Interest Points



Computer Vision Jia-Bin Huang, Virginia Tech

Many slides from N Snavely, K. Grauman & Leibe, and D. Hoiem

Administrative Stuffs

- HW 1 posted, due 11:59 PM Sept 19
 - Submission through Canvas
- Frequently Asked Questions for HW 1 posted on piazza
- Reading Szeliski: 4.1

What have we learned so far?

- Light and color
 - What an image records
- Filtering in spatial domain
 - Filtering = weighted sum of neighboring pixels
 - Smoothing, sharpening, measuring texture
- Filtering in frequency domain
 - Filtering = change frequency of the input image
 - Denoising, sampling, image compression
- Image pyramid and template matching
 - Filtering = a way to find a template
 - Image pyramids for coarse-to-fine search and multi-scale detection
- Edge detection
 - Canny edge = smooth -> derivative -> thin -> threshold -> link
 - Finding straight lines, binary image analysis











This module: Correspondence and Alignment

• Correspondence:

matching points, patches, edges, or regions across images



This module: Correspondence and Alignment

• Alignment/registration:

solving the transformation that makes two things match better



The three biggest problems in computer vision?

1. Registration

2. Registration

3. Registration



Takeo Kanade (CMU)

Example: automatic panoramas



Credit: Matt Brown

Example: fitting an 2D shape template



Example: fitting a 3D object model



Example: estimating "fundamental matrix" that corresponds two views



Example: tracking points



frame 0

frame 22

frame 49

Today's class

- What is interest point?
- Corner detection
- Handling scale and orientation
- Feature matching

Why extract features?

- Motivation: panorama stitching
 - We have two images how do we combine them?



Why extract features?

- Motivation: panorama stitching
 - We have two images how do we combine them?



Step 1: extract features Step 2: match features

Why extract features?

- Motivation: panorama stitching
 - We have two images how do we combine them?



Step 1: extract features Step 2: match features Step 3: align images

Image matching



by <u>Diva Sian</u>



by swashford

Harder case



by <u>Diva Sian</u>

by <u>scgbt</u>

Harder still?



Answer below (look for tiny colored squares...)



NASA Mars Rover images with SIFT feature matches

Applications

- Keypoints are used for:
 - Image alignment
 - 3D reconstruction
 - Motion tracking
 - Robot navigation
 - Indexing and database retrieval
 - Object recognition







Advantages of local features

Locality

• features are local, so robust to occlusion and clutter

Quantity

• hundreds or thousands in a single image

Distinctiveness:

• can differentiate a large database of objects

Efficiency

• real-time performance achievable

Overview of Keypoint Matching





 f_{B}

 $d(f_A, f_B) \! < \! T$

1. Find a set of distinctive keypoints

- 2. Define a region around each keypoint
- 3. Extract and normalize the region content
- 4. Compute a local descriptor from the normalized region

5. Match local descriptors

Goals for Keypoints





Detect points that are *repeatable* and *distinctive*

Key trade-offs

More Repeatable

Robust detection Precise localization

Description

More Distinctive

Minimize wrong matches

More Points

Robust to occlusion Works with less texture

More Flexible Robust to expected variations Maximize correct matches

Choosing interest points

Where would you tell your friend to meet you?

Choosing interest points

Where would you tell your friend to meet you?

Corner Detection: Basic Idea

- How does the window change when you shift it?
- Shifting the window in any direction causes a big change

"flat" region: no change in all directions

"edge": no change along the edge direction

"corner": significant change in all directions

Credit: S. Seitz, D. Frolova, D. Simakov

Harris corner detection: the math

Consider shifting the window W by (u, v)

- how do the pixels in W change?
- compare each pixel before and after by summing up the squared differences (SSD)
- this defines an SSD "error" E(u,v):

$$E(u,v) = \sum_{(x,y)\in W} (I(x+u,y+v) - I(x,y))^2$$

Small motion assumption

Taylor Series expansion of *I*:

$$I(x+u, y+v) = I(x, y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \text{higher order terms}$$

If the motion (u,v) is small, then first order approximation is good

$$I(x + u, y + v) \approx I(x, y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v$$
$$\approx I(x, y) + [I_x \ I_y] \begin{bmatrix} u \\ v \end{bmatrix}$$
shorthand: $I_x = \frac{\partial I}{\partial x}$

Plugging this into the formula on the previous slide...

