UiO : Department of Technology Systems
University of Oslo

Blob features

Trym Vegard Haavardsholm

2023

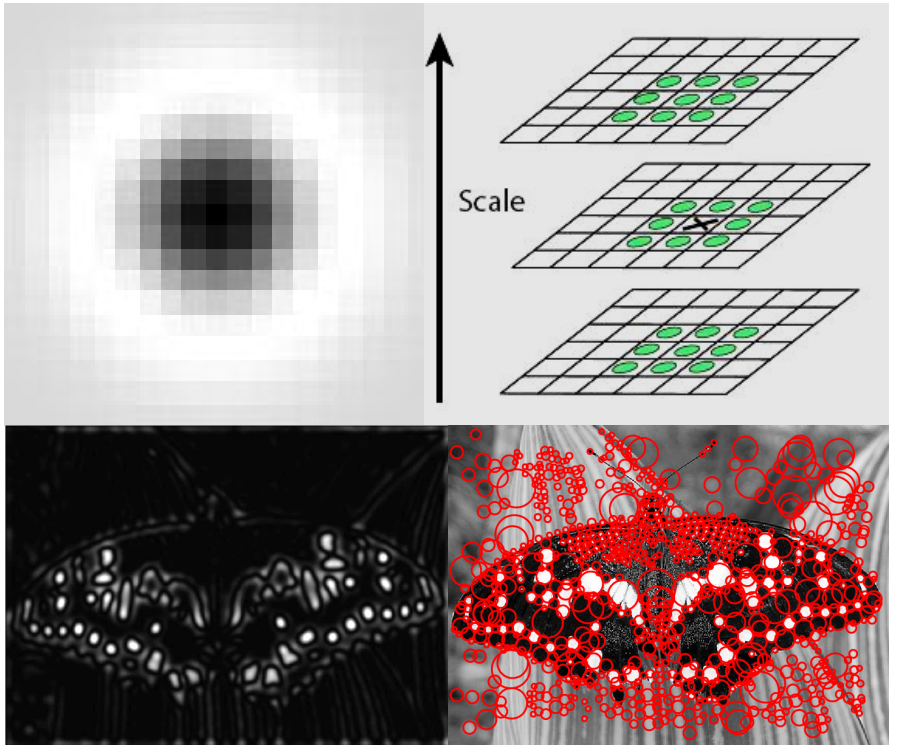

With illustrations from Svetlana Lazebnik, Grauman&Leibe, S. Seitz, James Hays and Noah Snavely

