

## Image Mosaics



Goal: Stitch together several images into a seamless composite
Richard Szeliski

## 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

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

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

- perspective?
... see interactive demo (VideoMosaic)



## Image Warping

image filtering: change range of image

image warping: change domain of image


## Image Warping

image filtering: change range of image

image warping: change domain of image


## Parametric (global) warping

Examples of parametric warps:

perspective
Image Stitching

cylindrical

## 2D coordinate transformations

translation: $\quad \boldsymbol{x}^{\prime}=\boldsymbol{x}+\boldsymbol{t} \quad \boldsymbol{x}=(x, y)$
rotation: $\quad \boldsymbol{x}^{\prime}=\boldsymbol{R} \boldsymbol{x}+\boldsymbol{t}$
similarity: $\quad x^{\prime}=s R x+t$
affine: $\quad \boldsymbol{x}^{\prime}=\boldsymbol{A} \boldsymbol{x}+\boldsymbol{t}$
perspective: $\quad \underline{\boldsymbol{x}}^{\prime} \cong \boldsymbol{H} \underline{\boldsymbol{x}} \quad \underline{\boldsymbol{x}}=(x, y, 1)$ ( $\underline{x}$ is a homogeneous coordinate)
These all form a nested group (closed w/ inv.)

## Image Warping

Given a coordinate transform $\boldsymbol{x}^{\prime}=\boldsymbol{h}(\boldsymbol{x})$ and a source image $\boldsymbol{f}(\boldsymbol{x})$, how do we compute a transformed image $\boldsymbol{g}\left(\boldsymbol{x}^{\prime}\right)=\boldsymbol{f}(\boldsymbol{h}(\boldsymbol{x}))$ ?


## Forward Warping

Send each pixel $\boldsymbol{f}(\boldsymbol{x})$ to its corresponding location $\boldsymbol{x}^{\prime}=\boldsymbol{h}(\boldsymbol{x})$ in $\boldsymbol{g}\left(\boldsymbol{x}^{\prime}\right)$

- What if pixel lands "between" two pixels?



## Forward Warping

Send each pixel $\boldsymbol{f}(\boldsymbol{x})$ to its corresponding location $\boldsymbol{x}^{\prime}=\boldsymbol{h}(\boldsymbol{x})$ in $\boldsymbol{g}\left(\boldsymbol{x}^{\prime}\right)$

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



## Inverse Warping

Get each pixel $\boldsymbol{g}\left(\boldsymbol{x}^{\prime}\right)$ from its corresponding location $\boldsymbol{x}^{\prime}=\boldsymbol{h}(\boldsymbol{x})$ in $\boldsymbol{f}(\boldsymbol{x})$

- What if pixel comes from "between" two pixels?
- Answer: resample color value from interpolated (prefiltered) source image




## Finding the transformation

Plane perspective mosaics

| 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?

## Rotational mosaics

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



## 3D $\rightarrow$ 2D Perspective Projection


$\left[\begin{array}{l}u \\ v \\ 1\end{array}\right] \sim\left[\begin{array}{c}U \\ V \\ W\end{array}\right]=\left[\begin{array}{llc}f & 0 & u_{c} \\ 0 & f & v_{c} \\ 0 & 0 & 1\end{array}\right]\left[\begin{array}{c}X_{c} \\ Y_{c} \\ Z_{c}\end{array}\right]$

## Rotational mosaic

Projection equations

1. Project from image to 3D ray
$\left(x_{0}, y_{0}, z_{0}\right)=\left(u_{0}-u_{c}, v_{0}-v_{c} f\right)$
2. Rotate the ray by camera motion
$\left(x_{1}, y_{1}, z_{1}\right) \quad=\boldsymbol{R}_{01}\left(x_{0}, y_{0}, z_{0}\right)$
3. Project back into new (source) image
$\left(u_{1}, v_{1}\right)=\left(f x_{1} / z_{1}+u_{c} f y_{1} / z_{1}+v_{c}\right)$

Image Mosaics (Stitching)
[Szeliski \& Shum, SIGGRAPH'97] [Szeliski, FnT CVCG, 2006]


## Mosaics for Video Coding

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


Richard Szeliski
Image Stitching
37

## Establishing correspondences

## Stitching demo

1. Direct method:

- Use generalization of affine motion model [Szeliski \& Shum '97]

2. Feature-based method

- Extract features, match, find consisten inliers [Lowe ICCV'99; Schmid ICCV'98, Brown\&Lowe ICCV'2003]
- Compute $\boldsymbol{R}$ from correspondences (absolute orientation)


## Panoramas

What if you want a $360^{\circ}$ field of view?


