Scale Invariance and Automatic Scale Selection

EECS 504
Foundations of Computer Vision

Instructor: Jason Corso (jjcorso)
web.eecs.umich.edu/~jjcorso/t/

Readings: SZ 4.2, 4.3; FP 5


www.EgyptSons.com

Scale invariant detection

Suppose you are looking for corners

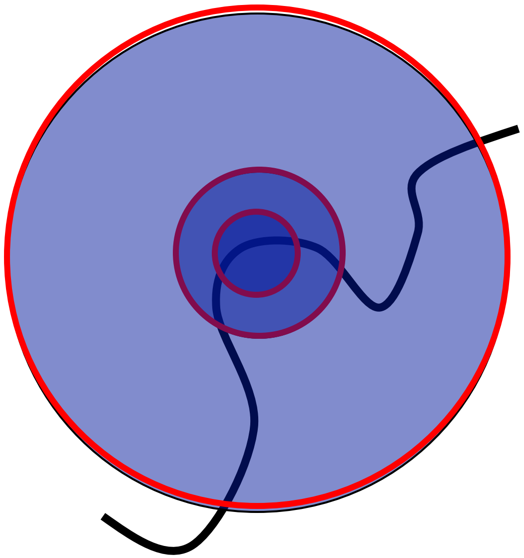

Key idea: find scale that gives local max/min of f

- f is a local maximum/minimum in both position and scale
- Common definition of f: Laplacian
  (or difference between two Gaussian filtered images with different sigmas)

Source: S. Seitz.

Scale-Invariant Feature Detection Example

- Recall: Edges…

f

g

$g \otimes f$

$\frac{d}{dx}(g \otimes f)$

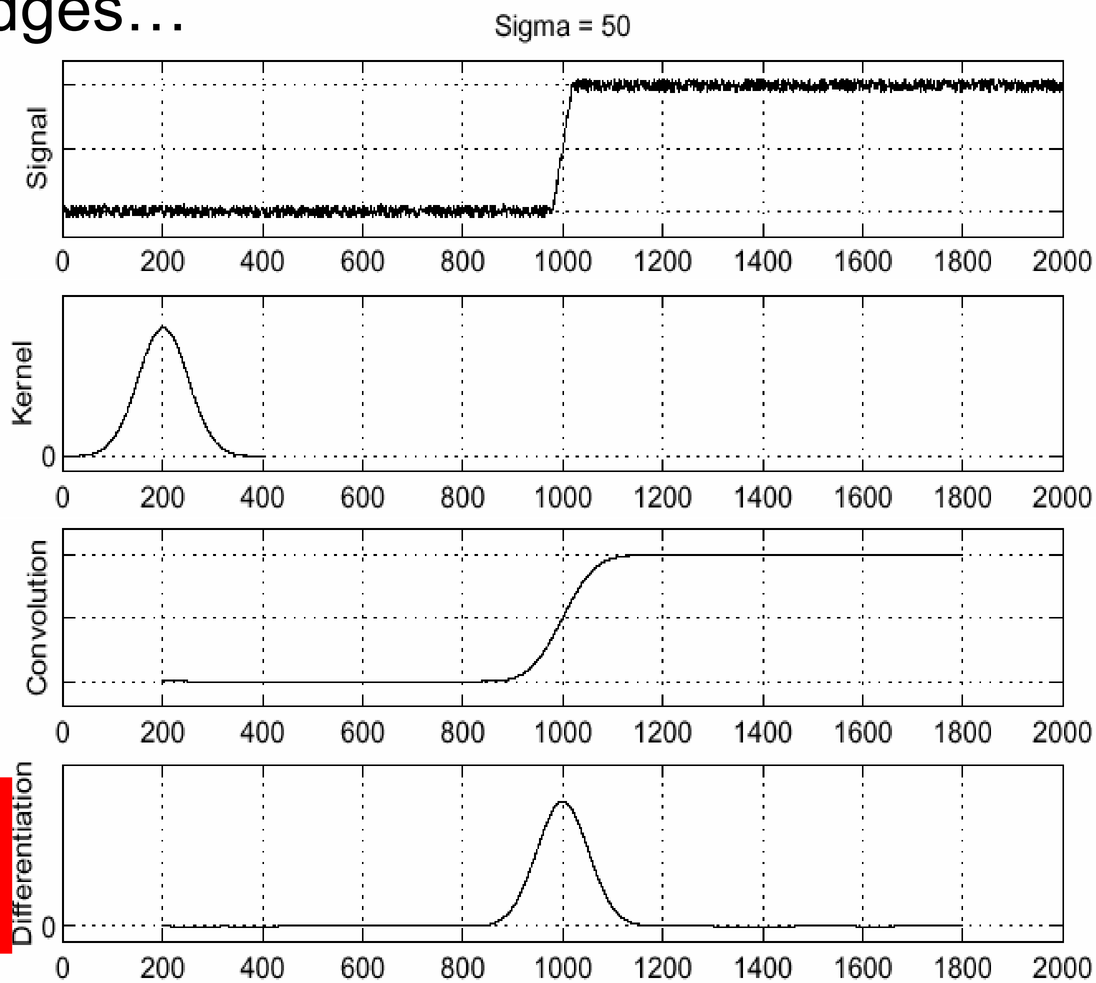

$= \frac{dg}{dx} \otimes f$ Derivative of a Gaussian filter

Source: S. Seitz.

Edge detection

f — Edge

$\frac{d}{dx}g$ — Derivative of Gaussian

$\frac{d}{dx}g \otimes f$ — Edge = maximum of derivative

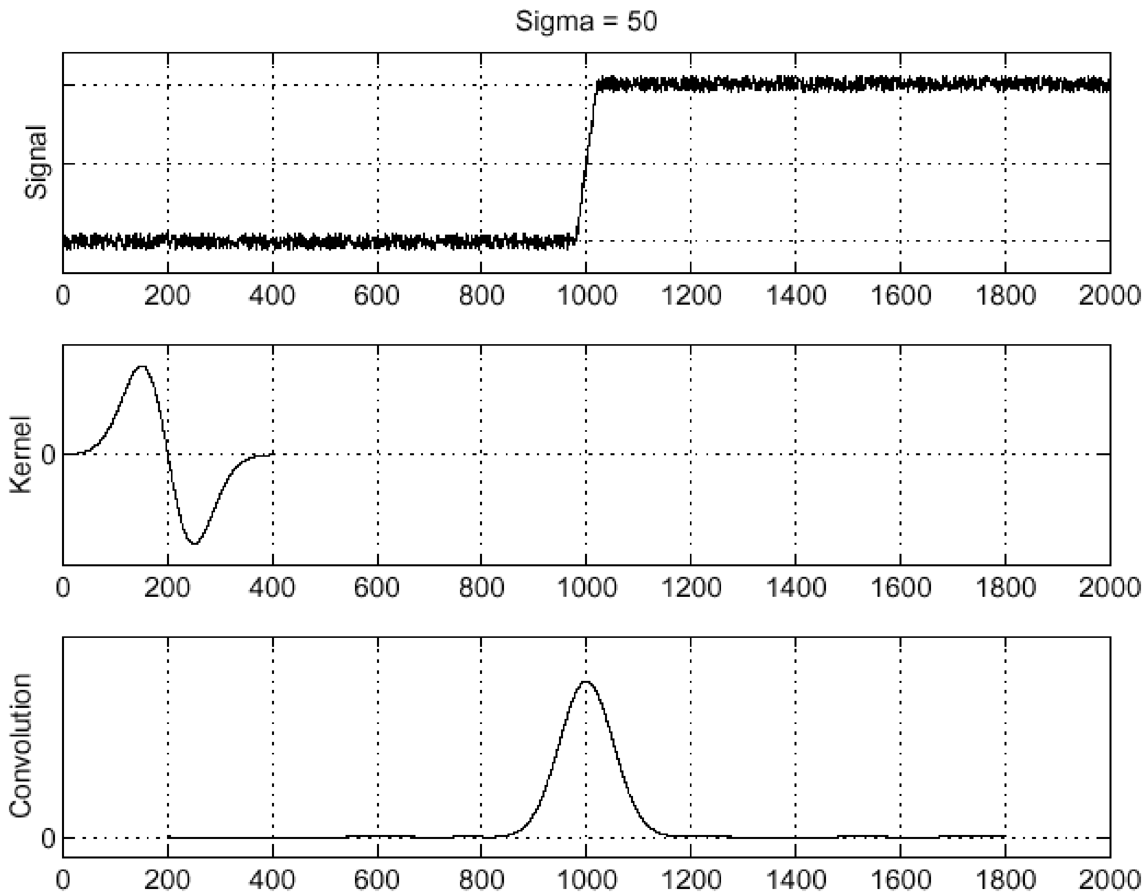

Source: S. Seitz.

Edge detection as zero crossing

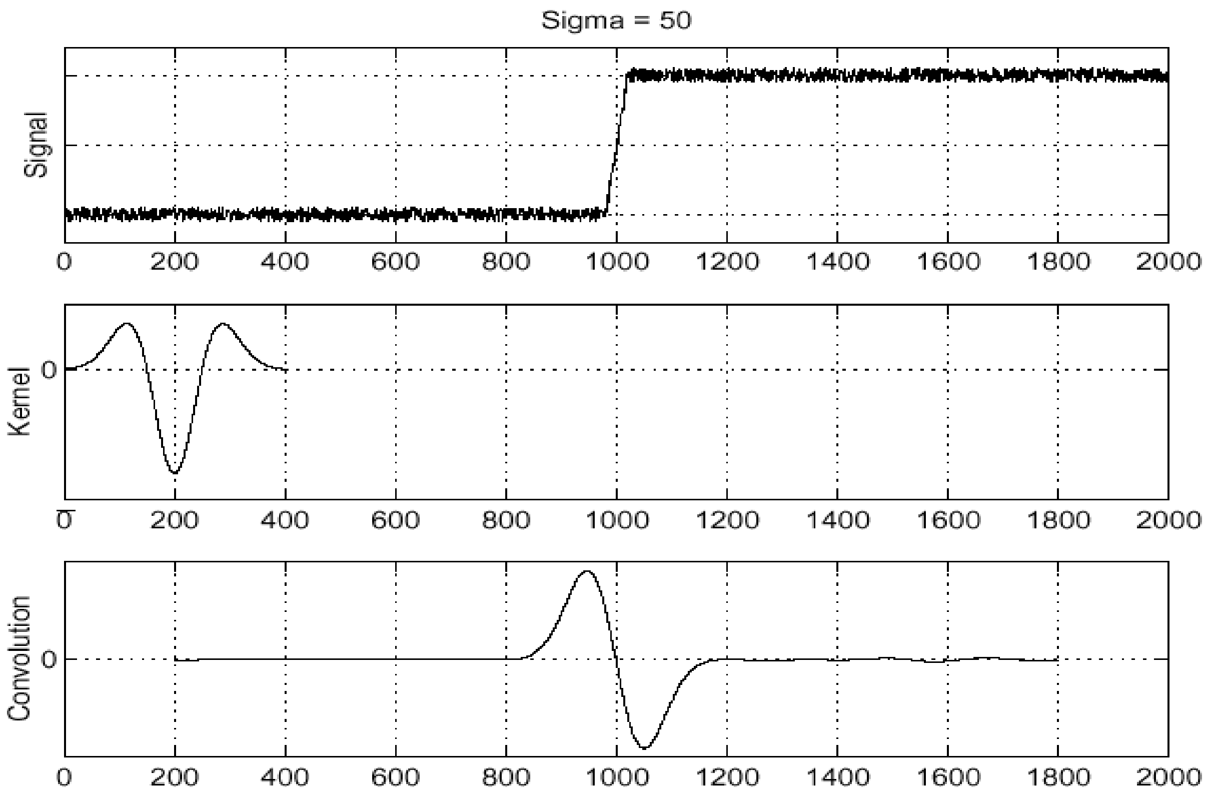

f — Edge

$$\frac{d^2}{dx^2}g$$ — Second derivative of Gaussian

$$\frac{d^2}{dx^2}g \otimes f$$ — Edge = zero crossing of second derivative

Source: S. Seitz.

Edge detection as zero crossing

Ridge? Blob? Two Step Edges?