Automatic scale selection

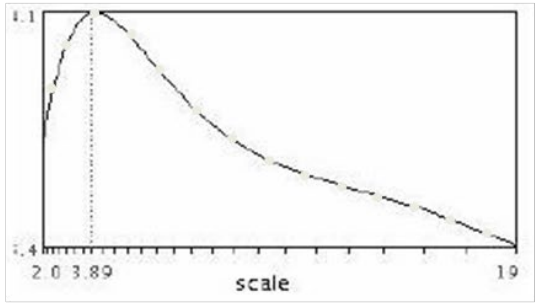

$f(I_{i_1 \ldots i_m}(x,\sigma))$

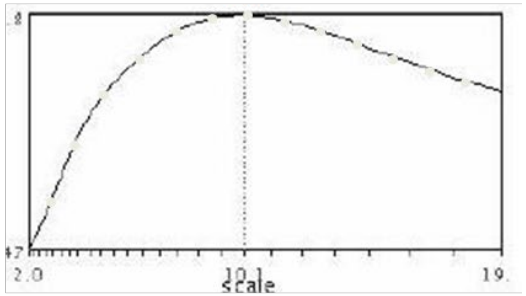

$f(I_{i_1 \ldots i_m}(x',\sigma'))$

An alternative to corner feature score functions

The *Laplacian of Gaussian (LoG)*

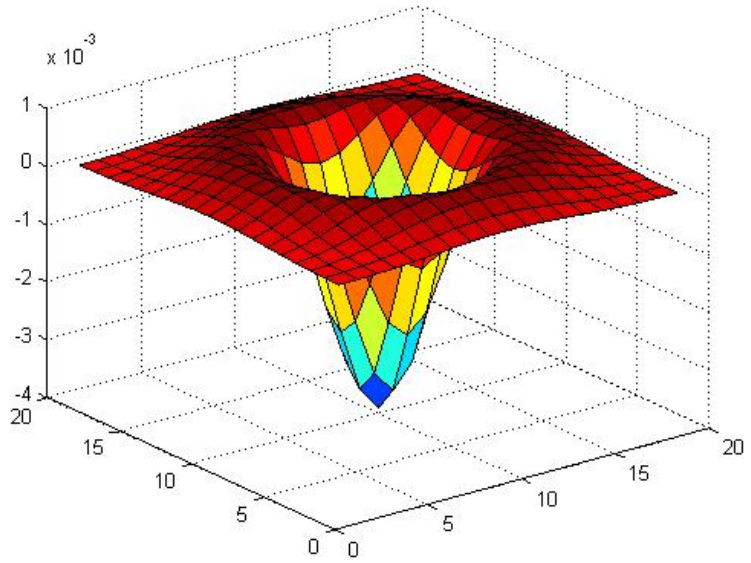

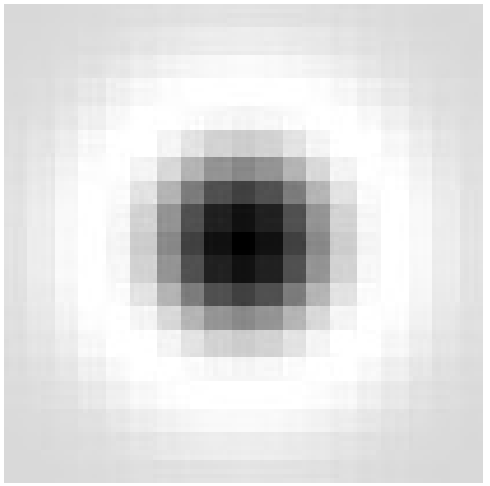

$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$$

Edges and blobs

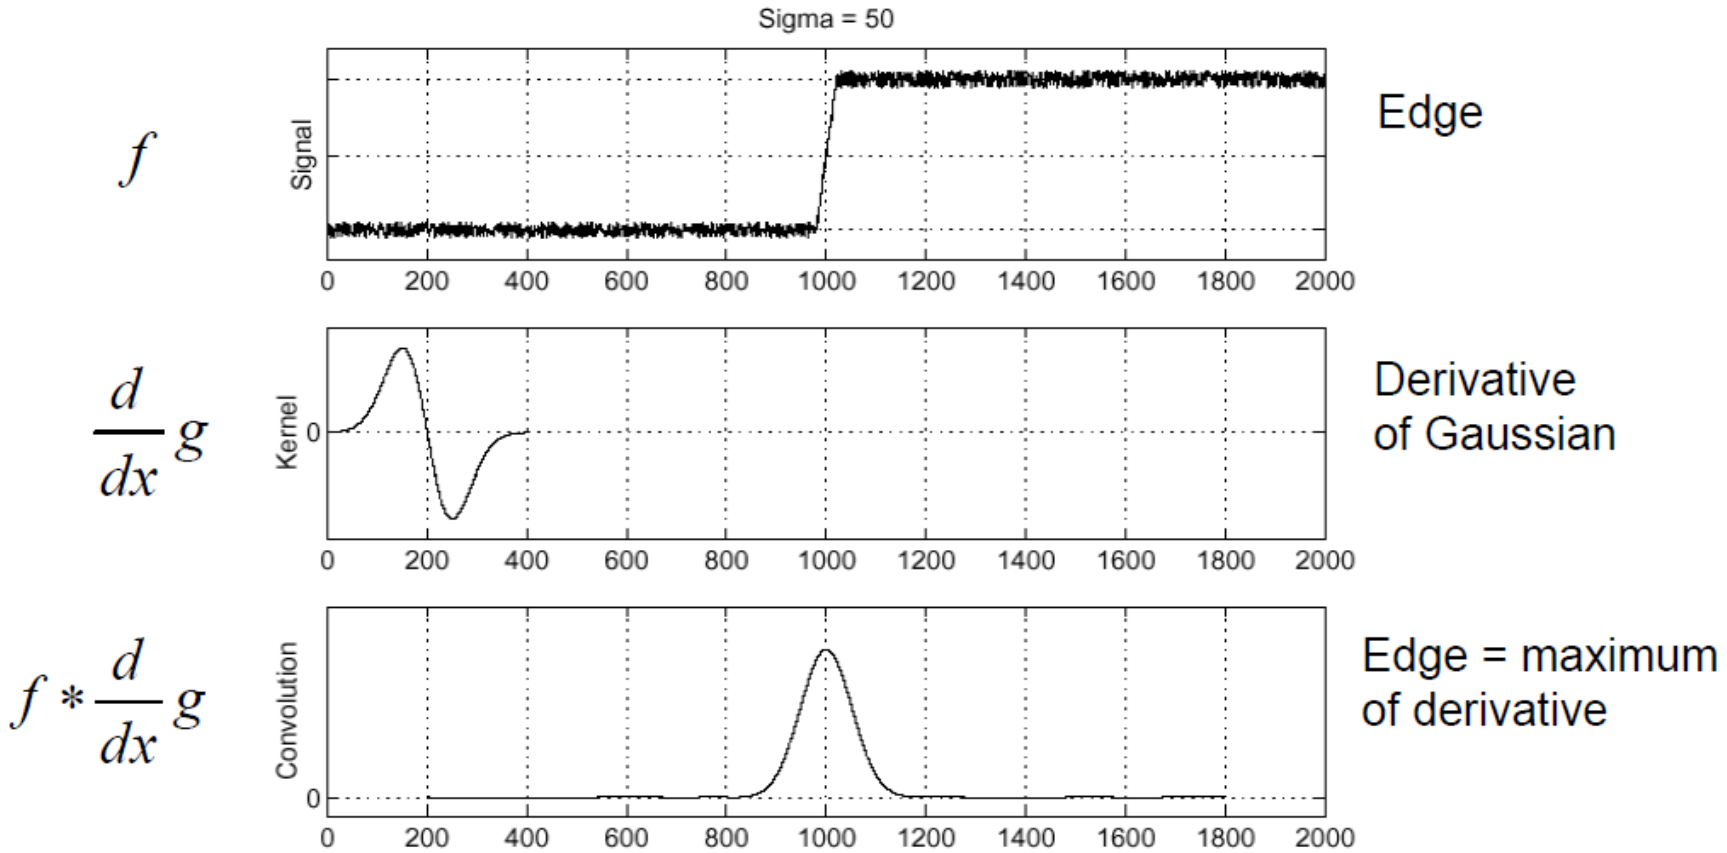

Edges and blobs

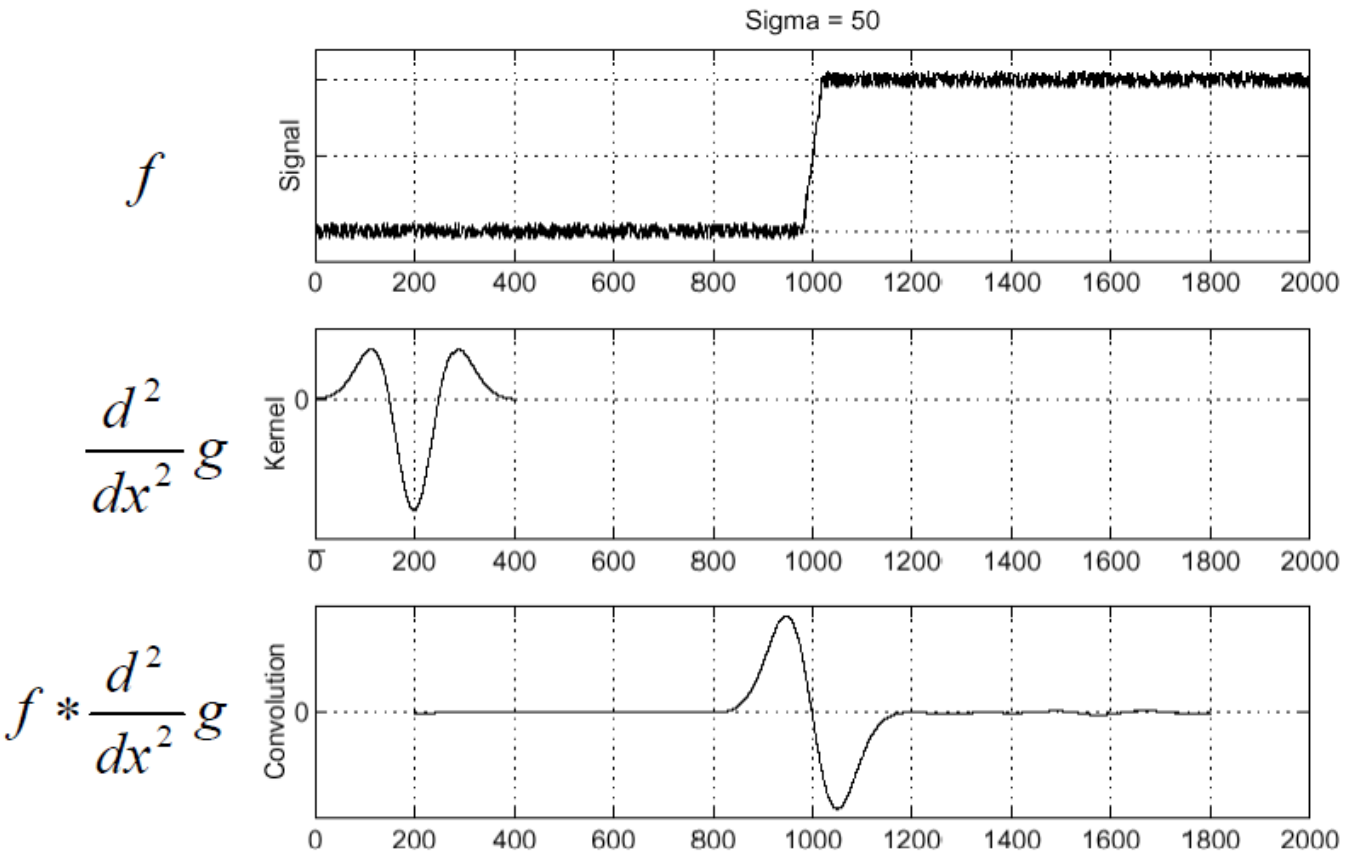

f

Edge

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

Second derivative of Gaussian (Laplacian)

