CSCI 497P/597P: Computer Vision

Lecture 9:

Image Features: Harris Corner Detection


9300
ONE
HARRIS CORNERS


9300
ONE
HARRIS CORNERS

Announcements

- How's P1 going?

Goals

- Understand the derivation of the Harris corner detector.
- Know how to efficiently compute the Harris score at every pixel in an image.
- Understand the concept of invariance as pertains to the Harris corner detectors

Conceptually:

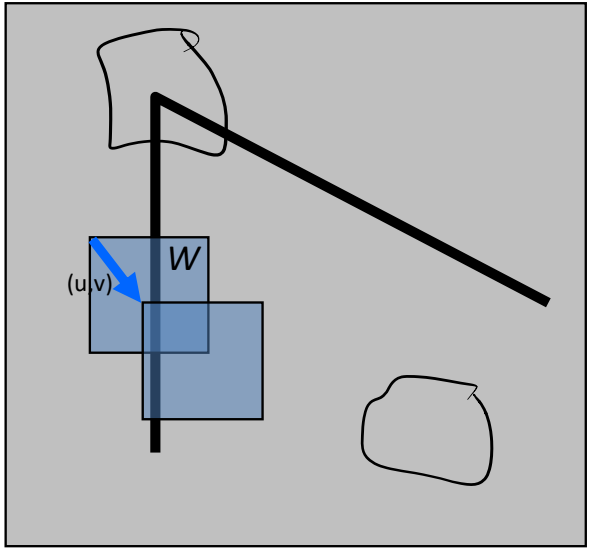

if you **nudge** the patch, how much does its appearance change?

Not at all in any direction?

A lot in some directions?

A lot in **all** directions?

If the **smallest** E(u,v) is large, **then** W is a corner.

Corner score = *minimum* E(u,v)

Mathily:

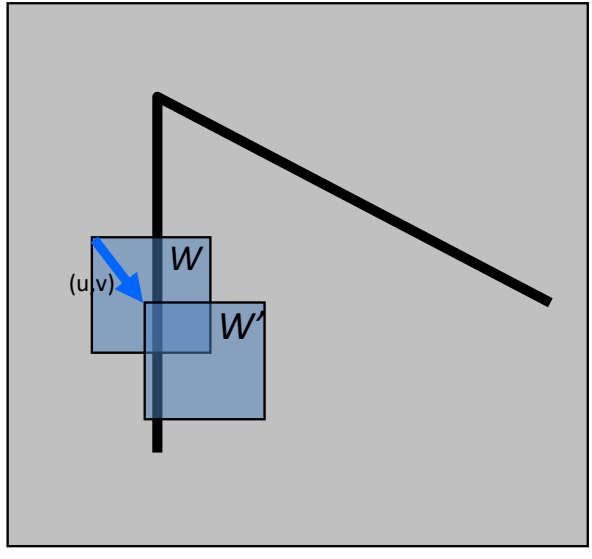

Start here: how different are W and W'?

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

Efficiently:

assume the image function is **locally linear**.

$$I(x+u, y+v) \approx I(x,y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \text{H.O.T}$$

Locally Linear Approximation

- E(u, v) tells us how much the window changes in the direction (u, v).
  - Because of our approximations, this is a **linear** model in intensity, and a **quadratic model** in the error function.

1D intuition:

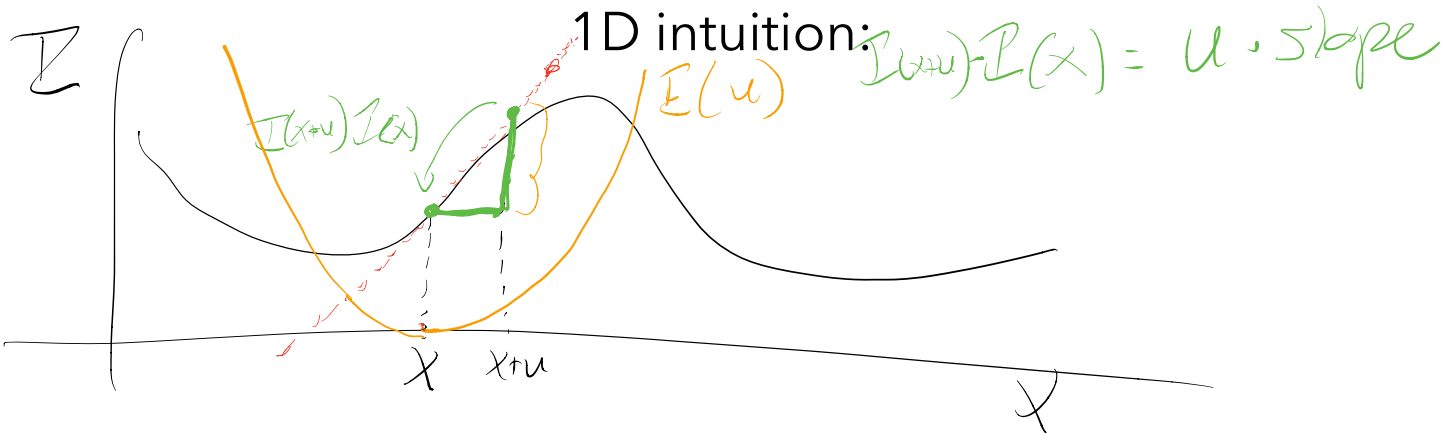

What does this mean?

