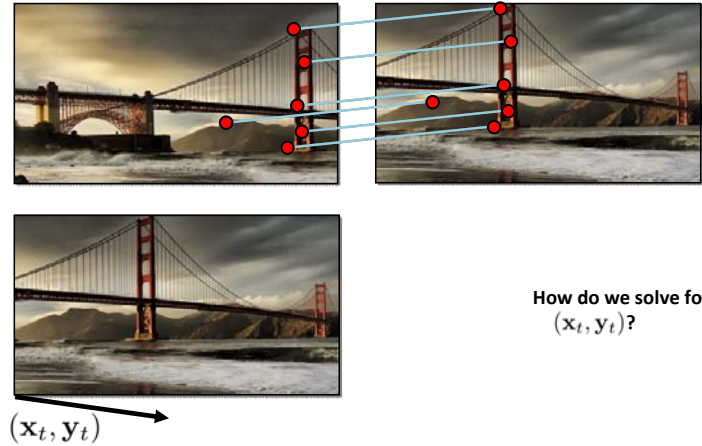


- Find transform T that best “agrees” with the matches


## Computing transformations



## Simple case: translations




## Simple case: translations



Displacement of match  $i = (x'_i - x_i, y'_i - y_i)$

$$(x_t, y_t) = \left( \frac{1}{n} \sum_{i=1}^n x'_i - x_i, \frac{1}{n} \sum_{i=1}^n y'_i - y_i \right)$$

## Simple case: translations




$$x_i + x_t = x'_i$$

$$y_i + y_t = y'_i$$

- System of linear equations
  - What are the knowns? Unknowns?
  - How many unknowns? How many equations (per match)?



## Simple case: translations



$$\mathbf{x}_i + \mathbf{x}_t = \mathbf{x}'_i$$

$$\mathbf{y}_i + \mathbf{y}_t = \mathbf{y}'_i$$

- Problem: more equations than unknowns
  - “Overdetermined” system of equations
  - We will find the *least squares* solution

## Least squares formulation

- For each point  $(\mathbf{x}_i, \mathbf{y}_i)$

$$\mathbf{x}_i + \mathbf{x}_t = \mathbf{x}'_i$$

$$\mathbf{y}_i + \mathbf{y}_t = \mathbf{y}'_i$$

- we define the *residuals* as

$$r_{\mathbf{x}_i}(\mathbf{x}_t) = (\mathbf{x}_i + \mathbf{x}_t) - \mathbf{x}'_i$$

$$r_{\mathbf{y}_i}(\mathbf{y}_t) = (\mathbf{y}_i + \mathbf{y}_t) - \mathbf{y}'_i$$

## Least squares formulation

- Goal: minimize sum of squared residuals

$$C(\mathbf{x}_t, \mathbf{y}_t) = \sum_{i=1}^n (r_{\mathbf{x}_i}(\mathbf{x}_t)^2 + r_{\mathbf{y}_i}(\mathbf{y}_t)^2)$$

- “Least squares” solution
- For translations, is equal to mean displacement

## Least squares formulation

- Can also write as a matrix equation

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ \vdots & \vdots \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} x'_1 - x_1 \\ y'_1 - y_1 \\ x'_2 - x_2 \\ y'_2 - y_2 \\ \vdots \\ x'_n - x_n \\ y'_n - y_n \end{bmatrix}$$

$$\mathbf{A}_{2n \times 2} \mathbf{t}_{2 \times 1} = \mathbf{b}_{2n \times 1}$$



## Least squares

$$\mathbf{A}\mathbf{t} = \mathbf{b}$$

- Find  $\mathbf{t}$  that minimizes


$$\|\mathbf{A}\mathbf{t} - \mathbf{b}\|^2$$

- To solve, form the *normal equations*


$$\mathbf{A}^T \mathbf{A} \mathbf{t} = \mathbf{A}^T \mathbf{b}$$

$$\mathbf{t} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$


## But not all matches are good




## Random Sample Consensus



## Random Sample Consensus



## Least squares fit



## RANSAC

- Idea:
  - All the inliers will agree with each other on the translation vector; the (hopefully small) number of outliers will (hopefully) disagree with each other
    - RANSAC only has guarantees if there are  $< 50\%$  outliers
  - "All good matches are alike; every bad match is bad in its own way."
    - Tolstoy via Alyosha Efras

## RANSAC

Same basic approach works for any transformation

- Translation, rotation, homographies, etc.
- Very useful tool

### General version


- Randomly choose a set of  $K$  correspondences
  - Typically  $K$  is the minimum size that lets you fit a model
- Fit a model (e.g., homography) to those correspondences
- Count the number of inliers that "approximately" fit the model
  - Need a threshold on the error
- Repeat as many times as you can
- Choose the model that has the largest set of inliers
- Refine the model by doing a least squares fit using ALL of the inliers

