

# Image Stitching

Computer Vision  
CSE P 576, Spring 2011  
Richard Szeliski  
Microsoft Research

# Panoramic Image Mosaics



Full screen panoramas (cubic): <http://www.panoramas.dk/>  
Mars: [http://www.panoramas.dk/fullscreen3/f2\\_mars97.html](http://www.panoramas.dk/fullscreen3/f2_mars97.html)  
2003 New Years Eve: <http://www.panoramas.dk/fullscreen3/f1.html>

# Gigapixel panoramas & images



# Image Mosaics



Goal: Stitch together several images into a seamless composite

## Today's lecture

---

### Image alignment and stitching

- motion models
- image warping
- point-based alignment
- complete mosaics (global alignment)
- compositing and blending
- ghost and parallax removal

## Readings

---

- Szeliski, CVAA:
  - Chapter 3.6: Image warping
  - Chapter 6.1: Feature-based alignment
  - Chapter 9.1: Motion models
  - Chapter 9.2: Global alignment
  - Chapter 9.3: Compositing
- Recognizing Panoramas, Brown & Lowe, ICCV'2003
- Szeliski & Shum, SIGGRAPH'97

---

## Motion models

## Motion models

---

What happens when we take two images with a camera and try to align them?

- translation?
- rotation?
- scale?
- affine?
- perspective?




... see interactive demo (VideoMosaic)

## Projective transformations

(aka *homographies*)

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} u \\ v \\ w \end{bmatrix} \quad \begin{aligned} x' &= u/w \\ y' &= v/w \end{aligned}$$

“keystone” distortions



## Image Warping

## Image Warping

image filtering: change *range* of image

$$g(x) = h(f(x))$$





image warping: change *domain* of image

$$g(x) = f(h(x))$$



Richard Szeliski

Image Stitching

11

## Image Warping

image filtering: change *range* of image

$$g(x) = h(f(x))$$





image warping: change *domain* of image

$$g(x) = f(h(x))$$




Richard Szeliski

Image Stitching

12

## Parametric (global) warping

Examples of parametric warps:



Richard Szeliski

Image Stitching

13

## 2D coordinate transformations

translation:  $\mathbf{x}' = \mathbf{x} + \mathbf{t}$        $\mathbf{x} = (x, y)$   
 rotation:  $\mathbf{x}' = \mathbf{R} \mathbf{x} + \mathbf{t}$   
 similarity:  $\mathbf{x}' = s \mathbf{R} \mathbf{x} + \mathbf{t}$   
 affine:  $\mathbf{x}' = \mathbf{A} \mathbf{x} + \mathbf{t}$   
 perspective:  $\underline{\mathbf{x}}' \cong \mathbf{H} \underline{\mathbf{x}}$        $\underline{\mathbf{x}} = (x, y, 1)$   
 ( $\underline{\mathbf{x}}$  is a *homogeneous* coordinate)

These all form a nested *group* (closed w/ inv.)


Richard Szeliski

Image Stitching

14

## Image Warping

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



Richard Szeliski


Image Stitching

15

## Forward Warping

Send each pixel  $\mathbf{f}(\mathbf{x})$  to its corresponding location  $\mathbf{x}' = \mathbf{h}(\mathbf{x})$  in  $\mathbf{g}(\mathbf{x}')$

- What if pixel lands “between” two pixels?



Richard Szeliski


Image Stitching

16

## Forward Warping

Send each pixel  $f(x)$  to its corresponding location  $x' = h(x)$  in  $g(x')$

- What if pixel lands “between” two pixels?
- Answer: add “contribution” to several pixels, normalize later (*splatting*)



Richard Szeliski


Image Stitching

17

## 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?



Richard Szeliski


Image Stitching

18

## 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



Richard Szeliski

Image Stitching

19

## Interpolation

Possible interpolation filters:

- nearest neighbor
- bilinear
- bicubic (interpolating)
- sinc / FIR

Needed to prevent “jaggies” and “texture crawl” (see [demo](#))



Richard Szeliski

Image Stitching

20

## Motion models (reprise)

## Motion models







Translation

Affine

Perspective

3D rotation



2 unknowns

6 unknowns

8 unknowns

3 unknowns

Richard Szeliski

Image Stitching

24

## Finding the transformation

Translation	=	2 degrees of freedom
Similarity	=	4 degrees of freedom
Affine	=	6 degrees of freedom
Homography	=	8 degrees of freedom

How many corresponding points do we need to solve?