- E(u, v) tells us how much the window changes in the direction (u, v).
  - Because of our approximations, this is a **linear** model in intensity, and a **quadratic model** in the error function.

(Demo)

Plug the approximation into E...

...and do some good ol' algebraic manipulation

Error function:

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

Linear image approximation:

$$I(x+u, y+v) \approx I(x,y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \text{H.O.T}$$

$$E(u,v) \approx \sum_{x,y\in W} \left[I(x,y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}\cdot v - I(x,y) \right]^2$$

$$\sum_{x,y \in W} \left(I_x \, u + I_y \cdot v \right)^2$$

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

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

$$\sum_{x,y \in W} \left(I_x^2 u^2 + 2 I_x I_y u v + I_y^2 v^2 \right)$$

$$E(u,v) = \sum_{x,y \in W} \begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

...rewrite in matrix form:

$$E(u,v) = (u, v) \begin{pmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}$$

What is this matrix thing

The four numbers that determine the shape of our (approximate) quadratic error surface.

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

$$A = \sum_{(x,y)\in W} I_x^2 \quad 0$$

$$B = \sum_{(x,y)\in W} I_x I_y \quad 0$$

$$C = \sum_{(x,y)\in W} I_y^2 \neq 0$$

One scenario:

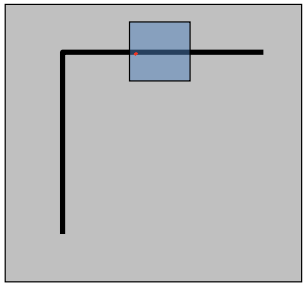

$$H = \begin{bmatrix} 0 & 0 \\ 0 & C \end{bmatrix}$$

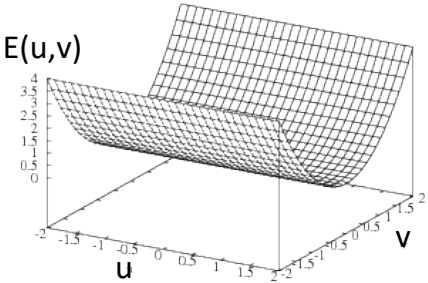

structure tensor

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

$$A = \sum_{(x,y)\in W} I_x^2 \neq 0$$

$$B = \sum_{(x,y)\in W} I_x I_y \; 0$$

$$C = \sum_{(x,y)\in W} I_y^2 \; 0$$

One scenario:

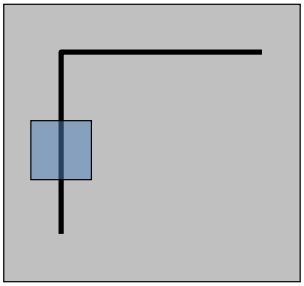

$$H = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$$

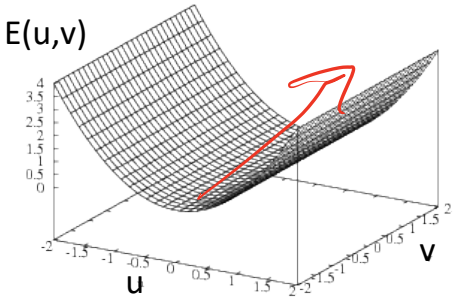

In general...

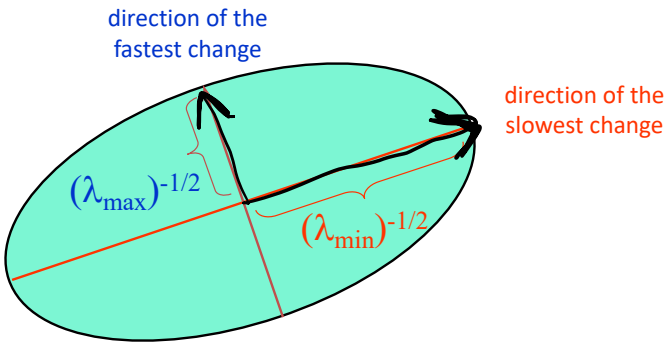
direction of the
fastest change
direction of the
slowest change
$(\lambda_{max})^{-1/2}$
$(\lambda_{min})^{-1/2}$

