



CPSC 425: Computer Vision

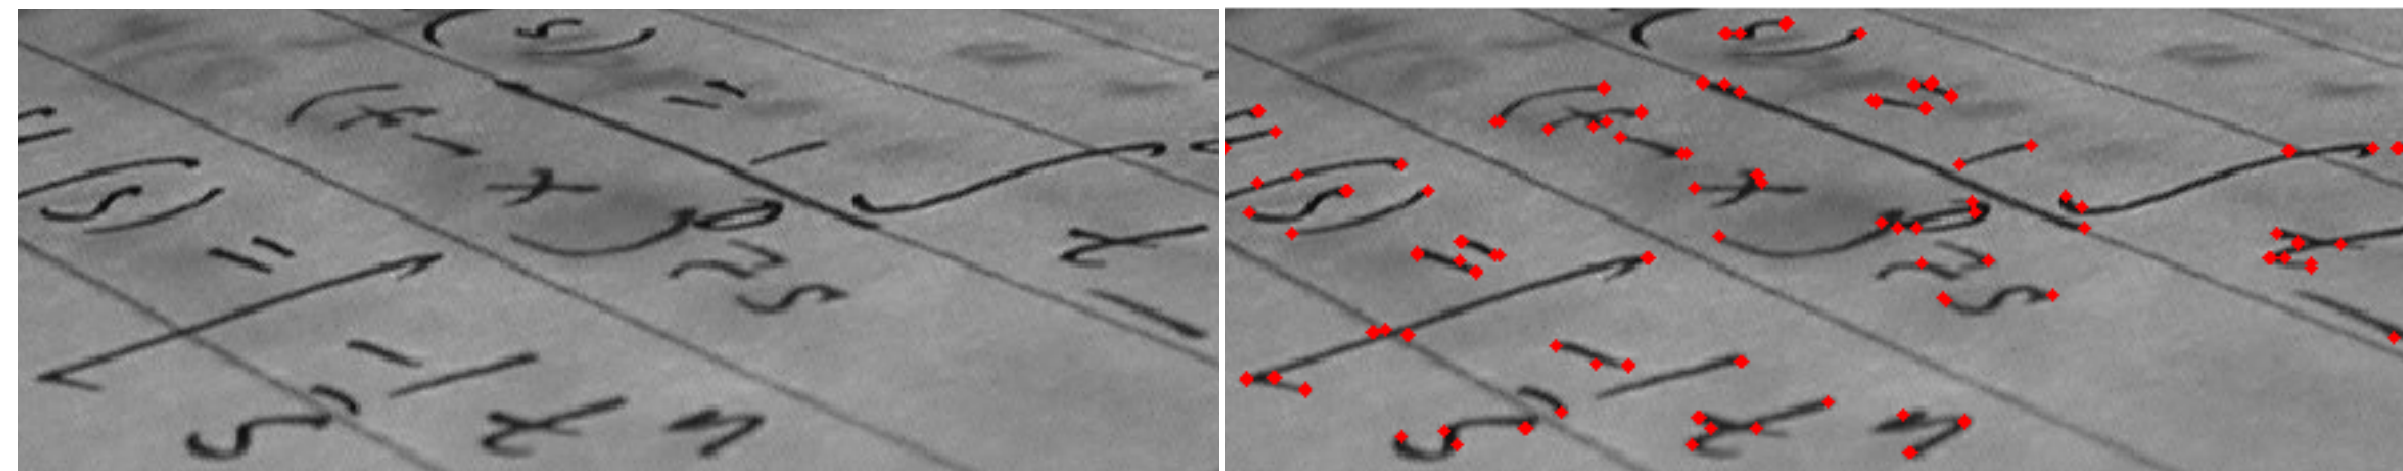


Image Credit: https://en.wikipedia.org/wiki/Corner_detection

Lecture 9: Corner Detection

(unless otherwise stated slides are taken or adopted from **Bob Woodham, Jim Little** and **Fred Tung**)

Menu for Today (February 4, 2020)

Topics:

- **Corner** Detection
- **Autocorrelation**
- **Harris** Corner Detector

Readings:

- **Today's** Lecture: Forsyth & Ponce (2nd ed.) 5.3.0 - 5.3.1
- **Next** Lecture: Forsyth & Ponce (2nd ed.) 6.1, 6.3

Reminders:

- **Assignment 2:** Face Detection in a Scaled Representation is **February 11th**

Today's **“fun”** Example:



Today's "fun" Example:



Today's **“fun”** Example:

Lecture 8: Re-cap

Physical properties of a 3D scene cause “**edges**” in an image:

- depth discontinuity
- surface orientation discontinuity
- reflectance discontinuity
- illumination boundaries

Two generic approaches to **edge detection**:

- local extrema of a first derivative operator → **Canny**
- zero crossings of a second derivative operator → **Marr/Hildreth**

Many algorithms consider “**boundary detection**” as a high-level recognition task and output a probability or confidence that a pixel is on a human-perceived boundary

Motivation: Template Matching

When might **template matching fail**?

— Different scales



— Different orientation



— Lighting conditions



— Left vs. Right hand



— Partial Occlusions



— Different Perspective

— Motion / blur

Motivation: Template Matching in Scaled Representation

When might **template matching** in scaled representation **fail**?

— ~~Different scales~~



— Different orientation



— Lighting conditions



— Left vs. Right hand



— Partial Occlusions



— Different Perspective

— Motion / blur

Motivation: Edge Matching in Scaled Representation

When might **edge matching** in scaled representation **fail**?

