

# Computer Vision

## Linear Filtering and Edge Detection

Professor Hager  
<http://www.cs.jhu.edu/~hager>

10/15/04

CS 461, Copyright G.D. Hager


## Outline

- Image noise
- Filtering by Convolution
- Properties of Convolution
- Derivative Operators
- The Canny Edge Detector

10/15/04

CS 461, Copyright G.D. Hager

## IMAGE NOISE



Cameras are not perfect sensors  
Scenes never quite match our expectations

10/15/04

CS 461, Copyright G.D. Hager

## Noise Models

- Noise is commonly modeled using the notion of “additive white noise.”
  - Scalar example:  $I(x) = I^*(x) + n(x)$
  - Images:  $I(i,j,t) = I^*(i,j,t) + n(i,j,t)$
  - Note that  $n(i,j,t)$  is independent of  $n(i',j',t')$  unless  $i'=i, j'=j, t'=t$ .
  - Typically we assume that  $n$  (noise) is independent of image location --- that is, it is i.i.d

10/15/04

CS 461, Copyright G.D. Hager

# Properties of Noise Processes

- Properties of temporal image noise:

$$\text{Mean } \mu(i,j) = \sum_t I(i,j,t)$$

$$\text{Standard Deviation } \sigma_{i,j} = \text{Sqrt} \left( \sum_t (\mu - I(i,j,t))^2 \right)$$


$$\text{Signal-to-noise Ratio } \frac{\sigma_{I^*}}{\sigma_{i,j}}$$

10/15/04

CS 461, Copyright G.D. Hager

## Image Noise

- An experiment: take several images of a static scene and look at the pixel values



mean = 38.6

std = 2.99

max snr = 255/3 = 85

10/15/04

CS 461, Copyright G.D. Hager

PROPERTIES OF TEMPORAL IMAGE NOISE  
(i.e., successive images)

**If standard deviation of grey values at a pixel is  $\sigma$  for a pixel for a single image, then the laws of statistics states that for independent sampling of grey values, for a temporal average of  $n$  images, the standard deviation is:**

$$\frac{\sigma}{\text{Sqrt}(n)}$$

10/15/04

CS 461, Copyright G.D. Hager

## Temporal vs. Spatial Noise

- It is common to assume that:
  - spatial noise in an image is consistent with the temporal image noise
  - the noise is independent and identically distributed
- Thus, we can think of the image itself as an additive noise process

10/15/04

CS 461, Copyright G.D. Hager

Gaussian  
Noise:  
 $\sigma=1$



10/15/04

Gaussian  
Noise:  
 $\sigma=16$



10/15/04

# How to reduce noise

- Averaging is a common way to reduce noise
  - instead of temporal averaging, how about spatial?
- For a pixel in image I at I,j

$$I'(i, j) = 1/9 \sum_{i'=i-1}^{i+1} \sum_{j'=j-1}^{j+1} I(i', j')$$

10/15/04

CS 461, Copyright G.D. Hager

## Convolution




Image (I)

$$* \begin{bmatrix} 1 & 2 & 1 \\ & & \\ -1 & -2 & -1 \end{bmatrix}$$


Kernel (K)

Note: Typically Kernel is relatively small in vision applications.

10/15/04

CS 461, Copyright G.D. Hager

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$



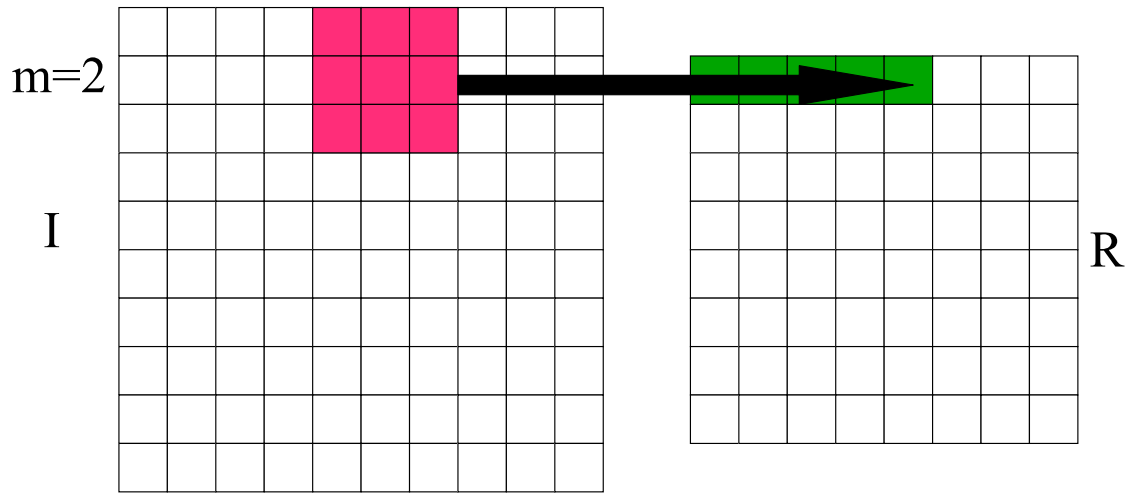
Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

Copyright © 2004, D. K. Regier

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

Copyright © 2004, D. K. Regier

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 461, Copyright © R. Rubinfeld