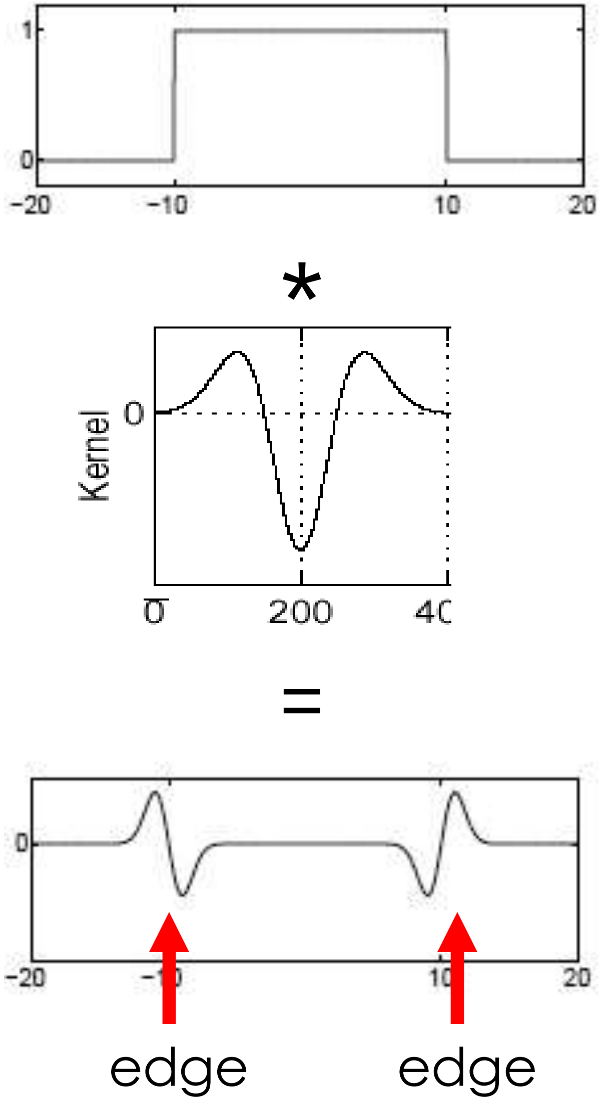

Source: S. Seitz.

From edges to blobs

- Blob = superposition of nearby edges

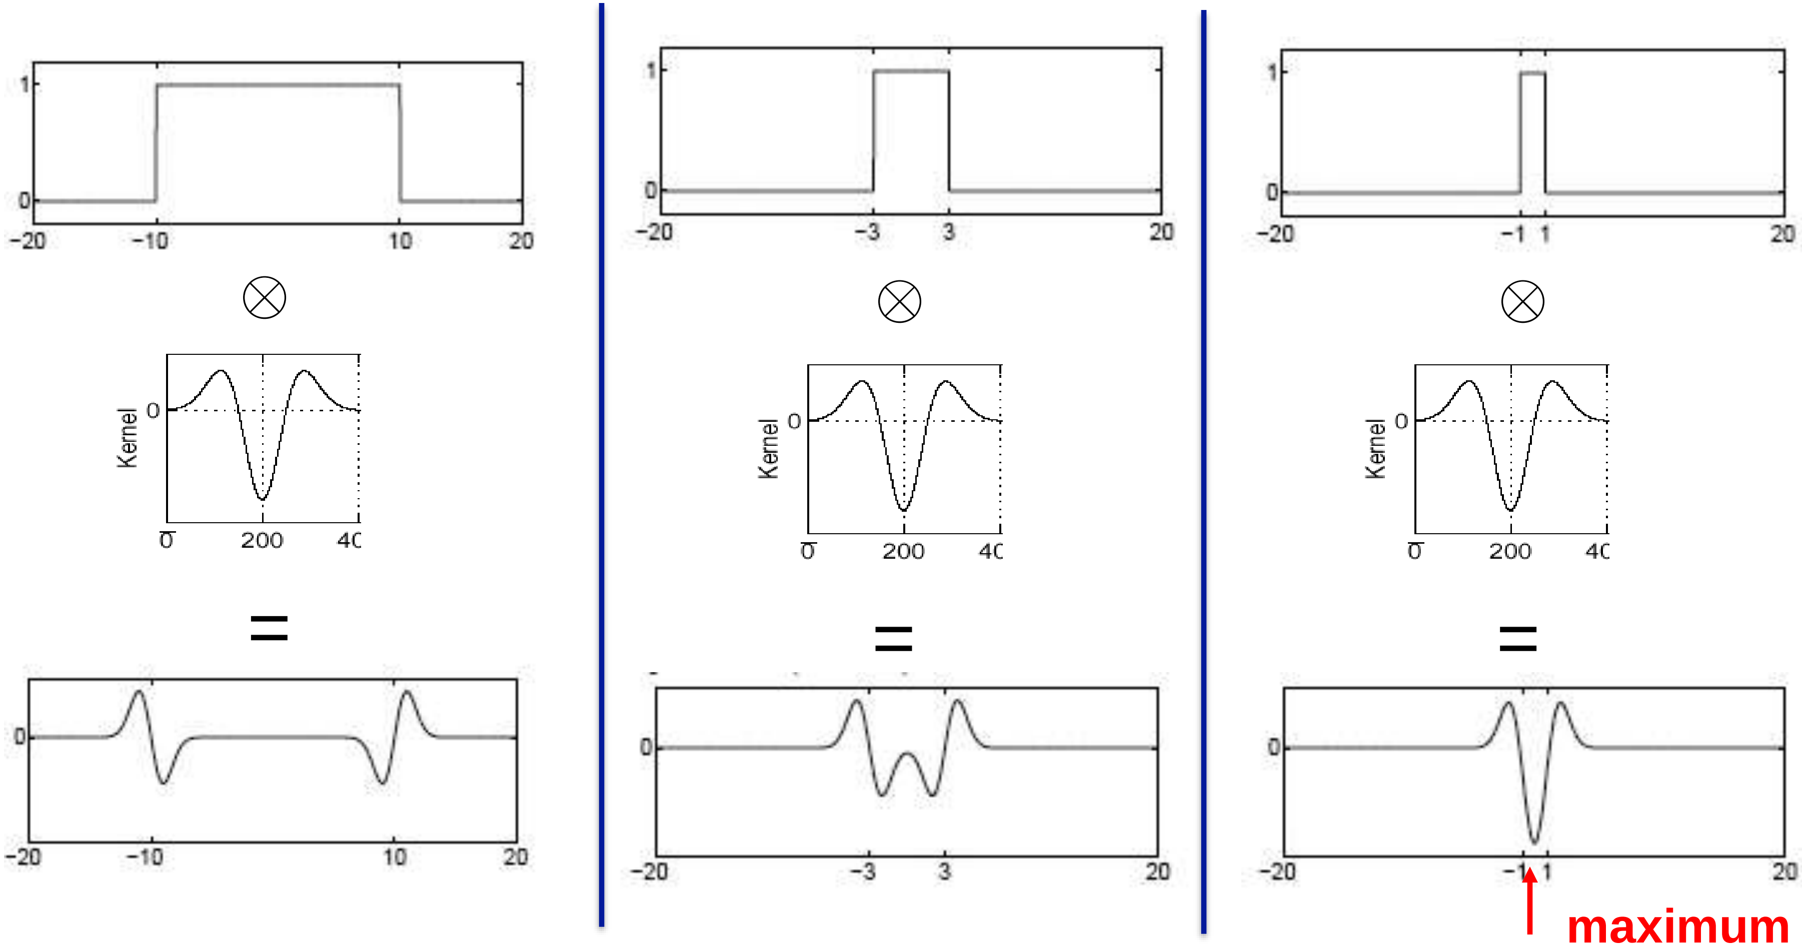

Magnitude of the Laplacian response achieves a maximum at the center of the blob, provided **the scale of the Laplacian is “matched” to the scale of the blob**

Source: S. Seitz.

From edges to blobs

- Blob = superposition of nearby edges

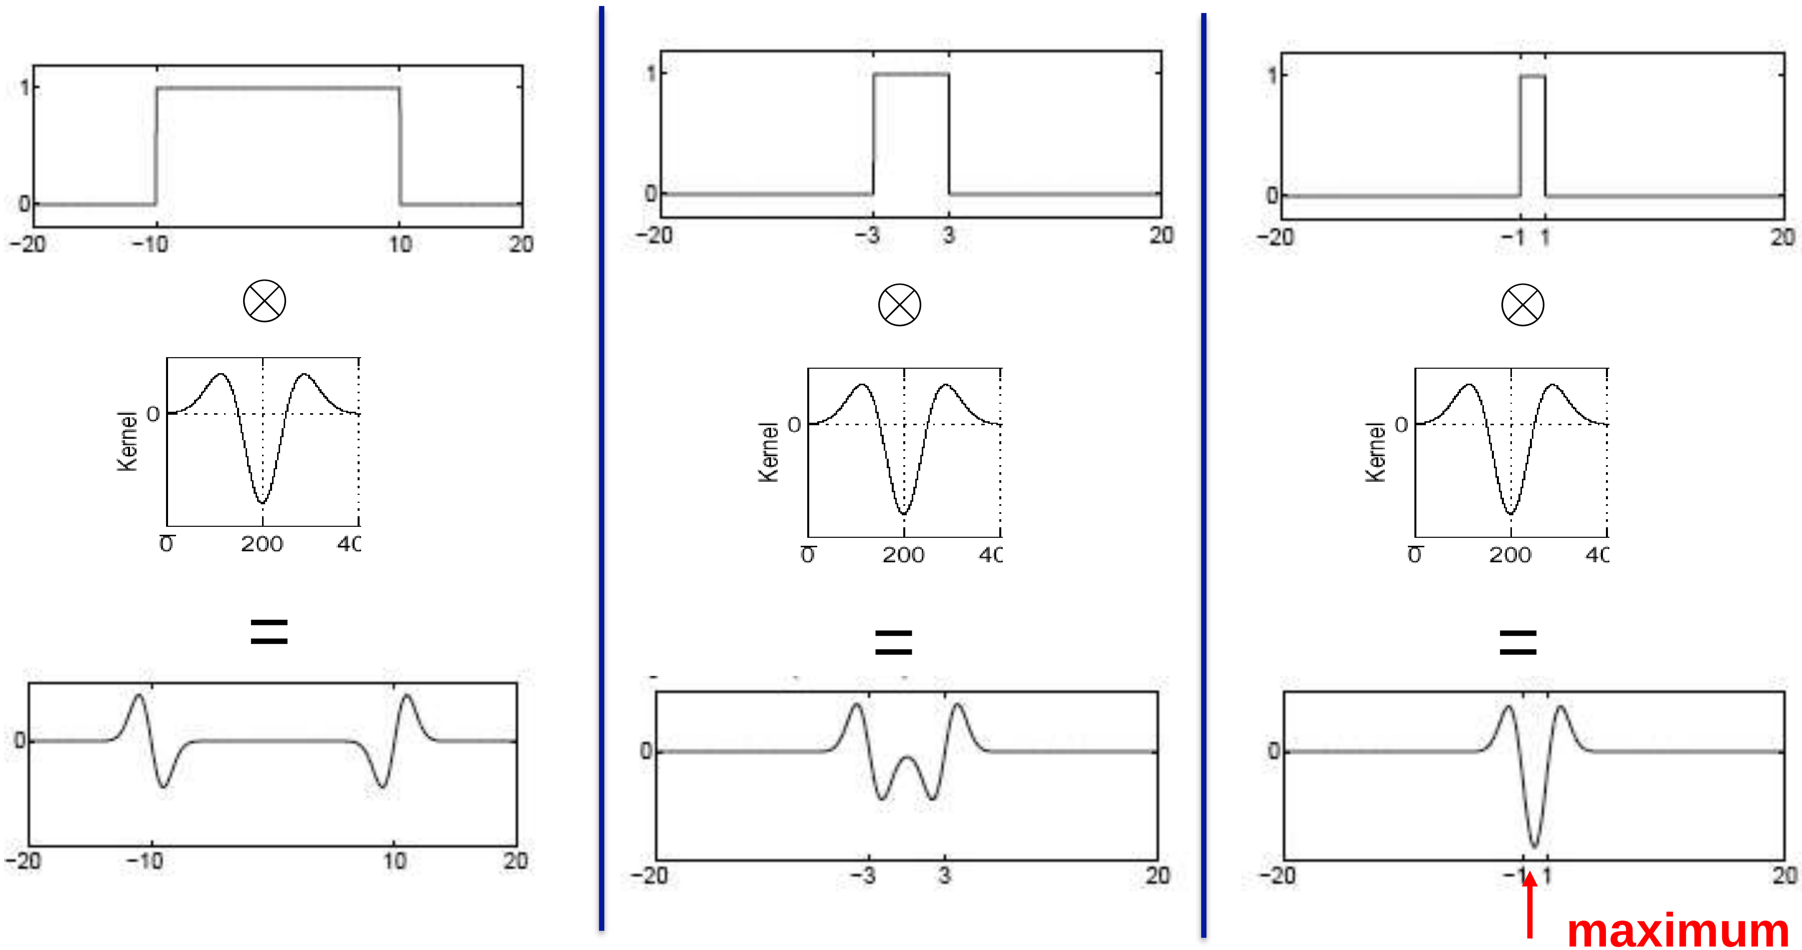

What if the blob is slightly thicker or slimmer?