— ~~Different scales~~




— Different orientation



— ~~Lighting conditions~~



— Left vs. Right hand



— Partial Occlusions



— Different Perspective

— Motion / blur

Planar Object Instance Recognition

Database of planar objects



Instance recognition






Recognition under **Occlusion**




Image Matching



Image Matching




Finding **Correspondences**




NASA Mars Rover images

Finding Correspondences




What is a **Good Feature**?



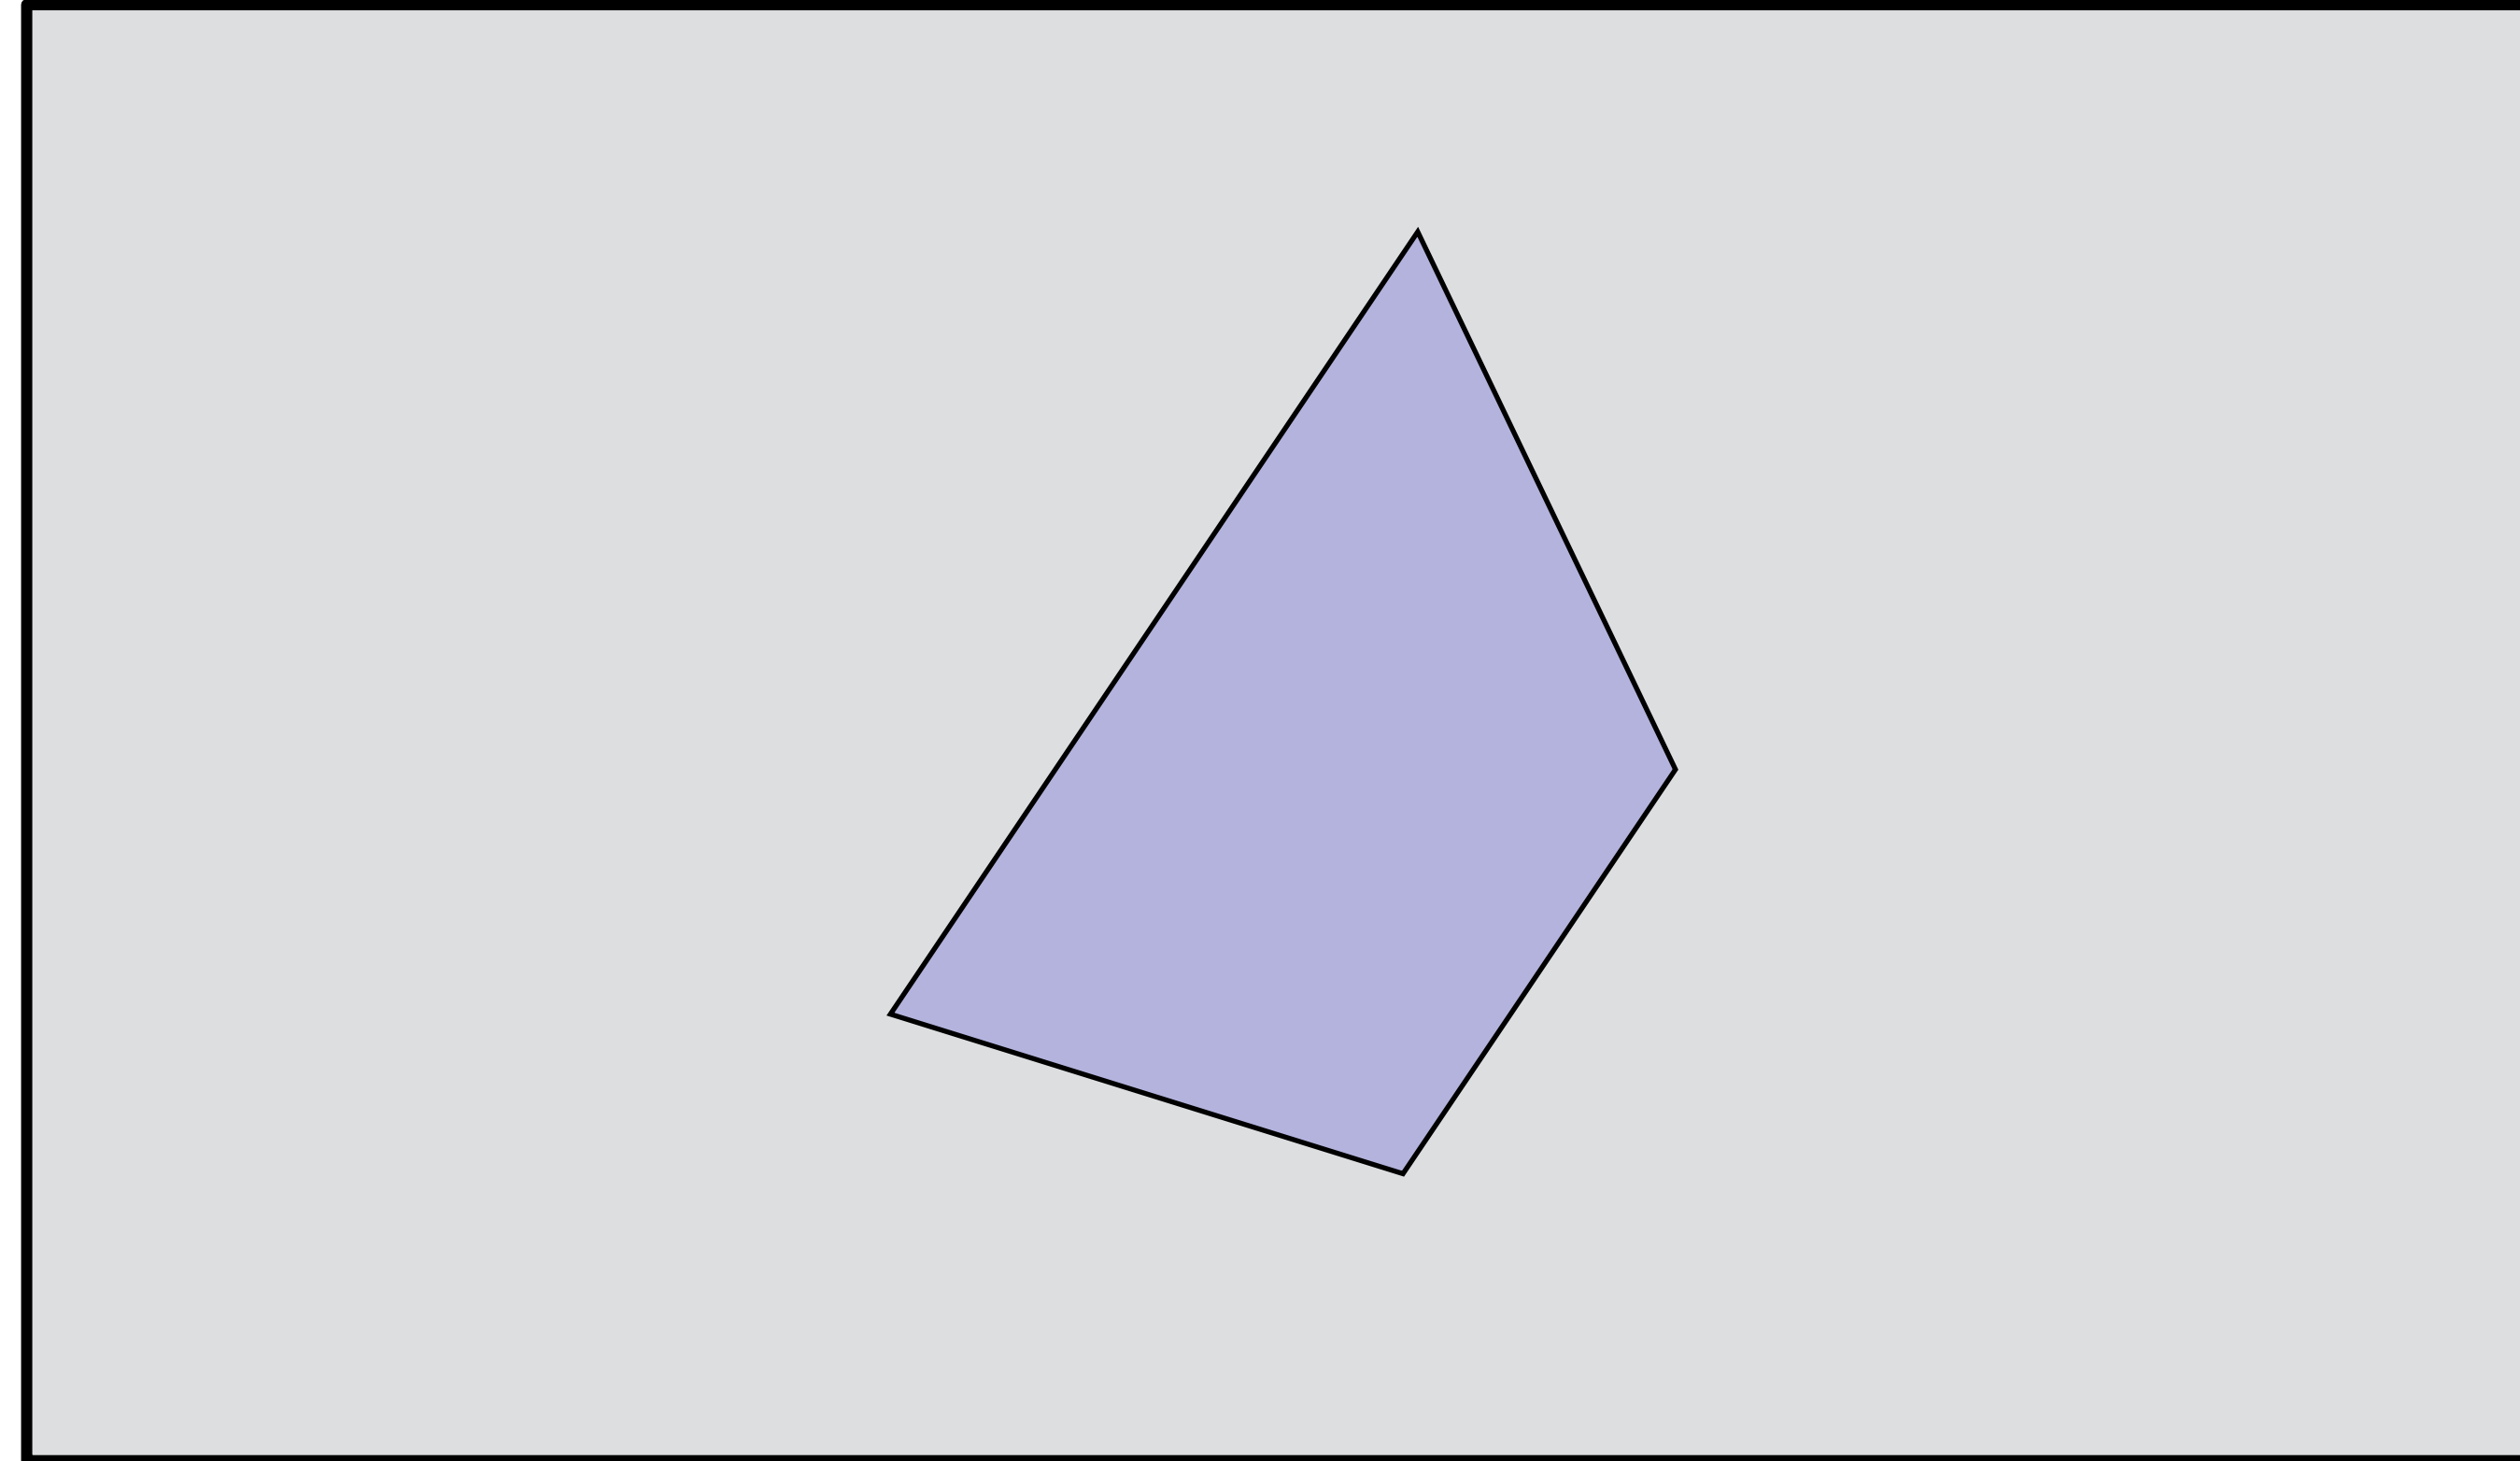
Pick a point in the image.
Find it again in the next image.