Harris corner detection: the math

Using the small motion assumption, replace *I* with a linear approximation

(Shorthand:
$$I_x = \frac{\partial I}{\partial x}$$
)

$$E(u, v) = \sum_{(x,y)\in W} (I(x+u, y+v) - I(x,y))^{2}$$

$$\approx \sum_{(x,y)\in W} (I(x,y) + I_{x}(x,y)u + I_{y}(x,y)v - I(x,y))^{2}$$

$$\approx \sum_{(x,y)\in W} (I_{x}(x,y)u + I_{y}(x,y)v)^{2}$$

Corner detection: the math

$$E(u, v) \approx \sum_{(x,y)\in W} \left(I_x(x, y)u + I_y(x, y)v \right)^2$$

$$\approx \sum_{(x,y)\in W} \left(I_x^2 u^2 + 2I_x I_y uv + I_y^2 v^2 \right)$$
$$\approx A u^2 + 2B uv + C v^2$$
$$A = \sum_{(x,y)\in W} I_x^2 \quad B = \sum_{(x,y)\in W} I_x I_y \quad C = \sum_{(x,y)\in W} I_y^2$$

• Thus, *E*(*u*,*v*) is locally approximated as a *quadratic form*

The second moment matrix

The surface E(u,v) is locally approximated by a quadratic form.

 $E(u,v) \approx Au^2 + 2Buv + Cv^2$ $\approx \left[\begin{array}{ccc} u & v \end{array} \right] \left| \begin{array}{ccc} A & B \\ B & C \end{array} \right| \left| \begin{array}{ccc} u \\ v \end{array} \right|$ $A = \sum I_x^2$ $(x,y) \in W$ $B = \sum I_x I_y$ $(x,y) \in W$ $C = \sum I_y^2$ $(x,y) \in W$ Let's try to understand its shape.

$$E(u,v) \approx \begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} A & B \\ B & C \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

$$H$$

$$E(u,v) \approx \begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} A & B \\ B & C \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

$$H$$

General case The shape of *H* tells us something about the *distribution* of gradients around a pixel

We can visualize *H* as an ellipse with axis lengths determined by the *eigenvalues* of *H* and orientation determined by the *eigenvectors* of *H*

Quick eigenvalue/eigenvector

review

The **eigenvectors** of a matrix **A** are the vectors **x** that satisfy:

$$Ax = \lambda x$$

The scalar λ is the **eigenvalue** corresponding to \boldsymbol{x}

• The eigenvalues are found by solving:

$$det(A - \lambda I) = 0$$

• In our case, **A** = **H** is a 2x2 matrix, so we have

$$det \left[\begin{array}{cc} h_{11} - \lambda & h_{12} \\ h_{21} & h_{22} - \lambda \end{array} \right] = 0$$

• The solution:

$$\lambda_{\pm} = \frac{1}{2} \left[(h_{11} + h_{22}) \pm \sqrt{4h_{12}h_{21} + (h_{11} - h_{22})^2} \right]$$

Once you know λ , you find **x** by solving

$$\begin{bmatrix} h_{11} - \lambda & h_{12} \\ h_{21} & h_{22} - \lambda \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 0$$


Eigenvalues and eigenvectors of H

- Define shift directions with the smallest and largest change in error
- x_{max} = direction of largest increase in *E*
- λ_{max} = amount of increase in direction x_{max}
- x_{min} = direction of smallest increase in *E*
- λ_{min} = amount of increase in direction x_{min}

Corner detection: the math

How are λ_{max} , x_{max} , λ_{min} , and x_{min} relevant for feature detection?

• What's our feature scoring function?

Corner detection: the math

How are λ_{max} , x_{max} , λ_{min} , and x_{min} relevant for feature detection?

• What's our feature scoring function?

Want E(u,v) to be large for small shifts in all directions

- the minimum of E(u,v) should be large, over all unit vectors [u v]
- this minimum is given by the smaller eigenvalue (λ_{min}) of H



Interpreting the eigenvalues

Classification of image points using eigenvalues of M:



 λ_1

Corner detection summary

Here's what you do

- Compute the gradient at each point in the image
- Create the *H* matrix from the entries in the gradient
- Compute the eigenvalues.
- Find points with large response (λ_{min} > threshold)
- Choose those points where λ_{min} is a local maximum as features



Corner detection summary

Here's what you do

- Compute the gradient at each point in the image
- Create the *H* matrix from the entries in the gradient
- Compute the eigenvalues.
- Find points with large response (λ_{min} > threshold)
- Choose those points where λ_{min} is a local maximum as features





The Harris operator

 λ_{min} is a variant of the "Harris operator" for feature detection

$$f = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2}$$
$$= \frac{determinant(H)}{trace(H)}$$

- The *trace* is the sum of the diagonals, i.e., *trace(H)* = $h_{11} + h_{22}$
- Very similar to λ_{min} but less expensive (no square root)
- Called the "Harris Corner Detector" or "Harris Operator"
- Lots of other detectors, this is one of the most popular

The Harris operator



Harris operator

min

Harris detector example





f value (red high, blue low)



Threshold (f > value)



Find local maxima of f

••• •

Harris features (in red)



