THE UNIVERSITY OF BRITISH COLUMBIA

CPSC 425: Computer Vision

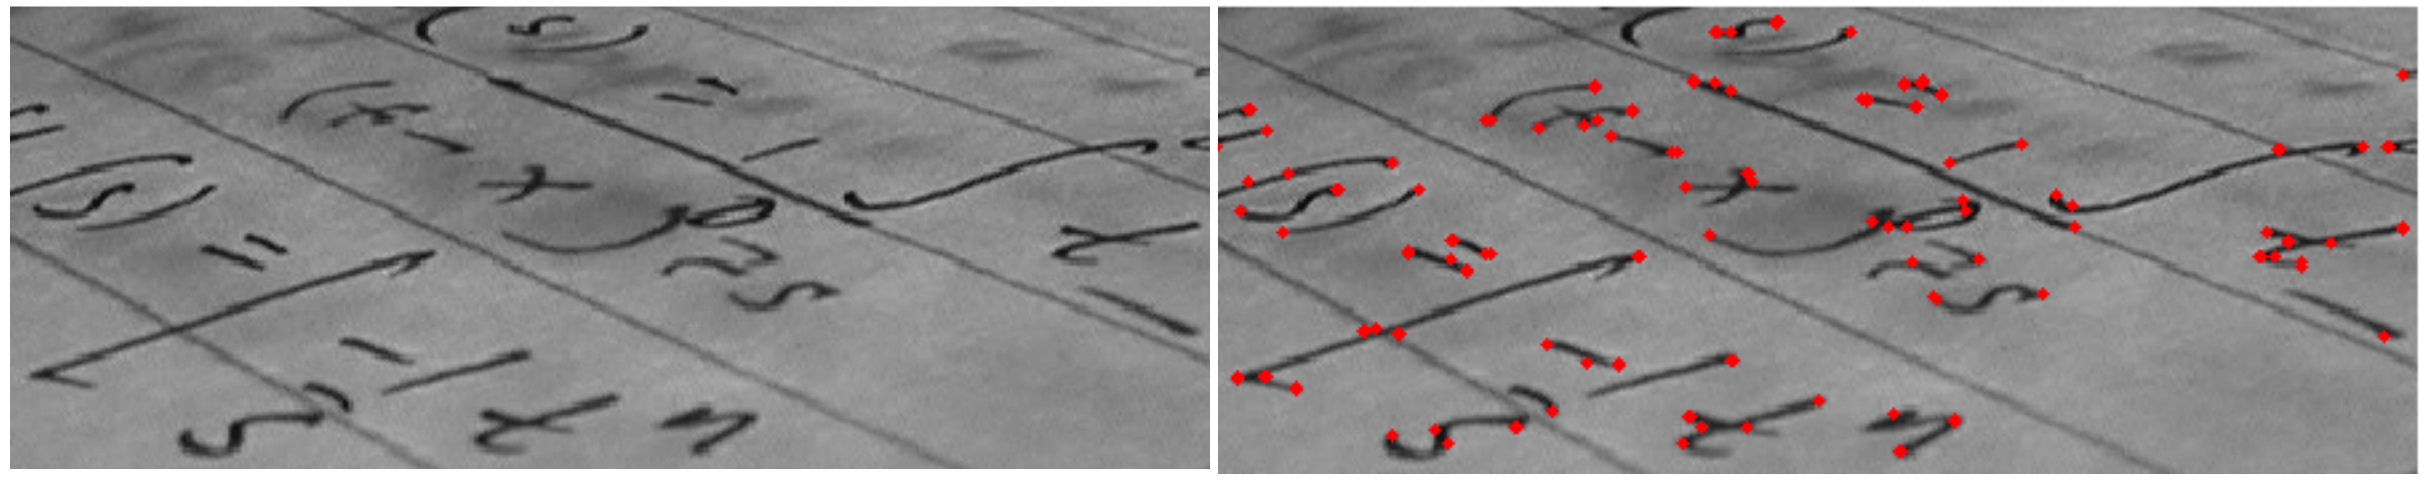

Image Credit: https://en.wikipedia.org/wiki/Corner_detection

Lecture 12: Corner Detection

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

Menu for Today (**October 1, 2018**)

Topics:

— Corner Detection

— Autocorrelation

— Harris Corner Detector

Redings:

— **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 **October 10th**

Today’s “**fun**” Example:

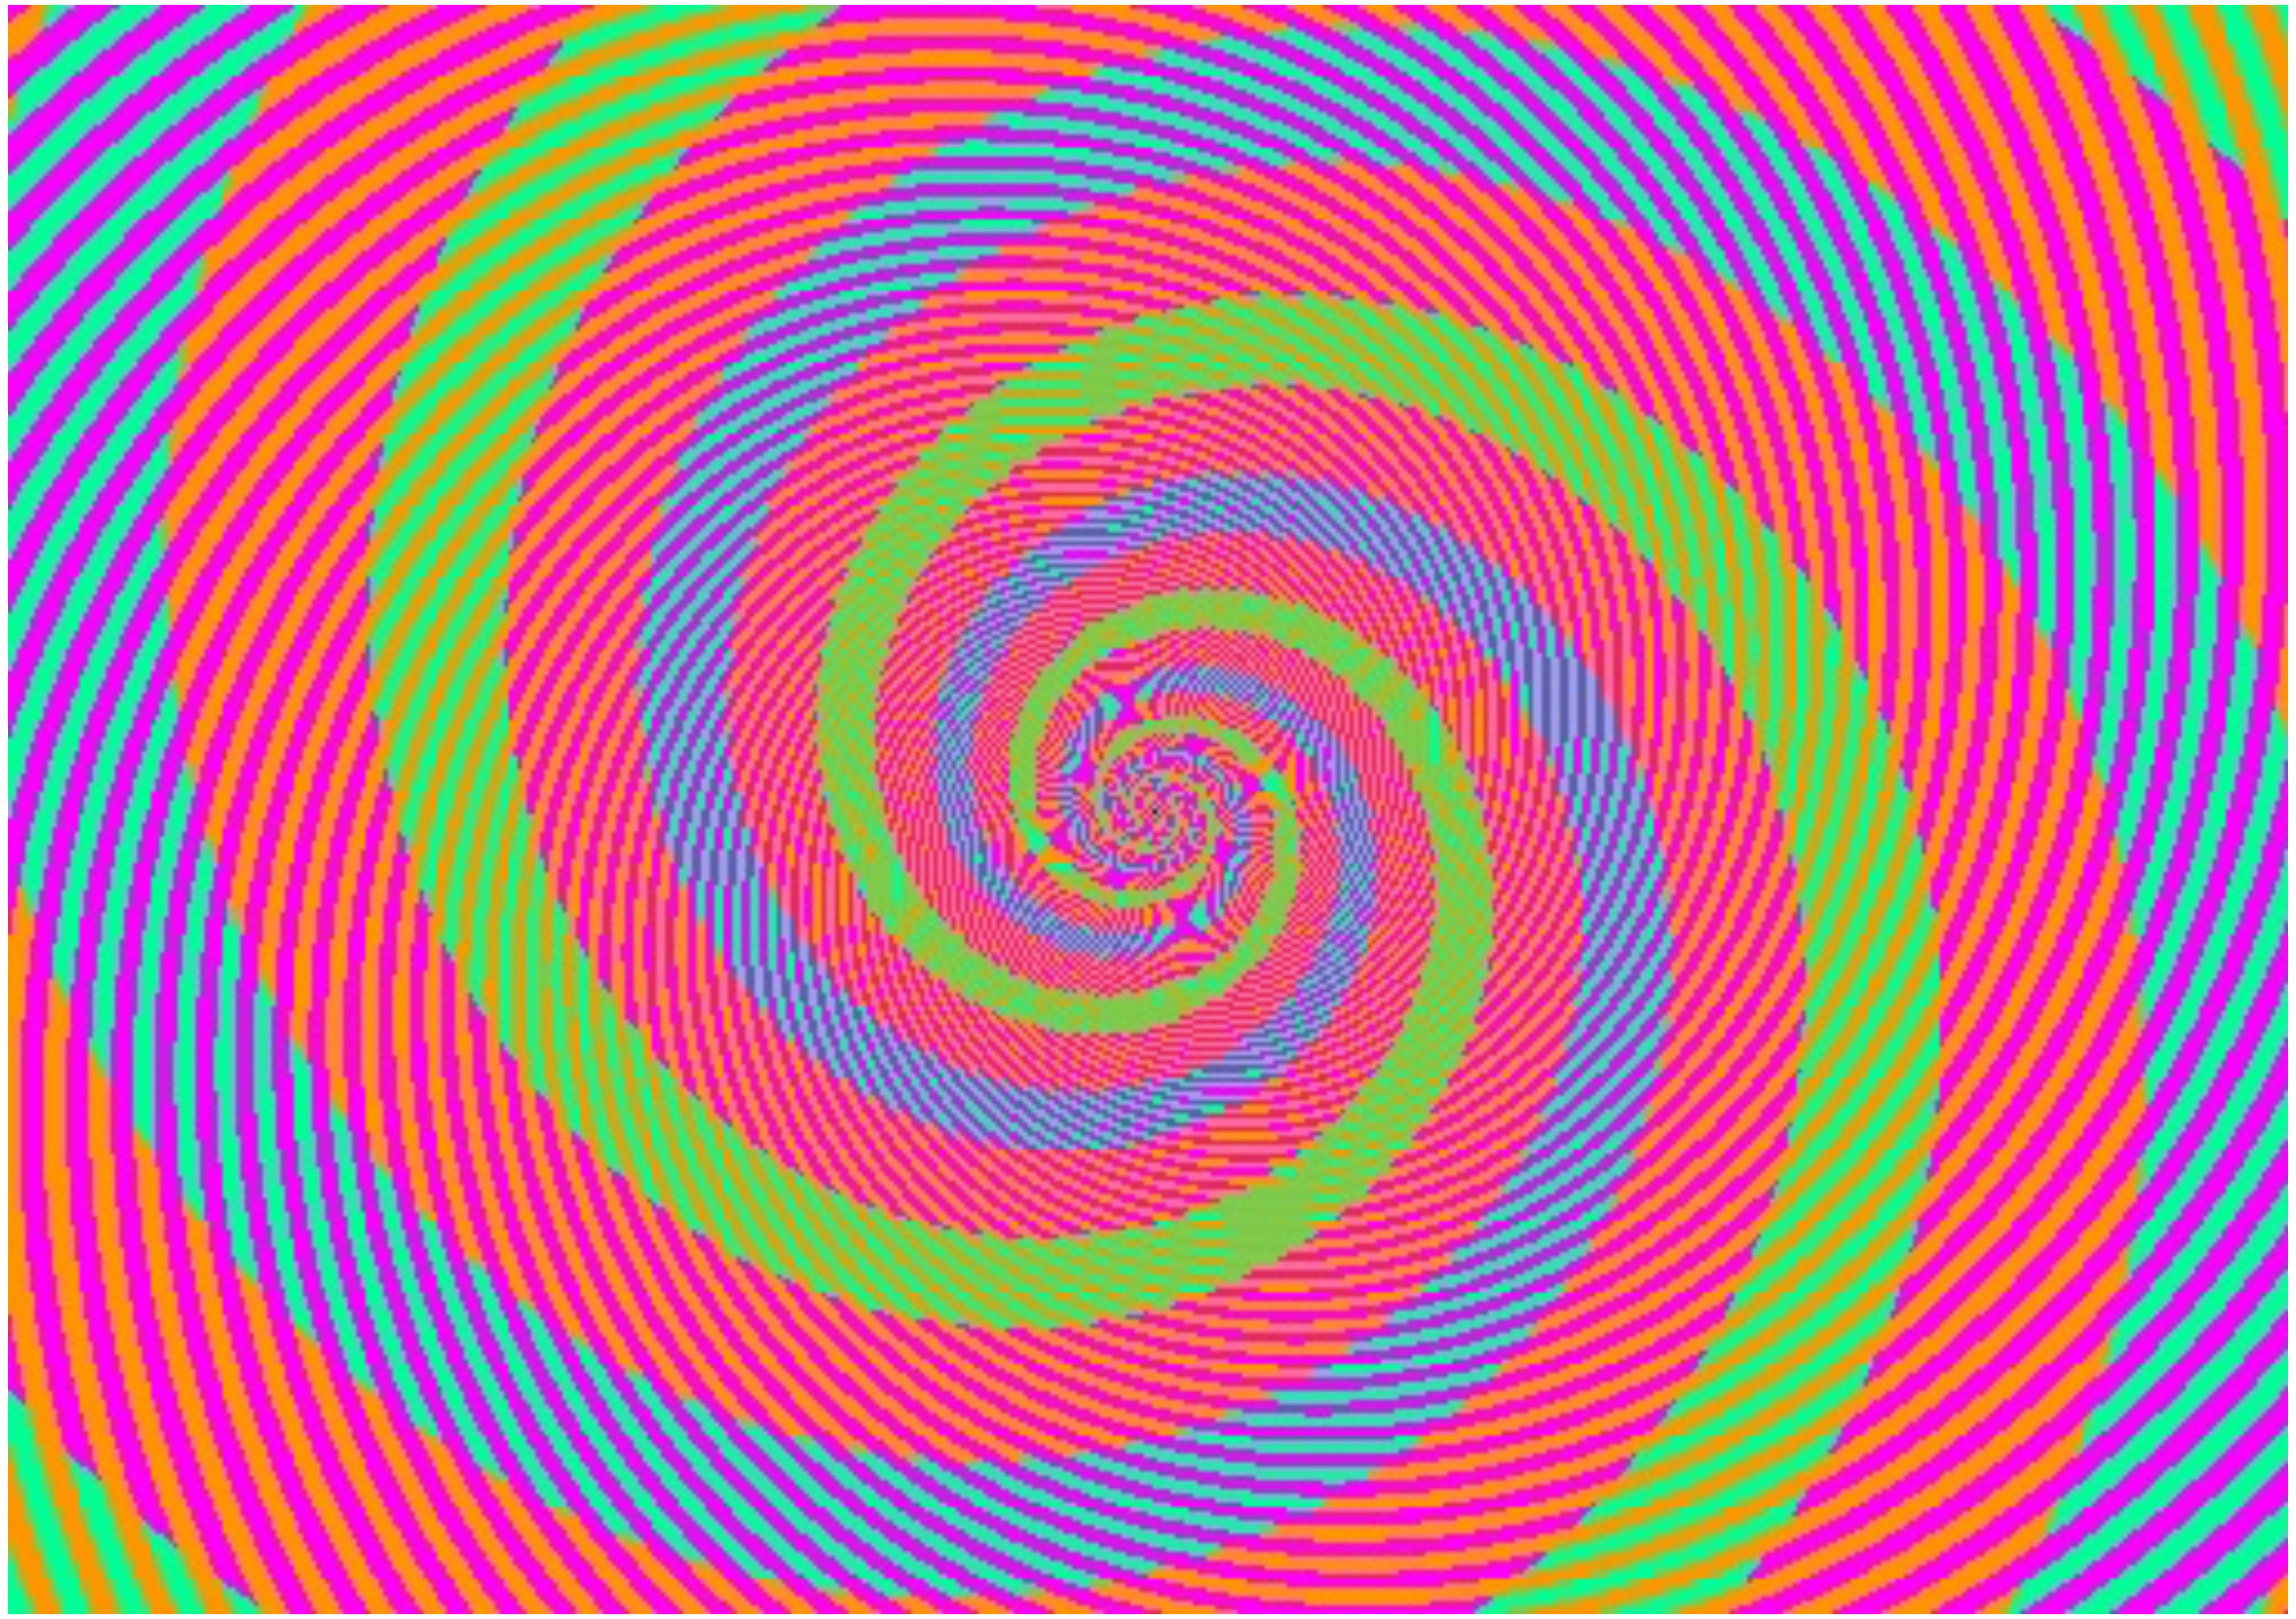

Today's "**fun**" Example:

Today's "**fun**" Example:

Lecture 11: 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 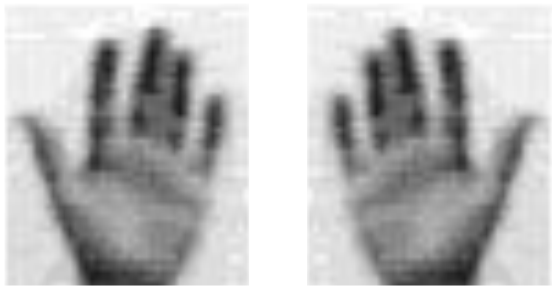

- Partial Occlusions 
- Different Perspective
- Motion / blur

Motivation: Template Matching in Scaled Representation

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

- ~~Different scales~~ 
- Different orientation 
- Lighting conditions 
- Left vs. Right hand 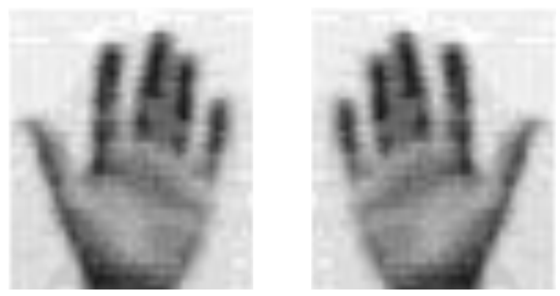
- Partial Occlusions 
- Different Perspective
- Motion / blur

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recognition under **Occlusion**

