Announcements

- PS1 grades up on SmartSite
- PS1 stats:
  - Mean: 83.26
  - Standard Dev: 28.51
- PS2 deadline extended to *Saturday, 11:59 am* due to SmartSite maintenance

2

William Field

3

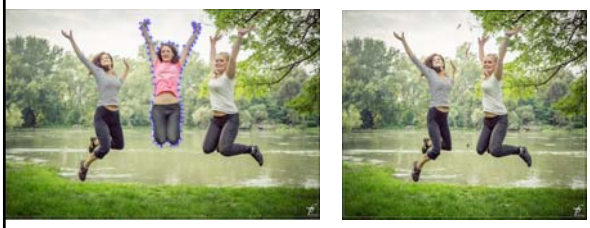

Miguel Covarrubias

4

Muhammad Elias

5

Thomas Rong

6

Last time

- Image warping based on homography
- Detecting corner-like points in an image

7

Today

- Local invariant features
  - Detection of interest points
    - (Harris corner detection)
    - Scale invariant blob detection: LoG
  - Description of local patches
    - SIFT: Histograms of oriented gradients

8

Local features: main components

1) Detection: Identify the interest points

2) Description: Extract vector feature descriptor surrounding each interest point.

$\mathbf{x}_1 = [x_1^{(1)}, \ldots, x_d^{(1)}]$

$\mathbf{x}_2 = [x_1^{(2)}, \ldots, x_d^{(2)}]$

3) Matching: Determine correspondence between descriptors in two views

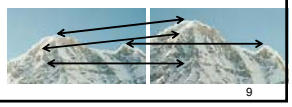

Kristen Grauman

9

Goal: interest operator repeatability

- We want to detect (at least some of) the same points in both images.

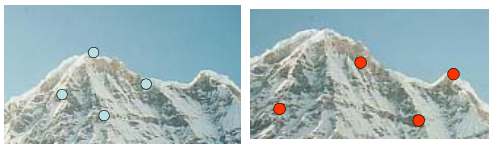

No chance to find true matches!

- Yet we have to be able to run the detection procedure *independently* per image.

10

Kristen Grauman

Goal: descriptor distinctiveness

- We want to be able to reliably determine which point goes with which.

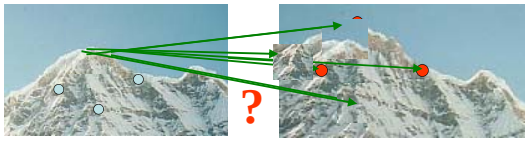

- Must provide some invariance to geometric and photometric differences between the two views.

11

Kristen Grauman

Local features: main components

1) Detection: Identify the interest points

2) Description:Extract vector feature descriptor surrounding each interest point.

3) Matching: Determine correspondence between descriptors in two views

12

Kristen Grauman

Recall: Corners as distinctive interest points

$$M = \sum w(x,y) \begin{bmatrix} I_x I_x & I_x I_y \\ I_x I_y & I_y I_y \end{bmatrix}$$

2 x 2 matrix of image derivatives (averaged in neighborhood of a point).

Notation: $I_x \Leftrightarrow \frac{\partial I}{\partial x}$ $\quad I_y \Leftrightarrow \frac{\partial I}{\partial y}$ $\quad I_x I_y \Leftrightarrow \frac{\partial I}{\partial x}\frac{\partial I}{\partial y}$

Kristen Grauman

13

Recall: Corners as distinctive interest points

Since *M* is symmetric, we have $M = X \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} X^T$

(Eigenvalue decomposition)

$$Mx_i = \lambda_i x_i$$

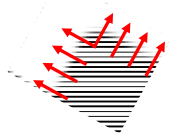

The *eigenvalues* of *M* reveal the amount of intensity change in the two principal orthogonal gradient directions in the window.

