Keypoint Detection: Harris Operator

Guido Gerig

CS 6320 Spring 2012

Credits: Ross Whitaker, Trevor Darrell, Lana Lazebnik, K. Grauman, B. Leibe, A. Efros

Keypoint Detection

- Where will the interest points come from?
  - What are salient features that we'll *detect* in multiple views?
- How to *describe* a local region?
- How to establish *correspondences*, i.e., compute matches?

Credit: Prof. Trevor Darrell, C280 Computer Vision, Lecture 6
trevor@eecs.berkeley.edu

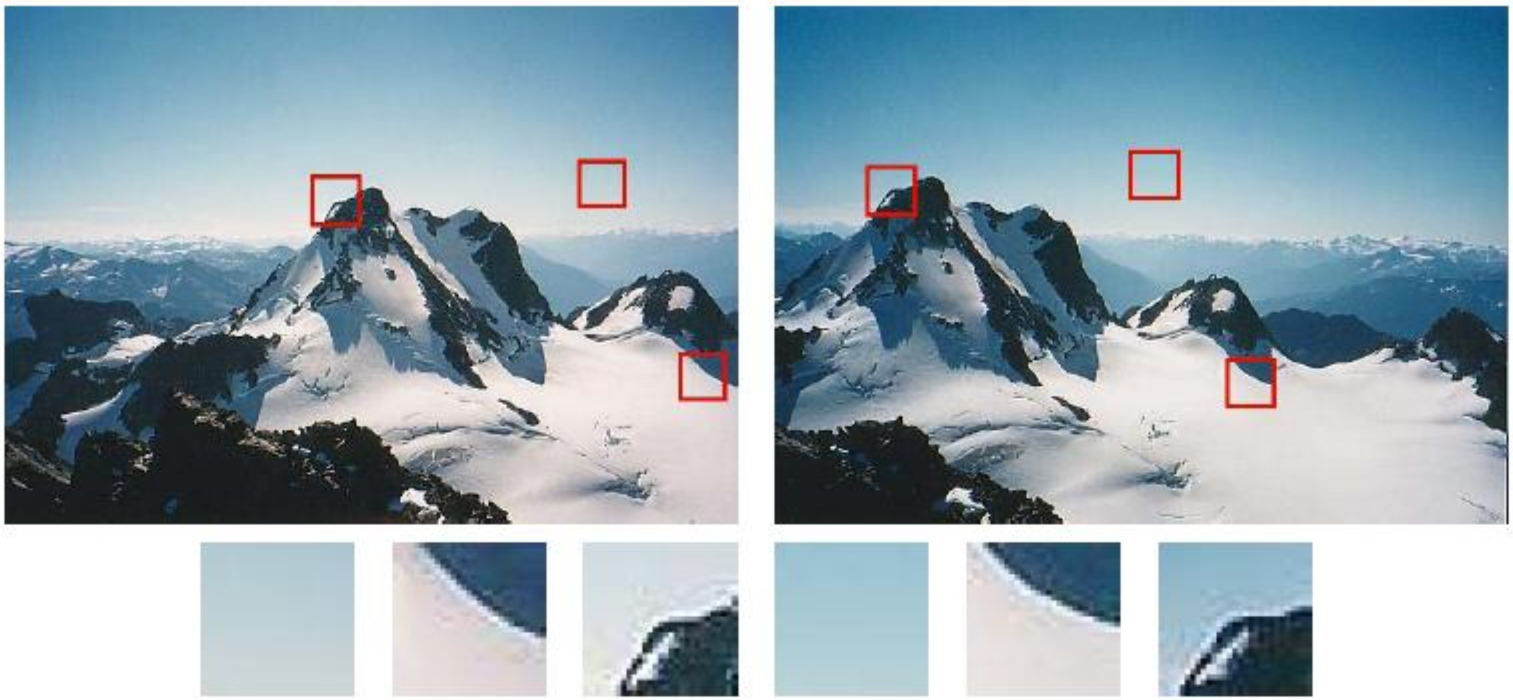

Figure 4.3: *Image pairs with extracted patches below. Notice how some patches can be localized or matched with higher accuracy than others.*

Finding Corners

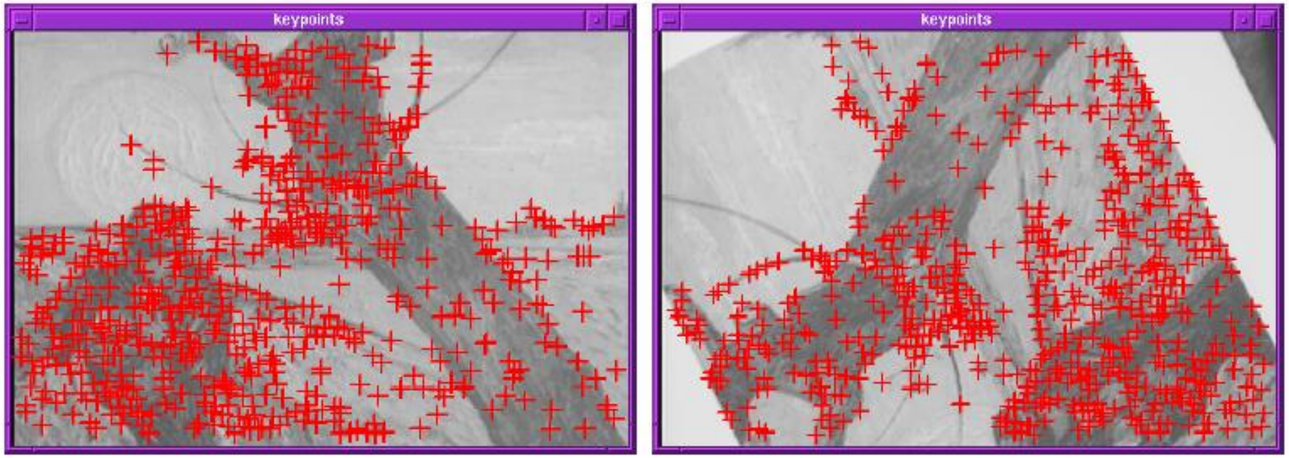

Key property: in the region around a corner, image gradient has two or more dominant directions