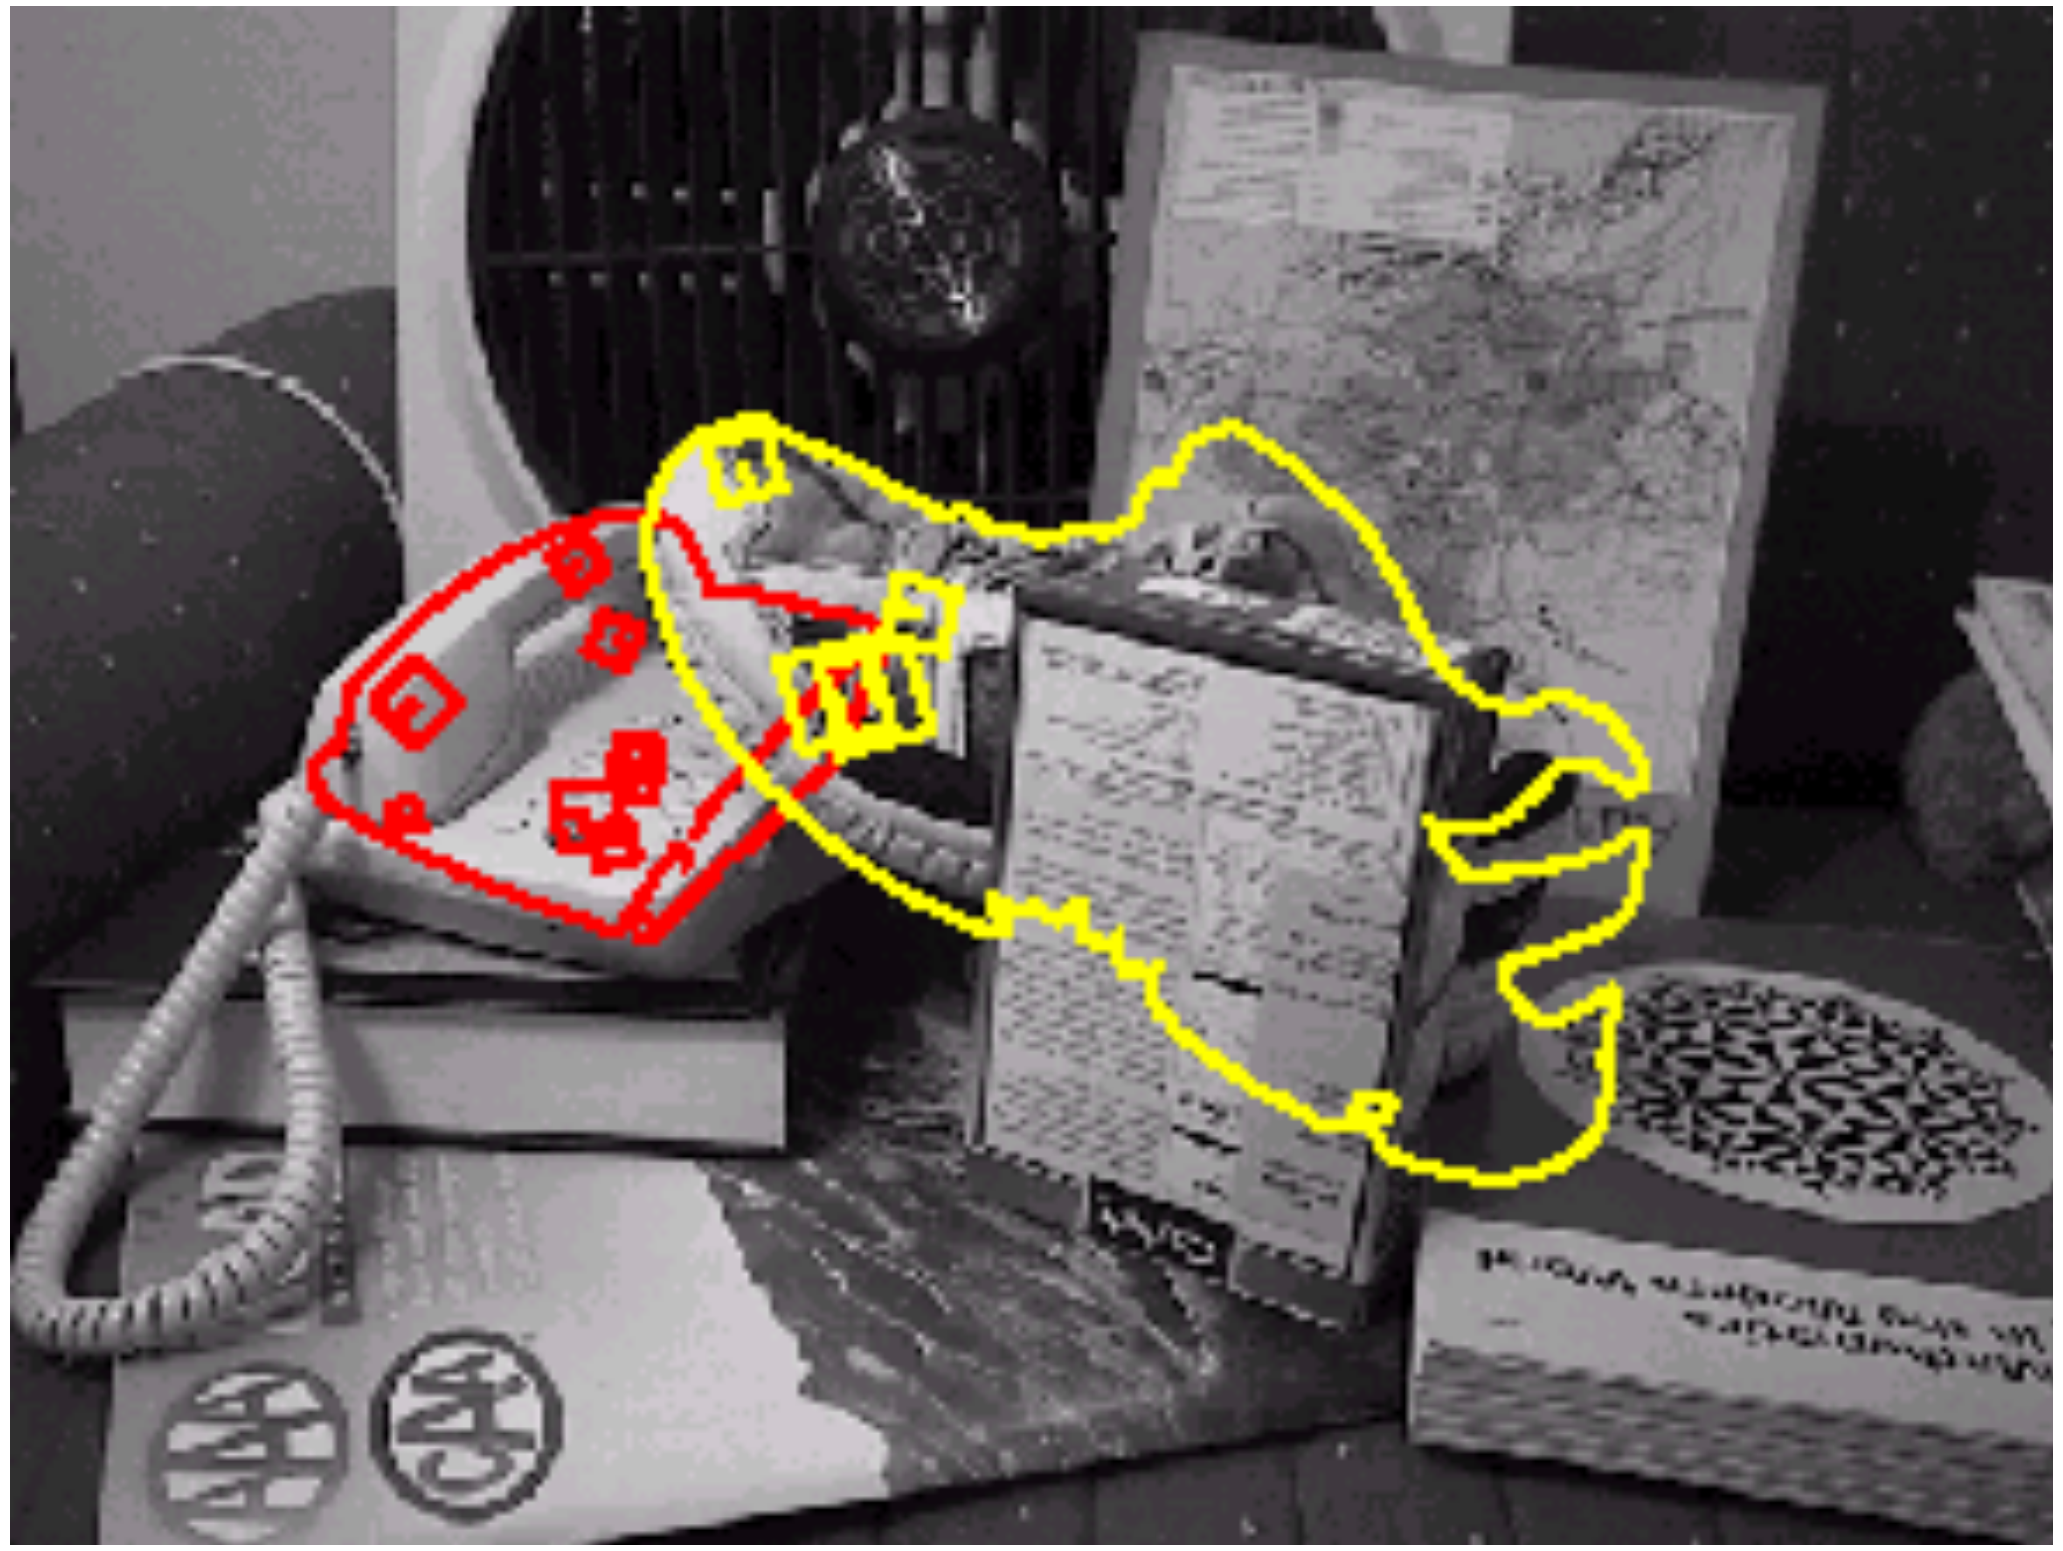

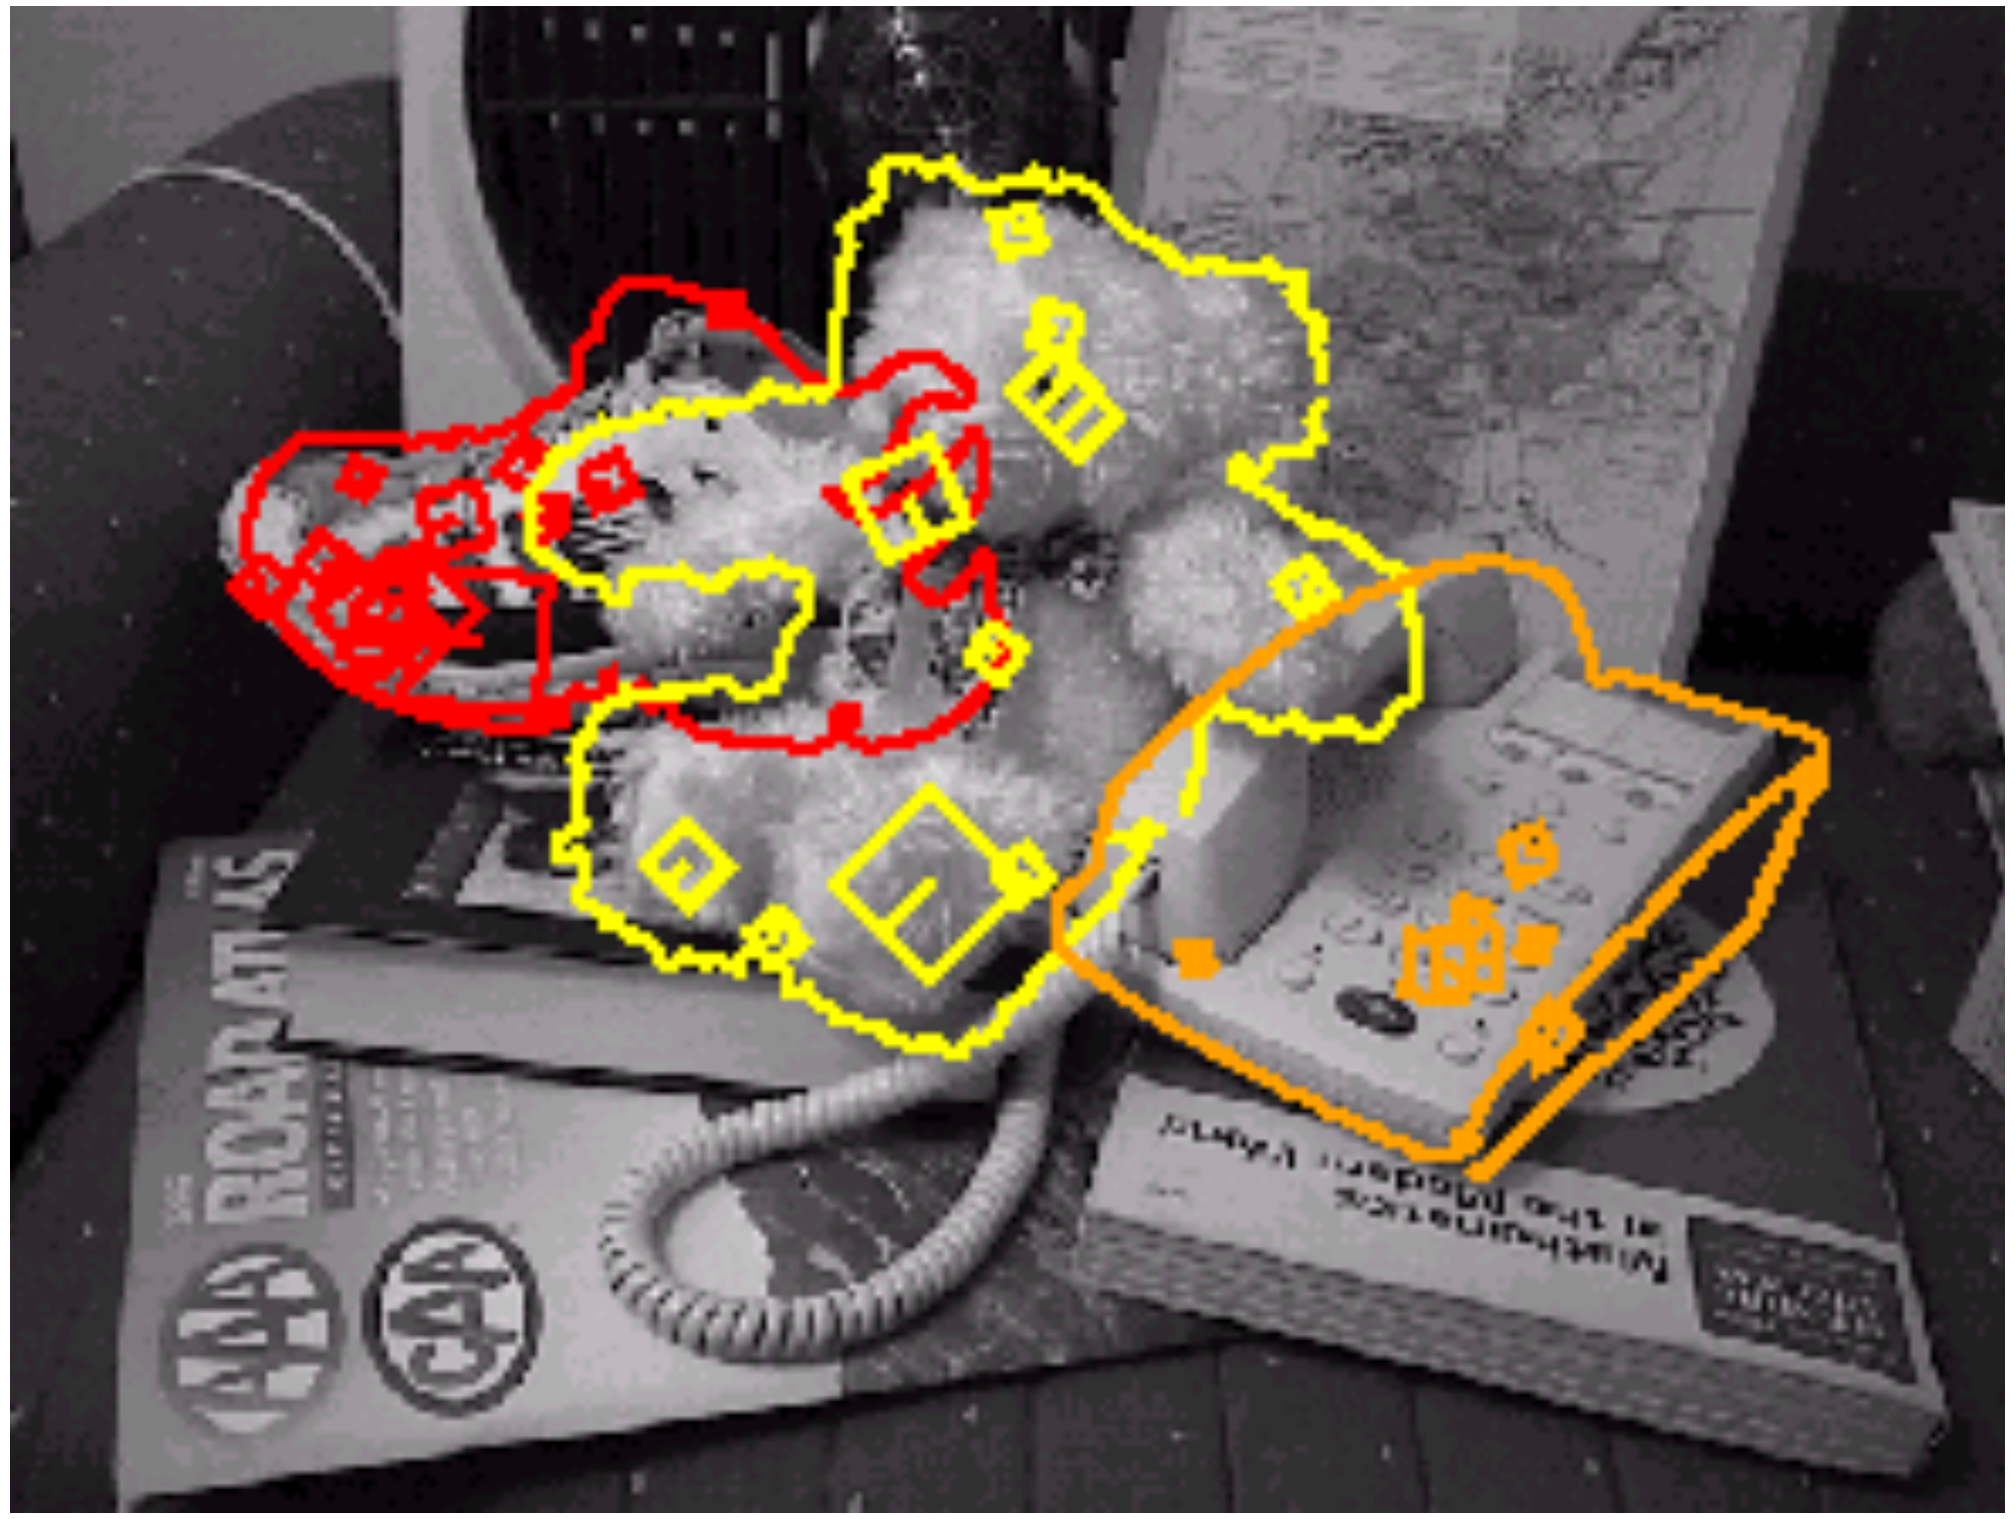

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Image **Matching**

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Image **Matching**

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Finding **Correspondences**

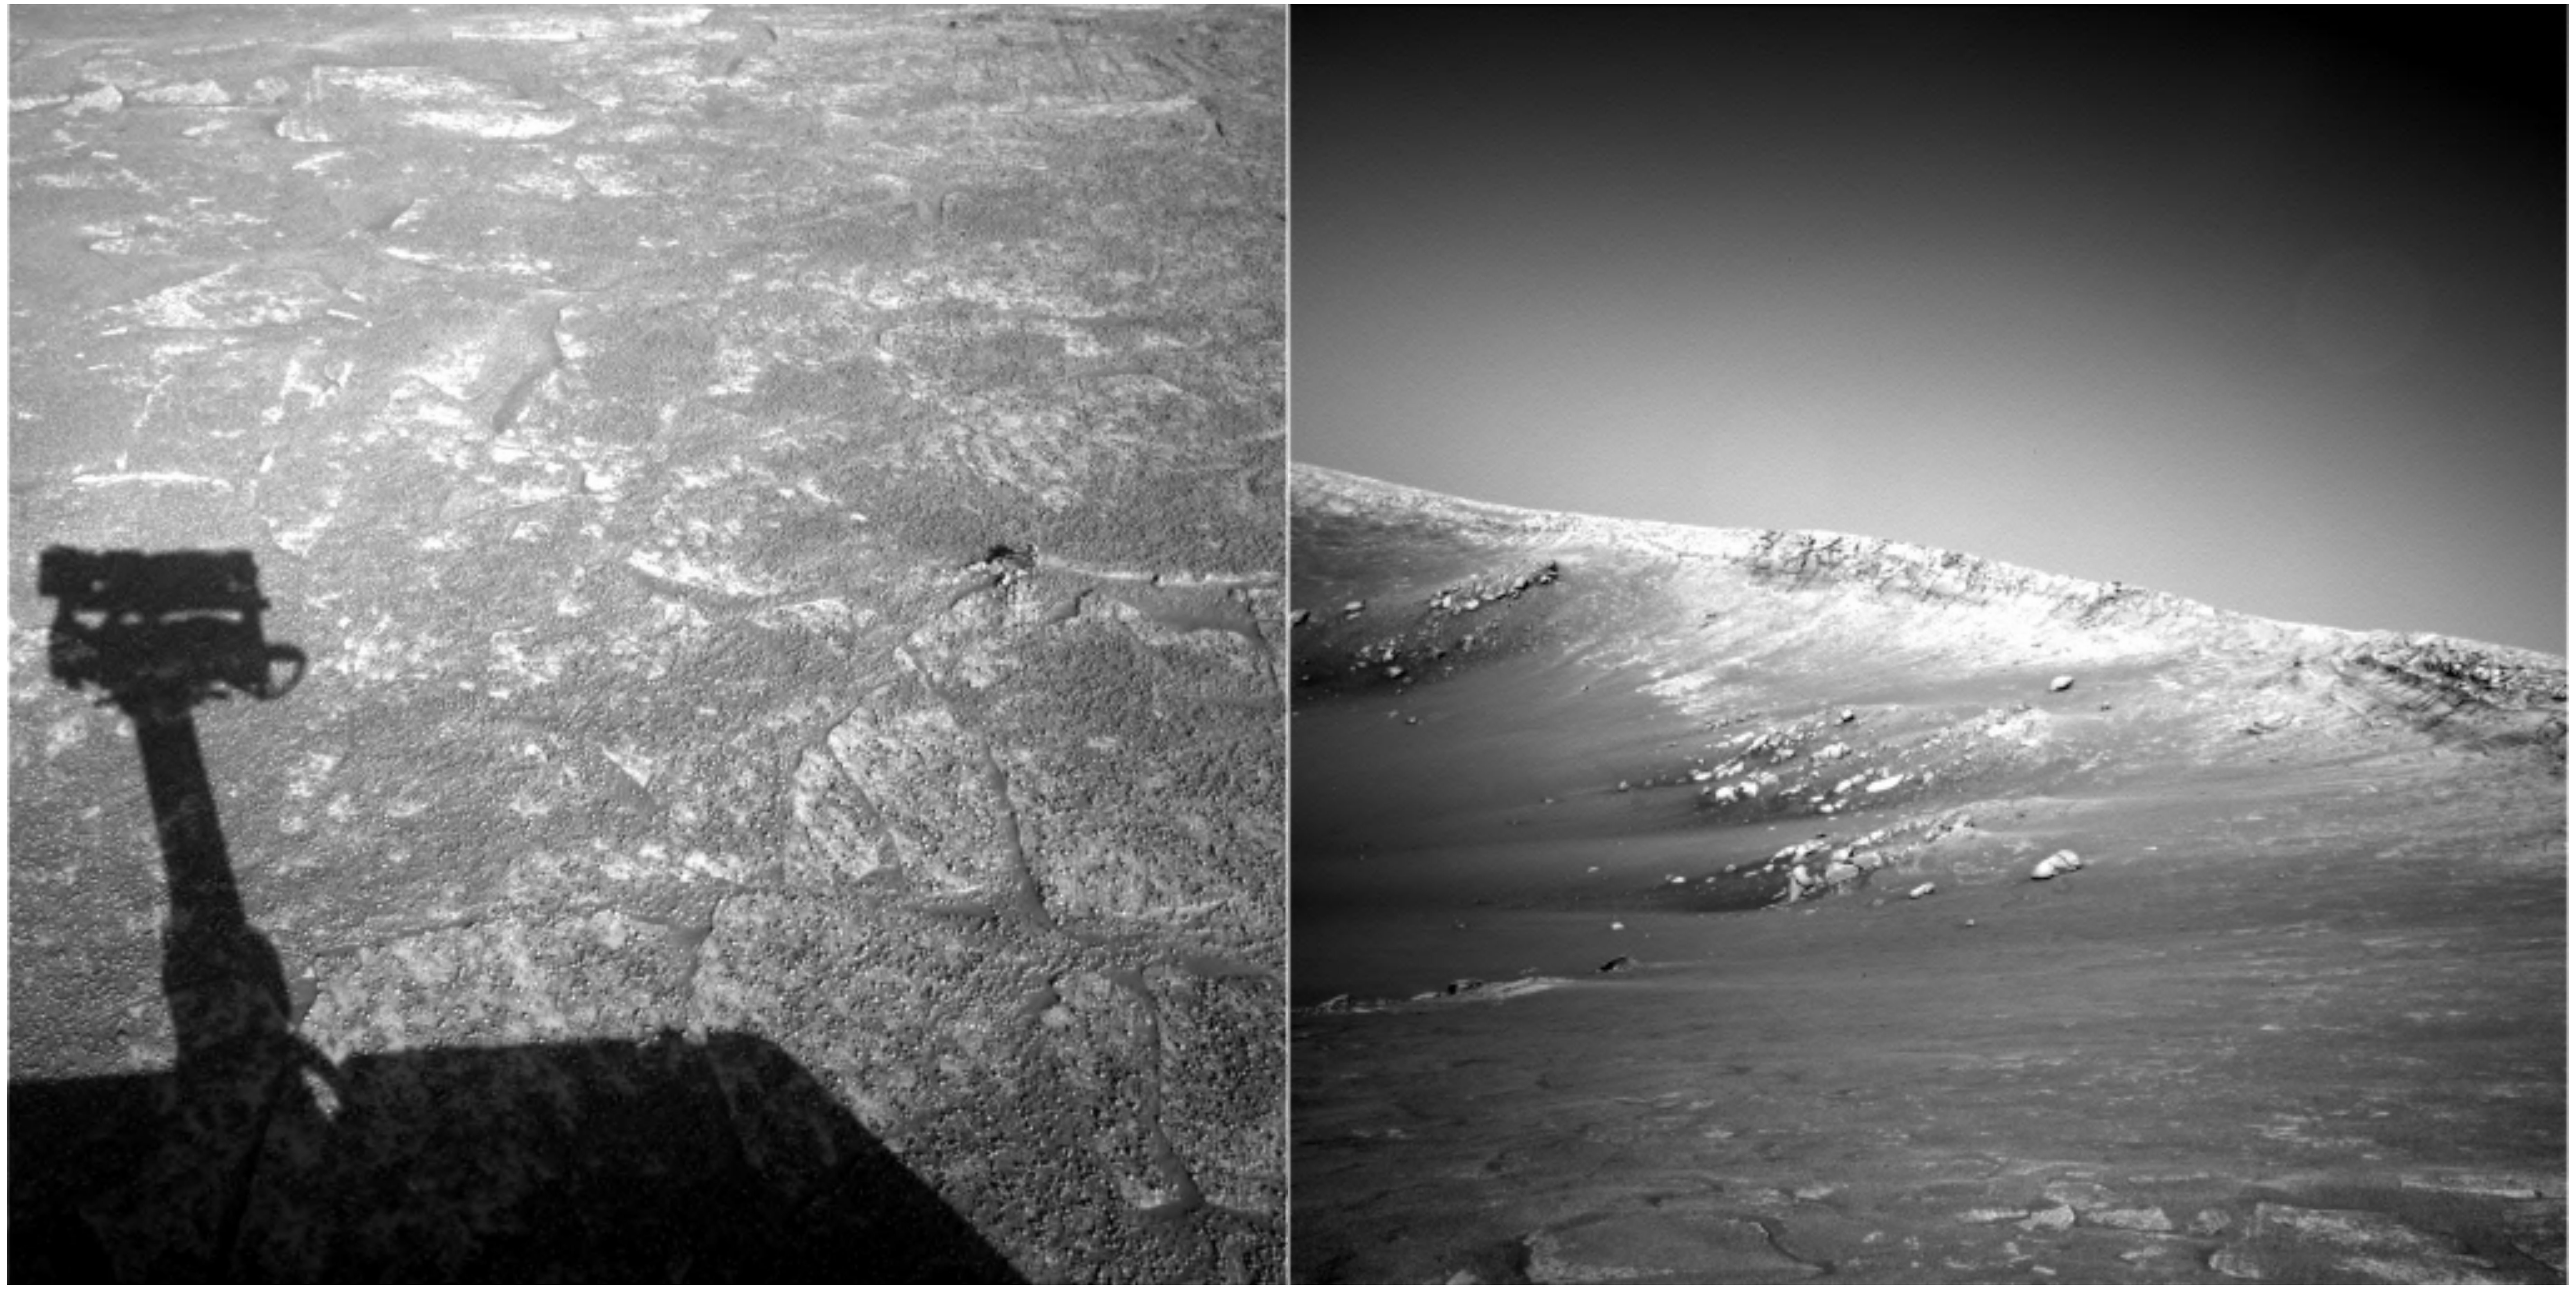

NASA Mars Rover images

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Finding **Correspondences**

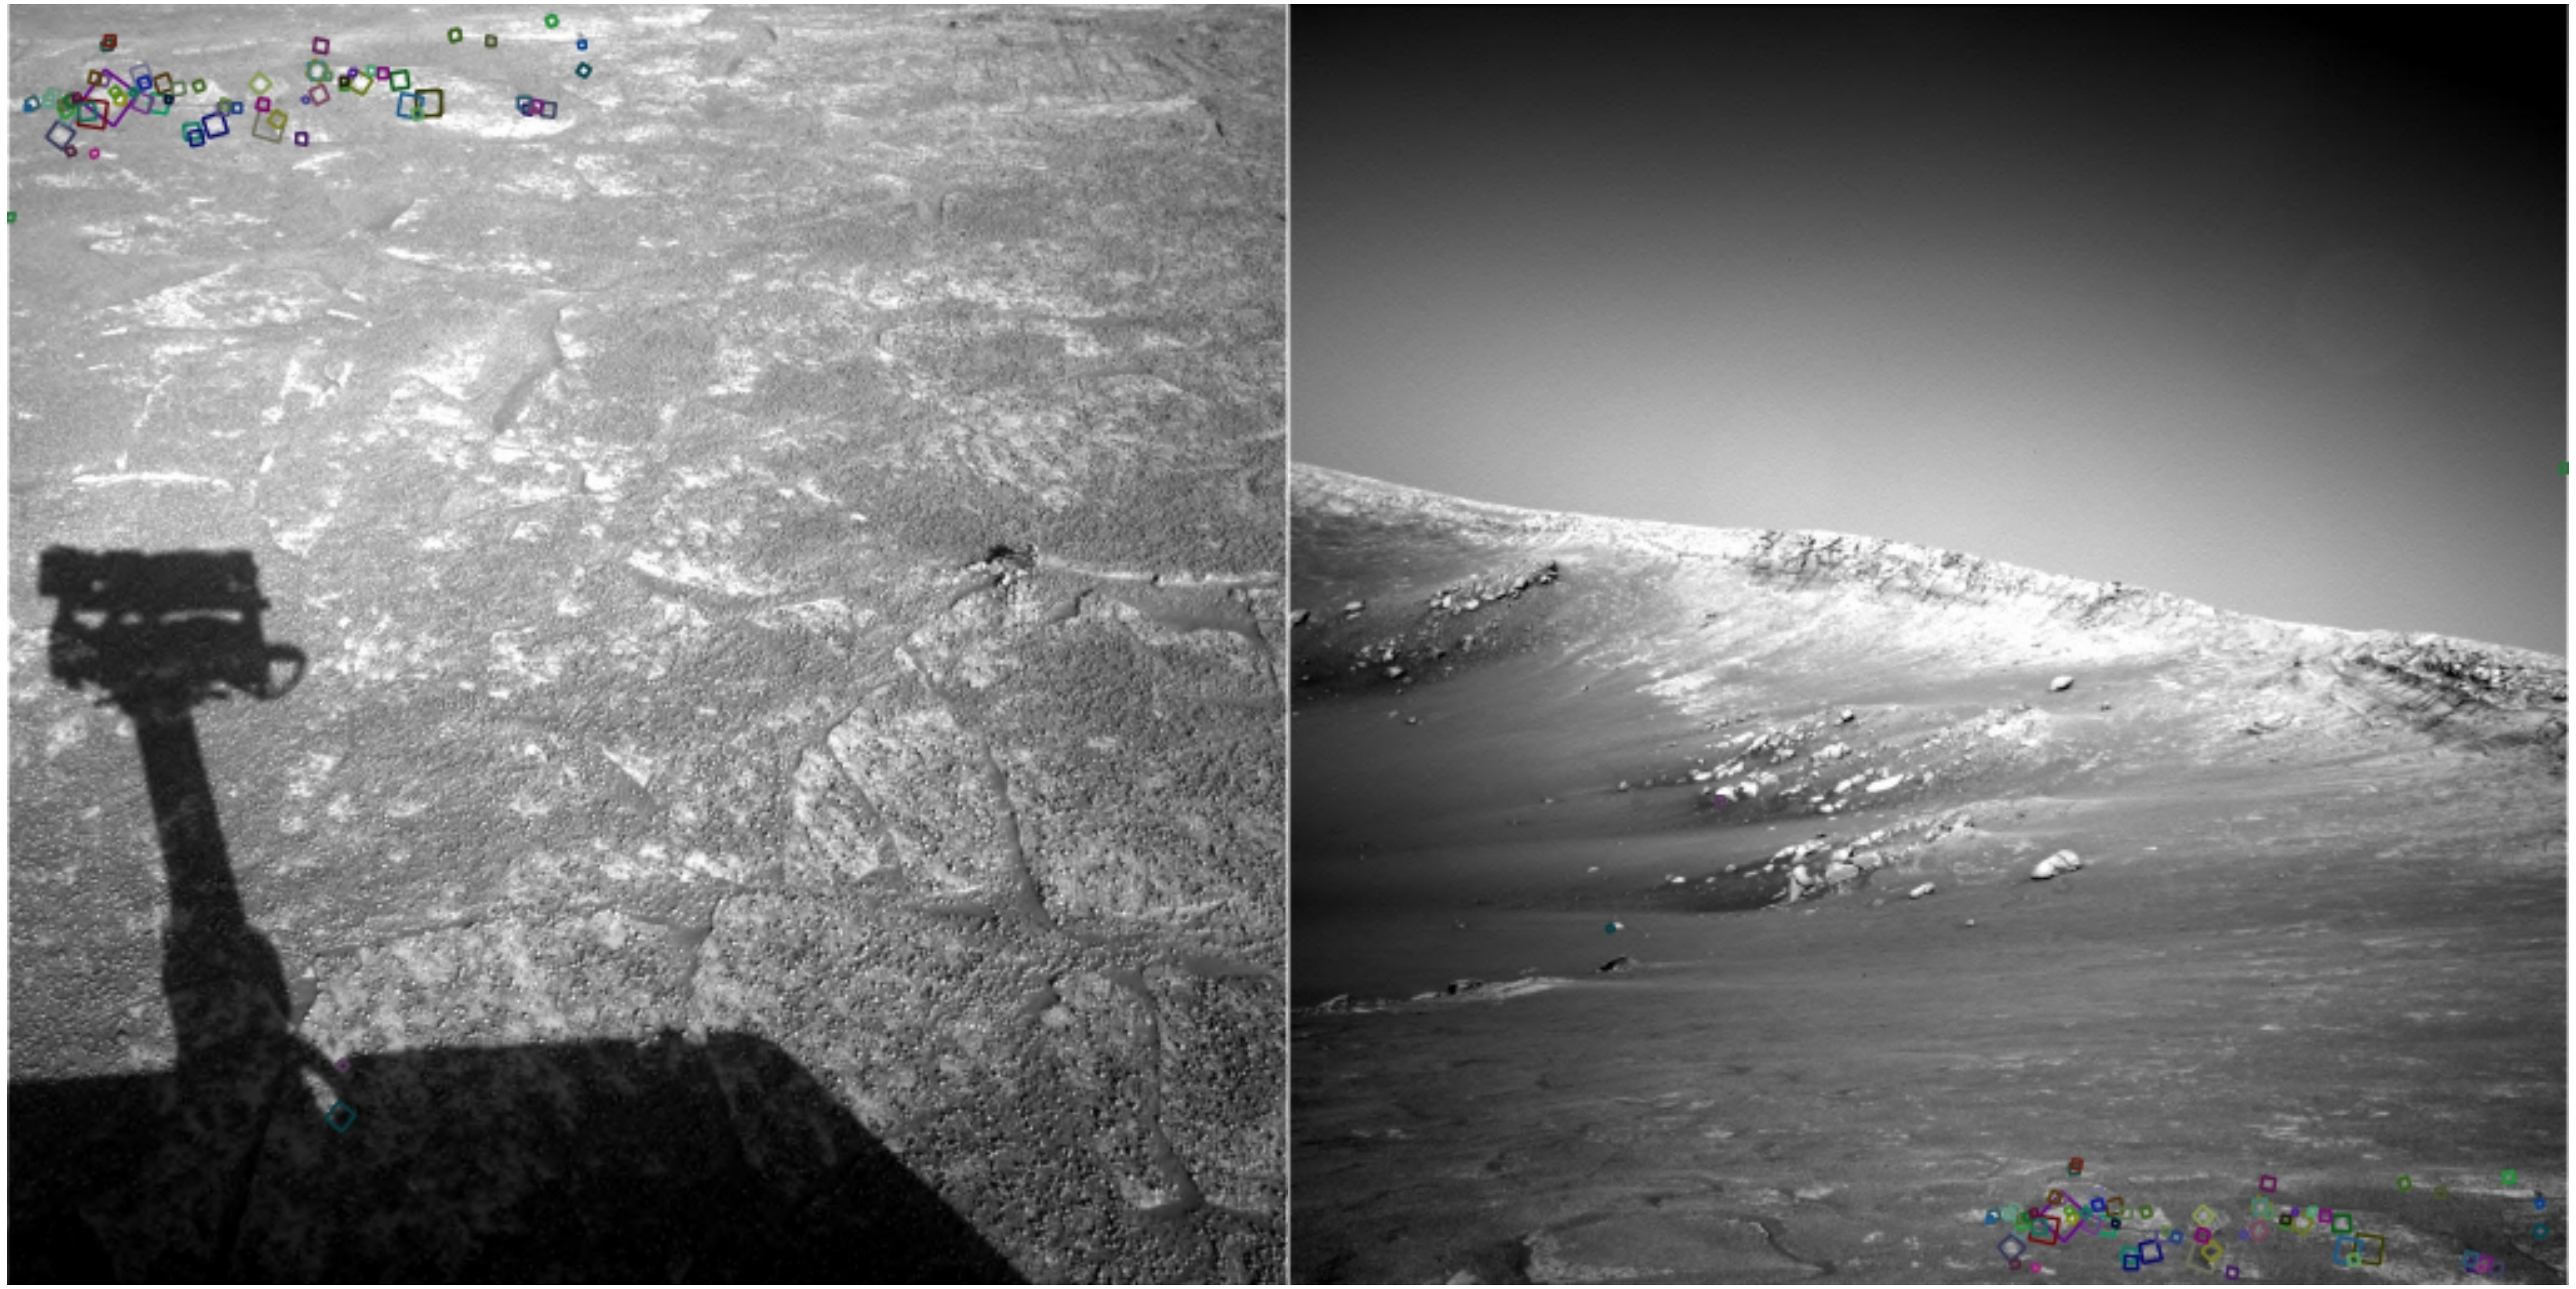

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

What is a **Good Feature**?

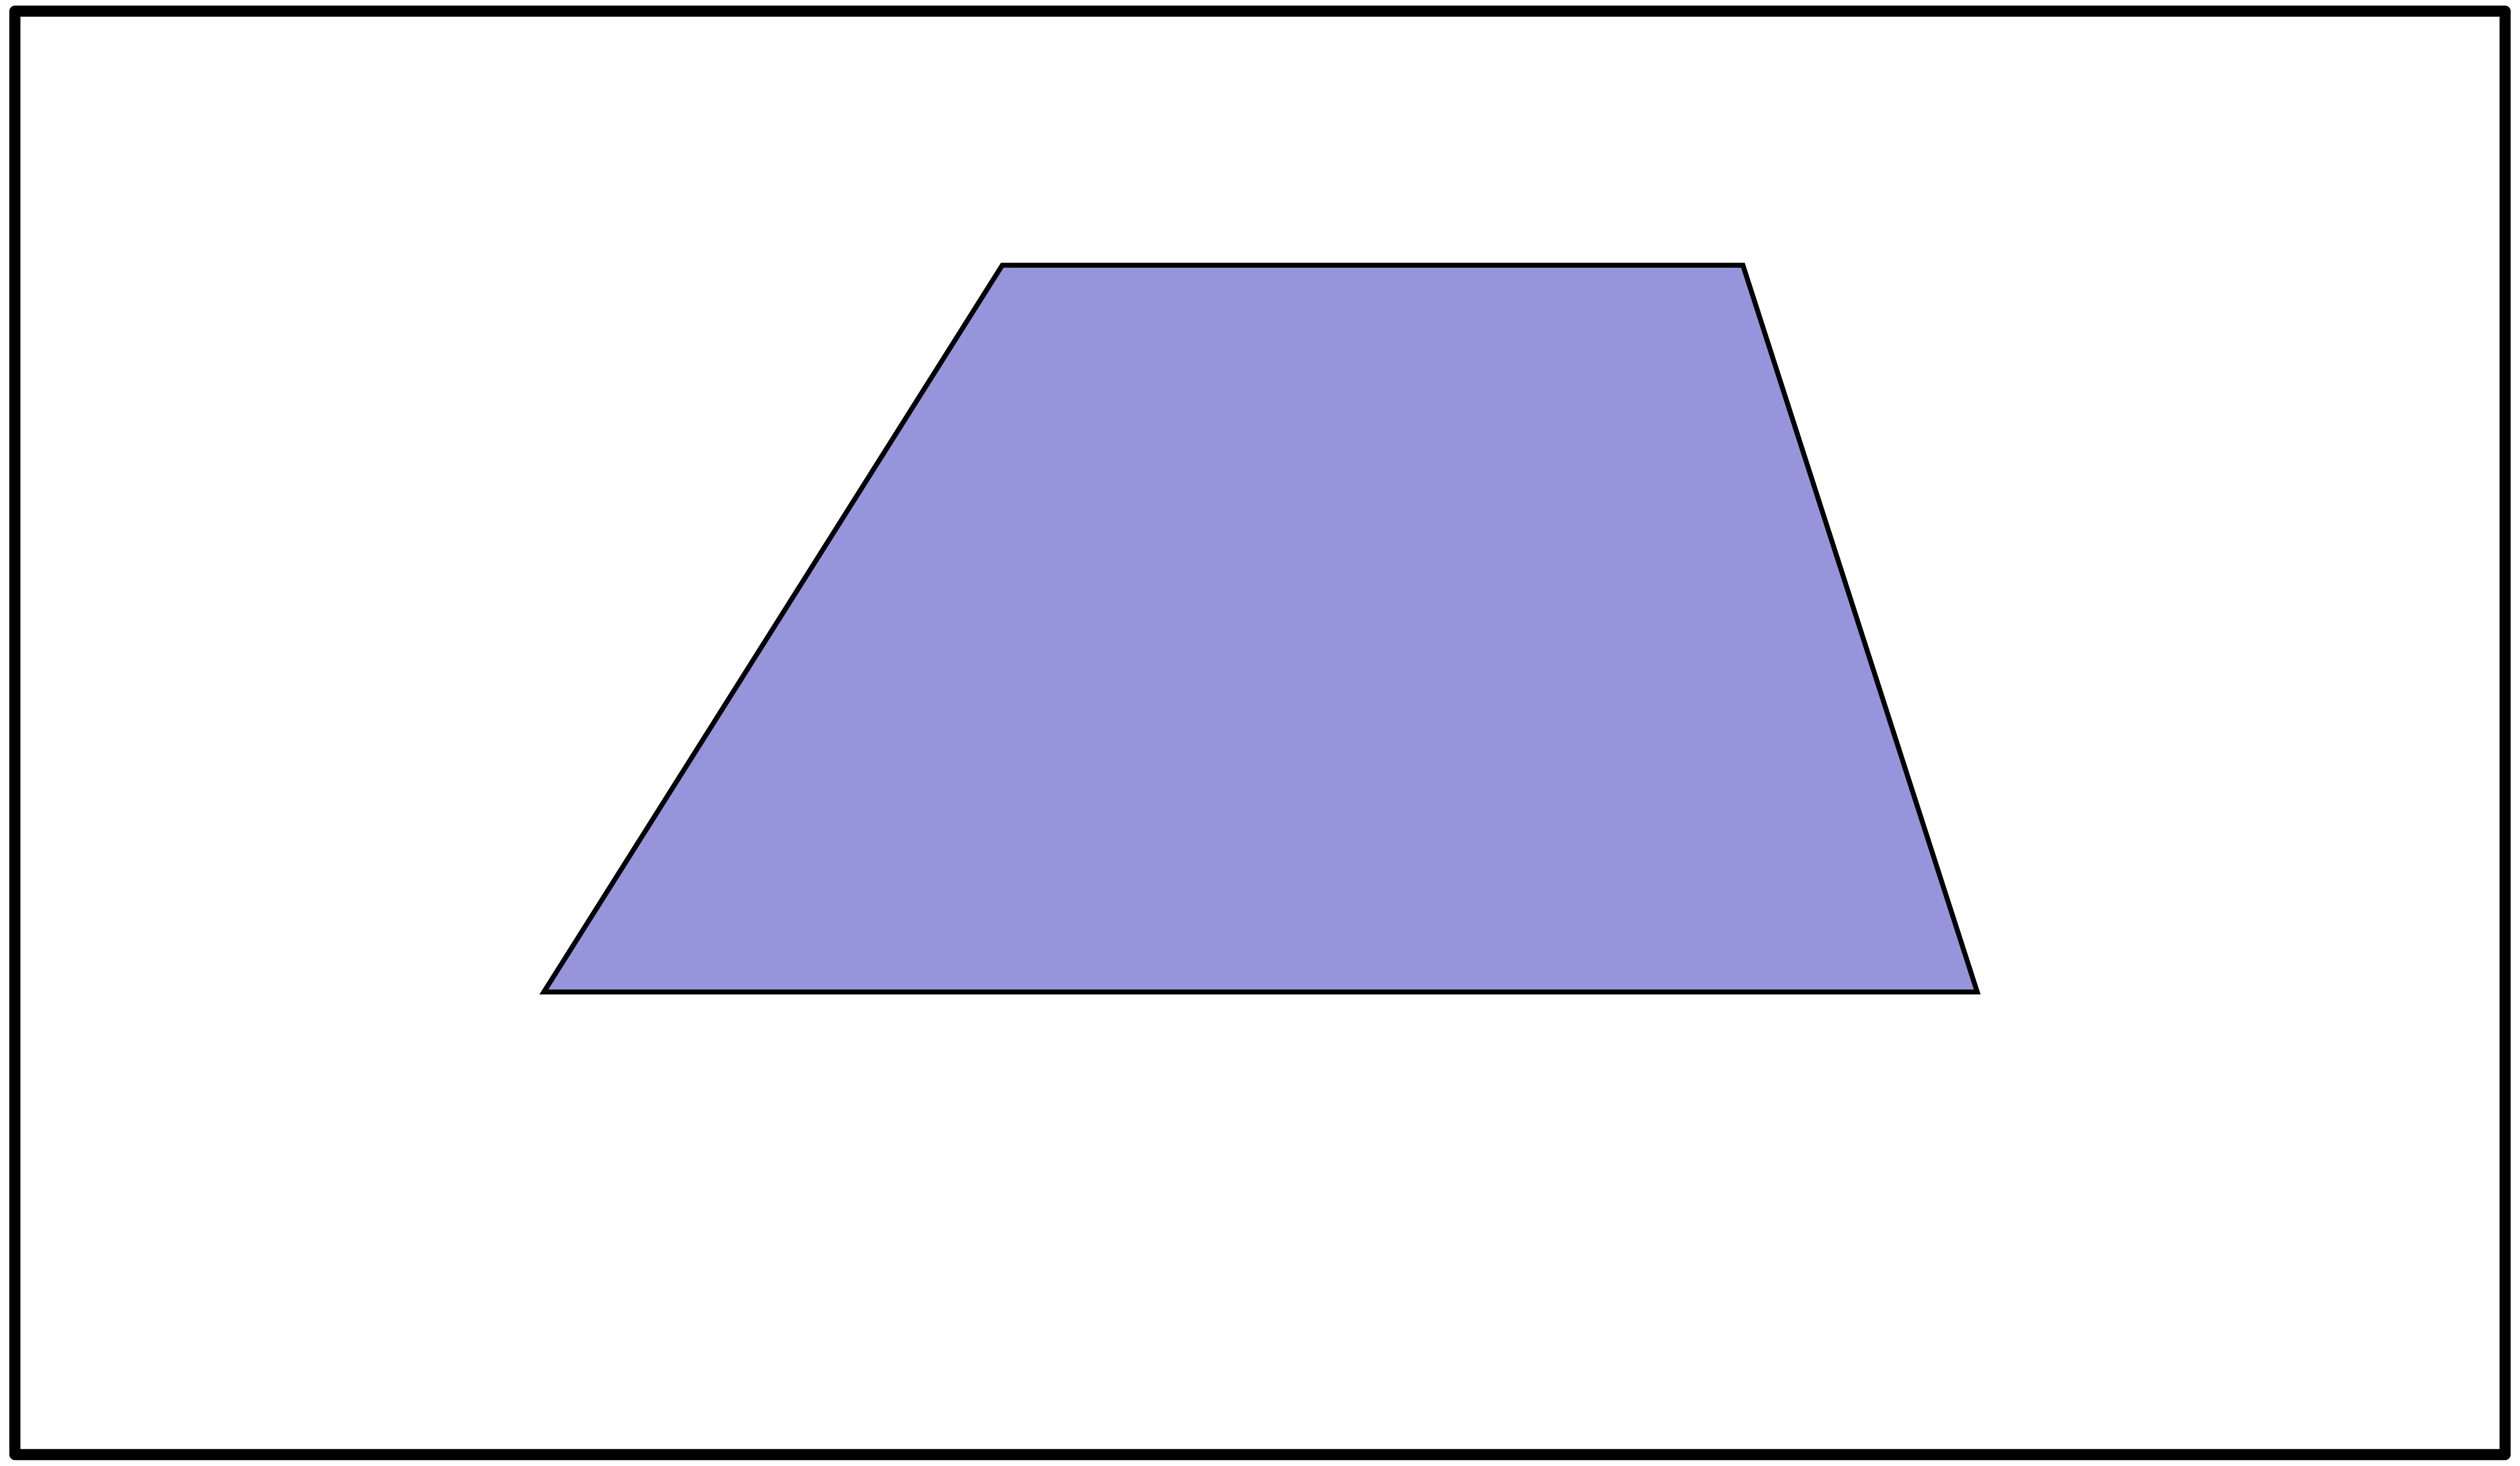

Pick a point in the image.
Find it again in the next image.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

What is a **Good Feature**?

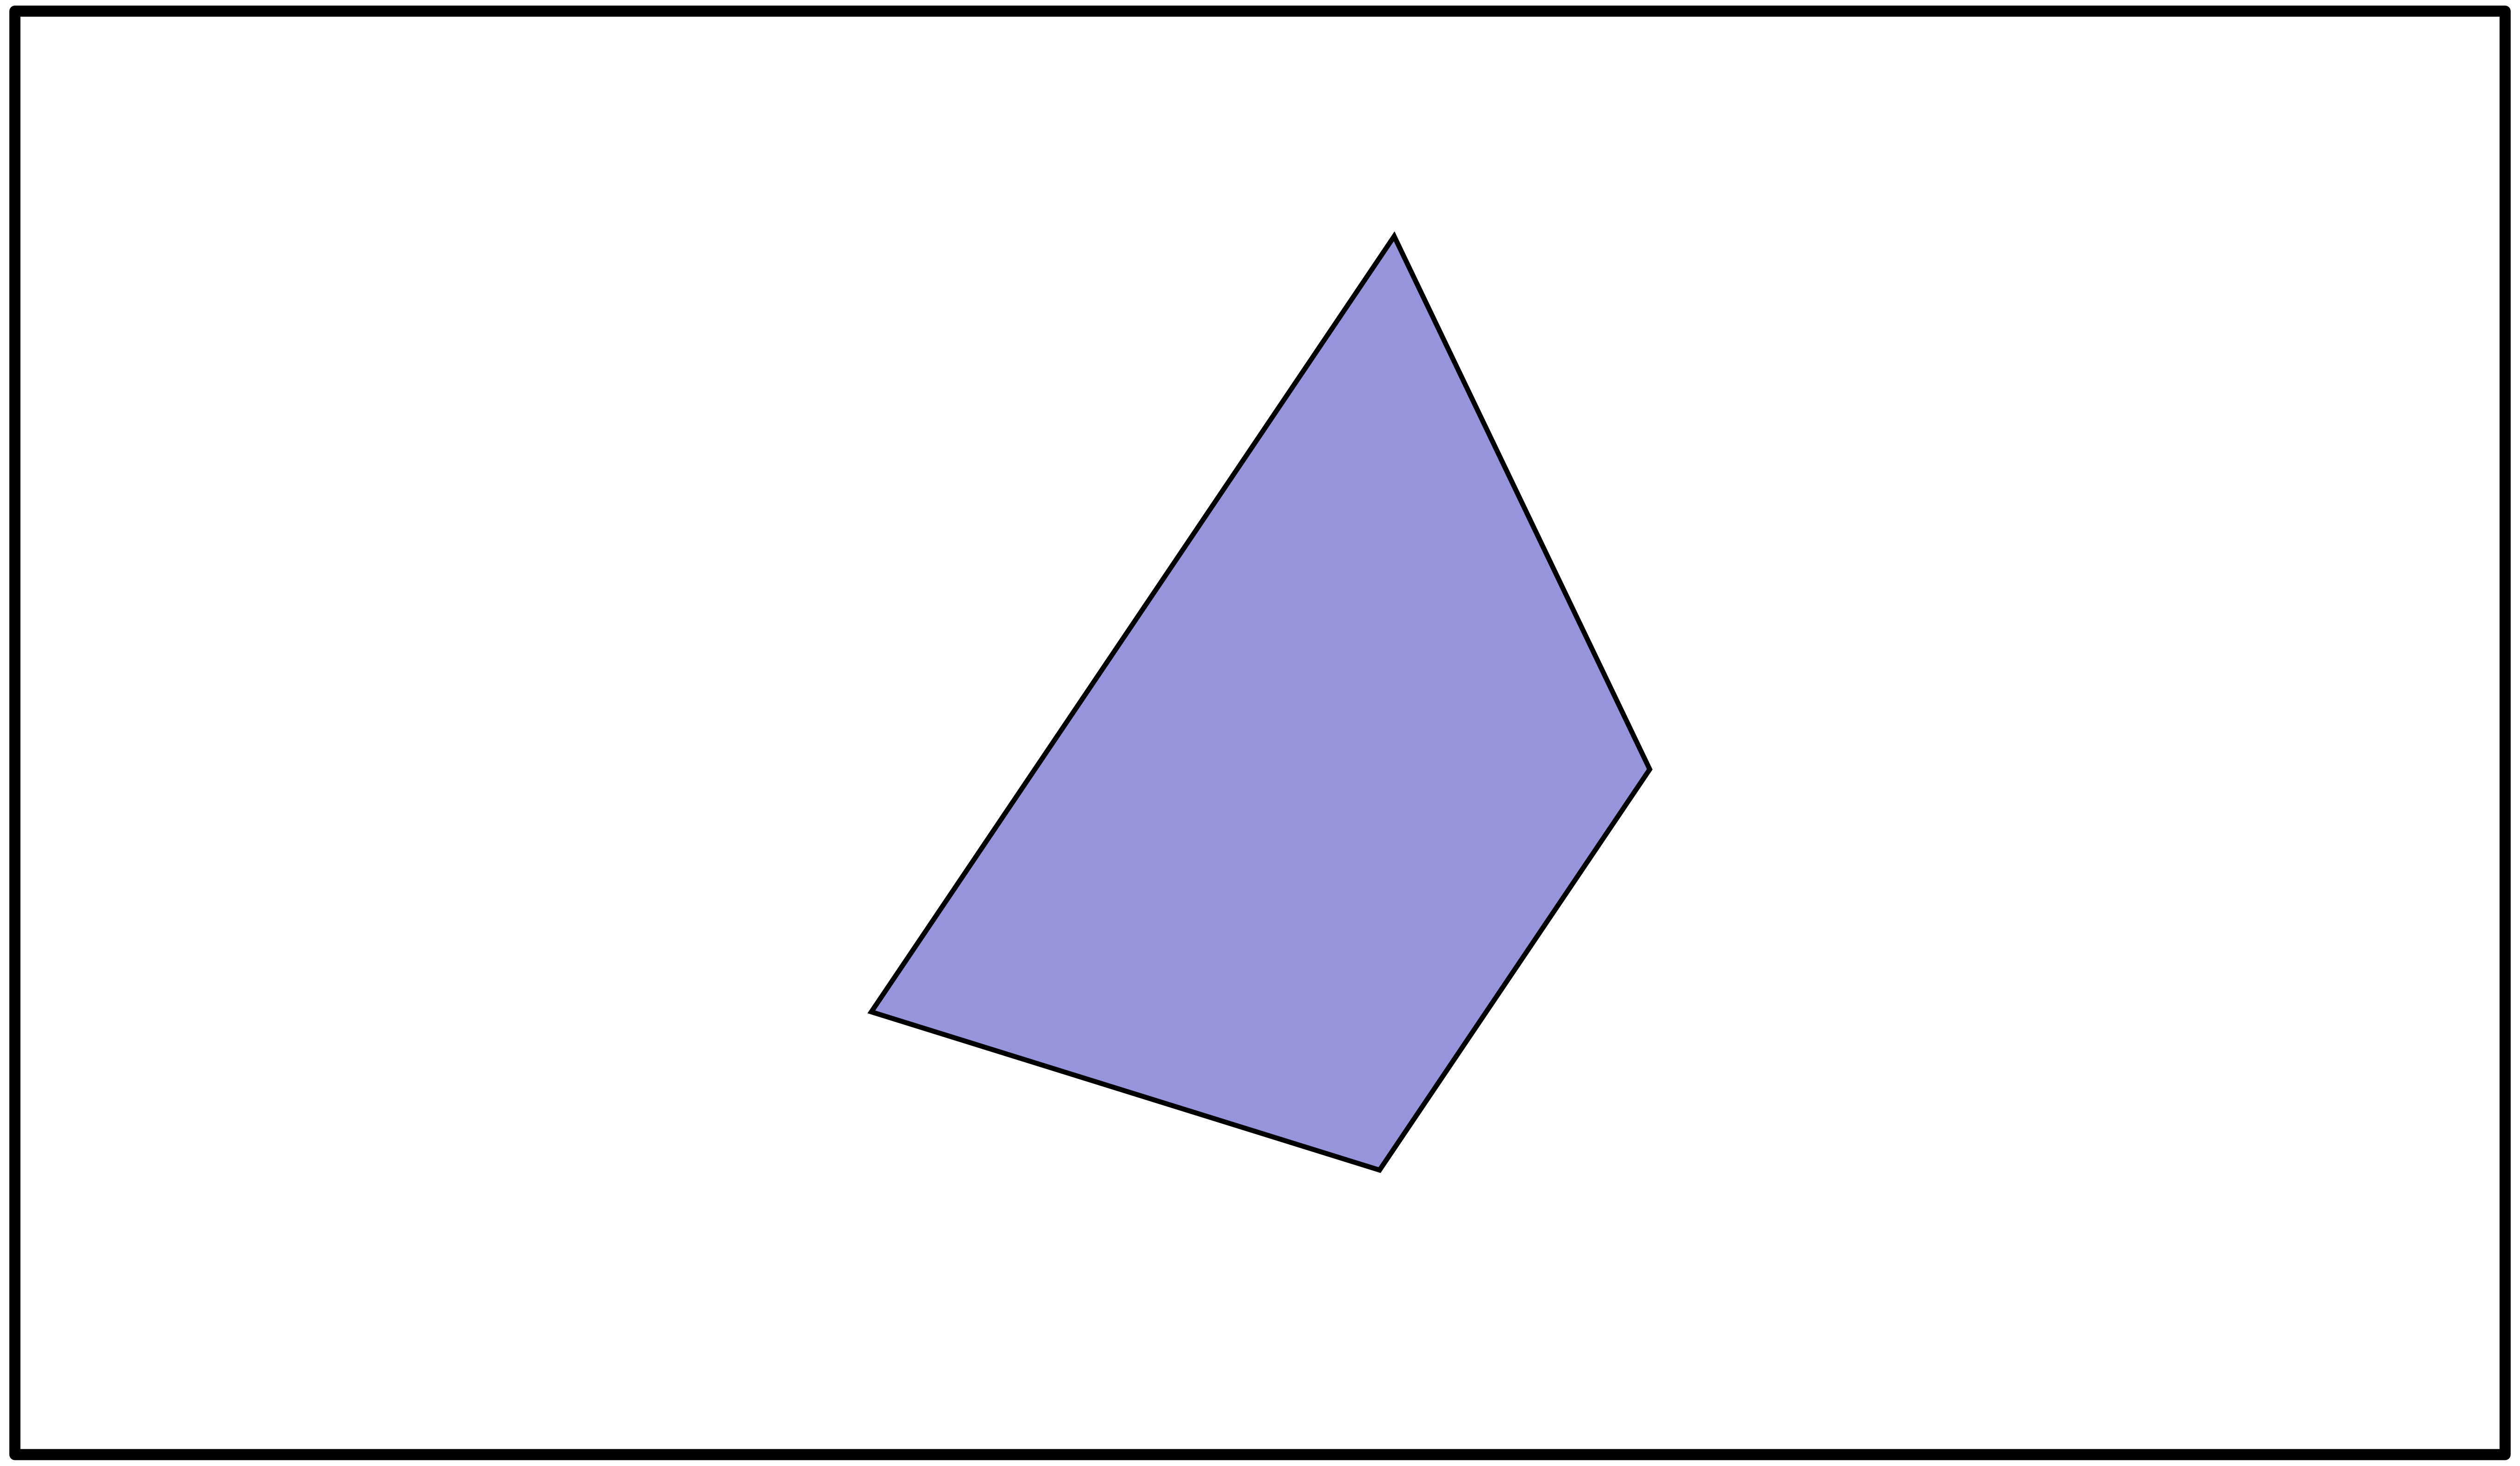

Pick a point in the image.
Find it again in the next image.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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.

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.

