



Filters, Features, Edges

Thursday, Sept 11

Last time

- Cross correlation
- Convolution
- Examples of smoothing filters
 - Box filter (averaging)
 - Gaussian

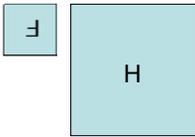
Convolution

- Convolution:
 - Flip the filter in both dimensions (bottom to top, right to left)
 - Then apply cross-correlation

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v]F[i - u, j - v]$$

$G = H \star F$

Notation for convolution operator



Smoothing with a Gaussian

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



...


```

for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
  
```

Predict the filtered outputs


 \star

0	0	0
0	1	0
0	0	0

 $= ?$


 \star

0	0	0
0	0	1
0	0	0

 $= ?$


 \star

0	0	0
0	2	0
0	0	0

 $= \frac{1}{9} \begin{matrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{matrix} = ?$

Practice with linear filters



0	0	0
0	1	0
0	0	0

?

Original

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	1	0
0	0	0



Filtered
(no change)

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0

?

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0



Shifted left
by 1 pixel
with
correlation

Source: D. Lowe

Practice with linear filters



Original

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

?

Source: D. Lowe

Practice with linear filters



Original

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



Blur (with a
box filter)

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	2	0
0	0	0

- $\frac{1}{9}$

1	1	1
1	1	1
1	1	1

?

Source: D. Lowe

Practice with linear filters

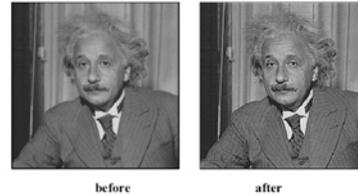
Original

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} \cdot \frac{1}{9} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Sharpening filter
- Accentuates differences with local average

Source: D. Lowe

Filtering examples: sharpening



before

after

Shift invariant linear system

- **Shift invariant:**
 - Operator behaves the same everywhere, i.e. the value of the output depends on the pattern in the image neighborhood, not the position of the neighborhood.
- **Linear:**
 - Superposition: $h * (f_1 + f_2) = (h * f_1) + (h * f_2)$
 - Scaling: $h * (k f) = k (h * f)$

Properties of convolution

- Linear & shift invariant
- Commutative:
 $f * g = g * f$
- Associative
 $(f * g) * h = f * (g * h)$
- Identity:
unit impulse $e = [\dots, 0, 0, 1, 0, 0, \dots]$. $f * e = f$
- Differentiation:
 $\frac{\partial}{\partial x} (f * g) = \frac{\partial f}{\partial x} * g$

Separability

- In some cases, filter is separable, and we can factor into two steps:
 - Convolve all rows
 - Convolve all columns

Separability

- In some cases, filter is separable, and we can factor into two steps: e.g.,

g

1	2	1
2	3	3
4	4	6

h

11		
18		
18		

f

1	11		
2	18		65
1	18		

What is the computational complexity advantage for a separable filter of size $k \times k$, in terms of number of operations per output pixel?

$$\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \times \begin{bmatrix} 1 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} \begin{bmatrix} 2 & 3 & 3 \\ 3 & 5 & 5 \\ 4 & 4 & 6 \end{bmatrix} = \begin{matrix} = 2 + 6 + 3 = 11 \\ = 6 + 20 + 10 = 36 \\ = 4 + 8 + 6 = 18 \\ \hline 65 \end{matrix}$$

$$f * (g * h) = (f * g) * h$$

Effect of smoothing filters

5x5

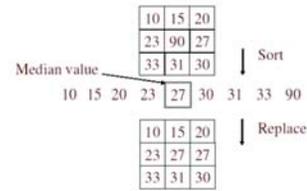


Additive Gaussian noise



Salt and pepper noise

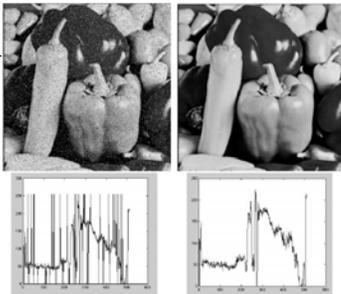
Median filter



- No new pixel values introduced
- Removes spikes: good for impulse, salt & pepper noise
- Linear?

