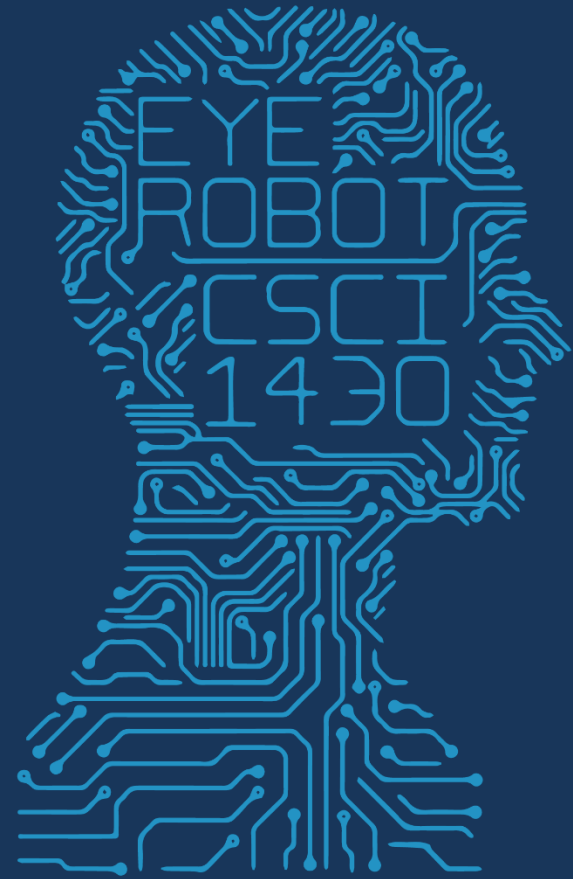




1950

FUTURE VISION





2017 MWF 1PM 368

COMPUTER VISION

Last lecture: Edges primer

- Edge detection to identify visual change in image
- Derivative of Gaussian and linear combination of convolutions
- What is an edge?
What is a good edge?



Canny edge detector

- Probably the most widely used edge detector in computer vision.
- Theoretical model: step-edges corrupted by additive Gaussian noise.
- Canny showed that first derivative of Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization.

J. Canny, [**A Computational Approach To Edge Detection**](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.


22,000 citations!

Examples: Controversy and Appropriateness



‘Lena’

Alexander Sawchuk @ USC, 1973



‘Fabio’

Deanna Needell @ Claremont McKenna, 2012

If it wasn't clear from class...

Use of Lena is now generally considered inappropriate, as it is not inclusive.


If you wish to include models, then please respect diversity.

We will approach diversity issues in more detail when we start to talk about machine learning and datasets.



Canny edge detector

1. Filter image with x, y derivatives of Gaussian


Derivative of Gaussian filter



x-direction



y-direction



Compute Gradients



X-Derivative of Gaussian



Y-Derivative of Gaussian



Canny edge detector

1. Filter image with x, y derivatives of Gaussian
2. Find magnitude and orientation of gradient

Compute Gradient Magnitude



$\text{sqrt}(\text{X-Deriv.ofGaussian}^2 + \text{Y-Deriv.ofGaussian}^2)$


= gradient



Compute Gradient Orientation

Threshold magnitude at minimum level


Get orientation via $\theta = \text{atan2}(g_y, g_x)$



Canny edge detector


1. Filter image with x, y derivatives of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel wide “ridges” to single pixel width

Non-maximum suppression for each orientation





At pixel q:
We have a maximum if the
value is larger than those at
both p and at r.

Interpolate along gradient
direction to get these values.






Sidebar: Bilinear Interpolation




$$f(x, y) \approx \begin{bmatrix} 1 - x & x \end{bmatrix} \begin{bmatrix} f(0, 0) & f(0, 1) \\ f(1, 0) & f(1, 1) \end{bmatrix} \begin{bmatrix} 1 - y \\ y \end{bmatrix}.$$

Sidebar: Interpolation options

- `imx2 = imresize(im, 2, interpolation_type)`
- 'nearest'
 - Copy value from nearest known
 - Very fast but creates blocky edges
- 'bilinear'
 - Weighted average from four nearest known pixels
 - Fast and reasonable results
- 'bicubic' (default)
 - Non-linear smoothing over larger area (4x4)
 - Slower, visually appealing, may create negative pixel values




Non-maximum suppression for each orientation



At pixel q:
We have a maximum if the
value is larger than those at
both p and at r.

Interpolate along gradient
direction to get these values.