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

Copyright © 2004, D. K. Regier

# Convolution: $R = K * I$




Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

Copyright © 2004, D. K. Regier

# Convolution: $R = K * I$



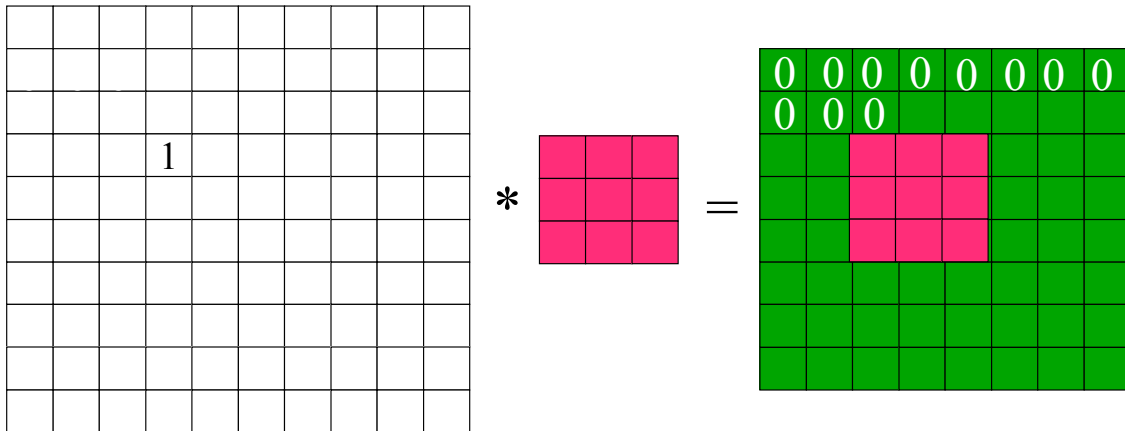
Kernel size  
is  $m+1$  by  $m+1$

$$R(i, j) = \sum_{h=-m/2}^{m/2} \sum_{k=-m/2}^{m/2} K(h, k) I(i-h, j-k)$$

10/15/04

CS 439, Stanford University, Copyright © 2004

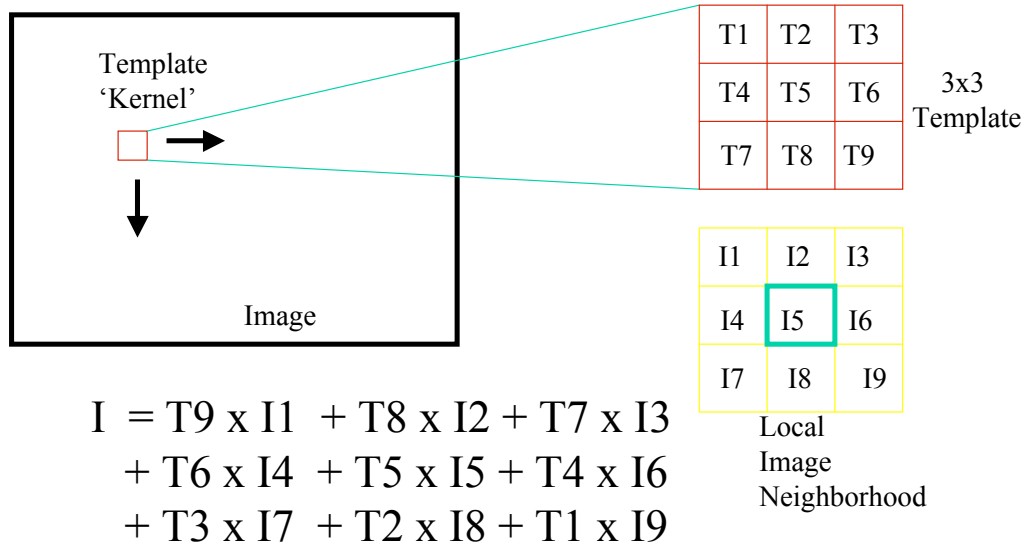
# Impulse Response



10/15/04

CS 439, Stanford University, Copyright © 2004

# DISCRETE CONVOLUTION



10/15/04

CS 461, Copyright G.D. Hager

## Convolution Formula

We often write  $I' = I * B$  to represent the convolution of  $I$  by  $B$ .  $B$  is referred to as the *kernel* of the convolution (or sometimes the “stencil” in the discrete case).


$$(I * B)(x,y) = \sum_i \sum_j I(x-i,y-j) * B(i,j)$$

Note these indices run backwards --- this can sometimes fool you down the road!


10/15/04

CS 461, Copyright G.D. Hager

# Linear filtering (warm-up slide)



original




?



10/15/04

CS46 © Copyright B. Eckel

# Linear filtering (warm-up slide)



original




Filtered  
(no change)


10/15/04

CS46 © Copyright B. Eckel

# Linear filtering



original




?



10/15/04

CS46 (Copyright © Eric Stenger)

## shift



original




shifted


10/15/04

CS46 (Copyright © Eric Stenger)

# Linear filtering



original




?



10/15/04

CS46 (Copyright © Ericnie)

# Blurring



original




Blurred (filter applied in both dimensions).

10/15/04

CS46 (Copyright © Ericnie)

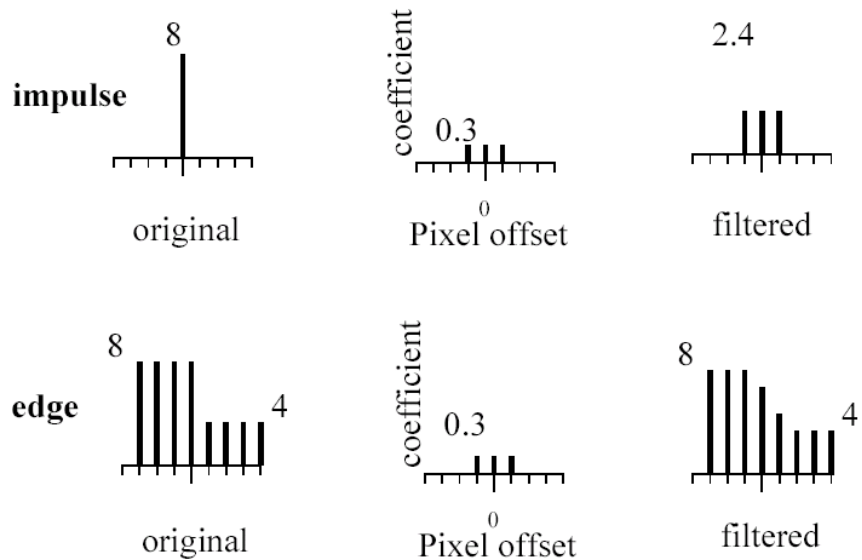
# Blur examples



10/15/04

CS 461 Copyright © K. Regina


# Blur examples



10/15/04

CS 461 Copyright © K. Regina


# Linear filtering (warm-up slide)



10/15/04

CS46 (© copyright B. Eckel)


# Linear filtering (no change)




10/15/04

CS46 (© copyright B. Eckel)

# Linear filtering



original




?