14

Recall: Corners as distinctive interest points

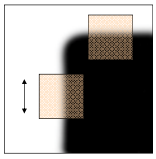

"edge":
$\lambda_1 >> \lambda_2$
$\lambda_2 >> \lambda_1$

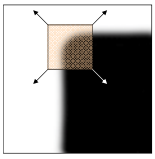

"corner":
λ_1 and λ_2 are large, $\lambda_1 \sim \lambda_2$;

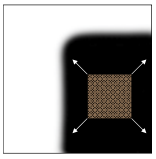

"flat" region
λ_1 and λ_2 are small;

One way to score the cornerness: $f = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2}$

15

Harris corner detector

1) Compute *M* matrix for image window surrounding each pixel to get its *cornerness* score.
2) Find points with large corner response (*f* > threshold)
3) Take the points of local maxima, i.e., perform non-maximum suppression

16

Kristen Grauman

Harris Detector: Steps

17

Kristen Grauman

Harris Detector: Steps

Compute corner response *f*

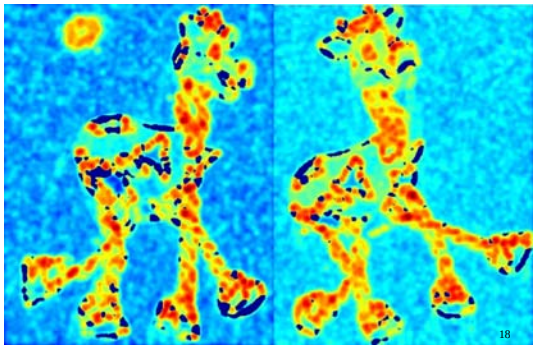

18

Kristen Grauman

Harris Detector: Steps

Find points with large corner response: f > threshold

Kristen Grauman 19

Harris Detector: Steps

Take only the points of local maxima of f

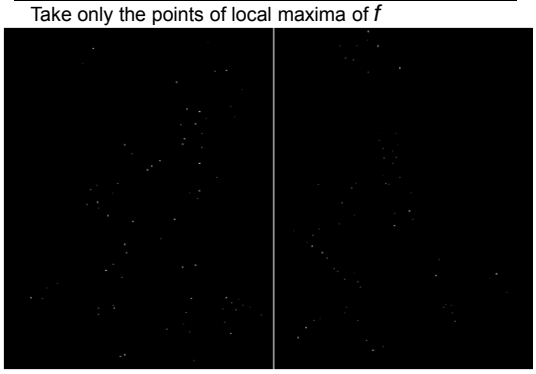

Kristen Grauman 20

Harris Detector: Steps

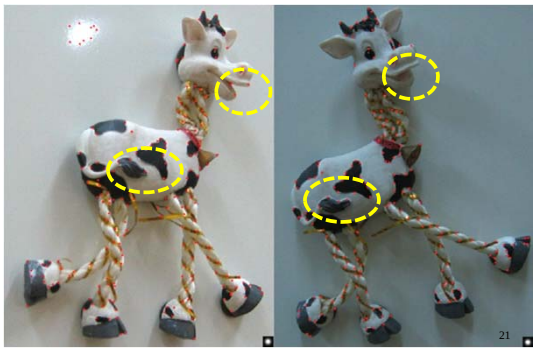

Kristen Grauman 21

Properties of the Harris corner detector

Rotation invariant? Yes

22

Properties of the Harris corner detector

Rotation invariant? Yes
Translation invariant? Yes

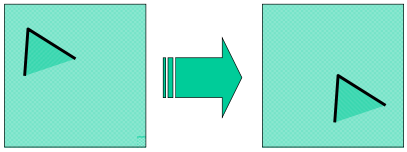

23

Properties of the Harris corner detector

Rotation invariant? Yes
Translation invariant? Yes
Scale invariant? No

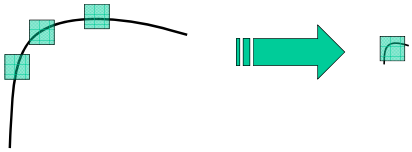

All points will be classified as edges

Corner !

24

Scale invariant interest points

How can we independently select interest points in each image, such that the detections are repeatable across different scales?