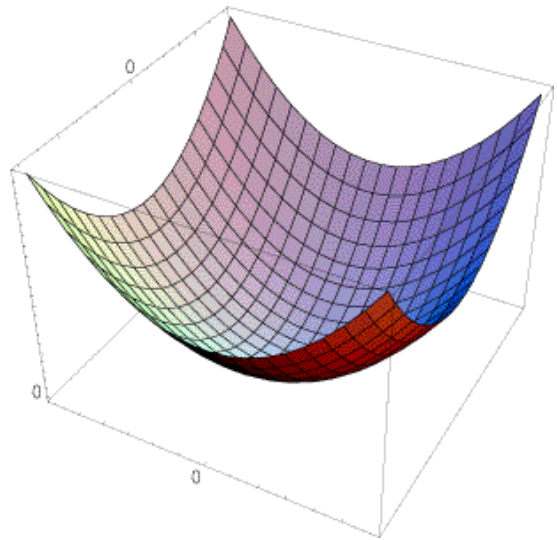
0
0
0

Eigenvectors!?

scalar

$$Ax = \lambda x$$

$$(x, \lambda)$$

eigvec eigval

$$\min_{x} \|x^T A x\|$$

Corner Detection: Upshot

- The **smaller** eigenvalue of H is **large** when the patch is centered on a corner.