CSE 576, Spring 2008

Structure from Motion

25

## Plane perspective mosaics

- 8-parameter generalization of affine motion
  - works for pure rotation or planar surfaces
- Limitations:
  - local minima
  - slow convergence
  - difficult to control interactively




Richard Szeliski

Image Stitching

26

## Rotational mosaics

- Directly optimize rotation and focal length
- Advantages:
  - ability to build full-view panoramas
  - easier to control interactively
  - more stable and accurate estimates



Richard Szeliski

Image Stitching

28

## 3D → 2D Perspective Projection

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} = [\mathbf{R}]_{3 \times 3} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + \mathbf{t}$$

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \begin{bmatrix} U \\ V \\ W \end{bmatrix} = \begin{bmatrix} f & 0 & u_c \\ 0 & f & v_c \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix}$$

Richard Szeliski

Image Stitching

29

## Rotational mosaic

Projection equations

1. Project from image to 3D ray
 
$$(x_0, y_0, z_0) = (u_0 - u_c, v_0 - v_c, f)$$
2. Rotate the ray by camera motion
 
$$(x_1, y_1, z_1) = \mathbf{R}_{01} (x_0, y_0, z_0)$$
3. Project back into new (source) image
 
$$(u_1, v_1) = (fx_1/z_1 + u_c, fy_1/z_1 + v_c)$$

Richard Szeliski

Image Stitching

30

## Image Mosaics (Stitching)

[Szeliski & Shum, SIGGRAPH'97]

[Szeliski, Fnt CVCG, 2006]

## Image Mosaics (stitching)

Blend together several overlapping images into one seamless *mosaic* (composite)




Richard Szeliski

Image Stitching

36

## Mosaics for Video Coding

Convert masked images into a background sprite for content-based coding



Richard Szeliski

Image Stitching

37

## Establishing correspondences


1. Direct method:
  - Use generalization of affine motion model [Szeliski & Shum '97]
2. Feature-based method
  - Extract features, match, find consistent *inliers* [Lowe ICCV'99; Schmid ICCV'98, Brown&Lowe ICCV'2003]
  - Compute  $R$  from correspondences (absolute orientation)

Richard Szeliski

Image Stitching

38

## Stitching demo



Richard Szeliski


Image Stitching

40



## Panoramas

What if you want a 360° field of view?




Richard Szeliski

Image Stitching

41

## Cylindrical panoramas



Steps

- Reproject each image onto a cylinder
- Blend
- Output the resulting mosaic

Richard Szeliski

Image Stitching

42

## Cylindrical Panoramas

Map image to cylindrical or spherical coordinates

- need *known* focal length



Image 384x300

f = 180 (pixels)

f = 280

f = 380


Richard Szeliski

Image Stitching

43

## Determining the focal length

1. Initialize from homography  $H$  (see text or [SzSh'97])
2. Use camera's EXIF tags (approx.)
3. Use a tape measure



4. Ask your instructor

Richard Szeliski

Image Stitching

44

### 3D → 2D Perspective Projection

$$\begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} = [R]_{3 \times 3} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + t$$


$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \begin{bmatrix} U \\ V \\ W \end{bmatrix} = \begin{bmatrix} f & 0 & u_c \\ 0 & f & v_c \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix}$$

### Cylindrical projection

- Map 3D point  $(X, Y, Z)$  onto cylinder  
 $(\hat{x}, \hat{y}, \hat{z}) = \frac{1}{\sqrt{X^2 + Y^2}}(X, Y, Z)$
- Convert to cylindrical coordinates  
 $(\sin\theta, h, \cos\theta) = (\hat{x}, \hat{y}, \hat{z})$
- Convert to cylindrical image coordinates  
 $(\tilde{x}, \tilde{y}) = (s\theta, sh) + (\tilde{x}_c, \tilde{y}_c)$   
 -  $s$  defines size of the final image