10/15/04

CS 464 Copyright © D. Crandall

(remember blurring)



original




Blurred (filter applied in both dimensions).

10/15/04

CS 464 Copyright © D. Crandall


# Sharpening



10/15/04

CS 461, Copyright © D. K. Regier

# Sharpening



10/15/04

CS 461, Copyright © D. K. Regier

# How to Reduce Noise

- For a pixel in image I at I,j

$$I'(i, j) = 1/9 \sum_{i'=i-1}^{i+1} \sum_{j'=j-1}^{j+1} I(i', j')$$

- Computing this for every pixel location is the *convolution* of the image I with the *template* (or *kernel*) consisting of a 3x3 array of 1's.
- Note that is this  $O(n^2m^2)$  for an nxn image and mxm template.
- Note we have to normalize the template to 1 to make sure we don't introduce any scaling into the image.

10/15/04

CS 461, Copyright G.D. Hager

# Smoothing by Averaging

Kernel: 



10/15/04

CS 461, Copyright G.D. Hager

# Some Convolution Facts

- We often write  $I' = I * B$  to represent the convolution of  $I$  by  $B$ .  $B$  is referred to as the *kernel* of the convolution (or sometimes the “stencil” in the discrete case).
- Note convolution is
  - Associative
  - Commutative
  - A linear operator
- We are using a *discrete convolution*; we will see this is not always consistent with an underlying *continuous convolution* that we may wish to implement
- Convolution is formally defined on unbounded images and kernels.
  - padding schemes:
    - pad with zeros (same size vs. full size)
    - compute only legal values

} time for a Matlab demo!

10/15/04

CS 461, Copyright G.D. Hager

# Understanding Convolution

- Another way to think about convolution is in terms of how it changes the *frequency distribution* in the image.
- Recall the *fourier* representation of a function
  - $F(u) = \iint f(x) e^{-2\pi i u x} dx$
  - recall that  $e^{-2\pi i u x} = \cos(2\pi u x) - i \sin(2\pi u x)$
  - Also we have  $f(x) = \iint F(u) e^{2\pi i u x} du$
  - $F(u) = |F(u)| e^{i \Phi(u)}$ 
    - a decomposition into magnitude and phase
  - $|F(u)|^2$  is the *power spectrum*
- Questions: what function takes many many many terms in the Fourier expansion?

10/15/04

CS 461, Copyright G.D. Hager

# Understanding Convolution

## Discrete Fourier Transform (DFT)

$$F[u, v] \equiv \frac{1}{N} \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} I[x, y] e^{-\frac{2\pi j}{N} (xu+yv)}$$

## Inverse DFT

$$I[x, y] \equiv \frac{1}{N} \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} F[u, v] e^{+\frac{2\pi j}{N} (ux+vy)}$$

Implemented via the “Fast Fourier Transform” algorithm (FFT)

10/15/04

CS 461, Copyright G.D. Hager



Fourier basis element

$$e^{-i2\pi(ux+vy)}$$


Transform is sum of orthogonal basis functions

Vector (u,v)

- Magnitude gives frequency
- Direction gives orientation.





Here  $u$  and  $v$  are larger than in the previous slide.



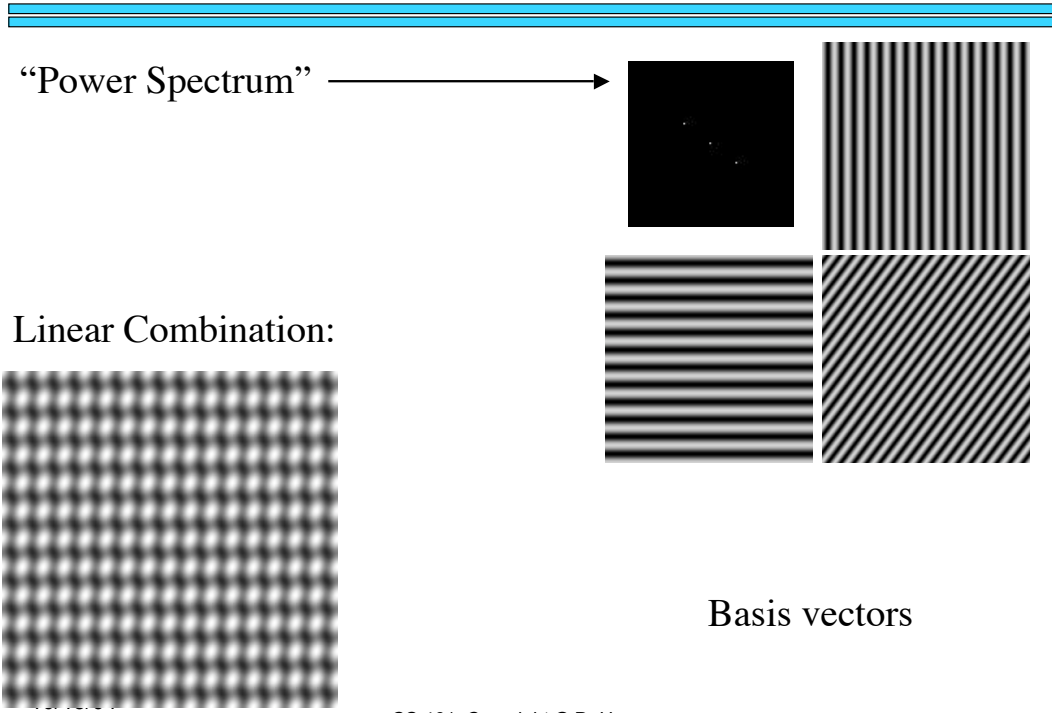
~ d f w e c o m i a t . k e g l a a r

And larger still...