Corners are repeatable and **distinctive**

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

Source: Lana Lazebnik

Corners as distinctive interest points

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

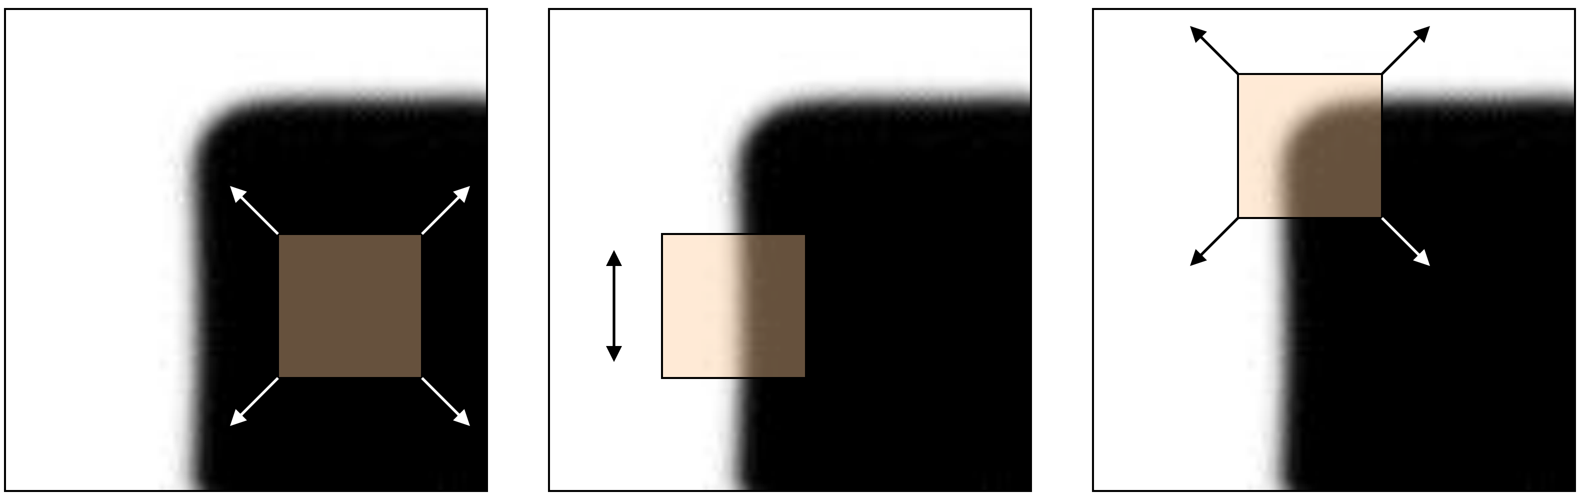

"flat" region:
no change in all directions

"edge":
no change along the edge direction

"corner":
significant change in all directions

Source: A. Efros

Local neighborhoods of 3 distinctive local patterns: Strong point features (points, corners) represent good landmarks.

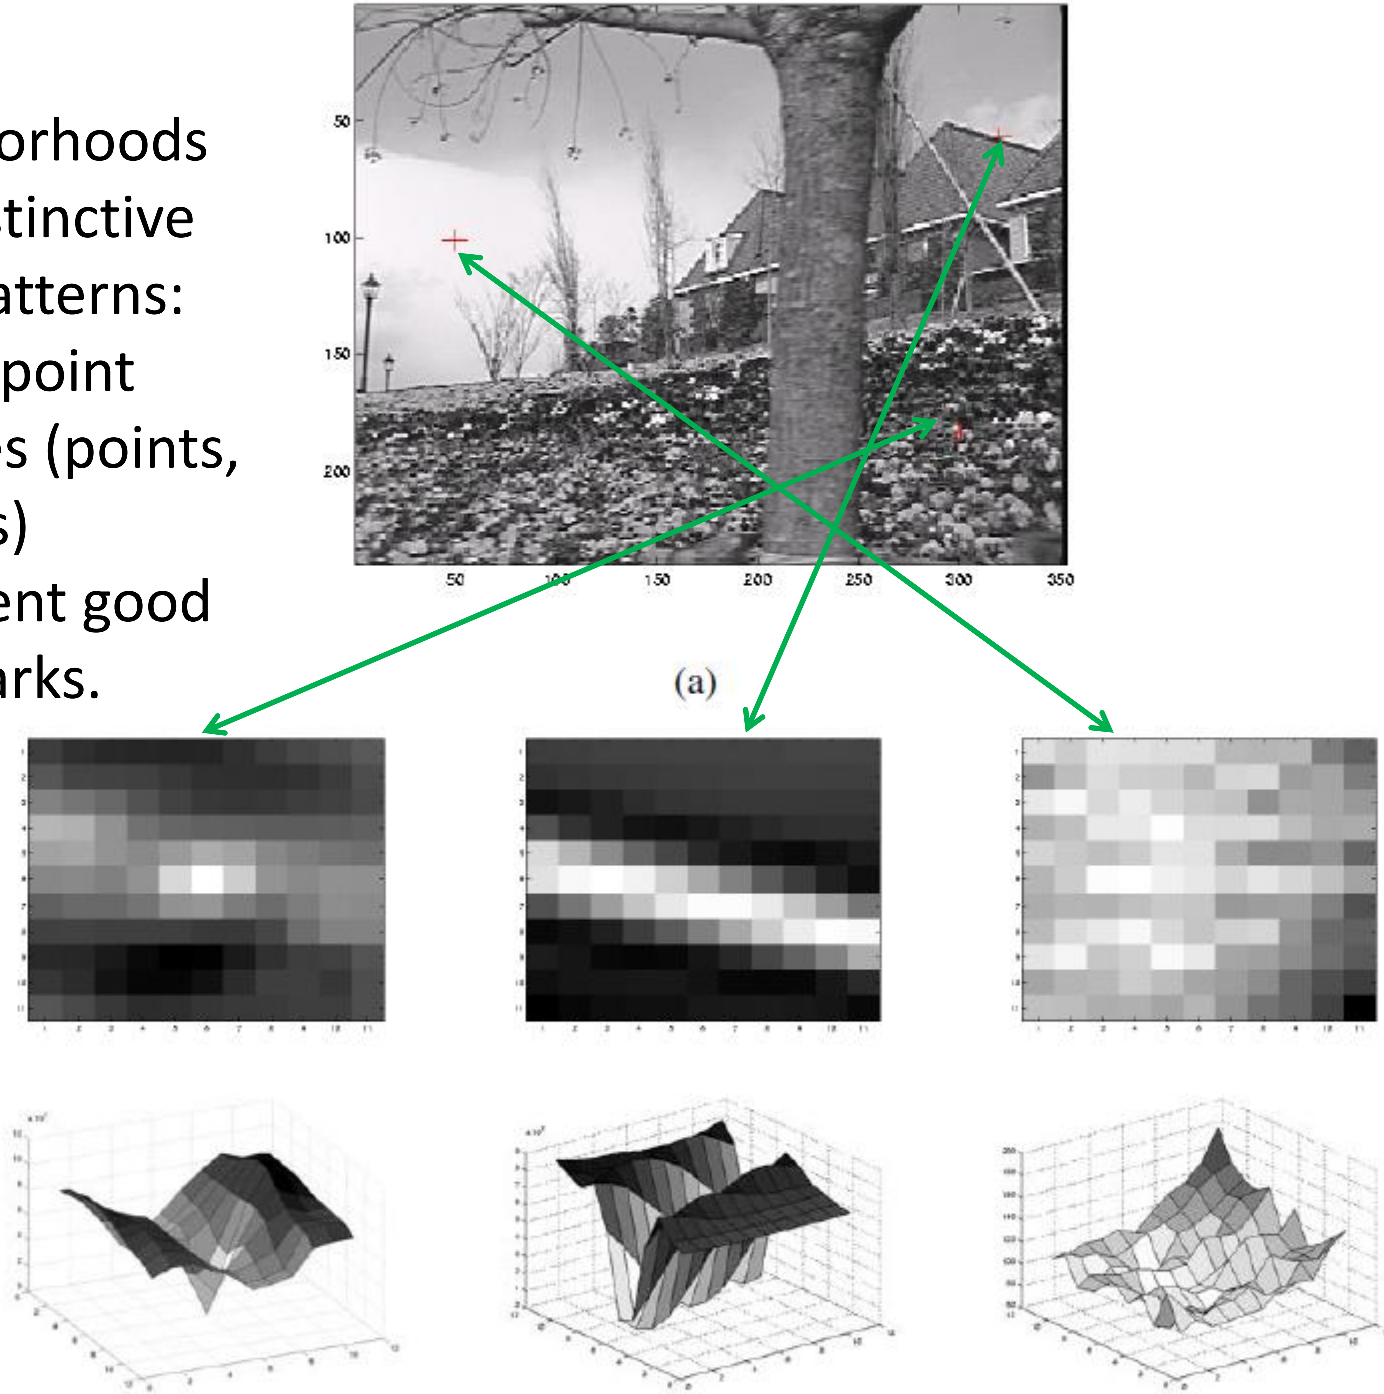

(a)