### Cylindrical warping


Given focal length  $f$  and image center  $(x_c, y_c)$



$$\begin{aligned} \theta &= (x_{cyl} - x_c) / f \\ h &= (y_{cyl} - y_c) / f \\ \hat{x} &= \sin \theta \\ \hat{y} &= h \\ \hat{z} &= \cos \theta \\ x &= f\hat{x} / \hat{z} + x_c \\ y &= f\hat{y} / \hat{z} + y_c \end{aligned}$$

### Spherical warping


Given focal length  $f$  and image center  $(x_c, y_c)$



$$\begin{aligned} \theta &= (x_{cyl} - x_c) / f \\ \varphi &= (y_{cyl} - y_c) / f \\ \hat{x} &= \sin \theta \cos \varphi \\ \hat{y} &= \sin \theta \\ \hat{z} &= \cos \theta \cos \varphi \\ x &= f\hat{x} / \hat{z} + x_c \\ y &= f\hat{y} / \hat{z} + y_c \end{aligned}$$

### 3D rotation

Rotate image before placing on unrolled sphere



$$\begin{aligned} \theta &= (x_{cyl} - x_c) / f \\ \phi &= (y_{cyl} - y_c) / f \\ \hat{x} &= \sin \theta \cos \phi \\ \hat{y} &= \sin \phi \\ \hat{z} &= \cos \theta \cos \phi \\ x &= f \hat{x} / \hat{z} + x_c \\ y &= f \hat{y} / \hat{z} + y_c \end{aligned}$$

### Radial distortion

Correct for "bending" in wide field of view lenses



**Project  $(\hat{x}, \hat{y}, \hat{z})$  to "normalized" image coordinates**

$$\begin{aligned} x'_n &= \hat{x} / \hat{z} \\ y'_n &= \hat{y} / \hat{z} \end{aligned}$$

**Apply radial distortion**

$$\begin{aligned} r^2 &= x'^2_n + y'^2_n \\ x'_d &= x'_n (1 + \kappa_1 r^2 + \kappa_2 r^4) \\ y'_d &= y'_n (1 + \kappa_1 r^2 + \kappa_2 r^4) \end{aligned}$$

**Apply focal length translate image center**

$$\begin{aligned} x' &= f x'_d + x_c \\ y' &= f y'_d + y_c \end{aligned}$$

To model lens distortion  
 • Use above projection operation instead of standard projection matrix multiplication

### Fisheye lens

Extreme "bending" in ultra-wide fields of view




$$\begin{aligned} \hat{r}^2 &= \hat{x}^2 + \hat{y}^2 \\ (\cos \theta \sin \phi, \sin \theta \sin \phi, \cos \phi) &= s(x, y, z) \end{aligned}$$


equations become

$$\begin{aligned} x' &= s \phi \cos \theta = s \frac{x}{r} \tan^{-1} \frac{r}{z} \\ y' &= s \phi \sin \theta = s \frac{y}{r} \tan^{-1} \frac{r}{z} \end{aligned}$$

### Matching features



## Random Sample Consensus




Richard Szeliski

Image Stitching

53

## Random Sample Consensus




Richard Szeliski

Image Stitching

54

## Least squares fit



Richard Szeliski

Image Stitching

55

## RANSAC for estimating homography

RANSAC loop:

1. Select four feature pairs (at random)
2. Compute homography  $\mathbf{H}$  (exact)
3. Compute inliers where  $\|p_i', \mathbf{H} p_i\| < \epsilon$

Keep largest set of inliers

Re-compute least-squares  $\mathbf{H}$  estimate using all of the inliers


CSE 576, Spring 2008

Structure from Motion

56


### Simple example: fit a line

Rather than homography H (8 numbers)  
fit  $y=ax+b$  (2 numbers a, b) to 2D pairs