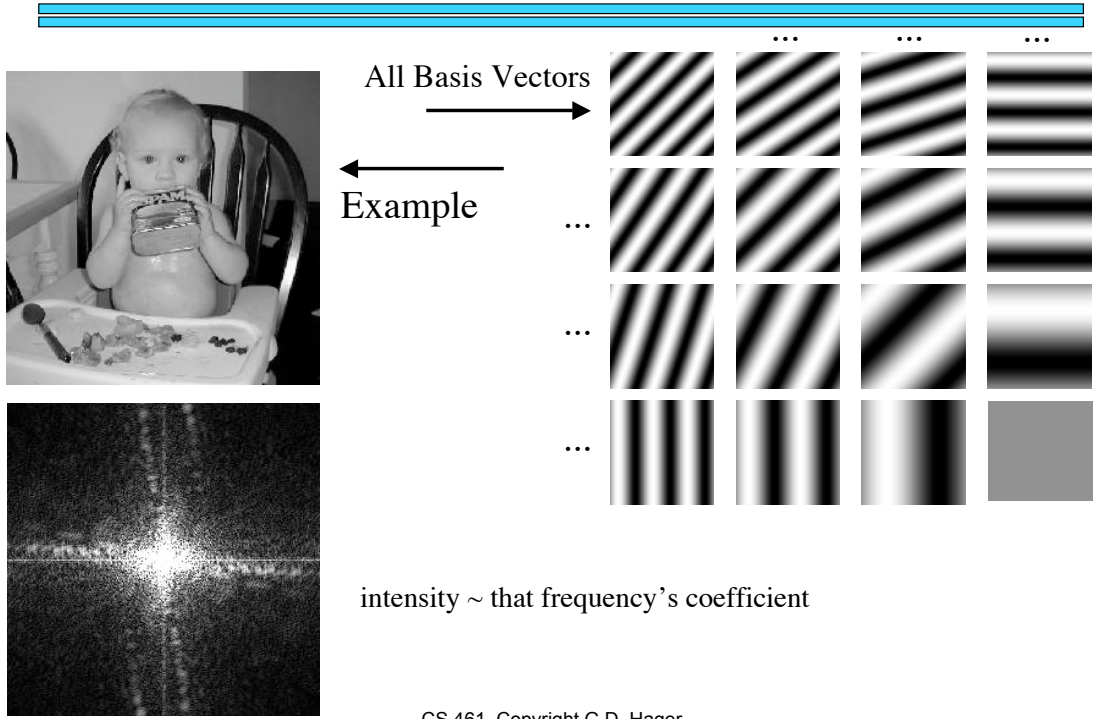
~ d f w e c o m i a t . k e g l a a r

# The Fourier “Hammer”



CS 461, Copyright G.D. Hager

# Frequency Decomposition




CS 461, Copyright G.D. Hager






# Using Fourier Representations

---

---



Smoothing



Data Reduction: only use *some* of the existing frequencies

10/15/04


CS 461, Copyright G.D. Hager

# Using Fourier Representations

---

---

Dominant Orientation



Limitations: not useful for local segmentation

10/15/04

CS 461, Copyright G.D. Hager

# Phase and Magnitude

$$e^{it} = \cos t + i \sin t$$

- Fourier transform of a real function is complex with real (R) and imaginary (I) components
  - difficult to plot, visualize
  - instead, we can think of the phase and magnitude of the transform
- Phase is the phase of the complex transform
  - $p(u) = \text{atan}(I(u)/R(u))$
- Magnitude is the magnitude of the complex transform
  - $m(u) = \text{sqrt}(R^2(u) + I^2(u))$
- Curious fact
  - all natural images have about the same magnitude transform
  - hence, phase seems to matter, but magnitude largely doesn't
- Demonstration
  - Take two pictures, swap the phase transforms, compute the inverse - what does the result look like?

10/15/04


CS 461, Copyright © R. D. Gregory



10/15/04

CS 461, Copyright © R. D. Gregory


This is the  
magnitude  
transform  
of the  
cheetah pic



10/15/04

CS 419, Copyright D. K. Regier

This is the  
phase  
transform  
of the  
cheetah pic



10/15/04


CS 419, Copyright D. K. Regier



10/15/04

CS 61C, © 2011 D. K. Regier

This is the  
magnitude  
transform  
of the  
zebra pic



10/15/04

CS 61C, © 2011 D. K. Regier


This is the phase transform of the zebra pic



10/15/04

Copyright © 2004, D. K. Regier


Reconstruction with zebra phase, cheetah magnitude



10/15/04

Copyright © 2004, D. K. Regier

Reconstruction  
with cheetah  
phase, zebra  
magnitude



10/15/04

CS 461, Copyright G.D. Hager

## The Fourier Transform and Convolution

- If  $H$  and  $G$  are images, and  $F(\cdot)$  represents Fourier transform, then

$$F(H * G) = F(H)F(G)$$


- Thus, one way of thinking about the properties of a convolution is by thinking of how it modifies the frequencies of the image to which it is applied.
- In particular, if we look at the power spectrum, then we see that convolving image  $H$  by  $G$  attenuates frequencies where  $G$  has low power, and amplifies those which have high power.
- This is referred to as the **Convolution Theorem**