25

Kristen Grauman

Automatic scale selection

Intuition:

- Find scale that gives local maxima of some function *f* in both position and scale.

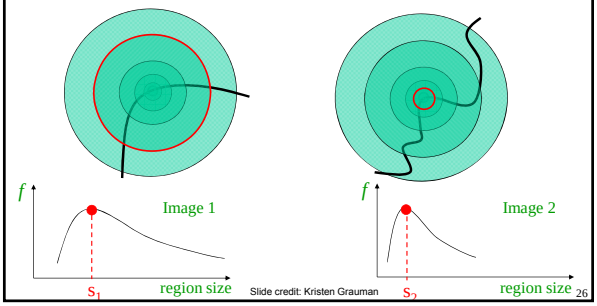

Slide credit: Kristen Grauman

26

What can be the "signature" function?

27

Kristen Grauman

Recall: Edge detection

f

$\frac{d}{dx} g$

$f * \frac{d}{dx} g$

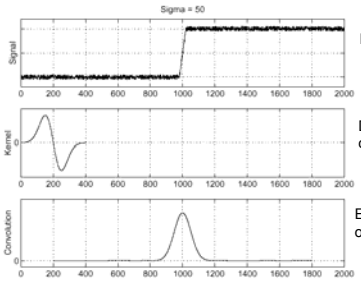

Edge

Derivative of Gaussian

Edge = maximum of derivative

28

Source: S. Seitz

Recall: Edge detection

f

$\frac{d^2}{dx^2} g$

$f * \frac{d^2}{dx^2} g$

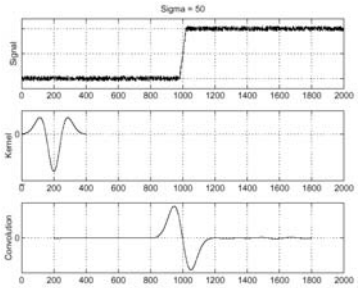

Edge

Second derivative of Gaussian (Laplacian)

Edge = zero crossing of second derivative

29

Source: S. Seitz

From edges to blobs

- Edge = ripple
- Blob = superposition of two ripples

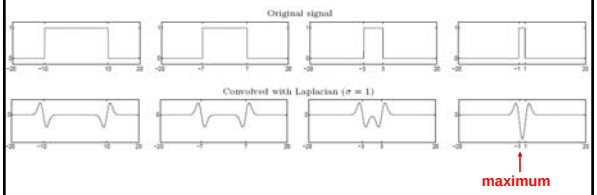

Spatial selection: the **magnitude** of the Laplacian response will achieve a maximum at the center of the blob, provided the scale of the Laplacian is "matched" to the scale of the blob

30

Slide credit: Lana Lazebnik

Blob detection in 2D

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

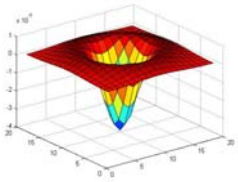

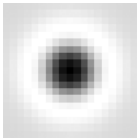

$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$$

31

Blob detection in 2D: scale selection

Laplacian-of-Gaussian = "blob" detector $\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$

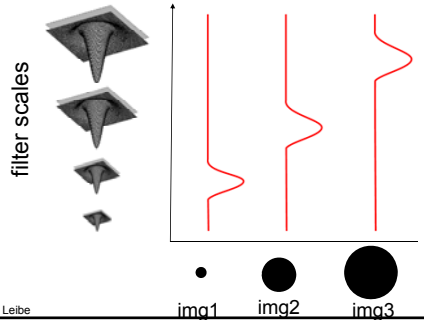

Bastian Leibe

32

Blob detection in 2D

We define the *characteristic scale* as the scale that produces peak of Laplacian response