Source: S. Seitz.

Stop Slides

Filter Pyramids

- Recall we can always filter with $\mathcal{G}(\sigma)$ for any σ
- As a result, we can think of a continuum of filtered images as σ grows.
  - This is referred to as the "scale space" of the images. We will see this show up several times.
- As a related note, suppose I want to subsample images
  - Subsampling reduces the highest frequencies
  - Averaging reduces noise
  - Pyramids are a way of doing both

Source: G Hager Slides

Gaussian Pyramid

- Algorithm:
  - 1. Filter with $\mathcal{G}(\sigma = 1)$
  - 2. Resample at every other pixel
  - 3. Repeat

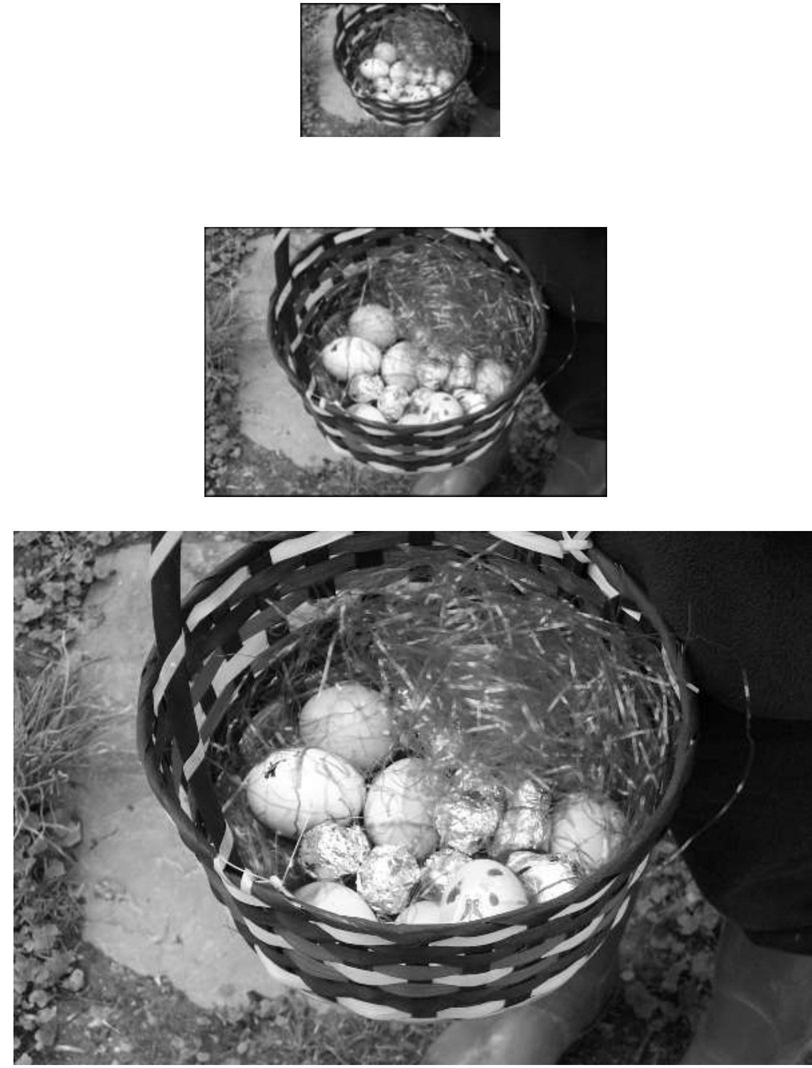

Source: G Hager Slides

Laplacian Pyramid Algorithm

- Create a Gaussian pyramid by successive smoothing with a Gaussian and down sampling

- Set the coarsest layer of the Laplacian pyramid to be the coarsest layer of the Gaussian pyramid