Harris Detector formulation

Change of intensity for the shift [*u,v*]:

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

Window function — Shifted intensity — Intensity

Window function $w(x,y)$ =

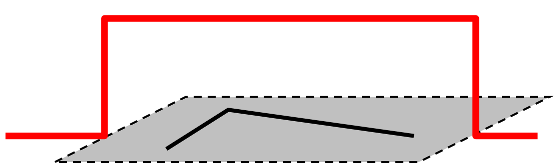

1 in window, 0 outside

or

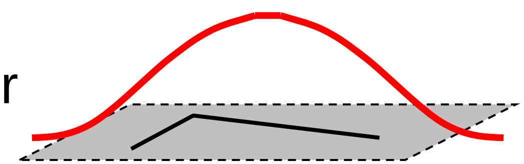

Gaussian

Source: R. Szeliski

Harris Detector formulation

This measure of change can be approximated by:

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

where *M* is a 2×2 matrix computed from image derivatives:

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

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

Sum over image region – area we are checking for corner

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

Harris Detector formulation

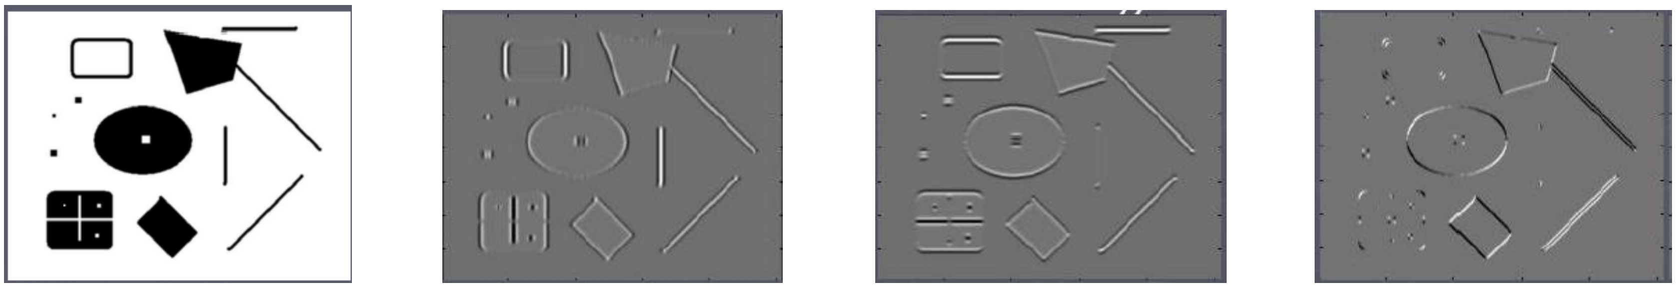

where M is a 2×2 matrix computed from image derivatives:

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

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

Sum over image region – area we are checking for corner

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

What does this matrix reveal?

First, consider an axis-aligned corner:

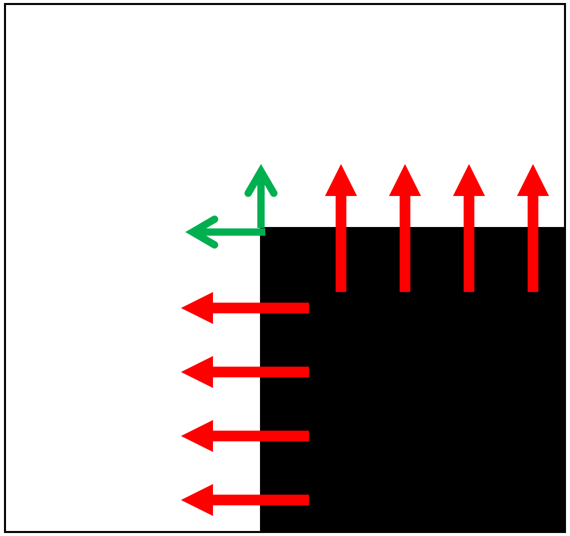

What does this matrix reveal?

First, consider an axis-aligned corner:

$$M = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

This means dominant gradient directions align with x or y axis

If either λ is close to 0, then this is **not** a corner, so look for locations where both are large.

What if we have a corner that is not aligned with the image axes?

Slide credit: David Jacobs

Arbitrary rotation

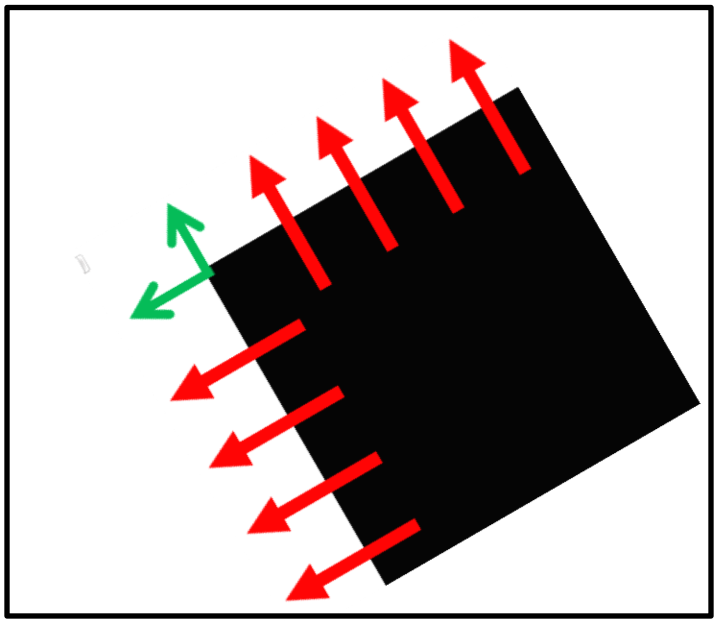

$$M = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix}$$

- M arbitrary, positive semidefinite
- Edges and corner no more aligned with image axes

General Case

Since M is symmetric, we have $M = R^{-1}\begin{bmatrix}\lambda_1 & 0 \\ 0 & \lambda_2\end{bmatrix}R$
(eigenvalue-eigenvector transformation)

