

# **CMPSCI 670: Computer Vision**

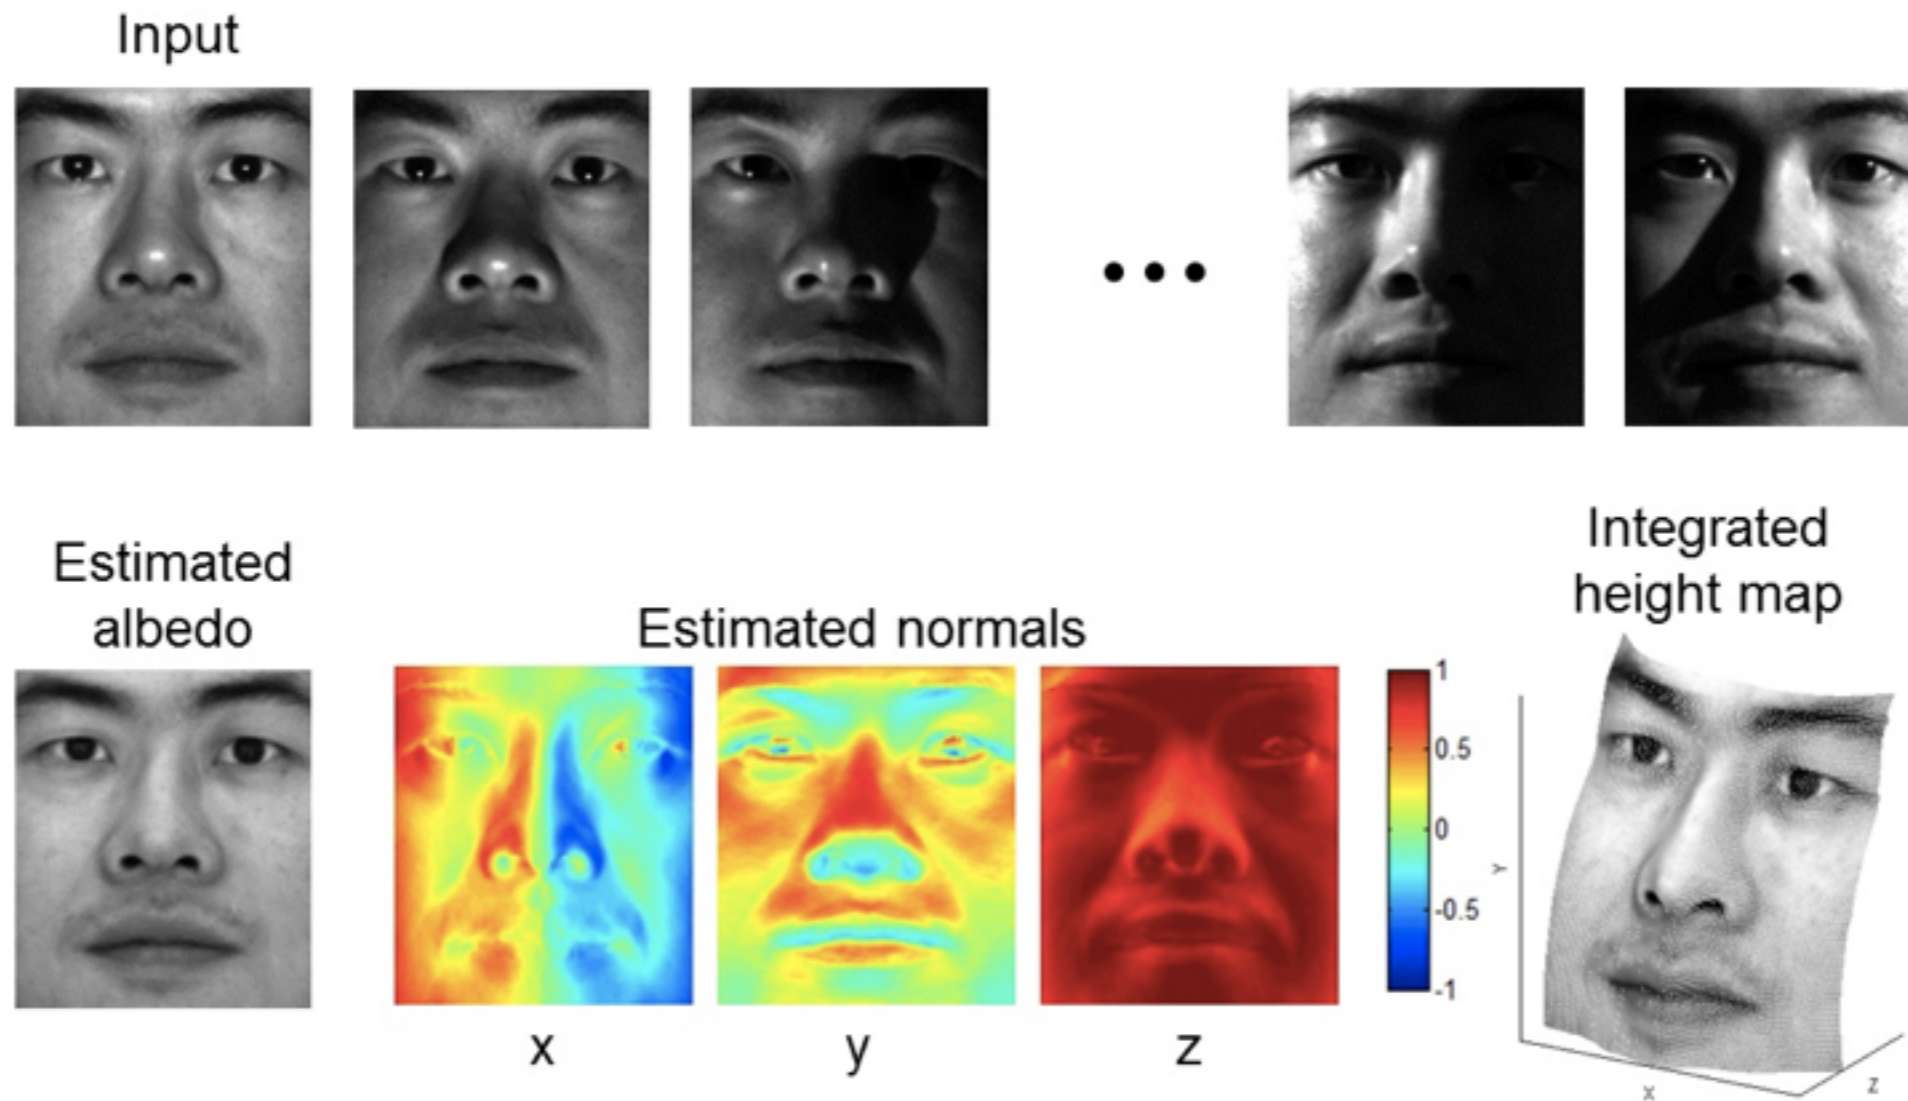
## Linear filtering cont., Edge detection

