


# CS4670/5670: Computer Vision

Kavita Bala

## Lecture 12: Scale Space Feature Detectors



# Announcements

- HW 1 and PA 2 out tonight or tomorrow
- Schedule will be updated shortly
- Artifact voting out today
  - Please vote

# Feature extraction: Corners



# The second moment matrix


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

$$E(u, v) \approx Au^2 + 2Buv + Cv^2$$
$$\approx \begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} A & B \\ B & C \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

$$A = \sum_{(x,y) \in W} I_x^2$$

$$B = \sum_{(x,y) \in W} I_x I_y$$

$$C = \sum_{(x,y) \in W} I_y^2$$



Let's try to understand its shape.




# Interpreting the second moment matrix

Consider a horizontal “slice” of  $E(u, v)$ :  $[u \ v] M \begin{bmatrix} u \\ v \end{bmatrix} = \text{const}$

This is the equation of an ellipse.


Diagonalization of  $M$ :  $M = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$

The axis lengths of the ellipse are determined by the eigenvalues and the orientation is determined by  $R$



# Corner detection: the math

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



$$M x_{\max} = \lambda_{\max} x_{\max}$$

$$M x_{\min} = \lambda_{\min} x_{\min}$$

## Eigenvalues and eigenvectors of M

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

# The Harris operator

$\lambda_{\min}$  is a variant of the “Harris operator” for feature detection


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

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

# Corner response function

$$R = \det(M) - \alpha \text{trace}(M)^2 = \lambda_1 \lambda_2 - \alpha(\lambda_1 + \lambda_2)^2$$


$\alpha$ : constant (0.04 to 0.15)



# Corner detection summary

Here's what you do

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




$\lambda_{\min}$

# Harris features (in red)



# Harris Corners – Why so complicated?

- Can't we just check for regions with lots of gradients in the x and y directions?
  - No! A diagonal line would satisfy that criteria







# Image transformations

- Geometric

**Rotation**



**Scale**



- Photometric

**Intensity change**




# Invariance and covariance

- We want corner locations to be *invariant* to photometric transformations and *covariant* to geometric transformations
  - **Invariance:** image is transformed and corner locations do not change
  - **Covariance:** if we have two transformed versions of the same image, features should be detected in corresponding locations