Median filter

Salt and pepper noise



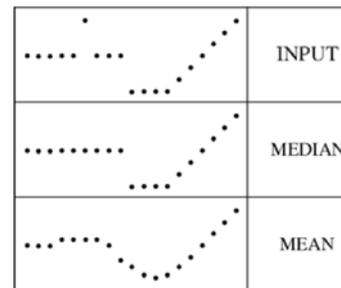
Median filtered

Plots of a row of the image

Source: M. Hebert

Median filter

- Median filter is edge preserving



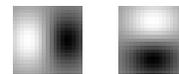
Filters for features

- Previously, thinking of filtering as a way to remove or reduce **noise**
- Now, consider how filters will allow us to abstract higher-level **“features”**.
 - Map raw pixels to an intermediate representation that will be used for subsequent processing
 - Goal: reduce amount of data, discard redundancy, preserve what's useful



Template matching

- Filters as **templates**:
Note that filters look like the effects they are intended to find --- “matched filters”



- Use normalized cross-correlation score to find a given pattern (template) in the image.
- Normalization needed to control for relative brightnesses.

Template matching

Scene

Template (mask)

A toy example

Template matching

Detected template

Template

Template matching

Detected template

Correlation map

Where's Waldo?

Scene

Template

Where's Waldo?

Detected template

Template

Where's Waldo?

Detected template

Correlation map

Template matching



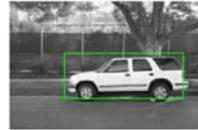
Scene



Template

What if the template is not identical to some subimage in the scene?

Template matching



Detected template



Template

Match can be meaningful, if scale, orientation, and general appearance is right.

Edge detection

- **Goal:** map image from 2d array of pixels to a set of curves or line segments or contours.
- **Why?**

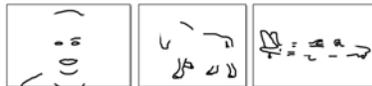
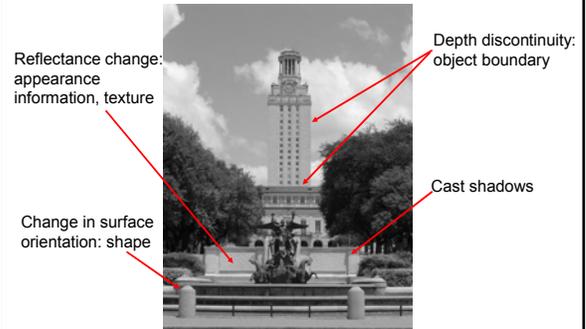


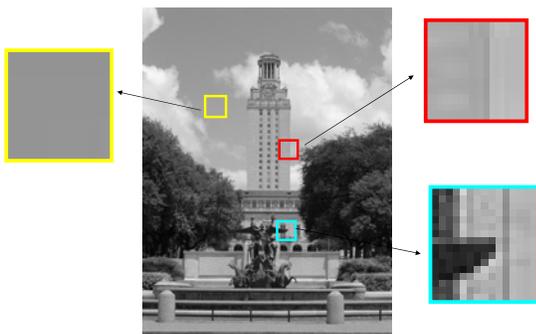
Figure from J. Shotton et al., PAMI 2007

- **Main idea:** look for strong gradients, post-process

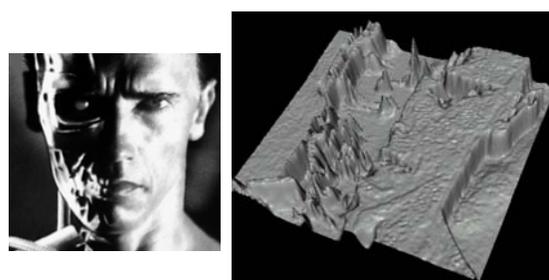
What can cause an edge?



Contrast and invariance



Recall : Images as functions



Edges look like steep cliffs

Source: S. Seitz

Derivatives and edges

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

Source: I. Lazebnik

Differentiation and convolution

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

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

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 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

Which shows changes with respect to x?

(showing flipped filters)

Assorted finite difference filters

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

```

>> My = fspecial('sobel');
>> outim = imfilter(double(im), My);
>> imagesc(outim);
>> colormap gray;
    
```

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:

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

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. Seltz

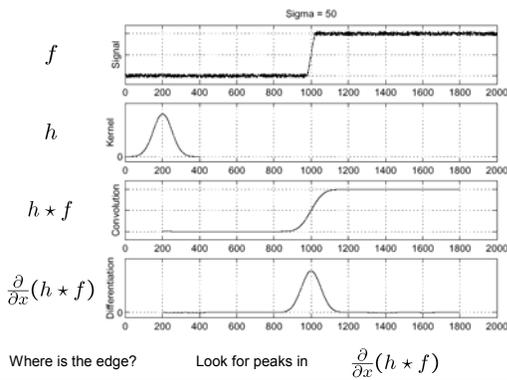
Effects of noise

Consider a single row or column of the image

- Plotting intensity as a function of position gives a signal

Where is the edge?