What is a **Good Feature**?



Pick a point in the image.
Find it again in the next image.

What is a **Good Feature**?



Pick a point in the image.
Find it again in the next image.

What is a **Good Feature**?

Local: features are local, robust to occlusion and clutter

Accurate: precise localization

Robust: noise, blur, compression, etc. do not have a big impact on the feature.

Distinctive: individual features can be easily matched

Efficient: close to real-time performance

What is a **corner**?



Image Credit: John Shakespeare, Sydney Morning Herald

We can think of a corner as any **locally distinct** 2D image feature that (hopefully) corresponds to a distinct position on an 3D object of interest in the scene.

What is a **corner**?

Corner

Interest Point



Image Credit: John Shakespeare, Sydney Morning Herald

We can think of a corner as any **locally distinct** 2D image feature that (hopefully) corresponds to a distinct position on an 3D object of interest in the scene.

Why are corners **distinct**?

A corner can be **localized reliably**.


Thought experiment:

Why are corners **distinct**?

A corner can be **localized reliably**.

Thought experiment:

- Place a small window over a patch of constant image value.




“**flat**” region:

Why are corners **distinct**?

A corner can be **localized reliably**.

Thought experiment:

— Place a small window over a patch of constant image value. If you slide the window in any direction, the image in the window will not change.



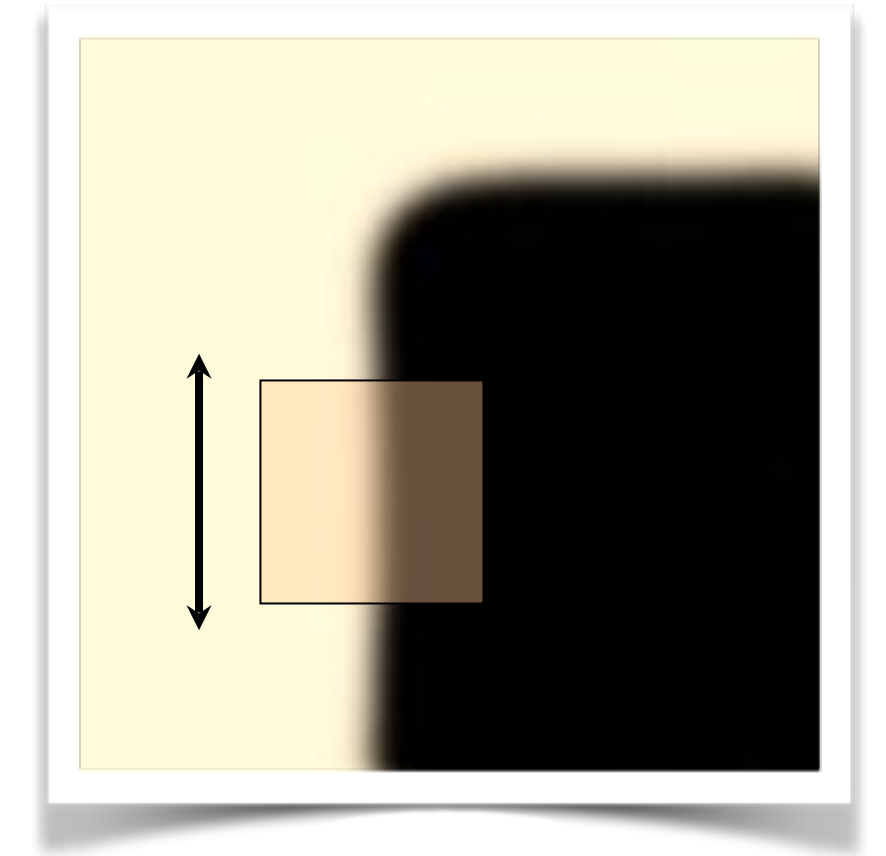
“**flat**” region:
no change in all
directions

Why are corners **distinct**?

