

# Administrivia

- And now the winners of homework #3 hybrid images
  - as decided by the graders ...



# Winner (#3): John Williams

Bernie Sanders + Hillary Clinton





**Bernie Clinton** 



# Winner (#2): Joshua Espinosa

Brain & Coral





# <image>

### Honorable mention: Matthew Lydigsen

lighthouse + dalek



### Honorable mention: Makenzie Schwartz

hillary + trump



7

# Honorable mention: Nathan Greenberg

6

8

?? + simon pegg



# <section-header><section-header><section-header>

# A framework for alignment

- Matching local features
  - Local information used, can contain outliers
  - But hopefully enough of these matches are good



# Image alignment

• But first are there any questions?

# A framework for alignment

10

- Matching local features
  - Local information used, can contain outliers
  - But hopefully enough of these matches are good
- Consensus building
  - Aggregate the good matches and find a transformation that explains these matches



# <section-header><section-header><section-header><section-header><image><image><image>

### Generating putative correspondences



### Feature detection with scale selection

• We want to extract features with characteristic scale that matches the image transformation such as scaling and translation (a.k.a. covariance)



Matching regions across scales



Source: L. Lazebnik 15

### Blob detection: basic idea

- Convolve the image with a "blob filter" at multiple scales
- Look for extrema (maxima or minima) of filter response in the resulting *scale space*
- This will give us a scale and space covariant detector





### Blob detection: basic idea

Find maxima *and minima* of blob filter response in space *and scale* minima



# Blob detection in 2D

Laplacian of Gaussian: Circularly symmetric operator for blob detection in 2D



# Scale selection

Source: N. Snavely 18

- At what scale does the Laplacian achieve a maximum response to a binary circle of radius r?
- To get maximum response, the zeros of the Laplacian have to be aligned with the circle
- The Laplacian is given by (up to scale):



### Characteristic scale

• We define the characteristic scale of a blob as the scale that produces peak of Laplacian response in the blob center



### Scale-space blob detector

1. Convolve image with scale-normalized Laplacian at several scales

Source: L. Lazebnik 22

24

# Scale-space blob detector: Example



# Scale-space blob detector: Example



sigma = 11.9912

### Scale-space blob detector

- 1. Convolve image with scale-normalized Laplacian at several scales
- 2. Find maxima of squared Laplacian response in scalespace



Source: L. Lazebnik 25

27

### Scale-space blob detector: Example



Source: L. Lazebnik 26



### From feature detection to description

- Scaled and rotated versions of the same neighborhood will give rise to blobs that are related by the same transformation
- What to do if we want to compare the appearance of these image regions?
  - Normalization: transform these regions into same-size circles
  - Problem: rotational ambiguity





Source: L. Lazebnik 28

# Eliminating rotation ambiguity

• To assign a unique orientation to circular image windows:

- Create histogram of local gradient directions in the patch
- Assign canonical orientation at peak of smoothed histogram



### SIFT features

• Detected features with characteristic scales and orientations:



From feature detection to description



how should we represent the patches?

# Feature descriptors

- Simplest descriptor: vector of raw intensity values
- How to compare two such vectors?
  - Sum of squared differences (SSD) this is a distance measure

$$SSD(\mathbf{u}, \mathbf{v}) = \sum_{i} (u_i - v_i)^2$$

- Not invariant to intensity change
- Normalized correlation this is a similarity measure

$$\rho(\mathbf{u}, \mathbf{v}) = \frac{(\mathbf{u} - \overline{\mathbf{u}})}{\|\mathbf{u} - \overline{\mathbf{u}}\|} \cdot \frac{(\mathbf{v} - \overline{\mathbf{v}})}{\|\mathbf{v} - \overline{\mathbf{v}}\|} = \frac{\sum_{i} (u_{i} - \overline{\mathbf{u}})(v_{i} - \overline{\mathbf{v}})}{\sqrt{\left(\sum_{j} (u_{j} - \overline{\mathbf{u}})^{2}\right)\left(\sum_{j} (v_{j} - \overline{\mathbf{v}})^{2}\right)}}$$

32

- Invariant to affine (translation + scaling) intensity change

Source: L. Lazebnik 31

### Problem with intensity vectors as descriptors

• Small deformations can affect the matching score a lot



### Feature descriptors: SIFT

- Descriptor computation:
  - Divide patch into 4x4 sub-patches
  - Compute histogram of gradient orientations (8 reference angles) inside each sub-patch
  - Resulting descriptor: 4x4x8 = 128 dimensions



David G. Lowe. "Distinctive image features from scale-invariant keypoints." *IJCV* 60 (2), pp. 91-110, 2004.

# Feature descriptors: SIFT

- Descriptor computation:
  - Divide patch into 4x4 sub-patches
  - Compute histogram of gradient orientations (8 reference angles) inside each sub-patch
  - Resulting descriptor: 4x4x8 = 128 dimensions
- Advantage over raw vectors of pixel values
  - Gradients less sensitive to illumination change
  - Pooling of gradients over the sub-patches achieves robustness to small shifts, but still preserves some spatial information

David G. Lowe. "Distinctive image features from scale-invariant keypoints." *IJCV* 60 (2), pp. 91-110, 2004.