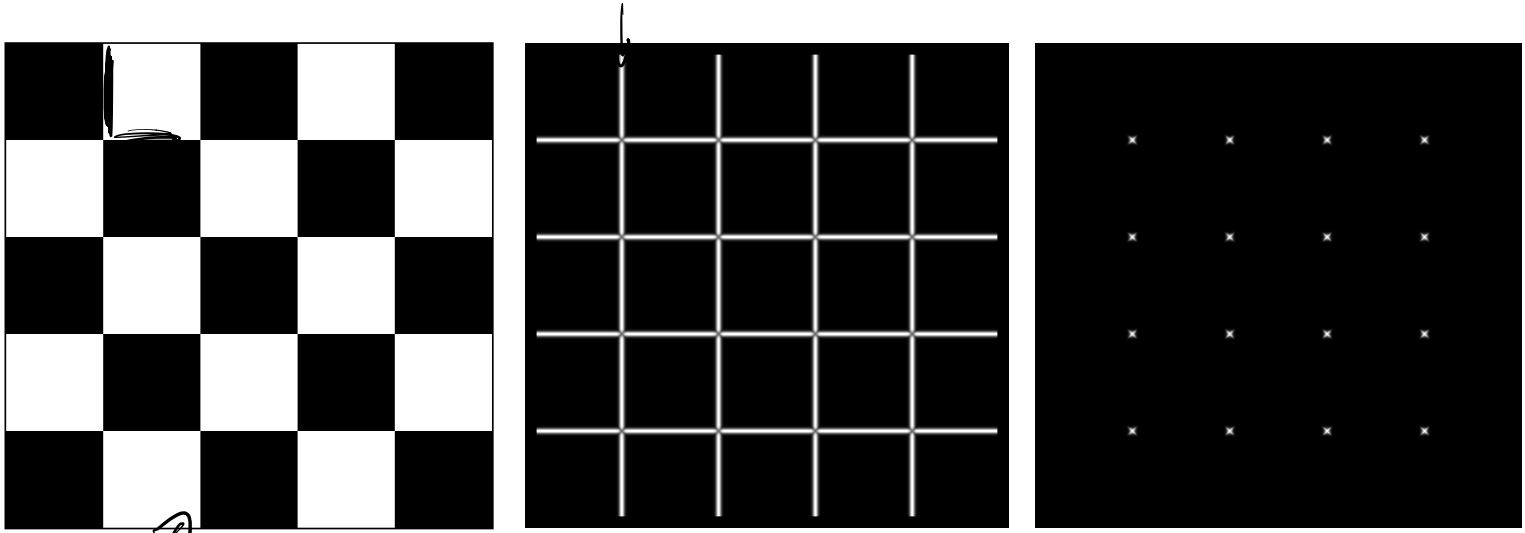

I $\lambda_{\max}$ $\lambda_{\min}$

Input

Smallest eigenvalue

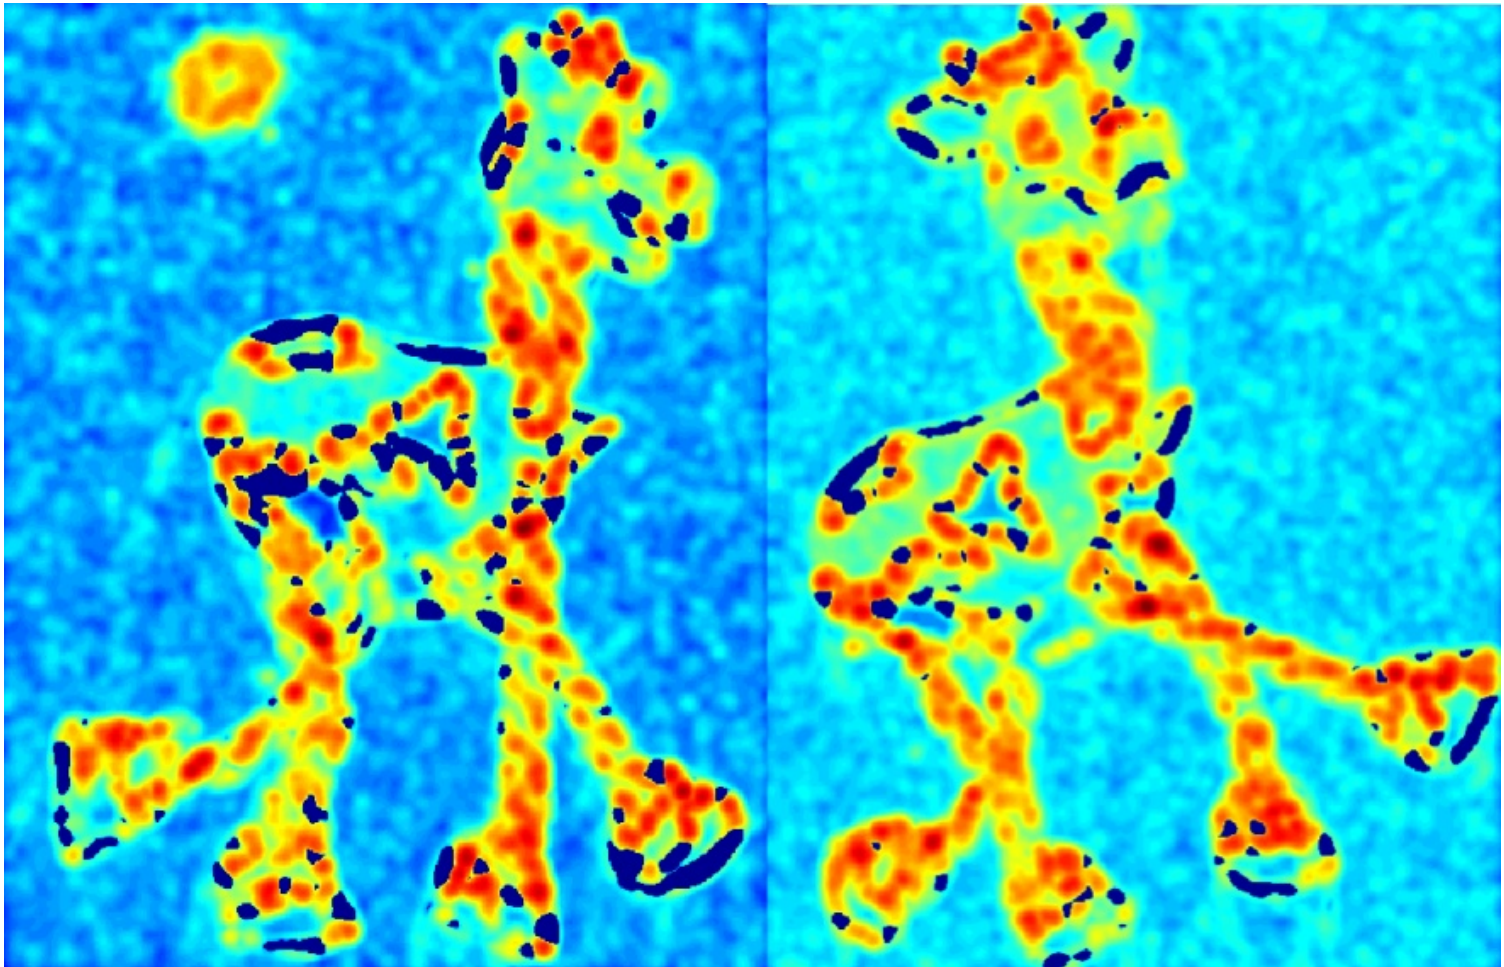

Thresholded

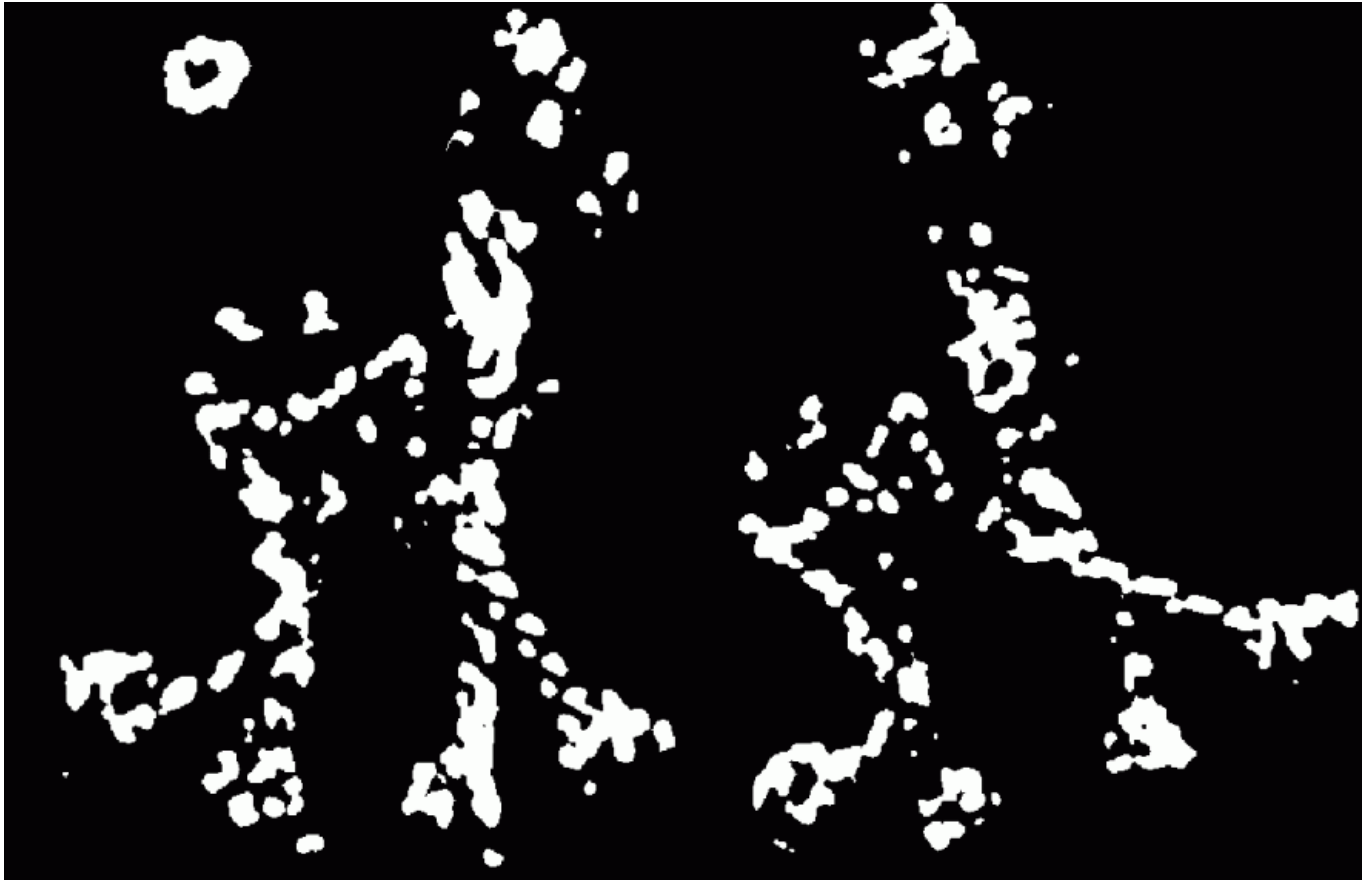

Keep only Local Maxima

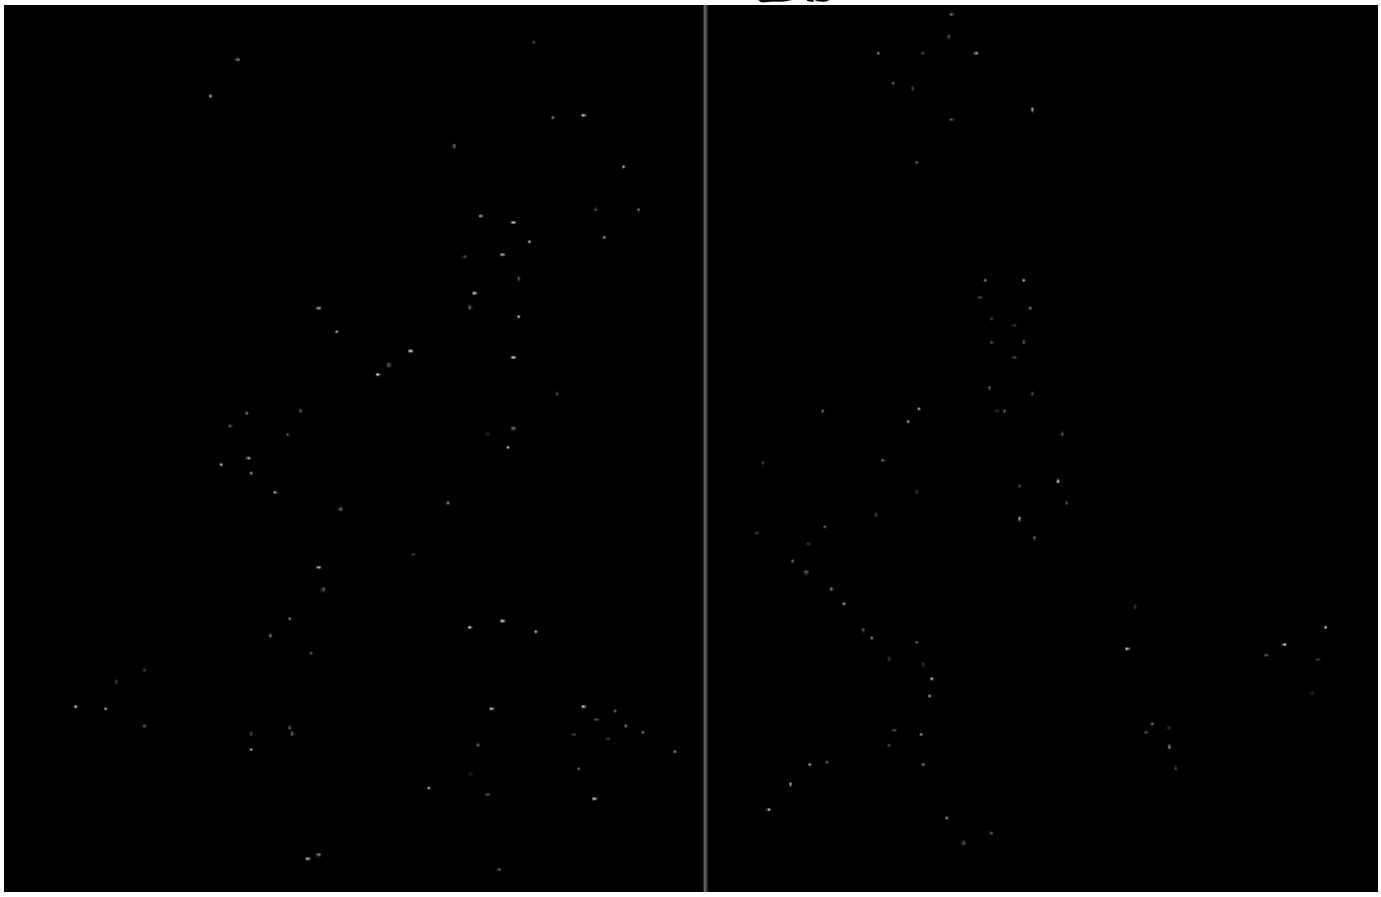

Resulting Corners

Too Long; Didn't Math

Harris Corners: TL;DM

- **Goal**: Find "unique" patches.
  Proxy for uniqueness: not like neighboring patches

- **Approach**: Eigenanalysis of SSD error on locally-linearized image windows.