A corner can be **localized reliably**.

Thought experiment:

- Place a small window over a patch of constant image value. If you slide the window in any direction, the image in the window will not change.
- Place a small window over an edge.



“edge”:

Why are corners **distinct**?

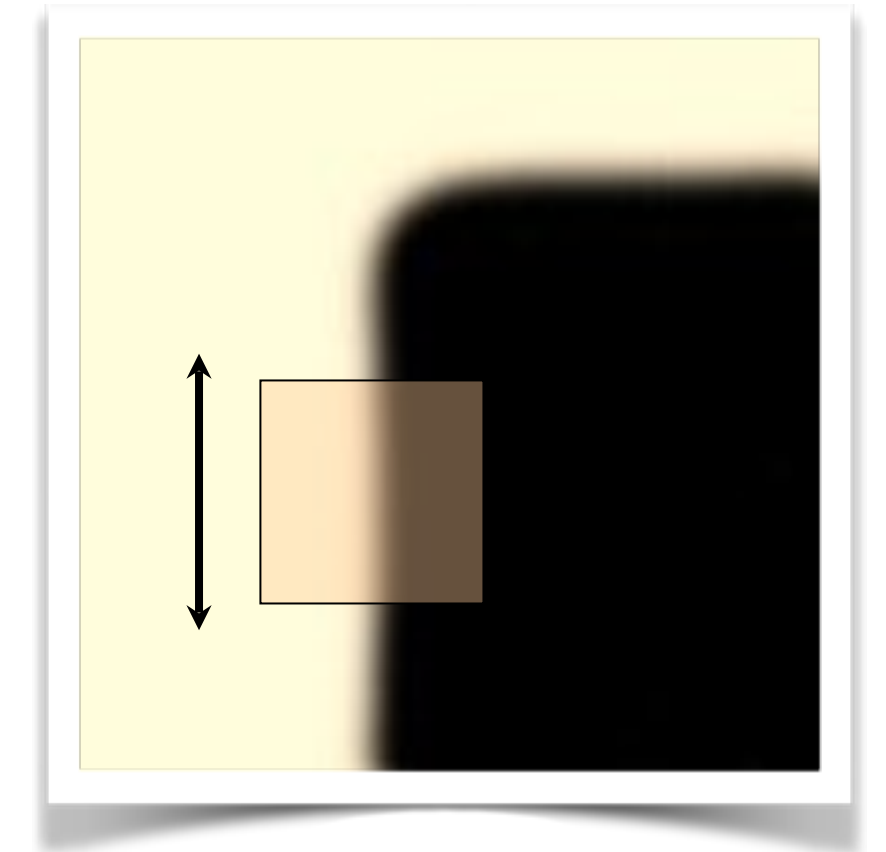
A corner can be **localized reliably**.

Thought experiment:

— Place a small window over a patch of constant image value. If you slide the window in any direction, the image in the window will not change.

— Place a small window over an edge. If you slide the window in the direction of the edge, the image in the window will not change

→ Cannot estimate location along an edge (a.k.a., **aperture** problem)




“edge”:
no change along
the edge direction

Why are corners **distinct**?

A corner can be **localized reliably**.

Thought experiment:

- Place a small window over a patch of constant image value. If you slide the window in any direction, the image in the window will not change.
- Place a small window over an edge. If you slide the window in the direction of the edge, the image in the window will not change
 - Cannot estimate location along an edge (a.k.a., **aperture** problem)
- Place a small window over a corner.



“corner”:

Why are corners **distinct**?

A corner can be **localized reliably**.


Thought experiment:

— Place a small window over a patch of constant image value. If you slide the window in any direction, the image in the window will not change.

— Place a small window over an edge. If you slide the window in the direction of the edge, the image in the window will not change

→ Cannot estimate location along an edge (a.k.a., **aperture** problem)

— Place a small window over a corner. If you slide the window in any direction, the image in the window changes.



“corner”:
significant change
in all directions

Corner Detection


Edge detectors perform poorly at corners

Observations:

- The gradient is ill defined exactly at a corner
- Near a corner, the gradient has two (or more) distinct values

How do you find a **corner**?

[Moravec 1980]



Easily recognized by looking through a small window

Shifting the window should give large change in intensity

Autocorrelation

Autocorrelation is the correlation of the image with itself.

— Windows centered on an edge point will have autocorrelation that falls off slowly in the direction along the edge and rapidly in the direction across (perpendicular to) the edge.

— Windows centered on a corner point will have autocorrelation that falls off rapidly in all directions.

Autocorrelation



Szeliski, Figure 4.5

Autocorrelation




Szeliski, Figure 4.5

Autocorrelation




Szeliski, Figure 4.5

Autocorrelation




Szeliski, Figure 4.5

Autocorrelation



Szeliski, Figure 4.5

Autocorrelation



Szeliski, Figure 4.5

Autocorrelation

Autocorrelation is the correlation of the image with itself.


— Windows centered on an edge point will have autocorrelation that falls off slowly in the direction along the edge and rapidly in the direction across (perpendicular to) the edge.

— Windows centered on a corner point will have autocorrelation that falls off rapidly in all directions.


