Edges

Image gradient

The gradient of an image:

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

The gradient points in the direction of most rapid change in intensity

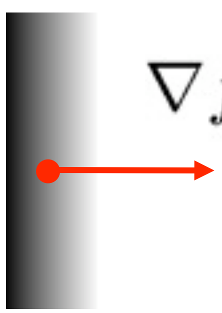

$$\nabla f = \left[\frac{\partial f}{\partial x}, 0\right]$$

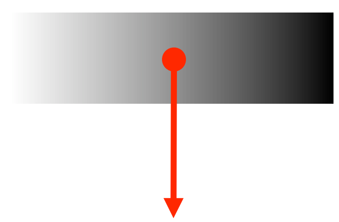

$$\nabla f = \left[0, \frac{\partial f}{\partial y}\right]$$

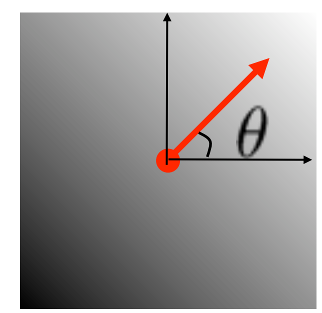

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

The gradient direction (orientation of edge normal) is given by:

$$\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)$$

The *edge strength* is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

Slide credit S. Seitz

Derivative of Gaussian filter

$$(I \otimes g) \otimes h = I \otimes (g \otimes h)$$

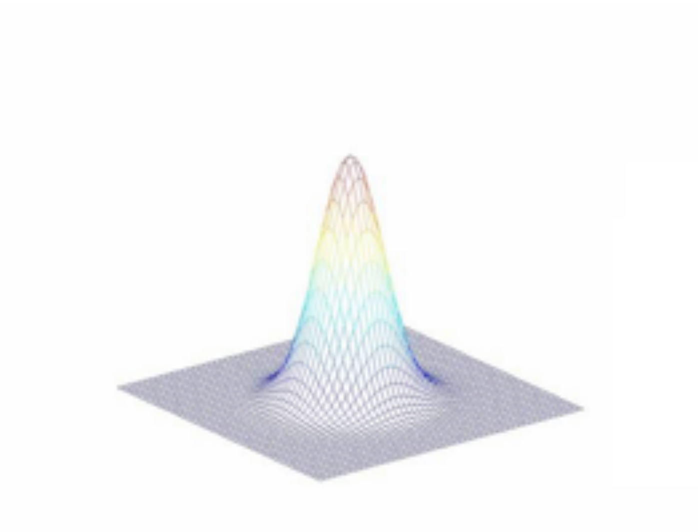

$$\frac{1}{159}\begin{bmatrix} 2 & 4 & 5 & 4 & 2 \\ 4 & 9 & 12 & 9 & 4 \\ 5 & 12 & 15 & 12 & 5 \\ 4 & 9 & 12 & 9 & 4 \\ 2 & 4 & 5 & 4 & 2 \end{bmatrix} \otimes \begin{bmatrix} 1 & -1 \end{bmatrix}$$

We can smooth and take the derivative with one filter, that is, with one convolution pass.

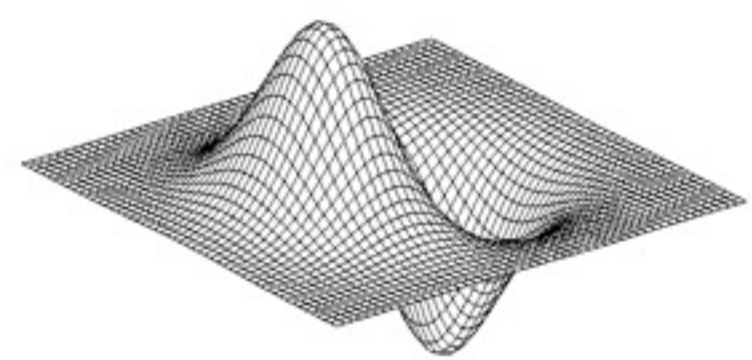

Effect of σ on Gaussian smoothing

Recall: parameter σ is the "scale" / "width" / "spread" of the Gaussian kernel, and controls the amount of smoothing.

