

## Announcements

---

- Project 2
  - Out today
  - Sign up for a panorama kit ASAP!
    - best slots (weekend) go quickly...

## Mosaics part 2

---



VR Seattle: <http://www.vrseattle.com/>  
Full screen panoramas (cubic): <http://www.panoramas.dk/>  
Mars: [http://www.panoramas.dk/fullscreen3/2\\_mars97.html](http://www.panoramas.dk/fullscreen3/2_mars97.html)

### Today's Readings

- Szeliski and Shum paper (sections 1 and 2, skim the rest)
  - <http://www.cs.washington.edu/education/courses/455/00w/readings/szeliskishum97.pdf>

## Project 2

---


1. Take pictures on a tripod (or handheld)
2. Warp to spherical coordinates
3. Extract features
4. Align neighboring pairs using RANSAC
5. Write out list of neighboring translations
6. Correct for drift
7. Read in warped images and blend them
8. 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 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

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

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


### 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

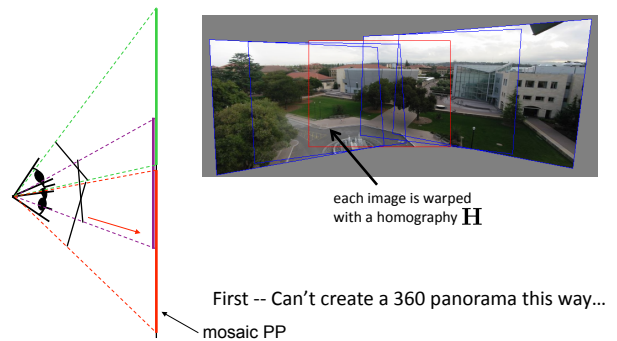
## Homography

A few examples on board

## Image warping with homographies




## Idea: projecting images onto a common plane




## Panoramas


What if you want a 360° field of view?




## Spherical projection



## Spherical reprojection



side view




top-down view


How to map sphere onto a flat image?

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

## Spherical reprojection



side view



top-down view

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



input

f = 200 (pixels)

f = 400

f = 800

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 image translations




Richard Szeliski

CSE 576 (Spring 2005): Computer Vision

19

## Random Sample Consensus




Richard Szeliski

CSE 576 (Spring 2005): Computer Vision

20

## Random Sample Consensus




Richard Szeliski

CSE 576 (Spring 2005): Computer  
Vision

21

## Least squares fit



Richard Szeliski

CSE 576 (Spring 2005): Computer  
Vision

22

## 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

## Computing transformations

Given a set of matches between images A and B

- How can we compute the transform  $T$  from A to B?



## Solving for translations

Using least squares

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

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

$2n \times 2$      $2 \times 1$      $2n \times 1$

## Affine transformations

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$



How many unknowns?

How many equations per match?

How many matches do we need?

## Affine transformations

Residuals:

$$r_{x_i}(a, b, c, d, e, f) = (ax_i + by_i + c) - x'_i$$

$$r_{y_i}(a, b, c, d, e, f) = (dx_i + ey_i + f) - y'_i$$

$$C(a, b, c, d, e, f) =$$

$$\sum_{i=1}^n (r_{x_i}(a, b, c, d, e, f)^2 + r_{y_i}(a, b, c, d, e, f)^2)$$

## Affine transformations


Matrix form

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_n & y_n & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_n & y_n & 1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} x'_1 \\ y'_1 \\ x'_2 \\ y'_2 \\ \vdots \\ x'_n \\ y'_n \end{bmatrix}$$

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

$2n \times 6$      $6 \times 1$      $2n \times 1$

## Homographies



To unwarped (rectify) an image

- solve for homography  $\mathbf{H}$  given  $\mathbf{p}$  and  $\mathbf{p}'$
- solve equations of the form:  $w\mathbf{p}' = \mathbf{H}\mathbf{p}$ 
  - linear in unknowns:  $w$  and coefficients of  $\mathbf{H}$
  - $\mathbf{H}$  is defined up to an arbitrary scale factor
  - how many points are necessary to solve for  $\mathbf{H}$ ?

## Solving for homographies

$$\begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} \cong \begin{bmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ h_{20} & h_{21} & h_{22} \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}$$

$$x'_i = \frac{h_{00}x_i + h_{01}y_i + h_{02}}{h_{20}x_i + h_{21}y_i + h_{22}}$$

$$y'_i = \frac{h_{10}x_i + h_{11}y_i + h_{12}}{h_{20}x_i + h_{21}y_i + h_{22}}$$

**Not linear!**

$$x'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{00}x_i + h_{01}y_i + h_{02}$$

$$y'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{10}x_i + h_{11}y_i + h_{12}$$

## Solving for homographies

$$x'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{00}x_i + h_{01}y_i + h_{02}$$

$$y'_i(h_{20}x_i + h_{21}y_i + h_{22}) = h_{10}x_i + h_{11}y_i + h_{12}$$

$$\begin{bmatrix} x_i & y_i & 1 & 0 & 0 & 0 & -x'_i x_i & -x'_i y_i & -x'_i \\ 0 & 0 & 0 & x_i & y_i & 1 & -y'_i x_i & -y'_i y_i & -y'_i \end{bmatrix} \begin{bmatrix} h_{00} \\ h_{01} \\ h_{02} \\ h_{10} \\ h_{11} \\ h_{12} \\ h_{20} \\ h_{21} \\ h_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