  math math math

- **Upshot**:
  The smaller eigenvalue of this 2x2 matrix indicates cornerishness:

$$H = \begin{bmatrix} \sum_{(x,y)\in W} I_x^2 & \sum_{(x,y)\in W} I_x I_y \\ \sum_{(x,y)\in W} I_x I_y & \sum_{(x,y)\in W} I_y^2 \end{bmatrix}$$

Harris Corners: TL;DM

input:

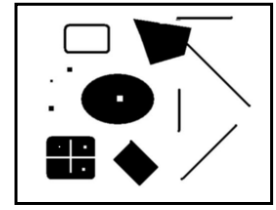

Algorithm:

1. Use Sobel filter to estimate gradients I_x, I_y

2. Compute I_x^2, I_y^2, I_xI_y

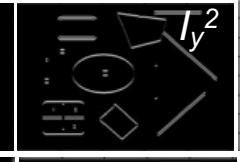

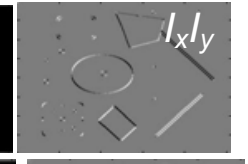

3. Filter each with a K x K mean filter

$$\begin{bmatrix} A & B \\ B & C \end{bmatrix}$$

$$\sum_{x,y \in W} I_x^2$$

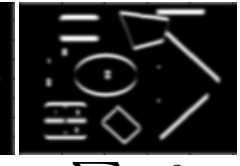

$$\sum_{x,y \in W} I_y^2$$

$$\sum_{x,y \in W} I_x I_y$$

4. Approximate smallest eigenvalue as:
   det(H) / tr(H)

det(H) = A*C - B*B
tr(H) = A+C

Corner Detection: Upshot

- The **smaller** eigenvalue of H is **large** when the patch is centered on a corner.

