


CS4670: Computer Vision

Noah Snavely


Lecture 6: Feature matching





Reading

- Szeliski: 4.1



Feature extraction: Corners and blobs



The Harris operator



Harris operator



λ_{\min}


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: 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: $I \rightarrow aI + b$
 - ✓ Only derivatives are used => invariance to intensity shift $I \rightarrow I + b$
 - ✓ Intensity scale: $I \rightarrow aI$



Partially invariant to affine intensity change

Harris Detector: Invariance Properties


- Scaling



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

Automatic scale selection


Lindeberg et al., 1996



$$f(I_{l_1 \dots l_m}(x, \sigma))$$


Slide from Tinne Tuytelaars

Automatic scale selection




$$f(I_{l_1...l_m}(x, \sigma))$$

Automatic scale selection




$$f(I_{l_1...l_m}(x, \sigma))$$

Automatic scale selection




$$f(I_{l_1...l_m}(x, \sigma))$$

Automatic scale selection




$$f(I_{l_1 \dots l_m}(x, \sigma))$$

Automatic scale selection




$$f(I_{l_1 \dots l_m}(x, \sigma))$$

Automatic scale selection



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



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



Automatic scale selection

Normalize: rescale to fixed size



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 in-between levels, e.g. a $\frac{3}{4}$ -size image)

Another common definition of f

- The *Laplacian of Gaussian (LoG)*



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


(very similar to a Difference of Gaussians (DoG) –
i.e. a Gaussian minus a slightly smaller Gaussian)

Laplacian of Gaussian

- “Blob” detector



$$* \text{LoG} =$$




minima

maximum


- Find maxima *and minima* of LoG operator in space and scale

Scale selection


- At what scale does the Laplacian achieve a maximum response for a binary circle of radius r ?



image




Laplacian



Characteristic scale

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



T. Lindeberg (1998). ["Feature detection with automatic scale selection."](#)
International Journal of Computer Vision **30** (2): pp 77--116.

Scale-space blob detector: Example




Scale-space blob detector: Example



sigma = 11.9912


Scale-space blob detector: Example



Questions?

Feature descriptors


We know how to detect good points
Next question: **How to match them?**



Answer: Come up with a *descriptor* for each point,
find similar descriptors between the two images

Feature descriptors

We know how to detect good points
Next question: **How to match them?**



Lots of possibilities (this is a popular research area)

- Simple option: match square windows around the point
- State of the art approach: SIFT
 - David Lowe, UBC <http://www.cs.ubc.ca/~lowe/keypoints/>

Invariance vs. discriminability

- Invariance:
 - Descriptor shouldn't change even if image is transformed
- Discriminability:
 - Descriptor should be highly unique for each point

Invariance

- Most feature descriptors are designed to be invariant to
 - Translation, 2D rotation, scale
- They can usually also handle
 - Limited 3D rotations (SIFT works up to about 60 degrees)
 - Limited affine transformations (some are fully affine invariant)
 - Limited illumination/contrast changes

How to achieve invariance

Need both of the following:

1. Make sure your detector is invariant
2. Design an invariant feature descriptor
 - Simplest descriptor: a single 0
 - What's this invariant to?
 - Next simplest descriptor: a square window of pixels
 - What's this invariant to?
 - Let's look at some better approaches...