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

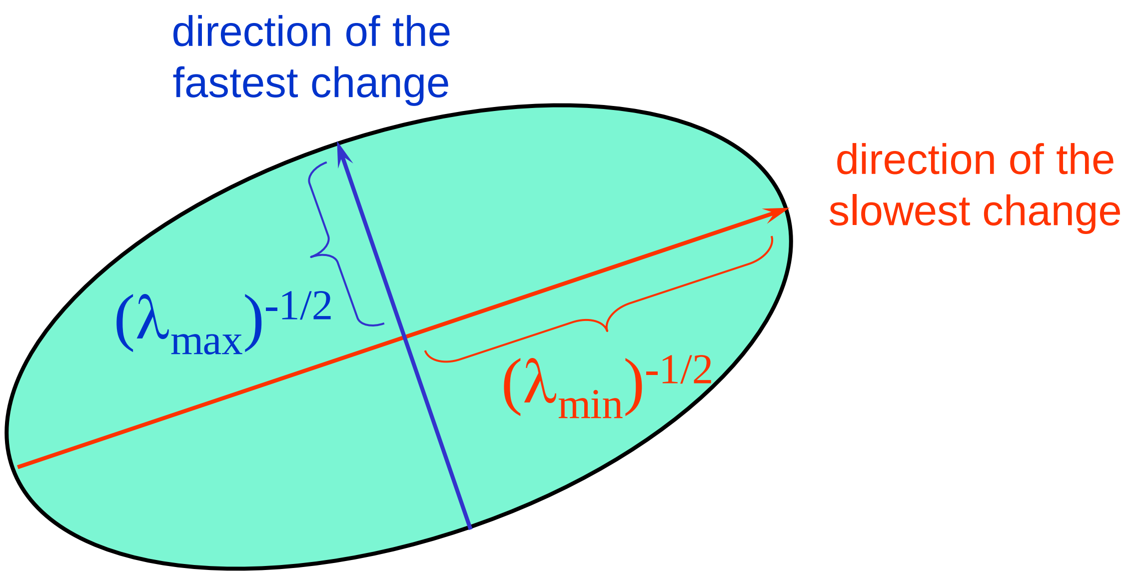

Slide adapted form Darya Frolova, Denis Simakov.

Interpreting the eigenvalues

Classification of image points using eigenvalues of *M*:

Edge (strong response across edge, small response along edge):
$\lambda_1 >> \lambda_2$ or $\lambda_2 >> \lambda_1$

Corner (strong responses in both directions):
λ_1 and $\lambda_2 >> 0$

Flat (small or no response in both directions):
λ_1 and λ_2 small

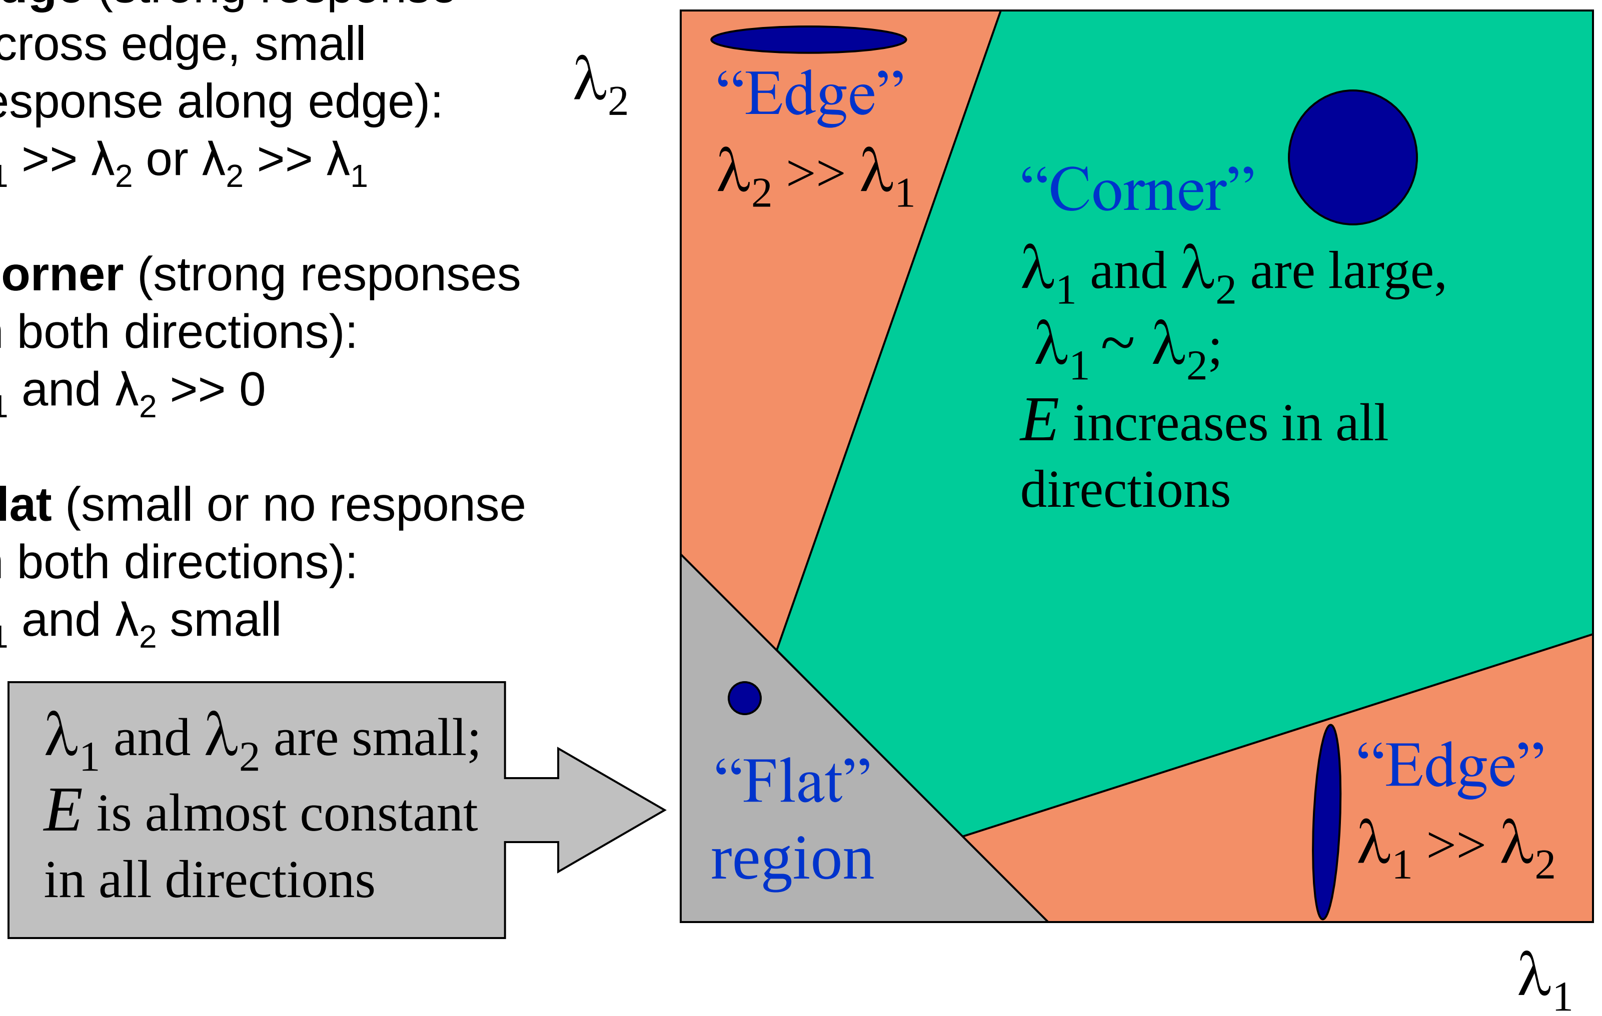

Corner response function: Harris

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

α: constant (0.04 to 0.06)