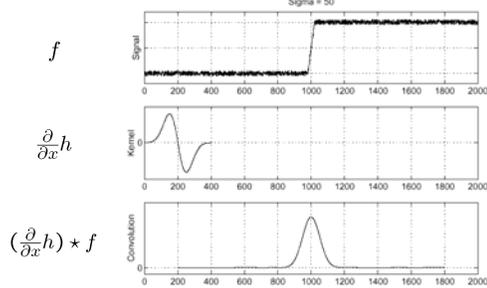
Solution: smooth first



Derivative theorem of convolution

$$\frac{\partial}{\partial x}(h * f) = (\frac{\partial}{\partial x}h) * f$$

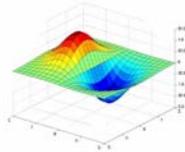
Differentiation property of convolution.



Derivative of Gaussian filter

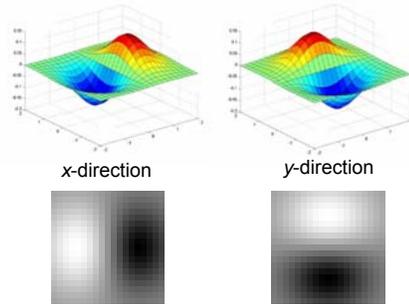
$$(I \otimes g) \otimes h = I \otimes (g \otimes h)$$

$$\begin{bmatrix} 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0219 & 0.0983 & 0.1621 & 0.0983 & 0.0219 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \end{bmatrix} \otimes \begin{bmatrix} 1 & -1 \end{bmatrix}$$



Why is this preferable?

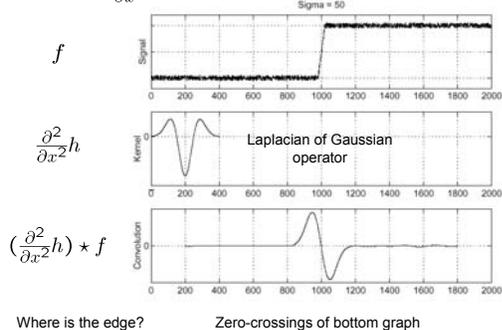
Derivative of Gaussian filters



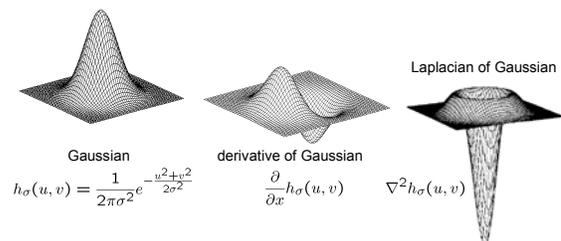
Source: L. Lazebnik

Laplacian of Gaussian

Consider $\frac{\partial^2}{\partial x^2}(h * f)$



2D edge detection filters



• ∇^2 is the **Laplacian** operator:

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

Mask properties

- Smoothing
 - Values positive
 - Sum to 1 → constant regions same as input
 - Amount of smoothing proportional to mask size
 - Remove "high-frequency" components; "low-pass" filter
- Derivatives
 - Opposite signs used to get high response in regions of high contrast
 - Sum to 0 → no response in constant regions
 - High absolute value at points of high contrast
- Filters act as templates
 - Highest response for regions that "look the most like the filter"
 - Dot product as correlation

Gradients -> edges

Primary edge detection steps:

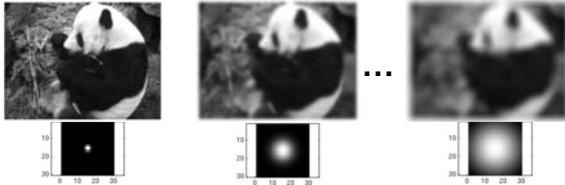
1. Smoothing: suppress noise
2. Edge enhancement: filter for contrast
3. Edge localization

Determine which local maxima from filter output are actually edges vs. noise

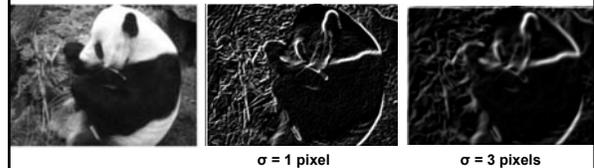
- Threshold, Thin

Smoothing with a Gaussian

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



Effect of σ on derivatives



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

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

So, what scale to choose?

It depends what we're looking for.



Too fine of a scale...can't see the forest for the trees.

Too coarse of a scale...can't tell the maple grain from the cherry.