Harris Corner Detection

1. Compute image gradients over small region
2. Compute the covariance matrix
3. Compute eigenvectors and eigenvalues
4. Use threshold on eigenvalues to detect corners

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




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



$$\begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$


1. Compute **image gradients** over a small region

(not just a single pixel)




array of x gradients

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




array of y gradients







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









Visualization of Gradients






What Does a **Distribution** Tells You About the **Region**?






What Does a **Distribution** Tells You About the **Region**?






What Does a **Distribution** Tells You About the **Region**?




Distribution reveals the **orientation** and **magnitude**




What Does a **Distribution** Tells You About the **Region**?




Distribution reveals the **orientation** and **magnitude**



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



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



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

How do we quantify the **orientation** and **magnitude**?

2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

Sum over small region
around the corner

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

Sum over small region
around the corner

Gradient with respect to x , times
gradient with respect to y

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

Sum over small region around the corner

Gradient with respect to x, times gradient with respect to y

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

$$\sum_{p \in P} I_x I_y = \text{sum} \left(\begin{matrix} I_x = \frac{\partial I}{\partial x} \\ \begin{array}{|c|c|c|c|c|} \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline \end{array} \end{matrix} \cdot \begin{matrix} I_y = \frac{\partial I}{\partial y} \\ \begin{array}{|c|c|c|c|c|} \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline & & & & \\ \hline \end{array} \end{matrix} \right)$$

array of x gradients

array of y gradients

2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

Sum over small region around the corner

Gradient with respect to x , times gradient with respect to y

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

Matrix is **symmetric**


2. Compute the **covariance matrix** (a.k.a. 2nd moment matrix)

By computing the **gradient covariance matrix** ...

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

we are fitting a **quadratic** to the gradients over a small image region


Simple Case



Local Image Patch


$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = ?$$

Simple Case




Local Image Patch

I_x




I_y




$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = ?$$

Simple Case




Local Image Patch



I_x


high value along vertical strip of pixels and 0 elsewhere



I_y


$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = ?$$

Simple Case




Local Image Patch

I_x



high value along vertical strip of pixels and 0 elsewhere


I_y



high value along horizontal strip of pixels and 0 elsewhere


$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = ?$$

Simple Case




Local Image Patch

I_x



high value along vertical strip of pixels and 0 elsewhere

I_y



high value along horizontal strip of pixels and 0 elsewhere

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

General Case

It can be shown that since every C is symmetric:



$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix} = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

... so general case is like a **rotated** version of the simple one

3. Computing **Eigenvalues** and **Eigenvectors**

Quick **Eigenvalue/Eigenvector** Review

Given a square matrix \mathbf{A} , a scalar λ is called an **eigenvalue** of \mathbf{A} if there exists a nonzero vector \mathbf{v} that satisfies

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$$

The vector \mathbf{v} is called an **eigenvector** for \mathbf{A} corresponding to the eigenvalue λ .

The eigenvalues of \mathbf{A} are obtained by solving

$$\det(\mathbf{A} - \lambda\mathbf{I}) = 0$$

3. Computing **Eigenvalues** and **Eigenvectors**

eigenvalue



$$Ce = \lambda e$$



eigenvector

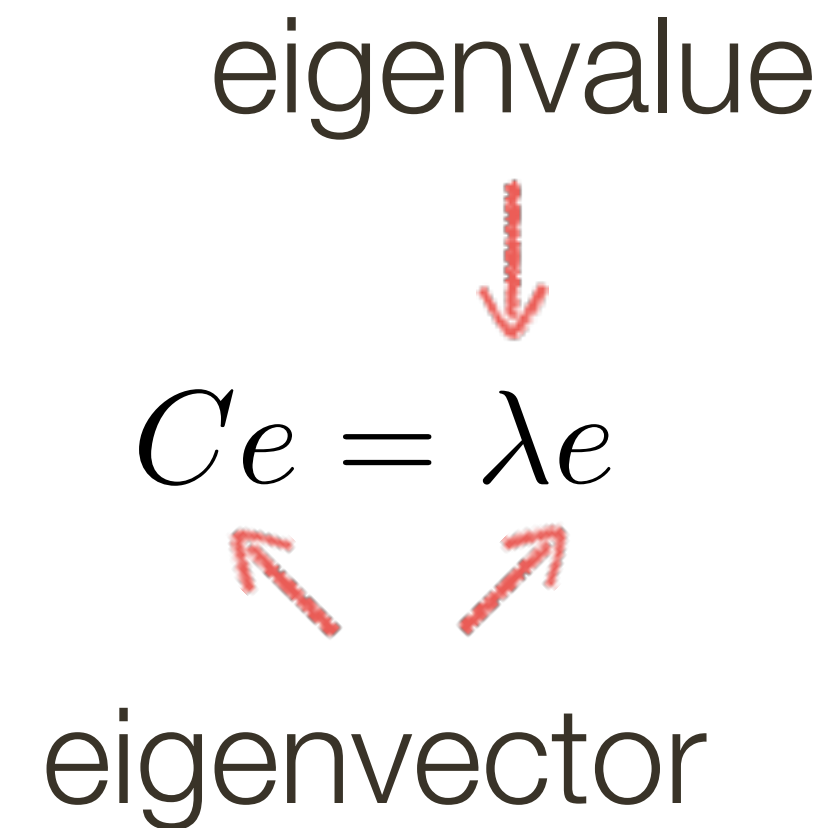
$$(C - \lambda I)e = 0$$

3. Computing **Eigenvalues** and **Eigenvectors**

eigenvalue

$$Ce = \lambda e$$

eigenvector



$$(C - \lambda I)e = 0$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

3. Computing **Eigenvalues** and **Eigenvectors**

eigenvalue

$$Ce = \lambda e$$

eigenvector

$$(C - \lambda I)e = 0$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. Computing **Eigenvalues** and **Eigenvectors**