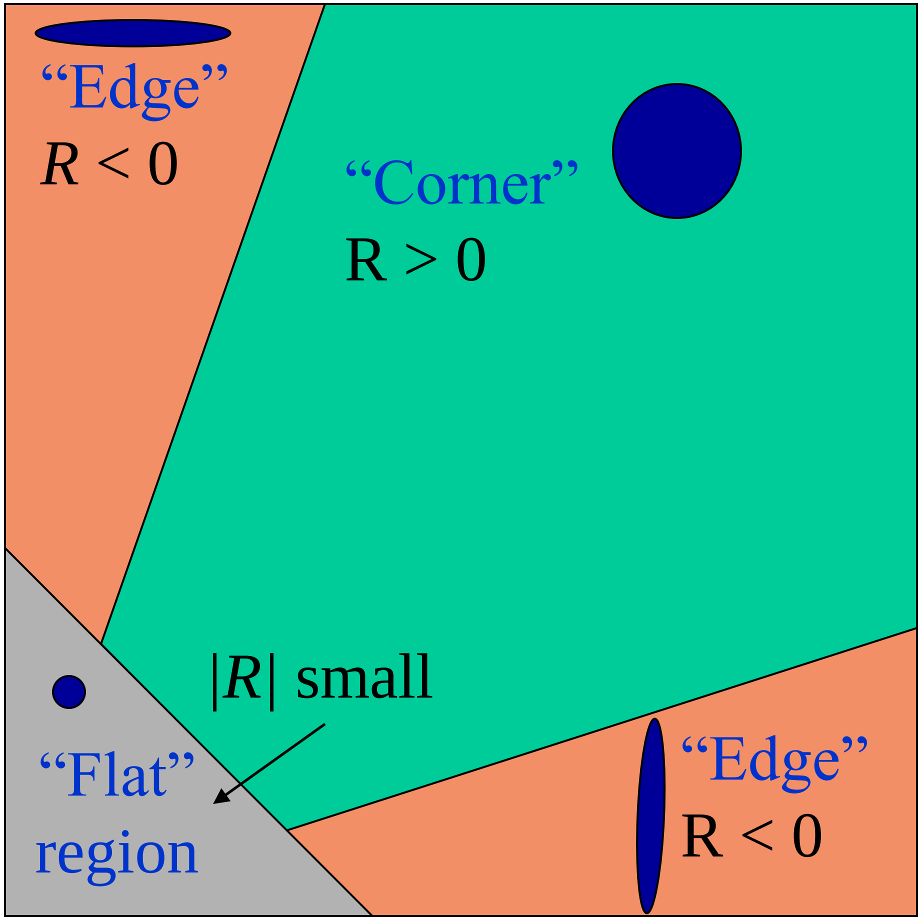

Harris Corner Detector

- Algorithm steps:
  - Compute M matrix within all image windows to get their Response scores
  - Find points with large corner response

    (*Response* > threshold)
  - Take the points of local maxima of *Response* *(search local neighborhoods, e.g. 3x3 or 5x5 for location of maximum response).*

