


CSC 589 Introduction to Computer Vision

Lecture 15 Feature Detection





Bei Xiao

Spring, 2014

American University

Motivation: Automatic panoramas



Microsoft puts some pizzazz into panoramic photos

The company's ICE software now can stitch video frames into panoramic images and fill in inevitable gaps. It shows the field of computational photography is still in its early days.

by **Stephen Shankland** [@stshank](#) / February 5, 2015 9:30 AM PST

[1](#) / [f](#) 88 / [t](#) 139 / [in](#) / [g+](#) / [more](#) +



Microsoft ICE stitches still photos or video frames into a single panoramic image. Version 2.0 can fill in gaps so you don't have to crop as much.

Screenshot by Stephen Shankland/CNET

http://gigapixelartzoom.com

- ICE software
- <http://research.microsoft.com/en-us/um/redmond/projects/ice/>

Approach

Feature detection: find it

Feature descriptor: represent it

Feature matching: match it

Feature tracking: track it, when motion

Why extract features?

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




Local features: main components

1) Detection: Identify the interest points




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

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




3) Matching: Determine correspondence between descriptors in two views

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



Characteristics of good features



- **Repeatability**
 - The same feature can be found in several images despite geometric and photometric transformations
- **Saliency**
 - Each feature is distinctive
- **Compactness and efficiency**
 - Many fewer features than image pixels
- **Locality**
 - A feature occupies a relatively small area of the image; robust to clutter and occlusion

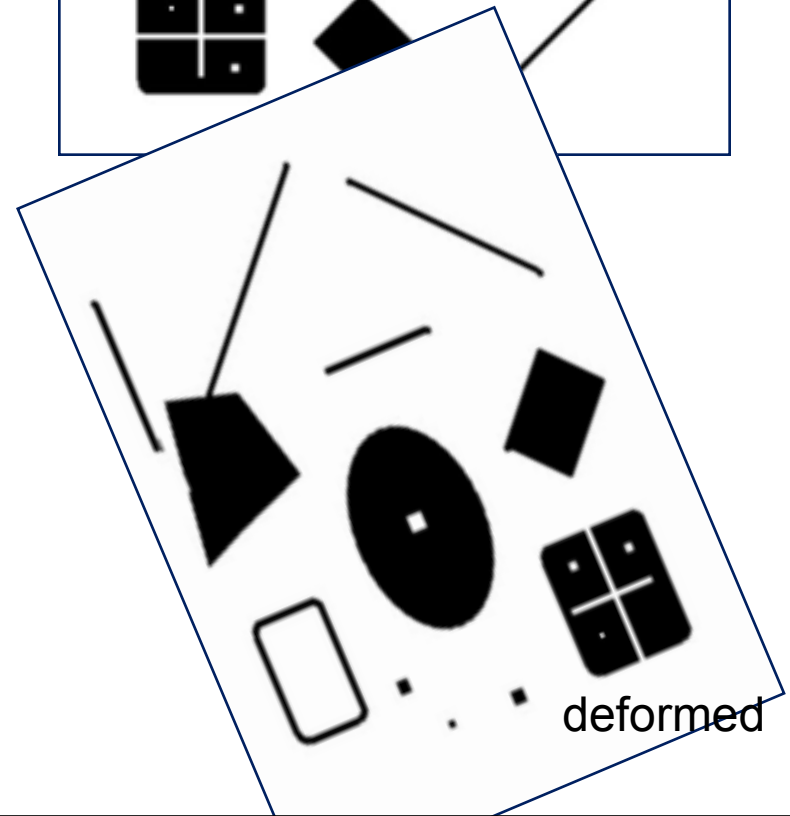

What makes a good feature?



What is a good feature?

- Suppose you have to click on some point, go away and come back after I deform the image, and click on the same points again.
 - Which points would you choose?

original



deformed

Want uniqueness

Look for image regions that are unusual

- Lead to unambiguous matches in other images

How to define “unusual”?


Choosing interest points

Where would you tell
your friend to meet
you?