### Simple example: fit a line

Pick 2 points  
Fit line  
Count inliers




### Simple example: fit a line

Pick 2 points  
Fit line  
Count inliers




### Simple example: fit a line

Pick 2 points  
Fit line  
Count inliers




## Simple example: fit a line

Pick 2 points  
Fit line  
Count inliers




## Simple example: fit a line

Use biggest set of inliers  
Do least-square fit



## RANSAC



red:  
rejected by 2nd nearest  
neighbor criterion  
blue:  
Ransac outliers  
yellow:  
inliers



## Image Stitching—review

1. Align the images over each other
  - camera pan  $\leftrightarrow$  translation on cylinder
2. Blend the images together ([demo](#))



Richard Szelski

Image Stitching

66

## Assignment 2 – Creating Panoramas


1. Implement Harris corner detector
2. Implement MatchInterestPoints
3. Compute homography using RANSAC
4. Compute size of stitched images from projected corners
5. Inverse sample image and average

Richard Szeliski

Image Stitching

67

## Resampling the final image



CSE 576, Spring 2008

Structure from Motion

68

## Full-view (360° spherical) panoramas



## Full-view Panorama



## Texture Mapped Model



## Global alignment

- Register *all* pairwise overlapping images
- Use a 3D rotation model (one R per image)
- Use direct alignment (patch centers) or feature based
- *Infer* overlaps based on previous matches (incremental)
- Optionally *discover* which images overlap other images using feature selection (RANSAC)

Richard Szeliski


Image Stitching

75

## Recognizing Panoramas

Matthew Brown & David Lowe  
ICCV'2003

## Recognizing Panoramas




[Brown & Lowe, ICCV'03]