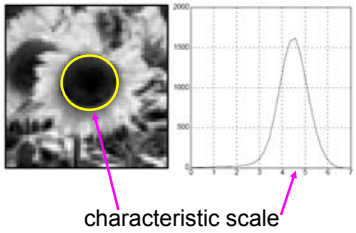

33

Slide credit: Lana Lazebnik

Example

Kristen Grauman

34

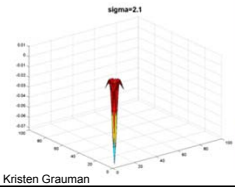

Kristen Grauman

35

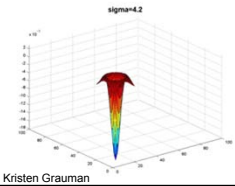

Kristen Grauman

36

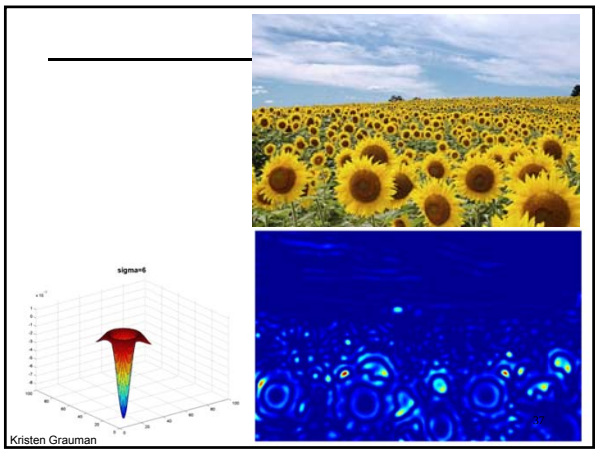
sigma=6
Kristen Grauman
37

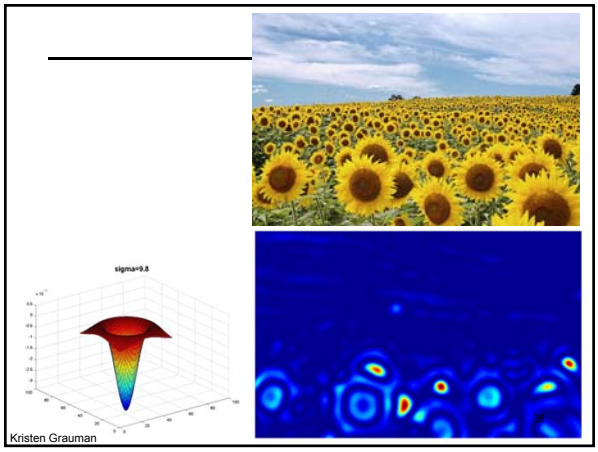
sigma=9.5
Kristen Grauman
38

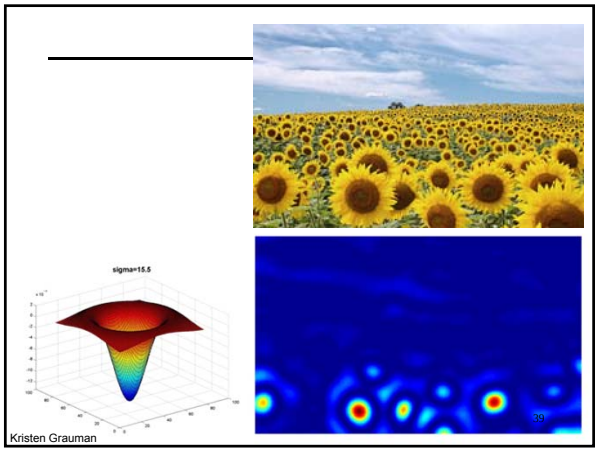
sigma=15.5
Kristen Grauman
39

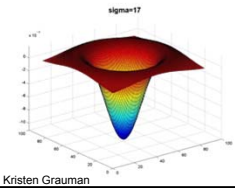

Kristen Grauman

40

Scale invariant interest points

Interest points are local maxima in both position and scale.