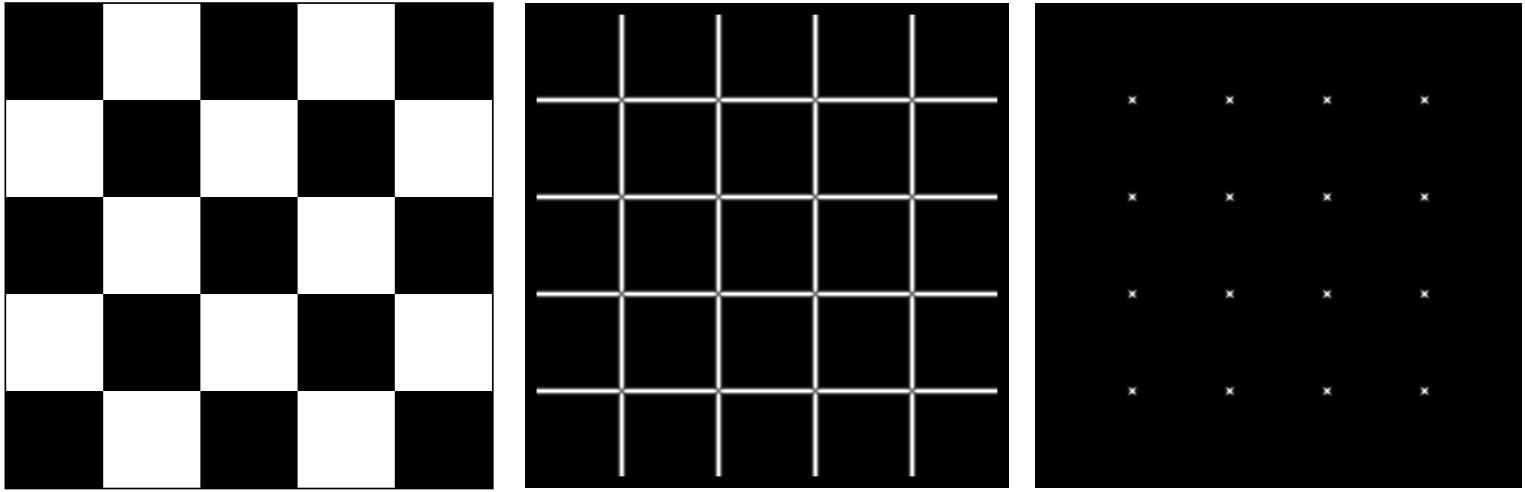

I $\lambda_{\max}$ $\lambda_{\min}$

Harris Corners: TL;DM

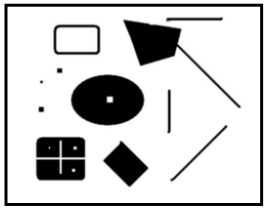

Algorithm:

1. Use Sobel filter to estimate gradients I_x, I_y
2. Compute I_x^2, I_y^2, I_xI_y
3. Filter each with a K x K mean filter
4. Approximate smallest eigenvalue as:
   det(H) / tr(H) Harris score
5. Threshold
6. Maximum filter

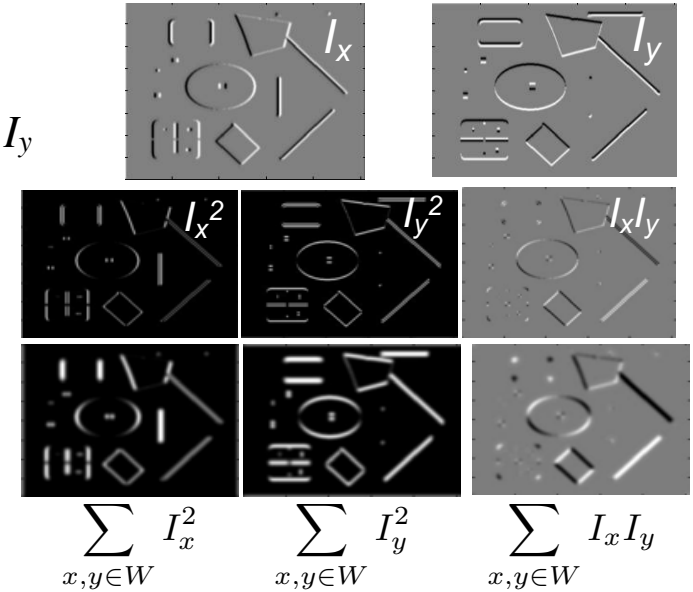

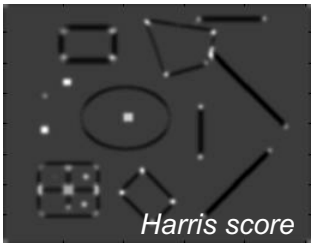

Two desirable properties:

- Uniqueness: features **shouldn't** match if they're from different points in the scene.