35

# Problem: Ambiguous putative matches



Source: Y. Furukawa 36

### Rejection of unreliable matches

- How can we tell which putative matches are more reliable?
- Heuristic: compare distance of **nearest** neighbor to that of **second** nearest neighbor
- Ratio of closest distance to second-closest distance will be *high* for features that are *not* distinctive



# RANSAC

- Random Sample Consensus
  - · Choose a small subset of points uniformly at random
  - · Fit a model to that subset
  - Find all remaining points that are "close" to the model and reject the rest as outliers
  - · Do this many times and choose the best model
- For rigid transformation we can estimate the parameters of the **transformation**, e.g., rotation angle, scaling, translation, etc, from **putative** correspondence matches
- Lets see how RANSAC works for a simple example.

M. A. Fischler, R. C. Bolles. <u>Random Sample Consensus: A Paradigm for Model Fitting with</u> <u>Applications to Image Analysis and Automated Cartography</u>. Comm. of the ACM, Vol 24, pp 381-395, 1981.



# RANSAC for line fitting example











### RANSAC for line fitting example



# RANSAC for line fitting example



### RANSAC for line fitting example





# RANSAC for line fitting

Repeat N times:

- Draw s points uniformly at random
- Fit line to these **s** points
- Find *inliers* to this line among the remaining points (i.e., points whose distance from the line is less than *t*)
- If there are *d* or more inliers and the number of inliers is higher than the *previous best*, accept the line and refit using all inliers.

# RANSAC for image matching

• Given a number of putative matches, select **inliers** based on their agreement with an underlying **transformation** 





• How do we represent image transformations?

# Families of transformation



Figure 2.4 Basic set of 2D planar transformations.

From Computer Vision: Algorithms and Applications, by Rick Szeliski

51

### How do we move pixels?

• Think about moving coordinates, not pixels.

### Points in 2D



56









### Identity: $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ Uniform scaling: $\begin{bmatrix} s & 0 \\ 0 & s \end{bmatrix}$ Scaling in x: $\left[ \begin{array}{cc} s_x & 0 \\ 0 & 1 \end{array} ight]$ Rotation by $\theta$ radians: $\cos(\theta) - \sin(\theta)$ $\sin(\theta) \quad \cos(\theta)$

 $\begin{bmatrix} 1 & sh_x \\ 0 & 1 \end{bmatrix}$ 

 $a \quad b$ c d

Shearing in x:

Arbitrary linear transformation:

Families of linear transforms

# Affine transformations

• Affine = linear + translation

 $\mathbf{x} \to \mathbf{M} \mathbf{x} + \mathbf{t}$ 

- Approximates viewpoint changes for roughly planar objects and roughly orthographic cameras
- Can be used to initialize fitting for more complex models



# Fitting an affine transformation

• Assume we know the correspondences, how do we get the transformation?



# Fitting an affine transformation

• Assume we know the correspondences, how do we get the transformation?



# Fitting an affine transformation

62

$$\begin{bmatrix} & \cdots & & & \\ x_i & y_i & 0 & 0 & 1 & 0 \\ 0 & 0 & x_i & y_i & 0 & 1 \\ & \cdots & & & & 1 \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \cdots \\ x'_i \\ y'_i \\ \cdots \end{bmatrix}$$

- Linear system with six unknowns
- Each match gives us two linearly independent equations: need at least three to solve for the transformation parameters

# Application: Panorama stitching



# Panoramic stitching

### • Approach

- Local feature matching
- RANSAC for alignment



# Panoramic stitching



# Panoramic stitching



- Extract features
  - corner/blob detector

### Panoramic stitching



- Extract features
- Compute *putative matches*

### Panoramic stitching



- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation *T*

# Panoramic stitching



- Extract features
- Compute *putative matches*
- Loop:
- *Hypothesize* transformation *T*
- Verify transformation (search for other matches consistent with T)

# Warping images

70

72

To shift an image (dx, dy)

- Let (ox, oy) be the coordinates of a pixel in the original image with a particular appearance
- Let (nx, ny) be the new coordinates, i.e., where we want that pixel
- For each pixel
  - nx = ox + dx
  - ny = oy + dy
  - newIm(ny, nx) = im(oy, ox);

### Mechanics of transformation

**procedure** forwardWarp(f, h, out g):

For every pixel  $\boldsymbol{x}$  in  $f(\boldsymbol{x})$ 

- 1. Compute the destination location x' = h(x).
- 2. Copy the pixel f(x) to g(x').

From Computer Vision: Algorithms and Applications, by Rick Szeliski

### Mechanics of transformation

**procedure** forwardWarp(f, h, out g):

For every pixel  $\boldsymbol{x}$  in  $f(\boldsymbol{x})$ 

- 1. Compute the destination location x' = h(x).
- 2. Copy the pixel f(x) to g(x').
- Problems with this transformation
  - Leaves gaps in the target image
  - Interpolation less intuitive

# Example of forward warp

Rotation by 45 degrees



# Mechanics of transformation

74

76

**procedure** inverseWarp(f, h, out g):

For every pixel  $\pmb{x}'$  in  $g(\pmb{x}')$ 

- 1. Compute the source location  $\boldsymbol{x} = \boldsymbol{\hat{h}}(\boldsymbol{x}')$
- 2. Resample f(x) at location x and copy to g(x')

From Computer Vision: Algorithms and Applications, by Rick Szeliski



# Panoramic stitching warping