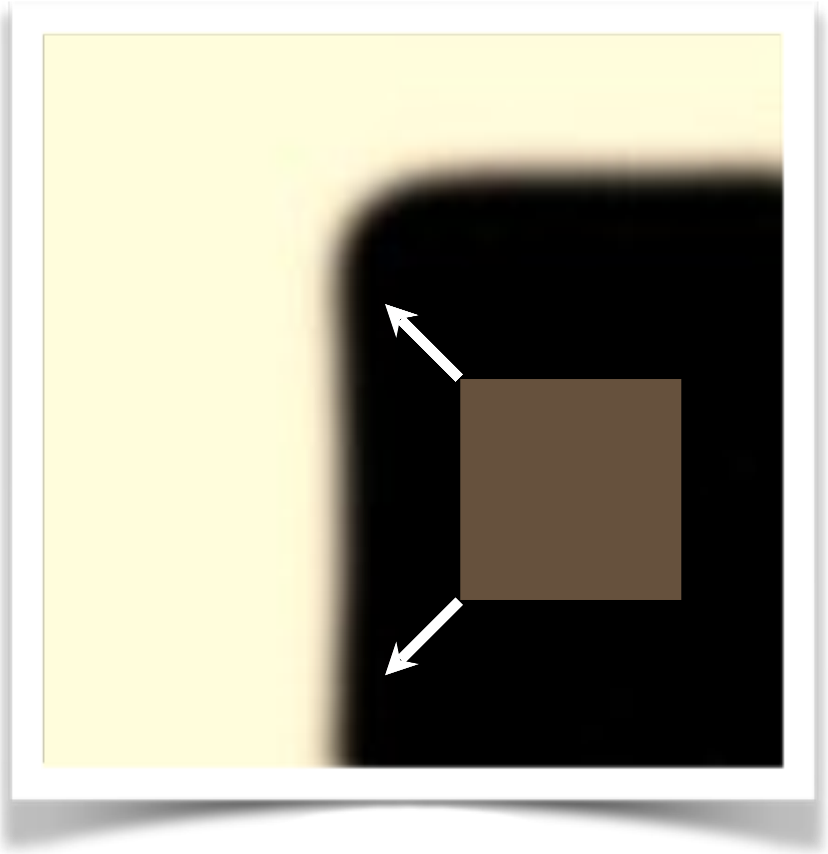

"**flat**" region:

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

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.

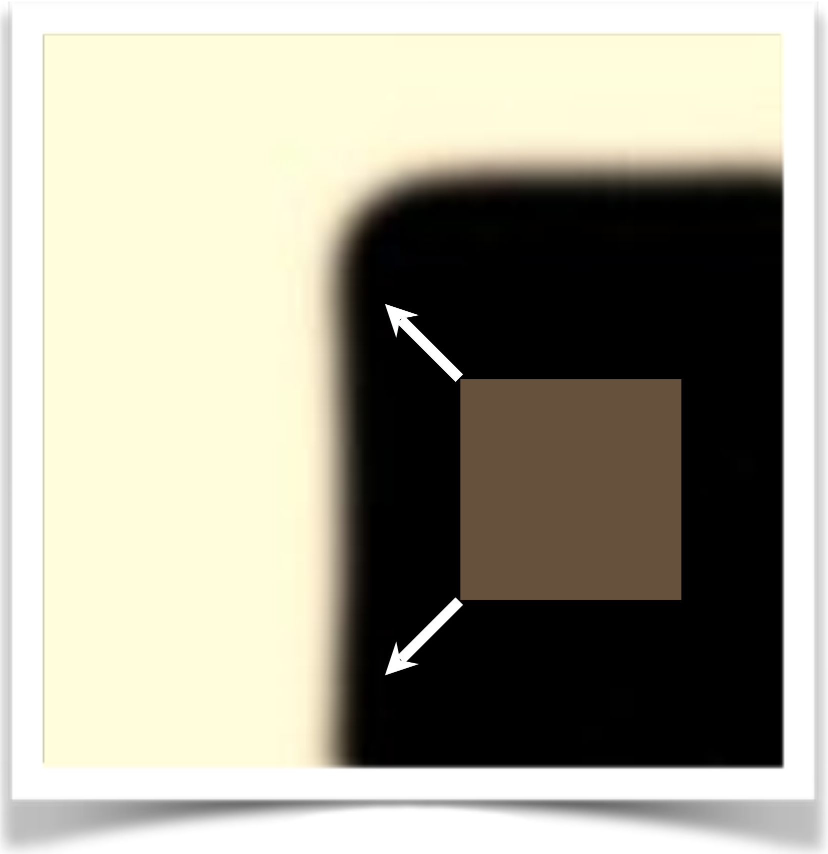

"**flat**" region:
no change in all
directions

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

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.

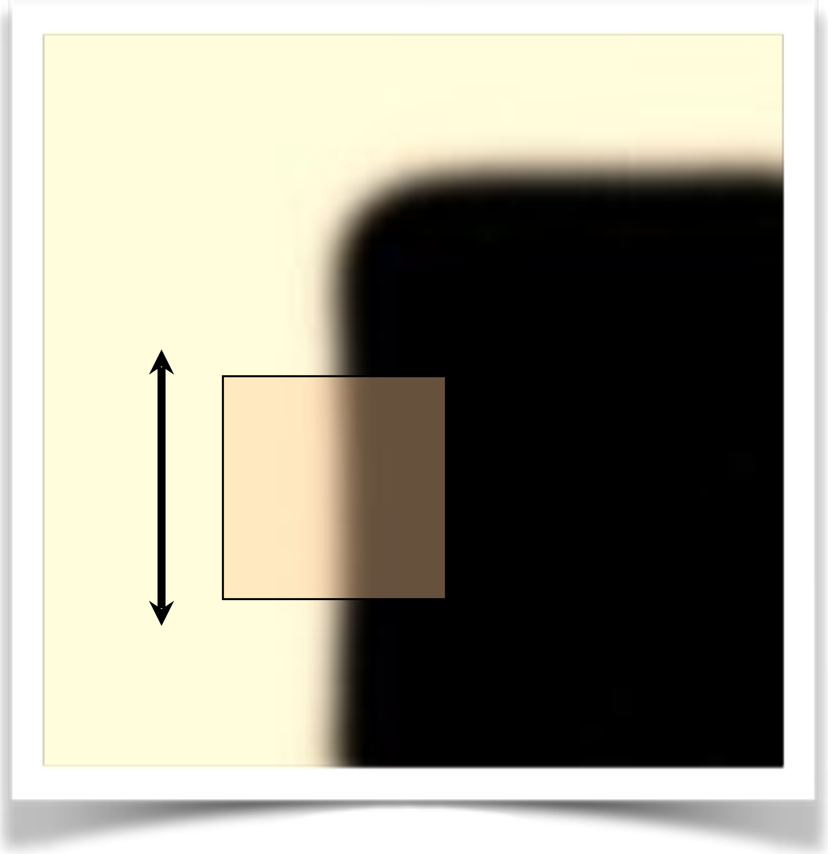

“**edge**”:

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

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.

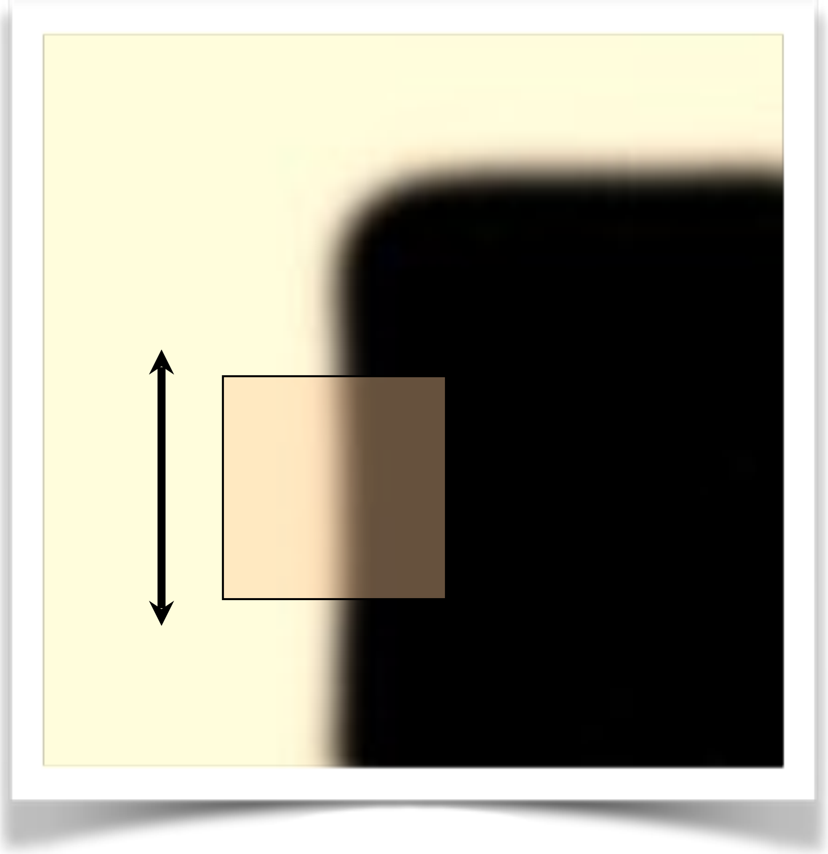

"**edge**": no change along the edge direction

— 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)

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

Why are corners **distinct**?

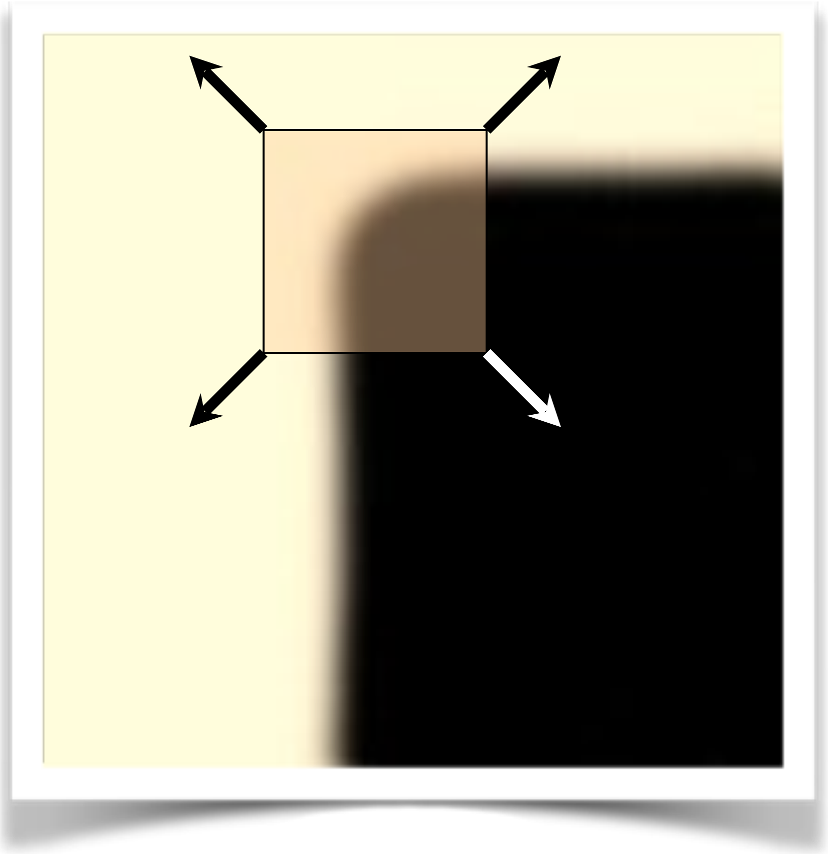

"**corner**":

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.

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

Why are corners **distinct**?

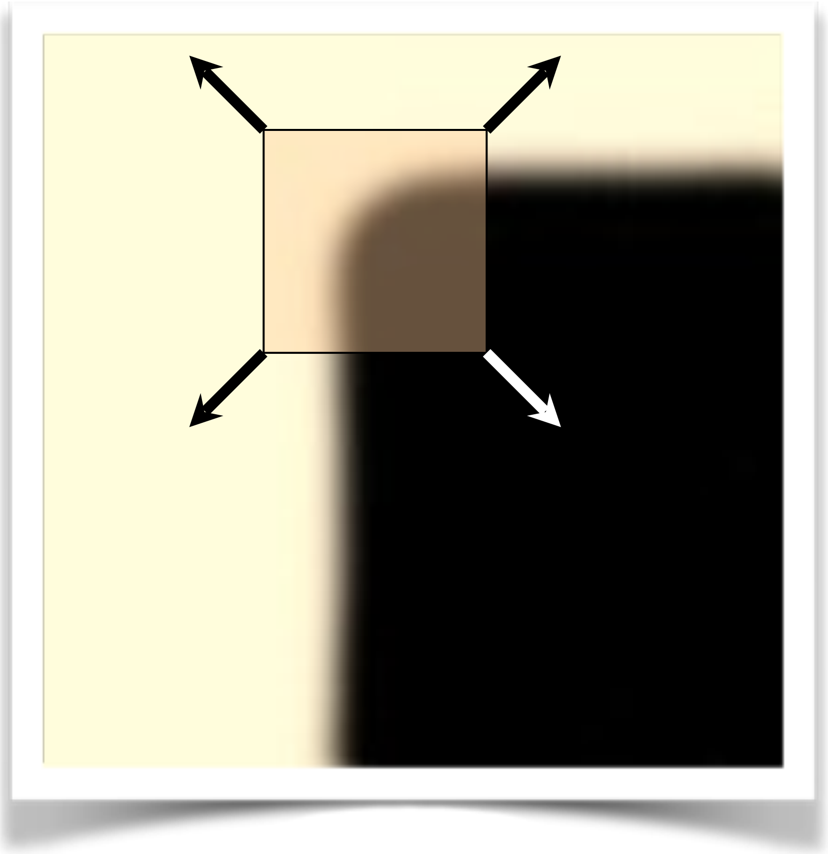

"**corner**":
significant change
in all directions

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.

Image Credit: Ioannis (Yannis) Gkioulekas (CMU)

How do you find a **corner**?

[Moravec 1980]

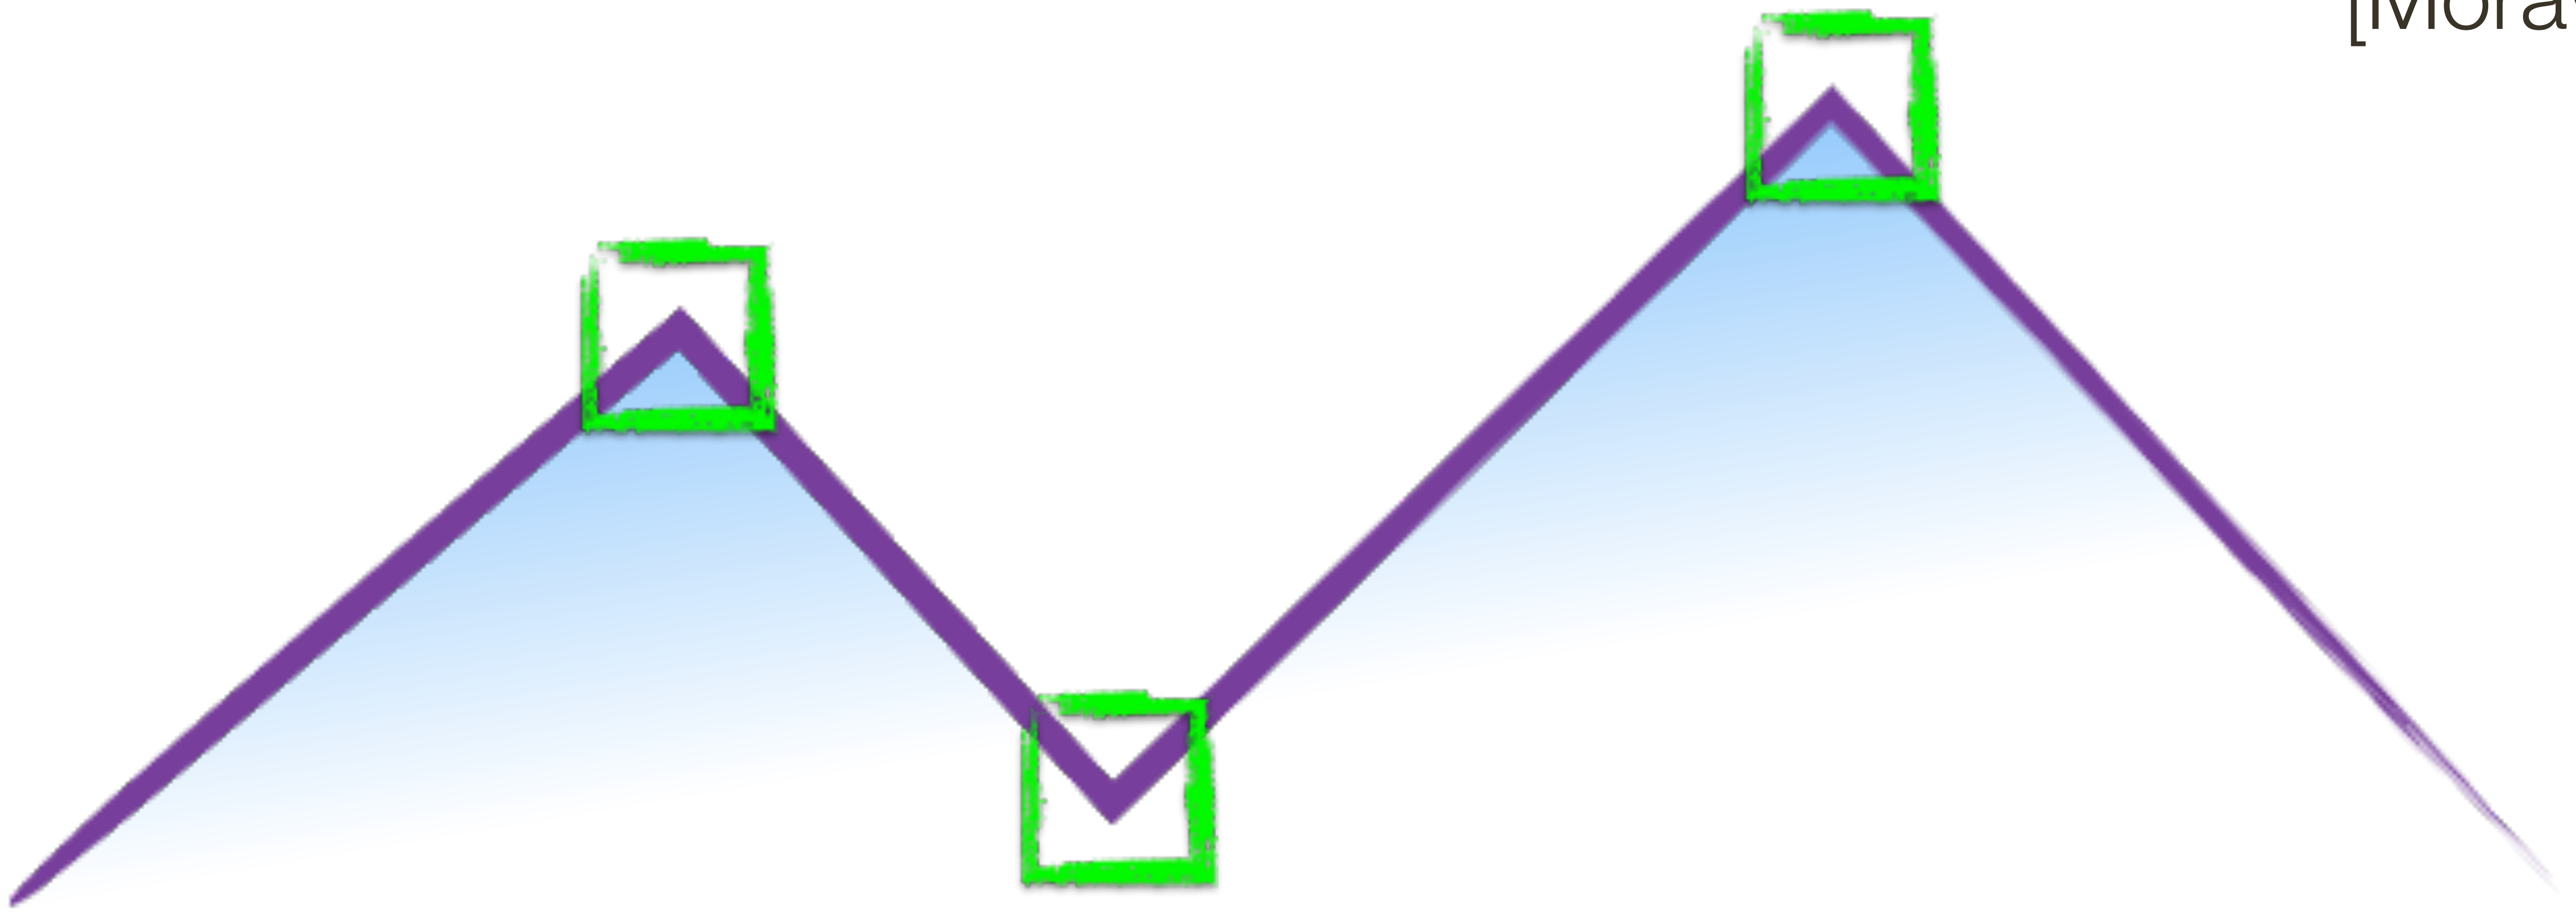

Easily recognized by looking through a small window

Shifting the window should give large change in intensity

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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 of rapidly in all directions.