- For each subsequent layer n+1, compute

$$L(n+1) = G(n+1) = \text{Upsample}(G(n))$$

Source: G Hager Slides

Laplacian of Gaussian Pyramid

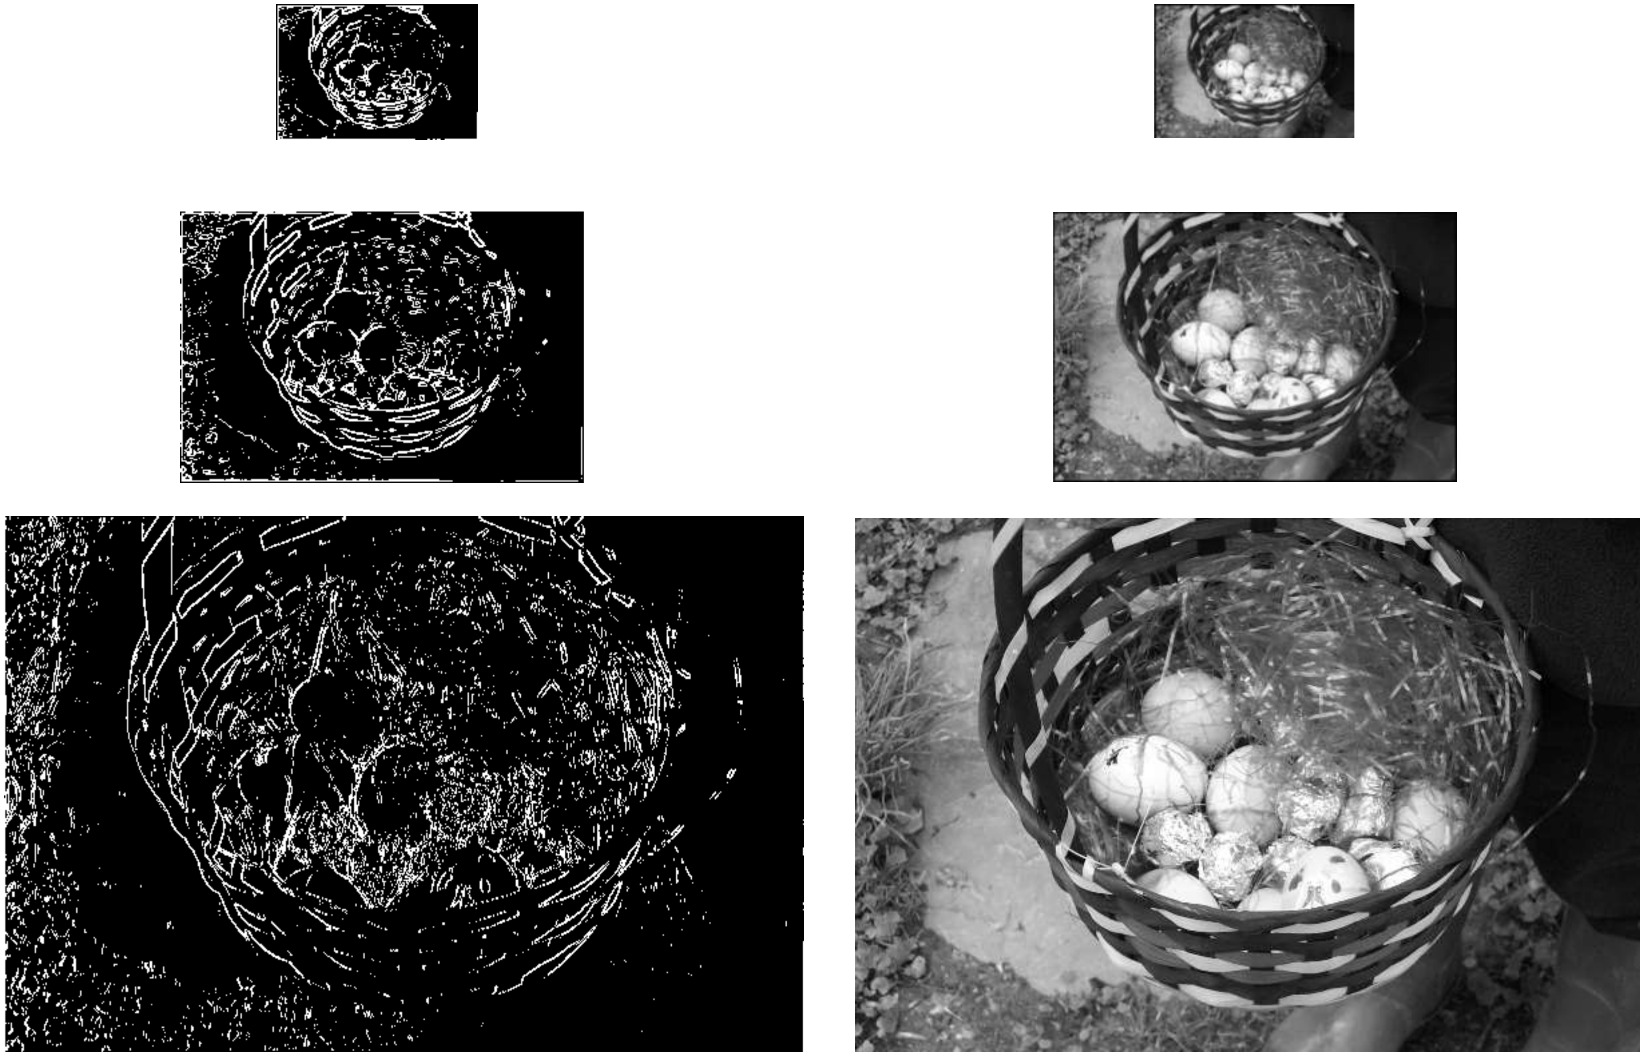

Source: G Hager Slides

Laplacian of Gaussian Pyramid

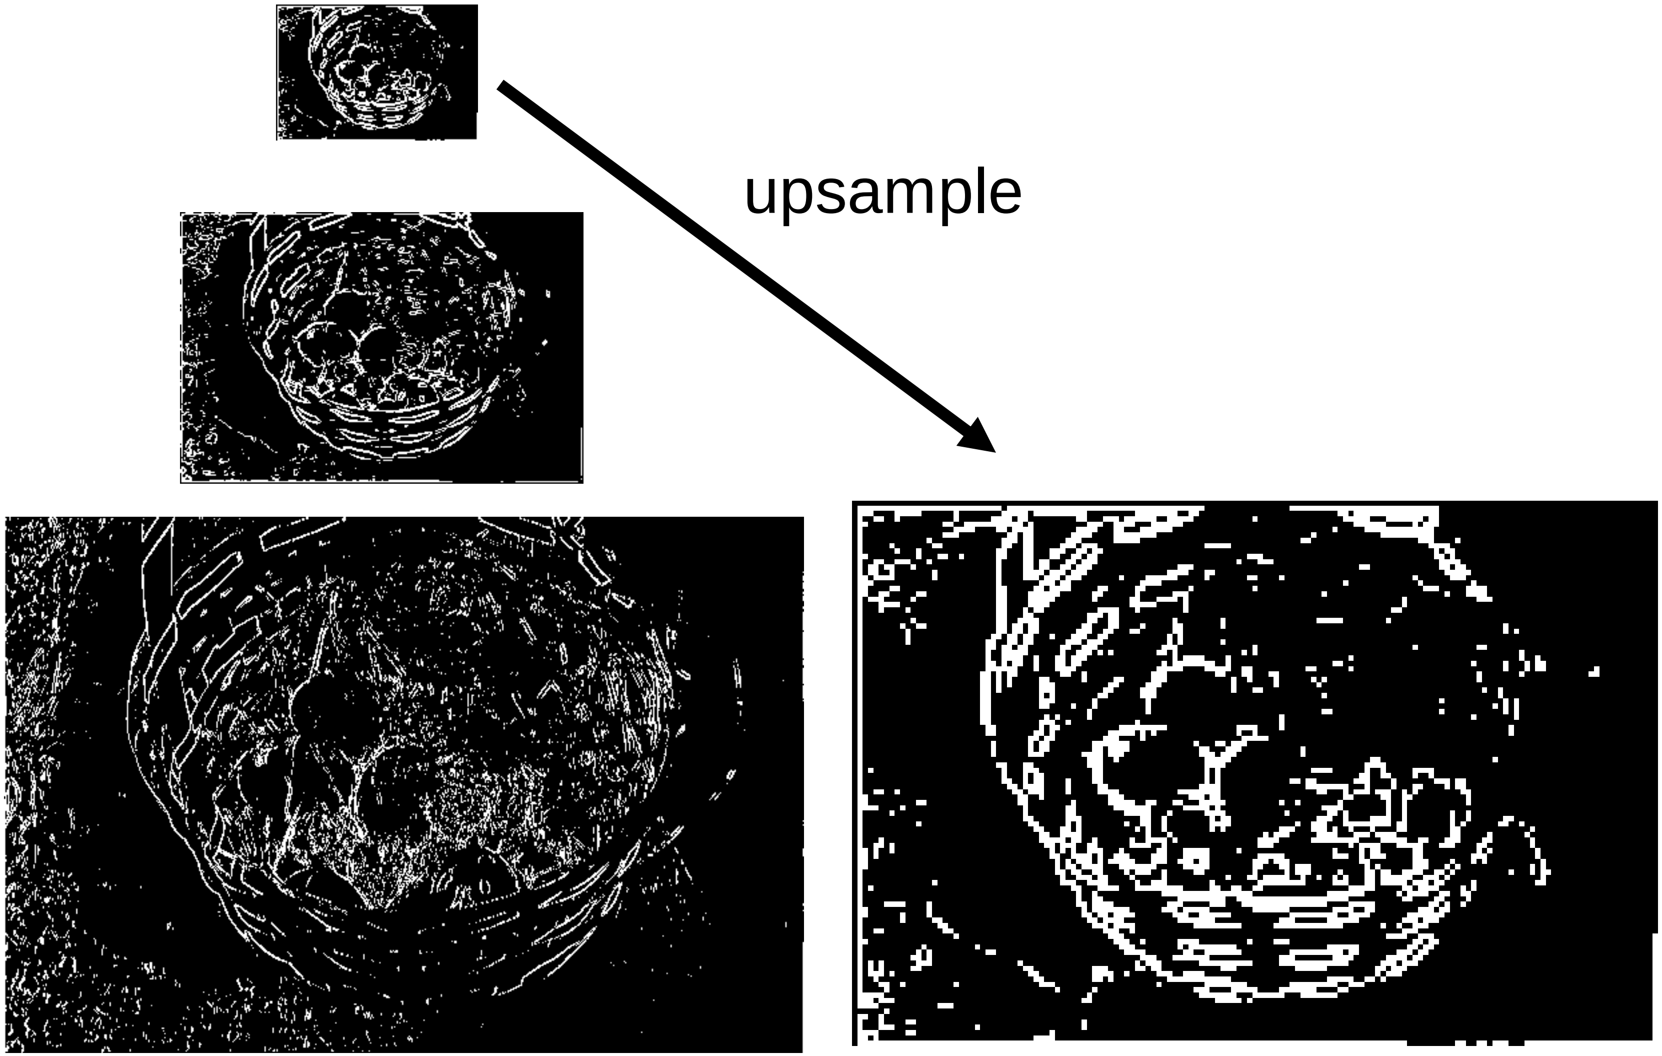

Source: G Hager Slides

Stop Slides

Scale selection

- We want to find the characteristic scale of the blob by convolving it with Laplacians at several scales and looking for the maximum response

increasing σ

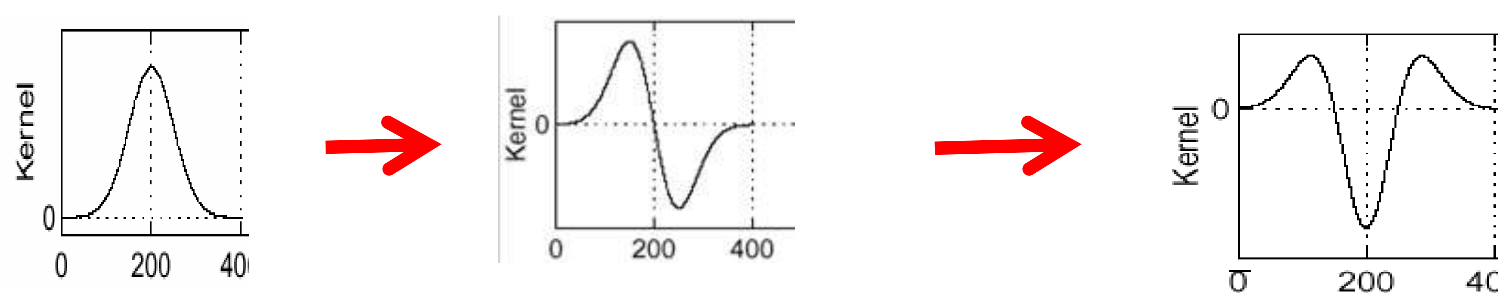

$$g \qquad \frac{d}{dx}g \qquad \frac{d^2}{dx^2}g$$

Source: S. Seitz.

Scale selection

- We want to find the characteristic scale of the blob by convolving it with Laplacians at several scales and looking for the maximum response

•However, Laplacian response decays as scale increases:

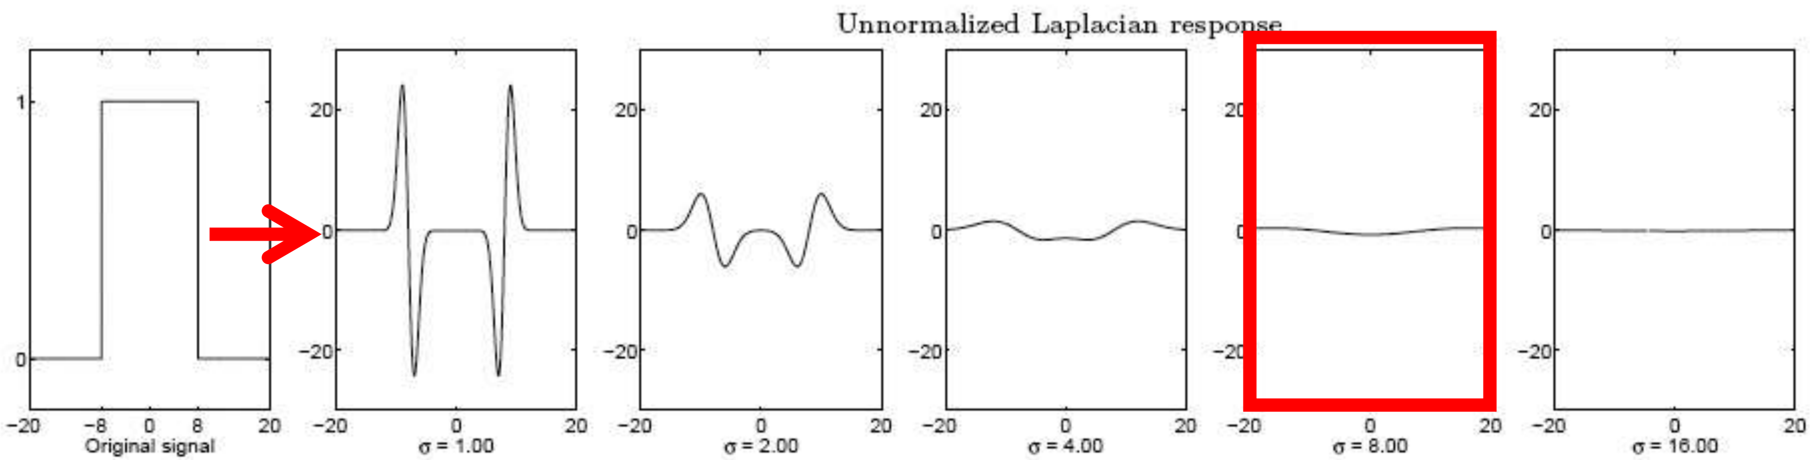

original signal
(radius=8)

This should give the max response ☹

increasing σ →

Source: S. Seitz.

Scale normalization

- The response of a derivative of Gaussian filter to a perfect step edge decreases as σ increases

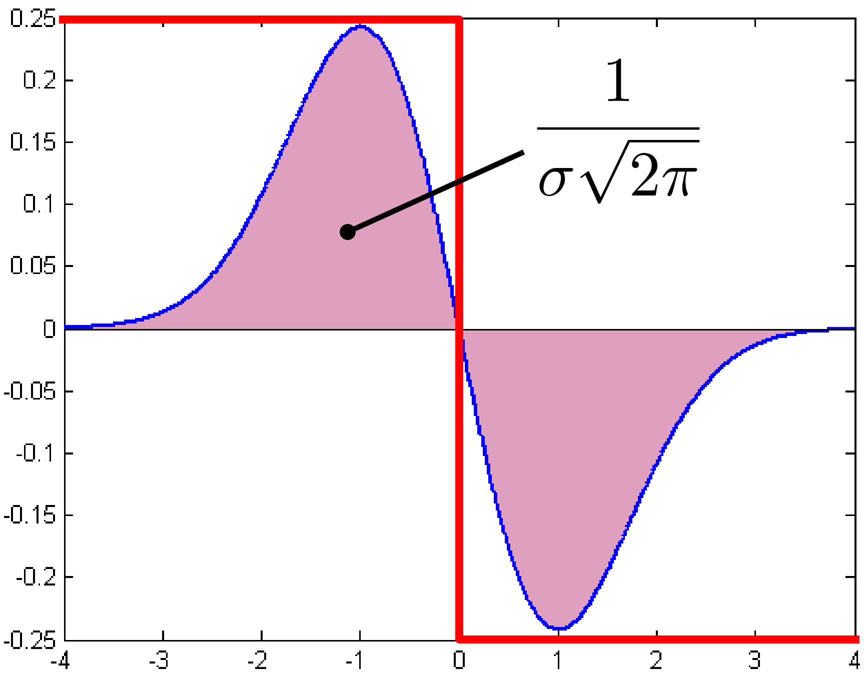

Source: S. Seitz.

Scale normalization

- To keep response the same (scale-invariant), must multiply Gaussian derivative by σ
- Laplacian is the second Gaussian derivative, so it must be multiplied by σ^2

Source: S. Seitz, S. Savarese.