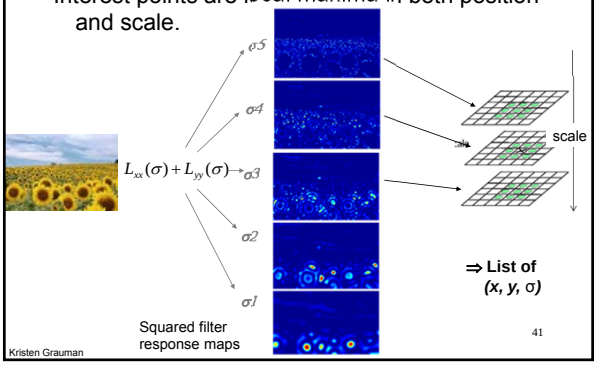

Squared filter response maps

⇒ **List of** ***(x, y, σ)***

Kristen Grauman

41

Scale-space blob detector: Example

42

Image credit: Lana Lazebnik

Technical detail

We can approximate the Laplacian with a difference of Gaussians; more efficient to implement.

$$L = \sigma^2 \left(G_{xx}(x,y,\sigma) + G_{yy}(x,y,\sigma) \right)$$

(Laplacian)

$$DoG = G(x,y,k\sigma) - G(x,y,\sigma)$$

(Difference of Gaussians)

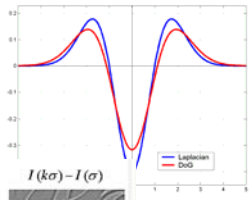

43

Slide credit: Kristen Grauman

Local features: main components

1) Detection: Identify the interest points

2) Description: Extract vector feature descriptor surrounding each interest point.

$$\mathbf{x}_1 = [x_1^{(1)}, \ldots, x_d^{(1)}]$$

$$\mathbf{x}_2 = [x_1^{(2)}, \ldots, x_d^{(2)}]$$

3) Matching: Determine correspondence between descriptors in two views

44

Slide credit: Kristen Grauman

Geometric transformations

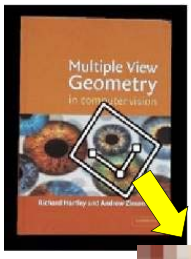

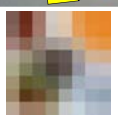

e.g. scale, translation, rotation

45

Slide credit: Kristen Grauman

Photometric transformations

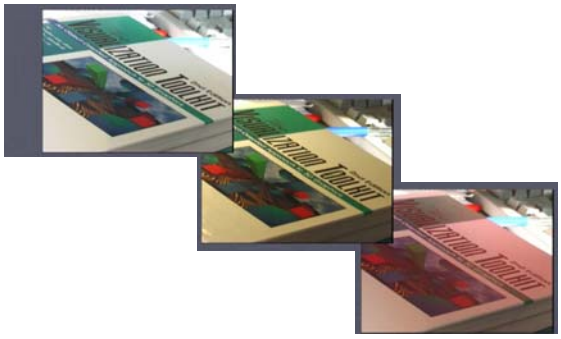

Figure from T. Tuytelaars ECCV 2006 tutorial

46

Raw patches as local descriptors

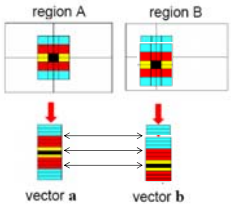

The simplest way to describe the neighborhood around an interest point is to write down the list of intensities to form a feature vector.

But this is very sensitive to even small shifts, rotations.

47

Slide credit: Kristen Grauman

SIFT descriptor [Lowe 2004]

- Use histograms to bin pixels within sub-patches according to their orientation.

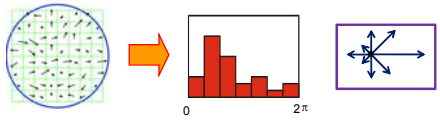

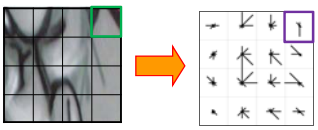

Why subpatches?

Why does SIFT have some illumination invariance?