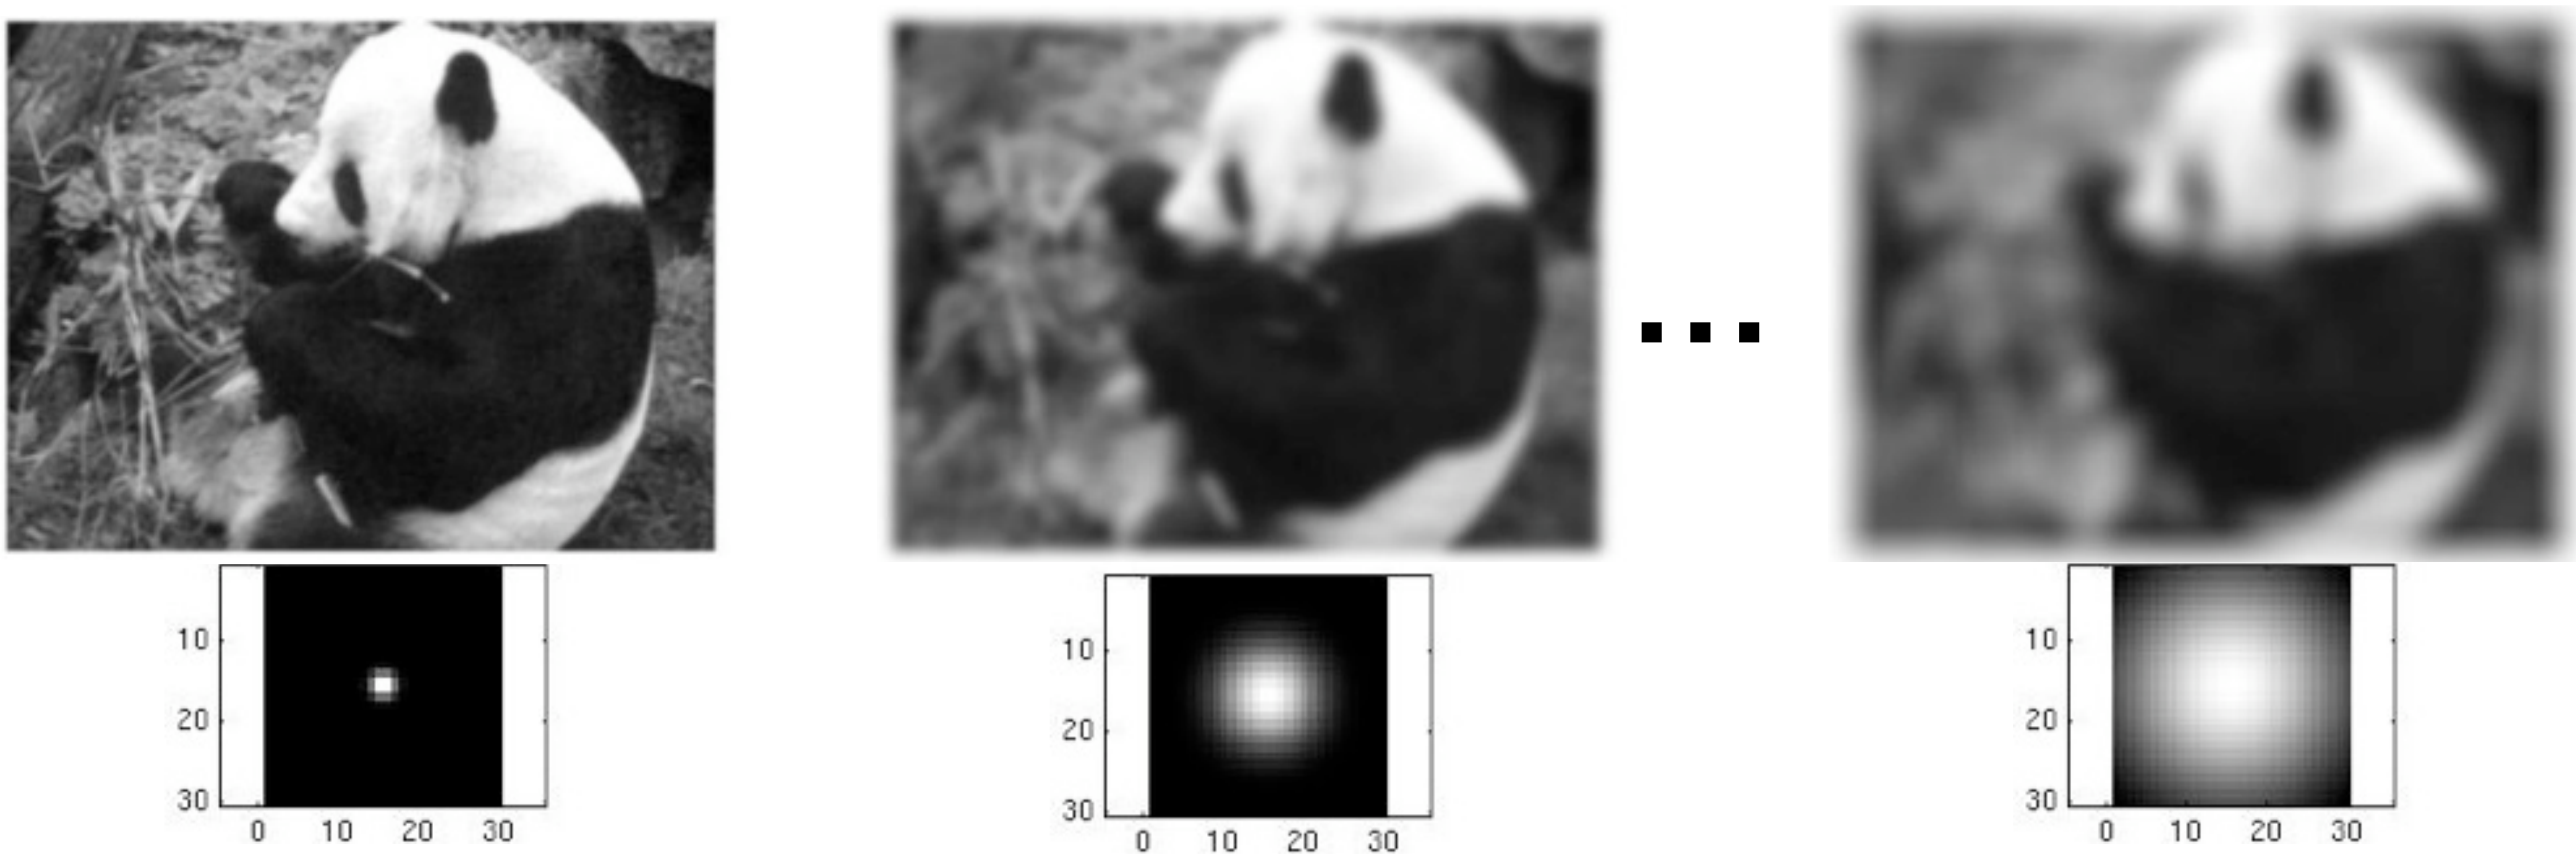

Effect of σ on derivatives

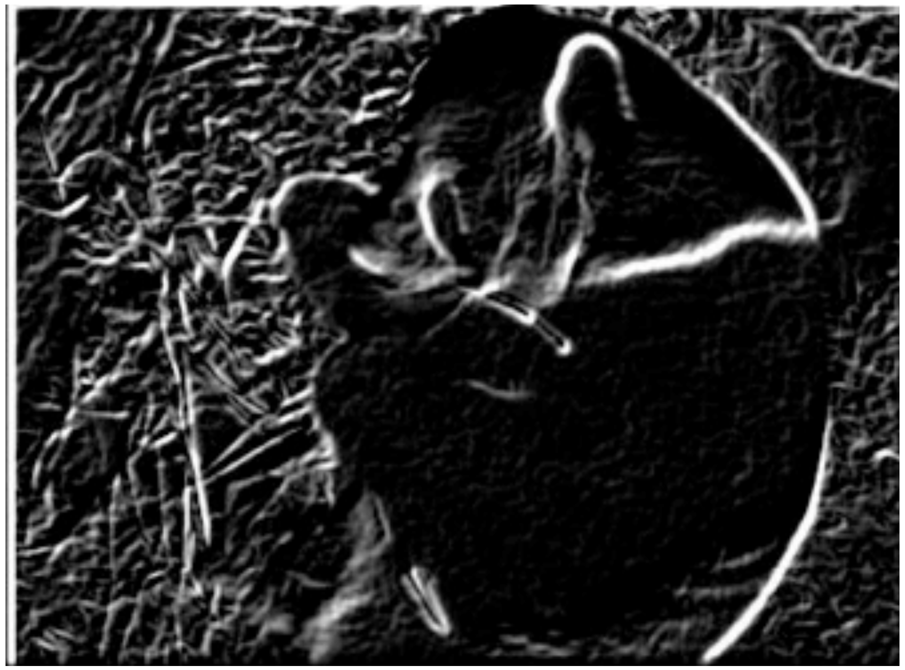

σ = 1 pixel

σ = 3 pixels

The apparent structures differ depending on Gaussian's scale parameter.