Autocorrelation

Szeliski, Figure 4.5

Autocorrelation

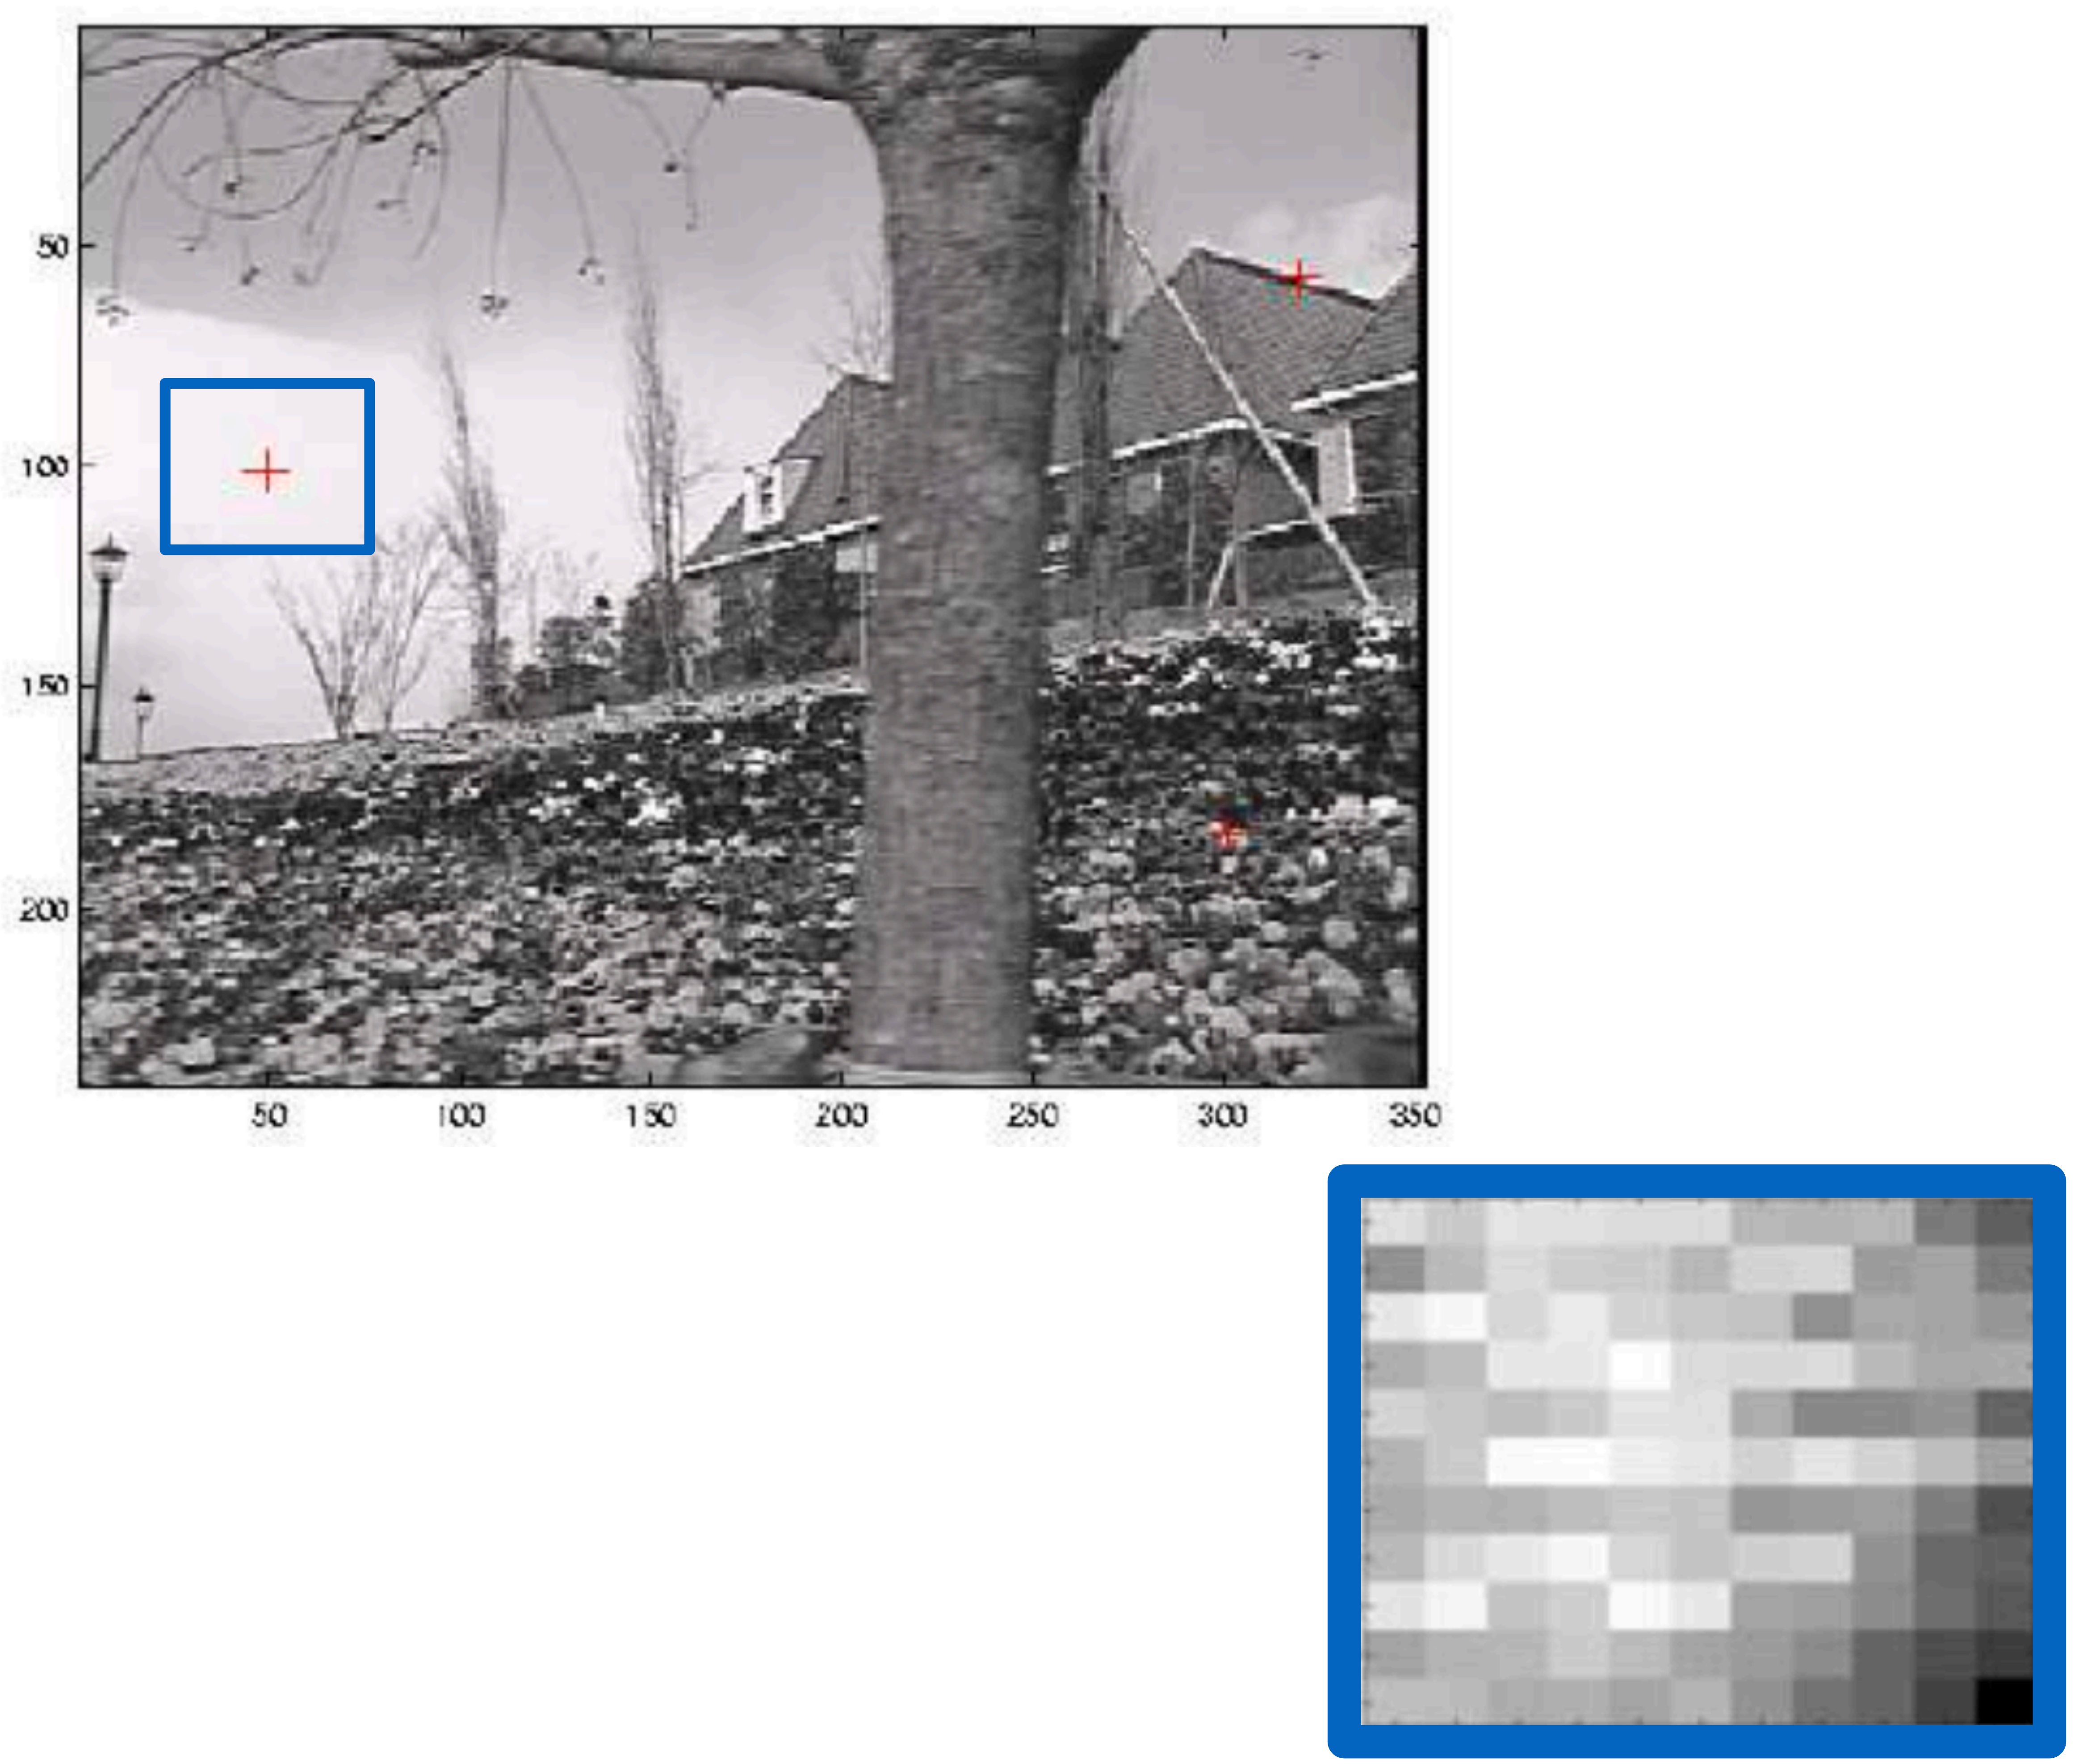

Szeliski, Figure 4.5

Autocorrelation

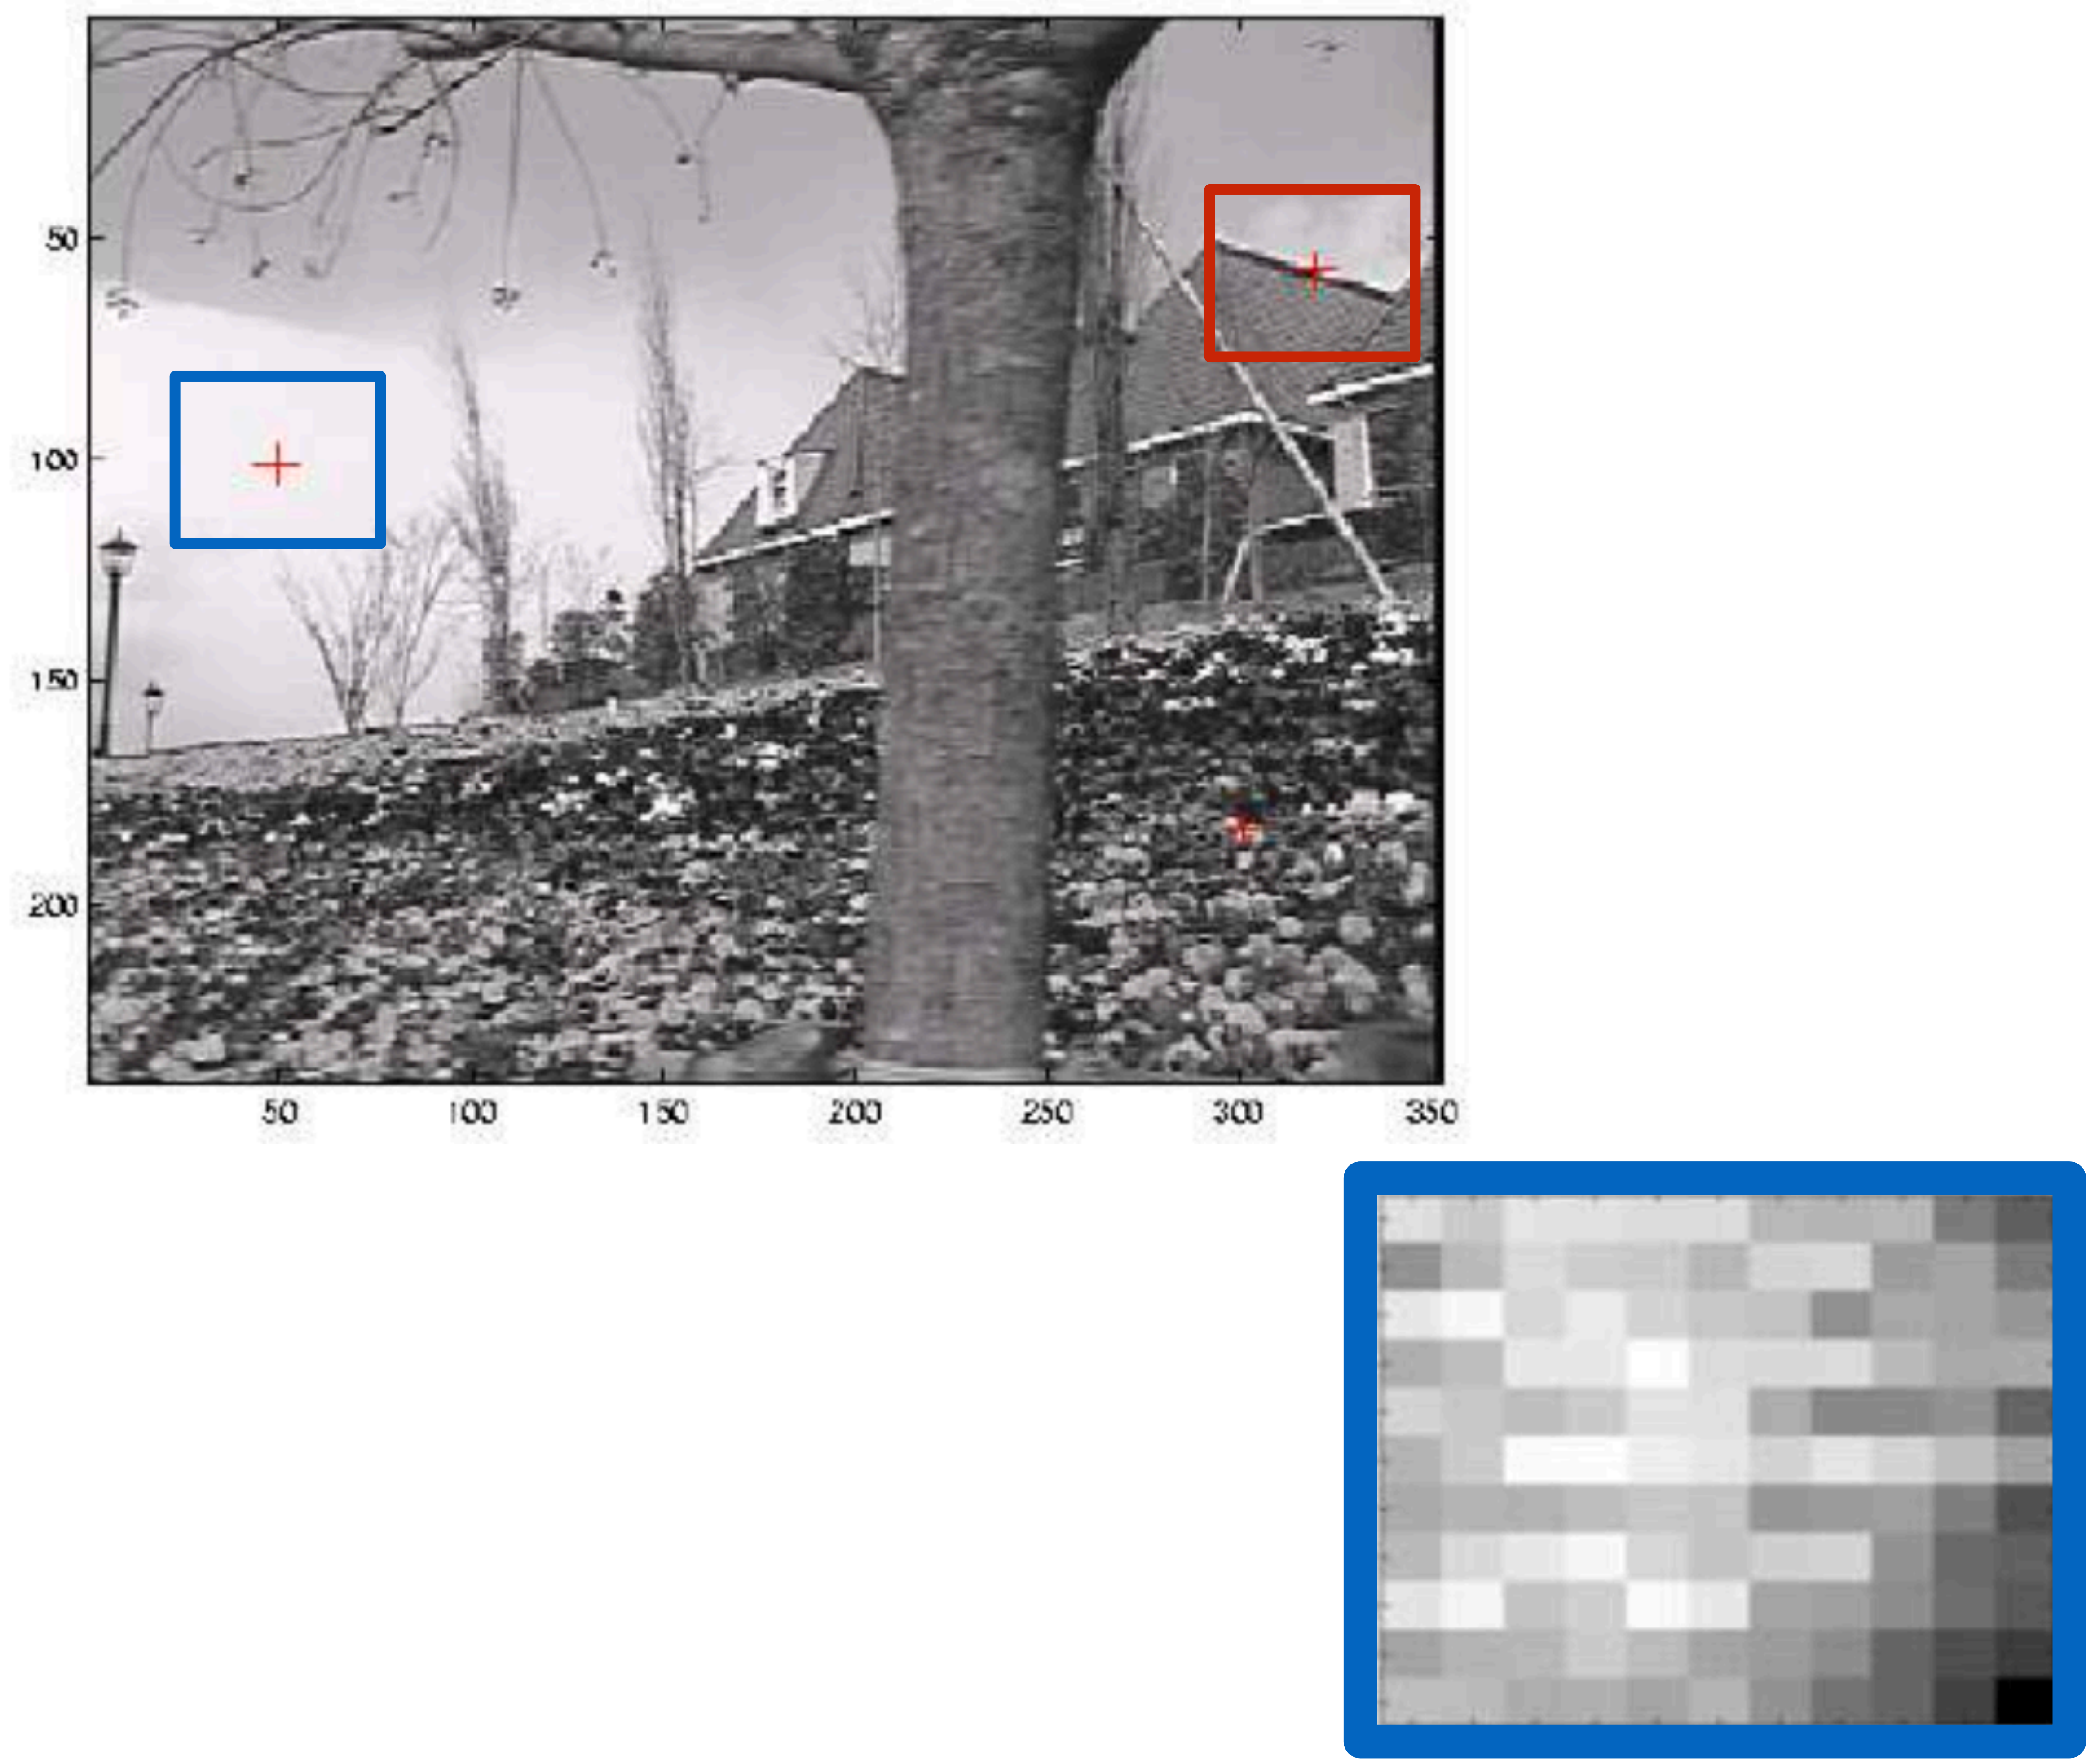

Szeliski, Figure 4.5

Autocorrelation

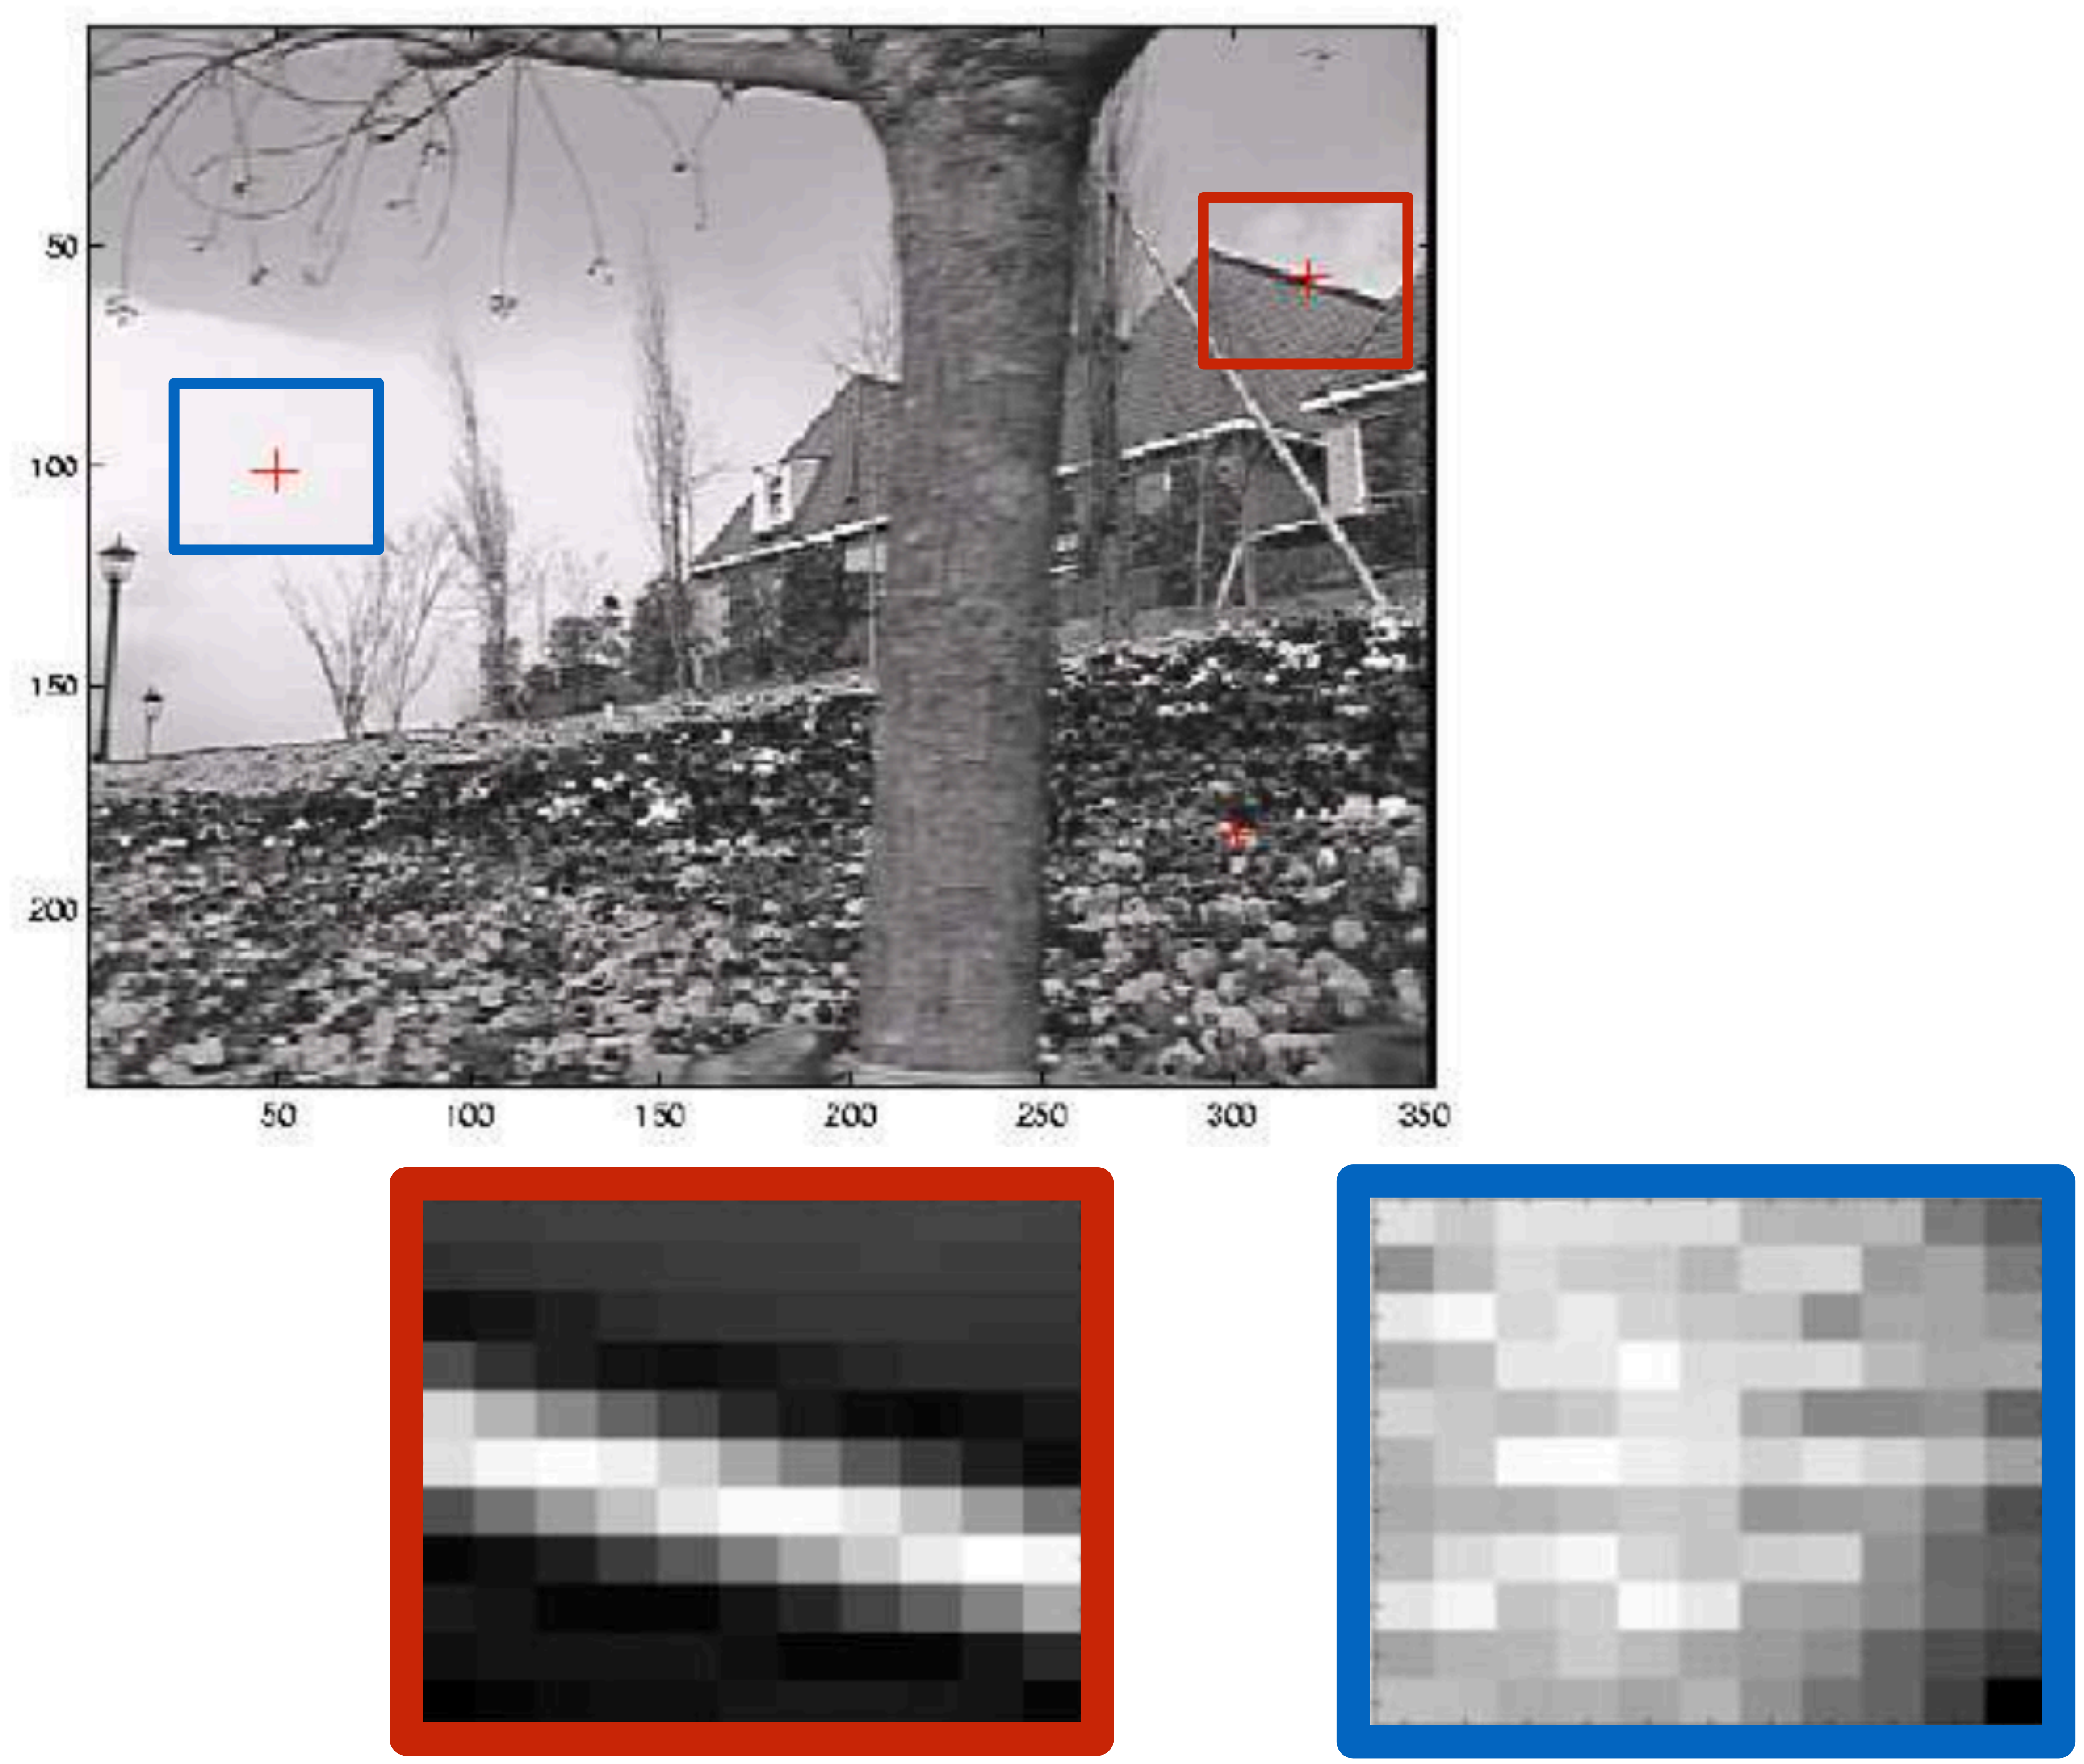

Szeliski, Figure 4.5

Autocorrelation

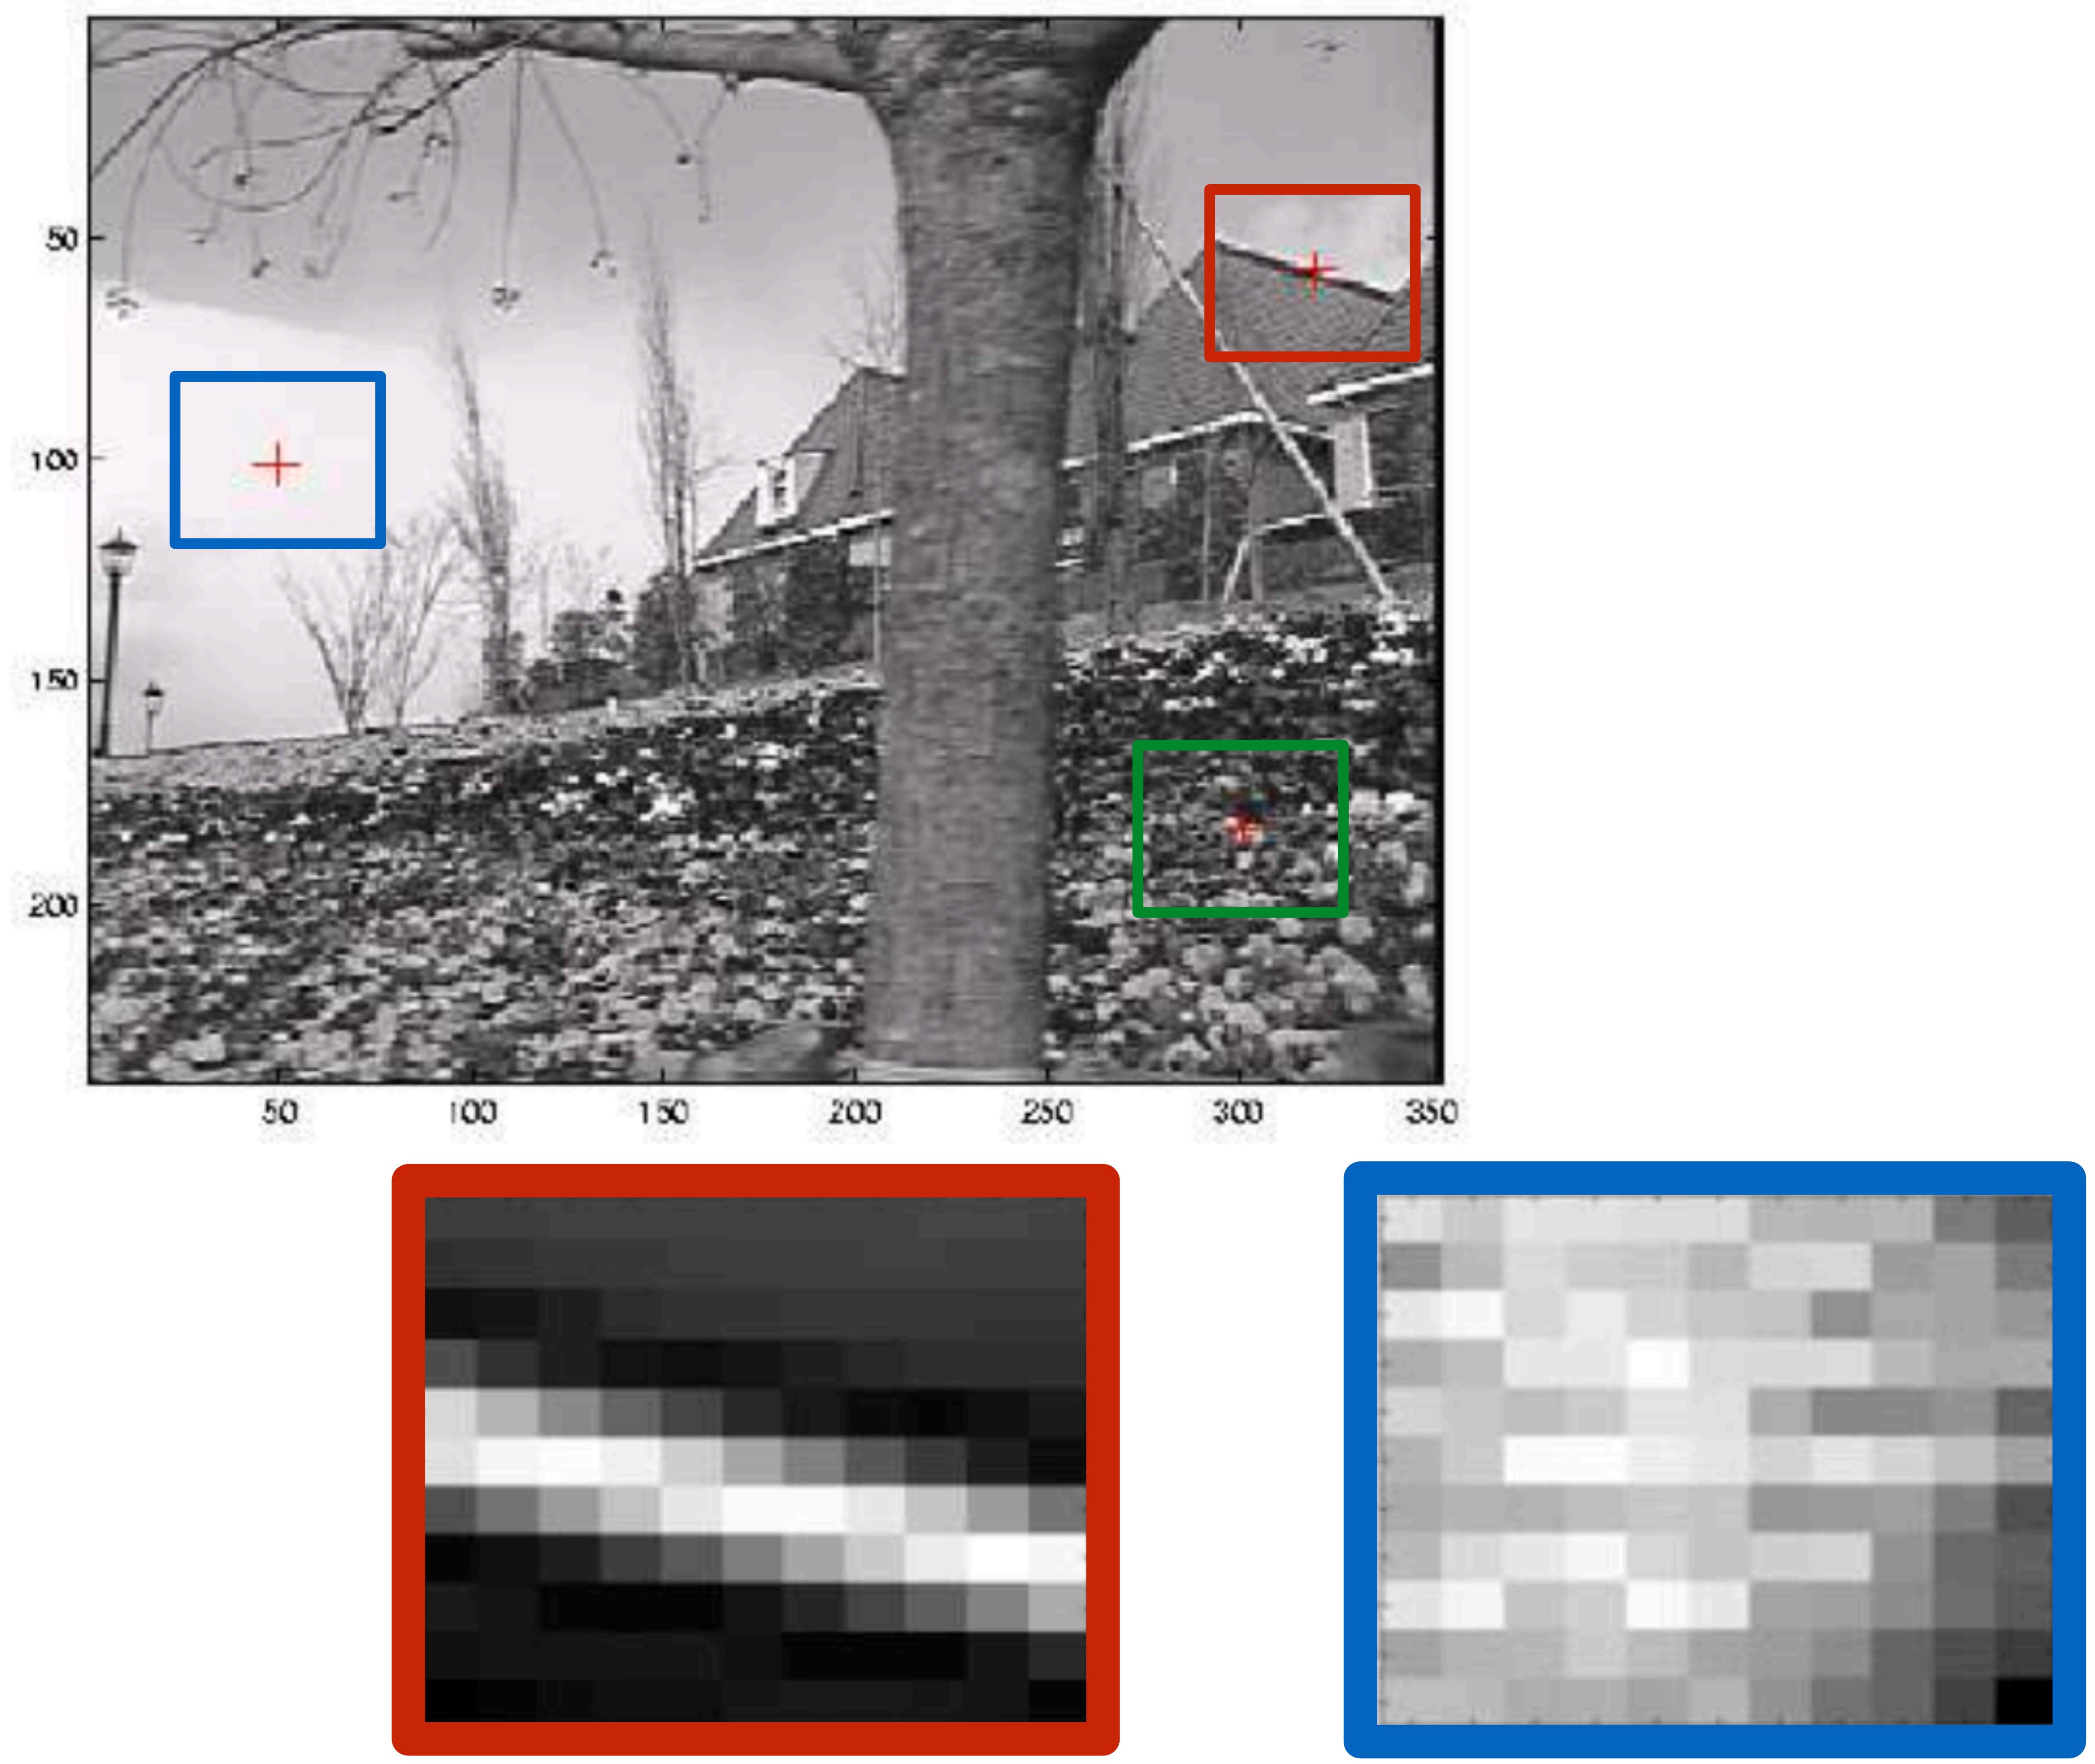

Szeliski, Figure 4.5

Autocorrelation

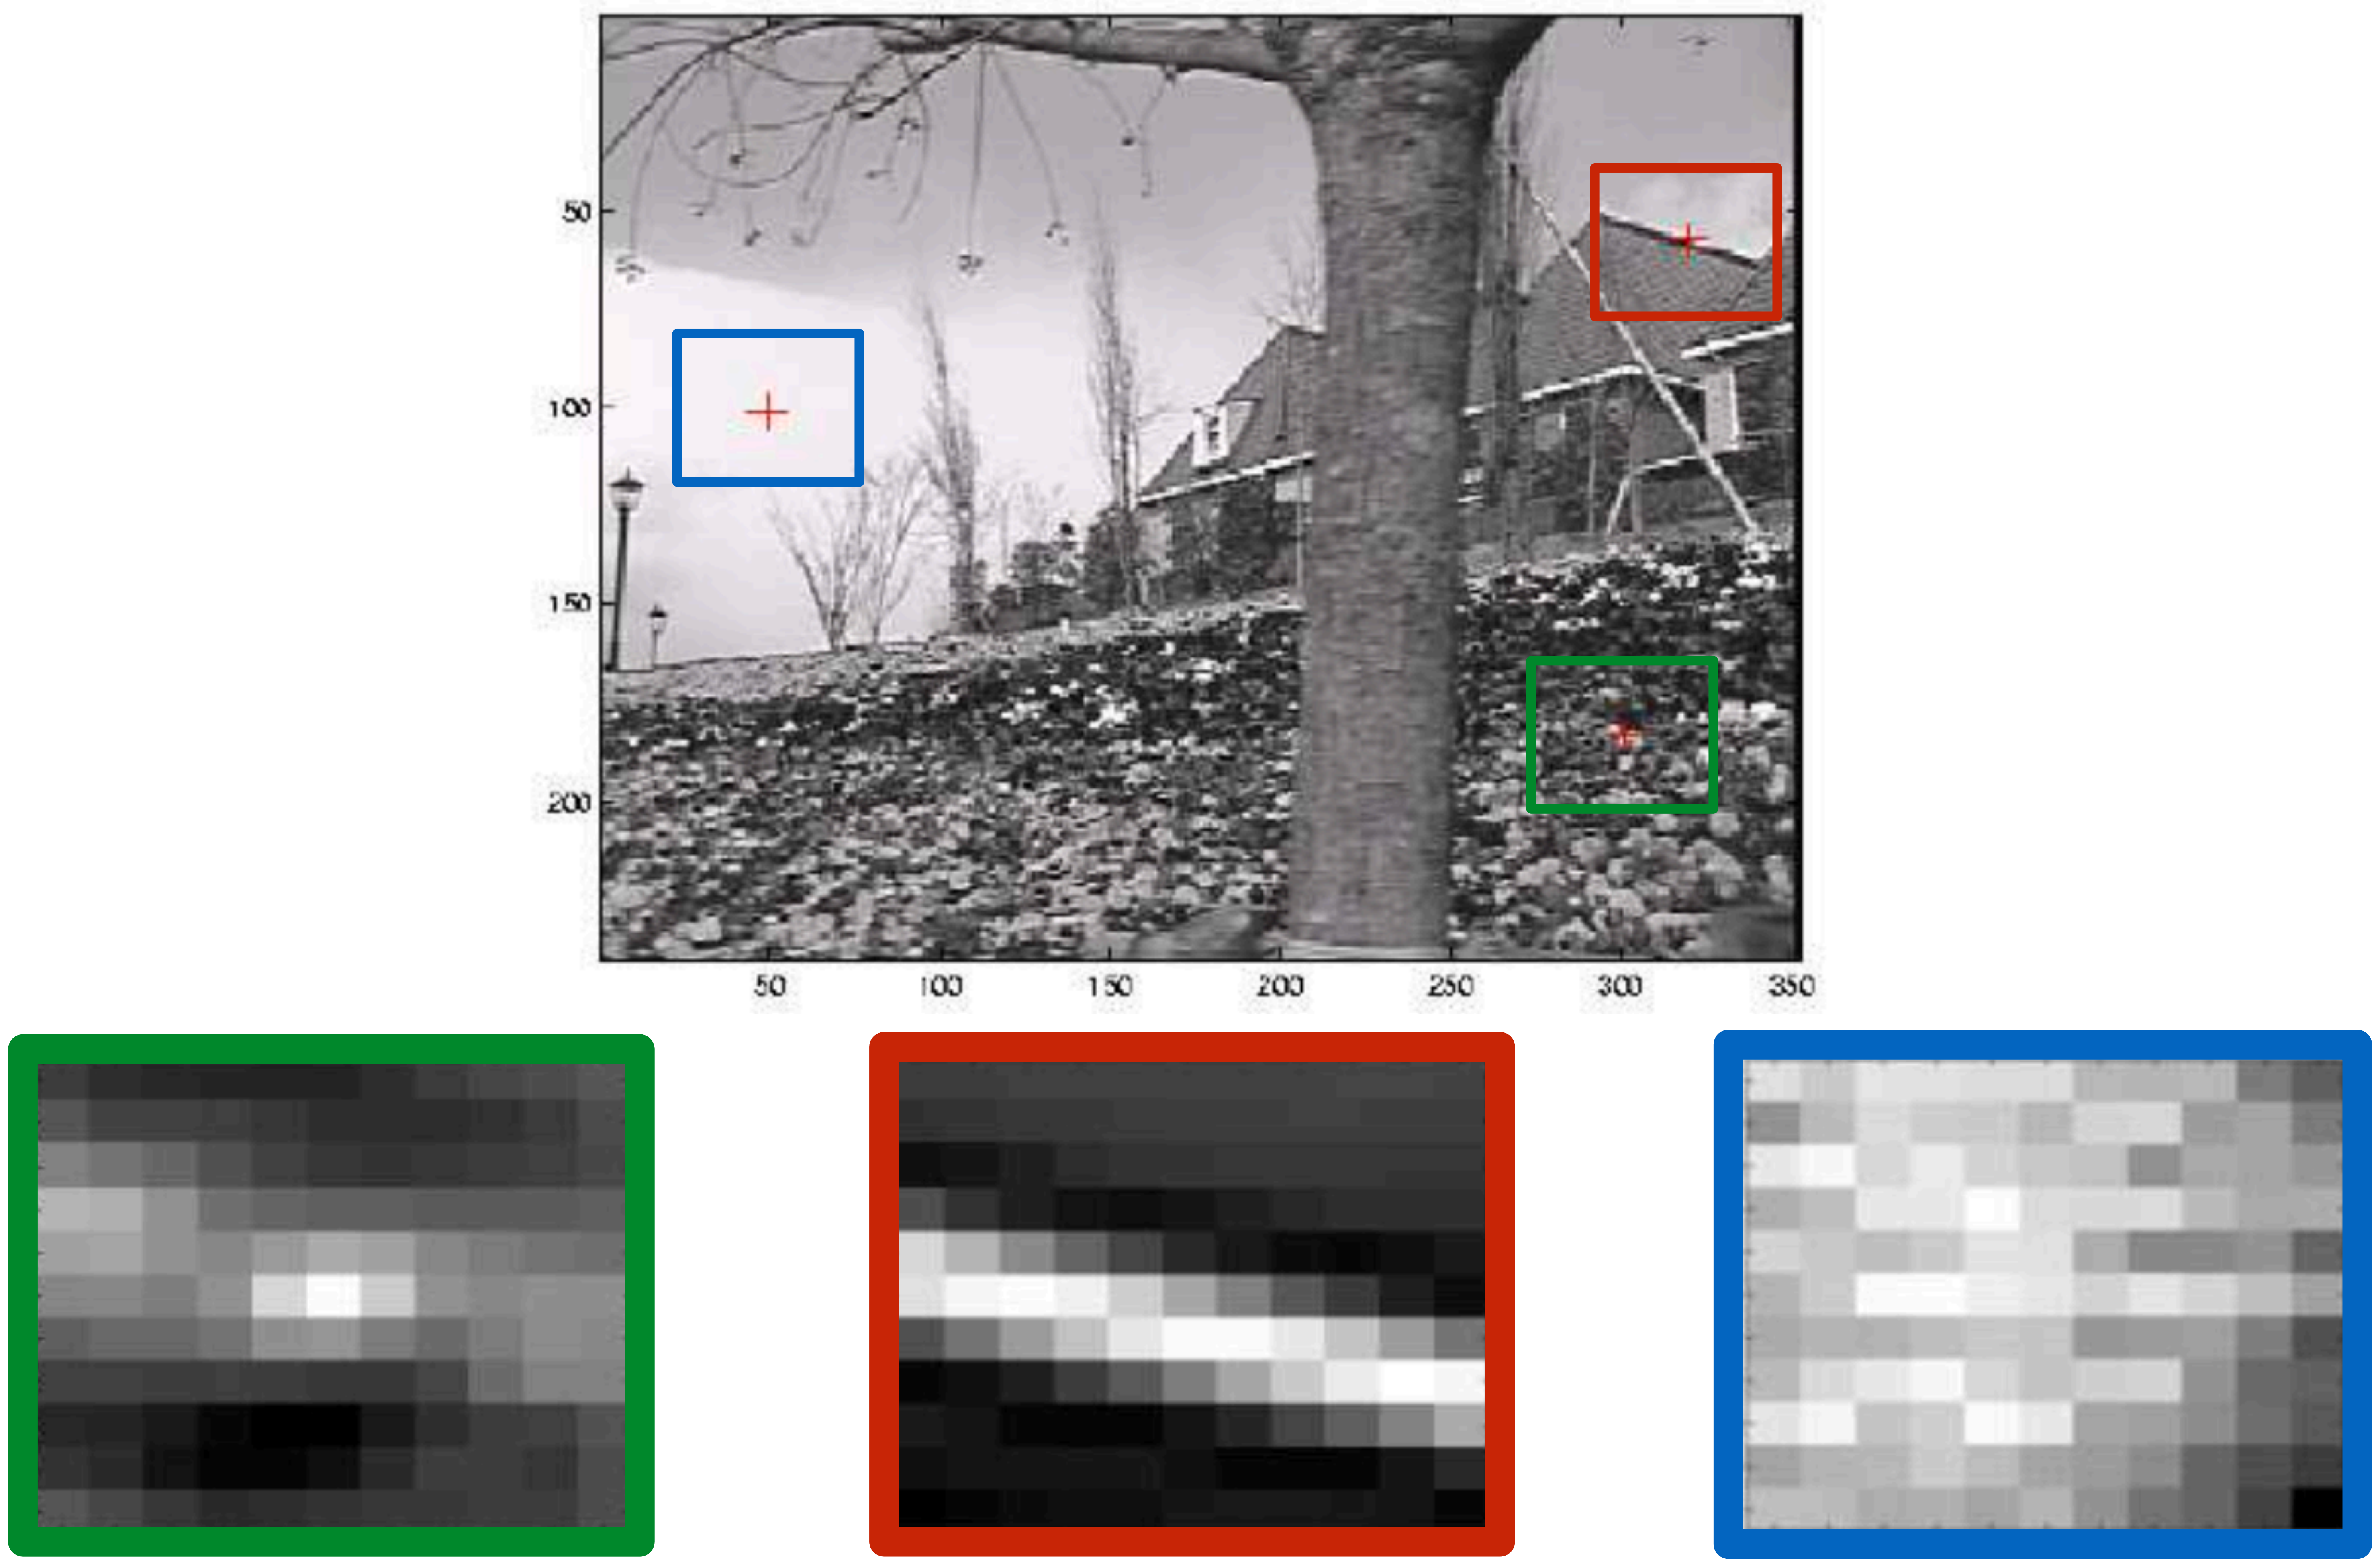

Szeliski, Figure 4.5

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 of rapidly 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

Harris Corner Detection

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

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

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

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

Slide Adopted: Ioannis (Yannis) Gkioulekas (CMU)

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

(not just a single pixel)