Effect of scale normalization

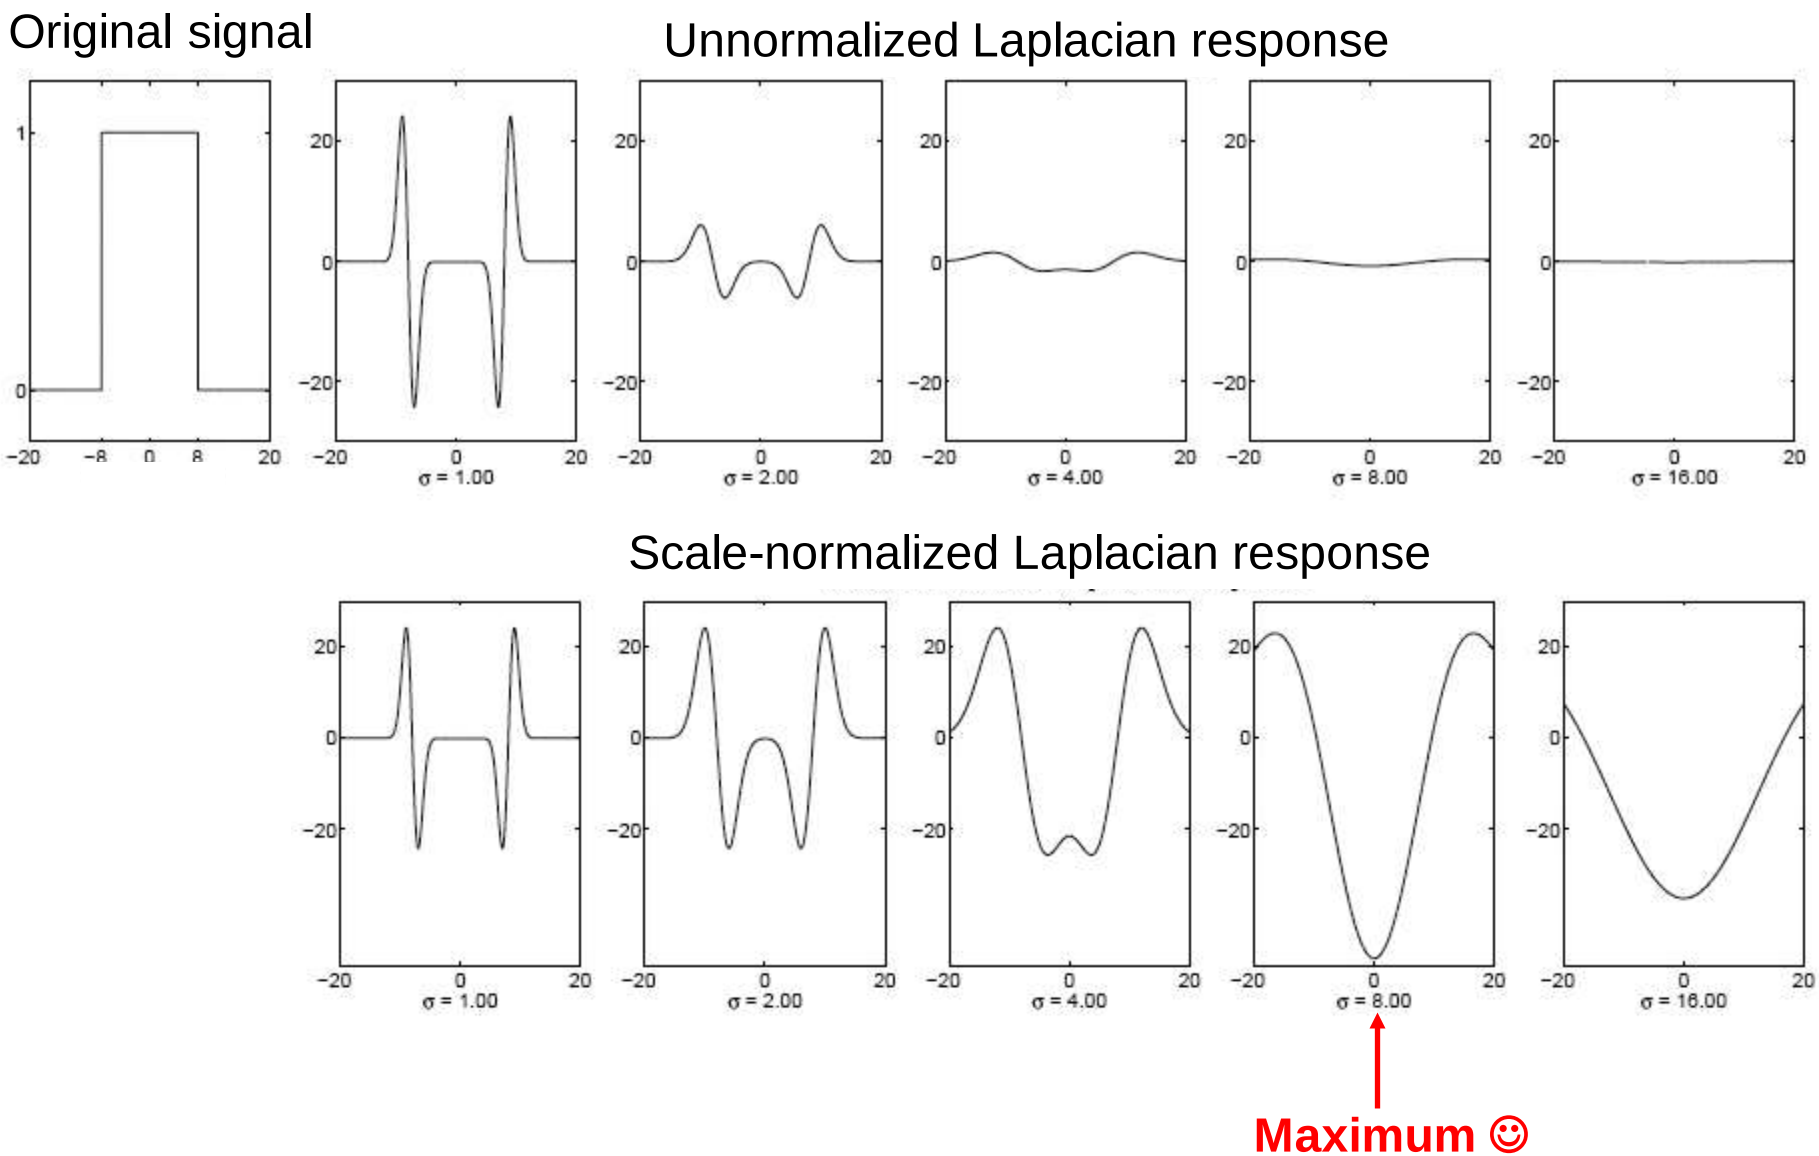

Source: S. Seitz, S. Savarese.

Example: Blob detection in 2D

- Laplacian of Gaussian: Circularly symmetric operator for blob detection in 2D

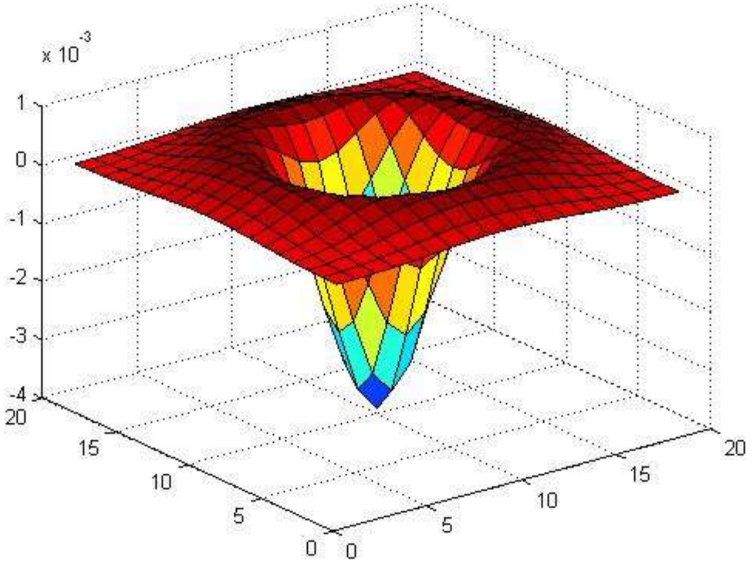

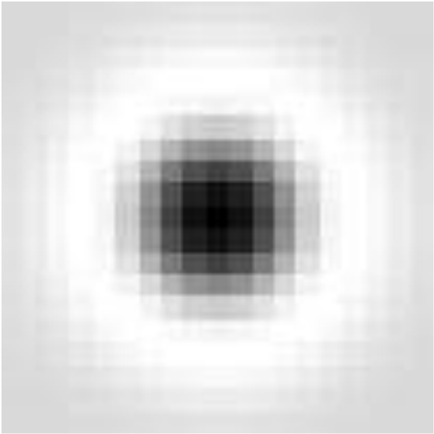

$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$$

Source: S. Seitz, S. Savarese.

Example: Blob detection in 2D

- Laplacian of Gaussian: Circularly symmetric operator for blob detection in 2D

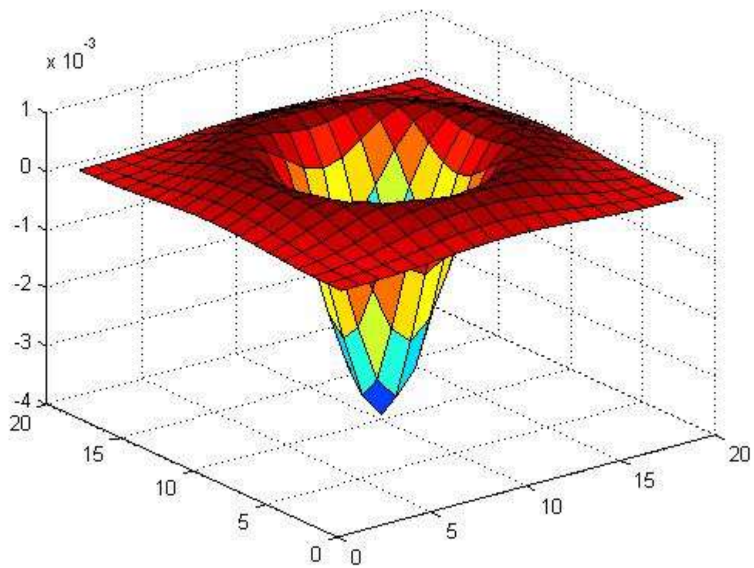

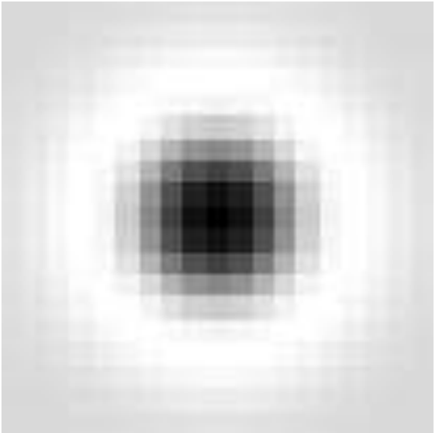

Scale-normalized: $\nabla^2_{\text{norm}} g = \sigma^2 \left(\frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2} \right)$

Source: S. Seitz, S. Savarese.

Scale selection

- For a binary circle of radius r , the Laplacian achieves a maximum at

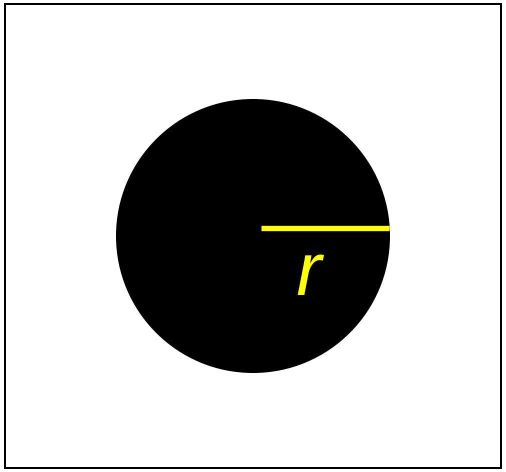

image

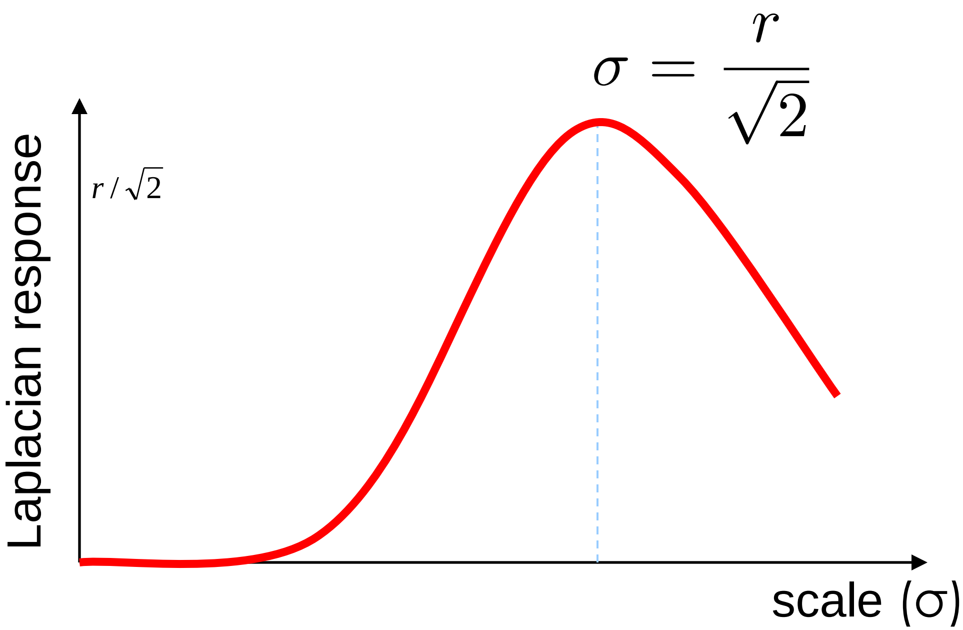

Source: S. Seitz, S. Savarese.