eigenvalue

$$Ce = \lambda e$$

eigenvector

$$(C - \lambda I)e = 0$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Example

$$C = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Example

$$C = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

$$\det \left(\begin{bmatrix} 2 - \lambda & 1 \\ 1 & 2 - \lambda \end{bmatrix} \right)$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Example

$$C = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

$$\det \left(\begin{bmatrix} 2 - \lambda & 1 \\ 1 & 2 - \lambda \end{bmatrix} \right)$$

$$(2 - \lambda)(2 - \lambda) - (1)(1)$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Example

$$C = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

$$\det \left(\begin{bmatrix} 2 - \lambda & 1 \\ 1 & 2 - \lambda \end{bmatrix} \right)$$

$$(2 - \lambda)(2 - \lambda) - (1)(1)$$

$$(2 - \lambda)(2 - \lambda) - (1)(1) = 0$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Example

$$C = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

$$\det \left(\begin{bmatrix} 2 - \lambda & 1 \\ 1 & 2 - \lambda \end{bmatrix} \right)$$

$$(2 - \lambda)(2 - \lambda) - (1)(1)$$

$$(2 - \lambda)(2 - \lambda) - (1)(1) = 0$$

$$\lambda^2 - 4\lambda + 3 = 0$$

$$(\lambda - 3)(\lambda - 1) = 0$$

$$\lambda_1 = 1, \lambda_2 = 3$$

1. Compute the determinant of
(returns a polynomial)

$$C - \lambda I$$

2. Find the roots of polynomial
(returns eigenvalues)

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

3. For each eigenvalue, solve
(returns eigenvectors)

$$(C - \lambda I)e = 0$$

Visualization as **Quadratic**

$$f(x, y) = x^2 + y^2$$

can be written in matrix form like this...

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

Visualization as **Quadratic**

$$f(x, y) = x^2 + y^2$$

can be written in matrix form like this...

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

Result of Computing **Eigenvalues** and **Eigenvectors** (using SVD)

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \boxed{1} & \boxed{0} \\ \boxed{0} & \boxed{1} \end{bmatrix} \begin{bmatrix} \boxed{1} & 0 \\ 0 & \boxed{1} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}^T$$

eigenvectors eigenvalues
along diagonal

axis of the 'ellipse slice' scaling of the quadratic along the axis

Visualization as **Ellipse**

Since C is symmetric, we have
$$C = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

We can visualize C as an ellipse with axis lengths determined by the eigenvalues and orientation determined by R

Ellipse equation:

$$f(x, y) = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \text{const}$$


Visualization as **Ellipse**

Since C is symmetric, we have $C = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$


We can visualize C as an ellipse with axis lengths determined by the eigenvalues and orientation determined by R

Ellipse equation:


$$f(x, y) = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \text{const}$$




Interpreting **Eigenvalues**




Interpreting Eigenvalues



Interpreting Eigenvalues



Interpreting Eigenvalues



Interpreting **Eigenvalues**






Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

4. Threshold on Eigenvalues to Detect Corners

4. Threshold on Eigenvalues to Detect Corners


(a function of λ_1)



Think of a function to score 'corneriness'

4. Threshold on Eigenvalues to Detect Corners


(a function of)



Think of a function to score 'corneriness'

4. Threshold on Eigenvalues to Detect Corners

(a function of)




Use the **smallest eigenvalue** as the response function

$$\min(\lambda_1, \lambda_2)$$

4. Threshold on Eigenvalues to Detect Corners


(a function of)



$$\lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2$$

4. Threshold on Eigenvalues to Detect Corners

(a function of)




$$\lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2$$
$$=$$
$$\det(C) - \kappa \text{trace}^2(C)$$

(more efficient)

4. Threshold on Eigenvalues to Detect Corners

$$\det(M) - \kappa \text{trace}^2(M) < 0 \quad \text{(a function of)}$$



$$\lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2 = \det(C) - \kappa \text{trace}^2(C)$$

(more efficient)

4. Threshold on Eigenvalues to Detect Corners

(a function of)

Harris & Stephens (1988)

$$\det(C) - \kappa \text{trace}^2(C)$$

Kanade & Tomasi (1994)

$$\min(\lambda_1, \lambda_2)$$

Nobel (1998)

$$\frac{\det(C)}{\text{trace}(C) + \epsilon}$$

Harris Corner Detection Review

- Filter image with **Gaussian**
- Compute magnitude of the x and y **gradients** at each pixel
- Construct C in a window around each pixel
 - Harris uses a **Gaussian window**
- Solve for product of the λ 's
- If λ 's both are big (product reaches local maximum above threshold) then we have a corner
 - Harris also checks that ratio of λ s is not too high

Compute the **Covariance Matrix**

Sum can be implemented as an (unnormalized) box filter with

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

Harris uses a **Gaussian** weighting instead

Compute the **Covariance Matrix**

Sum can be implemented as an (unnormalized) box filter with

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

↑ Error function ↑ Window function ↑ Shifted intensity ↑ Intensity

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

Harris uses a **Gaussian** weighting instead

(has to do with bilinear Taylor expansion of 2D function that measures change of intensity for small shifts ... remember AutoCorrelation)