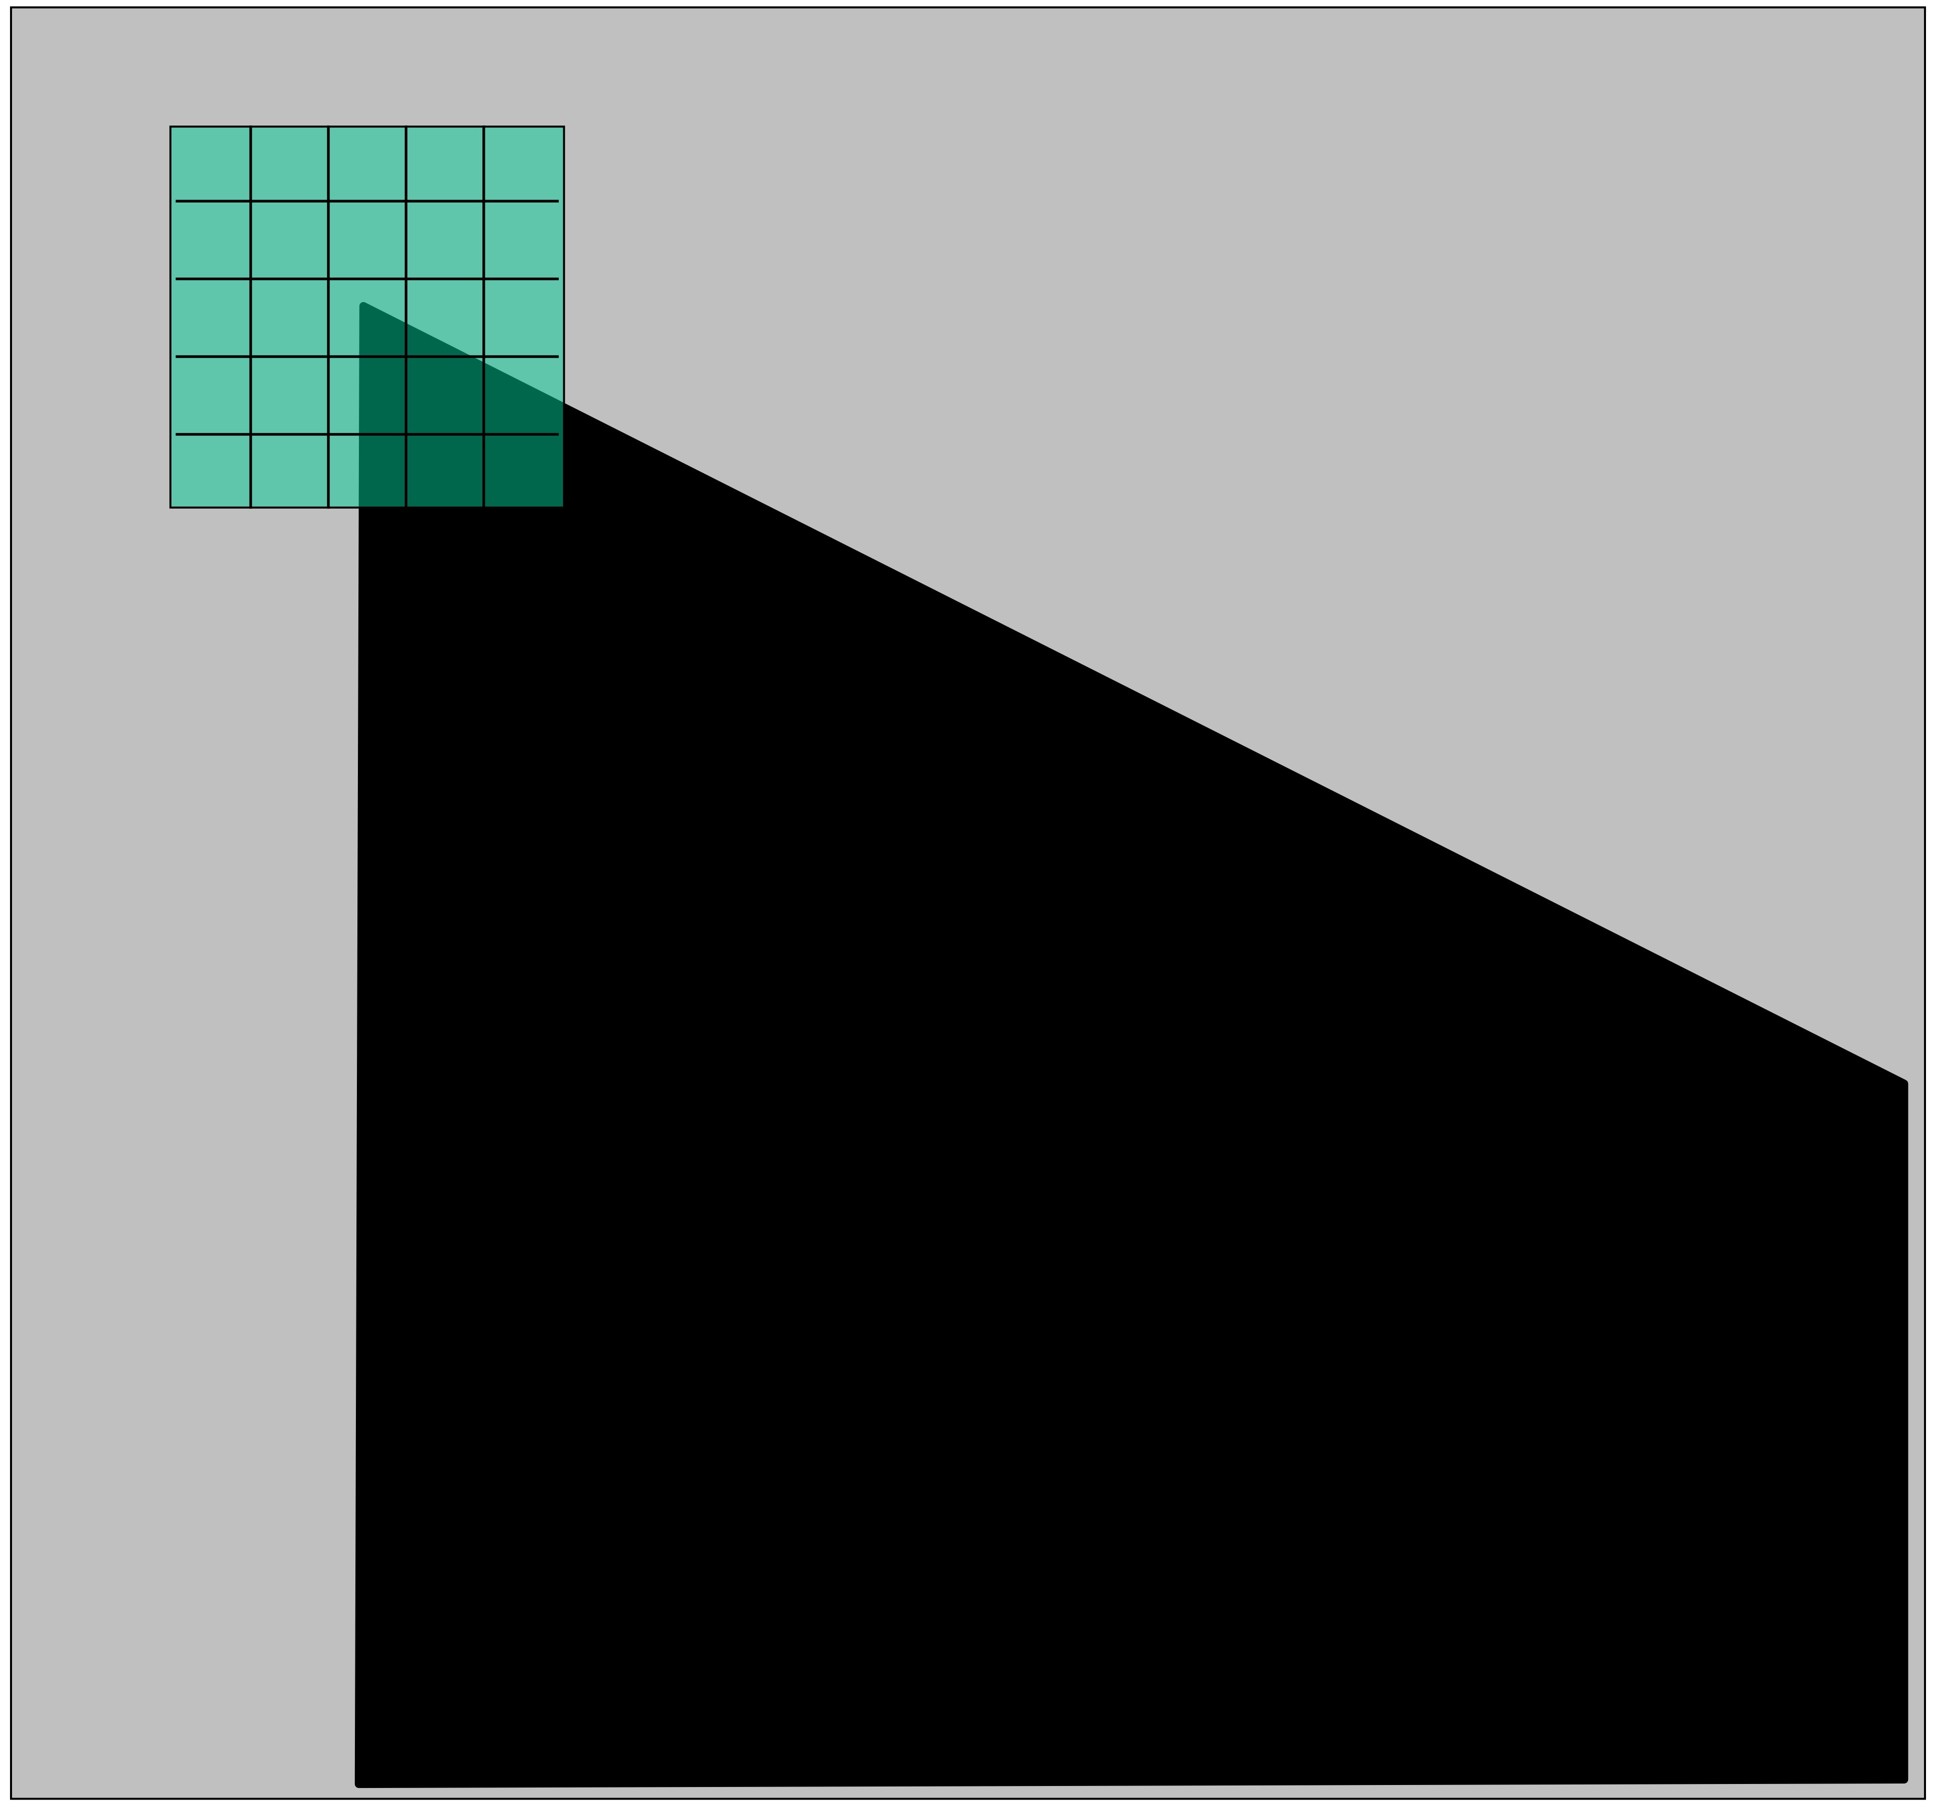

array of x gradients

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

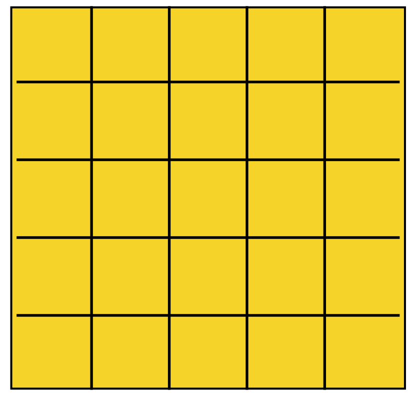

array of y gradients

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

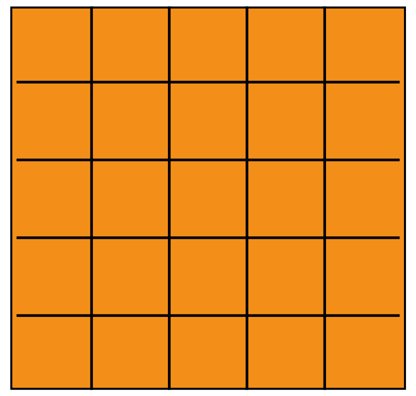

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Visualization of Gradients

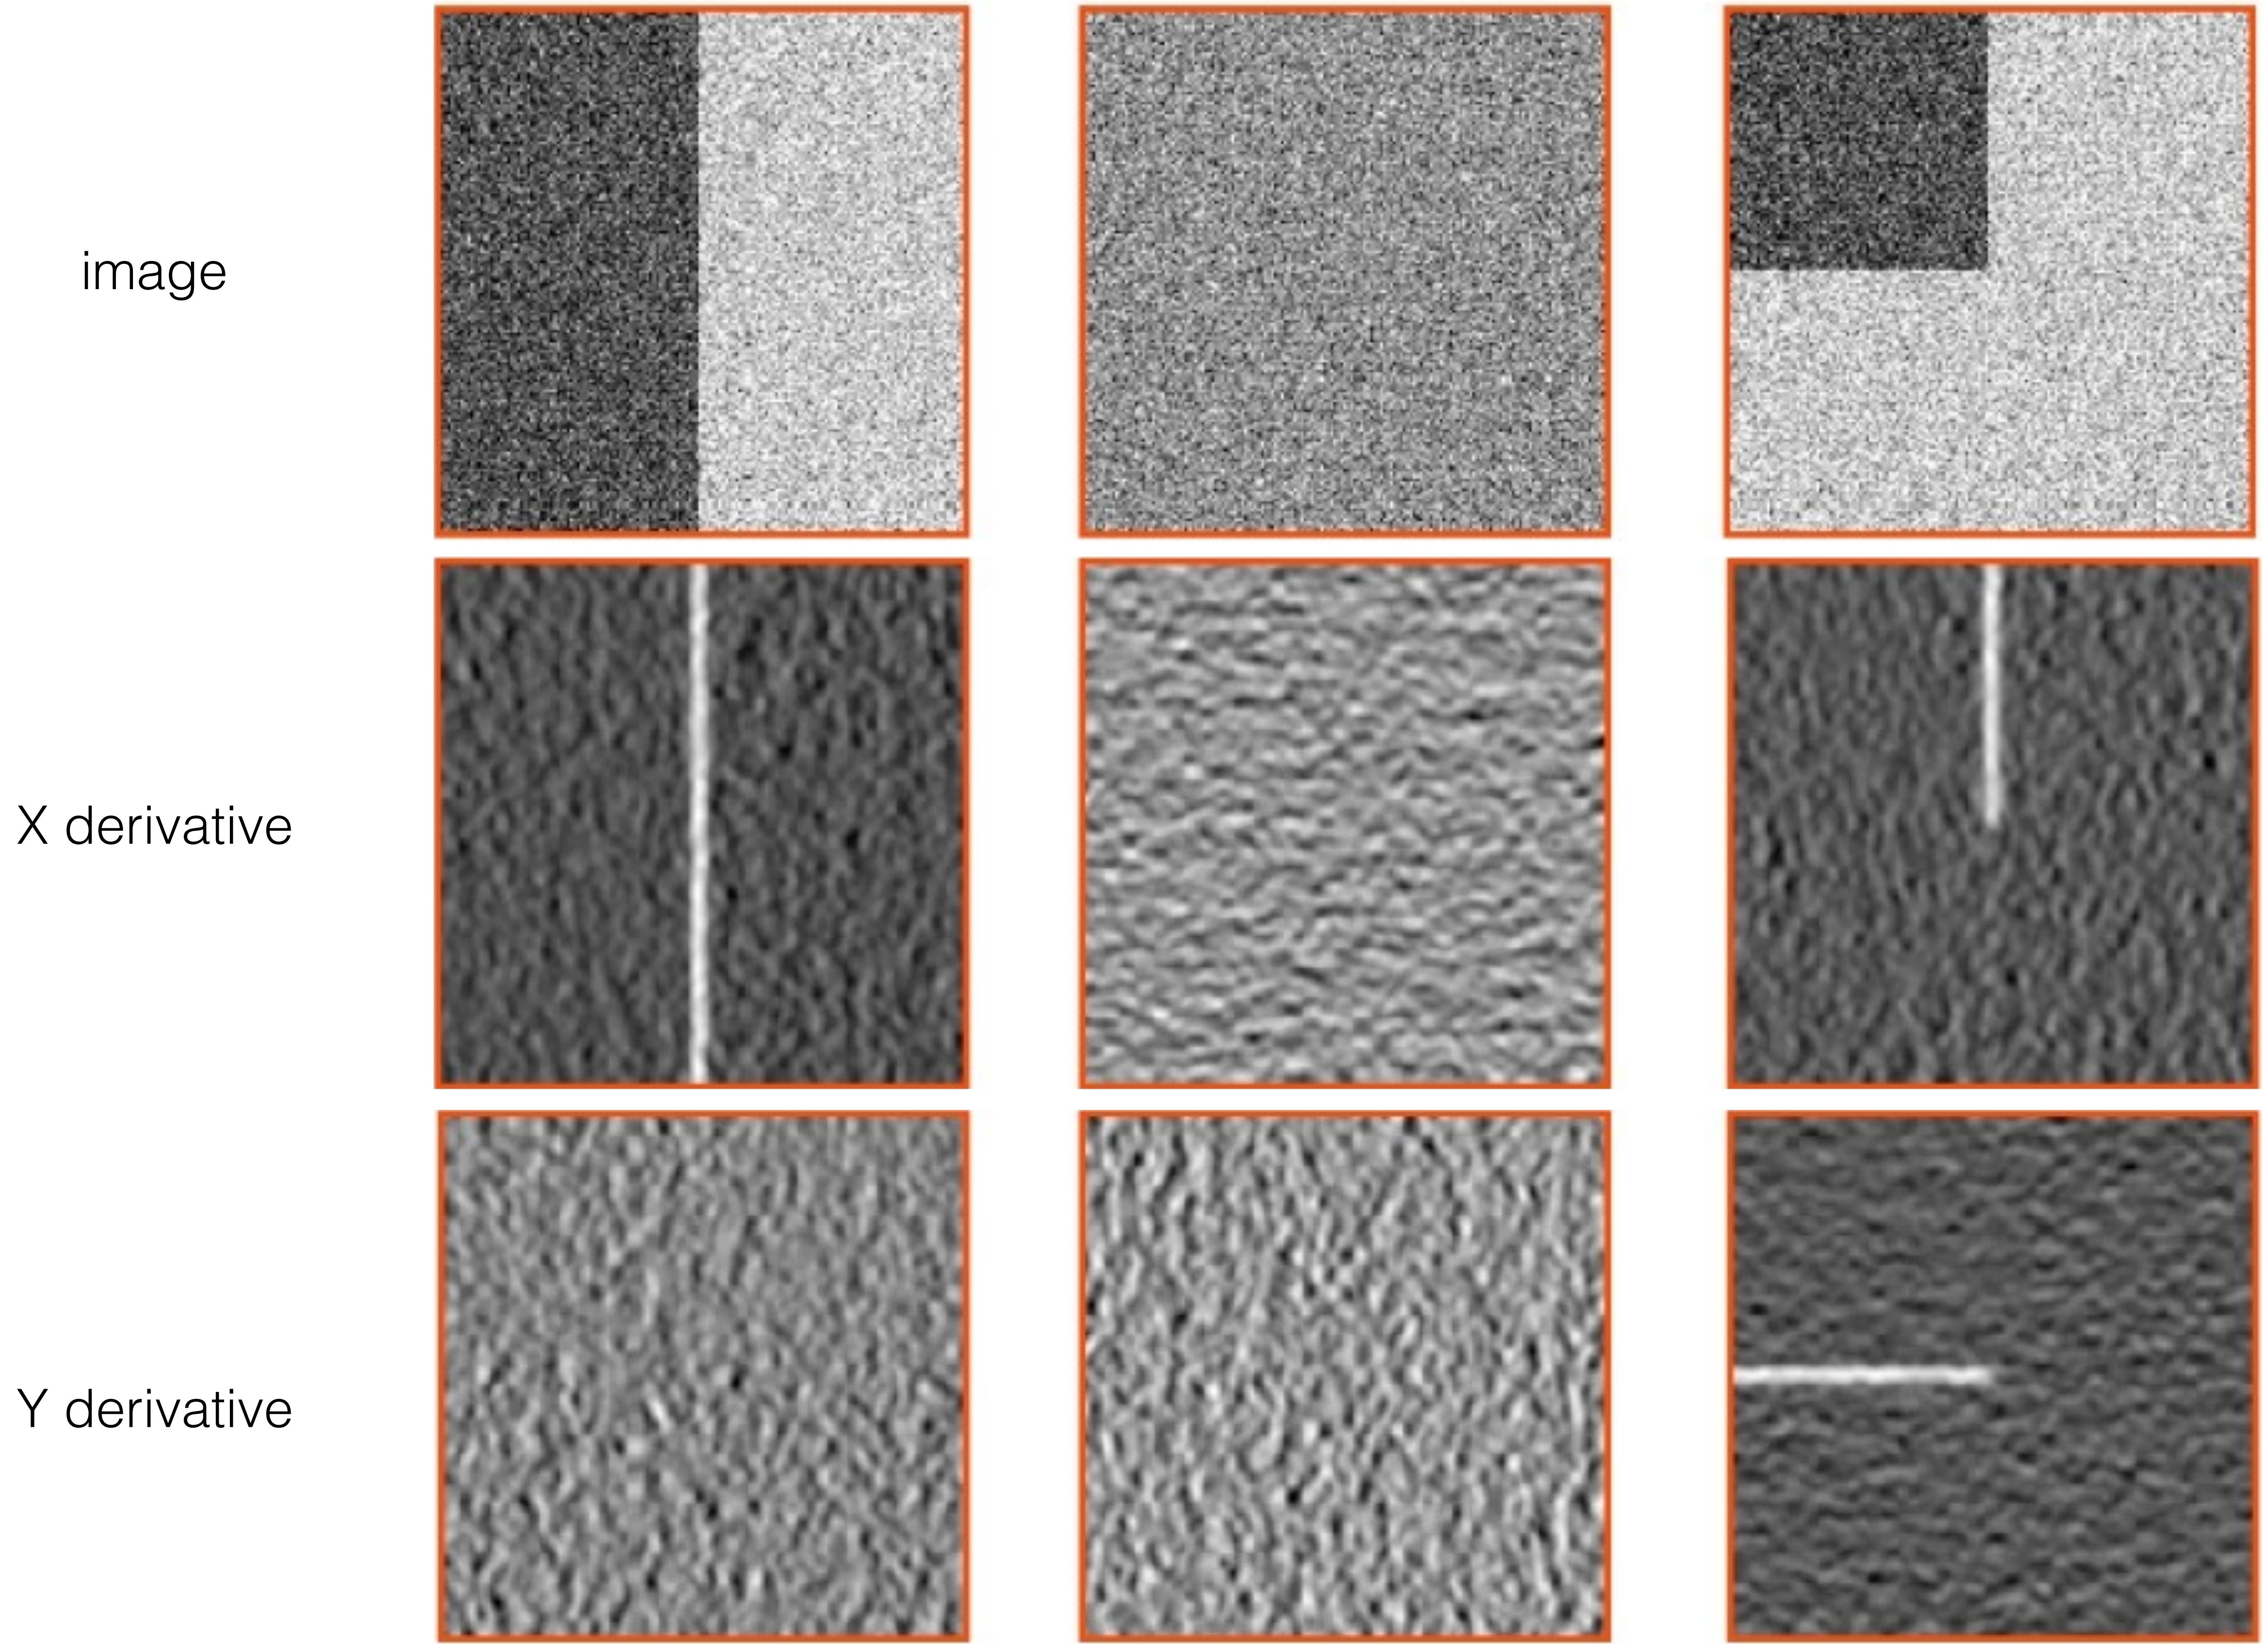

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

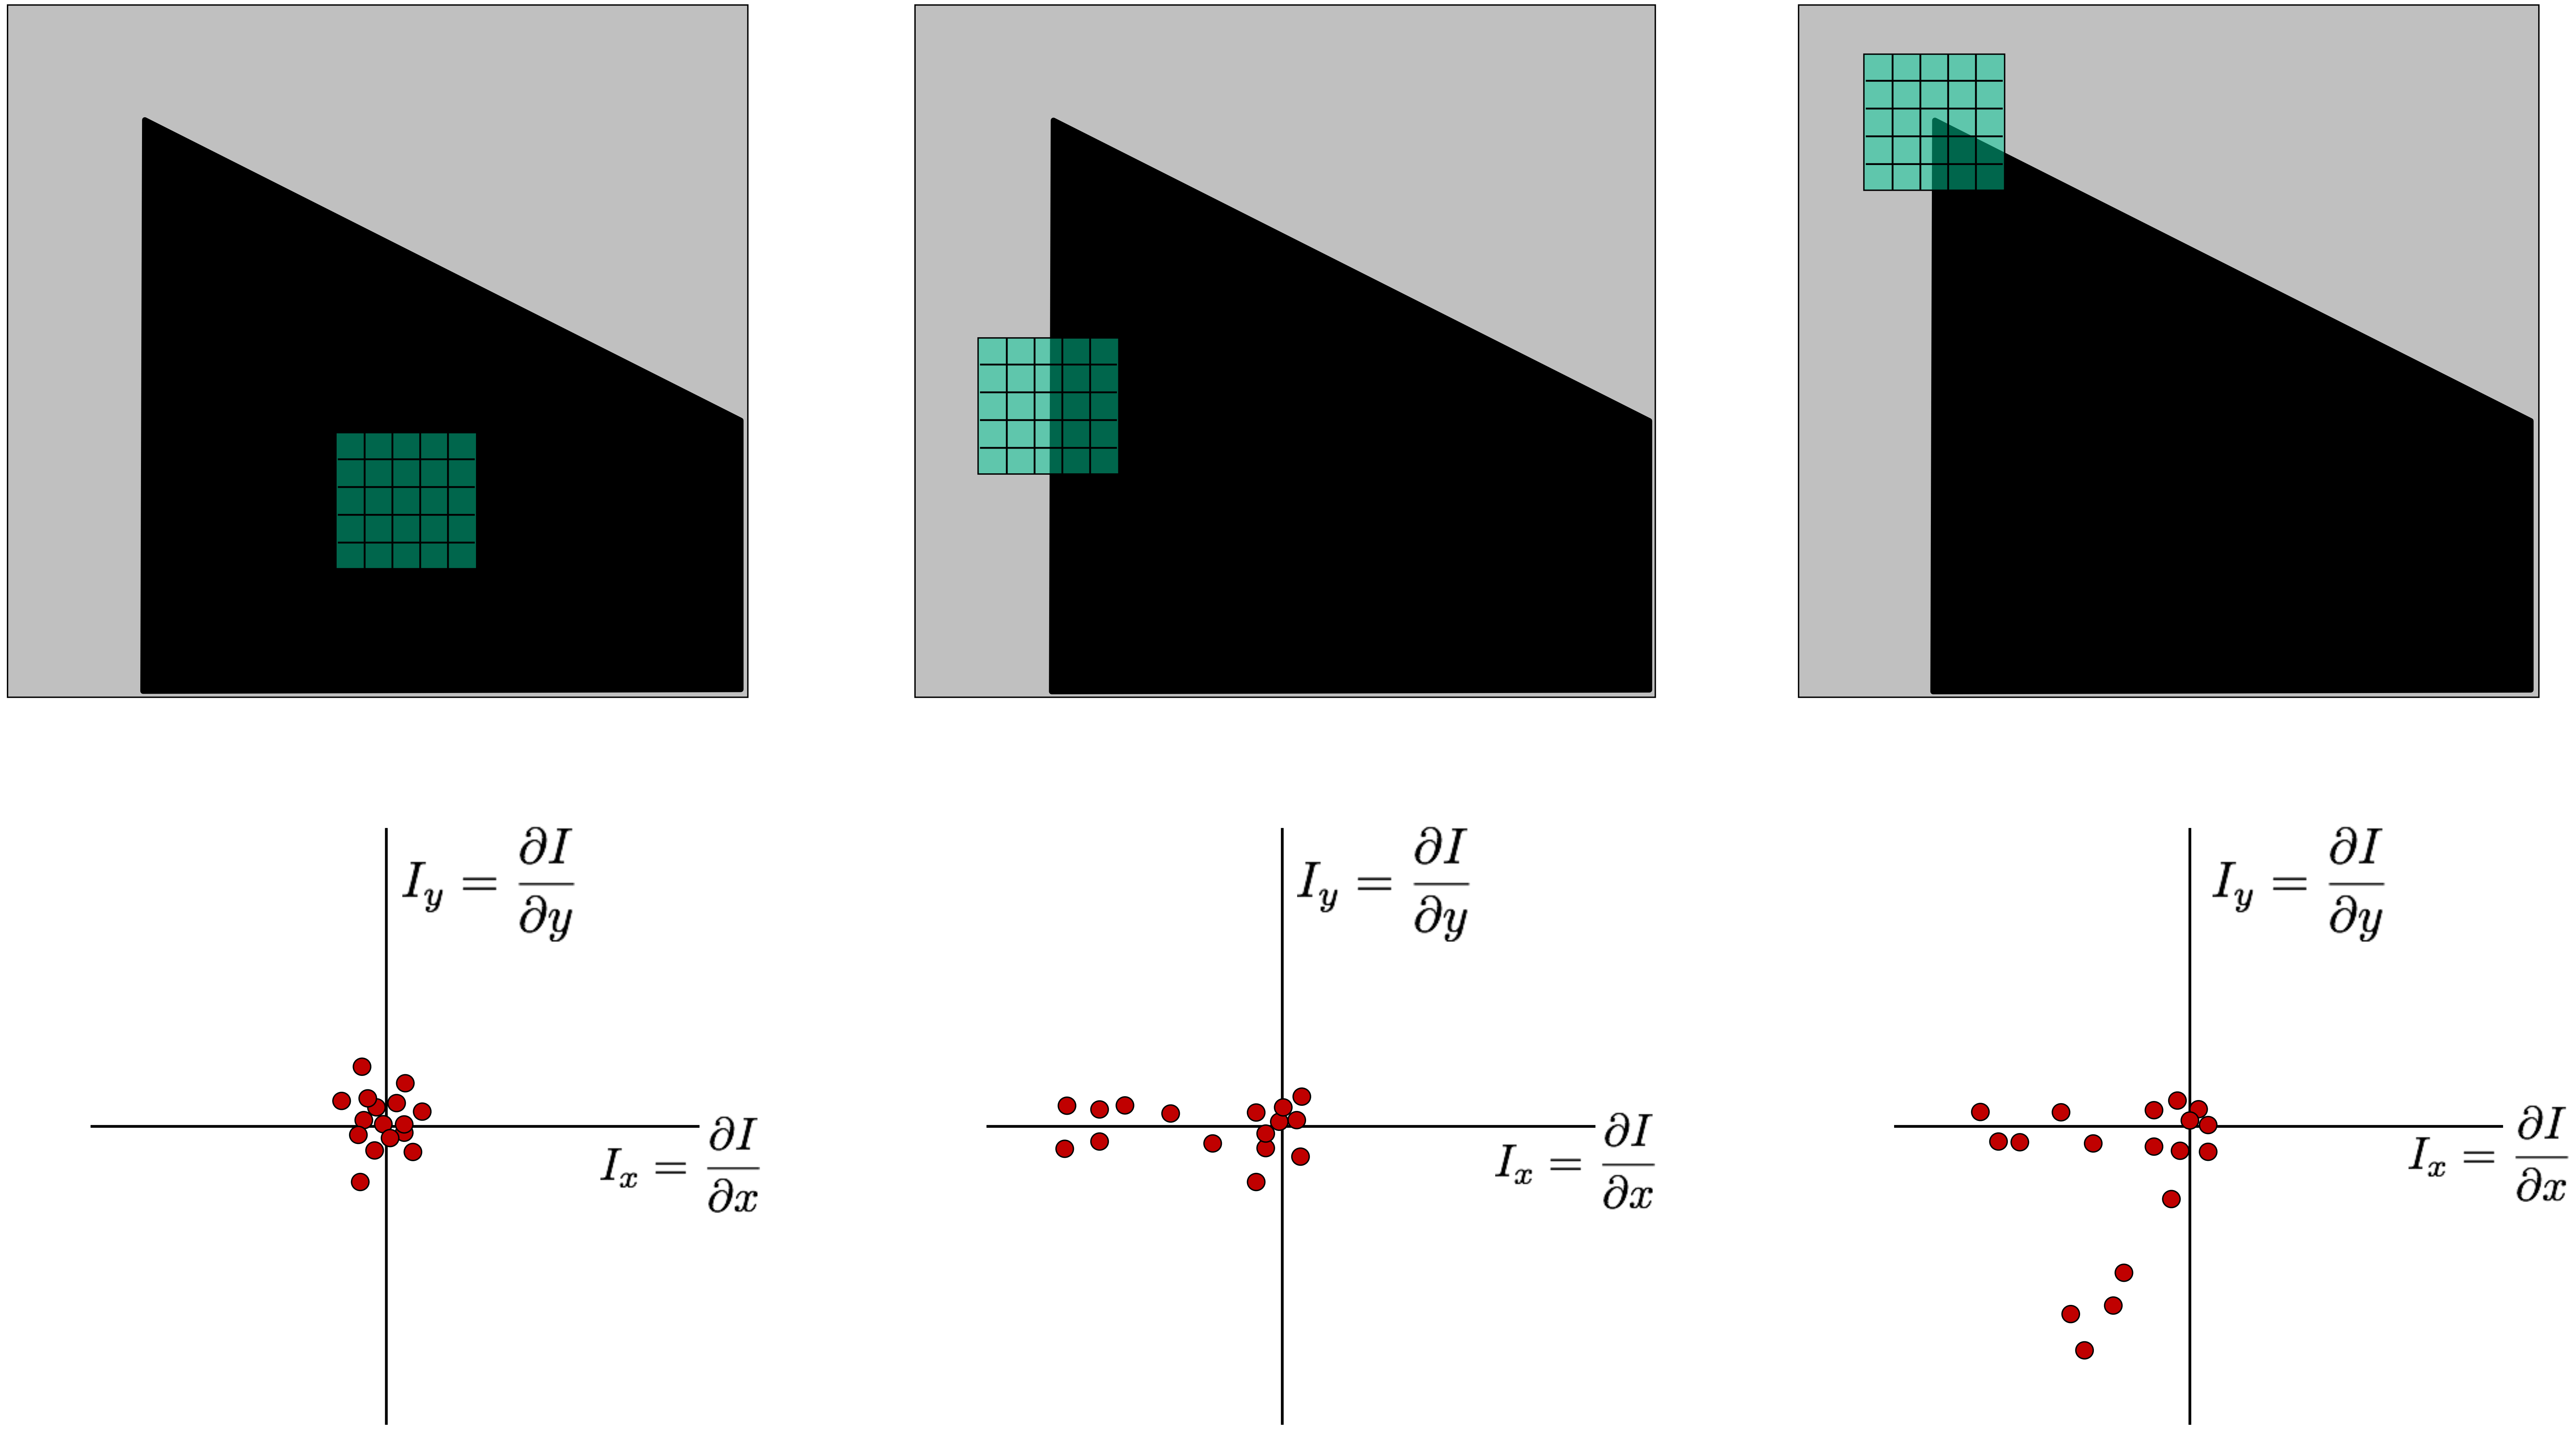

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

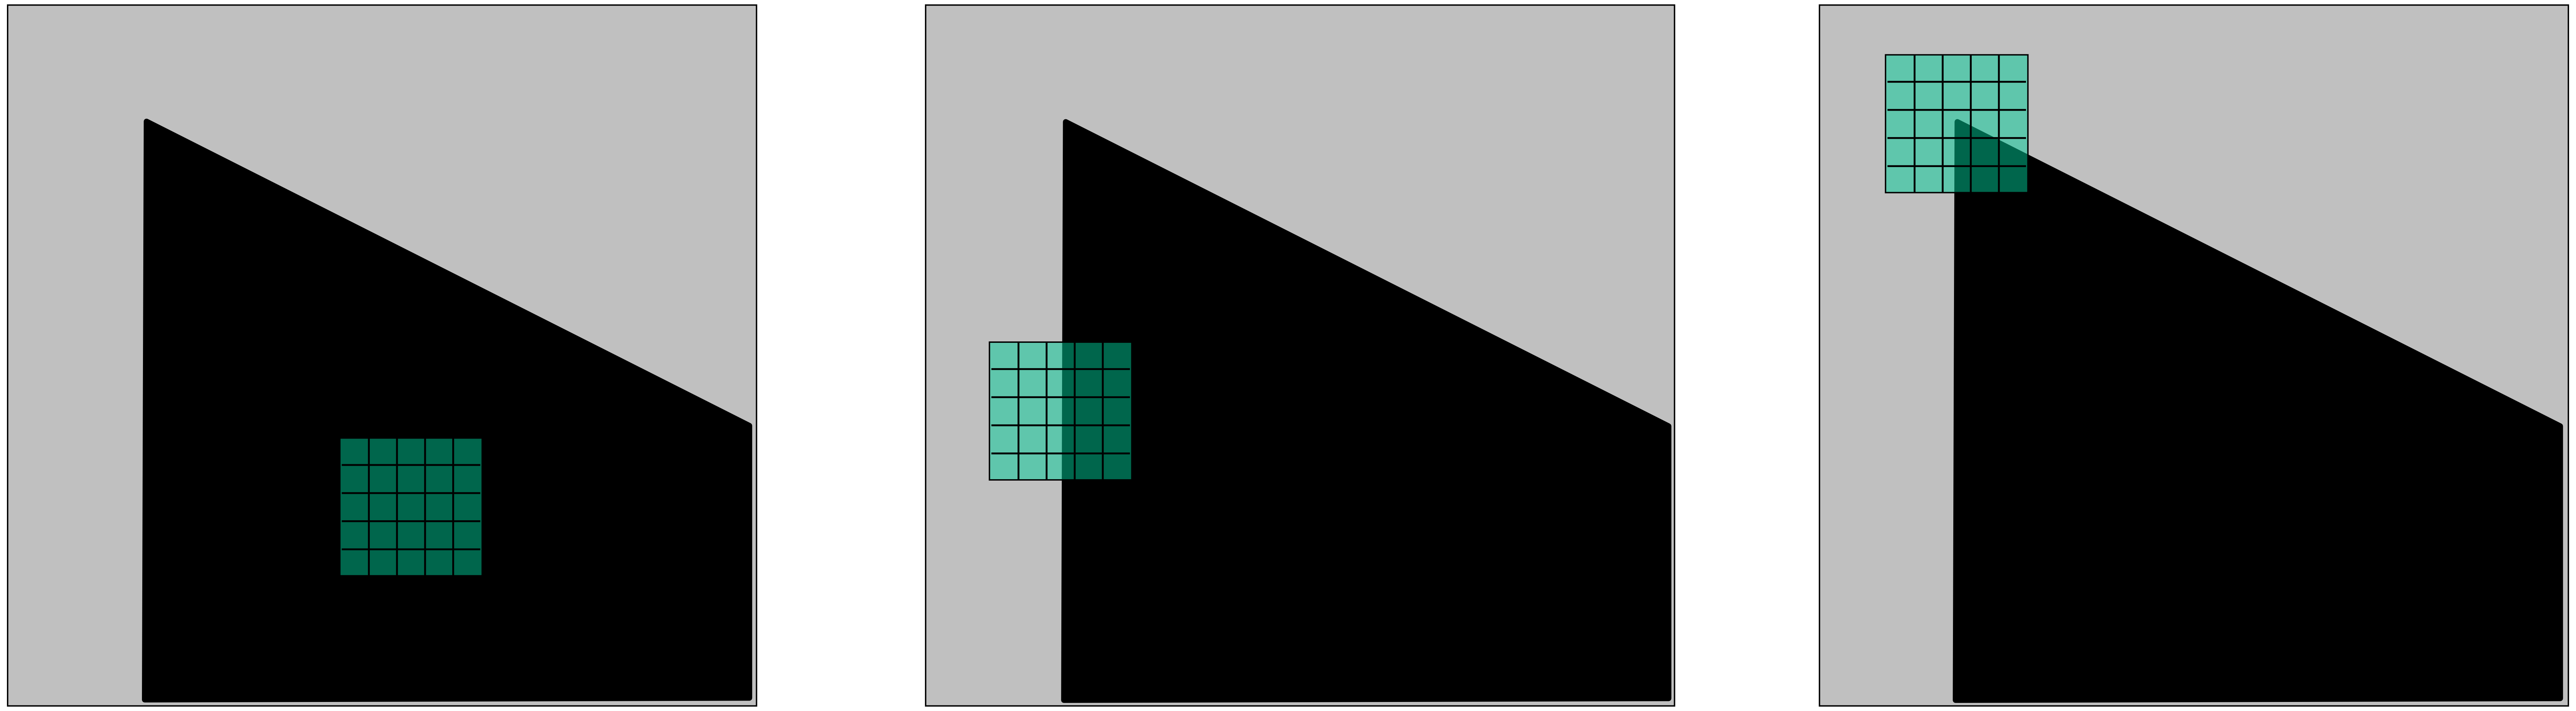

Distribution reveals the orientation and magnitude

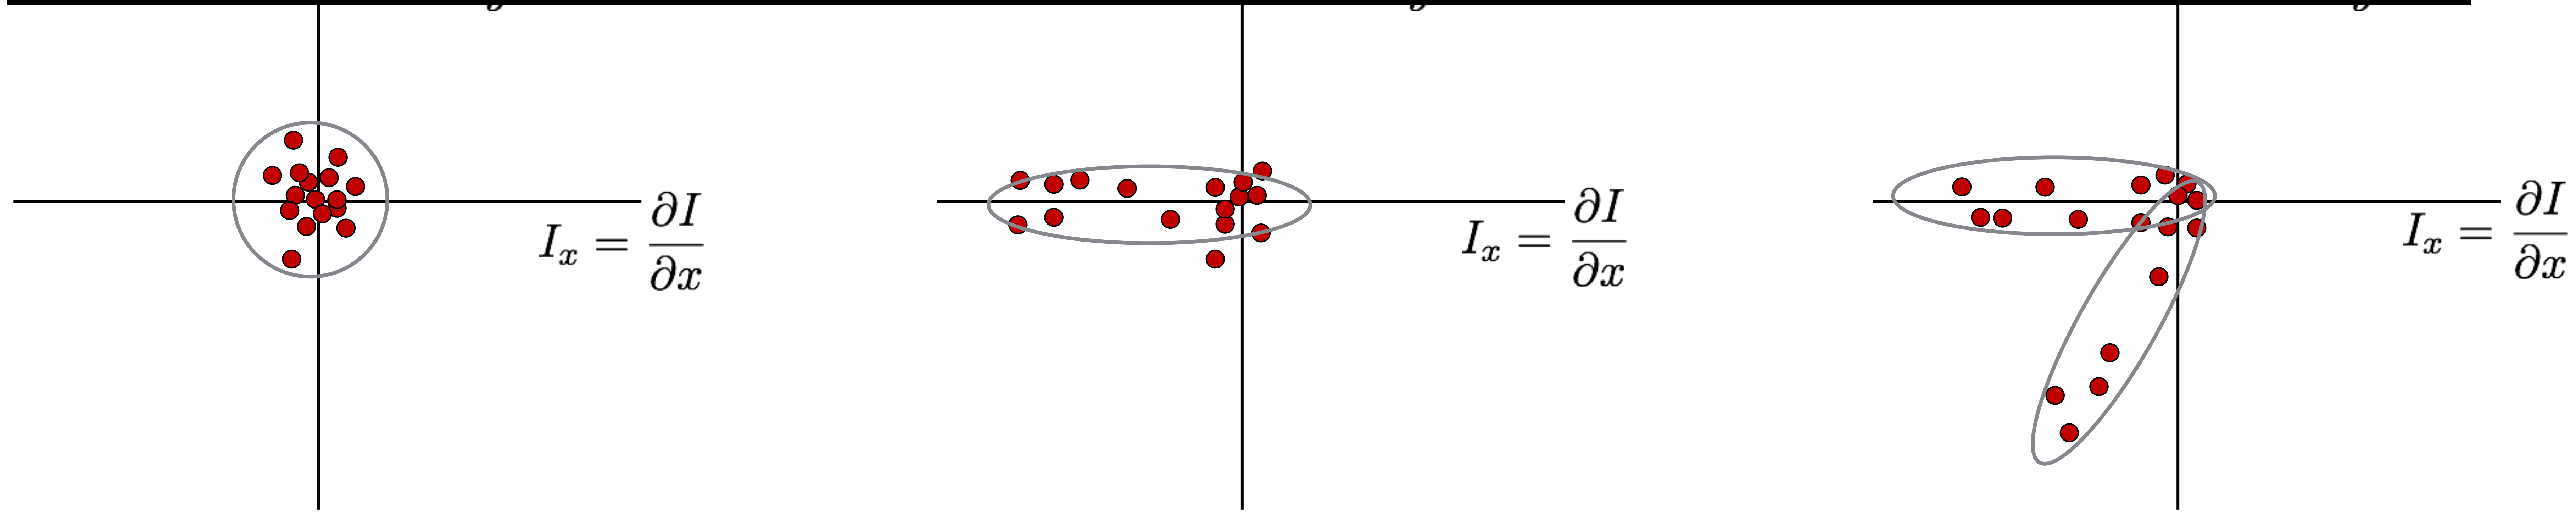

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

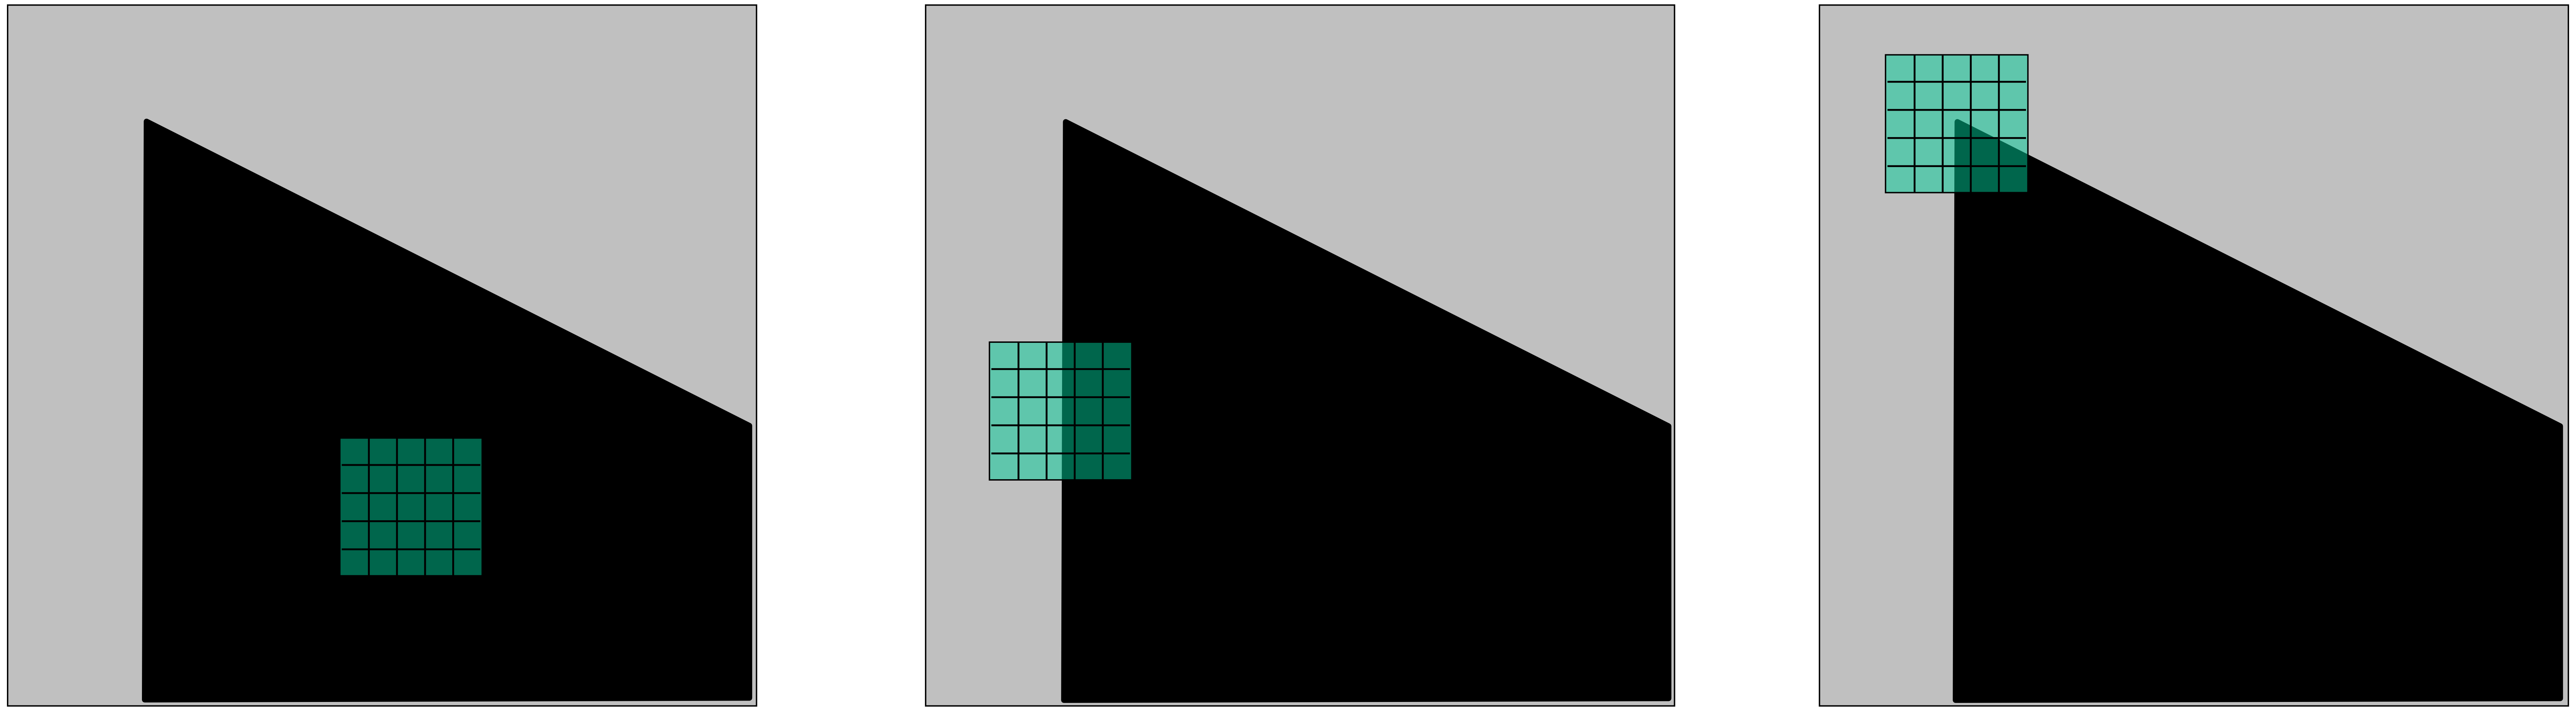

Distribution reveals the **orientation** and **magnitude**

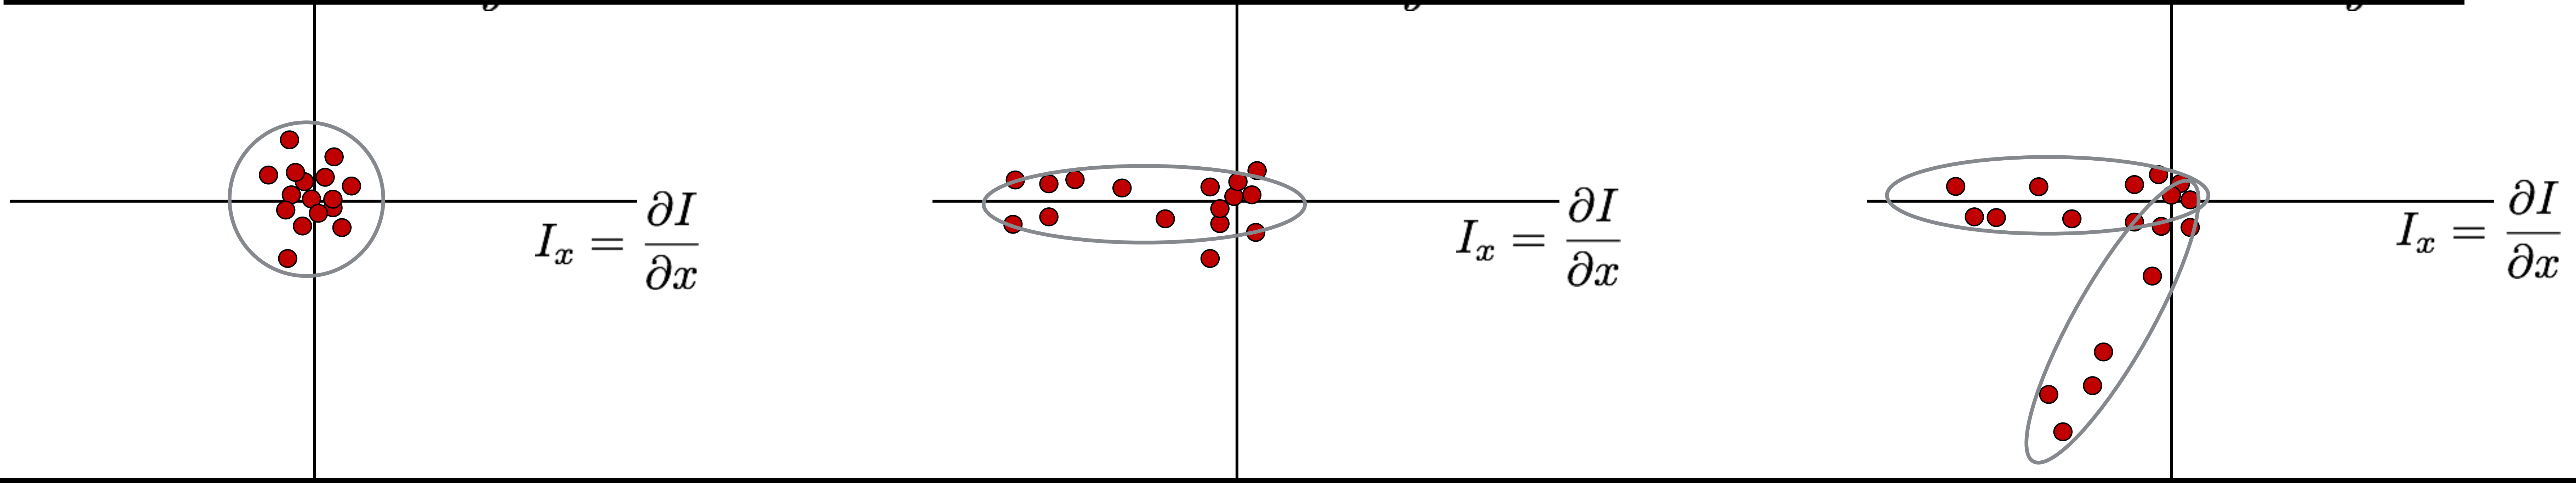

How do we quantify the **orientation** and **magnitude**?

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

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

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

Sum over small region around the corner

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

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

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