University of Massachusetts, Amherst  
September 24, 2014

Instructor: Subhransu Maji

# Administrivia


- Homework 2 is up (due October 6 before class starts)
- Photometric stereo



- Office hours this week (~~today~~ tomorrow) Th 3:45-4:45, CS274

# Joys of computer vision research

<http://xkcd.com/1425/>



IN CS, IT CAN BE HARD TO EXPLAIN  
THE DIFFERENCE BETWEEN THE EASY  
AND THE VIRTUALLY IMPOSSIBLE.

# Motivation: Image de-noising

- How can we reduce noise in a photograph?



# Moving average

- Let's replace each pixel with a *weighted* average of its neighborhood
- The weights are called the *filter kernel*
- What are the weights for the average of a 3x3 neighborhood?




$$\frac{1}{9}$$

1	1	1
1	1	1
1	1	1

“box filter”




# Gaussian vs. box filtering



$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$

# Alternative idea: Median filtering


- A **median filter** operates over a window by selecting the median intensity in the window



- Is median filtering linear?

# Median filter

- What advantage does median filtering have over Gaussian filtering?
  - Robustness to outliers







# Median filter

Salt-and-pepper noise



Median filtered



MATLAB: `medfilt2(image, [h w])`

# Gaussian vs. median filtering

3x3

5x5

7x7

Gaussian




Median






# Sharpening revisited



**before**



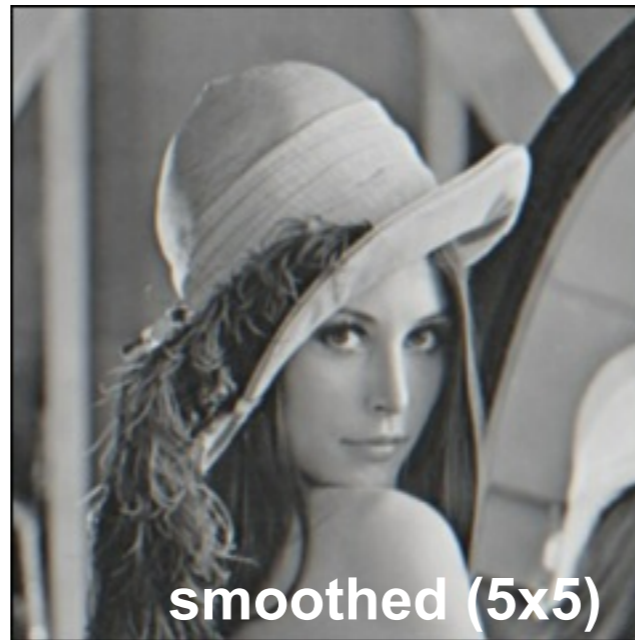
**after**

# Sharpening revisited

What does blurring take away?




−




=



Let's add it back:



+  $\alpha$



=






# Unsharp mask filter

$$f + \alpha(f - f * g) = (1 + \alpha)f - \alpha f * g = f * ((1 + \alpha)e - g)$$

↑  
image


↑  
blurred  
image

↑  
unit impulse  
(identity)



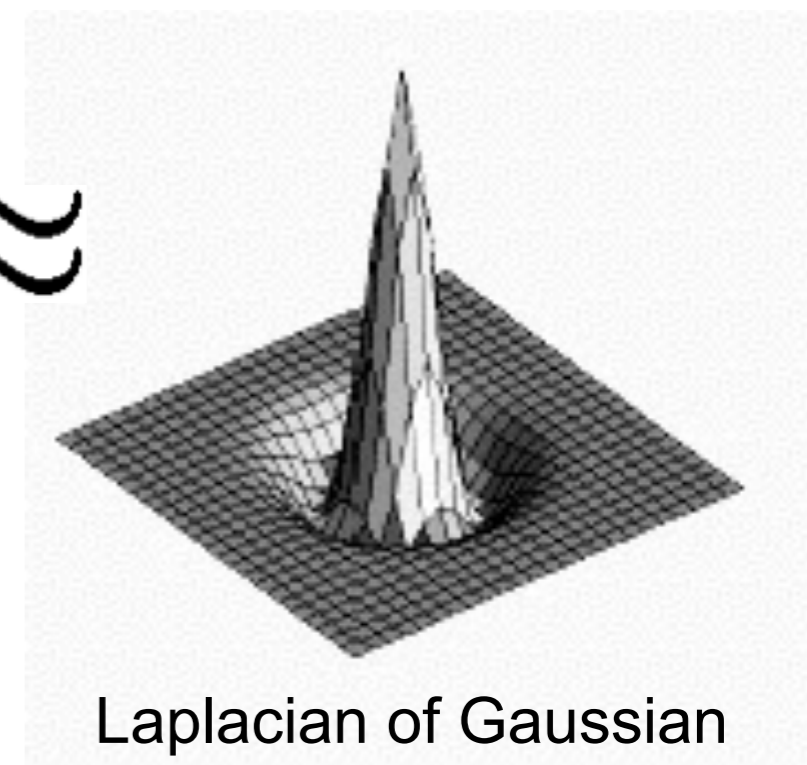
unit impulse

—



Gaussian

≈




Laplacian of Gaussian



# Application: Hybrid Images

Gaussian Filter



Laplacian Filter

A. Oliva, A. Torralba, P.G. Schyns,  
[“Hybrid Images,”](#) SIGGRAPH 2006



# Changing expression





SIGGRAPH2006

Sad



Surprised






motorcycle and bicycle





dolphin and car





# Edge detection



[Winter in Kraków photographed by Marcin Ryczek](#)




# Edge detection

- **Goal:** Identify sudden changes (discontinuities) in an image
  - Intuitively, most semantic and shape information from the image can be encoded in the edges
  - More compact than pixels
- **Ideal:** artist's line drawing (but artist is also using object-level knowledge)



# Origin of edges


Edges are caused by a variety of factors:




# Edge detection

- An edge is a place of rapid change in the image intensity function


image



intensity function  
(along horizontal scanline)



first derivative



edges correspond to  
extrema of derivative

# Derivatives with convolution

For 2D function  $f(x,y)$ , the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x, y)}{1}$$



To implement the above as convolution, what would be the associated filter?