Thresholding

- Choose a threshold value t
- Set any pixels less than t to zero (off)
- Set any pixels greater than or equal to t to one (on)

Original image



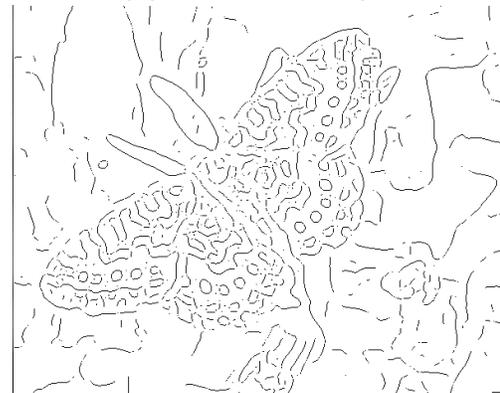
Gradient magnitude image



Thresholding gradient with a lower threshold



Thresholding gradient with a higher threshold



Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression:**
 - Thin multi-pixel wide "ridges" down to single pixel width
- 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');`
- `>>help edge`

Source: D. Lowe, L. Fei-Fei

The Canny edge detector



original image (Lena)

The Canny edge detector



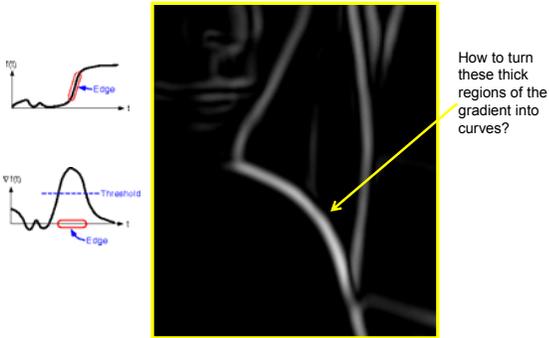
norm of the gradient

The Canny edge detector

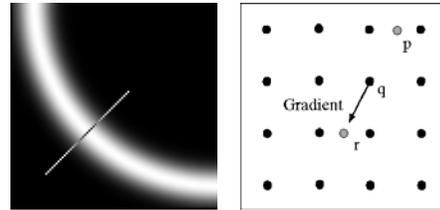


thresholding

The Canny edge detector



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.



Hysteresis thresholding



original image



high threshold
(strong edges)



low threshold
(weak edges)



hysteresis threshold

Source: L. Fei-Fei

Object boundaries vs. edges




Background




Texture




Shadows

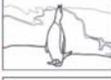
Edge detection is just the beginning...

image	human segmentation	gradient magnitude
		
		

Berkeley segmentation database:
<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>

Source: L. Lazebnik

Possible to learn from humans which combination of features is most indicative of a "good" contour?

[D. Martin et al. PAMI 2004] Human-marked segment boundaries

What features are responsible for perceived edges?

	Image	Intensity	OE	\overline{OE}	BG	CG	TG	\overline{TG}
Non-Boundaries	(a) 							
	(b) 							
	(c) 							
	(d) 							

Feature profiles (oriented energy, brightness, color, and texture gradients) along the patch's horizontal diameter

[D. Martin et al. PAMI 2004]

Image					
BG+CG+TG					
Human					

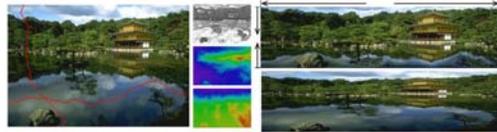
[D. Martin et al. PAMI 2004]

Summary

- Filters allow local image neighborhood to influence our description and features
 - Smoothing to reduce noise
 - Derivatives to locate contrast, gradient
- Filters have highest response on neighborhoods that “look like” it; can be thought of as template matching.
- Convolution properties will influence the efficiency with which we can process images.
 - Associative
 - Filter separability
- Edge detection processes the image gradient to find curves, or chains of edgels.

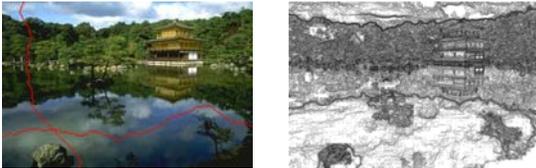
Next

- Tues 9/16 binary images
- Reminder: Pset 1 due Sept 18.



Seam Carving

- Energy function: $e_1(\mathbf{I}) = \left| \frac{\partial}{\partial x} \mathbf{I} \right| + \left| \frac{\partial}{\partial y} \mathbf{I} \right|$



- Want to remove or insert seams where they won't be very noticeable
- Choose seam based on minimum total energy path across image.