Harris Detector: Workflow

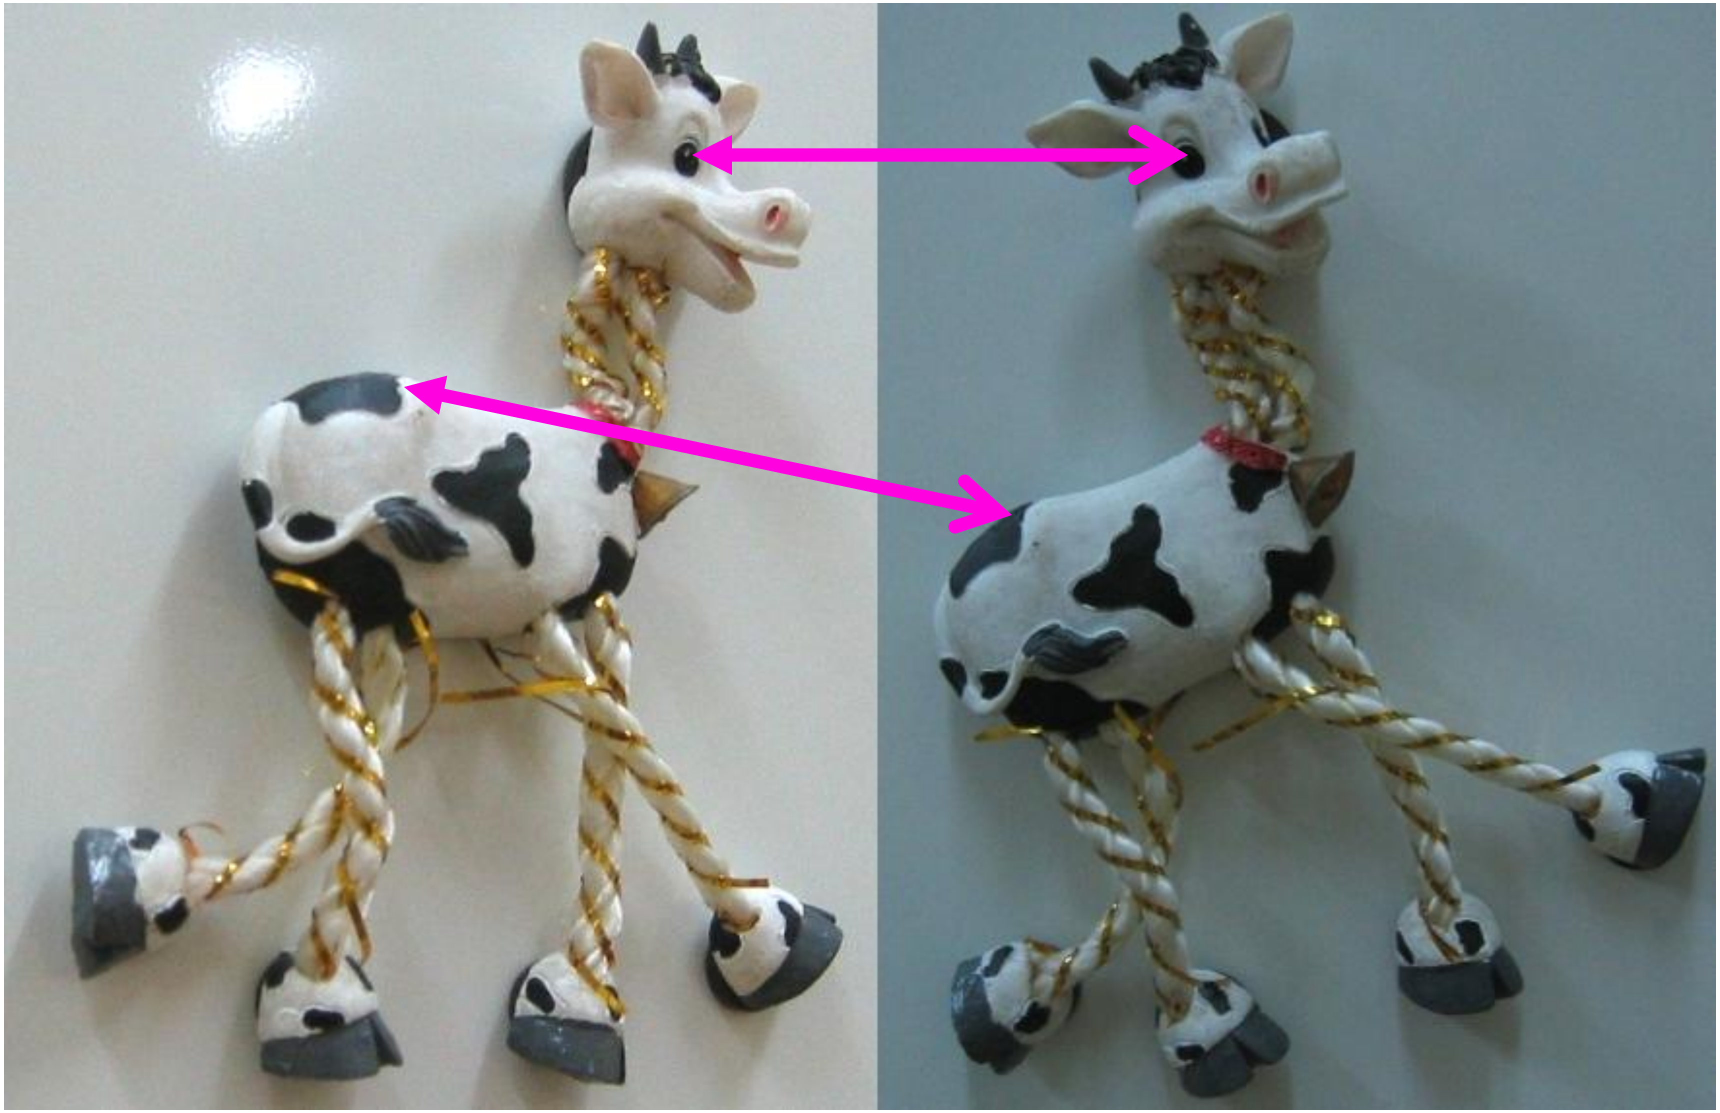

Slide adapted form Darya Frolova, Denis Simakov, Weizmann Institute.