$f * \frac{d^2}{dx^2} g$

Edge = zero crossing of second derivative

Edges and blobs

Edges → LoG returns ripples

Blobs → LoG returns superposition of two ripples

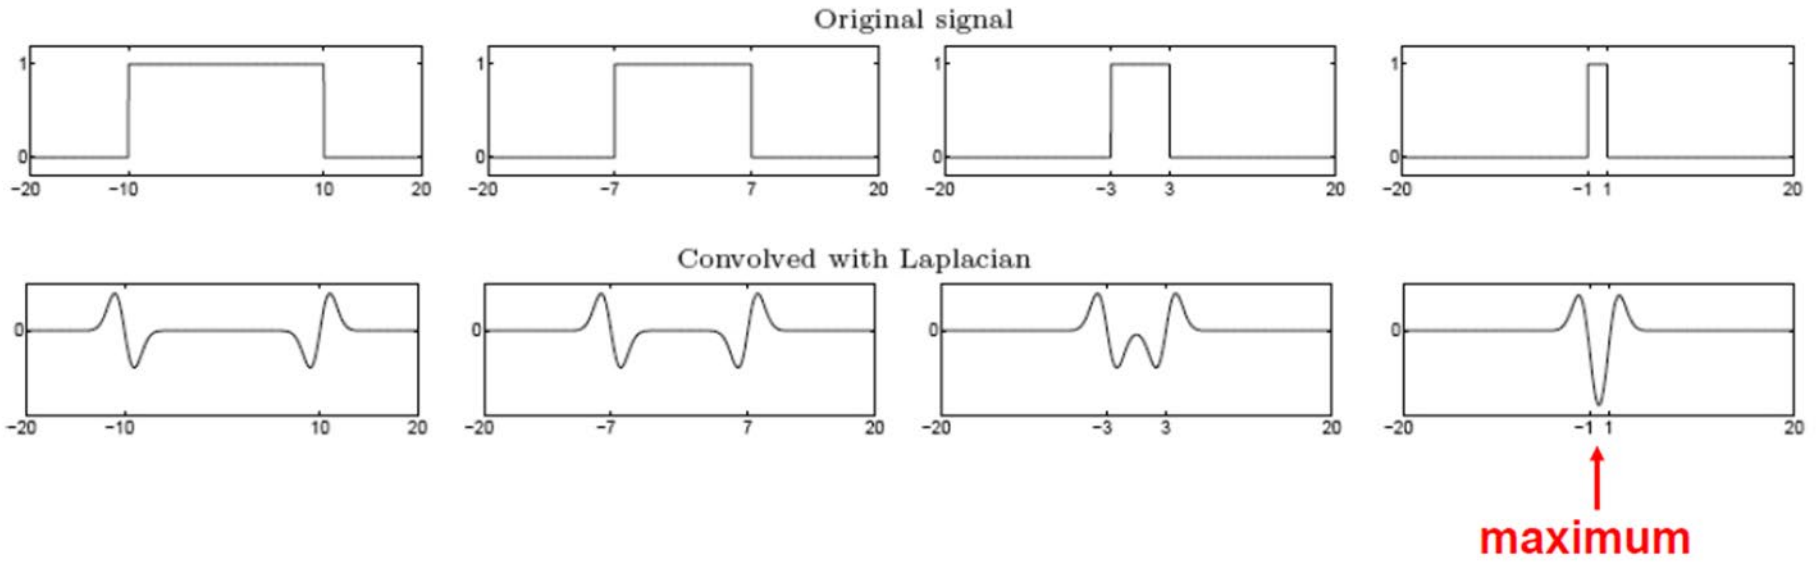

The magnitude of the Laplacian response is maximum at the centre of the blob provided the scale of the Laplacian matches the scale of the blob