Richar

77

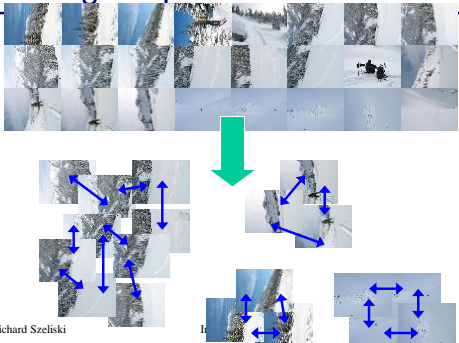


### Finding the panoramas



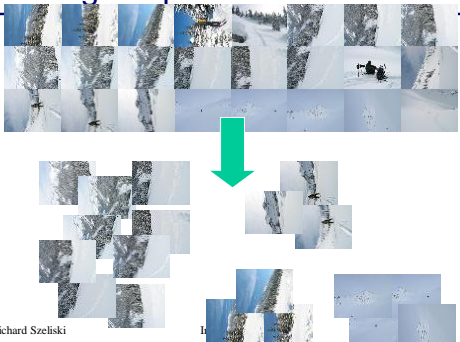
Richard Szeliski Image Stitching 78

### Finding the panoramas



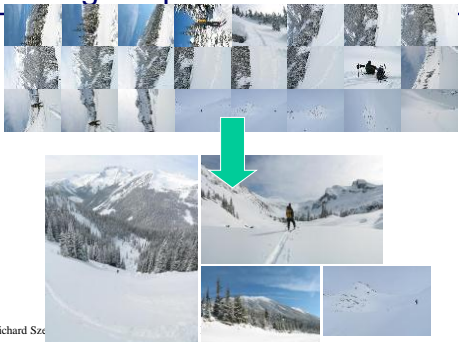
Richard Szeliski Image Stitching 79

### Finding the panoramas



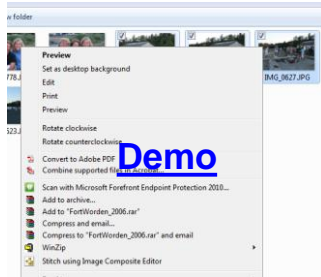
Richard Szeliski Image Stitching 80

### Finding the panoramas



Richard Szeliski Image Stitching 81

## Fully automated 2D stitching

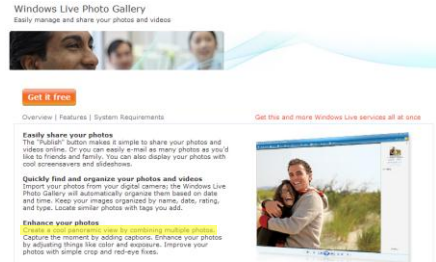


Richard Szeliski

Image Stitching

82

## Get you own copy



Richard Szeliski

Image Stitching

83

## Rec.pano.: system components

1. Feature detection and description
  - more uniform point density
2. Fast matching (hash table)
3. RANSAC filtering of matches
4. Intensity-based verification
5. Incremental bundle adjustment

[M. Brown, R. Szeliski, and S. Winder. Multi-image matching using multi-scale oriented patches, CVPR'2005]

Richard Szeliski

Image Stitching

84

## Multi-Scale Oriented Patches

### Interest points

- Multi-scale Harris corners
- Orientation from blurred gradient
- Geometrically invariant to similarity transforms

### Descriptor vector

- Bias/gain normalized sampling of local patch (8x8)
- Photometrically invariant to affine changes in intensity


Richard Szeliski

Image Stitching

85

## Feature irregularities

Distribute points evenly over the image




Richard Szeliski


Image Stitching

86


## Probabilistic Feature Matching

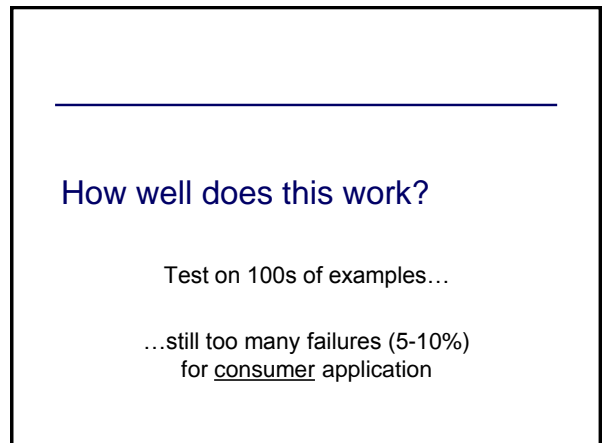
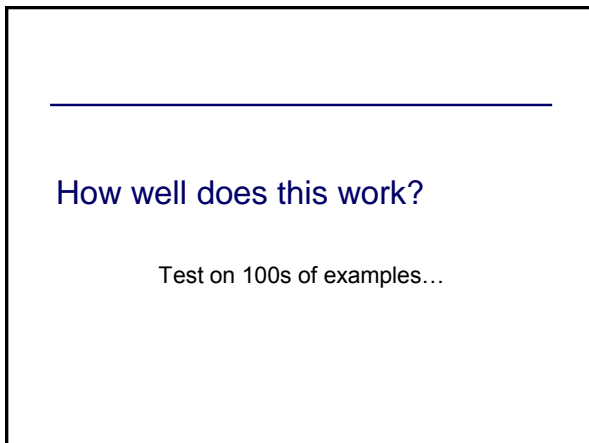
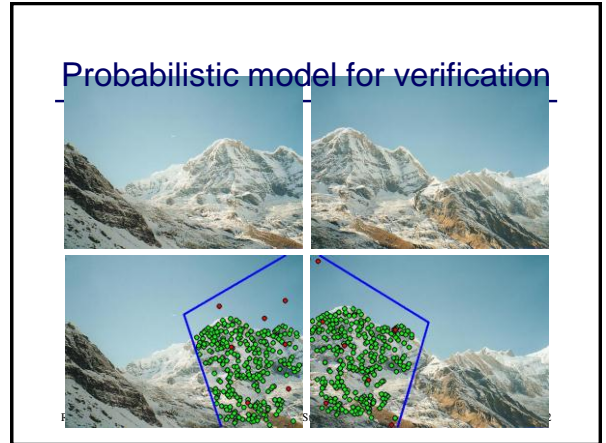



## RANSAC motion model



## RANSAC motion model





## Matching Mistakes: False Positive



## Matching Mistakes: False Positive



Richard Szeliski

Image Stitching

96

## Matching Mistake: False Negative

Moving objects: large areas of disagreement



Richard Szeliski

Image Stitching

97

## Matching Mistakes

Accidental alignment

- repeated / similar regions

Failed alignments

- moving objects / parallax
- low overlap
- "feature-less" regions (more variety?)