Before Non-max Suppression



Gradient magnitude

After non-max suppression




Gradient magnitude

Canny edge detector


1. Filter image with x, y derivatives of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel wide “ridges” to single pixel width
4. ‘Hysteresis’ Thresholding:
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them
 - ‘Follow’ edges starting from strong edge pixels
 - Connected components (Szeliski 3.3.4)

'Hysteresis' thresholding

- Two thresholds – high and low
- Grad. mag. $>$ high threshold? = strong edge
- Grad. mag. $<$ low threshold? noise
- In between = weak edge
- 'Follow' edges starting from strong edge pixels
- Continue them into weak edges
 - Connected components (Szeliski 3.3.4)



Final Canny Edges



Effect of σ (Gaussian kernel spread/size)



Original



Canny with $\sigma = 1$



Canny with $\sigma = 2$


The choice of σ depends on desired behavior


- large σ detects large scale edges
- small σ detects fine features

Canny edge detector

1. Filter image with x, y derivatives of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel wide “ridges” to single pixel width
4. ‘Hysteresis’ Thresholding:
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them
 - ‘Follow’ edges starting from strong edge pixels
 - Connected components (Szeliski 3.3.4)

MATLAB: `edge(image, ‘canny’)`






Interest Points and Corners

Szeliski 4.1



“Interest points” = “keypoints”
Sometimes called “features”

Correspondence across views

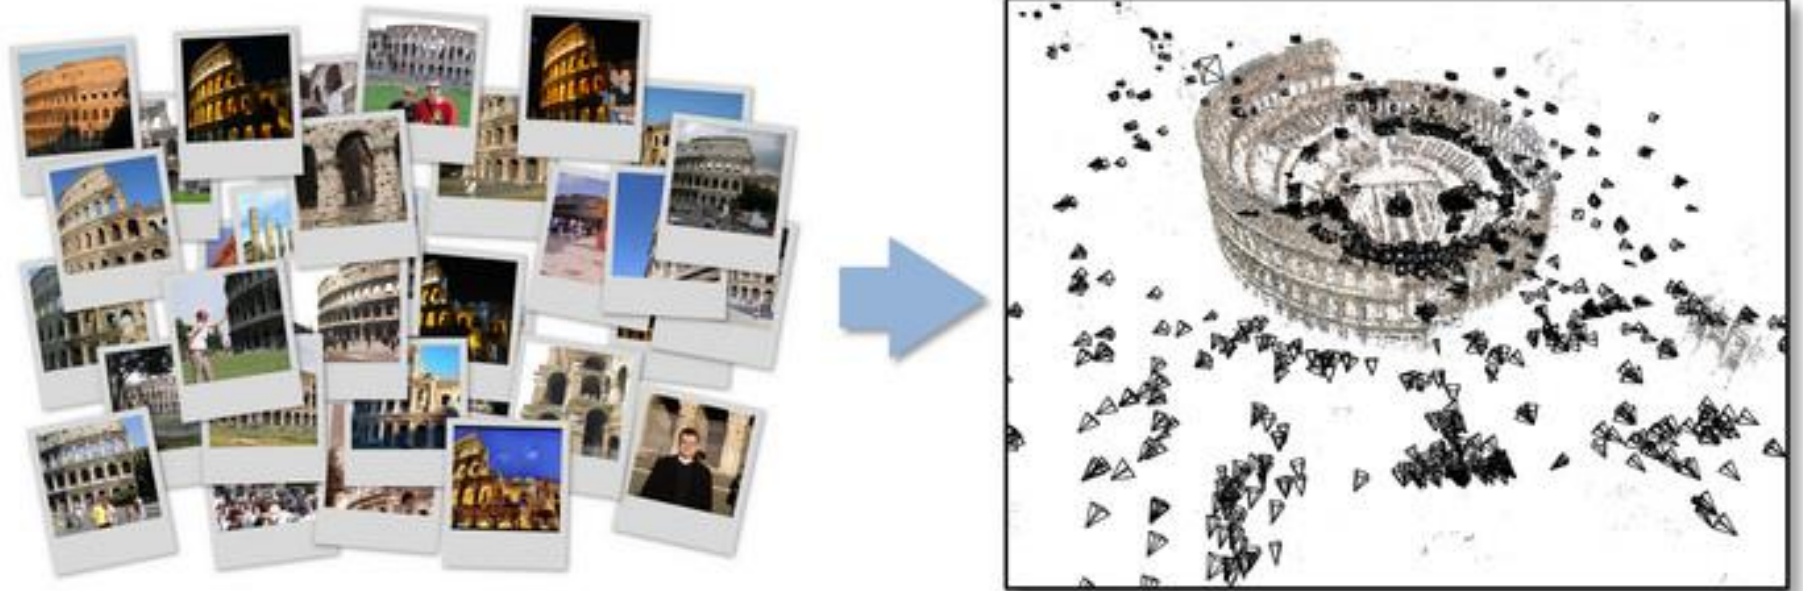
- Correspondence: matching points, patches, edges, or regions across images.



Example: estimate “fundamental matrix”
that corresponds two views




Example: structure from motion



Fundamental to Applications

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



Example application


- Panorama stitching
 - We have two images – how do we combine them?