## Solving for homographies


$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x'_1 x_1 & -x'_1 y_1 & -x'_1 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -y'_1 x_1 & -y'_1 y_1 & -y'_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_n & y_n & 1 & 0 & 0 & 0 & -x'_n x_n & -x'_n y_n & -x'_n \\ 0 & 0 & 0 & x_n & y_n & 1 & -y'_n x_n & -y'_n y_n & -y'_n \end{bmatrix} \begin{bmatrix} h_{00} \\ h_{01} \\ h_{02} \\ h_{10} \\ h_{11} \\ h_{12} \\ h_{20} \\ h_{21} \\ h_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

$$\mathbf{A}_{2n \times 9} \mathbf{h}_9 = \mathbf{0}_{2n}$$

Defines a least squares problem: minimize  $\|\mathbf{A}\mathbf{h} - \mathbf{0}\|^2$


- Since  $\mathbf{h}$  is only defined up to scale, solve for unit vector  $\hat{\mathbf{h}}$
- Solution:  $\hat{\mathbf{h}}$  = eigenvector of  $\mathbf{A}^T \mathbf{A}$  with smallest eigenvalue
- Works with 4 or more points

## 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...

## 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



## Good window size




"Optimal" window: smooth but not ghosted

- Doesn't always work...

## Image feathering


What if you're blending more than two images?




## Image feathering

What if you have more than two images?

- Generate weight map for each image
  - typically want large weight at center, small weight at edge
- Each output pixel is a weighted average of inputs
  - be sure to divide by sum of weights at the end



## Alpha Blending



Optional: see Blinn (CGA, 1994) for details:  
<http://www.explore.ieee.org/iel1/38/7531/00310740.pdf?iNumber=7531&prodNo=1&number=310740&arSI=83&arSI=87&arAuthor=Blinn%2C+J.F.>

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:

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


Q: what if  $\alpha = 0$ ?



## More advanced blending schemes

A quick survey...

## Pyramid blending




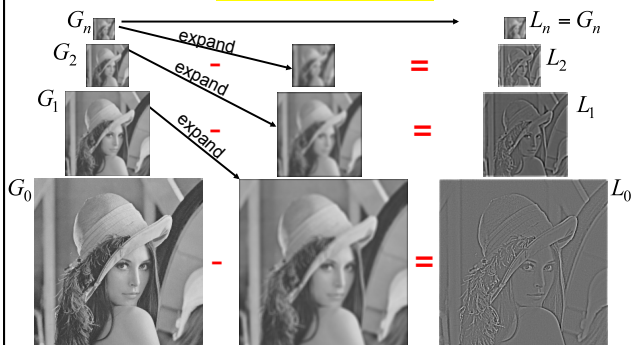
Create a Laplacian pyramid, blend each level

- Burt, P. J. and Adelson, E. H., [A multiresolution spline with applications to image mosaics](#), ACM Transactions on Graphics, 42(4), October 1983, 217-236.

## The Laplacian Pyramid

$$L_i = G_i - \text{expand}(G_{i+1})$$

Gaussian Pyramid  $G_i = L_i + \text{expand}(G_{i+1})$  Laplacian Pyramid



## Laplacian image blend

1. Compute Laplacian pyramid
2. Compute Gaussian pyramid on *weight* image (can put this in A channel)
3. Blend Laplacians using Gaussian blurred weights
4. Reconstruct the final image

Q: How do we compute the original weights?

A: For horizontal panorama, use *mid-lines*

Q: How about for a general "3D" panorama?

Richard Szeliski

Image Stitching

53

## Gradient-domain blending



Blend the gradients of the two images, then integrate  
For more info: Perez et al, SIGGRAPH 2003  
Also called "Poisson" blending

## De-Ghosting



## Local alignment (deghosting)

Use local optic flow to compensate for small motions  
[Shum & Szeliski, ICCV'98]



Figure 3: Deghosting a mosaic with motion parallax: (a) with parallax; (b) after single deghosting step (patch size 32); (c) multiple steps (sizes 32, 16 and 8).

## Photomontage [Agarwala et al., SIGGRAPH 2004]

- Each patch of the composite comes from a single image
- Solve for the seams that are hardest to detect (graph cuts)
- Blend across seams using gradient-domain blending



## Photomontage [Agarwala et al., SIGGRAPH 2004]



Figure 1 From a set of five source images of which four are shown on the left, we quickly create a composite family portrait in which everyone is smiling and looking at the camera (right). We simply flip through the stack and coarsely draw strokes using the designated source image objective over the people we wish to add to the composite. The user-applied strokes and computed regions are color-coded by the borders of the source images on the left (middle).

## Photomontage [Agarwala et al., SIGGRAPH 2004]



Figure 6 We use a set of portraits (first row) to mix and match facial features, to either improve a portrait, or create entirely new people. The faces are first hand-aligned, for example, to place all the noses in the same location. In the first two images in the second row, we replace the closed eyes of a portrait with the open eyes of another. The user paints strokes with the designated source objective to specify desired features. Next, we create a fictional person by combining three source portraits. Gradient-domain fusion is used to smooth out skin tone differences. Finally, we show two additional mixed portraits.


## Other types of mosaics



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

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

## Slit images



y-t slices of the video volume are known as *slit images*

- take a single column of pixels from each input image

Richard Szeliski

Image Stitching

61

## Slit images: cyclographs



## Slit images: photofinish