Weighting the derivatives

 In practice, using a simple window W doesn't work too well

$$H = \sum_{(x,y)\in W} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

 Instead, we'll weight each derivative value based on its distance from the center pixel

$$H = \sum_{(x,y)\in W} w_{x,y} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$



 $w_{x,y}$

Harris Detector – Responses [Harris88]



Harris Detector – Responses [Harris88]



Harris Detector: Invariance Properties

Rotation



Ellipse rotates but its shape (i.e. eigenvalues) remains the same

Corner response is invariant to image rotation

Harris Detector: Invariance Properties

• Affine intensity change: $I \rightarrow aI + b$



✓ Intensity scale: $I \rightarrow a I$



Partially invariant to affine intensity change

Harris Detector: Invariance PropertiesScaling



Not invariant to scaling

Scale invariant detection Suppose you're looking for corners



Key idea: find scale that gives local maximum of f

- in both position and scale
- One definition of *f*: the Harris operator



How to find corresponding patch sizes?

• Function responses for increasing scale (scale signature)







 $f(I_{i_1...i_m}(x',\sigma))$

• Function responses for increasing scale (scale signature)







• Function responses for increasing scale (scale signature)







• Function responses for increasing scale (scale signature)



2.0 3.89 scale 19

 $f(I_{i_1...i_m}(x,\sigma))$





• Function responses for increasing scale (scale signature)





 $f(I_{i_1...i_m}(x,\sigma))$



12.0 scale 19.

 $f(I_{i_1...i_m}(x',\sigma))$

K. Grauman, B. Leibe

• Function responses for increasing scale (scale signature)









K. Grauman, B. Leibe

Implementation

 Instead of computing *f* for larger and larger windows, we can implement using a fixed window size with a Gaussian pyramid







(sometimes need to create inbetween levels, e.g. a ³/₄-size image)

What Is A Useful Signature Function?

• Difference-of-Gaussian = "blob" detector



Difference-of-Gaussian (DoG)







K. Grauman, B. Leibe



DoG – Efficient Computation

Computation in Gaussian scale pyramid



Find local maxima in position-scale space of Difference-of-Gaussian



Results: Difference-of-Gaussian



Orientation Normalization

[Lowe, SIFT, 1999]

- Compute orientation histogram
- Select dominant orientation
- Normalize: rotate to fixed orientation



Maximally Stable Extremal Regions [Matas '02]

- Based on Watershed segmentation algorithm
- Select regions that stay stable over a large parameter range


Example Results: MSER



Available at a web site near you...

- For most local feature detectors, executables are available online:
 - -<u>http://www.robots.ox.ac.uk/~vgg/research/affine</u>
 - <u>http://www.cs.ubc.ca/~lowe/keypoints/</u>
 - http://www.vision.ee.ethz.ch/~surf

Local Descriptors

- The ideal descriptor should be
 - Robust
 - Distinctive
 - Compact
 - Efficient
- Most available descriptors focus on edge/gradient information
 - Capture texture information
 - Color rarely used

Scale Invariant Feature Transform

Basic idea:

- Take 16x16 square window around detected feature
- Compute edge orientation (angle of the gradient 90°) for each pixel
- Throw out weak edges (threshold gradient magnitude)
- Create histogram of surviving edge orientations



SIFT descriptor

Full version

- Divide the 16x16 window into a 4x4 grid of cells (2x2 case shown below)
- Compute an orientation histogram for each cell
- 16 cells * 8 orientations = 128 dimensional descriptor



Local Descriptors: SIFT Descriptor



Histogram of oriented gradients

- Captures important texture information
- Robust to small translations / affine deformations

[Lowe, ICCV 1999]

K. Grauman, B. Leibe

Details of Lowe's SIFT algorithm

- Run DoG detector
 - Find maxima in location/scale space
 - Remove edge points
- Find all major orientations
 - Bin orientations into 36 bin histogram
 - Weight by gradient magnitude
 - Weight by distance to center (Gaussian-weighted mean)
 - Return orientations within 0.8 of peak
 - Use parabola for better orientation fit
- For each (x,y,scale,orientation), create descriptor:
 - Sample 16x16 gradient mag. and rel. orientation
 - Bin 4x4 samples into 4x4 histograms
 - Threshold values to max of 0.2, divide by L2 norm
 - Final descriptor: 4x4x8 normalized histograms

 $\mathbf{H} = \begin{bmatrix} D_{xx} & D_{xy} \\ D_{xy} & D_{yy} \end{bmatrix}$ $\frac{\mathrm{Tr}(\mathbf{H})^2}{\mathrm{Det}(\mathbf{H})} < \frac{(r+1)^2}{r}$

SIFT Example



sift

4 DI GRADEN SCOT MEELOUD

868 SIFT features

Feature matching

Given a feature in I_1 , how to find the best match in I_2 ?

- 1. Define distance function that compares two descriptors
- 2. Test all the features in I_2 , find the one with min distance

Feature distance

How to define the difference between two features f_1, f_2 ?