48

Slide credit: Kristen Grauman

Making descriptor rotation invariant

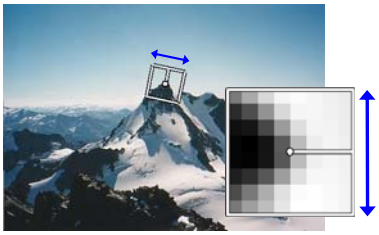

- Rotate patch according to its dominant gradient orientation
- This puts the patches into a canonical orientation.

49

Image from Matthew Brown

SIFT descriptor [Lowe 2004]

- Extraordinarily robust matching technique
  - Can handle changes in viewpoint
    - Up to about 60 degree out of plane rotation
  - Can handle significant changes in illumination
    - Sometimes even day vs. night (below)
  - Fast and efficient—can run in real time
  - Lots of code available
    - http://people.csail.mit.edu/albert/ladypack/wiki/index.php/Known_implementations_of_SIFT

50

Steve Seitz

Example

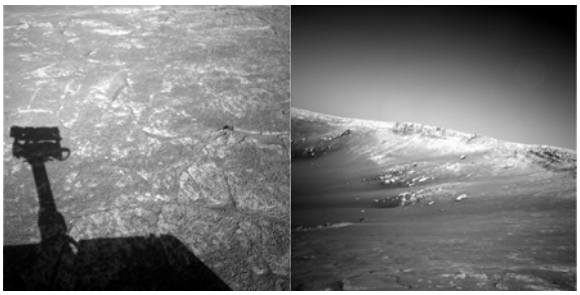

NASA Mars Rover images

51

Example

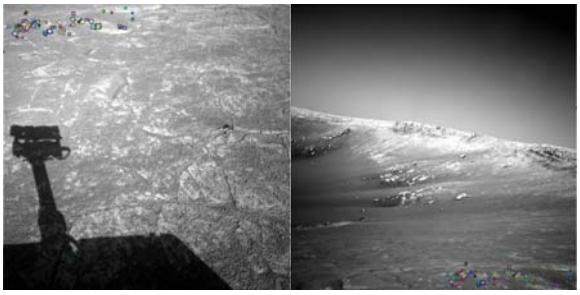

NASA Mars Rover images
with SIFT feature matches
Figure by Noah Snavely

52

SIFT descriptor properties

- Invariant to
  - Scale
  - Rotation
- Partially invariant to
  - Illumination changes
  - Camera viewpoint
  - Occlusion, clutter

53

Slide credit: Kristen Grauman

Local features: main components

1) Detection: Identify the interest points
2) Description: Extract vector feature descriptor surrounding each interest point.
3) Matching: Determine correspondence between descriptors in two views

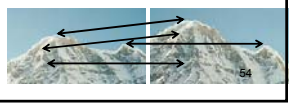

54

Slide credit: Kristen Grauman

Matching local features

Kristen Grauman

55

Matching local features

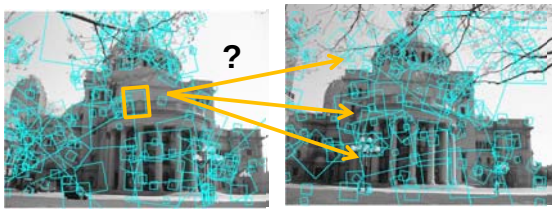

Image 1 Image 2

To generate **candidate matches**, find patches that have the most similar appearance (e.g., lowest SSD)

Simplest approach: compare them all, take the closest (or closest k, or within a thresholded distance)

Kristen Grauman

56

Ambiguous matches

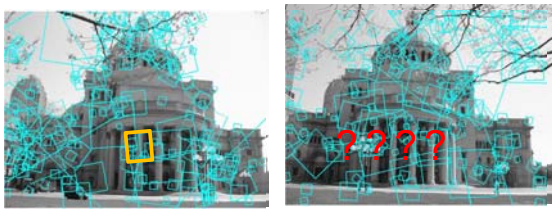

Image 1 Image 2

To add robustness to matching, can consider **ratio**: distance to best match / distance to second best match

If low, first match looks good.