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

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

$$I_x = \frac{\partial I}{\partial x} \qquad I_y = \frac{\partial I}{\partial y}$$

$\sum\limits_{p \in P} I_x I_y$ =sum(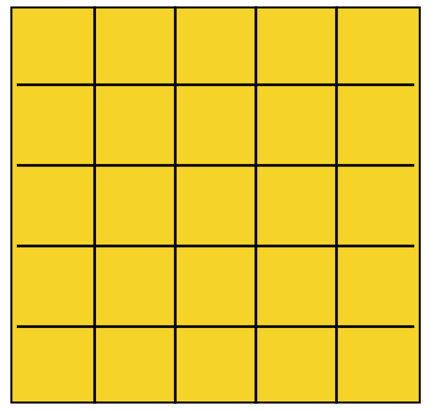 .* 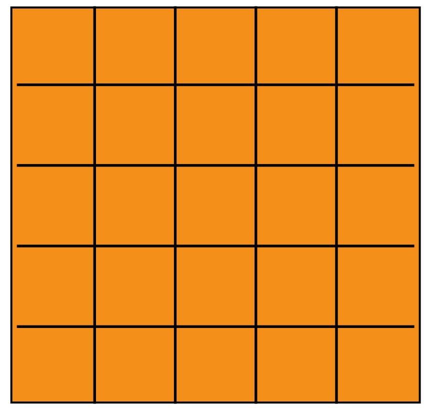)

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

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

Matrix is **symmetric**

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

By computing the **gradient covariance matrix** …

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

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

Simple Case

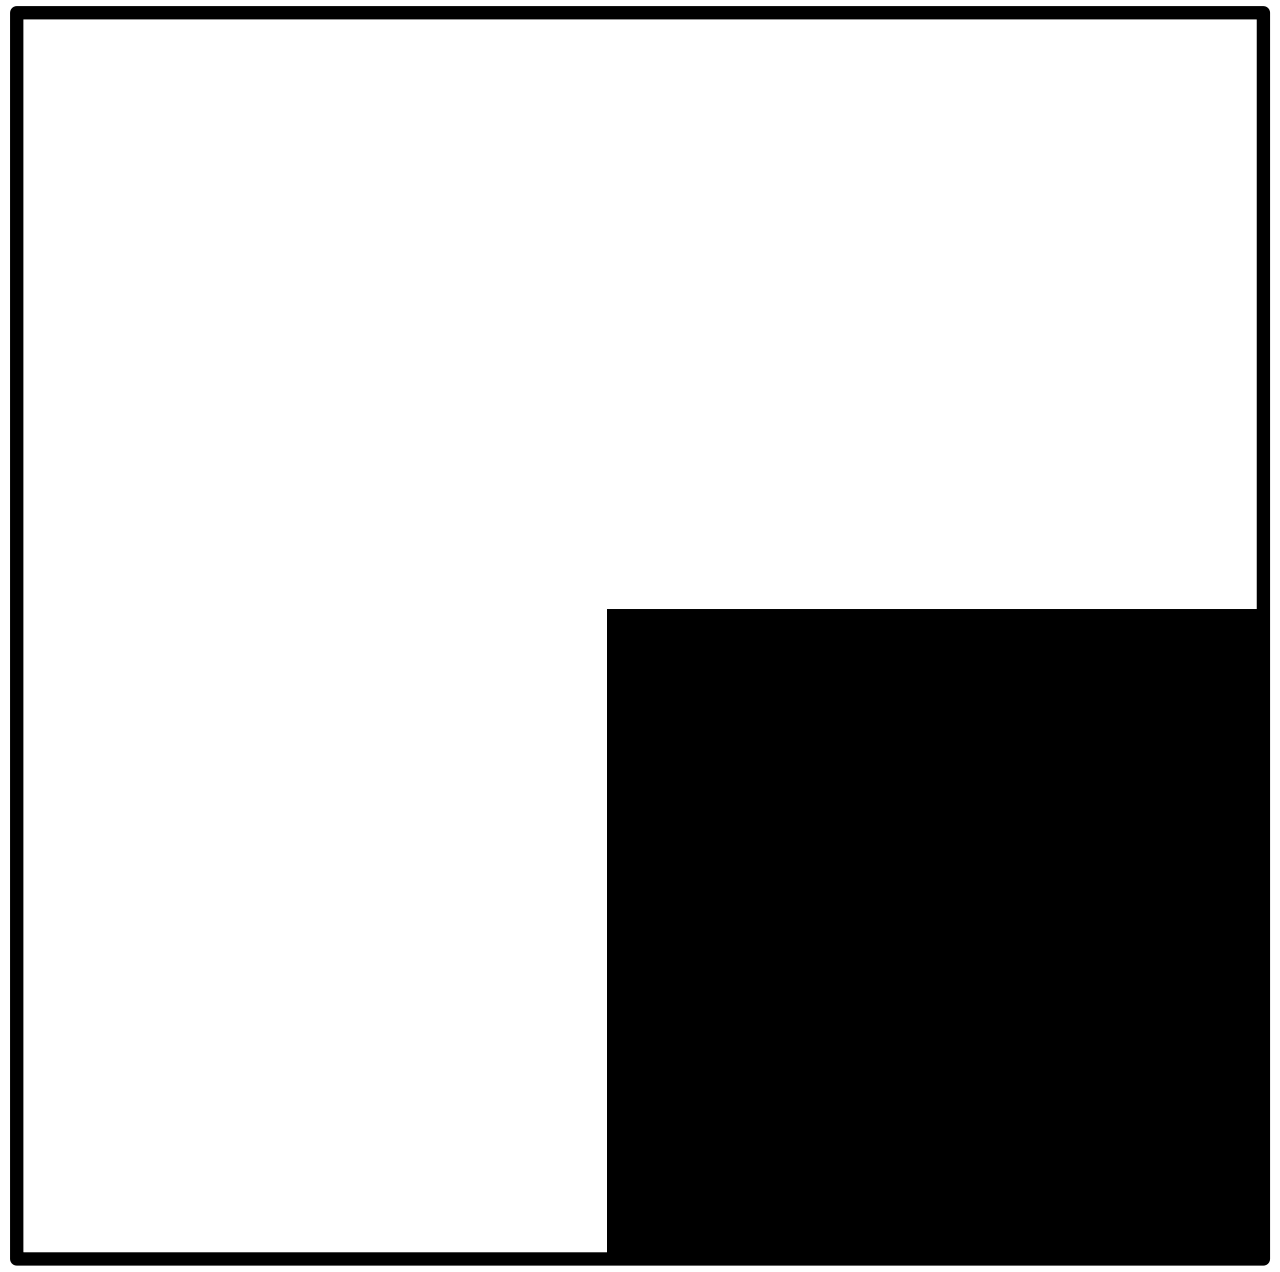

Local Image Patch

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

Simple Case

I_x I_y

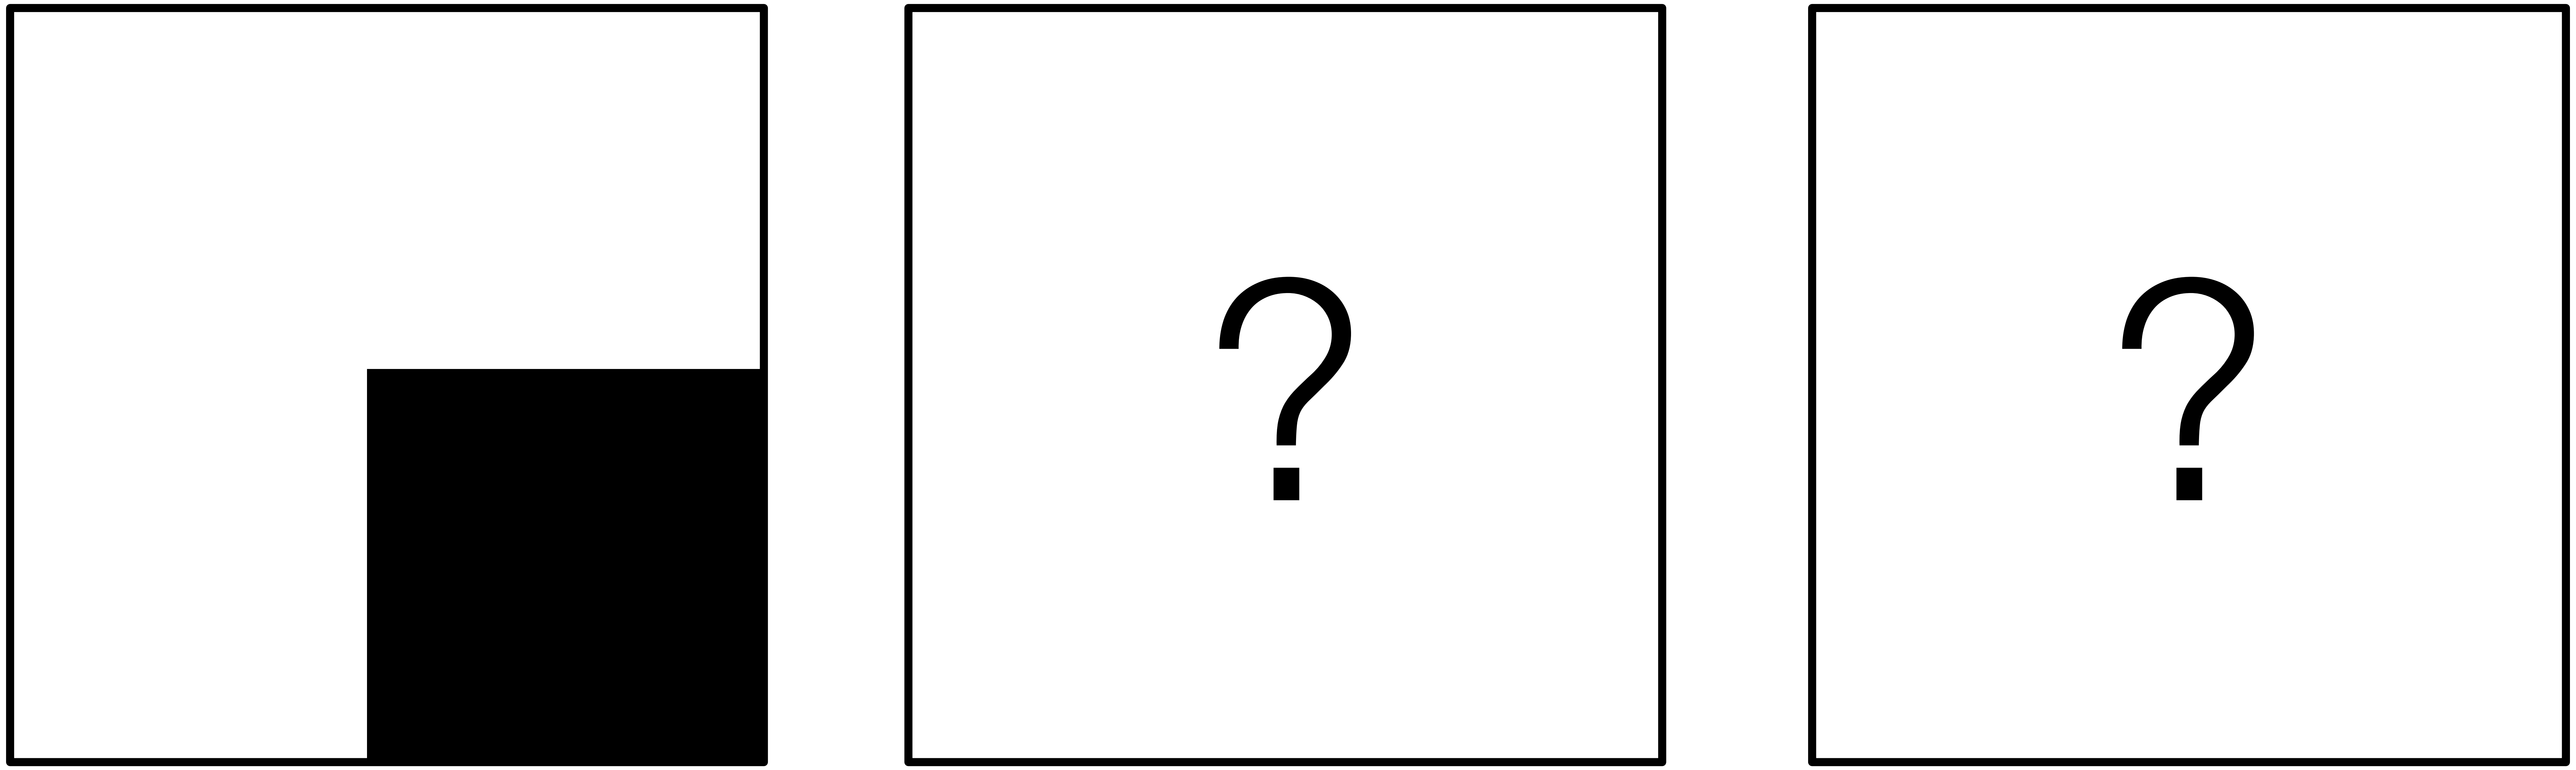

Local Image Patch

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

Simple Case

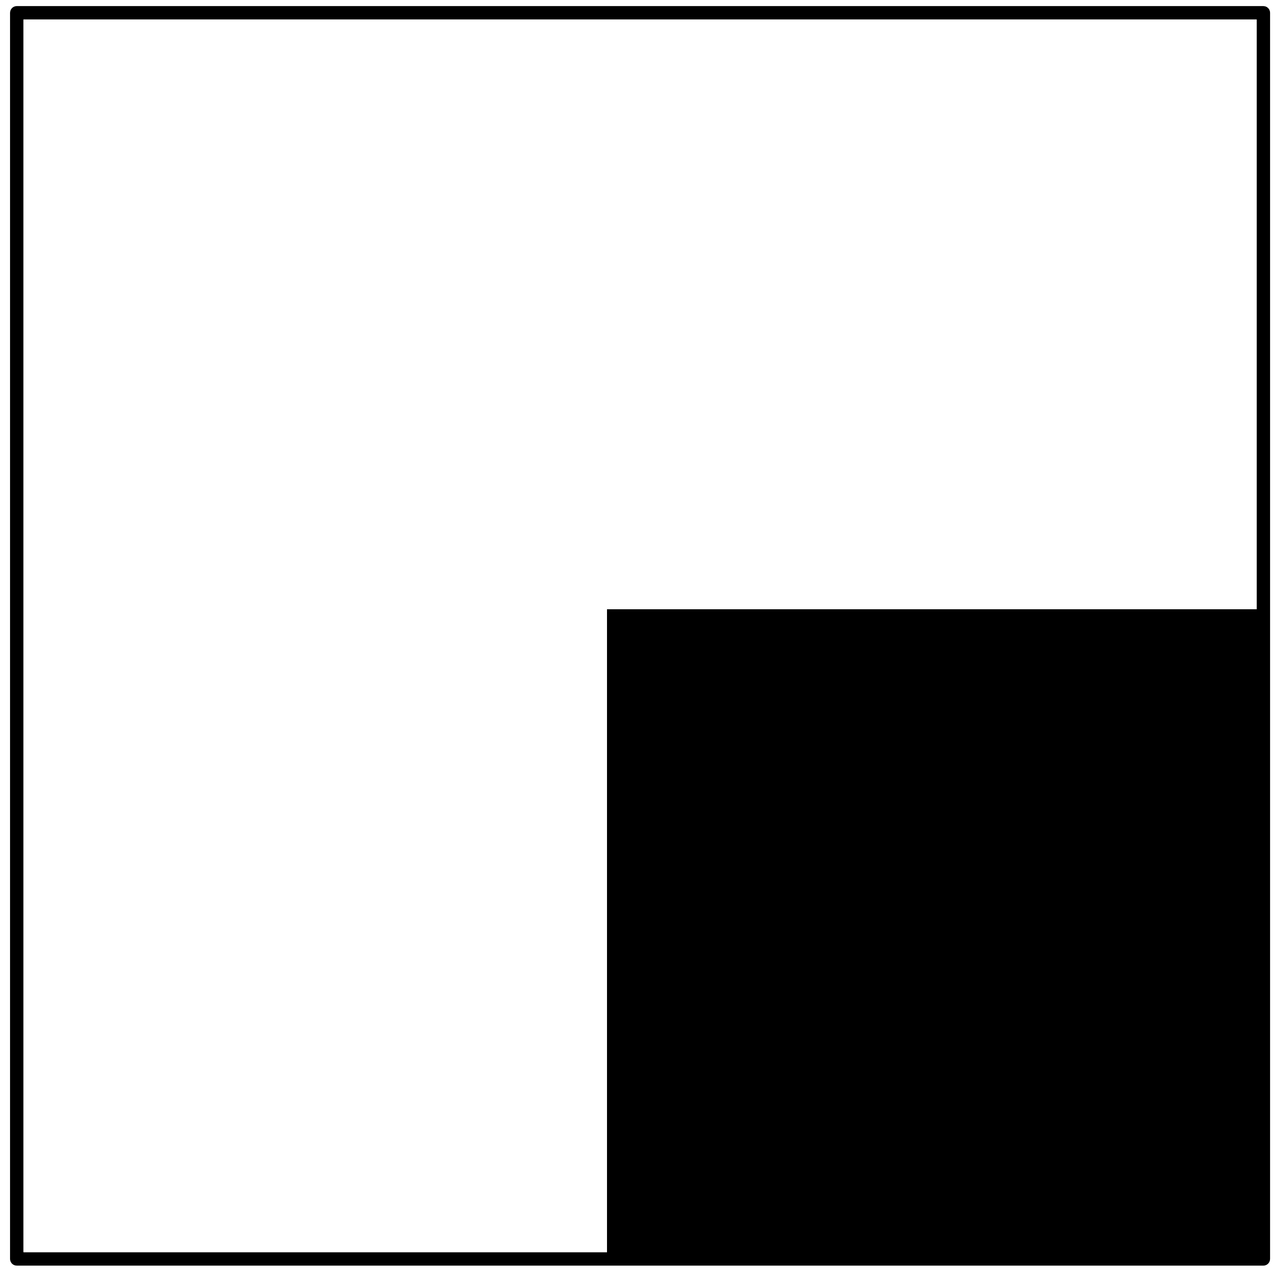

Local Image Patch

I_x

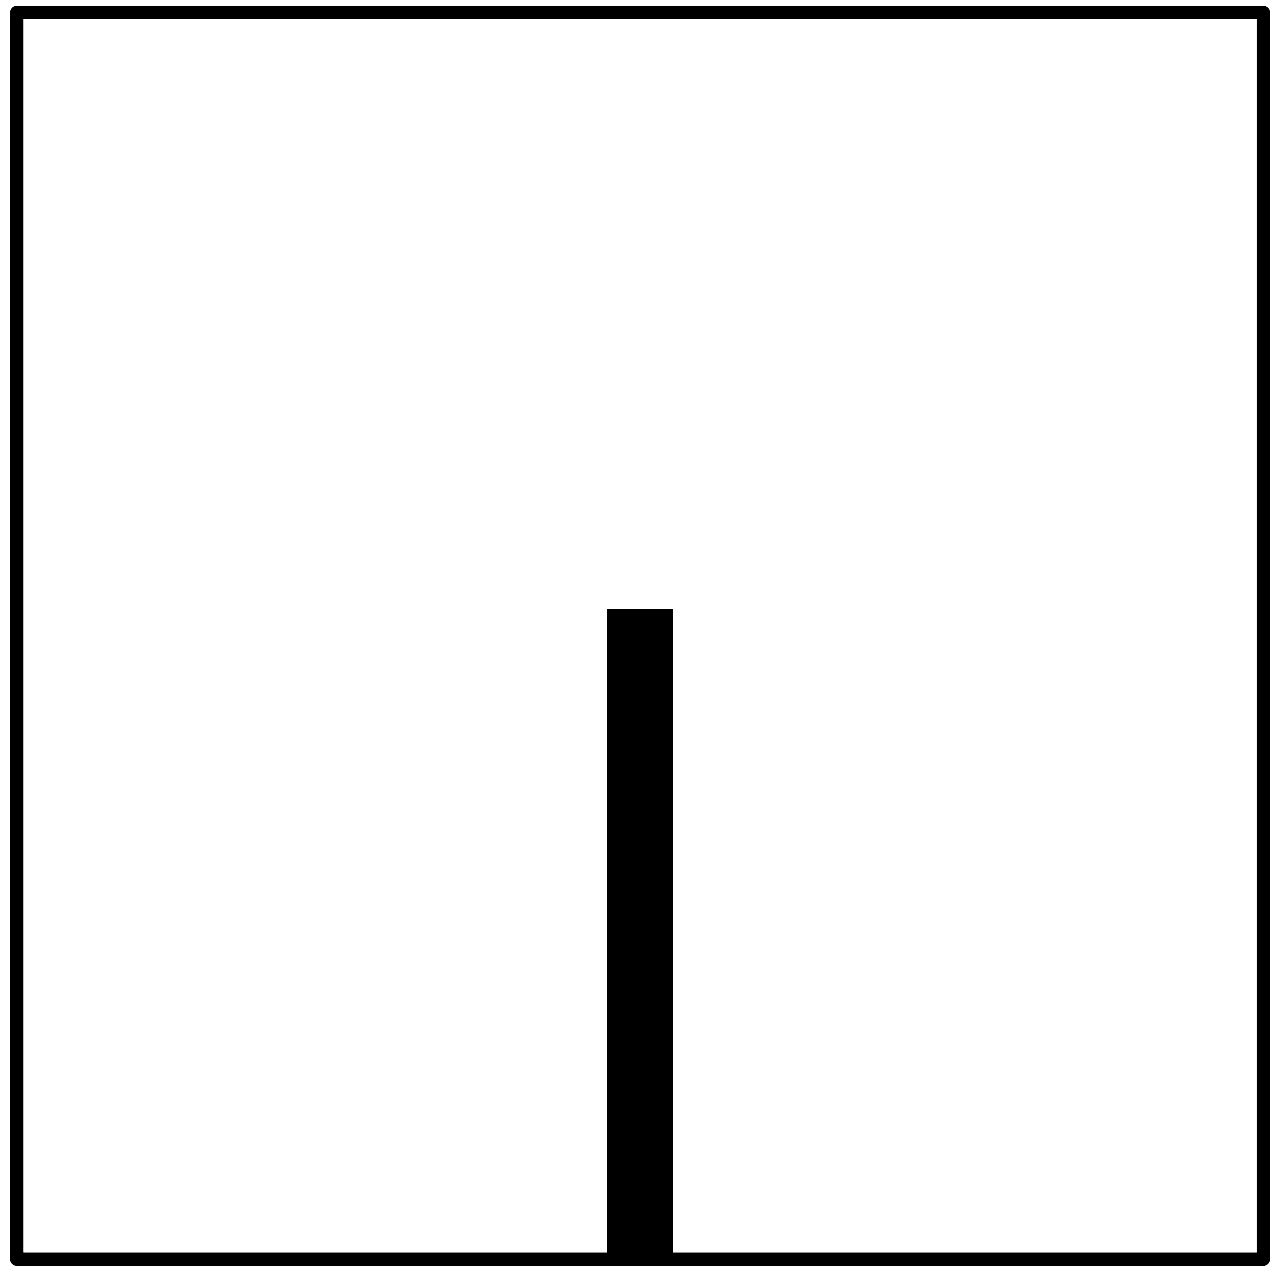

high value along vertical
strip of pixels and 0 elsewhere

I_y

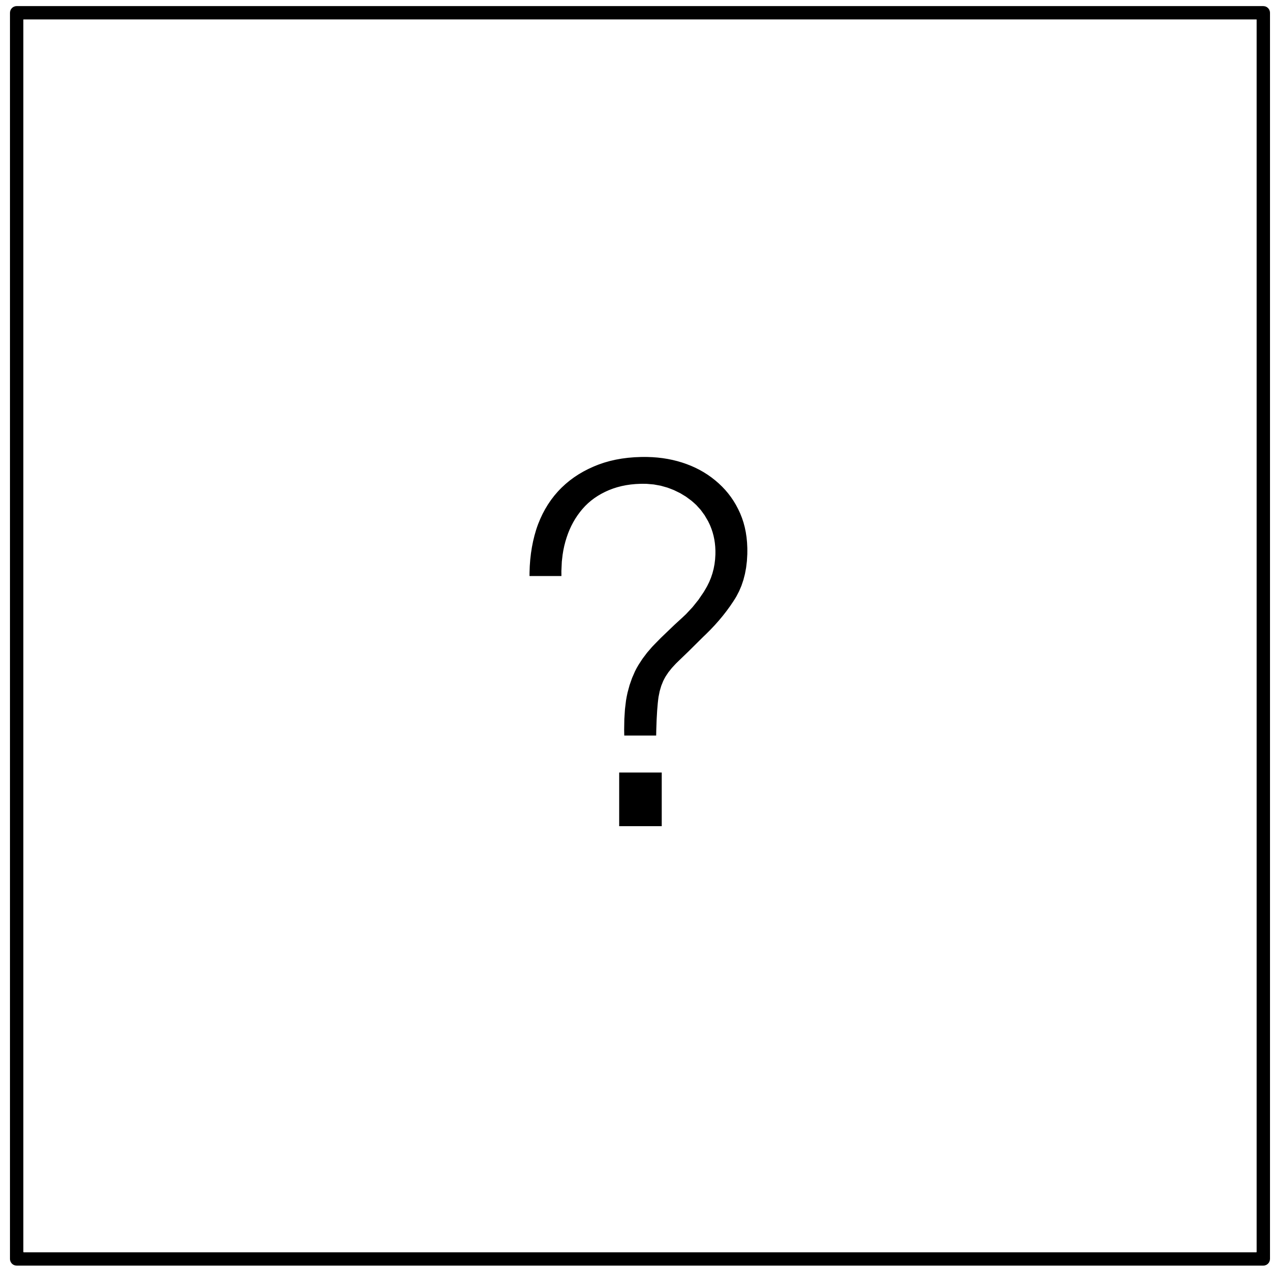

$$C = \begin{bmatrix} \sum\limits_{p \in P} I_x I_x & \sum\limits_{p \in P} I_x I_y \\ \sum\limits_{p \in P} I_y I_x & \sum\limits_{p \in P} I_y I_y \end{bmatrix} = ?$$