Selecting the characteristic scale

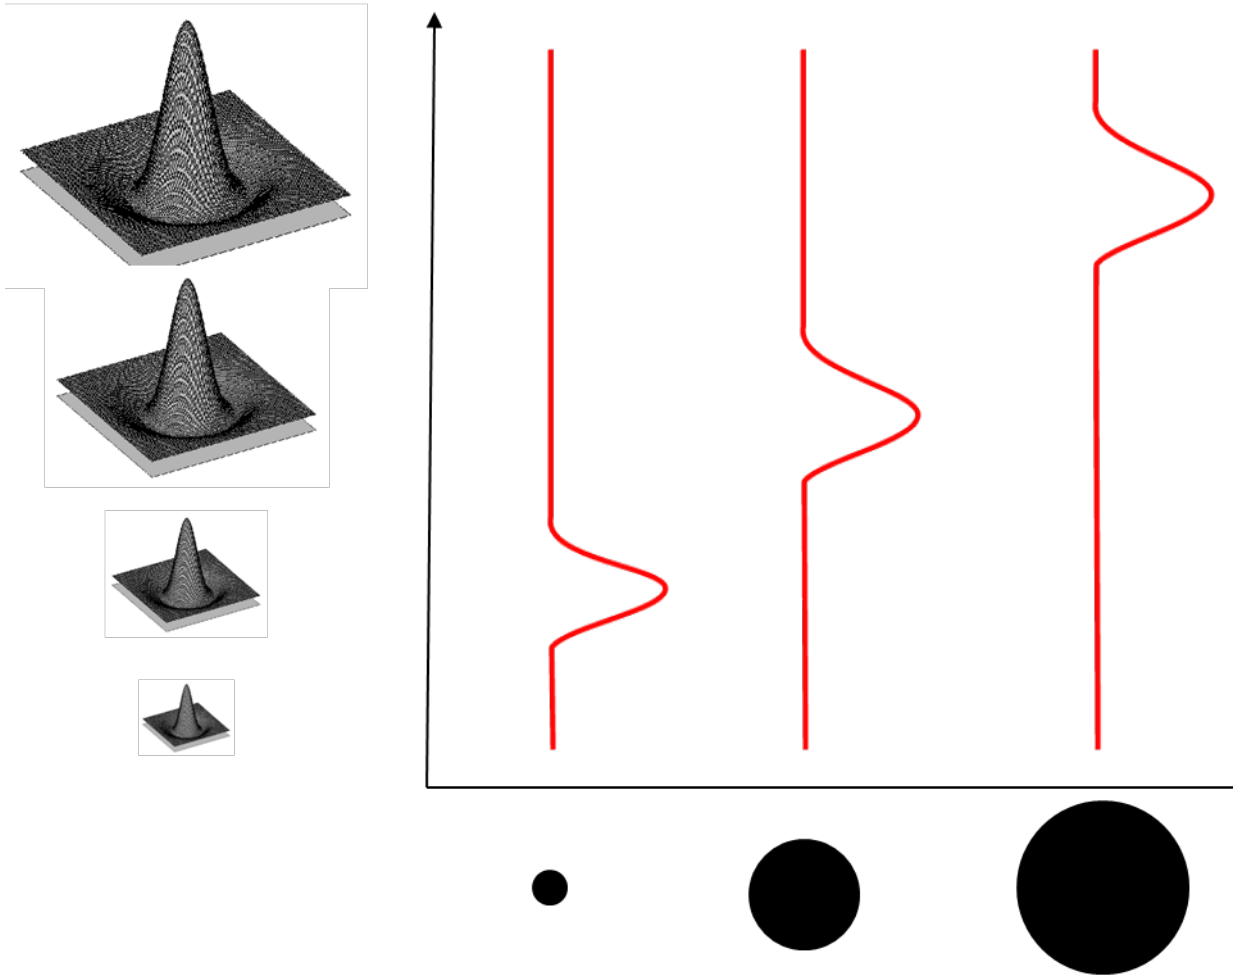

K. Grauman, B. Leibe

Scale-normalised Laplacian of Gaussian

Normalise to make the response independent of scale

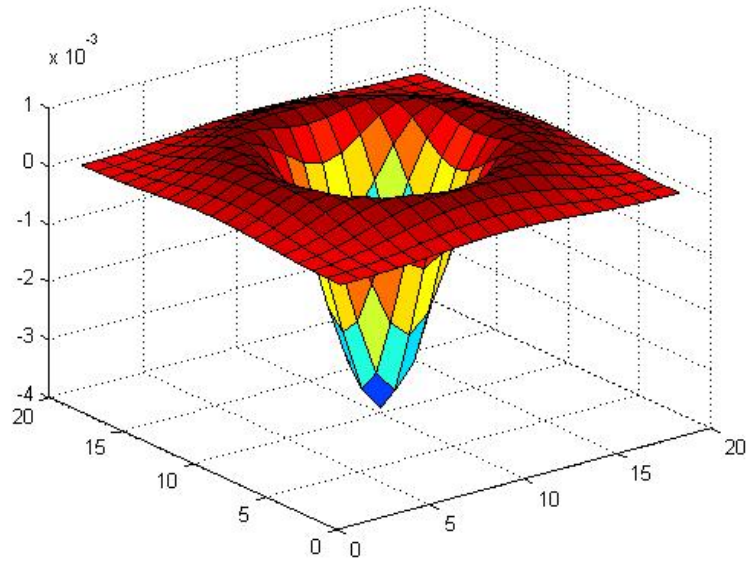

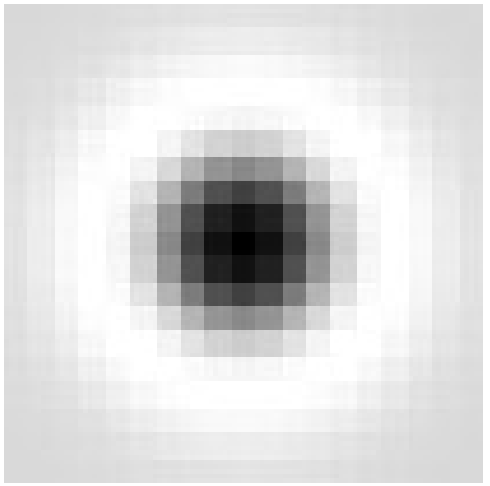

Scale-normalized: $\nabla^2_{\text{norm}} g = \sigma^2 \left(\frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2} \right)$

Selecting the characteristic scale

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

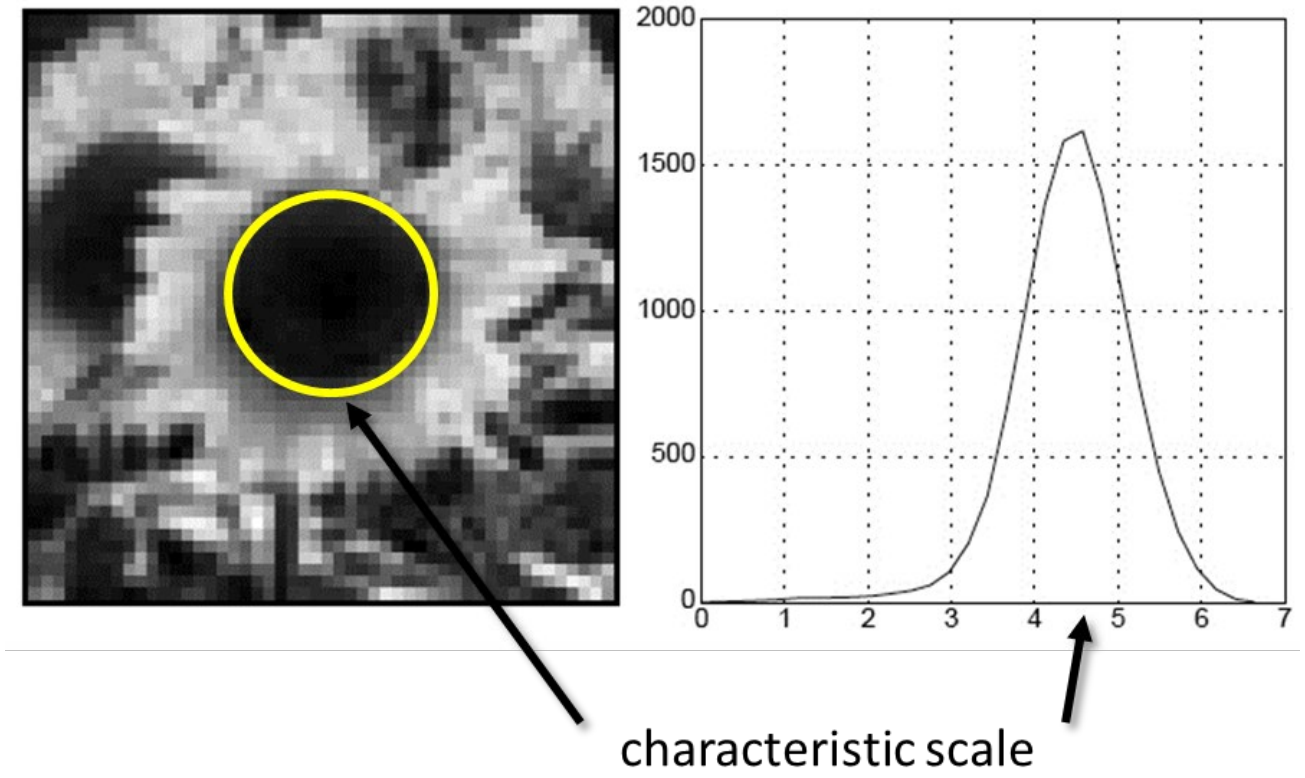

Scale selection

At what scale does the scale-normalised Laplacian
achieve a maximum response to a binary circle of radius *r*?