Characteristic scale

- We define the characteristic scale as the scale that produces peak of Laplacian response

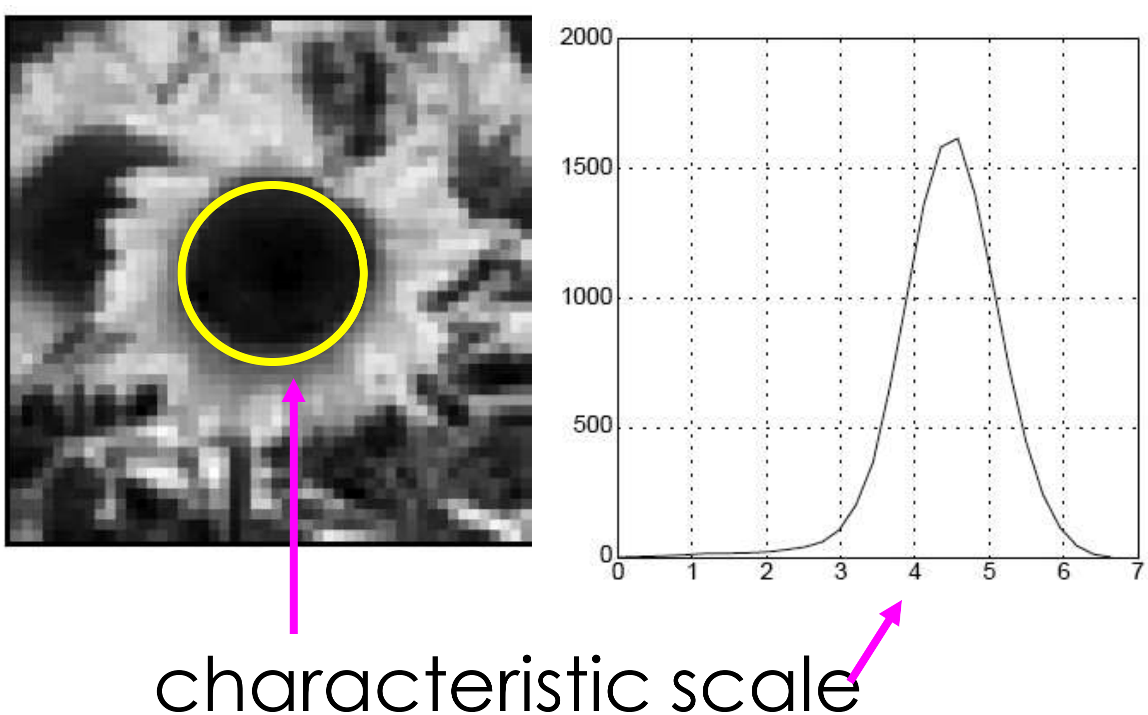

T. Lindeberg (1998). "Feature detection with automatic scale selection." *International Journal of Computer Vision* **30** (2): pp 77--116.

Source: S. Seitz, S. Savarese.

Scale-space blob detector

1. Convolve image with scale-normalized Laplacian at several scales
2. Find maxima of squared Laplacian response in scale-space
3. This indicates if a blob has been detected
4. And what is its intrinsic scale

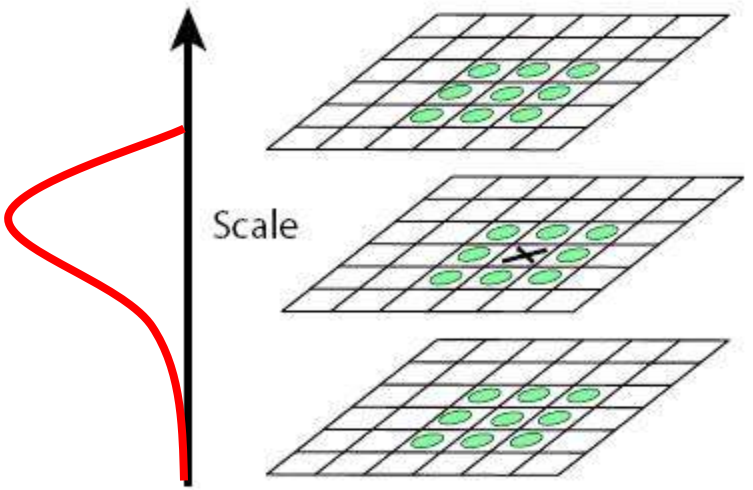

Source: S. Seitz, S. Savarese.

Scale-space blob detector: example

Source: S. Seitz, S. Savarese.

Scale-space blob detector: example

sigma = 11.9912

Source: S. Seitz, S. Savarese.

Scale-space blob detector: example

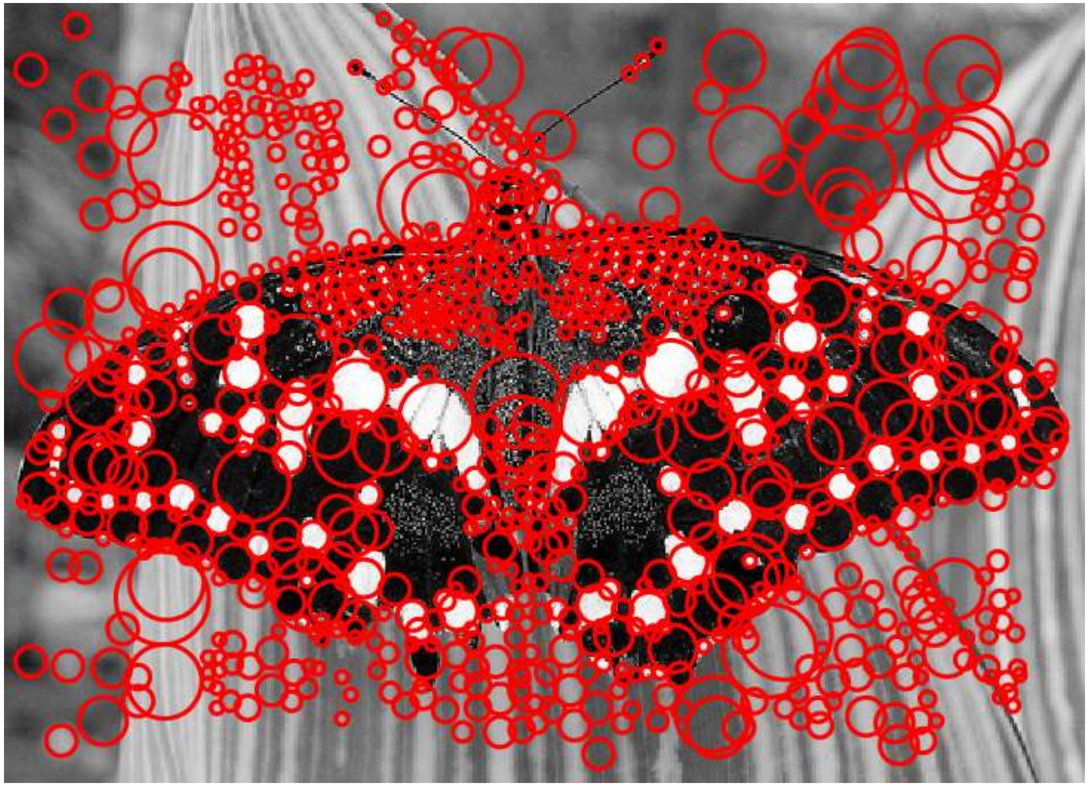

Source: S. Seitz, S. Savarese.