No 100% reliable algorithm?



Richard Szeliski

Image Stitching

98

## How can we fix these?

---

- Tune the feature detector
  - Tune the feature matcher (cost metric)
  - Tune the RANSAC stage (motion model)
  - Tune the verification stage
  - Use “higher-level” knowledge
    - e.g., typical camera motions
- Sounds like a big “learning” problem
- Need a large training/test data set (panoramas)

Richard Szeliski

Image Stitching

99

## Image Blending

---

## Image feathering

---

Weight each image proportional to its distance from the edge  
(distance map [Danielsson, CVGIP 1980])

1. Generate *weight map* for each image
2. Sum up all of the weights and divide by sum:  
weights sum up to 1:  $w_i' = w_i / (\sum_i w_i)$




Richard Szeliski

Image Stitching

101

## Image Feathering

---




Richard Szeliski

Image Stitching

102

## Feathering






Richard Szeliski

Image Stitching

103

## Effect of window size




Richard Szeliski

Image Stitching

104

## Effect of window size



Richard Szeliski

Image Stitching

105

## Good window size




**“Optimal” window: smooth but not ghosted**  
 • Doesn't always work...

Richard Szeliski


Image Stitching

106

## Pyramid Blending



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.



## 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

109

## Weight selection (3D panorama)

Idea: use original feather weights to select *strongest* contributing image



Can be implemented using L- $\infty$  norm: ( $p = 10$ )

$$w_i' = [w_i^p / (\sum_i w_i^p)]^{1/p}$$

Richard Szeliski

Image Stitching

110



## Poisson Image Editing



Blend the gradients of the two images, then integrate

For more info: Perez et al, SIGGRAPH 2003

- see [Computational Photography lecture](#)

Richard Szeliski

Image Stitching

111

## 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).

Richard Szeliski

Image Stitching

113

## Local alignment (deghosting)

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



Figure 4: Deghosting a mosaic with optical distortion: (a) with distortion; (b) after multiple steps.

Richard Szeliski

Image Stitching

114

## Region-based de-ghosting

Select only one image in *regions-of-difference* using weighted vertex cover  
[Uyttendaele *et al.*, CVPR'01]



Figure 5-- (A) Ghosted mosaic. (B) Result of de-ghosting algorithm.

Richard Szeliski

Image Stitching

115

## Region-based de-ghosting

Select only one image in *regions-of-difference* using weighted vertex cover  
[Uyttendaele *et al.*, CVPR'01]



Figure 6-- (A) Ghosted mosaic. (B) Result of de-ghosting algorithm.


Richard Szeliski

Image Stitching

116

## Cutout-based de-ghosting

- Select only one image per output pixel, using spatial continuity
- Blend across seams using gradient continuity ("Poisson blending")



[Agarwala *et al.*, SG'2004]

Richard Szeliski

Image Stitching

117

## Cutout-based compositing

Photomontage [Agarwala *et al.*, SG'2004]

- Interactively blend *different* images: group portraits



Figure 8 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 cursorily 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).

Richard Szeliski

Image Stitching

118

## PhotoMontage

Technical details:

- use Graph Cuts to optimize seam placement

Demo:

- Windows Live Photo Gallery
- Photo Fuse



Richard Szeliski

In

19

## Cutout-based compositing

Photomontage [Agarwala *et al.*, SG'2004]

- Interactively blend *different* images: focus settings




Figure 2. A set of source photographs of an old friend of ours is used shown on the left taken at different focal lengths. We use a global minimum cost seam algorithm to compute the optimal composite automatically (top row), with us back to focus detail, and the focusing plane directly below it, a small number of containing pixels to appear after pixel-wise fusion (top middle). The composites we show correspond to the PhotoMontage (top right), by blending the source images, middle, and to Laplacian pyramids (bottom right). All of these other approaches have artifacts. PhotoMontage's method makes seamless seams. PhotoMontage fails to actually seam faces to the body, and Laplacian pyramids create halos around center of the faces.