If high, could be ambiguous match.

57

Kristen Grauman

Matching SIFT Descriptors

- Nearest neighbor (Euclidean distance)
- Threshold ratio of nearest to 2nd nearest descriptor

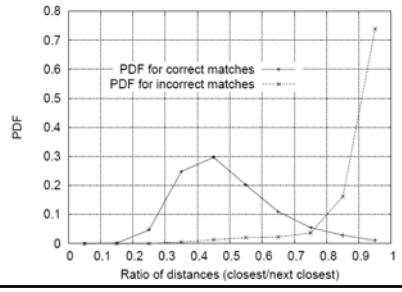

58

Lowe IJCV 2004

Recap: robust feature-based alignment

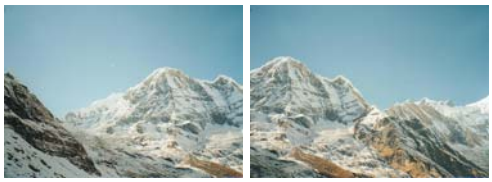

59

Source: L. Lazebnik

Recap: robust feature-based alignment

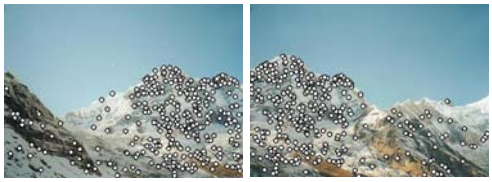

- Extract features

60

Source: L. Lazebnik

Recap: robust feature-based alignment

- Extract features
- Compute *putative matches*

61

Source: L. Lazebnik

Recap: robust feature-based alignment

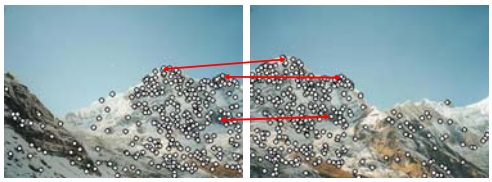

- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)

62

Source: L. Lazebnik

Recap: robust feature-based alignment

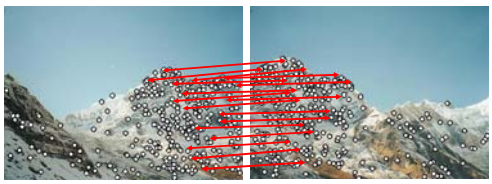

- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)
  - *Verify* transformation (search for other matches consistent with *T*)

63

Source: L. Lazebnik

Recap: robust feature-based alignment

- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)
  - *Verify* transformation (search for other matches consistent with *T*)

64

Source: L. Lazebnik

Applications of local invariant features

- Wide baseline stereo
- Motion tracking
- Panoramas
- Mobile robot navigation
- 3D reconstruction
- Recognition
- …

65

Slide credit: Kristen Grauman

Automatic mosaicing

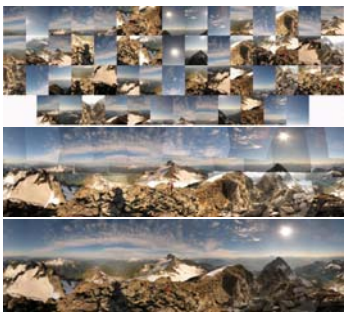

http://www.cs.ubc.ca/~mbrown/autostitch/autostitch.html

66

Slide credit: Kristen Grauman

Wide baseline stereo

[Image from T. Tuytelaars ECCV 2006 tutorial]

67

Recognition of specific objects, scenes

Schmid and Mohr 1997

Sivic and Zisserman, 2003

Rothganger et al. 2003

Lowe 2002

68

Kristen Grauman

Summary

- Interest point detection
  - Harris corner detector
  - Laplacian of Gaussian, automatic scale selection
- Invariant descriptors
  - Rotation according to dominant gradient direction
  - Histograms for robustness to small shifts and translations (SIFT descriptor)

69

Questions?

See you Thursday!

70