

CS4670: Computer Vision



Noah Snavely

Lecture 2: Convolution and edge detection

SHADOW



From [Sandlot Science](#)

Gaussian Kernel





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


Mean vs. Gaussian filtering




Gaussian filters



$\sigma = 1$ pixel



$\sigma = 5$ pixels




$\sigma = 10$ pixels



$\sigma = 30$ pixels


Gaussian filter

- Removes “high-frequency” components from the image (low-pass filter)
- Convolution with self is another Gaussian




- Convoluting two times with Gaussian kernel of width σ = convoluting once with kernel of width $\sigma\sqrt{2}$

Sharpening




before




after

Sharpening revisited

- What does blurring take away?



Let's add it back:




Sharpen filter

$$F + \alpha (F - \underbrace{F * H}_{\text{blurred image}})$$


↑ image

↑
unit impulse
(identity)




scaled impulse

−



Gaussian

≈



Laplacian of Gaussian

Sharpen filter



Convolution in the real world

Camera shake



Source: Fergus, *et al.* "Removing Camera Shake from a Single Photograph", SIGGRAPH 2006

Bokeh: Blur in out-of-focus regions of an image.



Source: <http://lullaby.homepage.dk/diy-camera/bokeh.html>

Questions?

Image noise



Original image

$$F[x, y]$$



White Gaussian noise

$$F[x, y] + \mathcal{N}(0, \sigma)$$



Salt and pepper noise

(each pixel has some chance of being switched to zero or one)

<http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/index.html>

Gaussian noise



$F[x, y] + \mathcal{N}(0, 5\%)$



$\sigma = 1$ pixel



$\sigma = 2$ pixels



$\sigma = 5$ pixels

Smoothing with larger standard deviations suppresses noise, but also blurs the image

Salt & pepper noise – Gaussian blur



$p = 10\%$



$\sigma = 1$ pixel



$\sigma = 2$ pixels




$\sigma = 5$ pixels

- What's wrong with the results?

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?

filters have width 5 :



Salt & pepper noise – median filtering



$p = 10\%$



$\sigma = 1$ pixel



$\sigma = 2$ pixels



$\sigma = 5$ pixels



3x3 window





5x5 window




7x7 window

Edge detection



- Convert a 2D image into a set of curves
 - Extracts salient features of the scene
 - More compact than pixels

Origin of Edges



- Edges are caused by a variety of factors

Characterizing edges

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




Image derivatives

- How can we differentiate a *digital* image $F[x,y]$?
 - Option 1: reconstruct a continuous image, f , then compute the derivative
 - Option 2: take discrete derivative (finite difference)

$$\frac{\partial f}{\partial x}[x, y] \approx F[x + 1, y] - F[x, y]$$

How would you implement this as a linear filter?

$$\frac{\partial f}{\partial x} : \begin{array}{|c|c|c|} \hline & & \\ \hline & & \\ \hline & & \\ \hline \end{array}$$

H_x




$$\frac{\partial f}{\partial y} : \begin{array}{|c|c|c|} \hline & & \\ \hline & & \\ \hline & & \\ \hline \end{array}$$

H_y

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



The *edge strength* is given by the gradient magnitude:



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

The gradient direction is given by:


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

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

Image gradient




Effects of noise




Noisy input image

$$f(x)$$

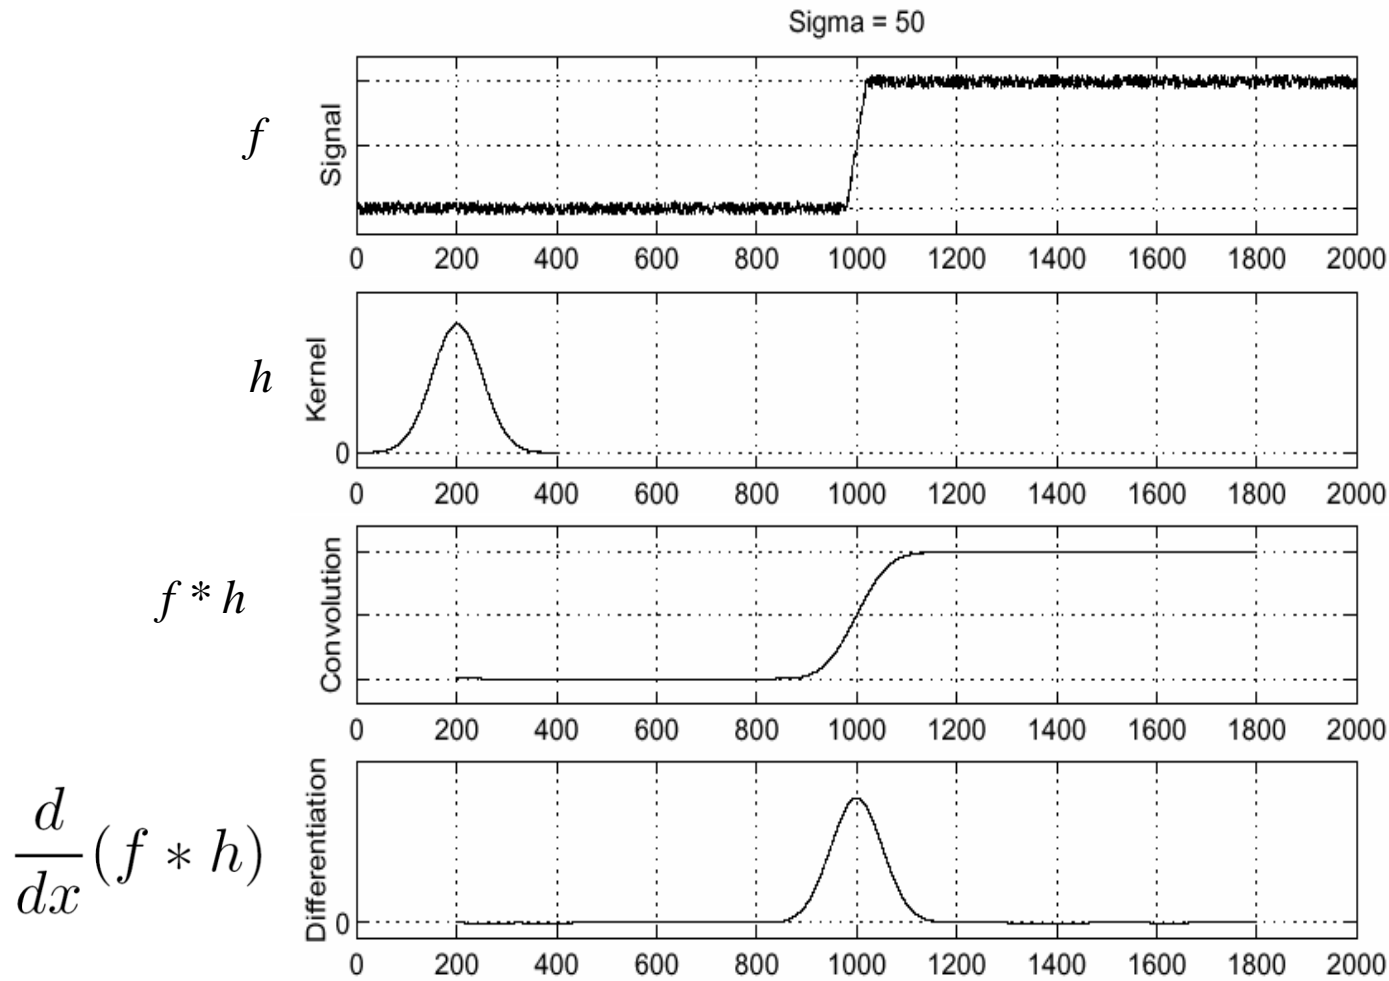


$$\frac{d}{dx} f(x)$$



Where is the edge?


Solution: smooth first




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

Associative property of convolution

- Differentiation is convolution, and convolution is associative: $\frac{d}{dx}(f * h) = f * \frac{d}{dx}h$
- This saves us one operation:




2D edge detection filters



Gaussian


$$h_{\sigma}(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$




derivative of Gaussian (x)

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



Derivative of Gaussian filter



x-direction



y-direction



The Sobel operator

- Common approximation of derivative of Gaussian

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline -2 & 0 & 2 \\ \hline -1 & 0 & 1 \\ \hline \end{array}$$

s_x

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 0 & 0 & 0 \\ \hline -1 & -2 & -1 \\ \hline \end{array}$$

s_y

- The standard defn. of the Sobel operator omits the $1/8$ term
 - doesn't make a difference for edge detection
 - the $1/8$ term **is** needed to get the right gradient value