# Partial derivatives of an image



$$\frac{\partial f(x, y)}{\partial x}$$

$$\frac{\partial f(x, y)}{\partial y}$$



-1	1
----	---

-1
1

 or 

1
-1

Which one shows changes with respect to x?



# Finite difference filters

Other approximations of derivative filters exist:

**Prewitt:**  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$




**Sobel:**  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

**Roberts:**  $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

# 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 increase in intensity

- How does this direction relate to the direction of the edge?



The gradient direction 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

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

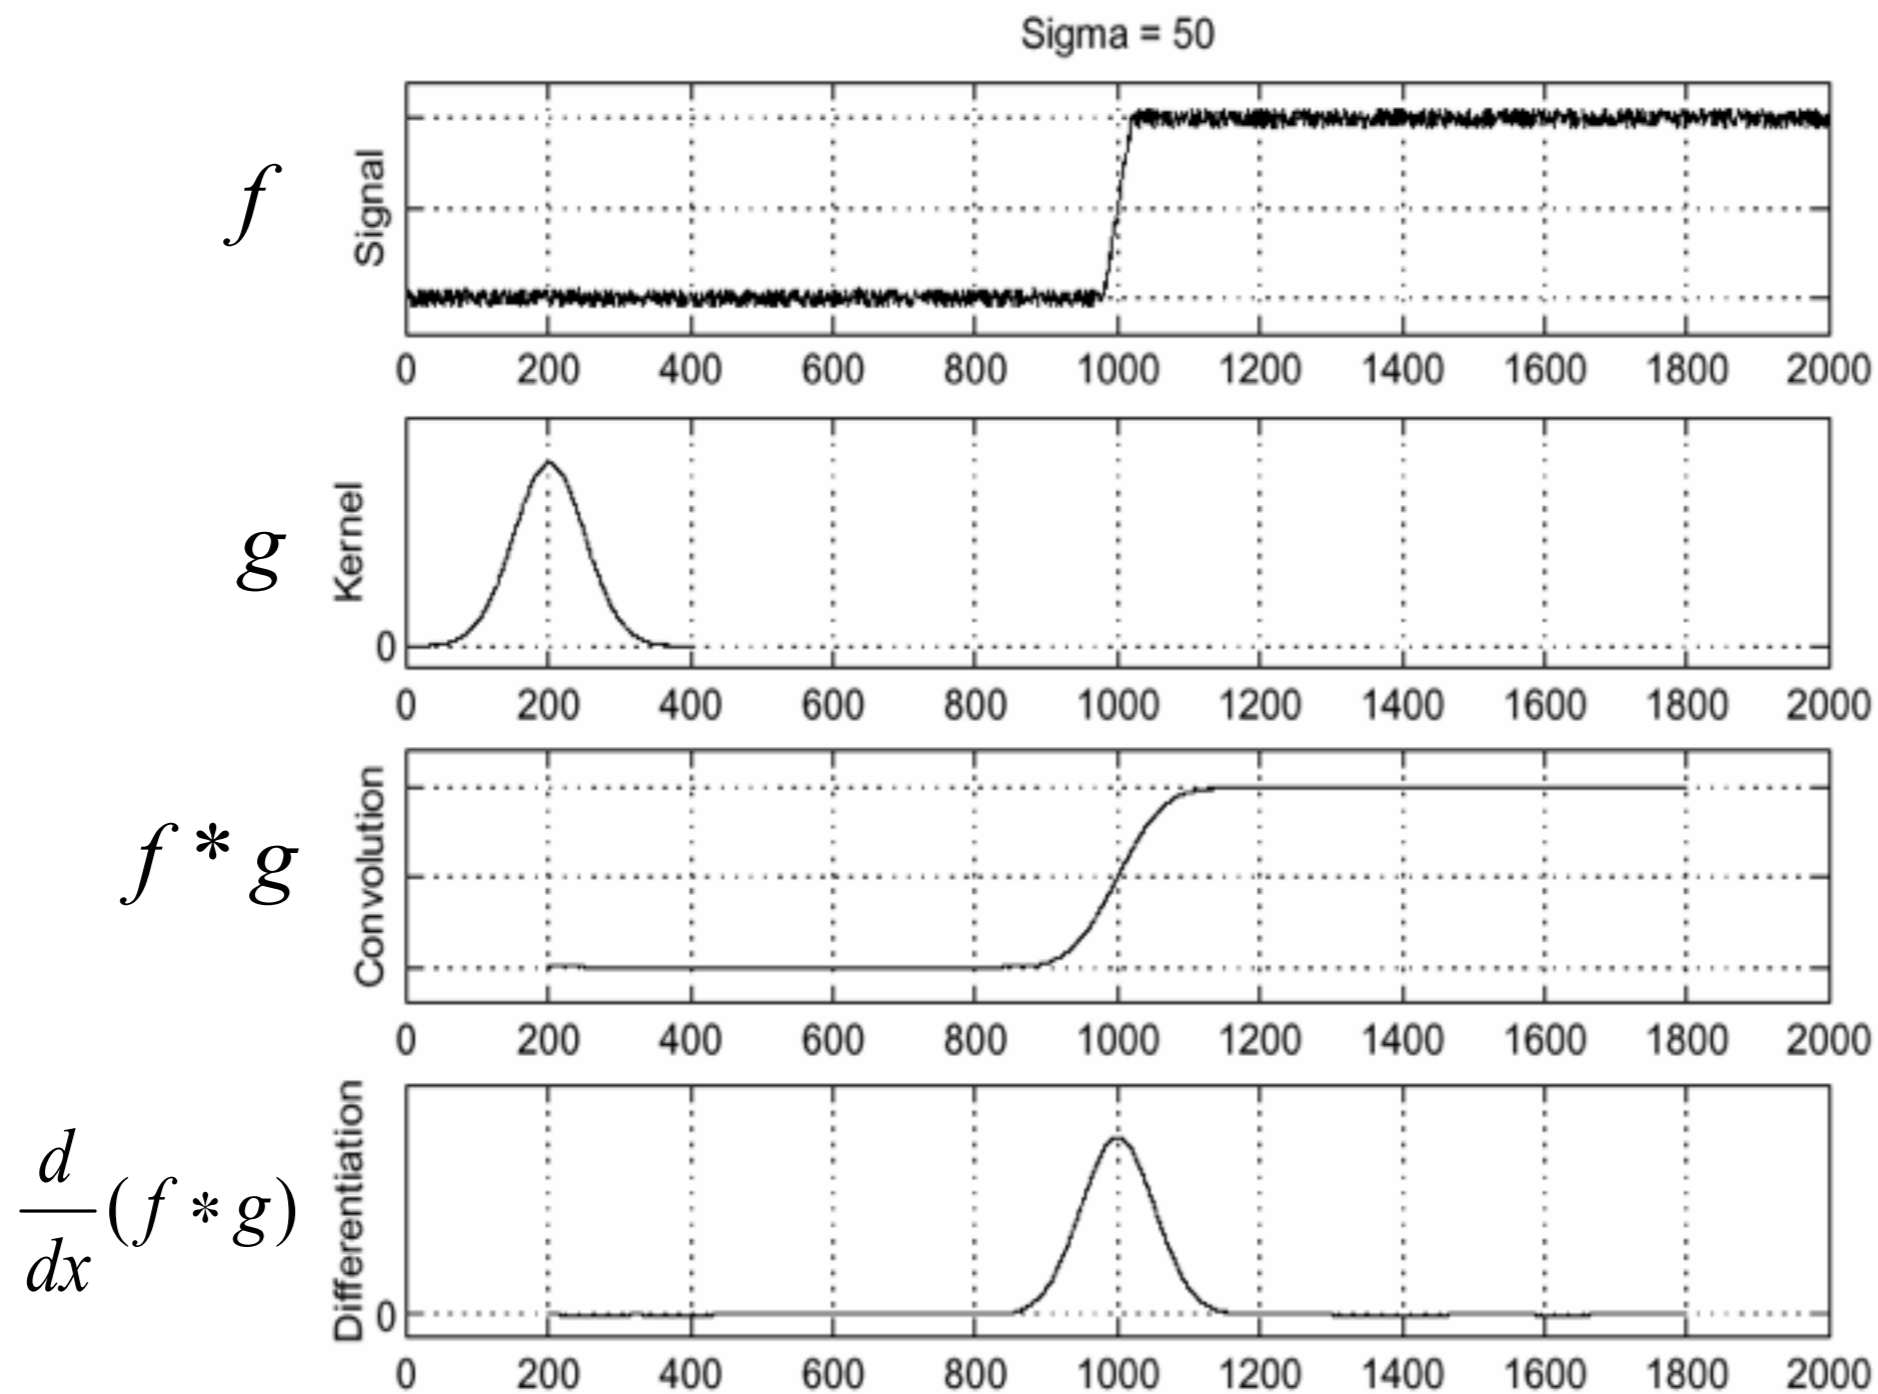
# Effects of noise

Consider a single row or column of the image



Where is the edge?

# Solution: smooth first



- To find edges, look for peaks in  $\frac{d}{dx}(f * g)$




# Derivative theorem of convolution


- Differentiation is convolution, and convolution is associative:

$$\frac{d}{dx}(f * g) = f * \frac{d}{dx}g$$



- This saves us one operation:




# Derivative of Gaussian filters



x-direction




y-direction





Which one finds horizontal/vertical edges?

# Derivative of Gaussian filters



x-direction



y-direction



Are these filters separable?




# Recall: Separability of the Gaussian filter

$$\begin{aligned} G_{\sigma}(x, y) &= \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right)\right) \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{y^2}{2\sigma^2}\right)\right) \end{aligned}$$