Larger values: larger scale edges detected
Smaller values: finer features detected

Laplacian of Gaussian

Consider smoothing (Gaussian) + second derivative [-1 2 1]

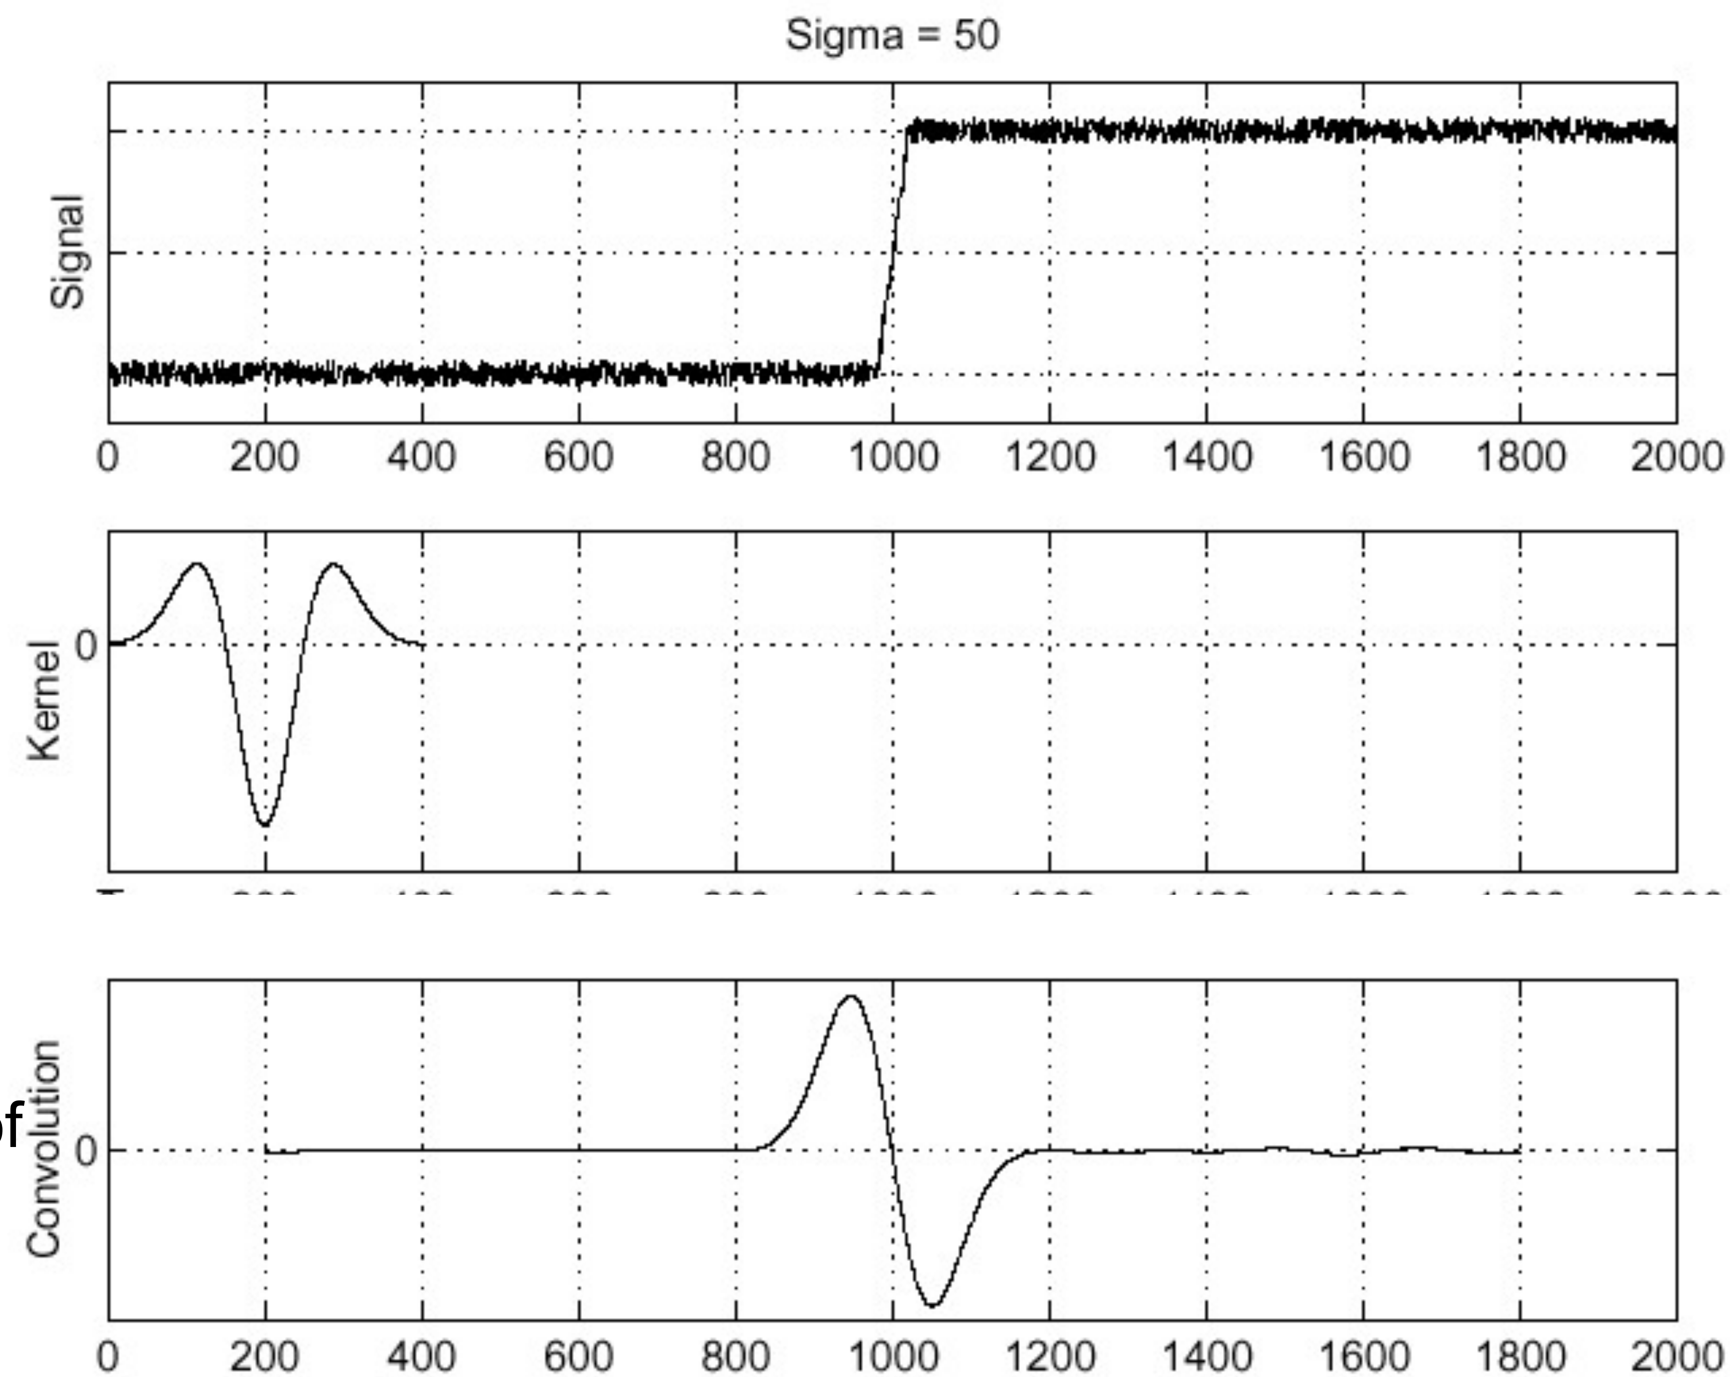

Laplacian of Gaussian (LOG) operator

Edges = zero-crossings of LOG

Where is the edge?