Simple Case

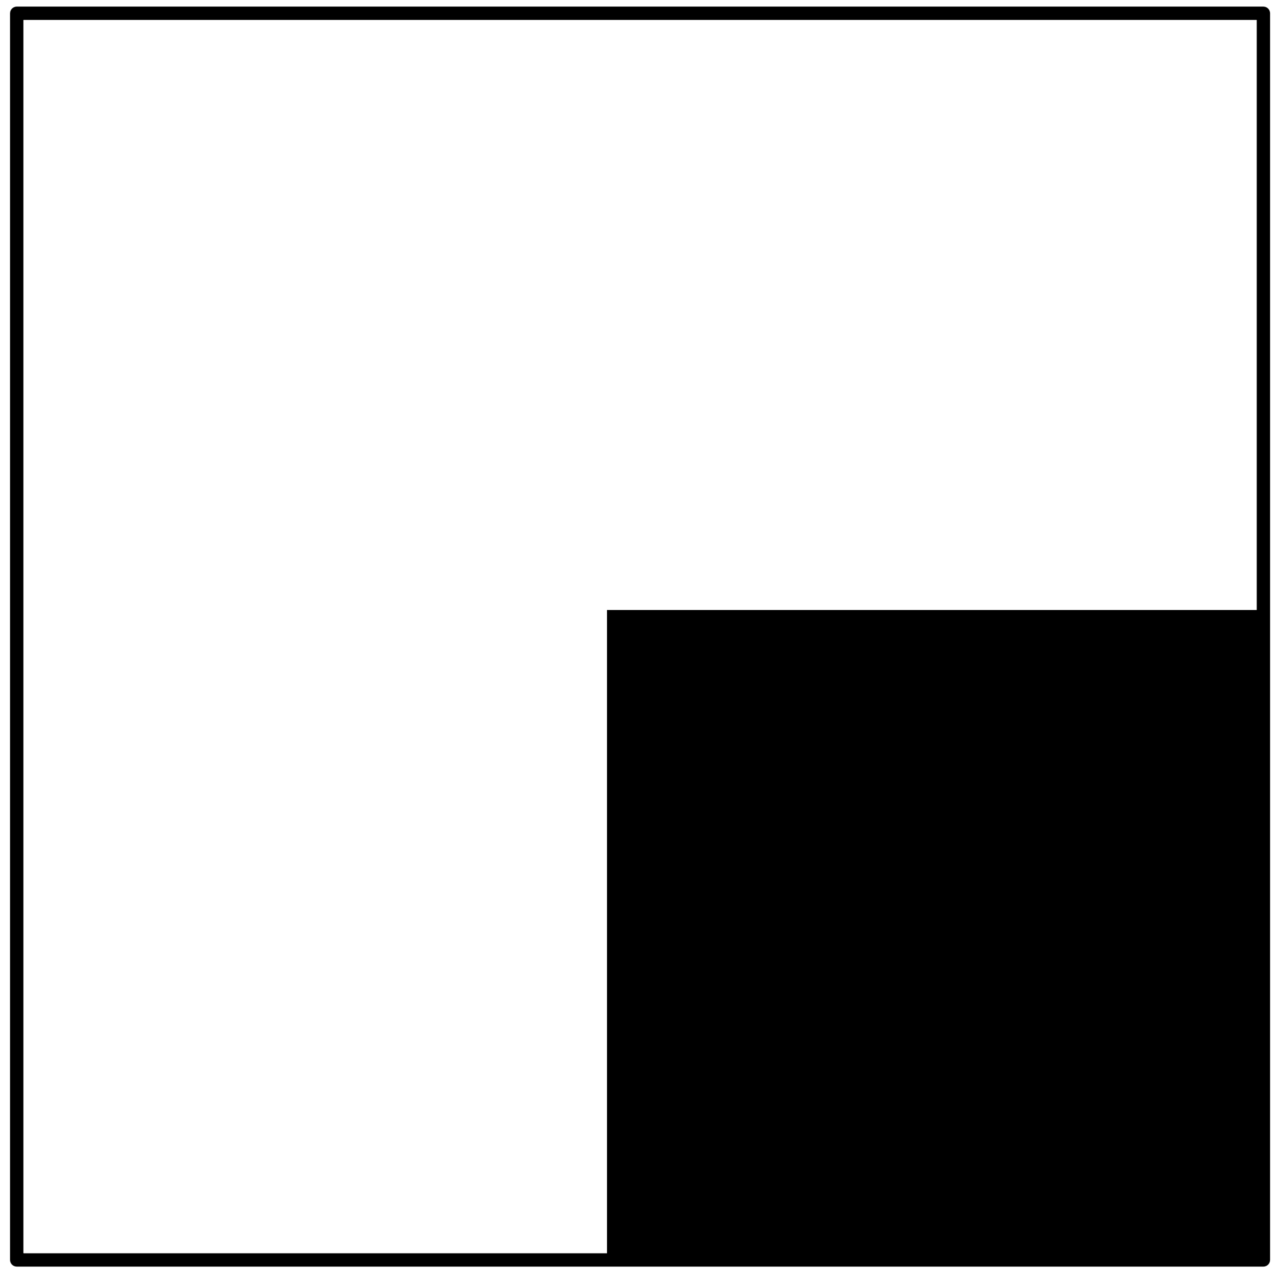

Local Image Patch

I_x

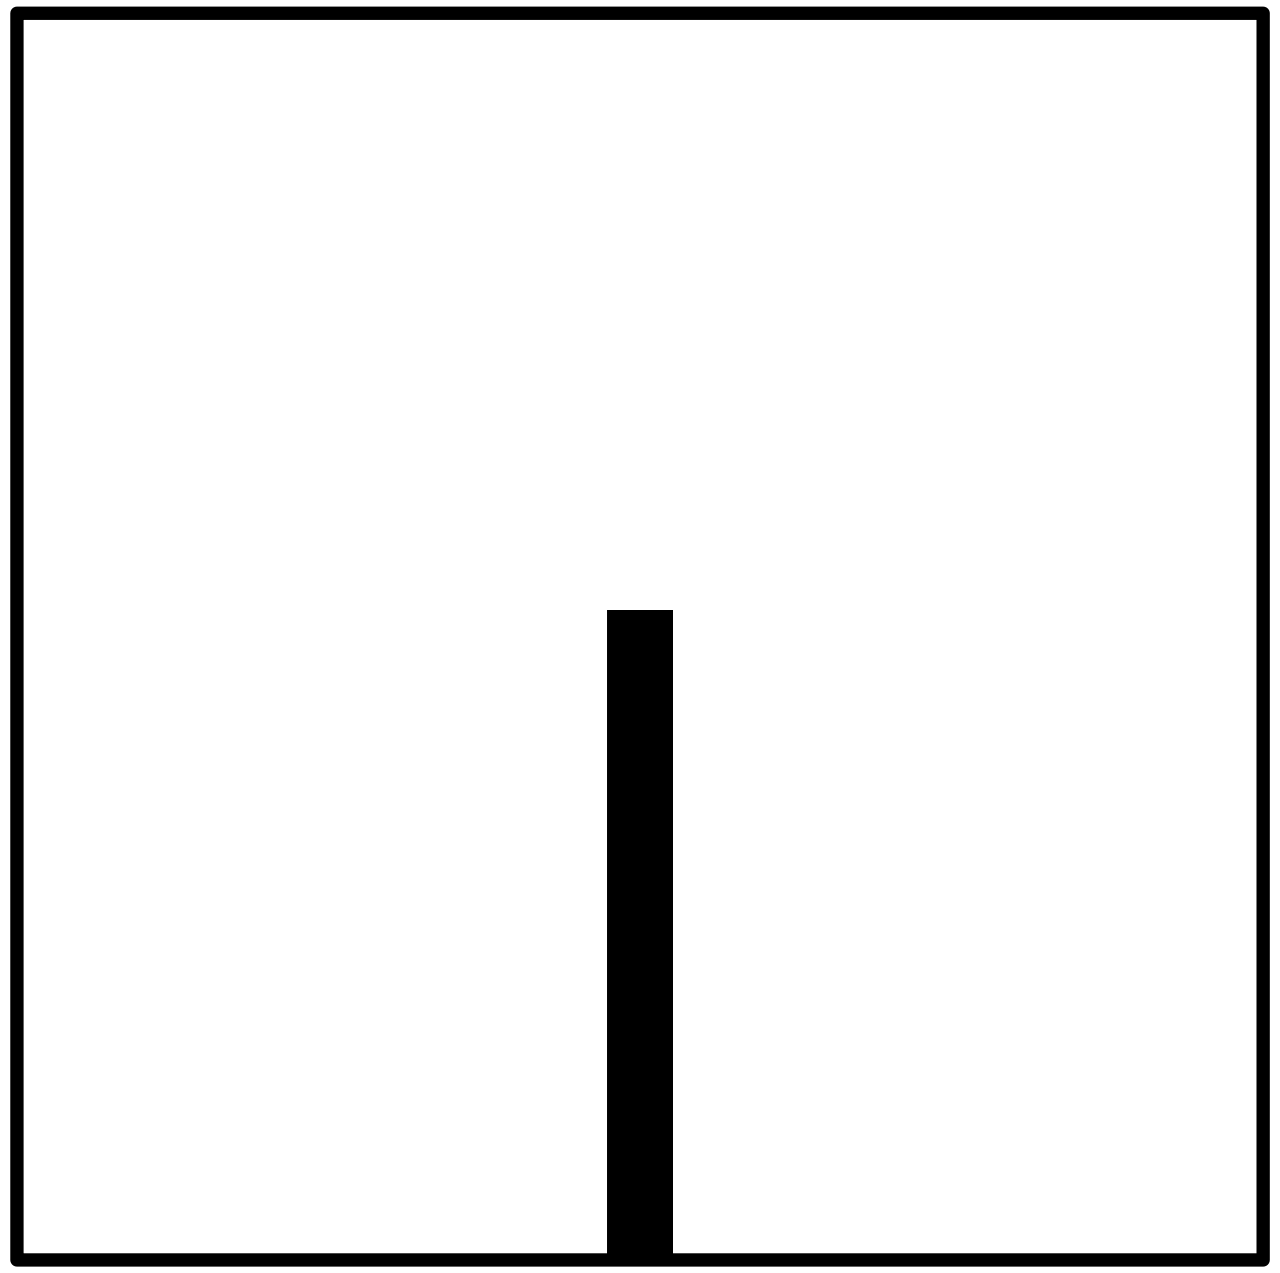

high value along vertical strip of pixels and 0 elsewhere

I_y

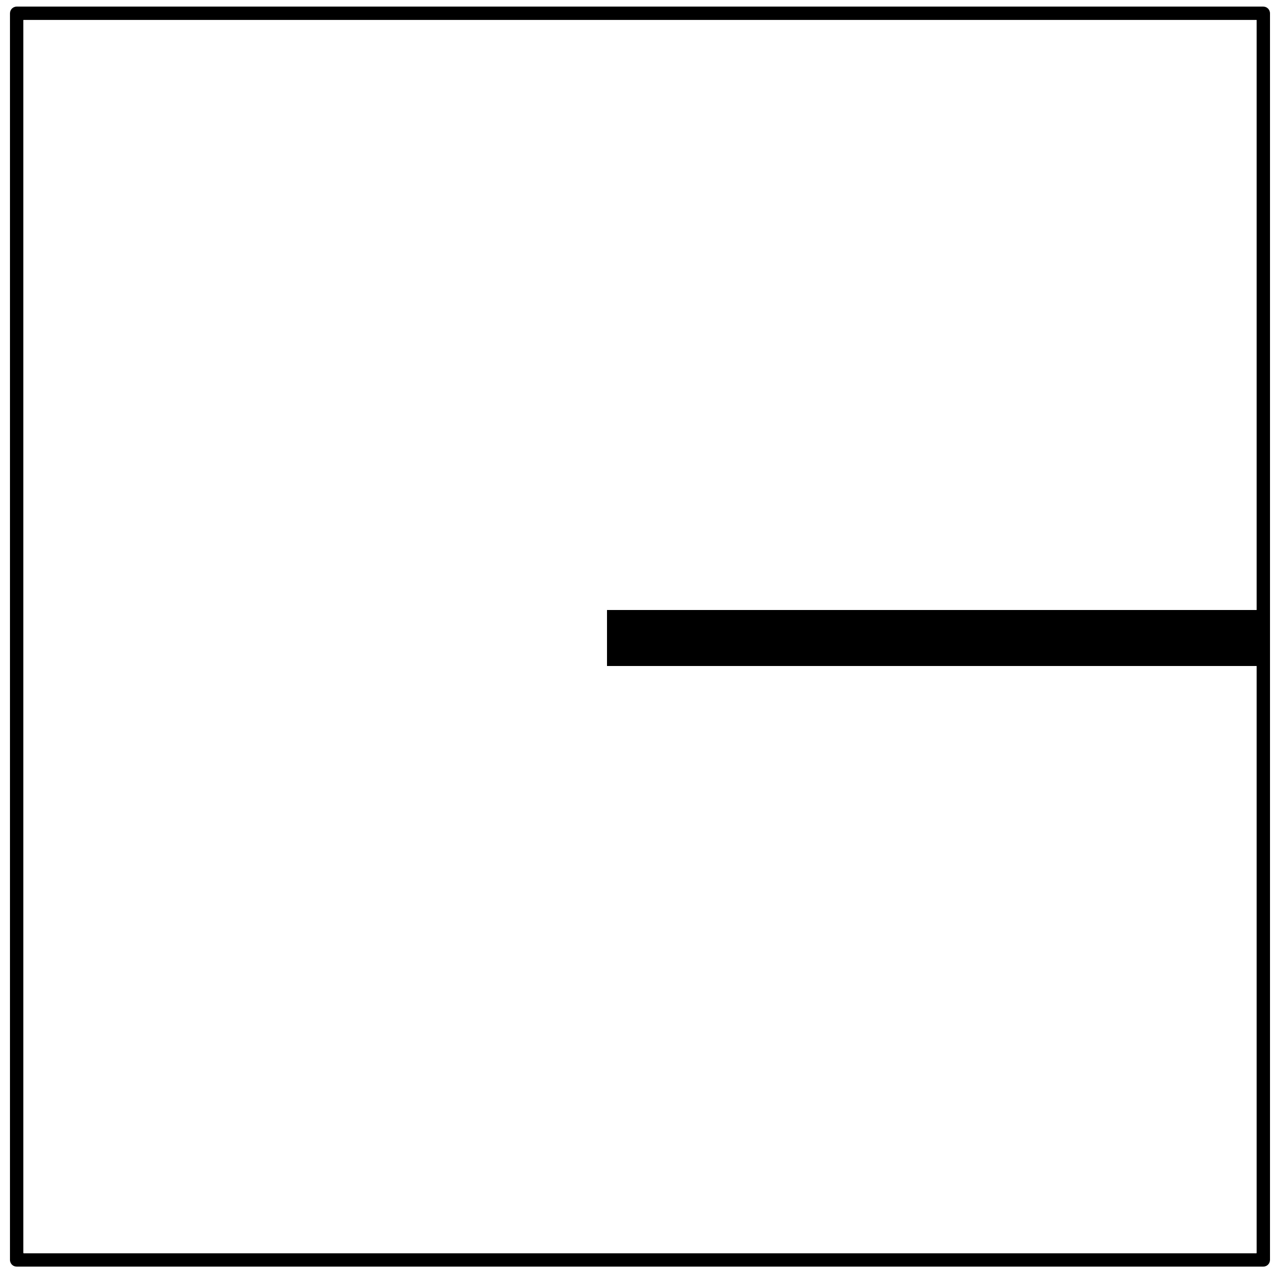

high value along horizontal strip of pixels and 0 elsewhere

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

Simple Case

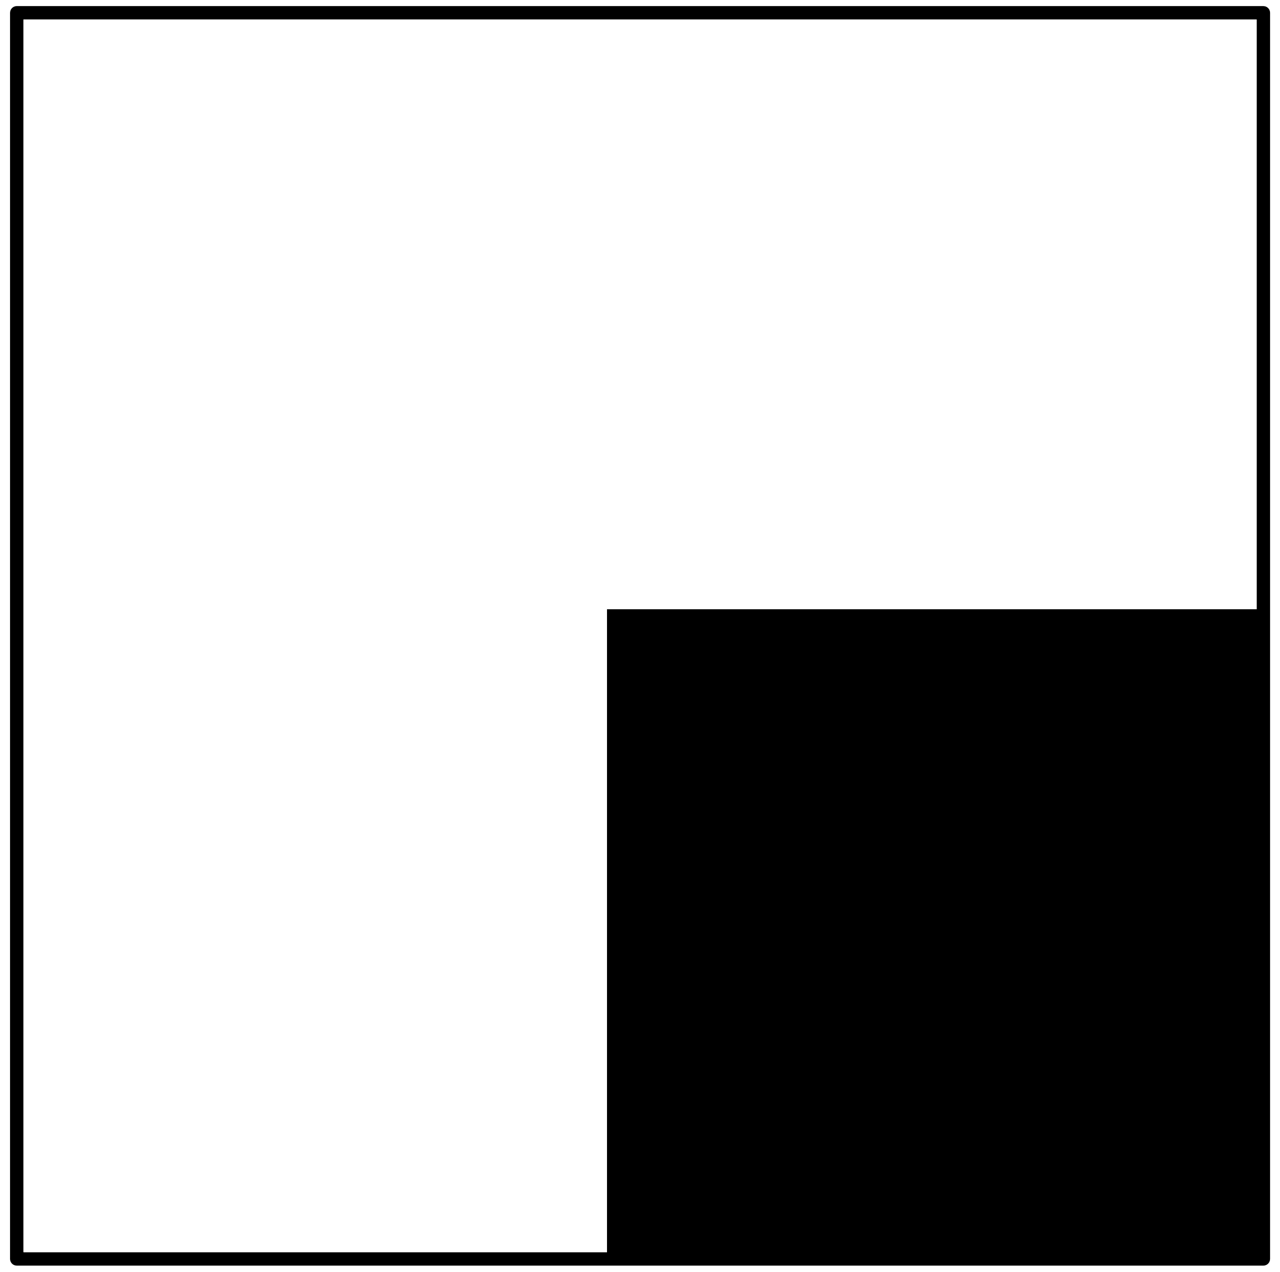

Local Image Patch

I_x

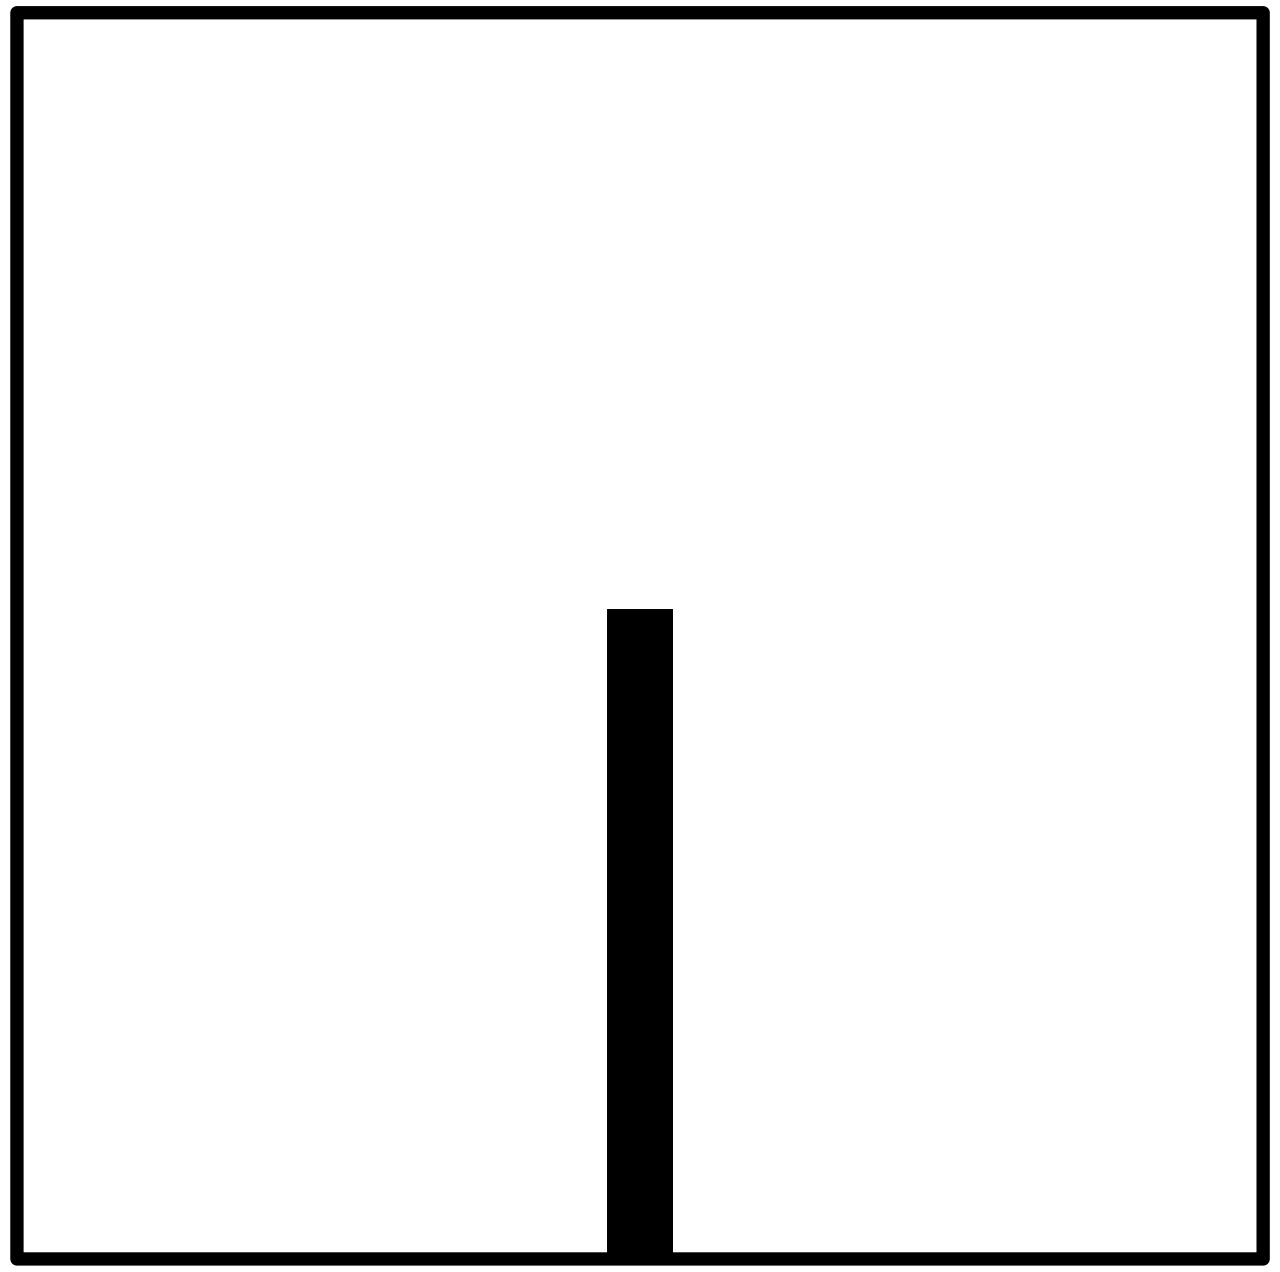

high value along vertical
strip of pixels and 0 elsewhere

I_y

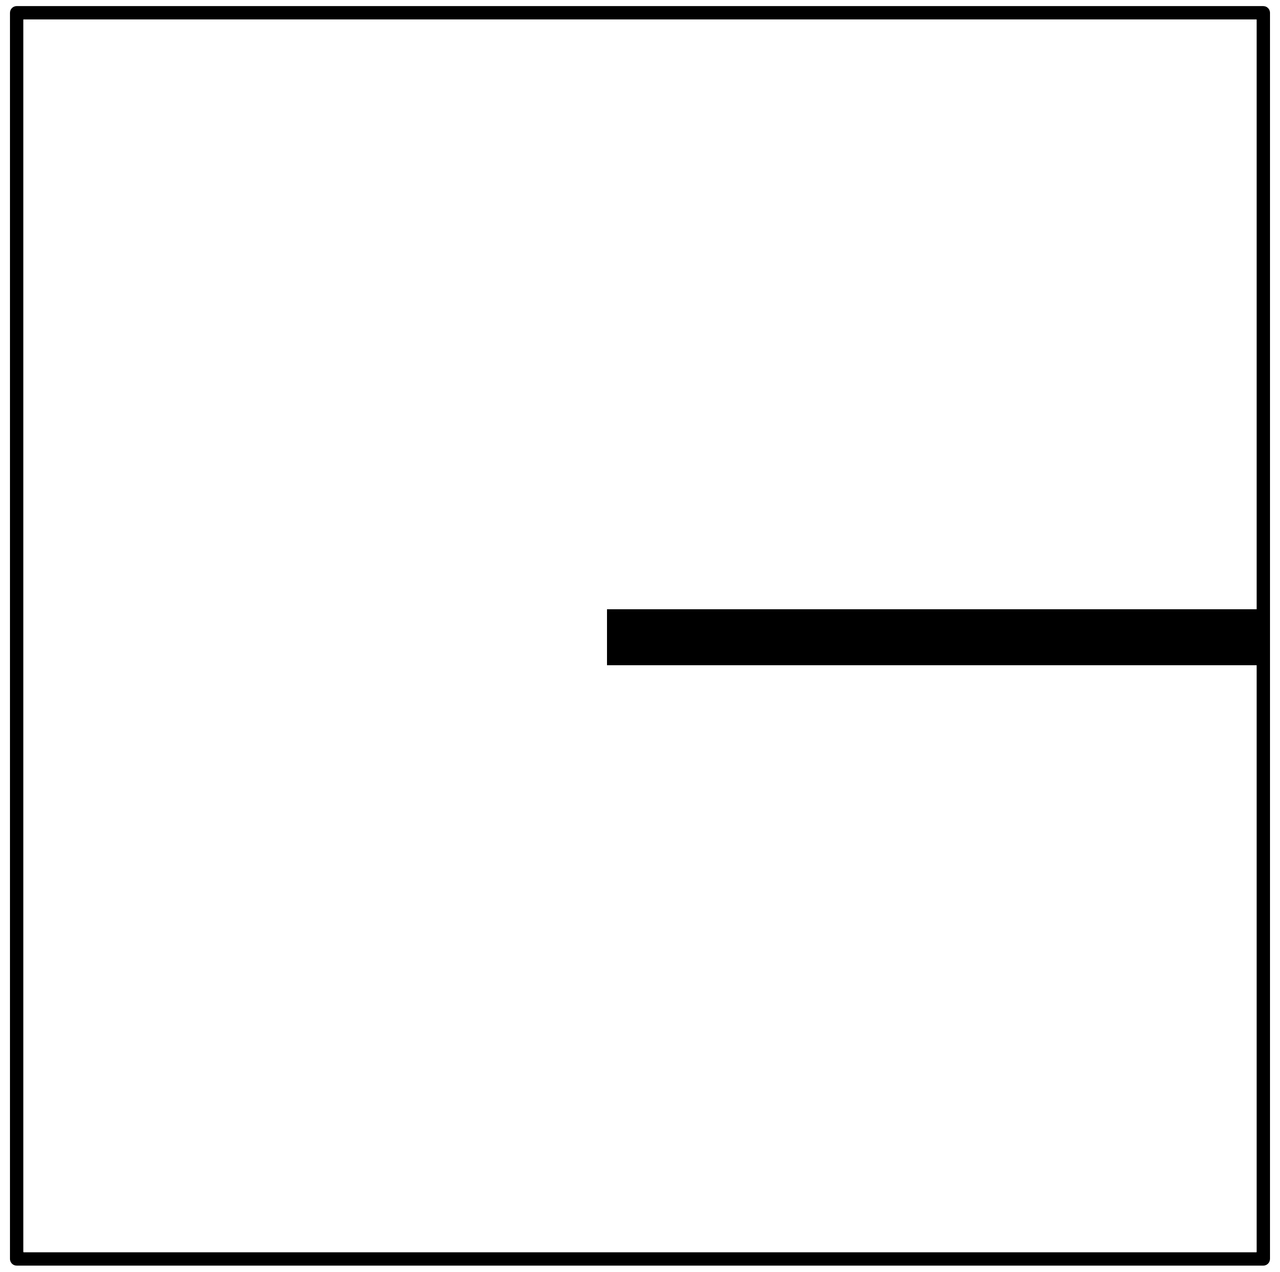

high value along horizontal
strip of pixels and 0 elsewhere

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

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**

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

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

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

eigenvalue

$$Ce = \lambda e$$

eigenvector

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

eigenvalue

$$Ce = \lambda e$$

eigenvector

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

1. Compute the determinant of (returns a polynomial)

$$C - \lambda I$$

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Visualization as **Quadratic**

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

can be written in matrix form like this…

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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)

eigenvectors

eigenvalues along diagonal

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

axis of the 'ellipse slice'

scaling of the quadratic along the axis

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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:

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

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

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:

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

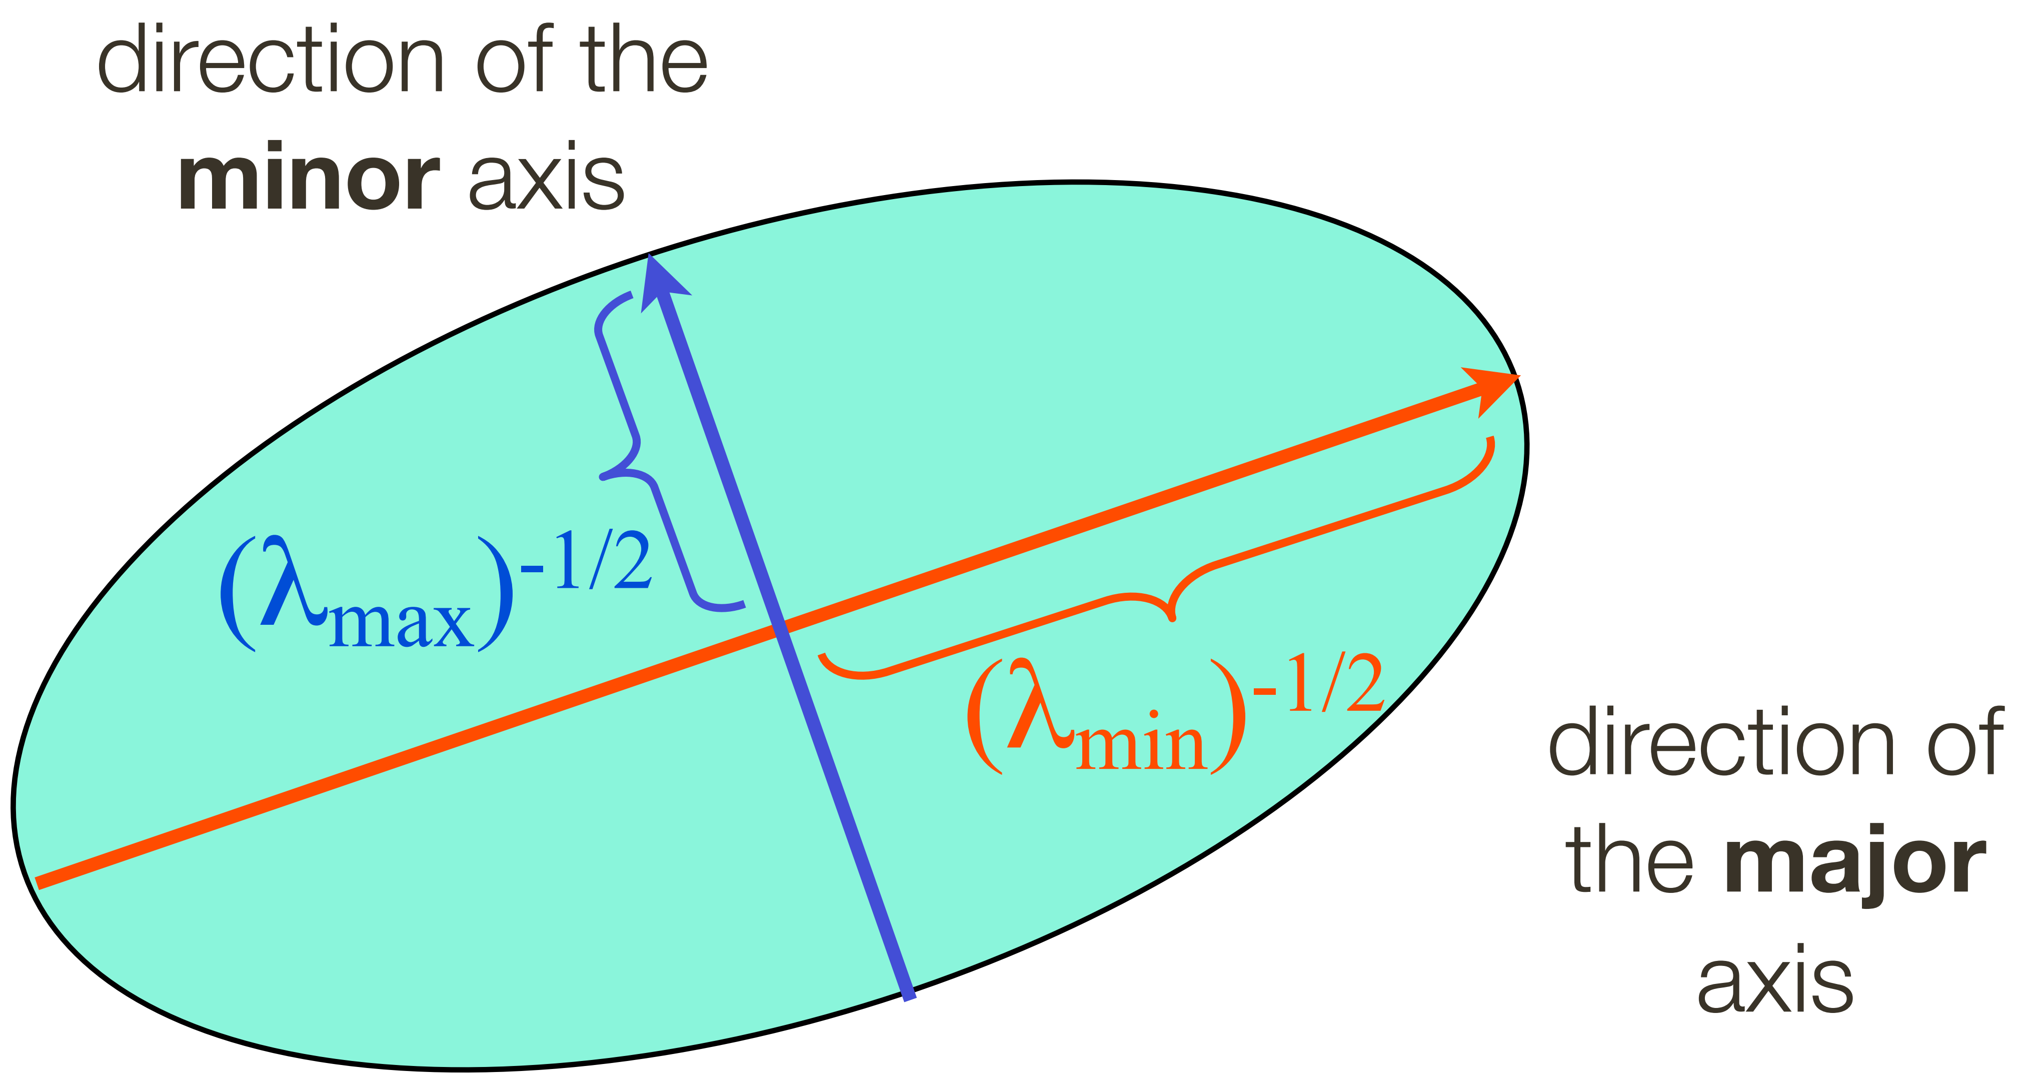

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Interpreting **Eigenvalues**