Automatic scale selection

Lindeberg et al., 1996

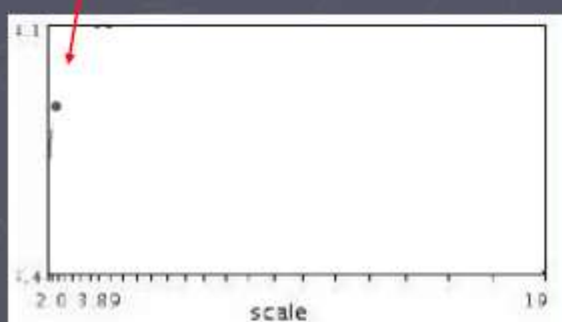

$f(I_{i_1 \ldots i_m}(x,\sigma))$

Slide from Tinne Tuytelaars

Automatic scale selection

$$f(I_{i_1 \ldots i_m}(x, \sigma))$$

Automatic scale selection

$$f(I_{i_1 \ldots i_m}(x,\sigma))$$

Automatic scale selection

$f(I_{i_1 \dots i_m}(x, \sigma))$

Automatic scale selection

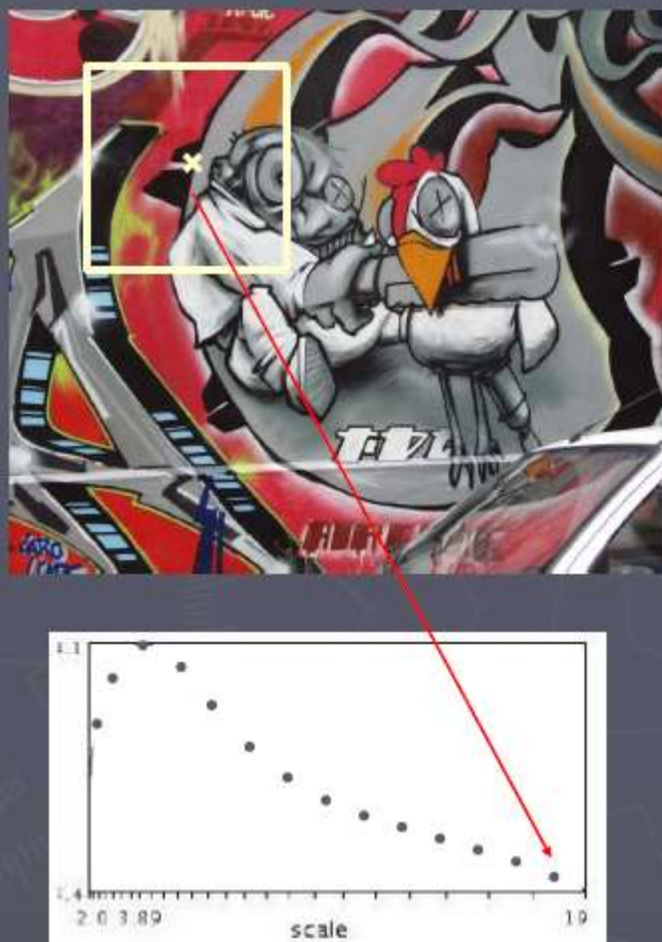

$f(I_{i_1 \ldots i_m}(x,\sigma))$

Automatic scale selection

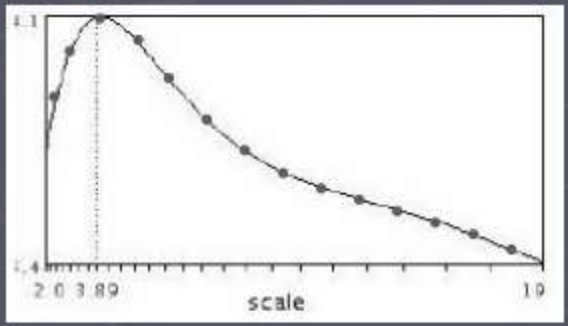

$f(I_{i_1 \ldots i_m}(x, \sigma))$

Automatic scale selection

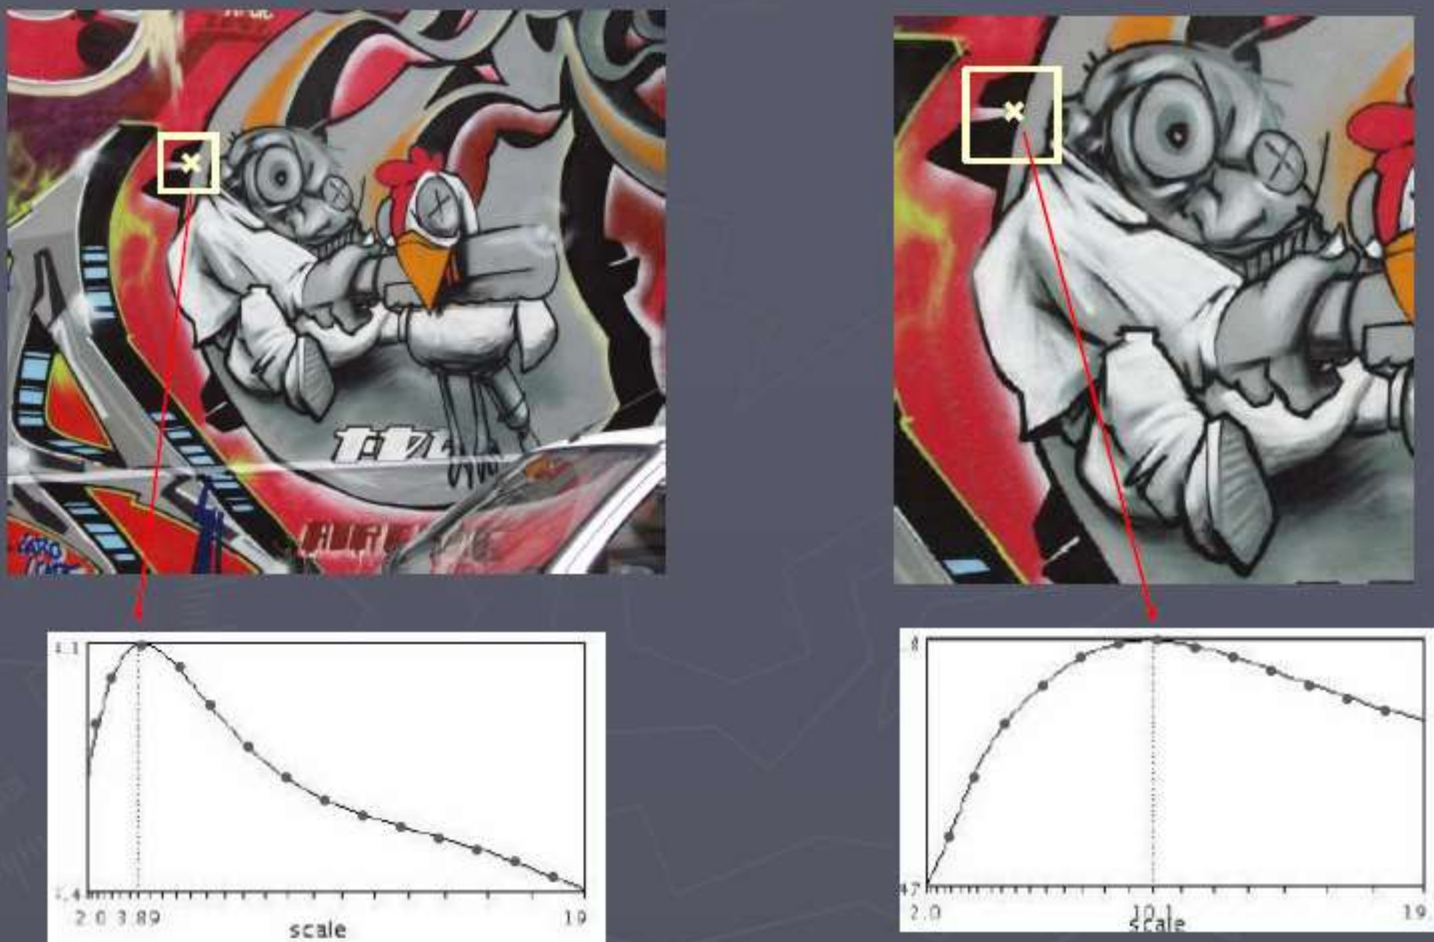

$f(I_{i_1 \ldots i_m}(x,\sigma))$ $f(I_{i_1 \ldots i_m}(x',\sigma'))$

Automatic scale selection

Normalize: rescale to fixed size

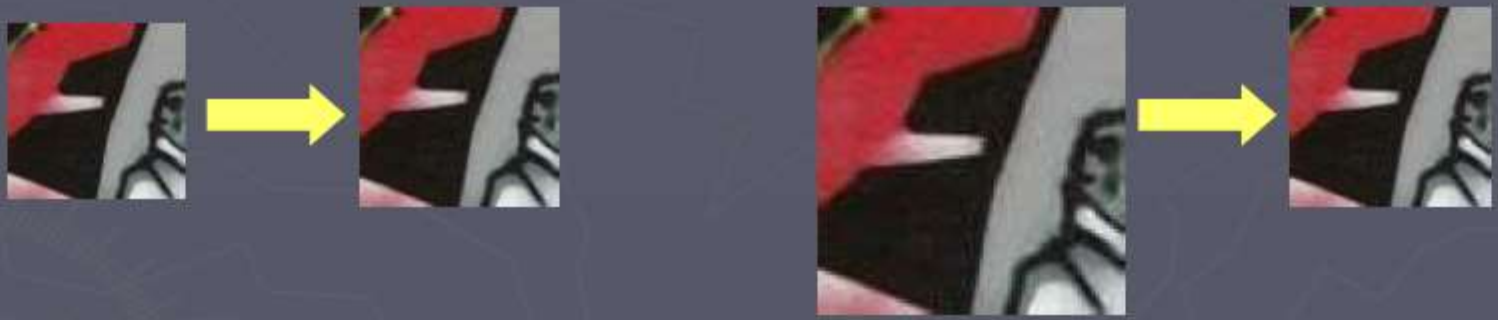

Difference of Gaussians Approximations to Laplacian

David G. Lowe. **"Distinctive image features from scale-invariant keypoints."** *IJCV* 60 (2), 04

- Approximating the Laplacian with a difference of Gaussians:

$$L = \sigma^2 \left(G_{xx}(x, y, \sigma) + G_{yy}(x, y, \sigma) \right)$$

Laplacian

$$DoG = G(x, y, k\sigma) - G(x, y, \sigma)$$

Difference of Gaussians

or

Difference of gaussian blurred images at scales k σ and σ

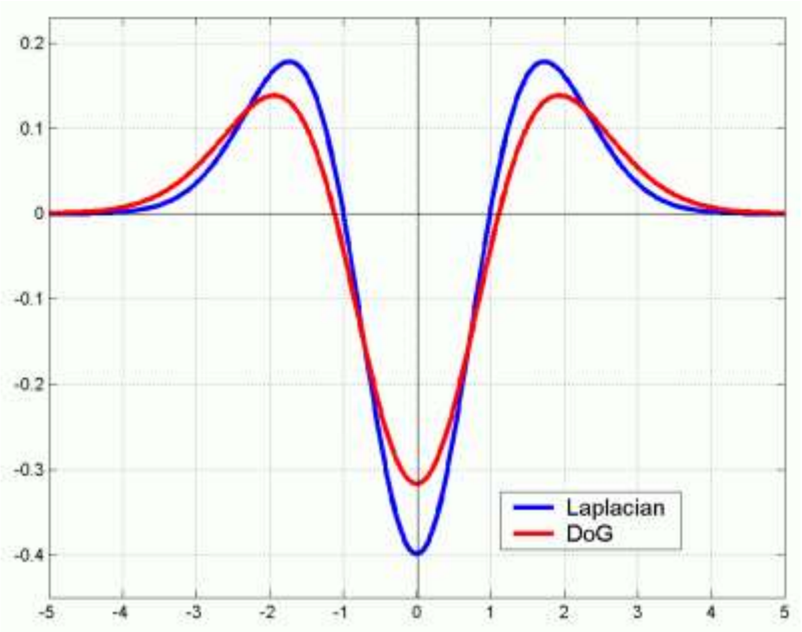

$$G(x, y, k\sigma) - G(x, y, \sigma) \approx (k-1)\sigma^2 L$$

Source: S. Seitz, S. Savarese.

Difference of Gaussians (DoG)

k

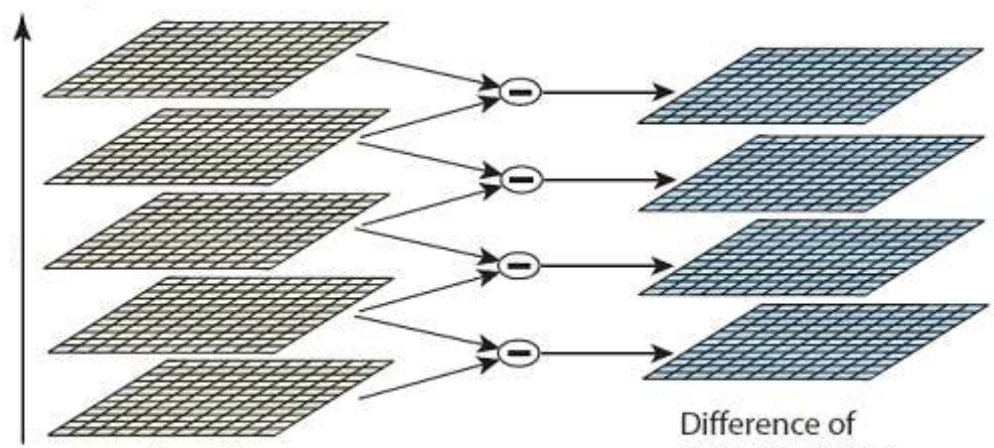

Output: location, scale, orientation

Source: S. Seitz, S. Savarese.

End

Invariance

Detector	Illumination	Rotation	Scale	View point
Harris corner	Yes	Yes	No	No
Lowe '99 (DoG)	Yes	Yes	Yes	No

Source: S. Savarese.

Harris-Laplace

[Mikolajczyk & Schmid '01]

- Collect locations (x,y) of detected Harris features for $\sigma = \sigma_1 \ldots \sigma_2$ (the sigma is here comes from $g_{x,}$ g_y)
- For each detected location (x,y) and for each σ, reject detection if Laplacian(x,y, σ) is not a local maximum

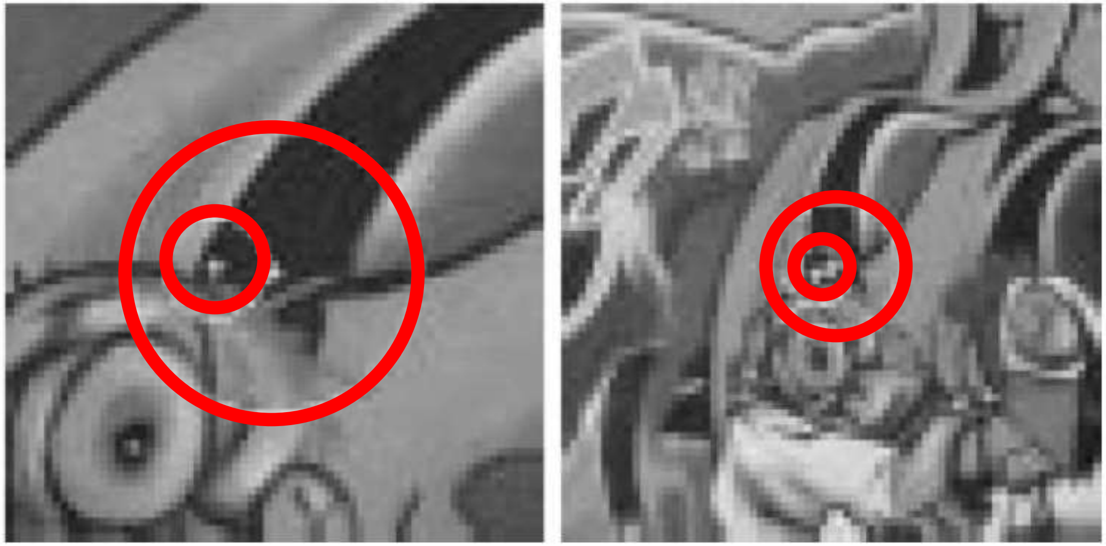

Output: location, scale

Source: S. Savarese.

Invariance

Detector	Illumination	Rotation	Scale	View point
Harris corner	Yes	Yes	No	No
Lowe '99 (DoG)	Yes	Yes	Yes	No
Mikolajczyk & Schmid '01	Yes	Yes	Yes	No

Source: S. Savarese.

Repeatability

Illumination invariance

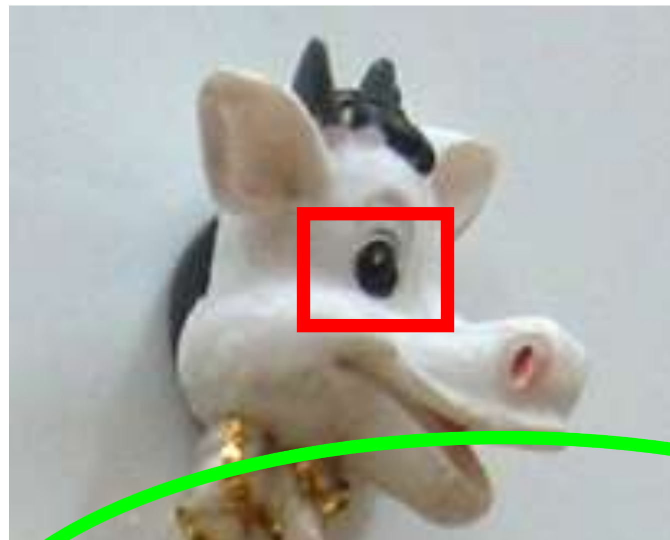

Scale invariance

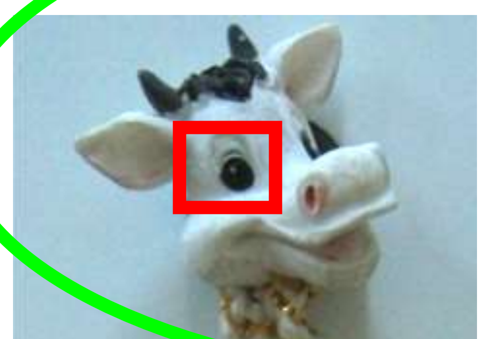

Pose invariance

- Rotation
- Affine

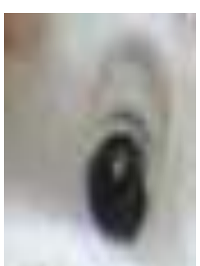

Source: S. Savarese.