## How many rounds?

- If we have to choose  $s$  samples each time
  - with an outlier ratio  $e$
  - and we want the right answer with probability  $p$

$s$	proportion of outliers $e$						
	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177


$p = 0.99$



Source: M. Pollefeys


## How big is $s$ ?

- For alignment, depends on the motion model
  - Here, each sample is a correspondence (pair of matching points)




Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} I & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} sR & t \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} A \\ 0 & 1 \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\tilde{H}_{3 \times 3}$	8	straight lines	

## Assembling the panorama



Stitch pairs together, blend, then crop


## Problem: Drift



### Error accumulation

- small errors accumulate over time

## Problem: Drift



### Solution

- add another copy of first image at the end
- this gives a constraint:  $y_n = y_1$
- there are a bunch of ways to solve this problem
  - add displacement of  $(y_1 - y_n)/(n - 1)$  to each image after the first
  - compute a global warp:  $y' = y + ax$
  - run a big optimization problem, incorporating this constraint
    - best solution, but more complicated
    - known as "bundle adjustment"

## Full-view Panorama

---



## Different projections are possible

---



## Project 2

---


Take pictures on a tripod (or handheld)  
Warp to spherical coordinates  
Extract features  
Align neighboring pairs using RANSAC  
Write out list of neighboring translations  
Correct for drift  
Read in warped images and blend them  
Crop the result and import into a viewer

Roughly based on **Autostitch**


- By Matthew Brown and David Lowe
- <http://www.cs.ubc.ca/~mbrown/autostitch/autostitch.html>

## Image Blending


---




## Feathering



## Effect of window (ramp-width) size



## Effect of window size




## Good window size



“Optimal” window: smooth but not ghosted

- Doesn't always work...

## Pyramid blending



Create a Laplacian pyramid, blend each level

- Burt, P. J. and Adelson, E. H., [A multiresolution spline with applications to image mosaics](http://www.cs.cmu.edu/~burt/papers/multires.pdf), ACM Transactions on Graphics, 42(4), October 1983, 217-236.


## Poisson Image Editing



For more info: Perez et al, SIGGRAPH 2003

- [http://research.microsoft.com/vision/cambridge/papers/perez\\_siggraph03.pdf](http://research.microsoft.com/vision/cambridge/papers/perez_siggraph03.pdf)

## Alpha Blending



Optional: see Blinn (CGA, 1994) for details:  
<http://www.cse.cmu.edu/~blinn/teaching/463/lectures/lec13/alpha.html>

Encoding blend weights:  $I(x,y) = (\alpha R, \alpha G, \alpha B, \alpha)$


color at  $p = \frac{(\alpha_1 R_1, \alpha_1 G_1, \alpha_1 B_1) + (\alpha_2 R_2, \alpha_2 G_2, \alpha_2 B_2) + (\alpha_3 R_3, \alpha_3 G_3, \alpha_3 B_3)}{\alpha_1 + \alpha_2 + \alpha_3}$

Implement this in two steps:

- accumulate: add up the ( $\alpha$  pre-multiplied) RGB $\alpha$  values at each pixel
- normalize: divide each pixel's accumulated RGB by its  $\alpha$  value

Q: what if  $\alpha = 0$ ?


## Image warping



Given a coordinate transform  $(x',y') = h(x,y)$  and a source image  $f(x,y)$ , how do we compute a transformed image  $g(x',y') = f(h(x,y))$ ?




## Forward warping



Send each pixel  $f(x,y)$  to its corresponding location  
 $(x',y') = h(x,y)$  in the second image

Q: what if pixel lands "between" two pixels?

## Forward warping




Send each pixel  $f(x,y)$  to its corresponding location  
 $(x',y') = h(x,y)$  in the second image

Q: what if pixel lands "between" two pixels?

A: distribute color among neighboring pixels  $(x',y')$   
– Known as "splatting"


## Inverse warping



Get each pixel  $g(x',y')$  from its corresponding location  
 $(x,y) = h^{-1}(x',y')$  in the first image

Q: what if pixel comes from "between" two pixels?

## Inverse warping



Get each pixel  $g(x',y')$  from its corresponding location  
 $(x,y) = h^{-1}(x',y')$  in the first image

Q: what if pixel comes from "between" two pixels?

A: *resample* color value

– We discussed resampling techniques before  
• nearest neighbor, bilinear, Gaussian, bicubic



## Forward vs. inverse warping

---

Q: which is better?

A: usually inverse—eliminates holes

- however, it requires an invertible warp function—not always possible...

## Other types of mosaics

---



Can mosaic onto *any* surface if you know the geometry

- See NASA's [Visible Earth project](#) for some stunning earth mosaics
  - <http://earthobservatory.nasa.gov/Newsroom/BlueMarble/>
  - Click for images...