Zero-crossings of bottom graph

2D edge detection filters

Let $h_\sigma(u,v)$ be a 2D Gaussian

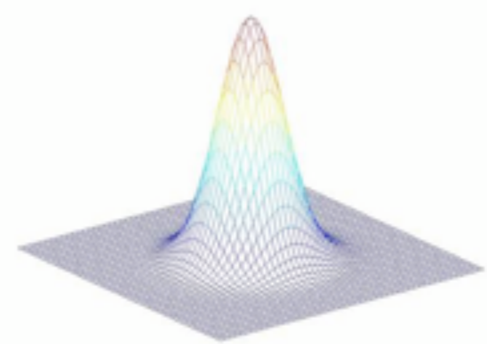

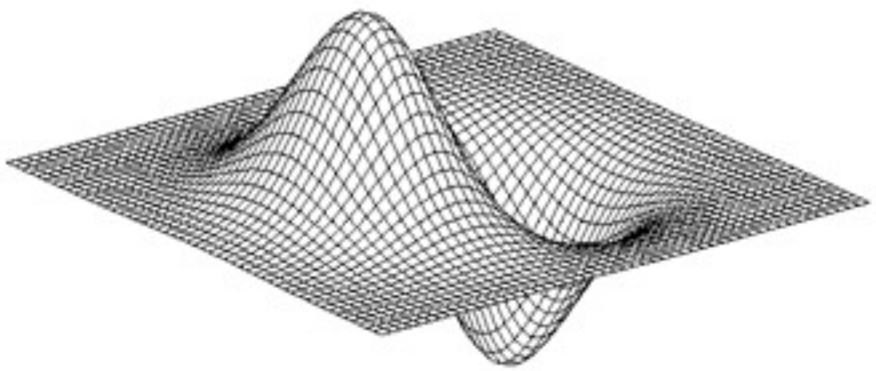

derivative of Gaussian

$$\frac{\partial}{\partial x} h_\sigma(u,v)$$

Laplacian of Gaussian

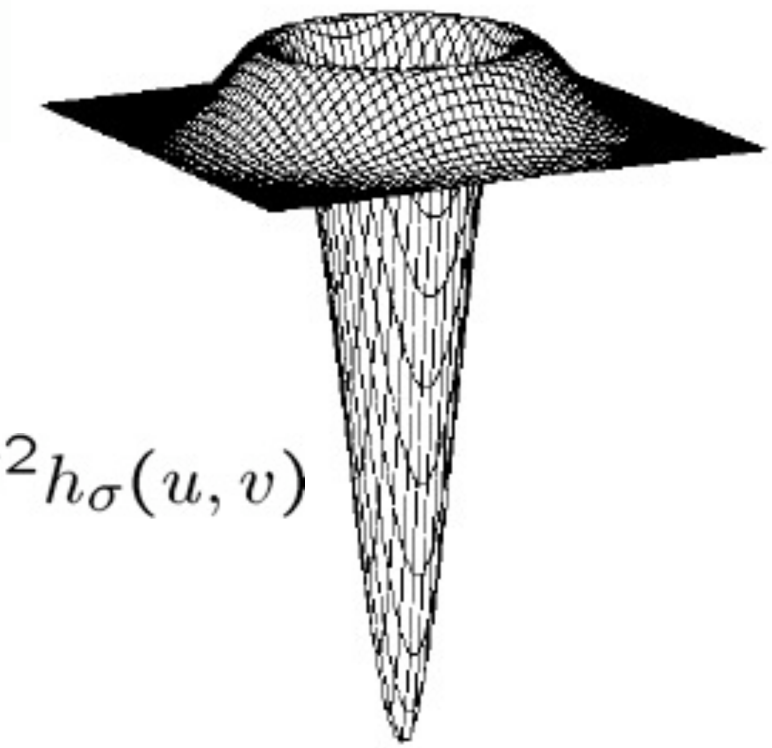

$$\nabla^2 h_\sigma(u,v)$$

∇^2 is the **Laplacian** operator: $\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$

The Canny edge detector

original image (Lena)

Source: S. Seitz

The Canny edge detector

norm of the gradient

The Canny edge detector

thresholding

The Canny edge detector

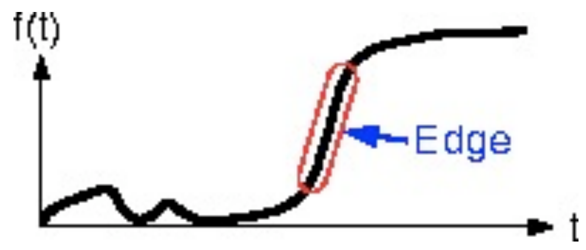

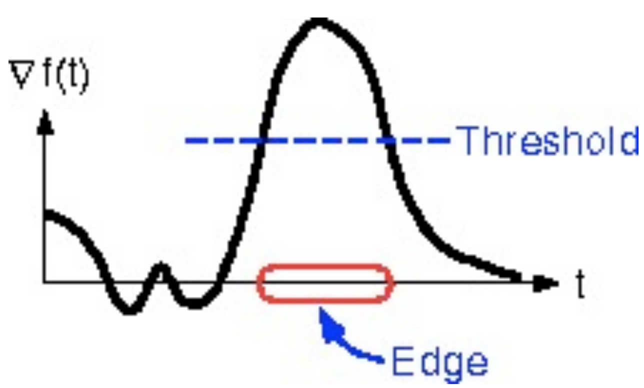

How to turn these thick regions of the gradient into curves?