# Image translation



- Derivatives and window function are shift-invariant

Corner location is covariant w.r.t. translation


# Image rotation



Second moment ellipse rotates but its shape (i.e. eigenvalues) remains the same



Corner location is covariant w.r.t. rotation

# Affine intensity change



$$I \rightarrow aI + b$$


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



*Partially invariant to affine intensity change*

# Harris Detector: Invariance Properties

- Scaling



*Not invariant to scaling*


So far: can localize in x-y, but not scale





# 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



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


Slide from Tinne Tuytelaars

# Automatic scale selection




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

# Automatic scale selection



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


# Automatic scale selection



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




# Automatic scale selection



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

# Automatic scale selection




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




# Implementation

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





# Feature extraction: Corners and blobs



# 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)

# Scale selection


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



image



Laplacian






# Laplacian of Gaussian

- “Blob” detector



$$* \text{LoG} =$$




minima

maximum

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

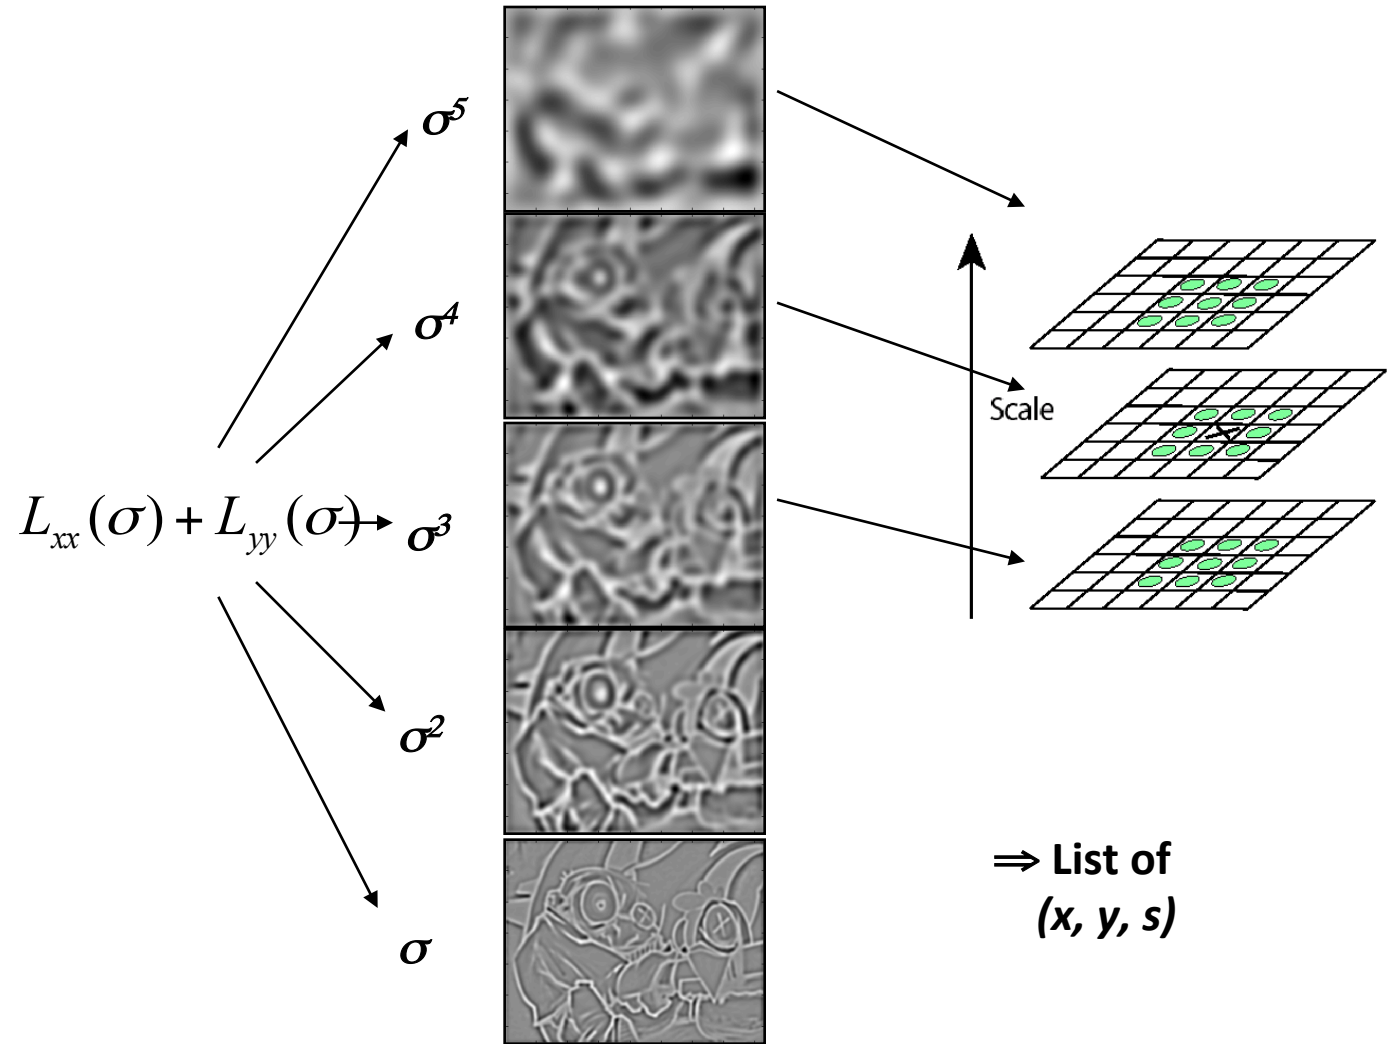

# Characteristic scale

- 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.

# Find local maxima in position-scale space





# Scale-space blob detector: Example




# Scale-space blob detector: Example



sigma = 11.9912


# Scale-space blob detector: Example



# 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




# Image transformations

- Geometric

**Rotation**




**Scale**



- Photometric

**Intensity change**



# 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

## 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...

# Rotation invariance for feature descriptors

- Find dominant orientation of the image patch
  - This is given by  $\mathbf{x}_{\max}$ , the eigenvector of  $\mathbf{M}$  corresponding to  $\lambda_{\max}$  (the *larger* eigenvalue)
  - Rotate the patch according to this angle




Figure by Matthew Brown


# Multiscale Oriented Patches descriptor

Take 40x40 square window around detected feature

- Scale to 1/5 size (using prefiltering)
- Rotate to horizontal
- Sample 8x8 square window centered at feature
- Intensity normalize the window by subtracting the mean, dividing by the standard deviation in the window



# Detections at multiple scales



*Figure 1. Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels from one of the Matier images. The boxes show the feature orientation and the region from which the descriptor vector is sampled.*



# 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