- Invariance: features **should** match if they do come from the same point in the scene.

Invariance

Invariance: Hard mode

Invariance: Mars Mode

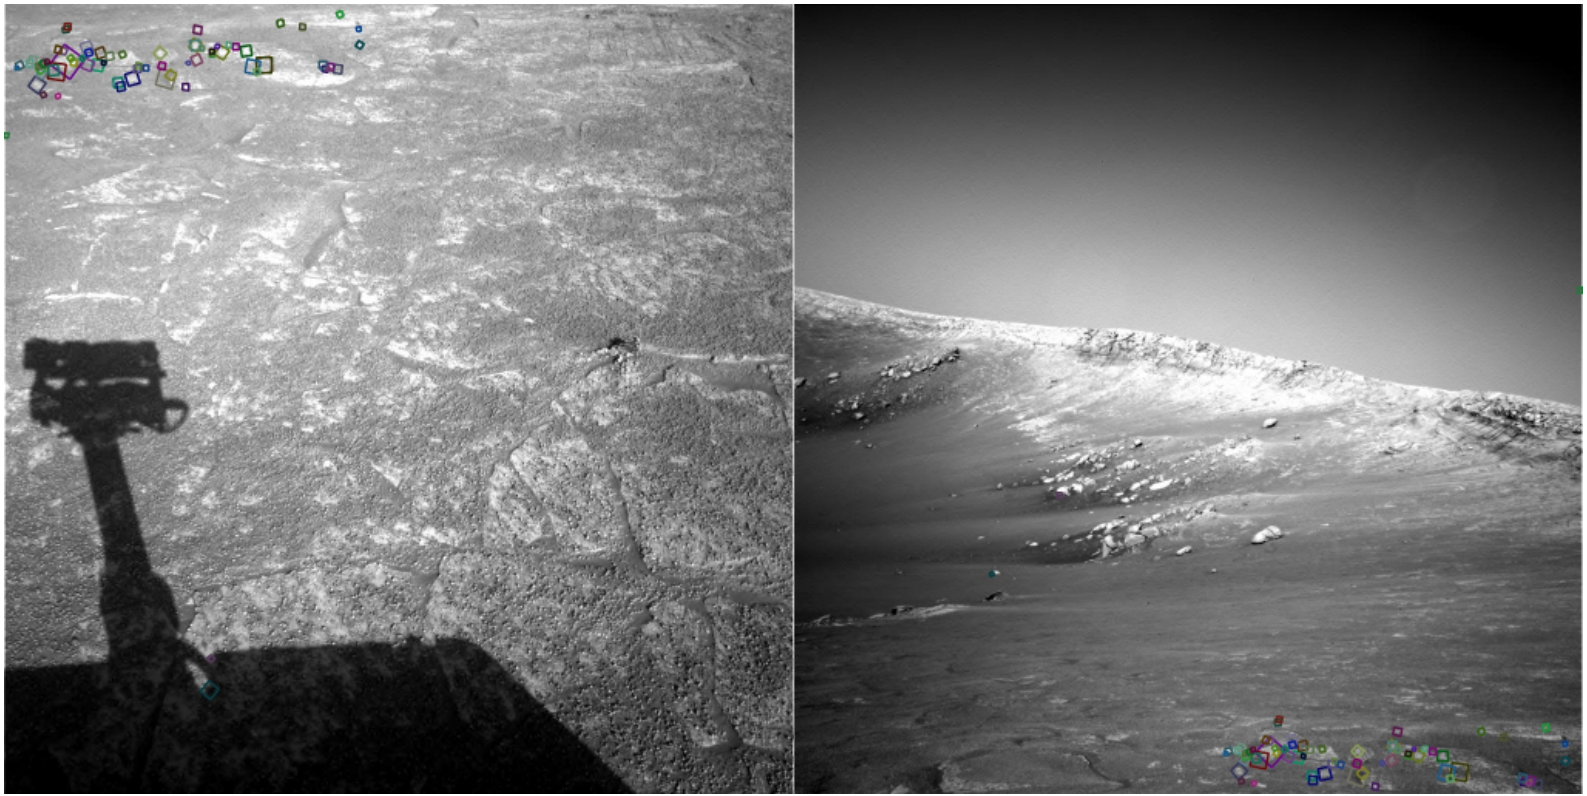

Invariance

- Suppose we're comparing two images of the same scene. What kinds of transformations could relate the two images if:
  - They are part of a panorama sequence?
  - They were taken at different times of day?
  - They were taken by different cameras?
  - They were taken from different viewpoints?