Non-maximum suppression

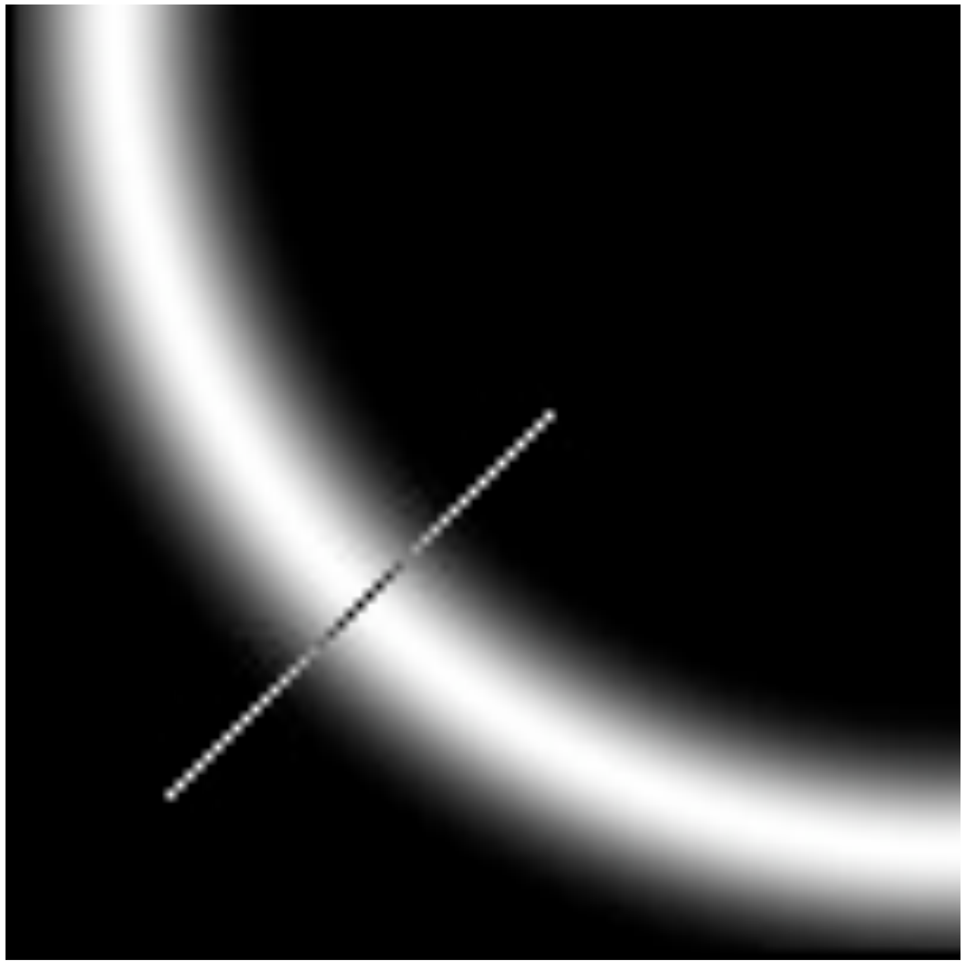

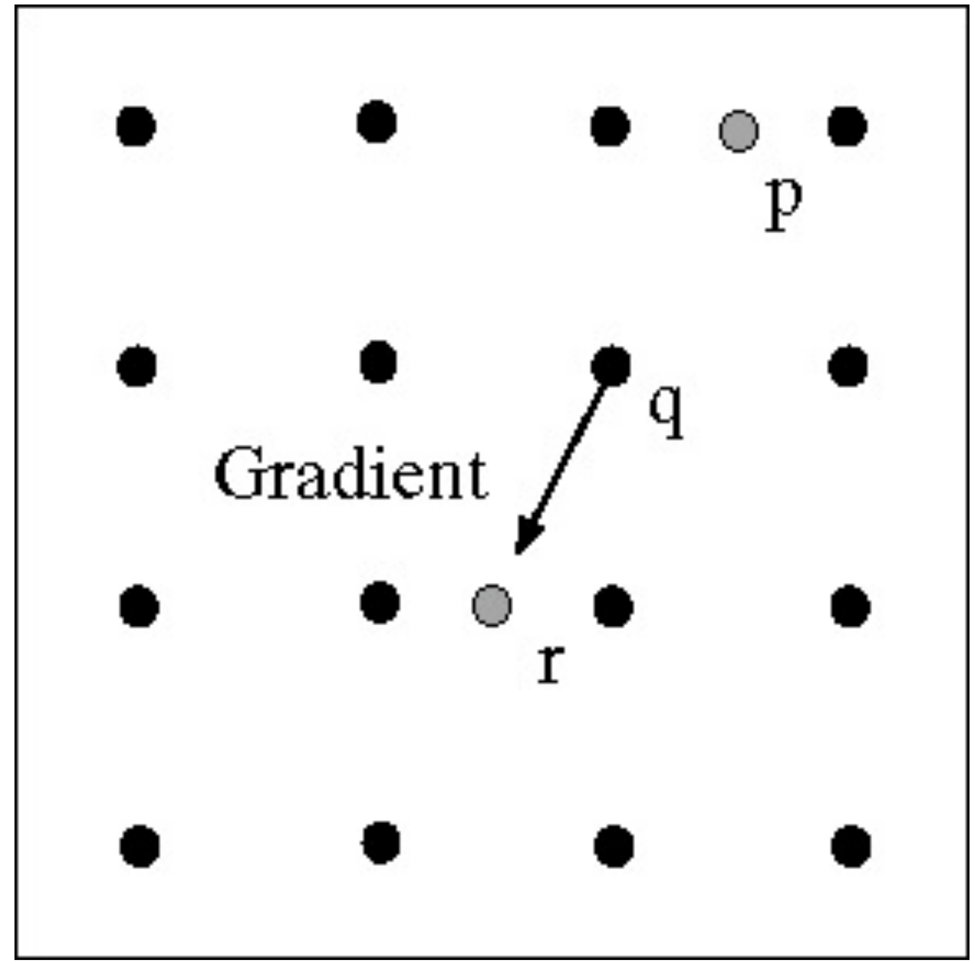

Check if pixel is local maximum along gradient direction, select single max across width of the edge