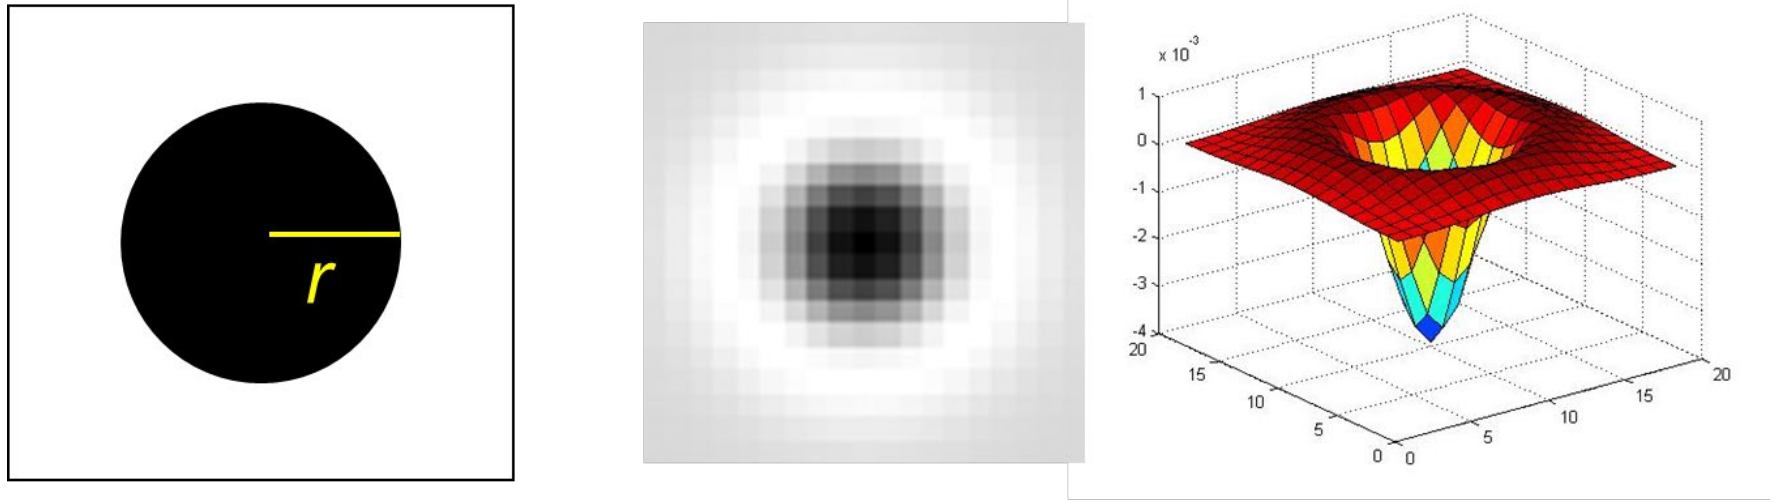

image Laplacian

Scale selection

We get the peak response when the zeros of the Laplacian are aligned with the circle

$$\sigma = r / \sqrt{2}.$$

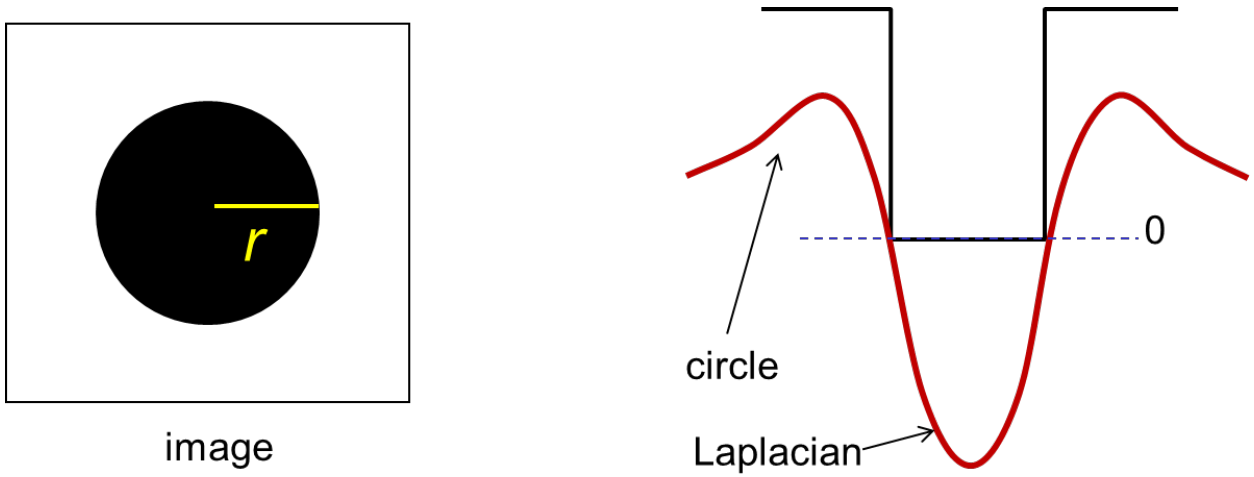

The LoG blob detector

Find *maxima* and *minima* of the scale-normalised LoG operator in space and scale