Richard Szeliski

Image Stitching

120

## Cutout-based compositing

Photomontage [Agarwala *et al.*, SG'2004]

- Interactively blend *different* images: people's faces



Figure 6. We use a set of patches (left row) to find and track facial features, to either improve a portrait, or create entirely new people. The faces are first face-aligned for example by using all the corners in the same location. For the first row images in the second row, we replace the closed eyes of a person with the open eyes of another. The new patches (left column) are also aligned to the eyes. The second row shows the result of replacing closed eyes with open eyes. The third row shows the result of replacing closed eyes with open eyes. The fourth row shows the result of replacing closed eyes with open eyes. The fifth row shows the result of replacing closed eyes with open eyes. The sixth row shows the result of replacing closed eyes with open eyes. The seventh row shows the result of replacing closed eyes with open eyes. The eighth row shows the result of replacing closed eyes with open eyes. The ninth row shows the result of replacing closed eyes with open eyes. The tenth row shows the result of replacing closed eyes with open eyes. The eleventh row shows the result of replacing closed eyes with open eyes. The twelfth row shows the result of replacing closed eyes with open eyes. The thirteenth row shows the result of replacing closed eyes with open eyes. The fourteenth row shows the result of replacing closed eyes with open eyes. The fifteenth row shows the result of replacing closed eyes with open eyes. The sixteenth row shows the result of replacing closed eyes with open eyes. The seventeenth row shows the result of replacing closed eyes with open eyes. The eighteenth row shows the result of replacing closed eyes with open eyes. The nineteenth row shows the result of replacing closed eyes with open eyes. The twentieth row shows the result of replacing closed eyes with open eyes. The twenty-first row shows the result of replacing closed eyes with open eyes. The twenty-second row shows the result of replacing closed eyes with open eyes. The twenty-third row shows the result of replacing closed eyes with open eyes. The twenty-fourth row shows the result of replacing closed eyes with open eyes. The twenty-fifth row shows the result of replacing closed eyes with open eyes. The twenty-sixth row shows the result of replacing closed eyes with open eyes. The twenty-seventh row shows the result of replacing closed eyes with open eyes. The twenty-eighth row shows the result of replacing closed eyes with open eyes. The twenty-ninth row shows the result of replacing closed eyes with open eyes. The thirtieth row shows the result of replacing closed eyes with open eyes. The thirty-first row shows the result of replacing closed eyes with open eyes. The thirty-second row shows the result of replacing closed eyes with open eyes. The thirty-third row shows the result of replacing closed eyes with open eyes. The thirty-fourth row shows the result of replacing closed eyes with open eyes. The thirty-fifth row shows the result of replacing closed eyes with open eyes. The thirty-sixth row shows the result of replacing closed eyes with open eyes. The thirty-seventh row shows the result of replacing closed eyes with open eyes. The thirty-eighth row shows the result of replacing closed eyes with open eyes. The thirty-ninth row shows the result of replacing closed eyes with open eyes. The fortieth row shows the result of replacing closed eyes with open eyes. The forty-first row shows the result of replacing closed eyes with open eyes. The forty-second row shows the result of replacing closed eyes with open eyes. The forty-third row shows the result of replacing closed eyes with open eyes. The forty-fourth row shows the result of replacing closed eyes with open eyes. The forty-fifth row shows the result of replacing closed eyes with open eyes. The forty-sixth row shows the result of replacing closed eyes with open eyes. The forty-seventh row shows the result of replacing closed eyes with open eyes. The forty-eighth row shows the result of replacing closed eyes with open eyes. The forty-ninth row shows the result of replacing closed eyes with open eyes. The fiftieth row shows the result of replacing closed eyes with open eyes. The fifty-first row shows the result of replacing closed eyes with open eyes. The fifty-second row shows the result of replacing closed eyes with open eyes. The fifty-third row shows