Corner Detection: Basic Idea


- We should easily recognize the point by looking through a small window



“flat” region




“edge”




“corner”

Corner Detection: Basic Idea


- We should easily recognize the point by looking through a small window
- Shifting a window in *any direction* should give *a large change* in intensity



“flat” region:
no change in
all directions



“edge”:
no change along
the edge
direction



“corner”:
significant
change in all
directions

Many Existing Detectors Available

Hessian & Harris	[Beaudet '78], [Harris '88]
Laplacian, DoG	[Lindeberg '98], [Lowe 1999]
Harris-/Hessian-Laplace	[Mikolajczyk & Schmid '01]
Harris-/Hessian-Affine	[Mikolajczyk & Schmid '04]
EBR and IBR	[Tuytelaars & Van Gool '04]
MSER	[Matas '02]
Salient Regions	[Kadir & Brady '01]
Others...	

Finding Corners

- Corners are repeatable and distinctive
- Key property: in the region around a corner, image gradient has two or more dominant directions


C.Harris and M.Stephens. "[A Combined Corner and Edge Detector.](#)" *Proceedings of the 4th Alvey Vision Conference*: pages 147--151.

Feature extraction: Corners

9300 Harris Corners Pkwy, Charlotte, NC



Corner Detection Results




http://docs.opencv.org/trunk/doc/py_tutorials/py_feature2d/py_features_harris/py_features_harris.html

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: compute sum of squared differences (SSD)
- this defines an SSD “error” $E(u, v)$:



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

Finding difference in intensity for a displacement (u, v) in ALL directions

Corner Detection: Mathematics

Change in appearance of window $w(x,y)$
for the shift $[u,v]$:

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

Shifted
intensity


Intensity

Corner Detection: Mathematics


Change in appearance of window $w(x,y)$
for the shift $[u,v]$:

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

$I(x, y)$





$E(u, v)$



Corner Detection: Mathematics

Change in appearance of window $w(x,y)$
for the shift $[u,v]$:

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




Corner Detection: Mathematics

Change in appearance of window W
for the shift $[u, v]$:

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

We want to find out how this function behaves for small shifts

$E(u, v)$



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

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


shorthand: $I_x = \frac{\partial I}{\partial x}$

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

Feature detection: the math

Consider shifting the window W by (u, v)

- define an SSD “error” $E(u, v)$:




$$\begin{aligned} E(u, v) &= \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2 \\ &\approx \sum_{(x, y) \in W} [I(x, y) + [I_x \ I_y] \begin{bmatrix} u \\ v \end{bmatrix} - I(x, y)]^2 \\ &\approx \sum_{(x, y) \in W} \left[[I_x \ I_y] \begin{bmatrix} u \\ v \end{bmatrix} \right]^2 \end{aligned}$$

Corner detection: the math

Consider shifting the window W by (u, v)

- define an SSD “error” $E(u, v)$:



$$\begin{aligned} E(u, v) &= \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2 \\ &\approx \sum_{(x, y) \in W} [I(x, y) + I_x u + I_y v - I(x, y)]^2 \\ &\approx \sum_{(x, y) \in W} [I_x u + I_y v]^2 \end{aligned}$$

Corner detection: the math


Consider shifting the window W by (u, v)

- define an SSD “error” $E(u, v)$:

$$E(u, v) \approx \sum_{(x, y) \in W} [I_x u + I_y v]^2$$
$$\approx Au^2 + 2Buv + Cv^2$$

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

- Thus, $E(u, v)$ is locally approximated as a quadratic error function



Corner Detection: Mathematics

The quadratic approximation simplifies to

$$E(u, v) \approx [u \ v] M \begin{bmatrix} u \\ v \end{bmatrix}$$

where M is a *second moment matrix* computed from image derivatives (aka structure tensor):


$$M = \sum_{x,y} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

$$M = \begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix} = \sum \begin{bmatrix} I_x \\ I_y \end{bmatrix} [I_x \ I_y] = \sum \nabla I (\nabla I)^T$$