10/15/04

CS 461, Copyright G.D. Hager

# The Properties of the Box Filter

$F(\text{mean filter}) =$



Thus, the mean filter enhances low frequencies but also has “side lobes” that admit higher frequencies

10/15/04

CS 461, Copyright G.D. Hager

## What a Box Filter Does



10/15/04

CS 461, Copyright G.D. Hager

# The Gaussian Filter: A Better Noise Reducer

- Ideally, we would like an averaging filter that removes (or at least attenuates) high frequencies beyond a given range

$$g(i, j; \sigma) = e^{-\frac{(i^2 + j^2)}{2\sigma^2}}$$

- It is not hard to show that the FT of a Gaussian is again a Gaussian. Hence, it operates as a low pass filter.
- Note that in general, we truncate --- a good general rule is that the width ( $w$ ) of the filter is at least such that  $w > 5 \sigma$ . Alternatively we can just stipulate that the width of the filter determines  $\sigma$  (or vice-versa).
- Note that in the discrete domain, we truncate the Gaussian, thus we are still subject to ringing like the box filter.

10/15/04

CS 461, Copyright G.D. Hager

## Smoothing by Averaging

Kernel: 




10/15/04

CS 461, Copyright G.D. Hager


# Smoothing with a Gaussian

Kernel: 



10/15/04

CS 439, Stanford University, © 2004



**The effects of smoothing**  
Each row shows smoothing with gaussians of different width; each column shows different realizations of an image of gaussian noise.

10/15/04

CS 439, Stanford University, © 2004

# Computational Issues: Separability

- Recall that convolution is associative. Suppose I use the templates  $g_x = \exp(-i^2/2 \sigma^2)$  and  $g_y = \exp(-j^2/2 \sigma^2)$ . Then
  - $g_x * (g_y * I) = (g_x * g_y) * I$
  - but, it is not hard to show that the first convolution is simply the 2-D Gaussian that we defined previously!
- In general, this means that we can “separate” the 2-D Gaussian convolution into 2 1-D convolutions with great computational cost savings.
- A good exercise is to show that the box filter is also separable.

10/15/04





CS 461, Copyright G.D. Hager

# Computational Issues: Minimizing Operations

- Note that for a 256 gray level image, we can *precompute* all values of the convolution and avoid multiplication.
- For the box filter, we can implement any size using  $4n$  additions per pixel.
- Also note that, by the central limit theorem, repeated box filter averaging yields approximations to a Gaussian filter.
- Finally, note that a sequence of filtering operations can be collapsed into one by associativity.
  - in general, this is not a win, but we’ll see examples where it is ...

10/15/04

CS 461, Copyright G.D. Hager



10/15/04

CS 461, Copyright G.D. Hager

## Other Types of Noise

- Impulsive noise
  - randomly pick a pixel and randomly set to a value
  - saturated version is called salt and pepper noise
- Quantization effects
  - Often called noise although it is not statistical
- Unanticipated image structures
  - Also often called noise although it is a real repeatable signal.

10/15/04

CS 461, Copyright G.D. Hager

# Digitization Effects

- The “diameter”  $d$  of a pixel determines the highest frequency representable in an image



$$l = 1/2d$$

- Real scenes may contain higher frequencies resulting in aliasing of the signal.
- In practice, this effect is often dominated by other digitization artifacts.
- One problem in particular is differing sampling rates between digitizer and camera readout of a row.

10/15/04

CS 461, Copyright G.D. Hager

# Limitations of Linear Operators



10/15/04

CS 461, Copyright G.D. Hager

# Nonlinear Filtering: The Median Filter

Suppose I look at the local statistics and replace each pixel with the *median* of it's neighbors:






10/15/04

CS 461, Copyright G.D. Hager

## Median filters : example

filters have width 5 :

	INPUT
	MEDIAN
	MEAN

10/15/04

CS 461, Copyright G.D. Hager

# Yet Another View of Convolution


- Suppose we consider the convolution template as a “vector”:
  - $T = [T_1, T_2, T_3 \dots T_n]$
- Likewise, consider a region of the image to which the convolution is applied as a vector
  - $I = [I_1, I_2, \dots, I_n]$
- Then the value of the convolution at a point is just the “dot product”
  - $v = T \cdot I$
- Thus, we can also think of convolution as a kind of “pattern match” where regions of the image that are “similar” to  $T$  respond more strongly than those that are dissimilar (up to a scale factor)
- Why does this make sense when thinking of Fourier transforms?

10/15/04

CS 461, Copyright G.D. Hager

## What Else Can You Do With Convolution?

- Thus far, we’ve only considered convolution kernels that are smoothing filters.
- Consider the following kernel:
  - $[1, -1]$




10/15/04

CS 461, Copyright G.D. Hager

# What Else Can You Do With Convolution?

- Thus far, we've only considered convolution kernels that are smoothing filters.
- Consider the following kernel:
  - $[1; -1]$



10/15/04

CS 461, Copyright G.D. Hager


## Physical causes of edges

1. Object boundaries
2. Surface normal discontinuities
3. Reflectance (albedo) discontinuities
4. Lighting discontinuities

10/15/04

CS 461, Copyright G.D. Hager

# Edge Types



Step




Ridge




Roof

10/15/04

CS 461, Copyright G.D. Hager




# Object Boundaries



CS 461, Copyright G.D. Hager



## Surface normal discontinuities



10/15/04

CS 461, Copyright G.D. Hager


## Boundaries of material properties



10/15/04

CS 461, Copyright G.D. Hager

# Boundaries of lighting



10/15/04

CS 461, Copyright G.D. Hager

## The Image Gradient

- Recall from calculus for a function of two variables  $f(x,y)$  :
  - The gradient: points in the direction of maximum increase.
  - Its magnitude is proportional to the rate of increase.
  - The total derivative in the direction  $n = n \cdot \nabla f$
- The kernel  $[-1,1]$  is a way of computing the x derivative
- The kernel  $[-1;1]$  is a way of computing the y derivative


10/15/04

CS 461, Copyright G.D. Hager

# Edge is Where Change Occurs


## 1-D

- Change is measured by derivative in 1D
  - Biggest change, derivative has maximum magnitude
  - Or 2<sup>nd</sup> derivative is zero.