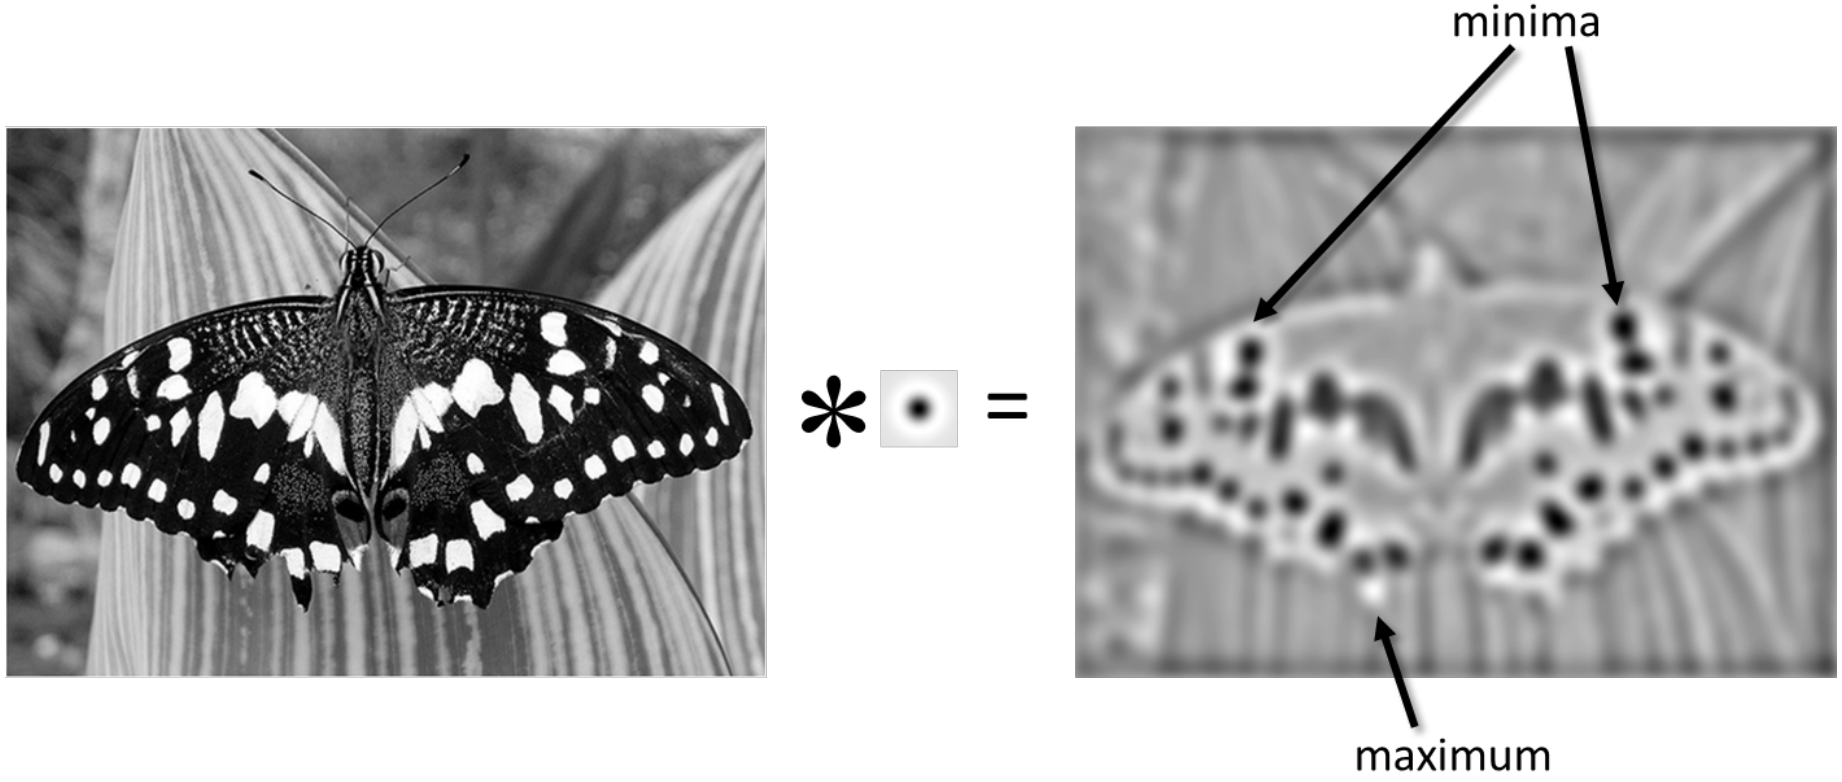

The LoG blob detector

1. Convolve the image with scale-normalised LoG at different scales
2. Find maxima of squared LoG response in scale-space