Local features: main components


1) Detection:

Find a set of distinctive key points.





2) Description:

Extract feature descriptor around each interest point as vector.



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




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


3) Matching:

Compute distance between feature vectors to find correspondence.

$$d(\mathbf{x}_1, \mathbf{x}_2) < T$$




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

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.

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.

Local features: main components

1) Detection:

Find a set of distinctive key points.



2) Description:

Extract feature descriptor around each interest point as vector.

3) Matching:

Compute distance between feature vectors to find correspondence.

Detection: Basic Idea

- We do not know which other image locations the feature will end up being matched against.
- But we can compute how stable a location is in appearance with respect to small variations in position u .
- *Compare image patch against local neighbors.*

Corner Detection: Mathematics

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


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

Window
function

Shifted
intensity


Intensity

Window function $w(x,y) =$



1 in window, 0 outside

or





Gaussian

Corner Detection: Mathematics

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



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



Corner Detection: Mathematics

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

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






TPS Time

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




(a)


$$E(u, v)$$




$$E(u, v)$$



(b)



(c)




(d)


Inverted, and in 3D!

Corner Detection: Basic Idea


- We might recognize the point by looking through a small window.
- We want a window shift in *any direction* to 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

Corner Detection: Mathematics

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


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

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

But this is very slow to compute naively.

$O(\text{window_width}^2 * \text{shift_range}^2 * \text{image_width}^2)$

$O(11^2 * 11^2 * 600^2) = 5.2$ billion of these
14.6 thousand per pixel in your image





Corner Detection: Mathematics

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

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

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

But we know the response in E that we are looking for – strong peak.




Recall: Taylor series expansion

A function f can be represented by an infinite series of its derivatives at a single point a :

$$f(a) + \frac{f'(a)}{1!}(x-a) + \frac{f''(a)}{2!}(x-a)^2 + \frac{f'''(a)}{3!}(x-a)^3 + \dots$$

Set $a = 0$
(MacLaurin series)
as window centered

Approximation of
 $f(x) = e^x$
centered at $f(0)$



Taylor expansion in 2D

Local quadratic approximation of $E(u,v)$ in the neighborhood of $(0,0)$ is given by the *second-order Taylor expansion*:

$$E(u,v) \approx E(0,0) + [u \ v] \begin{bmatrix} E_u(0,0) \\ E_v(0,0) \end{bmatrix} + \frac{1}{2} [u \ v] \begin{bmatrix} E_{uu}(0,0) & E_{uv}(0,0) \\ E_{uv}(0,0) & E_{vv}(0,0) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

Corner Detection: Mathematics

Local quadratic approximation of $E(u,v)$ in the neighborhood of $(0,0)$ is given by the *second-order Taylor expansion*:

$$E(u,v) \approx E(0,0) + [u \ v] \begin{bmatrix} E_u(0,0) \\ E_v(0,0) \end{bmatrix} + \frac{1}{2} [u \ v] \begin{bmatrix} E_{uu}(0,0) & E_{uv}(0,0) \\ E_{uv}(0,0) & E_{vv}(0,0) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$



Set to 0



First
derivative,
set to 0

Corner Detection: Mathematics

The quadratic approximation simplifies to

$$E(u, v) \approx [u \ v] \begin{bmatrix} E_{uu}(0,0) & E_{uv}(0,0) \\ E_{uv}(0,0) & E_{vv}(0,0) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

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

where M is a *second moment matrix* computed from image derivatives:




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

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

Corners as distinctive interest points

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

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



Notation:

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

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


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

Interpreting the second moment matrix

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

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


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



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.




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



Visualization of second moment matrices




Visualization of second moment matrices



Interpreting the eigenvalues

Classification of image points using eigenvalues of M :



Corner response function

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


α : constant (0.04 to 0.06)

Determinant ($\det(A)$):

$$|A| = \begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc.$$

Trace ($\text{trace}(A)$):

$$\text{tr}(A) = a_{11} + a_{22} + \cdots + a_{nn} = \sum_{i=1}^n a_{ii}$$



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 ($f > \text{threshold}$)
- 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




3. Gaussian filter $g(\sigma_I)$




4. Cornerness function – both eigenvalues are strong

$$\begin{aligned} har &= \det[\mu(\sigma_I, \sigma_D)] - \alpha [\text{trace}(\mu(\sigma_I, \sigma_D))]^2 \\ &= g(I_x^2)g(I_y^2) - [g(I_x I_y)]^2 - \alpha [g(I_x^2) + g(I_y^2)]^2 \end{aligned}$$

5. Non-maxima suppression




Harris Detector: Steps




Harris Detector: Steps

Compute corner response R



Harris Detector: Steps

Find points with large corner response: $R > \text{threshold}$



Harris Detector: Steps

Take only the points of local maxima of R



Harris Detector: Steps



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