## Noisy Step Edge

- Derivative is high everywhere.
- Must smooth before taking gradient.



## Some Other Interesting Kernels

The Roberts Operator  $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$

The Prewitt Operator  $\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{bmatrix}$

$$\begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix}$$

10/15/04

CS 461, Copyright G.D. Hager

## Some Other Interesting Kernals

The Sobel Operator  $\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$

The Laplacian Operator  $\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}$

A good exercise: derive the Laplacian from 1-D derivative filters.

Note the Laplacian is rotationally symmetric!

10/15/04

CS 461, Copyright G.D. Hager

# Smoothing Plus Derivatives


- One problem with differences is that they by definition reduce the signal to noise ratio (can you show this?)
- Recall smoothing operators (the Gaussian!) reduce noise.
- Hence, an obvious way of getting clean images with derivatives is to combine derivative filtering and smoothing: e.g.

$$- G * D_x * I = D_x * G * I$$

10/15/04


CS 461, Copyright G.D. Hager

## The Fourier Spectrum of DOG



Derivative of a Gaussian

PS of central slice



10/15/04

CS 461, Copyright G.D. Hager

# Properties of the DoG operator

- Now, going back to the directional derivative:
  - $D_u(f(x,y)) = f_x(x,y)u_1 + f_y(x,y)u_2$
- Now, including a Gaussian convolution, we see
  - $D_u[G*I] = D_u[G]*I = [u_1G_x + u_2G_y]*I = u_1G_x*I + u_2G_y*I$
- The two components  $G_x$  and  $G_y$  are the *image gradient*
- Note the directional derivative is maximized in the direction of the gradient
- (note some authors use DoG as “Difference of Gaussian” which we’ll run into soon ....)

10/15/04

CS 461, Copyright G.D. Hager

## Algorithm: Simple Edge Detection

- 1. Compute  $I_g = G(\sigma) * G(\sigma)' * I$
- 2. Compute  $I_x = [-1, 1; -1, 1] * I_g$
- 3. Compute  $I_y = [-1, 1; -1, 1]' * I_g$
- 4. Compute  $I_{mag} = \text{sqrt}(I_x \cdot I_x + I_y \cdot I_y)$
- 5. Threshold:  $I_{res} = I_{mag} > \tau$
  
- It is interesting to note that if we wanted an edge detector for a specific direction of edges, we can simply choose the appropriate projection (weighting) of the component derivatives.




10/15/04

CS 461, Copyright G.D. Hager


sigma = 1

## Example

sigma = 2



sigma = 5



sigma = 10

10/15/04

CS 461, Copyright G.D. Hager

## Types of Edge Operators

1. Operators approximating derivatives using differences.
  - directional: Roberts, Prewitt, DoG, etc.
  - Rotationally invariant: Laplacian (sum of second derivatives)
2. Operators based on the zero crossing of the second derivative (e.g. Canny).
3. Operators that attempt to match a specific image profile.

10/15/04

CS 461, Copyright G.D. Hager

# Filter Pyramids

- An Exercise:
  - Suppose I have  $G(\sigma)$  and I perform  $G(\sigma) * G(\sigma) * I$ 
    - Hint: think about the convolution theorem and the FFT
- Suppose I want to subsample images
  - subsampling reduces the highest frequencies
  - averaging reduces noise
  - Can I average and resample and reduce noise while not losing desirable frequencies?

10/15/04

CS 461, Copyright G.D. Hager

## Gaussian Pyramid

- Algorithm:
  - 1. Filter with  $G(\sigma)$
  - 2. Resample at every other pixel
  - 3. Repeat
- A common use of this is the Laplacian Pyramid



10/15/04

CS 461, Copyright G.D. Hager


# 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
  - $L(n+1) = G(n+1) - \text{upsample}(G(n))$

10/15/04

CS 461, Copyright G.D. Hager


## Laplacian of Gaussian Pyramid



10/15/04

CS 461, Copyright G.D. Hager

# Laplacian of Gaussian Pyramid



10/15/04

CS 461, Copyright G.D. Hager

## From Pixels to Edges

- Various operators can be used to *enhance* rapid contrast changes
- Detecting these contrast changes involves *thresholding* to separate noise from signal
- *Edges* are a result of *grouping* pixels (sometimes called “edgels”) into groups forming continuous curves.

Definitions:

*Edge normal*: Unit vector in direction of maximum intensity variation

*Edge direction*: Perpendicular to edge normal

*Edge position*: Image position of pixels of edge

*Edge strength*: Change in contrast along normal

10/15/04

CS 461, Copyright G.D. Hager

# From Pixels to Edges

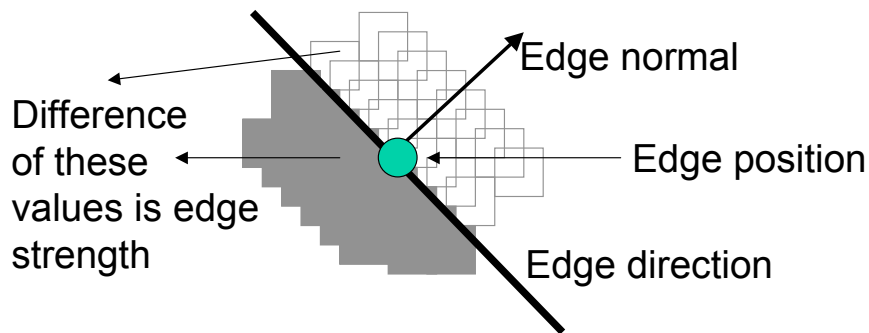
Definitions:

*Edge normal:* Unit vector in direction of maximum intensity variation

*Edge direction:* Perpendicular to edge normal

*Edge position:* Image position of pixels of edge

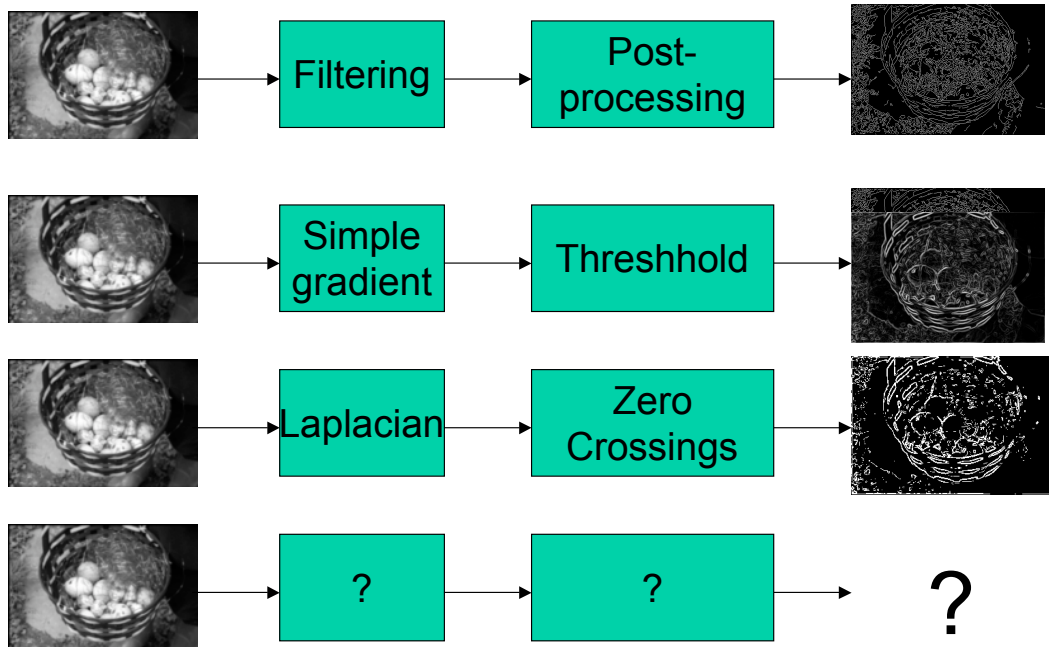
*Edge strength:* Change in contrast along normal



10/15/04

CS 461, Copyright G.D. Hager

## Edges: The Problem



What is optimal?

10/15/04

CS 461, Copyright G.D. Hager

# Canny Edge Detector

- The Plan:
  - Formulate an optimization problem for detection on 1-D signals
  - Generalize to 2D signals
  - Apply thresholding with hysteresis
  - Apply this operator at various scales
- The Assumptions:
  - Edge enhancement is linear
  - The edge model is step edges with amplitude A
  - Noise is additive, white and Gaussian

10/15/04

CS 461, Copyright G.D. Hager

## Optimization Criteria

- **Good Detection: *Minimize the probability of false positives and false negatives***

- Maximize the SNR

$$\frac{A}{n_0} \Sigma(f) = \left| \frac{A \int_{-W}^0 f(t) dt}{n_0 \sqrt{\int_{-W}^W f^2(t) dt}} \right|$$

- **Good Localization: *Edgels detected should lie as close as possible to the true edge***

- Maximize 1/distance to edge center which leads to maximizing LOC

$$\frac{A}{n_0} \Lambda(f') = \left| \frac{A |f'(0)|}{n_0 \sqrt{\int_{-W}^W f'^2(t) dt}} \right|$$

10/15/04

CS 461, Copyright G.D. Hager

# Optimization Cont'd

- Consider the maximizing the product of both criteria
  - result is itself a step filter
  - step filters are noise amplifying!
- Additional criterion: single response constraint:
  - detector should minimize the number of local maxima about an edge (recall what happens with step filter)
  - **RESULT1: localization vs. detection**

$$\Sigma(f_w) = \sqrt{w}\Sigma(f) \text{ and } \Lambda(f'_w) = \frac{1}{\sqrt{w}}\Lambda(f') \text{ where } f_w(x) = f(x/w)$$

- **RESULT2: optimal detector is very close to the first derivative of a Gaussian.**

10/15/04

CS 461, Copyright G.D. Hager

## The Procedure


- Enhancement:
  - compute x and y derivatives using DoG's.
  - compute direction and magnitude of gradient (two images)
- Nonmaximal Suppression:
  - Sample along the gradient direction
  - If given pixel is smaller than neighbor, set it to zero
- Hysteresis Thresholding:
  - Starting from upper left, visit pixels until one exceeds  $t_{\text{upper}}$
  - Follow chains of maxima in edge direction until value drops below  $t_{\text{lower}}$
  - Mark and save all visited values as a connected contour

10/15/04

CS 461, Copyright G.D. Hager


# Hysteresis

- Track edge points by starting at point where gradient magnitude  $> \tau_{\text{high}}$ .
- Follow edge in direction orthogonal to gradient.
- Stop when gradient magnitude  $< \tau_{\text{low}}$ .
  - i.e., use a high threshold to start edge curves and a low threshold to continue them.