- Simple approach: L₂ distance, ||f₁ f₂ ||
- can give good scores to ambiguous (incorrect) matches





Feature distance

How to define the difference between two features f_1, f_2 ?

- Better approach: ratio distance = $||f_1 f_2|| / ||f_1 f_2'||$
 - f_2 is best SSD match to f_1 in I_2
 - f_2' is 2nd best SSD match to f_1 in I_2
 - gives large values for ambiguous matches





Feature matching example



51 matches

Feature matching example



58 matches

Evaluating the results

How can we measure the performance of a feature matcher?



feature distance

True/false positives

How can we measure the performance of a feature matcher?



feature distance

The distance threshold affects performance

- True positives = # of detected matches that are correct
 - Suppose we want to maximize these—how to choose threshold?
- False positives = # of detected matches that are incorrect
 - Suppose we want to minimize these—how to choose threshold?

Evaluating the results

How can we measure the performance of a feature matcher?



Evaluating the results

How can we measure the performance of a feature matcher?



Matching SIFT Descriptors

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



Lowe LICV 2004

SIFT Repeatability



Lowe IJCV 2004

SIFT Repeatability



SIFT Repeatability



Lowe IJCV 2004

Local Descriptors: SURF



Fast approximation of SIFT idea

Efficient computation by 2D box filters & integral images ⇒ 6 times faster than SIFT Equivalent quality for object identification

Many other efficient descriptors are also available

GPU implementation available

Feature extraction @ 200Hz (detector + descriptor, 640×480 img) http://www.vision.ee.ethz.ch/~surf

[Bay, ECCV'06], [Cornelis, CVGPU'08]

K. Grauman, B. Leibe

Local Descriptors: Shape Context



Count the number of points inside each bin, e.g.:

- Count = 4 : Count = 10

Log-polar binning: more precision for nearby points, more flexibility for farther points.

Belongie & Malik, ICCV 2001

Local Descriptors: Geometric Blur



Choosing a detector

- What do you want it for?
 - Precise localization in x-y: Harris
 - Good localization in scale: Difference of Gaussian
 - Flexible region shape: MSER
- Best choice often application dependent
 - Harris-/Hessian-Laplace/DoG work well for many natural categories
 - MSER works well for buildings and printed things
- Why choose?
 - Get more points with more detectors
- There have been extensive evaluations/comparisons
 - [Mikolajczyk et al., IJCV'05, PAMI'05]
 - All detectors/descriptors shown here work well

Comparison of Keypoint Detectors

Table 7.1 Overview of feature detectors.

				Rotation	Scale	Affine		Localization		
Feature Detector	Corner	Blob	Region	invariant	invariant	invariant	Repeatability	accuracy	Robustness	Efficiency
Harris	\checkmark			\checkmark			+++	+++	+++	++
Hessian		\checkmark		\checkmark			++	++	++	+
SUSAN	\sim			\checkmark			++	++	++	+++
Harris-Laplace	\checkmark	(√)		\checkmark	\checkmark		+++	+++	++	+
Hessian-Laplace	()			\checkmark	\checkmark		+++	+++	+++	+
DoG	(\checkmark		\checkmark	\checkmark		++	++	++	++
SURF	()	\checkmark		\checkmark	\checkmark		++	++	++	+++
Harris-Affine	\checkmark	(√)		\checkmark	\checkmark	\checkmark	+++	+++	++	++
Hessian-Affine	(\checkmark	\checkmark		+++	+++	+++	++
Salient Regions	(\checkmark		\checkmark	\checkmark	()	+	+	++	+
Edge-based	\checkmark			\checkmark	\checkmark		+++	+++	+	+
MSER			\checkmark	\checkmark	\checkmark	\checkmark	+++	+++	++	+++
Intensity-based			\checkmark	\checkmark	\checkmark	\checkmark	++	++	++	++
Superpixels			\checkmark	\checkmark	(√)	()	+	+	+	+

Tuytelaars Mikolajczyk 2008

Choosing a descriptor

- Again, need not stick to one
- For object instance recognition or stitching, SIFT or variant is a good choice

Recent advances in interest points



Binary feature descriptors

- BRIEF: Binary Robust Independent Elementary Features, ECCV 10
- ORB (Oriented FAST and Rotated BRIEF), CVPR 11
- BRISK: Binary robust invariant scalable keypoints, ICCV 11
- Freak: Fast retina keypoint, CVPR 12
- LIFT: Learned Invariant Feature Transform, ECCV 16

Things to remember

- Keypoint detection: repeatable and distinctive
 - Corners, blobs, stable regions
 - Harris, DoG



- spatial histograms of orientation
- SIFT



Image gradients

Keypoint descriptor

Thank you

• Next class: feature tracking and optical flow