Harris Corner Detection Review

- Filter image with **Gaussian**
- Compute magnitude of the x and y **gradients** at each pixel
- Construct C in a window around each pixel
 - Harris uses a **Gaussian window**
- Solve for product of the λ 's
- If λ 's both are big (product reaches local maximum above threshold) then we have a corner
 - Harris also checks that ratio of λ s is not too high

Harris & Stephens (1988)

$$\det(C) - \kappa \text{trace}^2(C)$$

Example: Harris Corner Detection

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

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

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Example: Harris Corner Detection

Lets compute a measure of "corner-ness" for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

$$\sum \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 1 & 0 \end{bmatrix} \odot \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 1 & 0 \end{bmatrix} = 3$$

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

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

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

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

$$\mathbf{C} = \begin{bmatrix} 3 & 2 \\ 2 & 4 \end{bmatrix}$$

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

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

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

$$C = \begin{bmatrix} 3 & 2 \\ 2 & 4 \end{bmatrix} \Rightarrow \lambda_1 = 1.4384; \lambda_2 = 5.5616$$

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

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

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

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

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

$$\mathbf{C} = \begin{bmatrix} 3 & 2 \\ 2 & 4 \end{bmatrix} \Rightarrow \lambda_1 = 1.4384; \lambda_2 = 5.5616$$

$$\det(\mathbf{C}) - 0.04\text{trace}^2(\mathbf{C}) = 6.04$$

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

$$\mathbf{C} = \begin{bmatrix} 3 & 0 \\ 0 & 0 \end{bmatrix} \Rightarrow \lambda_1 = 3; \lambda_2 = 0$$

$$\det(\mathbf{C}) - 0.04\text{trace}^2(\mathbf{C}) = -0.36$$

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

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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

0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

Example: Harris Corner Detection

Lets compute a measure of “corner-ness” for the green pixel:

0	0	0	0	0	0	0
0	1	0	0	0	1	0
0	1	1	1	1	0	0
0	1	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0
0	0	1	1	1	0	0

$$\mathbf{C} = \begin{bmatrix} 3 & 0 \\ 0 & 2 \end{bmatrix} \Rightarrow \lambda_1 = 3; \lambda_2 = 2$$

$$\det(\mathbf{C}) - 0.04\text{trace}^2(\mathbf{C}) = 5$$

0	0	0	0	0	0	
-1	1	0	0	-1	1	
-1	0	0	0	1	0	
-1	0	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	
0	-1	0	0	1	0	

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


0	-1	0	0	0	-1	0
0	0	-1	-1	-1	1	0
0	0	0	0	0	0	0
0	1	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

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

Harris Corner Detection Review

- Filter image with **Gaussian**
- Compute magnitude of the x and y **gradients** at each pixel
- Construct C in a window around each pixel
 - Harris uses a **Gaussian window**
- Solve for product of the λ 's
- If λ 's both are big (product reaches local maximum above threshold) then we have a corner
 - Harris also checks that ratio of λ s is not too high

Properties: Rotational Invariance





Ellipse rotates but its shape
(**eigenvalues**) remains the same

Corner response is **invariant** to image rotation

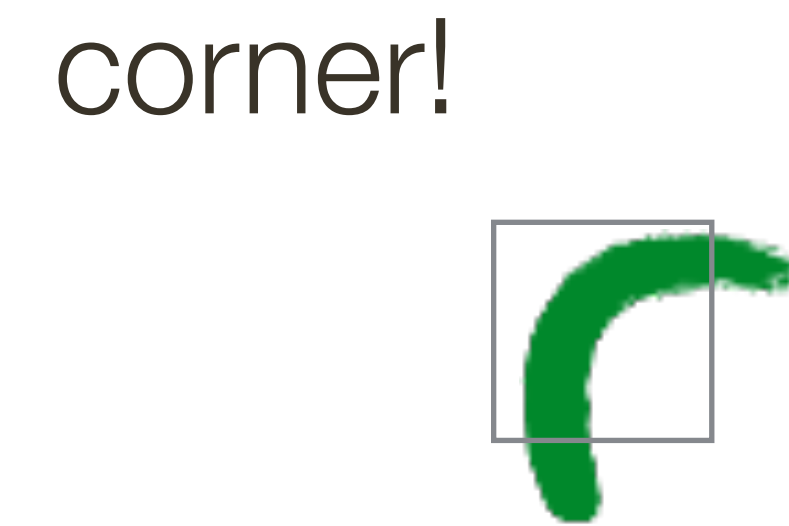

Properties: (partial) Invariance to Intensity Shifts and Scaling

Only derivatives are used -> Invariance to intensity shifts



Intensity scale could effect performance



Properties: NOT Invariant to Scale Changes



Example 1:



Example 2: Wagon Wheel (Harris Results)



$\sigma = 1$ (219 points)



$\sigma = 2$ (155 points)




$\sigma = 3$ (110 points)




$\sigma = 4$ (87 points)

Example 3: Crash Test Dummy (Harris Result)



corner response image



$\sigma = 1$ (175 points)

Original Image Credit: John Shakespeare, Sydney Morning Herald

Summary Table

Summary of what we have seen so far:

Representation	Result is...	Approach	Technique
intensity	dense	template matching	(normalized) correlation
edge	relatively sparse	derivatives	$\nabla^2 G$, Canny
corner	sparse	locally distinct features	Harris