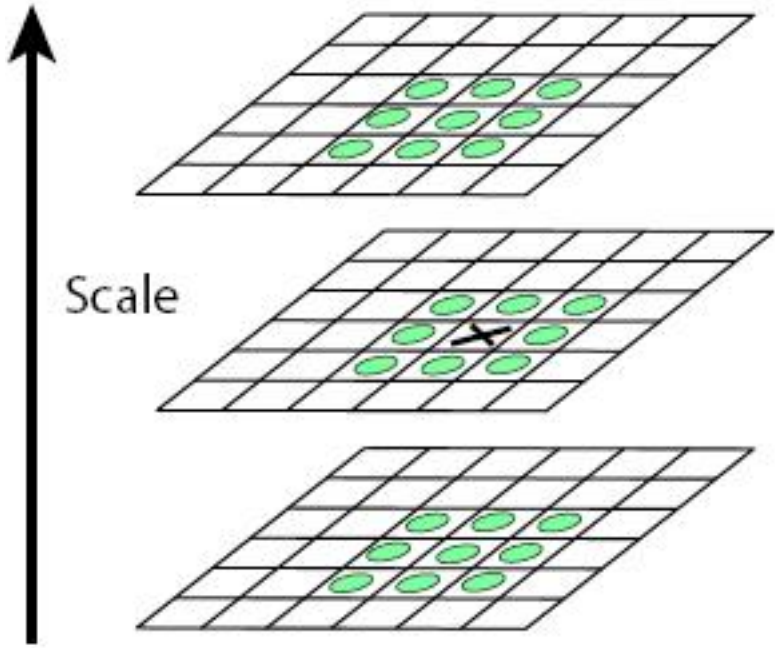

Blob detection example

Blob detection example

sigma = 2

Blob detection example

sigma = 2.5018

Blob detection example

sigma = 3.1296

Blob detection example

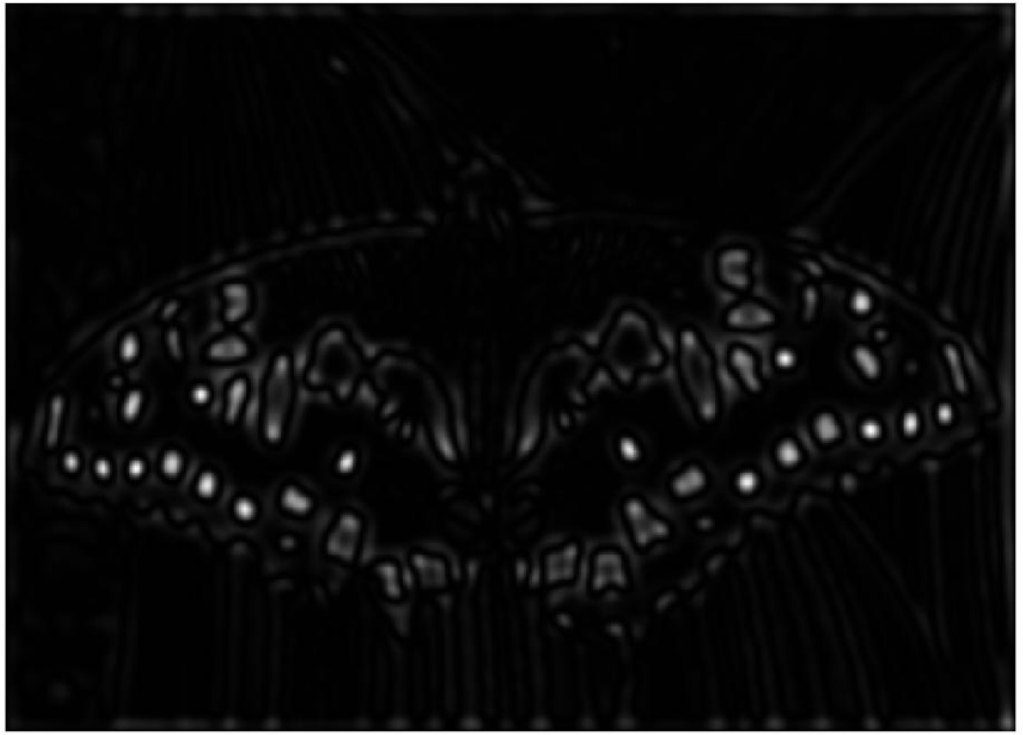

sigma = 3.9149

Blob detection example

sigma = 4.8972

Blob detection example

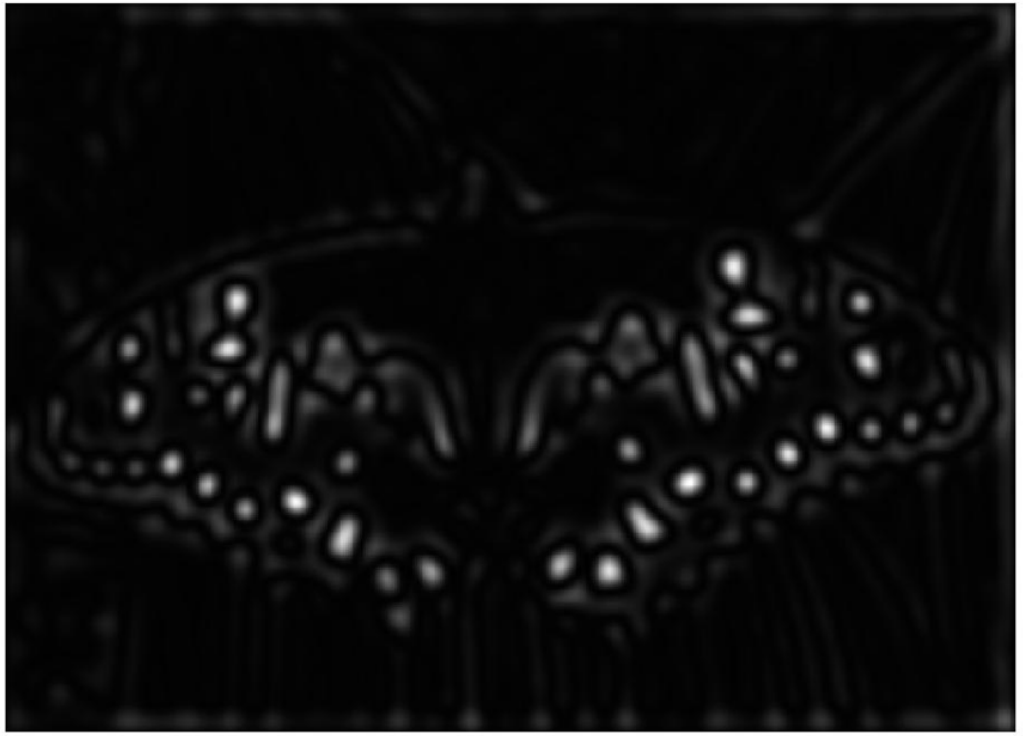

sigma = 6.126

Blob detection example

sigma = 7.6631

Blob detection example

sigma = 9.5859

Blob detection example

sigma = 9.5859

Blob detection example

sigma = 11.9912

Blob detection example

sigma = 15

Blob detection example

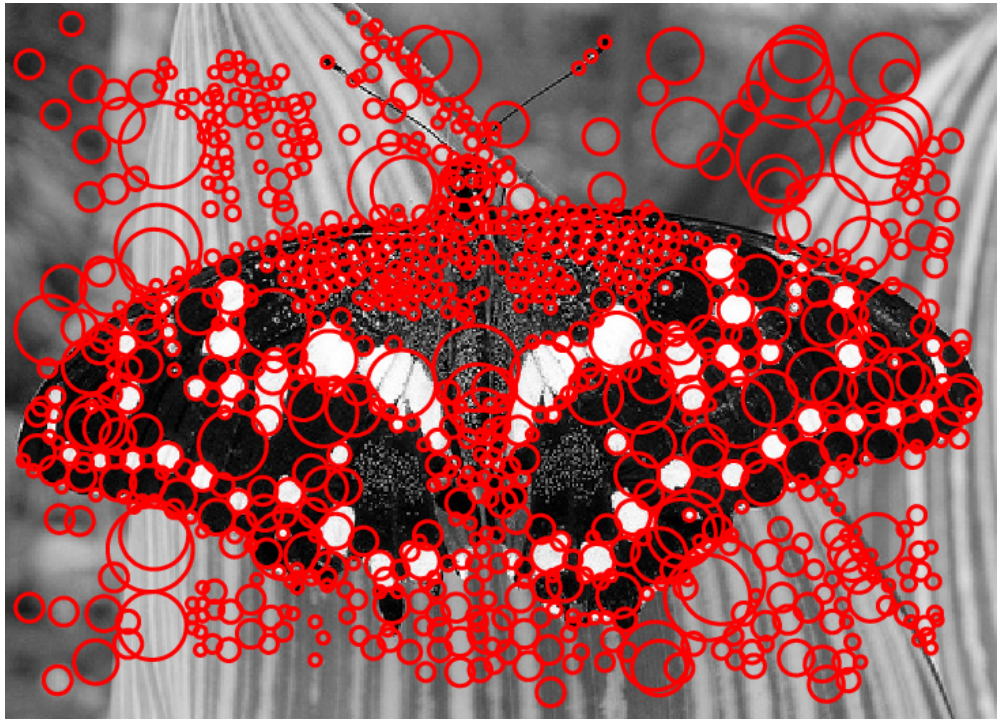

Efficient implementation

Approximate the normalised LoG
with a Difference of Gaussians (DoG):

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

(Laplacian of Gaussians)

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

(Difference of Gaussians)