the result of replacing closed eyes with open eyes. The fifty-fourth row shows the result of replacing closed eyes with open eyes. The fifty-fifth row shows the result of replacing closed eyes with open eyes. The fifty-sixth row shows the result of replacing closed eyes with open eyes. The fifty-seventh row shows the result of replacing closed eyes with open eyes. The fifty-eighth row shows the result of replacing closed eyes with open eyes. The fifty-ninth row shows the result of replacing closed eyes with open eyes. The sixtieth row shows the result of replacing closed eyes with open eyes. The sixty-first row shows the result of replacing closed eyes with open eyes. The sixty-second row shows the result of replacing closed eyes with open eyes. The sixty-third row shows the result of replacing closed eyes with open eyes. The sixty-fourth row shows the result of replacing closed eyes with open eyes. The sixty-fifth row shows the result of replacing closed eyes with open eyes. The sixty-sixth row shows the result of replacing closed eyes with open eyes. The sixty-seventh row shows the result of replacing closed eyes with open eyes. The sixty-eighth row shows the result of replacing closed eyes with open eyes. The sixty-ninth row shows the result of replacing closed eyes with open eyes. The seventieth row shows the result of replacing closed eyes with open eyes. The seventy-first row shows the result of replacing closed eyes with open eyes. The seventy-second row shows the result of replacing closed eyes with open eyes. The seventy-third row shows the result of replacing closed eyes with open eyes. The seventy-fourth row shows the result of replacing closed eyes with open eyes. The seventy-fifth row shows the result of replacing closed eyes with open eyes. The seventy-sixth row shows the result of replacing closed eyes with open eyes. The seventy-seventh row shows the result of replacing closed eyes with open eyes. The seventy-eighth row shows the result of replacing closed eyes with open eyes. The seventy-ninth row shows the result of replacing closed eyes with open eyes. The eightieth row shows the result of replacing closed eyes with open eyes. The eighty-first row shows the result of replacing closed eyes with open eyes. The eighty-second row shows the result of replacing closed eyes with open eyes. The eighty-third row shows the result of replacing closed eyes with open eyes. The eighty-fourth row shows the result of replacing closed eyes with open eyes. The eighty-fifth row shows the result of replacing closed eyes with open eyes. The eighty-sixth row shows the result of replacing closed eyes with open eyes. The eighty-seventh row shows the result of replacing closed eyes with open eyes. The eighty-eighth row shows the result of replacing closed eyes with open eyes. The eighty-ninth row shows the result of replacing closed eyes with open eyes. The ninetieth row shows the result of replacing closed eyes with open eyes. The ninety-first row shows the result of replacing closed eyes with open eyes. The ninety-second row shows the result of replacing closed eyes with open eyes. The ninety-third row shows the result of replacing closed eyes with open eyes. The ninety-fourth row shows the result of replacing closed eyes with open eyes. The ninety-fifth row shows the result of replacing closed eyes with open eyes. The ninety-sixth row shows the result of replacing closed eyes with open eyes. The ninety-seventh row shows the result of replacing closed eyes with open eyes. The ninety-eighth row shows the result of replacing closed eyes with open eyes. The ninety-ninth row shows the result of replacing closed eyes with open eyes. The one hundred row shows the result of replacing closed eyes with open eyes.

Richard Szeliski

Image Stitching

121

## More stitching possibilities

- Video stitching
- High dynamic range image stitching
  - see demo...
- Flash + Non-Flash
- Video-based rendering

Next-next week's lecture:  
Computational Photography

Richard Szeliski

Image Stitching

122

## 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/>

Richard Szeliski

Image Stitching

123

## 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

124

## Slit images: cyclographs




Richard Szeliski

Image Stitching

125

## Slit images: photofinish



Richard Szeliski

Image Stitching

126

Final thought:  
What is a "panorama"?

Tracking a subject




Repeated (best) shots



Multiple exposures



"Infer" what photographer wants?



Richard Szeliski

Image Stitching