Corners as distinctive interest points

$$M = \sum \begin{bmatrix} I_x I_x & I_x I_y \\ I_x I_y & I_y I_y \end{bmatrix}$$

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



Notation:

$$I_x \Leftrightarrow \frac{\partial I}{\partial x}$$

$$I_y \Leftrightarrow \frac{\partial I}{\partial y}$$


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

Interpreting the second moment matrix

The surface $E(u,v)$ is locally approximated by a quadratic form. Let's try to understand its shape.

$$E(u, v) \approx [u \ v] M \begin{bmatrix} u \\ v \end{bmatrix}$$

$$M = \sum_{x,y} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$




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

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


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

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



Horizontal edge: $I_x = 0$

$$M = \begin{bmatrix} 0 & 0 \\ 0 & C \end{bmatrix}$$




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

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


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

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



Vertical edge: $I_y = 0$


$$M = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$$



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.



Questions?


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



Interpreting the eigenvalues


Classification of image points using eigenvalues of M :



Corner response function

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

α : constant (0.04 to 0.06)



Questions?

Harris corner detector

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


C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#) *Proceedings of the 4th Alvey Vision Conference*: pages 147—151, 1988.

Harris Detector [Harris88]


- Second moment matrix

$$\mu(\sigma_I, \sigma_D) = g(\sigma_I) * \begin{bmatrix} I_x^2(\sigma_D) & I_x I_y(\sigma_D) \\ I_x I_y(\sigma_D) & I_y^2(\sigma_D) \end{bmatrix}$$

1. Image derivatives
(optionally, blur first)



2. Square of derivatives



$$\det M = \lambda_1 \lambda_2$$

$$\text{trace } M = \lambda_1 + \lambda_2$$

3. Gaussian filter $g(\sigma_I)$



4. Cornerness function – both eigenvalues are strong

$$har = \det[\mu(\sigma_I, \sigma_D)] - \alpha [\text{trace}(\mu(\sigma_I, \sigma_D))]^2 =$$

$$g(I_x^2)g(I_y^2) - [g(I_x I_y)]^2 - \alpha [g(I_x^2) + g(I_y^2)]^2$$

5. Non-maxima suppression



Harris Detector [Harris88]

1. Compute x and y derivatives of image

$$I_x = G_\sigma^x * I \quad I_y = G_\sigma^y * I$$

2. Compute products of derivatives at every pixel

$$I_{x2} = I_x \cdot I_x \quad I_{y2} = I_y \cdot I_y \quad I_{xy} = I_x \cdot I_y$$

3. Compute the sums of the products of derivatives at each pixel

$$S_{x2} = G_{\sigma^2} * I_{x2} \quad S_{y2} = G_{\sigma^2} * I_{y2} \quad S_{xy} = G_{\sigma^2} * I_{xy}$$

4. Define at each pixel (x, y) the matrix

$$H(x, y) = \begin{bmatrix} S_{x2}(x, y) & S_{xy}(x, y) \\ S_{xy}(x, y) & S_{y2}(x, y) \end{bmatrix}$$

5. Compute the response of the detector at each pixel

$$R = \text{Det}(H) - k(\text{Trace}(H))^2$$

6. Threshold on value of R . Compute nonmax suppression.

Implement Harris Corner in Python

```
from scipy.ndimage import filters
```

```
def compute_harris_response(im, sigma=3)
```

- Step 1: filter image in both x and y direction with a Gaussian filter
- Step 2: Compute Components of the Harris Matrix
- Step 3: Get determinant and trace of Harris matrix, Compute the ratio. This will be the response

```
gx = - x * exp(-(x**2/float((0.5*size)**2)+y**2/float((0.5*sizey)**2)))
```

Take-home reading

- Szeliski Chapter 4.1.1 Feature detector
- Solem, Chapter 2.1: Local Image Descriptors/
Harris Corner Detector.
- Harris Original Paper:
- [http://www.bmva.org/bmvc/1988/
avc-88-023.pdf](http://www.bmva.org/bmvc/1988/avc-88-023.pdf)