Affine invariance

K. Mikolajczyk and C. Schmid, Scale and Affine invariant interest point detectors, IJCV 60(1):63-86, 2004.

Similarly to characteristic scale selection, detect the characteristic shape of the local feature

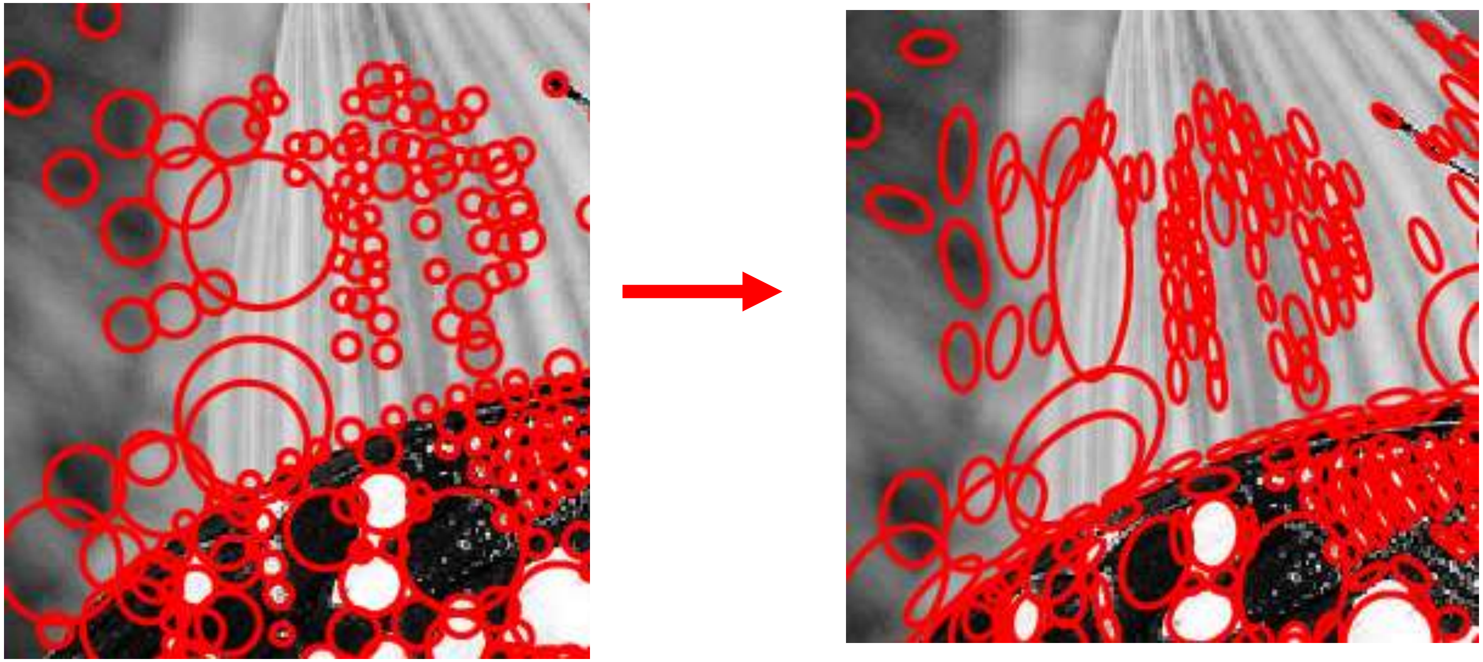

Source: S. Savarese.

Affine invariance

$$M = \sum_{x,y} w(x,y) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

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

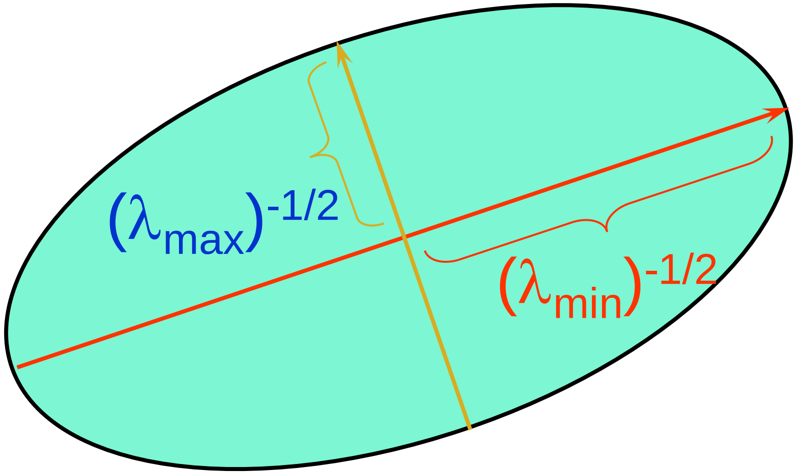

The second moment ellipse can be viewed as the "characteristic shape" of a region

Source: S. Savarese.

Affine adaptation

1. Detect initial region with Harris Laplace
2. Estimate affine shape with M
3. Normalize the affine region to a circular one
4. Re-detect the new location and scale in the normalized image
5. Go to step 2 if the eigenvalues of the M for the new point are not equal [detector not yet adapted to the characteristic shape]

Source: S. Savarese.

Without affine invariance

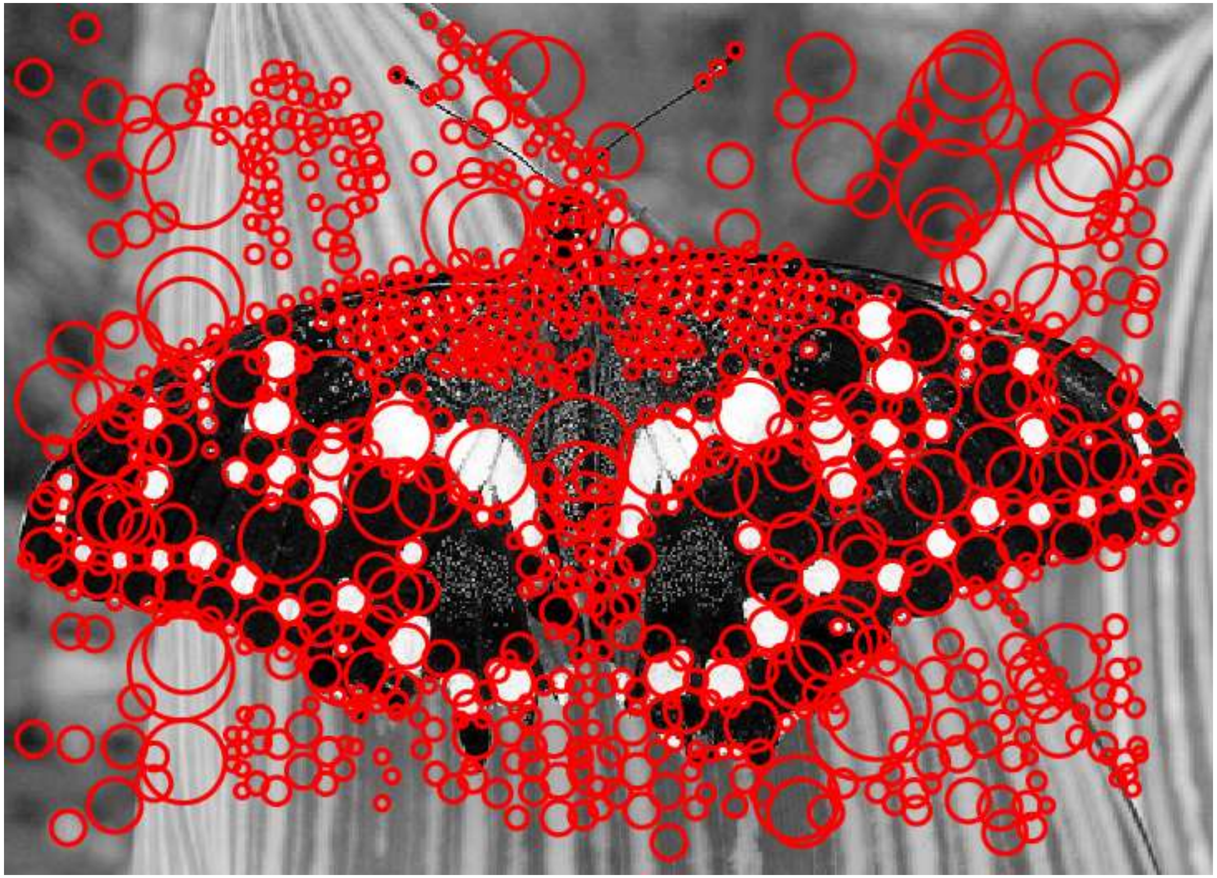

Scale-invariant regions (blobs)

Source: S. Savarese.

With affine invariance

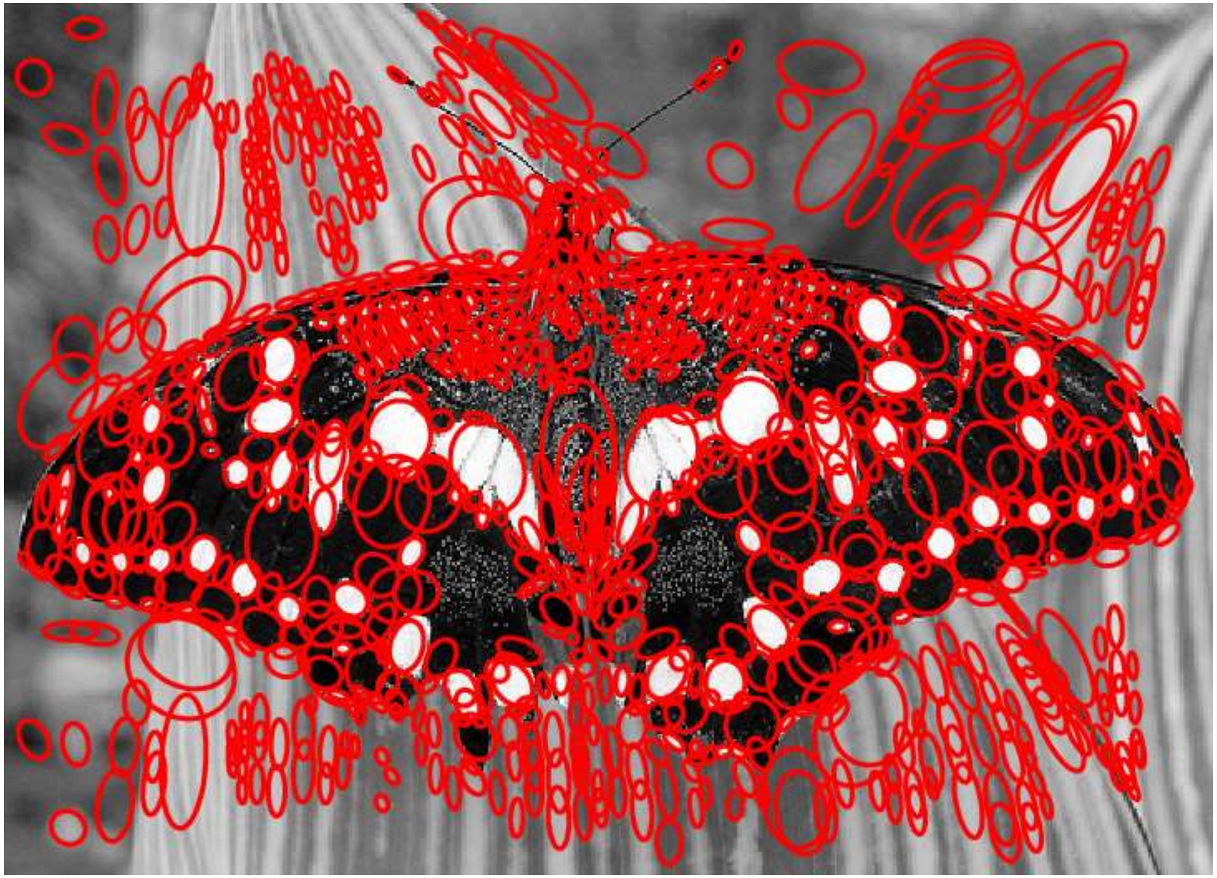

Affine-adapted blobs

Invariance

Detector	Illumination	Rotation	Scale	View point
Harris corner	Yes	Yes	No	No
Lowe '99 (DoG)	Yes	Yes	Yes	No
Mikolajczyk & Schmid '01	Yes	Yes	Yes	No
Mikolajczyk & Schmid '02	Yes	Yes	Yes	Yes

Source: S. Savarese.

Detector	Illumination	Rotation	Scale	View point
Harris corner	Yes	Yes	No	No
Lowe '99 (DoG)	Yes	Yes	Yes	Yes
Mikolajczyk & Schmid '01, '02	Yes	Yes	Yes	Yes
Tuytelaars, '00	Yes	Yes	No (Yes '04)	Yes
Kadir & Brady, 01	Yes	Yes	Yes	no
Matas, '02	Yes	Yes	Yes	no

Source: S. Savarese.