The 2D Gaussian can be expressed as the product of two functions, one a function of  $x$  and the other a function of  $y$

In this case, the two functions are the (identical) 1D Gaussian

# Scale of Gaussian derivative filter



1 pixel

3 pixels


7 pixels

Smoothed derivative removes noise, but blurs edge. Also finds edges at different “scales”

# Review: Smoothing vs. derivative filters



## Smoothing filters

- Gaussian: remove “high-frequency” components; “low-pass” filter
- Can the values of a smoothing filter be negative?
- What should the values sum to?
  - **One:** constant regions are not affected by the filter



## Derivative filters

- Derivatives of Gaussian
- Can the values of a derivative filter be negative?
- What should the values sum to?
  - **Zero:** no response in constant regions
- High absolute value at points of high contrast





# The Canny edge detector

1. Filter image with derivative of Gaussian
2. Find magnitude and orientation of gradient
3. **Non-maximum suppression:**
  - Thin wide “ridges” down to single pixel width
4. **Linking and thresholding (hysteresis):**
  - Define two thresholds: low and high
  - Use the high threshold to start edge curves and the low threshold to continue them

MATLAB: `edge(image, 'canny');`

J. Canny, [\*A Computational Approach To Edge Detection\*](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.

# The Canny edge detector

original image



# 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?

thresholding

# 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




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

thinning  
(non-maximum suppression)

# Hysteresis thresholding

Use a high threshold to start edge curves, and a low threshold to continue them.



# Hysteresis thresholding



**original image**



**high threshold  
(strong edges)**



**low threshold  
(weak edges)**



**hysteresis threshold**

# Recap: Canny edge detector

1. Compute x and y gradient images
2. Find magnitude and orientation of gradient
3. **Non-maximum suppression:**
  - Thin wide “ridges” down to single pixel width
4. **Linking and thresholding (hysteresis):**
  - Define two thresholds: low and high
  - Use the high threshold to start edge curves and the low threshold to continue them

MATLAB: `edge(image, 'canny');`

J. Canny, [\*A Computational Approach To Edge Detection\*](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.



# Further thoughts and readings ...

- Hybrid images project
  - <http://cvcl.mit.edu/hybridimage.htm>
- Canny edge detector
  - [www.limsi.fr/Individu/vezien/PAPIERS\\_ACS/canny1986.pdf](http://www.limsi.fr/Individu/vezien/PAPIERS_ACS/canny1986.pdf)
- Bilateral filtering for image denoising (and other application)
  - [http://people.csail.mit.edu/sparis/bf\\_course/](http://people.csail.mit.edu/sparis/bf_course/)
  
- If all else fails [www.xkcd.com](http://www.xkcd.com)