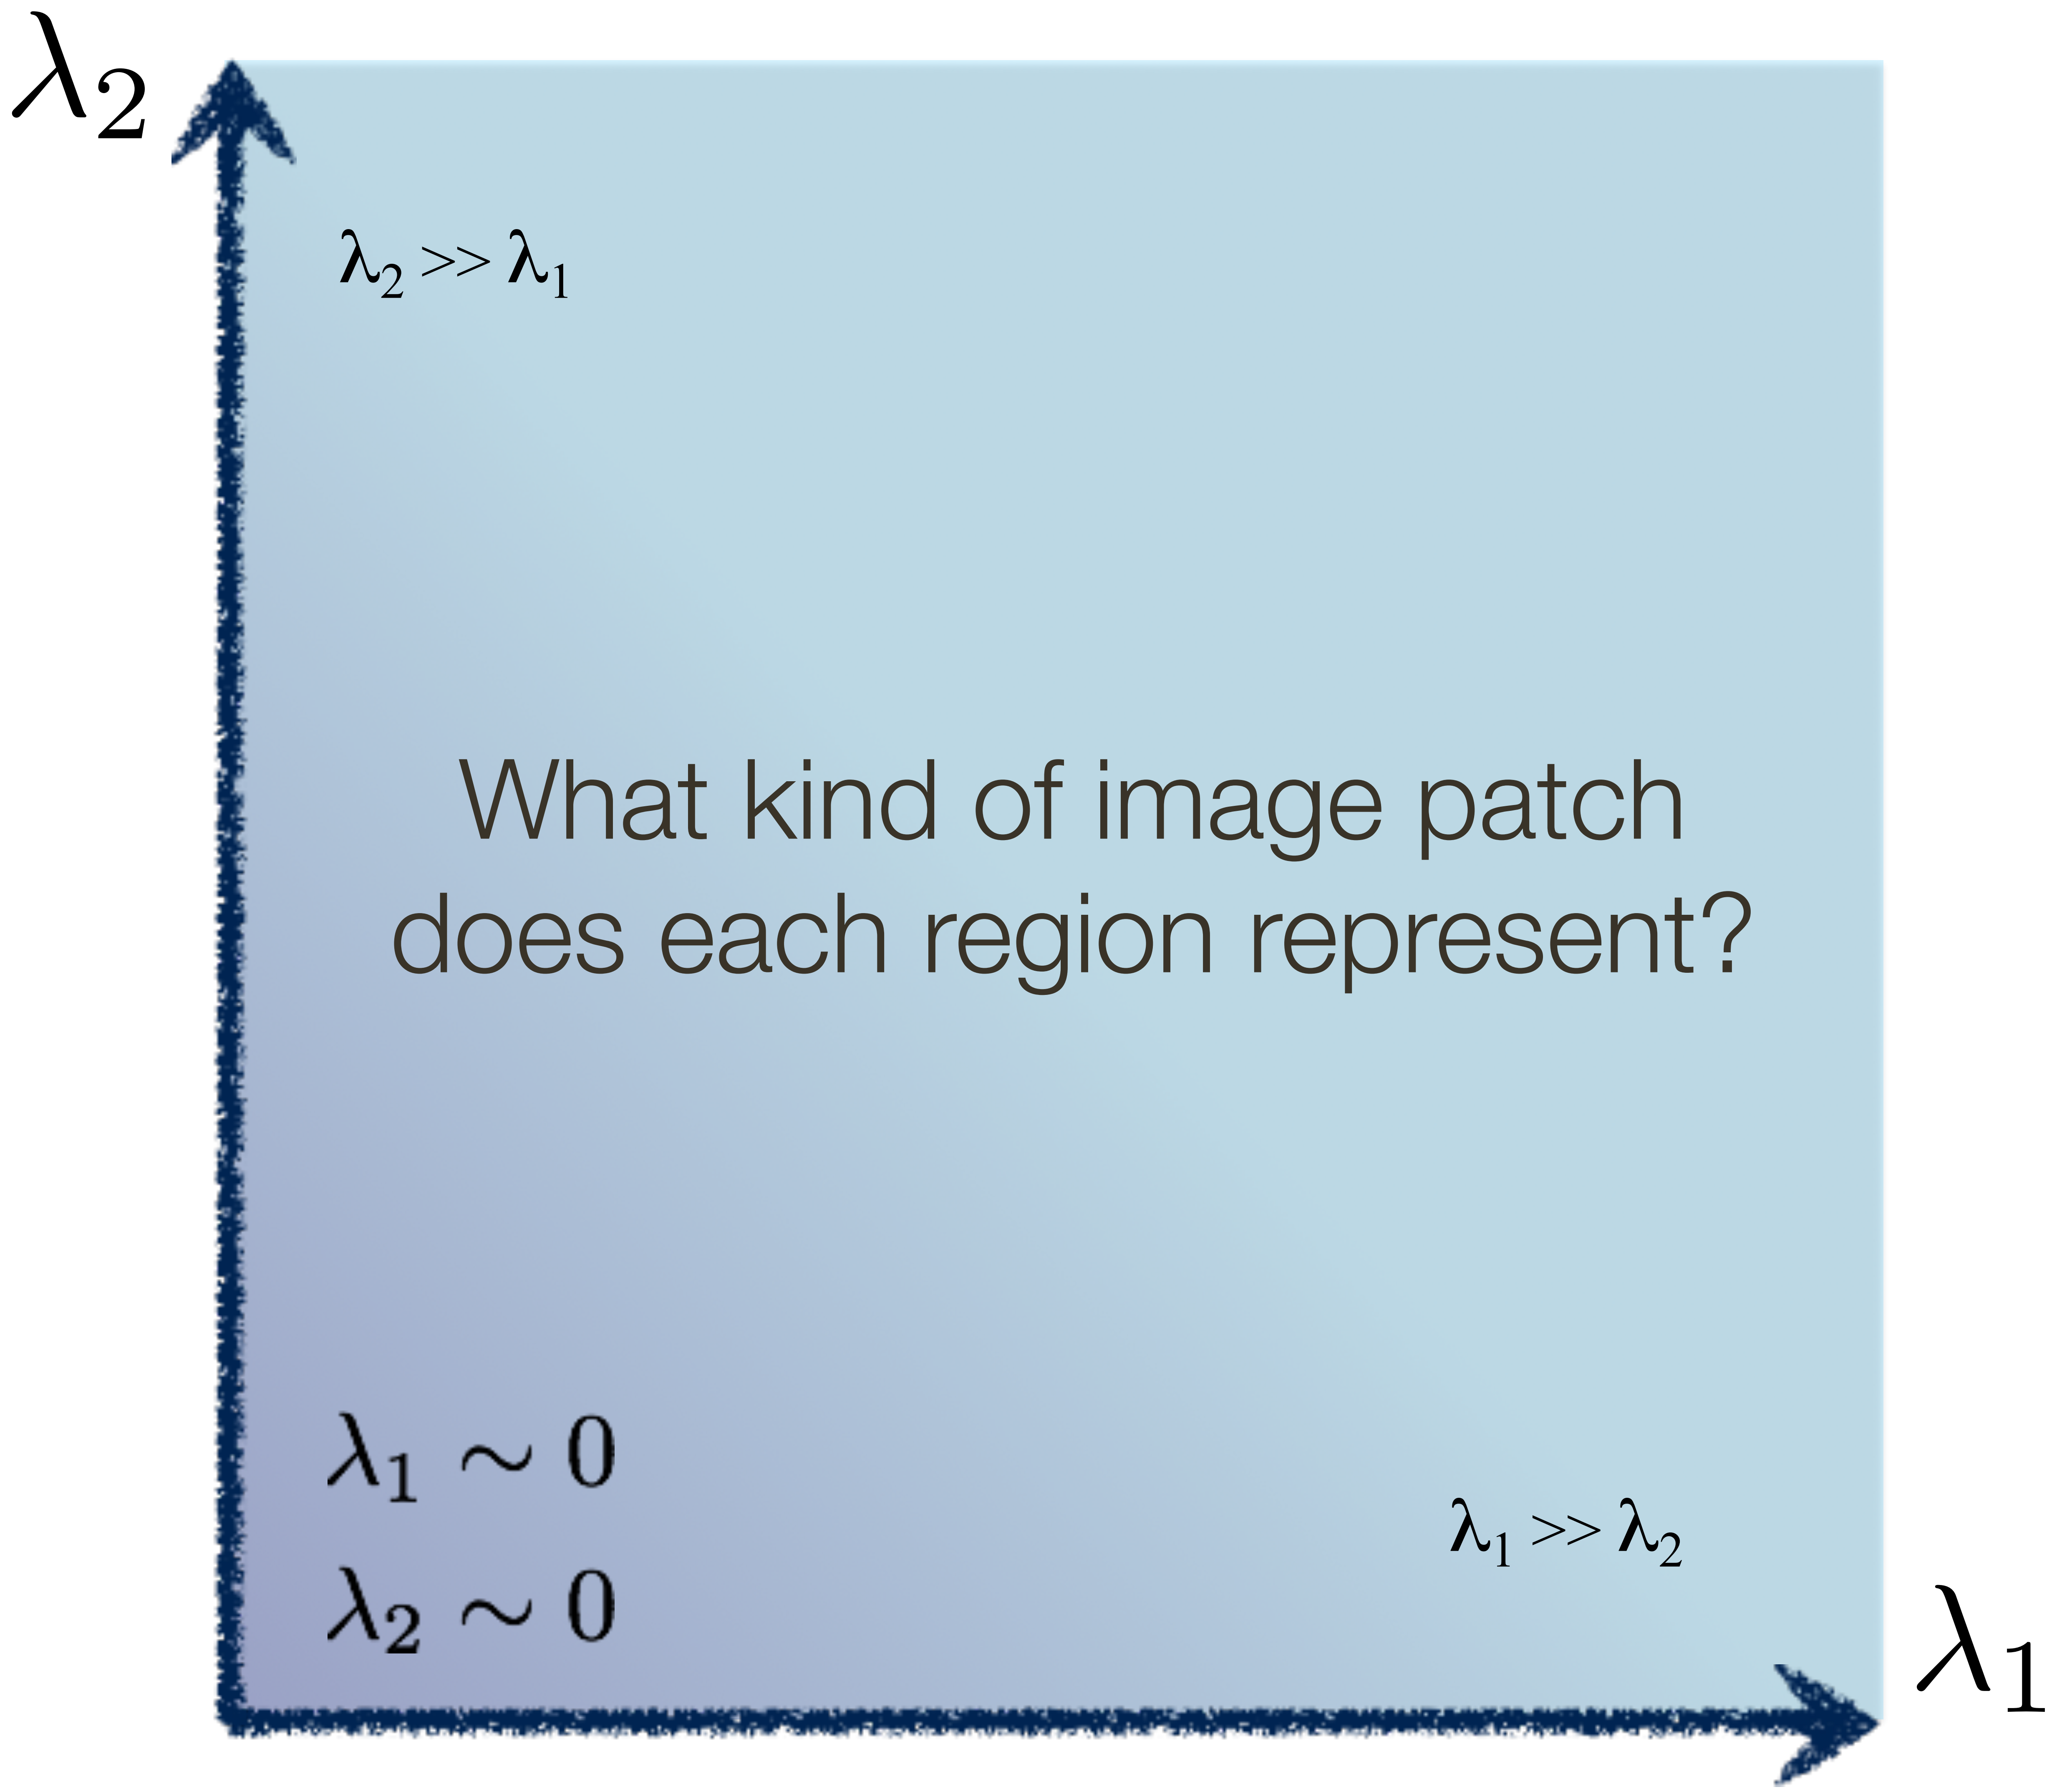

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Interpreting **Eigenvalues**

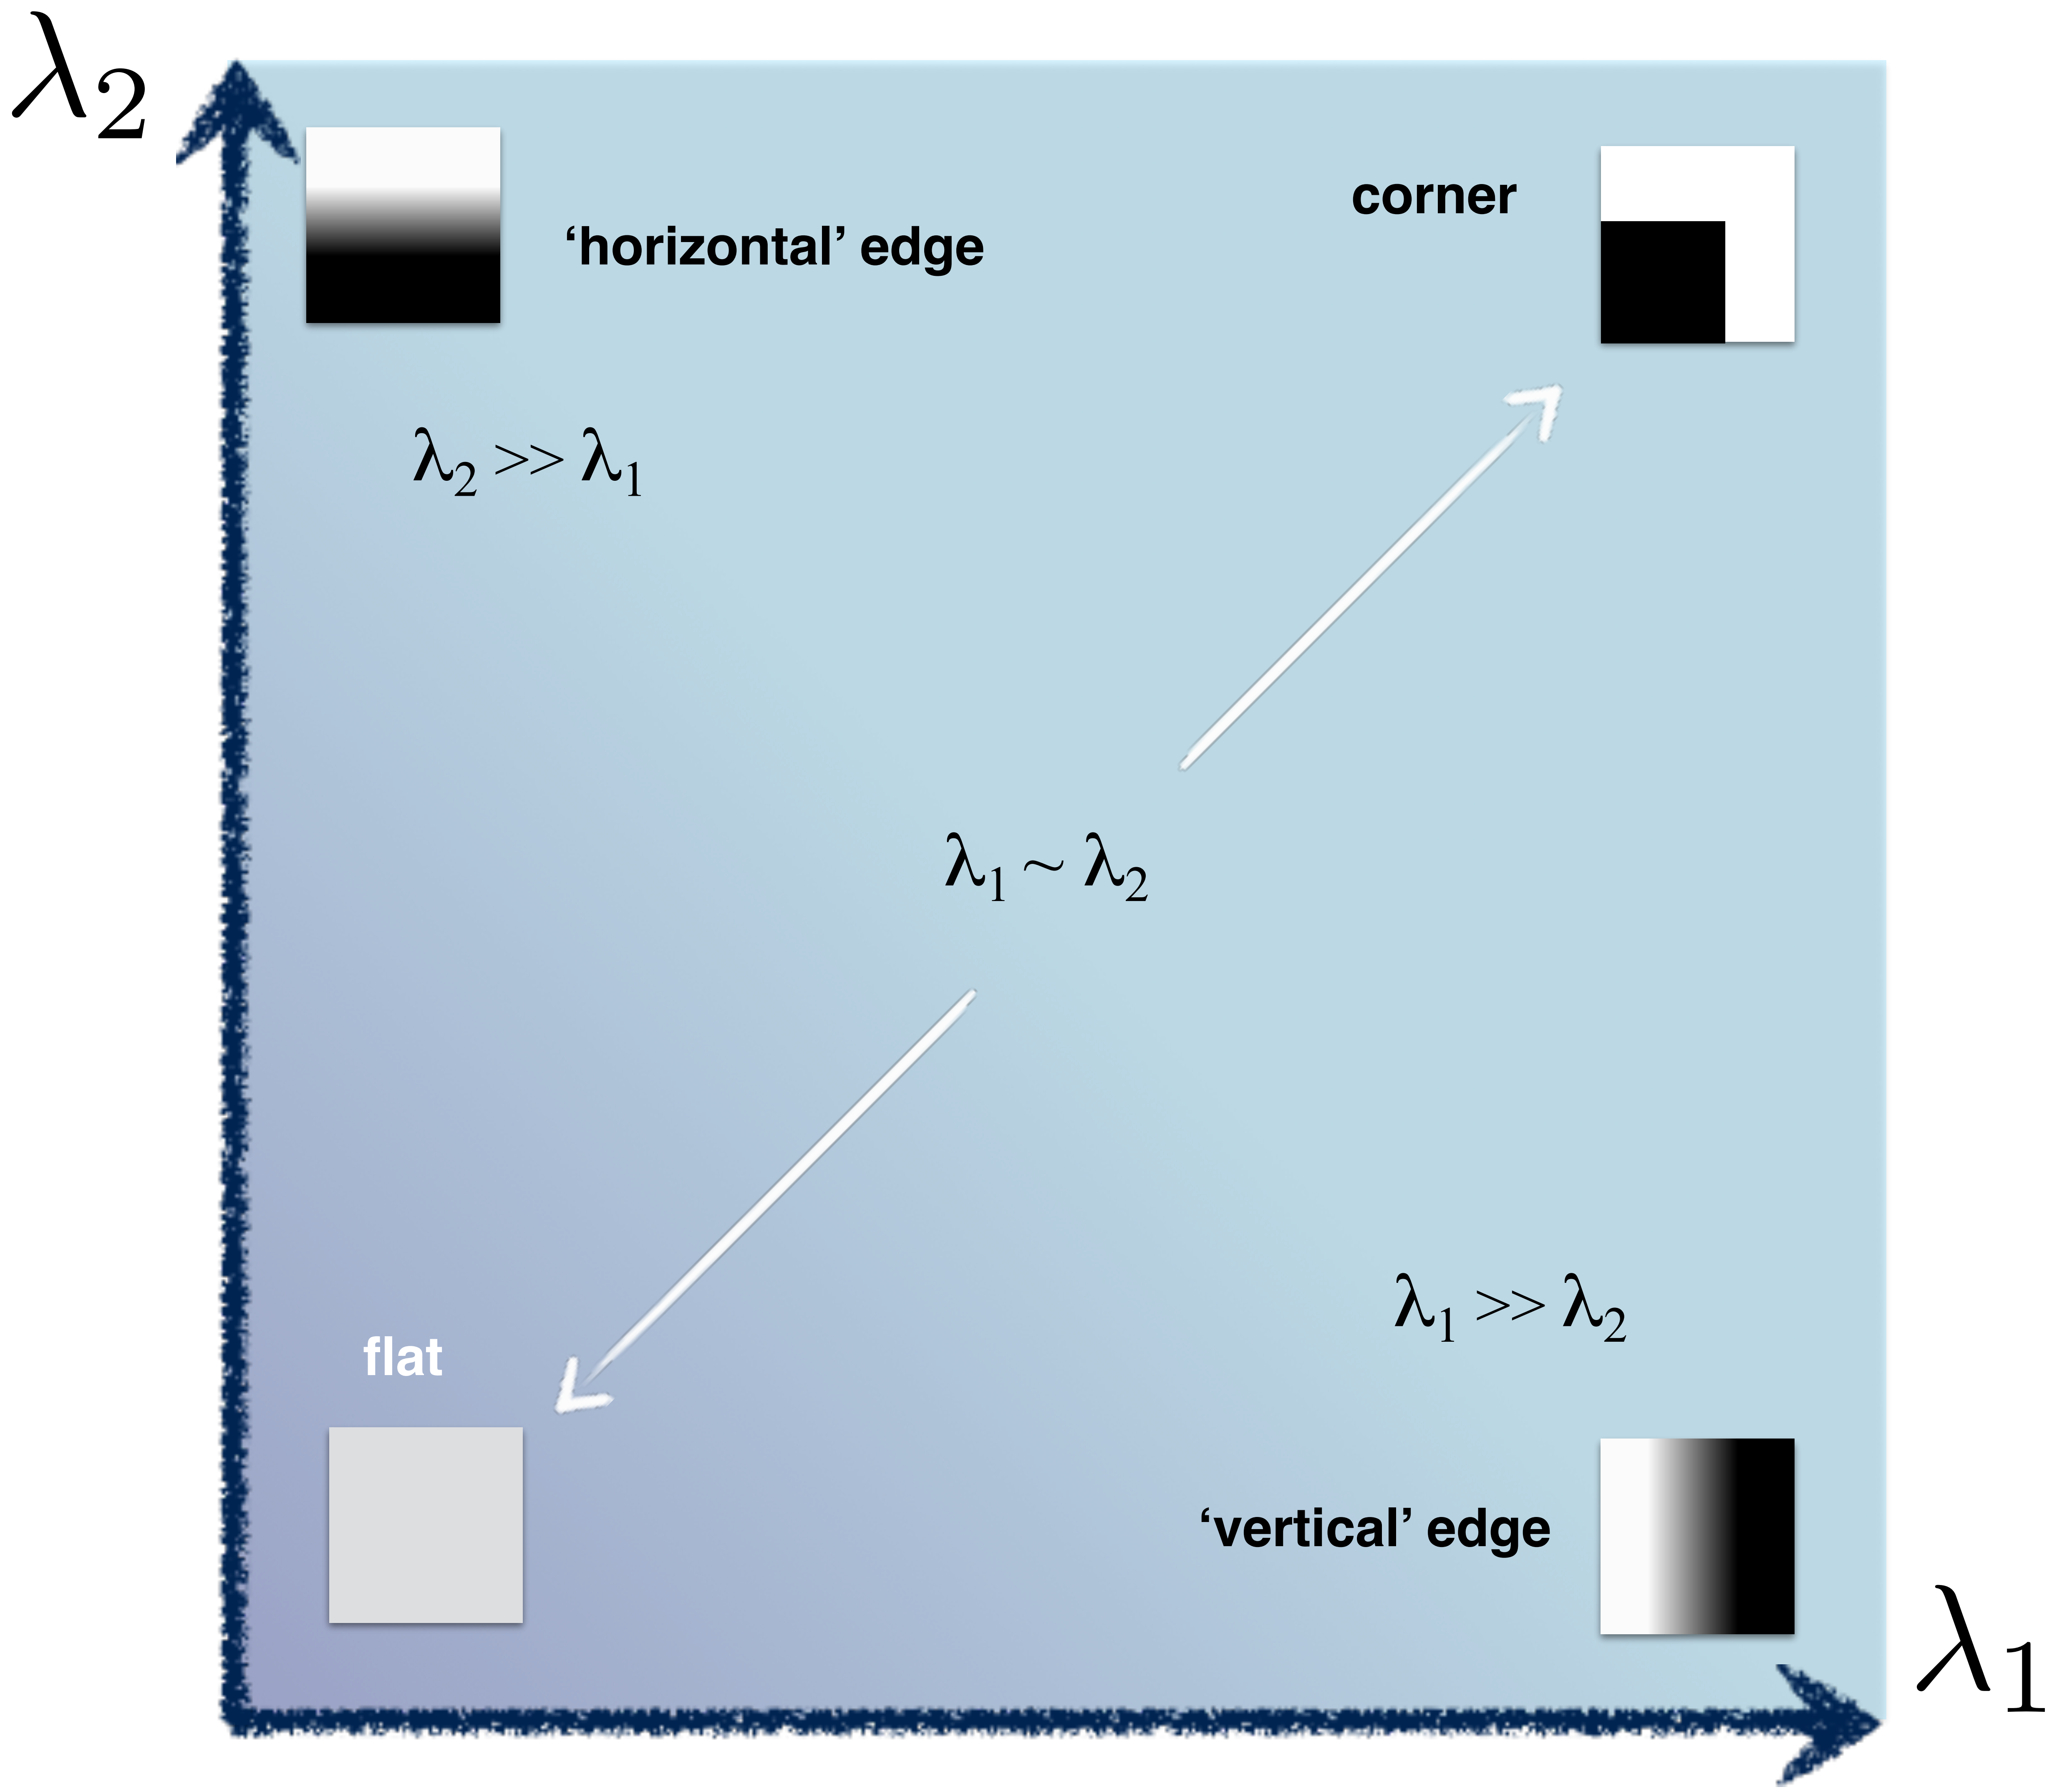

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Interpreting **Eigenvalues**

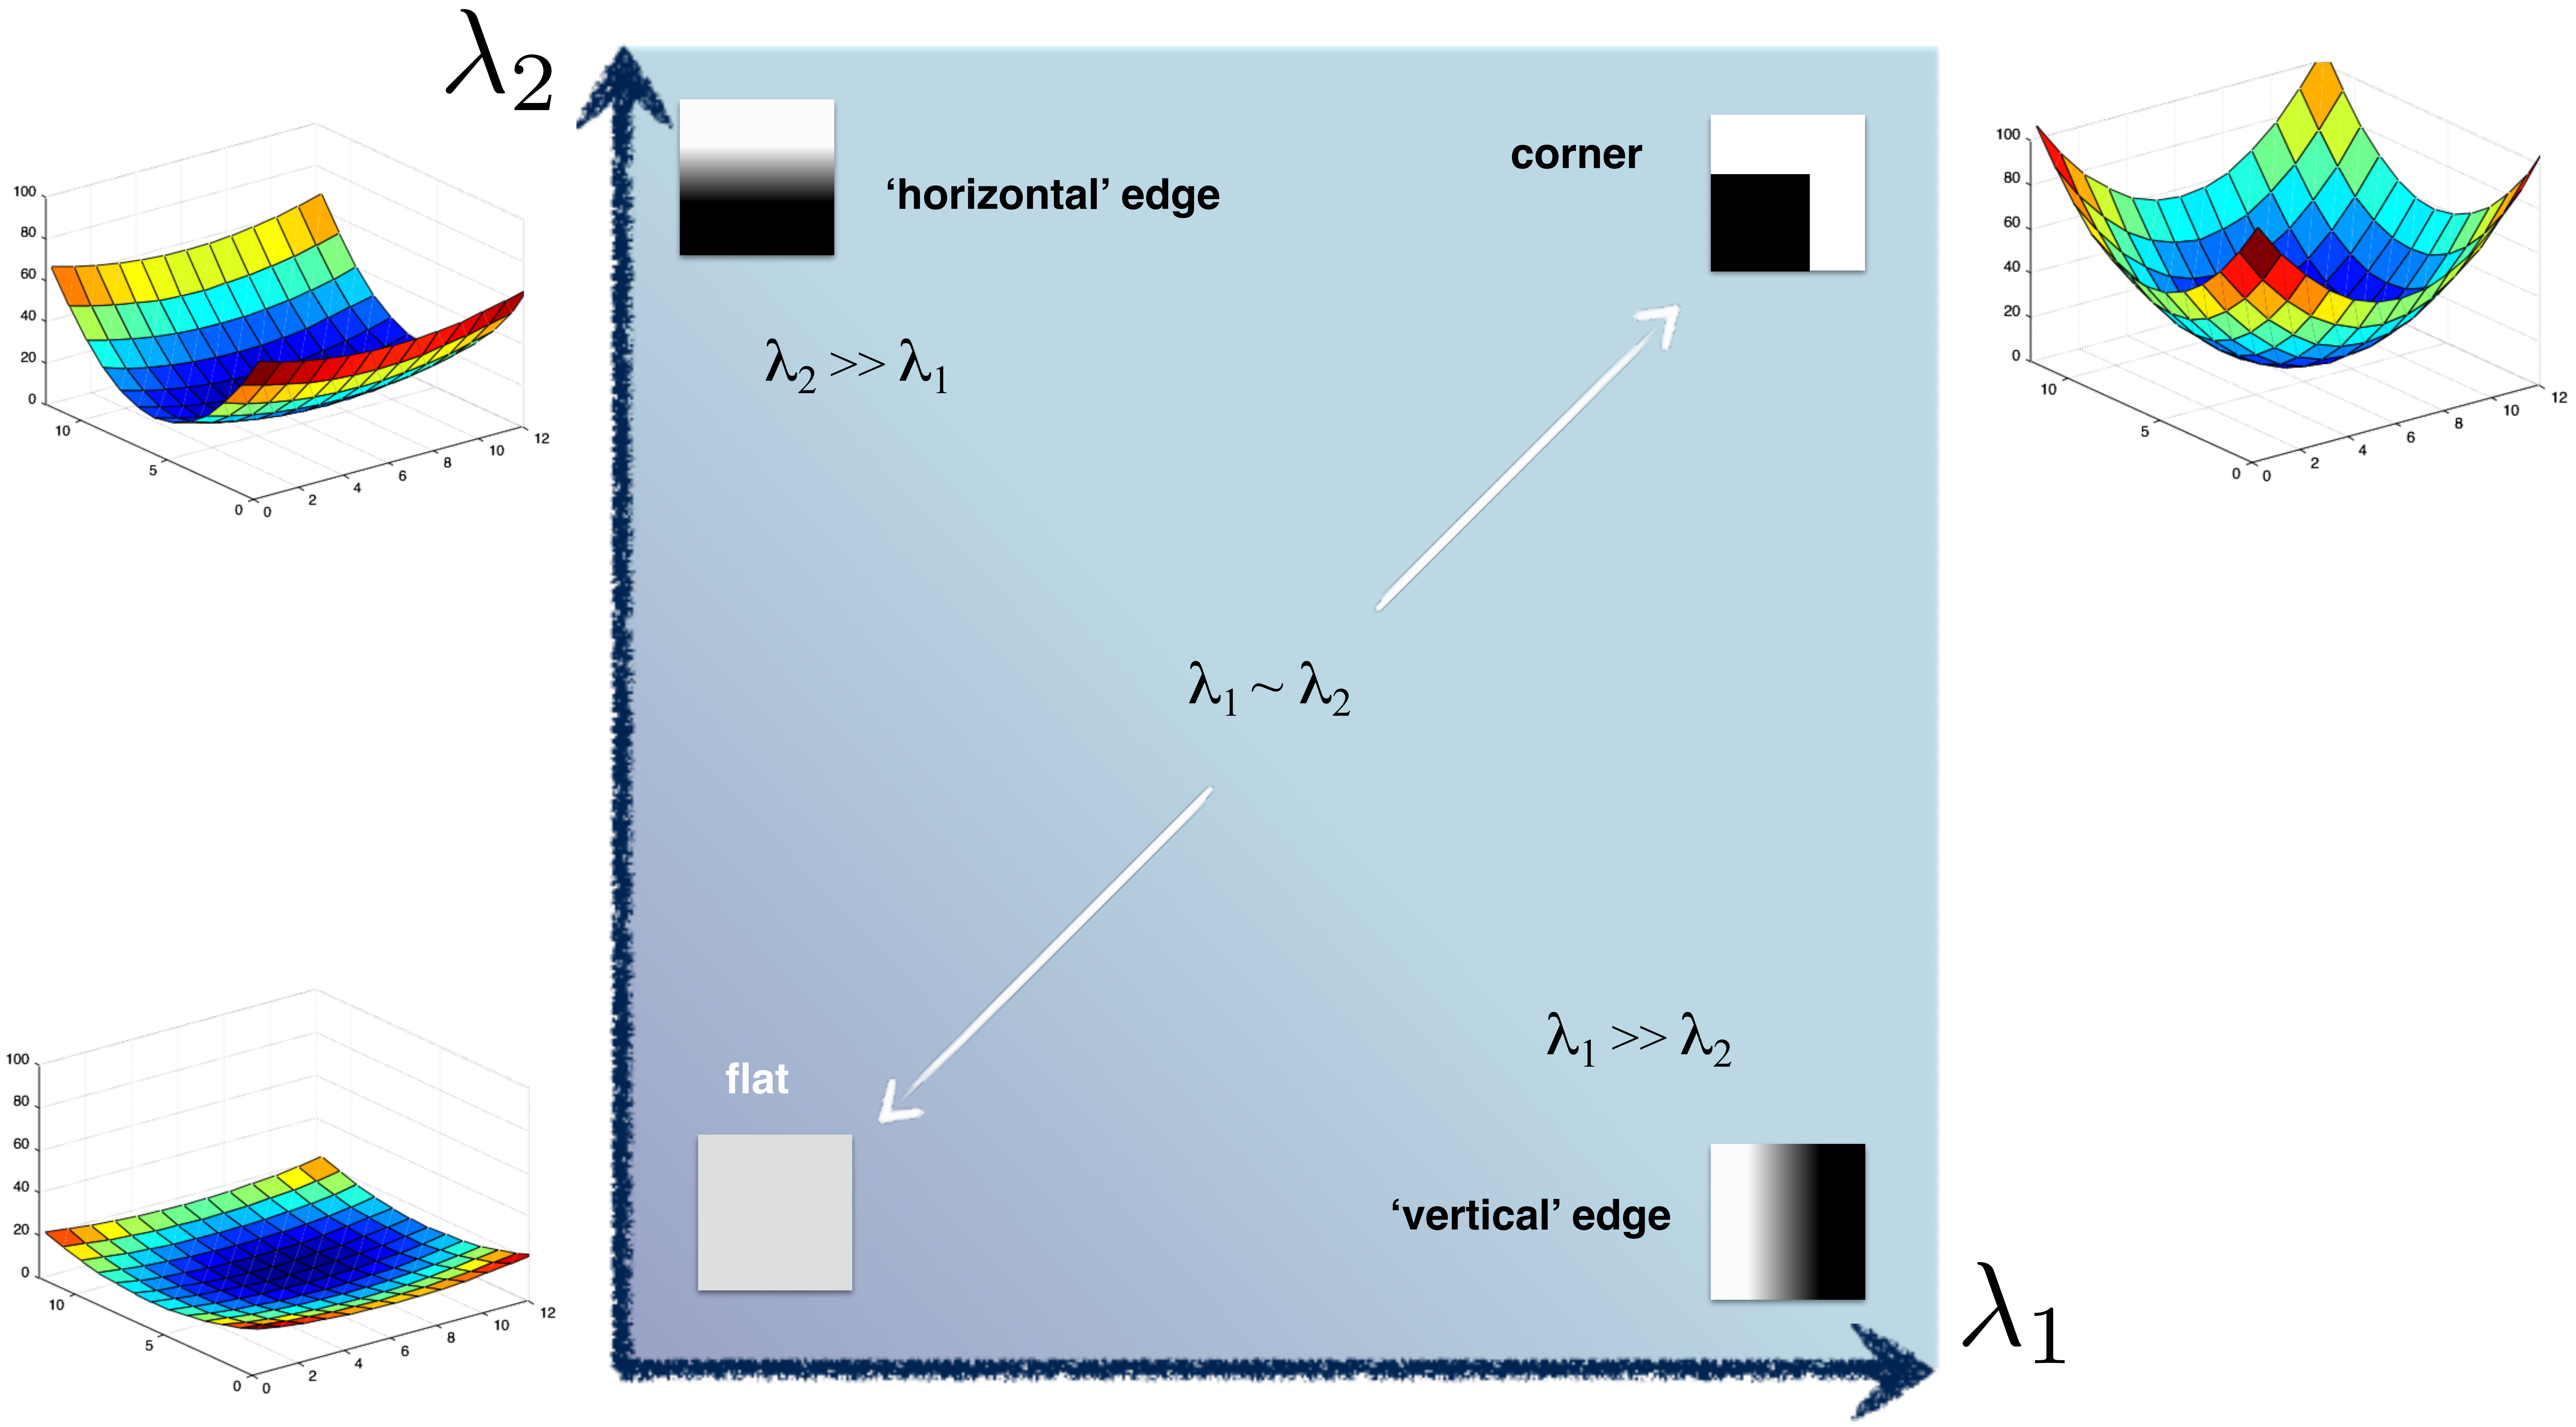

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Interpreting **Eigenvalues**

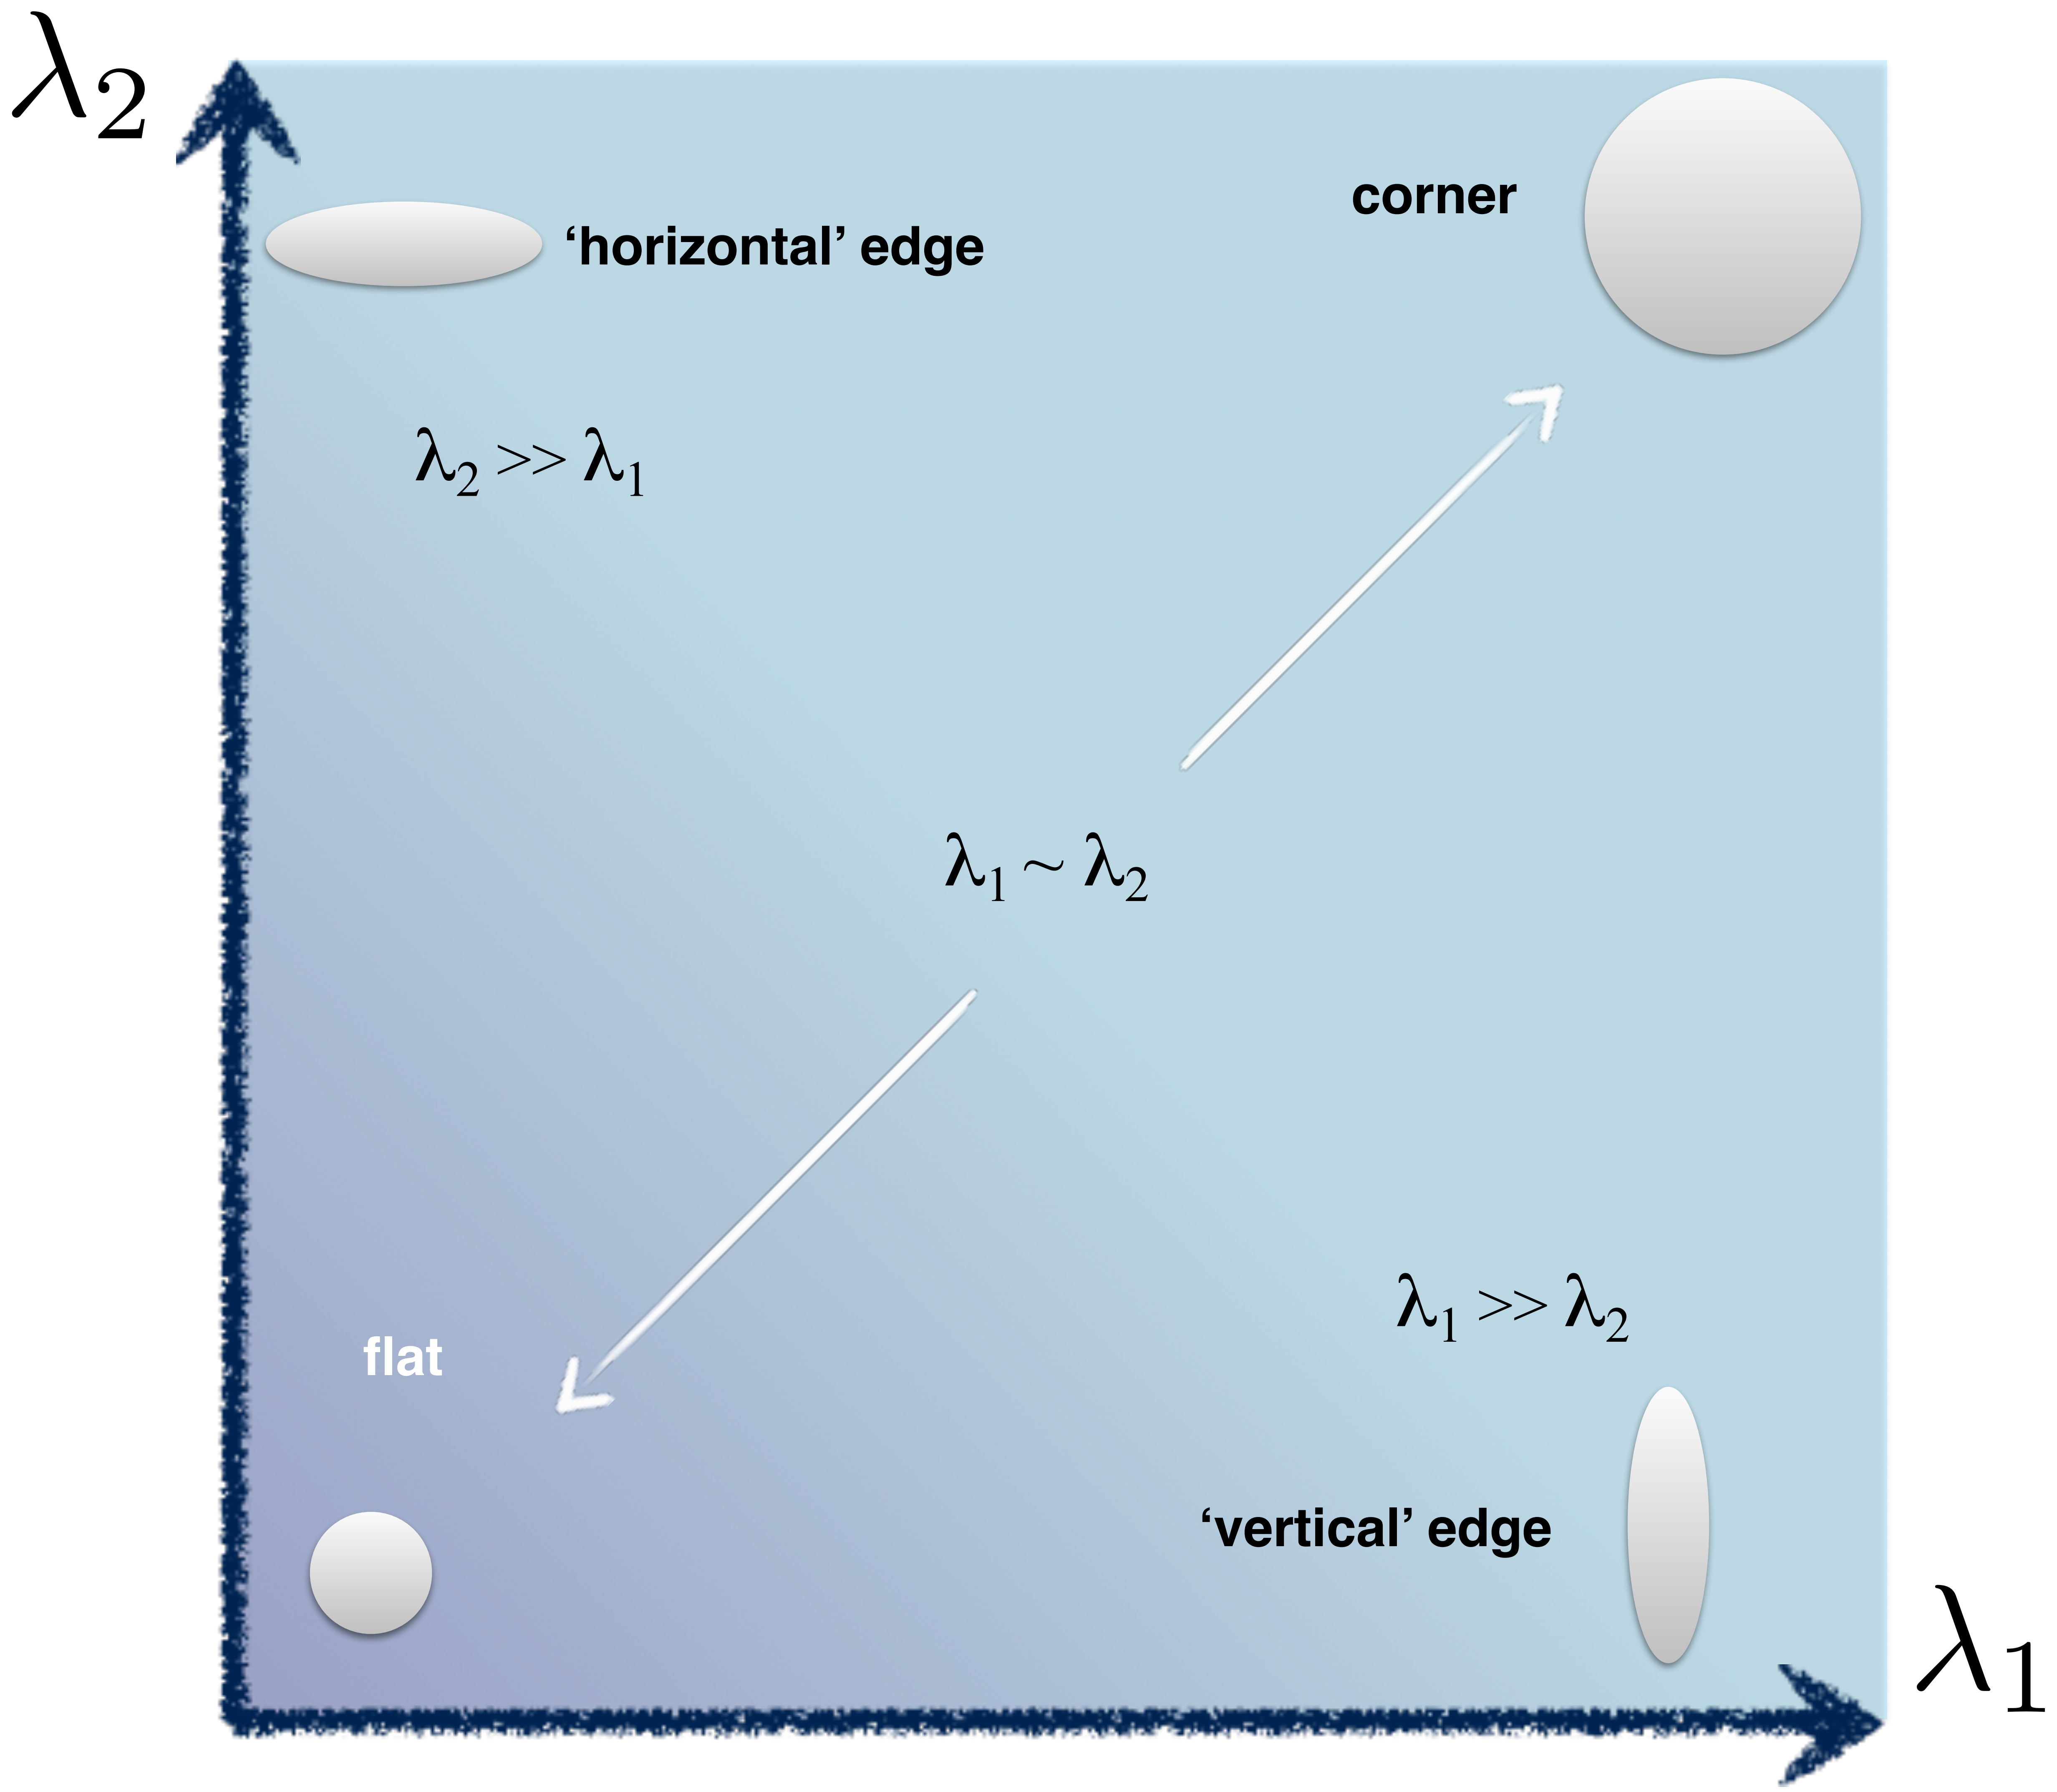

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)