- requires checking interpolated pixels p and r

The Canny edge detector

thinning
(non-maximum suppression)

Problem: pixels along this edge didn't survive the thresholding

Hysteresis thresholding

- Check that maximum value of gradient value is sufficiently large
  - drop-outs? use **hysteresis**
    - use a high threshold to start edge curves and a low threshold to continue them.

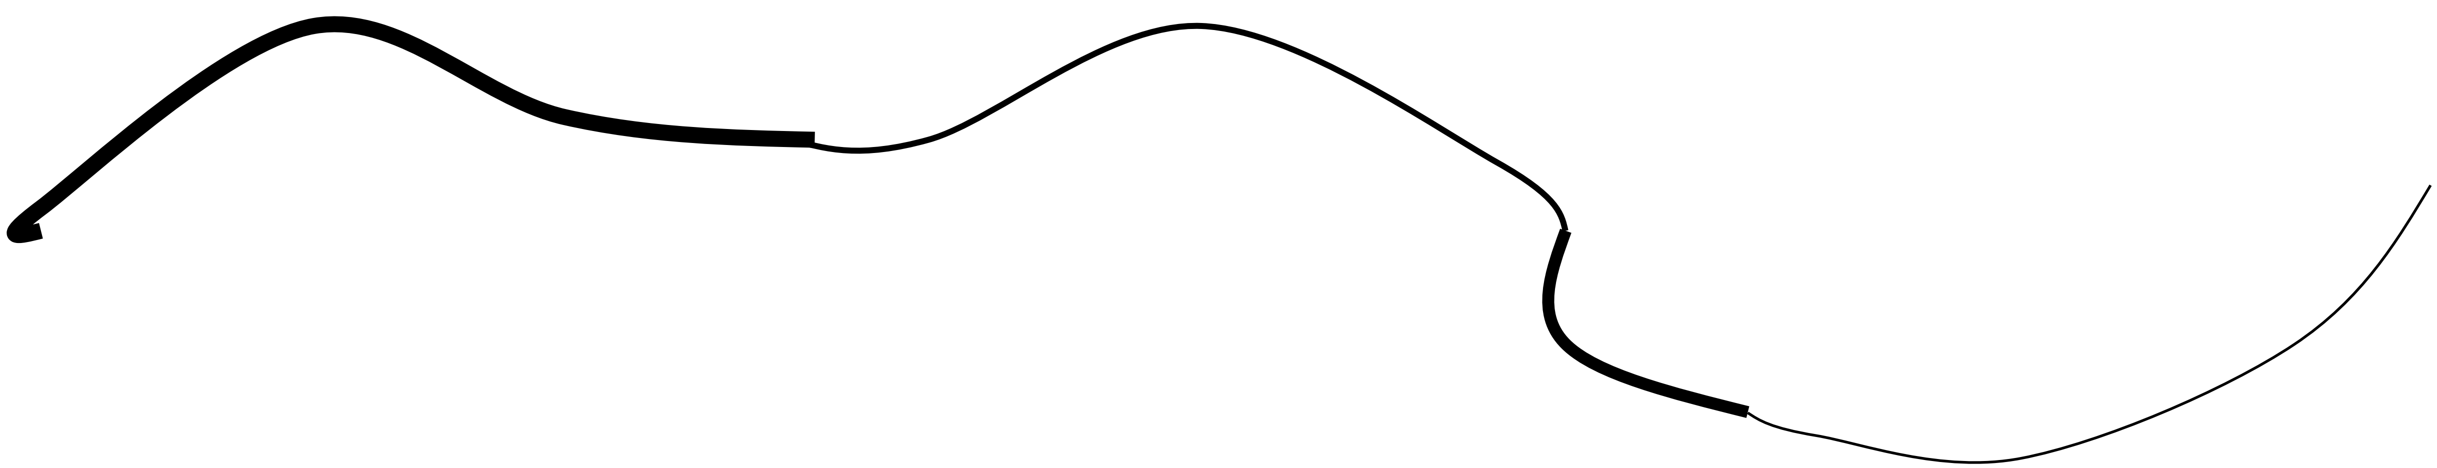

Source: S. Seitz

Hysteresis thresholding

original image

high threshold
(strong edges)

low threshold
(weak edges)

hysteresis threshold

Source: L. Fei-Fei