CS 461, Copyright G.D. Hager



# Canny Output



10/15/04

CS 461, Copyright G.D. Hager

# Canny Comparison




10/15/04

CS 461, Copyright G.D. Hager



10/15/04

CS 461, Copyright G.D. Hager



fine scale  
high  
threshold

10/15/04

CS 461, Copyright G.D. Hager



coarse  
scale,  
high high  
threshold

10/15/04

CS 461, Copyright G.D. Hager



10/15/04

CS 461, Copyright G.D. Hager

## Why is Canny so Dominant

- Still widely used after 20 years.
  1. Theory is nice (but end result same,).
  2. Details good (magnitude of gradient, non-max suppression).
  3. Hysteresis an important heuristic.
  4. Code was distributed.

10/15/04

CS 461, Copyright G.D. Hager

# SubPixel Precision

- It is often hard to exactly localize the maximum of a function
- Many algorithms need sub-pixel precision
- Thus, it is common to apply a \*second derivative\* operator locally to locate the edge
  - note we know the edge direction, so we can compute second directional derivatives!


DO IT IN MATLAB GREG!

10/15/04

CS 461, Copyright G.D. Hager

## Corners


- Why are they important?



CS 461, Copyright G.D. Hager

## Corners contain more info than lines.

- A point on a line is hard to match.




10/15/04

CS 461, Copyright G.D. Hager

## Corners contain more info than lines.

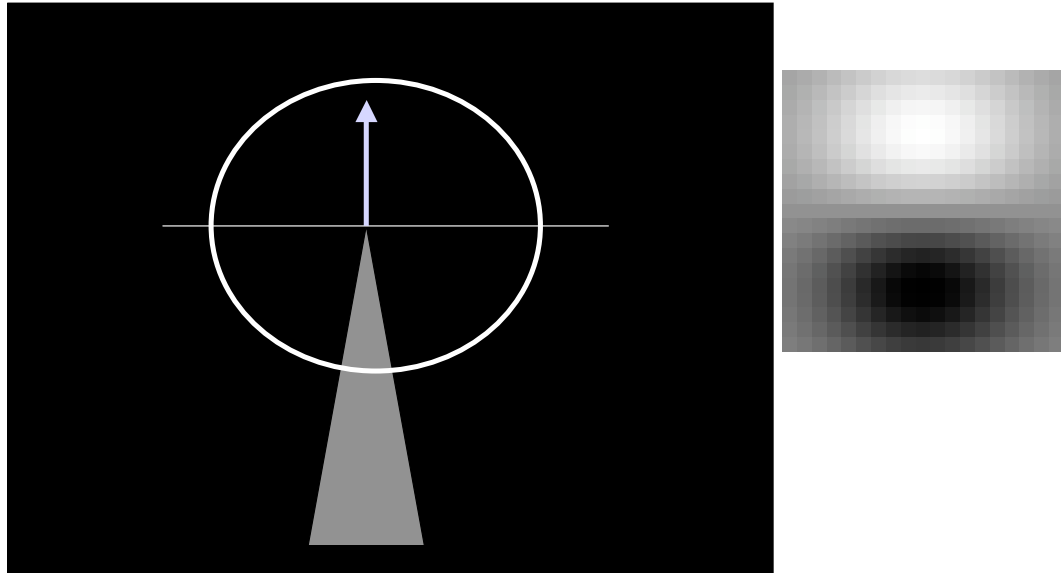
- A corner is easier to match



10/15/04

CS 461, Copyright G.D. Hager

## Edge Detectors Tend to Fail at Corners



10/15/04

CS 461, Copyright G.D. Hager

## Finding Corners

Intuition:

- Right at corner, gradient is ill defined.
- Near corner, gradient has two different values.

10/15/04

CS 461, Copyright G.D. Hager

# Another Use of Gradients: Detecting Corners

- Edges can be thought of as “1D” features
- Corners are “2D” features
- To make this precise, we need to think about the span of the gradients

$$C = \sum_N r I(u) (r I(u))' = R D R'$$

$$D = \text{diag}(\lambda_1, \lambda_2)$$

$$c = \min(\lambda_1, \lambda_2)$$

10/15/04

CS 461, Copyright G.D. Hager

## Formula for Finding Corners

We look at matrix:

Sum over a small region,  
the hypothetical corner

Gradient with respect to x,  
times gradient with respect to y

$$C = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix}$$

Matrix is symmetric

**WHY THIS?**

10/16/04

CS 461, Copyright G.D. Hager

First, consider case where:

$$C = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

This means all gradients in neighborhood are:

(k,0) or (0, c) or (0, 0) (or off-diagonals cancel).

What is region like if:

1.  $\lambda_1=0$ ?
2.  $\lambda_2=0$ ?
3.  $\lambda_1=0$  and  $\lambda_2=0$ ?

10/14/04 <sup>10/14/04</sup>  $\lambda_1 > 0$  and  $\lambda_2 > 0$ ? CS 461, Copyright G.D. Hager

General Case:

From Linear Algebra we haven't talked about it follows that since C is symmetric:

$$C = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

So every case is like one on last slide.

# Corners Algorithm

- For every point  $(u,v)$ , compute  $C$  for a neighborhood about  $(u,v)$
- Sort by the minimum singular value of  $C$
- Read off locations starting with highest values and working down until enough locations are found or we run out of locations
  - optional: as we read down, discard corners that are within a small distance of corners that have appeared higher in the list

10/15/04

CS 461, Copyright G.D. Hager

# Edges Summary

- Filtering is a way of removing noise or suppressing/enhancing frequency content
- Typically, we combine some type of image derivative with smoothing
- Image gradients are the basic tool in 2D images
- Derivative of Gaussian is generally the gradient operator of choice
- Canny detector is probably the most widely used algorithm for performing edge detection

10/15/04

CS 461, Copyright G.D. Hager