Harris Detector: Workflow

Compute corner response R

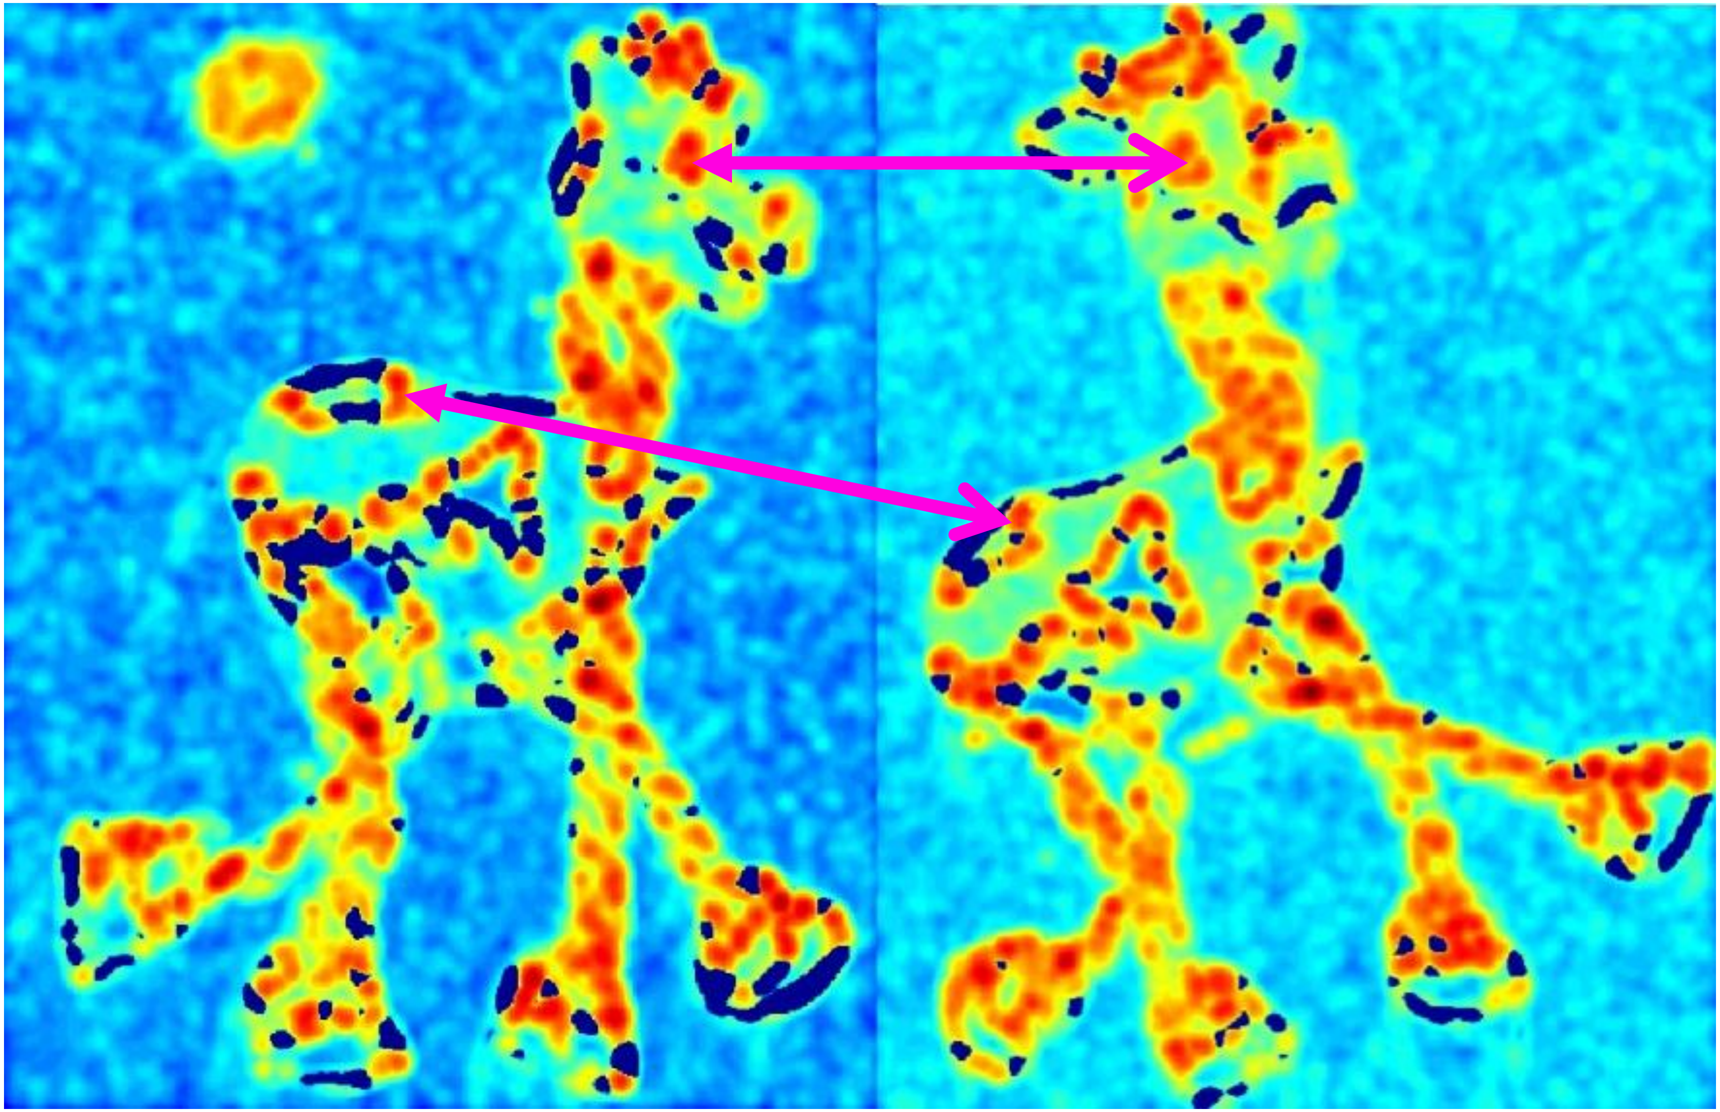

Harris Detector: Workflow

Find points with large corner response: R>threshold