## Cylindrical panoramas



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



## Determining the focal length

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

4. Ask your instructor

## 3D $\rightarrow$ 2D Perspective Projection

$$
\left.\begin{array}{c}
{\left[\begin{array}{c}
X_{c} \\
Y_{c} \\
Z_{c}
\end{array}\right]=[\mathbf{R}]_{3 \times 3}\left[\begin{array}{c}
X \\
Y \\
Z
\end{array}\right]+\mathbf{t}} \\
u \\
v \\
1
\end{array}\right] \sim\left[\begin{array}{c}
U \\
V \\
W
\end{array}\right]=\left[\begin{array}{ccc}
f & 0 & u_{c} \\
0 & f & v_{c} \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{c}
X_{c} \\
Y_{c} \\
Z_{c}
\end{array}\right] .
$$

Cylindrical projection


## Cylindrical warping

Given focal length $f$ and image center ( $x_{c}, y_{c}$ )


$$
\begin{aligned}
\theta & =\left(x_{c y l}-x_{c}\right) / f \\
h & =\left(y_{c y l}-y_{c}\right) / f \\
\widehat{x} & =\sin \theta \\
\widehat{y} & =h \\
\widehat{z} & =\cos \theta \\
x & =f \hat{x} / \hat{z}+x_{c} \\
y & =f \hat{y} / \hat{z}+y_{c}
\end{aligned}
$$

## Spherical warping

Given focal length $f$ and image center $\left(x_{c}, y_{c}\right)$

$$
\begin{aligned}
\theta & =\left(x_{c y l}-x_{c}\right) / f \\
\varphi & =\left(y_{c y l}-y_{c}\right) / f \\
\widehat{x} & =\sin \theta \cos \varphi \\
\widehat{y} & =\sin \varphi \\
\widehat{z} & =\cos \theta \cos \varphi \\
x & =f \hat{x} / \hat{z}+x_{c} \\
y & =f \hat{y} / \widehat{z}+y_{c}
\end{aligned}
$$

## 3D rotation

Rotate image before placing on unrolled sphere
$\theta=\left(x_{c y l}-x_{c}\right) / f$


Radial distortion
Correct for "bending" in wide field of view lenses


To model lens distortion

- Use above projection operation instead of standard projection matrix multiplication

Richard Szeliski

Fisheye lens
Extreme "bending" in ultra-wide fields of view

$\widehat{r}^{2}=\widehat{x}^{2}+\widehat{y}^{2}$
$(\cos \theta \sin \phi, \sin \theta \sin \phi, \cos \phi)=s(x, y, z)$
uations become

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

Matching features


Richard Szeliski


RANSAC for estimating homography
RANSAC loop:

1. Select four feature pairs (at random)
2. Compute homography $\boldsymbol{H}$ (exact)
3. Compute inliers where $\left\|p_{i}^{\prime}, \boldsymbol{H}_{p_{i}}\right\|<\varepsilon$

Keep largest set of inliers
Re-compute least-squares $\boldsymbol{H}$ estimate using all of the inliers


## Simple example: fit a line

Pick 2 points
Fit line
Count inliers


Simple example: fit a line
Simple example: fit a line
Pick 2 points
Fit line
Count inliers


Pick 2 points
Fit line
Count inliers



Simple example: fit a line
Use biggest set of inliers
Do least-square fit



## 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



## 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)


Finding the panoramas


Finding the panoramas




## Rec.pano.: system components

## Multi-Scale Oriented Patches

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]

## Feature irregularities

Distribute points evenly over the image




How well does this work?
How well does this work?

Test on 100s of examples...
...still too many failures (5-10\%)
for consumer application


## 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
$\rightarrow$ Sounds like a big "learning" problem
- Need a large training/test data set (panoramas)


## Image feathering

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

2. Sum up all of the weights and divide by sum: weights sum up to 1: $\quad w_{i}{ }^{\prime}=w_{i} /\left(\sum_{i} w_{i}\right)$




Effect of window size


## Good window size


"Optimal" window: smooth but not ghosted - Doesn't always work...


## 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?

## Weight selection (3D panorama)

Idea: use original feather weights to select strongest contributing image


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

$$
w_{i}^{\prime}=\left[w_{i}^{p} /\left(\sum_{i} w_{i}^{p}\right)\right]^{1 / p}
$$

Richard Szeliski
Image Stitching


## 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

## 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.


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


## 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]


## Cutout-based compositing

Photomontage [Agarwala et al., SG'2004]

- Interactively blend different images: group portraits



## PhotoMontage

Technical details:

- use Graph Cuts to optimize seam placement

Demo:

- Windows Live Photo Gallery Photo Fuse



## Cutout-based compositing

Photomontage [Agarwala et al., SG'2004]

- Interactively blend different images: focus settings



## Cutout-based compositing

Photomontage [Agarwala et al., SG'2004]

- Interactively blend different images:
people's faces



## 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

## 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/

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


Slit images: photofinish


Richard Szeliski
Image Stitching