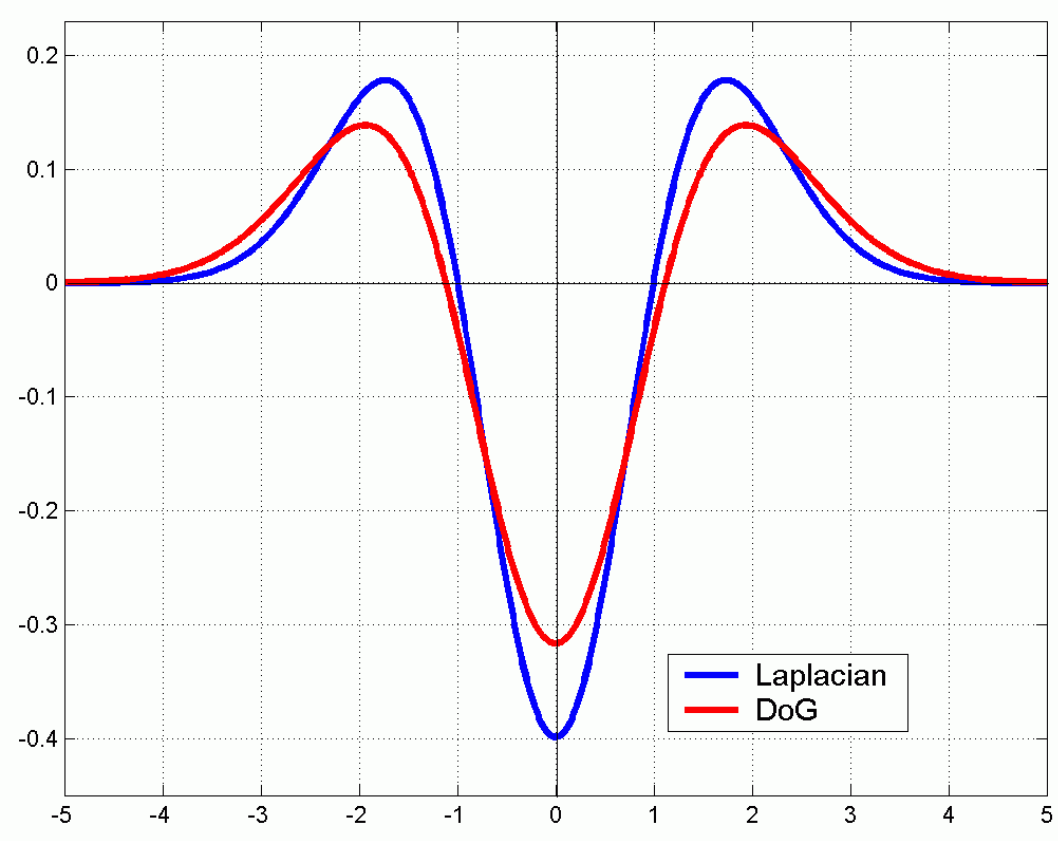

Efficient implementation

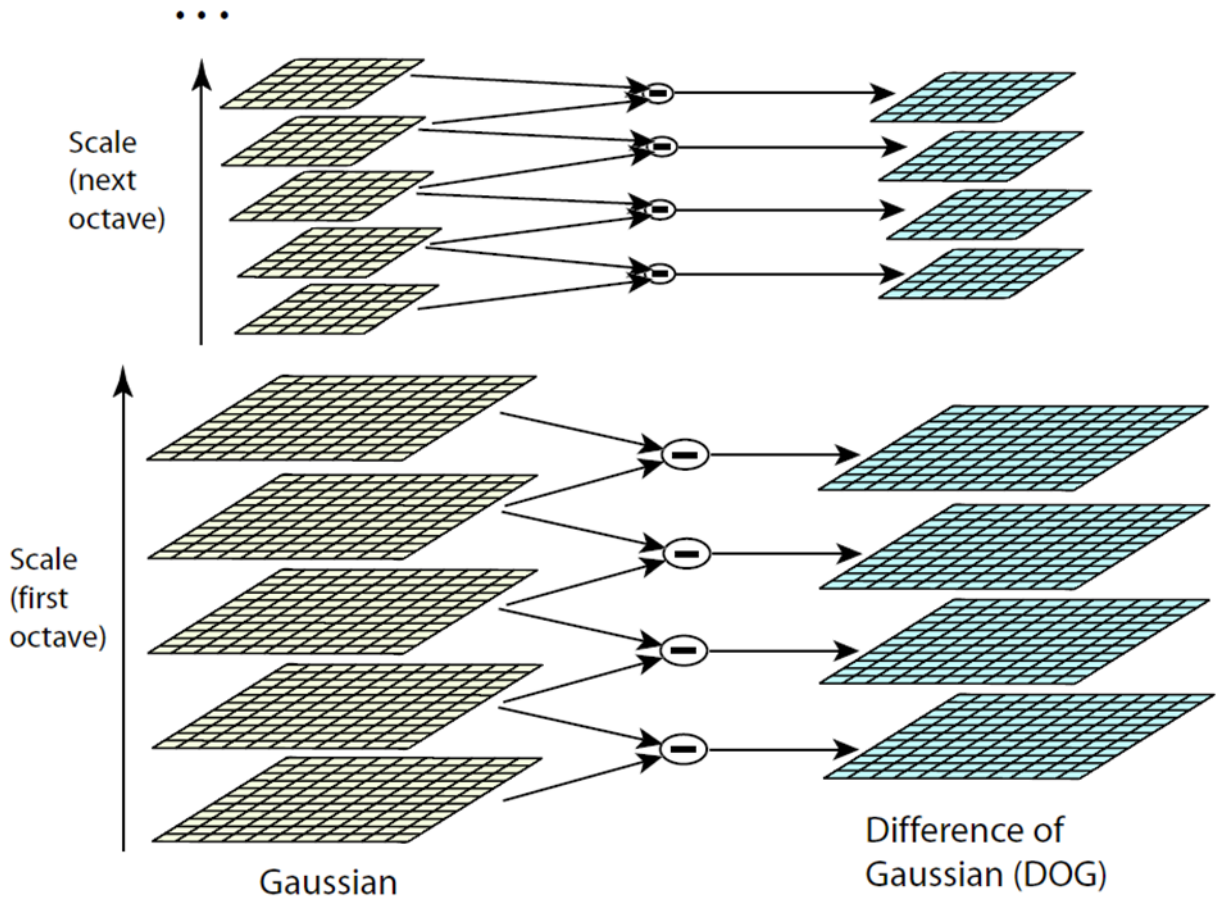

David G. Lowe. **"Distinctive image features from scale-invariant keypoints."**
IJCV 60 (2), pp. 91-110, 2004.

Efficient implementation

Detect local maxima and minima by comparing a pixel to its 26 neighbours in space and adjacent scales

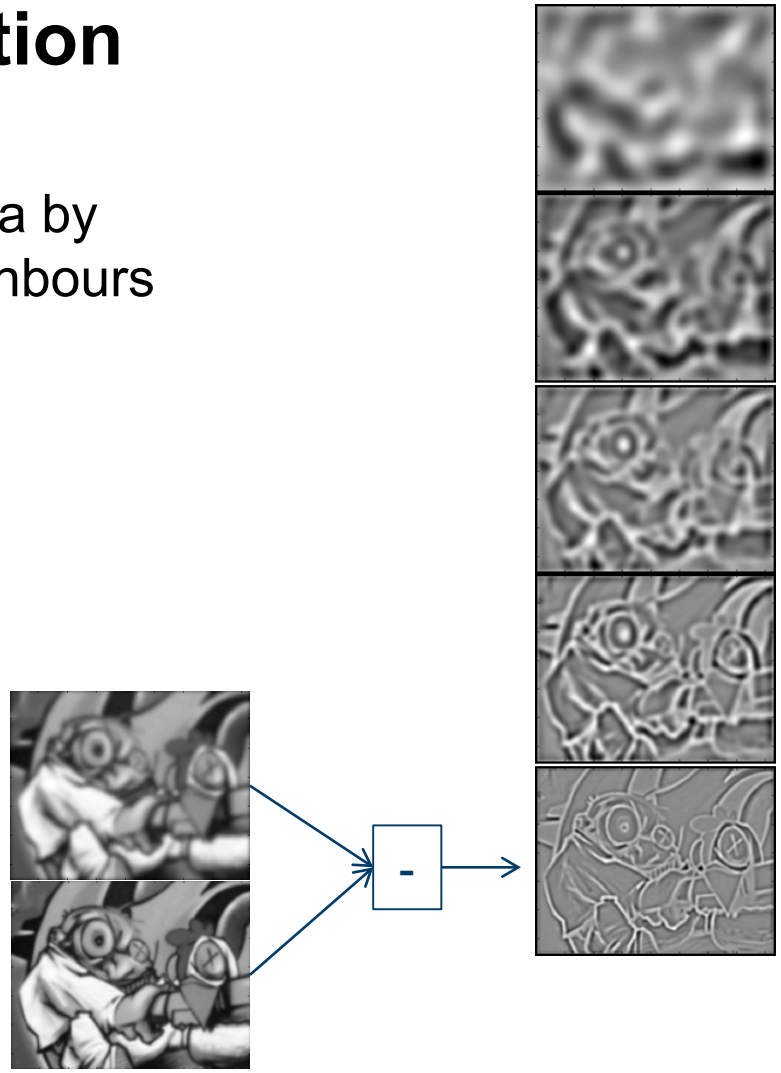

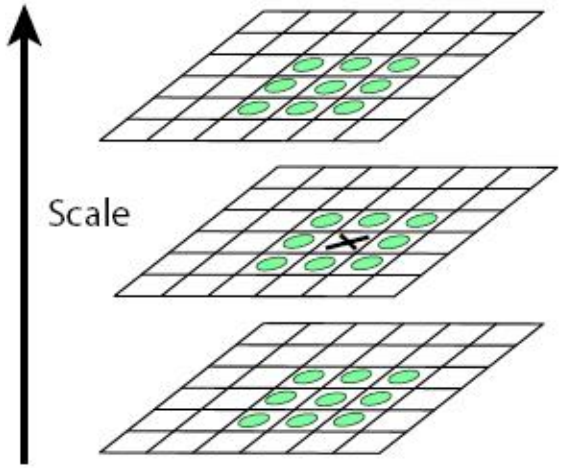

David G. Lowe. **"Distinctive image features from scale-invariant keypoints."**
IJCV 60 (2), pp. 91-110, 2004.

Summary – Keypoints

- Corner detectors
  - Distinct in space
  - Minimum eigenvalue, Harris
  - Properties
- Blob detectors
  - Distinct in space and scale
  - LoG, DoG, Lowe's
- Detected locations are often complementary
  - Combine methods!