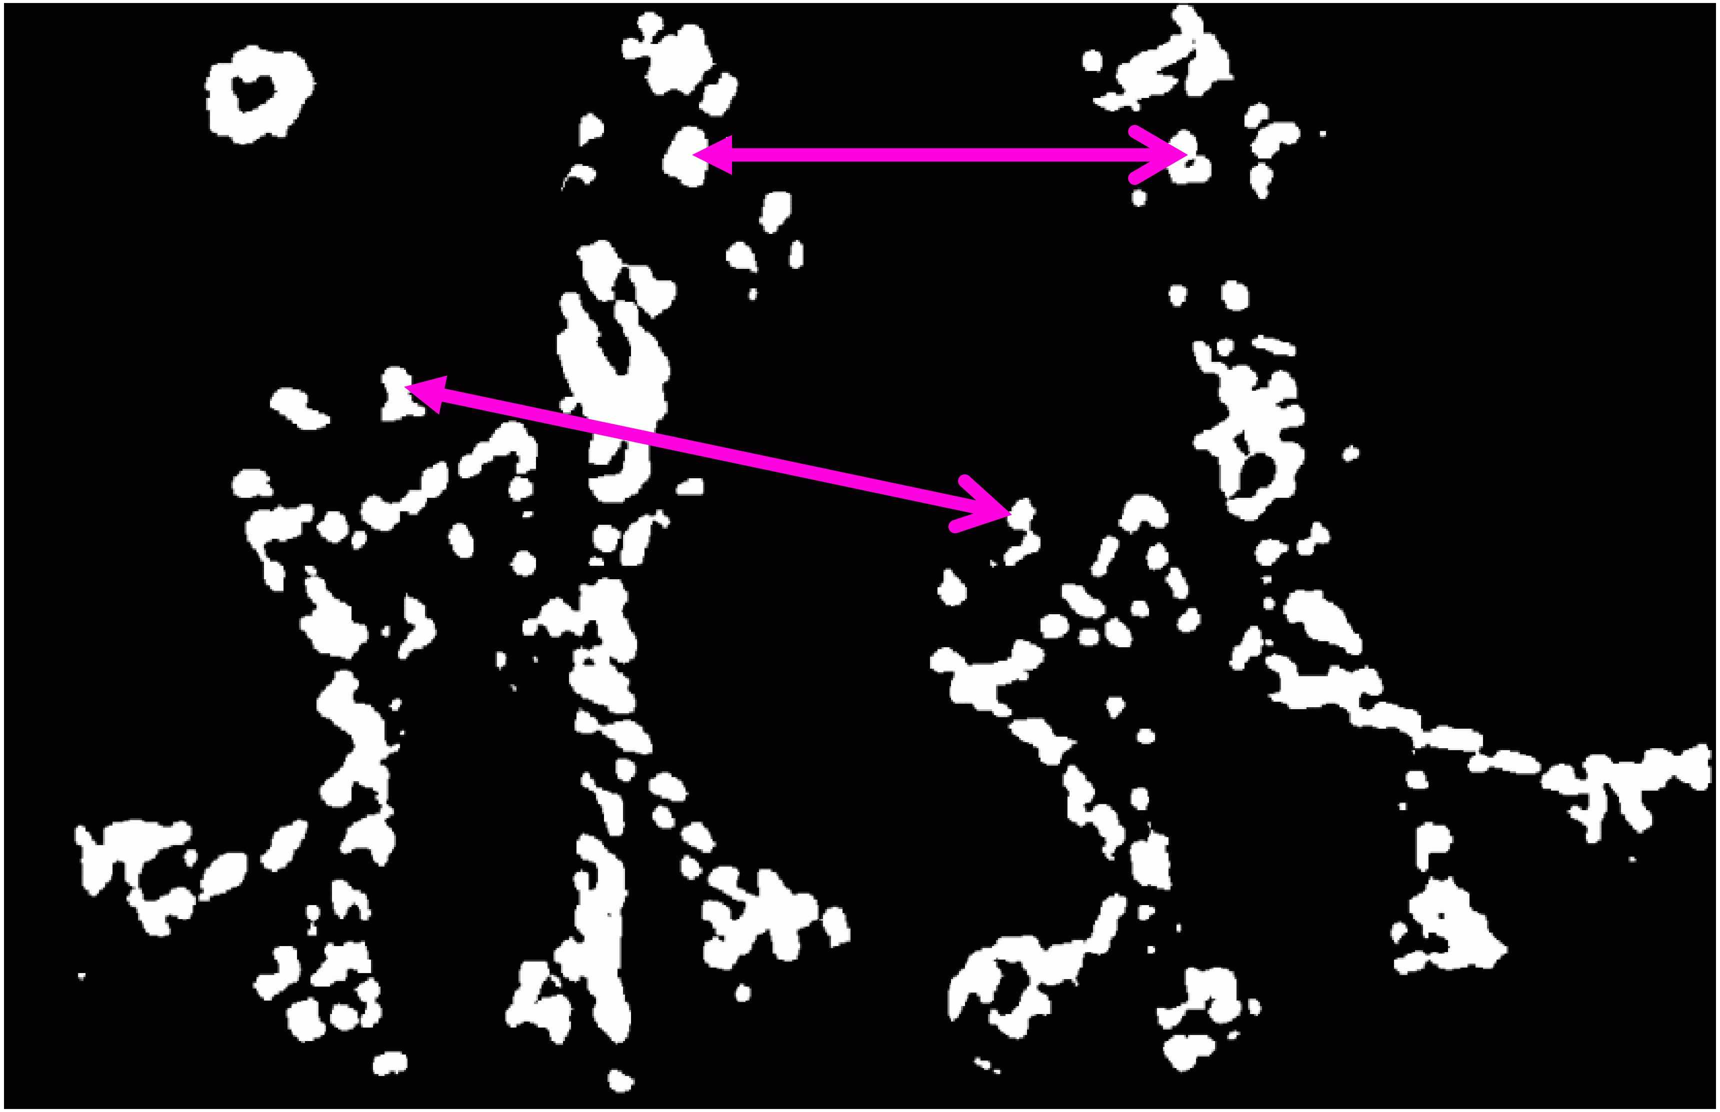

Harris Detector: Workflow

Take only the points of local maxima of R

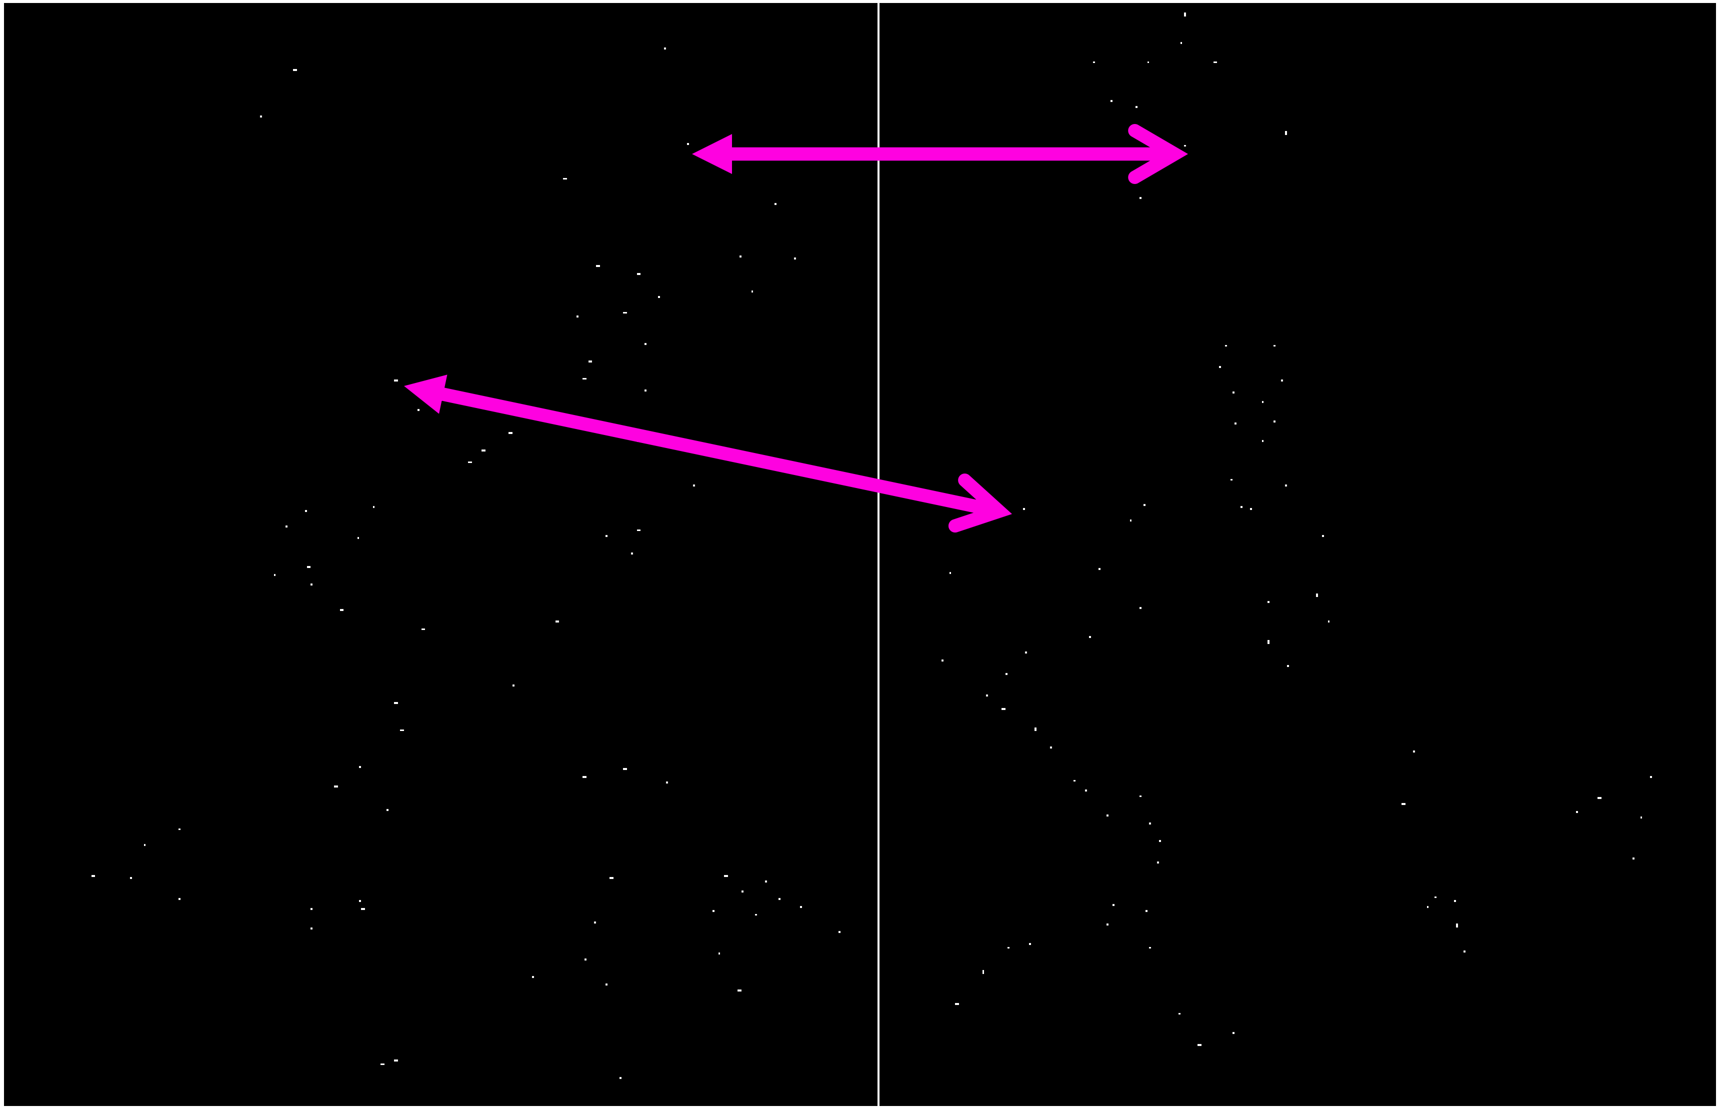

Harris Detector: Workflow